USC Viterbi School of Engineering

EE 599, Learning and Control for Safety-Critical Robotic Systems

Units: 04

Term: Spring 2023
Day—Time: MW, 4-5:30pm
Location: TBD

Instructor: Somil Bansal
Office: EEB 324
Office Hours: TBD
Contact Info: *somilban@usc.edu*

Teaching Assistant: TBD
Office:
Office Hours:
Contact Info:

IT Help: TBD
Hours of Service:
Contact Info:

Course Description

Machine learning has led to tremendous progress in domains such as computer vision, speech recognition, and natural language processing. Fueled by these advances, machine learning approaches are now being explored to develop intelligent physical systems that can operate reliably in unpredictable environments. These include not only robotic systems such as autonomous cars and drones, but also large-scale transportation and energy systems. However, learning techniques widely used today are extremely data hungry and lack the necessary mathematical framework to provide guarantees on correctness, causing safety concerns as data-driven physical systems are integrated in our society.

The course covers the mathematical foundations of dynamical system safety analysis and modern algorithmic approaches for robotic decision making in safety-critical contexts. The focus is on safe robot learning, paying special attention to robot model uncertainty and ensuring robot safety under vision and perception-based sensors.

The course will start with an overview of background material from relevant subfields: control theory, machine learning, and robotics. This will be followed by advanced techniques (reachability analysis, Lyapunov and barrier functions, etc.) in this area. These advanced techniques will also be demonstrated via reviewing recent research papers that develop and validate them on a variety of robotic applications, including navigation and manipulation. The course will conclude with an overview of recent work in ensuring and updating safety guarantees while learning, especially in the presence of vision and perception sensors. Project work as part of the course will provide a flavor of research in this new emerging area. The course is primarily intended for MS and PhD students in Robotics, Controls, Machine Learning, and Communications and Signal Processing Areas.

Learning Objectives

1. Providing students a review of the current state of safe control of autonomous systems and robots.
2. To provide a platform for students to explore projects related to this area.
3. Providing students an overview of recent advances in safe learning for robotic systems.
4. To introduce students to new research directions for safety-critical control.

Prerequisite(s): Students are recommended to have some background in control theory and/or robotics. They must have taken one of the following courses: EE 482 (Linear Control Systems) or EE 585 (Linear System Theory) or EE 445/CSCI 445 (Introduction to Robotics), or EE 545/CSCI 545 (Robotics).

Co-Requisite (s): N/A

Concurrent Enrollment: course(s) that must be taken simultaneously: N/A

Recommended Preparation: Students are recommended to have a solid background in linear algebra and calculus. Though not required, some familiarity with machine learning is highly recommended. Programming knowledge is also required.

Course Notes

Grading Type: letter grade
The course is Web-Enhanced **(Blackboard)**.
Syllabus, homeworks and other class information will be posted on Blackboard.

Technological Proficiency and Hardware/Software Required

Students will be assumed to have basic programming skills.

Required Readings and Supplementary Materials

Required textbook: There will be no required textbook. Course will be based on seminal papers from the past, recent papers on active areas of research, and technical reports. Lectures will cover selected material drawn from some of the additional texts listed below.

Additional recommended text:

1. *Sebastian Thrun,, Wolfram Burgard, and Dieter Fox. Probabilistic robotics. MIT press, 2005. [Thrun05]*
2. *Control System Design: An Introduction to State-Space Methods* by B. Friedland 1986. [Friedland86]
3. *Dynamic Non-cooperative Game Theory* by T. Başar and G. J. Olsder. SIAM Series in Classics in Applied Mathematics Philadelphia,1999. [Basar99]
4. *Reinforcement Learning and Optimal Control* by *D. P. Bertsekas, Athena Scientific, 2019. [Bertsekas19]*

Description and Assessment of Assignments

- ***Programming Assignments:*** Students will be assigned four programming assignments. The assignments will consist of solving problems based on the lecture topics discussed at the beginning of the course and will sometimes include a "research-oriented" problem to stimulate and probe students' creativity. Programming assignments will contrast different methods for safely controlling autonomous systems on the same set of environments, and will enable students to understand the trade-offs between different techniques.
- ***Class Presentation:*** Students will present papers (after an overview provided by the instructor). Paper presentations will be followed by open discussions. Students will be graded based on their presentations.
- ***Paper Reviews:*** Students will write reviews for a subset of papers that are being read in the course.
- ***Course Project:*** In addition, students will be required to work on a course project in groups of 2-3. Projects will involve designing safety controllers in simulation (and possibly but not required to) on real-robot platforms. Students will write a project proposal, present their findings in an oral presentation, and write a conference style paper. There will be multiple project milestones along the way to guide progress. Instructor will provide feedback on the various project milestones.

Grading Breakdown

Assignment	**% of Grade**	
Programming Assignments	40%	
Class Presentation	15%	
Paper Reviews	15%	
Course Project	30%	
TOTAL	100%	

Assignment Submission Policy

Homework to be submitted in class on the due date. Late homeworks will not be accepted unless prior approval for late submission has been obtained.

Additional Policies

Attendance of the lectures is expected.

Course Schedule: A Weekly Breakdown

Week	Topic & Research Papers
1	Scope of the course, outline of topics to be covered. Autonomous Systems and Safety 101. Summary of recent trends in safety in robotics.
	Part 1: Background
2	Dynamical systems review: modeling robotic systems as dynamical systems. - Reference: Chapter 2 of [Friedland86]
	Homework#1 assigned.
3	Controller design for known dynamical systems: optimal control problem, overview of dynamic programming and model predictive control. - Reference: Chapter 1 of [Bertsekas 19]
4	Safety analysis of dynamical system: safety as a reachability problem, Hamilton-Jacobi reachability. - Reference: Chapter 8 of [Basar99]
	Homework#1 due. *Homework#2 assigned.*
5	Safety analysis of dynamical system: computation of reachable sets.
6	Case studies in safety-critical robotics: robust walking, autonomous driving, and safe airspace design.
7	Structured and unstructured uncertainty models. Introduction to learning-based control.
	Homework#2 due. *Project Ideas Due.*
	Part 2: Learning-based control of robotic systems
8	Model building for unknown systems: indirect and direct learning of models. Sample reading list: - PILCO: A Model-Based and Data-Efficient Approach to Policy Search, ICML 2011. - Deep Reinforcement Learning in a Handful of Trials using Probabilistic Dynamics Models, NIPS 2018. - Goal-Driven Dynamics Learning via Bayesian Optimization, CDC 2017. - A Survey of Iterative Learning Control, IEEE Control Systems Magazine, 2006.
	Homework#3 assigned.

9	Learning for environment constraints and perception: a case study in visual navigation. Sample reading list: - Cognitive Mapping and Planning for Visual Navigation, CVPR 2017 - Beauty and the Beast: Optimal Methods Meet Learning for Drone Racing, ICRA 2019 - Learning to Fly by Crashing, IROS 2018
10	Learning for environment constraints and perception: a case study in robotic manipulation. Sample reading list: - Learning Hand-Eye Coordination for Robotic Grasping with Deep Learning and Large-Scale Data Collection, ISER 2016 - Learning Dexterous In-Hand Manipulation, arXiv - Dex-Net 2.0: Deep Learning to Plan Robust Grasps with Synthetic Point Clouds and Analytic Grasp Metrics, RSS 2017
	Homework#3 due. *Homework#4 assigned.* *Project Proposal (conference-style submission with mock results) due.*
	Part 3: Safe learning for robotic systems.
11	Safety guarantees during learning. Handling unknown safety specifications. Sample reading list: - A General Safety Framework for Learning-Based Control in Uncertain Robotic Systems, TAC 2019. - End-to-End Safe Reinforcement Learning through Barrier Functions for Safety-Critical Continuous Control Tasks, arXiv 2019. - Safe Model-based Reinforcement Learning with Stability Guarantees, arXiv 2017.
12	Updating safety assurances online. Handling partially observable safety specifications. Sample reading list: - Reachability-Based Safety Guarantees Using Efficient Initializations, CDC 2019. - An Efficient Reachability-Based Framework for Provably Safe Autonomous Navigation in Unknown Environments, CDC 2019.
	Homework#4 due.
13	Safety assurances for perception-driven robotic systems. Sample reading list: - Compositional Falsification of Cyber-Physical Systems with Machine Learning Components, Journal of Automated Reasoning 2019 - Formal Specification for Deep Neural Networks, arXiv 2018 - Verification of Image-based Neural Network Controllers Using Generative Models, arXiv 2021
14	Conclusions and Future Outlook

15	*Final Project Presentations and Conference-style Paper Due.*

Statement on Academic Conduct and Support Systems

Academic Integrity:
The University of Southern California is a learning community committed to developing successful scholars and researchers dedicated to the pursuit of knowledge and the dissemination of ideas. Academic misconduct, which includes any act of dishonesty in the production or submission of academic work, compromises the integrity of the person who commits the act and can impugn the perceived integrity of the entire university community. It stands in opposition to the university's mission to research, educate, and contribute productively to our community and the world.

All students are expected to submit assignments that represent their own original work, and that have been prepared specifically for the course or section for which they have been submitted. You may not submit work written by others or "recycle" work prepared for other courses without obtaining written permission from the instructor(s).

Other violations of academic integrity include, but are not limited to, cheating, plagiarism, fabrication (e.g., falsifying data), collusion, knowingly assisting others in acts of academic dishonesty, and any act that gains or is intended to gain an unfair academic advantage.

The impact of academic dishonesty is far-reaching and is considered a serious offense against the university. All incidences of academic misconduct will be reported to the Office of Academic Integrity and could result in outcomes such as failure on the assignment, failure in the course, suspension, or even expulsion from the university.

For more information about academic integrity see the student handbook or the Office of Academic Integrity's website, and university policies on Research and Scholarship Misconduct.

Please ask your instructor if you are unsure what constitutes unauthorized assistance on an exam or assignment, or what information requires citation and/or attribution.

Students and Disability Accommodations:

USC welcomes students with disabilities into all of the University's educational programs. The Office of Student Accessibility Services (OSAS) is responsible for the determination of appropriate accommodations for students who encounter disability-related barriers. Once a student has completed the OSAS process (registration, initial appointment, and submitted documentation) and accommodations are determined to be reasonable and appropriate, a Letter of Accommodation (LOA) will be available to generate for each course. The LOA must be given to each course instructor by the student and followed up with a discussion. This should be done as early in the semester as possible as accommodations are not retroactive. More information can be found at osas.usc.edu. You may contact OSAS at (213) 740-0776 or via email at osasfrontdesk@usc.edu.

Support Systems:

Counseling and Mental Health - (213) 740-9355 – 24/7 on call
Free and confidential mental health treatment for students, including short-term psychotherapy, group counseling, stress fitness workshops, and crisis intervention.

988 Suicide and Crisis Lifeline - 988 for both calls and text messages – 24/7 on call
The 988 Suicide and Crisis Lifeline (formerly known as the National Suicide Prevention Lifeline) provides free and confidential emotional support to people in suicidal crisis or emotional distress 24 hours a day, 7 days a week, across the United States. The Lifeline is comprised of a national network of over 200 local crisis centers, combining custom local care and resources with national standards and best practices. The new, shorter phone number makes it easier for people to remember and access mental health crisis services (though the previous 1 (800) 273-8255 number will continue to function indefinitely) and represents a continued commitment to those in crisis.

Relationship and Sexual Violence Prevention Services (RSVP) - (213) 740-9355(WELL) – 24/7 on call
Free and confidential therapy services, workshops, and training for situations related to gender- and power-based harm (including sexual assault, intimate partner violence, and stalking).

Office for Equity, Equal Opportunity, and Title IX (EEO-TIX) - (213) 740-5086
Information about how to get help or help someone affected by harassment or discrimination, rights of protected classes, reporting options, and additional resources for students, faculty, staff, visitors, and applicants.

Reporting Incidents of Bias or Harassment - (213) 740-5086 or (213) 821-8298
Avenue to report incidents of bias, hate crimes, and microaggressions to the Office for Equity, Equal Opportunity, and Title for appropriate investigation, supportive measures, and response.

The Office of Student Accessibility Services (OSAS) - (213) 740-0776
OSAS ensures equal access for students with disabilities through providing academic accommodations and auxiliary aids in accordance with federal laws and university policy.

USC Campus Support and Intervention - (213) 740-0411
Assists students and families in resolving complex personal, financial, and academic issues adversely affecting their success as a student.

Diversity, Equity and Inclusion - (213) 740-2101
Information on events, programs and training, the Provost's Diversity and Inclusion Council, Diversity Liaisons for each academic school, chronology, participation, and various resources for students.

USC Emergency - UPC: (213) 740-4321, HSC: (323) 442-1000 – 24/7 on call
Emergency assistance and avenue to report a crime. Latest updates regarding safety, including ways in which instruction will be continued if an officially declared emergency makes travel to campus infeasible.

USC Department of Public Safety - UPC: (213) 740-6000, HSC: (323) 442-1200 – 24/7 on call
Non-emergency assistance or information.

Office of the Ombuds - (213) 821-9556 (UPC) / (323-442-0382 (HSC)
A safe and confidential place to share your USC-related issues with a University Ombuds who will work with you to explore options or paths to manage your concern.

Occupational Therapy Faculty Practice - (323) 442-2850 or otfp@med.usc.edu
Confidential Lifestyle Redesign services for USC students to support health promoting habits and routines that enhance quality of life and academic performance.