

# Tracking

D.A. Forsyth, UIUC

# Tracking: Why?

- Motion capture
  - build models of moving people from video
- Recognition from motion
  - eg cyclists move differently than runners
- Surveillance
  - who is doing what?
    - for security (eg keep people out of sensitive areas in airports)
    - for HCI (eg kinect, eyetoy, etc.)

# Tracking

- Establish state of object using time sequence
  - state could be:
    - position; position+velocity; position+velocity+acceleration
    - or more complex, eg all joint angles for a person
  - Biggest problem -- Data Association
    - which image pixels are informative, which are not?
- Key ideas
  - Tracking by detection
    - if we know what an object looks like, that selects the pixels to use
  - Tracking through flow
    - if we know how an object moves, that selects the pixels to use

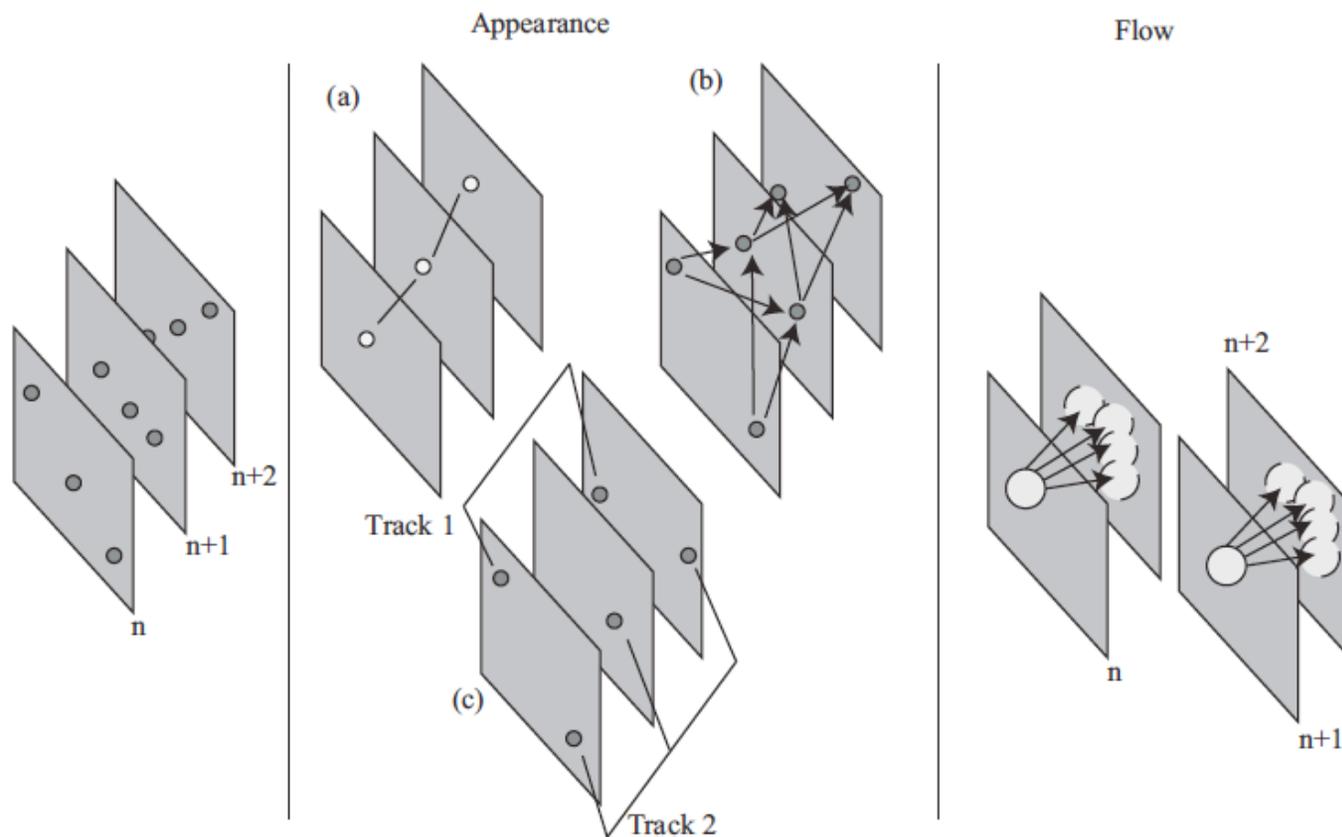


FIGURE 11.1: In tracking problems, we want to build space time paths followed by tokens—which might be objects, or regions, or interest points, or image windows—in an image sequence (left). There are two important sources of information; carefully used, they can resolve many tracking problems without further complexity. One is the appearance of the token being tracked. If there is only one token in each frame with a distinctive appearance, then we could detect it in each frame, then link the detector responses (a). Alternatively, if there is more than one instance per frame, a cost function together with weighted bipartite matching could be enough to build the track (b). If some instances drop out, we will need to link detector responses to abstract tracks (c); in the figure, track 1 has measurements for frames  $n$  and  $n + 2$ , but does not have a measurement for frame  $n + 1$ . Another important source of information is the motion of the token; if we have a manageable model of the flow, we could search for the flow that generates the best match in the next frame. We choose that match as the next location of the token, then iterate this procedure (right).

# Track by detection (simple form)

- Assume
  - a very reliable detector (e.g. faces; back of heads)
  - detections that are well spaced in images (or have distinctive properties)
    - e.g. news anchors; heads in public
- Link detects across time
  - only one - easy
  - multiple - weighted bipartite matching
  - but what if one is missing?
- Better: create abstract tracks
  - link detects to track
  - create tracks, reap tracks as required
  - clean up spacetime paths

**Notation:**

Write  $x_k(i)$  for the  $k$ 'th response of the detector in the  $i$ th frame

Write  $t(k, i)$  for the  $k$ 'th track in the  $i$ th frame

Write  $*t(k, i)$  for the detector response attached to the  $k$ 'th track in the  $i$ th frame  
(Think C pointer notation)

**Assumptions:** We have a detector which is reasonably reliable.

We know some distance  $d$  such that  $d(*t(k, i - 1), *t(k, i))$  is always small.

**First frame:** Create a track for each detector response.

**N'th frame:**

**Link** tracks and detector responses by solving a bipartite matching problem.

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

**Reap** any track that has not received a detector response for some number of frames.

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

**Algorithm 11.1:** Tracking by Detection.

# Tracking by known appearance

- Even if we don't have a detector
- Know rectangle in image (n)
- Want to find corresponding rectangle in (n+1)
- Search over nearby rectangles
  - to find one that minimizes SSD error (Sum of Squared Differences)

$$\sum_{i,j} (\mathcal{R}_{ij}^{(n)} - \mathcal{R}_{ij}^{(n+1)})^2.$$

- where sum is over pixels in rectangle
- Application
  - stabilize players in TV sport



FIGURE 11.2: A useful application of tracking is to stabilize an image box around a more interesting structure, in this case a football player in a television-resolution video. A frame from the video is shown on the left. Inset is a box around a player, zoomed to a higher resolution. Notice that the limbs of the player span a few pixels, are blurry, and are hard to resolve. A natural feature for inferring what the player is doing can be obtained by stabilizing the box around the player, then measuring the motion of the limbs with respect to the box. Players move relatively short distances between frames, and their body configuration changes a relatively small amount. This means the new box can be found by searching all nearby boxes of the same size to get the box whose pixels best match those of the original. On the right, a set of stabilized boxes; the strategy is enough to center the player in a box. *This figure was originally published as Figure 7 of “Recognizing Action at a Distance,” A. Efros, A.C. Berg, G. Mori, and J. Malik, Proc. IEEE ICCV, 2003, © IEEE, 2003.*

# Tracking by known appearance

Now write  $\mathcal{P}_t$  for the indices of the patch in the  $t$ th frame and  $I(x, t)$  for the  $t$ th frame. Assume that the patch is at  $x_t$  in the  $t$ th frame and it translates to  $x_t + h$  in the  $t + 1$ th frame. Then we can determine  $h$  by minimizing

$$E(h) = \sum_{u \in \mathcal{P}_t} [I(u, t) - I(u + h, t + 1)]^2$$

as a function of  $h$ . The minimum of the error occurs when

$$\nabla_h E(h) = 0.$$

Now if  $h$  is small, we can write  $I(u + h, t + 1) \approx I(u, t) + h^T \nabla I$ , where  $\nabla I$  is the image gradient. Substituting, and rearranging, we get

This matrix is part of the Harris corner detector —  $\left[ \sum_{u \in \mathcal{P}_t} (\nabla I)(\nabla I)^T \right] h = \sum_{u \in \mathcal{P}_t} [I(u, t) - I(u, t + 1)] \nabla I$ ,

which is a linear system we could solve directly for  $h$ . The solution of this system will be unreliable if the smallest eigenvalue of the symmetric positive semidefinite matrix  $\left[ \sum_{u \in \mathcal{P}_t} (\nabla I)(\nabla I)^T \right]$  is too small. This occurs when the image gradients in  $\mathcal{P}$  are all small—so the patch is featureless—or all point in one direction—so that we cannot localize the patch along that flow direction. If the estimate of  $h$  is unreliable, we must end the track. As Shi and Tomasi (1994) point out, this means that we can test the smallest eigenvalue of this matrix to tell whether a local window is worth tracking.

# Building tracks

- Start at scattered points in image 1
  - perhaps corner detector responses
- For each in image (n), compute position in (n+1)
  - as in previous slide
- Now check tracks
  - patch in (n+1) should look like an affine transform of patch in 1
  - Prune bad tracks

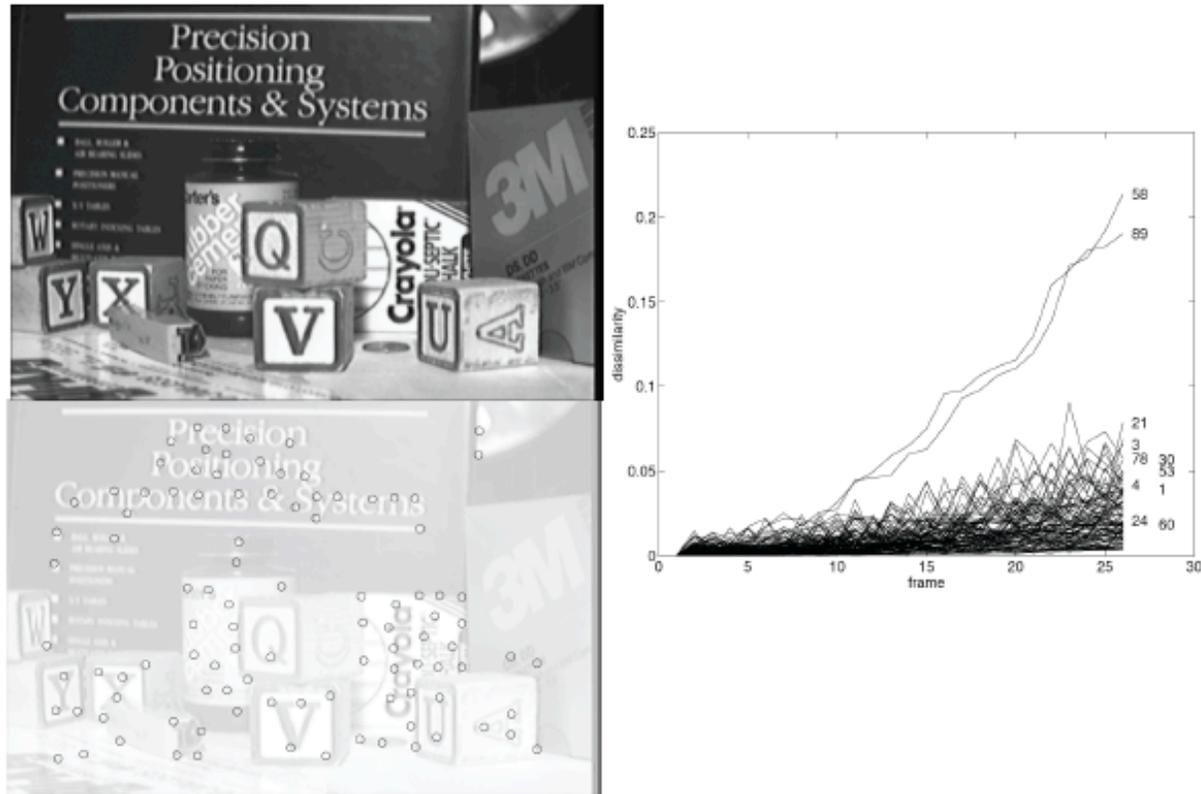


FIGURE 11.3: It is natural to track local neighborhoods, like those built in Section 5.3.2; however, for these neighborhoods to yield good tracks, they should pass a test of appearance complexity, shown in the text. This test checks that estimates of the translation of the neighborhood are stable. **Top left:** the first frame of an image sequence, with possible neighborhoods that pass this test shown on the **bottom left**. On the right, the sum-of-squared differences between the translated patch in frame  $n$  and the original in frame 1. Notice how this drifts up, meaning that the accumulated motion over many frames is *not* a translation; we need a better test to identify good tracks. *This figure was originally published as Figures 10, 11, 12 of “Good features to track,” by J. Shi and C. Tomasi, Proc. IEEE CVPR 1994, © IEEE, 1994.*

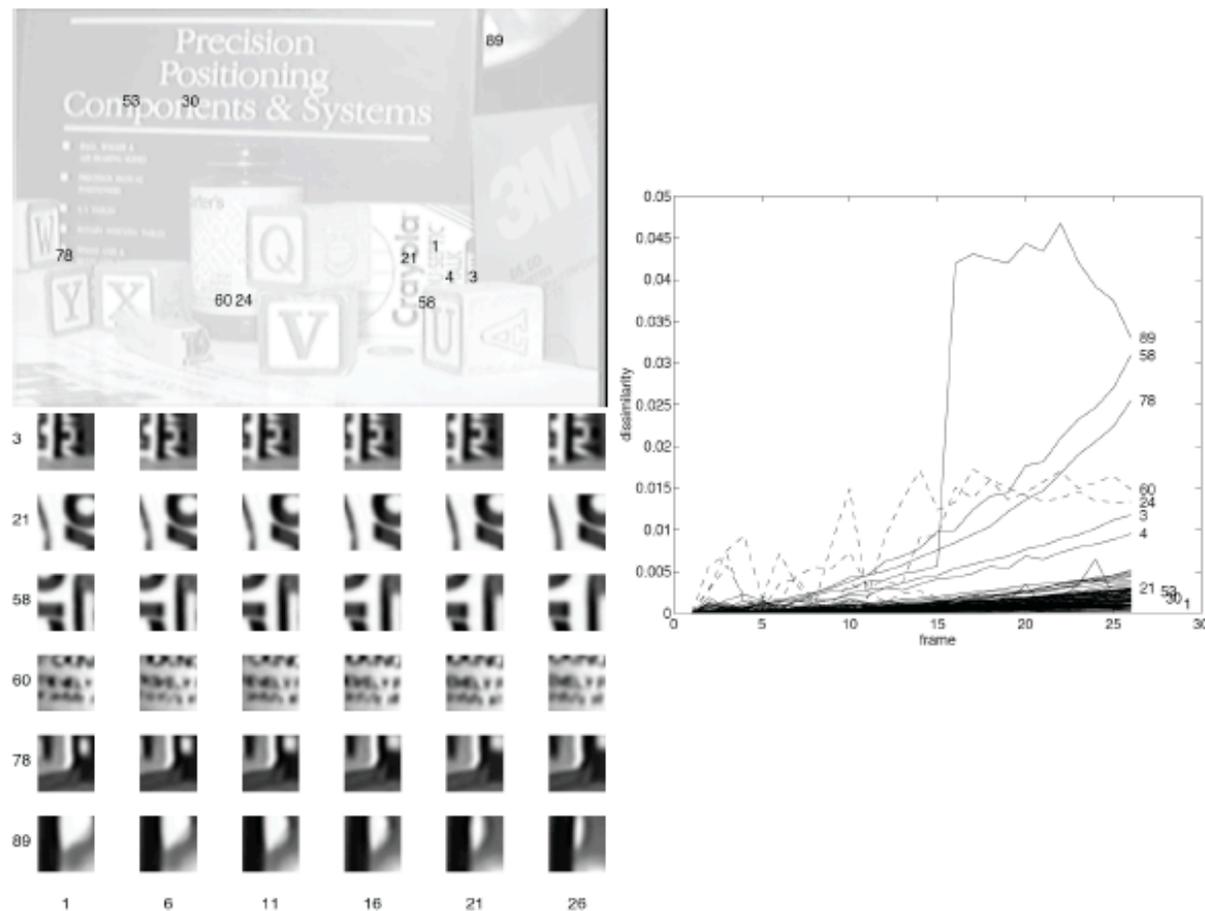


FIGURE 11.4: On the top left, the first frame of the sequence shown in Figure 11.3, with some neighborhoods overlaid. On the bottom left, the neighborhoods associated with these features (vertical) for different frames (horizontal). Notice how the pattern in the neighborhood deforms, perhaps because the object is rotating in 3D. This means that a translation model is good for the movement from frame  $n$  to frame  $n + 1$ , but does not explain the movement from frame 1 to frame  $n + 1$ . For this, we need to use an affine model. On the right, the value of the sum-of-squared differences between neighborhoods on a track in frame  $n$  and in frame 1, plotted against  $n$ . In this case, the neighborhood has been rectified by an affine transform, as in Section 11.1.3, before computing the SSD. Notice how some tracks are obviously good and others can be seen to have drifted. We could use this property to prune tracks. *This figure was originally published as Figures 13, 14, 15 of "Good features to track," by J. Shi and C. Tomasi, Proc. IEEE CVPR 1994, © IEEE, 1994.*

# But what if the patch deforms?

- Eg a football player's jersey
  - Colors are “similar” but SSD won't work
- Idea: patch histogram is stable
- To track:
  - repeat
    - predict location of new patch
    - search nearby for patch whose histogram matches original the best

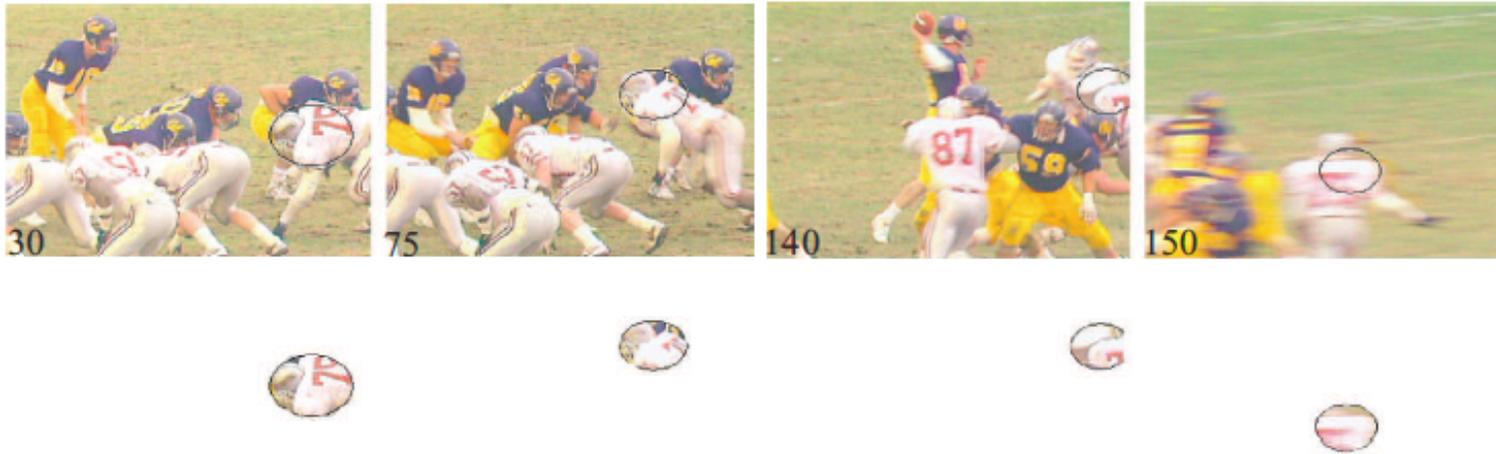


FIGURE 11.5: Four frames from a sequence depicting football players, with superimposed domains. The object to be tracked is the blob on top of player 78 (at the center right in frame 30). We have masked off these blobs (below) to emphasize just how strongly the pixels move around in the domain. Notice the motion blur in the final frame. These blobs can be matched to one another, and this is done by comparing histograms (in this case, color histograms), which are less affected by deformation than individual pixel values. *This figure was originally published as Figure 1 of “Kernel-Based Object Tracking” by D. Comaniciu, V. Ramesh, and P. Meer, IEEE Transactions on Pattern Analysis and Machine Intelligence, 2003, © IEEE 2003.*

Assume we have a sequence of  $N$  images; a domain  $\mathcal{D}_1$ , in the first image represented by parameters  $\mathbf{y}_1$  (for a circular domain of fixed size, these would be the location of the center; for a square, the center and edge length; and so on); a kernel function  $k$ ; a scale  $h$ ; and a feature representation  $\mathbf{f}$  of each pixel.

For  $n \in [1, \dots, N - 1]$

Obtain an initial estimate  $\mathbf{y}_{n+1}^{(0)}$  of the next domain either from a Kalman filter, or using  $\mathbf{y}_n$

Iterate until convergence

$$\mathbf{y}_{n+1}^{(j+1)} = \frac{\sum_i w_i \mathbf{x}_i g(\|\frac{\mathbf{x}_i - \mathbf{y}^{(j)}}{h}\|^2)}{\sum_i w_i g(\|\frac{\mathbf{x}_i - \mathbf{y}^{(j)}}{h}\|^2)}$$

where  $p_u$ ,  $k$ ,  $g$  are as given in the text

The track is the sequence of converged estimates  $\mathbf{y}_1, \dots, \mathbf{y}_N$ .

**Algorithm 11.2:** Tracking with the Mean Shift Algorithm.

# Q: Why no deep network methods?

- Silly A: I wrote these slides in 2011
- Better A: They've changed some feature mechanics
  - but nothing major
  - I'll work through some examples later

# When are motions “easy”?

- Current procedure
  - predict state
  - obtaining measurement from prediction by search
  - correct state
- Easy
  - When the object is close to where you expect it to be
    - eg Object guaranteed to move a little
- Large motions can be easy
  - When they’re “predictable”
    - e.g. ballistic motion
    - e.g. constant velocity
- Need a theory to fuse this procedure with motion model

# Q: why care about dynamics?

- Surely deep network methods make this go away?
- A: No they didn't
  - Some problems are “easy” with dynamical models, hard without
  - Example:
    - many similar, fast-moving cars

# Tracking - more formal view

- Very general model:
  - We assume there are moving objects, which have an underlying state  $X$
  - There are observations  $Y$ , some of which are functions of this state
  - There is a clock
    - at each tick, the state changes
    - at each tick, we get a new observation
- Examples
  - object is ball, state is 3D position+velocity, observations are stereo pairs
  - object is person, state is body configuration, observations are frames, clock is in camera (30 fps)

# Tracking - Probabilistic formulation

- Given
  - $P(X_{i-1}|Y_0, \dots, Y_{i-1})$ 
    - “Prior”
- We should like to know
  - $P(X_i|Y_0, \dots, Y_{i-1})$ 
    - “Predictive distribution”
  - $P(X_i|Y_0, \dots, Y_i)$ 
    - “Posterior”

## Key assumptions:

- **Only the immediate past matters:** formally, we require

$$P(\mathbf{X}_i | \mathbf{X}_1, \dots, \mathbf{X}_{i-1}) = P(\mathbf{X}_i | \mathbf{X}_{i-1})$$

This assumption hugely simplifies the design of algorithms, as we shall see; furthermore, it isn't terribly restrictive if we're clever about interpreting  $\mathbf{X}_i$  as we shall show in the next section.

- **Measurements depend only on the current state:** we assume that  $\mathbf{Y}_i$  is conditionally independent of all other measurements given  $\mathbf{X}_i$ . This means that

$$P(\mathbf{Y}_i, \mathbf{Y}_j, \dots, \mathbf{Y}_k | \mathbf{X}_i) = P(\mathbf{Y}_i | \mathbf{X}_i) P(\mathbf{Y}_j, \dots, \mathbf{Y}_k | \mathbf{X}_i)$$

Again, this isn't a particularly restrictive or controversial assumption, but it yields important simplifications.

# Tracking as Induction - base case

Firstly, we assume that we have  $P(\mathbf{X}_0)$

Then we have

$$\begin{aligned} P(\mathbf{X}_0 | \mathbf{Y}_0 = \mathbf{y}_0) &= \frac{P(\mathbf{y}_0 | \mathbf{X}_0) P(\mathbf{X}_0)}{P(\mathbf{y}_0)} \\ &= \frac{P(\mathbf{y}_0 | \mathbf{X}_0) P(\mathbf{X}_0)}{\int P(\mathbf{y}_0 | \mathbf{X}_0) P(\mathbf{X}_0) d\mathbf{X}_0} \\ &\propto P(\mathbf{y}_0 | \mathbf{X}_0) P(\mathbf{X}_0) \end{aligned}$$

# Tracking as induction - induction step

Given  $P(\mathbf{X}_{i-1}|\mathbf{y}_0, \dots, \mathbf{y}_{i-1})$ .

## Prediction

Prediction involves representing

$$P(\mathbf{X}_i|\mathbf{y}_0, \dots, \mathbf{y}_{i-1})$$

Notice this is  $i-1$   
current state based  
on previous  
measurements

Our independence assumptions make it possible to write

$$\begin{aligned} P(\mathbf{X}_i|\mathbf{y}_0, \dots, \mathbf{y}_{i-1}) &= \int P(\mathbf{X}_i, \mathbf{X}_{i-1}|\mathbf{y}_0, \dots, \mathbf{y}_{i-1})d\mathbf{X}_{i-1} \\ &= \int P(\mathbf{X}_i|\mathbf{X}_{i-1}, \mathbf{y}_0, \dots, \mathbf{y}_{i-1})P(\mathbf{X}_{i-1}|\mathbf{y}_0, \dots, \mathbf{y}_{i-1})d\mathbf{X}_{i-1} \\ &= \int P(\mathbf{X}_i|\mathbf{X}_{i-1})P(\mathbf{X}_{i-1}|\mathbf{y}_0, \dots, \mathbf{y}_{i-1})d\mathbf{X}_{i-1} \end{aligned}$$

# Tracking as induction - induction step

## Correction

Correction involves obtaining a representation of

$$P(\mathbf{X}_i | \mathbf{y}_0, \dots, \mathbf{y}_i)$$

Our independence assumptions make it possible to write

Notice this is i  
Prediction based on  
current measurement  
as well.

$$\begin{aligned} P(\mathbf{X}_i | \mathbf{y}_0, \dots, \mathbf{y}_i) &= \frac{P(\mathbf{X}_i, \mathbf{y}_0, \dots, \mathbf{y}_i)}{P(\mathbf{y}_0, \dots, \mathbf{y}_i)} \\ &= \frac{P(\mathbf{y}_i | \mathbf{X}_i, \mathbf{y}_0, \dots, \mathbf{y}_{i-1}) P(\mathbf{X}_i | \mathbf{y}_0, \dots, \mathbf{y}_{i-1}) P(\mathbf{y}_0, \dots, \mathbf{y}_{i-1})}{P(\mathbf{y}_0, \dots, \mathbf{y}_i)} \\ &= P(\mathbf{y}_i | \mathbf{X}_i) P(\mathbf{X}_i | \mathbf{y}_0, \dots, \mathbf{y}_{i-1}) \frac{P(\mathbf{y}_0, \dots, \mathbf{y}_{i-1})}{P(\mathbf{y}_0, \dots, \mathbf{y}_i)} \\ &= \frac{P(\mathbf{y}_i | \mathbf{X}_i) P(\mathbf{X}_i | \mathbf{y}_0, \dots, \mathbf{y}_{i-1})}{\int P(\mathbf{y}_i | \mathbf{X}_i) P(\mathbf{X}_i | \mathbf{y}_0, \dots, \mathbf{y}_{i-1}) d\mathbf{X}_i} \end{aligned}$$

# The Kalman Filter

- Assume that:
  - All state follows a linear dynamical model
  - Measurements are a linear function of state, plus noise
- Then (if first prior is Gaussian)
  - All PDF's are Gaussian
    - and so easy to represent
      - just need to keep track of mean and covariance
- The Kalman Filter correctly updates mean and covariance

# Linear models

Read this as:  $x_i$  is normally distributed.  
The mean is a linear function of  $x_{i-1}$  and  
whose variance is known (and can  
depend on  $i$ ).

$$x_i \sim N(\mathcal{D}_i x_{i-1}; \Sigma_{d_i})$$

$$y_i \sim N(\mathcal{M}_i x_i; \Sigma_{m_i})$$

Read this as:  $y_i$  is normally distributed.  
The mean is a linear function of  $x_i$  and  
whose variance is known (and can  
depend on  $i$ )

# Examples

- Dynamical models
  - Drifting points
    - new state = old state + gaussian noise
  - Points moving with constant velocity
    - new position = old position + (dt) old velocity + gaussian noise
    - new velocity = old velocity + gaussian noise
  - Points moving with constant acceleration
    - etc
- Measurement models
  - state = position; measurement = position + gaussian noise
  - state = position and velocity; measurement = position + gaussian noise
    - but we could infer velocity
  - state = position and velocity and acceleration; measurement = position + gaussian noise

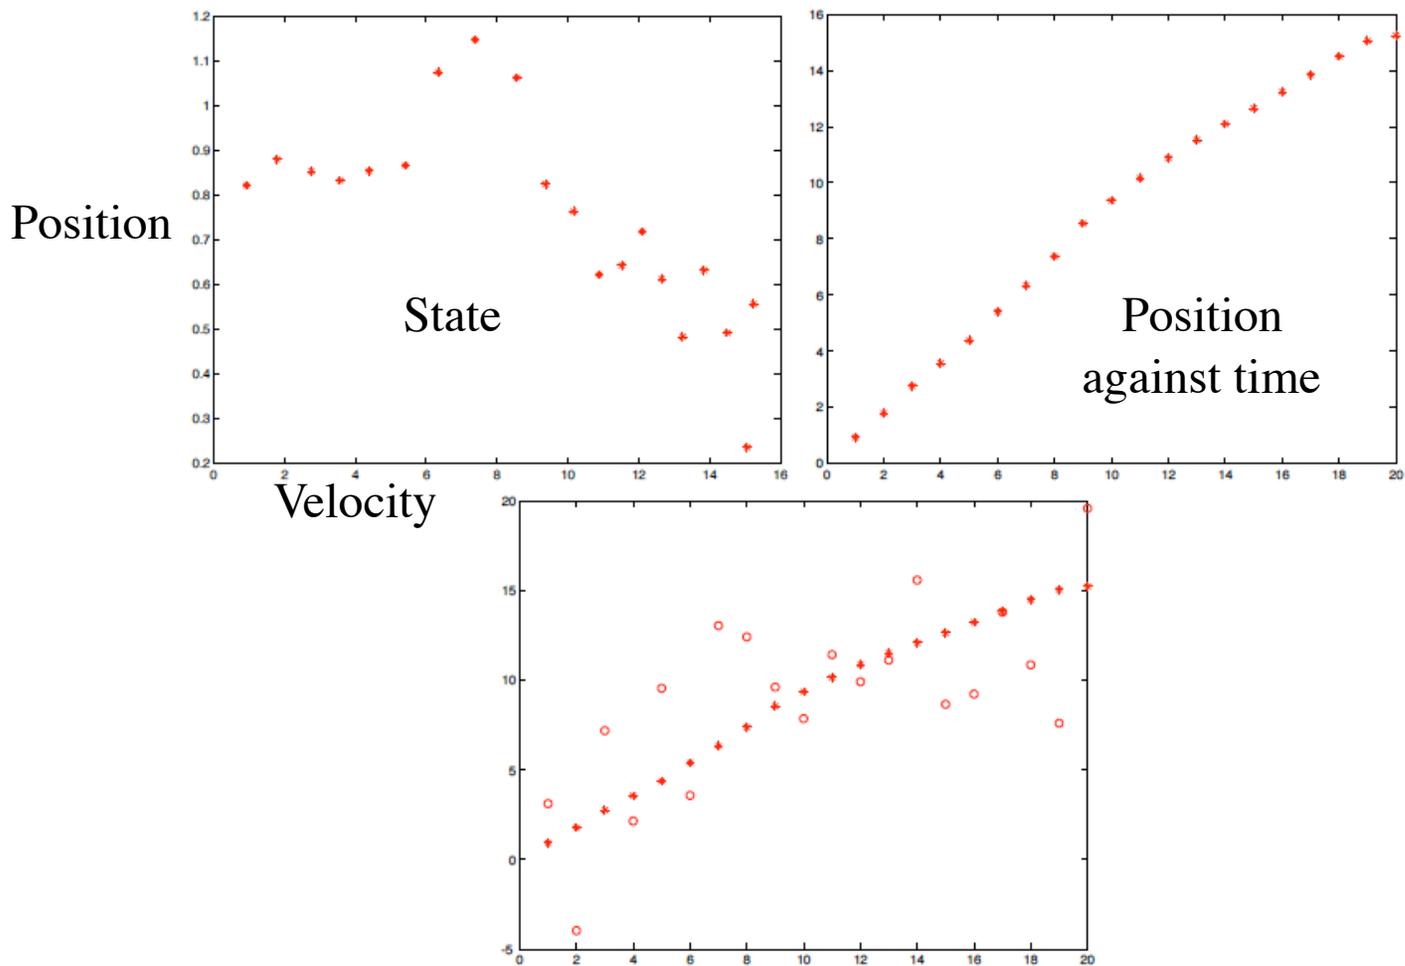


FIGURE 11.7: A constant velocity dynamic model for a point on the line. In this case, the state space is two dimensional, with one coordinate for position, one for velocity. The figure on the **top left** shows a plot of the state; each asterisk is a different state. Notice that the vertical axis (velocity) shows some small change compared with the horizontal axis. This small change is generated only by the random component of the model, so the velocity is constant up to a random change. The figure on the **top right** shows the first component of state (which is position) plotted against the time axis. Notice we have something that is moving with roughly constant velocity. The figure on the **bottom** overlays the measurements (the circles) on this plot. We are assuming that the measurements are of position only, and are quite poor; as we see, this doesn't significantly affect our ability to track.

# The Kalman Filter

- Dynamic Model

$$x_i \sim N(d_i x_{i-1}, \sigma_{d_i}^2)$$

$$y_i \sim N(m_i x_i, \sigma_{m_i}^2)$$

- Notation

mean of  $P(X_i | y_0, \dots, y_{i-1})$  as  $\bar{X}_i^-$

mean of  $P(X_i | y_0, \dots, y_i)$  as  $\bar{X}_i^+$

the standard deviation of  $P(X_i | y_0, \dots, y_{i-1})$  as  $\sigma_i^-$   
of  $P(X_i | y_0, \dots, y_i)$  as  $\sigma_i^+$ .

**Dynamic Model:**

$$\mathbf{x}_i \sim N(\mathcal{D}_i \mathbf{x}_{i-1}, \Sigma_{d_i})$$

$$\mathbf{y}_i \sim N(\mathcal{M}_i \mathbf{x}_i, \Sigma_{m_i})$$

**Start Assumptions:**  $\bar{\mathbf{x}}_0^-$  and  $\Sigma_0^-$  are known

**Update Equations: Prediction**

$$\bar{\mathbf{x}}_i^- = \mathcal{D}_i \bar{\mathbf{x}}_{i-1}^+$$

$$\Sigma_i^- = \Sigma_{d_i} + \mathcal{D}_i \Sigma_{i-1}^+ \mathcal{D}_i$$

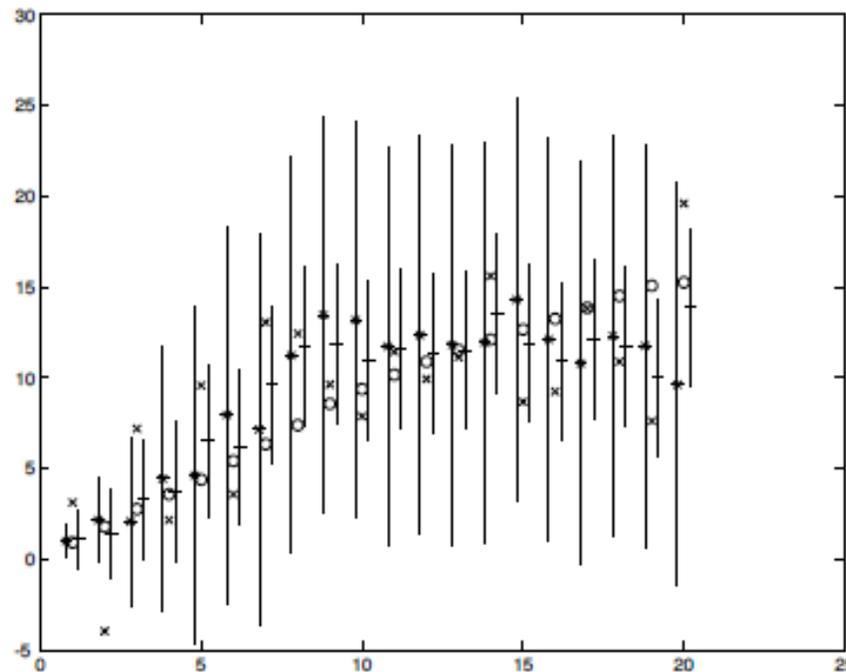
**Update Equations: Correction**

$$\mathcal{K}_i = \Sigma_i^- \mathcal{M}_i^T [\mathcal{M}_i \Sigma_i^- \mathcal{M}_i^T + \Sigma_{m_i}]^{-1}$$

$$\bar{\mathbf{x}}_i^+ = \bar{\mathbf{x}}_i^- + \mathcal{K}_i [\mathbf{y}_i - \mathcal{M}_i \bar{\mathbf{x}}_i^-]$$

$$\Sigma_i^+ = [Id - \mathcal{K}_i \mathcal{M}_i] \Sigma_i^-$$

**Algorithm 11.3:** The Kalman Filter.



Notice how uncertainty  
in state grows with  
movement and  
is reduced with  
measurement.

FIGURE 11.9: The Kalman filter for a point moving on the line under our model of constant velocity (compare with Figure 11.7). The state is plotted with open circles as a function of the step  $i$ . The \*s give  $\bar{x}_i^-$ , which is plotted slightly to the left of the state to indicate that the estimate is made before the measurement. The xs give the measurements, and the +s give  $\bar{x}_i^+$ , which is plotted slightly to the right of the state. The vertical bars around the \*s and the +s are three standard deviation bars, using the estimate of variance obtained before and after the measurement, respectively. When the measurement is noisy, the bars don't contract all that much when a measurement is obtained (compare with Figure 11.10).

# Data Association

- Nearest neighbours
  - choose the measurement with highest probability given predicted state
  - popular, but can lead to catastrophe
- Probabilistic Data Association
  - combine measurements, weighting by probability given predicted state
  - gate using predicted state

# Applying the Kalman filter

- Example: the jerseys
- Write a dynamical model
  - eg constant velocity
- Initialize
  - mark jersey in frame 1, or find interest points
- Track by iterating:
  - Predict state in frame  $n$  from previous state, dynamical model
  - Predict frame  $n$  location from frame  $n$  state
  - Search for best patch around that location - this is the measurement
  - Correct the state estimate using the measurement

# Tracking: Crucial points

- Careful image descriptions can make tracking easy
  - you track things with either
    - known appearance or
    - fixed appearance
- Clean probabilistic model for tracking with
  - Linear dynamics
  - Linear measurements
  - This is the Kalman Filter
- If dynamics or measurements are not linear
  - Representing probability distributions becomes very difficult