16-832 Spring'26
Integrated Planning and Learning

Introduction;
Why Integrate Planning and Learning

Maxim Likhachev

Robotics Institute

Carnegie Mellon University

Class Logistics

- Instructor:

  Maxim Likhachev – maxim@cs.cmu.edu

- TA:

  Gopal Venkitachalam – gopalakt@andrew.cmu.edu

- Website:

  http://www.cs.cmu.edu/~maxim/classes/integratedplanningandlearning

- Announcements, Questions:

  - on Piazza
  - should have received an email with access info

About Me

- My Research Interests:
  - Planning, Decision-making, Learning
  - Applications: planning for complex robotic systems including aerial and ground robots, manipulation platforms, small teams of heterogeneous robots

- More info: *http://www.cs.cmu.edu/~maxim*
- Search-based Planning Lab: *http://www.sbpl.net*



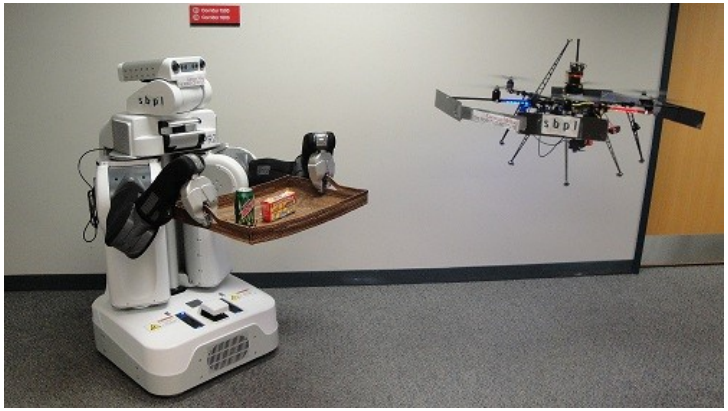
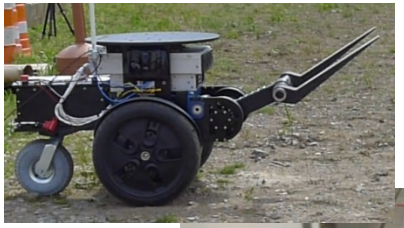
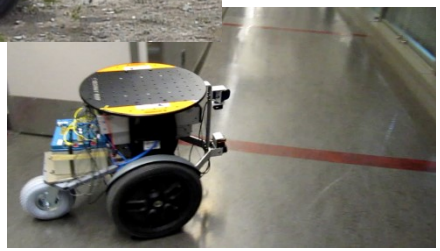

About Me

- Also, at Waymo, where I'm heavily involved in planning for self-driving vehicles

Planning vs. Learning

Planning-based control

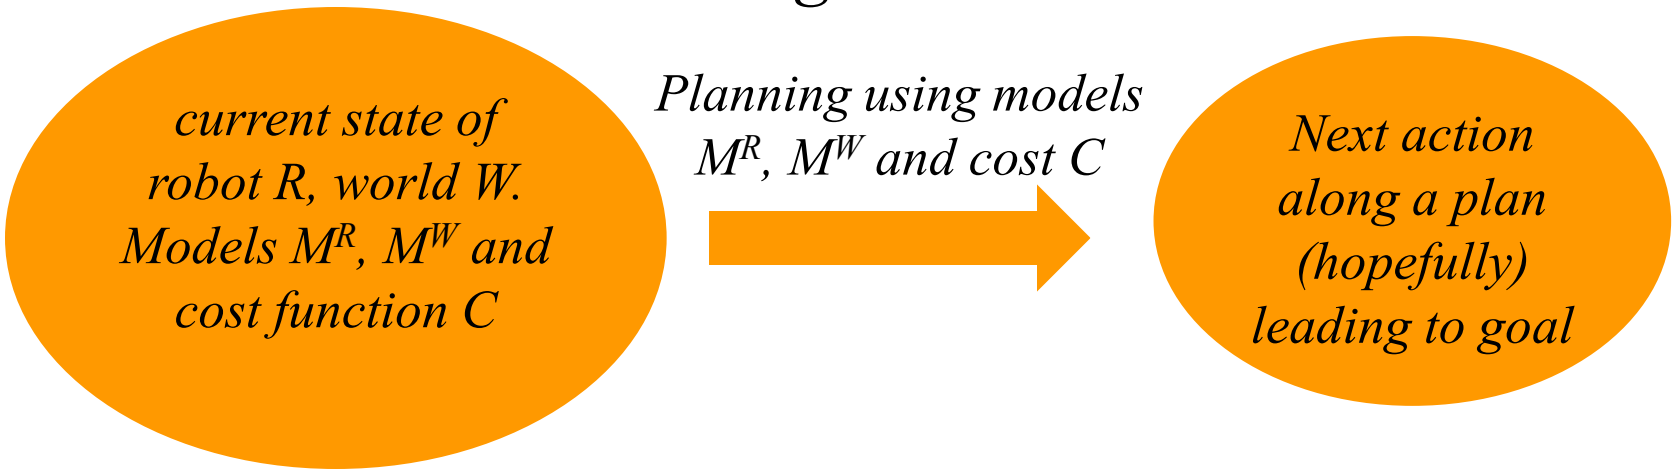

Learning-based control

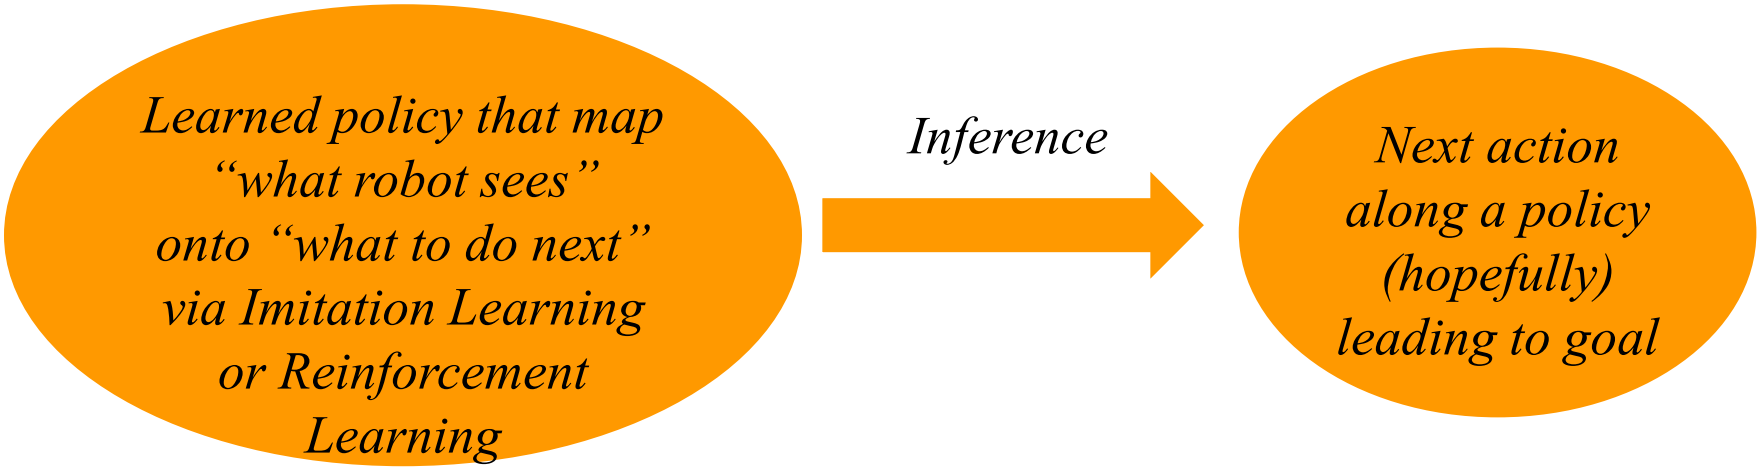

Class Objectives at High-level

- Understand the motivation behind integrating Planning and Learning
- Identify different sub-areas of Integrated Planning and Learning
- Study state-of-the-art approaches within each sub-area
- Research new state-of-the-art approaches within these sub-areas

Class Prerequisites

- Strong knowledge of planning
- Strong knowledge of learning
- Programming (e.g., C, C++, python)
- Exposure to robotics

Why Integrate Planning and Learning?

- From the planning point of view:

Why Integrate Planning and Learning?

- From the planning point of view:
  - speed up planning

planning with dynamics and contact modeling [Ram et al., TRO‘24]

Why Integrate Planning and Learning?

- From the planning point of view:
  - speed up planning;
  - hard-to-design constraints/goal condition

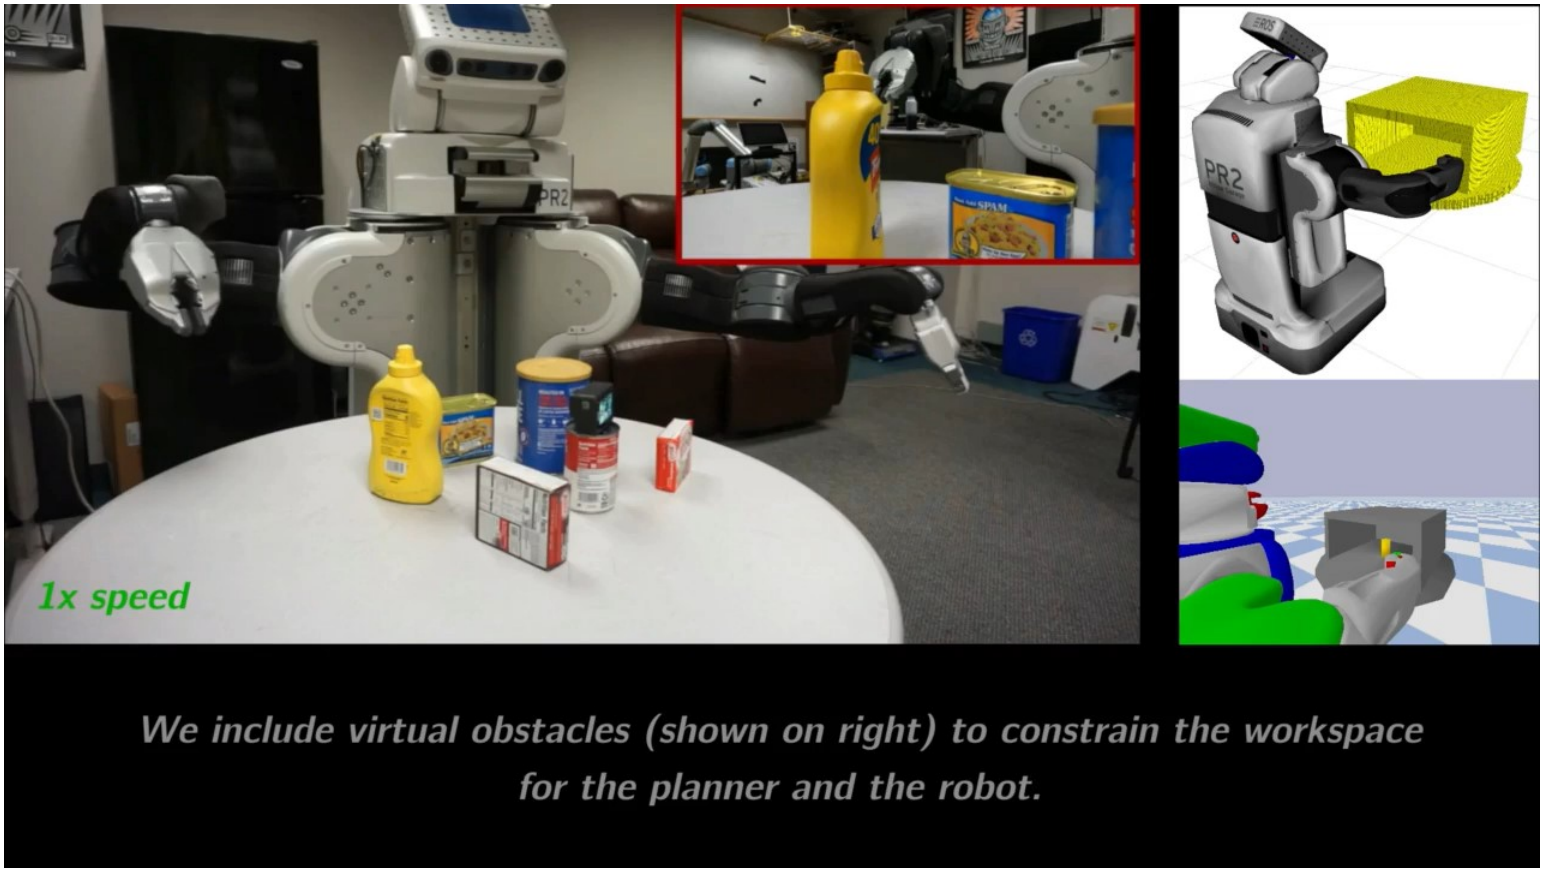

planning with physics-based simulation in the loop [Dhruv et al., ICRA'24]

Why Integrate Planning and Learning?

- From the planning point of view:
  - compensate for imperfect model of the world/robot

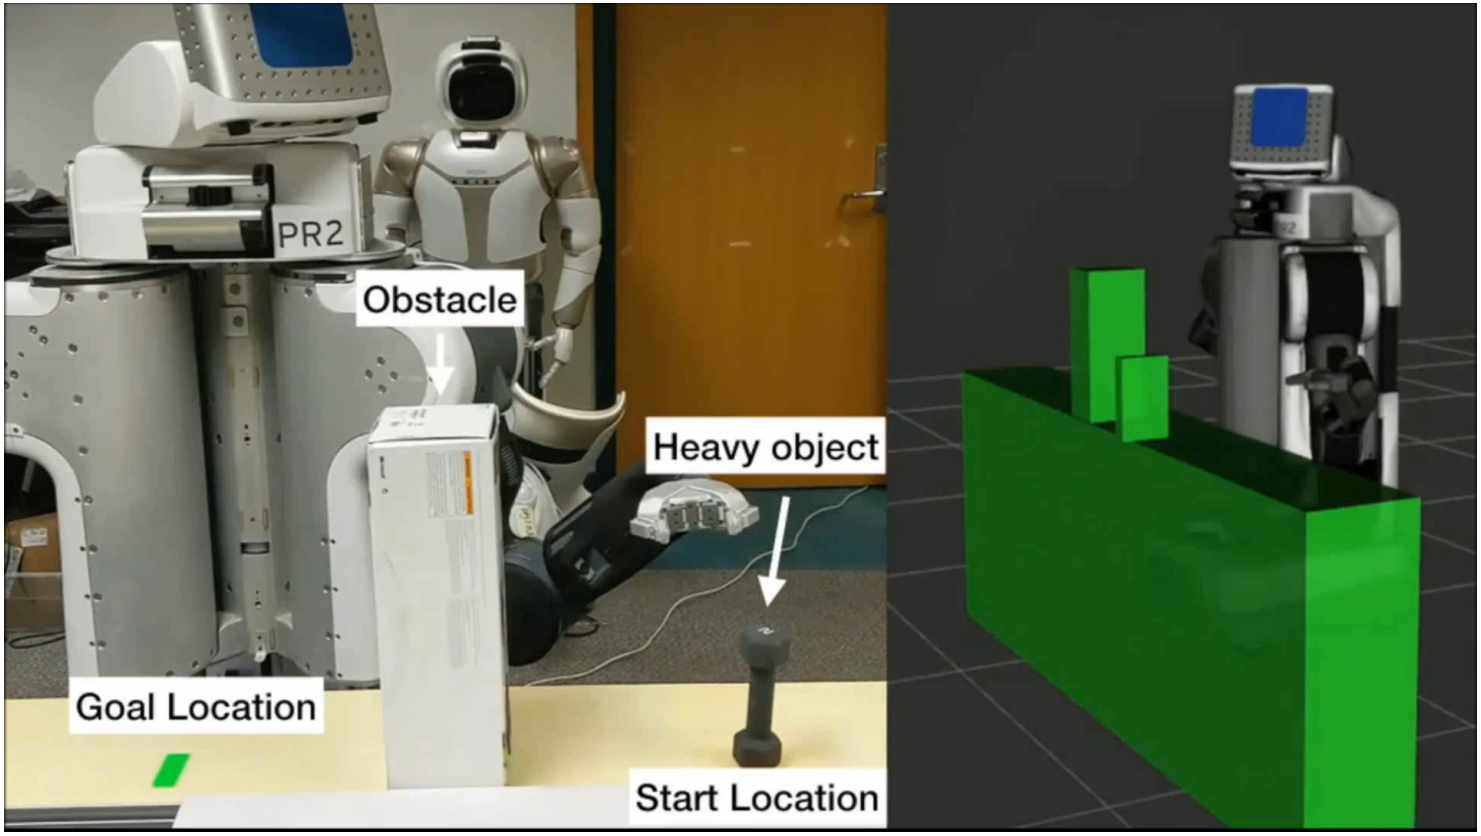

CMAX planning [Vemula et al., RSS'20]

Why Integrate Planning and Learning?

- From the planning point of view:
  - hard-to-design cost function

Planning for quadrupeds[Vernaza et al., ICRA '09]

Why Integrate Planning and Learning?

- From the learning point of view:

Why Integrate Planning and Learning?

- From the learning point of view:
  - scale up learning to long-horizon tasks
  - provide guarantees on solution quality and safety in long-tails
  - use planner to generate examples for training or during distillation
  - improve inference quality for autoregressive models

Tentative Class Schedule

Date	Day	Topic	Papers
12-Jan	Mon	Introduction; Why Integrate Planning and Learning?	
14-Jan	Wed	Integrating learning into planning: speeding up planning	"A Survey on the Integration of Machine Learning with Sampling-based Motion Planning" by McMahon et al. "Roadmaps with Gaps over Controllers: Achieving Efficiency in Planning under Dynamics" by Sivaramakrishnan et al. "Learning Heuristic Functions for Mobile Robot Path Planning Using Deep Neural Networks" by Takahashi et al. "Learning Local Heuristics for Search-Based Navigation Planning" by Veerapaneni et al.
19-Jan	Mon	NO CLASS	
21-Jan	Wed	Integrating learning into planning: learning cost functions for planning	"Learning to search: Functional gradient techniques for imitation learning" by Ratliff et al. "Learning Navigation Costs from Demonstration via Differentiable Planning" by Wang et al.
26-Jan	Mon	Integrating learning into planning: support for imprecise goal conditions for planning	"Open-World Task and Motion Planning via Vision-Language Model Inferred Constraints" by Kumar et al. "A Survey of Optimization-Based Task and Motion Planning: From Classical to Learning Approaches" by Zhao et al.
28-Jan	Wed	Integrating learning into planning: planning with imperfect world dynamics model	"CMAX++ : Leveraging Experience in Planning and Execution using Inaccurate Models" by Vemula et al. "SACHA: Soft Actor-Critic With Heuristic-Based Attention for Partially Observable Multi-Agent Path Finding" by Lin et al.
2-Feb	Mon	Integrating planning into learning: scaling up learning-based control to long-horizon tasks	"SPIN: distilling Skill-RRT for long-horizon prehensile and non-prehensile manipulation" by Jung et al. "Skill Discovery for Exploration and Planning using Deep Skill Graphs" by Bagaria et al.
4-Feb	Wed	Integrating planning into learning: ensuring safe task achievement in learning-based control	"Motion Planning Diffusion: Learning and Adapting Robot Motion Planning with Diffusion Models" by Carvalho et al.
9-Feb	Mon	Integrating planning into learning: learning from planning	"Planning-Guided Diffusion Policy Learning for Generalizable Contact-Rich Bimanual Manipulation" by Li et al. "Offline Imitation Learning Through Graph Search and Retrieval" by Yin et al.
11-Feb	Wed	Integrating planning into learning: improving inference via planning	"Monte Carlo Tree Diffusion for System 2 Planning" by Yoon et al. "Stream of Search (SoS): Learning to Search in Language" by Gandhi
16-Feb	Mon	Learning to plan	"Value Iteration Networks" by Tamar et al. "Path Planning using Neural A* Search" by Yonetani et al.
18-Feb	Wed	research project proposal presentations	
23-Feb	Mon	research project proposal presentations	
25-Feb	Wed	Integrating learning into planning: speeding up planning	"Using VLM Reasoning to Constrain Task and Motion Planning" by Yan et al. "MOSAIC: A Skill-Centric Algorithmic Framework for Long-Horizon Manipulation Planning" by Mishani et al.
2-Mar	Mon	SPRING BREAK; NO CLASS	
4-Mar	Wed	SPRING BREAK; NO CLASS	
9-Mar	Mon	Integrating learning into planning: learning cost functions for planning	"Learning Navigation Costs from Demonstration with Semantic Observations" by Wang et al. TBD
11-Mar	Wed	Integrating learning into planning: support for imprecise goal conditions for planning	TBD
16-Mar	Mon	Integrating learning into planning: planning with imperfect world dynamics model	"Learning Skills to Patch Plans Based on Inaccurate Models" by LaGrassa et al. "Guided Uncertainty-Aware Policy Optimization: Combining Learning and Model-Based Strategies for Sample-Efficient Policy Learning" by Lee et al.
18-Mar	Wed	Integrating planning into learning: scaling up learning-based control to long-horizon tasks	"Motion Planner Augmented Reinforcement Learning for Robot Manipulation in Obstructed Environments" by Yamada et al. "Search on the Replay Buffer: Bridging Planning and Reinforcement Learning" by Eysenbach et al. "Planning-Augmented Hierarchical Reinforcement Learning" by Gieselmann et al.
23-Mar	Mon	research project progress presentations	
25-Mar	Wed	research project progress presentations	
30-Mar	Mon	Integrating planning into learning: ensuring safe task achievement in learning-based control	TBD
1-Apr	Wed	Integrating planning into learning: learning from planning	TBD
6-Apr	Mon	Integrating planning into learning: improving inference via planning	TBD
8-Apr	Wed	Learning to plan	TBD
13-Apr	Mon	TBD	
15-Apr	Wed	TBD	
20-Apr	Mon	final project presentations	
22-Apr	Wed	final project presentations	

Students are responsible for:

- One paper presentation (individual, around 30 mins)
- Semester-long research project (2-3 people per group)
- In-class participation

Class Structure

- Grading

Research project	70%
In-class participation and paper presentations	30%

- No exam
- No late policies

Paper presentation

- About 30 min presentation

- Format of the presentation:
  - statement for **what** is being proposed
  - motivation for **why**
  - analysis of **related** work
  - **assumptions** being made
  - explanation of **how** it works (preferably with examples)
  - any **theoretical guarantees**
  - **experimental** analysis
  - **limitations/concerns**

- One person per paper – **please pick an available day/paper** in the sign-up list

Semester-long research project

- In groups of 2-3 people per project (individual projects possible but check with me first)

- Three presentations:
  - Initial proposal presentation
  - Mid-progress presentation
  - Final presentation

- At least two individual meetings with each group

- Outcome: 5+ page report of paper submission quality

- Immediate next step:
  - organize into groups and start thinking about a project
  - send me a paragraph by email with project proposal by **Feb 7**

Potential (pre-made) test domains for the project

- Push-T – push T-shape object into the desired configuration
  - can be configured to be:
    - without or with obstacles
    - fully observable (pose of T is known) or partially observable (only visual input)
  - potential issues to focus on:
    - too slow to simulate each potential action considered by planning
    - planning with inaccurate model
    - learning-based control for long-horizon tasks (when with obstacles)
    - …

Potential (pre-made) test domains for the project

- Grasping object amidst clutter
  - requires non-prehensile manipulation (e.g., pushing your way through to grasp the desired object)
  - potential issues to focus on:
    - too slow to simulate each potential action considered by planning
    - planning with inaccurate model
    - unspecified constraints
    - …

Potential (pre-made) test domains for the project

- Multi-agent path finding (MAPF)
  - generate collision-free plans for large number of agents operating in warehouse-like environments
  - potential issues to focus on:
    - how to speed up centralized planning
    - how to improve the quality of overall solution (e.g., learn roadmap-following behaviors)
    - how to achieve semi-decentralized control (e.g., learn individual policies yet provide guarantees on overall performance)
    - …

Questions about the class?