Autonomous Robotics

Kalman Filtering

Josiah Hanna
University of Wisconsin — Madison

Slides based on Kalman filter slides from “Probabilistic Robotics”

Programming Assignments

- Thoughts?

Learning Outcomes

After today's lecture, you will:

- Be able to specify the key assumptions for Kalman filters.
- Be able to specify the steps of a Kalman filter.
- Gain intuition for how the updates affect beliefs.

Kalman Filter Applications

Robot Localization

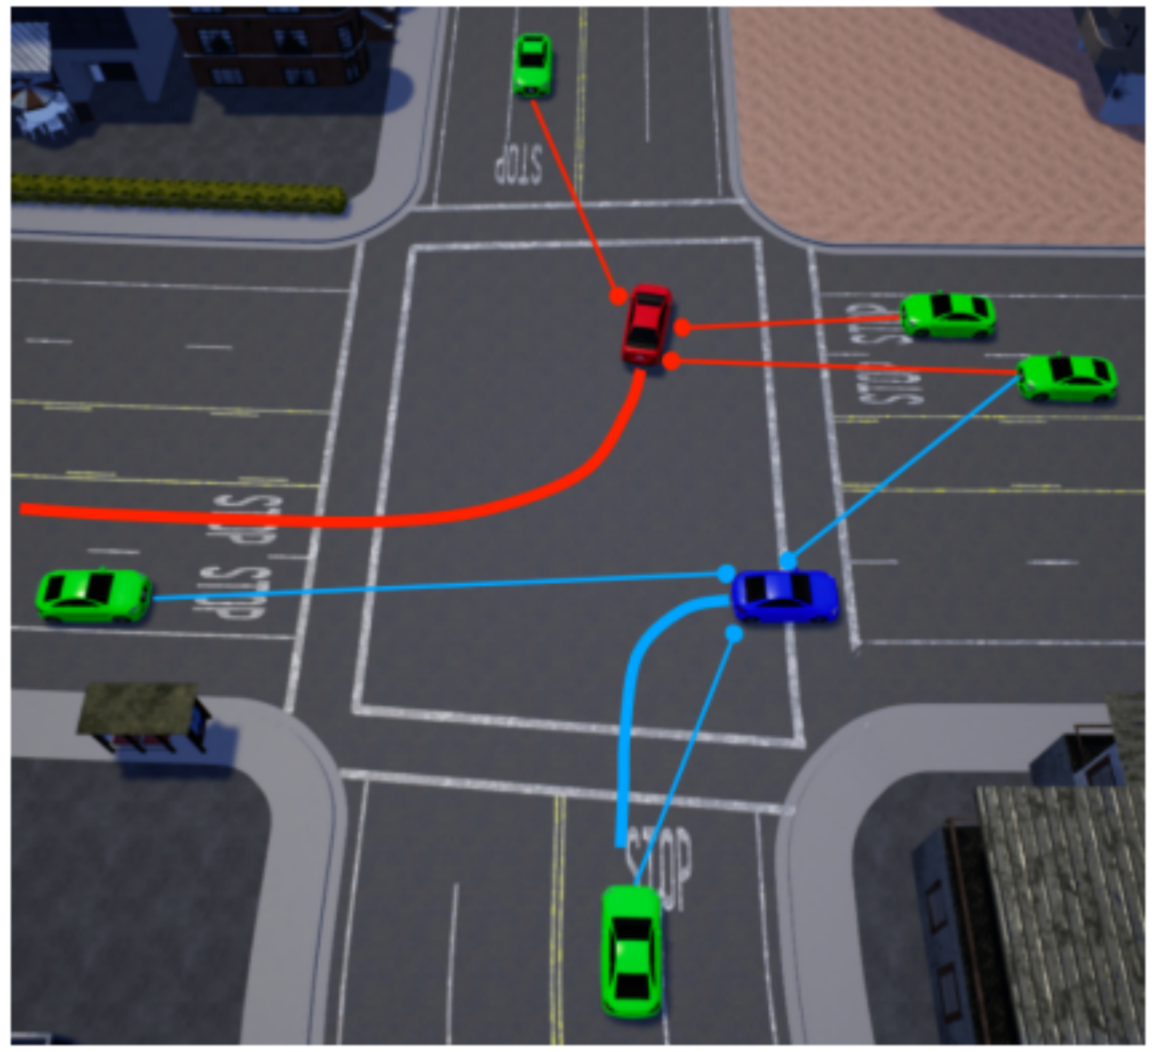

Autonomous driving [e.g., 1]

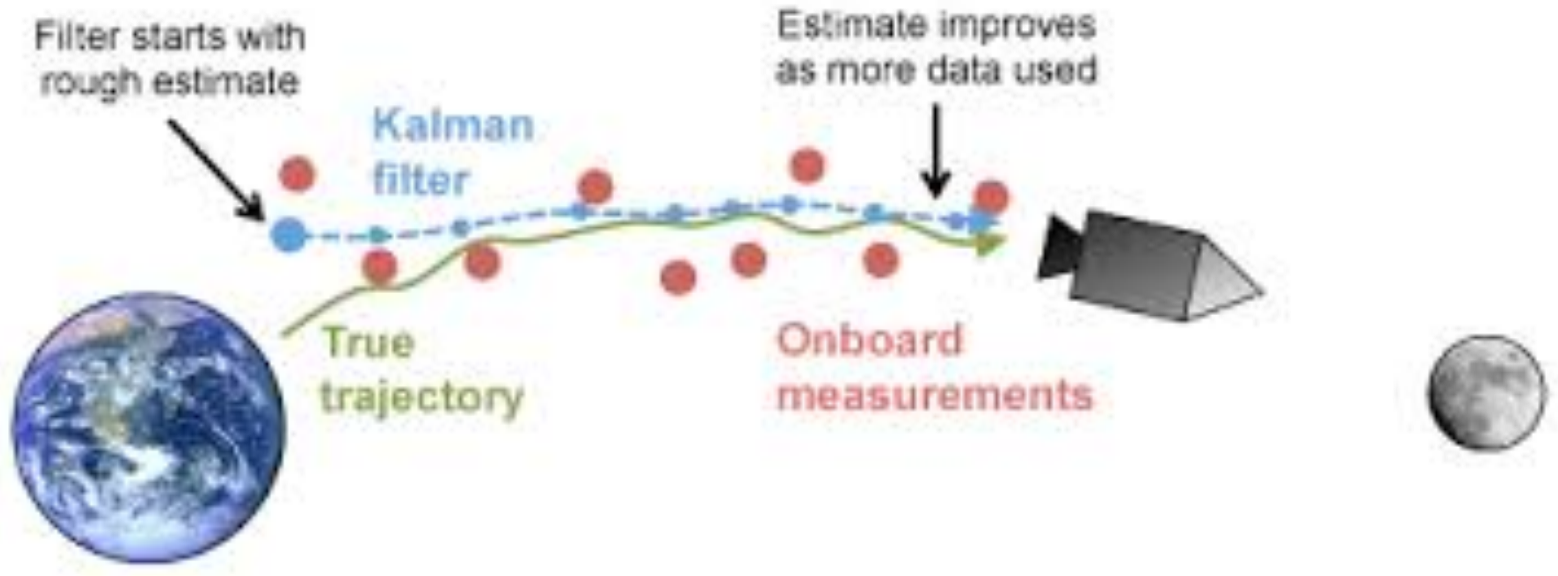

Object Tracking

[1] https://arxiv.org/pdf/2004.05965

Review: Gaussian Distributions

Univariate ($x \in \mathbf{R}$)

$$x \sim \mathcal{N}(\mu, \sigma^2) \qquad p(x) = \frac{1}{\sqrt{2\pi\sigma}} e^{-\frac{(x-\mu)^2}{2\sigma^2}}$$

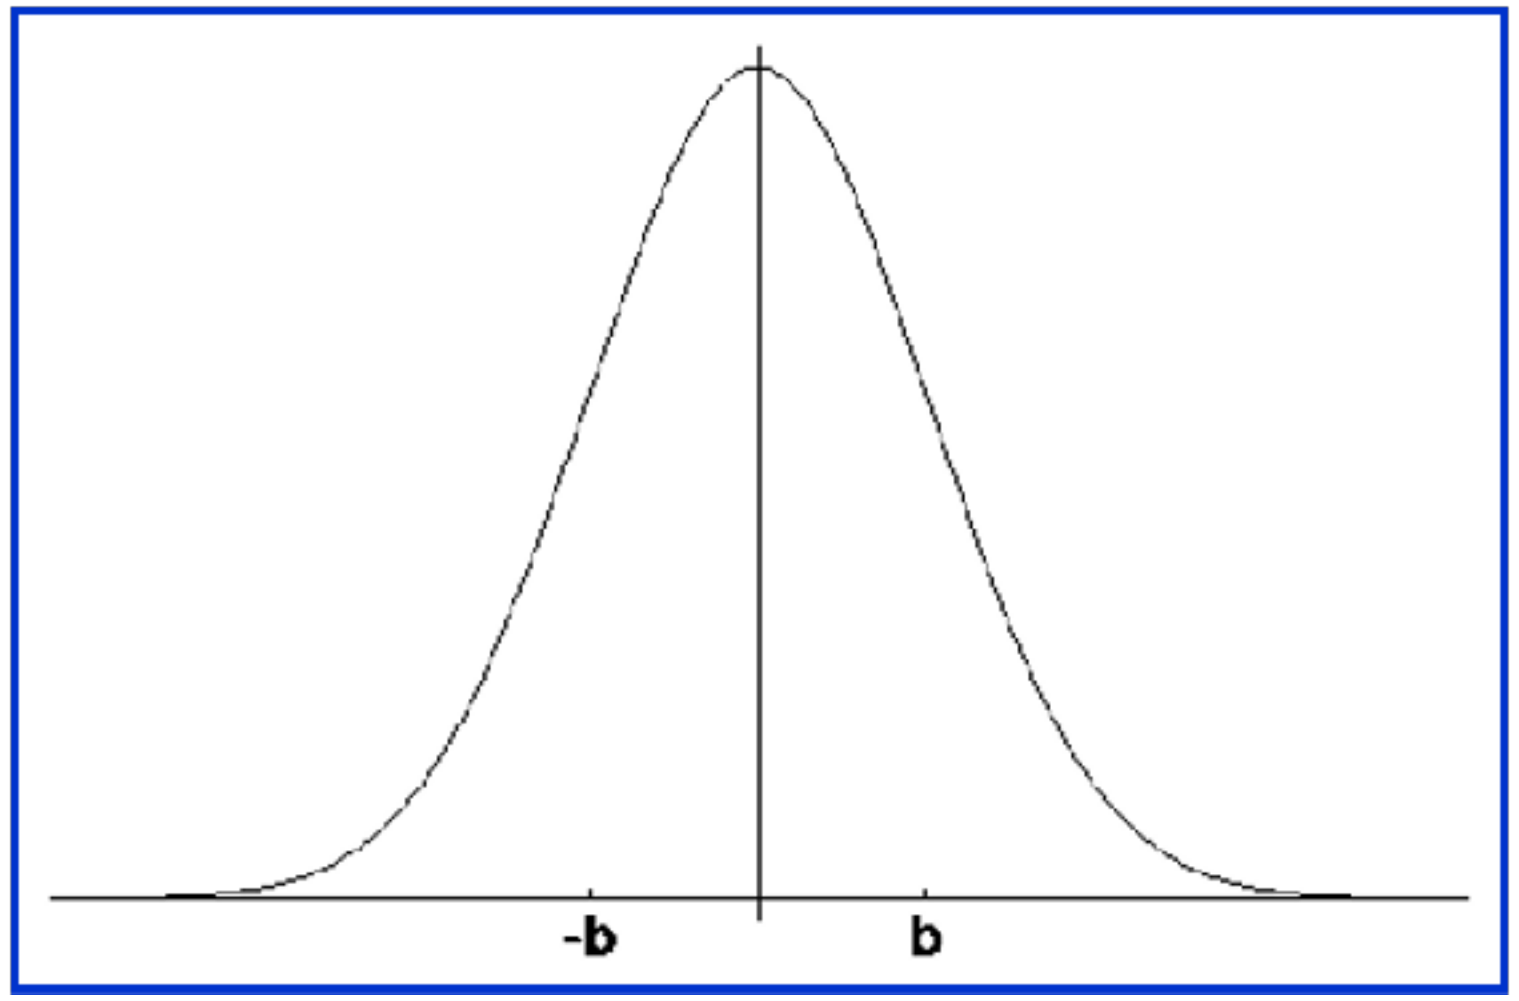

Multivariate ($x \in \mathbf{R}^d$)

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

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

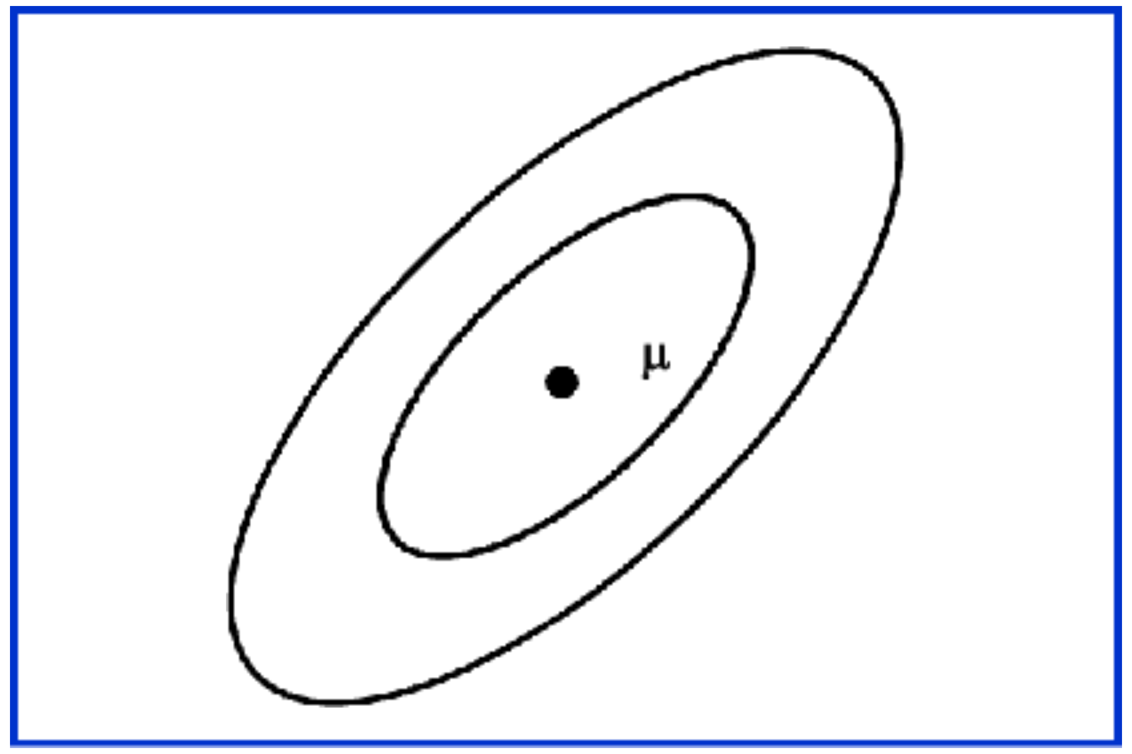

Properties of Gaussians (univariate case)

Fact 1: A linear function of a Gaussian random variable is Gaussian:

$$X \sim \mathcal{N}(\mu, \sigma^2) \text{ and } Y = aX + b \Longrightarrow Y \sim \mathcal{N}(a\mu + b, a^2\sigma^2)$$

Fact 2: If two independent random variables each have a Gaussian distribution, then the product of their distributions is Gaussian:

$$X \sim \mathcal{N}(\mu_1, \sigma_1^2) \text{ and } Y \sim \mathcal{N}(\mu_2, \sigma_2^2) \Longrightarrow p(X)p(Y) = \mathcal{N}(x; \bar{\mu}, \bar{\sigma}^2)$$

$$\bar{\mu} = \frac{\sigma_2^2}{\sigma_1^2 + \sigma_2^2}\mu_1 + \frac{\sigma_1^2}{\sigma_1^2 + \sigma_2^2}\mu_2 \qquad \bar{\sigma} = \frac{1}{\sigma_1^{-2} + \sigma_2^{-2}}$$

Properties of Gaussians (multi-variate case)

Fact 1: A linear function of a Gaussian random variable is Gaussian:

$$X \sim \mathcal{N}(\mu, \Sigma) \text{ and } Y = AX + B \Longrightarrow Y \sim \mathcal{N}(A\mu + B, A^\top \Sigma A)$$

Fact 2: If two independent random variables each have a Gaussian distribution, then the product of their distributions is Gaussian:

$$X \sim \mathcal{N}(\mu_1, \Sigma_1) \text{ and } Y \sim \mathcal{N}(\mu_2, \Sigma_2) \Longrightarrow p(X)p(Y) = \mathcal{N}(x; \bar{\mu}, \overline{\Sigma})$$

$$\bar{\mu} = \frac{\Sigma_2}{\Sigma_1 + \Sigma_2}\mu_1 + \frac{\Sigma_1}{\Sigma_1 + \Sigma_2}\mu_2 \qquad \overline{\Sigma} = \frac{1}{\Sigma_1^{-1} + \Sigma_2^{-1}}$$

Linear Gaussian Systems

We make the following assumptions on the robot's environment:

- States, controls, and observations are vectors: $x \in \mathbf{R}^d$ and $u \in \mathbf{R}^k$ and $z \in \mathbf{R}^m$.
- State transition and observation function are linear Gaussians:
  - $x_t = Ax_{t-1} + Bu_t + w_t$ where $w_t \sim \mathcal{N}(0,Q)$, $A \in \mathbf{R}^{d\times d}$, $B \in \mathbf{R}^{d\times k}$ and $Q \in \mathbf{R}^{d\times d}$. $\implies p(x_t \mid x_{t-1}, u_t) = \mathcal{N}(x; Ax_{t-1} + Bu_t, Q)$
  - $z_t = Hx_t + v_t$ where $v_t \sim \mathcal{N}(0,R)$, $H \in \mathbf{R}^{m\times d}$, and $R \in \mathbf{R}^{m\times m}$.
    $\implies g(z_t \mid x_t) = \mathcal{N}(z; Hx_t, R)$

Example of a Linear Gaussian System

- Consider a robot moving in a 2D plane.
  - State is $x = [x, \dot{x}, y, \dot{y}]$, i.e., position and velocity
  - Action is $u = [\ddot{x}, \ddot{y}]$, i.e., acceleration
  - Observation is noisy position: $z = [\tilde{x}, \tilde{y}]$.
  - $x_t = A_{1,1}x_{t-1} + A_{1,2}\dot{x}_{t-1} + B_{1,1}\ddot{x}_t + w_t(0)$ $\qquad \dot{x}_t = A_{2,2}\dot{x}_{t-1}B_{2,1}\ddot{x}_t + w_t(1)$
  - $\tilde{x}_t = H_{1,1}x_t + v_t(0)$

Similar transition and observation definitions for the y-coordinate.

Kalman Filter

- The Kalman filter is a Bayes filter that represents $\texttt{bel}(x_t)$ with a Gaussian distribution, $\mathcal{N}(\mu_t, \Sigma_t)$.

- The initial belief is Gaussian: $\texttt{bel}(x_0) = \mathcal{N}(x_0; \mu_0, \Sigma_0)$.

- Under our assumptions, the posterior remains a Gaussian distribution using the updates from the Bayes filter:

$$p(x_t \mid z_{1:t}, u_{1:t}) = \mathcal{N}(x_t; \mu_t, \Sigma_t)$$

- Intuition for correctness: plug Gaussian beliefs and linear Gaussian system state transitions and observations into Bayes filter updates.

The Kalman Filter as a Bayes Filter

- Initialize belief:

$$\texttt{bel}(x_0) = \mathcal{N}(x_0, \mu_0, \Sigma_0)$$

- Prediction:

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

$$\bar{\mu}_t = A\mu_{t-1} + Bu_t$$

$$\bar{\Sigma}_t = A^T \Sigma A + R$$

- Correction:

$$\texttt{bel}(x_t) = \eta g(z_t \mid x_t)\overline{\texttt{bel}}(x_t)$$

$$\mu_t = \bar{\mu}_t + K_t(z_t - H\bar{\mu}_t)$$

$$\Sigma_t = (I - K_t H)\bar{\Sigma}_t$$

The Kalman Gain

$$K_t = \bar{\Sigma}_t H^\top (H \bar{\Sigma}_t H^\top + R)^{-1}$$

- K_t is called the Kalman gain at time-step t.

- Use univariate case with $H = 1$ to build intuition:

$$K_t = \frac{\bar{\sigma}_t^2}{\bar{\sigma}_t^2 + R}$$

Uncertainty from prediction step

Total uncertainty

- The Kalman gain tells you how much to trust the prediction vs the observation.

- Small gain implies the measurement is less reliable and the belief is updated less from the prediction belief.

Illustration of Kalman Filter Updates

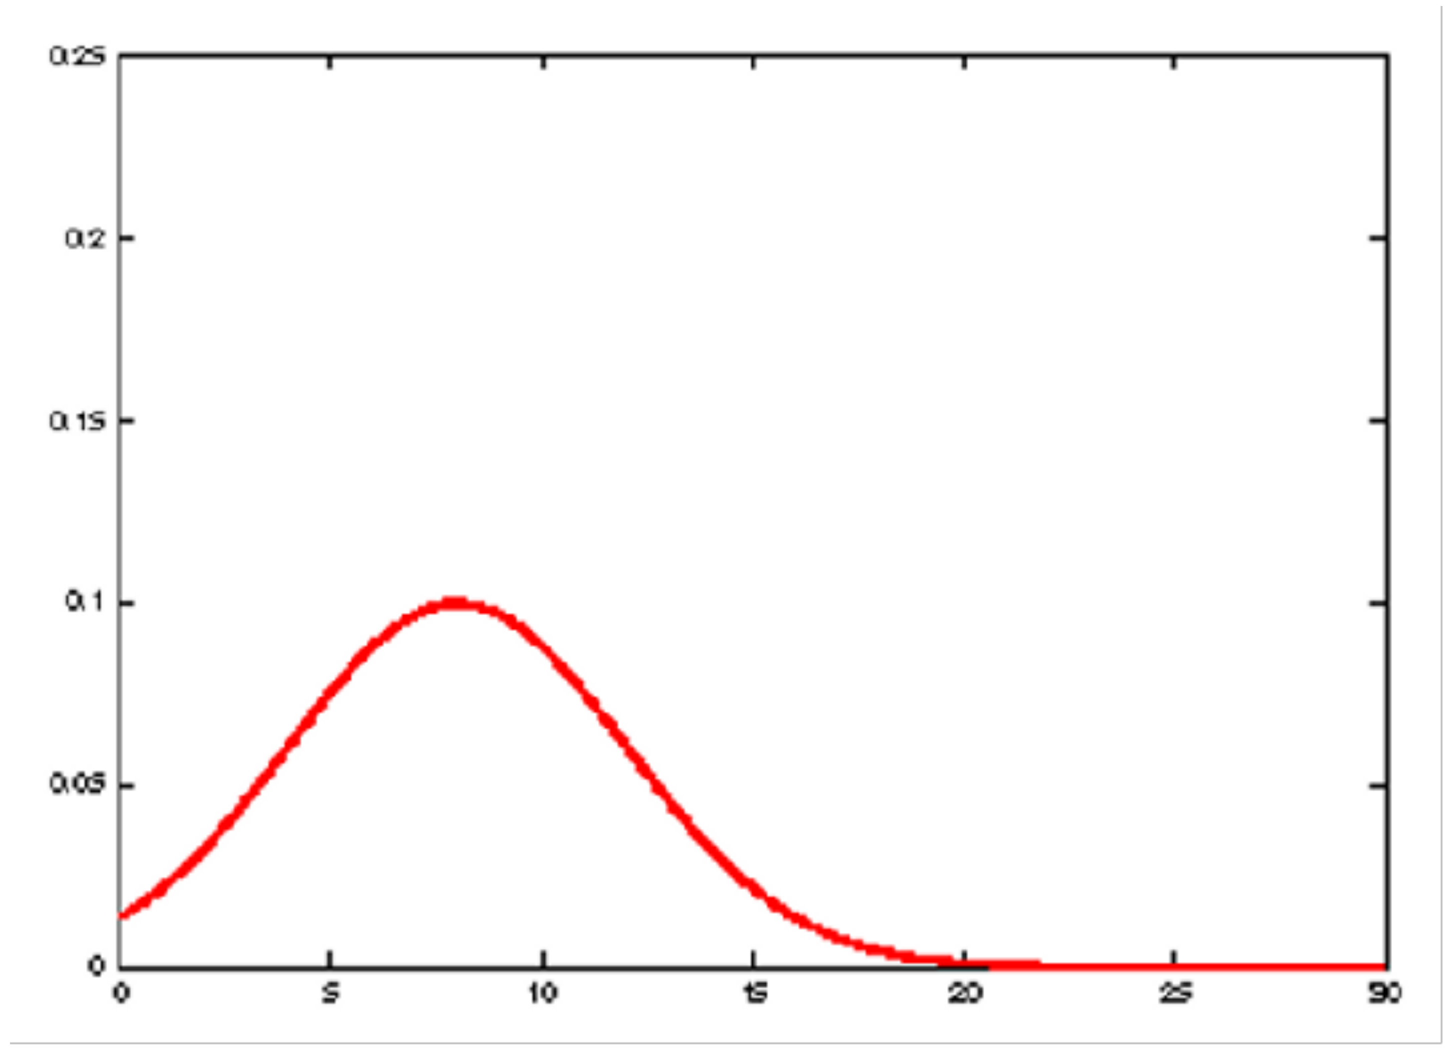

Belief after motion

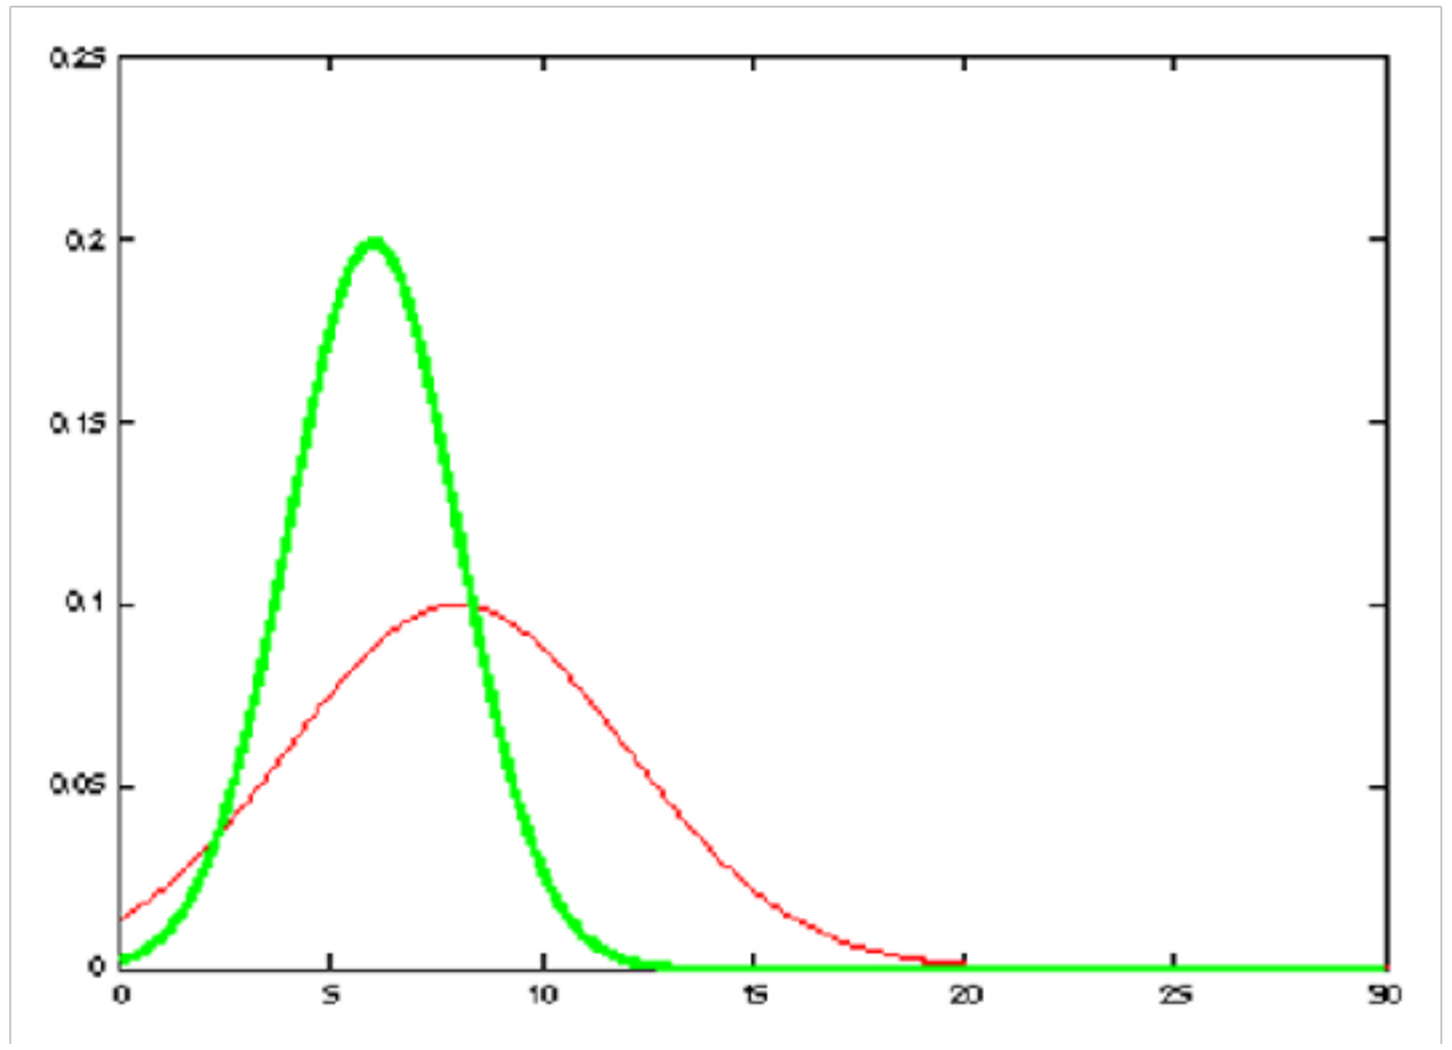

Observation Probability

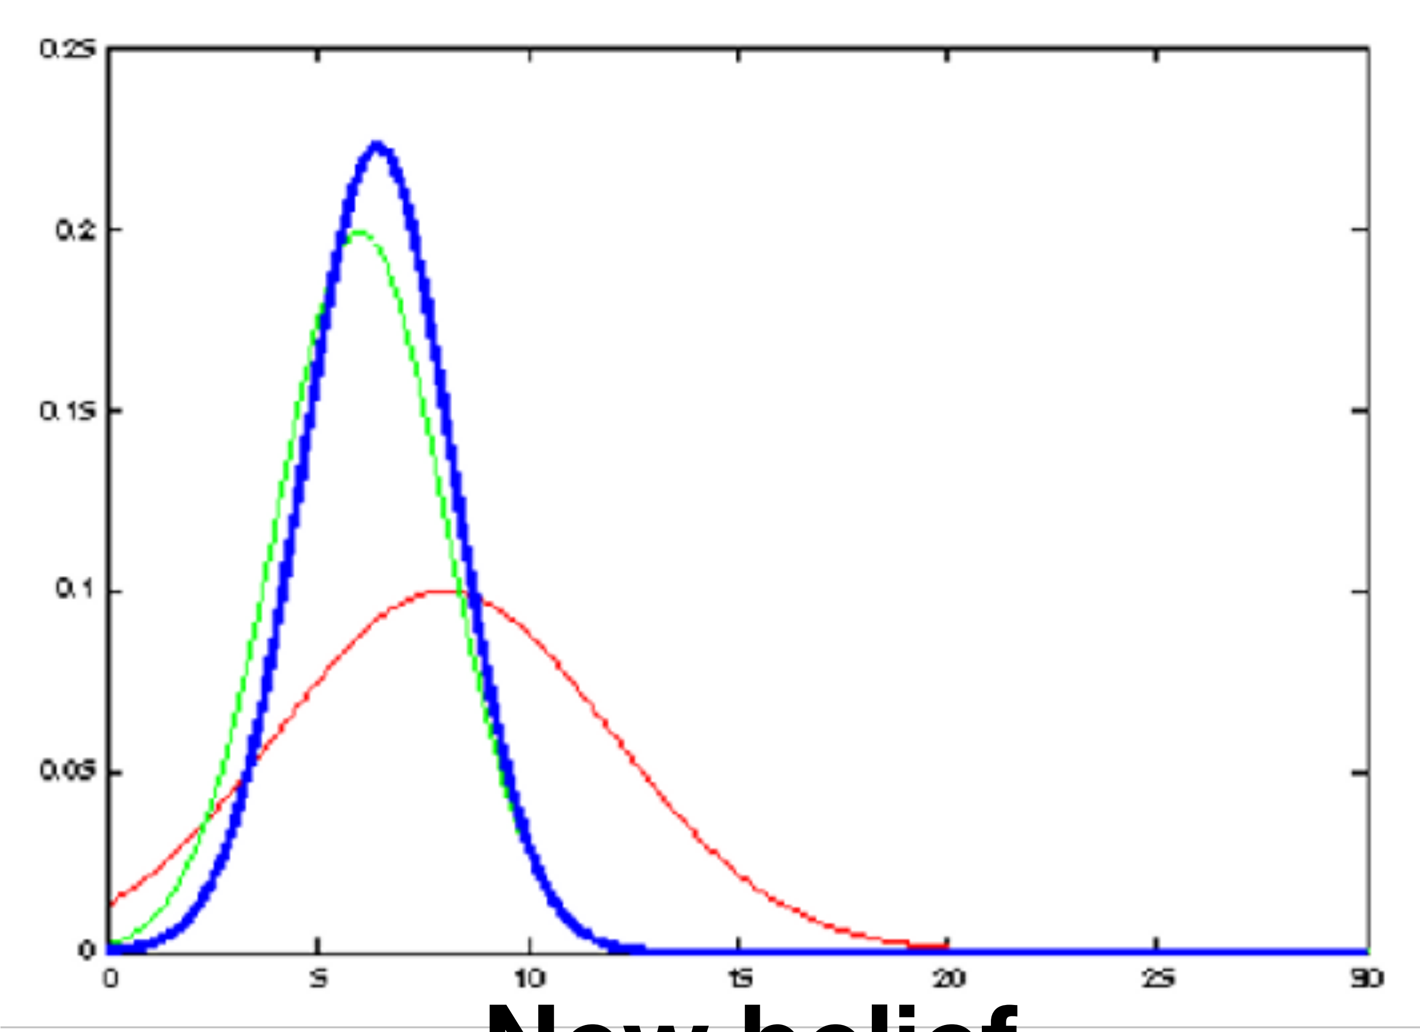

New belief

Advantages / Disadvantages

- Kalman filters:
  - Can be used for continuous state spaces.
  - Are optimal filters if our assumptions hold.
  - Are very efficient; polynomial in state and observation dimensionality.
- But…
  - Randomness may not be Gaussian.
  - Most robotics systems are nonlinear.

Practice

- Robot is moving along the x-axis and has state given by its x-coordinate. It's action is a desired velocity and it observes a noisy observation of its coordinate. The initial belief is $\mu_0 = 0$ and $\sigma_0 = 1$.

$$x_t = x_{t-1} + u_t + w_t \textbf{ where } w_t \sim \mathcal{N}(0,1)$$

$$z_t = x_t + v_t \textbf{ where } v_t \sim \mathcal{N}(0,2)$$

Compute the robot's belief about its location after it takes action $u_1 = 1$ and observes $z_t = 2$

Practice

- Robot is moving along the x-axis and has state given by its x-coordinate. It's action is a desired velocity and it observes a noisy observation of its coordinate. The initial belief is $\mu_0 = 0$ and $\sigma_0 = 1$.

$$x_t = x_{t-1} + u_t + w_t \textbf{ where } w_t \sim \mathcal{N}(0,1)$$

$$z_t = x_t + v_t \textbf{ where } v_t \sim \mathcal{N}(0,2)$$

Compute the robot's belief about its location after it takes action $u_1 = 1$ and observes $z_t = 2$

Prediction

$$\bar{\mu}_t = A\mu_0 + Bu_t = \mu_0 + u_t = 0 + 1 = 1$$

$$\bar{\Sigma}_1 = A^\top \Sigma_0 A + Q = (1)(1)(1) + 1 = 2$$

Correction

$$K_1 = \frac{\bar{\sigma}_1^2}{\bar{\sigma}_1^2 + R} = \frac{2}{2+2} = 1/2$$

$$\mu_1 = \bar{\mu}_0 + K_1(z - Hx_1) = 1 + \frac{1}{2}(2-1) = 3/2$$

$$\Sigma_1 = (I - K_1 H)\bar{\Sigma}_1 = (1 - 1/2)2 = 1$$

Summary

- Introduced the linear Gaussian model.
- Introduced the basic Kalman filter as an instantiation of the Bayes filter under a linear Gaussian assumption.
- Saw an example of how updates change the belief.

Action Items

- Read on particle filter for next week; send a reading response by 12 pm on Monday.