


Gaussian (or Normal) Distribution

Univariate

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


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



Multivariate

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

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



Properties of Gaussians

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

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

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



Covariance, covariance, covariance

- Know what's a covariance matrix
 - What it does, what it means, why it matters, how you can compute one, how you can visualize them

$$\bar{\mathbf{x}} = \frac{1}{N} \sum_{i=0}^N \mathbf{x}_i$$

$$\mathbf{e}_i = \mathbf{x}_i - \bar{\mathbf{x}}$$

$$\Sigma = E[\mathbf{e}_i \mathbf{e}_i^T]$$



Kalman Filter

- Recursive solution for discrete linear filtering problems
 - A state $x \in R^n$
 - A measurement $z \in R^m$
 - Discrete (i.e. for time $t = 1, 2, 3, \dots$)
 - Recursive process (i.e. $x_t = f(x_{t-1})$)
 - Linear system (i.e. $x_t = A x_{t-1}$)

- The system is defined by:

1) **linear** process model

$$x_t = A x_{t-1} + B u_{t-1} + w_{t-1}$$

state transition control Input (optional) Gaussian white noise

How a state transitions into another state

2) **linear** measurement model

$$z_t = H x_t + v_t$$

observation model Gaussian white noise

How a state relates to a measurement

Kalman Filter

- Recipe:

1. Start with an initial guess

$$\hat{\mathbf{x}}_0, \Sigma_0$$


2. Compute prior (prediction)

$$\begin{aligned}\hat{\mathbf{x}}'_t &= A\hat{\mathbf{x}}_{t-1} + B\mathbf{u}_{t-1} \\ \Sigma'_t &= A\Sigma_{t-1}A^T + Q\end{aligned}$$

3. Compute posterior (correction)


$$\begin{aligned}K_t &= \Sigma'_t H^T (H\Sigma'_t H^T + R)^{-1} \\ \hat{\mathbf{x}}_t &= \hat{\mathbf{x}}'_t + K_t(\mathbf{z}_t - H\hat{\mathbf{x}}'_t) \\ \Sigma_t &= (1 - K_t H)\Sigma'_t\end{aligned}$$

REPEAT



Minimizes the sum of the expected errors: $\sum e_i^2$


Initial
 $\hat{\mathbf{x}}_0, \Sigma_0$



Predict $\hat{\mathbf{x}}'_1$ from $\hat{\mathbf{x}}_0$ and \mathbf{u}_0

$$\hat{\mathbf{x}}'_1 = A\hat{\mathbf{x}}_0 + B\mathbf{u}_0$$

$$\Sigma'_1 = A\Sigma_0A^T + Q$$




Correction using \mathbf{z}_1

$$K_1 = \Sigma'_1H^T(H\Sigma'_1H^T + R)^{-1}$$

$$\hat{\mathbf{x}}_1 = \hat{\mathbf{x}}'_1 + K_1(\mathbf{z}_1 - H\hat{\mathbf{x}}'_1)$$


$$\Sigma_1 = (1 - K_1H)\Sigma'_1$$



Predict $\hat{\mathbf{x}}'_2$ from $\hat{\mathbf{x}}_1$ and \mathbf{u}_1

$$\hat{\mathbf{x}}'_2 = A\hat{\mathbf{x}}_1 + B\mathbf{u}_1$$

$$\Sigma'_2 = A\Sigma_1A^T + Q$$



Kalman Filter Limitations

- Assumptions:

- Linear process model

$$\mathbf{x}_{t+1} = A\mathbf{x}_t + B\mathbf{u}_t + \mathbf{w}_t$$

- Linear observation model

$$\mathbf{z}_t = H\mathbf{x}_t + \mathbf{v}_t$$

- White Gaussian noise

$$N(0, \Sigma)$$

- What can we do if system is not linear?

- Non-linear state dynamics

$$\mathbf{x}_{t+1} = f(\mathbf{x}_t, \mathbf{u}_t, \mathbf{w}_t)$$

- Non-linear observations

$$\mathbf{z}_t = h(\mathbf{x}_t, \mathbf{v}_t)$$

Extended Kalman Filter

- Kalman Filter Recipe:

- Given

$$\hat{\mathbf{x}}_0, \Sigma_0$$

- Prediction

$$\hat{\mathbf{x}}'_t = A\hat{\mathbf{x}}_{t-1} + B\mathbf{u}_{t-1}$$

$$\Sigma'_t = A\Sigma_{t-1}A^T + Q$$

- Measurement correction

$$K_t = \Sigma'_t H^T (H\Sigma'_t H^T + R)^{-1}$$

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

$$\Sigma_t = (I - K_t H)\Sigma'_t$$

- Extended Kalman Filter Recipe:

- Given

$$\hat{\mathbf{x}}_0, \Sigma_0$$

- Prediction

$$\hat{\mathbf{x}}'_t = \mathbf{f}(\hat{\mathbf{x}}_{t-1}, \mathbf{u}_{t-1}, 0)$$

$$\Sigma'_t = A_t \Sigma_{t-1} A_t^T + W_t Q W_t^T$$

- Measurement correction

$$K_t = \Sigma'_t H_t^T (H_t \Sigma'_t H_t^T + V_t R V_t^T)^{-1}$$


$$\hat{\mathbf{x}}_t = \hat{\mathbf{x}}'_t + K_t(\mathbf{z}_t - \mathbf{h}(\hat{\mathbf{x}}'_t, 0))$$

$$\Sigma_t = (I - K_t H_t)\Sigma'_t$$

EKF for Range-Bearing Localization

- State $\mathbf{s}_t = [x_t \quad y_t \quad \theta_t]^T$ 2D position and orientation
- Input $\mathbf{u}_t = [v_t \quad \omega_t]^T$ linear and angular velocity
- Process model

$$\mathbf{f}(\mathbf{s}_t, \mathbf{u}_t, \mathbf{w}_t) = \begin{bmatrix} x_{t-1} + (\Delta t)v_{t-1}\cos(\theta_{t-1}) \\ y_{t-1} + (\Delta t)v_{t-1}\sin(\theta_{t-1}) \\ \theta_{t-1} + (\Delta t)\omega_{t-1} \end{bmatrix} + \begin{bmatrix} w_{x_t} \\ w_y \\ w_{\theta_t} \end{bmatrix}$$



- Given a map, the robot sees N landmarks with coordinates

$$\mathbf{l}_1 = [x_{l_1} \quad y_{l_1}]^T, \dots, \mathbf{l}_N = [x_{l_N} \quad y_{l_N}]^T$$

The observation model is

$$\mathbf{z}_t = \begin{bmatrix} \mathbf{h}_1(\mathbf{s}_t, \mathbf{v}_1) \\ \vdots \\ \mathbf{h}_N(\mathbf{s}_t, \mathbf{v}_N) \end{bmatrix} \quad \mathbf{h}_i(\mathbf{s}_t, \mathbf{v}_t) = \begin{bmatrix} \sqrt{(x_t - x_{l_i})^2 + (y_t - y_{l_i})^2} \\ \tan^{-1}\left(\frac{y_t - y_{l_i}}{x_t - x_{l_i}}\right) - \theta_t \end{bmatrix} + \begin{bmatrix} v_r \\ v_b \end{bmatrix}$$

Kalman Filters and SLAM

- Localization: state is the location of the robot
- Mapping: state is the location of 2D landmarks
- SLAM: state combines both
- If the state is $\mathbf{s}_t = [x_t \ y_t \ \theta_t \ l_{1t}^T \ \dots \ l_{N_t}^T]^T$

then we can write a linear observation system

- note that if we don't have some fixed landmarks, our system is *unobservable* (we can't fully determine all unknown quantities)

- Covariance Σ is represented by <http://ais.informatik.uni-freiburg.de>

σ_{xx}	σ_{xy}	$\sigma_{x\theta}$	$\sigma_{xm_{1,x}}$	$\sigma_{xm_{1,y}}$	\dots	$\sigma_{xm_{n,x}}$	$\sigma_{xm_{n,y}}$
σ_{yx}	σ_{yy}	$\sigma_{y\theta}$	$\sigma_{ym_{1,x}}$	$\sigma_{ym_{1,y}}$	\dots	$\sigma_{ym_{n,x}}$	$\sigma_{ym_{n,y}}$
$\sigma_{\theta x}$	$\sigma_{\theta y}$	$\sigma_{\theta\theta}$	$\sigma_{\theta m_{1,x}}$	$\sigma_{\theta m_{1,y}}$	\dots	$\sigma_{\theta m_{n,x}}$	$\sigma_{\theta m_{n,y}}$
$\sigma_{m_{1,x}x}$	$\sigma_{m_{1,x}y}$	σ_{θ}	$\sigma_{m_{1,x}m_{1,x}}$	$\sigma_{m_{1,x}m_{1,y}}$	\dots	$\sigma_{m_{1,x}m_{n,x}}$	$\sigma_{m_{1,x}m_{n,y}}$
$\sigma_{m_{1,y}x}$	$\sigma_{m_{1,y}y}$	σ_{θ}	$\sigma_{m_{1,y}m_{1,x}}$	$\sigma_{m_{1,y}m_{1,y}}$	\dots	$\sigma_{m_{1,y}m_{n,x}}$	$\sigma_{m_{1,y}m_{n,y}}$
\vdots	\vdots	\vdots	\vdots	\vdots	\ddots	\vdots	\vdots
$\sigma_{m_{n,x}x}$	$\sigma_{m_{n,x}y}$	σ_{θ}	$\sigma_{m_{n,x}m_{1,x}}$	$\sigma_{m_{n,x}m_{1,y}}$	\dots	$\sigma_{m_{n,x}m_{n,x}}$	$\sigma_{m_{n,x}m_{n,y}}$
$\sigma_{m_{n,y}x}$	$\sigma_{m_{n,y}y}$	σ_{θ}	$\sigma_{m_{n,y}m_{1,x}}$	$\sigma_{m_{n,y}m_{1,y}}$	\dots	$\sigma_{m_{n,y}m_{n,x}}$	$\sigma_{m_{n,y}m_{n,y}}$

Σ

EKF Range Bearing SLAM

Prior State Estimation

- State $\mathbf{s}_t = [x_t \quad y_t \quad \theta_t \quad \mathbf{l}_1^T \quad \dots \quad \mathbf{l}_N^T]$ position/orientation of robot and landmarks coordinates
- Input $\mathbf{u}_t = [v_t \quad \omega_t]^T$ forward and angular velocity
- The process model for localization is

$$\mathbf{s}'_t = \begin{bmatrix} x_{t-1} \\ y_{t-1} \\ \theta_{t-1} \end{bmatrix} + \begin{bmatrix} \Delta t v_{t-1} \cos(\theta_{t-1}) \\ \Delta t v_{t-1} \sin(\theta_{t-1}) \\ \Delta t \omega_{t-1} \end{bmatrix}$$

This model is augmented for $2N+3$ dimensions to accommodate landmarks. This results in the process equation

$$\begin{bmatrix} x'_t \\ y'_t \\ \theta'_t \\ \mathbf{l}'_{1t} \\ \vdots \\ \mathbf{l}'_{Nt} \end{bmatrix} = \begin{bmatrix} x_{t-1} \\ y_{t-1} \\ \theta_{t-1} \\ \mathbf{l}_{1t-1} \\ \vdots \\ \mathbf{l}_{Nt-1} \end{bmatrix} + \begin{bmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \\ 0 & 0 & 0 \\ \vdots & \vdots & \vdots \\ 0 & 0 & 0 \end{bmatrix} \begin{bmatrix} \Delta t v_{t-1} \cos(\theta_{t-1}) \\ \Delta t v_{t-1} \sin(\theta_{t-1}) \\ \Delta t \omega_{t-1} \end{bmatrix}$$


Landmarks don't depend on external input

EKF Range Bearing SLAM

Prior Covariance Update

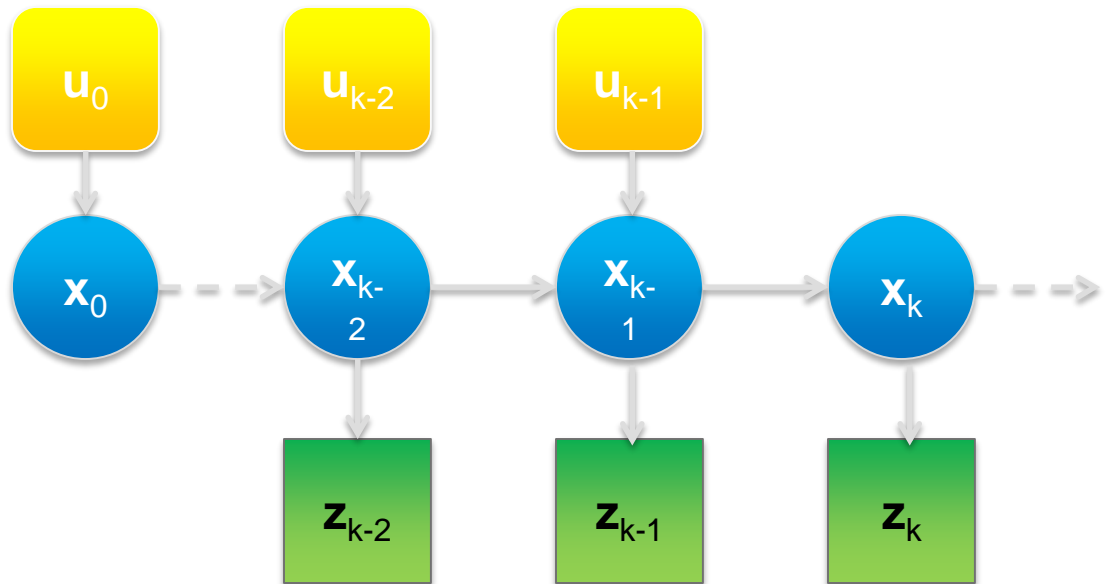
We assume static landmarks. Therefore, the function $f(\mathbf{s}, \mathbf{u}, \mathbf{w})$ only affects the robot's location and not the landmarks.

$$\Sigma'_t = A_t \Sigma_{t-1} A_t^T + W_t Q W_t^T$$



Bayes Filter

- Given a sequence of measurements $\mathbf{z}_1, \dots, \mathbf{z}_k$
- Given a sequence of commands $\mathbf{u}_0, \dots, \mathbf{u}_{k-1}$
- Given a sensor model $P(\mathbf{z}_k | \mathbf{x}_k)$
- Given a dynamic model $P(\mathbf{x}_k | \mathbf{x}_{k-1})$
- Given a prior probability $P(\mathbf{x}_0)$
- Find $P(\mathbf{x}_k | \mathbf{z}_{1:k}, \mathbf{u}_{0:k-1})$



Bayesian Localization

- Recall Bayes Theorem:

$$P(x | z) = \frac{P(z|x)P(x)}{P(z)}$$

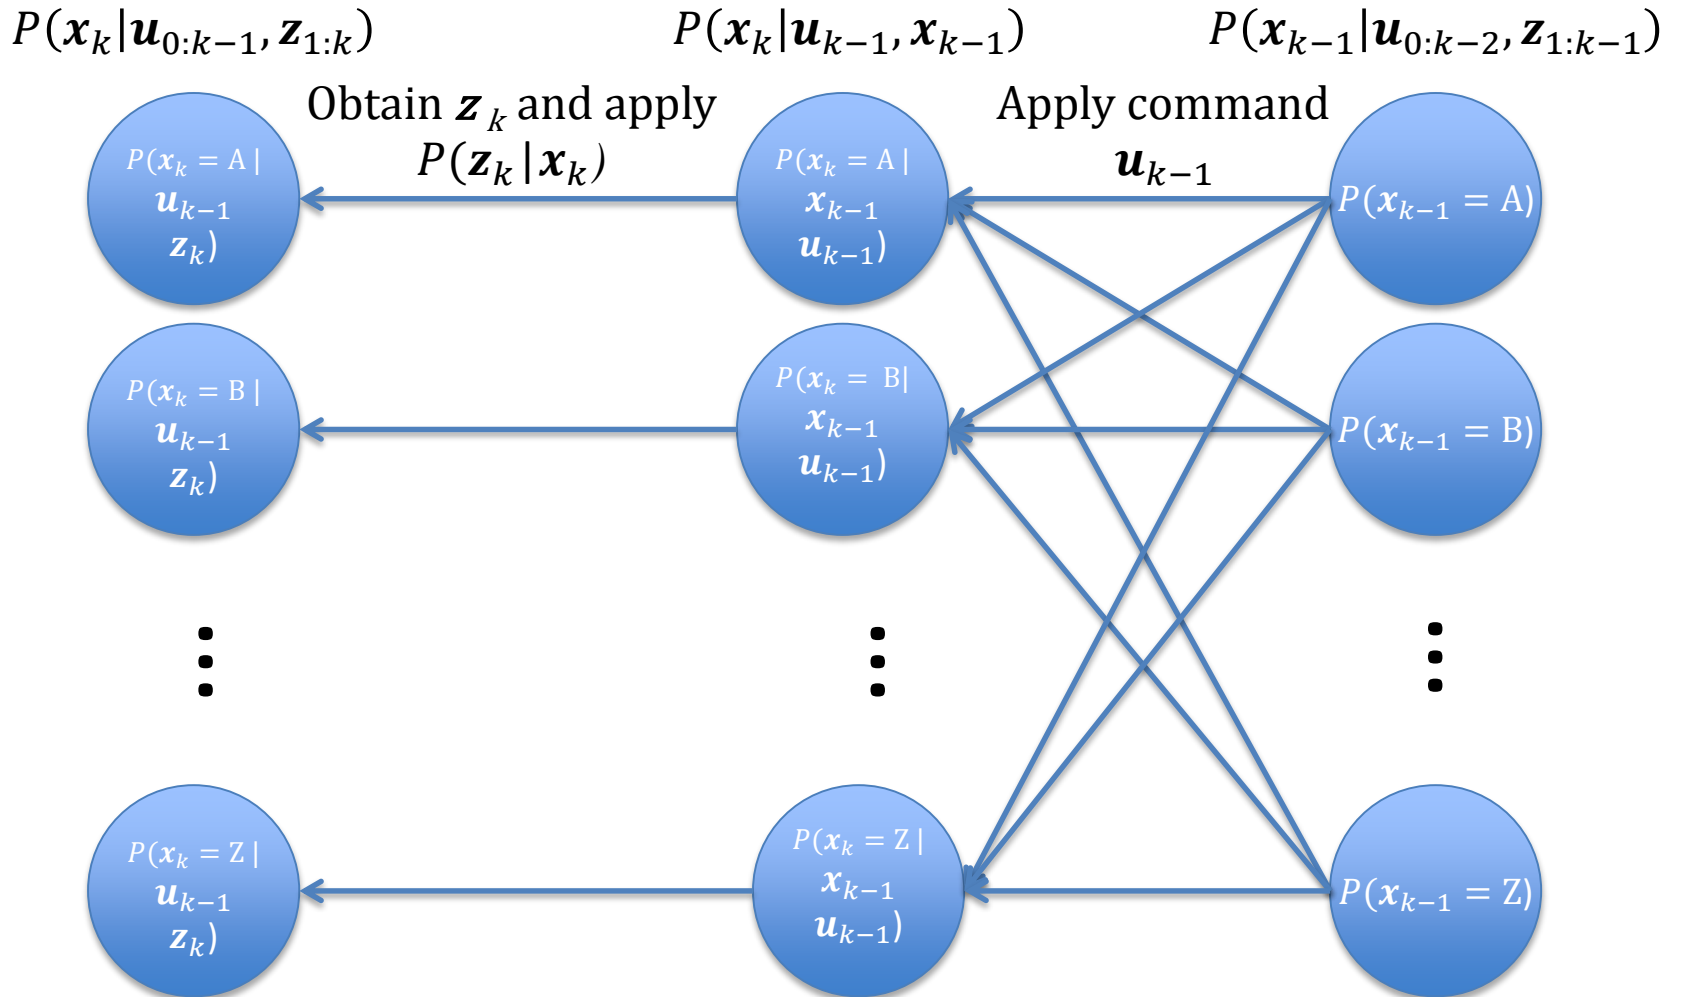
- Also remember conditional independence
- Think of \mathbf{x} as the state of the robot and \mathbf{z} as the data we know

$P(\mathbf{x}_k | \mathbf{u}_{0:k-1}, \mathbf{z}_{1:k}) \longrightarrow$ *Posterior Probability Distribution*

$$P(\mathbf{x}_k | \mathbf{u}_{0:k-1}, \mathbf{z}_{1:k}) = \frac{P(\mathbf{z}_k | \mathbf{x}_k, \mathbf{u}_{0:k-1}, \mathbf{z}_{1:k-1})P(\mathbf{x}_k | \mathbf{u}_{0:k-1}, \mathbf{z}_{1:k-1})}{P(\mathbf{z}_k | \mathbf{u}_{0:k-1}, \mathbf{z}_{1:k-1})}$$

$$= \underbrace{\eta_k P(\mathbf{z}_k | \mathbf{x}_k)}_{\text{observation}} \int_{\mathbf{x}_{k-1}} \underbrace{P(\mathbf{x}_k | \mathbf{u}_{k-1}, \mathbf{x}_{k-1})}_{\text{state prediction}} \underbrace{P(\mathbf{x}_{k-1} | \mathbf{u}_{0:k-2}, \mathbf{z}_{1:k-1})}_{\text{recursion}} d\mathbf{x}_{k-1}$$

Bayes Filter Recap



Predict Motion (Prior Distribution)


- Suppose we have $P(\mathbf{x}_k)$
- We have $P(\mathbf{x}_{k+1}|\mathbf{x}_k, \mathbf{u}_k)$
- Put together

$$P(\mathbf{x}_{k+1}) = \int_{\mathbf{x}_k} P(\mathbf{x}_{k+1}|\mathbf{x}_k, \mathbf{u}_k)P(\mathbf{x}_k)d\mathbf{x}_k$$

What is the probability distribution for \mathbf{x}_{k+1} given the command \mathbf{u}_k and all the previous states \mathbf{x}_k ?

System Model

Prior probability distribution $P(x_k)$



State space $X = \{ 1, 2, 3, 4 \}$


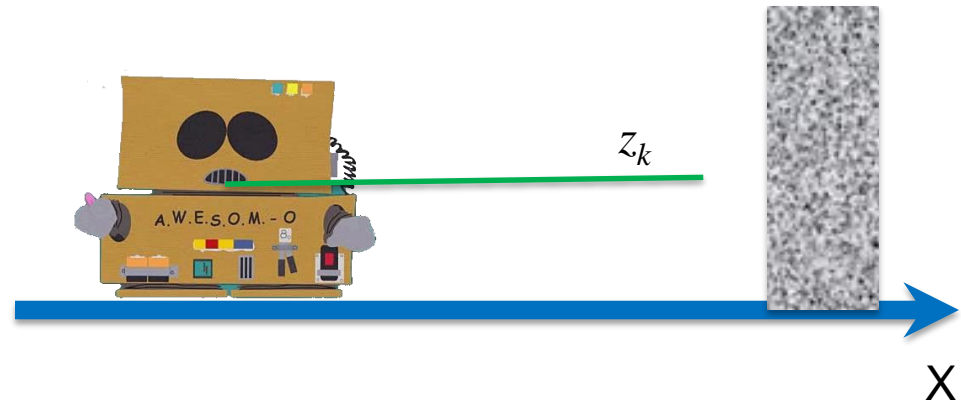
$P(x_{k+1} x_k, u_k)$	$X_{k+1}=1$	$X_{k+1}=2$	$X_{k+1}=3$	$X_{k+1}=4$
$X_k=1$	0.25	0.5	0.25	0
$X_k=2$	0	0.25	0.5	0.25
$X_k=3$	0	0	0.25	0.75
$X_k=4$	0	0	0	1

Transition matrix: The probability $P(j|i)$ of moving from i to j is given by $P_{i,j}$. Each row must sum

Compute $P(x_{k+1}) \stackrel{\text{to } 1}{=} \int_{x_k} P(x_{k+1} | x_k, u_k) dx_k$

Observation Model

- How likely is the measurement z_k given a state x_k ?
- The measurement z_k is what we get. So we have to stick by it
- However, not all states will likely produce the measurement
- We're not interested in finding the most likely state. *We want the whole distribution*



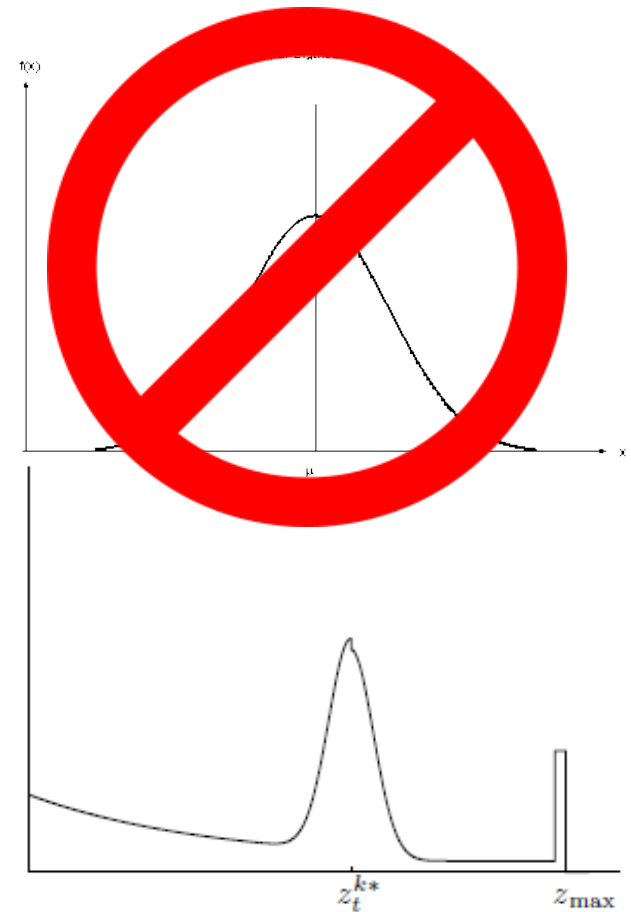
Observation Model Beam Model

- Mixing all these cases together we get

$$P(z_k | \mathbf{x}_k, m) = \begin{bmatrix} w_{\text{hit}} \\ w_{\text{short}} \\ w_{\text{max}} \\ w_{\text{rand}} \end{bmatrix}^T \begin{bmatrix} P_{\text{hit}}(z_k | \mathbf{x}_k, m) \\ P_{\text{short}}(z_k | \mathbf{x}_k, m) \\ P_{\text{max}}(z_k | \mathbf{x}_k, m) \\ P_{\text{rand}}(z_k | \mathbf{x}_k, m) \end{bmatrix}$$

where the weights are parameters

$$w_{\text{hit}} + w_{\text{short}} + w_{\text{max}} + w_{\text{rand}} = 1$$




Probabilistic Robotic



Likelihood Field Model

$$P_{\text{hit}}(z_k | \mathbf{x}_k, m) = \varepsilon_{\sigma_{\text{hit}}^2}(d^2)$$


(a) Gaussian distribution p_{hit}



Probabilistic Robotic



(a) $p_{\text{hit}}(z_t^k | x_t, m)$



Discrete Bayes Filter Algorithm

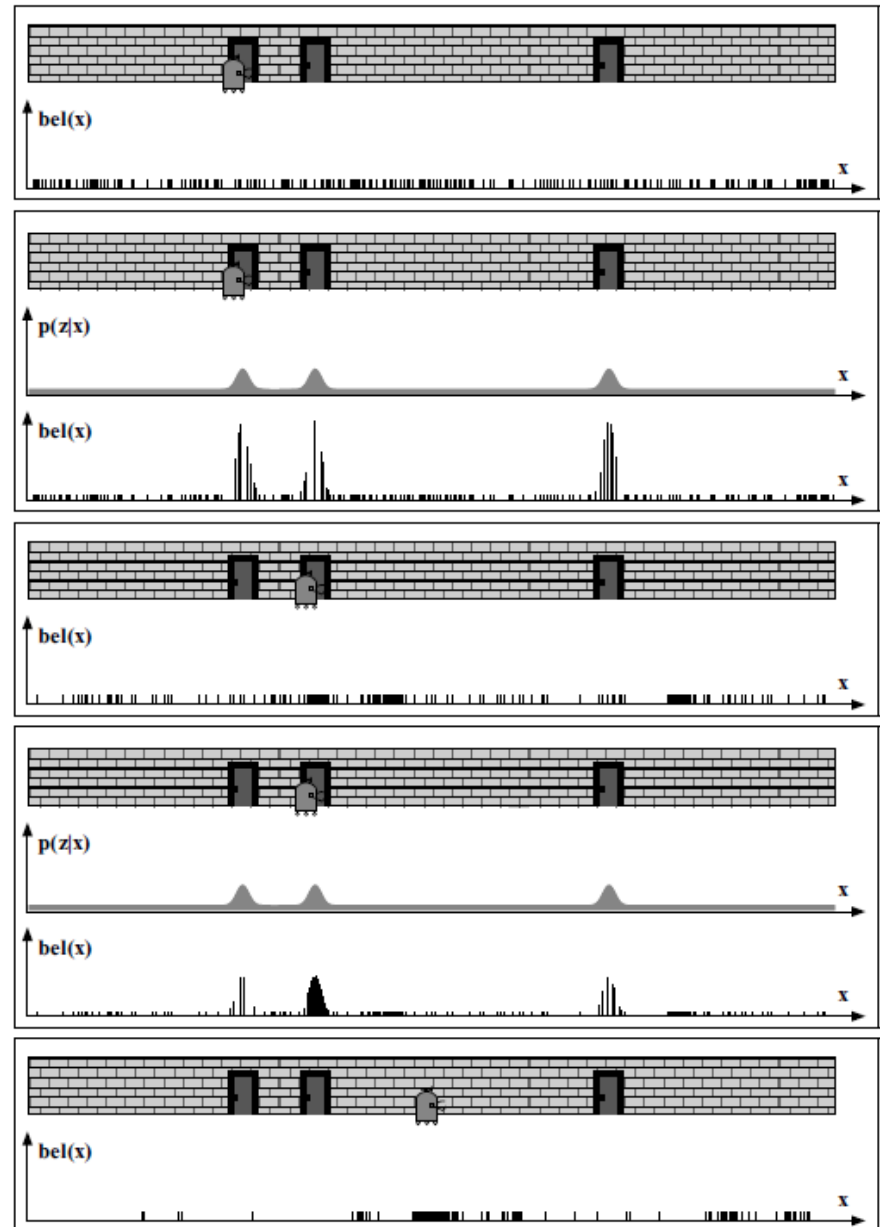
Algorithm **Discrete_Bayes_filter**($u_{0:k-1}, y_{1:k}, P(x_0)$)

1. $P(x) = P(x_0)$ (if you don't know: uniform distribution)
2. **for** $i=1:k$
3. **for all** states $x \in X$
4. $P'(x) = \sum_{x' \in X} P(x | u, x') P(x')$ Prediction given prior dist. and command
5. **end for**
6. $\eta = 0$ Normalization constant
7. **for all** states $x \in X$
8. $P(x) = P(z | x) P'(x)$ Update using measurement
9. $\eta = \eta + P(x)$
10. **end for**
11. **for all** states $x \in X$
12. $P(x) = P(x) / \eta$ Normalize to 1
13. **end for**
14. **end for**

Particle Filter

- Computing $P(\mathbf{z}_k | \mathbf{x}_k, m)$ and $P(\mathbf{x}_{k+1} | \mathbf{x}_k, \mathbf{u}_k)$ is not easy
 - In practice, it is never directly computable
 - Need to propagate an entire conditional distribution, not just one state like we did with Kalman,
- Represent probability distribution by random **samples**
- Estimation of **non-Gaussian, nonlinear** processes

Particle Filter



Particle Filter Algorithm

Algorithm **Particle_filter**($u_{0:k-1}$, $y_{1:k}$, $P(x_0)$, set of N samples $\mathcal{M} = \{x_j, w_j\}$)

1. **for** $i=1:k$
2. **for** $j=1:N$
3. compute a new state x by sampling according to $P(x | u_{i-1}, x_j)$
4. $x_j = x$
5. **end**
6. $\eta=0$
7. **for** $j=1:N$
8. $w_j = P(z_i | x_j)$
9. $\eta += w_j$
10. **end**
11. **for** $j=1:N$
12. $w_j = w_j / \eta$
13. **end**
14. resample (\mathcal{M}) according to weights w_j
15. **end**