

# Tracking

COS 429

Princeton University

# Tracking

- Feature tracking
- Object tracking

# Tracking

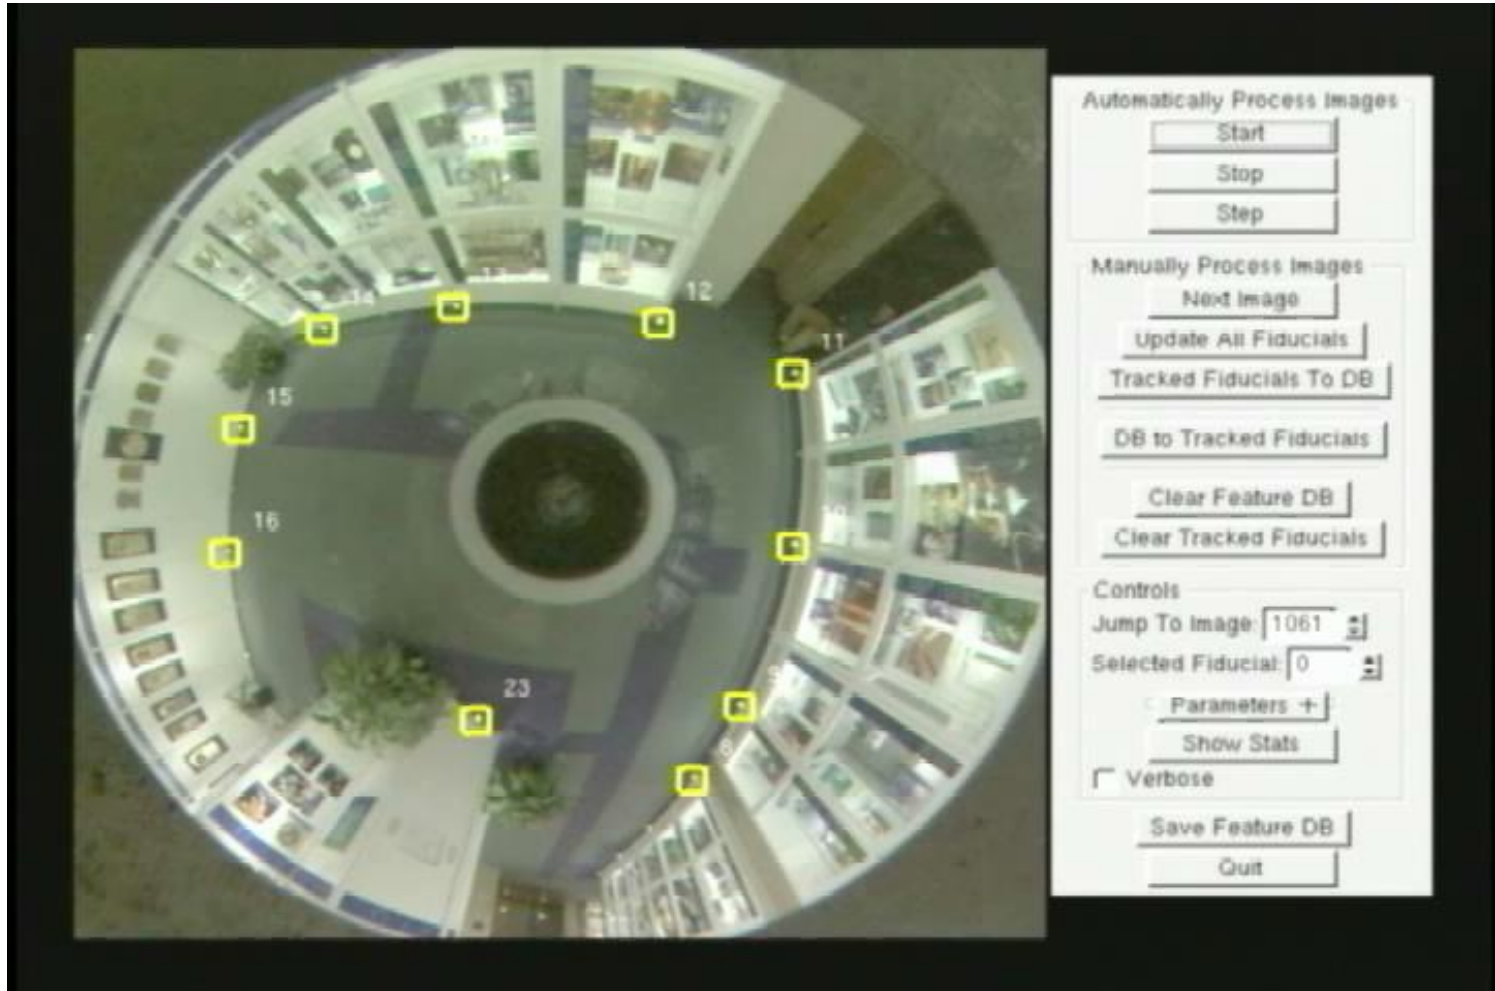
- Feature tracking ←
- Object tracking

# Feature Tracking

- Given sequence of images
- Find feature correspondences

# Applications?

# Feature Tracking Example



Sea of Images

# Feature Tracking Example



# Tracking

- Feature tracking
- Object tracking ←



# Object Tracking

- Given sequence of images
- Track moving foreground objects

# Object Tracking

Image 1




Image 2




Image 3





Image 4



- Can we estimate the position of the object?
- Can we predict future positions?


# Applications?

# Applications

- Astronomy, biology, etc.
- Surveillance
- Activity analysis
- Gesture recognition
- etc.

# Applications


- Astronomy, biology, etc.
- Surveillance
- Activity analysis
- Gesture recognition
- etc.



<http://www.youtube.com/watch?v=SLYgvHzAm2w>

# Applications


- Astronomy, biology, etc.
- **Surveillance**
- Activity analysis
- Gesture recognition
- etc.



<http://www.youtube.com/watch?v=bY8qGk45WxM>

# Applications

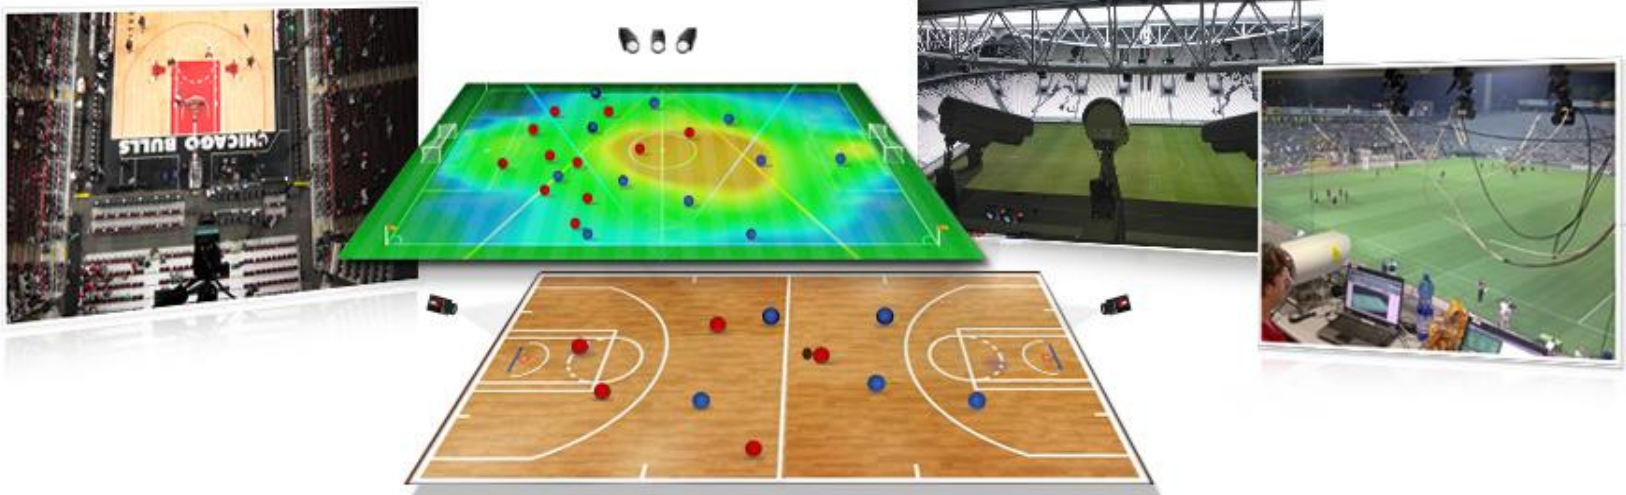
- Astronomy, biology, etc.
- Surveillance
- **Activity analysis**
- Gesture recognition
- etc.



wavestore

# Applications

- Astronomy, biology, etc.
- Surveillance
- **Activity analysis**
- Gesture recognition



SportVU




# Applications

- Astronomy, biology, etc.
- Security surveillance
- Activity analysis
- **Gesture recognition**
- etc.




(a)



(b)




(c)



# Methods?


# General Strategy

- Initialize *model* in the first frame
- Given model estimate for frame  $t-1$ :
  - *Predict* for frame  $t$ 
    - Use *dynamics model* of how the image changes
  - *Correct* for frame  $t$ 
    - Use foreground estimation in current frame to update model



# General Strategy

- Initialize *model* in the first frame
- Given model estimate for frame  $t-1$ :
  - *Predict* for frame  $t$ 
    - Use *dynamics model* of how the image changes
  - *Correct* for frame  $t$ 
    - Use *foreground estimation* in current frame to update model



# Outline

- Feature tracking
- Object tracking
  - Foreground estimation ←
  - Model update

# Foreground Estimation

Image at time  $t$ :

$$I(x, y, t)$$



Foreground at time  $t$ :

$$F(x, y, t)$$



How?

# Background Subtraction

Image at time  $t$ :

$$I(x, y, t)$$



Background at time  $t$ :

$$B(x, y, t)$$



—

|  $> Th$

1. Estimate the background for time  $t$ .
2. Subtract the estimated background from the input frame.
3. Apply a threshold,  $Th$ , to the absolute difference to get the **foreground mask**.

# Background Subtraction

Image at time  $t$ :

$$I(x, y, t)$$



Foreground at time  $t$ :

$$F(x, y, t)$$



1. Estimate the background for time  $t$ .
2. Subtract the estimated background from the input frame.
3. Apply a threshold,  $Th$ , to the absolute difference to get the **foreground mask**.



# Background Subtraction

Image at time  $t$ :

$$I(x, y, t)$$



Background at time  $t$ :

$$B(x, y, t)$$



How estimate the background?

# Background = Previous Frame

- ▶ Background is estimated to be the previous frame. Background subtraction equation then becomes:

$$B(x, y, t) = I(x, y, t - 1)$$



$$|I(x, y, t) - I(x, y, t - 1)| > Th$$

- ▶ Depending on the object structure, speed, frame rate and global threshold, this approach may or may **not** be useful (usually **not**).



—



| > Th

# Background = Previous Frame

$Th = 25$




$Th = 50$



$Th = 100$



$Th = 200$



# Background = Mean Filter

- ▶ In this case the background is the mean of the previous  $n$  frames:

$$B(x, y, t) = \frac{1}{n} \sum_{i=0}^{n-1} I(x, y, t - i)$$

↓

$$|I(x, y, t) - \frac{1}{n} \sum_{i=0}^{n-1} I(x, y, t - i)| > Th$$

- ▶ For  $n = 10$ :

Estimated Background



Foreground Mask



# Background = Mean Filter

- When won't this work?



# Background = Mean Filter

Test Image



Chair  
moved

Light  
gradually  
brightened

Light  
just  
switched  
on

Tree  
Waving

Foreground  
covers  
monitor  
pattern

No clean  
background  
training

Interior  
motion  
undetectable

Ideal  
Foreground



Adjacent  
Frame  
Difference




Mean &  
Threshold




- Toyama et al. 1999

# Background = Mixture Model



- Model each background pixel with a *mixture* of Gaussians; update parameters over time.




(a)



(b)



(c)



# Background = Median Filter

- ▶ Assuming that the background is more likely to appear in a scene, we can use the median of the previous  $n$  frames as the background model:

$$B(x, y, t) = \text{median}\{I(x, y, t - i)\}$$



$$|I(x, y, t) - \text{median}\{I(x, y, t - i)\}| > Th \text{ where } i \in \{0, \dots, n - 1\}.$$

- ▶ For  $n = 10$ :

Estimated Background



Foreground Mask





# Comparison



Mean



Median

# Background Subtraction Result



-



=



# Background Subtraction

## **Advantages:**

- Extremely easy to implement and use!
- All pretty fast.
- Corresponding background models need not be constant, they change over time.

## **Disadvantages:**

- Accuracy of frame differencing depends on object speed and frame rate
- Median background model: relatively high memory requirements.
- Setting global threshold Th...

*How could this approach be better?*

# Discriminative Models

- Use adaptive models of both foreground and background to estimate  $p(\text{foreground})$ 
  - Color histograms
  - Segmentation algorithms
  - etc.
- Potential problem: poor separation of foreground and background causes poor update to discriminative models
  - Drift


# Outline

- Feature tracking
- Object tracking
  - Foreground estimation
  - Model update

# Object Tracking

- Given Sequence of Images
- Track centers of moving foreground objects

Why isn't foreground estimation enough?




<http://www.youtube.com/watch?v=bY8qGk45WxM>



# Object Tracking

- Use *model* of object motion to compensate for noisy foreground estimation, handle occlusions, keep track of multiple foreground objects, etc.



<http://www.youtube.com/watch?v=bY8qGk45WxM>

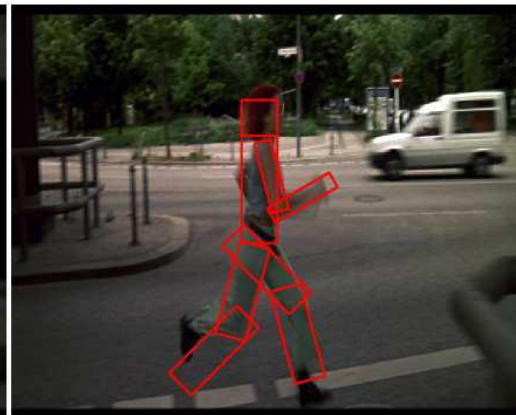
# Example Object Models



points



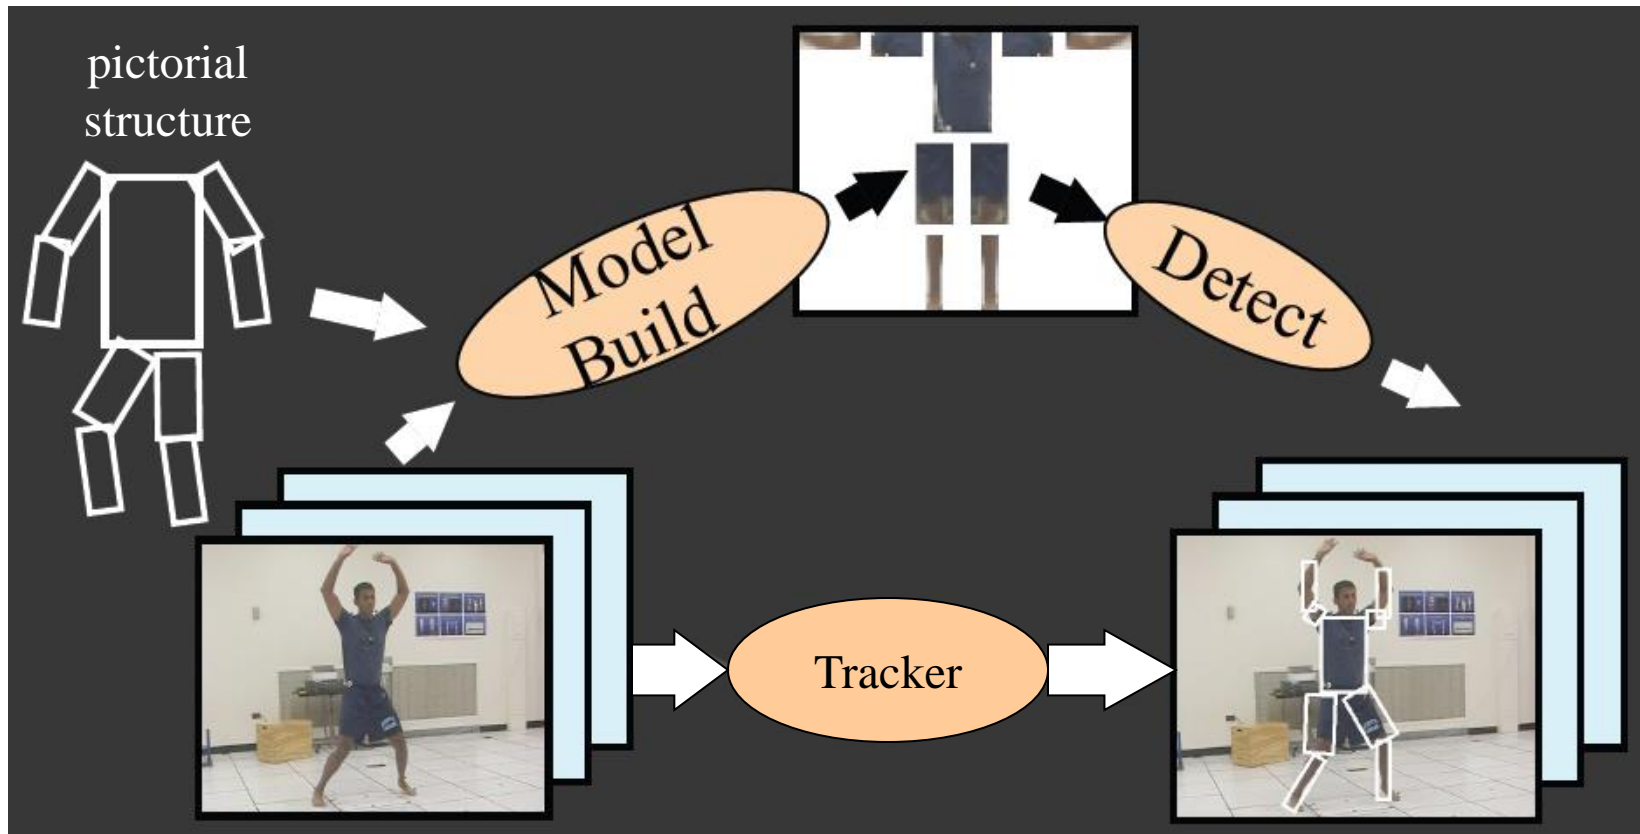
curves



Part-based structures



# Example Object Models




D. Ramanan, D. Forsyth, and A. Zisserman. [Tracking People by Learning their Appearance](#). PAMI 2007.

# Model Update

- Continually update probabilistic parameters of the object model based on observations (foreground estimates)


# Simple Example

- Measurement of a single point  $z_1$
- Variance  $\sigma_1^2$  (uncertainty  $\sigma_1$ )
  - Assuming Gaussian distribution




# Simple Example

- Measurement of a single point  $z_1$ , variance  $\sigma_1^2$ 
  - Assuming Gaussian distribution
- Best estimate of true position:  $\hat{x}_1 = z_1$
- Uncertainty in best estimate:  $\hat{\sigma}_1^2 = \sigma_1^2$




# Simple Example

- Second measurement  $z_2$ , variance  $\sigma_2^2$
- Best estimate of true position?
- Uncertainty in best estimate?



# Simple Example

- Second measurement  $z_2$ , variance  $\sigma_2^2$
- Best estimate of true position?
- Uncertainty in best estimate?



# Simple Example

- Best estimate of true position:  
(weighted average)

$$\begin{aligned}\hat{x}_2 &= \frac{\frac{1}{\sigma_1^2} z_1 + \frac{1}{\sigma_2^2} z_2}{\frac{1}{\sigma_1^2} + \frac{1}{\sigma_2^2}} \\ &= \hat{x}_1 + \frac{\hat{\sigma}_1^2}{\hat{\sigma}_1^2 + \sigma_2^2} (z_2 - \hat{x}_1)\end{aligned}$$

- Uncertainty in best estimate

$$\hat{\sigma}_2^2 = \frac{1}{\frac{1}{\hat{\sigma}_1^2} + \frac{1}{\sigma_2^2}}$$


# Possible Update Strategy

- Online Weighted Average
  - Combine successive measurements into constantly-improving estimate
  - Uncertainty decreases over time
  - Only need to keep current measurement, last estimate of state and uncertainty



# Kalman Filter

- Assume measurement errors are Gaussian
- Assume model parameters are Gaussian
- Assume model is linear
  
- Kalman filter provides optimal update strategy



**Rudolf Emil Kalman**

# Kalman Filter Terminology

- System model:

$$\mathbf{x}_k = \Phi_{k-1} \mathbf{x}_{k-1} + \boldsymbol{\xi}_{k-1}$$

- $\hat{\mathbf{x}}_k$  is estimate of state  $\hat{\mathbf{x}}$  with covariance  $P$
- The matrix  $\Phi_k$  is *state transition matrix*
- The vector  $\boldsymbol{\xi}_k$  represents *additive noise*, assumed to have covariance  $Q$

# Kalman Filter Terminology

- Measurement model:

$$z_k = H_k x_k + \mu_k$$

- Matrix  $H$  is *measurement matrix*
- The vector  $\mu$  is *measurement noise*, assumed to have covariance  $R$

# Kalman Filter Update

- Predict new state

$$\mathbf{x}'_k = \Phi_{k-1} \hat{\mathbf{x}}_{k-1}$$


$$\mathbf{P}'_k = \Phi_{k-1} \mathbf{P}_{k-1} \Phi_{k-1}^T + \mathbf{Q}_{k-1}$$

- Correct model based on new measurements

$$\hat{\mathbf{x}}_k = \mathbf{x}'_k + \mathbf{K}_k (\mathbf{z}_k - \mathbf{H}_k \mathbf{x}'_k)$$


$$\mathbf{P}_k = (\mathbf{I} - \mathbf{K}_k \mathbf{H}_k) \mathbf{P}'_k$$

# The Prediction-Correction-Cycle




$$\overline{bel}(x_t) = \begin{cases} \bar{\mu}_t = a_t \mu_{t-1} + b_t u_t \\ \bar{\sigma}_t^2 = a_t^2 \sigma_t^2 + \sigma_{act,t}^2 \end{cases}$$

$$\overline{bel}(x_t) = \begin{cases} \bar{\mu}_t = A_t \mu_{t-1} + B_t u_t \\ \bar{\Sigma}_t = A_t \Sigma_{t-1} A_t^T + R_t \end{cases}$$




# The Prediction-Correction-Cycle



$$bel(x_t) = \begin{cases} \mu_t = \bar{\mu}_t + K_t(z_t - \bar{\mu}_t) \\ \sigma_t^2 = (1 - K_t)\bar{\sigma}_t^2, K_t = \frac{\bar{\sigma}_t^2}{\bar{\sigma}_t^2 + \bar{\sigma}_{obs,t}^2} \end{cases}$$

$$bel(x_t) = \begin{cases} \mu_t = \bar{\mu}_t + K_t(z_t - C_t\bar{\mu}_t) \\ \Sigma_t = (I - K_tC_t)\bar{\Sigma}_t, K_t = \bar{\Sigma}_tC_t^T(C_t\bar{\Sigma}_tC_t^T + Q_t)^{-1} \end{cases}$$



Correction

# The Prediction-Correction-Cycle



Prediction

$$bel(x_t) = \begin{cases} \mu_t = \bar{\mu}_t + K_t(z_t - \bar{\mu}_t) \\ \sigma_t^2 = (1 - K_t)\bar{\sigma}_t^2 \end{cases}, K_t = \frac{\bar{\sigma}_t^2}{\bar{\sigma}_t^2 + \bar{\sigma}_{obs,t}^2}$$

$$\overline{bel}(x_t) = \begin{cases} \bar{\mu}_t = a_t \mu_{t-1} + b_t u_t \\ \bar{\sigma}_t^2 = a_t^2 \sigma_{t-1}^2 + \sigma_{act,t}^2 \end{cases}$$

$$bel(x_t) = \begin{cases} \mu_t = \bar{\mu}_t + K_t(z_t - C_t \bar{\mu}_t) \\ \Sigma_t = (I - K_t C_t) \bar{\Sigma}_t \end{cases}, K_t = \bar{\Sigma}_t C_t^T (C_t \bar{\Sigma}_t C_t^T + Q_t)^{-1}$$

$$\overline{bel}(x_t) = \begin{cases} \bar{\mu}_t = A_t \mu_{t-1} + B_t u_t \\ \bar{\Sigma}_t = A_t \Sigma_{t-1} A_t^T + R_t \end{cases}$$

Correction



# Kalman Filter Summary

- *Highly efficient*: Polynomial in measurement dimensionality  $k$  and state dimensionality  $n$ :

$$O(k^{2.376} + n^2)$$

- *Optimal for linear Gaussian systems!*



# Kalman Filter


- Problems:

What if model of motion is not linear?

What if it is not even parametric?

# Particle Filter

- Basic idea: model is represented by a population of particles ( $X_t$ )



# Particle Filter Algorithm

Initialization:

$$X_0 \leftarrow n \text{ particles } x_0^{[i]} \sim p(x_0)$$

particleFilters( $X_{t-1}$ ) {

  for  $i=1$  to  $n$

$$x_t^{[i]} \sim p(x_t | x_{t-1}^{[i]})$$

(prediction)

$$w_t^{[i]} = p(z_t | x_t^{[i]})$$

(importance weights)

  endfor

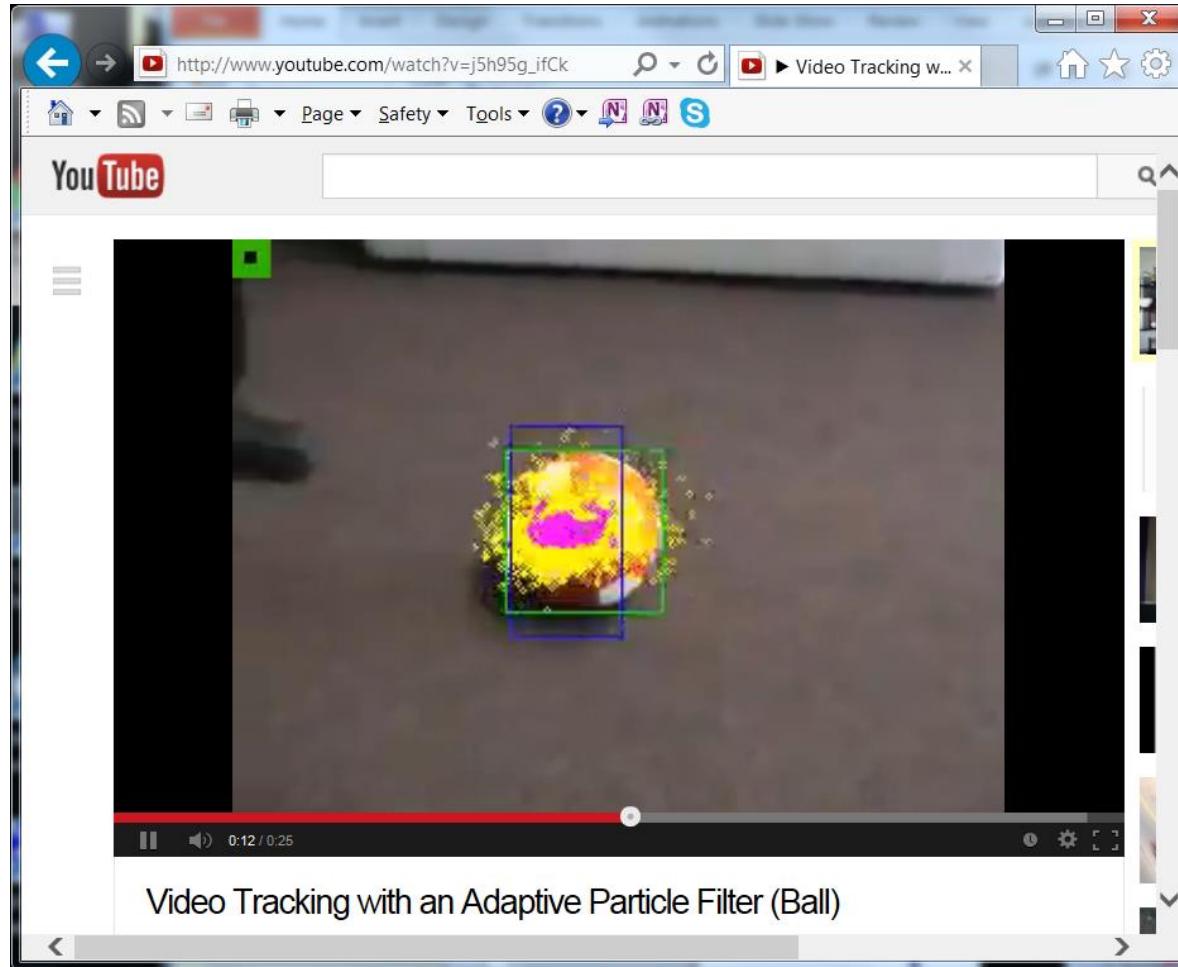
  for  $i=1$  to  $n$

    include  $x_t^{[i]}$  in  $X_t$  with probability  $\propto w_t^{[i]}$  (resampling)

  endfor

}

# Particle Filter Example



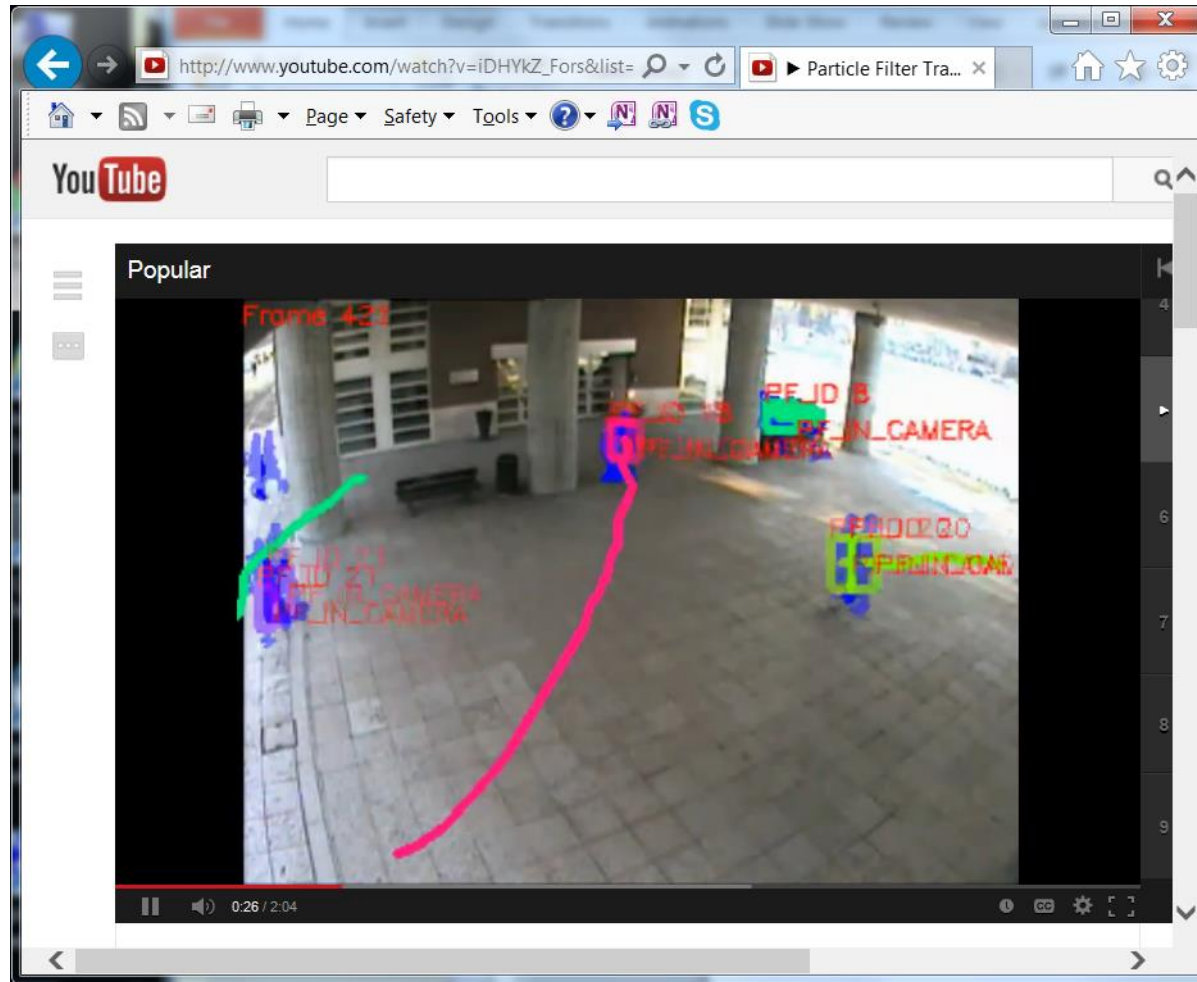
[http://www.youtube.com/watch?v=j5h95g\\_ifCk](http://www.youtube.com/watch?v=j5h95g_ifCk)

# Particle Filter Example



<http://www.youtube.com/watch?v=wCMk-pHzScE>

# Particle Filter Example



[http://www.youtube.com/watch?v=iDHYkZ\\_Fors&list=TLctNolqoeYKI\\_Mjy-26deFuPWKddjMI4P](http://www.youtube.com/watch?v=iDHYkZ_Fors&list=TLctNolqoeYKI_Mjy-26deFuPWKddjMI4P)

# Kalman Filter

- Estimates state of a system
  - Position
  - Velocity
  - Many other continuous state variables possible
- KF maintains
  - Mean vector for the state
  - Covariance matrix of state uncertainty
- Implements
  - Time update = prediction
  - Measurement update
- Standard Kalman filter is linear-Gaussian
  - Linear system dynamics, linear sensor model
  - Additive Gaussian noise (independent)
  - Nonlinear extensions: extended KF, unscented KF: linearize

# Particle Filter

- Estimates state of a system
    - Position
    - Velocity
    - Many other continuous state variables possible
  - KF maintains
    - ~~Mean vector for the state~~
    - ~~Covariance matrix of state uncertainty~~
  - Implements
    - Time update = prediction = predictive sampling
    - Measurement update = resampling, importance weights
  - ~~Standard Kalman filter is linear-Gaussian~~
    - ~~Linear system dynamics, linear sensor model~~
    - ~~Additive Gaussian noise (independent)~~
    - ~~Nonlinear extensions: extended KF, unscented KF: linearize~~
- and discrete
- set of particles  
(example states)
- fully nonlinear
- easy to implement



# Summary

- Feature tracking
- Object tracking
  - Foreground estimation
  - Model update
    - Kalman filter
    - Particle filter