

Autonomous Robotics / Intro. to Robotics

CS 485: Autonomous Robotics
ECE 370: Intro. to Robotics
Fall 2024 | Professor Gregory J. Stein

Course Information

CS 485 / ECE 370
3 Credits
Office Hours: [TBD] RSCH 169

Instructor Information

Prof. Gregory J. Stein
He, Him, His
gjstein@gmu.edu

Course Description

The course covers basic principles of design and practice of intelligent robotics systems. We will cover algorithms for the analysis of the data obtained by vision and range sensors, basic principles of modelling kinematics and dynamics, design of basic control strategies and motion planning. Issues of uncertainty modelling, state estimation, probabilistic inference will be introduced and examined in the context of localization and map making problems. The last part of the course covers the basics and examples of learning approaches where robotic agents can learn how to achieve complex goals in reinforcement learning framework.

This course is at the upper-class undergraduate level and aims to introduce students to various topics in robotics through hands-on programming-oriented projects. The grade will be based on homework projects and an open-ended final project. Projects will involve implementation of robot capabilities in simulation. By the end of the course, students will be familiar with the standard tools used by the robotics community and be able to comprehend the basics of research papers in the field.

See the detailed class schedule below for more details on what will be covered during the course.

Course Structure

Ensuring that the students are engaging with the material and with one another can be difficult. As such, most lectures will be broken into subparts: each class will consist primarily of a lecture “module” and occasionally include a ~10 minute interactive demonstration or breakout session during which students will collaborate to solve short programming assignments related to the course material. Students should bring a laptop to class if possible so that they can follow along with group or collaborative coding sessions. Remaining class time, if it exists, will consist of a shorter lecture covering a related topic or research area.

Prerequisite Knowledge Areas:

- [basics of] Algorithms and Data Structures: familiarity with object-oriented programming is expected.
- [basics of] Artificial Intelligence: aspects including curve fitting and some of the basics of machine learning are presumed. Also, students are expected to understand the basics of search algorithms: e.g., breadth first search.
- Linear Algebra: familiarity with matrix operations, including matrix multiplication, and what an eigenvector represents.
- Calculus: students should (intuitively) understand what it means to take the derivative of a function and how that can be used to relate one quantity to its rate of change.

In addition, students are expected to have some familiarity with the Python programming language (including numpy) in which all programming projects and assignments will be done. There may be some C++ programming, though I will provide you with starter code if that becomes necessary; I do not expect students to be fluent in C++ (though it doesn't hurt).

Assignments & Grading

The grading in this course will come from three main components:

- 5 programming assignments (65%) [13% each]
- 1 final project (35%)
- Participation (0%)

All assignments will be turned in via Blackboard.

Programming Assignments

In each programming assignment, students will be expected to implement some of the algorithms we will discuss during class. Assignments will be given in the form of Jupyter Notebooks and will often include some partially written code for students to complete.

Students will be expected to write up a report for each programming assignment in LaTeX. As this is an introductory graduate course, some creativity on the solutions may be expected; some of the problems will have open-ended prompts and students should expect to explore the parameter space of the algorithms they implement and report on their findings.

In addition, programming assignments may also have a couple of more "quiz-like" questions that do not involve programming and instead test conceptual knowledge related to the material. *The quiz-like questions do not need to be typed (though it is*

encouraged)! Feel free to upload scans of your hand-written answers instead. Just try to make sure your handwriting is readable.

Students may collaborate on the programming assignments in small groups (no more than 2 or 3 students), but solutions must be written up independently.

Final Project

While the details for the final project are still TBD (and will be provided in the middle of the term), the final project will be an open-ended project, the aim of which will be to explore some topic or application that goes beyond what we have done in the programming assignments.

Participation & Lectures

Lectures will occasionally include *breakout sessions* during which students will be expected to work collaboratively to tackle some small problem related to the lecture. I have chosen *not* to include participation as a part of the overall grade. You are still strongly encouraged to come to lecture, as the breakout sessions are an integral part of the course experience and regularly attending the lecture is the best way to ensure you don't fall behind. I will upload all lecture materials in advance of lecture and additionally record lectures, uploading them to Blackboard.

Course Resources

Textbook & Readings

There will be no assigned course textbook, though there will be periodic supplemental readings that will be given, mostly as optional additional content. The lecture slides are intended to be detailed and self-contained.

Lecture Slides

I will be providing lecture slides as PDF documents after each lecture via Blackboard.

Discussion

We will be using Piazza for class discussion. It is good to post public questions there, However, so that I do not become consumed by Piazza, I try to only log on every couple of days to reply to posts, so it may not be the best way to get urgent help on an issue should you need it. Please try to start your assignments early to avoid this potential problem.

Detailed Course Schedule

The following schedule is *tentative* and subject to change, though the subjects themselves are unlikely to change. Assignments will be distributed on Tuesdays and due 16 days later by end-of-day that Thursday (the week the assignment is due). The specific deadline for each assignment will be listed on Blackboard when they are posted. Solutions will be made live one week later, posted where the assignments were distributed.

WEEK	TOPICS	PROJECTS
1	<u>Course Introduction</u> Course Overview Software and the Robot Operating System (ROS)	
2	<u>Kinematics</u> Forward and Inverse Kinematics Transformations and Lie Groups	P1 Out
3	<u>Motion Algorithms I: Control and Graph Search</u> Control I Roadmaps and Planning as Graph Search	
4	<u>Motion Algorithms II: Control and Graph Search</u> Control II Roadmaps and Planning as Graph Search	P2 Out
5	<u>Motion Algorithms III: Continuous Motion Planning</u> Motion Primitives Monte Carlo Sampling RRT*	
6	<u>Robotic Systems and Hardware</u>	P3 Out
7	<u>Introduction to Probabilistic Robotics</u> Introduction to Factor Graphs and GTSAM The Kalman Filter Motion Tracking	
8	<u>Algorithms for Tracking and Localization</u> Overview of Robot Sensors Point Cloud Registration and Laser Scan Matching Iterative Closest Point Algorithms Particle Filtering	P4 Out
9	<u>Computer Vision for Robotics I: Sparse Features and Geometry</u> Object Detection Object Tracking Data Association	
10	<u>Computer Vision for Robotics II: Object Detection and Deep Learning</u> Deep Learning Crash Course Object Detection via Deep Learning Iterative Closest Point Algorithms	P5 Out

11	<u>Simultaneous Localization and Mapping I: Fundamentals</u> Overview of SLAM Factor Graphs Revisited Loop Closure (via Place Recognition)	Final Proposal Due
12	<u>Simultaneous Localization and Mapping II: Practical Considerations</u> Modern SLAM Semantic SLAM Visual Inertial Navigation	
13	<u>Markov Decision Processes</u>	
14	<u>Course Wrap Up</u> Open Problems and Applications Ethical Implications Course Wrap Up	Final Due

Lateness Policy

Every assignment (except The Final Project) can be turned in 2 days late without penalty. I ask that you make a best effort to turn in the assignments on time, and use the additional three days if you get stuck or if “life happens” (which it seems to do a lot these days). Any additional late days will result in a 10%/day penalty for all assignments, up to a maximum of 1 week after the *original deadline*, after which the assignment will not be accepted (so that I can distribute the solutions). If you feel you need an extension in addition to the 2 free late days, please try to ask me *before the original deadline*, as I would like to avoid emails asking for extensions with only hours before late penalties start to accumulate. **If you turn in the assignments on time, you will get a 2% bonus on that assignment’s grade (up to the maximum).** You do not need to ask permission to use these late days; they are applied automatically.

Inclusion & Integrity

I stand by Mason’s [commitment to diversity and inclusion](#) and hope to foster an inclusive environment in which all feel welcome in my class.

True diversity is defined not only as differences in individual backgrounds, personal identities, intellectual approaches, and demographics; it is also the removal of barriers and the creation of space that allow individuals to fully engage in the life of the university.

Every student in this class is exactly where they belong and it is our honor to welcome each of you to join us in learning throughout this semester. Every student in this class, regardless of background, sex, gender, race, ethnicity, class, political

affiliation, physical or mental ability, veteran status, nationality, or any other identity category, is an equal member of our class.

You have the right to be called by whatever name you wish, to be referred to by whatever pronoun you identify, and to adjust these at any point. If you feel uncomfortable in any aspect of our instruction that results in any barrier to your inclusion in this course, please contact me (your professor) directly.

Honor Code Statement

The [GMU Honor Code](#) is in effect at all times. In addition, the CS Department has further honor code policies regarding programming projects, which are detailed [here](#). Any deviation from the GMU or the CS department Honor Code is considered an Honor Code violation.

Disability Accommodation

If you have a documented learning disability or other condition which may affect academic performance, make sure this documentation is on file with the Office of Disability Services and then discuss with the professor about accommodations. Submitting the paperwork at the deadline for a project or quiz is far too late! Even if you don't know if you plan on utilizing the accommodations ahead of time, it's in your best interest to prepare them ahead of time.

Mental Wellness

Graduate School can be a stressful environment. My door is always open; if you are struggling with the course work or would like someone to talk to, feel free to reach out to me. GMU also provides [many mental health resources](#) that I encourage you to look at.

Sexual Harassment and Interpersonal Violence

As a faculty member and designated "Responsible Employee," I am required to report all disclosures of sexual assault, interpersonal violence, and stalking to Mason's Title IX Coordinator per university policy 1412. If you wish to speak with someone confidentially, please contact the Student Support and Advocacy Center (703-380-1434), Counseling and Psychological Services (703-993-2380), Student Health Services, or Mason's Title IX Coordinator (703-993-8730; cde@gmu.edu).

Privacy and Email

Students must use their Mason email account to receive important University information, including communications related to this class. I cannot respond to messages sent from or send messages to a non-Mason email address.

To protect your privacy, I also cannot list your Mason email address on any public forum or provide it to any other students. You may, of course, give your email address to any other students.

Any video recordings of class meetings that are shared only with the instructors and students officially enrolled in a class do not violate FERPA or any other privacy expectation.

All course materials posted to Blackboard or other course site are private; by federal law, any materials that identify specific students (via their name, voice, or image) must not be shared with anyone not enrolled in this class.