

16-350 Spring'26
Planning Techniques for Robotics

Introduction;
What is Planning for Robotics?

Maxim Likhachev
Robotics Institute
Carnegie Mellon University

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, currently split between CMU and [Waymo](#), where I'm heavily involved in planning for self-driving vehicles

Class Logistics

- Instructor:
Maxim Likhachev – maxim@cs.cmu.edu
- TA:
Barath Satheeshkumar - bsathees@andrew.cmu.edu
- Website:
<http://www.cs.cmu.edu/~maxim/classes/robotplanning>
- Piazza for Announcements and Questions:
You should have received an email

Class Logistics

- Books (optional):
 - Planning Algorithms by *Steven M. LaValle*
 - Heuristic Search, Theory and Applications by *Stefan Edelkamp and Stefan Schroedl*
 - Principles of Robot Motion, Theory, Algorithms, and Implementations by *Howie Choset, Kevin M. Lynch, Seth Hutchinson, George A. Kantor, Wolfram Burgard, Lydia E. Kavraki and Sebastian Thrun*
 - Artificial Intelligence: A Modern Approach by *Stuart Russell and Peter Norvig*

Class Prerequisites

- Knowledge of programming (e.g., C, C++)
- Knowledge of data structures
- Some prior exposure to robotics (e.g., Intro to Robotics class) is preferred

Class Objectives

- Understand and learn how to implement most popular planning algorithms in robotics including heuristic search-based planning algorithms, sampling-based planning algorithms, task planning, planning under uncertainty and multi-robot planning
- Learn basic principles behind the design of planning representations
- Understand core theoretical principles that many planning algorithms rely on and learn how to analyze theoretical properties of the algorithms
- Understand the challenges and basic approaches to interleaving planning and execution in robotic systems
- Learn common uses of planning in robotics

Tentative Class Schedule

Date	Day	Topic	HW out	HW due
12-Jan	Mon	Introduction; What is Planning?		
14-Jan	Wed	planning representations: explicit vs. implicit graphs, skeletonization-, grid- and lattice-based graphs		
19-Jan	Mon	NO CLASS		
21-Jan	Wed	planning representations: explicit vs. implicit graphs, skeletonization-, grid- and lattice-based graphs (cont'd)		
26-Jan	Mon	search algorithms: Uninformed A*	HW1	
28-Jan	Wed	search algorithms: A*, Multi-goal A*		
2-Feb	Mon	heuristics, weighted A*, Backward A*		
4-Feb	Wed	interleaving planning and execution: Anytime heuristic search		
9-Feb	Mon	interleaving planning and execution: Freespace assumption, Incremental heuristic search		
11-Feb	Wed	interleaving planning and execution: Limited Horizon search, LRTA*		
16-Feb	Mon	case study: planning for autonomous driving		HW1
18-Feb	Wed	planning representations: PRM for continuous spaces	HW2	
23-Feb	Mon	planning representations/search algorithms: RRT, RRT-Connect, RRT*		
25-Feb	Wed	planning representations/search algorithms: RRT, RRT-Connect, RRT* (cont'd)		
2-Mar	Mon	SPRING BREAK; NO CLASS		
4-Mar	Wed	SPRING BREAK; NO CLASS		
9-Mar	Mon	case study: planning for mobile manipulation and articulated robots		
11-Mar	Wed	search algorithms: Markov Property, dependent vs. independent variables		HW2
16-Mar	Mon	case study: planning for exploration and surveillance tasks		
18-Mar	Wed	final project proposal presentations		
23-Mar	Mon	planning representations: state-space vs. symbolic representation for task planning	HW3	
25-Mar	Wed	search algorithms: symbolic task planning algorithms		
30-Mar	Mon	planning under uncertainty: Minimax formulation		
1-Apr	Wed	planning under uncertainty: Expected Cost Minimization formulation		HW3
6-Apr	Mon	planning under uncertainty: Solving Markov Decision Processes		
8-Apr	Wed	exam		
13-Apr	Mon	multi-robot planning		
15-Apr	Wed	multi-robot planning		
20-Apr	Mon	TBD		
22-Apr	Wed	final project presentations		

Class Structure

- Grading

Three homeworks	33%
Exam	20%
In-class pop quizzes	10%
Final project	32%
Participation	5%

- Exam is tentatively scheduled for April 8 (no final exam)
- Late Policy
 - 3 free late days
 - No late days may be used for the final project!
 - Each additional late day incurs 10% penalty with 50% being the upper limit (grade of 90 becomes 81 for one additional late day)

Three Homeworks + Final Project

- All homeworks are individual (no groups)
- Final project are in groups of 2-3 students
- Homeworks are programming assignments
- Final project is a research-like project. For example:
 - to develop a planner for a robot planning problem of your choice
 - to extend an existing or develop a new planning algorithm
 - to prove novel properties of a planning algorithm
 - Get a feel for doing research: Individual meetings with groups, Two class presentations (initial idea and final)

Three Homeworks + Final Project

- Homework assignments for Masters students will have additional scope
- Undergraduate students will have an option to tackle this additional scope and receive bonus points

What is Planning?

- According to Wikipedia: “*Planning is the process of thinking about an organizing the activities required to achieve a desired goal.*”

What is Planning for Robotics?

- According to Wikipedia: “*Planning is the process of thinking about an organizing the activities required to achieve a desired goal.*”
- **Given**
 - model (states and actions) of the robot(s) $M^R = \langle S^R, A^R \rangle$
 - a model of the world M^W
 - current state of the robot s_{current}^R
 - current state of the world s_{current}^W
 - cost function C of robot actions
 - desired set of states for robot and world G
- **Compute a plan π that**
 - prescribes a set of actions a_1, \dots, a_K in A^R the robot should execute
 - reaches one of the desired states in G
 - (preferably) minimizes the cumulative cost of executing actions a_1, \dots, a_K

Example

- **Given**

- model (states and actions) of the robot(s) $M^R = \langle S^R, A^R \rangle$
- a model of the world M^W
- current state of the robot s_{current}^R
- current state of the world s_{current}^W
- cost function C of robot actions
- desired set of states for robot and world G

- **Compute a plan π that**

- prescribes a set of actions a_1, \dots, a_K in A^R the robot should execute
- reaches one of the desired states in G
- (preferably) minimizes the cumulative cost of executing actions a_1, \dots, a_K

Planning for omnidirectional robot:

What is M^R ?

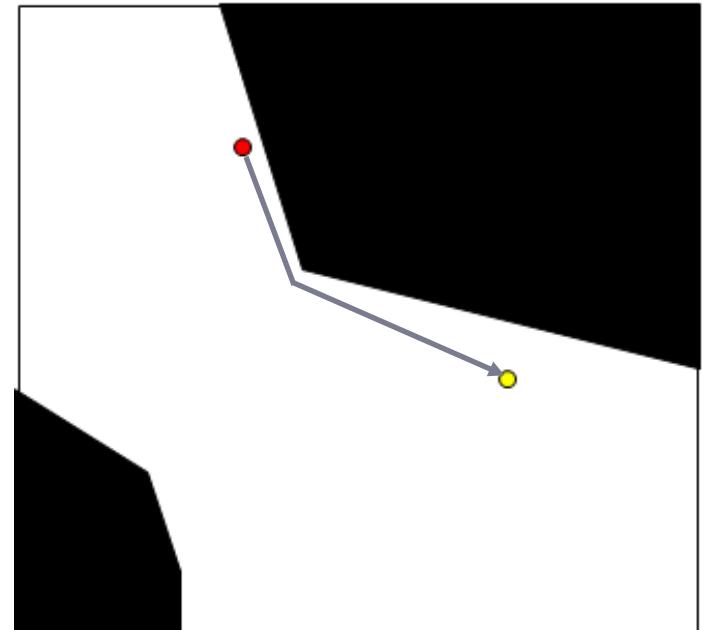
What is M^W ?

What is s_{current}^R ?

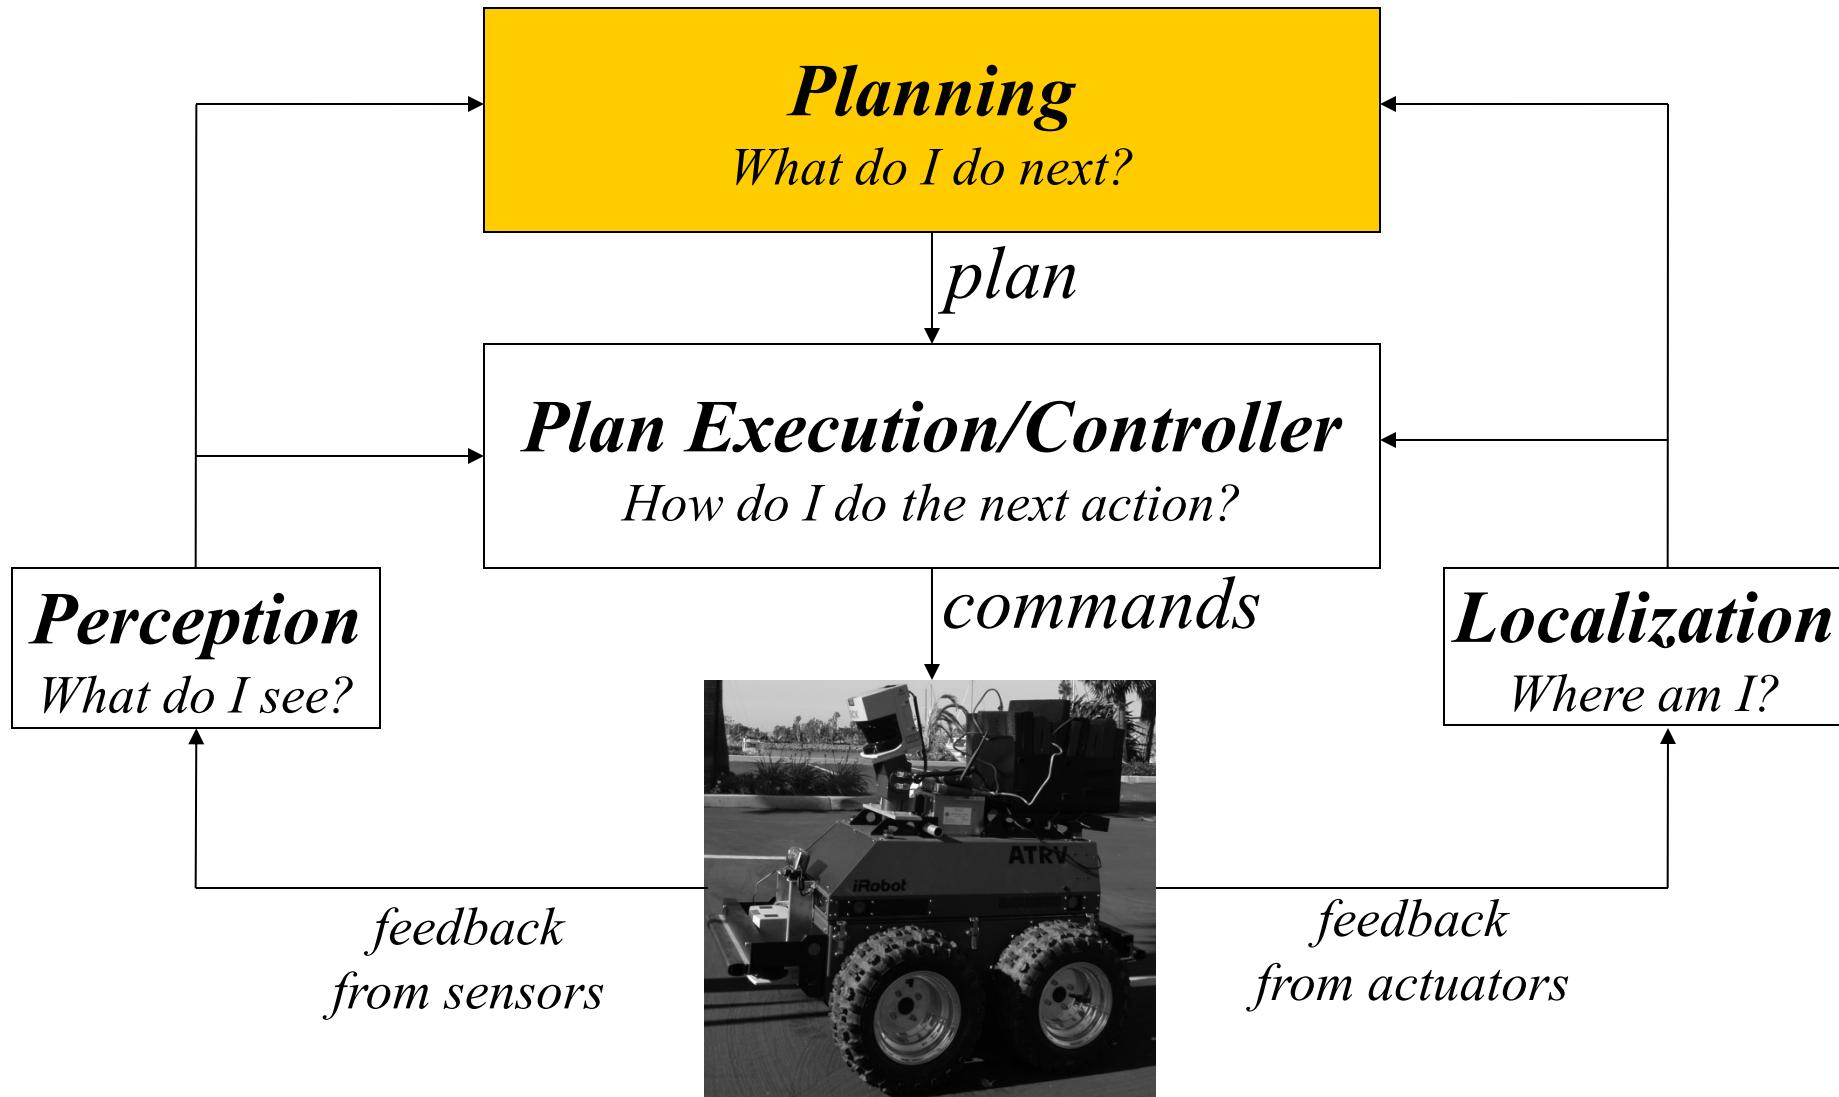
What is s_{current}^W ?

What is C ?

What is G ?



Planning within a Typical Autonomy Architecture



Few More Examples

- **Given**

- model (states and actions) of the robot(s) $M^R = \langle S^R, A^R \rangle$
- a model of the world M^W
- current state of the robot $s_{current}^R$
- current state of the world $s_{current}^W$
- cost function C of robot actions
- desired set of states for robot and world G

- **Compute a plan π that**

- prescribes a set of actions a_1, \dots, a_K in A^R the robot should execute
- reaches one of the desired states in G
- (preferably) minimizes the cumulative cost of executing actions a_1, \dots, a_K

Planning for omnidirectional drone:

What is M^R ?

What is M^W ?

What is $s_{current}^R$?

What is $s_{current}^W$?

What is C ?

What is G ?



MacAllister et al., 2013

Few More Examples

- **Given**

- model (states and actions) of the robot(s) $M^R = \langle S^R, A^R \rangle$
- a model of the world M^W
- current state of the robot $s_{current}^R$
- current state of the world $s_{current}^W$
- cost function C of robot actions
- desired set of states for robot and world G

- **Compute a plan π that**

- prescribes a set of actions a_1, \dots, a_K in A^R the robot should execute
- reaches one of the desired states in G
- (preferably) minimizes the cumulative cost of executing actions a_1, \dots, a_K

Planning for autonomous navigation:

What is M^R ?

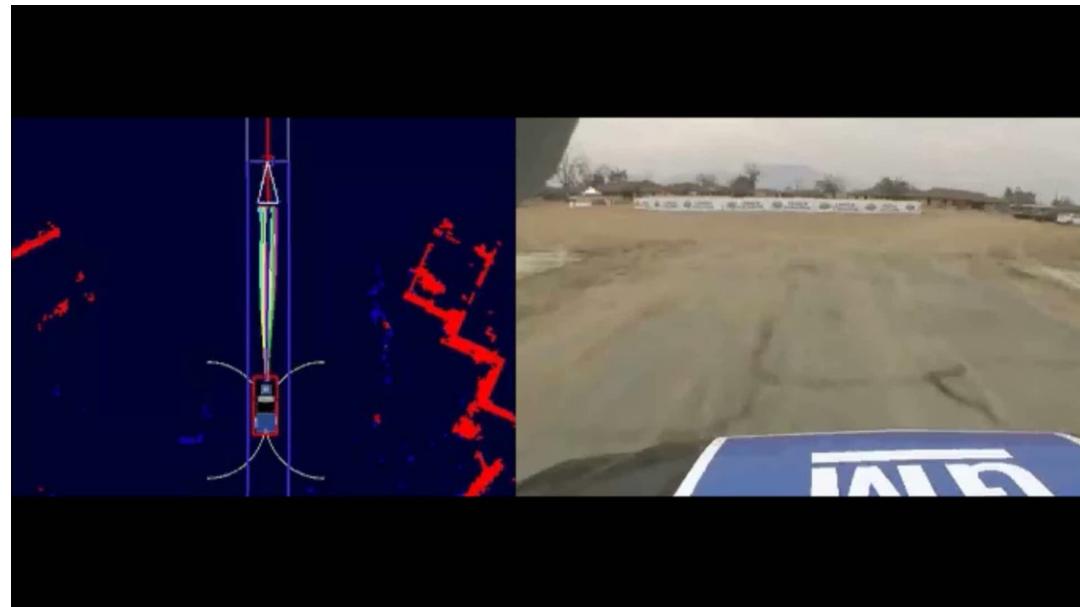
What is M^W ?

What is $s_{current}^R$?

What is $s_{current}^W$?

What is C ?

What is G ?



Likhachev & Ferguson, '09; part of Tartanracing team from CMU for the Urban Challenge 2007 race

Few More Examples

- **Given**

- model (states and actions) of the robot(s) $M^R = \langle S^R, A^R \rangle$
- a model of the world M^W
- current state of the robot $s_{current}^R$
- current state of the world $s_{current}^W$
- cost function C of robot actions
- desired set of states for robot and world G

- **Compute a plan π that**

- prescribes a set of actions a_1, \dots, a_K in A^R the robot should execute
- reaches one of the desired states in G
- (preferably) minimizes the cumulative cost of executing actions a_1, \dots, a_K

Planning for autonomous flight among people :

Narayanan et al., 2012

What is M^R ?

What is M^W ?

What is $s_{current}^R$?

What is $s_{current}^W$?

What is C ?

What is G ?



Few More Examples

- **Given**

- model (states and actions) of the robot(s) $M^R = \langle S^R, A^R \rangle$
- a model of the world M^W
- current state of the robot $s_{current}^R$
- current state of the world $s_{current}^W$
- cost function C of robot actions
- desired set of states for robot and world G

- **Compute a plan π that**

- prescribes a set of actions a_1, \dots, a_K in A^R the robot should execute
- reaches one of the desired states in G
- (preferably) minimizes the cumulative cost of executing actions a_1, \dots, a_K

Planning for a mobile manipulator robot opening a door:

Gray et al., 2013

What is M^R ?

What is M^W ?

What is $s_{current}^R$?

What is $s_{current}^W$?

What is C ?

What is G ?



Few More Examples

- **Given**

- model (states and actions) of the robot(s) $M^R = \langle S^R, A^R \rangle$
- a model of the world M^W
- current state of the robot $s_{current}^R$
- current state of the world $s_{current}^W$
- cost function C of robot actions
- desired set of states for robot and world G

- **Compute a plan π that**

- prescribes a set of actions a_1, \dots, a_K in A^R the robot should execute
- reaches one of the desired states in G
- (preferably) minimizes the cumulative cost of executing actions a_1, \dots, a_K

Planning for a mobile manipulator robot assembling a birdcage: Cohen et al., 2015

What is M^R ?

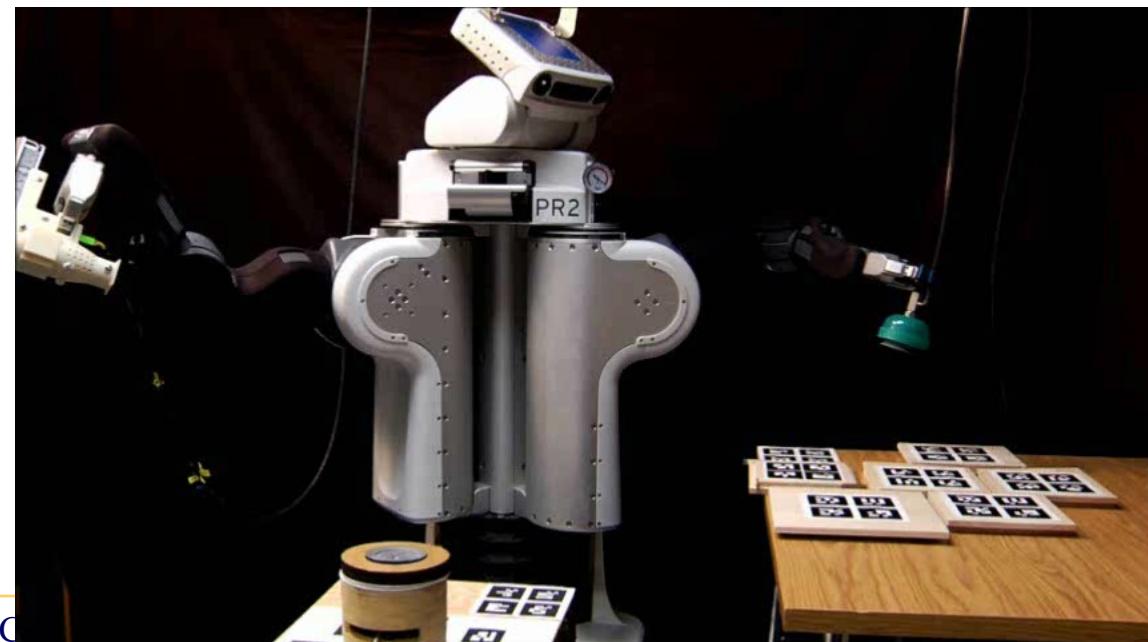
What is M^W ?

What is $s_{current}^R$?

What is $s_{current}^W$?

What is C ?

What is G ?



Few More Examples

- **Given**

- model (states and actions) of the robot(s) $M^R = \langle S^R, A^R \rangle$
- a model of the world M^W
- current state of the robot $s_{current}^R$
- current state of the world $s_{current}^W$
- cost function C of robot actions
- desired set of states for robot and world G

- **Compute a plan π that**

- prescribes a set of actions a_1, \dots, a_K in A^R the robot should execute
- reaches one of the desired states in G
- (preferably) minimizes the cumulative cost of executing actions a_1, \dots, a_K

Planning for a mobile manipulator unloading a truck:

What is M^R ?

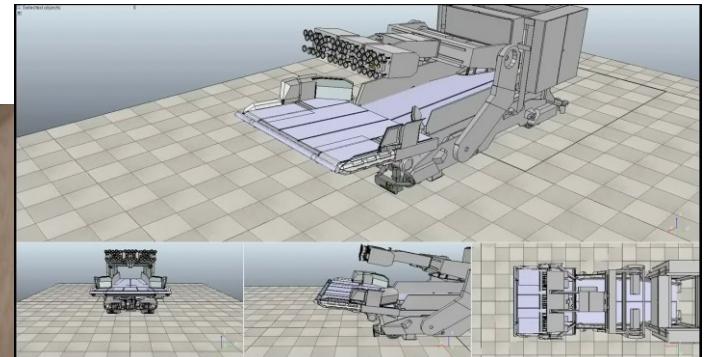
What is M^W ?

What is $s_{current}^R$?

What is $s_{current}^W$?

What is C ?

What is G ?



Assuming Infinite Computational Resources...

Where does Planning break?

Assuming Infinite Computational Resources...

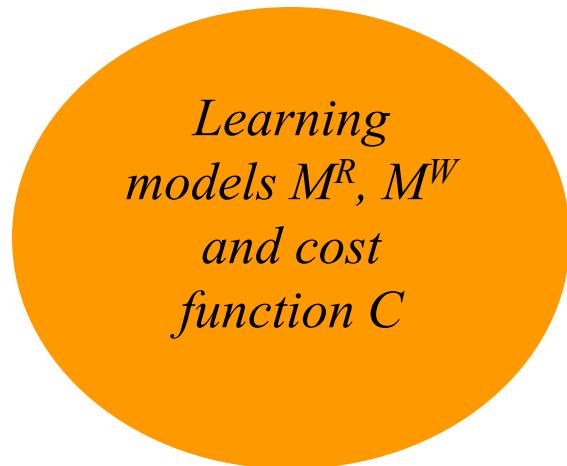
Where does Planning break?

Reliance on the knowledge/accuracy of the model!

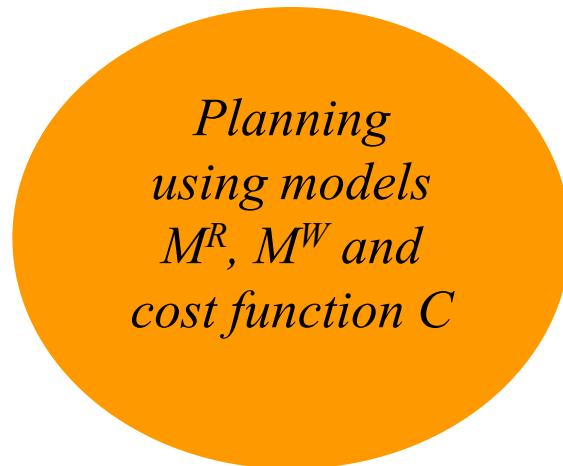
Role of Learning in Planning?

Planning vs. Learning

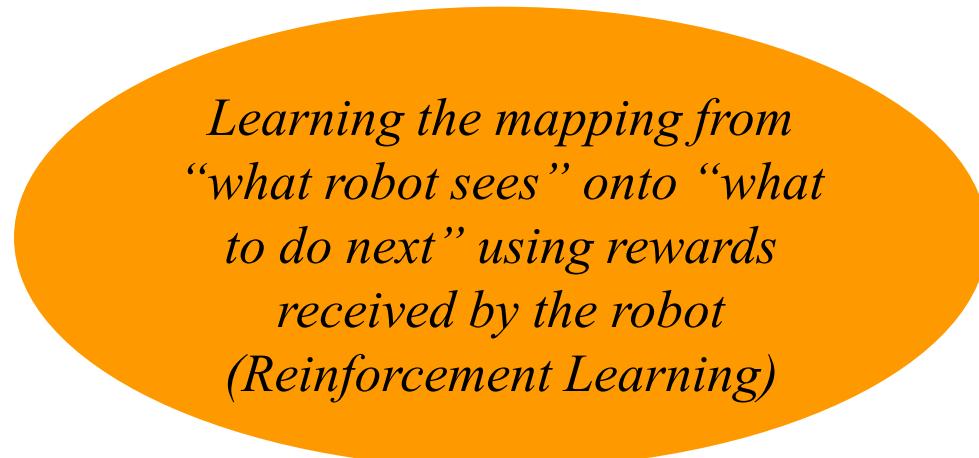
Model-based approach



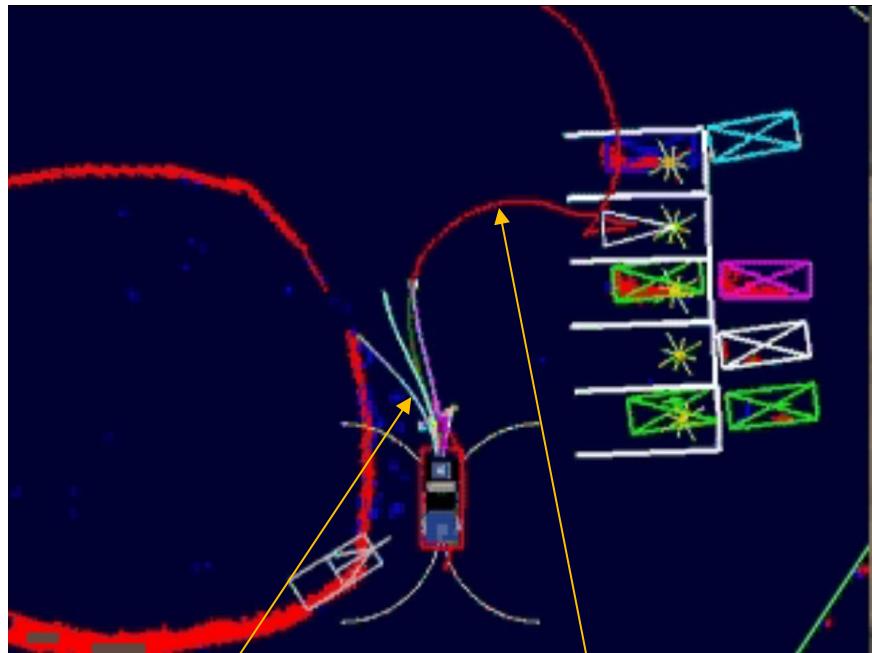
models M^R , M^W and cost function C



Model-free approach



Planning vs. Trajectory Following vs. Control

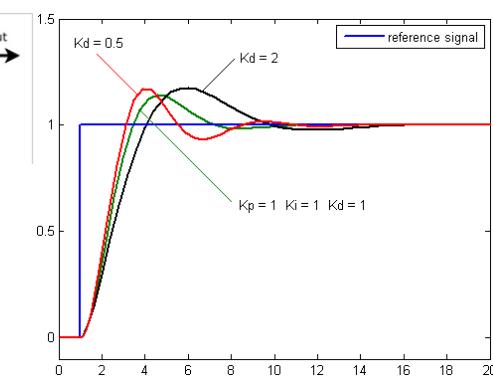
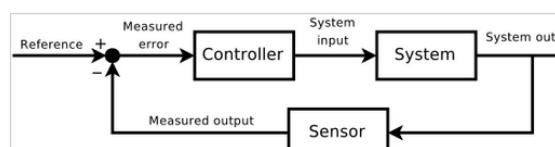


*local planning
(trajectory following)*

global planning



controller



Images from wikipedia

Questions about the class?