Autonomous Robotics
Winter 2024

Abhishek Gupta
TAs: Karthikeya Vemuri, Arnav Thareja
Marius Memmel, Yunchu Zhang

Class Outline

State Estimation
- Robotic System Design
- Filtering
- Localization
- SLAM

Control
- Feedback Control
- PID Control
- MPC
- LQR

Planning
- Search
- Heuristic Search
- Motion Planning
- Lazy Search

Learning
- Imitation Learning
- Policy Gradient
- Actor-Critic
- Model-Based RL

Logistics

- Teams and workstations assigned
- Project 1 released Saturday morning!
- Pick up cars next week:1/14: 4:30-5:30pm and1/15: 4:00-5:00pm

- Post questions, discuss any issues you are having on Ed.
- Students with **no** access to 002, e-mail us with your student ID.
- Students that have not been added to the class, email abhgupta@cs.washington.edu with the subject-line "Waitlisted for CSE478"

Recap

Our knowledge of how the world works **(Modeling)**

Helicopter Models

World Model

Planning

State estimation

How do we decide how to act to accomplish task? **(Planning/Control)**

How do we decide where we are in the world? **(Estimation)**

Control

Sensors

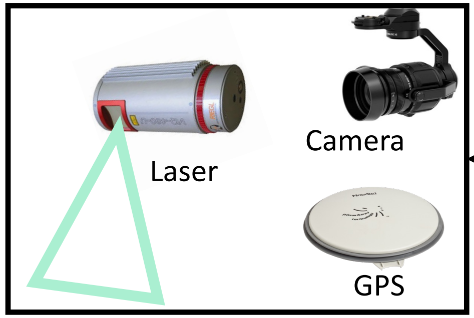

Robot interacts with environment

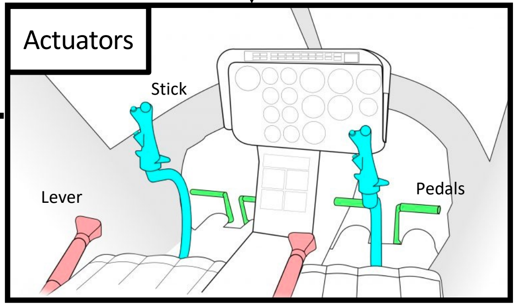

End-to-End Learning Based Control for Robots

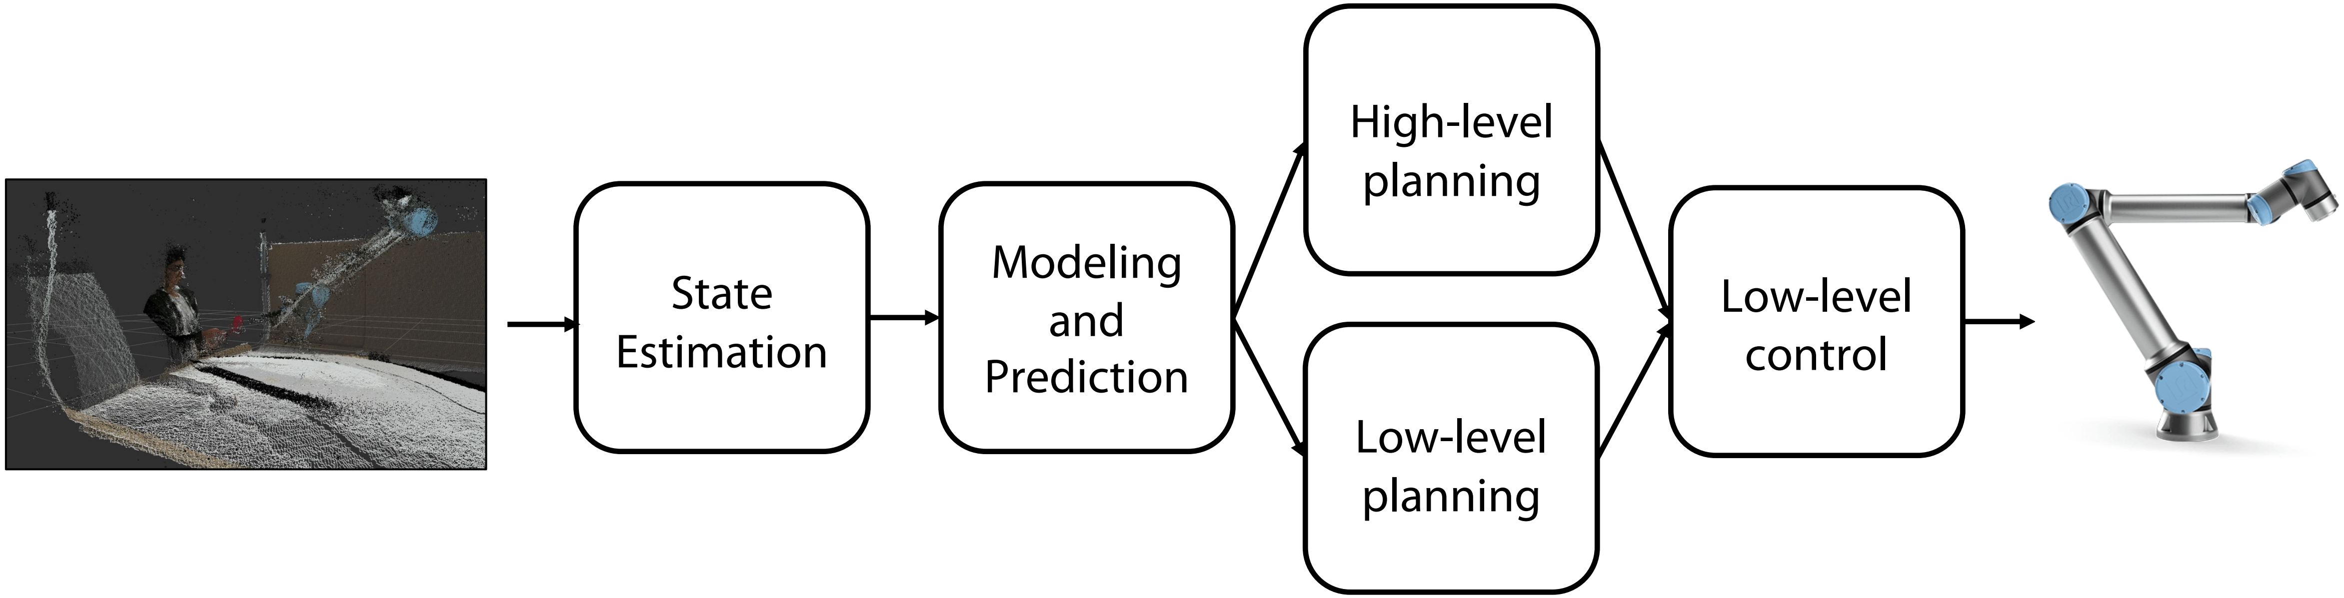

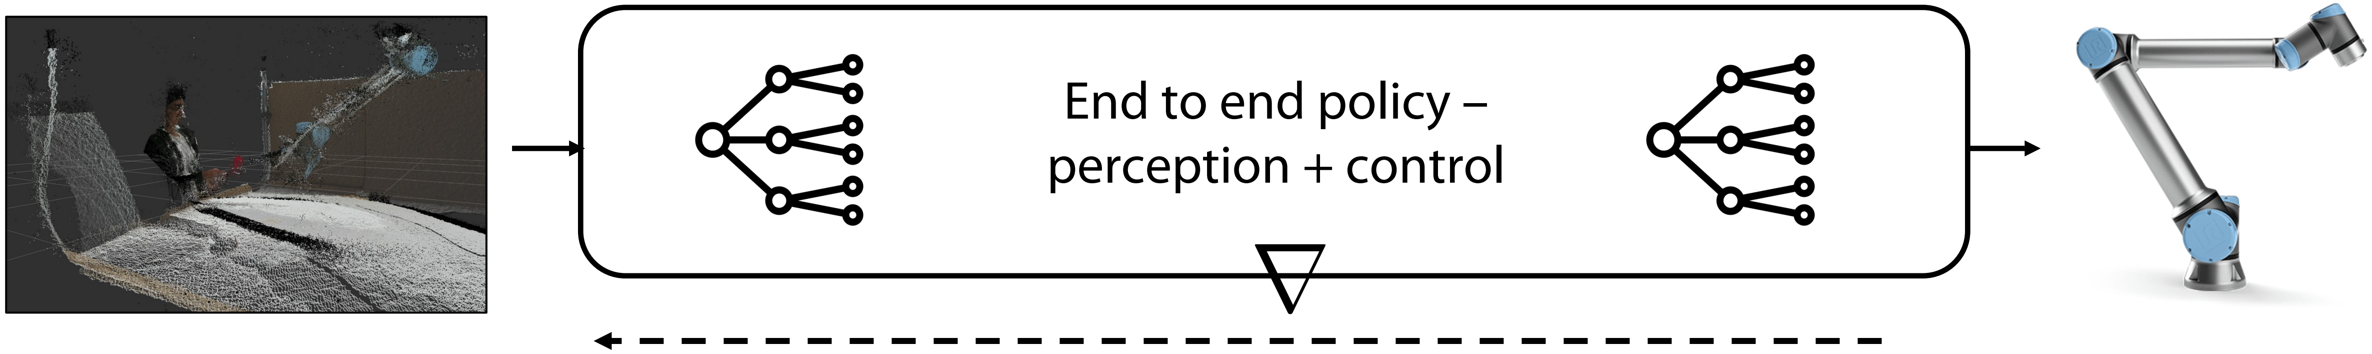

End-to-End Learning Based Control for Robots

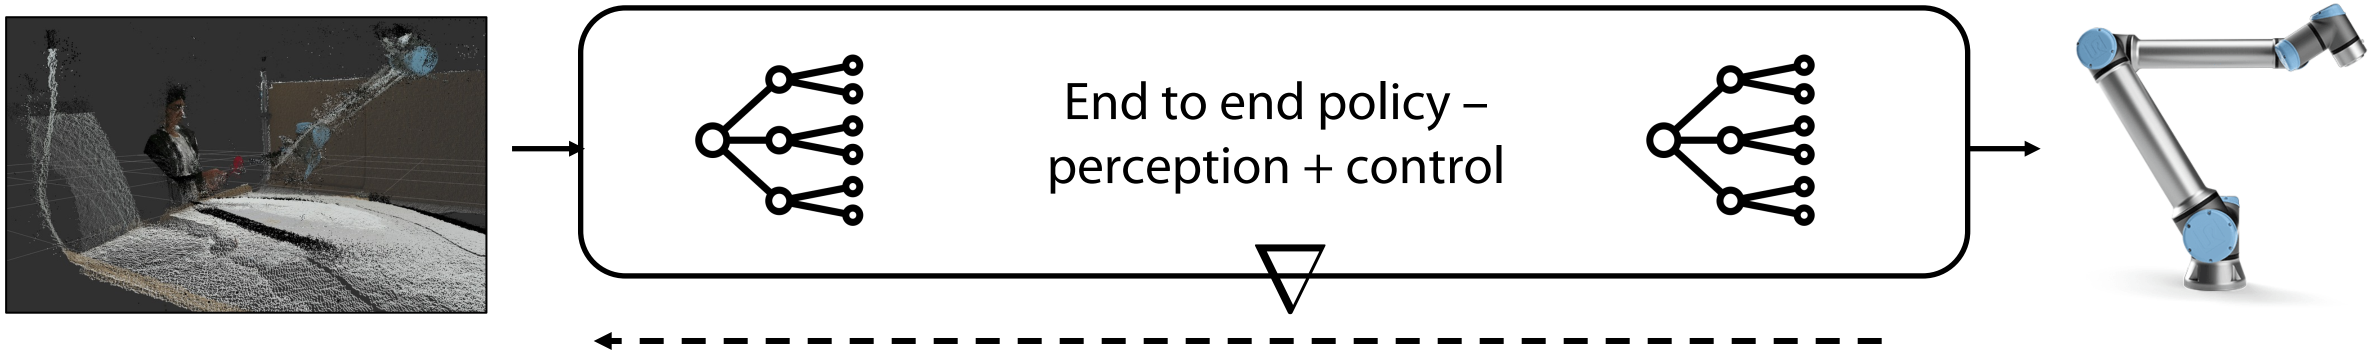

Option 1: Imitation Learning

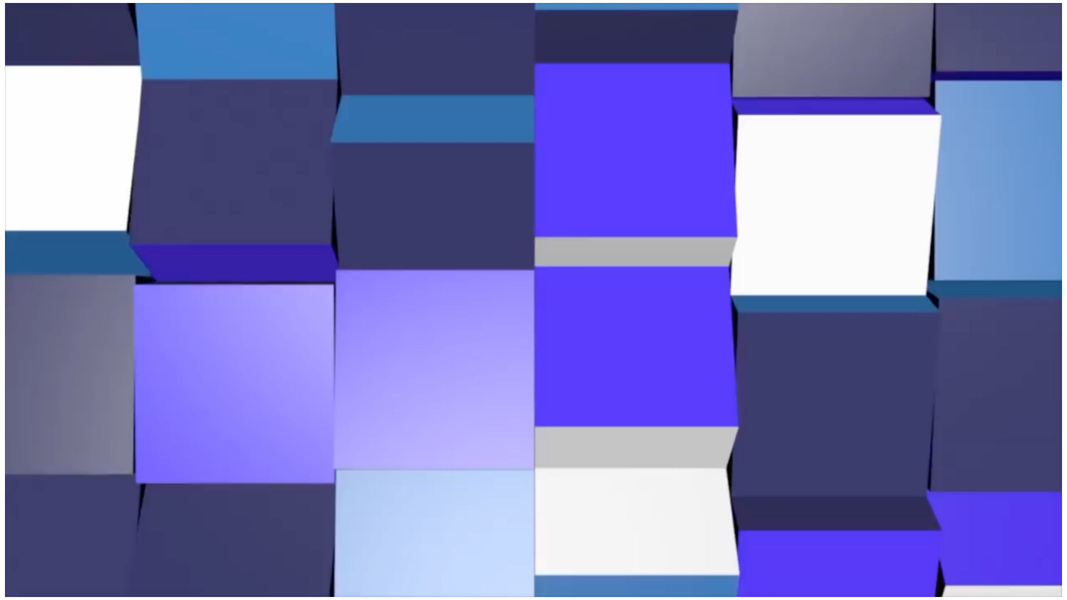

Learning by copying an expert

Option 2: Reinforcement Learning

Learning through trial and error

Why might we want/not want to do this?

Modules compensate for each other

Avoids hand-designing and supervising interfaces

Often more performant/less biased

Lack of Interpretability

Lack of Reusability

Often data inefficient

What are we going to talk about today?

A probabilistic approach to **state estimation**

Lecture Outline

What is state estimation?

↓

Probability Review and Bayes Rule

↓

Bayesian Filtering w/ Examples

Today's Objective: Understand how to formalize state estimation

What does it mean to **estimate** "state"?

State estimation

Sensors

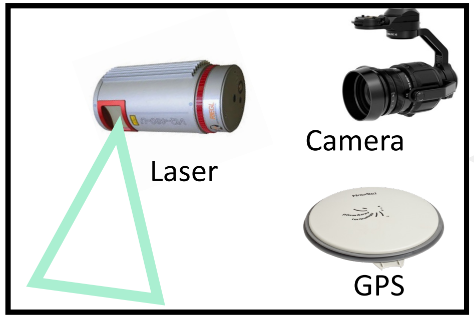

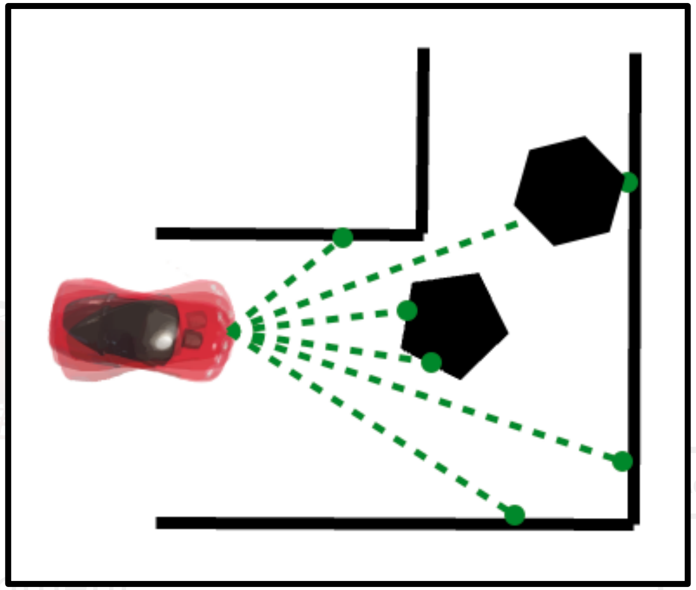

Why state estimation: A historical case study

Thrun et al. 1998

Why state estimation: A historical case study

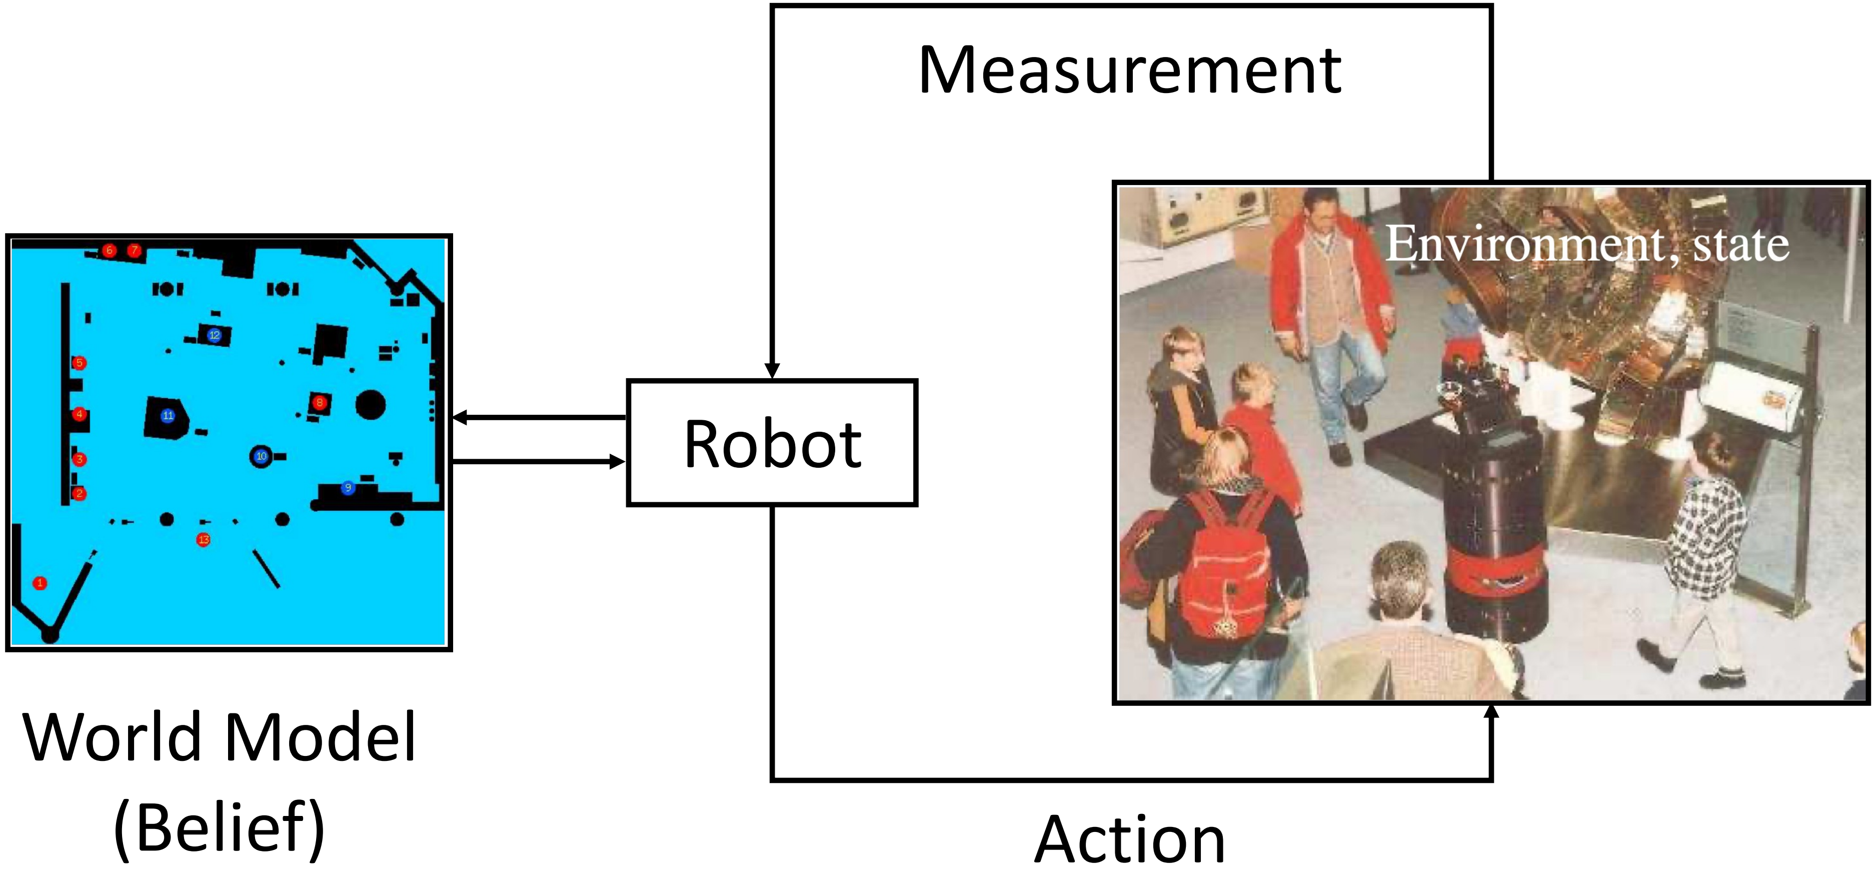

What is the problem of “state” estimation?

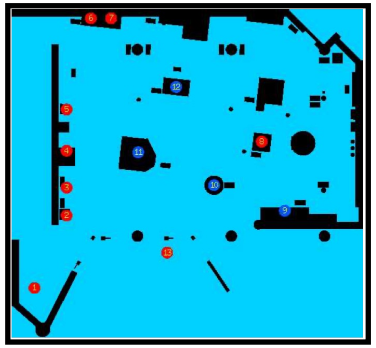

Given data (stream of measurements and actions) → ESTIMATE → Belief (probability of states)

Let us formally define this problem, and then ask why it is hard?

State: A very abstract definition

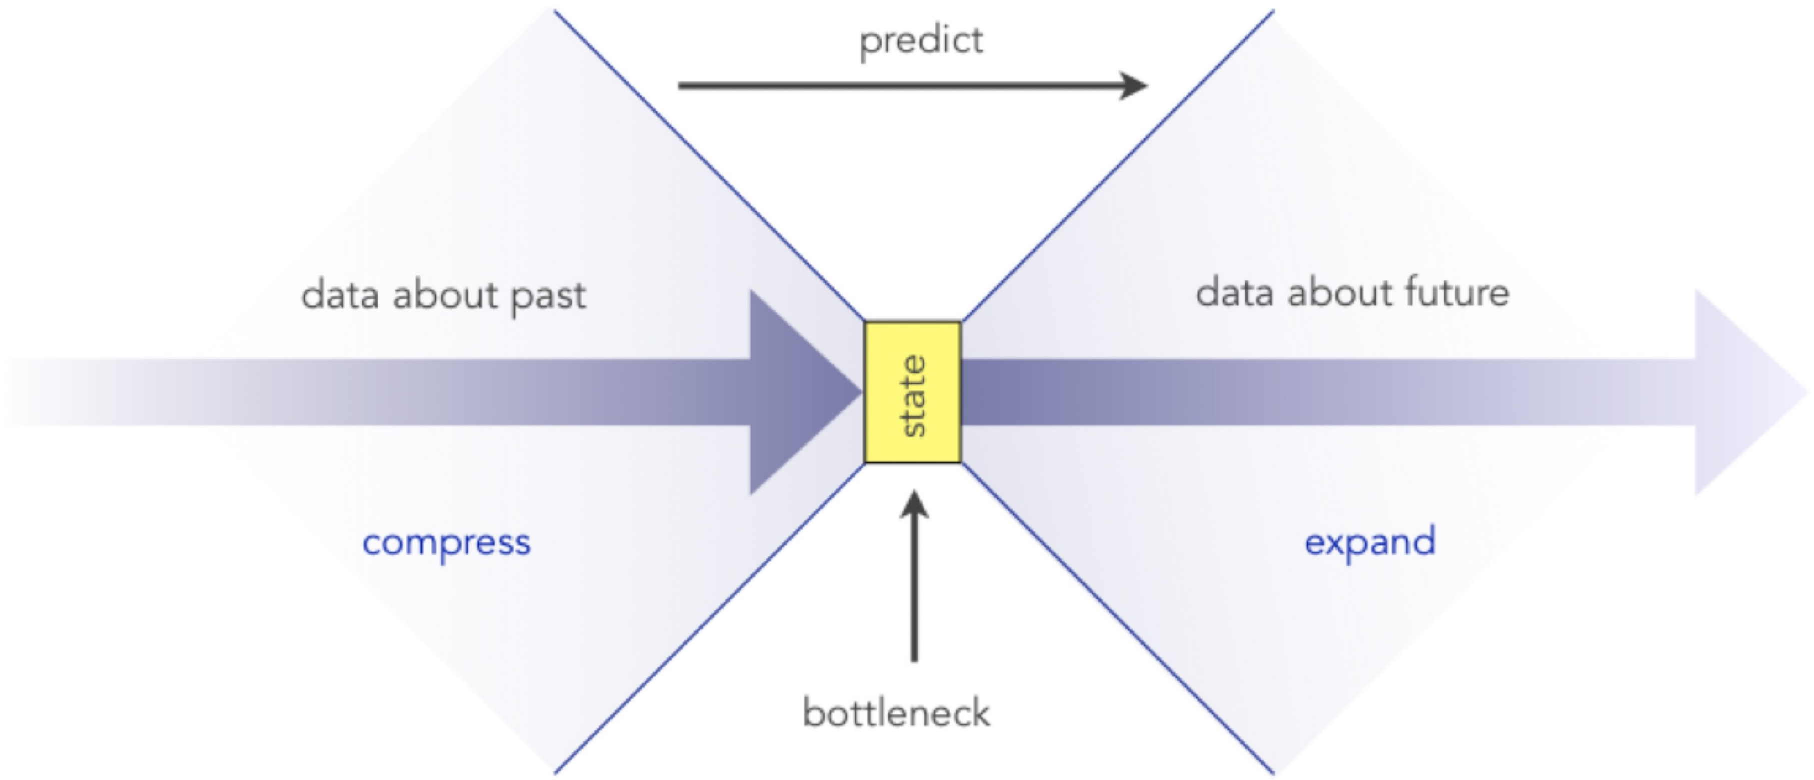

State: statistic of history sufficient to predict the future

State (x_t)

Collection of variables sufficient to predict the ~~future~~

(future that we care about)

What are some examples of state?

1. Pose of a robot - Usually 6 dof (3 position, 3 for orientation)
   - 3 dof for planar mobile robot (x, y, heading)

2. Configuration of a manipulator - Collection of joint angles

3. Location of objects in environment

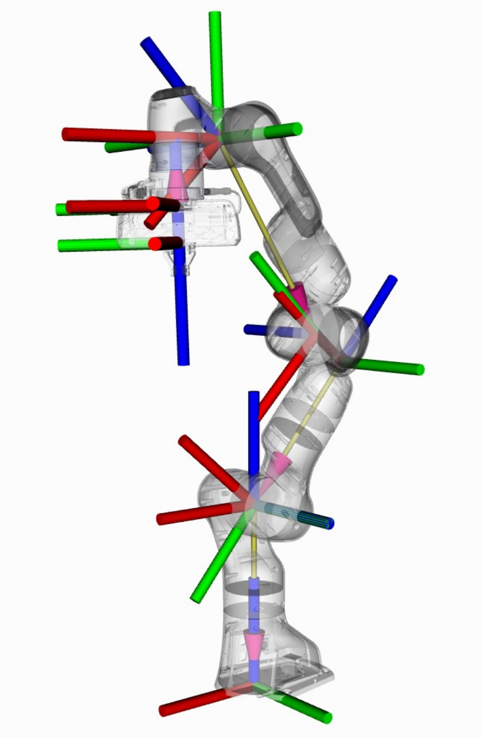

State can be static/dynamic, discrete/continuous/hybrid

Measurement (z_t)

Measurements are sensor values that provide information about state.

(Measurement does not always tell you state directly!! – why?)

What are some examples of measurements?

1. GPS - absolute information about robot pose

2. Laser scan - relative geometric information between pose and environment

3. Camera image - information about color / texture (harder to model)

Action (u_t)

Actions are what a robot uses to control how a state changes from one time to another

What are some examples of actions?

1. Active forces applied by the robot - (measure motor currents, force torque sensors, odometers)

2. NOP actions - doing nothing is also an action.

Fundamental Problem: State is hidden

All the robot sees is a stream of actions and measurements

$$u_1, z_1, u_2, z_2, u_3, z_3, \ldots$$

But robot never sees the state

$$x_1, x_2, x_3, \ldots$$

Fundamental Problem: State is hidden

But all decision making depends on knowing state

Solution: Estimate belief over state

$$bel(x_t) = P(x_t | z_{1:t}, u_{1:t})$$

Belief is a probability of each possible state given history

Also called Posterior / Information state / State of knowledge

Represent belief? Parametric (Gaussian), Non-parametric (Histogram)

Why be probabalistic?

x_t Pose/velocity of the object

- When state is abstracted/incomplete, this manifests as noise/uncertainty
- Being probabilistic allows for:
  - Robustness to external noise
  - Exploration to get better/gather information
  - Dealing with inherently stochastic systems
  - Accounting for inaccurate hardware/software

What happens if we are not estimating state?

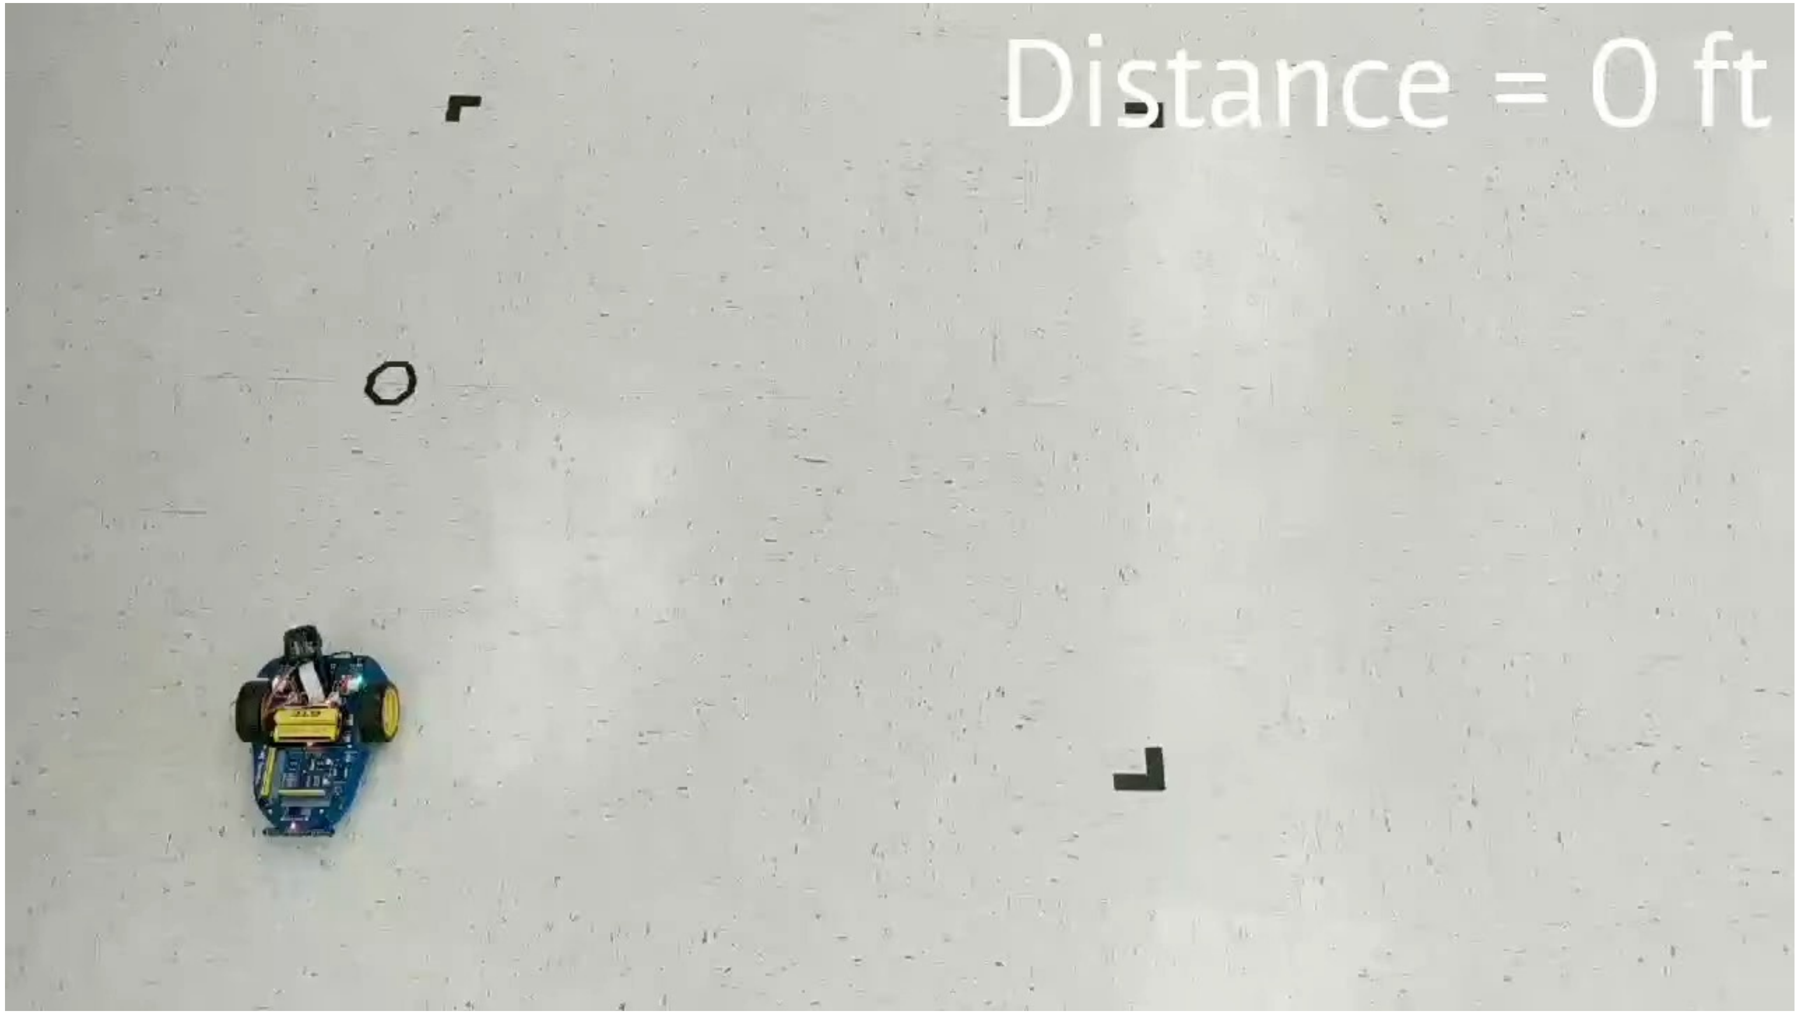

Errors clearly accumulate over time due to noise/unmodeled effects

Probabilistic Robotics

Key idea: Explicit representation of uncertainty

(using the calculus of probability theory)

- Perception = state estimation
- Action = utility optimization

Lecture Outline

What is state estimation?

↓

Probability Review and Bayes Rule

↓

Bayesian Filtering w/ Examples

Let’s brush up on probability!

Fundamental Axioms of Probability

$$0 \leq \Pr(A) \leq 1$$

$$\Pr(\Omega) = 1 \qquad \Pr(\phi) = 0$$

$$\Pr(A \cup B) = \Pr(A) + \Pr(B) - \Pr(A \cap B)$$

- Pr(A) denotes probability that the outcome
- ω is an element of the set of possible outcomes A.
- A is often called an event. Same for B.
- Ω is the set of all possible outcomes.
- φ is the empty set.

Useful Corollaries from Axioms

$$\begin{aligned}
\Pr(A \cup (\Omega \setminus A)) &= \Pr(A) + \Pr(\Omega \setminus A) - \Pr(A \cap (\Omega \setminus A)) \\
\Pr(\Omega) &= \Pr(A) + \Pr(\Omega \setminus A) - \Pr(\phi) \\
1 &= \Pr(A) + \Pr(\Omega \setminus A) - 0 \\
\Pr(\Omega \setminus A) &= 1 - \Pr(A)
\end{aligned}$$

If A and B have no overlap then

$$\Pr(A \cup B) = \Pr(A) + \Pr(B)$$

Discrete Random Variables

- ***X*** denotes a random variable.
- ***X*** can take on a countable number of values in $\{x_1, x_2, \ldots, x_n\}$.
- ***P(X=x$_i$)***, or ***P(x$_i$)***, is the probability that the random variable ***X*** takes on value ***x***$_i$.
- ***P(·)*** is called probability mass function.
- E.g. $$P(Room) = \langle 0.7, 0.2, 0.08, 0.02 \rangle$$

Examples of Discrete Random Variables

Binomial

$$\binom{n}{k} p^k q^{n-k}$$

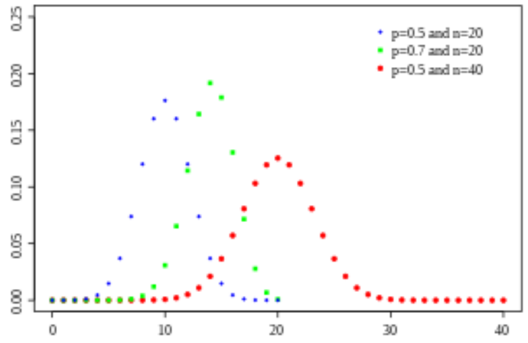

Multinomial

$$\frac{n!}{x_1! \cdots x_k!} p_1^{x_1} \cdots p_k^{x_k}$$

Bernoulli

$$\begin{cases} q = 1 - p & \text{if } k = 0 \\ p & \text{if } k = 1 \end{cases}$$

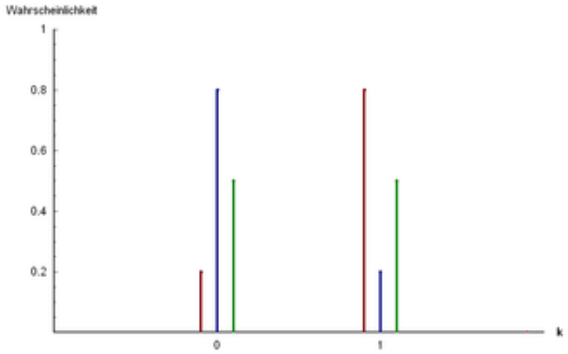

Poisson

$$\frac{\lambda^k e^{-\lambda}}{k!}$$

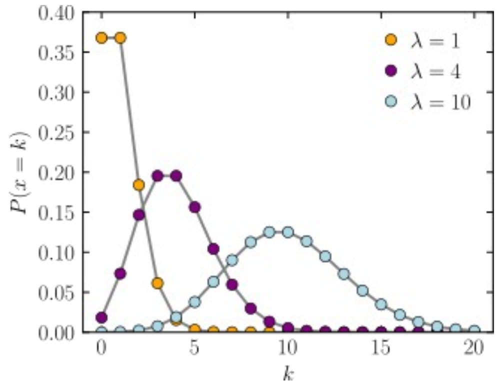

Continuous Random Variables

- ***X*** denotes a random variable.
- ***X*** can take on a continuum of values in the support of the probability density function
- ***P(X=x)***, or ***P(x)***, is the probability density function
  - Density function positive but not upper bounded by 1

$$\Pr(x \in (a, b)) = \int_a^b p(x)dx$$

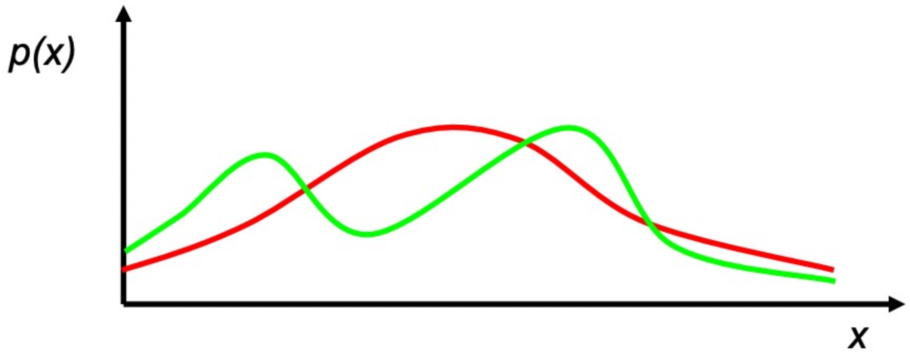

Examples of Continuous Random Variables

Multivariate Gaussian

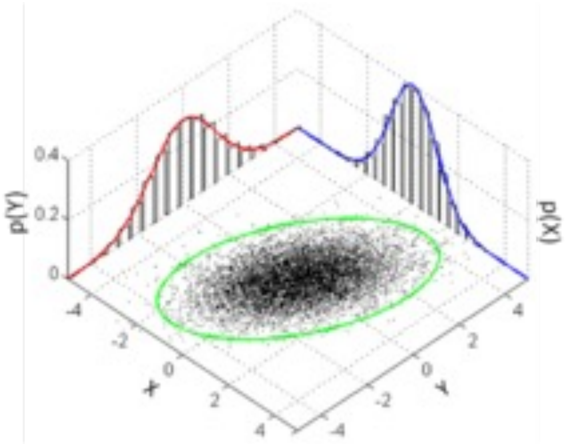

Beta Distribution

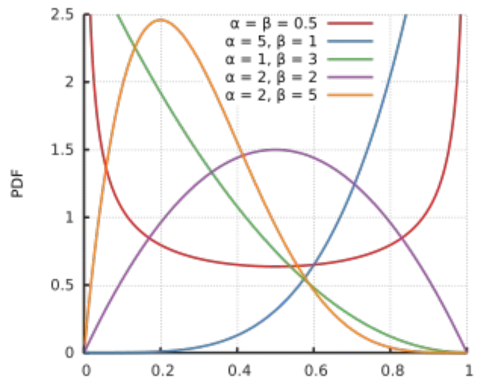

Beta Distribution

$$\begin{cases} \frac{1}{b-a} & \text{for } x \in [a,b] \\ 0 & \text{otherwise} \end{cases}$$

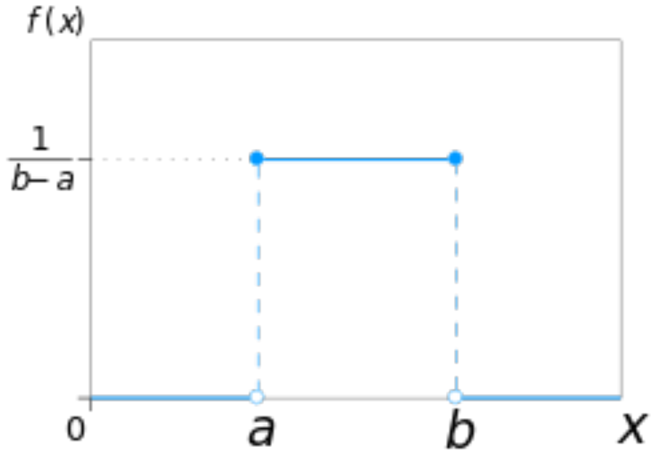

Joint and Conditional Probability

- $P(X{=}x \text{ and } Y{=}y) = P(x,y)$

- If X and Y are independent then

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

- $P(x \mid y)$ is the probability of ***x*** given ***y***

$$P(x \mid y) = P(x,y) \,/\, P(y)$$

$$P(x,y) = P(x \mid y)\,P(y)$$

- If X and Y are independent then

$$P(x \mid y) = P(x)$$

Law of Total Probability, Marginals

Discrete case

$$\sum_x P(x) = 1$$

$$P(x) = \sum_y P(x, y)$$

$$P(x) = \sum_y P(x \mid y) P(y)$$

Continuous case

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

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

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

Events

- P(+x, +y) ?
- P(+x) ?
- P(-y OR +x) ?

$P(X, Y)$

X	Y	P
+x	+y	0.2
+x	-y	0.3
-x	+y	0.4
-x	-y	0.1

Marginal Distributions

$P(X, Y)$

X	Y	P
+x	+y	0.2
+x	-y	0.3
-x	+y	0.4
-x	-y	0.1

$$P(x) = \sum_y P(x, y)$$

$$P(y) = \sum_x P(x, y)$$

$P(X)$

X	P
+x	
-x	

$P(Y)$

Y	P
+y	
-y	

Conditional Probabilities

$P(X,Y)$

X	Y	P
+x	+y	0.2
+x	-y	0.3
-x	+y	0.4
-x	-y	0.1

- P(+x | +y) ?
- P(-x | +y) ?
- P(-y | +x) ?

Bayes Formula

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

$$\Rightarrow$$

$$P(x|y) = \frac{P(y|x)\ P(x)}{P(y)} = \frac{\text{likelihood} \cdot \text{prior}}{\text{evidence}}$$

Bayes Formula

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

$$P(y) = \sum_{x'} P(y \mid x')P(x')$$

$$\Downarrow$$

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

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

Can replace with integral

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

Example of Bayes Formula in Action

Cancer \ Symptom	Yes	No	Total
Yes	1	0	1
No	10	99989	99999
Total	11	99989	100000

Just because everyone with cancer has the symptom, doesn't mean everyone with the symptom has cancer

$$\begin{aligned} P(\text{Cancer}|\text{Symptoms}) &= \frac{P(\text{Symptoms}|\text{Cancer})P(\text{Cancer})}{P(\text{Symptoms})} \\ &= \frac{P(\text{Symptoms}|\text{Cancer})P(\text{Cancer})}{P(\text{Symptoms}|\text{Cancer})P(\text{Cancer}) + P(\text{Symptoms}|\text{Non-Cancer})P(\text{Non-Cancer})} \\ &= \frac{1 \times 0.00001}{1 \times 0.00001 + (10/99999) \times 0.99999} = \frac{1}{11} \approx 9.1\% \end{aligned}$$

Why Bayes Formula?

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

$$P(y) = \sum_{x'} P(y \mid x')P(x')$$

Diagnostic

Causal

- Causal knowledge may be easier to obtain/estimate
- Which direction is causal is not always clear though!
- Allows us to estimate “beliefs” based on “measurements”

Lecture Outline

What is state estimation?

↓

Probability Review and Bayes Rule

↓

Bayesian Filtering w/ Examples

Let’s represent the state estimation problem graphically

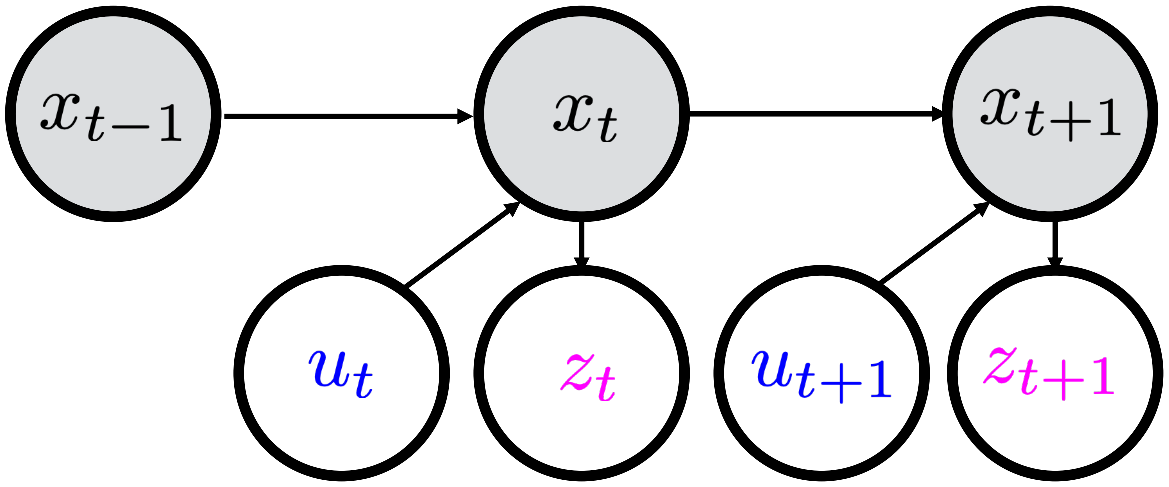

Assumptions:

1. Robot receives a stream of measurements / actions.

2. One measurement / action per time-step.

What is belief in this setting?

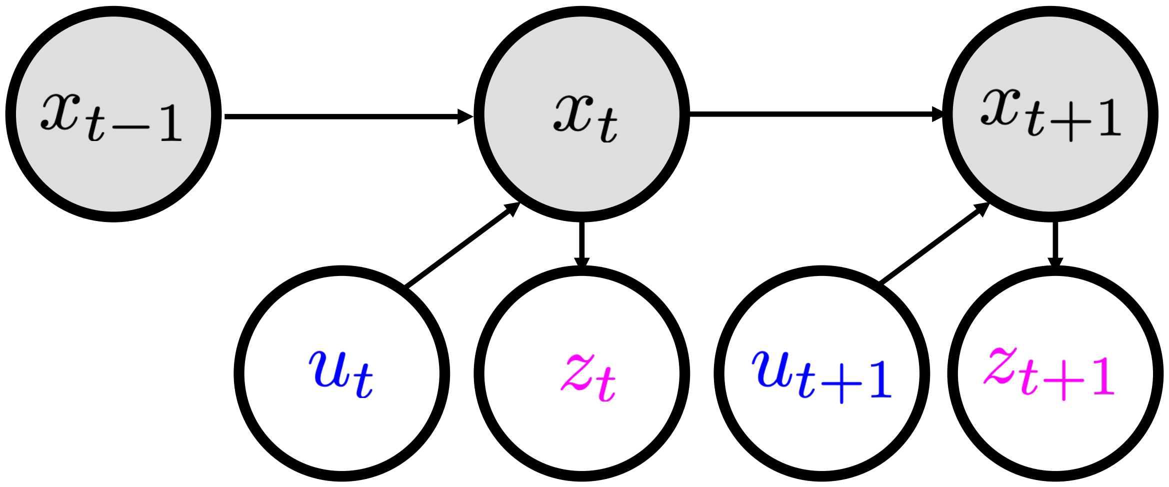

P(current state | all past information)

$$P(x_t | z_t, u_t, x_{t-1}, \ldots)$$

Can we estimate this?

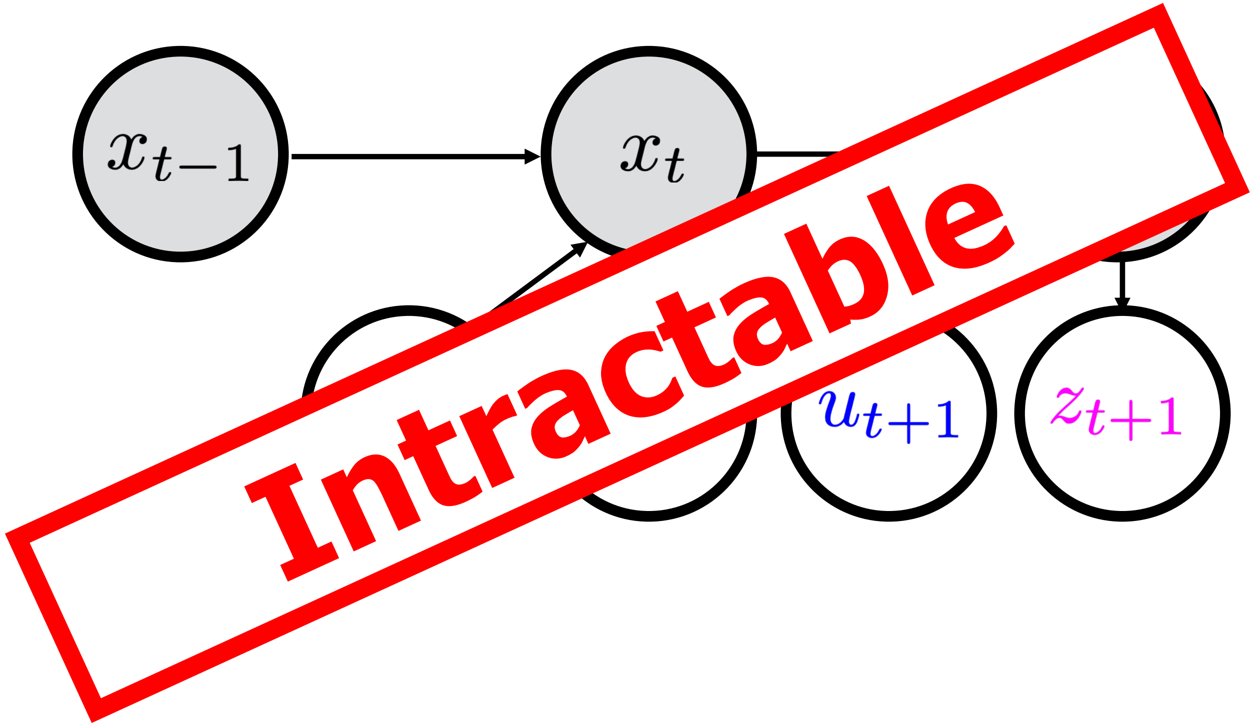

P(current state | all past information)

$$P(x_t | z_t, u_t, x_{t-1}, \ldots)$$

Good ol' Markov to the rescue

Andrey Andreyevich Markov (1856 - 1922)

Solution: Markov Assumption

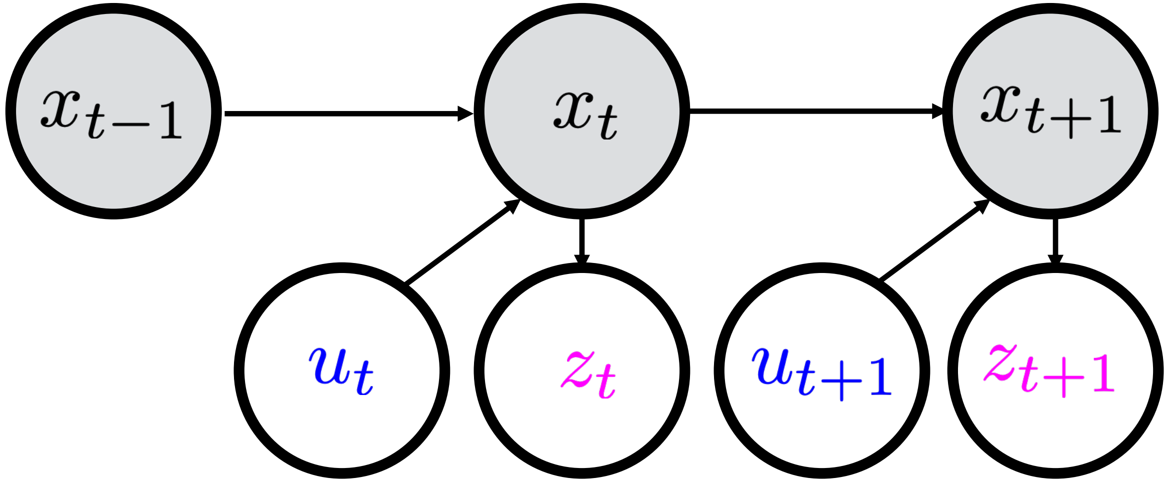

Markov assumption :

Future state ***conditionally independent*** *of past actions, measurements* ***given*** *present state.*

$$P(x_t|u_t, x_{t-1}, z_{t-1}, u_{t-1}, \dots) = P(x_t|u_t, x_{t-1})$$

$$P(z_t|x_t, u_t, x_{t-1}, z_{t-1}, u_{t-1}, \dots) = P(z_t|x_t)$$

Probabilistic models

State transition probability / dynamics / motion model

$$P(x_t | x_{t-1}, u_t)$$

Measurement probability / Observation model

$$P(z_t | x_t)$$

When does Markov not hold?

$$P(x_t | x_{t-1}, u_t) \qquad P(z_t | x_t)$$

whenever state doesn't capture all requisite information

- Camera images at different times of the day

- Unmodelled pedestrians in front of laser

- Steady gusts of wind

How do we tractably calculate belief?

Input data

Measurement $z_{1:t}$

Actions $u_{1:t}$

Belief

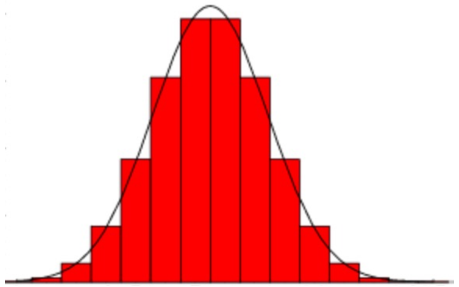

$$bel(x_t) = P(x_t | z_{1:t}, u_{1:t})$$

Ans: Bayes filter!

Bayes filter in a nutshell

Key Idea: Apply Markov to get a recursive update!

Bayes filter in a nutshell

Step 0. Start with the belief at time step t-1

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

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

Bayes filter in a nutshell

Step 1: Prediction - push belief through dynamics given action

(discrete) $$\overline{bel}(x_t) = \sum P(x_t|u_t, x_{t-1})bel(x_{t-1})$$

(total probability)

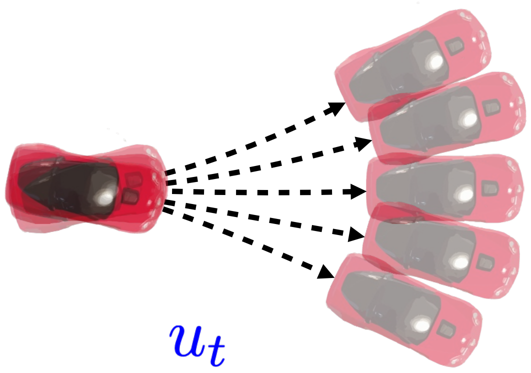

$bel(x_{t-1})$ $\overline{bel}(x_t)$

Bayes filter in a nutshell

Step 2: Correction - apply Bayes rule given measurement

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

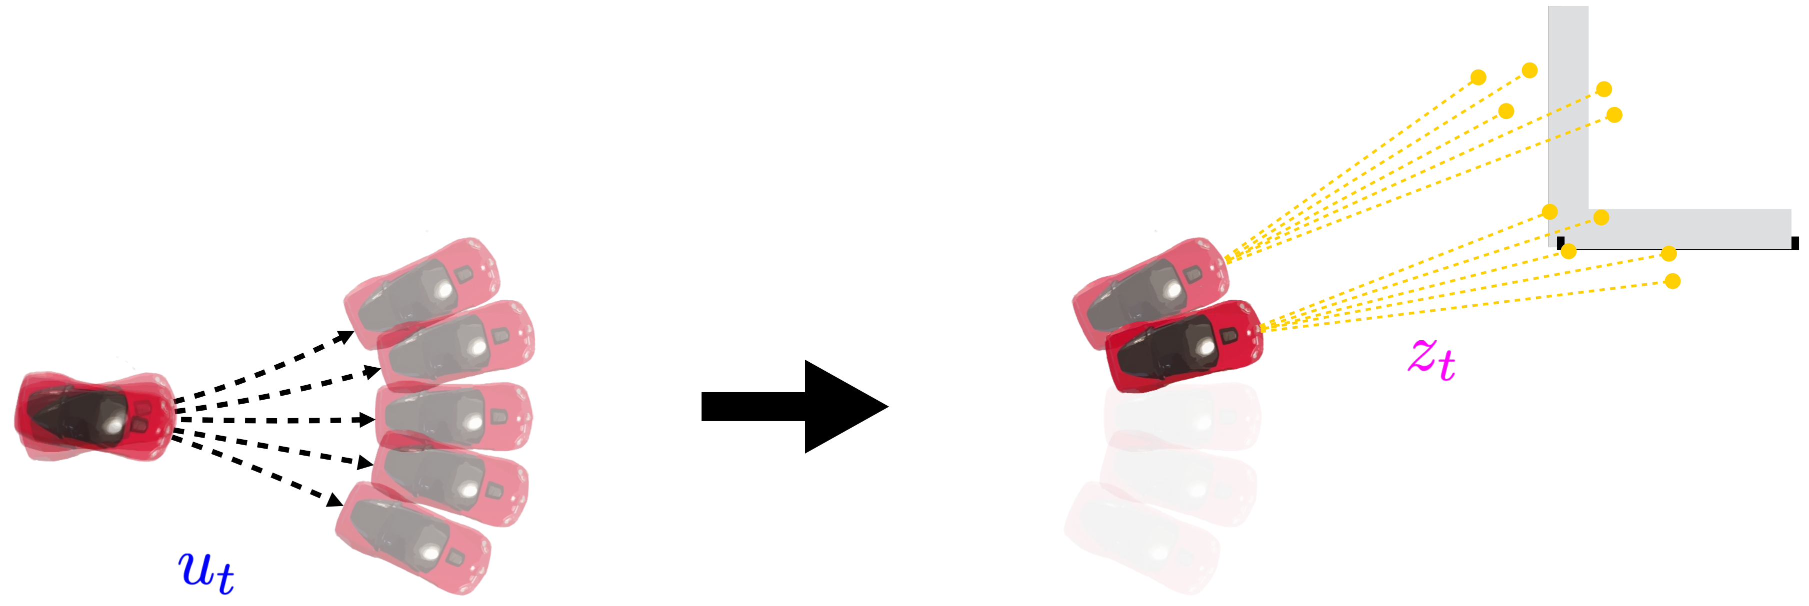

$bel(x_{t-1})$ $\overline{bel}(x_t)$ $bel(x_t)$

$$\eta = \frac{1}{\sum P(z_t|x_t)\overline{bel}(x_t)}$$

Bayes filter in a nutshell

Key Idea: Apply Markov to get a recursive update!

Step 0. Start with the belief at time step t-1

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

Step 1: Prediction - push belief through dynamics given action

$$\overline{bel}(x_t) = \sum P(x_t|u_t, x_{t-1})bel(x_{t-1})$$

Step 2: Correction - apply Bayes rule given measurement

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

Bayes filter is a powerful tool

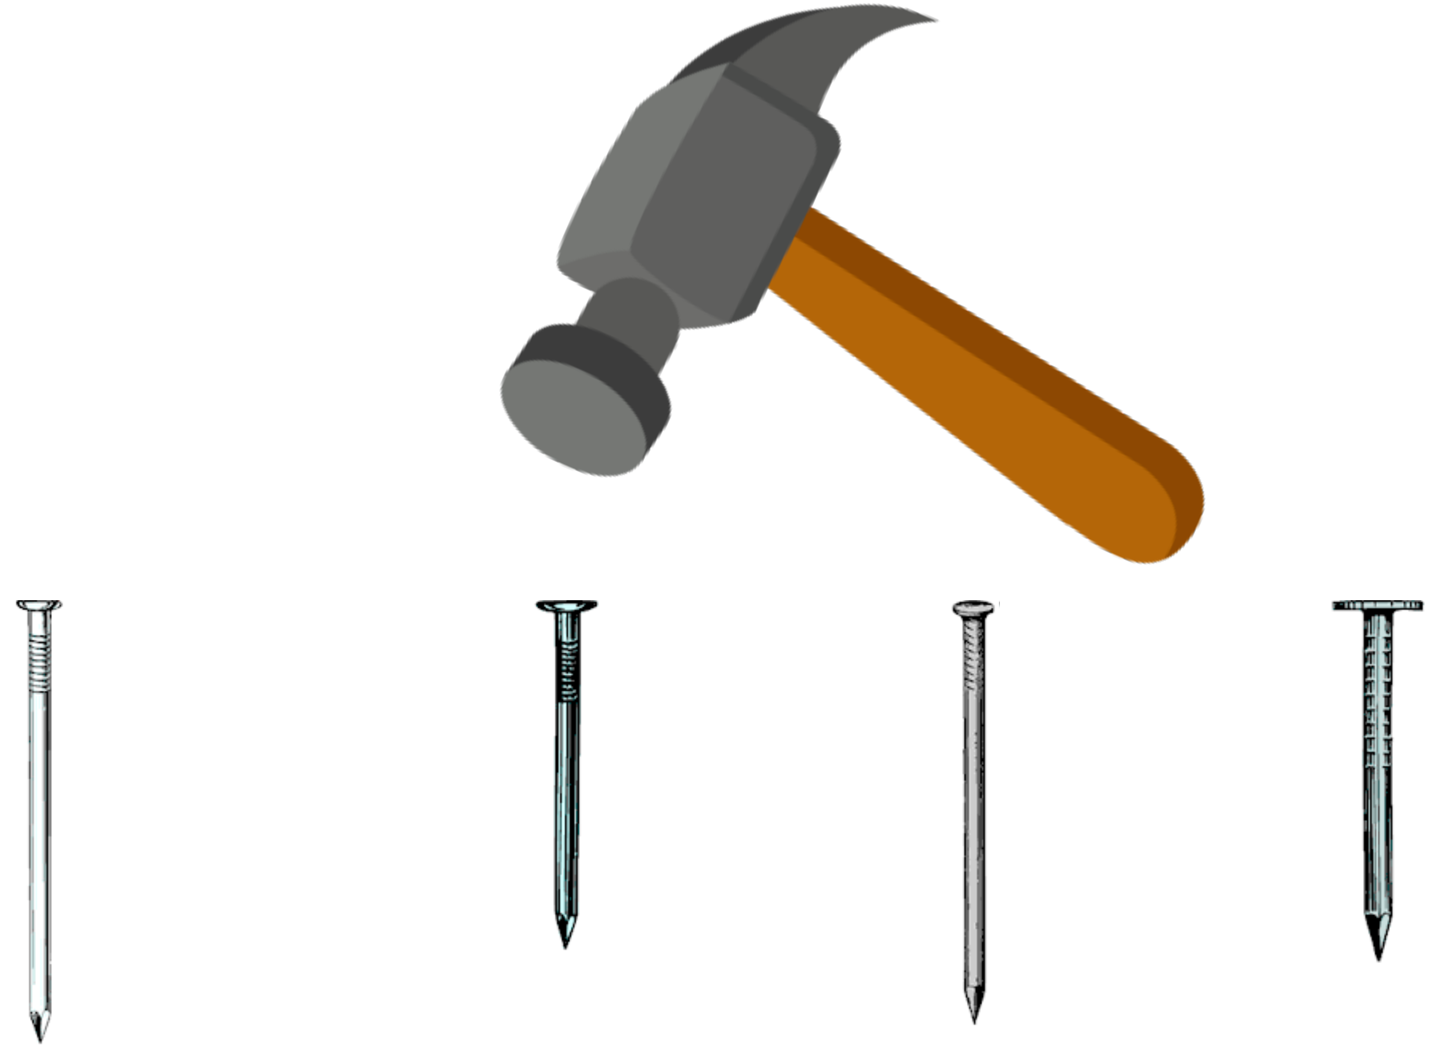

Localization Mapping SLAM POMDP

Example: Opening a Door

$\mathcal{X} =$ **O**PEN, **C**LOSED

$\mathcal{A} =$ **P**ULL, **L**EAVE $\quad P(x_t | x_{t-1}, u_t)$

$$P(O|C,P) = 0.7$$
$$P(C|C,P) = 0.3$$

HTTPS://YOUTU.BE/AFUA50H9UEK

Example: Opening a Door

$$\begin{bmatrix} P(x_t = \mathbf{O} | x_{t-1} = \mathbf{O}, u_t) & P(x_t = \mathbf{O} | x_{t-1} = \mathbf{C}, u_t) \\ P(x_t = \mathbf{C} | x_{t-1} = \mathbf{O}, u_t) & P(x_t = \mathbf{C} | x_{t-1} = \mathbf{C}, u_t) \end{bmatrix}$$

$$P(.|., \mathbf{P}) = \begin{bmatrix} 0.8 & 0.7 \\ 0.2 & 0.3 \end{bmatrix} \qquad P(.|., \mathbf{L}) = \begin{bmatrix} 0.5 & 0 \\ 0.5 & 1 \end{bmatrix}$$

HTTPS://YOUTU.BE/AFUA50H9UEK

Example: Opening a Door

$\mathcal{X} =$ **O**PEN, **C**LOSED

$\mathcal{A} =$ **P**ULL, **L**EAVE

$\mathcal{Z} =$ **O**PEN, **C**LOSED

$$P(z_t|x_t)$$

$$\begin{bmatrix} P(z_t|\mathbf{O}) \\ P(z_t|\mathbf{C}) \end{bmatrix}$$

$$P(\mathbf{O}|.) = \begin{bmatrix} 0.6 \\ 0.2 \end{bmatrix} \quad P(\mathbf{C}|.) = \begin{bmatrix} 0.4 \\ 0.8 \end{bmatrix}$$

HTTPS://YOUTU.BE/AFUA50H9UEK

Example: Opening a Door

$\mathcal{X} =$ **O**PEN, **C**LOSED

$\mathcal{A} =$ **P**ULL, **L**EAVE

$\mathcal{Z} =$ **O**PEN, **C**LOSED

$$Bel(x_0) = \begin{bmatrix} 0.4 \\ 0.6 \end{bmatrix}$$

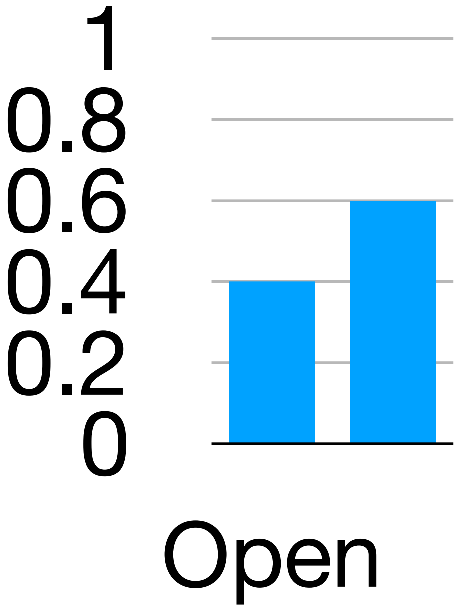

PULL

Example: Opening a Door

$\mathcal{X} =$ **O**PEN, **C**LOSED

$\mathcal{A} =$ **P**ULL, **L**EAVE

$\mathcal{Z} =$ **O**PEN, **C**LOSED

Prediction: Given action, propagate belief through dynamics

$$\overline{Bel}(x_t) = \sum_{x_{t-1}} P(x_t | u_t, x_{t-1}) Bel(x_{t-1})$$

$$\underbrace{\begin{bmatrix} P(x_t = \mathbf{O}) \\ P(x_t = \mathbf{C}) \end{bmatrix}}_{\overline{Bel}(x_t)} = \begin{bmatrix} P(x_t = \mathbf{O} | x_{t-1} = \mathbf{O}, u_t) & P(x_t = \mathbf{O} | x_{t-1} = \mathbf{C}, u_t) \\ P(x_t = \mathbf{C} | x_{t-1} = \mathbf{O}, u_t) & P(x_t = \mathbf{C} | x_{t-1} = \mathbf{C}, u_t) \end{bmatrix} \underbrace{\begin{bmatrix} P(x_{t-1} = \mathbf{O}) \\ P(x_{t-1} = \mathbf{C}) \end{bmatrix}}_{Bel(x_{t-1})}$$

Example: Opening a Door

$\mathcal{X} =$ **O**PEN, **C**LOSED

$\mathcal{A} =$ **P**ULL, **L**EAVE

$\mathcal{Z} =$ **O**PEN, **C**LOSED

Prediction: Given action, propagate belief through dynamics

$$\overline{Bel}(x_t) = \sum_{x_{t-1}} P(x_t | u_t, x_{t-1}) Bel(x_{t-1})$$

$$\underbrace{\begin{bmatrix} 0.74 \\ 0.26 \end{bmatrix}}_{\overline{Bel}(x_t)} = \underbrace{\begin{bmatrix} 0.8 & 0.7 \\ 0.2 & 0.3 \end{bmatrix}}_{P(.|., \mathbf{P})} \underbrace{\begin{bmatrix} 0.4 \\ 0.6 \end{bmatrix}}_{Bel(x_{t-1})}$$

Example: Opening a Door

$\mathcal{X} =$ **O**PEN, **C**LOSED

$\mathcal{A} =$ **P**ULL, **L**EAVE

$\mathcal{Z} =$ **O**PEN, **C**LOSED

$$\overline{Bel}(x_t) = \begin{bmatrix} 0.74 \\ 0.26 \end{bmatrix}$$

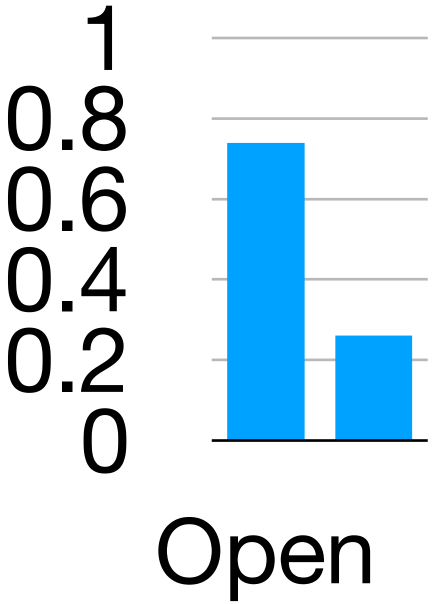

CLOSED

Example: Opening a Door

$\mathcal{X} =$ **O**PEN, **C**LOSED

$\mathcal{A} =$ **P**ULL, **L**EAVE

$\mathcal{Z} =$ **O**PEN, **C**LOSED

Correction: Given measurement, apply Bayes' rule

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

$$\begin{bmatrix} P(x_t = \mathbf{O}) \\ P(x_t = \mathbf{C}) \end{bmatrix} = \eta \begin{bmatrix} P(z_t|\mathbf{O}) \\ P(z_t|\mathbf{C}) \end{bmatrix} * \begin{bmatrix} P(x_t = \mathbf{O}) \\ P(x_t = \mathbf{C}) \end{bmatrix}$$

$Bel(x_t)$ $\qquad P(\mathbf{C}|.)$ $\qquad \overline{Bel}(x_t)$

Example: Opening a Door

$\mathcal{X} =$ **O**PEN, **C**LOSED

$\mathcal{A} =$ **P**ULL, **L**EAVE

$\mathcal{Z} =$ **O**PEN, **C**LOSED

Correction: Given measurement, apply Bayes' rule

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

$$\underbrace{\begin{bmatrix} P(x_t = \mathbf{O}) \\ P(x_t = \mathbf{C}) \end{bmatrix}}_{Bel(x_t)} = \eta \begin{bmatrix} 0.4 \\ 0.8 \end{bmatrix} * \underbrace{\begin{bmatrix} 0.74 \\ 0.26 \end{bmatrix}}_{\overline{Bel}(x_t)} = \eta \begin{bmatrix} 0.296 \\ 0.208 \end{bmatrix} = \begin{bmatrix} 0.58 \\ 0.42 \end{bmatrix}$$

Example: Opening a Door

$\mathcal{X} =$ **O**PEN, **C**LOSED

$\mathcal{A} =$ **P**ULL, **L**EAVE

$\mathcal{Z} =$ **O**PEN, **C**LOSED

$$Bel(x_t) = \begin{bmatrix} 0.58 \\ 0.42 \end{bmatrix}$$

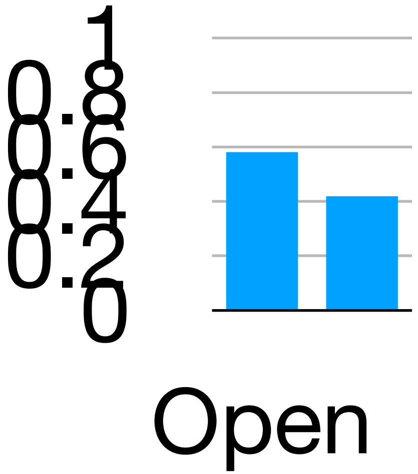

- Robot initially thought the door was open with 0.4 prob
- Robot took the PULL action, then thought the door was open with 0.74 prob
- Robot received a CLOSED measurement, now thinks open with 0.58 prob

Robot lost in a 1-D hallway

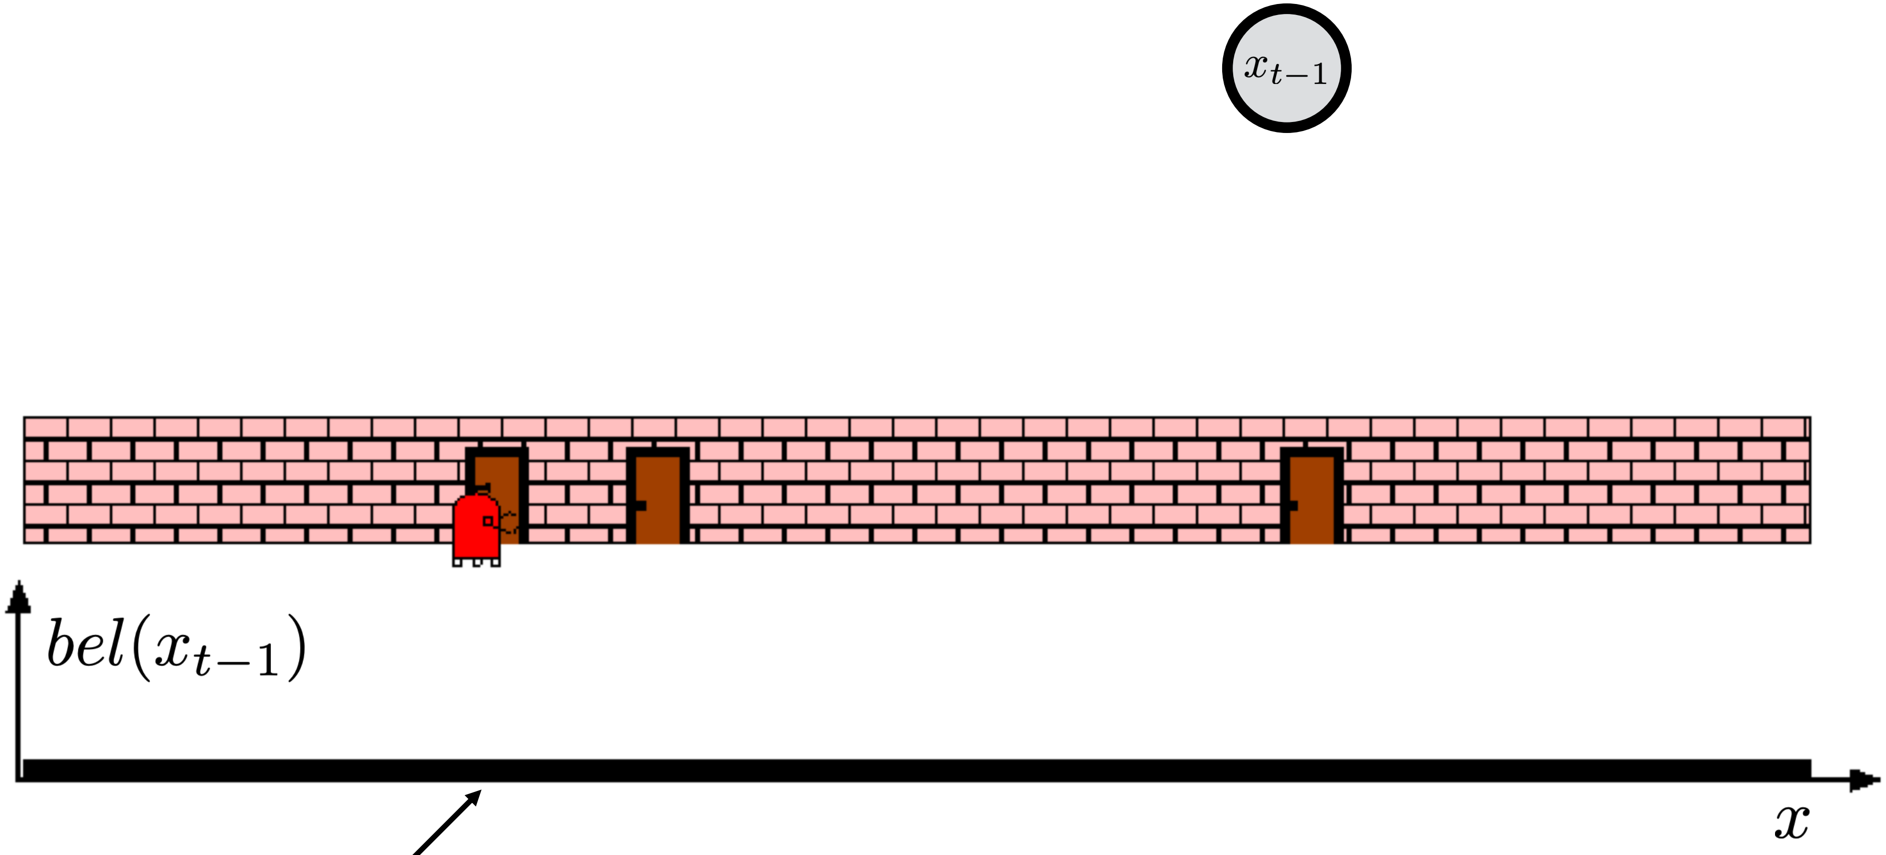

Uniform (robot could be anywhere)

Action at time t: NOP

$$u_t = \text{NOP}$$

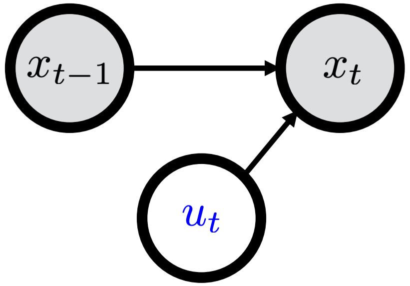

$$P(x_t|u_t, x_{t-1}) = \delta(x_t = x_{t-1})$$

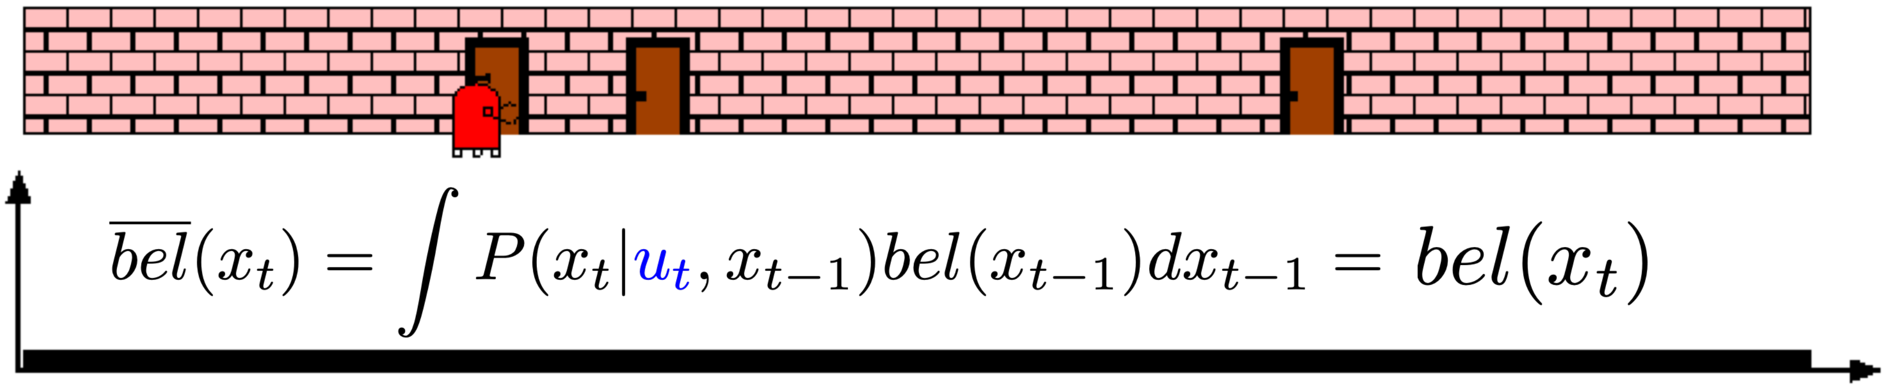

$$\overline{bel}(x_t) = \int P(x_t|u_t, x_{t-1}) bel(x_{t-1}) dx_{t-1} = bel(x_t)$$

x

NOP action implies belief remains the same!

(still uniform — no idea where I am)

Measurement at time t: “Door”

$$z_t = \text{Door}$$

$$P(z_t|x_t) = \mathcal{N}(\text{door centre}, 0.75m)$$

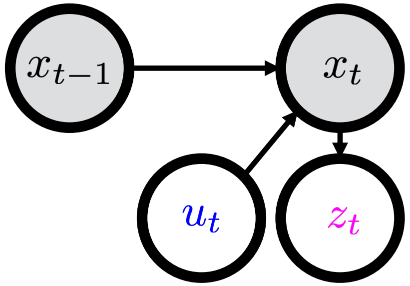

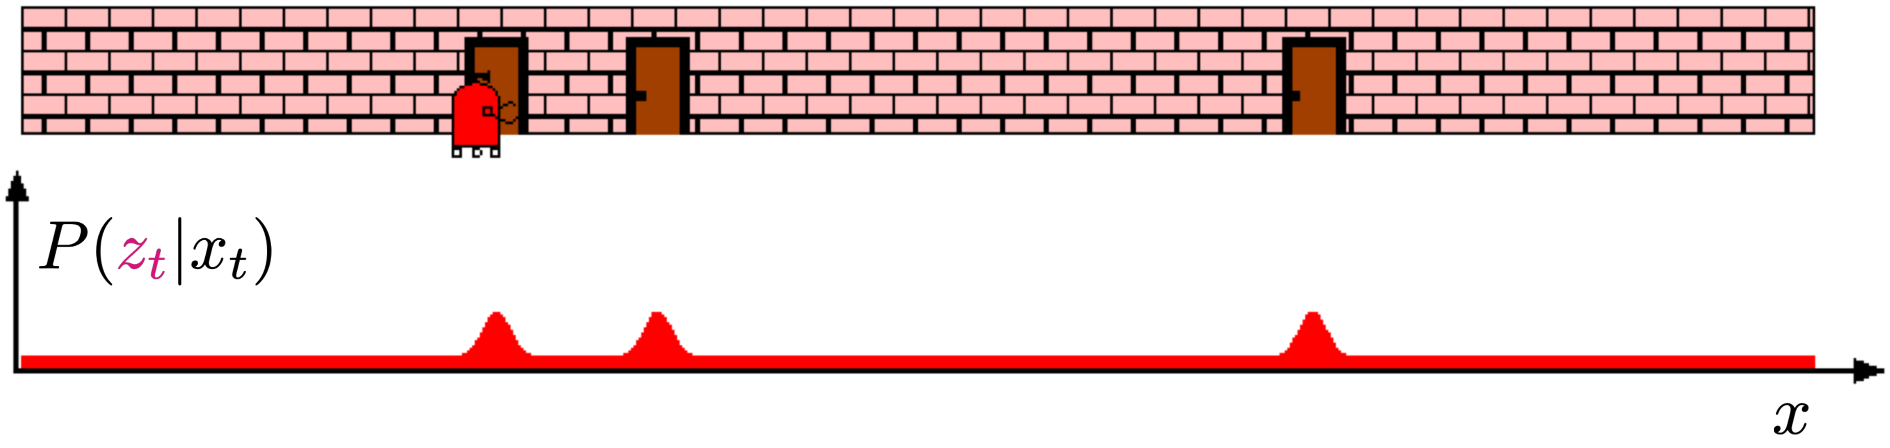

Action at time t+1: Move 3m right

$u_{t+1} = \text{3m right}$

$$P(x_{t+1}|u_{t+1}, x_t) = \mathcal{N}(x_t + u_{t+1}, 0.25m)$$

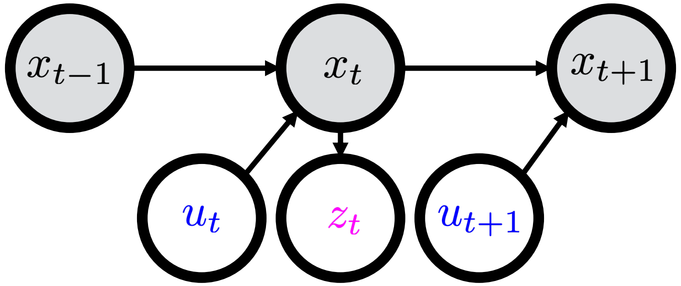

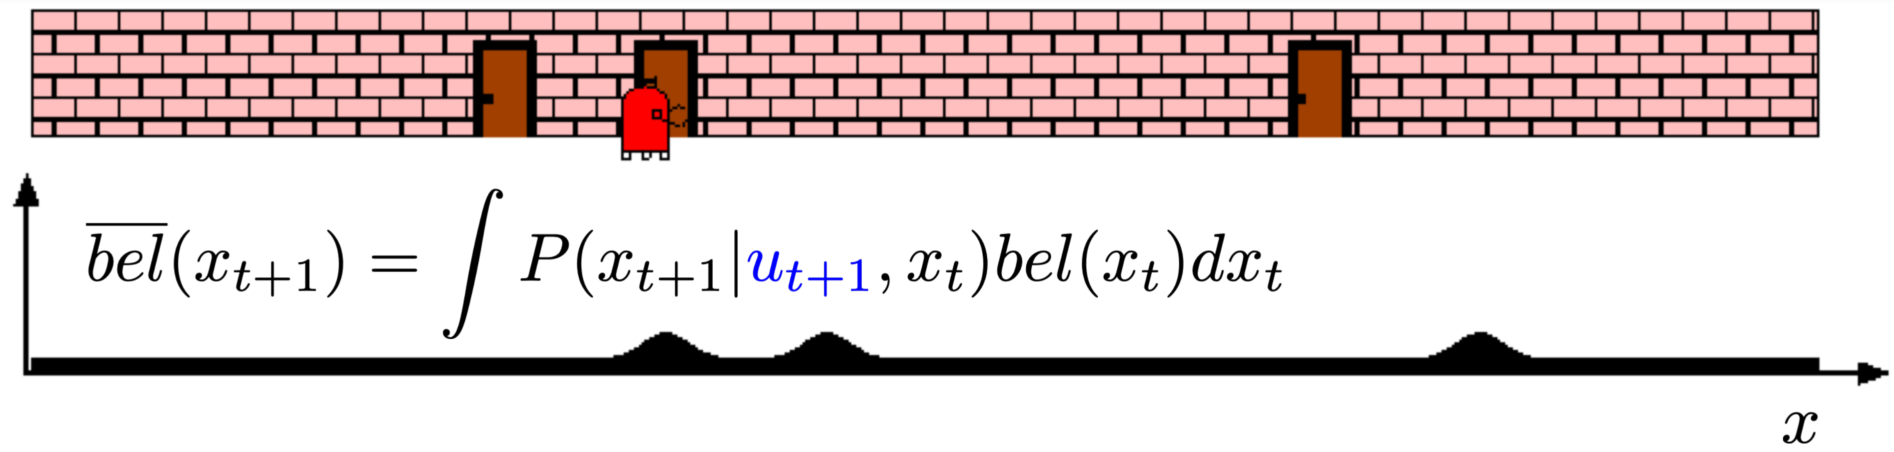

$$\overline{bel}(x_{t+1}) = \int P(x_{t+1}|u_{t+1}, x_t)bel(x_t)dx_t$$

Measurement at time t+1: "Door"

$z_{t+1} = \text{Door}$

$$P(z_{t+1}|x_{t+1}) = \mathcal{N}(\text{door centre}, 0.75m)$$

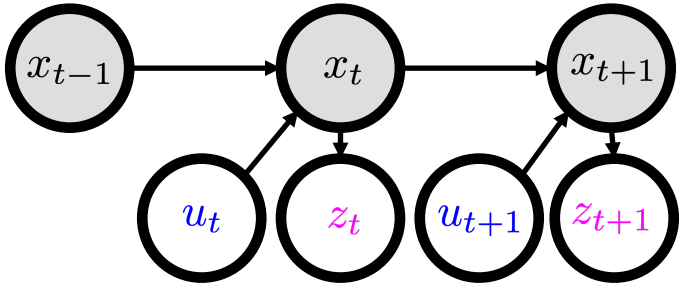

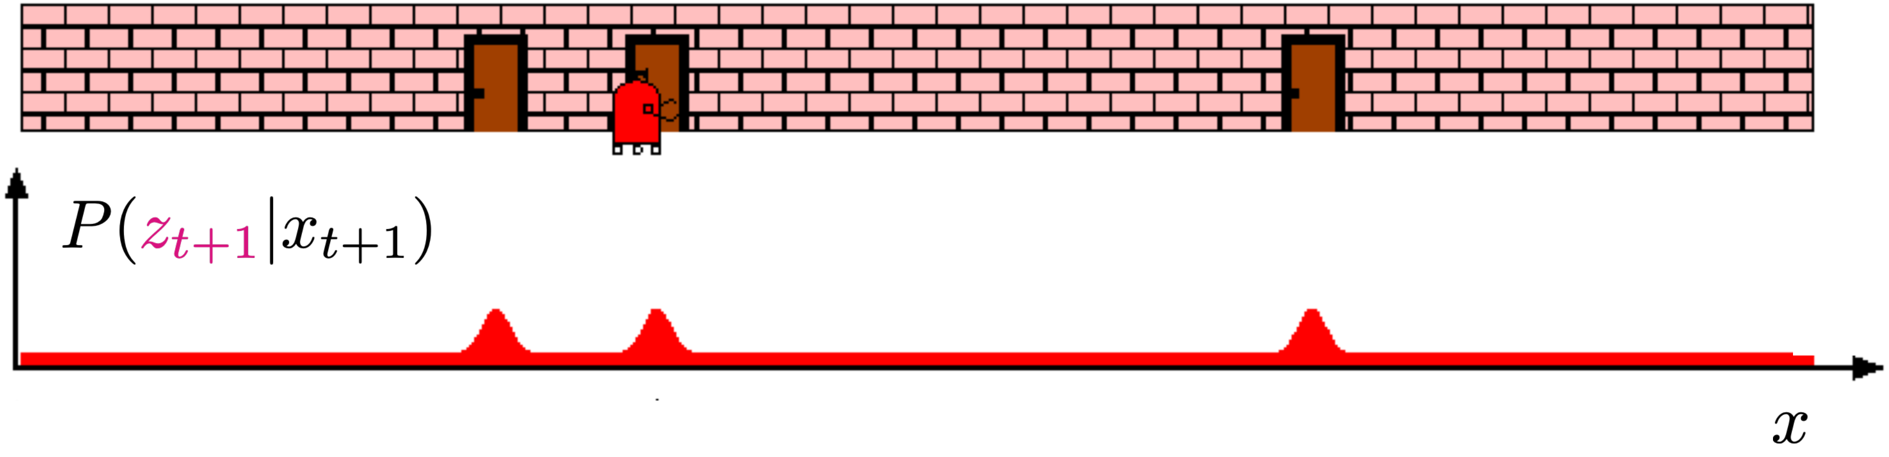

Do actions always increase uncertainty?

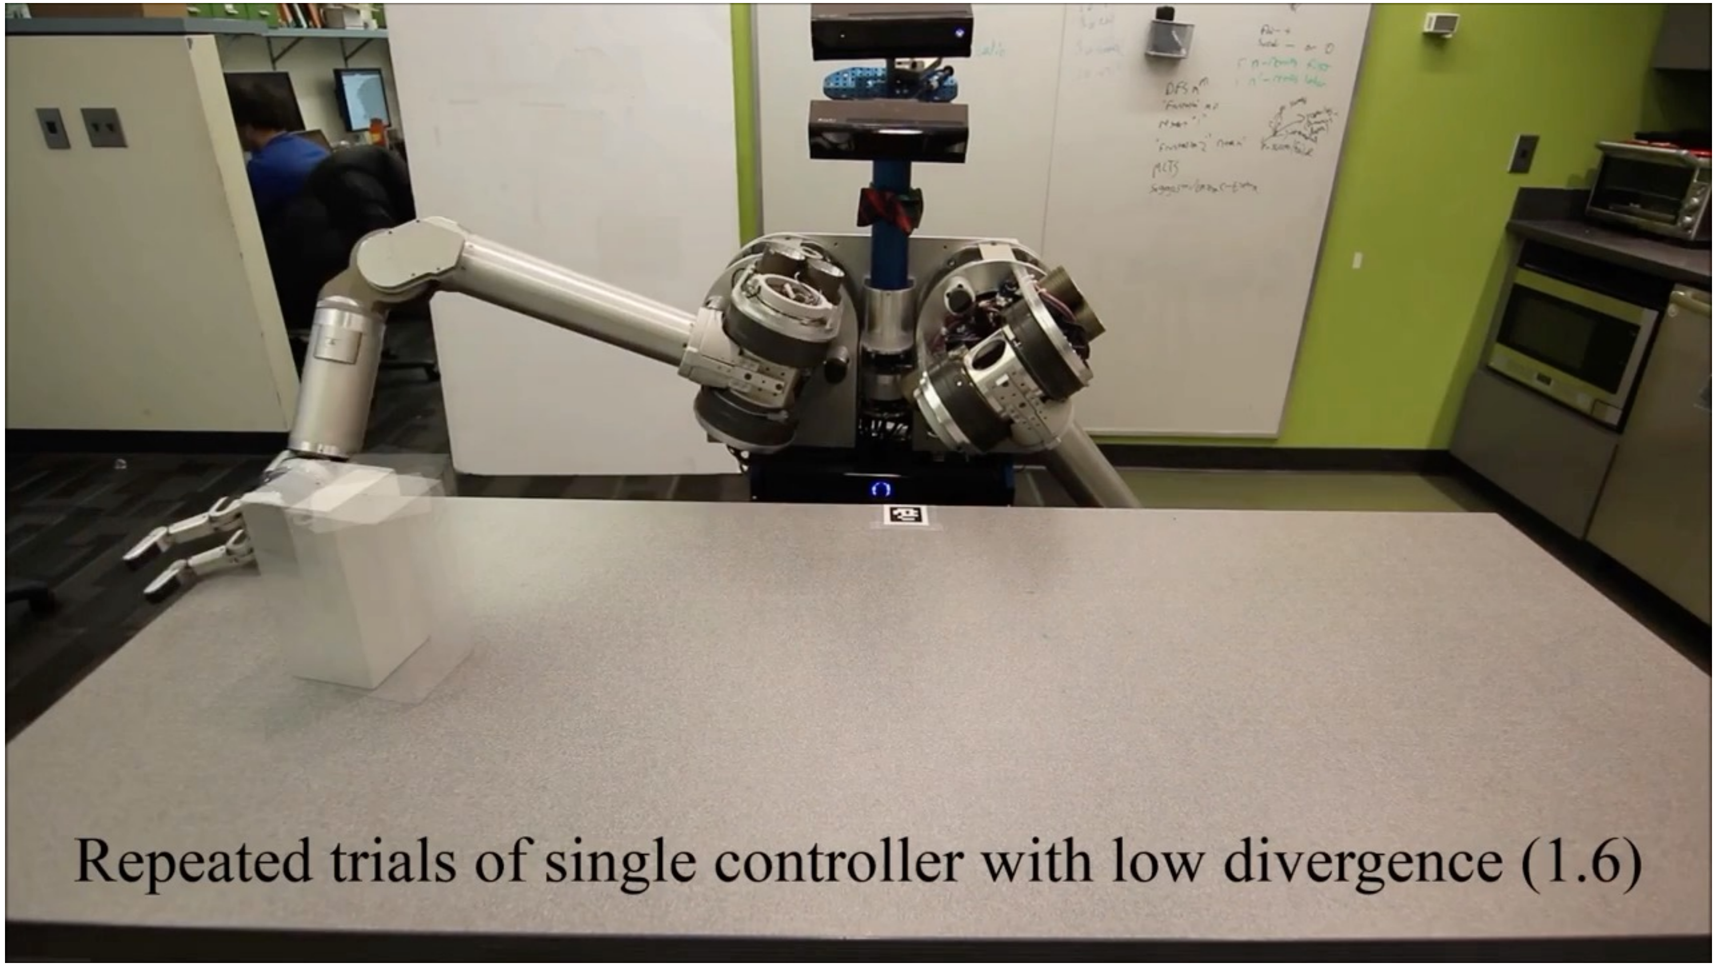

HTTPS://YOUTU.BE/BC1ATPRPPJC

Do measurements always reduce uncertainty?

- Level of uncertainty can be formalized as **entropy**
  - Low entropy if belief is tightly concentrated (e.g., concentrated on one state)
  - High entropy if belief is very spread out (e.g., uniform distribution)
- What if you reach into your pocket and can't find your keys?
  - Initially: low entropy (belief concentrated around pocket, some probability in other states around the house)
  - After: high entropy (very little probability in pocket, other states around the house have increased probability)

Ok this seems simple? What makes this hard!

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

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

How does this connect back to our racecar?

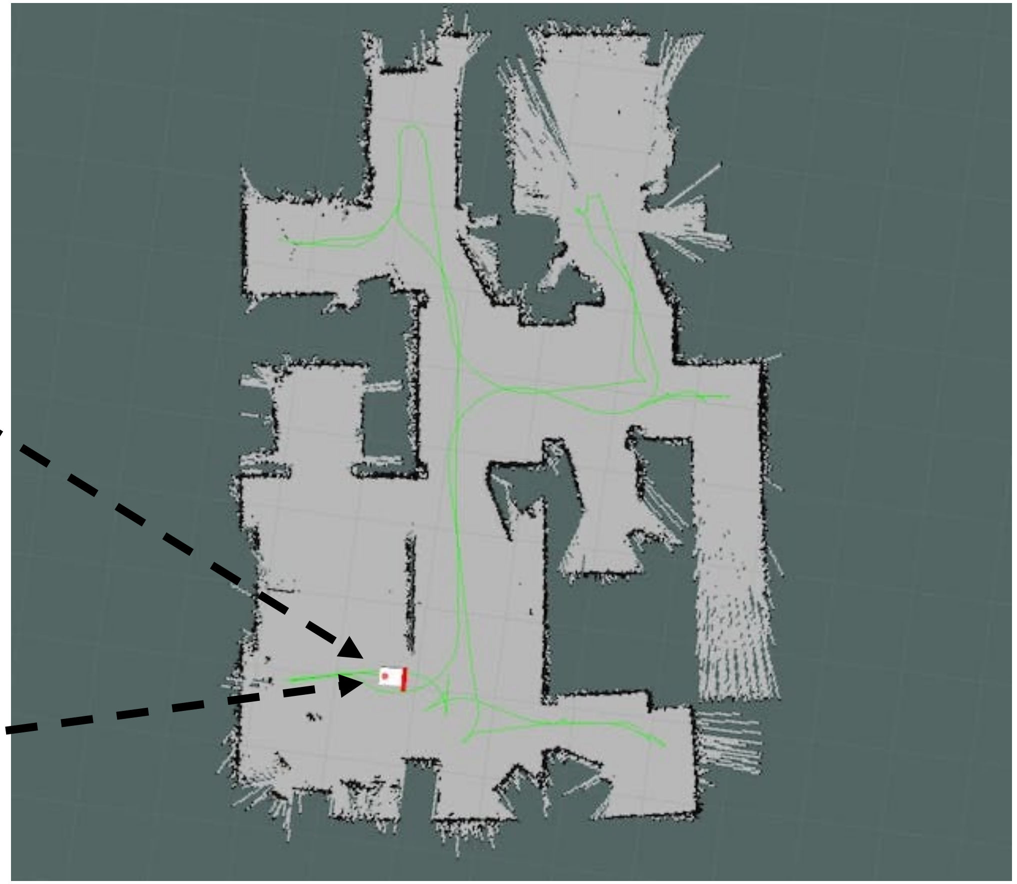

Where am I in the world?

Lecture Outline

What is state estimation?

↓

Probability Review and Bayes Rule

↓

Bayesian Filtering w/ Examples

Class Outline

State Estimation

- Robotic System Design
- Filtering
- Localization
- SLAM

Control

- Feedback Control
- PID Control
- MPC
- LQR

Planning

- Search
- Heuristic Search
- Motion Planning
- Lazy Search

Learning

- Imitation Learning
- Policy Gradient
- Actor-Critic
- Model-Based RL