

# Principles of Safe Autonomy: Lecture 13: More on filtering and SLAM

Sayan Mitra

Reference: Probabilistic Robotics by Sebastian Thrun, Wolfram Burgard, and Dieter Fox

Slides: From the book's website



# Filtering, estimation

- Applications of Particle filter
- Monte Carlo localization (MCL)
  - Advantages and disadvantages
- Kahoot
- Overview of SLAM
  - Scalability problem and solution
  - Data Association problem
  - Conditional Independence



# Particle filtering algorithm

$X_t = x_t^{[1]}, x_t^{[2]}, \dots, x_t^{[M]}$  particles

Algorithm Particle\_filter( $X_{t-1}, u_t, z_t$ ):

$\bar{X}_t = X_t = \emptyset$

for all  $m$  in  $[M]$  do:

    sample  $x_t^{[m]} \sim p(x_t | u_t, x_{t-1}^{[m]})$

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

$\bar{X}_t = \bar{X}_t + \langle x_t^{[m]}, w_t^{[m]} \rangle$

end for

for all  $m$  in  $[M]$  do:

    draw  $i$  with probability  $\propto w_t^{[i]}$

    add  $x_t^{[i]}$  to  $X_t$

end for

return  $X_t$

ideally,  $x_t^{[m]}$  is selected with probability prop. to  $p(x_t | z_{1:t}, u_{1:t})$

$\bar{X}_t$  is the temporary particle set

// sampling from state transition dist.

// calculates *importance factor*  $w_t$  or weight

// resampling or importance sampling; these are distributed according to  $\eta p(z_t | x_t^{[m]}) \overline{bel}(x_t)$

// survival of fittest: moves/adds particles to parts of the state space with higher probability




# Gaussians

$$p(x) \sim N(\mu, \sigma^2):$$

$$p(x) = \frac{1}{\sqrt{2\pi}\sigma} e^{-\frac{1}{2} \frac{(x-\mu)^2}{\sigma^2}}$$


Univariate



$$p(\mathbf{x}) \sim N(\boldsymbol{\mu}, \boldsymbol{\Sigma}):$$

$$p(\mathbf{x}) = \frac{1}{(2\pi)^{d/2} |\boldsymbol{\Sigma}|^{1/2}} e^{-\frac{1}{2} (\mathbf{x}-\boldsymbol{\mu})' \boldsymbol{\Sigma}^{-1} (\mathbf{x}-\boldsymbol{\mu})}$$

Multivariate



# Properties of Gaussians

$$\left. \begin{array}{l} X \sim N(\mu, \sigma^2) \\ Y = aX + b \end{array} \right\} \Rightarrow Y \sim N(a\mu + b, a^2\sigma^2)$$

$$\left. \begin{array}{l} X_1 \sim N(\mu_1, \sigma_1^2) \\ X_2 \sim N(\mu_2, \sigma_2^2) \end{array} \right\} \Rightarrow p(X_1) \cdot p(X_2) \sim N\left(\frac{\sigma_2^2}{\sigma_1^2 + \sigma_2^2} \mu_1 + \frac{\sigma_1^2}{\sigma_1^2 + \sigma_2^2} \mu_2, \frac{1}{\sigma_1^{-2} + \sigma_2^{-2}}\right)$$



# Multivariate Gaussians

$$\left. \begin{array}{l} X \sim N(\mu, \Sigma) \\ Y = AX + B \end{array} \right\} \Rightarrow Y \sim N(A\mu + B, A\Sigma A^T)$$

$$\left. \begin{array}{l} X_1 \sim N(\mu_1, \Sigma_1) \\ X_2 \sim N(\mu_2, \Sigma_2) \end{array} \right\} \Rightarrow p(X_1) \cdot p(X_2) \sim N\left(\frac{\Sigma_2}{\Sigma_1 + \Sigma_2} \mu_1 + \frac{\Sigma_1}{\Sigma_1 + \Sigma_2} \mu_2, \frac{1}{\Sigma_1^{-1} + \Sigma_2^{-1}}\right)$$

- We stay in the “Gaussian world” as long as we start with Gaussians and perform only linear transformations.



# Discrete Kalman Filter

Estimates the state  $x$  of a discrete-time controlled process that is governed by the linear stochastic difference equation

$$x_t = A_t x_{t-1} + B_t u_t + \varepsilon_t$$

with a measurement

$$z_t = C_t x_t + \delta_t$$



# What is a Kalman Filter?

Suppose we have a system that is governed by a linear difference equation:

$$x_t = A_t x_{t-1} + B_t u_t + \epsilon_t \quad \epsilon_t \sim N(0, R_t); \delta_t \sim N(0, Q_t)$$

with measurement

$$z_t = C_t x_t + \delta_t$$

- Tracks the estimated state of the system by the mean and variance of its state variables -- minimum mean-square error estimator
- Computes the *Kalman gain*, which is used to weight the impact of new measurements on the system state estimate against the predicted value from the process model
- Note that we no longer have discrete states or measurements!





# Components of a Kalman Filter

$A_t$

Matrix ( $n \times n$ ) that describes how the state evolves from  $t$  to  $t-1$  without controls or noise.

$B_t$

Matrix ( $n \times l$ ) that describes how the control  $u_t$  changes the state from  $t$  to  $t-1$ .

$C_t$

Matrix ( $k \times n$ ) that describes how to map the state  $x_t$  to an observation  $z_t$ .

$\varepsilon_t$

Random variables representing the process and measurement noise that are assumed to be independent and normally distributed with covariance  $R_t$  and  $Q_t$  respectively.

$\delta_t$



# Linear Gaussian System: Dynamics

$$x_t = A_t x_{t-1} + B_t u_t + \epsilon_t // p(x_t | x_{t-1}, u_t) = N(A_t x_{t-1} + B_t u_t, R_t)$$
$$bel(x_{t-1}) = N(x_{t-1}; \mu_{t-1}, \Sigma_{t-1})$$

Prediction step

$$\overline{bel}(x_t) = \int p(x_t | u_t, x_{t-1}) bel(x_{t-1}) dx_{t-1}$$

...

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



# Linear Gaussian System: Observations

- $z_t = C_t x_t + \delta_t // p(z_t | x_t) = N(z_t; C_t x_t, Q_t)$

Correction step

$$bel(x_t) = \eta p(z_t | x_t) \overline{bel}(x_t)$$

...

$$= \begin{cases} \mu_t = \bar{\mu}_t + K_t (z_t - C_t \bar{\mu}_t) \\ \Sigma_t = (I - K_t C_t) \bar{\Sigma}_t \end{cases}$$

- where  $K_t = \bar{\Sigma}_t C_t (C_t \bar{\Sigma}_t C_t^T + Q_t)^{-1}$  is called the Kalman gain



# Kalman Filter Algorithm

1. Algorithm `Kalman_Filter`( $\mu_{t-1}, \Sigma_{t-1}, u_t, z_t$ ):

2. Prediction

$$1. \bar{\mu}_t = A_t \mu_{t-1} + B_t u_t$$

$$2. \bar{\Sigma}_t = A_t \Sigma_{t-1} A_t^T + Q_t$$

3. Correction:


$$1. K_t = \bar{\Sigma}_t C_t^T (C_t \bar{\Sigma}_t C_t^T + R_t)^{-1}$$

$$2. \mu_t = \bar{\mu}_t + K_t (z_t - C_t \bar{\mu}_t)$$


$$3. \Sigma_t = (I - K_t C_t) \bar{\Sigma}_t$$

4. Return  $\mu_t, \Sigma_t$





Apply control action




Correction:


1.  $K_t = \bar{\Sigma}_t C_t^T (C_t \bar{\Sigma}_t C_t^T + R_t)^{-1}$
2.  $\mu_t = \bar{\mu}_t + K_t (z_t - C_t \bar{\mu}_t)$
3.  $\Sigma_t = (I - K_t C_t) \bar{\Sigma}_t$

Prediction:



1.  $\bar{\mu}_t = A_t \mu_{t-1} + B_t u_t$
2.  $\bar{\Sigma}_t = A_t \Sigma_{t-1} A_t^T + Q_t$



Get sensor measurement



# Kalman Filter Example



# Who was Rudolf Kalman?



Image Credit: Wikipedia

- Kálmán was one of the most influential people on control theory and is most known for his co-invention of the Kalman filter (or Kalman-Bucy Filter)
- The filtering approach was initially met with vast skepticism, so much so that he was forced to do the first publication of his results in mechanical engineering, rather than in electrical engineering or systems engineering
  - This worked out fine as some of the first use cases was with NASA on the Apollo spacecraft
- *Kalman filters are inside every robot, commercial airplanes, uses in seismic data processing, nuclear power plant instrumentation, and demographic models, as well as applications in econometrics*



# Monte Carlo Localization

- Represents beliefs by particles





# Importance Sampling

suppose we want to compute  $E_f[I(x \in A)]$  but we can only sample from density  $g$

$$E_f[I(x \in A)]$$

$$= \int f(x)I(x \in A)dx$$


$$= \int \frac{f(x)}{g(x)}g(x)I(x \in A)dx, \text{ provided } g(x) > 0$$

$$= \int w(x)g(x)I(x \in A)dx$$

$$= E_g[w(x)I(x \in A)]$$

We need  $f(x) > 0 \Rightarrow g(x) > 0$

**Weight samples:  $w = f/g$**



# Monte Carlo Localization (MCL)

$X_t = x_t^{[1]}, x_t^{[2]}, \dots, x_t^{[M]}$  particles

Algorithm MCL( $X_{t-1}, u_t, z_t, m$ ):

$\bar{X}_{t-1} = X_t = \emptyset$

for all  $m$  in  $[M]$  do:

$x_t^{[m]} = \mathbf{sample\_motion\_model}(u_t, x_{t-1}^{[m]})$

$w_t^{[m]} = \mathbf{measurement\_model}(z_t, x_t^{[m]}, m)$

$\bar{X}_t = \bar{X}_t + \langle x_t^{[m]}, w_t^{[m]} \rangle$

end for

for all  $m$  in  $[M]$  do:

draw  $i$  with probability  $\propto w_t^{[i]}$

add  $x_t^{[i]}$  to  $X_t$

end for

return  $X_t$

Plug in motion and measurement models  
in the particle filter



# Kahoot

- <https://play.kahoot.it/v2/?quizId=3f040019-06e6-4fbe-9c98-780be526f271>



# The SLAM Problem

- SLAM stands for simultaneous localization and mapping
- The task of **building a map** while estimating the pose of the robot relative to this map
- Robot does not have a map, unlike in localization
  
- Why is SLAM hard?  
Chicken and egg problem:  
a map is needed to localize the robot and  
a pose estimate is needed to build a map



# The SLAM Problem


**A robot moving through an unknown, static environment**

## Given:

- The robot's controls
- Observations of nearby features



## Estimate:

- Map of features
- Path of the robot





# SLAM Applications

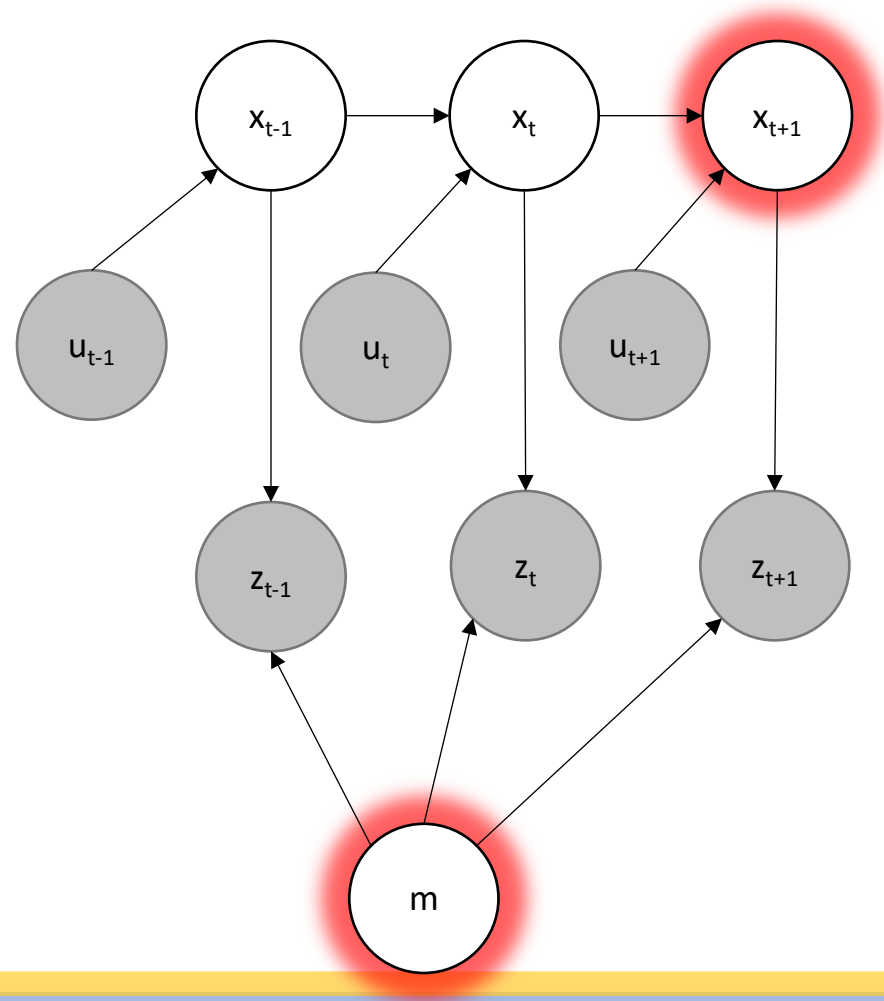


# Forms of SLAM

- State / history
  - Online SLAM:  $p(x_t, m | z_{1:t}, u_{1:t})$
  - Full SLAM:  $p(x_{1:t}, m | z_{1:t}, u_{1:t})$
- Continuous or discrete correspondence variables
  - $p(x_t, m, c_t | z_{1:t}, u_{1:t})$
- Many algorithms: EKFSLAM, GraphSLAM, FastSLAM



# Online SLAM



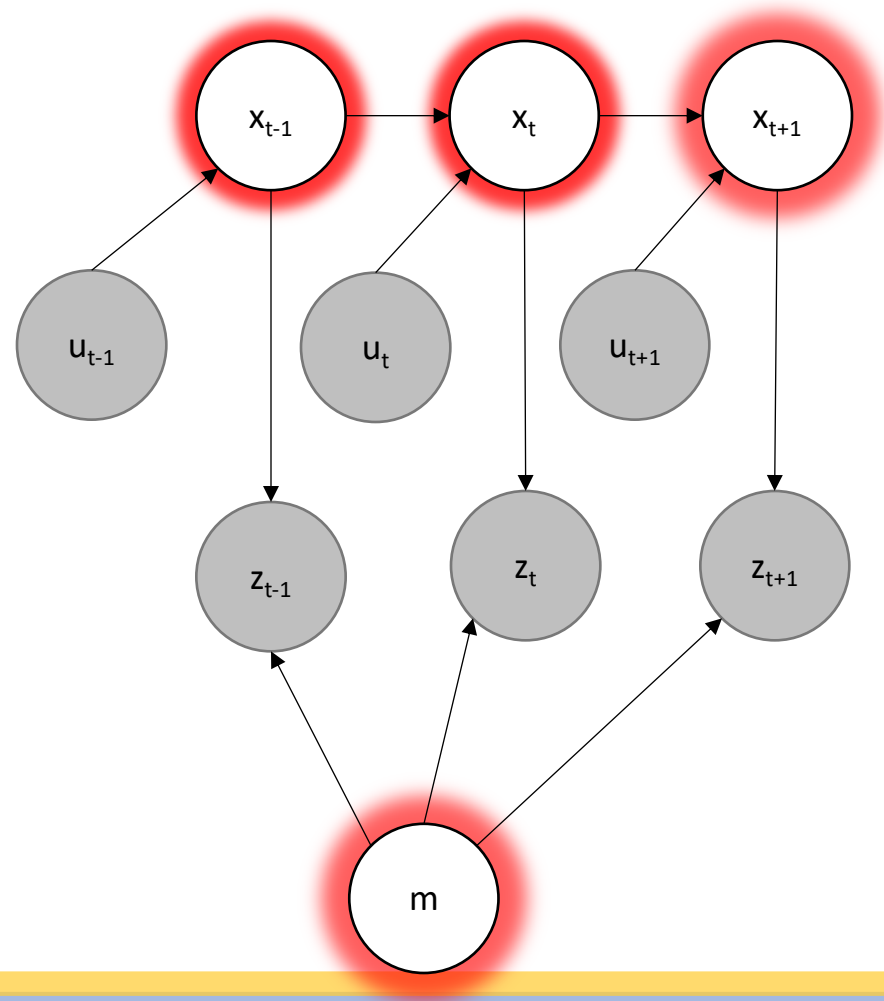
Shaded known:  
control inputs ( $u$ ),  
measurements ( $z$ ).  
White nodes to be  
determined ( $x, m$ )

want to calculate  
 $p(x_t, m | z_{1:t}, u_{1:t})$





# Full SLAM



Shaded known:  
control inputs ( $u$ ),  
measurements( $z$ ). White nodes to  
be determined ( $x, m$ )

want to calculate  
 $p(x_{1:t}, m | z_{1:t}, u_{1:t})$

Continuous  
unknowns:  $x_{1:t}, m$   
Discrete unknowns:  
Relationship of  
detected objects to  
new objects




$$p(x_{1:t}, c_t, m | z_{1:t}, u_{1:t})$$

$c_t$ : corrnspondence  
variable






# Representations

- Grid maps or scans



[Lu & Milios, 97; Gutmann, 98; Thrun 98; Burgard, 99; Konolige & Gutmann, 00; Thrun, 00; Arras, 99; Haehnel, 01;...]

- Landmark-based




[Leonard et al., 98; Castelanos et al., 99; Dissanayake et al., 2001; Montemerlo et al., 2002;...]



# Why is SLAM a hard problem?


**SLAM:** robot path and map are both **unknown**



Robot path error correlates errors in the map



# Why is SLAM a hard problem?



- In the real world, the mapping between observations and landmarks is unknown
- Picking wrong data associations can have catastrophic consequences
- Pose error correlates data associations



# SLAM:

Simultaneous Localization and Mapping

- Full SLAM:

Estimates entire path and map!

$$p(x_{1:t}, m \mid z_{1:t}, u_{1:t})$$

- Online SLAM:


$$p(x_t, m \mid z_{1:t}, u_{1:t}) = \int \int \dots \int p(x_{1:t}, m \mid z_{1:t}, u_{1:t}) dx_1 dx_2 \dots dx_{t-1}$$

Integrations typically done one at a time

Estimates most recent pose and map!



# Data Association Problem



- A data association is an assignment of observations to landmarks
- In general there are more than  $\binom{n}{m}$  (n observations, m landmarks) possible associations
- Also called “assignment problem”



# Localization vs. SLAM

- A particle filter can be used to solve both problems
- Localization: state space  $\langle x, y, \theta \rangle$
- SLAM: state space  $\langle x, y, \theta, map \rangle$ 
  - for landmark maps =  $\langle l_1, l_2, \dots, l_m \rangle$
  - for grid maps =  $\langle c_{11}, c_{12}, \dots, c_{1n}, c_{21}, \dots, c_{nm} \rangle$
- **Problem:** The number of particles needed to represent a posterior grows exponentially with the dimension of the state space!



- Naïve implementation of particle filters to SLAM will be crushed by the curse of dimensionality





# Dependencies

- Is there a dependency between the dimensions of the state space?
- If so, can we use the dependency to solve the problem more efficiently?



# Dependencies

- Is there a dependency between the dimensions of the state space?
- If so, can we use the dependency to solve the problem more efficiently?
  
- In the SLAM context
  - The map depends on the poses of the robot.
  - We know how to build a map given the position of the sensor is known.



# Conditional Independence

- A and B are conditionally independent given C if

$$P(A, B | C) = P(A | C) P(B | C)$$

- Height and vocabulary are not independent
- But they are conditionally independent given age





# Factored Posterior (Landmarks)


poses      map      observations & movements

$$p(x_{1:t}, l_{1:m} \mid z_{1:t}, u_{0:t-1}) = p(x_{1:t} \mid z_{1:t}, u_{0:t-1}) \cdot p(l_{1:m} \mid x_{1:t}, z_{1:t})$$

Factorization first introduced by Murphy in 1999



# Mapping using Landmarks



**Knowledge of the robot's true path renders landmark positions conditionally independent**



# Factored Posterior

$$\begin{aligned} & p(x_{1:t}, l_{1:m} \mid z_{1:t}, u_{0:t-1}) \\ &= p(x_{1:t} \mid z_{1:t}, u_{0:t-1}) \cdot p(l_{1:m} \mid x_{1:t}, z_{1:t}) \\ &= p(x_{1:t} \mid z_{1:t}, u_{0:t-1}) \cdot \prod_{i=1}^M p(l_i \mid x_{1:t}, z_{1:t}) \end{aligned}$$

Robot path posterior  
(localization problem)

Conditionally  
independent  
landmark positions



# Rao-Blackwellization

$$p(x_{1:t}, l_{1:m} \mid z_{1:t}, u_{0:t-1}) = p(x_{1:t} \mid z_{1:t}, u_{0:t-1}) \cdot \prod_{i=1}^M p(l_i \mid x_{1:t}, z_{1:t})$$

- This factorization is also called Rao-Blackwellization
- Given that the second term can be computed efficiently, particle filtering becomes possible!





# David Harold Blackwell (1919-2011)

Independently developed [dynamic](#) programming. Several theorems that bear his name, including the Blackwell renewal theorem, used in engineering, and the Rao-Blackwell theorem in statistics.

University of Illinois at Urbana-Champaign (BA, MA, PhD) and 12 honorary doctorates





photo from [stat.illinois](http://stat.illinois)

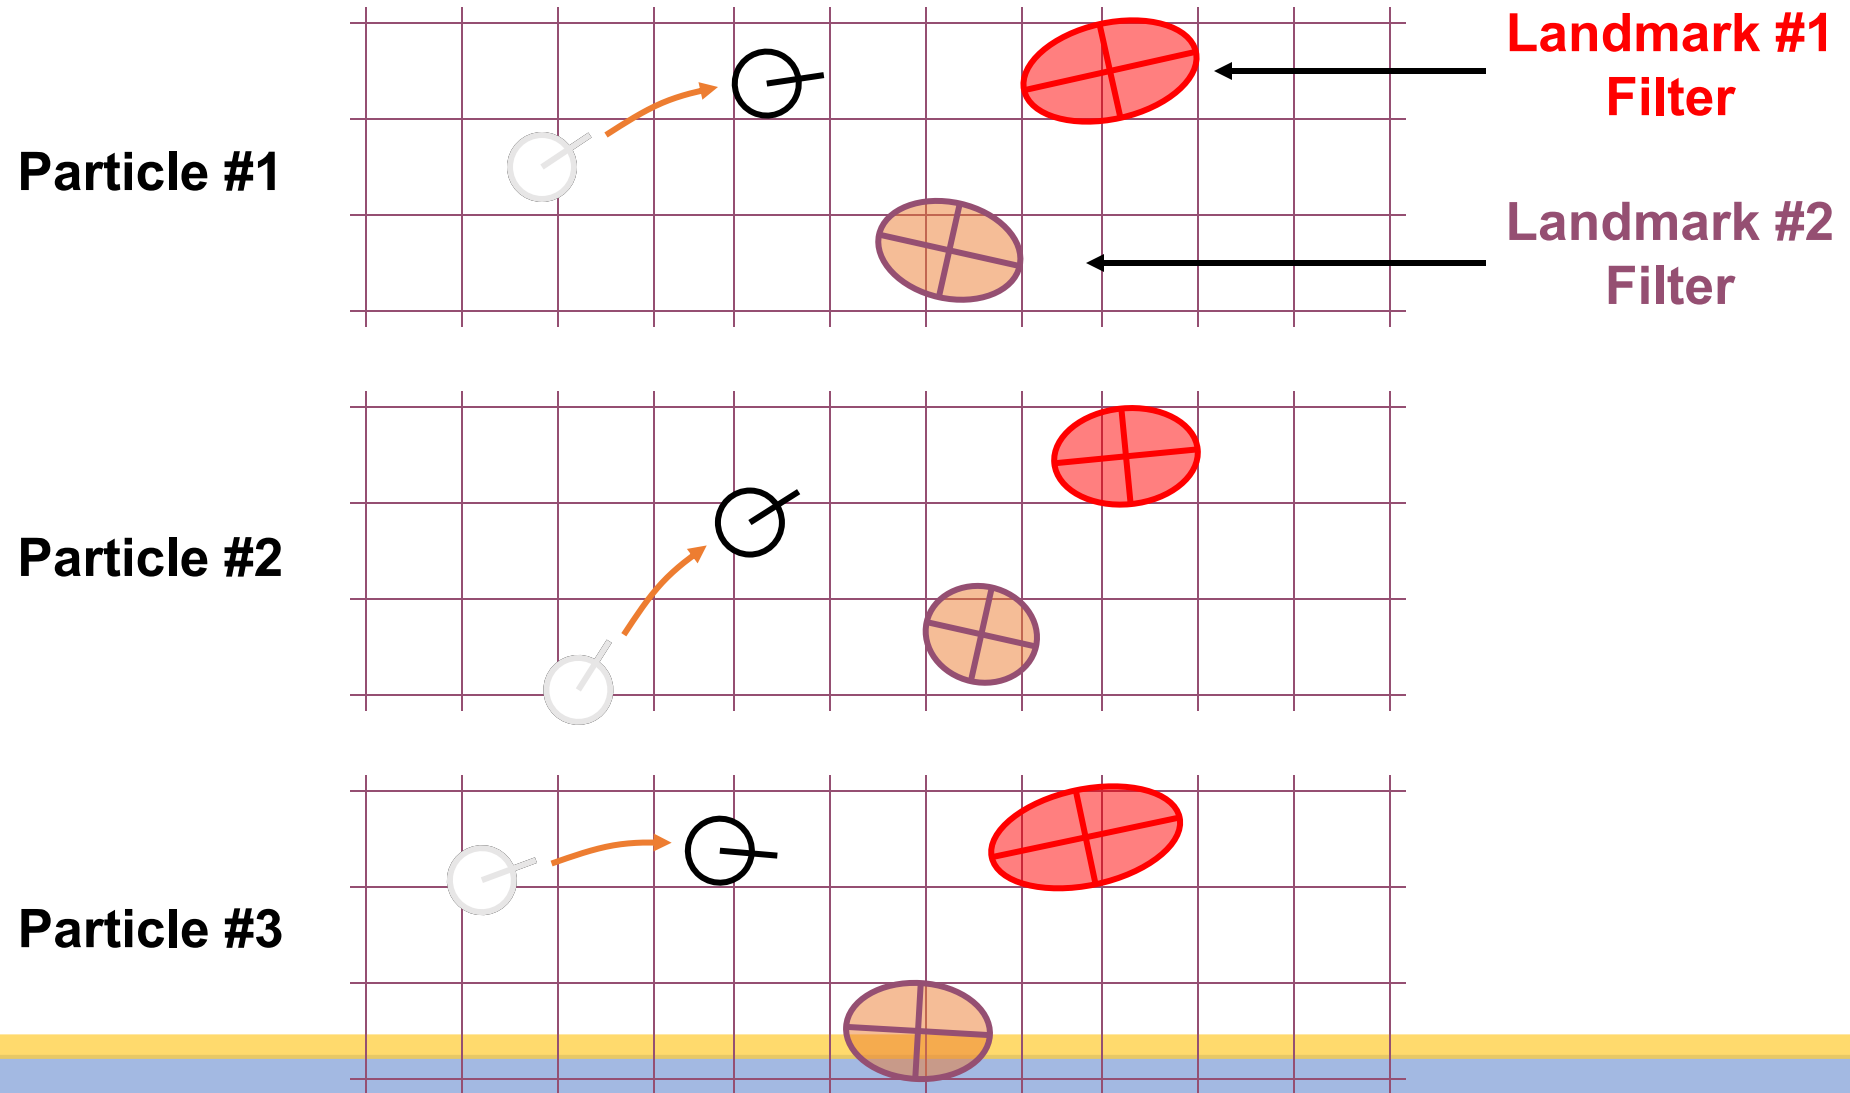


# FastSLAM


- [Rao-Blackwellized](#) particle filtering based on landmarks [Montemerlo et al., 2002]
- Each landmark is represented by a 2x2 Extended Kalman Filter (EKF)
- Each particle therefore has to maintain  $M$  EKFs



# FastSLAM – Action Update




# FastSLAM – Sensor Update




# FastSLAM – Sensor Update

Particle #1




Weight = 0.8

Particle #2



Weight = 0.4


Particle #3



Weight = 0.1



# FastSLAM - Video



# FastSLAM Complexity

- Update robot particles based on control  $u_{t-1}$
- Incorporate observation  $z_t$  into Kalman filters
- Resample particle set

**N = Number of particles**  
**M = Number of map features**

$O(N)$   
Constant time per particle

$O(N \cdot \log(M))$   
Log time per particle

$O(N \cdot \log(M))$   
Log time per particle


---

$O(N \cdot \log(M))$   
Log time per particle



# Data Association Problem

- Which observation belongs to which landmark?




- A robust SLAM must consider possible data associations
- Potential data associations depend also on the pose of the robot






# Multi-Hypothesis Data Association

- Data association is done on a per-particle basis
- Robot pose error is factored out of data association decisions



# Per-Particle Data Association



Was the observation generated by the red or the blue landmark?

$$P(\text{observation}|\text{red}) = 0.3$$

$$P(\text{observation}|\text{blue}) = 0.7$$

- Two options for per-particle data association
  - Pick the most probable match
  - Pick an random association weighted by the observation likelihoods
- If the probability is too low, generate a new landmark




# Results – Victoria Park

- 4 km traverse
- < 5 m RMS position error
- 100 particles

**Blue** = GPS

**Yellow** = FastSLAM



Dataset courtesy of University of Sydney



# Results – Victoria Park



<https://www.youtube.com/watch?v=BIOJSNHYSbc>



# Conclusions FastSLAM

- Maintain set of particles
  - Each particle contains  $s$  sampled robot path and a map
  - Each feature in the map represented by local gaussian
  - Result linear is size of map and number of particles
- Trick is to represent map as a set of separate Gaussians instead of a giant joint distribution
  - Possible because of conditional independence given a path
- Update rule similar to conventional particle filter
- Each particle can be based on a different data association



# More Details on FastSLAM

- M. Montemerlo, S. Thrun, D. Koller, and B. Wegbreit. FastSLAM: A factored solution to simultaneous localization and mapping, *AAAI02*
- D. Haehnel, W. Burgard, D. Fox, and S. Thrun. An efficient FastSLAM algorithm for generating maps of large-scale cyclic environments from raw laser range measurements, *IROS03*
- M. Montemerlo, S. Thrun, D. Koller, B. Wegbreit. FastSLAM 2.0: An Improved particle filtering algorithm for simultaneous localization and mapping that provably converges. *IJCAI-2003*
- G. Grisetti, C. Stachniss, and W. Burgard. Improving grid-based slam with rao-blackwellized particle filters by adaptive proposals and selective resampling, *ICRA05*
- A. Eliazar and R. Parr. DP-SLAM: Fast, robust simultaneous localization and mapping without predetermined landmarks, *IJCAI03*

