

Contents

Chapter 1	Introduction	7
Chapter 2	Tracking: Fundamental Notions	11
2.1	General Observations	12
2.2	Tracking by Detection	15
2.2.1	Background Subtraction	17
2.2.2	Deformable Templates	18
2.2.2.1	Robustness	21
2.2.2.2	The Hausdorff Distance	22
2.3	Tracking using Flow	24
2.3.1	Optic Flow	25
2.3.2	Image stabilization	26
2.3.3	Cardboard people	26
2.3.4	Building Flow Templates	29
2.3.5	Flow models from kinematic models	30
2.3.5.1	Tracking a Derivative Flow Model	32
2.3.5.2	The Flow Model from the Chain Rule	32
2.3.5.3	Rigid-body Transformations	34

2 Contents

2.4	Tracking with Probability	35
2.4.1	Linear Dynamics and the Kalman Filter	36
2.4.2	Data Association	36
2.4.3	Multiple Modes	38
Chapter 3 Tracking: Relations between 3D and 2D		43
3.1	Kinematic Inference with Multiple Views	44
3.2	Lifting to 3D	47
3.2.1	Geometric Ambiguity and Lifting by Kinematic Inference	48
3.2.2	Lifting by Minimization	52
3.2.3	Lifting by Regression	53
3.2.3.1	Lifting using the Nearest Neighbour	53
3.2.3.2	Snippets and Cameras	55
3.2.3.3	Regressing Pose against the Image	57
3.2.3.4	Disambiguation with the Immediate Past	58
3.3	Multiple Modes, Randomized Search and Human Tracking	59
3.3.1	Randomized Search with Particle Filters	60
3.3.1.1	Importance Sampling	60
3.3.1.2	Partitioned Sampling	63
3.3.1.3	Lower Dimensional State Models	64
3.3.1.4	Probabilistic Searches of the Posterior	65
3.3.1.5	Annealing	66
3.3.2	Multiple Probes from Covariance Analysis	67
Chapter 4 Tracking: Data Association for Human Tracking		77
4.1	Detecting Humans	78
4.1.1	Finding People by Matching Static Templates	79
4.1.2	Templates that include Motion	81
4.1.3	Traditional Parts	83
4.1.3.1	Discriminative Approaches	84
4.1.3.2	Generative Approaches	85
4.1.3.3	Mixed Approaches	87
4.1.4	Parts as Codebooks	88

4.2	Tracking by Matching Revisited	89
4.2.1	Likelihood	90
4.2.2	Whole Object Templates	91
4.2.3	Traditional Parts	93
4.2.3.1	Complex Spatio-Temporal Relations	93
4.2.3.2	Known Appearance Models	94
4.2.3.3	Inferred Generative Appearance Models	95
4.2.3.4	Inferred Discriminative Appearance Models	97
4.2.4	Parts as Codebooks	98
4.3	Evaluation	100
Chapter 5 Motion Synthesis		105
5.1	Fundamental Notions	106
5.1.1	The Motion Capture Process	106
5.1.1.1	Engineering Issues	107
5.1.1.2	Cleanup and Skeletonization	108
5.1.1.3	Configuration Representations	109
5.1.1.4	Skinning	109
5.1.2	Footskate	110
5.1.3	Inverse Kinematics	111
5.1.4	Resolving Kinematic Ambiguities with Examples	113
5.1.5	Specifying a Motion Demand	115
5.2	Motion Signal Processing	116
5.2.1	Temporal Scaling and Alignment	117
5.2.2	Blending, Transitions and Filtering	117
5.2.2.1	Multi-way Blends	118
5.2.2.2	Transitions	119
5.2.2.3	Filtering	119
5.2.2.4	Physical Blends	119
5.2.2.5	What to Blend	119
5.2.2.6	Finding Similar Motions	120
5.2.2.7	Difficulties with Blending	121
5.3	Motion Graphs	121
5.3.1	Building a Motion Graph	122

4 Contents

5.3.1.1	Links by transitions	123
5.3.2	Searching a Motion Graph	125
5.3.3	How good is a Motion Graph?	129
5.4	Motion Primitives	130
5.4.1	Primitives by Segmenting and Clustering	131
5.4.1.1	Difficulties with Dimension Reduction	132
5.4.2	Linking Segmentation to the Primitive Model	133
5.5	Enriching a Motion Collection	135
5.5.1	Rearranging Existing Motions	135
5.5.1.1	Motion Editing	136
5.5.1.2	New Motions by Cut and Paste	136
5.5.1.3	Motion Fill-in by Nonparametric Regression	137
5.5.1.4	Motion Interpolation	138
5.5.2	Motion from Physical Considerations	138
5.5.2.1	Simplified Characters	141
5.5.2.2	Modified Physics	142
5.5.2.3	Reduced Dimensions	144
5.5.2.4	Modifying Existing Motions	145
Chapter 6 Describing Activities		147
6.1	What should an Activity Representation do?	148
6.1.1	Necessary Properties of an Activity Representation	149
6.1.2	What Data is Available?	150
6.2	Miscellaneous Methods	151
6.2.1	Activity Representation Methods based around Temporal Logics	151
6.2.2	Activity Representation Methods based on Templates	151
6.3	Activity Representation using Hidden Markov Models and Finite State Representations	152
6.3.1	The Speech Analogy	152
6.3.1.1	Finite State Transducers	153
6.3.1.2	Why should we Care?	154
6.3.2	Activity Recognition Methods based around HMM's	155
6.3.3	Sign Language Recognition	156

Chapter 7 Discussion	159
7.1 Representations	159
7.1.1 Is 3D Configuration Ambiguous?	159
7.1.1.1 Single Frames	160
7.1.1.2 Short Timescales	161
7.1.1.3 Long Timescales	162
7.1.1.4 Summary	163
7.1.2 What Representation of 3D Configuration should be Adopted?	163
7.1.3 What is the Status of Dynamical Models?	166
7.1.4 The Space of Human Motions	168
7.1.5 Is Human Tracking Multimodal?	169
7.2 Generalization	170
7.2.1 Which Motions are Human?	171
7.3 Resources	174
References	179

¹Figure from “Kinematic jump processes for monocular 3D human tracking”, Sminchisescu and Triggs, *Proc. Computer Vision and Pattern Recognition*, 2003 © 2003 IEEE

²Figure from “Covariance Scaled Sampling for Monocular 3D Body Tracking”, Sminchisescu and Triggs, *Proc. Computer Vision and Pattern Recognition*, 2001 © 2001 IEEE

1

Introduction

What follows is an incomplete review of the literature on computational aspects of human motion. We cover tracking, motion synthesis, and action recognition. This review is intended to expose the main ideas of the subject, rather than to present citations of all relevant, which might not even be possible. In our opinion, the core material in the two fields of understanding human motion and simulating human motion is shared — the great need is for a better understanding of how to represent motion usefully.

This area uses a stunning array of different techniques, and we have not really been able to introduce each tool in as full detail as one might wish. We intend to expand this review to the length of a graduate text, which should (at least!) have a much more detailed review of action recognition, extended discussion of techniques, and more detailed and nuanced discussion of animation by physical principles. If time, energy and skill allow — a very big “if” indeed — we will review relevant biomechanical information as well. We encourage readers of this review with suggestions for revision, correction, incorporation, deletion, and so on, to contact D.A. Forsyth, whose home page can easily be found using a search engine. It would be helpful if you would include the term “Motion Book” in the email subject; he will attempt to address and acknowledge any suggestions.

This review is aimed somewhat more at a computer vision audience than at a computer graphics audience, a balance we intend to correct; one has to start somewhere. We review tracking in three chapters, covering respectively basic ideas, the complexities created by inferring 3D body configuration from 2D information, and detection and data association. We cover motion synthesis and activity recognition each in single chapters. We expect that these chapters will expand. Finally, we have a discussion of specific issues in motion representation that attempts to bring to bear all we have seen on some important, and murky, questions of representation.

We see three core problems in the representation of human motion, at least from the perspective of applied computer science: **What motions look like human motion and what do not?** Human observers are remarkably sensitive to the quality of motion. Quite small changes in a motion representation can lead to the perception that the motion “isn’t right” or “isn’t human”. It would be most useful to have a set of tools that can predict whether a motion will prove acceptable to a human observer. Such tools are of obvious significance to the animation community. One benefit to the vision community would be the ability to tell whether a track is human or not; this seems unlikely to be of major impact, because one is unlikely to be presented with tracks that look human to simple tests but are not, in fact, human. A more important result of studying this question would be an increased understanding of dynamical models of motion.

How should activities be described and represented? Typical applications of motion synthesis increasingly need be able to produce motions that can be controlled to appear purposeful, and do so quickly and simply. This means that whatever produces motion demands — the game AI, the simulation environment or the simulation author — needs some vocabulary with which to communicate to the motion synthesis machinery. No canonical vocabulary is known. This situation parallels a difficulty in the computer vision community, which would like to report what is happening in a video sequence but does not possess a vocabulary in which to report it. We believe that a good working hypothesis is that a vocabulary suitable for controlling motion is also suitable for describing it. This idea is informal — it’s hard to be formal in this area — but useful nonetheless.

How is human behaviour affected by nearby objects? Generally, a description of what one is doing is significantly affected by the objects nearby.

For example, standing next to a bus sign and standing outside a door are different — not because they look different, but because the intention, as revealed by what is nearby, is different. The actual motions generated in response to some demand are affected by what is nearby — for example, moving across broken ground is different from walking over clear terrain. We shall have relatively little to say on this point, which is not much addressed in either the animation or the computer vision literature. This is not because it's not important, but because making progress on this issue probably requires a better understanding of the first two matters.

We have listed these questions in order of perceived difficulty (from rather nasty to very hard to discuss with current methods). These big questions are attended by a host of lesser, but important questions. Generally, we can say more about the lesser questions, and the discussion chapter contains a host of speculation and recommendation on them.

2

Tracking: Fundamental Notions

In a tracking problem, one has some measurements that appear at each tick of a (notional) clock, and, from these measurements, one would like to determine the state of the world. There are two important sources of information. First, measurements constrain the possible state of the world. Second, there are dynamical constraints — the state of the world cannot change arbitrarily from time to time. Tracking problems are of great practical importance. There are very good reasons to want to, say, track aircraft using radar returns (good summary histories include [63, 65, 215]; comprehensive reviews of technique in this context include [27, 44, 147]).

Not all measurements are informative. For example, if one wishes to track an aircraft — where state might involve pose, velocity and acceleration variables, and measurements might be radar returns giving distance and angle to the aircraft from several radar aerials — some of the radar returns measured might not come from the aircraft. Instead, they might be the result of noise, of other aircraft, of strips of foil dropped to confuse radar apparatus (**chaff** or **window**; see [215]), or of other sources. The problem of determining which measurements are informative and which are not is known as **data association**.

Data association is the dominant difficulty in tracking objects in video.

This is because so few of the very many pixels in each frame lie on objects of interest. It can be spectacularly difficult to tell which pixels in an image come from an object of interest and which do not. There are a very wide variety of methods for doing so, the details of which largely depend on the specifics of the application problem. Surprisingly, data association is not usually explicitly discussed in the computer vision tracking literature. However, whether a method is useful rests pretty directly on its success at data association — differences in other areas tend not to matter all that much in practice.

2.1 General Observations

The literature on tracking people is immense. Furthermore, the problem has quite different properties depending on precisely what kind of representation one wishes to recover. The most important variable appears to be spatial scale. At a **coarse scale**, people are **blobs**. For example, we might view a plaza from the window of a building or a mall corridor from a camera suspended from the ceiling. Each person occupies a small block of pixels, perhaps 10-100 pixels in total. While we should be able to tell where a person is, there isn't much prospect of determining where the arms and legs are. At this scale, we can expect to recover representations of **occupancy** — where people spend time, for example [478] — or of **patterns of activity** — how people move from place to place, and at what time, for example [423].

At a **medium scale**, people can be thought of as blobs with attached **motion fields**. For example, a television program of a soccer match, where each individual is usually 50-100 pixels high. In this case, one can tell where a person is. Arms and legs are still difficult to localize, because they cover relatively few pixels, and there is motion blur. However, the motion fields around the body yield some information as to how the person is moving. One could expect to be able to tell where a runner is in the phase of the run from this information — are the legs extended away from the body, or crossing?

At a **fine scale**, the arms and legs cover enough pixels to be detected, and one wants to report the configuration of the body. We usually refer to this case as **kinematic tracking**. At a fine spatial scale, one may be able to report such details as whether a person is picking up or handling an object. There are a variety of ways in which one could encode and report configuration, depending on the model adopted — is one to report the configuration of the

arms? the legs? the fingers? — and on whether these reports should be represented in 2D or in 3D. We will discuss various representations in greater detail later. We do not regard the choice of a 2D representation or a 3D representation as significant to anything but detail in the tracking problem; in particular, as we shall show, passing from 2D to 3D representations is relatively straightforward under most conditions. In particular, the choice of 2D or 3D representation does not affect ambiguity in any way. If we track using a 2D representation and then lift to 3D, it may be difficult to recover 3D configuration unambiguously. These ambiguities are also present in the case where we track a 3D representation, though are often unremarked.

Each scale appears to be useful, but there are no reliable rules of thumb for determining what scale is most useful for what application. For example, one could see ways to tell whether people are picking up objects at a coarse scale. Equally, one could determine patterns of activity from a fine scale. Finally, as we shall see later, some quite complex determinations about activity can be made at a surprisingly coarse scale. Tracking tends to be much more difficult at the fine scale, because one must manage more degrees of freedom and because arms and legs can be small, and can move rather fast.

In this review, we focus almost entirely on the fine scale; even so, space will not allow detailed discussion of all that has been done. Our choice of scale is dictated by the intuition that good fine-scale tracking will be an essential component of any method that can give general reports on what people are doing in video. There are distinctive features of this problem that make fine scale tracking difficult:

- **State dimension:** One typically requires a high dimensional state vector to describe the configuration of the body in a frame. For example, assume we describe a person using a 2D representation. Each of ten body segments (torso, head, upper and lower arms and legs) will be represented by a rectangle of fixed size (that differs from segment to segment). This representation will use an absolute minimum of 12 state variables (position and orientation for one rectangle, and relative orientation for every other). A more practical version of the representation allows the rectangles to slide with respect to one another, and so needs 27 state variables. Considerably more variables are required for 3D models.

- **Nasty dynamics:** There is good evidence that such motions as walking have predictable, low-dimensional structure [378, 395]. However, the body can move extremely fast, with large accelerations. These large accelerations mean that one can stop moving predictably very quickly — for example, jumping in the air during a walk. For straightforward mechanical reasons, the body parts that move fastest tend to be small and on one end of a long lever which has big muscles at the other end (forearms, fingers and feet, for example). This means that the body segments that the dynamical model fails to predict are going to be hard to find because they are small. As a result, accurate tracking of forearms can be very difficult.
- **Complex appearance phenomena:** In most applications one is tracking clothed people. Clothing can change appearance dramatically as it moves, because the forces the body applies to the clothing change, and so the pattern of folds, caused by buckling, changes. There are two important results. First, the pattern of occlusions of texture changes, meaning that the apparent texture of the body segment can change. Second, each fold will have a typical shading pattern attached, and these patterns move in the image as the folds move on the surface. Again, the result is that the apparent texture of the body segment changes. These effects can be seen in figure 2.4.
- **Data association:** There is usually no distinctive color or texture that identifies a person (which is why people are notoriously difficult to find in static images). One possible cue is that many body segments appear at a distinctive scale as extended regions with rather roughly parallel sides. This isn't too helpful, as there are many other sources of such regions (for example, the spines of books on a shelf). Textured backgrounds are a particularly rich source of false structures in edge maps. Much of what follows is about methods to handle data association problems for people tracking.

2.2 Tracking by Detection

Assume we have some form of template that can detect objects reasonably reliably. A good example might be a face detector. Assume that faces don't move all that fast, and there aren't too many in any given frame. Furthermore, the relationship between our representation of the state of a face and the image is uncomplicated. This occurs, for example, when the faces we view are always frontal or close to frontal. In this case, we can represent the state of the face by what it looks like (which, in principle, doesn't change because the face is frontal) and where it is.

Under these circumstances, we can build a tracker quite simply. We maintain a pool of tracks. We detect all faces in each incoming frame. We match faces to tracks, perhaps using an appearance model built from previous instances and also — at least implicitly — a dynamical model. This is where our assumptions are important; we would like faces to be sufficiently well-spaced with respect to the kinds of velocities we expect that there is seldom any ambiguity in this matching procedure. This matching procedure should not require one-one matches, meaning that some tracks may not receive a face, and some faces may not be allocated a track. For every face that is not attached to a track, we create a new track. Any track that has not received a face for several frames is declared to have ended (algorithm 1 breaks out this approach).

This basic recipe for tracking by detection is worth remembering. In many situations, nothing more complex is required, and the recipe is used without comment in a variety of papers. As a simple example, at coarse scales and from the right view, background subtraction and looking for dark blobs of the right size is sufficient to identify human heads. Yan and Forsyth use this observation in a simple track-by-detection scheme, where heads are linked across frames using a greedy algorithm [478]. The method is effective for obtaining estimates of where people go in public spaces.

The method will need some minor improvements and significant technical machinery as the relationship between state and image measurements grows more obscure. However, in this simple form, the method gives some insight into general tracking problems. The trick of creating tracks promiscuously and then pruning any track that has not received a measurement for some time is a quite general and extremely effective trick. The process of

Assumptions: We have a detector which is reasonably reliable for all aspects that matter. Objects move relatively slowly with respect to the spacing of detector responses. As a result, a detector response caused either by another object or by a false positive tends to be far from the next true position of our object.

First frame:

Create a track for each detector response.

N'th frame:

Link tracks and detector responses. Typically, each track gets the closest detector response if it is not further away than some threshold. If the detector is capable of reporting some distinguishing feature (colour, texture, size, etc.), this can be used too.

Spawn a new track for each detector response not allocated to a track.

Reap any track that has not received a measurement for some number of frames.

Cleanup: We now have trajectories in space time. Link any where this is justified (perhaps by a more sophisticated dynamical or appearance model, derived from the candidates for linking).

Algorithm 1: *The simplest tracking by detection*

linking measurements to tracks is the aspect of tracking that will cause us the most difficulty (the other aspect, inferring states from measurements, is straightforward though technically involved). This process is made easier if measurements have features that distinctively identify the track from which they come. This can occur because, for example, a face will not change gender from frame to frame, or because tracks are widely spaced with respect to the largest practical speed (so that allocating a measurement to the closest track is effective).

All this is particularly useful for face tracking, because **face detection** — determining which parts of an image contain human faces, without reference to the individual identity of the faces — is one of the substantial successes

of computer vision. Neither space nor energy allow a comprehensive review of this topic here. However, the typical approach is: One searches either rectangular or circular image windows over translation, scale and sometimes rotation; corrects illumination within these windows by methods such as histogram equalization; then presents these windows to a classifier which determines whether a face is present or not. There is then some post-processing on the classifier output to ensure that only one detect occurs at each face. This general picture appears in relatively early papers [339, 373, 374, 428, 429]. Points of variation include: the details of illumination correction; appropriate search mechanisms for rotation (cf. [376] and [381]); appropriate classifiers (cf. [289], [320], [381], [375] and [429]); building an incremental classification procedure so that many windows are rejected early and so consume little computation (see [213, 214, 450, 452] and the huge derived literature). There are a variety of strategies for detecting faces using parts, an approach that is becoming increasingly common (compare [66], [199], [251], [283, 285] and [463]; faces are becoming a common category in so-called object category recognition, see, for example, [127]).

2.2.1 Background Subtraction

The simplest detection procedure is to have a good model of the background. In this case, everything that doesn't look like the background is worth tracking. The simplest background subtraction algorithm is to take an image of the background and then subtract it from each frame, thresholding the magnitude of the difference (there is a brief introduction to this area in [135]). Changes in illumination will defeat this approach. A natural improvement is to build a moving average estimate of the background, to keep track of illumination changes (e.g. see [471, 385]; gradients can be incorporated [280]). In outdoor scenes, this approach is defeated by such phenomena as leaves moving in the wind. More sophisticated background models keep track of maximal and minimal values at each pixel [168], or build local statistical models at each pixel [73, 140, 163, 202, 203, 421, 422].

Under some circumstances, background subtraction is sufficient to track people and perform a degree of kinematic inference. Wren *et al.* describe a system, Pfinder, that uses background subtraction to identify body pixels, then identifies arm, torso and leg pixels by building "blobby" clusters [471].

Haritaoglu *et al.* describe a system called W4, which uses background subtraction to segment people from an outdoor view [168]. Foreground regions are then linked in time by applying a second order dynamic model (velocity and acceleration) to propagate median coordinates (a robust estimate of the centroid) forward in time. Sufficiently close matches trigger a search process that matches the relevant foreground component in the previous frame to that in the current frame. Because people can pass one another or form groups, foreground regions can merge, split or appear. Regions appearing, splitting or merging are dealt with by creating (resp. fusing) tracks. Good new tracks can be distinguished from bad new tracks by looking forward in the sequence: a good track continues over time. Allowing a tracker to create new tracks fairly freely, and then telling good from bad by looking at the future in this way is a traditional, and highly useful, trick in the radar tracking community (e.g. see the comprehensive book by Blackman and Popoli [44]). The background subtraction scheme is fairly elaborate, using a range of thresholds to obtain a good blob (figure 2.1). The resulting blobs are sufficiently good that the contour can be parsed to yield a decomposition into body segments. The method then segments the contours using convexity criteria, and tags the segments using: distance to the head — which is at the top of the contour; distance to the feet — which are at the bottom of the contour; and distance to the median — which is reasonably stable. All this works because, for most configurations of the body, one will encounter body segments in the same order as one walks around the contour (figure 2.2). Shadows are a perennial nuisance for background subtraction, but this can be dealt with using a stereoscopic reconstruction, as Haritaoglu *et al.* show ([167]; see also [204]).

2.2.2 Deformable Templates

Image appearance or **appearance** is a flexible term used to refer to aspects of an image that are being encoded and should be matched. Appearance models might encode such matters as: Edge position; edge orientation; the distribution of color at some scale (perhaps as a histogram, perhaps as histograms for each of some set of spatially localized buckets); or texture (usually in terms of statistics of filter outputs).

A **deformable template** or **snake** is a parametric model of image appearance usually used to localize structures. For example, one might have a



Fig. 2.1 **Fig 1 of surveillance/00868683** Background subtraction identifies groups of pixels that differ significantly from a background model. The method is most useful for some some cases of surveillance, where one is guaranteed a fixed viewpoint and a static background changing slowly in appearance. On the **left**, a background model; in the **center**, a frame; and on the **right**, the resulting image blobs. The figure is taken from Haritaoglu et al. [168]; in this paper, authors use an elaborate method involving a combination of thresholds to obtain good blobs. Figure 2.2 illustrates a method due to these authors that obtains a kinematic configuration estimate by parsing the blob. Figure from “W4: Real-time surveillance of people and their activities”, Haritaoglu *et al.*, IEEE Trans. Pattern Analysis and Machine Intelligence, 2000, © 2000 IEEE

template that models the outline of a squash [218, 219] or the outline of a person [35], place the template on the image in about the right place, and let a fitting procedure figure out the best position, orientation and parameters.

We can write this out formally as follows. Assume we have some form of template that specifies image appearance as a function of some parameters. We write this template — which gives (say) image brightness (or color, or texture, and so on) as a function of space \mathbf{x} and some parameters θ — as $T(\mathbf{x}|\theta)$. We score a comparison between the image at frame n , which we write as $I(\mathbf{x}, t_n)$, and this template using the a scoring function ρ

$$\rho(T(\mathbf{x}|\theta), I(\mathbf{x}, t_n))$$

A **point template** is built as a set of **active sites** within a model coordinate frame. These sites are to match **keypoints** identified in the image. We now build a model of acceptable sets of active sites obtained as shape, location, etc., changes. Such models can be built with, for example, the methods of **principal component analysis** (see, for example, [212]). We can now identify a match by obtaining image keypoints, building a correspondence between image keypoints and active sites on the template, and identifying parameters that minimize the fitting error.

An alternative is a **curve template**, an idea originating with the **snakes** of [218, 219]. We choose a parametric family of image curves — for example,

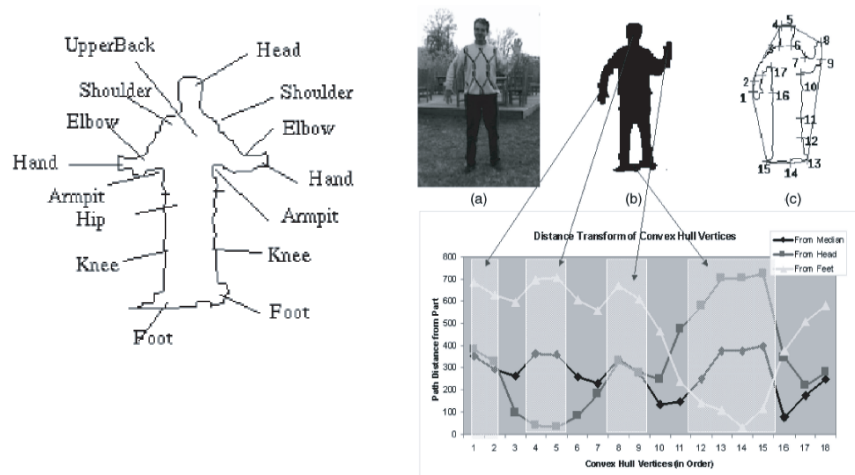


Fig. 2.2 For a given view of the body, body segments appear in the outline in a predictable manner. An example for a frontal view appears on the left. Haritaoglu et al identify vertices on the outline of a blob using a form of convexity reasoning (right (b) and right (c)), and then infer labels for these vertices by measuring the distance to head (at the top), feet (at the bottom) and median (right (d)). These distances give possibly ambiguous labels for each vertex; by applying a set of topological rules obtained using examples of multiple views like that on the left, they obtain an unambiguous labelling. **Fig 12, 15 of surveillance/00868683** Figure from “W4: Real-time surveillance of people and their activities”, Haritaoglu et al., IEEE Trans. Pattern Analysis and Machine Intelligence, 2000, © 2000 IEEE

a closed B-spline — and build a model of acceptable shapes, using methods like principal component analysis on the control points. There is an excellent account of methods in the book of Blake and Isard [45]. We can now identify a match by summing values of some image-based potential function over a set of sample points on the curve. A particularly important case occurs when we want the sample points to be close to image points where there is a strong feature response — say an edge point. It can be inconvenient to find every edge point in the image (a matter of speed) and this class of template allows us to search for edges only along short sections normal to the curve — an example of a **gate**.

Deformable templates have not been widely used as object detectors, because finding a satisfactory minimum — one that lies on the object of interest, most likely a global minimum — can be hard. The search is hard to initialize because one must identify the feature points that should lie within the gate of the template. However, in tracking problems this difficulty is mitigated if one has a dynamical model of some form. For example, the object might move

slowly, meaning that the minimum for frame n will be a good start point for frame $n + 1$. As another example, the object might move with a large, but near constant, velocity. This means that we can *predict* a good start point from frame $n + 1$ given frame n . A significant part of the difficulty is caused by image features that don't lie on the object, meaning that another useful case occurs in the near absence of clutter — perhaps background subtraction, or the imaging conditions, ensures that there are few or no extra features to confuse the fitting process.

Baumberg and Hogg track people with a deformable template built using a B-spline as above, with principal components used to determine \mathcal{W} [35]. They use background subtraction to obtain an outline for the figure, then sample the outline. For this kind of template, correspondence is generally a nuisance, but in some practical applications, this information can be supplied from quite simple considerations. For example, Baumberg and Hogg work with background subtracted data of pedestrians at fairly coarse scales from fixed views [35]. In this case, sampling the outline at fixed fractions of length, and starting at the lowest point on the principal axis yields perfectly acceptable correspondence information.

2.2.2.1 Robustness

We have presented scoring a deformable template as a form of least squares fitting problem. There is a basic difficulty in such problems. Points that are dramatically in error, usually called **outliers** and traditionally blamed on typist error [175, 372], can be overweighted in determining the fit. Outliers in vision problems tend to be unavoidable, because nature is so generous with visual data that there is usually something seriously misleading in any signal. There are a variety of methods for managing difficulties created by outliers that are used in building deformable template trackers. An estimator is called **robust** if the estimate tends to be only weakly affected by outliers. For example, the average of a set of observations is not a robust estimate of the mean of their source (because if one observation is, say, mistyped, the average could be wildly incorrect). The median *is* a robust estimate, because it will not be much affected by the mistyped observation.

Gating — the scheme of finding edge points by searching out some distance along the normal from a curve — is one strategy to obtain robustness. In

this case, one limits the distance searched. Ideally, there is only one edge point in the search window, but if there are more one takes the closest (strongest, *mutatis mutandis* depending on application details). If there is nothing, one accepts some fixed score, chosen to make the cost continuous. This means that the cost function, while strictly not differentiable, is not dominated by very distant edge points. These are not seen in the gate, and there is an upper bound on the error any one site can contribute.

An alternative is to use an **m-estimator**. One would like to score the template with a function of squared distance between site and measured point. This function should be close to the identity for small values (so that it behaves like the square) and close to some constant for large values (so that large values don't contribute large biases). A natural form is

$$\rho(u) = \frac{u}{u + \sigma}$$

so that, for d^2 small with respect to σ , we have $\rho(d^2) \approx d^2$ and for d^2 large with respect to σ we have $\rho(d^2) \approx 1$. The advantage of this approach is that nearby edgepoints dominate the fit; the disadvantage is that even fitting problems that are originally convex are no longer convex when the strategy is applied. Numerical methods are consequently more complex, and one must use multiple start points. There is little hope of having a convex problem, because different start points correspond to different splits of the data set into “important” points and outliers; there is usually more than one such split. Again, large errors no longer dominate the estimation process, and the method is almost universally applied for flow templates.

2.2.2.2 The Hausdorff Distance

The **Hausdorff distance** is a method to measure similarity between binary images (for example, edge maps; the method originates in Minkowski's work in convex analysis, where it takes a somewhat different form). Assume we have two sets of points P and Q ; typically, each point is an edge point in an image. We define the Hausdorff distance between the two sets to be

$$H(P, Q) = \max(h(P, Q), h(Q, P))$$

where

$$h(P, Q) = \max_{p \in P} \min_{q \in Q} \|p - q\|$$

The distance is small if there is a point in Q close to each point in P and a point in P close to each point in P . There is a difficulty with robustness, as the Hausdorff distance is large if there are points with no good matches. In practice, one uses a variant of the Hausdorff distance (the **generalized Hausdorff distance**) where the distance used is the k -th ranked of the available distances rather than the largest. Define F_k^{th} to be the operator that orders the elements of its input largest to smallest, then takes the k 'th largest. We now have

$$H_k(P, Q) = \max(h_k(P, Q), h_k(Q, P))$$

where

$$h_k(P, Q) = F_k^{th}(\min_{q \in Q} \|p - q\|)$$

(for example, if there are $2n$ points in P , then $h_n(P, Q)$ will give the median of the minimum distances). The advantage of all this is that some large distances get ignored.

Now we can compare a template P with an image Q by determining some family of transformations $\mathcal{T}(\theta)$ and then choosing the set of parameters $\hat{\theta}$ that minimizes

$$H_k(\mathcal{T}(\theta) \circ P, Q)$$

This will involve some form of search over θ . The search is likely to be simplified if — as applies in the case of tracking — we have a fair estimate of $\hat{\theta}$ to hand.

Huttenlocher *et al.* track using the Hausdorff distance [191]. The template, which consists of a set of edge points, is itself allowed to deform. Images are represented by edge points. They identify the instance of the latest template in the next frame by searching over translations θ of the template to obtain the smallest value of $H_k(\mathcal{T}(\theta) \circ P, Q)$. They then translate the template to that location, and identify all edge points that are within some distance of the current template's edge points. The resulting points form the template for the next frame. This process allows the template to deform to take into account, say, the deformation of the body as a person moves. Performance in heavily textured video must depend on the extent to which the edge detection process suppresses edges and the setting of this distance parameter (a large distance and lots of texture is likely to lead to catastrophe).

2.3 Tracking using Flow

The difficulty with tracking by detection is that one might not have a deformable template that fully specifies the appearance of an object. It is quite common to have a template that specifies the shape of the domain spanned by the object and the type of its transformation, but not what lies within. Typically, we don't know the pattern, but we do know how it moves. There are several important examples:

- **Human body segments** tend to look like a rectangle in any frame, and the motion of this rectangle is likely to be either Euclidean or affine, depending on imaging circumstances.
- **A face in a webcam** tends to fill a blob-like domain and undergo mainly Euclidean transformations. This is useful for those building user interfaces where the camera on the monitor views the user, and there are numerous papers dealing with this. The face is not necessarily frontal — computer users occasionally look away from their monitors — but tends to be large, blobby and centered.
- **Edge templates**, particularly those specifying outlines, are usually used because we don't know what the interior of the region looks like. Quite often, as we have seen, we know how the template can deform and move. However, we cannot score the interior of the domain because we don't know (say) the pattern of clothing being worn.

In each of these cases, we cannot use tracking by detection as above because we do not possess an appropriate template. As a matter of experience, objects don't change appearance much from frame to frame (alternatively, we should use the term appearance to apply to properties that don't change from frame to frame). All this implies that parts of the previous image could serve as a template if we have a motion model and domain model. We could use a correspondence model to link pixels in the domain in frame n with those in the domain in frame $n + 1$. A “good” linking should pair pixels that have similar appearances. Such considerations as camera properties, the motion of rigid objects, and computational expense suggest choosing the correspondence model from a small parametric family.

All this gives a formal framework. Write a pixel position in the n 'th frame as \mathbf{x}_n , the domain in the n 'th frame as \mathcal{D}_n , and the transformation from the n 'th frame to the $n+1$ 'th frame as $\mathcal{T}_{n \rightarrow n+1}(\cdot; \theta_n)$. In this notation θ_n represent parameters for the transformation from the n 'th frame to the $n+1$ 'th frame, and we have that $\mathbf{x}_{n+1} = \mathcal{T}_{n \rightarrow n+1}(\mathbf{x}_n; \theta_n)$.

We assume we know \mathcal{D}_n . We can obtain \mathcal{D}_{n+1} from \mathcal{D}_n as $\mathcal{T}_{n \rightarrow n+1}(\mathcal{D}_n; \theta_n)$. Now we can score the parameters θ_n representing the *change* in state between frames $n+1$ and n by comparing \mathcal{D}_n with \mathcal{D}_{n+1} (which is a function of θ_n). We compute some representation of image information $\mathbf{R}(\mathbf{x})$, and, within the domain \mathcal{D}_{n+1} compare $\mathbf{R}(\mathbf{x}_{n+1})$ with $\mathbf{R}(\mathcal{T}_{n \rightarrow n+1}(\mathbf{x}_n; \theta_n))$, where the transformation is applied to the domain \mathcal{D}_n .

2.3.1 Optic Flow

Generally, a frame-to-frame correspondence should be thought of as a **flow field** (or an **optic flow field**) — a vector field in the image giving local image motion at each pixel. A flow field is fairly clearly a correspondence, and a correspondence gives rise to a flow field (put the tail of the vector at the pixel position in frame n , and the head at the position in frame $n+1$). The notion of optic flow originates with Gibson (see, for example, [148]).

A useful construction in the optic flow literature assumes that image intensity is a continuous function of position and time, $I(\mathbf{x}, t)$. We then assume that the intensity of image patches does not change with movement. While this assumption may run into troubles with illumination models, specularities, etc., it is not outrageous for small movements. Furthermore, it underlies our willingness to compare pixel values in frames. Accepting this assumption, we have

$$\frac{dI}{dt} = \nabla I \cdot \frac{d\mathbf{x}}{dt} + \frac{\partial I}{\partial t} = 0$$

(known as the **optic flow equation**, e.g. see [186]). Flow is represented by $d\mathbf{x}/dt$. This is important, because if we confine our attention to an appropriate domain, comparing $I(\mathcal{T}(\mathbf{x}; \theta_n), t_{n+1})$ with $I(\mathbf{x}, t_n)$ involves, in essence, estimating the total derivative. In particular,

$$I(\mathcal{T}(\mathbf{x}; \theta_n), t_{n+1}) - I(\mathbf{x}, t_n) \approx \frac{dI}{dt}$$

Furthermore, the equivalence between correspondence and flow suggests a simpler form for the transformation of pixel values. We regard $\mathcal{T}(\mathbf{x}; \theta_n)$ as

taking \mathbf{x} from the tail of a flow arrow to the head. At short timescales, this justifies the view that $T(\mathbf{x}; \theta_n) = \mathbf{x} + \delta\mathbf{x}(\theta_n)$.

2.3.2 Image stabilization

This form of tracking can be used to build boxes around moving objects, a practice known as **image stabilization**. One has a moving object on a fairly uniform background, and would like to build a domain such that the moving object is centered on the domain. This has the advantage that one can look at relative, rather than absolute, motion cues. For example, one might take a soccer player running around a field, and build a box around the player. If one then fixes the box and its contents in one place, the vast majority of motion cues within the box are cues to how the player's body configuration is changing. As another example, one might stabilize a box around an aerial view of a moving vehicle; now the box contains all visual information about the vehicle's identity.

Efros *et al.* use a straightforward version of this method, where domains are rectangles and flow is pure translation, to stabilize boxes around people viewed at a medium scale (for example, in a soccer video) [116]. In some circumstances, good results can be obtained by matching a rectangle in frame n with the rectangle in frame $n + 1$ that has smallest sum-of-squared differences — which might be found by blank search, assisted perhaps by velocity constraints. This is going to work best if the background is relatively simple — say, the constant green of a soccer field — as then the background isn't a source of noise, so the figure need not be segmented (figure 2.3). For more complex backgrounds, the approach may still work if one performs background subtraction before stabilization. At a medium scale it is very difficult to localize arms and legs, but they do leave traces in the flow field. The stabilization procedure means that the flow information can be computed with respect to a torso coordinate system, resulting in a representation that can be used to match at a kinematic level, without needing an explicit representation of arm and leg configurations (figure 2.3).

2.3.3 Cardboard people

Flow based tracking has the advantage that one doesn't need an explicit model of the appearance of the template. Ju *et al.* build a model of legs in terms

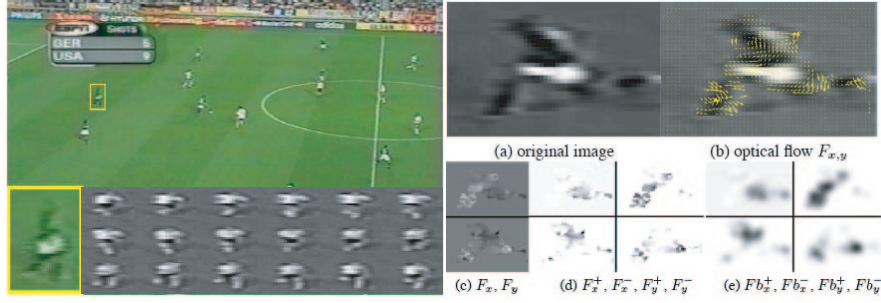


Fig. 2.3 **Figure 1, part of 3 and 4 from Efros et al. Recognizing action at a distance.** Flow based tracking can be useful for medium scale video. Efros et al. stabilize boxes around the torso of players in football video using a sum of squared differences (SSD) as a cost function and straightforward search to identify the best translation values. As the figure on the left shows, the resulting boxes are stable with respect to the torso. On the **top right**, larger versions of the boxes for some cases. Note that, because the video is at medium scale, it is difficult to resolve arms and legs, which are severely affected by motion blur. Nonetheless, one can make a useful estimate of what the body is doing by computing an estimate of optic flow (**bottom right**, F_x, F_y), rectifying this estimate (**bottom right**, $F_x^+, F_x^-, F_y^+, F_y^-$) and then smoothing the result (**bottom right**, $Fb_x^+,$ etc.). The result is a smoothed estimate of where particular velocity directions are distributed with respect to the torso, which can be used to match and label frames. Figure from “Recognizing Action at a Distance”, Efros *et al.*, IEEE Int. Conf. Computer Vision 2003, © 2003 IEEE

of a set of articulated rectangular patches (“cardboard people”) [216]. Assume we have a domain D in the n ’th image $I(\mathbf{x}, t_n)$ and a flow field $\delta\mathbf{x}(\theta)$ parametrized by θ . Now this flow field takes D to some domain in the $n+1$ ’th image, and establishes a correspondence between pixels in the n ’th and the $n+1$ ’th image. Ju *et al.* score

$$\sum_D \rho(I_{n+1}(\mathbf{x} + \delta\mathbf{x}(\theta)) - I_n(\mathbf{x}))$$

where ρ is some measure of image error, which is small when the two compare well and large when they are different. Notice that this is a very general approach to the tracking problem, with the difficulty that, unless one is careful about the flow model the problem of finding a minimum might be hard. To our knowledge, the image score is always applied to pixel values, and it seems interesting to wonder what would happen if one scored a difference in texture descriptors.

Typically, the score is not minimized directly, but is approximated with

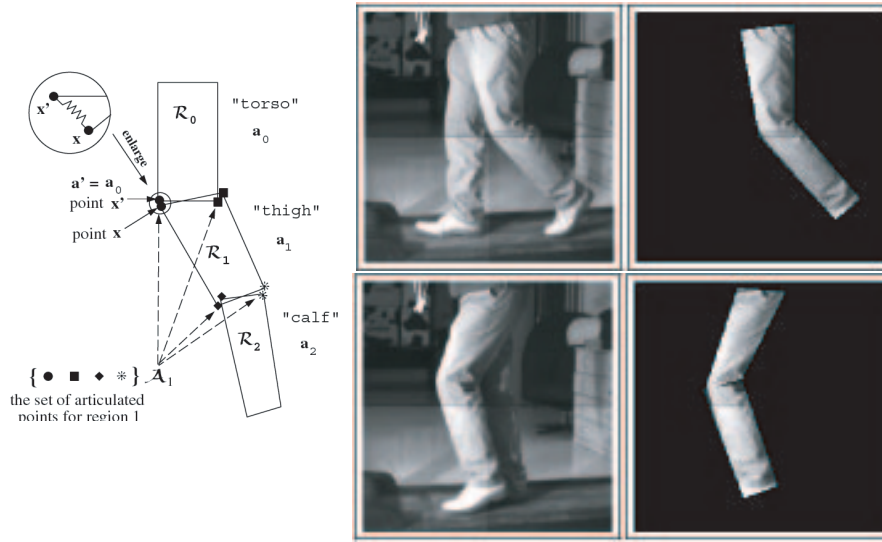


Fig. 2.4 **Figure 3, 4 from Ju, Black and Yacoob, “Cardboard People”, no permission yet** On the left, a 2D flow based model of a leg, called a “cardboard people” model by Ju et al [216]; there is a lower leg, an upper leg and a torso. Each domain is roughly rectangular, and the domains are coupled with an energy term to ensure they do not drift apart. The model is tracked by finding the set of deformation parameters that carve out a domain in the $n + 1$ 'th frame that is most like the known domain in the n 'th frame. On the right, two frames from a track, with the left column showing the original frame and the right column showing the track. Notice how the pattern of buckling folds on the trouser leg changes as the leg bends; this leads to quite significant changes in the texture and shading signal in the domain. These changes can be a significant nuisance. Figure from “Cardboard People: A Parameterized Model of Articulated Image Motion”, Ju *et al.*, IEEE Int. Conf. Face and Gesture, 1996 © 1996 IEEE

the optic flow equation and with a Taylor series. We have

$$\sum_D \rho(I(\mathbf{x} + \delta\mathbf{x}(\theta), t_{n+1}) - I_n(\mathbf{x}, t_n))$$

is approximately equal to

$$\sum_D \rho\left(\frac{dI}{dt}\right) = \sum_D \rho\left(\frac{\partial I}{\partial x} \delta x(\theta_n) + \frac{\partial I}{\partial y} \delta y(\theta_n) + \frac{\partial I}{\partial t}\right)$$

(this works because $\Delta t = 1$). Now assume that a patch has been marked out in a frame; then one can determine its configuration in the next by minimizing this error summed over the domain. The error itself is easily evaluated using smoothed derivative estimates. As we show below, we can further simplify error evaluation by building a flow model with convenient form. To track an

articulated figure, we attach a further term that encourages relevant vertices of each separate patch to stay close. Similarly, Black et al construct parametric families of flow fields and use them to track lips and legs, in the latter case yielding a satisfactory estimate of walk parameters [43]. In both cases, the flow model is view dependent. Yacoob and Davis build a view independent parametric flow field models to track views of walking humans [474]. As one would expect, this technique can be combined with others; for example, the W4S system of Haritaoglu *et al.* uses a “cardboard people” model to track torso configurations within the regions described above [167].

2.3.4 Building Flow Templates

We have seen how to construct tracks given parametric models of flow. But how is one to obtain good models? One strategy is to take a pool of examples of the types of flow one would like to track, and try to find a set of basis flows that explains most of the variation (for examples, see [216]). In this case, and writing θ_i for the i 'th component of the parameter vector and \mathbf{F}_i for the i 'th flow basis vector field, one has

$$\delta \mathbf{x} = \sum_i \theta_i \mathbf{F}_i$$

Now write ∇I for the image gradient and exploit the optic flow equation and a Taylor series as above. We get

$$\rho \left(\sum_i \theta_i ((\nabla I)^T \mathbf{F}_i) + \frac{\partial I}{\partial t} \right)$$

This can be done with a singular value decomposition (and is equivalent to principal components analysis). A second strategy is to assume that flows involve what are essentially 2D effects — this is particularly appropriate for lateral views of human limbs — so that a set of basis flows that encodes translation, rotation and some affine effects is probably sufficient. One can obtain such flows by writing

$$\delta \mathbf{x} = \begin{pmatrix} u(\mathbf{x}) \\ v(\mathbf{x}) \end{pmatrix} = \begin{pmatrix} a_0 + a_1x + a_2y + a_6x^2 + a_7xy \\ a_3 + a_4x + a_5y + a_6xy + a_7y^2 \end{pmatrix}$$

This model is linear in the parameters (the a_i), which is convenient; it provides a reasonable encoding of flows resulting from 3D motions of a 2D




Fig. 2.5 **Figure 2 from Ju, Black and Yacoob, “Cardboard People”, no permission yet** Typical flows generated by the model $(u(\mathbf{x}), v(\mathbf{x}))^T = (a_0 + a_1x + a_2y + a_6x^2 + a_7xy, a_3 + a_4x + a_5y + a_6xy + a_7y^2)$. Different values of the a_i give different flows, and the model can generate flows typical of a 2D figure moving in 3D. We write $\mathbf{a} = (a_0, a_1, a_2, a_3, a_4, a_5, a_6, a_7)$. **Divergence** occurs when the image is scaled; for example, $\mathbf{a} = (0, 1, 0, 0, 0, 1, 0, 0)$. **Deformation** occurs when one direction shrinks and another grows (for example, rotation about an axis parallel to the view plane in an orthographic camera); for example, $\mathbf{a} = (0, 1, 0, 0, 0, -1, 0, 0)$. **Curl** can result from in plane rotation; for example, $\mathbf{a} = (0, 0, -1, 0, 1, 0, 0, 0)$. **Yaw** models rotation about a vertical axis in a perspective camera; for example $\mathbf{a} = (0, 0, 0, 0, 0, 1, 0, 0)$. Finally, **pitch** models rotation about a horizontal axis in a perspective camera; for example $\mathbf{a} = (0, 0, 0, 0, 0, 0, 1)$. Figure from “Cardboard People: A Parameterized Model of Articulated Image Motion”, Ju *et al.*, IEEE Int. Conf. Face and Gesture, 1996 © 1996 IEEE

rectangle (see figure 2.5). One may also learn linearized flow models from example data [474].

2.3.5 Flow models from kinematic models

An alternative method to build such templates is to work in 3D, and exploit the chain rule, as in the work of Bregler and Malik [60, 61]. We start with a segment in 3D, which is in some configuration and viewed with some camera. Each point on the segment produces some image value. We could think of the image values as a function — the **appearance map** — defined on the segment. This allows us to see viewing the segment as building a mapping from the points on the segment to the image domain. The image values are obtained by taking each point in the image, finding the corresponding point (if any) on the segment, and then evaluating the appearance map at this point.




Fig. 2.6 **Figure 7** from Bregler and Malik CVPR paper “Tracking people with twists and exponential maps” **no permission yet** Bregler and Malik formulate parametric flow models by modelling a person as a kinematic chain and then differentiating the maps from segment to image. They then track by searching for the parameter update that best aligns the current image pixels with those of the previous frame under this flow model. There is no dynamical model. This means that complex legacy footage, like these frames from the photographs of Eduard Muybridge [304, 303], can be tracked. Muybridge’s frames are difficult to track because the frame-frame timing is not exact, and the figures can move in quite complex ways (see figure 4.7). Figure from “Tracking People with Twists and Exponential Maps”, Bregler and Malik, Proc. Computer Vision and Pattern Recognition, 1998, © 1998 IEEE

All this leads to an important formal model, again under the assumption that motions in 3D do not affect the appearance map in any significant way. We have a parametrized family of maps from points on the body to the image. A flow field in the image is a vector field induced by a change in the choice of parameters (caused by either a change in joint configuration or a camera movement). We will always assume that the change in parameters from frame to frame is small. At this point, we must introduce some nota-

tion. Write the map that takes points on the segment to points in the n 'th image as $\mathcal{T}_{s \rightarrow I}(\cdot; \theta_n)$, where θ_n are parameters representing camera configuration, intrinsics, etc. The point \mathbf{p} on the segment appears in image n at $\mathbf{x}_n = \mathcal{T}_{s \rightarrow I}(\mathbf{p}; \theta_n)$ and in image $n + 1$ at $\mathbf{x}_{n+1} = \mathcal{T}_{s \rightarrow I}(\mathbf{p}; \theta_{n+1})$. The tail of the flow arrow is at \mathbf{x}_n and the head is at \mathbf{x}_{n+1} . The change in parameters, $\Delta\theta = \theta_{n+1} - \theta_n$ is small. Then the flow is

$$\mathbf{x}_{n+1} - \mathbf{x}_n = \mathcal{T}_{s \rightarrow I}(\mathbf{p}; \theta_{n+1}) - \mathcal{T}_{s \rightarrow I}(\mathbf{p}; \theta_n) \approx \nabla_{\theta} \mathcal{T}_{s \rightarrow I} \cdot \Delta\theta$$

where the gradient, $\nabla_{\theta} \mathcal{T}_{s \rightarrow I}$, is evaluated at (\mathbf{p}, θ_n) .

2.3.5.1 Tracking a Derivative Flow Model

The main point here is that the flow at \mathbf{x}_n can be obtained by fixing the relevant point \mathbf{p} on the object, then considering the map taking the *parameters* to the image plane — the derivative of $\mathcal{T}_{s \rightarrow I}(\mathbf{p}; \cdot)$. This is important, because the flow $\nabla_{\theta} \mathcal{T}_{s \rightarrow I} \cdot \Delta\theta$ is a *linear* function of $\Delta\theta$. We now have the outline of a tracking algorithm:

- Start at frame $n = 0$ and some known configuration $\theta_0 = \hat{\theta}$.
- **Fit:** Fit the best value of $\Delta\theta$ to the flow between the frame n and frame $n + 1$ using the flow model given by the derivative evaluated at θ_n .
- **Update:** Update the parameters by $\theta_{n+1} = \theta_n + \Delta\theta$ and set n to $n + 1$.

This should be seen as a primitive integrator, using Euler's method and inheriting all the problems that come with it. This view confirms the reasonable suspicion that fast movements are unlikely to be tracked well unless that sampling rate is high.

2.3.5.2 The Flow Model from the Chain Rule

In the special case of segments lying on a **kinematic tree** — a series of links attached by joints of known parametric form, where there are no loops — the chain rule means that the derivative takes a special form. Each segment in a kinematic tree has its own coordinate system, and the joint is represented by a map from a link's world coordinate system to that of its parent. The parent of

segment k is segment $k - 1$. They are connected by a joint whose parameters at frame n are $\theta_{k,n}$. In general, in a kinematic tree, points on segments are affected by parameters at joints above them in the tree. Furthermore, we can obtain a transformation to the image by recursively concatenating transformations. Write the camera as $\mathcal{T}_{w \rightarrow i}$. Then the transformation taking a point of link k in frame n to the image can be written as

$$\mathcal{T}_{k \rightarrow i} = \mathcal{T}_{w \rightarrow i} \circ \mathcal{T}_{k-1 \rightarrow w} \circ \mathcal{T}_{k \rightarrow k-1}$$

Notice that the only transformation that depends on $\theta_{k,n}$ here is $\mathcal{T}_{k \rightarrow k-1}$.

There is an advantage to changing notation at this point. Write $\mathcal{T}_{k \rightarrow k-1}$ as \mathcal{T}_k . The root of the tree is at segment one, and we can write $\mathcal{T}_{1 \rightarrow w}$ as \mathcal{T}_1 and $\mathcal{T}_{w \rightarrow i}$ as \mathcal{T}_0 . We continue to divide up the parameters θ into components, $\theta_{k,n}$ being the components associated with segment k in the n 'th frame ($\theta_{k,n}$ are viewing parameters in frame n). We can now see the map from point \mathbf{p} on segment k to the image as

$$\mathcal{T}_{k \rightarrow i}(\mathbf{p}; \theta) = \mathcal{T}_0(\mathcal{T}_1(\mathcal{T}_2(\dots; \theta_2); \theta_1); \theta_0)$$

This is somewhat inconvenient to write out, and it is helpful to keep track of intermediate values. Introduce the notation

$$\mathbf{p}_l = \mathcal{T}_{k \rightarrow l}(\mathbf{p}; \theta)$$

for the point \mathbf{p} in the coordinate system of the l 'th link.

Our transformations have two types of argument: the points in space, and the camera parameters. It is useful to distinguish between two types of derivative. Write the partial derivative of a transformation \mathcal{T} with respect to its spatial arguments as $D\mathcal{T}$. In coordinates, \mathcal{T} would take the form $(f_1(x_1, x_2, x_3, \theta), f_2(x_1, x_2, x_3, \theta), f_3(x_1, x_2, x_3, \theta))$, and this derivative would be the matrix whose i, j 'th element is $\partial f_i / \partial x_j$. Similarly, write the partial derivative of a transformation \mathcal{T} with respect to parameters θ as D_θ . If we regard θ as a vector of parameters whose j 'th entry is θ_j , then in coordinates this derivative would be the matrix whose i, j 'th element is $\partial f_i / \partial \theta_j$.

This orgy of notation leads to a simple form for the flow. Write the flow at point \mathbf{x} — which is the image of point \mathbf{p} on segment k — in frame n as $\mathbf{v}(\mathbf{x}, \theta_n)$. Then

$$\mathbf{v}(\mathbf{x}, \theta_n) = D_\theta \mathcal{T}_0(\mathbf{p}_0; \theta_0) \cdot \Delta \theta_0 + D_x \mathcal{T}_0 \circ D_\theta \mathcal{T}_1(\mathbf{p}_1; \theta_1) \Delta \theta_1$$

$$\dots + D_x \mathcal{T}_0 \circ D_x \mathcal{T}_1 \circ \dots \circ D_x \mathcal{T}_{k-1} \circ D_\theta \mathcal{T}_k(\mathbf{p}; \theta_k) \Delta \theta_k$$

Our indexing scheme hasn't taken into account the fact that we're dealing with a tree, but this doesn't matter; we need care only about links on the path from the relevant segment to the root. Furthermore, there is a relatively efficient algorithm for computing this derivative. We pass from the leaves to the root computing intermediate configurations \mathbf{p} for each point \mathbf{p} and the relevant parameter derivatives. We then pass from the root to the leaves concatenating spatial derivatives and summing.

2.3.5.3 Rigid-body Transformations

All the above takes a convenient and simple form for rigid-body transformations (which are likely to be the main interest in human tracking). We use homogeneous coordinates to represent points in 3D, and so a rigid body transformation takes the form

$$\mathcal{T}(\mathbf{p}, \theta) = \begin{bmatrix} \mathcal{R} & \mathbf{t} \\ \mathbf{0} & 1 \end{bmatrix} \mathbf{p}$$

where \mathcal{R} is an orthonormal matrix with determinant one (a rotation matrix). The parameters are the parameters of the rotation matrix and the coefficients of the vector \mathbf{t} . This means the spatial derivative is the same as the transformation, which is convenient.

The derivatives with respect to the parameters are also relatively easily dealt with. Recall the definition of the **matrix exponential** as an infinite sum,

$$\exp(\mathcal{M}) = \mathcal{I} + \mathcal{M} + \frac{1}{2}\mathcal{M}^2 + \mathcal{M}^3 + \dots + \frac{1}{n}\mathcal{M}^n \dots$$

where the sum exists. Now it is straightforward to demonstrate that if

$$\mathcal{M} = \begin{bmatrix} \mathcal{A} & \mathbf{t} \\ \mathbf{0} & 0 \end{bmatrix}$$

and if \mathcal{A} is antisymmetric, then $\exp(\mathcal{M})$ is a rigid-body transformation. The elements of the antisymmetric matrix parametrize the rotation, and the right-most column is the translation. This is useful, because

$$\frac{\partial (\exp \mathcal{M}(\theta))}{\partial \theta} = \left(\frac{\partial \mathcal{M}(\theta)}{\partial \theta} \right) \exp \mathcal{M}(\theta)$$

which gives straightforward forms for the parameter derivatives.

2.4 Tracking with Probability

It is convenient to see tracking as a probabilistic inference problem. In particular, we have a sequence of states X_0, X_1, \dots, X_N produced by some dynamical process. These states are unknown — they are sometimes called **hidden states** for this reason — but there are measurements Y_0, Y_1, \dots, Y_N . Two problems follow naturally:

- **Tracking**, where we wish to determine some representation of $P(X_k|Y_0, \dots, Y_k)$;
- **Filtering**, where we wish to determine some representation of $P(X_k|Y_0, \dots, Y_N)$ (i.e. we get to use "future" measurements to infer the state).

These problems are massively simplified by two important assumptions.

- We assume measurements depend only the hidden state, that is, that $P(Y_k|X_0, \dots, X_N, Y_0, \dots, Y_N) = P(Y_k|X_k)$.
- We assume that the probability density for a new state is a function only of the previous state; that is, $P(X_k|X_0, \dots, X_{k-1}) = P(X_k|X_{k-1})$, or, equivalently, that X_i form a **Markov chain**.

Now tracking involves three steps:

Prediction: where we construct some prediction of the future state given past measurements, or equivalently, construct a representation of $P(X_k|Y_0, \dots, Y_{k-1})$. Straightforward manipulation of probability combined with the assumptions above yields that the **prior** or **predictive density** is

$$P(X_k|Y_0, \dots, Y_{k-1}) = \int P(X_k|X_{k-1})P(X_{k-1}|Y_0, \dots, Y_{k-1})dX_{k-1}$$

Data association: where we use the predictive density — which is sometimes called the **prior** — and anything else likely to be helpful, to determine which of a pool of measurements contribute to the value of Y_k .

Correction: where we incorporate the new measurement into what is known, or, equivalently, construct a representation of $P(X_k|Y_0, \dots, Y_k)$. Straightforward manipulation of probability combined with the assumptions above yields that the **posterior** is

$$P(X_k|Y_0, \dots, Y_k) = \frac{P(Y_k|X_k)P(X_k|Y_0, \dots, Y_{k-1})}{\int P(Y_k|X_k)P(X_k|Y_0, \dots, Y_{k-1})dX_k}$$

2.4.1 Linear Dynamics and the Kalman Filter

All this is much simplified in the case that the emission model is linear, the dynamic model is linear, and all noise is Gaussian. In this case, all densities are Normal and the mean and covariance are sufficient to represent them. Both tracking and filtering boil down to maintenance of these parameters. There is a simple set of update rules (given in algorithm 2; notation below), the **Kalman filter**.

Notation: We write $X \sim N(\mu; \Sigma)$ to mean that X is a normal random variable with mean μ and covariance Σ . Both dynamics and emission are linear, so we can write

$$X_k \sim N(\mathcal{A}_k X_{k-1}; \Sigma_k^{(d)})$$

and

$$Y_k \sim N(\mathcal{B}_k X_k; \Sigma_k^{(m)})$$

We will represent the mean of $P(X_i|y_0, \dots, y_{i-1})$ as \overline{X}_i^- and the mean of $P(X_i|y_0, \dots, y_i)$ as \overline{X}_i^+ — the superscripts suggest that they represent our belief about X_i immediately before and immediately after the i 'th measurement arrives. Similarly, we will represent the standard deviation of $P(X_i|y_0, \dots, y_{i-1})$ as Σ_i^- and of $P(X_i|y_0, \dots, y_i)$ as Σ_i^+ . In each case, we will assume that we know $P(X_{i-1}|y_0, \dots, y_{i-1})$, meaning that we know \overline{X}_{i-1}^+ and Σ_{i-1}^+ .

Filtering is straightforward. We obtain a backward estimate by running the filter backward in time, and treat this as another measurement. Extensive detail on the Kalman filter and derived methods appears in [147, 27].

2.4.2 Data Association

Data association involves determining which pixels or image measurements should contribute to a track. That data association is a nuisance is a persistent theme of this work. Data association is genuinely difficult to handle satisfactorily — after all, determining which pixels contribute to which decision seems to be a core — and often very difficult — computer vision problem. The problem is usually particularly difficult when one wishes to track people, for several reasons. First, standard data association techniques aren't really all that much help, as for almost every aspect the image domain covered by

Dynamic Model:

$$\begin{aligned}\mathbf{x}_i &\sim N(\mathcal{D}_i \mathbf{x}_{i-1}, \Sigma_{d_i}) \\ \mathbf{y}_i &\sim N(\mathcal{M}_i \mathbf{x}_i, \Sigma_{m_i})\end{aligned}$$

Start Assumptions: $\bar{\mathbf{x}}_0^-$ and Σ_0^- are known**Update Equations: Prediction**

$$\begin{aligned}\bar{\mathbf{x}}_i^- &= \mathcal{D}_i \bar{\mathbf{x}}_{i-1}^+ \\ \Sigma_i^- &= \Sigma_{d_i} + \mathcal{D}_i \Sigma_{i-1}^+ \mathcal{D}_i\end{aligned}$$

Update Equations: Correction

$$\begin{aligned}\mathcal{K}_i &= \Sigma_i^- \mathcal{M}_i^T [\mathcal{M}_i \Sigma_i^- \mathcal{M}_i^T + \Sigma_{m_i}]^{-1} \\ \bar{\mathbf{x}}_i^+ &= \bar{\mathbf{x}}_i^- + \mathcal{K}_i [\mathbf{y}_i - \mathcal{M}_i \bar{\mathbf{x}}_i^-] \\ \Sigma_i^+ &= [Id - \mathcal{K}_i \mathcal{M}_i] \Sigma_i^-\end{aligned}$$

Algorithm 2: *The Kalman filter updates estimates of the mean and covariance of the various distributions encountered while tracking a state variable of some fixed dimension using the given dynamic model.*

a person changes shape very aggressively, and can do so very fast. Second, there seem to be a lot of background objects that look like some human body parts; for example, kinematic tracking of humans in office scenes is very often complicated by the fact that many book spines (or book shelves) can look like forearms.

In tracking by detection, almost all computation is directed at data association, which is achieved by minimizing ρ with respect to the template's parameters — the support of ρ identifies the relevant pixels. Similarly, in tracking using flow, data association is achieved by choosing the parameters of a flow model to get a good match between domains in frames n and $n + 1$ — the definition of the domain cuts out the relevant pixels. When these meth-

ods run awry, it is because these data association methods have failed. Either one cannot find the template, or one cannot get good parameters for the flow model.

There are a variety of simple data association strategies which exploit the presence of probability models. In particular, we have an estimate of $P(X_n|Y_0, \dots, Y_{n-1})$ and we know $P(Y_n|X_n)$. From this we can obtain an estimate of $P(Y_n|Y_0, \dots, Y_{n-1})$, which gives us hints as to where the measurement might be.

One can use a **gate** — we look only at measurements that lie in a domain where $P(Y_n|Y_0, \dots, Y_{n-1})$ is big enough. This is a method with roots in radar tracking of missiles and aeroplanes, where one must deal with only a small number (compared with the number of pixels in an image!) of returns, but the idea has been useful in visual tracking applications.

One can use **nearest neighbours**. In the classical version, we have a small set of possible measurements, and we choose the measurement with the largest value of $P(Y_n|Y_0, \dots, Y_{n-1})$. This has all the dangers of wishful thinking — we are deciding that a measurement is valid because it is consistent with our track — but is often useful in practice (see chapter *** of [135] for an example of how badly the method can behave). This strategy doesn't apply to most cases of tracking people in video because the search to find the maximising Y_n — which would likely be an image region — could be too difficult (but see chapter 4). However, it could be applied when one is tracking markers attached to the body — in this case, we need to know which marker is which, and this information could be obtained by allocating a measurement to the marker whose predicted position is closest.

One can use **probabilistic data association**, where we use a weighted combination of measurements within a gate, weighted using (a) the predicted measurement and (b) the probability a measurement has dropped out. Again, this method has the dangers of wishful thinking, and again does not apply to most cases of tracking people; however, it could again be applied when one is tracking markers attached to the body.

2.4.3 Multiple Modes

The Kalman filter is the workhorse of estimation, and can give useful results under many conditions. One doesn't need a guarantee of linearity to use a

Kalman filter — if the logic of the application indicates that a linear model is reasonable, there is a good chance a Kalman filter will work. Rohr used a Kalman filter to track a walking person successfully, evaluating the measurement by matches to line segments on the outline [365, 364].

More recently, the method tends not to be used because of concerns about multiple modes. The representation adopted by a Kalman filter (the mean and covariance, sufficient statistics for a Gaussian distribution) tends to represent multimodal distributions poorly. There are several reasons one might encounter multiple modes.

First, nonlinear dynamics — or nonlinear measurement processes, or both — can create serious problems. The basic difficulty is that even quite innocuous looking setups can produce densities that are not normal, and are very difficult to represent and model. For example, let us look at only the hidden state. Assume that this is one dimensional. Now assume that state updates are *deterministic*, with $X_{i+1} = X_i + \epsilon \sin(X_i)$. If ϵ is sufficiently small, we have that for $0 < X_i < \pi$, $X_i < X_{i+1} < \pi$; for $-\pi < X_i < 0$, $-\pi < X_{i+1} < X_i$; and so on. Now assume that $P(X_0)$ is normal. For sufficiently large k , $P(X_k)$ will not be; there will be “clumps” of probability centered around the points $(2j+1)\pi$ for j an integer. These clumps will be very difficult to represent, particularly if $P(X_0)$ has very large variance so that many clumps are important. Notice that what is creating a problem here is that quite small non-linearities in dynamics can cause probability to be concentrated in ways that are very difficult to represent. In particular, nonlinear dynamics are likely to produce densities with complicated sufficient statistics. There are cases where nonlinear dynamics does lead to densities that can be guaranteed to have finite-dimensional sufficient statistics (see [39, 99, 98]); to our knowledge, these have not been applied to human tracking.

Second, there are practical phenomena in human tracking that tend to suggest that non-normal distributions are a significant part of the problem. Assume we wish to track a 3D model of an arm in a single image. The elbow is bent; as it straightens, it will eventually run into an **end-stop** — the forearm can’t rotate further without damage. At the end-stop, the posterior on state can’t be a normal distribution, because a normal distribution would have some support on the wrong side of the end-stop, and this has a significant effect on the shape of the posterior (see figure 3.5). Another case that is likely, but not guaranteed, to cause trouble is a **kinematic singularity**. For example,

if the elbow is bent, we will be able to observe rotation about the humerus, but current observation models will make this unobservable if the elbow is straight (because the outline of the arm will not change; no current method can use the changes in appearance of the hand that will result). The dimension of the state space has collapsed. The posterior might be a normal distribution in this reduced dimension space, but that would require explicitly representing the collapse. The alternative, a covariance matrix of reduced rank, creates unattractive problems of representation. Deutscher *et al.* produce evidence that, in both cases, posteriors are not, in fact, normal distributions, and show that an extended Kalman filter can lose track in these cases [106].

Third, kinematic ambiguity in the relations between 3D and 2D are a major source of multiple modes. Assume we are tracking a human figure using a 3D representation of the body in a single view. If, for example, many 3D configurations correspond exactly to a single 2D configuration, then we expect the posterior to have multiple modes. Chapter 3 discusses this issue in extensive detail.

Fourth, the richest source of multiple modes is data association problems. An easy example illustrates how nasty this problem can be. Assume we have a problem with linear dynamics and a linear measurement model. However, at each tick of the clock we receive more than one measurement, exactly one of which comes from the process being studied. We will continue to write the states as \mathbf{X}_i , the measurements as \mathbf{Y}_i ; but we now have δ_i , an indicator variable that tells which measurement comes from the process (and is unknown). $P(\mathbf{X}_N | \mathbf{Y}_{1..N}, \delta_{1..N})$ is clearly Gaussian. We want $P(\mathbf{X}_N | \mathbf{Y}_{1..N}) = \sum_{histories} P(\mathbf{X}_N | \mathbf{Y}_{1..N}, \delta_{1..N}) P(\delta_{1..N} | \mathbf{Y}_{1..N})$, which is clearly a mixture of Gaussians. The number of components is exponential in the number of frames — there is one component per history — meaning that $P(\mathbf{X}_N | \mathbf{Y}_{1..N})$ could have a very large number of modes.

The following two chapters discuss main potential sources of multi-modal behaviour in great detail. Chapter 3 discusses the relations between 2D and 3D models of the body, which are generally agreed to be a source of multiple modes. Chapter 4 discusses data association methods. In this chapter, there is a brief discussion of the particle filter, a current favorite method for dealing with multi-modal densities. There are other methods: Beneš describes a class of nonlinear dynamical model for which the posterior can be represented with a sufficient statistic of constant finite dimension [39]. Daum extends the class

of models for which this is the case ([98, 99]; see also [380] for an application and [122] for a comparison with the particle filter). Extensive accounts of particle filters appear in [109, 262, 361].

3

Tracking: Relations between 3D and 2D

Many applications require a representation of the body in three dimensions. Such a track could come from tracking with a 3D representation — perhaps a set of body segments in 3D, modelled as surfaces, triangle meshes or sample points — or by building a kinematic track in two dimensions, then “lifting” it to produce a 3D track. If there is only one camera, relations between the 2D figure and the 3D track are complicated, and may be ambiguous. Ambiguities appear to be less significant in the case where there are multiple cameras; we review this case only briefly (section 3.1).

The heart of the question is the number of possible 3D configurations that could explain a single image. There is no doubt that there are many *if* there is no motion information and *if* only geometric correspondence information is used. In other cases, whether there is any ambiguity is uncertain, and appears to depend quite precisely on the circumstances of measurement. When reconstruction is ambiguous, one expects to encounter multimodal distributions in a tracking problem built around 3D representations, because several distinct inferred 3D configurations could have the same likelihood.

We discuss methods for reconstructing body configuration in 3D from a single view (and perhaps a dynamical history) in considerable detail in section 3.2. This leads us to the (surprisingly complex and unresolved)

question of the extent and nature of ambiguities in 3D reconstruction (section 7.1.1). Furthermore, the lifting procedure adopted may have significant consequences for how we choose to represent body configuration (section 7.1.2). All this provides background information to understand tracking methods that work on 3D body models (section 3.3).

3.1 Kinematic Inference with Multiple Views

If one has multiple views of the body, the problem of reconstruction is considerably simplified. Ideally, the cameras are calibrated, in which case the main difficulty is localizing body parts. At least conceptually, one could lift from one frame using some method chosen from section 3.2, then search all ambiguities, evaluating by backprojecting into the other view. It is more sensible to search configurations with a cost function incorporating all views; this requires the cameras be calibrated. There are generally two questions: the score used to evaluate a particular reconstruction, and how to search for the best reconstruction.

Scores can be computed by explicitly reconstructing a three-dimensional structure from the views, then comparing the body representation to that structure. Cheung *et al.* use a volumetric reconstruction of the person — a quantized approximation to the visual hull — obtained using five views, and then encode kinematic configuration by fitting a set of ellipsoids to the 3D reconstruction with EM [78]; the process is realtime. Kehl *et al.* use an approximate visual hull, estimated by intersecting cones over foreground regions from between 4 and 8 calibrated cameras [221] (figure 3.1). The reconstruction is produced assuming a simple background, so that the cones can be obtained. The body model is a textured 3D mesh, controlled by a skeleton (section 5.1); texture maps are obtained from a modelling view. Tracking is by minimizing distance between the volumetric reconstruction and sample points on the mesh (which are a function of the skeleton's kinematic parameters). The minimization procedure itself is a sophisticated variant of stochastic gradient descent.

It is not necessary to construct the visual hull explicitly. There are numerous methods that use the visual hull implicitly, by comparing the reconstructed 3D model with the silhouette in each view. Carranza *et al.* use an implicit representation, comparing the silhouette of the 3D reconstruction with

silhouettes in each view using graphics hardware [72]. This yields a cost function that can be evaluated very fast, allowing real-time tracking.




Fig. 3.1 **figure 12 of kehl ea, full body tracking using multiple views, multipleview/01467432** Kehl et al. represent the body as a textured 3D mesh, controlled by a skeleton with a texture map obtained from a modelling view. They obtain a volumetric reconstruction from a set of calibrated cameras, then track the body by minimizing distance between sampling points on the mesh and the volumetric reconstruction. The **top row** shows frames from one camera with reprojected skeleton superimposed; the **bottom row** shows the surface reconstruction at the **left** of each frame and the original volumetric reconstruction at the **right**. The reconstruction is accurate, despite some difficulties in the volumetric measurement. Figure from “Full Body Tracking from Multiple Views Using Stochastic Sampling”, Kehl *et al.*, Proc. Computer Vision and Pattern Recognition, 2005 © 2005 IEEE

Stereo matches can give greater depth precision than the visual hull can provide. Plänkers and Fua estimate parameters for a model of the body consisting of a skeleton, metaball muscle model, and skin using stereo and, optionally, silhouette information [338]; the method appears to work with a complex background. Delamarre and Faugeras use a form of iterated closest point matching to produce forces that drive a 3D segment model into correspondence with the silhouette in three calibrated views [101, 102]. Drummond and Cipolla model the body with quadric segments, and track by applying a linearized flow model (as per section 2.3.5; [60, 61]) to a search for edge points close to projected sample points on the model [112] (see also [111] for more information on the formalism, and [110, 113] for information about tracking changes in camera parameters). Shahrokni *et al.* use a similar general approach, but employ a novel texture segmentation model to find silhouette points [387]. They search along a scan line near and approximately normal to the predicted silhouette to find points where there is a high posterior of a texture edge (see also an alternative method for finding texture silhouettes using a classifier in [388]; and using an entropy measure in [386]).

Texture information can be registered to the body model. Starck and Hilton obtain the best configuration of a 17 joint, meshed 3D model of the hu-

man body to fit stereo, silhouette and feature matches for each frame; texture is then reprojected onto the body (in [415]; see also [171, 418]). The texture is then backprojected onto the reconstruction and composited to give a single texture map. In recent work, Starck and Hilton show that correspondences between texture maps induced in separate frames yield temporal correspondences and so information on how relevant surfaces deform [416]. Models of this form allow relatively straightforward synthesis of new views [417]. These methods are oriented to performance capture, and appear to have been demonstrated for simple backgrounds only.

In principle, texture information registered to the body should yield a match score and improve matches, if the texture does not move with respect to the skeleton. We are not aware of methods that use this cue, though it may prove useful if one wants a detailed surface reconstruction of a model wearing tight garments. However, one can use a flow model to register texture from frame to frame. Yamamoto *et al.* use a linear flow model derived from the kinematic model (cf section 2.3.5) with three cameras to obtain good tracks from hand-initialized data; they use three calibrated cameras [475]. The paper describes no difficulties resulting from movement of texture with respect to the body, but we expect that this effect significantly limits the precision of available reconstructions (see also figure 2.4, and the discussion in section 2.1). Theobalt *et al.* describe improved configurations obtained from the method of Carranza *et al.* ([72]) by incorporating an optic flow model to correct the estimates of configuration [436]. Subjects are not wearing very tight clothing, and there again seem to be no difficulties resulting from movement of texture with respect to the body.

Generally, search methods involve either standard optimization techniques or fairly standard variants. However, Deutscher *et al.* use a form of randomized search, described in greater detail in section 3.3.1, to align a 3D model with silhouette edges [104, 107]. Sigal *et al.* use a form of belief propagation, described in greater detail in section 3.3.1.1, to infer configuration in three or four views; the method uses detectors to guide a form of search [396]. Carranza *et al.* use a surface model, controlled by a 17 joint skeleton [72]. The search for a reconstruction at a time instant uses the reconstruction at the previous instant as a start point; however, because motion can be fast, and the sampling rate is relatively slow (15 Hz, p 571), a form of grid search at each limb separately is necessary to avoid local minima. A texture estimate is ob-

tained by rectifying all images to the surface model, and blending.

The most comprehensive and recent discussion of 3D reconstruction issues appears in two papers. Cheung *et al.* give an extensive discussion of representations of the visual hull and methods of obtaining them; the methods they describe can incorporate temporal information, color information, stereopsis and silhouette information [79]. Cheung *et al.* then use these methods to build a body model from a series of calibration sequences, which give both surface and skeleton information [80]. This model is then tracked by minimizing the sum of two scores. The first compares the deformed body model with the silhouettes in each image at a given timestep. The second compares an object reconstruction obtained at a given timestep with the silhouettes in each modelling frame. As authors note, there are 3D situations that are either kinematically ambiguous or at least very difficult for a tracking algorithm of this form. The first occurs when body parts are close together (for example, an arm pressed against the torso) and may lead to a self-intersecting reconstruction. This difficulty appears to be intrinsic to the use of silhouette features. The second occurs when the arm is straight, making rotation about the axis of the humerus ambiguous. The difficulty is that the photometric detail is too weak to force the method to the right configuration of the hand. Curiously, although Mori and Malik have shown that one can obtain landmark positions automatically [296], there appears to be no multiple view reconstruction work that identifies landmarks in several views (with, for example, the method of Mori and Malik, section 3.2.1) and builds a geometric reconstruction this way.

Reducing configuration ambiguity is one reason to use multiple cameras; another is to keep track of individuals who move out of view of a particular camera (e.g. [287, 226, 67]). Currently, this is done at a coarse scale (people are blobs).

3.2 Lifting to 3D

There are a variety of methods for lifting a 2D representation of the body to 3D. Different methods draw from different bodies of technique (kinematics, statistics, computational geometry, optimization, etc.), but the geometry of lifting gives clear bounds to what ambiguities may appear (subsection 3.2.1). The extent of ambiguity appears to depend on whether the ambiguous reconstructions violate kinematic constraints, and whether a dynamical history is

available. The remarkable fact is that reconstruction ambiguity seems to be either quite easily evaded or not to manifest itself at all. Thus, while many papers advocate methods to manage ambiguity, almost any method appears to work — one doesn't see many records of systems failing due to ambiguity. This may be because experiments are poorly conducted; but it is more likely that the implicit folk mythology — that ambiguous reconstructions are quite easily avoided — is true. We discuss this point in section 7.1.1.

3.2.1 Geometric Ambiguity and Lifting by Kinematic Inference

The way that people are imaged means that there are very few cases where a scaled orthographic camera model is not appropriate. One such case to keep in mind is a person pointing towards the camera; if the hand is quite close, compared with the length of the arm, one may see distinct perspective effects over the hand and arm and in extreme cases the hand can occlude much of the body.

Regard each body segment as a cylinder, for the moment of known length. If we know the camera scale, and can mark each end of the body segment — we might do this by hand, as Taylor [433, 434] does and Barrón and Kakadiaris [31, 32] do, or by a strategy of matching image patches to marked up images as Mori and Malik do [296, 297] — then we know the cosine of the angle between the image plane and the axis of the segment, which means we have the segment in 3D up to a twofold ambiguity and translation in depth (figure 3.2 gives examples). We can reconstruct each separate segment and obtain an ambiguity of translation in depth (which is important and often forgotten) and a two-fold ambiguity at each segment.

For the moment, assume we know all segment lengths and the camera scale. We can now reconstruct the body by obtaining a reconstruction for each segment, and joining them up. Each segment has a single missing degree of freedom (depth), but the segments must join up, meaning that we have a discrete set of ambiguities. Depending on circumstances, one might work with from seven to nine body segments (the head is often omitted; the torso can reasonably be modelled with several segments), yielding from 128 to 512 possible reconstructions. These ambiguities persist for perspective images; examples appear in figure 3.4.

Barrón and Kakadiaris show that anthropometric parameters can be esti-




Fig. 3.2 [2d3Dlift/cvpr00.pdf](#), figure 5, p 682, [CJpaper](#) Two 3D reconstructions obtained by Taylor [433], for single orthographic views of human figures. The image appears **left**, with joint vertices on the body identified by hand (the user also identifies which vertex on each segment is closer to the camera). **Center** shows a rendered reconstruction in the viewing camera, and **right** shows a rendering from a different view direction. Figure from “Reconstruction of articulated objects from point correspondences in a single uncalibrated image”, Taylor, Proc. Computer Vision and Pattern Recognition, 2000© 2000 IEEE

mated as well [31, 32]. They do this by constructing a multivariate Gaussian prior on segment lengths, which do not vary much in size (a factor of 1.5 covers the range of human heights from four foot six to six foot nine, which deals with the vast majority of adults). Ratios of body segment lengths vary even less (e.g. see [314, 31, 32]). Barrón and Kakadiaris assume that, in any view, two segments are close to parallel to the image plane, meaning that the ratio of their image lengths is very close to the actual length ratio. They construct a discrete set of possible bodies, and use image length ratios to index into this set to obtain a start point for an optimization procedure that obtains the actual anthropometric parameters by choosing the set that agrees with the image, meets joint limit constraints, and has highest prior probability (this could be seen as a MAP estimate).

The discrete ambiguities can be dealt with in a number of ways. One could ask the user to identify the closer endpoint of each segment (Taylor [433], p. 681). One could simply choose, as Barrón and Kakadiaris appear to do. In detail, their method uses each kinematically acceptable 3D reconstruction as a start point for the minimization procedure described above, and chooses the reconstruction with best value of the objective function. Since this procedure enforces kinematic constraints but does not apply distinct weights to distinct kinematic reconstructions, the unconstrained objective function must have a symmetry corresponding to the reconstruction ambiguity, and so the choice depends largely on random factors within the optimization procedure. It is important to notice that this doesn’t seem to cause any problems, which suggests that substantial kinematic ambiguities might be rather rare. We will pick




Fig. 3.3 Mori+Malik, 2d3dlift/morimecv01.pdf, figures 6 and 7 Mori and Malik deal with discrete ambiguities by matching test image outlines to exemplars, which have keypoints marked [296, 297]. The keypoint markup includes which end of the segment is closer to the view. The images on the **left** show example test images, with keypoints established by the matching strategy superimposed. The resulting reconstruction appears on the **right**. Figure from “Estimating Human Body Configurations using Shape Context Matching”, Mori and Malik, IEEE Workshop on Models versus Exemplars in Computer Vision 2001 © 2001 IEEE

up on this point in section 7.

Mori and Malik deal with discrete ambiguities by matching [296, 297]. They have a set of example images with joint positions marked. The outline of the body in each example is sampled, and each sample point is encoded with a **shape context** (an encoding that represents local image structure at high resolution and longer scale image structure at a lower resolution). Keypoints are marked in the examples by hand, and this marking includes a representation of which end of the body segment is closer to the camera. The outline of the body is identified in a test image (Mori and Malik use an edge detector; a cluttered background might present issues here), and sample points on the outline are matched to sample points in examples. A global matching procedure then identifies appropriate exemplars for each body segment and an appropriate 2D configuration. The body is represented as a set of segments, allowing (a) kinematic deformations in 2D and (b) different body segments in the test image to be matched to segments in different training images. The best matching example keypoint can be extracted from the matching procedure, and an estimate of the position of that keypoint in the test image is obtained from a least-squares fit transformation which aligns a number of sample points around that keypoint. The result is a markup of the test image with labelled joint positions and with which end of the segment is closest to the camera. A 3D reconstruction follows, as above (figure 3.3 gives some

examples).




Fig. 3.4 **figure 2 of sminchisescu+triggs, kinematic jump processes, kinematicambiguity/01211339** Ambiguous reconstructions of a 3D figure, all consistent with a single view, from Sminchisescu and Triggs [408]. The ambiguities are most easily visualized by an argument about scaled orthographic cameras, given in the text, but persist for perspective views as these authors show. Note that the cocked wrist in the leftmost figure violates kinematic constraints — no person with an undamaged wrist can take this configuration. 1

fix this par

Camera scale can be calibrated, and anthropometric parameters estimated, but the remaining discrete set of ambiguities appears to have substantial practical consequences. Furthermore, they appear to apply to whatever particular scheme of reconstruction is adopted, unless the measurement process is more informative than locating segment endpoints. Schemes that use the silhouette (section 3.2.2) and image matches (section 3.2.3) appear to be subject to this set of ambiguities (section 7.1.1 discusses further observations that might help).

Current likelihood models compare some set of predicted with observed image features (typically, silhouette edges), and so must have multiple peaks corresponding to the ambiguities described. These peaks appear in the posterior (figure 3.5). While this makes the multiple peaks predictable, they are still a major nuisance. Typically, at each peak in the likelihood there are some directions where the value of the likelihood varies slowly (small eigenvalues in the Hessian). This is because localization of either landmarks or silhouette points is difficult, and large changes in the estimate of depth to a joint or of a limb angle can result in small changes to image positions. The problem directions tend to move a joint in depth (figure 3.4).




Fig. 3.5 **Figure 4, 2a of Sminchisescu and Triggs, Covariance scaled..., particlefilter/009090509** Several nasty phenomena result from kinematic ambiguities and from kinematic limits, as Sminchisescu and Triggs [404, 407]. Ambiguities in reconstruction — which are caused because body segments can be oriented in 3D either toward or away from the camera, as described in the text — result in multiple modes in the posterior. The two graphs on the left (a and b) show the fitting cost (which can be thought of as a log-likelihood) as a function of the value of two state variables (scaled by their standard deviation). The state variables refer to the kinematic configuration of the 3D model. Note the significant “bumps” from the second mode (the vertical arrows). For reference, there is a quadratic approximation shown as well. Note also the significant deformations of modes resulting from a kinematic limit (the horizontal arrows). This is caused by the fact that no probability can lie on the other side of the limit, so the mode must be “squashed”. 2

3.2.2 Lifting by Minimization

As we have seen (section 3.1), if one has multiple views, the body configuration can be reconstructed by minimizing an error between the image and projected configuration in each view. A wide variety of view errors are available, though most involve a comparison between inferred outline points and an image silhouette. Sminchisescu and Telea show that this approach can produce a reconstruction from a single view ([403]; see also [402]). Their error function includes a term to force the projected body to cover as much silhouette as possible and a term to force the projected body inside the silhouette. It is important to smooth the silhouette (from background subtraction), because noise components on the silhouette boundary can produce a difficult optimization problem. The silhouette is skeletonized and the skeleton is then pruned and “inflated” using a form of distance transform. The method produces good reconstructions, but must experience at least reconstruction ambiguities similar to those experienced by kinematic inference.

Randomized search is a reasonable strategy for attacking the minimization. Sminchisescu and Triggs describe various methods to bias the likelihood function searched by a sampler so that the state will move freely between local minima [406, 405, 410]. Sminchisescu and Triggs exploit an explicit representation of kinematic ambiguities to help this search, by making proposals

for large changes of state that have a strong likelihood of being good [409]. Lee and Cohen use a markov chain Monte Carlo method to search the likelihood, using both a set of image detectors and a model of kinematic ambiguities to propose moves; this gives a set of possible reconstructions for the upper body [245] and the whole body [246].

3.2.3 Lifting by Regression

Assume we are given a set of example pairs $(\mathbf{x}_i, \mathbf{y}_i)$, where \mathbf{x}_i is a vector of measurements of image properties and \mathbf{y}_i is the known 3D configuration of the body for that measurement vector.

We can regard lifting as a **regression problem** — predict \mathbf{y} for a new set of image measurements \mathbf{x} , using the training data. This regression problem has some nasty properties.

- **Dimension:** We expect \mathbf{x} to be drawn from a high-dimensional space. Worse, we expect that the possible \mathbf{x} that we can observe lie on a relatively low-dimensional subspace of the original space. For example, we expect to see arms and legs in a limited range of configurations; we expect to see people with arms of similar appearance; we expect to see people with legs of similar appearance; and so on.
- **Metric distortion:** We do not expect that the distance between \mathbf{x}_i and \mathbf{x}_j necessarily reflects the distance between \mathbf{y}_i and \mathbf{y}_j . For example, two quite distinct body configurations could have very similar images (as a result of geometric ambiguities section 3.2.1).
- **Multiple values:** Worse, we could have two distinct values of \mathbf{y} that are correctly associated with a single value of \mathbf{x} , (as a result of the discrete ambiguity of section 3.2.1).

Notation: To avoid dealing with isolated constants, we will assume that one component of \mathbf{x} always has the value 1.

3.2.3.1 Lifting using the Nearest Neighbour

The simplest regression method is to use the value associated with the nearest neighbour. Athitsos and Sclaroff determine 20 kinematic configuration

parameters from an image of a hand by matching the image to a set of examples [21, 22]. Examples cover a wide range of viewing conditions, and the cost of obtaining the best match (in a total of 107,328 images) limits the number of distinct hand configurations to 26.

One can incorporate dynamical information into the distance cost matching entire 3D motion paths to 2D image tracks. Howe computes a match cost frame by frame, by comparing rendered motion capture data from the CMU Motion Capture collection (<http://whereisit>) with image silhouettes [187]. Views are again assumed lateral and orthographic, and are sampled every 10° around the body. Translation and scale could be handled either by sampling, or by obtaining estimates from a bounding box. The comparison is scored with a chamfer distance. Write

$$H(S_1, S_2) = \sum_{p \in S_1} \min_{q \in S_2} d(p, q)$$

(noting a similarity with the Hausdorff distance, section 2.2.2.2), θ^l for the 3D configuration of the l 'th frame of motion capture data with respect to the camera (meaning that rotation, translation and scale are encoded here), P_{θ^l} for the set of pixels covered by a rendering of θ^l , and P_{S_j} for the pixels lying in the j 'th silhouette. The comparison between θ^l and S_j is now scored as

$$M(\theta^l, S_j) = H(P_{\theta^l}, P_{S_j}) + H(P_{S_j}, P_{\theta^l})$$

Now write the (unknown) value of θ at time i as Θ_i — this value could be any one of the available θ^l . Howe then constructs a cost linking frames of motion capture $\Delta(\Theta_i, \Theta_{i-1})$; this cost could include a charge for extreme camera motions, though the paper does not explicitly describe this (the cost used charges for large changes in body configuration). The motion is lifted by applying dynamic programming to

$$C(\Theta_1, \dots, \Theta_N) = \sum_{i=2}^N \Delta(\Theta_i, \Theta_{i-1}) + \sum_{i=1}^N M(\Theta_i, S_i)$$

There are too many frames of motion capture to implement an exact dynamic programming solution, and we allow only values θ^l of Θ_i such that $M(\theta^l, S_i)$ is less than some threshold. The method appears to produce solutions that are unambiguous, which is consistent with the view that 3D reconstruction ambiguities are probably a phenomenon of short, rather than long, time-scales.

There is also some useful evidence that reconstruction errors or uncertainties do not propagate over long time-scales (figure 3.6). However, there is no attempt to use either N-best dynamic programming or beam search to identify 3D reconstructions that have cost comparable to the best cost, but are significantly different.




Fig. 3.6 01384804, figure 7 (8 is incomprehensible, and so omitted) Howe's formulation lifts to 3D by comparing projected motion capture data with image silhouettes [187]. There is a frame-frame cost for the reconstruction, and the final 3D lift is obtained by dynamic programming. In a formalism like this, one could reasonably fear that a mistaken reconstruction in one frame might result in an entirely wrong path. In practice, this does not occur. The graph is obtained by constraining the first lifted frame of a sequence to each of a 1000 different (incorrect) states; the plot shows the number of distinct states found in the succeeding frames for each path, as a function of frame. The local image evidence quickly overwhelms the effect of history; by the 10'th frame, there are only two distinct states. Figure from "Silhouette Lookup for Automatic Pose Tracking", Howe, Proc. IEEE Workshop on Articulated and Non-Rigid Motion, 2004 © 2004 IEEE

3.2.3.2 Snippets and Cameras

This work suggests that, while a single frame reconstruction might be ambiguous, a match from a short 2D track to a short 3D track might not be (in section 7, we lay out evidence it is not). Howe *et al.* compare projected motion

capture data with image tracks, but now use posterior inference to estimate dynamic parameters [188]. These parameters are an encoding of “snippets” — 11 frames of motion capture data — which are clustered using a mixture of Gaussians. Each 11 frame section of the track produces a snippet with maximal posterior, and the snippets are blended into one another to give a 3D reconstruction. While authors acknowledge the tracker loses track after a while, the lifting procedure appears to be robust and effective.

Ramanan and Forsyth use a similar approach, but apply constraints to camera dynamics, too ([355]; see also [354]). They assume that views are lateral, estimate scale and translation from the image, and sample the remaining camera parameter (rotation about the vertical axis). They constrain the camera speed, and charge for large motions in three dimensions. The best matching sequence can then be obtained by dynamic programming. The method cannot recover the motion in depth of the root, but successfully recovers the configuration of the body with respect to the root and all root parameters but depth.

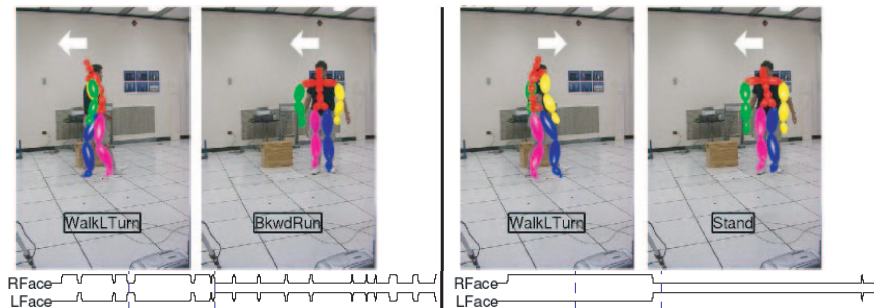


Fig. 3.7 **Figure 5.4 from Ramanan’s thesis** Left frames are taken from a walking sequence, matched to motion capture data using the method of Ramanan and Forsyth [355]. Matches are independent from frame to frame. Note that the lateral view of the body (**far left**) is ambiguous, and can be reconstructed inaccurately. This ambiguity does not persist, because the camera cannot move freely from frame to frame. **Right frames** show reconstructions obtained using dynamic programming to enforce a model of camera cost. The correct reconstruction is usually available, because the person does not stay in an ambiguous configuration. The frames are taken from a time sequence, and the graphs **below** show an automatically computed annotation sequence — facing left vs. facing right — as a function of time. Note that the case on the **left** shows an essentially random choice of direction when the ambiguity is present (the person appears to flip from facing left to facing right regularly). This is because the free rotation of the camera means the ambiguity appears on a per-frame basis. For the case on the **right**, the smoothing created by charging for fast camera rotations means that the labels change seldom (and are, in fact, correct). Figure from Ramanan’s UC Berkeley PhD thesis, “Tracking People and Recognizing their Activities”, 2005© 2005 D. Ramanan

The discrete ambiguity in configuration is handled by incorporating information about surrounding frames into the match cost. In particular, the cost of matching a given image frame with a given motion capture frame is averaged over a window of image (resp. motion capture) frames centered around the frame under consideration. This means that the match uses an implicit (in the collection of motion capture) dynamical model to resolve these discrete ambiguities, at the cost of not being able to lift configurations that are not in the motion capture data.

The charge for camera rotation is reasonable, because cameras do not usually swing around the body by very large amounts, but it is also important, because Ramanan and Forsyth’s model does not match heads and so has difficulty telling which way the body is facing for lateral views, particularly when the limbs are in line with the body (figure 3.7). This results in a lateral view of a standing person can be interpreted as facing either right or left; the camera rotation charge means that, if the person walks off — and so reveals the direction in which they are facing — this information can be propagated.

3.2.3.3 Regressing Pose against the Image

Rosales and Sclaroff use of a collection of local experts (“specialized mappings”) to regress hand configuration against image appearance [367]. Shakhnarovich *et al.* train with a data set of 3D configurations and rendered frames, obtained using POSER (a program that renders human figures, from Creative Labs). They show error rates on held out data for a variety of regression methods applied to the pool of neighbours obtained using parameter sensitive hashing. Generally, performance improves with more neighbours, with using a linear (rather than constant) locally weighted regression, and if the method is robust. The best is a robust linear locally weighted regression. Their method produces estimates of joint angles with RMS errors of approximately 20° for a 13 degree of freedom upper body model [389]; a version of this approach can produce full 3D shape estimates [162]. Liu *et al.* demonstrate a full body reconstruction from silhouettes in five views using a similar regression model; the reconstruction is not evaluated directly, but is used to control motion synthesis [360].

3.2.3.4 Disambiguation with the Immediate Past

A major difficulty with this procedure is the possibility that a single set of image features may predict multiple poses. This could be a result of weaknesses in image features — for example, it is hard to tell which way the actor is facing in a lateral view of a standing person with current image features — but is more likely the consequence of the kinematic ambiguities described above. Reconstructions performed in the past could disambiguate the current reconstruction. Brand links images with motion capture by fitting HMM’s to both motion capture data and image data; these HMM’s share a dynamical model [56]. The HMM’s are fitted with a variant fitting algorithm which tends to obtain models with relatively low entropy (there is some discussion in [56]; more in [55, 57]). Reconstruction in 3D is obtained by inferring a state sequence from image data, then choosing a sequence of emitted states from the motion capture model, using a smoothed approximation rather than the Viterbi sequence.

We could think of pose as lying on a set of distinct “sheets”, each of which is a single valued function of image features, and then build distinct models for each sheet. This leads to tricky problems in identifying the sheets, however. Agarwal and Triggs observe that the pose in the previous frames, if correctly computed, should give a good guide to the current pose — one is unlikely to jump from sheet to sheet in a single frame [3, 6]. This observation implies that, while $\mathbf{y}_t(\mathbf{x}_t)$ might be a multiple valued function, $\mathbf{y}_t(\mathbf{x}_t, \mathbf{y}_{t-1}, \mathbf{y}_{t-2})$ is not. At reasonable sampling rates, the pose in the last two frames should give a fair estimate of the pose in the current frame. Agarwal and Triggs first construct a regressed estimate of the pose in frame t , $\hat{\mathbf{y}}_t$ from \mathbf{y}_{t-1} and \mathbf{y}_{t-2} using a linear regression. They then compute a regression estimate of \mathbf{y}_t from \mathbf{x}_t and $\hat{\mathbf{y}}_t$, using a regression vector machine trained with a variant algorithm. The method produces estimates of joint angles with RMS errors of 4° for 55 degrees of freedom (3 angles per joint for an 18 joint skeleton, and 1 orientation DOF with respect to the camera). We expect the method to behave badly at singularities of the pose (figure 7.1). In a more recent paper, Agarwal and Triggs encode the “sheets” implicitly with a latent variable, and obtain improved reconstructions [5].

3.3 Multiple Modes, Randomized Search and Human Tracking

We have clear evidence that tracking a 3D representation of the body can result in multiple modes in the posterior and that these modes do not look Gaussian locally (figure 3.5; but see section 7). The need to manage these modes has spawned a number of methods, all of which are forms of randomized search.

Particle filters should be seen as a form of randomized search. One starts a set of points that tend to be concentrated around large values of the posterior. These are pushed through the dynamical model, to predict possible configurations in the data. The result is a sampled representation of the prior. The predictions are compared to the data, and those that compare well are given higher weights, yielding a sampled representation of the posterior. This simple view provides some insight into why particle filters in their most basic form are not particularly well adapted to kinematic tracking.

There is a problem with dimension. The state vector for most kinematic tracking problems must be high dimensional. One expects to encounter at least 20 degrees of freedom (one at each knee, two at each hip, three at each shoulder, one at each elbow and six for the root) and quite possibly many more. This means that mismatches between the prior and the likelihood can generate serious problems. Such mismatches are likely for three reasons.

First, the body can move quickly and unexpectedly, meaning that probability must be quite widely spread in the prior to account for large accelerations. It is hard to be clear on how much uncertainty there is in the state of the body at some time given the past, and there are fair arguments either way (section 7.1.3). However, fast movements do occur, and current methods are forced to have fairly diffuse dynamical models to cope with them.

Second, the likelihood has multiple peaks, which can be very narrow. Narrow peaks occur because some body segments — forearms are a particularly nasty example — have relatively small cross-section in the image, and so only a small range of body states will place these segments in about the right image configuration. Multiple peaks occur because there tend to be numerous objects that look somewhat like body segments (long, narrow, parallel sides, constant colour). We are now using the predictions of the prior to find the largest narrow peak in a high-dimensional likelihood — for this to have any

hope of success, the predictions need to be good or to occur in very large numbers. But we know the predictions will be poor, because we know people can generate fast, unexpected movements.

Third, detectors used to produce a likelihood model may be inaccurate. This can result in small errors in inferred state, which in turn produce potentially large changes in state from frame to frame. As Sminchiescu and Triggs point out ([407], p. 372), this suggests using a relatively diffuse dynamical model as an insurance policy.

The key idea in particle filters is the randomized search. One might abandon, or at least de-emphasize, probabilistic semantics, and focus on building an effective search of the likelihood. The key difficulties are that the peaks in the likelihood are narrow (and so easy to miss) and that the configuration space is high-dimensional (so that useful search probes may be difficult to find). The narrow peaks in the likelihood could be dealt with by annealing, and good search probes may be found by considering the ambiguity of 3D reconstructions. We review these approaches in section 3.3.1.5.

3.3.1 Randomized Search with Particle Filters

There are a series of approaches to deal with problems created by the dimension of the state space. First, we could refine the search using importance sampling methods. Second, we could use sequential inference methods to obtain more efficient samples of the prior. Third, we could build lower-dimensional dynamical models. Finally, we could build more complex searches of the likelihood.

3.3.1.1 Importance Sampling

Importance sampling is a method to concentrate samples in places that seem likely to be useful. Assume we have a distribution $g(\mathbf{X}_t)$ from which we can draw samples, and which is a better guide to the likelihood than the prior $P(\mathbf{X}_t | \mathbf{Y}_0, \dots, \mathbf{Y}_{t-1})$ is. We can then draw samples \mathbf{X}_{ti} from $g(\mathbf{X}_t)$. Then the set of samples

$$(\mathbf{X}_{ti}, \frac{P(\mathbf{X}_t = \mathbf{X}_{ti} | \mathbf{Y}_0, \dots, \mathbf{Y}_{t-1}) P(\mathbf{Y}_t | \mathbf{X}_t = \mathbf{X}_{ti})}{g(\mathbf{X}_{ti})})$$

is a representation of the posterior, as can be established by pattern matching to the expressions above. Given several plausible importance functions, one could use a mixture of these functions and the prior as an importance function. Drawing samples from this mixture is straightforward; one draws a sample according to the mixing weights, and uses this to choose a sampling strategy. Image observations are a natural source of importance functions. Isard and Blake use this approach to track hands and forearms [201], using a skin detector to build an importance function. Rittscher and Blake use importance sampling methods to track contours of motions drawn from two classes (pure jump and half star jump); the tracker maintains a representation of posterior on the motion class, which can be used to distinguish between motion classes successfully [362]. Forsyth uses edge detector responses as a source of proposal mechanisms to find simple boundaries [130], and Zhu *et al.* — who call the approach data driven MCMC — use image observations to propose segmentations [445, 444, 486]. We are not aware of the method being used for kinematic tracking; however, it is a way to unify the more successful kinematic tracking methods of section 4.2 with particle filter based inference.

If one models a person with a tree-structured kinematic model, then identifying each body segment in the image is a matter of dynamic programming (we discuss this issue in greater detail in section 4.2.3). However, adding temporal dependencies produces a structure that does not allow for simple exact inference, because the state of a limb in frame t has two parents: the state in time $t - 1$, and the state of its parent in frame t (figure 3.8). Loopy propagation is a method for approximate inference on graphical models which are not trees. One constructs a spanning tree, passes messages with the usual algorithm on that spanning tree, and then repeats for other choices of spanning tree. This is an approximation, because some probabilities are overestimated as a result of cycles in the graph; experiment shows that, under many circumstances, the approximation gives usable and helpful results. Useful accounts of this method appear in [481, 302, 464].

Sigal *et al.* use loopy propagation, representing messages passed between nodes using a set of particles [396]. Their template is a 3D model of a person with links both in time and in space learned from data. The likelihood is




Fig. 3.8 If one models a person with a kinematic chain, then determining where a given person appears in a static image involves inference on a tree structured graphical model. On the **left**, a graphical model illustrating this point. In the usual language of graphical models, open nodes represent unknowns, arrows represent dependencies, and shaded nodes represent measurements. Each open node encodes the state (for example, image position; image position and orientation; 3D position, orientation and scale; and so on) of the body segment implied by the label (t: torso; lul: left upper leg; and so on). The arrow represents a model of $P(\text{variable at head}|\text{variable at tail})$. The filled nodes represent various detector responses. Notice that each open node has at most one parent, so the open nodes form a tree, so that inference is a matter of dynamic programming (or, equivalently, message passing; section 4.2.3 or a text such as [135, 273]). On the **right**, we show what happens when one has temporal dependencies. We show only two frames (there's enough clutter in the drawing), and the gray arcs are temporal links. The graphical model becomes much more complex. Most open nodes now have two parents, a spatial parent and a temporal parent, and this means that exact inference is impractical.

modelled with a conditional exponential model, where

$$P(\mathbf{Y}|\mathbf{X}) \propto \exp\left(-\sum_i \lambda_i g_i(\mathbf{X}, \mathbf{Y})\right)$$

with parameters λ_i learned from data. Such models, often called **maximum entropy models** and quite popular in the language modelling community, are commonly fitted by maximizing likelihood (which requires computing the partition function), using an algorithm known as **iterative scaling** (see [96, 335, 40, 370, 210]). Sigal *et al.* use a series of detectors which are tuned to body parts (but not, in the nature of such detectors, particularly reliable; otherwise there'd be nothing to do) to produce an importance function. Some percentage of messages passed to limb nodes are drawn from this importance function, giving strong suggestions about the configuration in 3D of a particular body segment. They demonstrate tracks of people in 3D from three views. Unusually, there is a strong evaluation component, which we describe in section 4.3.




Fig. 3.9 **Figures 1 and 7 from Particlefilters/013150163** Sigal et al. build a 3D model of a person as a set of segments [396]. Again, the state of each segment but the root has two parents — the corresponding segment in the previous frame, and that segment's parent in the model (left). This yields an inference problem that is too difficult in general to do exactly. Sigal et al. track in multiple views using a form of particle filter adapted for loopy belief propagation. The image likelihood is a conditional exponential model. Authors use a combination of segment detectors and uniformly distributed samples to propose likely configurations of limbs in the image; these are incorporated in the inference procedure as importance functions. The figure on the right shows camera outputs with superimposed information for two of four views (rows); column (a) shows limb segments proposed by the detector; (b) shows proposals from a uniform distribution; (c) shows samples from the belief distribution after 30 frames of belief propagation; and (d) shows the state with the highest belief.

3.3.1.2 Partitioned Sampling

Partitioned sampling is a variant of importance sampling that uses a sequence of samples within each time slice. Assume that the state vector \mathbf{X} has several components; notation etc. is much simpler if we assume only two, and the more general case follows, so we shall work with two and write $\mathbf{X} = (\mathbf{x}_1, \mathbf{x}_2)$. We will also drop the subscript for time to simplify notation. Now assume that we have an importance function $I(\mathbf{X})$ that is a good guide to the likelihood (what this means will become apparent), and can be factored as $I(\mathbf{X}) = I_1(\mathbf{x}_1)I_2(\mathbf{x}_1, \mathbf{x}_2)$ Now if \mathbf{u}_i is a set of IID samples of $P(\mathbf{x}_1)$, then

$$(\mathbf{u}_i, I_1(\mathbf{u}_i))$$

represents a probability distribution proportional to $P(\mathbf{x}_1)I_1(\mathbf{x}_1)$. Take this representation and resample with replacement according to the weights, to obtain $(\mathbf{u}_j, 1)$ which must also be a representation of that distribution. Now obtain \mathbf{v}_{kj} , which are IID samples of $P(\mathbf{x}_2|\mathbf{x}_1 = \mathbf{u}_j)$. Then

$$((\mathbf{v}_{kj}, \mathbf{u}_j), I_2(\mathbf{u}_j, \mathbf{v}_{kj}))$$

represents a probability distribution proportional to

$P(\mathbf{x}_2|\mathbf{x}_1)P(\mathbf{x}_1)I_1(\mathbf{x}_1)I_2(\mathbf{x}_1, \mathbf{x}_2)$. Take this representation and in turn resample with replacement according to the weights, to obtain $(\mathbf{v}_j, 1)$, which is also a representation of that distribution. Finally,

$$\left((\mathbf{v}_{lj}, \mathbf{u}_j), \frac{P(Y|\mathbf{X} = (\mathbf{v}_{lj}, \mathbf{u}_j))}{I_1(\mathbf{u}_j)I_2(\mathbf{u}_j, \mathbf{v}_{lj})} \right)$$

represents the posterior. Notice that we have omitted various if's, and's and but's to do with the support of the importance function and so on, to get to this point. The advantage of this strategy is that we have guided the search of the likelihood using our importance function; in particular, the first resampling step discards particles that lie in spots where there is evidence — supplied by the importance function — that the marginal of the posterior will be small. Throwing these particles away means that, when we elaborate the particles to represent the whole state, the resulting particles should tend to lie in places where the likelihood is large. Of course, all this depends on the quality of our likelihood functions. MacCormick and Isard track hands using partitioned sampling [271]. MacCormick and Blake use this method to track multiple objects [270, 269], where one needs a method to avoid both tracks lying on the same object. The importance functions are obtained by considering each object separately, and the likelihood function is a mixture of three cases: no objects in the tracker gate, one object in the tracker gate, and two objects in the tracker gate. Again, we are aware of no kinematic trackers of humans that use this method, but see it as a way to unify the more successful kinematic tracking methods of section 4.2 with particle filter based inference.

3.3.1.3 Lower Dimensional State Models

Sidenbladh *et al.* build a 3D model of a human as a kinematic chain, with state encoded as the configuration and velocity of each element of this chain with respect to its parent, and the root with respect to the camera. Each segment of the model has an attached encoding of appearance, and the likelihood is computed by comparing a rendering of the state with the image, using the appearance encoding. There is a separate constant likelihood term for self-occluded segments, and a discount term for foreshortened segments, because foreshortening of a segment causes texture foreshortening. The tracker is initialized by hand. Tracks are obtained using a straightforward particle filter,




Fig. 3.10 **Figure 6 and 7 from Particlefilters/eccv00.pdf** Sidenbladh et al. use particle filters to track a 3D model of a walking person, using a reduced dimensional dynamical model fitted to motion capture data of walking people. This means that the dynamics are more predictable, and so the search of the likelihood is more effective; the difficulty is that one must know the activity before being able to track. On the **left**, a track of a walking person who turns during the walk. The 3D reconstruction of this track is shown **below left**. On the **right**, a “track” of a walking person, initialized as on the left, but now ignoring image data; this illustrates the strength of the prior. In particular, the “track” continues to walk, but does not turn when the subject turns.

using a random walk dynamical model and also using a dynamical model specialized to walking. This walking model is obtained by principal components analysis on motion captured walk data. The appearance model appears to have dynamics to account for changes in illumination; the authors point out that this advantage over a fixed appearance template comes at the cost of potentially increased tracker drift. The random walk model is shown to track a two segment arm with reasonable success, but authors indicate that more complex kinematic models are difficult to track this way. The advantage of a low dimensional model of walking dynamics is that the effective dimension of the state space at the $k + 1$ 'th frame is relatively small, and this relatively tight motion prior allows quite good tracking of a walking figure (figure 3.10). The difficulty with this approach is that one might need to choose which activity is occurring to be able to track it, and that seems difficult to do.

3.3.1.4 Probabilistic Searches of the Posterior

Choo and Fleet implement a more extensive search of the posterior using a Markov chain Monte Carlo (MCMC) method [82]. They interpret the particles at a particular step as a set of initial states for an MCMC sampler; this sampler then runs independently on each state. Any such sampler will eventually produce a fair sample of the posterior. It is reasonable to expect that

running an MCMC sampler on a set of particles will produce IID samples of the posterior. These can, in turn, be passed through the particle filter and refined again. Choo and Fleet use Duane *et al.*'s hybrid Monte Carlo method to obtain samples (see [115, 306]; there is a brief account in [134]), but other methods might be used. The method is used to compute 3D configurations from images of markers. It has not been shown to cope with the dramatic problems with local maxima one associates with texture and clutter, and it seems unlikely that it can. The difficulty here is that it may take very many steps of the MCMC method to produce samples that have “forgotten” their start point. In practice, it is extremely difficult for such a sampler to pass from one local maximum of the posterior to another; this means that such a sampler is unlikely to overcome the problems created by a posterior with many narrow peaks (see [149]; in some applications, for example where there is a symmetry in the posterior, this may not be a nuisance [134], but one cannot rely on MCMC methods to discover all peaks in a posterior without quite strong proofs of good mixing behaviour).

3.3.1.5 Annealing

A variety of search strategies are available. One strategy is to launch an **annealed search** of the likelihood. We do this by defining a set of intermediate weighting functions, to obtain $w_0(\mathbf{X}) = P(\mathbf{Y}|\mathbf{X})$, $w_1(\mathbf{X})$, \dots , $w_M(\mathbf{X})$, where w_k is a somewhat smoother version of w_{k+1} . At any time step we have \mathbf{u}_i a set of IID samples of $P(\mathbf{X})$. Instead of weighting these samples by the likelihood, we weight by w_M . We resample with replacement according to the weights and reset the weights to one, yielding \mathbf{u}_j . We take each sample and add noise drawn from a normal distribution with zero mean. We now weight the resulting samples using w_{M-1} . This process continues until each sample is weighted using the likelihood. Deutscher *et al* use this scheme to track a person moving using a 3D model viewed with multiple cameras [104, 107] (figure 3.11). The likelihood is evaluated using both image values within and edge points near the projected outline; annealing in effect uses a smoothed version of this (very peaky) likelihood function to guide samples toward peaks in the likelihood. This method can be given exact probabilistic semantics by interpreting the annealing procedure as an importance function, an observation due to Neal [307, 308, 309]. Deutscher




Fig. 3.11 **Figure 8 of DeutscherBlakeReid** Deutscher et al. [104] track a moving person in 3D using an annealed particle filter. In effect, particles are passed through the dynamic model, then weighted with a smoothed version of the likelihood. They are resampled according to weights, then perturbed randomly and weighted using a less heavily smoothed likelihood. This concentrates particles in regions where the likelihood is likely to be high. The process continues for some number of layers of annealing. The figure shows tracks for a particular set of frames using three different algorithms. On the **left**, a straightforward particle filter, which loses track fairly quickly because searching a peaky likelihood using a smooth prior doesn't work well. In the **center**, the results of one layer of annealing. Notice that the right leg is poorly tracked, but the track has improved. On the **right**, the results from ten layers of annealing. Notice the much improved track. The particles no longer have any probabilistic semantics, however, and the ability of the method to deal with clutter and texture — which can hugely complicate the likelihood function — is not proven. Figure from “Articulated Body Motion Capture by Annealed Particle Filtering”, Proc. Computer Vision and Pattern Recognition, Deutscher *et al.*, 2000, © 2000 IEEE

et al have shown that performance improvements are available by using partitioning methods together with an annealed particle filter (figure 3.12). All examples show isolated persons on black backgrounds; there is no evidence that the annealing is powerful enough to cope with the rich range of local likelihood peaks that can result from, say, texture or clutter.

3.3.2 Multiple Probes from Covariance Analysis

One difficulty with a sampled model of the posterior is that we don't know if there are larger values of the posterior close to each sample. We could regard each sample as a plausible start point for a search of the posterior. We are now no longer building a set of particles that explicitly represents the posterior in the sense above, but are using multiple states to represent the prospect that the posterior is multi-modal. Each state lies on a mode in the posterior, and we attempt to ensure that all modes have a state. The origins of this approach lie with Cham and Rehg [75], who use it to track a 2D kinematic model of the body.

Sminchisescu and Triggs elaborate this search by analysis of the Hessian of the log-posterior [404, 407]. They track a 3D model of a person, which has




Fig. 3.12 **Figure 10 and 11 of Particlefilters/deutscheretalcvpr2001** Deutscher et al ([107]; see also [105]) show that one can use partitioning methods with the annealed particle filter. They track a 3D model of a person in a single view. On the **left**, a track and the inferred 3D configuration for a running person. On the **right**, a track and the inferred 3D configuration for a person doing a handstand. Again, there are no probabilistic semantics, and again the ability of the method to deal with clutter and texture is not proven. Figure from Automatic Partitioning of High Dimensional Search Spaces Associated with Articulated Body Motion Capture, Proc. Computer Vision and Pattern Recognition, Deutscher *et al.*, 2001, © 2001 IEEE

parameters giving the kinematic configuration, relative proportions of segments, and deformations of the surface skin. Sminchisescu and Triggs do not use a dynamical model. However, they do encode joint limits, and so must represent a model of $P(\mathbf{X}_k | \mathbf{Y}_k)$ (which we call the posterior in what follows; note that only the current measurement is involved). They regard the reconstruction at frame $k - 1$ as an initial point for a search of the posterior at frame k . The likelihood is evaluated by comparing projected model points with image points, using values of edges and other image features. What is known about state is represented by a collection of tuples; the i 'th tuple (c_i, μ_i, Σ_i) , contains a weight c_i , a state value μ_i and a covariance matrix Σ_i . Each state value gives the state at a mode of the likelihood. The covariance is the Hessian of the negative log-posterior at the mode, and the weight is the value of the posterior at the mode. Weights are normalized to sum to one.

This information is propagated from the $k - 1$ 'th frame to the k 'th frame by using these tuples to launch searches of the likelihood. The search proceeds by:

- **Choosing a tuple to propagate** by drawing one of the initial tuples randomly according to the weight. Assume we have drawn tuple i .
- **Computing a local covariance scaling** by obtaining the k directions in the Σ_i where the least change in posterior is likely —

these directions are those in which the state is most uncertain — by a singular value decomposition, and computing the restriction of Σ_i to this space; call the result Σ'_i .

- **Generating a new tuple** by generating a sample \mathbf{s} distributed as $N(\mu_i, s\Sigma'_i)$ for some scale parameter s (it is wise to have $s > 1$). We start an optimization procedure for the posterior at \mathbf{s} ; this produces \mathbf{s}' . The weight for the new tuple is the value of the posterior at this point; the mean is \mathbf{s}' ; and the covariance is the Hessian of the negative log-posterior at \mathbf{s}' .

These steps are repeated multiple times, to produce a set of tuples representing the posterior. This set is pruned to remove tuples that represent the same mode — the states will be the same — and the result represents the new posterior. Numerous variants of this method are possible; for example, it is natural to produce a large pool of tuples, prune duplicates, and then keep only the K best. Performance comparisons between these methods appear in [407].

Appendix I: Particle Filter Basics

Probability distributions should be thought of as devices that are used to compute expectations; this means that any procedure that can produce good estimates of expectations with respect to a distribution is a representation of that distribution. Assume we wish to represent a distribution $Q(X)$, and we can sample some other distribution $P(X)$. Take a set of points \mathbf{s}_i which are independent identically distributed samples of $P(\mathbf{X})$. If we attach a weight $w_i = kQ(\mathbf{s}_i)/P(\mathbf{s}_i)$ (where k is a constant, perhaps unknown) to each point, the resulting set of pairs $\{(\mathbf{s}_i, w_i)\}$ is a representation of $Q(\mathbf{X})$. This is because

$$\frac{\sum_i f(\mathbf{s}_i)w_i}{\sum_i w_i}$$

is an estimate of

$$E_Q(f) = \int f(X)Q(X)dX$$

(by the weak law of large numbers). Samples are commonly thought of as **particles**. Notice, by the way, that it is usual in the vision literature to normalize weights so that $\sum_i w_i = 1$, and in the statistics literature to divide by the sum of weights as we have done; Liu gives an argument that the latter is better practice [262].

Priors into posteriors: A useful trick allows us to obtain a representation of a posterior $P(\mathbf{X}|\mathbf{Y} = \mathbf{y})$ from a representation of a prior $P(\mathbf{X})$. Assume the prior is represented by $\{(\mathbf{s}_i, w_i)\}$. Now

$$\int f(\mathbf{X})P(\mathbf{X}|\mathbf{Y} = \mathbf{y})d\mathbf{X} = \frac{1}{K} \int f(\mathbf{X})P(\mathbf{Y} = \mathbf{y}|\mathbf{X})P(\mathbf{X})d\mathbf{X}$$

where

$$K = P(\mathbf{Y} = \mathbf{y}) = \int P(\mathbf{Y} = \mathbf{y}|\mathbf{X})P(\mathbf{X})d\mathbf{X}$$

What this means is that

$$\frac{\sum P(\mathbf{Y} = \mathbf{y}|\mathbf{X} = \mathbf{s}_i)w_i}{\sum w_i}$$

is an estimate of K . In turn, this means that $(\mathbf{s}_i, w_i P(\mathbf{Y} = \mathbf{y}|\mathbf{X} = \mathbf{s}_i))$ is a representation of the posterior.

Resampling: Finally, we can turn one representation of a distribution $P(\mathbf{X})$ into another representation of the same distribution by resampling the particles. The most important case is if we form a distribution π on the index from w_i , where the probability of drawing the i 'th index is $w_i/\sum w_i$. We now produce a new representation $\{(\mathbf{s}_l, 1)\}$, where the l are independent identically distributed (henceforth IID) samples from π . This is a representation of $P(\mathbf{X})$. The easiest way to see this is to assume we have drawn a very large number of samples. Now if we form $\sum f(\mathbf{s}_l)$, we will see samples from the original set of particles with a frequency proportional to the original weights.

Prediction: Now assume that we have a set of points and weights (\mathbf{s}_i, w_i) which represents the distribution $P(X_{k-1}|Y_0, \dots, Y_{k-1})$. Write $(\mathbf{q}_{j(i)})$ for a set of $N(i)$ points, indexed by j , drawn from $P(X_k|X_{k-1} = \mathbf{s}_i)$. We can represent $P(X_k|Y_0, \dots, Y_{k-1})$ using the set $(\mathbf{q}_{j(i)}, w_i/N(i))$. This is because

$$\frac{\sum_i \sum_j f(\mathbf{q}_{j(i)})w_i/N(i)}{\sum_i w_i/N(i)}$$

is an estimate of

$$\int \left[\int f(X_k)P(X_k|X_{k-1})dX_k \right] P(X_{k-1}|Y_0, \dots, Y_{k-1})$$

(one can see this by pattern matching to the previous paragraph). This could be impractical, because the number of samples might grow — perhaps

$N(i) > 1$ for all i — but we ignore this possibility for the moment. Notice that \mathbf{X}_t might live in a space of complex form — for example, multiple components of different dimension, the product of two spaces encoding very different types of state etc. — but as long as we can obtain samples, no difficulty results (e.g. [76, 170]; various formulations and applications appear in [109, 262, 361]).

Correction: The result of the step above is a set of weighted samples, $(\mathbf{q}_{j(i)}, w_i)$, which represent $P(\mathbf{X}_k | Y_0, \dots, Y_{k-1})$. This is a prior. We apply the recipe above, and obtain a representation of the posterior as $(\mathbf{q}_{j(i)}, w_i P(\mathbf{Y}_k = \mathbf{y}_k | \mathbf{X}_k = \mathbf{q}_{j(i)}))$.

Practicalities

The recipe above is not much use. First, we cannot manage a pool of particles that either shrinks or grows too fast. The easiest way to deal with this problem is to ensure that $N(i) = 1$, though this may not be the best. Second, as one can discover with a simple experiment, the variance of the weights tends to increase as time goes on (there is a theorem to this effect, in [262]). What this means is that one weight gets very much larger than all the others, which means that, in effect, only one sample is being used to represent the distribution. This doesn't affect the expected value of our estimates, but it does mean that their variance might be extremely high, and so they are of no practical use. This effect — known as **sample impoverishment** — can be countered by using an estimate of the variance of the weights to determine when it is occurring, and then resampling the particles. This is effective because we will tend to drop particles with low weight and keep multiple copies of particles with high weight. These latter will spawn multiple distinct samples in the next prediction stage. As a result, we can reasonably hope that our particles tend to be concentrated near spots where the posterior is large. Algorithm 3 gives a practical particle filter; one should realize that numerous variants are possible (see [109, 262, 361] for examples).

The particle filter can be a powerful and effective inference tool. It should be seen as a form of randomized search. One starts a set of points that tend to be concentrated around large values of the posterior. These are pushed through the dynamical model, to predict possible configurations in the data. The predictions are compared to the data, and those that compare well

Initialization: Represent $P(\mathbf{X}_0)$ by a set of N samples

$$\{(\mathbf{s}_0^{k,-}, w_0^{k,-})\}$$

where

$$\mathbf{s}_0^{k,-} \sim P_s(\mathbf{S}) \text{ and } w_0^{k,-} = P(\mathbf{s}_0^{k,-})/P_s(\mathbf{S} = \mathbf{s}_0^{k,-})$$

Ideally, $P(\mathbf{X}_0)$ has a simple form and $\mathbf{s}_0^{k,-} \sim P(\mathbf{X}_0)$ and $w_0^{k,-} = 1$.

Prediction: Represent $P(\mathbf{X}_i|\mathbf{y}_0, \mathbf{y}_{i-1})$ by

$$\{(\mathbf{s}_i^{k,-}, w_i^{k,-})\}$$

where

$$\mathbf{s}_i^{k,-} = f(\mathbf{s}_{i-1}^{k,+}) + \xi_i^k \text{ and } w_i^{k,-} = w_{i-1}^{k,+} \text{ and } \xi_i^k \sim N(0, \Sigma_{d_i})$$

Correction: Represent $P(\mathbf{X}_i|\mathbf{y}_0, \mathbf{y}_i)$ by

$$\{(\mathbf{s}_i^{k,+}, w_i^{k,+})\}$$

where

$$\mathbf{s}_i^{k,+} = \mathbf{s}_i^{k,-} \text{ and } w_i^{k,+} = P(\mathbf{Y}_i = \mathbf{y}_i|\mathbf{X}_i = \mathbf{s}_i^{k,-})w_i^{k,-}$$

Resampling: Normalise the weights so that $\sum_i w_i^{k,+} = 1$ and compute the variance of the normalised weights. If this variance exceeds some threshold, then construct a new set of samples by drawing, with replacement, N samples from the old set, using the weights as the probability that a sample will be drawn. The weight of each sample is now $1/N$.

Algorithm 3: A practical particle filter resamples the posterior.

are given higher weights. This simple view provides some insight into the strengths of the method, and into what could go wrong.

Appendix II: Regression

The Nearest Neighbours

One could obtain the k -nearest neighbours, where for a new set of measurements \mathbf{x} , we obtain the k examples \mathbf{x}_i that lie closest to \mathbf{x} (using an appropriate distance function, which we discuss below), then average the \mathbf{y}_i associated with the examples. There is no reason to expect good results from this method unless the examples are relatively uniformly spaced — in particular, if the k nearest neighbours that we obtain are far away from our query example, the estimate could be very poor. Worse, if \mathbf{y} is a multi-valued function of \mathbf{x} , the method might average values associated with distinct “sheets”, and produce a meaningless result.

Assume, for the moment, that \mathbf{y} is not a multi-valued function of \mathbf{x} . We must still obtain the k nearest neighbours for a query \mathbf{x} . We expect to require sufficient examples that blank search through the examples is unattractive. An approximation called (r, ϵ) **nearest-neighbours** is attractive, because one can obtain a solution in time less than linear in the number of examples using an approach known as **locality sensitive hashing**. Assume a query point \mathbf{x} . If there is an example \mathbf{x}_j such that $d(\mathbf{x}, \mathbf{x}_j) < r$, then with high probability the algorithm returns \mathbf{x}_k such that $d(\mathbf{x}, \mathbf{x}_k) < (1 + \epsilon)r$, otherwise it reports no point. It does so by computing a hash key from a randomly selected subset of a **locality sensitive** family of functions; this key has the property that similar points have a high probability of colliding, and dissimilar points have a low probability of colliding. One then obtains a list of possible nearest neighbours for a query by computing a hash key for the query, obtaining all examples with which the query collides, then searching this set of examples (details appear in [150, 95, 196, 195]).

The usual difficulty in practical applications of nearest neighbours problems is that it isn’t clear what distance to use. Shakhnarovich *et al.* argue that in regression problems, it is important that the distance reflect the distance between predicted parameters (rather than, say, some canonical distance between examples) [389]. While their argument applies to the general case, the application is to lifting. In this case, two quite similar image configurations \mathbf{x}_i and \mathbf{x}_j may predict rather different 3D configurations \mathbf{y}_i and \mathbf{y}_j — a query close to \mathbf{x}_i should not be seen as close to \mathbf{x}_j , because the two predictions are different. This means that we would like the hashing functions we choose to

tend to generate collisions for examples that have close values of \mathbf{y} .

Shakhnarovich *et al.* achieve this by noting that a hash function acts like a classifier that determines whether two points are close (when they collide) or not. By searching thresholds T and **decision stumps** — functions of the image configuration \mathbf{x} with the property

$$\phi(\mathbf{x}) = \left\{ \begin{array}{ll} 1 & \text{if } \phi(\mathbf{x}) \geq T \\ -1 & \text{otherwise} \end{array} \right\}$$

they obtain a set of hash functions, on the set of training examples, tend to cause examples with similar \mathbf{y} to collide and those with distinct \mathbf{y} not to collide.

Some Regression Technique

Assume we have a set of examples likely to be relevant for a regression problem, perhaps as a result of k -nearest neighbours. How are we to obtain an estimate? in general, we must fit some approximation and obtain $\mathbf{y} = \mathbf{f}(\mathbf{x}, \beta)$, where β is a vector of parameters chosen by minimizing an error predicted and observed values. For our purposes, \mathbf{y} is a high dimensional vector. This means that we will not, in general, be able to build a good model of correlations between elements of \mathbf{y} , and so will predict its components independent of one another.

We will concentrate on the l 'th component. Our problem is now to predict a single value from some input vector. Because the components are assumed independent, we can ignore the question of which component we are dealing with, and so write (y_i, \mathbf{x}_i) for the i 'th example. We write

$$\mathbf{Y} = (y_1, \dots, y_N)^T$$

(i.e. the values to be predicted for each example) \mathcal{X} for the matrix

$$\begin{bmatrix} \mathbf{x}_1^T \\ \dots \\ \mathbf{x}_N^T \end{bmatrix}$$

Recall that we assumed that one component of \mathbf{x} has the value 1. We can then write the predictions produced by a vector β as $\mathcal{X}\beta$. The simplest method of obtaining a prediction is to solve for β that minimizes

$$(\mathbf{Y} - \mathcal{X}\beta)^T(\mathbf{Y} - \mathcal{X}\beta)$$

(this is **linear regression**). This is a maximum likelihood solution assuming that $p(y|\mathbf{x}, \beta, \sigma^2)$ is normal with mean zero and variance σ^2 . Linear regression tends to overfit. One problem is that, if terms in \mathbf{x} are strongly correlated with one another, an appropriate linear combination of these terms is noise; and multiplying this noise by a large number may lead to a good fit. This suggests solving for β that minimizes

$$(\mathbf{Y} - \mathcal{X}\beta)^T(\mathbf{Y} - \mathcal{X}\beta) + \lambda\beta^T\beta$$

(a practice known as **ridge regression**). Here λ is usually chosen by cross-validation [169]. An alternative method to control the size of the coefficients is to solve for β that minimizes

$$(\mathbf{Y} - \mathcal{X}\beta)^T(\mathbf{Y} - \mathcal{X}\beta) + \lambda \sum_i |\beta_i|$$

(a practice known as the **lasso**). The effect of using an L_1 penalty is first to make solutions non-linear in the example values, and — for sufficiently large λ — compelling some components of β to be zero.

A **relevance vector machine** (RVM) uses a Bayesian model to force weights in the regression to be small or zero. One has a model of $P(y|\mathbf{x}, \beta, \sigma^2)$ (normal, mean zero and variance σ^2 , as above). We express a preference for smaller weights by writing a prior probability distribution on β , conditioned on hyperparameters α . This distribution, $p(\beta|\alpha)$, is normal with mean zero and diagonal covariance; the variance of each component is $1/\alpha$. The posterior over weights, $p(\beta|y, \mathbf{x}, \sigma^2, \alpha)$, is obtained from Bayes' rule. We then obtain $p(y|\sigma^2, \alpha)$ by marginalization, and choose a value of α that maximizes this likelihood (say $\hat{\alpha}$). We then choose β to maximize $p(\beta|y, \mathbf{x}, \sigma^2, \hat{\alpha})$. Algorithms for each step appear in [123, 437, 438] (with a variant in [3]); an analysis in [123] shows that some components of α will tend to be infinite, implying a zero component in β .

In each of these cases, one is building a linear model that applies over the whole of the set of examples. This is unattractive for our application (for example, there is fair prospect that there is more than one \mathbf{y} for a given \mathbf{x} , something that doesn't happen in linear models). It is natural to try to fit **local linear models**, where the model is linear, but the β varies with the input. We can do this by weighting errors with weights that depend on how close the query is to the examples.

For some input \mathbf{x} , we compute a diagonal weight matrix $\mathcal{W}(\mathbf{x})$. There are a variety of possibilities for the diagonal, but using

$$\exp \left[\left(\frac{1}{2\sigma^2} \right) (\mathbf{x} - \mathbf{x}_j)^T (\mathbf{x} - \mathbf{x}_j) \right]$$

for the j 'th entry is a sensible choice (errors close to the example have large weights, those far have small weights; σ may be chosen by cross-validation). We now use β obtained by minimizing

$$(\mathbf{Y} - \mathcal{X}\beta)^T \mathcal{W}(\mathbf{Y} - \mathcal{X}\beta) + \lambda \beta^T \beta$$

for that example. At the cost of solving a linear system for every query, we have a fit that smoothes local linear models into one another without needing to solve non-linear equations. The method has a good reputation for high-dimensional \mathbf{x} . Notice that the weight penalties described above can be applied to this problem as well; one computes a ridge regression, lasso, or relevance vector machine for each query as appropriate.

We are assuming that the \mathbf{x} we have computed form an appropriate set of features. However, a more extensive feature representation can be obtained by comparing the query against examples using a **kernel**. Assume that we have a satisfactory **kernel function** $K(\mathbf{u}, \mathbf{v})$. A popular kernel function is

$$K(\mathbf{u}, \mathbf{v}) = \exp \left(\left(\frac{1}{2\sigma^2} \right) (\|\mathbf{u} - \mathbf{v}\|) \right)$$

(for a choice of σ usually obtained by cross-validation). We can then consider a regression of the form

$$y(\mathbf{x}) = \beta_0 + \sum_{i \in \text{examples}} \beta_i K(\mathbf{x}, \mathbf{x}_i)$$

Note that this remains linear in β , so that the methods expounded above apply. Various details and other applications of this useful trick appear in [411].

4

Tracking: Data Association for Human Tracking

Tracking people is a means to an end, and trackers should be assessed in that way. Human trackers should be reasonably accurate, start automatically (hardly any practical application can use trackers that can't be started automatically), run for long times without any particular difficulties, and not rely excessively on implausible assumptions about background, etc. These are the correct criteria by which to judge.

In our opinion, the literature has, until quite recently, placed too much emphasis on probabilistic inference machinery, while paying insufficient attention to the (possibly dull but certainly essential) vision problems implied by data association. Furthermore, this inference machinery may, in fact, be being used to solve a non-problem (section 7.1.5).

Early human trackers, which used quite straightforward matching methods, (for example, Hogg's 1983 paper [181]; Rohr's 1994 tracker [365]) could produce kinematic tracks for people moving without sudden accelerations on reasonably simple, high-contrast backgrounds if started manually. The advantages of a known, simple background have been thoroughly explored (section 2.2.1). The more recent trackers we have described use more complex inference machinery, but without any great change in competence.

Improvements in competence seem to have come with increased attention

paid to tracking by detection schemes. These are well established in, say, face tracking. For example, one can build a fairly satisfactory face tracker by simply running a face detector on frames, and linking over time; smart linking schemes built around affine invariant feature patches can result in very satisfactory tracks [401]. Tracking by detection is now capable of building good human kinematic tracks, without relying on background subtraction.

4.1 Detecting Humans

Human detection is difficult, and important. It is difficult because people (usually!) wear clothing of widely varying appearance; because changes in body configuration can result in dramatic changes in appearance; and because different views result in dramatic changes in appearance. There are several important applications. A huge literature now deals with methods to detect pedestrians automatically, because this is a function that autonomous or semi-autonomous motor-cars will need. There is a substantial literature on detecting and interpreting gestures for human-computer interaction purposes. There is a smaller but growing literature on using various human detection and description methods for understanding the content of various multi-media datasets. There is a small but occasionally startling literature on methods for detecting sexually explicit images. Interest in these areas is not confined to academia; in each of these areas, there are both research efforts by established companies and start-up companies appearing regularly.

No published method can find clothed people wearing unknown clothing in arbitrary configurations in complex images reliably, though, as we shall see, there is reason to believe that this situation will change. The first standard approach to this problem involves matching to one or a family of templates, which might use either spatial or temporal information (or both). We review this area in section 4.1.1 and section 4.1.2. The second standard approach is to identify parts of a person and then reasoning about an assembly of these parts to identify the person. We distinguish between two types of method, according to the type of part: First, one may use parts that are semantic in origin (“arms”, “legs”, “faces”, and so on), and we review this approach in section 4.1.3. Second, one may use parts that are defined by statistical criteria (for example, they form a good codebook for representing the image of the person), and we review this approach in section 4.1.4.

4.1.1 Finding People by Matching Static Templates

Approximately half-a-million pedestrians are killed by cars each year (1997 figures, in [144]). Car manufacturers and governments have an interest in ensuring that cars are less dangerous, and there is a considerable body of research on automated pedestrian detection. Gavrila gives an overview of the subject in [144], which covers cues such as radar, infrared, and so on, which have practical importance but are of no interest to us. For our purposes, this is an example of person detection that may be simpler than the general problem, and is certainly important.

At relatively low resolution, pedestrians tend to have a characteristic appearance. Generally, one must cope with lateral or frontal views of a walk. In these cases, one will see either a “lollipop” shape — the torso is wider than the legs, which are together in the stance phase of the walk — or a “scissor” shape — where the legs are swinging in the walk. This encourages the use of template matching.

Papageorgiou and Poggio represent 128x64 image windows with a modified wavelet expansion, and present the expansion to a **support vector machine** (SVM), which determines whether a pedestrian is present [330]. SVM's are classifiers, trained with positive and negative examples. For a brief informative discussion of SVM's see [448] or [85]. More extensive information appears in [382, 390, 447], and discussion in the context of a variety of other classifiers is in [169]. The training data consists of windows with and without people in them; each positive example is scaled such that the person spans approximately 80 pixels from shoulder to foot. A variety of image representations are tested, with the modified wavelet expansion applied to colour images performing significantly better than wavelet expansions applied to grey-level images, low resolution pixel values for grey-level images, principal components analysis representations of grey-level images, and the like. The strength of these wavelet features appears to be that they emphasize points that are, rather roughly, outline points. This yields a method for exploiting the restricted range of contours without explicitly encoding contour templates. The wavelet expansion can be reduced in dimension to obtain a faster, though somewhat less accurate, matcher. There are several variants of this approach in the literature [317, 318, 325, 326, 327, 329].

Zhao and Thorpe use stereopsis to segment the image into blocks, then

present each block to a neural network [484]. The stereo cue acts as a variant of background subtraction, because there are typically substantial discontinuities in depth between pedestrian and background. A comparison of this system with that of Papageorgiou *et al.* (the version in [326]) suggests it is more accurate, possibly because the stereo segmentation reduces the number of windows that must be searched.

There are a variety of systems that use edge templates explicitly. Gavrilu describes an approach that matches image contours against a hierarchy of contour templates using a chamfer distance [142]. The method is oriented to real-time detection. The image is passed through an edge detector, and then passed through a smoothed distance transform (see [33]); a template is evaluated by computing the sum of distance transform values at template feature points, so that a small value results in a match. One needs numerous templates for such a method to be successful (distinct views; distinct phases in the walk), and Gavrilu organizes the set of templates into a hierarchy using an agglomerative clustering method rather like k-means. Each node of the hierarchy contains a summary template (summaries at nodes deeper in the hierarchy encode more spatial detail), and a representation of the distance of the examples from that summary. Matching proceeds by computing a cost to the representative node at the current level, and testing this against a threshold to determine whether to expand that node or not. A verification step uses radial basis functions to classify those image windows that appear to match edge templates. Gavrilu *et al.* describe an improved version of this method, using stereo cues and temporal integration [145]. Broggi *et al.* describe a method that uses vertical edges, the characteristic appearance of the head and shoulders, and background subtraction to identify pedestrians [62].

Wu *et al.* build random field models of image windows with and without a pedestrian, and then detect using a likelihood ratio [473]. Shape is encoded with a random field, and measurements are assumed to be conditionally independent given the shape and some deformation parameters. There is a search over scale, translation and orientation. The considerable technical difficulties involved in evaluating the likelihood are dealt with using a variational approximation. One would expect a performance penalty for using a generative formalism in what is, in essence, a discriminative problem (does this window contain a pedestrian or not?), but ROC curves suggest the method is comparable with strong recent discriminative methods in performance.

Dalal and Triggs give a comprehensive study of features and their effects on performance for the pedestrian detection problem [94]. The method that performs best involves a histogram of oriented gradient responses (a **HOG** descriptor). This is a variant of Lowe’s SIFT feature [267]. Each window is decomposed into blocks (large spatial domains) and cells (smaller spatial domains). A histogram of gradient directions (or edge orientations) is computed for each cell. In each block, a measure of histogram “energy” is computed, and used to normalize the histogram for each cell in the block. This supplies a modicum of illumination invariance. The detection window is tiled with an overlapping grid, within each cell of which HOG descriptors are computed, and the resulting feature vector is presented to an SVM. Dalal and Triggs show this method produces no errors on the 709 image MIT dataset of [330]; they describe an expanded dataset of 1805 images. The paper compares HOG descriptors with the original method of Papageorgiou and Poggio [330]; with an extended version of the Haar wavelets of Mohan *et al.* [290]; with the PCA-Sift of Ke and Sukthankar ([220]; see also [284]); and with the shape contexts of Belongie *et al.* [36]. There is considerable detailed information on tuning of features.

4.1.2 Templates that include Motion

Static templates most likely work because the outlines of pedestrians tend to be of limited complexity. While it would be nice to have a formal notion of what this meant, the appropriate comparison is with arbitrary views of people in arbitrary configurations (say, the figure skater of figure 4.10). Pedestrians also tend to move in quite restricted ways — they are typically either standing or walking. Niyogi and Adelson point out that, if one forms an **XYT image** — a stack of frames, registered as to camera motion, originally due to Baker [26] — these motions produce quite distinctive structures (figure 4.1), which can be used to identify motions [313] or recover some gait parameters [312]. Polana and Nelson consider spatial patterns of motion energy, which also have a characteristic structure [346]. There is a substantial literature on the characteristic appearance of human motion fields; a good start is [53, 340, 341, 342, 344, 345, 347, 254, 255, 256, 255]. Particular efforts have been directed to periodic motion; one might consult [77, 88, 89, 90, 91, 160, 159, 252, 260, 261, 383, 384, 252, 435].




Fig. 4.1 **figure 2 of polana nelson recognizing activities** , **figure 2 of Niyogi Adelson recognizing gait**
 On the **left**, an XYT image of a human walker. The axes are as shown; the stack has been sliced at values of Y , to show the pattern that appears in the cross section. Notice that, at the torso there is a straight line (whose slope gives an estimate of velocity) and at the lower legs there is a characteristic “braid” pattern, first pointed out by Niyogi and Adelson [313]. On the **right**, a series of estimates of the spatial distribution of motion energy (larger white blocks are more energy) for different frames of a walk (**top**) and a run (**bottom**); the frame is rectified to the human figure by translation, and one image frame from each sequence is shown. Notice that, as Polana and Nelson point out, this spatial distribution is quite characteristic [346]. Figure from “Recognizing Activities”, Polana and Nelson, Proc. Int. Conf. Pattern Recognition, 1994, © 1994 IEEE Figure from Analyzing Gait with Spatiotemporal Surfaces, Niyogi and Adelson, Proc. IEEE Workshop on Nonrigid and Articulated Motion, 1994, © 1994 IEEE

This characteristic structure can be used to detect pedestrians in a variety of ways. Papageorgiou and Poggio compute spatial wavelet features for the frame of interest and the four previous frames, stack these into a feature vector, and present this feature vector to an SVM, as above [328]. The result is a fairly significant improvement in detection rate for given false positive rate. The performance improvements that Dalal and Triggs obtain by careful feature engineering (as above) are probably available here, too. The features encode motion implicitly (by presenting the frames in sequence), but not explicitly.

Viola *et al.* use explicit motion features — obtained by computing spatial averages of differences between a frame and a previous frame, possibly shifted spatially — and obtain dramatic improvements in detection rates over static features ([453, 454]; see also the explicit use of spatial features in [87, 321, 322], which prunes detect hypotheses by looking for walking

cues). This work uses a **cascade architecture**, where detection is by a sequence of classifiers, each of which operates only on windows accepted by the previous classifier. The classifiers are engineered so that they each have a low false negative rate, so that classifiers early in the cascade reject many windows, and so that the overall cascade is accurate. Features are sums of spatial averages over box-shaped windows in space and time, and so can be evaluated in large numbers extremely quickly; the techniques of classifier and features are due to Viola and Jones [450, 451, 452].

Dimitrijevic *et al.* build a spatio-temporal template as a list of spatial templates in time-order [108]. The spatial templates are edge templates giving the silhouette of the figure, and are matched with a chamfer distance, as above. The spatial templates and the spatio-temporal templates (which are acceptable sequences of spatial templates) are obtained by rendering skinned motion capture data against a blue background from a wide variety of views. The match is scored by computing the time average of chamfer distances. The detector is trained to detect the portion of the walk cycle where both feet are on the ground (other frames could be handled by various forms of temporal interpolation or tracking; see also section 4.2.3.2). The paper describes a variety of optimizations helpful to obtain a reasonable speed.

4.1.3 Traditional Parts

Detecting pedestrians with templates most likely works because pedestrians appear in a relatively limited range of configurations and views. It appears certain that using the architecture of constructing features for whole image windows and then throwing the result into a classifiers could be used to build a person-finder for arbitrary configurations and arbitrary views only with a major engineering effort. The set of examples required would be spectacularly large, for example. This is unattractive, because this set of examples implicitly encodes a set of facts that are relatively easy to make explicit. In particular, people are made of body segments which individually have a quite simple structure, and these segments are connected into a kinematic structure which is quite well understood.

All this suggests finding people by finding the parts and then reasoning about their layout — essentially, building templates with complex internal kinematics. The core idea is very old (for example, one might con-

sult [9, 10, 42, 174, 275, 319]) but the details are hard to get right and important novel formulations are a regular feature of the current research literature. It is currently usual to approach this question in terms of 2D representations, which represent a view of a person as a set of body segments — which could be represented as image rectangles — linked by rotary (and perhaps translational) joints.

The advantage of these 2D kinematic templates is that they are relatively easy to learn. Learning 2D kinematic templates requires the relative scale of body segments, link probabilities, and an appearance encoding for each body segment. It is relatively straightforward to obtain scale information from static images. Link probabilities can be modelled in a variety of ways. It is usually better to represent translation as well as rotation of a link with respect to another; if we now use a distribution that is flat, or near to, within a useful range, we are preferring no legal kinematic configuration over any other. This isn't in accord with reality — most of the time in most footage, people are walking — but is convenient because it doesn't lock us into any particular activity. In this form, link probabilities can be modelled using either static images or anthropometric information.

4.1.3.1 Discriminative Approaches

The first difficulty is that simply identifying the body parts can be hard. This is simplified if people are not wearing clothing, because skin has a quite distinctive appearance in images. Forsyth *et al.* then search for naked people by finding extended skin regions, and testing them to tell whether they are consistent with body kinematics [132, 133]. The method is effective on their dataset (and can be extended to find horses [131]), but is not competitive with more recent methods for finding “adult” images (which typically use whole-image features [16, 51, 209, 479]). Ioffe and Forsyth formalize this process of testing, and apply it to relatively simple images of clothed people [197, 200]. Their procedure builds a classifier that accepts or rejects whole assemblies of body components; this is then projected onto factors to obtain derived classifiers that can reject partial assemblies that could never result in acceptable complete assemblies. Sprague and Luo use this approach to find clothed people in more complex images, by reasoning about image segments [414].

Mohan *et al.* use a discriminative approach not only to identify good as-



Fig. 4.2 from Mohan et al, figures 5 and 10 Mohan et al. use SVM's to find major body parts (left arm, right arm, head/shoulders and legs) as in the training examples shown on the **top**. They then use these SVM's to search frames for components; the response of all part SVM's in each window is pooled and then presented to an SVM which identifies whole figures. On the **bottom**, examples showing good detects; the whole body window is outlined with lines, and the part windows with dashed lines.

semblies of parts (as above), but also to find body parts [290]. SVM's are trained to detect the whole left arm, the whole right arm, the legs and the head/shoulders (see figure 4.2); because these body components are relatively large, and because the work focuses on pedestrians, it is possible to search for them in an image centered frame — one can inspect vertical boxes of the right size and aspect ratio to tell whether an arm is present. The SVM part detectors produce a score (distance to the separating hyperplane). For each 128x64 window, the top score for each type of part is placed in a slot in a vector, which is presented to a further SVM. Geometric consistency is enforced by finding the top score for each type of part over a subset of the window to be classified. The approach is applied to pedestrian images, and outperforms the methods of [317, 329].

4.1.3.2 Generative Approaches

Naked people are easier, because identifying body parts is easier. If we had an encoding of the appearance of the individual parts, this would simplify finding people, because identifying an instance involves dynamic programming; but, done in a straightforward fashion, this is slow because the likelihood evaluation is slow. Felzenswalb and Huttenlocher show how one may use **distance transforms** to speed this process up substantially [124, 125]. In particular, assume that the model is built out of a set of components, the i 'th of which has some configuration l . We assume that the components are




Fig. 4.3 **DT's, F+H CVPR 00, fig 4 part** A pictorial structure is a 2D model of appearance as a kinematic tree of segments. Each segment has configuration variables which encode the spatial support of the segment — for example, position and orientation — a local appearance model — for example, the color of a segment — and there is a cost associated with each edge in the tree — for example, the cost of finding a lower leg far from an upper leg. One can find the best instance of such a structure by discretizing the configuration variables for each segment, then using dynamic programming. Felzenswalb and Huttenlocher show that, for properly defined segment-segment costs, the cost-to-go function in the dynamic programming can be evaluated more cheaply than one would expect, meaning that localization can be fast [124, 125].

linked in a tree of n nodes. Then to find the best instance, we can discretize the configurations — assume that we use m sample points — and do dynamic programming. However, this will cost $O(nm^2)$, which is unattractive because m is likely to be quite big, particularly if the configurations are high-dimensional. Felzenswalb and Huttenlocher show that, as long as the link cost has a particular form, the cost-to-go functions encountered in the dynamic programming problem are, in fact, generalized distance transforms, and so can be computed in $O(m)$ time (so that the whole thing costs $O(nm)$, which is a useful improvement). The paper demonstrates these models being used in two contexts: finding people and finding cars. People are modelled with rectangles of fixed size and known color (appearance is modelled with image color) and can be localized quite effectively (figure 4.3). Kumar *et al.* extend this model to incorporate boundaries into the likelihood and use loopy belief propagation to apply it to arbitrary graphs (rather than trees); the method is applied to pictures of cows and horses [240].

4.1.3.3 Mixed Approaches

Ronfard *et al.* use a discriminative model to identify body parts, and then a form of generative model to construct and evaluate assemblies [366]. Their approach searches for parts that are on a finer scale than those of Mohan *et al.* (upper arms vs. arms), and these can't be found by looking for boxes of a fixed size, orientation and aspect ratio. This makes it a good idea to search for body parts over scales and orientations — in effect, a search in a part-centered coordinate system. They compare an SVM part detector and an RVM (section 3) part detector, both applied to features that consist of filtered image grey levels within the window; authors suggest that more sophisticated features, for example those of Dalal and Triggs (section 4.1.1), might give improvements. Each of the detectors produces a detection score. People are modelled as a 2D kinematic chain of parts, with link scores depending on a weighted sum of position, angle and detector scores. The chain is detected with dynamic programming, but the savings obtained by Felzenswalb and Huttenlocher (section 4.1.3.2) do not appear to be available. The weights used in the sum are obtained by a novel application of SVM's. The authors collect a large number of positive and negative examples of links, use a linear SVM with link terms as features to classify them, then use the weights produced by that linear SVM as weights in the link cost. Detection performance is strong; however, there are no standard datasets for evaluating detection of people in arbitrary configurations so comparisons are difficult.

Mikolajczyk *et al.* use discriminative part detectors, applied to orientation images and built using methods similar to those of Viola and Jones (see section 4.1.2), to identify faces, head-and-shoulders, and legs [286]. Non-maximum suppression isolates detected parts. Once a part is found, it predicts possible locations for other parts, which are used to drive a search. Finally, the assemblies that are found are presented to a likelihood ratio classifier. Micilotta *et al.* use discriminative methods to detect hands, face and legs; a randomized search through assemblies is used to identify one with a high likelihood, which is tested against a threshold [282]. Similarly, Roberts *et al.* use a randomized search to assemble parts; parts are scored with a generative model, which is used to obtain a proposal distribution for joints [363].

4.1.4 Parts as Codebooks

Representing a body by segments may not, in fact, be natural; our goal is effective encoding for recognition, rather than disarticulation. One might represent people by image patches chosen to be good at representing people. Leibe *et al.* have built the best known pedestrian detection system using this approach [249]. They first obtain multiple frames of pedestrians, segmented from the background using a form of background subtraction (section 2.2.1), to serve as training data. They build similarity covariant neighbourhoods at interest points (using the methods of [267]), and rectify windows to a fixed scale. These rectified windows are clustered, the cluster centers yielding a codebook. They now build a representation of the probability of encountering a codebook entry at a particular location in the object frame by counting matches to codebook entries for each example.

Write λ for position and scale of the object, o_n for the class of object (we may be interested in detecting pedestrians and dogs, for example; the background is one such class), c_i for the i 'th codebook entry and \mathbf{l} for the location and scale of the codebook entry. We can build a model of $P(o_n, \lambda | c_i, \mathbf{l})$ — the probability that an object of class o_n occurs at location and scale λ conditioned on a codebook entry of type i observed at \mathbf{l} — by counting. Write \mathbf{e} for an image patch.

We obtain a model of $P(o_n, \lambda | \mathbf{e}, \mathbf{l})$ because we know $P(\mathbf{e} | c_i)$ and can marginalize. Local maxima of this model may be instances of the object; they are obtained with the mean-shift algorithm [84].

Given an hypothesis, we can now determine a probability map giving whether each pixel lies on that hypothesized object or not. Write \mathbf{p} for the location of a pixel. For $P(\mathbf{p} = \text{figure} | o_n, \lambda)$, we obtain

$$\sum_{(\mathbf{e}, \mathbf{l}) \ni \mathbf{p}} \sum_i P(\mathbf{p} = \text{figure} | o_n, \lambda, c_i, \mathbf{l}) \frac{p(o_n, \lambda | c_i, \mathbf{l}) p(c_i | \mathbf{e}) p(\mathbf{e}, \mathbf{l})}{p(o_n, \lambda)}$$

where the \ni sign refers to windows and locations that cover the pixel. The relevant densities can be obtained by counting. All this means that, associated with each plausible detect, we have a map of pixels that might lie on a pedestrian. In turn, we can search for collections of pedestrian hypotheses that explain these pixel maps best by evaluating the description length. The search works by evaluating the change in description length obtained by

changing hypotheses (the particular greedy search used is a variant of that in [248, 250]). The hypotheses are refined with a form of chamfer matching applied within the area segmented as belonging to a pedestrian, and a further description length search applied only to silhouettes yields a final count of pedestrians.




Fig. 4.4 leibeScheile, persondetection/parts/leibe-crowded-scenes The detection and verification process of Leibe and Scheile [249] begins by using image patches to obtain a posterior on pedestrian position and scale. In turn, this leads to a putative segmentation (see **a**, which shows the support map for an image from this stage, with hypotheses leading to the support map shown as green boxes superimposed on the image). However, because the consistency model is local, these putative segmentations could, for example, have extra limbs (see the extra legs in **b**). Obtaining an accurate count and segmentation requires the use of global data, supplied by contours and chamfer matching. However, as we see in **c**, some false positives lie on top of regions with multiple edges, which could defeat contour matching; if one matches only to the pixels covered by the support map (**d** and **e**), this effect is less pronounced, and only the correct hypotheses are confirmed (**f**). The **bottom** row shows a series of results. The red box in the **center right** image is a true false positive; in the **right**, the red box is a detect of a pedestrian who does not appear in the annotation of the image (because marking up example images accurately is very difficult).

4.2 Tracking by Matching Revisited

Methods for tracking humans by detection follow, in rough outline, methods for detecting humans. One may use either whole person templates, or collections of parts, which might be traditional or form a codebook. However, there are some important variations in the question of what appearance model (gen-

erative vs. discriminative; inferred or provided) one uses and how one scores the comparison of the model with the image.

4.2.1 Likelihood

Most probabilistic tracking algorithms must compute the likelihood of some image patch conditioned on the presence of a model at some point. The easy model to adopt is to produce a template for the patch from the model parameters, subtract that template from the image, and assume that the result consists of independent noise — that is, that the value at each pixel is independent. Whether it is wise to use this model or not depends on how the template is produced — for example, a template that does not encode illumination effects is going to result in a residual whose pixel values are not independent from one another (see Sullivan *et al.* for this example [426]), and so the likelihood model is going to significantly misestimate the image likelihood.

The problem occurs in a variety of forms. For example, if one represents an image patch with a series of filter outputs (after, say, [393, 394]), each element is unlikely to be independent and errors are unlikely to be independent. Sullivan *et al.* describe the problem, and demonstrate a set of actions (including building an illumination model and estimating correlation between filter outputs) that tend to ameliorate it, in the context of face finding [426]. Roth *et al.* build likelihood models for vectors of filter outputs using a **Gibbs model** (known in other circles as a **maximum entropy model** or a **conditional exponential model**; see section 3.3.1.1) [371]. Their method is trained using an algorithm due to Liu *et al.* ([263]; see also [257], and one might compare variants of iterative scaling [40, 96, 210, 335, 370]). There is some evidence that the likelihood produced using this model is more tightly tuned to — in their example — the presence and location of a leg. The model is used by Sigal *et al.* ([396]; section 3.3.1.1) to obtain tracks of people in 3D from three views.

While it is clear that there is an issue here, it is a bit uncertain how significant it is. I am not aware of clear evidence that better tracking or localization results from being careful about this point, and am inclined to believe that the rough-and-ready nature of current likelihood models is not a major problem.

4.2.2 Whole Object Templates

Toyama and Blake encode image likelihoods using a mixture built out of templates, which they call **exemplars**. Assume we have a single template — which could be a curve, or an edge map, or some such. These templates may be subject to the action of some (perhaps local) group, for example translations, rotations, scale or deformations. We model the likelihood of an image patch given a template and its deformation with an exponential distribution on distance between the image patch and the deformed template (one could regard this as a simplified maximum entropy model; we are not aware of successful attempts to add complexity at this point). The normalizing constant is estimated with Laplace’s method. Multiple templates can be used to encode the important possible appearances of the foreground object. State is now (a) the template and (b) the deformation parameters, and the likelihood can be evaluated conditioned on state as above.

We can think of this method as a collection of template matchers linked over time with a dynamical model. The templates, and the dynamical model, are learned from training sequences. Because we are modelling the foreground, the training sequences can be chosen so that their background is simple, so that responses from (say) edge, curve, and the like detectors all originate on the moving person. Choosing templates now becomes a matter of clustering. Once templates have been chosen, a dynamical model is estimated by counting; authors do not discuss this point, but it seems likely that some form of smoothing would be useful, because if one has many templates and relatively short training sequences, observing that one template never follows another does not establish the probability of the event is zero. Smoothing techniques for problems of this form are a popular tool in the statistical natural language community, and several appear in Manning and Schütze’s book [274].

What makes the resulting method attractive is that it relies on *foreground enhancement* — the template groups together image components that, taken together, imply a person is present. The main difficulty with the method is that many templates may be needed to cover all views of a moving person. Furthermore, inferring state may be quite difficult. Authors use a particle filter; but if one views a particle filter as a type of randomized search started using dynamics, as above, then it is clear that this search will be more difficult as the




Fig. 4.5 **Figure 3, 6 and 7 from Toyama-Blake IJCV paper** Toyama and Blake [443, 442] track a 2D model of a person by learning a set of templates — which they call exemplars — from other sequences of moving people. The image consists of a deformed template and noise, and state is given by which template is rendered, and the deformation through which the template is rendered. The likelihood is obtained from a comparison between the template and the image. Tracking uses a particle filter. On the **top**, a typical set of templates, consisting of edge points (one may also use curves, region textures, and so on). On the **lower left**, a track displayed by rendering the template deformation pair with the largest posterior. On the **lower right**, a track of the same sequence obtained with some frames blank; notice that the dynamical model fills in reasonable templates, suggesting that such a tracker could be robust to brief occlusions.

movement is less predictable and as the number of templates increases. Part of the difficulty is that the likelihood may change quite sharply with relatively small changes in transformation parameters.

Spatial templates can be used to identify key points on the body. Sullivan and Carlsson encode a motion sequence (of a tennis player) using a small set of templates, chosen to represent many frames well [427]. These templates are then marked up with key points on the body, and matched to frames using a score of edge distance that yields pointwise correspondence; they show that a rough face and torso track, obtained using a particle filter, improves the correspondence. The key points are transferred to the markup, and the correspondence between edge points is used to deform the matched template to line up with the image; this deformation carries the keypoints along (figure 4.6). Finally, the configuration of the keypoints is significantly improved using a particle filter for backward smoothing. Loy *et al.* show that such transferred keypoints can be used to produce a three dimensional reconstruction of the configuration of the body [268].




Fig. 4.6 **Figure 3, part of 7, 8 part of 11 from Sullivan and Carlsson Devatracker/23500629.pdf** Sullivan and Carlsson encode a motion sequence (of a tennis player) using a small set of templates, chosen to represent many frames well [427]. Matching uses a topological criterion that can generate large numbers of matches (a); but by obtaining a rough face and torso track (b), they can localize matches (c). This makes it possible to transfer key body points marked on the templates by hand to the frames, and by deforming the templates to obtain good estimates of the image location of these key points (d). Finally, the configuration of the keypoints is significantly improved using a particle filter for backward smoothing.

4.2.3 Traditional Parts

We have already discussed tree-structured models of the body (section 4.1.3). There are two areas in which tracking humans by detection varies from human finding. The first is in how one models temporal and spatial relations, which can easily lead to intractable models. The second is in whether the appearance model is supplied or inferred.

4.2.3.1 Complex Spatio-Temporal Relations

The advantage of a tree-structured kinematic model, that one can use dynamic programming for detection, extends to a mixture of such trees. However, adding temporal dependencies produces a structure that does not allow for simple exact inference, because the state of a limb in frame t has two parents: the state in time $t - 1$, and the state of its parent in frame t (recall figure 3.8). Ioffe and Forsyth attack this problem with a form of coordinate ascent on $P(\mathbf{X}_0, \dots, \mathbf{X}_k | \mathbf{Y}_0, \dots, \mathbf{Y}_k)$ [198]. They use a mixture of trees as a template. Spatial links are learned from static images and temporal links simply apply a velocity bound. The posterior is maximised by an iterative procedure, which interleaves two steps maximising over space in a particular frame while fixing all others, and maximising over time for a particular limb segment, while fixing all other segments. Each step uses dynamic programming. Segments are assumed to be white, or close; the model doesn't

encode the head position, which occasionally leads to arms and legs getting confused. As figure 4.7 indicates, fair tracks are possible without a dynamical model. One should see the work of Sigal *et al.* (section 3.3.1.1; figure 3.9) as involving a similar, but more sophisticated, inference procedure.

4.2.3.2 Known Appearance Models

This difficulty is quite often ignored, apparently without major consequences. Mori and Malik use no dynamical model, detecting joints repeatedly in each frame using the method described in section 3.2.1; the result is a fair track of a fast-moving skater [296]. Lee and Nevatia use a Markov model of configuration (but not of appearance), where each body configuration depends only on the previous configuration [247]. The model uses the known appearance of skin to identify faces and hands, and contrast with the background to identify major limbs and torso. Markov chain Monte Carlo is used to give a randomized search for good matches between configuration and image, with proposals using both forward and backward dynamics.

Agarwal and Triggs build a set of dynamical models, each of which explains a cluster of motion data well; a mixture of these models is then used to propose the 2D configuration in the $i + 1$ 'th frame from the state in the i 'th frame [4]. The models are fit to a reduced dimensional representation. The question of how one knows which model to use is dealt with by mixing the models, mixture weights being set by the current frame. The entropy of these weights tends to be low, as many 2D configurations can arise from only one of their motions. The result is a model that can make quite accurate dynamical predictions for their example sequences. The predictions are refined by an optimization method, as in [407]. The model is a 2D kinematic tree, and likelihood is evaluated by warping the image backwards, using the current state estimate, and comparing that warped image to body part reference templates that are part of the initialization ([4], p.61). There is no information about what appearance model the templates encode; one could see this method as an extension of the people-finding approach of Felzenswalb and Huttenlocher [125] that finds local minima suggested by the dynamical model.




Fig. 4.7 MOT, part of figure 3 Ioffe and Forsyth build a 2D model of a person as a set of segments, modelled using a mixture of trees to capture aspect phenomena [198]. In an image sequence, each segment except the root has two parents — the corresponding segment in the previous frame, and that segment’s parent in the model. The appearance model of each individual segment is crude — segments are light bars of fixed scale. Authors find the best sequence of models by interleaving optimization over time with optimization over space; the result is a fair track, despite significant changes in aspect.

4.2.3.3 Inferred Generative Appearance Models

This leaves us with building a model of appearance. We must choose an encoding of appearance, and determine what appearance each segment has. The trackers we have described up to this point train models of appearance using one or another form of training data; but one could try to build these models on the sequence being tracked. The advantage of doing so is that these appearance models can be specialized to the individual being tracked — rather than attempt to encode human appearance generally, which appears to be difficult. This is the only place where, for example, we can clearly tell what color clothing is being worn by the subject.

Ramanan and Forsyth encode appearance using color — the texture changes produced by shading on folds in clothing make texture descriptors unhelpful — and determine appearance for each segment by clustering. Their algorithm assumes known scale and known link probabilities. Since individuals don’t change clothing in track sequences, one can expect that body segments look the same over the sequence, and so there should be many instances of the true segments in a long sequence. Furthermore, the correct segments lie in distinctive configurations with respect to one another in each frame, if detected. This constraint is more easily exploited by looking for torso segments first, because they’re larger and tend to move more slowly. Ramanan and Forsyth use a filter tuned to parallel edges separated by a particular image distance to identify candidate torso segments; they then cluster these and

prune clusters that are stationary. They look for arm and leg segments near each instance of a candidate torso segment, and if enough are found, declare that the candidate represents a true torso in appearance. Now the appearance of each arm and leg segment can be determined by finding segments near the torso that lie in the correct configuration and have coherent appearance (this is simplified by the useful observation that left and right arms and left and right legs typically look the same). Tracking now becomes a straightforward matter of detecting instances of each model in each frame, and linking those that meet a velocity constraint.

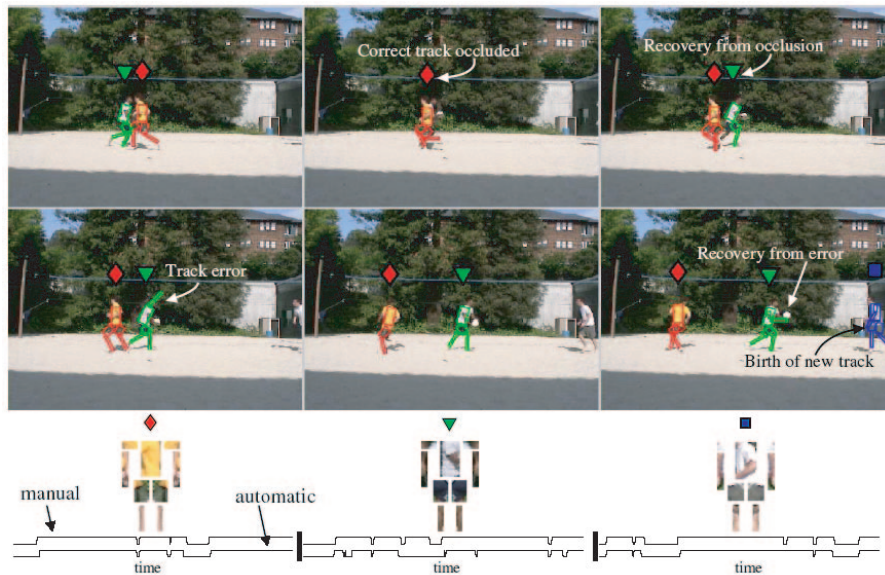


Fig. 4.8 **Fig 7, CVPR 03, RamananForsyth** Ramanan and Forsyth build an appearance model for segments in a 2D model of a person automatically, using methods described in the text. They then track by detecting instances of this appearance model in frames and linking instances across time. The advantages of this tracking by detection strategy are that one can identify particular individuals, recover from occlusions, from errors in the track and from individuals leaving the frame. The **top** shows frames from a tracked sequence; on the **bottom**, appearance models for each of the three individuals identified by their appearance modelling strategy.

This displays some advantages of a tracking by detection framework, and the difficulties that result from relying on a dynamical model. First, recovery from occlusion, people leaving frame or dropped frames is straightforward; because we know what each individual looks like, we can detect the indi-

vidual when they reappear and link the tracks (this point is widely acknowledged; see, for example, [95, 301]). Second, track errors don't propagate; when a segment is misidentified in a frame, this doesn't fatally contaminate the appearance model. Difficulties occur if different individuals look the same (although one may be able to deal with this by instancing) or if we fail to build a model.

4.2.3.4 Inferred Discriminative Appearance Models

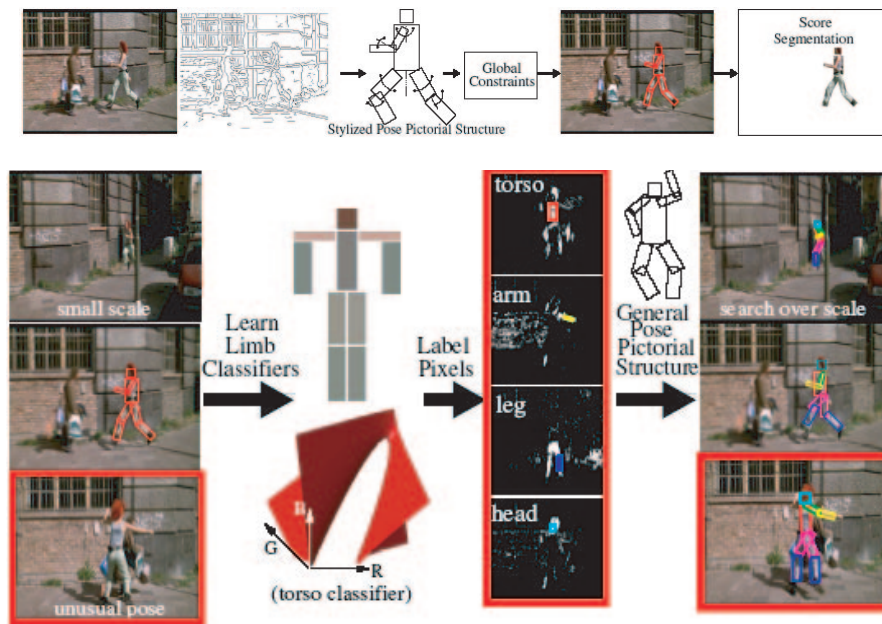


Fig. 4.9 **Deva's thesis**, also CVPR05 Ramanan et al. demonstrate that one can build appearance models by looking for human configurations that show all limbs and are easily detected. It turns out that, even in quite short sequences of people engaging in quite extreme behaviour, one can find lateral walking views. **Top:** These views can be detected by using a pictorial structure model on an edge-based representation, using quite low entropy links to impose the requirement that one has a lateral view of walking. This detector is tuned to produce no false positives — false negatives are quite acceptable, as long as one instance is found. **Bottom:** Once an instance has been found, we have the basis of a discriminative appearance model, because we know what each limb segment looks like and we have a lot of pixels that do not lie on a limb segment. Ramanan et al. build a discriminative appearance model for each body segment using logistic regression, then apply a pictorial structure model to the output of this process — so that a good segment match contains many pixels where $P(\text{segment}|\text{pixel values})$ are high. The resulting tracker is illustrated in figure 4.10.

Ramanan *et al.* demonstrate an alternative method of building a model. Assume that people occasionally adopt a pose that is (a) highly stylized (and therefore easy to detect) and (b) displays arms and legs clearly (so that appearance is easy to read off). Then, if we reliably detect at least one instance of this pose without false positives, we can read off an appearance model from the detection. Furthermore, we can make this appearance model discriminative, because we have a set of pixels that clearly do lie on the segments, and others that clearly do not. It is an empirical property that people do seem to adopt such poses, even in sequences of quite complex motions. They are relatively straightforward to detect by matching an edge template using a pictorial structure model. Notice that we are helped by the detection regime here — we don't need to detect every instance, just enough to build an appearance model, but we don't want false positives. Ramanan *et al.* use logistic regression to build discriminative models for each limb segment, then a pictorial structure model to detect. Again, tracking is a simple matter of detecting instances of the model and linking those that meet a velocity constraint. These discriminative models significantly reduce the difficulty of searching for an instance of a person, because much of the image is discarded by the models. In particular, the models can emphasize aspects of appearance that distinguish a particular individual from that individual's background. In his thesis, Ramanan shows that a discriminative model of appearance results in significantly better tracking behaviour (figure 4.11).

4.2.4 Parts as Codebooks

Song *et al.* use a variant of tree-structured models to identify human motion. They identify local image flows at interest points in an image, using the Lucas-Tomasi-Kanade procedure for identifying and tracking localizable points [412, 413]. For a fixed view of a fixed activity, flows at various interest points on the body are strongly related, and discriminative. They build a triangulated graph, whose nodes represent the state of each interest point on the body and whose edges represent the existence of a probabilistic relation between the nodes. Because this graph is triangulated, the junction tree is straightforward to find and inference is relatively simple (see, for example, [208]). They then detect human motion by identifying the best correspondence between image flow features and graph nodes and testing against



Fig. 4.10 **Deva's thesis** Frames from sequences tracked with the methods of Ramanan et al., where a discriminative appearance model is built using a specialized detector (figure 4.9), and then detected in each frame using a pictorial structures model. The figure shows commercial sports footage with fast and extreme motions. On the **top**, results from a 300 frame sequence of a baseball pitch from the 2002 World Series. On the **bottom**, results from the complete medal-winning performance of Michelle Kwan from the 1998 Winter Olympics. We label frame numbers from the 7600-frame sequence. For each sequence, the system first runs a walking pose finder on each frame, and uses the single frame with the best score (shown in the **left insets**) to train the discriminative appearance models. In the baseball sequence, the system is able to track through frames with excessive motion blur and interlacing effects (the **center inset**). In the skating sequence, the system is able to track through extreme poses for thousands of frames. The process is fully automatic.

a threshold. One requires multiple models for multiple activities, though how many models might be needed to cover a wide range of activities and aspects is a difficult question. The method is effective at identifying human motion; note that frames are explicitly not linked over time, something that doesn't seem to cause any real difficulties for the method, which should be seen as an early track-by-detection method.

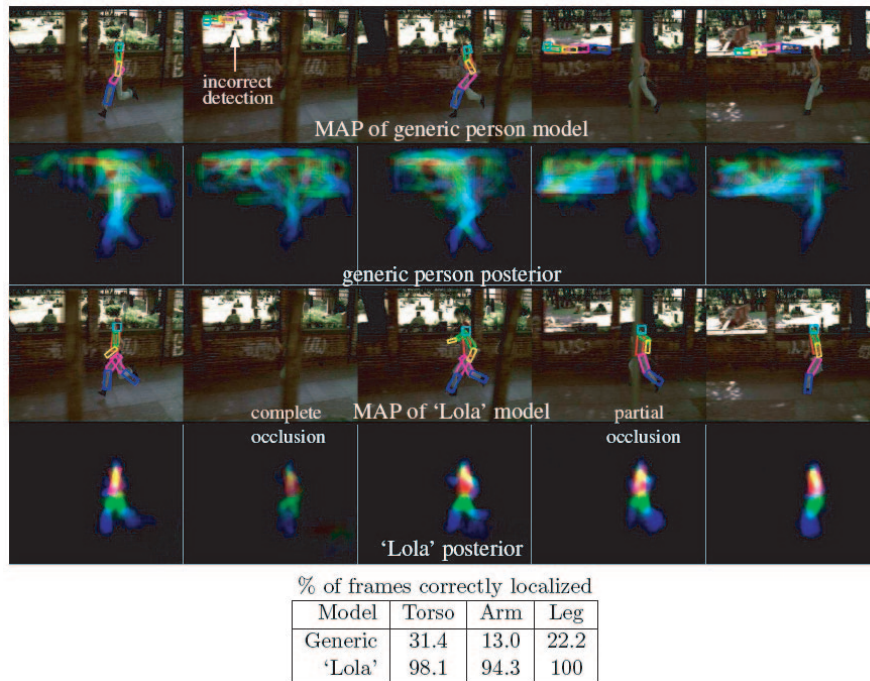


Fig. 4.11 **Deva's thesis** Ramanan shows that tracking people is easier with an instance-specific model as opposed to a generic model []. The **top two rows** show detections of a pictorial structure where parts are modeled with edge templates. The figure shows both the MAP pose — as boxes — and a visualization of the entire posterior obtained by overlaying translucent, lightly colored samples (so major peaks in the posterior give strong coloring). Note that the generic edge model is confused by the texture in the background, as evident by the bumpy posterior map. The **bottom two rows** show results using a model specialized to the subject of the sequence, using methods described above (part appearances are learned from a stylized detection). This model does a much better job of data association; it eliminates most of the background pixels. The table quantifies this phenomenon by recording the percentage of frames where limbs are accurately localized — clearly the specialized model does a much better job.

4.3 Evaluation

There is no current consensus on how to evaluate a tracker, and numerical evaluations are relatively rare; figure 4.13 shows results from all evaluations of which I am aware. There are several numerical evaluations of lifting to 3D; see, for example, [3, 246, 245]. In our opinion, it is insufficient to simply apply it to several video sequences and show some resulting frames (a practice fairly widespread until recently). Counting the number of frames until the tracker fails is unhelpful: First, the tracker may not fail. Second, the causes

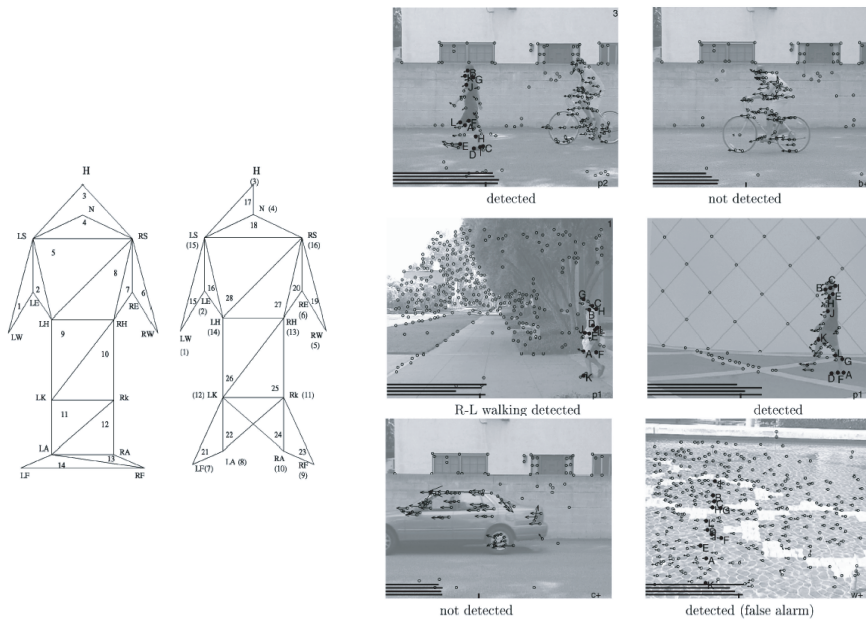


Fig. 4.12 **devatracker/01206511, figs 1 and 8** On the **left**, two triangulated graph models of the human figure. Each node represents the state of some interest point on the body; because the graph has a triangulated form and simple cliques, the junction tree is easy to obtain and inference is relatively straightforward (one could use dynamic programming on the junction tree). Song et al. use this representation to detect people engaged in known activities, using learned models to infer the form of the distributions represented by the edges of the graph [412, 413]. They detect flow at interest points in the image, then use these models to identify the maximum likelihood labelling of the image interest points in terms of the body interest points; detection is by threshold on the likelihood. On the **right**, some detection examples. Note the method is generally successful.

of failure are more interesting than the implicit estimate of their frequency, which may be poor. Third, this sort of test should be conducted on a very large scale to be informative, and that is seldom practical. Trackers are — or should be — a means to a larger end, and evaluation should most likely focus on this point. In this respect, trackers are probably like edge-detectors, in that detailed evaluation is both very difficult and not wholly relevant. What matters is whether one can use the resulting representation for other purposes without too much inconvenience.

A fair proxy for this criterion is to regard the tracker as a detector, and test its accuracy at detection and localization. In particular, if one has a pool of frames each containing a known number of instances of a person, one

can (a) compare the correct count with the tracker's count and (b) check that the inferred figure is in the right place. The first test can be conducted on a large scale without making unreasonable demands on human attention, but the second test is difficult to do on a large scale. Ramanan and Forsyth use these criteria; their criterion for whether a particular body segment is in the right place is to check the predicted segment intersects the image segment (which is a generous test) [356, 396].

Lee and Nevatia evaluate reprojection error for the tracked person [247]. There might be some difficulty in using this approach on a large scale. Sigal *et al* construct a 3D reconstruction, and so can report the distance in millimetres between the true and expected positions (predicted from the posterior) of markers [396]. Agarwal and Triggs give the RMS error in joint angles compared to motion capture on a 500 frame sequence [5].

There is little consensus on what RMS errors actually mean in terms of the quality of reported motion. There is some information in [?], which evaluates compression of motion capture; this boils down to the fact that very small RMS errors in joint position indicate that the motion is acceptable, but quite large errors are hard to evaluate. There is no information on what errors in joint angle mean.

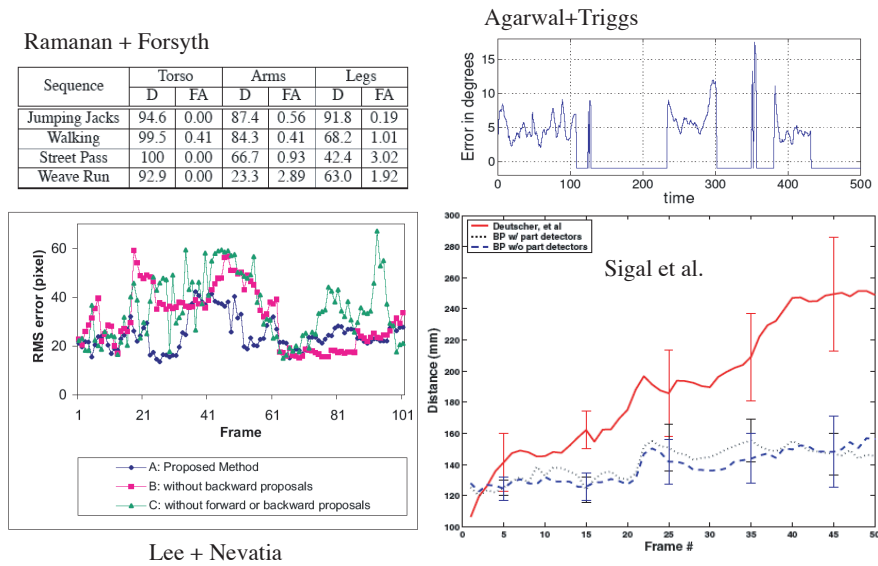


Fig. 4.13 Part of Figure 8 from devatrackers/0135063 and Table 1 of devatrackers/cvprfinal, fig 8 from 2d3dliftnew/01565379 and fig from devatrackers/leemotion05 On the top left, reports of the percentage of limb segments in the track that overlay the actual limb segments (D) and that are false alarms (FA) for a series of tracks using the methods of Ramanan and Forsyth, reported in [356]. On the bottom left, reports of RMS error of backprojected pose in pixels from the work of Lee and Nevatia [247]. On the top right, RMS error in joint angle for 500 tracked frames from Agarwal and Triggs; the zero error indicates a person was not present [5]. On the bottom right, distance between points on reconstructed 3D models obtained using the methods of Sigal et al. ([396]) and tracked motion capture markers supplying ground truth; there are two baselines, the method of Deutscher et al. ([104]), which fairly quickly loses track, and belief propagation without part detectors, which is surprisingly good.

5

Motion Synthesis

There are a variety of reasons to synthesize convincing looking human motion. Game platforms are now very powerful and players demand games with very rich, complex environments, which might include large numbers of non-player characters (NPC's — which are controlled by the game engine) engaged in a variety of activities. These figures need to move purposefully, react convincingly to impacts, and be able to change their activities on demand. Ideally, the motions are clean and look human; players can control characters smoothly; and there are no jumps or jerks resulting from sudden, unanticipated demands — which might originate either with a player or with game AI. Typically, this industry is willing to sacrifice a degree of quality if it can produce a very large volume of motions, and do so quickly. The film industry has traditionally been less interested in computational motion synthesis, largely because human animators — or, for that matter, actors — are still the best way to get high quality motion. This trend appears to be changing.

Another, perhaps less frivolous in purpose, is the simulation industry. Commodity graphics hardware has advanced to the point where many of the “immersive virtual reality” simulation and training applications which were proposed during the 1990's are now actually becoming quite practical. Many of these applications require environments that must be populated with hu-

mans. Currently, most such applications make do with minimally realistic human figures, but as recent computer games have demonstrated it is now possible to render humans with very realistic static appearance. Variations in rendering style alter a viewer's perception of motions [177, 178]. As the characters' appearance improves so too does viewer expectations concerning the characters' motion. More realistic characters with a more interesting range of behaviors present substantial challenges.

Situation simulations used for training military, rescue, and other hazardous-duty personnel are currently predominantly populated by unrealistic human characters. While these characters suffice for some aspects of training, they still place strong limitations of the simulation's potential effectiveness: a fire-rescue worker's response to a mannequin with the word "victim" is fundamentally different to the response that would be elicited by a character that behaves and appears like a frightened 10 year old child. Similar, but more gruesome, arguments can be advanced concerning the need for realistic humans in combat simulations [490, 491]. In either case, the desired goal is that the user become immersed in the simulation to the point where they behave as if the situation were real, and we believe that realistic simulated humans are required for this to happen.

Finally, an understanding of synthesizing accurate looking human motions may yield insight into the structure of motions. Possible benefits for the computer vision community include dynamical models for tracking humans, methods for determining whether a motion is human or not, and insights into action representation.

5.1 Fundamental Notions

5.1.1 The Motion Capture Process

Motion capture refers to special arrangements made to measure the configuration of a human body with (relatively) non-invasive processes. Early systems involved instrumented exoskeletons (the method is now usually seen as too invasive to be useful except in special cases) or magnetic transducers in a calibrated magnetic field (the method is now usually seen as unreliable in large spaces). More recent systems involve optical markers. One can use either **passive markers** (for example, make people wear tight-fitting black clothing with small white spots on them) or **active markers** (for example,

flashing infrared lights attached to the body). A collection of cameras views some open space within which people wearing markers move around. The 3D configuration of the markers is reconstructed for each individual; this is then cleaned up (to remove bad matches, etc.; see below) and mapped to an appropriate skeleton.

resolution; data association; missing markers

5.1.1.1 Engineering Issues

Motion capture is a complex and sophisticated technology; typical modern motion capture setups require a substantial quantity of skilled input to produce data, and there have been many unsuccessful attempts to build systems (or even to use commercial systems) within the academic community.

There are three main sources of difficulty. First, one requires **high resolution**, both in time and in space. High temporal resolution is required to localize in time the sharp accelerations caused both by contacts and by some kinds of motion — hitting, jumping, etc. Insufficient temporal resolution results in “squashy” motions, and 120Hz cameras are now common. There are attendant difficulties of getting pixels out of the camera fast enough. It is now typical to use cameras that produce only reports of marker position, rather than full frames of video. High spatial resolution is required to avoid “pops” — a fast snapping movement from one frame to the next that can be the result of spatial noise — and jittery looking movement. The result is significant demands on the camera system, because it is desirable that each marker is seen by at least two cameras. This is difficult to achieve when there are many people because the body parts occlude one another. Furthermore, actors must move in a relatively large space, particularly if one wants to capture fast movements like running. The result of all this is that there must be many cameras, all kept very well calibrated; and because they are far from the markers, the cameras must have high resolution.

The second difficulty is **data association**. One must determine which reported marker position in which frame corresponds to a particular marker location on the body. Helpful cues include: camera calibration (which gives epipolar constraints); the relatively fast frame rate (meaning that nearest neighbours often propagates marker identities well); and the fact that some

aspects of the geometry of the figure to be reconstructed are known. Difficulties include: the sheer volume of measurements (meaning that there is a good chance that there are many views in which epipolar constraints are not as helpful as one would want); the possibility that some markers are seen in only some views (meaning one cannot afford to simply throw away reports); the fact that some movements are fast (meaning that nearest neighbours can be misleading). There is a tendency, in our opinion premature, to feel this problem is solved (but see, for example, [230]).

The third difficulty is **missing markers**. To be reconstructed, a marker needs to be visible to at least two cameras with sufficient baseline, and the correspondence needs to be unambiguous. Occasionally, this isn't the case, usually as a result of occlusion by other bodies or body parts.

5.1.1.2 Cleanup and Skeletonization

Typical workflow involves capturing 3D point positions for markers, discounting or possibly correcting any errors in correspondence by hand, then using software to link markers across time. There are usually errors, which are again discounted or corrected by hand. Motions are almost always captured to animate particular, known models. This means that one must map the representation of motion from the 3D position of markers to the configuration space of the model, which is typically abstracted as a **skeleton** — a kinematic tree of joints of known properties and modelled as points separated by segments of fixed, known lengths, that *approximates* the kinematics of the human body. The anatomy of the major joints of the body is extremely complex, and accurate physical modelling of a body joint may require many revolute and prismatic joints with many small segments linking them (the shoulder is a particularly nasty example [117, 446], but, for example, the drawings in [68] emphasize the complex kinematics of human joints). This complexity is unmanageable for most purposes, and so one must choose a much lower dimensional approximation. Different approximations have different properties — the details are a matter of folklore — and one chooses based on the needs of the application and the number of degrees of freedom of the skeleton. Skeletonization is not innocent, and it is usual to use artists to clean up skeletonized data, essentially by adjusting it until it looks good. The pernicious practice of discarding point data once it has been skeletonized is widespread, and it re-

mains the case that data represented using one skeleton cannot necessarily be transferred to a different skeleton reliably. **Reviews** of available techniques in motion capture appear in, for example [50, 154, 266, 281, 288, 397].

5.1.1.3 Configuration Representations

For the moment, fix a skeleton. While this isn't usually an exact representation of the body's kinematics, we will assume that giving the configuration of this skeleton gives the configuration of the body. The configuration of the skeleton can be specified either in terms of its **joint angles**, or in terms of the position in 3D of the segment endpoints (**joint positions**). Not every set of points in 3D is a legal set of segment endpoints (the segments are of fixed lengths), so sets of points that are a legal set of segment endpoints must meet some **skeletal constraints**. The set of all legal configurations of the body is termed the **configuration space**; the joint angles are an explicit parametrization of this space, and sets of points in 3D taken with constraints can be seen as an implicit representation.

5.1.1.4 Skinning

I think there are many skinning methods. The most standard one is linear blend skinning. I don't think I've ever seen example based skinning in a realtime character (for the full body).

In animation applications, one wants the motion capture data to drive some rendered figure — when the actor moves an arm, the virtual character should do the same. The virtual character is represented as a pool of textured polygons, and one must determine how the vertices of these polygons change when the arm is lifted. The process of building a mapping from configuration — always represented as joint angles for this purpose — to polygon vertices is referred to as **skinning**. Skinning methods typically determine an appropriate configuration for the skin for each of a set of example poses, then interpolate [291]. One represents configuration as joint angles for skinning purposes because using joint positions is unwieldy (one would have to manage the constraints; we are not aware of any advantage to be obtained by doing so).

5.1.2 Footskate

An important practical problem is **footskate**, where the feet of a rendered motion appear to slide on the ground plane. In the vast majority of actual motions, the feet of the actor stay fixed when they are in contact with the floor (there are exceptions — skating, various sliding movements). This property is quite sensitive to measurement problems, which tend to result in reconstructions where some point quite close to, but not on, the bottom of the foot is stationary with respect to the ground. The result is that the reconstructed foot appears to slide on the ground (and sometimes penetrates it). The effect can be both noticeable and offensive visually. Footskate can be the result of: poorly placed markers; markers slipping; errors in correspondence across space or time; reconstruction errors; or attempts to edit, clean up or modify the motion. Part of the difficulty is that the requirement that the base of the foot lie on the ground results in complex and delicate constraints on the structure of the motion signal at many joints. These constraints appear to have the property that quite small, quite local changes in the signal violate them. It is likely that these properties are shared by other kinds of contact constraint (for example, moving with a hand on the wall), but the issue has not arisen that much in practice to date.

There are methods for cleaning up footskate. Kovar *et al.* assume that constraints that identify whether heel or toe of which foot is planted in which frame (but not where it is planted) are available [237]. Kovar *et al.* then: choose positions for each planted point, determine ankle poses to meet these constraints; adjust the root position and orientation so that the legs can meet the resulting ankles; compute legs that join the root and the ankle mainly by adjusting angles, but occasionally by adjusting leg lengths slightly; and then smooth the adjustment over multiple frames. The method is effective and successful.

Ikemoto *et al.* demonstrate that one can clean up footskate introduced by editing and so on automatically [192]. They build a classifier that can annotate frames in a collection with toe and heel plant annotations. These annotations are preserved through editing, blending, etc. When a motion has been assembled from edited frames, the annotations are smoothed over time, and the method then identifies possible footplant positions automatically by looking at the foot position over the time period of the footplant. Finally, inverse

kinematic methods (section 5.1.3) are used to clean up the frames.

5.1.3 Inverse Kinematics

Footskate cleanup is an example of a more general problem — adjust the joint angles of a motion so that it meets some constraints on joint positions. Assume we have a fixed skeleton; we now wish to clean up a motion referred to this skeleton, perhaps moving a foot position or ensuring that a contact occurs between a hand and a doorhandle. This creates a difficulty for either representation of configuration: if we work with joint angles, we must obtain joint angles such that the constraint is met; if we work with joint positions, we must obtain a set of joint positions that meet both this constraint and the skeletal constraints. We will confine our discussion to the case of joint positions, which is more important in practice.

For the moment, let us consider only a single frame of motion. Write the vector of joint angles as θ , and the joint positions as a function of joint angles as $\mathbf{x}(\theta)$. Assume that we would like to meet a set of constraints on joint positions $\mathbf{g}(\mathbf{x}) = 0$. The problem of **inverse kinematics** is to obtain a θ such that $\mathbf{g}(\mathbf{x}(\theta)) = 0$. The constraint is important in the formulation, because we hardly ever wish to specify a change in every joint position. For example, assume we wish to move the elbow of a figure so it rests on a windowsill — we would like to adjust the kinematic configuration so that the elbow lies at a point, but we don't wish to specify every joint position to achieve this. Notice there is room for some confusion here. In the robotics and theoretical kinematics literature, the problem is almost always discussed in terms of choosing joint angles to constrain the endpoint configuration of a manipulator. In graphics applications, the term refers to meeting any kinematic constraint.

Under some conditions, closed form solutions are available for at least some parameters (e.g. see [232, 233, 439, 440]). Alternatively, a solution can be interpolated: D'Souza *et al.* learn inverse kinematics for a humanoid robot with locally weighted regression [114], and Schaal *et al.* describe learning methods for a variety of robot problems, including inverse dynamics [379].

More usually, one must see this as a numerical root finding problem. The update for Newton-Raphson method involves finding a small change in configuration $\delta\theta$ such that $\mathbf{g}(\mathbf{x}(\theta_0 + \delta\theta)) = 0$. We may be able to obtain $\delta\theta$

from

$$\begin{aligned} 0 &= \mathbf{g}(\mathbf{x}(\theta_0 + \delta\theta)) \\ &\approx \mathbf{g}(\mathbf{x}_0 + \mathcal{J}_{\mathbf{x},\theta}\delta\theta) \\ &\approx \mathbf{g}(\mathbf{x}_0) + \mathcal{J}_{\mathbf{g},\mathbf{x}}\mathcal{J}_{\mathbf{x},\theta}\delta\theta \end{aligned}$$

where $\mathcal{J}_{\mathbf{x},\theta}$ is the jacobian of \mathbf{x} with respect to θ , etc. In the ideal case, the product of jacobians is square and of full rank, but this seldom happens. For almost every point in the configuration space, the rank of the jacobian $\mathcal{J}_{\mathbf{x},\theta}$ should be the dimension of the configuration space (if this isn't the case, then we have a redundant angle in our parametrization; we assume that this does not happen). At some points, the rank of this jacobian will go down — these are the kinematic singularities of section 2.4.1. The practical consequence of this is that some position updates may not be attainable (for example, consider the straightened elbow of section 2.4.1; the only instantaneous hand velocity attainable is perpendicular to the forearm). The rank of $\mathcal{J}_{\mathbf{g},\mathbf{x}}$ may be small. For example, if our constraint requires that a point be in a particular place, the rank will be three. This is a manifestation of **kinematic redundancy**, which is a major nuisance. A natural strategy to deal with constraint ambiguity is to obtain a least squares solution for $\delta\theta$ — but the resulting pose may not be natural (one can use other norms, see [103]). A second source of difficulties in the optimization problem are joint limits, which mean that our optimization problem is subject to some inequality constraints on the θ . The **feasible set** of solutions that meet this constraints is not necessarily convex, which can mean the general optimization problem is hard.

Kinematic redundancy is a global rather than local matter. There may be more than one θ such that $\mathbf{g}(\mathbf{x}(\theta)) = 0$. For example, assume that we wish to constrain a figure to stand with its feet on the floor in given spots, and a hand on a given spot on a wall. Typically, there is either no solution to these constraints — the wall is too far away — or many. The collection of solutions is rather rich (stand next to a wall with your hand on the wall; you can move in all sorts of ways without having your feet move or your hand leave the wall), and could be continuous or discrete.

All this creates a nasty problem. Applying inverse kinematics on a frame-by-frame basis may produce solutions at each frame that are inconsistent (as a result of kinematic redundancy). This is complicated by the presence of multiple solutions, and the vagaries of root finding. For example, assume we want

a solution where the hand is against the wall, as above. In frame n , the root finder converges to a solution where the elbow is below the shoulder; but the start point for frame $n + 1$ is slightly different from that for frame n , and the root finder could find a solution where the elbow is above the shoulder. This sort of behaviour results in noticeable and annoying “pops” in the motion. The effect can be countered by adjusting multiple frames simultaneously, but this is expensive computationally; much of the recent literature is a search for efficient approximation methods.

The use of inverse kinematics in animation dates to at least the work of Girard and Maciejewski [152]; see also [151] and [166]. Methods for handling singularities are discussed in [272]. A good summary of early work in animation is [24]. Tolani *et al.* contains a considerable body of helpful background and review material [440]. Zhao and Badler approach inverse kinematics as a nonlinear programming problem — using our notation, find $\arg \min |g(\mathbf{x}(\theta))|$ subject to joint constraints, etc. — and use a variant of a standard optimization method; it is not possible to guarantee a global minimum (neither the objective function nor the constraints are convex) [482]. Incompatible constraints can be handled by a scheme allocating different priorities to constraints [25]. Shin *et al.* obtain a real-time solver for a puppetry application by linking a fast frame-by-frame solver using a mixed analytical-numerical strategy with a Kalman filter smoother [392]. Joints other than the shoulder have been studied in some detail [295, 331].

5.1.4 Resolving Kinematic Ambiguities with Examples

The danger here is that one may obtain poses that do not look human. Motion editing deals with this by being interactive, so that an animator who doesn’t like the results can fiddle with the constraints until something better appears (see also [334]). An alternative is to allow relatively few degrees of freedom — for example, allow the animator to adjust only one limb at a time — or to require similarity to some reference pose [430, 476, 483]. This isn’t always practical. An alternative, as Grochow *et al.* demonstrate, is to build a probabilistic model of poses and then obtain the best pose [164].

One can do this as follows (for consistency within this review, our notation differs from that of Grochow *et al.*). Write \mathbf{y} for a feature vector describing a pose \mathbf{x} (the feature vector could contain such information as joint posi-

tions, velocities, accelerations, etc.). Write \mathbf{u} for the (unknown) values of a low dimensional parametrization of the space of poses. Use the subscript i to identify values associated with the i 'th example. Now assume we have a regression model $P(\mathbf{y}|\mathbf{u}, \theta)$ for θ some parameters (which in this case choose a model and weight components with respect to one another). We could obtain an inverse kinematic solution by maximizing

$$P(\mathbf{y}(\mathbf{x}), \mathbf{u}|\theta)$$

with respect to \mathbf{x} and \mathbf{u} , subject to some kinematic constraints $g(\mathbf{x}) = 0$. Notice we need \mathbf{x} (the configuration of the body), \mathbf{y} (the feature vector) and \mathbf{u} (the low dimensional representation) here. This is because \mathbf{u} does not predict a unique \mathbf{y} — we may need to choose a body configuration that is close to, but not on, the low dimensional structure predicted by the model — and because many poses might have the same feature representation. The regression model is built using N examples \mathbf{y}_i (note we do not know \mathbf{u}_i for these examples). We assume the examples are independent and identically distributed (note the independence assumption needs care with motion data; frames may be correlated over quite long timescales), and obtain \mathbf{u}_i, θ to maximise

$$P(\mathbf{u}_i, \theta|\mathbf{y}_i)$$

Grochow *et al.* use a scaled gaussian process latent variable model as a regression model, and note that some simpler models tend to overfit dramatically. The method produces very good results; authors note that a form of rough-and-ready smoothing (obtained by interpolating between parameters obtained with clean training data and training data with added noise) seems to produce useful models that allow a greater range of legal poses.

While motion editing does not offer direct insight into representing motion, the artifacts produced by this work have been useful, and it has produced several helpful insights. The first is that it is quite dangerous to require large changes in a motion signal; typically, the resulting motion path does not look human (e.g. [154]). The second is that enforcing some criteria — for example, conservation of momentum and angular momentum [391]; requiring the zero-moment point lies within the support polygon [97, 231, 391] — can improve motion editing results quite significantly. However, note that one can generate bad motions without violating any of these constraints, because motion

is the result of extremely complex considerations. The third is that requiring motion lie close to examples can help produce quite good results.

5.1.5 Specifying a Motion Demand

One must specify what is desired to a motion synthesis algorithm. While synthesis algorithms tend to vary quite widely, there are not many options for constraints. **Geometric constraints** may constrain: the position or position and orientation of the root; the position or position and orientation of one or more body segments; or, in extreme cases, the exact configuration of the body (in which case the **frame constraint** can be thought of as a **keyframe**). Geometric constraints may take various forms involving either equalities or inequalities. For example, one may constrain a point to lie on a plane, a line, or a point (which are all equality constraints), or to lie within a region (an inequality constraint).

Depending on algorithmic details, constraints may be either exact or represented as a penalty function. Constraints may be either **summary constraints**, applying to the position and orientation of a summary of configuration such as the overall center of gravity or the root, or **detailed constraints**, applying to individual body segments or particular points on the body. One can apply either **instantaneous constraints**, which constrain at a particular time, or **path constraints**, which constrain to a path over a period of time. It is common, but not universal, to assume that a path constraint comes with timing information. It is usual to assume that impossible constraints are not supplied.

Such constraints can be used to sketch out the structure of a motion in greater or lesser detail, depending on what an algorithm requires. In most cases, however, they don't determine the motion. For example, in some cases quite a precise temporal parametrization of a path may not determine whether a figure must run or walk. Usually, one would like to supply relatively few constraints (authoring constraints is a nuisance), meaning that the resulting motion is usually dramatically ambiguous. There are almost always very many ways to meet instantaneous summary constraints for the start and the end of a motion (i.e. start here at this time, end there at that time). One might dawdle at the start, then sprint; walk very slowly; run, walk, then run, then dawdle, and so on.

Annotation constraints are intended to reduce this ambiguity. These constraints are demands that a motion be of a particular type, that are painted on the timeline. The interesting issue is how one encodes the type of a motion. Arikan *et al.* choose a set of 13 terms (“run”, “walk”, “jump”, “wave”, “pick up”, “crouch”, “stand”, “turn left”, “turn right”, “backwards”, “reach”, “catch”, “carry”) that appear to be useful for their dataset [17]. It is desirable to respect the fact that motions can compose — for example, one can run while carrying — and they do so by allowing any combination of these terms to be an annotation. One can visualize an annotation as a bit vector, with 13 entries, one per term. This model ignores the fact that most combinations of annotations — e.g. “stand” and “run” — are meaningless; this is deliberate, because there isn’t a principled way to build a space of legal annotations and dependencies between annotations may result in nasty inference problems. Arikan *et al.* then mark up a collection of motion capture data using classifiers. The features are a representation of a pool of motion frames spanning the frame to be classified. The classifiers are trained independently, one per term, by marking up some frames, fitting a classifier, and then repeatedly classifying all frames, viewing and correcting a sample of labelled motions, and fitting a new classifier. This process converges quickly, allowing a large pool of motion to be marked up relatively quickly, probably because it is easy to view a large pool of animations and correctly identify mislabelled motions. The result is a pool of frames of motion capture data, each carrying a vector of 13 bits, each of which is determined independent of the others. Interestingly, Arikan *et al.* point out that, although their model does not exclude inconsistent annotations, relatively few of the 2^{13} available annotations are actually applied, and they observe no inconsistent annotations.

5.2 Motion Signal Processing

Think of a motion as a time-parametrized path on some space describing kinematic configuration. Now assume we have two such paths that are “close”. Assume we have a good correspondence between the paths. We expect that a convex combination of corresponding frames may result in a good motion, and that this may still be true if the weights are time-varying. It turns out that these expectations are largely met. In fact, a variety of such operations on motion are successful, an observation originating with Bruderlin and

Williams [64].

5.2.1 Temporal Scaling and Alignment

As Bruderlin and Williams point out, if one runs a motion slightly faster or slightly slower, the result is still usually an acceptable motions [64]. The advantage of this observation is that we can **time align** motions. Assume we have two motions which are sampled at the same rate. An **alignment** between the motions is a pair of functions, $c_1(i)$ and $c_2(i)$, which identifies which frame from the first (resp. second) motion to use at the i 'th time-step. Generally, we want to align motions so that, for some norm,

$$\| \mathbf{X}_{c_1(i)}^{(1)} - \mathbf{X}_{c_2(i)}^{(2)} \|$$

is small. As Kovar and Gleicher show, such an alignment can be computed with dynamic programming [234]. For most reasonable applications, the norm should be invariant to the root coordinate system of the frames, and this can be achieved most easily by representing the frames with joint positions, and computing the minimum sum-of-squared distances between corresponding points over all Euclidean transformations using the method of Horn [184, 185]. In practice this alignment should be thought of as inserting (resp. deleting) frames from each motion so that the sequences align best. Typically, we are interested in i running from 1 to k ; if we reindex each motion so that the first frames of each correspond, we typically want constraints on the magnitude of $c_1(i) - i$ and $c_2(i) - i$. Furthermore, we want each correspondence to advance time, so that for each, $c(i) - c(i - 1) \geq 0$. We are not aware of any alignment methods that interpolate and resample motions to obtain corresponding frames, but this is a natural extension of the general blending framework.

5.2.2 Blending, Transitions and Filtering

Now assume we have two motions with a time alignment. At each timestep t_i , we have a pair of frames that could (we believe) be blended. To produce a blend effectively, we must determine (a) the root coordinate system of the blended frame and (b) where the two source frames should lie in that root coordinate system. Assume, for the moment, that these problems have been

solved. Producing a blend is then straightforward — we form $\mathbf{X}_{c_1(i)}^{(1)}\phi(t_i) + (1 - \phi(t_i))\mathbf{X}_{c_2(i)}^{(2)}$ in the appropriate coordinate system.

Doing all this requires careful handling of the root. If our representation contains the root, then we will be able to blend very few motions because even motions that are similar may occur in different places, which is clearly a waste of data. However, we cannot simply strip every frame of root information, because the root path is often quite strongly correlated with the body pose. For example, people use different gaits for fast and slow translational movements. As another example, an actor trying to move quickly along a root path with a sharp kink in it typically makes a form of braking and pivoting step.

The solution seems to be to (a) ignore root information for the whole sequence (rather than per frame; as a result, we preserve velocity and angular velocity information) and (b) allow small deformations of the root paths so they line up with one another. It is difficult to be precise about what “small” means here, though a moderate degree of warping in both time and space still results in a good motion [234].

The root coordinate system of the blended motion is typically obtained from the motion demand. One may simply rotate and translate the frames into this coordinate system (as [193] do), or one may interpolate and smooth the transformations that do so (as [234] do). Furthermore, as Safonova and Hodgins show, linear blends can produce motions that are physically inoffensive [377].

5.2.2.1 Multi-way Blends

Bruderlin and Williams envisage blending more than two sequences; doing so leads to better motions [64] (see also [368, 466]). Kovar and Gleicher give methods to find motions that are similar, and so can be blended, and to create parametrized blend spaces involving these examples [235]. Locomotion is a particularly important and common form of motion. There are several methods to create parameterized blend spaces for each of walking, running, and standing [241, 332, 333]. These methods can generate realistic transitions between these three types of motion.

5.2.2.2 Transitions

A particularly important application of blending is to produce **transitions** — motions that “link” activities, for example, the slowing pace one takes when moving from a run to a walk. Lee *et al.* blend to produce links in a motion graph (section 5.3) [243]. Their method assumes one has two frames known to be similar, and blends a window in the future of one frame with a window in the future of the other.

5.2.2.3 Filtering

As Bruderlin and Williams establish, one may apply a filter to joint angle or joint position signals and obtain a good motion by doing so [64]. Ikemoto and Forsyth show that constant offsets to joint angles sometimes result in good motions as well [194]. As far as we know, there are no guidelines about what is likely to be successful here.

5.2.2.4 Physical Blends

Arikan *et al.* need to produce transitions between distinct motion sequences on-line to meet realtime demands [19]. At the point of transition, there is a discontinuity as the motion jumps from the last frame of the working sequence to the first frame of the next sequence. Arikan *et al.* produce a final frame by adding an offset vector to the measured frames. This offset vector decays with time as a second order linear system; the discontinuity is avoided by subtracting from the offset at the transition point, so that the sum of frames and offset has no discontinuity.

5.2.2.5 What to Blend

It is important to know when two sequences can be blended successfully. Lee *et al.* choose to blend when the distance between a pair of frames, evaluated as a weighted sum of differences in joint angles, is small [243]. Wang and Bodenheimer demonstrate that the choice of weights this algorithm for identifying similar frames is important, and show that better weights than those used in the original paper can be learned from data [460]. One could reasonably hope for a more extensive criterion than just requiring some frames to be close

and this seems like a productive area of study. Wang and Bodenheimer show that the size of the difference between frames gives some cue to the length of an appropriate transition, as does the velocity [461]. Arikan *et al.* wish to produce motions that look like responses to pushes or shoves [19]. To do so, they produce many possible transitions to many distinct sequences, each with a physical deformation, then use a regression method to determine which best serves the motion demand encoded by the push. This strategy of searching multiple cases for a good motion is extended to blends by Ikemoto *et al.*, who produce a very large range of blends, then test the resulting blended motions to see which is good [193]. Doing so successfully requires a good method to evaluate motions, a difficult problem we discuss in section 7.2.1.

5.2.2.6 Finding Similar Motions

Current blending methods blend motions that are “close”; this means we need methods to find such motions. Kovar and Gleicher describe a method to build fast searches of a motion collection for matching motions, where a match is defined by time-aligning a pair of motion sequences and then scoring frame-frame differences in a root-invariant manner [235]. They use a combinatorial structure (a “match web”) to encode possible search results, so that search is fast. Time-alignment and scoring may not reveal good matches — for example, two walks that are out of phase might look very different, but be good matches — and Kovar and Gleicher deal with this by repeated matching to match results.

Forbes and Fiume represent frames of motion on a basis obtained with weighted PCA (the weights are necessary because small changes in hip angle can generate large variance in toe position, which gives PCA basis that behaves badly), and match by first searching for discriminative seed points, then time-aligning the query signal with the motion dataset [129].

These methods work well at matching motions to motion queries, but the need for a query motion can become burdensome, for example, if an animator is searching for a motion in a collection. Müller *et al.* encode motion frames with binary predicates — for example, is a foot ahead of, or behind, the plane of the body — and then search for either precise or soft matches to a predicate [299]. These predicates can be surprisingly expressive; for example, an appropriate combination can recover frames at any phase of a walk (left foot

forward and right foot back, or right foot forward and left foot back; and so on).

Wu *et al.* cluster frames, then match a sequence of cluster centers under dynamic time warping [472]. Keogh *et al.* point out that general time warping can lead to serious alignment problems, and argue that a uniform time scale is a generally better model ([224]; see also [223, 70] and see [222] for an exact indexing method under dynamic time warping). They demonstrate an extremely fast method for finding sequences within a uniform time scale of a given sequence, using a combination of bounds and R-trees.

5.2.2.7 Difficulties with Blending

Assume we have two motions both captured at the same frequency. Both contain temporally localized large accelerations (for example, they might be grabbing or hitting motions). The temporal parametrization of the motions is slightly different, meaning that the samples are aligned slightly differently in time with respect to the motions. Even at the best possible time alignment, if we blend these motions we expect to lose some of the structure at high temporal frequencies — which would be the large accelerations. The result is a motion that can be “squashy” in appearance and can lose its temporal crispness. This problem doesn’t always occur, and might be manageable if one is careful (for example, it might be worth reconstructing motions using some form of interpolation, resampling at very high frequencies, then aligning the resampled motions). A version of this problem occurs for dimension reduction methods, too.

5.3 Motion Graphs

Motion capture data is used in very large quantities by, for example, the movie and computer game industries. For each title that will contain human motion, an appropriate script of motions is produced; typically, this involves a relatively small set of “complete” motions that can be joined up in a variety of different ways. This script is captured, and then motions are generated within the game by attaching an appropriate set of these motion building blocks together. Motions captured for a particular title are then usually discarded as re-use presents both economic and legal difficulties.

This suggests a form of directed graph structure encoding legal transitions between motions. The attraction is that if we have such a graph, then any path is a legal motion; thus, with some luck, much of the work of motion synthesis could be done in advance. Furthermore, it may be possible to issue quality guarantees for any synthesized motion if we can do so locally within the graph. This hope has not yet materialized, but remains an attraction of the representation. Another attraction of this approach is that it can be used to synthesize more than just motions; for example, Stone *et al.* show that one can use a similar approach to synthesize both motion and synchronized audio for utterances from a synthetic character [424].

There are several ways to implement this graph structure, but the important matter here is a representation of legal motion transitions. The simplest, which we favour as a conceptual (but not necessarily computational) device is to regard every frame of motion as a node and insert a directed edge from a frame to any frame that could succeed it. We will call this object a **motion graph**, and always have this representation in mind when we use the term. An alternative representation is to build a set of unique clips (runs of frames where there is no choice of successor — one could build these by clumping together nodes in the previous representation that have only one successor), use the unique clips as edges and make **choice points** into nodes. In this representation, one thinks of running one clip which ends in a node where we can choose which clip to run next. Finally, we could make each clip be a node, and then insert edges between nodes that allow a cut. Here we must be careful with the semantics, because there could be more than one edge from node to node — it may be possible to cut from clip *A* to clip *B* in different ways — and our edges need to carry information about where they leave the source clip and where they arrive at in the target clip. There is no difference of substance between the representations; we favour the first, as we find it easier to think about.

5.3.1 Building a Motion Graph

A set of observed motion sequences is a motion graph (there is a pool of frames, and a set of **observed edges**). This graph can be made significantly more useful by adding directed edges — which we call **computed edges** — from each frame to any frame that could succeed it in some sequence. Typ-

ically, we do so by identifying places where we can build a **transition** — a sequence of frames that starts at one frame in the graph (say, frame A_i in sequence A), ends at another (B_j in sequence B), and joins the frames preceding the start to those succeeding finish in a natural motion — and blending as in section 5.2.2 to build these transitions. This involves adding frames of interpolated motion.

5.3.1.1 Links by transitions

Kovar *et al.* build links by testing pairs of frames A_i and B_j to tell whether a transition of fixed length is possible between them, then building that transition [236]. They compare a window of fixed length into the future of frame A with a window of the same length into the past of frame B . Each window is represented as a set of points in 3D, and there are implicit correspondences. The distance is then the minimum sum of weighted squared distances between corresponding points available by choice of rigid-body transformation applied to one sequence. The weights are necessary because errors in some joint positions appear to be more noticeable than errors in other positions. This distance is computed for every pair of frames for which it exists (the future or the past might be too short). They build transitions between pairs of frames where the distance is a local minimum (the topology being supplied by the order of frames in the original sequences) and is lower than a threshold. The transition is built by aligning the windows with a rigid body transformation, then blending them. Footskate is avoided by identifying frames with footplant constraints, and blending in such a way as to preserve these constraints. There is no time or space deformation. If the motion graph is to be used in game applications, there is real value in allowing a designer to interact with this process, as Gleicher *et al.* show [158]. In this work, the designer can choose among possible “to” frames for a given “from” frame, and can disallow (resp. allow) transitions suggested (resp. discouraged) by the criterion above.

Links by similarity: Lee *et al.* test for a possible link from A_i to B_j by testing a distance between A_i and B_{j-1} , the logic being that if these two are sufficiently similar, then their futures could be interchanged ([243]; see also section 5.2.2). Notice that this suggests that if A_i can be linked with B_j , then we should be able to link from B_j to A_{i+1} . The distance is obtained as

a weighted sum of differences in joint angles, summed with differences in velocities at various points across the body (the choice of weights is important; see below). Two frames can then be linked if the distance is sufficiently small, the velocity term ensuring that the temporal ordering of motion is respected. Links between frames with dissimilar contact states, or that are not local maxima (again, the topology is given by the order of frames in the original sequences), are pruned.

Arikan and Forsyth represent frames as a sets of points in 3D in a coordinate frame centered on the torso, and obtain a distance by summing squared differences in positions and velocities in that frame taken together with the differences in velocity and acceleration of the torso frame itself [18]. Any edge where that distance lies below a threshold is inserted as a computed edge, with the direction being obtained from considerations of smoothness as below. They do not require any particular combinatorial structure in their graph, and so do not post process.

There are some tricks to building motion graphs that are not mentioned in the literature. It is important to keep carefully in mind that edges are directed. One should not confuse directedness of edges with symmetry in distances. If A_i and B_j are similar, that means that four motion sequences are acceptable: $\dots A_{i-1} A_i A_{i+1} \dots$, $\dots A_{i-1} A_i A_{i+1} \dots$, $\dots A_{i-1} A_i B_{j+1} \dots$, and $\dots B_{j-1} B_j A_{i+1} \dots$ (we do not count motions where A_i and B_j are substituted for one another).

Cleanup: Some applications require a fast decision at each choice point, meaning it may be hard to look far ahead in the graph when making that decision. In these cases, it is helpful to remove nodes that lack outgoing edges and graph components that cannot be escaped (see figure 5.1). This is best achieved by computing the strongly connected components of the graph (components such that, for any pair of nodes in the component, there is a directed path between them) and keeping the largest [236, 243].

Open issues: The methods we have described have generally been successful at producing usable motion graphs. There remain a number of open issues in building a motion graph. Identifying pairs of frames that allow one to build a transition is probably the right approach, but one could quibble with current implementations. It remains difficult to know whether one can or can't build transitions between a pair of frames (see section 5.2.2 above). One has no control over the **diameter** — the average length of the shortest path connecting two points in the graph — of the resulting graph. The di-




Fig. 5.1 *Examples of bad motion graphs. On the left, a motion graph where it is possible to get stuck in one component. This problem can be avoided by computing strongly connected components and taking the largest, at the possible cost of excluding some frames. The graph on the right has the difficulty that it is possible to get caught in a motion where no alternatives are available for many frames. This presents a difficulty if one wishes the motion to be responsive. Typically, there is a tension between obtaining high quality motions — which tend to require relatively few edges in the graph — and responsive motions — which tend to need as many edges leaving nodes. One would like a graph where the shortest path between two nodes is guaranteed to be (a) short and (b) good. No current method can guarantee to produce such a graph.*

ameter is important, because it affects the responsiveness of the motion — a synthesis program could reasonably demand a fast transition from one frame to another. Because current methods evaluate the goodness of an edge locally (but not the effect on the graph of incorporating it), they tend not to produce graphs with good combinatorial properties. Ikemoto *et al.* investigated a graph built competently with recent methods, and show that, for a reasonable choice of threshold, one has both that the shortest path between some quite similar frames can be very long, and that some pairs of frames are connected with very bad short paths [193]. It would be most attractive to have automatic methods that produce graphs of low diameter.

5.3.2 Searching a Motion Graph

We assume that our method of constructing edges is satisfactory, which means that any path in the motion graph is a motion. We can construct paths in the motion graph using local or global properties. A **local search** involves looking ahead some fixed number of frames. This means that the motion can respond to inputs, but may mean that some constraints can't be met. A **global search** involves looking at entire paths. The resulting motion is less responsive, but more easily constrained.

Local search methods: Kovar *et al* concentrate on choosing the next frame of motion, or, equivalently, choosing one of the outgoing edges at a choice node in the motion graph [236]. Nodes without outgoing edges are a problem; they can be removed with a simple graph algorithm. The choice of edge can be made in a variety of ways: one could look at a game controller, look at the local tangent direction of the desired root path, look at an annotation constraint (figure 5.2), or use a random variable. This latter approach can generate very good background motion when used with care. The trick is not to cut between motion sequences too often (because all methods of constructing motion graphs have flaws, and a path that contains mainly computed edges in the motion graph will tend to explore those flaws, and look bad). This can be achieved, for example, by choosing observed edges with rather higher probability than computed edges.




Fig. 5.2 **Fig 8, Mographs/mograph.pdf** These figures show motions synthesized using the motion graph method of Kovar *et al.* to meet path constraints and annotation constraints. The demand path is the coloured path on the ground plane; this is yellow for “walking”, green for “sneaking” and blue for “martial arts move”. The black path shows the projected root path, and the figures are frames sampled at even intervals to give a sense of the motion.

Local searches can run into problems. The motion graph might contain some frames that can be reached only by making the right choice at a choice point many frames away. In this case, choosing based only on a local criterion could make it impossible to meet some constraints, or at least meet them in a timely fashion. This is the **horizon problem** — a choice now might lead to trouble that is invisible, because it is on the other side of the horizon separating the future cases we consider from those we don’t. If the graph were guaranteed by the method of construction to have a short diameter, this problem would be much easier to handle. Other methods of coping with the horizon problem include: using a representation of available futures when

making a choice; choosing paths using some form of global search; and enriching the motion graph (the reasoning is that, with enough frames of motion in the graph, the diameter will be short without any explicit construction).

Taking the future into account: The body is capable of very fast accelerations. This suggests that, in a motion graph built with enough data, there is a fairly short path from any one frame to any other. In turn, this suggests that the horizon problem wouldn't be a problem if the horizon looked forward sufficiently far in time. However, in this case the range of futures available from a particular frame must be very large. Lee *et al.* encode the future in terms of clusters of frames [243]. These clusters form a graph, where each cluster is a node and there is an edge from one node to another if there is an edge from a frame in the cluster represented by the “from” node to a frame in the cluster represented by the “to” node. A given frame in the motion graph is associated with some node in this **cluster graph**. For any node in the cluster graph, we can construct a **cluster tree** — a tree, rooted at the node under consideration, that gives the nodes in the cluster graph accessible with a fixed number of hops. We now represent the available futures at a given frame by the cluster tree associated with that frame (there is a **cluster path** from the root to each leaf). Motions are controlled using either a choice based interface (where the animator chooses at each choice point), a sketch interface — where the sketch provides a demand signal — or a vision interface — where background subtracted frames from multiple viewpoints provide a demand signal. In both the sketch and the vision interface, frames are chosen by scoring the available cluster paths against the demand signal.

For this method, the choice of clustering criterion depends on the application. The alternatives are to represent the body relative to the root of the body, relative to the root of the body in the frame at the root of the cluster path, or in absolute coordinates. The first case is appropriate in uncluttered environments, where one can reasonably expect that any frame can occur at any location and in any orientation. The second can be appropriate when one needs anticipation — for example, synthesizing the run-up to a jump which must leave the ground at a point chosen during the synthesis procedure; this is a need one associates with animations in computer games that emphasize complex movements like jumps. The third case is appropriate to a cluttered environment, where a frame may be usable in only one spot in the motion domain.

Motion ambiguity: The family of acceptable paths through a motion graph that meet a given set of motion constraints is usually very large, a phenomenon we refer to as **motion ambiguity**. **Local motion ambiguity** arises because most motion data collections contain multiple copies of some motions — typically, walking and running — and that there is a rich collection of links between frames in these motions. As a result, there is a spectacular number of walking motion paths available. One could deal with this issue by clustering, but it isn't the major source of difficulty. The real problem is an important general peculiarity, which we call **global motion ambiguity**, which occurs because it is very seldom possible to author constraints on a motion animation that are unambiguous — the number of constraints required would be unnaturally large. This seems to be a result of the ways in which people find it natural to think about human motion (this issue will re-surface in our discussion of activity representations). For example, if I am instructed to go from point A to point B in some period of time, I can do so in a very large number of ways unless the constraints imply maximum velocity at all times. Some property of my motor control system is able to “fill in” sensible choices, so that the ambiguity is not apparent. One consequence of all this is that the horizon problem should not be a problem in practice because there are lots of paths that meet a set of constraints. Another is that searches for a global motion path can be complicated, because of the number of paths available.

Global Search Methods: Arikan and Forsyth search for complete motion paths that meet given constraints [18]. Such searches are intrinsically off-line so one must sacrifice the goal of interaction, but if the search is fast enough it can be used for authoring animations. Motion ambiguity means that simply applying Dijkstra's algorithm doesn't work, because the algorithm must manage too many intermediate paths. Arikan and Forsyth use a variant of the motion graph where each clip of observed motion is a node, and edges represent acceptable cuts. This means that edges need to be tagged with “from” and “to” frames within the node, and that there are typically multiple self-edges and multiple edges between any pair of nodes. They produce a sequence of compressed version of this graph by clustering edges, so that a pool of edges with similar “from” and “to” frames can be replaced by a single edge with approximating “from” and “to” frames in the more heavily clustered version. They then use a randomized search to find a pool of paths in the most heavily compressed version of the graph; these paths are either refined locally to pro-

duce paths in less heavily compressed graphs, or modified. The best resulting path is then reported. They report a trick that can be used to make synthesized motion paths look as though actors are interacting. One obtains measurements of an interaction, then uses frame constraints to construct paths into and out of the interaction.

Low entropy: Human motion appears to be quite predictable in the sense that one can predict the frame that will occur a short while in the future rather well using the current frame — we use the term **low temporal entropy** to refer to this property. This is in tension with what we have seen already (that motion constraints are ambiguous, and that it is generally fairly easy to move between any two frames in the motion graph quite quickly). We discuss this point in much greater detail in section 7.1.3. This entropy property allows useful approximations for search algorithms.

Annotation based synthesis: One method to control motion ambiguity is to require the synthesis process to produce motions that meet annotation constraints (described in section 5.1.5). Arikan *et al.* use demands that either require the annotation to be present, to be absent, or are “don’t care” [17]. The annotations are painted on the timeline. Frames in the motion graph carry annotations, and we must produce a path that meets position and frame constraints, and carries the required annotation at the required time. For the moment, assume that the only geometric constraint is on the start point. Then building a path that meets annotation constraints is a matter of dynamic programming (there are local costs for failing to meet annotation demands, and frame-frame costs for continuity). The dynamic programming problem is too hard to solve in that form, because there are too many frames of motion. Instead, Arikan *et al.* coarsely quantize the graph into blocks of frames that form sequences and then use dynamic programming on a random subset of these blocks. There are then two search activities: refining blocks, and changing the (randomly chosen) working set of blocks. This works well, because ambiguity means that one doesn’t miss much structure by random sampling and low entropy means that a quantized path represents the actual solution quite well.

5.3.3 How good is a Motion Graph?

Methods of producing motion graphs are hard to assess, because it is quite difficult to tell whether a motion graph is good or bad. Reasonable criteria

include: that there be “few” bad paths (which may not be the same as having “few” bad links); that most paths are acceptable; that the diameter of the graph is small; that almost any spatial path can be synthesized; that there are only short sequences joined at few, well connected choice points (an extremely useful property; see [158]). There is little detailed work on this topic — among other things, the criteria above are mutually contradictory and it isn’t clear how to build algorithms to that do well on some of them.

Reitsma and Pollard have shown how to determine how well a motion graph makes goals in an environment reachable [358]. They discretize the state space (environment and rotation of the figure on the plane), then build a graph on the nodes by recording which node is reached by leaving each node in the discretized state space using each clip in the motion graph. Links that pass through obstacles can be pruned. By building a strongly connected component of this graph, one can count how many states in the environment are reachable with the current motion graph. Specific problems can then be identified: for example, a shortage of stopping and turning motions in Reitsma and Pollard’s motion graph made it difficult to get their character into tight spaces.

5.4 Motion Primitives

Our very rough model of the space of motions above doesn’t really take long time structure of motions into account. Such structure is evident in how people move on a daily basis. One can walk backward for long distances, but one doesn’t; one can intersperse; for that matter, some can walk on their hands, but few do for long periods. This sort of structure needs to be thought of in terms that are probabilistic, rather than deterministic (because the semantics are that one could but one tends not to).

A natural method for building models of motion on these time scales is to identify clusters of motion of the same type and then consider the statistics of how these **motion primitives** are strung together. There are pragmatic advantages to this approach: we can avoid blending between motions that are obviously different; we can model and account for long term temporal structure in motion; and we may be able to compress our representation of motion with the right choice of primitive model. Finally, a primitive based representation has some advantages for recognition, and Feng and Perona describe a

method that first matches motor primitives at short timescales, then identifies the activity by temporal relations between primitives [126]. In animation, the idea dates at least to the work of Rose *et al.*, who describe motion **verbs** — our primitives — and **adverbs** — parameters that can be supplied to choose a particular instance from a scattered data interpolate [368]. The verbs appear to be chosen by hand; within a particular primitive, motions are aligned (c.f. section 5.2.2) and then a scattered data interpolate produces an instance. There is a **verb graph** which gives the combinatorial structure of how verbs can be joined up.

5.4.1 Primitives by Segmenting and Clustering

Primitives are sometimes called **movemes**. Mataric *et al.* represent motor primitives with force fields used to drive controllers for joint torque on a rigid-body model of the upper body [276, 277]. These force fields have a stationary point at a desired hand configuration; different force fields can be superposed to obtain different endpoints. The primitives appear to be chosen by hand. The motions are 3D motion captured arm movement; segment boundaries are obtained by looking for points where the sum of squares of velocity at all joints is small. Del Vecchio *et al.* define primitives by considering all possible motions generated by a parametric family of linear time-invariant systems; if a split of the parameter space results in two sets of motions that are always distinct, that split can be used to derive primitives [449]. The definition of the primitives results in a segmentation algorithm, and authors show that reaching and drawing motions can be distinguished in this framework.

There is quite a lot of evidence that motions segment and cluster well — meaning that one can use various segmentation and clustering processes as intermediate steps in motion synthesis, without serious difficulties resulting. This is not something one would expect, given the dimension of most motion representations. Barbič *et al.* compare three motion segmenters, each using a purely kinematic representation of motion [28]. Their method moves along a sequence of frames adding frames to the pool, computing a representation of the pool using the first k principal components, and looking for sharp increases in the residual error of this representation. Their Gaussian mixture model segmenter regards frames as IID samples from a Gaussian mixture model, then computes the mixture component from which a frame arises.

Their probabilistic PCA segmenter works like the PCA segmenter, but obtains a normal probability density from the principal component analysis and then compute the Mahalanobis of new frames from the mean of this model; this segmenter appears to be the best of the three. While there is no agreed way to evaluate a motion segmentation, Barbič *et al.* report segmentations that look good. For our purposes, the most significant point here is that distinct movements tend to be dramatically distinct — one doesn't need to look at fine details of dynamics to segment such motions as “*walk*”, “*stand*”, “*sit down*” and “*run*”.

Dimension Reduction: It is natural to expect that any primitive structure in motions could be exposed by reducing the dimension of the data. Furthermore, dimension reduction methods could yield a conveniently compressed encoding of a motion primitive. Fod *et al.* construct primitives by segmenting motions at points of low total velocity, then subjecting the segments to principal component analysis and clustering [128]. Jenkins and Mataric segment motions using kinematic considerations, then use a variant of Isomap (detailed in [206]) that incorporates temporal information by reducing distances between frames that have similar temporal neighbours to obtain an embedding for kinematic variables [207]. They cluster in the resulting space to obtain motion primitives over short temporal scales, then apply isomap again to obtain primitives on longer temporal scales; they report plausible motions.

There is other evidence that relatively few measurements can yield the kinematic configuration of the body — that is, that a low dimensional representation of configuration applies. Chai and Hodgins demonstrate a form of **video puppetry** — where an animated figure is controlled by observations of an actor — using relatively few markers; this approach most likely works because motions tend to be confined to a low dimensional subspace [74]. Safonova *et al.* are able to produce plausible figure animations using optimization techniques confined to a low-dimensional space (see [378], figure 5.3 and section 5.5.2.3).

5.4.1.1 Difficulties with Dimension Reduction

Dimension reduction methods can be subject to the same problems that occur with blending methods. It is hard to ensure that all sequences used in building a model are time aligned sufficiently precisely that the high-frequency struc-

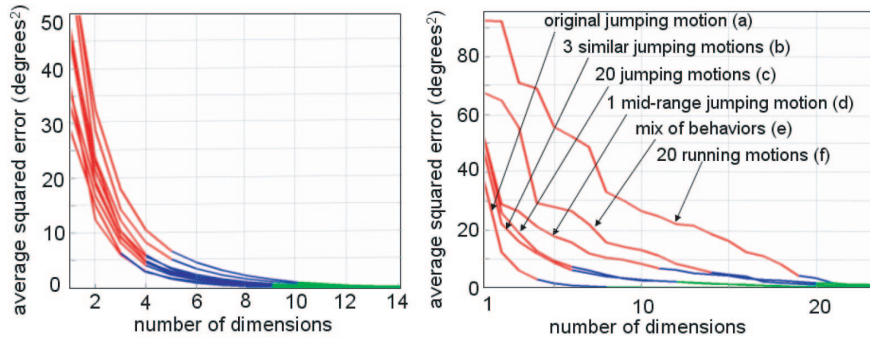


Fig. 5.3 **dimensionreduction/lowoptimization, figures 2 and 3** Safonova et al. synthesize motions in low-dimensional spaces, constructed by taking a fixed number of principal components of static frames [378]. Their work contains extensive evidence that low-dimensional representations of motion are useful, and appear to fit data well. On the **left**, a graph comparing the error (RMS in angles) between the original motion data and a projection onto this low-dimensional space, for different numbers of principal components, and averaged over between ten and twenty motions. The curves are coded with red (for a reconstruction that is not acceptable visually), blue (for a reconstruction that has minor visual artifacts) and green (for a good reconstruction). Motions are of several types, including: running, walking, jumping, climbing, stretching, boxing, drinking, playing football, lifting objects, sitting down and getting up. Each type of motion is encoded with a different set of principal components. The method appears to display quite good generalization, as the graph on the **right** suggests. This shows RMS error in joint angle for reconstructions with different numbers of principal components for a jumping motion, with the basis estimated on: (a) the frames being reconstructed (which, not unnaturally, gives the best result); (b) a set of three similar jumping motions; (c) a set of 20 jumping motions; (d) a single jumping motion; (e) a mix of behaviours and (f) 20 running motions. The color coding is the same as for the graph on the left. Notice that, while the basis chosen clearly should depend on the behaviour (because (f) yields a poor basis), once one has that accounted for, a basis chosen on a different instance still gives quite a good reconstruction with a relatively low dimension — the generalization is quite good.

ture associated with fast, definite movements doesn't average out. Squashy-looking motions can result, as can footskate. It is most likely that one should separate out these components and then synthesize them independently once the overall structure of the motion has been established.

5.4.2 Linking Segmentation to the Primitive Model

Segmentation and encoding should interact — we can reasonably expect a good segmentation results in good primitives, but the other way works, too; if one has a good representation of each particular primitive, that could drive segmentation. This is now a commonplace in the machine learning community. Li *et al.* segment and model motions simultaneously using a linear dy-

namical system model of each separate primitive and a Markov model to string the primitives together by specifying the likelihood of encountering a primitive given the previous primitive [253]. For the moment, assume the segmentation is known and we wish to identify a primitive from some set of observations that have been determined to come from that primitive. We assume that each primitive consists of a sequence of observations \mathbf{Y}_t , each generated by a hidden state \mathbf{x}_t . We would like the system to have second order dynamics so that the model takes accelerations into account; this is equivalent to assuming that x_t is a linear function of x_{t-1} and x_{t-2} . We can obtain a Markovian model by stacking two state vectors to obtain $\mathbf{X}_t = [\mathbf{x}_t, \mathbf{x}_{t-1}]^T$. The model of each primitive now takes the form

$$\begin{aligned}\mathbf{X}_t &= \mathcal{A}_t \mathbf{X}_{t-1} + \mathbf{V}_t \\ \mathbf{Y}_t &= \mathcal{B}_t \mathbf{X}_t + \mathbf{W}_t\end{aligned}$$

where \mathbf{V}_t and \mathbf{W}_t are normal random variables with known mean and variance. Notice that \mathcal{A}_t will have the form

$$\begin{pmatrix} \mathcal{U}_t \mathcal{U}_{t-1} \mathcal{I} & 0 \end{pmatrix}$$

(so that one has the right behaviour from the stacked components of the state vector). You should compare this model to the HMM's used for tracking; we have the same model, but now we wish to obtain the values of \mathcal{A} and \mathcal{B} from observations of \mathbf{Y}_t , rather than estimate the states. The difficulty here is that the model is not uniquely specified in this form. For example, assume that \mathcal{C}_t is a sequence of matrices of full rank, then the state sequence $\hat{\mathbf{X}}_t = \mathcal{C}_t \mathbf{X}_t$ taken with matrices $\mathcal{A}_t \mathcal{C}_t^{-1}$ and $\mathcal{B}_t \mathcal{C}_t^{-1}$, has the same likelihood. Li *et al* deal with this by insisting that the states be the projection of the observations on to a subset of the principal components of the observations, and can then estimate \mathcal{A}_t and \mathcal{B}_t with maximum likelihood.

Of course, the segmentation is not known. We will estimate the segmentation and primitives together with an iterative procedure: fix the primitives, estimate the best segmentation; now re-estimate the primitives with that segmentation; etc. This mirrors EM, but one is now using the maximum likelihood segmentation conditioned on the primitive parameters as an estimate of the expected segmentation conditioned on the parameters. The segmentation can be obtained with dynamic programming (Li *et al*. assume that each primitive emits at least 60 frames, which complicates the representation only very

slightly). To see that the best segmentation of some sequence of length N into M primitives of length no shorter than L is available using dynamic programming, we build a graph whose nodes consist of statements that frames i to $i + k$ of the sequence were produced by primitive j ; there can be no more than N^2M such nodes. Each node is labelled with the negative log-likelihood of the relevant sequence under the relevant dynamical model. There is a directed edge from each node to any node that can succeed it, labelled with the negative log-likelihood that the one primitive follows the other under the Markov model. We now obtain the minimum value path through this (acyclic, directed) graph using dynamic programming.

The resulting model can be used generatively to produce new motions. Li *et al.* obtain their best results by specifying the body configuration at each change of primitive — so that the model interpolates between these frames. This avoids phenomena like drift (which must occur because of the random noise component) causing minor but annoying effects like the feet floating above or below the ground.

5.5 Enriching a Motion Collection

All the methods we have discussed involve “small” changes to existing motions to obtain new motions that are basically similar. The ideal is to have methods that can produce completely new, and good, motions from constraints and, perhaps, data. Approaches to this area involve reasoning about the fundamental considerations that produce motion (as opposed to processes for synthesizing motion to meet immediate needs). There are two types of method: methods that attempt to obtain new motions by “large” operations on existing examples, and methods that use physical and variational criteria to produce novel motions.

5.5.1 Rearranging Existing Motions

Human motions quite clearly have some properties allowing composition over the body and over time. These properties are a formidable source of complexity of a form that will defeat naive data-driven methods — for example, to synthesize an actor walking while scratching with the left hand, do we really need to see this particular action? does this mean we need to see walking while scratching with the right hand to synthesize that, too? must we

observe scratching different locations with each hand, too?

5.5.1.1 Motion Editing

Gleicher shows that one can usefully edit motions — typically, so that they meet constraints that are a small revision of constraints met by the original motion — by adding a displacement [153]. Gleicher minimizes a measure of the size of the displacement subject to the new constraints. There is no guarantee that the resulting motion will necessarily look human, but for small displacements it tends to; this means that the motion author can manage constraints and update process so that the resulting motion looks human. The optimization problem is nasty. Lee and Shin obtain a more manageable optimization problem by representing the motion as a hierarchical B-spline [244]. The displacement is also a hierarchical B-spline, and they engage in a coarse-to-fine search across the hierarchy. The IK solver at the k 'th frame at the n 'th level now has the $k-1$ 'th frame at that level and all frames at the $n-1$ 'th level available to generate a start point and to constrain the solution. Witkin and Popović modify motions using parametric warps, so that they pass through keyframes specified by an animator [470]. Shin *et al.* use similar methods to touchup motion to meet physical constraints (for example, motion not in contact is ballistic and preserves angular momentum), while sacrificing physical rigor in the formulation for speed [391]; see also [431, 432]). Motion editing in this way is useful, and there are several other systems; a review appears in [156].

5.5.1.2 New Motions by Cut and Paste

Simple methods can produce good results for some composition across the body, but not for all cases. Ikemoto and Forsyth build new motions from old by cutting arms or upper bodies off one motion and attaching them to another [194]. Pairs of motions are selected by several different randomized proposal mechanisms, components transplanted between them, and the two results then presented to a classifier which attempts to tag sequences that do not look human. The classifier is quite reliable when presented with motions that are reasonably similar to examples, but tends to be less reliable when presented with dramatically different motions; this is a difficulty, because the

whole point of understanding composition is to synthesize good motions that are dramatically different from examples.

What is important here is that the classifier is necessary; many such transplants are successful, but some apparently innocuous transplants generate motions that are extremely bad. It is difficult to be precise about the source of difficulty, but at least one kind of problem appears to result from passive reactions. For example, assume the actor punches his left arm in the air very hard; then there is typically a small transient wiggle in the right arm. If one transplants the right arm to another sequence where there is no such punch, the resulting sequence often looks very bad, with the right arm apparently the culprit. One might speculate that humans can identify movements that both don't look like as though they have been commanded by the central nervous system and can't be explained as a passive phenomenon.

5.5.1.3 Motion Fill-in by Nonparametric Regression

The idea that motion of one part of the body leaves a signature in the motion of other parts of the body is confirmed by work of Pullen and Bregler [351], who built a motion synthesis system that allows animators to sketch part of the motion of the body, and then uses a non-parametric regression method to fill in the details. Joint angle signals are segmented at local extrema. The segments are represented at multiple temporal scales. Animators can then sketch part of a motion — for example, hip and knee angles at a coarse temporal scale — and the system then obtains fragments of joint angle for the other joints and other scales. These are found by matching the fragments of sketched motion to a motion capture dataset (allowing a degree of scaling in both time and angle in the matching process). Typically, there are multiple matches for each fragment. The set of resulting fragments is searched to produce signals that tend to have as many **consecutive fragments** — fragments that succeed one another in the observed data — as possible. These signals may not be continuous (and usually are not, unless the fragments are consecutive), so discontinuous joins are smoothed using a blending technique. Multiple motions can result from this process, and it is up to the animator to choose the best.

The method produces rather good motions, using examples and motion demands from the same “type” of activity. Conditioning on the kind of mo-

tion appears to be important — one couldn't reasonably expect that it would be possible to synthesize good football motions from observations of dance — but it is difficult to be precise about what one is actually conditioning on. The fact the method works can be used as evidence in support of the idea that motions have some form of structure that takes in the whole of the body. It is probably unwise to use this view to argue against a compositional representation of motion, because the experiments in the paper don't establish that there is only one possible path for, say, the upper body given a particular set of lower body motions.

5.5.1.4 Motion Interpolation

In **motion interpolation**, one attempts to produce motions that interpolate between, or extrapolate from, existing motion-capture measurements. A natural procedure is to produce a controller that can track the measurements and then, when measurements are no longer available, produce motions by controlling some body parameters. A variety of approaches that make use of physical simulation have been developed along these lines. Controllers that track motion data provide a useful mechanism for smoothing recorded errors while also adjusting for disturbances not present in the recorded motion [120, 348, 488, 489]. Other approaches make use of hand designed or optimized controllers that operate independently from recorded motion [118, 119, 165, 180, 350]. Building controllers that generate human-like motion remains an open research problem.

5.5.2 Motion from Physical Considerations

The motion editing methods we have seen do not require that deformed motions be physical. In fact, these methods are simplifications that originate in a body of research to generate human motion from considerations of physical constraint and energy. This work originates with Witkin and Kass, who introduced the use of **variational methods**, widely known as **spacetime constraints** [469].

We have a jointed figure, whose configuration can be represented by some set of parameters \mathbf{q} . These coordinates can be **reduced coordinates**, where any set of values represents a legal configuration of the figure — these could be, for example, root coordinates and joint angles. An alternative is to use

generalized coordinates, where not every choice of values represents a legal configuration of the figure — these could be, for example, the pose of each separate limb segment; in this case we need constraints to ensure that the limbs don't fly apart. The configuration of this figure is subject to some constraints. For example, a figure that is sliding on the floor will be constrained to have each foot on the floor. This figure is subjected to a set of forces and torques \mathbf{f} . Assume the figure is moving for the time interval I . From mechanics, the motion of this figure achieves an extremal value of the time integral of the **Lagrangian** (see, for example [2, 20, 161]). We write the Lagrangian as $L(\mathbf{f}(t), \mathbf{q}(t), \lambda, t)$, where λ are the **Lagrange multipliers** (which can be interpreted as the coefficients of generalized workless constraint forces that ensure the motion meets the constraints). Some constraints are **dynamical constraints** (which refer to forces, torques, momenta and the like); we shall write this set as $\mathbf{D}_e(\mathbf{f}, \mathbf{q}, \lambda, t) = 0$ and $\mathbf{D}_i(\mathbf{f}, \mathbf{q}, \lambda, t) \leq 0$. Others are **kinematic constraints** (which constrain configuration); we shall write this set of constraints as $\mathbf{K}_e(\mathbf{q}, t) = 0$ and $\mathbf{K}_i(\mathbf{q}, t) \leq 0$.

Let us confine our attention to an interval where we know which kinematic constraints are active (i.e. which components of \mathbf{K}_i are equal to 0), and write the set of active kinematic constraints including all the equality constraints as $\mathbf{P}(\mathbf{q}, t) = 0$. Write the remaining set of kinematic inequality constraints as $\mathbf{P}_i(\mathbf{q}, t) < 0$. Any physical motion extremizes the Lagrangian subject to these constraints, and, from variational calculus, we obtain the **Euler-Lagrange equations**, which are differential equations satisfied by any motion that does extremize the Lagrangian. We adopt the notation where differentiating by a vector results in a vector of derivatives with respect to each component. Write the Euler-Lagrange equations as

$$\mathbf{E}(\mathbf{f}, \mathbf{q}, \lambda, t) = 0 = \begin{pmatrix} \frac{d(\frac{\partial L}{\partial \dot{\mathbf{q}}})}{dt} - \frac{\partial L}{\partial \mathbf{q}} - \lambda^T \frac{\partial P}{\partial \mathbf{q}} - \mathbf{f} \\ \mathbf{P}(\mathbf{q}, t) \end{pmatrix}$$

Notice that we now have algebraic equations that constrain derivatives. Equations of this form are known as **differential-algebraic equations**; they have a (well-deserved) reputation for creating nasty numerical problems (a fair place to start is [172, 173]).

Now we wish to choose a motion that meets the dynamical constraints, and where some other criterion — which might measure, for example, work

— is extremised. Write this criterion as

$$\int G(\mathbf{q}, \mathbf{f}, \lambda, t) dt$$

The problem becomes

$$\text{Maximize } \int G(\mathbf{q}, \mathbf{f}, \lambda, t) dt$$

Subject to:

$$\mathbf{E}(\mathbf{f}, \mathbf{q}, \lambda, t) = 0$$

$$\mathbf{D}_e(\mathbf{f}, \mathbf{q}, \lambda, t) = 0$$

$$\mathbf{D}_i(\mathbf{f}, \mathbf{q}, \lambda, t) \leq 0$$

$$\mathbf{P}_i(\mathbf{q}, t) < 0$$

Witkin and Kass did not use the idea to generate human motions, but demonstrated very attractive animations of a bouncing lamp produced using this method. There are very serious practical difficulties in producing animations of human motion like this. The actual minimization process might be extremely difficult. In fact, there is no prospect of getting a useful result by simply dropping this problem into a commercial optimization package. The state space has complex geometry caused by the internal degrees of freedom, joint limits and the like. Contact and frame constraints can produce unpleasant feasible sets, and one should expect the problem not to be convex. One must encode the function $\mathbf{x}(t)$ with some finite dimensional parameter space, and the choice of encoding may create difficulties; for example, contact constraints tend to produce quite high frequency terms in the motion signal (or, equivalently but rather easier to observe, smoothing the motion signal tends to lead to footskate). There is some reason to believe that a coarse-to-fine representation is useful [265]. One may simplify optimization difficulties by choosing simplified characters (e.g. [120, 348, 350, 441]; freefall diving is a particular interest [86, 264]) or by exploiting interaction with an animator (e.g. [83]). Ngo and Marks produce motions for quite complex characters using spacetime optimization by building motions out of **stimulus-response** pairs — parametric packets of motion that are triggered by some parametric test ([310, 311]; see also [279] for other motions built out of packets). The precise set of packets, and the parameters of those packets, are chosen using search by a genetic algorithm (see also the work of Sims [398]). There is no claim that these motions necessarily appear human.

The choice of objective function can affect the resulting motion and is by no manner of means obvious. It is occasionally asserted that human motion should minimize some choice of mechanical energy. One should place little weight on this idea for most motions because there are too many other important considerations that shape how we move. For example, Wu and Popović need a specially crafted objective function that allows for the enormous energy expenditure required at takeoff to obtain convincing bird flights [81]. As another example, the energy saved by using a slow reaching motion might be far outweighed by that lost by getting to the target fruit too late. For that matter, even more energy could be saved by not moving at all; but at some cost. Liu *et al.* show a method to obtain simulation parameters from examples [258].

For these reasons, spacetime optimization has not to our knowledge been used to generate complete human motions over long periods. Rose *et al.*, who generate **motion transitions** — short sequences of motion that join specified frames “naturally” — using an optimization procedure that minimizes the total squared torque moving the upper body [369]. The legs are controlled kinematically, using either manual or automatically supplied constraints for footplants. Anderson and Pandy describe a simulation of one step of a walk for a highly detailed dynamic model that produces (using months of supercomputer time) a pattern of muscle activations that minimize an effort criterion and also look like human muscle activation patterns ([15]; see also [323]). Spacetime optimization has, however, been of tremendous value in deforming existing motions.

5.5.2.1 Simplified Characters

Popović and Witkin use characters with simplified kinematics, and model muscle forces explicitly (the muscle is modelled as a **proportional-derivative controller** attempting to drive a degree of freedom to a set-point) [350]. Their method produces physically plausible motions that meet constraints and are close to observations. They represent major features of motion using **handles** — vector functions of configuration, typically a map onto some lower dimensional space, the details of which vary between applications. For example, if one wished to ensure a motion preserved contact, appropriate handles might be the position of points on the figure. A space-

time optimization is used to fit the simplified model to observed motion data, resulting in handles $\mathbf{h}_s(\mathbf{q}_s)$; a second spacetime optimization produces a simplified model that meets the constraints with handles $\mathbf{h}_t(\mathbf{q}_t)$; and the handles for the observed data are $\mathbf{h}_o(\mathbf{q}_o)$. They now seek to produce a final motion \mathbf{q}_f with handles $\mathbf{h}_f(\mathbf{q}_f) = \mathbf{h}_o(\mathbf{q}_o) + (\mathbf{h}_t(\mathbf{q}_t) - \mathbf{h}_s(\mathbf{q}_s))$ (that is, displace the handles of the original motion with a displacement computed from the simplified figure).

They do this by optimizing an objective function that penalizes **mass displacement**, which is computed as a sum of squared magnitudes of differences in positions between corresponding sample points on the final motion and the observed motion, weighted by the mass at that sample point. As a result, degrees of freedom in the final animation that are not constrained by the handles are derived from the original motion. The optimization is constrained by the requirement on the handles (above) and physical constraints on the motion. The parameters are configuration and muscle demands. The spacetime method appears to benefit considerably from the relatively few degrees of freedom in the simplified character and the presence of an initial point (the observed motion).

5.5.2.2 Modified Physics

Liu and Popović produce character animations from rough initial sketches using an optimization method by breaking the motion into phases, simplifying the physical constraints, and, where necessary, exploiting the animator's input [259]. They then identify transitions — where the figure moves from one set of constraints applying to another — and require the animator to provide frames for these transitions, which tend to be a particular source of difficulty for optimization methods. They must now produce a series of motion clips to fill in between these transitions. There are two important cases: **ballistic motion**, where there is no contact — the body is in flight, as in jumping, diving, etc. — and **constrained motion**, where there is some contact. In ballistic motion, if we use reduced coordinates, then all external forces are due to gravity (so the acceleration of the center of mass is \mathbf{g}) and angular momentum is conserved. Constrained motions are required to have a momentum curve of a particular form (figure 5.4), which is consistent with biomechanical observations.




Fig. 5.4 **Physicssmoothing/paper1.pdf, figure 6** The angular momentum curve for a whole motion for the method of Liu and Popović [259], showing total angular momentum as a function of time. The motion before p_1 and after p_4 is ballistic, so the total angular momentum is a constant. The form of the momentum curve is taken from biomechanical models [324, 229]. The form is imposed by smoothly interpolating p_1 , p_2 , p_3 and p_4 , requiring that $p_2 < p_1$, $d_2 < d_1$ and $(p_2 - p_4)(p_3 - p_4) < 0$. Figure 5.5 shows a motion obtained using this method.

The objective function is a sum of three terms: a measure of mass displacement ; a measure of **coordinate velocity**, which penalizes large changes in the degrees of freedom to enforce frame-frame coherence; and a measure of **static balance**, which penalizes large distances between the center of mass and the location of point constraints. The objective function and the constraints are functions of $\mathbf{q}(t)$ (and its derivatives) and the control points for the momentum curve. The method does not constrain forces or torques at joint, and they do not participate in the objective function, which means that they can be ignored (this doesn't mean the motion isn't physical; it means that we assume that the body will supply whatever internal forces or torques are required to follow the motion path). Abe *et al.* drop the mass displacement and coordinate velocity terms in favour of a similarity term, and use a variety of different momentum profiles to produce further variations on motion




Fig. 5.5 **Physicssmoothing/paper1.pdf, figure 6** **Top:** a motion demand supplied by an animator and **bottom** a motion synthesized using the procedure of Liu and Popović [259]. The motion is obtained by (a) inferring constraints from the demand; (b) extracting transitions from an animator; and then (c) computing a set of clips that meet these transitions and the inferred constraints, have angular momentum curves of the form of figure 5.4 and extremize an objective function that penalizes mass-displacement, coordinate velocities, and out-of-balance configurations.

capture data [1].

There is a real advantage to not constraining forces and torques and not allowing them to participate in the objective function: one does not need to compute them. This means that computing various Jacobians that arise in the optimization procedure can be made linear (rather than quadratic) in the number of degrees of freedom, as Fang and Pollard show [121].

5.5.2.3 Reduced Dimensions

Safonova *et al.* describe a method for synthesizing motions from variational considerations using a dimension reduced representation of configuration [378]. For each “type” of motion (for example, running, walking, jumping, climbing, stretching, boxing, drinking, playing football, lifting objects, sitting down and getting up), Safonova *et al.* construct a basis of principal components for the frames from that sequence. New motions are now represented using coefficients on this basis. Motions are obtained by optimizing a sum of three terms: the first, the integral of summed squared torques, penalizes effort; the second penalizes, the integral of summed squared velocities and accelerations, penalizes high-frequency wobbles; and the third, the summed Mahalanobis distance of coefficients from the mean, penalizes frames that are strongly unlike examples. This optimization problem is considerably simplified, because effort is focused on a small set of dimensions that are clearly significant and independent. Motions are specified with ini-

tial, final and key frame constraints; time, contact and pose constraints are also possible. The method imposes torque limits. The method produces good motions from relatively limited constraints. To obtain a motion, the “type” of the motion required must be known, and sufficient constraints must be provided, so the method is most useful in a situation where an animator can interact with the synthesis procedure.

5.5.2.4 Modifying Existing Motions

Hodgins and Pollard describe scaling rules that allow a motion that applies to one character to be transferred to another character, using methods of dimensional analysis ([179]; for dimensional analysis, see [30]). Sulejmanpašić and Popović modify existing motions to obtain revised motions that meet animator demands using a full dynamical model ([425]; see also [349], which describes a search method to obtain parameters of a rigid body simulation that is similar to a sketch). The method produces physical motions; each step of the iteration computes an update direction for positions, torques, reaction forces, etc. that is the smallest update to meet a linearized version of the demand. There is then a line search along the chosen direction to obtain an update that gives the smallest constraint error. Authors demonstrate that a poor choice of scaling for the variables significantly complicates obtaining a solution, and describe an experimental procedure for choosing a scaling. The method produces good motions efficiently, if the demand is not too far from the original motion.

6

Describing Activities

Understanding what people are doing is one of the great unsolved problems of computer vision. A fair solution opens tremendous application possibilities, including: improved surveillance systems; a better understanding of what people do in public; better architectural design; and better human computer interfaces. While there has been extensive study of this topic, it still isn't terribly well understood. One can obtain statistics of some behaviours from coarse scale tracks (e.g. for car parks, see [423]; for architectural domains, see [478]). But understanding activities that depend on detailed information about the body is still hard. We contend that the major difficulties have been (a) that good kinematic tracking is hard; (b) that models typically have too many parameters to be learned directly from data; and (c) for much everyday behaviour, there isn't a clear taxonomy into which to classify observations.

There is a long tradition of research on interpreting activities in the vision community (see, for example, the extensive survey in [189]). There are three major threads. First, one can use temporal logics to represent crucial order relations between states that constrain activities. Second, one can use spatio-temporal templates to identify instances of activities. Third, one can use (typically, hidden Markov) models of dynamics.

6.1 What should an Activity Representation do?

There appear to be a series of quite different cases in activity recognition. First, we distinguish between short, medium and long timescales. Second, we distinguish between motions that can be sustained (walking, running, waving) and motions that have a localizable character (catch, throw, punch, kick). Since we want our complex, composite motions to share a vocabulary of base units, we use the kinematic configuration of the body, limb velocities, and perhaps accelerations as distinctive features at **short** timescales — which might be of the order of a small number of frames. We define **acts** to be frame labels that can be decided on such very short timescale features — such labels, (for example, walk-right-leg-stance-left-leg-swing) tend not to have directly useful semantics.

At **medium** timescales, we have **activities** — motions like walking, running, jumping, standing, waving, whose temporal extent can be short (but may be long); such motions are typically composites of multiple acts. Furthermore, activities can be sustained for long periods. We use the term **actions** for motions that have a localizable character and require medium timescales to identify. Both actions and activities may be difficult to identify with only a few frames but are relatively easy to identify from hundreds to thousands of frames. Both actions and activities allow a degree of composition — for example, one could walk and scratch at the same time.

One's interpretation of a view of moving humans is strongly affected by objects nearby. For example, a person standing in an isolated field may be behaving strangely; the same person in the same configuration next to a bus stop is waiting for a bus. We believe that the most natural level at which to start inserting considerations of context into activity recognition is that of activities — where one can pool object detector responses over a long enough sequence of frames to expect quite good behaviour — and define the next layer of the representation to be **motions in context**. Context applies to both activities and actions. These occur at medium timescales, but the nature of the motion in context is determined by both the actions in the sequence and the response of object detectors. We use the term **behaviour** to cover motions at **long** timescales — typically, behaviours such as fighting, exercising or visiting an ATM might be composed of a selection of different motions in context, linked up by activities, and organized in a variety of possible ways and

meeting a variety of constraints on temporal ordering. It has been recognized for some time that there are other helpful distinctions (e.g. Bobick [46] distinguishes between movements, activity and actions, corresponding to longer timescales and increasing complexity of representation; some variants are described in two useful review papers [7, 146]).

6.1.1 Necessary Properties of an Activity Representation

The big goal is a theory and mechanism for recognizing a wide range of behaviours. There are some important constraints on solutions to this problem. First, we expect that typical behaviours are a composite of many activities, and this composite is not unique — the same behaviour may be represented by multiple sequences of actions, as long as these sequences observe an internal structure. For example, one may scratch or groom at any time while visiting an ATM, but one must type a PIN before retrieving money, and insert a card before typing a PIN (notice that one can't retrieve a card before inserting it, but at some machines one might retrieve the card before typing a PIN; at others, the card is retrieved after typing the PIN and before recovering money). Each activity may itself be one of several different composites of multiple actions, in the same way. Each action might also have compositional properties — for example, one may walk with three-quarters of one's body while scratching with the fourth limb. The modelling strategy must respect both this hierarchical structure and the compositional nature of motion.

Second, we expect that there is not labelled data for each possible case; we cannot simply learn models without any human interaction. This applies to models of actions, activities and behaviours. This difficulty is created by the compositional nature of human motion; the sheer richness of available motions defeats pure data-driven strategies. An important criterion for choosing a modelling strategy is that it be easy for humans to author and to assess rich models quickly. Such models should be amenable to parameter learning from data, but it should not be necessary to see an example of every possible instance of a behaviour to build a model.

Third, we expect that the supervised data that is available may be marked up somewhat inaccurately. Typically, a behaviour will be marked up with activity names (an activity with actions, respectively), but the boundaries of the markup are unlikely to be accurate. We expect the learning algorithm to

be robust to some segmentation noise.

Finally, we models should have the property that basic activities with the model — model building, composition, and inference — is relatively straightforward. In this, we follow the experience of the statistical natural language community, that trading expressiveness in models for simplicity of authoring and inference is often advantageous.

6.1.2 What Data is Available?

An important part of design here is to keep into account what kinds of data are easy to obtain and what difficult, so as to plan model authoring around what is practical. Experience suggests that it is possible to get from minutes to hours of reasonable quality motion capture data; relatively few minutes of video labelled as to actions (these labels are very difficult to produce because they require frame accuracy); minutes to hours of video labelled in reasonable detail with respect to activities and behaviours, accepting poor temporal resolution in the labels; and of the order of months of public observation video. It is relatively straightforward to look at large volumes of labelled motion capture data and correct labels, not least because one can observe many frames simultaneously (e.g. see [17]).

One important source of difficulty is that it is hard to tell which aspects of behaviour should be modelled accurately in order to perform useful tasks. Resolving this requires (a) study of ideas in sufficient generality that they transfer between tasks and (b) some example tasks. But the selection of example tasks is not innocuous. In particular, a distinctive feature of everyday activity is the number of behaviours that appear familiar, but for which the observer may not know a word or even a compact description. In contrast, in some domains (e.g. ballet [69]; gymnastics [116]; tai chi [54, 59]; tennis [427]; walking [71]) there are quite specific vocabularies that refer to very precisely delineated behaviours. This is an advantage for building demonstration systems, because one can evaluate them, but may avoid the real difficulty, which is that for most activities we want to classify the activity without knowing a precise or canonical set of classes.

6.2 Miscellaneous Methods

6.2.1 Activity Representation Methods based around Temporal Logics

Pinhanez and Bobick [336, 337] describe a method for detecting what we have called behaviours using a representation derived from Allen's interval algebra [14], a method for representing temporal relations between a set of intervals. One authors a description of the behaviour in terms of primitives, which are indivisible and occupy temporal intervals. The description incorporates a set of legal relations between the primitive intervals; a description is consistent if at least one set of intervals, together with an allocation of those intervals to primitives, satisfies it. One determines whether an event is past, now or future by solving a consistent labelling problem, allowing temporal propagation. There is no dynamical model — sets of intervals produced by processes with quite different dynamics could be a consistent labelling; this can be an advantage at the behaviour level, but probably is a source of difficulties at the action/activity level. These papers do not show the method applied to noisy detectors; there are results using simulated detectors on real data.

Siskind [400, 399] describes methods to infer activities related to objects — such as throw, pick up, carry, and so on — from an event logic formulated around a set of physical primitives — such as translation, support relations, contact relations, and the like — from a representation of video. A combination of spatial and temporal criteria are required to infer both relations and events, using a form of logical inference. The methods are focussed on activity representation, and do not use real video data; there is no mechanism to account for missing or noisy interpretations of video.

6.2.2 Activity Representation Methods based on Templates

The notion that a motion produces a characteristic spatio-temporal pattern dates at least to Polana and Nelson [340, 341, 343, 342, 347]. Spatio-temporal patterns are used to recognize actions in work by Bobick and Davis [47] and Davis and Bobick [100]. Ben-Arie et al [38, 37] recognize actions by first finding and tracking body parts using a form of template matcher and voting on lifted tracks; the tracks are lifted to 3D and a spatio-temporal representation of each body segment votes separately for an action. The action with

the most votes is chosen. The method is successful, and has the advantage that it is robust to composition — if all but the left arm is walking, the action will still be recognized. However, the vocabulary consists of eight items (jump, kneel, pick, put, run, sit, stand, walk) and the vocabulary cannot be composed. An alternative is to match gestural information directly, incorporating a timewarp to improve the match. Bobick and Wilson [48, 49] use a state-based method that encodes gestures as a string of vector-quantized observation segments; this preserves order, but drops dynamical information. The advantage is relatively fast training.

6.3 Activity Representation using Hidden Markov Models and Finite State Representations

6.3.1 The Speech Analogy

Hidden Markov models (HMM's) pervade studies of motion, gesture and activity, and a complete review of their applications here may now be impossible. HMM's are models of sequences, and at their heart is a clock. One has a set of hidden states; at each tick of the clock, a Markov process chooses a new state, dependent on the previous state and nothing else; and an emission process produces an observation from the new state. There are clean solutions for the standard problems of learning (determining an appropriate state transition model and emission model for a given state model) and inference (determine which hidden states occurred given a set of observed states). HMM's have been used for understanding human behaviour but typically with quite small state models.

Very large state models are common in speech recognition, where HMM's have been hugely influential. This area is a useful source of inspiration by analogy. Viewed from a great height, a typical speech system has a series of components: a language model showing how words are built up into sentences; a pronunciation dictionary, giving sequences of context independent phones that correspond to words; a context dependency model, showing how local influences produce context dependent phones (cphones hereafter) from context independent phones; an acoustic observation model showing how acoustic observations result from context dependent phones (this is an extremely compact description of a highly sophisticated area; more extensive descriptions appear in [205, 353]). The resulting object is a vast HMM —

in our example, states can be thought of as being tagged with word-cphone-phone-sample — to explain each sample.

This HMM has some important, attractive features. Learning and authoring can be broken into tractable subproblems — the language model might be learned with one kind of dataset, the pronunciation dictionary with another — and as a result, we obtain an HMM on a massive scale, but with little difficulty in authoring it. While the state space is so big that dynamic programming must be sacrificed for a beam search, the state transition model is not impossible to learn, because most state transitions don't occur. Furthermore, the model is forced to share parameters in important ways — a phoneme in one word has the same model as that phoneme in a different word.

6.3.1.1 Finite State Transducers

Finite state models have had considerable success in the speech and language community. We introduce some terminology here, from the reviews by Mohri and others [294, 293, 292]. A *finite state automaton* is a directed graph, whose nodes are known as states. There is at least one final state and one initial state; each edge is labelled with an element of an alphabet. The automaton accepts any string corresponding to a path from an initial state to a final state. In a *finite state transducer*, transitions are labelled with both an element of an input alphabet and an element of an output alphabet; any string accepted by the transducer results in a string of output symbols, and so the transducer can be seen as representing a relation between families of strings. Transducers (representing relations between strings) can be composed, and there are efficient algorithms for computing the composition of two transducers.

In a *string-to-weight transducer*, the output alphabet consists of weights (typically, in a semiring or better; non-negative reals with addition and min is common, because it corresponds to the case of Viterbi and negative log-probabilities); there are initial and final weights. If a string-to-weight transducer accepts some string, its output for that string is defined as the minimum sum of weights over the paths accepting the string. Particularly attractive are subsequential string-to-weight transducers, where there is only one path accepting any given string. Not all transducers can be transformed to this form; there are algorithms for this process, known as *determinization* when it is possible. Furthermore, there are *minimization* algorithms, that can produce the

unique (up to automorphism) smallest transducer that implements the same set of mappings as a given transducer.

6.3.1.2 Why should we Care?

Each of the components of a speech architecture (language model; a pronunciation dictionary; context dependency model; acoustic observation model) is a string-to-weight transducer. In principle, one could compose the lot to produce a single, enormous string-to-weight transducer, determinize it, minimize the result, and search that (this is equivalent to recognizing that, in the final analysis, the composition of each component produces an HMM with an enormous state space). In practice, the object involved is far too large. Instead, one uses a *beam search* to produce a reduced string-to-weight transducer (*the word lattice*) that contains a reduced pool of higher probability paths. Determinizing and minimizing this transducer is practical and useful; the result is very much faster searches.

There are two reasons that this material is of interest to us. First, the trick of reducing a speech signal to a (determinized and minimized) word lattice produces a highly compact representation of a large number of different transcriptions (each corresponding to a path through the string-to-weight transducer) that is easy to search and manage. We argue below that we can produce act, action and activity models which will allow reduction of video to an action/activity lattice with the same attractive properties. Second, a finite state automaton (whose states represent actions and activities) is a reasonable representation for a behaviour. If one determinizes and minimizes this, standard algorithms allow one to identify weights associated with instances of such a transducer in a word lattice extremely fast. This means we could be able to engage in fast searches for behaviours.

LeCun et al identify other useful building blocks associated with finite state models [242]. Their *graph transformers* take (weighted directed) graphs as inputs and produce graphs as outputs; an example of a transformer would be composition with a fixed transducer. Particularly useful is the idea of a Viterbi transformer, a process that (using our terminology) takes a string-to-weight transducer and applies a beam search to produce a reduced string-to-weight transducer which is effectively a word lattice. They demonstrate that gradient based learning can usefully be applied to architectures of such

objects.

6.3.2 Activity Recognition Methods based around HMM's

HMM's have been very widely adopted in activity recognition, but the models used have tended to be small (for example, one sees three and five state models in [54, 59]). Yamato *et al.* describe recognizing tennis strokes with HMM's [477]. Wilson and Bobick describe the use of HMM's for recognizing gestures such as pushes [467]. Yang et al use HMM's to recognize handwriting gestures [480]. Feng and Perona [126] call actions "movelets", and build a vocabulary by vector quantizing a representation of image shape, as a collection of rectangle, varying over time. These codewords are then strung together by an HMM, representing activities; there is one HMM per activity. We can then identify a new video by computing the image representation for each frame, obtaining the movelets, and choosing the particular model that generated the keyword sequence by a form of maximum likelihood. The method is not view invariant, depending on an image centered representation.

There has been a great deal of interest in models obtained by modifying the HMM structure. The intention is to improve the expressive power of the model without complicating the processes of learning or inference. Brand et al use coupled HMM's (CHMM's), which involve some number of simultaneous HMM's operating to the same clock, where the choice of a particular model's hidden state is affected by all other model's states [54, 59]. Such an object is clearly itself an HMM, but authors demonstrate a training method that reduces the number of parameters to learn by coupling but with very much enlarged state space; however, instead of estimating the parameters of that object, one projects the parameter estimates to transition parameters for each separate model. This means that one learns parameters for each separate model that tend to couple the two models. They show these models can distinguish between a set of T'ai Chi moves.

Oliver et al [316, 315] represent behaviours using layered hidden Markov models (LHMM's). These models involve a bank of HMM's at the lowest level, each generating some portion of the observation. The observations at higher levels are the maximum likelihood hidden state sequences for the lower levels. One then obtains for each HMM the maximum likelihood hidden state sequence. At the next level, the observations are these states, and this

continues recursively. The resulting object is an HMM, but of complex structure; the LHMM form offers authoring advantages. This representation outperforms a straightforward HMM in recognizing such activities as phone conversation from both vision and acoustic data. Similarly, Mori et al build a hierarchical representation out of HMM's to recognize everyday gesture [298].

Wilson and Bobick [468] use a form of HMM where an unknown, global parameter applies to all emission models (which they call a parametric hidden Markov model or PHMM) to model gestures with a parametric form (such as might accompany “it was *this* big”). Data is from stereo or a Polhemus. There are recognition results for classes of gesture such as pointing. Kettner and Brand [225](also, Brand and Kettner, [58]) fit an HMM while penalizing model entropy; this tends to reduce the number of non-zero parameters, so that one can fit models with quite large state spaces satisfactorily (such models are sometimes known as Entropic HMM's or EHMM's). Galata *et al.* use variable length Markov models (VLMM's: a model that generates a state stochastically based on a variable but bounded length history) to encode behaviour and obtain a reduction in perplexity by doing so [138, 139].

Building variant HMM's is a way to simplify learning the state transition process from data (if the state space is large, the number of parameters is a problem). But there is an alternative — one could author the state transition process in such a way that it has relatively few free parameters, despite a very large state space, and then learn those parameters.

Finite state methods have been used directly. Hongeng *et al.* demonstrate recognition of multiperson activities from video of people at coarse scales (few kinematic details are available); activities include conversing and blocking [183]. Zhao and Nevatia use a finite-state model of walking, running and standing, built from motion capture [485]. Hong *et al.* use finite state machines to model gesture [182]. We are not aware of material that attempts to build large hierarchical finite state machines, patterned after speech recognition programs, and using opportunistic learning, as we propose to do.

6.3.3 Sign Language Recognition

The best-known system for sign matching is due to Starner and Pentland [419, 420]. Features are image moments of the hand region; signers either wear coloured gloves, or hands are identified using a skin filter. A Hid-

den Markov Model (HMM) is used to model individual signs; signs are strung together with a rigid language model (pronoun verbnoun adjective pronoun). Authors report a recognition rate of 90% with a vocabulary of 40 signs. Grobel and Assan recognize isolated signs under similar conditions for a 262-word vocabulary using HMM's [227]. This work was extended to recognize continuous German sign language with a vocabulary of 97 signs by Bauer and Hienz [34]. Vogler and Metaxas have built a system that uses estimates of arm position, recovered either from a physical sensor mounted on the body or from a system of three cameras that measures arm position fairly accurately [455, 456, 459]. For a vocabulary of 53 words, and an independent word language model, they report a word recognition accuracy of the order of 90%. A more recent system attempted to recognize phonemes with HMM's; Vogler and Metaxas were able to recognize signs from a 22 word vocabulary with similar recognition rates for phoneme and word models (without handshapes in [457], with handshapes in [458]).

Kadous transduced isolated Australian sign language signs with a power-glove, reporting a recognition rate of 80% using decision trees [305]. Matsuo *et al* transduced Japanese sign language with stereo cameras, using decision tree methods to recognize a vocabulary of 38 signs [278]. Kim *et al.* transduce Korean sign language using datagloves, reporting 94% accuracy in recognition for 131 Korean signs [228]. Al-Jarrah and Halawani report high recognition accuracy for 30 Arabic manual alphabet signs recognized from monocular views of a signer using a fuzzy inference system [12]. Gao *et al.* describe recognizing isolated signs drawn from a vocabulary of 5177 using datagloves and an HMM model [141, 465]. Their system is not speaker-independent: they describe relatively high accuracy for the original signer, and a significant reduction in performance for other signers. Similarly, Zieren and Kraiss report high, but not speaker independent, accuracy for monocular recognition of German sign language drawn from a vocabulary of 152 signs [487]. Akyol and Canzler describe an information terminal which can recognize 16 signs with a high, user-independent, recognition rate; their system uses HMM's to infer signs from monocular views of users wearing coloured gloves [11]. Bowden *et al.* use independent component analysis to obtain state estimates from a set of discriminative visual features; each sign is encoded as a Markov chain, learned from a single example [52]. They report high accuracy recognition from a lexicon of 49 signs using a very small training set.

7

Discussion

7.1 Representations

The question of how one represents the configuration of the body appears to be important. In tracking applications, the important choice seems to be whether one tracks 3D or 2D representations, and we discuss this in some detail below. In many applications one predicts the configuration of the body from some evidence. Examples include a generative model of motion for animation; a dynamical model for tracking; a regression model for lifting from 2D to 3D. There is some reason to believe that the choice of the coordinates one predicts is important, and we discuss this point below. Finally, the reader will have noticed alarmingly contradictory evidence on the usefulness of dynamical models in understanding human motion; we try to resolve this contradiction.

7.1.1 Is 3D Configuration Ambiguous?

The work of both Taylor and of Barrón and Kakadiaris suggests that there are discrete ambiguities in 3D body configuration inferred from a single 2D view (in [31, 32]; see section 3.2.1). Depending on how many body segments one accepts, this ambiguity might be from 16-fold to 1024-fold. Some ambiguous

reconstructions might be ruled out by kinematic constraints, but one expects these ambiguities to manifest themselves in any attempt to recover the body in 3D.

We have seen a variety of strategies to disambiguate reconstructions. One might have more than one camera (section 3.1). One might observe local features that distinguish between a limb pointed towards and away from the camera (section 3.2.1; I have in mind the work of Mori and Malik [296, 297]). One might maintain a potentially multimodal representation of the posterior (section 3.3). One might use reconstructions in previous frames (section 3.2.3.4).

The problem with all this is that there is a body of work that does not explicitly respect these ambiguities and that does not suffer as a result. Under just what circumstances is 3D configuration ambiguous? I believe the picture is complex, and we need to break out cases.

7.1.1.1 Single Frames

First, it is clear that the ambiguities exist in a single frame view of the whole body. One must censor ambiguities using known kinematic limits, and this means that the extent of the ambiguity is, to an important degree, dependent on both the body configuration and the view direction. The situation is generally worse than one might expect from our account of geometric methods, because it is usually not possible to tell the difference between the left and right arms (resp. legs). These ambiguities are important in practice.

However, in a single frame *frontal* view of the upper body, there may be no ambiguity. Left and right arms are easily distinguished. There are several cases for the arms. First, if the hands are visible and occlude the torso then the shoulder is not significantly ambiguous — either the elbow is approximately on the plane of the torso (the one that passes through shoulders and navel), or it is in front; there is no need to attend to small angles here. It is possible to get into a configuration where the elbow is well behind the torso and the hands do not occlude the torso, but it is neither easy nor natural — one may be able to simply rule this out as possible but unlikely. Because the forearm is about as long as the upper arm, either the forearm is approximately parallel to the plane of the torso, or it is extended toward the camera. Often, this is required because otherwise the hands would be embedded in the torso; when it is not a kinematic necessity, it is uncomfortable. All this — loose but plausible —

argument suggests that 3D reconstruction ambiguities cannot contribute substantial errors to 3D reconstructions of frontal views of the upper body, and so explains why Shakhnarovich *et al.* don't need to deal with ambiguities.

7.1.1.2 Short Timescales

The tracking literature generally sees the posterior on 3D position given past 2D measurements (i.e. $P(\mathbf{X}_i | \mathbf{Y}_0, \dots, \mathbf{Y}_i)$) as multimodal, implying the presence of ambiguities. Managing these modes is the main thrust of that literature. However, there is some evidence that knowledge of future frames causes these ambiguities to disappear. In Howe's work [187], possessing the whole 2D track leads to an unambiguous 3D reconstruction (via dynamic programming, section 3.2.3.1). Howe *et al.* reconstruct 3D by matching to snippets of motion capture, as do Ramanan and Forsyth (section 3.2.3.2).

One possible resolution is as follows. 3D configuration (\mathbf{X}_i) is a multiple valued function of 2D configuration (\mathbf{Y}_i) (which is best thought of by considering the graph of the function, figure 7.1). A snippet — a short run — of frames corresponds to several possible paths on this graph. However, the process of censoring kinematically unacceptable reconstructions leads to a complicated structure, where parts of the graph are excluded. In turn, for most motions, very short runs of frames are ambiguous, but longer runs are not, because the incorrect paths wander into parts of the graph that are not available. This point is remarked on by Sminchisescu and Triggs ([407], p 372, “In practice, choosing the wrong minimum rapidly leads to mistracking...”), and may explain the “glitches” of [6] (p. 49).

This model explains why reconstruction on very short timescales may be ambiguous, while reconstruction on short timescales is not. The effect depends on the dynamic model, which must make narrow enough predictions (see figure 7.1). In turn this may explain why the 3D tracking literature (which uses either no dynamic model or a rather high entropy dynamic model) finds 3D reconstructions ambiguous.

We must recognize the limits of the available evidence here. All motion capture collections are small (and, for the foreseeable future, all possible motion capture collections will be small compared with the range of available motions). Furthermore, ambiguities are view dependent in important ways. It is entirely conceivable that research that finds 3D reconstructions by matching




Fig. 7.1 Some understanding of the behaviour of ambiguities in reconstructing the 3D configuration of the body (\mathbf{X}) from 2D image configuration (\mathbf{Y}) can be obtained by thinking about the graph of the multivalued function $\mathbf{X}(\mathbf{Y})$. The shape of this graph depends on the viewing direction, but it must have singularities and we expect that the process of censoring kinematically unacceptable reconstructions carves out holes in some of the sheets. While any single reconstruction may be ambiguous, as in the case shown here, sequences may not be, assuming that the dynamical model prohibits skips between sheets, etc. This model suggests several points. First, notice that reconstructing $(\mathbf{X}_1, \dots, \mathbf{X}_7)$ given $(\mathbf{Y}_1, \dots, \mathbf{Y}_7)$ is not ambiguous. Neither is reconstructing $(\mathbf{X}_4, \mathbf{X}_5)$ from $(\mathbf{Y}_4, \mathbf{Y}_5)$, for that matter. However, $(\mathbf{X}_1, \dots, \mathbf{X}_4)$ from $(\mathbf{Y}_1, \dots, \mathbf{Y}_4)$ is ambiguous. Second, the model does not suggest any difference between reconstructions that use only the past and those that use both the past and the future. Third, reconstructing \mathbf{X}_2 conditioned on \mathbf{Y}_2 and \mathbf{X}_1 — the procedure of Agarwal and Triggs [3, 6] — is still ambiguous. This method may encounter serious problems if it makes the wrong choice at this time, because it will then not be able to explain the measurement \mathbf{Y}_5 , (which may explain the “glitches” of [6], p. 49). Again, all these observations require a dynamical model that is relatively tight, something that matching to snippets — which implements a non-parametric dynamical model — supplies.

to motion capture simply hasn’t used enough motion capture data to observe ambiguities, or hasn’t used the right views of the body. I don’t believe this to be the case, but the matter needs clearer resolution.

7.1.1.3 Long Timescales

Some ambiguities persist over long timescales, and should be resolved by better image measurement. The best example is the left-leg/right-leg ambiguity in lateral views of actions like walking. In principle, there is an ambiguity at

each frame of such a sequence, but dynamic constraints and camera motion constraints mean that this ambiguity is of the order of one bit per sequence. Other cases are left-arm/right-arm ambiguity in lateral views and front/back ambiguity in some frontal views. I do not know of a complete list of such cases, though it appears that such a list would be short.

7.1.1.4 Summary

There is a body of evidence, strongly suggestive though not absolutely conclusive, that 3D reconstruction from 2D frames has few ambiguities that persist over any but the shortest timescales. Those that do, can persist over quite long timescales, and have to do with left/right labels rather than the configuration. This property depends on the dynamic model adopted, and requires that one use a snippet of frames.

If I have interpreted this evidence correctly, it suggests that the proper approach to tracking in 3D is to track in 2D, and then report an estimate of 3D by matching the 2D track to 3D body configuration snippets. This is because one does not then need to deal explicitly with multiple modes in the posterior. There are two cases: one could reconstruct \mathbf{X}_i from $(\mathbf{Y}_{i-k}, \dots, \mathbf{Y}_{i+k})$, or from $(\mathbf{Y}_{i-2k}, \dots, \mathbf{Y}_i)$. I do not see strong evidence to distinguish between these cases, but I believe that the first may be better where possible, as noise in the dynamical model may accumulate.

7.1.2 What Representation of 3D Configuration should be Adopted?

There are two major options for representing the configuration of the body in 3D. One might use joint angles, or one might use joint positions. But which is better? Apart from the mild complications in passing from joint positions to joint angles (an entire subject, section 5.1.3), the question is basically an empirical one. It is an important empirical question that hasn't received enough attention. This is because, if one wishes to regress the 3D configuration against some variable (section 3.2.3), one needs information about covariance in the 3D coordinate system.

This need appears in a number of ways, some less obvious than others. If one is building a straightforward regression, then for something as high-dimensional as body configuration one is forced to assume a reduced form

(diagonal; a constant scaling of the identity; or some such) for the error covariance, which is too big to estimate accurately. If one is building a nearest neighbour method, it is a good idea to work in coordinates that are largely independent (or, which is the same thing, to weight distances with an inverse covariance matrix). There is some evidence that quite low-dimensional representations of motion are tolerable for some synthesis applications (for example, Safonova *et al.* synthesize motion in low-dimensional spaces without major costs in quality [378]; and see section 5.4.1). One might think that joint angles are a better coordinate system, because joint positions are clearly correlated (some points are a fixed distance apart). There isn't much evidence on this point, and all we have favours joint positions as a representation.

Arikan describes a method to compress motion signals by fitting a parametric curve to joint position information, clustering the results, representing each cluster with principal components, and then using a discrete cosine transform to represent fast phenomena that occur as a result of contacts [?]. This method is much more effective than compressing any joint angle representation, so much so that the overhead of inverse kinematics in decompression (which is simplified by attaching an extra vertex to each segment and compressing the overcomplete representation) presents no problem. The process of clustering motions, then applying PCA within a cluster, produces a form of decorrelated representation.

We expect that correlation structure within a motion varies between types of motion. It isn't currently possible to be precise about what the term "type" means here, but a stroll and a walk might be the same type of motion, whereas a walk and a throw are not. In walking, there is a characteristic oscillatory motion of both upper and lower body, 180° out of phase with one another. A good parametric representation of a particular walk might require very few parameters — frequency and phase might do it. If the intention is to perform kinematic reconstructions for configurations whose frequencies are well represented by typical motion capture or video data (in outdoor data, lots of walking, some running, and other activities very infrequent), then the correlations between joint angles typical of walking are very important.

The correlations observed in walking are very different from those observed in, say, throwing or striking motions. In walking, the arm moves, rather roughly, like a pendulum. In some throwing or striking motions, there is a clear proximal-distal sequence by which the joints are activated, leading to a

whip-like motion (for some cases, see [13, 23, 352]; the effect does not occur in all sports [136]; it can be used in animation [41]). This means that, for example, a near-straight elbow implies a particular shoulder position quite accurately, which is quite unlike walking (where the elbow is always close to straight, whatever the shoulder position).

There are two issues here. First, talking about correlation requires some sensible theory about frequencies of events within motions, which appears to be hard to obtain. We discuss this point below. Second, assuming that we have some such theory, we should respect it in choice of regression coordinate system. In particular, we expect that regression predictions of 3D configuration from 2D will perform better or worse with different choices of coordinates. This point doesn't appear to have been much discussed in the literature, though it may help motivate to Shakhnarovich *et al.*'s work on making locality sensitive hashing sensitive to sharp changes in predicted parameters [389].

The frequency with which different motions occur is not much discussed in the literature, but it's a difficult point with some nasty consequences. For example, one could produce an (apparently) very effective outdoor surveillance system by simply labelling every activity observed as walking. This system would be wrong an infinitesimal percentage of the time, because most of the time people are walking; but its output would be unhelpful. Recovering accurate labellings of relatively uncommon events is what is required, and this means collecting data is tricky and model-building is important. For example, in years of informal observation of people outside I have never seen a flasher, and can so presume that the phenomenon is relatively uncommon, but we know that it represents a significant nuisance that engages authority. Should we represent what a flasher does with models built from data? if so, where is the data to come from? if not, how?

The likely differences in covariance structure of different types of motion suggests that we should impose some sort of hierarchical structure on motion data. We know this can be done for at least some kinds of structure and some collections of data, because motion capture data appears to cluster very well; in many systems, clustering the motion capture data is a first step, and no bad consequences appear to result. But there hasn't been much investigation of what sort of structures are good. A good structure might make 3D from 2D easier, by using gross motion phenomena to predict the type of motion and then operating in an advantageous 3D representation. A good structure might

lead to better dynamical models (section 7.1.3). And a good structure might make at least some aspects of a vocabulary for actions or activities apparent.

7.1.3 What is the Status of Dynamical Models?

The literature contains a series of positions on dynamical models of motion. The idea that dynamical models are not helpful, or are even harmful, in tracking is suggested by, among others, the work of Sminchisescu and Triggs [407], of Mori and Malik [296, 297], and of Ramanan *et al.* [356, 357] (sections 3.3.2, 3.2.1 and 4.2). This work simply dispenses with dynamical models as an unreliable guide to the future configuration of a person. In fact, Sminchisescu and Triggs suggest that such dynamical models as have been adopted in the particle filter tracking literature, have been built more to compensate for weaknesses in the search process than as predictive models ([407], p. 373). Furthermore, it has been remarkably difficult to build methods that can reliably tell whether a given motion is a good human motion or not (a point we discuss in section 7.2.1).

Dynamical models of human motion have tended to lead to animations of relatively poor quality (it is unfair to name names). One difficulty seems to be that, when one fits a parametric model to motion capture data, the inevitable slight errors in temporal alignment smooth out some high frequency structure, so that motions that should have fast definition (hitting, jumping, etc.) become “squashy” in appearance. Sometimes important physical properties of motions are preserved [377] and sometimes they are not. Often, the motion that results is ugly.

However, we have quite strong evidence that the dynamics of motion is constrained. There are several points. First, Sidenbladh *et al.* check obtain very good tracks from low-dimensional parametric models fitted to motion capture data; of course, one must be sure that the person being tracked engages in the activity to which the model was fit, but this difficulty doesn’t erase the usefulness of the dynamical model. Second, the fact that both Howe *et al.* [188] and Ramanan and Forsyth [355] can lift to 3D by matching multiple frames of motion capture to multiple frames of image data suggests that some form of dynamical constraint is present. If there wasn’t much constraint at the relevant time scales (approx 1/6 second), then some of the video snippets would not find a good match and the lift would be grossly inaccurate.

This suggests (but does not establish) that motion at short time scales has a fairly rigid structure. The difficulty regarding this point as comprehensive is that, in both cases, the collections of video and of motion capture are quite small — perhaps both sets of authors were lucky. Third, motion capture data appears to cluster rather well, as we have said. And fourth, Arikan can compress motion successfully, by clustering and then compressing snippets of motion about a second long [?].

There is something here that isn't as well understood as it needs to be. I believe the resolution is as follows: at very short time scales (say, 1/60 s), the number of kinematic configurations that will ever follow a given configuration under any circumstances is probably small. There is some difficulty being formal, because we don't really know what a fair sample of motion is, so it is difficult to talk about the frequency of events. This difficulty isn't so significant at very short timescales; I claim that, whatever one's model of the frequency of activities, at very short timescales the conditional entropy of the next frame of motion given the past frames is very small. At longer timescales — say, a second — this language is more difficult to use because one probably can get from any one body configuration to any other in a second (if one ignores the root), and one has to deal with the question of how often particular transitions arise. In current motion capture collections, given that we think of ourselves as quite mobile, opportunistic movers, the notable feature is how seldom most transitions occur. In fact, they occur with frequency zero, which suggests some interesting questions about smoothing here, quite like those that arise in natural language problems where most pairs of words do not occur (see [274, 217]).

This view of motion as highly constrained at short timescales is consistent with the evidence that motion is constrained. But I believe it is also consistent with the evidence that dynamical models, as currently practiced, aren't particularly helpful. If motion behaves as I have described it, most current dynamical models put almost all of their probability in the wrong place, which would be the problem. This problem may be quite difficult to fix. Body configuration appears to occupy a fairly high-dimensional space, though it is probably confined to a low-dimensional subset of that space. It is technically quite difficult to build models that make predictions that are confined to a “small” subset of a low-dimensional subset of a high-dimensional space, particularly when one doesn't know what the low-dimensional subset is. Nonetheless, the ef-

fort may be worthwhile, because good dynamical models of motion would be valuable in animation. The situation in tracking is less clear — one might get a better result from efforts to improve appearance models than from efforts to incorporate improved dynamical models. The increased understanding of motion that would result from an attempt to build improved dynamical models is certainly worthwhile.

7.1.4 The Space of Human Motions

As we have seen (section 5.1.4), one can obtain good kinematic reconstructions in the presence of ambiguous constraints by requiring that the reconstruction be close to a space learned from data. This suggests that relatively few available body positions are actually occupied. This, the fact that motion clusters well (section 5.3.2; section 5.4.1), and the fact that motion can be compressed effectively ([?]), suggest it may be helpful to think about the space of human motions as a geometrical object.

For the moment, let us adopt some encoding of the state of the body (the details don't matter for this discussion, but we'd expect to see the configuration of the root, the configuration of the body relative to the root, velocities and most likely accelerations in this encoding). Because segment lengths don't vary, because velocities are limited and because there are torque limits, not every point in this state space represents a legal motion. It is useful to think of the legal motions as forming a "sheet" in this space. We make no claim on the topology of this object, not even that it is a manifold. We can think of motions as functions from time to this space. These functions must meet some obvious constraints — for example, velocities computed as time derivatives of kinematic configuration need to be the same as corresponding velocities recorded in the state vector. We expect other local constraints, too, resulting from torque limits and the like. We can represent the space of human motions by all acceptable functions from time to our space. There should be some form of structure at long time scales — we know, for example, that it is possible to walk backwards for long distances, but that it is very seldom done — but shorter time scales are easier to handle at present.

This object is intimately related to blending. Assume we have two legal states x_1 and x_2 that are close. For many such pairs, we can expect that states that lie on the line segment joining them are also legal. Another way

to put this point is that, if the two states are sufficiently close, then the vector $\mathbf{x}_2 - \mathbf{x}_1$ should lie on the tangent space at \mathbf{x}_1 . Now assume we have two observed motions $\mathbf{f}_1(t)$ and $\mathbf{f}_2(t)$, which run for similar time periods and are sufficiently similar to one another (I do not know how to be precise what this means). We expect — and can observe in data — that repeated versions of the same movement have slightly different temporal parametrizations. For example, each step of a walk can take a slightly different span of time. This means that we will need to massage the temporal parametrization, which is what time alignment does. Assume we can place states in correspondence by a small — again, it isn't currently possible to be precise — change in temporal parametrization $\tau(t)$, so that $\mathbf{f}_1(t)$ is close to $\mathbf{f}_2(\tau(t))$. Under these circumstances, we can expect that $\mathbf{f}_1(t) - \mathbf{f}_2(\tau(t))$ lies close to the tangent space to the space of human motions.

We know that good blends can be obtained from nearby motions, and that viable deformations include filtering angles, adding constant offsets, deforming the root path, applying a global rigid-body transformation and applying small time deformations. Can we infer others from seeing blends as tangents to the space of motions? There is some reason to hope that we can, because all the deformations I have described form actions of a local group, and this implies a structure to the tangent space.

7.1.5 Is Human Tracking Multimodal?

Time, space, and available energy have limited the number of citations to the vast literature on human tracking. Much of this literature is about a single point: how to manage inference in the presence of multimodal posteriors, possibly in a high-dimensional space. The main methods are variants of the particle filter, section 3.3.1. However, other methods are possible. For example, in a multiple hypothesis tracker, one Kalman filter keeps track of each of a fixed number mode, and we must determine methods to prune the number of modes. This method was used by Cham and Rehg [75], to track a 2D kinematic model of the body — note that this was adopted explicitly to cope with data association difficulties, and they make no argument that posteriors for 2D human tracking are intrinsically multimodal. One might maintain a mixture of gaussians, a mixture of other densities, or some form of kernel representation; all the options have been thoroughly explored.

A core thesis of this work is that all this effort is unnecessary. We do not need to deal with multiple modes resulting from data association problems, if we deal with data association directly. There are now excellent tools for doing so, as chapter 4 has demonstrated. These tools support only 2D tracking, however, and we might reasonably wish to report the configuration of the body in 3D. To do so, we may have to deal with ambiguities in the likelihood, which will result in multiple modes. As we have argued above (section 7.1.1), there are several ways to avoid ambiguity. First, better image measurement may reduce ambiguity (after Mori and Malik [296, 297]). Second, it is very likely that 2D tracks allow unambiguous lifts to 3D for “snippets.” Third, for some situations the ambiguity may not, in fact, appear.

The point is a general one: good features combined with simple inference (resp. classification) methods seem to be better than bad features combined with sophisticated methods. Given finite resources, we should pay more attention to visual features and phenomena than to the alluring world of statistical algorithms.

7.2 Generalization

Many of the methods we have described can loosely be described as a statistical view of motion — in essence, we are expecting, usually implicitly, that a model that is good at representing the motions that one has seen will be good at representing the motions that one will see. This property of a model and a dataset is known as **generalization** in the machine learning community, where quite strong guarantees are available if one has an appropriately representative data set and if the model adopted meets certain criteria (e.g. see [447, 448]). There is no reason to believe that these guarantees are available in the case of human motion; it appears likely that they never will be.

This is a problem that has to do with both data and models. There is an important issue of datasets here that clouds the picture somewhat. In our opinion, it probably is the case that many significant motion distinctions are “large” — in the sense that they involve huge changes in kinematic configuration — and so quite simple clustering and dimension reduction methods can expose much structure in motion. What remains uncertain is the extent to which the vocabulary of motions that are well-behaved in this way can be used to encode what one does every day — current experimental work covers

relatively small ranges of motion, because motion data is difficult to collect in large volumes. Furthermore, it isn't currently possible to collect data without being intrusive — there are no collections of motion data that can be said to represent “what people do”. Finally, there is a significant difficulty with rare motions. In some applications, not encoding a motion that people do relatively seldom is entirely appropriate (for most animation applications, for example, relatively small amounts of sensibly collected motion data is quite sufficient). In other applications, one should be able to encode even very rare behaviours (think contortionist), so that they can be reported.

This difficulty manifests itself in two important, and related technical problems that are largely unsolved. First, all automatic methods for scoring motions generalize poorly. Second, data-driven methods for generating motion cannot produce satisfactory motions that are significantly different from the input data (or, equivalently, generalize poorly). Motions appear to have structural properties — like composition across the body — that produce a very large range of motions, too large to sample and observe with current methods. The problem seems to be that good generalization will require good models for these properties, and we don't have them.

7.2.1 Which Motions are Human?

People are often extremely sensitive to the detailed structure of a motion. Several researchers have used light-dot displays, also referred to as biological motion stimuli, to study perception of human movements [137]. The light-dot displays show only dots or patches of light that move with the main joints of walking figures, but even these minimal cues have been shown to be sufficient for viewers to make detailed assessments of the nature of both the motion and the underlying figure[211]. Work by Cutting and Kozlowski showed that viewers easily recognized friends by their walking gaits on light-dot displays[92]. They also reported that the gender of unfamiliar walkers was readily identifiable, even after the number of lights had been reduced to just two located on the ankles[238]. In a published note, they later explained that the two light-dot decisions were probably attributable to stride length[239]. Continuing this work, Barclay, Cutting, and Kozlowski showed that gender recognition based on walking gait required between 1.6 and 2.7 seconds of display, or about two step cycles[29, 93]. Not much is known about

what inclines people toward or away from the judgement that a motion is “good” or “natural”. It is known that the choice of rendering has an effect, with more naturalistic renderings making people more inclined to reject motions [178, 177].

A device that could tell good, human-like motions from bad ones would be very useful. One could animate new motions using hypothesize-and-test using such a device. Ideally, the device might produce some information about what looks good or bad about the animation. One could use it to test tracks of activities that had never before been seen to tell whether the track represented a human motion or a tracker failure. Ideally, the device might produce some probability that the observation had come from a person.

Building one is difficult. There have been several attempts. Generalization — giving an accurate score to motions very different from the training motions — is a notoriously difficult problem. Ikemoto and Forsyth use a classifier to evaluate motions produced by a cut-and-paste method, and find the classifier significantly less accurate on novel motions [194]. The classifier is trained using both positive and negative examples. There is some advantage to not using negative examples, which can be both difficult to obtain and inaccurate. Ren et al. fit an ensemble of generative models to positive examples; motion is scored by taking the lowest likelihood over all models to obtain a conservative score [359]. While the combined model gives the best behaviour in practice, their hidden Markov model (HMM) is almost as accurate as the combined model. There is no information on generalization behaviour.

Arikan *et al.* use a regression method (built using scattered data interpolation) to predict the goodness of applying a particular deformation to a particular motion to represent a push or a shove [19]. Their oracle agrees roughly with the behaviour of human observers in a two-alternative forced-choice test. In particular, the probability that a human will say a motion is good when the oracle says it is bad, is low. The probability that a human will say a motion is good when the oracle says it is good is around 50% (the exact value depends on the study group). This needs to be compared with the probability that a human will say that pure motion capture is good, which is approximately the same. The logic of their application means that the oracle is never presented with examples that are strongly different from the training set.

However, if negative examples are available, we expect that models trained discriminatively are likely to perform better, because they possess

more information about the location of the boundary between good and bad motion. Ikemoto *et al.* train several scoring functions on 400 short motion transitions, annotated as good or bad motions by hand [193]. Methods include: likelihood under an HMM fitted to positive motion examples represented by an acceleration feature vector; logistic regression applied to an acceleration feature vector; the minimum score of this logistic regression and another applied to a feature that encodes footskate; and a score of footskate. The scoring methods that encode footskate outperform the others, and the pure footskate score is the best. There is no information on generalization.

There are several reasons it is difficult to build this device. There is a painful shortage of useful data. While there is a lot of motion capture data available, no collection of practical size can explore all that a body can do. At least in part, the relative poverty of collected motion data is because one can compose motions across time and across the body — it is possible to walk while scratching with either hand. The structural models necessary to encode this property do not yet exist. Some data may not be as useful as it looks. Arikan *et al.*'s subjects were inclined to regard rendered motion capture data as unnatural about half the time in two-alternative (good/bad) forced choice tests [19]. This may have to do with the motion capture pipeline. It is hard to get good marker placements and good measurements, and quite often motion is cleaned up.

Furthermore, data is encoded in a high-dimensional space, with all the attendant difficulties, and while we know there are correlations between dimensions (above), we don't know much about what they are, even at the level of practical scientific folklore. To make things worse, it isn't clear what features expose the phenomena that make a motion look good or bad. For example, intuition might suggest that whether a motion was "physical" is an important criterion, but there is little evidence in support of this view. Similarly, footskate doesn't appear to be much more than a detail, but the presence of footskate seems to be quite a good test for whether a motion is good or not. Finally, the relatively constrained structure of motion (if one accepts the argument above), means that building a good classifier or scoring function might be quite difficult, because it must cut out a very small and complicated portion of a spatio-temporal encoding of motion.

7.3 Resources

Getting good motion capture data requires considerable effort, skill and expense. Relatively few groups have found it useful to have their own motion capture studio. For those who wish to, major manufacturers of motion capture equipment include Vicon (<http://www.vicon.com>) and Motion Analysis (<http://www.motionanalysis.com>). Hodgins' group at CMU has done a great service to the research community by collecting and publishing some 1700 motion sequences, available at <http://mocap.cs.cmu.edu/>.

There are several other reviews of aspects of human motion. In animation, Hodgins *et al.* give a general review of computer animation [176]; Multon *et al.* survey computer animation of human walking [300]; and Gleicher gives a brief survey of animation from examples, motion capture and motion editing [155].

There are more reviews of tracking methods, none particularly recent. There is a special issue of Computer Vision and Image Understanding dedicated to vision based understanding of shape, appearance and movement (volume 81, 2001). Moeslund and Granum give an extensive survey of computer vision based methods for human motion capture [288]. Gleicher and Ferrier give a critical review of methods to recover 3D configuration from video, concentrating on single views [157]. Aggarwal *et al.* review articulated motion understanding [8]; Aggarwal and Cai review human motion analysis [7]. Gavrilu surveys visual analysis of human movement [143]. Hu *et al.* survey visual surveillance of object motion and behaviour [190]. Wang and Singh survey video analysis of human dynamics [462].

Index

- (r, ϵ) nearest-neighbours, 69
- active markers, 102
- adverbs, 125
- alignment, 112
- annealed search, 61
- Annotation constraints, 111
- appearance, 14
- appearance map, 26
- ballistic motion, 137
- basic recipe
 - track by detection, 11
- blobs, *see* spatial scale
- cascade architecture, 78
- chaff, 7
- choice points, 116
- cluster graph, 121
- cluster path, 121
- cluster tree, 121
- coarse scale, *see* spatial scale
- computed edges, 116
- configuration, 82
- configuration space, 104
- consecutive fragments, 132
- constrained motion, 137
- coordinate velocity, 137
- data association, 7
- decision stumps, 69
- deformable template, *see* snakes
- deformable templates
 - active sites, 15
 - curve template, 15
 - keypoints, 15
 - point template, 15
 - snakes, 15
- detailed constraints, 110
- diameter, 119
- differential-algebraic equations,

- 134
- difficulties at fine scale
 - spatial scale, 9
- distance transforms, 82
- dynamical constraints, 133
- end-stop, 35
- Euler-Lagrange equations, 134
- exemplars, 87
- face detection, 75
- feasible set, 107
- fine scale, *see* spatial scale
- Finite state models, 147
- flow field, 21
 - optic flow equation, 21
- footskate, 105
- frame constraint, 110
- gate, 16, 33
- general points
 - appearance phenomena, 10
 - data association, 10
 - dynamics of motion, 10
 - state dimension, 9
- generalization, 164
- generalized coordinates, 133
- generalized Hausdorff distance, 18
- Geometric constraints, 110
- global search, 119
- handles, 136
- Hausdorff distance, 18
- hidden states, 30
- HOG, 77
- horizon problem, 120
- Image appearance, *see* appearance
- image stabilization, 22
- Importance sampling, 55
- instantaneous constraints, 110
- inverse kinematics, 106
- iterative scaling, 56
- joint angles, 104
- joint positions, 104
- Kalman filter, 31
- kernel, 72
- kernel function, 72
- keyframe, 110
- kinematic constraints, 134
- kinematic redundancy, 107
- kinematic singularity, 35
- kinematic tracking
 - definition, 8
- kinematic tree, 28
- Lagrange multipliers, 133
- Lagrangian, 133
- lasso, 71
- linear regression, 70
- local linear models, 71
- local search, 119
- locality sensitive, 69
- locality sensitive hashing, 69
- low temporal entropy, 123
- Markov chain, 30
- mass displacement, 136, 137
- matrix exponential, 30
- maximum entropy models, 56
- medium scale, *see* spatial scale
- motion ambiguity, 122
 - global motion ambiguity, 122

- Local motion ambiguity, 122
- Motion capture, 102
- motion fields, *see* spatial scale
- motion graph, 116
- motion primitives, 125
- motion transitions, 136
- movemes, 125
- nearest neighbours, 33
- observed edges, 116
- occupancy, 8
- optic flow field, *see* flow field
- outliers, 17
- particles, 65
- Partitioned sampling, 57
- passive markers, 102
- patterns of activity, 8
- posterior, 31
- predictive density, 31
- principal component analysis, 15
- prior, 31
- probabilistic data association, 34
- probabilistic tracking
 - prior, 31
- proportional-derivative controller, 136
- reduced coordinates, 133
- regression problem, 48
- relevance vector machine, 71
- ridge regression, 71
- robustness, 17
 - m-estimator, 17
- sample impoverishment, 67
- shape context, 44
- skeletal constraints, 104
- skeleton, 103
- skinning, 104
- spacetime constraints, 133
- spatial scale
 - blobs, 8
 - coarse scale, 8
 - fine scale, 8
 - medium scale, 8
 - motion fields, 8
- static balance, 137
- stimulus-response, 135
- summary constraints, 110
- support vector machine, 76
- time align, 112
- transition, 116
- transitions, 113
- variational methods, 133
- verb graph, 125
- verbs, 125
- video puppetry, 127
- window, *see* chaff
- XYT image, 78

References

- [1] Y. Abe, C. K. Liu, and Z. Popović. Momentum-based parameterization of dynamic character motion. In *SCA '04: Proceedings of the 2004 ACM SIGGRAPH/Eurographics symposium on Computer animation*, pages 173–182, New York, NY, USA, 2004. ACM Press.
- [2] R. Abraham and J. E. Marsden. *Foundations of Mechanics*. Addison-Wesley, 1978.
- [3] A. Agarwal and B. Triggs. Learning to track 3d human motion from silhouettes. In *ICML '04: Proceedings of the twenty-first international conference on Machine learning*, page 2, New York, NY, USA, 2004. ACM Press.
- [4] A. Agarwal and B. Triggs. Tracking articulated motion using a mixture of autoregressive models. In *European Conference on Computer Vision*, pages Vol III: 54–65, 2004.
- [5] A. Agarwal and B. Triggs. Monocular human motion capture with a mixture of regressors. In *IEEE Conf. on Computer Vision and Pattern Recognition*, 2005.
- [6] A. Agarwal and B. Triggs. Recovering 3d human pose from monocular images. *IEEE T. Pattern Analysis and Machine Intelligence*, 28(1):44–58, 2006.
- [7] J. Aggarwal and Q. Cai. Human motion analysis: A review. *Computer Vision and Image Understanding*, 73(3):428–440, March 1999.
- [8] J. Aggarwal, Q. Cai, W. Liao, and B. Sabata. Nonrigid motion analysis: Articulated and elastic motion. *Computer Vision and Image Understanding*, 70(2):142–156, May 1998.
- [9] G. Agin and T. Binford. Computer description of curved objects. In *Int. Joint Conf. Artificial Intelligence*, pages 629–640, 1973.
- [10] G. Agin and T. Binford. Computer description of curved objects. *TC*, 25(4):439–449, April 1976.

- [11] S. Akyol and U. Canzler. An information terminal using vision based sign language recognition. In *ITEA Workshop on Virtual Home Environments, VHE Middleware Consortium*, pages 61–68, 2002.
- [12] O. Al-Jarrah and A. Halawani. Recognition of gestures in arabic sign language using neuro-fuzzy systems. *Artificial Intelligence*, 133(1-2):117–138, December 2001.
- [13] R. M. Alexander. Optimum timing of muscle activation for simple models of throwing. *J. Theor. Biol.*, 150:349–372, 1991.
- [14] J. F. Allen. Towards a general theory of action and time. *Artif. Intell.*, 23(2):123–154, 1984.
- [15] F. Anderson and M. Pandy. A dynamic optimization solution for vertical jumping in three dimensions. *Computer Methods in Biomechanics and Biomedical Engineering*, 2:201–231, 1999.
- [16] W. Arentz and B. Olstad. Classifying offensive sites based on image content. *Computer Vision and Image Understanding*, 94(1-3):295–310, April 2004.
- [17] O. Arikan, D. Forsyth, and J. O’Brien. Motion synthesis from annotations. In *Proceedings of SIGGRAPH 95*, 2003.
- [18] O. Arikan and D. A. Forsyth. Interactive motion generation from examples. In *Proceedings of the 29th annual conference on Computer graphics and interactive techniques*, pages 483–490. ACM Press, 2002.
- [19] O. Arikan, D. A. Forsyth, and J. F. O’Brien. Pushing people around. In *SCA '05: Proceedings of the 2005 ACM SIGGRAPH/Eurographics symposium on Computer animation*, pages 59–66, New York, NY, USA, 2005. ACM Press.
- [20] V. I. Arnold. *Mathematical Methods of Classical Mechanics*. Springer-Verlag, 1989.
- [21] V. Athitsos and S. Sclaroff. An appearance-based framework for 3d hand shape classification and camera viewpoint estimation. In *Int. Conf. Automatic Face and Gesture Recognition*, pages 40–45, 2002.
- [22] V. Athitsos and S. Sclaroff. Estimating 3d hand pose from a cluttered image. In *IEEE Conf. on Computer Vision and Pattern Recognition*, pages II: 432–439, 2003.
- [23] A. Atwater. Biomechanics of overarm throwing movements and of throwing injuries. *Exerc. Sport. Sci. Rev.*, 7:43–85, 1979.
- [24] N. I. Badler, B. A. Barsky, and D. Zeltzer, editors. *Making them move: mechanics, control, and animation of articulated figures*. Morgan Kaufmann Publishers Inc., San Francisco, CA, USA, 1991.
- [25] P. Baerlocher and R. Boulic. An inverse kinematics architecture enforcing an arbitrary number of strict priority levels. *The Visual Computer*, 20(6):402–417, 2004.
- [26] H. Baker. Building surfaces of evolution: The weaving wall. *Int. J. Computer Vision*, 3(1. May 1989):51–72, May 1989.
- [27] Y. Bar-Shalom and X.-R. Li. *Estimation with Applications to Tracking and Navigation*. John Wiley & Sons, Inc., New York, NY, USA, 2001.
- [28] J. Barbič, A. Safonova, J.-Y. Pan, C. Faloutsos, J. K. Hodgins, and N. S. Pollard. Segmenting motion capture data into distinct behaviors. In *GI '04: Proceedings of the 2004 conference on Graphics interface*, pages 185–194, School of Computer Science, University of Waterloo, Waterloo, Ontario, Canada, 2004. Canadian Human-Computer Communications Society.

- [29] C. D. Barclay, J. E. Cutting, and L. T. Kozlowski. Temporal and spatial factors in gait perception that influence gender recognition. *Perception & Psychophysics*, 23(2):145–152, 1978.
- [30] G. Barenblatt. *Scaling*. Cambridge University Press, 2003.
- [31] C. Barron and I. Kakadiaris. Estimating anthropometry and pose from a single image. In *IEEE Conf. on Computer Vision and Pattern Recognition*, pages I: 669–676, 2000.
- [32] C. Barron and I. Kakadiaris. Estimating anthropometry and pose from a single uncalibrated image. *Computer Vision and Image Understanding*, 81(3):269–284, March 2001.
- [33] H. Barrow, J. Tenenbaum, R. Bolles, and H. Wolf. Parametric correspondence and chamfer matching: Two new techniques for image matching. In *Int. Joint Conf. Artificial Intelligence*, pages 659–663, 1977.
- [34] B. Bauer and H. Hienz. Relevant features for video-based continuous sign language recognition. In *IEEE Workshop on Automatic Face and Gesture Recognition*, pages 440–445, 2000.
- [35] A. Baumberg and D. Hogg. Learning flexible models from image sequences. In *European Conference on Computer Vision*, pages 299–308, 1994.
- [36] S. Belongie, J. Malik, and J. Puzicha. Shape matching and object recognition using shape contexts. *IEEE T. Pattern Analysis and Machine Intelligence*, 24(4):509–522, April 2002.
- [37] J. Ben-Arie, P. Pandit, and S. Rajaram. View-based human activity recognition by indexing and sequencing. In *IEEE Conf. on Computer Vision and Pattern Recognition*, pages II:78–83, 2001.
- [38] J. Ben-Arie, Z. Wang, P. Pandit, and S. Rajaram. Human activity recognition using multidimensional indexing. *IEEE T. Pattern Analysis and Machine Intelligence*, 24(8):1091–1104, August 2002.
- [39] V. Beneš. Exact finite-dimensional filters with certain diffusion non linear drift. *Stochastics*, 5:65–92, 1981.
- [40] A. Berger, S. D. Pietra, and V. D. Pietra. A maximum entropy approach to natural language processing. *Computational Linguistics*, 22(1), 1996.
- [41] D. Bhat and J. Kearney. On animating whip-type motions. *The Journal of Visualization and Computer Animation*, 5:229–249, 1996.
- [42] T. Binford. Inferring surfaces from images. *Artificial Intelligence*, 17(1-3):205–244, August 1981.
- [43] M. Black, Y. Yacoob, A. Jepson, and D. Fleet. Learning parameterized models of image motion. In *IEEE Conf. on Computer Vision and Pattern Recognition*, pages 561–567, 1997.
- [44] S. Blackman and R. Popoli. *Design and Analysis of Modern Tracking Systems*. Artech House, 1999.
- [45] A. Blake and M. Isard. *Active Contours: The Application of Techniques from Graphics, Vision, Control Theory and Statistics to Visual Tracking of Shapes in Motion*. Springer-Verlag New York, Inc., Secaucus, NJ, USA, 1998.
- [46] A. Bobick. Movement, activity, and action: The role of knowledge in the perception of motion. *Proc. Roy. Soc. B*, 352:1257–1265, 1997.
- [47] A. Bobick and J. Davis. The recognition of human movement using temporal templates. *IEEE T. Pattern Analysis and Machine Intelligence*, 23(3):257–267, March 2001.

- [48] A. Bobick and A. Wilson. A state based technique for the summarization and recognition of gesture. In *Int. Conf. on Computer Vision*, pages 382–388, 1995.
- [49] A. Bobick and A. Wilson. A state based approach to the representation and recognition of gesture. *IEEE T. Pattern Analysis and Machine Intelligence*, 19(12):1325–1337, December 1997.
- [50] B. Bodenheimer, C. Rose, S. Rosenthal, and J. Pella. The process of motion capture: Dealing with the data. In *Computer Animation and Simulation '97. Proceedings of the Eurographics Workshop*, 1997.
- [51] A. Bosson, G. Cawley, Y. Chan, and R. Harvey. Non-retrieval: Blocking pornographic images. In *Int. Conf. Image Video Retrieval*, pages 50–59, 2002.
- [52] R. Bowden, D. Windridge, T. Kadir, A. Zisserman, and M. Brady. A linguistic feature vector for the visual interpretation of sign language. In *European Conference on Computer Vision*, pages Vol I: 390–401, 2004.
- [53] J. Boyd and J. Little. Phase in model-free perception of gait. In *IEEE Workshop on Human Motion*, pages 3–10, 2000.
- [54] M. Brand. Coupled hidden markov models for complex action recognition. Media lab vision and modelling tr-407, MIT, 1997.
- [55] M. Brand. An entropic estimator for structure discovery. In *Proceedings of the 1998 conference on Advances in neural information processing systems II*, pages 723–729, Cambridge, MA, USA, 1999. MIT Press.
- [56] M. Brand. Shadow puppetry. In *Int. Conf. on Computer Vision*, pages 1237–1244, 1999.
- [57] M. Brand. Structure learning in conditional probability models via an entropic prior and parameter extinction. *Neural Comput.*, 11(5):1155–1182, 1999.
- [58] M. Brand and V. Kettner. Discovery and segmentation of activities in video. *IEEE T. Pattern Analysis and Machine Intelligence*, 22(8):844–851, August 2000.
- [59] M. Brand, N. Oliver, and A. Pentland. Coupled hidden markov models for complex action recognition. In *IEEE Conf. on Computer Vision and Pattern Recognition*, pages 994–999, 1997.
- [60] C. Bregler and J. Malik. Tracking people with twists and exponential maps. In *IEEE Conf. on Computer Vision and Pattern Recognition*, pages 8–15, 1998.
- [61] C. Bregler, J. Malik, and K. Pullen. Twist based acquisition and tracking of animal and human kinematics. *Int. J. Computer Vision*, 56(3):179–194, February 2004.
- [62] A. Broggi, M. Bertozzi, A. Fascioli, and M. Sechi. Shape-based pedestrian detection. In *Proc. IEEE Intelligent Vehicles Symposium*, pages 215–220, 2000.
- [63] L. Brown. *A Radar History of World War II: Technical and Military Imperatives*. Institute of Physics Press, 2000.
- [64] A. Bruderlin and L. Williams. Motion signal processing. In *SIGGRAPH '95: Proceedings of the 22nd annual conference on Computer graphics and interactive techniques*, pages 97–104, New York, NY, USA, 1995. ACM Press.
- [65] R. Buder. *The Invention that Changed the World*. Touchstone Press, 1998. reprint.
- [66] M. Burl, T. Leung, and P. Perona. Face localisation via shape statistics. In *Int. Workshop on Automatic Face and Gesture Recognition*, 1995.
- [67] Q. Cai and J. K. Aggarwal. Automatic tracking of human motion in indoor scenes across multiple synchronized video streams. In *ICCV '98: Proceedings of the Sixth*

- International Conference on Computer Vision*, page 356, Washington, DC, USA, 1998. IEEE Computer Society.
- [68] B. Calais-Germain. *Anatomy of Movement*. Eastland Press, 1993.
- [69] L. W. Campbell and A. F. Bobick. Recognition of human body motion using phase space constraints. In *ICCV*, pages 624–630, 1995.
- [70] M. Cardle, M. Vlachos, S. Brooks, E. Keogh, and D. Gunopulos. Fast motion capture matching with replicated motion editing. In *Proceedings of SIGGRAPH 2003 - Sketches and Applications*, 2003.
- [71] S. Carlsson. Recognizing walking people. In *European Conference on Computer Vision*, pages I: 472–486, 2000.
- [72] J. Carranza, C. Theobalt, M. A. Magnor, and H.-P. Seidel. Free-viewpoint video of human actors. *ACM Trans. Graph.*, 22(3):569–577, 2003.
- [73] A. Cavallaro and T. Ebrahimi. Video object extraction based on adaptive background and statistical change detection. In *SPIE*, pages 465–475, 4310.
- [74] J. Chai and J. K. Hodgins. Performance animation from low-dimensional control signals. *ACM Trans. Graph.*, 24(3):686–696, 2005.
- [75] T. Cham and J. Rehg. A multiple hypothesis approach to figure tracking. In *IEEE Conf. on Computer Vision and Pattern Recognition*, pages II: 239–245, 1999.
- [76] N. Checka, K. Wilson, M. Siracusa, and T. Darrell. Multiple person and speaker activity tracking with a particle filter. In *IEEE Conf. on Acoustics, Speech, and Signal Processing*, volume 5, pages 881–4, 2004.
- [77] F. Cheng, W. Christmas, and J. Kittler. Periodic human motion description for sports video databases. In *Proceedings IAPR International Conference on Pattern Recognition*, pages III: 870–873, 2004.
- [78] K. M. Cheung, T. Kanade, J.-Y. Bouguet, and M. Holler. A real time system for robust 3d voxel reconstruction of human motions. In *Proceedings of the 2000 IEEE Conference on Computer Vision and Pattern Recognition (CVPR '00)*, volume 2, pages 714 – 720, June 2000.
- [79] K.-M. G. Cheung, S. Baker, and T. Kanade. Shape-from-silhouette across time part i: Theory and algorithms. *Int. J. Comput. Vision*, 62(3):221–247, 2005.
- [80] K.-M. G. Cheung, S. Baker, and T. Kanade. Shape-from-silhouette across time part ii: Applications to human modeling and markerless motion tracking. *Int. J. Comput. Vision*, 63(3):225–245, 2005.
- [81] J. chi Wu and Z. Popović. Realistic modeling of bird flight animations. *ACM Trans. Graph.*, 22(3):888–895, 2003.
- [82] K. Choo and D. Fleet. People tracking using hybrid monte carlo filtering. In *Int. Conf. on Computer Vision*, pages II: 321–328, 2001.
- [83] M. F. Cohen. Interactive spacetime control for animation. In *SIGGRAPH '92: Proceedings of the 19th annual conference on Computer graphics and interactive techniques*, pages 293–302, New York, NY, USA, 1992. ACM Press.
- [84] D. Comaniciu and P. Meer. Distribution free decomposition of multivariate data. *Pattern analysis and applications*, 2(1):22–30, 1999.
- [85] C. Cortes and V. Vapnik. Support-vector networks. *Machine Learning*, 20(3):273–97, 1995.
- [86] L. Crawford and S. Sastry. Biological motor control approaches for a planar diver. In *IEEE Conf. on Decision and Control*, pages 3881–3886, 1995.

- [87] C. Curio, J. Edelbrunner, T. Kalinke, C. Tzomakas, and W. von Seelen. Walking pedestrian recognition. *Intelligent Transportation Systems*, 1(3):155–163, September 2000.
- [88] R. Cutler and L. Davis. View-based detection and analysis of periodic motion. In *Proceedings IAPR International Conference on Pattern Recognition*, pages Vol I: 495–500, 1998.
- [89] R. Cutler and L. Davis. Real-time periodic motion detection, analysis, and applications. In *IEEE Conf. on Computer Vision and Pattern Recognition*, pages II: 326–332, 1999.
- [90] R. Cutler and L. Davis. Robust periodic motion and motion symmetry detection. In *IEEE Conf. on Computer Vision and Pattern Recognition*, pages II: 615–622, 2000.
- [91] R. Cutler and L. Davis. Robust real-time periodic motion detection, analysis, and applications. *IEEE T. Pattern Analysis and Machine Intelligence*, 22(8):781–796, August 2000.
- [92] J. E. Cutting and L. T. Kozlowski. Recognizing friends by their walk: Gait perception without familiarity cues. *Bulletin of the Psychonomic Society*, 9(5):353–356, 1977.
- [93] J. E. Cutting, D. R. Proffitt, and L. T. Kozlowski. A biomechanical invariant for gait perception. *Journal of Experimental Psychology: Human Perception and Performance*, 4(3):357–372, 1978.
- [94] N. Dalal and B. Triggs. Histograms of oriented gradients for human detection. In *IEEE Conf. on Computer Vision and Pattern Recognition*, pages I: 886–893, 2005.
- [95] T. Darrell, G. Gordon, M. Harville, and J. Woodfill. Integrated person tracking using stereo, color, and pattern detection. *Int. J. Computer Vision*, 37(2):175–185, June 2000.
- [96] J. Darroch and D. Ratcliff. Generalized iterative scaling for log-linear models. *Ann. Math. Statistics*, 43:1470–1480, 1972.
- [97] A. Dasgupta and Y. Nakamura. Making feasible walking motion of humanoid robots from human motion capture data. In *1999 IEEE International Conference on Robotics & Automation*, pages 1044–1049, 1999.
- [98] F. Daum. Beyond kalman filters: practical design of nonlinear filters. In *Proc. SPIE*, volume 2561, pages 252–262, 1995.
- [99] F. Daum. Exact finite dimensional nonlinear filters. *IEEE. Trans. Automatic Control*, 31:616–622, 1995.
- [100] J. Davis and A. Bobick. The representation and recognition of action using temporal templates. In *IEEE Conf. on Computer Vision and Pattern Recognition*, pages 928–934, 1997.
- [101] Q. Delamarre and O. Faugeras. 3d articulated models and multi-view tracking with silhouettes. In *ICCV '99: Proceedings of the International Conference on Computer Vision-Volume 2*, page 716, Washington, DC, USA, 1999. IEEE Computer Society.
- [102] Q. Delamarre and O. Faugeras. 3d articulated models and multiview tracking with physical forces. *Comput. Vis. Image Underst.*, 81(3):328–357, 2001.
- [103] A. S. Deo and I. D. Walker. Minimum effort inverse kinematics for redundant manipulators. *IEEE Transactions on Robotics and Automation*, 13(6), 1997.
- [104] J. Deutscher, A. Blake, and I. Reid. Articulated body motion capture by annealed particle filtering. In *IEEE Conf. on Computer Vision and Pattern Recognition*, pages II: 126–133, 2000.
- [105] J. Deutscher, A. Davison, and I. Reid. Automatic partitioning of high dimensional search spaces associated with articulated body motion capture. In *IEEE Conf. on Computer Vision and Pattern Recognition*, pages II:669–676, 2001.

- [106] J. Deutscher, B. North, B. Bascle, and A. Blake. Tracking through singularities and discontinuities by random sampling. In *Int. Conf. on Computer Vision*, pages 1144–1149, 1999.
- [107] J. Deutscher and I. Reid. Articulated body motion capture by stochastic search. *Int. J. Computer Vision*, 61(2):185–205, February 2005.
- [108] M. Dimitrijevic, V. Lepetit, and P. Fua. Human body pose recognition using spatio-temporal templates. In *ICCV workshop on Modeling People and Human Interaction*, 2005.
- [109] A. Doucet, N. D. Freitas, and N. Gordon. *Sequential Monte Carlo Methods in Practice*. Springer-Verlag, 2001.
- [110] T. Drummond and R. Cipolla. Real-time tracking of complex structures with on-line camera calibration. In T. P. Pridmore and D. Elliman, editors, *Proceedings of the British Machine Vision Conference 1999, BMVC 1999, Nottingham, 13-16 September 1999*, 1999.
- [111] T. Drummond and R. Cipolla. Real-time tracking of multiple articulated structures in multiple views. In *ECCV '00: Proceedings of the 6th European Conference on Computer Vision-Part II*, pages 20–36, London, UK, 2000. Springer-Verlag.
- [112] T. Drummond and R. Cipolla. Real-time tracking of highly articulated structures in the presence of noisy measurements. In *ICCV*, pages 315–320, 2001.
- [113] T. Drummond and R. Cipolla. Real-time visual tracking of complex structures. *IEEE T. Pattern Analysis and Machine Intelligence*, 24(7):932–946, July 2002.
- [114] A. D’Souza, S. Vijayakumar, and S. Schaal. Learning inverse kinematics. In *Int. Conf. Intelligent Robots and Systems*, pages 298–303, 2001.
- [115] S. Duane, A. Kennedy, B. Pendleton, and D. Roweth. Hybrid monte carlo. *Physics Letters B*, 195:216–222, 1987.
- [116] A. A. Efros, A. C. Berg, G. Mori, and J. Malik. Recognizing action at a distance. In *ICCV '03: Proceedings of the Ninth IEEE International Conference on Computer Vision*, pages 726–733, Washington, DC, USA, 2003. IEEE Computer Society.
- [117] A. Engin and S. Tumer. Three-dimensional kinematic modelling of the human shoulder complex - part i: Physical model and determination of joint sinus cones. *ASME Journal of Biomechanical Engineering*, 111:107–112, 1989.
- [118] P. Faloutsos, M. van de Panne, and D. Terzopoulos. Composable controllers for physics-based character animation. In *Proceedings of ACM SIGGRAPH 2001*, Computer Graphics Proceedings, Annual Conference Series, pages 251–260, Aug. 2001.
- [119] P. Faloutsos, M. van de Panne, and D. Terzopoulos. The virtual stuntman: dynamic characters with a repertoire of autonomous motor skills. *Computers & Graphics*, 25(6):933–953, Dec. 2001.
- [120] A. C. Fang and N. S. Pollard. Efficient synthesis of physically valid human motion. *ACM Transactions on Graphics*, 22(3):417–426, July 2003.
- [121] A. C. Fang and N. S. Pollard. Efficient synthesis of physically valid human motion. *ACM Trans. Graph.*, 22(3):417–426, 2003.
- [122] A. Farina, D. Benvenuti, and B. Ristic. A comparative study of the benes filtering problem. *Signal Processing*, 82:133–147, 2002.
- [123] A. Faul and M. Tipping. Analysis of sparse bayesian learning. In *Advances in Neural Information Processing Systems 14*, pages 383–389. MIT Press, 2002.

- [124] P. Felzenszwalb and D. Huttenlocher. Efficient matching of pictorial structures. In *IEEE Conf. on Computer Vision and Pattern Recognition*, pages II: 66–73, 2000.
- [125] P. Felzenszwalb and D. Huttenlocher. Pictorial structures for object recognition. *Int. J. Computer Vision*, 61(1):55–79, January 2005.
- [126] X. Feng and P. Perona. Human action recognition by sequence of wavelet codewords. In *3D Data Processing Visualization and Transmission, 2002. Proceedings. First International Symposium on*, pages 717–721, 2002.
- [127] R. Fergus, P. Perona, and A. Zisserman. Object class recognition by unsupervised scale-invariant learning. In *IEEE Conf. on Computer Vision and Pattern Recognition*, pages II: 264–271, 2003.
- [128] A. Fod, M. J. Matarić, and O. C. Jenkins. Automated derivation of primitives for movement classification. *Auton. Robots*, 12(1):39–54, 2002.
- [129] K. Forbes and E. Fiume. An efficient search algorithm for motion data using weighted pca. In *Symposium on Computer Animation*, 2005.
- [130] D. Forsyth. Sampling, resampling and colour constancy. In *IEEE Conf. on Computer Vision and Pattern Recognition*, pages I: 300–305, 1999.
- [131] D. Forsyth and M. Fleck. Body plans. In *IEEE Conf. on Computer Vision and Pattern Recognition*, pages 678–683, 1997.
- [132] D. Forsyth and M. Fleck. Automatic detection of human nudes. *Int. J. Computer Vision*, 32(1):63–77, August 1999.
- [133] D. Forsyth, M. Fleck, and C. Bregler. Finding naked people. In *European Conference on Computer Vision*, pages 593–602, 1996.
- [134] D. Forsyth, J. Haddon, and S. Ioffe. The joy of sampling. *Int. J. Computer Vision*, 2001.
- [135] D. Forsyth and J. Ponce. *Computer Vision: a modern approach*. Prentice-Hall, 2002.
- [136] L. Fradet, M. Botcazou, C. Durocher, A. Cretual, F. Multon, J. Prioux, and P. Delamarche. Do handball throws always exhibit a proximal-to-distal segmental sequence? *Journal of Sports Sciences*, 22(5):439–447, 2004.
- [137] J. Freyd. Dynamic mental representations. *Psychological Review*, 94(4):427–438, 1987.
- [138] A. Galata, N. Johnson, and D. Hogg. Learning behaviour models of human activities. In *British Machine Vision Conference*, page Modelling Human Behaviour, 1999.
- [139] A. Galata, N. Johnson, and D. Hogg. Learning structured behavior models using variable length markov models. In *IEEE Workshop on Modelling People*, 1999.
- [140] D. Gao, J. Zhou, and L. Xin. A novel algorithm of adaptive background estimation. In *IEEE Int. Conf. Image Processing*, pages II: 395–398, 2001.
- [141] W. Gao, J. Ma, X. Chen, et al. Handtalker: a multimodal dialog system using sign language and 3d virtual human. In *Proc. Third Int. Conf. Multimodal Interface*, pages 564–571, 2000.
- [142] D. Gavrila. Pedestrian detection from a moving vehicle. In *European Conference on Computer Vision*, pages II: 37–49, 2000.
- [143] D. Gavrila. Pedestrian detection from a moving vehicle. In *European Conference on Computer Vision*, pages II: 37–49, 2000.
- [144] D. Gavrila. Sensor-based pedestrian protection. *Intelligent Transportation Systems*, pages 77–81, 2001.

- [145] D. Gavrilu, J. Giebel, and S. Munder. Vision-based pedestrian detection: the protector system. In *Intelligent Vehicle Symposium*, pages 13–18, 2004.
- [146] D. M. Gavrilu. The visual analysis of human movement: A survey. *Computer Vision and Image Understanding: CVIU*, 73(1):82–98, 1999.
- [147] A. Gelb. *Applied Optimal Estimation*. MIT Press, 1974. written together with Staff of the Analytical Sciences Corporation.
- [148] J. Gibson. *The Perception of the Visual World*. Houghton Mifflin, 1955.
- [149] W. Gilks, S. Richardson, and D. Spiegelhalter, editors. *Markov chain Monte Carlo in practice*. Chapman and Hall, 1996.
- [150] A. Gionis, P. Indyk, and R. Motwani. Similarity search in high dimensions via hashing. In *25th International Conference on Very Large Data Bases (VLDB 99)*, pages 518–529, 1999.
- [151] M. Girard. Interactive design of 3-d computer-animated legged animal motion. In *SI3D '86: Proceedings of the 1986 workshop on Interactive 3D graphics*, pages 131–150, New York, NY, USA, 1987. ACM Press.
- [152] M. Girard and A. A. Maciejewski. Computational modeling for the computer animation of legged figures. In *SIGGRAPH '85: Proceedings of the 12th annual conference on Computer graphics and interactive techniques*, pages 263–270, New York, NY, USA, 1985. ACM Press.
- [153] M. Gleicher. Motion editing with spacetime constraints. In *Proceedings of the 1997 Symposium on Interactive 3D Graphics*, 1997.
- [154] M. Gleicher. Animation from observation: Motion capture and motion editing. *Computer Graphics*, 1999.
- [155] M. Gleicher. Animation from observation: Motion capture and motion editing. *SIGGRAPH Comput. Graph.*, 33(4):51–54, 2000.
- [156] M. Gleicher. Comparing constraint-based motion editing methods. *Graphical Models*, 2001.
- [157] M. Gleicher and N. Ferrier. Evaluating video-based motion capture. In *CA '02: Proceedings of the Computer Animation*, page 75, Washington, DC, USA, 2002. IEEE Computer Society.
- [158] M. Gleicher, H. J. Shin, L. Kovar, and A. Jepsen. Snap-together motion: assembling run-time animations. In *SI3D '03: Proceedings of the 2003 symposium on Interactive 3D graphics*, pages 181–188, New York, NY, USA, 2003. ACM Press.
- [159] R. Goldenberg, R. Kimmel, E. Rivlin, and M. Rudzsky. ”dynamism of a dog on a leash” or behavior classification by eigen-decomposition of periodic motions. In *European Conference on Computer Vision*, page I: 461 ff., 2002.
- [160] R. Goldenberg, R. Kimmel, E. Rivlin, and M. Rudzsky. Behavior classification by eigendecomposition of periodic motions. *Pattern Recognition*, 38(7):1033–1043, July 2005.
- [161] H. Goldstein. *Classical Mechanics*. Addison Wesley, Reading, MA, 1950.
- [162] K. Grauman, G. Shakhnarovich, and T. Darrell. Virtual visual hulls: Example-based 3d shape inference from silhouettes. In *SMVP04*, pages 26–37, 2004.
- [163] W. Grimson, L. Lee, R. Romano, and C. Stauffer. Using adaptive tracking to classify and monitor activities in a site. In *IEEE Conf. on Computer Vision and Pattern Recognition*, pages 22–29, 1998.

- [164] K. Grochow, S. L. Martin, A. Hertzmann, and Z. Popović. Style-based inverse kinematics. *ACM Trans. Graph.*, 23(3):522–531, 2004.
- [165] R. Grzeszczuk, D. Terzopoulos, and G. Hinton. Neuroanimator: Fast neural network emulation and control of physics-based models. In *Proceedings of SIGGRAPH 98*, Computer Graphics Proceedings, Annual Conference Series, pages 9–20, July 1998.
- [166] J. K. Hahn. Realistic animation of rigid bodies. In *SIGGRAPH '88: Proceedings of the 15th annual conference on Computer graphics and interactive techniques*, pages 299–308, New York, NY, USA, 1988. ACM Press.
- [167] I. Haritaoglu, D. Harwood, and L. Davis. W4s: A real-time system for detecting and tracking people in 2 1/2-d. In *European Conference on Computer Vision*, page I: 877, 1998.
- [168] I. Haritaoglu, D. Harwood, and L. S. Davis. W4: Real-time surveillance of people and their activities. *IEEE T. Pattern Analysis and Machine Intelligence*, 22:809–830, 2000.
- [169] T. Hastie, R. Tibshirani, and J. Friedman. *The Elements of Statistical Learning: Data Mining, Inference and Prediction*. Springer Verlag, 2001.
- [170] T. Heap and D. Hogg. Wormholes in shape space: Tracking through discontinuous changes in shape. In *Int. Conf. on Computer Vision*, pages 344–349, 1998.
- [171] A. Hilton and J. Starck. Multiple view reconstruction of people. In *2nd International Symposium on 3D Data Processing, Visualization and Transmission (3DPVT 2004)*, 6-9 September 2004, Thessaloniki, Greece, pages 357–364, 2004.
- [172] A. C. Hindmarsh and L. R. Petzold. Algorithms and software for ordinary differential equations and differential-algebraic equations, part i: Euler methods and error estimation. *Comput. Phys.*, 9(1):34–41, 1995.
- [173] A. C. Hindmarsh and L. R. Petzold. Algorithms and software for ordinary differential equations and differential-algebraic equations, part ii: higher-order methods and software packages. *Comput. Phys.*, 9(2):148–155, 1995.
- [174] G. Hinton. Relaxation and its role in vision. Technical report, University of Edinburgh, 1978. PhD Thesis.
- [175] D. C. Hoaglin, F. Mosteller, and J. W. Tukey, editors. *Understanding Robust and Exploratory Data Analysis*. John Wiley, 1983.
- [176] J. Hodgins, J. O'Brien, and R. Bodenheimer. Computer animation. In J. G. Webster, editor, *Wiley Encyclopedia of Electrical and Electronics Engineering*, volume 3, pages 686–690, 1999.
- [177] J. K. Hodgins, J. F. O'Brien, and J. Tumblin. Do geometric models affect judgments of human motion? In *Graphics Interface '97*, pages 17–25, May 1997.
- [178] J. K. Hodgins, J. F. O'Brien, and J. Tumblin. Perception of human motion with different geometric models. *IEEE Transactions on Visualization and Computer Graphics*, 4(4):307–316, Oct. 1998.
- [179] J. K. Hodgins and N. S. Pollard. Adapting simulated behaviors for new characters. In *SIGGRAPH '97: Proceedings of the 24th annual conference on Computer graphics and interactive techniques*, pages 153–162, New York, NY, USA, 1997. ACM Press/Addison-Wesley Publishing Co.
- [180] J. K. Hodgins, W. L. Wooten, D. C. Brogan, and J. F. O'Brien. Animating human athletics. In *Proceedings of SIGGRAPH 95*, Computer Graphics Proceedings, Annual Conference Series, pages 71–78, Aug. 1995.

- [181] D. Hogg. Model-based vision: A program to see a walking person. *Image and Vision Computing*, 1(1):5–20, 1983.
- [182] P. Hong, M. Turk, and T. Huang. Gesture modeling and recognition using finite state machines. In *Int. Conf. Automatic Face and Gesture Recognition*, pages 410–415, 2000.
- [183] S. Hongeng, R. Nevatia, and F. Bremond. Video-based event recognition: activity representation and probabilistic recognition methods. *Computer Vision and Image Understanding*, 96(2):129–162, November 2004.
- [184] B. Horn. Closed form solutions of absolute orientation using orthonormal matrices. *J. Opt. Soc. America - A*, 5(7):1127–1135, 1987.
- [185] B. Horn. Closed form solutions of absolute orientation using unit quaternions. *J. Opt. Soc. America - A*, 4(4):629–642, April 1987.
- [186] B. Horn and B. Schunck. Determining optical flow. *Artificial Intelligence*, 17(1-3):185–203, August 1981.
- [187] N. Howe. Silhouette lookup for automatic pose tracking. In *IEEE Workshop on Articulated and Non-Rigid Motion*, page 15, 2004.
- [188] N. R. Howe, M. E. Leventon, and W. T. Freeman. Bayesian reconstruction of 3d human motion from single-camera video. In S. Solla, T. Leen, and K.-R. Müller, editors, *Advances in Neural Information Processing Systems 12*, pages 820–26. MIT Press, 2000.
- [189] W. Hu, T. Tan, L. Wang, and S. Maybank. A survey on visual surveillance of object motion and behaviors. *IEEE transactions on systems, man, and cybernetics part c: applications and reviews*, 34(3), 2004.
- [190] W. Hu, T. Tan, L. Wang, and S. Maybank. A survey on visual surveillance of object motion and behaviors. *IEEE Trans. Systems, Man and Cybernetics - Part C: Applications and Reviews*, 34(3):334–352, 2004.
- [191] D. Huttenlocher, J. Noh, and W. Rucklidge. Tracking non-rigid objects in complex scenes. In *Int. Conf. on Computer Vision*, pages 93–101, 1993.
- [192] L. Ikemoto, O. Arikani, and D. Forsyth. Knowing when to put your foot down. In *Proc Symp. Interactive 3D graphics and Games*, 2006.
- [193] L. Ikemoto, O. Arikani, and D. Forsyth. Quick motion transitions with cached multi-way blends. Technical Report UCB/EECS-2006-14, EECS Department, University of California, Berkeley, February 13 2006.
- [194] L. Ikemoto and D. Forsyth. Enriching a motion collection by transplanting limbs. In *Proc. Symposium on Computer Animation*, 2004.
- [195] P. Indyk and R. Motwani. Approximate nearest neighbors: towards removing the curse of dimensionality. In *STOC '98: Proceedings of the thirtieth annual ACM symposium on Theory of computing*, pages 604–613, New York, NY, USA, 1998. ACM Press.
- [196] P. Indyk, R. Motwani, P. Raghavan, and S. Vempala. Locality-preserving hashing in multidimensional spaces. In *STOC '97: Proceedings of the twenty-ninth annual ACM symposium on Theory of computing*, pages 618–625, New York, NY, USA, 1997. ACM Press.
- [197] S. Ioffe and D. Forsyth. Learning to find pictures of people. In *Proc. Neural Information Processing Systems*, 1998.
- [198] S. Ioffe and D. Forsyth. Human tracking with mixtures of trees. In *Int. Conf. on Computer Vision*, pages I: 690–695, 2001.

- [199] S. Ioffe and D. Forsyth. Mixtures of trees for object recognition. In *IEEE Conf. on Computer Vision and Pattern Recognition*, 2001.
- [200] S. Ioffe and D. Forsyth. Probabilistic methods for finding people. *Int. J. Computer Vision*, 43(1):45–68, June 2001.
- [201] M. Isard and A. Blake. Icondensation: Unifying low-level and high-level tracking in a stochastic framework. In *European Conference on Computer Vision*, page I: 893, 1998.
- [202] Y. Ivanov, A. Bobick, and J. Liu. Fast lighting independent background subtraction. In *In Proc. of the IEEE Workshop on Visual Surveillance – VS’98*, pages 49–55, 1998.
- [203] Y. Ivanov, A. Bobick, and J. Liu. Fast lighting independent background subtraction. *Int. J. Computer Vision*, 37(2):199–207, June 2000.
- [204] O. Javed and M. Shah. Tracking and object classification for automated surveillance. In *European Conference on Computer Vision*, page IV: 343 ff., 2002.
- [205] F. Jelinek. *Statistical Methods for Speech Recognition*. MIT Press, Cambridge, MA, 1998.
- [206] O. C. Jenkins and M. J. Matarić. A spatio-temporal extension to isomap non-linear dimension reduction. In *ICML ’04: Proceedings of the twenty-first international conference on Machine learning*, page 56, New York, NY, USA, 2004. ACM Press.
- [207] O. C. Jenkins and M. J. Mataric. Automated derivation of behavior vocabularies for autonomous humanoid motion. In *AAMAS ’03: Proceedings of the second international joint conference on Autonomous agents and multiagent systems*, pages 225–232, New York, NY, USA, 2003. ACM Press.
- [208] F. V. Jensen. *An Introduction to Bayesian Networks*. UCL Press, London, 1996.
- [209] C. Jeong, J. Kim, and K. Hong. Appearance-based nude image detection. In *Proceedings IAPR International Conference on Pattern Recognition*, pages IV: 467–470, 2004.
- [210] R. Jin, R. Yan, J. Zhang, and A. Hauptmann. A faster iterative scaling algorithm for conditional exponential models. In *Proc. International Conference on Machine learning*, 2003.
- [211] G. Johansson. Visual perception of biological motion and a model for its analysis. *Perception & Psychophysics*, 14(2):201–211, 1973.
- [212] I. Joliffe. *Principal Component Analysis*. Springer-Verlag, 2002.
- [213] M. Jones and P. Viola. Face recognition using boosted local features. In *IEEE International Conference on Computer Vision (ICCV)*, 2003.
- [214] M. Jones and P. Viola. Fast multi-view face detection. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2003.
- [215] R. V. Jones. *Most Secret War*. Wordsworth Military Library, 1998. reprint.
- [216] S. X. Ju, M. J. Black, and Y. Yacoob. Cardboard people: A parameterized model of articulated image motion. In *Proc. Int. Conference on Face and Gesture*, pages 561–567, 1996.
- [217] D. Jurafsky and J. Martin. *Speech and Language Processing: An Introduction to Natural Language Processing, Computational Linguistics and Speech Recognition*. Prentice-Hall, 2000.
- [218] M. Kass, A. Witkin, and D. Terzopoulos. Snakes: Active contour models. In *Int. Conf. on Computer Vision*, pages 259–268, 1987.
- [219] M. Kass, A. Witkin, and D. Terzopoulos. Snakes: Active contour models. *Int. J. Computer Vision*, 1(4):321–331, January 1988.

- [220] Y. Ke and R. Sukthankar. Pca-sift: a more distinctive representation for local image descriptors. In *IEEE Conf. on Computer Vision and Pattern Recognition*, pages II: 506–513, 2004.
- [221] R. Kehl, M. Bray, and L. V. Gool. Full body tracking from multiple views using stochastic sampling. In *CVPR '05: Proceedings of the 2005 IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR'05) - Volume 2*, pages 129–136, Washington, DC, USA, 2005. IEEE Computer Society.
- [222] E. J. Keogh. Exact indexing of dynamic time warping. In *VLDB*, pages 406–417, 2002.
- [223] E. J. Keogh. Efficiently finding arbitrarily scaled patterns in massive time series databases. In *7th European Conference on Principles and Practice of Knowledge Discovery in Databases*, pages 253–265, 2003.
- [224] E. J. Keogh, T. Palpanas, V. B. Zordan, D. Gunopulos, and M. Cardle. Indexing large human-motion databases. In *Proc. 30th VLDB Conf.*, pages 780–791, 2004.
- [225] V. Kettner and M. Brand. Minimum-entropy models of scene activity. In *IEEE Conf. on Computer Vision and Pattern Recognition*, pages I:281–286, 1999.
- [226] V. Kettner and R. Zabih. Counting people from multiple cameras. In *ICMCS '99: Proceedings of the IEEE International Conference on Multimedia Computing and Systems Volume II-Volume 2*, page 267, Washington, DC, USA, 1999. IEEE Computer Society.
- [227] K. Grobel and M. Assan. Isolated sign language recognition using hidden markov models. In *Proc. Int. Conf. System Man and Cybernetics*, pages 162–167, 1997.
- [228] J. Kim, W. Jang, and Z. Bien. A dynamic gesture recognition system for the korean sign language (ksl). *Systems, Man and Cybernetics-B*, 26(2):354–359, April 1996.
- [229] D. King. Generating vertical velocity and angular momentum during skating jumps. In *23rd Annual Meeting of the American Society of Biomechanics*, 1999.
- [230] A. Kirk, J. O'Brien, and D. Forsyth. Skeletal parameter estimation from optical motion capture data. In *IEEE Conf. on Computer Vision and Pattern Recognition*, 2005.
- [231] H. Ko and N. Badler. Animating human locomotion with inverse dynamics. *IEEE Computer Graphics and Application*, 16(2):50–59, 1996.
- [232] K. Kondo. Inverse kinematics of a human arm. Technical report, Stanford University, Stanford, CA, USA, 1994.
- [233] J. U. Korein and N. I. Badler. Techniques for generating the goal-directed motion of articulated structures. *IEEE Computer Graphics and Applications*, pages 71–81, 1982.
- [234] L. Kovar and M. Gleicher. Flexible automatic motion blending with registration curves. In *SCA '03: Proceedings of the 2003 ACM SIGGRAPH/Eurographics symposium on Computer animation*, pages 214–224, Aire-la-Ville, Switzerland, Switzerland, 2003. Eurographics Association.
- [235] L. Kovar and M. Gleicher. Automated extraction and parameterization of motions in large data sets. *ACM Trans. Graph.*, 23(3):559–568, 2004.
- [236] L. Kovar, M. Gleicher, and F. Pighin. Motion graphs. In *Proceedings of the 29th annual conference on Computer graphics and interactive techniques*, pages 473–482. ACM Press, 2002.
- [237] L. Kovar, J. Schreiner, and M. Gleicher. Footskate cleanup for motion capture editing. In *Proceedings of the 2002 ACM SIGGRAPH/Eurographics symposium on Computer animation*, pages 97–104. ACM Press, 2002.

- [238] L. T. Kozlowski and J. E. Cutting. Recognizing the sex of a walker from a dynamic point-light display. *Perception & Psychophysics*, 21(6):575–580, 1977.
- [239] L. T. Kozlowski and J. E. Cutting. Recognizing the gender of walkers from point-lights mounted on ankles: Some second thoughts. *Perception & Psychophysics*, 23(5):459, 1978.
- [240] M. P. Kumar, P. H. S. Torr, and A. Zisserman. Extending pictorial structures for object recognition. In *Proceedings of the British Machine Vision Conference*, 2004.
- [241] T. Kwon and S. Y. Shin. Motion modeling for on-line locomotion synthesis. In *SCA '05: Proceedings of the 2005 ACM SIGGRAPH/Eurographics symposium on Computer animation*, pages 29–38, New York, NY, USA, 2005. ACM Press.
- [242] Y. Lecun, L. Bottou, Y. Bengio, and P. Haffner. Gradient-based learning applied to document recognition. *Proceedings of the IEEE*, 86(11):2278–2324, 1998.
- [243] J. Lee, J. Chai, P. Reitsma, J. Hodgins, and N. Pollard. Interactive control of avatars animated with human motion data. In *Proceedings of SIGGRAPH 95*, 2002.
- [244] J. Lee and S. Y. Shin. A hierarchical approach to interactive motion editing for human-like figures. In *Proceedings of the 26th annual conference on Computer graphics and interactive techniques*, pages 39–48. ACM Press/Addison-Wesley Publishing Co., 1999.
- [245] M. Lee and I. Cohen. Human upper body pose estimation in static images. In *European Conference on Computer Vision*, pages Vol II: 126–138, 2004.
- [246] M. Lee and I. Cohen. Proposal maps driven mcmc for estimating human body pose in static images. In *IEEE Conf. on Computer Vision and Pattern Recognition*, pages II: 334–341, 2004.
- [247] M. Lee and R. Nevatia. Dynamic human pose estimation using markov chain monte carlo approach. In *IEEE Workshop on Motion and Video Computing*, pages 168–175, 2005.
- [248] B. Leibe, A. Leonardis, and B. Schiele. Combined object categorization and segmentation with an implicit shape model. In *ECCV-04 Workshop on Stat. Learn. in Comp. Vis.*, pages 17–32, 2004.
- [249] B. Leibe, E. Seemann, and B. Schiele. Pedestrian detection in crowded scenes. In *IEEE Conf. on Computer Vision and Pattern Recognition*, pages I: 878–885, 2005.
- [250] A. Leonardis, A. Gupta, and R. Bajcsy. Segmentation of range images as the search for geometric parametric models. *Int. J. Computer Vision*, 14(3):253–277, April 1995.
- [251] T. Leung, M. Burl, and P. Perona. Finding faces in cluttered scenes using random labelled graph matching. In *Int. Conf. on Computer Vision*, 1995.
- [252] B. Li and H. Holstein. Recognition of human periodic motion: A frequency domain approach. In *Proceedings IAPR International Conference on Pattern Recognition*, pages I: 311–314, 2002.
- [253] Y. Li, T. Wang, and H.-Y. Shum. Motion texture: a two-level statistical model for character motion synthesis. In *Proceedings of the 29th annual conference on Computer graphics and interactive techniques*, pages 465–472. ACM Press, 2002.
- [254] J. Little and J. Boyd. Describing motion for recognition. In *International Symposium on Computer Vision*, pages 235–240, 1995.
- [255] J. Little and J. Boyd. Recognizing people by their gait: The shape of motion. *Videre*, 1(2), 1998.

- [256] J. Little and J. Boyd. Shape of motion and the perception of human gaits. In *IEEE Workshop on Empirical Evaluation Methods in Computer Vision*, 1998.
- [257] C. Liu, S. Zhu, and H. Shum. Learning inhomogeneous gibbs model of faces by mini-max entropy. In *Int. Conf. on Computer Vision*, pages I: 281–287, 2001.
- [258] C. K. Liu, A. Hertzmann, and Z. Popović. Learning physics-based motion style with nonlinear inverse optimization. *ACM Trans. Graph.*, 24(3):1071–1081, 2005.
- [259] C. K. Liu and Z. Popović. Synthesis of complex dynamic character motion from simple animations. In *SIGGRAPH '02: Proceedings of the 29th annual conference on Computer graphics and interactive techniques*, pages 408–416, New York, NY, USA, 2002. ACM Press.
- [260] F. Liu and R. Picard. Detecting and segmenting periodic motion. Media lab vision and modelling tr-400, MIT, 1996.
- [261] F. Liu and R. Picard. Finding periodicity in space and time. In *Int. Conf. on Computer Vision*, pages 376–383, 1998.
- [262] J. S. Liu. *Monte Carlo strategies in scientific computing*. Springer, 2001.
- [263] Z. Liu, H. Chen, and H. Shum. An efficient approach to learning inhomogeneous gibbs model. In *IEEE Conf. on Computer Vision and Pattern Recognition*, pages I: 425–431, 2003.
- [264] Z. Liu and M. F. Cohen. Decomposition of linked figure motion: Diving. In *5th Eurographics Workshop on Animation and Simulation*, 1994.
- [265] Z. Liu, S. J. Gortler, and M. F. Cohen. Hierarchical spacetime control. In *SIGGRAPH '94: Proceedings of the 21st annual conference on Computer graphics and interactive techniques*, pages 35–42, New York, NY, USA, 1994. ACM Press.
- [266] M. Liverman. *The Animator's Motion Capture Guide : Organizing, Managing, Editing*. Charles River Media, 2004.
- [267] D. Lowe. Distinctive image features from scale-invariant keypoints. *Int. J. Computer Vision*, 60(2):91–110, November 2004.
- [268] G. Loy, M. Eriksson, J. Sullivan, and S. Carlsson. Monocular 3d reconstruction of human motion in long action sequences. In *European Conference on Computer Vision*, pages Vol IV: 442–455, 2004.
- [269] J. MacCormick and A. Blake. A probabilistic exclusion principle for tracking multiple objects. In *Int. Conf. on Computer Vision*, pages 572–578, 1999.
- [270] J. MacCormick and A. Blake. A probabilistic exclusion principle for tracking multiple objects. *Int. J. Computer Vision*, 39(1):57–71, August 2000.
- [271] J. MacCormick and M. Isard. Partitioned sampling, articulated objects, and interface-quality hand tracking. In *European Conference on Computer Vision*, pages II: 3–19, 2000.
- [272] A. A. Maciejewski. Motion simulation: Dealing with the ill-conditioned equations of motion for articulated figures. *IEEE Comput. Graph. Appl.*, 10(3):63–71, 1990.
- [273] D. J. C. MacKay. *Information Theory, Inference & Learning Algorithms*. Cambridge University Press, New York, NY, USA, 2002.
- [274] C. Manning and H. Schütze. *Foundations of Statistical Natural Language Processing*. MIT Press, 1999.
- [275] D. Marr and H. Nishihara. Representation and recognition of the spatial organization of three-dimensional shapes. *Proc. Roy. Soc. B*, 200:269–294, 1978.

- [276] M. J. Matarić, V. B. Zordan, and Z. Mason. Movement control methods for complex, dynamically simulated agents: Adonis dances the macarena. In *AGENTS '98: Proceedings of the second international conference on Autonomous agents*, pages 317–324, New York, NY, USA, 1998. ACM Press.
- [277] M. J. Matarić, V. B. Zordan, and M. M. Williamson. Making complex articulated agents dance. *Autonomous Agents and Multi-Agent Systems*, 2(1):23–43, 1999.
- [278] H. Matsuo, S. Igi, S. Lu, Y. Nagashima, Y. Takata, and T. Teshima. The recognition algorithm with non-contact for japanese sign language using morphological analysis. In *Proc. Int. Gesture Workshop*, pages 273–284, 1997.
- [279] M. McKenna and D. Zeltzer. Dynamic simulation of autonomous legged locomotion. In *SIGGRAPH '90: Proceedings of the 17th annual conference on Computer graphics and interactive techniques*, pages 29–38, New York, NY, USA, 1990. ACM Press.
- [280] S. McKenna, S. Jabri, Z. Duric, A. Rosenfeld, and H. Wechsler. Tracking groups of people. *Computer Vision and Image Understanding*, 80(1):42–56, October 2000.
- [281] A. Menache. *Understanding Motion Capture for Computer Animation and Video Games*. Morgan-Kaufmann, 1999.
- [282] A. Micilotta, E. Ong, and R. Bowden. Detection and tracking of humans by probabilistic body part assembly. In *British Machine Vision Conference*, volume 1, pages 429–438, 2005.
- [283] K. Mikolajczyk. Face detector. Technical report, INRIA Rhone-Alpes. Ph.D report.
- [284] K. Mikolajczyk and C. Schmid. A performance evaluation of local descriptors. *IEEE T. Pattern Analysis and Machine Intelligence*, 2004. accepted.
- [285] K. Mikolajczyk, C. Schmid, and A. Zisserman. Human detection based on a probabilistic assembly of robust part detectors. In *European Conference on Computer Vision*, pages Vol I: 69–82, 2004.
- [286] K. Mikolajczyk, C. Schmid, and A. Zisserman. Human detection based on a probabilistic assembly of robust part detectors. In *European Conference on Computer Vision*, pages Vol I: 69–82, 2004.
- [287] A. Mittal and L. S. Davis. M2tracker: A multi-view approach to segmenting and tracking people in a cluttered scene. *Int. J. Comput. Vision*, 51(3):189–203, 2003.
- [288] T. Moeslund. Summaries of 107 computer vision-based human motion capture papers. Technical Report LLA 99-01, University of Aalborg, 1999.
- [289] B. Moghaddam and A. Pentland. Probabilistic visual learning for object detection. In *Int. Conf. on Computer Vision*, pages 786–793, 1995.
- [290] A. Mohan, C. Papageorgiou, and T. Poggio. Example-based object detection in images by components. *IEEE T. Pattern Analysis and Machine Intelligence*, 23(4):349–361, April 2001.
- [291] A. Mohr and M. Gleicher. Building efficient, accurate character skins from examples. *ACM Trans. Graphics*, 22(3):562–568, 2003.
- [292] M. Mohri. Finite-state transducers in language and speech processing. *Computational Linguistics*, 23(2), 1997.
- [293] M. Mohri, F. C. Pereira, and M. Riley. Weighted finite-state transducers in speech recognition. *Computer Speech and Language*, 16(1):69–88, 2002.
- [294] M. Mohri and M. Riley. Weighted finite-state transducers in speech recognition (tutorial). In *Proceedings of the International Conference on Spoken Language Processing 2002 (ICSLP '02)*, 2002.

- [295] G. Monheit and N. I. Badler. A kinematic model of the human spine and torso. *IEEE Comput. Graph. Appl.*, 11(2):29–38, 1991.
- [296] G. Mori and J. Malik. Estimating human body configurations using shape context matching. In *European Conference on Computer Vision LNCS 2352*, volume 3, pages 666–680, 2002.
- [297] G. Mori and J. Malik. Recovering 3d human body configurations using shape contexts. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2005. to appear.
- [298] T. Mori, Y. Segawa, M. Shimosaka, and T. Sato. Hierarchical recognition of daily human actions based on continuous hidden markov models. In *Int. Conf. Automatic Face and Gesture Recognition*, pages 779–784, 2004.
- [299] M. Müller, T. Röder, and M. Clausen. Efficient content-based retrieval of motion capture data. *ACM Trans. Graph.*, 24(3):677–685, 2005.
- [300] F. Multon, L. France, M.-P. Cani, and G. Debunne. Computer animation of human walking: a survey. *Journal of Visualization and Computer Animation (JVCA)*, 10:39–54, 1999. Published under the name Marie-Paule Cani-Gascuel.
- [301] J. Mundy and C.-F. Chang. Fusion of intensity, texture, and color in video tracking based on mutual information. In *Applied Imagery Pattern Recognition Workshop*, pages 10–15, 2004.
- [302] K. Murphy, Y. Weiss, and M. Jordan. Loopy belief propagation for approximate inference: An empirical study. In *Proceedings of the Annual Conference on Uncertainty in Artificial Intelligence*, pages 467–475, 1999.
- [303] E. Muybridge. *Animals in Motion*. Dover, 1957.
- [304] E. Muybridge. *The Human Figure in Motion*. Dover, 1989.
- [305] M.W.Kadous. Machine recognition of auslan signs using powergloves: towards large lexicon integration of sign language. In *Proc. Workshop on the Integration of Gesture in Language and Speech*, pages 165–174, 1996.
- [306] R. Neal. Probabilistic inference using markov chain monte carlo methods. Computer science tech report crg-tr-93-1, University of Toronto, 1993.
- [307] R. Neal. Sampling from multimodal distributions using tempered transitions. *Statistics and Computing*, 6:353–366, 1996.
- [308] R. Neal. Annealed importance sampling. Technical report, Technical Report No. 9805 (revised), Dept. of Statistics, University of Toronto, 1998.
- [309] R. M. Neal. Annealed importance sampling. *Statistics and Computing*, 11(2):125–139, 2001.
- [310] J. T. Ngo and J. Marks. Physically realistic motion synthesis in animation. *Evol. Comput.*, 1(3):235–268, 1993.
- [311] J. T. Ngo and J. Marks. Spacetime constraints revisited. In *SIGGRAPH '93: Proceedings of the 20th annual conference on Computer graphics and interactive techniques*, pages 343–350, New York, NY, USA, 1993. ACM Press.
- [312] S. Niyogi and E. Adelson. Analyzing gait with spatiotemporal surfaces. In *Proc. IEEE Workshop on Nonrigid and Articulated Motion*, pages 64–69, 1994.
- [313] S. Niyogi and E. Adelson. Analyzing and recognizing walking figures in xyt. Media lab vision and modelling tr-223, MIT, 1995.
- [314] S. of Anthropology Research Project, editor. *Anthropometric Source Book*. Webb Associates, 1978. NASA reference publication 1024, 3 Vols.

- [315] N. Oliver, A. Garg, and E. Horvitz. Layered representations for learning and inferring office activity from multiple sensory channels. *Computer Vision and Image Understanding*, 96(2):163–180, November 2004.
- [316] N. Oliver, E. Horvitz, and A. Garg. Layered representations for human activity recognition. In *Multimodal Interfaces, 2002. Proceedings. Fourth IEEE International Conference on*, pages 3–8, 2002.
- [317] M. Oren, C. Papageorgiou, P. Sinha, E. Osuna, and T. Poggio. Pedestrian detection using wavelet templates. In *IEEE Conf. on Computer Vision and Pattern Recognition*, pages 193–199, 1997.
- [318] M. Oren, C. Papageorgiou, P. Sinha, E. Osuna, and T. Poggio. A trainable system for people detection. In *DARPA IU Workshop*, pages 207–214, 1997.
- [319] J. O’Rourke and N. Badler. Model-based image analysis of human motion using constraint propagation. *IEEE T. Pattern Analysis and Machine Intelligence*, 2(6):522–536, November 1980.
- [320] E. Osuna, R. Freund, and F. Girosi. Training support vector machines: an application to face detection. In *IEEE Conf. on Computer Vision and Pattern Recognition*, pages 130–6, 1997.
- [321] C. Pai, H. Tyan, Y. Liang, H. Liao, and S. Chen. Pedestrian detection and tracking at crossroads. In *IEEE Int. Conf. Image Processing*, pages II: 101–104, 2003.
- [322] C. Pai, H. Tyan, Y. Liang, H. Liao, and S. Chen. Pedestrian detection and tracking at crossroads. *Pattern Recognition*, 37(5):1025–1034, May 2004.
- [323] M. Pandey and F. Anderson. Dynamic simulation of human movement using large-scale models of the body. In *Proc. IEEE Intl. Conference on Robotics and Automation*, pages 676–681, 2000.
- [324] M. Pandey, F. Anderson, and D. Hull. A parameter optimization approach for the optimal control of large-scale musculoskeletal systems. *J. of Biomech. Eng.*, pages 450–460, 1992.
- [325] C. Papageorgiou. A trainable system for object detection in images and video sequences constantine. Technical report, MIT, 2000. Ph. D.
- [326] C. Papageorgiou, T. Evgeniou, and T. Poggio. A trainable object detection system. In *DARPA IU Workshop*, pages 1019–1024, 1998.
- [327] C. Papageorgiou, M. Oren, and T. Poggio. A general framework for object detection. In *Int. Conf. on Computer Vision*, pages 555–562, 1998.
- [328] C. Papageorgiou and T. Poggio. A pattern classification approach to dynamical object detection. In *Int. Conf. on Computer Vision*, pages 1223–1228, 1999.
- [329] C. Papageorgiou and T. Poggio. Trainable pedestrian detection. In *IEEE Int. Conf. Image Processing*, pages IV:35–39, 1999.
- [330] C. Papageorgiou and T. Poggio. A trainable system for object detection. *Int. J. Computer Vision*, 38(1):15–33, June 2000.
- [331] V. Parenti-Castelli, A. Leardini, R. D. Gregorio, and J. J. O’Connor. On the modeling of passive motion of the human knee joint by means of equivalent planar and spatial parallel mechanisms. *Auton. Robots*, 16(2):219–232, 2004.
- [332] S. I. Park, H. J. Shin, T. H. Kim, and S. Y. Shin. On-line motion blending for real-time locomotion generation: Research articles. *Comput. Animat. Virtual Worlds*, 15(3-4):125–138, 2004.

- [333] S. I. Park, H. J. Shin, and S. Y. Shin. On-line locomotion generation based on motion blending. In *SCA '02: Proceedings of the 2002 ACM SIGGRAPH/Eurographics symposium on Computer animation*, pages 105–111, 2002.
- [334] C. B. Phillips, J. Zhao, and N. I. Badler. Interactive real-time articulated figure manipulation using multiple kinematic constraints. In *SI3D '90: Proceedings of the 1990 symposium on Interactive 3D graphics*, pages 245–250, New York, NY, USA, 1990. ACM Press.
- [335] S. D. Pietra, V. D. Pietra, and J. Lafferty. Inducing features of random fields. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 19(4):380–393, 1997.
- [336] C. Pinhanez and A. Bobick. Pnf propagation and the detection of actions described by temporal intervals. In *DARPA IU Workshop*, pages 227–234, 1997.
- [337] C. Pinhanez and A. Bobick. Human action detection using pnf propagation of temporal constraints. In *IEEE Conf. on Computer Vision and Pattern Recognition*, pages 898–904, 1998.
- [338] R. Plänkers and P. Fua. Tracking and modeling people in video sequences. *Comput. Vis. Image Underst.*, 81(3):285–302, 2001.
- [339] T. Poggio and K.-K. Sung. Finding human faces with a gaussian mixture distribution-based face model. In *Asian Conf. on Computer Vision*, pages 435–440, 1995.
- [340] R. Polana and R. Nelson. Detecting activities. In *DARPA IU Workshop*, pages 569–574, 1993.
- [341] R. Polana and R. Nelson. Detecting activities. In *IEEE Conf. on Computer Vision and Pattern Recognition*, pages 2–7, 1993.
- [342] R. Polana and R. Nelson. Detecting activities. *J. Visual Communication Image Representation*, 5:172–180, 1994.
- [343] R. Polana and R. Nelson. Detecting activities. In *Proceedings IAPR International Conference on Pattern Recognition*, pages A:815–818, 1994.
- [344] R. Polana and R. Nelson. Low level recognition of human motion. In *IEEE Workshop on Articulated and Non-Rigid Motion*, 1994.
- [345] R. Polana and R. Nelson. Recognition of nonrigid motion. In *ARPA94*, pages II:1219–1224, 1994.
- [346] R. Polana and R. Nelson. Recognizing activities. In *Proceedings IAPR International Conference on Pattern Recognition*, pages A:815–818, 1994.
- [347] R. Polana and R. Nelson. Detection and recognition of periodic, nonrigid motion. *Int. J. Computer Vision*, 23(3):261–282, 1997.
- [348] N. S. Pollard and F. Behmaram-Mosavat. Force-based motion editing for locomotion tasks. In *In Proceedings of the IEEE International Conference on Robotics and Automation*, 2000.
- [349] J. Popović, S. M. Seitz, and M. Erdmann. Motion sketching for control of rigid-body simulations. *ACM Trans. Graph.*, 22(4):1034–1054, 2003.
- [350] Z. Popović and A. Witkin. Physically based motion transformation. In *SIGGRAPH '99: Proceedings of the 26th annual conference on Computer graphics and interactive techniques*, pages 11–20, New York, NY, USA, 1999. ACM Press/Addison-Wesley Publishing Co.
- [351] K. Pullen and C. Bregler. Motion capture assisted animation: Texturing and synthesis. *Proceedings of SIGGRAPH 95*, 2002.

- [352] C. Putnam. A segment interaction analysis of proximal-to-distal sequential segment motion patterns. *Med. Sci. Sports. Exerc.*, 23:130–144, 1991.
- [353] L. Rabiner and B.-H. Juang. *Fundamentals of Speech Recognition*. Prentice-Hall, 1993.
- [354] D. Ramanan. *Tracking People and Recognizing their Activities*. PhD thesis, U.C. Berkeley, 2005.
- [355] D. Ramanan and D. Forsyth. Automatic annotation of everyday movements. In *Proc. Neural Information Processing Systems*, 2003.
- [356] D. Ramanan and D. Forsyth. Finding and tracking people from the bottom up. In *IEEE Conf. on Computer Vision and Pattern Recognition*, pages II: 467–474, 2003.
- [357] D. Ramanan, D. Forsyth, and A. Zisserman. Strike a pose: Tracking people by finding stylized poses. In *IEEE Conf. on Computer Vision and Pattern Recognition*, pages I: 271–278, 2005.
- [358] P. Reitsma and N. Pollard. Evaluating motion graphs for character navigation. In *Eurographics/ACM Symposium on Computer Animation*, pages 89–98, 2004.
- [359] L. Ren, A. Patrick, A. A. Efros, J. K. Hodgins, and J. M. Rehg. A data-driven approach to quantifying natural human motion. *ACM Trans. Graph.*, 24(3):1090–1097, 2005.
- [360] L. Ren, G. Shakhnarovich, J. K. Hodgins, H. Pfister, and P. Viola. Learning silhouette features for control of human motion. *ACM Trans. Graph.*, 24(4):1303–1331, 2005.
- [361] B. Ristic, S. Arulampalam, and N. Gordon. *Beyond the Kalman Filter: Particle Filters for Tracking Applications*. Artech House, 2004.
- [362] J. Rittscher and A. Blake. Classification of human body motion. In *Int. Conf. on Computer Vision*, pages 634–639, 1999.
- [363] T. Roberts, S. McKenna, and I. Ricketts. Human pose estimation using learnt probabilistic region similarities and partial configurations. In *European Conference on Computer Vision*, pages Vol IV: 291–303, 2004.
- [364] K. Rohr. Incremental recognition of pedestrians from image sequences. In *IEEE Conf. on Computer Vision and Pattern Recognition*, pages 9–13, 1993.
- [365] K. Rohr. Towards model-based recognition of human movements in image sequences. *CVGIP: Image Understanding*, 59(1):94–115, 1994.
- [366] R. Ronfard, C. Schmid, and B. Triggs. Learning to parse pictures of people. In *European Conference on Computer Vision*, page IV: 700 ff., 2002.
- [367] R. Rosales, V. Athitsos, L. Sigal, and S. Sclaroff. 3d hand pose reconstruction using specialized mappings. In *Int. Conf. on Computer Vision*, pages I: 378–385, 2001.
- [368] C. Rose, M. F. Cohen, and B. Bodenheimer. Verbs and adverbs: Multidimensional motion interpolation. *IEEE Comput. Graph. Appl.*, 18(5):32–40, 1998.
- [369] C. Rose, B. Guenter, B. Bodenheimer, and M. F. Cohen. Efficient generation of motion transitions using spacetime constraints. In *SIGGRAPH '96: Proceedings of the 23rd annual conference on Computer graphics and interactive techniques*, pages 147–154, New York, NY, USA, 1996. ACM Press.
- [370] R. Rosenfeld. A maximum entropy approach to adaptive statistical language modelling. *Computer, Speech and Language*, 10:187–228, 1996.
- [371] S. Roth, L. Sigal, and M. Black. Gibbs likelihoods for bayesian tracking. In *CVPR04*, pages I: 886–893, 2004.
- [372] P. Rousseeuw. *Robust Regression and Outlier Detection*. Wiley, 1987.

- [373] H. Rowley, S. Baluja, and T. Kanade. Human face detection in visual scenes. In D. Touretzky, M. Mozer, and M. Hasselmo, editors, *Advances in Neural Information Processing* 8, pages 875–881, 1996.
- [374] H. Rowley, S. Baluja, and T. Kanade. Neural network-based face detection. In *IEEE Conf. on Computer Vision and Pattern Recognition*, pages 203–8, 1996.
- [375] H. Rowley, S. Baluja, and T. Kanade. Neural network-based face detection. *IEEE T. Pattern Analysis and Machine Intelligence*, 20(1):23–38, 1998.
- [376] H. Rowley, S. Baluja, and T. Kanade. Rotation invariant neural network-based face detection. In *IEEE Conf. on Computer Vision and Pattern Recognition*, pages 38–44, 1998.
- [377] A. Safonova and J. K. Hodgins. Analyzing the physical correctness of interpolated human motion. In *SCA '05: Proceedings of the 2005 ACM SIGGRAPH/Eurographics symposium on Computer animation*, pages 171–180, New York, NY, USA, 2005. ACM Press.
- [378] A. Safonova, J. K. Hodgins, and N. S. Pollard. Synthesizing physically realistic human motion in low-dimensional, behavior-specific spaces. *ACM Trans. Graph.*, 23(3):514–521, 2004.
- [379] S. Schaal, C. G. Atkeson, and S. Vijayakumar. Scalable techniques from nonparametric statistics for real time robot learning. *Applied Intelligence*, 17(1):49–60, 2002.
- [380] G. Schmidt. Designing nonlinear filters based on daum’s theory. *Journal of Guidance, Control and Dynamics*, 16:371–376, 1993.
- [381] H. Schneiderman and T. Kanade. A statistical method for 3d object detection applied to faces and cars. In *IEEE Conf. on Computer Vision and Pattern Recognition*, pages I: 746–751, 2000.
- [382] B. Scholkopf and A. J. Smola. *Learning with Kernels: Support Vector Machines, Regularization, Optimization, and Beyond*. MIT Press, Cambridge, MA, USA, 2001.
- [383] S. Seitz and C. Dyer. Affine invariant detection of periodic motion. In *IEEE Conf. on Computer Vision and Pattern Recognition*, pages 970–975, 1994.
- [384] S. Seitz and C. Dyer. View invariant analysis of cyclic motion. *Int. J. Computer Vision*, 25(3):231–251, December 1997.
- [385] A. Senior. Tracking people with probabilistic appearance models. In *IEEE Workshop on Performance Evaluation Tracking Surveillance*, pages 48–55, 2002.
- [386] A. Shahrokni, T. Drummond, and P. Fua. Fast Texture-Based Tracking and Delineation Using Texture Entropy. In *International Conference on Computer Vision*, 2005.
- [387] A. Shahrokni, T. Drummond, V. Lepetit, and P. Fua. Markov-based Silhouette Extraction for Three-Dimensional Body Tracking in Presence of Cluttered Background. In *British Machine Vision Conference*, Kingston, UK, 2004.
- [388] A. Shahrokni, F. Fleuret, and P. Fua. Classifier-based Contour Tracking for Rigid and Deformable Objects. In *British Machine Vision Conference*, Oxford, UK, 2005.
- [389] G. Shakhnarovich, P. Viola, and T. Darrell. Fast pose estimation with parameter-sensitive hashing. In *Int. Conf. on Computer Vision*, pages 750–757, 2003.
- [390] J. Shawe-Taylor and N. Cristianini. *Kernel Methods for Pattern Analysis*. Cambridge University Press, New York, NY, USA, 2004.
- [391] H. J. Shin, L. Kovar, and M. Gleicher. Physical touch-up of human motions. In *PG '03: Proceedings of the 11th Pacific Conference on Computer Graphics and Applications*, page 194, Washington, DC, USA, 2003. IEEE Computer Society.

- [392] H. J. Shin, J. Lee, S. Y. Shin, and M. Gleicher. Computer puppetry: An importance-based approach. *ACM Trans. Graph.*, 20(2):67–94, 2001.
- [393] H. Sidenbladh and M. Black. Learning image statistics for bayesian tracking. In *Int. Conf. on Computer Vision*, pages II: 709–716, 2001.
- [394] H. Sidenbladh and M. Black. Learning the statistics of people in images and video. *Int. J. Computer Vision*, 54(1):181–207, September 2003.
- [395] H. Sidenbladh, M. J. Black, and D. J. Fleet. Stochastic tracking of 3d human figures using 2d image motion. In *European Conference on Computer Vision*, 2000.
- [396] L. Sigal, S. Bhatia, S. Roth, M. Black, and M. Isard. Tracking loose-limbed people. In *IEEE Conf. on Computer Vision and Pattern Recognition*, pages I: 421–428, 2004.
- [397] M.-C. Silaghi, R. Plänkers, R. Boulic, P. Fua, and D. Thalmann. Local and global skeleton fitting techniques for optical motion capture. In *Modelling and Motion Capture Techniques for Virtual Environments*, pages 26–40, Nov. 1998. Proceedings of CAPTECH '98.
- [398] K. Sims. Evolving virtual creatures. In *SIGGRAPH '94: Proceedings of the 21st annual conference on Computer graphics and interactive techniques*, pages 15–22, New York, NY, USA, 1994. ACM Press.
- [399] J. M. Siskind. Grounding language in perception. *Artificial Intelligence Review*, 8:371–391, 1995.
- [400] J. M. Siskind. Reconstructing force-dynamic models from video sequences. *Artificial Intelligence*, 151:91–154, 2003.
- [401] J. Sivic, M. Everingham, and A. Zisserman. Person spotting: video shot retrieval for face sets. In *International Conference on Image and Video Retrieval (CIVR 2005)*, Singapore, 2005.
- [402] C. Sminchisescu. Consistency and coupling in human model likelihoods. In *Proceedings International Conference on Automatic Face and Gesture Recognition*, pages 22–27, 2002.
- [403] C. Sminchisescu and A. Telea. Human pose estimation from silhouettes: A consistent approach using distance level sets. In *WSCG02*, page 413, 2002.
- [404] C. Sminchisescu and B. Triggs. Covariance scaled sampling for monocular 3d body tracking. In *IEEE Conf. on Computer Vision and Pattern Recognition*, pages I:447–454, 2001.
- [405] C. Sminchisescu and B. Triggs. Building roadmaps of local minima of visual models. In *European Conference on Computer Vision*, page I: 566 ff., 2002.
- [406] C. Sminchisescu and B. Triggs. Hyperdynamics importance sampling. In *European Conference on Computer Vision*, page I: 769 ff., 2002.
- [407] C. Sminchisescu and B. Triggs. Estimating articulated human motion with covariance scaled sampling. *The International Journal of Robotics Research*, 22(6):371–391, 2003.
- [408] C. Sminchisescu and B. Triggs. Kinematic jump processes for monocular 3d human tracking. In *IEEE Conf. on Computer Vision and Pattern Recognition*, pages I: 69–76, 2003.
- [409] C. Sminchisescu and B. Triggs. Kinematic jump processes for monocular 3d human tracking. In *IEEE Conf. on Computer Vision and Pattern Recognition*, pages I: 69–76, 2003.

- [410] C. Sminchisescu and B. Triggs. Building roadmaps of minima and transitions in visual models. *Int. J. Computer Vision*, 61(1):81–101, January 2005.
- [411] A. J. Smola and B. Schölkopf. A tutorial on support vector regression. *Statistics and Computing*, 14(3):199–222, 2004.
- [412] Y. Song, X. Feng, and P. Perona. Towards detection of human motion. In *IEEE Conf. on Computer Vision and Pattern Recognition*, pages I: 810–817, 2000.
- [413] Y. Song, L. Goncalves, and P. Perona. Unsupervised learning of human motion. *IEEE T. Pattern Analysis and Machine Intelligence*, 25(7):814–827, July 2003.
- [414] N. Sprague and J. Luo. Clothed people detection in still images. In *Proceedings IAPR International Conference on Pattern Recognition*, pages III: 585–589, 2002.
- [415] J. Starck and A. Hilton. Model-based multiple view reconstruction of people. In *Int. Conf. on Computer Vision*, pages 915–922, 2003.
- [416] J. Starck and A. Hilton. Spherical matching for temporal correspondence of non-rigid surfaces. In *Int. Conf. on Computer Vision*, 2005.
- [417] J. Starck and A. Hilton. Virtual view synthesis of people from multiple view video sequences. *Graphical Models*, 67(6):600–620, 2005.
- [418] J. Starck, A. Hilton, and J. Illingworth. Human shape estimation in a multi-camera studio. In *BMVC*, 2001.
- [419] T. Starner and A. Pentland. Visual recognition of american sign language using hidden markov models. Technical report, MIT, 1996.
- [420] T. Starner, J. Weaver, and A. Pentland. Real-time american sign language recognition using desk and wearable computer based video. *IEEE T. Pattern Analysis and Machine Intelligence*, 20(12):1371–1375, 1998.
- [421] C. Stauffer and W. Grimson. Adaptive background mixture models for real-time tracking. In *IEEE Conf. on Computer Vision and Pattern Recognition*, pages II: 246–252, 1999.
- [422] C. Stauffer and W. Grimson. Adaptive background mixture models for real-time tracking. In *IEEE Conf. on Computer Vision and Pattern Recognition*, pages II: 246–252, 1999.
- [423] C. Stauffer and W. Grimson. Learning patterns of activity using real-time tracking. *IEEE T. Pattern Analysis and Machine Intelligence*, 22(8):747–757, August 2000.
- [424] M. Stone, D. DeCarlo, I. Oh, C. Rodriguez, A. Stere, A. Lees, and C. Bregler. Speaking with hands: creating animated conversational characters from recordings of human performance. *ACM Trans. Graph.*, 23(3):506–513, 2004.
- [425] A. Sulejmanpašić and J. Popović. Adaptation of performed ballistic motion. *ACM Trans. Graph.*, 24(1):165–179, 2005.
- [426] J. Sullivan, A. Blake, and J. Rittscher. Statistical foreground modelling for object localisation. In *European Conference on Computer Vision*, pages II: 307–323, 2000.
- [427] J. Sullivan and S. Carlsson. Recognizing and tracking human action. In *European Conference on Computer Vision*, page I: 629 ff., 2002.
- [428] K.-K. Sung and T. Poggio. Example based learning for view based face detection. Ai memo 1521, MIT, 1994.
- [429] K.-K. Sung and T. Poggio. Example-based learning for view-based human face detection. *IEEE T. Pattern Analysis and Machine Intelligence*, 20:39–51, 1998.
- [430] S. Tak and H. Ko. Example guided inverse kinematics. In *International Conference on Computer Graphics and Imaging*, pages 19–23, 2000.

- [431] S. Tak and H.-S. Ko. A physically-based motion retargeting filter. *ACM Trans. Graph.*, 24(1):98–117, 2005.
- [432] S. Tak, O. Song, and H. Ko. Motion balance filtering. 19(3):437–446, 2000.
- [433] C. Taylor. Reconstruction of articulated objects from point correspondences in a single uncalibrated image. In *IEEE Conf. on Computer Vision and Pattern Recognition*, pages 677–84, 2000.
- [434] C. Taylor. Reconstruction of articulated objects from point correspondences in a single uncalibrated image. *Computer Vision and Image Understanding*, 80(3):349–363, December 2000.
- [435] A. Thangali and S. Sclaroff. Periodic motion detection and estimation via space-time sampling. In *Motion05*, pages II: 176–182, 2005.
- [436] C. Theobalt, J. Carranza, M. A. Magnor, and H.-P. Seidel. Enhancing silhouette-based human motion capture with 3d motion fields. In *PG '03: Proceedings of the 11th Pacific Conference on Computer Graphics and Applications*, page 185, Washington, DC, USA, 2003. IEEE Computer Society.
- [437] M. E. Tipping. The relevance vector machine. In *In Advances in Neural Information Processing Systems 12*, pages 332–388. MIT Press, 2000.
- [438] M. E. Tipping. Sparse bayesian learning and the relevance vector machine. *J. Mach. Learn. Res.*, 1:211–244, 2001.
- [439] D. Tolani and N. I. Badler. Real-time inverse kinematics of the human arm. *Presence*, 5(4):393–401, 1996.
- [440] D. Tolani, A. Goswami, and N. I. Badler. Real-time inverse kinematics techniques for anthropomorphic limbs. *Graphical Models*, 62:353–388, 2000.
- [441] N. Torkos and M. V. de Panne. Footprint-based quadruped motion synthesis. In *Graphics Interface 98*, pages 151–160, 1998.
- [442] K. Toyama and A. Blake. Probabilistic tracking in a metric space. In *Int. Conf. on Computer Vision*, pages II: 50–57, 2001.
- [443] K. Toyama and A. Blake. Probabilistic tracking with exemplars in a metric space. *Int. J. Computer Vision*, 48(1):9–19, June 2002.
- [444] Z. Tu and S. Zhu. Image segmentation by data-driven markov chain monte carlo. In *Int. Conf. on Computer Vision*, pages II: 131–138, 2001.
- [445] Z. Tu and S. Zhu. Image segmentation by data-driven markov chain monte carlo. *IEEE T. Pattern Analysis and Machine Intelligence*, 24(5):657–673, May 2002.
- [446] S. Tumer and A. Engin. Three-dimensional kinematic modelling of the human shoulder complex - part ii: Mathematical modelling and solution via optimization. *ASME Journal of Biomechanical Engineering*, 111:113–121, 1989.
- [447] V. Vapnik. *The Nature of Statistical Learning Theory*. Springer Verlag, 1996.
- [448] V. N. Vapnik. *Statistical Learning Theory*. John Wiley and Sons, 1998.
- [449] D. D. Vecchio, R. Murray, and P. Perona. Decomposition of human motion into dynamics-based primitives with application to drawing tasks. *Automatica*, 39(12):2085–2098, 2003.
- [450] P. Viola and M. Jones. Rapid object detection using a boosted cascade of simple features. In *IEEE Conf. on Computer Vision and Pattern Recognition*, pages I:511–518, 2001.
- [451] P. Viola and M. Jones. Robust real-time face detection. In *Int. Conf. on Computer Vision*, page II: 747, 2001.

- [452] P. Viola and M. Jones. Robust real-time face detection. *Int. J. Computer Vision*, 57(2):137–154, May 2004.
- [453] P. Viola, M. Jones, and D. Snow. Detecting pedestrians using patterns of motion and appearance. In *Int. Conf. on Computer Vision*, pages 734–741, 2003.
- [454] P. Viola, M. Jones, and D. Snow. Detecting pedestrians using patterns of motion and appearance. *Int. J. Computer Vision*, 63(2):153–161, July 2005.
- [455] C. Vogler and D. Metaxas. ASL recognition based on a coupling between HMMs and 3D motion analysis. In *Int. Conf. on Computer Vision*, pages 363–369, 1998.
- [456] C. Vogler and D. Metaxas. Parallel hidden markov models for American sign language recognition. In *Int. Conf. on Computer Vision*, pages 116–122, 1999.
- [457] C. Vogler and D. Metaxas. Toward scalability in asl recognition: breaking down signs into phonemes. In *Gesture workshop 99*, 1999.
- [458] C. Vogler and D. Metaxas. Handshapes and movements: Multiple-channel american sign language recognition. In *Proc. Gesture Workshop*, pages 247–258, 2003.
- [459] C. Vogler, H. Sun, and D. Metaxas. A framework for motion recognition with applications to American sign language and gait recognition. In *IEEE Workshop on Human Motion*, 2000.
- [460] J. Wang and B. Bodenheimer. An evaluation of a cost metric for selecting transitions between motion segments. In *SCA '03: Proceedings of the 2003 ACM SIGGRAPH/Eurographics symposium on Computer animation*, pages 232–238, Aire-la-Ville, Switzerland, Switzerland, 2003. Eurographics Association.
- [461] J. Wang and B. Bodenheimer. Computing the duration of motion transitions: an empirical approach. In *SCA '04: Proceedings of the 2004 ACM SIGGRAPH/Eurographics symposium on Computer animation*, pages 335–344, New York, NY, USA, 2004. ACM Press.
- [462] J. J. Wang and S. Singh. Video analysis of human dynamics - a survey. *Real-Time Imaging*, 9(5):321–346, 2003.
- [463] M. Weber, W. Einhauser, M. Welling, and P. Perona. Viewpoint-invariant learning and detection of human heads. In *IEEE International Conference on Automatic Face and Gesture Recognition*, pages 20–7, 2000.
- [464] Y. Weiss. Belief propagation and revision in networks with loops. Technical report, Massachusetts Institute of Technology, Cambridge, MA, USA, 1997.
- [465] W. Gao, J. Ma, J. Wu, and C. Wang. Sign language recognition based on hmm/ann/dp. *Int. J. Pattern Recognition and Artificial Intelligence*, 14(5):587–602, 2000.
- [466] D. J. Wiley and J. K. Hahn. Interpolation synthesis of articulated figure motion. *IEEE Comput. Graph. Appl.*, 17(6):39–45, 1997.
- [467] A. Wilson and A. Bobick. Learning visual behavior for gesture analysis. In *IEEE Symposium on Computer Vision*, pages 229–234, 1995.
- [468] A. Wilson and A. Bobick. Parametric hidden markov models for gesture recognition. *IEEE T. Pattern Analysis and Machine Intelligence*, 21(9):884–900, September 1999.
- [469] A. Witkin and M. Kass. Spacetime constraints. In *SIGGRAPH '88: Proceedings of the 15th annual conference on Computer graphics and interactive techniques*, pages 159–168, New York, NY, USA, 1988. ACM Press.
- [470] A. Witkin and Z. Popović. Motion warping. In *SIGGRAPH '95: Proceedings of the 22nd annual conference on Computer graphics and interactive techniques*, pages 105–108, New York, NY, USA, 1995. ACM Press.

- [471] C. Wren, A. Azarbayejani, T. Darrell, and A. Pentland. Pfinder: Real-time tracking of the human body. *IEEE T. Pattern Analysis and Machine Intelligence*, 19(7):780–785, July 1997.
- [472] M.-Y. Wu, C.-Y. Chiu, S.-P. Chao, S.-N. Yang, and H.-C. Lin. Content-based retrieval for human motion data. In *16th IPPR Conference on Computer Vision, Graphics and Image Processing*, pages 605–612, 2003.
- [473] Y. Wu, T. Yu, and G. Hua. A statistical field model for pedestrian detection. In *IEEE Conf. on Computer Vision and Pattern Recognition*, pages I: 1023–1030, 2005.
- [474] Y. Yacoob and L. Davis. Learned models for estimation of rigid and articulated human motion from stationary or moving camera. *Int. J. Computer Vision*, 36(1):5–30, January 2000.
- [475] M. Yamamoto, A. Sato, S. Kawada, T. Kondo, and Y. Osaki. Incremental tracking of human actions from multiple views. In *CVPR '98: Proceedings of the IEEE Computer Society Conference on Computer Vision and Pattern Recognition*, page 2, Washington, DC, USA, 1998. IEEE Computer Society.
- [476] K. Yamane and Y. Nakamura. Natural motion animation through constraining and deconstraining at will. *IEEE Transactions on Visualization and Computer Graphics*, 9(3):352–360, 2003.
- [477] J. Yamato, J. Ohya, and K. Ishii. Recognising human action in time sequential images using hidden markov model. In *IEEE Conf. on Computer Vision and Pattern Recognition*, pages 379–385, 1992.
- [478] W. Yan and D. Forsyth. Learning the behaviour of users in a public space through video tracking. In *CVPR*, 2004. In review.
- [479] J. Yang, Z. Fu, T. Tan, and W. Hu. A novel approach to detecting adult images. In *Proceedings IAPR International Conference on Pattern Recognition*, pages IV: 479–482, 2004.
- [480] J. Yang, Y. Xu, and C. S. Chen. Human action learning via hidden markov model. *IEEE Transactions on Systems Man and Cybernetics*, 27:34–44, 1997.
- [481] J. S. Yedidia, W. T. Freeman, and Y. Weiss. Understanding belief propagation and its generalizations. In *Exploring artificial intelligence in the new millennium*, pages 239–269. Morgan Kaufmann Publishers Inc., San Francisco, CA, USA, 2003.
- [482] J. Zhao and N. I. Badler. Inverse kinematics positioning using nonlinear programming for highly articulated figures. *ACM Trans. Graph.*, 13(4):313–336, 1994.
- [483] L. Zhao and N. Badler. Gesticulation behaviors for virtual humans. In *PG '98: Proceedings of the 6th Pacific Conference on Computer Graphics and Applications*, page 161, Washington, DC, USA, 1998. IEEE Computer Society.
- [484] L. Zhao and C. Thorpe. Stereo- and neural network-based pedestrian detection. *Intelligent Transportation Systems*, 1(3):148–154, September 2000.
- [485] T. Zhao and R. Nevatia. Tracking multiple humans in complex situations. *IEEE T. Pattern Analysis and Machine Intelligence*, 26(9):1208–1221, September 2004.
- [486] S. Zhu, R. Zhang, and Z. Tu. Integrating bottom-up/top-down for object recognition by data driven markov chain monte carlo. In *IEEE Conf. on Computer Vision and Pattern Recognition*, pages I: 738–745, 2000.
- [487] J. Zieren and K.-F. Kraiss. Non-intrusive sign language recognition for human computer interaction. In *Proc. IFAC/IFIP/IFORS/IEA symposium on analysis, design and evaluation of human machine systems*, 2004.

- [488] V. B. Zordan and J. K. Hodgins. Tracking and modifying upper-body human motion data with dynamic simulation. In *Computer Animation and Simulation '99*, Sept. 1999.
- [489] V. B. Zordan and J. K. Hodgins. Motion capture-driven simulations that hit and react. In *ACM SIGGRAPH Symposium on Computer Animation*, pages 89–96, July 2002.
- [490] M. Zyda, J. Hiles, A. Mayberry, C. Wardynski, M. Capps, B. Osborn, R. Shilling, M. Robaszewski, and M. Davis. Entertainment r&d for defense. *IEEE Computer Graphics and Applications*, pages 28–36, 2003.
- [491] M. Zyda, A. Mayberry, C. Wardynski, R. Shilling, and M. Davis. The moves institute's america's army operations game. In *Proceedings of the ACM SIGGRAPH 2003 Symposium on Interactive 3D Graphics*, pages 217–218, 2003.