

Winter 2020

Stanford Center for Professional Development CS 237B: Principles of Robot Autonomy II

Instructors:

Prof. Jeannette Bohg

Office: Gates 140

Email: bohg@stanford.edu

Prof. Marco Pavone

Office: Durand 261

Email: pavone@stanford.edu

Prof. Dorsa Sadigh

Office: Gates 142

Email: dorsa@cs.stanford.edu

Course Assistants:

Ashar Alam

Email: ashar.alam@stanford.edu

Erdem Biyik

Email: ebiyik@stanford.edu

Jenna Lee

Email: yunjlee@stanford.edu

Toki Migimatsu

Email: takatoki@stanford.edu

Location and time: Packard 101, Monday and Wednesday, 1:30pm – 2:50pm.

Sections:

- Mondays, 10:30am – 12:30pm, Skilling Lab space.
- Tuesdays, 12:00 – 2:00pm, Skilling Lab space.
- Thursdays, 5:00 – 7:00pm, Skilling Lab space.

Office Hours:

Prof. Bohg: Fridays, 1:00 – 2:00pm (Gates 140), after class, and by appointment.

Prof. Pavone: Tuesdays, 1:00 – 2:00pm (Durand 261), after class, and by appointment.

Prof. Sadigh: Fridays, 9:00 – 10:00am (Gates 142), after class, and by appointment.

Course assistants: Mondays 9:00 – 10:00am, online; Tuesdays 10:00 – 12:00pm, and Fridays 3:00 – 5:00pm, in Durand 023.

Units: 3 or 4. Taking this class for 4 units entails additionally completing a paper review at the end of the quarter, with details to be announced later.

Prerequisites:

- Familiarity with ROS and basic techniques for robot autonomy (e.g., AA274A or equivalent).
- Familiarity with programming (e.g., CS 106A or equivalent) and Python.
- College calculus, linear algebra (e.g., CME 100 or equivalent).
- Basic probability and statistics (e.g., CME 106 or equivalent).

Course websites:

- For course content and announcements: <http://cs237b.stanford.edu>
- For course-related questions: <http://piazza.com/stanford/winter2020/cs237b>
- For homework submissions: <http://www.gradescope.com/courses/77478>
- For lecture videos: <http://canvas.stanford.edu/courses/112347>
- For urgent questions: cs237b-win1920-staff@lists.stanford.edu

Textbooks: There is no required textbook.

Course Content: This course teaches advanced principles for endowing mobile autonomous robots with capabilities to autonomously learn new skills and to physically interact with the environment and with humans. It also provides an overview of different robot system architectures. Concepts that will be covered in the course are: Reinforcement Learning (RL) and its relationship to optimal control, contact and dynamics models for prehensile and non-prehensile robot manipulation, imitation learning and human intent inference, as well as different system architectures and their verification. Students will learn the theoretical foundations for these concepts and implement them on mobile manipulation platforms.

Course Goals: With this course, students will:

- obtain a fundamental understanding of advanced principles of robot autonomy, including robot learning, system architectures, physical interaction with the environment, and interaction with humans;
- implement these concepts on real robot platforms.

Course Structure and Homework Policy: The class comprises four modules, roughly of equal length, namely:

1. learning-based control and perception (01/06 – 01/15);
2. system architectures, verification & validation (01/22 – 02/03);
3. interaction with the physical environment (02/05 – 02/19);
4. interaction with humans (02/24 – 03/04).

There will be a total of **four** problem sets. Rules:

- Because of the multiple topics that will be pursued in the course, it is important to keep up with the assignments. To account for unforeseen extraordinary circumstances, students are given a total of 6 free late days that may be used for the homeworks; a maximum of 3 late days will be allowed on a given assignment.
- Cooperation is allowed in doing the homework. You are encouraged to discuss approaches to solving homework problems with your classmates, however **you must always prepare the solutions on your own**. You **must** write on your problem set the names of the classmates you worked with. Copying solutions, in whole or in part, from other students or any other source will be considered a case of **academic dishonesty**.
- **Homework submissions must be typeset** (e.g., in L^AT_EX or Word.)

Sections: In addition to lectures, sections will be set up to provide a chance for students to develop skills necessary for the final project. You must sign up for a section time using the link posted on Piazza. These sign ups are first-come first-serve.

SCPD Students: Sections are not filmed, SCPD students are expected to participate via video conference.

Participation on Piazza: Piazza will be the main tool for class discussion. A student will get an extra point each time he/she (1) asks a question about lecture material; (2) answers a question about lecture material; or (3) answers a question about homework. Questions or answers should be endorsed by one of the CAs in order to receive credit. A student can accrue a maximum of five extra points. Additional details will be provided in a pinned note on Piazza.

Final Project: For the final project, students will be assigned the task of deploying some of the robot autonomy features covered in the class on a TurtleBot robot. More details about the final project will be given during the week of 02/17/2020.

SCPD Students: SCPD students will have two options for the final project:

1. If local, SCPD students may simply come to Stanford and work on the hardware platform in the same way non-SCPD students do.
2. If not local, the autonomy stack built for the final project will only be tested in simulation (namely, Gazebo).

Course Grade Calculation:

- (60%) homework.
- (20%) final exam.
- (20%) final project.
- (extra 5%) participation on Piazza.

Schedule: subject to some slippage

Date	Topic	Assignment
01/06	Course overview, intro to ML for robotics	
01/08	Markov decision processes, intro to RL	HW1 out
01/13	Model-based and model-free RL for robot control	
01/15	Learning-based perception	
01/20	Martin Luther King, Jr., Day (no classes)	
01/22	System architectures	HW1 due, HW2 out
01/27	Specifications and model checking	
01/29	Formal verification of neural networks	
02/03	System-level verification via stress testing	
02/05	Fundamentals of grasping	HW2 due, HW3 out
02/10	Grasp force optimization and planar pushing	
02/12	Learning-based grasping and manipulation	
02/17	Presidents' Day (no classes)	
02/19	Interactive perception	Final project released
02/24	Imitation learning I	
02/26	Imitation learning II	HW3 due, HW4 out
03/02	Learning from pairwise comparisons and physical feedback	
03/04	Interaction-aware control, intent inference, and shared autonomy	Final project check-in
03/09	Final exam (in class)	HW4 due
03/11	Final Project Demo (no class)	

Students with Documented Disabilities: Students who may need an academic accommodation based on the impact of a disability must initiate the request with the Office of Accessible Education (OAE). Professional staff will evaluate the request with required documentation, recommend reasonable accommodations, and prepare an Accommodation Letter for faculty dated in the current quarter in which the request is made. Students should contact the OAE as soon as possible since timely notice is needed to coordinate accommodations. The OAE is located at 563 Salvatierra Walk (phone: 723-1066, URL: <http://studentaffairs.stanford.edu/oea>).

Lecture Recordings: Video cameras located in the back of the room will record all lectures for this course. For your convenience, you can access these recordings by logging into the course Canvas site. These recordings might be reused in other Stanford courses, viewed by other Stanford students, faculty, or staff, or used for other education and research purposes. Note that while the cameras are positioned with the intention of recording only the instructor, occasionally a part of your image or voice might be incidentally captured. If you have questions, please contact a member of the teaching team.