Why state estimation?

- "State" is an extremely hard thing to define and measure

  Usually unobservable (only "measurements" are observable)

- State can be a choice

  More detailed state, less uncertainty, but harder to estimate

  Less detailed state, more uncertainty, harder to estimate

x_t

Pose/velocity of the object

Position and momentum of all particles

vs

- Definition of state is a modeling **choice**

Fundamental Problem: State is hidden

All the robot sees is a stream of actions and measurements

But robot never sees the state

Fundamental Problem: State is hidden

But all decision making depends on knowing state

Estimate belief over state

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?

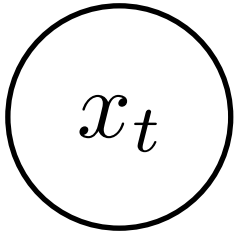

Pose/velocity of the object

- 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

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

- 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

If A and B have no overlap then

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

Discrete Random Variables

- denotes a random variable.
- can take on a countable number of values in $\{x_1, x_2, \ldots, x_n\}$.
- , or , is the probability that the random variable X takes on value x_i.
- 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

Multinomial

Bernoulli

Poisson

Continuous Random Variables

- denotes a random variable.
- can take on a continuum of values in the support of the probability density function
- , or , 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

Uniform Distribution

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	Continuous case
$\sum_x P(x) = 1$	$\int p(x)\, dx = 1$
$P(x) = \sum_y P(x, y)$	$p(x) = \int p(x, y)\, dy$
$P(x) = \sum_y P(x \mid y) P(y)$	$p(x) = \int p(x \mid y)\, p(y)\, dy$

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

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

$$\Rightarrow$$

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

Bayes Formula

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

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

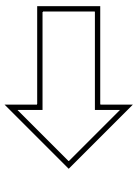

$$P(y, x) = P(y|x)p(x)$$

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

Can replace with integral

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

Example of Bayes Formula in Action

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

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

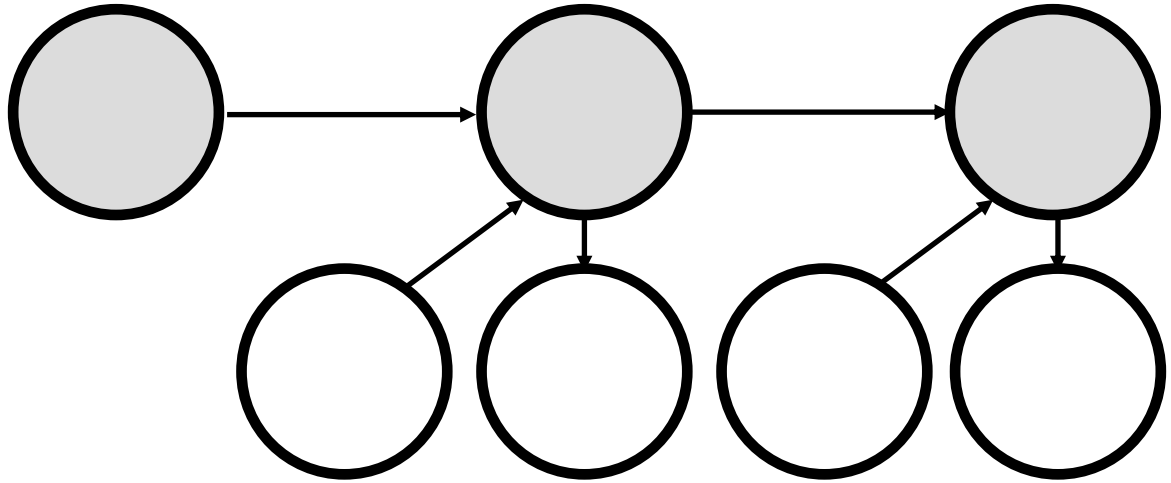

Assumptions:

1. Robot receives a stream of measurements / actions.

2. One measurement / action per time-step.

What is belief in this setting?

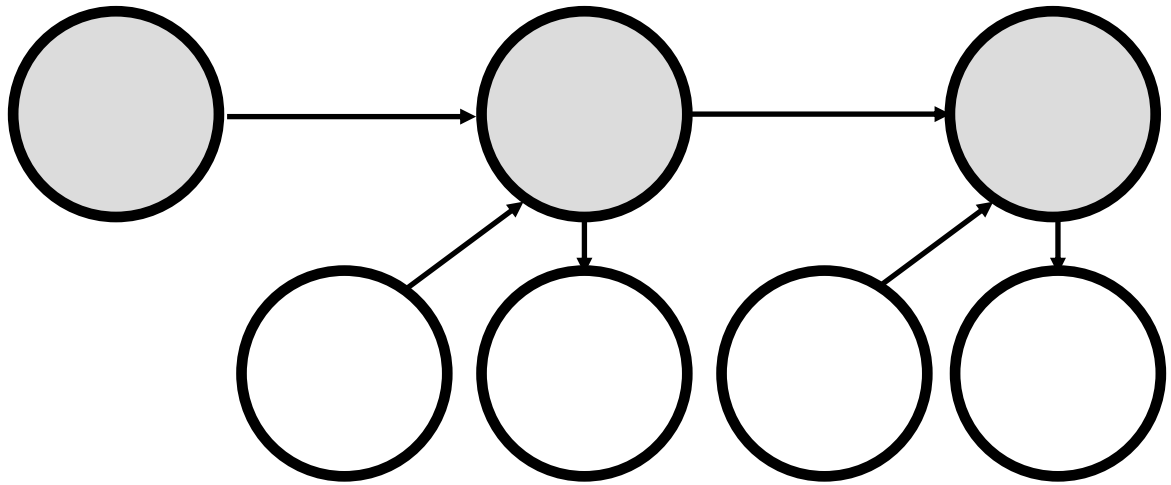

P(current state | all past information)

Can we estimate this?

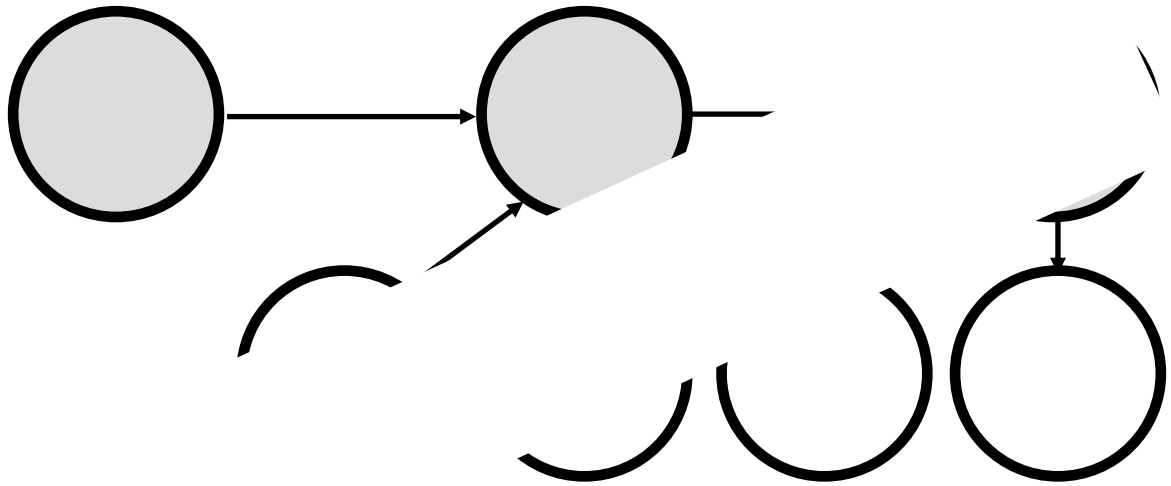

P(current state | all past information)

Good ol’ Markov to the rescue

Andrey Andreyevich Markov (1856 - 1922)

Solution: Markov Assumption

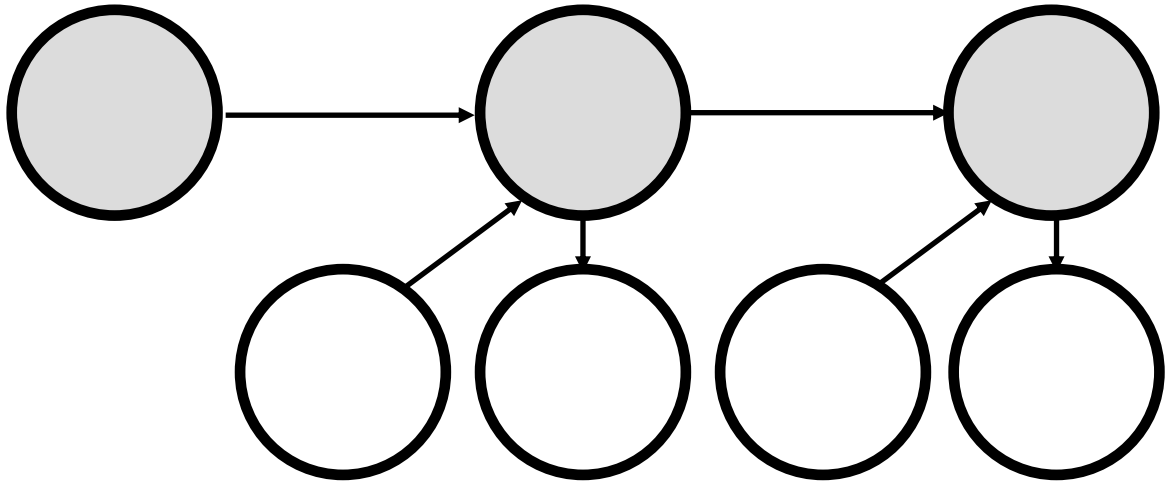

Markov assumption :

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

Probabilistic models

State transition probability / dynamics / motion model

Measurement probability / Observation model

When does Markov not hold?

whenever state doesn’t capture all requisite information

- Unmodelled pedestrians in front of laser

- Steady gusts of wind

How do we tractably calculate belief?

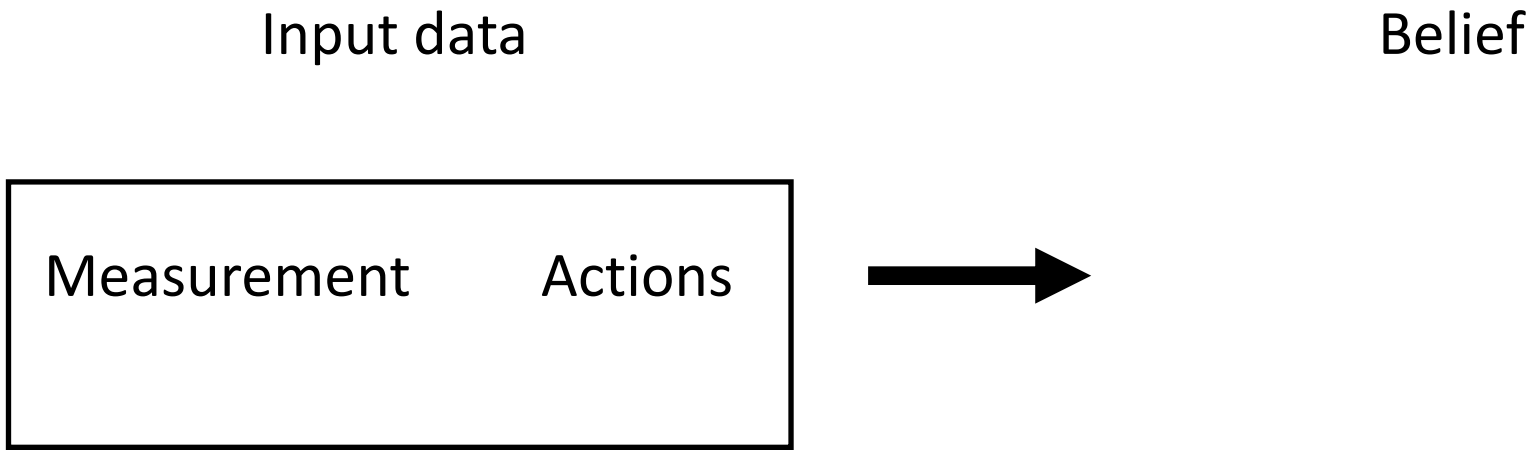

Bayes filter!

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

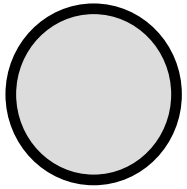

Bayes filter in a nutshell

Step 1: Prediction - push belief through dynamics given action

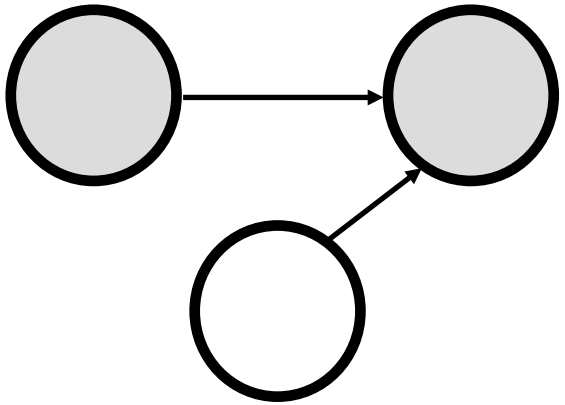

(discrete)

(total probability)

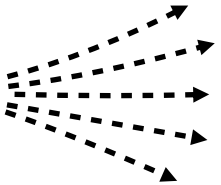

Bayes filter in a nutshell

Step 2: Correction - apply Bayes rule given measurement

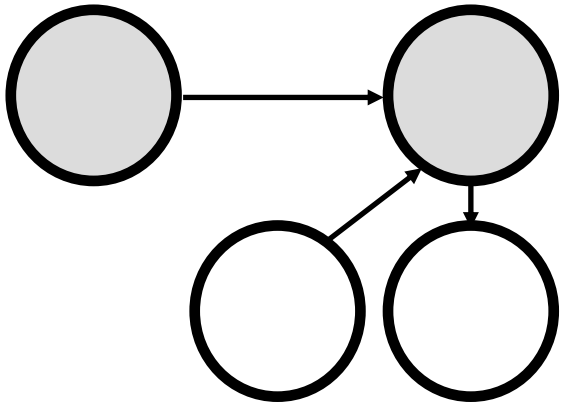

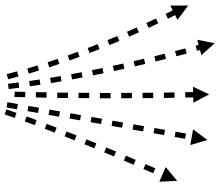

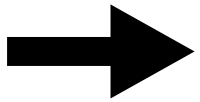

Bayes filter in a nutshell

Key Idea: Apply Markov to get a recursive update!

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

Step 1: Prediction - push belief through dynamics given action

Step 2: Correction - apply Bayes rule given

Bayes filter is a powerful tool

Localization Mapping SLAM POMDP

Example: Opening a Door

OPEN, **C**LOSED

PULL, **L**EAVE

HTTPS://YOUTU.BE/AFUA50H9UEK

Example: Opening a Door

HTTPS://YOUTU.BE/AFUA50H9UEK

Example: Opening a Door

OPEN, **C**LOSED

PULL, **L**EAVE

OPEN, **C**LOSED

HTTPS://YOUTU.BE/AFUA50H9UEK

Example: Opening a Door

OPEN, **C**LOSED

PULL, **L**EAVE

OPEN, CLOSED

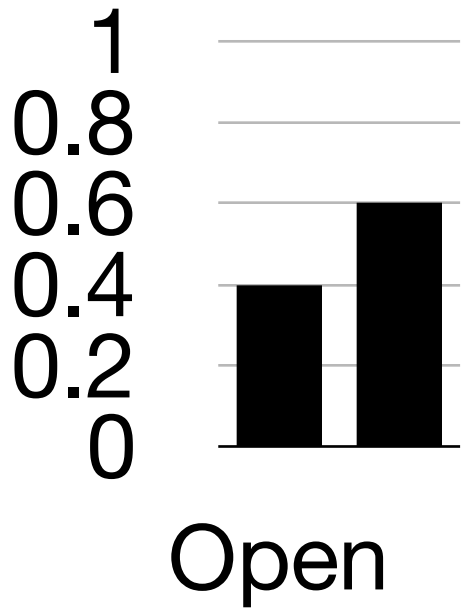

PULL

Example: Opening a Door

OPEN, **C**LOSED

PULL, **L**EAVE

OPEN, **C**LOSED

Prediction: Given action, propagate belief through dynamics

Example: Opening a Door

OPEN, **C**LOSED

PULL, **L**EAVE

OPEN, **C**LOSED

Prediction: Given action, propagate belief through dynamics

Example: Opening a Door

OPEN, **C**LOSED

PULL, **L**EAVE

OPEN, CLOSED

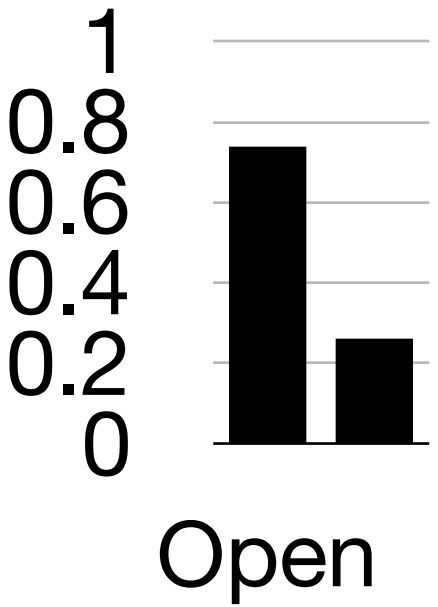

CLOSED

Example: Opening a Door

OPEN, **C**LOSED

PULL, **L**EAVE

OPEN, CLOSED

Correction: Given measurement, apply Bayes' rule

*

Example: Opening a Door

OPEN, **C**LOSED

PULL, **L**EAVE

OPEN, CLOSED

Correction: Given measurement, apply Bayes' rule

*

Example: Opening a Door

OPEN, **C**LOSED

PULL, **L**EAVE

OPEN, CLOSED

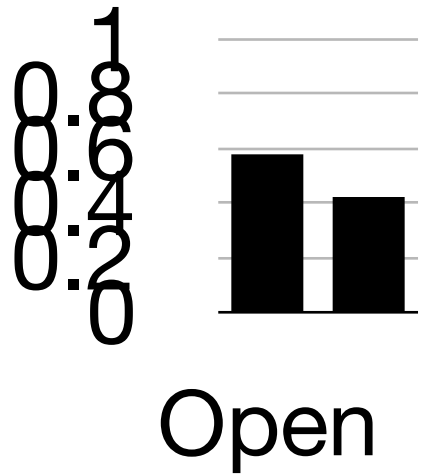

- Robot initially thought the door was open with 0.4 prob
- Robot took the PULL action, then thought the door was open with 0.74 prob
- Robot received a CLOSED measurement, now thinks open with 0.58 prob

Robot lost in a 1-D hallway

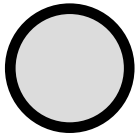

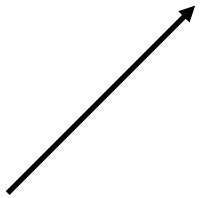

Uniform (robot could be anywhere)

Action at time t: NOP

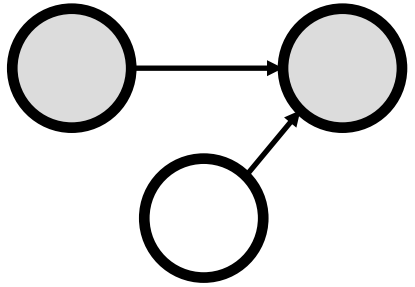

NOP action implies belief remains the same!

(still uniform — no idea where I am)

Measurement at time t: “Door”

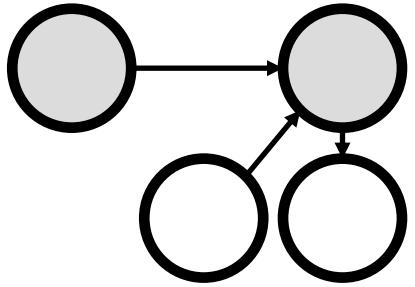

Action at time t+1: Move 3m right

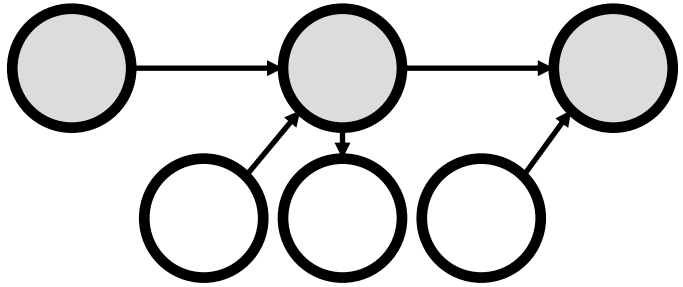

Measurement at time t+1: “Door”

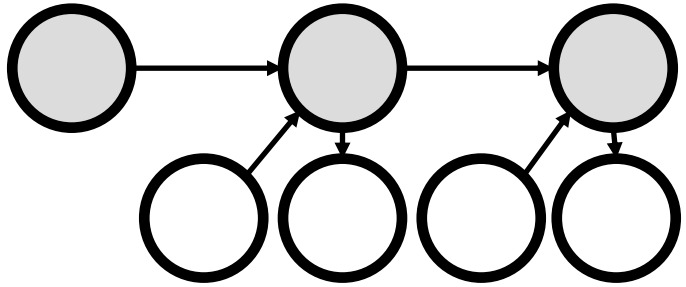

Do actions always increase uncertainty?

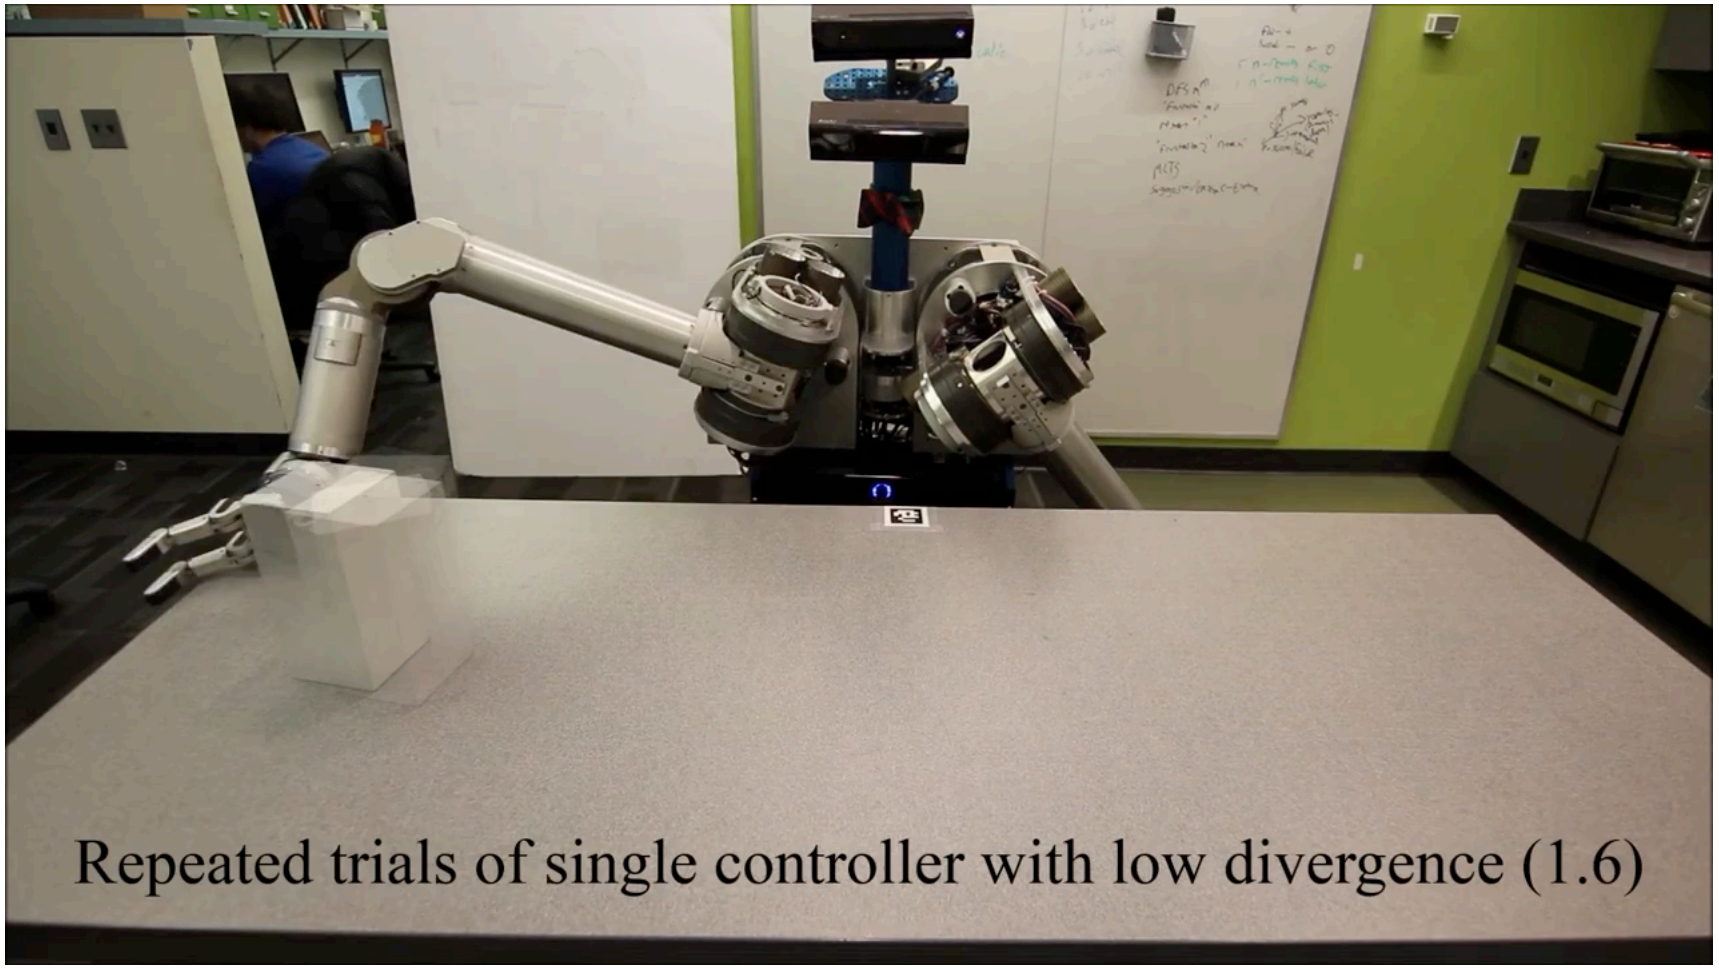

HTTPS://YOUTU.BE/BC1ATPRPPJC

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)

Uniform Distribution

Gaussians

Practical Concerns: Overconfidence

- Once the belief collapses to either 0 or 1, only the motion model can shake it loose
- Too many measurements will collapse belief
- Correlated incorrect measurements are dangerous

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

[redacted]

Imitation Learning

Policy Gradient

Actor-Critic

Model-Based RL