

# Contents

<b>1</b>	<b>Notation and conventions</b>	<b>7</b>
1.0.1	Background Information . . . . .	8
1.1	Some Useful Mathematical Facts . . . . .	9
1.2	Acknowledgements . . . . .	9
<b>2</b>	<b>First Tools for Looking at Data</b>	<b>10</b>
2.1	Datasets . . . . .	10
2.2	What's Happening? - Plotting Data . . . . .	12
2.2.1	Bar Charts . . . . .	13
2.2.2	Histograms . . . . .	13
2.2.3	How to Make Histograms . . . . .	14
2.2.4	Conditional Histograms . . . . .	16
2.3	Summarizing 1D Data . . . . .	16
2.3.1	The Mean . . . . .	16
2.3.2	Standard Deviation and Variance . . . . .	19
2.3.3	Variance . . . . .	23
2.3.4	The Median . . . . .	24
2.3.5	Interquartile Range . . . . .	26
2.3.6	Using Summaries Sensibly . . . . .	27
2.4	Plots and Summaries . . . . .	28
2.4.1	Some Properties of Histograms . . . . .	28
2.4.2	Standard Coordinates and Normal Data . . . . .	31
2.4.3	Boxplots . . . . .	34
2.5	Whose is bigger? Investigating Australian Pizzas . . . . .	35
2.6	You should . . . . .	40
2.6.1	be able to: . . . . .	40
2.6.2	remember: . . . . .	40
<b>3</b>	<b>Looking at Relationships</b>	<b>44</b>
3.1	Plotting 2D Data . . . . .	44
3.1.1	Categorical Data, Counts, and Charts . . . . .	44
3.1.2	Series . . . . .	48
3.1.3	Scatter Plots for Spatial Data . . . . .	50
3.1.4	Exposing Relationships with Scatter Plots . . . . .	53
3.2	Correlation . . . . .	56
3.2.1	The Correlation Coefficient . . . . .	58
3.2.2	Using Correlation to Predict . . . . .	63
3.2.3	Confusion caused by correlation . . . . .	67
3.3	Sterile Males in Wild Horse Herds . . . . .	68
3.4	You should . . . . .	71
3.4.1	be able to: . . . . .	71
3.4.2	remember: . . . . .	71

<b>4</b>	<b>Basic ideas in probability</b>	<b>77</b>
4.1	Experiments, Outcomes, Events, and Probability . . . . .	77
4.1.1	The Probability of an Outcome . . . . .	79
4.1.2	Events . . . . .	81
4.1.3	The Probability of Events . . . . .	84
4.1.4	Computing Probabilities by Counting Outcomes . . . . .	87
4.1.5	Computing Probabilities by Reasoning about Sets . . . . .	91
4.2	Independence . . . . .	94
4.2.1	Example: Airline Overbooking . . . . .	98
4.3	Conditional Probability . . . . .	102
4.3.1	Evaluating Conditional Probabilities . . . . .	102
4.3.2	The Prosecutors Fallacy . . . . .	109
4.3.3	Independence and Conditional Probability . . . . .	110
4.4	Example: The Monty Hall Problem . . . . .	112
4.5	You should . . . . .	116
4.5.1	be able to: . . . . .	116
4.5.2	remember: . . . . .	116
<b>5</b>	<b>Random Variables and Expectations</b>	<b>122</b>
5.1	Random Variables . . . . .	122
5.1.1	Joint and Conditional Probability for Random Variables . . . . .	125
5.1.2	Just a Little Continuous Probability . . . . .	128
5.2	Expectations and Expected Values . . . . .	130
5.2.1	Expected Values . . . . .	131
5.2.2	Mean, Variance and Covariance . . . . .	134
5.2.3	Expectations and Statistics . . . . .	139
5.3	Two Inequalities, and a Law . . . . .	139
5.3.1	Indicator Functions . . . . .	141
5.3.2	Proving the Inequalities . . . . .	142
5.3.3	IID Samples and the Weak Law of Large Numbers . . . . .	143
5.4	Using Expectations . . . . .	145
5.4.1	Should you accept a bet? . . . . .	146
5.4.2	Odds, Expectations and Bookmaking — a Cultural Diversion . . . . .	148
5.4.3	Ending a Game Early . . . . .	149
5.4.4	Making a Decision with Decision Trees and Expectations . . . . .	150
5.4.5	Utility . . . . .	152
5.5	You should . . . . .	155
5.5.1	be able to: . . . . .	155
5.5.2	remember: . . . . .	155
<b>6</b>	<b>Useful Probability Distributions</b>	<b>161</b>
6.1	Discrete Distributions . . . . .	162
6.1.1	The Discrete Uniform Distribution . . . . .	162
6.1.2	Bernoulli Random Variables . . . . .	162
6.1.3	The Geometric Distribution . . . . .	163
6.1.4	The Binomial Probability Distribution . . . . .	164
6.1.5	Multinomial probabilities . . . . .	166

6.1.6	The Poisson Distribution . . . . .	167
6.2	Continuous Distributions . . . . .	169
6.2.1	The Continuous Uniform Distribution . . . . .	169
6.2.2	The Beta Distribution . . . . .	169
6.2.3	The Gamma Distribution . . . . .	170
6.2.4	The Exponential Distribution . . . . .	171
6.3	The Normal Distribution . . . . .	172
6.3.1	The Standard Normal Distribution . . . . .	172
6.3.2	The Normal Distribution . . . . .	174
6.3.3	Properties of the Normal Distribution . . . . .	175
6.4	Approximating Binomials with Large $N$ . . . . .	176
6.4.1	Large $N$ . . . . .	177
6.4.2	Getting Normal . . . . .	179
6.4.3	So What? . . . . .	181
6.5	You should . . . . .	183
6.5.1	remember: . . . . .	183
<b>7</b>	<b>Markov Chains and Simulation</b>	<b>189</b>
7.1	Markov Chains . . . . .	189
7.1.1	Motivating Example: Multiple Coin Flips . . . . .	189
7.1.2	Motivating Example: The Gambler's Ruin . . . . .	191
7.1.3	Motivating Example: A Virus . . . . .	193
7.1.4	Markov Chains . . . . .	193
7.1.5	Example: Particle Motion as a Markov Chain . . . . .	197
7.2	Simulation . . . . .	198
7.2.1	Obtaining Random Numbers . . . . .	198
7.2.2	Computing Expectations with Simulations . . . . .	199
7.2.3	Computing Probabilities with Simulations . . . . .	200
7.2.4	Simulation Results as Random Variables . . . . .	201
7.3	Simulation Examples . . . . .	203
7.3.1	Simulating Experiments . . . . .	203
7.3.2	Simulating Markov Chains . . . . .	204
7.3.3	Example: Ranking the Web by Simulating a Markov Chain . . . . .	206
7.3.4	Example: Simulating a Complicated Game . . . . .	209
7.4	You should . . . . .	213
7.4.1	remember . . . . .	213
<b>8</b>	<b>Inference: Making Point Estimates</b>	<b>221</b>
8.1	Estimating Model Parameters with Maximum Likelihood . . . . .	222
8.1.1	The Maximum Likelihood Principle . . . . .	223
8.1.2	Cautions about Maximum Likelihood . . . . .	230
8.2	Incorporating Priors with Bayesian Inference . . . . .	231
8.2.1	Constructing the Posterior . . . . .	232
8.2.2	Normal Prior and Normal Likelihood . . . . .	236
8.2.3	MAP Inference . . . . .	239
8.2.4	Filtering . . . . .	240
8.2.5	Cautions about Bayesian Inference . . . . .	243

8.3	Samples, Urns and Populations . . . . .	243
8.3.1	Estimating the Population Mean from a Sample . . . . .	244
8.3.2	The Variance of the Sample Mean . . . . .	245
8.3.3	The Probability Distribution of the Sample Mean . . . . .	249
8.3.4	When The Urn Model Works . . . . .	249
8.4	You should . . . . .	251
8.4.1	be able to: . . . . .	251
8.4.2	remember: . . . . .	251
<b>9</b>	<b>Inference: Intervals and Testing</b>	<b>255</b>
9.1	Straightforward Constructions of Confidence Intervals . . . . .	255
9.1.1	Confidence Intervals for Population Means . . . . .	256
9.1.2	Bayesian Confidence Intervals . . . . .	259
9.2	Using Simulation to Construct Intervals . . . . .	262
9.2.1	Constructing Confidence Intervals for Parametric Models . . . . .	262
9.2.2	Estimating Standard Error . . . . .	265
9.3	Using the Standard Error to Evaluate Hypotheses . . . . .	269
9.3.1	Does this Population have this Mean? . . . . .	271
9.3.2	Do Two Populations have the same Mean? . . . . .	275
9.3.3	Variants on the Basic Test . . . . .	280
9.4	$\chi^2$ Tests: Is Data Consistent with a Model? . . . . .	281
9.5	You should . . . . .	284
9.5.1	remember . . . . .	284
<b>10</b>	<b>Extracting Important Relationships in High Dimensions</b>	<b>287</b>
10.1	Summaries and Simple Plots . . . . .	287
10.1.1	The Mean . . . . .	287
10.1.2	Parallel Plots and Stem Plots . . . . .	288
10.1.3	Blobs . . . . .	290
10.1.4	Understanding Blobs with Scatterplot Matrices . . . . .	290
10.1.5	Using Covariance to encode Variance and Correlation . . . . .	292
10.2	Blob Analysis of High-Dimensional Data . . . . .	297
10.2.1	Transforming High Dimensional Data . . . . .	297
10.2.2	Transforming Blobs . . . . .	298
10.2.3	Whitening Data . . . . .	301
10.3	Principal Components Analysis . . . . .	303
10.3.1	The Blob Coordinate System and Smoothing . . . . .	303
10.3.2	The Low-Dimensional Representation of a Blob . . . . .	306
10.3.3	Smoothing Data with a Low-Dimensional Representation . . . . .	308
10.3.4	The Error of the Low-Dimensional Representation . . . . .	309
10.3.5	Example: Representing Spectral Reflectances . . . . .	311
10.3.6	Example: Representing Faces with Principal Components . . . . .	313
10.4	Multi-Dimensional Scaling . . . . .	315
10.4.1	Principal Coordinate Analysis . . . . .	315
10.4.2	Example: Mapping with Multidimensional Scaling . . . . .	318
10.5	Example: Understanding Height and Weight . . . . .	320
10.6	You should . . . . .	324

10.6.1	be able to: . . . . .	324
10.6.2	remember: . . . . .	324
10.7	Appendix: Tables of Data . . . . .	326
<b>11</b>	<b>Learning to Classify</b>	<b>328</b>
11.1	Classification, Error, and Loss . . . . .	328
11.1.1	What is “Best”? . . . . .	329
11.1.2	What will happen on Test Data? . . . . .	329
11.2	The Support Vector Machine . . . . .	331
11.2.1	Choosing a Classifier with the Hinge Loss . . . . .	332
11.2.2	Finding a Minimum: General Points . . . . .	334
11.2.3	Finding a Minimum: Stochastic Gradient Descent . . . . .	334
11.2.4	Example: Training a Support Vector Machine with Stochastic Gradient Descent	336
11.2.5	Multi-Class Classifiers . . . . .	338
11.3	Classifying with Random Forests . . . . .	340
11.3.1	Building a Decision Tree . . . . .	341
11.3.2	Entropy and Information Gain . . . . .	344
11.3.3	Entropy and Splits . . . . .	345
11.3.4	Choosing a Split with Information Gain . . . . .	346
11.3.5	Forests . . . . .	348
11.3.6	Building and Evaluating a Decision Forest . . . . .	349
11.3.7	Classifying Data Items with a Decision Forest . . . . .	349
11.4	Classifying with Nearest Neighbors . . . . .	352
11.5	You should . . . . .	355
11.5.1	remember: . . . . .	355
<b>12</b>	<b>Clustering: Models of High Dimensional Data</b>	<b>357</b>
12.1	The Curse of Dimension . . . . .	357
12.1.1	The Curse: Data isn’t Where You Think it is . . . . .	357
12.1.2	Minor Banes of Dimension . . . . .	359
12.2	The Multivariate Normal Distribution . . . . .	359
12.2.1	Affine Transformations and Gaussians . . . . .	360
12.2.2	Plotting a 2D Gaussian: Covariance Ellipses . . . . .	361
12.3	Agglomerative and Divisive Clustering . . . . .	362
12.3.1	Clustering and Distance . . . . .	363
12.4	The K-Means Algorithm and Variants . . . . .	367
12.4.1	How to choose K . . . . .	371
12.4.2	Soft Assignment . . . . .	372
12.4.3	General Comments on K-Means . . . . .	374
12.5	Describing Repetition with Vector Quantization . . . . .	375
12.5.1	Vector Quantization . . . . .	376
12.5.2	Efficient Clustering and Hierarchical K Means . . . . .	378
12.5.3	Example: Activity from Accelerometer Data . . . . .	379
12.6	You should . . . . .	382
12.6.1	remember: . . . . .	382

<b>13 Regression</b>	<b>384</b>
13.1 Linear Regression and Least Squares . . . . .	386
13.1.1 Linear Regression . . . . .	386
13.1.2 Residuals and R-squared . . . . .	390
13.2 Producing Good Linear Regressions . . . . .	394
13.2.1 Checking a Linear Regression . . . . .	395
13.2.2 Finding Problem Data Points . . . . .	396
13.2.3 Adding Explanatory Variables . . . . .	397
13.2.4 Regularizing Linear Regressions . . . . .	398
13.2.5 Transforming Variables . . . . .	399
13.3 Exploiting Your Neighbors for Regression . . . . .	401
13.3.1 Using your Neighbors to Predict More than a Number . . . . .	402
13.4 You should . . . . .	406
13.4.1 remember: . . . . .	406
<b>14 Math Resources</b>	<b>410</b>
14.1 Useful Material about Matrices . . . . .	410
14.1.1 Approximating A Symmetric Matrix . . . . .	411

## CHAPTER 1

# Notation and conventions

A dataset as a collection of  $d$ -tuples (a  $d$ -tuple is an ordered list of  $d$  elements). Tuples differ from vectors, because we can always add and subtract vectors, but we cannot necessarily add or subtract tuples. There are always  $N$  items in any dataset. There are always  $d$  elements in each tuple in a dataset. The number of elements will be the same for every tuple in any given tuple. Sometimes we may not know the value of some elements in some tuples.

We use the same notation for a tuple and for a vector. Most of our data will be vectors. We write a vector in bold, so  $\mathbf{x}$  could represent a vector or a tuple (the context will make it obvious which is intended).

The entire data set is  $\{\mathbf{x}\}$ . When we need to refer to the  $i$ 'th data item, we write  $\mathbf{x}_i$ . Assume we have  $N$  data items, and we wish to make a new dataset out of them; we write the dataset made out of these items as  $\{\mathbf{x}_i\}$  (the  $i$  is to suggest you are taking a set of items and making a dataset out of them). If we need to refer to the  $j$ 'th component of a vector  $\mathbf{x}_i$ , we will write  $x_i^{(j)}$  (notice this isn't in bold, because it is a component not a vector, and the  $j$  is in parentheses because it isn't a power). Vectors are always column vectors.

When I write  $\{kx\}$ , I mean the dataset created by taking each element of the dataset  $\{x\}$  and multiplying by  $k$ ; and when I write  $\{x + c\}$ , I mean the dataset created by taking each element of the dataset  $\{x\}$  and adding  $c$ .

### Terms:

- $\text{mean}(\{x\})$  is the mean of the dataset  $\{x\}$  (definition 2.1, page 16).
- $\text{std}(\{x\})$  is the standard deviation of the dataset  $\{x\}$  (definition 2.2, page 19).
- $\text{var}(\{x\})$  is the standard deviation of the dataset  $\{x\}$  (definition 2.3, page 23).
- $\text{median}(\{x\})$  is the standard deviation of the dataset  $\{x\}$  (definition 2.4, page 24).
- $\text{percentile}(\{x\}, k)$  is the  $k\%$  percentile of the dataset  $\{x\}$  (definition 2.5, page 26).
- $\text{iqr}\{x\}$  is the interquartile range of the dataset  $\{x\}$  (definition 2.7, page 27).
- $\{\hat{x}\}$  is the dataset  $\{x\}$ , transformed to standard coordinates (definition 2.8, page 31).
- Standard normal data is defined in definition 2.9, (page 32).
- Normal data is defined in definition 2.10, (page 32).
- $\text{corr}(\{(x, y)\})$  is the correlation between two components  $x$  and  $y$  of a dataset (definition 3.1, page 58).

- $\emptyset$  is the empty set.
- $\Omega$  is the set of all possible outcomes of an experiment.
- Sets are written as  $\mathcal{A}$ .
- $\mathcal{A}^c$  is the complement of the set  $\mathcal{A}$  (i.e.  $\Omega - \mathcal{A}$ ).
- $\mathcal{E}$  is an event (page 392).
- $P(\{\mathcal{E}\})$  is the probability of event  $\mathcal{E}$  (page 392).
- $P(\{\mathcal{E}\}|\{\mathcal{F}\})$  is the probability of event  $\mathcal{E}$ , conditioned on event  $\mathcal{F}$  (page 392).
- $p(x)$  is the probability that random variable  $X$  will take the value  $x$ ; also written  $P(\{X = x\})$  (page 392).
- $p(x, y)$  is the probability that random variable  $X$  will take the value  $x$  and random variable  $Y$  will take the value  $y$ ; also written  $P(\{X = x\} \cap \{Y = y\})$  (page 392).
- $\operatorname{argmax}_x f(x)$  means the value of  $x$  that maximises  $f(x)$ .
- $\operatorname{argmin}_x f(x)$  means the value of  $x$  that minimises  $f(x)$ .
- $\max_i(f(x_i))$  means the largest value that  $f$  takes on the different elements of the dataset  $\{x_i\}$ .
- $\hat{\theta}$  is an estimated value of a parameter  $\theta$ .

### 1.0.1 Background Information

*Cards:* A standard deck of playing cards contains 52 cards. These cards are divided into four suits. The suits are: spades and clubs (which are black); and hearts and diamonds (which are red). Each suit contains 13 cards: Ace, 2, 3, 4, 5, 6, 7, 8, 9, 10, Jack (sometimes called Knave), Queen and King. It is common to call Jack, Queen and King *court cards*.

*Dice:* If you look hard enough, you can obtain dice with many different numbers of sides (though I've never seen a three sided die). We adopt the convention that the sides of an  $N$  sided die are labeled with the numbers  $1 \dots N$ , and that no number is used twice. Most dice are like this.

*Fairness:* Each face of a fair coin or die has the same probability of landing upmost in a flip or roll.

*Roulette:* A roulette wheel has a collection of slots. There are 36 slots numbered with the digits  $1 \dots 36$ , and then one, two or even three slots numbered with zero. There are no other slots. A ball is thrown at the wheel when it is spinning, and it bounces around and eventually falls into a slot. If the wheel is properly balanced, the ball has the same probability of falling into each slot. The number of the slot the ball falls into is said to "come up". There are a variety of bets available.

## 1.1 SOME USEFUL MATHEMATICAL FACTS

The gamma function  $\Gamma(x)$  is defined by a series of steps. First, we have that for  $n$  an integer,

$$\Gamma(n) = (n - 1)!$$

and then for  $z$  a complex number with positive real part (which includes positive real numbers), we have

$$\Gamma(z) = \int_0^\infty t^z \frac{e^{-t}}{t} dt.$$

By doing this, we get a function on positive real numbers that is a smooth interpolate of the factorial function. We won't do any real work with this function, so won't expand on this definition. In practice, we'll either look up a value in tables or require a software environment to produce it.

## 1.2 ACKNOWLEDGEMENTS

Typos spotted by: Han Chen (numerous!), Henry Lin (numerous!), Eric Huber, Brian Lunt, Yusuf Sobh, Scott Walters, — Your Name Here — TA's for this course have helped improve the notes. Thanks to Zicheng Liao, Michael Sittig, Nikita Spirin, Saurabh Singh, Daphne Tsatsoulis, Henry Lin, Karthik Ramaswamy.

## CHAPTER 2

# First Tools for Looking at Data

The single most important question for a working scientist — perhaps the single most useful question anyone can ask — is: “what’s going on here?” Answering this question requires creative use of different ways to make pictures of datasets, to summarize them, and to expose whatever structure might be there. This is an activity that is sometimes known as “Descriptive Statistics”. There isn’t any fixed recipe for understanding a dataset, but there is a rich variety of tools we can use to get insights.

### 2.1 DATASETS

A dataset is a collection of descriptions of different instances of the same phenomenon. These descriptions could take a variety of forms, but it is important that they are descriptions of the same thing. For example, my grandfather collected the daily rainfall in his garden for many years; we could collect the height of each person in a room; or the number of children in each family on a block; or whether 10 classmates would prefer to be “rich” or “famous”. There could be more than one description recorded for each item. For example, when he recorded the contents of the rain gauge each morning, my grandfather could have recorded (say) the temperature and barometric pressure. As another example, one might record the height, weight, blood pressure and body temperature of every patient visiting a doctor’s office.

The descriptions in a dataset can take a variety of forms. A description could be **categorical**, meaning that each data item can take a small set of prescribed values. For example, we might record whether each of 100 passers-by preferred to be “Rich” or “Famous”. As another example, we could record whether the passers-by are “Male” or “Female”. Categorical data could be **ordinal**, meaning that we can tell whether one data item is larger than another. For example, a dataset giving the number of children in a family for some set of families is categorical, because it uses only non-negative integers, but it is also ordinal, because we can tell whether one family is larger than another.

Some ordinal categorical data appears not to be numerical, but can be assigned a number in a reasonably sensible fashion. For example, many readers will recall being asked by a doctor to rate their pain on a scale of 1 to 10 — a question that is usually relatively easy to answer, but is quite strange when you think about it carefully. As another example, we could ask a set of users to rate the usability of an interface in a range from “very bad” to “very good”, and then record that using -2 for “very bad”, -1 for “bad”, 0 for “neutral”, 1 for “good”, and 2 for “very good”.

Many interesting datasets involve **continuous** variables (like, for example, height or weight or body temperature) when you could reasonably expect to encounter any value in a particular range. For example, we might have the heights of

all people in a particular room; or the rainfall at a particular place for each day of the year; or the number of children in each family on a list.

You should think of a dataset as a collection of  $d$ -tuples (a  $d$ -tuple is an ordered list of  $d$  elements). Tuples differ from vectors, because we can always add and subtract vectors, but we cannot necessarily add or subtract tuples. We will always write  $N$  for the number of tuples in the dataset, and  $d$  for the number of elements in each tuple. The number of elements will be the same for every tuple, though sometimes we may not know the value of some elements in some tuples (which means we must figure out how to predict their values, which we will do much later).

Index	net worth	Index	Taste score	Index	Taste score
1	100, 360	1	12.3	11	34.9
2	109, 770	2	20.9	12	57.2
3	96, 860	3	39	13	0.7
4	97, 860	4	47.9	14	25.9
5	108, 930	5	5.6	15	54.9
6	124, 330	6	25.9	16	40.9
7	101, 300	7	37.3	17	15.9
8	112, 710	8	21.9	18	6.4
9	106, 740	9	18.1	19	18
10	120, 170	10	21	20	38.9

TABLE 2.1: *On the left, net worths of people you meet in a bar, in US \$; I made this data up, using some information from the US Census. The index column, which tells you which data item is being referred to, is usually not displayed in a table because you can usually assume that the first line is the first item, and so on. On the right, the taste score (I'm not making this up; higher is better) for 20 different cheeses. This data is real (i.e. not made up), and it comes from <http://lib.stat.cmu.edu/DASL/Datafiles/Cheese.html>.*

Each element of a tuple has its own type. Some elements might be categorical. For example, one dataset we shall see several times has entries for Gender; Grade; Age; Race; Urban/Rural; School; Goals; Grades; Sports; Looks; and Money for 478 children, so  $d = 11$  and  $N = 478$ . In this dataset, each entry is categorical data. Clearly, these tuples are not vectors because one cannot add or subtract (say) Gender, or add Age to Grades.

Most of our data will be vectors. We use the same notation for a tuple and for a vector. We write a vector in bold, so  $\mathbf{x}$  could represent a vector or a tuple (the context will make it obvious which is intended).

The entire data set is  $\{\mathbf{x}\}$ . When we need to refer to the  $i$ 'th data item, we write  $\mathbf{x}_i$ . Assume we have  $N$  data items, and we wish to make a new dataset out of them; we write the dataset made out of these items as  $\{\mathbf{x}_i\}$  (the  $i$  is to suggest you are taking a set of items and making a dataset out of them).

In this chapter, we will work mainly with continuous data. We will see a variety of methods for plotting and summarizing 1-tuples. We can build these plots from a dataset of  $d$ -tuples by extracting the  $r$ 'th element of each  $d$ -tuple.

All through the book, we will see many datasets downloaded from various web sources, because people are so generous about publishing interesting datasets on the web. In the next chapter, we will look at 2-dimensional data, and we look at high dimensional data in chapter 10.

## 2.2 WHAT'S HAPPENING? - PLOTTING DATA

The very simplest way to present or visualize a dataset is to produce a table. Tables can be helpful, but aren't much use for large datasets, because it is difficult to get any sense of what the data means from a table. As a continuous example, table 2.1 gives a table of the net worth of a set of people you might meet in a bar (I made this data up). You can scan the table and have a rough sense of what is going on; net worths are quite close to \$ 100, 000, and there aren't any very big or very small numbers. This sort of information might be useful, for example, in choosing a bar.

People would like to measure, record, and reason about an extraordinary variety of phenomena. Apparently, one can score the goodness of the flavor of cheese with a number (bigger is better); table 2.1 gives a score for each of thirty cheeses (I did not make up this data, but downloaded it from <http://lib.stat.cmu.edu/DASL/Datafiles/Cheese.html>). You should notice that a few cheeses have very high scores, and most have moderate scores. It's difficult to draw more significant conclusions from the table, though.

Gender	Goal	Gender	Goal
boy	Sports	girl	Sports
boy	Popular	girl	Grades
girl	Popular	boy	Popular
girl	Popular	boy	Popular
girl	Popular	boy	Popular
girl	Popular	girl	Grades
girl	Popular	girl	Sports
girl	Grades	girl	Popular
girl	Sports	girl	Grades
girl	Sports	girl	Sports

TABLE 2.2: Chase and Dunner (?) collected data on what students thought made other students popular. As part of this effort, they collected information on (a) the gender and (b) the goal of students. This table gives the gender (“boy” or “girl”) and the goal (to make good grades — “Grades”; to be popular — “Popular”; or to be good at sports — “Sports”). The table gives this information for the first 20 of 478 students; the rest can be found at <http://lib.stat.cmu.edu/DASL/Datafiles/PopularKids.html>. This data is clearly categorical, and not ordinal.

Table 2.2 shows a table for a set of categorical data. Psychologists collected data from students in grades 4-6 in three school districts to understand what factors students thought made other students popular. This fascinating data set can be found at <http://lib.stat.cmu.edu/DASL/Datafiles/PopularKids.html>, and was prepared by Chase and Dunner (?). Among other things, for each student they asked whether the student's goal was to make good grades (“Grades”, for short); to be

popular (“Popular”); or to be good at sports (“Sports”). They have this information for 478 students, so a table would be very hard to read. Table 2.2 shows the gender and the goal for the first 20 students in this group. It’s rather harder to draw any serious conclusion from this data, because the full table would be so big. We need a more effective tool than eyeballing the table.

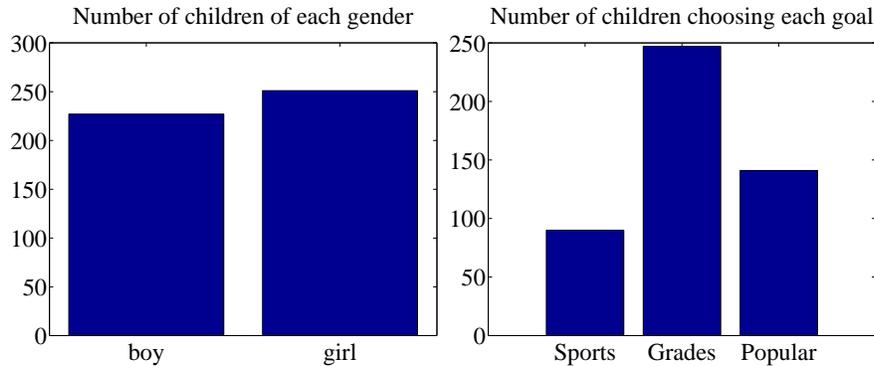


FIGURE 2.1: *On the left*, a bar chart of the number of children of each gender in the Chase and Dunner study (). Notice that there are about the same number of boys and girls (the bars are about the same height). *On the right*, a bar chart of the number of children selecting each of three goals. You can tell, at a glance, that different goals are more or less popular by looking at the height of the bars.

### 2.2.1 Bar Charts

A **bar chart** is a set of bars, one per category, where the height of each bar is proportional to the number of items in that category. A glance at a bar chart often exposes important structure in data, for example, which categories are common, and which are rare. Bar charts are particularly useful for categorical data. Figure 2.1 shows such bar charts for the genders and the goals in the student dataset of Chase and Dunner (). You can see at a glance that there are about as many boys as girls, and that there are more students who think grades are important than students who think sports or popularity is important. You couldn’t draw either conclusion from Table 2.2, because I showed only the first 20 items; but a 478 item table is very difficult to read.

### 2.2.2 Histograms

Data is continuous when a data item could take any value in some range or set of ranges. In turn, this means that we can reasonably expect a continuous dataset contains few or no pairs of items that have *exactly* the same value. Drawing a bar chart in the obvious way — one bar per value — produces a mess of unit height bars, and seldom leads to a good plot. Instead, we would like to have fewer bars, each representing more data items. We need a procedure to decide which data items count in which bar.

A simple generalization of a bar chart is a **histogram**. We divide the range

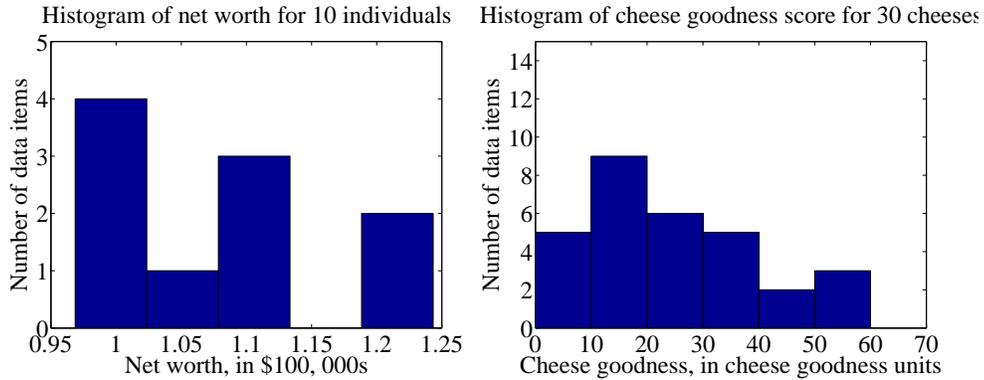


FIGURE 2.2: On the **left**, a histogram of net worths from the dataset described in the text and shown in table 2.1. On the **right**, a histogram of cheese goodness scores from the dataset described in the text and shown in table 2.1.

of the data into intervals, which do not need to be equal in length. We think of each interval as having an associated pigeonhole, and choose one pigeonhole for each data item. We then build a set of boxes, one per interval. Each box sits on its interval on the horizontal axis, and its height is determined by the number of data items in the corresponding pigeonhole. In the simplest histogram, the intervals that form the bases of the boxes are equally sized. In this case, the height of the box is given by the number of data items in the box.

Figure 2.2 shows a histogram of the data in table 2.1. There are five bars — by my choice; I could have plotted ten bars — and the height of each bar gives the number of data items that fall into its interval. For example, there is one net worth in the range between \$102,500 and \$107,500. Notice that one bar is invisible, because there is no data in that range. This picture suggests conclusions consistent with the ones we had from eyeballing the table — the net worths tend to be quite similar, and around \$100,000.

Figure 2.2 also shows a histogram of the data in table 2.1. There are six bars (0-10, 10-20, and so on), and the height of each bar gives the number of data items that fall into its interval — so that, for example, there are 9 cheeses in this dataset whose score is greater than or equal to 10 and less than 20. You can also use the bars to estimate other properties. So, for example, there are 14 cheeses whose score is less than 20, and 3 cheeses with a score of 50 or greater. This picture is much more helpful than the table; you can see at a glance that quite a lot of cheeses have relatively low scores, and few have high scores.

### 2.2.3 How to Make Histograms

Usually, one makes a histogram by finding the appropriate command or routine in your programming environment. I use Matlab, and chapter ?? sketches some useful Matlab commands. However, it is useful to understand the procedures used.

**Histograms with Even Intervals:** The easiest histogram to build uses equally sized intervals. Write  $x_i$  for the  $i$ 'th number in the dataset,  $x_{\min}$  for the

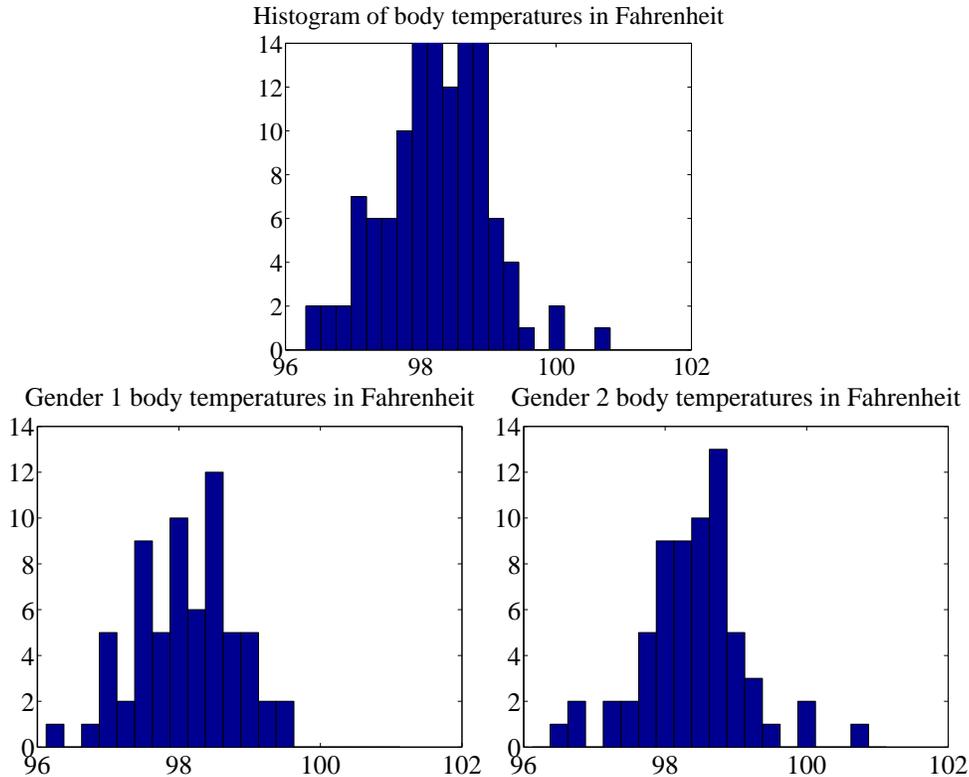


FIGURE 2.3: On **top**, a histogram of body temperatures, from the dataset published at <http://www2.stetson.edu/~jrasp/data.htm>. These seem to be clustered fairly tightly around one value. The **bottom row** shows histograms for each gender (I don't know which is which). It looks as though one gender runs slightly cooler than the other.

smallest value, and  $x_{\max}$  for the largest value. We divide the range between the smallest and largest values into  $n$  intervals of even width  $(x_{\max} - x_{\min})/n$ . In this case, the height of each box is given by the number of items in that interval. We could represent the histogram with an  $n$ -dimensional vector of counts. Each entry represents the count of the number of data items that lie in that interval. Notice we need to be careful to ensure that each point in the range of values is claimed by exactly one interval. For example, we could have intervals of  $[0 - 1)$  and  $[1 - 2)$ , or we could have intervals of  $(0 - 1]$  and  $(1 - 2]$ . We could *not* have intervals of  $[0 - 1]$  and  $[1 - 2]$ , because then a data item with the value 1 would appear in two boxes. Similarly, we could not have intervals of  $(0 - 1)$  and  $(1 - 2)$ , because then a data item with the value 1 would not appear in any box.

**Histograms with Uneven Intervals:** For a histogram with even intervals, it is natural that the height of each box is the number of data items in that box. But a histogram with even intervals can have empty boxes (see figure 2.2). In this case, it can be more informative to have some larger intervals to ensure that

each interval has some data items in it. But how high should we plot the box? Imagine taking two consecutive intervals in a histogram with even intervals, and fusing them. It is natural that the height of the fused box should be the average height of the two boxes. This observation gives us a rule.

Write  $dx$  for the width of the intervals;  $n_1$  for the height of the box over the first interval (which is the number of elements in the first box); and  $n_2$  for the height of the box over the second interval. The height of the fused box will be  $(n_1 + n_2)/2$ . Now the *area* of the first box is  $n_1 dx$ ; of the second box is  $n_2 dx$ ; and of the fused box is  $(n_1 + n_2) dx$ . For each of these boxes, the *area* of the box is proportional to the number of elements in the box. This gives the correct rule: plot boxes such that the area of the box is proportional to the number of elements in the box.

#### 2.2.4 Conditional Histograms

Most people believe that normal body temperature is  $98.4^\circ$  in Fahrenheit. If you take other people's temperatures often (for example, you might have children), you know that some individuals tend to run a little warmer or a little cooler than this number. I found data giving the body temperature of a set of individuals at <http://www2.stetson.edu/~jrasp/data.htm>. As you can see from the histogram (figure 2.3), the body temperatures cluster around a small set of numbers. But what causes the variation?

One possibility is gender. We can investigate this possibility by comparing a histogram of temperatures for males with histogram of temperatures for females. Such histograms are sometimes called **conditional histograms** or **class-conditional histograms**, because each histogram is conditioned on something (in this case, the histogram uses only data that comes from gender).

The dataset gives genders (as 1 or 2 - I don't know which is male and which female). Figure 2.3 gives the class conditional histograms. It does seem like individuals of one gender run a little cooler than individuals of the other, although we don't yet have mechanisms to test this possibility in detail (chapter 13.1).

### 2.3 SUMMARIZING 1D DATA

For the rest of this chapter, we will assume that data items take values that are continuous real numbers. Furthermore, we will assume that values can be added, subtracted, and multiplied by constants in a meaningful way. Human heights are one example of such data; you can add two heights, and interpret the result as a height (perhaps one person is standing on the head of the other). You can subtract one height from another, and the result is meaningful. You can multiply a height by a constant — say,  $1/2$  — and interpret the result (A is half as high as B).

#### 2.3.1 The Mean

One simple and effective summary of a set of data is its **mean**. This is sometimes known as the **average** of the data.

**Definition: 2.1** *Mean*

Assume we have a dataset  $\{x\}$  of  $N$  data items,  $x_1, \dots, x_N$ . Their mean is

$$\text{mean}(\{x\}) = \frac{1}{N} \sum_{i=1}^{i=N} x_i.$$

Indexing: Mean

Indexing: Mean

For example, assume you're in a bar, in a group of ten people who like to talk about money. They're average people, and their net worth is given in table 2.1 (you can choose who you want to be in this story). The mean of this data is \$107, 903.

**Properties of the Mean** The mean has several important properties you should remember:

- Scaling data scales the mean: or  $\text{mean}(\{kx_i\}) = k\text{mean}(\{x_i\})$ .
- Translating data translates the mean: or  $\text{mean}(\{x_i + c\}) = \text{mean}(\{x_i\}) + c$ .
- The sum of signed differences from the mean is zero. This means that

$$\sum_{i=1}^N (x_i - \text{mean}(\{x_i\})) = 0.$$

- Choose the number  $\mu$  such that the sum of squared distances of data points to  $\mu$  is minimized. That number is the mean. In notation

$$\arg \min_{\mu} \sum_i (x_i - \mu)^2 = \text{mean}(\{x_i\})$$

These properties are easy to prove (and so easy to remember). I have broken these out into a box of useful facts below, to emphasize them. All but one proof is relegated to the exercises.

**Proposition:**  $\arg \min_{\mu} \sum_i (x_i - \mu)^2 = \text{mean}(\{x\})$

**Proof:** Choose the number  $\mu$  such that the sum of squared distances of data points to  $\mu$  is minimized. That number is the mean. In notation:

$$\arg \min_{\mu} \sum_i (x_i - \mu)^2 = \text{mean}(\{x\})$$

We can show this by actually minimizing the expression. We must have that the derivative of the expression we are minimizing is zero at the value of  $\mu$  we are seeking. So we have

$$\begin{aligned} \frac{d}{d\mu} \sum_{i=1}^N (x_i - \mu)^2 &= \sum_{i=1}^N 2(x_i - \mu) \\ &= 2 \sum_{i=1}^N (x_i - \mu) \\ &= 0 \end{aligned}$$

so that  $2N\text{mean}(\{x\}) - 2N\mu = 0$ , which means that  $\mu = \text{mean}(\{x\})$ .

**Property 2.1:** The Average Squared Distance to the Mean is Minimized

Now this result means that the mean is the single number that is closest to all the data items. The mean tells you where the overall blob of data lies. For this reason, it is often referred to as a “location parameter”. If you choose to summarize the dataset with a number that is as close as possible to each data item, the mean is the number to choose. The mean is also a guide to what new values will look like, if you have no other information. For example, in the case of the bar, a new person walks in, and I must guess that person’s net worth. Then the mean is the best guess, because it is closest to all the data items we have already seen. In the case of the bar, if a new person walked into this bar, and you had to guess that person’s net worth, you should choose \$107, 903.

**Useful Facts: 2.1** *Mean*

- $\text{mean}(\{kx_i\}) = k\text{mean}(\{x_i\})$ .
- $\text{mean}(\{x_i + c\}) = \text{mean}(\{x_i\}) + c$ .
- $[\sum_{i=1}^N (x_i - \text{mean}(\{x_i\}))] = 0$ .

•

$$\arg \min_{\mu} \sum_i (x_i - \mu)^2 = \text{mean}(\{x_i\})$$

The mean is a location parameter; it tells you where the data lies along a number line.

Indexing: Mean

Indexing: Mean

## 2.3.2 Standard Deviation and Variance

We would also like to know the extent to which data items are close to the mean. This information is given by the **standard deviation**, which is the root mean square of the offsets of data from the mean.

**Definition: 2.2** *Standard deviation*

Assume we have a dataset  $\{x\}$  of  $N$  data items,  $x_1, \dots, x_N$ . The standard deviation of this dataset is:

$$\text{std}(\{x_i\}) = \sqrt{\frac{1}{N} \sum_{i=1}^N (x_i - \text{mean}(\{x\}))^2} = \sqrt{\text{mean}(\{(x_i - \text{mean}(\{x\}))^2\})}$$

Indexing: Standard deviation

Indexing: Standard deviation

You should think of the standard deviation as a scale. It measures the size of the average deviation from the mean for a dataset, or how wide the spread of data is. For this reason, it is often referred to as a scale parameter. When the standard deviation of a dataset is large, there are many items with values much larger than, or much smaller than, the mean. When the standard deviation is small, most data items have values close to the mean. This means it is helpful to talk about how many standard deviations away from the mean a particular data item is. Saying that data item  $x_j$  is “within  $k$  standard deviations from the mean” means that

$$\text{abs}(x_j - \text{mean}(\{x\})) \leq k\text{std}(\{x_i\}).$$

Similarly, saying that data item  $x_j$  is “more than  $k$  standard deviations from the mean” means that

$$\text{abs}(x_i - \text{mean}(\{x\})) > k\text{std}(\{x\}).$$

As I will show below, there must be some data at least one standard deviation away from the mean, and there can be very few data items that are many standard deviations away from the mean.

**Properties of the Standard Deviation** Standard deviation has very important properties:

- Translating data does not change the standard deviation, i.e.  $\text{std}(\{x_i + c\}) = \text{std}(\{x_i\})$ .
- Scaling data scales the standard deviation, i.e.  $\text{std}(\{kx_i\}) = k\text{std}(\{x_i\})$ .
- For any dataset, there can be only a few items that are many standard deviations away from the mean. In particular, assume we have  $N$  data items,  $x_i$ , whose standard deviation is  $\sigma$ . Then there are at most  $\frac{1}{k^2}$  data points lying  $k$  or more standard deviations away from the mean.
- For any dataset, there must be at least one data item that is at least one standard deviation away from the mean.

The first two properties are easy to prove, and are relegated to the exercises. I prove the others below. Again, for emphasis, I have broken these properties out in a box below.

**Proposition:** Assume we have a dataset  $\{x\}$  of  $N$  data items,  $x_1, \dots, x_N$ . Assume the standard deviation of this dataset is  $\text{std}(\{x\}) = \sigma$ . Then there are at most  $\frac{1}{k^2}$  data points lying  $k$  or more standard deviations away from the mean.

**Proof:** Assume the mean is zero. There is no loss of generality here, because translating data translates the mean, but doesn't change the standard deviation. Now we must construct a dataset with the largest possible fraction  $r$  of data points lying  $k$  or more standard deviations from the mean. To achieve this, our data should have  $N(1 - r)$  data points each with the value 0, because these contribute 0 to the standard deviation. It should have  $Nr$  data points with the value  $k\sigma$ ; if they are further from zero than this, each will contribute more to the standard deviation, so the fraction of such points will be fewer. Because

$$\text{std}(\{x\}) = \sigma = \sqrt{\frac{\sum_i x_i^2}{N}}$$

we have that, for this rather specially constructed dataset,

$$\sigma = \sqrt{\frac{Nrk^2\sigma^2}{N}}$$

so that

$$r = \frac{1}{k^2}.$$

We constructed the dataset so that  $r$  would be as large as possible, so

$$r \leq \frac{1}{k^2}$$

for any kind of data at all.

**Property 2.2:** For any dataset, it is hard for data items to get many standard deviations away from the mean.

This bound (box 2.2) is true for *any kind of data*. This bound implies that, for example, at most 100% of *any* dataset could be one standard deviation away from the mean, 25% of *any* dataset is 2 standard deviations away from the mean and at most 11% of *any* dataset could be 3 standard deviations away from the mean. But the configuration of data that achieves this bound is very unusual. This means the bound tends to wildly *overstate* how much data is far from the mean for most practical datasets. Most data has more random structure, meaning that we expect to see very much *less* data far from the mean than the bound predicts. For example, much data can reasonably be modelled as coming from a normal distribution (a topic we'll go into later). For such data, we expect that about 68% of the data is within one standard deviation of the mean, 95% is within two standard deviations of the mean, and 99.7% is within three standard deviations of the mean, and the percentage of data that is within ten standard deviations of the mean is essentially indistinguishable from 100%. This kind of behavior is quite common; the crucial point about the standard deviation is that you won't see much

data that lies many standard deviations from the mean, because you can't.

**Proposition:**

$$(\text{std}(\{x\}))^2 \leq \max_i (x_i - \text{mean}(\{x\}))^2.$$

**Proof:** You can see this by looking at the expression for standard deviation. We have

$$\text{std}(\{x\}) = \sqrt{\frac{1}{N} \sum_{i=1}^{i=N} (x_i - \text{mean}(\{x\}))^2}.$$

Now, this means that

$$N(\text{std}(\{x\}))^2 = \sum_{i=1}^{i=N} (x_i - \text{mean}(\{x\}))^2.$$

But

$$\sum_{i=1}^{i=N} (x_i - \text{mean}(\{x\}))^2 \leq N \max_i (x_i - \text{mean}(\{x\}))^2$$

so

$$(\text{std}(\{x\}))^2 \leq \max_i (x_i - \text{mean}(\{x\}))^2.$$

**Property 2.3:** For any dataset, there must be at least one data item that is at least one standard deviation away from the mean.

Boxes 2.2 and 2.3 mean that the standard deviation is quite informative. Very little data is many standard deviations away from the mean; similarly, at least some of the data should be one or more standard deviations away from the mean. So the standard deviation tells us how data points are scattered about the mean.

**Useful Facts: 2.2** *Standard deviation*

- $\text{std}(\{x + c\}) = \text{std}(\{x\})$ .
- $\text{std}(\{kx\}) = k \text{std}(\{x\})$ .
- For  $N$  data items,  $x_i$ , whose standard deviation is  $\sigma$ , there are at most  $\frac{1}{k^2}$  data points lying  $k$  or more standard deviations away from the mean.
- $(\text{std}(\{x\}))^2 \leq \max_i (x_i - \text{mean}(\{x\}))^2$ .

The standard deviation is often referred to as a scale parameter; it tells you how broadly the data spreads about the mean.

Indexing: Standard deviation

Indexing: Standard deviation

There is an ambiguity that comes up often here because two (very slightly) different numbers are called the standard deviation of a dataset. One — the one we use in this chapter — is an estimate of the scale of the data, as we describe it. The other differs from our expression very slightly; one computes

$$\sqrt{\frac{\sum_i (x_i - \text{mean}(\{x\}))^2}{N - 1}}$$

(notice the  $N - 1$  for our  $N$ ). If  $N$  is large, this number is basically the same as the number we compute, but for smaller  $N$  there is a difference that can be significant. Irritatingly, this number is also called the standard deviation; even more irritatingly, we will have to deal with it, but not yet. I mention it now because you may look up terms I have used, find this definition, and wonder whether I know what I'm talking about. In this case, I do (although I would say that).

The confusion arises because sometimes the datasets we see are actually samples of larger datasets. For example, in some circumstances you could think of the net worth dataset as a sample of all the net worths in the USA. In such cases, we are often interested in the standard deviation *of the dataset that was sampled*. The second number is a slightly better way to estimate this standard deviation than the definition we have been working with. Don't worry - the  $N$  in our expressions is the right thing to use for what we're doing.

### 2.3.3 Variance

It turns out that thinking in terms of the square of the standard deviation, which is known as the **variance**, will allow us to generalize our summaries to apply to higher dimensional data.

**Definition: 2.3** *Variance*

Assume we have a dataset  $\{x\}$  of  $N$  data items,  $x_1, \dots, x_N$ . where  $N > 1$ . Their variance is:

$$\text{var}(\{x\}) = \frac{1}{N} \left( \sum_{i=1}^{i=N} (x_i - \text{mean}(\{x\}))^2 \right) = \text{mean}(\{(x_i - \text{mean}(\{x\}))^2\}).$$

Indexing: Variance

Indexing: Variance

One good way to think of the variance is as the mean-square error you would incur if you replaced each data item with the mean. Another is that it is the square of the standard deviation.

**Properties of the Variance** The properties of the variance follow from the fact that it is the square of the standard deviation. I have broken these out in a box, for emphasis.

**Useful Facts: 2.3** *Variance*

- $\text{var}(\{x + c\}) = \text{var}(\{x\})$ .
- $\text{var}(\{kx\}) = k^2 \text{var}(\{x\})$ .

Indexing: Variance

Indexing: Variance

While one could restate the other two properties of the standard deviation in terms of the variance, it isn't really natural to do so. The standard deviation is in the same units as the original data, and should be thought of as a scale. Because the variance is the square of the standard deviation, it isn't a natural scale (unless you take its square root!).

#### 2.3.4 The Median

One problem with the mean is that it can be affected strongly by extreme values. Go back to the bar example, of section 2.3.1. Now Warren Buffett (or Bill Gates, or your favorite billionaire) walks in. What happened to the average net worth?

Assume your billionaire has net worth \$ 1, 000, 000, 000. Then the mean net worth suddenly has become

$$\frac{10 \times \$107,903 + \$1,000,000,000}{11} = \$91,007,184$$

But this mean isn't a very helpful summary of the people in the bar. It is probably more useful to think of the net worth data as ten people together with one billionaire. The billionaire is known as an **outlier**.

One way to get outliers is that a small number of data items are very different, due to minor effects you don't want to model. Another is that the data was misrecorded, or mistranscribed. Another possibility is that there is just too much variation in the data to summarize it well. For example, a small number of extremely wealthy people could change the average net worth of US residents dramatically, as the example shows. An alternative to using a mean is to use a **median**.

**Definition: 2.4** *Median*

The median of a set of data points is obtained by sorting the data points, and finding the point halfway along the list. If the list is of even length, it's usual to average the two numbers on either side of the middle. We write

$$\text{median}(\{x_i\})$$

for the operator that returns the median.

Indexing: Median

Indexing: Median

For example,

$$\begin{aligned}\text{median}(\{3, 5, 7\}) &= 5, \\ \text{median}(\{3, 4, 5, 6, 7\}) &= 5,\end{aligned}$$

and

$$\text{median}(\{3, 4, 5, 6\}) = 4.5.$$

For much, but not all, data, you can expect that roughly half the data is smaller than the median, and roughly half is larger than the median. Sometimes this property fails. For example,

$$\text{median}(\{1, 2, 2, 2, 2, 2, 2, 3\}) = 2.$$

With this definition, the median of our list of net worths is \$107,835. If we insert the billionaire, the median becomes \$108,930. Notice by how little the number has changed — it remains an effective summary of the data.

**Properties of the median** You can think of the median of a dataset as giving the “middle” or “center” value. It is another way of estimating where the dataset lies on a number line (and so is another location parameter). This means it is rather like the mean, which also gives a (slightly differently defined) “middle” or “center” value. The mean has the important properties that if you translate the dataset, the mean translates, and if you scale the dataset, the mean scales. The median has these properties, too, which I have broken out in a box. Each is easily proved, and proofs are relegated to the exercises.

**Useful Facts: 2.4** *Median*

- $\text{median}(\{x + c\}) = \text{median}(\{x\}) + c.$
- $\text{median}(\{kx\}) = k\text{median}(\{x\}).$

Indexing: Median

Indexing: Median

## 2.3.5 Interquartile Range

Outliers are a nuisance in all sorts of ways. Plotting the histogram of the net worth data with the billionaire included will be tricky. Either you leave the billionaire out of the plot, or all the histogram bars are tiny. Visualizing this plot shows outliers can affect standard deviations severely, too. For our net worth data, the standard deviation without the billionaire is \$9265, but if we put the billionaire in there, it is  $\$3.014 \times 10^8$ . When the billionaire is in the dataset, the mean is about 91M\$ and the standard deviation is about 300M\$ — so all but one of the data items lie about a third of a standard deviation away from the mean on the small side. The other data item (the billionaire) is about three standard deviations away from the mean on the large side. In this case, the standard deviation has done its work of informing us that there are huge changes in the data, but isn't really helpful as a description of the data.

The problem is this: describing the net worth data with billionaire as a having a mean of  $\$9.101 \times 10^7$  with a standard deviation of  $\$3.014 \times 10^8$  really isn't terribly helpful. Instead, the data really should be seen as a clump of values that are near \$100,000 and moderately close to one another, and one massive number (the billionaire outlier).

One thing we could do is simply remove the billionaire and compute mean and standard deviation. This isn't always easy to do, because it's often less obvious which points are outliers. An alternative is to follow the strategy we did when we used the median. Find a summary that describes scale, but is less affected by outliers than the standard deviation. This is the **interquartile range**; to define it, we need to define percentiles and quartiles, which are useful anyway.

**Definition: 2.5** *Percentile*

The  $k$ 'th percentile is the value such that  $k\%$  of the data is less than or equal to that value. We write  $\text{percentile}(\{x\}, k)$  for the  $k$ 'th percentile of dataset  $\{x\}$ .

Indexing: Percentile

Indexing: Percentile

**Definition: 2.6** *Quartiles*

The first quartile of the data is the value such that 25% of the data is less than or equal to that value (i.e.  $\text{percentile}(\{x\}, 25)$ ). The second quartile of the data is the value such that 50% of the data is less than or equal to that value, which is usually the median (i.e.  $\text{percentile}(\{x\}, 50)$ ). The third quartile of the data is the value such that 75% of the data is less than or equal to that value (i.e.  $\text{percentile}(\{x\}, 75)$ ).

Indexing: Quartiles

Indexing: Quartiles

**Definition: 2.7** *Interquartile Range*

The interquartile range of a dataset  $\{x\}$  is  $\text{iqr}\{x\} = \text{percentile}(\{x\}, 75) - \text{percentile}(\{x\}, 25)$ .

Indexing: Interquartile Range

Indexing: Interquartile Range

Like the standard deviation, the interquartile range gives an estimate of how widely the data is spread out. But it is quite well-behaved in the presence of outliers. For our net worth data without the billionaire, the interquartile range is \$12350; with the billionaire, it is \$17710.

**Properties of the interquartile range** You can think of the interquartile range of a dataset as giving an estimate of the scale of the difference from the mean. This means it is rather like the standard deviation, which also gives a (slightly differently defined) scale. The standard deviation has the important properties that if you translate the dataset, the standard deviation translates, and if you scale the dataset, the standard deviation scales. The interquartile range has these properties, too, which I have broken out into a box. Each is easily proved, and proofs are relegated to the exercises.

**Useful Facts: 2.5** *Interquartile range*

- $\text{iqr}\{x + c\} = \text{iqr}\{x\}$ .
- $\text{iqr}\{kx\} = k^2 \text{iqr}\{x\}$ .

Indexing: Interquartile range

Indexing: Interquartile range

For most datasets, interquartile ranges tend to be somewhat larger than standard deviations. This isn't really a problem. Each is a method for estimating the scale of the data — the range of values above and below the mean that you are likely to see. It is neither here nor there if one method yields slightly larger estimates than another, as long as you don't compare estimates across methods.

## 2.3.6 Using Summaries Sensibly

One should be careful how one summarizes data. For example, the statement that “the average US family has 2.6 children” invites mockery (the example is from Andrew Vickers' book *What is a p-value anyway?*), because you can't have fractions

of a child — no family has 2.6 children. A more accurate way to say things might be “the average of the number of children in a US family is 2.6”, but this is clumsy. What is going wrong here is the 2.6 is a mean, but the number of children in a family is a categorical variable. Reporting the mean of a categorical variable is often a bad idea, because you may never encounter this value (the 2.6 children). For a categorical variable, giving the median value and perhaps the interquartile range often makes much more sense than reporting the mean.

For continuous variables, reporting the mean is reasonable because you could expect to encounter a data item with this value, even if you haven’t seen one in the particular data set you have. It is sensible to look at both mean and median; if they’re significantly different, then there is probably something going on that is worth understanding. You’d want to plot the data using the methods of the next section before you decided what to report.

You should also be careful about how precisely numbers are reported (equivalently, the number of significant figures). Numerical and statistical software will produce very large numbers of digits freely, but not all are always useful. This is a particular nuisance in the case of the mean, because you might add many numbers, then divide by a large number; in this case, you will get many digits, but some might not be meaningful. For example, Vickers (ibid) describes a paper reporting the mean length of pregnancy as 32.833 weeks. That fifth digit suggests we know the mean length of pregnancy to about 0.001 weeks, or roughly 10 minutes. Neither medical interviewing nor people’s memory for past events is that detailed. Furthermore, when you interview them about embarrassing topics, people quite often lie. There is no prospect of knowing this number with this precision.

People regularly report silly numbers of digits because it is easy to miss the harm caused by doing so. But the harm is there: you are implying to other people, and to yourself, that you know something more accurately than you do. At some point, someone will suffer for it.

## 2.4 PLOTS AND SUMMARIES

Knowing the mean, standard deviation, median and interquartile range of a dataset gives us some information about what its histogram might look like. In fact, the summaries give us a language in which to describe a variety of characteristic properties of histograms that are worth knowing about (Section 2.4.1). Quite remarkably, many different datasets have histograms that have about the same shape (Section 2.4.2). For such data, we know roughly what percentage of data items are how far from the mean.

Complex datasets can be difficult to interpret with histograms alone, because it is hard to compare many histograms by eye. Section 2.4.3 describes a clever plot of various summaries of datasets that makes it easier to compare many cases.

### 2.4.1 Some Properties of Histograms

The **tails** of a histogram are the relatively uncommon values that are significantly larger (resp. smaller) than the value at the peak (which is sometimes called the **mode**). A histogram is **unimodal** if there is only one peak; if there are more than one, it is **multimodal**, with the special term **bimodal** sometimes being used for

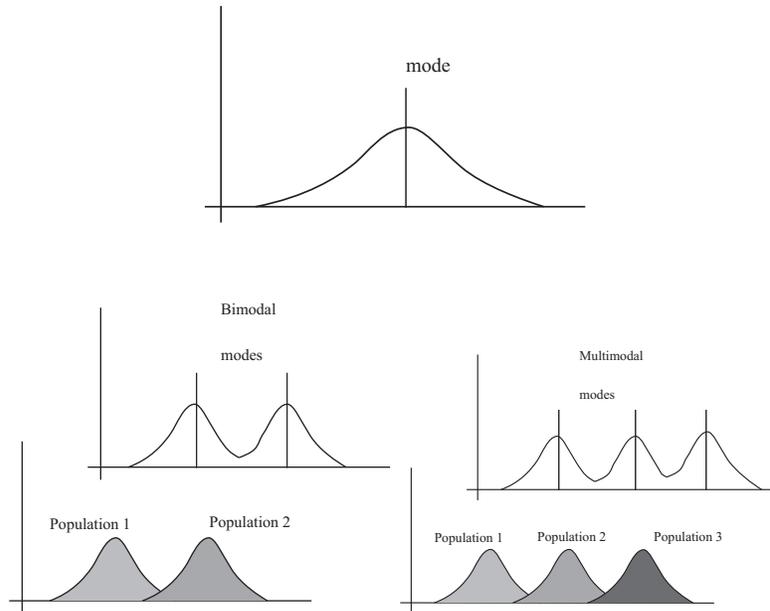


FIGURE 2.4: Many histograms are unimodal, like the example on the **top**; there is one peak, or mode. Some are bimodal (two peaks; **bottom left**) or even multimodal (two or more peaks; **bottom right**). One common reason (but not the only reason) is that there are actually two populations being conflated in the histograms. For example, measuring adult heights might result in a bimodal histogram, if male and female heights were slightly different. As another example, measuring the weight of dogs might result in a multimodal histogram if you did not distinguish between breeds (eg chihuahua, terrier, german shepherd, pyrenean mountain dog, etc.).

the case where there are two peaks (Figure 2.4). The histograms we have seen have been relatively symmetric, where the left and right tails are about as long as one another. Another way to think about this is that values a lot larger than the mean are about as common as values a lot smaller than the mean. Not all data is symmetric. In some datasets, one or another tail is longer (figure 2.5). This effect is called **skew**.

Skew appears often in real data. SOCR (the Statistics Online Computational Resource) publishes a number of datasets. Here we discuss a dataset of citations to faculty publications. For each of five UCLA faculty members, SOCR collected the number of times each of the papers they had authored had been cited by other authors (data at [http://wiki.stat.ucla.edu/socr/index.php/SOCR\\_Data\\_Dinov\\_072108\\_H\\_Index\\_Pubs](http://wiki.stat.ucla.edu/socr/index.php/SOCR_Data_Dinov_072108_H_Index_Pubs)). Generally, a small number of papers get many citations, and many papers get few citations. We see this pattern in the histograms of citation numbers (figure 2.6). These are very different from (say) the body temperature pictures. In the citation histograms, there are many data items that have very few citations, and few that have many citations. This means that the right tail of the histogram is longer, so the histogram is skewed to the right.

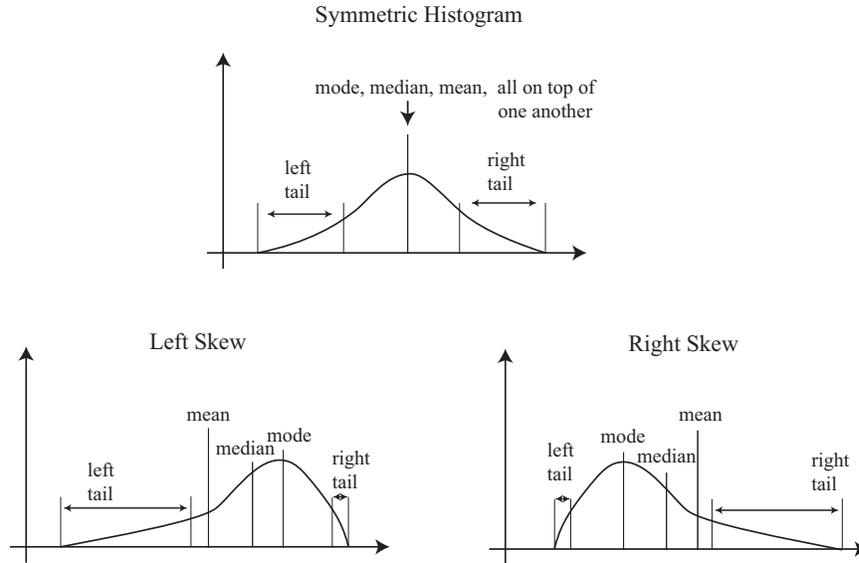


FIGURE 2.5: On the **top**, an example of a symmetric histogram, showing its tails (relatively uncommon values that are significantly larger or smaller than the peak or mode). **Lower left**, a sketch of a left-skewed histogram. Here there are few large values, but some very small values that occur with significant frequency. We say the left tail is “long”, and that the histogram is left skewed. You may find this confusing, because the main bump is to the right — one way to remember this is that the left tail has been stretched. **Lower right**, a sketch of a right-skewed histogram. Here there are few small values, but some very large values that occur with significant frequency. We say the right tail is “long”, and that the histogram is right skewed.

One way to check for skewness is to look at the histogram; another is to compare mean and median (though this is not foolproof). For the first citation histogram, the mean is 24.7 and the median is 7.5; for the second, the mean is 24.4, and the median is 11. In each case, the mean is a lot bigger than the median. Recall the definition of the median (form a ranked list of the data points, and find the point halfway along the list). For much data, the result is larger than about half of the data set and smaller than about half the dataset. So if the median is quite small compared to the mean, then there are many small data items and a small number of data items that are large — the right tail is longer, so the histogram is skewed to the right.

Left-skewed data also occurs; figure 2.6 shows a histogram of the birth weights of 44 babies born in Brisbane, in 1997 (from [http://www.amstat.org/publications/jse/jse\\_data\\_archive.htm](http://www.amstat.org/publications/jse/jse_data_archive.htm)). This data appears to be somewhat left-skewed, as birth weights can be a lot smaller than the mean, but tend not to be much larger than the mean.

Skewed data is often, but not always, the result of constraints. For example,

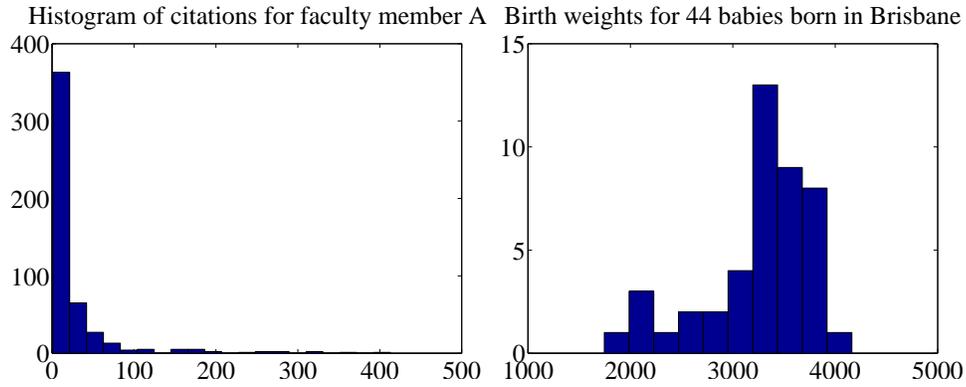


FIGURE 2.6: *On the left*, a histogram of citations for a faculty member, from data at [http://wiki.stat.ucla.edu/socr/index.php/SOCR\\_Data\\_Dinov\\_072108\\_H\\_Index\\_Pubs](http://wiki.stat.ucla.edu/socr/index.php/SOCR_Data_Dinov_072108_H_Index_Pubs). Very few publications have many citations, and many publications have few. This means the histogram is strongly right-skewed. *On the right*, a histogram of birth weights for 44 babies borne in Brisbane in 1997. This histogram looks slightly left-skewed.

good obstetrical practice tries to ensure that very large birth weights are rare (birth is typically induced before the baby gets too heavy), but it may be quite hard to avoid some small birth weights. This could skew birth weights to the left (because large babies will get born, but will not be as heavy as they could be if obstetricians had not interfered). Similarly, income data can be skewed to the right by the fact that income is always positive. Test mark data is often skewed — whether to right or left depends on the circumstances — by the fact that there is a largest possible mark and a smallest possible mark.

#### 2.4.2 Standard Coordinates and Normal Data

It is useful to look at lots of histograms, because it is often possible to get some useful insights about data. However, in their current form, histograms are hard to compare. This is because each is in a different set of units. A histogram for length data will consist of boxes whose horizontal units are, say, metres; a histogram for mass data will consist of boxes whose horizontal units are in, say, kilograms. Furthermore, these histograms typically span different ranges.

We can make histograms comparable by (a) estimating the “location” of the plot on the horizontal axis and (b) estimating the “scale” of the plot. The location is given by the mean, and the scale by the standard deviation. We could then normalize the data by subtracting the location (mean) and dividing by the standard deviation (scale). The resulting values are unitless, and have zero mean. They are often known as **standard coordinates**.

**Definition: 2.8** *Standard coordinates*

Assume we have a dataset  $\{x\}$  of  $N$  data items,  $x_1, \dots, x_N$ . We represent these data items in standard coordinates by computing

$$\hat{x}_i = \frac{(x_i - \text{mean}(\{x\}))}{\text{std}(\{x\})}.$$

We write  $\{\hat{x}\}$  for a dataset that happens to be in standard coordinates.

Indexing: Standard coordinates

Indexing: Standard coordinates

Standard coordinates have some important properties. Assume we have  $N$  data items. Write  $x_i$  for the  $i$ 'th data item, and  $\hat{x}_i$  for the  $i$ 'th data item in standard coordinates (I sometimes refer to these as “normalized data items”). Then we have

$$\text{mean}(\{\hat{x}\}) = 0.$$

We also have that

$$\text{std}(\{\hat{x}\}) = 1.$$

An extremely important fact about data is that, for many kinds of data, histograms of these standard coordinates look the same. Many completely different datasets produce a histogram that, in standard coordinates, has a very specific appearance. It is symmetric, unimodal; it looks like a narrow bump. If there were enough data points and the histogram boxes were small enough, the curve would look like the curve in figure 2.7. This phenomenon is so important that data of this form has a special name.

**Definition: 2.9** *Standard normal data*

Data is **standard normal data** if, when we have a great deal of data, the histogram is a close approximation to the **standard normal curve**.

This curve is given by

$$y(x) = \frac{1}{\sqrt{2\pi}} e^{(-x^2/2)}$$

(which is shown in figure 2.7).

Indexing: Standard normal data

Indexing: Standard normal data

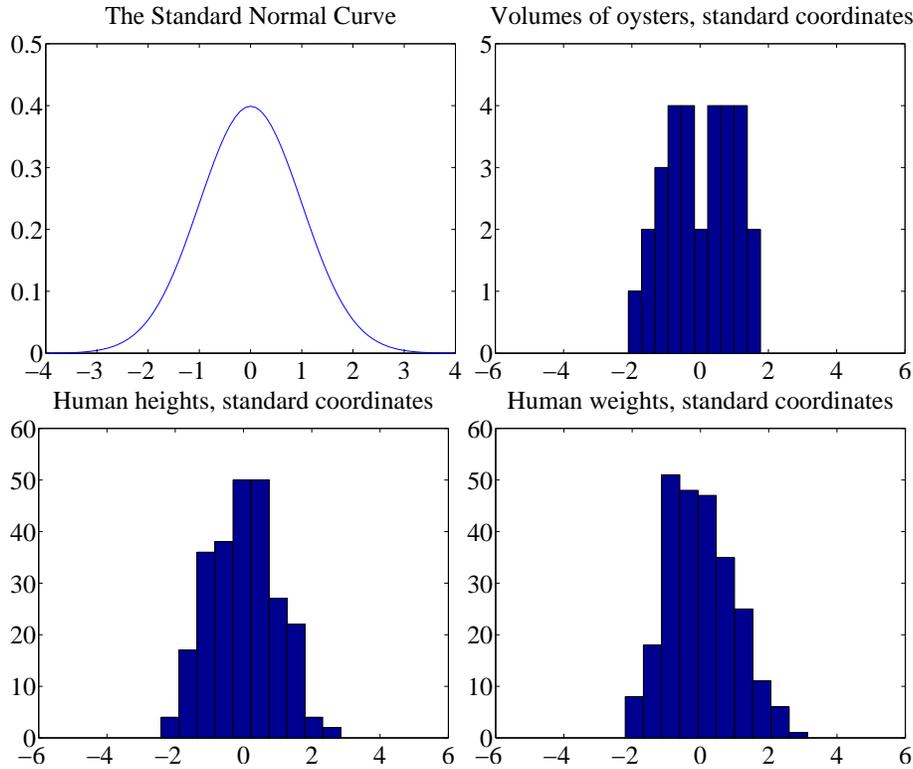


FIGURE 2.7: Data is standard normal data when its histogram takes a stylized, bell-shaped form, plotted above. One usually requires a lot of data and very small histogram boxes for this form to be reproduced closely. Nonetheless, the histogram for normal data is unimodal (has a single bump) and is symmetric; the tails fall off fairly fast, and there are few data items that are many standard deviations from the mean. Many quite different data sets have histograms that are similar to the normal curve; I show three such datasets here.

**Definition: 2.10** *Normal data*

Data is **normal data** if, when we subtract the mean and divide by the standard deviation (i.e. compute standard coordinates), it becomes standard normal data.

Indexing: Normal data

Indexing: Normal data

It is not always easy to tell whether data is normal or not, and there are a variety of tests one can use, which we discuss later. However, there are many examples of normal data. Figure 2.7 shows a diverse variety of data sets, plotted

as histograms in standard coordinates. These include: the volumes of 30 oysters (from [http://www.amstat.org/publications/jse/jse\\_data\\_archive.htm](http://www.amstat.org/publications/jse/jse_data_archive.htm); look for 30oysters.dat.txt); human heights (from <http://www2.stetson.edu/~jrasp/data.htm>; look for bodyfat.xls, with two outliers removed); and human weights (from <http://www2.stetson.edu/~jrasp/data.htm>; look for bodyfat.xls, with two outliers removed).

**Properties of normal data** For the moment, assume we know that a dataset is normal. Then we expect it to have the following properties:

- If we normalize it, its histogram will be close to the standard normal curve. This means, among other things, that the data is not significantly skewed.
- About 68% of the data lie within one standard deviation of the mean. We will prove this later.
- About 95% of the data lie within two standard deviations of the mean. We will prove this later.
- About 99% of the data lie within three standard deviations of the mean. We will prove this later.

In turn, these properties imply that data that contains outliers (points many standard deviations away from the mean) is not normal. This is usually a very safe assumption. It is quite common to model a dataset by excluding a small number of outliers, then modelling the remaining data as normal. For example, if I exclude two outliers from the height and weight data from <http://www2.stetson.edu/~jrasp/data.htm>, the data looks pretty close to normal.

### 2.4.3 Boxplots

It is usually hard to compare multiple histograms by eye. One problem with comparing histograms is the amount of space they take up on a plot, because each histogram involves multiple vertical bars. This means it is hard to plot multiple overlapping histograms cleanly. If you plot each one on a separate figure, you have to handle a large number of separate figures; either you print them too small to see enough detail, or you have to keep flipping over pages.

A **boxplot** is a way to plot data that simplifies comparison. A boxplot displays a dataset as a vertical picture. There is a vertical box whose height corresponds to the interquartile range of the data (the width is just to make the figure easy to interpret). Then there is a horizontal line for the median; and the behavior of the rest of the data is indicated with whiskers and/or outlier markers. This means that each dataset makes is represented by a vertical structure, making it easy to show multiple datasets on one plot *and interpret the plot* (Figure 2.8).

To build a boxplot, we first plot a box that runs from the first to the third quartile. We then show the median with a horizontal line. We then decide which data items should be outliers. A variety of rules are possible; for the plots I show, I used the rule that data items that are larger than  $q_3 + 1.5(q_3 - q_1)$  or smaller than  $q_1 - 1.5(q_3 - q_1)$ , are outliers. This criterion looks for data items that are more than one and a half interquartile ranges above the third quartile, or more than one and a half interquartile ranges below the first quartile.

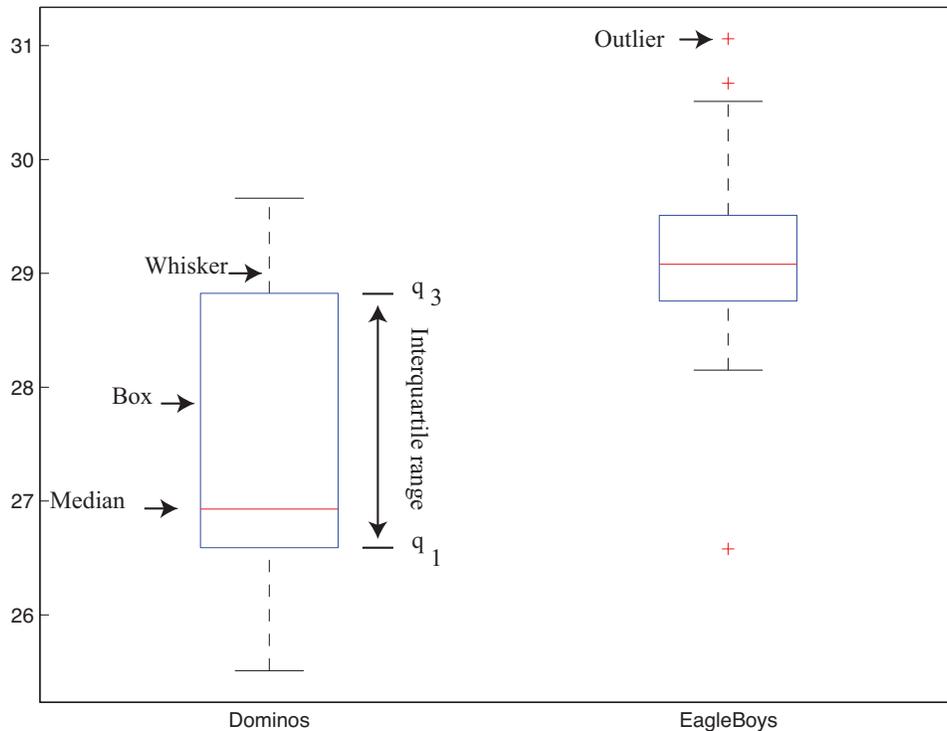


FIGURE 2.8: A boxplot showing the box, the median, the whiskers and two outliers. Notice that we can compare the two datasets rather easily; the next section explains the comparison.

Once we have identified outliers, we plot these with a special symbol (crosses in the plots I show). We then plot whiskers, which show the range of non-outlier data. We draw a whisker from  $q_1$  to the smallest data item that is not an outlier, and from  $q_3$  to the largest data item that is not an outlier. While all this sounds complicated, any reasonable programming environment will have a function that will do it for you. Figure 2.8 shows an example boxplot. Notice that the rich graphical structure means it is quite straightforward to compare two histograms.

## 2.5 WHOSE IS BIGGER? INVESTIGATING AUSTRALIAN PIZZAS

At [http://www.amstat.org/publications/jse/jse\\_data\\_archive.htm](http://www.amstat.org/publications/jse/jse_data_archive.htm), you will find a dataset giving the diameter of pizzas, measured in Australia (search for the word “pizza”). This website also gives the backstory for this dataset. Apparently, EagleBoys pizza claims that their pizzas are always bigger than Dominos pizzas, and published a set of measurements to support this claim (the measurements were available at <http://www.eagleboys.com.au/realsizepizza> as of Feb 2012, but seem not to be there anymore).

Whose pizzas are bigger? and why? A histogram of all the pizza sizes appears

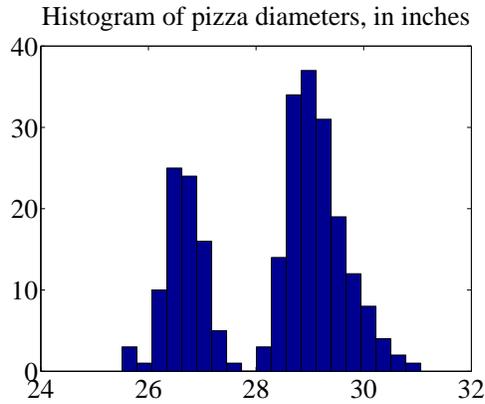


FIGURE 2.9: A histogram of pizza diameters from the dataset described in the text. Notice that there seem to be two populations.

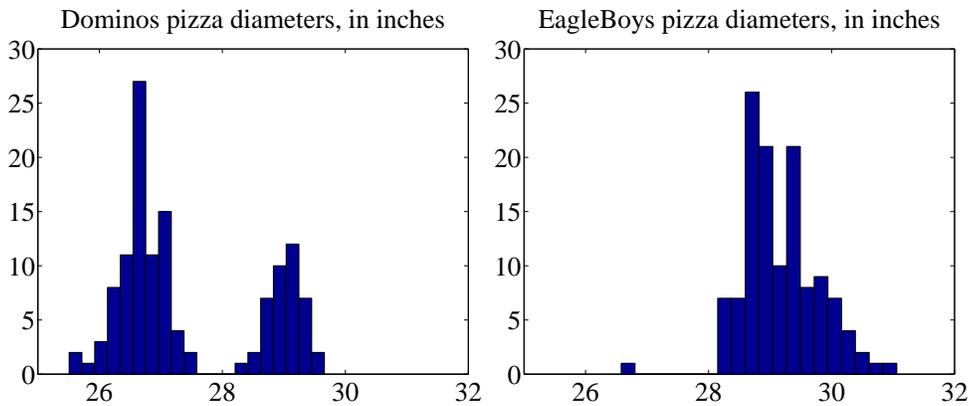


FIGURE 2.10: On the **left**, the class-conditional histogram of Dominos pizza diameters from the pizza data set; on the **right**, the class-conditional histogram of EagleBoys pizza diameters. Notice that EagleBoys pizzas seem to follow the pattern we expect — the diameters are clustered tightly around a mean, and there is a small standard deviation — but Dominos pizzas do not seem to be like that. There is more to understand about this data.

in figure 2.9. We would not expect every pizza produced by a restaurant to have exactly the same diameter, but the diameters are probably pretty close to one another, and pretty close to some standard value. This would suggest that we'd expect to see a histogram which looks like a single, rather narrow, bump about a mean. This is not what we see in figure 2.9 — instead, there are two bumps, which suggests two populations of pizzas. This isn't particularly surprising, because we know that some pizzas come from EagleBoys and some from Dominos.

If you look more closely at the data in the dataset, you will notice that each data item is tagged with the company it comes from. We can now easily plot

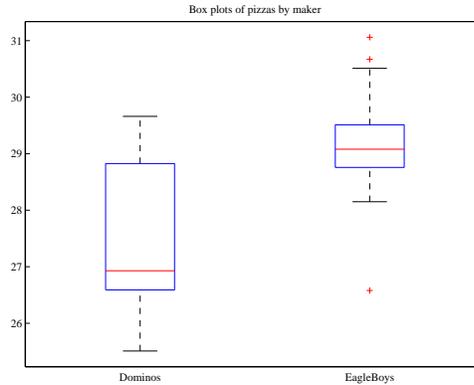


FIGURE 2.11: *Boxplots of the pizza data, comparing EagleBoys and Dominos pizza. There are several curiosities here: why is the range for Dominos so large (25.5-29)? EagleBoys has a smaller range, but has several substantial outliers; why? One would expect pizza manufacturers to try and control diameter fairly closely, because pizzas that are too small present risks (annoying customers; publicity; hostile advertising) and pizzas that are too large should affect profits.*

conditional histograms, conditioning on the company that the pizza came from. These appear in figure 2.10. Notice that EagleBoys pizzas seem to follow the pattern we expect — the diameters are clustered tightly around one value — but Dominos pizzas do not seem to be like that. This is reflected in a boxplot (figure 2.11), which shows the range of Dominos pizza sizes is surprisingly large, and that EagleBoys pizza sizes have several large outliers. There is more to understand about this data. The dataset contains labels for the type of crust and the type of topping — perhaps these properties affect the size of the pizza?

EagleBoys produces DeepPan, MidCrust and ThinCrust pizzas, and Dominos produces DeepPan, ClassicCrust and ThinNCrispy pizzas. This may have something to do with the observed patterns, but comparing six histograms by eye is unattractive. A boxplot is the right way to compare these cases (figure 2.12). The boxplot gives some more insight into the data. Dominos thin crust appear to have a narrow range of diameters (with several outliers), where the median pizza is rather larger than either the deep pan or the classic crust pizza. EagleBoys pizzas all have a range of diameters that is (a) rather similar across the types and (b) rather a lot like the Dominos thin crust. There are outliers, but few for each type.

Another possibility is that the variation in size is explained by the topping. We can compare types and toppings by producing a set of conditional boxplots (i.e. the diameters for each type and each topping). This leads to rather a lot of boxes (figure 2.13), but they’re still easy to compare by eye. The main difficulty is that the labels on the plot have to be shortened. I made labels using the first letter from the manufacturer (“D” or “E”); the first letter from the crust type (previous paragraph); and the first and last letter of the topping. Toppings for Dominos are: Hawaiian; Supreme; BBQMeatlovers. For EagleBoys, toppings are: Hawaiian; SuperSupremo; and BBQMeatlovers. This gives the labels: ‘DCBs’; (Dominos; Clas-

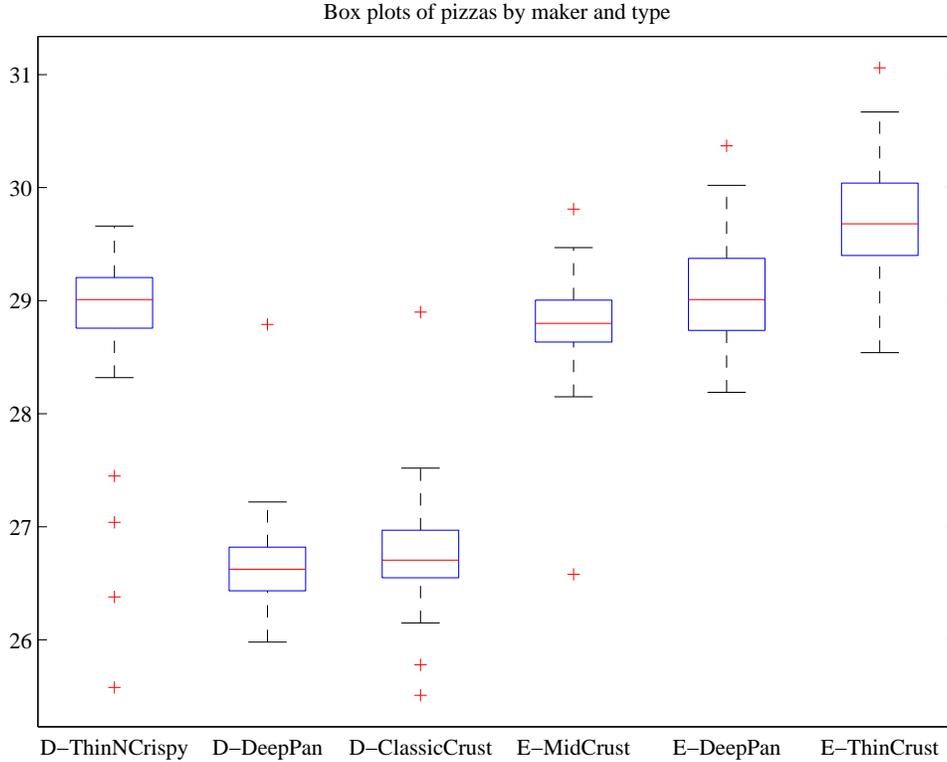


FIGURE 2.12: *Boxplots for the pizza data, broken out by type (thin crust, etc.).*

sicCrust; BBQMeatlovers); 'DCHn'; 'DCSe'; 'DDBs'; 'DDHn'; 'DDSe'; 'DTBs'; 'DTHn'; 'DTSe'; 'EDBs'; 'EDHn'; 'EDSo'; 'EMBs'; 'EMHn'; 'EMSo'; 'ETBs'; 'ETHn'; 'ETSo'. Figure 2.13 suggests that the topping isn't what is important, but the crust (group the boxplots by eye).

What could be going on here? One possible explanation is that Eagleboys have tighter control over the size of the final pizza. One way this could happen is that all EagleBoys pizzas start the same size and shrink the same amount in baking, whereas all Dominos pizzas start a standard diameter, but different Dominos crusts shrink differently in baking. Another way is that Dominos makes different size crusts for different types, but that the cooks sometimes get confused. Yet another possibility is that Dominos controls portions by the mass of dough (so thin crust diameters tend to be larger), but Eagleboys controls by the diameter of the crust.

You should notice that this is more than just a fun story. If you were a manager at a pizza firm, you'd need to make choices about how to control costs. Labor costs, rent, and portion control (i.e. how much pizza, topping, etc. a customer gets for their money) are the main thing to worry about. If the same kind of pizza has a wide range of diameters, you have a problem, because some customers are getting too much (which affects your profit) or too little (which means they might call someone

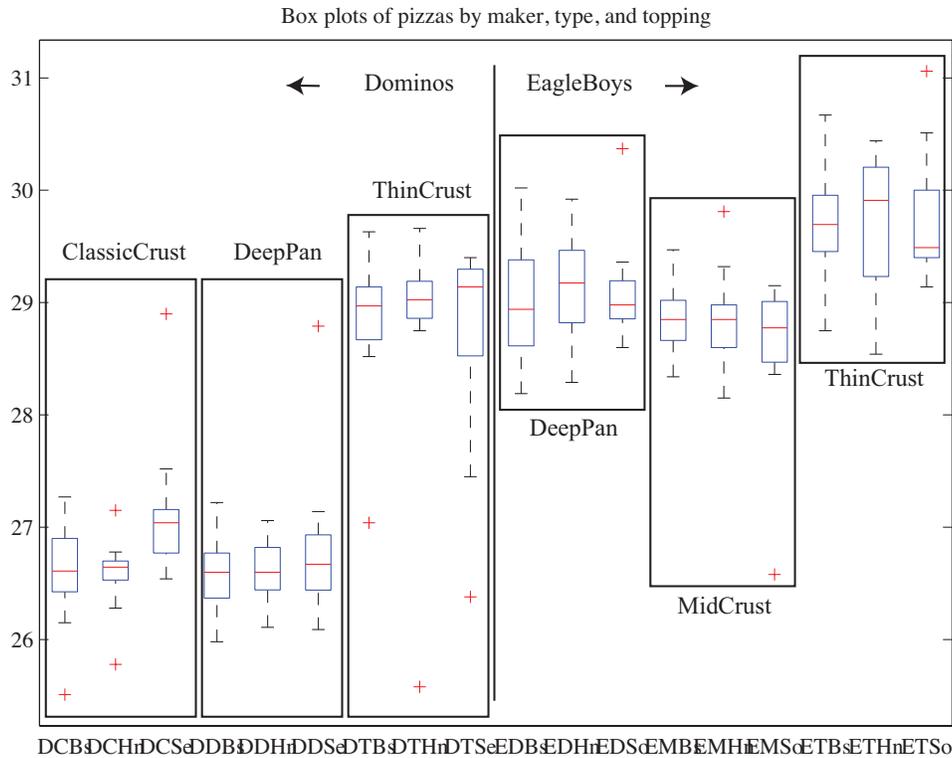


FIGURE 2.13: The pizzas are now broken up by topping as well as crust type (look at the source for the meaning of the names). I have separated Dominos from EagleBoys with a vertical line, and grouped each crust type with a box. It looks as though the issue is not the type of topping, but the crust. EagleBoys seems to have tighter control over the size of the final pizza.

else next time). But making more regular pizzas might require more skilled (and so more expensive) labor. The fact that Dominos and EagleBoys seem to be following different strategies successfully suggests that more than one strategy might work. But you can't choose if you don't know what's happening. As I said at the start, "what's going on here?" is perhaps the single most useful question anyone can ask.

## 2.6 YOU SHOULD

## 2.6.1 be able to:

- Plot a bar chart for a dataset.
- Plot a histogram for a dataset.
- Tell whether the histogram is skewed or not, and in which direction.
- Plot a box plot for one or several datasets.
- Interpret a box plot.

## 2.6.2 remember:

New term: categorical . . . . .	10
New term: ordinal . . . . .	10
New term: continuous . . . . .	10
New term: bar chart . . . . .	13
New term: histogram . . . . .	13
New term: conditional histograms . . . . .	16
New term: class-conditional histograms . . . . .	16
New term: mean . . . . .	16
Definition: Mean . . . . .	17
Useful facts: Mean . . . . .	19
New term: standard deviation . . . . .	19
Definition: Standard deviation . . . . .	19
Useful facts: Standard deviation . . . . .	22
New term: variance . . . . .	23
Definition: Variance . . . . .	23
Useful facts: Variance . . . . .	24
New term: outlier . . . . .	24
New term: median . . . . .	24
Definition: Median . . . . .	25
Useful facts: Median . . . . .	25
New term: interquartile range . . . . .	26
Definition: Percentile . . . . .	26
Definition: Quartiles . . . . .	27
Definition: Interquartile Range . . . . .	27
Useful facts: Interquartile range . . . . .	27
New term: tails . . . . .	28
New term: mode . . . . .	28
New term: unimodal . . . . .	28
New term: multimodal . . . . .	28
New term: bimodal . . . . .	28
New term: skew . . . . .	29
New term: standard coordinates . . . . .	31
Definition: Standard coordinates . . . . .	32
New term: standard normal data . . . . .	32

New term: standard normal curve . . . . .	32
Definition: Standard normal data . . . . .	32
New term: normal data . . . . .	33
Definition: Normal data . . . . .	33
New term: boxplot . . . . .	34

## PROBLEMS

- 2.1. Show that  $\text{mean}(\{kx\}) = k\text{mean}(\{x\})$  by substituting into the definition.
- 2.2. Show that  $\text{mean}(\{x + c\}) = \text{mean}(\{x\}) + c$  by substituting into the definition.
- 2.3. Show that  $\sum_{i=1}^N (x_i - \text{mean}(\{x\})) = 0$  by substituting into the definition.
- 2.4. Show that  $\text{std}(\{x + c\}) = \text{std}(\{x\})$  by substituting into the definition (you'll need to recall the properties of the mean to do this).
- 2.5. Show that  $\text{std}(\{kx\}) = k\text{std}(\{x\})$  by substituting into the definition (you'll need to recall the properties of the mean to do this).
- 2.6. Show that  $\text{median}(\{x + c\}) = \text{median}(\{x\}) + c$  by substituting into the definition.
- 2.7. Show that  $\text{median}(\{kx\}) = k\text{median}(\{x\})$  by substituting into the definition.
- 2.8. Show that  $\text{iqr}\{x + c\} = \text{iqr}\{x\}$  by substituting into the definition.
- 2.9. Show that  $\text{iqr}\{kx\} = k\text{iqr}\{x\}$  by substituting into the definition.

## PROGRAMMING EXERCISES

- 2.10. You can find a data set showing the number of barrels of oil produced, per year for the years 1880-1984 at <http://lib.stat.cmu.edu/DASL/Datafiles/Oilproduction.html>. Is a mean a useful summary of this dataset? Why?
- 2.11. You can find a dataset giving the cost (in 1976 US dollars), number of megawatts, and year of construction of a set of nuclear power plants at <http://lib.stat.cmu.edu/DASL/Datafiles/NuclearPlants.html>.
  - (a) Are there outliers in this data?
  - (b) What is the mean cost of a power plant? What is the standard deviation?
  - (c) What is the mean cost per megawatt? What is the standard deviation?
  - (d) Plot a histogram of the cost per megawatt. Is it skewed? Why?
- 2.12. You can find a dataset giving the sodium content and calorie content of three types of hot dog at <http://lib.stat.cmu.edu/DASL/Datafiles/Hotdogs.html>. The types are Beef, Poultry, and Meat (a rather disturbingly vague label). Use class-conditional histograms to compare these three types of hot dog with respect to sodium content and calories.
- 2.13. You will find a dataset giving (among other things) the number of 3 or more syllable words in advertising copy appearing in magazines at <http://lib.stat.cmu.edu/DASL/Datafiles/magadsdat.html>. The magazines are grouped by the education level of their readers; the groups are 1, 2, and 3 (the variable is called GRP in the data).
  - (a) Use a boxplot to compare the number of three or more syllable words for the ads in magazines in these three groups. What do you see?
  - (b) Use a boxplot to compare the number of sentences appearing in the ads in magazines in these three groups. What do you see?
- 2.14. You can find a dataset giving various properties of public colleges ranked by Kiplingers at <http://www.kiplinger.com/tool/college/T014-S001-kiplinger-s-best-values-in-public-colleges/index.php>. To obtain the data, just select the area on your web browser and paste it into an editor window. I had to then do a little fooling around to remove commas in long numbers, get rid of \$ signs and the like, but then I had a comma separated list. Similar data for private colleges can be found at <http://www.kiplinger.com/tool/college/T014-S001-kiplinger-s-best-values-in-private-colleges/index.php>. You should look at the R code fragment ?? to get started.
  - (a) If you run that fragment on your files, then look at the figures, they should suggest that there isn't much difference in student debt between top, mid

and bottom third public colleges. But there is for private colleges. What happens if you subdivide the top third into smaller pieces? What do you think is going on here?

- (b) If you run that fragment on your files, then look at the figures, they should suggest that there isn't much difference in faculty per student between top, mid and bottom third public colleges. But there is for private colleges. What happens if you subdivide the top third into smaller pieces? What do you think is going on here?
- (c) Use a set of box plots to come to conclusions about how cost varies with quality for the private colleges.
- (d) Use a set of box plots to come to conclusions about how cost varies with quality for the public colleges.

## CHAPTER 3

# Looking at Relationships

We think of a dataset as a collection of  $d$ -tuples (a  $d$ -tuple is an ordered list of  $d$  elements). For example, the Chase and Dunner dataset had entries for Gender; Grade; Age; Race; Urban/Rural; School; Goals; Grades; Sports; Looks; and Money (so it consisted of 11-tuples). The previous chapter explored methods to visualize and summarize a set of values obtained by extracting a single element from each tuple. For example, I could visualize the heights or the weights of a population (as in Figure 2.7). But I could say nothing about the relationship between the height and weight. In this chapter, we will look at methods to visualize and summarize the relationships between pairs of elements of a dataset.

### 3.1 PLOTTING 2D DATA

We take a dataset, choose two different entries, and extract the corresponding elements from each tuple. The result is a dataset consisting of 2-tuples, and we think of this as a two dimensional dataset. The first step is to plot this dataset in a way that reveals relationships. The topic of how best to plot data fills many books, and we can only scratch the surface here. Categorical data can be particularly tricky, because there are a variety of choices we can make, and the usefulness of each tends to depend on the dataset and to some extent on one's cleverness in graphic design (section 3.1.1).

For some continuous data, we can plot the one entry as a function of the other (so, for example, our tuples might consist of the date and the number of robberies; or the year and the price of lynx pelts; and so on, section 3.1.2).

Mostly, we use a simple device, called a scatter plot. Using and thinking about scatter plots will reveal a great deal about the relationships between our data items (section 3.1.3).

#### 3.1.1 Categorical Data, Counts, and Charts

Categorical data is a bit special. Assume we have a dataset with several categorical descriptions of each data item. One way to plot this data is to think of it as belonging to a richer set of categories. Assume the dataset has categorical descriptions, which are not ordinal. Then we can construct a new set of categories by looking at each of the cases for each of the descriptions. For example, in the Chase and Dunner data of table 2.2, our new categories would be: “boy-sports”; “girl-sports”; “boy-popular”; “girl-popular”; “boy-grades”; and “girl-grades”. A large set of categories like this can result in a poor bar chart, though, because there may be too many bars to group the bars successfully. Figure 3.1 shows such a bar chart. Notice that it is hard to group categories by eye to compare; for example, you can see that slightly more girls think grades are important than boys do, but to do so you need to compare two bars that are separated by two other bars. An alternative is a **pie**

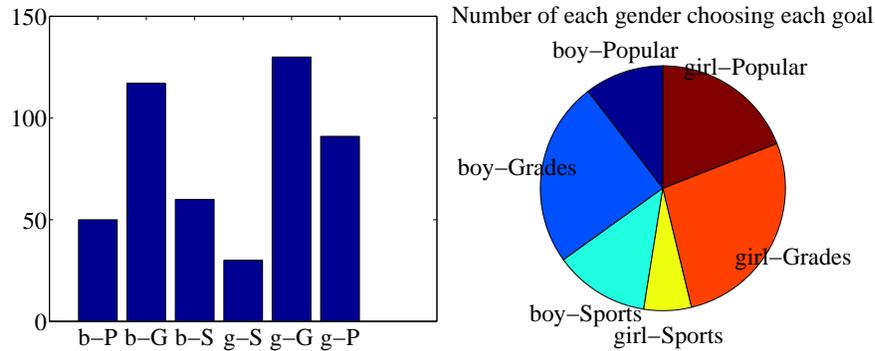


FIGURE 3.1: I sorted the children in the Chase and Dunner study into six categories (two genders by three goals), and counted the number of children that fell into each cell. I then produced the bar chart on the **left**, which shows the number of children of each gender, selecting each goal. On the **right**, a pie chart of this information. I have organized the pie chart so it is easy to compare boys and girls by eye — start at the top; going down on the left side are boy goals, and on the right side are girl goals. Comparing the size of the corresponding wedges allows you to tell what goals (resp. girls) identify with more or less often.

**chart**, where a circle is divided into sections whose angle is proportional to the size of the data item. You can think of the circle as a pie, and each section as a slice of pie. Figure 3.1 shows a pie chart, where each section is proportional to the number of students in its category. In this case, I’ve used my judgement to lay the categories out in a way that makes comparisons easy. I’m not aware of any tight algorithm for doing this, though.

Pie charts have problems, because it is hard to judge small differences in area accurately by eye. For example, from the pie chart in figure 3.1, it’s hard to tell that the “boy-sports” category is slightly bigger than the “boy-popular” category (try it; check using the bar chart). For either kind of chart, it is quite important to think about *what* you plot. For example, the plot of figure 3.1 shows the total number of respondents, and if you refer to figure 2.1, you will notice that there are slightly more girls in the study. Is the *percentage* of boys who think grades are important smaller (or larger) than the *percentage* of girls who think so? you can’t tell from these plots, and you’d have to plot the percentages instead.

An alternative is to use a **stacked bar chart**. You can (say) regard the data as of two types, “Boys” and “Girls”. Within those types, there are subtypes (“Popularity”, “Grades” and “Sport”). The height of the bar is given by the number of elements in the type, and the bar is divided into sections corresponding to the number of elements of that subtype. Alternatively, if you want the plot to show relative frequencies, the bars could all be the same height, but the shading corresponds to the fraction of elements of that subtype. This is all much harder to say than to see or to do (Figure 3.2).

An alternative to a pie chart that is very useful for two dimensional data is a **heat map**. This is a method of displaying a matrix as an image. Each entry of

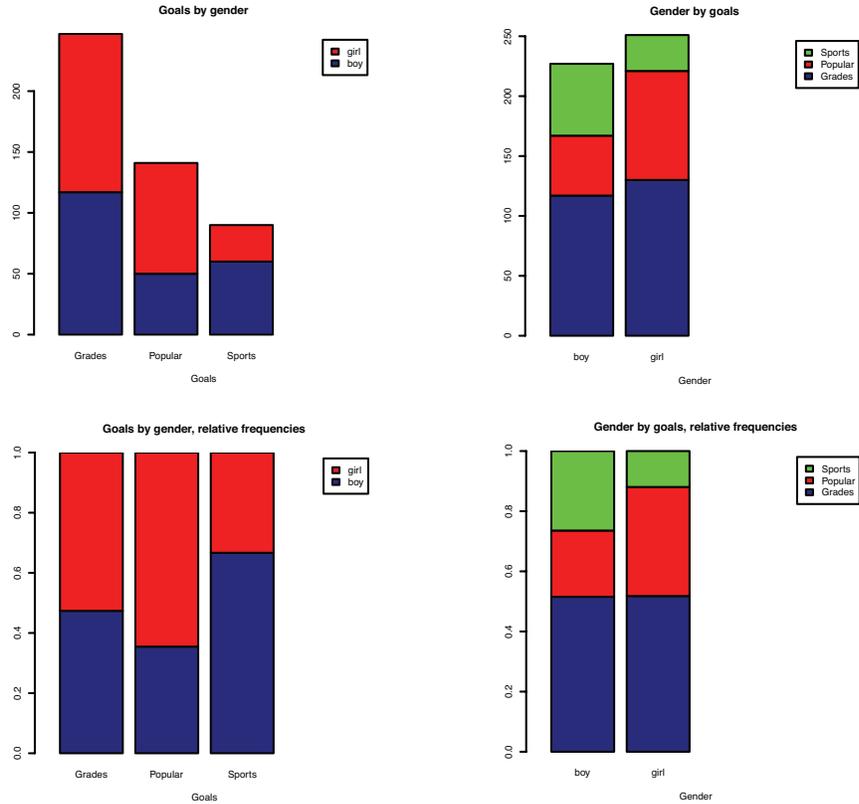


FIGURE 3.2: These bar charts use stacked bars. In the **top row**, the overall height of the bar is given by the number of elements of that type but each different subtype is identified by shading, so you can tell by eye, for example, how many of the “Grades” in the study were “Boys”. This layout makes it hard to tell what fraction of, say, “Boys” aspire to “Popularity”. In the **bottom row**, all bars have the same height, but the shading of the bar identifies the fraction of that type that has a corresponding subtype. This means you can tell by eye what fraction of “Girls” aspire to “Sports”.

the matrix is mapped to a color, and the matrix is represented as an image. For the Chase and Dunner study, I constructed a matrix where each row corresponds to a choice of “sports”, “grades”, or “popular”, and each column corresponds to a choice of “boy” or “girl”. Each entry contains the count of data items of that type. Zero values are represented as white; the largest values as red; and as the value increases, we use an increasingly saturated pink. This plot is shown in figure 3.3

If the categorical data is ordinal, the ordering offers some hints for making a good plot. For example, imagine we are building a user interface. We build an initial version, and collect some users, asking each to rate the interface on scales for “ease of use” (-2, -1, 0, 1, 2, running from bad to good) and “enjoyability” (again, -2, -1, 0, 1, 2, running from bad to good). It is natural to build a 5x5 table, where

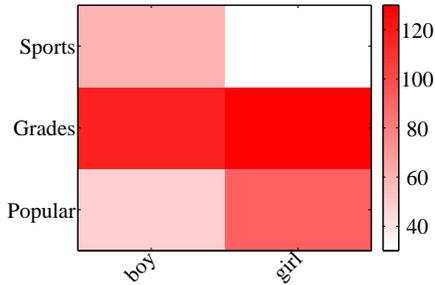


FIGURE 3.3: A heat map of the Chase and Danner data. The color of each cell corresponds to the count of the number of elements of that type. The colorbar at the side gives the correspondence between color and count. You can see at a glance that the number of boys and girls who prefer grades is about the same; that about the same number of boys prefer sports and popularity, with sports showing a mild lead; and that more girls prefer popularity to sports.

	-2	-1	0	1	2
-2	24	5	0	0	1
-1	6	12	3	0	0
0	2	4	13	6	0
1	0	0	3	13	2
2	0	0	0	1	5

TABLE 3.1: I simulated data representing user evaluations of a user interface. Each cell in the table on the **left** contains the count of users rating “ease of use” (horizontal, on a scale of -2 -very bad- to 2 -very good) vs. “enjoyability” (vertical, same scale). Users who found the interface hard to use did not like using it either. While this data is categorical, it’s also ordinal, so that the order of the cells is determined. It wouldn’t make sense, for example, to reorder the columns of the table or the rows of the table.

each cell represents a pair of “ease of use” and “enjoyability” values. We then count the number of users in each cell, and build graphical representations of this table. One natural representation is a **3D bar chart**, where each bar sits on its cell in the 2D table, and the height of the bars is given by the number of elements in the cell. Table 3.1 shows a table and figure 3.4 shows a 3D bar chart for some simulated data. The main difficulty with a 3D bar chart is that some bars are hidden behind others. This is a regular nuisance. You can improve things by using an interactive tool to rotate the chart to get a nice view, but this doesn’t always work. Heatmaps don’t suffer from this problem (Figure 3.4), another reason they are a good choice.

Counts of user responses for a user interface

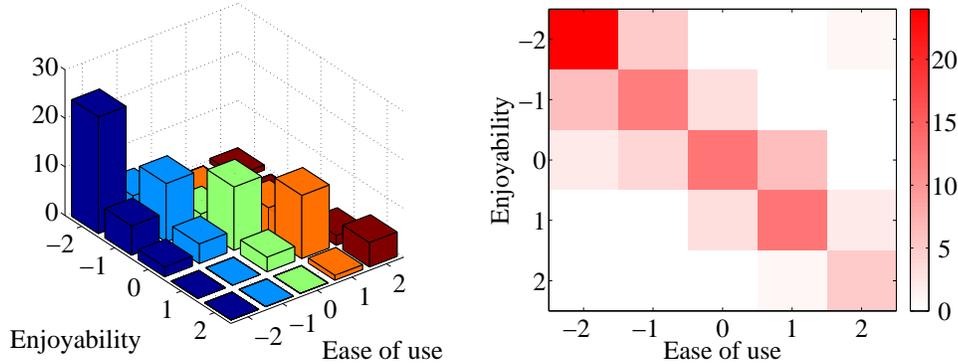


FIGURE 3.4: *On the left*, a 3D bar chart of the data. The height of each bar is given by the number of users in each cell. This figure immediately reveals that users who found the interface hard to use did not like using it either. However, some of the bars at the back are hidden, so some structure might be hard to infer. *On the right*, a heat map of this data. Again, this figure immediately reveals that users who found the interface hard to use did not like using it either. It's more apparent that everyone disliked the interface, though, and it's clear that there is no important hidden structure.

**Remember this:** There are a variety of tools for plotting categorical data. It's difficult to give strict rules for which to use when, but usually one tries to avoid pie charts (angles are hard to judge by eye) and 3D bar charts (where occlusion can hide important effects).

### 3.1.2 Series

Sometimes one component of a dataset gives a natural ordering to the data. For example, we might have a dataset giving the maximum rainfall for each day of the year. We could record this either by using a two-dimensional representation, where one dimension is the number of the day and the other is the temperature, or with a convention where the  $i$ 'th data item is the rainfall on the  $i$ 'th day. For example, at <http://lib.stat.cmu.edu/DASL/Datafiles/timeseriesdat.html>, you can find four datasets indexed in this way. It is natural to plot data like this as a function of time. From this dataset, I extracted data giving the number of burglaries each month in a Chicago suburb, Hyde Park. I have plotted part this data in Figure 3.5 (I left out the data to do with treatment effects). It is natural to plot a graph of the burglaries as a function of time (in this case, the number of the month). The plot shows each data point explicitly. I also told the plotting software to draw lines joining data points, because burglaries do not all happen on a specific day.

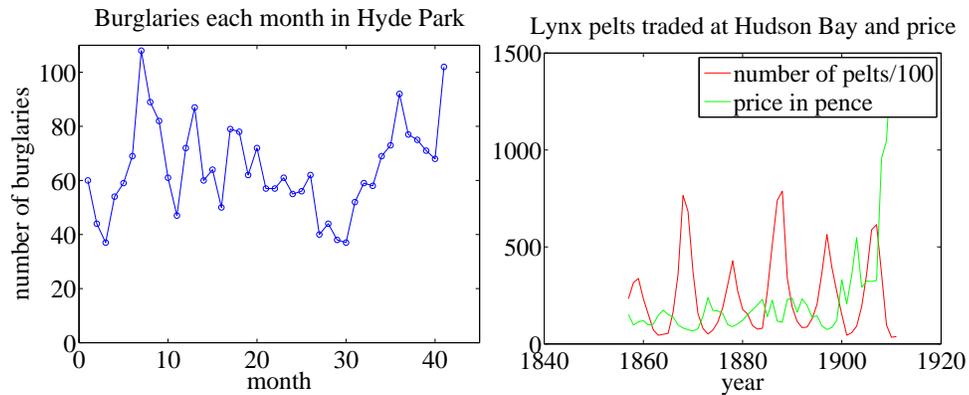


FIGURE 3.5: **Left**, the number of burglaries in Hyde Park, by month. **Right**, a plot of the number of lynx pelts traded at Hudson Bay and of the price paid per pelt, as a function of the year. Notice the scale, and the legend box (the number of pelts is scaled by 100).

The lines suggest, reasonably enough, the rate at which burglaries are happening between data points.

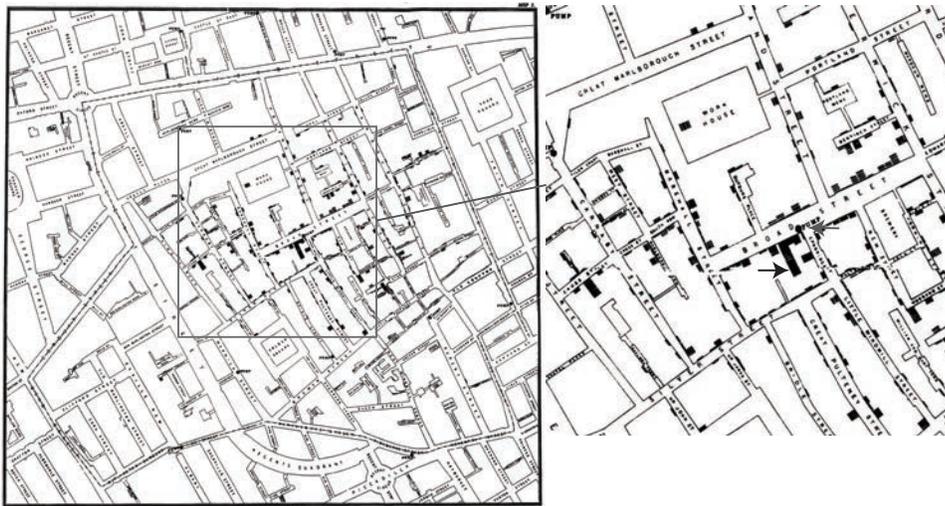


FIGURE 3.6: Snow's scatter plot of cholera deaths on the **left**. Each cholera death is plotted as a small bar on the house in which the bar occurred (for example, the black arrow points to one stack of these bars, indicating many deaths, in the detail on the **right**). Notice the fairly clear pattern of many deaths close to the Broad street pump (grey arrow in the detail), and fewer deaths further away (where it was harder to get water from the pump).

As another example, at <http://lib.stat.cmu.edu/datasets/Andrews/> you can

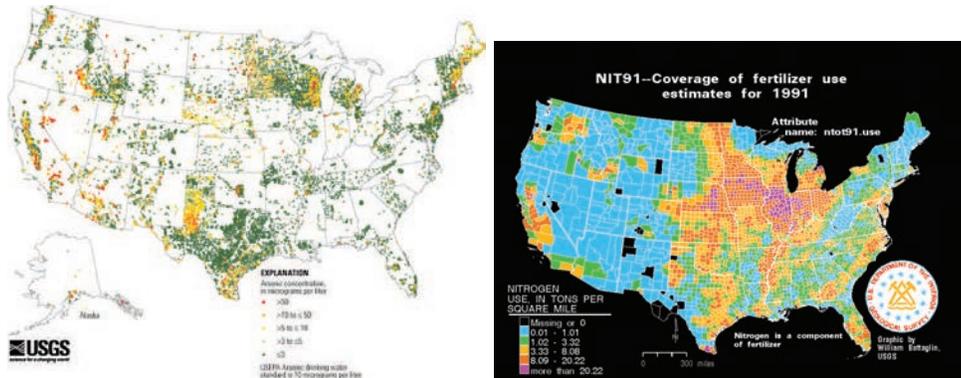


FIGURE 3.7: **Left**, a scatter plot of arsenic levels in US groundwater, prepared by the US Geological Survey (you can find the data at [http://water.usgs.gov/GIS/metadata/usgswrd/XML/arsenic\\_map.xml](http://water.usgs.gov/GIS/metadata/usgswrd/XML/arsenic_map.xml)). Here the shape and color of each marker shows the amount of arsenic, and the spatial distribution of the markers shows where the wells were sampled. **Right**, the usage of Nitrogen (a component of fertilizer) by US county in 1991, prepared by the US Geological Survey (you can find the data at <http://water.usgs.gov/GIS/metadata/usgswrd/XML/nit91.xml>). In this variant of a scatter plot (which usually takes specialized software to prepare) one fills each region with a color indicating the data in that region.

find a dataset that records the number of lynx pelts traded to the Hudson's Bay company and the price paid for each pelt. This version of the dataset appeared first in table 3.2 of *Data: a Collection of Problems from many Fields for the Student and Research Worker* by D.F. Andrews and A.M. Herzberg, published by Springer in 1985. I have plotted it in figure 3.5. The dataset is famous, because it shows a periodic behavior in the number of pelts (which is a good proxy for the number of lynx), which is interpreted as a result of predator-prey interactions. Lynx eat rabbits. When there are many rabbits, lynx kittens thrive, and soon there will be many lynx; but then they eat most of the rabbits, and starve, at which point the rabbit population rockets. You should also notice that after about 1900, prices seem to have gone up rather quickly. I don't know why this is. There is also some suggestion, as there should be, that prices are low when there are many pelts, and high when there are few.

### 3.1.3 Scatter Plots for Spatial Data

It isn't always natural to plot data as a function. For example, in a dataset containing the temperature and blood pressure of a set of patients, there is no reason to believe that temperature is a function of blood pressure, or the other way round. Two people could have the same temperature, and different blood pressures, or vice-versa. As another example, we could be interested in what causes people to die of cholera. We have data indicating *where* each person died in a particular outbreak. It isn't helpful to try and plot such data as a function.

The **scatter plot** is a powerful way to deal with this situation. In the first

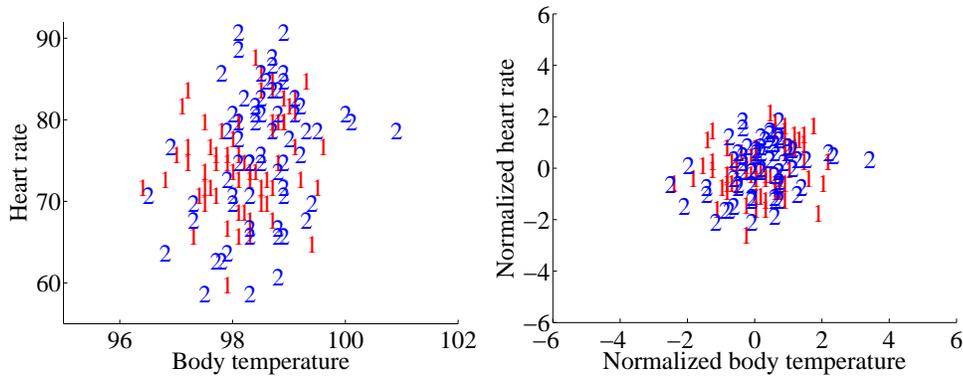


FIGURE 3.8: A scatter plot of body temperature against heart rate, from the dataset at <http://www2.stetson.edu/~jrasp/data.htm>; *normtemp.xls*. I have separated the two genders by plotting a different symbol for each (though I don't know which gender is indicated by which letter); if you view this in color, the differences in color makes for a greater separation of the scatter. This picture suggests, but doesn't conclusively establish, that there isn't much dependence between temperature and heart rate, and any dependence between temperature and heart rate isn't affected by gender.

instance, assume that our data points actually describe points on the a real map. Then, to make a scatter plot, we make a mark on the map at a place indicated by each data point. What the mark looks like, and how we place it, depends on the particular dataset, what we are looking for, how much we are willing to work with complex tools, and our sense of graphic design.

Figure 3.6 is an extremely famous scatter plot, due to John Snow. Snow — one of the founders of epidemiology — used a scatter plot to reason about a cholera outbreak centered on the Broad Street pump in London in 1854. At that time, the mechanism that causes cholera was not known. Snow plotted cholera deaths as little bars (more bars, more deaths) on the location of the house where the death occurred. More bars means more deaths, fewer bars means fewer deaths. There are more bars per block close to the pump, and few far away. This plot offers quite strong evidence of an association between the pump and death from cholera. Snow used this scatter plot as evidence that cholera was associated with water, and that the Broad Street pump was the source of the tainted water.

Figure 3.7 shows a scatter plot of arsenic levels in groundwater for the United States, prepared by the US Geological Survey. The data set was collected by Focazio and others in 2000; by Welch and others in 2000; and then updated by Ryker 2001. It can be found at [http://water.usgs.gov/GIS/metadata/usgswrd/XML/arsenic\\_map.xml](http://water.usgs.gov/GIS/metadata/usgswrd/XML/arsenic_map.xml). One variant of a scatter plot that is particularly useful for geographic data occurs when one fills regions on a map with different colors, following the data in that region. Figure 3.7 shows the nitrogen usage by US county in 1991; again, this figure was prepared by the US Geological Survey.

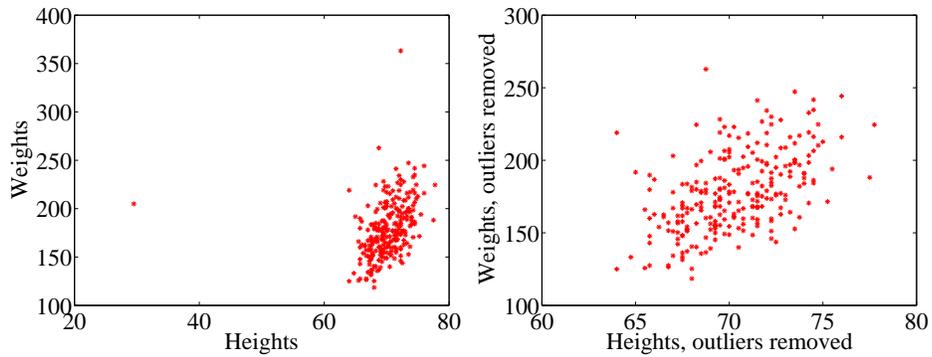


FIGURE 3.9: A scatter plots of weight against height, from the dataset at <http://www2.stetson.edu/~jrasp/data.htm>. **Left:** Notice how two outliers dominate the picture, and to show the outliers, the rest of the data has had to be bunched up. **Right** shows the data with the outliers removed. The structure is now somewhat clearer.

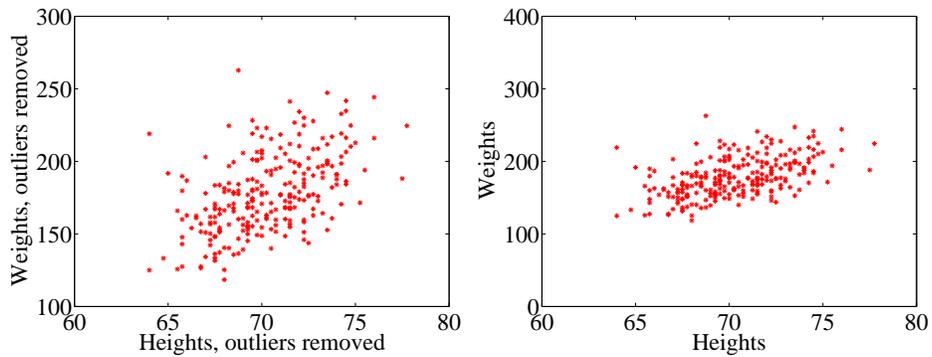


FIGURE 3.10: Scatter plots of weight against height, from the dataset at <http://www2.stetson.edu/~jrasp/data.htm>. **Left:** data with two outliers removed, as in figure 3.9. **Right:** this data, rescaled slightly. Notice how the data looks less spread out. But there is no difference between the datasets. Instead, your eye is easily confused by a change of scale.

**Remember this:** Scatter plots are a most effective tool for geographic data and 2D data in general. A scatter plot should be your first step with a new 2D dataset.

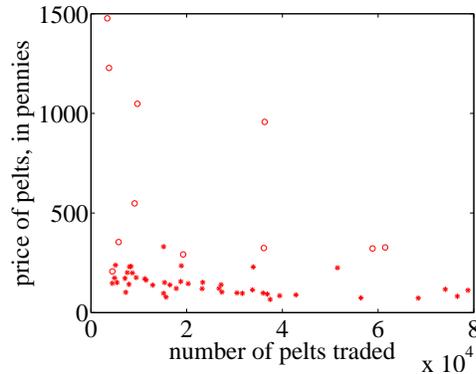


FIGURE 3.11: A scatter plot of the price of lynx pelts against the number of pelts. I have plotted data for 1901 to the end of the series as circles, and the rest of the data as \*'s. It is quite hard to draw any conclusion from this data, because the scale is confusing. Furthermore, the data from 1900 on behaves quite differently from the other data.

#### 3.1.4 Exposing Relationships with Scatter Plots

Scatter plots are natural for geographic data, but a scatter plot is a useful, simple tool for ferreting out associations in other kinds of data as well. Now we need some notation. Assume we have a dataset  $\{x\}$  of  $N$  data items,  $x_1, \dots, x_N$ . Each data item is a  $d$  dimensional vector (so its components are numbers). We wish to investigate the relationship between two components of the dataset. For example, we might be interested in the 7'th and the 13'th component of the dataset. We will produce a two-dimensional plot, one dimension for each component. It does not really matter which component is plotted on the  $x$ -coordinate and which on the  $y$ -coordinate (though it will be some pages before this is clear). But it is very difficult to write sensibly without talking about the  $x$  and  $y$  coordinates.

We will make a two-dimensional dataset out of the components that interest us. We must choose which component goes first in the resulting 2-vector. We will plot this component on the  $x$ -coordinate (and we refer to it as the  $x$ -coordinate), and to the other component as the  $y$ -coordinate. This is just to make it easier to describe what is going on; there's no important idea here. It really will not matter which is  $x$  and which is  $y$ . The two components make a dataset  $\{\mathbf{x}_i\} = \{(x_i, y_i)\}$ . To produce a scatter plot of this data, we plot a small shape at the location of each data item.

Such scatter plots are very revealing. For example, figure 3.8 shows a scatter plot of body temperature against heart rate for humans. In this dataset, the gender of the subject was recorded (as "1" or "2" — I don't know which is which), and so I have plotted a "1" at each data point with gender "1", and so on. Looking at the data suggests there isn't much difference between the blob of "1" labels and the blob of "2" labels, which suggests that females and males are about the same in this respect.

The scale used for a scatter plot matters. For example, plotting lengths in

meters gives a very different scatter from plotting lengths in millimeters. Figure 3.9 shows two scatter plots of weight against height. Each plot is from the same dataset, but one is scaled so as to show two outliers. Keeping these outliers means that the rest of the data looks quite concentrated, just because the axes are in large units. In the other plot, the axis scale has changed (so you can't see the outliers), but the data looks more scattered. This may or may not be a misrepresentation. Figure 3.10 compares the data with outliers removed, with the same plot on a somewhat different set of axes. One plot looks as though increasing height corresponds to increasing weight; the other looks as though it doesn't. This is purely due to deceptive scaling — each plot shows the same dataset.

Dubious data can also contribute to scaling problems. Recall that, in figure 3.5, price data before and after 1900 appeared to behave differently. Figure 3.11 shows a scatter plot of the lynx data, where I have plotted number of pelts against price. I plotted the post-1900 data as circles, and the rest as asterisks. Notice how the circles seem to form a quite different figure, which supports the suggestion that something interesting happened around 1900. We can reasonably choose to analyze data after 1900 separately from before 1900. A choice like this should be made with care. If you exclude every data point that might disagree with your hypothesis, you may miss the fact that you are wrong. Leaving out data is an essential component of many kinds of fraud. You should always reveal whether you have excluded data, and why, to allow the reader to judge the evidence.

When you look at Figure 3.11, you should notice the scatter plot does not seem to support the idea that prices go up when supply goes down. This is puzzling because it's generally a pretty reliable idea. In fact, the plot is just hard to interpret because it is poorly scaled. Scale is an important nuisance, and it's easy to get misled by scale effects.

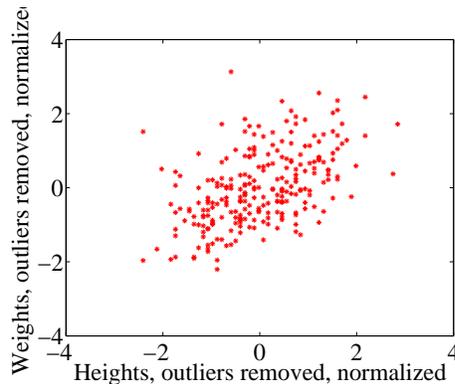


FIGURE 3.12: A normalized scatter plot of weight against height, from the dataset at <http://www2.stetson.edu/~jrasp/data.htm>. Now you can see that someone who is a standard deviation taller than the mean will tend to be somewhat heavier than the mean too.

The way to avoid the problem is to plot in standard coordinates. We can normalize without worrying about the dimension of the data — we normalize each

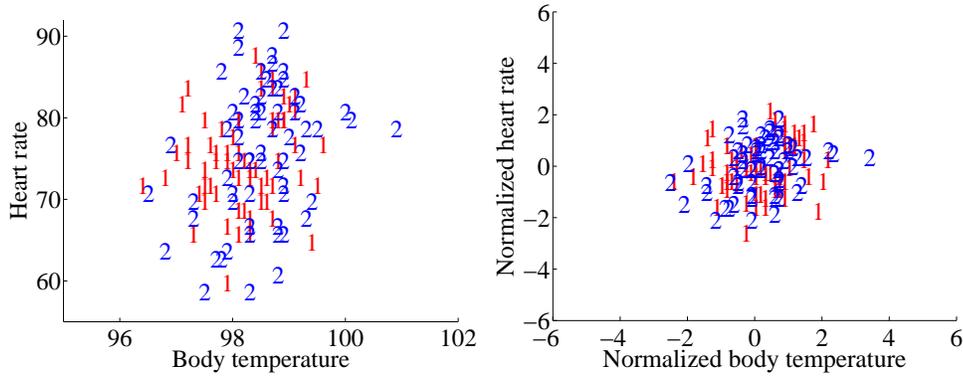


FIGURE 3.13: **Left:** A scatter plot of body temperature against heart rate, from the dataset at <http://www2.stetson.edu/~jrasp/data.htm>; normtemp.xls. I have separated the two genders by plotting a different symbol for each (though I don't know which gender is indicated by which letter); if you view this in color, the differences in color makes for a greater separation of the scatter. This picture suggests, but doesn't conclusively establish, that there isn't much dependence between temperature and heart rate, and any dependence between temperature and heart rate isn't affected by gender. The scatter plot of the normalized data, in standard coordinates, on the right supports this view.

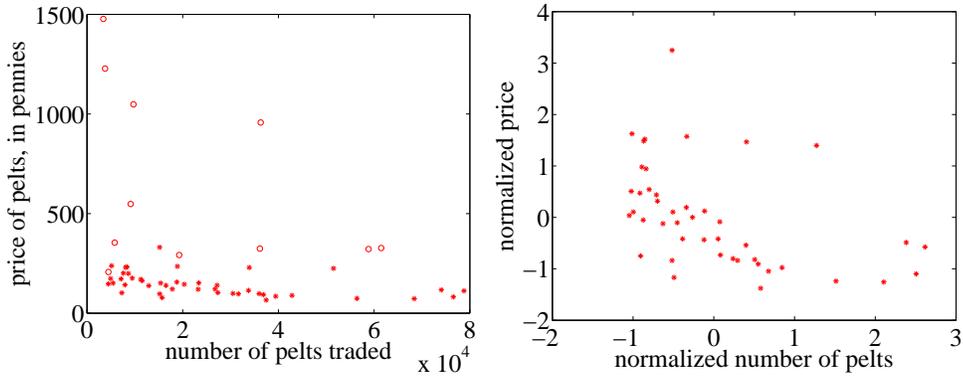


FIGURE 3.14: **Left:** A scatter plot of the price of lynx pelts against the number of pelts (this is a repeat of figure 3.11 for reference). I have plotted data for 1901 to the end of the series as circles, and the rest of the data as \*'s. It is quite hard to draw any conclusion from this data, because the scale is confusing. **Right:** A scatter plot of the price of pelts against the number of pelts for lynx pelts. I excluded data for 1901 to the end of the series, and then normalized both price and number of pelts. Notice that there is now a distinct trend; when there are fewer pelts, they are more expensive, and when there are more, they are cheaper.

dimension independently by subtracting the mean of that dimension and dividing by the standard deviation of that dimension. This means we can normalize the  $x$  and  $y$  coordinates of the two-dimensional data separately. We continue to use the convention of writing the normalized  $x$  coordinate as  $\hat{x}$  and the normalized  $y$  coordinate as  $\hat{y}$ . So, for example, we can write  $\hat{x}_j = (x_j - \text{mean}(\{x\}))/\text{std}(\{x\})$  for the  $\hat{x}$  value of the  $j$ 'th data item in normalized coordinates. Normalizing shows us the dataset on a standard scale. Once we have done this, it is quite straightforward to read off simple relationships between variables from a scatter plot.

**Remember this:** *The plot scale can mask effects in scatter plots, and it's usually a good idea to plot in standard coordinates.*

### 3.2 CORRELATION

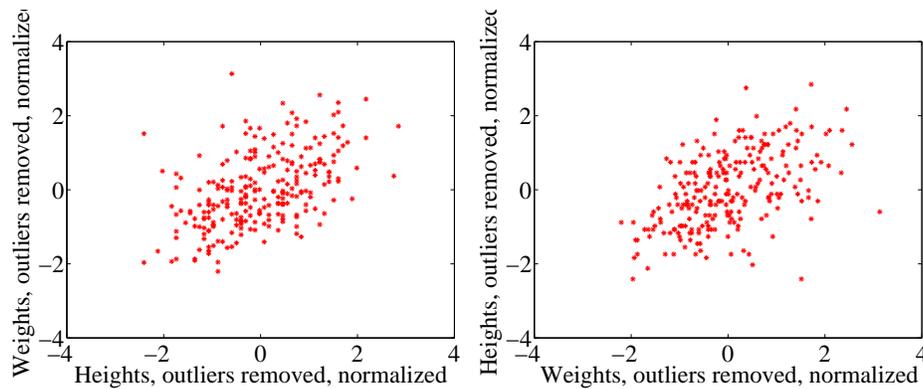


FIGURE 3.15: *On the left, a normalized scatter plot of weight ( $y$ -coordinate) against height ( $x$ -coordinate). On the right, a scatter plot of height ( $y$ -coordinate) against weight ( $x$ -coordinate). I've put these plots next to one another so you don't have to mentally rotate (which is what you should usually do).*

The simplest, and most important, relationship to look for in a scatter plot is this: when  $\hat{x}$  increases, does  $\hat{y}$  tend to increase, decrease, or stay the same? This is straightforward to spot in a normalized scatter plot, because each case produces a very clear shape on the scatter plot. Any relationship is called **correlation** (we will see later how to measure this), and the three cases are: positive correlation, which means that larger  $\hat{x}$  values tend to appear with larger  $\hat{y}$  values; zero correlation, which means no relationship; and negative correlation, which means that larger  $\hat{x}$  values tend to appear with smaller  $\hat{y}$  values. I have shown these cases together in one figure using a real data example (Figure 3.16), so you can compare the appearance of the plots.

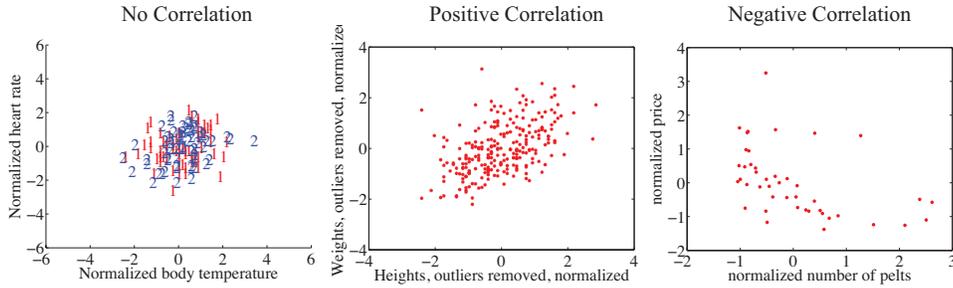


FIGURE 3.16: *The three kinds of scatter plot are less clean for real data than for our idealized examples. Here I used the body temperature vs heart rate data for the zero correlation; the height-weight data for positive correlation; and the lynx data for negative correlation. The pictures aren't idealized — real data tends to be messy — but you can still see the basic structures.*

**Positive correlation** occurs when larger  $\hat{x}$  values tend to appear with larger  $\hat{y}$  values. This means that data points with with small (i.e. negative with large magnitude)  $\hat{x}$  values must have small  $\hat{y}$  values, otherwise the mean of  $\hat{x}$  (resp.  $\hat{y}$ ) would be too big. In turn, this means that the scatter plot should look like a “smear” of data from the bottom left of the graph to the top right. The smear might be broad or narrow, depending on some details we’ll discuss below. Figure 3.12 shows normalized scatter plots of weight against height, and of body temperature against heart rate. In the weight-height plot, you can clearly see that individuals who are higher tend to weigh more. The important word here is “tend” — taller people could be lighter, but mostly they tend not to be. Notice, also, that I did NOT say that they weighed more *because* they were taller, but only that they tend to be heavier.

**Negative correlation** occurs when larger  $\hat{x}$  values tend to appear with smaller  $\hat{y}$  values. This means that data points with with small  $\hat{x}$  values must have large  $\hat{y}$  values, otherwise the mean of  $\hat{x}$  (resp.  $\hat{y}$ ) would be too big. In turn, this means that the scatter plot should look like a “smear” of data from the top left of the graph to the bottom right. The smear might be broad or narrow, depending on some details we’ll discuss below. Figure 3.14 shows a normalized scatter plot of the lynx pelt-price data, where I have excluded the data from 1901 on. I did so because there seemed to be some other effect operating to drive prices up, which was inconsistent with the rest of the series. This plot suggests that when there were more pelts, prices were lower, as one would expect.

**Zero correlation** occurs when there is no relationship. This produces a characteristic shape in a scatter plot, but it takes a moment to understand why. If there really is no relationship, then knowing  $\hat{x}$  will tell you nothing about  $\hat{y}$ . All we know is that  $\text{mean}(\{\hat{y}\}) = 0$ , and  $\text{var}(\{\hat{y}\}) = 1$ . This is enough information to predict what the plot will look like. We know that  $\text{mean}(\{\hat{x}\}) = 0$  and  $\text{var}(\{\hat{x}\}) = 1$ ; so there will be many data points with  $\hat{x}$  value close to zero, and few with a much larger or much smaller  $\hat{x}$  value. The same applies to  $\hat{y}$ . Now consider the data points in a strip of  $\hat{x}$  values. If this strip is far away from the origin, there will

be few data points in the strip, because there aren't many big  $\hat{x}$  values. If there is no relationship, we don't expect to see large or small  $\hat{y}$  values in this strip, because there are few data points in the strip and because large or small  $\hat{y}$  values are uncommon — we see them only if there are many data points, and then seldom. So for a strip with  $\hat{x}$  close to zero, we might see some  $\hat{y}$  values that are far from zero because we will see many  $\hat{y}$  values. For a strip with  $\hat{x}$  that is far from zero, we expect to see few  $\hat{y}$  values that are far from zero, because we see few points in this strip. This reasoning means the data should form a round blob, centered at the origin. In the temperature-heart rate plot of figure 3.13, it looks as though nothing of much significance is happening. The average heart rate seems to be about the same for people who run warm or who run cool. There is probably not much relationship here.

The correlation is not affected by which variable is plotted on the  $x$ -axis and which is plotted on the  $y$ -axis. Figure 3.15 compares a plot of height against weight to one of weight against height. Usually, one just does this by rotating the page, or by imagining the new picture. The left plot tells you that data points with higher height value tend to have higher weight value; the right plot tells you that data points with higher weight value tend to have higher height value — i.e. the plots tell you the same thing. It doesn't really matter which one you look at. Again, the important word is “tend” — the plot doesn't tell you anything about *why*, it just tells you that when one variable is larger the other tends to be, too.

### 3.2.1 The Correlation Coefficient

Consider a normalized data set of  $N$  two-dimensional vectors. We can write the  $i$ 'th data point *in standard coordinates*  $(\hat{x}_i, \hat{y}_i)$ . We already know many important summaries of this data, because it is in standard coordinates. We have  $\text{mean}(\{\hat{x}\}) = 0$ ;  $\text{mean}(\{\hat{y}\}) = 0$ ;  $\text{std}(\{\hat{x}\}) = 1$ ; and  $\text{std}(\{\hat{y}\}) = 1$ . Each of these summaries is itself the mean of some monomial. So  $\text{std}(\{\hat{x}\})^2 = \text{mean}(\{\hat{x}^2\}) = 1$ ;  $\text{std}(\{\hat{y}\})^2 = \text{mean}(\{\hat{y}^2\})$  (the other two are easy). We can rewrite this information in terms of means of monomials, giving  $\text{mean}(\{\hat{x}\}) = 0$ ;  $\text{mean}(\{\hat{y}\}) = 0$ ;  $\text{mean}(\{\hat{x}^2\}) = 1$ ; and  $\text{mean}(\{\hat{y}^2\}) = 1$ . There is one monomial missing here, which is  $\hat{x}\hat{y}$ .

The term  $\text{mean}(\{\hat{x}\hat{y}\})$  captures correlation between  $x$  and  $y$ . The term is known as the **correlation coefficient** or **correlation**.

**Definition: 3.1** *Correlation coefficient*

Assume we have  $N$  data items which are 2-vectors  $(x_1, y_1), \dots, (x_N, y_N)$ , where  $N > 1$ . These could be obtained, for example, by extracting components from larger vectors. We compute the correlation coefficient by first normalizing the  $x$  and  $y$  coordinates to obtain  $\hat{x}_i = \frac{(x_i - \text{mean}(\{x\}))}{\text{std}(x)}$ ,  $\hat{y}_i = \frac{(y_i - \text{mean}(\{y\}))}{\text{std}(y)}$ . The correlation coefficient is the mean value of  $\hat{x}_i \hat{y}_i$ , and can be computed as:

$$\text{corr}(\{(x, y)\}) = \frac{\sum_i \hat{x}_i \hat{y}_i}{N}$$

Indexing: Correlation coefficient

Indexing: Correlation coefficient

Correlation is a measure of our ability to predict one value from another. The correlation coefficient takes values between  $-1$  and  $1$  (we'll prove this below). If the correlation coefficient is close to  $1$ , then we are likely to predict very well. Small correlation coefficients (under about  $0.5$ , say, but this rather depends on what you are trying to achieve) tend not to be all that interesting, because (as we shall see) they result in rather poor predictions. Figure 3.17 gives a set of scatter plots of different real data sets with different correlation coefficients. These all come from data set of age-height-weight, which you can find at <http://www2.stetson.edu/~jrasp/data.htm> (look for bodyfat.xls). In each case, two outliers have been removed. Age and height are hardly correlated, as you can see from the figure. Younger people do tend to be slightly taller, and so the correlation coefficient is  $-0.25$ . You should interpret this as a small correlation. However, the variable called “adiposity” (which isn't defined, but is presumably some measure of the amount of fatty tissue) is quite strongly correlated with weight, with a correlation coefficient is  $0.86$ . Average tissue density is quite strongly negatively correlated with adiposity, because muscle is much denser than fat, so these variables are negatively correlated — we expect high density to appear with low adiposity, and vice versa. The correlation coefficient is  $-0.86$ . Finally, density is very strongly correlated with body weight. The correlation coefficient is  $-0.98$ .

It's not always convenient or a good idea to produce scatter plots in standard coordinates (among other things, doing so hides the units of the data, which can be a nuisance). Fortunately, scaling or translating data does not change the value of the correlation coefficient (though it can change the sign if one scale is negative). This means that it's worth being able to spot correlation in a scatter plot that isn't in standard coordinates (even though correlation is always *defined* in standard coordinates). Figure 3.18 shows different correlated datasets plotted in their original units. These data sets are the same as those used in figure 3.17

**Properties of the Correlation Coefficient**

You should memorize the following properties of the correlation coefficient:

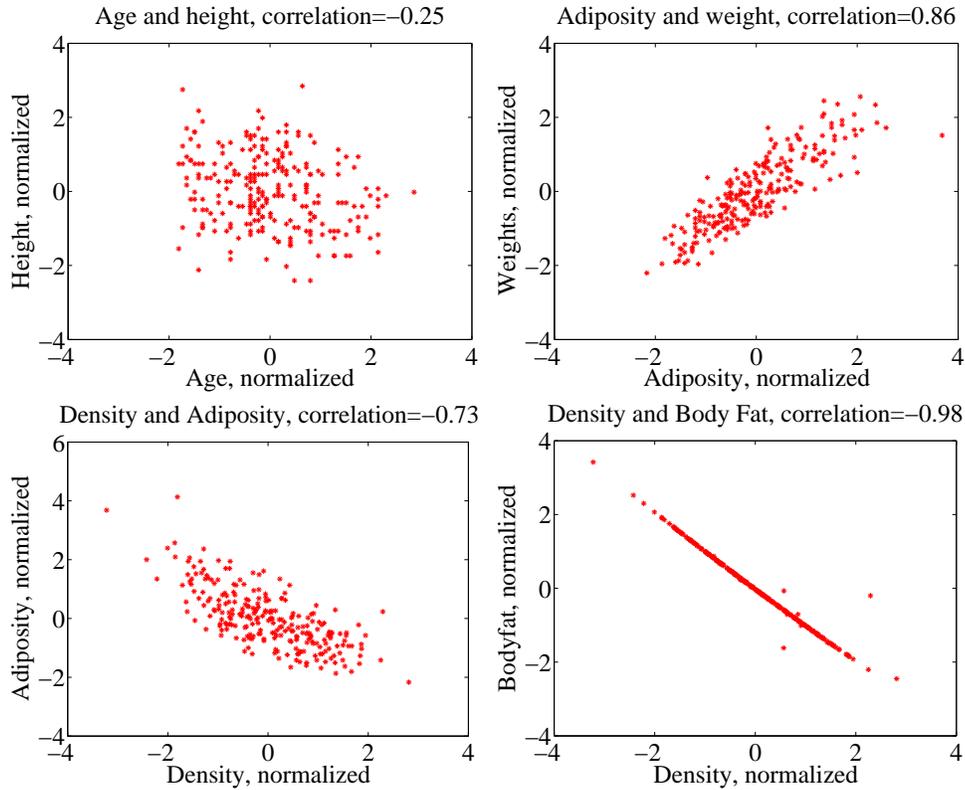


FIGURE 3.17: Scatter plots for various pairs of variables for the age-height-weight dataset from <http://www2.stetson.edu/~jrasp/data.htm>; *bodyfat.xls*. In each case, two outliers have been removed, and the plots are in standard coordinates (compare to figure 3.18, which shows these data sets plotted in their original units). The legend names the variables.

- The correlation coefficient is symmetric (it doesn't depend on the order of its arguments), so

$$\text{corr}(\{(x, y)\}) = \text{corr}(\{(y, x)\})$$

- The value of the correlation coefficient is not changed by translating the data. Scaling the data can change the sign, but not the absolute value. For constants  $a \neq 0$ ,  $b$ ,  $c \neq 0$ ,  $d$  we have

$$\text{corr}(\{(ax + b, cx + d)\}) = \text{sign}(ac)\text{corr}(\{(x, y)\})$$

- If  $\hat{y}$  tends to be large (resp. small) for large (resp. small) values of  $\hat{x}$ , then the correlation coefficient will be positive.
- If  $\hat{y}$  tends to be small (resp. large) for large (resp. small) values of  $\hat{x}$ , then the correlation coefficient will be negative.

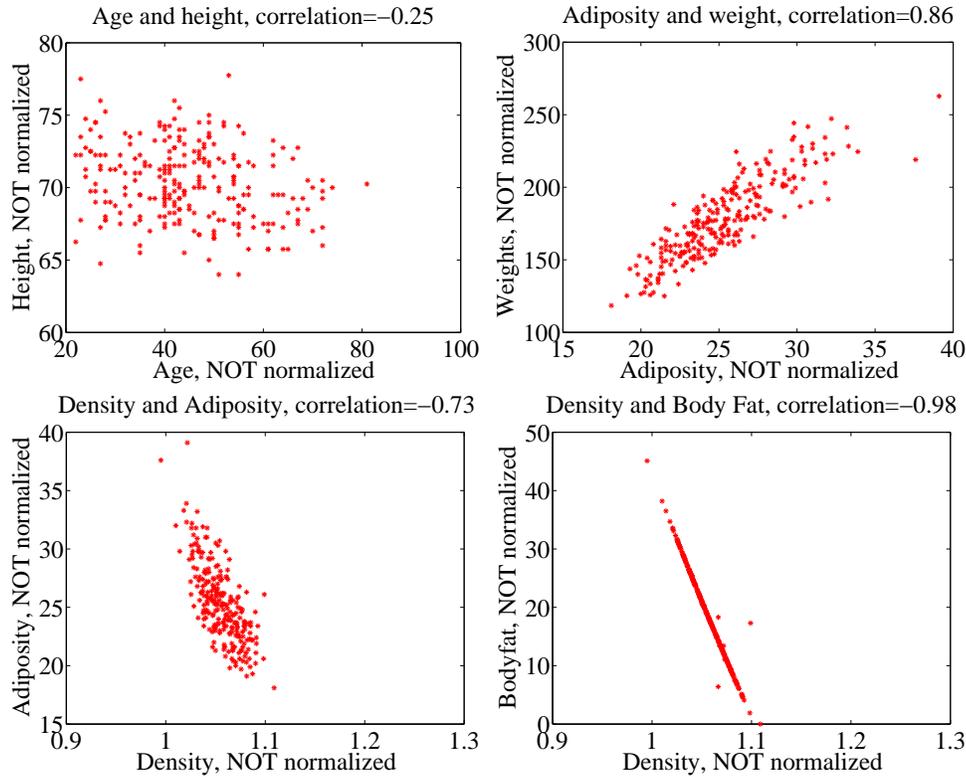


FIGURE 3.18: Scatter plots for various pairs of variables for the age-height-weight dataset from <http://www2.stetson.edu/~jrasp/data.htm>; *bodyfat.xls*. In each case, two outliers have been removed, and the plots are NOT in standard coordinates (compare to figure 3.17, which shows these data sets plotted in normalized coordinates). The legend names the variables.

- If  $\hat{y}$  doesn't depend on  $\hat{x}$ , then the correlation coefficient is zero (or close to zero).
- The largest possible value is 1, which happens when  $\hat{x} = \hat{y}$ .
- The smallest possible value is -1, which happens when  $\hat{x} = -\hat{y}$ .

The first property is easy, and we relegate that to the exercises. One way to see that the correlation coefficient isn't changed by translation or scale is to notice that it is defined in standard coordinates, and scaling or translating data doesn't change those. Another way to see this is to scale and translate data, then write out the equations; notice that taking standard coordinates removes the effects of the scale and translation. In each case, notice that if the scale is negative, the sign of the correlation coefficient changes.

The property that, if  $\hat{y}$  tends to be large (resp. small) for large (resp. small) values of  $\hat{x}$ , then the correlation coefficient will be positive, doesn't really admit

a formal statement. But it's relatively straightforward to see what's going on. Because  $\text{mean}(\{\hat{x}\}) = 0$ , small values of  $\text{mean}(\{\hat{x}\})$  must be negative and large values must be positive. But  $\text{corr}(\{(x, y)\}) = \frac{\sum_i \hat{x}_i \hat{y}_i}{N}$ ; and for this sum to be positive, it should contain mostly positive terms. It can contain few or no hugely positive (or hugely negative) terms, because  $\text{std}(\hat{x}) = \text{std}(\hat{y}) = 1$  so there aren't many large (or small) numbers. For the sum to contain mostly positive terms, then the sign of  $\hat{x}_i$  should be the same as the sign  $\hat{y}_i$  for most data items. Small changes to this argument work to show that if  $\hat{y}$  tends to be small (resp. large) for large (resp. small) values of  $\hat{x}$ , then the correlation coefficient will be negative.

Showing that no relationship means zero correlation requires slightly more work. Divide the scatter plot of the dataset up into thin vertical strips. There are  $S$  strips. Each strip is narrow, so the  $\hat{x}$  value does not change much for the data points in a particular strip. For the  $s$ 'th strip, write  $N(s)$  for the number of data points in the strip,  $\hat{x}(s)$  for the  $\hat{x}$  value at the center of the strip, and  $\bar{\hat{y}}(s)$  for the mean of the  $\hat{y}$  values within that strip. Now the strips are narrow, so we can approximate all data points within a strip as having the same value of  $\hat{x}$ . This yields

$$\text{mean}(\{\hat{x}\hat{y}\}) \approx \frac{1}{S} \sum_{s \in \text{strips}} \hat{x}(s) [N(s)\bar{\hat{y}}(s)]$$

(where you could replace  $\approx$  with  $=$  if the strips were narrow enough). Now assume that  $\bar{\hat{y}}(s)$  does not change from strip to strip, meaning that there is no relationship between  $\hat{x}$  and  $\hat{y}$  in this dataset (so the picture is like the left hand side in figure 3.16). Then each value of  $\bar{\hat{y}}(s)$  is the same — we write  $\bar{\hat{y}}$  — and we can rearrange to get

$$\text{mean}(\{\hat{x}\hat{y}\}) \approx \bar{\hat{y}} \frac{1}{S} \sum_{s \in \text{strips}} \hat{x}(s).$$

Now notice that

$$0 = \text{mean}(\{\hat{y}\}) \approx \frac{1}{S} \sum_{s \in \text{strips}} N(s)\bar{\hat{y}}(s)$$

(where again you could replace  $\approx$  with  $=$  if the strips were narrow enough). This means that if every strip has the same value of  $\bar{\hat{y}}(s)$ , then that value must be zero. In turn, if there is no relationship between  $\hat{x}$  and  $\hat{y}$ , we must have  $\text{mean}(\{\hat{x}\hat{y}\}) = 0$ .

**Proposition:**

$$-1 \leq \text{corr}(\{(x, y)\}) \leq 1$$

**Proof:** Writing  $\hat{x}$ ,  $\hat{y}$  for the normalized coefficients, we have

$$\text{corr}(\{(x, y)\}) = \frac{\sum_i \hat{x}_i \hat{y}_i}{N}$$

and you can think of the value as the inner product of two vectors. We write

$$\mathbf{x} = \frac{1}{\sqrt{N}} [\hat{x}_1, \hat{x}_2, \dots, \hat{x}_N] \quad \text{and} \quad \mathbf{y} = \frac{1}{\sqrt{N}} [\hat{y}_1, \hat{y}_2, \dots, \hat{y}_N]$$

and we have  $\text{corr}(\{(x, y)\}) = \mathbf{x}^T \mathbf{y}$ . Notice  $\mathbf{x}^T \mathbf{x} = \text{std}(x)^2 = 1$ , and similarly for  $\mathbf{y}$ . But the inner product of two vectors is at its maximum when the two vectors are the same, and this maximum is 1. This argument is also sufficient to show that smallest possible value of the correlation is  $-1$ , and this occurs when  $\hat{x}_i = -\hat{y}_i$  for all  $i$ .

**Property 3.1:** The largest possible value of the correlation is 1, and this occurs when  $\hat{x}_i = \hat{y}_i$  for all  $i$ . The smallest possible value of the correlation is  $-1$ , and this occurs when  $\hat{x}_i = -\hat{y}_i$  for all  $i$ .

### 3.2.2 Using Correlation to Predict

Assume we have  $N$  data items which are 2-vectors  $(x_1, y_1), \dots, (x_N, y_N)$ , where  $N > 1$ . These could be obtained, for example, by extracting components from larger vectors. As usual, we will write  $\hat{x}_i$  for  $x_i$  in normalized coordinates, and so on. Now assume that we know the correlation coefficient is  $r$  (this is an important, traditional notation). What does this mean?

One (very useful) interpretation is in terms of prediction. Assume we have a data point  $(x_0, ?)$  where we know the  $x$ -coordinate, but not the  $y$ -coordinate. We can use the correlation coefficient to predict the  $y$ -coordinate. First, we transform to standard coordinates. Now we must obtain the best  $\hat{y}_0$  value to predict, using the  $\hat{x}_0$  value we have.

We want to construct a prediction function which gives a prediction for any value of  $\hat{x}$ . This predictor should behave as well as possible on our existing data. For each of the  $(\hat{x}_i, \hat{y}_i)$  pairs in our data set, the predictor should take  $\hat{x}_i$  and produce a result as close to  $\hat{y}_i$  as possible. We can choose the predictor by looking at the errors it makes at each data point.

We write  $\hat{y}_i^p$  for the value of  $\hat{y}_i$  predicted at  $\hat{x}_i$ . The simplest form of predictor is linear. If we predict using a linear function, then we have, for some unknown  $a, b$ , that  $\hat{y}_i^p = a\hat{x}_i + b$ . Now think about  $u_i = \hat{y}_i - \hat{y}_i^p$ , which is the error in our prediction. We would like to have  $\text{mean}(\{u\}) = 0$  (otherwise, we could reduce the

error of the prediction just by subtracting a constant).

$$\begin{aligned}
 \text{mean}(\{u\}) &= \text{mean}(\{\hat{y} - \hat{y}^p\}) \\
 &= \text{mean}(\{\hat{y}\}) - \text{mean}(\{a\hat{x}_i + b\}) \\
 &= \text{mean}(\{\hat{y}\}) - a\text{mean}(\{\hat{x}\}) + b \\
 &= 0 - a0 + b \\
 &= 0.
 \end{aligned}$$

This means that we must have  $b = 0$ .

To estimate  $a$ , we need to think about  $\text{var}(\{u\})$ . We should like  $\text{var}(\{u\})$  to be as small as possible, so that the errors are as close to zero as possible (remember, small variance means small standard deviation which means the data is close to the mean). We have

$$\begin{aligned}
 \text{var}(\{u\}) &= \text{var}(\{\hat{y} - \hat{y}^p\}) \\
 &= \text{mean}(\{(\hat{y} - a\hat{x})^2\}) \quad \text{because } \text{mean}(\{u\}) = 0 \\
 &= \text{mean}(\{(\hat{y})^2 - 2a\hat{x}\hat{y} + a^2(\hat{x})^2\}) \\
 &= \text{mean}(\{(\hat{y})^2\}) - 2a\text{mean}(\{\hat{x}\hat{y}\}) + a^2\text{mean}(\{(\hat{x})^2\}) \\
 &= 1 - 2ar + a^2,
 \end{aligned}$$

which we want to minimize by choice of  $a$ . At the minimum, we must have

$$\frac{d\text{var}(\{u_i\})}{da} = 0 = -2r + 2a$$

so that  $a = r$  and the correct prediction is

$$\hat{y}_0^p = r\hat{x}_0$$

You can use a version of this argument to establish that if we have  $(\hat{x}_0, \hat{y}_0)$ , then the best prediction for  $\hat{x}_0$  (*which is in standard coordinates*) is  $r\hat{y}_0$ . It is important to notice that the coefficient of  $\hat{y}_i$  is NOT  $1/r$ ; you should work this example, which appears in the exercises. We now have a prediction procedure, outlined below.

**Procedure: 3.1** *Predicting a value using correlation*

Assume we have  $N$  data items which are 2-vectors  $(x_1, y_1), \dots, (x_N, y_N)$ , where  $N > 1$ . These could be obtained, for example, by extracting components from larger vectors. Assume we have an  $x$  value  $x_0$  for which we want to give the best prediction of a  $y$  value, based on this data. The following procedure will produce a prediction:

- Transform the data set into standard coordinates, to get

$$\hat{x}_i = \frac{1}{\text{std}(x)}(x_i - \text{mean}(\{x\}))$$

$$\hat{y}_i = \frac{1}{\text{std}(y)}(y_i - \text{mean}(\{y\}))$$

$$\hat{x}_0 = \frac{1}{\text{std}(x)}(x_0 - \text{mean}(\{x\})).$$

- Compute the correlation

$$r = \text{corr}(\{(x, y)\}) = \text{mean}(\{\hat{x}\hat{y}\}).$$

- Predict  $\hat{y}_0 = r\hat{x}_0$ .
- Transform this prediction into the original coordinate system, to get

$$y_0 = \text{std}(y)r\hat{x}_0 + \text{mean}(\{y\})$$

Now assume we have a  $y$  value  $y_0$ , for which we want to give the best prediction of an  $x$  value, based on this data. The following procedure will produce a prediction:

- Transform the data set into standard coordinates.
- Compute the correlation.
- Predict  $\hat{x}_0 = r\hat{y}_0$ .
- Transform this prediction into the original coordinate system, to get

$$x_0 = \text{std}(x)r\hat{y}_0 + \text{mean}(\{x\})$$

There is another way of thinking about this prediction procedure, which is often helpful. Assume we need to predict a value for  $x_0$ . In normalized coordinates, our prediction is  $\hat{y}^p = r\hat{x}_0$ ; if we revert back to the original coordinate system, the

prediction becomes

$$\frac{(y^p - \text{mean}(\{y\}))}{\text{std}(y)} = r \left( \frac{(x_0 - \text{mean}(\{x\}))}{\text{std}(x)} \right).$$

This gives a really useful rule of thumb, which I have broken out in the box below.

**Procedure: 3.2** *Predicting a value using correlation: Rule of thumb - 1*

If  $x_0$  is  $k$  standard deviations from the mean of  $x$ , then the predicted value of  $y$  will be  $rk$  standard deviations away from the mean of  $y$ , and the sign of  $r$  tells whether  $y$  increases or decreases.

An even more compact version of the rule of thumb is in the following box.

**Procedure: 3.3** *Predicting a value using correlation: Rule of thumb - 2*

The predicted value of  $y$  goes up by  $r$  standard deviations when the value of  $x$  goes up by one standard deviation.

We can compute the average root mean square error that this prediction procedure will make. The square of this error must be

$$\begin{aligned} \text{mean}(\{u^2\}) &= \text{mean}(\{y^2\}) - 2r\text{mean}(\{xy\}) + r^2\text{mean}(\{x^2\}) \\ &= 1 - 2r^2 + r^2 \\ &= 1 - r^2 \end{aligned}$$

so the root mean square error will be  $\sqrt{1 - r^2}$ . This is yet another interpretation of correlation; if  $x$  and  $y$  have correlation close to one, then predictions could have very small root mean square error, and so might be very accurate. In this case, knowing one variable is about as good as knowing the other. If they have correlation close to zero, then the root mean square error in a prediction might be as large as the root mean square error in  $\hat{y}$  — which means the prediction is nearly a pure guess.

The prediction argument means that we can spot correlations for data in other kinds of plots — one doesn't have to make a scatter plot. For example, if we were to observe a child's height from birth to their 10'th year (you can often find these observations in ballpen strokes, on kitchen walls), we could plot height as a function of year. If we also had their weight (less easily found), we could plot weight as a function of year, too. The prediction argument above says that, if you can predict the weight from the height (or vice versa) then they're correlated. One way to spot this is to look and see if one curve goes up when the other does (or goes down when the other goes up). You can see this effect in figure 3.5, where (before 19h00), prices go down when the number of pelts goes up, and vice versa. These two variables are negatively correlated.

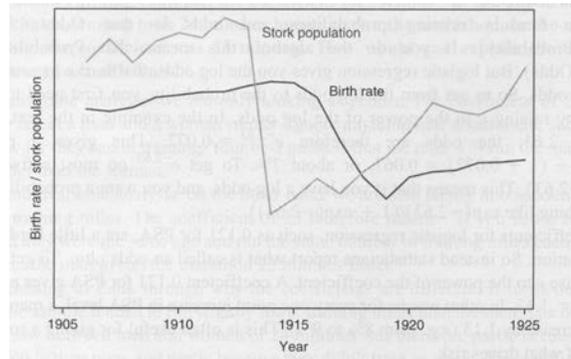


FIGURE 3.19: This figure, from Vickers (*ibid*, p184) shows a plot of the stork population as a function of time, and the human birth rate as a function of time, for some years in Germany. The correlation is fairly clear; but this does not mean that reducing the number of storks means there are fewer able to bring babies. Instead, this is the impact of the first world war — a hidden or latent variable.

### 3.2.3 Confusion caused by correlation

There is one very rich source of potential (often hilarious) mistakes in correlation. When two variables are correlated, they change together. If the correlation is positive, that means that, in typical data, if one is large then the other is large, and if one is small the other is small. In turn, this means that one can make a reasonable prediction of one from the other. However, correlation DOES NOT mean that changing one variable causes the other to change (sometimes known as causation).

Two variables in a dataset could be correlated for a variety of reasons. One important reason is pure accident. If you look at enough pairs of variables, you may well find a pair that appears to be correlated just because you have a small set of observations. Imagine, for example, you have a dataset consisting of only two vectors — there is a pretty good chance that there is some correlation between the coefficients. Such accidents can occur in large datasets, particularly if the dimensions are high.

Another reason variables could be correlated is that there is some causal relationship — for example, pressing the accelerator tends to make the car go faster, and so there will be some correlation between accelerator position and car acceleration. As another example, adding fertilizer does tend to make a plant grow bigger. Imagine you record the amount of fertilizer you add to each pot, and the size of the resulting potplant. There should be some correlation.

Yet another reason variables could be correlated is that there is some other background variable — often called a **latent variable** — linked causally to each of the observed variables. For example, in children (as Freedman, Pisani and Purves note in their excellent *Statistics*), shoe size is correlated with reading skills. This DOES NOT mean that making your feet grow will make you read faster, or that you can make your feet shrink by forgetting how to read. The real issue here is

the age of the child. Young children tend to have small feet, and tend to have weaker reading skills (because they've had less practice). Older children tend to have larger feet, and tend to have stronger reading skills (because they've had more practice). You can make a reasonable prediction of reading skills from foot size, because they're correlated, even though there is no direct connection.

This kind of effect can mask correlations, too. Imagine you want to study the effect of fertilizer on potplants. You collect a set of pots, put one plant in each, and add different amounts of fertilizer. After some time, you record the size of each plant. You expect to see correlation between fertilizer amount and plant size. But you might not if you had used a different species of plant in each pot. Different species of plant can react quite differently to the same fertilizer (some plants just die if over-fertilized), so the species could act as a latent variable. With an unlucky choice of the different species, you might even conclude that there was a negative correlation between fertilizer and plant size. This example illustrates why you need to take great care in setting up experiments and interpreting their results.

This sort of thing happens often, and it's an effect you should look for. Another nice example comes from Vickers (*ibid*). The graph, shown in Figure 3.19, shows a plot of (a) a dataset of the stork population in Europe over a period of years and (b) a dataset of the birth rate over those years. This isn't a scatter plot; instead, the data has been plotted on a graph. You can see by eye that these two datasets are quite strongly correlated. Even more disturbing, the stork population dropped somewhat before the birth rate dropped. Is this evidence that storks brought babies in Europe during those years? No (the usual arrangement seems to have applied). For a more sensible explanation, look at the dates. The war disturbed both stork and human breeding arrangements. Storks were disturbed immediately by bombs, etc., and the human birth rate dropped because men died at the front.

### 3.3 STERILE MALES IN WILD HORSE HERDS

Large herds of wild horses are (apparently) a nuisance, but keeping down numbers by simply shooting surplus animals would provoke outrage. One strategy that has been adopted is to sterilize males in the herd; if a herd contains sufficient sterile males, fewer foals should result. But catching stallions, sterilizing them, and reinserting them into a herd is a performance — does this strategy work?

We can get some insight by plotting data. At <http://lib.stat.cmu.edu/DASL/Datafiles/WildHorses.html>, you can find a dataset covering herd management in wild horses. I have plotted part of this dataset in figure 3.20. In this dataset, there are counts of all horses, sterile males, and foals made on each of a small number of days in 1986, 1987, and 1988 for each of two herds. I extracted data for one herd. I have plotted this data as a function of the count of days since the first data point, because this makes it clear that some measurements were taken at about the same time, but there are big gaps in the measurements. In this plot, the data points are shown with a marker. Joining them leads to a confusing plot because the data points vary quite strongly. However, notice that the size of the herd drifts down slowly (you could hold a ruler against the plot to see the trend), as does the number of foals, when there is a (roughly) constant number of sterile males.

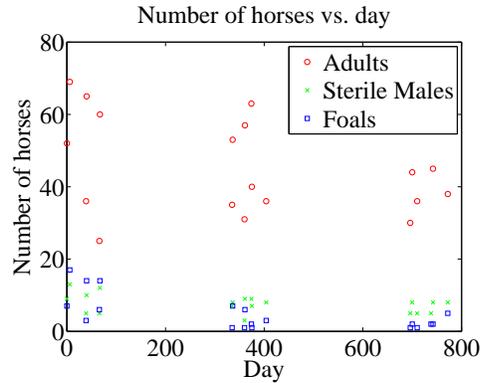


FIGURE 3.20: A plot of the number of adult horses, sterile males, and foals in horse herds over a period of three years. The plot suggests that introducing sterile males might cause the number of foals to go down. Data from <http://lib.stat.cmu.edu/DASL/Datafiles/WildHorses.html>.

Does sterilizing males result in fewer foals? This is likely hard to answer for this dataset, but we could ask whether herds with more sterile males have fewer foals. A scatter plot is a natural tool to attack this question. However, the scatter plots of figure 3.21 suggest, rather surprisingly, that when there are more sterile males there are more adults (and vice versa), and when there are more sterile males there are more foals (and vice versa). This is borne out by a correlation analysis. The correlation coefficient between foals and sterile males is 0.74, and the correlation coefficient between adults and sterile males is 0.68. You should find this very surprising — how do the horses know how many sterile males there are in the herd? You might think that this is an effect of scaling the plot, but there is a scatter plot in normalized coordinates in figure 3.21 that is entirely consistent with the conclusions suggested by the unnormalized plot. What is going on here?

The answer is revealed by the scatter plots of figure 3.22. Here, rather than plotting a '\*' at each data point, I have plotted the day number of the observation. This is in days from the first observation. You can see that the whole herd is shrinking — observations where there are many adults (resp. sterile adults, foals) occur with small day numbers, and observations where there are few have large day numbers. Because the whole herd is shrinking, it is true that when there are more adults and more sterile males, there are also more foals. Alternatively, you can see the plots of figure 3.20 as a scatter plot of herd size (resp. number of foals, number of sterile males) against day number. Then it becomes clear that the whole herd is shrinking, as is the size of each group. To drive this point home, we can look at the correlation coefficient between adults and days (-0.24), between sterile adults and days (-0.37), and between foals and days (-0.61). We can use the rule of thumb in box 3.3 to interpret this. This means that every 282 days, the herd loses about three adults; about one sterile adult; and about three foals. For the herd to have a stable size, it needs to gain by birth as many foals as it loses both to growing up and to death. If the herd is losing three foals every 282 days, then if they all grow

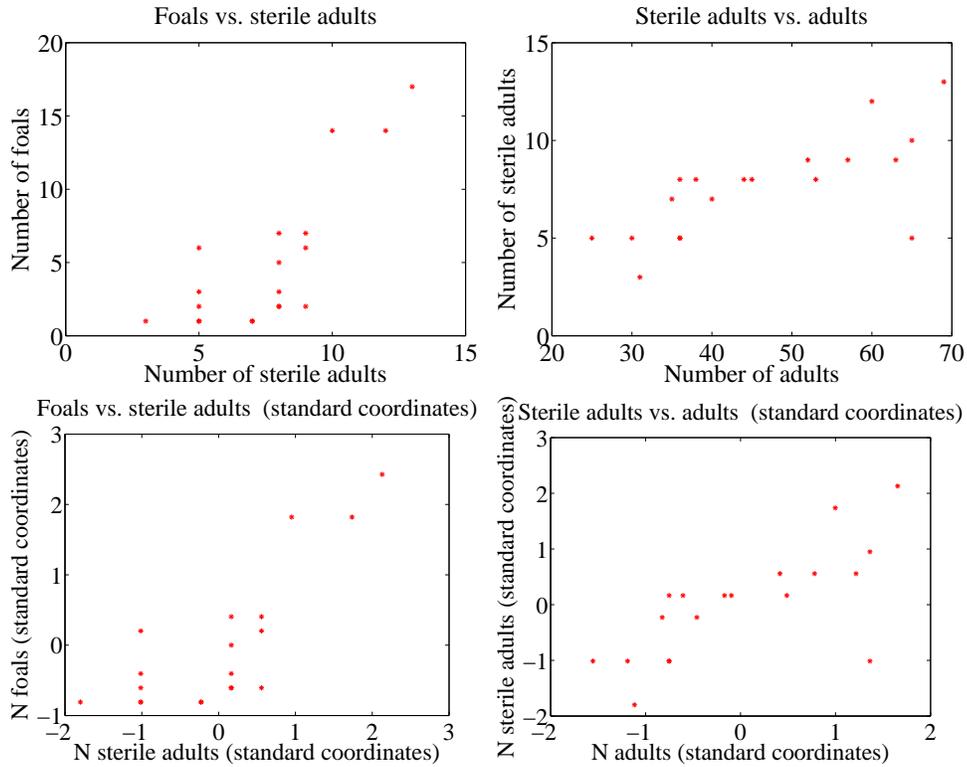


FIGURE 3.21: Scatter plots of the number of sterile males in a horse herd against the number of adults, and the number of foals against the number of sterile males, from data of <http://lib.stat.cmu.edu/DASL/Datafiles/WildHorses.html>. **Top:** unnormalized; **bottom:** standard coordinates.

up to replace the missing adults, the herd will be shrinking slightly (because it is losing four adults in this time); but if it loses foals to natural accidents, etc., then it is shrinking rather fast.

The message of this example is important. To understand a simple dataset, you might need to plot it several ways. You should make a plot, look at it and ask what it says, and then try to use another type of plot to confirm or refute what you think might be going on.

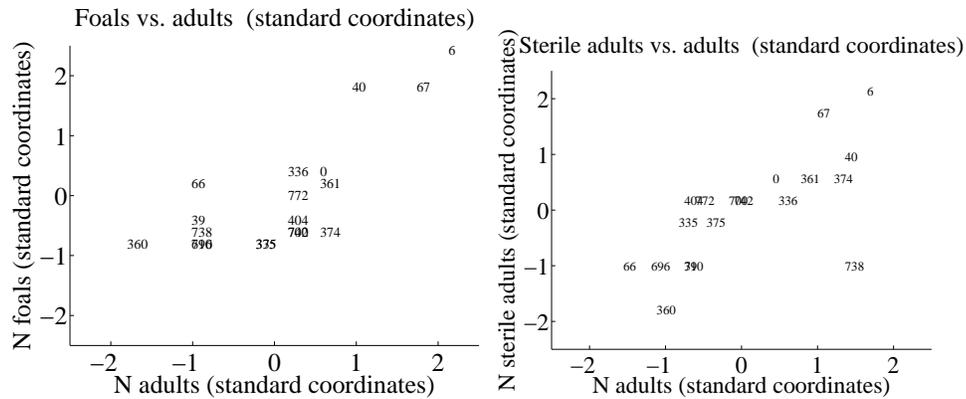


FIGURE 3.22: Scatter plots of the number of foals vs. the number of adults and the number of adults vs. the number of sterile adults for the wild horse herd, from <http://lib.stat.cmu.edu/DASL/Datafiles/WildHorses.html>. Rather than plot data points as dots, I have plotted the day on which the observation was made. Notice how the herd starts large, and then shrinks.

### 3.4 YOU SHOULD

#### 3.4.1 be able to:

- Plot a bar chart, a heat map, and a pie chart for a categorical dataset.
- Plot a dataset as a graph, making sensible choices about markers, lines and the like.
- Plot a scatter plot for a dataset.
- Plot a normalized scatter plot for a dataset.
- Interpret the scatter plot to tell the sign of the correlation between two variables, and estimate the size of the correlation coefficient.
- Compute a correlation coefficient.
- Interpret a correlation coefficient.
- Use correlation to make predictions.

#### 3.4.2 remember:

New term: pie chart . . . . .	45
New term: stacked bar chart . . . . .	45
New term: heat map . . . . .	45
New term: 3D bar chart . . . . .	47
There are a variety of tools for plotting categorical data. . . . .	48
New term: scatter plot . . . . .	50
A scatter plot should be your first step with a new 2D dataset. . . . .	52

It's usually a good idea to plot in standard coordinates. . . . .	56
New term: correlation . . . . .	56
New term: correlation coefficient . . . . .	58
New term: correlation . . . . .	58
Definition: Correlation coefficient . . . . .	59
New term: latent variable . . . . .	67

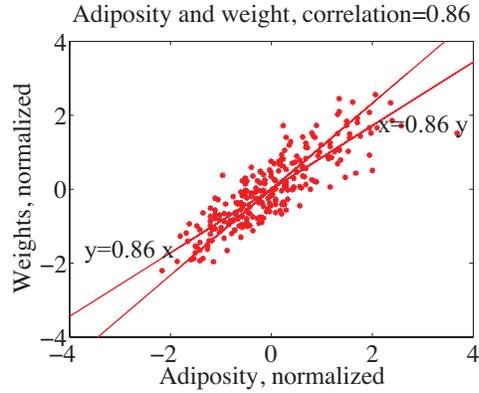


FIGURE 3.23: This figure shows two lines,  $y = 0.86x$  and  $x = 0.86y$ , superimposed on the normalized adiposity-weight scatter plot.

## PROBLEMS

- 3.1.** In a population, the correlation coefficient between weight and adiposity is 0.9. The mean weight is 150lb. The standard deviation in weight is 30lb. Adiposity is measured on a scale such that the mean is 0.8, and the standard deviation is 0.1.
- Using this information, predict the expected adiposity of a subject whose weight is 170lb
  - Using this information, predict the expected weight of a subject whose adiposity is 0.75
  - How reliable do you expect this prediction to be? Why? (your answer should be a property of correlation, not an opinion about adiposity or weight)
- 3.2.** In a population, the correlation coefficient between family income and child IQ is 0.30. The mean family income was \$60,000. The standard deviation in income is \$20,000. IQ is measured on a scale such that the mean is 100, and the standard deviation is 15.
- Using this information, predict the expected IQ of a child whose family income is \$70,000
  - How reliable do you expect this prediction to be? Why? (your answer should be a property of correlation, not an opinion about IQ)
  - The family income now rises—does the correlation predict that the child will have a higher IQ? Why?
- 3.3.** Show that  $\text{corr}(\{(x, y)\}) = \text{corr}(\{(y, x)\})$  by substituting into the definition.
- 3.4.** Show that if  $\hat{y}$  tends to be small (resp. large) for large (resp. small) values of  $\hat{x}$ , then the correlation coefficient will be negative.
- 3.5.** We have a 2D dataset consisting of  $N$  pairs  $(\hat{x}_i, \hat{y}_i)$  in normalized coordinates. This data has correlation coefficient  $r$ . We observe a new  $\hat{y}$  value  $\hat{y}_0$ , and wish to predict the (unknown)  $x$  value. We will do so with a linear prediction, choosing  $a, b$ , to predict an  $\hat{x}$  for any  $\hat{y}$  using the rule  $\hat{x}^P = a\hat{y}^P + b$ . Write  $u_i = \hat{x}_i - \hat{x}_i^P$  for the error that this rule makes on each data item.
- We require that  $\text{mean}(\{u\}) = 0$ . Show that this means that  $b = 0$ .
  - We require that  $\text{var}(\{u\})$  is minimized. Show that this means that  $a = r$ .

- (c) We now have a result that seems paradoxical — if I have  $(\hat{x}_0, ?)$  I predict  $(\hat{x}_0, r\hat{x}_0)$  and if I have  $(?, y_0)$ , I predict  $(r\hat{y}_0, \hat{y}_0)$ . Use figure 3.23 to explain why this is right. The important difference between the two lines is that one lies (approximately) in the middle of each vertical span of data, and the other lies (approximately) in the middle of each horizontal span of data.
- 3.6. I did the programming exercise about the earth temperature below. I looked at the years 1965-2012. Write  $\{(y, T)\}$  for the dataset giving the temperature ( $T$ ) of the earth in year  $y$ . I computed:  $\text{mean}(\{y\}) = 1988.5$ ,  $\text{std}(y) = 14$ ,  $\text{mean}(\{T\}) = 0.175$ ,  $\text{std}(T) = 0.231$  and  $\text{corr}(\{y\})T = 0.892$ . What is the best prediction using this information for the temperature in mid 2014? in mid 2028? in mid 2042?
- 3.7. I did the programming exercise about the earth temperature below. It is straightforward to build a dataset  $\{(T, n_t)\}$  where each entry contains the temperature of the earth ( $T$ ) and the number of counties where FEMA declared tornadoes  $n_t$  (for each year, you look up  $T$  and  $n_t$ , and make a data item). I computed:  $\text{mean}(\{T\}) = 0.175$ ,  $\text{std}(T) = 0.231$ ,  $\text{mean}(\{n_t\}) = 31.6$ ,  $\text{std}(n_t) = 30.8$ , and  $\text{corr}(\{T\})n_t = 0.471$ . What is the best prediction using this information for the number of tornadoes if the global earth temperature is 0.5? 0.6? 0.7?

## PROGRAMMING EXERCISES

- 3.8. At <http://lib.stat.cmu.edu/DASL/Datafiles/cigcancerdat.html>, you will find a dataset recording per capita cigarette sales and cancer deaths per 100K population for a variety of cancers, recorded for 43 states and the District of Columbia in 1960.
- (a) Plot a scatter plot of lung cancer deaths against cigarette sales, using the two letter abbreviation for each state as a marker. You should see two fairly obvious outliers. The backstory at <http://lib.stat.cmu.edu/DASL/Stories/cigcancer.html> suggests that the unusual sales in Nevada are generated by tourism (tourists go home, and die there) and the unusual sales in DC are generated by commuting workers (who also die at home).
- (b) What is the correlation coefficient between per capita cigarette sales and lung cancer deaths per 100K population? Compute this with, and without the outliers. What effect did the outliers have? Why?
- (c) What is the correlation coefficient between per capita cigarette sales and bladder cancer deaths per 100K population? Compute this with, and without the outliers. What effect did the outliers have? Why?
- (d) What is the correlation coefficient between per capita cigarette sales and kidney cancer deaths per 100K population? Compute this with, and without the outliers. What effect did the outliers have? Why?
- (e) What is the correlation coefficient between per capita cigarette sales and leukemia deaths per 100K population? Compute this with, and without the outliers. What effect did the outliers have? Why?
- (f) You should have computed a positive correlation between cigarette sales and lung cancer deaths. Does this mean that smoking causes lung cancer? Why?
- (g) You should have computed a negative correlation between cigarette sales and leukemia deaths. Does this mean that smoking cures leukemia? Why?
- 3.9. At <http://www.cru.uea.ac.uk/cru/info/warming/gtc.csv>, you can find a dataset of global temperature by year. When I accessed this, the years spanned 1880-

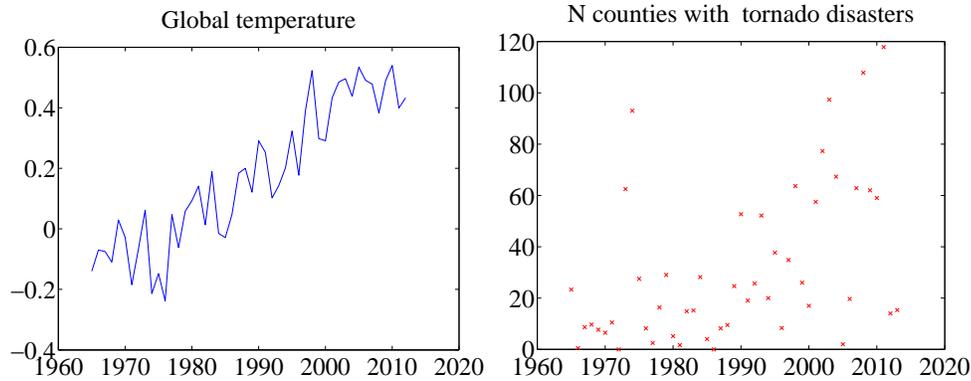


FIGURE 3.24: Plots I prepared from **left** *uea* data on temperature and **right** *FEMA* data on tornadoes by county. These should help you tell if you're on the right track.

2012. I don't know what units the temperatures are measured in. Keep in mind that measuring the temperature of the earth has non-trivial difficulties (you can't just insert an enormous thermometer!), and if you look at <http://www.cru.uea.ac.uk/cru> and <http://www.cru.uea.ac.uk/cru/data/temperature/> you can see some discussion of the choices made to get these measurements. There are two kinds of data in this dataset, smoothed and unsmoothed. I used the unsmoothed data, which should be fine for our purposes. The government publishes a great deal of data at <http://data.gov>. From there, I found a dataset, published by the Federal Emergency Management Agency (FEMA), of all federally declared disasters (which I found at <http://www.fema.gov/media-library/assets/documents/28318?id=6292>). We would like to see whether weather-related disasters are correlated to global temperature.

(a) The first step is preprocessing the data. The FEMA data has all sorts of information. From 1965 on, a disaster was declared per county (so you get one line in the data set for each county), but before 1965, it seems to have been by state. We divide the disasters into four types: TORNADO, FLOOD, STORM, HURRICANE. (FEMA seems to have a much richer type system). We want to know how many counties declare a disaster of each type, in each year. This is a really rough estimate of the number of people affected by the disaster. If a disaster has two types (in some rows, you will see "SEVERE STORMS, HEAVY RAINS & FLOODING" we will allocate the credit evenly between the types (i.e. for this case we would count 1/2 for STORM and 1/2 for FLOOD). You should write code that will (a) read the dataset and then (b) compute a table with the count of the number of counties where a disaster of each type has occurred for each year. This takes a bit of work. Notice you only need to deal with two columns of the data (the date it was declared, and the type). Notice also that FEMA changed the way it represented dates somewhere through the first column (they added times), which can cause problems. You can tell the type of the disaster by just using a string match routine with the four keywords. Figure 3.24 shows the plot of temperature and of number of counties where FEMA declared a tornado disaster for this data.

- (b) Plot a normalized scatter plot of the number of counties where FEMA declared the disaster against temperature, for each kind.
  - (c) For each kind of disaster, compute the correlation coefficient between the number of counties where FEMA declared the disaster and the year. For each kind of disaster, use this correlation coefficient to predict the number of disasters of this kind for 2013. Compare this to the true number, and explain what you see.
  - (d) For each kind of disaster, compute the correlation coefficient between the number of counties where FEMA declared the disaster and the global temperature. For each kind of disaster, use this correlation coefficient to predict the number of disasters of this kind when the earth reaches 0.6 temperature units and 0.7 temperature units (on the absolute temperature scale).
  - (e) Does this data show that warming of the earth causes weather disasters? Why?
  - (f) Does this data suggest that more people will be affected by disasters in the US in the future? Why?
  - (g) Does this data suggest that the earth will be warmer in the future? Why?
- 3.10.** If you go to <https://github.com/TheUpshot/Military-Surplus-Gear>, you will find data on purchases of military weapons by US police departments. This data is organized by state and county. There's a fair amount of data here, and you'll need to do some data jockeying.
- (a) Prepare a plot showing how much each Illinois county spent under this program.
  - (b) Now look up population levels in the counties. Prepare a plot showing how much each county spent *per capita*.
  - (c) Prepare a graphic illustrating what the overall most popular items were — i.e., those items counties bought the most of.
  - (d) Prepare a graphic illustrating on what items the most money was spent — for example, was more money spent on “RIFLE,5.56 MILLIMETER” or on “MINE RESISTANT VEHICLE”?
  - (e) Prepare a graphic illustrating the pattern of purchases across counties for the ten overall most popular items.
  - (f) Can you draw any interesting conclusions?

## CHAPTER 4

# Basic ideas in probability

We will perform experiments — which could be pretty much anything, from flipping a coin, to eating too much saturated fat, to smoking, to crossing the road without looking — and reason about the outcomes (mostly bad for the examples I gave). But these outcomes are uncertain, and we need to weigh those uncertainties against one another. If I flip a coin, I could get heads or tails, and there's no reason to expect to see one more often than the other. If I eat too much saturated fat or smoke, I will very likely have problems, though I might not. If I cross the road without looking, I may be squashed by a truck or I may not. Our methods need also to account for information. If I look before I cross the road, I am much less likely to be squashed.

Probability is the machinery we use to predict what will happen in an experiment. Probability measures the tendency of events to occur frequently or seldom when we repeat experiments. Building this machinery involves a formal model of potential outcomes and sets of outcomes. Once we have this, a set of quite simple axioms allows us, at least in principle, to compute probabilities in most situations.

### 4.1 EXPERIMENTS, OUTCOMES, EVENTS, AND PROBABILITY

If we flip a fair coin many times, we expect it to come up heads about as often as it comes up tails. If we toss a fair die many times, we expect each number to come up about the same number of times. We are performing an experiment each time we flip the coin, and each time we toss the die. We can formalize this experiment by describing the set of **outcomes** that we expect from the experiment. Every run of the experiment produces one of the set of possible outcomes. There are never two or more outcomes, and there is never no outcome.

In the case of the coin, the set of outcomes is:

$$\{H, T\}.$$

In the case of the die, the set of outcomes is:

$$\{1, 2, 3, 4, 5, 6\}.$$

We are making a modelling choice by specifying the outcomes of the experiment. For example, we are assuming that the coin can only come up heads or tails (but doesn't stand on its edge; or fall between the floorboards; or land behind the bookcase; or whatever).

**Definition: 4.1** *Sample space*

The sample space is the set of all outcomes, which we usually write  $\Omega$ .

Indexing: Sample space

Indexing: Sample space

**Worked example 4.1** *Find the lady*

We have three playing cards. One is a queen; one is a king, and one is a knave. All are shown face down, and one is chosen at random and turned up. What is the set of outcomes?

**Solution:** Write Q for queen, K for king, N for knave; the outcomes are  $\{Q, K, N\}$

**Worked example 4.2** *Find the lady, twice*

We play Find the Lady twice, replacing the card we have chosen. What is the sample space?

**Solution:** We now have  $\{QQ, QK, QN, KQ, KK, KN, NQ, NK, NN\}$

**Worked example 4.3** *Children*

A couple decides to have children until either (a) they have both a boy and a girl or (b) they have three children. What is the set of outcomes?

**Solution:** Write B for boy, G for girl, and write them in birth order; we have  $\{BG, GB, BBG, BBB, GGB, GGG\}$ .

**Worked example 4.4** *Monty Hall (sigh!) with indistinguishable goats*

There are three boxes. There is a goat, a second goat, and a car. These are placed into the boxes at random. The goats are indistinguishable for our purposes; equivalently, we do not care about the difference between goats. What is the sample space?

**Solution:** Write G for goat, C for car. Then we have  $\{CGG, GCG, GGC\}$ .

**Worked example 4.5** *Monty Hall with distinguishable goats*

There are three boxes. There is a goat, a second goat, and a car. These are placed into the boxes at random. One goat is male, the other female, and the distinction is important. What is the sample space?

**Solution:** Write M for male goat, F for female goat, C for car. Then we have  $\{CFM, CMF, FCM, MCF, FMC, MFC\}$ . Notice how the number of outcomes has increased, because we now care about the distinction between goats.

**Worked example 4.6** *A poor choice of strategy for planning a family*

A couple decides to have children. As they know no mathematics, they decide to have children until a girl then a boy are born. What is the sample space? Does this strategy bound the number of children they could be planning to have?

**Solution:** Write B for boy, G for girl. The sample space looks like any string of B's and G's that (a) ends in GB and (b) does not contain any other GB. In regular expression notation, you can write such strings as  $B^*G^+B$ . There is a lower bound on the length of the string (two), but no upper bound. As a family planning strategy, this is unrealistic, but it serves to illustrate the point that sample spaces don't have to be finite to be tractable.

**Useful Facts: 4.1** *Sample space*

You must specify a sample space to build a probability model. Sample spaces do not need to be finite.

Indexing: Sample space

Indexing: Sample space

## 4.1.1 The Probability of an Outcome

We represent our model of how often a particular outcome will occur in a repeated experiment with a **probability**, a non-negative number. This number gives the relative frequency of the outcome of interest, when an experiment is repeated a very large number of times.

Assume that we repeat an experiment  $N$  times. Assume also that the coins, dice, whatever involved in each repetition of the experiment don't communicate

with one another from experiment to experiment (or, equivalently, that experiments don't "know" about one another). We describe these experiments as **independent**. We say that an outcome  $A$  has probability  $P$  if (a) in about  $N \times P$  of those experiments the outcome will occur and (b) as  $N$  gets larger, the fraction where the outcome occurs will get closer to  $P$ . We write  $\#(A)$  for the number of times outcome  $A$  occurs. We interpret  $P$  as

$$\lim_{N \rightarrow \infty} \frac{\#(A)}{N}.$$

We can draw two important conclusions immediately.

- For any outcome  $A$ ,  $0 \leq P(A) \leq 1$ .
- $\sum_{A_i \in \Omega} P(A_i) = 1$ .

Remember that every run of the experiment produces exactly one outcome. The probabilities add up to one because each experiment must have one of the outcomes in the sample space.

**Worked example 4.7** *A biased coin*

Assume we have a coin where the probability of getting heads is  $P(H) = \frac{1}{3}$ , and so the probability of getting tails is  $P(T) = \frac{2}{3}$ . We flip this coin three million times. How many times do we see heads?

**Solution:**  $P(H) = \frac{1}{3}$ , so we expect this coin will come up heads in  $\frac{1}{3}$  of experiments. This means that we will very likely see very close to a million heads. Later on, we will be able to be more precise.

Some problems can be handled by building a set of outcomes and reasoning about the probability of each outcome. This is particularly useful when the outcomes *must* have the same probability, which happens rather a lot.

**Worked example 4.8** *Find the Lady, with even probabilities*

Recall the problem of worked example ???. Assume that the card that is chosen is chosen fairly — that is, each card is chosen with the same probability. What is the probability of turning up a Queen?

**Solution:** There are three outcomes, and each is chosen with the same probability, so the probability is  $1/3$ .

**Worked example 4.9** *Monty Hall, indistinguishable goats, even probabilities*

Recall the problem of worked example 4.5. Each outcome has the same probability. We choose to open the first box. With what probability will we find a goat (any goat)?

**Solution:** There are three outcomes, each has the same probability, and two give a goat, so  $2/3$

**Worked example 4.10** *Monty Hall, yet again*

Each outcome has the same probability. We choose to open the first box. With what probability will we find the car?

**Solution:** There are three places the car could be, each has the same probability, so  $1/3$

**Worked example 4.11** *Monty Hall, with distinct goats, again*

Each outcome has the same probability. We choose to open the first box. With what probability will we find a female goat?

**Solution:** Using the reasoning of the previous example, but substituting “female goat” for “car”,  $1/3$ . The point of this example is that the sample space matters. If you care about the gender of the goat, then it’s important to keep track of it; if you don’t, it’s a good idea to omit it from the sample space.

**Remember this:** *The probability of an outcome is the frequency of that outcome in a very large number of repeated experiments. The sum of probabilities over all outcomes must be one.*

#### 4.1.2 Events

Assume we run an experiment and get an outcome. We know what the outcome is (that’s the whole point of a sample space). This means we can tell whether the

outcome we get belongs to some particular known *set* of outcomes. We just look in the set and see if our outcome is there. This means that we should be able to predict the probability of a *set* of outcomes from any reasonable model of an experiment. For example, we might roll a die and ask what the probability of getting an even number is. As another example, we might flip a coin ten times, and ask what the probability of getting three heads is.

An **event** is a set of outcomes. The set of all outcomes, which we wrote  $\Omega$ , must be an event. It is not a particularly interesting event, because we must have  $P(\Omega) = 1$  (because we said that every run of an experiment produces one outcome, and that outcome must be in  $\Omega$ ). In principle, there could be no outcome, although this never happens. This means that the empty set, which we write  $\emptyset$ , is an event, and we have  $P(\emptyset) = 0$ . The space of events has a rich structure, which makes it possible to compute probabilities for other, more interesting, events.

**Notation:** Generally, we write sets like  $\mathcal{A}$ ; in principle, you could confuse this notation with the matrix notation, but it's clear from context which is meant. We write  $\mathcal{A} \cup \mathcal{B}$  for the union of two sets,  $\mathcal{A} \cap \mathcal{B}$  for the intersection of two sets, and  $\mathcal{A} - \mathcal{B}$  for the set theoretic difference (i.e.  $\mathcal{A} - \mathcal{B} = \{x \in \mathcal{A} | x \notin \mathcal{B}\}$ ). We will write  $\Omega - \mathcal{U}$  as  $\mathcal{U}^c$ ; read “the complement of  $\mathcal{U}$ ”.

Events have three important properties that follow from their nature as sets of outcomes:

- If  $\mathcal{U}$  and  $\mathcal{V}$  are events — sets of outcomes — then so is  $\mathcal{U} \cap \mathcal{V}$ . You should interpret this as the event that we have an outcome that is in  $\mathcal{U}$  and also in  $\mathcal{V}$ .
- If  $\mathcal{U}$  and  $\mathcal{V}$  are events, then  $\mathcal{U} \cup \mathcal{V}$  is also an event. You should interpret this as the event that we have an outcome that is either in  $\mathcal{U}$  or in  $\mathcal{V}$  (or in both).
- If  $\mathcal{U}$  is an event, then  $\mathcal{U}^c = \Omega - \mathcal{U}$  is also an event. You should think of this as the event we get an outcome that is not in  $\mathcal{U}$ .

This means that the set of all possible events has a very important structure.

#### Useful Facts: 4.2 *Event*

An event is a set of outcomes. We usually write the set of all possible events  $\Sigma$ . This set has a very important structure:

- $\emptyset$  is in  $\Sigma$ .
- $\Omega$  is in  $\Sigma$ .
- If  $\mathcal{U} \in \Sigma$  and  $\mathcal{V} \in \Sigma$  then  $\mathcal{U} \cup \mathcal{V} \in \Sigma$ .
- If  $\mathcal{U} \in \Sigma$  and  $\mathcal{V} \in \Sigma$  then  $\mathcal{U} \cap \mathcal{V} \in \Sigma$ .
- If  $\mathcal{U} \in \Sigma$  then  $\mathcal{U}^c \in \Sigma$ .

Indexing: Event

Indexing: Event

This means that the space of events can be quite big. For a single flip of a coin, the only possible space of events looks like:

$$\{\emptyset, \{H\}, \{T\}, \{H, T\}\}$$

Many experiments admit more than one space of events. For example, if we flip two coins, one natural event space is

$$\left\{ \begin{array}{l} \emptyset, \Omega, \\ \{HH\}, \{HT\}, \{TH\}, \{TT\}, \\ \{HH, HT\}, \{HH, TH\}, \{HH, TT\}, \{HT, TT\}, \{HT, TH\}, \{TT, HH\}, \\ \{HT, TH, TT\}, \{HH, TH, TT\}, \{HH, HT, TT\}, \{HH, HT, TH\} \end{array} \right\}$$

which can represent any possible event that can be built out of two coin flips. But this is not the only event space possible with these outcomes.

**Worked example 4.12** *The structure of event spaces*

I flip two coins. Is the following collection of sets an event space?

$$\Sigma = \{\emptyset, \Omega, \{HH\}, \{TT, HT, TH\}\}$$

**Solution:** Yes, because:  $\emptyset \in \Sigma$ ;  $\Omega \in \Sigma$ ; if  $\mathcal{A} \in \Sigma$ ,  $\mathcal{A}^c \in \Sigma$ ; if  $\mathcal{A} \in \Sigma$  and  $\mathcal{B} \in \Sigma$ ,  $\mathcal{A} \cup \mathcal{B} \in \Sigma$ ; and if  $\mathcal{A} \in \Sigma$  and  $\mathcal{B} \in \Sigma$ ,  $\mathcal{A} \cap \mathcal{B} \in \Sigma$ . You can interpret the event space as encoding the events “two heads” and “anything other than two heads”.

So, as example 4.12 illustrates, we can have different consistent collections of events built on top of the same set of outcomes. This makes sense, because it allows us to reason about different kinds of result obtained with the same experimental equipment.

For a single throw of the die, the set of every possible event is

$$\left\{ \begin{array}{cccccc} \emptyset, & \{1, 2, 3, 4, 5, 6\}, & & & & \\ \{1\}, & \{2\}, & \{3\}, & \{4\}, & \{5\}, & \{6\}, \\ & \{1, 2\}, & \{1, 3\}, & \{1, 4\}, & \{1, 5\}, & \{1, 6\}, \\ & & \{2, 3\}, & \{2, 4\}, & \{2, 5\}, & \{2, 6\}, \\ & & & \{3, 4\}, & \{3, 5\}, & \{3, 6\}, \\ & & & & \{4, 5\}, & \{4, 6\}, \\ & & & & & \{5, 6\}, \\ & \{1, 2, 3\}, & \{1, 2, 4\}, & \{1, 2, 5\}, & \{1, 2, 6\}, & \\ & & \{1, 3, 4\}, & \{1, 3, 5\}, & \{1, 3, 6\}, & \\ & & & \{1, 4, 5\}, & \{1, 4, 6\}, & \\ & & & & \{1, 5, 6\}, & \\ & & \{2, 3, 4\}, & \{2, 3, 5\}, & \{2, 3, 6\}, & \\ & & \{2, 4, 5\}, & \{2, 4, 6\}, & \{2, 5, 6\}, & \\ & & & \{3, 4, 5\}, & \{3, 4, 6\}, & \\ & & & & \{3, 5, 6\}, & \\ & & & & \{4, 5, 6\}, & \\ & \{1, 2, 3, 4\}, & \{1, 2, 3, 5\}, & \{1, 2, 3, 6\}, & & \\ & & \{1, 3, 4, 5\}, & \{1, 3, 4, 6\}, & & \\ & & \{2, 3, 4, 5\}, & \{2, 3, 4, 6\}, & & \\ & & & \{3, 4, 5, 6\}, & & \\ \{2, 3, 4, 5, 6\}, & \{1, 3, 4, 5, 6\}, & \{1, 2, 4, 5, 6\}, & \{1, 2, 3, 5, 6\}, & \{1, 2, 3, 4, 6\}, & \{1, 2, 3, 4, 5\} \end{array} \right\}$$

(which gives some explanation as to why we don't usually write out the whole space of events). In fact, it is seldom necessary to explain which event space one is working with. We usually assume that the event space consists of all events that can be obtained with the outcomes (as in the event space shown for a die).

### 4.1.3 The Probability of Events

So far, we have described the probability of each outcome with a non-negative number. This number represents the relative frequency of the outcome. Because we can tell when an event has occurred, we can compute the relative frequency of events, too. Because it is a relative frequency, the probability of an event is a non-negative number, and is no greater than one. But the probability of events must be consistent with the probability of outcomes. This implies a set of quite straightforward properties:

- **The probability of every event is non-negative**, which we write  $P(\mathcal{A}) \geq 0$  for all  $\mathcal{A}$  in the collection of events.
- **Every experiment has an outcome**, which we write  $P(\Omega) = 1$ .
- **The probability of disjoint events is additive**, which requires more notation. Assume that we have a collection of events  $\mathcal{A}_i$ , indexed by  $i$ . We require that these have the property  $\mathcal{A}_i \cap \mathcal{A}_j = \emptyset$  when  $i \neq j$ . This means that there is no outcome that appears in more than one  $\mathcal{A}_i$ . In turn, if we interpret probability as relative frequency, we must have that  $P(\cup_i \mathcal{A}_i) = \sum_i P(\mathcal{A}_i)$ .

Any function  $P$  taking events to numbers that has these properties is a probability. These very simple properties imply a series of other very important properties.

**Useful Facts: 4.3** *Probability of events*

- $P(\mathcal{A}^c) = 1 - P(\mathcal{A})$
- $P(\emptyset) = 0$
- $P(\mathcal{A} - \mathcal{B}) = P(\mathcal{A}) - P(\mathcal{A} \cap \mathcal{B})$
- $P(\mathcal{A} \cup \mathcal{B}) = P(\mathcal{A}) + P(\mathcal{B}) - P(\mathcal{A} \cap \mathcal{B})$
- $P(\cup_1^n \mathcal{A}_i) = \sum_i P(\mathcal{A}_i) - \sum_{i < j} P(\mathcal{A}_i \cap \mathcal{A}_j) + \sum_{i < j < k} P(\mathcal{A}_i \cap \mathcal{A}_j \cap \mathcal{A}_k) + \dots (-1)^{(n+1)} P(\mathcal{A}_1 \cap \mathcal{A}_2 \cap \dots \cap \mathcal{A}_n)$

Indexing: Probability of events

Indexing: Probability of events

I prove each of these below. Looking at the useful facts should suggest a helpful analogy. Think about the probability of an event as the “size” of that event. This “size” is relative to  $\Omega$ , which has “size” 1. I find this a good way to remember equations. Some people find Venn diagrams a useful way to keep track of this argument, and Figure 4.1 is for them.

This analogy allows us to explain  $P(\mathcal{A}) + P(\mathcal{A}^c) = 1$  by noticing that  $\mathcal{A}$  and  $\mathcal{A}^c$  don’t overlap, and together make up all of  $\Omega$ . So the “size” of  $\mathcal{A}$  and the “size” of  $\mathcal{A}^c$  should add to the “size” of  $\Omega$ .

We explain  $P(\mathcal{A} - \mathcal{B}) = P(\mathcal{A}) - P(\mathcal{A} \cap \mathcal{B})$  by noticing the “size” of the part of  $\mathcal{A}$  that isn’t  $\mathcal{B}$  is obtained by taking the “size” of  $\mathcal{A}$  and subtracting the “size” of the part that is also in  $\mathcal{B}$ . This analogy, which can be made precise by thinking about “size” in the right way, also explains  $P(\mathcal{A} \cup \mathcal{B}) = P(\mathcal{A}) + P(\mathcal{B}) - P(\mathcal{A} \cap \mathcal{B})$ . You can get the “size” of  $\mathcal{A} \cup \mathcal{B}$  by adding the two “sizes”, then subtracting the “size” of the intersection because otherwise you would double-count the part where the two sets overlap.

**Proposition:**  $P(\mathcal{A}^c) = 1 - P(\mathcal{A})$

**Proof:**  $\mathcal{A}^c$  and  $\mathcal{A}$  are disjoint, so that  $P(\mathcal{A}^c \cup \mathcal{A}) = P(\mathcal{A}^c) + P(\mathcal{A}) = P(\Omega) = 1$ .

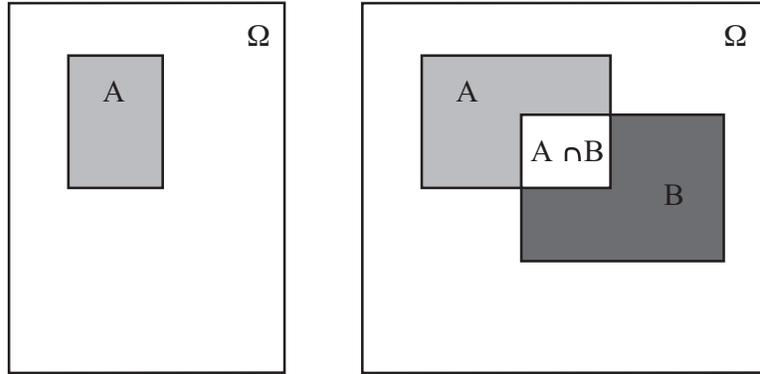


FIGURE 4.1: If you think of the probability of an event as measuring its “size”, many of the rules are quite straightforward to remember. Venn diagrams can sometimes help. On the **left**, a Venn diagram to help remember that  $P(\mathcal{A}) + P(\mathcal{A}^c) = 1$ . The “size” of  $\Omega$  is 1, outcomes lie either in  $\mathcal{A}$  or  $\mathcal{A}^c$ , and the two don’t intersect. On the **right**, you can see that  $P(\mathcal{A} - \mathcal{B}) = P(\mathcal{A}) - P(\mathcal{A} \cap \mathcal{B})$  by noticing that  $P(\mathcal{A} - \mathcal{B})$  is the “size” of the part of  $\mathcal{A}$  that isn’t  $\mathcal{B}$ . This is obtained by taking the “size” of  $\mathcal{A}$  and subtracting the “size” of the part that is also in  $\mathcal{B}$ , i.e. the “size” of  $\mathcal{A} \cap \mathcal{B}$ . Similarly, you can see that  $P(\mathcal{A} \cup \mathcal{B}) = P(\mathcal{A}) + P(\mathcal{B}) - P(\mathcal{A} \cap \mathcal{B})$  by noticing that you can get the “size” of  $\mathcal{A} \cup \mathcal{B}$  by adding the “sizes” of  $\mathcal{A}$  and  $\mathcal{B}$ , then subtracting the “size” of the intersection to avoid double counting.

**Proposition:**  $P(\emptyset) = 0$

**Proof:**  $P(\emptyset) = P(\Omega^c) = P(\Omega - \Omega) = 1 - P(\Omega) = 1 - 1 = 0$ .

**Proposition:**  $P(\mathcal{A} - \mathcal{B}) = P(\mathcal{A}) - P(\mathcal{A} \cap \mathcal{B})$

**Proof:**  $\mathcal{A} - \mathcal{B}$  is disjoint from  $\mathcal{A} \cap \mathcal{B}$ , and  $(\mathcal{A} - \mathcal{B}) \cup (\mathcal{A} \cap \mathcal{B}) = \mathcal{A}$ . This means that  $P(\mathcal{A} - \mathcal{B}) + P(\mathcal{A} \cap \mathcal{B}) = P(\mathcal{A})$ .

**Proposition:**  $P(\mathcal{A} \cup \mathcal{B}) = P(\mathcal{A}) + P(\mathcal{B}) - P(\mathcal{A} \cap \mathcal{B})$

**Proof:**  $P(\mathcal{A} \cup \mathcal{B}) = P(\mathcal{A} \cup (\mathcal{B} \cap \mathcal{A}^c)) = P(\mathcal{A}) + P((\mathcal{B} \cap \mathcal{A}^c))$ . Now  $\mathcal{B} = (\mathcal{B} \cap \mathcal{A}) \cup (\mathcal{B} \cap \mathcal{A}^c)$ . Furthermore,  $(\mathcal{B} \cap \mathcal{A})$  is disjoint from  $(\mathcal{B} \cap \mathcal{A}^c)$ , so we have  $P(\mathcal{B}) = P((\mathcal{B} \cap \mathcal{A}) + P((\mathcal{B} \cap \mathcal{A}^c))$ . This means that  $P(\mathcal{A} \cup \mathcal{B}) = P(\mathcal{A}) + P((\mathcal{B} \cap \mathcal{A}^c)) = P(\mathcal{A}) + P(\mathcal{B}) - P((\mathcal{B} \cap \mathcal{A}))$ .

**Proposition:**  $P(\cup_1^n \mathcal{A}_i) = \sum_i P(\mathcal{A}_i) - \sum_{i < j} P(\mathcal{A}_i \cap \mathcal{A}_j) + \sum_{i < j < k} P(\mathcal{A}_i \cap \mathcal{A}_j \cap \mathcal{A}_k) + \dots + (-1)^{(n+1)} P(\mathcal{A}_1 \cap \mathcal{A}_2 \cap \dots \cap \mathcal{A}_n)$

**Proof:** This can be proven by repeated application of the previous result. As an example, we show how to work the case where there are three sets (you can get the rest by induction).

$$\begin{aligned}
 P(\mathcal{A}_1 \cup \mathcal{A}_2 \cup \mathcal{A}_3) &= P(\mathcal{A}_1 \cup (\mathcal{A}_2 \cup \mathcal{A}_3)) \\
 &= P(\mathcal{A}_1) + P(\mathcal{A}_2 \cup \mathcal{A}_3) - P(\mathcal{A}_1 \cap (\mathcal{A}_2 \cup \mathcal{A}_3)) \\
 &= P(\mathcal{A}_1) + (P(\mathcal{A}_2) + P(\mathcal{A}_3) - P(\mathcal{A}_2 \cap \mathcal{A}_3)) - \\
 &\quad P((\mathcal{A}_1 \cap \mathcal{A}_2) \cup (\mathcal{A}_1 \cap \mathcal{A}_3)) \\
 &= P(\mathcal{A}_1) + (P(\mathcal{A}_2) + P(\mathcal{A}_3) - P(\mathcal{A}_2 \cap \mathcal{A}_3)) - \\
 &\quad P(\mathcal{A}_1 \cap \mathcal{A}_2) - P(\mathcal{A}_1 \cap \mathcal{A}_3) \\
 &\quad - (-P((\mathcal{A}_1 \cap \mathcal{A}_2) \cap (\mathcal{A}_1 \cap \mathcal{A}_3))) \\
 &= P(\mathcal{A}_1) + P(\mathcal{A}_2) + P(\mathcal{A}_3) - \\
 &\quad P(\mathcal{A}_2 \cap \mathcal{A}_3) - P(\mathcal{A}_1 \cap \mathcal{A}_2) - P(\mathcal{A}_1 \cap \mathcal{A}_3) + \\
 &\quad P(\mathcal{A}_1 \cap \mathcal{A}_2 \cap \mathcal{A}_3)
 \end{aligned}$$

#### 4.1.4 Computing Probabilities by Counting Outcomes

The rule  $P(\mathcal{A} \cup \mathcal{B}) = P(\mathcal{A}) + P(\mathcal{B}) - P(\mathcal{A} \cap \mathcal{B})$  yields a very useful procedure for computing the probability of some events. Imagine  $\mathcal{A}$  and  $\mathcal{B}$  are disjoint (so  $\mathcal{A} \cap \mathcal{B} = \emptyset$ ). Then  $P(\mathcal{A} \cup \mathcal{B})$  is just  $P(\mathcal{A}) + P(\mathcal{B})$ . So imagine some event  $\mathcal{E}$  that consists of a set of outcomes. Each singleton set — containing just one outcome — is disjoint from each other one. I can make the set  $\mathcal{E}$  by: starting with an empty set; unioning this with one of the outcomes in  $\mathcal{E}$ ; then repeatedly unioning the result a new outcome until I get  $\mathcal{E}$ . I am always unioning disjoint sets, so the probability of this event is the sum of probabilities of the outcomes. So we have

$$P(\mathcal{E}) = \sum_{O \in \mathcal{E}} P(O)$$

where  $O$  ranges over the outcomes in  $\mathcal{E}$ .

This is particularly useful when you know each outcome in  $\Omega$  has the same probability. In this case, you can show

$$P(\mathcal{E}) = \frac{\text{Number of outcomes in } \mathcal{E}}{\text{Total number of outcomes in } \Omega}$$

(look at the exercises). Such problems become, basically, advanced counting exercises.

**Worked example 4.13** *Odd numbers with fair dice*

We throw a fair (each number has the same probability) die twice, then add the two numbers. What is the probability of getting an odd number?

**Solution:** There are 36 outcomes. Each has the same probability ( $1/36$ ). 18 of them give an odd number, and the other 18 give an even number, so the probability is  $18/36 = 1/2$

**Worked example 4.14** *Drawing a red ten*

I shuffle a standard pack of cards, and draw one card. What is the probability that it is a red ten?

**Solution:** There are 52 cards, and each is an outcome. Two of these outcomes are red tens; so we have  $2/52 = 1/26$ .

**Worked example 4.15** *Numbers divisible by five with fair dice*

We throw a fair (each number has the same probability) die twice, then add the two numbers. What is the probability of getting a number divisible by five?

**Solution:** There are 36 outcomes. Each has the same probability ( $1/36$ ). For this event, the spots must add to either 5 or to 10. There are 4 ways to get 5. There are 3 ways to get 10, so the probability is  $7/36$ .

**Worked example 4.16** *Birthdays in succession*

We stop three people at random, and ask the day of the week on which they are born. What is the probability that they are born on three days of the week in succession (for example, the first on Monday; the second on Tuesday; the third on Wednesday; or Saturday-Sunday-Monday; and so on).

**Solution:** We assume that births are equally common on each day of the week. The space of outcomes consists of triples of days, and each outcome has the same probability. The event is the set of triples of three days in succession (which has seven elements, one for each starting day). The space of outcomes has  $7^3$  elements in it, so the probability is

$$\frac{\text{Number of outcomes in the event}}{\text{Total number of outcomes}} = \frac{7}{7^3} = \frac{1}{49}.$$

**Worked example 4.17** *Shared birthdays*

We stop two people at random. What is the probability that they were born on the same day of the week?

**Solution:** The day the first person was born doesn't matter; the probability the second person was born on that day is  $1/7$ . Or you could count outcomes explicitly to get

$$\frac{\text{Number of outcomes in the event}}{\text{Total number of outcomes}} = \frac{7}{7 \times 7} = \frac{1}{7}.$$

Sometimes a bit of fiddling with the space of outcomes makes it easy to compute what we want. Examples 4.19 and 4.20 show cases where you can use fictitious outcomes as an accounting device to simplify a computation.

**Worked example 4.18** *Children - 1*

*This example is a version of of example 1.12, p44, in Stirzaker, “Elementary Probability”. A couple decides to have children. They decide simply to have three children. Assume that boys and girls are equally likely at each birth. Let  $B_i$  be the event that there are  $i$  boys, and  $C$  be the event there are more girls than boys. Compute  $P(B_1)$  and  $P(C)$ .*

**Solution:** There are eight outcomes. Each has the same probability. Three of them have a single boy, so  $P(B_1) = 3/8$ .  $P(C) = P(C^c)$  (because  $C^c$  is the event there are more boys than than girls, AND the number of children is odd), so that  $P(C) = 1/2$ ; you can also get this by counting outcomes.

**Worked example 4.19** *Children - 2*

*This example is a version of of example 1.12, p44, in Stirzaker, “Elementary Probability”. A couple decides to have children. They decide to have children until the first girl is born, or until there are three, and then stop. Assume that each gender is equally likely at each birth. Let  $B_i$  be the event that there are  $i$  boys, and  $C$  be the event there are more girls than boys. Compute  $P(B_1)$  and  $P(C)$ .*

**Solution:** In this case, we could write the outcomes as  $\{G, BG, BBG\}$ , but if we think about them like this, we have no simple way to compute their probability. Instead, we could use the sample space from the previous answer, but assume that some of the later births are fictitious. This gives us natural collection of events for which it is easy to compute probabilities. Having one girl corresponds to the event  $\{Gbb, Gbg, Ggb, Ggg\}$ , where I have used lowercase letters to write the fictitious later births; the probability is  $1/2$ . Having a boy then a girl corresponds to the event  $\{BGb, BGg\}$  (and so has probability  $1/4$ ). Having two boys then a girl corresponds to the event  $\{BBG\}$  (and so has probability  $1/8$ ). Finally, having three boys corresponds to the event  $\{BBB\}$  (and so has probability  $1/8$ ). This means that  $P(B_1) = 1/4$  and  $P(C) = 1/2$ .

**Worked example 4.20** *Children - 2*

*This example is a version of of example 1.12, p44, in Stirzaker, “Elementary Probability”. A couple decides to have children. They decide to have children until there is one of each gender, or until there are three, and then stop. Assume that each gender is equally likely at each birth. Let  $B_i$  be the event that there are  $i$  boys, and  $C$  be the event there are more girls than boys. Compute  $P(B_1)$  and  $P(C)$ .*

**Solution:** We could write the outcomes as  $\{GB, BG, GGB, GGG, BBG, BBB\}$ . Again, if we think about them like this, we have no simple way to compute their probability; so we use the sample space from the previous example with the device of the fictitious births again. The important events are  $\{GBb, GBg\}$ ;  $\{BGb, BGg\}$ ;  $\{GGB\}$ ;  $\{GGG\}$ ;  $\{BBG\}$ ; and  $\{BBB\}$ . Like this, we get  $P(B_1) = 5/8$  and  $P(C) = 1/4$ .

**Remember this:** *In some problems, you can compute the probabilities of events by counting outcomes.*

## 4.1.5 Computing Probabilities by Reasoning about Sets

The rule  $P(\mathcal{A}^c) = 1 - P(\mathcal{A})$  is occasionally useful for computing probabilities on its own. More commonly, you need other reasoning as well. The next problem illustrates an important feature of questions in probability: your intuition can be quite misleading. One problem is that the number of outcomes can be bigger or smaller than you expect.

**Worked example 4.21** *Shared birthdays*

What is the probability that, in a room of 30 people, there is a pair of people who have the same birthday?

**Solution:** We simplify, and assume that each year has 365 days, and that none of them are special (i.e. each day has the same probability of being chosen as a birthday). This model isn't perfect (there tend to be slightly more births roughly 9 months after: the start of spring; blackouts; major disasters; and so on) but it's workable. The easy way to attack this question is to notice that our probability,  $P(\{\text{shared birthday}\})$ , is

$$1 - P(\{\text{all birthdays different}\}).$$

This second probability is rather easy to compute. Each outcome in the sample space is a list of 30 days (one birthday per person). Each outcome has the same probability. So

$$P(\{\text{all birthdays different}\}) = \frac{\text{Number of outcomes in the event}}{\text{Total number of outcomes}}.$$

The total number of outcomes is easily seen to be  $365^{30}$ , which is the total number of possible lists of 30 days. The number of outcomes in the event is the number of lists of 30 days, all different. To count these, we notice that there are 365 choices for the first day; 364 for the second; and so on. So we have

$$P(\{\text{shared birthday}\}) = 1 - \frac{365 \times 364 \times \dots \times 336}{365^{30}} = 1 - 0.2937 = 0.7063$$

which means there's really a pretty good chance that two people in a room of 30 share a birthday.

If we change the birthday example slightly, the problem changes drastically. If you stand up in a room of 30 people and bet that two people in the room have the same birthday, you have a probability of winning of about 0.71. If you bet that there is someone else in the room who has the same birthday that you do, your probability of winning is very different.

**Worked example 4.22** *Shared birthdays*

You bet there is someone else in a room of 30 people who has the same birthday that you do. Assuming you know nothing about the other 29 people, what is the probability of winning?

**Solution:** The easy way to do this is

$$P(\{\text{winning}\}) = 1 - P(\{\text{losing}\}).$$

Now you will lose if everyone has a birthday different from you. You can think of the birthdays of the others in the room as a list of 29 days of the year. If your birthday is on the list, you win; if it's not, you lose. The number of losing lists is the number of lists of 29 days of the year such that your birthday is not in the list. This number is easy to get. We have 364 days of the year to choose from for each of 29 locations in the list. The total number of lists is the number of lists of 29 days of the year. Each list has the same probability. So

$$P(\{\text{losing}\}) = \frac{364^{29}}{365^{29}}$$

and

$$P(\{\text{winning}\}) \approx 0.0765.$$

There is a wide variety of problems like this; if you're so inclined, you can make a small but quite reliable profit off people's inability to estimate probabilities for this kind of problem correctly (examples 4.21 and 4.22 are reliably profitable; you could probably do quite well out of examples 4.16 and 4.17).

The rule  $P(\mathcal{A} - \mathcal{B}) = P(\mathcal{A}) - P(\mathcal{A} \cap \mathcal{B})$  is also occasionally useful for computing probabilities on its own; more commonly, you need other reasoning as well.

**Worked example 4.23** *Dice*

You flip two fair six-sided dice, and add the number of spots. What is the probability of getting a number divisible by 2, but not by 5?

**Solution:** There is an interesting way to work the problem. Write  $\mathcal{D}_n$  for the event the number is divisible by  $n$ . Now  $P(\mathcal{D}_2) = 1/2$  (count the cases; or, more elegantly, notice that each die has the same number of odd and even faces, and work from there). Now  $P(\mathcal{D}_2 - \mathcal{D}_5) = P(\mathcal{D}_2) - P(\mathcal{D}_2 \cap \mathcal{D}_5)$ . But  $\mathcal{D}_2 \cap \mathcal{D}_5$  contains only three outcomes (6, 4, 5, 5 and 4, 6), so  $P(\mathcal{D}_2 - \mathcal{D}_5) = 18/36 - 3/36 = 5/12$

Sometimes it is easier to reason about unions than to count outcomes directly.

**Worked example 4.24** *Two fair dice*

I roll two fair dice. What is the probability that the result is divisible by either 2 or 5, or both?

**Solution:** Write  $\mathcal{D}_n$  for the event the number is divisible by  $n$ . We want  $P(\mathcal{D}_2 \cup \mathcal{D}_5) = P(\mathcal{D}_2) + P(\mathcal{D}_5) - P(\mathcal{D}_2 \cap \mathcal{D}_5)$ . From example 4.23, we know  $P(\mathcal{D}_2) = 1/2$  and  $P(\mathcal{D}_2 \cap \mathcal{D}_5) = 3/36$ . By counting outcomes,  $P(\mathcal{D}_5) = 7/36$ . So  $P(\mathcal{D}_2 \cup \mathcal{D}_5) = (18 + 7 - 3)/36 = 22/36$ .

## 4.2 INDEPENDENCE

Some experimental results do not affect others. For example, if I flip a coin twice, whether I get heads on the first flip has no effect on whether I get heads on the second flip. As another example, I flip a coin; the outcome does not affect whether I get hit on the head by a falling apple later in the day. We refer to events with this property as **independent**.

Here is a pair of events that is not independent. Imagine I throw a die. Write  $\mathcal{A}$  for the event that the die comes up with an odd number of spots, and write  $\mathcal{B}$  for the event that the number of spots is either 3 or 5. Now these events are interrelated in an important way. If I *know* that  $\mathcal{A}$  has occurred, I also know that  $\mathcal{B}$  has occurred — I don't need to check separately, because  $\mathcal{A}$  implies  $\mathcal{B}$ .

Here is an example of a weaker interaction that results in events not being independent. Write  $\mathcal{C}$  for the event that the die comes up with an odd number of spots, and write  $\mathcal{D}$  for the event that the number of spots is larger than 3. These events are interrelated. The probability of each event separately is  $1/2$ . If I *know* that  $\mathcal{C}$  has occurred, then I know that the die shows either 1, 3, or 5 spots. One of these outcomes belongs to  $\mathcal{D}$ , and two do not. This means that knowing that  $\mathcal{C}$  has occurred tells you something about whether  $\mathcal{D}$  has occurred. Independent events do *not* have this property. This means that the probability that they occur together has an important property, given in the box below.

**Useful Facts: 4.4** *Independent events*

Two events  $\mathcal{A}$  and  $\mathcal{B}$  are **independent** if and only if

$$P(\mathcal{A} \cap \mathcal{B}) = P(\mathcal{A})P(\mathcal{B})$$

Indexing: Independent events

Indexing: Independent events

The “size” analogy helps motivate this expression. We think of  $P(\mathcal{A})$  as the “size” of  $\mathcal{A}$  relative to  $\Omega$ , and so on. Now  $P(\mathcal{A} \cap \mathcal{B})$  measures the “size” of  $\mathcal{A} \cap \mathcal{B}$

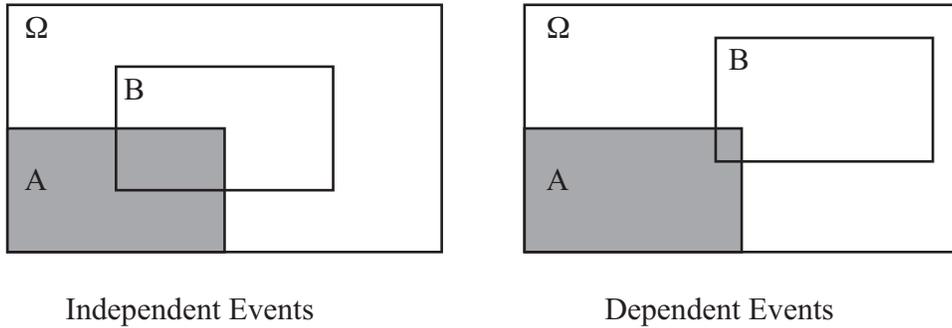


FIGURE 4.2: On the **left**,  $\mathcal{A}$  and  $\mathcal{B}$  are independent.  $\mathcal{A}$  spans  $1/4$  of  $\Omega$ , and  $\mathcal{A} \cap \mathcal{B}$  spans  $1/4$  of  $\mathcal{B}$ . This means that  $\mathcal{B}$  can't affect  $\mathcal{A}$ .  $1/4$  of the outcomes of  $\Omega$  lie in  $\mathcal{A}$ , and  $1/4$  of the outcomes in  $\mathcal{B}$  lie in  $\mathcal{A} \cap \mathcal{B}$ . On the **right**, they are not. Very few of the outcomes in  $\mathcal{B}$  lie in  $\mathcal{B} \cap \mathcal{A}$ , so that observing  $\mathcal{B}$  means that  $\mathcal{A}$  becomes less likely, because very few of the outcomes in  $\mathcal{B}$  also lie in  $\mathcal{A} \cap \mathcal{B}$ .

— that is, the part of  $\mathcal{A}$  that lies inside  $\mathcal{B}$ . But if  $\mathcal{A}$  and  $\mathcal{B}$  are independent, then the size of  $\mathcal{A} \cap \mathcal{B}$  relative to  $\mathcal{B}$  should be the same as the size of  $\mathcal{A}$  relative to  $\Omega$  (Figure 4.2). Otherwise,  $\mathcal{B}$  affects  $\mathcal{A}$ , because  $\mathcal{A}$  is more (or less) likely when  $\mathcal{B}$  has occurred.

So for  $\mathcal{A}$  and  $\mathcal{B}$  to be independent, we must have

$$\text{“Size” of } \mathcal{A} = \frac{\text{“Size” of piece of } \mathcal{A} \text{ in } \mathcal{B}}{\text{“Size” of } \mathcal{B}},$$

or, equivalently,

$$P(\mathcal{A}) = \frac{P(\mathcal{A} \cap \mathcal{B})}{P(\mathcal{B})}$$

which yields our expression.

**Worked example 4.25** *A fair die*

The space of outcomes for a fair six-sided die is

$$\{1, 2, 3, 4, 5, 6\}.$$

The die is fair means that each outcome has the same probability. Now we toss two fair dice. The outcome for each die is independent of that for the other. With what probability do we get two threes?

**Solution:**

$$\begin{aligned} P(\text{first die yields 3} \cap \text{second die yields 3}) &= P(\text{first die yields 3}) \times \\ &\quad P(\text{second die yields 3}) \\ &= (1/6)(1/6) \\ &= 1/36 \end{aligned}$$

**Worked example 4.26** *Find the Lady, twice*

Recall the setup of worked example ???. Assume that the card that is chosen is chosen fairly — that is, each card is chosen with the same probability. The game is played twice, and the cards are reshuffled between games. What is the probability of turning up a Queen and then a Queen again?

**Solution:** The events are independent, so  $1/9$ .

**Worked example 4.27** *Children*

A couple decides to have two children. Genders are assigned to children at random, fairly, at birth and independently at each birth (our models have to abstract a little!). What is the probability of having a boy and then a girl?

**Solution:**

$$\begin{aligned} P(\text{first is boy} \cap \text{second is girl}) &= P(\text{first is boy})P(\text{second is girl}) \\ &= (1/2)(1/2) \\ &= 1/4 \end{aligned}$$

You can use definition 4.4 (i.e.  $\mathcal{A}$  and  $\mathcal{B}$  are independent if and only if  $P(\mathcal{A} \cap \mathcal{B}) = P(\mathcal{A})P(\mathcal{B})$ ) to tell whether events are independent or not. Quite small changes to a problem affect whether events are independent, as in the worked example below.

**Worked example 4.28** *Cards and Independence*

We shuffle a standard deck of 52 cards and draw one card. The event  $\mathcal{A}$  is “the card is a red suit” and the event  $\mathcal{B}$  is “the card is a 10”. (1): Are  $\mathcal{A}$  and  $\mathcal{B}$  independent?

Now we take a standard deck of cards, and remove the ten of hearts. We shuffle this deck, and draw one card. The event  $\mathcal{C}$  is “the card drawn from the modified deck is a red suit” and the event  $\mathcal{D}$  is “the card drawn from the modified deck is a 10”. (2): Are  $\mathcal{C}$  and  $\mathcal{D}$  independent?

**Solution:** (1):  $P(\mathcal{A}) = 1/2$ ,  $P(\mathcal{B}) = 1/13$  and in example 4.14 we determined  $P(\mathcal{A} \cap \mathcal{B}) = 2/52$ . But  $2/52 = 1/26 = P(\mathcal{A})P(\mathcal{B})$ , so they are independent.  
 (2): These are not independent because  $P(\mathcal{C}) = 25/51$ ,  $P(\mathcal{D}) = 3/51$  and  $P(\mathcal{C} \cap \mathcal{D}) = 1/51 \neq P(\mathcal{C})P(\mathcal{D}) = 75/(51^2)$

The probability of a sequence of independent events can become very small very quickly, and this often misleads people.

**Worked example 4.29** *Accidental DNA Matches*

I search a DNA database with a sample. Each time I attempt to match this sample to an entry in the database, there is a probability of an accidental chance match of  $1e - 4$ . Chance matches are independent. There are 20, 000 people in the database. What is the probability I get at least one match, purely by chance?

**Solution:** This is  $1 - P(\text{no chance matches})$ . But  $P(\text{no chance matches})$  is much smaller than you think. We have

$$\begin{aligned}
 P(\text{no chance matches}) &= P\left(\begin{array}{l} \text{no chance match to record 1} \cap \\ \text{no chance match to record 2} \cap \\ \dots \cap \\ \text{no chance match to record 20, 000} \end{array}\right) \\
 &= P(\text{no chance match to a record})^{20,000} \\
 &= (1 - 1e - 4)^{20,000} \\
 &\approx 0.14
 \end{aligned}$$

so the probability is about 0.86 that you get at least one match by chance. If you're surprised, look at the exponent. Notice that if the database gets bigger, the probability grows; so at 40, 000 the probability of one match by chance is 0.98.

**Remember this:** *The probability of a set of independent events can become very small, and this often misleads the intuition.*

## 4.2.1 Example: Airline Overbooking

Counting outcomes in an event can require pretty elaborate combinatorial arguments. One form of argument that is particularly important is to reason about permutations. You should recall that the number of permutations of  $k$  items is  $k!$ .

**Worked example 4.30** *Counting outcomes with permutations*

I flip a coin  $N$  times. How many of the possible outcomes have exactly  $k$  heads?

**Solution:** Here is one natural way to think about this problem. Each outcome we are interested in is a string of  $k$  H's and  $N - k$  T's, and we need to count the number of such strings. We can do so with permutations. Write down any string of  $k$  H's and  $N - k$  T's. Any other string of  $k$  H's and  $N - k$  T's is a permutation of this one. However, many of these permutations simply swap one H with another, or one T with another. The total number of permutations of a string of  $N$  entries is  $N!$ . The number of permutations that swap H's with one another is  $k!$  (because there are  $k$  H's), and the number of permutations that swap T's with one another is  $(N - k)!$ . The total number of strings must then be

$$\frac{N!}{k!(N - k)!} = \binom{N}{k}$$

There is another way to think about example 4.30, which is more natural if you've seen combinations before. You start with a string of  $N$  T's, then you choose  $k$  distinct elements to turn into H's. The number of choices of  $k$  distinct elements in  $N$  items is:

$$\frac{N!}{k!(N - k)!} = \binom{N}{k}$$

so there must be this number of strings.

We can now quite easily study airline overbooking. Airlines generally sell more tickets for a flight than there are seats on the aircraft, because some passengers don't turn up on time, usually for random reasons. If the airline only sold one ticket per seat, their planes would likely have empty seats — which are lost profit — on each flight. If too many passengers turn up for a flight, someone will accept money to take the next flight. The amount of money that an airline has to give a passenger who doesn't get a seat is usually prescribed by regulators. Now an airline must compute with what probability they will have to pay out as a function of the number of tickets sold.

**Worked example 4.31** *Overbooking - 1*

An airline has a regular flight with six seats. It always sells seven tickets. Passengers turn up for the flight with probability  $p$ , and do so independent of other passengers. What is the probability that the flight is overbooked?

**Solution:** This is like a coin-flip problem; think of each passenger as a biased coin. With probability  $p$ , it comes up  $T$  (for *turn up*) and with probability  $(1 - p)$ , it turns up  $N$  (for *no-show*). This coin is flipped seven times, and we are interested in the probability that there are seven  $T$ 's. This is  $p^7$ , because the flips are independent.

**Worked example 4.32** *Overbooking - 2*

An airline has a regular flight with six seats. It always sells eight tickets. Passengers turn up for the flight with probability  $p$ , and do so independent of other passengers. What is the probability that six passengers arrive? (i.e. the flight is not overbooked or underbooked).

**Solution:** Now we flip the coin eight times, and are interested in the probability of getting exactly six  $T$ 's. We get, by the reasoning of worked example 4.30,

$$\frac{8!}{2!6!}p^6(1-p)^2 = 28p^6(1-p)^2$$

**Worked example 4.33** *Overbooking - 3*

An airline has a regular flight with six seats. It always sells eight tickets. Passengers turn up for the flight with probability  $p$ , and do so independent of other passengers. What is the probability that the flight is overbooked?

**Solution:** Now we flip the coin eight times, and are interested in the probability of getting more than six  $T$ 's. This is the union of two disjoint events (seven  $T$ 's and eight  $T$ 's). For the case of seven  $T$ 's, one flip must be  $N$ ; there are eight choices. For the case of eight  $T$ 's, all eight flips must be  $T$ , and there is only one way to achieve this. So the probability the flight is overbooked is

$$\begin{aligned} P(\text{overbooked}) &= P(7 T\text{'s} \cup 8 T\text{'s}) \\ &= P(7 T\text{'s}) + P(8 T\text{'s}) \\ &= 8p^7(1-p) + p^8 \end{aligned}$$

**Worked example 4.34** *Overbooking - 4*

An airline has a regular flight with  $s$  seats. It always sells  $t$  tickets. Passengers turn up for the flight with probability  $p$ , and do so independent of other passengers. What is the probability that  $u$  passengers turn up?

**Solution:** Now we flip the coin  $t$  times, and are interested in the probability of getting  $u$   $T$ 's. By the argument of worked example 4.30, there are

$$\frac{t!}{u!(t-u)!}$$

disjoint outcomes with  $u$   $T$ 's and  $t-u$   $N$ 's. Each such outcome is independent, and has probability  $p^u(1-p)^{t-u}$ . So

$$P(u \text{ passengers turn up}) = \frac{t!}{u!(t-u)!} p^u (1-p)^{t-u}$$

**Worked example 4.35** *Overbooking - 5*

An airline has a regular flight with  $s$  seats. It always sells  $t$  tickets. Passengers turn up for the flight with probability  $p$ , and do so independent of other passengers. What is the probability that the flight is oversold?

**Solution:** We need  $P(\{s + 1 \text{ turn up}\} \cup \{s + 2 \text{ turn up}\} \cup \dots \cup \{t \text{ turn up}\})$ . But the events  $\{i \text{ turn up}\}$  and  $\{j \text{ turn up}\}$  are disjoint if  $i \neq j$ . So we can exploit example 4.34, and write

$$\begin{aligned} P(\text{overbooked}) &= P(\{s + 1 \text{ turn up}\}) + P(\{s + 2 \text{ turn up}\}) + \\ &\quad \dots P(\{t \text{ turn up}\}) \\ &= \sum_{i=s+1}^t P(\{i \text{ turn up}\}) \\ &= \sum_{i=s+1}^t \frac{t!}{i!(t-i)!} p^i (1-p)^{t-i} \end{aligned}$$

## 4.3 CONDITIONAL PROBABILITY

As we saw in section 4.2, observing one event may change the probability another event will occur. Conditional probability is a mechanism to keep track of this effect.

There were several examples of events that weren't independent in that section (and you should look at these again). Here is another example. If you throw a fair die twice and add the numbers, then the probability of getting a number less than six is  $\frac{10}{36}$ . Now imagine you *know* that the first die came up three. In this case, the probability that the sum will be less than six is  $\frac{1}{3}$ , which is slightly larger than  $\frac{10}{36}$ . If the first die came up four, then the probability the sum will be less than six is  $\frac{1}{6}$ , which is rather less than  $\frac{10}{36}$ . If the first die came up one, then the probability that the sum is less than six becomes  $\frac{2}{3}$ , which is much larger than  $\frac{10}{36}$ .

Each of these probabilities is an example of a **conditional probability**. We assume we have a space of outcomes and a collection of events. The conditional probability of  $\mathcal{B}$ , conditioned on  $\mathcal{A}$ , is the probability that  $\mathcal{B}$  occurs given that  $\mathcal{A}$  has definitely occurred. We write this as

$$P(\mathcal{B}|\mathcal{A})$$

From the examples, it should be clear to you that for some cases  $P(\mathcal{B}|\mathcal{A})$  is the same as  $P(\mathcal{B})$ , and for other cases it is not.

## 4.3.1 Evaluating Conditional Probabilities

To get an expression for  $P(\mathcal{B}|\mathcal{A})$ , notice that, because  $\mathcal{A}$  is known to have occurred, our space of outcomes or sample space is now reduced to  $\mathcal{A}$ . We know that our outcome lies in  $\mathcal{A}$ ;  $P(\mathcal{B}|\mathcal{A})$  is the probability that it also lies in  $\mathcal{B} \cap \mathcal{A}$ .

The outcome lies in  $\mathcal{A}$ , and so it must lie in either  $P(\mathcal{B} \cap \mathcal{A})$  or in  $P(\mathcal{B}^c \cap \mathcal{A})$ . This means that

$$P(\mathcal{B}|\mathcal{A}) + P(\mathcal{B}^c|\mathcal{A}) = 1.$$

Now recall the idea of probabilities as relative frequencies. If  $P(\mathcal{C} \cap \mathcal{A}) = kP(\mathcal{B} \cap \mathcal{A})$ , this means that outcomes in  $\mathcal{C} \cap \mathcal{A}$  will appear  $k$  times as often as outcomes in  $\mathcal{B} \cap \mathcal{A}$ . But this must apply even if we know in advance that the outcome is in  $\mathcal{A}$ . This means that, if  $P(\mathcal{C} \cap \mathcal{A}) = kP(\mathcal{B} \cap \mathcal{A})$ , then  $P(\mathcal{C}|\mathcal{A}) = kP(\mathcal{B}|\mathcal{A})$ . In turn, we must have

$$P(\mathcal{B}|\mathcal{A}) \propto P(\mathcal{B} \cap \mathcal{A}).$$

Now we need to determine the constant of proportionality; write  $c$  for this constant, meaning

$$P(\mathcal{B}|\mathcal{A}) = cP(\mathcal{B} \cap \mathcal{A}).$$

We have that

$$P(\mathcal{B}|\mathcal{A}) + P(\mathcal{B}^c|\mathcal{A}) = cP(\mathcal{B} \cap \mathcal{A}) + cP(\mathcal{B}^c \cap \mathcal{A}) = cP(\mathcal{A}) = 1,$$

so that

$$P(\mathcal{B}|\mathcal{A}) = \frac{P(\mathcal{B} \cap \mathcal{A})}{P(\mathcal{A})}.$$

I find the “size” metaphor helpful here. We have that  $P(\mathcal{B}|\mathcal{A})$  measures the probability that an outcome is in  $\mathcal{B}$ , given we *know* it is in  $\mathcal{A}$ . From the “size” perspective,  $P(\mathcal{B}|\mathcal{A})$  measures the size of  $(\mathcal{A} \cap \mathcal{B})$  *relative to*  $\mathcal{A}$ . So our expression makes sense, because the fraction of the event  $\mathcal{A}$  that is also part of the event  $\mathcal{B}$  is given by the size of the intersection divided by the size of  $\mathcal{A}$ .

Another, very useful, way to write the expression  $P(\mathcal{B}|\mathcal{A}) = \frac{P(\mathcal{B} \cap \mathcal{A})}{P(\mathcal{A})}$  is:

$$P(\mathcal{B}|\mathcal{A})P(\mathcal{A}) = P(\mathcal{B} \cap \mathcal{A}).$$

Now, since  $\mathcal{B} \cap \mathcal{A} = \mathcal{A} \cap \mathcal{B}$ , we must have that

$$P(\mathcal{B}|\mathcal{A}) = \frac{P(\mathcal{A}|\mathcal{B})P(\mathcal{B})}{P(\mathcal{A})}$$

**Worked example 4.36** *Car factories*

There are two car factories,  $A$  and  $B$ . Each year, factory  $A$  produces 1000 cars, of which 10 are lemons. Factory  $B$  produces 2 cars, each of which is a lemon. All cars go to a single lot, where they are thoroughly mixed up. I buy a car.

- What is the probability it is a lemon?
- What is the probability it came from factory  $B$ ?
- The car is now revealed to be a lemon. What is the probability it came from factory  $B$ , conditioned on the fact it is a lemon?

**Solution:**

- Write the event the car is a lemon as  $\mathcal{L}$ . There are 1002 cars, of which 12 are lemons. The probability that I select any given car is the same, so we have  $P(\mathcal{L}) = 12/1002$ .
- Same argument yields  $P(\mathcal{B}) = 2/1002$ .
- Write  $\mathcal{B}$  for the event the car comes from factory  $B$ . I need  $P(\mathcal{B}|\mathcal{L}) = \frac{P(\mathcal{L} \cap \mathcal{B})}{P(\mathcal{L})} = P(\mathcal{L}|\mathcal{B})P(\mathcal{B})/P(\mathcal{L})$ . I have  $P(\mathcal{L}|\mathcal{B})P(\mathcal{B})/P(\mathcal{L}) = (1 \times 2/1002)/(12/1002) = 1/6$ .

**Worked example 4.37** *Royal flushes in poker - 1*

*This exercise is after Stirzaker, p. 51.*

You are playing a straightforward version of poker, where you are dealt five cards face down. A royal flush is a hand of AKQJ10 all in one suit. What is the probability that you are dealt a royal flush?

**Solution:** This is

$$\frac{\text{number of hands that are royal flushes, ignoring card order}}{\text{total number of different five card hands, ignoring card order}}$$

There are four hands that are royal flushes (one for each suit). Now the total number of five card hands is

$$\binom{52}{5} = 2598960$$

so we have

$$\frac{4}{2598960} = \frac{1}{649740}.$$

**Worked example 4.38** *Royal flushes in poker - 2*

*This exercise is after Stirzaker, p. 51.*

You are playing a straightforward version of poker, where you are dealt five cards face down. A royal flush is a hand of AKQJ10 all in one suit. The fifth card that you are dealt lands face up. It is the nine of spades. What now is the probability that you have been dealt a royal flush? (i.e. what is the conditional probability of getting a royal flush, conditioned on the event that one card is the nine of spades)

**Solution:** No hand containing a nine of spades is a royal flush, so this is easily zero.

**Worked example 4.39** *Royal flushes in poker - 3*

*This exercise is after Stirzaker, p. 51.*

You are playing a straightforward version of poker, where you are dealt five cards face down. A royal flush is a hand of AKQJ10 all in one suit. The fifth card that you are dealt lands face up. It is the Ace of spades. What now is the probability that you have been dealt a royal flush? (i.e. what is the conditional probability of getting a royal flush, conditioned on the event that one card is the Ace of spades)

**Solution:** Now consider the events

$\mathcal{A}$  = you get a royal flush *and* the last card is the ace of spades

and

$\mathcal{B}$  = the last card you get is the ace of spades,

and the expression

$$P(\mathcal{A}|\mathcal{B}) = \frac{P(\mathcal{A} \cap \mathcal{B})}{P(\mathcal{B})}.$$

Now  $P(\mathcal{B}) = \frac{1}{52}$ .  $P(\mathcal{A} \cap \mathcal{B})$  is given by

$$\frac{\text{number of five card royal flushes where card five is Ace of spades}}{\text{total number of different five card hands}}.$$

where we DO NOT ignore card order. This is

$$\frac{4 \times 3 \times 2 \times 1}{52 \times 51 \times 50 \times 49 \times 48}$$

yielding

$$P(\mathcal{A}|\mathcal{B}) = \frac{1}{249900}.$$

Notice the interesting part: the conditional probability is rather larger than the probability. If you see this ace, the conditional probability is  $\frac{13}{5}$  times the probability that you will get a flush if you don't. Seeing this card has really made a difference.

**Worked example 4.40** *Two dice*

We throw two fair dice. What is the conditional probability that the sum of spots on both dice is greater than six, conditioned on the event that the first die comes up five?

**Solution:** Write the event that the first die comes up 5 as  $\mathcal{F}$ , and the event the sum is greater than six as  $\mathcal{S}$ . There are five outcomes where the first die comes up 5 and the number is greater than 6, so  $P(\mathcal{F} \cap \mathcal{S}) = 5/36$ . Now

$$P(\mathcal{S}|\mathcal{F}) = P(\mathcal{F} \cap \mathcal{S})/P(\mathcal{F}) = (5/36)/(1/6) = 5/6.$$

Notice that  $\mathcal{A} \cap \mathcal{B}$  and  $\mathcal{A} \cap \mathcal{B}^c$  are disjoint sets, and that  $\mathcal{A} = (\mathcal{A} \cap \mathcal{B}) \cup (\mathcal{A} \cap \mathcal{B}^c)$ . So, because  $P(\mathcal{A}) = P(\mathcal{A} \cap \mathcal{B}) + P(\mathcal{A} \cap \mathcal{B}^c)$ , we have

$$P(\mathcal{A}) = P(\mathcal{A}|\mathcal{B})P(\mathcal{B}) + P(\mathcal{A}|\mathcal{B}^c)P(\mathcal{B}^c)$$

a tremendously important and useful fact. Another version of this fact is also very useful. Assume we have a collection of disjoint sets  $\mathcal{B}_i$ . These sets must have the property that (a)  $\mathcal{B}_i \cap \mathcal{B}_j = \emptyset$  for  $i \neq j$  and (b) they cover  $\mathcal{A}$ , meaning that  $\mathcal{A} \cap (\cup_i \mathcal{B}_i) = \mathcal{A}$ . Then, because  $P(\mathcal{A}) = \sum_i P(\mathcal{A} \cap \mathcal{B}_i)$ , so we have

$$P(\mathcal{A}) = \sum_i P(\mathcal{A}|\mathcal{B}_i)P(\mathcal{B}_i)$$

It is wise to be suspicious of your intuitions when thinking about problems in conditional probability. There is a really big difference between  $P(\mathcal{A}|\mathcal{B})$  and  $P(\mathcal{B}|\mathcal{A})P(\mathcal{A})$ . Not respecting this difference can lead to serious problems (section ??), and seems to be easy to do. The division sign in the expression  $P(\mathcal{A}|\mathcal{B}) = P(\mathcal{B}|\mathcal{A})P(\mathcal{A})/P(\mathcal{B})$  can have alarming effects; as a result, most people have quite poor intuitions about conditional probability. Here is one helpful example. If you buy a lottery ticket ( $\mathcal{L}$ ), the probability of winning ( $\mathcal{W}$ ) is small. So  $P(\mathcal{W}|\mathcal{L})$  may be very small. But  $P(\mathcal{L}|\mathcal{W})$  is 1 — the winner is always someone who bought a ticket.

**Useful Facts: 4.5** *Conditional probability formulas*

You should remember the following formulas:

- $P(\mathcal{B}|\mathcal{A}) = \frac{P(\mathcal{A}|\mathcal{B})P(\mathcal{B})}{P(\mathcal{A})}$
- $P(\mathcal{A}) = P(\mathcal{A}|\mathcal{B})P(\mathcal{B}) + P(\mathcal{A}|\mathcal{B}^c)P(\mathcal{B}^c)$
- Assume (a)  $\mathcal{B}_i \cap \mathcal{B}_j = \emptyset$  for  $i \neq j$  and (b)  $\mathcal{A} \cap (\cup_i \mathcal{B}_i) = \mathcal{A}$ ; then  $P(\mathcal{A}) = \sum_i P(\mathcal{A}|\mathcal{B}_i)P(\mathcal{B}_i)$

Indexing: Conditional probability formulas  
 Indexing: Conditional probability formulas

**Worked example 4.41** *False positives*

After Stirzaker, p55. You have a blood test for a rare disease that occurs by chance in 1 person in 100,000. If you have the disease, the test will report that you do with probability 0.95 (and that you do not with probability 0.05). If you do not have the disease, the test will report a false positive with probability  $1e-3$ . If the test says you do have the disease, what is the probability it that you actually have the disease?

**Solution:** Write  $S$  for the event you are sick and  $R$  for the event the test reports you are sick. We need  $P(S|R)$ . We have

$$\begin{aligned}
 P(S|R) &= \frac{P(R|S)P(S)}{P(R)} \\
 &= \frac{P(R|S)P(S)}{P(R|S)P(S) + P(R|S^c)P(S^c)} \\
 &= \frac{0.95 \times 1e-5}{0.95 \times 1e-5 + 1e-3 \times (1 - 1e-5)} \\
 &= 0.0094
 \end{aligned}$$

which should strike you as being a bit alarming. Notice what is happening here. There are two ways that the test could come back positive: either you have the disease, or the test is producing a false positive. But the disease is so rare that it's much more likely you have a false positive result than you have the disease.

**Worked example 4.42** *False positives -2*

After Stirzaker, p55. You want to *design* a blood test for a rare disease that occurs by chance in 1 person in 100,000. If you have the disease, the test will report that you do with probability  $p$  (and that you do not with probability  $(1 - p)$ ). If you do not have the disease, the test will report a false positive with probability  $q$ . You want to choose the value of  $p$  so that if the test says you have the disease, there is at least a 50% probability that you do.

**Solution:** Write  $S$  for the event you are sick and  $R$  for the event the test reports you are sick. We need  $P(S|R)$ . We have

$$\begin{aligned} P(S|R) &= \frac{P(R|S)P(S)}{P(R)} \\ &= \frac{P(R|S)P(S)}{P(R|S)P(S) + P(R|S^c)P(S^c)} \\ &= \frac{p \times 1e-5}{p \times 1e-5 + q \times (1 - 1e-5)} \\ &\geq 0.5 \end{aligned}$$

which means that  $p \geq 99999q$  which should strike you as being very alarming indeed, because  $p \leq 1$  and  $q \geq 0$ . One plausible pair of values is  $q = 1e-5$ ,  $p = 1 - 1e-5$ . The test has to be spectacularly accurate to be of any use.

## 4.3.2 The Prosecutors Fallacy

It is quite easy to make mistakes in conditional probability. Several such mistakes have names, because they're so common. It's worth looking at mistakes as an example of how one's intuition can mislead. One important named mistake is the **prosecutor's fallacy**. A prosecutor has evidence  $\mathcal{E}$  against a suspect. Write  $\mathcal{I}$  for the event that the suspect is innocent. We know that  $P(\mathcal{E}|\mathcal{I})$  is extremely small. The prosecutor argues that the suspect must be guilty, because  $P(\mathcal{E}|\mathcal{I})$  is so small. This argument is misleading, because the conditional probability of interest is  $P(\mathcal{I}|\mathcal{E})$ . It does not matter how unlikely the evidence is if you are innocent; what matters is the probability you are innocent, given the evidence.

The distinction is very important, because  $P(\mathcal{I}|\mathcal{E})$  could be big even if  $P(\mathcal{E}|\mathcal{I})$  is small. In the expression

$$P(\mathcal{I}|\mathcal{E}) = \frac{P(\mathcal{E}|\mathcal{I})P(\mathcal{I})}{P(\mathcal{E})} = \frac{P(\mathcal{E}|\mathcal{I})P(\mathcal{I})}{(P(\mathcal{E}|\mathcal{I})P(\mathcal{I}) + P(\mathcal{E}|\mathcal{I}^c)(1 - P(\mathcal{I})))}$$

notice that if  $P(\mathcal{I})$  is large or if  $P(\mathcal{E}|\mathcal{I}^c)$  is much smaller than  $P(\mathcal{E}|\mathcal{I})$ , then  $P(\mathcal{I}|\mathcal{E})$  could be close to one even if  $P(\mathcal{E}|\mathcal{I})$  is small. The question to look at is not how unlikely the evidence is if the subject is innocent; instead, the question is how likely the subject is to be guilty compared to some other source of the evidence. These are two very different questions.

The prosecutor's fallacy has contributed to a variety of miscarriages of justice. One famous incident occurred in the UK, involving a mother, Sally Clark, who was convicted of murdering two of her children. Expert evidence by paediatrician Roy Meadow argued that the probability of both deaths resulting from Sudden Infant Death Syndrome was extremely small. Her first appeal cited, among other grounds, statistical error in the evidence (you should spot the prosecutors fallacy; others were involved, too). The appeals court rejected this appeal, calling the statistical point "a sideshow". This prompted a great deal of controversy, both in the public press and various professional journals, including a letter from the then president of the Royal Statistical Society to the Lord Chancellor, pointing out that "statistical evidence . . . (should be) . . . presented only by appropriately qualified statistical experts". A second appeal (on other grounds) followed, and was successful. The appellate judges specifically criticized the statistical evidence, although it was not a point of appeal. Clark never recovered from this horrific set of events and died in tragic circumstances shortly after the second appeal. Roy Meadow was then struck off the rolls for serious professional misconduct as an expert witness, a ruling he appealed successfully. You can find a more detailed account of this case, with pointers to important documents including the letter to the Lord Chancellor (which is well worth reading), at [http://en.wikipedia.org/wiki/Roy\\_Meadow](http://en.wikipedia.org/wiki/Roy_Meadow); there is further material on the prosecutors fallacy at [http://en.wikipedia.org/wiki/Prosecutor%27s\\_fallacy](http://en.wikipedia.org/wiki/Prosecutor%27s_fallacy).

**Remember this:** *Conditional probability misleads intuition so regularly that some errors have names. Be very careful.*

#### 4.3.3 Independence and Conditional Probability

As we have seen, two events are **independent** if

$$P(\mathcal{A} \cap \mathcal{B}) = P(\mathcal{A})P(\mathcal{B}).$$

If two events  $\mathcal{A}$  and  $\mathcal{B}$  are independent, then

$$P(\mathcal{A}|\mathcal{B}) = P(\mathcal{A})$$

and

$$P(\mathcal{B}|\mathcal{A}) = P(\mathcal{B}).$$

Again, this means that knowing that  $\mathcal{A}$  occurred tells you nothing about  $\mathcal{B}$  — the probability that  $\mathcal{B}$  will occur is the same whether you know that  $\mathcal{A}$  occurred or not.

Events  $\mathcal{A}_1 \dots \mathcal{A}_n$  are **pairwise independent** if each pair is independent (i.e.  $\mathcal{A}_1$  and  $\mathcal{A}_2$  are independent, etc.). They are **independent** if for any collection of distinct indices  $i_1 \dots i_k$  we have

$$P(\mathcal{A}_{i_1} \cap \dots \cap \mathcal{A}_{i_k}) = P(\mathcal{A}_{i_1}) \dots P(\mathcal{A}_{i_k})$$

Notice that independence is a much stronger assumption than pairwise independence.

**Worked example 4.43** *Cards and pairwise independence*

We draw three cards from a properly shuffled standard deck, with replacement and reshuffling (i.e., draw a card, make a note, return to deck, shuffle, draw the next, make a note, shuffle, draw the third). Let  $\mathcal{A}$  be the event that “card 1 and card 2 have the same suit”; let  $\mathcal{B}$  be the event that “card 2 and card 3 have the same suit”; let  $\mathcal{C}$  be the event that “card 1 and card 3 have the same suit”. Show these events are pairwise independent, but not independent.

**Solution:** By counting, you can check that  $P(\mathcal{A}) = 1/4$ ;  $P(\mathcal{B}) = 1/4$ ; and  $P(\mathcal{A} \cap \mathcal{B}) = 1/16$ , so that these two are independent. This argument works for other pairs, too. But  $P(\mathcal{C} \cap \mathcal{A} \cap \mathcal{B}) = 1/16$  which is not  $1/4^3$ , so the events are not independent; this is because the third event is logically implied by the first two.

Events  $\mathcal{A}_1 \dots \mathcal{A}_k$  are **conditionally independent** conditioned on event  $\mathcal{B}$  if

$$P(\mathcal{A}_{i_1} \cap \dots \cap \mathcal{A}_{i_k} | \mathcal{B}) = P(\mathcal{A}_{i_1} | \mathcal{B}) \dots P(\mathcal{A}_{i_k} | \mathcal{B})$$

**Worked example 4.44** *Cards and conditional independence*

We remove a red 10 and a red 6 from a standard deck of playing cards. We shuffle the remaining cards, and draw one card. Write  $\mathcal{A}$  for the event that the card drawn is a 10,  $\mathcal{B}$  for the event the card drawn is red, and  $\mathcal{C}$  for the event that the card drawn is either a 10 or a 6. Show that  $\mathcal{A}$  and  $\mathcal{B}$  are not independent, but are conditionally independent conditioned on  $\mathcal{C}$ .

**Solution:** We have  $P(\mathcal{A}) = 3/50$ ,  $P(\mathcal{B}) = 24/50$ ,  $P(\mathcal{A} \cap \mathcal{B}) = 1/50$ , so

$$P(\mathcal{A} | \mathcal{B}) = \frac{1/50}{24/50} = \frac{1}{24} \neq P(\mathcal{A})$$

so  $\mathcal{A}$  and  $\mathcal{B}$  are not independent. We have also that  $P(\mathcal{A} | \mathcal{C}) = 1/2$  and  $P(\mathcal{B} | \mathcal{C}) = 2/6 = 1/3$ . Now

$$P(\mathcal{A} \cap \mathcal{B} | \mathcal{C}) = 1/6 = P(\mathcal{A} | \mathcal{C})P(\mathcal{B} | \mathcal{C})$$

so  $\mathcal{A}$  and  $\mathcal{B}$  are conditionally independent conditioned on  $\mathcal{C}$ .

**Definition: 4.2** *Pairwise independent, independent and conditionally independent events*

Events  $\mathcal{A}_1 \dots \mathcal{A}_n$  are **pairwise independent** if each pair is independent (i.e.  $\mathcal{A}_1$  and  $\mathcal{A}_2$  are independent, etc.). They are **independent** if for any collection of distinct indices  $i_1 \dots i_k$  we have

$$P(\mathcal{A}_{i_1} \cap \dots \cap \mathcal{A}_{i_k}) = P(\mathcal{A}_{i_1}) \dots P(\mathcal{A}_{i_k})$$

They are **conditionally independent** conditioned on event  $\mathcal{B}$  if

$$P(\mathcal{A}_{i_1} \cap \dots \cap \mathcal{A}_{i_k} | \mathcal{B}) = P(\mathcal{A}_{i_1} | \mathcal{B}) \dots P(\mathcal{A}_{i_k} | \mathcal{B})$$

Indexing: Pairwise independent, independent and conditionally independent events

Indexing: Pairwise independent, independent and conditionally independent events

We usually do not have the information required to prove that events are independent. Instead, we use intuition (for example, two flips of the same coin are likely to be independent unless there is something very funny going on) or simply choose to apply models in which some variables are independent.

Independent events can lead very quickly to very small probabilities, as we saw in example 4.29. This can mislead intuition quite badly, and lead to serious problems. In particular, these small probabilities can interact with the prosecutor's fallacy in a dangerous way. In example 4.29, we saw how the probability of getting a chance match in a large DNA database could be quite big, even though the probability of a single match is small. One version of the prosecutor's fallacy is to argue that, because the probability of a single match is small, the person who matched the DNA must have committed the crime. The fallacy is to ignore the fact that the probability of a chance match to a large database is quite high.

People quite often reason poorly about independent events. The most common problem is known as the **gambler's fallacy**. This occurs when you reason that the probability of an independent event has been changed by previous outcomes. For example, imagine I toss a coin that is known to be fair 20 times and get 20 heads. The probability that the next toss will result in a head has not changed at all — it is still 0.5 — but many people will believe that it has changed. This idea is also sometimes referred to as **antichance**.

#### 4.4 EXAMPLE: THE MONTY HALL PROBLEM

Careless thinking about probability, particularly conditional probability, can cause wonderful confusion. The Monty Hall problem is a relatively simple exercise in conditional probability. Nonetheless, it has been the subject of extensive, lively, and often quite inaccurate correspondence in various national periodicals — it seems to catch the attention, which is why we describe it in some detail. The problem works like this: There are three doors. Behind one is a car. Behind each of the others is

a goat. The car and goats are placed randomly and fairly, so that the probability that there is a car behind each door is the same. You will get the object that lies behind the door you choose at the end of the game. The goats are interchangeable, and, for reasons of your own, you would prefer the car to a goat.

The game goes as follows. You select a door. The host then opens a door and shows you a goat. You must now choose to either keep your door, or switch to the other door. What should you do?

You cannot tell what to do *using the information provided*, by the following argument. Label the door you chose at the start of the game 1; the other doors 2 and 3. Write  $C_i$  for the event that the car lies behind door  $i$ . Write  $G_m$  for the event that a goat is revealed behind door  $m$ , where  $m$  is the number of the door where the goat was revealed (which could be 1, 2, or 3). You need to know  $P(C_1|G_m)$ . But

$$P(C_1|G_m) = \frac{P(G_m|C_1)P(C_1)}{P(G_m|C_1)P(C_1) + P(G_m|C_2)P(C_2) + P(G_m|C_3)P(C_3)}$$

and you do not know  $P(G_m|C_1)$ ,  $P(G_m|C_2)$ ,  $P(G_m|C_3)$ , because you don't know the rule by which the host chooses which door to open to reveal a goat. Different rules lead to quite different analyses.

There are several possible rules for the host to show a goat:

- **Rule 1:** choose a door uniformly at random.
- **Rule 2:** choose from the doors with goats behind them *that are not door 1* uniformly and at random.
- **Rule 3:** if the car is at 1, then choose 2; if at 2, choose 3; if at 3, choose 1.
- **Rule 4:** choose from the doors with goats behind them uniformly and at random.

We should keep track of the rules in the conditioning, so we write  $P(G_m|C_1, r_1)$  for the conditional probability that a goat was revealed behind door  $m$  when the car is behind door 1, using rule 1 (and so on). This means we are interested in

$$P(C_1|G_m, r_n) = \frac{P(G_m|C_1, r_n)P(C_1)}{P(G_m|C_1, r_n)P(C_1) + P(G_m|C_2, r_n)P(C_2) + P(G_m|C_3, r_n)P(C_3)}.$$

Notice that each of these rules is consistent with your observations — what you saw could have occurred under any of these rules. You have to know which rule the host uses to proceed. You should be aware that in many of the discussions of this problem, people assume without comment that the host uses rule 2, then proceed with this assumption.

**Worked example 4.45** *Monty Hall, rule one*

Assume the host uses rule one, and shows you a goat behind door two. What is  $P(C_1|G_2, r_1)$ ?

**Solution:** To work this out, we need to know  $P(G_2|C_1, r_1)$ ,  $P(G_2|C_2, r_1)$  and  $P(G_2|C_3, r_1)$ . Now  $P(G_2|C_2, r_1)$  must be zero, because the host could not reveal a goat behind door two if there was a car behind that door. Write  $O_2$  for the event the host *chooses* to open door two, and  $B_2$  for the event there happens to be a goat behind door two. These two events are independent — the host chose the door uniformly at random. We can compute

$$\begin{aligned} P(G_2|C_1, r_1) &= P(O_2 \cap B_2|C_1, r_1) \\ &= P(O_2|C_1, r_1)P(B_2|C_1, r_1) \\ &= (1/3)(1) \\ &= 1/3 \end{aligned}$$

where  $P(B_2|C_1, r_1) = 1$  because we conditioned on the fact there was a car behind door one, so there is a goat behind each other door. This argument establishes  $P(G_2|C_3, r_1) = 1/3$ , too. So  $P(C_1|G_2, r_1) = 1/2$  — the host showing you the goat does not motivate you to do anything, because if  $P(C_1|G_2, r_1) = 1/2$ , then  $P(C_3|G_2, r_1) = 1/2$ , too — there's nothing to choose between the two closed doors.

**Worked example 4.46** *Monty Hall, rule two*

Assume the host uses rule two, and shows you a goat behind door two. What is  $P(C_1|G_2, r_2)$ ?

**Solution:** To work this out, we need to know  $P(G_2|C_1, r_2)$ ,  $P(G_2|C_2, r_2)$  and  $P(G_2|C_3, r_2)$ . Now  $P(G_2|C_2, r_2) = 0$ , because the host chooses from doors with goats behind them.  $P(G_2|C_1, r_2) = 1/2$ , because the host chooses uniformly and at random from doors with goats behind them that are not door one; if the car is behind door one, there are two such doors.  $P(G_2|C_3, r_2) = 1$ , because there is only one door that (a) has a goat behind it and (b) isn't door one. Plug these numbers into the formula, to get  $P(C_1|G_2, r_2) = 1/3$ . This is the source of all the fuss. It says that, if you know the host is using rule two, you should switch doors if the host shows you a goat behind door two (because  $P(C_3|G_2, r_2) = 2/3$ ).

Notice what is happening: if the car is behind door three, then the *only* choice of goat for the host is the goat behind two. So by choosing a door under rule two,

the host is signalling some information to you, which you can use. By using rule three, the host can tell you precisely where the car is (exercises).

Many people find the result of example 4.46 counterintuitive, and object (sometimes loudly, in newspaper columns, letters to the editor, etc.). One example that some people find helpful is an extreme case. Imagine that, instead of three doors, there are 1002. The host is using rule two, modified in the following way: open all but one of the doors that are not door one, choosing only doors that have goats behind them to open. You choose door one; the host opens 1000 doors — say, all but doors one and 1002. What would you do?

4.5 YOU SHOULD

4.5.1 be able to:

- Write out a set of outcomes for an experiment.
- Construct an event space.
- Compute the probabilities of outcomes and events.
- Determine when events are independent.
- Compute the probabilities of outcomes by counting events, when the count is straightforward.
- Compute a conditional probability.

4.5.2 remember:

Definition: Sample space . . . . .	78
Useful facts: Sample space . . . . .	79
New term: independent . . . . .	80
Probability of outcomes is frequency of outcomes. . . . .	81
New term: event . . . . .	82
Useful facts: Event . . . . .	83
Useful facts: Probability of events . . . . .	85
Compute the probability of events by counting outcomes. . . . .	91
New term: independent . . . . .	94
New term: independent . . . . .	94
Useful facts: Independent events . . . . .	94
Warning: independence can mislead . . . . .	98
Useful facts: Conditional probability formulas . . . . .	108
New term: prosecutor’s fallacy . . . . .	109
Be careful with conditional probability. . . . .	110
Definition: Pairwise independent, independent and conditionally independent events	112
New term: gambler’s fallacy . . . . .	112

## PROBLEMS

## Outcomes

- 4.1. You roll a four sided die. What is the space of outcomes?
- 4.2. King Lear decides to allocate three provinces (1, 2, and 3) to his daughters (Goneril, Regan and Cordelia - read the book) at random. Each gets one province. What is the space of outcomes?
- 4.3. You randomly wave a flyswatter at a fly. What is the space of outcomes?
- 4.4. You read the book, so you know that King Lear had family problems. As a result, he decides to allocate two provinces to one daughter, one province to another daughter, and no provinces to the third. Because he's a bad problem solver, he does so at random. What is the space of outcomes?

## The Probability of an Outcome

- 4.5. You roll a fair four sided die. What is the probability of getting a 3?
- 4.6. You shuffle a standard deck of playing cards and draw a card. What is the probability that this is the king of hearts?
- 4.7. A roulette wheel has 36 slots numbered 1-36, and two slots numbered zero. The croupier spins the wheel, and throws a ball onto the surface; the ball bounces around and ends up in a slot (which is chosen fairly and at random). What is the probability the ball ends up in slot 2?

## Events

- 4.8. At a particular University,  $1/2$  of the students drink alcohol and  $1/3$  of the students smoke cigarettes.
  - (a) What is the largest possible fraction of students who do neither?
  - (b) It turns out that, in fact,  $1/3$  of the students do neither. What fraction of the students does both?
- 4.9. I flip two coins. What one set needs to be added to this collection of sets to form an event space?

$$\Sigma = \{\emptyset, \Omega, \{TH\}, \{HT, TH, TT\}, \{HH\}, \{HT, TT\}, \{HH, TH\}\}$$

## The Probability of Events

- 4.10. You flip a fair coin three times. What is the probability of seeing HTH? (i.e. Heads, then Tails, then Heads)
- 4.11. You flip a fair coin three times. What is the probability of seeing two heads and one tail?
- 4.12. You remove the king of hearts from a standard deck of cards, then shuffle it and draw a card.
  - (a) What is the probability this card is a king?
  - (b) What is the probability this card is a heart?
- 4.13. You shuffle a standard deck of cards, then draw four cards.
  - (a) What is the probability all four are the same suit?
  - (b) What is the probability all four are red?
  - (c) What is the probability each has a different suit?
- 4.14. You roll three fair six-sided dice and add the numbers. What is the probability the result is even?

- 4.15. You roll three fair six-sided dice and add the numbers. What is the probability the result is even *and* not divisible by 20?
- 4.16. You shuffle a standard deck of cards, then draw seven cards. What is the probability that you see no aces?
- 4.17. Show that  $P(\mathcal{A} - (\mathcal{B} \cup \mathcal{C})) = P(\mathcal{A}) - P(\mathcal{A} \cap \mathcal{B}) - P(\mathcal{A} \cap \mathcal{C}) + P(\mathcal{A} \cap \mathcal{B} \cap \mathcal{C})$ .
- 4.18. You draw a single card from a standard 52 card deck. What is the probability that it is red?
- 4.19. You remove all heart cards from a standard 52 card deck, then draw a single card from the result. What is the probability that the card you draw is red?

### Computing Probabilities by Counting Outcomes

- 4.20. Assume each outcome in  $\Omega$  has the same probability. In this case, show

$$P(\mathcal{E}) = \frac{\text{Number of outcomes in } \mathcal{E}}{\text{Total number of outcomes in } \Omega}$$

- 4.21. You roll a fair four sided die, and then a fair six sided die. You add the numbers on the two dice. What is the probability the result is even?
- 4.22. You roll a fair 20 sided die. What is the probability of getting an even number?
- 4.23. You roll a fair five sided die. What is the probability of getting an even number?
- 4.24. *I am indebted to Amin Sadeghi for this exercise.* You must sort four balls into two buckets. There are two white, one red and one green ball.
- (a) For each ball, you choose a bucket independently and at random, with probability  $\frac{1}{2}$ . Show that the probability each bucket has a colored ball in it is  $\frac{1}{2}$ .
- (b) You now choose to sort these balls in such a way that each bucket has two balls in it. You can do so by generating a permutation of the balls uniformly and at random, then placing the first two balls in the first bucket and the second two balls in the second bucket. Show that there are 16 permutations where there is one colored ball in each bucket.
- (c) Use the results of the previous step to show that, using the sorting procedure of that step, the probability of having a colored ball in each bucket is  $\frac{2}{3}$ .
- (d) Why do the two sorting procedures give such different outcomes?

### Permutations and Combinations

- 4.25. You flip a fair coin seven times. What is the probability that you see three H's and two T's?
- 4.26. An airline sells  $T$  tickets for a flight with  $S$  seats, where  $T > S$ . Passengers turn up for the flight independently, and the probability that a passenger with a ticket will turn up for a flight is  $p_t$ . The pilot is eccentric, and will fly only if precisely  $E$  passengers turn up, where  $E < S$ . Write an expression for the probability the pilot will fly.
- 4.27. You shuffle a standard deck of playing cards, and deal a hand of 10 cards. With what probability does this hand have five red cards?
- 4.28. Magic the Gathering is a popular card game. Cards can be land cards, or other cards. We consider a game with two players. Each player has a deck of 40 cards. Each player shuffles their deck, then deals seven cards, called their *hand*.
- (a) Assume that player one has 10 land cards in their deck and player two has

20. With what probability will *each* player have four lands in their hand?
- (b) Assume that player one has 10 land cards in their deck and player two has 20. With what probability will player one have two lands and player two have three lands in hand?
- (c) Assume that player one has 10 land cards in their deck and player two has 20. With what probability will player two have more lands in hand than player one?
- 4.29. The previous exercise divided Magic the Gathering cards into lands vs. other. We now recognize four kinds of cards: land, spell, creature and artifact. We consider a game with two players. Each player has a deck of 40 cards. Each player shuffles their deck, then deals seven cards, called their *hand*.
- (a) Assume that player one has 10 land cards, 10 spell cards, 10 creature cards and 10 artifact cards in their deck. With what probability will player one have at least one of each kind of card in hand?
- (b) Assume that player two has 20 land cards, 5 spell cards, 7 creature cards and 8 artifact cards in their deck. With what probability will player two have at least one of each kind of card in hand?
- (c) Assume that player one has 10 land cards, 10 spell cards, 10 creature cards and 10 artifact cards in their deck; and player two has 20 land cards, 5 spell cards, 7 creature cards and 8 artifact cards in their deck. With what probability will at least one of the players have at least one of each kind card in hand?

#### Independence

- 4.30. Event  $\mathcal{A}$  has  $P(\mathcal{A}) = 0.5$ . Event  $\mathcal{B}$  has  $P(\mathcal{B}) = 0.2$ . We also know that  $P(\mathcal{A} \cup \mathcal{B}) = 0.65$ . Are A and B independent? Why
- 4.31. You take a standard deck of cards, shuffle it, and remove both red kings. You then draw a card.
- (a) Is the event {card is red} independent of the event {card is a queen}?
- (b) Is the event {card is black} independent of the event {card is a king}?

#### Conditional Probability

- 4.32. You roll two fair six-sided dice. What is the conditional probability the sum of numbers is greater than three, conditioned on the first die coming up even.
- 4.33. You take a standard deck of cards, shuffle it, and remove one card. You then draw a card.
- (a) What is the conditional probability that the card you draw is a red king, conditioned on the removed card being a king?
- (b) What is the conditional probability that the card you draw is a red king, conditioned on the removed card being a red king?
- (c) What is the conditional probability that the card you draw is a red king, conditioned on the removed card being a black ace?
- 4.34. A royal flush is a hand of five cards, consisting of Ace, King, Queen, Jack and 10 of a single suit. Poker players like this hand, but don't see it all that often.
- (a) You draw five cards from a standard deck of playing cards. What is the probability of getting a royal flush?
- (b) You draw three cards from a standard deck of playing cards. These are Ace, King, Queen of hearts. What is the probability that the next two cards you draw will result in a getting a royal flush? (this is the conditional

probability of getting a royal flush, conditioned on the first three cards being AKQ of hearts).

- 4.35.** You roll a fair five-sided die, and a fair six-sided die.
- What is the probability that the sum of numbers is even?
  - What is the conditional probability that the sum of numbers is even, conditioned on the six-sided die producing an odd number?
- 4.36.** You take a standard deck of playing cards, shuffle it, and remove 13 cards without looking at them. You then shuffle the resulting deck of 39 cards, and draw three cards. Each of these three cards is red. What is the conditional probability that every card you removed is black?
- 4.37.** Magic the Gathering is a popular card game. Cards can be land cards, or other cards. We will consider a deck of 40 cards, containing 10 land cards and 30 other cards. A player shuffles that deck, and draws seven cards but does not look at them. The player then chooses one of these cards at random; it is a land.
- What is the conditional probability that the original hand of seven cards is all lands?
  - What is the conditional probability that the original hand of seven cards contains only one land?
- 4.38.** Magic the Gathering is a popular card game. Cards can be land cards, or other cards. We will consider a deck of 40 cards, containing 10 land cards and 30 other cards. A player shuffles that deck, and draws seven cards but does not look at them. The player then chooses three of these cards at random; each of these three is a land.
- What is the conditional probability that the original hand of seven cards is all lands?
  - What is the conditional probability that the original hand of seven cards contains only three lands?
- 4.39.** You take a standard deck of playing cards, and remove one card at random. You then draw a single card. Write  $\mathcal{S}$  for the event that the card you remove is a six. Write  $\mathcal{N}$  for the event that the card you remove is not a six. Write  $\mathcal{R}$  for the event that the card you remove is red. Write  $\mathcal{B}$  for the event the card you remove is black.
- Write  $\mathcal{A}$  for the event you draw a 6. What is  $P(\mathcal{A}|\mathcal{S})$ ?
  - Write  $\mathcal{A}$  for the event you draw a 6. What is  $P(\mathcal{A}|\mathcal{N})$ ?
  - Write  $\mathcal{A}$  for the event you draw a 6. What is  $P(\mathcal{A})$ ?
  - Write  $\mathcal{D}$  for the event you draw a red six. Are  $\mathcal{D}$  and  $\mathcal{A}$  independent? why?
  - Write  $\mathcal{D}$  for the event you draw a red six. What is  $P(\mathcal{D})$ ?
- 4.40.** A student takes a multiple choice test. Each question has  $N$  answers. If the student knows the answer to a question, the student gives the right answer, and otherwise guesses uniformly and at random. The student knows the answer to 70% of the questions. Write  $\mathcal{K}$  for the event a student knows the answer to a question and  $\mathcal{R}$  for the event the student answers the question correctly.
- What is  $P(\mathcal{K})$ ?
  - What is  $P(\mathcal{R}|\mathcal{K})$ ?
  - What is  $P(\mathcal{K}|\mathcal{R})$ , as a function of  $N$ ?
  - What values of  $N$  will ensure that  $P(\mathcal{K}|\mathcal{R}) > 99\%$ ?

## The Monty Hall Problem

- 4.41. Monty Hall, Rule 3:** If the host uses rule 3, then what is  $P(C_1|G_2, r_3)$ ? Do this by computing conditional probabilities.
- 4.42. Monty Hall, Rule 4:** If the host uses rule 4, and shows you a goat behind door 2, what is  $P(C_1|G_2, r_4)$ ? Do this by computing conditional probabilities.

## CHAPTER 5

# Random Variables and Expectations

We have machinery to describe experiments with random outcomes, but we mostly care about numbers that are random. It is straightforward to link a number to the outcome of an experiment. The result is a random variable, a useful new idea. Random variables turn up in all sorts of places. For example, the amount of money you win or lose on a bet is a random variable. Now if you take the same bet repeatedly, you could wonder how much money will change hands in total, per bet. The answer to questions like this is the expected value of a random variable.

Expected values have strong properties. When one knows some expected values, you can bound various probabilities. This phenomenon parallels the property of data that we saw earlier – you don't find a large fraction of the dataset many standard deviations away from the mean. Particularly important to computer scientists (and gamblers!) is the weak law of large numbers. This law says, loosely, that the value, per bet, of repeating a bet many times will almost certainly be the expected value. Among other things, this law legitimizes estimating expectations and probabilities that are hard to calculate by using a simulation. This turns out to be really useful, because simulations are often quite easy programs to write.

### 5.1 RANDOM VARIABLES

Quite commonly, we would like to deal with numbers that are random. We can do so by linking numbers to the outcome of an experiment. We define a **random variable**:

**Definition: 5.1** *Discrete random variable*

Given a sample space  $\Omega$ , a set of events  $\mathcal{F}$ , and a probability function  $P$ , and a countable set of real numbers  $D$ , a discrete random variable is a function with domain  $\Omega$  and range  $D$ .

Indexing: Discrete random variable

Indexing: Discrete random variable

This means that for any outcome  $\omega$  there is a number  $X(\omega)$ .  $P$  will play an important role, but first we give some examples.

**Example: 5.1** *Numbers from coins*

We flip a coin. Whenever the coin comes up heads, we report 1; when it comes up tails, we report 0. This is a random variable.

**Example: 5.2** *Numbers from coins II*

We flip a coin 32 times. We record a 1 when it comes up heads, and when it comes up tails, we record a 0. This produces a 32 bit random number, which is a random variable.

**Example: 5.3** *The number of pairs in a poker hand*

(from Stirzaker). We draw a hand of five cards. The number of pairs in this hand is a random variable, which takes the values 0, 1, 2 (depending on which hand we draw)

A function that takes a discrete random variable to a set of numbers is also a discrete random variable.

**Example: 5.4** *Parity of coin flips*

We flip a coin 32 times. We record a 1 when it comes up heads, and when it comes up tails, we record a 0. This produces a 32 bit random number, which is a random variable. The parity of this number is also a random variable.

Associated with any value  $x$  of the random variable  $X$  are a series of events. The most important is the set of outcomes  $\omega$  such that  $X(\omega) = x$ , which we can write  $\{\omega : X(\omega) = x\}$ ; it is usual to simplify to  $\{X = x\}$ , and we will do so. The probability that a random variable  $X$  takes the value  $x$  is given by  $P(\{\omega : X(\omega) = x\})$ , which is more usually written  $P(\{X = x\})$ . This is sometimes written as  $P(X = x)$ , and rather often written as  $P(x)$ .

We could also be interested in the set of outcomes  $\omega$  such that  $X(\omega) \leq x$  (i.e. in  $\{\omega : X(\omega) \leq x\}$ ), which we will write  $\{X \leq x\}$ ; The probability that  $X$  takes a

value less than or equal to  $x$  is given by  $P(\{\omega : X(\omega) \leq x\})$ , which is more usually written  $P(\{X \leq x\})$ . Similarly, we could be interested in  $\omega$  such that  $\{X(\omega) > x\}$ , and so on.

**Definition: 5.2** *Probability distribution of a discrete random variable*

The probability distribution of a discrete random variable is the set of numbers  $P(\{X = x\})$  for each value  $x$  that  $X$  can take. The distribution takes the value 0 at all other numbers. Notice that the distribution is non-negative. **Notation warning:** probability notation can be quirky. You may encounter  $p(x)$  with the meaning “some probability distribution” or  $p(x)$  meaning “the value of the probability distribution  $P(\{X = x\})$  at the point  $x$ ” or  $p(x)$  with the meaning “the probability distribution  $P(\{X = x\})$ ”. Context may help disambiguate these uses.

Indexing: Probability distribution of a discrete random variable

Indexing: Probability distribution of a discrete random variable

**Definition: 5.3** *Cumulative distribution of a discrete random variable*

The cumulative distribution of a discrete random variable is the set of numbers  $P(\{X \leq x\})$  for each value  $x$  that  $X$  can take. Notice that this is a non-decreasing function of  $x$ . **Notation warning:** Cumulative distributions are often written with an  $f$ , so that  $f(x)$  might mean  $P(\{X \leq x\})$ .

Indexing: Cumulative distribution of a discrete random variable

Indexing: Cumulative distribution of a discrete random variable

**Worked example 5.1** *Numbers from coins III*

We flip a biased coin 2 times. The flips are independent. The coin has  $P(H) = p$ ,  $P(T) = 1 - p$ . We record a 1 when it comes up heads, and when it comes up tails, we record a 0. This produces a 2 bit random number, which is a random variable taking the values 0, 1, 2, 3. What is the probability distribution and cumulative distribution of this random variable?

**Solution:** Probability distribution:  $P(0) = (1 - p)^2$ ;  $P(1) = (1 - p)p$ ;  $P(2) = p(1 - p)$ ;  $P(3) = p^2$ . Cumulative distribution:  $f(0) = (1 - p)^2$ ;  $f(1) = (1 - p)$ ;  $f(2) = p(1 - p) + (1 - p) = (1 - p)^2$ ;  $f(3) = 1$ .

**Worked example 5.2** *Betting on coins*

One way to get a random variable is to think about the reward for a bet. We agree to play the following game. I flip a coin. The coin has  $P(H) = p$ ,  $P(T) = 1 - p$ . If the coin comes up heads, you pay me  $q$ ; if the coin comes up tails, I pay you  $r$ . The number of dollars that change hands is a random variable. What is its probability distribution?

**Solution:** We see this problem from my perspective. If the coin comes up heads, I get  $q$ ; if it comes up tails, I get  $-r$ . So we have  $P(X = q) = p$  and  $P(X = -r) = (1 - p)$ , and all other probabilities are zero.

## 5.1.1 Joint and Conditional Probability for Random Variables

All the concepts of probability that we described for events carry over to random variables. This is as it should be, because random variables are really just a way of getting numbers out of events. However, terminology and notation change a bit.

Assume we have two random variables  $X$  and  $Y$ . The probability that  $X$  takes the value  $x$  and  $Y$  takes the value  $y$  could be written as  $P(\{X = x\} \cap \{Y = y\})$ . It is more usual to write it as  $P(x, y)$ . You can think of this as a table of values, one for each possible pair of  $x$  and  $y$  values. This table is usually referred to as the **joint probability distribution** of the random variables. Nothing (except notation) has really changed here, but the change of notation is useful.

We will simplify notation further. Usually, we are interested in random variables, rather than potentially arbitrary outcomes or sets of outcomes. We will write  $P(X)$  to denote the probability distribution of a random variable, and  $P(x)$  or  $P(X = x)$  to denote the probability that that random variable takes a particular value. This means that, for example, the rule we could write as

$$P(\{X = x\} | \{Y = y\})P(\{Y = y\}) = P(\{X = x\} \cap \{Y = y\})$$

will be written as

$$P(x|y)P(y) = P(x, y).$$

Recall the rule from section 4.3.1:

$$P(\mathcal{A}|\mathcal{B}) = \frac{P(\mathcal{B}|\mathcal{A})P(\mathcal{A})}{P(\mathcal{B})}.$$

This rule can be rewritten in our notation for random variables. This is the most familiar form of **Bayes' rule**, which is important enough to appear in its own box.

**Definition: 5.4** *Bayes' rule*

$$P(x|y) = \frac{P(y|x)P(x)}{P(y)}$$

Indexing: Bayes' rule

Indexing: Bayes' rule

Random variables have another useful property. If  $x_0 \neq x_1$ , then the event  $\{X = x_0\}$  must be disjoint from the event  $\{X = x_1\}$ . This means that

$$\sum_x P(x) = 1$$

and that, for any  $y$ ,

$$\sum_x P(x|y) = 1$$

(if you're uncertain on either of these points, check them by writing them out in the language of events).

Now assume we have the joint probability distribution of two random variables,  $X$  and  $Y$ . Recall that we write  $P(\{X = x\} \cap \{Y = y\})$  as  $P(x, y)$ . Now consider the sets of outcomes  $\{Y = y\}$  for each different value of  $y$ . These sets must be disjoint, because  $y$  cannot take two values at the same time. Furthermore, each element of the set of outcomes  $\{X = x\}$  must lie in one of the sets  $\{Y = y\}$ . So we have

$$\sum_y P(\{X = x\} \cap \{Y = y\}) = P(\{X = x\})$$

which is usually written as

$$\sum_y P(x, y) = P(x)$$

and is often referred to as the **marginal probability** of  $X$ .

**Definition: 5.5** *Independent random variables*

The random variables  $X$  and  $Y$  are **independent** if the events  $\{X = x\}$  and  $\{Y = y\}$  are independent. This means that

$$P(\{X = x\} \cap \{Y = y\}) = P(\{X = x\})P(\{Y = y\}),$$

which we can rewrite as

$$P(x, y) = P(x)P(y)$$

Indexing: Independent random variables

Indexing: Independent random variables

**Worked example 5.3** *Sums and differences of dice*

You throw two dice. The number of spots on the first die is a random variable (call it  $X$ ); so is the number of spots on the second die ( $Y$ ).  $X$  and  $Y$  are independent. Now define  $S = X + Y$  and  $D = X - Y$ . What is the probability distribution of  $S$  and of  $D$ ?

**Solution:**  $S$  can have values in the range  $2, \dots, 12$ . There is only one way to get a  $S = 2$ ; two ways to get  $S = 3$ ; and so on. Using the methods of chapter 13.1 for each case, the probabilities for  $[2, 3, 4, 5, 6, 7, 8, 9, 10, 11, 12]$  are  $[1, 2, 3, 4, 5, 6, 5, 4, 3, 2, 1]/36$ . Similarly,  $D$  can have values in the range  $-5, \dots, 5$ . Again, using the methods of chapter 13.1, the probabilities for  $[-5, -4, -3, -2, -1, 0, 1, 2, 3, 4, 5]$  are  $[1, 2, 3, 4, 5, 6, 5, 4, 3, 2, 1]/36$ .

**Worked example 5.4** *Sums and differences of dice, II*

Using the terminology of example 5.3, what is the joint probability distribution of  $S$  and  $D$ ?

**Solution:** This is more interesting to display, because it's an  $11 \times 11$  table. Each entry of the table represents a pair of  $S, D$  values. Many pairs can't occur (for example, for  $S = 2$ ,  $D$  can only be zero; if  $S$  is even, then  $D$  must be even; and so on). You can work out the table by checking each case; it's in Table 5.1.

**Worked example 5.5** *Sums and differences of dice, III*

Using the terminology of example 5.3, are  $X$  and  $Y$  independent? are  $S$  and  $D$  independent?

**Solution:**  $X$  and  $Y$  are clearly independent. But  $S$  and  $D$  are not. There are several ways to see this. One way is to notice that, if you know  $S = 2$ , then you know the value of  $D$  precisely; but if you know  $S = 3$ ,  $D$  could be either 1 or  $-1$ . This means that  $P(S|D)$  depends on  $D$ , so they're not independent. Another way is to notice that the rank of the table, as a matrix, is 6, which means that it can't be the outer product of two vectors.

$$\frac{1}{36} \times \begin{pmatrix} 0 & 0 & 0 & 0 & 0 & 1 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 1 & 0 & 1 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 & 0 & 1 & 0 & 1 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 & 1 & 0 & 1 & 0 & 1 & 0 & 0 \\ 0 & 1 & 0 & 1 & 0 & 1 & 0 & 1 & 0 & 1 & 0 \\ 1 & 0 & 1 & 0 & 1 & 0 & 1 & 0 & 1 & 0 & 1 \\ 0 & 1 & 0 & 1 & 0 & 1 & 0 & 1 & 0 & 1 & 0 \\ 0 & 0 & 1 & 0 & 1 & 0 & 1 & 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 1 & 0 & 1 & 0 & 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 1 & 0 & 1 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 1 & 0 & 0 & 0 & 0 & 0 \end{pmatrix}$$

TABLE 5.1: A table of the joint probability distribution of  $S$  (vertical axis; scale  $2, \dots, 12$ ) and  $D$  (horizontal axis; scale  $-5, \dots, 5$ ) from example 5.4

**Worked example 5.6** *Sums and differences of dice, IV*

Using the terminology of example 5.3, what is  $P(S|D = 0)$ ? what is  $P(D|S = 11)$ ?

**Solution:** You could work it out either of these from the table, or by first principles. If  $D = 0$ ,  $S$  can have values  $2, 4, 6, 8, 10, 12$ , and each value has conditional probability  $1/6$ . If  $S = 11$ ,  $D$  can have values  $1$ , or  $-1$ , and each value has conditional probability  $1/2$ .

### 5.1.2 Just a Little Continuous Probability

Our random variables take values from a discrete set of numbers  $D$ . This makes the underlying machinery somewhat simpler to describe, and is often, but not always, enough for model building. Some phenomena are more naturally modelled as being continuous — for example, human height; human weight; the mass of a distant star; and so on. Giving a complete formal description of probability on a continuous space is surprisingly tricky, and would involve us in issues that do not arise much in practice.

These issues are caused by two interrelated facts: real numbers have infinite precision; and you can't count real numbers. A continuous random variable is still a random variable, and comes with all the stuff that a random variable comes with. We will not speculate on what the underlying sample space is, nor on the underlying events. This can all be sorted out, but requires moderately heavy lifting that isn't particularly illuminating for us. The most interesting thing for us is specifying the probability distribution. Rather than talk about the probability that a real number takes a particular value (which we can't really do satisfactorily most of the time), we will instead talk about the probability that it lies in some interval. So we can specify a probability distribution for a continuous random variable by giving a set

of (very small) intervals, and for each interval providing the probability that the random variable lies in this interval.

The easiest way to do this is to supply a **probability density function**. Let  $p(x)$  be a probability density function for a continuous random variable  $X$ . We interpret this function by thinking in terms of small intervals. Assume that  $dx$  is an infinitesimally small interval. Then

$$p(x)dx = P(\{\text{event that } X \text{ takes a value in the range } [x, x + dx]\}).$$

Important properties of probability density functions follow from this definition.

**Useful Facts: 5.1** *Probability density functions*

- Probability density functions are non-negative. This follows from the definition; a negative value at some  $u$  would imply that  $P(\{x \in [u, u + du]\})$  was negative, and this cannot occur.
- For  $a < b$

$$P(\{X \text{ takes a value in the range } [a, b]\}) = \int_a^b p(x)dx.$$

which we obtain by summing  $p(x)dx$  over all the infinitesimal intervals between  $a$  and  $b$ .

- We must have that

$$\int_{-\infty}^{\infty} p(x)dx = 1.$$

This is because

$$P(\{X \text{ takes a value in the range } [-\infty, \infty]\}) = 1 = \int_{-\infty}^{\infty} p(x)dx$$

- Probability density functions are usually called pdf's.
- It is quite usual to write all pdf's as lower-case  $p$ 's. If one specifically wishes to refer to probability, one writes an upper case  $P$ , as in the previous points.

Indexing: Probability density functions

Indexing: Probability density functions

The property that  $\int_{-\infty}^{\infty} p(x)dx = 1$  is useful, because when we are trying to determine a probability density function, we can ignore a constant factor. So if  $g(x)$  is a non-negative function that is proportional to the probability density function (often pdf) we are interested in, we can recover the pdf by computing

$$p(x) = \frac{1}{\int_{-\infty}^{\infty} g(x)dx} g(x).$$

This procedure is sometimes known as **normalizing**, and  $\int_{-\infty}^{\infty} g(x)dx$  is the **normalizing constant**.

One good way to think about pdf's is as the limit of a histogram. Imagine you collect an arbitrarily large dataset of data items, each of which is independent. You build a histogram of that dataset, using arbitrarily narrow boxes. You scale the histogram so that the sum of the box areas is one. The result is a probability density function.

The pdf doesn't represent the probability that a random variable takes a value. Instead, you should think of  $p(x)$  as being the limit of a ratio (which is why it's called a density):

$$\frac{\text{the probability that the random variable will lie in a small interval centered on } x}{\text{the length of the small interval centered on } x}$$

Notice that, while a pdf has to be non-negative, and it has to integrate to 1, it does *not* have to be smaller than one. A ratio like this could be a lot larger than one, as long as it isn't larger than one for too many  $x$  (because the integral must be one). In fact, probability density functions can be moderately strange functions (exercises).

**Worked example 5.7**    *A probability density function that is larger than one*

Assume we have a physical system that can produce random numbers. It produces numbers in the range 0 to  $\epsilon$ , where  $\epsilon > 0$ . Each number has the same probability of appearing. No number larger than  $\epsilon$  or smaller than 0 can ever appear. What is the probability density function?

**Solution:** Write  $p(x)$  for the probability density function. We must have that  $p(x) = 0$  for  $x < 0$  and  $p(x) = 0$  for  $x > \epsilon$ . We must have that  $p(x)$  is constant between 0 and  $\epsilon$  and that  $\int p(x)dx = 1$ . So

$$p(x) = \begin{cases} 0 & \text{if } x < 0 \\ 0 & \text{if } x > \epsilon \\ \frac{1}{\epsilon} & \text{otherwise} \end{cases}$$

Notice that if  $\epsilon < 1$ , we have that  $p(x) > 1$  for all  $x$ .

## 5.2 EXPECTATIONS AND EXPECTED VALUES

Example 5.2 described a simple game. I flip a coin. The coin has  $P(H) = p$ ,  $P(T) = 1 - p$ . If the coin comes up heads, you pay me  $q$ ; if the coin comes up tails, I pay you  $r$ . Now imagine we play this game many times. Our frequency definition of probability means that in  $N$  games, we expect to see about  $pN$  heads and  $(1 - p)N$  tails. In turn, this means that my total income from these  $N$  games

should be about  $(pN)q - ((1-p)N)r$ . The  $N$  in this expression is inconvenient; instead, we could say that for any single game, my expected income is

$$pq - (1-p)r.$$

This isn't the actual income from a single game (which would be either  $q$  or  $-r$ , depending on what the coin did). Instead, it's an estimate of what would happen over a large number of games, on a per-game basis. This is an example of an **expected value**.

### 5.2.1 Expected Values

**Definition: 5.6** *Expected value*

Given a discrete random variable  $X$  which takes values in the set  $\mathcal{D}$  and which has probability distribution  $P$ , we define the expected value

$$\mathbb{E}[X] = \sum_{x \in \mathcal{D}} xP(X = x).$$

This is sometimes written which is  $\mathbb{E}_P[X]$ , to clarify which distribution one has in mind

Indexing: Expected value

Indexing: Expected value

Notice that an expected value could take a value that the random variable doesn't take.

**Example: 5.5** *Betting on coins*

We agree to play the following game. I flip a fair coin (i.e.  $P(H) = P(T) = 1/2$ ). If the coin comes up heads, you pay me 1; if the coin comes up tails, I pay you 1. The expected value of my income is 0, even though the random variable never takes that value.

**Worked example 5.8**    *Betting on coins, again*

We agree to play the following game. I flip a fair coin (i.e.  $P(H) = P(T) = 1/2$ ). If the coin comes up heads, you pay me 2; if the coin comes up tails, I pay you 1. What is the expected value of this game?

**Solution:** The expected value of my income is

$$\left(\frac{1}{2}\right) \times 2 - \left(\frac{1}{2}\right) \times 1 = \frac{1}{2}.$$

Notice this isn't even an integer, and there's no way that any one instance of the game would yield a payoff of  $1/2$ . But this is what I would get, per game, if I played many times.

Your intuition is likely to tell you that the game of example 5.8 is good for me and bad for you. This intuition is correct. It turns out that an even stronger statement is possible: playing this game repeatedly is pretty much guaranteed to be excellent for me and disastrous for you. It'll take some pages before I can be crisp about precisely what I mean here and why it is true.

**Definition: 5.7**    *Expectation*

Assume we have a function  $f$  that maps a discrete random variable  $X$  into a set of numbers  $\mathcal{D}_f$ . Then  $f(x)$  is a discrete random variable, too, which we write  $F$ . The expected value of this random variable is written

$$\mathbb{E}[f] = \sum_{u \in \mathcal{D}_f} uP(F = u) = \sum_{x \in \mathcal{D}} f(x)P(X = x)$$

which is sometimes referred to as “the expectation of  $f$ ”. The process of computing an expected value is sometimes referred to as “taking expectations”.

Indexing: Expectation

Indexing: Expectation

We can compute expectations for continuous random variables, too, though summing over all values now turns into an integral. Assume I have a continuous random variable  $X$  with probability density function  $p(x)$ . Remember I interpret the probability density function as meaning that, for an infinitesimal interval size  $dx$ ,  $p(x)dx = P(\{X \in [x, x + dx]\})$ . Divide the set of possible values that  $X$  can take into small intervals of width  $\Delta x$ , centered on  $x_i$ . We can construct a discrete random variable  $\hat{X}$  which takes values  $x_i$ . We have that  $P(\{\hat{X} = x_i\}) \approx p(x_i)\Delta x$ , where I used the approximation sign because  $\Delta x$  may not be infinitesimally small.

Now write  $\mathbb{E}[\hat{X}]$  for the expected value of  $\hat{X}$ . We have

$$\mathbb{E}[\hat{X}] = \sum_{x_i} x_i P(x_i) \approx \sum_{x_i} x_i p(x_i) \Delta x.$$

As the intervals limit to infinitesimal intervals,  $\hat{X}$  limits to  $X$  (think of a picture of a histogram with infinitely narrow boxes). Then  $\mathbb{E}[\hat{X}]$  has a limit which is an integral, and this defines the expected value. So we have the expressions in the boxes below.

**Definition: 5.8** *Expected value*

Given a continuous random variable  $X$  which takes values in the set  $\mathcal{D}$  and which has probability distribution  $P$ , we define the expected value

$$\mathbb{E}[X] = \int_{x \in \mathcal{D}} xp(x)dx.$$

This is sometimes written  $\mathbb{E}_p[X]$ , to clarify which distribution one has in mind.

Indexing: Expected value

Indexing: Expected value

The expected value of a continuous random variable could be a value that the random variable doesn't take, too. Notice one attractive feature of the  $\mathbb{E}[X]$  notation; we don't need to make any commitment to whether  $X$  is a discrete random variable (where we would write a sum) or a continuous random variable (where we would write an integral). The reasoning by which we turned a sum into an integral works for functions of continuous random variables, too.

**Definition: 5.9** *Expectation*

Assume we have a function  $f$  that maps a continuous random variable  $X$  into a set of numbers  $\mathcal{D}_f$ . Then  $f(x)$  is a continuous random variable, too, which we write  $F$ . The expected value of this random variable is

$$\mathbb{E}[f] = \int_{x \in \mathcal{D}} f(x)p(x)dx$$

which is sometimes referred to as “the expectation of  $f$ ”. The process of computing an expected value is sometimes referred to as “taking expectations”.

Indexing: Expectation

Indexing: Expectation Under some circumstances the expected value may not exist. The integral needs to exist, and be finite, for us to interpret the expected value meaningfully, and that isn't guaranteed for every continuous random variable. Nothing we do will encounter this issue, and so we will ignore it.

You can see an expectation as an operation you apply to a random variable. It doesn't matter whether the random variable is discrete or continuous; that just changes the recipe for computing the value of the expectation. The crucial property of this operation is that it is linear; this is so important I have put it in its own box.

**Useful Facts: 5.2** *Expectations are linear.*

Write  $f, g$  for functions of random variables.

- $\mathbb{E}[0] = 0$
- for any constant  $k$ ,  $\mathbb{E}[kf] = k\mathbb{E}[f]$
- $\mathbb{E}[f + g] = \mathbb{E}[f] + \mathbb{E}[g]$ .

Indexing: Expectations are linear.

Indexing: Expectations are linear.

I have written this box in a rather compact form. This is because the expression  $\mathbb{E}[X]$  for the expected value of a random variable is actually a special case of  $\mathbb{E}[f]$  — one just uses the identity function for  $f$ . So the box also tells us that  $\mathbb{E}[X + Y] = \mathbb{E}[X] + \mathbb{E}[Y]$ , and so on.

### 5.2.2 Mean, Variance and Covariance

There are three very important expectations with special names.

**Definition: 5.10** *Mean or expected value*

The mean or expected value of a random variable  $X$  is

$$\mathbb{E}[X]$$

Indexing: Mean or expected value

Indexing: Mean or expected value

**Definition: 5.11**    *Variance*

The variance of a random variable  $X$  is

$$\text{var}[X] = \mathbb{E}[(X - \mathbb{E}[X])^2]$$

Indexing: Variance

Indexing: Variance

**Useful Facts: 5.3**    *Variance*

1. For any constant  $k$ ,  $\text{var}[k] = 0$
2.  $\text{var}[X] \geq 0$
3.  $\text{var}[kX] = k^2 \text{var}[X]$
4. if  $X$  and  $Y$  are independent, then  $\text{var}[X + Y] = \text{var}[X] + \text{var}[Y]$

1 and 2 are obvious. You will prove 3 and 4 in the exercises.

Indexing: Variance

Indexing: Variance

**Definition: 5.12**    *Covariance*

The covariance of two random variables  $X$  and  $Y$  is

$$\text{cov}(X, Y) = \mathbb{E}[(X - \mathbb{E}[X])(Y - \mathbb{E}[Y])]$$

Indexing: Covariance

Indexing: Covariance

You should notice that  $\text{var}[X] = \text{cov}(X, X)$ .

**Worked example 5.9**    *Mean of a coin flip*

We flip a biased coin, with  $P(H) = p$ . The random variable  $X$  has value 1 if the coin comes up heads, 0 otherwise. What is the mean of  $X$ ? (i.e.  $\mathbb{E}[X]$ ).

**Solution:**  $\mathbb{E}[X] = \sum_{x \in D} xP(X = x) = 1p + 0(1 - p) = p$

**Worked example 5.10**    *Variance*

Can a random variable have  $\mathbb{E}[X] > \sqrt{\mathbb{E}[X^2]}$ ?

**Solution:** No, because that would mean that  $\mathbb{E}[(X - \mathbb{E}[X])^2] < 0$ . But this is the expected value of a non-negative quantity; it must be non-negative.

**Worked example 5.11**    *More variance*

We just saw that a random variable can't have  $\mathbb{E}[X] > \sqrt{\mathbb{E}[X^2]}$ . But I can easily have a random variable with large mean and small variance - isn't this a contradiction?

**Solution:** No, you're confused. Your question means you think that the variance of  $X$  is given by  $\mathbb{E}[X^2]$ ; but actually  $\text{var}[X] = \mathbb{E}[X^2] - \mathbb{E}[X]^2$

**Useful Facts: 5.4**    *A useful expression for variance*

$$\begin{aligned} \text{var}[X] &= \mathbb{E}[(X - \mathbb{E}[X])^2] \\ &= \mathbb{E}\left[X^2 - 2X\mathbb{E}[X] + \mathbb{E}[X]^2\right] \\ &= \mathbb{E}[X^2] - 2\mathbb{E}[X]\mathbb{E}[X] + \mathbb{E}[X]^2 \\ &= \mathbb{E}[X^2] - (\mathbb{E}[X])^2 \end{aligned}$$

Indexing: A useful expression for variance

Indexing: A useful expression for variance

**Useful Facts: 5.5** *A useful expression for covariance*

$$\begin{aligned}
 \text{cov}(X, Y) &= \mathbb{E}[(X - \mathbb{E}[X])(Y - \mathbb{E}[Y])] \\
 &= \mathbb{E}[(XY - Y\mathbb{E}[X] - X\mathbb{E}[Y] + \mathbb{E}[X]\mathbb{E}[Y])] \\
 &= \mathbb{E}[XY] - 2\mathbb{E}[Y]\mathbb{E}[X] + \mathbb{E}[X]\mathbb{E}[Y] \\
 &= \mathbb{E}[XY] - \mathbb{E}[X]\mathbb{E}[Y].
 \end{aligned}$$

Indexing: A useful expression for covariance

Indexing: A useful expression for covariance

Now assume that we have a probability distribution  $P(X)$  defined on some discrete set of numbers. There is some random variable that produced this probability distribution. This means that we could talk about the mean of a probability distribution  $P$  (rather than the mean of a random variable whose probability distribution is  $P(X)$ ). It is quite usual to talk about the mean of a probability distribution. Furthermore, we could talk about the variance of a probability distribution  $P$  (rather than the variance of a random variable whose probability distribution is  $P(X)$ ).

**Useful Facts: 5.6** *Variance and Covariance*

1. if  $X$  and  $Y$  are independent, then  $\mathbb{E}[XY] = \mathbb{E}[X]\mathbb{E}[Y]$ .
2. if  $X$  and  $Y$  are independent, then  $\text{cov}(X, Y) = 0$ .

If 5 is true, then 6 is obviously true (apply the expression of useful facts 5.5). I prove 5 below.

Indexing: Variance and Covariance

Indexing: Variance and Covariance

**Proposition:**    *If  $X$  and  $Y$  are independent random variables, then  $\mathbb{E}[XY] = \mathbb{E}[X]\mathbb{E}[Y]$ .*

**Proof:** Recall that  $\mathbb{E}[X] = \sum_{x \in D} xP(X = x)$ , so that

$$\begin{aligned} \mathbb{E}[XY] &= \sum_{(x,y) \in D_x \times D_y} xyP(X = x, Y = y) \\ &= \sum_{x \in D_x} \sum_{y \in D_y} (xyP(X = x, Y = y)) \\ &= \sum_{x \in D_x} \sum_{y \in D_y} (xyP(X = x)P(Y = y)) \\ &\quad \text{because } X \text{ and } Y \text{ are independent} \\ &= \sum_{x \in D_x} \sum_{y \in D_y} (xP(X = x)) (yP(Y = y)) \\ &= \left( \sum_{x \in D_x} xP(X = x) \right) \left( \sum_{y \in D_y} yP(Y = y) \right) \\ &= (\mathbb{E}[X])(\mathbb{E}[Y]). \end{aligned}$$

This is certainly not true when  $X$  and  $Y$  are not independent (try  $Y = -X$ ).

**Worked example 5.12**    *Variance of a coin flip*

We flip a biased coin, with  $P(H) = p$ . The random variable  $X$  has value 1 if the coin comes up heads, 0 otherwise. What is the variance of  $X$ ? (i.e.  $\text{var}[X]$ ).

**Solution:**  $\text{var}[X] = \mathbb{E}[(X - \mathbb{E}[X])^2] = \mathbb{E}[X^2] - \mathbb{E}[X]^2 = (1p - 0(1-p)) - p^2 = p(1-p)$

The variance of a random variable is often inconvenient, because its units are the square of the units of the random variable. Instead, we could use the **standard deviation**.

**Definition: 5.13** *Standard deviation*

The **standard deviation** of a random variable  $X$  is defined as

$$\text{std}(\{X\}) = \sqrt{\text{var}[X]}$$

Indexing: Standard deviation

Indexing: Standard deviation

You do need to be careful with standard deviations. If  $X$  and  $Y$  are independent random variables, then  $\text{var}[X + Y] = \text{var}[X] + \text{var}[Y]$ , but  $\text{std}(\{X + Y\}) = \sqrt{\text{std}(\{X\})^2 + \text{std}(\{Y\})^2}$ . One way to avoid getting mixed up is to remember that variances add, and derive expressions for standard deviations from that.

### 5.2.3 Expectations and Statistics

You should have noticed we now have two notions each for mean, variance, covariance, and standard deviation. One, which we expounded in section 2.3, describes datasets. We will call these **descriptive statistics**. The other, described above, is a property of probability distributions. We will call these **expectations**. In each case, the reason we have one name for two notions is that the notions are not really all that different.

Imagine we have a dataset  $\{x\}$  of  $N$  items, where the  $i$ 'th item is  $x_i$ . We can build a probability distribution out of this dataset, by placing a probability on each data item. We will give each data item the same probability (which must be  $1/N$ , so all probabilities add to 1). Write  $\mathbb{E}[x]$  for the mean of this distribution. We have

$$\mathbb{E}[x] = \sum_i x_i p(x_i) = \frac{1}{N} \sum_i x_i = \text{mean}(\{x\}).$$

The variances, standard deviations and covariance have the same property: For this particular distribution (sometimes called the **empirical distribution**), the expectations have the same value as the descriptive statistics (exercises).

In section 5.3.3, we will see a form of converse to this fact. Imagine we have a dataset that consists of independent, identically distributed samples from a probability distribution (i.e. we know that each data item was obtained independently from the distribution). For example, we might have a count of heads in each of a number of coin flip experiments. Then the descriptive statistics will turn out to be accurate estimates of the expectations.

## 5.3 TWO INEQUALITIES, AND A LAW

Mean and variance tell us quite a lot about a random variable, as two important inequalities show.

**Definition: 5.14** *Markov's inequality*

**Markov's inequality** is

$$P(\{|X| \geq a\}) \leq \frac{\mathbb{E}[|X|]}{a}.$$

Indexing: Markov's inequality

Indexing: Markov's inequality

Markov's inequality means that a random variable is most unlikely to have an absolute value a lot larger than the mean of its absolute value. This should seem fairly intuitive from the definition of expectation. Recall that

$$\mathbb{E}[X] = \sum_{x \in D} xP(\{X = x\})$$

Assume that  $D$  contains only non-negative numbers (that absolute value). Then the only way to have a small value of  $\mathbb{E}[X]$  is to be sure that, when  $x$  is large,  $P(\{X = x\})$  is small. The proof is a rather more formal version of this observation, and appears below.

**Definition: 5.15** *Chebyshev's inequality*

**Chebyshev's inequality** is

$$P(\{|X - \mathbb{E}[X]| \geq a\}) \leq \frac{\text{var}[X]}{a^2}.$$

It is common to see this in another form, obtained by writing  $\sigma$  for the standard deviation of  $X$ , substituting  $k\sigma$  for  $a$ , and rearranging

$$P(\{|X - \mathbb{E}[X]| \geq k\sigma\}) \leq \frac{1}{k^2}$$

Indexing: Chebyshev's inequality

Indexing: Chebyshev's inequality

Chebyshev's inequality means that the probability of a random variable taking a particular value must fall off rather fast as that value moves away from the mean, in units scaled to the standard deviation. This probably doesn't seem intuitive from the definition of expectation. But think about it this way: values of a random variable that are many standard deviations above the mean must have

low probability, otherwise the standard deviation would be bigger. The proof is a rather more formal version of this observation, and appears below. We introduce indicator functions next, because an easy proof of these inequalities uses indicator functions.

### 5.3.1 Indicator Functions

An **indicator function** is a function that is one when some condition is true, and zero otherwise. The reason indicator functions are useful is that their expected values have interesting properties.

**Definition: 5.16** *Indicator functions*

An indicator function for an event is a function that takes the value zero for values of  $X$  where the event does not occur, and one where the event occurs. For the event  $\mathcal{E}$ , we write

$$\mathbb{I}_{[\mathcal{E}]}(X)$$

for the relevant indicator function.

Indexing: Indicator functions

Indexing: Indicator functions

For example,

$$\mathbb{I}_{\{|X| \leq a\}}(X) = \begin{cases} 1 & \text{if } -a < X < a \\ 0 & \text{otherwise} \end{cases}$$

Indicator functions have one useful property.

$$\mathbb{E}[\mathbb{I}_{[\mathcal{E}]}] = P(\mathcal{E})$$

which you can establish by checking the definition of expectations.

## 5.3.2 Proving the Inequalities

**Proposition:** *Markov's inequality*

$$P(\{|X| \geq a\}) \leq \frac{\mathbb{E}[|X|]}{a}.$$

**Proof:** (from Wikipedia). Notice that, for  $a > 0$ ,

$$a\mathbb{I}_{\{|X| \geq a\}}(X) \leq |X|$$

(because if  $|X| \geq a$ , the LHS is  $a$ ; otherwise it is zero). Now we have

$$\mathbb{E}[a\mathbb{I}_{\{|X| \geq a\}}] \leq \mathbb{E}[|X|]$$

but, because expectations are linear, we have

$$\mathbb{E}[a\mathbb{I}_{\{|X| \geq a\}}] = a\mathbb{E}[\mathbb{I}_{\{|X| \geq a\}}] = aP(\{|X| \geq a\})$$

and so we have

$$aP(\{|X| \geq a\}) \leq \mathbb{E}[|X|]$$

and we get the inequality by division, which we can do because  $a > 0$ .

**Proposition:** *Chebyshev's inequality*

$$P(\{|X - \mathbb{E}[X]| \geq a\}) \leq \frac{\text{var}[X]}{a^2}.$$

**Proof:** Write  $U$  for the random variable  $(X - \mathbb{E}[X])^2$ . Markov's inequality gives us

$$P(\{|U| \geq w\}) \leq \frac{\mathbb{E}[|U|]}{w}$$

Now notice that, if  $a^2 = w$ ,

$$P(\{|U| \geq w\}) = P(\{|X - \mathbb{E}[X]| \geq a\})$$

so we have

$$P(\{|U| \geq w\}) = P(\{|X - \mathbb{E}[X]| \geq a\}) \leq \frac{\mathbb{E}[|U|]}{w} = \frac{\text{var}[X]}{a^2}$$

## 5.3.3 IID Samples and the Weak Law of Large Numbers

Imagine a random variable  $X$ , obtained by flipping a fair coin and reporting 1 for an  $H$  and  $-1$  for a  $T$ . We can talk about the probability distribution  $P(X)$  of this random variable; we can talk about the expected value that the random variable takes; but the random variable itself doesn't have a value. However, if we actually flip a coin, we get either a 1 or a  $-1$ . Observing a value is sometimes called a **trial**. The resulting value is often called a **sample** of the random variable (or of its probability distribution). So flipping the coin is a trial, and the number you get is a sample. If we flipped a coin many times, we'd have a set of numbers (or samples). These numbers would be independent. Their histogram would look like  $P(X)$ . Collections of data items like this are important enough to have their own name.

Assume we have a set of data items  $x_i$  such that (a) they are independent; (b) each is produced from the same process; and (c) the histogram of a very large set of data items looks increasingly like the probability distribution  $P(X)$  as the number of data items increases. Then we refer to these data items as **independent identically distributed samples** of  $P(X)$ ; for short, **iid samples** or even just **samples**. It's worth knowing that it can be a difficult computational problem to get IID samples from some given probability distribution. For all of the cases we will deal with, it will be obvious how to get IID samples. Usually, they were generated for us — i.e. somebody flipped the coin, etc.

Now assume we have a set of  $N$  IID samples  $x_i$  of a probability distribution  $P(X)$ . Write

$$X_N = \frac{\sum_{i=1}^N x_i}{N}.$$

Now  $X_N$  is a random variable (the  $x_i$  are IID samples, and for a different set of samples you will get a different, random,  $X_N$ ). Notice that  $P(X = x_1, X = x_2, \dots, X = x_n) = P(X = x_1)P(X = x_2) \dots P(X = x_n)$ , because the samples are independent and each is a sample of  $P(X)$ . This means that

$$\mathbb{E}[X_N] = \mathbb{E}[X]$$

because

$$\mathbb{E}[X_N] = \left(\frac{1}{N}\right) \sum_{i=1}^N \mathbb{E}[X].$$

This means that

$$\frac{\sum_{i=1}^N x_i}{N}$$

should be an accurate estimate of  $\mathbb{E}[X]$ . In fact, as  $N$  gets large, the estimate becomes more accurate.

**Definition: 5.17** *Weak Law of Large Numbers*

The **weak law of large numbers** states that, if  $P(X)$  has finite variance, then for any positive number  $\epsilon$

$$\lim_{N \rightarrow \infty} P(\{|X_N - \mathbb{E}[X]| \geq \epsilon\}) = 0.$$

Equivalently, we have

$$\lim_{N \rightarrow \infty} P(\{|X_N - \mathbb{E}[X]| < \epsilon\}) = 1.$$

Each form means that, for a large enough set of IID samples, the average of the samples (i.e.  $X_N$ ) will, with high probability, be very close to the expectation  $\mathbb{E}[X]$ .

Indexing: Weak Law of Large Numbers

Indexing: Weak Law of Large Numbers

**Proposition:** *Weak law of large numbers*

$$\lim_{N \rightarrow \infty} P(\{|X_N - \mathbb{E}[X]| \geq \epsilon\}) = 0.$$

**Proof:** Assume that  $P(X)$  has finite variance; that is,  $\text{var}(\{X\}) = \sigma^2$ . Choose  $\epsilon > 0$ . Now we have that

$$\begin{aligned} \text{var}(\{X_N\}) &= \text{var}\left(\left\{\frac{\sum_{i=1}^N x_i}{N}\right\}\right) \\ &= \left(\frac{1}{N^2}\right)\text{var}\left(\left\{\sum_{i=1}^N x_i\right\}\right) \\ &= \left(\frac{1}{N^2}\right)(N\sigma^2) \\ &\quad \text{because the } x_i \text{ are independent} \\ &= \frac{\sigma^2}{N} \end{aligned}$$

and that

$$\mathbb{E}[X_N] = \mathbb{E}[X].$$

Now Chebyshev's inequality gives

$$P(\{|X_N - \mathbb{E}[X]| \geq \epsilon\}) \leq \frac{\sigma^2}{N\epsilon^2}$$

so

$$\lim_{N \rightarrow \infty} P(\{|X_N - \mathbb{E}[X]| \geq \epsilon\}) = \lim_{N \rightarrow \infty} \frac{\sigma^2}{N\epsilon^2} = 0.$$

Notice that

$$1 - P(\{|X_N - \mathbb{E}[X]| \geq \epsilon\}) = P(\{|X_N - \mathbb{E}[X]| < \epsilon\}) \geq 1 - \frac{\sigma^2}{N\epsilon^2}.$$

so that

$$\lim_{N \rightarrow \infty} P(\{|X_N - \mathbb{E}[X]| < \epsilon\}) = \lim_{N \rightarrow \infty} \left(1 - \frac{\sigma^2}{N\epsilon^2}\right) = 1.$$

## 5.4 USING EXPECTATIONS

The weak law of large numbers gives us a very valuable way of thinking about expectations. Assume we have a random variable  $X$ . Then the weak law says

that, if you observe this random variable over a large number of trials, the average of the values you observe should be very close to  $\mathbb{E}[X]$ . Notice that this extends to functions of random variables (because they are random variables, too). For example, I observe values  $x_i$  of a random variable  $X$  over a large number  $N$  of trials, and compute

$$\frac{1}{N} \sum_{i=1}^N f(x_i).$$

The weak law says that the value I get should be very close to  $\mathbb{E}[f]$ . You can show this by defining a new random variable  $F = f(X)$ . This has a probability distribution  $P(F)$ , which might be difficult to know — but we don't need to.  $\mathbb{E}[f]$ , the expected value of the function  $f$  under the distribution  $P(X)$ . This is the same as  $\mathbb{E}[F]$ , and the weak law applies.

Remember: the average over repeated trials of a random variable is very close to the expectation. You can use this information to make many kinds of decision in uncertain circumstances.

#### 5.4.1 Should you accept a bet?

We can't answer this as a moral question, but we can as a practical question, using expectations. Generally, a bet involves an agreement that amounts of money will change hands, depending on the outcome of an experiment. Mostly, you are interested in how much you get from the bet, so it is natural to give sums of money you receive a positive sign, and sums of money you pay out a negative sign. The weak law says that if you repeat a bet many times, you are increasingly likely to receive the expected value of the bet, per bet. Under this convention, the practical answer is easy: accept a bet enthusiastically if its expected value is positive, otherwise decline it. It is interesting to notice how poorly this advice describes actual human behavior.

In the hope that not every reader is familiar with roulette, I will sketch how the game works, because I can then use it for examples and exercises. A roulette wheel has a collection of slots. There are 36 slots numbered with the digits  $1 \dots 36$ , and then one, two or even three slots numbered with zero. There are no other slots. A ball is thrown at the wheel when it is spinning, and it bounces around and eventually falls into a slot. If the wheel is properly balanced, the ball has the same probability of falling into each slot. The number of the slot the ball falls into is said to “come up”. There are a variety of bets available.

**Worked example 5.13** *Red or Black?*

On a roulette wheel, 18 of the numbers  $1 \dots 36$  are red, and the others are black. The zero (or zeros) are colorless. You can bet on (among other things) whether a red number or a black number comes up. If you bet 1 on red, and a red number comes up, you keep your stake and get 1; if a black number or a zero comes up, you get  $-1$  (i.e. the house keeps your bet).

- On a wheel with one zero, what is the expected value of a 1 bet on red?
- On a wheel with two zeros, what is the expected value of a 1 bet on red?
- On a wheel with three zeros, what is the expected value of a 1 bet on red?

**Solution:** Write  $p_r$  for the probability a red number comes up. The expected value is  $1 \times p_r + (-1)(1 - p_r)$  which is  $2p_r - 1$ .

- In this case,  $p_r = (\text{number of red numbers})/(\text{total number of numbers}) = 18/37$ . So the expected value is  $-1/37$  (you lose about three cents each time you bet).
- In this case,  $p_r = 18/38$ . So the expected value is  $-2/38 = -1/19$  (you lose slightly more than five cents each time you bet).
- In this case,  $p_r = 18/39$ . So the expected value is  $-3/39 = -1/13$  (you lose slightly less than eight cents each time you bet).

Notice that in the roulette game, the money you lose will go to the house. So the expected value to the house is just the negative of the expected value to you. You might not play the wheel often, but the house plays the wheel very often when there are many players. The weak law means a house with many players can rely on receiving three, five, or eight cents per bet. This is a partial explanation of why there are lots of roulette wheels, and usually free food nearby. Not all bets are like this, though.

**Worked example 5.14** *Coin game*

In this game, P1 flips a fair coin and P2 calls “H” or “T”. If P2 calls right, then P1 throws the coin into the river; otherwise, P1 keeps the coin. What is the expected value of this game to P2? and to P1?

**Solution:** To P2, which we do first, because it’s easiest: P2 gets 0 if P2 calls right, and 0 if P2 calls wrong; these are the only cases, so the expected value is 0. To P1: P1 gets  $-1$  if P2 calls right, and 0 if P1 calls wrong. The coin is fair, so the probability P2 calls right is  $1/2$ . The expected value is  $-1/2$ . While I can’t explain why people would play such a game, I’ve actually seen this done.

We call a bet **fair** when its expected value is zero. Taking a bet with a negative expected value is unwise, because, on average, you will lose money. Worse, the more times you play, the more you lose. Similarly, repeatedly taking a bet with a positive expected value is reliably profitable. However, you do need to be careful you computed the expected value right.

**Worked example 5.15** *Birthdays in succession*

P1 and P2 agree to the following bet. P1 gives P2 a stake of 1. If three people, stopped at random on the street, have birthdays in succession (i.e. Mon-Tue-Wed, and so on), then P2 gives P1 100. Otherwise, P1 loses the stake. What is the expected value of this bet to P1?

**Solution:** Write  $p$  for the probability of winning. Then the expected value is  $p \times 100 - (1 - p) \times 1$ . We computed  $p$  in example 13.1 (it was  $1/49$ ). So the bet is worth  $(52/49)$ , or slightly more than a dollar, to P1. P1 should be happy to agree to this as often as possible.

The reason P2 agrees to bets like that of example 5.15 is most likely that P2 can’t compute the probability exactly. P2 thinks the event is quite unlikely, so the expected value is negative; but it isn’t as unlikely as P2 thought it was, and this is how P1 makes a profit. This is one of the reasons you should be careful accepting a bet from a stranger: they might be able to compute better than you.

## 5.4.2 Odds, Expectations and Bookmaking — a Cultural Diversion

Gamblers sometimes use a terminology that is a bit different from ours. In particular, the term **odds** is important. The term comes from the following idea: P1 pays a bookmaker  $b$  (the stake) to make a bet; if the bet is successful, P1 receives  $a$  and the stake back, and if not, loses the original stake. This bet is referred to as odds of  $a : b$  (read “odds of  $a$  to  $b$ ”).

Assume the bet is fair, so that the expected value is zero. Write  $p$  for the probability of winning. The net income to P1 is  $ap - b(1 - p)$ . If this is zero, then

$p = b/(a + b)$ . So you can interpret odds in terms of probability, *if* you assume the bet is fair.

A bookmaker sets odds at which to accept bets from gamblers. The bookmaker does not wish to lose money at this business, and so must set odds which are potentially profitable. Doing so is not simple (bookmakers can, and occasionally do, lose catastrophically, and go out of business). In the simplest case, assume that the bookmaker knows the probability  $p$  that a particular bet will win. Then the bookmaker could set odds of  $(1 - p)/p : 1$ . In this case, the expected value of the bet is zero; this is fair, but not attractive business, so the bookmaker will set odds assuming that the probability is a bit higher than it really is. There are other bookmakers out there, so there is some reason for the bookmaker to try to set odds that are close to fair.

In some cases, you can tell when you are dealing with a bookmaker who is likely to go out of business soon. For example, imagine there are two horses running in a race, both at 10 : 1 odds — whatever happens, you could win by betting 1 on each. There is a more general version of this phenomenon. Assume the bet is placed on a horse race, and that bets pay off only for the winning horse. Assume also that exactly one horse will win (i.e. the race is never scratched, there aren't any ties, etc.), and write the probability that the  $i$ 'th horse will win as  $p_i$ . Then

$$\sum_{i \in \text{horses}} p_i$$

must be 1. Now if the bookmaker's odds yield a set of probabilities that is less than 1, their business should fail, because there is at least one horse on which they are paying out too much. Bookmakers deal with this possibility by writing odds so that  $\sum_{i \in \text{horses}} p_i$  is larger than one.

But this is not the only problem a bookmaker must deal with. The bookmaker doesn't actually know the probability that a particular horse will win, and must account for errors in this estimate. One way to do so is to collect as much information as possible (talk to grooms, jockeys, etc.). Another is to look at the pattern of bets that have been placed already. If the bookmaker and the gamblers agree on the probability that each horse will win, then there should be no expected advantage to choosing one horse over another — each should pay out slightly less than zero to the gambler (otherwise the bookmaker doesn't eat). But if the bookmaker has underestimated the probability that a particular horse will win, a gambler may get a positive expected payout by betting on that horse. This means that if one particular horse attracts a lot of money from bettors, it is wise for the bookmaker to offer less generous odds on that horse. There are two reasons: first, the bettors might know something the bookmaker doesn't, and they're signalling it; second, if the bets on this horse are very large and it wins, the bookmaker may not have enough capital left to pay out or to stay in business. All this means that real bookmaking is a complex, skilled business.

#### 5.4.3 Ending a Game Early

Imagine two people are playing a game for a stake, but must stop early — who should get what percentage of the stake? One way to do this is to give each player

what they put in at the start, but this is (mildly) unfair if one has an advantage over the other. The alternative is to give each player the expected value of the game at that state for that player. Sometimes one can compute that expectation quite easily.

**Worked example 5.16** *Ending a game early*

(from Durrett), two players each pay 25 to play the following game. They toss a fair coin. If it comes up heads, player H wins that toss; if tails, player T wins. The first player to reach 10 wins takes the stake of 50. But one player is called away when the state is 8-7 (H-T) — how should the stake be divided?

**Solution:** In this state, each player can either win — and so get 50 — or lose — and so get 0. The expectation for H is  $50P(\{\text{H wins from 8-7}\}) + 0P(\{\text{T wins from 8-7}\})$ , so we need to compute  $P(\{\text{H wins from 8-7}\})$ . Similarly, the expectation for T is  $50P(\{\text{T wins from 8-7}\}) + 0P(\{\text{H wins from 8-7}\})$ , so we need to compute  $P(\{\text{T wins from 8-7}\})$ ; but  $P(\{\text{T wins from 8-7}\}) = 1 - P(\{\text{H wins from 8-7}\})$ . Now it is slightly easier to compute  $P(\{\text{T wins from 8-7}\})$ , because T can only win in two ways: 8-10 or 9-10. These are independent. For T to win 8-10, the next three flips must come up T, so that event has probability  $1/8$ . For T to win 9-10, the next four flips must have one H in them, but the last flip may not be H (or else H wins); so the next four flips could be H T T T, T H T T, or T T H T. The probability of this is  $3/16$ . This means the total probability that T wins is  $5/16$ . So T should get 16.625 and H should get the rest (although they might have to flip for the odd half cent).

#### 5.4.4 Making a Decision with Decision Trees and Expectations

Imagine we have to choose an action. Once we have chosen, a sequence of random events occurs, and we get a reward with some probability. Which action should we choose? A good answer is to choose the action with the best expected outcome. In fact, choosing any other action is unwise, because if we encounter this situation repeatedly and make a choice that is even only slightly worse than the best, we could lose heavily. This is a very common recipe, and it can be applied to many situations. Usually, but not always, the reward is in money, and we will compute with money rewards for the first few examples.

For such problems, it can be useful to draw a **decision tree**. A decision tree is a drawing of possible outcomes of decisions, which makes costs, benefits and random elements explicit. Each node of the tree represents a test of an attribute (which could be either a decision, or a random variable), and each edge represents a possible outcome of a test. The final outcomes are leaves. Usually, decision nodes are drawn as squares, chance elements as circles, and leaves as triangles.

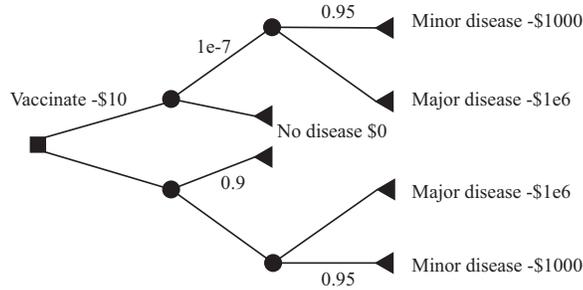


FIGURE 5.1: A decision tree for the vaccination problem. The only decision is whether to vaccinate or not (the box at the root of the tree). I have only labelled edges where this is essential, so I did not annotate the “no vaccination” edge with zero cost. Once you decide whether to vaccinate or not, there is a random node (whether you get the disease or not), and, if you get it, another (minor or major).

**Worked example 5.17** *Vaccination*

It costs 10 to be vaccinated against a common disease. If you have the vaccination, the probability you will get the disease is  $1e - 7$ . If you do not, the probability is 0.1. The disease is unpleasant; with probability 0.95, you will experience effects that cost you 1000 (eg several days in bed), but with probability 0.05, you will experience effects that cost you  $1e6$ . Should you be vaccinated?

**Solution:** Figure 5.1 shows a decision tree for this problem. I have annotated some edges with the choices represented, and some edges with probabilities; the sum of probabilities over all rightward (downgoing) edges leaving a random node is 1. It is straightforward to compute expectations. The expected cost of the disease is  $0.95 \times 1000 + 0.05 \times 1e6 = 50,950$ . If you are vaccinated, your expected income will be  $-(10 + 1e - 7 \times 50,950) = -10.01$  (rounding to the nearest cent). If you are not, your expected income is  $-5,095$ . You should be vaccinated.

Sometimes there is more than one decision. We can still do simple examples, though drawing a decision tree is now quite important, because it allows us to keep track of cases and avoid missing anything. For example, assume I wish to buy a cupboard. Two nearby towns have used furniture shops (usually called antique shops these days). One is further away than the other. If I go to town A, I will have time to look in two (of three) shops; if I go to town B, I will have time to look in one (of two) shops. I could lay out this sequence of decisions (which town to go to; which shop to visit when I get there) as Figure 5.2.

You should notice that this figure is missing a lot of information. What is the probability that I will find what I’m looking for in the shops? What is the value of finding it? What is the cost of going to each town? and so on. This information is

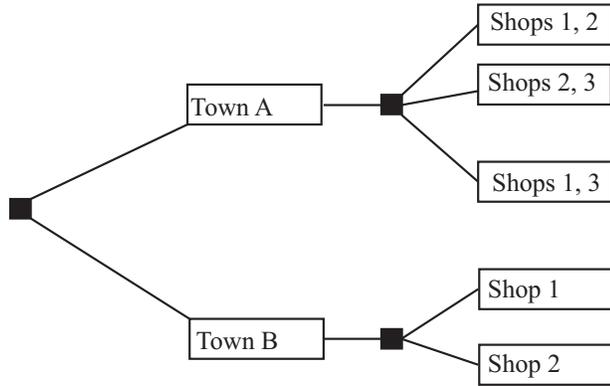


FIGURE 5.2: *The decision tree for the example of visiting furniture shops. Town A is nearer than town B, so if I go there I can choose to visit two of the three shops there; if I go to town B, I can visit only one of the two shops there. To decide what to do, I could fill in the probabilities and values of outcomes, compute the expected value of each pair of decisions, and choose the best. This could be tricky to do (where do I get the probabilities from?) but offers a rational and principled way to make the decision.*

not always easy to obtain. In fact, I might simply need to give my best subjective guess of these numbers. Furthermore, particularly if there are several decisions, computing the expected value of each possible sequence could get difficult. There are some kinds of model where one can compute expected values easily, but a good viable hypothesis about why people don't make optimal decisions is that optimal decisions are actually too hard to compute.

### 5.4.5 Utility

Sometimes it is hard to work with money. For example, in the case of a serious disease, choosing treatments often boils down to expected survival times, rather than money.

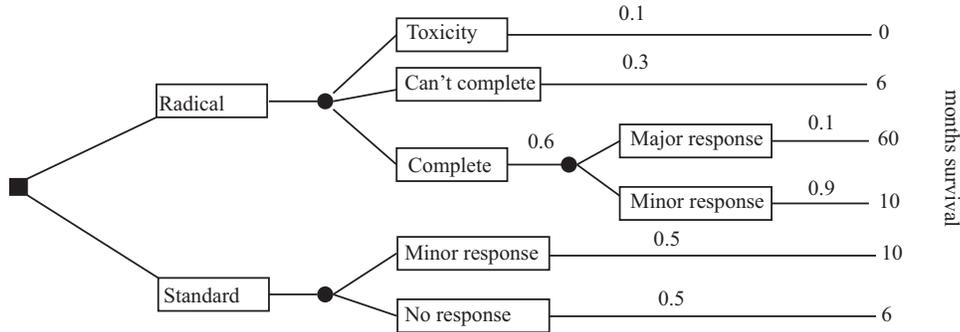


FIGURE 5.3: A decision tree for example 5.18. Notations vary a bit, and here I have put boxes around the labels for the edges.

**Worked example 5.18** *Radical treatment*

(This example largely after Vickers, p97). Imagine you have a nasty disease. There are two kinds of treatment: standard, and radical. Radical treatment might kill you (with probability 0.1); might be so damaging that doctors stop (with probability 0.3); but otherwise you will complete the treatment. If you do complete radical treatment, there could be a major response (probability 0.1) or a minor response. If you follow standard treatment, there could be a major response (probability 0.5) or a minor response, but the outcomes are less good. All this is best summarized in a decision tree (Figure 5.3). What gives the longest expected survival time?

**Solution:** In this case, expected survival time with radical treatment is  $(0.1 \times 0 + 0.3 \times 6 + 0.6 \times (0.1 \times 60 + 0.9 \times 10)) = 10.8$  months; expected survival time without radical treatment is  $0.5 \times 10 + 0.5 \times 6 = 8$  months.

Working with money values is not always a good idea. For example, many people play state lotteries. The expected value of a 1 bet on a state lottery is well below 1 — why do people play? It’s easy to assume that all players just can’t do sums, but many players are well aware that the expected value of a bet is below the cost. It seems to be the case that people value money in a way that doesn’t depend linearly on the amount of money. So, for example, people may value a million dollars rather more than a million times the value they place on one dollar. If this is true, we need some other way to keep track of value; this is sometimes called **utility**. It turns out to be quite hard to know how people value things, and there is quite good evidence that (a) human utility is complicated and (b) it is difficult to explain human decision making in terms of expected utility.

**Worked example 5.19** *Human utility is not expected payoff*

Here are four games:

- **Game 1:** The player is given 1. A biased coin is flipped, and the money is taken back with probability  $p$ ; otherwise, the player keeps it.
- **Game 2:** The player stakes 1, and a fair coin is flipped; if the coin comes up heads, the player gets  $r$  and the stake back, but otherwise loses the original stake.
- **Game 3:** The player bets nothing; a biased coin is flipped, and if it comes up heads (probability  $q$ ), the player gets  $1e6$ .
- **Game 4:** The player stakes 1000; a fair coin is flipped, and if it comes up heads, the player gets  $s$  and the stake back, but otherwise loses the original stake.

In particular, what happens if  $r = 3 - 2p$  and  $q = (1 - p)/1e6$  and  $s = 2 - 2p + 1000$ ?

**Solution:** Game 1 has expected value  $(1 - p)1$ . Game 2 has expected value  $(1/2)(r - 1)$ . Game 3 has expected value  $q1e6$ . Game 4 has expected value  $(1/2)s - 500$ .

In the case given, each game has the same expected value. Nonetheless, people usually have decided preferences for which game they would play. Generally, 4 is unattractive (seems expensive to play); 3 seems like free money, and so a good thing; 2 might be OK but is often seen as uninteresting; and 1 is unattractive. This should suggest to you that people's reasoning about money and utility is not what simple expectations can predict.

5.5 YOU SHOULD

5.5.1 be able to:

- Interpret notation for joint and conditional probability for random variables; in particular, understand notation such as:  $P(\{X\})$ ,  $P(\{X = x\})$ ,  $p(x)$ ,  $p(x, y)$ ,  $p(x|y)$
- Interpret a probability density function  $p(x)$  as  $P(\{X \in [x, x + dx]\})$ .
- Interpret the expected value of a discrete random variable.
- Interpret the expected value of a continuous random variable.
- Compute expected values of random variables for straightforward cases.
- Write down expressions for mean, variance and covariance for random variables.
- Write out a decision tree.

5.5.2 remember:

Definition: Discrete random variable . . . . .	122
Definition: Probability distribution of a discrete random variable . . . . .	124
Definition: Cumulative distribution of a discrete random variable . . . . .	124
New term: joint probability distribution . . . . .	125
New term: Bayes' rule . . . . .	125
Definition: Bayes' rule . . . . .	126
New term: marginal probability . . . . .	126
New term: independent . . . . .	126
Definition: Independent random variables . . . . .	127
New term: probability density function . . . . .	129
Useful facts: Probability density functions . . . . .	129
New term: normalizing . . . . .	130
New term: normalizing constant . . . . .	130
New term: expected value . . . . .	131
Definition: Expected value . . . . .	131
Definition: Expectation . . . . .	132
Definition: Expected value . . . . .	133
Definition: Expectation . . . . .	134
Useful facts: Expectations are linear. . . . .	134
Definition: Mean or expected value . . . . .	134
Definition: Variance . . . . .	135
Useful facts: Variance . . . . .	135
Definition: Covariance . . . . .	135
Useful facts: A useful expression for variance . . . . .	136
Useful facts: A useful expression for covariance . . . . .	137
Useful facts: Variance and Covariance . . . . .	137
New term: standard deviation . . . . .	138
Definition: Standard deviation . . . . .	139

New term: descriptive statistics . . . . . 139  
 New term: empirical distribution . . . . . 139  
 Definition: Markov’s inequality . . . . . 140  
 Definition: Chebyshev’s inequality . . . . . 140  
 New term: indicator function . . . . . 141  
 Definition: Indicator functions . . . . . 141  
 New term: trial . . . . . 143  
 New term: sample . . . . . 143  
 New term: independent identically distributed samples . . . . . 143  
 New term: iid samples . . . . . 143  
 New term: samples . . . . . 143  
 Definition: Weak Law of Large Numbers . . . . . 144  
 New term: odds . . . . . 148  
 New term: decision tree . . . . . 150  
 New term: utility . . . . . 153

## PROBLEMS

## Joint and Conditional Probability for Random Variables

- 5.1. A roulette wheel has one zero. Write  $X$  for the random variable representing the number that will come up on the wheel. What is the probability distribution of  $X$ ?
- 5.2. Define a random variable  $X$  by the following procedure. Draw a card from a standard deck of playing cards. If the card is knave, queen, or king, then  $X = 11$ . If the card is an ace, then  $X = 1$ ; otherwise,  $X$  is the number of the card (i.e. two through ten). Now define a second random variable  $Y$  by the following procedure. When you evaluate  $X$ , you look at the color of the card. If the card is red, then  $Y = X - 1$ ; otherwise,  $Y = X + 1$ .
- What is  $P(\{X \leq 2\})$ ?
  - What is  $P(\{X \geq 10\})$ ?
  - What is  $P(\{X \geq Y\})$ ?
  - What is the probability distribution of  $Y - X$ ?
  - What is  $P(\{Y \geq 12\})$ ?
- 5.3. Define a random variable by the following procedure. Flip a fair coin. If it comes up heads, the value is 1. If it comes up tails, roll a die: if the outcome is 2 or 3, the value of the random variable is 2. Otherwise, the value is 3.
- What is the probability distribution of this random variable?
  - What is the cumulative distribution of this random variable?
- 5.4. Define three random variables,  $X$ ,  $Y$  and  $Z$  by the following procedure. Roll a six-sided die and a four-sided die. Now flip a coin. If the coin comes up heads, then  $X$  takes the value of the six-sided die and  $Y$  takes the value of the four-sided die. Otherwise,  $X$  takes the value of the four-sided die and  $Y$  takes the value of the six-sided die.  $Z$  always takes the value of the sum of the dice.
- What is  $P(X)$ , the probability distribution of this random variable?
  - What is  $P(X, Y)$ , the joint probability distribution of these two random variables?
  - Are  $X$  and  $Y$  independent?
  - Are  $X$  and  $Z$  independent?
- 5.5. Define two random variables  $X$  and  $Y$  by the following procedure. Flip a fair coin; if it comes up heads, then  $X = 1$ , otherwise  $X = -1$ . Now roll a six-sided die, and call the value  $U$ . We define  $Y = U + X$ .
- What is  $P(Y|X = 1)$ ?
  - What is  $P(X|Y = 0)$ ?
  - What is  $P(X|Y = 7)$ ?
  - What is  $P(X|Y = 3)$ ?
  - Are  $X$  and  $Y$  independent?
- 5.6. Magic the Gathering is a popular card game. Cards can be land cards, or other cards. We consider a game with two players. Each player has a deck of 40 cards. Each player shuffles their deck, then deals seven cards, called their *hand*. The rest of each player's deck is called their *library*. Assume that player one has 10 land cards in their deck and player two has 20. Write  $L_1$  for the number of lands in player one's hand and  $L_2$  for the number of lands in player two's hand. Write  $L_t$  for the number of lands in the top 10 cards of player one's library.
- Write  $S = L_1 + L_2$ . What is  $P(\{S = 0\})$ ?

- (b) Write  $D = L_1 - L_2$ . What is  $P(\{D = 0\})$ ?
- (c) What is the probability distribution for  $L_1$ ?
- (d) Write out the probability distribution for  $P(L_1|L_t = 10)$ .
- (e) Write out the probability distribution  $P(L_1|L_t = 5)$ .

### Continuous Random Variables

- 5.7. A continuous random variable has probability density function  $p(x)$  which is proportional to  $g(x)$ , where

$$g(x) = \begin{cases} 0 & \text{if } x < -\frac{\pi}{2} \\ 0 & \text{if } x > \frac{\pi}{2} \\ \cos(x) & \text{otherwise} \end{cases} .$$

Write  $c$  for the constant of proportionality, so that  $p(x) = cg(x)$ .

- (a) What is  $c$ ? (you can look up the integral if you want)
  - (b) What is  $P(\{X \geq 0\})$  (i.e. the probability you will observe a value greater than 0)? (you can look up the integral if you want)
  - (c) What is  $P(\{|X| \leq 1\})$ ? (you can look up the integral if you want)
- 5.8. There is some (small!) voltage over the terminals of a warm resistor caused by noise (electrons moving around in the heat and banging into one another). This is a good example of a continuous random variable, and we can assume there is some probability density function for it, say  $p(x)$ . We assume that  $p(x)$  has the property that

$$\lim_{\epsilon \rightarrow 0} \int_{v-\epsilon}^{v+\epsilon} p(x) dx = 0$$

which is what you'd expect for any function you're likely to have dealt with. Now imagine I define a new random variable by the following procedure: I flip a coin; if it comes up heads, I report 0; if tails, I report the voltage over the resistor. This random variable,  $u$ , has a probability 1/2 of taking the value 0, and 1/2 of taking a value from  $p(x)$ . Write this random variable's probability density function  $q(u)$ . Show that

$$\lim_{\epsilon \rightarrow 0} \int_{-\epsilon}^{\epsilon} q(u) du = \frac{1}{2}$$

### Expected Values

- 5.9. Magic the Gathering is a popular card game. Cards can be land cards, or other cards. We consider a game with two players. Each player has a deck of 40 cards. Each player shuffles their deck, then deals seven cards, called their *hand*. The rest of each player's deck is called their *library*. Assume that player one has 10 land cards in their deck and player two has 20. Write  $L_1$  for the number of lands in player one's hand and  $L_2$  for the number of lands in player two's hand. Write  $L_t$  for the number of lands in the top 10 cards of player one's library.
- (a) What is  $\mathbb{E}[L_1]$ ?
  - (b) What is  $\mathbb{E}[L_2]$ ?
  - (c) What is  $\text{var}[L_1]$ ?
- 5.10. A simple coin game is as follows: we have a box, which starts empty. P1 flips a fair coin. If it comes up heads, P2 gets the contents of the box, and the game

ends. If it comes up tails, P1 puts a dollar in the box and they flip again; this repeats until it comes up heads

- (a) With what probability will P2 win exactly 10 units?  
 (b) Write  $S_\infty = \sum_{i=0}^{\infty} r^i$ . Show that  $(1-r)S_\infty = 1$ , so that

$$S_\infty = \frac{1}{1-r}$$

- (c) Show that

$$\sum_{i=0}^{\infty} ir^i = \left(\sum_{i=1}^{\infty} r^i\right) + r\left(\sum_{i=1}^{\infty} r^i\right) + r^2\left(\sum_{i=1}^{\infty} r^i\right) + \dots$$

(look carefully at the limits of the sums!) and so show that

$$\sum_{i=0}^{\infty} ir^i = \frac{r}{(1-r)^2}.$$

- (d) What is the expected value of the game? (you may find the results of the two previous subexercises helpful; they're not there just for show).  
 (e) How much should P2 pay to play, to make the game fair?
- 5.11.** A simple card game is as follows. P1 pays a stake of 1 to play. P1 and P2 then each draw a card. If both cards are the same color, P2 keeps the stake and the game ends. If they are different colors, P2 pays P1 the stake and 1 extra (a total of 2).
- (a) What is the expected value of the game to P1?  
 (b) P2 modifies the game, as follows. If both cards are court cards (that is, knave, queen, king), then P2 keeps the stake and the game ends; otherwise, the game works as before. Now what is the expected value of the game to P1?
- 5.12.** An airline company runs a flight that has six seats. Each passenger who buys a ticket has a probability  $p$  of turning up for the flight. These events are independent.
- (a) The airline sells six tickets. What is the expected number of passengers, if  $p = 0.9$ ?  
 (b) How many tickets should the airline sell to ensure that the expected number of passengers is greater than six, if  $p = 0.7$ ? **Hint:** The easiest way to do this is to write a quick program that computes the expected value of passengers that turn up for each the number of tickets sold, then search the number of tickets sold.
- 5.13.** An airline company runs a flight that has 10 seats. Each passenger who buys a ticket has a probability  $p$  of turning up for the flight. The gender of the passengers is not known until they turn up for a flight, and women buy tickets with the same frequency that men do. The pilot is eccentric, and will not fly unless at least two women turn up.
- (a) How many tickets should the airline sell to ensure that the expected number of passengers that turn up is greater than 10?  
 (b) The airline sells 10 tickets. What is the expected number of passengers on the aircraft, given that it flies? (i.e. that at least two women turn up). Estimate this value with a simulation.

- 5.14.** A coin game that is occasionally played is odd person out. In this game, there are rounds. In a round, each person flips a coin. There is an odd person out in that round if all but one have H and the other has T, OR all but one have T and the last has H.
- (a) Three people play one round. What is the probability that there is an odd person out?
  - (b) Now four people play one round. What is the probability that there is an odd person out?
  - (c) Five people play until there is an odd person out. What is the expected number of rounds that they will play? (you can save yourself quite a lot of calculation by reading section ??, if you don't mind skipping ahead a bit).

#### Mean, Variance and Covariance

- 5.15.** Show that  $\text{var}[kX] = k^2\text{var}[X]$ .
- 5.16.** Show that if  $X$  and  $Y$  are independent random variables, then  $\text{var}[X + Y] = \text{var}[X] + \text{var}[Y]$ . You will find it helpful to remember that, for  $X$  and  $Y$  independent,  $\mathbb{E}[XY] = \mathbb{E}[X]\mathbb{E}[Y]$ .

#### Markov and Chebyshev Inequalities

- 5.17.** The random variable  $X$  takes the values  $-2, -1, 0, 1, 2$ , but has an unknown probability distribution. You know that  $\mathbb{E}[\|X\|] = 0.2$ . Use Markov's inequality to give a *lower* bound on  $P(\{X = 0\})$ . *Hint:* Notice that  $P(\{X = 0\}) = 1 - P(\{\|X\| = 1\}) - P(\{\|X\| = 2\})$ .
- 5.18.** The random variable  $X$  takes the values  $1, 2, 3, 4, 5$ , but has unknown probability distribution. You know that  $\mathbb{E}[X] = 2$  and  $\text{var}(\{X\}) = 0.01$ . Use Chebyshev's inequality to give a *lower* bound on  $P(\{X = 2\})$ .
- 5.19.** You have a biased random number generator. This generator produces a random number with mean value  $-1$ , and standard deviation  $0.5$ . Write  $\mathcal{A}$  for the event that the number generator produces a non-negative number. Use Chebyshev's inequality to bound  $P(\mathcal{A})$ .
- 5.20.** You observe a random number generator. You know that it can produce the values  $-2, -1, 0, 1, 2$ . You are told that it has been adjusted so that: (1) the mean value it produces is zero and; (2) the standard deviation of the numbers it produces is  $1$ .
- (a) Write  $\mathcal{A}$  for the event that the number generator produces a number that is not  $0$ . Use Chebyshev's inequality to bound  $P(\mathcal{A})$ .
  - (b) Write  $\mathcal{B}$  for the event that the number generator produces  $-2$  or  $2$ . Use Chebyshev's inequality to bound  $P(\mathcal{B})$ .

#### Using Expectations

- 5.21.** Imagine we have a game with two players, who are playing for a stake. There are no draws, the winner gets the whole stake, and the loser gets nothing. The game must end early. We decide to give each player the expected value of the game for that player, from that state. Show that the expected values add up to the value of the stake (i.e. there won't be too little or too much money in the stake).

## Useful Probability Distributions

We will use probability as a tool to resolve practical questions about data. I describe some forms these questions take below, for concreteness. Generally, resolving these questions requires some form of model. This model gives an abstract representation of the problem that is useful for problem solving, and (typically) comes with recipes for attacking the major types of problem.

We could ask **what process produced the data?** For example, I observe a set of independent coin flips. I would now like to know the probability of observing a head when the coin is flipped. This might seem a bit empty as a problem, but an analogous problem is: are male children or female children more frequent? Notice that the answers to these questions are typically not exact. Instead, they are estimates. We will see a variety of methods to estimate the probability of observing a head from a sequence of independent coin flips, but none is guaranteed to return the “right” answer (because you can’t). Instead, we have to live with information about how accurate the estimate is.

We could ask **what sort of data can we expect in the future?** For example, we could ask: is gender assigned independently? equivalently, can you predict the gender of the next child born to a couple more accurately if you look at the genders of the previous children? In reliability engineering, one asks: how long will it be until this product fails? One version of this question that occupies many people is: how long until I die? By the way, the best answer seems to be subtract your age from a number that seems to be close to 85. Again, these are questions which don’t lend themselves to the “right” answer, as opposed to the best possible estimate. You might get hit by a truck tomorrow.

We could ask **what labels should we attach to unlabelled data?** For example, we might see a large number of credit card transactions, some known to be legitimate and others known to be fraudulent. We now see a new transaction: is it legitimate? You can see that versions of this question appear in many applications. As another example, you see many programs downloaded from the web, some known to be legitimate and others known to be malware. You now see a new program: is it safe to run? It may be possible in some circumstances to know the “right” answer to this question, but we are usually stuck with the best answer.

We could ask **is an effect easily explained by chance variations, or is it real?** For example, you believe that the average weight of a mouse is 15 grams. You could test this by catching 100 mice (a sample) and weighing them; but, when you do this, the answer won’t be 15 grams *even if* the average weight of a mouse is 15 grams. This is because you have a random selection of mice. But we will be able to build a model of the random variations in the sample average. We can use this model to tell whether the difference between sample average and 15 grams is easily explained by random variations in the sample, or is significant. This allows

us to tell how forcefully the evidence contradicts your original belief.

Building a model requires understanding the specific problem you want to solve, then choosing from a vocabulary of many different models that might apply to the problem. Long experience shows that even if the model you choose does not match the problem exactly, it can still be useful. In this chapter, I describe the properties of some probability distributions that are used again and again in model building.

## 6.1 DISCRETE DISTRIBUTIONS

### 6.1.1 The Discrete Uniform Distribution

If every value of a discrete random variable has the same probability, then the probability distribution is the discrete uniform distribution. We have seen this distribution before, numerous times. For example, I define a random variable by the number that shows face-up on the throw of a fair die. This has a uniform distribution. As another example, write the numbers 1-52 on the face of each card of a standard deck of playing cards. The number on the face of the first card drawn from a well-shuffled deck is a random variable with a uniform distribution.

One can construct expressions for the mean and variance of a discrete uniform distribution, but they're not usually much use (too many terms, not often used). Keep in mind that if two random variables have a uniform distribution, their sum and difference will not (recall example 5.3).

### 6.1.2 Bernoulli Random Variables

A Bernoulli random variable models a biased coin with probability  $p$  of coming up heads in any one flip.

**Definition: 6.1** *Bernoulli random variable*

A Bernoulli random variable takes the value 1 with probability  $p$  and 0 with probability  $1 - p$ . This is a model for a coin toss, among other things

Indexing: Bernoulli random variable

Indexing: Bernoulli random variable

**Useful Facts: 6.1** *Bernoulli random variable*

1. A Bernoulli random variable has mean  $p$ .
2. A Bernoulli random variable has variance  $p(1 - p)$ .

Proofs are easy, and in the exercises.

Indexing: Bernoulli random variable

Indexing: Bernoulli random variable

### 6.1.3 The Geometric Distribution

We have a biased coin. The probability it will land heads up,  $P(\{H\})$  is given by  $p$ . We flip this coin until the first head appears. The number of flips required is a discrete random variable which takes integer values greater than or equal to one, which we shall call  $X$ . To get  $n$  flips, we must have  $n - 1$  tails followed by 1 head. This event has probability  $(1 - p)^{(n-1)}p$ . We can now write out the probability distribution that  $n$  flips are required.

**Definition: 6.2** *Geometric distribution*

We have an experiment with a binary outcome (i.e. heads or tails; 0 or 1; and so on), with  $P(H) = p$  and  $P(T) = 1 - p$ . We repeat this experiment until the first head occurs. The probability distribution for  $n$ , the number of repetitions, is the geometric distribution. It has the form

$$P(\{X = n\}) = (1 - p)^{(n-1)}p.$$

for  $0 \leq p \leq 1$  and  $n \geq 1$ ; for other  $n$  it is zero.  $p$  is called the **parameter** of the distribution.

Indexing: Geometric distribution

Indexing: Geometric distribution

Notice that the geometric distribution is non-negative everywhere. It is straightforward to show that it sums to one, and so is a probability distribution (exercises).

**Useful Facts: 6.2** *Geometric distribution*

1. The mean of the geometric distribution is  $\frac{1}{p}$ .
2. The variance of the geometric distribution is  $\frac{1-p}{p^2}$ .

The proof of these facts requires some work with series, and is relegated to the exercises.

Indexing: Geometric distribution

Indexing: Geometric distribution

It should be clear that this model isn't really about coins, but about repeated trials. The trial could be anything that has some probability of failing. Each trial is independent, and the rule for repeating is that you keep trying until the first success. Textbooks often set exercises involving missiles and aircraft; I'll omit these on the grounds of taste.

## 6.1.4 The Binomial Probability Distribution

Assume we have a biased coin with probability  $p$  of coming up heads in any one flip. The binomial probability distribution gives the probability that it comes up heads  $h$  times in  $N$  flips.

Worked example 4.30 yields one way of deriving this distribution. In that example, I showed that there are

$$N!/(h!(N-h)!)^2$$

outcomes of  $N$  coin flips that have  $h$  heads. These outcomes are disjoint, and each has probability  $p^h(1-p)^{(N-h)}$ . As a result, we must have the probability distribution below.

**Definition: 6.3** *Binomial distribution*

In  $N$  independent repetitions of an experiment with a binary outcome (ie heads or tails; 0 or 1; and so on) with  $P(H) = p$  and  $P(T) = 1 - p$ , the probability of observing a total of  $h$   $H$ 's and  $N - h$   $T$ 's is

$$P_b(h; N, p) = \binom{N}{h} p^h (1-p)^{(N-h)}$$

as long as  $0 \leq h \leq N$ ; in any other case, the probability is zero.

Indexing: Binomial distribution

Indexing: Binomial distribution

**Useful Fact: 6.3** *Binomial distribution*

Write  $P_b(i; N, p)$  for the binomial distribution that one observes  $i$   $H$ 's in  $N$  trials.

$$\sum_{i=0}^N P_b(i; N, p) = (p + (1-p))^N = (1)^N = 1$$

by pattern matching to the binomial theorem. As a result,

$$\sum_{i=0}^N P_b(i; N, p) = 1$$

Indexing: Binomial distribution

Indexing: Binomial distribution

The binomial distribution satisfies a recurrence relation. We must have that

$$P_b(h; N, p) = pP_b(h - 1; N - 1, p) + (1 - p)P_b(h; N - 1, p).$$

This is because can get  $h$  heads in  $N$  flips either by having  $h - 1$  heads in  $N - 1$  flips, then flipping another, or by having  $h$  heads in  $N$  flips then flipping a tail. You can verify by induction that the binomial distribution satisfies this recurrence relation.

**Useful Facts: 6.4** *Binomial distribution*

1. The mean of  $P_b(i; N, p)$  is  $Np$ .
2. The variance of  $P_b(i; N, p)$  is  $Np(1 - p)$

The proofs are informative, and so are not banished to the exercises.

Indexing: Binomial distribution

Indexing: Binomial distribution

**Proof: 6.1** *The binomial distribution*

Notice that the number of heads in  $N$  coin tosses is can be obtained by adding the number of heads in each toss. Write  $Y_i$  for the Bernoulli random variable representing the  $i$ 'th toss. If the coin comes up heads,  $Y_i = 1$ , otherwise  $Y_i = 0$ . The  $Y_i$  are independent. Now

$$\begin{aligned}\mathbb{E}[X] &= \mathbb{E}\left[\sum_{j=1}^N Y_j\right] \\ &= \sum_{j=1}^N \mathbb{E}[Y_j] \\ &= N\mathbb{E}[Y_1] && \text{because the } Y_i \text{ are independent} \\ &= Np.\end{aligned}$$

The variance is easy, too. Each coin toss is independent, so the variance of the sum of coin tosses is the sum of the variances. This gives

$$\begin{aligned}\text{var}[X] &= \text{var}\left[\sum_{j=1}^N Y_j\right] \\ &= N\text{var}[Y_1] \\ &= Np(1-p)\end{aligned}$$

## 6.1.5 Multinomial probabilities

The binomial distribution describes what happens when a coin is flipped multiple times. But we could toss a die multiple times too. Assume this die has  $k$  sides, and we toss it  $N$  times. The distribution of outcomes is known as the **multinomial distribution**.

We can guess the form of the multinomial distribution in rather a straightforward way. The die has  $k$  sides. We toss the die  $N$  times. This gives us a sequence of  $N$  numbers. Each toss of the die is independent. Assume that side 1 appears  $n_1$  times, side 2 appears  $n_2$  times, ... side  $k$  appears  $n_k$  times. Any single sequence with this property will appear with probability  $p_1^{n_1} p_2^{n_2} \dots p_k^{n_k}$ , because the tosses are independent. However, there are

$$\frac{N!}{n_1! n_2! \dots n_k!}$$

such sequences. Using this reasoning, we arrive at the distribution below

**Definition: 6.4** *Multinomial distribution*

I perform  $N$  independent repetitions of an experiment with  $k$  possible outcomes. The  $i$ 'th such outcome has probability  $p_i$ . I see outcome 1  $n_1$  times, outcome 2  $n_2$  times, etc. Notice that  $n_1 + n_2 + n_3 + \dots + n_k = N$ . The probability of observing this set of outcomes is

$$P_m(n_1, \dots, n_k; N, p_1, \dots, p_k) = \frac{N!}{n_1! n_2! \dots n_k!} p_1^{n_1} p_2^{n_2} \dots p_k^{n_k}.$$

Indexing: Multinomial distribution

Indexing: Multinomial distribution

**Worked example 6.1** *Dice*

I throw five fair dice. What is the probability of getting two 2's and three 3's?

**Solution:**  $\frac{5!}{2!3!} (\frac{1}{6})^2 (\frac{1}{6})^3$

## 6.1.6 The Poisson Distribution

Assume we are interested in counts that occur in an interval of time (e.g. within a particular hour). Because they are counts, they are non-negative and integer valued. We know these counts have two important properties. First, they occur with some fixed average rate. Second, an observation occurs independent of the interval since the last observation. Then the Poisson distribution is an appropriate model.

There are numerous such cases. For example, the marketing phone calls you receive during the day time are likely to be well modelled by a Poisson distribution. They come at some average rate — perhaps 5 a day as I write, during the last phases of an election year — and the probability of getting one clearly doesn't depend on the time since the last one arrived. Classic examples include the number of Prussian soldiers killed by horse-kicks each year; the number of calls arriving at a call center each minute; the number of insurance claims occurring in a given time interval (outside of a special event like a hurricane, etc.).

**Definition: 6.5** *Poisson distribution*

A non-negative, integer valued random variable  $X$  has a Poisson distribution when its probability distribution takes the form

$$P(\{X = k\}) = \frac{\lambda^k e^{-\lambda}}{k!},$$

where  $\lambda > 0$  is a parameter often known as the **intensity** of the distribution.

Indexing: Poisson distribution

Indexing: Poisson distribution

Notice that the Poisson distribution is a probability distribution, because it is non-negative and because

$$\sum_{i=0}^{\infty} \frac{\lambda^i}{i!} = e^{\lambda}$$

so that

$$\sum_{k=0}^{\infty} \frac{\lambda^k e^{-\lambda}}{k!} = 1$$

**Useful Facts: 6.5** *Poisson distribution*

1. The mean of a Poisson distribution with intensity  $\lambda$  is  $\lambda$ .
2. The variance of a Poisson distribution with intensity  $\lambda$  is  $\lambda$  (no, that's not an accidentally repeated line or typo).

The proof of these facts requires some work with series, and is relegated to the exercises.

Indexing: Poisson distribution

Indexing: Poisson distribution

I described the Poisson distribution as a natural model for counts of randomly distributed points along a time axis. But it doesn't really matter that this is a time axis — it could be a space axis instead. For example, you could take a length of road, divide it into even intervals, then count the number of road-killed animals in each interval. If the location of each animal is independent of the location of any other animal, then you could expect a Poisson model to apply to the count data. Assume that the Poisson model that best describes the data has parameter  $\lambda$ . One property of such models is that if you doubled the length of the intervals, then the resulting dataset would be described by a Poisson model with parameter

$2\lambda$ ; similarly, if you halved the length of the intervals, the best model would have parameter  $\lambda/2$ . This corresponds to our intuition about such data; roughly, the number of road-killed animals in two miles of road should be twice the number in one mile of road. This property means that no pieces of the road are “special” — each behaves the same as the other.

We can build a really useful model of spatial randomness by observing this fact and generalizing very slightly. A **Poisson point process** with intensity  $\lambda$  is a set of random points with the property that the number of points in an interval of length  $s$  is a Poisson random variable with parameter  $\lambda s$ . Notice how this captures our intuition that if points are “very randomly” distributed, there should be twice as many of them in an interval that is twice as long.

This model is easily, and very usefully, extended to points on the plane, on surfaces, and in 3D. In each case, the process is defined on a domain  $D$  (which has to meet some very minor conditions that are of no interest to us). The number of points in any subset  $s$  of  $D$  is a Poisson random variable, with intensity  $\lambda m(s)$ , where  $m(s)$  is the area (resp. volume) of  $s$ . These models are useful, because they capture the property that (a) the points are random and (b) the probability you find a point doesn’t depend on where you are. You could reasonably believe models like this apply to, say, dead flies on windcreens; the places where you find acorns at the foot of an oak tree; the distribution of cowpats in a field; the distribution of cherries in a fruitcake; and so on.

## 6.2 CONTINUOUS DISTRIBUTIONS

### 6.2.1 The Continuous Uniform Distribution

Some continuous random variables have a natural upper bound and a natural lower bound but otherwise we know nothing about them. For example, imagine we are given a coin of unknown properties by someone who is known to be a skillful maker of unfair coins. The manufacturer makes no representations as to the behavior of the coin. The probability that this coin will come up heads is a random variable, about which we know nothing except that it has a lower bound of zero and an upper bound of one.

If we know nothing about a random variable apart from the fact that it has a lower and an upper bound, then a **uniform distribution** is a natural model. Write  $l$  for the lower bound and  $u$  for the upper bound. The probability density function for the uniform distribution is

$$p(x) = \begin{cases} 0 & x < l \\ 1/(u-l) & l \leq x \leq u \\ 0 & x > u \end{cases}$$

A continuous random variable whose probability distribution is the uniform distribution is often called a **uniform random variable**.

### 6.2.2 The Beta Distribution

It’s hard to explain now why the Beta (or  $\beta$ ) distribution is useful, but it will come in useful later (section 13.1). The Beta distribution is a probability distribution for

a continuous random variable  $x$  in the range  $0 \leq x \leq 1$ . There are two parameters,  $\alpha > 0$  and  $\beta > 0$ . Recall the definition of the  $\Gamma$  function from section 1.1. We have that

$$P_\beta(x|\alpha, \beta) = \frac{\Gamma(\alpha + \beta)}{\Gamma(\alpha)\Gamma(\beta)} x^{(\alpha-1)}(1-x)^{(\beta-1)}.$$

From this expression, you can see that:

- $P_\beta(x|1, 1)$  is a uniform distribution on the unit interval.
- $P_\beta(x|\alpha, \beta)$  has a single maximum at  $x = (\alpha - 1)/(\alpha + \beta - 2)$  for  $\alpha > 1, \beta > 1$  (differentiate and set to zero).
- Generally, as  $\alpha$  and  $\beta$  get larger, this peak gets narrower.
- For  $\alpha = 1, \beta > 1$  the largest value of  $P_\beta(x|\alpha, \beta)$  is at  $x = 0$ .
- For  $\alpha > 1, \beta = 1$  the largest value of  $P_\beta(x|\alpha, \beta)$  is at  $x = 1$ .

Figure 6.1 shows plots of the probability density function of the Beta distribution for a variety of different values of  $\alpha$  and  $\beta$ .

**Useful Facts: 6.6** *Beta distribution*

For a Beta distribution with parameters  $\alpha, \beta$

1. The mean is  $\frac{\alpha}{\alpha + \beta}$ .
2. The variance is  $\frac{\alpha\beta}{(\alpha + \beta)^2(\alpha + \beta + 1)}$ .

Indexing: Beta distribution

Indexing: Beta distribution

### 6.2.3 The Gamma Distribution

The Gamma (or  $\gamma$ ) distribution will also come in useful later on (section 13.1). The Gamma distribution is a probability distribution for a non-negative continuous random variable  $x \geq 0$ . There are two parameters,  $\alpha > 0$  and  $\beta > 0$ . The probability density function is

$$P_\gamma(x|\alpha, \beta) = \frac{\beta^\alpha}{\Gamma(\alpha)} x^{(\alpha-1)} e^{-\beta x}.$$

Figure 6.2 shows plots of the probability density function of the Gamma distribution for a variety of different values of  $\alpha$  and  $\beta$ .

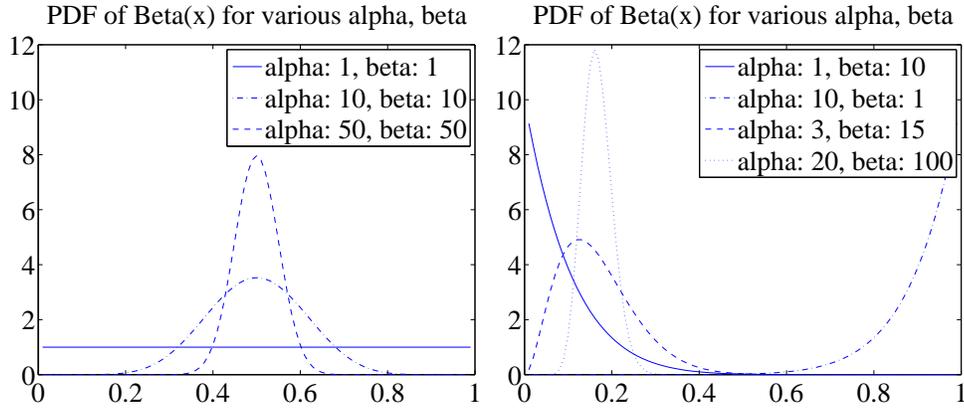


FIGURE 6.1: Probability density functions for the Beta distribution with a variety of different choices of  $\alpha$  and  $\beta$ .

**Useful Facts: 6.7** Gamma distribution

For a Gamma distribution with parameters  $\alpha, \beta$

1. The mean is  $\frac{\alpha}{\beta}$ .
2. The variance is  $\frac{\alpha}{\beta^2}$ .

Indexing: Gamma distribution

Indexing: Gamma distribution

#### 6.2.4 The Exponential Distribution

Assume we have an infinite interval of time or space, with points distributed on it. Assume these points form a Poisson point process, as above. For example, we might consider the times at which email arrives; or the times at which phone calls arrive at a large telephone exchange; or the locations of roadkill on a road. The distance (or span of time) between two consecutive points is a random variable  $X$ . This random variable takes an **exponential distribution**. There is a single parameter,  $\lambda$ . We have that

$$P_{\text{exp}}(x|\lambda) = \begin{cases} \lambda \exp^{-\lambda x} & \text{for } x \geq 0 \\ 0 & \text{otherwise} \end{cases} .$$

This distribution is often useful in modelling the failure of objects. We assume that failures form a Poisson process in time; then the time to the next failure is exponentially distributed.

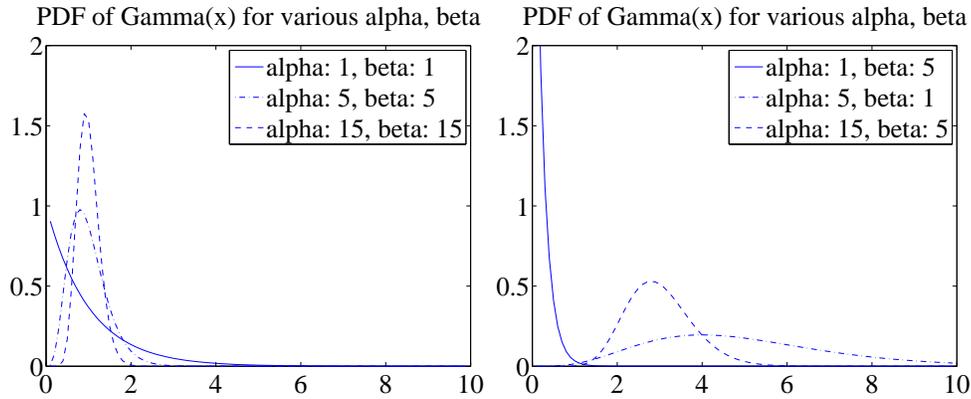


FIGURE 6.2: Probability density functions for the Gamma distribution with a variety of different choices of  $\alpha$  and  $\beta$ .

**Useful Facts: 6.8** *Exponential distribution*

For an exponential distribution with parameter  $\lambda$

1. The mean is  $\frac{1}{\lambda}$ .
2. The variance is  $\frac{1}{\lambda^2}$ .

Indexing: Exponential distribution

Indexing: Exponential distribution

Notice the relationship between this parameter and the parameter of the Poisson distribution. If (say) the phone calls are distributed with Poisson distribution with intensity  $\lambda$  (per hour), then your expected number of calls per hour is  $\lambda$ . The time between calls will be exponentially distributed with parameter  $\lambda$ , and the expected time to the next call is  $1/\lambda$  (in hours).

## 6.3 THE NORMAL DISTRIBUTION

### 6.3.1 The Standard Normal Distribution

**Definition: 6.6** *Standard Normal distribution*

The probability density function

$$p(x) = \left( \frac{1}{\sqrt{2\pi}} \right) \exp \left( \frac{-x^2}{2} \right).$$

is known as the **standard normal distribution**

Indexing: Standard Normal distribution

Indexing: Standard Normal distribution

The first step is to plot this probability density function (Figure 6.3). You should notice it is quite familiar from work on histograms, etc. in Chapter 13.1. It has the shape of the histogram of standard normal data, or at least the shape that the histogram of standard normal data aspires to.

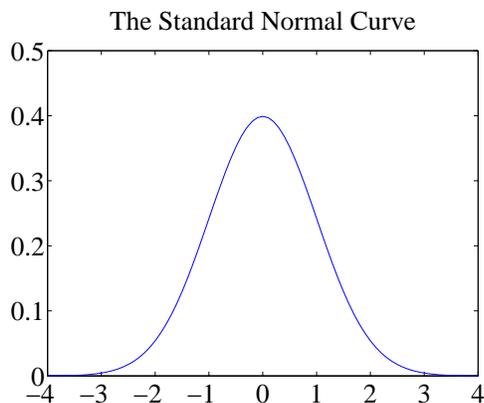


FIGURE 6.3: A plot of the probability density function of the standard normal distribution. Notice how probability is concentrated around zero, and how there is relatively little probability density for numbers with large absolute values.

**Useful Facts: 6.9** *standard normal distribution*

1. The mean of the standard normal distribution is 0.
2. The variance of the standard normal distribution is 1.

These results are easily established by looking up (or doing!) the relevant integrals; they are relegated to the exercises.

Indexing: standard normal distribution

Indexing: standard normal distribution

A continuous random variable is a **standard normal random variable** if its probability density function is a standard normal distribution.

### 6.3.2 The Normal Distribution

Any probability density function that is a standard normal distribution *in standard coordinates* is a **normal distribution**. Now write  $\mu$  for the mean of a random variable and  $\sigma$  for its standard deviation; we are saying that, if

$$\frac{x - \mu}{\sigma}$$

has a standard normal distribution, then  $p(x)$  is a normal distribution. We can work out the form of the probability density function of a general normal distribution in two steps: first, we notice that for any normal distribution, we must have

$$p(x) \propto \exp \left[ -\frac{(x - \mu)^2}{2\sigma^2} \right].$$

But, for this to be a probability density function, we must have  $\int_{-\infty}^{\infty} p(x)dx = 1$ . This yields the constant of proportionality, and we get

**Definition: 6.7** *Normal distribution*

The probability density function

$$p(x) = \left( \frac{1}{\sqrt{2\pi}\sigma} \right) \exp \left( \frac{-(x - \mu)^2}{2\sigma^2} \right).$$

is a normal distribution.

Indexing: Normal distribution

Indexing: Normal distribution

**Useful Facts: 6.10** *normal distribution*

The probability density function

$$p(x) = \left( \frac{1}{\sqrt{2\pi}\sigma} \right) \exp \left( \frac{-(x - \mu)^2}{2\sigma^2} \right).$$

has

1. mean  $\mu$
2. and variance  $\sigma$ .

These results are easily established by looking up (or doing!) the relevant integrals; they are relegated to the exercises.

Indexing: normal distribution

Indexing: normal distribution

A continuous random variable is a **normal random variable** if its probability density function is a **normal distribution**. Notice that it is quite usual to call normal distributions **gaussian distributions**.

### 6.3.3 Properties of the Normal Distribution

Normal distributions are important, because one often runs into data that is well described by a normal distribution. It turns out that anything that behaves like a binomial distribution with a lot of trials — for example, the number of heads in many coin tosses; as another example, the percentage of times you get the outcome of interest in a simulation in many runs — should produce a normal distribution (Section 6.4). For this reason, pretty much any experiment where you perform a simulation, then count to estimate a probability or an expectation, should give you an answer that has a normal distribution.

It is a remarkable and deep fact, known as the **central limit theorem**, that adding many independent random variables produces a normal distribution pretty much *whatever* the distributions of those random variables. I've not shown this in detail because it's a nuisance to prove. However, if you add together many random variables, each of pretty much any distribution, then the answer has a distribution close to the normal distribution. It turns out that many of the processes we observe add up subsidiary random variables. This means that you will see normal distributions very often in practice.

A normal random variable tends to take values that are quite close to the mean, measured in standard deviation units. We can demonstrate this important fact by computing the probability that a standard normal random variable lies between  $u$  and  $v$ . We form

$$\int_u^v \frac{1}{\sqrt{2\pi}} \exp \left( -\frac{u^2}{2} \right) du.$$

It turns out that this integral can be evaluated relatively easily using a special function. The **error function** is defined by

$$\operatorname{erf}(x) = \frac{2}{\sqrt{\pi}} \int_0^x \exp(-t^2) dt$$

so that

$$\frac{1}{2} \operatorname{erf}\left(\frac{x}{\sqrt{2}}\right) = \int_0^x \frac{1}{\sqrt{2\pi}} \exp\left(-\frac{u^2}{2}\right) du.$$

Notice that  $\operatorname{erf}(x)$  is an odd function (i.e.  $\operatorname{erf}(-x) = -\operatorname{erf}(x)$ ). From this (and tables for the error function, or Matlab) we get that, for a standard normal random variable

$$\frac{1}{\sqrt{2\pi}} \int_{-1}^1 \exp\left(-\frac{x^2}{2}\right) dx \approx 0.68$$

and

$$\frac{1}{\sqrt{2\pi}} \int_{-2}^2 \exp\left(-\frac{x^2}{2}\right) dx \approx 0.95$$

and

$$\frac{1}{\sqrt{2\pi}} \int_{-3}^3 \exp\left(-\frac{x^2}{2}\right) dx \approx 0.99.$$

These are very strong statements. They measure how often a standard normal random variable has values that are in the range  $-1, 1$ ,  $-2, 2$ , and  $-3, 3$  respectively. But these measurements apply to normal random variables if we recognize that they now measure how often the normal random variable is some number of standard deviations away from the mean. In particular, it is worth remembering that:

**Useful Facts: 6.11** *Normal Random Variables*

- About 68% of the time, a normal random variable takes a value within one standard deviation of the mean.
- About 95% of the time, a normal random variable takes a value within one standard deviation of the mean.
- About 99% of the time, a normal random variable takes a value within one standard deviation of the mean.

Indexing: Normal Random Variables

Indexing: Normal Random Variables

## 6.4 APPROXIMATING BINOMIALS WITH LARGE $N$

The Binomial distribution appears to be a straightforward thing. We assume we flip a coin  $N$  times, where  $N$  is a very large number. The coin has probability  $p$  of

coming up heads, and so probability  $q = 1 - p$  of coming up tails. The number of heads  $h$  follows the binomial distribution, so

$$P(h) = \frac{N!}{h!(N-h)!} p^h q^{(N-h)}$$

The mean of this distribution is  $Np$ , the variance is  $Npq$ , and the standard deviation is  $\sqrt{Npq}$ .

Evaluating this probability distribution for large  $N$  is very difficult, because factorials grow fast. We will construct an approximation to the binomial distribution for large  $N$  that allows us to evaluate the probability that  $h$  lies in some range. This approximation will show that the probability that  $h$  is within one standard deviation of the mean is approximately 68%.

This is important, because it shows that our model of probability as frequency is consistent. Consider the probability that the number of heads you see lies within one standard deviation of the mean. The size of that interval is  $2\sqrt{Npq}$ . As  $N$  gets bigger, the size of that interval, *relative to the total number of flips*, gets smaller. If I flip a coin  $N$  times, in principle I could see a number of heads  $h$  that ranges from 0 to  $N$ . However, we will establish that about 68% of the time,  $h$  will lie in the interval within one standard deviation of the mean. The size of this interval, *relative to the total number of flips* is

$$2\frac{\sqrt{Npq}}{N} = 2\sqrt{\frac{pq}{N}}.$$

As a result, as  $N \rightarrow \infty$ ,

$$\frac{h}{N} \rightarrow p$$

because  $h$  will tend to land in an interval around  $pN$  that gets narrower as  $N$  gets larger.

The main difficulty with Figure 6.4 (and with the argument above) is that the mean and standard deviation of the binomial distribution tends to infinity as the number of coin flips tends to infinity. This can confuse issues. For example, the plots of Figure 6.4 show narrowing probability distributions — but is this because the scale is compacted, or is there a real effect? It turns out there is a real effect, and a good way to see it is to consider the normalized number of heads.

#### 6.4.1 Large $N$

Recall that to normalize a dataset, you subtract the mean and divide the result by the standard deviation. We can do the same for a random variable. We now consider

$$x = \frac{h - Np}{\sqrt{Npq}}.$$

The probability distribution of  $x$  can be obtained from the probability distribution for  $h$ , because  $h = Np + x\sqrt{Npq}$ , so

$$P(x) = \left( \frac{N!}{(Np + x\sqrt{Npq})!(Nq - x\sqrt{Npq})!} \right) p^{(Np + x\sqrt{Npq})} q^{(Nq - x\sqrt{Npq})}.$$

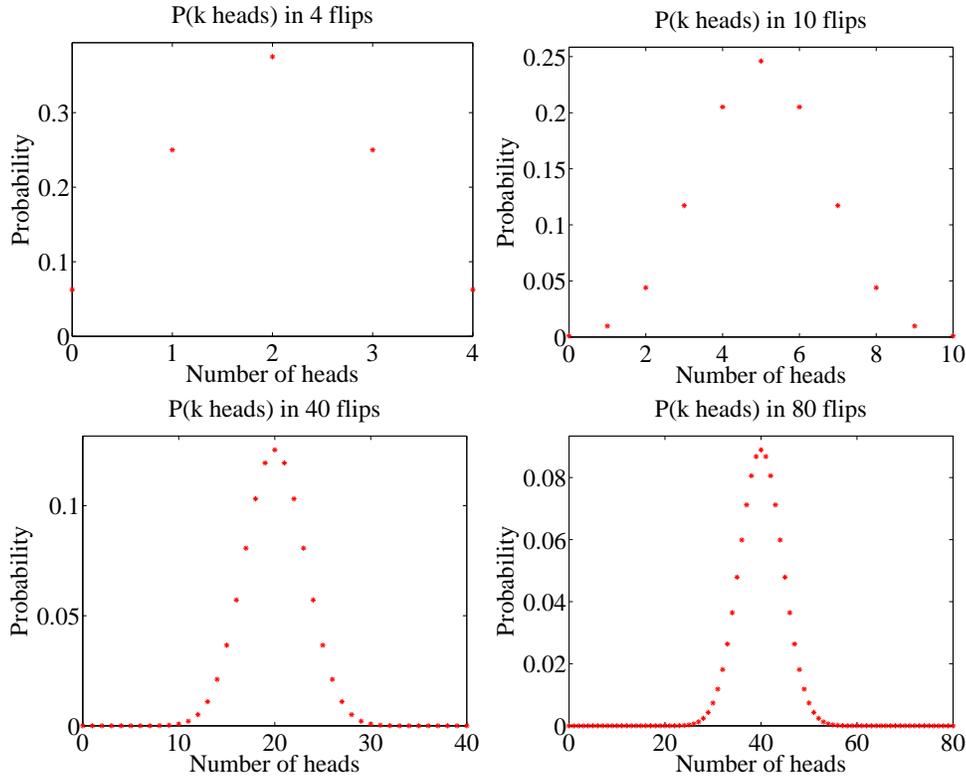


FIGURE 6.4: Plots of the binomial distribution for  $p = q = 0.5$  for different values of  $N$ . You should notice that the set of values of  $h$  (the number of heads) that have substantial probability is quite narrow compared to the range of possible values. This set gets narrower as the number of flips increases. This is because the mean is  $pN$  and the standard deviation is  $\sqrt{Npq}$  — so the fraction of values that is within one standard deviation of the mean is  $O(1/\sqrt{N})$ .

I have plotted this probability distribution for various values of  $N$  in Figure 6.5.

But it is hard to work with this distribution for very large  $N$ . The factorials become very difficult to evaluate. Second, it is a discrete distribution on  $N$  points, spaced  $1/\sqrt{Npq}$  apart. As  $N$  becomes very large, the number of points that have non-zero probability becomes very large, and  $x$  can be very large, or very small. For example, there is some probability, though there may be very little indeed, on the point where  $h = N$ , or, equivalently,  $x = N(p + \sqrt{Npq})$ . For sufficiently large  $N$ , we think of this probability distribution as a probability density function. We can do so, for example, by spreading the probability for  $x_i$  (the  $i$ 'th value of  $x$ ) evenly over the interval between  $x_i$  and  $x_{i+1}$ . We then have a probability density function that looks like a histogram, with bars that become narrower as  $N$  increases. But what is the limit?

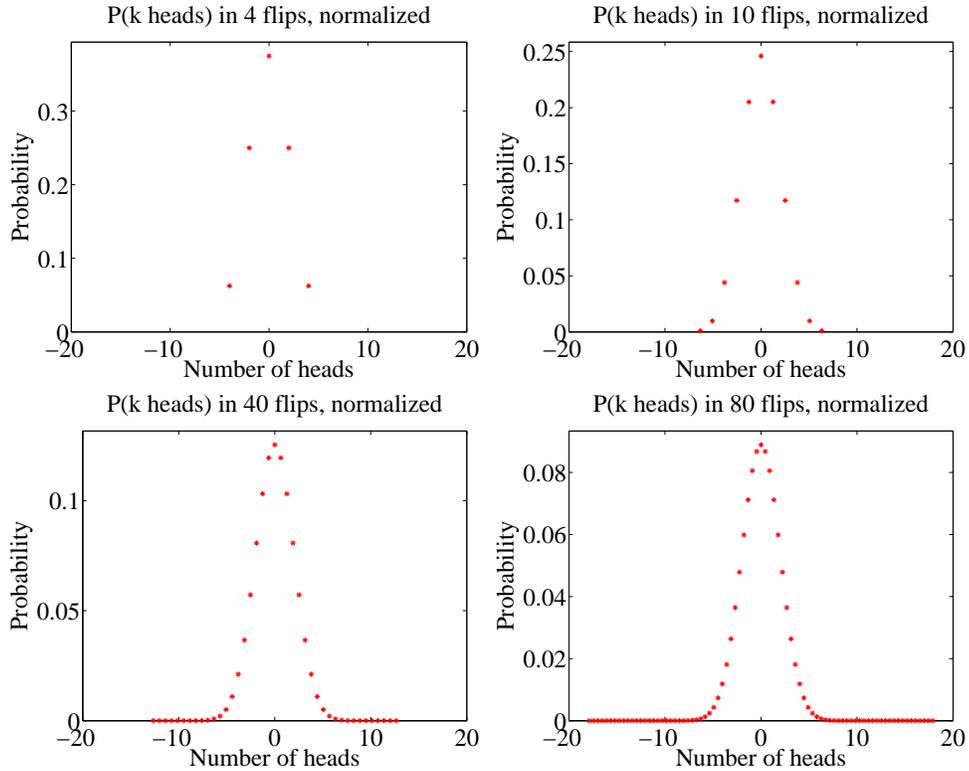


FIGURE 6.5: Plots of the distribution for the normalized variable  $x$ , with  $P(x)$  given in the text, obtained from the binomial distribution with  $p = q = 0.5$  for different values of  $N$ . These distributions are normalized (mean 0, variance 1. They look increasingly like a standard normal distribution EXCEPT that the value at their mode gets smaller as  $N$  gets bigger (there are more possible outcomes). In the text, we will establish that the standard normal distribution is a limit, in a useful sense.

### 6.4.2 Getting Normal

To proceed, we need Stirling’s approximation, which says that, for large  $N$ ,

$$N! \approx \sqrt{2\pi} \sqrt{N} \left(\frac{N}{e}\right)^N.$$

This yields

$$P(h) \approx \left(\frac{Np}{h}\right)^h \left(\frac{Nq}{N-h}\right)^{(N-h)} \sqrt{\frac{N}{2\pi h(N-h)}}$$

Recall we used the normalized variable

$$x = \frac{h - Np}{\sqrt{Npq}}.$$

We will encounter the term  $\sqrt{Npq}$  often, and we use  $\sigma = \sqrt{Npq}$  as a shorthand. We can compute  $h$  and  $N - h$  from  $x$  by the equalities

$$h = Np + \sigma x \qquad N - h = Nq - \sigma x.$$

So the probability distribution written in this new variable  $x$  is

$$P(x) \approx \left( \frac{Np}{Np + \sigma x} \right)^{(Np + \sigma x)} \left( \frac{Nq}{Nq - \sigma x} \right)^{(Nq - \sigma x)} \sqrt{\frac{N}{2\pi(Np + \sigma x)(Nq - \sigma x)}}$$

There are three terms to deal with here. It is easiest to work with  $\log P$ . Now

$$\log(1 + x) = x - \frac{1}{2}x^2 + O(x^3)$$

so we have

$$\begin{aligned} \log \left( \frac{Np}{Np + \sigma x} \right) &= -\log \left( 1 + \frac{\sigma x}{Np} \right) \\ &\approx -\frac{\sigma x}{Np} + \left( \frac{1}{2} \right) \left( \frac{\sigma x}{Np} \right)^2 \end{aligned}$$

and

$$\log \left( \frac{Nq}{Nq - \sigma x} \right) \approx \frac{\sigma x}{Nq} + \left( \frac{1}{2} \right) \left( \frac{\sigma x}{Nq} \right)^2.$$

From this, we have that

$$\begin{aligned} \log \left[ \left( \frac{Np}{Np + \sigma x} \right)^{(Np + \sigma x)} \left( \frac{Nq}{Nq - \sigma x} \right)^{(Nq - \sigma x)} \right] &\approx [Np + \sigma x] \left[ -\frac{\sigma x}{Np} + \left( \frac{1}{2} \right) \left( \frac{\sigma x}{Np} \right)^2 \right] + \\ &\quad [Nq - \sigma x] \left[ \frac{\sigma x}{Nq} + \left( \frac{1}{2} \right) \left( \frac{\sigma x}{Nq} \right)^2 \right] \\ &= -\left( \frac{1}{2} \right) x^2 + O((\sigma x)^3) \end{aligned}$$

(recall  $\sigma = \sqrt{Npq}$  if you're having trouble with the last step). Now we look at the square-root term. We have

$$\begin{aligned} \log \sqrt{\frac{N}{2\pi(Np + \sigma x)(Nq - \sigma x)}} &= -\frac{1}{2} (\log [Np + \sigma x] + \log [Nq - \sigma x] - \log N + \log 2\pi) \\ &= -\frac{1}{2} \left( \begin{array}{l} \log Np + O\left(\left(\frac{\sigma x}{Np}\right)\right) \\ + \log Nq - O\left(\left(\frac{\sigma x}{Nq}\right)\right) \\ - \log N + \log 2\pi \end{array} \right) \end{aligned}$$

but, since  $N$  is very large compared to  $\sigma x$ , we can ignore the  $O\left(\left(\frac{\sigma x}{Np}\right)\right)$  terms. Then this term is not a function of  $x$ . So we have

$$\log P(x) \approx \frac{-x^2}{2} + \text{constant}.$$

Now because  $N$  is very large, our probability distribution  $P$  limits to a probability density function  $p$ , with

$$p(x) \propto \exp\left(\frac{-x^2}{2}\right).$$

We can get the constant of proportionality from integrating, to

$$p(x) = \left(\frac{1}{\sqrt{2\pi}}\right) \exp\left(\frac{-x^2}{2}\right).$$

This constant of proportionality deals with the effect in figure 6.5, where the mode of the distribution gets smaller as  $N$  gets bigger. It does so because there are more points with non-zero probability to be accounted for. But we are interested in the limit where  $N$  tends to infinity. This must be a probability density function, so it must integrate to one.

Review this blizzard of terms. We started with a binomial distribution, but standardized the variables so that the mean was zero and the standard deviation was one. We then assumed there was a very large number of coin tosses, so large that that the distribution started to look like a continuous function. The function we get is the standard normal distribution.

### 6.4.3 So What?

I have proven an extremely useful fact, which I shall now put in a box.

**Useful Fact: 6.12** *Binomial distribution for large  $N$*

Assume  $h$  follows the binomial distribution with parameters  $p$  and  $q$ . Write  $x = \frac{h-Np}{\sqrt{Npq}}$ . Then, for sufficiently large  $N$ , the probability distribution  $P(x)$  can be approximated by the probability density function

$$\left(\frac{1}{\sqrt{2\pi}}\right) \exp\left(\frac{-x^2}{2}\right)$$

in the sense that

$$P(\{x \in [a, b]\}) \approx \int_a^b \left(\frac{1}{\sqrt{2\pi}}\right) \exp\left(\frac{-u^2}{2}\right) du$$

Indexing: Binomial distribution for large  $N$

Indexing: Binomial distribution for large  $N$

This justifies our model of probability as frequency. I interpreted an event having probability  $p$  to mean that, if I had a large number  $N$  of independent repetitions of the experiment, the number that produced the event would be close to  $Np$ , and would get closer as  $N$  got larger. We know that, for example, 68% of

the time a standard normal random variable takes a value between 1 and  $-1$ . In this case, the standard normal random variable is

$$\frac{h - (Np)}{\sqrt{Npq}}$$

so that 68% of the time,  $h$  must take a value in the range  $[Np - \sqrt{Npq}, Np + \sqrt{Npq}]$ . Equivalently, the relative frequency  $h/N$  must take a value in the range

$$\left[ p - \frac{pq}{\sqrt{N}}, p + \frac{pq}{\sqrt{N}} \right]$$

but as  $N \rightarrow \infty$  this range gets smaller and smaller, and  $h/N$  limits to  $p$ . So our view of probability as a frequency is consistent.

To obtain  $h$ , we added  $N$  independent Bernoulli random variables. So you can interpret the box as saying that the sum of many independent Bernoulli random variables has a probability distribution that limits to the normal distribution as the number added together gets larger. Remember that I have stated, though not precisely, but not proved the deep and useful fact that the sum of pretty much any independent random variables has a distribution that gets closer to a normal distribution as the number added together gets larger.

6.5 YOU SHOULD

6.5.1 remember:

Definition: Bernoulli random variable . . . . . 162  
 Useful facts: Bernoulli random variable . . . . . 163  
 Definition: Geometric distribution . . . . . 163  
 Useful facts: Geometric distribution . . . . . 163  
 Definition: Binomial distribution . . . . . 164  
 Useful fact: Binomial distribution . . . . . 164  
 Useful facts: Binomial distribution . . . . . 165  
 Definition: Multinomial distribution . . . . . 167  
 New term: intensity . . . . . 168  
 Definition: Poisson distribution . . . . . 168  
 Useful facts: Poisson distribution . . . . . 168  
 New term: Poisson point process . . . . . 169  
 New term: uniform distribution . . . . . 169  
 New term: uniform random variable . . . . . 169  
 Useful facts: Beta distribution . . . . . 170  
 Useful facts: Gamma distribution . . . . . 171  
 New term: exponential distribution . . . . . 171  
 Useful facts: Exponential distribution . . . . . 172  
 New term: standard normal distribution . . . . . 173  
 Definition: Standard Normal distribution . . . . . 173  
 Useful facts: standard normal distribution . . . . . 173  
 New term: standard normal random variable . . . . . 174  
 New term: normal distribution . . . . . 174  
 Definition: Normal distribution . . . . . 174  
 Useful facts: normal distribution . . . . . 175  
 New term: normal random variable . . . . . 175  
 New term: normal distribution . . . . . 175  
 New term: gaussian distributions . . . . . 175  
 New term: central limit theorem . . . . . 175  
 New term: error function . . . . . 176  
 Useful facts: Normal Random Variables . . . . . 176  
 Useful fact: Binomial distribution for large  $N$  . . . . . 181

## PROBLEMS

## Sums and Differences of Discrete Random Variables

**6.1.** Assume  $X$  and  $Y$  are discrete random variables which take integer values in the range  $1 \dots 100$  (inclusive). Write  $S = X + Y$ .

(a) Show that

$$P(S = k) = \sum_{u=1}^{u=100} P(\{X = k - u\} \cap \{Y = u\}).$$

(b) Show that

$$P(D = k) = \sum_{u=1}^{u=100} P(\{X = k + u\})P(\{Y = u\}).$$

(c) Now assume that both  $X$  and  $Y$  are uniform random variables. Show that  $S$  is not uniform by considering  $P(S = 2)$ ,  $P(S = 3)$ , and  $P(S = 100)$ .

## The Geometric Distribution

**6.2.** Write  $S_\infty = \sum_{i=0}^{\infty} r^i$ . Show that  $(1 - r)S_\infty = 1$ , so that

$$S_\infty = \frac{1}{1 - r}$$

**6.3.** Use the result of the previous exercise to show that

$$\begin{aligned} \sum_{n=1}^{\infty} P(\{X = n\}) &= p \sum_{n=1}^{\infty} (1 - p)^{(n-1)} \\ &= 1 \end{aligned}$$

**6.4.** Show that

$$\sum_{i=0}^{\infty} ir^i = \left(\sum_{i=1}^{\infty} r^i\right) + r\left(\sum_{i=1}^{\infty} r^i\right) + r^2\left(\sum_{i=1}^{\infty} r^i\right) + \dots$$

(look carefully at the limits of the sums!) and so show that

$$\sum_{i=0}^{\infty} ir^i = \frac{r}{(1 - r)^2}.$$

**6.5.** Write  $S_\infty = \sum_{i=0}^{\infty} r^i$ . Show that

$$\sum_{i=0}^{\infty} i^2 r^i = (\gamma - 1) + 3r(\gamma - 1) + 5r^2(\gamma - 1) + \dots$$

and so that

$$\sum_{i=0}^{\infty} i^2 r^i = \frac{r(1 + r)}{(1 - r)^3}$$

**6.6.** Show that, for a geometric distribution with parameter  $p$ , the mean is

$$\sum_{i=1}^{\infty} i(1-p)^{(i-1)}p = \sum_{i=0}^{\infty} (i+1)(1-p)^i p.$$

Now by rearranging and using the previous results, show that the mean is

$$\sum_{i=1}^{\infty} i(1-p)^{(i-1)}p = \frac{1}{p}$$

**6.7.** Show that, for a geometric distribution with parameter  $p$ , the variance is  $\frac{1-p}{p^2}$ . To do this, note the variance is  $\mathbb{E}[X^2] - \mathbb{E}[X]^2$ . Now use the results of the previous exercises to show that

$$\mathbb{E}[X^2] = \sum_{i=1}^{\infty} i^2(1-p)^{(i-1)}p = \frac{p}{1-p} \frac{(1-p)(2-p)}{p^3},$$

then rearrange to get the expression for variance.

#### Bernoulli Random Variables

- 6.8.** Write  $X$  for a Bernoulli random variable which takes the value 1 with probability  $p$  (and 0 with probability  $(1-p)$ ).
- (a) Show that  $\mathbb{E}[X] = p$ .
- (b) Show that  $\mathbb{E}[X^2] - \mathbb{E}[X]^2 = p(1-p)$

#### The Binomial Distribution

- 6.9.** Show that  $P_b(N-i; N, p) = P_b(i; N, p)$  for all  $i$ .
- 6.10.** Write  $h_r$  for the number of heads obtained in  $r$  flips of a coin which has probability  $p$  of coming up heads. Compare the following two ways to compute the probability of getting  $i$  heads in five coin flips:
- Flip the coin three times, count  $h_3$ , then flip the coin twice, count  $h_2$ , then form  $w = h_3 + h_2$ .
  - Flip the coin five times, and count  $h_5$ .

Show that the probability distribution for  $w$  is the same as the probability distribution for  $h_5$ . Do this by showing that

$$P(\{w = i\}) = \sum_{j=0}^5 P(\{h_3 = j\} \cap \{h_2 = i - j\}) = P(\{h_5 = i\}).$$

- 6.11.** Now we will do the previous exercise in a more general form. Again, write  $h_r$  for the number of heads obtained in  $r$  flips of a coin which has probability  $p$  of coming up heads. Compare the following two ways to compute the probability of getting  $i$  heads in  $N$  coin flips:
- Flip the coin  $t$  times, count  $h_t$ , then flip the coin  $N-t$  times, count  $h_{N-t}$ , then form  $w = h_t + h_{N-t}$ .
  - Flip the coin  $N$  times, and count  $h_N$ .

Show that the probability distribution for  $w$  is the same as the probability distribution for  $h_N$ . Do this by showing that

$$P(\{w = i\}) = \sum_{j=0}^N P(\{h_t = j\} \cap \{h_{N-t} = i - j\}) = P(\{h_N = i\}).$$

You will likely find the recurrence relation

$$P_b(i; N, p) = pP_b(i - 1; N - 1, p) + (1 - p)P_b(i; N - 1, p).$$

is useful.

- 6.12.** An airline runs a regular flight with six seats on it. The airline sells six tickets. The gender of the passengers is unknown at time of sale, but women are as common as men in the population. All passengers always turn up for the flight. The pilot is eccentric, and will not fly a plane unless at least one passenger is female. What is the probability that the pilot flies?
- 6.13.** An airline runs a regular flight with  $s$  seats on it. The airline always sells  $t$  tickets for this flight. The probability a passenger turns up for departure is  $p$ , and passengers do this independently. What is the probability that the plane travels with exactly 3 empty seats?
- 6.14.** An airline runs a regular flight with  $s$  seats on it. The airline always sells  $t$  tickets for this flight. The probability a passenger turns up for departure is  $p$ , and passengers do this independently. What is the probability that the plane travels with 1 or more empty seats?
- 6.15.** An airline runs a regular flight with 10 seats on it. The probability that a passenger turns up for the flight is 0.95. What is the smallest number of seats the airline should sell to ensure that the probability the flight is full (i.e. 10 or more passengers turn up) is bigger than 0.99? (you'll probably need to use a calculator or write a program for this).

### The Multinomial Distribution

- 6.16.** Show that the multinomial distribution

$$P_m(n_1, \dots, n_k; N, p_1, \dots, p_k) = \frac{N!}{n_1!n_2!\dots n_k!} p_1^{n_1} p_2^{n_2} \dots p_k^{n_k}$$

must satisfy the recurrence relation

$$\begin{aligned} P_m(n_1, \dots, n_k; N, p_1, \dots, p_k) &= p_1 P_m(n_1 - 1, \dots, n_k; N - 1, p_1, \dots, p_k) + \\ & p_2 P_m(n_1, n_2 - 1, \dots, n_k; N - 1, p_1, \dots, p_k) + \dots \\ & p_k P_m(n_1, n_2, \dots, n_k - 1; N - 1, p_1, \dots, p_k) \end{aligned}$$

### The Poisson Distribution

- 6.17.** Compute the Taylor series for  $xe^x$  around  $x = 0$ . Use this and pattern matching to show that the mean of the Poisson distribution with intensity parameter  $\lambda$  is  $\lambda$ .
- 6.18.** Compute the Taylor series for  $(x^2 + x)e^x$  around  $x = 0$ . Use this and pattern matching to show that the variance of the Poisson distribution with intensity parameter  $\lambda$  is  $\lambda$ .

## Sums of Continuous Random Variables

- 6.19.** Write  $p_x$  for the probability density function of a continuous random variable  $X$  and  $p_y$  for the probability density function of a continuous random variable  $Y$ . Show that the probability density function of  $S = X + Y$  is

$$p(s) = \int_{-\infty}^{\infty} p_x(s-u)p_y(u)du = \int_{-\infty}^{\infty} p_x(u)p_y(s-u)du$$

## The Normal Distribution

- 6.20.** Write

$$f(x) = \left( \frac{1}{\sqrt{2\pi}} \right) \exp \left( \frac{-x^2}{2} \right).$$

- (a) Show that  $f(x)$  is non-negative for all  $x$ .  
 (b) By integration, show that

$$\int_{-\infty}^{\infty} f(x)dx = 1,$$

so that  $f(x)$  is a probability density function.

- (c) Show that

$$\int_{-\infty}^{\infty} xf(x)dx = 0.$$

The easiest way to do this is to notice that  $f(x) = f(-x)$

- (d) Show that

$$\int_{-\infty}^{\infty} xf(x-\mu)dx = \mu.$$

The easiest way to do this is to change variables, and use the previous two exercises.

- (e) Show that

$$\int_{-\infty}^{\infty} x^2 f(x)dx = 1.$$

You'll need to either do, or look up, the integral to do this exercise.

- 6.21.** Write

$$g(x) = \exp \left[ -\frac{(x-\mu)^2}{2\sigma^2} \right]$$

Show that

$$\int_{-\infty}^{\infty} g(x)dx = \sqrt{2\pi}\sigma.$$

You can do this by a change of variable, and the results of the previous exercises.

- 6.22.** Write

$$p(x) = \left( \frac{1}{\sqrt{2\pi}\sigma} \right) \exp \left( \frac{-(x-\mu)^2}{2\sigma^2} \right).$$

- (a) Show that

$$\int_{-\infty}^{\infty} xp(x)dx = \mu$$

using the results of the previous exercises.

(b) Show that

$$\int_{-\infty}^{\infty} (x - \mu)^2 p(x) dx = \sigma^2$$

using the results of the previous exercises.

### The Binomial Distribution for Large $N$

- 6.23.** I flip a fair coin  $N$  times and count heads. We consider the probability that  $h$ , the fraction of heads, is in some range of numbers. For each of these questions, you should just write an expression, rather than evaluate the integral. *Hint:* If you know the range of numbers for  $h$ , you know the range for  $h/N$ .
- (a) For  $N = 1e6$ , use the normal approximation to estimate  $P(\{h \in [49500, 50500]\})$ .
  - (b) For  $N = 1e4$ , use the normal approximation to estimate  $P(\{h > 9000\})$ .
  - (c) For  $N = 1e2$ , use the normal approximation to estimate  $P(\{h > 60\} \cup \{h < 40\})$ ?

## CHAPTER 7

# Markov Chains and Simulation

There are many situations where one must work with a sequence of random variables. For example, consider a bus queue; people arrive at random, and so do buses. What is the probability that the queue gets to a particular length? and what is the expected length of the queue? As another example, we could have a random model of purchases from a shop — under any particular rule for restoring inventory, what is the largest (resp. smallest) amount of stock on the shelf? what is the expected amount of stock on the shelf? It turns out that there is a simple and powerful model that applies to many sequences of random variables. One can often use this model to make closed-form predictions. In cases where closed-form predictions aren't available, one can use simulation methods to estimate probabilities and expectations.

### 7.1 MARKOV CHAINS

A Markov chain is a sequence of random variables which has important independence properties. We will describe these properties in some detail below; first, we give some examples. Markov chains are easily represented with the language of finite state machines. For some Markov chains, it is easy to determine probabilities and expectations of interest in closed form using simple methods. For others, it is tricky (straightforward, but unreasonable amounts of straightforward work). In these cases, we can estimate the relevant probabilities and expectations by simulating the finite state machine.

#### 7.1.1 Motivating Example: Multiple Coin Flips

We start with three examples, each of which is easy to work. Each suggests a much harder question, however, which we need new machinery to handle.

**Worked example 7.1**    *Multiple Coin Flips - 1*

You choose to flip a fair coin until you see two heads in a row, and then stop. What is the probability that you flip the coin twice?

**Solution:** Because you stopped after two flips, you must have seen two heads. So  $P(2 \text{ flips}) = P(\{HH\}) = P(\{H\})^2 = 1/4$ .

**Worked example 7.2** *Multiple Coin Flips - 2*

You choose to flip a fair coin until you see two heads in a row, and then stop. What is the probability that you flip the coin three times?

**Solution:** Because you stopped after three flips, you must have seen  $T$ , then  $H$ , then  $H$ ; any other sequence either doesn't stop, or stops too early. We write this with the last flip last, so  $THH$ . So  $P(3 \text{ flips}) = P(\{THH\}) = P(\{T\})P(\{H\})^2 = 1/8$ .

**Worked example 7.3** *Multiple Coin Flips - 3*

You choose to flip a fair coin until you see two heads in a row, and then stop. What is the probability that you flip the coin four times?

**Solution:** This is more interesting. The last three flips must have been  $THH$  (otherwise you'd go on too long, or end too early). But, because the second flip must be a  $T$ , the first could be either  $H$  or  $T$ . This means there are two sequences that work:  $HTHH$  and  $TTHH$ . So  $P(4 \text{ flips}) = 2/8 = 1/4$ .

The harder question here is to ask what is  $P(N)$ , for  $N$  some number (larger than 4, because we know the answers in the other cases). It is unattractive to work this case by case (you could try this, if you don't believe me). One very helpful way to think about the coin flipping experiment is as a finite state machine (Figure 7.1). If you think of this machine as a conventional finite state machine, it accepts any string of  $T$ ,  $H$ , that (a) ends with  $HH$ , and (b) has no other  $HH$  in it. Alternatively, you could think of this machine as encoding a probabilistic process. At each state, an event occurs with some probability (in this case, a coin comes up  $H$  or  $T$ , with probability  $(1/2)$ ). The event causes the machine to follow the appropriate edge. If we take this view, this machine stops when we have flipped two heads in succession. It encodes our problem of computing the probability of flipping a coin  $N$  times then stopping — this is just the probability that the machine hits the end state after  $N$  transitions.

It turns out to be straightforward to construct a recurrence relation for  $P(N)$  (i.e. an equation for  $P(N)$  in terms of  $P(N-1)$ ,  $P(N-2)$ , and so on). This property is quite characteristic of repeated experiments. For this example, first, notice that  $P(1) = 0$ . Now imagine you have flipped the coin  $N$  times and stopped. We can assume that  $N > 3$ , because we know what happens for the other cases. The *last* three flips must have been  $THH$ . Write  $\Sigma_N$  for a sequence that is: (a)  $N$  flips long; (b) ends in  $HH$ ; and (c) contains no other subsequence  $HH$ . Equivalently, this is a string accepted by the finite state machine of figure 7.1. We must have that any

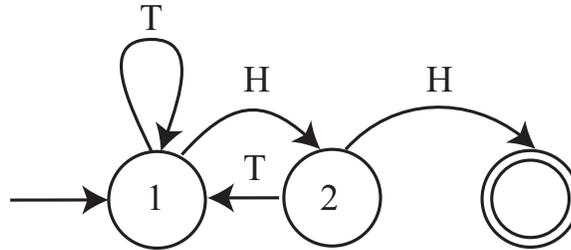


FIGURE 7.1: A finite state machine representing the coin flip example. By convention, the end state is a double circle, and the start state has an incoming arrow. I've labelled the arrows with the event that leads to the transition, but haven't bothered to put in the probabilities, because each is 0.5.

$\Sigma_N$  has either the form  $T\Sigma_{N-1}$  or the form  $HT\Sigma_{N-2}$ . But this means that

$$\begin{aligned} P(N) &= P(T)P(N-1) + P(HT)P(N-2) \\ &= (1/2)P(N-1) + (1/4)P(N-2) \end{aligned}$$

It is possible to solve this recurrence relation to get an explicit formula, but doing so would take us out of our way. You will check this recurrence relation in an exercise.

One really useful way to think of this recurrence relation is that represents an exercise in counting. We want to count all sequences that are of length  $N$ , and are accepted by the finite state machine of figure 7.1. This means they: (a) are  $N$  flips long; (b) end in  $HH$ ; and (c) contain no other subsequence  $HH$ . Now work backward along the FSM. The only way to arrive at the final state is to be in state 1, then see  $HH$ . So you can obtain an acceptable  $N$  element sequence by (a) prepending a  $T$  to an acceptable  $N-1$  element sequence or (b) prepending  $TH$  (which takes you to 2, then back to 1) to an acceptable  $N-2$  element sequence. This line of reasoning can be made much more elaborate. There are a few examples in the exercises.

### 7.1.2 Motivating Example: The Gambler's Ruin

Another useful example is known as the **gambler's ruin**. Assume you bet \$1 a tossed coin will come up heads. If you win, you get \$1 and your original stake back. If you lose, you lose your stake. But this coin has the property that  $P(H) = p < 1/2$ . We will study what happens when you bet repeatedly.

Assume you have \$ $s$  when you start. You will keep betting until either (a) you have \$0 (you are ruined; you can't borrow money) or (b) the amount of money you have accumulated is \$ $j$ , where  $j > s$ . The coin tosses are independent. We will compute  $P(\text{ruined, starting with } s|p)$  the probability that you leave the table with nothing, when you start with \$ $s$ . For brevity, we write  $P(\text{ruined, starting with } s|p) = p_s$ . You can represent the gambler's ruin problem with a finite state machine as well. Figure 7.2 shows a representation of the gambler's ruin problem.

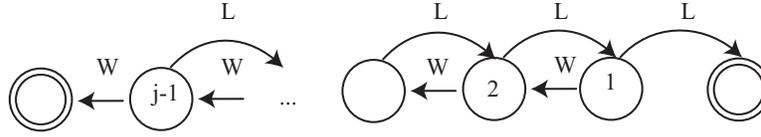


FIGURE 7.2: A finite state machine representing the gambler's ruin example. I have labelled each state with the amount of money the gambler has at that state. There are two end states, where the gambler has zero (is ruined), or has  $j$  and decides to leave the table. The problem we discuss is to compute the probability of being ruined, given the start state is  $s$ . This means that any state except the end states could be a start state. I have labelled the state transitions with "W" (for win) and "L" for lose, but have omitted the probabilities.

**Worked example 7.4** *The gambler's ruin - 1*

Using the notation above, determine  $p_0$  and  $p_j$

**Solution:** We must have  $p_0 = 1$ , because if you have \$0, you leave the table. Similarly, if you have \$ $j$ , you leave the table with \$ $j$ , so you don't leave the table with nothing, so  $p_j = 0$ .

**Worked example 7.5** *The gambler's ruin - 2*

Using the notation above, write a recurrence relation for  $p_s$  (the probability that you leave the table with nothing when you started with \$ $s$ ).

**Solution:** Assume that you win the first bet. Then you have \$ $s + 1$ , so your probability of leaving the table with nothing now becomes  $p_{s+1}$ . If you lose the first bet, then you have \$ $s - 1$ , so your probability of leaving the table with nothing now becomes  $p_{s-1}$ . The coin tosses are independent, so we can write

$$p_s = pp_{s+1} + (1-p)p_{s-1}.$$

Some fairly lively work with series, relegated to the end of the chapter as exercises, yields

$$p_s = \frac{\left(\frac{1-p}{p}\right)^j - \left(\frac{1-p}{p}\right)^s}{\left(\frac{1-p}{p}\right)^j - 1}.$$

This expression is quite informative. Notice that, if  $p < 1/2$ , then  $(1-p)/p > 1$ . This means that as  $j \rightarrow \infty$ , we have  $p_s \rightarrow 1$ . If you gamble repeatedly on an unfair

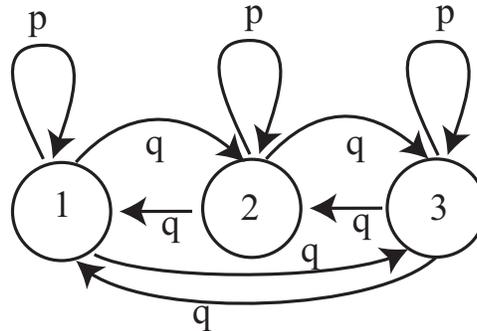


FIGURE 7.3: A virus can exist in one of 3 strains. At the end of each year, the virus mutates. With probability  $\alpha$ , it chooses uniformly and at random from one of the 2 other strains, and turns into that; with probability  $1 - \alpha$ , it stays in the strain it is in. For this figure, we have transition probabilities  $p = (1 - \alpha)$  and  $q = (\alpha/2)$ .

coin, the probability that you run out of money before you hit some threshold ( $j$  in this case) tends to one.

### 7.1.3 Motivating Example: A Virus

Problems represented with a finite state machine don't need to have an end state. As one example, (which I got from ACC Coolen's lecture notes), we have a virus that can exist in one of  $k$  strains. At the end of each year, the virus mutates. With probability  $\alpha$ , it chooses uniformly and at random from one of the  $k - 1$  other strains, and turns into that; with probability  $1 - \alpha$ , it stays in the strain it is in (Figure 7.3 shows an example with three strains). For that figure, we have  $p = (1 - \alpha)$  and  $q = (\alpha/2)$ . The virus just keeps on changing, and there is no end state. But there are a variety of very interesting questions we can try to answer. We would like to know, for example, the expected time to see a strain a second time, i.e. if the virus is in strain 1 at time 1, what is the expected time before it is in strain 1 again? If the virus has mutated many times, what is the probability that it is in each strain? and do these probabilities depend on the start strain?

### 7.1.4 Markov Chains

In Figure 7.1 and Figure 7.2, I showed the event that caused the state transitions. It is more usual to write a probability on the figure, as I did in Figure 7.3, because the probability of a state transition (rather than what caused it) is what really matters. The underlying object is now a weighted directed graph, because we have removed the events and replaced them with probabilities. These probabilities are known as **transition probabilities**; notice that the sum of transition probabilities over *outgoing* arrows must be 1.

We can now think of our process as a **biased random walk** on a weighted directed graph. A bug (or any other small object you prefer) sits on one of the graph's nodes. At each time step, the bug chooses one of the outgoing edges at

random. The probability of choosing an edge is given by the probabilities on the drawing of the graph (equivalently, the transition probabilities). The bug then follows that edge. The bug keeps doing this until it hits an end state.

This bug produces a sequence of random variables. If there are  $k$  states in the finite state machine, we can label each state with a number,  $1 \dots k$ . At the  $n$ 'th time step, the state of the process — the node that the bug is sitting on — is a random variable, which we write  $X_n$ . These random variables have an important property. The probability that  $X_n$  takes some particular value depends only on  $X_{n-1}$ , and not on any other previous state. If we know where the bug is at step  $n - 1$  in our model, we know where it could go, and the probability of each transition. Where it was at previous times does not affect this, *as long as we know its state at step  $n - 1$* .

A sequence of random variables  $X_n$  is a **Markov chain** if it has the property that,  $P(X_n = j | \text{values of all previous states}) = P(X_n = j | X_{n-1})$ , or, equivalently, only the last state matters in determining the probability of the current state. The probabilities  $P(X_n = j | X_{n-1} = i)$  are the **transition probabilities**. Any model built by taking the transitions of a finite state machine and labelling them with probabilities must be a Markov chain. However, this is not the only way to build or represent a Markov chain.

One representation of a Markov chain uses a matrix of transition probabilities. We define the matrix  $\mathcal{P}$  with  $p_{ij} = P(X_n = j | X_{n-1} = i)$ . Notice that this matrix has the properties that  $p_{ij} \geq 0$  and

$$\sum_j p_{ij} = 1$$

because at the end of each time step the model must be in some state. Equivalently, the sum of transition probabilities for outgoing arrows is one. Non-negative matrices with this property are **stochastic matrices**. By the way, you should look very carefully at the  $i$ 's and  $j$ 's here — Markov chains are usually written in terms of *row* vectors, and this choice makes sense in that context.

#### Worked example 7.6 Viruses

Write out the transition probability matrix for the virus of Figure 7.3, assuming that  $\alpha = 0.2$ .

**Solution:** We have  $P(X_n = 1 | X_{n-1} = 1) = (1 - \alpha) = 0.8$ , and  $P(X_n = 2 | X_{n-1} = 1) = \alpha/2 = P(X_n = 3 | X_{n-1} = 1)$ ; so we get

$$\begin{pmatrix} 0.8 & 0.1 & 0.1 \\ 0.1 & 0.8 & 0.1 \\ 0.1 & 0.1 & 0.8 \end{pmatrix}$$

Now imagine we do not know the initial state of the chain, but instead have a probability distribution. This gives  $P(X_0 = i)$  for each state  $i$ . It is usual to take these  $k$  probabilities and place them in a  $k$ -dimensional row vector, which is usually written  $\pi$ . For example, we might not know what the initial strain of the virus is, but just that each strain is equally likely. So for a 3-strain virus, we would have  $\pi = [1/3, 1/3, 1/3]$ . From this information, we can compute the probability distribution over the states at time 1 by

$$\begin{aligned} P(X_1 = j) &= \sum_i P(X_1 = j, X_0 = i) \\ &= \sum_i P(X_1 = j | X_0 = i) P(X_0 = i) \\ &= \sum_i p_{ij} \pi_i. \end{aligned}$$

If we write  $\mathbf{p}^{(n)}$  for the row vector representing the probability distribution of the state at step  $n$ , we can write this expression as

$$\mathbf{p}^{(1)} = \pi \mathcal{P}.$$

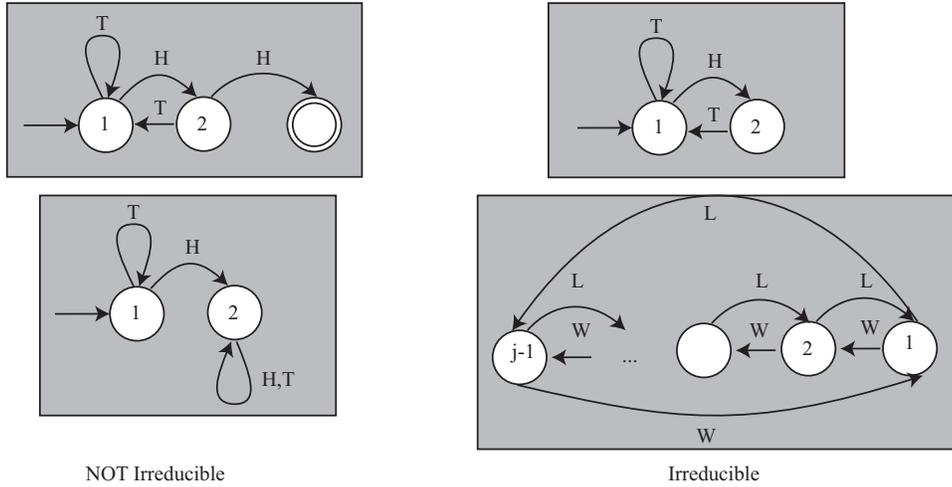
Now notice that

$$\begin{aligned} P(X_2 = j) &= \sum_i P(X_2 = j, X_1 = i) \\ &= \sum_i P(X_2 = j | X_1 = i) P(X_1 = i) \\ &= \sum_i p_{ij} \left( \sum_k p_{ki} \pi_k \right). \end{aligned}$$

so that

$$\mathbf{p}^{(n)} = \pi \mathcal{P}^n.$$

This expression is useful for simulation, and also allows us to deduce a variety of interesting properties of Markov chains.



NOT Irreducible

Irreducible

FIGURE 7.4: Examples of finite state machines that give rise to Markov chains that are NOT irreducible (left) and irreducible (right). To obtain Markov chains from these drawings, we would have to give the probabilities of the events that lead to state transitions. We'll assume that none of the probabilities are zero; after that, the values don't matter for irreducibility analysis. The **top left** FSM is not irreducible, because it has an end state; once the bug reaches this point, it can't go anywhere else. The **bottom left** FSM is not irreducible, because the bug can get stuck in state 2.

**Worked example 7.7** Viruses

We know that the virus of Figure 7.3 started in strain 1. After two state transitions, what is the distribution of states when  $\alpha = 0.2$ ? when  $\alpha = 0.9$ ? What happens after 20 state transitions? If the virus starts in strain 2, what happens after 20 state transitions?

**Solution:** If the virus started in strain 1, then  $\pi = [1, 0, 0]$ . We must compute  $\pi(\mathcal{P}(\alpha))^2$ . This yields  $[0.66, 0.17, 0.17]$  for the case  $\alpha = 0.2$  and  $[0.4150, 0.2925, 0.2925]$  for the case  $\alpha = 0.9$ . Notice that, because the virus with small  $\alpha$  tends to stay in whatever state it is in, the distribution of states after two years is still quite peaked; when  $\alpha$  is large, the distribution of states is quite uniform. After 20 transitions, we have  $[0.3339, 0.3331, 0.3331]$  for the case  $\alpha = 0.2$  and  $[0.3333, 0.3333, 0.3333]$  for the case  $\alpha = 0.9$ ; you will get similar numbers even if the virus starts in strain 2. After 20 transitions, the virus has largely “forgotten” what the initial state was.

In example 7.7, the distribution of virus strains after a long interval appeared not to depend much on the initial strain. This property is true of many Markov chains. Assume that any state can be reached from any other state, by some

sequence of transitions. Such chains are called **irreducible**; notice this means there is no end state, like the virus example. Irreducibility also means that the chain cannot get “stuck” in a state or a collection of states (Figure 7.4). Then there is a unique vector  $\mathbf{s}$ , usually referred to as the **stationary distribution**, such that for *any* initial state distribution  $\pi$ ,

$$\lim_{n \rightarrow \infty} \pi \mathcal{P}^{(n)} = \mathbf{s}.$$

Equivalently, if the chain has run through many steps, it no longer matters what the initial distribution is. You expect that the probability distribution over states is  $\mathbf{s}$ .

### 7.1.5 Example: Particle Motion as a Markov Chain

One can find Markov chains in quite unexpected places, often with useful consequences. In this example, I will obtain a Markov chain without reasoning about graphs or finite state machines. We will investigate a particle moving under gravity, in 3 dimensions. Write the position of the particle as  $\mathbf{p}(t)$ , its velocity as  $\mathbf{v}(t)$ , its acceleration as  $\mathbf{a}(t)$ , its mass as  $m$ , and the gravitational force as  $\mathbf{g}$ . Then we know that

$$\begin{aligned} \mathbf{v}(t) &= \frac{d\mathbf{p}}{dt} \\ \mathbf{a}(t) &= \frac{d\mathbf{v}}{dt} \\ &= \mathbf{g}. \end{aligned}$$

Now stack the position and the acceleration into a single vector  $X(t) = (\mathbf{p}(t), \mathbf{v}(t))^T$ . We could write these equations as

$$\frac{dX}{dt} = \mathcal{A}X + \mathbf{b}$$

where

$$\mathcal{A} = \begin{pmatrix} 0 & \mathcal{I} \\ 0 & 0 \end{pmatrix}$$

and

$$\mathbf{b} = \begin{pmatrix} 0 \\ \mathbf{g} \end{pmatrix}.$$

Now imagine that we look at the position, velocity and acceleration of the particle at fixed time instants, so we are really interested in  $X_i = X(t_0 + i\Delta t)$ . In this case, we can approximate  $\frac{dX}{dt}$  by  $(X_{i+1} - X_i)/\Delta t$ . We then have

$$X_{i+1} = X_i + (\Delta t)(\mathcal{A}X_i + \mathbf{b}).$$

This is beginning to look like a Markov chain, because  $X_{i+1}$  depends only on  $X_i$  but not on any previous  $X$ . But there isn't any randomness here. We can fix that by assuming that the particle is moving in, say, a light turbulent wind. Then the acceleration at any time consists of (a) the acceleration of gravity and (b) a small,

random force produced by the wind. Write  $\mathbf{w}_i$  for this small, random force at time  $i$ , and  $P(\mathbf{w}_i|i)$  for its probability distribution, which could reasonably depend on time. We can then rewrite our equation by writing

$$X_{i+1} = X_i + (\Delta t)(\mathcal{A}X_i + \mathbf{b}_r).$$

where

$$\mathbf{b}_r = \begin{pmatrix} 0 \\ \mathbf{g} + \mathbf{w}_i \end{pmatrix}.$$

Now the  $X_i$  are clearly random variables. We could get  $P(X_{i+1}|X_i)$  by rearranging terms. If I know  $X_{i+1}$  and  $X_i$ , I know the value of  $\mathbf{w}_i$ ; I could plug this into  $P(\mathbf{w}_i|i)$  to yield  $P(X_{i+1}|X_i)$ .

Using a finite state machine in this example would be a bit unnatural. This example turns out to be surprisingly useful in applications, because (with other algorithmic machinery we can't go into here) it offers the basis for algorithms that can track moving objects by predicting where they will go next. Applications are widespread. Military applications include tracking aircraft, UFO's, missiles, etc. Civilian applications include surveillance in public places; games console interfaces that track moving people; and methods that can follow experimental mice as they move around their cages (useful for testing medicines).

## 7.2 SIMULATION

Many problems in probability can be worked out in closed form if one knows enough combinatorial mathematics, or can come up with the right trick. Textbooks are full of these, and we've seen some. Explicit formulas for probabilities are often extremely useful. But it isn't always easy or possible to find a formula for the probability of an event in a model. An alternative strategy is to build a simulation, run it many times, and count the fraction of outcomes where the event occurs. This is a simulation experiment.

### 7.2.1 Obtaining Random Numbers

Simulation is at its most useful when we try to estimate probabilities that are hard to calculate. Usually we have processes with some random component, and we want to estimate the probability of a particular outcome. To do so, we will need a supply of random numbers to simulate the random component.

#### Uniform Random Numbers

I will describe features in Matlab, because I'm used to Matlab. It's a good programming environment for numerical work, particularly for working with matrices and vectors. You can expect pretty much any programming environment to provide a random number generator that returns a number uniformly distributed in the range  $[0 - 1]$ . If you know anything about representing numbers, you'll already have spotted something funny. The number that comes out of the random number generator must be represented by a small set of bits, but almost all numbers in the interval  $[0 - 1]$  require an infinite number of bits to represent. We'll sweep this issue under the general carpet of floating point representations; it doesn't matter

for anything we need to do. The Matlab function to do this is called `rand`. It can also return matrices; for example, `rand(10, 20)` will get you a  $10 \times 20$  table of independent uniformly distributed random numbers in the range  $[0 - 1]$ .

It is useful to get a uniformly distributed random integer in a particular range (for example, you might want to choose a random element in an array). You can do so with a random number generator and a function (like Matlab's `floor`) that returns the largest integer smaller than its argument. If I want a discrete random variable with uniform distribution, maximum value 100 and minimum value 7, I habitually choose a very tiny number (for this example, say  $1e - 7$ ) and do `floor((100-7-1e-7)*rand()+7)`. I use the  $1e - 7$  because I can never remember whether `rand` produces a number no larger than one, or one that is guaranteed to be smaller than one, and I never need to care about the very tiny differences in probability caused by the  $1e-7$ . You might be able to do something cleaner if you bothered checking this point.

### Normal Random Numbers

In pretty much any programming environment, you would also expect to find a random number generator that returns a normal random variable, with mean zero and standard deviation one. In Matlab, this function is called `randn`. Conveniently, `randn(3, 4)` will give you a  $3 \times 4$  table of such numbers, which are independent of each other. As you would expect from section 13.1, to change the mean of this random number, you add a constant; to change the variance, you multiply by a constant. So in Matlab, if you want a normal random variable with mean 3 and standard deviation 4, you use `4*randn()+3`.

## 7.2.2 Computing Expectations with Simulations

Simulation is also a very good way to estimate expectations. Imagine we have a random variable  $X$  with probability distribution  $P(X)$  that takes values in some domain  $D$ . Assume that we can easily produce independent simulations, and that we wish to know  $\mathbb{E}[f]$ , the expected value of the function  $f$  under the distribution  $P(X)$ .

The weak law of large numbers tells us how to proceed. Define a new random variable  $F = f(X)$ . This has a probability distribution  $P(F)$ , which might be difficult to know. We want to estimate  $\mathbb{E}[f]$ , the expected value of the function  $f$  under the distribution  $P(X)$ . This is the same as  $\mathbb{E}[F]$ . Now if we have a set of IID samples of  $X$ , which we write  $x_i$ , then we can form a set of IID samples of  $F$  by forming  $f(x_i) = f_i$ . Write

$$F_N = \frac{\sum_{i=1}^N f_i}{N}.$$

This is a random variable, and the weak law of large numbers gives that, for any positive number  $\epsilon$

$$\lim_{N \rightarrow \infty} P(\{\|F_N - \mathbb{E}[F]\| > \epsilon\}) = 0.$$

You can interpret this as saying that, that for a set of IID random samples  $x_i$ , the

probability that

$$\frac{\sum_{i=1}^N f(x_i)}{N}$$

is very close to  $\mathbb{E}[f]$  is high for large  $N$

**Worked example 7.8** *Computing an Expectation*

Assume the random variable  $X$  is uniformly distributed in the range  $[0 - 1]$ , and the random variable  $Y$  is uniformly distributed in the range  $[0 - 10]$ .  $X$  and  $Z$  are independent. Write  $Z = (Y - 5X)^3 - X^2$ . What is  $\text{var}(\{Z\})$ ?

**Solution:** With enough work, one could probably work this out in closed form. An easy program will get a good estimate. We have that  $\text{var}(\{Z\}) = \mathbb{E}[Z^2] - \mathbb{E}[Z]^2$ . My program computed 1000 values of  $Z$  (by drawing  $X$  and  $Y$  from the appropriate random number generator, then evaluating the function). I then computed  $\mathbb{E}[Z]$  by averaging those values, and  $\mathbb{E}[Z]^2$  by averaging their squares. For a run of my program, I got  $\text{var}(\{Z\}) = 2.76 \times 10^4$ .

### 7.2.3 Computing Probabilities with Simulations

You can compute a probability using a simulation, too, because a probability can be computed by taking an expectation. Recall the property of indicator functions that

$$\mathbb{E}[\mathbb{I}_{\{\mathcal{E}\}}] = P(\mathcal{E})$$

(Section 5.3.1). This means that computing the probability of an event  $\mathcal{E}$  involves writing a function that is 1 when the event occurs, and 0 otherwise; we then estimate the expected value of that function.

The weak law of large numbers justifies this procedure. An experiment involves drawing a sample from the relevant probability distribution  $P(X)$ , then determining whether event  $\mathcal{E}$  occurred or not. We define a new random variable  $E$ , which takes the value 1 when the event  $\mathcal{E}$  occurs during our experiment (i.e. with probability  $P(\mathcal{E})$ ) and 0 when it doesn't (i.e. with probability  $P(\mathcal{E}^c)$ ). Our experiment yields a sample of this random variable. We perform  $N$  experiments yielding a set of IID samples of this distribution  $e_i$  and compute  $E_N = \frac{\sum_{i=1}^N e_i}{N}$ . From the weak law of large numbers, we have that for any  $\epsilon > 0$ ,

$$\lim_{N \rightarrow \infty} P(\{\|E_N - \mathbb{E}[E]\| \geq \epsilon\}) = 0$$

meaning that for large enough  $N$ ,  $E_N$  will be very close to  $\mathbb{E}[E] = P(\mathcal{E})$ .

**Worked example 7.9** *Computing a Probability for Multiple Coin Flips*

You flip a fair coin three times. Use a simulation to estimate the probability that you see three  $H$ 's.

**Solution:** You really should be able to work this out in closed form. But it's amusing to check with a simulation. I wrote a simple program that obtained a 1000x3 table of uniformly distributed random numbers in the range  $[0 - 1]$ . For each number, if it was greater than 0.5 I recorded an  $H$  and if it was smaller, I recorded a  $T$ . Then I counted the number of rows that had 3  $H$ 's (i.e. the expected value of the relevant indicator function). This yielded the estimate 0.127, which compares well to the right answer.

**Worked example 7.10** *Computing a Probability*

Assume the random variable  $X$  is uniformly distributed in the range  $[0 - 1]$ , and the random variable  $Y$  is uniformly distributed in the range  $[0 - 10]$ . Write  $Z = (Y - 5X)^3 - X^2$ . What is  $P(\{Z > 3\})$ ?

**Solution:** With enough work, one could probably work this out in closed form. An easy program will get a good estimate. My program computed 1000 values of  $Z$  (by drawing  $X$  and  $Y$  from the appropriate random number generator, then evaluating the function) and counted the fraction of  $Z$  values that was greater than 3 (which is the relevant indicator function). For a run of my program, I got  $P(\{Z > 3\}) \approx 0.619$

## 7.2.4 Simulation Results as Random Variables

The estimate of a probability or of an expectation that comes out of a simulation experiment is a random variable, because it is a function of random numbers. If you run the simulation again, you'll get a different value (unless you did something silly with the random number generator). Generally, you should expect this random variable to behave like a normal random variable. You can check this by constructing a histogram over a large number of runs. The mean of this random variable is the parameter you are trying to estimate. It is useful to know that this random variable tends to be normal, because it means the standard deviation of the random variable tells you a lot about the likely values you will observe.

Another helpful rule of thumb, which is almost always right, is that the standard deviation of this random variable behaves like

$$\frac{C}{\sqrt{N}}$$

where  $C$  is a constant that depends on the problem and can be very hard to evaluate, and  $N$  is the number of runs of the simulation. What this means is that if you want

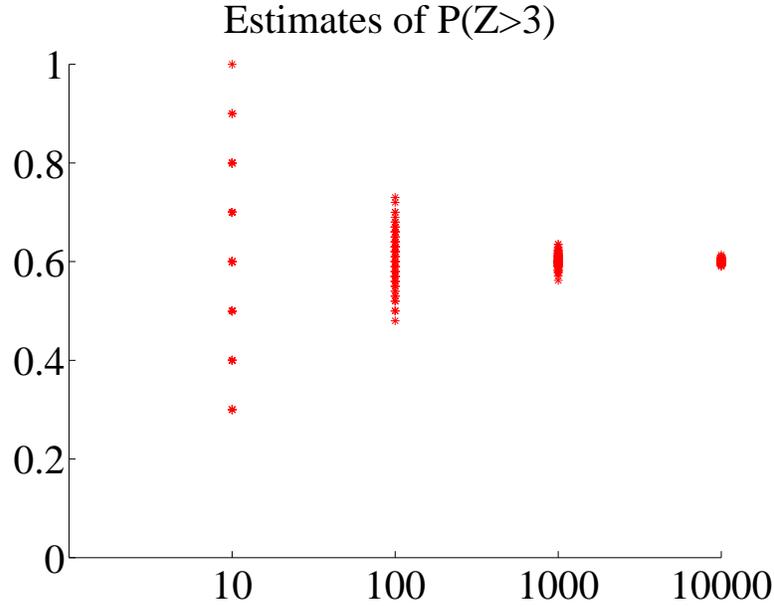


FIGURE 7.5: *Estimates of the probability from example 7.10, obtained from different runs of my simulator using different numbers of samples. In each case, I used 100 runs; the number of samples is shown on the horizontal axis. You should notice that the estimate varies pretty widely when there are only 10 samples, but the variance (equivalently, the size of the spread) goes down sharply as the number of samples increases to 1000. Because we expect these estimates to be roughly normally distributed, the variance gives a good idea of how accurate the original probability estimate is.*

to (say) double the accuracy of your estimate of the probability or the expectation, you have to run four times as many simulations. Very accurate estimates are tough to get, because they require immense numbers of simulation runs.

Figure 7.5 shows how the result of a simulation behaves when the number of runs changes. I used the simulation of example 7.10, and ran multiple experiments for each of a number of different samples (i.e. 100 experiments using 10 samples; 100 using 100 samples; and so on).

Small probabilities can be rather hard to estimate, as we would expect. In the case of example 7.10, let us estimate  $P(\{Z > 950\})$ . A few moments with a computer will show that this probability is of the order of  $1e-3$  to  $1e-4$ . I obtained a million different simulated values of  $Z$  from my program, and saw 310 where  $Z > 950$ . This means that to know this probability to, say, three digits of numerical accuracy might involve a daunting number of samples. Notice that this does not contradict the rule of thumb that the standard deviation of the random variable defined by a simulation estimate behaves like  $\frac{C}{\sqrt{N}}$ ; it's just that in this case,  $C$  is very large indeed.

## 7.3 SIMULATION EXAMPLES

Computing probabilities and expectations from simulations should be so natural to a computer science student that it can be hard to see the magic. You estimate the probability of an event  $\mathcal{E}$  by writing a program that runs  $N$  independent simulations of an experiment, counts how many times the event occurs (which we write  $\#(\mathcal{E})$ ), and reports

$$\frac{\#(\mathcal{E})}{N}$$

as an estimate of  $P(\mathcal{E})$ . This estimate will not be exact, and may be different for different runs of the program. It's often a good, simple estimate.

Choosing  $N$  depends on a lot of considerations. Generally, a larger  $N$  means a more accurate answer, and also a slower program. If  $N$  is too small, you may find that you report 1 or 0 for the probability (think of what would happen if you measured  $P(H)$  for a coin with one flip). One strategy is to run several simulations, report the mean, and use the standard deviation as some guide to the accuracy.

## 7.3.1 Simulating Experiments

**Worked example 7.11** *Getting 14's with 20-sided dice*

You throw 3 fair 20-sided dice. Estimate the probability that the sum of the faces is 14 using a simulation. Use  $N = [1e1, 1e2, 1e3, 1e4, 1e5, 1e6]$ . Which estimate is likely to be more accurate, and why?

**Solution:** You need a fairly fast computer, or this will take a long time. I ran ten versions of each experiment for  $N = [1e1, 1e2, 1e3, 1e4, 1e5, 1e6]$ , yielding ten probability estimates for each  $N$ . These were different for each version of the experiment, because the simulations are random. I got means of  $[0, 0.0030, 0.0096, 0.0100, 0.0096, 0.0098]$ , and standard deviations of  $[0.00670, 0.00330, 0.00090, 0.00020, 0.0001]$ . This suggests the true value is around 0.0098, and the estimate from  $N = 1e6$  is best. The reason that the estimate with  $N = 1e1$  is 0 is that the probability is very small, so you don't usually observe this case at all in only ten trials.

**Worked example 7.12** *Comparing simulation with computation*

You throw 3 fair six-sided dice. You wish to know the probability the sum is 3. Compare the true value of this probability with estimates from six runs of a simulation using  $N = 10000$ . What conclusions do you draw?

**Solution:** I ran six simulations with  $N = 10000$ , and got [0.0038, 0.0038, 0.0053, 0.0041, 0.0056, 0.0049]. The mean is 0.00458, and the standard deviation is 0.0007, which suggests the estimate isn't that great, but the right answer should be in the range [0.00388, 0.00528] with high probability. The true value is  $1/216 \approx 0.00463$ . The estimate is tolerable, but not super accurate.

## 7.3.2 Simulating Markov Chains

**Worked example 7.13** *Coin Flips with End Conditions*

I flip a coin repeatedly until I encounter a sequence HTHT, at which point I stop. What is the probability that I flip the coin nine times?

**Solution:** You might well be able to construct a closed form solution to this if you follow the details of example 13.1 and do quite a lot of extra work. A simulation is really straightforward to write; notice you can save time by not continuing to simulate coin flips once you've flipped past nine times. I got 0.0411 as the mean probability over 10 runs of a simulation of 1000 experiments each, with a standard deviation of 0.0056.

**Worked example 7.14** *A Queue*

A bus is supposed to arrive at a bus stop every hour for 10 hours each day. The number of people who arrive to queue at the bus stop each hour has a Poisson distribution, with intensity 4. If the bus stops, everyone gets on the bus and the number of people in the queue becomes zero. However, with probability 0.1 the bus driver decides not to stop, in which case people decide to wait. If the queue is ever longer than 15, the waiting passengers will riot (and then immediately get dragged off by the police, so the queue length goes down to zero). What is the expected time between riots?

**Solution:** I'm not sure whether one could come up with a closed form solution to this problem. A simulation is completely straightforward to write. I get a mean time of 441 hours between riots, with a standard deviation of 391. It's interesting to play around with the parameters of this problem; a less conscientious bus driver, or a higher intensity arrival distribution, lead to much more regular riots.

**Worked example 7.15** *Inventory*

A store needs to control its stock of an item. It can order stocks on Friday evenings, which will be delivered on Monday mornings. The store is old-fashioned, and open only on weekdays. On each weekday, a random number of customers comes in to buy the item. This number has a Poisson distribution, with intensity 4. If the item is present, the customer buys it, and the store makes \$100; otherwise, the customer leaves. Each evening at closing, the store loses \$10 for each unsold item on its shelves. The store's supplier insists that it order a fixed number  $k$  of items (i.e. the store must order  $k$  items each week). The store opens on a Monday with 20 items on the shelf. What  $k$  should the store use to maximise profits?

**Solution:** I'm not sure whether one could come up with a closed form solution to this problem, either. A simulation is completely straightforward to write. To choose  $k$ , you run the simulation with different  $k$  values to see what happens. I computed accumulated profits over 100 weeks for different  $k$  values, then ran the simulation five times to see which  $k$  was predicted. Results were 21, 19, 23, 20, 21. I'd choose 21 based on this information.

For example 7.15, you should plot accumulated profits. If  $k$  is small, the store doesn't lose money by storing items, but it doesn't sell as much stuff as it could; if  $k$  is large, then it can fill any order but it loses money by having stock on the shelves. A little thought will convince you that  $k$  should be near 20, because that is the expected number of customers each week, so  $k = 20$  means the store can expect to sell all its new stock. It may not be exactly 20, because it must depend a little

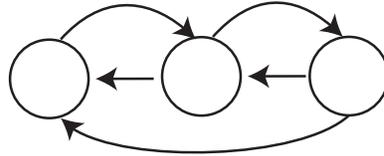


FIGURE 7.6: A very small fraction of the web, drawn to suggest a finite state machine; each state represents a page, and each directed edge represents an outgoing link. A random web surfer could either (a) follow an outgoing link, chosen at random or (b) type in the URL of any page, chosen at random. Such a surfer would see lots of pages that have many incoming links from pages that have lots of incoming links, and so on. Pages like this are likely important, so that finding important pages is analogous to simulating a random web surfer.

on the balance between the profit in selling an item and the cost of storing it. For example, if the cost of storing items is very small compared to the profit, a very large  $k$  might be a good choice. If the cost of storage is sufficiently high, it might be better to never have anything on the shelves; this point explains the absence of small stores selling PC's.

### 7.3.3 Example: Ranking the Web by Simulating a Markov Chain

Perhaps the most valuable technical question of the last thirty years has been: Which web pages are interesting? Some idea of the importance of this question is that it was only really asked about 20 years ago, and at least one gigantic technology company has been spawned by a partial answer. This answer, due to Larry Page and Sergey Brin, and widely known as PageRank, starts with a Markov chain.

Figure 7.6 shows a picture of (a very small fraction of) the world wide web. I have drawn the web using a visual metaphor that should strongly suggest a finite state machine, and so a Markov chain. Each page is a state. Directed edges from page to page represent links. I count only the first link from a page to another page. Some pages are linked, others are not. I want to know how important each page is.

One way to think about importance is to think about what a random web surfer would do. The surfer can either (a) choose one of the outgoing links on a page at random, and follow it or (b) type in the URL of a new page, and go to that instead. As Figure 7.6 suggests, it is useful to think of this as a random walk on a finite state machine. We expect that this random surfer should see a lot of pages that have lots of incoming links from other pages that have lots of incoming links that (and so on). These pages are important, because lots of pages have linked to them.

For the moment, ignore the surfer's option to type in a URL. Write  $r(i)$  for the importance of the  $i$ 'th page. We model importance as leaking from page to page across outgoing links (the same way the surfer jumps). Page  $i$  receives importance down each incoming link. The amount of importance is proportional to the amount of importance at the other end of the link, and inversely proportional to the number of links leaving that page. So a page with only one outgoing link transfers all its

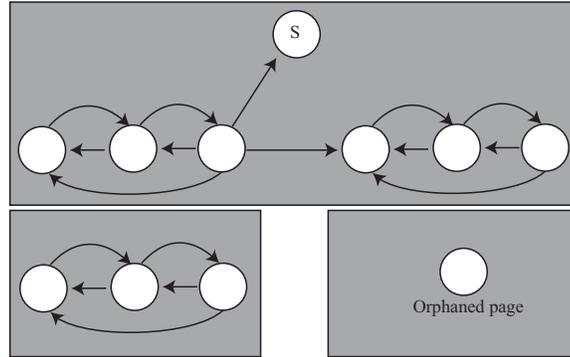


FIGURE 7.7: *The web isn't really like Figure 7.6. It's more like this figure (only bigger). In this figure, we see sections of the web that can't be reached by following links from other sections (the gray boxes); orphaned pages; and pages where a random walker would get stuck ( $S$ ). Any algorithm for assigning importance must be able to deal with these effects.*

importance down that link; and the way for a page to receive a lot of importance is for it to have a lot of important pages link to it alone. We write

$$r(j) = \sum_{i \rightarrow j} \frac{r(i)}{|i|}$$

where  $|i|$  means the total number of links pointing *out* of page  $i$ . We can stack the  $r(j)$  values into a *row* vector  $\mathbf{r}$ , and construct a matrix  $\mathcal{P}$ , where

$$p_{ij} = \begin{cases} \frac{1}{|i|} & \text{if } i \text{ points to } j \\ 0 & \text{otherwise} \end{cases}$$

With this notation, the importance vector has the property

$$\mathbf{r} = \mathbf{r}\mathcal{P}$$

and should look a bit like the stationary distribution of a random walk to you, except that  $\mathcal{P}$  isn't stochastic — there may be some rows where the row sum of  $\mathcal{P}$  is zero, because there are *no* outgoing links from that page. We can fix this easily by replacing each row that sums to zero with  $(1/n)\mathbf{1}$ , where  $n$  is the total number of pages. Call the resulting matrix  $\mathcal{G}$  (it's quite often called the **raw Google matrix**). Notice that doing this doesn't really change anything significant; pages with no outgoing links leak a tiny fraction of their importance to every other page.

Figure 7.6 isn't a particularly good model of the web (and not because it's tiny, either). Figure 7.7 shows some of the problems that occur. There are pages with no outgoing links (which we've dealt with), pages with no incoming links, and even pages with no links at all. Worse, a random walk can get trapped (in one of the gray boxes). One way to fix all this would be to construct a process that wandered around the web inserting links that clean up the structure of the

graph. This strategy is completely infeasible, because the real web is much too big. Allowing the surfer to randomly enter a URL sorts out all of these problems, because it inserts an edge of small weight from every node to every other node. Now the random walk cannot get trapped.

There are a variety of possible choices for the weight of these inserted edges. The original choice was to make each inserted edge have the same weight. Write  $\mathbf{1}$  for the  $n$  dimensional column vector containing a 1 in each component, and let  $0 < \alpha < 1$ . We can write the matrix of transition probabilities as

$$\mathcal{G}(\alpha) = \alpha \frac{(\mathbf{1}\mathbf{1}^T)}{n} + (1 - \alpha)\mathcal{G}$$

where  $\mathcal{G}$  is the original Google matrix. An alternative choice is to choose a weight for each web page, using anything from advertising revenues to page visit statistics to thaumaturgy (Google keeps quiet about the details). Write this weight vector  $\mathbf{v}$ , and require that  $\mathbf{1}^T \mathbf{v} = 1$  (i.e. the coefficients sum to one). Then we could have

$$\mathcal{G}(\alpha, \mathbf{v}) = \alpha \frac{(\mathbf{1}\mathbf{v}^T)}{n} + (1 - \alpha)\mathcal{G}.$$

Now the importance vector  $\mathbf{r}$  is the (unique, though I won't prove this) *row* vector  $\mathbf{r}$  such that

$$\mathbf{r} = \mathbf{r}\mathcal{G}(\alpha, \mathbf{v}).$$

How do we compute this vector? One natural algorithm is to start with some initial estimate, and propagate it. We write  $\mathbf{r}^{(k)}$  for the estimated importance after the  $k$ 'th step. We define updates by

$$(\mathbf{r}^{(k)}) = (\mathbf{r}^{(k-1)})\mathcal{G}(\alpha, \mathbf{v}).$$

We can't compute this directly, either, because  $\mathcal{G}(\alpha, \mathbf{v})$  is unreasonably big so (a) we can't form or store  $\mathcal{G}(\alpha, \mathbf{v})$  and (b) we can't multiply by it either. But we could estimate  $\mathbf{r}$  with a random walk, because  $\mathbf{r}$  is the stationary distribution of a Markov chain. If we simulate this walk for many steps, the probability that the simulation is in state  $j$  should be  $r(j)$ , at least approximately.

This simulation is easy to build. Imagine our random walking bug sits on a web page. At each time step, it transitions to a new page by either (a) picking from all existing pages at random, using  $\mathbf{v}$  as a probability distribution on the pages (which it does with probability  $\alpha$ ); or (b) chooses one of the outgoing links uniformly and at random, and follows it (which it does with probability  $1 - \alpha$ ). The stationary distribution of this random walk is  $\mathbf{r}$ . Another fact that I shall not prove is that, when  $\alpha$  is sufficiently large, this random walk very quickly "forgets" its initial distribution. As a result, you can estimate the importance of web pages by starting this random walk in a random location; letting it run for a bit; then stopping it, and collecting the page you stopped on. The pages you see like this are independent, identically distributed samples from  $\mathbf{r}$ ; so the ones you see more often are more important, and the ones you see less often are less important.

## 7.3.4 Example: Simulating a Complicated Game

I will build several examples around a highly simplified version of a real card game. This game is Magic: The Gathering, and is protected by a variety of trademarks, etc. My version — MTGDFAF — isn't very interesting as a game, but is simple enough to be studied, and interesting enough it casts some light on the real game. The game is played with decks of 60 cards. There are two types of card: Lands, and Spells. Lands can be placed on the play table and stay there permanently; Spells are played and then disappear. A Land on the table can be “tapped” or “untapped”. Players take turns. Each player draws a hand of seven cards from a shuffled deck. In each turn, a player first untaps any Lands on the table, then draws a card, then plays a land onto the table (if the player has one in hand to play), then finally can play one or more spells. Each spell has a fixed cost (of  $1, \dots, 10$ ), and this cost is played by “tapping” a land (which is not untapped until the start of the next turn). This means that the player can cast only cheap spells in the early turns of the game, and expensive spells in the later turns.

**Worked example 7.16** *MTGDFAF — The number of lands*

Assume a deck of 60 cards has 24 Lands. It is properly shuffled, and you draw seven cards. You could draw  $0, \dots, 7$  Lands. Estimate the probability for each, using a simulation. Furthermore, estimate the error in your estimates.

**Solution:** The matlab function `randperm` produces a random permutation of given length. This means you can use it to simulate a shuffle of a deck, as in listing 7.1. I then drew 10, 000 random hands of seven cards, and counted how many times I got each number. Finally, to get an estimate of the error, I repeated this experiment 10 times and computed the standard deviation of each estimate of probability. This produced

0.0218 0.1215 0.2706 0.3082 0.1956 0.0686 0.0125 0.0012

for the probabilities (for 0 to 7, increasing number of lands to the right) and

0.0015 0.0037 0.0039 0.0058 0.0027 0.0032 0.0005 0.0004

for the standard deviations of these estimates.

Listing 7.1: Matlab code used to simulate the number of lands

```

simcards=[ones(24, 1); zeros(36, 1)]
% 1 if land, 0 otherwise
ninsim=10000;
nsims=10;
counts=zeros(nsims, 8);
for i=1:10
    for j=1:10000
        shuffle=randperm(60);
        hand=simcards(shuffle(1:7));
        %useful matlab trick here
        nlands=sum(hand);
        %ie number of lands
        counts(i, 1+nlands)=...
            counts(i, 1+nlands)+1;
        % number of lands could be zero
    end
end
probs=counts/ninsim;
mean(probs)
std(probs)
%%

```

**Worked example 7.17** *MTGDAF — The number of lands*

What happens to the probability of getting different numbers of lands if you put only 15 Lands in a deck of 60? It is properly shuffled, and you draw seven cards. You could draw 0, . . . , 7 Lands. Estimate the probability for each, using a simulation. Furthermore, estimate the error in your estimates.

**Solution:** You can change one line in the listing to get

```
0.1159 0.3215 0.3308 0.1749 0.0489 0.0075 0.0006 0.0000
```

for the probabilities (for 0 to 7, increasing number of lands to the right) and

```
0.0034 0.0050 0.0054 0.0047 0.0019 0.0006 0.0003 0.0000
```

for the standard deviations of these estimates.

**Worked example 7.18** *MTGDAF — Playing spells*

Assume you have a deck of 24 Lands, 10 Spells of cost 1, 10 Spells of cost 2, 10 Spells of cost 3, 2 Spells of cost 4, 2 Spells of cost 5, and 2 Spells of cost 6. Assume you always only play the cheapest spell in your hand (i.e. you never play two spells). What is the probability you will be able to play at least one spell on each of the first four turns?

**Solution:** This simulation requires just a little more care. You draw the hand, then simulate the first four turns. In each turn, you can only play a spell whose cost you can pay, and only if you have it. I used the matlab of listing 7.2 and listing 7.3; I found the probability to be 0.64 with standard deviation 0.01. Of course, my code might be wrong....

**Worked example 7.19** *MTGDAF — Playing spells*

Now we use a different distribution of cards. Assume you have a deck of 20 Lands, 9 Spells of cost 1, 5 Spells of cost 2, 5 Spells of cost 3, 5 Spells of cost 4, 5 Spells of cost 5, and 11 Spells of cost 6. Assume you always only play the cheapest spell in your hand (i.e. you never play two spells). What is the probability you will be able to play at least one spell on each of the first four turns?

**Solution:** This simulation requires just a little more care. You draw the hand, then simulate the first four turns. In each turn, you can only play a spell whose cost you can pay, and only if you have it. I found the probability to be 0.33 with standard deviation 0.05. Of course, my code might be wrong....

One engaging feature of the real game that is revealed by these very simple simulations is the tension between a players goals. The player would like to have few lands — so as to have lots of spells — but doing so means that there’s a bigger chance of not being able to play a spell. Similarly, a player would like to have lots of powerful (=expensive) spells, but doing so means there’s a bigger chance of not being able to play a spell. Players of the real game spend baffling amounts of time arguing with one another about the appropriate choices for a good set of cards.

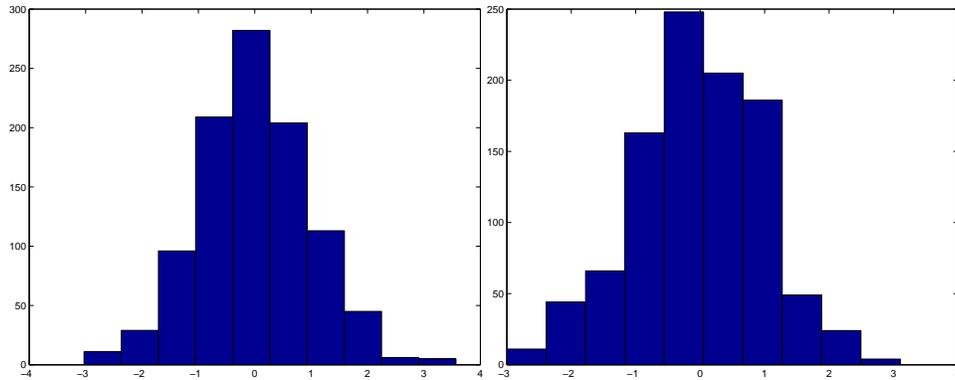


FIGURE 7.8: Estimates of probabilities produced by simulation typically behave like normal data. On the **left**, I show a histogram of probabilities of having a hand of 3 Lands in the simulation of example 7.16; these are plotted in standard coordinates. On the **right**, I show a histogram of probability of playing a spell in each of the first four turns (example 7.19), from 1000 simulation experiments; again, these are plotted in standard coordinates. Compare these to the standard normal histograms of the previous chapter.

**Worked example 7.20** *MTGDAF* — How long before you can play a spell of cost 3?

Assume you have a deck of 15 Lands, 15 Spells of cost 1, 14 Spells of cost 2, 10 Spells of cost 3, 2 Spells of cost 4, 2 Spells of cost 5, and 2 Spells of cost 6. What is the expected number of turns before you can play a spell of cost 3? Assume you always play a land if you can.

**Solution:** I get 6.3, with a standard deviation of 0.1. The problem is it can take quite a large number of turns to get three lands out. I used the code of listings 7.4 and 7.5

Recall that you can reasonably expect that the probability you compute from a simulation behaves like normal data. Run a simulation experiment a large number of times and construct a data set whose entries are the probability from each experiment. This data should be normal. This means that, if you subtract the mean and divide by the standard deviation, you should get a histogram that looks like the standard normal curve. Figure 7.8 shows some examples of this effect. This effect is important, because it means that the right answer should be very few standard deviations away from the mean you compute — the standard deviation gives you quite a good idea of the accuracy of your estimate.

## 7.4 YOU SHOULD

## 7.4.1 remember

- A Markov chain is a sequence of random variables, with particular independence properties. It is valuable to think of a Markov chain as being produced by a biased random walk on a directed graph. In this view, the values of the random variables are given by the labels at each node of the graph.
- In some Markov chains, there is a special state that the chain cannot leave (an absorbing state). Not all chains have such states.
- Markov chains that satisfy reasonable structural assumptions have a stationary distribution. The random variables emitted by the chain will be samples (but not necessarily independent samples) of that stationary distribution once the chain has changed state often enough.
- One can often write expressions for various probabilities associated with chains; these are usually recurrence relations.
- Simulating Markov chains is usually straightforward, and often extremely useful.
- Simulation is a powerful method to estimate the probabilities of particular events. The method can be used to estimate expectations as well.
- In general, one simulates a random process with a program, then counts the different outputs encountered in a large number of runs of the simulation. In most cases, for a sufficiently large number of runs, the number produced by a simulation is the value of a normal random variable whose expected value is the number sought and whose variance is  $k/N$  where  $k$  is an unknown constant and  $N$  is the number of simulation runs. This means that the accuracy of simulation based estimates improves increasingly slowly as  $N$  increases.
- Simulation methods can be used to estimate a wide range of useful properties, including the importance of web documents and the value of various strategies in games.

## PROBLEMS

**7.1. Multiple coin flips:**

- (a) For example ??, show that  $P(5) = (3/32)$  by directly writing out the sequences of flips, and summing their probabilities.
- (b) Now check the recurrence relation  $P(N) = (1/2)P(N-1) + (1/4)P(N-2)$  for the case  $N = 5$ .
- (c) What is  $P(7)$ ?

**7.2. Multiple die rolls:** You roll a fair die until you see a 5, then a 6; after that, you stop. Write  $P(N)$  for the probability that you roll the die  $N$  times.

- (a) What is  $P(1)$ ?
- (b) Show that  $P(2) = (1/36)$ .

- (c) Draw a finite state machine encoding all the sequences of die rolls that you could encounter. Don't write the events on the edges; instead, write their probabilities. There are 5 ways not to get a 5, but only one probability, so this simplifies the drawing.
- (d) Show that  $P(3) = (1/36)$ .
- (e) Now use your finite state machine to argue that  $P(N) = (5/6)P(N-1) + (25/36)P(N-2)$ .
- 7.3. More complicated multiple coin flips:** You flip a fair coin until you see either *HTH* or *THT*, and then you stop. We will compute a recurrence relation for  $P(N)$ .
- (a) Figure ?? shows a finite state machine. Check that this finite state machine represents all coin sequences that you will encounter.
- (b) Write  $\Sigma_N$  for some string of length  $N$  accepted by this finite state machine. Use this finite state machine to argue that  $\text{Sigma}_N$  has one of four forms:
1.  $TT\Sigma_{N-2}$
  2.  $HH\Sigma_{N-3}$
  3.  $THH\Sigma_{N-2}$
  4.  $HTT\Sigma_{N-3}$
- (c) Now use this argument to show that  $P(N) = (1/2)P(N-2) + (1/4)P(N-3)$ .

**7.4. The gambler's ruin:**

- (a) Show that you can rearrange the recurrence relation of example 13.1 to get

$$p_{s+1} - p_s = \frac{(1-p)}{p} (p_s - p_{s-1}).$$

Now show that this means that

$$p_{s+1} - p_s = \left( \frac{(1-p)}{p} \right)^2 (p_{s-1} - p_{s-2})$$

so that

$$\begin{aligned} p_{s+1} - p_s &= \left( \frac{(1-p)}{p} \right)^s (p_1 - p_0) \\ &= \left( \frac{(1-p)}{p} \right)^s (p_1 - 1). \end{aligned}$$

- (b) Now we need a simple result about series. Assume I have a series  $u_k$ ,  $k \geq 0$ , with the property that

$$u_k - u_{k-1} = cr^{k-1}.$$

Show that

$$u_k - u_0 = c \left( \frac{r^k - 1}{r - 1} \right).$$

(c) Use the results of the last two steps to show that

$$p_s - 1 = (p_1 - 1) \left( \frac{\left(\frac{1-p}{p}\right)^s - 1}{\left(\frac{1-p}{p}\right) - 1} \right)$$

(d) Use the fact that  $p_j = 0$  and the result of the last exercise to show

$$(p_1 - 1) = \frac{-1}{\left( \frac{\left(\frac{1-p}{p}\right)^j - 1}{\left(\frac{1-p}{p}\right) - 1} \right)}$$

(e) Use the results of the previous exercises to show that

$$p_s = \frac{\left(\frac{1-p}{p}\right)^j - \left(\frac{1-p}{p}\right)^s}{\left(\frac{1-p}{p}\right)^j - 1}$$

## PROGRAMMING EXERCISES

- 7.5.** Assume you have an MTG-DAF deck of 24 Lands, 10 Spells of cost 1, 10 Spells of cost 2, 10 Spells of cost 3, 2 Spells of cost 4, 2 Spells of cost 5, and 2 Spells of cost 6. It is properly shuffled, and you draw seven cards. At each turn, you play a land if it is in your hand, and you always only play the most expensive spell in your hand that you are able to play and you never play two spells.
- Write a program to estimate the probability that the first time you play a spell is on turn 4.
  - Write a program to estimate the probability that the first spell you play has cost 4.
  - Extend your program to prepare a 6 x 6 table of turns by spells. In the  $t$ ,  $s$ 'th cell of the table, your program should place the probability that the first spell you play is played on turn  $t$ , and has cost  $s$ . Notice that many cells easily contain a zero.
- 7.6.** Assume you have an MTG-DAF deck of 10 Lands, 10 Spells of cost 1, 10 Spells of cost 2, 10 Spells of cost 3, 2 Spells of cost 4, 2 Spells of cost 5, and 16 Spells of cost 6. It is properly shuffled, and you draw seven cards. At each turn, you play a land if it is in your hand, and you always only play the cheapest spell in your hand that you are able to play and you never play two spells.
- Write a program to estimate the probability that the first time you play a spell is on turn 4.
  - Write a program to estimate the probability that the first spell you play has cost 4.
  - Extend your program to prepare a 6 x 6 table of turns by spells. In the  $t$ ,  $s$ 'th cell of the table, your program should place the probability that the first spell you play is played on turn  $t$ , and has cost  $s$ . Notice that many cells easily contain a zero.
- 7.7.** You have three MTG-DAF decks. Deck 1 has 24 Lands, 10 Spells of cost 1, 10 Spells of cost 2, 10 Spells of cost 3, 2 Spells of cost 4, 2 Spells of cost 5, and 2 Spells of cost 6. Deck 2 has 10 Lands, 10 Spells of cost 1, 10 Spells of cost 2, 10 Spells of cost 3, 2 Spells of cost 4, 2 Spells of cost 5, and 16 Spells

of cost 6. Deck 3 has 44 Lands, and 16 Spells of cost 6. These decks will be played by players who rigorously implement the following rules: At each turn, a player plays a land if it is in hand, and always plays only the most expensive spell in hand that can be played, and never plays two spells. The first player to play a spell of cost greater than 4 wins. Write a program to estimate the probability of winning for each deck in each possible matchup. In the matchups, you should consider (for example) Deck 1 vs Deck 2, 1 plays first, AND Deck 1 vs Deck 2, 2 plays first.

Listing 7.2: Matlab code used to simulate the four turns

```

simcards=[zeros(24, 1); ones(10, 1);...
          2*ones(10, 1);3*ones(10, 1); ...
          4*ones(2, 1); 5*ones(2, 1); 6*ones(2, 1)];
nsims=10;
ninsim=1000;
counts=zeros(nsims, 1);
for i=1:nsims
    for j=1:ninsim
        % draw a hand
        shuffle=randperm(60);
        hand=simcards(shuffle(1:7));
        %reorganize the hand
        cleanhand=zeros(7, 1);
        for k=1:7
            cleanhand(hand(k)+1)=cleanhand(hand(k)+1)+1;
            % ie count of lands, spells, by cost
        end
        landsontable=0;
        [playedspell1, landsontable, cleanhand]=...
            playround(landsontable, cleanhand, shuffle, ...
                simcards, 1);
        [playedspell2, landsontable, cleanhand]=...
            playround(landsontable, cleanhand, shuffle, ...
                simcards, 2);
        [playedspell3, landsontable, cleanhand]=...
            playround(landsontable, cleanhand, shuffle, ...
                simcards, 3);
        [playedspell4, landsontable, cleanhand]=...
            playround(landsontable, cleanhand, shuffle, ...
                simcards, 4);
        counts(i)=counts(i)+...
            playedspell1*playedspell2*...
            playedspell3*playedspell4;
    end
end
end

```

Listing 7.3: Matlab code used to simulate playing a turn

```

function [playedspell, landsontable, cleanhand]=...
    playround(landsontable, cleanhand, shuffle, simcards, ...
    turn)
    % draw
ncard=simcards( shuffle(7+turn) );
cleanhand(ncard+1)=cleanhand(ncard+1)+1;
% play land
if cleanhand(1)>0
    landsontable=landsontable+1;
    cleanhand(1)=cleanhand(1)-1;
end
playedspell=0;
if landsontable>0
    i=1; done=0;
    while done==0
        if cleanhand(i)>0
            cleanhand(i)=cleanhand(i)-1;
            playedspell=1;
            done=1;
        else
            i=i+1;
            if i>landsontable
                done=1;
            end
        end
    end
end
end

```

Listing 7.4: Matlab code used to estimate number of turns before you can play a spell of cost 3

```

simcards=[zeros(15, 1); ones(15, 1);...
          2*ones(14, 1);3*ones(10, 1); ...
          4*ones(2, 1); 5*ones(2, 1); 6*ones(2, 1)];
nsims=10;
ninsim=1000;
counts=zeros(nsims, 1);
for i=1:nsims
    for j=1:ninsim
        % draw a hand
        shuffle=randperm(60);
        hand=simcards(shuffle(1:7));
        %reorganize the hand
        cleanhand=zeros(7, 1);
        for k=1:7
            cleanhand(hand(k)+1)=cleanhand(hand(k)+1)+1;
            % ie count of lands, spells, by cost
        end
        landsontable=0;
        k=0; played3spell=0;
        while played3spell==0;
            [played3spell, landsontable, cleanhand]=...
                play3round(landsontable, cleanhand, shuffle, ...
                    simcards, k+1);
            k=k+1;
        end
        counts(i)=counts(i)+k;
    end
counts(i)=counts(i)/ninsim;
end

```

Listing 7.5: Matlab code used to simulate a turn to estimate the number of turns before you can play a spell of cost 3

```
function [played3spell , landsontable , cleanhand]=...
    play3round(landsontable , cleanhand , shuffle , simcards , ...
    turn)
    % draw
ncard=simcards( shuffle(7+turn) );
cleanhand(ncard+1)=cleanhand(ncard+1)+1;
% play land
if cleanhand(1)>0
    landsontable=landsontable+1;
    cleanhand(1)=cleanhand(1)-1;
end
played3spell=0;
if (landsontable >=3)&&(cleanhand(4)>0)
    played3spell=1;
end
```

# Inference: Making Point Estimates

**Inference** is the process of drawing conclusions from data. One form of inference is to estimate a number, or set of numbers, that describes a dataset. The result is known as a **point estimate**. An alternative is to estimate an interval within which a number lies, with some degree of certainty. Such estimates are known as **interval estimates**. Finally, one might wish to assess the extent to which a body of evidence supports rejecting an hypothesis — known as **hypothesis testing**. In this chapter, we deal with point estimates. In the following chapter, we deal with interval estimates and hypothesis testing, which require an understanding of point estimates. There are two, somewhat distinct, situations in which we could make point estimates.

In the first, we have a dataset  $\{\mathbf{x}\}$ , and a probability model we believe applies to that dataset. But we need to select appropriate values of the parameters to ensure that the model describes the data. For example, we might have a set of  $N$  coin flips which we believe to be independent and identically distributed. Of these,  $k$  flips came up  $H$ . We know that a binomial distribution with  $p(H) = p$  is a good model — but what value of  $p$  should we use? Your intuition is likely to suggest using  $k/N$ , but we'd like a more robust procedure than guessing. We need an inference procedure to obtain the unknown parameter from the data. Notice that this will be an estimate, rather than the “true” value. As we shall see, there is more than one possible procedure to apply, depending to some extent on the problem. In some cases (section 8.1), we estimate parameter values based solely on data; in others (section 8.2), we are able to use prior information about the parameters to affect the estimate.

In the second situation, we want to know some property of a population. For example, we may wish to know the mean weight of a person, or the mean response of a mouse to a drug. It would be a stretch to describe this population with one of the probability models that we have seen. In principle, the number we want is not even necessarily random; in principle, we could measure everyone on the planet and average the weights. In practice, this doesn't make sense, for quite straightforward reasons. You can't really weigh every person, or dose every mouse, on the planet. Instead, to estimate this property, we obtain a sample (some people; some mice; etc.) of the population, and estimate the property from the sample. There is now an important problem. Different samples lead to different estimates of the property. We will arrange to have the sample drawn randomly from the population, so the sample we see represents the value of a set of random variables. If you followed the proof of the weak law of large numbers, you should suspect that the mean of this sample could be a good estimate of the population mean. This turns out to be the case (section 13.1). However, there is some random error in the estimate, and we can tell (on average) how large the error caused by random sampling could be.

## 8.1 ESTIMATING MODEL PARAMETERS WITH MAXIMUM LIKELIHOOD

Assume we have a dataset  $\mathcal{D} = \{\mathbf{x}\}$ , and a probability model we believe applies to that dataset. Generally, application logic suggests the type of model (i.e. normal probability density; Poisson probability; geometric probability; and so on). But usually, we do not know the parameters of the model — for example, the mean and standard deviation of a normal distribution; the intensity of a poisson distribution; and so on. Notice that this situation is unlike what we have seen to date. In chapter 6, we assumed that we knew  $\theta$ , and could then use the model to assign a probability to a set of data items  $\mathcal{D}$ . Here we *know* the value of  $\mathcal{D}$ , but don't know  $\theta$ . Our model will be better or worse depending on how well we choose the parameters. We need a strategy to estimate the parameters of a model from a sample dataset. Notice how each of the following examples fits this pattern.

**Example: 8.1** *Inferring  $p$  from repeated flips — binomial*

We could flip the coin  $N$  times, and count the number of heads  $k$ . We know that an appropriate probability model for a set of independent coin flips is the binomial model  $P(k; N, p)$ . But we do not know  $p$ , which is the parameter — we need a strategy to extract a value of  $p$  from the data.

**Example: 8.2** *Inferring  $p$  from repeated flips — geometric*

We could flip the coin repeatedly until we see a head. We know that, in this case, the number of flips has the geometric distribution with parameter  $p$ . In this case, the data is a sequence of  $T$ 's with a final  $H$  from the coin flips. There are  $N$  flips (or terms) and the last flip is a head. We know that an appropriate probability model is the geometric distribution  $P_g(N; p)$ . But we do not know  $p$ , which is the parameter — we need a strategy to extract a value of  $p$  from the data.

**Example: 8.3** *Inferring the intensity of spam — poisson*

It is reasonable to assume that the number of spam emails one gets in an hour has a Poisson distribution. But what is the intensity parameter  $\lambda$ ? We could count the number of spam emails that arrive in each of a set of distinct hours, giving a dataset of counts  $\mathcal{D}$ . We need a strategy to wrestle an estimate of  $\lambda$  from this dataset.

**Example: 8.4** *Inferring the mean and standard deviation of normal data*

Imagine we know for some reason that our data is well described by a normal distribution. We could ask what is the mean and standard deviation of the normal distribution that best represents the data?

Write  $\theta$  for the parameters of our model. If we knew  $\theta$ , then the probability of observing the data  $\mathcal{D}$  would be  $P(\mathcal{D}|\theta)$ . We know  $\mathcal{D}$ , and we don't know  $\theta$ , so the value of  $P(\mathcal{D}|\theta)$  is a function of  $\theta$ .

**Definition: 8.1** *Likelihood*

The function  $P(\mathcal{D}|\theta)$ , which is a function of  $\theta$ , is known as the **likelihood** of the data  $D$ , and is often written  $\mathcal{L}(\theta)$  (or  $\mathcal{L}(\theta; \mathcal{D})$  if you want to remember that data is involved).

Indexing: Likelihood

Indexing: Likelihood

## 8.1.1 The Maximum Likelihood Principle

We need a “reasonable” procedure to choose a value of  $\theta$  to report. One — and I stress this is not the only one — is the **maximum likelihood principle**.

**Definition: 8.2** *The maximum likelihood principle*

Choose  $\theta$  such that  $\mathcal{L}(\theta) = P(\mathcal{D}|\theta)$  is maximised, as a function of  $\theta$ .

Indexing: The maximum likelihood principle

Indexing: The maximum likelihood principle

In words, this means: Choose the parameter such that the probability of observing the data you actually see, is maximised. This should strike you as being a reasonable choice. You should also be aware that this is not the only possible choice (we'll see another one in section 8.2).

For the examples we work with, the data will be **independent and identically distributed** or **IID**. This means that each data item is an independently obtained sample from the same probability distribution (see section ??). In turn, this means that the likelihood is a product of terms, one for each data item, which we can write as

$$\mathcal{L}(\theta) = P(\mathcal{D}|\theta) = \prod_{i \in \text{dataset}} P(d_i|\theta).$$

It is traditional to write  $\theta$  for any set of parameters that are unknown. There are two, distinct, important concepts we must work with. One is the unknown parameter(s), which we will write  $\theta$ . The other is the *estimate* of the value of that parameter, which we will write  $\hat{\theta}$ . This estimate is the best we can do — it may not be the “true” value of the parameter.

**Worked example 8.1** *Inferring  $p(H)$  for a coin from flips using a binomial model*

In  $N$  independent coin flips, you observe  $k$  heads. Use the maximum likelihood principle to infer  $p(H)$ .

**Solution:** The coin has  $\theta = p(H)$ , which is the unknown parameter. We know that an appropriate probability model is the binomial model  $P(k; N, \theta)$ . We have that

$$\mathcal{L}(\theta) = P(\mathcal{D}|\theta) = P_b(k; N, \theta) = \binom{N}{k} \theta^k (1 - \theta)^{(N-k)}$$

which is a function of  $\theta$  — the unknown probability that a coin comes up heads;  $k$  and  $N$  are known. We must find the value of  $\theta$  that maximizes this expression. Now the maximum occurs when

$$\frac{\partial \mathcal{L}(\theta)}{\partial \theta} = 0.$$

We have

$$\frac{\partial \mathcal{L}(\theta)}{\partial \theta} = \binom{N}{k} \left( k\theta^{k-1}(1 - \theta)^{(N-k)} - \theta^k(N - k)(1 - \theta)^{(N-k-1)} \right)$$

and this is zero when

$$k\theta^{k-1}(1 - \theta)^{(N-k)} = \theta^k(N - k)(1 - \theta)^{(N-k-1)}$$

so the maximum occurs when

$$k(1 - \theta) = \theta(N - k).$$

This means the maximum likelihood estimate is

$$\hat{\theta} = \frac{k}{N}$$

which is what we guessed would happen, but now we know why that guess “makes sense”.

**Worked example 8.2** *Inferring  $p(H)$  from coin flips using a geometric model*

You flip a coin  $N$  times, stopping when you see a head. Use the maximum likelihood principle to infer  $p(H)$  for the coin.

**Solution:** The coin has  $\theta = p(H)$ , which is the unknown parameter. We know that an appropriate probability model is the geometric model  $P_g(N; \theta)$ . We have that

$$\mathcal{L}(\theta) = P(\mathcal{D}|\theta) = P_g(N; \theta) = (1 - \theta)^{(N-1)}\theta$$

which is a function of  $\theta$  — the unknown probability that a coin comes up heads;  $N$  is known. We must find the value of  $\theta$  that maximizes this expression. Now the maximum occurs when

$$\frac{\partial \mathcal{L}(\theta)}{\partial \theta} = 0 = ((1 - \theta)^{(N-1)} - (N - 1)(1 - \theta)^{(N-2)}\theta)$$

So the maximum likelihood estimate is

$$\hat{\theta} = \frac{1}{N}.$$

We didn't guess this.

**Worked example 8.3** *Inferring die probabilities from multiple rolls and a multinomial distribution*

You throw a die  $N$  times, and see  $n_1$  ones, ... and  $n_6$  sixes. Write  $p_1, \dots, p_6$  for the probabilities that the die comes up one, ..., six. Use the maximum likelihood principle to estimate  $p_1, \dots, p_6$ .

**Solution:** The data are  $n, n_1, \dots, n_6$ . The parameters are  $\theta = (p_1, \dots, p_6)$ .  $P(\mathcal{D}|\theta)$  comes from the multinomial distribution. In particular,

$$\mathcal{L}(\theta) = P(\mathcal{D}|\theta) = \frac{n!}{n_1! \dots n_6!} p_1^{n_1} p_2^{n_2} \dots p_6^{n_6}$$

which is a function of  $\theta = (p_1, \dots, p_6)$ . Now we want to maximize this function by choice of  $\theta$ . Notice that we could do this by simply making all  $p_i$  very large — but this omits a fact, which is that  $p_1 + p_2 + p_3 + p_4 + p_5 + p_6 = 1$ . So we substitute using  $p_6 = 1 - p_1 - p_2 - p_3 - p_4 - p_5$  (there are other, neater, ways of dealing with this issue, but they take more background knowledge). At the maximum, we must have that for all  $i$ ,

$$\frac{\partial \mathcal{L}(\theta)}{\partial p_i} = 0$$

which means that, for  $p_i$ , we must have

$$n_i p_i^{(n_i-1)} (1 - p_1 - p_2 - p_3 - p_4 - p_5)^{n_6} - p_i^{n_i} n_6 (1 - p_1 - p_2 - p_3 - p_4 - p_5)^{(n_6-1)} = 0$$

so that, for each  $p_i$ , we have

$$n_i (1 - p_1 - p_2 - p_3 - p_4 - p_5) - n_6 p_i = 0$$

or

$$\frac{p_i}{1 - p_1 - p_2 - p_3 - p_4 - p_5} = \frac{n_i}{n_6}.$$

You can check that this equation is solved by

$$\hat{\theta} = \frac{1}{(n_1 + n_2 + n_3 + n_4 + n_5 + n_6)} (n_1, n_2, n_3, n_4, n_5, n_6)$$

The logarithm is a monotonic function (i.e. if  $x > 0$ ,  $y > 0$ ,  $x > y$ , then  $\log(x) > \log(y)$ ). This means that the values of  $\theta$  that maximise the log-likelihood are the same as the values that maximise the likelihood. This observation is very useful, because it allows us to transform a product into a sum. The derivative of a product involves numerous terms; the derivative of a sum is easy to take. We have

$$\log P(\mathcal{D}|\theta) = \log \prod_{i \in \text{dataset}} P(d_i|\theta) = \sum_{i \in \text{dataset}} \log P(d_i|\theta)$$

and in some cases,  $\log P(d_i|\theta)$  takes a convenient, easy form.

**Definition: 8.3** *The log-likelihood of a dataset under a model*

is a function of the unknown parameters, and you will often see it written as

$$\log \mathcal{L}(\theta) = \sum_{i \in \text{dataset}} \log P(d_i|\theta).$$

Indexing: The log-likelihood of a dataset under a model

Indexing: The log-likelihood of a dataset under a model

**Worked example 8.4** *Poisson distributions*

You observe  $N$  intervals, each of the same, fixed length (in time, or space). You know that, in these intervals, events occur with a Poisson distribution (for example, you might be observing Prussian officers being kicked by horses, or telemarketer calls...). You know also that the intensity of the Poisson distribution is the same for each observation. The number of events you observe in the  $i$ 'th interval is  $n_i$ . What is the intensity,  $\lambda$ ?

**Solution:** The likelihood is

$$\mathcal{L}(\theta) = \prod_{i \in \text{intervals}} P(\{n_i \text{ events}\}|\theta) = \prod_{i \in \text{intervals}} \frac{\theta^{n_i} e^{-\theta}}{n_i!}.$$

It will be easier to work with logs. The log-likelihood is

$$\log \mathcal{L}(\theta) = \sum_i (n_i \log \theta - \theta - \log n_i!)$$

so that we must solve

$$\frac{\partial \log \mathcal{L}(\theta)}{\partial \theta} = \sum_i \left( \frac{n_i}{\theta} - 1 \right) = 0$$

which yields a maximum likelihood estimate of

$$\hat{\theta} = \frac{\sum_i n_i}{N}$$

**Worked example 8.5** *The intensity of swearing*

A famously swear-y politician gives a talk. You listen to the talk, and for each of 30 intervals 1 minute long, you record the number of swearwords. You record this as a histogram (i.e. you count the number of intervals with zero swear words, with one, etc.). For the first 10 intervals, you see

no. of swear words	no. of intervals
0	5
1	2
2	2
3	1
4	0

and for the following 20 intervals, you see

no. of swear words	no. of intervals
0	9
1	5
2	3
3	2
4	1

Assume that the politician's use of swearwords is Poisson. What is the intensity using the first 10 intervals? the second 20 intervals? all the intervals? why are they different?

**Solution:** Use the expression from worked example 8.4 to find

$$\begin{aligned}\hat{\lambda}_{10} &= \frac{\text{total number of swearwords}}{\text{number of intervals}} \\ &= \frac{9}{10}\end{aligned}$$

$$\begin{aligned}\hat{\lambda}_{20} &= \frac{\text{total number of swearwords}}{\text{number of intervals}} \\ &= \frac{21}{20}\end{aligned}$$

$$\begin{aligned}\hat{\lambda}_{30} &= \frac{\text{total number of swearwords}}{\text{number of intervals}} \\ &= \frac{30}{30}.\end{aligned}$$

These are different because the maximum likelihood estimate is an *estimate* — we can't expect to recover the exact value from a dataset. Notice, however, that the estimates are quite close.

**Worked example 8.6** *Normal distributions*

Assume we have  $x_1, \dots, x_N$ , and we wish to model these data with a normal distribution. Use the maximum likelihood principle to estimate the mean of that normal distribution.

**Solution:** The likelihood of a set of data values under the normal distribution with unknown mean  $\theta$  and standard deviation  $\sigma$  is

$$\begin{aligned}\mathcal{L}(\theta) &= P(x_1, \dots, x_N | \theta, \sigma) \\ &= P(x_1 | \theta, \sigma) P(x_2 | \theta, \sigma) \dots P(x_N | \theta, \sigma) \\ &= \prod_{i=1}^N \frac{1}{\sqrt{2\pi}\sigma} \exp\left(-\frac{(x_i - \theta)^2}{2\sigma^2}\right)\end{aligned}$$

and this expression is a moderate nuisance to work with. The log of the likelihood is

$$\log \mathcal{L}(\theta) = \left( \sum_{i=1}^N -\frac{(x_i - \theta)^2}{2\sigma^2} \right) + \text{term not depending on } \theta.$$

We can find the maximum by differentiating wrt  $\theta$  and setting to zero, which yields

$$\begin{aligned}\frac{\partial \log \mathcal{L}(\theta)}{\partial \theta} &= \sum_{i=1}^N \frac{2(x_i - \theta)}{2\sigma^2} \\ &= 0 \\ &= \frac{1}{\sigma^2} \left( \sum_{i=1}^N x_i - N\theta \right)\end{aligned}$$

so the maximum likelihood estimate is

$$\hat{\theta} = \frac{\sum_{i=1}^N x_i}{N}$$

which probably isn't all that surprising. Notice we did not have to pay attention to  $\sigma$  in this derivation — we did not assume it was known, it just doesn't do anything.

**Worked example 8.7** *Normal distributions -II*

Assume we have  $x_1, \dots, x_N$  which are data that can be modelled with a normal distribution. Use the maximum likelihood principle to estimate the standard deviation of that normal distribution.

**Solution:** Now we have to write out the log of the likelihood in more detail. Write  $\mu$  for the mean of the normal distribution and  $\theta$  for the unknown standard deviation of the normal distribution. We get

$$\log \mathcal{L}(\theta) = \left( \sum_{i=1}^N -\frac{(x_i - \mu)^2}{2\theta^2} \right) - N \log \theta + \text{Term not depending on } \theta$$

We can find the maximum by differentiating wrt  $\theta$  and setting to zero, which yields

$$\frac{\partial \log \mathcal{L}(\theta)}{\partial \theta} = \frac{-2}{\theta^3} \sum_{i=1}^N \frac{-(x_i - \mu)^2}{2} - \frac{N}{\theta} = 0$$

so the maximum likelihood estimate is

$$\hat{\theta} = \sqrt{\frac{\sum_{i=1}^N (x_i - \mu)^2}{N}}$$

which probably isn't all that surprising, either.

You should notice that one could maximize the likelihood of a normal distribution with respect to mean and standard deviation in one go (i.e. I could have done worked examples 8.6 and 8.7 in one worked example, instead of two). I did this example in two parts because I felt it was more accessible that way; if you object, you're likely to be able to fill in the details yourself very easily.

The maximum likelihood principle has a variety of neat properties we cannot expound. One worth knowing about is **consistency**; for our purposes, this means that the maximum likelihood estimate of parameters can be made arbitrarily close to the right answer by having a sufficiently large dataset. Now assume that our data doesn't actually come from the underlying model. This is the usual case, because we can't usually be sure that, say, the data truly is normal or truly comes from a Poisson distribution. Instead we choose a model that we think will be useful. When the data doesn't come from the model, maximum likelihood produces an estimate of  $\theta$  that corresponds to the model that is (in quite a strong sense, which we can't explore here) the closest to the source of the data. Maximum likelihood is very widely used because of these neat properties. But there are some difficulties.

### 8.1.2 Cautions about Maximum Likelihood

One important problem is that it might be hard to find the maximum of the likelihood exactly. There are strong numerical methods for maximizing functions, and

these are very helpful, but even today there are likelihood functions where it is very hard to find the maximum.

The second is that small amounts of data can present nasty problems. For example, in the binomial case, if we have only one flip we will estimate  $p$  as either 1 or 0. We should find this report unconvincing. In the geometric case, with a fair coin, there is a probability 0.5 that we will perform the estimate and then report that the coin has  $p = 1$ . This should also worry you. As another example, if we throw a die only a few times, we could reasonably expect that, for some  $i$ ,  $n_i = 0$ . This doesn't necessarily mean that  $p_i = 0$ , though that's what the maximum likelihood inference procedure will tell us.

This creates a very important technical problem — how can I estimate the probability of events that haven't occurred? This might seem like a slightly silly question to you, but it isn't. Solving this problem has really significant practical consequences. As one example, consider a biologist trying to count the number of butterfly species on an island. The biologist catches and classifies a lot of butterflies, then leaves. But are there more butterfly species on the island? To get some sense that we can reason successfully about this problem, compare two cases. In the first, the biologist catches many individuals of each of the species observed. In this case, you should suspect that catching more butterflies is unlikely to yield more species. In the second case, there are many species where the biologist sees only one individual of that species. In this case, you should suspect that catching more butterflies might very well yield new species.

As another example, a really important part of natural language processing involves estimating the probability of groups of three words. These groups are usually known as “trigrams”. People typically know an awful lot of words (tens to hundreds of thousands, depending on what you mean by a word). This means that there are a tremendous number of trigrams, and you can expect that any real dataset lacks almost all of them, because the dataset isn't big enough for there to be even just one of each trigram. Some are missing because they don't occur in real life, but others are not there simply because they are unusual (eg “Atom Heart Mother” actually occurs in real life, but you may not have seen it — try a web search if the phrase doesn't ring a bell). Modern speech recognition systems need to know how probable *every* trigram is. If the speech system thinks a trigram has zero probability and the trigram actually occurs, the system will make a mistake. We can't solve this problem just by giving each trigram a very small non-zero probability, because there are too many trigrams — it is important to distinguish between rare ones, and ones that don't ever occur. But what probability should I use for a rare trigram? Maximum likelihood would say use zero, but this would generate problems. Formalizing all this gets difficult quite quickly,

## 8.2 INCORPORATING PRIORS WITH BAYESIAN INFERENCE

Another important issue with maximum likelihood is that there is no mechanism to incorporate prior beliefs. For example, imagine you get a new die from a reliable store, roll it six times and see a one once. You would be happy to believe that  $p(6) = 1/6$  for this die. Now imagine you borrow a die from a friend with a long history of making weighted dice. Your friend tells you this die is weighted so that

$p(1) = 1/2$ . You roll the die six times and see a one once; in this case, you might worry that  $p(6)$  isn't  $1/6$ , and you just happened to get a slightly unusual set of rolls. You'd worry because you have good reason to believe the die isn't fair, and you'd want more evidence to believe  $p(6) = 1/6$ . Maximum likelihood can't distinguish between these two cases.

The difference lies in prior information — information we possess before we look at the data. We would like to take this information into account when we estimate the model. One way to do so is to place a **prior probability distribution**  $p(\theta)$  on the parameters  $\theta$ . Then, rather than working with the likelihood  $p(\mathcal{D}|\theta)$ , we could apply Bayes' rule, and form the **posterior**  $p(\theta|\mathcal{D})$ . This posterior represents the probability that  $\theta$  takes various values, given the data  $\mathcal{D}$ .

**Definition: 8.4** *Bayesian inference*

Extracting information from the posterior  $p(\theta|\mathcal{D})$  is usually called **Bayesian inference**

Indexing: Bayesian inference

Indexing: Bayesian inference

**Definition: 8.5** *MAP estimate*

A natural estimate of  $\theta$  is the value that maximizes the posterior  $p(\theta|\mathcal{D})$ . This estimate is known as a **maximum a posteriori estimate** or **MAP estimate**.

Indexing: MAP estimate

Indexing: MAP estimate

### 8.2.1 Constructing the Posterior

Bayes' rule tells us that

$$p(\theta|\mathcal{D}) = \frac{P(\mathcal{D}|\theta)P(\theta)}{P(\mathcal{D})}$$

but (as we shall see) it can be hard to work out  $P(\mathcal{D})$ . For some problems, we might not need to know it.

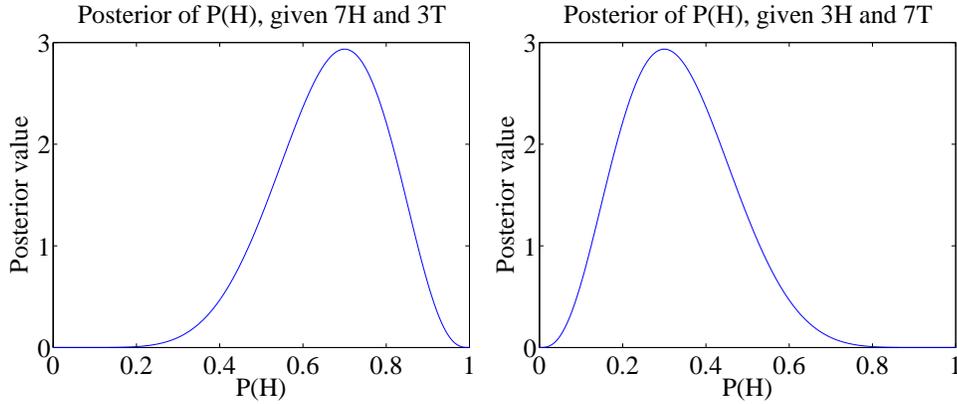


FIGURE 8.1: The curves show a function proportional to the posterior on  $\theta$ , for the two cases of example ???. Notice that this information is rather richer than the single value we would get from maximum likelihood inference.

**Worked example 8.8** *Flipping a coin*

We have a coin with probability  $\theta$  of coming up heads when flipped. We start knowing nothing about  $\theta$ . We then flip the coin 10 times, and see 7 heads (and 3 tails). Plot a function proportional to  $p(\theta | \{7 \text{ heads and } 3 \text{ tails}\})$ . What happens if there are 3 heads and 7 tails?

**Solution:** We know nothing about  $p$ , except that  $0 \leq \theta \leq 1$ , so we choose a uniform prior on  $p$ . We have that  $p(\{7 \text{ heads and } 3 \text{ tails}\} | \theta)$  is binomial. The joint distribution is  $p(\{7 \text{ heads and } 3 \text{ tails}\} | \theta) \times p(\theta)$  but  $p(\theta)$  is uniform, so doesn't depend on  $\theta$ . So the posterior is *proportional to*:  $\theta^7(1 - \theta)^3$  which is graphed in figure 8.1. The figure also shows  $\theta^3(1 - \theta)^7$  which is *proportional to* the posterior for 3 heads and 7 tails. In each case, the evidence does not rule out the possibility that  $\theta = 0.5$ , but tends to discourage the conclusion. Maximum likelihood would give  $\theta = 0.7$  or  $\theta = 0.3$ , respectively.

In Example 8.8, it is interesting to follow how the posterior on  $p$  changes as evidence come in, which is easy to do because the posterior is proportional to a binomial distribution. Figure 8.2 shows a set of these posteriors for different sets of evidence.

For other problems, we will need to marginalize out  $\theta$ , by computing

$$P(\mathcal{D}) = \int_{\theta} P(\mathcal{D}|\theta)P(\theta)d\theta.$$

It is usually impossible to do this in closed form, so we would have to use a numerical integral or a trick. The next section expounds one useful trick.

**Conjugate priors**

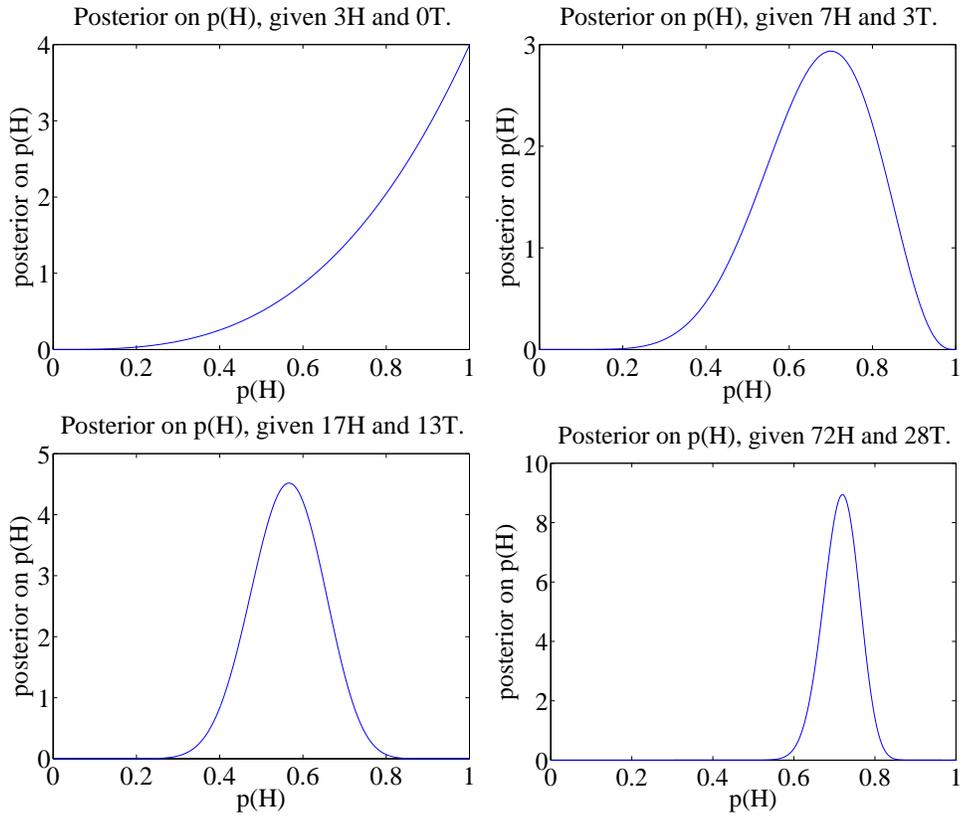


FIGURE 8.2: The probability that an unknown coin will come up heads when flipped is  $p(H)$ . For these figures, I simulated coin flips from a coin with  $p = 0.75$ . I then plotted the posterior for various data. Notice how, as we see more flips, we get more confident about  $p$ . The vertical axis changes substantially between plots in this figure.

In some cases,  $P(\theta)$  and  $P(\mathcal{D}|\theta)$ , when multiplied together, take a familiar form. This happens when  $P(\mathcal{D}|\theta)$  and  $P(\theta)$  each belong to parametric families where there is a special relationship between the families. When a prior has this property, it is called a **conjugate** prior. There are some cases worth knowing, given in the worked examples.

**Worked example 8.9** *Flipping a coin - II*

We have a coin with probability  $\theta$  of coming up heads when flipped. We model the prior on  $\theta$  with a Beta distribution, with parameters  $\alpha > 0$ ,  $\beta > 0$ . We then flip the coin  $N$  times, and see  $h$  heads. What is  $P(\theta|N, h, \alpha, \beta)$ ?

**Solution:** We have that  $P(N, h|\theta)$  is binomial, and that  $P(\theta|N, h, \alpha, \beta) \propto P(N, h|\theta)P(\theta|\alpha, \beta)$ . This means that

$$P(\theta|N, h, \alpha, \beta) \propto \binom{N}{h} \theta^h (1-\theta)^{(N-h)} \frac{\Gamma(\alpha+\beta)}{\Gamma(\alpha)\Gamma(\beta)} \theta^{\alpha-1} (1-\theta)^{\beta-1}.$$

and we can write

$$P(\theta|N, h, \alpha, \beta) \propto \theta^{(\alpha+h-1)} (1-\theta)^{(\beta+N-h-1)}.$$

Notice this has the form of a Beta distribution, so it is easy to recover the constant of proportionality. We have

$$P(\theta|N, h, \alpha, \beta) = \frac{\Gamma(\alpha+\beta+N)}{\Gamma(\alpha+h)\Gamma(\beta+N-h)} \theta^{(\alpha+h-1)} (1-\theta)^{(\beta+N-h-1)}.$$

**Worked example 8.10** *More swearsy politicians*

Example 8.5 gives some data from a swearsy politician. Assume we have only the first 10 intervals of observations, and we wish to estimate the intensity using a Poisson model. Write  $\theta$  for this parameter. Use a Gamma distribution as a prior, and write out the posterior.

**Solution:** We have that

$$\begin{aligned} p(\mathcal{D}|\theta) &= \left(\frac{\theta^0 e^{-\theta}}{0!}\right)^5 \left(\frac{\theta^1 e^{-\theta}}{1!}\right)^2 \times \\ &\quad \left(\frac{\theta^2 e^{-\theta}}{2!}\right)^2 \left(\frac{\theta^3 e^{-\theta}}{3!}\right)^1 \\ &= \frac{\theta^9 e^{-10\theta}}{12} \end{aligned}$$

and

$$p(\theta|\alpha, \beta) = \frac{\beta^\alpha}{\Gamma(\alpha)} \theta^{(\alpha-1)} e^{-\beta\theta}$$

This means that

$$p(\theta|\mathcal{D}) \propto \theta^{(\alpha-1+9)} e^{-(\beta+10)\theta}.$$

Notice this has the form of another Gamma distribution, so we can write

$$p(\theta|\mathcal{D}) = \frac{(\beta+10)^{(\alpha+9)}}{\Gamma(\alpha+9)} \theta^{(\alpha-1+9)} e^{-(\beta+10)\theta}$$

### 8.2.2 Normal Prior and Normal Likelihood

When both  $P(\mathcal{D}|\theta)$  and  $P(\theta)$  are normal, some important simplifications occur. First, the prior is conjugate to the likelihood. Second, the posterior is also normal. And third, the mean and standard deviation of the posterior take a simple form.

We start with an example. Assume we drop a measuring device down a borehole. It is designed to stop falling and catch onto the side of the hole after it has fallen  $\mu_0$  meters. On board is a device to measure its depth. This device reports a known constant times the correct depth plus a zero mean normal random variable, which we call “noise”. The device reports depth every second.

The first question to ask is what depth do we believe the device is at *before* we receive any measurement? We designed the device to stop at  $\mu_0$  meters, so we are not completely ignorant about where it is. However, it may not have worked absolutely correctly. We choose to model the depth at which it stops as  $\mu_0$  meters plus a zero mean normal random variable. The second term could be caused by error in the braking system, etc. We could estimate the standard deviation of the second term (which we write  $\sigma_0$ ) either by dropping devices down holes, then measuring

with tape measures, or by analysis of likely errors in our braking system. The depth of the object is the unknown parameter of the model; we write this depth  $\theta$ . Now the model says that  $\theta$  is a normal random variable with mean  $\mu_0$  and standard deviation  $\sigma_0$ .

Notice that this model probably isn't exactly right — for example, there must be some probability in the model that the object falls beyond the bottom of the hole, which it can't do — but it captures some important properties of our system. The device should stop at or close to  $\mu_0$  meters most of the time, and it's unlikely to be too far away.

Now assume we receive a single measurement — what do we now know about the device's depth? The first thing to notice is that there is something to do here. Ignoring the prior and taking the measurement might not be wise. For example, imagine that the noise in the wireless system is large, so that the measurement is often corrupted — our original guess about the device's location might be better than the measurement. Write  $x_1$  for the measurement. Notice that the scale of the measurement may not be the same as the scale of the depth, so the mean of the measurement is  $c_1\theta$ , where  $c_1$  is a change of scale (for example, from inches to meters). We have that  $p(x_1|\theta)$  is normal with mean  $c_1\theta$  and standard deviation  $\sigma_{n1}$ . We would like to know  $p(\theta|x_1)$ .

We have that

$$\begin{aligned} \log p(\theta, x_1) &= \log p(x_1|\theta) + \log p(\theta) \\ &= -\frac{(x_1 - c_1\theta)^2}{2\sigma_{n1}^2} - \frac{(\theta - \mu_0)^2}{2\sigma_0^2} \\ &\quad + \text{terms not depending on } \theta \text{ or } x. \end{aligned}$$

We have two estimates of the position,  $\theta$ , and we wish to come up with a representation of what we know about  $\theta$ . One is  $x_1$ , which is a *measurement* — we know its value. The expected value of  $x_1$  is  $c_1\theta$ , so we could infer  $\theta$  from  $x_1$ . But we have another estimate of the position, which is  $\mu_0$ . The posterior,  $p(\theta|x_1)$ , is a probability distribution on the variable  $\theta$ ; it depends on the known values  $x_1$ ,  $\mu_0$ ,  $\sigma_0$  and  $\sigma_{n1}$ . We need to determine its form. We can do so by some rearrangement of the expression for  $\log p(\theta, x_1)$ .

Notice first that this expression is of degree 2 in  $\theta$  (i.e. it has terms  $\theta^2$ ,  $\theta$  and things that don't depend on  $\theta$ ). This means that  $p(\theta|x_1)$  must be a normal distribution, because we can rearrange its log into the form of the log of a normal distribution. This yields a fact of crucial importance.

**Useful Fact: 8.1** *Normal distributions are conjugate*

A normal prior and a normal likelihood yield a normal posterior.

Indexing: Normal distributions are conjugate

Indexing: Normal distributions are conjugate

Write  $\mu_1$  for the mean of this distribution, and  $\sigma_1$  for its standard deviation. The log of the distribution must be

$$-\frac{(\theta - \mu_1)^2}{2\sigma_1^2} + \text{terms not depending on } \theta.$$

The terms not depending on  $\theta$  are not interesting, because if we know  $\sigma_1$  those terms must add up to

$$\log\left(\frac{1}{\sqrt{2\pi}\sigma_1}\right)$$

so that the probability density function sums to one. Our goal is to rearrange terms into the form above. Notice that

$$-\frac{(\theta - \mu_1)^2}{2\sigma_p^2} = -\theta^2 \left(\frac{1}{2\sigma_1^2}\right) + 2\theta \frac{\mu_1}{2\sigma_p^2} + \text{term not depending on } \theta$$

We have

$$\begin{aligned} \log p(\theta|x_1) &= -\frac{(c_1\theta - x_1)^2}{2\sigma_{n1}^2} - \frac{(\theta - \mu_0)^2}{2\sigma_0^2} + \text{terms not depending on } \theta \\ &= -\theta^2 \left( \frac{1}{2 \left( \frac{\sigma_{n1}^2 \sigma_0^2}{\sigma_{n1}^2 + c_1^2 \sigma_0^2} \right)} \right) + 2\theta \left( c_1 \frac{x_1}{2\sigma_{n1}^2} + \frac{\mu_0}{2\sigma_0^2} \right) \\ &+ \text{terms not depending on } \theta \end{aligned}$$

which means that

$$\sigma_1^2 = \frac{\sigma_{n1}^2 \sigma_0^2}{\sigma_{n1}^2 + c_1^2 \sigma_0^2}$$

and

$$\begin{aligned} \mu_1 &= 2 \left( c_1 \frac{x_1}{2\sigma_{n1}^2} + \frac{\mu_0}{2\sigma_0^2} \right) \frac{\sigma_{n1}^2 \sigma_0^2}{\sigma_{n1}^2 + c_1^2 \sigma_0^2} \\ &= \left( \frac{c_1 x_1 \sigma_0^2 + \mu_0 \sigma_{n1}^2}{\sigma_{n1}^2 \sigma_0^2} \right) \frac{\sigma_{n1}^2 \sigma_0^2}{\sigma_{n1}^2 + c_1^2 \sigma_0^2} \\ &= \frac{c_1 x_1 \sigma_0^2 + \mu_0 \sigma_{n1}^2}{\sigma_{n1}^2 + c_1^2 \sigma_0^2}. \end{aligned}$$

These equations “make sense”. Imagine that  $\sigma_0$  is very small, and  $\sigma_{n1}$  is very big; then our new expected value of  $\theta$  — which is  $\mu_1$  — is about  $\mu_0$ . Equivalently, because our prior was very accurate, and the measurement was unreliable, our expected value is about the prior value. Similarly, if the measurement is reliable (i.e.  $\sigma_{n1}$  is small) and the prior has high variance (i.e.  $\sigma_0$  is large), then our expected value of  $\theta$  is about  $x_1/c_1$  — i.e. the measurement, rescaled. I have put these equations, in a more general form, in a box below.

**Useful Fact: 8.2** *Normal posteriors*

Assume we wish to estimate a parameter  $\theta$ . The prior distribution for  $\theta$  is normal, with known mean  $\mu_\pi$  and known standard deviation  $\sigma_\pi$ . We receive a single data item  $x$ . The likelihood of this data item is normal with mean  $c\theta$  and standard deviation  $\sigma_m$ , where  $c$  and  $\sigma_m$  are known. Then the posterior,  $p(\theta|x, c, \sigma_m, \mu_\pi, \sigma_\pi)$ , is normal, with mean

$$\frac{cx\sigma_\pi^2 + \mu_\pi\sigma_m^2}{\sigma_m^2 + c^2\sigma_\pi^2}$$

and standard deviation

$$\sqrt{\frac{\sigma_m^2\sigma_\pi^2}{\sigma_m^2 + c^2\sigma_\pi^2}}.$$

Indexing: Normal posteriors

Indexing: Normal posteriors

## 8.2.3 MAP Inference

Look at example 8.1, where we estimated the probability a coin would come up heads with maximum likelihood. We could not change our estimate just by knowing the coin was fair, but we could come up with a number for  $\theta = p(H)$  (rather than, say, a posterior distribution). A natural way to produce a point estimate for  $\theta$  that incorporates prior information is to choose  $\hat{\theta}$  such that

$$\hat{\theta} = \operatorname{argmax}_{\theta} P(\theta|\mathcal{D}) = \operatorname{argmax}_{\theta} \frac{P(\theta, \mathcal{D})}{P(\mathcal{D})}$$

This is the MAP estimate. If we wish to perform MAP inference,  $P(\mathcal{D})$  doesn't matter (it changes the value, but not the location, of the maximum). This means we can work with  $P(\theta, \mathcal{D})$ , often called the **joint** distribution.

**Worked example 8.11** *Flipping a coin - II*

We have a coin with probability  $\theta$  of coming up heads when flipped. We model the prior on  $\theta$  with a Beta distribution, with parameters  $\alpha > 0$ ,  $\beta > 0$ . We then flip the coin  $N$  times, and see  $h$  heads. What is the MAP estimate of  $\theta$ ?

**Solution:** We have that

$$P(\theta|N, h, \alpha, \beta) = \frac{\Gamma(\alpha + \beta + N)}{\Gamma(\alpha + h)\Gamma(\beta + N - h)} \theta^{(\alpha+h-1)}(1-\theta)^{(\beta+N-h-1)}.$$

You can get the MAP estimate by differentiating and setting to 0, yielding

$$\hat{\theta} = \frac{\alpha - 1 + h}{\alpha + \beta - 2 + N}.$$

This has rather a nice interpretation. You can see  $\alpha$  and  $\beta$  as extra counts of heads (resp. tails) that are added to the observed counts. So, for example, if you were fairly sure that the coin should be fair, you might make  $\alpha$  and  $\beta$  large and equal. When  $\alpha = 1$  and  $\beta = 1$ , we have a uniform prior as in the previous examples.

**Worked example 8.12** *More swears politicians*

We observe our swearing politician for  $N$  intervals, seeing  $n_i$  swear words in the  $i$ 'th interval. We model the swearing with a Poisson model. We wish to estimate the intensity, which we write  $\theta$ . We use a Gamma distribution for the prior on  $\theta$ . What is the MAP estimate of  $\theta$ ?

**Solution:** Write  $T = \sum_{i=1}^N n_i$ . We have that

$$p(\theta|\mathcal{D}) = \frac{(\beta + N)^{(\alpha+T)}}{\Gamma(\alpha + T)} \theta^{(\alpha-1+T)} e^{-(\beta+T)\theta}$$

and the MAP estimate is

$$\hat{\theta} = \frac{(\alpha - 1 + T)}{(\beta + N)}$$

(which you can get by differentiating with respect to  $\theta$ , then setting to zero). Notice that if  $\beta$  is close to zero, you can interpret  $\alpha$  as extra counts; if  $\beta$  is large, then it strongly discourages large values of  $\hat{\theta}$ , even if the counts are large.

## 8.2.4 Filtering

We can make online estimates of the maximum likelihood value of mean and standard deviation for a normal distribution. Assume, rather than seeing  $N$  elements

of a dataset in one go, you get to see each one once, and you cannot store them. Assume that this dataset is modelled as normal data. Write  $\hat{\mu}_k$  for the maximum likelihood estimate of the mean based on data items  $1 \dots k$  (and  $\hat{\sigma}_k$  for the maximum likelihood estimate of the standard deviation, etc.). Notice that

$$\hat{\mu}_{k+1} = \frac{(k\hat{\mu}_k) + x_{k+1}}{(k+1)}$$

and that

$$\hat{\sigma}_{k+1} = \sqrt{\frac{(k\hat{\sigma}_k^2) + (x_{k+1} - \hat{\mu}_{k+1})^2}{(k+1)}}$$

This means that you can incorporate new data into your estimate as it arrives without keeping all the data. This process of updating a representation of a dataset as new data arrives is known as **filtering**.

This is particularly useful in the case of normal posteriors. Recall that if we have a normal prior and a normal likelihood, the posterior is normal. Now consider a stream of incoming measurements. Our initial representation of the parameters we are trying to estimate is the prior, which is normal. We allow ourselves to see one measurement, which has normal likelihood; then the posterior is normal. You can think of this posterior as a prior for the parameter estimate based on the *next* measurement. But we know what to do with a normal prior, a normal likelihood, and a measurement; so we can incorporate the measurement and go again. This means we can exploit our expression for the posterior mean and standard deviation in the case of normal likelihood and normal prior and a single measurement to deal with multiple measurements very easily.

Assume a second measurement,  $x_2$  arrives. We know that  $p(x_2|\theta, c_2, \sigma_{n2})$  is normal with mean  $c_2\theta$  and standard deviation  $\sigma_{n2}$ . In the example, we have a new measurement of depth — perhaps in a new, known, scale — with new noise (which might have larger, or smaller, standard deviation than the old noise) added. Then we can use  $p(\theta|x_1, c_1, \sigma_{n1})$  as a prior to get a posterior  $p(\theta|x_1, x_2, c_1, c_2, \sigma_{n1}, \sigma_{n2})$ . Each is normal, by useful fact 8.1. Not only that, but we can easily obtain the expressions for the mean  $\mu_2$  and the standard deviation  $\sigma_2$  *recursively* as functions of  $\mu_1$  and  $\sigma_1$ .

Applying useful fact 8.2, we have

$$\mu_2 = \frac{c_2x_2\sigma_1^2 + \mu_1\sigma_{n2}^2}{\sigma_{n2}^2 + c_2^2\sigma_1^2}$$

and

$$\sigma_2^2 = \frac{\sigma_{n2}^2\sigma_1^2}{\sigma_{n2}^2 + c_2^2\sigma_1^2}.$$

But what works for 2 and 1 will work for  $k+1$  and  $k$ . We know the posterior after  $k$  measurements will be normal, with mean  $\mu_k$  and standard deviation  $\sigma_k$ . The  $k+1$ 'th measurement  $x_{k+1}$  arrives, and we have  $p(x_{k+1}|\theta, c_{k+1}, \sigma_{n(k+1)})$  is normal. Then the posterior is normal, and we can write the mean  $\mu_{k+1}$  and the standard deviation  $\sigma_{k+1}$  recursively as functions of  $\mu_k$  and  $\sigma_k$ . The result is worth putting in a box.

**Useful Fact: 8.3** *Online updating of normal posteriors*

Assume we wish to estimate a parameter  $\theta$ . The prior distribution for  $\theta$  is normal, with known mean  $\mu_\pi$  and known standard deviation  $\sigma_\pi$ . All data is normal conditioned on  $\theta$ . We have already received  $k$  data items. The posterior  $p(\theta|x_1, \dots, x_k, c_1, \dots, c_k, \sigma_{n1}, \dots, \sigma_{nk}, \mu_\pi, \sigma_\pi)$  is normal, with mean  $\mu_k$  and standard deviation  $\sigma_k$ . We receive a new data item  $x_{k+1}$ . The likelihood of this data item is normal with mean  $c\theta$  and standard deviation  $\sigma_{n(k+1)}$ , where  $c_{k+1}$  and  $\sigma_{n(k+1)}$  are known. Then the posterior,  $p(\theta|x_1, \dots, x_{k+1}, c_1, \dots, c_k, c_{k+1}, \sigma_{n1}, \dots, \sigma_{n(k+1)}, \mu_\pi, \sigma_\pi)$ , is normal, with mean

$$\mu_{k+1} = \frac{c_{k+1}x_{k+1}\sigma_k^2 + \mu_k\sigma_{n(k+1)}^2}{\sigma_{n(k+1)}^2 + c_{k+1}^2\sigma_k^2}$$

and

$$\sigma_{k+1}^2 = \frac{\sigma_{n(k+1)}^2\sigma_k^2}{\sigma_{n(k+1)}^2 + c_{k+1}^2\sigma_k^2}.$$

Indexing: Online updating of normal posteriors

Indexing: Online updating of normal posteriors

Again, notice the very useful fact that, if everything is normal, we can update our posterior representation when new data arrives using a very simple recursive form.

**Worked example 8.13**    *Normal data*

Assume you see  $N$  datapoints  $x_i$  which are modelled by a normal distribution with unknown mean  $\theta$  and with known standard deviation  $\sigma$ . You model the prior on  $\theta$  using a normal distribution with mean  $\mu_0$  and standard deviation  $\sigma_0$ . What is the MAP estimate of the mean?

**Solution:** Recall that the maximum value of a normal distribution occurs at its mean. Now problem is covered by useful fact 8.2, but in this case we have  $c_i = 1$  for each data point, and  $\sigma_i = \sigma$ . We can write

$$\mu_N = \frac{x_N \sigma_{N-1}^2 + \mu_{N-1} \sigma^2}{\sigma^2 + \sigma_{N-1}^2}$$

and

$$\sigma_N^2 = \frac{\sigma^2 \sigma_{N-1}^2}{\sigma^2 + \sigma_{N-1}^2}.$$

and evaluate the recursion down to  $\mu_0, \sigma_0$ .

## 8.2.5 Cautions about Bayesian Inference

Just like maximum likelihood inference, bayesian inference is not a recipe that can be applied without thought. It turns out that, when there is a lot of data, the prior has little inference on the outcome of the inference, and the MAP solution looks a lot like the maximum likelihood solution. So the difference between the two approaches is most interesting when there is little data, where the prior matters. The difficulty is that it might be hard to know what to use as a good prior. In the examples, I emphasized mathematical convenience, choosing priors that lead to clean posteriors. There is no reason to believe that nature uses conjugate priors (even though conjugacy is a neat property). How should one choose a prior for a real problem?

This isn't an easy point. If there is little data, then the choice could really affect the inference. Sometimes we're lucky, and the logic of the problem dictates a choice of prior. Mostly, we have to choose and live with the consequences of the choice. Often, doing so is succesful in applications.

The fact we can't necessarily justify a choice of prior seems to be one of life's inconveniences, but it represents a significant philosophical problem. It's been at the core of a long series of protracted, often quite intense, arguments about the philosophical basis of statistics. I haven't followed these arguments closely enough to summarize them; they seem to have largely died down without any particular consensus being reached.

## 8.3 SAMPLES, URNS AND POPULATIONS

Very often the data we see is a small part of the data we could have seen, if we'd been able to collect enough data. We need to know how the measurements we make

on the dataset relate to the measurements we could have made, if we had all the data. This situation occurs very often. For example, imagine we wish to know the average weight of a rat. This isn't random; you could weigh every rat on the planet, and then average the answers. But doing so would be absurd (among other things, you'd have to weigh them all at the same time, which would be tricky). Instead, we weigh a small set of rats, chosen rather carefully. If we have chosen sufficiently carefully, then the answer from the small set is quite a good representation of the answer from the whole set.

The data we could have observed, if we could have seen everything, is the **population**. The data we actually have is the **sample**. We would like to know the mean of the population, but can see only the sample; surprisingly, we can say a great deal from the sample alone, assuming that it is chosen appropriately.

### 8.3.1 Estimating the Population Mean from a Sample

Assume we have a population  $\{x\}$ , for  $i = 1, \dots, N_p$ . Notice the subscript here — this is the number of items in the population. The population could be unreasonably big: for example, it could consist of all the people in the world. We want to know the mean of this dataset, but we do not get to see the whole dataset. Instead, we see the sample.

How the sample is obtained is key to describing the population. We will focus on only one model (there are lots of others). In our model, the sample is obtained by choosing a fixed number of data items. Write  $k$  for the number of data items in the sample. We expect  $k$  is a lot smaller than  $N_p$ . Each item is chosen independently, and fairly. This means that each time we choose, we choose one from the entire set of  $N_p$  data items, and each has the same probability of being chosen. This is sometimes referred to as “sampling with replacement”.

One natural way to think about sampling with replacement is to imagine the data items as being written on tickets, which are placed in an urn (old-fashioned word for a jar, now used mainly by statisticians and morticians). You obtain the sample by repeating the following experiment  $k$  times: shake the urn; take a ticket from the urn and write down the data on the ticket; put it back in the urn. Notice that, in this case, each sample is drawn from the same urn. This is important, and makes the analysis easier. If we had not put the ticket back, the urn would change between samples.

We summarize the whole dataset with its mean, which we write  $\text{popmean}(\{x\})$ . This is known as the **population mean**. The notation is just to drive home the facts that it's the mean of the whole population, and that we don't, and can't, know it. The whole point of this exercise is to *estimate* this mean.

We would like to estimate the mean of the whole dataset from the items that we actually see. Imagine we draw  $k$  tickets from the urn as above, and average the values. The result is a random variable, because different draws of  $k$  tickets will give us different values. Write  $X^{(k)}$  for this random variable, which is referred to as the **sample mean**. Because expectations are linear, we must have that

$$\mathbb{E}[X^{(k)}] = \frac{1}{k} \left( \mathbb{E}[X^{(1)}] + \dots + \mathbb{E}[X^{(1)}] \right) = \mathbb{E}[X^{(1)}]$$

(where  $X^{(1)}$  is the random variable whose value is obtained by drawing one ticket

from the urn). Now

$$\begin{aligned}\mathbb{E}[X^{(1)}] &= \sum_{i \in 1, \dots, N_p} x_i p(i) \\ &= \sum_{i \in 1, \dots, N_p} x_i \frac{1}{N_p} && \text{because we draw fairly from the urn} \\ &= \frac{\sum_{i \in 1, \dots, N_p} x_i}{N_p} \\ &= \text{popmean}(\{x\})\end{aligned}$$

which is the mean value of the items in the urn. This means that

$$\mathbb{E}[X^{(k)}] = \text{popmean}(\{x_i\}).$$

Under our sampling model, the expected value of the sample mean is the population mean.

**Useful Facts: 8.4** *Sample means and population means*

The sample mean is a random variable. It is random, because different samples from the population will have different values of the sample mean. The expected value of this random variable is the population mean.

Indexing: Sample means and population means

Indexing: Sample means and population means

We will not get the same value of  $X^{(k)}$  each time we perform the experiment, because we see different data items in each sample. So  $X^{(k)}$  has variance, and this variance is important. If it is large, then each estimate is quite different. If it is small, then the estimates cluster. Knowing the variance of  $X^{(k)}$  would tell us how accurate our estimate of the population mean is.

### 8.3.2 The Variance of the Sample Mean

We write  $\text{popsd}(\{x\})$  for the standard deviation of the whole population of  $\{x\}$ . Again, we write it like this to keep track of the facts that (a) it's for the whole population and (b) we don't — and usually can't — know it.

We can compute the variance of  $X^{(k)}$  (the sample mean) easily. We have

$$\text{var}[X^{(k)}] = \mathbb{E}[(X^{(k)})^2] - \mathbb{E}[X^{(k)}]^2 = \mathbb{E}[(X^{(k)})^2] - (\text{popmean}(\{x\}))^2$$

so we need to know  $\mathbb{E}[(X^{(k)})^2]$ . We can compute this by writing

$$X^{(k)} = \frac{1}{k}(X_1 + X_2 + \dots + X_k)$$

where  $X_1$  is the value of the first ticket drawn from the urn, etc. We then have

$$X^{(k)2} = \left(\frac{1}{k}\right)^2 \begin{pmatrix} X^2_1 + X^2_2 + \dots X^2_k + X_1X_2 + \dots \\ X_1X_k + X_2X_1 + \dots X_2X_k + \dots X_{k-1}X_k \end{pmatrix}.$$

Expectations are linear, so we have that

$$\mathbb{E}\left[(X^{(k)})^2\right] = \left(\frac{1}{k}\right)^2 \begin{pmatrix} \mathbb{E}[X^2_1] + \mathbb{E}[X^2_2] + \dots \mathbb{E}[X^2_k] + \mathbb{E}[X_1X_2] + \\ \dots \mathbb{E}[X_1X_k] + \mathbb{E}[X_2X_1] + \dots \mathbb{E}[X_2X_k] + \dots \mathbb{E}[X_{k-1}X_k] \end{pmatrix}.$$

The *order* in which the tickets are drawn from the urn doesn't matter, because each time we draw a ticket we draw from the same urn. This means that  $\mathbb{E}[X^2_1] = \mathbb{E}[X^2_2] = \dots \mathbb{E}[X^2_k]$ . You can think of this term as the expected value of the random variable generated by: drawing a single number out of the urn; squaring that number; and reporting the square. Notice that  $\mathbb{E}[X^2_1] = \mathbb{E}[(X^{(1)})^2]$  (look at the definition of  $X^{(1)}$ ).

Because the *order* doesn't matter, we also have that  $\mathbb{E}[X_1X_2] = \mathbb{E}[X_1X_3] = \dots \mathbb{E}[X_{k-1}X_k]$ . You can think of this term as the expected value of the random variable generated by: drawing a number out of the urn; writing it down; returning it to the urn; then drawing a second number from the urn; and reporting the product of these two numbers. So we can write

$$\mathbb{E}\left[X^{(k)2}\right] = \left(\frac{1}{k}\right)^2 \left(k\mathbb{E}[(X^{(1)})^2] + k(k-1)\mathbb{E}[X_1X_2]\right)$$

and these two terms are quite easy to evaluate.

#### Worked example 8.14    *Urn variances*

Show that

$$\mathbb{E}\left[(X^{(1)})^2\right] = \frac{\sum_{i=1}^{N_p} x_i^2}{N_p} = \text{popstd}(\{x\})^2 + \text{popmean}(\{x\})^2$$

**Solution:** First, we have  $(X^{(1)})^2$  is the number obtained by taking a ticket out of the urn and squaring its data item. Now

$$\text{popstd}(\{x\})^2 = \mathbb{E}\left[(X^{(1)})^2\right] - \mathbb{E}\left[X^{(1)}\right]^2 = \mathbb{E}\left[(X^{(1)})^2\right] - \text{popmean}(\{x\})^2$$

so

$$\mathbb{E}\left[(X^{(1)})^2\right] = \text{popstd}(\{x\})^2 + \text{popmean}(\{x\})^2$$

**Worked example 8.15**    *Urn variances*

Show that

$$\mathbb{E}[X_1 X_2] = \text{popmean}(\{x\})^2$$

**Solution:** This looks hard, but isn't. Recall from the facts in chapter 5 (useful facts 5.6, page 137) that if  $X$  and  $Y$  are independent random variables,  $\mathbb{E}[XY] = \mathbb{E}[X]\mathbb{E}[Y]$ . But  $X_1$  and  $X_2$  are independent — they are different random draws from the same urn. So

$$\mathbb{E}[X_1 X_2] = \mathbb{E}[X_1]\mathbb{E}[X_2]$$

but  $\mathbb{E}[X_1] = \mathbb{E}[X_2]$  (they are draws from the same urn) and  $\mathbb{E}[X] = \text{popmean}(\{x\})$ . So

$$\mathbb{E}[X_1 X_2] = \text{popmean}(\{x\})^2.$$

Now

$$\begin{aligned} \mathbb{E}[(X^{(k)})^2] &= \frac{k\mathbb{E}[(X^{(1)})^2] + k(k-1)\mathbb{E}[X_1 X_2]}{k^2} \\ &= \frac{\mathbb{E}[(X^{(1)})^2] + (k-1)\mathbb{E}[X_1 X_2]}{k} \\ &= \frac{(\text{popsd}(\{x\})^2 + \text{popmean}(\{x\})^2) + (k-1)\text{popmean}(\{x\})^2}{k} \\ &= \frac{\text{popsd}(\{x\})^2}{k} + \text{popmean}(\{x\})^2 \end{aligned}$$

so we have

$$\begin{aligned} \text{var}[X^{(k)}] &= \mathbb{E}[(X^{(k)})^2] - \mathbb{E}[X^{(k)}]^2 \\ &= \frac{\text{popsd}(\{x\})^2}{k} + \text{popmean}(\{x\})^2 - \text{popmean}(\{x\})^2 \\ &= \frac{\text{popsd}(\{x\})^2}{k}. \end{aligned}$$

This is a very useful result which is well worth remembering together with our facts on the sample mean, so we'll put them in a box together.

**Useful Fact: 8.5** *The sample mean*

The sample mean is a random variable. Write  $X^{(k)}$  for the mean of  $k$  samples. We have that:

$$\begin{aligned}\mathbb{E}[X^{(k)}] &= \text{popmean}(\{x\}) \\ \text{var}[X^{(k)}] &= \frac{\text{popsd}(\{x\})^2}{k} \\ \text{std}(\{X^{(k)}\}) &= \frac{\text{popsd}(\{x\})}{\sqrt{k}}\end{aligned}$$

Indexing: The sample mean

Indexing: The sample mean

The consequence is this: If you draw  $k$  samples, the standard deviation of your estimate of the mean is

$$\frac{\text{popsd}(\{x\})}{\sqrt{k}}$$

which means that (a) the more samples you draw, the better your estimate becomes and (b) the estimate improves rather slowly — for example, to halve the standard deviation in your estimate, you need to draw four times as many samples. The standard deviation of the estimate of the mean is often known as the **standard error** of the mean. This allows us to draw a helpful distinction: the population has a standard deviation, and our estimate of its mean (or other things — but we won't go into this) has a standard error.

Notice we cannot state the standard error of our estimate exactly, because we do not know  $\text{popsd}(\{x\})$ . But we could make a good estimate of  $\text{popsd}(\{x\})$ , by computing the standard deviation of the examples that we have. It is now helpful to have some notation for the particular sample we have. I will write  $\sum_{i \in \text{sample}}$  for a sum over the sample items, and we will use

$$\text{mean}(\{x\}) = \frac{\sum_{i \in \text{sample}} x_i}{k}$$

for the mean of the sample — that is, the mean of the data we actually see; this is consistent with our old notation, but there's a little reindexing to keep track of the fact we don't see all of the population. Similarly, I will write

$$\text{std}(\{x\}) = \sqrt{\frac{\sum_{i \in \text{sample}} (x_i - \text{mean}(\{x_i\}))^2}{k}}$$

for the sample standard deviation. Again, this is the standard deviation of the data we actually see; and again, this is consistent with our old notation, again with a

little reindexing to keep track of the fact we don't see all of the population. We could estimate

$$\text{popstd}(\{x\}) \approx \text{std}(\{x\})$$

and as long as we have enough examples, this estimate is good. If the number of samples  $k$  is small, it is better to use

$$\text{popstd}(\{x\}) \approx \sqrt{\frac{\sum_{i \in \text{sample}} (x_i - \text{mean}(\{x\}))^2}{k-1}}.$$

In fact, much more is known about the distribution of  $X^{(k)}$ .

### 8.3.3 The Probability Distribution of the Sample Mean

The sample mean is a random variable. We know an expression for its mean, and we can estimate its variance. In fact, we can determine its probability distribution, though I won't do this rigorously. In section 6.4.3, I mentioned that adding a number of independent random variables almost always got you a normal random variable, a fact sometimes known as the central limit theorem. I didn't prove it, and I'm not going to now. But when we form  $X^{(k)}$ , we're adding random variables. This means that  $X^{(k)}$  is a normal random variable, for sufficiently big  $k$  (for some reason,  $k > 30$  is usually seen as right).

This is important, because it has the following consequence. Draw a large number of different samples of  $k$  elements from the population. Each is a dataset of  $k$  items. Compute  $\text{mean}(\{x\})$  for each, and regard the resulting numbers  $e_1, \dots, e_r$  as data items. Convert the  $e_i$  to standard coordinates  $s_i$ , where

$$s_i = \frac{(e_i - \text{mean}(\{e_i\}))}{\text{std}(e_i)}$$

(i.e. by subtracting the mean of the  $e_i$ , and dividing by their standard deviation). Now construct a histogram of the  $s$ . If  $r$  is sufficiently large, the histogram will be close to the standard normal curve.

### 8.3.4 When The Urn Model Works

In our model, there was a population of  $N_p$  data items  $x_i$ , and we saw  $k$  of them, chosen at random. In particular, each choice was fair (in the sense that each data item had the same probability of being chosen) and independent. These assumptions are very important for our analysis to apply. If our data does not have these properties, bad things can happen.

For example, assume we wish to estimate the percentage of the population that has beards. This is a mean (the data items take the value 1 for a person with a beard, and 0 without a beard). If we select people according to our model, then ask them whether they have a beard, then our estimate of the percentage of beards should behave as above.

The first thing that should strike you is that it isn't at all easy to select people according to this model. For example, we might select phone numbers at random, then call and ask the first person to answer the phone whether they have a beard;

but many children won't answer the phone because they are too small. The next important problem is that errors in selecting people can lead to massive errors in your estimate. For example, imagine you decide to survey all of the people at a kindergarten on a particular day; or all of the people in a women's clothing store; or everyone attending a beard growing competition (they do exist). In each case, you will get an answer that is a very poor estimate of the right answer, and the standard error might look very small. Of course, it is easy to tell that these cases are a bad choice.

It may not be easy to tell what a good choice is. You should notice the similarity between estimating the percentage of the population that wears a beard, and estimating the percentage that will vote for a particular candidate. There is a famous example of a survey that mispredicted the result of the Dewey-Truman presidential election in 1948; poll-takers phoned random phone numbers, and asked for an opinion. But at that time, telephones tended to be owned by a small percentage of rather comfortable households, who tended to prefer one candidate, and so the polls mispredicted the result rather badly.

Sometimes, we don't really have a choice of samples. For example, we might be presented with a small dataset of (say) human body temperatures. If we can be satisfied that the people were selected rather randomly, we might be able to use this dataset to predict expected body temperature. But if we knew that the subjects had their temperatures measured because they presented themselves at the doctor with a suspected fever, then we most likely cannot use it to predict expected body temperature.

One important and valuable case where this model works is in simulation. If you can guarantee that your simulations are independent (which isn't always easy), this model applies to estimates obtained from a simulation. Notice that it is usually straightforward to build a simulation so that the  $i$ 'th simulation reports an  $x_i$  where  $\text{popmean}(\{x\})$  gives you the thing you want to measure. For example, imagine you wish to measure the probability of winning a game; then the simulation should report one when the game is won, and zero when it is lost. As another example, imagine you wish to measure the expected number of turns before a game is won; then your simulation should report the number of turns elapsed before the game was won.

8.4 YOU SHOULD

8.4.1 be able to:

- Write out the likelihood for a set of independent data items produced by models from chapter 6 (at least Normal, Binomial, Multinomial, Poisson, Beta, Gamma, Exponential).
- Write out the log likelihood for a set of independent data items produced by models from chapter 6 (at least Normal, Binomial, Multinomial, Poisson, Beta, Gamma, Exponential).
- Find maximum likelihood solutions for parameters of these models from a set of independent data items.
- Describe situations where maximum likelihood estimates might not be reliable.
- Describe the difference between maximum likelihood estimation and Bayesian inference.
- Write an expression for the posterior or log-posterior of model parameters given a set of independent data items.
- Compute the MAP estimate for the cases shown in the worked examples.
- Compute on-line estimates of the maximum likelihood estimate of the mean and standard deviation of a normal model.
- Compute on-line estimates of the MAP estimate of the mean and standard deviation in the case of a normal prior and a normal likelihood.
- Estimate the population mean from a sample mean.
- Estimate the standard error of the estimate of a population mean.

8.4.2 remember:

New term: Inference . . . . .	221
New term: point estimate . . . . .	221
New term: interval estimates . . . . .	221
New term: hypothesis testing . . . . .	221
New term: likelihood . . . . .	223
Definition: Likelihood . . . . .	223
New term: maximum likelihood principle . . . . .	223
Definition: The maximum likelihood principle . . . . .	223
New term: independent and identically distributed . . . . .	223
New term: IID . . . . .	223
Definition: The log-likelihood of a dataset under a model . . . . .	227
New term: consistency . . . . .	230
New term: prior probability distribution . . . . .	232
New term: posterior . . . . .	232

New term: Bayesian inference . . . . . 232  
 Definition: Bayesian inference . . . . . 232  
 New term: maximum a posteriori estimate . . . . . 232  
 New term: MAP estimate . . . . . 232  
 Definition: MAP estimate . . . . . 232  
 New term: conjugate . . . . . 234  
 Useful fact: Normal distributions are conjugate . . . . . 237  
 Useful fact: Normal posteriors . . . . . 239  
 New term: joint . . . . . 239  
 New term: filtering . . . . . 241  
 Useful fact: Online updating of normal posteriors . . . . . 242  
 New term: population . . . . . 244  
 New term: sample . . . . . 244  
 New term: population mean . . . . . 244  
 New term: sample mean . . . . . 244  
 Useful facts: Sample means and population means . . . . . 245  
 Useful fact: The sample mean . . . . . 248  
 New term: standard error . . . . . 248

## PROBLEMS

## Maximum Likelihood Methods

- 8.1. Fitting a Normal Distribution:** You are given a dataset of  $N$  numbers. Write  $x_i$  for the  $i$ 'th number. You wish to model this dataset with a normal distribution.
- Show the maximum likelihood estimate of the mean of this distribution is  $\text{mean}(\{x\})$ .
  - Show the maximum likelihood estimate of the standard deviation of this distribution is  $\text{std}(x)$ .
  - Now assume that all of these numbers take the same value - what happens to your estimate of the standard deviation?
- 8.2. Fitting a Poisson Distribution:** You count the number of times that the annoying “MacSweeper” popup window appears per hour when you surf the web. You wish to model these counts with a Poisson distribution. On day 1, you surf for 4 hours, and see counts of 3, 1, 4, 2 (in hours 1 through 4 respectively). On day 2, you surf for 3 hours, and observe counts of 2, 1, 2. On day 3, you surf for 5 hours, and observe counts of 3, 2, 2, 1, 4. On day 4, you surf for 6 hours, but keep only the count for all six hours, which is 13. You wish to model the intensity in counts per hour.
- What is the maximum likelihood estimate of the intensity for each of days 1, 2, and 3 *separately*?
  - What is the maximum likelihood estimate of the intensity for day 4?
  - What is the maximum likelihood estimate of the intensity for all days taken together?
- 8.3. Fitting a Geometric Model:** You wish to determine the number of zeros on a roulette wheel without looking at the wheel. You will do so with a geometric model. Recall that when a ball on a roulette wheel falls into a non-zero slot, odd/even bets are paid; when it falls into a zero slot, they are not paid. There are 36 non-zero slots on the wheel.
- Assume you observe a total of  $r$  odd/even bets being paid before you see a bet not being paid. What is the maximum likelihood estimate of the number of slots on the wheel?
  - How reliable is this estimate? Why?
  - You decide to watch the wheel  $k$  times to make an estimate. In the first experiment, you see  $r_1$  odd/even bets being paid before you see a bet not being paid; in the second,  $r_2$ ; and in the third,  $r_3$ . What is the maximum likelihood estimate of the number of slots on the wheel?
- 8.4. Fitting a Binomial Model:** You encounter a deck of Martian playing cards. There are 87 cards in the deck. You cannot read Martian, and so the meaning of the cards is mysterious. However, you notice that some cards are blue, and others are yellow.
- You shuffle the deck, and draw one card. It is yellow. What is the maximum likelihood estimate of the fraction of blue cards in the deck?
  - You repeat the previous exercise 10 times, replacing the card you drew each time before shuffling. You see 7 yellow and 3 blue cards in the deck. What is the maximum likelihood estimate of the fraction of blue cards in the deck?

## Bayesian Methods

- 8.5. Zeros on a Roulette Wheel:** We now wish to make a more sophisticated estimate of the number of zeros on a roulette wheel without looking at the wheel. We will do so with Bayesian inference. Recall that when a ball on a roulette wheel falls into a non-zero slot, odd/even bets are paid; when it falls into a zero slot, they are not paid. There are 36 non-zero slots on the wheel. We assume that the number of zeros is one of  $\{0, 1, 2, 3\}$ . We assume that these cases have prior probability  $\{0.1, 0.2, 0.4, 0.3\}$ .
- Write  $n$  for the event that, in a single spin of the wheel, an odd/even bet will not be paid (equivalently, the ball lands in one of the zeros). Write  $z$  for the number of zeros in the wheel. What is  $P(n|z)$  for each of the possible values of  $z$  (i.e. each of  $\{0, 1, 2, 3\}$ )?
  - Under what circumstances is  $P(z = 0|\text{observations})$  NOT 0?
  - You observe 36 independent spins of the same wheel. A zero comes up in 2 of these spins. What is  $P(z|\text{observations})$ ?
- 8.6. A Normal Distribution:** You are given a dataset of 3 numbers,  $-1, 0, 20$ . You wish to model this dataset with a normal distribution with unknown mean  $\mu$  and standard deviation 1. You will make a MAP estimate of  $\mu$ . The prior on  $\mu$  is normal, with mean 0 and standard deviation 10.
- What is the MAP estimate of  $\mu$ ?
  - A new datapoint, with value 1, arrives. What is the new MAP estimate of  $\mu$ ?

## Samples and Populations

- 8.7. The Average Mouse:** You wish to estimate the average weight of a mouse. You obtain 10 mice, sampled uniformly at random and with replacement from the mouse population. Their weights are 21, 23, 27, 19, 17, 18, 20, 15, 17, 22 grams respectively.
- What is the best estimate for the average weight of a mouse, from this data?
  - What is the standard error of this estimate?
  - How many mice would you need to reduce the standard error to 0.1?
- 8.8. Sample Variance and Standard Error:** You encounter a deck of Martian playing cards. There are 87 cards in the deck. You cannot read Martian, and so the meaning of the cards is mysterious. However, you notice that some cards are blue, and others are yellow.
- You shuffle the deck, and draw one card. You repeat this exercise 10 times, replacing the card you drew each time before shuffling. You see 7 yellow and 3 blue cards in the deck. As you know, the maximum likelihood estimate of the fraction of blue cards in the deck is 0.3. What is the standard error of this estimate?
  - How many times would you need to repeat the exercise to reduce the standard error to 0.05?

# Inference: Intervals and Testing

Although a point estimate is the best estimate available, it could still be wrong. This doesn't always matter. But in some applications, it can be important to know what range a parameter could take, and still be consistent with the data. We have seen this possibility before, in example 8.8, where I graphed the posterior — you could look at that graph and see that  $p$  could be in a fairly large range of values.

Being able to estimate an interval of values is important, particularly when there are safety or legal considerations to worry about. Imagine you have a machine that fills cereal boxes. Each box gets a quantity of cereal that is random, but has low variance. If the weight of cereal in any box is below the amount printed on the label, you might be in trouble. When you choose the amount to print, a point estimate of the mean weight of cereal might not be particularly helpful. If that estimate is a little low, you could have problems. Instead, what you'd like to know is an interval that the mean lies in with very high probability, so you can print a number that is smaller than the smallest in the interval.

Such intervals are usually known as **confidence intervals**. There are a variety of techniques for estimating them, depending on the particular problem. In section 9.1, I describe straightforward methods to construct confidence intervals. Confidence intervals for sample means are easily constructed using our understanding of standard error. Similarly, the posterior yields confidence intervals in a straightforward way when we use Bayesian inference. In other cases we might not have a good estimate of standard error; but we can use simulation to construct one (section 9.2). That section also shows how to exploit simulation methods to estimate confidence intervals for maximum likelihood parameter estimates.

Now imagine we hypothesize that (say) a population has a given mean. We can then ask how large a confidence interval we have to draw around the sample mean to include the hypothesis. If that interval is very large, then we would have to have drawn a very unusual sample to see the sample mean that we see. This means that our data does not support the hypothesis. This is an extremely fruitful line of reasoning, because it allows us to evaluate evidence in support of a hypothesis. Section 9.3 describes methods to do so. In fact, we can evaluate evidence in support of a *model*, too, and section 9.4 shows how to do this.

## 9.1 STRAIGHTFORWARD CONSTRUCTIONS OF CONFIDENCE INTERVALS

A **statistic** is a function of a dataset. One example of a statistic is the mean of a sample. Yet another is the MAP estimate of a parameter (which you get by computing some numbers from the dataset alone). We observe the value of a statistic, which we can compute from our dataset. But the dataset is random, because it is either a sample from a population or an IID sample from a distribution mode. This means that we should think of the statistic as the observed value of a

random variable — if we had a different sample from the same population (resp. IID sample from the same distribution) we would compute a different value. In these two cases, we know the observed value of the statistic *and* a model for the probability distribution of the statistic, as a random variable. This means we can construct confidence intervals relatively easily.

### 9.1.1 Confidence Intervals for Population Means

Assume we have a population, and we wish to know a confidence interval for its mean. We draw a sample  $\{x\}$  of  $k$  items, and compute the sample mean. This is the value of a random variable — random, because it depends on the randomly drawn sample — whose probability distribution we know. This knowledge is very powerful, because it tells us how close to the true mean our estimate is likely to be. The reasoning looks like this. Our estimate of the unknown number  $\text{popmean}(\{x\})$  is the mean of the sample we have, which we write  $\text{mean}(\{x\})$ . We know that the error — which is

$$\text{mean}(\{x\}) - \text{popmean}(\{x\})$$

— is a normal random variable, with mean 0 and standard deviation  $\frac{\text{popstd}(\{x\})}{\sqrt{k}}$ .

We cannot evaluate this standard deviation because we do not — and, in practice, can not — know  $\text{popstd}(\{x\})$ . But we can *estimate*  $\text{popstd}(\{x\})$ . The population standard deviation is the square root of the population variance, and the population variance is the average of  $(x_i - \text{mean}(\{x\}))^2$  over the population. As a result, when we sample the population, the sample variance behaves in a familiar way. The sample variance is a random variable (because the sample was chosen randomly). Its expected value is the population variance (because it's the mean of something). Rather than determine the standard deviation of the sample variance, we will assume that it is small (which is true for a reasonably sized sample), and then pretend that there is no *significant* variation in the value from sample to sample. If we make these assumptions, then we can estimate  $\text{popstd}(\{x\}) \approx \text{std}(\{x\})$ .

With this approximation, we estimate the standard deviation of the sample mean to be  $\frac{\text{std}(\{x\})}{\sqrt{k}}$ . We can scale the error by the standard deviation to get

$$\hat{M} = \frac{\text{mean}(\{x\}) - \text{popmean}(\{x\})}{\left(\frac{\text{popstd}(\{x\})}{\sqrt{k}}\right)} \approx \frac{\text{mean}(\{x\}) - \text{popmean}(\{x\})}{\left(\frac{\text{std}(\{x\})}{\sqrt{k}}\right)}$$

which is a standard normal random variable. But we know rather a lot about the behaviour of standard normal random variables. In particular, we have

**Useful Facts: 9.1** *Normal random variables*

- About 68% of the time, the value of a standard normal random variable is within one standard deviation of the mean.
- About 95% of the time, the value of a standard normal random variable is within two standard deviations of the mean.
- About 99% of the time, the value of a standard normal random variable is within three standard deviations of the mean.

In turn, this means that

- About 68% of the time,  $\text{abs}(\text{mean}(\{x\}) - \text{popmean}(\{x\})) \leq \frac{\text{popstd}(\{x\})}{\sqrt{k}}$ .
- About 95% of the time,  $\text{abs}(\text{mean}(\{x\}) - \text{popmean}(\{x\})) \leq 2 \frac{\text{popstd}(\{x\})}{\sqrt{k}}$ .
- About 99% of the time,  $\text{abs}(\text{mean}(\{x\}) - \text{popmean}(\{x\})) \leq 3 \frac{\text{popstd}(\{x\})}{\sqrt{k}}$ .

Indexing: Normal random variables

Indexing: Normal random variables

This means that we can construct a **confidence interval** for our estimate. For about 68% of samples, the sample mean will lie in the interval between  $\text{mean}(\{x\}) - \frac{\text{std}(\{x\})}{\sqrt{k}}$  and  $\text{mean}(\{x\}) + \frac{\text{std}(\{x\})}{\sqrt{k}}$ , and so on.

We can plot the confidence interval by drawing **error bars** — draw a bar one (or two, or three) standard deviations up and down from the estimate. We interpret this interval as representing the effect of sampling uncertainty on our estimate. If the urn model really did apply, then the confidence intervals have the property that the true mean lies inside the interval for about 68% of possible samples (for one standard error error bars; or 95% for two; etc.).

**Worked example 9.1**    *Speed of light*

Give a 95% confidence interval for the time of a light transit across a fixed distance, based on the dataset at <http://lib.stat.cmu.edu/DASL/Datafiles/SpeedofLight.html>.

**Solution:** The first thing to notice about this dataset is it seems to contain two negative times; but this is illusory. Look closely by tracking down on the web the paper referred to in the story (Stigler, S.M., “Do robust estimators work with real data?” *Annals of Statistics*, 5 (1977), pp. 1055-1078.). You will find that Newcomb measured the the passage of time for light to cross a fixed distance. He then subtracted a constant, and reported the result. Write  $t$  for the true measurement of time, in microseconds, and  $n$  for the number in the dataset. We have

$$t = 24.8 + n \times 1e - 3.$$

Newcomb omitted the smallest value in the dataset, which is very different from all others. **Omitting** the smallest value gives a mean of 24.82729, a standard deviation of 0.0062, and a standard error of 0.00078. So with 95% confidence the value lies between 24.82573 and 24.82885. **Keeping** the smallest value gives a mean of 24.8262, a standard deviation of 0.0107 and a standard error of 0.0013. So with 95% confidence the value lies between 24.8236 and 24.8288. Notice the temptation to drop the outlier — doing so really makes the measurement look better. These days, the speed of light is known precisely. It’s used to define the metre, so we would take the position that Newcomb was measuring the distance of the transit.

It is quite straightforward, and informative, to verify these ideas with a simulation. I used the heights column from the bodyfat dataset (from <http://www2.stetson.edu/~jrasp/data.htm>; look for bodyfat.xls). I removed the single height outlier. I simulated the population using the whole dataset (251 items), then drew numerous samples of various sizes, with replacement. I computed the mean of each of these sets of samples. Figure 9.1 shows a scatter plot of sample means for different samples, using a set of sizes (2, 4, 9, 16, . . . , 100). I have also plotted the population mean, and the true 1-standard error bars (i.e. using the population standard deviation) for each of these sample sizes. Notice how most sample means lie within the 1-standard error bars, as they should.

Figure 9.2 shows a sample mean and the computed standard error bars for a set of different sized samples. I have also shown the population mean. Notice how the population mean is usually within the computed standard error bars. Notice also how the computed standard error bars (i.e. using the sample standard deviation) are a bit small for small samples, but compare well with the true values for large samples. The problem here is that, for small samples, the sample standard deviation can be a fairly poor estimate of the population standard deviation. In the figure, I show a scatter plot of standard error estimated using sample standard deviation against standard error evaluated using population standard deviation. When the

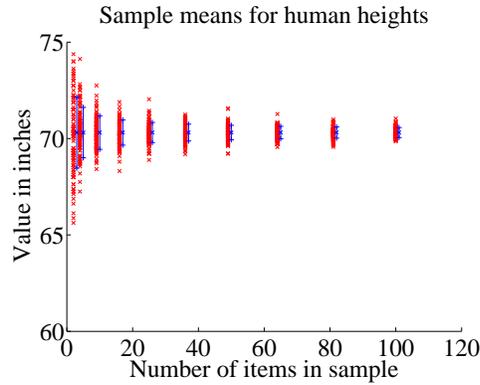


FIGURE 9.1: I took the heights dataset. I then formed sampled elements with replacement to form random subsets of sizes  $(2, 4, 9, 16, \dots, 100)$ . For each of 100 subsets of each size, I computed the sample mean — these are shown as  $x$ 's on the plot. I then computed the population mean, and the standard error as measured by the population standard deviation. The  $x$  to the side of each column is the population mean, and the vertical bars are one standard error above and below the population mean. Notice how (a) the sample means vary less as the sample gets bigger and (b) the sample means largely lie within the error bars.

standard error is small (because the sample is large), the two are quite close; when it is large (because the sample is small), they can be quite different. This suggests, correctly, that accurate inference from small samples can be tricky. The right thing to do is get more data — it is optimistic to hope for a good estimate of average human height from two people, after all.

Now we might reasonably ask another question. We should like to know a confidence interval such that the true mean lies within it for  $p$  of the samples. This is equivalent to asking what is the  $u$  such that

$$\int_{-u}^u \frac{1}{\sqrt{2\pi}} \exp(-x^2/2) dx = p$$

(i.e. what is the range of values about zero such that  $p\%$  of standard normal random variables lies in this range). Such numbers can be extracted from the inverse of the error function (which is known as the **inverse error function**). An alternative is to look in statistical tables, which usually give  $u$  values for a useful set of  $p$ 's.

### 9.1.2 Bayesian Confidence Intervals

We can use the posterior to evaluate the probability that the true value of a parameter lies within an interval. Assume we want to know for  $a < b$ , the probability that the parameter  $\theta$  lies in that interval. In many cases, we can compute

$$\int_a^b p(\theta|\mathcal{D}) d\theta,$$

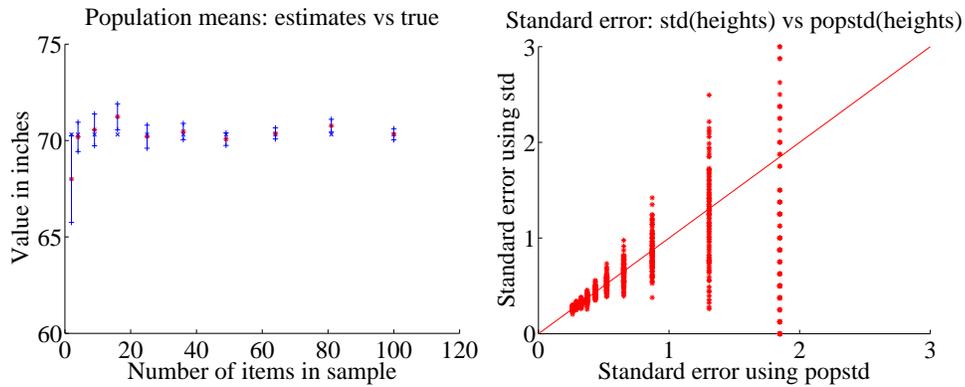


FIGURE 9.2: On the **right**, I chose one sample at random of each size; the sample mean is shown as a \*. There are error bars (one standard error above and below) around the sample mean. These error bars are computed from the sample standard deviation. The population mean is the  $x$ . Notice how the population mean is within the error bars most, but not all, of the time (about 68% of the time, as they should be). The sample mean is rather a good estimate of the population mean, and the standard error is quite a reliable estimate of how well the sample mean represents the population mean when the sample is large enough. When the sample is small, the estimate is poor, as you can see on the **right**. This is a scatter plot of the true standard error (using the population standard deviation) against the estimate (using the sample standard deviation). When the sample is small, the standard error is big, and the estimate is poor.

which supplies the relevant probability.

**Worked example 9.2** *Flipping a coin - II*

We have a coin with probability  $\theta$  of coming up heads when flipped. We start knowing nothing about  $\theta$  except that, being a probability,  $0 \leq \theta \leq 1$ . We use a uniform prior on  $\theta$ . We then flip the coin 10 times, and see 7 heads (and 3 tails). What is  $P(\{\theta \in [0.5, 0.8]\} | \mathcal{D})$ ?

**Solution:** We could answer this question by computing

$$\int_{0.5}^{0.8} p(\theta | \mathcal{D}) d\theta.$$

Recall that a Beta distribution with  $\alpha = \beta = 1$  is uniform, and look at example 8.11. The posterior is

$$P(\theta | 10, 7, 1, 1) = \frac{\Gamma(12)}{\Gamma(8)\Gamma(4)} \theta^7 (1 - \theta)^3.$$

I evaluated this integral numerically, and got 0.73.

In example 9.2, I determined that  $P(\{\theta \in [0.5, 0.8]\} | \mathcal{D})$  was 0.73. Instead, I might wish to specify  $0 \leq u < 0.5$ , and then find an interval  $[a, b]$  such that  $P(\{\theta \in [a, b]\} | \mathcal{D}) = 1 - 2u$ . Notice that this interval is not unique. Notice that

$$P(\{\theta \in [a, b]\} | \mathcal{D}) = 1 - P(\{\theta \leq a\} | \mathcal{D}) - P(\{\theta \geq b\} | \mathcal{D})$$

We will choose  $a$  such that

$$P(\{\theta \leq a\} | \mathcal{D}) = \int_{-\infty}^a P(\theta | \mathcal{D}) d\theta = u$$

and  $b$  such that

$$P(\{\theta \geq b\} | \mathcal{D}) = \int_b^{\infty} P(\theta | \mathcal{D}) d\theta = u.$$

Actually obtaining values for  $a$  and  $b$  might be quite difficult. One strategy is to search a range of values.

**Worked example 9.3** *Flipping a coin - II*

We have a coin with probability  $\theta$  of coming up heads when flipped. We start knowing nothing about  $\theta$  except that, being a probability,  $0 \leq \theta \leq 1$ . We use a uniform prior on  $\theta$ . We then flip the coin 10 times, and see 7 heads (and 3 tails). Construct an interval  $[a, b]$  such that  $P(\{\theta \leq a\}|\mathcal{D}) \approx 0.05$  and  $P(\{\theta \geq b\}|\mathcal{D}) \approx 0.05$ . This is a 90% confidence interval.

**Solution:** The way to do this is to find an  $a$  such that  $\int_{-\infty}^a p(\theta|\mathcal{D})d\theta \approx 0.05$ , and a  $b$  such that  $\int_b^{\infty} p(\theta|\mathcal{D})d\theta \approx 0.05$ . I was only able to do these integrals numerically. I wrote a brief program to compute the relevant integrals numerically, then searched for  $a$  and  $b$  using interval halving. I found  $P(\{\theta \leq 0.435\}|\mathcal{D}) \approx 0.05$  and  $P(\{\theta \geq 0.865\}|\mathcal{D}) \approx 0.05$ ; this means the interval is  $[0.435, 0.865]$ .

## 9.2 USING SIMULATION TO CONSTRUCT INTERVALS

The analysis of some problems is difficult. But the idea of a confidence interval is naturally associated with repeated experiments. We can use simulations to estimate confidence intervals when analysis is difficult. Generally, we see the dataset as one of many possible dataset we could have seen. We must then assess what our estimate would be like for the other possible datasets. But what are those possible datasets to be?

When the data is explained by a parametric probability model, we can use that model to produce other possible datasets. If we compute a maximum likelihood estimate  $\hat{\theta}$  of the parameters of the model, we can draw IID samples *from the model*. We then look at the estimate from that new dataset; the spread of the estimates yields our confidence interval.

When we have no applicable probability model of the dataset, we can resample the dataset to simulate the effects of sampling error. This strategy allows us to build confidence intervals for properties like medians.

## 9.2.1 Constructing Confidence Intervals for Parametric Models

Assume we have a dataset  $\mathcal{D} = \{\mathbf{x}\}$ , and a probability model we believe applies to that dataset. We know how to estimate appropriate parameter values by maximizing the likelihood function  $L(\theta)$ . But we do not yet have a way to think about how accurately the data determines the best choice of parameter value. For example, if we change the dataset slightly, will the estimated parameter change a lot? or a little?

A  $(1 - 2\alpha)$  confidence interval for a parameter is an interval  $[c_\alpha, c_{(1-\alpha)}]$ . We construct the interval such that, if we were to perform a very large number of repetitions of our original experiment then estimate a parameter value for each,  $c_\alpha$  would be the  $\alpha$  quantile and  $c_{(1-\alpha)}$  would be the  $(1 - \alpha)$  quantile of those parameter values. We interpret this to mean that, with confidence  $(1 - 2\alpha)$ , the correct value

of our parameter lies in this interval. This definition isn't really watertight. How do we perform a very large number of repetitions? If we don't, how can we tell how good the confidence interval is? Nonetheless, we can construct intervals that illustrate how sensitive our original inference is to the data that we have.

There are a variety of methods to construct confidence intervals. We will focus on simulation. The algorithm should feel so natural to you that you may already have guessed what to do. First, we compute the maximum likelihood estimate of the parameters,  $\hat{\theta}$ . We assume this estimate is right, but need to see how our estimates of  $\hat{\theta}$  would vary with different collections of data from our model with that parameter value. So we compute a collection of simulated datasets,  $\mathcal{D}_i$ , each the same size as the original dataset. We obtain these by simulating  $P(d|\hat{\theta})$ . Next, we compute a maximum likelihood estimate for each of these simulated datasets,  $\hat{\theta}_i$ . Finally, we compute the relevant percentiles of these datasets. The result is our interval.

**Procedure: 9.1** *Estimating Confidence Intervals for Maximum Likelihood Estimates using Simulation*

Assume we have a dataset  $\mathcal{D} = \{x\}$  of  $N$  items. We have a parametric model of this data  $p(x|\theta)$ , and write the likelihood  $\mathcal{L}(\theta; \mathcal{D}) = P(\mathcal{D}|\theta)$ . We construct  $(1 - 2\alpha)$  confidence intervals using the following steps.

1. Compute the maximum likelihood estimate of the parameter,  $\hat{\theta} = \underset{\theta}{\operatorname{argmax}} \mathcal{L}(\theta; \mathcal{D})$ .
2. Construct  $R$  simulated datasets  $\mathcal{D}_i$ , each consisting of  $N$  IID samples drawn from  $p(x|\hat{\theta})$ .
3. For each such dataset, compute  $\hat{\theta}_i = \underset{\theta}{\operatorname{argmax}} \mathcal{L}(\theta; \mathcal{D}_i)$ .
4. Obtain  $c_\alpha(\hat{\theta}_i)$ , the  $\alpha$ 'th quantile of the collection  $\hat{\theta}_i$  and  $c_{(1-\alpha)}(\hat{\theta}_i)$ , the  $1 - \alpha$ 'th quantile of the collection  $\hat{\theta}_i$ .

The confidence interval is  $[c_\alpha, c_{(1-\alpha)}]$ .

Figure 9.3 shows an example. In this case, I worked with simulated data from a normal distribution. In each case, the normal distribution had mean 0, but there are four different standard deviations (1, 4, 7, and 10). I simulated 10 different datasets from each of these distributions, containing 10, 40, 90, ..., 810, 1000 data items. For each, I computed the maximum likelihood estimate of the mean. This isn't zero, *even though the data was drawn from a zero-mean distribution*, because the dataset is finite. I then estimated the confidence intervals for each using 10,000 simulated datasets of the same size. I show 95% confidence intervals for the cases  $\sigma = 1$ ,  $\sigma = 4$ ,  $\sigma = 10$  in the figure, plotted against dataset size. Notice that these aren't symmetric about zero, because the maximum likelihood estimate isn't zero.

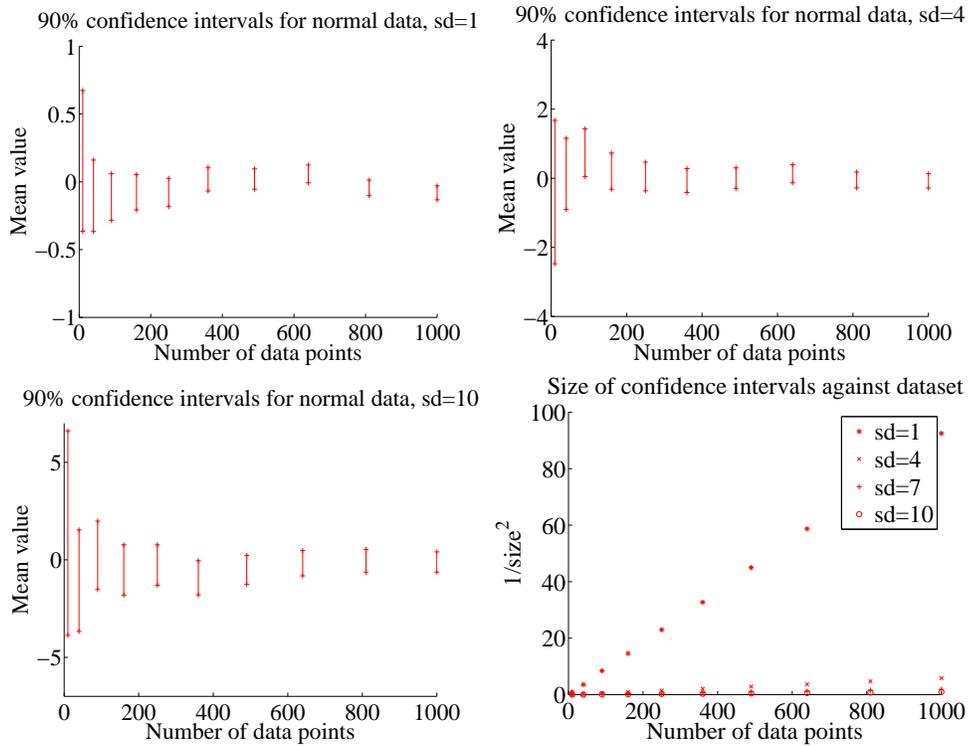


FIGURE 9.3: Confidence intervals computed for simulated normal data; details in the text.

They shrink as the dataset grows, but slowly. They are bigger when the standard deviation is bigger. It should seem reasonable that you can't expect an accurate estimate of the mean of a normal distribution with large standard deviation using only a few data points. One can demonstrate using statistical theory that would take us rather out of our way that the size of these intervals should behave like

$$\frac{\sigma}{\sqrt{N}}$$

and Figure 9.3 suggests that they do. I plotted  $1/\text{size}^2$ ; you can see this grows linearly as  $N$  grows, and falls very fast as  $\sigma$  grows.

**Worked example 9.4** *Confidence Intervals by Simulation - II*

Construct a 90% confidence interval for the intensity estimate for the data of example 8.5 for the cases of 10 observations, 20 observations, and all 30 observations.

**Solution:** Recall from that example the maximum likelihood estimates of the intensity are  $7/10$ ,  $22/20$ , and  $29/30$  in the three cases. I used the Matlab function `poissrnd` to get 10000 replicates of a dataset of 10 (resp. 20, 30) items from a Poisson distribution with the relevant intensities. I then used `prctile` to get the 5% and 95% percentiles, yielding the intervals

[0.3, 1.2]	for 10 observations
[0.75, 1.5]	for 20 observations
[0.6667, 1.2667]	for 30 observations

Notice how having more observations makes the confidence interval smaller.

### 9.2.2 Estimating Standard Error

We were able to produce convenient and useful estimates of standard error for sample means. But what happens if we want to, say, reason about the median of a population? It turns out that estimating standard error becomes difficult mathematically. This is an important problem, because if we had standard error estimates, we could use our methods for building confidence intervals and for testing hypotheses. It turns out that quite simple simulation methods give very good estimates.

For Figure 9.4, I assumed that all 253 weight measurements represented the entire population, then simulated what would happen for different random samples (with replacement) of different sizes. Figure 9.4 suggests that the sample median behaves quite like the sample mean as the random sample changes. Different samples have different medians, but the distribution of values looks fairly normal. When there are more elements in the sample, the standard deviation of median values is smaller. But we have no method to compute this standard deviation. The method I used here doesn't apply in practice, because we don't usually have the whole population.

There is a method, known as the **bootstrap**, which gives a very good estimate of the standard error of any statistic. Assume we wish to estimate the standard error of a statistic  $S(\{\mathbf{x}\})$ , which is a function of our dataset  $\{\mathbf{x}\}$  of  $N$  data items. We compute  $r$  **bootstrap replicates** of this sample. Each replicate is obtained by sampling the dataset uniformly, and with replacement. One helpful way to think of this is that we are modelling our dataset as a sample of a probability distribution. This distribution, sometimes known as the **empirical distribution**, has probability  $1/N$  at each of the data items we see, and zero elsewhere. Now to obtain replicates, we simply draw new sets of IID samples from this probability distribution.

Notice that the bootstrap replicates are *not* a random permutation of the

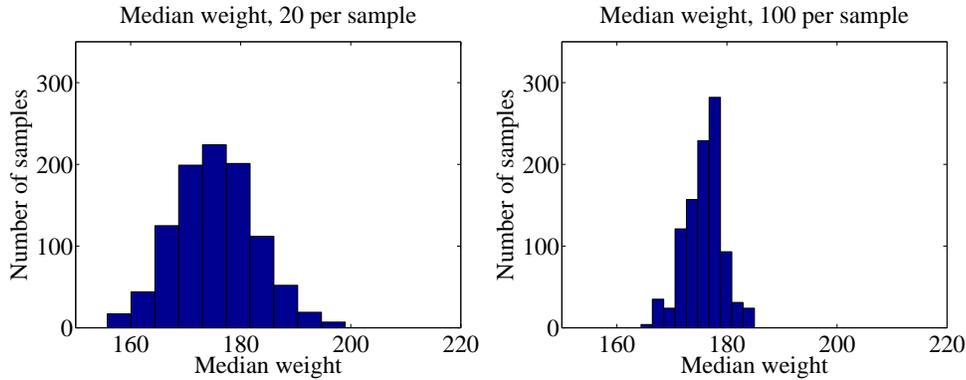


FIGURE 9.4: I took the weights dataset used all 253 measurements to represent a population. Rather than compute the median of the whole population, I chose to compute the median of a randomly chosen sample. The figures show a histogram of 1000 different values of the median, computed for 1000 different samples (of size 20 on the **left**, and of size 100 on the **right**). Notice that (a) there is a moderate amount of variation in the median of the sample; (b) these histograms look normal, and appear to have about the same median; (c) increasing the size of the sample has reduced the spread of the histogram.

dataset; instead, we select one data item fairly and at random from the whole dataset  $N$  times. This means we expect a particular bootstrap replicate will have multiple copies of some data items, and no copies of others.

We write  $\{\mathbf{x}\}_i$  for the  $i$ 'th bootstrap replicate of the dataset. We now compute

$$\bar{S} = \frac{\sum_i S(\{\mathbf{x}\}_i)}{r}$$

and the standard error estimate for  $S$  is given by:

$$\text{stderr}(\{S\}) = \sqrt{\frac{\sum_i [S(\{\mathbf{x}\}_i) - \bar{S}]^2}{r - 1}}$$

**Procedure: 9.2** *The bootstrap*

Estimate the standard error for a statistic  $S$  evaluated on a dataset of  $N$  items  $\{\mathbf{x}\}$ .

1. Compute  $r$  bootstrap replicates of the dataset. Write the  $i$ 'th replicate  $\{\mathbf{x}\}_i$ . Obtain each by:

- (a) Building a uniform probability distribution on the numbers  $1 \dots N$ .
- (b) Drawing  $N$  independent samples from this distribution. Write  $s(i)$  for the  $i$ 'th such sample.
- (c) Building a new dataset  $\{x_{s(1)}, \dots, x_{s(N)}\}$ .

2. For each replicate, compute  $\sum_i S(\{\mathbf{x}\}_i)$ .

3. Compute

$$\bar{S} = \frac{\sum_i S(\{\mathbf{x}\}_i)}{r}$$

4. The standard error estimate for  $S$  is given by:

$$\text{stderr}(\{S\}) = \sqrt{\frac{\sum_i [S(\{\mathbf{x}\}_i) - \bar{S}]^2}{r - 1}}$$

**Worked example 9.5** *The bootstrap standard error of the median*

You can find a dataset giving the salaries of CEO's at small firms in 1993 at <http://lib.stat.cmu.edu/DASL/Datafiles/ceodat.html>. Construct a 90% confidence interval for the median salary.

**Solution:** Salaries are in thousands of dollars, and one salary isn't given (we omit this value in what follows). Figure 9.5 shows a histogram of the salaries; notice there are some values that look like outliers. This justifies using a median. The median of the dataset is 350 (i.e. \$ 350,000 — this is 1993 data!). I constructed 10,000 bootstrap replicates. Figure 9.5 shows a histogram of the medians of the replicates. I used the matlab `prctile` function to extract the 5% and 95% percentiles of these medians, yielding the interval between 298 and 390. This means that we can expect that, for 90% of samples of CEO salaries for small companies, the median salary will be in the given range.

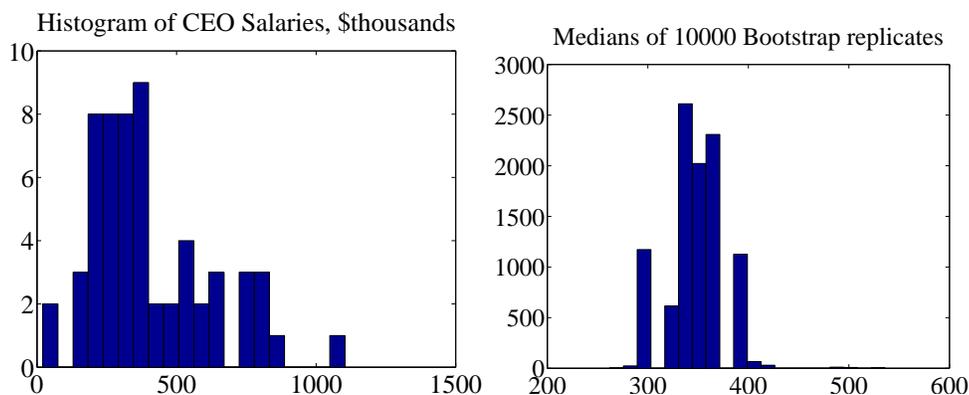


FIGURE 9.5: On the **left**, a histogram of salaries for CEO's of small companies in 1993, from the dataset of <http://lib.stat.cmu.edu/DASL/Datafiles/ceodat.html>. On the **right**, a histogram of the medians of 10, 000 bootstrap replicates of this data. This simulates the effect of sampling variation on the median; see worked example 9.5.

**Worked example 9.6** *Are Brinks collectors different from city collectors?*

Using the dataset of <http://lib.stat.cmu.edu/DASL/Datafiles/brinkdat.html>, test whether Brinks coin collectors are different from city coin collectors (read the story, at that URL).

**Solution:** There are two types of parking meter in this story (which you should have read!). One is close to the city hall, and is always collected by city employees. The other is far, and for some months was collected by Brinks employees and for some months collected by city employees. I ignored the figures for month 10, because there was a missing digit in the number (don't know how or why). For the meters that are far, coins were collected for 22 months by city employees and for 24 months by Brinks employees.

For the far meters, I constructed 10, 000 bootstrap replicates of each case. Figure 9.6 compares the two sets of replicates, showing the histogram of means for city collections vs Brinks collections. Notice how the means of the two histograms are quite separated, compared to their variance. This suggests, quite strongly, that the Brinks collections are different (the difference between the means is about 1.6 standard deviations). I used the nearby meter figures to check whether there was something special about the months on which Brinks employees collected coins. Again, I constructed 10, 000 bootstrap replicates of each case (nearby coins in a Brinks month vs. nearby coins in a city month). This is plotted in figure 9.6. There's nothing special about the months, from the figure.

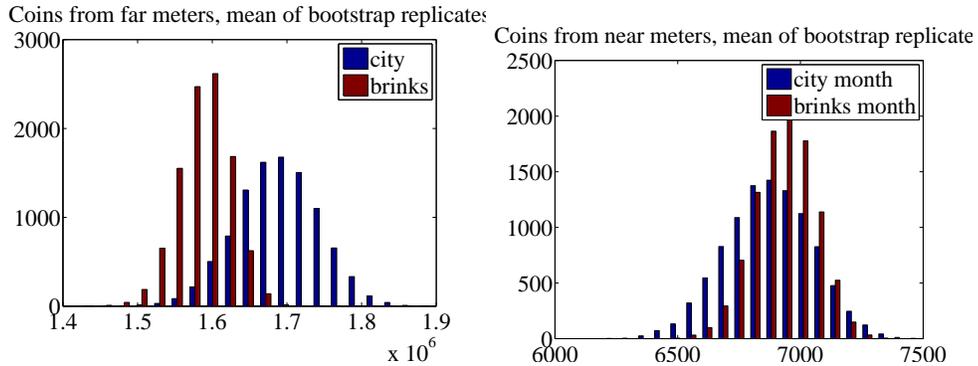


FIGURE 9.6: *On the left*, two histograms comparing the means of bootstrap replicates of the coins collected by Brinks employees and those collected by city employees. Notice how the two histograms are quite separated, compared to their width. This is fairly strong evidence that the two cases are different. The difference is not explained by the months in which the coins are collected. *On the right*, two histograms comparing the means of bootstrap replicates of coins collected from nearby meters. These coins were collected by city employees, but I have separated Brinks months from city months. There is no evidence the months are different.

### 9.3 USING THE STANDARD ERROR TO EVALUATE HYPOTHESES

Assume we have a **hypothesis** to work with. For example, we believe that women and men have the same body temperature. We should like to assess the extent to which the evidence we have contradicts the hypothesis. At first glance, this may seem strange to you — surely one wants to assess the extent to which the evidence *supports* the hypothesis — but in fact it's natural. You can't prove that a scientific hypothesis is true; you can only fail to show that it's false. Just one piece of evidence can destroy a scientific hypothesis, but no amount of evidence can remove all doubt.

In our example, we measure the temperatures of a sample of women and of a sample of men. We cannot expect that the sample means are the same, because these are affected by the choice of sample. However, we have a good model of these sample means as random variables; in particular, we know the standard error. We can use this to assess the extent to which the difference in temperatures can be explained solely by the choice of sample. If it is easy to explain the difference in temperatures by choice of sample, then the evidence does not contradict the hypothesis; alternatively, if the difference in temperatures is hard to explain by choice of sample, the evidence contradicts the hypothesis.

There is an important, quite general, line of reasoning here. It is a bad idea to try and explain data using a hypothesis that makes the data you observed a rare event.

**Example: 9.1** *Patriot missiles*

I got this example from “Dueling idiots”, a nice book by P.J. Nahin, Princeton University Press. Apparently in 1992, the Boston Globe of Jan 24 reported on this controversy. The pentagon claimed that the patriot missile successfully engaged SCUD missiles in 80% of encounters. An MIT physicist, Theodore Postol, pointed out there was a problem. He viewed tapes of 14 patriot/SCUD encounters, with one hit and 13 misses. We can reasonably assume each encounter is independent. We can extract the probability of getting one hit and 13 misses if  $P(\text{hit}) = 0.8$  from the binomial model, to get a number around  $1e-8$ . Now you could look at this information and make several arguments: (a) the pentagon is right and the probability really is 0.8, but Postol looked at a really unlucky set of videotapes; (b) the probability is not 0.8, because you would need to fire 14 patriots at 14 SCUD missiles about  $1e8$  times to see this set of videotapes once; (c) for some reason, the videotapes are not independent — perhaps only unsuccessful encounters get filmed. If Postol viewed tapes at random (i.e. he didn’t select only unsuccessful tapes, etc.), then argument (a) is easily dismissed, because the pentagon would have had to be unreasonably unlucky — it’s a bad idea to try to explain data with a hypothesis that makes the data very unlikely.

**Example: 9.2** *MTG-DAF and Shuffling*

You build a deck of 24 lands and 36 spells. You shuffle this deck, and draw a hand of seven cards. You get no lands. You repeat the experiment, and still get no lands. On a third try, you still get no lands. By example 7.16, this event (three shuffles and draws with no lands in each hand) has probability about  $8e-6$ . You could draw several conclusions: (a) you got unlucky; (b) you miscounted lands somehow; (c) you’re shuffling the cards badly (i.e. the cards stick together, or you are bad at shuffling). Conclusion (a) is unattractive, because you would need to have been very unlucky (you’d see this outcome once in roughly  $1e5$  experiments, where each consists of shuffling and drawing three hands). It’s a bad idea to try to explain data with a hypothesis that makes the data very unlikely, so conclusions (b) and (c) have become more plausible.

## 9.3.1 Does this Population have this Mean?

Imagine we hypothesize that the average human body temperature is  $95^\circ$ . Write  $\bar{T}$  for the random variable evaluated by: collecting a random sample of people, measuring their temperatures, and computing the average of these temperatures. The mean of this random variable is the population mean, and the standard deviation of this random variable is the standard error, which we write  $s$ . Now consider the random variable

$$G = \frac{(\bar{T} - 95^\circ)}{s}.$$

This is a standard normal random variable, if our hypothesis is true.

We now have a way to tell whether the evidence supports our hypothesis. We assess how strange the sample would have to be to yield the value that we actually see, *if* the hypothesis is true. We can do this, because we know the value we observe is the value of a standard normal random variable *if* the hypothesis is true. If that value is far away from zero, for the hypothesis to be true the sample would have to be extremely unlikely.

We can evaluate our standard normal random variable by computing the mean temperature of the sample, which we write  $\bar{t}$ . This represents an observed value of  $\bar{T}$ . If you're puzzled by the relationship between the two, use an analogy with rolling a die. The random variable is what you could get if you rolled a die; the value is the face that came up when you did roll it. Similarly,  $\bar{T}$  is what you could get by choosing a sample, and  $\bar{t}$  is the value you did get with the sample you have. In turn, we obtain a value for  $G$ , which is

$$g = \frac{(\bar{t} - 95^\circ)}{s}.$$

Now  $G$  is a standard normal random variable, if our hypothesis is true. We can interpret this to mean that for 68% of possible samples, the value of  $g$  will lie between  $-1$  and  $1$  (etc.), if our hypothesis is true. Equivalently  $G$  will take a value between  $|g|$  and  $-|g|$ , for a fraction  $f$  of all possible samples of the population, where

$$f = \frac{1}{\sqrt{2\pi}} \int_{-|g|}^{|g|} \exp\left(\frac{-u^2}{2}\right) du.$$

Now assume we see a very large (or very small) value of  $g$ . The value of  $f$  will be close to one, which implies that most samples from the population will have a  $g$  value closer to zero if the hypothesis were true. Equivalently, this says that, if the hypothesis were true, our sample is highly unusual, which implies the data fails to support the hypothesis.

It is easy to think about  $p = 1 - f$  than about  $f$ . You should think of  $p$  as representing the fraction of samples that would give a larger absolute value of  $g$ , *if* the null hypothesis was true. If this fraction is very small, then there is significant evidence against the null hypothesis. We have that  $p$  — the fraction of experiments that would give us a larger absolute value of  $g$  than the one we saw — is given by

$$p = (1 - f) = \left(1 - \frac{1}{\sqrt{2\pi}} \int_{-|g|}^{|g|} \exp\left(\frac{-u^2}{2}\right) du\right) = P(\{G > |g|\}) \cup P(\{G < -|g|\})$$

and we can compute this number easily, because we know the distribution of  $G$  is normal (or nearly so). The fraction is sometimes referred to as a **p-value**.

Once we have the p-value, testing is straightforward. A small value of the fraction means that the outcome we see would be rare *if* the null hypothesis is true. The fraction that we compute should be interpreted as an assessment of the significance of the evidence against the null hypothesis. The p-value is smaller when the evidence against the null hypothesis is stronger; we get to decide how small a p-value means we should reject the null hypothesis.

It is conventional to reject the null hypothesis when the p-value is less than 0.05. This is sometimes called “a significance level of 5%”. You should interpret a p-value of 0.05 as meaning that you would see evidence this unusual in about one experiment in twenty if the null hypothesis was true. Sometimes, the p-value is even smaller, and this can be interpreted as very strong evidence the null hypothesis is wrong. A p-value of less than 0.01 allows one to reject the null hypothesis at “a significance level of 1%”. Similarly, you should interpret a p-value of 0.01 as meaning that you would see evidence this unusual in about one experiment in a hundred if the null hypothesis was true. We now have a general recipe for testing whether a population has a particular mean, which belongs in a box:

**Procedure: 9.3** *Testing whether a population has a given mean, based on a sample*

The initial hypothesis is that the population has a known mean, which we write  $\mu$ . Write  $\{x\}$  for the sample, and  $k$  for the sample size.

- Compute the sample mean, which we write  $\text{mean}(\{x\})$ .
- Estimate the standard error  $s_e$  using

$$s_e = \frac{\text{std}(x)}{\sqrt{k}}.$$

- Compute the **test statistic** using

$$s = \frac{(\mu - \text{mean}(\{x\}))}{s_e}.$$

- Compute the p-value, using

$$p = (1 - f) = \left(1 - \frac{1}{\sqrt{2\pi}} \int_{-|s|}^{|s|} \exp\left(\frac{-u^2}{2}\right) du\right)$$

Equivalently, compute the probability that a standard normal random variable  $X$  takes a value greater than  $s$  or less than  $-s$ , using the expression

$$P(\{X > |s|\}) \cup P(\{X < -|s|\}).$$

- The p-value summarizes the extent to which the data contradicts the hypothesis. A small p-value implies that, *if* the hypothesis is true, the sample is very unusual. The smaller the p-value, the more strongly the evidence contradicts the hypothesis.

It is common to think that a hypothesis can be rejected only if the p-value is less than 5% (or some number). You should not think this way; the p-value summarizes the extent to which the data contradicts the hypothesis, and your particular application logic affects how you interpret it.

**Worked example 9.7** *Average Human Body Weight*

Assess the null hypothesis that the average human body weight is 175 lb, using the height and weight data set of <http://www2.stetson.edu/~jrasp/data.htm> (called bodyfat.xls).

**Solution:** The dataset contains 252 samples; the average weight is 178.9lb. We know this average is an estimate of the mean of the original large set of all people. This estimate is a normal random variable whose mean is the mean of the population, and whose standard deviation is given by the standard error. Our sample is large (which usually means over 30 elements) and so we can estimate the standard error as

$$\frac{\text{standard deviation of sample}}{\sqrt{\text{number in sample}}} = \frac{29.4}{15.9} = 1.9,$$

where the units are lb. The test statistic is

$$\frac{\text{sample mean} - \text{hypothesized population mean}}{\text{standard error}} = \frac{178.9 - 175}{1.9} = 2.05.$$

This is the value of a standard normal random variable. We must compute the fraction of outcomes where the test statistic would be larger than 2.05, or smaller than -2.05. This is

$$2 * \frac{1}{\sqrt{2\pi}} \int_{2.05}^{\infty} \exp\left(\frac{-x^2}{2}\right) dx = 0.02$$

so the p-value is 0.02. We can interpret this as quite strong evidence that the average human body weight is not, in fact, 175lb. This p-value says that, if (a) the average human body weight is 175lb and (b) we repeat the experiment (weigh 252 people and average their weights) 50 times, we would see a number as big as the one we see about once.

**Worked example 9.8** *Average BMI*

The BMI (body mass index) is a number intended to capture how much a person's weight deviates from the usual for their height. Assess the null hypothesis that the average human BMI is 27, using the height and weight data set of <http://www2.stetson.edu/~jrasp/data.htm>; `bodyfat.xls`.

**Solution:** BMI is computed using

$$BMI = \frac{\text{weight in lb}}{(\text{height in in})^2} \times 703.$$

I did excluded two possible outliers in this dataset (entry 39 has weight 363 and height 72.25; entry 42 has weight 205 and height 29.5). I found a mean BMI of 25.3, and a standard deviation of 3.35. There are 250 items in the dataset, so the standard error is 0.21. The test statistic is

$$\frac{\text{sample mean} - \text{hypothesized population mean}}{\text{standard error}} = \frac{25.3 - 27}{2.1} = -8.1$$

so the p-value is the probability that a standard normal random variable would take a value larger than -8.1 or smaller than 8.1. This number is very small — there is no particular reason to care about the difference between this number and zero. In turn, the evidence is very strongly against the hypothesis.

If one keeps the outliers, one gets a mean BMI of 25.9, and a standard deviation of 9.56. There are 252 items in the dataset, so the standard error is 0.60. Now the p-value is 0.08, suggesting that the evidence is against the hypothesis, but not overwhelmingly. Notice the significant effect of just two datapoints. We might need more data to be sure we could reject the hypothesis.

You should notice an important relationship between our test and the material of section 9.1. When we constructed a confidence interval, we knew the distribution that the sample mean would take as we randomly chose different samples from the population. In turn, this meant we could plot (for example) the range of possible values the population mean would take for (say) 95% of possible samples. When we test a hypothesis about the mean, we are asking what kind of confidence interval we would have to draw around the hypothesized mean to encompass the observed sample mean.

### 9.3.2 Do Two Populations have the same Mean?

We have two samples, and we need to know whether they come from the same, or different, populations. For example, we might observe people using two different interfaces, measure how quickly they perform a task, then ask are their performances different? As another example, we might run an application with no other applications running, and test how long it takes to run some standard tasks. Because we don't know what the operating system, cache, etc. are up to, this number behaves a bit like a random variable, so it is worthwhile to do several experiments,

yielding one set of samples. We now do this with other applications running as well, yielding another set of samples — is this set different from the first set? For realistic datasets, the answer is always yes, because they're random samples. A better question is could the differences be the result of chance, or are these datasets really samples of two different populations?

We will ask if the population means are different. It turns out we can answer this rather easily, using some tricks about normal random variables. Assume that the samples are large enough (30 seems to be the magic number here). Write  $\{x\}$  for the first dataset, which has size  $k_x$ , and  $\{y\}$  for the second, which has size  $k_y$ . These datasets need not be of the same size. Each dataset is a random sample, drawn with replacement, from a population. We should like to assess the evidence that these two populations have the same mean. To do so, we need some simple facts about normal random variables.

**Useful Facts: 9.2** *Sums and differences of normal random variables*

Let  $X_1$  be a normal random variable with mean  $\mu_1$  and standard deviation  $\sigma_1$ . Let  $X_2$  be a normal random variable with mean  $\mu_2$  and standard deviation  $\sigma_2$ . Let  $X_1$  and  $X_2$  be independent. Then we have that:

- for any constant  $c_1 \neq 0$ ,  $c_1 X_1$  is a normal random variable with mean  $c_1 \mu_1$  and standard deviation  $c_1 \sigma_1$ ;
- for any constant  $c_2$ ,  $X_1 + c_2$  is a normal random variable with mean  $\mu_1 + c_2$  and standard deviation  $\sigma_1$ ;
- $X_1 + X_2$  is a normal random variable with mean  $\mu_1 + \mu_2$  and standard deviation  $\sqrt{\sigma_1^2 + \sigma_2^2}$ .

I will not prove these facts; we already know the expressions for means and standard deviations from our results on expectations. The only open question is to show that the distributions are normal. This is easy for the first two results. The third requires a bit of integration that isn't worth our trouble; you could do reconstruct the proof from section 13.1's notes on sums of random variables and some work with tables of integrals.

Indexing: Sums and differences of normal random variables

Indexing: Sums and differences of normal random variables

Write  $X^{(k_x)}$  for the random variable obtained by: drawing a random sample with replacement of  $k_x$  elements from the first population, then averaging this sample. Write  $Y^{(k_y)}$  for the random variable obtained by: drawing a random sample with replacement of  $k_y$  elements from the first population, then averaging this sample. From section 8.3, we know that each random variable is normal, and we know the means and standard deviations of these random variables. In turn,

this means that  $X^{(k_x)} - Y^{(k_y)}$  is a normal random variable.

Now write  $D$  for  $X^{(k_x)} - Y^{(k_y)}$ . If the two populations have the same mean, then

$$\mathbb{E}[D] = 0.$$

Furthermore,

$$\text{std}(D) = \sqrt{\text{std}(X^{(k_x)})^2 + \text{std}(X^{(k_y)})^2}.$$

We know how to estimate  $\text{std}(X^{(k_x)})$  — it is the standard error of the sample, and we can estimate it as  $\text{std}(x)/\sqrt{k_x}$ . So we can estimate  $\text{std}(D)$  as

$$\text{std}(D) \approx \sqrt{\left(\frac{\text{std}(x)}{\sqrt{k_x}}\right)^2 + \left(\frac{\text{std}(y)}{\sqrt{k_y}}\right)^2}.$$

We can now use the same reasoning we used to test the hypothesis that a population had a particular, known mean. We have identified a number we can compute from the two samples. We know how this number would vary under random choices of sample. If the value we observe is too many standard deviations away from the mean, the evidence is against our hypothesis. If we wanted to believe the hypothesis, we would be forced to believe that the sample is extremely strange. I have summarized this reasoning in box 9.4.

**Procedure: 9.4** *Testing whether two populations have a given mean, based on a sample of each*

The initial hypothesis is that the populations have the same, unknown, mean. Write  $\{x\}$  for the sample of the first population,  $\{y\}$  for the sample of the second population, and  $k_x, k_y$  for the sample sizes.

- Compute the sample means for each population,  $\text{mean}(\{x\})$  and  $\text{mean}(\{y\})$ .
- Estimate the standard error  $s_e$  for each population, using

$$s_{ex} = \frac{\text{std}(x)}{\sqrt{k}}, \quad s_{ey} = \frac{\text{std}(y)}{\sqrt{k}}.$$

- Compute the standard error for the difference between the means,

$$s_{ed} = \sqrt{s_{ex}^2 + s_{ey}^2}.$$

- Compute the **test statistic** using

$$s = \frac{(\text{mean}(\{x\}) - \text{mean}(\{y\}))}{s_{ed}}.$$

- Compute the p-value, using

$$p = (1 - f) = (1 - \int_{-|s|}^{|s|} \exp\left(\frac{-u^2}{2}\right) du)$$

Equivalently, compute the probability that a standard normal random variable  $X$  takes a value greater than  $|s|$  or less than  $-|s|$ , using the expression

$$P(\{X > |s|\}) \cup P(\{X < -|s|\}).$$

- The p-value summarizes the extent to which the data contradicts the hypothesis. A small p-value implies that, *if* the hypothesis is true, the sample is very unusual. The smaller the p-value, the more strongly the evidence contradicts the hypothesis.

It is common to think that a hypothesis can be rejected only if the p-value is less than 5% (or some number). You should not think this way; the p-value summarizes the extent to which the data contradicts the hypothesis, and your particular application logic affects how you interpret it.

**Worked example 9.9** *Are US and Japanese cars different*

At <http://www.itl.nist.gov/div898/handbook/eda/section3/eda3531.htm> you can find a dataset, published by NIST, giving measurements of miles per gallon for Japanese and US cars. Assess the evidence these two populations have the same mean MPG.

**Solution:** There are 249 measurements for Japanese cars, and 79 for US cars. The mean for Japanese cars is 20.1446, and for US cars is 30.4810. The standard error for Japanese cars is 0.4065, and for US cars is 0.6872. The value of the test statistic is

$$\frac{\text{mean}(\{x\}) - \text{mean}(\{y\})}{s_{ed}} = 10.33/0.7984 = 12.94$$

and the p-value is the probability of encountering a standard normal random variable of this value or greater. This is so close to zero I had trouble getting sensible numbers out of MATLAB; the evidence very strongly rejects this hypothesis. A version of this example, using the more complex two-sample t-test, is worked in the NIST/SEMATECH e-Handbook of Statistical Methods, at <http://www.itl.nist.gov/div898/handbook/>, as of 2013.

**Worked example 9.10** *Do women and men have the same body temperature?*

Using the dataset of human temperatures at , assess the evidence against the hypothesis that gender 1 has the same mean temperature as gender 2

**Solution:** We compute the mean temperatures and standard errors shown

Gender:	1	2
Mean:	98.10	98.39
Std Error:	0.0867	0.0922

The null hypothesis is that these

two are the same, so the test statistic is

$$\frac{\text{difference}}{\text{std error}} = \frac{98.39 - 98.10}{\sqrt{0.0867^2 + 0.0922^2}} = 2.285$$

and we must ask what is the probability of getting a number with absolute value this big, or bigger, from a normal distribution (two-sided test). This is 0.0223. The evidence in this data set does not support the null hypothesis.

## 9.3.3 Variants on the Basic Test

The basic recipe we described is: find a statistic that summarizes the relationship between the sample and the hypothesis; figure out the probability distribution of that statistic under random variation caused by sampling; and then determine how rare the sample would have to be to observe the value of the statistic that was actually observed. There are a wide range of different instances of the recipe. I will sketch some variants here, but be aware that this subject is very big indeed.

**One-sided Tests**

The test of procedure 9.3 is usually referred to as a **two-sided** test. This is because it computes

$$P(\{X > |s|\}) \cup P(\{X < -|s|\}).$$

This is the fraction of samples that would produce a value that is larger than  $|s|$  or smaller than  $-|s|$ . In some cases, we can expect that only the larger (or smaller) value is of interest. This yields a **one-sided** test; the relevant expression is either

$$P(\{X > s\})$$

or

$$P(\{X < s\}).$$

Generally, it is more conservative to use a two-sided test, and one should do so unless there is a good reason not to. Very often, authors use one-sided tests because they result in smaller p-values, and small p-values are often a requirement for publication.

**Correcting the Standard Deviation**

When the sample is small, the sample standard deviation is a poor estimate of the population standard deviation. The value of the sample mean that we compute minimizes the sample standard deviation, which means that the estimate tends to be a little too small. In turn, the standard error is a little too small, and there is slightly more probability that the sample mean is far from the population mean than the normal model allows for. This can be corrected for. Instead of using the standard deviation of the sample to form the standard error, we use

$$\sqrt{\frac{\sum_i (x_i - \text{mean}(\{x_i\}))^2}{k - 1}}.$$

**Z-Tests and T-Tests**

When we have a large sample, the sample mean is a normal random variable whose mean is the population mean and whose standard deviation is the standard error. This means that we know its distribution. In this case, the test is known as a **z-test**. An important difficulty with this procedure is we assumed we know the standard error. As figure 9.2 suggested, estimating standard error with  $\text{std}(x)/\sqrt{k}$  works fine for large  $k$ , but is less successful for small  $k$ . Usually, practical advice suggests that one should use a Z-test only if the sample has at least 30 elements.

If the sample is smaller, the distribution of the sample mean is not really normal. It turns out to have a form known as **Student's t-distribution**. This is a

family of probability distributions. There are two parameters; the observed value  $s$ , and the number of degrees of freedom. The number of degrees of freedom is  $k - 1$  for our purposes. When the number of degrees of freedom is small, the t-distribution has rather heavier tails than the normal distribution, so the test takes into account that the standard error may be larger than we think (because the population standard deviation is larger than we expected). When the number of degrees of freedom is large, the t-distribution is very similar to the normal distribution. One can get p-values from tables, or by the Matlab function `ttest`.

#### 9.4 $\chi^2$ TESTS: IS DATA CONSISTENT WITH A MODEL?

Sometimes we have a model, and we would like to know whether the data is consistent with that model. For example, imagine we have a six-sided die. We throw it many times, and record which number comes up each time. We would like to know if the die is fair (i.e. is the data consistent with the model that the die is fair). It is highly unlikely that each face comes up the same number of times, even if the die is fair. Instead, there will be some variation in the frequencies observed; with what probability is that variation, or bigger, the result of chance effects?

As another example, we decide that the number of calls by a telemarketer in each hour is distributed with a Poisson distribution. We don't know the intensity. We could collect call data, and use maximum likelihood to determine the intensity. Once we have the best estimate of the intensity, we still want to know whether the model is consistent with the data.

In each case, the model predicts the frequencies of events. For the six-sided die case, the model tells us how often we expect to see each side. For the call case, the model predicts how often we would see no calls, one call, two calls, three calls, etc. in each hour. To tell whether the model fits the data, we need to compare the frequencies we observed with theoretical frequencies.

The appropriate test is a  $\chi^2$  (say "khi-squared") test. Assume we have a set of disjoint events  $\mathcal{E}_1, \dots, \mathcal{E}_k$  which cover the space of outcomes (i.e. any outcome lies in one of these events). Assume we perform  $k$  experiments, and record the number of times each event occurs. We have a null hypothesis regarding the probability of events. We can take the probability of each event and multiply by  $k$  to get a frequency under the null hypothesis. Now write  $f_o(\mathcal{E}_i)$  for the observed frequency of event  $i$ ;  $f_t(\mathcal{E}_i)$  for the theoretical frequency of the event under the null hypothesis. We form the statistic

$$\sum_i \frac{(f_o(\mathcal{E}_i) - f_t(\mathcal{E}_i))^2}{f_t(\mathcal{E}_i)}$$

which compares the observed and actual frequency of events. It turns out that this statistic has a distribution very close to a known form, called the  $\chi^2$  distribution, as long as each count is 5 or more. The distribution has two parameters; the statistic, and the number of degrees of freedom. The number of degrees of freedom to use for a straightforward test is  $k - 1$ ; if one has to estimate a total of  $p$  parameters for the null hypothesis, this number is  $k - p - 1$ .

After this, things follow the usual recipe. We compute the statistic; we then look at tables, or use the matlab function `chi2cdf`, to find the probability that the statistic takes this value or greater under the null hypothesis. If this is small, then

we reject the null hypothesis.

**Worked example 9.11**  $\chi^2$  test for dice

I throw a die 100 times. I record the outcomes, in the table below. Is this a fair die?

face	count
1	46
2	13
3	12
4	11
5	9
6	9

**Solution:** The expected frequency is  $100/6$  for each face. The  $\chi^2$  statistic has the value 62.7, and there are 5 degrees of freedom. We get the significance as `1-chi2cdf(62.7, 5)`, which is (basically)  $3e-12$ . You would have to run this experiment  $3e11$  times to see a table as skewed as this once, by chance. The die is not fair. Notice the most important bit of this example, which is how to get the number out of matlab.

**Worked example 9.12** *Is gender 2 temperature normal?*

Recall we used body temperature data for two genders in earlier examples. The gender 2 data looked as though it might not be normal. Use a  $\chi^2$  test to tell whether it is normal with mean  $98.4^\circ$  and standard deviation 0.743 or not.

**Solution:** The null hypothesis is that the data is normal with mean  $98.4^\circ$  and standard deviation 0.743. We break up the range into five buckets (less than  $97.65 = 98.4 - 0.74$ ; between 97.65 and  $98 = 98.4 - 0.743/2$ ; between 98 and  $98.76 = 98.4 + 0.743/2$ ; and greater than 99.14). With a little work with error functions, we get that the theoretical frequency in each bucket under the null hypothesis is (10.3126; 9.7423; 24.8901; 9.7423; 10.3126). The actual frequencies are (7; 13; 26; 12; 7). The  $\chi^2$  statistic is about  $4e3$ , and the significance level is essentially zero. This data isn't normal with the parameters given (though it might be normal with a different mean and a different standard deviation). Looking at the frequencies suggests the problem; there are far too few temperatures far away from the mean for this to be normal, and far too many in the center bucket.

**Worked example 9.13** *Is swearing Poisson?*

A famously swearsy politician gives a talk. You listen to the talk, and for each of 30 intervals 1 minute long, you record the number of swearwords. You record this as a histogram (i.e. you count the number of intervals with zero swear words, with one, etc.), obtaining the table below.

no. of swear words	no. of intervals
0	13
1	9
2	8
3	5
4	5

The null hypothesis is that the politician's swearing is Poisson distributed, with intensity ( $\lambda$ ) one. Can you reject this null hypothesis?

**Solution:** If the null hypothesis is true, then the probability of getting  $n$  swear words in a fixed length interval would be  $\frac{\lambda^n e^{-\lambda}}{n!}$ . There are 10 intervals, so the theoretical frequencies are 10 times the following probabilities

number of swear words	probability
0	0.368
1	0.368
2	0.184
3	0.061
4	0.015

so the  $\chi^2$  statistic takes the value 243.1 and there are 4 degrees of freedom. The significance `1-chi2cdf(243.1, 4)` is indistinguishable from zero by matlab, so you can firmly reject the null hypothesis. Of course, it may just be that the intensity is wrong; but we don't know how to deal with that.

## 9.5 YOU SHOULD

## 9.5.1 remember

New term: confidence intervals . . . . .	255
New term: statistic . . . . .	255
Useful facts: Normal random variables . . . . .	257
New term: confidence interval . . . . .	257
New term: error bars . . . . .	257
New term: inverse error function . . . . .	259
New term: bootstrap . . . . .	265
New term: bootstrap replicates . . . . .	265
New term: empirical distribution . . . . .	265
New term: hypothesis . . . . .	269
New term: p-value . . . . .	272
New term: test statistic . . . . .	273
Useful facts: Sums and differences of normal random variables . . . . .	276
New term: test statistic . . . . .	278
New term: two-sided . . . . .	280
New term: one-sided . . . . .	280
New term: z-test . . . . .	280
New term: Student's t-distribution . . . . .	280

## PROBLEMS

## Confidence Intervals

- 9.1. The Weight of Mice** You wish to estimate the average weight of a mouse. You obtain 10 mice, sampled uniformly at random and with replacement from the mouse population. Their weights are 21, 23, 27, 19, 17, 18, 20, 15, 17, 22 grams respectively.
- (a) Give a 68% confidence interval for the weight of a mouse, from this data.  
 (b) Give a 99% confidence interval for the weight of a mouse, from this data.
- 9.2. The Weight of Rats** You wish to estimate the average weight of a pet rat. You obtain 40 rats (easily and cheaply done; keep them, because they're excellent pets), sampled uniformly at random and with replacement from the pet rat population. The mean weight is 340 grams, with a standard deviation of 75 grams.
- (a) Give a 68% confidence interval for the weight of a pet rat, from this data.  
 (b) Give a 99% confidence interval for the weight of a pet rat, from this data.

## Hypothesis Testing

- 9.3. Yet more Mouse-weighing** I claim the average weight of a mouse is 25 grams. You decide to evaluate the evidence in support of this claim. You obtain 10 mice, sampled uniformly at random and with replacement from the mouse population. Their weights are 21, 23, 27, 19, 17, 18, 20, 15, 17, 22 grams respectively. Does the evidence support my claim? to what extent? Why?
- 9.4. How big are Parktown Prawns?** The Parktown prawn is an impressively repellent large insect, common in Johannesburg (look them up on the Web). I claim that their average length is 10cm. You collect 100 Parktown prawns (this will take about 10 mins, in the right places in Johannesburg; more difficult from here). The mean length of these prawns is 7cm. The standard deviation is 1cm. Assess the evidence against my claim.
- 9.5. Two Populations of Rats** Zucker rats are specially bred to have curious weight properties, related to their genetics (look them up on the Web). You measure 30 lean Zucker rats, obtaining an average weight of 500 grams with a standard deviation of 50 grams. You measure 20 fatty Zucker rats, obtaining an average weight of 1000 grams with a standard deviation of 100 grams. Assess the evidence against the claim that these populations have the same weight.
- 9.6. Male and Female pet Rats** You measure 35 female pet rats, obtaining an average weight of 300 grams with a standard deviation of 30 grams. You measure 30 male pet rats, obtaining an average weight of 400 grams with a standard deviation of 100 grams. Assess the evidence against the claim that these populations have the same weight.
- 9.7. Lean and Fatty Zucker Rats** Zucker rats are specially bred to have curious weight properties, related to their genetics (look them up on the Web). You measure 30 lean Zucker rats, obtaining an average weight of 500 grams with a standard deviation of 50 grams. You measure 35 fatty Zucker rats, obtaining an average weight of 1000 grams with a standard deviation of 100 grams. In steps, you will assess the evidence against the claim that a fatty Zucker rat has exactly twice the weight of a lean Zucker rat. You know that the product of a normal random variable and a constant is a normal random variable. You should assume (and accept, because I won't prove it) that the sum of two

normal random variables is a normal random variable.

(a) Write  $L^{(k)}$  for the random variable obtained by drawing a uniform sample of  $k$  lean rats and averaging their weights. You can assume that  $k$  is large enough that this is normal.

- What is  $\mathbb{E}\left[L^{(k)}\right]$ ? (write an expression, no need to prove anything)
- What is  $\text{std}\left(L^{(k)}\right)$ ? (write an expression, no need to prove anything)

(b) Now write  $F^{(s)}$  for the random variable obtained by drawing a uniform sample of  $s$  fatty rats and averaging their weights. You can assume that  $s$  is large enough that this is normal.

- What is  $\mathbb{E}\left[F^{(s)}\right]$ ? (write an expression, no need to prove anything)
- What is  $\text{std}\left(F^{(s)}\right)$ ? (write an expression, no need to prove anything)

(c) Write  $\text{popmean}(\{L\})$  for the population mean weight of lean rats, and  $\text{popmean}(\{F\})$  for the population mean weight of fatty rats. Assume that  $2\text{popmean}(\{L\}) = \text{popmean}(\{F\})$ .

- In this case, what is  $\mathbb{E}\left[F^{(s)} - 2L^{(k)}\right]$ ?
- In this case, what is  $\text{std}\left(F^{(s)} - 2L^{(k)}\right)$ ?
- Your expression for  $\text{std}\left(F^{(s)} - 2L^{(k)}\right)$  will have contained terms in the population standard deviation of  $F$  and  $L$ . What is the standard error of  $F^{(s)} - 2L^{(k)}$ ?

(d) Now assess the evidence against the hypothesis that a fatty Zucker rat weighs exactly twice as much as a lean Zucker rat.

**9.8. Are boys and girls equiprobable?** In Carcelle-le-Grignon at the end of the 18<sup>th</sup> century, there were 2,009 births. There were 983 boys and 1026 girls. You can regard this as a fair random sample (with replacement, though try not to think too hard about what that means) of births. Assess the evidence against the hypothesis that a boy is born with probability exactly 0.5.

**9.9.** In 1998, the average height of an adult male in South Africa was estimated to be 169cm. Assume that this estimate is exact; assume also that the population standard deviation is 10cm. What fraction of samples consisting of 50 adult males from South Africa (selected uniformly at random, and with replacement) will have average height greater than 200cm?

**9.10.** Assume the average weight of an adult male short-hair house cat is 5 kg, and the standard deviation is 0.7 kg (these numbers are reasonable, but there's quite a lively fight between cat fanciers about the true numbers).

- (a) What fraction of samples consisting of 30 adult male short-hair house cats (selected uniformly at random, and with replacement) will have average weight less than 4kg?
- (b) What fraction of samples consisting of 300 adult male short-hair house cats (selected uniformly at random, and with replacement) will have average weight less than 4kg?
- (c) Why are these numbers different?

# Extracting Important Relationships in High Dimensions

Chapter 3 described methods to explore the relationship between two elements in a dataset. We could extract a pair of elements and construct various plots. For vector data, we could also compute the correlation between different pairs of elements. But if each data item is  $d$ -dimensional, there could be a lot of pairs to deal with.

We will think of our dataset as a collection of  $d$  dimensional vectors. It turns out that there are easy generalizations of our summaries. However, it is hard to plot  $d$ -dimensional vectors. We need to find some way to make them fit on a 2-dimensional plot. Some simple methods can offer insights, but to really get what is going on we need methods that can at all pairs of relationships in a dataset in one go.

These methods visualize the dataset as a “blob” in a  $d$ -dimensional space. Many such blobs are flattened in some directions, because components of the data are strongly correlated. Finding the directions in which the blobs are flat yields methods to compute lower dimensional representations of the dataset.

## 10.1 SUMMARIES AND SIMPLE PLOTS

In this chapter, we assume that our data items are vectors. This means that we can add and subtract values and multiply values by a scalar without any distress. This is an important assumption, but it doesn’t necessarily mean that data is continuous (for example, you can meaningfully add the number of children in one family to the number of children in another family). It does rule out a lot of discrete data. For example, you can’t add “sports” to “grades” and expect a sensible answer.

**Notation:** Our data items are vectors, and we write a vector as  $\mathbf{x}$ . The data items are  $d$ -dimensional, and there are  $N$  of them. The entire data set is  $\{\mathbf{x}\}$ . When we need to refer to the  $i$ ’th data item, we write  $\mathbf{x}_i$ . We write  $\{\mathbf{x}_i\}$  for a new dataset made up of  $N$  items, where the  $i$ ’th item is  $\mathbf{x}_i$ . If we need to refer to the  $j$ ’th component of a vector  $\mathbf{x}_i$ , we will write  $x_i^{(j)}$  (notice this isn’t in bold, because it is a component not a vector, and the  $j$  is in parentheses because it isn’t a power). Vectors are always column vectors.

### 10.1.1 The Mean

For one-dimensional data, we wrote

$$\text{mean}(\{x\}) = \frac{\sum_i x_i}{N}.$$

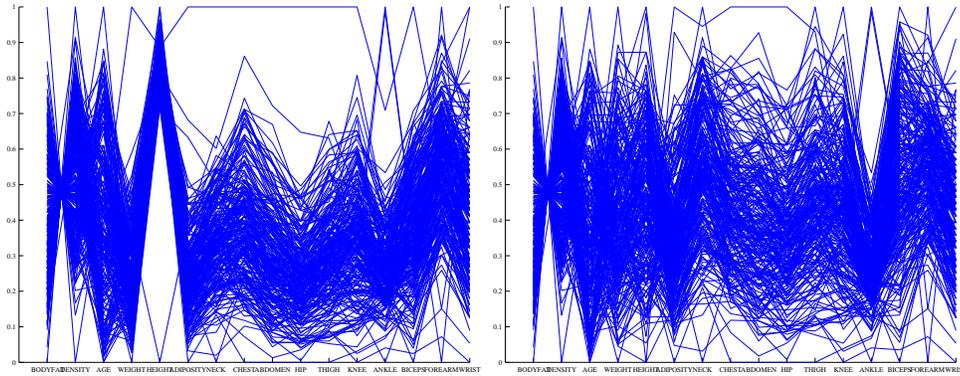


FIGURE 10.1: **Left:** A parallel plot of the bodyfat dataset, including all data points. I have named the components on the horizontal axis. It is easy to see that large values of bodyfat correspond to small values of density, and vice versa. Notice that one datapoint has height very different from all others; similarly, one datapoint has weight very different from all others. **Right:** A plot with those data items removed, so that those components are renormalized. Two datapoints have rather distinct ankle measurements. Generally, you can see that large knees go with large ankles and large biceps (the  $v$  structure).

This expression is meaningful for vectors, too, because we can add vectors and divide by scalars. We write

$$\text{mean}(\{\mathbf{x}\}) = \frac{\sum_i \mathbf{x}_i}{N}$$

and call this the mean of the data. Notice that each component of  $\text{mean}(\{\mathbf{x}\})$  is the mean of that component of the data. There is not an easy analogue of the median, however (how do you order high dimensional data?) and this is a nuisance. Notice that, just as for the one-dimensional mean, we have

$$\text{mean}(\{\mathbf{x} - \text{mean}(\{\mathbf{x}\})\}) = 0$$

(i.e. if you subtract the mean from a data set, the resulting data set has zero mean).

### 10.1.2 Parallel Plots and Stem Plots

Parallel plots can sometimes reveal information, particularly when the dimension of the dataset is low. To construct a parallel plot, you compute a normalized representation of each component of each data item. The component is normalized by translating and scaling so that the minimum value over the dataset is zero, and the maximum value over the dataset is one. Now write the  $i$ 'th normalised data item as  $(n_1, n_2, \dots, n_d)$ . For this data item, you plot a broken line joining  $(1, n_1)$  to  $(2, n_2)$  to  $(3, n_3)$ , etc. These plots are superimposed on one another. In the case of the bodyfat dataset, this yields the plot of figure 10.1.

Some structures in the parallel plot are revealing. Outliers often stick out (in figure 10.1, it's pretty clear that there's a data point with a very low height value,

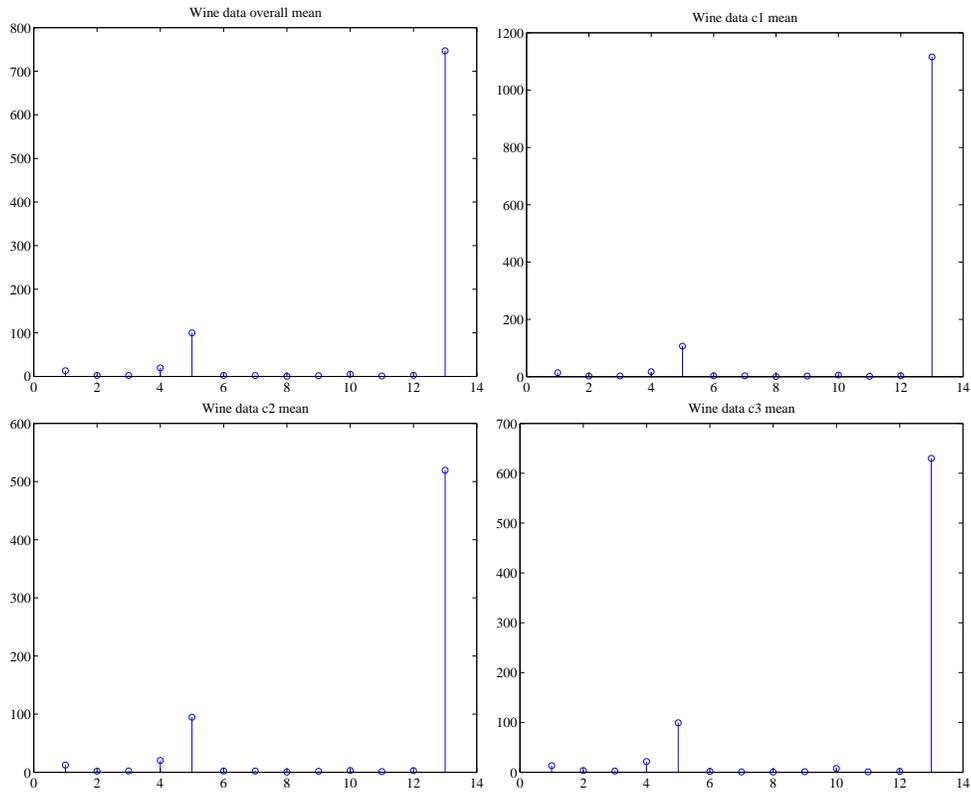


FIGURE 10.2: A stem plot of various means of the wine dataset, from <http://archive.ics.uci.edu/ml/datasets/Wine>. **Top left:** The overall mean of the dataset. **Top right:** The mean of the class 1 data items. **Bottom left:** The mean of the class 2 data items. **Bottom right:** The mean of the class 3 data items. You should notice that each class has a slightly different mean (look at the 13'th component).

and also one with a very large weight value). Outliers affect the scaling, and so make other structures difficult to spot. I have removed them for figure 10.1. In this figure, you can see that two negatively correlated components next to one another produce a butterfly like shape (bodyfat and density). In this plot, you can also see that there are still several data points that are very different from others (two data items have ankle values that are very different from the others, for example).

Another simple but useful plotting mechanism is the stem plot. This is can be a useful way to plot a high dimensional data point. One plots each component of the vector as a vertical line, typically with a circle on the end (easier seen that said; look at figure 10.2). The dataset I used for this is the wine dataset, from the UC Irvine machine learning data repository. You can find this dataset at <http://archive.ics.uci.edu/ml/datasets/Wine>. For each of three types of wine, the data records the values of 13 different attributes. In the figure, I show the overall mean of the dataset, and also the mean of each type of wine (also known as the class means, or

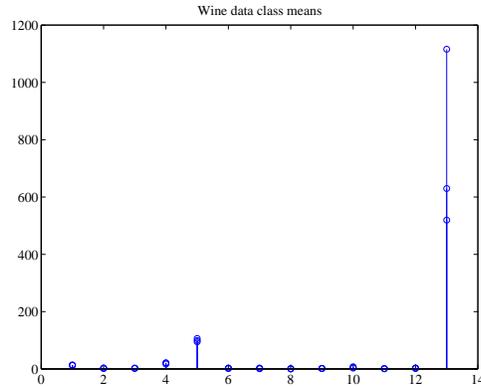


FIGURE 10.3: *In this figure, I have overlaid stem plots of each class mean from the wine dataset, from <http://archive.ics.uci.edu/ml/datasets/Wine>, so that you can see the differences between class means.*

class conditional means). A natural way to compare class means is to plot them on top of one another in a stem plot (figure 10.3).

### 10.1.3 Blobs

When we plotted histograms, we saw that mean and variance were a very helpful description of data that had a unimodal histogram. If the histogram had more than one mode, one needed to be somewhat careful to interpret the mean and variance; in the pizza example, we plotted diameters for different manufacturers to try and see the data as a collection of unimodal histograms. In higher dimensions, the analogue of a unimodal histogram is a “blob” — a group of data points that clusters nicely together and should be understood together.

Generally, mean and covariance are a good description of data that lies in a “blob” (Figure 10.4). You might not believe that “blob” is a technical term, but it’s quite widely used. This is because it is relatively easy to understand a single blob of data. There are good summary representations (mean and covariance, which I describe below). If a dataset forms multiple blobs, we can usually coerce it into a representation as a collection of blobs (using the methods of chapter 12). But many datasets really are single blobs, and we concentrate on such data here. There are quite useful tricks for understanding blobs of low dimension by plotting them, which I describe below. To understand a high dimensional blob, we will need to think about the coordinate transformations that places it into a particularly convenient form; I’ll cover this in the next section.

### 10.1.4 Understanding Blobs with Scatterplot Matrices

Plotting high dimensional data is tricky. If there are relatively few dimensions, you could just choose three of them and produce a scatterplot. Figure 10.5 shows such a scatterplot, for data that was originally four dimensional. This dataset is a famous dataset to do with the botanical classification of irises. I found a copy at the UC

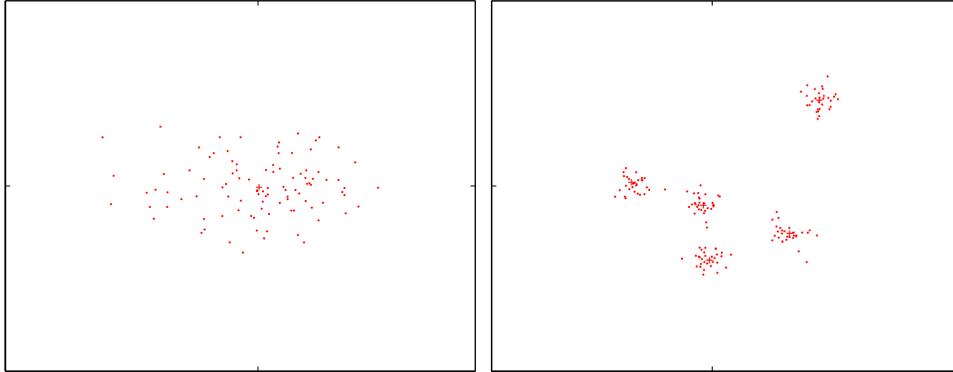


FIGURE 10.4: *On the left*, a “blob” in two dimensions. This is a set of data points that lie somewhat clustered around a single center, given by the mean. I have plotted the mean of these data points with a ‘+’. *On the right*, a data set that is best thought of as a collection of five blobs. I have plotted the mean of each with a ‘+’. We could compute the mean and covariance of this data, but it would be less revealing than the mean and covariance of a single blob. In chapter 13.1, I will describe automatic methods to describe this dataset as a series of blobs.

Irvine repository of datasets that are important in machine learning. You can find the repository at <http://archive.ics.uci.edu/ml/index.html>.

Another strategy that is very useful when there aren’t too many dimensions is to use a scatterplot matrix. To build one, you lay out scatterplots for each pair of variables in a matrix. On the diagonal, you name the variable that is the vertical axis for each plot in the row, and the horizontal axis in the column. This sounds more complicated than it is; look at the example of figure 10.5, which shows a scatterplot matrix for the iris data of figure 10.5.

Figure 10.6 shows a scatter plot matrix for four of the variables in the height weight dataset of <http://www2.stetson.edu/~jrasp/data.htm>; look for bodyfat.xls at that URL). This is originally a 16-dimensional dataset, but a 16 by 16 scatterplot matrix is squashed and hard to interpret. What is nice about this kind of plot is that it’s quite easy to spot correlations between pairs of variables, though you do need to take into account the coordinates have not been normalized. For figure 10.6, you can see that weight and adiposity appear to show quite strong correlations, but weight and age are pretty weakly correlated. Height and age seem to have a low correlation. It is also easy to visualize unusual data points. Usually one has an interactive process to do so — you can move a “brush” over the plot to change the color of data points under the brush. To show what might happen, figure 10.7 shows a scatter plot matrix with some points shown as circles. Notice how they lie inside the “blob” of data in some views, and outside in others. This is an effect of projection.

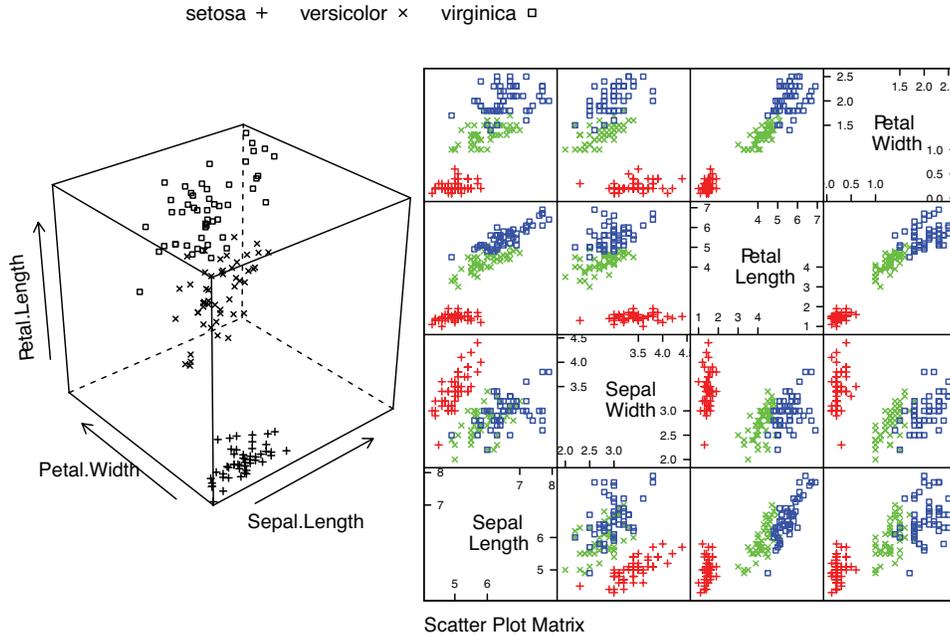


FIGURE 10.5: **Left:** a 3D scatterplot for the famous Iris data, originally due to *\*\*\**. I have chosen three variables from the four, and have plotted each species with a different marker. You can see from the plot that the species cluster quite tightly, and are different from one another. **Right:** a scatterplot matrix for the famous Iris data, originally due to *\*\*\**. There are four variables, measured for each of three species of iris. I have plotted each species with a different marker. You can see from the plot that the species cluster quite tightly, and are different from one another.

### 10.1.5 Using Covariance to encode Variance and Correlation

Variance, standard deviation and correlation can each be seen as an instance of a more general operation on data. Assume that we have two one dimensional data sets  $\{x\}$  and  $\{y\}$ . Then we can define the covariance of  $\{x\}$  and  $\{y\}$ .

**Definition: 10.1** *Covariance*

Assume we have two sets of  $N$  data items,  $\{x\}$  and  $\{y\}$ . We compute the covariance by

$$\text{cov}(\{x\}, \{y\}) = \frac{\sum_i (x_i - \text{mean}(\{x\}))(y_i - \text{mean}(\{y\}))}{N}$$

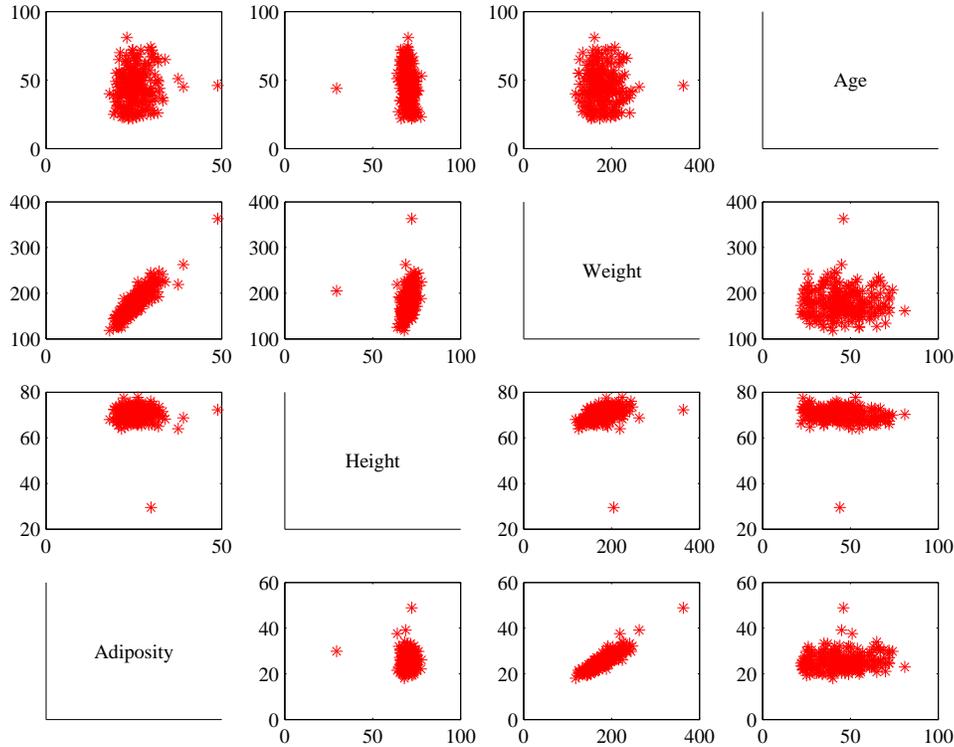


FIGURE 10.6: This is a scatterplot matrix for four of the variables in the height weight dataset of <http://www2.stetson.edu/~jrasp/data.htm>. Each plot is a scatterplot of a pair of variables. The name of the variable for the horizontal axis is obtained by running your eye down the column; for the vertical axis, along the row. Although this plot is redundant (half of the plots are just flipped versions of the other half), that redundancy makes it easier to follow points by eye. You can look at a column, move down to a row, move across to a column, etc. Notice how you can spot correlations between variables and outliers (the arrows).

#### Indexing: Covariance

Indexing: Covariance Covariance measures the tendency of corresponding elements of  $\{x\}$  and of  $\{y\}$  to be larger than (resp. smaller than) the mean. Just like mean, standard deviation and variance, covariance can refer either to a property of a dataset (as in the definition here) or a particular expectation (as in chapter 5). The correspondence is defined by the order of elements in the data set, so that  $x_1$  corresponds to  $y_1$ ,  $x_2$  corresponds to  $y_2$ , and so on. If  $\{x\}$  tends to be larger (resp. smaller) than its mean for data points where  $\{y\}$  is also larger (resp. smaller) than its mean, then the covariance should be positive. If  $\{x\}$  tends to be larger (resp. smaller) than its mean for data points where  $\{y\}$  is smaller (resp. larger) than its mean, then the covariance should be negative.

From this description, it should be clear we have seen examples of covariance

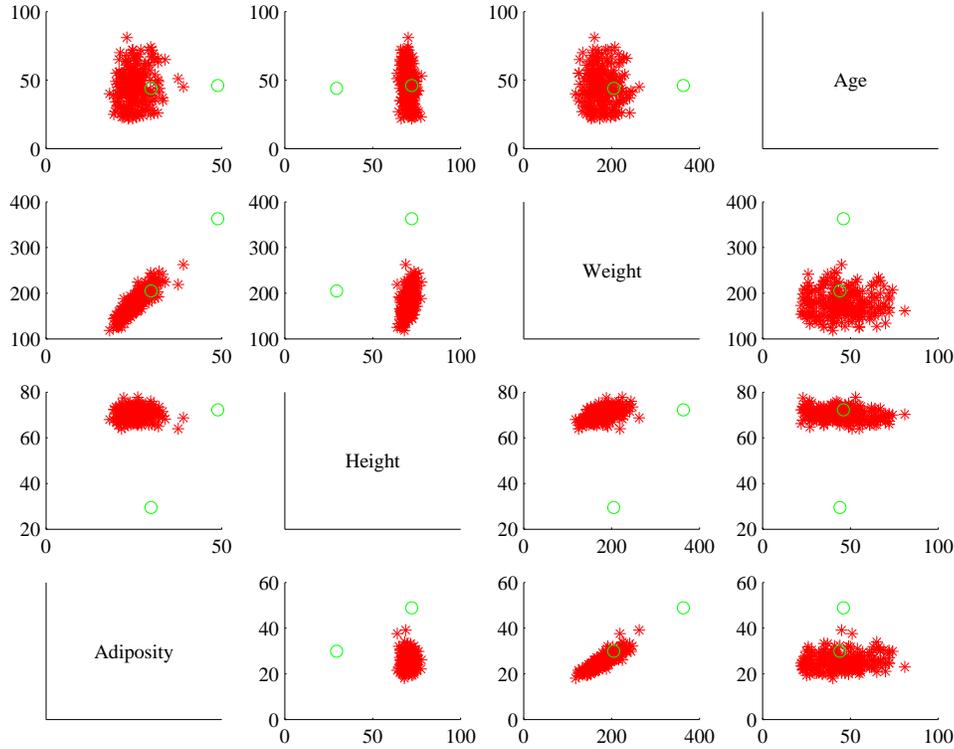


FIGURE 10.7: You should compare this figure with figure 10.6. I have marked two data points with circles in this figure; notice that in some panels these are far from the rest of the data, in others close by. A “brush” in an interactive application can be used to mark data like this to allow a user to explore a dataset.

already. Notice that

$$\text{std}(x)^2 = \text{var}(\{x\}) = \text{cov}(\{x\}, \{x\})$$

which you can prove by substituting the expressions. Recall that variance measures the tendency of a dataset to be different from the mean, so the covariance of a dataset with itself is a measure of its tendency not to be constant.

More important, notice that

$$\text{corr}(\{(x, y)\}) = \frac{\text{cov}(\{x\}, \{y\})}{\sqrt{\text{cov}(\{x\}, \{x\})}\sqrt{\text{cov}(\{y\}, \{y\})}}.$$

This is occasionally a useful way to think about correlation. It says that the correlation measures the tendency of  $\{x\}$  and  $\{y\}$  to be larger (resp. smaller) than their means for the same data points, compared to how much they change on their own.

Working with covariance (rather than correlation) allows us to unify some ideas. In particular, for data items which are  $d$  dimensional vectors, it is straight-

forward to compute a single matrix that captures all covariances between all pairs of components — this is the covariance matrix.

**Definition: 10.2** *Covariance Matrix*

The covariance matrix is:

$$\text{Covmat}(\{\mathbf{x}\}) = \frac{\sum_i (\mathbf{x}_i - \text{mean}(\{\mathbf{x}\}))(\mathbf{x}_i - \text{mean}(\{\mathbf{x}\}))^T}{N}$$

Notice that it is quite usual to write a covariance matrix as  $\Sigma$ , and we will follow this convention.

Indexing: Covariance Matrix

Indexing: Covariance Matrix

**Properties of the Covariance Matrix** Covariance matrices are often written as  $\Sigma$ , whatever the dataset (you get to figure out precisely which dataset is intended, from context). Generally, when we want to refer to the  $j$ ,  $k$ 'th entry of a matrix  $\mathcal{A}$ , we will write  $\mathcal{A}_{jk}$ , so  $\Sigma_{jk}$  is the covariance between the  $j$ 'th and  $k$ 'th components of the data.

- The  $j$ ,  $k$ 'th entry of the covariance matrix is the covariance of the  $j$ 'th and the  $k$ 'th components of  $\mathbf{x}$ , which we write  $\text{cov}(\{x^{(j)}\}, \{x^{(k)}\})$ .
- The  $j$ ,  $j$ 'th entry of the covariance matrix is the variance of the  $j$ 'th component of  $\mathbf{x}$ .
- The covariance matrix is symmetric.
- The covariance matrix is always positive semi-definite; it is positive definite, *unless* there is some vector  $\mathbf{a}$  such that  $\mathbf{a}^T(\mathbf{x}_i - \text{mean}(\{\mathbf{x}_i\})) = 0$  for all  $i$ .

**Proposition:**

$$\text{Covmat}(\{\mathbf{x}\})_{jk} = \text{cov}\left(\{x^{(j)}\}, \{x^{(k)}\}\right)$$

**Proof:** Recall

$$\text{Covmat}(\{\mathbf{x}\}) = \frac{\sum_i (\mathbf{x}_i - \text{mean}(\{\mathbf{x}\}))(\mathbf{x}_i - \text{mean}(\{\mathbf{x}\}))^T}{N}$$

and the  $j, k$ 'th entry in this matrix will be

$$\frac{\sum_i (x_i^{(j)} - \text{mean}(\{x^{(j)}\}))(x_i^{(k)} - \text{mean}(\{x^{(k)}\}))}{N}$$

which is  $\text{cov}(\{x^{(j)}\}, \{x^{(k)}\})$ .

**Proposition:**

$$\text{Covmat}(\{\mathbf{x}_i\})_{jj} = \Sigma_{jj} = \text{var}\left(\{x^{(j)}\}\right)$$

**Proof:**

$$\begin{aligned} \text{Covmat}(\{\mathbf{x}\})_{jj} &= \text{cov}\left(\{x^{(j)}\}, \{x^{(j)}\}\right) \\ &= \text{var}\left(\{x^{(j)}\}\right) \end{aligned}$$

**Proposition:**

$$\text{Covmat}(\{\mathbf{x}\}) = \text{Covmat}(\{\mathbf{x}\})^T$$

**Proof:** We have

$$\begin{aligned} \text{Covmat}(\{\mathbf{x}\})_{jk} &= \text{cov}\left(\left\{x^{(j)}\right\}, \left\{x^{(k)}\right\}\right) \\ &= \text{cov}\left(\left\{x^{(k)}\right\}, \left\{x^{(j)}\right\}\right) \\ &= \text{Covmat}(\{\mathbf{x}\})_{kj} \end{aligned}$$

**Proposition:** Write  $\Sigma = \text{Covmat}(\{\mathbf{x}\})$ . If there is no vector  $\mathbf{a}$  such that  $\mathbf{a}^T(\mathbf{x}_i - \text{mean}(\{\mathbf{x}\})) = 0$  for all  $i$ , then for any vector  $\mathbf{u}$ , such that  $\|\mathbf{u}\| > 0$ ,

$$\mathbf{u}^T \Sigma \mathbf{u} > 0.$$

If there is such a vector  $\mathbf{a}$ , then

$$\mathbf{u}^T \Sigma \mathbf{u} \geq 0.$$

**Proof:** We have

$$\begin{aligned} \mathbf{u}^T \Sigma \mathbf{u} &= \frac{1}{N} \sum_i [\mathbf{u}^T(\mathbf{x}_i - \text{mean}(\{\mathbf{x}\}))] [(\mathbf{x}_i - \text{mean}(\{\mathbf{x}\}))^T \mathbf{u}] \\ &= \frac{1}{N} \sum_i [\mathbf{u}^T(\mathbf{x}_i - \text{mean}(\{\mathbf{x}\}))]^2. \end{aligned}$$

Now this is a sum of squares. If there is some  $\mathbf{a}$  such that  $\mathbf{a}^T(\mathbf{x}_i - \text{mean}(\{\mathbf{x}\})) = 0$  for every  $i$ , then the covariance matrix must be positive semidefinite (because the sum of squares could be zero in this case). Otherwise, it is positive definite, because the sum of squares will always be positive.

## 10.2 BLOB ANALYSIS OF HIGH-DIMENSIONAL DATA

### 10.2.1 Transforming High Dimensional Data

Assume we apply an affine transformation to our data set  $\{\mathbf{x}\}$ , to obtain a new dataset  $\{\mathbf{u}\}$ , where  $\mathbf{u}_i = \mathcal{A}\mathbf{x}_i + \mathbf{b}$ . Here  $\mathcal{A}$  is any matrix (it doesn't have to be

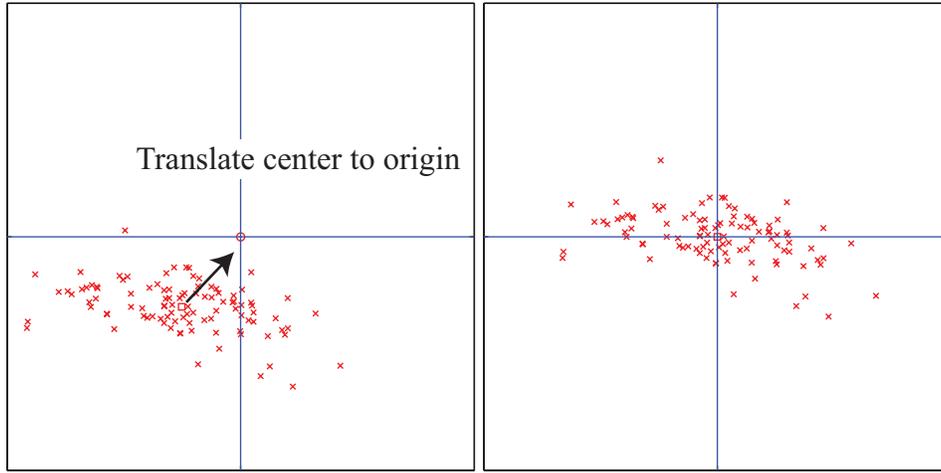


FIGURE 10.8: On the **left**, a “blob” in two dimensions. This is a set of data points that lie somewhat clustered around a single center, given by the mean. I have plotted the mean of these data points with a hollow square (it’s easier to see when there is a lot of data). To translate the blob to the origin, we just subtract the mean from each datapoint, yielding the blob on the **right**.

square, or symmetric, or anything else; it just has to have second dimension  $d$ ). It is easy to compute the mean and covariance of  $\{\mathbf{u}\}$ . We have

$$\begin{aligned}\text{mean}(\{\mathbf{u}\}) &= \text{mean}(\{\mathcal{A}\mathbf{x} + \mathbf{b}\}) \\ &= \mathcal{A}\text{mean}(\{\mathbf{x}\}) + \mathbf{b},\end{aligned}$$

so you get the new mean by multiplying the original mean by  $\mathcal{A}$  and adding  $\mathbf{b}$ . The new covariance matrix is easy to compute as well. We have:

$$\begin{aligned}\text{Covmat}(\{\mathbf{u}\}) &= \text{Covmat}(\{\mathcal{A}\mathbf{x} + \mathbf{b}\}) \\ &= \frac{\sum_i (\mathbf{u}_i - \text{mean}(\{\mathbf{u}\}))(\mathbf{u}_i - \text{mean}(\{\mathbf{u}\}))^T}{N} \\ &= \frac{\sum_i (\mathcal{A}\mathbf{x}_i + \mathbf{b} - \mathcal{A}\text{mean}(\{\mathbf{x}\}) - \mathbf{b})(\mathcal{A}\mathbf{x}_i + \mathbf{b} - \mathcal{A}\text{mean}(\{\mathbf{x}\}) - \mathbf{b})^T}{N} \\ &= \frac{\mathcal{A} \sum_i (\mathbf{x}_i - \text{mean}(\{\mathbf{x}\}))(\mathbf{x}_i - \text{mean}(\{\mathbf{x}\}))^T \mathcal{A}^T}{N} \\ &= \mathcal{A} \text{Covmat}(\{\mathbf{x}\}) \mathcal{A}^T.\end{aligned}$$

### 10.2.2 Transforming Blobs

The trick to interpreting high dimensional data is to use the mean and covariance to understand the blob. Figure 10.8 shows a two-dimensional data set. Notice that there is obviously some correlation between the  $x$  and  $y$  coordinates (it’s a diagonal blob), and that neither  $x$  nor  $y$  has zero mean. We can easily compute the mean and subtract it from the data points, and this translates the blob so that the origin

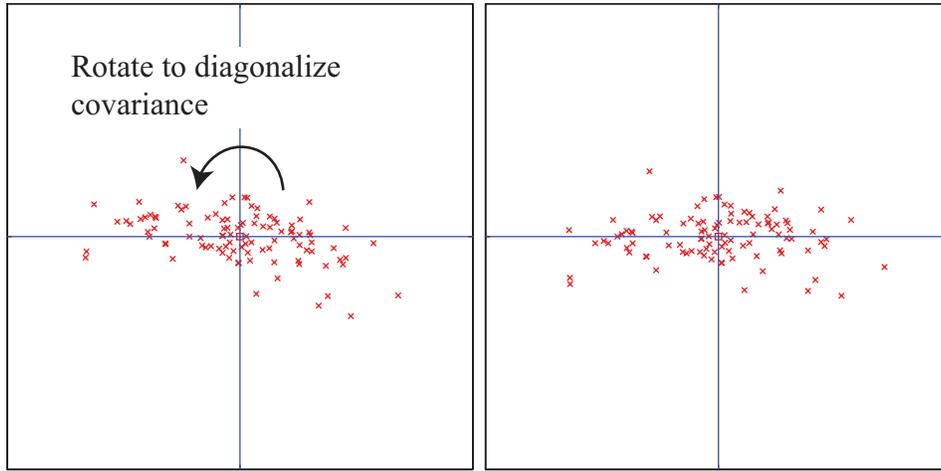


FIGURE 10.9: *On the left*, the translated blob of figure 10.8. This blob lies somewhat diagonally, because the vertical and horizontal components are correlated. *On the right*, that blob of data rotated so that there is no correlation between these components. We can now describe the blob by the vertical and horizontal variances alone, as long as we do so in the new coordinate system. In this coordinate system, the vertical variance is significantly larger than the horizontal variance — the blob is short and wide.

is at the center (Figure 10.8). In coordinates, this means we compute the new dataset  $\{\mathbf{u}\}$  from the old dataset  $\{\mathbf{x}\}$  by the rule  $\mathbf{u}_i = \mathbf{x}_i - \text{mean}(\{\mathbf{x}\})$ . This new dataset has been translated so that the mean is zero.

Once this blob is translated (Figure 10.9, left), we can rotate it as well. It is natural to try to rotate the blob so that there is no correlation between distinct pairs of dimensions. We can do so by diagonalizing the covariance matrix. In particular, let  $\mathcal{U}$  be the matrix formed by stacking the eigenvectors of  $\text{Covmat}(\{\mathbf{x}\})$  into a matrix (i.e.  $\mathcal{U} = [\mathbf{v}_1, \dots, \mathbf{v}_d]$ , where  $\mathbf{v}_j$  are eigenvectors of the covariance matrix). We now form the dataset  $\{\mathbf{n}\}$ , using the rule

$$\mathbf{n}_i = \mathcal{U}^T \mathbf{u}_i = \mathcal{U}^T (\mathbf{x}_i - \text{mean}(\{\mathbf{x}\})).$$

The mean of this new dataset is clearly  $\mathbf{0}$ . The covariance of this dataset is

$$\begin{aligned} \text{Covmat}(\{\mathbf{n}\}) &= \text{Covmat}(\{\mathcal{U}^T \mathbf{x}\}) \\ &= \mathcal{U}^T \text{Covmat}(\{\mathbf{x}\}) \mathcal{U} \\ &= \Lambda, \end{aligned}$$

where  $\Lambda$  is a diagonal matrix of eigenvalues of  $\text{Covmat}(\{\mathbf{x}\})$ . Remember that, in describing diagonalization, we adopted the convention that the eigenvectors of the matrix being diagonalized were ordered so that the eigenvalues are sorted in descending order along the diagonal of  $\Lambda$ . We now have two very useful facts about  $\{\mathbf{n}\}$ : (a) every pair of distinct components has covariance zero, and so has correlation zero; (b) the first component has the highest variance, the second component

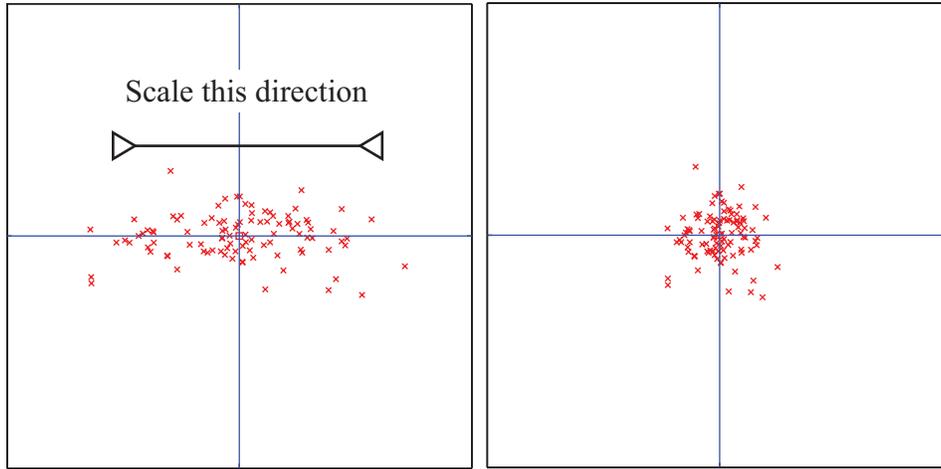


FIGURE 10.10: *On the left, the translated and rotated blob of figure 10.9. This blob is stretched — one direction has more variance than another. Because all covariances are zero, it is easy to scale the blob so that all variances are one (the blob on the right). You can think of this as a standard blob. All blobs can be reduced to a standard blob, by relatively straightforward linear algebra.*

has the second highest variance, and so on. We can rotate and translate any blob into a coordinate system that has these properties. *In this coordinate system*, we can describe the blob simply by giving the variances of each component — the covariances are zero.

Translating a blob of data doesn't change the scatterplot matrix in any interesting way (the axes change, but the picture doesn't). Rotating a blob produces really interesting results, however. Figure 10.11 shows the dataset of figure 10.6, translated to the origin and rotated to diagonalize it. Now we do not have names for each component of the data (they're linear combinations of the original components), but each pair is now not correlated. This blob has some interesting shape features. Figure 10.11 shows the gross shape of the blob best. Each panel of this figure has the same scale in each direction. You can see the blob extends about 80 units in direction 1, but only about 15 units in direction 2, and much less in the other two directions. You should think of this blob as being rather cigar-shaped; it's long in one direction, but there isn't much in the others. The cigar metaphor isn't perfect because there aren't any 4 dimensional cigars, but it's helpful. You can think of each panel of this figure as showing views down each of the four axes of the cigar.

Now look at figure 10.12. This shows the same rotation of the same blob of data, but now the scales on the axis have changed to get the best look at the detailed shape of the blob. First, you can see that blob is a little curved (look at the projection onto direction 2 and direction 4). There might be some effect here worth studying. Second, you can see that some points seem to lie away from the main blob. I have plotted each data point with a dot, and the interesting points

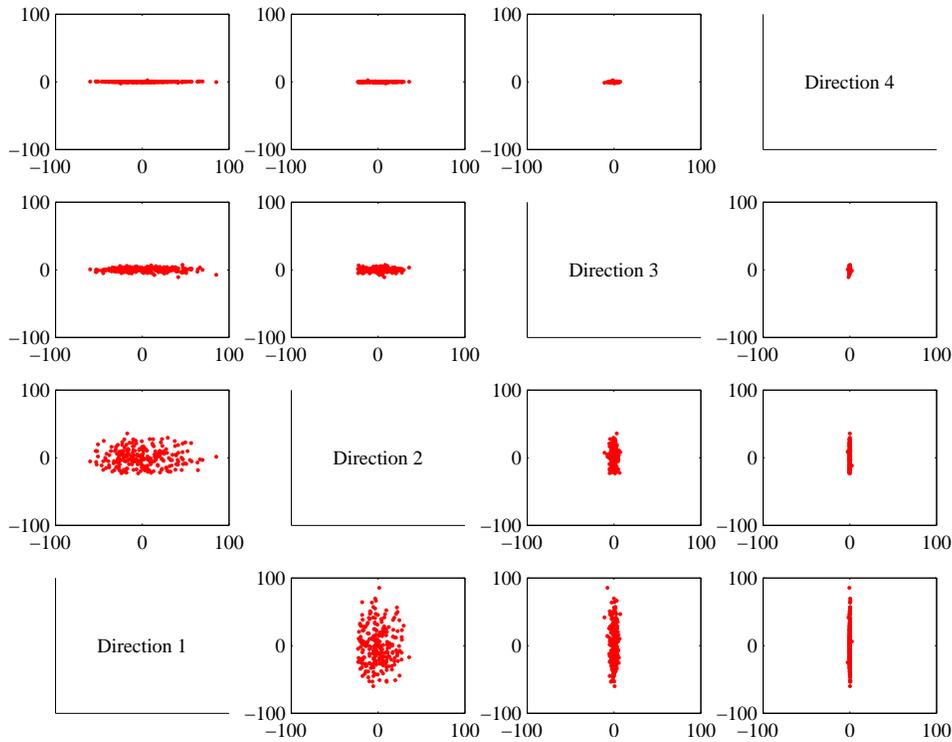


FIGURE 10.11: A panel plot of the bodyfat dataset of figure 10.6, now rotated so that the covariance between all pairs of distinct dimensions is zero. Now we do not know names for the directions — they’re linear combinations of the original variables. Each scatterplot is on the same set of axes, so you can see that the dataset extends more in some directions than in others.

with a number. These points are clearly special in some way.

We could now scale the data in this new coordinate system so that all the variances are either one (if there is any variation in that direction) or zero (directions where the data doesn’t vary — these occur only if some directions are functions of others). Figure 10.10 shows the final scaling. The result is a standard blob. Our approach applies to any dimension — I gave 2D figures because they’re much easier to understand. There is a crucial point here: we can reduce any blob of data, in any dimension, to a standard blob of that dimension. All blobs are the same, except for some stretching, some rotation, and some translation. This is why blobs are so well-liked.

### 10.2.3 Whitening Data

It is sometimes useful to actually reduce a dataset to a standard blob. Doing so is known as whitening the data (for reasons I find obscure). This can be a sensible thing to do when we don’t have a clear sense of the relative scales of the components

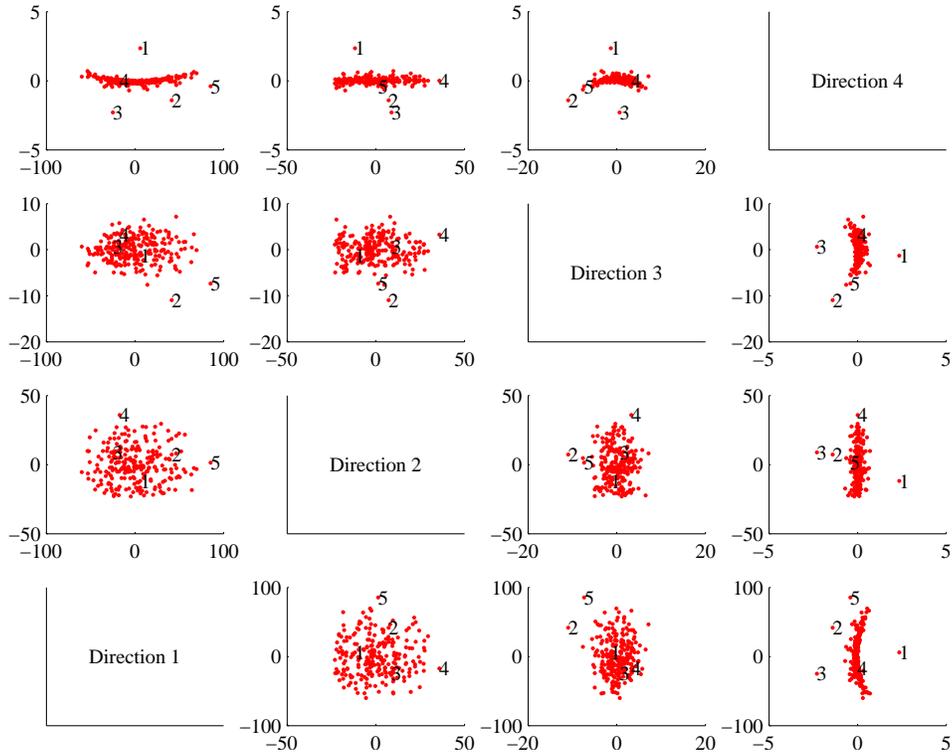


FIGURE 10.12: A panel plot of the bodyfat dataset of figure 10.6, now rotated so that the covariance between all pairs of distinct dimensions is zero. Now we do not know names for the directions — they’re linear combinations of the original variables. I have scaled the axes so you can see details; notice that the blob is a little curved, and there are several data points that seem to lie some way away from the blob, which I have numbered.

of each data vector. For example, if we have a dataset where one component ranges from  $1e5$  to  $2e5$ , and the other component ranges from  $-1e-5$  to  $1e-5$ , we are likely to face numerical problems in many computations (adding small numbers to big numbers is often unwise). Often, this kind of thing follows from a poor choice of units, rather than any kind of useful property of the data. In such a case, it could be quite helpful to whiten the data. Another reason to whiten the data might be that we know relatively little about the meaning of each component. In this case, the original choice of coordinate system was somewhat arbitrary anyhow, and transforming data to a uniform blob could be helpful.

**Useful Facts: 10.1** *Whitening a dataset*

For a dataset  $\{\mathbf{x}\}$ , compute:

- $\mathcal{U}$ , the matrix of eigenvectors of  $\text{Covmat}(\{\mathbf{x}\})$ ;
- and  $\text{mean}(\{\mathbf{x}\})$ .

Now compute  $\{\mathbf{n}\}$  using the rule

$$\mathbf{n}_i = \mathcal{U}^T(\mathbf{x}_i - \text{mean}(\{\mathbf{x}\})).$$

Then  $\text{mean}(\{\mathbf{n}\}) = \mathbf{0}$  and  $\text{Covmat}(\{\mathbf{n}\})$  is diagonal.

Now write  $\Lambda$  for the diagonal matrix of eigenvalues of  $\text{Covmat}(\{\mathbf{x}\})$  (so that  $\text{Covmat}(\{\mathbf{x}\})\mathcal{U} = \mathcal{U}\Lambda$ ). Assume that each of the diagonal entries of  $\Lambda$  is greater than zero (otherwise there is a redundant dimension in the data). Write  $\lambda_i$  for the  $i$ 'th diagonal entry of  $\Lambda$ , and write  $\Lambda^{-(1/2)}$  for the diagonal matrix whose  $i$ 'th diagonal entry is  $1/\sqrt{\lambda_i}$ . Compute  $\{\mathbf{z}\}$  using the rule

$$\mathbf{z}_i = \Lambda^{-(1/2)}\mathcal{U}(\mathbf{x}_i - \text{mean}(\{\mathbf{x}\})).$$

We have that  $\text{mean}(\{\mathbf{z}\}) = \mathbf{0}$  and  $\text{Covmat}(\{\mathbf{z}\}) = \mathcal{I}$ . The dataset  $\{\mathbf{z}\}$  is often known as whitened data.

Indexing: Whitening a dataset

Indexing: Whitening a dataset

It isn't always a good idea to whiten data. In some circumstances, each separate component is meaningful, and in a meaningful set of units. For example, one of the components might be a length using a natural scale and the other might be a time on a natural scale. When this happens, we might be reluctant to transform the data, either because we don't want to add lengths to times or because we want to preserve the scales.

### 10.3 PRINCIPAL COMPONENTS ANALYSIS

Mostly, when one deals with high dimensional data, it isn't clear which individual components are important. As we have seen with the height weight dataset (for example, in the case of density and weight) some components can be quite strongly correlated. Equivalently, as in Figure 13.1, the blob is not aligned with the coordinate axes.

#### 10.3.1 The Blob Coordinate System and Smoothing

We can use the fact that we *could* rotate, translate and scale the blob to define a coordinate system within the blob. The origin of that coordinate system is the mean of the data, and the coordinate axes are given by the eigenvectors of the

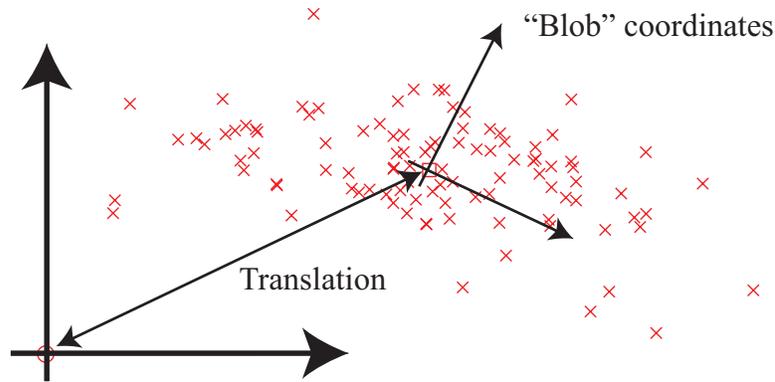


FIGURE 10.13: A 2D blob, with its natural blob coordinate system. The origin of this coordinate system is at the mean of the data. The coordinate axes are (a) at right angles to one another and (b) are directions that have no covariance.

covariance matrix. These are orthonormal, so they form a set of unit basis vectors at right angles to one another (i.e. a coordinate system). You should think of these as blob coordinates; Figure 10.13 illustrates a set of blob coordinates.

The blob coordinate system is important because, *once we know the blob coordinate system*, we can identify important scales of variation in the data. For example, if you look at Figure 10.13, you can see that this blob is extended much further along one direction than along the other. We can use this information to identify the most significant forms of variation in very high dimensional data. In some directions in the blob coordinate system, the blob will be spread out — ie have large variance — but in others, it might not be.

Equivalently, imagine we choose to represent each data item *in blob coordinates*. Then the mean over the dataset will be zero. Each pair of distinct coordinates will be uncorrelated. Some coordinates — corresponding to directions where the blob is spread out — will have a large range of values. Other coordinates — directions in which the blob is small — will have a small range of values. We could choose to replace these coordinates with zeros, with little significant loss in accuracy. The advantage of doing so is that we would have lower dimensional data to deal with.

However, it isn't particularly natural to work in blob coordinates. Each component of a data item may have a distinct meaning and scale (i.e. feet, pounds, and so on), but this is not preserved in any easy way in blob coordinates. Instead, we should like to (a) compute a lower dimensional representation in blob coordinates then (b) transform that representation into the original coordinate system of the data item. Doing so is a form of **smoothing** — suppressing small, irrelevant variations by exploiting multiple data items.

For example, look at Figure 10.14. Imagine we transform the blob on the left to blob coordinates. The covariance matrix in these coordinates is a  $3 \times 3$  diagonal matrix. One of the values on the diagonal is large, because the blob is extended on one direction; but the other two are small. This means that, in blob coordinates, the

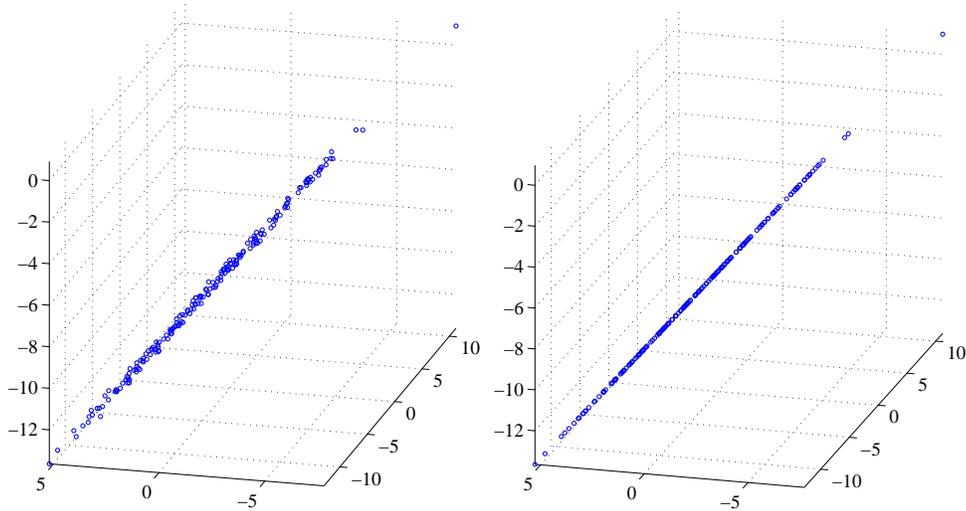


FIGURE 10.14: On the **left**, a blob of 3D data that has very low variance in two directions in the blob coordinates. As a result, all the data points are very close to a 1D blob. Experience shows that this is a common phenomenon. Although there might be many components in the data items, all data points are very close to a much lower dimensional object in the high dimensional space. When this is the case, we could obtain a lower dimensional representation of the data by working in blob coordinates, or we could smooth the data (as on the **right**), by projecting each data point onto the lower dimensional space.

data varies significantly in one direction, but very little in the other two directions.

Now imagine we project the data points onto the high-variation direction; equivalently, we set the other two directions to zero for each data point. Each of the new data points is very close to the corresponding old data point, because by setting the small directions to zero we haven't moved the point very much. In blob coordinates, the covariance matrix of this new dataset has changed very little. It is again a  $3 \times 3$  diagonal matrix, but now two of the diagonal values are zero, because there isn't any variance in those directions. The third value is large, because the blob is extended in that direction. We take the new dataset, and rotate and translate it into the original coordinate system. Each point must lie close to the corresponding point in the original dataset. However, the new dataset lies along a straight line (because it lay on a straight line in the blob coordinates). This process gets us the blob on the right in Figure 10.14. This blob is a smoothed version of the original blob.

Smoothing works because when two data items are strongly correlated, the value of one is a good guide to the value of the other. This principle works for more than two data items. Section 13.1 describes an example where the data items have dimension 101, but all values are extremely tightly correlated. In a case like this, there may be very few dimensions in blob coordinates that have any significant variation (3-6 for this case, depending on some details of what one

believes is a small number, and so on). The components are so strongly correlated in this case that the 101-dimensional blob really looks like a slightly thickened 3 (or slightly more) dimensional blob that has been inserted into a 101-dimensional space (Figure 10.14). If we project the 101-dimensional data onto that structure in the original, 101-dimensional space, we may get much better estimates of the components of each data item than the original measuring device can supply. This occurs because each component is now estimated using correlations between all the measurements.

### 10.3.2 The Low-Dimensional Representation of a Blob

We wish to construct an  $r$  dimensional representation of a blob, where we have chosen  $r$  in advance. First, we compute  $\{\mathbf{v}\}$  by translating the blob so its mean is at the origin, so that  $\mathbf{v}_i = \mathbf{x}_i - \text{mean}(\{\mathbf{x}\})$ . Now write  $\mathcal{V} = [\mathbf{v}_1, \mathbf{v}_2, \dots, \mathbf{v}_N]$ . The covariance matrix of  $\{\mathbf{v}\}$  is then

$$\text{Covmat}(\{\mathbf{v}\}) = \frac{1}{N} \mathcal{V} \mathcal{V}^T = \text{Covmat}(\{\mathbf{x}\}).$$

Now write  $\Lambda$  for the diagonal matrix of eigenvalues of  $\text{Covmat}(\{\mathbf{x}\})$  and  $\mathcal{U}$  for the matrix of eigenvectors, so that  $\text{Covmat}(\{\mathbf{x}\})\mathcal{U} = \mathcal{U}\Lambda$ . We assume that the elements of  $\Lambda$  are sorted in decreasing order along the diagonal. The covariance matrix for the dataset transformed into blob coordinates will be  $\Lambda$ . Notice that

$$\begin{aligned} \Lambda &= \mathcal{U}^T \text{Covmat}(\{\mathbf{x}\}) \mathcal{U} \\ &= \mathcal{U}^T \mathcal{V} \mathcal{V}^T \mathcal{U} \\ &= (\mathcal{U}^T \mathcal{V})(\mathcal{U}^T \mathcal{V})^T. \end{aligned}$$

This means we can interpret  $(\mathcal{U}^T \mathcal{V})$  as a new dataset  $\{\mathbf{b}\}$ . This is our data, rotated into blob coordinates.

Now write  $\Pi_r$  for the  $d \times d$  matrix

$$\begin{bmatrix} \mathcal{I}_r & 0 \\ 0 & 0 \end{bmatrix}$$

which projects a  $d$  dimensional vector onto its first  $r$  components, and replaces the others with zeros. Then we have that

$$\Lambda_r = \Pi_r \Lambda \Pi_r^T$$

is the covariance matrix for the reduced dimensional data in blob coordinates. Notice that  $\Lambda_r$  keeps the  $r$  largest eigenvalues on the diagonal of  $\Lambda$ , and replaces all others with zero.

We have

$$\begin{aligned} \Lambda_r &= \Pi_r \Lambda \Pi_r^T \\ &= \Pi_r \mathcal{U}^T \text{Covmat}(\{\mathbf{x}\}) \mathcal{U} \Pi_r^T \\ &= (\Pi_r \mathcal{U}^T \mathcal{V})(\mathcal{V}^T \mathcal{U} \Pi_r^T) \\ &= \mathcal{P} \mathcal{P}^T \end{aligned}$$

where  $\mathcal{P} = (\Pi_r \mathcal{U}^T \mathcal{V})$ . This represents our data, rotated into blob coordinates, and then projected down to  $r$  dimensions, with remaining terms replaced by zeros. Write  $\{\mathbf{b}_r\}$  for this new dataset.

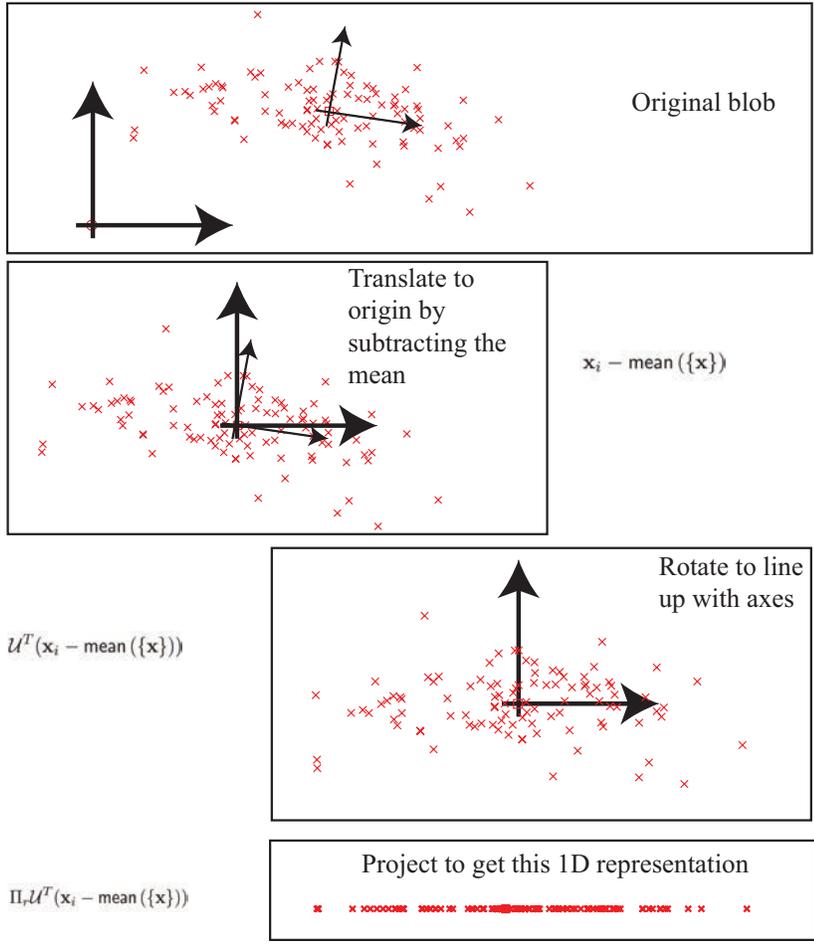


FIGURE 10.15: Computing a low dimensional representation for principal components analysis.

Occasionally, we need to refer to this representation, and we give it a special name. Write

$$\text{pcaproj}(\mathbf{x}_i, r, \{\mathbf{x}\}) = \Pi_r \mathcal{U}^T(\mathbf{x}_i - \text{mean}(\{\mathbf{x}\}))$$

where the notation seeks to explicitly keep track of the fact that the low dimensional representation of a particular data item *depends on the whole dataset* (because you have to be able to compute the mean, and the eigenvectors of the covariance). Notice that  $\text{pcaproj}(\mathbf{x}_i, r, \{\mathbf{x}\})$  is a representation of the dataset with important properties:

- The representation is  $r$ -dimensional (i.e. the last  $d - r$  components are zero).

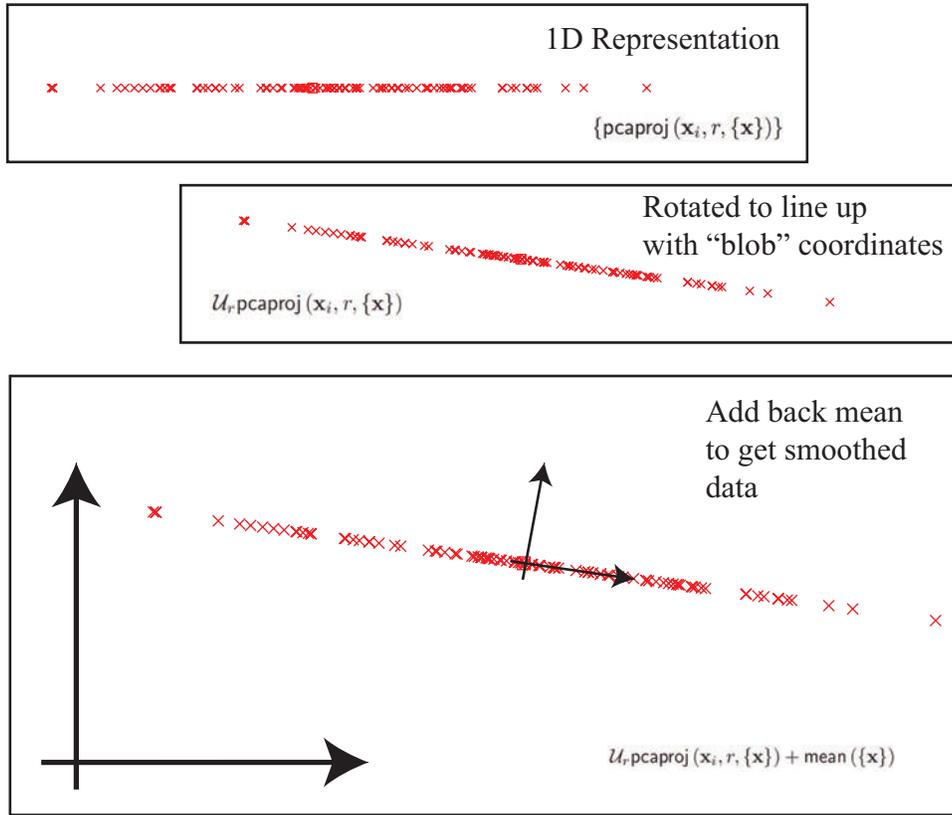


FIGURE 10.16: Smoothing data with principal components analysis.

- Each pair of distinct components of  $\{\text{pcproj}(\mathbf{x}_i, r, \{\mathbf{x}\})\}$  has zero covariance.
- The first component of  $\{\text{pcproj}(\mathbf{x}_i, r, \{\mathbf{x}\})\}$  has largest variance; the second component has second largest variance; and so on.

### 10.3.3 Smoothing Data with a Low-Dimensional Representation

We would now like to construct a low dimensional representation of the blob, in the original coordinates. We do so by rotating the low-dimensional representation back to the original coordinate system, then adding back the mean to translate the origin back to where it started. We can write this as

$$\begin{aligned} \text{pcsmooth}(\mathbf{x}_i, r, \{\mathbf{x}\}) &= \mathcal{U}\Pi_r^T (\Pi_r \mathcal{U}^T (\mathbf{x}_i - \text{mean}(\{\mathbf{x}\}))) + \text{mean}(\{\mathbf{x}\}) \\ &= \mathcal{U}\Pi_r^T \text{pcproj}(\mathbf{x}_i, r, \{\mathbf{x}\}) + \text{mean}(\{\mathbf{x}\}) \end{aligned}$$

we have a new representation of the  $i$ 'th data item *in the original space* (Figure 13.1). Now consider the dataset obtained by smoothing each of our data items. We write this dataset as  $\{\text{pcsmooth}(\mathbf{x}_i, r, \{\mathbf{x}\})\}$ .

You should think of  $\{\text{pcsmooth}(\mathbf{x}_i, r, \{\mathbf{x}\})\}$  as a smoothed version of the original data. One way to think of this process is that we have chosen a low-dimensional basis that represents the main variance in the data rather well. It is quite usual to think of a data item as being given by the mean plus a weighted sum of these basis elements. In this view, the first weight has larger variance than the second, and so on. By construction, this dataset lies in an  $r$  dimensional affine subspace of the original space. We constructed this  $r$ -dimensional space to preserve the largest variance directions of the data. Each column of this matrix is known as a **principal component**. In particular, we have constructed this dataset so that

- $\text{mean}(\{\text{pcsmooth}(\mathbf{x}_i, r, \{\mathbf{x}\})\}) = \text{mean}(\{\mathbf{x}\})$ ;
- $\text{Covmat}(\{\text{pcsmooth}(\mathbf{x}_i, r, \{\mathbf{x}\})\})$  has rank  $r$ ;
- $\text{Covmat}(\{\text{pcsmooth}(\mathbf{x}_i, r, \{\mathbf{x}\})\})$  is the best approximation of  $\text{Covmat}(\{\mathbf{x}\})$  with rank  $r$ .

Figure 10.16 gives a visualization of the smoothing process. By comparing figures 10.11 and 10.17, you can see that a real dataset can lose several dimensions without much significant going wrong. As we shall see in the examples, some datasets can lose many dimensions without anything bad happening.

**Procedure: 10.1** *Principal Components Analysis*

Assume we have a general data set  $\mathbf{x}_i$ , consisting of  $N$   $d$ -dimensional vectors. Now write  $\Sigma = \text{Covmat}(\{\mathbf{x}\})$  for the covariance matrix.

Form  $\mathcal{U}$ ,  $\Lambda$ , such that  $\Sigma\mathcal{U} = \mathcal{U}\Lambda$  (these are the eigenvectors and eigenvalues of  $\Sigma$ ). Ensure that the entries of  $\Lambda$  are sorted in decreasing order. Choose  $r$ , the number of dimensions you wish to represent. Typically, we do this by plotting the eigenvalues and looking for a “knee” (Figure ??). It is quite usual to do this by hand.

**Constructing a low-dimensional representation:** Form  $\mathcal{U}_r$ , a matrix consisting of the first  $r$  columns of  $\mathcal{U}$ . Now compute  $\{\text{pcproj}(\mathbf{x}_i, r, \{\mathbf{x}\})\} = \{(\Pi_r \mathcal{U}^T (\mathbf{x}_i - \text{mean}(\{\mathbf{x}\})))\}$ . This is a set of data vectors which are  $r$  dimensional, and where each component is independent of each other component (i.e. the covariances of distinct components are zero).

**Smoothing the data:** Form  $\{\text{pcsmooth}(\mathbf{x}_i, r, \{\mathbf{x}\})\} = \{(\mathcal{U}_r \text{pcproj}(\mathbf{x}_i, r, \{\mathbf{x}\}) + \text{mean}(\{\mathbf{x}\}))\}$ . These are  $d$  dimensional vectors that lie in a  $r$ -dimensional subspace of  $d$ -dimensional space. The “missing dimensions” have the lowest variance, and are independent.

#### 10.3.4 The Error of the Low-Dimensional Representation

We took a dataset,  $\{\mathbf{x}\}$ , and constructed a  $d$ -dimensional dataset  $\{\mathbf{b}\}$  in blob coordinates. We did so by translating, then rotating, the data, so no information was

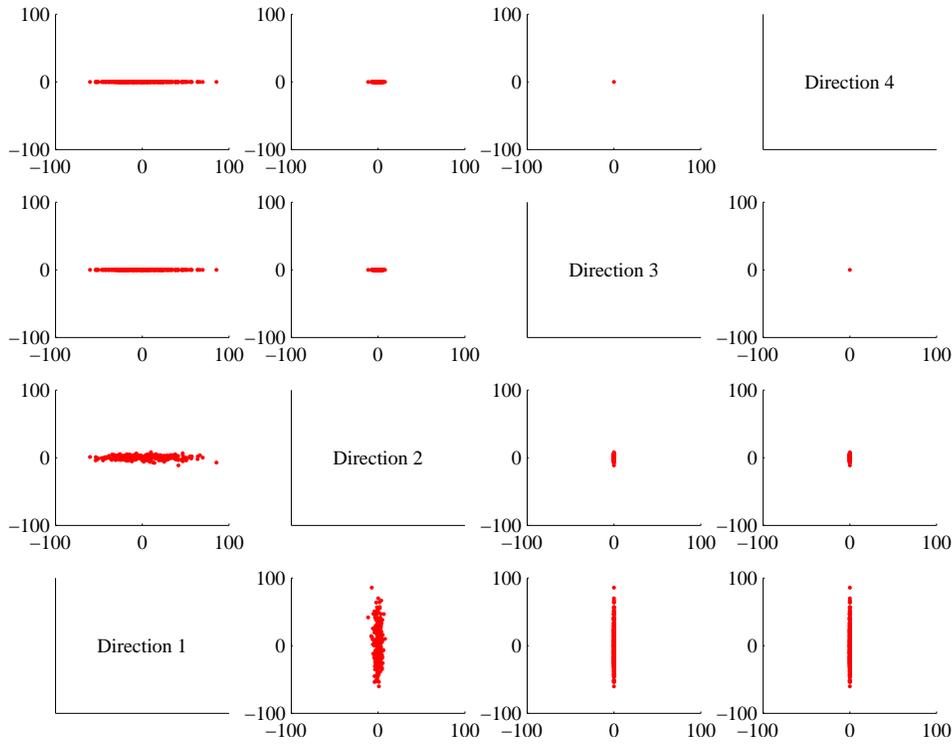


FIGURE 10.17: A panel plot of the bodyfat dataset of figure 10.6, with the dimension reduced to two using principal components analysis. Compare this figure to figure 10.11, which is on the same set of axes. You can see that the blob has been squashed in direction 3 and direction 4. But not much has really happened, because there wasn't very much variation in those directions in the first place.

lost; we could reconstruct our original dataset by rotating, then translating  $\{\mathbf{b}\}$ . But in blob coordinates we projected each data item down to the first  $r$  components to get an  $r$ -dimensional dataset  $\{\mathbf{b}_r\}$ . We then reconstructed a smoothed dataset by rotating, then translating,  $\{\mathbf{b}_r\}$ . Information has been lost here, but how much?

The answer is easy to get if you recall that rotations and translations do not change lengths. This means that

$$\|\mathbf{x}_i - \text{pcsmooth}(\mathbf{x}_i, r, \{\mathbf{x}\})\|^2 = \|\mathbf{b}_i - \mathbf{b}_{r,i}\|^2.$$

This expression is easy to evaluate, because  $\mathbf{b}_i$  and  $\mathbf{b}_{r,i}$  agree in their first  $r$  components. The remaining  $d - r$  components of  $\mathbf{b}_{r,i}$  are zero. So we can write

$$\|\mathbf{x}_i - \text{pcsmooth}(\mathbf{x}_i, r, \{\mathbf{x}\})\|^2 = \sum_{u=r+1}^d (\mathbf{b}_i^{(u)})^2.$$

Now a natural measure of error is the average over the dataset of this term. We

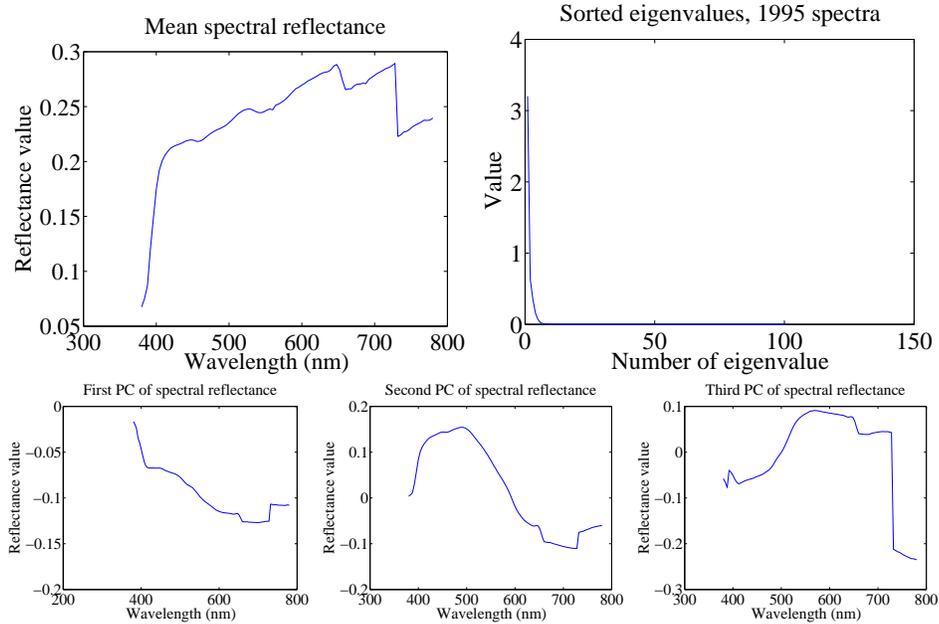


FIGURE 10.18: On the **top left**, the mean spectral reflectance of a dataset of 1995 spectral reflectances, collected by Kobus Barnard (at <http://www.cs.sfu.ca/~colour/data/>). On the **top right**, eigenvalues of the covariance matrix of spectral reflectance data, from a dataset of 1995 spectral reflectances, collected by Kobus Barnard (at <http://www.cs.sfu.ca/~colour/data/>). Notice how the first few eigenvalues are large, but most are very small; this suggests that a good representation using few principal components is available. The **bottom row** shows the first three principal components. A linear combination of these, with appropriate weights, added to the mean of figure ??, gives a good representation of the dataset.

have that

$$\frac{1}{N} \sum_{u=r+1}^d (\mathbf{b}_i^{(u)})^2 = \sum_{u=r+1}^d \text{var}(\{\mathbf{b}^{(u)}\})$$

which is easy to evaluate, because we know these variances — they are the values of the  $d - r$  eigenvalues that we decided to ignore. So the mean error can be written as

$$\mathbf{1}^T (\Lambda - \Lambda_r) \mathbf{1}.$$

Now we could choose  $r$  by identifying how much error we can tolerate. More usual is to plot the eigenvalues of the covariance matrix, and look for a “knee”, like that in Figure 13.1. You can see that the sum of remaining eigenvalues is small.

### 10.3.5 Example: Representing Spectral Reflectances

Diffuse surfaces reflect light uniformly in all directions. Examples of diffuse surfaces include matte paint, many styles of cloth, many rough materials (bark, cement,

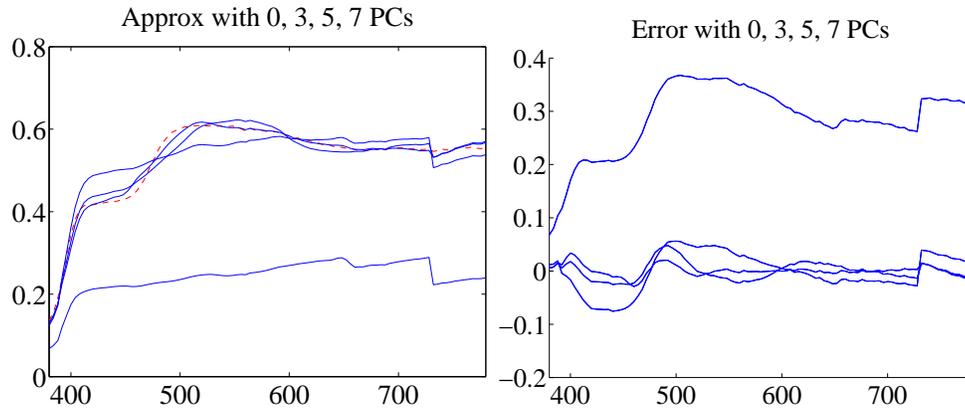


FIGURE 10.19: *On the left*, a spectral reflectance curve (dashed) and approximations using the mean, the mean and 3 principal components, the mean and 5 principal components, and the mean and 7 principal components. Notice the mean is a relatively poor approximation, but as the number of principal components goes up, the error falls rather quickly. *On the right* is the error for these approximations. Figure plotted from a dataset of 1995 spectral reflectances, collected by Kobus Barnard (at <http://www.cs.sfu.ca/~colour/data/>).

stone, etc.). One way to tell a diffuse surface is that it does not look brighter (or darker) when you look at it along different directions. Diffuse surfaces can be colored, because the surface reflects different fractions of the light falling on it at different wavelengths. This effect can be represented by measuring the spectral reflectance of a surface, which is the fraction of light the surface reflects as a function of wavelength. This is usually measured in the visual range of wavelengths (about 380nm to about 770 nm). Typical measurements are every few nm, depending on the measurement device. I obtained data for 1995 different surfaces from <http://www.cs.sfu.ca/~colour/data/> (there are a variety of great datasets here, from Kobus Barnard).

Each spectrum has 101 measurements, which are spaced 4nm apart. This represents surface properties to far greater precision than is really useful. Physical properties of surfaces suggest that the reflectance can't change too fast from wavelength to wavelength. It turns out that very few principal components are sufficient to describe almost any spectral reflectance function. Figure 10.18 shows the mean spectral reflectance of this dataset, and Figure 10.18 shows the eigenvalues of the covariance matrix.

This is tremendously useful in practice. One should think of a spectral reflectance as a function, usually written  $\rho(\lambda)$ . What the principal components analysis tells us is that we can represent this function rather accurately on a (really small) finite dimensional basis. This basis is shown in figure 10.18. This means

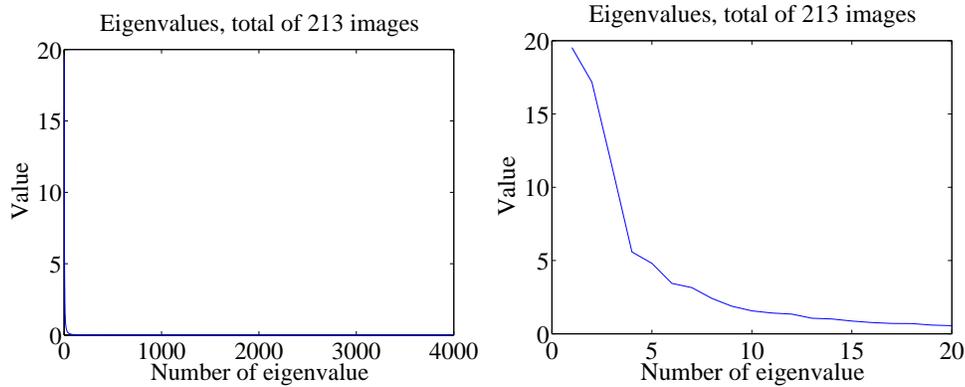


FIGURE 10.20: On the **left**, the eigenvalues of the covariance of the Japanese facial expression dataset; there are 4096, so it's hard to see the curve (which is packed to the left). On the **right**, a zoomed version of the curve, showing how quickly the values of the eigenvalues get small.

that there is a mean function  $r(\lambda)$  and  $k$  functions  $\phi_m(\lambda)$  such that, for any  $\rho(\lambda)$ ,

$$\rho(\lambda) = r(\lambda) + \sum_{i=1}^k c_i \phi_i(\lambda) + e(\lambda)$$

where  $e(\lambda)$  is the error of the representation, which we know is small (because it consists of all the other principal components, which have tiny variance). In the case of spectral reflectances, using a value of  $k$  around 3-5 works fine for most applications (Figure 10.19). This is useful, because when we want to predict what a particular object will look like under a particular light, we don't need to use a detailed spectral reflectance model; instead, it's enough to know the  $c_i$  for that object. This comes in useful in a variety of rendering applications in computer graphics. It is also the key step in an important computer vision problem, called **color constancy**. In this problem, we see a picture of a world of colored objects under unknown colored lights, and must determine what color the objects are. Modern color constancy systems are quite accurate, even though the problem sounds underconstrained. This is because they are able to exploit the fact that relatively few  $c_i$  are enough to accurately describe a surface reflectance.

### 10.3.6 Example: Representing Faces with Principal Components

An image is usually represented as an array of values. We will consider intensity images, so there is a single intensity value in each cell. You can turn the image into a vector by rearranging it, for example stacking the columns onto one another (use `reshape` in Matlab). This means you can take the principal components of a set of images. Doing so was something of a fashionable pastime in computer vision for a while, though there are some reasons that this is not a great representation of pictures. However, the representation yields pictures that can give great intuition into a dataset.

Mean image from Japanese Facial Expression dataset



First sixteen principal components of the Japanese Facial Expression dat



FIGURE 10.21: *The mean and first 16 principal components of the Japanese facial expression dataset.*

Figure ?? shows the mean of a set of face images encoding facial expressions of Japanese women (available at <http://www.kasrl.org/jaffe.html>; there are tons of face datasets at <http://www.face-rec.org/databases/>). I reduced the images to  $64 \times 64$ , which gives a 4096 dimensional vector. The eigenvalues of the covariance of this dataset are shown in figure 10.20; there are 4096 of them, so it's hard to see a trend, but the zoomed figure suggests that the first couple of hundred contain most of the variance. Once we have constructed the principal components, they can be rearranged into images; these images are shown in figure 10.21. Principal components give quite good approximations to real images (figure 10.22).

The principal components sketch out the main kinds of variation in facial expression. Notice how the mean face in Figure 10.21 looks like a relaxed face, but with fuzzy boundaries. This is because the faces can't be precisely aligned, because each face has a slightly different shape. The way to interpret the components is to remember one adjusts the mean towards a data point by adding (or subtracting) some scale times the component. So the first few principal components have to

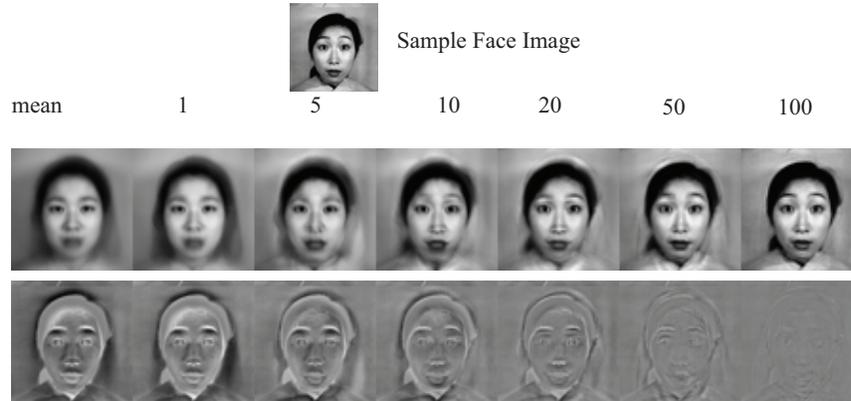


FIGURE 10.22: *Approximating a face image by the mean and some principal components; notice how good the approximation becomes with relatively few components.*

do with the shape of the haircut; by the fourth, we are dealing with taller/shorter faces; then several components have to do with the height of the eyebrows, the shape of the chin, and the position of the mouth; and so on. These are all images of women who are not wearing spectacles. In face pictures taken from a wider set of models, moustaches, beards and spectacles all typically appear in the first couple of dozen principal components.

## 10.4 MULTI-DIMENSIONAL SCALING

One way to get insight into a dataset is to plot it. But choosing what to plot for a high dimensional dataset could be difficult. Assume we must plot the dataset in two dimensions (by far the most common choice). We wish to build a scatter plot in two dimensions — but where should we plot each data point? One natural requirement is that the points be laid out in two dimensions in a way that reflects how they sit in many dimensions. In particular, we would like points that are far apart in the high dimensional space to be far apart in the plot, and points that are close in the high dimensional space to be close in the plot.

### 10.4.1 Principal Coordinate Analysis

We will plot the high dimensional point  $\mathbf{x}_i$  at  $\mathbf{v}_i$ , which is a two-dimensional vector. Now the squared distance between points  $i$  and  $j$  in the high dimensional space is

$$D_{ij}^{(2)}(\mathbf{x}) = (\mathbf{x}_i - \mathbf{x}_j)^T (\mathbf{x}_i - \mathbf{x}_j)$$

(where the superscript is to remind you that this is a squared distance). We could build an  $N \times N$  matrix of squared distances, which we write  $\mathcal{D}^{(2)}(\mathbf{x})$ . The  $i, j$ 'th entry in this matrix is  $D_{ij}^{(2)}(\mathbf{x})$ , and the  $\mathbf{x}$  argument means that the distances are

between points in the high-dimensional space. Now we could choose the  $\mathbf{v}_i$  to make

$$\sum_{ij} \left( D_i^{(2)} j(\mathbf{x}) - D_{ij}^{(2)}(\mathbf{v}) \right)^2$$

as small as possible. Doing so should mean that points that are far apart in the high dimensional space are far apart in the plot, and that points that are close in the high dimensional space are close in the plot.

In its current form, the expression is difficult to deal with, but we can refine it. Because translation does not change the distances between points, it cannot change either of the  $\mathcal{D}^{(2)}$  matrices. So it is enough to solve the case when the mean of the points  $\mathbf{x}_i$  is zero. We can assume that  $\frac{1}{N} \sum_i \mathbf{x}_i = \mathbf{0}$ . Now write  $\mathbf{1}$  for the  $n$ -dimensional vector containing all ones, and  $\mathcal{I}$  for the identity matrix. Notice that

$$D_{ij}^{(2)} = (\mathbf{x}_i - \mathbf{x}_j)^T (\mathbf{x}_i - \mathbf{x}_j) = \mathbf{x}_i \cdot \mathbf{x}_i - 2\mathbf{x}_i \cdot \mathbf{x}_j + \mathbf{x}_j \cdot \mathbf{x}_j.$$

Now write

$$\mathcal{A} = \left[ \mathcal{I} - \frac{1}{N} \mathbf{1}\mathbf{1}^T \right].$$

Using this expression, you can show that the matrix  $\mathcal{M}$ , defined below,

$$\mathcal{M}(\mathbf{x}) = -\frac{1}{2} \mathcal{A} \mathcal{D}^{(2)}(\mathbf{x}) \mathcal{A}^T$$

has  $i, j$ th entry  $\mathbf{x}_i \cdot \mathbf{x}_j$  (exercises). I now argue that, to make  $\mathcal{D}^{(2)}(\mathbf{v})$  is close to  $\mathcal{D}^{(2)}(\mathbf{x})$ , it is enough to make  $\mathcal{M}(\mathbf{v})$  close to  $\mathcal{M}(\mathbf{x})$ . Proving this will take us out of our way unnecessarily, so I omit a proof.

We can choose a set of  $\mathbf{v}_i$  that makes  $\mathcal{D}^{(2)}(\mathbf{v})$  close to  $\mathcal{D}^{(2)}(\mathbf{x})$  quite easily, using the method of the previous section. Take the dataset of  $N$   $d$ -dimensional column vectors  $\mathbf{x}_i$ , and form a matrix  $\mathcal{X}$  by stacking the vectors, so

$$\mathcal{X} = [\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_N].$$

In this notation, we have

$$\mathcal{M}(\mathbf{x}) = \mathcal{X}^T \mathcal{X}.$$

This matrix is symmetric, and it is positive semidefinite. It can't be positive definite, because the data is zero mean, so  $\mathcal{M}(\mathbf{x})\mathbf{1} = 0$ . The  $\mathcal{M}(\mathbf{v})$  we seek must (a) be as close as possible to  $\mathcal{M}(\mathbf{x})$  and (b) have rank 2. It must have rank 2 because there must be some  $\mathcal{V}$  which is  $2 \times N$  so that  $\mathcal{M}(\mathbf{v}) = \mathcal{V}^T \mathcal{V}$ . The columns of this  $\mathcal{V}$  are our  $\mathbf{v}_i$ .

We can use the method of section 14.1.1 to construct  $\mathcal{M}(\mathbf{v})$  and  $\mathcal{V}$ . As usual, we write  $\mathcal{U}$  for the matrix of eigenvectors of  $\mathcal{M}(\mathbf{x})$ ,  $\Lambda$  for the diagonal matrix of eigenvalues sorted in descending order,  $\Lambda_2$  for the  $2 \times 2$  upper left hand block of  $\Lambda$ , and  $\Lambda_2^{(1/2)}$  for the matrix of positive square roots of the eigenvalues. Then our methods yield

$$\mathcal{M}(\mathbf{v}) = \mathcal{U}_2 \Lambda_2^{(1/2)} \Lambda_2^{(1/2)} \mathcal{U}_2^T$$

and

$$\mathcal{V} = \Lambda_2^{(1/2)} \mathcal{U}_2^T$$

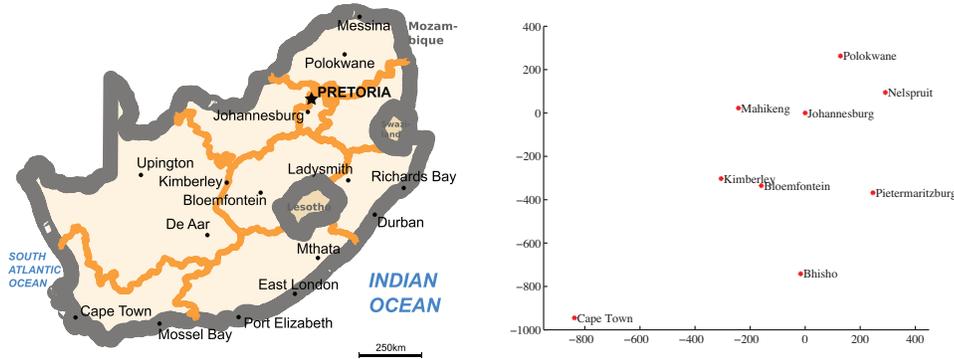


FIGURE 10.23: On the left, a public domain map of South Africa, obtained from [http://commons.wikimedia.org/wiki/File:Map\\_of\\_South\\_Africa.svg](http://commons.wikimedia.org/wiki/File:Map_of_South_Africa.svg), and edited to remove surrounding countries. On the right, the locations of the cities inferred by multidimensional scaling, rotated, translated and scaled to allow a comparison to the map by eye. The map doesn't have all the provincial capitals on it, but it's easy to see that MDS has placed the ones that are there in the right places (use a piece of ruled tracing paper to check).

and we can plot these  $\mathbf{v}_i$  (example in section 13.1). This method for constructing a plot is known as **principal coordinate analysis**.

This plot might not be perfect, because reducing the dimension of the data points should cause some distortions. In many cases, the distortions are tolerable. In other cases, we might need to use a more sophisticated scoring system that penalizes some kinds of distortion more strongly than others. There are many ways to do this; the general problem is known as **multidimensional scaling**.

**Procedure: 10.2** *Principal Coordinate Analysis*

Assume we have a matrix  $D^{(2)}$  consisting of the squared differences between each pair of  $N$  points. We do not need to know the points. We wish to compute a set of points in  $r$  dimensions, such that the distances between these points are as similar as possible to the distances in  $D^{(2)}$ . Form  $\mathcal{A} = [\mathcal{I} - \frac{1}{N}\mathbf{1}\mathbf{1}^T]$ . Form  $\mathcal{W} = \frac{1}{2}\mathcal{A}D^{(2)}\mathcal{A}^T$ . Form  $\mathcal{U}, \Lambda$ , such that  $\mathcal{W}\mathcal{U} = \mathcal{U}\Lambda$  (these are the eigenvectors and eigenvalues of  $\mathcal{W}$ ). Ensure that the entries of  $\Lambda$  are sorted in decreasing order. Choose  $r$ , the number of dimensions you wish to represent. Form  $\Lambda_r$ , the top left  $r \times r$  block of  $\Lambda$ . Form  $\Lambda_r^{(1/2)}$ , whose entries are the positive square roots of  $\Lambda_r$ . Form  $\mathcal{U}_r$ , the matrix consisting of the first  $r$  columns of  $\mathcal{U}$ . Then  $\mathcal{V} = \Lambda_r^{(1/2)}\mathcal{U}_r^T = [\mathbf{v}_1, \dots, \mathbf{v}_N]$  is the set of points to plot.

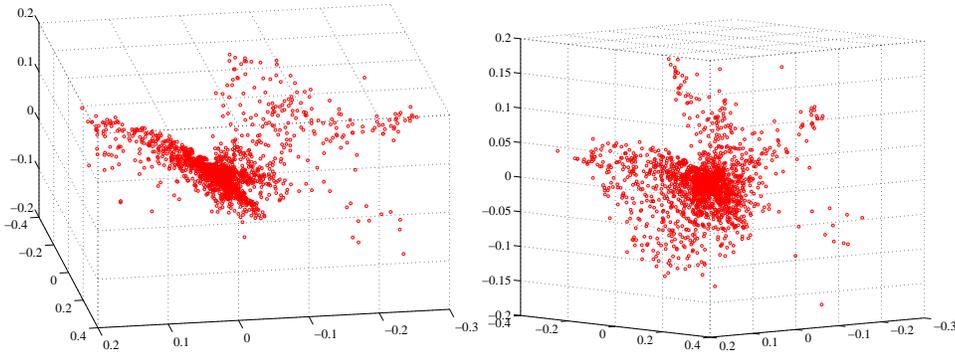


FIGURE 10.24: Two views of the spectral data of section 10.3.5, plotted as a scatter plot by applying principal coordinate analysis to obtain a 3D set of points. Notice that the data spreads out in 3D, but seems to lie on some structure; it certainly isn't a single blob. This suggests that further investigation would be fruitful.

#### 10.4.2 Example: Mapping with Multidimensional Scaling

Multidimensional scaling gets positions (the  $\mathcal{V}$  of section 10.4.1) from distances (the  $\mathcal{D}^{(2)}(\mathbf{x})$  of section 10.4.1). This means we can use the method to build maps from distances alone. I collected distance information from the web (I used <http://www.distancefromto.net>, but a google search on “city distances” yields a wide range of possible sources), then apply multidimensional scaling. Table 10.2 shows distances between the South African provincial capitals, in kilometers, rounded to the nearest kilometer. I then used principal coordinate analysis to find positions for each capital, and rotated, translated and scaled the resulting plot to check it against a real map (Figure 10.23).

One natural use of principal coordinate analysis is to see if one can spot any structure in a dataset. Does the dataset form a blob, or is it clumpy? This isn't a perfect test, but it's a good way to look and see if anything interesting is happening. In figure 10.24, I show a 3D plot of the spectral data, reduced to three dimensions using principal coordinate analysis. The plot is quite interesting. You should notice that the data points are spread out in 3D, but actually seem to lie on a complicated curved surface — they very clearly don't form a uniform blob. To me, the structure looks somewhat like a butterfly. I don't know why this occurs, but it certainly suggests that something worth investigating is going on. Perhaps the choice of samples that were measured is funny; perhaps the measuring instrument doesn't make certain kinds of measurement; or perhaps there are physical processes that prevent the data from spreading out over the space.

Our algorithm has one really interesting property. In some cases, we do not actually know the datapoints as vectors. Instead, we *just* know distances between the datapoints. This happens often in the social sciences, but there are important cases in computer science as well. As a rather contrived example, one could survey people about breakfast foods (say, eggs, bacon, cereal, oatmeal, pancakes, toast, muffins, kippers and sausages for a total of 9 items). We ask each person to rate the

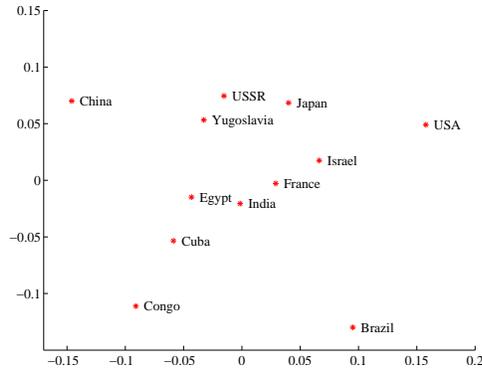


FIGURE 10.25: A map of country similarity, prepared from the data of figure 10.1. The map is often interpreted as showing a variation in development or wealth (poorest at bottom left to richest at top right); and freedom (most repressed at top left and freest at bottom right). I haven't plotted these axes, because the interpretation wouldn't be consistent with current intuition (the similarity data is forty years old, and quite a lot has happened in that time).

similarity of each pair of distinct items on some scale. We advise people that similar items are ones where, if they were offered both, they would have no particular preference; but, for dissimilar items, they would have a strong preference for one over the other. The scale might be “very similar”, “quite similar”, “similar”, “quite dissimilar”, and “very dissimilar” (scales like this are often called **Likert scales**). We collect these similarities from many people for each pair of distinct items, and then average the similarity over all respondents. We compute distances from the similarities in a way that makes very similar items close and very dissimilar items distant. Now we have a table of distances between items, and can compute a  $\mathbf{V}$  and produce a scatter plot. This plot is quite revealing, because items that most people think are easily substituted appear close together, and items that are hard to substitute are far apart. The neat trick here is that we did not start with a  $\mathcal{X}$ , but with just a set of distances; but we were able to associate a vector with “eggs”, and produce a meaningful plot.

Table 10.1 shows data from one such example. Students were interviewed (in 1971! things may have changed since then) about their perceptions of the similarity of countries. The averaged perceived similarity is shown in table 10.1. Large numbers reflect high similarity, so we can't use these numbers directly. It is reasonable to turn these numbers into distances by (a) using 0 as the distance between a country and itself and (b) using  $e^{-s_{ij}}$  as the distance between countries  $i$  and  $j$  (where  $s_{ij}$  is the similarity between them). Once we have distances, we can apply the procedure of section 10.4.1 to get points, then plot a scatter plot (Figure 10.25).

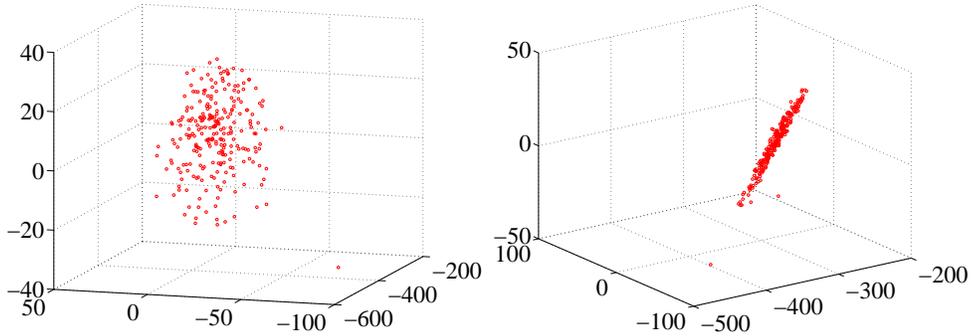


FIGURE 10.26: Two views of a multidimensional scaling to three dimensions of the height-weight dataset. Notice how the data seems to lie in a flat structure in 3D, with one outlying data point. This means that the distances between data points can be (largely) explained by a 2D representation.

10.5 EXAMPLE: UNDERSTANDING HEIGHT AND WEIGHT

Recall the height-weight data set of section ?? (from <http://www2.stetson.edu/~jrasp/data.htm>; look for bodyfat.xls at that URL). This is, in fact, a 16-dimensional dataset. The entries are (in this order): *bodyfat*; *density*; *age*; *weight*; *height*; *adiposity*; *neck*; *chest*; *abdomen*; *hip*; *thigh*; *knee*; *ankle*; *biceps*; *forearm*; *wrist*. We know already that many of these entries are correlated, but it’s hard to grasp a 16 dimensional dataset in one go. The first step is to investigate with a multidimensional scaling.

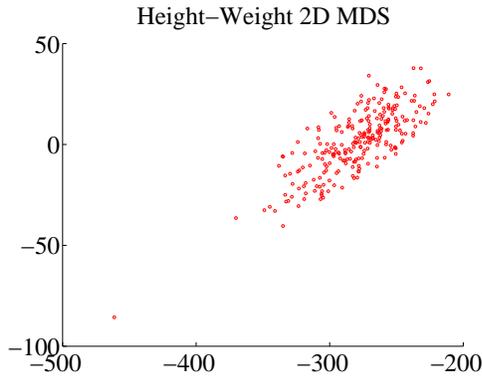


FIGURE 10.27: A multidimensional scaling to two dimensions of the height-weight dataset. One data point is clearly special, and another looks pretty special. The data seems to form a blob, with one axis quite a lot more important than another.

Figure ?? shows a multidimensional scaling of this dataset down to three dimensions. The dataset seems to lie on a (fairly) flat structure in 3D, meaning that inter-point distances are relatively well explained by a 2D representation. Two points seem to be special, and lie far away from the flat structure. The structure

isn't perfectly flat, so there will be small errors in a 2D representation; but it's clear that a lot of dimensions are redundant. Figure 10.27 shows a 2D representation of these points. They form a blob that is stretched along one axis, and there is no sign of multiple blobs. There's still at least one special point, which we shall ignore but might be worth investigating further. The distortions involved in squashing this dataset down to 2D seem to have made the second special point less obvious than it was in figure ??.

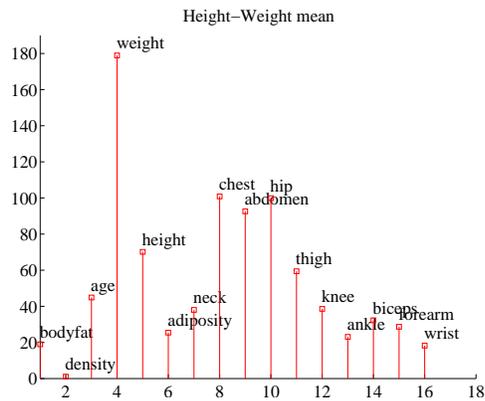


FIGURE 10.28: *The mean of the bodyfat.xls dataset. Each component is likely in a different unit (though I don't know the units), making it difficult to plot the data without being misleading. I've adopted one solution here, by plotting each component as a vertical bar, and labelling the bar. You shouldn't try to compare the values to one another. Instead, think of this plot as a compact version of a table.*

The next step is to try a principal component analysis. Figure 10.28 shows the mean of the dataset. The components of the dataset have different units, and shouldn't really be compared. But it is difficult to interpret a table of 16 numbers, so I have plotted the mean by showing a vertical bar for each component. Figure 10.29 shows the eigenvalues of the covariance for this dataset. Notice how one dimension is very important, and after the third principal component, the contributions become small. Of course, I could have said “fourth”, or “fifth”, or whatever — the precise choice depends on how small a number you think is “small”.

Figure 10.29 also shows the first principal component. The eigenvalues justify thinking of each data item as (roughly) the mean plus some weight times this principal component. From this plot you can see that data items with a larger value of *weight* will also have larger values of most other measurements, except *age* and *density*. You can also see how much larger; if the weight goes up by 8.5 units, then the abdomen will go up by 3 units, and so on. This explains the main variation in the dataset.

In the rotated coordinate system, the components are not correlated, and they have different variances (which are the eigenvalues of the covariance matrix). You can get some sense of the data by adding these variances; in this case, we get 1404. This means that, in the translated and rotated coordinate system, the average data

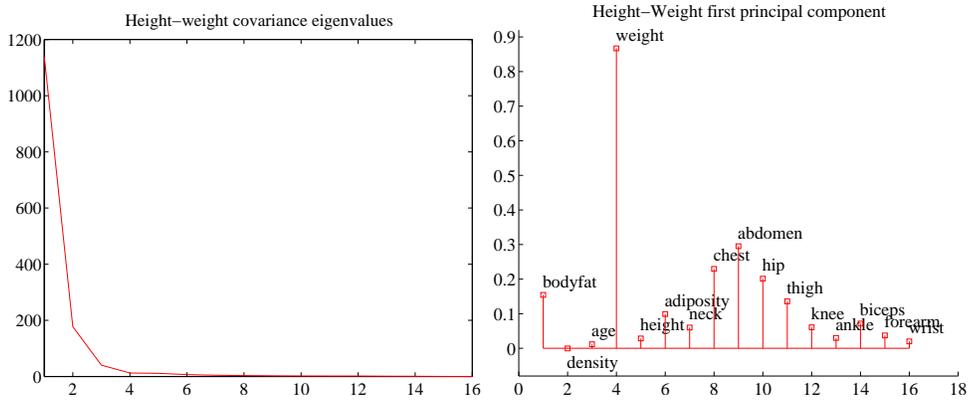


FIGURE 10.29: On the **left**, the eigenvalues of the covariance matrix for the bodyfat data set. Notice how fast the eigenvalues fall off; this means that most principal components have very small variance, so that data can be represented well with a small number of principal components. On the **right**, the first principal component for this dataset, plotted using the same convention as for figure 10.28.

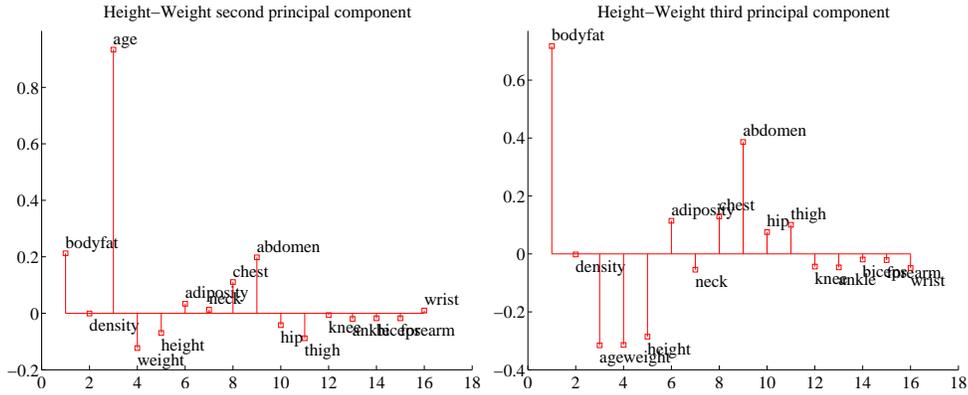


FIGURE 10.30: On the **left**, the second principal component, and on the **right** the third principal component of the height-weight dataset.

point is about  $37 = \sqrt{1404}$  units away from the center (the origin). Translations and rotations do not change distances, so the average data point is about 37 units from the center in the original dataset, too. If we represent a datapoint by using the mean and the first three principal components, there will be some error. We can estimate the average error from the component variances. In this case, the sum of the first three eigenvalues is 1357, so the mean square error in representing a datapoint by the first three principal components is  $\sqrt{(1404 - 1357)}$ , or 6.8. The relative error is  $6.8/37 = 0.18$ . Another way to represent this information, which is more widely used, is to say that the first three principal components explain all but  $(1404 - 1357)/1404 = 0.034$ , or 3.4% of the variance; notice that this is the square

of the relative error, which will be a much smaller number.

All this means that explaining a data point as the mean and the first three principal components produces relatively small errors. Figure 10.31 shows the second and third principal component of the data. These two principal components suggest some further conclusions. As *age* gets larger, *height* and *weight* get slightly smaller, but the weight is redistributed; *abdomen* gets larger, whereas *thigh* gets smaller. A smaller effect (the third principal component) links *bodyfat* and *abdomen*. As *bodyfat* goes up, so does *abdomen*.

10.6 YOU SHOULD

10.6.1 be able to:

- Create, plot and interpret the first few principal components of a dataset.
- Compute the error resulting from ignoring some principal components.

10.6.2 remember:

Definition: Covariance . . . . .	292
Definition: Covariance Matrix . . . . .	295
Useful facts: Whitening a dataset . . . . .	303
New term: smoothing . . . . .	304
New term: principal component . . . . .	309
New term: color constancy . . . . .	313
New term: principal coordinate analysis . . . . .	317
New term: multidimensional scaling . . . . .	317
New term: Likert scales . . . . .	319

## PROBLEMS

## Summaries

- 10.1.** You have a dataset  $\{\mathbf{x}\}$  of  $N$  vectors,  $\mathbf{x}_i$ , each of which is  $d$ -dimensional. We will consider a linear function of this dataset. Write  $\mathbf{a}$  for a constant vector; then the value of this linear function evaluated on the  $i$ 'th data item is  $\mathbf{a}^T \mathbf{x}_i$ . Write  $f_i = \mathbf{a}^T \mathbf{x}_i$ . We can make a new dataset  $\{f\}$  out of the values of this linear function.
- (a) Show that  $\text{mean}(\{f\}) = \mathbf{a}^T \text{mean}(\{\mathbf{x}\})$  (easy).
- (b) Show that  $\text{var}(\{f\}) = \mathbf{a}^T \text{Covmat}(\{\mathbf{x}\}) \mathbf{a}$  (harder, but just push it through the definition).
- (c) Assume the dataset has the special property that there exists some  $\mathbf{a}$  so that  $\mathbf{a}^T \text{Covmat}(\{\mathbf{x}\}) \mathbf{a}$ . Show that this means that the dataset lies on a hyperplane.

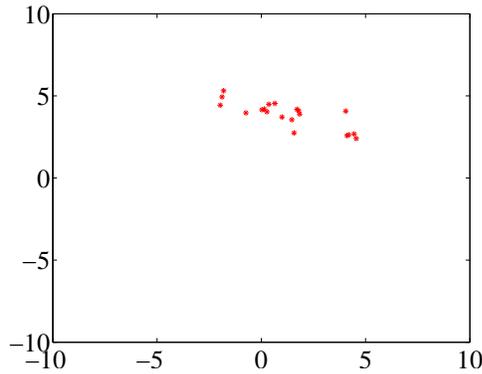


FIGURE 10.31: Figure for the question

- 10.2.** On Figure 10.31, mark the mean of the dataset, the first principal component, and the second principal component.
- 10.3.** You have a dataset  $\{\mathbf{x}\}$  of  $N$  vectors,  $\mathbf{x}_i$ , each of which is  $d$ -dimensional. Assume that  $\text{Covmat}(\{\mathbf{x}\})$  has one non-zero eigenvalue. Assume that  $\mathbf{x}_1$  and  $\mathbf{x}_2$  do not have the same value.
- (a) Show that you can choose a set of  $t_i$  so that you can represent *every* data item  $\mathbf{x}_i$  *exactly*
- $$\mathbf{x}_i = \mathbf{x}_1 + t_i(\mathbf{x}_2 - \mathbf{x}_1).$$
- (b) Now consider the dataset of these  $t$  values. What is the relationship between (a)  $\text{std}(t)$  and (b) the non-zero eigenvalue of  $\text{Covmat}(\{\mathbf{x}\})$ ? Why?

## PROGRAMMING EXERCISES

- 10.4.** Obtain the iris dataset from the UC Irvine machine learning data repository at <http://archive.ics.uci.edu/ml/machine-learning-databases/iris/iris.data>.
- (a) Plot a scatterplot matrix of this dataset, showing each species with a different marker. The fragment of R code in listing 10.1 should take you most of the way.

Listing 10.1: R code for iris example.

```

# r code for scatterplot of iris data
irisdat<-read.csv('iris.dat', header=FALSE);
library('lattice')
numiris=irisdat[,c(1, 2, 3, 4)]
postscript("irisscatterplot.eps")
# so that I get a postscript file
speciesnames<-c('setosa', 'versicolor', 'virginica')
pchr<-c(1, 2, 3)
colr<-c('red', 'green', 'blue', 'yellow', 'orange')
ss<-expand.grid(species=1:3)
parset<-with(ss, simpleTheme(pch=pchr[species],
                             col=colr[species]))
splom(irisdat[, c(1:4)], groups=irisdat$V5,
      par.settings=parset,
      varnames=c('Sepal\nLength', 'Sepal\nWidth',
                 'Petal\nLength', 'Petal\nWidth'),
      key=list(text=list(speciesnames),
               points=list(pch=pchr, columns=3))
dev.off()

```

- (b) Now obtain the first two principal components of the data. Plot the data on those two principal components alone, again showing each species with a different marker. Has this plot introduced significant distortions? Explain
- 10.5.** Take the wine dataset from the UC Irvine machine learning data repository at <http://archive.ics.uci.edu/ml/datasets/seeds>.
- (a) Plot the eigenvalues of the covariance matrix in sorted order. How many principal components should be used to represent this dataset? Why?
- (b) Construct a stem plot of each of the first 3 principal components (i.e. the eigenvectors of the covariance matrix with largest eigenvalues). What do you see?
- (c) Compute the first two principal components of this dataset, and project it onto those components. Now produce a scatter plot of this two dimensional dataset, where data items of class 1 are plotted as a '1', class 2 as a '2', and so on.
- 10.6.** Take the wheat kernel dataset from the UC Irvine machine learning data repository at <http://archive.ics.uci.edu/ml/datasets/seeds>. Compute the first two principal components of this dataset, and project it onto those components.
- (a) Produce a scatterplot of this projection. Do you see any interesting phenomena?
- (b) Plot the eigenvalues of the covariance matrix in sorted order. How many principal components should be used to represent this dataset? why?

	B	Co	Cu	E	F	In	Is	J	C	USSR	USA	Y
Brazil												
Congo	4.83											
Cuba	5.28	4.56										
Egypt	3.44	5.00	5.17									
France	4.72	4.00	4.11	4.78								
India	4.50	4.83	4.00	5.83	3.44							
Israel	3.83	3.33	3.61	4.67	4.00	4.11						
Japan	3.50	3.39	2.94	3.83	4.22	4.50	4.83					
China	2.39	4.00	5.50	4.39	3.67	4.11	3.00	4.17				
USSR	3.06	3.39	5.44	4.39	5.06	4.50	4.17	4.61	5.72			
USA	5.39	2.39	3.17	3.33	5.94	4.28	5.94	6.06	2.56	5.00		
Yugoslavia	3.17	3.50	5.11	4.28	4.72	4.00	4.44	4.28	5.06	6.67	3.56	

TABLE 10.1: *Similarity data for pairs of countries collected in 1971 by Wise (). Because similarity is symmetric, I show only one side of each pair.*

	C	K	M	N	Po	Pi	J	Bl	Bh
Cape Town	0	834	1136	1535	1547	1228	1263	912	846
Kimberley	834	0	332	716	713	555	430	150	526
Mahikeng	1136	332	0	537	441	625	242	367	797
Nelspruit	1535	716	537	0	234	464	306	622	891
Polokwane	1547	713	441	234	0	642	293	664	1015
Pietermaritzburg	1228	555	625	464	642	0	442	407	457
Johannesburg	1263	430	242	306	293	442	0	371	742
Bloemfontein	912	150	367	622	664	407	371	0	431
Bhisho	846	526	797	891	1015	457	742	431	0

TABLE 10.2: *Distances between South African provincial capitals, in kilometers, rounded to the nearest kilometer, and obtained from <http://www.distancefromto.net>.*

# Learning to Classify

A **classifier** is a procedure that accepts a set of features and produces a class label for them. There could be two, or many, classes. Classifiers are immensely useful, and find wide application, because many problems are naturally classification problems. For example, if you wish to determine whether to place an advert on a web-page or not, you would use a classifier (i.e. look at the page, and say yes or no according to some rule). As another example, if you have a program that you found for free on the web, you would use a classifier to decide whether it was safe to run it (i.e. look at the program, and say yes or no according to some rule). As yet another example, credit card companies must decide whether a transaction is good or fraudulent.

All these examples are two class classifiers, but in many cases it is natural to have more classes. You can think of sorting laundry as applying a multi-class classifier. You can think of doctors as complex multi-class classifiers. In this (crude) model, the doctor accepts a set of features, which might be your complaints, answers to questions, and so on, and then produces a response which we can describe as a class. The grading procedure for any class is a multi-class classifier: it accepts a set of features — performance in tests, homeworks, and so on — and produces a class label (the letter grade).

Classifiers are built by taking a set of labeled examples and using them to come up with a procedure that assigns a label to any new example. In the general problem, we have a training dataset  $(\mathbf{x}_i, y_i)$ ; each of the **feature vectors**  $\mathbf{x}_i$  consists of measurements of the properties of different types of object, and the  $y_i$  are labels giving the type of the object that generated the example. We will then use the training dataset to find a procedure that will predict an accurate label ( $y$ ) for any new object ( $\mathbf{x}$ ).

**Definition: 11.1** *Classifier*

A classifier is a procedure that accepts a feature vector and produces a label.

Indexing: Classifier

Indexing: Classifier

## 11.1 CLASSIFICATION, ERROR, AND LOSS

You should think of a classifier as a procedure — we pass in a feature vector, and get a class label in return. We want to use the training data to find the procedure

that is “best” when used on the test data. This problem has two tricky features. First, we need to be clear on what a good procedure is. Second, we really want the procedure to be good on test data, which we haven’t seen and won’t see; we only get to see the training data. These two considerations shape much of what we do.

### 11.1.1 What is “Best”?

The choice of procedure must depend on the cost of making a mistake. This cost can be represented with a **loss function**, which specifies the cost of making each type of mistake. For example, a two-class classifier can make two kinds of mistake. Because two-class classifiers are so common, there is a special name for each. A **false positive** occurs when a negative example is classified positive; a **false negative** occurs when a positive example is classified negative. By convention, the loss of getting the right answer is zero, and the loss for any wrong answer is non-negative.

The choice of procedure should depend quite strongly on the cost of each mistake. For example, pretend there is only one disease; then doctors would be classifiers, deciding whether a patient had it or not. If this disease is dangerous, but is safely and easily treated, false negatives are expensive errors, but false positives are cheap. In this case, procedures that tend to make more false positives than false negatives are better. Similarly, if the disease is not dangerous, but the treatment is difficult and unpleasant, then false positives are expensive errors and false negatives are cheap, and so we prefer false negatives to false positives.

You might argue that the best choice of classifier makes no mistake. But for most practical cases, the best choice of classifier is guaranteed to make mistakes. As an example, consider an alien who tries to classify humans into male and female, using only height as a feature. However the alien’s classifier uses that feature, it will make mistakes. This is because the classifier must choose, for each value of height, whether to label the humans with that height male or female. But for the vast majority of heights, there are some males and some females with that height, and so the alien’s classifier must make some mistakes whatever gender it chooses for that height.

For many practical problems, it is difficult to know what loss function to use. There is seldom an obvious choice. One common choice is to assume that all errors are equally bad. This gives the **0-1 loss** — every error has loss 1, and all right answers have loss zero.

### 11.1.2 What will happen on Test Data?

What we really want from a classifier is to have small loss on test data. But this is difficult to measure or achieve directly. For example, think about the case of classifying credit-card transactions as “good” or “bad”. We could certainly obtain a set of examples that have been labelled for training, because the card owner often complains some time after a fraudulent use of their card. But what is important here is to see a new transaction and label it without holding it up for a few months to see what the card owner says. The classifier may never know if the label is right or not.

Generally, we will assume that the training data is “like” the test data, and so we will try to make the classifier perform well on the training data. Classifiers

that have small training error might not have small test error. One example of this problem is the (silly) classifier that takes any data point and, if it is the same as a point in the training set, emits the class of that point and otherwise chooses randomly between the classes. This classifier has been learned from data, and has a zero error rate on the training dataset; it is likely to be unhelpful on any other dataset, however.

Test error is usually worse than training error, because of an effect that is sometimes called **overfitting**, so called because the classification procedure fits the training data better than it fits the test data. Other names include **selection bias**, because the training data has been selected and so isn't exactly like the test data, and **generalizing badly**, because the classifier fails to generalize. The effect occurs because the classifier has been trained to perform well *on the training dataset*, and the training dataset is not the same as the test dataset. First, it is quite likely smaller. Second, it might be biased through a variety of accidents. This means that small training error may have to do with quirks of the training dataset that don't occur in other sets of examples. One consequence of overfitting is that classifiers should always be evaluated on data that was not used in training.

**Remember this:** *Classifiers should always be evaluated on data that was not used in training.*

Now assume that we are using the 0-1 loss, so that the loss of using a classifier is the same as the **error rate**, that is, the percentage of classification attempts on a test set that result in the wrong answer. We cannot estimate the error rate of the classifier using training data, because the classifier has been trained to do well on that data, which will mean our error rate estimate will be too low. An alternative is to separate out some training data to form a **validation set** (confusingly, this is often called a **test set**), then train the classifier on the rest of the data, and evaluate on the validation set. This has the difficulty that the classifier will not be the best estimate possible, because we have left out some training data when we trained it. This issue can become a significant nuisance when we are trying to tell which of a set of classifiers to use—did the classifier perform poorly on validation data because it is not suited to the problem representation or because it was trained on too little data?

We can resolve this problem with **cross-validation**, which involves repeatedly: splitting data into training and validation sets uniformly and at random, training a classifier on the training set, evaluating it on the validation set, and then averaging the error over all splits. This allows an estimate of the likely future performance of a classifier, at the expense of substantial computation. You should notice that cross-validation, in some sense, looks at the sensitivity of the classifier to a change in the training set. The most usual form of this algorithm involves omitting single items from the dataset and is known as **leave-one-out cross-validation**.

## 11.2 THE SUPPORT VECTOR MACHINE

Assume we have a set of  $N$  example points  $\mathbf{x}_i$  that belong to two classes, which we indicate by 1 and  $-1$ . These points come with their class labels, which we write as  $y_i$ ; thus, our dataset can be written as

$$\{(\mathbf{x}_1, y_1), \dots, (\mathbf{x}_N, y_N)\}.$$

We wish to predict the sign of  $y$  for any point  $\mathbf{x}$ . We will use a linear classifier, so that for a new data item  $\mathbf{x}$ , we will predict

$$\text{sign}(\mathbf{a} \cdot \mathbf{x} + b)$$

and the particular classifier we use is given by our choice of  $\mathbf{a}$  and  $b$ .

You should think of  $\mathbf{a}$  and  $b$  as representing a hyperplane, given by the points where  $\mathbf{a} \cdot \mathbf{x} + b = 0$ . This hyperplane separates the positive data from the negative data, and is known as the **decision boundary**. Notice that the magnitude of  $\mathbf{a} \cdot \mathbf{x} + b$  grows as the point  $\mathbf{x}$  moves further away from the hyperplane.

**Example: 11.1** *A linear model with a single feature*

Assume we use a linear model with one feature. Then the model has the form  $y_i^{(p)} = \text{sign}(ax_i + b)$ . For any particular example which has the feature value  $x^*$ , this means we will test whether  $x^*$  is larger than, or smaller than,  $-b/a$ .

**Example: 11.2** *A linear model with two features*

Assume we use a linear model with two features. Then the model has the form  $y_i^{(p)} = \text{sign}(\mathbf{a}^T \mathbf{x}_i + b)$ . The sign changes along the line  $\mathbf{a}^T \mathbf{x} + b = 0$ . You should check that this is, indeed, a line. On one side of this line, the model makes positive predictions; on the other, negative. Which side is which can be swapped by multiplying  $\mathbf{a}$  and  $b$  by  $-1$ .

This family of classifiers may look bad to you, and it is easy to come up with examples that it misclassifies badly. In fact, the family is extremely strong. First, it is easy to estimate the best choice of rule for very large datasets. Second, linear classifiers have a long history of working very well in practice on real data. Third, linear classifiers are fast to evaluate.

In fact, examples that are classified badly by the linear rule usually are classified badly because there are too few features. Remember the case of the alien who classified humans into male and female by looking at their heights; if that alien

had looked at their chromosomes as well, the error rate would be extremely small. In practical examples, experience shows that the error rate of a poorly performing linear classifier can usually be improved by adding features to the vector  $\mathbf{x}$ .

### 11.2.1 Choosing a Classifier with the Hinge Loss

We will choose  $\mathbf{a}$  and  $b$  by choosing values that minimize a cost function. We will adopt a cost function of the form:

Training error cost + penalty term.

For the moment, we will ignore the penalty term and focus on the training error cost. Write

$$\gamma_i = \mathbf{a}^T \mathbf{x}_i + b$$

for the value that the linear function takes on example  $i$ . Write  $C(\gamma_i, y_i)$  for a function that compares  $\gamma_i$  with  $y_i$ . The training error cost will be of the form

$$(1/N) \sum_{i=1}^N C(\gamma_i, y_i).$$

A good choice of  $C$  should have some important properties. If  $\gamma_i$  and  $y_i$  have different signs, then  $C$  should be large, because the classifier will make the wrong prediction for this training example. Furthermore, if  $\gamma_i$  and  $y_i$  have different signs *and*  $\gamma_i$  has large magnitude, then the classifier will very likely make the wrong prediction for test examples that are close to  $\mathbf{x}_i$ . This is because the magnitude of  $(\mathbf{a} \cdot \mathbf{x} + b)$  grows as  $\mathbf{x}$  gets further from the decision boundary. So  $C$  should get larger as the magnitude of  $\gamma_i$  gets larger in this case.

If  $\gamma_i$  and  $y_i$  have the same signs, but  $\gamma_i$  has small magnitude, then the classifier will classify  $\mathbf{x}_i$  correctly, but might not classify points that are nearby correctly. This is because a small magnitude of  $\gamma_i$  means that  $\mathbf{x}_i$  is close to the decision boundary. So  $C$  should not be zero in this case. Finally, if  $\gamma_i$  and  $y_i$  have the same signs and  $\gamma_i$  has large magnitude, then  $C$  can be zero because  $\mathbf{x}_i$  is on the right side of the decision boundary and so are all the points near to  $\mathbf{x}_i$ .

The choice

$$C(y_i, \gamma_i) = \max(0, 1 - y_i \gamma_i)$$

has these properties. If  $y_i \gamma_i > 1$  (so the classifier predicts the sign correctly *and*  $\mathbf{x}_i$  is far from the boundary) there is no cost. But in any other case, there is a cost. The cost rises if  $\mathbf{x}_i$  moves toward the decision boundary from the correct side, and grows linearly as  $\mathbf{x}_i$  moves further away from the boundary on the wrong side (Figure 11.1). This means that minimizing the loss will encourage the classifier to (a) make strong positive (or negative) predictions for positive (or negative) examples and (b) for examples it gets wrong, make the most positive (negative) prediction that it can. This choice is known as the **hinge loss**.

Now we think about future examples. We don't know what their feature values will be, and we don't know their labels. But we do know that an example with feature vector  $\mathbf{x}$  will be classified with the rule  $\text{sign}((\mathbf{a} \cdot \mathbf{x} + b))$ . Imagine we classify this example wrongly. If  $\|\mathbf{a}\|$  is small, then this example is at least close

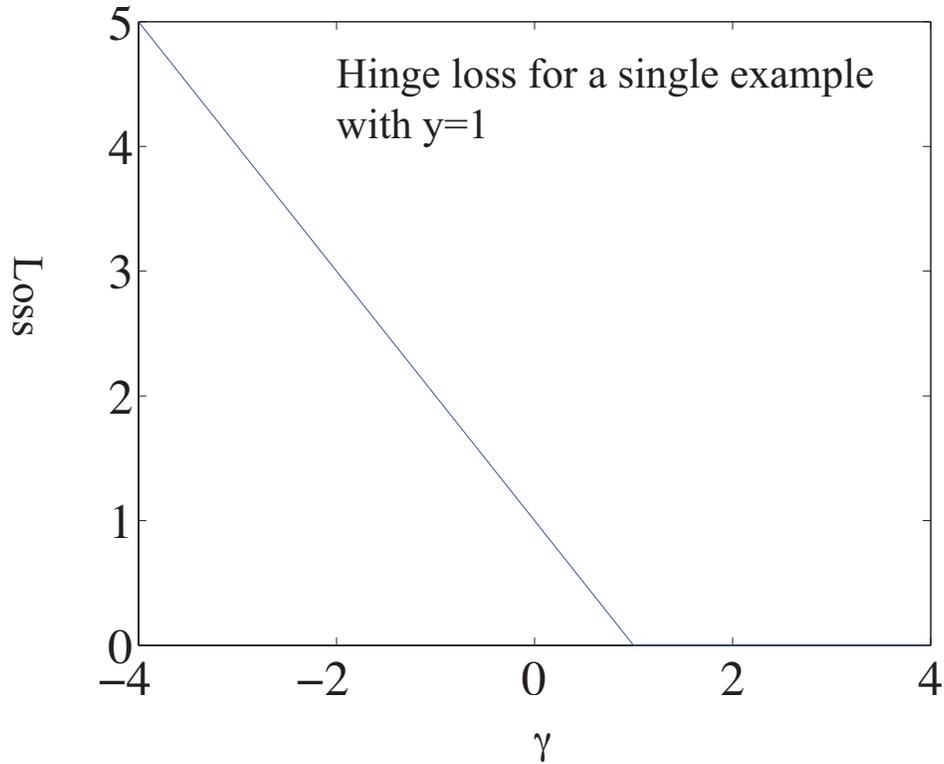


FIGURE 11.1: The hinge loss, plotted for the case  $y_i = 1$ . The horizontal variable is the  $\gamma_i = \mathbf{a} \cdot \mathbf{x}_i + b$  of the text. Notice that giving a strong negative response to this positive example causes a loss that grows linearly as the magnitude of the response grows. Notice also that giving an insufficiently positive response also causes a loss. Giving a strongly positive response is free.

to the decision boundary, and so examples nearby might be classified correctly. By this argument, we would like to achieve a small value of the hinge loss using a small value of  $\|\mathbf{a}\|$ .

The way to do this is to minimize

$$S(\mathbf{a}, b; \lambda) = \left[ (1/N) \sum_{i=1}^N \max(0, 1 - y_i (\mathbf{a} \cdot \mathbf{x}_i + b)) \right] + \frac{\lambda}{2} \mathbf{a}^T \mathbf{a}$$

where  $\lambda$  is some weight that balances the importance of a small hinge loss against the importance of a small  $\|\mathbf{a}\|$ . There are now two problems to solve. First, assume we know  $\lambda$ ; we will need to find  $\mathbf{a}$  and  $b$  that minimize  $S(\mathbf{a}, b; \lambda)$ . Second, we will need to estimate  $\lambda$ .

## 11.2.2 Finding a Minimum: General Points

I will first summarize general recipes for finding a minimum. Write  $\mathbf{u} = [\mathbf{a}, b]$  for the vector obtained by stacking the vector  $\mathbf{a}$  together with  $b$ . We have a function  $g(\mathbf{u})$ , and we wish to obtain a value of  $\mathbf{u}$  that achieves the minimum for that function. Sometimes we can solve this problem in closed form by constructing the gradient and finding a value of  $\mathbf{u}$  that makes the gradient zero. This happens mainly for specially chosen problems that occur in textbooks. For practical problems, we tend to need a numerical method.

Typical methods take a point  $\mathbf{u}^{(i)}$ , update it to  $\mathbf{u}^{(i+1)}$ , then check to see whether the result is a minimum. This process is started from a start point. The choice of start point may or may not matter for general problems, but for our problem it won't matter. The update is usually obtained by computing a direction  $\mathbf{p}^{(i)}$  such that for small values of  $h$ ,  $g(\mathbf{u}^{(i)} + h\mathbf{p}^{(i)})$  is smaller than  $g(\mathbf{u}^{(i)})$ . Such a direction is known as a **descent direction**. We must then determine how far to go along the descent direction, a process known as **line search**.

One method to choose a descent direction is **gradient descent**, which uses the negative gradient of the function. Recall our notation that

$$\mathbf{u} = \begin{pmatrix} u_1 \\ u_2 \\ \dots \\ u_d \end{pmatrix}$$

and that

$$\nabla g = \begin{pmatrix} \frac{\partial g}{\partial u_1} \\ \frac{\partial g}{\partial u_2} \\ \dots \\ \frac{\partial g}{\partial u_d} \end{pmatrix}.$$

We can write a Taylor series expansion for the function  $g(\mathbf{u}^{(i)} + h\mathbf{p}^{(i)})$ . We have that

$$g(\mathbf{u}^{(i)} + h\mathbf{p}^{(i)}) = g(\mathbf{u}^{(i)}) + h(\nabla g)^T \mathbf{p}^{(i)} + O(h^2)$$

This means that we can expect that if

$$\mathbf{p}^{(i)} = -\nabla g(\mathbf{u}^{(i)}),$$

we expect that, at least for small values of  $h$ ,  $g(\mathbf{u}^{(i)} + h\mathbf{p}^{(i)})$  will be less than  $g(\mathbf{u}^{(i)})$ . This works (as long as  $g$  is differentiable, and quite often when it isn't) because  $g$  must go down for at least small steps in this direction.

## 11.2.3 Finding a Minimum: Stochastic Gradient Descent

Assume we wish to minimize some function  $g(\mathbf{u}) = g_0(\mathbf{u}) + (1/N) \sum_{i=1}^N g_i(\mathbf{u})$ , as a function of  $\mathbf{u}$ . Gradient descent would require us to form

$$-\nabla g(\mathbf{u}) = -\left( \nabla g_0(\mathbf{u}) + (1/N) \sum_{i=1}^N \nabla g_i(\mathbf{u}) \right)$$

and then take a small step in this direction. But if  $N$  is large, this is unattractive, as we might have to sum a lot of terms. This happens a lot in building classifiers, where you might quite reasonably expect to deal with millions of examples. For some cases, there might be trillions of examples. Touching each example at each step really is impractical.

Instead, assume that, at each step, we choose a number  $k$  in the range  $1 \dots N$  uniformly and at random, and form

$$\mathbf{p}_k = -(\nabla g_0(\mathbf{u}) + \nabla g_k(\mathbf{u}))$$

and then take a small step along  $\mathbf{p}_k$ . Our new point becomes

$$\mathbf{a}^{(i+1)} = \mathbf{a}^{(i)} + \eta \mathbf{p}_k^{(i)},$$

where  $\eta$  is called the **steplength** (even though it very often isn't the length of the step we take!). It is easy to show that

$$\mathbb{E}[\mathbf{p}_k] = \nabla g(\mathbf{u})$$

(where the expectation is over the random choice of  $k$ ). This implies that if we take many small steps along  $\mathbf{p}_k$ , they should average out to a step backwards along the gradient. This approach is known as **stochastic gradient descent** (because we're not going along the gradient, but along a random vector which is the gradient only in expectation). It isn't obvious that stochastic gradient descent is a good idea. Although each step is easy to take, we may need to take more steps. The question is then whether we gain in the increased speed of the step what we lose by having to take more steps. Not much is known theoretically, but in practice the approach is hugely successful for training classifiers.

Choosing a steplength  $\eta$  takes some work. We can't search for the step that gives us the best value of  $g$ , because we don't want to evaluate the function  $g$  (doing so involves looking at each of the  $g_i$  terms). Instead, we use a steplength that is large at the start — so that the method can explore large changes in the values of the classifier parameters — and small steps later — so that it settles down. One useful strategy is to divide training into **epochs**. Each epoch is a block of a fixed number of iterations. Each iteration is one of the steps given above, with fixed steplength. However, the steplength changes from epoch to epoch. In particular, in the  $r$ 'th epoch, the steplength is

$$\eta^{(r)} = \frac{a}{r + b}$$

where  $a$  and  $b$  are constants chosen by experiment with small subsets of the dataset.

One cannot really test whether stochastic gradient descent has converged to the right answer. A better approach is to plot the error as a function of epoch on a validation set. This should vary randomly, but generally go down as the epochs proceed. I have summarized this discussion in box 11.1. You should be aware that the recipe there admits many useful variations, though.

**Procedure: 11.1** *Stochastic Gradient Descent*

We have a dataset containing  $N$  pairs  $(\mathbf{x}_i, y_i)$ . Each  $x_i$  is a  $d$ -dimensional feature vector, and each  $y_i$  is a label, either 1 or  $-1$ . Choose a set of possible values of the regularization weight  $\lambda$ . We wish to train a model that minimizes a cost function of the form  $g(\mathbf{u}) = \frac{\lambda}{2} \mathbf{u}^T \mathbf{u} + (\frac{1}{N}) \sum_{i=1}^N g_i(\mathbf{u})$ . Separate the data into three sets: test, training and validation. For each value of the regularization weight, train a model, and evaluate the model on validation data. Keep the model that produces the lowest error rate on the validation data, and report its performance on the test data.

Train a model by choosing a fixed number of epochs  $N_e$ , and the number of steps per epoch  $N_s$ . Choose a random start point,  $\mathbf{u}_0 = [\mathbf{a}, b]$ . For each epoch, first compute the steplength. In the  $e$ 'th epoch, the steplength is typically  $\eta = \frac{1}{ae+b}$  for constants  $a$  and  $b$  chosen by small-scale experiments (you try training a model with different values and see what happens). For the  $e$ 'th epoch, choose a subset of the training set for validation for that epoch. Now repeat until the model has been updated  $N_s$  times:

- Take  $k$  steps. Each step is taken by selecting a single data item uniformly and at random. Assume we select the  $i$ 'th data item. We then compute  $\mathbf{p} = -\nabla g_i(\mathbf{u}) - \lambda \mathbf{u}$ , and update the model by computing

$$\mathbf{u}_{n+1} = \mathbf{u}_n + \eta \mathbf{p}$$

- Evaluate the current model by computing the accuracy on the validation set for that epoch. Plot the accuracy as a function of step number.

## 11.2.4 Example: Training a Support Vector Machine with Stochastic Gradient Descent

I have summarized stochastic gradient descent in algorithm 11.1, but here is an example in more detail. We need to choose  $\mathbf{a}$  and  $b$  to minimize

$$C(\mathbf{a}, b) = (1/N) \sum_{i=1}^N \max(0, 1 - y_i (\mathbf{a} \cdot \mathbf{x}_i + b)) + \frac{\lambda}{2} \mathbf{a}^T \mathbf{a}.$$

This is a support vector machine, because it uses hinge loss. For a support vector machine, stochastic gradient descent is particularly easy. We have estimates  $\mathbf{a}^{(n)}$  and  $b^{(n)}$  of the classifier parameters, and we want to improve the estimates. We pick the  $k$ 'th example at random. We must now compute

$$\nabla \left( \max(0, 1 - y_k (\mathbf{a} \cdot \mathbf{x}_k + b)) + \frac{\lambda}{2} \mathbf{a}^T \mathbf{a} \right).$$

Assume that  $y_k(\mathbf{a} \cdot \mathbf{x}_k + b) > 1$ . In this case, the classifier predicts a score with the right sign, and a magnitude that is greater than one. Then the first term is zero, and the gradient of the second term is easy. Now if  $y_k(\mathbf{a} \cdot \mathbf{x}_k + b) < 1$ , we can ignore the max, and the first term is  $1 - y_k(\mathbf{a} \cdot \mathbf{x}_k + b)$ ; the gradient is again easy. But what if  $y_k(\mathbf{a} \cdot \mathbf{x}_k + b) = 1$ ? there are two distinct values we could choose for the gradient, because the max term isn't differentiable. It turns out not to matter which term we choose (Figure ??), so we can write the gradient as

$$p_k = \begin{cases} \begin{bmatrix} \lambda \mathbf{a} \\ 0 \end{bmatrix} & \text{if } y_k(\mathbf{a} \cdot \mathbf{x}_k + b) \geq 1 \\ \begin{bmatrix} \lambda \mathbf{a} - y_k \mathbf{x} \\ -y_k \end{bmatrix} & \text{otherwise} \end{cases}$$

We choose a steplength  $\eta$ , and update our estimates using this gradient. This yields:

$$\mathbf{a}^{(n+1)} = \mathbf{a}^{(n)} - \eta \begin{cases} \lambda \mathbf{a} & \text{if } y_k(\mathbf{a} \cdot \mathbf{x}_k + b) \geq 1 \\ \lambda \mathbf{a} - y_k \mathbf{x} & \text{otherwise} \end{cases}$$

and

$$b^{(n+1)} = b^{(n)} - \eta \begin{cases} 0 & \text{if } y_k(\mathbf{a} \cdot \mathbf{x}_k + b) \geq 1 \\ -y_k & \text{otherwise} \end{cases}.$$

To construct figures, I downloaded the dataset at <http://archive.ics.uci.edu/ml/datasets/Adult>. This dataset apparently contains 48,842 data items, but I worked with only the first 32,000. Each consists of a set of numeric and categorical features describing a person, together with whether their annual income is larger than or smaller than 50K\$. I ignored the categorical features to prepare these figures. This isn't wise if you want a good classifier, but it's fine for an example. I used these features to predict whether income is over or under 50K\$. I split the data into 5,000 test examples, and 27,000 training examples. It's important to do so at random. There are 6 numerical features. I subtracted the mean (which doesn't usually make much difference) and rescaled each so that the variance was 1 (which is often very important). I used two different training regimes.

In the first training regime, there were 100 epochs. In each epoch, I applied 426 steps. For each step, I selected one data item uniformly at random (sampling with replacement), then stepped down the gradient. This means the method sees a total of 42,600 data items. This means that there is a high probability it has touched each data item once (27,000 isn't enough, because we are sampling with replacement, so some items get seen more than once). I chose 5 different values for the regularization parameter and trained with a steplength of  $1/(0.01 * e + 50)$ , where  $e$  is the epoch. At the end of each epoch, I computed  $\mathbf{a}^T \mathbf{a}$  and the accuracy (fraction of examples correctly classified) of the current classifier on the held out test examples. Figure 11.2 shows the results. You should notice that the accuracy changes slightly each epoch; that for larger regularizer values  $\mathbf{a}^T \mathbf{a}$  is smaller; and that the accuracy settles down to about 0.8 very quickly.

In the second training regime, there were 100 epochs. In each epoch, I applied 50 steps. For each step, I selected one data item uniformly at random (sampling with replacement), then stepped down the gradient. This means the method sees

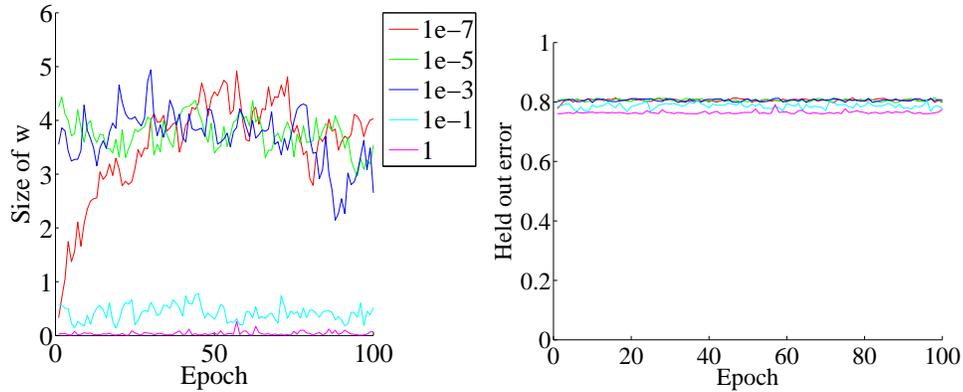


FIGURE 11.2: On the **left**, the magnitude of the weight vector  $\mathbf{a}$  at the end of each epoch for the first training regime described in the text. On the **right**, the accuracy on held out data at the end of each epoch. Notice how different choices of regularization parameter lead to different magnitudes of  $\mathbf{a}$ ; how the method isn't particularly sensitive to choice of regularization parameter (they change by factors of 100); how the accuracy settles down fairly quickly; and how overlarge values of the regularization parameter do lead to a loss of accuracy.

a total of 5,000 data items, and about 3, 216 unique data items — it hasn't seen the whole training set. I chose 5 different values for the regularization parameter and trained with a steplength of  $1/(0.01 * e + 50)$ , where  $e$  is the epoch. At the end of each epoch, I computed  $\mathbf{a}^T \mathbf{a}$  and the accuracy (fraction of examples correctly classified) of the current classifier on the held out test examples. Figure 11.3 shows the results. You should notice that the accuracy changes slightly each epoch; that for larger regularizer values  $\mathbf{a}^T \mathbf{a}$  is smaller; and that the accuracy settles down to about 0.8 very quickly; and that there isn't much difference between the two training regimes. All of these points are relatively typical of stochastic gradient descent with very large datasets.

**Remember this:** *Linear SVM's are a go-to classifier. When you have a binary classification problem, the first step should be to try a linear SVM. There is an immense quantity of good software available.*

### 11.2.5 Multi-Class Classifiers

I have shown how one trains a linear SVM to make a binary prediction (i.e. predict one of two outcomes). But what if there are three, or more, labels? In principle, you could write a binary code for each label, then use a different SVM to predict

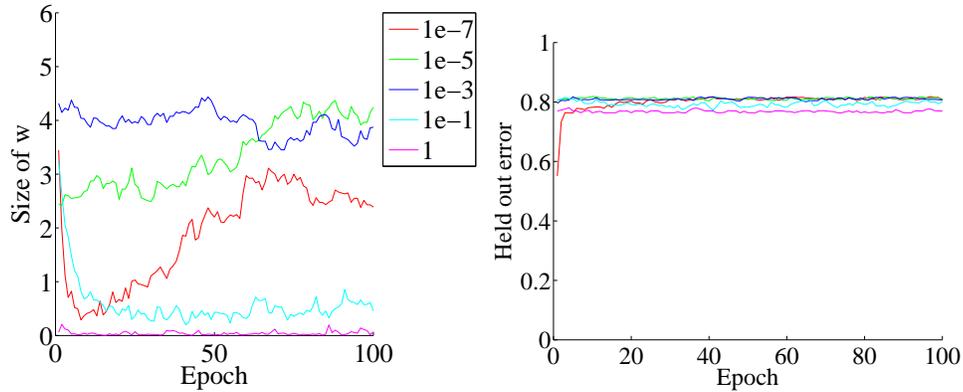


FIGURE 11.3: On the **left**, the magnitude of the weight vector  $\mathbf{a}$  at the end of each epoch for the second training regime described in the text. On the **right**, the accuracy on held out data at the end of each epoch. Notice how different choices of regularization parameter lead to different magnitudes of  $\mathbf{a}$ ; how the method isn't particularly sensitive to choice of regularization parameter (they change by factors of 100); how the accuracy settles down fairly quickly; and how overlarge values of the regularization parameter do lead to a loss of accuracy.

each bit of the code. It turns out that this doesn't work terribly well, because an error by one of the SVM's is usually catastrophic.

There are two methods that are widely used. In the **all-vs-all** approach, we train a binary classifier for each pair of classes. To classify an example, we present it to each of these classifiers. Each classifier decides which of two classes the example belongs to, then records a vote for that class. The example gets the class label with the most votes. This approach is simple, but scales very badly with the number of classes (you have to build  $O(N^2)$  different SVM's for  $N$  classes).

In the **one-vs-all** approach, we build a binary classifier for each class. This classifier must distinguish its class from all the other classes. We then take the class with the largest classifier score. One can think up quite good reasons this approach shouldn't work. For one thing, the classifier isn't told that you intend to use the score to tell similarity between classes. In practice, the approach works rather well and is quite widely used. This approach scales a bit better with the number of classes ( $O(N)$ ).

Evaluating a multi-class classifier is more complex than evaluating a binary classifier. A multi-class classifier can make many more kinds of mistake than a binary classifier can. It is useful to know the total error rate of the classifier, which is the percentage of classification attempts that produce the wrong answer. An alternative is the accuracy, which is the percentage of classification attempts that produce the *right* answer. If the error rate is low enough, or the accuracy is high enough, there's not much to worry about. But if it's not, you can look at the **class confusion matrix** to see what's going on.

Table 11.1 gives an example. This is a class confusion matrix from a classifier

	Predict 0	Predict 1	Predict 2	Predict 3	Predict 4	Class error
True 0	151	7	2	3	1	7.9%
True 1	32	5	9	9	0	91%
True 2	10	9	7	9	1	81%
True 3	6	13	9	5	2	86%
True 4	2	3	2	6	0	100%

TABLE 11.1: The class confusion matrix for a multiclass classifier. Further details about the dataset and this example appear in worked example 11.1.

built on a dataset where one tries to predict the degree of heart disease from a collection of physiological and physical measurements. There are five classes (0 . . . 4). The  $i, j$ 'th cell of the table shows the number of data points of true class  $i$  that were classified to have class  $j$ . As I find it hard to recall whether rows or columns represent true or predicted classes, I have marked this on the table. For each row, there is a class error rate, which is the percentage of data points of that class that were misclassified. The first thing to look at in a table like this is the diagonal; if the largest values appear there, then the classifier is working well. This clearly isn't what is happening for table 11.1. Instead, you can see that the method is very good at telling whether a data point is in class 0 or not (the class error rate is rather small), but cannot distinguish between the other classes. This is a strong hint that the data can't be used to draw the distinctions that we want. It might be a lot better to work with a different set of classes.

**Remember this:** *It is straightforward to build a multi-class classifier out of binary classifiers. Any decent SVM package will do this for you.*

### 11.3 CLASSIFYING WITH RANDOM FORESTS

I described a classifier as a rule that takes a feature, and produces a class. One way to build such a rule is with a sequence of simple tests, where each test is allowed to use the results of all previous tests. This class of rule can be drawn as a tree (Figure ??), where each node represents a test, and the edges represent the possible outcomes of the test. To classify a test item with such a tree, you present it to the first node; the outcome of the test determines which node it goes to next; and so on, until the example arrives at a leaf. When it does arrive at a leaf, we label the test item with the most common label in the leaf. This object is known as a **decision tree**. Notice one attractive feature of this decision tree: it deals with multiple class labels quite easily, because you just label the test item with the most common label in the leaf that it arrives at when you pass it down the tree.

Figure 11.5 shows a simple 2D dataset with four classes, next to a decision tree that will correctly classify at least the training data. Actually classifying data

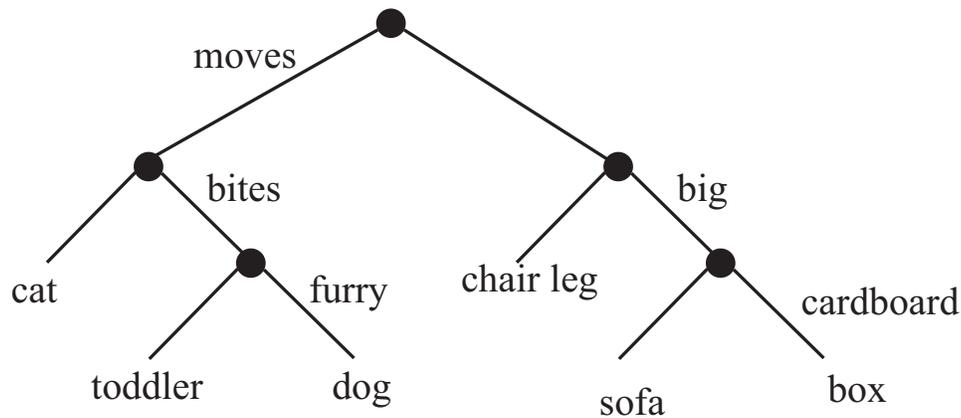


FIGURE 11.4: *This — the household robot’s guide to obstacles — is a typical decision tree. I have labelled only one of the outgoing branches, because the other is the negation. So if the obstacle moves, bites, but isn’t furry, then it’s a toddler. In general, an item is passed down the tree until it hits a leaf. It is then labelled with the leaf’s label.*

with a tree like this is straightforward. We take the data item, and pass it down the tree. Notice it can’t go both left and right, because of the way the tests work. This means each data item arrives at a single leaf. We take the most common label at the leaf, and give that to the test item. In turn, this means we can build a geometric structure on the feature space that corresponds to the decision tree. I have illustrated that structure in figure 11.5, where the first decision splits the feature space in half (which is why the term *split* is used so often), and then the next decisions split each of those halves into two.

The important question is how to get the tree from data. It turns out that the best approach for building a tree incorporates a great deal of randomness. As a result, we will get a different tree each time we train a tree on a dataset. None of the individual trees will be particularly good (they are often referred to as “weak learners”). The natural thing to do is to produce many such trees (a **decision forest**), and allow each to vote; the class that gets the most votes, wins. This strategy is extremely effective.

### 11.3.1 Building a Decision Tree

There are many algorithms for building decision trees. We will use an approach chosen for simplicity and effectiveness; be aware there are others. We will always use a binary tree, because it’s easier to describe and because that’s usual (it doesn’t change anything important, though). Each node has a **decision function**, which takes data items and returns either 1 or -1.

We train the tree by thinking about its effect on the training data. We pass the whole pool of training data into the root. Any node splits its incoming data into two pools, left (all the data that the decision function labels 1) and right (ditto,

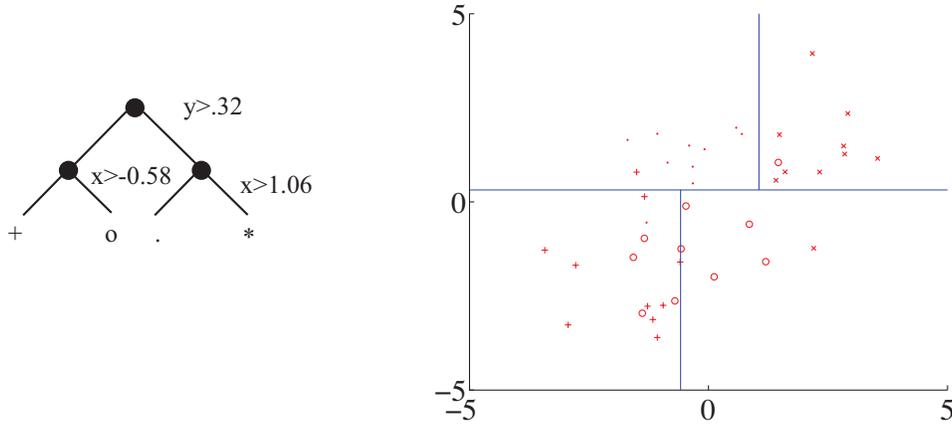


FIGURE 11.5: A straightforward decision tree, illustrated in two ways. On the **left**, I have given the rules at each split; on the **right**, I have shown the data points in two dimensions, and the structure that the tree produces in the feature space.

-1). Finally, each leaf contains a pool of data, which it can't split because it is a leaf.

Training the tree uses a straightforward algorithm. First, we choose a class of decision functions to use at each node. It turns out that a very effective algorithm is to choose a single feature at random, then test whether its value is larger than, or smaller than a threshold. For this approach to work, one needs to be quite careful about the choice of threshold, which is what we describe in the next section. Some minor adjustments, described below, are required if the feature chosen isn't ordinal. Surprisingly, being clever about the choice of *feature* doesn't seem add a great deal of value. We won't spend more time on other kinds of decision function, though there are lots.

Now assume we use a decision function as described, and we know how to choose a threshold. We start with the root node, then recursively either split the pool of data at that node, passing the left pool left and the right pool right, or stop splitting and return. Splitting involves choosing a decision function from the class to give the "best" split for a leaf. The main questions are how to choose the best split (next section), and when to stop.

Stopping is relatively straightforward. Quite simple strategies for stopping are very good. It is hard to choose a decision function with very little data, so we must stop splitting when there is too little data at a node. We can tell this is the case by testing the amount of data against a threshold, chosen by experiment. If all the data at a node belongs to a single class, there is no point in splitting. Finally, constructing a tree that is too deep tends to result in generalization problems, so we usually allow no more than a fixed depth  $D$  of splits. Choosing the best splitting threshold is more complicated.

Figure 11.6 shows two possible splits of a pool of training data. One is quite obviously a lot better than the other. In the good case, the split separates the pool

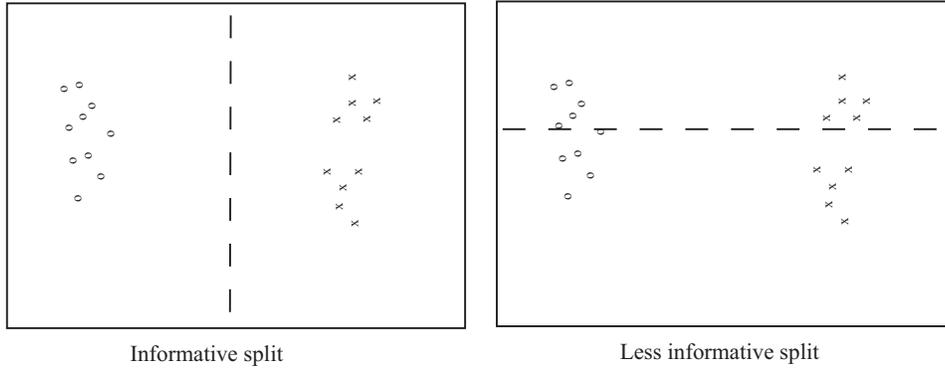


FIGURE 11.6: *Two possible splits of a pool of training data. Positive data is represented with an 'x', negative data with a 'o'. Notice that if we split this pool with the informative line, all the points on the left are 'o's, and all the points on the right are 'x's. This is an excellent choice of split — once we have arrived in a leaf, everything has the same label. Compare this with the less informative split. We started with a node that was half 'x' and half 'o', and now have two nodes each of which is half 'x' and half 'o' — this isn't an improvement, because we do not know more about the label as a result of the split.*

into positives and negatives. In the bad case, each side of the split has the same number of positives and negatives. We cannot usually produce splits as good as the good case here. What we are looking for is a split that will make the proper label more certain.

Figure 11.7 shows a more subtle case to illustrate this. The splits in this figure are obtained by testing the horizontal feature against a threshold. In one case, the left and the right pools contain about the same fraction of positive ('x') and negative ('o') examples. In the other, the left pool is all positive, and the right pool is mostly negative. This is the better choice of threshold. If we were to label any item on the left side positive and any item on the right side negative, the error rate would be fairly small. If you count, the best error rate for the informative split is 20% on the training data, and for the uninformative split it is 40% on the training data.

But we need some way to score the splits, so we can tell which threshold is best. Notice that, in the uninformative case, knowing that a data item is on the left (or the right) does not tell me much more about the data than I already knew. We have that  $p(1|\text{left pool, uninformative}) = 2/3 \approx 3/5 = p(1|\text{parent pool})$  and  $p(1|\text{right pool, uninformative}) = 1/2 \approx 3/5 = p(1|\text{parent pool})$ . For the informative pool, knowing a data item is on the left classifies it completely, and knowing that it is on the right allows us to classify it an error rate of  $1/3$ . The informative split means that my uncertainty about what class the data item belongs to is significantly reduced if I know whether it goes left or right. To choose a good threshold, we need to keep track of how informative the split is.

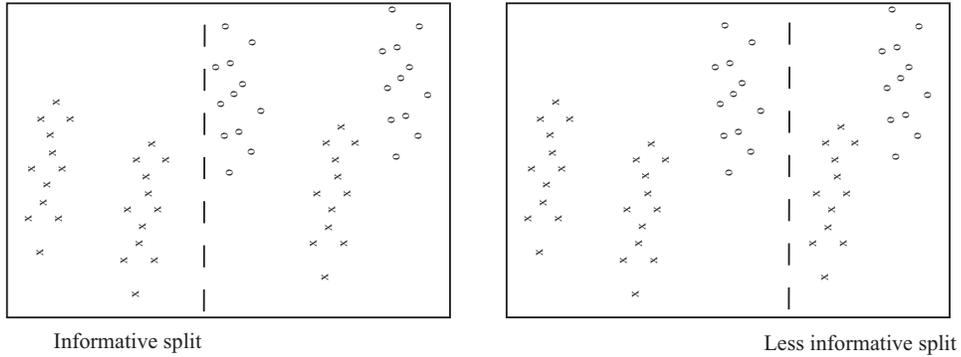


FIGURE 11.7: *Two possible splits of a pool of training data. Positive data is represented with an 'x', negative data with a 'o'. Notice that if we split this pool with the informative line, all the points on the left are 'x's, and two-thirds of the points on the right are 'o's. This means that knowing which side of the split a point lies would give us a good basis for estimating the label. In the less informative case, about two-thirds of the points on the left are 'x's and about half on the right are 'x's — knowing which side of the split a point lies is much less useful in deciding what the label is.*

### 11.3.2 Entropy and Information Gain

It turns out to be straightforward to keep track of information, in simple cases. We will start with an example. Assume I have 4 classes. There are 8 examples in class 1, 4 in class 2, 2 in class 3, and 2 in class 4. How much information *on average* will you need to send me to tell me the class of a given example? Clearly, this depends on how you communicate the information. You could send me the complete works of Edward Gibbon to communicate class 1; the Encyclopaedia for class 2; and so on. But this would be redundant. The question is how little can you send me. Keeping track of the amount of information is easier if we encode it with bits (i.e. you can send me sequences of '0's and '1's).

Imagine the following scheme. If an example is in class 1, you send me a '1'. If it is in class 2, you send me '01'; if it is in class 3, you send me '001'; and in class 4, you send me '101'. Then the expected number of bits you will send me is

$$p(\text{class} = 1)1 + p(2)2 + p(3)3 + p(4)3 = \frac{1}{2}1 + \frac{1}{4}2 + \frac{1}{8}3 + \frac{1}{8}3$$

which is 1.75 bits. This number doesn't have to be an integer, because it's an expectation.

Notice that for the  $i$ 'th class, you have sent me  $-\log_2 p(i)$  bits. We can write the expected number of bits you need to send me as

$$-\sum_i p(i) \log_2 p(i).$$

This expression handles other simple cases correctly, too. You should notice that it isn't really important *how many* objects appear in each class. Instead, the *fraction*

of all examples that appear in the class is what matters. This fraction is the prior probability that an item will belong to the class. You should try what happens if you have two classes, with an even number of examples in each; 256 classes, with an even number of examples in each; and 5 classes, with  $p(1) = 1/2$ ,  $p(2) = 1/4$ ,  $p(3) = 1/8$ ,  $p(4) = 1/16$  and  $p(5) = 1/16$ . If you try other examples, you may find it hard to construct a scheme where you can send as few bits *on average* as this expression predicts. It turns out that, in general, the smallest number of bits you will need to send me is given by the expression

$$-\sum_i p(i) \log_2 p(i)$$

under all conditions, though it may be hard or impossible to determine what representation is required to achieve this number.

The **entropy** of a probability distribution is a number that scores how many bits, on average, would need to be known to identify an item sampled from that probability distribution. For a discrete probability distribution, the entropy is computed as

$$-\sum_i p(i) \log_2 p(i)$$

where  $i$  ranges over all the numbers where  $p(i)$  is not zero. For example, if we have two classes and  $p(1) = 0.99$ , then the entropy is 0.0808, meaning you need very little information to tell which class an object belongs to. This makes sense, because there is a very high probability it belongs to class 1; you need very little information to tell you when it is in class 2. If you are worried by the prospect of having to send 0.0808 bits, remember this is an average, so you can interpret the number as meaning that, if you want to tell which class each of  $10^4$  independent objects belong to, you could do so in principle with only 808 bits.

Generally, the entropy is larger if the class of an item is more uncertain. Imagine we have two classes and  $p(1) = 0.5$ , then the entropy is 1, and this is the largest possible value for a probability distribution on two classes. You can always tell which of two classes an object belongs to with just one bit (though you might be able to tell with even less than one bit).

### 11.3.3 Entropy and Splits

Now we return to the splits. Write  $\mathcal{P}$  for the set of all data at the node. Write  $\mathcal{P}_l$  for the left pool, and  $\mathcal{P}_r$  for the right pool. The **entropy** of a pool  $\mathcal{C}$  that scores how many bits would be required to represent the class of an item in that pool, on average. Write  $n(i; \mathcal{C})$  for the number of items of class  $i$  in the pool, and  $N(\mathcal{C})$  for the number of items in the pool. Then the entropy is  $H(\mathcal{C})$  of the pool  $\mathcal{C}$  is

$$-\sum_i \frac{n(i; \mathcal{C})}{N(\mathcal{C})} \log_2 \frac{n(i; \mathcal{C})}{N(\mathcal{C})}.$$

It is straightforward that  $H(\mathcal{P})$  bits are required to classify an item in the parent pool  $\mathcal{P}$ . For an item in the left pool, we need  $H(\mathcal{P}_l)$  bits; for an item in the right

pool, we need  $H(\mathcal{P}_r)$  bits. If we split the parent pool, we expect to encounter items in the left pool with probability

$$\frac{N(\mathcal{P}_l)}{N(\mathcal{P})}$$

and items in the right pool with probability

$$\frac{N(\mathcal{P}_r)}{N(\mathcal{P})}.$$

This means that, on average, we must supply

$$\frac{N(\mathcal{P}_l)}{N(\mathcal{P})}H(\mathcal{P}_l) + \frac{N(\mathcal{P}_r)}{N(\mathcal{P})}H(\mathcal{P}_r)$$

bits to classify data items if we split the parent pool. Now a good split is one that results in left and right pools that are informative. In turn, we should need fewer bits to classify once we have split than we need before the split. You can see the difference

$$I(\mathcal{P}_l, \mathcal{P}_r; \mathcal{P}) = H(\mathcal{P}) - \left( \frac{N(\mathcal{P}_l)}{N(\mathcal{P})}H(\mathcal{P}_l) + \frac{N(\mathcal{P}_r)}{N(\mathcal{P})}H(\mathcal{P}_r) \right)$$

as the **information gain** caused by the split. This is the average number of bits that you *don't* have to supply if you know which side of the split an example lies. Better splits have larger information gain.

#### 11.3.4 Choosing a Split with Information Gain

Recall that our decision function is to choose a feature at random, then test its value against a threshold. Any data point where the value is larger goes to the left pool; where the value is smaller goes to the right. This may sound much too simple to work, but it is actually effective and popular. Assume that we are at a node, which we will label  $k$ . We have the pool of training examples that have reached that node. The  $i$ 'th example has a feature vector  $\mathbf{x}_i$ , and each of these feature vectors is a  $d$  dimensional vector.

We choose an integer  $j$  in the range  $1 \dots d$  uniformly and at random. We will split on this feature, and we store  $j$  in the node. Recall we write  $x_i^{(j)}$  for the value of the  $j$ 'th component of the  $i$ 'th feature vector. We will choose a threshold  $t_k$ , and split by testing the sign of  $x_i^{(j)} - t_k$ . Choosing the value of  $t_k$  is easy. Assume there are  $N_k$  examples in the pool. Then there are  $N_k - 1$  possible values of  $t_k$  that lead to different splits. To see this, sort the  $N_k$  examples by  $x_i^{(j)}$ , then choose values of  $t_k$  halfway between example values (Figure 11.8). For each of these values, we compute the information gain of the split. We then keep the threshold with the best information gain.

We can elaborate this procedure in a useful way, by choosing  $m$  features at random, finding the best split for each, then keeping the feature and threshold value that is best. It is important that  $m$  is a lot smaller than the total number of features — a usual root of thumb is that  $m$  is about the square root of the total

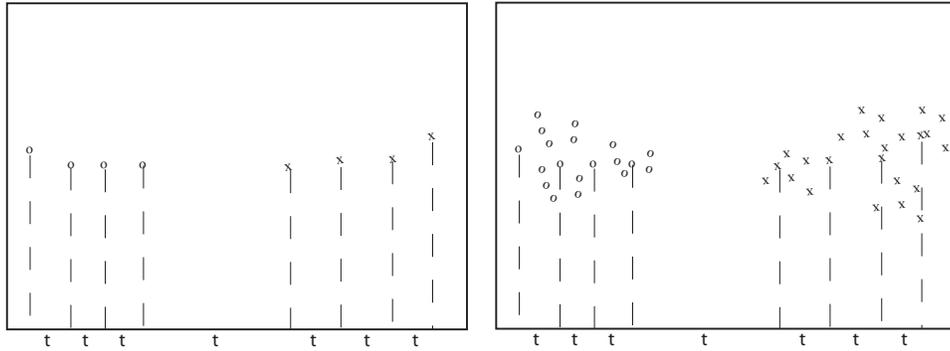


FIGURE 11.8: We search for a good splitting threshold by looking at values of the chosen component that yield different splits. On the **left**, I show a small dataset and its projection onto the chosen splitting component (the horizontal axis). For the 8 data points here, there are only 7 threshold values that produce interesting splits, and these are shown as 't's on the axis. On the **right**, I show a larger dataset; in this case, I have projected only a subset of the data, which results in a small set of thresholds to search.

number of features. It is usual to choose a single  $m$ , and choose that for all the splits.

Now assume we happen to have chosen to work with a feature that isn't ordinal, and so can't be tested against a threshold. A natural, and effective, strategy is as follows. We can split such a feature into two pools by flipping an unbiased coin for each value — if the coin comes up  $H$ , any data point with that value goes left, and if it comes up  $T$ , any data point with that value goes right. We chose this split at random, so it might not be any good. We can come up with a good split by repeating this procedure  $F$  times, computing the information gain for each split, then keeping the one that has the best information gain. We choose  $F$  in advance, and it usually depends on the number of values the categorical variable can take.

We now have a relatively straightforward blueprint for an algorithm, which I have put in a box. It's a blueprint, because there are a variety of ways in which it can be revised and changed.

**Procedure: 11.2** *Building a decision tree*

We have a dataset containing  $N$  pairs  $(\mathbf{x}_i, y_i)$ . Each  $x_i$  is a  $d$ -dimensional feature vector, and each  $y_i$  is a label. Call this dataset a **pool**. Now recursively apply the following procedure:

- If the pool is too small, or if all items in the pool have the same label, or if the depth of the recursion has reached a limit, stop.
- Otherwise, search the features for a good split that divides the pool into two, then apply this procedure to each child.

We search for a good split by the following procedure:

- Choose a subset of the feature components at random. Typically, one uses a subset whose size is about the square root of the feature dimension.
- For each component of this subset, search for the best splitting threshold. Do so by selecting a set of possible values for the threshold, then for each value splitting the dataset (every data item with a value of the component below the threshold goes left, others go right), and computing the information gain for the split. Keep the threshold that has the largest information gain.

A good set of possible values for the threshold will contain values that separate the data “reasonably”. If the pool of data is small, you can project the data onto the feature component (i.e. look at the values of that component alone), then choose the  $N - 1$  distinct values that lie between two data points. If it is big, you can randomly select a subset of the data, then project that subset on the feature component and choose from the values between data points.

## 11.3.5 Forests

A single decision tree tends to yield poor classifications. One reason is because the tree is not chosen to give the best classification of its training data. We used a random selection of splitting variables at each node, so the tree can’t be the “best possible”. Obtaining the best possible tree presents significant technical difficulties. It turns out that the tree that gives the best possible results on the training data can perform rather poorly on test data. The training data is a small subset of possible examples, and so must differ from the test data. The best possible tree on the training data might have a large number of small leaves, built using carefully chosen splits. But the choices that are best for training data might not be best for test data.

Rather than build the best possible tree, we have built a tree efficiently, but

with number of random choices. If we were to rebuild the tree, we would obtain a different result. This suggests the following extremely effective strategy: build many trees, and classify by merging their results.

### 11.3.6 Building and Evaluating a Decision Forest

There are two important strategies for building and evaluating decision forests. I am not aware of evidence strongly favoring one over the other, but different software packages use different strategies, and you should be aware of the options. In one strategy, we separate labelled data into a training and a test set. We then build multiple decision trees, training each using the whole training set. Finally, we evaluate the forest on the test set. In this approach, the forest has not seen some fraction of the available labelled data, because we used it to test. However, each tree has seen every training data item.

**Procedure: 11.3** *Building a decision forest*

We have a dataset containing  $N$  pairs  $(\mathbf{x}_i, y_i)$ . Each  $\mathbf{x}_i$  is a  $d$ -dimensional feature vector, and each  $y_i$  is a label. Separate the dataset into a test set and a training set. Train multiple distinct decision trees on the training set, recalling that the use of a random set of components to find a good split means you will obtain a distinct tree each time.

In the other strategy, sometimes called **bagging**, each time we train a tree we randomly subsample the labelled data with replacement, to yield a training set the same size as the original set of labelled data. Notice that there will be duplicates in this training set, which is like a bootstrap replicate. This training set is often called a **bag**. We keep a record of the examples that do not appear in the bag (the “out of bag” examples). Now to evaluate the forest, we evaluate each tree on its out of bag examples, and average these error terms. In this approach, the entire forest has seen all labelled data, and we also get an estimate of error, but no tree has seen all the training data.

**Procedure: 11.4** *Building a decision forest using bagging*

We have a dataset containing  $N$  pairs  $(\mathbf{x}_i, y_i)$ . Each  $\mathbf{x}_i$  is a  $d$ -dimensional feature vector, and each  $y_i$  is a label. Now build  $k$  bootstrap replicates of the training data set. Train one decision tree on each replicate.

### 11.3.7 Classifying Data Items with a Decision Forest

Once we have a forest, we must classify test data items. There are two major strategies. The simplest is to classify the item with each tree in the forest, then

take the class with the most votes. This is effective, but discounts some evidence that might be important. For example, imagine one of the trees in the forest has a leaf with many data items with the same class label; another tree has a leaf with exactly one data item in it. One might not want each leaf to have the same vote.

**Procedure: 11.5**    *Classification with a decision forest*

Given a test example  $\mathbf{x}$ , pass it down each tree of the forest. Now choose one of the following strategies.

- Each time the example arrives at a leaf, record one vote for the label that occurs most often at the leaf. Now choose the label with the most votes.
- Each time the example arrives at a leaf, record  $N_l$  votes for each of the labels that occur at the leaf, where  $N_l$  is the number of times the label appears in the training data at the leaf. Now choose the label with the most votes.

An alternative strategy that takes this observation into account is to pass the test data item down each tree. When it arrives at a leaf, we record one vote for each of the training data items in that leaf. The vote goes to the class of the training data item. Finally, we take the class with the most votes. This approach allows big, accurate leaves to dominate the voting process. Both strategies are in use, and I am not aware of compelling evidence that one is always better than the other. This may be because the randomness in the training process makes big, accurate leaves uncommon in practice.

**Worked example 11.1** *Classifying heart disease data*

Build a random forest classifier to classify the “heart” dataset from the UC Irvine machine learning repository. The dataset is at <http://archive.ics.uci.edu/ml/datasets/Heart+Disease>. There are several versions. You should look at the processed Cleveland data, which is in the file “processed.cleveland.data.txt”.

**Solution:** I used the R random forest package. This uses a bagging strategy. There is sample code in listing ???. This package makes it quite simple to fit a random forest, as you can see. In this dataset, variable 14 (V14) takes the value 0, 1, 2, 3 or 4 depending on the severity of the narrowing of the arteries. Other variables are physiological and physical measurements pertaining to the patient (read the details on the website). I tried to predict all five levels of variable 14, using the random forest as a multivariate classifier. This works rather poorly, as the out-of-bag class confusion matrix below shows. The total out-of-bag error rate was 45%.

	Predict 0	Predict 1	Predict 2	Predict 3	Predict 4	Class error
True 0	151	7	2	3	1	7.9%
True 1	32	5	9	9	0	91%
True 2	10	9	7	9	1	81%
True 3	6	13	9	5	2	86%
True 4	2	3	2	6	0	100%

This is the example of a class confusion matrix from table 11.1. Fairly clearly, one can predict narrowing or no narrowing from the features, but not the degree of narrowing (at least, not with a random forest). So it is natural to quantize variable 14 to two levels, 0 (meaning no narrowing), and 1 (meaning any narrowing, so the original value could have been 1, 2, or 3). I then built a random forest to predict this from the other variables. The total out-of-bag error rate was 19%, and I obtained the following out-of-bag class confusion matrix

	Predict 0	Predict 1	Class error
True 0	138	26	16%
True 1	31	108	22%

Notice that the false positive rate (16%, from 26/164) is rather better than the false negative rate (22%). Looking at these class confusion matrices, you might wonder whether it is better to predict 0, . . . , 4, then quantize. But this is not a particularly good idea. While the false positive rate is 7.9%, the false negative rate is much higher (36%, from 50/139). In this application, a false negative is likely more of a problem than a false positive, so the tradeoff is unattractive.

Listing 11.1: R code used for the random forests of worked example 11.1

```

setwd('/users/daf/Current/courses/Probcourse/Trees/RCode');
install.packages('randomForest')
library(randomForest)
heart<-read.csv('processed.cleveland.data.txt', header=FALSE)
heart$levels<-as.factor(heart$V14)
heartforest.allvals<-
  randomForest(formula=levels~V1+V2+V3+V4+V5+V6
               +V7+V8+V9+V10+V11+V12+V13,
               data=heart, type='classification', mtry=5)
# this fits to all levels
# I got the CCM by typing
heartforest.allvals
heart$yesno<-cut(heart$V14, c(-Inf, 0.1, Inf))
heartforest<-
  randomForest(formula=yesno~V1+V2+V3+V4+V5+V6
               +V7+V8+V9+V10+V11+V12+V13,
               data=heart, type='classification', mtry=5)
# this fits to the quantized case
# I got the CCM by typing
heartforest

```

**Remember this:** *Random forests are straightforward to build, and very effective. They can predict any kind of label. Good software implementations are easily available.*

## 11.4 CLASSIFYING WITH NEAREST NEIGHBORS

Generally we expect that, if two points  $\mathbf{x}_i$  and  $\mathbf{x}_j$  are close, then their labels will be the same most of the time. If we were unlucky, the two points could lie on either side of a decision boundary, but most pairs of nearby points will have the same labels.

This observation suggests the extremely useful and general strategy of exploiting a data item's neighbors. If you want to classify a data item, find the closest example, and report the class of that example. Alternatively, you could find the closest  $k$  examples, and vote. Imagine we have a data point  $\mathbf{x}$  that we wish to classify (a query point). Our strategy will be to find the closest training example, and report its class.

How well can we expect this strategy to work? Remember that any classifier will slice up the space of examples into cells (which might be quite complicated) where every point in a cell has the same label. The boundaries between cells are decision boundaries — when a point passes over the decision boundary, the label changes. Now assume we have a large number of labelled training examples, and

we know the best possible set of decision boundaries. If there are many training examples, there should be at least one training example that is close to the query point. If there are enough training examples, then the closest point should be inside the same cell as the query point.

You may be worried that, if the query point is close to a decision boundary, the closest point might be on the other side of that boundary. But if it were, we could improve things by simply having more training points. All this suggests that, *with enough training points*, our classifier should work about as well as the best possible classifier. This intuition turns out to be correct, though the number of training points required is wholly impractical, particularly for high-dimensional feature vectors.

One important generalization is to find the  $k$  nearest neighbors, then choose a label from those. A  $(k, l)$  nearest neighbor classifier finds the  $k$  example points closest to the point being considered, and classifies this point with the class that has the highest number of votes, as long as this class has more than  $l$  votes (otherwise, the point is classified as unknown). A  $(k, 0)$ -nearest neighbor classifier is usually known as a **k-nearest neighbor classifier**, and a  $(1, 0)$ -nearest neighbor classifier is usually known as a **nearest neighbor classifier**. In practice, one seldom uses more than three nearest neighbors. Finding the  $k$  nearest points for a particular query can be difficult, and Section ?? reviews this point.

There are three practical difficulties in building nearest neighbor classifiers. You need a lot of labelled examples. You need to be able to find the nearest neighbors for your query point. And you need to use a sensible choice of distance. For features that are obviously of the same type, such as lengths, the usual metric may be good enough. But what if one feature is a length, one is a color, and one is an angle? One possibility is to whiten the features (section 10.1). This may be hard if the dimension is so large that the covariance matrix is hard to estimate. It is almost always a good idea to scale each feature independently so that the variance of each feature is the same, or at least consistent; this prevents features with very large scales dominating those with very small scales. Notice that nearest neighbors (fairly obviously) doesn't like categorical data. If you can't give a clear account of how far apart two things are, you shouldn't be doing nearest neighbors. It is possible to fudge this point a little, by (say) putting together a distance between the levels of each factor, but it's probably unwise.

Nearest neighbors is wonderfully flexible about the labels the classifier predicts. Nothing changes when you go from a two-class classifier to a multi-class classifier.

**Worked example 11.2** *Classifying using nearest neighbors*

Build a nearest neighbor classifier to classify the digit data originally constructed by Yann Lecun. You can find it at several places. The original dataset is at <http://yann.lecun.com/exdb/mnist/>. The version I used was used for a Kaggle competition (so I didn't have to decompress Lecun's original format). I found it at <http://www.kaggle.com/c/digit-recognizer>.

**Solution:** As you'd expect, R has nearest neighbor code that seems quite good (I haven't had any real problems with it, at least). There isn't really all that much to say about the code. I used the R FNN package. This uses a bagging strategy. There is sample code in listing ???. I trained on 1000 of the 42000 examples, so you could see how in the code. I tested on the next 200 examples. For this (rather small) case, I found the following class confusion matrix

	P									
	0	1	2	3	4	5	6	7	8	9
0	12	0	0	0	0	0	0	0	0	0
1	0	20	4	1	0	1	0	2	2	1
2	0	0	20	1	0	0	0	0	0	0
3	0	0	0	12	0	0	0	0	4	0
4	0	0	0	0	18	0	0	0	1	1
5	0	0	0	0	0	19	0	0	1	0
6	1	0	0	0	0	0	18	0	0	0
7	0	0	1	0	0	0	0	19	0	2
8	0	0	1	0	0	0	0	0	16	0
9	0	0	0	2	3	1	0	1	1	14

There are no class error rates here, because I was in a rush and couldn't recall the magic line of R to get them. However, you can see the classifier works rather well for this case.

**Remember this:** *Nearest neighbor classifiers are often very effective. They can predict any kind of label. You do need to be careful to have enough data, and to have a meaningful distance function.*

## 11.5 YOU SHOULD

## 11.5.1 remember:

New term: classifier . . . . .	328
New term: feature vectors . . . . .	328
Definition: Classifier . . . . .	328
New term: overfitting . . . . .	330
New term: selection bias . . . . .	330
New term: generalizing badly . . . . .	330
Do not evaluate a classifier on training data. . . . .	330
New term: validation set . . . . .	330
New term: test set . . . . .	330
New term: decision boundary . . . . .	331
New term: descent direction . . . . .	334
New term: line search . . . . .	334
New term: gradient descent . . . . .	334
New term: steplength . . . . .	335
New term: stochastic gradient descent . . . . .	335
Linear SVM's are a go-to classifier. . . . .	338
New term: class confusion matrix . . . . .	339
Any SVM package should build a multi-class classifier for you. . . . .	340
New term: decision tree . . . . .	340
New term: decision forest . . . . .	341
New term: decision function . . . . .	341
New term: entropy . . . . .	345
New term: entropy . . . . .	345
New term: information gain . . . . .	346
New term: bagging . . . . .	349
New term: bag . . . . .	349
Random forests are good and easy. . . . .	352
Nearest neighbors are good and easy. . . . .	354

## PROBLEMS

## PROGRAMMING EXERCISES

- 11.1.** The UC Irvine machine learning data repository hosts a collection of data on breast cancer diagnostics, donated by Olvi Mangasarian, Nick Street, and William H. Wolberg. You can find this data at [http://archive.ics.uci.edu/ml/datasets/Breast+Cancer+Wisconsin+\(Diagnostic\)](http://archive.ics.uci.edu/ml/datasets/Breast+Cancer+Wisconsin+(Diagnostic)). For each record, there is an id number, 10 continuous variables, and a class (benign or malignant). There are 569 examples. Separate this dataset randomly into 100 validation, 100 test, and 369 training examples.

Write a program to train a support vector machine on this data using stochastic gradient descent. You should not use a package to train the classifier (you don't really need one), but your own code. You should ignore the id number, and use the continuous variables as a feature vector. You should search for an appropriate value of the regularization constant, trying at least the values  $\lambda = [1e - 3, 1e - 2, 1e - 1, 1]$ . Use the validation set for this search

You should use at least 50 epochs of at least 100 steps each. In each epoch, you should separate out 50 training examples at random for evaluation. You should compute the accuracy of the current classifier on the set held out for the epoch every 10 steps. You should produce:

- (a) A plot of the accuracy every 10 steps, for each value of the regularization constant.
  - (b) Your estimate of the best value of the regularization constant, together with a brief description of why you believe that is a good value.
  - (c) Your estimate of the accuracy of the best classifier on held out data
- 11.2.** The UC Irvine machine learning data repository hosts a collection of data on whether a mushroom is edible, donated by Jeff Schlimmer and to be found at <http://archive.ics.uci.edu/ml/datasets/Mushroom>. This data has a set of categorical attributes of the mushroom, together with two labels (poisonous or edible). Use the R random forest package (as in the example in the chapter) to build a random forest to classify a mushroom as edible or poisonous based on its attributes.
- (a) Produce a class-confusion matrix for this problem. If you eat a mushroom based on your classifier's prediction it is edible, what is the probability of being poisoned?
  - (b) (Optional) By playing with the many options in the package, can you get a more accurate classifier?

# Clustering: Models of High Dimensional Data

High-dimensional data comes with problems. Data points tend not to be where you think; they can scattered quite far apart, and can be quite far from the mean. Except in special cases, the only really reliable probability model is the Gaussian (or Gaussian blob, or blob).

There is an important rule of thumb for coping with high dimensional data: **Use simple models.** A blob is a good simple model. Modelling data as a blob involves computing its mean and its covariance. Sometimes, as we shall see, the covariance can be hard to compute. Even so, a blob model is really useful. It is natural to try and extend this model to cover datasets that don't obviously consist of a single blob.

One very good, very simple, model for high dimensional data is to assume that it consists of multiple blobs. To build models like this, we must determine (a) what the blob parameters are and (b) which datapoints belong to which blob. Generally, we will collect together data points that are close and form blobs out of them. This process is known as **clustering**.

Clustering is a somewhat puzzling activity. It is extremely useful to cluster data, and it seems to be quite important to do it reasonably well. But it surprisingly hard to give crisp criteria for a good (resp. bad) clustering of a dataset. Usually, clustering is part of building a model, and the main way to know that the clustering algorithm is bad is that the model is bad.

## 12.1 THE CURSE OF DIMENSION

High dimensional models display unintuitive behavior (or, rather, it can take years to make your intuition see the true behavior of high-dimensional models as natural). In these models, most data lies in places you don't expect. We will do several simple calculations with an easy high-dimensional distribution to build some intuition.

### 12.1.1 The Curse: Data isn't Where You Think it is

Assume our data lies within a cube, with edge length two, centered on the origin. This means that each component of  $\mathbf{x}_i$  lies in the range  $[-1, 1]$ . One simple model for such data is to assume that each dimension has uniform probability density in this range. In turn, this means that  $P(x) = \frac{1}{2^d}$ . The mean of this model is at the origin, which we write as  $\mathbf{0}$ .

The first surprising fact about high dimensional data is that most of the data can lie quite far away from the mean. For example, we can divide our dataset into two pieces.  $\mathcal{A}(\epsilon)$  consists of all data items where *every* component of the data has

a value in the range  $[-(1 - \epsilon), (1 - \epsilon)]$ .  $\mathcal{B}(\epsilon)$  consists of all the rest of the data. If you think of the data set as forming a cubical orange, then  $\mathcal{B}(\epsilon)$  is the rind (which has thickness  $\epsilon$ ) and  $\mathcal{A}(\epsilon)$  is the fruit.

Your intuition will tell you that there is more fruit than rind. This is true, for three dimensional oranges, but not true in high dimensions. The fact that the orange is cubical just simplifies the calculations, but has nothing to do with the real problem.

We can compute  $P(\{\mathbf{x} \in \mathcal{A}(\epsilon)\})$  and  $P(\{\mathbf{x} \in \mathcal{B}(\epsilon)\})$ . These probabilities tell us the probability a data item lies in the fruit (resp. rind).  $P(\{\mathbf{x} \in \mathcal{A}(\epsilon)\})$  is easy to compute as

$$P(\{\mathbf{x} \in \mathcal{A}(\epsilon)\}) = (2(1 - \epsilon))^d \left(\frac{1}{2^d}\right) = (1 - \epsilon)^d$$

and

$$P(\{\mathbf{x} \in \mathcal{B}(\epsilon)\}) = 1 - P(\{\mathbf{x} \in \mathcal{A}(\epsilon)\}) = 1 - (1 - \epsilon)^d.$$

But notice that, as  $d \rightarrow \infty$ ,

$$P(\{\mathbf{x} \in \mathcal{A}(\epsilon)\}) \rightarrow 0.$$

This means that, for large  $d$ , we expect most of the data to be in  $\mathcal{B}(\epsilon)$ . Equivalently, for large  $d$ , we expect that at least one component of each data item is close to either 1 or  $-1$ .

This suggests (correctly) that much data is quite far from the origin. It is easy to compute the average of the squared distance of data from the origin. We want

$$\mathbb{E}[\mathbf{x}^T \mathbf{x}] = \int_{\text{box}} \left( \sum_i x_i^2 \right) P(\mathbf{x}) d\mathbf{x}$$

but we can rearrange, so that

$$\mathbb{E}[\mathbf{x}^T \mathbf{x}] = \sum_i \mathbb{E}[x_i^2] = \sum_i \int_{\text{box}} x_i^2 P(\mathbf{x}) d\mathbf{x}.$$

Now each component of  $\mathbf{x}$  is independent, so that  $P(\mathbf{x}) = P(x_1)P(x_2)\dots P(x_d)$ . Now we substitute, to get

$$\mathbb{E}[\mathbf{x}^T \mathbf{x}] = \sum_i \mathbb{E}[x_i^2] = \sum_i \int_{-1}^1 x_i^2 P(x_i) dx_i = \sum_i \frac{1}{2} \int_{-1}^1 x_i^2 dx_i = \frac{d}{3},$$

so as  $d$  gets bigger, most data points will be further and further from the origin. Worse, as  $d$  gets bigger, data points tend to get further and further from one another. We can see this by computing the average of the squared distance of data points from one another. Write  $\mathbf{u}$  for one data point and  $\mathbf{v}$ ; we can compute

$$\mathbb{E}[d(\mathbf{u}, \mathbf{v})^2] = \int_{\text{box}} \int_{\text{box}} \sum_i (u_i - v_i)^2 d\mathbf{u} d\mathbf{v} = \mathbb{E}[\mathbf{u}^T \mathbf{u}] + \mathbb{E}[\mathbf{v}^T \mathbf{v}] - \mathbb{E}[\mathbf{u}^T \mathbf{v}]$$

but since  $\mathbf{u}$  and  $\mathbf{v}$  are independent, we have  $\mathbb{E}[\mathbf{u}^T \mathbf{v}] = \mathbb{E}[\mathbf{u}]^T \mathbb{E}[\mathbf{v}] = 0$ . This yields

$$\mathbb{E}[d(\mathbf{u}, \mathbf{v})^2] = 2\frac{d}{3}.$$

This means that, for large  $d$ , we expect our data points to be quite far apart.

### 12.1.2 Minor Banes of Dimension

High dimensional data presents a variety of important practical nuisances which follow from the curse of dimension. It is hard to estimate covariance matrices, and it is hard to build histograms.

Covariance matrices are hard to work with for two reasons. The number of entries in the matrix grows as the square of the dimension, so the matrix can get big and so difficult to store. More important, the amount of data we need to get an accurate estimate of all the entries in the matrix grows fast. As we are estimating more numbers, we need more data to be confident that our estimates are reasonable. There are a variety of straightforward work-arounds for this effect. In some cases, we have so much data there is no need to worry. In other cases, we assume that the covariance matrix has a particular form, and just estimate those parameters. There are two strategies that are usual. In one, we assume that the covariance matrix is diagonal, and estimate only the diagonal entries. In the other, we assume that the covariance matrix is a scaled version of the identity, and just estimate this scale. You should see these strategies as acts of desperation, to be used only when computing the full covariance matrix seems to produce more problems than using these approaches.

It is difficult to build histogram representations for high dimensional data. The strategy of dividing the domain into boxes, then counting data into them, fails miserably because there are too many boxes. In the case of our cube, imagine we wish to divide each dimension in half (i.e. between  $[-1, 0]$  and between  $[0, 1]$ ). Then we must have  $2^d$  boxes. This presents two problems. First, we will have difficulty representing this number of boxes. Second, unless we are exceptionally lucky, most boxes must be empty because we will not have  $2^d$  data items.

However, one representation is extremely effective. We can represent data as a collection of **clusters** — coherent blobs of similar datapoints that could, under appropriate circumstances, be regarded as the same. We could then represent the dataset by, for example, the center of each cluster and the number of data items in each cluster. Since each cluster is a blob, we could also report the covariance of each cluster, if we can compute it.

## 12.2 THE MULTIVARIATE NORMAL DISTRIBUTION

All the nasty facts about high dimensional data, above, suggest that we need to use quite simple probability models. By far the most important model is the multivariate normal distribution, which is quite often known as the multivariate gaussian distribution. There are two sets of parameters in this model, the mean  $\mu$  and the covariance  $\Sigma$ . For a  $d$ -dimensional model, the mean is a  $d$ -dimensional column vector and the covariance is a  $d \times d$  dimensional matrix. The covariance is a symmetric matrix. For our definitions to be meaningful, the covariance matrix must be

positive definite.

The form of the distribution  $p(\mathbf{x}|\mu, \Sigma)$  is

$$p(\mathbf{x}|\mu, \Sigma) = \frac{1}{\sqrt{(2\pi)^d \det(\Sigma)}} \exp\left(-\frac{1}{2}(\mathbf{x} - \mu)^T \Sigma^{-1} (\mathbf{x} - \mu)\right).$$

The following facts explain the names of the parameters:

**Useful Facts: 12.1** *Parameters of a Multivariate Normal Distribution*

Assuming a multivariate normal distribution, we have

- $\mathbb{E}[\mathbf{x}] = \mu$ , meaning that the mean of the distribution is  $\mu$ .
- $\mathbb{E}[(\mathbf{x} - \mu)(\mathbf{x} - \mu)^T] = \Sigma$ , meaning that the entries in  $\Sigma$  represent covariances.

Indexing: Parameters of a Multivariate Normal Distribution

Indexing: Parameters of a Multivariate Normal Distribution

Assume I know have a dataset of items  $\mathbf{x}_i$ , where  $i$  runs from 1 to  $N$ , and we wish to model this data with a multivariate normal distribution. The maximum likelihood estimate of the mean,  $\hat{\mu}$ , is

$$\hat{\mu} = \frac{\sum_i \mathbf{x}_i}{N}$$

(which is quite easy to show). The maximum likelihood estimate of the covariance,  $\hat{\Sigma}$ , is

$$\hat{\Sigma} = \frac{\sum_i (\mathbf{x}_i - \hat{\mu})(\mathbf{x}_i - \hat{\mu})^T}{N}$$

(which is rather a nuisance to show, because you need to know how to differentiate a determinant). These facts mean that we already know most of what is interesting about multivariate normal distributions (or gaussians).

### 12.2.1 Affine Transformations and Gaussians

Gaussians behave very well under affine transformations. In fact, we've already worked out all the math. Assume I have a dataset  $\mathbf{x}_i$ . The mean of the maximum likelihood gaussian model is  $\text{mean}(\{\mathbf{x}_i\})$ , and the covariance is  $\text{Covmat}(\{\mathbf{x}_i\})$ . I can now transform the data with an affine transformation, to get  $\mathbf{y}_i = \mathcal{A}\mathbf{x}_i + \mathbf{b}$ . The mean of the maximum likelihood gaussian model for the transformed dataset is  $\text{mean}(\{\mathbf{y}_i\})$ , and we've dealt with this; similarly, the covariance is  $\text{Covmat}(\{\mathbf{y}_i\})$ , and we've dealt with this, too.

A very important point follows in an obvious way. I can apply an affine transformation to any multivariate gaussian to obtain one with (a) zero mean and (b) independent components. In turn, this means that, *in the right coordinate system*, any gaussian is a product of zero mean one-dimensional normal distributions.

This fact is quite useful. For example, it means that simulating multivariate normal distributions is quite straightforward — you could simulate a standard normal distribution for each component, then apply an affine transformation.

### 12.2.2 Plotting a 2D Gaussian: Covariance Ellipses

There are some useful tricks for plotting a 2D Gaussian, which are worth knowing both because they're useful, and they help to understand Gaussians. Assume we are working in 2D; we have a Gaussian with mean  $\mu$  (which is a 2D vector), and covariance  $\Sigma$  (which is a 2x2 matrix). We could plot the collection of points  $\mathbf{x}$  that has some fixed value of  $p(\mathbf{x}|\mu, \Sigma)$ . This set of points is given by:

$$\frac{1}{2}((\mathbf{x} - \mu)^T \Sigma^{-1} (\mathbf{x} - \mu)) = c^2$$

where  $c$  is some constant. I will choose  $c^2 = \frac{1}{2}$ , because the choice doesn't matter, and this choice simplifies some algebra. You might recall that a set of points  $\mathbf{x}$  that satisfies a quadratic like this is a conic section. Because  $\Sigma$  (and so  $\Sigma^{-1}$ ) is positive definite, the curve is an ellipse. There is a useful relationship between the geometry of this ellipse and the Gaussian.

This ellipse — like all ellipses — has a major axis and a minor axis. These are at right angles, and meet at the center of the ellipse. We can determine the properties of the ellipse in terms of the Gaussian quite easily. The geometry of the ellipse isn't affected by rotation or translation, so we will translate the ellipse so that  $\mu = \mathbf{0}$  (i.e. the mean is at the origin) and rotate it so that  $\Sigma^{-1}$  is diagonal. Writing  $\mathbf{x} = [x, y]$  we get that the set of points on the ellipse satisfies

$$\frac{1}{2} \left( \frac{1}{k_1^2} x^2 + \frac{1}{k_2^2} y^2 \right) = \frac{1}{2}$$

where  $\frac{1}{k_1^2}$  and  $\frac{1}{k_2^2}$  are the diagonal elements of  $\Sigma^{-1}$ . We will assume that the ellipse has been rotated so that  $k_1 < k_2$ . The points  $(k_1, 0)$  and  $(-k_1, 0)$  lie on the ellipse, as do the points  $(0, k_2)$  and  $(0, -k_2)$ . The major axis of the ellipse, in this coordinate system, is the x-axis, and the minor axis is the y-axis. In this coordinate system,  $x$  and  $y$  are independent. If you do a little algebra, you will see that the standard deviation of  $x$  is  $\text{abs}(k_1)$  and the standard deviation of  $y$  is  $\text{abs}(k_2)$ . So the ellipse is longer in the direction of largest standard deviation and shorter in the direction of smallest standard deviation.

Now rotating the ellipse means we will pre- and post-multiply the covariance matrix with some rotation matrix. Translating it will move the origin to the mean. As a result, the ellipse has its center at the mean, its major axis is in the direction of the eigenvector of the covariance with largest eigenvalue, and its minor axis is in the direction of the eigenvector with smallest eigenvalue. A plot of this ellipse, which can be coaxed out of most programming environments with relatively little effort, gives us a great deal of information about the underlying Gaussian. These ellipses are known as **covariance ellipses**.

## 12.3 AGGLOMERATIVE AND DIVISIVE CLUSTERING

There are two natural algorithms for clustering. In **divisive clustering**, the entire data set is regarded as a cluster, and then clusters are recursively split to yield a good clustering (Algorithm 12.2). In **agglomerative clustering**, each data item is regarded as a cluster, and clusters are recursively merged to yield a good clustering (Algorithm 12.1).

```

Make each point a separate cluster
Until the clustering is satisfactory
    Merge the two clusters with the
        smallest inter-cluster distance
end

```

**Algorithm 12.1:** *Agglomerative Clustering or Clustering by Merging.*

```

Construct a single cluster containing all points
Until the clustering is satisfactory
    Split the cluster that yields the two
        components with the largest inter-cluster distance
end

```

**Algorithm 12.2:** *Divisive Clustering, or Clustering by Splitting.*

There are two major issues in thinking about clustering:

- *What is a good inter-cluster distance?* Agglomerative clustering uses an inter-cluster distance to fuse nearby clusters; divisive clustering uses it to split insufficiently coherent clusters. Even if a natural distance between data points is available (which might not be the case for vision problems), there is no canonical inter-cluster distance. Generally, one chooses a distance that seems appropriate for the data set. For example, one might choose the distance between the closest elements as the inter-cluster distance, which tends to yield extended clusters (statisticians call this method **single-link clustering**). Another natural choice is the maximum distance between an element of the first cluster and one of the second, which tends to yield rounded clusters (statisticians call this method **complete-link clustering**). Finally, one could use an average of distances between elements in the cluster, which also tends to yield “rounded” clusters (statisticians call this method **group average clustering**).
- *How many clusters are there?* This is an intrinsically difficult task if there is no model for the process that generated the clusters. The algorithms we have described generate a hierarchy of clusters. Usually, this hierarchy is displayed to a user in the form of a **dendrogram**—a representation of the

structure of the hierarchy of clusters that displays inter-cluster distances—and an appropriate choice of clusters is made from the dendrogram (see the example in Figure 12.1).

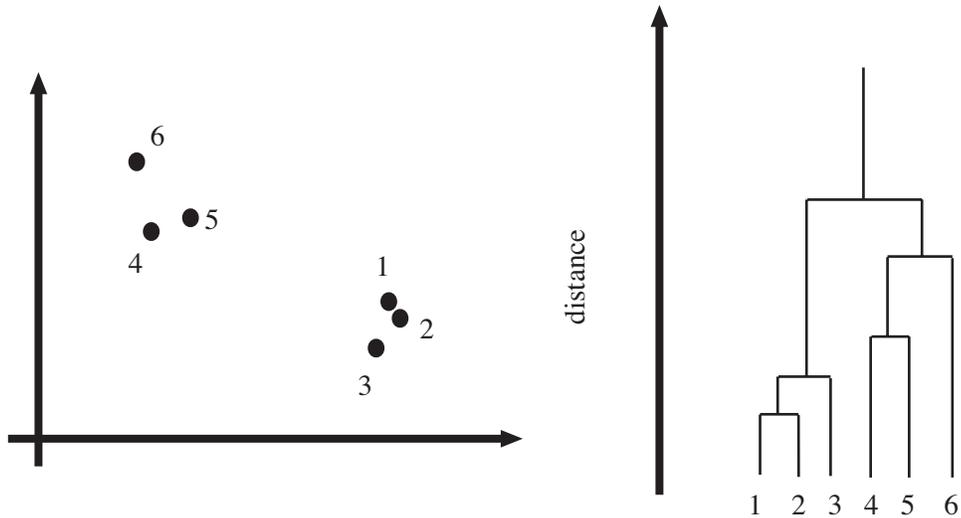


FIGURE 12.1: **Left**, a data set; **right**, a dendrogram obtained by agglomerative clustering using single-link clustering. If one selects a particular value of distance, then a horizontal line at that distance splits the dendrogram into clusters. This representation makes it possible to guess how many clusters there are and to get some insight into how good the clusters are.

The main difficulty in using a divisive model is knowing where to split. This is sometimes made easier for particular kinds of data. For example, we could segment an image by clustering pixel values. In this case, you can sometimes find good splits by constructing a histogram of intensities, or of color values.

Another important thing to notice about clustering from the example of figure 12.1 is that there is no right answer. There are a variety of different clusterings of the same data. For example, depending on what scales in that figure mean, it might be right to zoom out and regard all of the data as a single cluster, or to zoom in and regard each data point as a cluster. Each of these representations may be useful.

### 12.3.1 Clustering and Distance

In the algorithms above, and in what follows, we assume that the features are scaled so that distances (measured in the usual way) between data points are a good representation of their similarity. This is quite an important point. For example, imagine we are clustering data representing brick walls. The features might contain several distances: the spacing between the bricks, the length of the wall, the height of the wall, and so on. If these distances are given in the same set of units, we could have real trouble. For example, assume that the units are centimeters. Then the

spacing between bricks is of the order of one or two centimeters, but the heights of the walls will be in the hundreds of centimeters. In turn, this means that the distance between two datapoints is likely to be completely dominated by the height and length data. This could be what we want, but it might also not be a good thing.

There are some ways to manage this issue. One is to know what the features measure, and know how they should be scaled. Usually, this happens because you have a deep understanding of your data. If you don't (which happens!), then it is often a good idea to try and normalize the scale of the data set. There are two good strategies. The simplest is to translate the data so that it has zero mean (this is just for neatness - translation doesn't change distances), then scale each direction so that it has unit variance. More sophisticated is to translate the data so that it has zero mean, then transform it so that each direction is independent and has unit variance. Doing so is sometimes referred to as **decorrelation** or **whitening** (because you make the data more like white noise); I described how to do this in section 10.2.3.

**Worked example 12.1** *Agglomerative clustering in Matlab*

Cluster the height-weight dataset of <http://www2.stetson.edu/~jrasp/data.htm> (look for bodyfat.xls) using an agglomerative clusterer, and describe the results.

**Solution:** Matlab provides some tools that are useful for agglomerative clustering. These functions use a scheme where one first builds the whole tree of merges, then analyzes that tree to decide which clustering to report. `linkage` will determine which pairs of clusters should be merged at which step (there are arguments that allow you to choose what type of inter-cluster distance it should use); `dendrogram` will plot you a dendrogram; and `cluster` will extract the clusters from the linkage, using a variety of options for choosing the clusters. I used these functions to prepare the dendrogram of figure 12.2 for the height-weight dataset of section ?? (from <http://www2.stetson.edu/~jrasp/data.htm>; look for bodyfat.xls). I deliberately forced Matlab to plot the whole dendrogram, which accounts for the crowded look of the figure (you can allow it to merge small leaves, etc.). I used a single-link strategy. In particular, notice that many data points are about the same distance from one another, which suggests a single big cluster with a smaller set of nearby clusters. The clustering of figure 12.3 supports this view. I plotted the data points on the first two principal components, using different colors and shapes of marker to indicate different clusters. There are a total of 30 clusters here, though most are small.

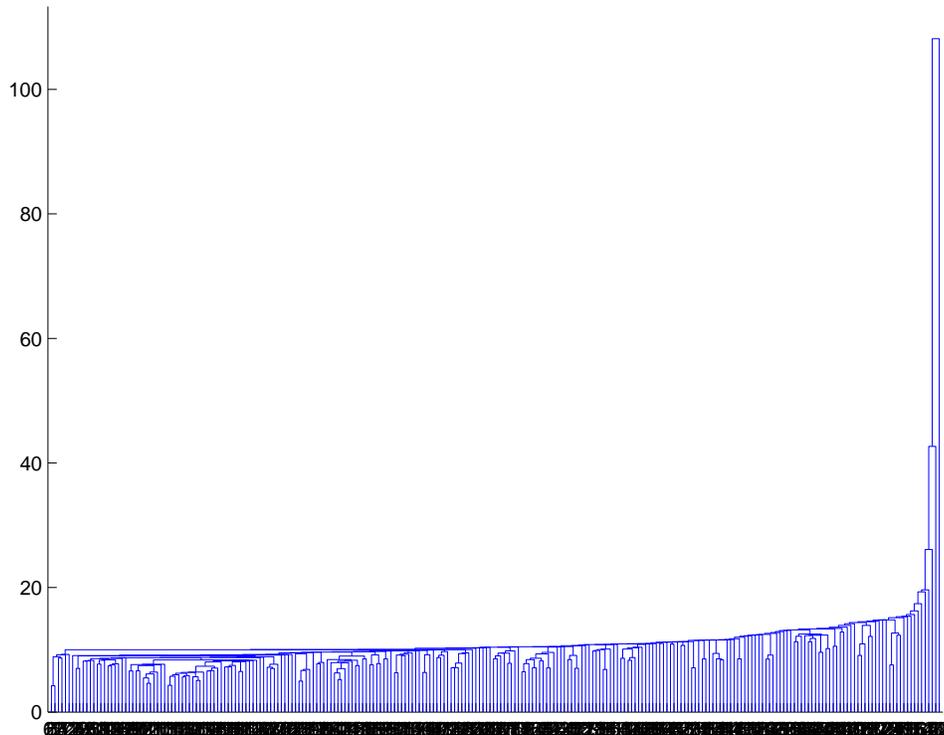


FIGURE 12.2: A dendrogram obtained from the body fat dataset, using single link clustering. Recall that the data points are on the horizontal axis, and that the vertical axis is distance; there is a horizontal line linking two clusters that get merged, established at the height at which they're merged. I have plotted the entire dendrogram, despite the fact it's a bit crowded at the bottom, because it shows that most data points are relatively close (i.e. there are lots of horizontal branches at about the same height).

**Worked example 12.2** *Agglomerative clustering in Matlab, 2*

Cluster the seed dataset from the UC Irvine Machine Learning Dataset Repository (you can find it at <http://archive.ics.uci.edu/ml/datasets/seeds>).

**Solution:** Each item consists of seven measurements of a wheat kernel; there are three types of wheat represented in this dataset. As you can see in figures 12.4 and 12.5, this data clusters rather well.



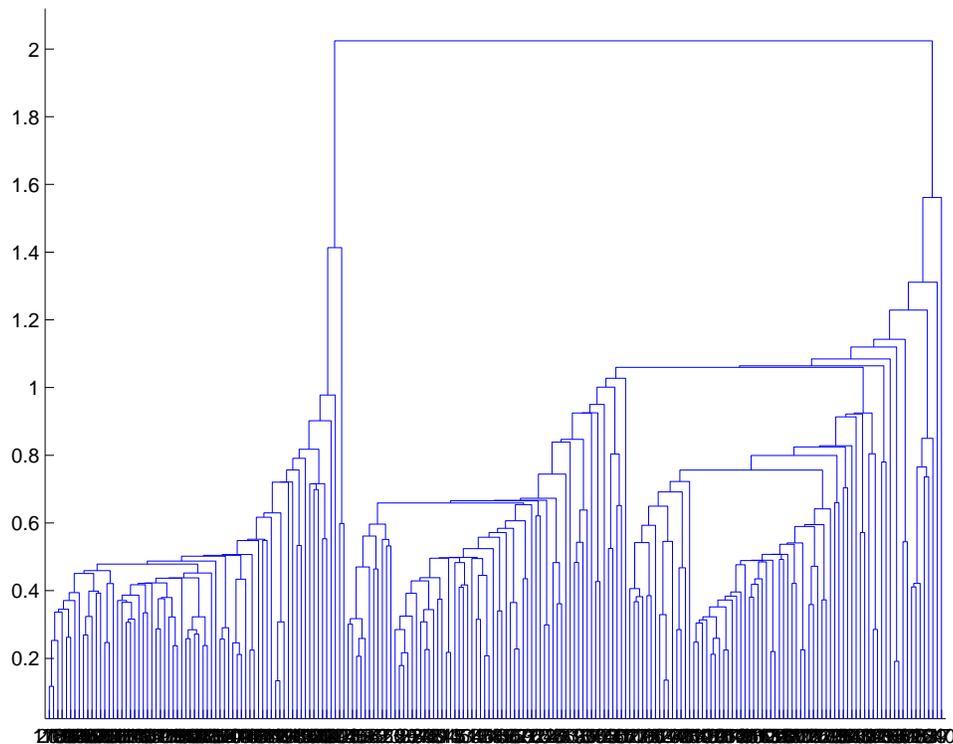


FIGURE 12.4: A dendrogram obtained from the seed dataset, using single link clustering. Recall that the data points are on the horizontal axis, and that the vertical axis is distance; there is a horizontal line linking two clusters that get merged, established at the height at which they're merged. I have plotted the entire dendrogram, despite the fact it's a bit crowded at the bottom, because you can now see how clearly the data set clusters into a small set of clusters — there are a small number of vertical “runs”.

## 12.4 THE K-MEANS ALGORITHM AND VARIANTS

Assume we have a dataset that, we believe, forms many clusters that look like blobs. If we knew where the center of each of the clusters was, it would be easy to tell which cluster each data item belonged to — it would belong to the cluster with the closest center. Similarly, if we knew which cluster each data item belonged to, it would be easy to tell where the cluster centers were — they'd be the mean of the data items in the cluster. This is the point closest to every point in the cluster.

We can turn these observations into an algorithm. Assume that we know how many clusters there are in the data, and write  $k$  for this number. The  $j$ th data item to be clustered is described by a feature vector  $\mathbf{x}_j$ . We write  $\mathbf{c}_i$  for the center of the  $i$ th cluster. We assume that most data items are close to the center of their

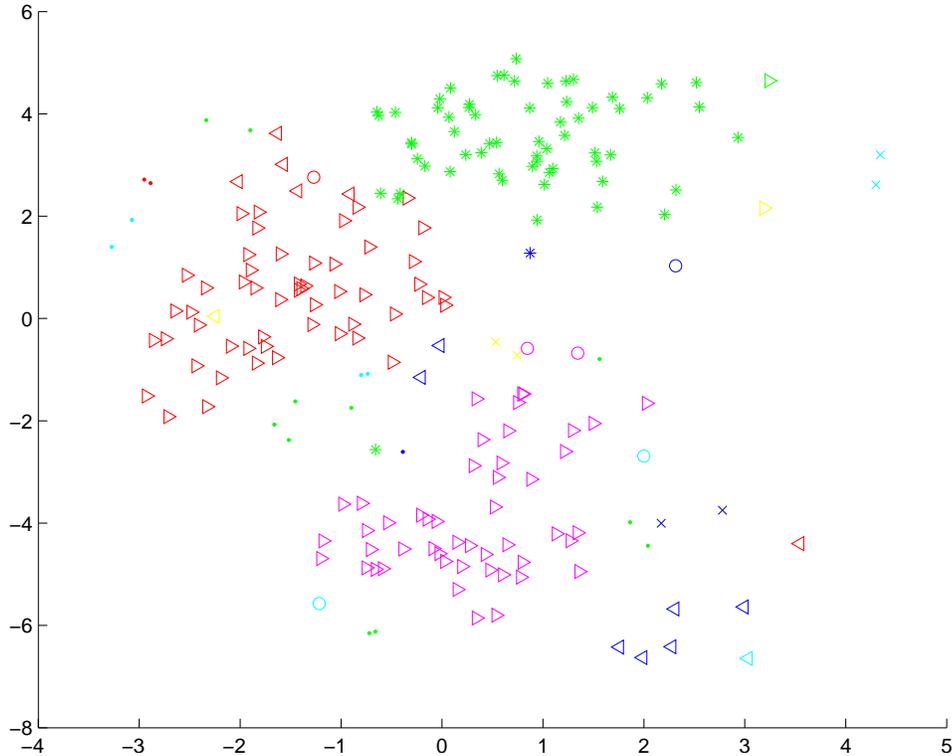


FIGURE 12.5: A clustering of the seed dataset, using agglomerative clustering, single link distance, and requiring a maximum of 30 clusters. I have plotted each cluster with a distinct marker (though some markers differ only by color; you might need to look at the PDF version to see this figure at its best). Notice that there are a set of fairly natural isolated clusters. The original data is 8 dimensional, which presents plotting problems; I show a scatter plot on the first two principal components (though I computed distances for clustering in the original 8 dimensional space).

cluster. This suggests that we cluster the data by minimizing the the cost function

$$\Phi(\text{clusters, data}) = \sum_{i \in \text{clusters}} \left\{ \sum_{j \in i\text{th cluster}} (\mathbf{x}_j - \mathbf{c}_i)^T (\mathbf{x}_j - \mathbf{c}_i) \right\}.$$

Notice that if we know the center for each cluster, it is easy to determine which cluster is the best choice for each point. Similarly, if the allocation of points to clusters is known, it is easy to compute the best center for each cluster. However, there are far too many possible allocations of points to clusters to search this space for a minimum. Instead, we define an algorithm that iterates through two activities:

- Assume the cluster centers are known and, allocate each point to the closest cluster center.

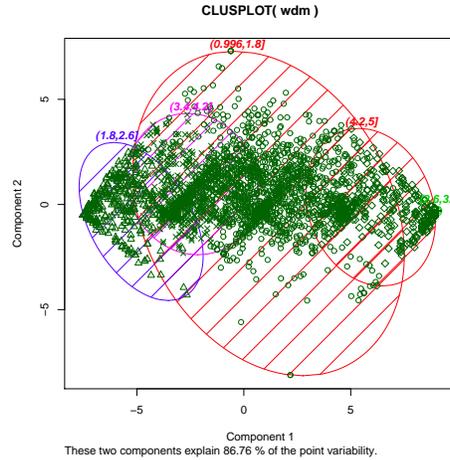


FIGURE 12.6: A clustering of the student evaluation dataset, using agglomerative clustering and 5 clusters. I produced this clustering with the code fragment in listing 12.1. One difficulty with this dataset is there are 28 questions (Q1-Q28), meaning the data is quite difficult to plot. You should be able to see the clusters on the figure, but notice that they are smeared on top of one another; this is because I had to project off 26 dimensions to produce the plot. The plotting function I used (`clusplot`) plots a covariance ellipse for each cluster.

- Assume the allocation is known, and choose a new set of cluster centers. Each center is the mean of the points allocated to that cluster.

We then choose a start point by randomly choosing cluster centers, and then iterate these stages alternately. This process eventually converges to a local minimum of the objective function (the value either goes down or is fixed at each step, and it is bounded below). It is not guaranteed to converge to the global minimum of the objective function, however. It is also not guaranteed to produce  $k$  clusters, unless we modify the allocation phase to ensure that each cluster has some nonzero number of points. This algorithm is usually referred to as **k-means** (summarized in Algorithm 12.3).

Usually, we are clustering high dimensional data, so that visualizing clusters can present a challenge. If the dimension isn't too high, then we can use panel plots. An alternative is to project the data onto two principal components, and plot the clusters there. A natural dataset to use to explore k-means is the iris data, where we know that the data should form three clusters (because there are three species). Recall this dataset from section ???. I reproduce figure 10.5 from that section as figure 12.11, for comparison. Figures 12.8, ??? and ??? show different k-means clusterings of that data.

Listing 12.1: R code for student example.

```

setwd('/users/daf/Current/courses/Probcourse/Clustering/RCode')
wdat<-read.csv('turkiye-student-evaluation_R-Specific.csv')
wdm<-wdat[,c(6:33)] # choose the questions
d <- dist(wdm, method = "euclidean") # distance matrix
fit <- hclust(d, method="ward") # the clustering
cm<-cov(wdm)
ev<-eigen(cm)
princ<-ev$vector
proj<-t(rbind(t(princ[, 1]), t(princ[, 2])))
wdzm<-scale(wdm, center=TRUE, scale=FALSE) #subtract the mean
wdzm<-data.matrix(wdzm)
wdpcs<-wdzm%*%proj
kv<-c(5)
groups<-cutree(fit, k=kv)
wdbig<-wdm
wdbig$group<-groups
cg<-cut(groups, kv)
p<-ggplot(as.data.frame(wdpcs))+geom_point(aes(y=V1, x=V2, shape=cg))
setEPS()
postscript("clusters.eps")
p
dev.off()
cm1<-colMeans(subset(wdbig, group==1)) # mean of the 1st cluster

```

**Worked example 12.4** *K-means clustering in R*

Cluster the iris dataset into two clusters using k-means, then plot the results on the first two principal components

**Solution:** I used the code fragment in listing 12.2, which produced figure ??

Choose  $k$  data points to act as cluster centers  
 Until the cluster centers change very little  
   Allocate each data point to cluster whose center is nearest.  
   Now ensure that every cluster has at least  
     one data point; one way to do this is by  
     supplying empty clusters with a point chosen at random from  
     points far from their cluster center.  
 Replace the cluster centers with the mean of the elements  
   in their clusters.  
 end

**Algorithm 12.3:** *Clustering by K-Means.*

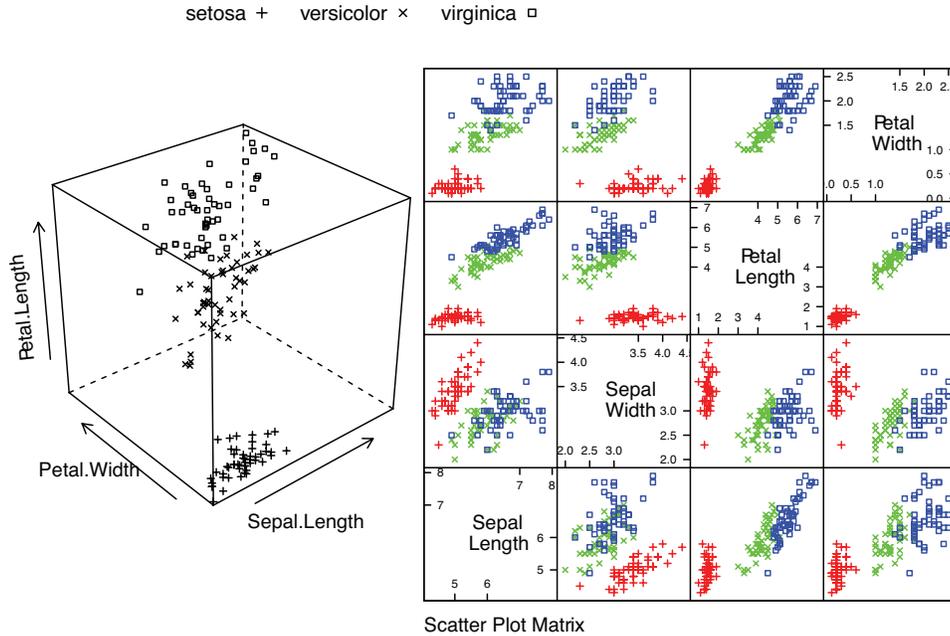


FIGURE 12.7: **Left:** a 3D scatterplot for the famous Iris data, originally due to *\*\*\**. I have chosen three variables from the four, and have plotted each species with a different marker. You can see from the plot that the species cluster quite tightly, and are different from one another. **Right:** a scatterplot matrix for the famous Iris data, originally due to *\*\*\**. There are four variables, measured for each of three species of iris. I have plotted each species with a different marker. You can see from the plot that the species cluster quite tightly, and are different from one another.

### 12.4.1 How to choose K

The iris data is just a simple example. We know that the data forms clean clusters, and we know there should be three of them. Usually, we don't know how many clusters there should be, and we need to choose this by experiment. One strategy is to cluster for a variety of different values of  $k$ , then look at the value of the cost function for each. If there are more centers, each data point can find a center that is close to it, so we expect the value to go down as  $k$  goes up. This means that looking for the  $k$  that gives the smallest value of the cost function is not helpful, because that  $k$  is always the same as the number of data points (and the value is then zero). However, it can be very helpful to plot the value as a function of  $k$ , then look at the “knee” of the curve. Figure 12.11 shows this plot for the iris data. Notice that  $k = 3$  — the “true” answer — doesn't look particularly special, but  $k = 2$ ,  $k = 3$ , or  $k = 4$  all seem like reasonable choices. It is possible to come up with a procedure that makes a more precise recommendation by penalizing clusterings that use a large  $k$ , because they may represent inefficient encodings of the data. However, this is often not worth the bother.

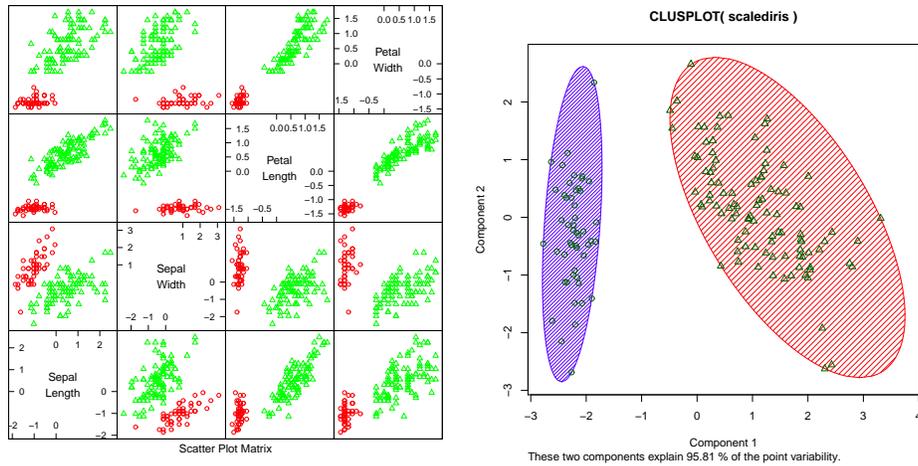


FIGURE 12.8: *On the left*, a panel plot of the iris data clustered using  $k$ -means with  $k = 2$ . By comparison with figure 12.11, notice how the versicolor and virginica clusters appear to have been merged. *On the right*, this data set projected onto the first two principal components, with one blob drawn over each cluster.

In some special cases (like the iris example), we might know the right answer to check our clustering against. In such cases, one can evaluate the clustering by looking at the number of different labels in a cluster (sometimes called the purity), and the number of clusters. A good solution will have few clusters, all of which have high purity.

Mostly, we don't have a right answer to check against. An alternative strategy, which might seem crude to you, for choosing  $k$  is extremely important in practice. Usually, one clusters data to use the clusters in an application (one of the most important, vector quantization, is described in section 12.5). There are usually natural ways to evaluate this application. For example, vector quantization is often used as an early step in texture recognition or in image matching; here one can evaluate the error rate of the recognizer, or the accuracy of the image matcher. One then chooses the  $k$  that gets the best evaluation score on validation data. In this view, the issue is not how good the clustering is; it's how well the system that uses the clustering works.

#### 12.4.2 Soft Assignment

One difficulty with  $k$ -means is that each point must belong to exactly one cluster. But, given we don't know how many clusters there are, this seems wrong. If a point is close to more than one cluster, why should it be forced to choose? This reasoning suggests we assign points to cluster centers with weights.

We allow each point to carry a total weight of 1. In the conventional  $k$ -means algorithm, it must choose a single cluster, and assign its weight to that cluster alone. In soft-assignment  $k$ -means, the point can allocate some weight to each cluster center, as long as (a) the weights are all non-negative and (b) the weights

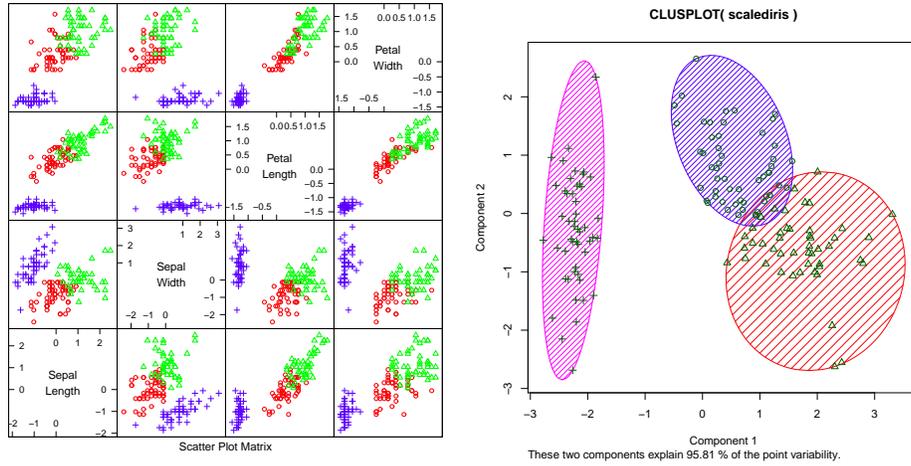


FIGURE 12.9: On the left, a panel plot of the iris data clustered using  $k$ -means with  $k = 3$ . By comparison with figure 12.11, notice how the clusters appear to follow the species labels. On the right, this data set projected onto the first two principal components, with one blob drawn over each cluster.

sum to one. Write  $w_{i,j}$  for the weight connecting point  $i$  to cluster center  $j$ . We interpret these weights as the degree to which the point participates in a particular cluster. We require  $w_{i,j} \geq 0$  and  $\sum_j w_{i,j} = 1$ .

We would like  $w_{i,j}$  to be large when  $\mathbf{x}_i$  is close to  $\mathbf{c}_j$ , and small otherwise. Write  $d_{i,j}$  for the distance  $\|\mathbf{x}_i - \mathbf{c}_j\|$  between these two. Write

$$s_{i,j} = e^{-\frac{d_{i,j}^2}{2\sigma^2}}$$

where  $\sigma > 0$  is a choice of scaling parameter. This is often called the affinity between the point  $i$  and the center  $j$ . Now a natural choice of weights is

$$w_{i,j} = \frac{s_{i,j}}{\sum_{l=1}^k s_{i,l}}.$$

All these weights are non-negative, they sum to one, and the weight is large if the point is much closer to one center than to any other. The scaling parameter  $\sigma$  sets the meaning of “much closer” — we measure distance in units of  $\sigma$ .

Once we have weights, re-estimating the cluster centers is easy. We use a weights to compute a weighted average of the points. In particular, we re-estimate the  $j$ 'th cluster center by

$$\frac{\sum_i w_{i,j} \mathbf{x}_i}{\sum_i w_{i,j}}.$$

Notice that  $k$ -means is a special case of this algorithm where  $\sigma$  limits to zero. In this case, each point has a weight of one for some cluster, and zero for all others, and the weighted mean becomes an ordinary mean. I have collected the description into Algorithm 12.4 for convenience.

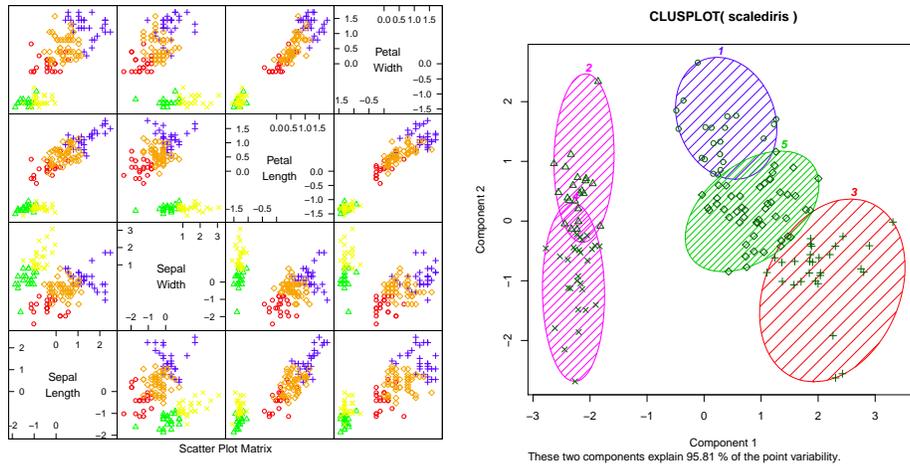


FIGURE 12.10: On the **left**, a panel plot of the iris data clustered using  $k$ -means with  $k = 5$ . By comparison with figure 12.11, notice how setosa seems to have been broken in two groups, and versicolor and virginica into a total of three. On the **right**, this data set projected onto the first two principal components, with one blob drawn over each cluster.

### 12.4.3 General Comments on K-Means

If you experiment with  $k$ -means, you will notice one irritating habit of the algorithm. It almost always produces either some rather spread out clusters, or some single element clusters. Most clusters are usually rather tight and blobby clusters, but there is usually one or more bad cluster. This is fairly easily explained. Because every data point must belong to some cluster, data points that are far from all others (a) belong to some cluster and (b) very likely “drag” the cluster center into a poor location. This applies even if you use soft assignment, because every point must have total weight one. If the point is far from all others, then it will be assigned to the closest with a weight very close to one, and so may drag it into a poor location, or it will be in a cluster on its own.

There are ways to deal with this. If  $k$  is very big, the problem is often not significant, because then you simply have many single element clusters that you can ignore. It isn’t always a good idea to have too large a  $k$ , because then some larger clusters might break up. An alternative is to have a junk cluster. Any point that is too far from the closest true cluster center is assigned to the junk cluster, and the center of the junk cluster is not estimated. Notice that points should not be assigned to the junk cluster permanently; they should be able to move in and out of the junk cluster as the cluster centers move.

In some cases, we want to cluster objects that can’t be averaged. For example, you can compute distances between two trees but you can’t meaningfully average them. In some cases, you might have a table of distances between objects, but not know vectors representing the objects. For example, one could collect data on the similarities between countries (as in Section 10.4.2, particularly Figure 10.25),

Listing 12.2: R code for iris example.

```

setwd('/users/daf/Current/courses/Probcourse/Clustering/RCode')
#library('lattice')
# work with iris dataset this is famous, and included in R
# there are three species
head(iris)
#
library('cluster')
numiris=iris[,c(1, 2, 3, 4)] #the numerical values
#scalediris<-scale(numiris) # scale to unit variance
nclus<-2
sfit<-kmeans(numiris, nclus)
colr<-c('red', 'green', 'blue', 'yellow', 'orange')
clusplot(numiris, sfit$cluster, color=TRUE, shade=TRUE,
         labels=0, lines=0)

```

then try and cluster using this data (similarities can be turned into distances by, for example, taking the negative logarithm). A variant of k-means, known as k-medoids, applies to this case.

In k-medoids, the cluster centers are data items rather than averages, but the rest of the algorithm has a familiar form. We assume the number of medoids is known, and initialize these randomly. We then iterate two procedures. In the first, we allocate data points to medoids. In the second, we choose the best medoid for each cluster by finding the medoid that minimizes the sum of distances of points in the cluster to that medoid (blank search is fine).

## 12.5 DESCRIBING REPETITION WITH VECTOR QUANTIZATION

Classifying images is an important problem, with a variety of applications. We have seen powerful classification methods. But what features should one use to classify images? or sounds? or accelerometer signals?

In this section, I will describe a construction that extends to a variety of signals (i.e. images, video, sounds, accelerometer output, ultrasound, etc.), because it exploits repetition. Repetition is an important feature of many interesting signals. For example, images contain *textures*, which are orderly patterns that look like large numbers of small structures that are repeated. Examples include the spots of animals such as leopards or cheetahs; the stripes of animals such as tigers or zebras; the patterns on bark, wood, and skin. Similarly, speech signals contain *phonemes* — characteristic, stylised sounds that people assemble together to produce speech (for example, the “ka” sound followed by the “tuh” sound leading to “cat”). An important part of representing signals that repeat is identifying the components that repeat, then describing the signal in terms of those components.

Repetition occurs in subtle forms. The essence is that a small number of local patterns can be used to represent a large number of examples. You see this effect in pictures of scenes. If you collect many pictures of, say, a beach scene, you will expect most to contain some waves, some sky, and some sand. The individual

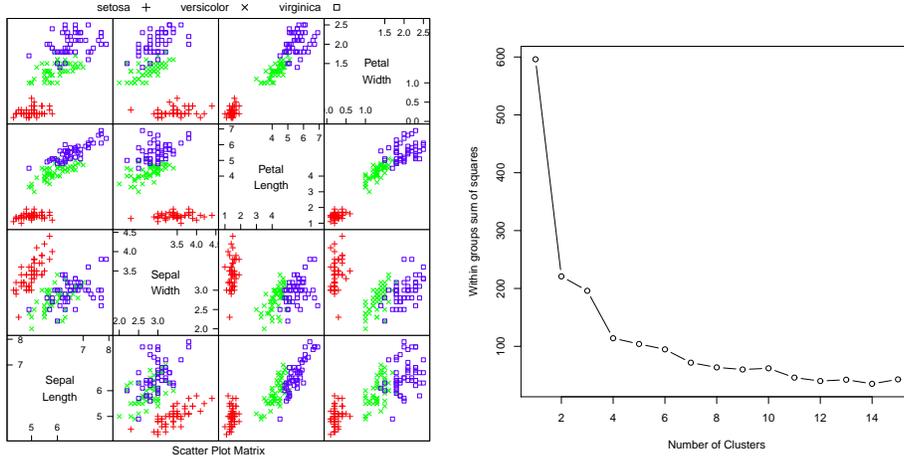


FIGURE 12.11: On the left, the scatterplot matrix for the Iris data, for reference. On the right, a plot of the value of the cost function for each of several different values of  $k$ . Notice how there is a sharp drop in cost going from  $k = 1$  to  $k = 2$ , and again at  $k = 4$ ; after that, the cost falls off slowly. This suggests using  $k = 2$ ,  $k = 3$ , or  $k = 4$ , depending on the precise application.

patches of wave, sky or sand can be surprisingly similar, and different images are made by selecting some patches from a vocabulary of patches, then placing them down to form an image. Similarly, pictures of living rooms contain chair patches, TV patches, and carpet patches. Many different living rooms can be made from small vocabularies of patches; but you won't often see wave patches in living rooms, or carpet patches in beach scenes.

This view suggests that we should (a) develop a vocabulary of patterns and then (b) construct a feature that describes which vocabulary elements appear in the signal. It is often the case that where the vocabulary elements appear is less important than which ones appear. This is true for images, because different pictures of a living room may have the same patches in different locations (for example, you might just move the camera to the left). As another example, zebras have stripes, but precisely where the stripes lie on a zebra doesn't seem to matter much in practice. As we shall see below, it is true for other kinds of data, too.

We must now develop a vocabulary of patterns. We could find these patterns by looking for small pieces of signal of fixed size that are common. In an image, a piece of signal might be a 10x10 patch; in a sound file, which is likely represented as a vector, it might be a subvector of fixed size. But finding common patterns is hard to do, because the signal is continuous — each pattern will be slightly different, so we cannot simply count how many times a particular pattern occurs.

### 12.5.1 Vector Quantization

**Vector quantization** is a strategy to deal with these difficulties. Vector quantization is a procedure that can take a vector in a continuous space, and replace

```

Choose  $k$  data points to act as cluster centers
Until the cluster centers change very little
  First, we estimate the weights

  For  $i$  indexing data points
    For  $j$  indexing cluster centers
      Compute  $s_{i,j} = e^{\frac{-\|\mathbf{x}_i - \mathbf{c}_j\|}{2\sigma^2}}$ 

  For  $i$  indexing data points
    For  $j$  indexing cluster centers
      Compute  $w_{i,j} = s_{i,j} / \sum_{l=1}^k s_{i,l}$ 

  Now, we re-estimate the centers

  For  $j$  indexing cluster centers
    Compute  $\mathbf{c}_j = \frac{\sum_i w_{i,j} \mathbf{x}_i}{\sum_i w_{i,j}}$ 

end

```

**Algorithm 12.4:** *Soft Clustering by K-Means.*

it with a number in the range  $[1, \dots, k]$  (where you get to choose  $k$ ). To vector quantize, we first build a set of clusters out of a training set of vectors; this set of clusters is often thought of as a dictionary. We now replace any new vector with the cluster center closest to that vector. Many different clusterers can be used for vector quantization, but it is most common to use k-means or one of its variants.

We can now represent a collection of vectors as a histogram of cluster centers.

At <http://archive.ics.uci.edu/ml/datasets/Wholesale+customers>, you will find a dataset giving sums of money spent annually on different commodities by customers in Portugal. These customers are divided by channel (two channels) and by region (three regions). You can think of the data as being divided into six groups (one for each channel-region pair). There are 440 records, and so many customers per class. The commodities are divided into a set of categories (fresh; milk; grocery; frozen; detergents and paper; and delicatessen) relevant for the study. Figure 12.13 shows a panel plot of the customer data. Relatively little structure is apparent here — you can't, for example, see any evidence of six clean clusters.

If you think about it, it's too much to expect that all customers in a class are similar. Instead, we expect that there might be different “types” of customer — some might spend more money on milk, others on grocery, and so on. Different groups likely contain different numbers of each type of customer. This would make it hard to see the groups in the panel plot.

Vector quantization can expose this kind of structure rather effectively. I used k-means clustering to cluster the customer data to 20 different clusters (Figure 12.13). Then I described the each group of data by the histogram of customer

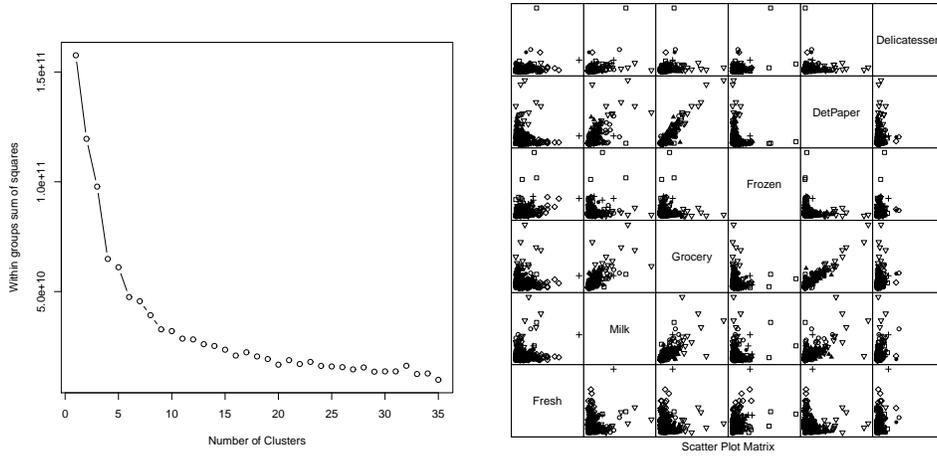


FIGURE 12.12: *On the left*, the sum of squared error for clusterings of the customer data with  $k$ -means for  $k$  running from 2 to 35. This suggests using a  $k$  somewhere in the range 10-30; I chose 20. *On the right*, I have clustered this data to 20 cluster centers with  $k$ -means. The clusters do seem to be squashed together, but the plot on the left suggests that clusters do capture some important information. Using too few clusters will clearly lead to problems. Notice that I did not scale the data, because each of the measurements is in a comparable unit. For example, wouldn't make sense to scale expenditures on fresh and expenditures on grocery with a different scale.

types that appeared in that group (Figure ??). Notice how the distinction between the groups is now apparent — the groups do appear to contain quite different distributions of customer type. To be more confident in this analysis, we would need to be sure that different types of customer really are different. We could do this by repeating the analysis for fewer clusters, or by looking at the similarity of customer types.

### 12.5.2 Efficient Clustering and Hierarchical K Means

One important difficulty occurs in applications. We might need to have an enormous dataset (millions of image patches are a real possibility), and so a very large  $k$ . In this case,  $k$  means clustering becomes difficult because identifying which cluster center is closest to a particular data point scales linearly with  $k$  (and we have to do this for every data point at every iteration). There are two useful strategies for dealing with this problem.

The first is to notice that, if we can be reasonably confident that each cluster contains many data points, some of the data is redundant. We could randomly subsample the data, cluster that, then keep the cluster centers. This works, but doesn't scale particularly well.

A more effective strategy is to build a hierarchy of  $k$ -means clusters. We randomly subsample the data (typically, quite aggressively), then cluster this with

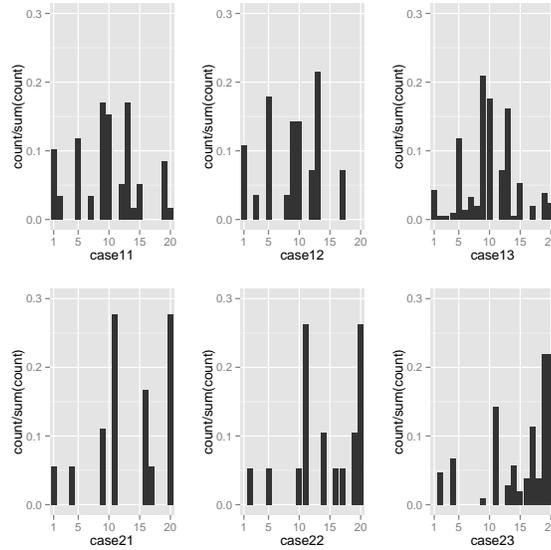


FIGURE 12.13: *The histogram of different types of customer, by group, for the customer data. Notice that different groups do really appear to have different types of customer represented in different amounts.*

a small value of  $k$ . Each data item is then allocated to the closest cluster center, and the data in each cluster is clustered again with k-means. We now have something that looks like a two-level tree of clusters. Of course, this process can be repeated to produce a multi-level tree of clusters. It is easy to use this tree to vector quantize a query data item. We vector quantize at the first level. Doing so chooses a branch of the tree, and we pass the data item to this branch. It is either a leaf, in which case we report the number of the leaf, or it is a set of clusters, in which case we vector quantize, and pass the data item down. This procedure is efficient both when one clusters and at run time.

### 12.5.3 Example: Activity from Accelerometer Data

A complex example dataset appears at <https://archive.ics.uci.edu/ml/datasets/Dataset+for+ADL+Recognition+>. This dataset consists of examples of the signal from a wrist mounted accelerometer, produced as different subjects engaged in different activities of daily life. Activities include: brushing teeth, climbing stairs, combing hair, descending stairs, and so on. Each is performed by sixteen volunteers. The accelerometer samples the data at 32Hz (i.e. this data samples and reports the acceleration 32 times per second). The accelerations are in the x, y and z-directions. Figure 12.14 shows the x-component of various examples of toothbrushing.

There is an important problem with using data like this. Different subjects take quite different amounts of time to perform these activities. For example, some subjects might be more thorough tooth-brushers than other subjects. As another example, people with longer legs walk at somewhat different frequencies than people

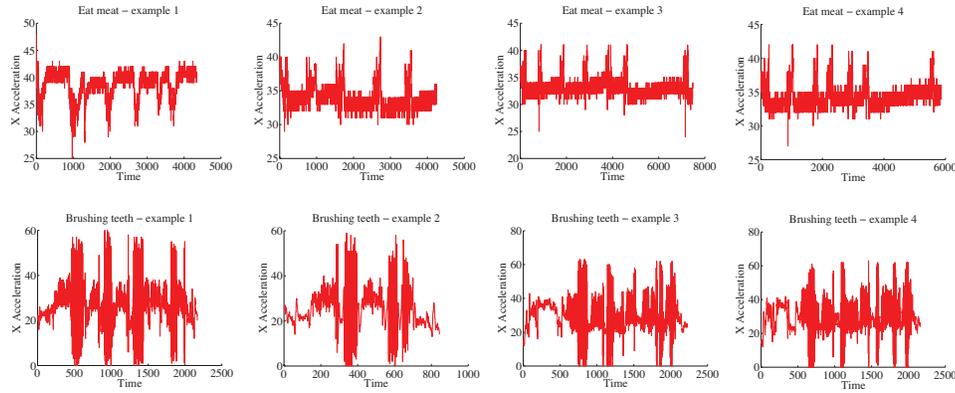


FIGURE 12.14: Some examples from the accelerometer dataset at <https://archive.ics.uci.edu/ml/datasets/Dataset+for+ADL+Recognition+with+Wrist-worn+Accelerometer>. I have labelled each signal by the activity. These show acceleration in the X direction (Y and Z are in the dataset, too). There are four examples for **brushing teeth** and four for **eat meat**. You should notice that the examples don't have the same length in time (some are slower and some faster eaters, etc.), but that there seem to be characteristic features that are shared within a category (brushing teeth seems to involve faster movements than eating meat).

with shorter legs. This means that the same activity performed by different subjects will produce data vectors *that are of different lengths*. It's not a good idea to deal with this by warping time and resampling the signal. For example, doing so will make a thorough toothbrusher look as though they are moving their hands very fast (or a careless toothbrusher look ludicrously slow: think speeding up or slowing down a movie). So we need a representation that can cope with signals that are a bit longer or shorter than other signals.

Another important property of these signals is that all examples of a particular activity should contain repeated patterns. For example, brushing teeth should show fast accelerations up and down; walking should show a strong signal at somewhere around 2 Hz; and so on. These two points should suggest vector quantization to you. Representing the signal in terms of stylized, repeated structures is probably a good idea because the signals probably contain these structures. And if we represent the signal in terms of the relative frequency with which these structures occur, the representation will have a fixed length, even if the signal doesn't. To do so, we need to consider (a) over what time scale we will see these repeated structures and (b) how to ensure we segment the signal into pieces so that we see these structures.

Generally, repetition in activity signals is so obvious that we don't need to be smart about segment boundaries. I broke these signals into 32 sample segments, one following the other. Each segment represents one second of activity. This is long enough for the body to do something interesting, but not so long that our representation will suffer if we put the segment boundaries in the wrong place. This resulted in about 40,000 segments. I then used hierarchical k-means to cluster these

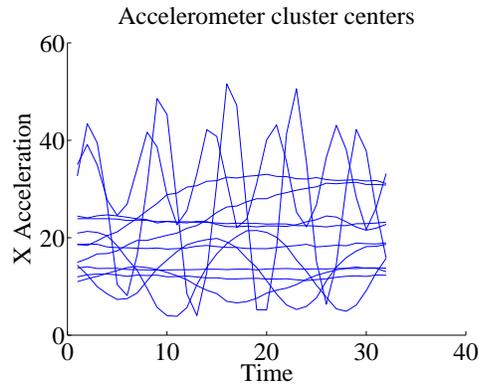


FIGURE 12.15: *Some cluster centers from the accelerometer dataset. Each cluster center represents a one-second burst of activity. There are a total of 480 in my model, which I built using hierarchical  $k$ -means. Notice there are a couple of centers that appear to represent movement at about 5Hz; another few that represent movement at about 2Hz; some that look like 0.5Hz movement; and some that seem to represent much lower frequency movement. These cluster centers are samples (rather than chosen to have this property).*

segments. I used two levels, with 40 cluster centers at the first level, and 12 at the second. Figure 12.15 shows some cluster centers at the second level.

I then computed histogram representations for different example signals (Figure 12.16). You should notice that when the activity label is different, the histogram looks different, too.

Another useful way to check this representation is to compare the average within class chi-squared distance with the average between class chi-squared distance. I computed the histogram for each example. Then, for each pair of examples, I computed the chi-squared distance between the pair. Finally, for each pair of *activity labels*, I computed the average distance between pairs of examples where one example has one of the activity labels and the other example has the other activity label. In the ideal case, all the examples with the same label would be very close to one another, and all examples with different labels would be rather different. Table 12.1 shows what happens with the real data. You should notice that for some pairs of activity label, the mean distance between examples is smaller than one would hope for (perhaps some pairs of examples are quite close?). But generally, examples of activities with different labels tend to be further apart than examples of activities with the same label.

Yet another way to check the representation is to try classification with nearest neighbors, using the chi-squared distance to compute distances. I split the dataset into 80 test pairs and 360 training pairs; using 1-nearest neighbors, I was able to get a held-out error rate of 0.79. This suggests that the representation is fairly good at exposing what is important.

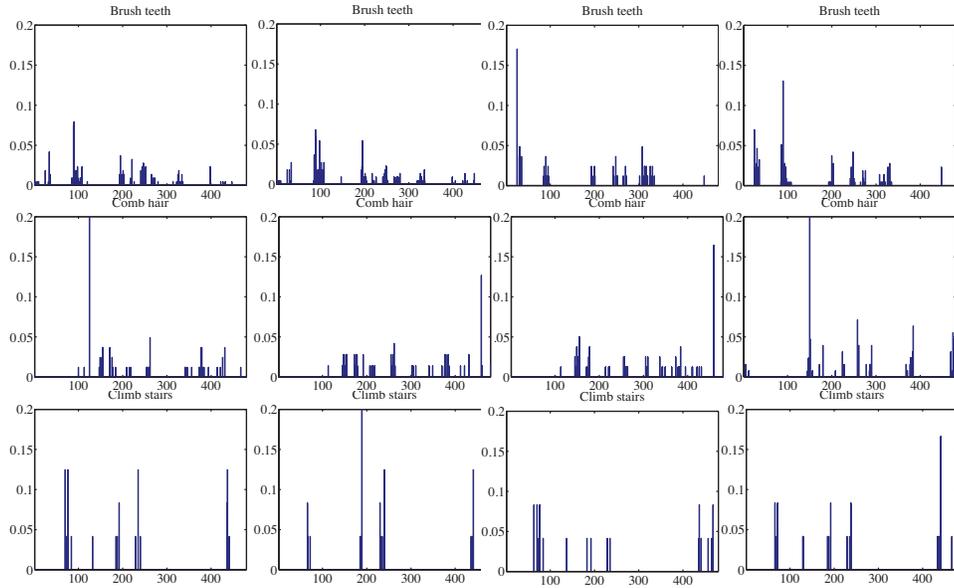


FIGURE 12.16: Histograms of cluster centers for the accelerometer dataset, for different activities. You should notice that (a) these histograms look somewhat similar for different actors performing the same activity and (b) these histograms look somewhat different for different activities.

12.6 YOU SHOULD

12.6.1 remember:

- New term: clustering . . . . . 357
- New term: clusters . . . . . 359
- Useful facts: Parameters of a Multivariate Normal Distribution . . . 360
- New term: covariance ellipses . . . . . 361
- New term: decorrelation . . . . . 364
- New term: whitening . . . . . 364
- New term: k-means . . . . . 369
- New term: Vector quantization . . . . . 376

0.9	2.0	1.9	2.0	2.0	2.0	1.9	2.0	1.9	1.9	2.0	2.0	2.0	2.0
	1.6	2.0	1.8	2.0	2.0	2.0	1.9	1.9	2.0	1.9	1.9	2.0	1.7
		1.5	2.0	1.9	1.9	1.9	1.9	1.9	1.9	1.9	1.9	1.9	2.0
			1.4	2.0	2.0	2.0	2.0	2.0	2.0	2.0	2.0	2.0	1.8
				1.5	1.8	1.7	1.9	1.9	1.8	1.9	1.9	1.8	2.0
					0.9	1.7	1.9	1.9	1.8	1.9	1.9	1.9	2.0
						0.3	1.9	1.9	1.5	1.9	1.9	1.9	2.0
							1.8	1.8	1.9	1.9	1.9	1.9	1.9
								1.7	1.9	1.9	1.9	1.9	1.9
									1.6	1.9	1.9	1.9	2.0
										1.8	1.9	1.9	1.9
											1.8	2.0	1.9
												1.5	2.0
													1.5

TABLE 12.1: Each column of the table represents an activity for the activity dataset <https://archive.ics.uci.edu/ml/datasets/Dataset+for+ADL+Recognition+with+Wrist-worn+Accelerometer>, as does each row. In each of the upper diagonal cells, I have placed the average chi-squared distance between histograms of examples from that pair of classes (I dropped the lower diagonal for clarity). Notice that in general the diagonal terms (average within class distance) are rather smaller than the off diagonal terms. This quite strongly suggests we can use these histograms to classify examples successfully.

PROGRAMMING EXERCISES

- 12.1.** Obtain the activities of daily life dataset from the UC Irvine machine learning website (<https://archive.ics.uci.edu/ml/datasets/Dataset+for+ADL+Recognition+with+Wrist-worn+Accelerometer>) data provided by Barbara Bruno, Fulvio Mastrogiovanni and Antonio Sgorbissa).
- (a) Build a classifier that classifies sequences into one of the 14 activities provided. To make features, you should vector quantize, then use a histogram of cluster centers (as described in the subsection; this gives a pretty explicit set of steps to follow). You will find it helpful to use hierarchical k-means to vector quantize. You may use whatever multi-class classifier you wish, though I'd start with R's decision forest, because it's easy to use and effective. You should report (a) the total error rate and (b) the class confusion matrix of your classifier.
  - (b) Now see if you can improve your classifier by (a) modifying the number of cluster centers in your hierarchical k-means and (b) modifying the size of the fixed length samples that you use.

# Regression

Classification tries to predict a class from a data item. **Regression** tries to predict a value. For example, we know the zip code of a house, the square footage of its lot, the number of rooms and the square footage of the house, and we wish to predict its likely sale price. As another example, we know the cost and condition of a trading card for sale, and we wish to predict a likely profit in buying it and then reselling it. As yet another example, we have a picture with some missing pixels – perhaps there was text covering them, and we want to replace it – and we want to fill in the missing values. As a final example, you can think of classification as a special case of regression, where we want to predict either  $+1$  or  $-1$ ; this isn't usually the best way to proceed, however.

Predicting values is very useful, and so there are many examples like this. Another reason to build a regression model is to compare trends in data. Doing so could make it clear what is really happening. Here is an example from Efron (“Computer-Intensive methods in statistical regression”, B. Efron, SIAM Review, 1988). The table in the appendix shows some data from medical devices, which sit in the body and release a hormone. The data shows the amount of hormone currently in a device after it has spent some time in service, and the time the device spent in service. The data describes devices from three production lots (A, B, and C). Each device, from each lot, is supposed to have the same behavior. The important question is: Are the lots the same? The amount of hormone changes over time, so we can't just compare the amounts currently in each device. Instead, we need to determine the relationship between time in service and hormone, and see if this relationship is different between batches. We can do so by regressing hormone against time.

Figure 13.1 shows how a regression can help. In this case, we have modelled the amount of hormone in the device as

$$a \times (\text{time in service}) + b$$

for  $a$ ,  $b$  chosen to get the best fit (much more on this point later!). This means we can plot each data point on a scatter plot, together with the best fitting line. This plot allows us to ask whether any particular batch behaves differently from the overall model in any interesting way.

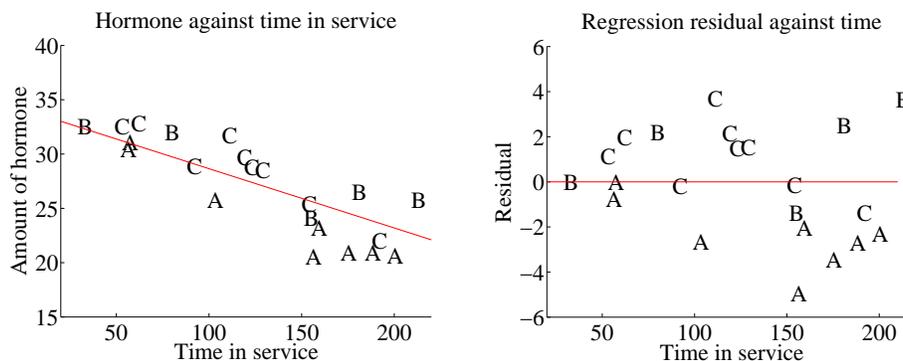


FIGURE 13.1: On the **left**, a scatter plot of hormone against time for devices from tables 13.1 and 13.1. Notice that there is a pretty clear relationship between time and amount of hormone (the longer the device has been in service the less hormone there is). The issue now is to understand that relationship so that we can tell whether lots A, B and C are the same or different. The best fit line to all the data is shown as well, fitted using the methods of section 13.1. On the **right**, a scatter plot of residual — the distance between each data point and the best fit line — against time for the devices from tables 13.1 and 13.1. Now you should notice a clear difference; some devices from lots B and C have positive and some negative residuals, but all lot A devices have negative residuals. This means that, when we account for loss of hormone over time, lot A devices still have less hormone in them. This is pretty good evidence that there is a problem with this lot.

**Definition: 13.1** *Regression*

Regression accepts a feature vector and produces a prediction, which is usually a number, but can sometimes have other forms. It is possible, but not usually particularly helpful, to see classification as a form of regression.

Indexing: Regression

Indexing: Regression

However, it is hard to evaluate the distances between data points and the best fitting line by eye. A sensible alternative is to subtract the amount of hormone predicted by the model from the amount that was measured. Doing so yields a **residual** — the difference between a measurement and a prediction. We can then plot those residuals (Figure 13.1). In this case, the plot suggests that lot A is special — all devices from this lot contain less hormone than our model predicts.

## 13.1 LINEAR REGRESSION AND LEAST SQUARES

Assume we have a dataset consisting of a set of  $N$  pairs  $(\mathbf{x}_i, y_i)$ . We think of  $y_i$  as the value of some function evaluated at  $\mathbf{x}_i$ , with some random component added. This means there might be two data items where the  $\mathbf{x}_i$  are the same, and the  $y_i$  are different. We refer to the  $\mathbf{x}_i$  as **explanatory variables** and the  $y_i$  is a **dependent variable**. We want to use the examples we have — the **training examples** — to build a model of the dependence between  $y$  and  $\mathbf{x}$ . This model will be used to predict values of  $y$  for new values of  $\mathbf{x}$ , which are usually called **test examples**. It can also be used to understand the relationships between the  $\mathbf{x}$ . The model needs to have some probabilistic component; we do not expect that  $y$  is a function of  $\mathbf{x}$ , and there is likely some error in evaluating  $y$  anyhow.

### 13.1.1 Linear Regression

A good, simple model is to assume that the dependent variable is obtained by evaluating a linear function of the explanatory variables, then adding a zero-mean normal random variable. We can write this model as

$$y = \mathbf{x}^T \beta + \xi$$

where  $\xi$  is a zero mean normal random variable with unknown variance. In this expression,  $\beta$  is a vector of weights, which we must estimate. When we use this model to predict a value of  $y$  for a particular set of explanatory variables  $\mathbf{x}^*$ , we cannot predict the value that  $\xi$  will take. Our best available prediction is the mean value (which is zero). Notice that if  $\mathbf{x} = 0$ , the model predicts  $y = 0$ . This may seem like a problem to you — you might be concerned that we can fit only lines through the origin — but remember that  $\mathbf{x}$  contains explanatory variables, and we can choose what appears in  $\mathbf{x}$ . The two examples show how a sensible choice of  $\mathbf{x}$  allows us to fit a line with an arbitrary  $y$ -intercept.

**Definition: 13.2** *Linear regression*

A linear regression takes the feature vector  $\mathbf{x}$  and predicts  $\mathbf{x}^T \beta$ , for some vector of coefficients  $\beta$ . The coefficients are adjusted, using data, to produce the best predictions.

Indexing: Linear regression

Indexing: Linear regression

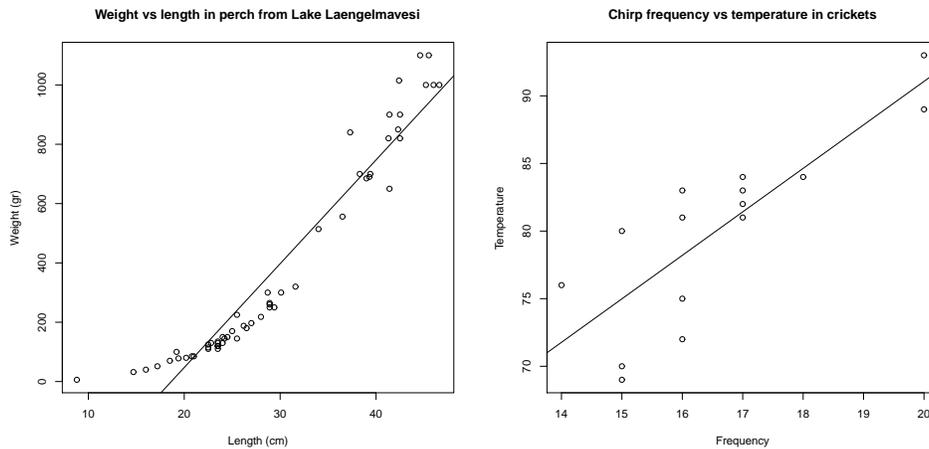


FIGURE 13.2: On the **left**, a regression of weight against length for perch from a Finnish lake (you can find this dataset, and the back story at [http://www.amstat.org/publications/jse/jse\\_data\\_archive.htm](http://www.amstat.org/publications/jse/jse_data_archive.htm); look for “fishcatch” on that page). Notice that the linear regression fits the data fairly well, meaning that you should be able to predict the weight of a perch from its length fairly well. On the **right**, a regression of air temperature against chirp frequency for crickets. The data is fairly close to the line, meaning that you should be able to tell the temperature from the pitch of cricket’s chirp fairly well. This data is from <http://mste.illinois.edu/patel/amar430/keyprob1.html>.

**Example: 13.1** A linear model fitted to a single explanatory variable

Assume we fit a linear model to a single explanatory variable. Then the model has the form  $y = x\beta + \xi$ , where  $\xi$  is a zero mean random variable. For any value  $x^*$  of the explanatory variable, our best estimate of  $y$  is  $\beta x^*$ . In particular, if  $x^* = 0$ , the model predicts  $y = 0$ , which is unfortunate. We can draw the model by drawing a line through the origin with slope  $\beta$  in the  $x, y$  plane. The  $y$ -intercept of this line must be zero.

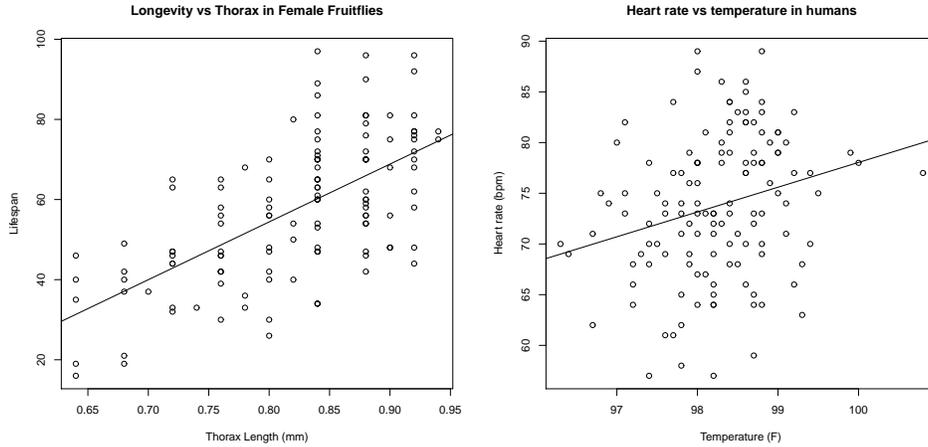


FIGURE 13.3: Regressions do not necessarily yield good predictions or good model fits. On the **left**, a regression of the lifespan of female fruitflies against the length of their torso as adults (apparently, this doesn't change as a fruitfly ages; you can find this dataset, and the back story at [http://www.amstat.org/publications/jse/jse\\_data\\_archive.htm](http://www.amstat.org/publications/jse/jse_data_archive.htm); look for “fruitfly” on that page). The figure suggests you can make some prediction of how long your fruitfly will last by measuring its torso, but not a particularly accurate one. On the **right**, a regression of heart rate against body temperature for adults. You can find the data at [http://www.amstat.org/publications/jse/jse\\_data\\_archive.htm](http://www.amstat.org/publications/jse/jse_data_archive.htm) as well; look for “temperature” on that page. Notice that predicting heart rate from body temperature isn't going to work that well, either.

**Example: 13.2** A linear model with a non-zero  $y$ -intercept

Assume we have a single explanatory variable, which we write  $u$ . We can then create a vector  $\mathbf{x} = [u, 1]^T$  from the explanatory variable. We now fit a linear model to this vector. Then the model has the form  $y = \mathbf{x}^T \beta + \xi$ , where  $\xi$  is a zero mean random variable. For any value  $\mathbf{x}^* = [u^*, 1]^T$  of the explanatory variable, our best estimate of  $y$  is  $(\mathbf{x}^*)^T \beta$ , which can be written as  $y = \beta_1 u^* + \beta_2$ . If  $x^* = 0$ , the model predicts  $y = \beta_2$ . We can draw the model by drawing a line through the origin with slope  $\beta_1$  and  $y$ -intercept  $\beta_2$  in the  $x, y$  plane.

We must determine  $\beta$ . Because we have that  $P(y|x, \beta)$  is normal, with mean  $\mathbf{x}^T \beta$ , we can write out the log-likelihood of the data. Write  $\sigma^2$  for the variance of

$\xi$ , which we don't know, but will not need to worry about right now. We have that

$$\begin{aligned}\log \mathcal{L}(\beta) &= \sum_i \log P(y_i | \mathbf{x}_i, \beta) \\ &= \frac{1}{2\sigma^2} \sum_i (y_i - \mathbf{x}_i^T \beta)^2 + \text{term not depending on } \beta\end{aligned}$$

Maximizing the log-likelihood of the data is equivalent to minimizing the negative log-likelihood of the data. Furthermore, the term  $\frac{1}{2\sigma^2}$  does not affect the location of the minimum. We must have that  $\beta$  minimizes

$$\sum_i (y_i - \mathbf{x}_i^T \beta)^2.$$

We can write all this more conveniently using vectors and matrices. Write  $\mathbf{y}$  for the vector

$$\begin{pmatrix} y_1 \\ y_2 \\ \dots \\ y_n \end{pmatrix}$$

and  $\mathcal{X}$  for the matrix

$$\begin{pmatrix} \mathbf{x}_1^T \\ \mathbf{x}_2^T \\ \dots \mathbf{x}_n^T \end{pmatrix}.$$

Then we want to minimize

$$(\mathbf{y} - \mathcal{X}\beta)^T (\mathbf{y} - \mathcal{X}\beta)$$

which means that we must have

$$\mathcal{X}^T \mathcal{X}\beta - \mathcal{X}^T \mathbf{y} = 0.$$

For reasonable choices of features, we could expect that  $\mathcal{X}^T \mathcal{X}$  — which should strike you as being a lot like a covariance matrix — has full rank. If it does, which is the usual case, this equation is easy to solve. If it does not, there is more to do, which we will do in section 13.2.4.

**Remember this:** *The vector of coefficients  $\beta$  is usually estimated using a least-squares procedure.*

Listing 13.1: R code used for the linear regression example of worked example 13.1

```
efd<-read.table('efrontable.txt',header=TRUE)
# the table has the form
#N1 Ah Bh Ch N2 At Bt Ct
# now we need to construct a new dataset
hor<-stack(efd, select=2:4)
tim<-stack(efd, select=6:8)
foo<-data.frame(time=tim[, c("values")],
                hormone=hor[, c("values")])
foo.lm<-lm(hormone~time, data=foo)
plot(foo)
abline(foo.lm)
```

#### Worked example 13.1 *Simple Linear Regression with R*

Regress the hormone data against time for all the devices in the Efron example.

**Solution:** This example is mainly used to demonstrate how to regress in R. There is sample code in listing 13.1. The summary in the listing produces a great deal of information (try it). Most of it won't mean anything to you yet. You can get a figure by doing `plot(foo.lm)`, but these figures will not mean anything yet, either. In the code, I've shown how to plot the data and a line on top of it.

#### 13.1.2 Residuals and R-squared

Assume we have produced a regression by solving

$$\mathcal{X}^T \mathcal{X} \beta - \mathcal{X}^T \mathbf{y} = 0.$$

We cannot expect that  $\mathcal{X}\beta$  is the same as  $\mathbf{y}$ . Instead, there is likely to be some error. The **residual** is the vector

$$\mathbf{e} = \mathbf{y} - \mathcal{X}\beta$$

which gives the difference between the true value and the model's prediction at each point. The **mean square error** is

$$m = \frac{\mathbf{e}^T \mathbf{e}}{N}$$

and this gives the average of the squared error of prediction on the training examples.

**Procedure: 13.1** *Linear Regression with a Normal Model*

We have a dataset containing  $N$  pairs  $(\mathbf{x}_i, y_i)$ . Each  $x_i$  is a  $d$ -dimensional explanatory vector, and each  $y_i$  is a single dependent variable. We assume that each data point conforms to the model

$$y_i = \mathbf{x}_i^T \beta + \xi_i$$

where  $\xi_i$  is a normal random variable with mean 0 and unknown variance. Write  $\mathbf{y}$  for the vector

$$\begin{pmatrix} y_1 \\ y_2 \\ \dots \\ y_n \end{pmatrix}$$

and  $\mathcal{X}$  for the matrix

$$\begin{pmatrix} \mathbf{x}_1^T \\ \mathbf{x}_2^T \\ \dots \\ \mathbf{x}_n^T \end{pmatrix}.$$

We estimate  $\hat{\beta}$  (the value of  $\beta$ ) by solving the linear system

$$\mathcal{X}^T \mathcal{X} \hat{\beta} - \mathcal{X}^T \mathbf{y} = 0.$$

The residuals are

$$\mathbf{e} = \mathbf{y} - \mathcal{X} \hat{\beta}.$$

We have that  $\mathbf{e}^T \mathbf{1} = 0$ . The **mean square error** is given by

$$m = \frac{\mathbf{e}^T \mathbf{e}}{N}.$$

There is an important quantitative measure of how good a regression is. The dependent variable has some variance (unless it's a constant, which makes prediction easy). If our model is of any use, it should explain some aspects of the value of the dependent variable. This means that the variance of the residual should be smaller than the variance of the dependent variable. If the model made perfect predictions, then the variance of the residual should be zero.

We can formalize all this in a relatively straightforward way. We will ensure that  $\mathcal{X}$  always has a column of ones in it, so that the regression can have a non-zero y-intercept. We now fit a model

$$\mathbf{y} = \mathcal{X} \beta + \mathbf{e}$$

(where  $\mathbf{e}$  is the vector of residual values) by choosing  $\beta$  such that  $\mathbf{e}^T \mathbf{e}$  is minimized.

Then we get some useful technical results.

**Useful Facts: 13.1** *Regression*

We write  $\mathbf{y} = \mathcal{X}\beta + \mathbf{e}$ , where  $\mathbf{e}$  is the residual. Assume  $\mathcal{X}$  has a column of ones, and  $\beta$  is chosen to minimize  $\mathbf{e}^T \mathbf{e}$ . Then we have

1.  $\mathbf{e}^T \mathcal{X} = \mathbf{0}$ , i.e. that  $\mathbf{e}$  is orthogonal to any column of  $\mathcal{X}$ . This is because, if  $\mathbf{e}$  is not orthogonal to some column of  $\mathcal{X}$ , we can increase or decrease the  $\beta$  term corresponding to that column to make the error smaller. Another way to see this is to notice that  $\beta$  is chosen to minimize  $\mathbf{e}^T \mathbf{e}$ , which is  $(\mathbf{y} - \mathcal{X}\beta)^T (\mathbf{y} - \mathcal{X}\beta)$ . Now because this is a minimum, the gradient with respect to  $\beta$  is zero, so  $(\mathbf{y} - \mathcal{X}\beta)^T (-\mathcal{X}) = -\mathbf{e}^T \mathcal{X} = \mathbf{0}$ .
2.  $\mathbf{e}^T \mathbf{1} = 0$  (recall that  $\mathcal{X}$  has a column of all ones, and apply the previous result).
3.  $\mathbf{1}^T (\mathbf{y} - \mathcal{X}\beta) = 0$  (same as previous result).
4.  $\mathbf{e}^T \mathcal{X}\beta = 0$  (first result means that this is true).

Indexing: Regression

Indexing: Regression

Now  $\mathbf{y}$  is a one dimensional dataset arranged into a vector, so we can compute  $\text{mean}(\{y\})$  and  $\text{var}[y]$ . Similarly,  $\mathcal{X}\beta$  is a one dimensional dataset arranged into a vector (its elements are  $\mathbf{x}_i^T \beta$ ), as is  $\mathbf{e}$ , so we know the meaning of mean and variance for each. We have a particularly important result:

$$\text{var}[y] = \text{var}[\mathcal{X}\beta] + \text{var}[e].$$

This is quite easy to show, with a little more notation. Write  $\bar{\mathbf{y}} = (\frac{1}{N})(\mathbf{1}^T \mathbf{y})\mathbf{1}$  for the vector whose entries are all  $\text{mean}(\{y\})$ ; similarly for  $\bar{\mathbf{e}}$  and for  $\bar{\mathcal{X}\beta}$ . We have

$$\text{var}[y] = (1/N)(\mathbf{y} - \bar{\mathbf{y}})^T (\mathbf{y} - \bar{\mathbf{y}})$$

and so on for  $\text{var}[e_i]$ , etc. Notice from the facts that  $\bar{\mathbf{y}} = \bar{\mathcal{X}\beta}$ . Now

$$\begin{aligned} \text{var}[y] &= (1/N) \left( [\mathcal{X}\beta - \bar{\mathcal{X}\beta}] + [\mathbf{e} - \bar{\mathbf{e}}] \right)^T \left( [\mathcal{X}\beta - \bar{\mathcal{X}\beta}] + [\mathbf{e} - \bar{\mathbf{e}}] \right) \\ &= (1/N) \left( [\mathcal{X}\beta - \bar{\mathcal{X}\beta}]^T [\mathcal{X}\beta - \bar{\mathcal{X}\beta}] + 2[\mathbf{e} - \bar{\mathbf{e}}]^T [\mathcal{X}\beta - \bar{\mathcal{X}\beta}] + [\mathbf{e} - \bar{\mathbf{e}}]^T [\mathbf{e} - \bar{\mathbf{e}}] \right) \\ &= (1/N) \left( [\mathcal{X}\beta - \bar{\mathcal{X}\beta}]^T [\mathcal{X}\beta - \bar{\mathcal{X}\beta}] + [\mathbf{e} - \bar{\mathbf{e}}]^T [\mathbf{e} - \bar{\mathbf{e}}] \right) \\ &\quad \text{because } \bar{\mathbf{e}} = \mathbf{0} \text{ and } \mathbf{e}^T \mathcal{X}\beta = 0 \text{ and } \mathbf{e}^T \mathbf{1} = 0 \\ &= \text{var}[\mathcal{X}\beta] + \text{var}[e]. \end{aligned}$$

This is extremely important, because it allows us to think about a regression as explaining variance in  $\mathbf{y}$ . As we are better at explaining  $\mathbf{y}$ ,  $\text{var}[e]$  goes down. In

turn, a natural measure of the goodness of a regression is what percentage of the variance of  $y$  it explains. This is known as  $R^2$  (the r-squared measure). We have

$$R^2 = \frac{\text{var}[\mathbf{x}_i^T \beta]}{\text{var}[y_i]}$$

which gives some sense of how well the regression explains the training data.

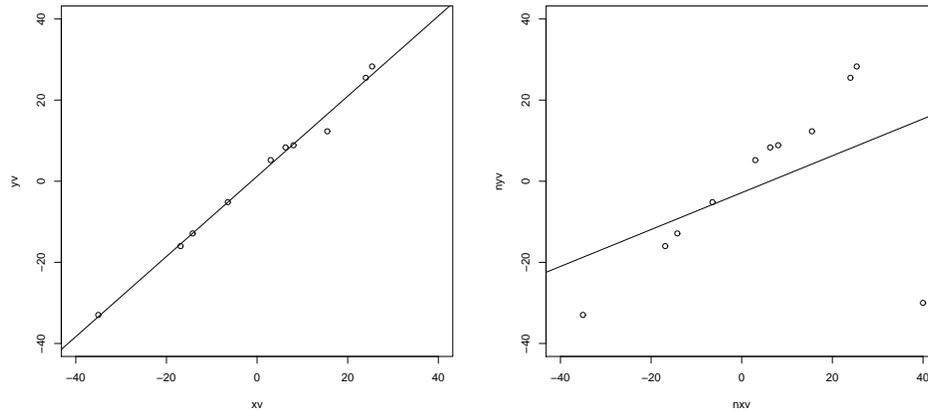


FIGURE 13.4: *On the left*, a synthetic dataset with one independent and one explanatory variable, with the regression line plotted. Notice the line is close to the data points, and its predictions seem likely to be reliable. *On the right*, the result of adding a single outlying datapoint to that dataset. The regression line has changed significantly, because the regression line tries to minimize the sum of squared vertical distances between the data points and the line. Because the outlying datapoint is far from the line, the squared vertical distance to this point is enormous. The line has moved to reduce this distance, at the cost of making the other points further from the line.

Good predictions result in high values of  $R^2$ , and a perfect model will have  $R^2 = 1$  (which doesn't usually happen). For example, the regression of figure 13.1 has an  $R^2$  value of 0.87. Figure 262 shows some regressions we have seen before with  $R^2$  values on top. Notice that if you look at the summary that R provides for a linear regression, it will offer you *two* estimates of the value for  $R^2$ . These estimates are obtained in ways that try to account for (a) the amount of data in the regression, and (b) the number of variables in the regression. For our purposes, the differences between these numbers and the  $R^2$  I defined are not significant. For the figures, I computed  $R^2$  as I described in the text above, but if you substitute one of R's numbers nothing terrible will happen.

**Remember this:** *The quality of predictions made by a regression can be evaluated by looking at the fraction of the variance in the dependent variable that is explained by the regression. This number is called  $R^2$ , and lies between zero and one; regressions with larger values make better predictions.*

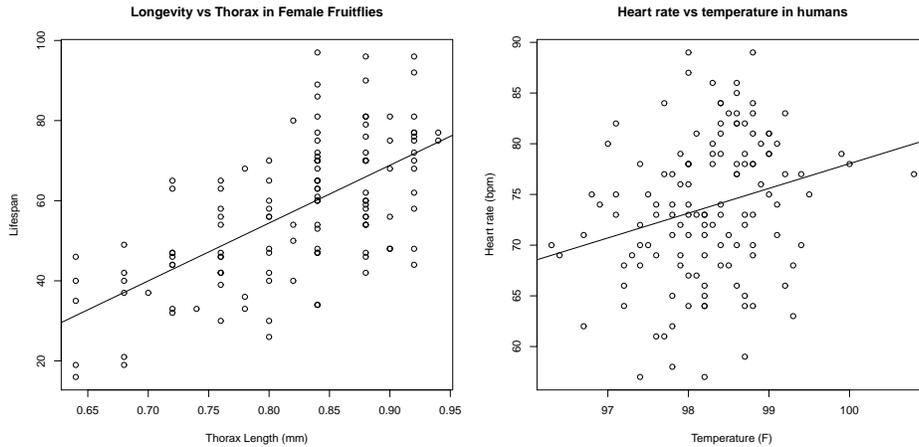


FIGURE 13.5: The **top row** shows the regressions of figure 13.2, which generally give quite good predictions, with  $R^2$  superimposed. The **bottom row** shows the regressions of figure 13.3, which generally give quite bad predictions, with  $R^2$  superimposed.

## 13.2 PRODUCING GOOD LINEAR REGRESSIONS

Linear regression is useful, but it isn't magic. Some regressions make poor predictions (recall the regressions of figure 13.3). Here is an extreme example, where a linear regression is very bad, and also very easily fixed. This requires some context. Some words are used very often in text; most are used seldom. The dataset for this example consists of counts of the number of times a word appears for the 100 most common words in Shakespeare's printed works. It was originally collected from a concordance, and has been used to attack a variety of interesting questions, including an attempt to assess how many words Shakespeare knew. This is hard, because he likely knew many words that he didn't use in his works, so one can't just count. If you look at the plot of Figure 13.8, you can see that a linear regression of count (the number of times a word is used) against rank (how common a word is, 1-100) is not really useful. The most common words are used very often, and the number of times a word is used falls off very sharply as one looks at less common words. You can see this effect in the scatter plot of residual against dependent variable in Figure 13.8 — the residual depends rather strongly on the dependent variable.

This is an extreme example that illustrates how poor linear regressions can be. As we shall see in section ??, a quite simple trick markedly improves the regression.

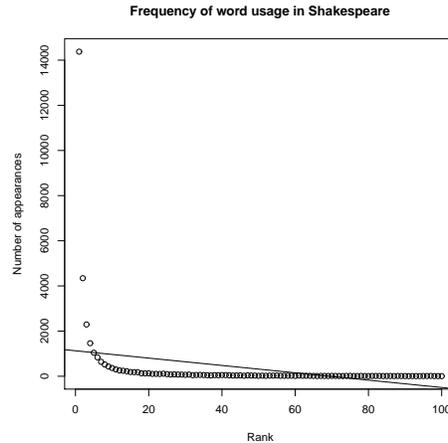


FIGURE 13.6: *On the left*, word count plotted against rank for the 100 most common words in Shakespeare, using a dataset that comes with R (called “bard”, and quite likely originating in an unpublished report by J. Gani and I. Saunders). I show a regression line too. This is a poor fit by eye, and the  $R^2$  is poor, too ( $R^2 = 0.1$ ).

### 13.2.1 Checking a Linear Regression

Some regressions will work well, but others might not. For example, regressing the first digit of your telephone number against the length of your foot won’t work that well. There are several straightforward tests to tell whether a regression is working.

You can **look at a plot** for a dataset with one explanatory variable and one dependent variable. You plot the data on a scatter plot, then plot the model as a line on that scatterplot. Just looking at the picture can be informative (compare Figure 13.2 and Figure 13.3). But we need more powerful tools, so we can deal with models that are harder to plot — for example, models that use more than one explanatory variable.

You can check if the regression **predicts a constant**. This is usually a bad sign. You can check this by looking at the predictions for each of the training data items. If the variance of these predictions is small compared to the variance of the independent variable, the regression isn’t working well. If you have only one explanatory variable, then you can plot the regression line. If the line is horizontal, or close, then the value of the explanatory variable makes very little contribution to the prediction. This suggests that there is no particular relationship between the explanatory variable and the independent variable.

You can check if **the residual isn’t random**. We assumed that  $y - \mathbf{x}^T \beta$  was a zero-mean normal random variable with fixed variance. Our model means that the value of the residual vector should not depend on the corresponding  $y$ -value. If it does, this is good evidence that there is a problem. Looking at a scatter plot

of  $\mathbf{e}$  against  $\mathbf{y}$  will often reveal trouble in a regression (Figure ??). In the case of Figure ??, the trouble is caused by a few outlying data points severely affecting the regression. We will discuss how to identify and deal with such points in Section 13.2. Once they have been removed, the regression improves markedly (Figure ??).

**Remember this:** *Linear regressions can make bad predictions. You can check for trouble by: evaluating  $R^2$ ; looking at a plot; looking to see if the regression makes a constant prediction; or checking whether the residual is random. Other strategies exist, but are beyond the scope of this book.*

### 13.2.2 Finding Problem Data Points

Outlying data points can significantly weaken the usefulness of a regression. We need to identify data points that might be a problem, and then resolve how to deal with them. One possibility is that they are true outliers — someone recorded a data item wrong, or they represent an effect that just doesn't occur all that often. Another is that they are important data, and our linear model may not be good enough. If the data points really are outliers, we can ignore them; if they aren't, we may be able to improve the regression by transforming features or by finding a new explanatory variable.

When we construct a regression, we are solving for the  $\beta$  that minimizes  $\sum_i (y_i - \mathbf{x}_i^T \beta)^2$ , equivalently for the  $\beta$  that produces the smallest value of  $\sum_i e_i^2$ . This means that residuals with large value can have a very strong influence on the outcome — we are squaring that large value, resulting in an enormous value. Generally, many residuals of medium size will have a smaller cost than one large residual and the rest tiny. As figure 265 illustrates, this means that a data point that lies far from the others can swing the regression line significantly.

This creates a problem, because data points that are clearly wrong (sometimes called **outliers**) can also have the highest influence on the outcome of the regression. Compare Figure ?? and Figure ?? to see this effect for a simple case. When we have only one explanatory variable, there's an easy method to spot problem data points. We produce a scatter plot and a regression line, and the difficulty is usually obvious. In particularly tricky cases, printing the plot and using a see-through ruler to draw a line by eye can help (if you use an opaque ruler, you may not see some errors). When there are more explanatory variables, we need more powerful visualization tools, which are beyond the scope of this text.

**Remember this:** *Outliers can affect linear regressions significantly. Usually, if you can plot the regression, you can look for outliers by eyeballing the plot. Other methods exist, but are beyond the scope of this text.*

## 13.2.3 Adding Explanatory Variables

Plotting a linear regression against one explanatory variable can suggest that some variables are missing. For example, if you look at the perch data of Figure 13.2, the fitted line looks quite good, but the data points look as though they might be willing to follow a curve. We can get a curve quite easily. Our current model gives the weight as a linear function of the length with a noise term (which we wrote  $y_i = \beta_1 x_i + \beta_0 + \xi_i$ ). But we could expand this model to incorporate other functions of the length. In fact, it's quite surprising that the weight of a fish should be predicted by its length. If the fish doubled in each direction, say, its weight should go up by a factor of eight. The success of our regression suggests that fish do not just scale in each direction as they grow. But we might try the model  $y_i = \beta_2 x_i^2 + \beta_1 x_i + \beta_0 + \xi_i$ . This is easy to do. The  $i$ 'th row of the matrix  $\mathcal{X}$  currently looks like  $[x_i, 1]$ . We build a new matrix  $\mathcal{X}^{(b)}$ , where the  $i$ 'th row is  $[x_i^2, x_i, 1]$ , and proceed as before. This gets us a new model, which is plotted in Figure 13.7.

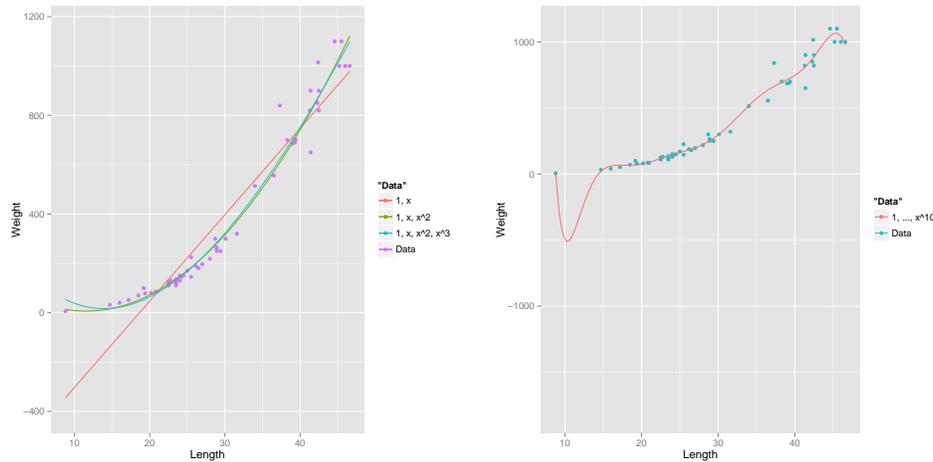


FIGURE 13.7: *On the left, several different models predicting fish weight from length. The line uses the explanatory variables 1 and  $x_i$ ; and the curves use other monomials in  $x_i$  as well, as shown by the legend. This allows the models to predict curves that lie closer to the data. It is important to understand that, while you can make a curve go closer to the data by inserting monomials, that doesn't mean you necessarily have a better model. On the right, I have used monomials up to  $x_i^{10}$ . This curve lies very much closer to the data points than any on the other side, at the cost of some very odd looking wiggles inbetween data points. I can't think of any reason that these structures would come from true properties of fish, and it would be hard to trust predictions from this model.*

You should notice that it can be quite easy to add a lot of explanatory variables (in the case of the fish, I tried  $x_i^3$  as well). However, it's hard to decide

whether the regression has actually gotten better. The least-squares error *on the training data* will *never* go up when you add new explanatory variables. This is easy to see, because you could always use a coefficient of zero with the new variables and get back the previous regression. However, the models that you choose are likely to produce worse and worse predictions as you add explanatory variables. Knowing when to stop can be tough, though it's sometimes obvious that the model is untrustworthy (Figure 13.7). Dealing with all this is beyond the scope of this book.

When there is more than one explanatory variable, it's tough to plot the regression nicely, so you can't rely on your eye to detect "curviness" in data. An alternative is to plot the value of the residual against the value predicted by the model. Remember that we modelled data as a linear function of the explanatory variables, together with noise. For the training data, we know the values taken by the noise (this is the residual). If these values *depend* on the value predicted by the model, then there might be a missing explanatory variable.

**Remember this:** *You can insert explanatory variables into  $\mathbf{x}$ . Doing so is likely to improve the performance of the regression on the training data, but may well make it perform worse on test data. Knowing when to stop is hard.*

#### 13.2.4 Regularizing Linear Regressions

One occasionally important difficulty is that the explanatory variables might be significantly correlated. If they are, then it will generally be easy to predict one explanatory variable from another. This means that  $\mathcal{X}^T \mathcal{X}$  may have some very small eigenvalues (because there is a vector  $\mathbf{u}$  so that  $\mathcal{X}\mathbf{u}$  is small; this means that  $\mathbf{u}^T \mathcal{X}^T \mathcal{X} \mathbf{u}$  must be small).

These small eigenvalues lead to bad predictions. If  $\mathcal{X}^T \mathcal{X}$  has a small eigenvalue, then there is some vector  $\mathbf{v}$  such that  $\mathcal{X}^T \mathcal{X} \mathbf{v}$  is small, or, equivalently, that the matrix can turn large vectors into small ones; but that means that  $(\mathcal{X}^T \mathcal{X})^{-1}$  will turn some small vectors into big ones. In turn, this means that small errors in  $\mathbf{y}$  — which are likely inevitable — will result in big errors in  $\beta$ . This could cause trouble in two ways. If we are looking at  $\beta$  to tell which explanatory variables are important, then large errors in  $\beta$  will be a problem. And if we are trying to predict new  $y$  values, we expect that errors in  $\beta$  turn into errors in prediction.

An important and useful way to suppress these errors is to regularize, using the same trick we saw in the case of classification. Instead of choosing  $\beta$  to minimize

$$\sum_i (y_i - \mathbf{x}_i^T \beta)^2 = (\mathbf{y} - \mathcal{X}\beta)^T (\mathbf{y} - \mathcal{X}\beta)$$

we minimize

$$\sum_i (y_i - \mathbf{x}_i^T \beta)^2 + \lambda \beta^T \beta = (\mathbf{y} - \mathcal{X}\beta)^T (\mathbf{y} - \mathcal{X}\beta) + \lambda \beta^T \beta$$

where  $\lambda > 0$  is a constant. We choose  $\lambda$  in the same way we used for classification; split the training set into a training piece and a validation piece, train for different values of  $\lambda$ , and test the resulting regressions on the validation piece. We choose the  $\lambda$  that yields the smallest validation error. Notice we could use multiple splits, and average over the splits.

This helps, because to solve for  $\beta$  we must solve the equation

$$(\mathcal{X}^T \mathcal{X} + \lambda \mathcal{I})\beta = \mathcal{X}^T \mathbf{y}$$

(obtained by differentiating with respect to  $\beta$  and setting to zero) and the smallest eigenvalue of the matrix  $(\mathcal{X}^T \mathcal{X} + \lambda \mathcal{I})$  will be at least  $\lambda$ .

**Remember this:** *The performance of a regression can be improved by regularizing, particularly if some explanatory variables are correlated. The procedure is similar to that used for classification.*

### 13.2.5 Transforming Variables

Sometimes the data isn't in a form that leads to a good linear regression. In this case, transforming explanatory variables, the dependent variable, or both can lead to big improvements. Figure ?? shows one example, based on the idea of word frequencies. Some words are used very often in text; most are used seldom. The dataset for this figure consists of counts of the number of times a word occurred for the 100 most common words in Shakespeare's printed works. It was originally collected from a concordance, and has been used to attack a variety of interesting questions, including an attempt to assess how many words Shakespeare knew. This is hard, because he likely knew many words that he didn't use in his works, so one can't just count.

Notice that a linear regression of count (the number of times a word is used) against rank (how common a word is, 1-100) is not really useful. It doesn't model the very high frequencies with which common words are used. However, if we regress log-count against log-rank, we get a very good fit indeed. This suggests that Shakespeare's word usage (at least for the 100 most common words) is consistent with **Zipf's law**. This gives the relation between frequency  $f$  and rank  $r$  for a word as

$$f \propto \frac{1}{r^s}$$

where  $s$  is a constant characterizing the distribution. Our linear regression suggests that  $s$  is approximately 1.67 for this data.

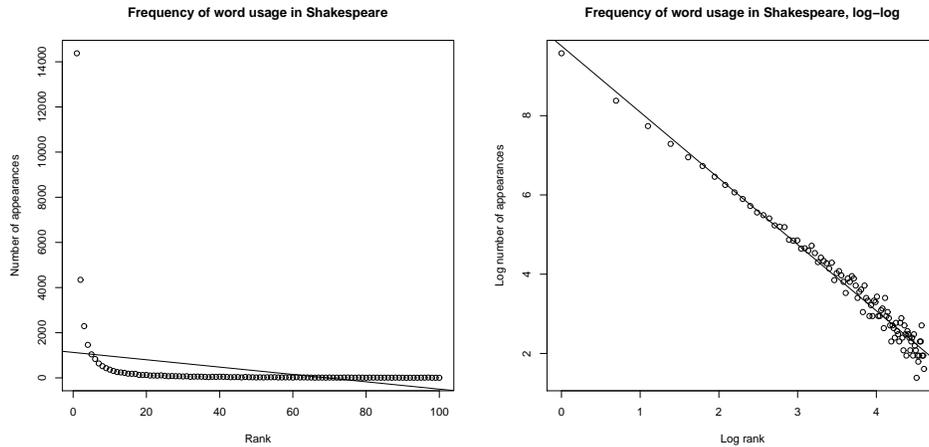


FIGURE 13.8: *On the left*, word count plotted against rank for the 100 most common words in Shakespeare, using a dataset that comes with R (called “bard”, and quite likely originating in an unpublished report by J. Gani and I. Saunders). I show a regression line too. This is a poor fit by eye, confirmed by Figure ???. *On the right*, log word count plotted against log rank for the 100 most common words in Shakespeare, using a dataset that comes with R (called “bard”, and quite likely originating in an unpublished report by J. Gani and I. Saunders). The regression line is very close to the data.

In some cases, the natural logic of the problem will suggest variable transformations that improve regression performance. For example, one could argue that humans have approximately the same density, and so that weight should scale as the cube of height; in turn, this suggests that one regress weight against the cube root of height. Generally, shorter people tend not to be scaled versions of taller people, so the cube root might be too aggressive. For example, the body mass index divides weight by the square (rather than the cube) of height. An appropriate transformation has attracted a great deal of interest over the last century and a half, perhaps because there are many confounding variables. For example, the body mass index tends to regard muscular individuals as obese.

**Remember this:** *The performance of a regression can be improved by transforming variables. Transformations usually follow from looking at plots, or thinking about the logic of the problem*

## 13.3 EXPLOITING YOUR NEIGHBORS FOR REGRESSION

Nearest neighbors can clearly predict a number for a query example — you find the closest training example, and report its number. This would be one way to use nearest neighbors for regression, but it isn't terribly effective. One important difficulty is that the regression prediction is piecewise constant (Figure 13.9). If there is an immense amount of data, this may not present major problems, because the steps in the prediction will be small and close together. But it's not generally an effective use of data.

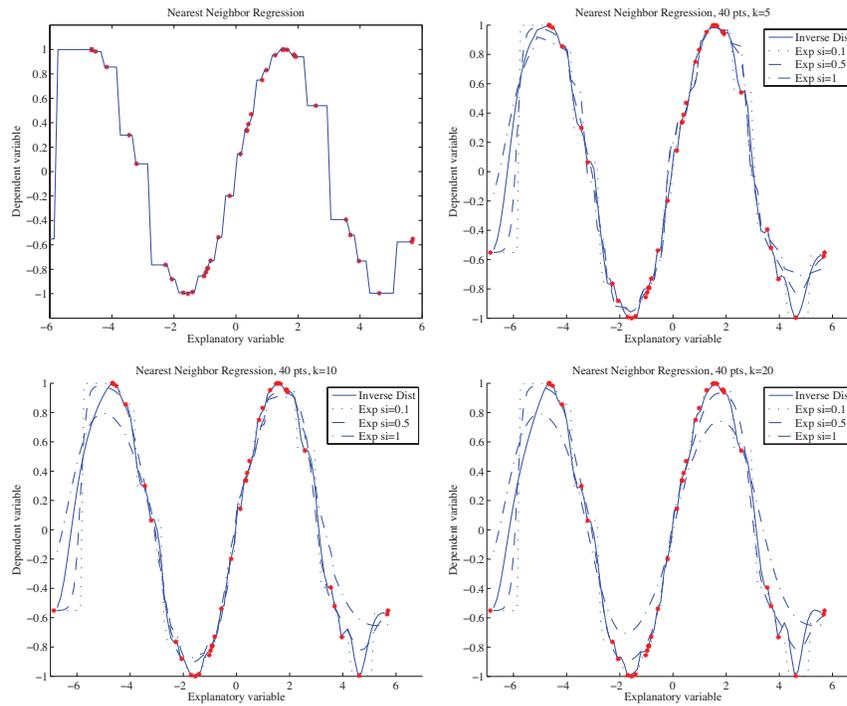


FIGURE 13.9: Different forms of nearest neighbors regression, predicting  $y$  from a one-dimensional  $x$ , using a total of 40 training points. **Top left:** reporting the nearest neighbor leads to a piecewise constant function. **Top right:** improvements are available by forming a weighted average of the five nearest neighbors, using inverse distance weighting or exponential weighting with three different scales. Notice if the scale is small, then the regression looks a lot like nearest neighbors, and if it is too large, all the weights in the average are nearly the same (which leads to a piecewise constant structure in the regression). **Bottom left and bottom right** show that using more neighbors leads to a smoother regression.

A more effective strategy is to find several nearby training examples, and use them to produce an estimate. This approach can produce very good regression estimates, because every prediction is made by training examples that are near to the query example. However, producing a regression estimate is expensive, because

for every query one must find the nearby training examples.

Write  $\mathbf{x}$  for the query point, and assume that we have already collected the  $N$  nearest neighbors, which we write  $\mathbf{x}_i$ . Write  $y_i$  for the value of the dependent variable for the  $i$ 'th of these points. Notice that some of these neighbors could be quite far from the query point. We don't want distant points to make as much contribution to the model as nearby points. This suggests forming a weighted average of the predictions of each point. Write  $w_i$  for the weight at the  $i$ 'th point. Then the estimate is

$$y_{pred} = \frac{\sum_i w_i y_i}{\sum_i w_i}.$$

A variety of weightings are reasonable choices. Write  $d_i = \|\mathbf{x} - \mathbf{x}_i\|$  for the distance between the query point and the  $i$ 'th nearest neighbor. Then inverse distance weighting uses  $w_i = 1/d_i$ . Alternatively, we could use an exponential function to strongly weight down more distant points, using

$$w_i = \exp\left(\frac{-d_i^2}{2\sigma^2}\right).$$

We will need to choose a scale  $\sigma$ , which can be done by cross-validation. Hold out some examples, make predictions at the held out examples using a variety of different scales, and choose the scale that gives the best held-out error. Alternatively, if there are enough nearest neighbors, we could form a distance weighted linear regression, then predict the value at the query point from that regression.

Each of these strategies presents some difficulties when  $\mathbf{x}$  has high dimension. In that case, it is usual that the nearest neighbor is a lot closer than the second nearest neighbor. If this happens, then each of these weighted averages will boil down to evaluating the dependent variable at the nearest neighbor (because all the others will have very small weight in the average).

**Remember this:** *Nearest neighbors can be used for regression. In the simplest approach, you find the nearest neighbor to your feature vector, and take that neighbor's number as your prediction. More complex approaches smooth predictions over multiple neighbors.*

### 13.3.1 Using your Neighbors to Predict More than a Number

Linear regression takes some features and predicts a number. But in practice, one often wants to predict something more complex than a number. For example, I might want to predict a parse tree (which has combinatorial structure) from a sentence (the explanatory variables). As another example, I might want to predict a map of the shadows in an image (which has spatial structure) against an image (the explanatory variables). As yet another example, I might want to predict which direction to move the controls on a radio-controlled helicopter (which have to be

moved together) against a path plan and the current state of the helicopter (the explanatory variables).

Looking at neighbors is a very good way to solve such problems. The general strategy is relatively simple. We find a large collection of pairs of training data. Write  $\mathbf{x}_i$  for the explanatory variables for the  $i$ 'th example, and  $\mathbf{y}_i$  for the dependent variable in the  $i$ 'th example. This dependent variable could be anything — it doesn't need to be a single number. It might be a tree, or a shadow map, or a word, or anything at all. I wrote it as a vector because I needed to choose some notation.

In the simplest, and most general, approach, we obtain a prediction for a new set of explanatory variables  $\mathbf{x}$  by (a) finding the nearest neighbor and then (b) producing the dependent variable for that neighbor. We might vary the strategy slightly by using an approximate nearest neighbor. If the dependent variables have enough structure that it is possible to summarize a collection of different dependent variables, then we might recover the  $k$  nearest neighbors and summarize their dependent variables. How we summarize rather depends on the dependent variables. For example, it is a bit difficult to imagine the average of a set of trees, but quite straightforward to average images. If the dependent variable was a word, we might not be able to average words, but we can vote and choose the most popular word. If the dependent variable is a vector, we can compute either distance weighted averages or a distance weighted linear regression.

**Example: Filling Large Holes with Whole Images** Many different kinds of user want to remove things from images or from video. Art directors might like to remove unattractive telephone wires; restorers might want to remove scratches or marks; there's a long history of government officials removing people with embarrassing politics from publicity pictures (see the fascinating examples in ?); and home users might wish to remove a relative they dislike from a family picture. All these users must then find something to put in place of the pixels that were removed.

If one has a large hole in a large image, we may not be able to just extend a texture to fill the hole. Instead, entire objects might need to appear in the hole (Figure 13.10). There is a straightforward, and extremely effective, way to achieve this. We match the image to a large collection of images, to find the nearest neighbors (the details of the distance function are below). This yields a set of example images where all the pixels we didn't want to replace are close to those of the query image. From these, we choose one, and fill in the pixels from that image.

There are several ways to choose. If we wish to do so automatically, we could use the example with the smallest distance to the image. Very often, an artist is involved, and then we could prepare a series of alternatives — using, perhaps, the  $k$  closest examples — then show them to the artist, who will choose one. This method, which is very simple to describe, is extremely effective in practice.

It is straightforward to get a useful distance between images. We have an image with some missing pixels, and we wish to find nearby images. We will assume that all images are the same size. If this isn't in fact the case, we could either crop or resize the example images. A good measure of similarity between two images  $\mathcal{A}$  and  $\mathcal{B}$  can be measured by forming the **sum of squared differences** (or

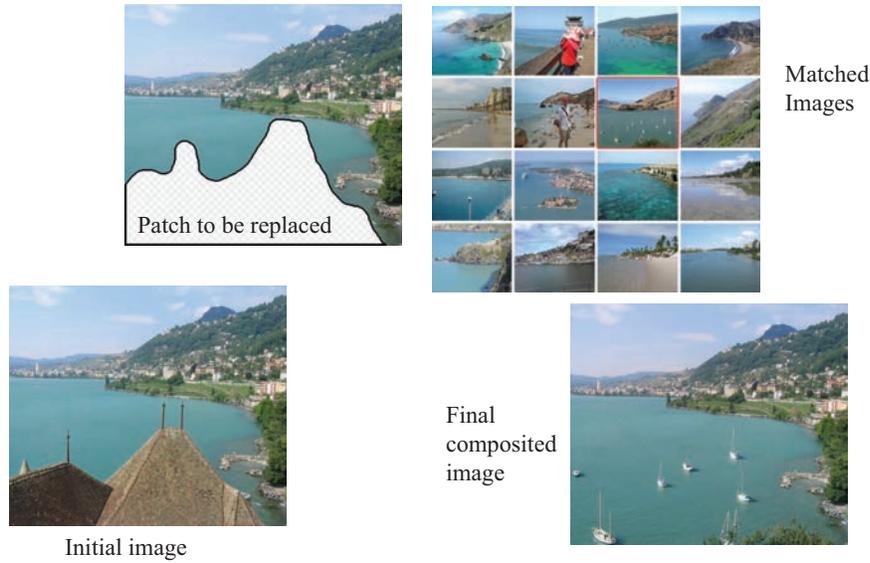


FIGURE 13.10: We can fill large holes in images by matching the image to a collection, choosing one element of the collection, then cutting out an appropriate block of pixels and putting them into the hole in the query image. In this case, the hole has been made by an artist, who wishes to remove the roofline from the view. Notice how there are a range of objects (boats, water) that have been inserted into the hole. These objects are a reasonable choice, because the overall structures of the query and matched image are largely similar — getting an image that matches most of the query image supplies enough context to ensure that the rest “makes sense”.

SSD) of corresponding pixel values. You should think of an image as an array of pixels. If the images are grey-level images, then each pixel contains a single number, encoding the grey-level. If the images are color images, then each pixel (usually!) contains three values, one encoding the red-level, one encoding the green-level, and one encoding the blue-level. The SSD is computed as

$$\sum_{(i,j)} (\mathcal{A}_{ij} - \mathcal{B}_{ij})^2$$

where  $i$  and  $j$  range over all pixels. If the images are grey-level images, then by  $(\mathcal{A}_{ij} - \mathcal{B}_{ij})^2$ , I mean the squared difference between grey levels; if they are color images, then this means the sum of squared differences between red, green and blue values. This distance is small when the images are similar, and large when they are different (it is essentially the length of the difference vector).

Now we don't know some of the pixels in the query image. Write  $\mathcal{K}$  for the set of pixels around a point whose values are known, and  $\#\mathcal{K}$  for the size of this set. We can now use

$$\frac{1}{\#\mathcal{K}} \sum_{(i,j) \in \mathcal{K}} (\mathcal{A}_{ij} - \mathcal{B}_{ij})^2.$$

Filling in the pixels requires some care. One does not usually get the best results by just copying the missing pixels from the matched image into the hole. Instead, it is better to look for a good **seam**. We search for a curve enclosing the missing pixels which (a) is reasonably close to the boundary of the missing pixels and (b) gives a good boundary between the two images. A good boundary is one where the query image (on one side) is “similar to” the matched image (on the other side). A good sense of similarity requires that pixels match well, and that image gradients crossing the boundary tend to match too.

**Remember this:** *Nearest neighbors can be used to predict more than numbers. Examples include parse trees, blocks of pixels, and so on.*

## 13.4 YOU SHOULD

## 13.4.1 remember:

New term: Regression . . . . .	384
Definition: Regression . . . . .	385
New term: residual . . . . .	385
New term: explanatory variables . . . . .	386
New term: dependent variable . . . . .	386
New term: training examples . . . . .	386
New term: test examples . . . . .	386
Definition: Linear regression . . . . .	386
Estimating $\beta$ . . . . .	389
New term: residual . . . . .	390
New term: mean square error . . . . .	390
New term: mean square error . . . . .	391
Useful facts: Regression . . . . .	392
$R^2$ evaluates the quality of predictions made by a regression . . . . .	394
Linear regressions can fail. . . . .	396
New term: outliers . . . . .	396
Outliers can affect linear regressions significantly. . . . .	397
You can increase the number of explanatory variables. . . . .	398
You can regularize a regression . . . . .	399
New term: Zipf's law . . . . .	399
Transforming variables is useful . . . . .	400
Nearest neighbors can be used for regression. . . . .	402
New term: sum of squared differences . . . . .	403
New term: seam . . . . .	405
Nearest neighbors can predict structures. . . . .	405

## APPENDIX: DATA

Batch A		Batch B		Batch C	
Amount of Hormone	Time in Service	Amount of Hormone	Time in Service	Amount of Hormone	Time in Service
25.8	99	16.3	376	28.8	119
20.5	152	11.6	385	22.0	188
14.3	293	11.8	402	29.7	115
23.2	155	32.5	29	28.9	88
20.6	196	32.0	76	32.8	58
31.1	53	18.0	296	32.5	49
20.9	184	24.1	151	25.4	150
20.9	171	26.5	177	31.7	107
30.4	52	25.8	209	28.5	125

TABLE 13.1: A table showing the amount of hormone remaining and the time in service for devices from lot A, lot B and lot C. The numbering is arbitrary (i.e. there's no relationship between device 3 in lot A and device 3 in lot B). We expect that the amount of hormone goes down as the device spends more time in service, so cannot compare batches just by comparing numbers.

PROBLEMS

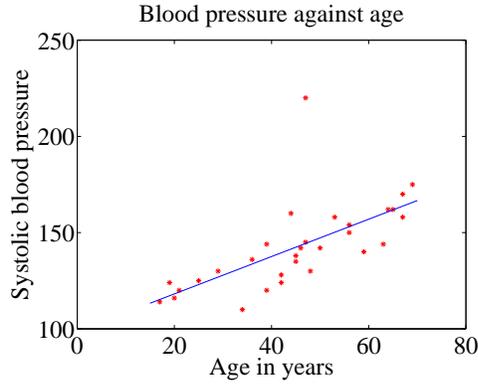


FIGURE 13.11: A regression of blood pressure against age, for 30 data points.

13.1. Figure 13.11 shows a linear regression of systolic blood pressure against age. There are 30 data points.

- (a) Write  $e_i = y_i - \mathbf{x}_i^T \beta$  for the residual. What is the  $\text{mean}(\{e\})$  for this regression?
- (b) For this regression,  $\text{var}(\{y\}) = 509$  and the  $R^2$  is 0.4324. What is  $\text{var}(\{e\})$  for this regression?
- (c) How well does the regression explain the data?
- (d) What could you do to produce better predictions of blood pressure (without actually measuring blood pressure)?

13.2. In this exercise, I will show that the prediction process of chapter 3 (see page 63) is a linear regression with two independent variables. Assume we have  $N$  data items which are 2-vectors  $(x_1, y_1), \dots, (x_N, y_N)$ , where  $N > 1$ . These could be obtained, for example, by extracting components from larger vectors. As usual, we will write  $\hat{x}_i$  for  $x_i$  in normalized coordinates, and so on. The correlation coefficient is  $r$  (this is an important, traditional notation).

- (a) Show that  $r = \text{mean}(\{(x - \text{mean}(\{x\}))(y - \text{mean}(\{y\}))\}) / (\text{std}(x)\text{std}(y))$ .
- (b) Now write  $s = \frac{\text{std}(y)}{\text{std}(x)}$ . Now assume that we have an  $x_o$ , for which we wish to predict a  $y$  value. Show that the value of the prediction obtained using the method of page 64 is

$$sr(x_o - \text{mean}(\{x\})) + \text{mean}(\{y\}).$$

- (c) Show that  $sr = \text{mean}(\{(xy)\}) - \text{mean}(\{x\})\text{mean}(\{y\})$ .
- (d) Now write

$$\mathcal{X} = \begin{pmatrix} x_1 & 1 \\ x_2 & 1 \\ \dots & \dots \\ x_n & 1 \end{pmatrix} \text{ and } \mathbf{Y} = \begin{pmatrix} y_1 \\ y_2 \\ \dots \\ y_n \end{pmatrix}.$$

The coefficients of the linear regression will be  $\hat{\beta}$ , where  $\mathcal{X}^T \mathcal{X} \hat{\beta} = \mathcal{X}^T \mathbf{Y}$ . Show that

$$\mathcal{X}^T \mathcal{X} = N \begin{pmatrix} \text{mean}(\{x^2\}) & \text{mean}(\{x\}) \\ \text{mean}(\{x\}) & 1 \end{pmatrix}$$

- (e) Now show that  $\text{var}(\{x\}) = \text{mean}(\{(x - \text{mean}(\{x}))^2\}) = \text{mean}(\{x^2\}) - \text{mean}(\{x\})^2$ .
- (f) Now show that  $\text{std}(x)\text{std}(y)\text{corr}(\{(x, y)\}) = \text{mean}(\{(x - \text{mean}(\{x\}))(y - \text{mean}(\{y\}))\})$ .

## CHAPTER 14

# Math Resources

### 14.1 USEFUL MATERIAL ABOUT MATRICES

#### Terminology:

- A matrix  $\mathcal{M}$  is **symmetric** if  $\mathcal{M} = \mathcal{M}^T$ . A symmetric matrix is necessarily square.
- We write  $\mathcal{I}$  for the identity matrix.
- A matrix is **diagonal** if the only non-zero elements appear on the diagonal. A diagonal matrix is necessarily symmetric.
- A symmetric matrix is **positive semidefinite** if, for any  $\mathbf{x}$  such that  $\mathbf{x}^T \mathbf{x} > 0$  (i.e. this vector has at least one non-zero component), we have  $\mathbf{x}^T \mathcal{M} \mathbf{x} \geq 0$ .
- A symmetric matrix is **positive definite** if, for any  $\mathbf{x}$  such that  $\mathbf{x}^T \mathbf{x} > 0$ , we have  $\mathbf{x}^T \mathcal{M} \mathbf{x} > 0$ .
- A matrix  $\mathcal{R}$  is **orthonormal** if  $\mathcal{R}^T \mathcal{R} = \mathcal{I} = \mathcal{I}^T = \mathcal{R} \mathcal{R}^T$ . Orthonormal matrices are necessarily square.

**Orthonormal matrices:** You should think of orthonormal matrices as rotations, because they do not change lengths or angles. For  $\mathbf{x}$  a vector,  $\mathcal{R}$  an orthonormal matrix, and  $\mathbf{u} = \mathcal{R}\mathbf{x}$ , we have  $\mathbf{u}^T \mathbf{u} = \mathbf{x}^T \mathcal{R}^T \mathcal{R} \mathbf{x} = \mathbf{x}^T \mathcal{I} \mathbf{x} = \mathbf{x}^T \mathbf{x}$ . This means that  $\mathcal{R}$  doesn't change lengths. For  $\mathbf{y}, \mathbf{z}$  both unit vectors, we have that the cosine of the angle between them is  $\mathbf{y}^T \mathbf{x}$ ; but, by the same argument as above, the inner product of  $\mathcal{R}\mathbf{y}$  and  $\mathcal{R}\mathbf{x}$  is the same as  $\mathbf{y}^T \mathbf{x}$ . This means that  $\mathcal{R}$  doesn't change angles, either.

**Eigenvectors and Eigenvalues:** Assume  $\mathcal{S}$  is a  $d \times d$  symmetric matrix,  $\mathbf{v}$  is a  $d \times 1$  vector, and  $\lambda$  is a scalar. If we have

$$\mathcal{S}\mathbf{v} = \lambda\mathbf{v}$$

then  $\mathbf{v}$  is referred to as an **eigenvector** of  $\mathcal{S}$  and  $\lambda$  is the corresponding **eigenvalue**. Matrices don't have to be symmetric to have eigenvectors and eigenvalues, but the symmetric case is the only one of interest to us.

In the case of a symmetric matrix, the eigenvalues are real numbers, and there are  $d$  distinct eigenvectors that are normal to one another, and can be scaled to have unit length. They can be stacked into a matrix  $\mathcal{U} = [\mathbf{v}_1, \dots, \mathbf{v}_d]$ . This matrix is orthonormal, meaning that  $\mathcal{U}^T \mathcal{U} = \mathcal{I}$ . This means that there is a diagonal matrix  $\Lambda$  such that

$$\mathcal{S}\mathcal{U} = \mathcal{U}\Lambda.$$

In fact, there is a large number of such matrices, because we can reorder the eigenvectors in the matrix  $\mathcal{U}$ , and the equation still holds with a new  $\Lambda$ , obtained by reordering the diagonal elements of the original  $\Lambda$ . There is no reason to keep track of this complexity. Instead, we adopt the convention that the elements of  $\mathcal{U}$  are always ordered so that the elements of  $\Lambda$  are sorted along the diagonal, with the largest value coming first.

**Diagonalizing a symmetric matrix:** This gives us a particularly important procedure. We can convert any symmetric matrix  $\mathcal{S}$  to a diagonal form by computing

$$\mathcal{U}^T \mathcal{S} \mathcal{U} = \Lambda.$$

This procedure is referred to as **diagonalizing** a matrix. Again, we assume that the elements of  $\mathcal{U}$  are always ordered so that the elements of  $\Lambda$  are sorted along the diagonal, with the largest value coming first. Diagonalization allows us to show that positive definiteness is equivalent to having all positive eigenvalues, and positive semidefiniteness is equivalent to having all non-negative eigenvalues.

**Factoring a matrix:** Assume that  $\mathcal{S}$  is symmetric and positive semidefinite. We have that

$$\mathcal{S} = \mathcal{U} \Lambda \mathcal{U}^T$$

and all the diagonal elements of  $\Lambda$  are non-negative. Now construct a diagonal matrix whose diagonal entries are the positive square roots of the diagonal elements of  $\Lambda$ ; call this matrix  $\Lambda^{(1/2)}$ . We have  $\Lambda^{(1/2)} \Lambda^{(1/2)} = \Lambda$  and  $(\Lambda^{(1/2)})^T = \Lambda^{(1/2)}$ . Then we have that

$$\mathcal{S} = (\mathcal{U} \Lambda^{(1/2)}) (\Lambda^{(1/2)} \mathcal{U}^T) = (\mathcal{U} \Lambda^{(1/2)}) (\mathcal{U} \Lambda^{(1/2)})^T$$

so we can factor  $\mathcal{S}$  into the form  $\mathcal{X} \mathcal{X}^T$  by computing the eigenvectors and eigenvalues.

#### 14.1.1 Approximating A Symmetric Matrix

Assume we have a  $k \times k$  symmetric matrix  $\mathcal{T}$ , and we wish to construct a matrix  $\mathcal{A}$  that approximates it. We require that (a) the rank of  $\mathcal{A}$  is precisely  $r < k$  and (b) the approximation should minimize the **Frobenius norm**, that is,

$$\|(\mathcal{T} - \mathcal{A})\|_F^2 = \sum_{ij} (T_{ij} - A_{ij})^2.$$

It turns out that there is a straightforward construction that yields  $\mathcal{A}$ .

The first step is to notice that if  $\mathcal{U}$  is orthonormal and  $\mathcal{M}$  is any matrix, then

$$\|\mathcal{U} \mathcal{M}\|_F = \|\mathcal{M} \mathcal{U}\|_F = \|\mathcal{M}\|_F.$$

This is true because  $\mathcal{U}$  is a rotation (as is  $\mathcal{U}^T = \mathcal{U}^{-1}$ ), and rotations do not change the length of vectors. So, for example, if we write  $\mathcal{M}$  as a table of row vectors  $\mathcal{M} = [\mathbf{m}_1, \mathbf{m}_2, \dots, \mathbf{m}_k]$ , then  $\mathcal{U} \mathcal{M} = [\mathcal{U} \mathbf{m}_1, \mathcal{U} \mathbf{m}_2, \dots, \mathcal{U} \mathbf{m}_k]$ . Now  $\|\mathcal{M}\|_F^2 = \sum_{j=1}^k \|\mathbf{m}_j\|^2$ , so  $\|\mathcal{U} \mathcal{M}\|_F^2 = \sum_{i=1}^k \|\mathcal{U} \mathbf{m}_k\|^2$ . But rotations do not change lengths, so  $\|\mathcal{U} \mathbf{m}_k\|^2 = \|\mathbf{m}_k\|^2$ , and so  $\|\mathcal{U} \mathcal{M}\|_F = \|\mathcal{M}\|_F$ . To see the result for the case of  $\mathcal{M} \mathcal{U}$ , just think of  $\mathcal{M}$  as a table of row vectors.

Notice that, if  $\mathcal{U}$  is the orthonormal matrix whose columns are eigenvectors of  $\mathcal{T}$ , then we have

$$\|(\mathcal{T} - \mathcal{A})\|_F^2 = \|\mathcal{U}^T(\mathcal{T} - \mathcal{A})\mathcal{U}\|_F^2.$$

Now write  $\Lambda_r$  for  $\mathcal{U}^T \mathcal{A} \mathcal{U}$ , and  $\Lambda$  for the diagonal matrix of eigenvalues of  $\mathcal{T}$ . Then we have

$$\|(\mathcal{T} - \mathcal{A})\|_F^2 = \|\Lambda - \Lambda_A\|_F^2,$$

an expression that is easy to solve for  $\Lambda_A$ . We know that  $\Lambda$  is diagonal, so the best  $\Lambda_A$  is diagonal, too. The rank of  $\mathcal{A}$  must be  $r$ , so the rank of  $\Lambda_A$  must be  $r$  as well. To get the best  $\Lambda_A$ , we keep the  $r$  largest diagonal values of  $\Lambda$ , and set the rest to zero;  $\Lambda_A$  has rank  $r$  because it has only  $r$  non-zero entries on the diagonal, and every other entry is zero.

Now to recover  $\mathcal{A}$  from  $\Lambda_A$ , we know that  $\mathcal{U}^T \mathcal{U} = \mathcal{U} \mathcal{U}^T = \mathcal{I}$  (remember,  $\mathcal{I}$  is the identity). We have  $\Lambda_A = \mathcal{U}^T \mathcal{A} \mathcal{U}$ , so

$$\mathcal{A} = \mathcal{U} \Lambda_A \mathcal{U}^T.$$

We can clean up this representation in a useful way. Notice that only the first  $r$  columns of  $\mathcal{U}$  (and the corresponding rows of  $\mathcal{U}^T$ ) contribute to  $\mathcal{A}$ . The remaining  $k - r$  are each multiplied by one of the zeros on the diagonal of  $\Lambda_A$ . Remember that, by convention,  $\Lambda$  was sorted so that the diagonal values are in descending order (i.e. the largest value is in the top left corner). We now keep only the top left  $r \times r$  block of  $\Lambda_A$ , which we write  $\Lambda_r$ . We then write  $\mathcal{U}_r$  for the  $k \times r$  matrix consisting of the first  $r$  columns of  $\mathcal{U}$ . Then

$$\mathcal{A} = \mathcal{U}_r \Lambda_r \mathcal{U}_r^T$$

This is so useful a result, I have displayed it in a box; you should remember it.

**Procedure: 14.1** *Approximating a symmetric matrix with a low rank matrix*

Assume we have a symmetric  $k \times k$  matrix  $\mathcal{T}$ . We wish to approximate  $\mathcal{T}$  with a matrix  $\mathcal{A}$  that has rank  $r < k$ . Write  $\mathcal{U}$  for the matrix whose columns are eigenvectors of  $\mathcal{T}$ , and  $\Lambda$  for the diagonal matrix of eigenvalues of  $\mathcal{A}$  (so  $\mathcal{A} \mathcal{U} = \mathcal{U} \Lambda$ ). Remember that, by convention,  $\Lambda$  was sorted so that the diagonal values are in descending order (i.e. the largest value is in the top left corner).

Now construct  $\Lambda_r$  from  $\Lambda$  by setting the  $k - r$  smallest values of  $\Lambda$  to zero, and keeping only the top left  $r \times r$  block. Construct  $\mathcal{U}_r$ , the  $k \times r$  matrix consisting of the first  $r$  columns of  $\mathcal{U}$ . Then

$$\mathcal{A} = \mathcal{U}_r \Lambda_r \mathcal{U}_r^T$$

is the best possible rank  $r$  approximation to  $\mathcal{T}$  in the Frobenius norm.

Now if  $\mathcal{A}$  is positive semidefinite (i.e. if at least the  $r$  largest eigenvalues of  $\mathcal{T}$  are non-negative), then we can factor  $\mathcal{A}$  as in the previous section. This yields a procedure to approximate a symmetric matrix by factors. This is so useful a result, I have displayed it in a box; you should remember it.

**Procedure: 14.2** *Approximating a symmetric matrix with low dimensional factors*

Assume we have a symmetric  $k \times k$  matrix  $\mathcal{T}$ . We wish to approximate  $\mathcal{T}$  with a matrix  $\mathcal{A}$  that has rank  $r < k$ . We assume that at least the  $r$  largest eigenvalues of  $\mathcal{T}$  are non-negative. Write  $\mathcal{U}$  for the matrix whose columns are eigenvectors of  $\mathcal{T}$ , and  $\Lambda$  for the diagonal matrix of eigenvalues of  $\mathcal{A}$  (so  $\mathcal{A}\mathcal{U} = \mathcal{U}\Lambda$ ). Remember that, by convention,  $\Lambda$  was sorted so that the diagonal values are in descending order (i.e. the largest value is in the top left corner).

Now construct  $\Lambda_r$  from  $\Lambda$  by setting the  $k - r$  smallest values of  $\Lambda$  to zero and keeping only the top left  $r \times r$  block. Construct  $\Lambda_r^{(1/2)}$  by replacing each diagonal element of  $\Lambda$  with its positive square root. Construct  $\mathcal{U}_r$ , the  $k \times r$  matrix consisting of the first  $r$  columns of  $\mathcal{U}$ . Then write  $\mathcal{V} = (\mathcal{U}_r \Lambda_r^{(1/2)})$

$$\mathcal{A} = \mathcal{V}\mathcal{V}^T$$

is the best possible rank  $r$  approximation to  $\mathcal{T}$  in the Frobenius norm.

**Error functions and Gaussians:** The **error function** is defined by

$$\operatorname{erf}(x) = \frac{2}{\sqrt{\pi}} \int_0^x e^{-t^2} dt$$

and programming environments can typically evaluate the error function. This fact is made useful to us by a simple change of variables. We get

$$\frac{1}{\sqrt{2\pi}} \int_0^x e^{-\frac{u^2}{2}} du = \frac{1}{\sqrt{\pi}} \int_0^{\frac{x}{\sqrt{2}}} e^{-t^2} dt = \frac{1}{2} \operatorname{erf}\left(\frac{x}{\sqrt{2}}\right).$$

A particularly useful manifestation of this fact comes by noticing that

$$\frac{1}{\sqrt{2\pi}} \int_{-\infty}^0 e^{-\frac{t^2}{2}} dt = 1/2$$

(because  $\frac{1}{\sqrt{2\pi}} e^{-\frac{u^2}{2}}$  is a probability density function, and is symmetric about 0). As a result, we get

$$\frac{1}{\sqrt{2\pi}} \int_{-\infty}^x e^{-\frac{t^2}{2}} dt = 1/2(1 + \operatorname{erf}\left(\frac{x}{\sqrt{2}}\right)).$$

**Inverse error functions:** We sometimes wish to know the value of  $x$  such that

$$\frac{1}{\sqrt{2\pi}} \int_{-\infty}^x e^{-t^2/2} dt = p$$

for some given  $p$ . The relevant function of  $p$  is known as the **probit function** or the **normal quantile function**. We write

$$x = \Phi(p).$$

The probit function  $\Phi$  can be expressed in terms of the **inverse error function**. Most programming environments can evaluate the inverse error function (which is the inverse of the error function). We have that

$$\Phi(p) = \sqrt{2} \operatorname{erf}^{-1}(2p - 1).$$

One problem we solve with some regularity is: choose  $u$  such that

$$\int_{-u}^u \frac{1}{\sqrt{2\pi}} \exp(-x^2/2) dx = p.$$

Notice that

$$\begin{aligned} \frac{p}{2} &= \frac{1}{\sqrt{2\pi}} \int_0^u e^{-t^2/2} dt \\ &= \frac{1}{2} \operatorname{erf}\left(\frac{u}{\sqrt{2}}\right) \end{aligned}$$

so that

$$u = \sqrt{2} \operatorname{erf}^{-1}(p).$$

# Index

- 0-1 loss, *see* classifier
- 3D bar chart, 47
- a useful expression for covariance
  - useful facts, 137
- a useful expression for variance
  - useful facts, 136
- all-vs-all, *see* classifier
- antichance, *see* gambler's fallacy
- average, *see* mean
- bag, 349
- bagging, 349
- bar chart, 13
- Bayes' rule, 125
- bayes' rule
  - definition, 126
- Bayesian inference, 232
- bayesian inference
  - definition, 232
- bernoulli random variable
  - definition, 162
  - useful facts, 163
- beta distribution
  - useful facts, 170
- biased random walk, 193
- bimodal, 28
- binomial distribution
  - definition, 164
  - useful facts, 164, 165
- binomial distribution for large  $n$ 
  - useful facts, 181
- bootstrap, 265
- bootstrap replicates, 265
- boxplot, 34
- categorical, 10
- central limit theorem, 175
- chebyshev's inequality
  - definition, 140
- class confusion matrix, 339
- class-conditional histograms, 16
- classifier, 328
  - 0-1 loss, 329
  - cost of misclassification, 329
  - cross-validation, 330
  - definition, 328
  - error rate, 330
  - estimating performance, 329
    - leave-one-out cross-validation, 330
  - false negative, 329
  - false positive, 329
  - loss function, 329
  - multiclass
    - all-vs-all, 339
    - one-vs-all, 339
  - nearest neighbor classifier
    - k-nearest neighbor classifier, 353
  - nearest neighbors, 352
- classifiers
  - hinge loss, 332
- clustering, 357
  - complete-link clustering, 362
  - dendrogram, 362, 363
  - group average clustering, 362
  - grouping and agglomeration, 362
  - partitioning and division, 362
  - single-link clustering, 362
  - using K-means, 367
- clusters, 359
- color constancy, 313
- conditional histograms, 16
- conditional probability formulas
  - useful facts, 108
- confidence interval, 257
- confidence intervals, 255

- conjugate, 234
- consistency, 230
- continuous, 10
- correlation, 56, 58
- correlation coefficient, 58
  - definition, 59
- covariance
  - definition, 135, 293
- covariance ellipses, 361
- covariance matrix
  - definition, 295
- cross-validation, *see* classifier
- cumulative distribution of a discrete random variable
  - definition, 124
- decision boundary, 331
- decision forest, 341
- decision function, 341
- decision tree, 150, 340
- decorrelation, 364
- definition
  - bayes' rule, 126
  - bayesian inference, 232
  - bernoulli random variable, 162
  - binomial distribution, 164
  - chebyshev's inequality, 140
  - classifier, 328
  - correlation coefficient, 59
  - covariance, 135, 293
  - covariance matrix, 295
  - cumulative distribution of a discrete random variable, 124
  - discrete random variable, 122
  - expectation, 132, 134
  - expected value, 131, 133
  - geometric distribution, 163
  - independent random variables, 127
  - indicator functions, 141
  - interquartile range, 27
  - likelihood, 223
  - linear regression, 386
  - map estimate, 232
  - markov's inequality, 140
  - mean, 17
  - mean or expected value, 134
  - median, 25
  - multinomial distribution, 167
  - normal data, 33
  - normal distribution, 174
  - pairwise independent, independent and conditionally independent events, 112
  - percentile, 26
  - poisson distribution, 168
  - probability distribution of a discrete random variable, 124
  - quartiles, 27
  - regression, 385
  - sample space, 78
  - standard coordinates, 32
  - standard deviation, 19, 139
  - standard normal data, 32
  - standard normal distribution, 173
  - the log-likelihood of a dataset under a model, 227
  - the maximum likelihood principle, 223
  - variance, 23, 135
  - weak law of large numbers, 144
- dendrogram, *see* clustering
- dependent variable, 386
- descent direction, 334
- descriptive statistics, 139
- diagonal, 410
- diagonalizing, 411
- discrete random variable
  - definition, 122
- eigenvalue, 410
- eigenvector, 410
- empirical distribution, 139, 265
- entropy, 345
- error bars, 257
- error function, 176, 413
- error rate, *see* classifier
- event, 82
  - useful facts, 83
- expectation
  - definition, 132, 134
- expectations are linear.
  - useful facts, 134
- expected value, 131
  - definition, 131, 133

- explanatory variables, 386
- exponential distribution, 171
  - useful facts, 172
- false negative, *see* classifier
- false positive, *see* classifier
- feature vectors, 328
- filtering, 241
- Frobenius norm, 411
- gambler's fallacy, 112
  - antichance, 112
- gambler's ruin, 191
- gamma distribution
  - useful facts, 171
- gaussian distributions, 175
- generalizing badly, 330
- geometric distribution
  - definition, 163
  - useful facts, 163
- gradient descent, 334
- heat map, 45
- hinge loss, *see* classifiers
- histogram, 13
- hypothesis, 269
- hypothesis testing, 221
- IID, 223
- iid samples, 143
- independent, 80, 94, 126
- independent and identically distributed, 223
- independent events
  - useful facts, 94
- independent identically distributed samples, 143
- independent random variables
  - definition, 127
- indicator function, 141
- indicator functions
  - definition, 141
- Inference, 221
- information gain, 346
- intensity, 168
- interquartile range, 26
  - definition, 27
  - useful facts, 27
- interval estimates, 221
- inverse error function, 259, 414
- irreducible, 197
- joint, 239
- joint probability distribution, 125
- k-means, *see* clustering, 369
- k-nearest neighbor classifier, *see* classifier
- latent variable, 67
- leave-one-out cross-validation, *see* classifier
- likelihood, 223
  - definition, 223
- Likert scales, 319
- line search, 334
- linear regression
  - definition, 386
- loss function, *see* classifier
- MAP estimate, 232
- map estimate
  - definition, 232
- marginal probability, 126
- Markov chain, 194
- markov's inequality
  - definition, 140
- maximum a posteriori estimate, 232
- maximum likelihood principle, 223
- mean, 16
  - average, 16
  - definition, 17
  - useful facts, 19
- mean or expected value
  - definition, 134
- mean square error, 390, 391
- median, 24
  - definition, 25
  - useful facts, 25
- mode, 28
- multidimensional scaling, 317
- multimodal, 28
- multinomial distribution
  - definition, 167
- nearest neighbor classifier, *see* classifier

- normal data, 33
  - definition, 33
- normal distribution, 174, 175
  - definition, 174
  - useful facts, 175
- normal distributions are conjugate
  - useful facts, 237
- normal posteriors
  - useful facts, 239
- normal quantile function, 414
- normal random variable, 175
- normal random variables
  - useful facts, 176, 257
- normalizing, 130
- normalizing constant, 130
  
- odds, 148
- one-sided, 280
- one-vs-all, *see* classifier
- online updating of normal posteriors
  - useful facts, 242
- ordinal, 10
- orthonormal, 410
- outlier, 24
- outliers, 396
- overfitting, 330
  
- p-value, 272
- pairwise independent, independent and conditionally independent events
  - definition, 112
- parameters of a multivariate normal distribution
  - useful facts, 360
- percentile
  - definition, 26
- pie chart, 45
- point estimate, 221
- poisson distribution
  - definition, 168
  - useful facts, 168
- Poisson point process, 169
- population, 244
- population mean, 244
- positive definite, 410
- positive semidefinite, 410
- posterior, 232
  
- principal component, 309
- principal coordinate analysis, 317
- prior probability distribution, 232
- probability density function, 129
- probability density functions
  - useful facts, 129
- probability distribution of a discrete random variable
  - definition, 124
- probability of events
  - useful facts, 85
- probit function, 414
- prosecutor's fallacy, 109
  
- quartiles
  - definition, 27
  
- raw Google matrix, 207
- Regression, 384
- regression
  - definition, 385
  - useful facts, 392
- residual, 385, 390
  
- sample, 143, 244
- sample mean, 244
- sample means and population means
  - useful facts, 245
- sample space
  - definition, 78
  - useful facts, 79
- samples, 143
- scatter plot, 50
- seam, 405
- selection bias, 330
- skew, 29
- smoothing, 304
- SSD, *see* sum of squared differences
- stacked bar chart, 45
- standard coordinates, 31
  - definition, 32
- standard deviation, 19, 138
  - definition, 19, 139
  - useful facts, 23
- standard error, 248
- standard normal curve, 32
- standard normal data, 32

- definition, 32
- standard normal distribution, 173
  - definition, 173
  - useful facts, 174
- standard normal random variable, 174
- stationary distribution, 197
- statistic, 255
- steplength, 335
- stochastic gradient descent, 335
- stochastic matrices, 194
- Student's t-distribution, 280
- sum of squared differences, 403
- sums and differences of normal random variables
  - useful facts, 276
- symmetric, 410
- tails, 28
- test examples, 386
- test set, 330
- test statistic, 273, 278
- the log-likelihood of a dataset under a model
  - definition, 227
- the maximum likelihood principle
  - definition, 223
- the sample mean
  - useful facts, 248
- training examples, 386
- transition probabilities, 193, 194
- trial, 143
- two-sided, 280
- uniform distribution, 169
- uniform random variable, 169
- unimodal, 28
- useful facts
  - a useful expression for covariance, 137
  - a useful expression for variance, 136
  - bernoulli random variable, 163
  - beta distribution, 170
  - binomial distribution, 164, 165
  - binomial distribution for large  $n$ , 181
  - conditional probability formulas, 108
  - event, 83
  - expectations are linear., 134
  - exponential distribution, 172
  - gamma distribution, 171
  - geometric distribution, 163
  - independent events, 94
  - interquartile range, 27
  - mean, 19
  - median, 25
  - normal distribution, 175
  - normal distributions are conjugate, 237
  - normal posteriors, 239
  - normal random variables, 176, 257
  - online updating of normal posteriors, 242
  - parameters of a multivariate normal distribution, 360
  - poisson distribution, 168
  - probability density functions, 129
  - probability of events, 85
  - regression, 392
  - sample means and population means, 245
  - sample space, 79
  - standard deviation, 23
  - standard normal distribution, 174
  - sums and differences of normal random variables, 276
  - the sample mean, 248
  - variance, 24, 135
  - variance and covariance, 137
  - whitening a dataset, 303
- utility, 153
- validation set, 330
- variance, 23
  - definition, 23, 135
  - useful facts, 24, 135
- variance and covariance
  - useful facts, 137
- Vector quantization, 376
- weak law of large numbers
  - definition, 144
- whitening, 364
- whitening a dataset
  - useful facts, 303
- z-test, 280

Zipf's law, 399