Robotics
Spring 2023

Abhishek Gupta

TAs: Yi Li, Srivatsa GS

Recap: Course Overview

- Filtering/Smoothing
- Localization
- Mapping
- SLAM

- Search
- Motion Planning
- TrajOpt
- Stability/Certification

- MDPs and RL
- Imitation Learning
- Solving POMDPs

Lecture Outline

Recap

↓

Bayesian Filtering

↓

Gaussian Properties

↓

Kalman Filtering

Recap: Bayes Rule and Recursive Bayesian Updates

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

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

Recursive update to integrate in multiple measurements

$$P(x|z_1,\ldots,z_n) = \eta_{1:n} \prod_{i=1,\ldots,n} P(z_i|x)P(x)$$

Recap: Bayes Rule and Recursive Bayesian Updates

$$P(x|z_1,\ldots,z_n) = \frac{P(z_n|x,z_1,\ldots,z_{n-1})P(x|z_1,\ldots,z_{n-1})}{P(z_n|z_1,\ldots,z_{n-1})}$$

Markov assumption: z_n is conditionally independent of $z_1,\ldots,z_{n-1}$ given x.

$$p(z_n|x,z_1,\ldots,z_{n-1}) = p(z_n|x)$$

$$P(x|z_1,\ldots,z_n) = \frac{P(z_n|x)P(x|z_1,\ldots,z_{n-1})}{P(z_n|z_1,\ldots,z_{n-1})}$$

$$= \eta P(z_n|x)P(x|z_1,\ldots,z_{n-1})$$

$$= \eta_{1:n} \prod_{i=1,\ldots,n} P(z_i|x)P(x)$$

Lecture Outline

Recap

↓

Bayesian Filtering

↓

Gaussian Properties

↓

Kalman Filtering

Let’s estimate “state” of our robot

- What affects uncertainty:
  - Robot actions (increase uncertainty typically)
  - Sensor measurements (decrease uncertainty typically)

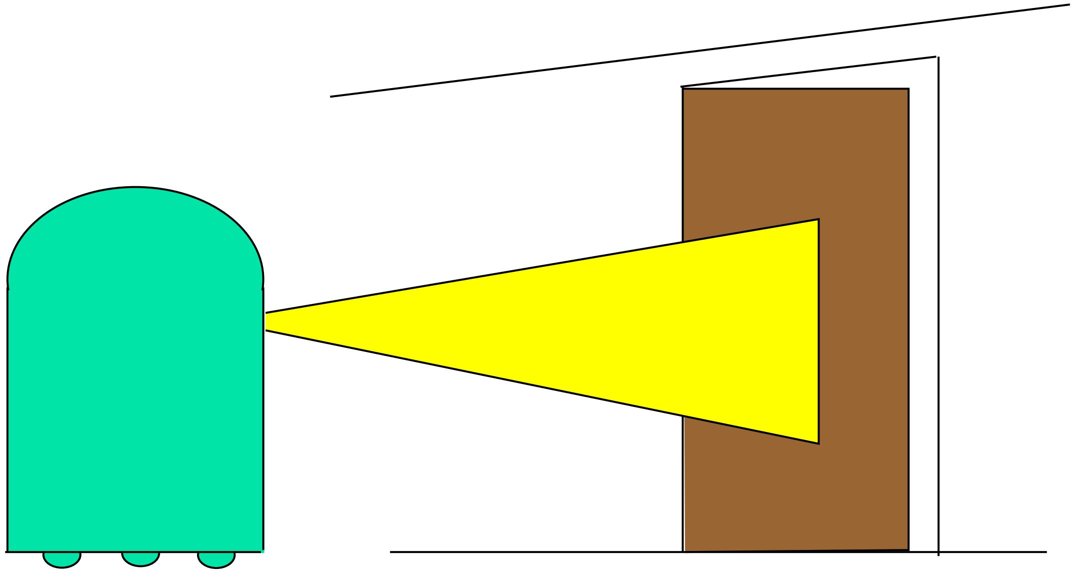

Bayes Filters: Framework

- **Given:**
  - Stream of observations ***z*** and action data ***u:***

$$d_t = \{z_0, u_0, z_1, u_1, \ldots, z_t\}$$

  - Sensor model ***P(z|x).***
  - Action model ***P(x'| u, x)***
  - Prior probability of the initial system state ***P(x).***
- **Wanted:**
  - Estimate of the state ***X*** of a dynamical system.
  - The posterior of the state is also called **Belief**:

$$Bel(x_t) = P(x_t | u_{0:t-1}, z_{0:t})$$

How do actions increase uncertainty?

- Actions transition the state of the system forward x → x’
  - But they may (and usually) do so with errors/noise!
- Robot wheels have slippage/noise, joints have stochasticity, environment introduces noise

How do sensors reduce uncertainty?

- Measurements usually convey more information about the state of the world
- Sensor readings can range from images to laser scans to tactile sensing, each of which has a different effect on uncertainty

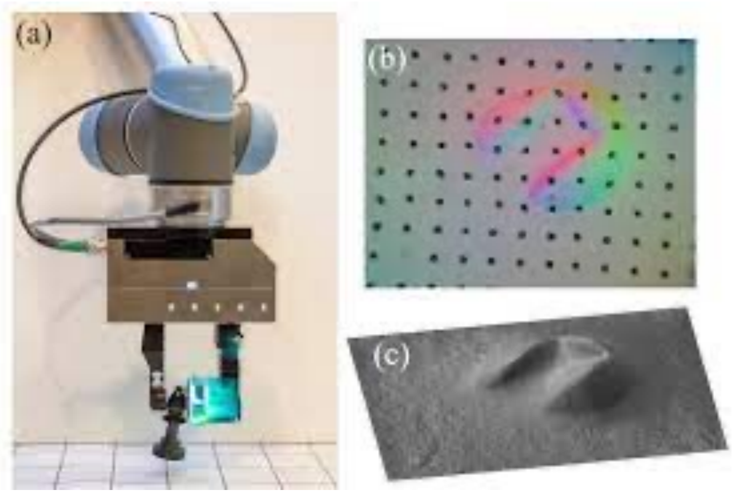

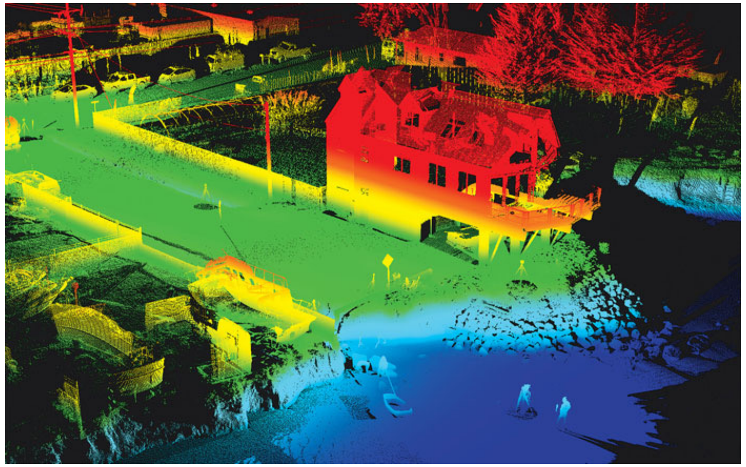

Filtering

- Filtering is the process of making sense ("filtering") of sensor measurements and actions to estimate the system state
- Many different types of filters:
  - Matched filters (known signal)
  - Wiener filters (signal from noise)
  - Bayesian filters (bayesian state estimation)
    - Kalman
    - EKF / UKF
    - ….

Example Situation for Filtering

"Where is my robot?"

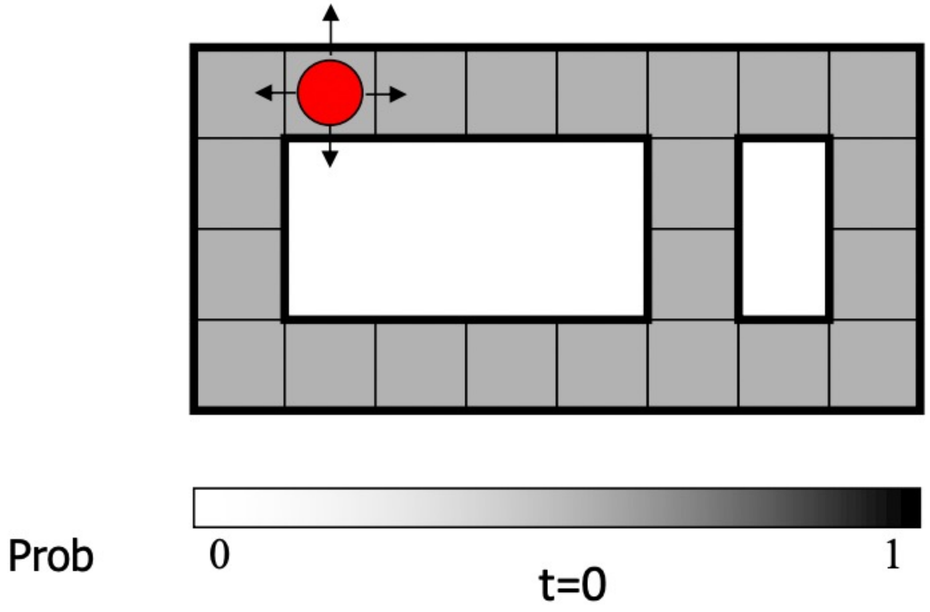

- Sensor model: never more than 1 mistake
- Know the heading (North, East, South or West)
- Motion model: may not execute action with small prob.

Markov Assumption

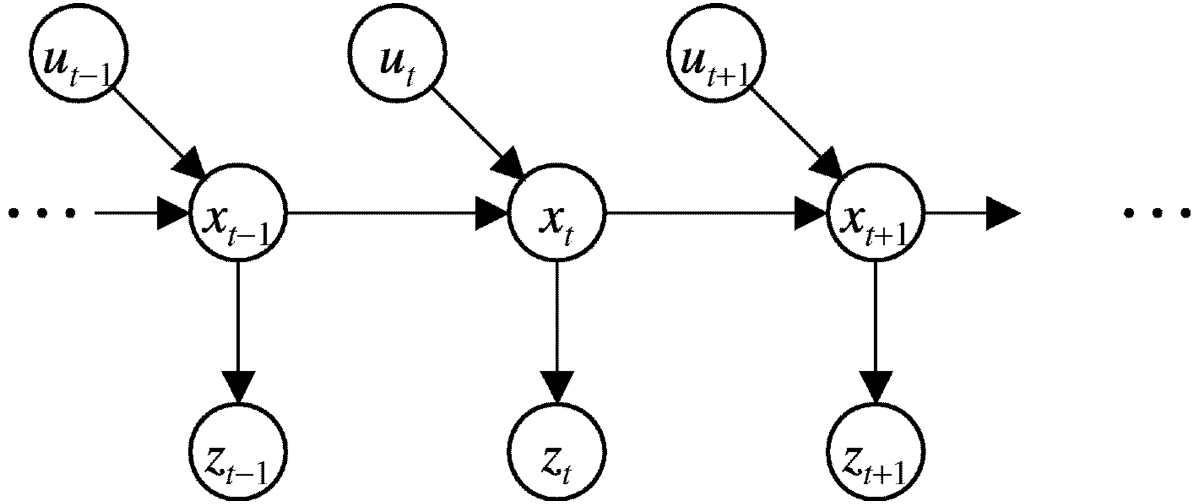

$$p(x_t \mid z_{0:t-1}, u_{0:t-1}, x_{0:t-1}) = p(x_t \mid x_{t-1}, u_{t-1})$$
$$p(z_t \mid x_{0:t}, u_{0:t-1}, z_{0:t-1}) = p(z_t \mid x_t)$$

Underlying Assumptions

- Static world
- Independent noise
- Perfect model, no approximation errors

Bayes Filters

z = observation
u = action
x = state

$$Bel(x_t) = P(x_t|u_{0:t-1}, z_{0:t})$$

We want to recursively express $Bel(x_t)$ in terms of three entities

$$p(z_t|x_t)$$

Measurement

$$p(x_t|x_{t-1}, u_{t-1})$$

Dynamics

$$Bel(x_{t-1})$$

Previous Belief

Bayes Filters: Intuition

z = observation
u = action
x = state

$$Bel(x_t) = P(x_t | u_{0:t-1}, z_{0:t})$$

We want to recursively express Bel(x_t) in terms of three entities

Integrate in effect of action

$$Bel(x_{t-1}) \quad + \quad p(x_t | x_{t-1}, u_{t-1}) \longrightarrow \overline{Bel}(x_t)$$

Previous Belief

Dynamics

With integration → understand the effect of taking an action

Bayes Filters: Intuition

z = observation
u = action
x = state

$$Bel(x_t) = P(x_t | u_{0:t-1}, z_{0:t})$$

We want to recursively express Bel(x_t) in terms of three entities

Integrate in Measurement

$$\overline{Bel}(x_t) \quad + \quad p(z_t | x_t) \longrightarrow Bel(x_t)$$

Previous Belief

Measurement

With normalization → understand the effect of your latest measurement

Bayes Filters

z = observation
u = action
x = state

$$Bel(x_t) = P(x_t | u_{0:t-1}, z_{0:t})$$

Bayes $$= \eta \; p(z_t | x_t, u_{0:t-1}, z_{0:t-1}) P(x_t | u_{0:t-1}, z_{0:t-1})$$

Remember: Bayes Rule

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

$$\eta = \frac{1}{\sum_x P(y, x)}$$

$$P(x|y) = \eta P(y, x)$$

Bayes Filters

z = observation
u = action
x = state

$$Bel(x_t) = P(x_t | u_{0:t-1}, z_{0:t})$$

Bayes $$= \eta \, p(z_t | x_t, u_{0:t-1}, z_{0:t-1}) P(x_t | u_{0:t-1}, z_{0:t-1})$$

Markov $$= \eta \, p(z_t | x_t) P(x_t | u_{0:t-1}, z_{0:t-1})$$

Remember: Markov Property

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

Bayes Filters

z = observation
u = action
x = state

$$Bel(x_t) = P(x_t | u_{0:t-1}, z_{0:t})$$

Bayes $$= \eta \, p(z_t | x_t, u_{0:t-1}, z_{0:t-1}) P(x_t | u_{0:t-1}, z_{0:t-1})$$

Markov $$= \eta \, p(z_t | x_t) P(x_t | u_{0:t-1}, z_{0:t-1})$$

Total prob.

$$= \eta \, p(z_t | x_t) \int P(x_t | u_{0:t-1}, z_{0:t-1}, x_{t-1}) P(x_{t-1} | u_{0:t-1}, z_{0:t-1}) dx_{t-1}$$

Remember: Marginalization

$$p(x) = \int p(x, y) dy$$

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

Bayes Filters

z = observation
u = action
x = state

$$Bel(x_t) = P(x_t | u_{0:t-1}, z_{0:t})$$

Bayes $$= \eta \, p(z_t | x_t, u_{0:t-1}, z_{0:t-1}) P(x_t | u_{0:t-1}, z_{0:t-1})$$

Markov $$= \eta \, p(z_t | x_t) P(x_t | u_{0:t-1}, z_{0:t-1})$$

Total prob.

$$= \eta \, p(z_t | x_t) \int P(x_t | u_{0:t-1}, z_{0:t-1}, x_{t-1}) P(x_{t-1} | u_{0:t-1}, z_{0:t-1}) dx_{t-1}$$

Markov $$= \eta \, p(z_t | x_t) \int P(x_t | u_{t-1}, x_{t-1}) P(x_{t-1} | u_{0:t-1}, z_{0:t-1}) dx_{t-1}$$

$$= \eta \, p(z_t | x_t) \int P(x_t | u_{t-1}, x_{t-1}) Bel(x_{t-1}) dx_{t-1}$$

Understanding Bayes Filters

z = observation
u = action
x = state

$$Bel(x_t) = P(x_t | u_{0:t-1}, z_{0:t})$$
$$= \eta \, p(z_t | x_t) \int P(x_t | u_{t-1}, x_{t-1}) Bel(x_{t-1}) dx_{t-1}$$

↑

Step 1: Dynamics Update

Incorporate the effect of motion on uncertainty (typically increases)

Understanding Bayes Filters

z = observation
u = action
x = state

$$Bel(x_t) = P(x_t|u_{0:t-1}, z_{0:t})$$
$$= \eta \, p(z_t|x_t) \int P(x_t|u_{t-1}, x_{t-1}) Bel(x_{t-1}) dx_{t-1}$$

Step 2: Measurement Update

Incorporate the effect of new measurements on uncertainty (typically decreases)

Understanding Bayes Filters

z = observation
u = action
x = state

$$Bel(x_t) = P(x_t|u_{0:t-1}, z_{0:t})$$
$$= \eta\, p(z_t|x_t) \int P(x_t|u_{t-1}, x_{t-1}) Bel(x_{t-1}) dx_{t-1}$$

All Bayes filter iterate between performing the dynamics (prediction) step and the measurement (correction) step

Bayes Filter Algorithm

Initial Prior

$p(x_0)$

Dynamics/Prediction

Estimate $\overline{Bel}(x_t)$

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

Measurement/Correction

Estimate $Bel(x_t)$

$$Bel(x_t) = \eta p(z_t | x_t) \overline{Bel}(x_t) = \frac{p(z_t | x_t) \overline{Bel}(x_t)}{\sum_{x'} p(z_t | x_t = x') \overline{Bel}(x_t = x')}$$

Example Run for Localization

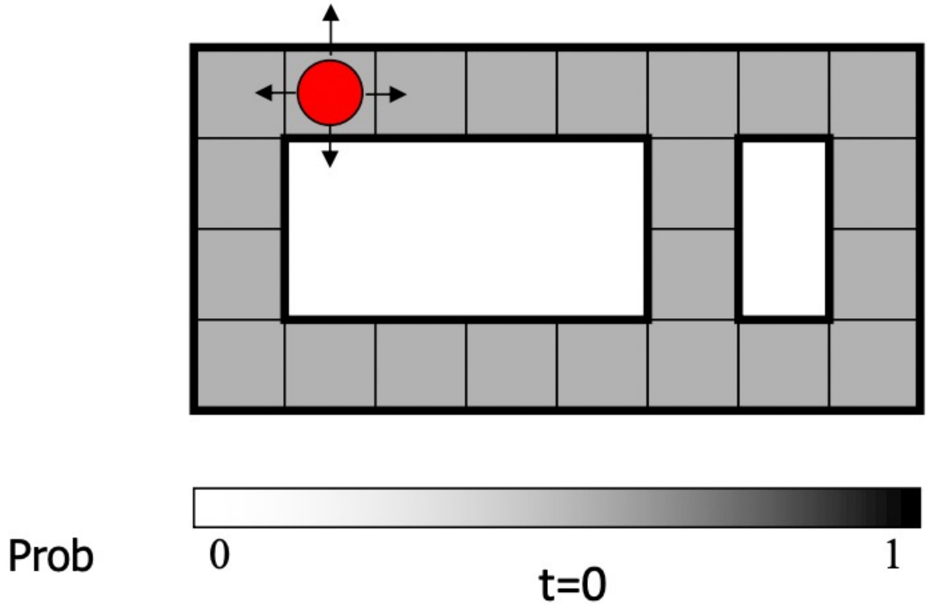

- Sensor model: never more than 1 mistake
- Know the heading (North, East, South or West)
- Motion model: may not execute action with small prob.

Example Run for Localization

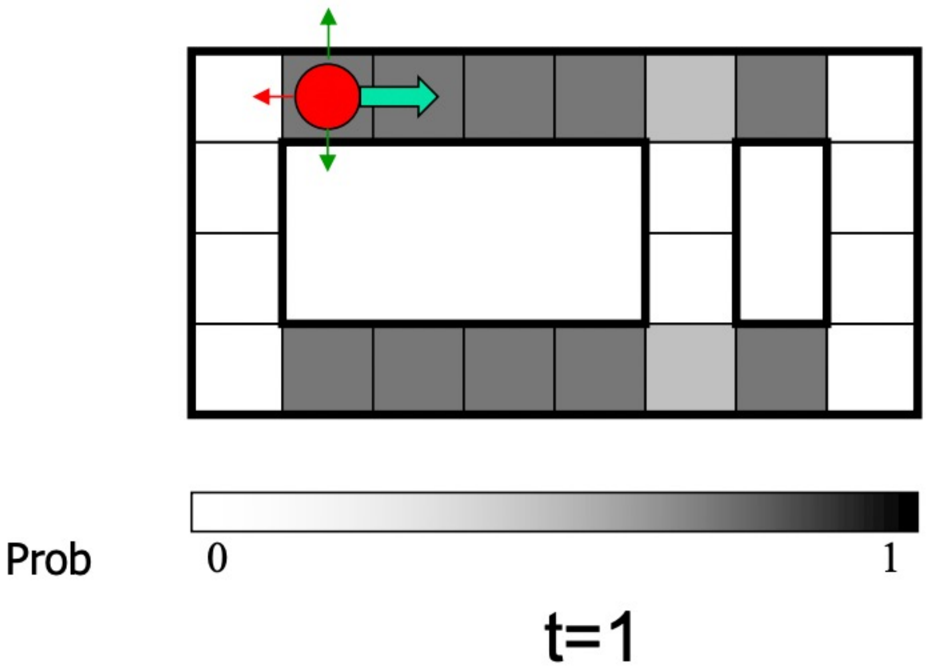

- Sensor model: never more than 1 mistake
- Know the heading (North, East, South or West)
- Motion model: may not execute action with small prob.

Example Run for Localization

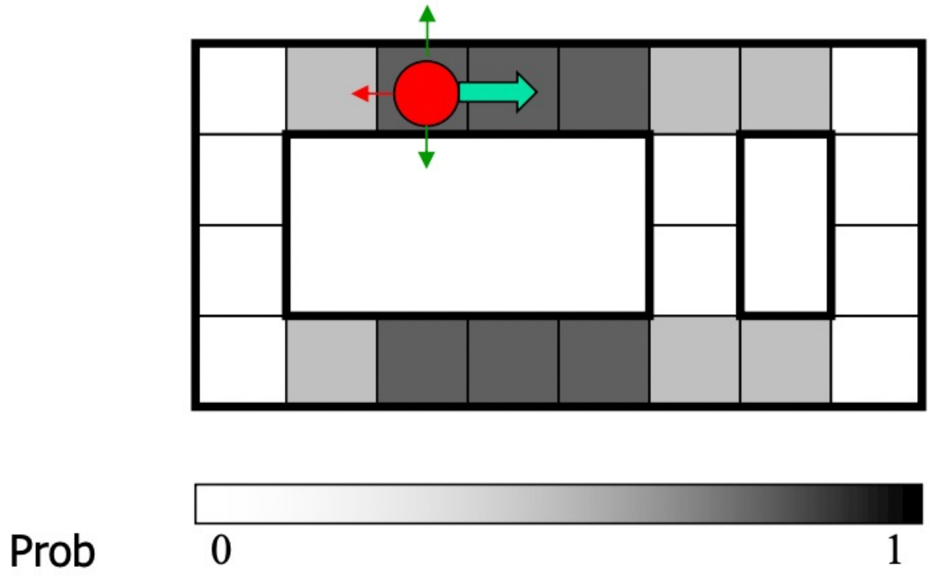

t=2

- Sensor model: never more than 1 mistake
- Know the heading (North, East, South or West)
- Motion model: may not execute action with small prob.

Example Run for Localization

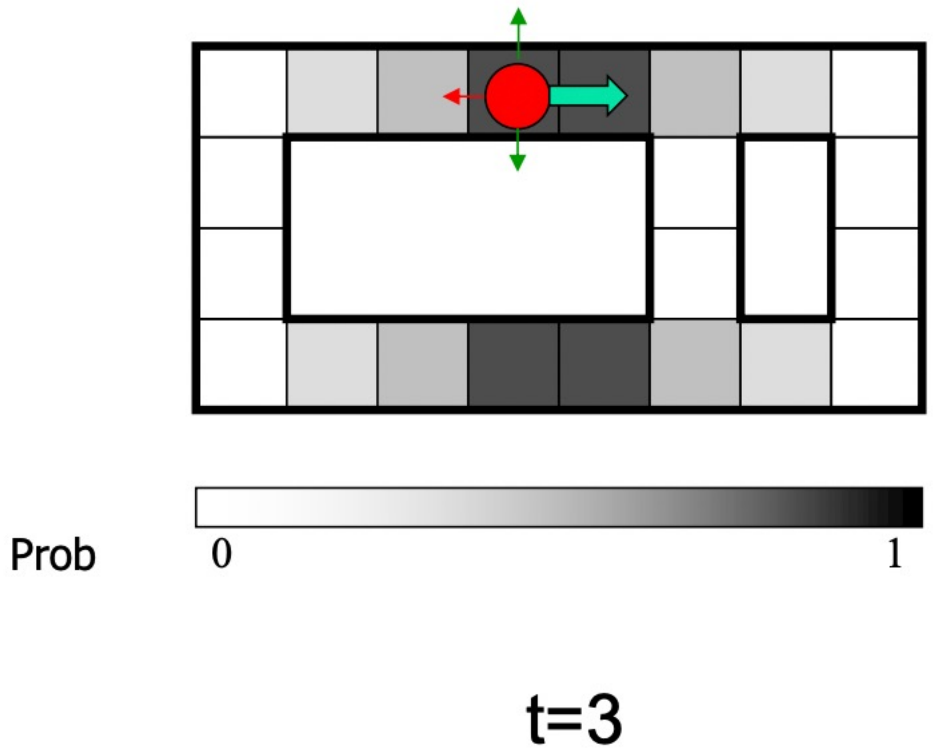

- Sensor model: never more than 1 mistake
- Know the heading (North, East, South or West)
- Motion model: may not execute action with small prob.

Example Run for Localization

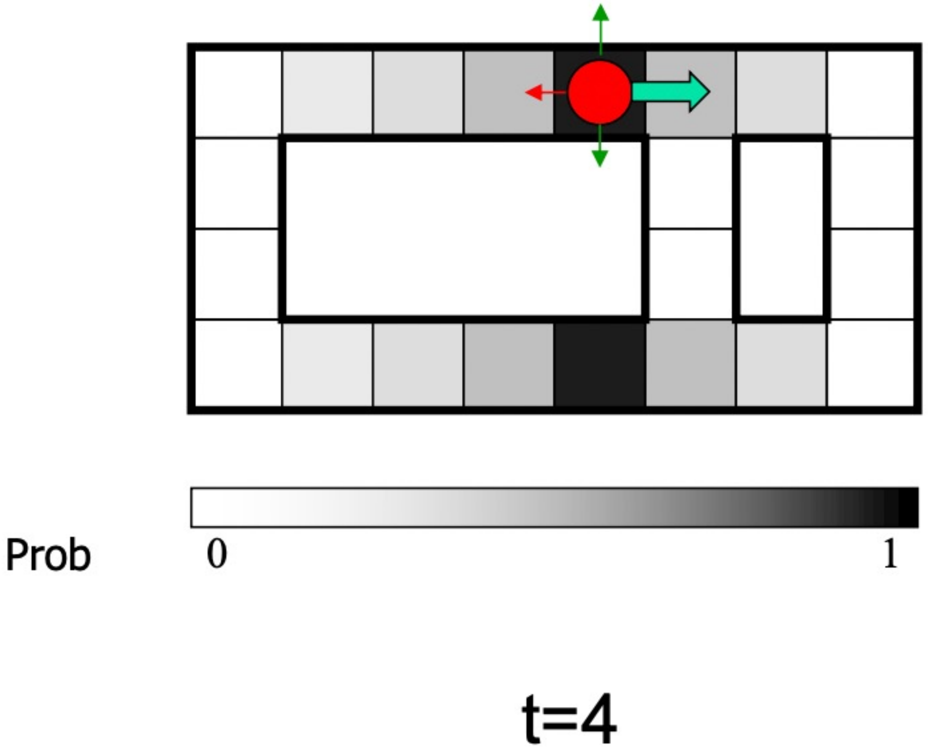

- Sensor model: never more than 1 mistake
- Know the heading (North, East, South or West)
- Motion model: may not execute action with small prob.

Example Run for Localization

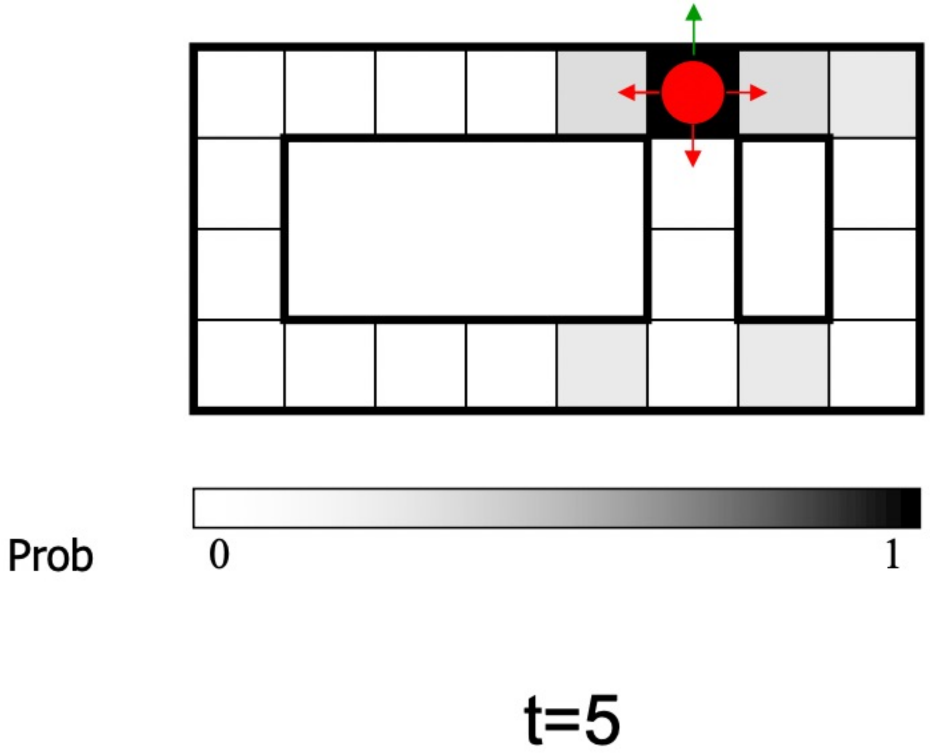

- Sensor model: never more than 1 mistake
- Know the heading (North, East, South or West)
- Motion model: may not execute action with small prob.

Bayes Filters are Familiar!

$$Bel(x_t) = \eta \; P(z_t \mid x_t) \int P(x_t \mid u_t, x_{t-1}) \; Bel(x_{t-1}) \; dx_{t-1}$$

- Kalman filters
- Particle filters
- Hidden Markov models
- Dynamic Bayesian networks
- Partially Observable Markov Decision Processes (POMDPs)

Why is this difficult?

$$Bel(x_t) = \eta \ P(z_t \mid x_t) \int P(x_t \mid u_t, x_{t-1}) \ Bel(x_{t-1}) \ dx_{t-1}$$

Tractable Bayesian inference is challenging in the general case

We will work out the conjugate prior and discrete case,
leaving the MCMC/VI cases as an exercise

Lecture Outline

Recap

↓

Bayesian Filtering

↓

Gaussian Properties

↓

Kalman Filtering

Recap: Bayesian Filters

$$Bel(x_t) = P(x_t|u_{0:t-1}, z_{0:t})$$
$$= \eta \, p(z_t|x_t) \int P(x_t|u_{t-1}, x_{t-1}) Bel(x_{t-1}) dx_{t-1}$$

- Dynamics (Prediction)

$$\overline{Bel}(x_t) = \int P(x_t|u_{t-1}, x_{t-1}) Bel(x_{t-1}) dx_{t-1}$$

- Measurement (Correction)

$$Bel(x_t) = \eta P(z_t|x_t) \overline{Bel}(x_t)$$

What makes this challenging?

- Dynamics (Prediction)

$$\overline{Bel}(x_t) = \int P(x_t|u_{t-1}, x_{t-1}) Bel(x_{t-1}) dx_{t-1}$$

- Measurement (Correction)

$$Bel(x_t) = \eta P(z_t|x_t)\overline{Bel}(x_t)$$

Tractable computation of Bayesian posteriors

How can we make this more tractable?

- Dynamics (Prediction)

$$\overline{Bel}(x_t) = \int P(x_t | u_{t-1}, x_{t-1}) Bel(x_{t-1}) dx_{t-1}$$

- Measurement (Correction)

$$Bel(x_t) = \eta P(z_t | x_t) \overline{Bel}(x_t)$$

Model as Linear Gaussian

Let’s take a little Gaussian detour

Gaussians (1D)

- Gaussian with mean (μ) and standard deviation (σ)

$$X \sim \mathcal{N}(\mu, \sigma^2)$$

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

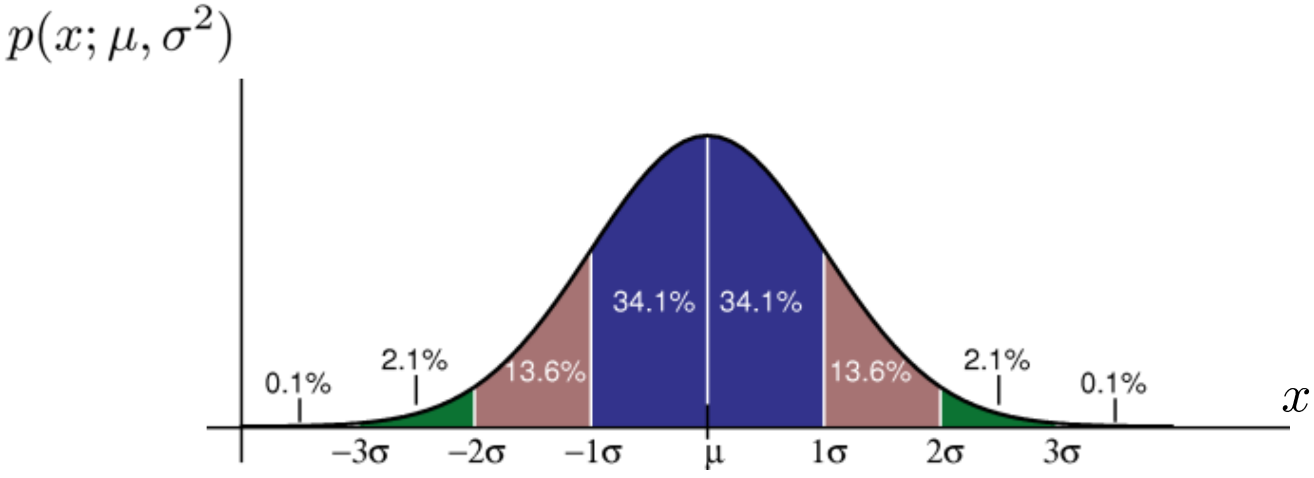

Gaussians (2D)

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

$$\mathbf{x} = \begin{pmatrix} x_a \\ x_b \end{pmatrix}, \quad \mu = \begin{pmatrix} \mu_a \\ \mu_b \end{pmatrix}$$

$$\Sigma = \begin{pmatrix} \Sigma_{aa} & \Sigma_{ab} \\ \Sigma_{ba} & \Sigma_{bb} \end{pmatrix}$$

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

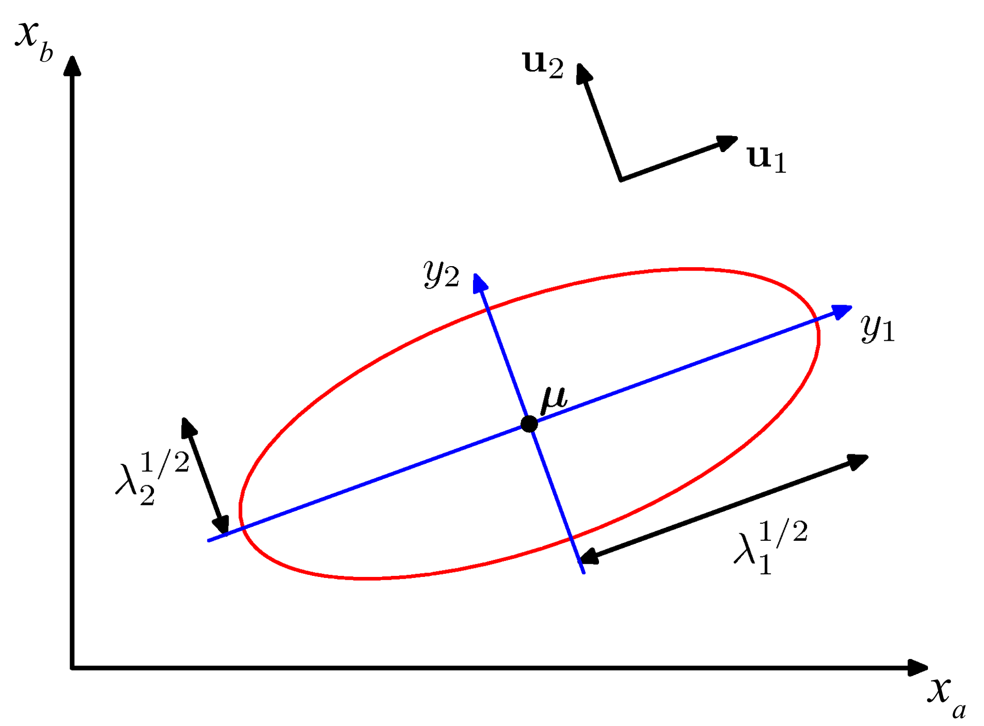

2D examples

Slide from Pieter Abbeel

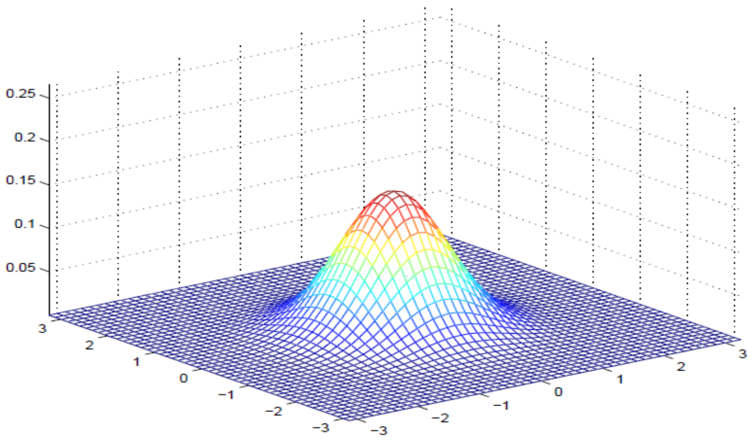

- μ = [0; 0]
- Σ = [1 0 ; 0 1]

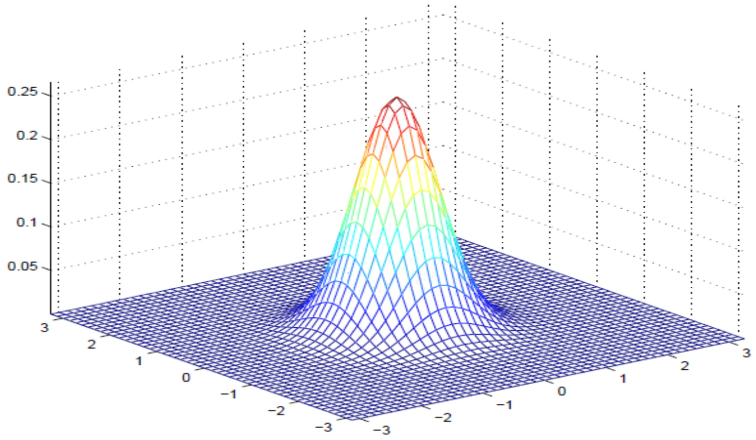

- μ = [0; 0]
- Σ = [.6 0 ; 0 .6]

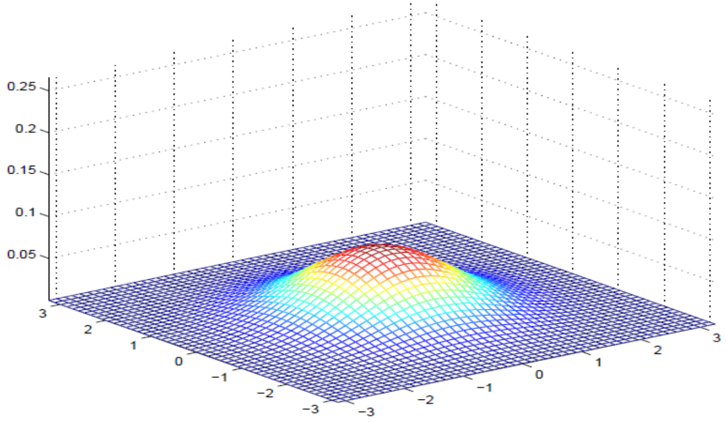

- μ = [0; 0]
- Σ = [2 0 ; 0 2]

2D examples

Slide from Pieter Abbeel

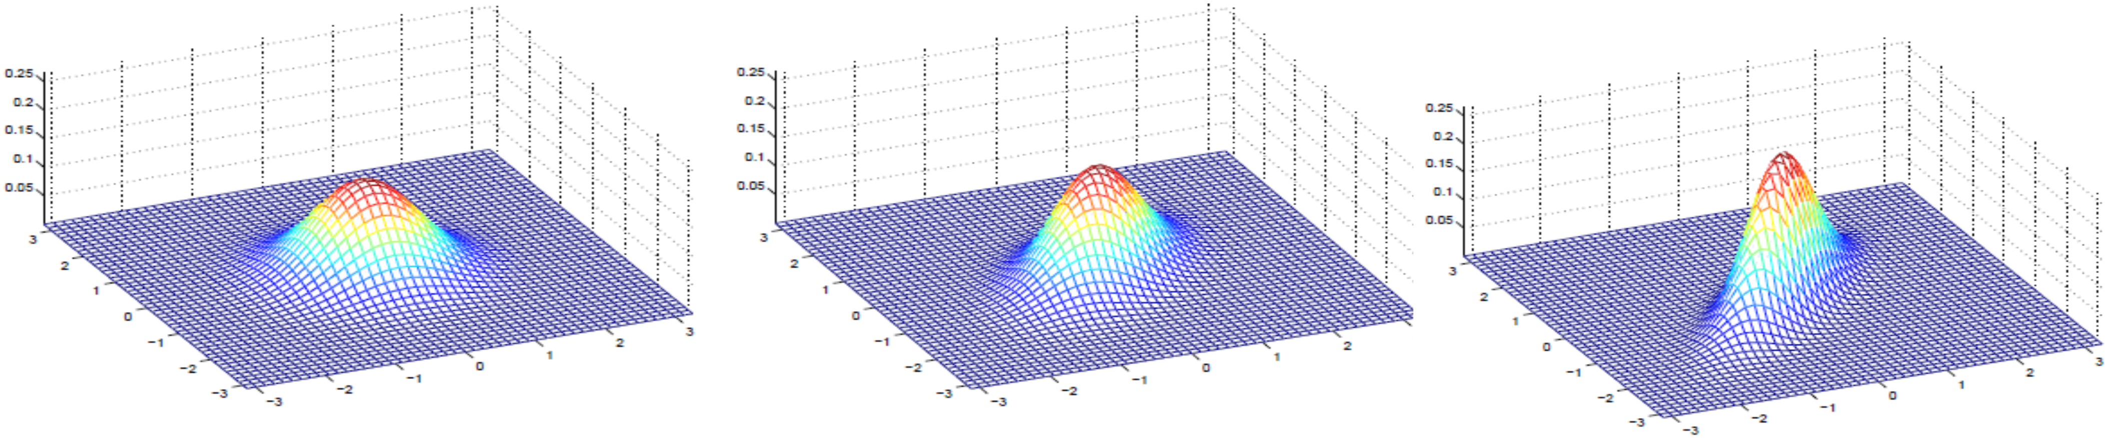

- $\mu = [0; 0]$
- $\Sigma = [1 \;\; 0; 0 \;\; 1]$

- $\mu = [0; 0]$
- $\Sigma = [1 \;\; 0.5; 0.5 \;\; 1]$

- $\mu = [0; 0]$
- $\Sigma = [1 \;\; 0.8; 0.8 \;\; 1]$

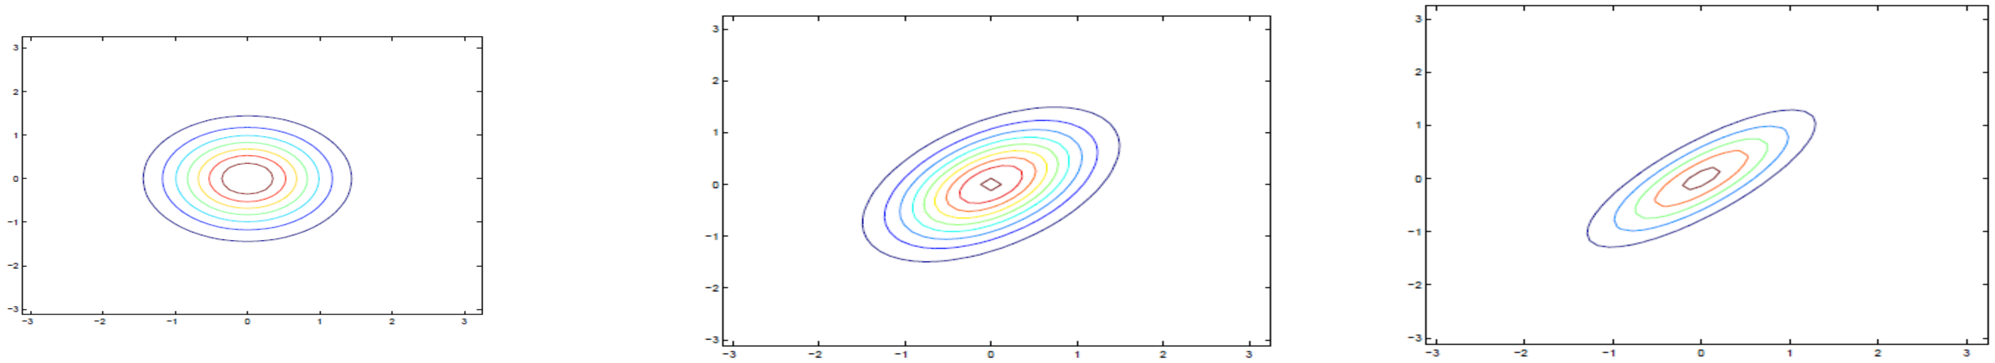

2D examples

Slide from Pieter Abbeel

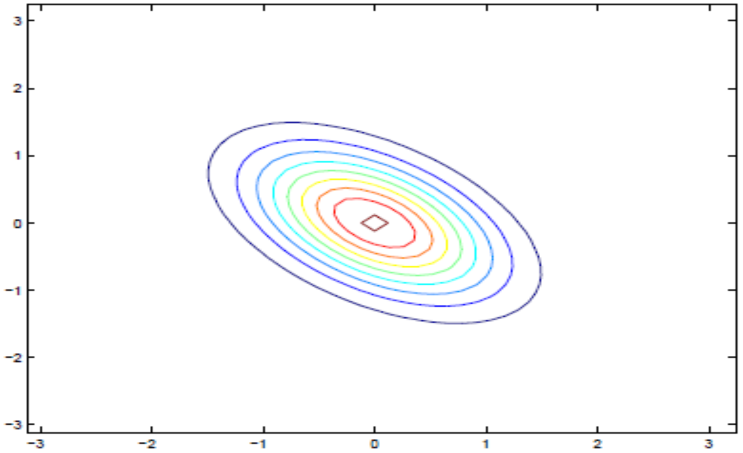

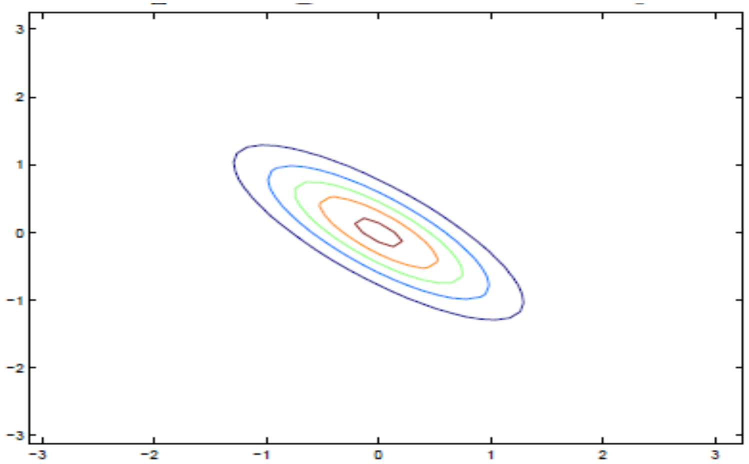

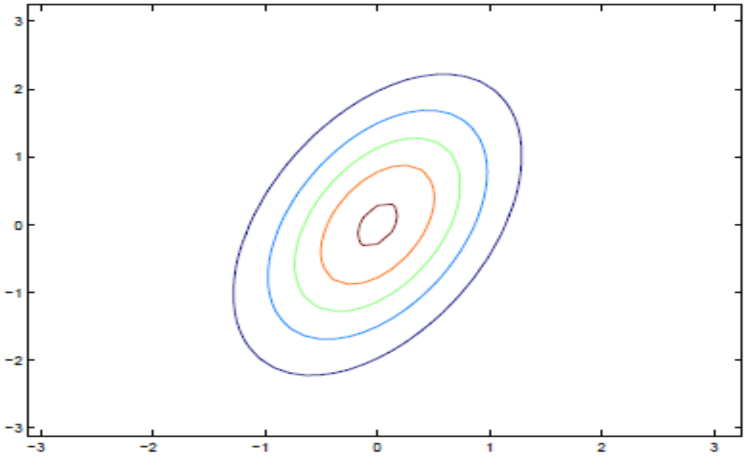

- μ = [0; 0]
- Σ = [1 -0.5 ; -0.5 1]

- μ = [0; 0]
- Σ = [1 -0.8 ; -0.8 1]

- μ = [0; 0]
- Σ = [3 0.8 ; 0.8 1]

Multivariate Gaussians

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

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

- Marginalization and conditioning in Gaussians results in Gaussians
- We stay in the "Gaussian world" as long as we start with Gaussians and perform only linear transformations.

Let's show how!

Partitioned Multivariate Gaussians

$$\mathcal{N}(\mu, \Sigma) = \mathcal{N}\left(\begin{bmatrix} \mu_X \\ \mu_Y \end{bmatrix}, \begin{bmatrix} \Sigma_{XX} & \Sigma_{XY} \\ \Sigma_{YX} & \Sigma_{YY} \end{bmatrix} \right)$$

$$p(\begin{bmatrix} x \\ y \end{bmatrix}; \mu, \Sigma) = \frac{1}{2\pi^{\frac{n}{2}} |\Sigma|^{0.5}} \exp\left(-\frac{1}{2} \left(\begin{bmatrix} x \\ y \end{bmatrix} - \begin{bmatrix} \mu_X \\ \mu_Y \end{bmatrix} \right)^T \begin{bmatrix} \Sigma_{XX} & \Sigma_{XY} \\ \Sigma_{YX} & \Sigma_{YY} \end{bmatrix}^{-1} \left(\begin{bmatrix} x \\ y \end{bmatrix} - \begin{bmatrix} \mu_X \\ \mu_Y \end{bmatrix} \right) \right)$$

Let us consider an arbitrary partitioning of a multivariate Gaussian

What is p(x) and p(x|y=y_0)?

Marginal (→ p(x))

Conditional (→ p(x|y=y_0))

Important Gaussian Identities

- Marginalization:

$$p(X) = \mathcal{N}(\mu_X, \Sigma_{XX})$$

$$p(Y) = \mathcal{N}(\mu_Y, \Sigma_{YY})$$

- Conditioning

$$p(X|Y = y_0) = \mathcal{N}(\mu_X + \Sigma_{XY}\Sigma_{YY}^{-1}(y_0 - \mu_Y), \Sigma_{XX} - \Sigma_{XY}\Sigma_{YY}^{-1}\Sigma_{YX})$$

Marginalization of Gaussian

$$p(\begin{bmatrix} x \\ y \end{bmatrix}; \mu, \Sigma) = \frac{1}{2\pi^{\frac{n}{2}} |\Sigma|^{0.5}} \exp\left(-\frac{1}{2} \left(\begin{bmatrix} x \\ y \end{bmatrix} - \begin{bmatrix} \mu_X \\ \mu_Y \end{bmatrix} \right)^T \begin{bmatrix} \Gamma_{XX} & \Gamma_{XY} \\ \Gamma_{YX} & \Gamma_{YY} \end{bmatrix} \left(\begin{bmatrix} x \\ y \end{bmatrix} - \begin{bmatrix} \mu_X \\ \mu_Y \end{bmatrix} \right) \right)$$

$$p(x) = \int p(\begin{bmatrix} x \\ y \end{bmatrix}; \mu, \Sigma) dy = \frac{1}{2\pi^{\frac{n}{2}} |\Sigma|^{0.5}} \int \exp\left(-\frac{1}{2} \left(\begin{bmatrix} x \\ y \end{bmatrix} - \begin{bmatrix} \mu_X \\ \mu_Y \end{bmatrix} \right)^T \begin{bmatrix} \Gamma_{XX} & \Gamma_{XY} \\ \Gamma_{YX} & \Gamma_{YY} \end{bmatrix} \left(\begin{bmatrix} x \\ y \end{bmatrix} - \begin{bmatrix} \mu_X \\ \mu_Y \end{bmatrix} \right) \right) dy$$

Sketch:

1. We will write out the whole likelihood in quadratic form
2. We will add and subtract terms to complete the square
3. We will use the fact that Gaussians integrate to 1 to simplify
4. We will show that the marginal is p(x)

Conditioning with Gaussians

$$p\left(\begin{bmatrix} x \\ y \end{bmatrix}; \mu, \Sigma\right) = \frac{1}{2\pi^{\frac{n}{2}} |\Sigma|^{0.5}} \exp\left(-\frac{1}{2}\left(\begin{bmatrix} x \\ y \end{bmatrix} - \begin{bmatrix} \mu_X \\ \mu_Y \end{bmatrix}\right)^T \begin{bmatrix} \Gamma_{XX} & \Gamma_{XY} \\ \Gamma_{YX} & \Gamma_{YY} \end{bmatrix} \left(\begin{bmatrix} x \\ y \end{bmatrix} - \begin{bmatrix} \mu_X \\ \mu_Y \end{bmatrix}\right)\right)$$

$$p(x|Y = y_0) \quad \propto \quad p\left(\begin{bmatrix} x \\ y_0 \end{bmatrix}; \mu, \Sigma\right)$$

Sketch:

1. We will write out the whole likelihood in quadratic form
2. Separate out the y and x terms
3. Complete the square and back out to a conditional Gaussian likelihood

Proofs (Skipped in Lecture)

Recap of linear algebra lemmas

- Matrix Inversion Lemma

$$\begin{bmatrix} A & B \\ C & D \end{bmatrix}^{-1} = \begin{bmatrix} (M \setminus D)^{-1} & -A^{-1}B(M \setminus A)^{-1} \\ -D^{-1}C(M \setminus D)^{-1} & (M \setminus A)^{-1} \end{bmatrix}$$

Schur Complement

$$(M \setminus A) = D - CA^{-1}B$$

$$(M \setminus D) = A - BD^{-1}C$$

Partitioned Multivariate Gaussians: Dual

$$\Gamma = \begin{bmatrix} \Gamma_{XX} & \Gamma_{XY} \\ \Gamma_{YX} & \Gamma_{YY} \end{bmatrix} = \Sigma^{-1} = \begin{bmatrix} \Sigma_{XX} & \Sigma_{XY} \\ \Sigma_{YX} & \Sigma_{YY} \end{bmatrix}^{-1}$$

Remember matrix lemmas

$$\begin{bmatrix} A & B \\ C & D \end{bmatrix}^{-1} = \begin{bmatrix} (M \setminus D)^{-1} & -A^{-1}B(M \setminus A)^{-1} \\ -D^{-1}C(M \setminus D)^{-1} & (M \setminus A)^{-1} \end{bmatrix}$$

$$(M \setminus A) = D - CA^{-1}B$$

$$(M \setminus D) = A - BD^{-1}C$$

Partitioned Multivariate Gaussians: Dual

$$\Gamma = \begin{bmatrix} \Gamma_{XX} & \Gamma_{XY} \\ \Gamma_{YX} & \Gamma_{YY} \end{bmatrix} = \Sigma^{-1} = \begin{bmatrix} \Sigma_{XX} & \Sigma_{XY} \\ \Sigma_{YX} & \Sigma_{YY} \end{bmatrix}^{-1}$$

Remember matrix lemmas

$$\begin{bmatrix} A & B \\ C & D \end{bmatrix}^{-1} = \begin{bmatrix} (M \setminus D)^{-1} & -A^{-1}B(M \setminus A)^{-1} \\ -D^{-1}C(M \setminus D)^{-1} & (M \setminus A)^{-1} \end{bmatrix}$$

$$\Gamma_{XX} = (\Sigma_{XX} - \Sigma_{XY}\Sigma_{YY}^{-1}\Sigma_{YX})^{-1}$$

$$\Gamma_{YY} = (\Sigma_{YY} - \Sigma_{YX}\Sigma_{XX}^{-1}\Sigma_{XY})^{-1}$$

$$\Gamma_{XY} = -\Sigma_{XX}^{-1}\Sigma_{XY}(\Sigma_{YY} - \Sigma_{YX}\Sigma_{XX}^{-1}\Sigma_{XY})^{-1}$$

$$\Gamma_{YX} = -\Sigma_{YY}^{-1}\Sigma_{YX}(\Sigma_{XX} - \Sigma_{XY}\Sigma_{YY}^{-1}\Sigma_{YX})^{-1}$$

Important Gaussian Identities

- Marginalization:

$$p(X) = \mathcal{N}(\mu_X, \Sigma_{XX})$$

$$p(Y) = \mathcal{N}(\mu_Y, \Sigma_{YY})$$

- Conditioning

$$p(X|Y = y_0) = \mathcal{N}(\mu_X + \Sigma_{XY}\Sigma_{YY}^{-1}(y_0 - \mu_Y), \Sigma_{XX} - \Sigma_{XY}\Sigma_{YY}^{-1}\Sigma_{YX})$$

Marginalization of Gaussian

$$p(\begin{bmatrix} x \\ y \end{bmatrix}; \mu, \Sigma) = \frac{1}{2\pi^{\frac{n}{2}} |\Sigma|^{0.5}} \exp\left(-\frac{1}{2} \left(\begin{bmatrix} x \\ y \end{bmatrix} - \begin{bmatrix} \mu_X \\ \mu_Y \end{bmatrix} \right)^T \begin{bmatrix} \Gamma_{XX} & \Gamma_{XY} \\ \Gamma_{YX} & \Gamma_{YY} \end{bmatrix} \left(\begin{bmatrix} x \\ y \end{bmatrix} - \begin{bmatrix} \mu_X \\ \mu_Y \end{bmatrix} \right) \right)$$

$$p(x) = \int p(\begin{bmatrix} x \\ y \end{bmatrix}; \mu, \Sigma) dy = \frac{1}{2\pi^{\frac{n}{2}} |\Sigma|^{0.5}} \int \exp\left(-\frac{1}{2} \left(\begin{bmatrix} x \\ y \end{bmatrix} - \begin{bmatrix} \mu_X \\ \mu_Y \end{bmatrix} \right)^T \begin{bmatrix} \Gamma_{XX} & \Gamma_{XY} \\ \Gamma_{YX} & \Gamma_{YY} \end{bmatrix} \left(\begin{bmatrix} x \\ y \end{bmatrix} - \begin{bmatrix} \mu_X \\ \mu_Y \end{bmatrix} \right) \right) dy$$

Sketch:

1. We will write out the whole likelihood in quadratic form
2. We will add and subtract terms to complete the square
3. We will use the fact that Gaussians integrate to 1 to simplify
4. We will show that the marginal is p(x)

Marginalization of Gaussian

$$p(x) = \int p(\begin{bmatrix} x \\ y \end{bmatrix}; \mu, \Sigma) dy = \frac{1}{2\pi^{\frac{n}{2}}|\Sigma|^{0.5}} \int \exp\left(-\frac{1}{2}\left(\begin{bmatrix} x \\ y \end{bmatrix} - \begin{bmatrix} \mu_X \\ \mu_Y \end{bmatrix}\right)^T \begin{bmatrix} \Gamma_{XX} & \Gamma_{XY} \\ \Gamma_{YX} & \Gamma_{YY} \end{bmatrix} \left(\begin{bmatrix} x \\ y \end{bmatrix} - \begin{bmatrix} \mu_X \\ \mu_Y \end{bmatrix}\right)\right) dy$$

$$= \frac{1}{2\pi^{\frac{n}{2}}|\Sigma|^{0.5}} \int \exp\Bigg(-\frac{1}{2}((x-\mu_X)^T\Gamma_{XX}(x-\mu_X) + (y-\mu_Y)^T\Gamma_Y(y-\mu_Y) + 2(y-\mu_Y)^T\Gamma_{XY}(x-\mu_X))dy$$

$$\begin{aligned} = \frac{1}{2\pi^{\frac{n}{2}}|\Sigma|^{0.5}} \int \exp\Bigg(-\frac{1}{2}((x-\mu_X)^T\Gamma_{XX}(x-\mu_X) &+ (y-\mu_Y)^T\Gamma_Y(y-\mu_Y) \\ &+ 2(y-\mu_Y)^T\Gamma_{YY}\Gamma_{YY}^{-1}\Gamma_{XY}(x-\mu_X) \\ &+ (x-\mu_X)^T\Gamma_{XY}\Gamma_{YY}^{-1}\Gamma_{YY}\Gamma_{YY}^{-1}\Gamma_{YX}(x-\mu_X) \\ &-(x-\mu_X)^T\Gamma_{XY}\Gamma_{YY}^{-1}\Gamma_{YY}\Gamma_{YY}^{-1}\Gamma_{YX}(x-\mu_X))dy \end{aligned}$$

$$\begin{aligned} = \frac{1}{2\pi^{\frac{n}{2}}|\Sigma|^{0.5}} \exp\left(-\frac{1}{2}((x-\mu_X)^T\Gamma_{XX}(x-\mu_X) - (x-\mu_X)^T\Gamma_{XY}\Gamma_{YY}^{-1}\Gamma_{YY}\Gamma_{YY}^{-1}\Gamma_{YX}(x-\mu_X))\right) \int \exp\Bigg(-\frac{1}{2}((y-\mu_Y)^T\Gamma_Y(y-\mu_Y) \\ +2(y-\mu_Y)^T\Gamma_{YY}\Gamma_{YY}^{-1}\Gamma_{XY}(x-\mu_X) \\ +(x-\mu_X)^T\Gamma_{XY}\Gamma_{YY}^{-1}\Gamma_{YY}\Gamma_{YY}^{-1}\Gamma_{YX}(x-\mu_X))dy \end{aligned}$$

$$= \frac{2\pi^{\frac{n_Y}{2}}|\Gamma_{YY}^{-1}|^{0.5}}{2\pi^{\frac{n}{2}}|\Sigma|^{0.5}} \exp\left(-\frac{1}{2}((x-\mu_X)^T(\Gamma_{XX} - \Gamma_{XY}\Gamma_{YY}^{-1}\Gamma_{YY}\Gamma_{YY}^{-1}\Gamma_{YX})(x-\mu_X))\right)$$

$$= \mathcal{N}(\mu_X, (\Gamma_{XX} - \Gamma_{XY}\Gamma_{YY}^{-1}\Gamma_{YX})^{-1}) = \mathcal{N}(\mu_X, \Sigma_{XX})$$

Marginalization Recap

$$p(\begin{bmatrix} x \\ y \end{bmatrix}; \mu, \Sigma) = \frac{1}{2\pi^{\frac{n}{2}}|\Sigma|^{0.5}} \exp\left(-\frac{1}{2}\left(\begin{bmatrix} x \\ y \end{bmatrix} - \begin{bmatrix} \mu_X \\ \mu_Y \end{bmatrix}\right)^T \begin{bmatrix} \Gamma_{XX} & \Gamma_{XY} \\ \Gamma_{YX} & \Gamma_{YY} \end{bmatrix} \left(\begin{bmatrix} x \\ y \end{bmatrix} - \begin{bmatrix} \mu_X \\ \mu_Y \end{bmatrix}\right)\right)$$

$$p(x) = \mathcal{N}(\mu_X, (\Gamma_{XX} - \Gamma_{XY}\Gamma_{YY}^{-1}\Gamma_{YX})^{-1}) = \mathcal{N}(\mu_X, \Sigma_{XX})$$

Very simple result for marginalization

Simply grab the appropriate partitioned matrix, same holds for Y

Conditioning with Gaussians

$$p(\begin{bmatrix} x \\ y \end{bmatrix}; \mu, \Sigma) = \frac{1}{2\pi^{\frac{n}{2}} |\Sigma|^{0.5}} \exp\left(-\frac{1}{2} \left(\begin{bmatrix} x \\ y \end{bmatrix} - \begin{bmatrix} \mu_X \\ \mu_Y \end{bmatrix} \right)^T \begin{bmatrix} \Gamma_{XX} & \Gamma_{XY} \\ \Gamma_{YX} & \Gamma_{YY} \end{bmatrix} \left(\begin{bmatrix} x \\ y \end{bmatrix} - \begin{bmatrix} \mu_X \\ \mu_Y \end{bmatrix} \right) \right)$$

$$p(x|Y = y_0) \quad \propto \quad p(\begin{bmatrix} x \\ y_0 \end{bmatrix}; \mu, \Sigma)$$

Sketch:

1. We will write out the whole likelihood in quadratic form
2. Separate out the y and x terms
3. Complete the square and back out to a conditional Gaussian likelihood

Conditioning with Gaussians

$$p(\begin{bmatrix} x \\ y \end{bmatrix}; \mu, \Sigma) = \frac{1}{2\pi^{\frac{n}{2}}|\Sigma|^{0.5}} \exp\left(-\frac{1}{2}\left(\begin{bmatrix} x \\ y \end{bmatrix} - \begin{bmatrix} \mu_X \\ \mu_Y \end{bmatrix} \right)^T \begin{bmatrix} \Gamma_{XX} & \Gamma_{XY} \\ \Gamma_{YX} & \Gamma_{YY} \end{bmatrix} \left(\begin{bmatrix} x \\ y \end{bmatrix} - \begin{bmatrix} \mu_X \\ \mu_Y \end{bmatrix} \right) \right)$$

$$p(x|Y = y_0) \quad \propto \quad p(\begin{bmatrix} x \\ y_0 \end{bmatrix}; \mu, \Sigma)$$

$$\propto \quad \exp\left(-\frac{1}{2}(x - \mu_X)^\top \Gamma_{XX}(x - \mu_X) - (x - \mu_X)^\top \Gamma_{XY}(y_0 - \mu_Y) - \frac{1}{2}(y_0 - \mu_Y)^\top \Gamma_{YY}(y_0 - \mu_Y) \right)$$

$$\propto \quad \exp\left(-\frac{1}{2}(x - \mu_X)^\top \Gamma_{XX}(x - \mu_X) - (x - \mu_X)^\top \Gamma_{XY}(y_0 - \mu_Y) \right)$$

$$= \quad \exp\left(-\frac{1}{2}(x - \mu_X)^\top \Gamma_{XX}(x - \mu_X) - (x - \mu_X)^\top \Gamma_{XX}\Gamma_{XX}^{-1}\Gamma_{XY}(y_0 - \mu_Y) - \frac{1}{2}(y_0 - \mu_Y)\Gamma_{YX}\Gamma_{XX}^{-1}\Gamma_{XX}\Gamma_{XX}^{-1}\Gamma_{XY}(y_0 - \mu_Y) + \frac{1}{2}(y_0 - \mu_Y)\Gamma_{YX}\Gamma_{XX}^{-1}\Gamma_{XX}\Gamma_{XX}^{-1}\Gamma_{XY}(y_0 - \mu_Y) \right)$$

$$= \quad \exp\left(-\frac{1}{2}(x - \mu_X + \Gamma_{XX}^{-1}\Gamma_{XY}(y_0 - \mu_Y)^\top \Gamma_{XX}(x - \mu_X + \Gamma_{XX}^{-1}\Gamma_{XY}(y_0 - \mu_Y) \right) \exp\left(\frac{1}{2}(y_0 - \mu_Y)\Gamma_{YX}\Gamma_{XX}^{-1}\Gamma_{XX}\Gamma_{XX}^{-1}\Gamma_{XY}(y_0 - \mu_Y) \right)$$

$$\propto \quad \exp\left(-\frac{1}{2}(x - \mu_X + \Gamma_{XX}^{-1}\Gamma_{XY}(y_0 - \mu_Y)^\top \Gamma_{XX}(x - \mu_X + \Gamma_{XX}^{-1}\Gamma_{XY}(y_0 - \mu_Y) \right)$$

$$= \mathcal{N}(\mu_X + \Gamma_{XX}^{-1}\Gamma_{XY}(y_0 - \mu_Y), \Gamma_{XX}^{-1})$$
$$= \mathcal{N}(\mu_X + \Sigma_{XY}\Sigma_{YY}^{-1}(y_0 - \mu_Y), \Sigma_{XX} - \Sigma_{XY}\Sigma_{YY}^{-1}\Sigma_{YX})$$

- Conditional mean shifted depending on y_0
- Covariance not dependent on y_0

Important Gaussian Identities

- Marginalization:

$$p(X) = \mathcal{N}(\mu_X, \Sigma_{XX})$$

$$p(Y) = \mathcal{N}(\mu_Y, \Sigma_{YY})$$

- Conditioning

$$\begin{aligned} p(X|Y = y_0) &= \mathcal{N}(\mu_X + \Gamma_{XX}^{-1}\Gamma_{XY}(y_0 - \mu_Y), \Gamma_{XX}^{-1}) \\ &= \mathcal{N}(\mu_X + \Sigma_{XY}\Sigma_{YY}^{-1}(y_0 - \mu_Y), \Sigma_{XX} - \Sigma_{XY}\Sigma_{YY}^{-1}\Sigma_{YX}) \end{aligned}$$

Lecture Outline

Recap

↓

Bayesian Filtering

↓

Gaussian Properties

↓

Kalman Filtering

Discrete Kalman Filter

Kalman filter = Bayes filter with Linear Gaussian dynamics and sensor models

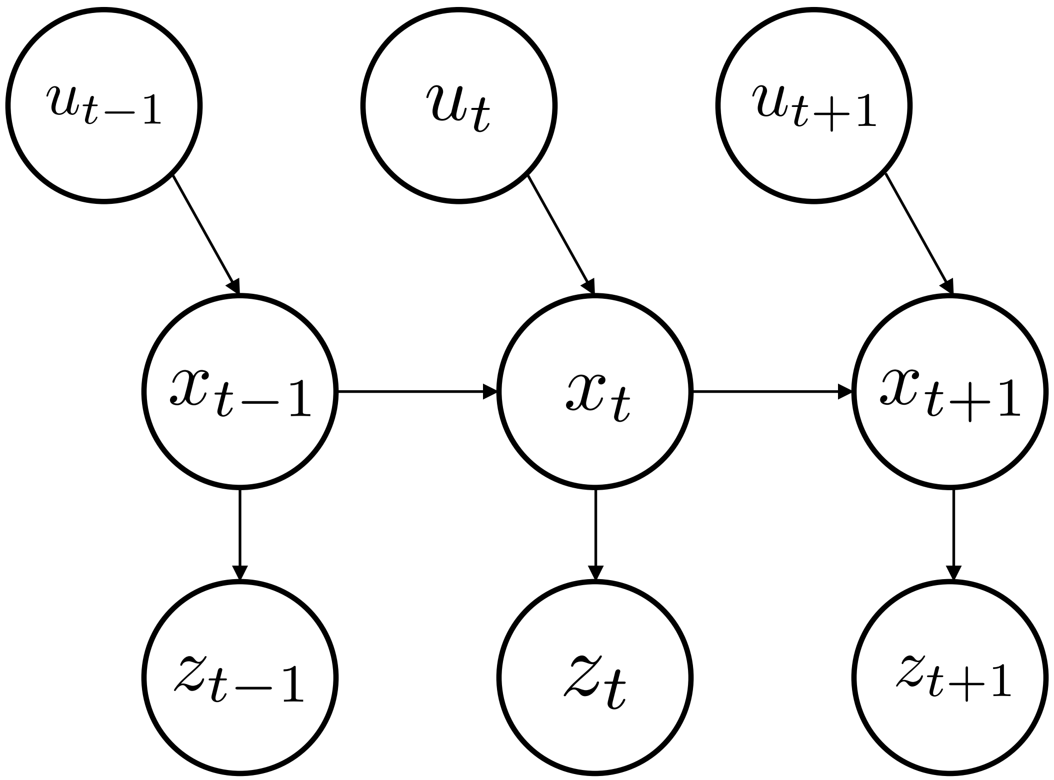

Discrete Kalman Filter

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

$$x_{t+1} = Ax_t + Bu_t + \epsilon_t$$
$$\epsilon_t \sim \mathcal{N}(0, Q)$$

with a measurement

$$z_{t+1} = Cx_{t+1} + \delta_t$$
$$\delta_t \sim \mathcal{N}(0, R)$$

Linear Gaussian

Components of a Kalman Filter

A Matrix (n x n) that describes how the state evolves from ***t-1*** to ***t*** without controls or noise.

B Matrix (n x l) that describes how the control $\boldsymbol{u_{t-1}}$ changes the state from ***t-1*** to ***t***

C Matrix (k x n) that describes how to map the state $\boldsymbol{x_t}$ to an observation $\boldsymbol{z_t}$.

ϵ_t δ_t Random variables representing the process and measurement noise that are assumed to be independent and normally distributed with covariance ***R*** and ***Q*** respectively.

Goal of the Kalman Filter

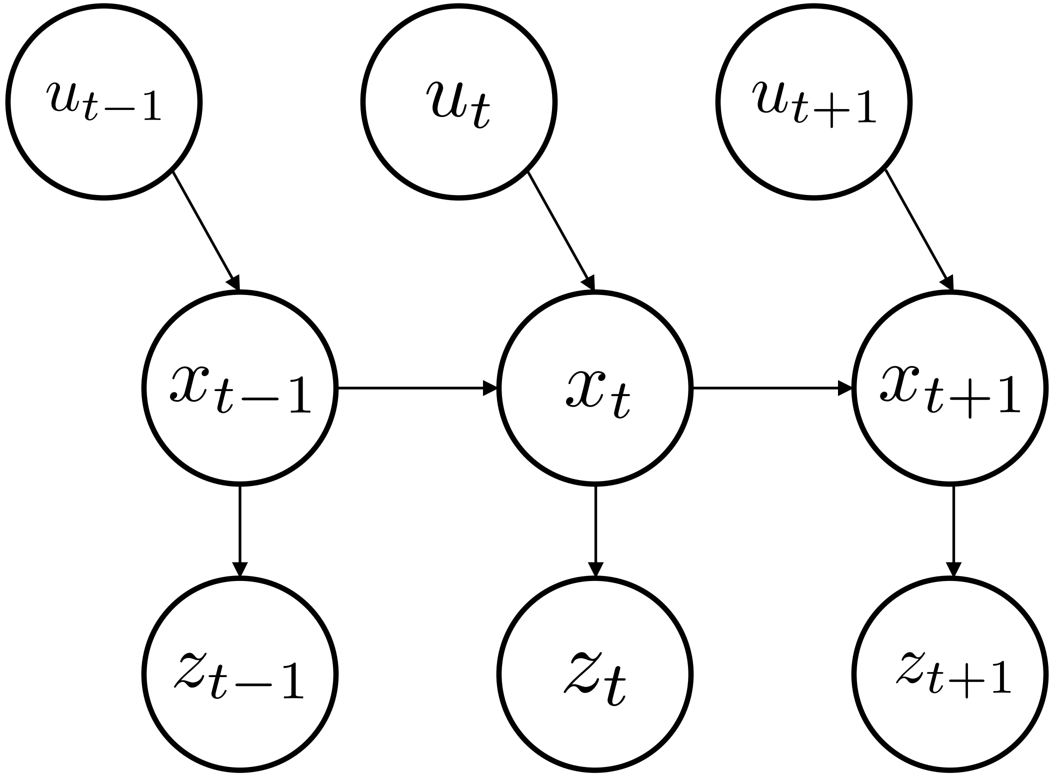

Belief

$$p(x_t | z_{0:t}, u_{0:t-1})$$

Idea: recursive update for Bayes filter

$$\propto p(z_t|x_t) \int p(x_t|x_{t-1}, u_{t-1}) p(x_{t-1}|z_{0:t-1}, u_{0:t-2}) dx_{t-1}$$

Measurement

Dynamics

Recursive Belief

2 step process:

- Dynamics update (incorporate action)
- Measurement update (incorporate sensor reading)

Linear Gaussian Systems: Initialization

- Initial belief is normally distributed:

$$bel(x_0) = N(x_0; \mu_0, \Sigma_0)$$

Linear Gaussian Systems: Dynamics

Initial Prior

$p(x_0)$

Dynamics/Prediction

Estimate $\overline{Bel}(x_t)$

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

Measurement/Correction

Estimate $Bel(x_t)$

$$Bel(x_t) = \eta p(z_t|x_t)\overline{Bel}(x_t) = \frac{p(z_t|x_t)\overline{Bel}(x_t)}{\sum_{x'} p(z_t|x_t = x')\overline{Bel}(x_t = x')}$$

Linear Gaussian Systems: Dynamics

- Integrate the effect of one action under the dynamics, before measurement comes in

$$x_{t+1} = Ax_t + Bu_t + \epsilon_t \qquad \epsilon_t \sim \mathcal{N}(0, Q_t)$$

$$p(x_{t+1}|x_t, u_t) = \mathcal{N}(Ax_t + Bu_t, Q_t)$$

$$p(x_{t+1}|z_{0:t}, u_{0:t}) = \int p(x_t|u_{0:t-1}, z_{0:t})p(x_{t+1}|x_t, u_t)dx_t$$

Gaussian, easy!

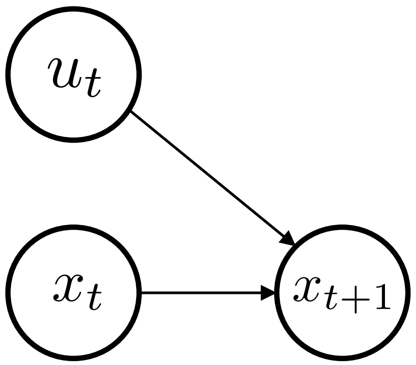

Linear Gaussian Systems: Dynamics Intuition

Previous belief $p(x_t|u_{0:t}, z_{0:t}) \sim \mathcal{N}(\mu_{t|0:t}, \Sigma_{t|0:t})$

Belief Update $p(x_{t+1}|z_{0:t}, u_{0:t+1}) \sim \mathcal{N}(A\mu_{t|0:t} + Bu_t, A\Sigma_{t|0:t}A^T + Q_t)$

Intuition: Scale and shift the mean according to dynamics, uncertainty grows quadratically!

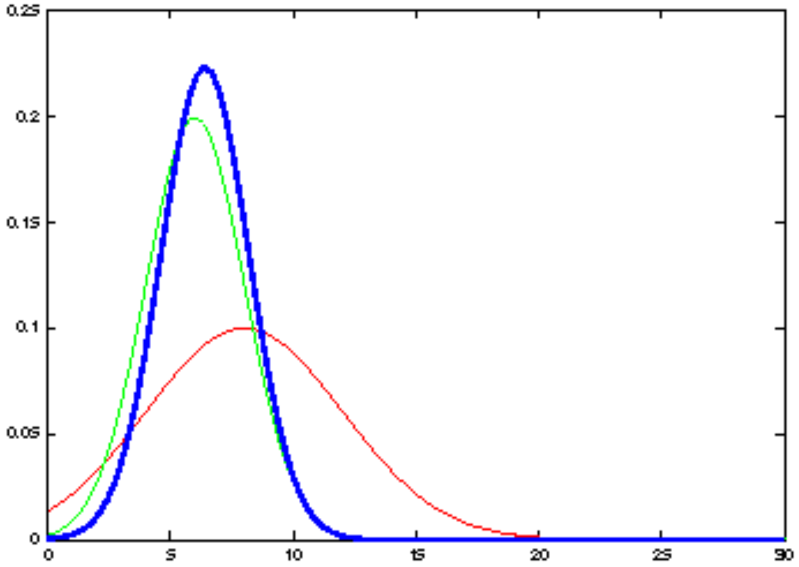

Belief at x_t

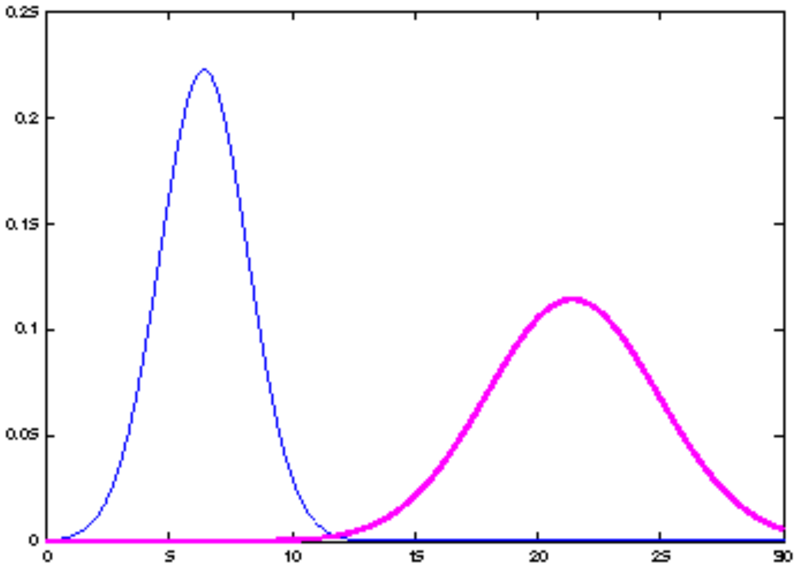

Belief post dynamics → shifted mean, scaled and shifted variance

Linear Gaussian Systems: Dynamics

- Integrate the effect of one action under the dynamics, before measurement comes in

$$x_{t+1} = Ax_t + Bu_t + \epsilon_t \qquad \epsilon_t \sim \mathcal{N}(0, Q_t)$$

$$p(x_{t+1}|x_t, u_t) = \mathcal{N}(Ax_t + Bu_t, Q_t)$$

$$p(x_{t+1}|z_{0:t}, u_{0:t}) = \int p(x_t|u_{0:t-1}, z_{0:t})p(x_{t+1}|x_t, u_t)dx_t$$

Previous belief $p(x_t|u_{0:t-1}, z_{0:t}) \sim \mathcal{N}(\mu_{t|0:t}, \Sigma_{t|0:t})$

Belief Update $p(x_{t+1}|z_{0:t}, u_{0:t}) \sim \mathcal{N}(A\mu_{t|0:t} + Bu_t, A\Sigma_{t|0:t}A^T + Q_t)$

How??

Linear Gaussian Systems: Dynamics

- Integrate the effect of one action under the dynamics, before measurement comes in

$$p(x_{t+1}|z_{0:t}, u_{0:t}) = \int p(x_t|u_{0:t-1}, z_{0:t})p(x_{t+1}|x_t, u_t)dx_t$$

Stays in Gaussian world

$$(X_{t+1}, X_t)|_{z_{0:t}, u_{0:t}} \sim \mathcal{N}\left(\begin{bmatrix} \mu_{t|0:t} \\ \mu_{t+1|0:t} \end{bmatrix}, \begin{bmatrix} \Sigma_{t|0:t} & \Sigma_{t,t+1|0:t} \\ \Sigma_{t+1,t|0:t} & \Sigma_{t+1|0:t} \end{bmatrix} \right)$$

Current belief at time t

Now compute the mean and covariance and then marginalize

Linear Gaussian Systems: Dynamics

$$(X_{t+1}, X_t)|_{z_{0:t}, u_{0:t}} \sim \mathcal{N}\left(\begin{bmatrix} \mu_{t|0:t} \\ \mu_{t+1|0:t} \end{bmatrix}, \begin{bmatrix} \Sigma_{t|0:t} & \Sigma_{t,t+1|0:t} \\ \Sigma_{t+1,t|0:t} & \Sigma_{t+1|0:t} \end{bmatrix} \right)$$

$$\mu_{t+1|0:t} = \mathbb{E}\left[X_{t+1}|z_{0:t}, u_{0:t}\right]$$

Mean

$$\Sigma_{t+1|0:t} = \mathbb{E}\left[(X_{t+1|0:t} - \mu_{t+1|0:t})(X_{t+1|0:t} - \mu_{t+1|0:t})^T\right]$$

Diagonal Covariance

$$\Sigma_{t,t+1|0:t} = \mathbb{E}\left[(X_{t|0:t} - \mu_{t|0:t})(X_{t+1|0:t} - \mu_{t+1|0:t})^T\right]$$

Cross Covariance

Linear Gaussian Systems: Dynamics

$$(X_{t+1}, X_t)|_{z_{0:t}, u_{0:t}} \sim \mathcal{N}\left(\begin{bmatrix} \mu_{t|0:t} \\ \mu_{t+1|0:t} \end{bmatrix}, \begin{bmatrix} \Sigma_{t|0:t} & \Sigma_{t,t+1|0:t} \\ \Sigma_{t+1,t|0:t} & \Sigma_{t+1|0:t} \end{bmatrix} \right)$$

Mean

$$\begin{aligned} \mu_{t+1|0:t} &= \mathbb{E}\left[X_{t+1} | z_{0:t}, u_{0:t}\right] \\ &= \mathbb{E}\left[AX_t + Bu_t + \epsilon_t | z_{0:t}, u_{0:t}\right] \\ &= A\mathbb{E}\left[X_t | z_{0:t}, u_{0:t}\right] + Bu_t + \mathbb{E}\left[\epsilon_t | z_{0:t}, u_{0:t}\right] \\ &= A\mu_{t|0:t} + Bu_t \end{aligned}$$

Diagonal Covariance

$$\begin{aligned} \Sigma_{t+1|0:t} &= \mathbb{E}\left[(X_{t+1|0:t} - \mu_{t+1|0:t})(X_{t+1|0:t} - \mu_{t+1|0:t})^T\right] \\ &= \mathbb{E}\left[(AX_{t|0:t} + Bu_t + \epsilon_t - A\mu_{t|0:t} - Bu_t)(AX_{t|0:t} + Bu_t + \epsilon_t - A\mu_{t|0:t} - Bu_t)^T\right] \\ &= A\mathbb{E}\left[(X_{t|0:t} - \mu_{t|0:t})(X_{t|0:t} - \mu_{t|0:t})^T\right] A^T + Q_t \\ &= A\Sigma_{t|0:t}A^T + Q_t \end{aligned}$$

Cross Covariance

$$\Sigma_{t,t+1|0:t} = \mathbb{E}\left[(X_{t|0:t} - \mu_{t|0:t})(X_{t+1|0:t} - \mu_{t+1|0:t})^T\right]$$

$$\Sigma_{t,t+1|0:t} = \Sigma_{t|0:t}A^T$$

Linear Gaussian Systems: Dynamics

$$(X_{t+1}, X_t)|_{z_{0:t}, u_{0:t}} \sim \mathcal{N}\left(\begin{bmatrix} \mu_{t|0:t} \\ \mu_{t+1|0:t} \end{bmatrix}, \begin{bmatrix} \Sigma_{t|0:t} & \Sigma_{t,t+1|0:t} \\ \Sigma_{t+1,t|0:t} & \Sigma_{t+1|0:t} \end{bmatrix} \right)$$

Mean $$\mu_{t+1|0:t} = A\mu_{t|0:t} + Bu_t$$

Diagonal Covariance $$\Sigma_{t+1|0:t} = A\Sigma_{t|0:t}A^T + Q_t$$

Previous belief $$p(x_t|u_{0:t-1}, z_{0:t}) \sim \mathcal{N}(\mu_{t|0:t}, \Sigma_{t|0:t})$$

Belief Update $$p(x_{t+1}|z_{0:t}, u_{0:t}) \sim \mathcal{N}(A\mu_{t|0:t} + Bu_t, A\Sigma_{t|0:t}A^T + Q_t)$$

Intuition: Scale and shift the mean according to dynamics, uncertainty grows quadratically!

Linear Gaussian Systems: Dynamics

Previous belief $$p(x_t | u_{0:t}, z_{0:t}) \sim \mathcal{N}(\mu_{t|0:t}, \Sigma_{t|0:t})$$

Belief Update $$p(x_{t+1} | z_{0:t}, u_{0:t+1}) \sim \mathcal{N}(A\mu_{t|0:t} + Bu_t, A\Sigma_{t|0:t}A^T + Q_t)$$

Intuition: Scale and shift the mean according to dynamics, uncertainty grows!

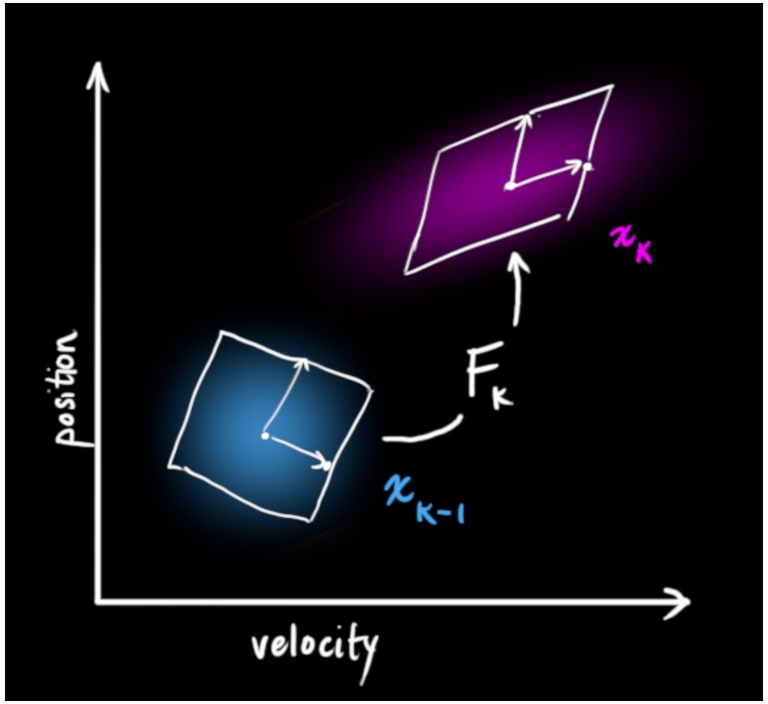

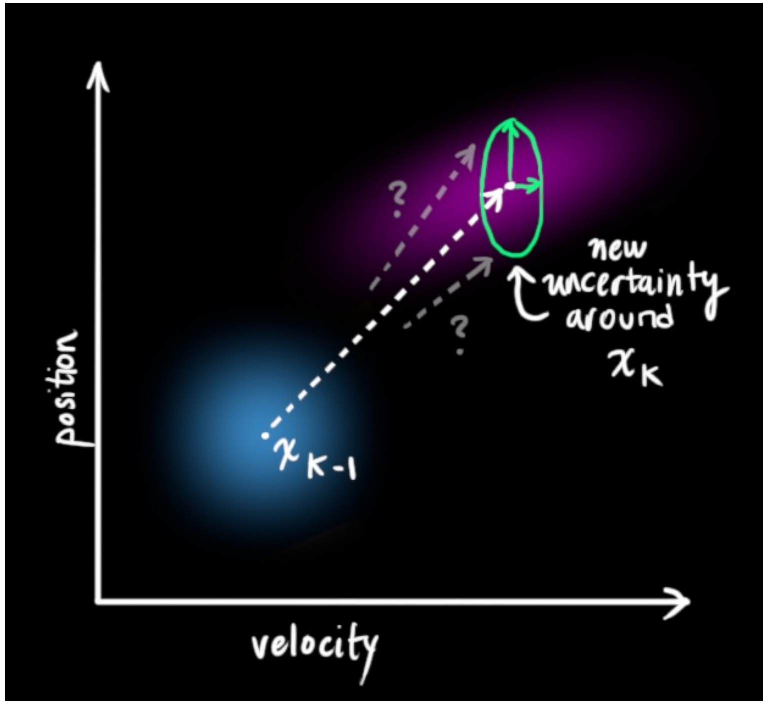

Intuition Behind Prediction Step

Previous belief $p(x_t | u_{0:t}, z_{0:t}) \sim \mathcal{N}(\mu_{t|0:t}, \Sigma_{t|0:t})$

Belief Update $p(x_{t+1} | z_{0:t}, u_{0:t+1}) \sim \mathcal{N}(A\mu_{t|0:t} + Bu_t, A\Sigma_{t|0:t}A^T + Q_t)$

Intuition: Scale and shift the mean according to dynamics, uncertainty grows!

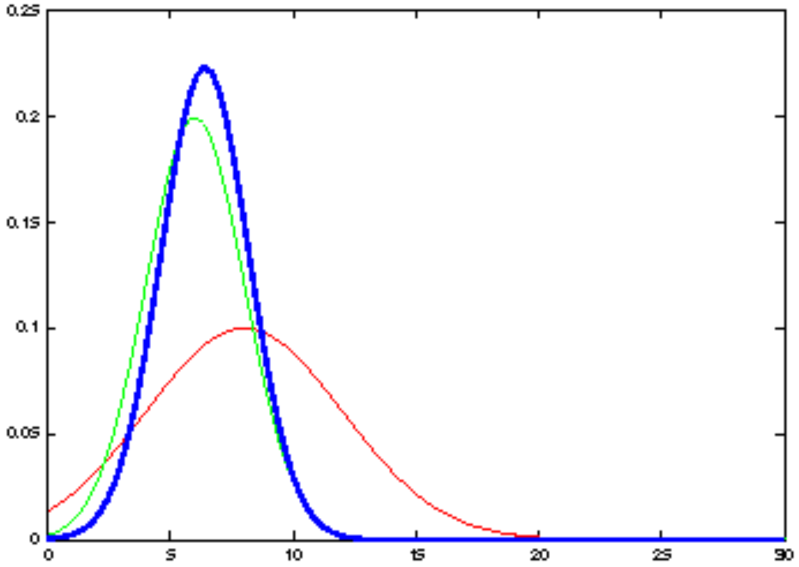

Belief at x_t

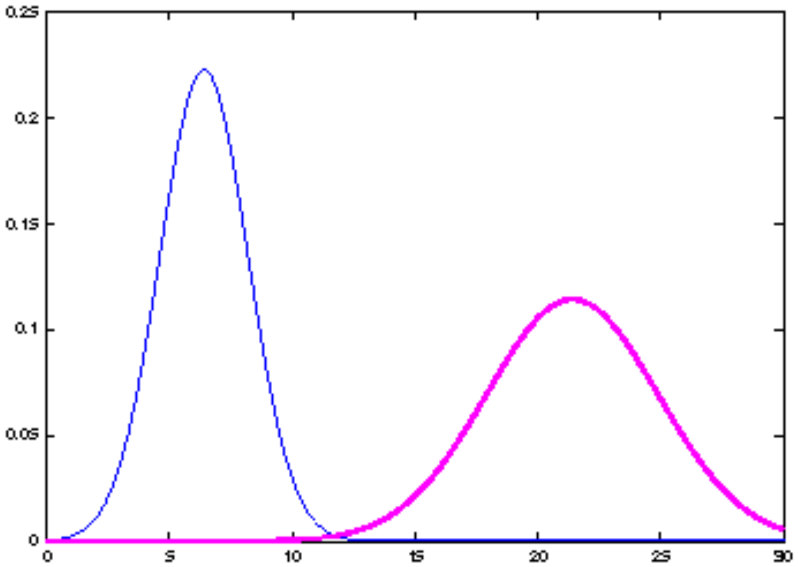

Belief post dynamics → shifted mean, scaled and shifted variance

Kalman Filter: Where are we?

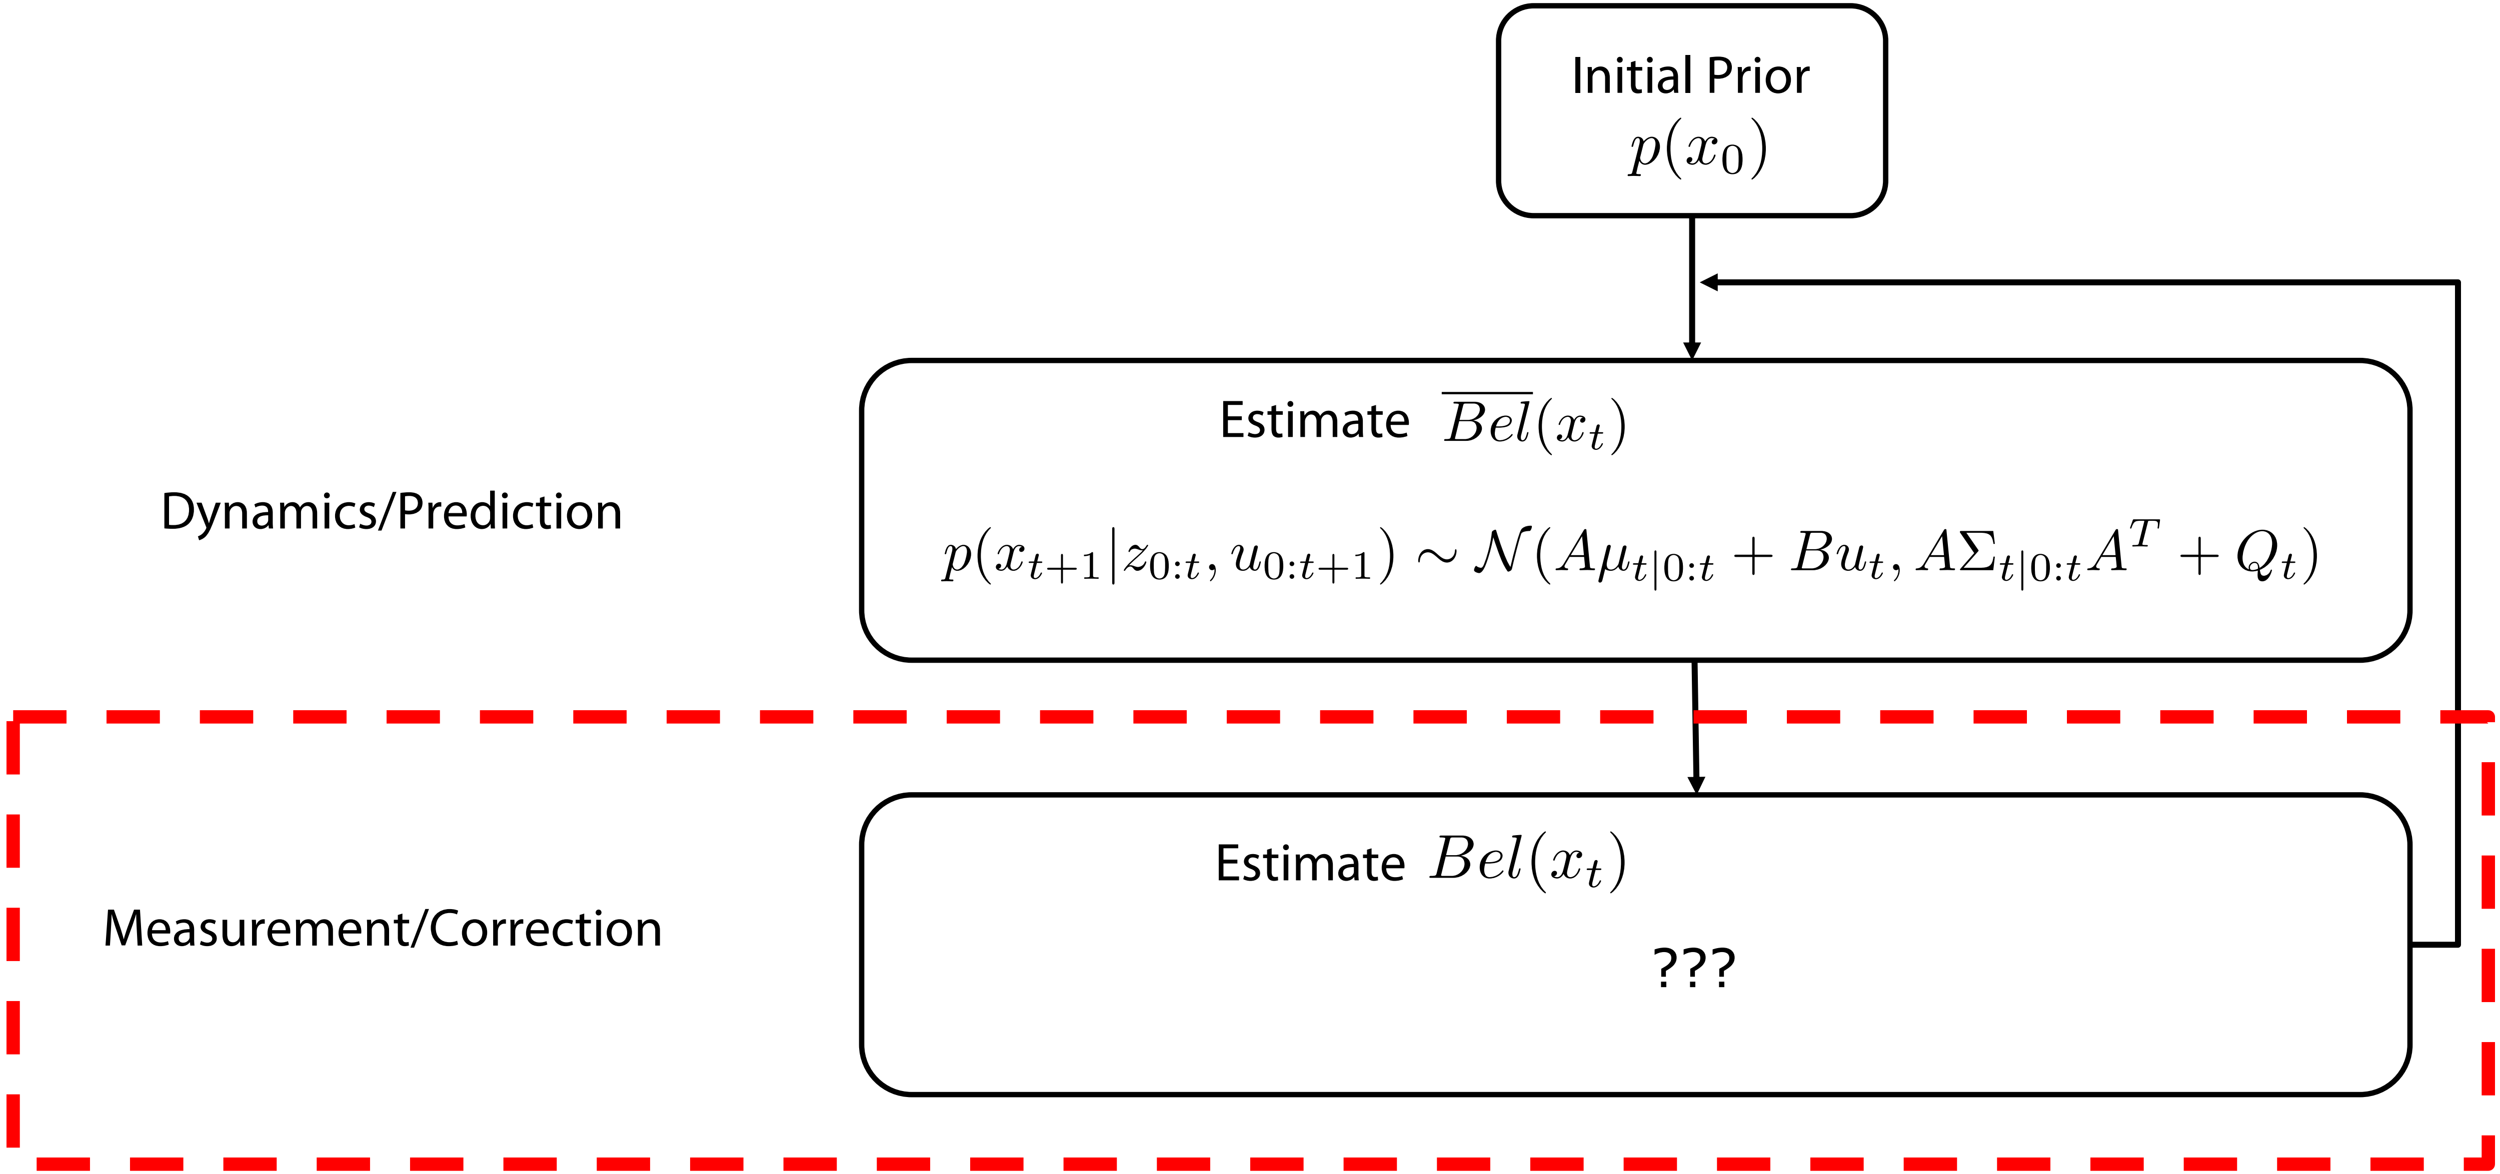

Linear Gaussian Systems: Observations

Measurement/Correction

Estimate $Bel(x_t)$

???

$$p(x_{t+1}|z_{0:t+1}, u_{0:t}) = \eta\, p(z_t|x_t)p(x_{t+1}|z_{0:t}, u_{0:t})$$

Need to do conditioning/normalization with Linear Gaussians

Linear Gaussian Systems: Observations

- Integrate the effect of an observation using sensor model, after dynamics

$$z_{t+1} = Cx_{t+1} + \delta_t \qquad \delta_t \sim \mathcal{N}(0, R_t)$$

$$p(z_{t+1}|x_{t+1}) = \mathcal{N}(Cx_{t+1}, R_t)$$

$$p(x_{t+1}|z_{0:t+1}, u_{0:t}) \propto p(z_{t+1}|x_{t+1})p(x_{t+1}|z_{0:t}, u_{0:t})$$

Gaussian, easy to normalize

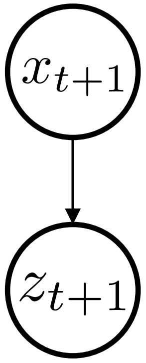

Linear Gaussian Systems: Observations Intuition

Previous belief (post dynamics)

New Measurement

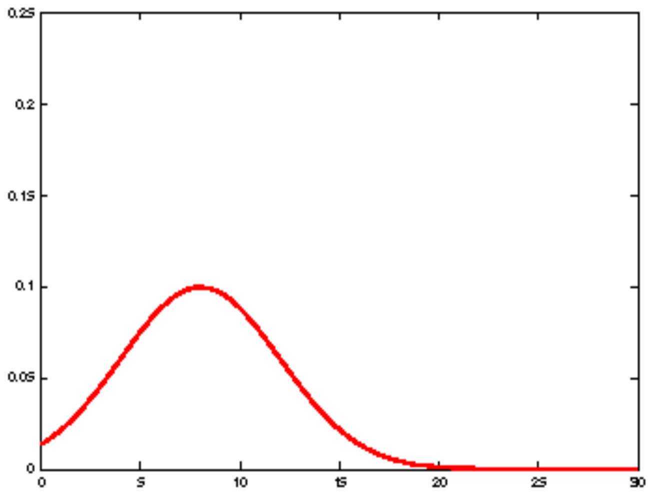

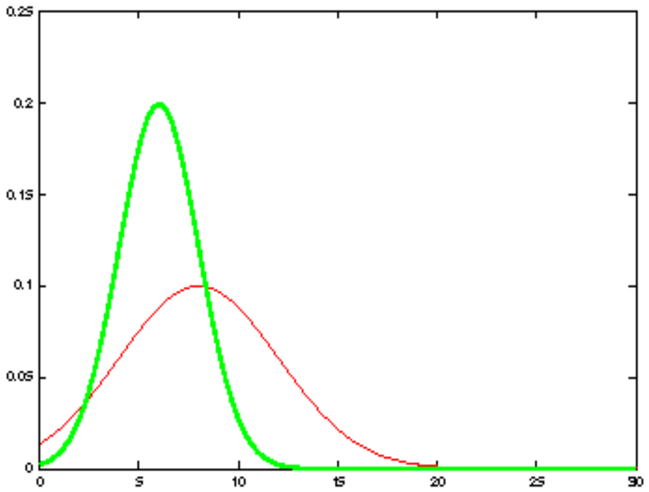

For the sake of simplicity, let's say C = I

Previous belief

$$p(x_{t+1}|z_{0:t}, u_{0:t}) = \mathcal{N}(\mu_{t+1|0:t}, \Sigma_{t+1|0:t})$$

Updated belief

$$p(x_{t+1}|z_{0:t+1}, u_{0:t})$$

$$= \mathcal{N}(\mu_{t+1|0:t} + K_{t+1}(z_{t+1} - \mu_{t+1|0:t}), (I - K_{t+1})\Sigma_{t+1|0:t})$$

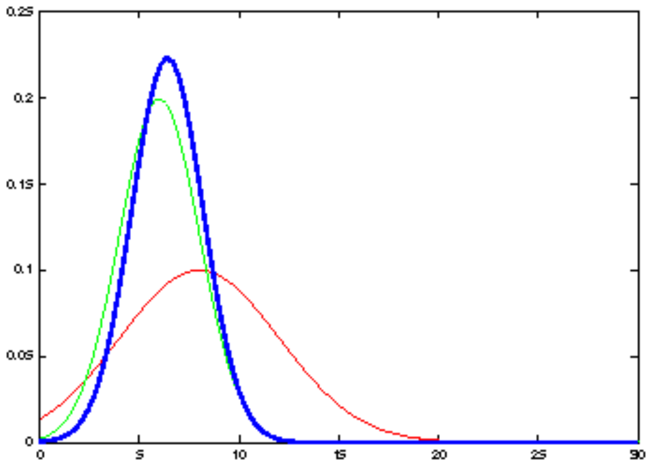

Linearly interpolate between measurement and previous mean based on K

Scale down uncertainty based on K

Linear Gaussian Systems: Observations

- Integrate the effect of an observation using sensor model, after dynamics

$$z_{t+1} = Cx_{t+1} + \delta_t \qquad \delta_t \sim \mathcal{N}(0, R_t)$$

$$p(z_{t+1}|x_{t+1}) = \mathcal{N}(Cx_{t+1}, R_t)$$

$$p(x_{t+1}|z_{0:t+1}, u_{0:t}) \propto p(z_{t+1}|x_{t+1})p(x_{t+1}|z_{0:t}, u_{0:t})$$

Previous belief $p(x_{t+1}|z_{0:t}, u_{0:t}) = \mathcal{N}(\mu_{t+1|0:t}, \Sigma_{t+1|0:t})$ Computed from dynamics step

Updated belief $p(x_{t+1}|z_{0:t+1}, u_{0:t})$

$$= \mathcal{N}(\mu_{t+1|0:t} + K_{t+1}(z_{t+1} - C\mu_{t+1|0:t}), (I - K_{t+1}C)\Sigma_{t+1|0:t})$$

How??

Linear Gaussian Systems: Observations

- Integrate the effect of an observation using sensor model, after dynamics

$$p(x_{t+1}|z_{0:t+1}, u_{0:t}) \propto p(z_{t+1}|x_{t+1})p(x_{t+1}|z_{0:t}, u_{0:t})$$

Stays in Gaussian world, but now conditioning instead of marginalization

Sketch:

1. Construct the joint of x_{t+1} and z_{t+1} conditioned on the past
2. Solve for the mean and covariances of this joint
3. Perform conditioning with z_{t+1} equaling a particular value

Linear Gaussian Systems: Observations

- Integrate the effect of an observation using sensor model, after dynamics

$$p(x_{t+1}|z_{0:t+1}, u_{0:t}) \propto p(z_{t+1}|x_{t+1})p(x_{t+1}|z_{0:t}, u_{0:t})$$

Stays in Gaussian world, but now conditioning instead of marginalization

Linear Gaussian Systems: Observations

- Integrate the effect of an observation using sensor model, after dynamics

$$p(x_{t+1}|z_{0:t+1}, u_{0:t}) \propto p(z_{t+1}|x_{t+1})p(x_{t+1}|z_{0:t}, u_{0:t})$$

Stays in Gaussian world, but now conditioning instead of marginalization

$$X_{t+1|0:t}, Z_{t+1|0:t} \sim \mathcal{N}\left(\begin{bmatrix} \mu^X_{t+1|0:t} \\ \mu^Z_{t+1|0:t} \end{bmatrix}, \begin{bmatrix} \Sigma^{XX}_{t+1|0:t} & \Sigma^{XZ}_{t+1|0:t} \\ \Sigma^{ZX}_{t+1|0:t} & \Sigma^{ZZ}_{t+1|0:t} \end{bmatrix} \right)$$

Belief from the dynamics step

Following the same procedure as last time

$$\begin{bmatrix} \mu^X_{t+1|0:t} \\ \\ \mu^Z_{t+1|0:t} = C\mu^X_{t+1|0:t} \end{bmatrix} \begin{bmatrix} \Sigma^{XX}_{t+1|0:t} & \Sigma^{XZ}_{t+1|0:t} = \Sigma^{XX}_{t+1|0:t}C^T \\ \\ \Sigma^{ZX}_{t+1|0:t} = (\Sigma^{XX}_{t+1|0:t}C^T)^T & \Sigma^{ZZ}_{t+1|0:t} = C\Sigma^{XX}_{t+1|0:t}C^T + R_{t+1} \end{bmatrix}$$

Linear Gaussian Systems: Observations

- Integrate the effect of an observation using sensor model, after dynamics

$$\begin{bmatrix} \mu^X_{t+1|0:t} \\ \\ \mu^Z_{t+1|0:t} = C\mu^X_{t+1|0:t} \end{bmatrix} \begin{bmatrix} \Sigma^{XX}_{t+1|0:t} & \Sigma^{XZ}_{t+1|0:t} = \Sigma^{XX}_{t+1|0:t} C^T \\ \\ \Sigma^{ZX}_{t+1|0:t} = (\Sigma^{XX}_{t+1|0:t} C^T)^T & \Sigma^{ZZ}_{t+1|0:t} = C\Sigma^{XX}_{t+1|0:t} C^T + R_{t+1} \end{bmatrix}$$

$$\mu^Z_{t+1|0:t} = \mathbb{E}\left[Z_{t+1} | X_{t+1}\right]$$

$$\Sigma^{ZZ}_{t+1|0:t} = \mathbb{E}\left[(Z_{t+1} - \mu^Z_{t+1|0:t})(Z_{t+1} - \mu^Z_{t+1|0:t})^T\right]$$

$$\Sigma^{ZZ}_{t+1|0:t} = \mathbb{E}\left[(Z_{t+1} - \mu^Z_{t+1|0:t})(X_{t+1|0:t} - \mu_{t+1|0:t})^T\right]$$

Linear Gaussian Systems: Observations

- Integrate the effect of an observation using sensor model, after dynamics

$$\begin{bmatrix} \mu^X_{t+1|0:t} \\ \mu^Z_{t+1|0:t} = C\mu^X_{t+1|0:t} \end{bmatrix} \quad \begin{bmatrix} \Sigma^{XX}_{t+1|0:t} & \Sigma^{XZ}_{t+1|0:t} = \Sigma^{XX}_{t+1|0:t} C^T \\ \Sigma^{ZX}_{t+1|0:t} = (\Sigma^{XX}_{t+1|0:t} C^T)^T & \Sigma^{ZZ}_{t+1|0:t} = C\Sigma^{XX}_{t+1|0:t} C^T + R_{t+1} \end{bmatrix}$$

Now we just condition on $Z_{t+1} = z_{t+1}$

Remember

$$p(X|Y = y_0) = \mathcal{N}(\mu_X + \Gamma^{-1}_{XX}\Gamma_{XY}(y_0 - \mu_Y), \Gamma^{-1}_{XX})$$

$$= \mathcal{N}(\mu_X + \Sigma_{XY}\Sigma^{-1}_{YY}(y_0 - \mu_Y), \Sigma_{XX} - \Sigma_{XY}\Sigma^{-1}_{YY}\Sigma_{YX})$$

$$p(x_{t+1}|z_{0:t+1}, u_{0:t}) = \mathcal{N}(\mu_{t+1|0:t} + \Sigma_{t+1|0:t}C^T(C\Sigma_{t+1|0:t}C^T + R_{t+1})^{-1}(z_{t+1} - C\mu_{t+1|0:t}),$$
$$\Sigma_{t+1|0:t} - \Sigma_{t+1|0:t}C^T(C\Sigma_{t+1|0:t}C^T + R_{t+1})^{-1}C\Sigma_{t+1|0:t})$$

Linear Gaussian Systems: Observations

- Integrate the effect of an observation using sensor model, after dynamics

$$z_{t+1} = Cx_{t+1} + \delta_t \qquad \delta_t \sim \mathcal{N}(0, R_t)$$

$$p(z_{t+1} | x_{t+1}) = \mathcal{N}(Cx_{t+1}, R_t)$$

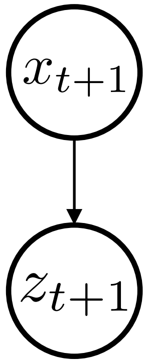

$$\begin{aligned} p(x_{t+1} | z_{0:t+1}, u_{0:t}) &= \mathcal{N}(\mu_{t+1|0:t} + \overbrace{\Sigma_{t+1|0:t} C^T (C \Sigma_{t+1|0:t} C^T + R_{t+1})^{-1}}^{\text{Kalman Gain}} (z_{t+1} - C\mu_{t+1|0:t}), \\ &\qquad \Sigma_{t+1|0:t} - \Sigma_{t+1|0:t} C^T (C \Sigma_{t+1|0:t} C^T + R_{t+1})^{-1} C \Sigma_{t+1|0:t}) \end{aligned}$$

$$K_{t+1} = \Sigma_{t+1|0:t} C^T (C \Sigma_{t+1|0:t} C^T + R_{t+1})^{-1}$$

$$p(x_{t+1} | z_{0:t+1}, u_{0:t}) = \mathcal{N}(\mu_{t+1|0:t} + K_{t+1}(z_{t+1} - C\mu_{t+1|0:t}), (I - K_{t+1}C)\Sigma_{t+1|0:t})$$

Linear Gaussian Systems: Observations

Previous belief $$p(x_{t+1}|z_{0:t}, u_{0:t}) = \mathcal{N}(\mu_{t+1|0:t}, \Sigma_{t+1|0:t})$$ Computed from dynamics step

Updated belief $$p(x_{t+1}|z_{0:t+1}, u_{0:t})$$

$$= \mathcal{N}(\mu_{t+1|0:t} + K_{t+1}(z_{t+1} - C\mu_{t+1|0:t}), (I - K_{t+1}C)\Sigma_{t+1|0:t})$$

Intuition: Correct the update linearly according to measurement error from expectation, shrink uncertainty accordingly

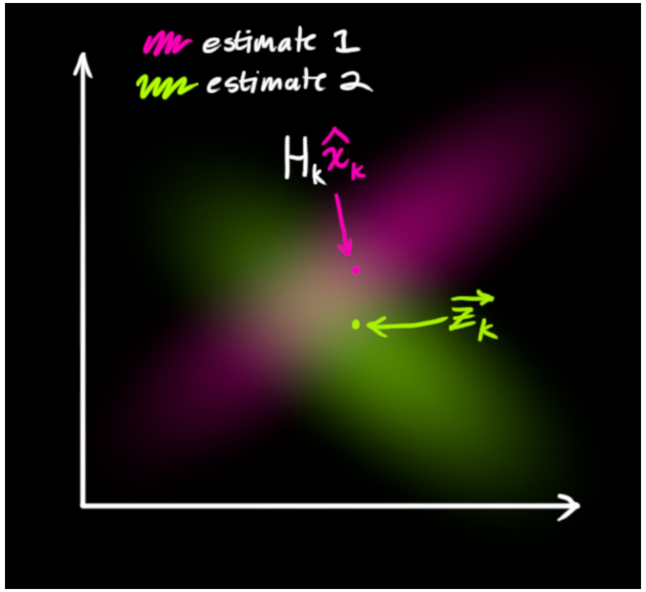

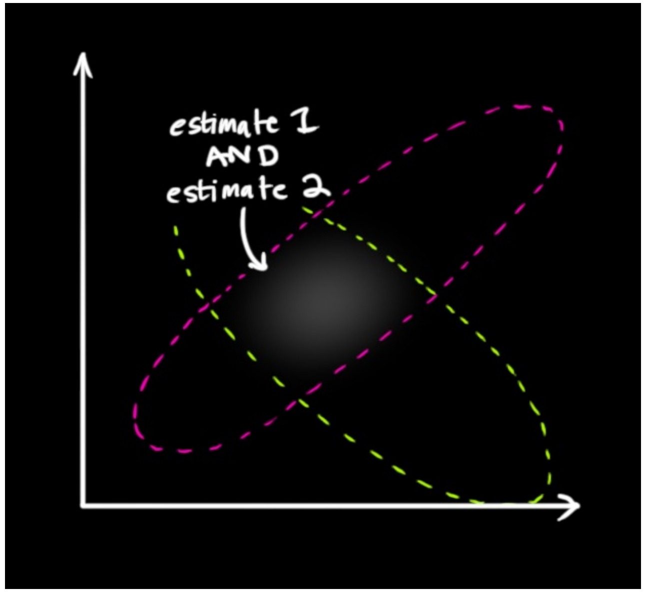

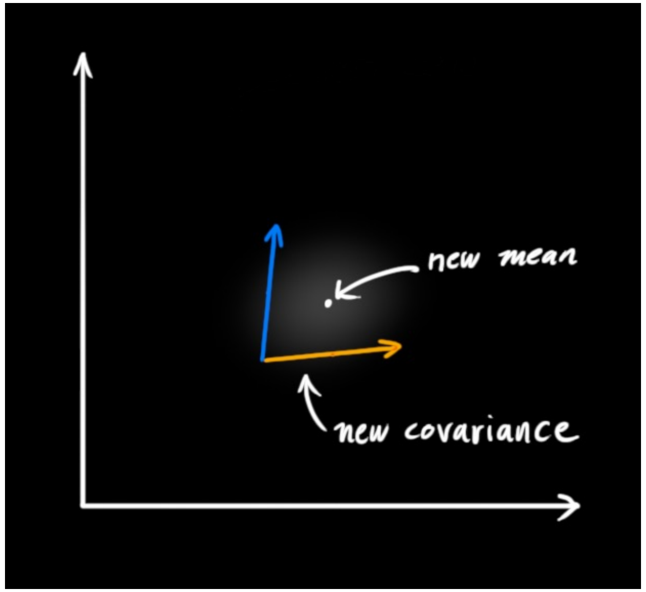

Unpacking the Kalman Gain

Previous belief $$p(x_{t+1}|z_{0:t}, u_{0:t}) = \mathcal{N}(\mu_{t+1|0:t}, \Sigma_{t+1|0:t})$$ Computed from dynamics step

Updated belief $$p(x_{t+1}|z_{0:t+1}, u_{0:t})$$

$$= \mathcal{N}(\mu_{t+1|0:t} + K_{t+1}(z_{t+1} - C\mu_{t+1|0:t}), (I - K_{t+1}C)\Sigma_{t+1|0:t})$$

$$K_{t+1} = \Sigma_{t+1|0:t}C^T(C\Sigma_{t+1|0:t}C^T + R)^{-1}$$

For the sake of simplicity, let's say C = I
in $z_{t+1} = Cx_{t+1} + \delta_t$

$$K_{t+1} = \frac{\Sigma_{t+1|0:t}}{\Sigma_{t+1|0:t} + R}$$

Case 1: Very noisy sensor, R>>Σ

Case 2: Deterministic sensor, R = 0

Intuition Behind Correction Step

- Previous belief (red)
- New Measurement (green)

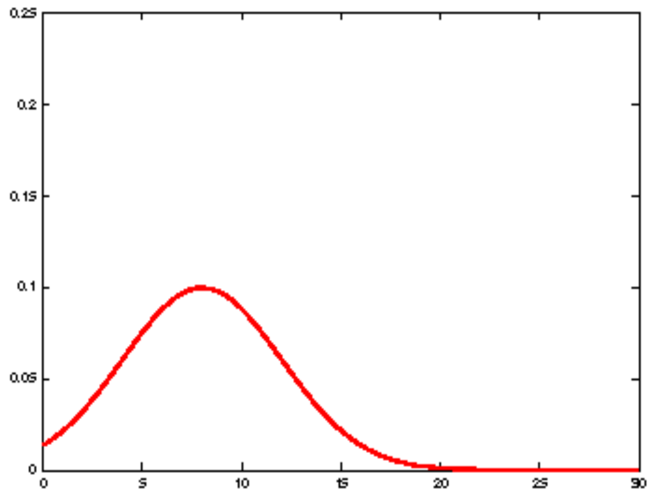

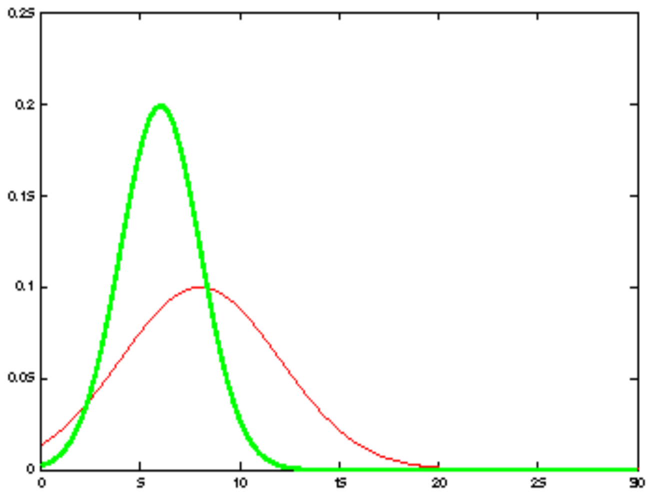

For the sake of simplicity, let's say C = I

$$p(x_{t+1}|z_{0:t+1}, u_{0:t}) = \mathcal{N}(\mu_{t+1|0:t} + K_{t+1}(z_{t+1} - \mu_{t+1|0:t}), (I - K_{t+1})\Sigma_{t+1|0:t})$$

$$K_{t+1} = \frac{\Sigma_{t+1|0:t}}{\Sigma_{t+1|0:t} + R}$$

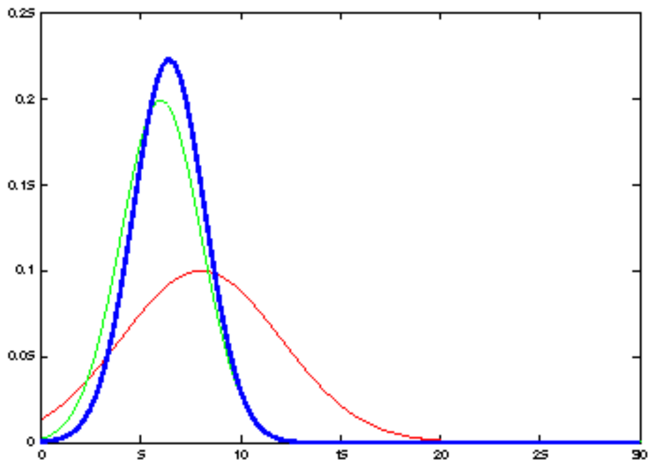

Corrects belief based on measurement
→ Average between mean and measurement based on K
→ Scale down uncertainty based on K

Kalman Filter Pseudocode

1. def Kalman_filter($\mu_{t|0:t}$, $\Sigma_{t|0:t}$, u_t, z_{t+1}):
2. Prediction:

$$\mu_{t+1|0:t} = A\mu_{t|0:t} + Bu_t$$
$$\Sigma_{t+1|0:t} = A\Sigma_{t|0:t}A^T + Q_t$$

3. Correction:

$$K_{t+1} = \Sigma_{t+1|0:t}C^T(C\Sigma_{t+1|0:t}C^T + R_{t+1})^{-1}$$
$$\mu_{t+1|0:t+1} = \mu_{t+1|0:t} + K_{t+1}(z_{t+1} - C\mu_{t+1|0:t})$$
$$\Sigma_{t+1|0:t+1} = (I - K_{t+1}C)\Sigma_{t+1|0:t}$$

4. Return $\mu_{t+1|0:t+1}$, $\Sigma_{t+1|0:t+1}$

$$x_{t+1} = Ax_t + Bu_t + \epsilon_t$$
$$\epsilon_t \sim \mathcal{N}(0, Q)$$

$$z_{t+1} = Cx_{t+1} + \delta_t$$
$$\delta_t \sim \mathcal{N}(0, R)$$

Reminder of the model

Kalman Filter Algorithm

Initial Prior
$p(x_0)$

Dynamics/Prediction
(given some u)

Estimate $\overline{Bel}(x_t)$

$$p(x_{t+1}|z_{0:t}, u_{0:t+1}) \sim \mathcal{N}(A\mu_{t|0:t} + Bu_t, A\Sigma_{t|0:t}A^T + Q_t)$$

Measurement/Correction
(given some z)

Estimate $Bel(x_t)$

$$p(x_{t+1}|z_{0:t+1}, u_{0:t})$$
$$= \mathcal{N}(\mu_{t+1|0:t} + K_{t+1}(z_{t+1} - C\mu_{t+1|0:t}), (I - K_{t+1}C)\Sigma_{t+1|0:t})$$

Kalman Filter in Action

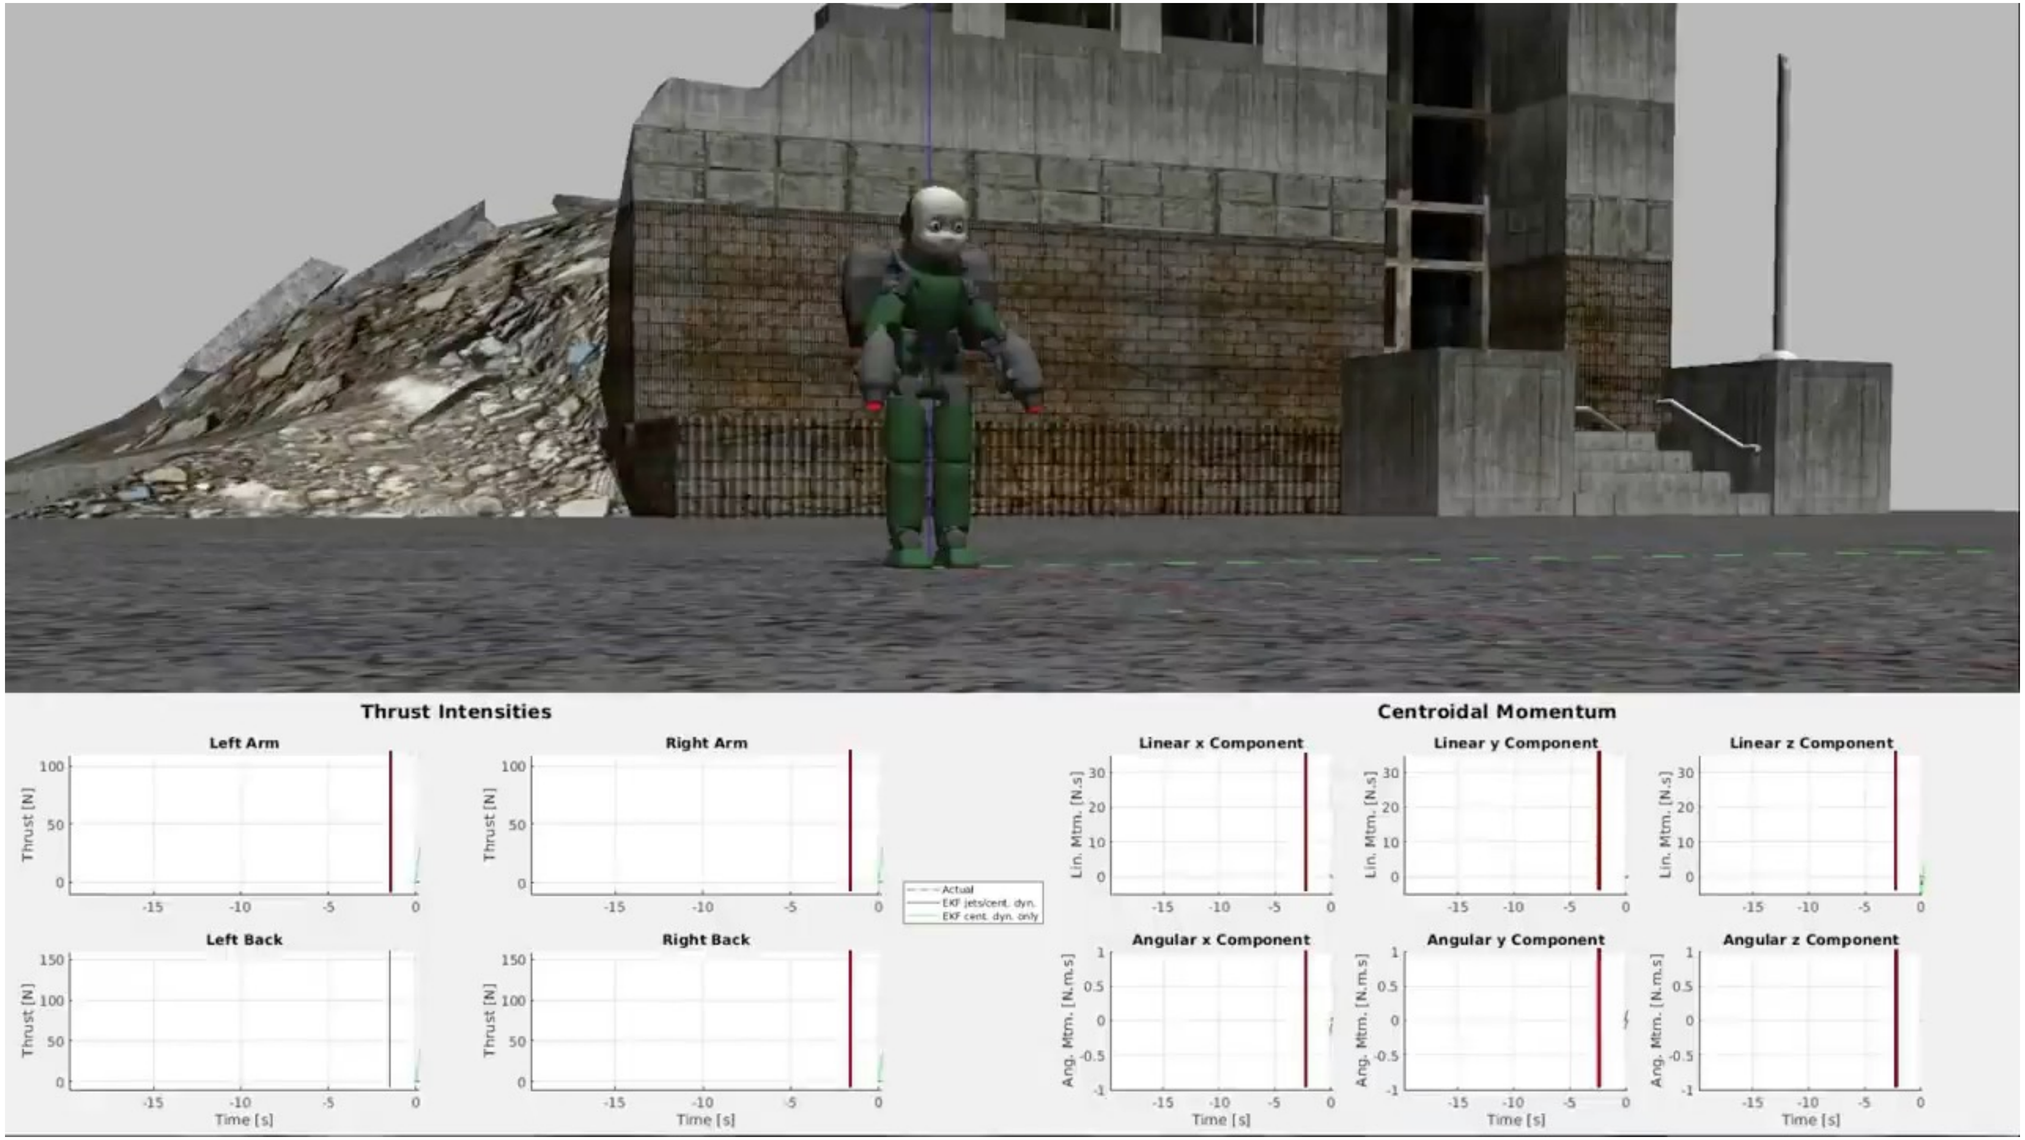

Kalman Filter Summary

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

  $O(k^{2.8} + n^2)$

Matrix Inversion (Correction)

$$K_{t+1} = \Sigma_{t+1|0:t} C^T (C \Sigma_{t+1|0:t} C^T + R_{t+1})^{-1}$$

Matrix Multiplication (Prediction)

$$p(x_{t+1}|z_{0:t}, u_{0:t+1}) \sim \mathcal{N}(A\mu_{t|0:t} + Bu_t, A\Sigma_{t|0:t}A^T + Q_t)$$

- Optimal for linear Gaussian systems!

- Most robotics systems are nonlinear! → next time

Why should we care in 2023?

Still a very widely used technique for estimation/localization/mapping in real problems

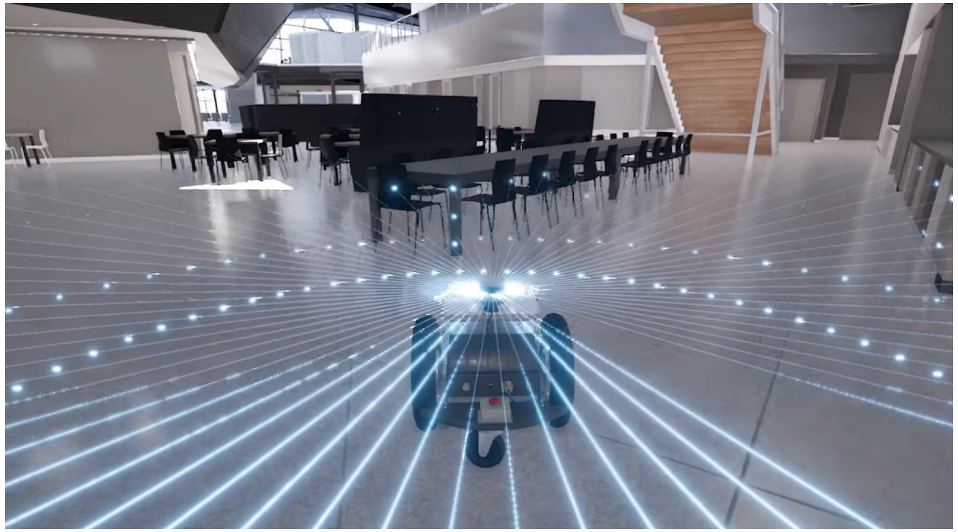

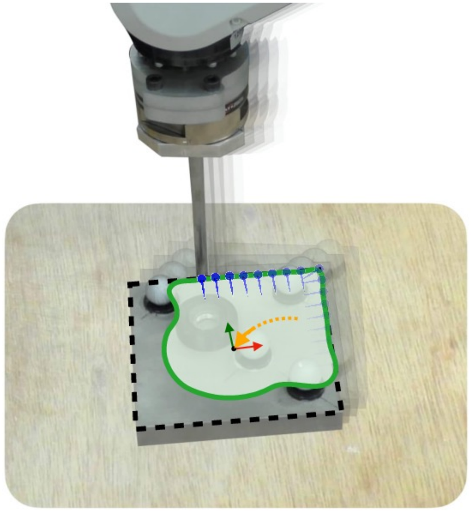

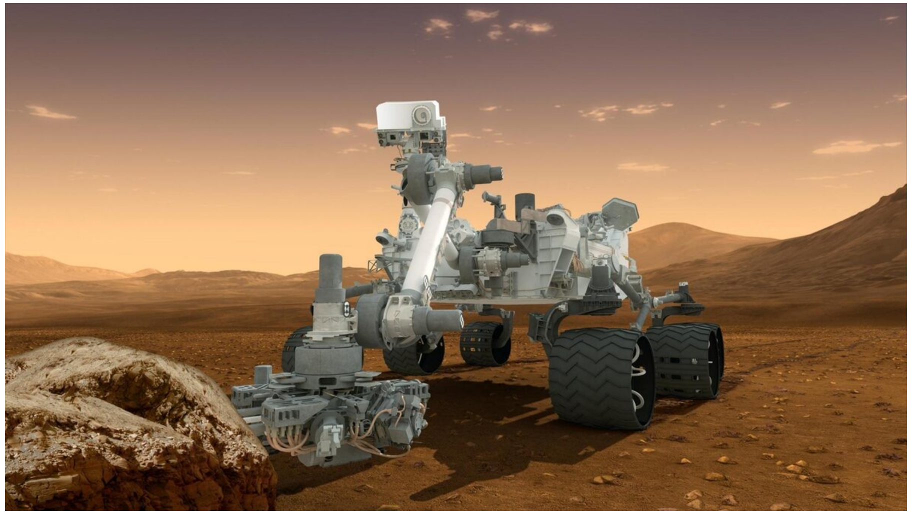

Why should we care in 2023?

Embed to Control: A Locally Li
Dynamics Model for Control from

Manuel Watter* **Jost Tobias Springenberg*** **Martin Riedmiller**
Joschka Boedecker Google DeepMind
University of Freiburg, Germany
{watterm,springj,jboedeck}@cs.uni-freib

Mastering Diverse Domains through World Models

Danijar Hafner[1,2] Jurgis Pasukonis[1] Jimmy Ba[2] Timothy Lillicrap[1]

[1]DeepMind [2]University of Toronto

SOLAR: Deep Structured Representations for Model-Based Reinforcement Learning

Marvin Zhang [*1] **Sharad Vikram** [*2] **Laura Smith** [1] **Pieter Abbeel** [1] **Matthew J. Johnson** [3] **Sergey Levine** [1]

Lecture Outline

Recap

↓

Bayesian Filtering

↓

Gaussian Properties

↓

Kalman Filtering