AA 598B Special Topics

Decision-Making & Control for Safe Interactive Autonomy

Instructor: Prof. Karen Leung

Autumn 2024

https://faculty.washington.edu/kymleung/aa598/

Announcements

- Guest lecture on Wednesday by Dr. Boris Ivanovic, Senior Research Scientist and Manager in NVIDIA Autonomous Vehicle Research Group
  - Submit talk review/reflection
- Homework 2 out tomorrow (will be light)
- Start thinking about project proposals
  - Due Nov 1 Friday
- Next Wednesday long paper discussion
- Useful video from IROS: https://www.youtube.com/watch?v=QYbAvOPcy0s

Project proposal

- **Research project**—A research project that is connected to topics covered in this course. It is encouraged for it be to connected to your PhD/MS research or other course projects. But the connection to the course must be evident and the contributions distinct.
- **Literature survey**—A deep dive into several papers on a chosen topic area, including your inclusion criteria, motivating questions, and insights.
- The grading for the project is as follows.
  - Project proposal due week 5 (5%)
  - Project presentation in week 10 (10%)
  - Project presentation peer review in week 10 (5%)
  - Project report due finals week (15%)

Last time

- Wrapped up behavior prediction models
  - Ontological "theory of mind" approaches
  - Phenomenological "deep generative model" approaches
  - Many different datasets and open-source code available
  - Many different metrics are used to evaluate prediction performance
  - Prediction models are being adapted for generate realistic human behaviors for simulators
    - Still a hard problem (stability, controllability, multi-agents, etc)

Today

- Start planning module!
  - Defining the problem
  - Techniques to solve the problem

Interaction-aware planning

Module #2


Traffic in the right lane is exiting a freeway.

Goal: Compute $u_{robot} = \pi(x_{robot}, x_{humans}, e)$

What makes interaction-aware planning challenging?

(discussion)

- multiple agent interaction → combinatorial explosion → hard to find tractable planning alg.
  - selecting a view cone.
- sensing/tracking multiple agents can be hard. → affect planning frequency
- uncertainty in human tendencies. → models for uncertainty, ideally avoid strong assumptions.
- conservative vs efficiency
- accounting for out of distribution events.
- adapting to new/unseen settings./OOD (out of dis.)
- planning horizon?
- feedback from human↔robot
  - behaviors may amplify

Multimodal uncertainty

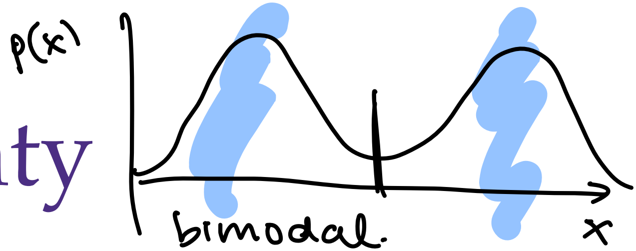

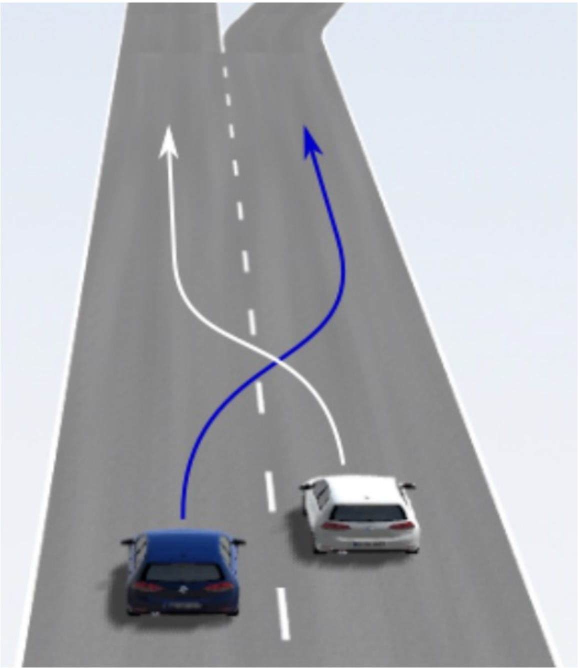

Schmerling et al 2018

GMM: gaussian mixture model

- the mean is not a good measure of outcome.
- outcomes can be very distinct → lead to very different plans.

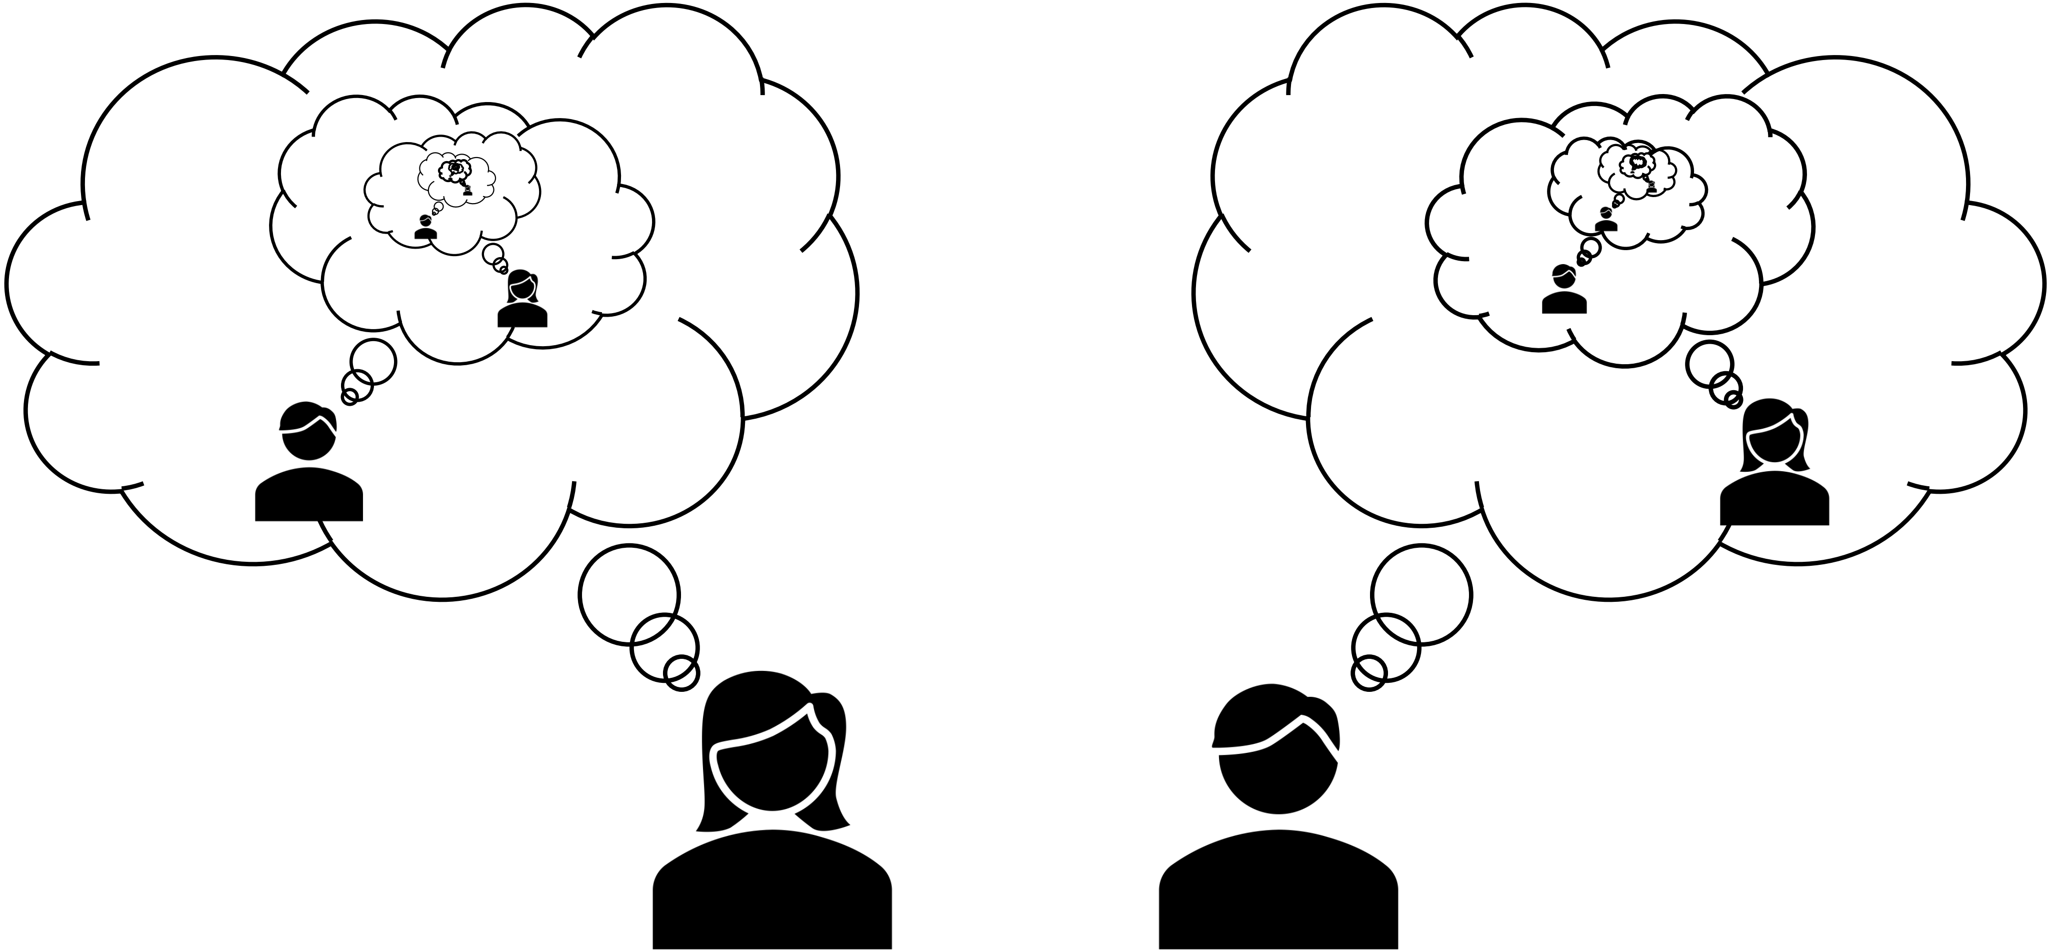

Need to account for how others may respond to your own actions

Decision-making under uncertainty

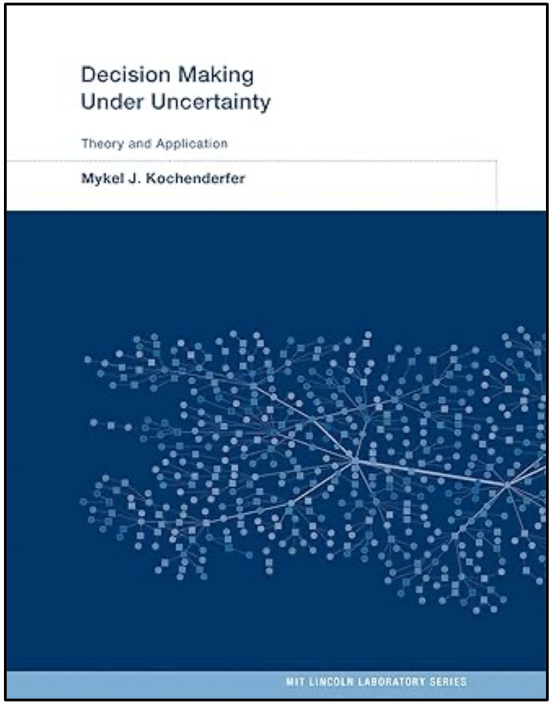

https://mykel.kochenderfer.com/textbooks/

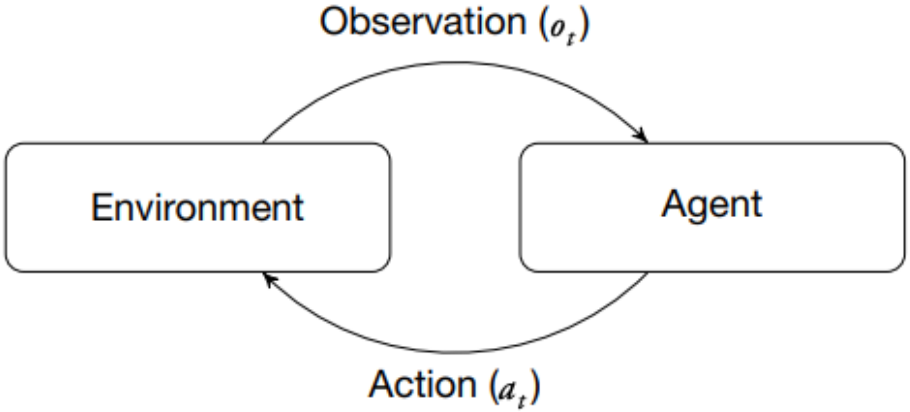

Planning problem: Find **actions** that accomplish the desired **task**

- Actions: How are actions represented?

- next robot state for the next timestep
  + tracking controller (PID, LQR)
- control inputs. $u_R = \pi(x_R, \ldots)$
- desired trajectory / sequence of waypoints + tracking controller.

Planning problem: Find **actions** that accomplish the desired **task**

Task: How is the task defined?

- objective functions + constraints → mathematical functions describing these.
- indicator function for task success/failure.
- demonstrations of success/failure
- LLMs / define through language.
- temporal logic, a formal language to express specifications.
  STL, LTL

Approaches to solving a planning problem

- **Search-based**: Enumerate over all possible options and pick the best one BFS, DFS, A*, sampling-based motion planning
  PRM, RRT*, FMT*

Approaches to solving a planning problem

- **Search-based**: Enumerate over all possible options and pick the best one
- **Supervised learning**: Mimic what an expert did (i.e., behavior cloning)

Approaches to solving a planning problem

- **Search-based**: Enumerate over all possible options and pick the best one
- **Supervised learning**: Mimic what an expert did (i.e., behavior cloning)
- **Optimization-based**: Assume problem dynamics and frame as optimization problem fmincon, cvxpy, IPOPT...

Approaches to solving a planning problem

- **Search-based**: Enumerate over all possible options and pick the best one
- **Supervised learning**: Mimic what an expert did (i.e., behavior cloning)
- **Optimization-based**: Assume problem dynamics and frame as optimization problem
- **Reinforcement learning**: Learn from interactions & optimize

CS599 @ Abhishek Gupta.

General optimal planning problem

(assume no human) $x_R^{(t)} \in \mathbb{R}^n$, $u_R^{(t)} \in \mathbb{R}^m$, $t \in \mathbb{R}_+$

goal is to find $u_R^{(0)}, \ldots, u_R^{(T)}$ that accomplishes the task.

$$\min_{u_R^{(0:T)}} \quad J\left(x_R^{(0:T+1)}, u_R^{(0:T)}\right)$$

s.t. $x_R^{(t+1)} = f\left(x_R^{(t)}, u_R^{(t)}\right)$ dynamics

$g_i\left(x_R^{(0:T+1)}, u_R^{(0:T)}\right) \geq 0 \quad i = 1, \ldots, G$ inequality eg. control limits, obstacles, speed..

$h_j\left(x_R^{(0:T+1)}, u_R^{(0:T)}\right) = 0 \quad j = 1, \ldots, H$ equality initial state. goal state

eg. STL.