



# Filtering and Robot Localization

Robert Platt  
Northeastern University



# Robot localization example


Robot is actually located here, but it doesn't know it.



- Goal: localize the robot based on sequential observations
- robot is given a map of the world; robot could be in any square
  - initially, robot doesn't know which square it's in

# Robot localization example


Robot perceives that there are walls above and below, but no walls either left or right



On each time step, the robot moves, and then observes the directions in which there are walls.

- observes a four-bit binary number
- observations are noisy: there is a small chance that each bit will be flipped.

# Robot localization example




Prob




0

1



# Robot localization example




Prob




0

1



# Robot localization example




Prob





0

1




# Robot localization example



Question: how do we update this probability distribution from time  $t$  to  $t+1$ ?

# Hidden Markov Models (HMMs)




Called an “emission”

State,  $X_t$ , is assumed to be unobserved

However, you get to make one observation,  $E_t$ , on each timestep.




# Hidden Markov Models (HMMs)



Process dynamics:  $P(X_t | X_{t-1})$  ← How the system changes from one time step to the next

Observation dynamics:  $P(E_t | X_t)$  ← What gets observed as a function of what state the system is in

# Hidden Markov Models (HMMs)



Process dynamics:

$$P(X_t | X_{t-1})$$

How the system changes from one time step to the next


Observation dynamics:

$$P(E_t | X_t)$$

What gets observed as a function of what state the system is in

Let's assume (for now) that these probability distributions are given to us.


# Hidden Markov Models (HMMs)



Process dynamics:  $P(X_t | X_{t-1}) = P(X_t | X_{t-1}, \dots, X_1)$

Observation dynamics:  $P(E_t | X_t) = P(E_t | X_t, X_{t-1}, \dots, X_1)$

# Hidden Markov Models (HMMs)




Process dynamics:  $P(X_t|X_{t-1}) = P(X_t|X_{t-1}, \dots, X_1)$

Observation dynamics:  $P(E_t|X_t) = P(E_t|X_t, X_{t-1}, \dots, X_1)$

Markov assumptions


# HMM example



$R_t$	$R_{t+1}$	$P(R_{t+1}   R_t)$
+r	+r	0.7
+r	-r	0.3
-r	+r	0.3
-r	-r	0.7



$R_t$	$U_t$	$P(U_t   R_t)$
+r	+u	0.9
+r	-u	0.1
-r	+u	0.2
-r	-u	0.8

# Bayes Filtering

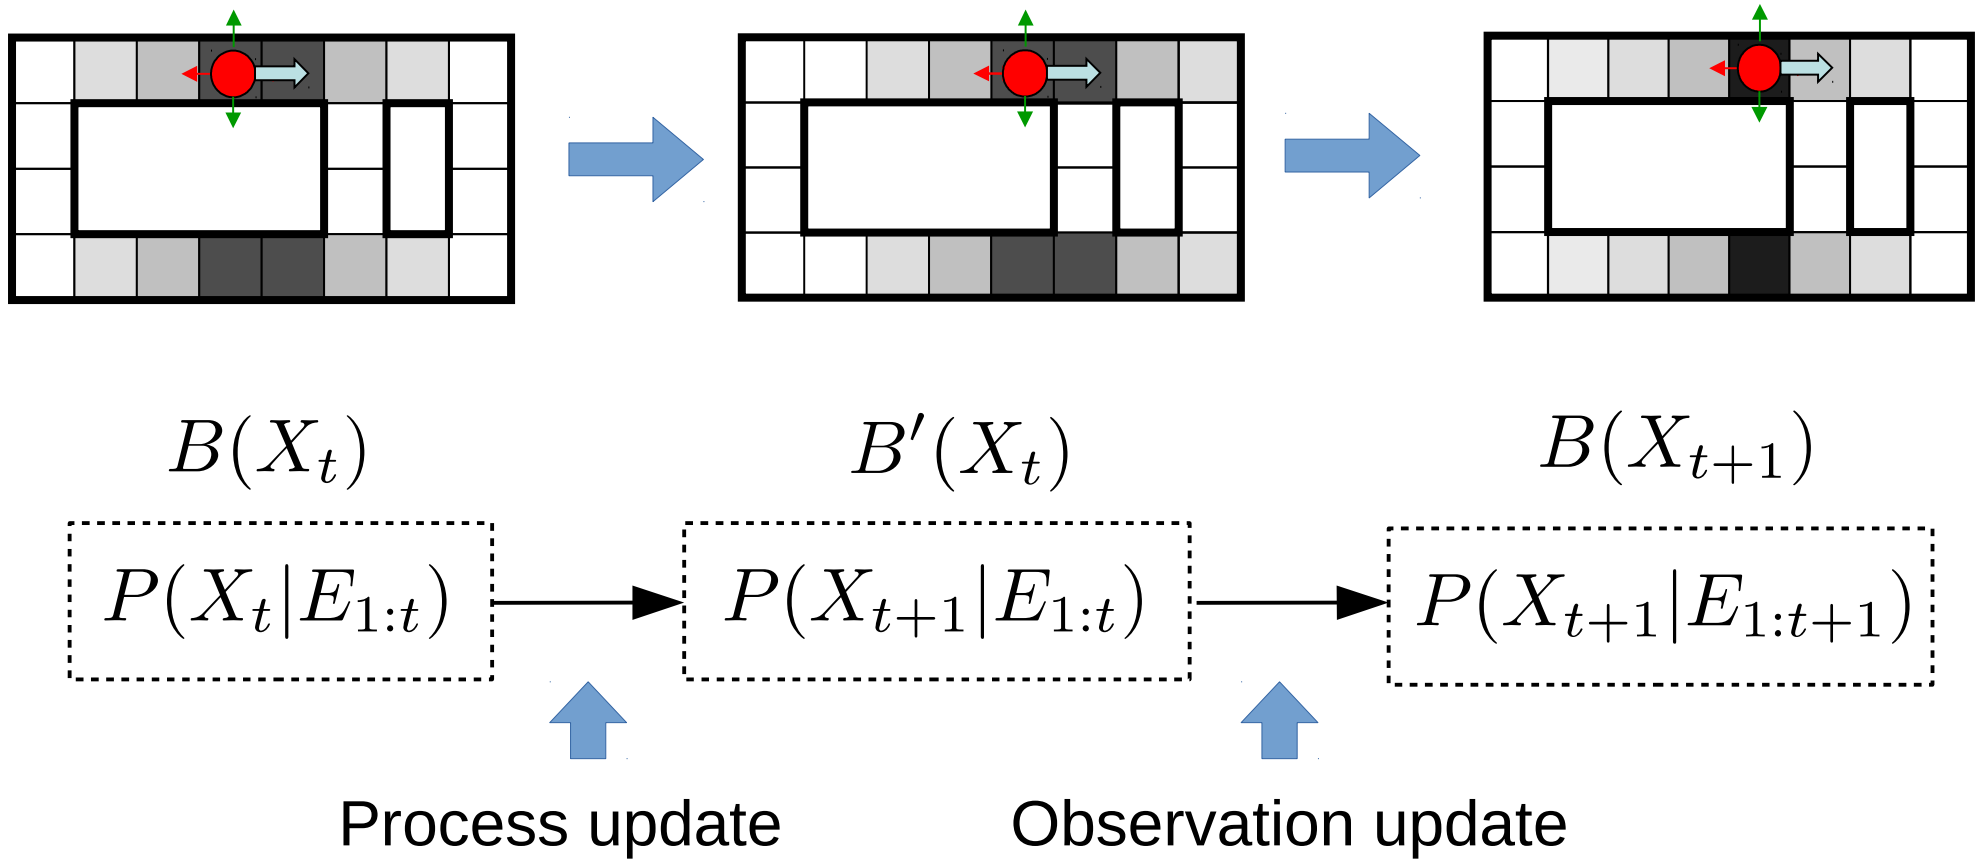


How do we go from this distribution to this distribution?

# Bayes Filtering




# Bayes Filtering






# Process update



$B(X_t)$



$B'(X_t)$



$P(X_t | e_{1:t})$

$P(X_{t+1} | e_{1:t})$

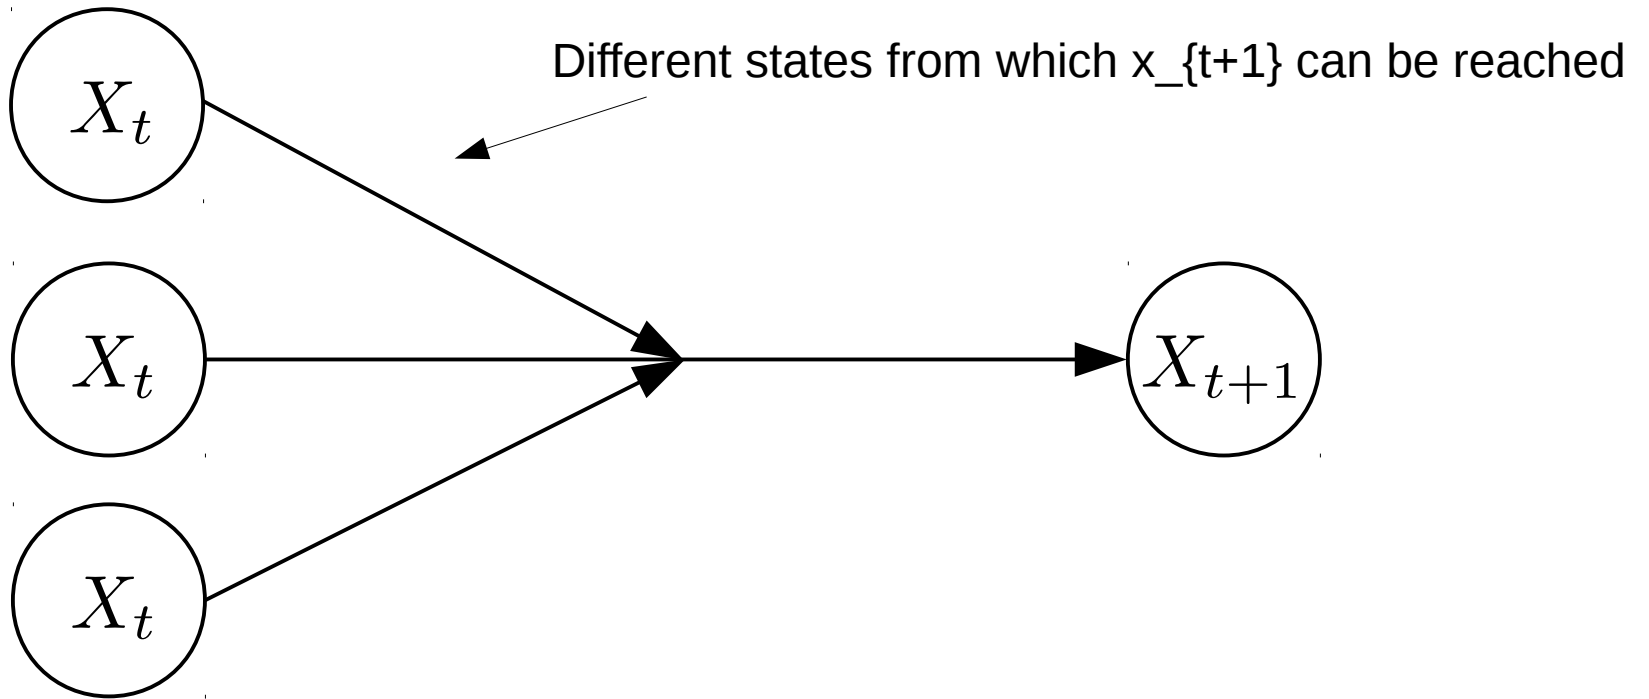

# Process update



$$P(X_{t+1} | e_{1:t}) = \sum_{X_t} P(X_{t+1} | X_t, e_{1:t}) P(X_t | e_{1:t})$$

Marginalize over next states

# Process update

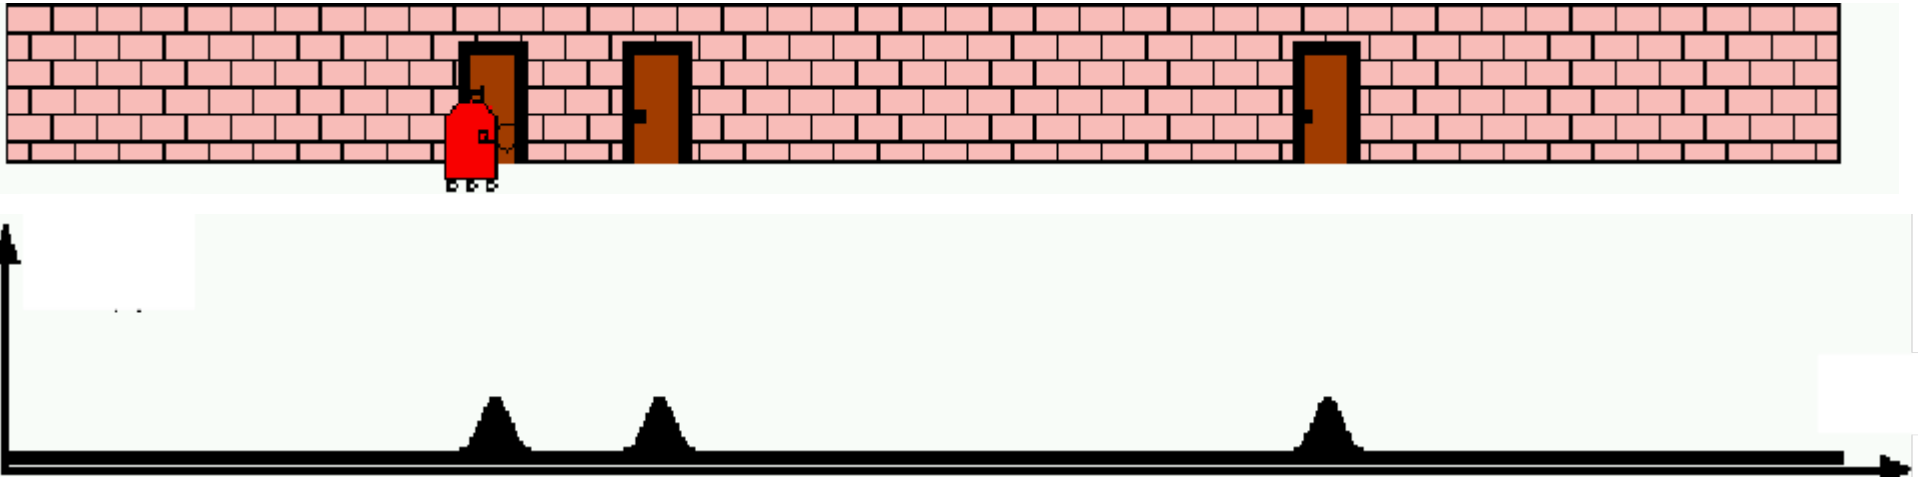


$$B'(X_{t+1}) = \sum_{X_t} P(X_{t+1}|X_t, e_{1:t})B(X_t)$$

Marginalize over next states


# Process update

Before process update



# Process update


After process update



$$B'(X_{t+1}) = \sum_{X_t} P(X_{t+1} | X_t, e_{1:t}) B(X_t) \leftarrow \text{This is a little like convolution...}$$

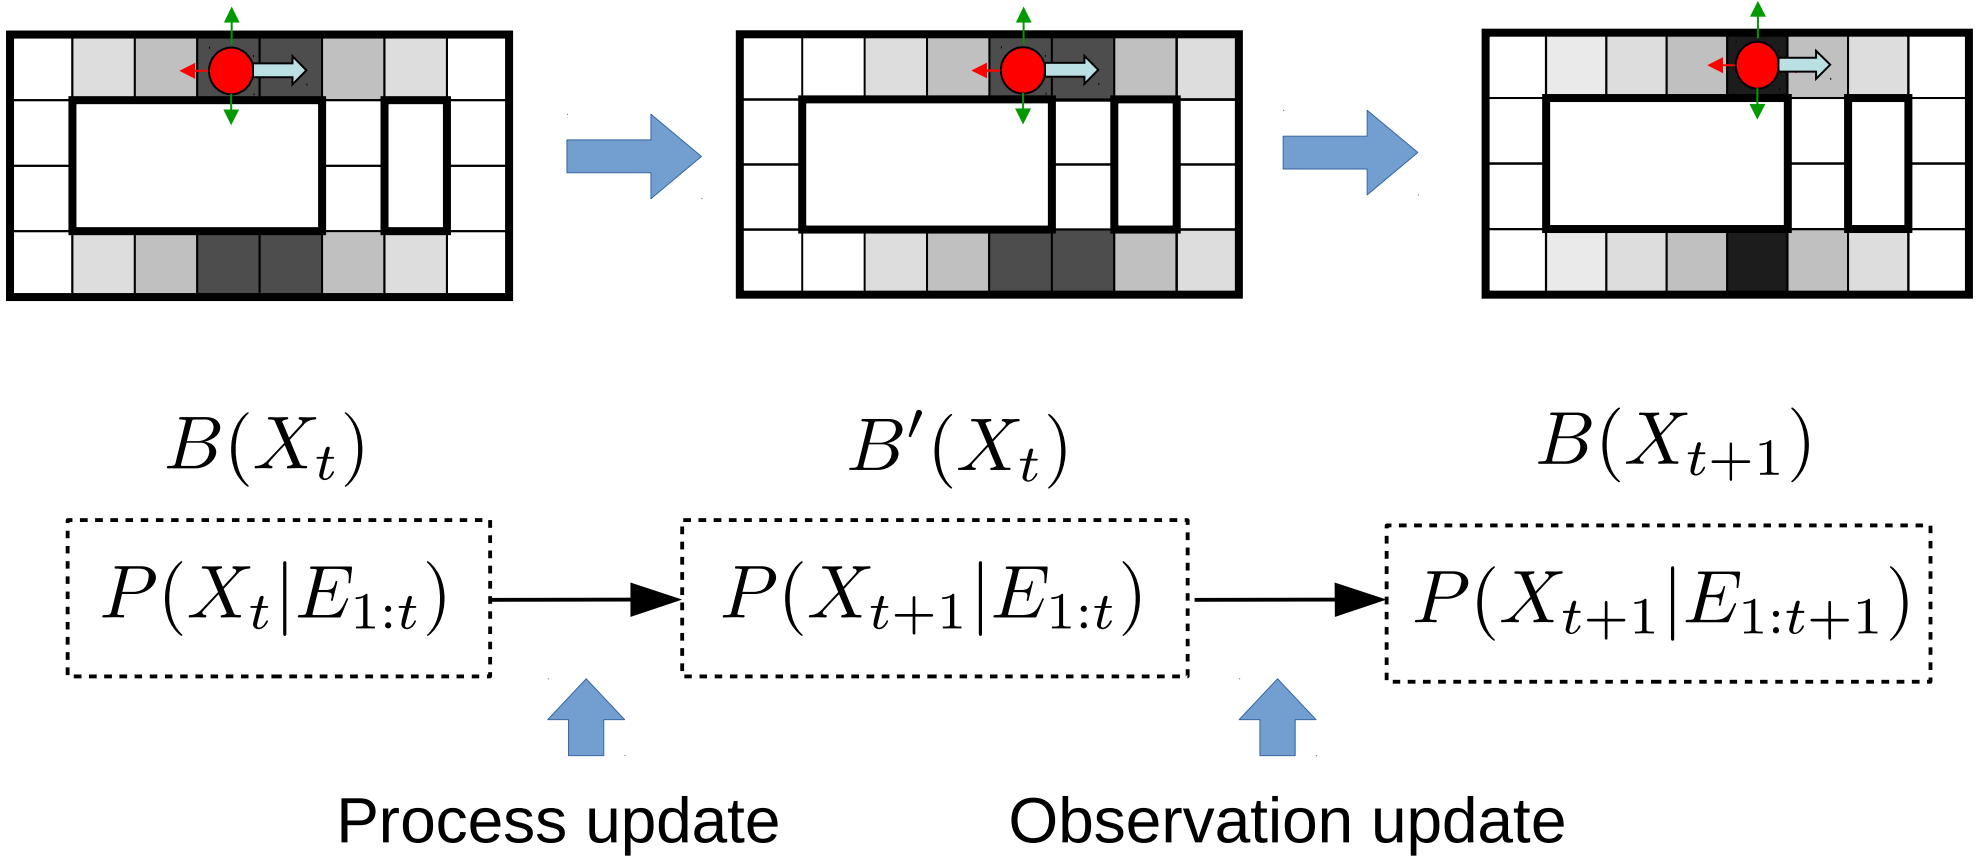
# Process update

After process update




- Each time you execute a process update, belief gets more disbursed
- *i.e.* Shannon entropy increases
  - this makes sense: as you predict state further into the future, your uncertainty grows.

# Bayes Filtering




# Observation update



$$P(X_{t+1}|e_{1:t+1}) = \eta P(e_{t+1}|X_{t+1})P(X_{t+1}|e_{1:t})$$




# Observation update



$$P(X_{t+1}|e_{1:t+1}) = \eta \underbrace{P(e_{t+1}|X_{t+1})}_{\text{Probability of seeing observation } e_{t+1} \text{ from state } X_{t+1}} P(X_{t+1}|e_{1:t})$$

Probability of seeing observation  $e_{t+1}$  from state  $X_{t+1}$

# Observation update




$$P(X_{t+1}|e_{1:t+1}) = \eta P(e_{t+1}|X_{t+1})P(X_{t+1}|e_{1:t})$$

$$B(X_{t+1}) = \eta P(e_{t+1}|X_{t+1})B'(X_{t+1})$$


Where  $\eta = \frac{1}{P(e_{t+1})}$  is a normalization factor

# Observation update

Before observation update



After observation update




# Weather HMM example

$R_t$	$R_{t+1}$	$P(R_{t+1} R_t)$
+r	+r	0.7
+r	-r	0.3
-r	+r	0.3
-r	-r	0.7

$$B(+r) = 0.5$$

$$B(-r) = 0.5$$



$R_t$	$U_t$	$P(U_t R_t)$
+r	+u	0.9
+r	-u	0.1
-r	+u	0.2
-r	-u	0.8

# Weather HMM example


$$B'( +r) = 0.5$$

$$B'( -r) = 0.5$$

$$B( +r) = 0.5$$

$$B( -r) = 0.5$$

$R_t$	$R_{t+1}$	$P(R_{t+1}   R_t)$
+r	+r	0.7
+r	-r	0.3
-r	+r	0.3
-r	-r	0.7



$R_t$	$U_t$	$P(U_t   R_t)$
+r	+u	0.9
+r	-u	0.1
-r	+u	0.2
-r	-u	0.8


# Weather HMM example

$R_t$	$R_{t+1}$	$P(R_{t+1} R_t)$
+r	+r	0.7
+r	-r	0.3
-r	+r	0.3
-r	-r	0.7

$B(+r) = 0.5$   
 $B(-r) = 0.5$

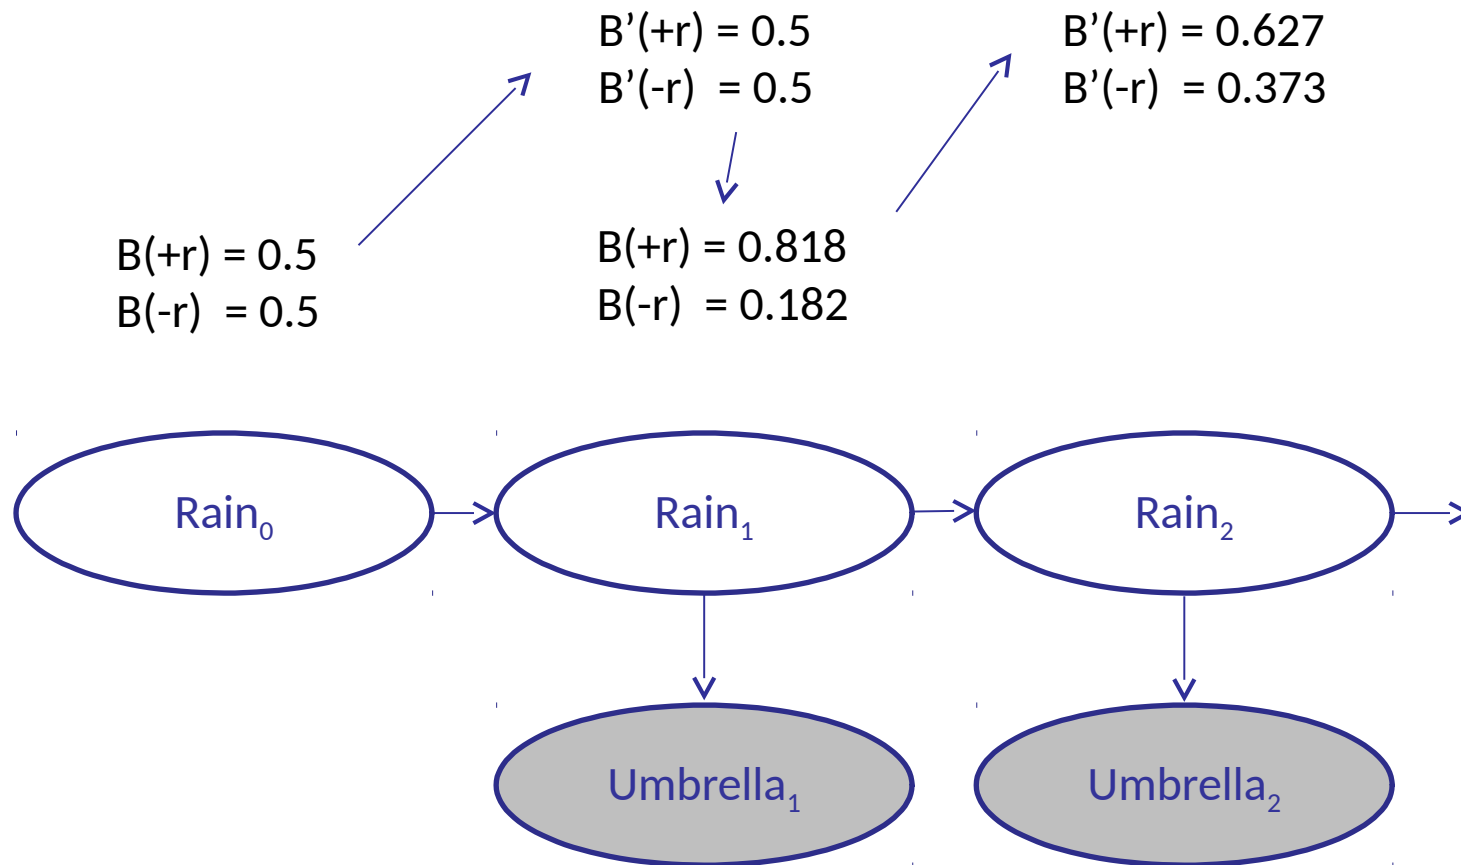
$B'(+r) = 0.5$   
 $B'(-r) = 0.5$

$B(+r) = 0.818$   
 $B(-r) = 0.182$



$R_t$	$U_t$	$P(U_t R_t)$
+r	+u	0.9
+r	-u	0.1
-r	+u	0.2
-r	-u	0.8

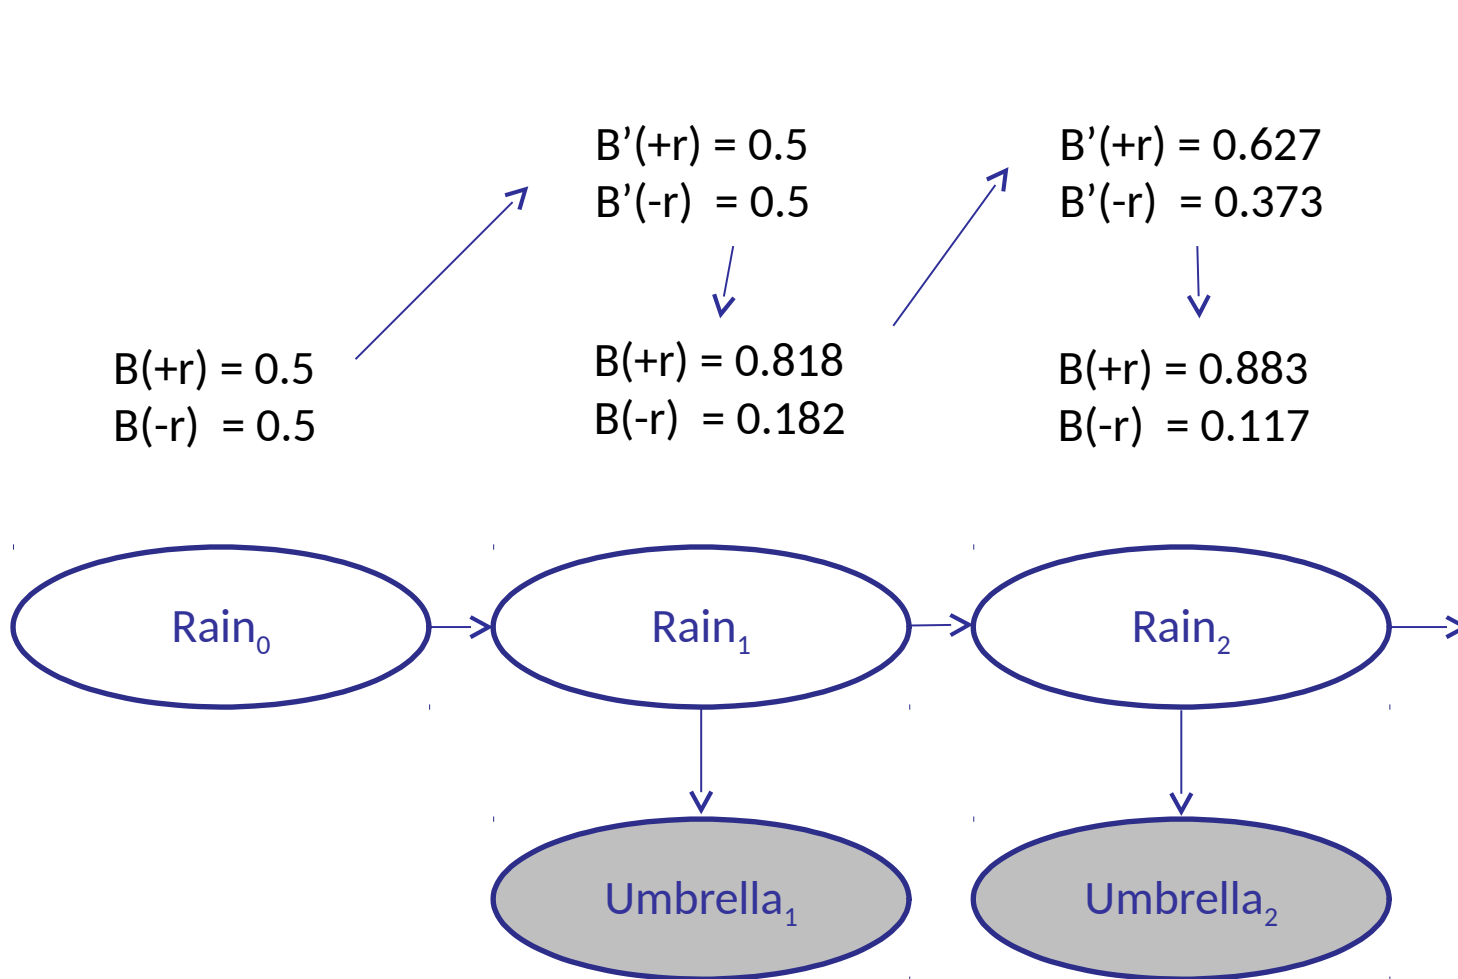
# Weather HMM example



$R_t$	$R_{t+1}$	$P(R_{t+1} R_t)$
+r	+r	0.7
+r	-r	0.3
-r	+r	0.3
-r	-r	0.7

$R_t$	$U_t$	$P(U_t R_t)$
+r	+u	0.9
+r	-u	0.1
-r	+u	0.2
-r	-u	0.8

# Weather HMM example




$R_t$	$R_{t+1}$	$P(R_{t+1} R_t)$
+r	+r	0.7
+r	-r	0.3
-r	+r	0.3
-r	-r	0.7


$R_t$	$U_t$	$P(U_t R_t)$
+r	+u	0.9
+r	-u	0.1
-r	+u	0.2
-r	-u	0.8




# Robot localization example




# Robot localization example




# Robot localization example



# Robot localization example



# Robot localization example




Prob




0


1




# Robot localization example



Prob



# Robot localization example




Prob




0

1



# Robot localization example




Prob




0

1





# Robot localization example




Prob




0

1



# Robot localization example




Prob



0

1





# Applications of HMMs

- **Speech recognition HMMs:**
  - Observations are acoustic signals (continuous valued)
  - States are specific positions in specific words (so, tens of thousands)
- **Machine translation HMMs:**
  - Observations are words (tens of thousands)
  - States are translation options
- **Robot tracking:**
  - Observations are range readings (continuous)
  - States are positions on a map (continuous)

# Particle Filter

Why must I be confined to this grid?




Standard Bayes filtering requires discretizing state space into grid cells

Can do Bayes filtering w/o discretizing?

– yes: particle filtering or Kalman filtering

# Particle Filter



Sequential Bayes Filtering is great, but it's not great for continuous state spaces.  
– you need to discretize the state space (e.g. a grid) in order to use Bayes filtering  
– but, doing filtering on a grid is not efficient...


Therefore:

- particle filters
- Kalman filters



Two different ways of filtering in continuous state spaces

# Particle Filter




Key idea: represent a probability distribution as a finite set of points

– density of points encodes probability mass.

– particle filtering is an adaptation of Bayes filtering to this particle representation

# Monte Carlo Sampling




Suppose you are given an unknown probability distribution,  $P(x)$

Suppose you can't evaluate the distribution analytically, but you can draw samples from it

What can you do with this information?

$$E_{x \sim P(x)}(h(x)) = \int_x h(x) P(x)$$
$$\approx \frac{1}{k} \sum_{i=1}^k h(x^i) \quad \text{where } x^i \text{ are samples drawn from } P(x)$$

# Importance Sampling



Suppose you are given an unknown probability distribution,  $P(x)$

Suppose you can't evaluate the distribution analytically, but you can draw samples from it

What can you do with this information?


Suppose you can't even sample from it?

Suppose that all you can do is evaluate the function at a given point?




# Importance Sampling

Question: how estimate expected values if cannot draw samples from  $f(x)$   
– suppose all we can do is evaluate  $f(x)$  at a given point...



# Importance Sampling


Question: how estimate expected values if cannot draw samples from  $f(x)$   
– suppose all we can do is evaluate  $f(x)$  at a given point...



Answer: draw samples from a different distribution and weight them

# Importance Sampling

Question: how estimate expected values if cannot draw samples from  $f(x)$   
– suppose all we can do is evaluate  $f(x)$  at a given point...



Answer: draw samples from a different distribution and weight them

$$E_{x \sim f(x)}(h(x)) = \int_x h(x) \frac{f(x)}{g(x)} g(x)$$
$$\approx \frac{1}{k} \sum_{i=1}^k h(x^i) w_i \quad \text{where } x^i \text{ are samples drawn from } g(x)$$

and  $w_i = f(x^i)/g(x^i)$

Proposal distribution

# Particle Filter

Prior distribution

$$x_t^1, \dots, x_t^n \quad w_t^1, \dots, w_t^n = 1$$

$$B(X_t)$$

$$P(X_t | E_{1:t})$$



$$B'(X_t)$$

$$P(X_{t+1} | E_{1:t})$$



$$B(X_{t+1})$$

$$P(X_{t+1} | E_{1:t+1})$$

# Particle Filter

Prior distribution

$$x_t^1, \dots, x_t^n \quad w_t^1, \dots, w_t^n = 1$$

$$B(X_t)$$

$$P(X_t | E_{1:t})$$



Process update

$$\bar{x}_{t+1}^i \sim P(X_{t+1} | x_t^i, e_{1:t})$$

$$B'(X_t)$$

$$P(X_{t+1} | E_{1:t})$$



$$B(X_{t+1})$$

$$P(X_{t+1} | E_{1:t+1})$$

# Particle Filter

Prior distribution

$$x_t^1, \dots, x_t^n \quad w_t^1, \dots, w_t^n = 1$$

Process update

$$\bar{x}_{t+1}^i \sim P(X_{t+1} | x_t^i, e_{1:t})$$

Observation update

$$w_{t+1}^i = P(e_{t+1} | \bar{x}_{t+1}^i) w_t^i$$

$$B(X_t)$$

$$P(X_t | E_{1:t})$$



$$B'(X_t)$$

$$P(X_{t+1} | E_{1:t})$$



$$B(X_{t+1})$$

$$P(X_{t+1} | E_{1:t+1})$$

# Particle Filter

$$B(X_t) \quad \boxed{P(X_t | E_{1:t})}$$



$$B'(X_t) \quad \boxed{P(X_{t+1} | E_{1:t})}$$



$$B(X_{t+1}) \quad \boxed{P(X_{t+1} | E_{1:t+1})}$$

Do this  $n$  times

Prior distribution

$$x_t^1, \dots, x_t^n \quad w_t^1, \dots, w_t^n = 1$$

Process update

$$\bar{x}_{t+1}^i \sim P(X_{t+1} | x_t^i, e_{1:t})$$

Observation update


$$w_{t+1}^i = P(e_{t+1} | \bar{x}_{t+1}^i) w_t^i$$

Resample

$$X_{t+1} = \{\}$$




$$\longrightarrow X_{t+1} = X_{t+1} \cup \bar{x}_{t+1}^i \text{ w/ prob } w_{t+1}^i$$

# Particle Filter








# Particle Filter





# Particle Filter




# Particle Filter




# Particle Filter



# Particle Filter Example



# Particle Filter Example



# Particle Filtering

## Pros:

- works in continuous spaces
- can represent multi-modal distributions

## Cons:

- parameters to tune
- sample impoverishment

# Sample Impoverishment


## Pros:

- works in continuous spaces
- can represent multi-modal distributions

## Cons:

- parameters to tune
- sample impoverishment

No particles nearby the true system state





# Sample Impoverishment

If there aren't enough samples, then we might "resample away" the true state...

Prior distribution

$$x_t^n \quad w_t^1, \dots, w_t^n = 1$$

Process update

$$P(X_{t+1} | x_t^i, e_{1:t})$$

Observation update

$$w_{t+1}^i = P(e_{t+1} | \bar{x}_{t+1}^i) w_t^i$$

$$B(X_{t+1}) \quad P(X_{t+1} | E_{1:t+1})$$

Do this  $n$  times

Resample

$$X_{t+1} = \{ \}$$

$$\rightarrow X_{t+1} = X_{t+1} \cup \bar{x}_{t+1}^i \text{ w/ prob } w_{t+1}^i$$

# Sample Impoverishment

If there aren't enough samples, then we might "resample away" the true state...

One solution: add an additional  $k$  samples drawn completely at random

Prior distribution

$$x_t^n \quad w_t^1, \dots, w_t^n = 1$$

Process update

$$P(X_{t+1} | x_t^i, e_{1:t})$$

Observation update

$$w_{t+1}^i = P(e_{t+1} | \bar{x}_{t+1}^i) w_t^i$$

$$B(X_{t+1}) \quad P(X_{t+1} | E_{1:t+1})$$

Do this  $n$  times

Resample

$$X_{t+1} = \{ \}$$

$$\rightarrow X_{t+1} = X_{t+1} \cup \bar{x}_{t+1}^i \text{ w/ prob } w_{t+1}^i$$

# Sample Impoverishment

If there aren't enough samples, then we might "resample away" the true state...

One solution: add an additional  $k$  samples drawn completely at random

BUT: there's always a chance that the true state won't be represented well by the particles...

Prior distribution

$$x_t^n \quad w_t^1, \dots, w_t^n = 1$$

Process update

$$P(X_{t+1} | x_t^i, e_{1:t})$$

Observation update

$$w_{t+1}^i = P(e_{t+1} | \bar{x}_{t+1}^i) w_t^i$$

$$B(X_{t+1}) \quad P(X_{t+1} | E_{1:t+1})$$


Do this  $n$  times

Resample

$$X_{t+1} = \{\}$$

$$\rightarrow X_{t+1} = X_{t+1} \cup \bar{x}_{t+1}^i \text{ w/ prob } w_{t+1}^i$$

# Kalman Filtering



Another way to adapt Sequential Bayes Filtering to continuous state spaces

- relies on representing the probability distribution as a Gaussian
- first developed in the early 1960s (before general Bayes filtering); used in Apollo program



# Kalman Idea




Image: Thrun *et al.*, CS233B course notes

# Kalman Idea




Image: Thrun et al., CS233B course notes

$$P(x_{t+1} | z_{0:t}) = \sum_{x_t} P(x_{t+1} | x_t) P(x_t | z_{0:t})$$

$$P(x_{t+1} | z_{0:t+1}) = \eta P(z_{t+1} | x_{t+1}) P(x_{t+1} | z_{0:t})$$




Image: Thrun et al., CS233B course notes

# Gaussians

- Univariate  
Gaussian:

$$P(x) = \eta e^{-\frac{1}{2} \frac{(x-\mu)^2}{\sigma^2}}$$

- Multivariate  
Gaussian:

$$P(x) = \eta e^{-\frac{1}{2} (x-\mu)^T \Sigma^{-1} (x-\mu)}$$


$$P(x) = N(x; \mu, \Sigma)$$

# Playing w/ Gaussians

- Suppose:  $P(x) = N(x; \mu, \Sigma)$   
 $y = Ax + b$

- Calculate:  $P(y) = ?$

$$P(y) = N(y; A\mu + b, A\Sigma A^T)$$





# In fact

- Suppose:  $P(x) = N(x; \mu, \Sigma)$   
 $y = Ax + b$

- Then:

$$P \begin{pmatrix} x \\ y \end{pmatrix} = N \left[ \begin{array}{c} x \\ y \end{array} ; \begin{array}{c} \mu \\ A\mu + b \end{array}, \begin{pmatrix} \Sigma & \Sigma A^T \\ A\Sigma & A\Sigma A^T \end{pmatrix} \right]$$

# Illustration




Image: Thrun *et al.*, CS233B course notes

# And

Suppose:  $P(x) = N(x; \mu, \Sigma)$

$$P(y|x) = N(y; Ax + b, R)$$

Then:

$$P \begin{pmatrix} x \\ y \end{pmatrix} = N \left[ \begin{matrix} x \\ y \end{matrix} ; \begin{matrix} \mu \\ A\mu + b \end{matrix}, \begin{pmatrix} \Sigma & \Sigma A^T \\ A\Sigma & A\Sigma A^T + R \end{pmatrix} \right]$$

$$P(y) = N(y; A\mu + b, A\Sigma A^T + R)$$

Marginal distribution



Does this remind us of anything?

# Does this remind us of anything?

Process update  
(discrete):

Process update  
(continuous):  $P(x_{t+1}|z_{0:t}) = \int_{x_t} P(x_{t+1}|x_t)P(x_t|z_{0:t})$

# Does this remind us of anything?

Process update

(discrete): 
$$P(x_{t+1}|z_{0:t}) = \sum_{x_t} P(x_{t+1}|x_t)P(x_t|z_{0:t})$$

Process update

(continuous):

$$N(x_{t+1}|Ax_t, Q)$$

transition dynamics

$$N(x_t|\mu_t, \Sigma_t)$$

prior

# Does this remind us of anything?

Process update

(discrete): 
$$P(x_{t+1}|z_{0:t}) = \sum_{x_t} P(x_{t+1}|x_t)P(x_t|z_{0:t})$$

Process update

(continuous):

$$N(x_{t+1}|Ax_t, Q)$$

transition dynamics

$$N(x_t|\mu_t, \Sigma_t)$$

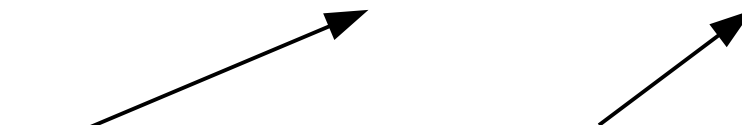
prior

$$P(x_{t+1}|z_{0:t}) = \int_{x_t} N(x_{t+1}|Ax_t, Q)N(x_t; \mu_t, \Sigma_t)$$

$$P(x_{t+1}|z_{0:t}) = N(x_{t+1}|A\mu_t, A\Sigma_t A^T + Q)$$

# Observation update

Observation  
update:


$$N(z_{t+1} | Cx_{t+1}, R)$$

$$N(x_t | \mu'_t, \Sigma'_t)$$

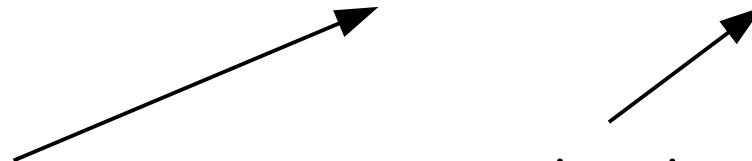
Where:  $\mu'_t = A\mu_t$

$$\Sigma'_t = A\Sigma_t A^T + Q$$



# Observation update

Observation  
update:


$$N(x_t | \mu'_t, \Sigma'_t)$$

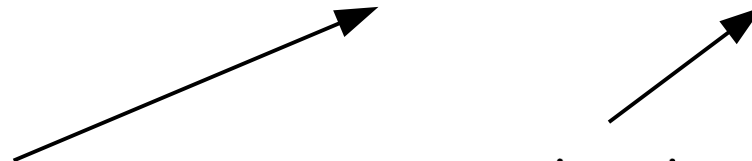
Where:

$$\Sigma'_t = A\Sigma_t A^T + Q$$

$$P(z_{t+1}, x_{t+1} | z_{0:t}) = \eta N(z_{t+1} | Cx_t, R) N(x_t; \mu'_t, \Sigma'_t)$$

# Observation update

Observation  
update:


$$N(x_t | \mu'_t, \Sigma'_t)$$

Where:

$$\Sigma'_t = A\Sigma_t A^T + Q$$

$$P(z_{t+1}, x_{t+1} | z_{0:t}) = N \left[ \begin{array}{c} x_{t+1} \\ z_{t+1} \end{array} \mid \begin{array}{c} \mu'_t \\ C\mu'_t \end{array}, \left( \begin{array}{cc} \Sigma'_t & \Sigma'_t C^T \\ C\Sigma'_t & C\Sigma'_t A^T + R \end{array} \right) \right]$$

# Observation update

But we need:  $P(x_{t+1} | z_{0:t+t}) = ?$

## Another Gaussian identity...

Suppose:  $N \left[ \begin{array}{c} x \\ y \end{array} \mid \begin{array}{c} a \\ b \end{array}, \left( \begin{array}{cc} A & C \\ C^T & B \end{array} \right) \right]$

Calculate:  $P(y|x) = ?$

$$P(y|x) = N(y | b + C^T A^{-1}(x - a), B - C^T A^{-1}C)$$

# Observation update

$$P(z_{t+1}, x_{t+1} | z_{0:t}) = N \left[ \begin{array}{c} x_{t+1} \\ z_{t+1} \end{array} ; \begin{array}{c} \mu'_t \\ C\mu'_t \end{array}, \left( \begin{array}{cc} \Sigma & \Sigma C^T \\ C\Sigma & C\Sigma A^T + R \end{array} \right) \right]$$

But we need:  $P(x_{t+1} | z_{0:t+1}) = ?$

$$P(x_{t+1} | z_{0:t+1}) = N(x_{t+1}; \mu_{t+1}, \Sigma_{t+1})$$

$$\mu_{t+1} = \mu'_t + \Sigma'_t C^T (R + C\Sigma'_t C^T)^{-1} (z_{t+1} - C\mu'_t)$$

$$\Sigma_{t+1} = \Sigma'_t - \Sigma'_t C^T (R + C\Sigma'_t C^T)^{-1} C\Sigma'_t$$

# To summarize the Kalman filter

System:  $P(x_{t+1}|x_t) = N(x_{t+1}|Ax_t, Q)$   
 $P(z_{t+1}|x_{t+1}) = N(z_{t+1}|Cx_{t+1}, R)$

Prior:  $\mu_t$   
 $\Sigma_t$

Process update:  $\mu'_t = A\mu_t$

Measurement  
update:

# Suppose there is an action term...

System:  $P(x_{t+1}|x_t) = N(x_{t+1}|Ax_t + u_t, Q)$

Prior:

$$\Sigma_t$$

Process update:  $\mu'_t = A\mu_t + u_t$

$$\Sigma'_t = A\Sigma_t A^T + Q$$

Measurement  
update:

# To summarize the Kalman filter

Prior:

$$\Sigma_t$$

Process update:  $\mu'_t = A\mu_t$

$$\Sigma'_t = A\Sigma_t A^T + Q$$

Measurement  
update:



This factor is often  
called the “Kalman  
gain”





# Things to note about the Kalman filter

Process update:  $\mu'_t = A\mu_t$


Measurement  
update:

- covariance update is independent of observation
- Kalman is only optimal for linear-Gaussian systems
- the distribution “stays” Gaussian through this update
- the error term can be thought of as the different between the observation and the prediction

# Kalman in 1D

System:

$$P(x_{t+1}|x_t) = N(x_{t+1} : x_t + u_t, q)$$
$$P(z_{t+1}|x_{t+1}) = N(z_{t+1} | 2x_{t+1}, r)$$




Process update:


$$\bar{\mu}_t = \mu_t + u_t$$
$$\bar{\sigma}_t^2 = \sigma_t^2 + q$$

Measurement update:

$$\mu_{t+1} = \bar{\mu}_t + \frac{2\bar{\sigma}_t^2}{r + 4\bar{\sigma}_t^2} (z_{t+1} - \bar{\mu}_t)$$
$$\sigma_{t+1}^2 = \bar{\sigma}_t^2 - \frac{4(\bar{\sigma}_t^2)^2}{r + 4\bar{\sigma}_t^2}$$



# Kalman Idea



# Example: estimate velocity




Image: Thrun *et al.*, CS233B course notes

# Example: filling a tank

$$x = \begin{pmatrix} l \\ f \end{pmatrix} \begin{array}{l} \leftarrow \text{Level of} \\ \text{tank} \\ \leftarrow \text{Fill rate} \end{array}$$

$$l_{t+1} = l_t + f dt$$

Process: 
$$x_{t+1} = \begin{pmatrix} 1 & dt \\ 0 & 1 \end{pmatrix} x_t + q$$

Observation: 
$$z_{t+1} = \begin{pmatrix} 1 & 0 \end{pmatrix} x_{t+1} + r$$

# Example: estimate velocity

$$x_{t+1} = Ax_t + w_t$$

$$\begin{pmatrix} x_{t+1} \\ y_{t+1} \\ \dot{x}_{t+1} \\ \dot{y}_{t+1} \end{pmatrix} = \begin{pmatrix} 1 & 0 & dt & 0 \\ 0 & 1 & 0 & dt \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 \end{pmatrix} \begin{pmatrix} x_t \\ y_t \\ \dot{x}_t \\ \dot{y}_t \end{pmatrix} + w_t$$

$$z_{t+1} = Cx_{t+1} + r_{t+1}$$

$$\begin{pmatrix} x_{t+1} \\ y_{t+1} \end{pmatrix} = \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \end{pmatrix} \begin{pmatrix} x_{t+1} \\ y_{t+1} \\ \dot{x}_{t+1} \\ \dot{y}_{t+1} \end{pmatrix} + r_{t+1}$$

# But, my system is NON-LINEAR!

$$\begin{aligned}x_{t+1} &= f(x_t, u_t) \\ &\neq Ax_t + Bu_t\end{aligned}$$

What should I do?

# But, my system is NON-LINEAR!

$$x_{t+1} = f(x_t, u_t)$$

- What should I do?

Well, there are some options...

-



# But, my system is NON-LINEAR!

$$x_{t+1} = f(x_t, u_t)$$

- What should I do?

Well, there are some options...

- But none of them are great.
-

# But, my system is NON-LINEAR!

$$x_{t+1} = f(x_t, u_t)$$

- What should I do?

Well, there are some options...

But none of them are great.

Here's one: the Extended Kalman Filter

# Extended Kalman filter

Take a Taylor expansion:

$$\begin{aligned}x_{t+1} &= f(x_t, u_t) \\ &\approx f(\mu_t, u_t) + A_t(x_t - \mu_t)\end{aligned}$$

$$\text{Where: } A_t = \frac{\partial f}{\partial x}(\mu_t, u_t)$$

$$\begin{aligned}z_{t+1} &= h(x_t) \\ &\approx h(\mu_t) + C_t(x_t - \mu_t)\end{aligned}$$

$$\text{Where: } C_t = \frac{\partial h}{\partial x}(\mu_t)$$

# Extended Kalman filter

Take a Taylor expansion:

$$x_{t+1} = f(x_t, u_t)$$

$$\text{Where: } A_t = \frac{\partial f}{\partial x}(\mu_t, u_t)$$

$$z_{t+1} = h(x_t)$$

$$\text{Where: } C_t = \frac{\partial h}{\partial x}(\mu_t)$$

Then use the same equations...

# To summarize the EKF

Prior:

$$\Sigma_t$$

Process update:

$$\mu'_t = f(\mu_t, u_t)$$

$$\Sigma'_t = A_t \Sigma_t A_t^T + Q$$

Measurement  
update:

$$\mu_{t+1} = \mu'_t + \Sigma'_t C^T (R + C \Sigma'_t C^T)^{-1} (z_{t+1} - h(\mu'_t))$$

# Extended Kalman filter

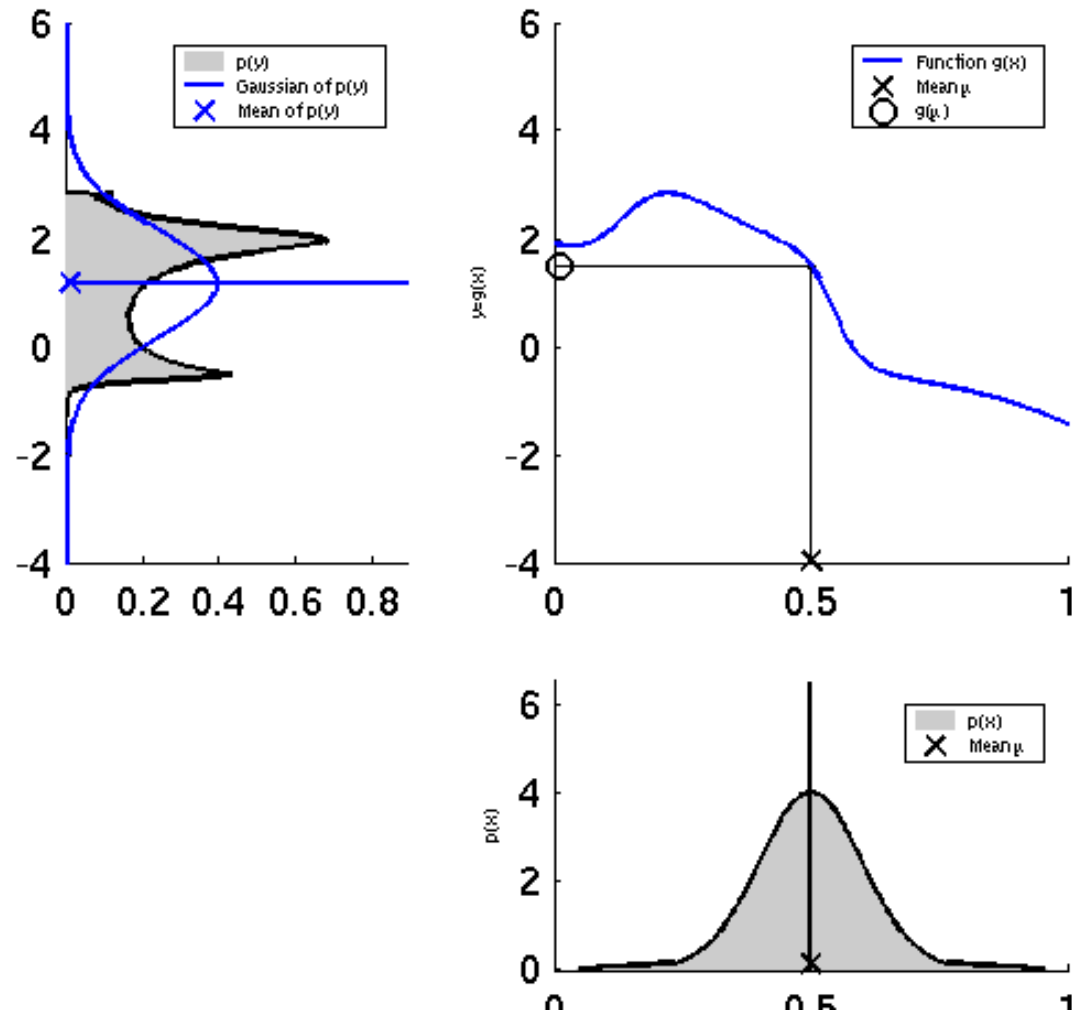


Image: Thrun *et al.*, CS233B course notes

# Extended Kalman filter

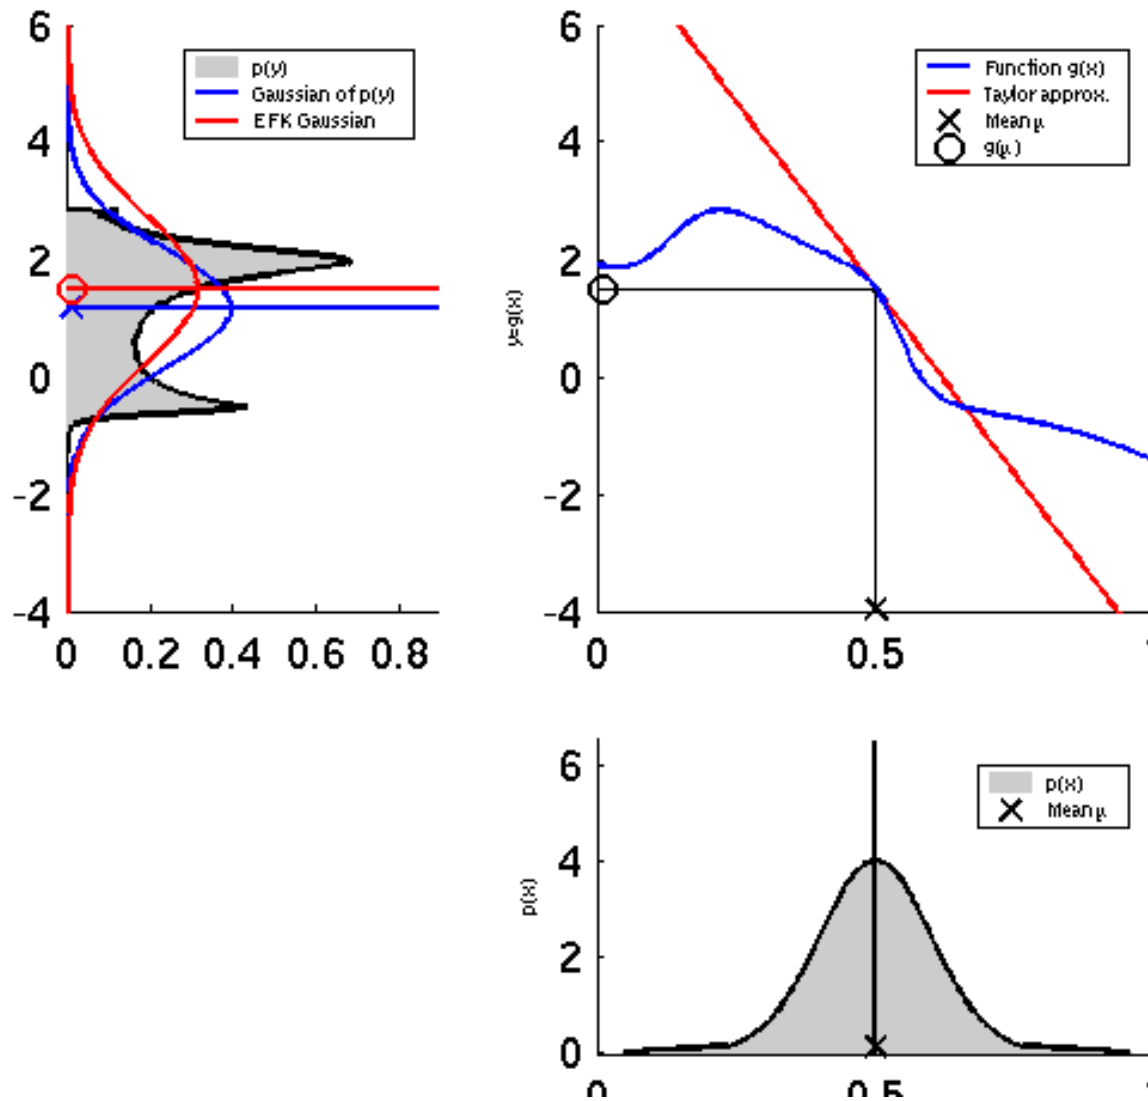



Image: Thrun et al., CS233B course notes

# EKF Mobile Robot Localization



Suppose we have a mobile robot wandering around in a 2-d world ...

Process dynamics:  $\mathbf{x}\langle k+1 \rangle = \mathbf{f}(\mathbf{x}\langle k \rangle, \delta\langle k \rangle, \mathbf{v}\langle k \rangle)$

state

Odometry  
measurement

noise

Process noise is assumed to be Gaussian:  $\mathbf{v} = (v_d, v_\theta) \sim N(0, \mathbf{V})$



# EKF Mobile Robot Localization

Process dynamics:


The diagram illustrates the process dynamics equation for EKF Mobile Robot Localization. The equation is presented within a light brown rectangular box. Above the box, the word "noise" is centered, with two arrows pointing downwards to the noise terms  $v_d$  and  $v_\theta$  in the equation. Below the box, the text "Odometry measurement" is centered, with two arrows pointing upwards to the distance  $\delta_d$  and angle  $\delta_\theta$  terms in the equation.

$$\xi\langle k+1\rangle = \begin{pmatrix} x\langle k\rangle + (\delta_d\langle k\rangle + v_d)\cos(\theta\langle k\rangle + \delta_\theta + v_\theta) \\ y\langle k\rangle + (\delta_d\langle k\rangle + v_d)\sin(\theta\langle k\rangle + \delta_\theta + v_\theta) \\ \theta\langle k\rangle + \delta_\theta + v_\theta \end{pmatrix}$$

noise

Odometry measurement


# EKF Mobile Robot Localization



But, wheels slip – odometry is not always correct...

How do we localize? Extended Kalman Filter!

# EKF Process Update



Dynamics:

Linearized dynamics:

$$\hat{\mathbf{x}}\langle k+1\rangle = \hat{\mathbf{x}}\langle k\rangle + \mathbf{F}_x(\mathbf{x}\langle k\rangle - \hat{\mathbf{x}}\langle k\rangle) + \mathbf{F}_v\mathbf{v}\langle k\rangle$$

Where:

$$\mathbf{F}_x = \left. \frac{\partial \mathbf{f}}{\partial \mathbf{x}} \right|_{\mathbf{v}=0} = \begin{pmatrix} 1 & 0 & -\delta_d\langle k\rangle - \sin(\theta\langle k\rangle + \delta_\theta) \\ 0 & 1 & \delta_d\langle k\rangle \cos(\theta\langle k\rangle + \delta_\theta) \\ 0 & 0 & 1 \end{pmatrix}$$
$$\mathbf{F}_v = \left. \frac{\partial \mathbf{f}}{\partial \mathbf{v}} \right|_{\mathbf{v}=0} = \begin{pmatrix} \cos(\theta\langle k\rangle + \delta_\theta) & -\delta_d\langle k\rangle \sin(\theta\langle k\rangle + \delta_\theta) \\ \sin(\theta\langle k\rangle + \delta_\theta) & \delta_d\langle k\rangle \cos(\theta\langle k\rangle + \delta_\theta) \\ 0 & 1 \end{pmatrix}$$

# EKF Process Update

EKF uncertainty estimate



Dynamics:

Linearized dynamics:

$$\hat{\mathbf{x}}\langle k+1\rangle = \hat{\mathbf{x}}\langle k\rangle + \mathbf{F}_x(\mathbf{x}\langle k\rangle - \hat{\mathbf{x}}\langle k\rangle) + \mathbf{F}_v\mathbf{v}\langle k\rangle$$

Where:

$$\mathbf{F}_v = \left. \frac{\partial \mathbf{f}}{\partial \mathbf{v}} \right|_{\mathbf{v}=0} = \begin{pmatrix} \cos(\theta\langle k\rangle + \delta_\theta) & -\delta_d\langle k\rangle \sin(\theta\langle k\rangle + \delta_\theta) \\ \sin(\theta\langle k\rangle + \delta_\theta) & \delta_d\langle k\rangle \cos(\theta\langle k\rangle + \delta_\theta) \\ 0 & 1 \end{pmatrix}$$


Process update:

$$\hat{\mathbf{x}}\langle k+1|k\rangle = \mathbf{f}(\hat{\mathbf{x}}\langle k\rangle, \delta\langle k\rangle, \mathbf{0})$$

$$\hat{\mathbf{P}}\langle k+1|k\rangle = \mathbf{F}_x\langle k\rangle \hat{\mathbf{P}}\langle k|k\rangle \mathbf{F}_x\langle k\rangle^T + \mathbf{F}_v\langle k\rangle \hat{\mathbf{V}}\mathbf{F}_v\langle k\rangle^T$$


With no observations,  
uncertainty grows  
over time...

# Observations



## Observations:

– range and bearing of a landmark



Observations:  $z = h(x_v, x_f, w)$

$$z = \begin{pmatrix} \sqrt{(y_i - y_v)^2 + (x_i - x_v)^2} \\ \tan^{-1}(y_i - y_v)/(x_i - x_v) - \theta_v \end{pmatrix} + \begin{pmatrix} w_r \\ w_\beta \end{pmatrix}$$

← range  
← bearing


$$\begin{pmatrix} w_r \\ w_\beta \end{pmatrix} \sim N(0, W), \quad W = \begin{pmatrix} \sigma_r^2 & 0 \\ 0 & \sigma_\beta^2 \end{pmatrix}$$

# Observations

Observations:

$$\mathbf{z} = \mathbf{h}(\mathbf{x}_v, \mathbf{x}_f, \mathbf{w})$$

Landmarks  
(i.e. features of the env't)



$$\mathbf{z}\langle k \rangle = \hat{\mathbf{h}} + \mathbf{H}_x (\mathbf{x}\langle k \rangle - \hat{\mathbf{x}}\langle k \rangle) + \mathbf{H}_w \mathbf{w}\langle k \rangle$$

where:

$$\mathbf{H}_{x_i} = \left. \frac{\partial \mathbf{h}}{\partial \mathbf{x}_v} \right|_{\mathbf{w}=0} = \begin{pmatrix} -\frac{x_i - x_v\langle k \rangle}{r} & -\frac{y_i - y_v\langle k \rangle}{r} & 0 \\ \frac{x_i - x_v\langle k \rangle}{r^2} & -\frac{y_i - y_v\langle k \rangle}{r^2} & -1 \end{pmatrix}$$

$$\mathbf{H}_w = \left. \frac{\partial \mathbf{h}}{\partial \mathbf{w}} \right|_{\mathbf{w}=0} = \begin{pmatrix} 1 & 0 \\ 0 & 1 \end{pmatrix}$$

$$\begin{pmatrix} w_r \\ w_\beta \end{pmatrix} \sim N(0, \mathbf{W}) \quad \mathbf{W} = \begin{pmatrix} \sigma_r^2 & 0 \\ 0 & \sigma_\beta^2 \end{pmatrix}$$

# EKF Mobile Robot Localization



Process Update:

Observation Update:

$$\hat{\mathbf{x}}\langle k+1|k+1\rangle = \hat{\mathbf{x}}\langle k+1|k\rangle + \mathbf{K}\langle k+1\rangle \nu\langle k+1\rangle$$


$$\hat{\mathbf{P}}\langle k+1|k+1\rangle = \hat{\mathbf{P}}\langle k+1|k\rangle \mathbf{F}_x\langle k\rangle^T - \mathbf{K}\langle k+1\rangle \mathbf{H}_x\langle k+1\rangle \hat{\mathbf{P}}\langle k+1|k\rangle$$

$$\nu\langle k+1\rangle = \mathbf{z}\langle k+1\rangle - \mathbf{h}(\hat{\mathbf{x}}\langle k+1|k\rangle, \mathbf{x}_f, 0)$$

$$\mathbf{S}\langle k+1\rangle = \mathbf{H}_x\langle k+1\rangle \hat{\mathbf{P}}\langle k+1|k\rangle \mathbf{H}_x\langle k+1\rangle^T + \mathbf{H}_w\langle k+1\rangle \hat{\mathbf{W}}\langle k+1\rangle \mathbf{H}_w\langle k+1\rangle^T$$

$$\mathbf{K}\langle k+1\rangle = \hat{\mathbf{P}}\langle k+1|k\rangle \mathbf{H}_x\langle k+1\rangle^T \mathbf{S}\langle k+1\rangle^{-1}$$

# Mapping using the EKF



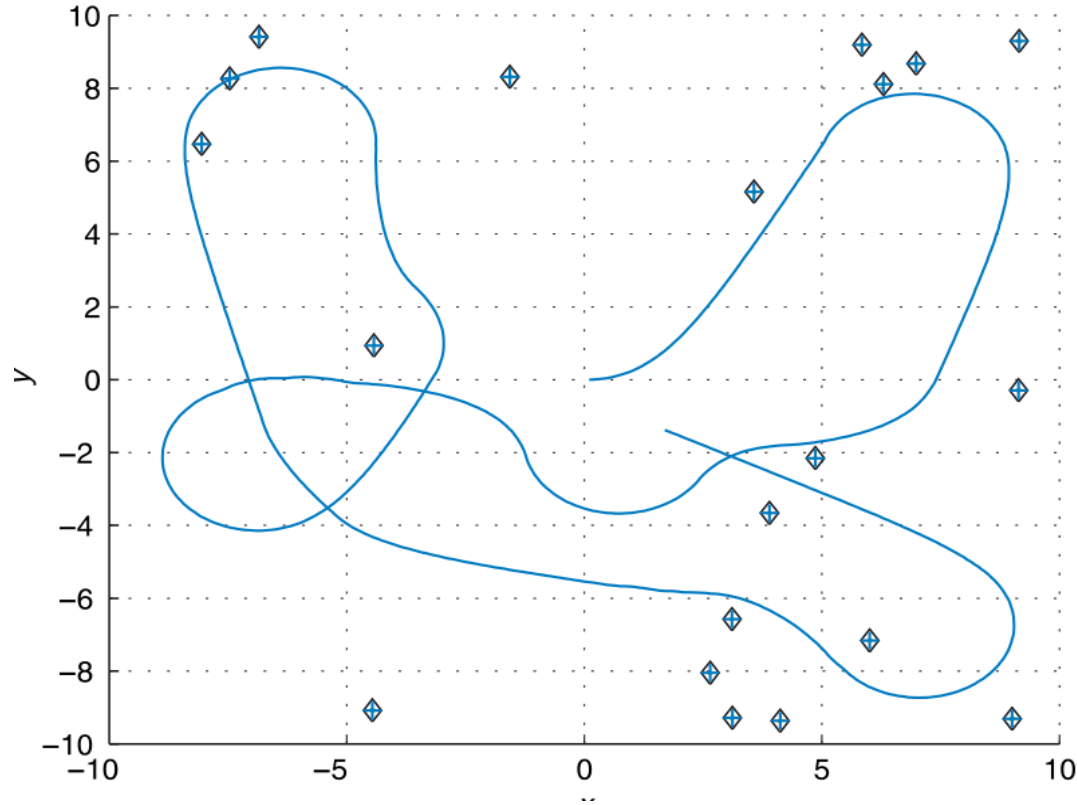
How do we use the EKF to estimate landmark positions?

State:  $\hat{\mathbf{x}} = (x_1, y_1, x_2, y_2, \dots, x_M, y_M)^T$

↑  
Positions of each of the M landmarks (base frame)



# Mapping using the EKF



Process update (no new detections):

$$\hat{\mathbf{x}}\langle k+1|k\rangle = \hat{\mathbf{x}}\langle k|k\rangle$$

$$\hat{\mathbf{P}}\langle k+1|k\rangle = \hat{\mathbf{P}}\langle k|k\rangle$$

Process update (new detections):

$$\mathbf{x}\langle k|k\rangle^* = \begin{pmatrix} \mathbf{x}\langle k|k\rangle \\ \mathbf{g}(\mathbf{x}_v\langle k|k\rangle, \mathbf{z}\langle k\rangle) \end{pmatrix}$$

$$\hat{\mathbf{P}}\langle k|k\rangle^* = \mathbf{Y}_z \begin{pmatrix} \hat{\mathbf{P}}\langle k|k\rangle & \mathbf{0} \\ \mathbf{0} & \hat{\mathbf{W}} \end{pmatrix} \mathbf{Y}_z^T$$

est position of new landmark

covariance of new landmark


where:

$$\mathbf{Y}_z = \frac{\partial \mathbf{y}}{\partial \mathbf{z}} = \begin{pmatrix} \mathbf{I}_{n \times n} & \mathbf{0}_{n \times 2} \\ \mathbf{G}_x & \mathbf{0}_{2 \times n-3} & \mathbf{G}_z \end{pmatrix}$$

$$\mathbf{G}_x = \frac{\partial \mathbf{g}}{\partial \mathbf{x}_v} = \begin{pmatrix} 0 & 0 & 0 \\ 0 & 0 & 0 \end{pmatrix}$$

$$\mathbf{G}_z = \frac{\partial \mathbf{g}}{\partial \mathbf{z}} = \begin{pmatrix} \cos(\theta_v + \theta_z) & -r_z \sin(\theta_v + \theta_z) \\ \sin(\theta_v + \theta_z) & r_z \cos(\theta_v + \theta_z) \end{pmatrix}$$

# Mapping using the EKF




Observation update:

$$\mathbf{H}_x = (0 \cdots \mathbf{H}_{x_i} \cdots 0)$$

$$\mathbf{H}_{x_i} = \frac{\partial \mathbf{h}}{\partial \mathbf{x}_i} = \begin{pmatrix} \frac{x_i - x_v}{r} & \frac{y_i - y_v}{r} \\ -\frac{x_i - x_v}{r^2} & \frac{y_i - y_v}{r^2} \end{pmatrix}$$

# SLAM using the EKF



Estimate both robot position and landmark positions:

$$\hat{\mathbf{x}} = (\underbrace{x_v, y_v, \theta_v}_{\text{robot position}}, \underbrace{x_1, y_1, x_2, y_2, \dots, x_M, y_M}_{\text{Landmark positions}})$$

Vehicle/landmark covariance


Vehicle covariance

$$\hat{\mathbf{P}} = \begin{pmatrix} \hat{\mathbf{P}}_{vv} & \hat{\mathbf{P}}_{vm} \\ \hat{\mathbf{P}}_{vm}^T & \hat{\mathbf{P}}_{mm} \end{pmatrix}$$

Landmark covariance

# SLAM using the EKF

Process update:




$$\mathbf{G}_x = \frac{\partial \mathbf{g}}{\partial \mathbf{x}_v} = \begin{pmatrix} 1 & 0 & -r_z \sin(\theta_v + \theta_z) \\ 0 & 1 & r_z \cos(\theta_v + \theta_z) \end{pmatrix}$$




Same observation update, but using:  $\mathbf{H}_x = (\mathbf{H}_{x_v} \cdots 0 \cdots \mathbf{H}_{x_i} \cdots 0)$


# SLAM using the EKF




(a)



(b)



(c)



(d)

Landmark covariance drops significantly as soon as "loop closure" occurs.

# SLAM using the EKF

$$H_x = (H_{x_v} \cdots 0 \cdots H_{x_i} \cdots 0)$$


$$Y_z = \frac{\partial \mathbf{y}}{\partial \mathbf{z}} = \begin{pmatrix} I_{n \times n} & \mathbf{0}_{n \times 2} \\ \mathbf{G}_x & \mathbf{G}_z \end{pmatrix}$$

$$\mathbf{G}_x = \frac{\partial \mathbf{g}}{\partial \mathbf{x}_v} = \begin{pmatrix} 1 & 0 & -r_z \sin(\theta_v + \theta_z) \\ 0 & 1 & r_z \cos(\theta_v + \theta_z) \end{pmatrix}$$


# Mapping using the EKF

$$H_x = (H_{x_v} \cdots 0 \cdots H_{x_i} \cdots 0)$$


# Process update



T = 1



T = 2



T = 5


- Each time you execute a process update, belief gets more disbursed
- *i.e.* Shannon entropy increases
  - this makes sense: as you predict state further into the future, your uncertainty grows.

$$B'(X_{t+1}) = \sum_{X_t} P(X_{t+1}|X_t, e_{1:t})B(X_t)$$


This is a little like convolution...



# Observation update



Before observation



After observation

Process update *increases* uncertainty

Observation update *decreases* uncertainty

– observations give you more information