CS231A

Computer Vision:

From 3D Reconstruction to Recognition

Optimal Estimation

Perception as a Continuous Process

Perception as a Multi-Modal Experience

Perception as Inference

Recursive State Estimation

Mathematical Formalism to:

- continuously integrate measurements
- from different sensor sources
- to infer the state of a latent variable

What is a state? What is a representation?

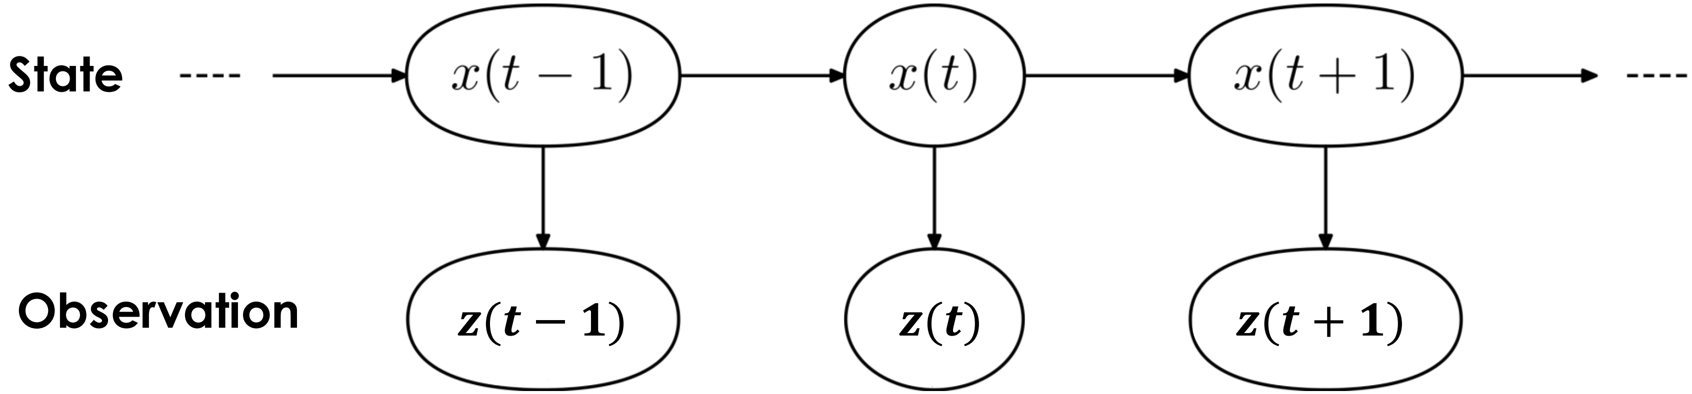

Hidden Markov Model

Representations for Autonomous Driving

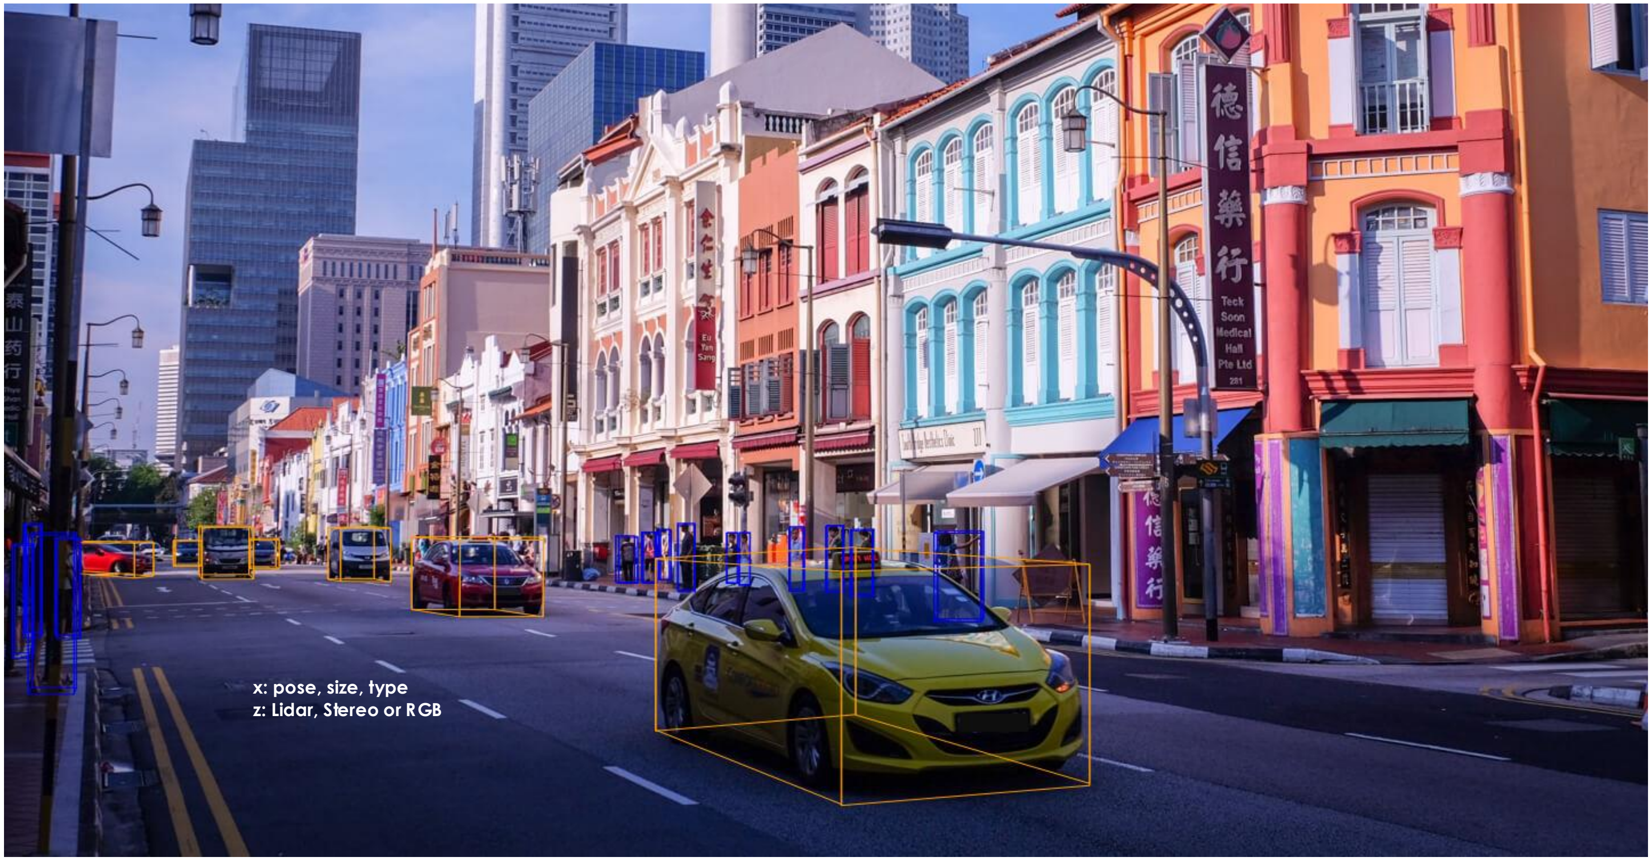

Image adapted from NuScenes by Motional. nuscenes.org

Representations for Manipulation

Manuel Wühtrich et al. "Probabilistic Object Tracking using a Depth Camera", IROS 2013

Why do we care about state estimation in Robotics?

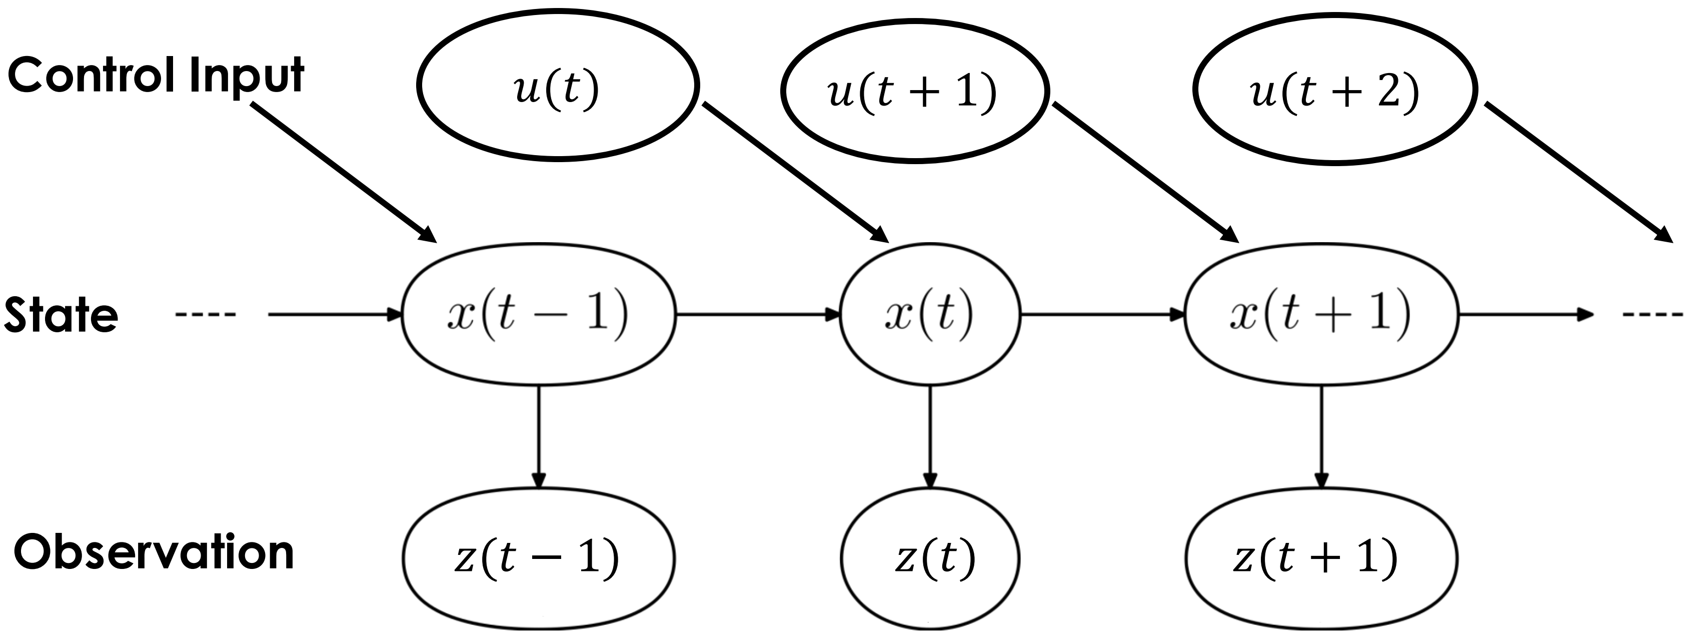

Partially Observable Markov Decision Process

Today

- Intro: Why state estimation?
- Bayes Filter
- Kalman Filter
- Extended Kalman Filter

- For more depth:
  - **AA 273: State Estimation and Filtering for Robotic Perception – Mac Schwager**

The Agent and the Environment

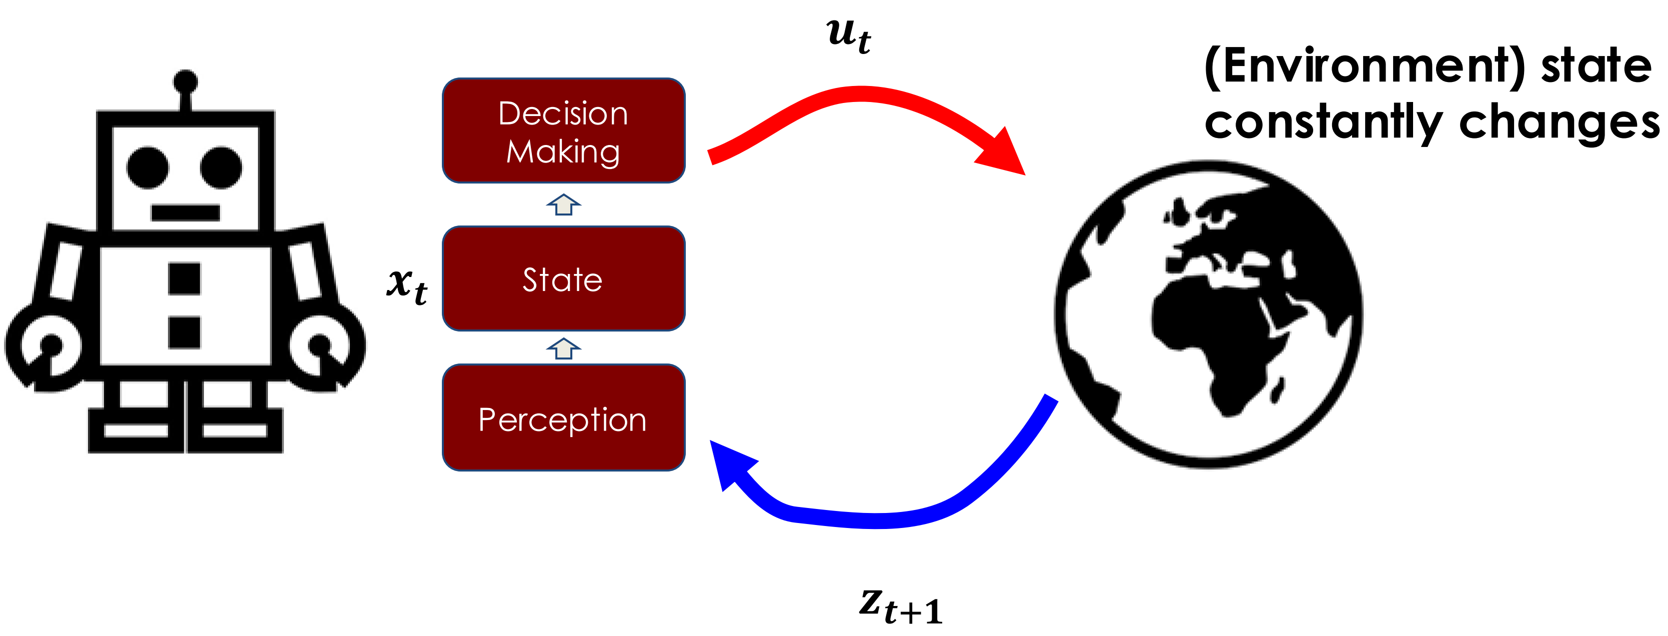

Notation

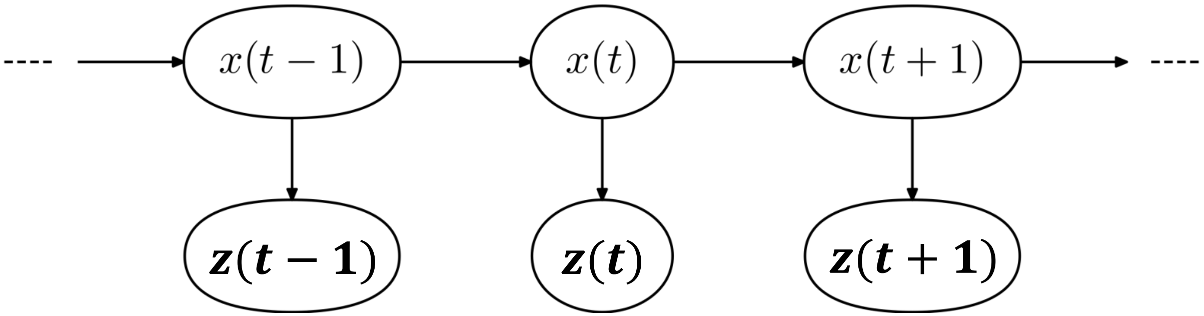

x State of dynamical system, dim n

x_t Instantiation of system state at time t

z Sensor Observation Vector, dim k

z_t Specific Observation at time t

u Robot action / control input, dim m

u_t Robot action / control input at time t

$p(x_t | z_{0:t}, u_{0:t})$ Probability distribution

Markov Assumption

State is complete

Probability Theory Refresh

Let X denote a random variable and x denote a specific event that X might take on.

$p(X = x)$ Probability that X takes on value x

$p(X = x) \geq 0.$

$p(X = x) = p(x)$ Notation shorthand

$\int p(x)\, dx = 1\,.$

Gaussians

$$p(x) = \left(2\pi\sigma^2\right)^{-\frac{1}{2}} \exp\left\{-\frac{1}{2}\frac{(x-\mu)^2}{\sigma^2}\right\}$$

$$p(x) = \det\left(2\pi\Sigma\right)^{-\frac{1}{2}} \exp\left\{-\tfrac{1}{2}(x-\mu)^T\Sigma^{-1}(x-\mu)\right\}$$

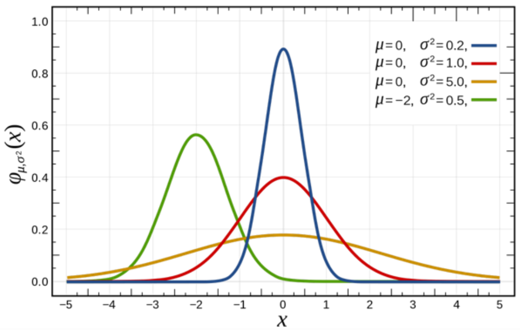

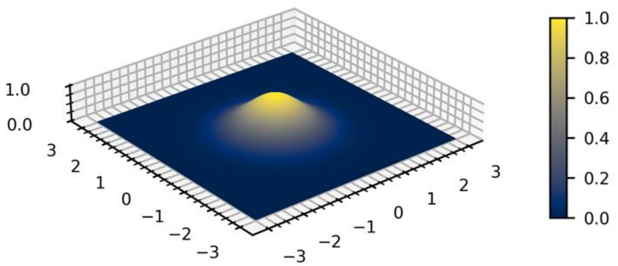

Probability Theory Refresher

If X and Y are *independent*, we have

$$p(x,y) = p(x)\,p(y)\,.$$

If X and Y are not independent, we have

$$p(x,y) = p(x|y)p(y)$$

$$p(x,y) = p(y|x)p(x)$$

$$p(x \mid y) = \frac{p(x,y)}{p(y)}$$

If X and Y are *independent*, we have $p(x \mid y) = \frac{p(x)\,p(y)}{p(y)} = p(x)\,.$

theorem of total probability:

$$p(x) = \sum_y p(x \mid y)\,p(y) \qquad \text{(discrete case)}$$

$$p(x) = \int p(x \mid y)\,p(y)\,dy \qquad \text{(continuous case)}$$

Bayes rule,

$$p(x \mid y) = \frac{p(y \mid x)\,p(x)}{p(y)} = \frac{p(y \mid x)\,p(x)}{\sum_{x'} p(y \mid x')\,p(x')} \qquad \text{(discrete)}$$

$$p(x \mid y) = \frac{p(y \mid x)\,p(x)}{p(y)} = \frac{p(y \mid x)\,p(x)}{\int p(y \mid x')\,p(x')\,dx'} \qquad \text{(continuous)}$$

Probability Theory Refresher

Expectations

Discrete $$E[X] = \sum_x x\, p(x)\,,$$

Continuous $$E[X] = \int x\, p(x)\, dx\,.$$

$$\mathrm{Cov}[X] = E[X - E[X]]^2 = E[X^2] - E[X]^2$$

Probabilistic Generative Laws

- Evolution of state and measurement governed by probabilistic laws
- x_t generated stochastically

State Transition Model

- Probability distribution conditioned on all previous states, measurements and controls

$$p(x_t \mid x_{0:t-1}, z_{1:t-1}, u_{1:t})$$

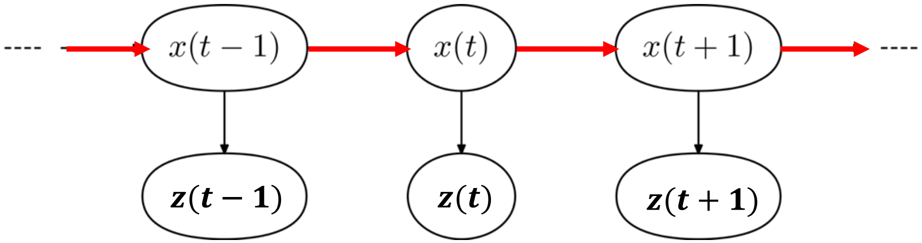

- Assumption: State complete

$$p(x_t \mid x_{0:t-1}, z_{1:t-1}, u_{1:t}) \quad = \quad p(x_t \mid x_{t-1}, u_t)$$

Measurement Model

- Probability distribution conditioned on all previous states, measurements and controls

$$p(z_t \mid x_{0:t}, z_{1:t-1}, u_{1:t})$$

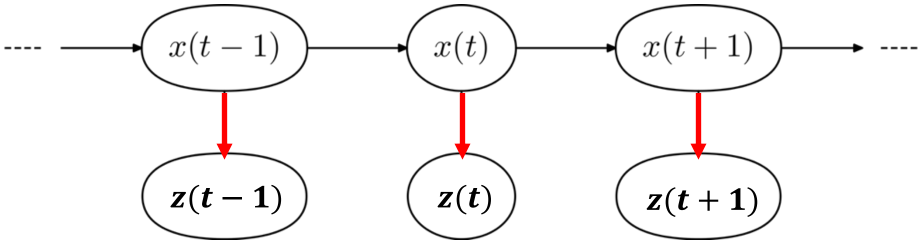

- Assumption: State complete

$$p(z_t \mid x_{0:t}, z_{1:t-1}, u_{1:t}) \quad = \quad p(z_t \mid x_t)$$

Belief Distribution

- Assigns probability to each possible hypothesis about what the true state may be
- Posterior distributions over state conditioned on all the data

$$bel(x_t) \quad = \quad p(x_t \mid z_{1:t}, u_{1:t})$$

- Before incorporating measurement $\boldsymbol{z_t}$ = prediction

$$\overline{bel}(x_t) \quad = \quad p(x_t \mid z_{1:t-1}, u_{1:t})$$

The Bayes Filter

- Recursive filter for estimating x_t only from x_{t-1}, z_t and u_t and not from the ever-growing history $z_{1:t}, u_{1:t}$

Algorithm Bayes_filter$(bel(x_{t-1}), u_t, z_t)$**:** Transition/Dynamics model

for all x_t do

$\overline{bel}(x_t) = \int \boxed{p(x_t \mid u_t, x_{t-1})}\, bel(x_{t-1})\, dx$ Predict Step

$bel(x_t) = \eta\, \boxed{p(z_t \mid x_t)}\, \overline{bel}(x_t)$ Update Step

endfor

return $bel(x_t)$

Measurement Model

Simple example – Belief & Measurement Model

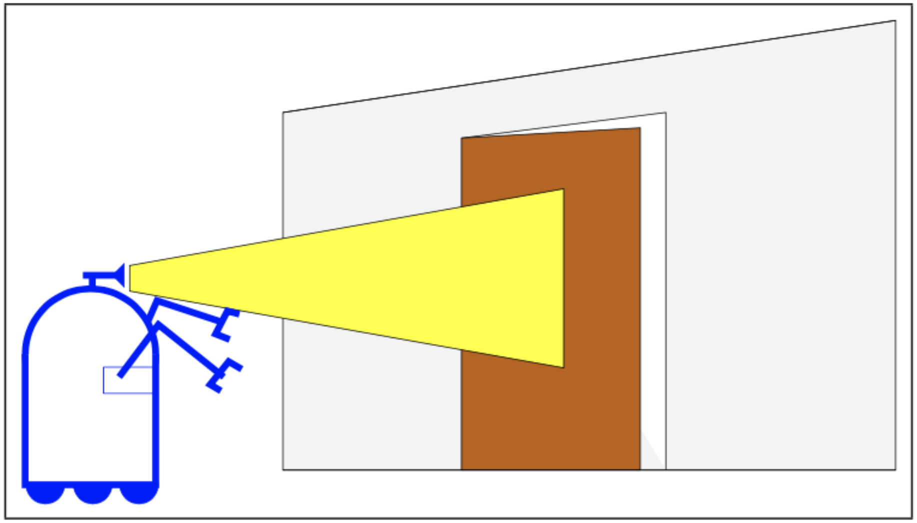

Figure 2.2 A mobile robot estimating the state of a door.

$$bel(X_0 = \textbf{open}) = 0.5$$
$$bel(X_0 = \textbf{closed}) = 0.5$$

$$p(Z_t = \textbf{sense\_open} \mid X_t = \textbf{is\_open}) = 0.6$$
$$p(Z_t = \textbf{sense\_closed} \mid X_t = \textbf{is\_open}) = 0.4$$

$$p(Z_t = \textbf{sense\_open} \mid X_t = \textbf{is\_closed}) = 0.2$$
$$p(Z_t = \textbf{sense\_closed} \mid X_t = \textbf{is\_closed}) = 0.8$$

Simple example – Transition Model

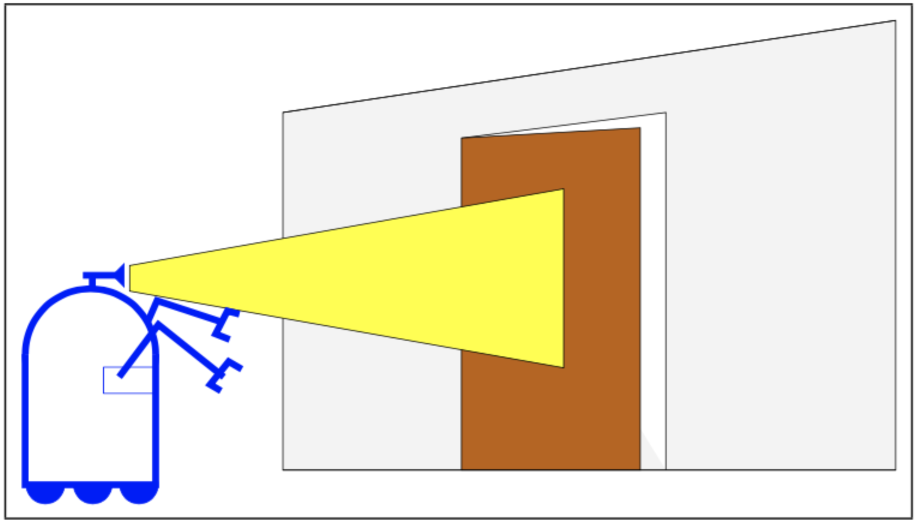

Figure 2.2 A mobile robot estimating the state of a door.

$$p(X_t = \textbf{is\_open} \mid U_t = \textbf{push}, X_{t\_1} = \textbf{is\_open}) = 1$$
$$p(X_t = \textbf{is\_closed} \mid U_t = \textbf{push}, X_{t\_1} = \textbf{is\_open}) = 0$$
$$p(X_t = \textbf{is\_open} \mid U_t = \textbf{push}, X_{t\_1} = \textbf{is\_closed}) = 0.8$$
$$p(X_t = \textbf{is\_closed} \mid U_t = \textbf{push}, X_{t\_1} = \textbf{is\_closed}) = 0.2$$

$$p(X_t = \textbf{is\_open} \mid U_t = \textbf{do\_nothing}, X_{t\_1} = \textbf{is\_open}) = 1$$
$$p(X_t = \textbf{is\_closed} \mid U_t = \textbf{do\_nothing}, X_{t\_1} = \textbf{is\_open}) = 0$$
$$p(X_t = \textbf{is\_open} \mid U_t = \textbf{do\_nothing}, X_{t\_1} = \textbf{is\_closed}) = 0$$
$$p(X_t = \textbf{is\_closed} \mid U_t = \textbf{do\_nothing}, X_{t\_1} = \textbf{is\_closed}) = 1$$

The Bayes Filter - Derivation

- Bayes Rule

$$\boldsymbol{p(a,b) = p(a|b)p(b) = p(b|a)p(a)}$$

$$\boldsymbol{p(a|b) = \frac{p(b|a)p(a)}{p(b)}}$$

$$p(x_t \mid z_{1:t}, u_{1:t}) \quad = \quad \frac{p(z_t \mid x_t, z_{1:t-1}, u_{1:t})\; p(x_t \mid z_{1:t-1}, u_{1:t})}{\boxed{p(z_t \mid z_{1:t-1}, u_{1:t})}}$$

Normalization

The Bayes Filter - Derivation

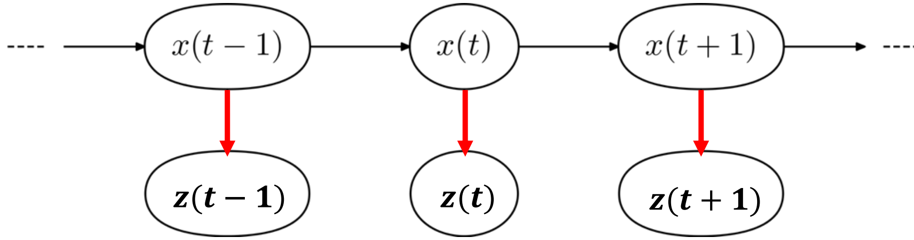

- State is complete

$$p(z_t \mid x_t, z_{1:t-1}, u_{1:t}) \quad = \quad p(z_t \mid x_t)$$

- Simplify

$$\begin{aligned}
p(x_t \mid z_{1:t}, u_{1:t}) &= \frac{p(z_t \mid x_t, z_{1:t-1}, u_{1:t})\, p(x_t \mid z_{1:t-1}, u_{1:t})}{p(z_t \mid z_{1:t-1}, u_{1:t})} \\
&= \eta\, p(z_t \mid x_t, z_{1:t-1}, u_{1:t})\, p(x_t \mid z_{1:t-1}, u_{1:t}) \\
&= \eta\, \underbrace{\boxed{p(z_t \mid x_t)}}_{\text{simplified}}\, p(x_t \mid z_{1:t-1}, u_{1:t})
\end{aligned}$$

The Bayes Filter - Derivation

$$p(x_t \mid z_{1:t}, u_{1:t}) \quad = \quad \eta \, p(z_t \mid x_t) \, \boxed{p(x_t \mid z_{1:t-1}, u_{1:t})}$$

$$bel(x_t) \quad = \quad \eta \, p(z_t \mid x_t) \, \overline{bel}(x_t)$$

This still depends on entire history

Algorithm Bayes_filter$(bel(x_{t-1}), u_t, z_t)$**:**
 for all x_t do
  $\overline{bel}(x_t) = \int p(x_t \mid u_t, x_{t-1}) \, bel(x_{t-1}) \, dx$
  $bel(x_t) = \eta \, \boxed{p(z_t \mid x_t)} \, \overline{bel}(x_t)$ **Update Step**
 endfor
 return $bel(x_t)$

Measurement Model

The Bayes Filter - Derivation

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

- Total probability $p(a) = \int p(a|b)p(b)db$

$$\overline{bel}(x_t) = p(x_t \mid z_{1:t-1}, u_{1:t})$$
$$= \int \underline{p(x_t \mid x_{t-1}, z_{1:t-1}, u_{1:t})}\, \boxed{p(x_{t-1} \mid z_{1:t-1}, u_{1:t})}\, dx_{t-1}$$

Previous Belief over x

- State is complete

$$\underline{p(x_t \mid x_{t-1}, z_{1:t-1}, u_{1:t})} = p(x_t \mid x_{t-1}, u_t)$$

The Bayes Filter - Derivation

$$p(x_t \mid x_{t-1}, z_{1:t-1}, u_{1:t}) \quad = \quad p(x_t \mid x_{t-1}, u_t)$$

$$\begin{aligned}
\overline{bel}(x_t) &= p(x_t \mid z_{1:t-1}, u_{1:t}) \\
&= \int p(x_t \mid x_{t-1}, z_{1:t-1}, u_{1:t})\, p(x_{t-1} \mid z_{1:t-1}, u_{1:t})\, dx_{t-1} \\
&= \int \boxed{p(x_t \mid x_{t-1}, u_t)}\, p(x_{t-1} \mid z_{1:t-1}, u_{1:t-1})\, dx_{t-1}
\end{aligned}$$

simplified

1: **Algorithm Bayes_filter**$(bel(x_{t-1}), u_t, z_t)$**:**
2: for all x_t do
3: $\overline{bel}(x_t) = \int \boxed{p(x_t \mid u_t, x_{t-1})}\, bel(x_{t-1})\, dx$
4: $bel(x_t) = \eta\, p(z_t \mid x_t)\, \overline{bel}(x_t)$
5: endfor
6: return $bel(x_t)$

Transition/Dynamics model

Predict Step

Limitations

1. $p(x)$ is defined $\forall x$ – intractable
   - Discrete and small spaces
   - Continuous and/or large spaces – Moments, Finite # of samples
2. The integral term -> costly to compute

Re-Iterate Example

- Is door open or not?

$$
\begin{aligned}
bel(X_0 = \textbf{open}) &= 0.5 \\
bel(X_0 = \textbf{closed}) &= 0.5
\end{aligned}
$$

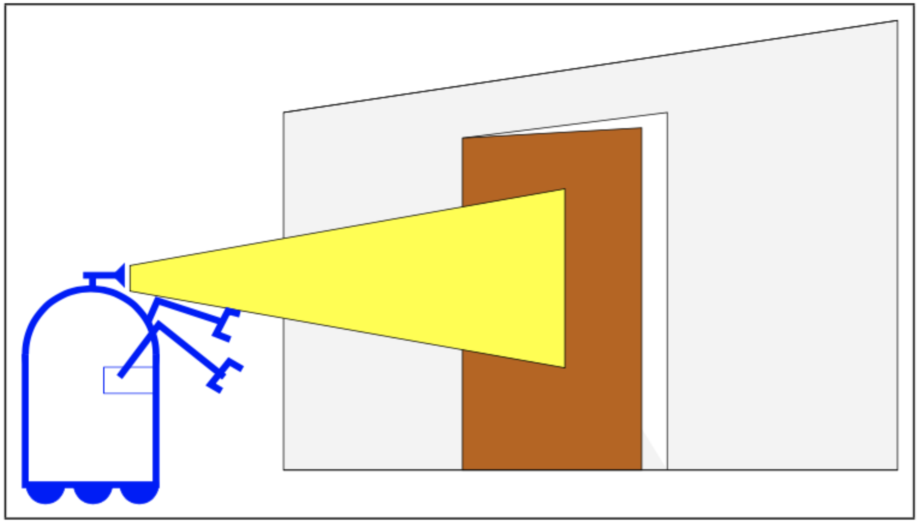

Figure 2.2 A mobile robot estimating the state of a door.

Measurement Model

$$
\begin{aligned}
p(Z_t = \textbf{sense\_open} \mid X_t = \textbf{is\_open}) &= 0.6 \\
p(Z_t = \textbf{sense\_closed} \mid X_t = \textbf{is\_open}) &= 0.4
\end{aligned}
$$

$$
\begin{aligned}
p(Z_t = \textbf{sense\_open} \mid X_t = \textbf{is\_closed}) &= 0.2 \\
p(Z_t = \textbf{sense\_closed} \mid X_t = \textbf{is\_closed}) &= 0.8
\end{aligned}
$$

Transition Model for do_nothing

$$
\begin{aligned}
p(X_t = \textbf{is\_open} \mid U_t = \textbf{do\_nothing}, X_{t\_1} = \textbf{is\_open}) &= 1 \\
p(X_t = \textbf{is\_closed} \mid U_t = \textbf{do\_nothing}, X_{t\_1} = \textbf{is\_open}) &= 0 \\
p(X_t = \textbf{is\_open} \mid U_t = \textbf{do\_nothing}, X_{t\_1} = \textbf{is\_closed}) &= 0 \\
p(X_t = \textbf{is\_closed} \mid U_t = \textbf{do\_nothing}, X_{t\_1} = \textbf{is\_closed}) &= 1
\end{aligned}
$$

Re-Iterate Example

- Is door open or not?

$$bel(X_0 = \textbf{open}) = 0.5$$
$$bel(X_0 = \textbf{closed}) = 0.5$$

Received Sensor Measurement:

$$bel(x_1) = \eta\, p(Z_1 = \textbf{sense\_open} \mid x_1)\, \overline{bel}(x_1)$$

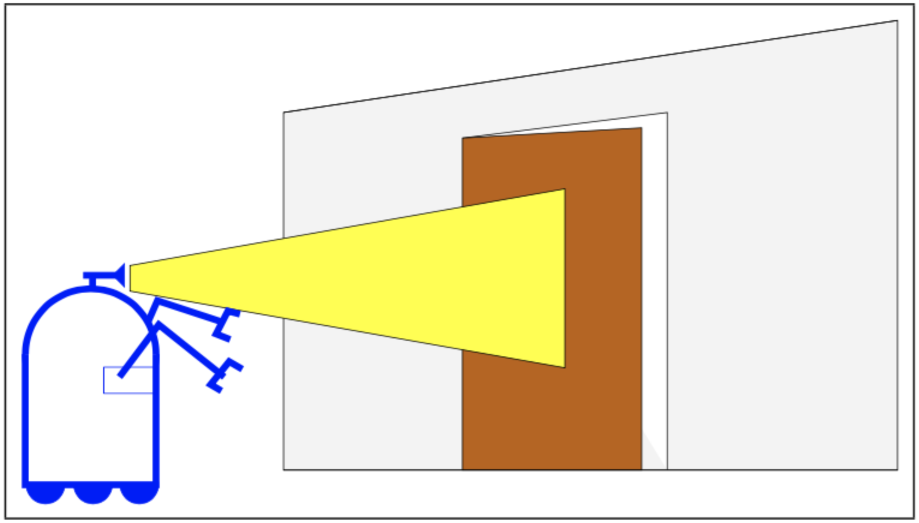

Figure 2.2 A mobile robot estimating the state of a door.

Measurement Model

$$\boxed{p(Z_t = \textbf{sense\_open} \mid X_t = \textbf{is\_open}) = 0.6}$$
$$p(Z_t = \textbf{sense\_closed} \mid X_t = \textbf{is\_open}) = 0.4$$

$$\boxed{p(Z_t = \textbf{sense\_open} \mid X_t = \textbf{is\_closed}) = 0.2}$$
$$p(Z_t = \textbf{sense\_closed} \mid X_t = \textbf{is\_closed}) = 0.8$$

Transition Model for do_nothing

$$p(X_t = \textbf{is\_open} \mid U_t = \textbf{do\_nothing}, X_{t\_1} = \textbf{is\_open}) = 1$$
$$p(X_t = \textbf{is\_closed} \mid U_t = \textbf{do\_nothing}, X_{t\_1} = \textbf{is\_open}) = 0$$
$$p(X_t = \textbf{is\_open} \mid U_t = \textbf{do\_nothing}, X_{t\_1} = \textbf{is\_closed}) = 0$$
$$p(X_t = \textbf{is\_closed} \mid U_t = \textbf{do\_nothing}, X_{t\_1} = \textbf{is\_closed}) = 1$$

$\boldsymbol{bel(x_t) = (0.75, 0.25)}$ **Don't forget normalization!**

The Bayes Filter

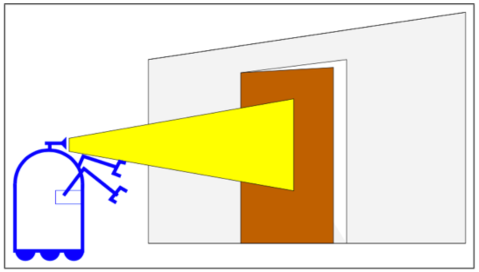

Figure 2.2 A mobile robot estimating the state of a door.

- Recursive filter for estimating x_t only from x_{t-1}, z_t and u_t and not from the ever-growing history $z_{1:t}, u_{1:t}$

Algorithm Bayes_filter($bel(x_{t-1}), u_t, z_t$): Transition/Dynamics model
  for all x_t do
    $\overline{bel}(x_t) = \int \boxed{p(x_t \mid u_t, x_{t-1})}\, bel(x_{t-1})\, dx$ Predict Step
    $bel(x_t) = \eta\, \boxed{p(z_t \mid x_t)}\, \overline{bel}(x_t)$ Update Step
  endfor
  return $bel(x_t)$

Measurement Model

Gaussian Filters - Kalman Filter

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

$$p(x) \quad = \quad \det\left(2\pi\Sigma\right)^{-\frac{1}{2}} \; \exp\left\{-\tfrac{1}{2}(x-\mu)^T \Sigma^{-1} (x-\mu)\right\}$$

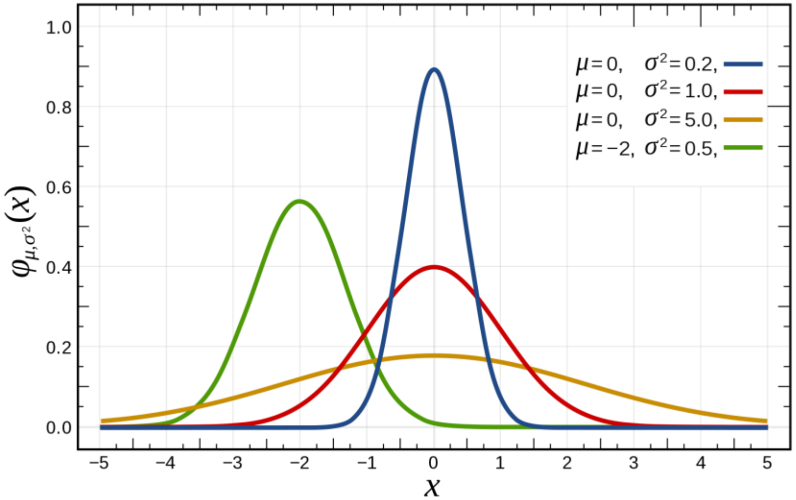

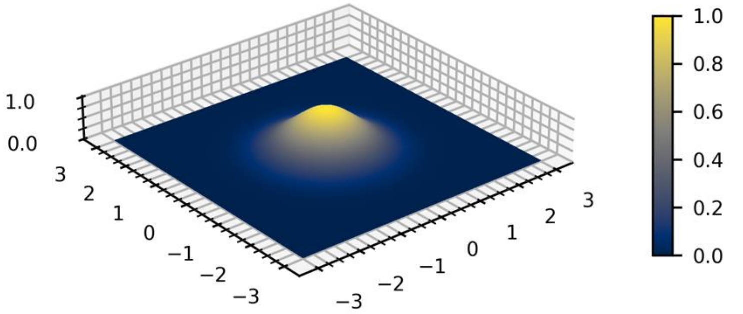

Kalman Filter

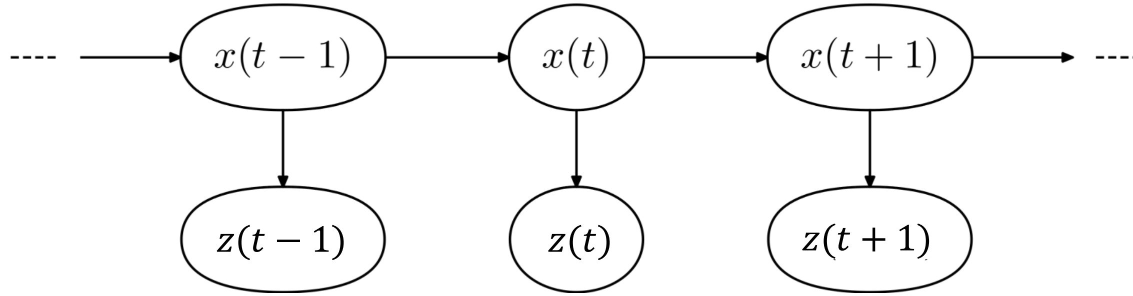

- Gaussian Belief
- Linear Transition Model

$$x_t = A_t x_{t-1} + B_t u_t + \varepsilon_t \qquad x_t = \begin{pmatrix} x_{1,t} \\ x_{2,t} \\ \vdots \\ x_{n,t} \end{pmatrix} \qquad u_t = \begin{pmatrix} u_{1,t} \\ u_{2,t} \\ \vdots \\ u_{m,t} \end{pmatrix}$$

Process Noise $\boldsymbol{\varepsilon \sim N(0, R)}$

- Linear Measurement Model

$$z_t = C_t x_t + \delta_t$$

Measurement Noise $\boldsymbol{\delta \sim N(0, Q)}$

Kalman Filter

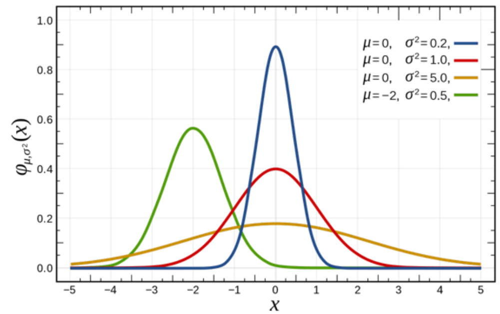

- Initial Belief $\mathbf{x_0} \sim N(\boldsymbol{\mu_0}, \boldsymbol{\Sigma_0})$

$$bel(x_0) \quad = \quad p(x_0) \quad = \quad \det(2\pi\Sigma_0)^{-\frac{1}{2}} \; \exp\left\{-\tfrac{1}{2}\,(x_0 - \mu_0)^T \Sigma_0^{-1} (x_0 - \mu_0)\right\}$$

- Distribution over next state

$$p(x_t \mid u_t, x_{t-1}) = \det(2\pi R_t)^{-\frac{1}{2}} \; \exp\left\{-\tfrac{1}{2}(x_t - \underline{A_t x_{t-1} - B_t u_t})^T R_t^{-1} (x_t - \underline{A_t x_{t-1} - B_t u_t})\right\}$$

Process Noise

Transition Model

- Likelihood of Measurement

Measurement Noise

$$p(z_t \mid x_t) \quad = \quad \det(2\pi Q_t)^{-\frac{1}{2}} \; \exp\left\{-\tfrac{1}{2}(z_t - \underline{C_t \; x_t})^T Q_t^{-1} (z_t - \underline{C_t \; x_t})\right\}$$

Measurement Mode

The Kalman Filter Algorithm

1: **Algorithm Kalman_filter**($\mu_{t-1}, \Sigma_{t-1}, u_t, z_t$):
2: $\bar{\mu}_t = A_t \, \mu_{t-1} + B_t \, u_t$
3: $\bar{\Sigma}_t = A_t \, \Sigma_{t-1} \, A_t^T + R_t$
4: $K_t = \bar{\Sigma}_t \, C_t^T (C_t \, \bar{\Sigma}_t \, C_t^T + Q_t)^{-1}$
5: $\mu_t = \bar{\mu}_t + K_t(z_t - C_t \, \bar{\mu}_t)$
6: $\Sigma_t = (I - K_t \, C_t) \, \bar{\Sigma}_t$
7: return μ_t, Σ_t

Uncertainty increases

K = Kalman Gain $K \approx \frac{R}{Q}$

Uncertainty decreases

1: **Algorithm Bayes_filter**($bel(x_{t-1}), u_t, z_t$):
2: for all x_t do
3: $\overline{bel}(x_t) = \int p(x_t \mid u_t, x_{t-1}) \, bel(x_{t-1}) \, dx$ Predict Step
4: $bel(x_t) = \eta \, p(z_t \mid x_t) \, \overline{bel}(x_t)$ Update Step
5: endfor
6: return $bel(x_t)$

If R large, then K is large. Update dominated by innovation.

If Q large, then K is small. Update dominated by prediction.

Example

$\boldsymbol{p(z_0|x_0)}$

Measurement

$\boldsymbol{p(x_0)}$

$\boldsymbol{bel(x_0)}$

After Update

$\overline{\boldsymbol{bel}}\,\boldsymbol{(x_1)}$

After Prediction

$\boldsymbol{p(z_1|x_1)}$

Measurement

$\boldsymbol{bel(x_1)}$

After Update

Propagating a Gaussian through a Linear Model

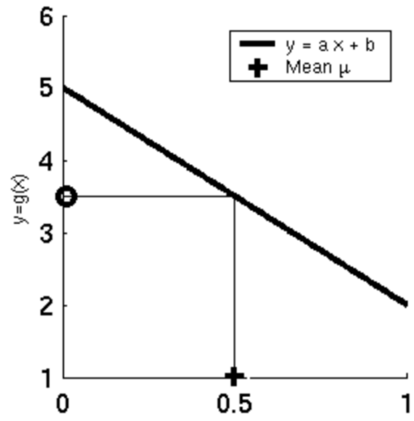

Propagating a Gaussian through a Non-Linear Model

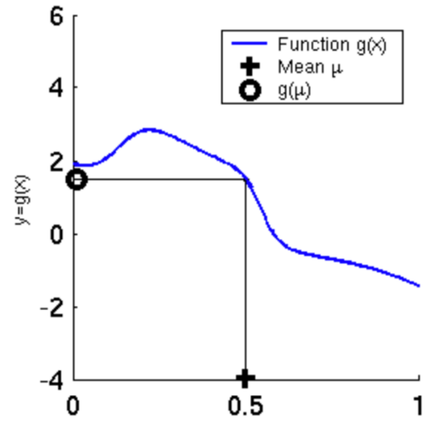

Linearizing the Non-Linear Model

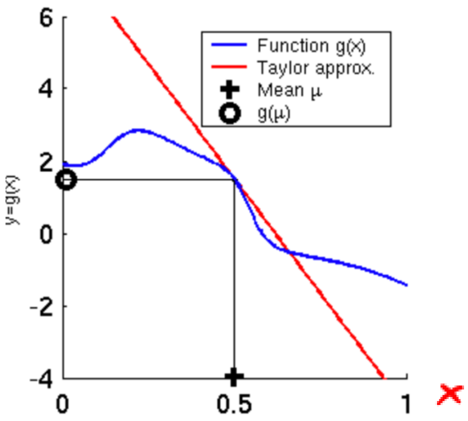

Representations for Manipulation

Manuel Wühtrich et al. “Probabilistic Object Tracking using a Depth Camera”, IROS 2013

Extended Kalman filter - Process Model

$$\begin{aligned} x_t &= g(u_t, x_{t-1}) + \varepsilon_t \\ z_t &= h(x_t) + \delta_t \; . \end{aligned}$$

Process Model

Measurement Model

First order Taylor Expansion – linear approximation around value and slope

$$g'(u_t, x_{t-1}) \quad := \quad \frac{\partial g(u_t, x_{t-1})}{\partial x_{t-1}}$$

Gradient of Nonlinear function around x_{t-1}

Taylor Expansion

$$\begin{aligned} g(u_t, x_{t-1}) &\approx g(u_t, \mu_{t-1}) + \underbrace{g'(u_t, \mu_{t-1})}_{=:\, G_t} (x_{t-1} - \mu_{t-1}) \\ &= g(u_t, \mu_{t-1}) + \underline{G_t \, (x_{t-1} - \mu_{t-1})} \end{aligned}$$

Jacobian

Extended Kalman filter - Process Model

$$\begin{aligned} g(u_t, x_{t-1}) &\approx g(u_t, \mu_{t-1}) + \underbrace{g'(u_t, \mu_{t-1})}_{=:\, G_t} (x_{t-1} - \mu_{t-1}) \\ &= g(u_t, \mu_{t-1}) + G_t \, (x_{t-1} - \mu_{t-1}) \end{aligned}$$

Same equations as in previous slide

Written as Gaussian:

$$\begin{aligned} & p(x_t \mid u_t, x_{t-1}) \\ \approx \quad & \det(2\pi R_t)^{-\frac{1}{2}} \exp\left\{-\tfrac{1}{2} \left[x_t - g(u_t, \mu_{t-1}) - G_t \, (x_{t-1} - \mu_{t-1})\right]^T \right. \\ & \left. R_t^{-1} \left[x_t - g(u_t, \mu_{t-1}) - G_t \, (x_{t-1} - \mu_{t-1})\right]\right\} \end{aligned}$$

Extended Kalman Filter – Measurement Model

$$
\begin{aligned}
x_t &= g(u_t, x_{t-1}) + \varepsilon_t \\
z_t &= h(x_t) + \delta_t \,.
\end{aligned}
$$

Process Model

Measurement Model

First order Taylor Expansion – linear approximation around value and slope

$$
\begin{aligned}
h(x_t) &\approx h(\bar{\mu}_t) + \underbrace{h'(\bar{\mu}_t)}_{=:\, H_t} (x_t - \bar{\mu}_t) \\
&= h(\bar{\mu}_t) + \underline{H_t\, (x_t - \bar{\mu}_t)}
\end{aligned}
$$

Jacobian

Written as Gaussian:

$$
\begin{aligned}
p(z_t \mid x_t) = \; & \det(2\pi Q_t)^{-\frac{1}{2}} \exp\left\{-\tfrac{1}{2}\left[z_t - h(\bar{\mu}_t) - H_t\,(x_t - \bar{\mu}_t)\right]^T \right. \\
& \left. Q_t^{-1} \left[z_t - h(\bar{\mu}_t) - H_t\,(x_t - \bar{\mu}_t)\right]\right\}
\end{aligned}
$$

The Extended Kalman Filter Algorithm

1: **Algorithm Extended_Kalman_filter**($\mu_{t-1}, \Sigma_{t-1}, u_t, z_t$):
2: $\bar{\mu}_t = g(u_t, \mu_{t-1})$
3: $\bar{\Sigma}_t = G_t \, \Sigma_{t-1} \, G_t^T + R_t$
4: $K_t = \bar{\Sigma}_t \, H_t^T (H_t \, \bar{\Sigma}_t \, H_t^T + Q_t)^{-1}$
5: $\mu_t = \bar{\mu}_t + K_t (z_t - h(\bar{\mu}_t))$
6: $\Sigma_t = (I - K_t \, H_t) \, \bar{\Sigma}_t$
7: return μ_t, Σ_t

Predict

Update

	Kalman filter	EKF
state prediction (Line 2)	$A_t \, \mu_{t-1} + B_t \, u_t$	$g(u_t, \mu_{t-1})$
measurement prediction (Line 5)	$C_t \, \bar{\mu}_t$	$h(\bar{\mu}_t)$

CS231

Introduction to Computer Vision

Next lecture:

Optimal Estimation cont'