

Stanford University
AA203: Optimal and Learning-based Control

Instructors

Dr. Marco Pavone
Office: Durand 261
Email: pavone@stanford.edu
Dr. Daniele Gammelli
Email: gammelli@stanford.edu

Course Assistants

Daniel Morton
Email: dmorton@stanford.edu
Matt Foutter
Email: mfoutter@stanford.edu

Logistics

- For course content: <https://stanfordasl.github.io/aa203>
- For lecture recordings and announcements:
<https://canvas.stanford.edu/courses/188274>
- For course-related questions: <https://edstem.org/us/courses/57727>
- For homework submissions: <https://www.gradescope.com/courses/760531>
- For urgent questions: aa203-spr2324-staff@lists.stanford.edu

Location and time

Huang Engineering Center 18. In-person lectures will take place Mondays and Wednesdays 1:30pm-2:50pm; lecture recordings will be made available to all students on [Canvas](#). In-person attendance is not required to participate in this course. All office hours in this course are in-person. Discussion sections will be held for the first four weeks of the quarter, on Fridays from 9:00am to 10:30am, in Durand 251. You have the option to attend in-person or online with a zoom link made available on [Canvas](#). Also, a recording of each recitation will be made available on [Canvas](#) after each discussion.

Office Hours

Marco Pavone: Tuesdays, 1-2pm in Durand 261 and by appointment.
Daniele Gammelli: Thursdays, 1:30-2:30pm in Durand 217 (project discussion)
Daniel Morton: Fridays, 3:00-5:00pm in Durand 217
Matt Foutter: Mondays, 3:00-4:00pm and Thursdays, 5:00-6:00pm in Durand 251 and online through the recitation zoom link. Priority will be given to in-person students.

Prerequisites

Familiarity with a standard undergraduate engineering mathematics curriculum (CME 100-106; vector calculus, ordinary differential equations, introductory probability theory) and strong familiarity with linear algebra (e.g., EE 263 or CME 200). Students will ideally have had at least one course in optimization (e.g., EE 364A, CME 307, CS 205L, CS 269O, AA 222) and/or at least one course in machine learning (e.g., CS229, CS230, CS231n) and/or control (e.g., ENGR 205, AA 212). If you have questions about the prerequisites, please ask the instructor.

Course Notes

An evolving set of partial course notes for AA 203 is available at <https://github.com/StanfordASL/AA203-Notes>. In addition to these notes, the textbooks below may be valuable for context or further reference.

Textbooks (Optional)

- D. P. Bertsekas. Dynamic Programming and Optimal Control, Vol. I and II, Athena Scientific, 2012, ISBN-10: 188652908. Price: \$134.50.
- D. P. Bertsekas. Nonlinear Programming, Athena Scientific, 2016, ISBN-10: 1886529051. Price: \$89.00.
- D. P. Bertsekas. Reinforcement Learning and Optimal Control, Athena Scientific, 2019, ISBN-10: 1886529396. Price: \$89.00.
- F. Borrelli, A. Bemporad, M. Morari. Predictive Control for Linear and Hybrid Systems, 2017, ISBN-10: 1107652871. Price: \$52.25.
- D. K. Kirk. Optimal Control Theory: An Introduction. Dover Publications, 2004, ISBN-10: 0486434842 . Price: \$23.49.
- J. B. Rawlings, D. Q. Mayne, M. M. Diehl. Model Predictive Control: Theory, Computation, and Design, 2nd Edition, Nob Hill Publishing LLC, 2017, ISBN-10: 0975937731. Available online at: <https://sites.engineering.ucsb.edu/~jbraw/mpc/MPC-book-2nd-edition-1st-printing.pdf>
- R. S. Sutton and A. G. Barto. Reinforcement Learning: An Introduction. MIT Press, 2018, ISBN-10: 0262039249. Available online at: <http://www.incompleteideas.net/book/RLbook2018.pdf>

Course Content

Optimal control solution techniques for systems with known and unknown dynamics. Dynamic programming, Hamilton-Jacobi reachability, and direct and indirect methods for trajectory optimization. Introduction to model predictive control. Robust control, model-based and model-free reinforcement learning, and connections between modern reinforcement learning and fundamental optimal control ideas.

Course Goals

To learn the theoretical and implementation aspects of main techniques in optimal control and model-based reinforcement learning. In particular, dynamic programming, Hamilton-Jacobi reachability, direct and indirect methods for optimal control, model predictive control (MPC), regression models used in model-based RL, practical aspects of model-based RL, and the basics of model-free RL. To learn how to use such techniques in applications and research work with tools such as Python, CVX, JAX, and PyTorch. At the end of the class the student will be able to:

- Apply optimal control techniques to optimize the operations of physical, social, and economic processes (e.g., aerospace vehicles, autonomous cars, robotic systems, financial systems, etc.).
- Design learning-based control schemes and apply them to the aforementioned applications.

Target audience

Undergraduate and graduate students interested in achieving an advanced knowledge of optimal control, learning-based control, and reinforcement learning. Specifically, this course should benefit anyone who performs research or plans to become a professional in the following fields of engineering: Electrical Engineering (control of electro-mechanical systems); Aeronautics & Astronautics (guidance, navigation, and control of aerospace systems), Mechanical & Civil Engineering (especially robotics, automotive), Computer Science (especially machine learning, robotics), Chemical Engineering (control of complex chemical plants). The course may be useful to students and researchers in several other fields including Neuroscience, Mathematics, Political Science, Finance, Economics.

Course Grade Calculation:

- (60%) homework; there will be 4 homework assignments, assigned roughly every two weeks.
- (40%) final project. Projects can be done in teams of up to three students. Guidelines for final projects will be posted on course website.
- (Bonus up to 5%) Endorsed posts on Ed Discussion. Students will receive a bonus of 0.5% to their final grade for each teaching-staff-endorsed answer on Ed, up to a maximum of 5%.

Homework Policy

- There will be a total of four graded problem sets.
- There will be a non-graded homework zero. This serves as a diagnostic to check if you have seen the necessary prerequisite material.
- 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 write up 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.

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/oae>).

Lecture recordings

Operated by the Stanford Center for Professional Development (SCPD), 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](#). 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 the course staff at aa203-spr2324-staff@lists.stanford.edu.

Schedule (subject to some change)

Available on the course website at <http://asl.stanford.edu/aa203>.