CS 3630!

Lecture 13:
A Logistics Robot:
Sensing

Logistics Robots

Sensing

So far, we’ve seen pretty simple sensor models:

- Discrete measurements (conductivity, light)
- Univariate Gaussians (weight/scale)

In this chapter, we’ll see more realistic sensor models:

- Proximity (object detection, binary)
- Range (distance to a beacon, Gaussian)
- Pseudo-GPS (2D coordinates, bi-variate Gaussian)

For Perception, we’ll require more sophisticated computational tools that exploit efficient and effective approximation schemes.

Warehouse Environment

- Sensors measure various features of the environment.
  - Geometric aspects of the environment (e.g., location of obstacles)
  - Artifacts placed in the environment (e.g., QR Code, RFID transmitters, GPS)
  - Visual features in the environment.

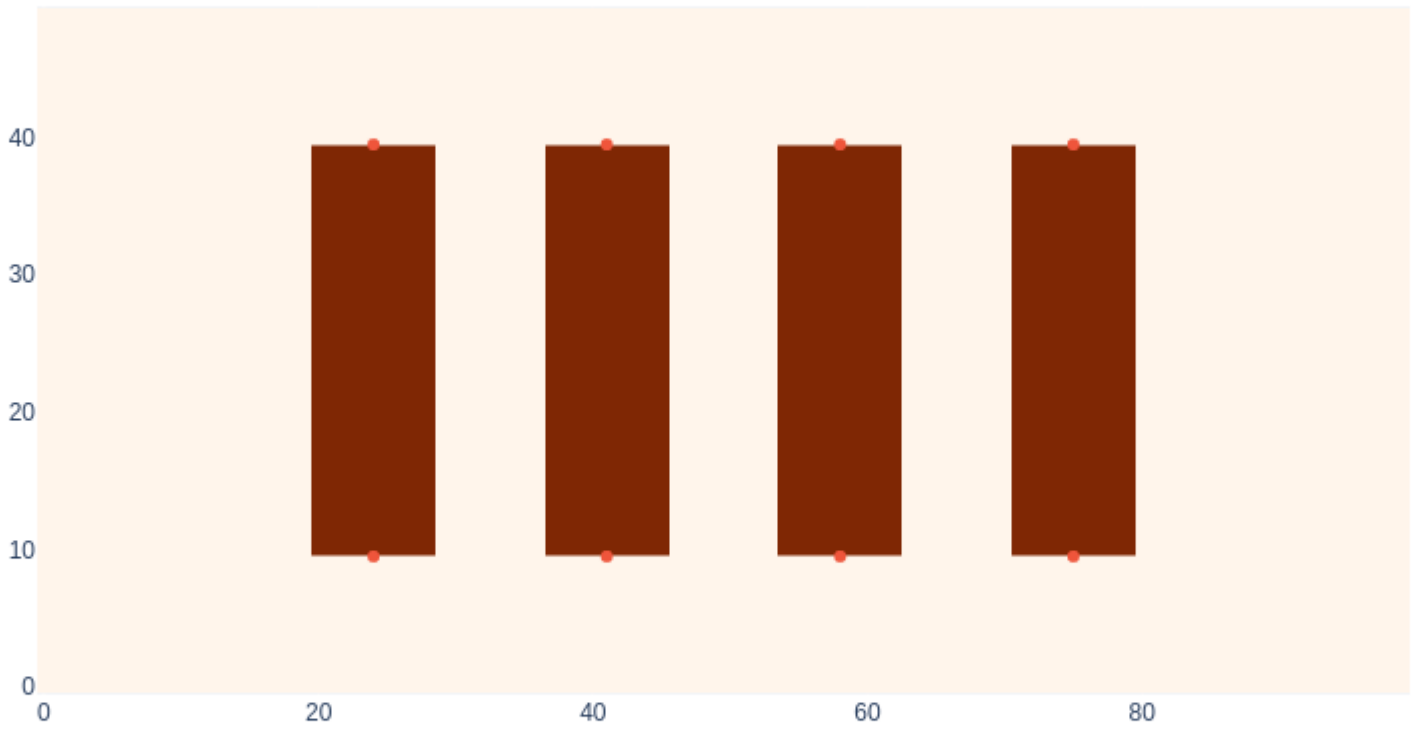

Environment:

- Warehouse is an enclosed 100x50m space
- Four shelving units
- Eight beacons (for range sensor)

Sensors:

- Proximity sensor detects walls and shelves
- Range sensor measures distance to the nearest beacon
- Pseudo-GPS sensor gives 2D coordinates of the robot in the warehouse.

An Ideal Proximity Sensor

- Binary sensor that detects obstacles.
- Sensor returns measurement $z_k \in \{ON, OFF\}$
- Denote by X_{obs} the obstacle region (includes shelves and walls)
- Distance to nearest obstacle is defined by

$$d(x) = \min_{x' \in X_{obs}} \|x - x'\|^{\frac{1}{2}}$$

- If $d(x_k) \leq d_0$ (for some predetermined distance d_0), the sensor triggers:

$$z_k = ON$$

- If $d(x_k) > d_0$, $z_k = OFF$.

Ideal Proximity Sensor

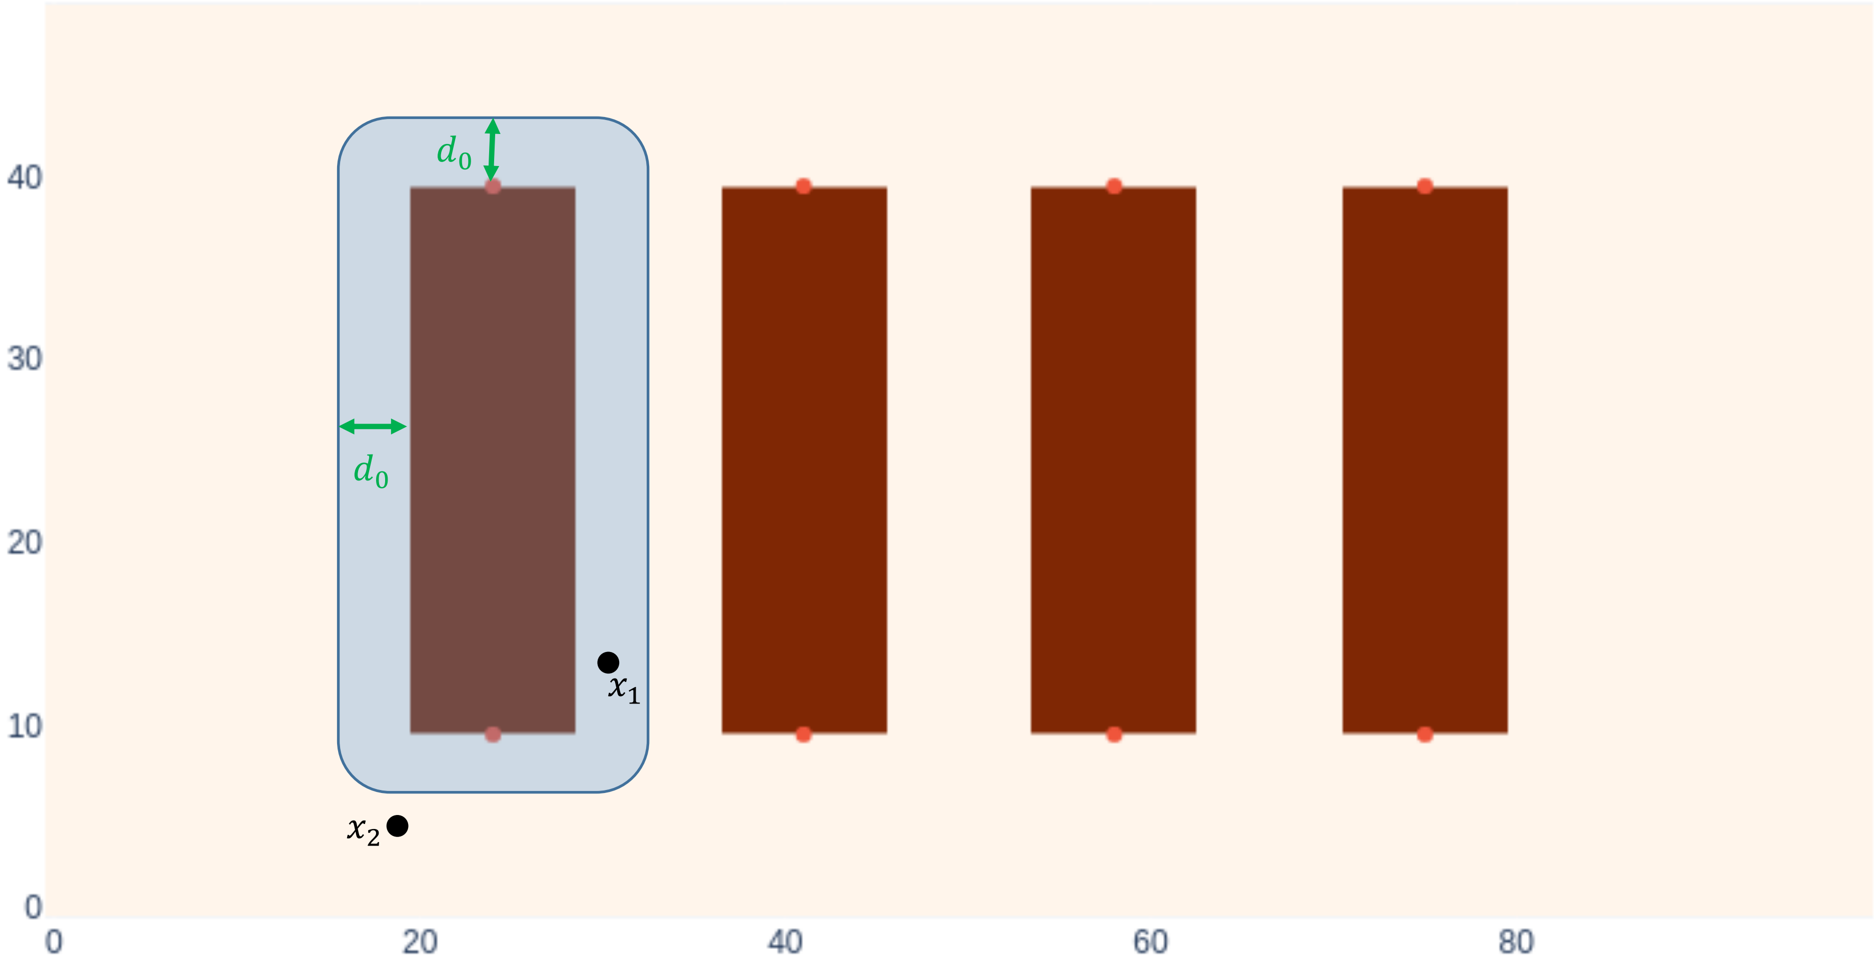

In this example,

- $z_1 = ON$
- $z_2 = OFF$

- This figure illustrates how the proximity sensor works for one of the shelves.
- Similar blue regions exist for all four shelves and the four walls.

A Noisy Proximity Sensor

- Real proximity sensors have variations in their ability to detect obstacles.
- Often the ability to detect obstacles degrades as the distance to the obstacle increases.
- One possible measurement model:

$$P(z_k = ON \mid x_k) = \begin{cases} 1 & d(x_k) \leq d_0 \\ e^{-\alpha(d(x_k)-d_0)} & otherwise \end{cases}$$

- This sensor gives a false positive with probability that exponentially decreases to zero as the robot moves away from an obstacle.
- Since $P(Z_k | x_k)$ is a conditional probability, it follows that

$$P(z_k = OFF \mid x_k) = 1 - P(z_k = ON \mid x_k)$$

Likelihood for the Ideal Proximity Sensor

- We can model an ideal proximity sensor using the measurement model:

$$P(z_k = ON \mid x_k) = \begin{cases} 1 & d(x_k) \leq d_0 \\ 0 & otherwise \end{cases}$$

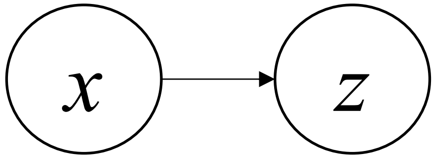

- The ***likelihood*** for this sensor is given by:

$$\mathcal{L}(x_k; z_k = ON) = \begin{cases} 1 & d(x_k) \leq d_0 \\ 0 & otherwise \end{cases}$$

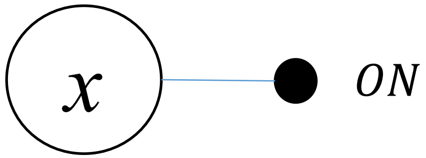

$$\mathcal{L}(x_k; z_k = OFF) = \begin{cases} 0 & d(x_k) \leq d_0 \\ 1 & otherwise \end{cases}$$

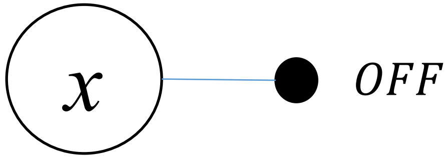

Likelihood for the Ideal Proximity Sensor

- We can plot the ***likelihood*** for for each possible value of z_k.

$$\mathcal{L}(x_k; z_k = ON) = \begin{cases} 1 & d(x_k) \leq d_0 \\ 0 & otherwise \end{cases}$$

$$\mathcal{L}(x_k; z_k = OFF) = \begin{cases} 0 & d(x_k) \leq d_0 \\ 1 & otherwise \end{cases}$$

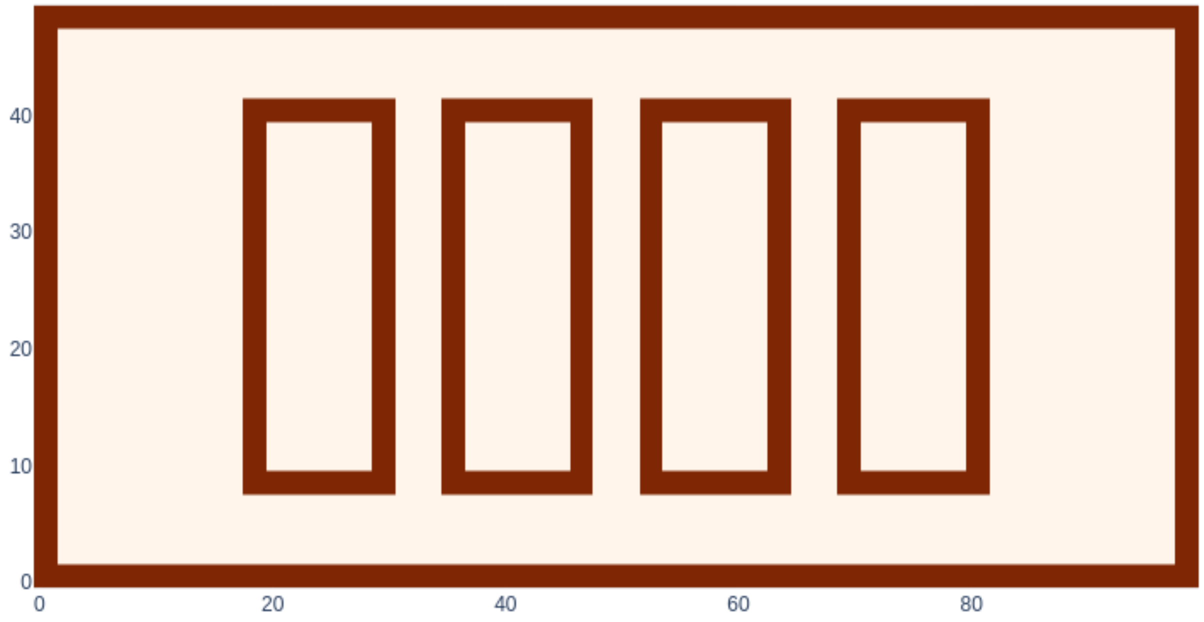

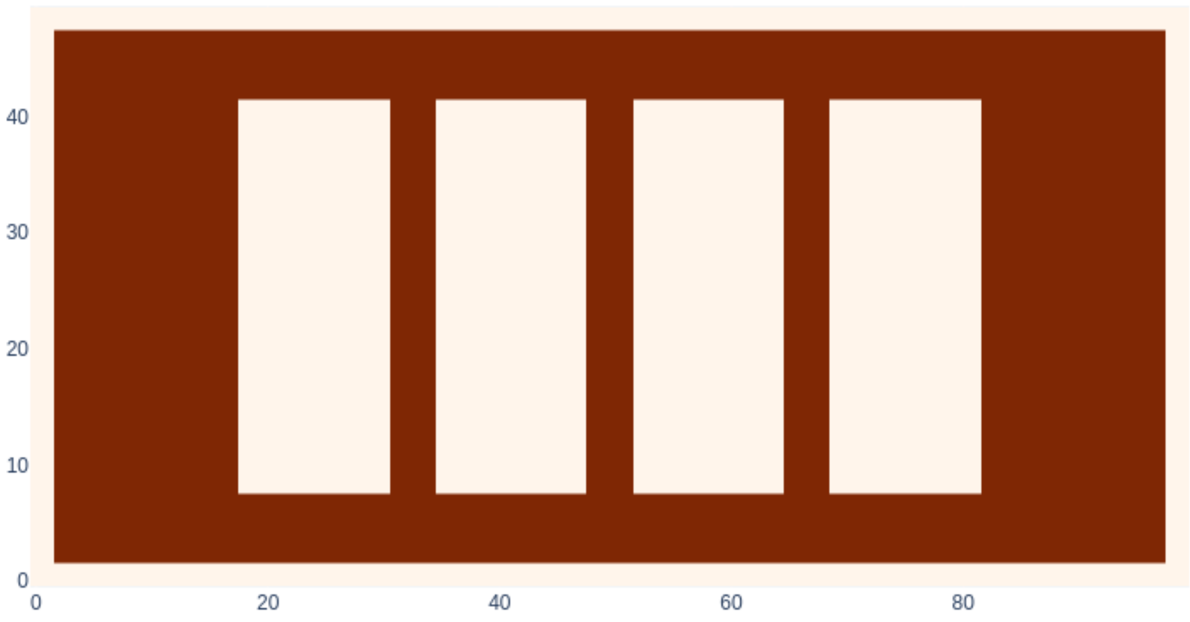

- **The likelihood is a function of x_k. It is *not* a probability distribution!**
- **The specific form of the likelihood depends on which value of z_k was observed.**

An Ideal Range Sensor

- Eight beacons have been placed in the warehouse, at locations $b_0, \ldots, b_7$.
- The range sensor is a nonlinear sensor that returns the distance to the beacons:

$$h(x_k; b_i) = \|x_k - b_i\| = \sqrt{(x_k - b_i)^T (x_k - b_i)}$$

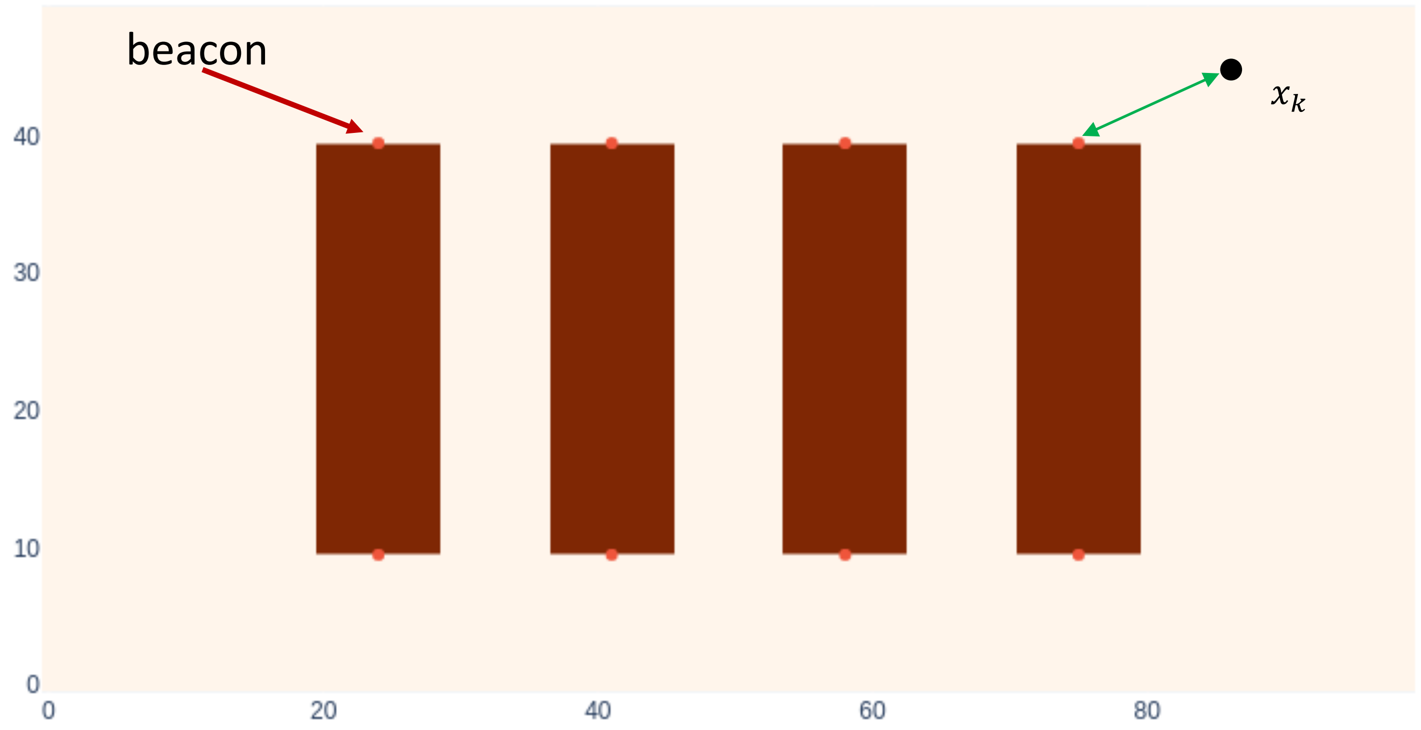

- This sensor can be realized using RFID technology.
- Of course the beacon range is finite, so when $\|x_k - b_i\| > d_{\max}$ for all i, we set

$$h(x_k; b_i) = \mathbf{inf}$$

A Noisy Range Sensor

- We often assume that sensor measurements are corrupted by additive noise. In this case, our range sensor returns a noisy measurement:

$$z_k = h(x_k; b_i) + w_k = \|x_k - b_i\| + w_k$$

in which w_k is the noise term.

- We'll assume i.i.d. zero-mean Gaussian noise, $f_{W_k}(w_k) = \frac{1}{\sigma\sqrt{2\pi}} e^{-\frac{w_k^2}{2\sigma^2}}$
- The resulting conditional pdf for the measurement (given x_k and b_i) is given by

$$f_{Z_k}(z_k | x_k, b_i) = \frac{1}{\sigma\sqrt{2\pi}} e^{-\frac{\left(z_k - h(x_k; b_i)\right)^2}{2\sigma^2}}$$

➢ Given the state and the beacon ID, the range measurement is a Gaussian r.v. whose mean is equal to the true range.

Measurement Model

- The sensor measurement model is a conditional pdf:

$$f_{Z_k}(z_k|x_k, b_i) = \frac{1}{\sigma\sqrt{2\pi}} e^{-\frac{\left(z_k - h(x_k; b_i)\right)^2}{2\sigma^2}}$$

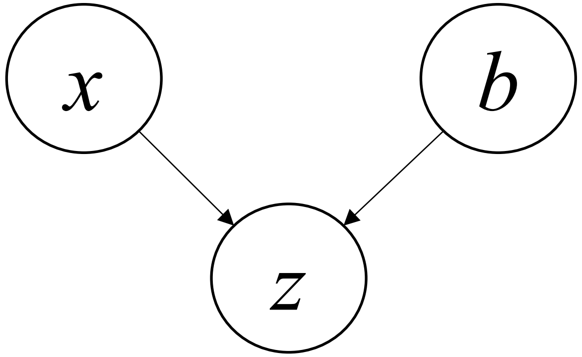

- This pdf describes the behavior the a r.v. z_k when x_k and b_i are known.
- As such, we can expect f_{Z_k} to behave like any other pdf, e.g.,

$$\int_{-\infty}^{\infty} f_{Z_k}(z_k|x_k, b_i) dz_k = 1$$

Measurement Likelihood

- The measurement ***likelihood*** is a function of x_k

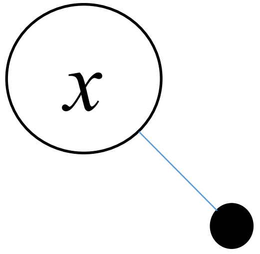

$$\mathcal{L}(x_k; z_k, b_i) = \frac{1}{\sigma\sqrt{2\pi}} e^{-\frac{\left(z_k - h(x_k; b_i)\right)^2}{2\sigma^2}}$$

- This likelihood is ***not*** a probability. For example,

$$\int_{-\infty}^{\infty} \mathcal{L}(x_k; z_k, b_i) dx_k \neq 1$$

- The likelihood tells us something about how likely it would be to see various values for x_k, but it does not tell us probabilities.

Measurement Likelihood

- For a given measurement z_k and specific beacon b_i, we can plot the likelihood function on our warehouse map.
- For the case $b_i = b_0$ and $z_k = 4.03$, we obtain the plot for $\mathcal{L}(x_k; 4.03, b_0)$ shown below (also shown in the book).

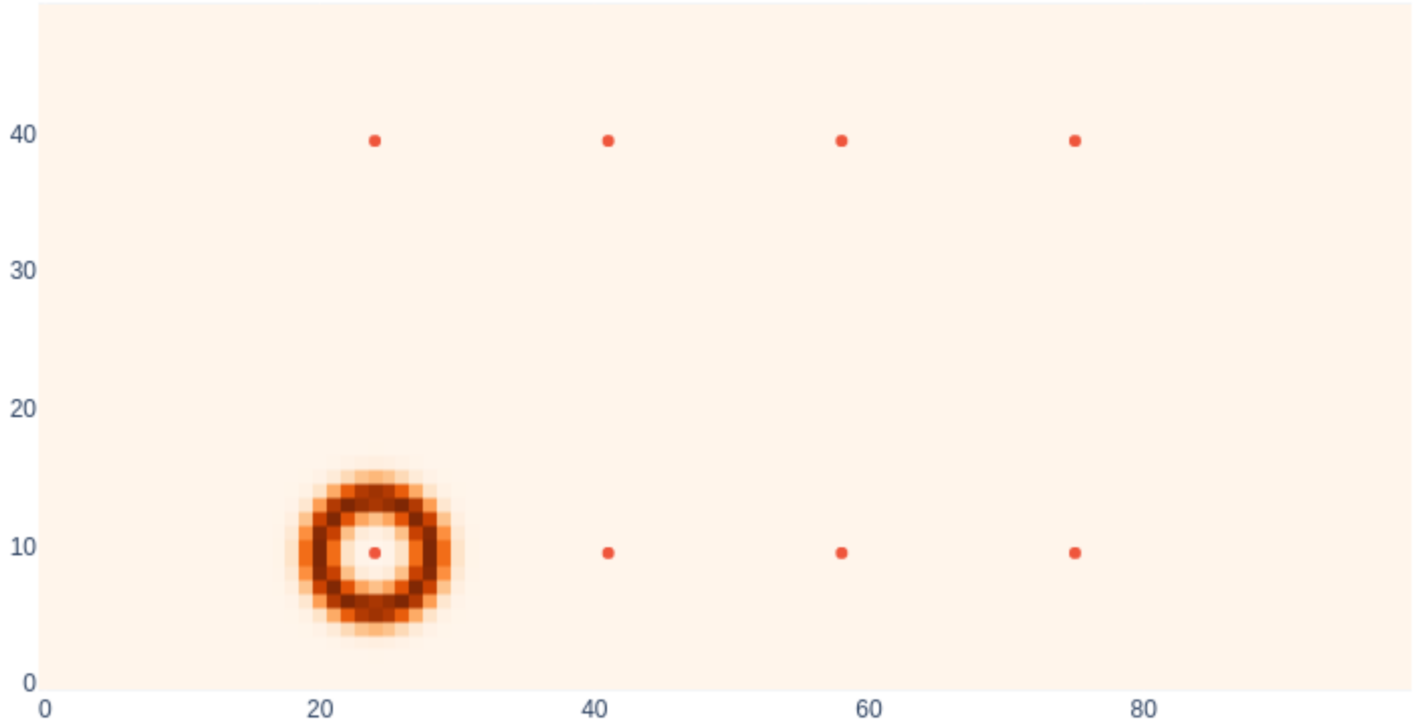

- The likelihood achieves its maximum on the circle of radius 4.03, centered on beacon b_0.
- The value of $\mathcal{L}(x_k; 4.03, b_0)$ looks like a Gaussian curve along any radial line extended from beacon b_0.

Out-of-range Measurements

- Clearly the range sensor provides valuable information when it is able to return a distance to a specific beacon.
- Suppose all beacons are out of range, i.e., $\|x_k - b_i\| > d_{max}$ for all i, and therefore $h(x_k; b_i) = \mathbf{inf}$.
- If we assume that the cutoff at d_{max} is sharp (a nice assumption mathematically, even if it is unrealistic in practice), we can construct a likelihood for this case: $\mathcal{L}(x_k; z_k = \text{inf}, b_i = \text{NONE})$

$$\mathcal{L}(x_k; z_k = \text{inf}, b_i = \text{NONE}) = \begin{cases} 1 & h(x_k; b_i) > d_{max}, i = 0 \dots 7 \\ 0 & otherwise \end{cases}$$

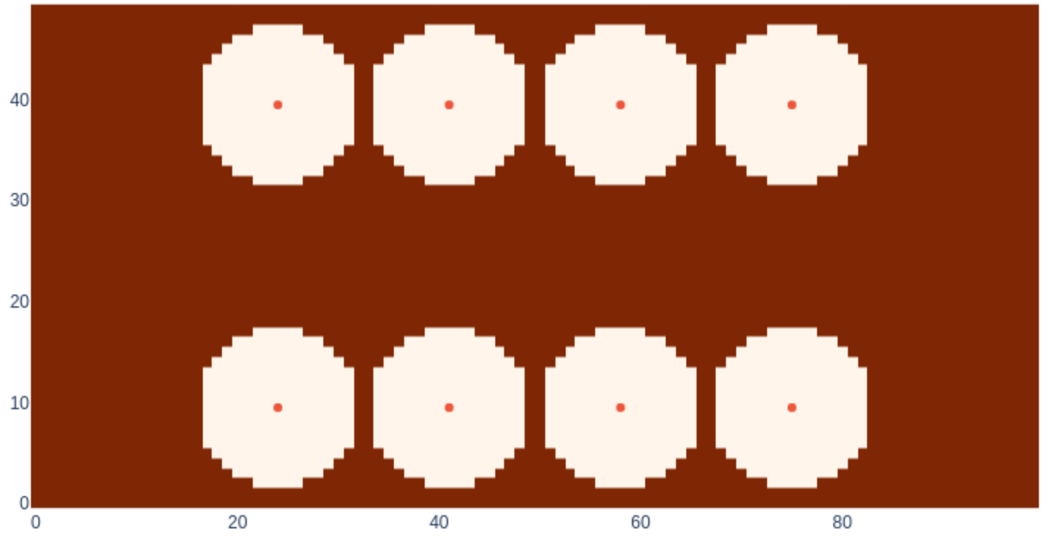

If the cutoff at d_{max} is sharp, the likelihood of being within sensing range of a beacon is zero when the sensor returns $z_k = \text{inf}$.

A Pseudo-GPS Sensor

- GPS-like sensors return the coordinates of the sensor relative to some fixed, global reference frame.
- In the simplest case, we have $z_k = h(x_k) = x_k$.
- It is not unusual to define measurements in units that are different from those used by the robot, e.g., the robot might measure its coordinates in meters while the GPS returns coordinates in centimeters.
- In these cases, we simply scale the measurement appropriately: $z_k = h(x_k) = Cx_k$
- If we now consider additive noise, we obtain our measurement model for noisy GPS-like sensors:

$$z_k = h(x_k) + w_k = Cx_k + w_k$$

- If w_k is i.i.d. zero-mean Gaussian noise (as usual), the measurements are governed by a conditional Gaussian probability density:

$$f_{Z_k}(z_k|x_k) = \frac{1}{\sqrt{|2\pi\Sigma|}} exp\left\{-\frac{1}{2}(z_k - Cx_k)^T \Sigma^{-1}(z_k - Cx_k)\right\}$$

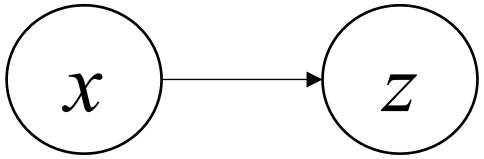

GPS-style Likelihoods

- The likelihood for our GPS-like sensor is given by

$$\mathcal{L}(x_k; z_k) = \frac{1}{\sqrt{|2\pi\Sigma|}} exp\left\{-\frac{1}{2}(z_k - Cx_k)^T\Sigma^{-1}(z_k - Cx_k)\right\}$$

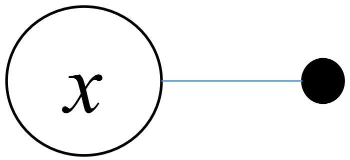

- Let's work on the exponent: $(z_k - Cx_k)^T$

$$(z_k - Cx_k) = C(C^{-1}z_k - x_k) \rightarrow (z_k - Cx_k)^T = [C(C^{-1}z_k - x_k)]^T = (C^{-1}z_k - x_k)^T C^T$$

- Therefore, we can write the likelihood as:

$$\mathcal{L}(x_k; z_k) = \frac{1}{\sqrt{|2\pi\Sigma|}} exp\left\{-\frac{1}{2}(x_k - C^{-1}z_k)^T C^T\Sigma^{-1}C(x_k - C^{-1}z_k)\right\}$$

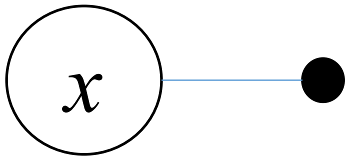

which has the form of a Gaussian with mean $C^{-1}z_k$ and inverse covariance $C^T\Sigma^{-1}C$.

Simulating States and Measurements

- Given a control tape $u_1, \ldots, u_{n-1}$ and a prior distribution for X_1, it's easy to generate a sample trajectory $x_1, \ldots, x_n$ along with a sample measurement history $z_1, \ldots, z_n$.

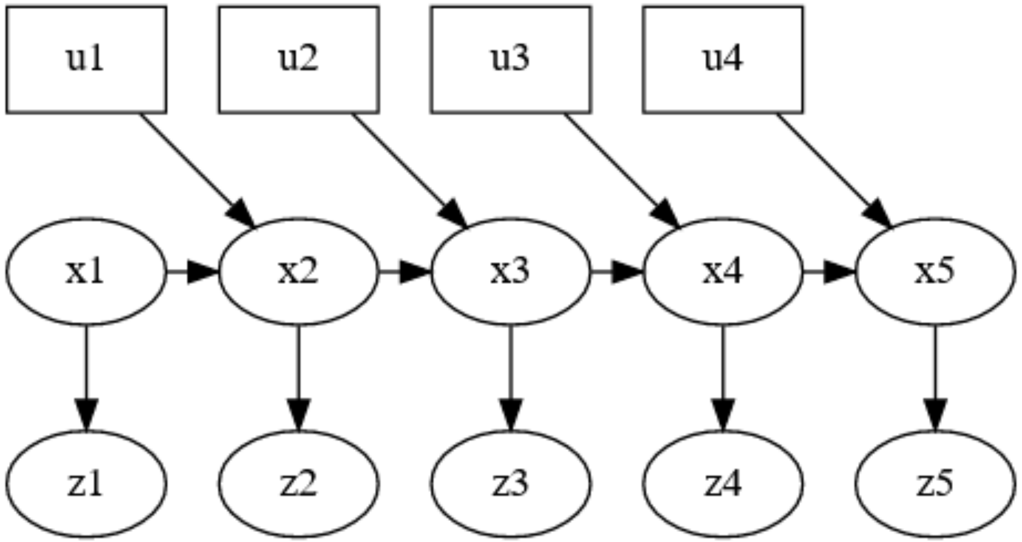

1. Generate a sample for x_1 by sampling from the prior $P(X_1 = x_1)$.
2. Generate a sample measurement z_1 by sampling from the measurement model $p(Z_1|x_1)$
3. For each i:
   1. Generate a sample for x_i by sampling from the transition distribution $p(X_i|x_{i-1}, u_{i-1})$
   2. Generate a measurement sample z_i by sampling from the measurement model $p(Z_i|x_i)$

Next Time…

Perception

- Bayes Filter
- Markov Localization
- Monte Carlo Localization