Robotics
Spring 2023

Abhishek Gupta

TAs: Yi Li, Srivatsa GS

Course Logistics

- Where: MEB 242
- When: 11:30-1 Tu/Thu
- Who:
  - Abhishek Gupta (Instructor)
  - Yi Li (TA)
  - Srivatsa GS (TA)
- Office hours:
  - Abhishek: Gates 215, Friday 4-5pm
  - Yi: Gates 152, Thursday 3-4pm
  - Srivatsa: Gates 152, Monday 3-4pm

Course Logistics - Grading

- Grading:
  - 50% of grade is on Final Project
  - 15% of grade is on HW1
  - 15% of grade is on HW2
  - 15% of grade is on HW3
  - 5% participation
- Communications through EdStem/e-mail
- Every week: 2 lectures by Abhishek
- Final projects will be presented in a poster session.
  - Intermediate project proposals and milestone check ins.
- Please participate, otherwise it will be boring for all of us!

Course Logistics - Project

- Final project (50% of grade):
  - Project proposal (1 page)
  - Milestone report (3-4 pages)
  - Final report (6-8 pages)
- Project can be investigating any question related to robotics:
  - New algorithm
  - Performant/stable implementation
  - Empirical investigation
  - New robotic application
  - …
- Can be done in groups of 1-2 students.

Course Logistics – Homeworks

- 3 homeworks covering 3 class modules:
  - Estimation: EKF/UKF/Particle filtering for localization
  - Control and Planning: RRT/RRT*/A*/D* for motion planning
  - End to end learning: Behavior cloning, Dagger, policy gradient, actor critic
- Homeworks are all in Python, using pybullet or pytorch

Course Logistics - Integrity

- Late policy

  You are allowed to use **6 late days throughout the quarter**. After this, assignments turned in late will incur a penalty of 20%, for each day. Please plan ahead.

- Academic Honesty Policy

  While we **encourage students to discuss homeworks, each student must write up their own solution.** It's fine to use a source for generic algorithms (with attribution), but it is not allowed to copy solutions to the problems. Additionally, **students may not post their code online**. If we determine that a student posted their code online, they will get an automatic 50% reduction on the entire assignment (math + code) and if they copy code for the problems from another student or from online, they will get an automatic 0% for the entire assignment (and possibly reported to the college).

Please don't cheat, make my life easier

Who am I?

- New assistant professor in CSE
- Grew up in Oregon/India, last 10 years in Berkeley
- Undergrad Berkeley, Ph.D. Berkeley, Postdoc MIT.
- Interests: RL/robotics/optimization and control/robustness and generalization
- Outside of work: Tennis/soccer/sketching/dog enthusiast

Who is Yi?

TA: Yi Li

- PhD student in RSElab
- Office hour: Thursday 3-4 pm
- Location: Gates 152
- Email: yili18@cs.washington.edu
- Research Experience:
  - Unseen Object Instance Segmentation and Tracking
  - Object 6D Pose Estimation
  - Instance Segmentation
  - Object Detection

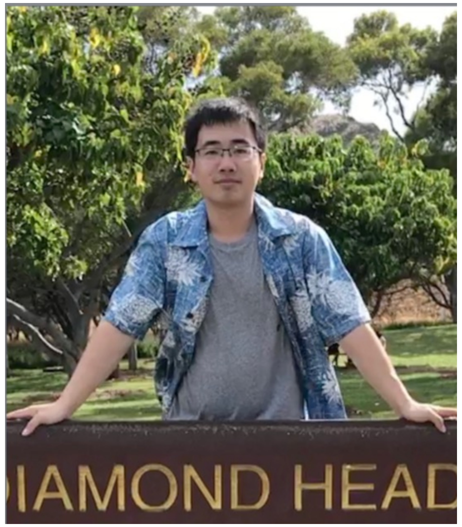

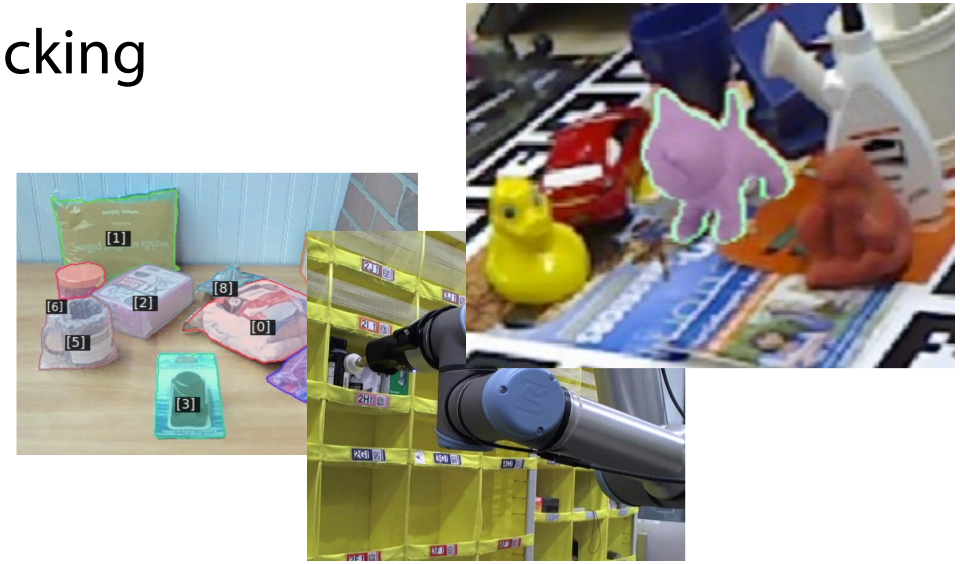

Who is Srivatsa?

- 2nd Year Masters Student in ME
  - Working on Visual Odometry and SLAM on the RACER project.
  - Previously worked on 6Dof Grasp generation and packing in UW-Amazon manipulation research and at Voaige Inc.
  - Office hours: Gates 152, Monday 3-4pm

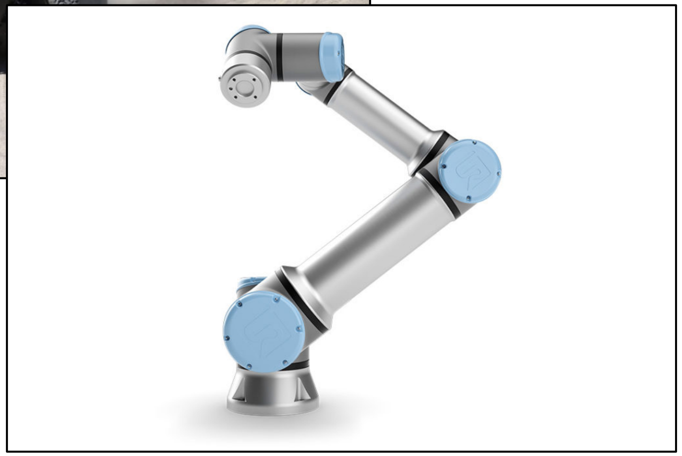

What is a robot?

- First definitions:
  - Karel Capek → robots were biological beings performing unpleasant labor.

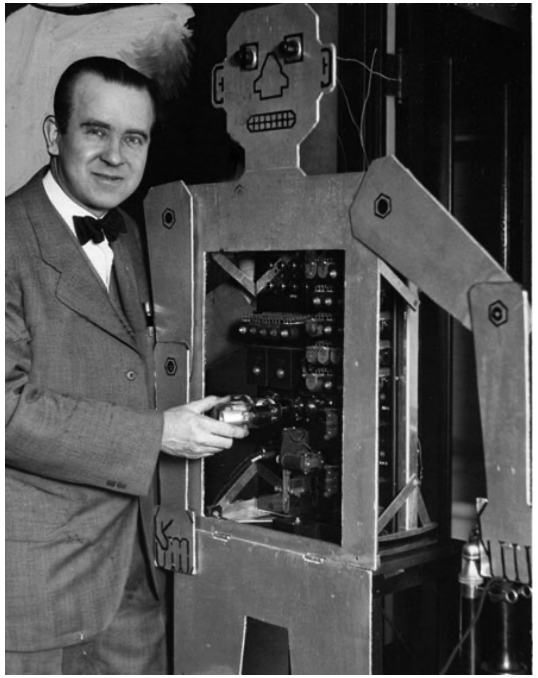

Herbert Televox(1927)

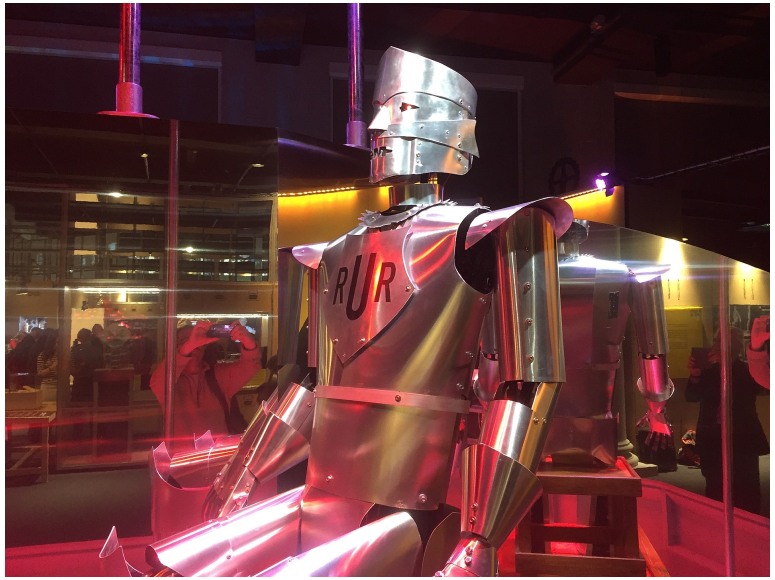

Eric (1928)

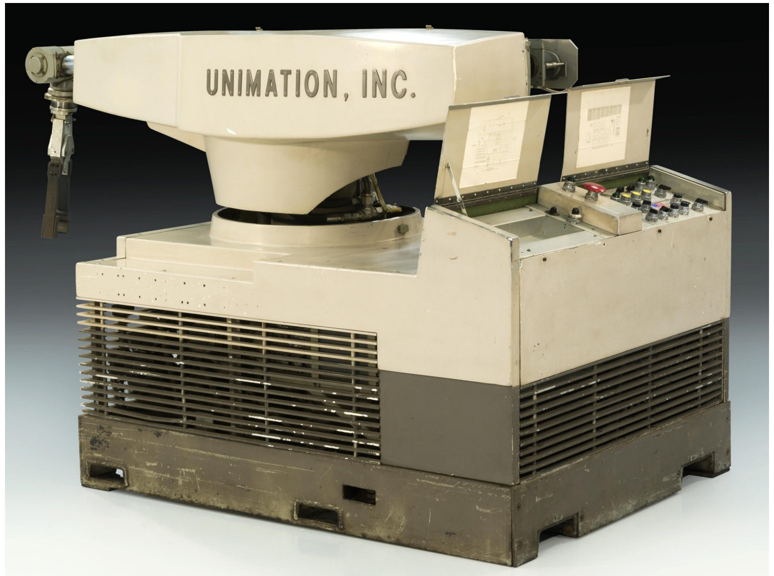

Unimate (1961)

The first wave of robots

Shakey

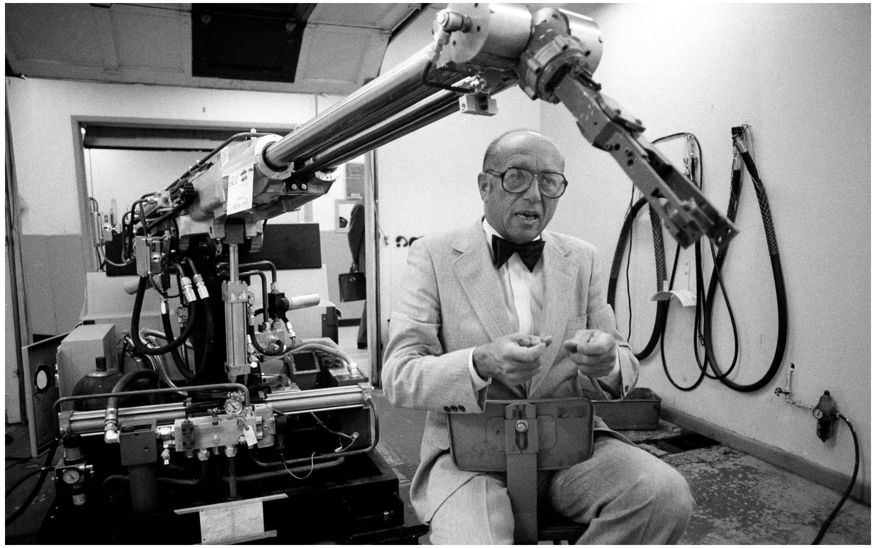

Engelberger
(Unimate ++)

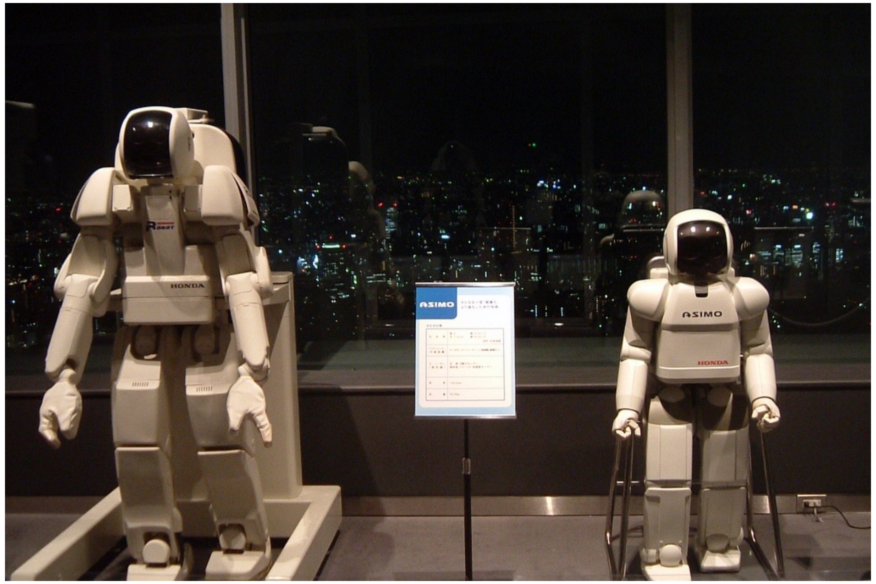

Honda P series

The second wave of robots

DARPA Grand Challenge

PR1 Robot

Robots Today

Everyday Robotics - Google

Atlas – Boston Dynamics

Robotics Spans Applications and Industries

- Applicable in a variety of industries and spaces:
  - Industry:
    - Industrial manufacturing
    - Warehouse navigation
  - Outdoor navigation/locomotion:
    - Legged locomotion
    - Outdoor navigation
    - Last mile delivery
    - Self driving cars
  - Home and office manipulation
    - Mobile manipulation
    - Dexterous manipulation

Industrial Robotics

Industrial Robotics Today

Industrial Pick and Place

Robots in Warehouses

Navigation

CMU NavLab

DARPA Urban Challenge 2007

Self-Driving Cars

Outdoor Off-Road Autonomy

Locomotion

Boston Dynamics BigDog (2008)

RoboCup

Boston Dynamics Spot

Humanoid Parkour

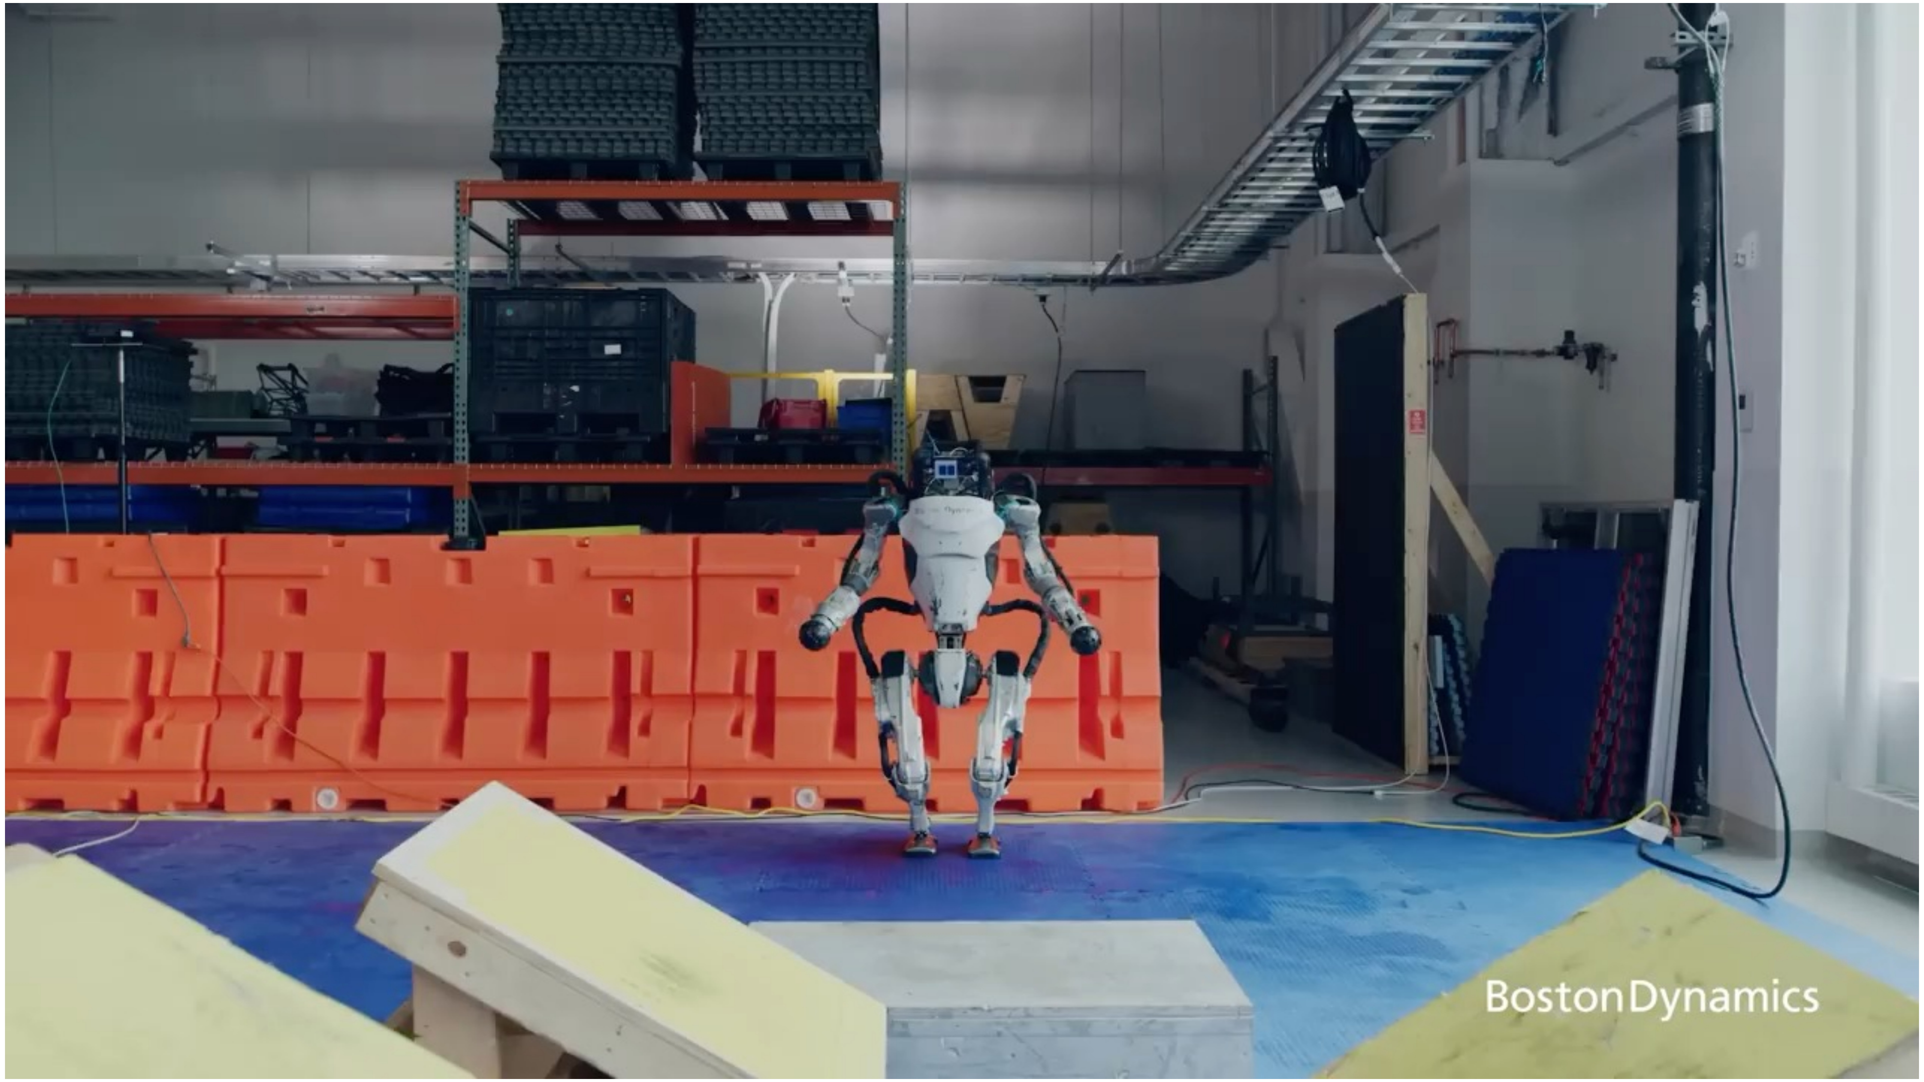

Drone Delivery

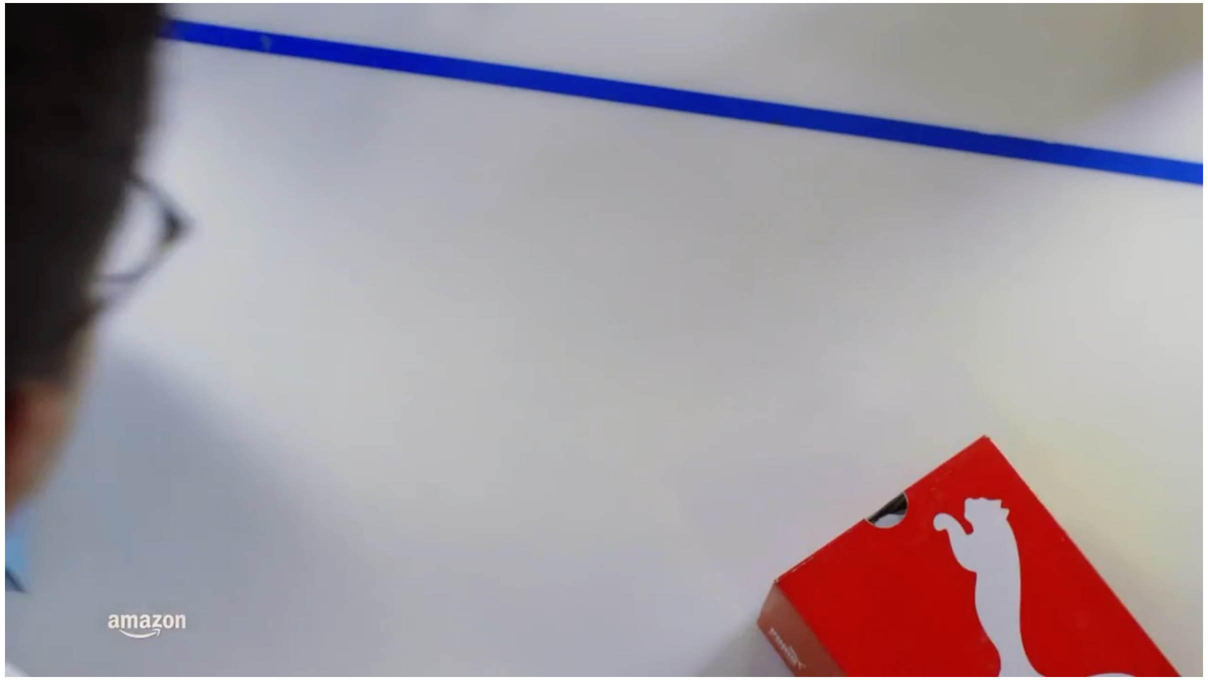

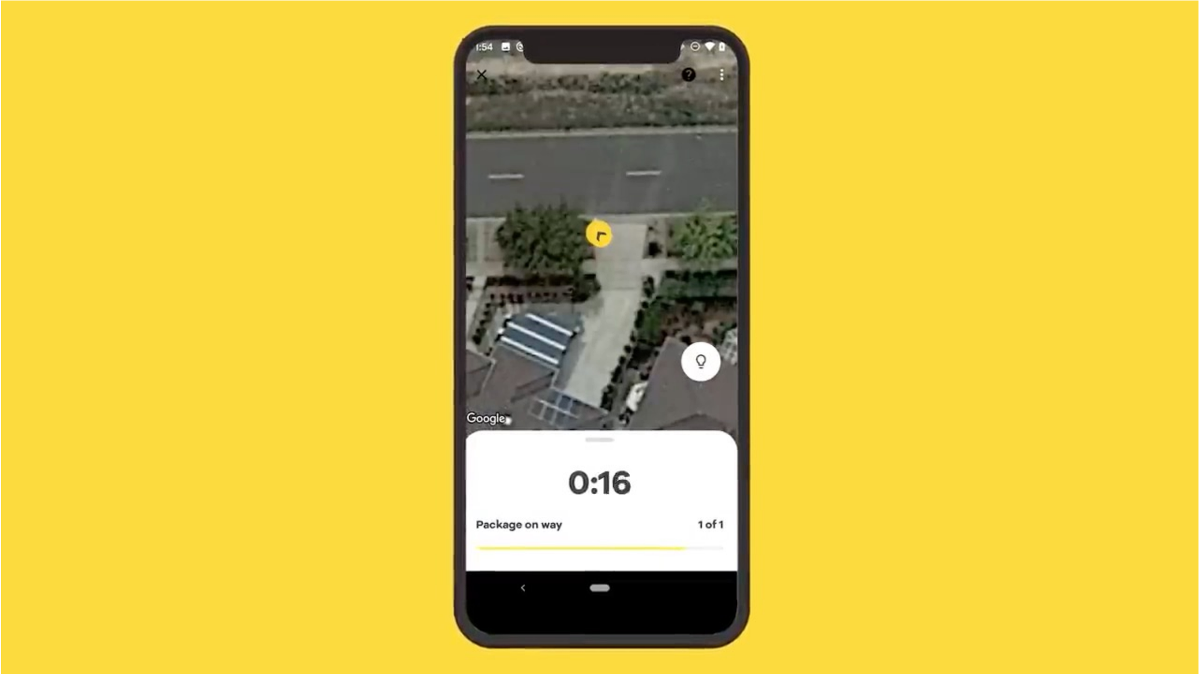

Indoor Manipulation

Dexterous Manipulation

Mobile Manipulation

HaptX Dataglove

How should we start to formalize the robotics problem?

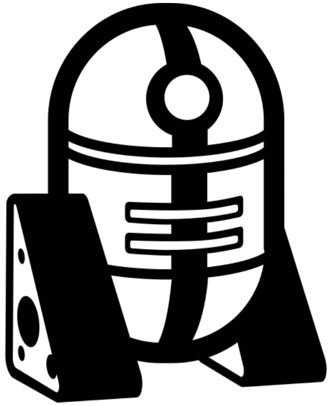

- Agent: Rational entity equipped with sensors and actuators
- Environment: accepts actuation commands and steps forward according to some dynamics

Graphical Model of Robotics

- State: Minimum sufficient statistic encapsulating the world, sufficient for prediction
- Measurement/Observation: Current sensor readings, potentially partially observed
- Action: Actuators that agent can use to affect the state

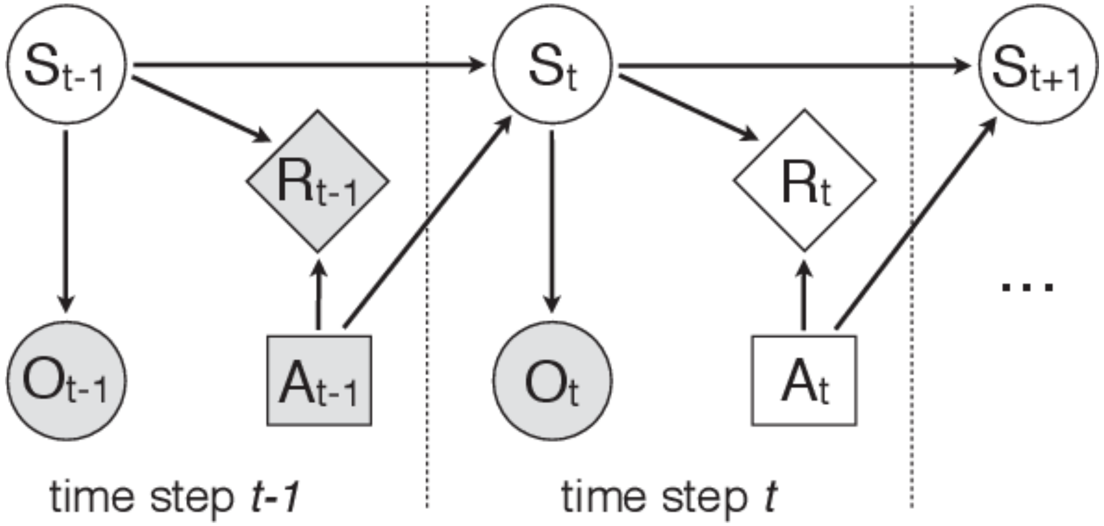

Graphical Model of Robotics

- State: sufficient statistic encapsulating the world, sufficient for prediction (this is a choice)
- Measurement/Observation: Current sensor readings, potentially partially observed (this is a choice)
- Action: Actuators that agent can use to affect the state (this is a choice)

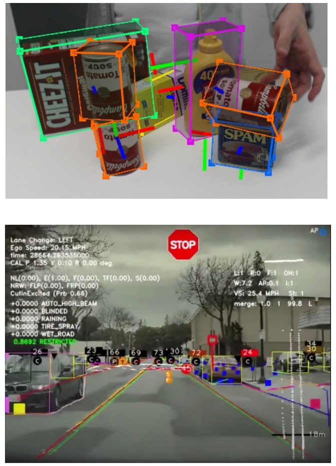

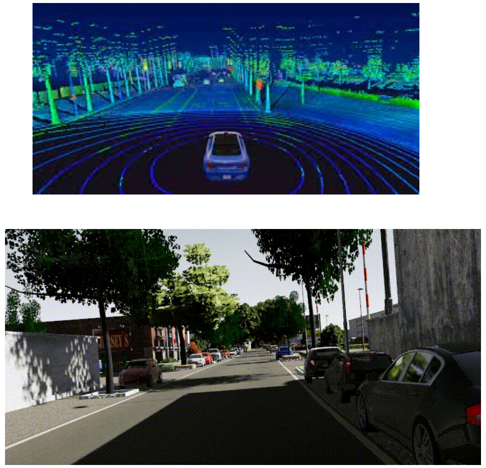

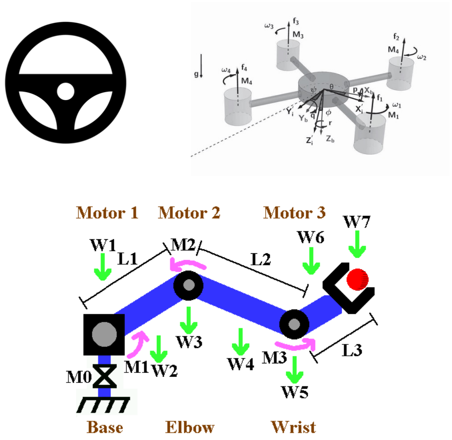

Sense-Plan-Act Framework

Robotics has three primary subpieces:

1. Sensing → from measurements
2. Planning → from models
3. Acting → in the world

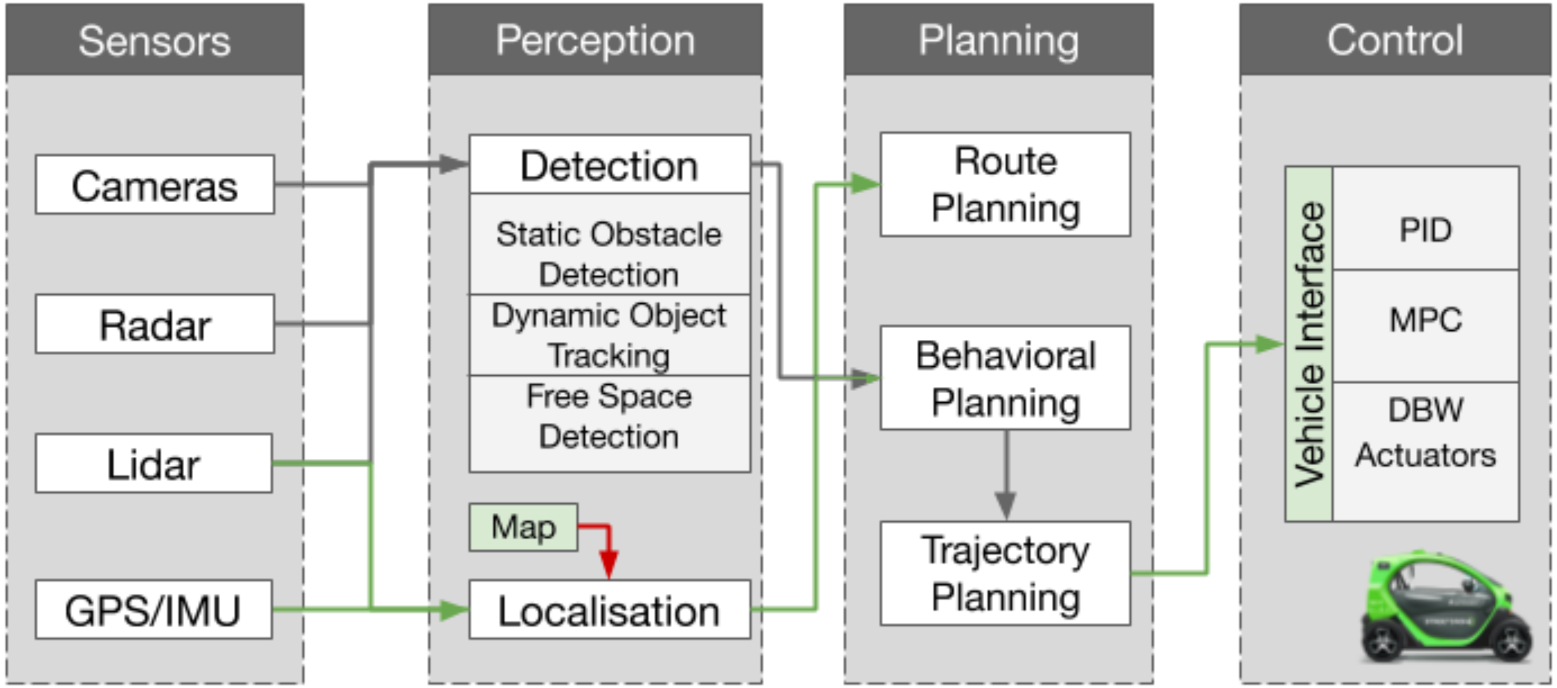

Sensing: Why is it nontrivial?

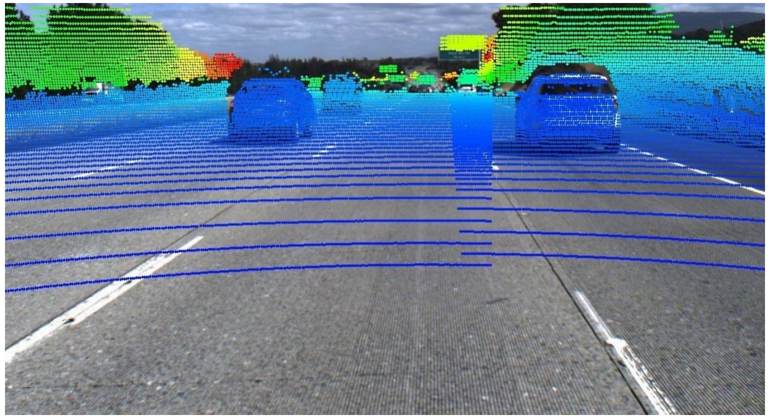

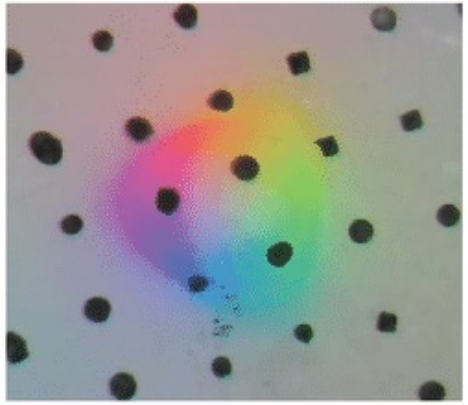

- Sensors have overwhelming amounts of information
- Partially observed
- Noisy and prone to drift

Sensing: Why is it nontrivial?

State is never perfectly known, only a **belief** over state can be estimated

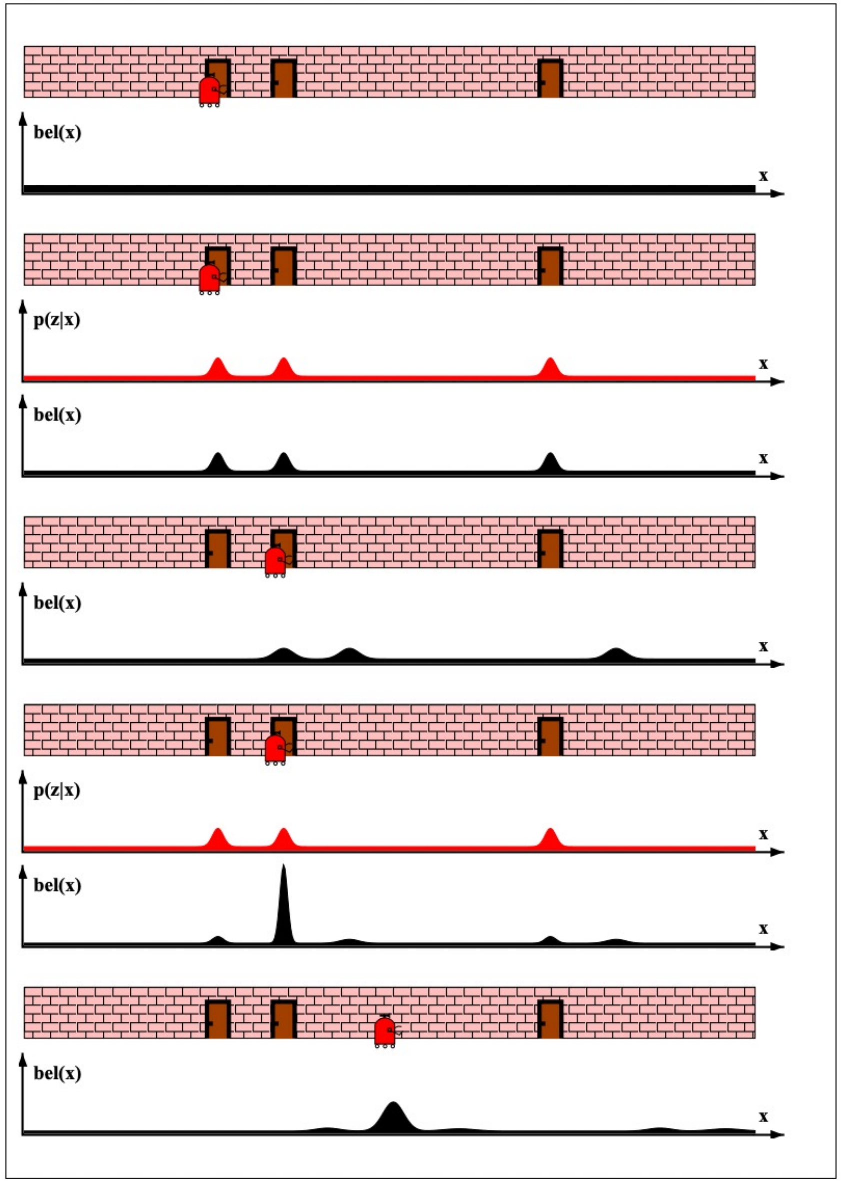

Probabilistic Robotics

A robot that carries a notion of its own uncertainty and that acts accordingly is superior to one that does not.

- Maintaining uncertainty allows for
  - Information gathering
  - Robustness to imperfect sensing and actuation
  - Interpretability
  - Exploration

Planning: Why is it nontrivial?

- Searching/Optimization through a complex non-convex space
- Combination of discrete/continuous optimization

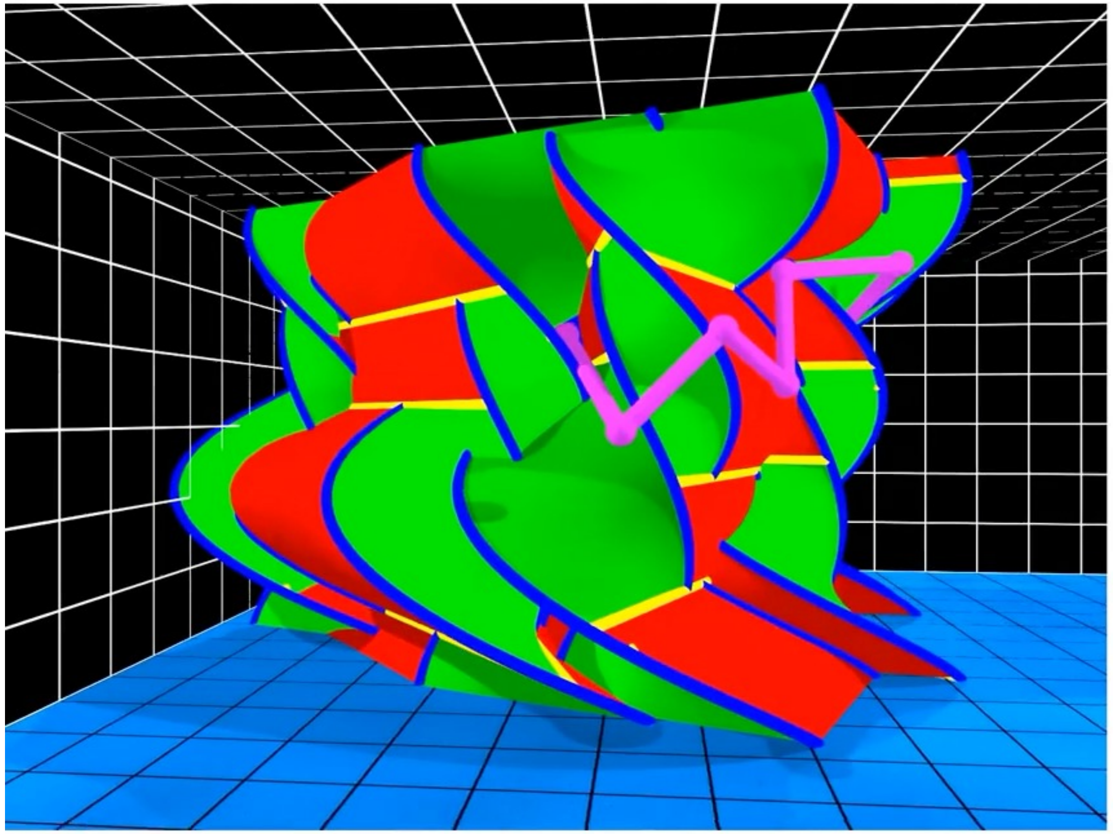

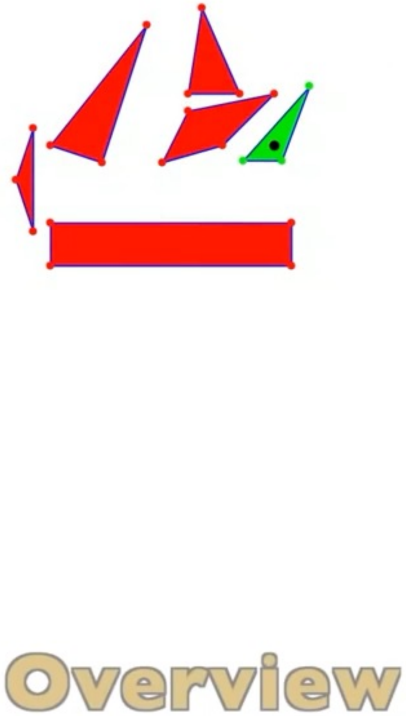

Rational Agents and Utility Maximization

- How do we even formulate planning?
  - Utility maximization: trajectory optimization
  - Search: shortest path finding
- Viewing planning through the lens of rationality allows us to use tools from optimization

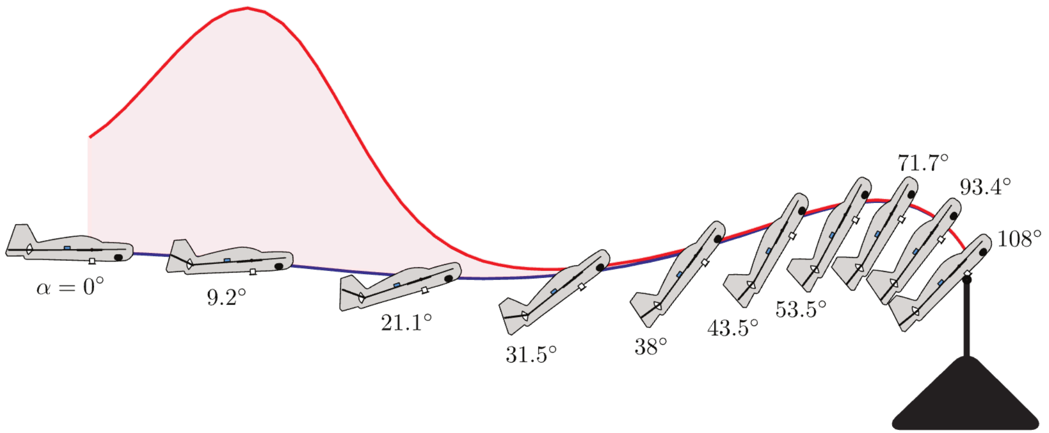

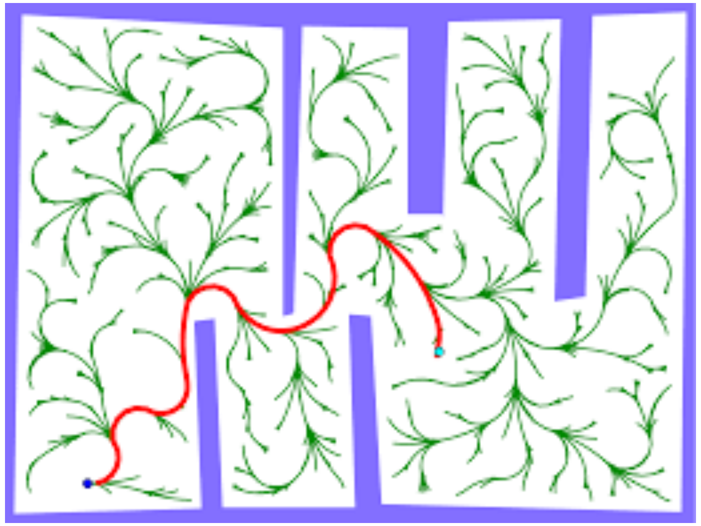

Acting: Why is it nontrivial?

- Robot systems in the real world are subject to significant perturbations/noise → need to be stable in the face of these perturbations

Low-level control and Stabilization

- Unstable/suboptimal systems can be catastrophic

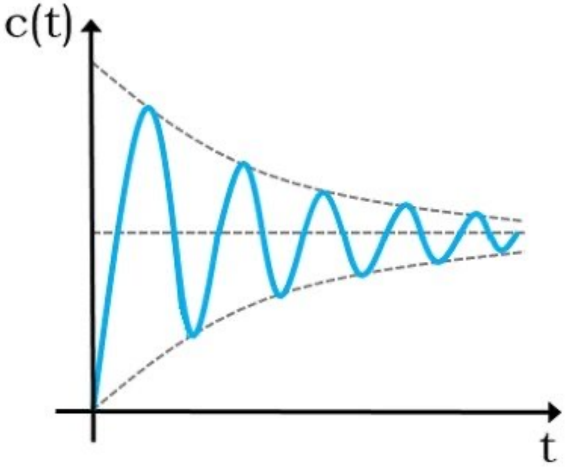

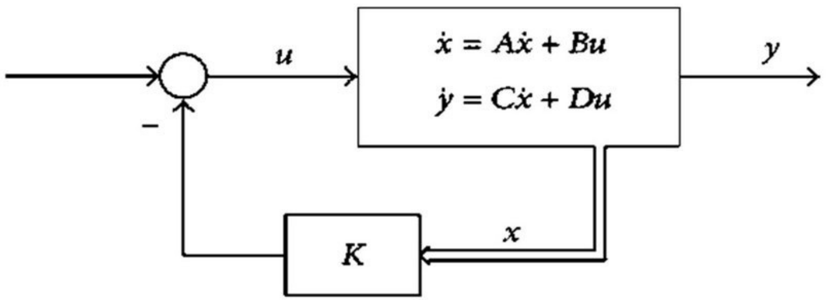

Robotics: Integrated System Research

Focus on addressing all problems at once

$$\mathbf{M}(\mathbf{q})\ddot{\mathbf{q}} + \mathbf{C}(\mathbf{q}, \dot{\mathbf{q}})\dot{\mathbf{q}} = \tau_g(\mathbf{q}) + \mathbf{Bu},$$

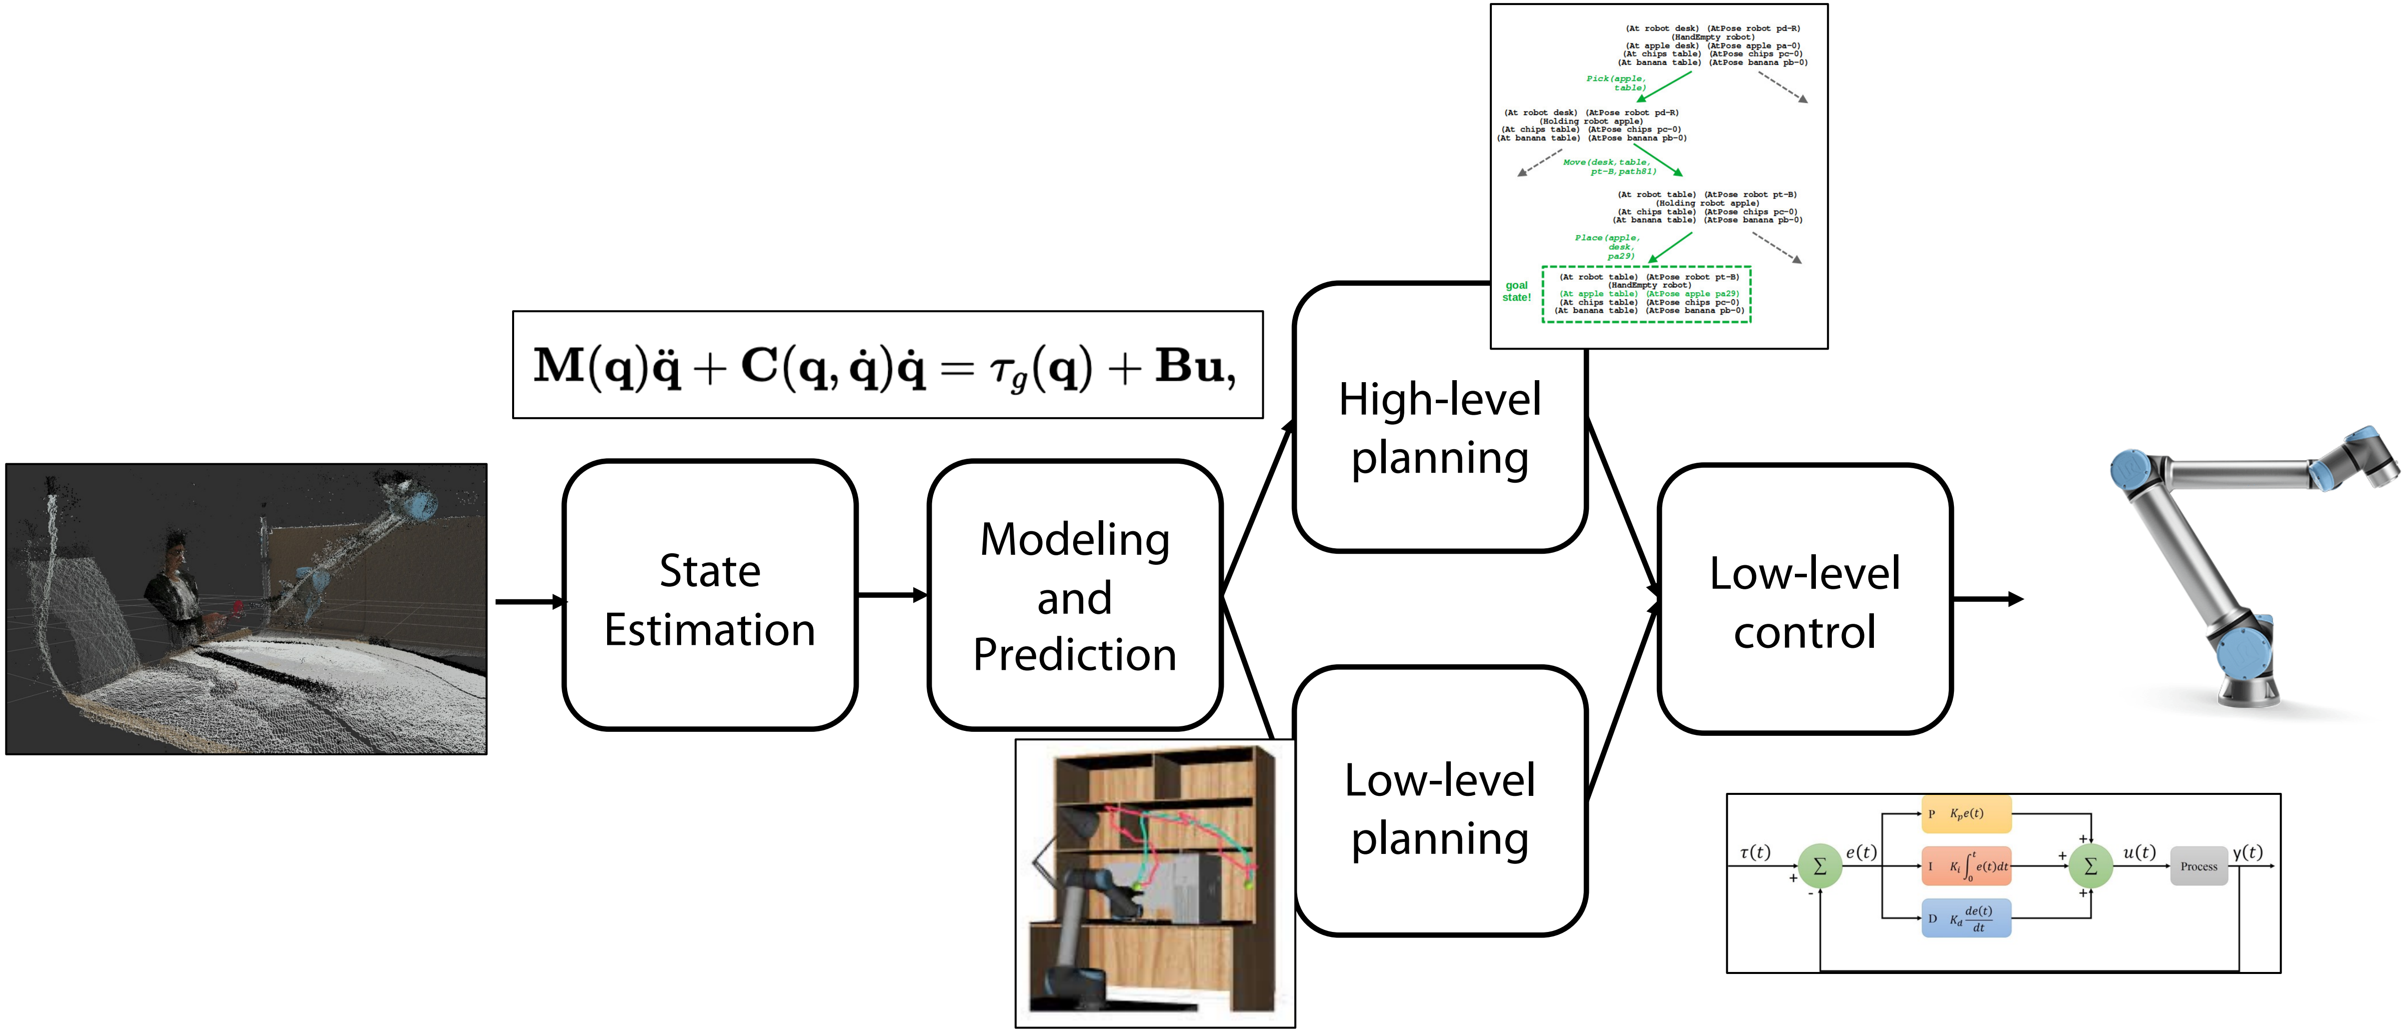

What makes robotics hard?

Unique interplay of expectations and assumptions

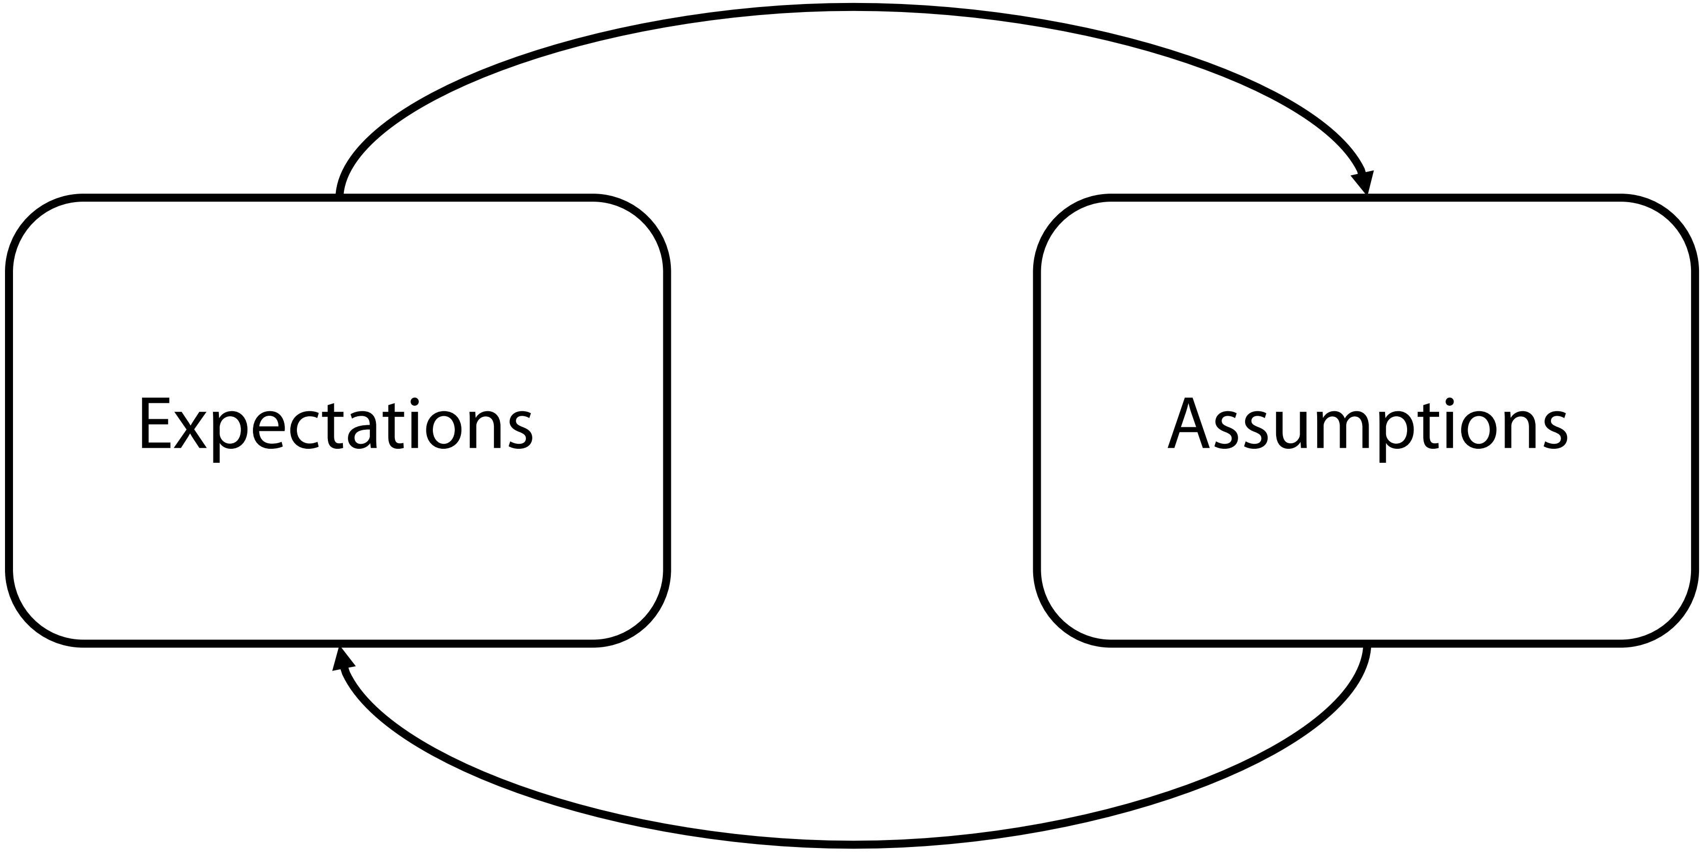

Expectations

- Significantly larger set of tasks being performed
- Failure can be catastrophic and unsafe
- Precision required may be much larger than other decision making problems
- Multi-step decision making problems

Assumptions

- Limited data due to the real world
- No perfect simulator
- Chicken and egg problem on deployment
- Evaluation is difficult

Why now?

Hardware is getting cheaper and more accessible

400K

30K

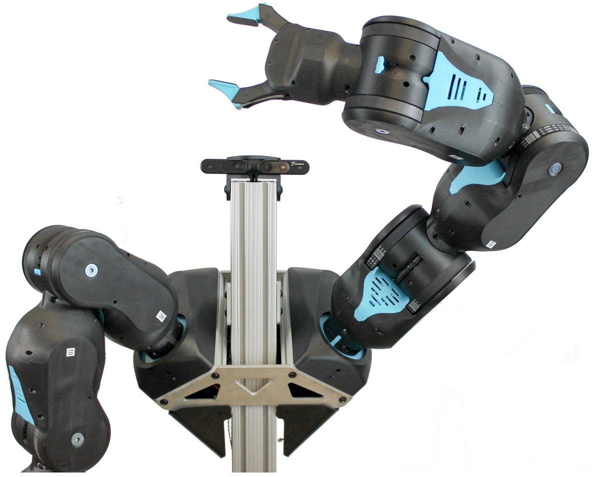

5K

Why now?

Algorithms/models have started maturing/stabilizing

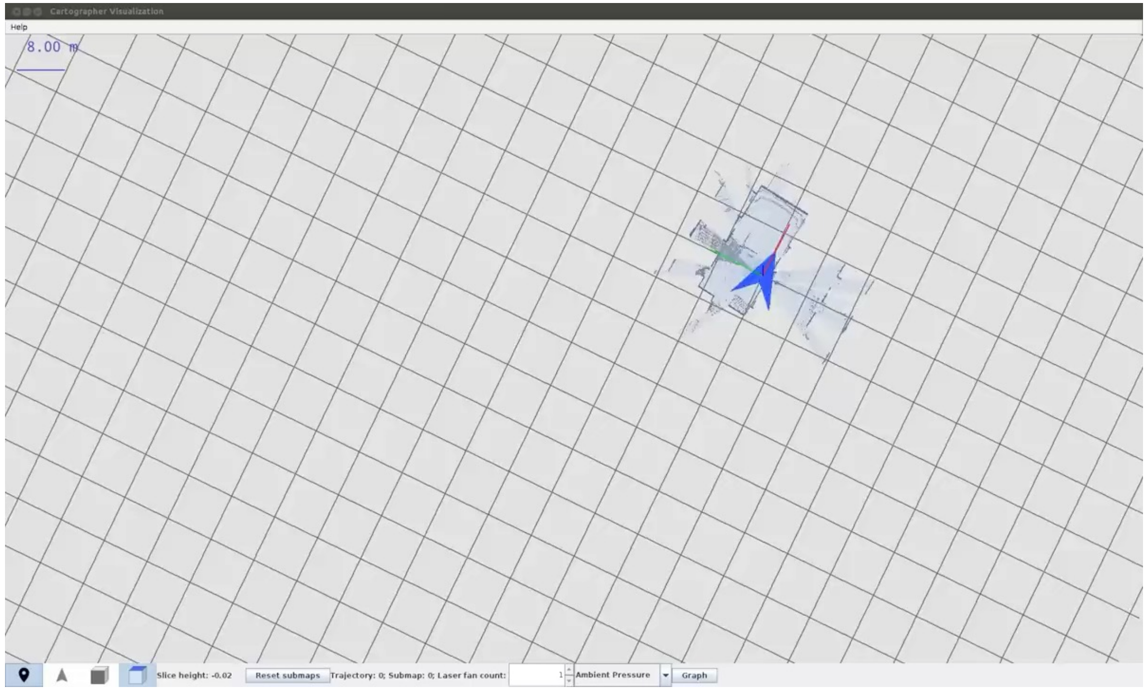

Why now?

Data/compute is now available at a scale not possible before

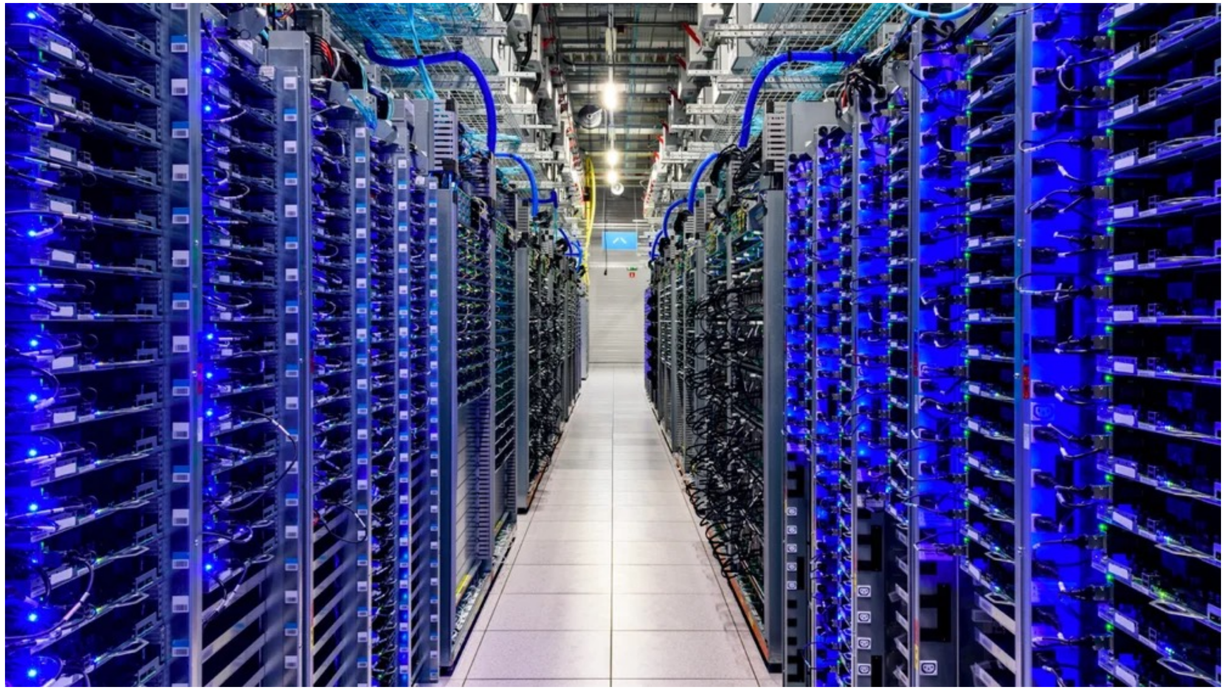

Why now?

Adjacent fields are showing remarkable progress

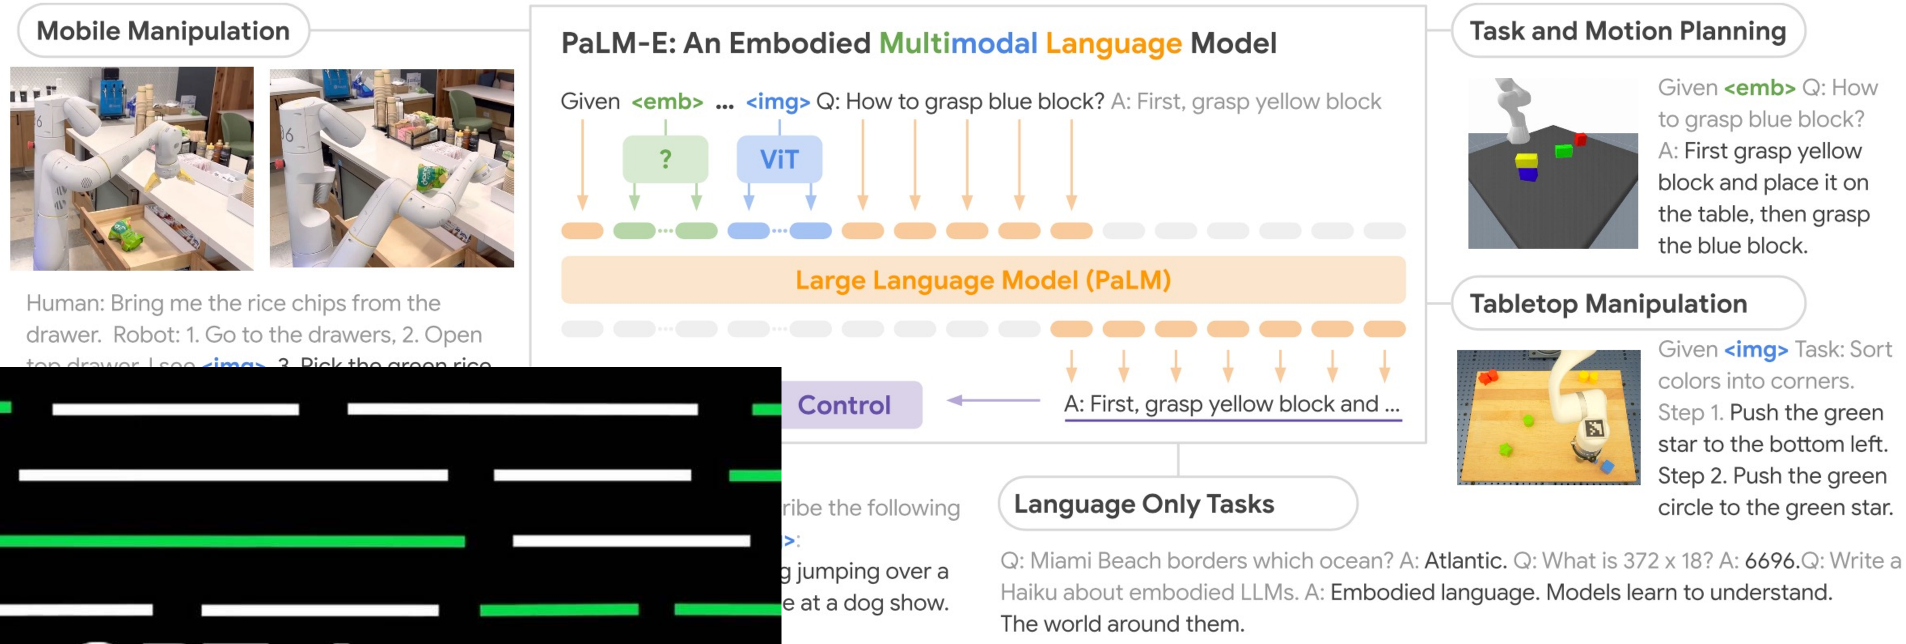

Why now?

- Fast growing investment into automation/robotics

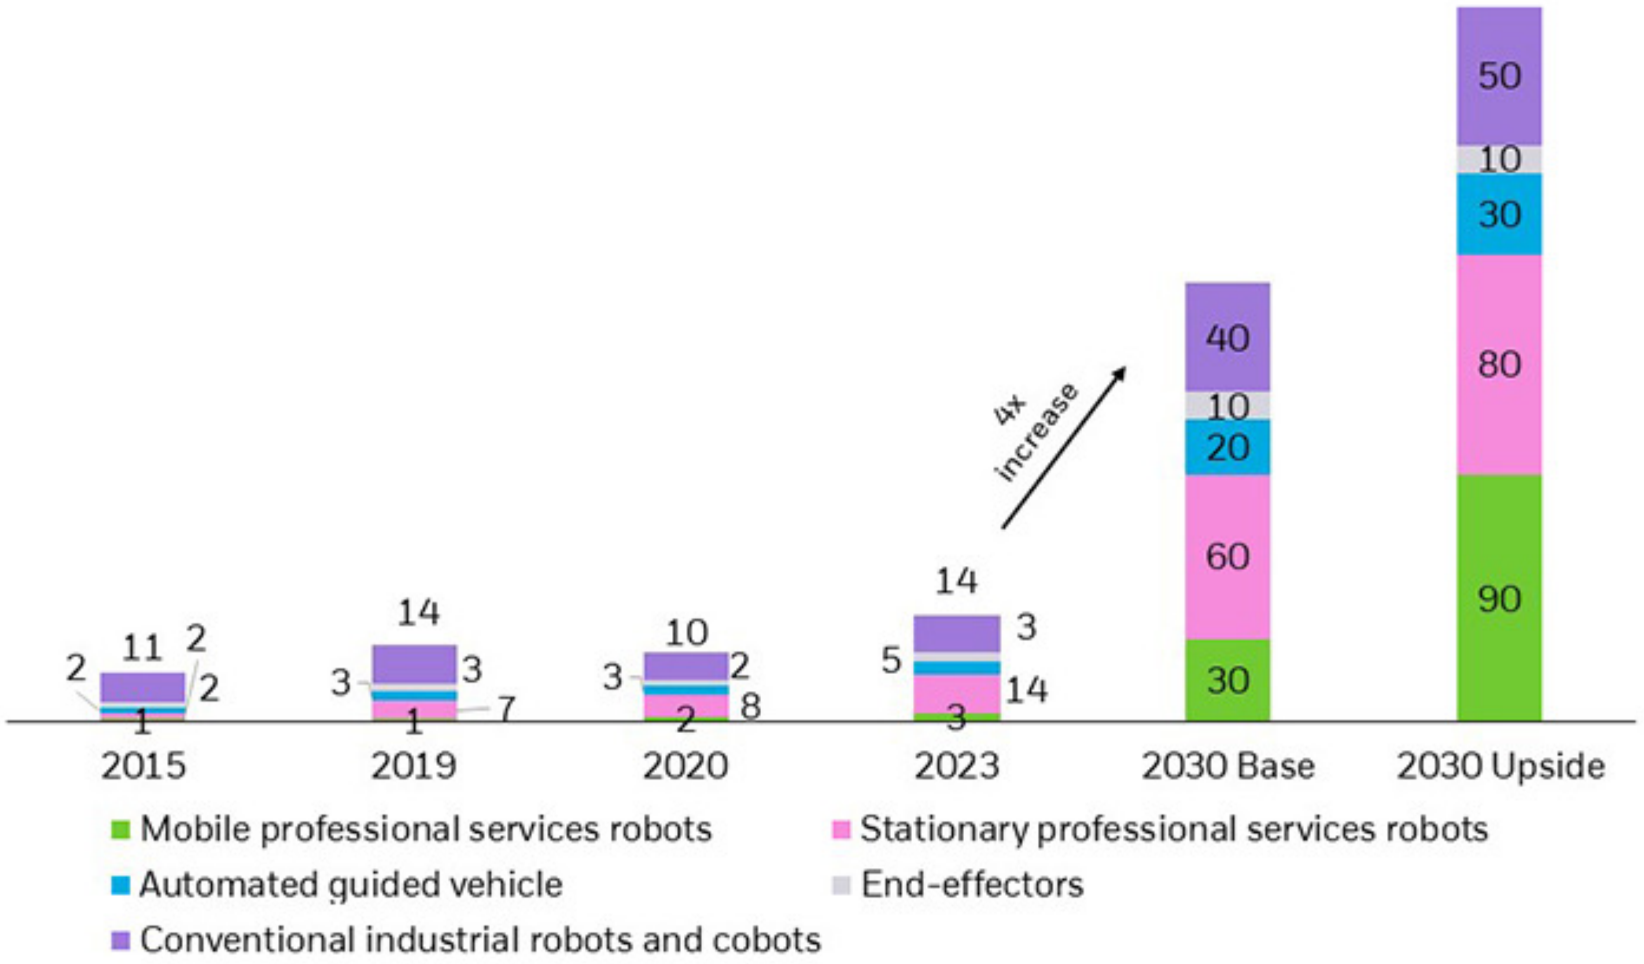

Is robotics useful to study beyond applications?

Arguably intelligence needs embodiment

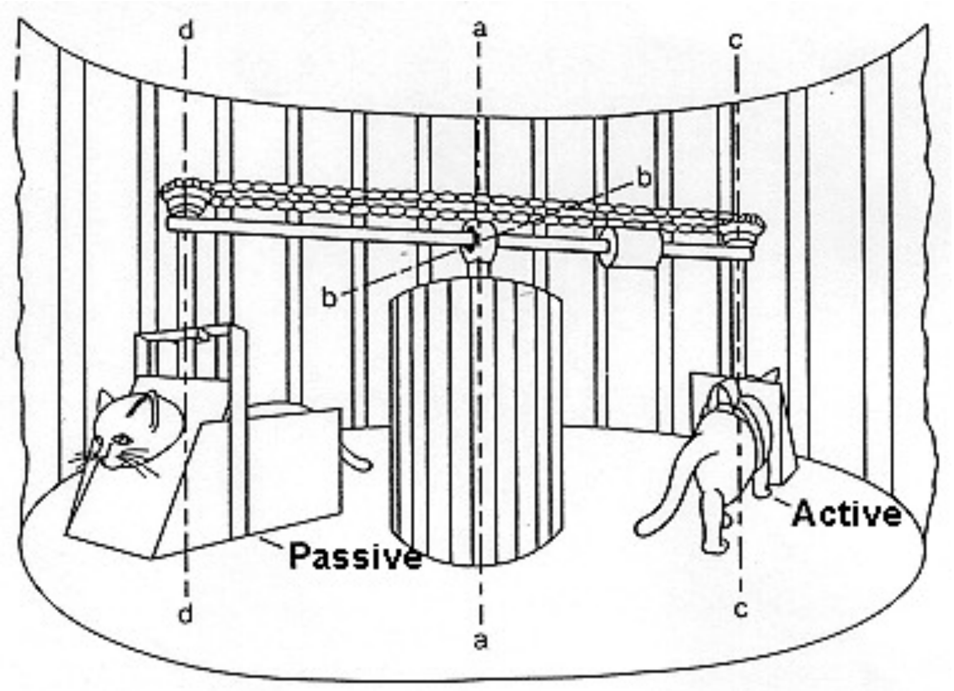

- visually-guided paw placement:
- avoidance of a visual cliff
- blink to an approaching object:
- visual pursuit of a moving object:
- pupillary reflex to light:
- tactual placing response

Table 1 Responses on the visual cliff

Active kittens		**Passive kittens**	
shallow	deep	shallow	deep
12	0	6	6
12	0	4	8
12	0	7	5
12	0	6	6
12	0	7	5
12	0	7	5
12	0	5	7
12	0	8	4
12	0	6	6
12	0	8	4

Robotics may be a way to study fundamental intelligence

Zooming out – why this matters for the study of intelligence?

Hypothesis: Intelligence with and without embodiment looks drastically different

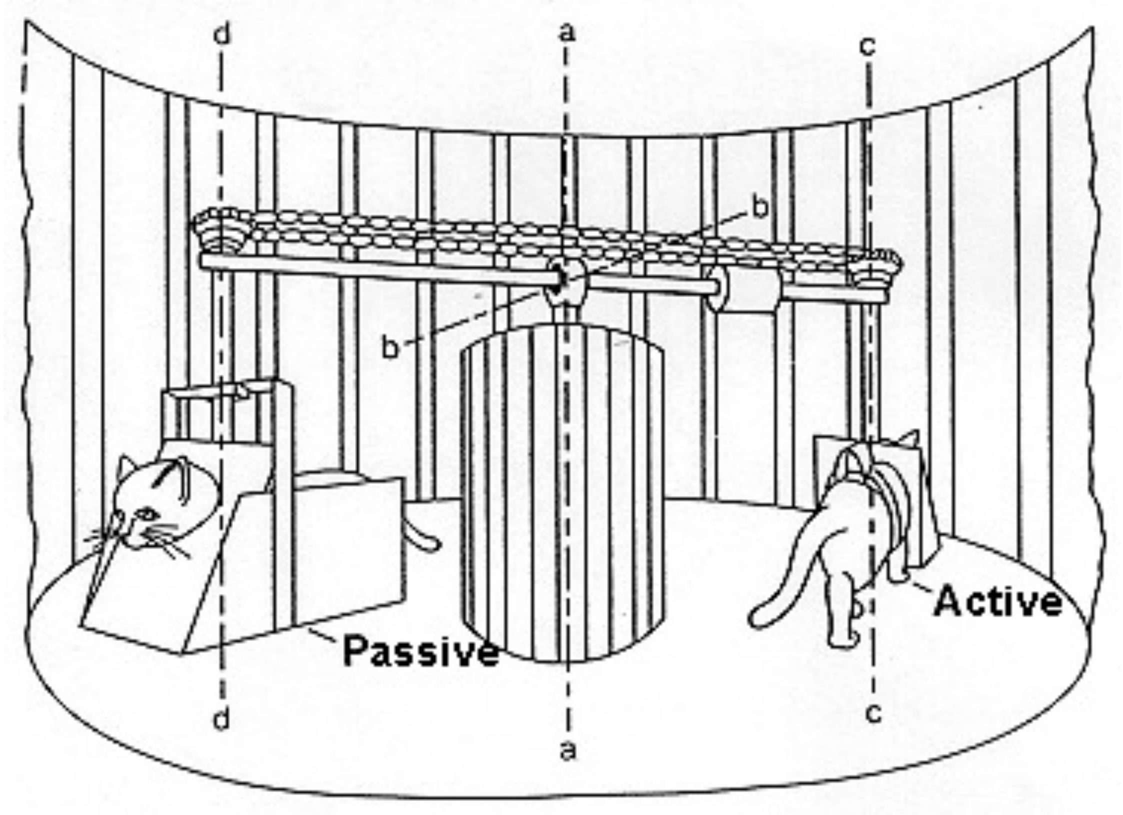

Elephants don't play chess!

Separation Principle

Under certain assumptions, the optimal state estimation + optimal deterministic control yields an optimal system

$$\dot{x} = Ax + Bu$$
$$y = Cx$$

$$u = -K\hat{x}$$
$$u = -Kx$$

Not always true!

$$\dot{\hat{x}} = A\hat{x} + Bu + L(y - C\hat{x})$$
$$\dot{e} = (A - LC)e$$

Control as if you had perfect state

Estimate as if you didn't care about control

Why do we care?

Module 1: Estimation

$$\dot{\hat{x}} = A\hat{x} + Bu + L(y - C\hat{x})$$
$$\dot{e} = (A - LC)e$$

Filtering/Smoothing

Localization

Mapping

SLAM

Why do we care?

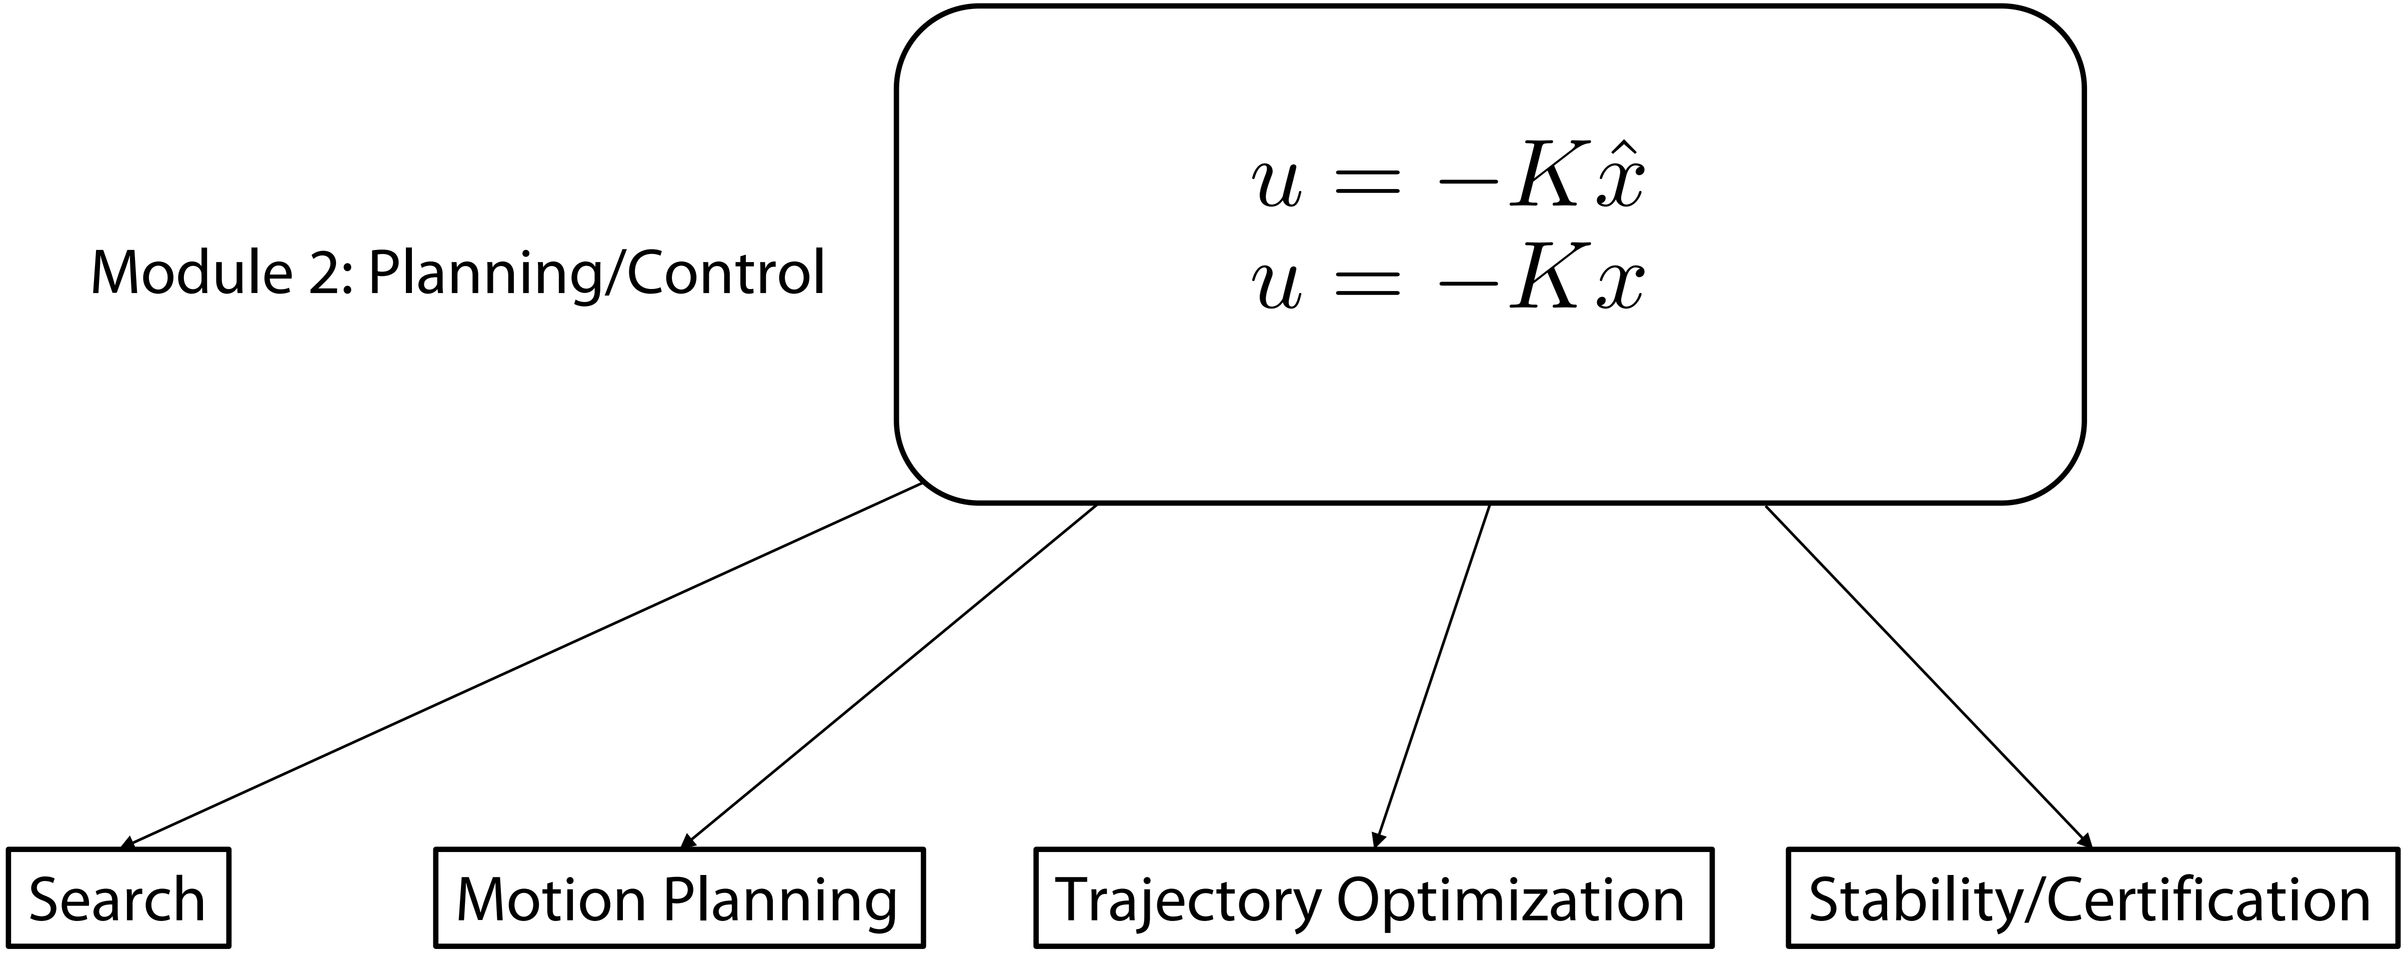

When does this not hold?

$$u = -K\hat{x}$$
$$u = -Kx$$

$$\dot{\hat{x}} = A\hat{x} + Bu + L(y - C\hat{x})$$
$$\dot{e} = (A - LC)e$$

Not always true!

Imperfect and arbitrarily non-linear systems

MDPs and Reinforcement Learning

Imitation Learning

Solving POMDPs

How does a typical robotics pipeline look?

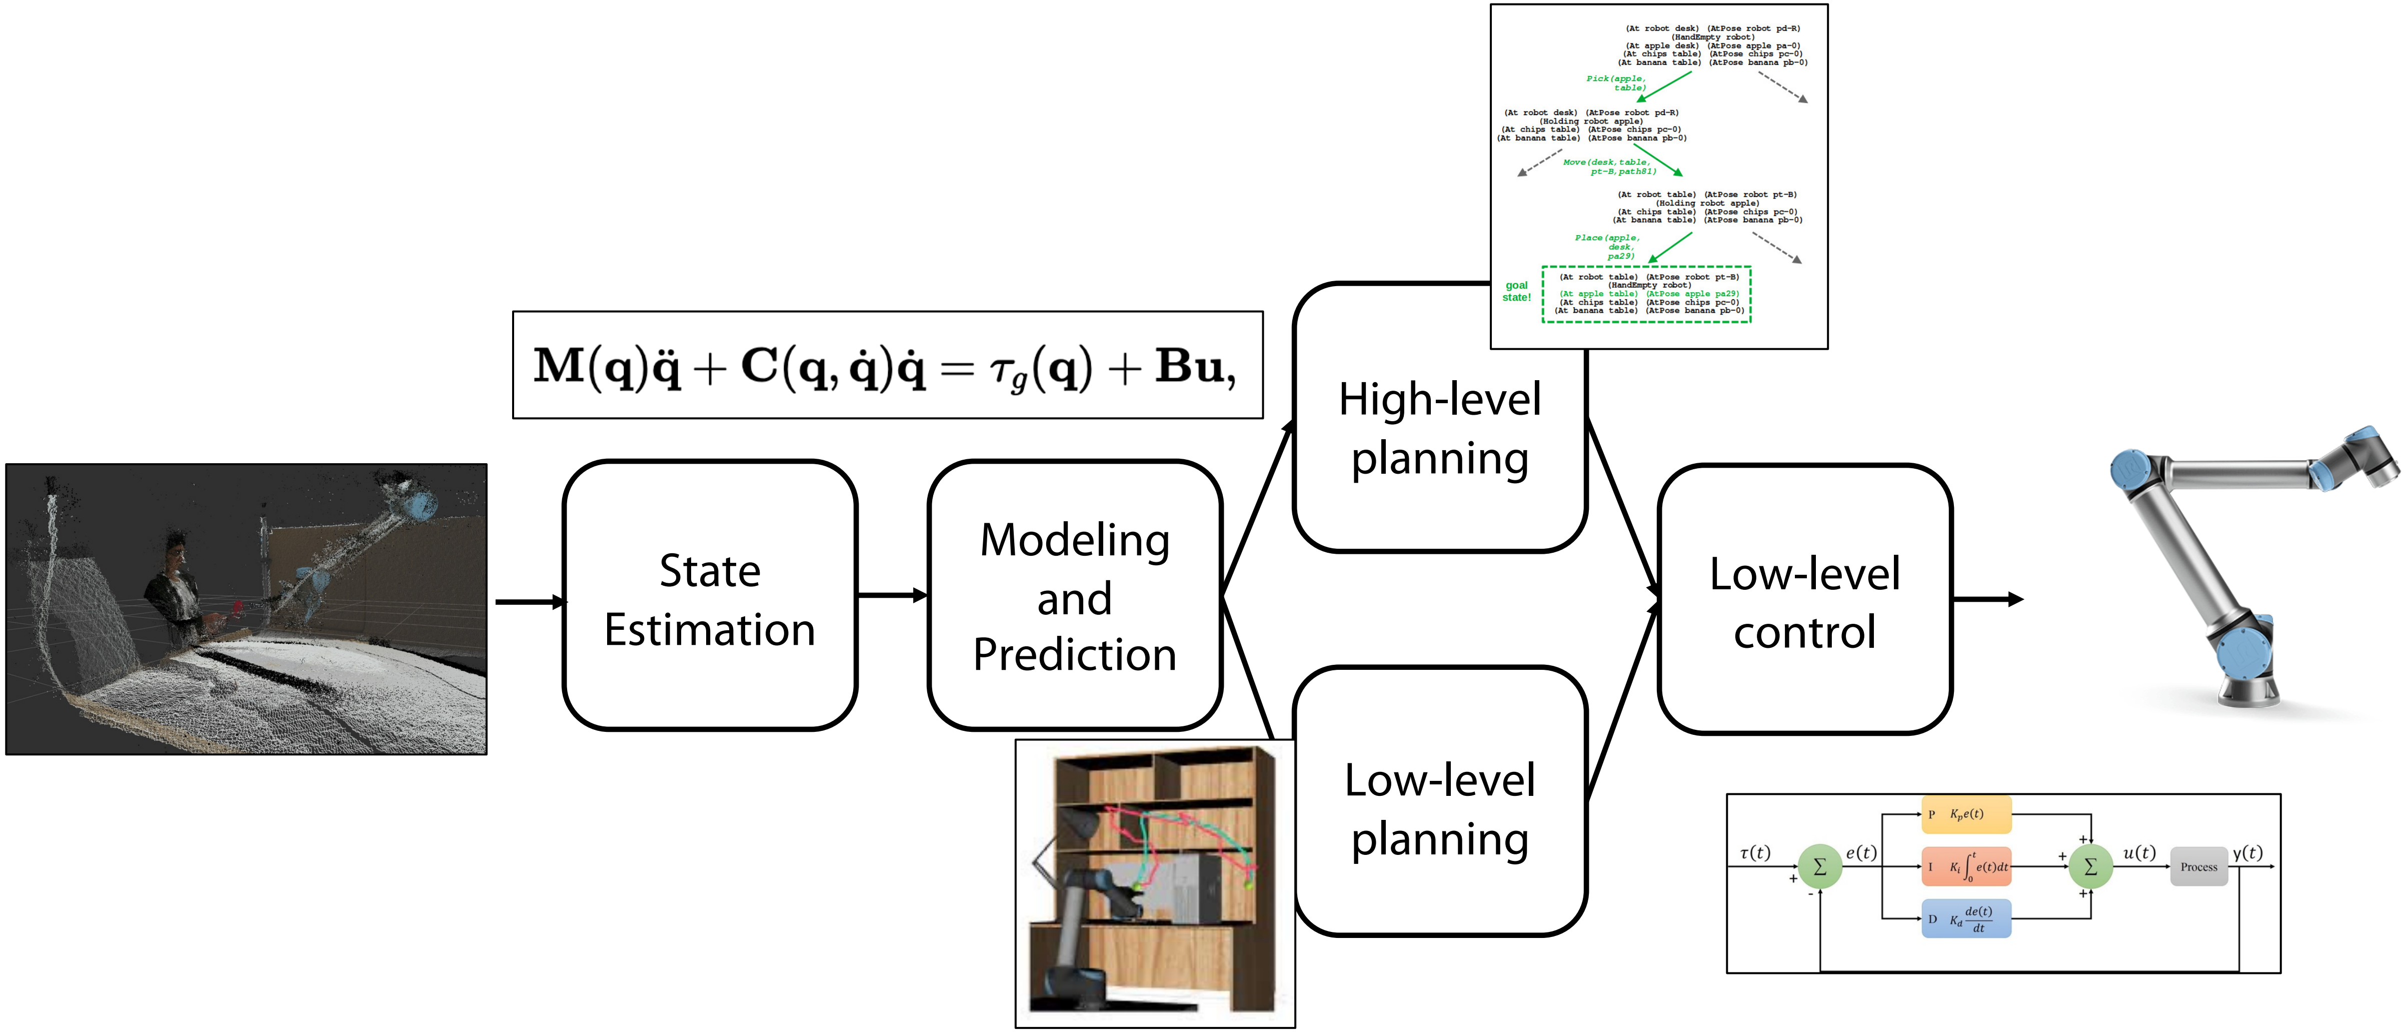

Deep reinforcement learning pipeline for robotics

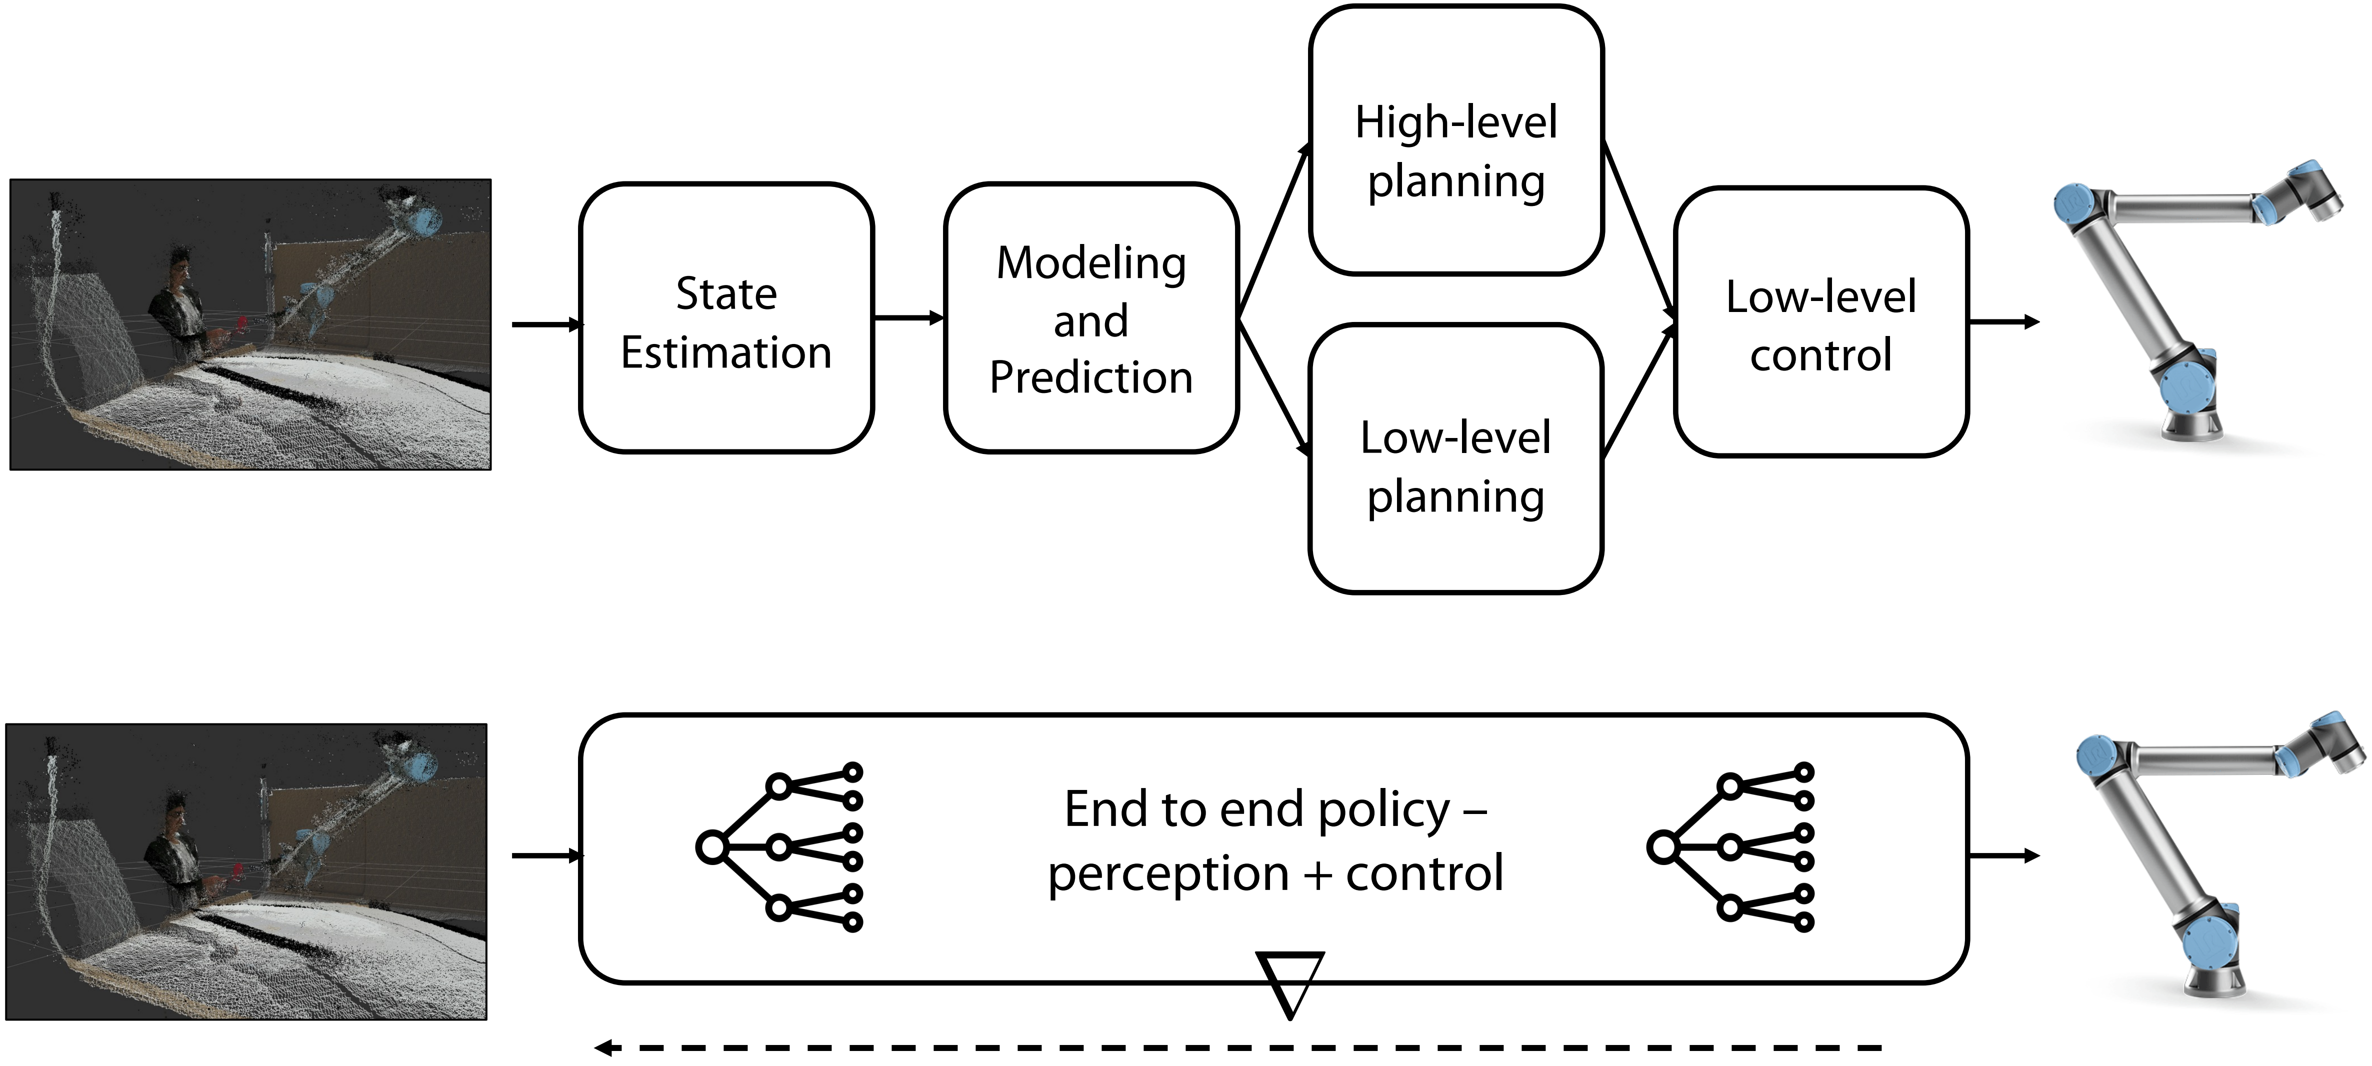

Not quite so simple, agent environment interface must be chosen!

Why might we 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

Course Outline

- Filtering/Smoothing
- Localization
- Mapping
- SLAM

- Search
- Motion Planning
- TrajOpt
- Stability/Certification

- MDPs and RL
- Imitation Learning
- Solving POMDPs

Goal of this course

- Understand what makes robotics so challenging
- Cover fundamental techniques and provide historical context on methods in robotics estimation and planning
- Provide exposure to state of the art and modern techniques in robotics and control

What broader tools will we learn?

Estimation:

Bayesian Inference

Maximum likelihood inference

Control:

Discrete search

Convex optimization

Dynamic Programming

End to end:

Statistical inference

Deep neural networks

Reinforcement Learning

Useful beyond robotics across decision making problems

What we will not cover?

Kinematics or Dynamics Modeling

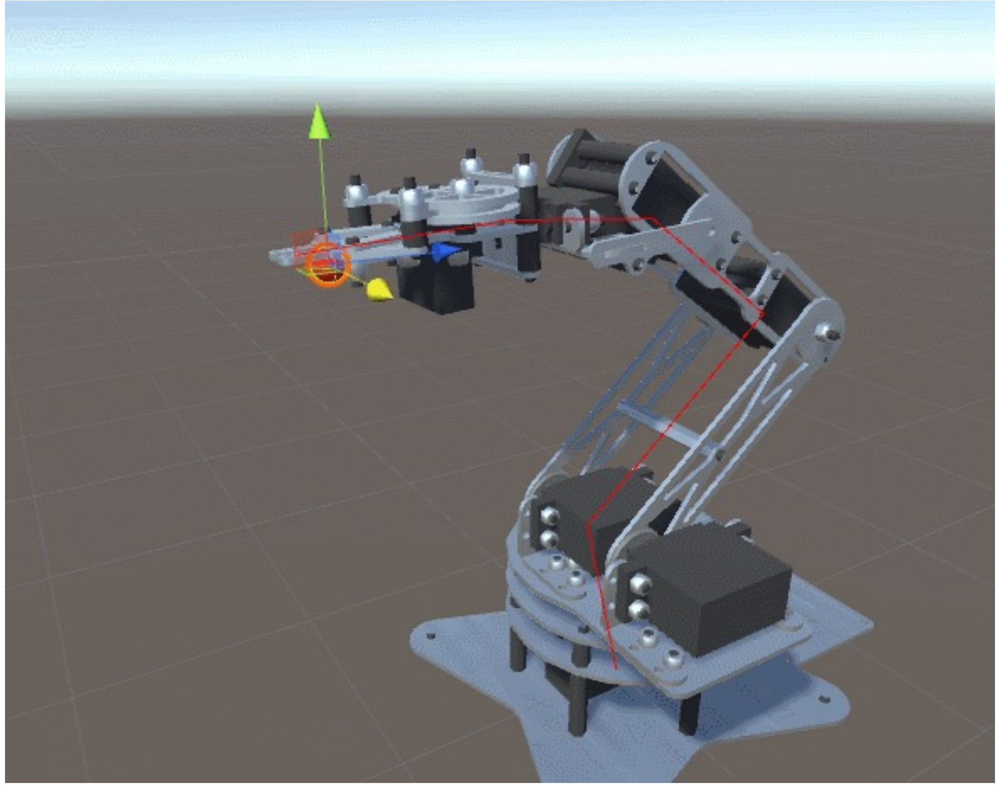

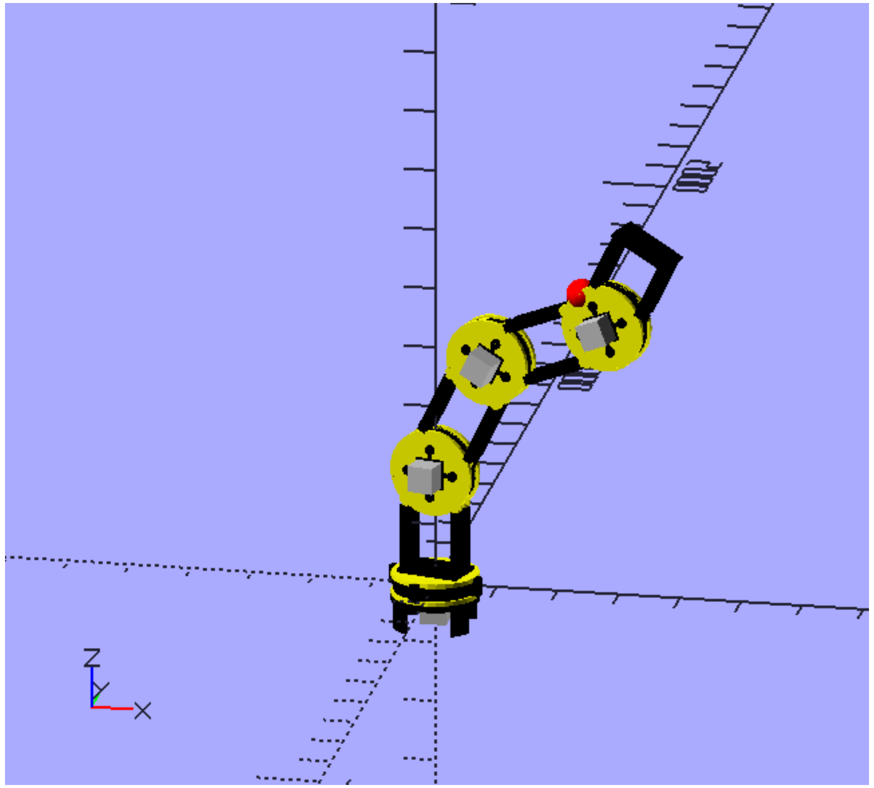

$$\mathbf{M}(\mathbf{q})\ddot{\mathbf{q}} + \mathbf{C}(\mathbf{q}, \dot{\mathbf{q}})\dot{\mathbf{q}} = \tau_g(\mathbf{q}) + \mathbf{B}\mathbf{u},$$

What we will not cover?

Advances in Computer Vision

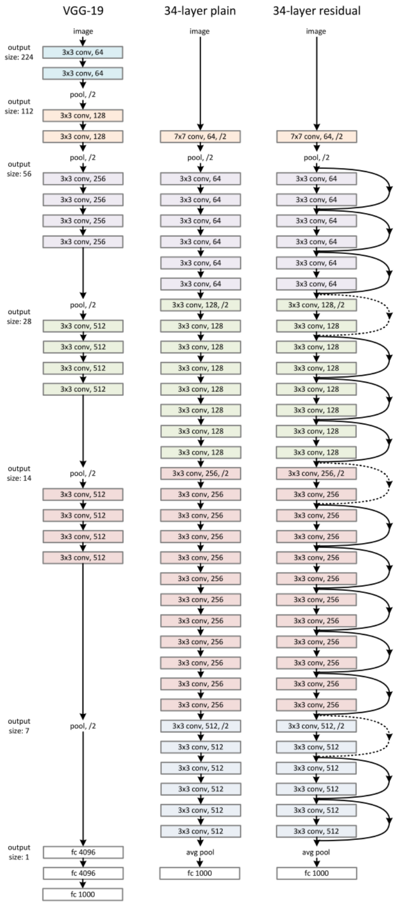

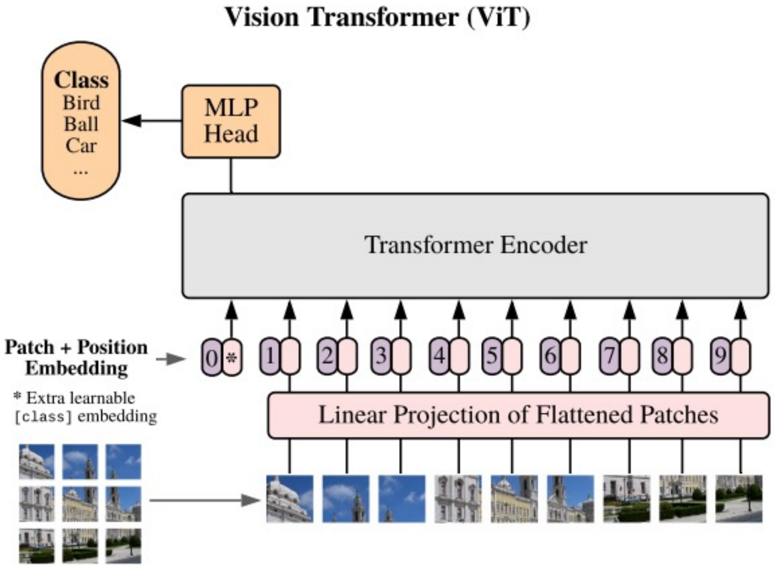

What we will not cover?

Task and Motion Planning