

Parametric (Global) Motion Models

Global motion models offer

- more constrained solutions than smoothness (Horn-Schunck)
- integration over a larger area than a translation-only model can accommodate (Lucas-Kanade)

Parametric (Global) Motion Models

2D Models:

(Translation)

Affine

Quadratic

Planar projective transform (Homography)

3D Models:

Instantaneous camera motion models

Homography+epipole

Plane+Parallax

- Transformations/warping of image

$$E(\mathbf{h}) = \sum_{\mathbf{x} \in R} [I(\mathbf{x} + \mathbf{h}) - I_0(\mathbf{x})]^2$$

Translations:


$$\mathbf{h} = \begin{bmatrix} \delta x \\ \delta y \end{bmatrix}$$

Generalizations

$$E(\mathbf{A}, \mathbf{h}) = \sum_{\mathbf{x} \in \mathbb{R}^2} [I(\mathbf{Ax} + \mathbf{h}) - I_0(\mathbf{x})]^2$$

Affine: $\mathbf{A} = \begin{bmatrix} a & b \\ c & d \end{bmatrix} \quad \mathbf{h} = \begin{bmatrix} \delta x \\ \delta y \end{bmatrix}$

Generalization



Affine: $\mathbf{A} = \begin{bmatrix} a & b \\ c & d \end{bmatrix} \quad \mathbf{h} = \begin{bmatrix} \delta x \\ \delta y \end{bmatrix}$

Example: Affine Motion

$$\begin{aligned} u(x, y) &= a_1 + a_2x + a_3y \\ v(x, y) &= a_4 + a_5x + a_6y \end{aligned}$$

Substituting into the B.C. Equation:

$$I_x \cdot u + I_y \cdot v + I_t \approx 0$$

$$I_x(a_1 + a_2x + a_3y) + I_y(a_4 + a_5x + a_6y) + I_t \approx 0$$


Each pixel provides 1 linear constraint in 6 *global* unknowns
(*minimum 6 pixels necessary*)

Least Square Minimization (over all pixels):

$$Err(\vec{a}) = \sum \left[I_x(a_1 + a_2x + a_3y) + I_y(a_4 + a_5x + a_6y) + I_t \right]$$

$$u(x, y) = a_0 + a_1 x + a_2 y$$

$$v(x, y) = a_3 + a_4 x + a_5 y$$



$$\mathbf{u}(\mathbf{x}; \mathbf{a}) = (u(x, y), v(x, y))$$

$I(\mathbf{x}, t-1)$



$\mathbf{x} + \mathbf{u}(\mathbf{x}; \mathbf{a})$

Warp





\mathbf{x}

$$I(\mathbf{x} + \mathbf{u}(\mathbf{x}; \mathbf{a}), t-1) = I(\mathbf{x}, t)$$

(Brightness Constancy Assumption)

Rigid pose estimation

- Head pose model: 6 DOF



Optic flow for rigid motion

- 3-D velocity:

$$V = T + \Omega \times P = T - \hat{\mathbf{P}}\Omega = \begin{bmatrix} \mathbf{I} & -\hat{\mathbf{P}} \end{bmatrix} \begin{bmatrix} T \\ \Omega \end{bmatrix}$$

$$V = \begin{bmatrix} V_x \\ V_y \\ V_z \end{bmatrix} = \begin{bmatrix} 1 & 0 & 0 & 0 & Z & -Y \\ 0 & 1 & 0 & -Z & 0 & X \\ 0 & 0 & 1 & Y & -X & 0 \end{bmatrix} \begin{bmatrix} T \\ \Omega \end{bmatrix}$$

$\hat{\mathbf{P}} = [\mathbf{P}_x]$ (skew-sym.)
--

Optic flow for rigid motion

- Rigid Motion (for small \mathbf{v}): $\begin{bmatrix} v_x \\ v_y \end{bmatrix} = \mathbf{H} \begin{bmatrix} T \\ \Omega \end{bmatrix}$

$$\mathbf{H} = \begin{bmatrix} f & 0 & -x \\ 0 & f & -y \end{bmatrix} \frac{1}{Z'} \begin{bmatrix} 1 & 0 & 0 & 0 & Z & -Y \\ 0 & 1 & 0 & -Z & 0 & X \\ 0 & 0 & 1 & Y & -X & 0 \end{bmatrix}$$

Perspective projection
of 3-D velocity

3-D velocity at point P

Direct Rigid Motion Estimation

- Brightness Change Constraint

$$I(x, y, t) = I(x + v_x, y + v_y, t + 1)$$

$$\frac{dI}{dx} v_x + \frac{dI}{dy} v_y + \frac{dI}{dt} = 0$$

$$\left[-\frac{dI}{dt} \right] = \begin{bmatrix} \frac{dI}{dx} & \frac{dI}{dy} \end{bmatrix} \begin{bmatrix} v_x \\ v_y \end{bmatrix}$$

Direct Rigid Motion Estimation

- Brightness Change Constraint

$$\begin{bmatrix} -\frac{dI}{dt} \end{bmatrix} = \begin{bmatrix} \frac{dI}{dx} & \frac{dI}{dy} \end{bmatrix} \begin{bmatrix} v_x \\ v_y \end{bmatrix}$$

- Rigid Motion Model

$$\begin{bmatrix} v_x \\ v_y \end{bmatrix} = \mathbf{H} \begin{bmatrix} T \\ \Omega \end{bmatrix}$$

$$\mathbf{H} = \begin{bmatrix} f & 0 & -x \\ 0 & f & -y \end{bmatrix} \frac{1}{Z'} \begin{bmatrix} 1 & 0 & 0 & 0 & Z & -Y \\ 0 & 1 & 0 & -Z & 0 & X \\ 0 & 0 & 1 & Y & -X & 0 \end{bmatrix}$$

Direct Motion Estimation

- One equation per pixel:

$$\begin{bmatrix} -\frac{dI}{dt} \end{bmatrix} = \begin{bmatrix} \frac{dI}{dx} & \frac{dI}{dy} \end{bmatrix} \begin{bmatrix} f & 0 & -x \\ 0 & f & -y \end{bmatrix} \frac{1}{Z'} \begin{bmatrix} 1 & 0 & 0 & 0 & Z & -Y \\ 0 & 1 & 0 & -Z & 0 & X \\ 0 & 0 & 1 & Y & -X & 0 \end{bmatrix} \begin{bmatrix} T \\ \Omega \end{bmatrix}$$

First, convert X,Y from screen coordinates to pixel coordinates....

Direct Motion Estimation

- One equation per pixel:

$$\begin{bmatrix} -\frac{dI}{dt} \end{bmatrix} = \begin{bmatrix} \frac{dI}{dx} & \frac{dI}{dy} \end{bmatrix} \begin{bmatrix} f & 0 & -x \\ 0 & f & -y \end{bmatrix} \frac{1}{Z'} \begin{bmatrix} 1 & 0 & 0 & 0 & Z & -yZ'/f \\ 0 & 1 & 0 & -Z & 0 & xZ'/f \\ 0 & 0 & 1 & yZ'/f & -xZ'/f & 0 \end{bmatrix} \begin{bmatrix} T \\ \Omega \end{bmatrix}$$

- Still hard!
- Z unknown; assume surface shape...
 - Negahdaripour & Horn - Planar
 - Black and Yacoob - Affine
 - Basu and Pentland; Bregler and Malik - Ellipsoidal
 - Essa et al. - Polygonal approximation
 - ...

“Direct Depth”

Use real-time stereo!

- Gives Z directly; no approximate model needed
- Express Direct Constraint on Depth Gradient

$$I(x, y, t) = I(x + v_x, y + v_y, t + 1)$$

$$Z(x, y, t) = Z(x + v_x, y + v_y, t + 1) - v_z$$

$$\frac{dZ}{dx} v_x + \frac{dZ}{dy} v_y + \frac{dZ}{dt} - v_z = 0$$

Direct Depth

Combined with rigid motion model:

- Orthographic

$$\begin{bmatrix} -dI/dt \\ -dZ/dt \end{bmatrix} = \begin{bmatrix} dI/dx & dI/dy & 0 \\ dZ/dx & dZ/dy & -1 \end{bmatrix} \begin{bmatrix} 1 & 0 & 0 & 0 & Z & -y \\ 0 & 1 & 0 & -Z & 0 & x \\ 0 & 0 & 1 & y & -x & 0 \end{bmatrix} \begin{bmatrix} T \\ \Omega \end{bmatrix}$$


- Perspective

$$\begin{bmatrix} -dI/dt \\ -dZ/dt \end{bmatrix} = \begin{bmatrix} f dI/dx & f dI/dy & -y dI/dy - x dI/dx \\ f dZ/dx & f dZ/dy & -1 \end{bmatrix} \frac{1}{Z'} \begin{bmatrix} 1 & 0 & 0 & 0 & Z & -yZ'/f \\ 0 & 1 & 0 & -Z & 0 & xZ'/f \\ 0 & 0 & 1 & yZ'/f & -xZ'/f & 0 \end{bmatrix} \begin{bmatrix} T \\ \Omega \end{bmatrix}$$

One system per pixel, same T, Ω . Solve with QR or SVD.

Application

Track users' head gaze for hands-free pointing...



The three main issues in tracking

- **Prediction:** we have seen $\mathbf{y}_0, \dots, \mathbf{y}_{i-1}$ — what state does this set of measurements predict for the i 'th frame? to solve this problem, we need to obtain a representation of $P(\mathbf{X}_i | \mathbf{Y}_0 = \mathbf{y}_0, \dots, \mathbf{Y}_{i-1} = \mathbf{y}_{i-1})$.
- **Data association:** Some of the measurements obtained from the i -th frame may tell us about the object's state. Typically, we use $P(\mathbf{X}_i | \mathbf{Y}_0 = \mathbf{y}_0, \dots, \mathbf{Y}_{i-1} = \mathbf{y}_{i-1})$ to identify these measurements.
- **Correction:** now that we have \mathbf{y}_i — the relevant measurements — we need to compute a representation of $P(\mathbf{X}_i | \mathbf{Y}_0 = \mathbf{y}_0, \dots, \mathbf{Y}_i = \mathbf{y}_i)$.

Tracking

- Very general model:
 - We assume there are moving objects, which have an underlying state X
 - There are measurements 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, measurements are stereo pairs
 - object is person, state is body configuration, measurements are frames, clock is in camera (30 fps)

Three main steps

- **Prediction:** we have seen $\mathbf{y}_0, \dots, \mathbf{y}_{i-1}$ — what state does this set of measurements predict for the i 'th frame? to solve this problem, we need to obtain a representation of $P(\mathbf{X}_i | \mathbf{Y}_0 = \mathbf{y}_0, \dots, \mathbf{Y}_{i-1} = \mathbf{y}_{i-1})$.
- **Data association:** Some of the measurements obtained from the i -th frame may tell us about the object's state. Typically, we use $P(\mathbf{X}_i | \mathbf{Y}_0 = \mathbf{y}_0, \dots, \mathbf{Y}_{i-1} = \mathbf{y}_{i-1})$ to identify these measurements.
- **Correction:** now that we have \mathbf{y}_i — the relevant measurements — we need to compute a representation of $P(\mathbf{X}_i | \mathbf{Y}_0 = \mathbf{y}_0, \dots, \mathbf{Y}_i = \mathbf{y}_i)$.

Simplifying 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.

Assumptions allow recursive solutions

Decompose estimation problem

- part that depends on new observation
- part that can be computed from previous history

E.g., running average:

$$a_t = \frac{1}{t} \sum_{i=1:t} y_i = \frac{1}{t} \sum_{i=1:(t-1)} y_i + \frac{1}{t} y_t = \frac{t-1}{t} \frac{1}{t-1} \sum_{i=1:(t-1)} y_i + \frac{1}{t} y_t$$

$$a_t = \frac{t-1}{t} a_{t-1} + \frac{1}{t} y_t$$

Now in form that allows recursive application

Tracking as induction

- Assume data association is done
 - a dangerous assumption--assumes good segmentation
- Do correction for the 0'th frame
- Assume we have corrected estimate for i 'th frame
 - show we can do prediction for $i+1$, correction for $i+1$

Base case

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

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

Induction step

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

Prediction

Prediction involves representing

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

Our independence assumptions make it possible to write

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

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

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

Linear dynamic models

- Use notation \sim to mean “has the pdf of”, $N(a, b)$ is a normal distribution with mean a and covariance b .
- Then a linear dynamic model has the form

State Dynamics

$$\mathbf{x}_i = N(\mathbf{D}_{i-1}\mathbf{x}_{i-1}; \Sigma_{d_i})$$

$$\mathbf{y}_i = N(\mathbf{M}_i\mathbf{x}_i; \Sigma_{m_i})$$

Measurement Dynamics

- This is much, much more general than it looks, and extremely powerful

Examples

- Drifting points
 - we assume that the new position of the point is the old one, plus noise.
 - For the measurement model, we may not need to observe the whole state of the object
 - e.g. a point moving in 3D, at the $3k$ 'th tick we see x , $3k+1$ 'th tick we see y , $3k+2$ 'th tick we see z
 - in this case, we can still make decent estimates of **all three** coordinates at each tick.
 - This property, which does not apply to every model, is called **Observability**

Examples

- Points moving with constant velocity
- Periodic motion
- Etc.
- Points moving with constant acceleration

Points moving with constant velocity

- We have

$$u_i = u_{i-1} + \Delta t v_{i-1} + \epsilon_i$$

$$v_i = v_{i-1} + \epsilon_i$$

– (the Greek letters denote noise terms)

- Stack (u, v) into a single state vector

$$\begin{bmatrix} u \\ v \end{bmatrix}_i = \begin{bmatrix} 1 \\ 0 \end{bmatrix} \Delta t \begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix} \begin{bmatrix} u \\ v \end{bmatrix}_{i-1} + \text{noise}$$

– which is the form we had above

Points moving with constant acceleration

- We have

$$u_i = u_{i-1} + \Delta t v_{i-1} + \epsilon_i$$

$$v_i = v_{i-1} + \Delta t a_{i-1} + \epsilon_i$$



$$a_i = a_{i-1} + \epsilon_i$$

– (the Greek letters denote noise terms)

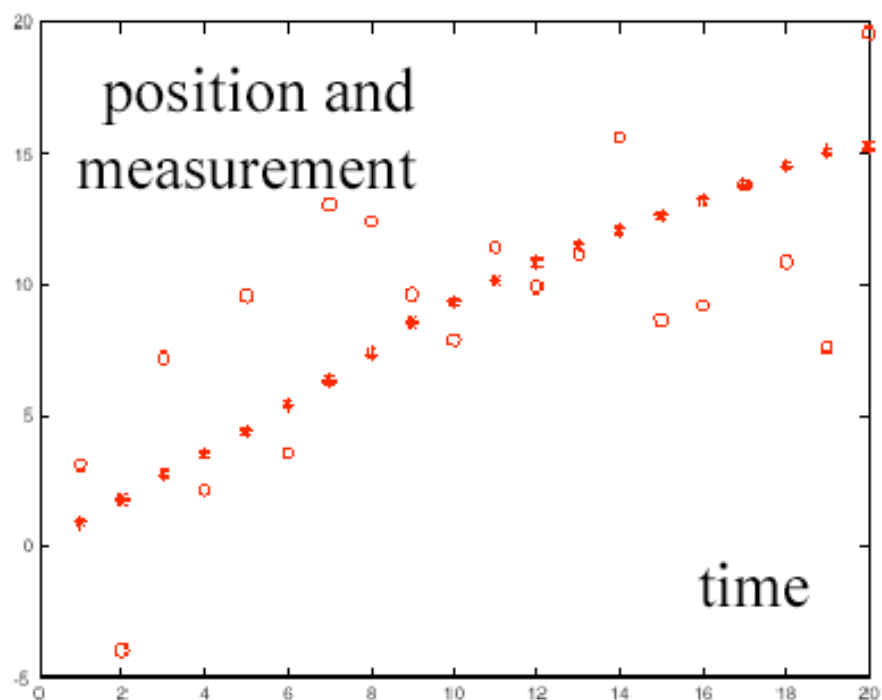
- Stack (u, v) into a single state vector



$$\begin{bmatrix} u \\ v \\ a \end{bmatrix}_i = \begin{bmatrix} 1 & \Delta t & 0 \\ 0 & 1 & \Delta t \\ 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} u \\ v \\ a \end{bmatrix}_{i-1} + \text{noise}$$

– which is the form we had above



Constant
Velocity
Model






Constant
Acceleration
Model

The Kalman Filter

- Key ideas:
 - Linear models interact uniquely well with Gaussian noise - make the prior Gaussian, everything else Gaussian and the calculations are easy
 - Gaussians are really easy to represent --- once you know the mean and covariance, you're done

The Kalman Filter



The Kalman Filter in 1D

- Dynamic Model

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

- Notation

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


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

Corrected mean — 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 σ_i^-

of $P(X_i | y_0, \dots, y_i)$ as σ_i^+

The Kalman Filter



Prediction for 1D Kalman filter

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

- The new state is obtained by
 - multiplying old state by known constant
 - adding zero-mean noise
- Therefore, predicted mean for new state is
 - constant times mean for old state
- Old variance is normal random variable
 - variance is multiplied by square of constant
 - and variance of noise is added.

$$\bar{X}_i^- = d_i \bar{X}_{i-1}^+$$

$$(\sigma_i^-)^2 = \sigma_{d_i}^2 + (d_i \sigma_{i-1}^+)^2$$

Dynamic Model:

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

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

Start Assumptions: \bar{x}_0^- and σ_0^- are known

Update Equations: Prediction

$$\bar{x}_i^- = d_i \bar{x}_{i-1}^+$$


$$\sigma_i^- = \sqrt{\sigma_{d_i}^2 + (d_i \sigma_{i-1}^+)^2}$$

Update Equations: Correction

$$x_i^+ = \left(\frac{\bar{x}_i^- \sigma_{m_i}^2 + m_i y_i (\sigma_i^-)^2}{\sigma_{m_i}^2 + m_i^2 (\sigma_i^-)^2} \right)$$

$$\sigma_i^+ = \sqrt{\left(\frac{\sigma_{m_i}^2 (\sigma_i^-)^2}{(\sigma_{m_i}^2 + m_i^2 (\sigma_i^-)^2)} \right)}$$

The Kalman Filter



Correction for 1D Kalman filter

We have

$$\begin{aligned}P(X_i|y_0, \dots, y_i) &= \frac{P(y_i|X_i)P(X_i|y_0, \dots, y_{i-1})}{\int P(y_i|X_i)P(X_i|y_0, \dots, y_{i-1})dX_i} \\&\propto P(y_i|X_i)P(X_i|y_0, \dots, y_{i-1})\end{aligned}$$

$$g(x; \mu, v) = \exp\left(-\frac{(x - \mu)^2}{2v}\right)$$

$$\begin{aligned}P(X_i|y_0, \dots, y_i) &\propto g(y_i; m_i X_i, \sigma_{m_i}^2)g(X_i; \bar{X}_i^-, (\sigma_i^-)^2) \\&= g(m_i X_i; y_i, \sigma_{m_i}^2)g(X_i; \bar{X}_i^-, (\sigma_i^-)^2) \\&= g(X_i; \frac{y_i}{m_i}, \frac{\sigma_{m_i}^2}{m_i^2})g(X_i; \bar{X}_i^-, (\sigma_i^-)^2)\end{aligned}$$

Dynamic Model:

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

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

Start Assumptions: \bar{x}_0^- and σ_0^- are known

Update Equations: Prediction

$$\bar{x}_i^- = d_i \bar{x}_{i-1}^+$$

$$\sigma_i^- = \sqrt{\sigma_{d_i}^2 + (d_i \sigma_{i-1}^+)^2}$$

Update Equations: Correction

$$x_i^+ = \left(\frac{\bar{x}_i^- \sigma_{m_i}^2 + m_i y_i (\sigma_i^-)^2}{\sigma_{m_i}^2 + m_i^2 (\sigma_i^-)^2} \right)$$

$$\sigma_i^+ = \sqrt{\left(\frac{\sigma_{m_i}^2 (\sigma_i^-)^2}{(\sigma_{m_i}^2 + m_i^2 (\sigma_i^-)^2)} \right)}$$

Correction for 1D Kalman filter

- Pattern match to identities given in book
 - basically, guess the integrals, get:

$$x_i^+ = \left(\frac{\bar{x}_i^- \sigma_{m_i}^2 + m_i y_i (\sigma_i^-)^2}{\sigma_{m_i}^2 + m_i^2 (\sigma_i^-)^2} \right)$$

- Notice:
 - if measurement noise is small
we rely mainly on the measurer
if it's large, mainly on the
prediction

$$\sigma_i^+ = \sqrt{\left(\frac{\sigma_{m_i}^2 (\sigma_i^-)^2}{(\sigma_{m_i}^2 + m_i^2 (\sigma_i^-)^2)} \right)}$$

In higher dimensions, derivation follows the same lines, but isn't as easy. Expressions here.

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: $\overline{\mathbf{x}}_0^-$ and Σ_0^- are known

Update Equations: Prediction

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



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

Update Equations: Correction


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


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



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




Constant
Velocity
Model









Constant
Acceleration
Model




Smoothing

- Idea
 - We don't have the best estimate of state - what about the future?
 - Run two filters, one moving forward, the other backward in time.
 - Now combine state estimates
 - The crucial point here is that we can obtain a smoothed estimate by viewing the backward filter's prediction as yet another measurement for the forward filter
 - so we've already done the equations








n-D

Generalization to n-D is straightforward but more complex.


n-D

Generalization to n-D is straightforward but more complex.



n-D Prediction

Generalization to n-D is straightforward but more complex.



Prediction:

- Multiply estimate at prior time with forward model:


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

- Propagate covariance through model and add new noise:

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

n-D Correction

Generalization to n-D is straightforward but more complex.



Correction:

- Update *a priori* estimate with measurement to form *a posteriori*

n-D correction

Find linear filter on innovations

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

which minimizes *a posteriori* error covariance:

$$E\left[\left(x - \overline{x}^+\right)^T \left(x - \overline{x}^+\right)\right]$$

\mathbf{K} is the *Kalman Gain* matrix. A solution is

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

Kalman Gain Matrix

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

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

As measurement becomes more reliable, K weights residual more heavily,

$$\lim_{\Sigma_m \rightarrow 0} K_i = M^{-1}$$

As prior covariance approaches 0, measurements are ignored:

$$\lim_{\Sigma_i^- \rightarrow 0} K_i = 0$$

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: $\overline{\mathbf{x}}_0^-$ and Σ_0^- are known

Update Equations: Prediction

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


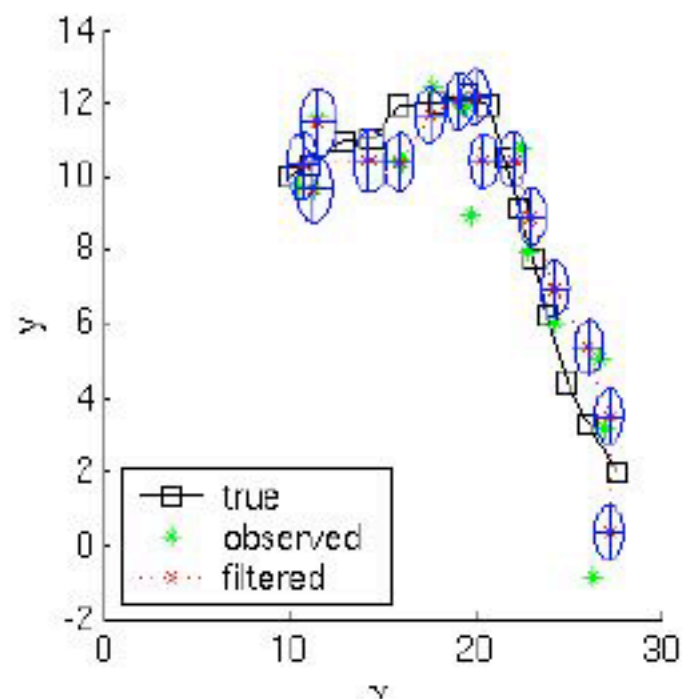
$$\Sigma_i^- = \Sigma_{d_i} + \mathcal{D}_i \Sigma_{i-1}^+ \mathcal{D}_i$$

Update Equations: Correction

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

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


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







2-D constant velocity example from Kevin Murphy's Matlab toolbox

- MSE of filtered estimate is 4.9; of smoothed estimate. 3.2.
- Not only is the smoothed estimate better, but we know that it is better, as illustrated by the smaller uncertainty ellipses
- Note how the smoothed ellipses are larger at the ends, because these points have seen less data.
- Also, note how rapidly the filtered ellipses reach their steady-state (“Ricatti”) values.


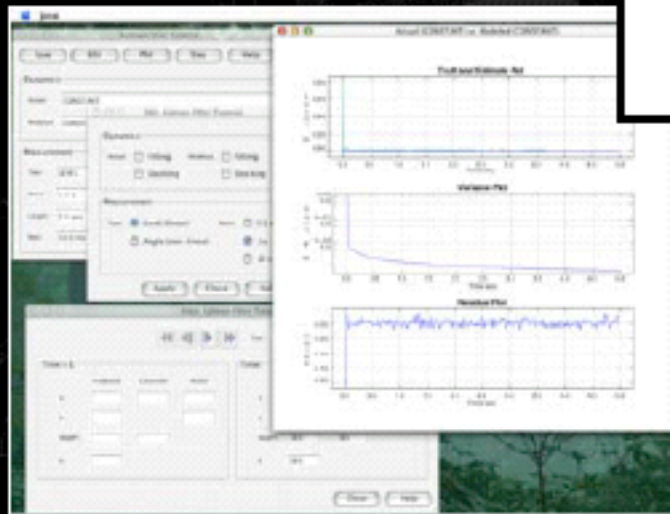
[figure from <http://www.ai.mit.edu/~murphyk/Software/Kalman/kalman.html>]





Online demo

- On-line 1D simulation
- Linear and non-linear
- Variable dynamics



<http://www.cs.unc.edu/~welch/kalman/>

SIGGRAPH
2001
EXPLORE INTERACTION
AND DIGITAL IMAGES

[figure from Welsh and Bishop 2001]


Online demo

The *Kalman Filter Learning Tool* tool simulates a relatively simple example setup involving estimation of the water level in a tank.

Water dynamics. The user can independently choose both the actual and modeled dynamics of the water. The choices include no motion (the default), filling, sloshing, or both filling and sloshing.

Measurement model. The user can also choose the method of measurement. The measurement model choices include two options that are commonly used (for example) in toilet tanks: a vertical level (linear) float-type sensor, or an angular (non-linear) float-type sensor. A diagram depicting the two case is shown below. The user is also allowed to increase or decrease (by a factor of 10) the magnitude of the random linear or angular measurement noise.

Online demo



[figure from <http://www.cs.unc.edu/~welch/kalman/kftool/index.html>]

Online demo

<http://www.cs.unc.edu/~welch/kalman/kftool/KalmanFilterApplet.html>

Abrupt changes

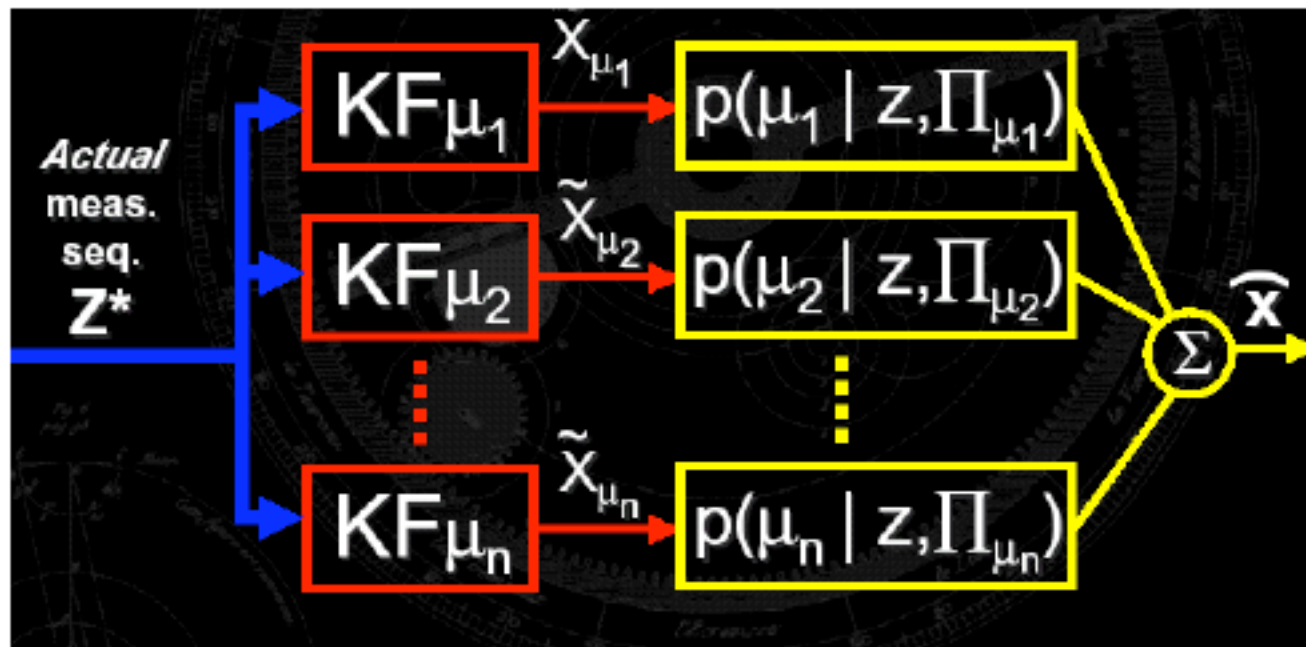
What if environment is sometimes unpredictable?

Do people move with constant velocity?

Test several models of assumed dynamics, use the best.

Multiple model filters

Test several models of assumed dynamics



[figure from Welsh and Bishop 2001]

Resources

- Kalman filter homepage

<http://www.cs.unc.edu/~welch/kalman/>

- Kevin Murphy's Matlab toolbox:

<http://www.ai.mit.edu/~murphyk/Software/Kalman/kalman.html>

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

