Spring 2026

Cornell University

CS 5757: Optimization Methods for Robotics

Instructors

Preston Culbertson
Office: CIS 459
Email: pdc79@cornell.edu

Course Assistants

Shengmiao (Samuel) Jin
Email: sj887@cornell.edu

Logistics

- Course website https://www.cs.cornell.edu/courses/cs5757/2026sp/ - all course information and materials will be posted here.
- Ed Discussion forum https://edstem.org/us/courses/94233/discussion - for asking and answering questions about the course material.
- Gradescope https://www.gradescope.com/courses/1218746 - for submitting assignments and viewing grades/feedback.
- Canvas https://canvas.cornell.edu/courses/86534 - final grades and assignment solutions will be posted here.

Location and Time

Gates Hall 114. Lectures take place in-person on Tuesdays and Thursdays from 8:40am to 9:55am.

Office Hours

Preston Culbertson: Thursdays, 2:00-4:00pm, CIS 459, or by appointment
Samuel Jin: Thursdays, 5:00-7:00pm, Malott 301H

Course Description

Robotics requires translating high-level goals into precise physical motion under real-world dynamics, sensing, and constraints. Optimization offers a principled mathematical framework for this challenge, unifying methods for planning, control, and state estimation. This course develops the theory and practice of numerical optimization, spanning sampling-based methods, Newton-style algorithms, and constrained nonlinear programming. Students will implement these methods in core robotics applications such as trajectory optimization and state estimation on manifolds. Coursework will emphasize hands-on experience: problem sets will guide students through implementing core algorithms and applying them to realistic robotics tasks, while a semester-long project allows students to explore the application of optimization-based techniques in robotics, AI, or other domains of interest.

Course Goals

To learn the theoretical foundations and practical details of numerical optimization as applied to problems in robotics. In particular, students will learn the fundamentals of nonlinear programming (gradient descent, Newton's method, sequential quadratic programming), trajectory optimization (transcription, LQR, MPC), estimation (factor graphs, SLAM), and modern optimization applications in robotics (system ID, kinematic retargeting, safety filtering). Students will learn how to practically implement these methods using Python and modern tools for automatic differentiation and optimization, including OSQP, JAX, and CVX. By the end of the course students will be able to:

- Formulate planning, control, and state estimation tasks in robotics as well-posed numerical optimization problems, specifying objectives, constraints, and problem structure.
- Implement and compare zero-, first-, and second-order optimization methods, and evaluate their convergence properties, computational tradeoffs, and suitability for different classes of problems.
- Apply optimization algorithms to robotics research problems and analyze how system dynamics, constraints, and geometric structure influence solvability and efficiency.

Prerequisites

Strong familiarity with linear algebra (e.g., MATH 2940) and vector calculus (e.g., MATH 1920). Proficiency in Python. Familiarity with basic probability theory (e.g., ENGRD 2700) recommended.

Resources

This course does not have a required textbook but will have selected readings and lecture materials provided throughout. The course will draw on material from several textbooks, including:

- Jorge Nocedal and Stephen Wright. *Numerical Optimization*, Springer New York, 2006. [link]
- Russ Tedrake. *Underactuated Robotics*, lecture notes. [link]
- Timothy D. Barfoot, *State Estimation for Robotics*, Cambridge University Press, 2017. [link]

Assignments and Grading

Course grades are determined by homework and a final project.

- **Homework (50%)** Four, equally-weighted (12.5%) assignments will be given at roughly three week intervals.
- **Final Project (45%)** Work may be done individually or in teams of up to three. Projects should investigate a problem related to optimal or learning based control or estimation. Evaluation is divided across a project proposal (5%), a milestone report (10%), and a final report with a three minute video (30%).
- **Participation (5%)** Participation credit will be awarded based on completing (pre- and mid-)course surveys, and participation in lecture activities (e.g., PollEverywhere questions).

- **Bonus (up to 5%)** Students earn 0.5% of extra credit for each staff-endorsed answer on Ed Discussion, up to a total of 5%.

Course Policies

Inclusiveness

You should expect and demand to be treated with respect in this course by both your instructors and classmates. You belong in this classroom and we take our responsibility seriously to create an environment that supports your learning and well-being. If you have an experience in this course that does not align with these values, or your needs are not being accommodated, please reach out to Prof. Culbertson.

Collaboration

Cooperation is allowed and even encouraged on assignments, but we ask that you report any classmates with whom you consulted to complete the assignment. While discussing problems, approaches to proofs, and concepts is appropriate, sharing code and exact solutions is not. Copying solutions from other students is harmful to your learning, and will be considered a violation of Academic Integrity.

Late Policy

Assignments must be submitted by the posted due date on Gradescope. To account for unforeseen circumstances (illness, travel, heavy workload, etc.), you are provided with 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.

If you are experiencing serious difficulties or extenuating circumstances (e.g., medical or family emergency) that prevent you from completing work on time, please notify Prof. Culbertson as soon as possible to discuss a solution. Extensions beyond the assigned late days will be granted only in exceptional circumstances, such as documented illness, not for situations such as job interviews or heavy workloads in other courses.

Regrades

Regrade requests, if the case is strong and a significant number of points are at stake, should be submitted via Gradescope within one week of grades being released to students for that deliverable. You must provide a justification for the regrade request.

Student Disability Accommodations

Your access in this course is important to us. Please request your accommodation letter early in the semester, or as soon as you are registered with Student Disability Services (SDS), so there is enough time to arrange the academic accommodations you are approved for.

Once SDS issues your accommodation letter, it will be emailed to both you and me. After you receive it, reach out to the instructional staff so we can coordinate the details of your accommodations in a way that fits the structure of the course.

If you encounter any access barriers in the course, whether in printed materials, graphics, online content, or communication, let Prof. Culbertson or SDS know right away so the issue can be resolved promptly.

If you need an immediate accommodation, speak with Prof. Culbertson after class or email both Prof. Culbertson and SDS at sds_cu@cornell.edu.

If you have a disability or think you may have one, contact SDS for a confidential discussion at sds_cu@cornell.edu or visit sds.cornell.edu.

Mental health and well-being

Your mental health and well-being are important to me. Your ability to engage fully in this course depends on your overall well-being, and Cornell provides resources intended to support you throughout the semester. Mental and emotional strain can build quickly, and acknowledging that early is often the most effective step toward staying healthy.

Information about campus support options is available at `https://mentalhealth.cornell.edu/get-support/support-students`. Cornell also offers 24/7 phone consultation at 607-255-5155.

You may also reach out to Prof. Culbertson, your college student services office, or Cornell Health if you need additional support.

Generative AI

You may use generative AI tools (ChatGPT, Copilot, and so on) in this course but should treat them as you would a human collaborator. This means that their use should be clearly disclosed (assignments will require a statement on AI usage), and that copying solutions exactly from them is not acceptable. We ***do not*** recommend the use of IDEs with integrated AI tools (e.g., Cursor) for this class, and will treat the use of agentic workflows (i.e., automatic file creation and editing) as a form of copying code.

Tentative Schedule

Date	Topic	Assignment
Module 1: Foundations		
01/20	What is an optimization problem? Costs, constraints, dynamics	HW0 out
01/22	Dynamics: Manipulator equations, integration	
01/27	Jacobians, linearization and automatic differentiation	
01/29	$SE(3)$: Representing and perturbing orientations	HW1 out
Module 2: Numerical methods		
02/03	Unconstrained optimization: Gradient descent and line search	
02/05	Second-order methods: Newton's method, trust regions	
02/10	Nonlinear least-squares, Gauss Newton, IK case study	
02/12	Constraints and KKT conditions	HW1 due, HW2 out
02/17	*February break, no classes*	
02/19	Convexity, Quadratic Programs (QPs)	Project proposal due
02/24	Penalty methods and Augmented Lagrangian	
02/26	Sequential Quadratic Programming (SQP)	
Module 3: Trajectory optimization		
03/03	Optimal control overview, Linear Quadratic Regulator (LQR)	
03/05	Iterative LQR (iLQR), DDP	HW2 due, HW3 out
03/10	Handling constraints in MPC (via SQP)	
03/12	Zero-order methods: sampling-based MPC	
03/17	Sampling-based MPC continued: MPPI, DIAL, CEM	
03/19	Case study: legged locomotion	Project milestone due
Module 4: Measurement uncertainty and state estimation		
03/24	The dual problem: MAP inference and least squares	
03/26	(Extended) Kalman Filtering as Gauss-Newton	HW3 due, HW4 out
03/31	*Spring break, no classes*	
04/02	*Spring break, no classes*	
04/07	Factor graphs and sparse least squares	
04/09	SLAM, bundle adjustment	
Module 5: Dynamics uncertainty and robustness		
04/14	System identification: estimating model parameters	
04/16	System ID continued: residual physics and model-based RL	HW4 due
04/21	Safety filters: Control Barrier Functions (CBFs) and shielding	
04/23	Sim-to-real: Domain randomization as robust optimization	
04/28	Guest lecture: TBA	
04/30	Project workshop	
05/05	Course summary, final project videos	Project reports due