Introduction to State Estimation

Instructor: Chris Mavrogiannis

TAs: Kay Ke, Gilwoo Lee, Matt Schmittle

*Slides based on or adapted from Sanjiban Choudhury

Logistics

- Begin working on Assignment 0!
- Post questions, discuss any issues you are having on piazza.
- Students with **no** access to 022, e-mail cardkey@cs.washington.edu with your student ID and CC me.
- Students that have not been added to the class, come talk to me after the lecture.
- If you did not attend recitation yesterday, come talk to me after the lecture.

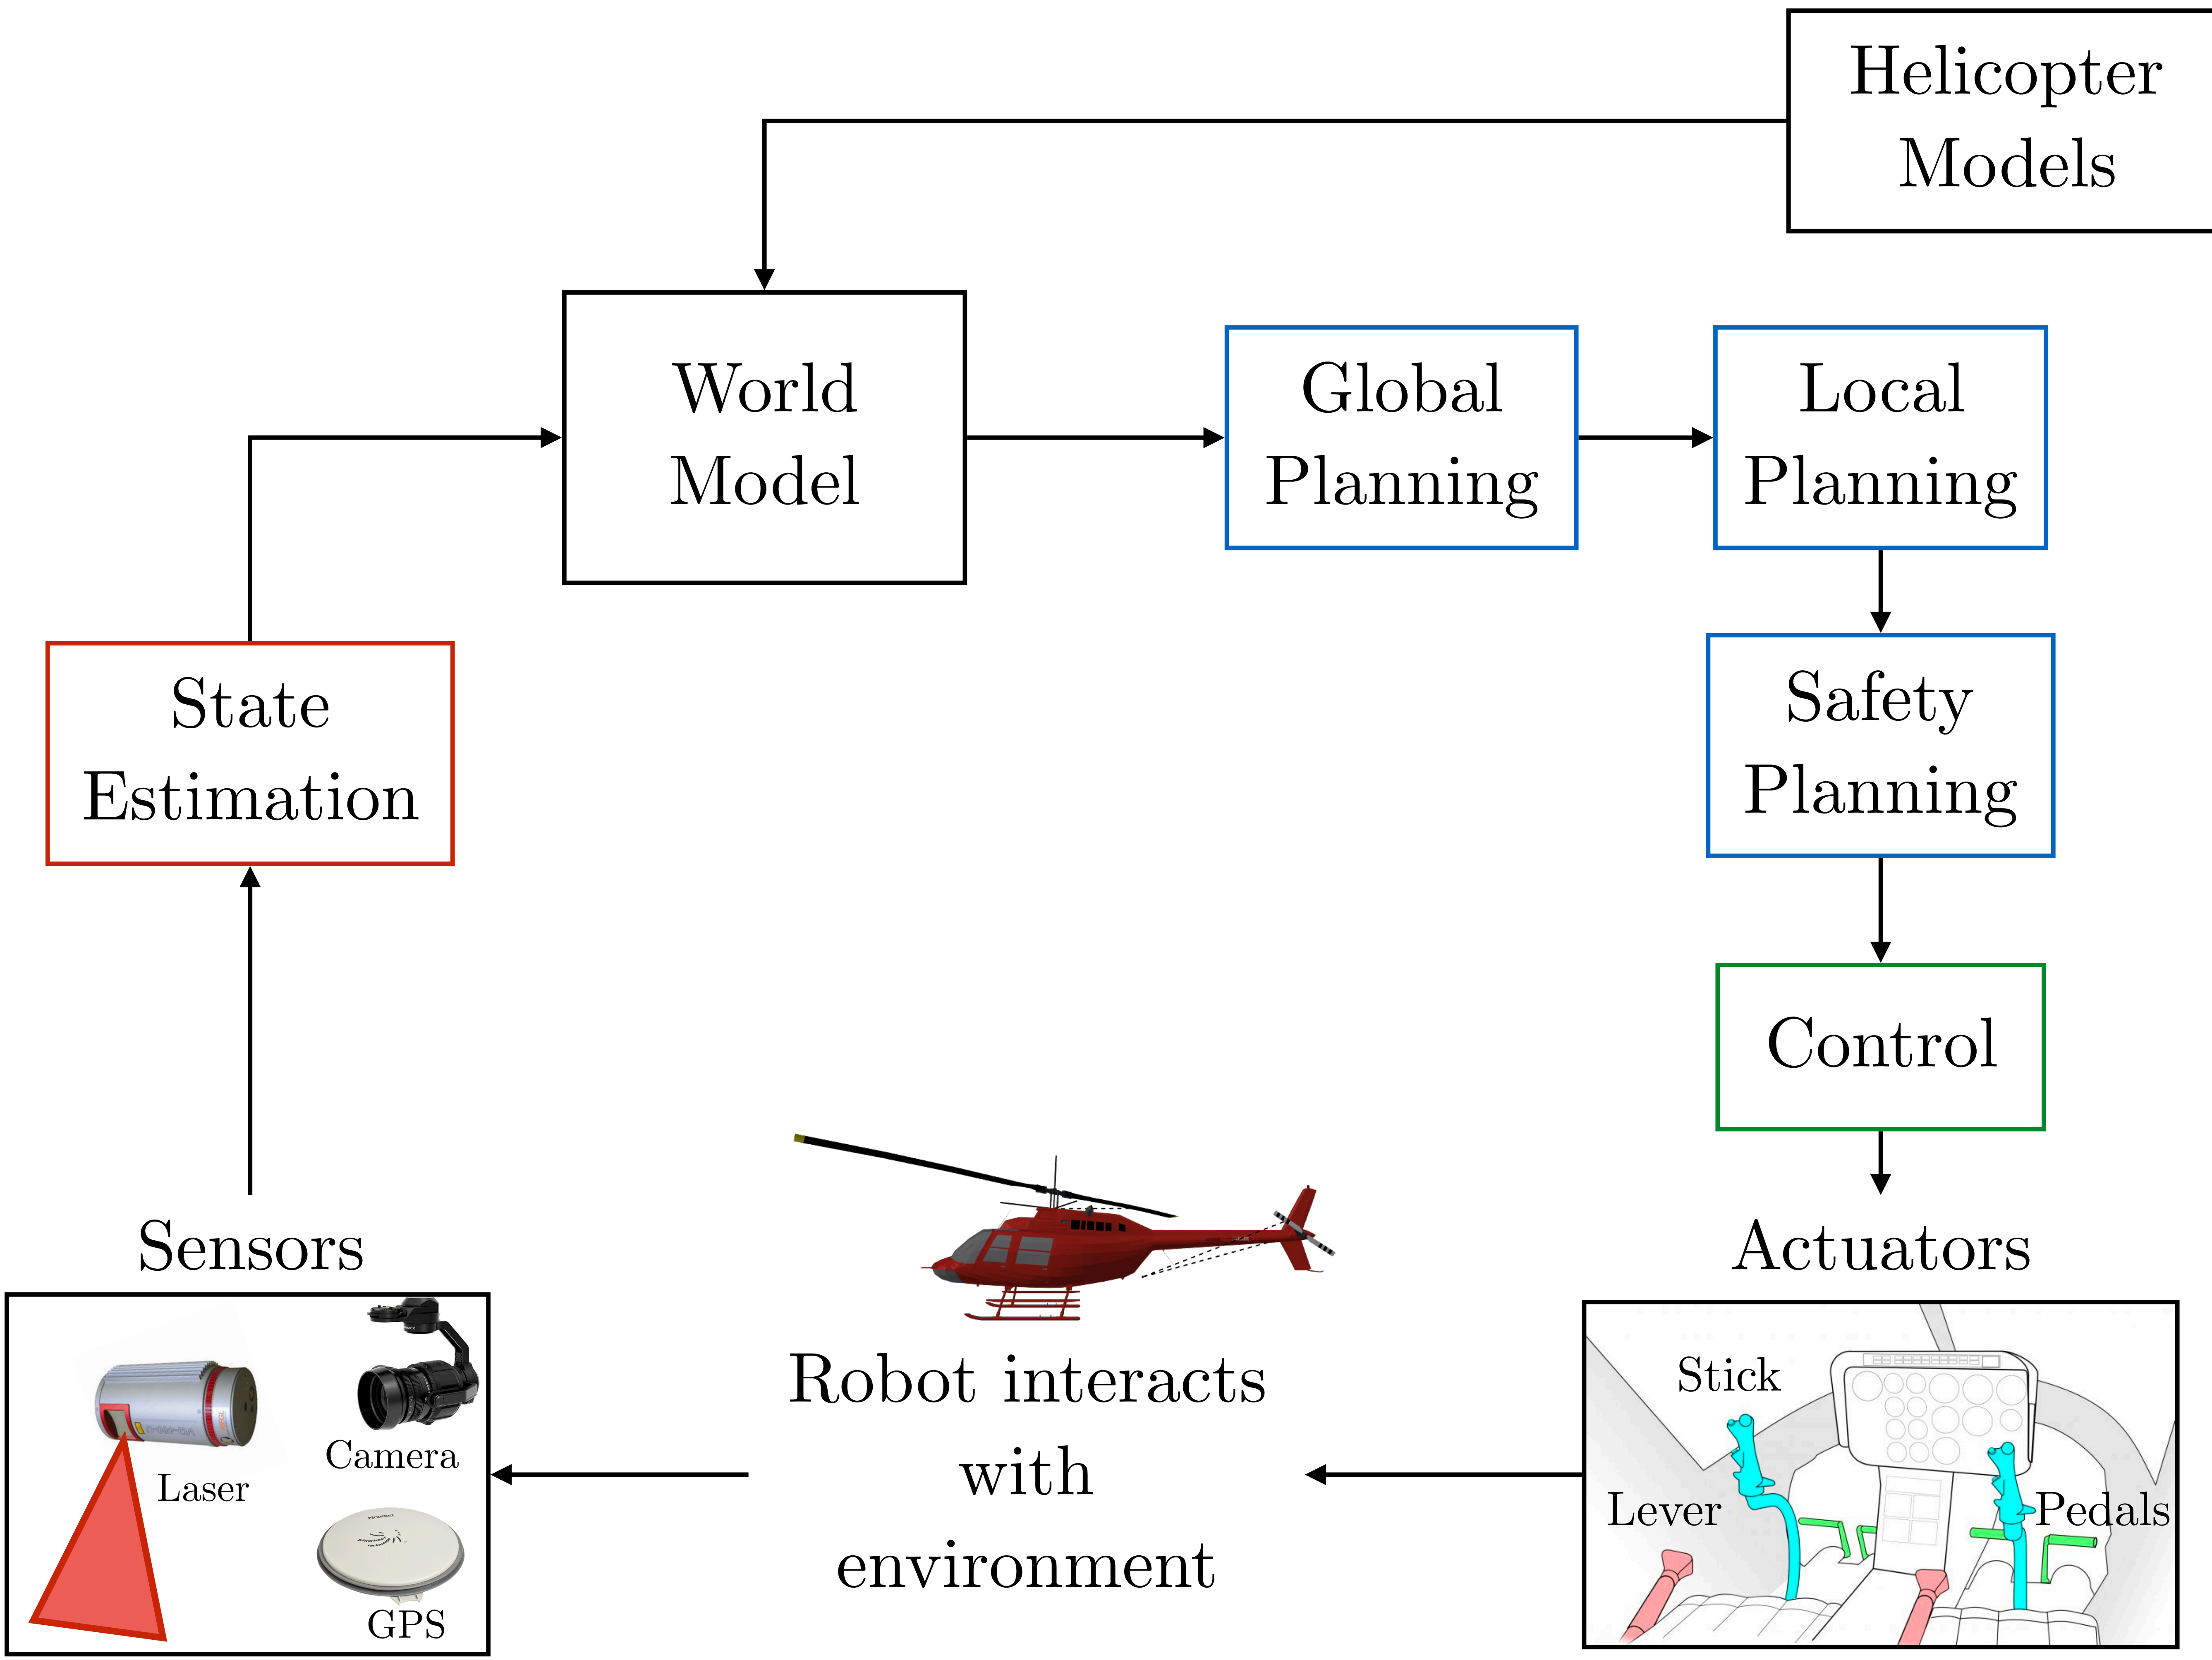
Helicopter Models
World Model
Global Planning
Local Planning
Safety Planning
Control
State Estimation
Sensors
Laser
Camera
GPS
Robot interacts with environment
Actuators
Stick
Lever
Pedals

Estimate state

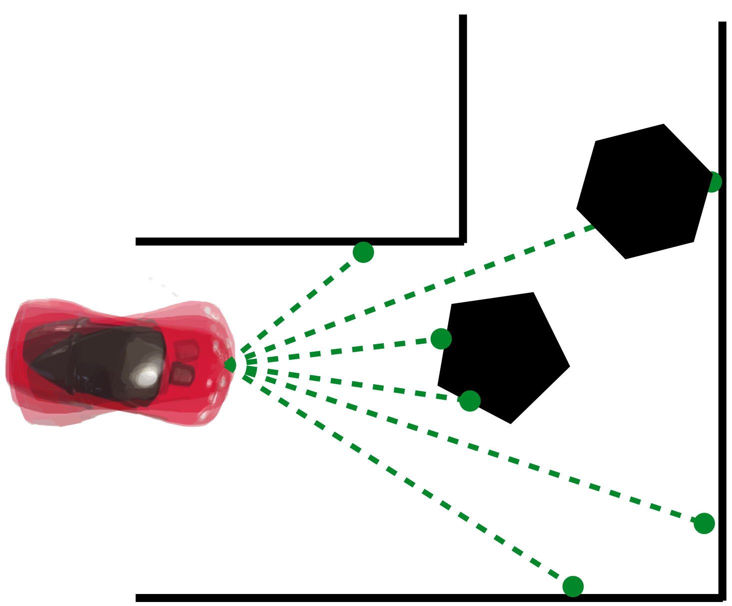

Plan a sequence of motions

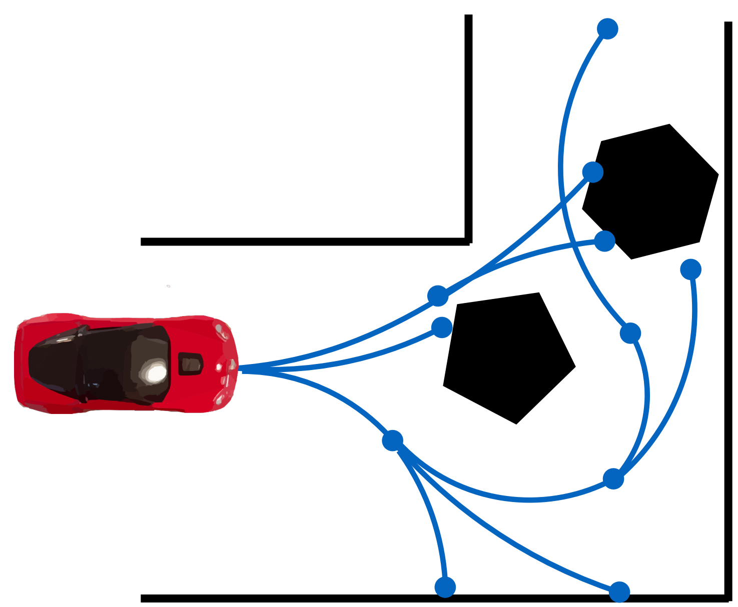

Control robot to follow plan

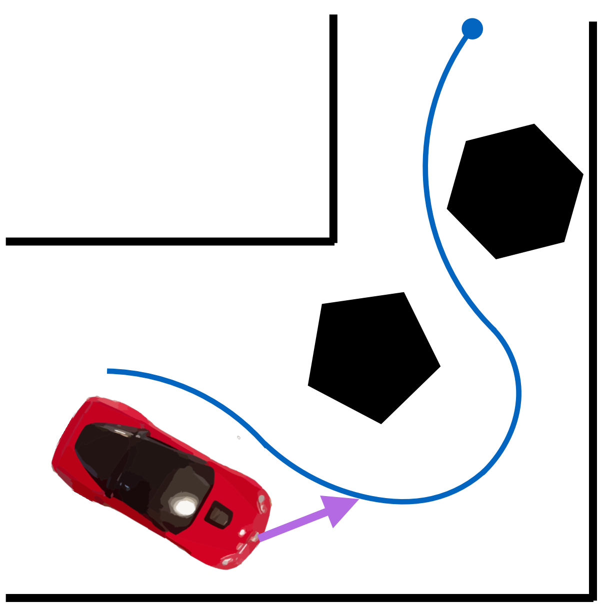

Today's objective

1. Formulate state estimation as a Bayes filtering problem

2. Discuss various components of a Bayes filter

3. Intuition behind Bayes filtering

A robot interacting with the world

Thrun et al. 1998

A robot interacting with the world

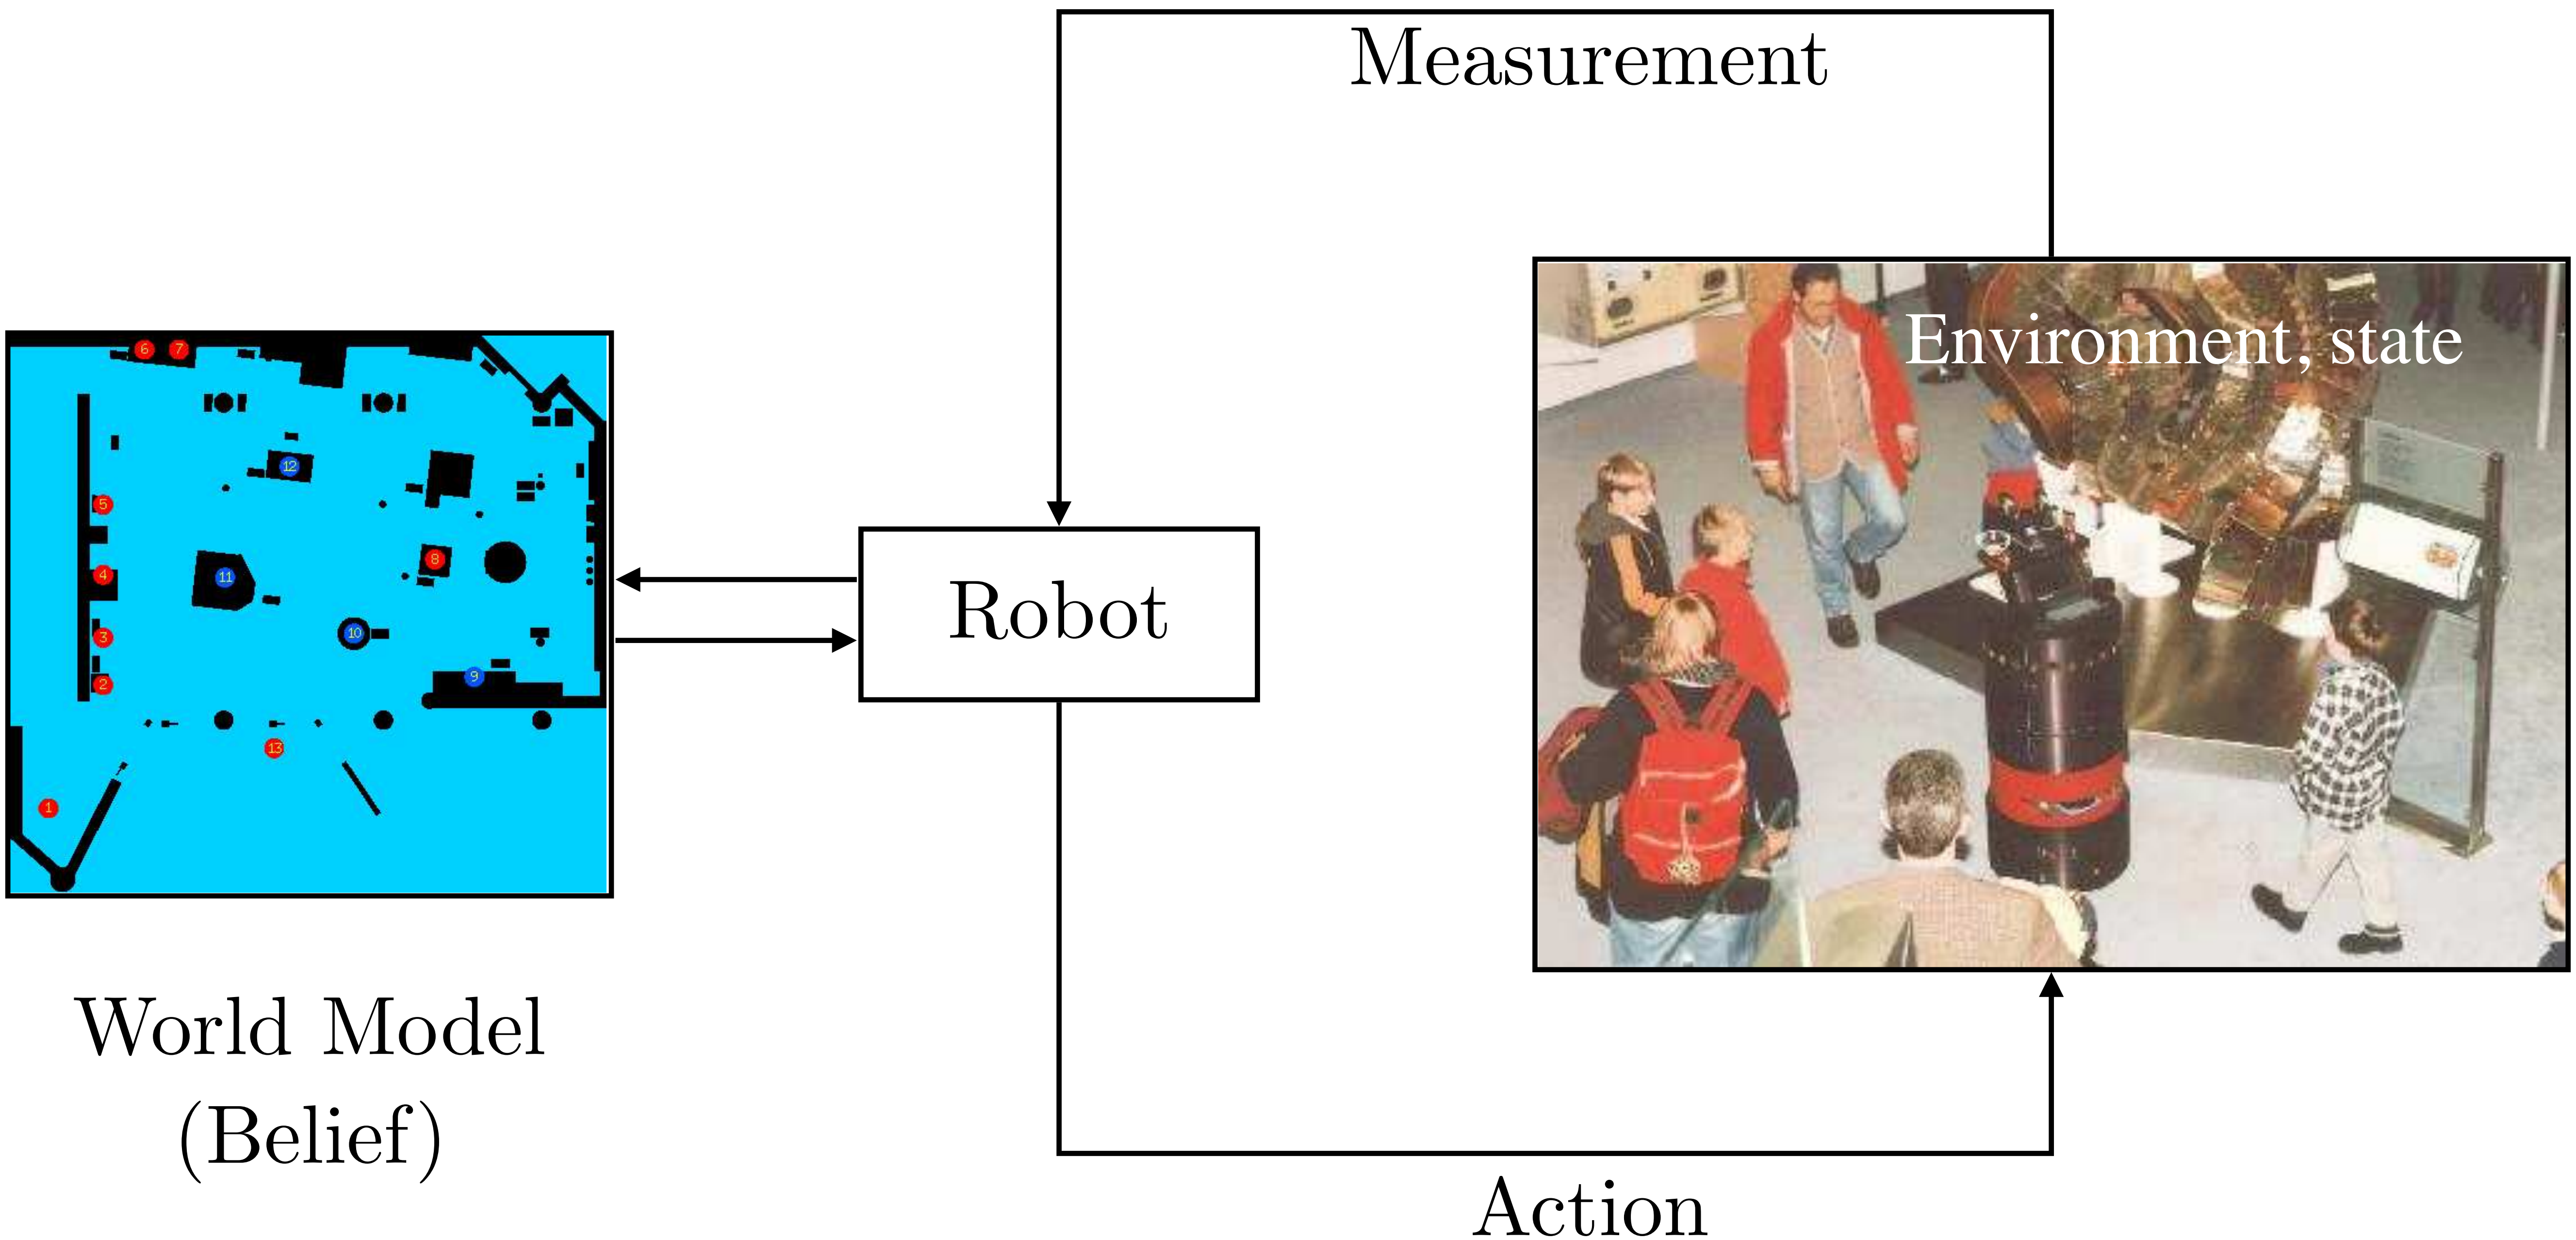

Bayes filtering in action!

Urmson "How a driverless car sees the road" TED (2015)

Example 1: Rich information

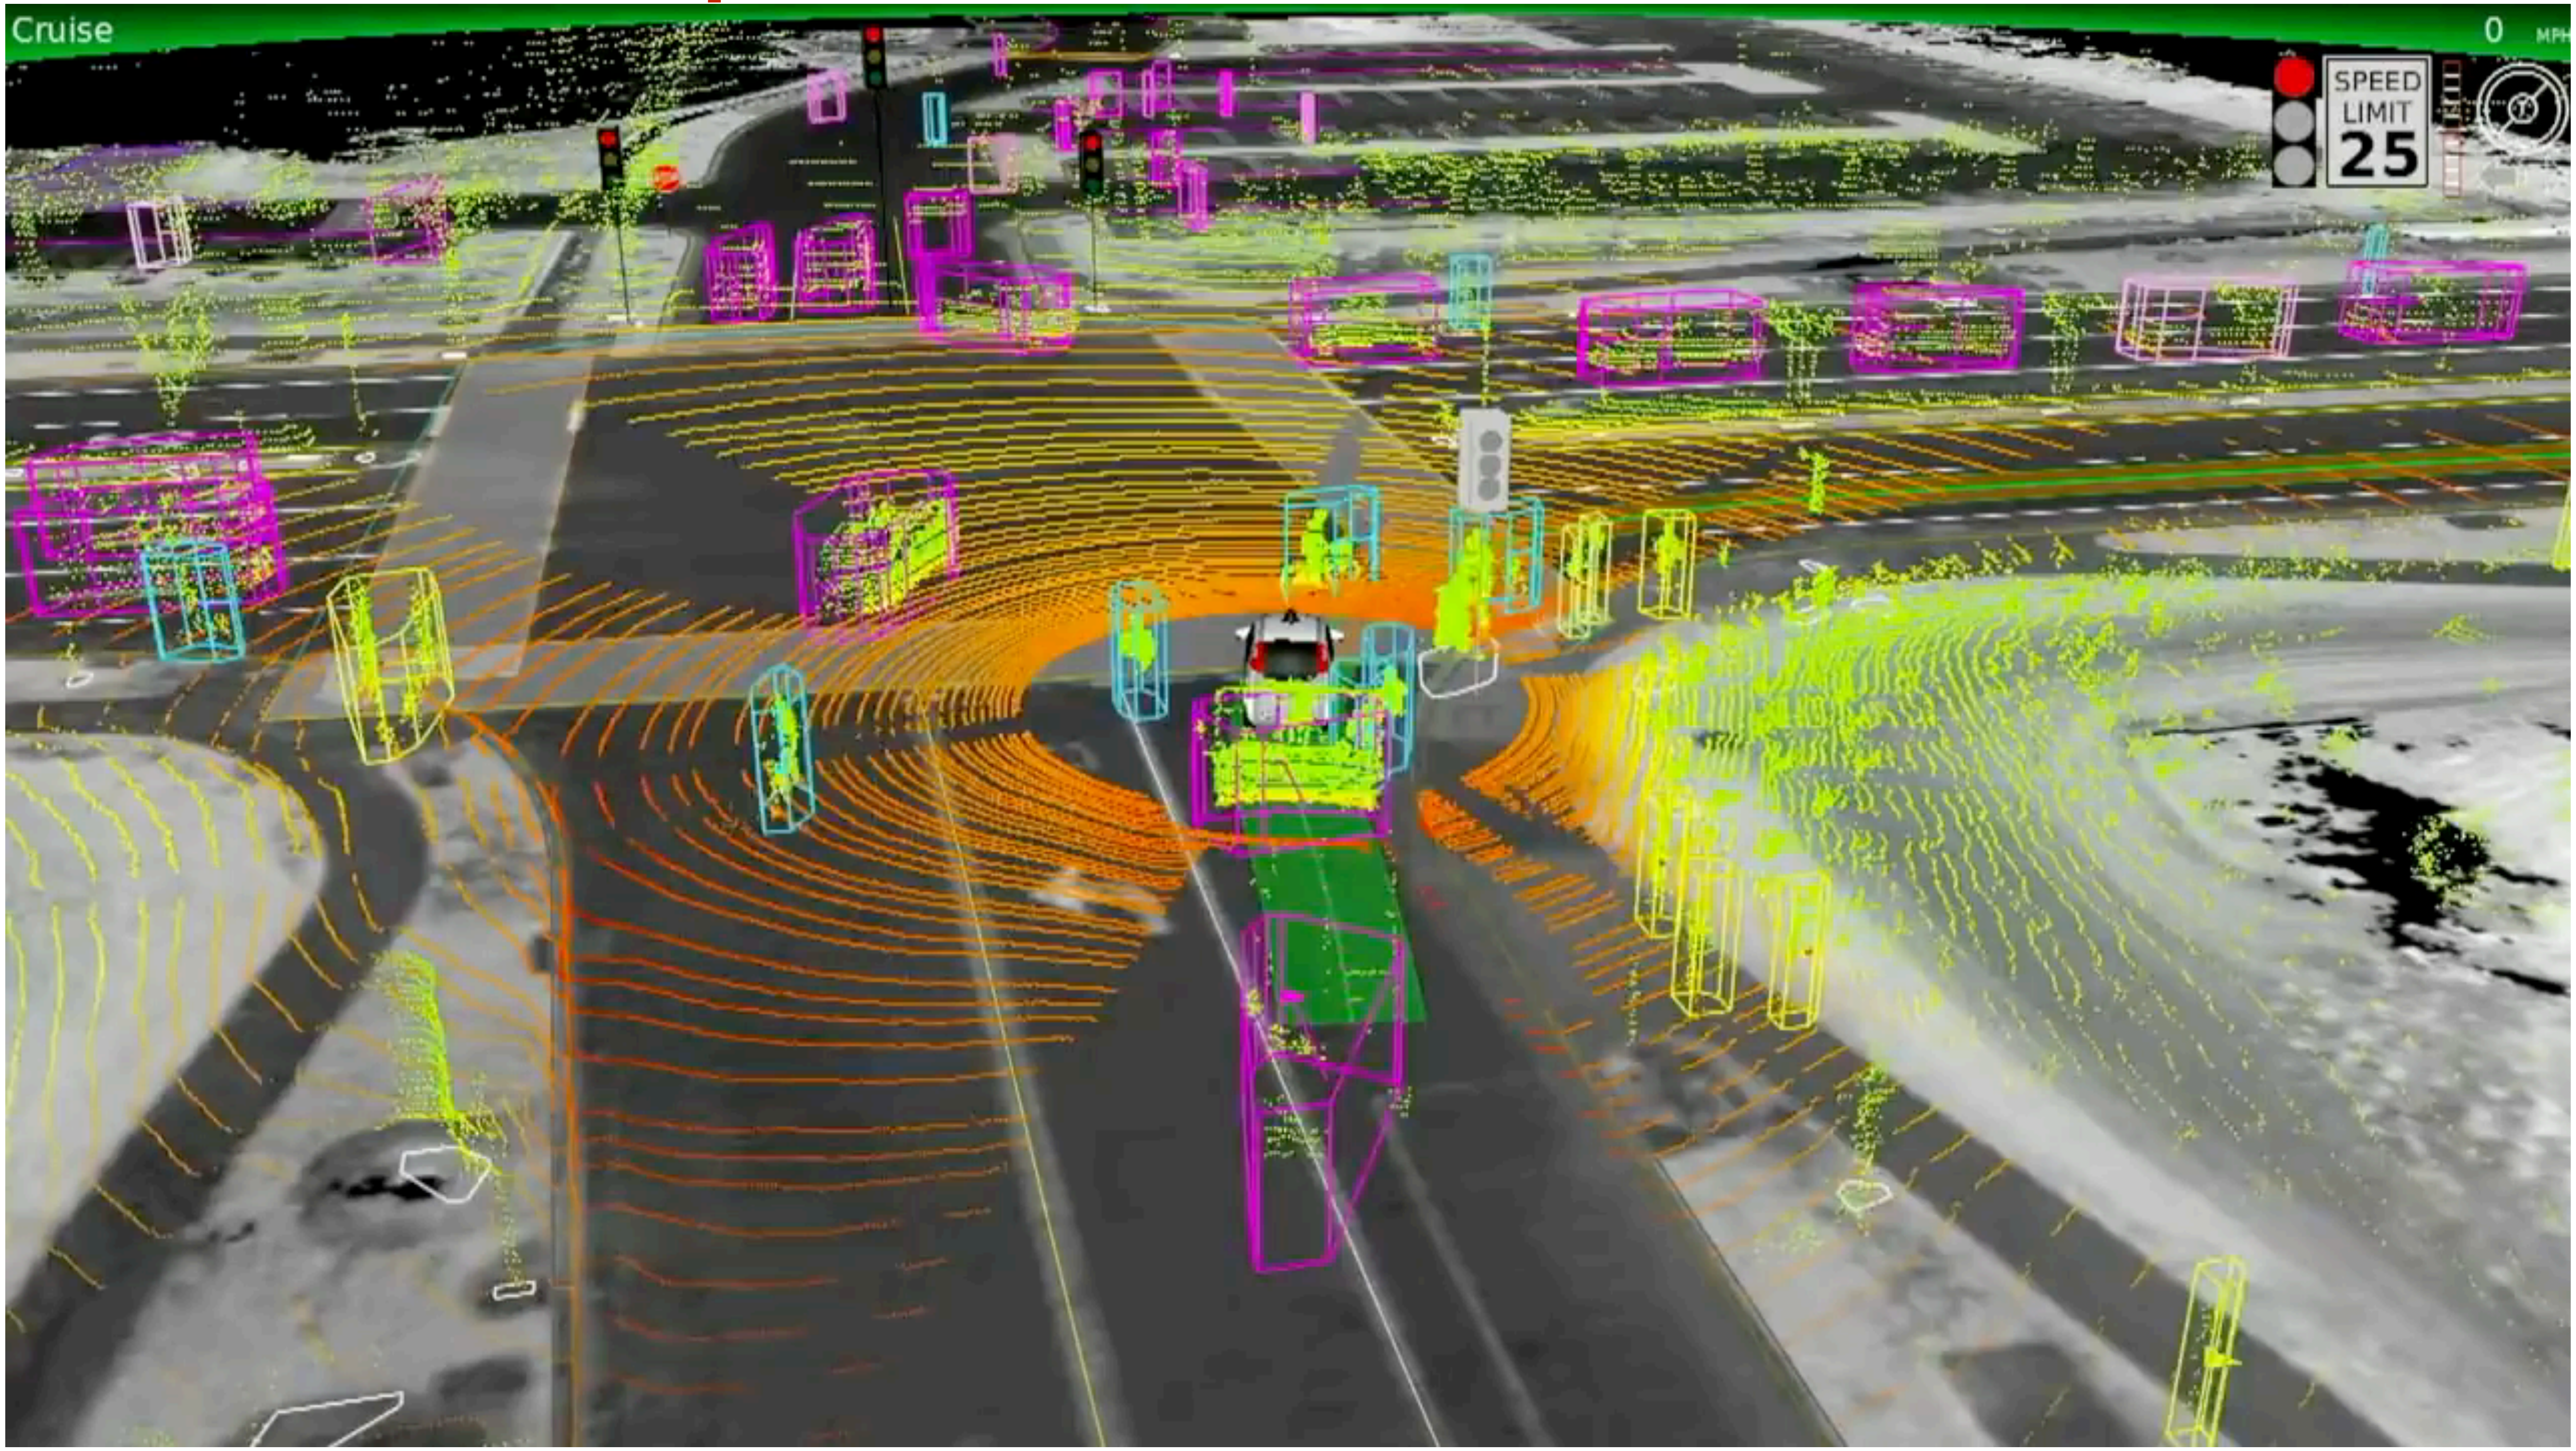

Measurement:	World Model (Belief):	Action:
LiDAR, camera, GPS	Location, cars, people..	Steering, speed

Example 2: Medium information

Measurement: Planar LiDAR

World Model (Belief): State of robot

Action: Pitch, Roll

Guilliard, et al.. Autonomous thermalling as a partially observable Markov decision process. In RSS, 2018

Example 3: Very low information

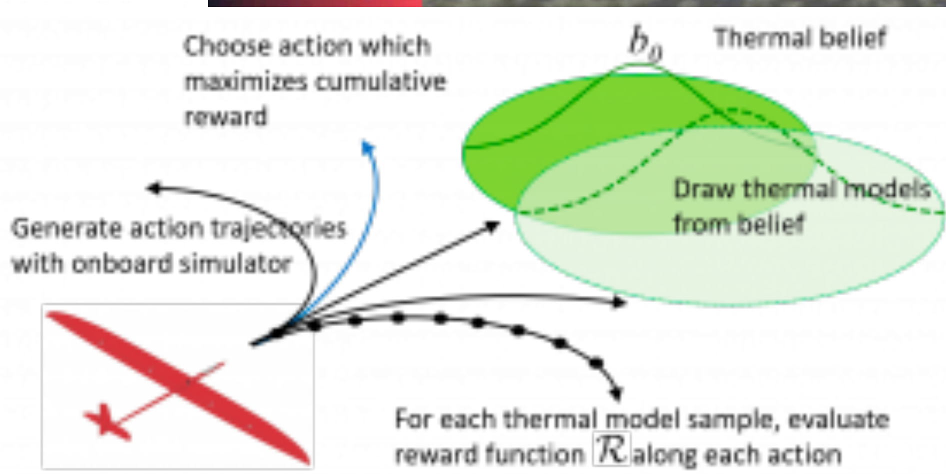

Measurement: Pitot tube (airpseed), GPS

World Model (Belief): Location of wind thermal sensor

Action: Pitch, Roll

A robot interacting with the world

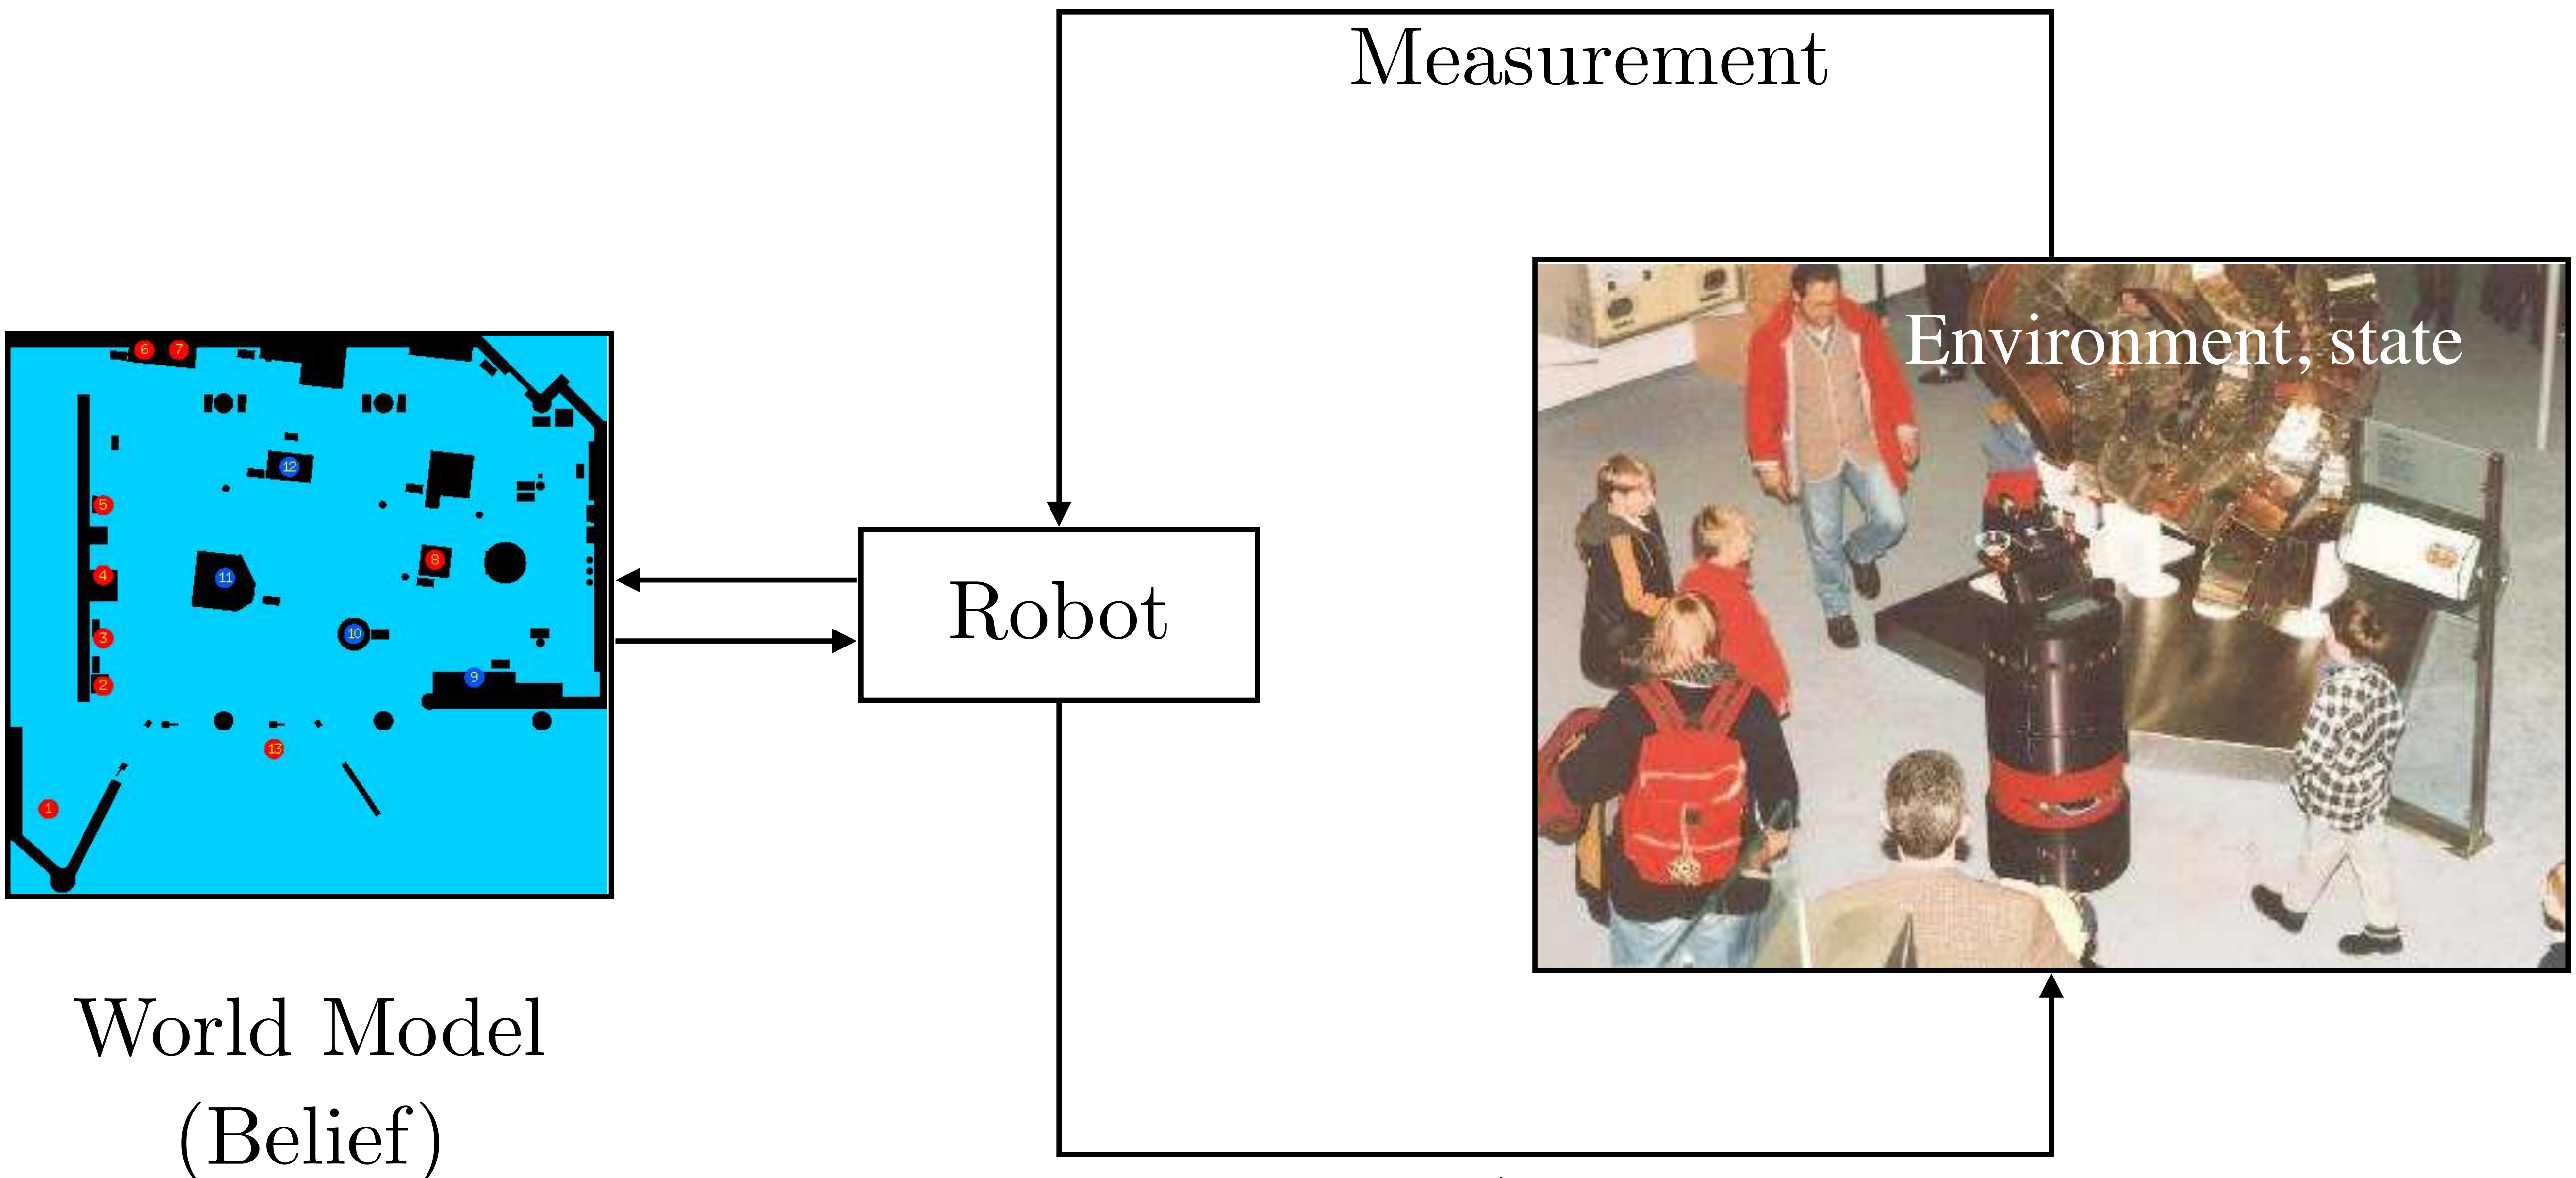

Given data (stream of measurements and actions) ⟶ ESTIMATE ⟶ Belief (probability of states)

How do we formally define this problem?

State: A very abstract definition

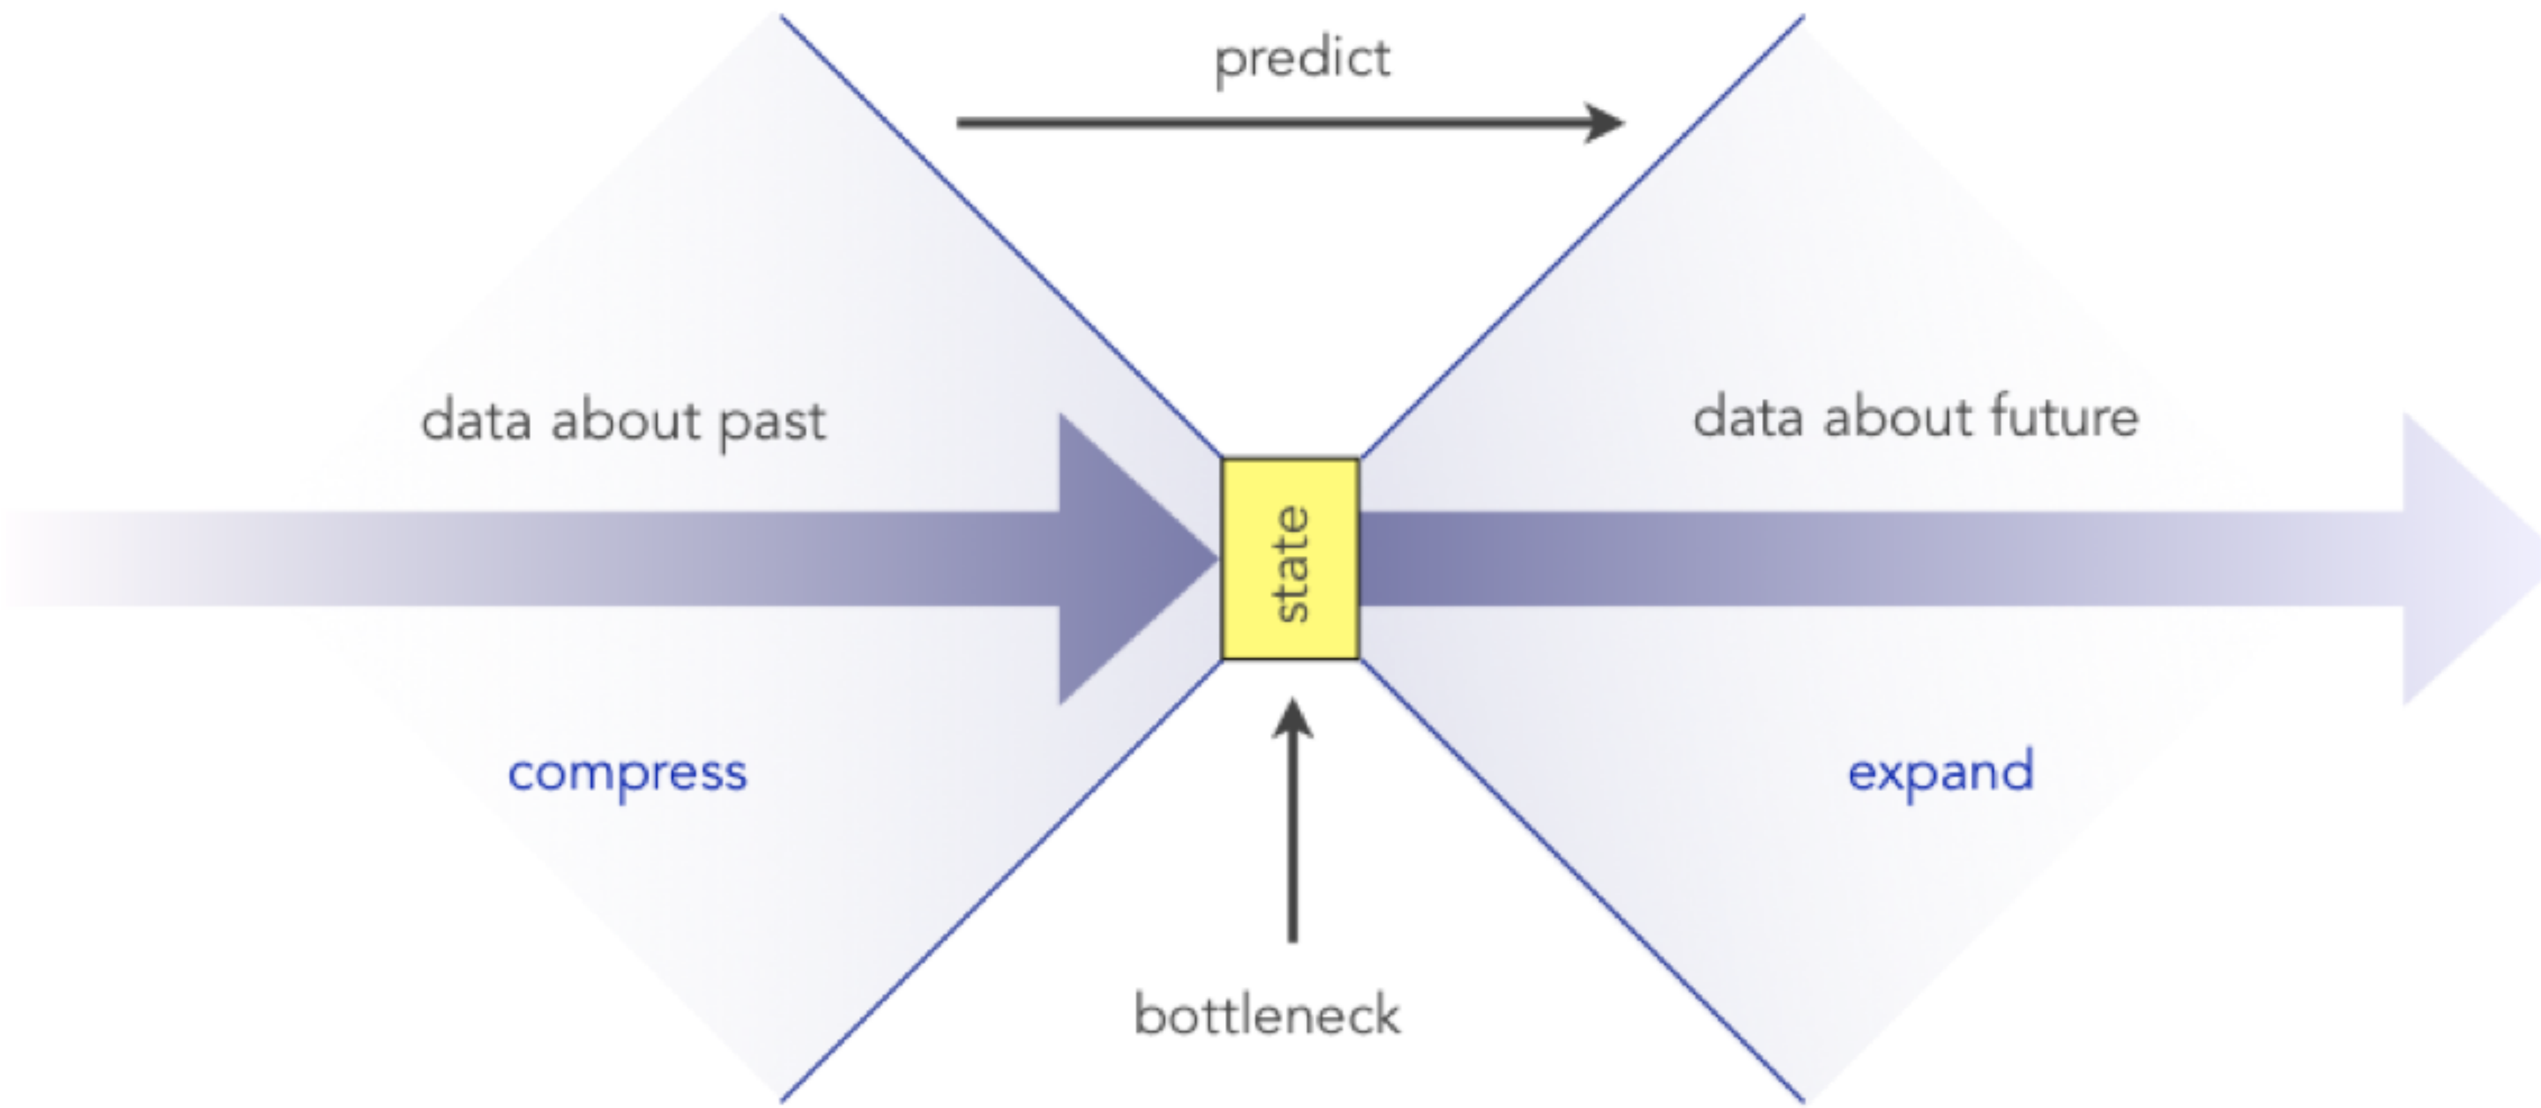

State: statistic of history sufficient to predict the future

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 the current state.

(Measurement does not always tell you state directly!!)

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 affect 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. Passive actions that change environment - weather (can detect with sensors)

3. NOP actions - doing nothing is also an action. State does not change.

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)

Let's think about causality of events

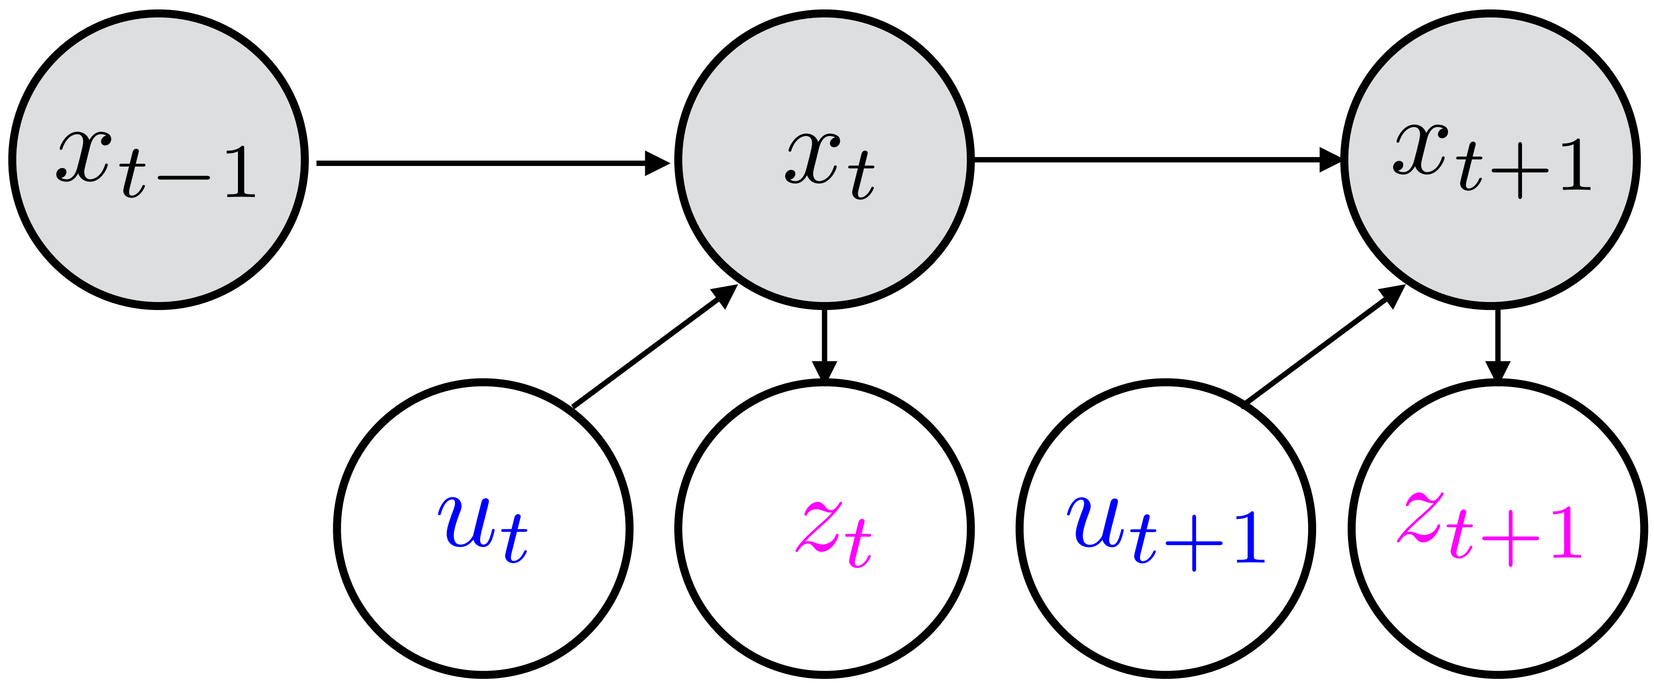

Assumptions:

1. Robot receives a stream of measurements / actions.

2. One measurement / action per time-step.

Problem: How do we estimate belief?

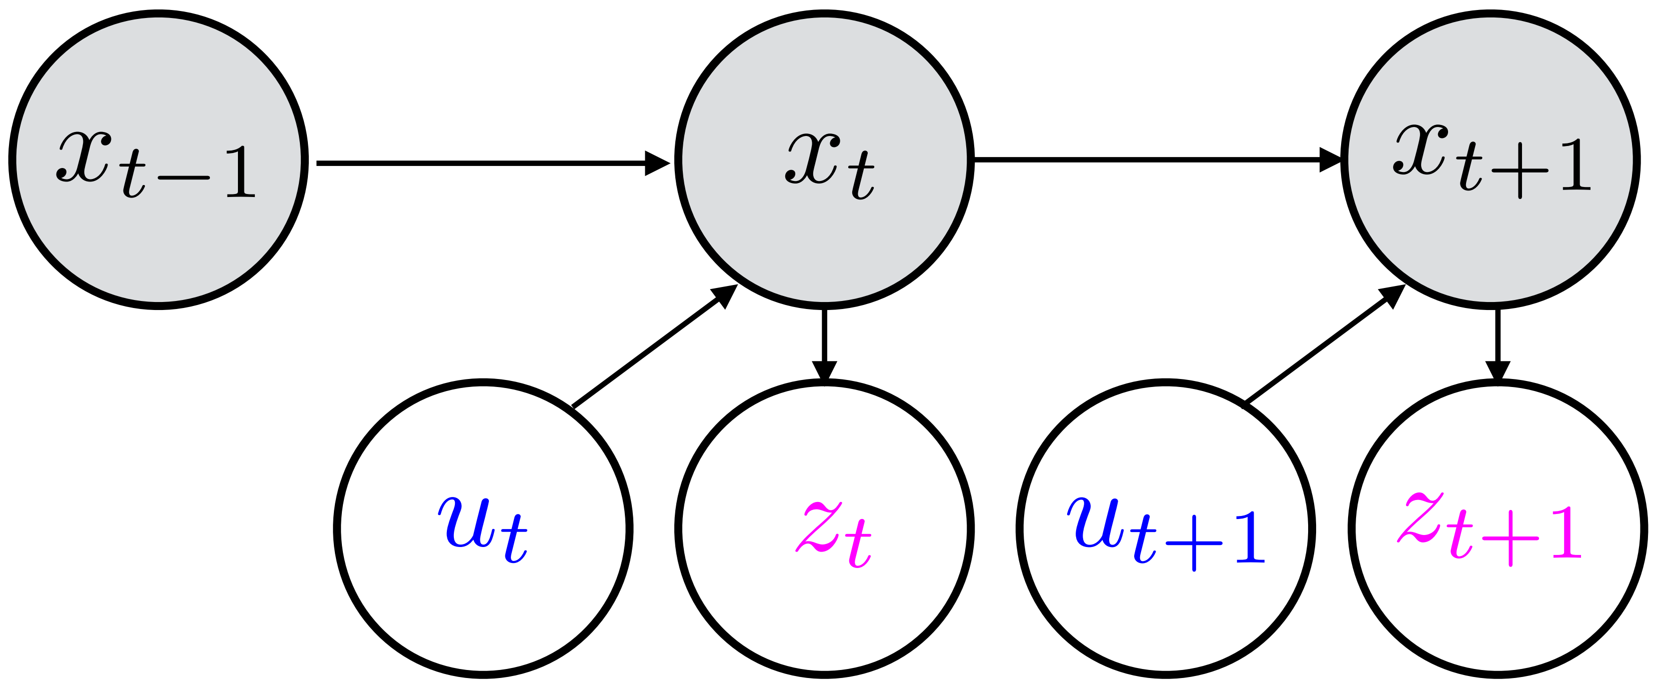

P(current state | all past information)

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

Problem: How do we estimate belief?

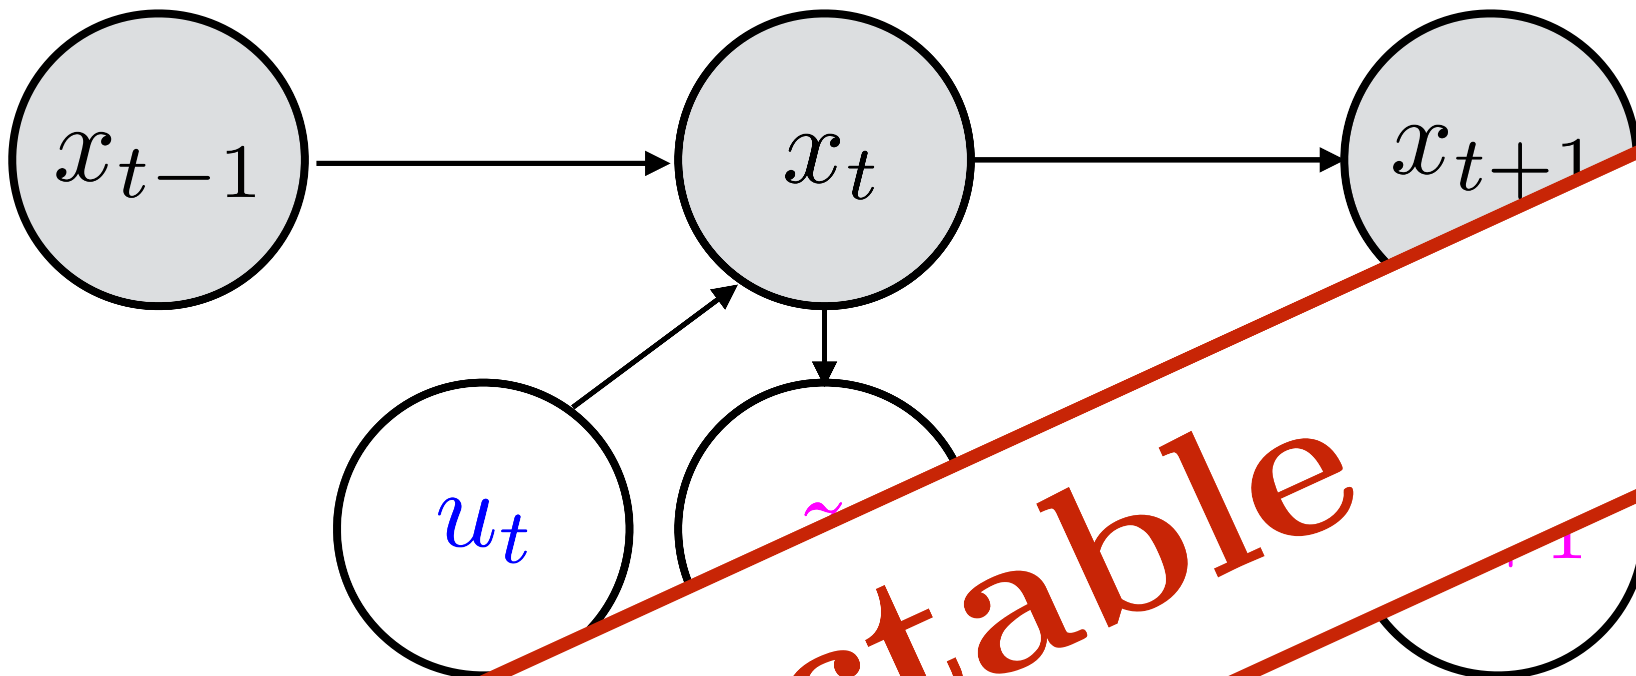

Intractable

ent state | all past information)

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

Solution

Andrey Andreyevich Markov (1856 - 1922)

Solution: Markov assumption

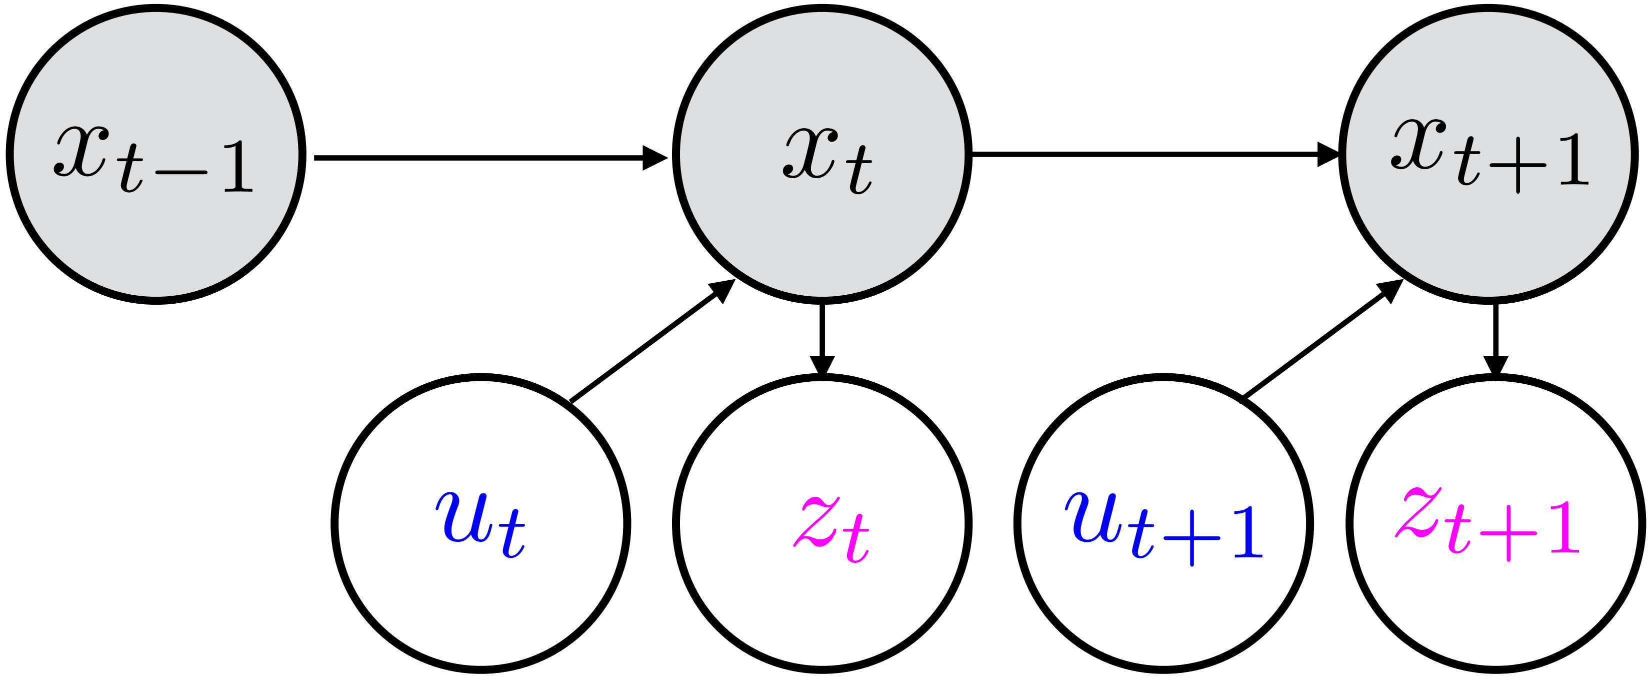

Markov assumption :

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

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

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

Reminder: Conditional Independence

$$P(A|B,C) = P(A|C)$$

$$\text{iff} \quad A, B \quad \text{conditionally independent given} \quad C$$

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

Central Question:
How do we tractably calculate belief?

Input data

Measurement $z_{1:t}$

Actions $u_{1:t}$

Belief

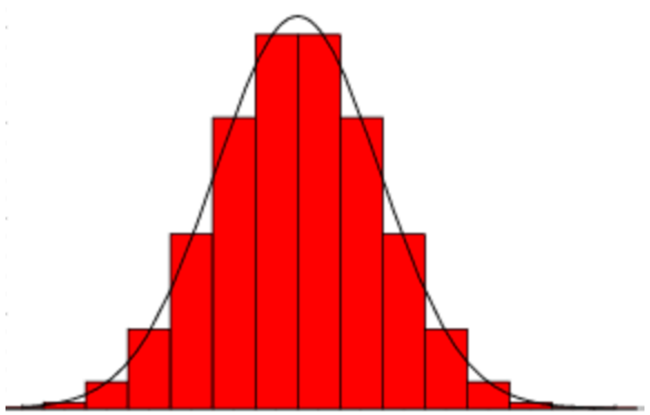

$$bel(x_t) = P(x_t | z_{1:t}, u_{1:t})$$

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

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

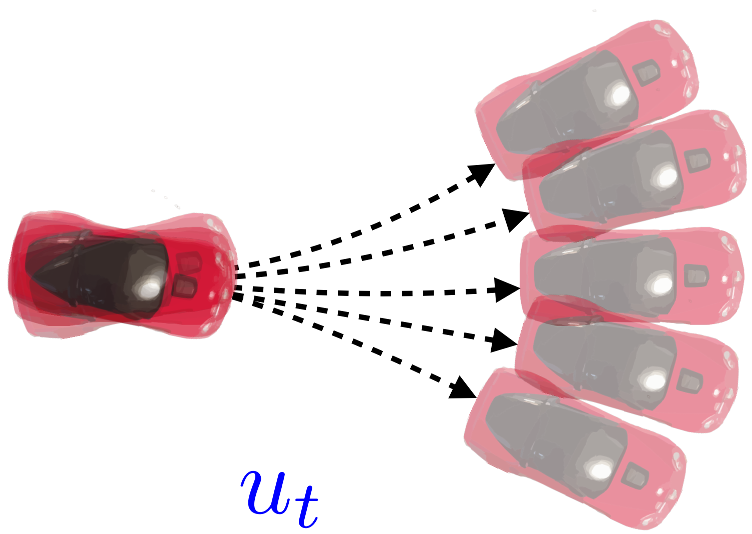

${\color{blue}u_t}$

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

(total probability)

Bayes filter in a nutshell

Step 2: Correction - apply Bayes rule given measurement

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

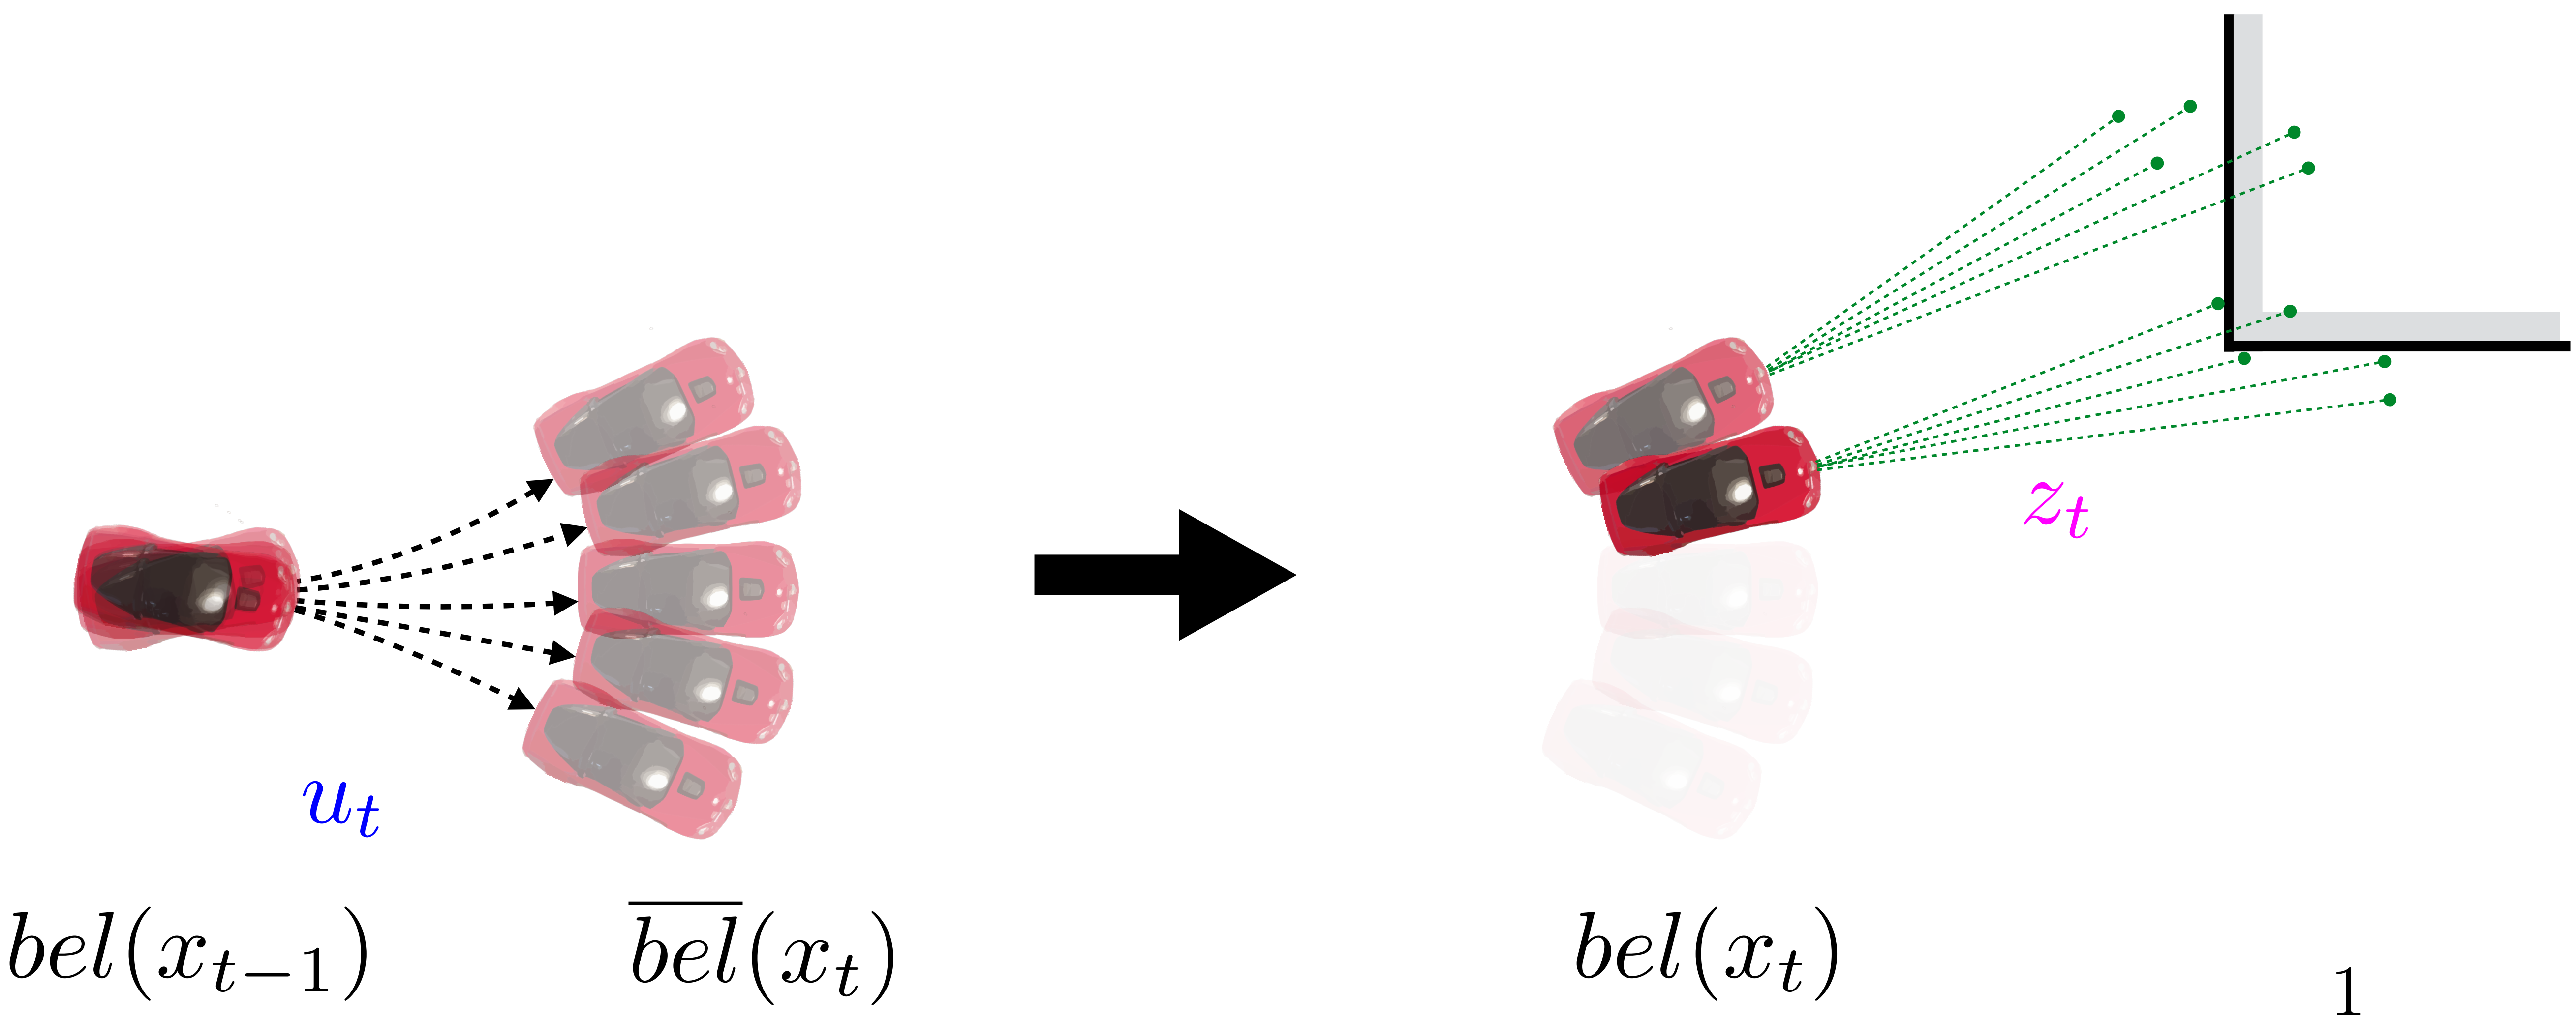

$$\eta = \frac{1}{\sum P(z_t|x_t)\overline{bel}(x_t)}$$

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