University of Illinois Urbana-Champaign
Department of Aerospace Engineering
Spring 2022

AE 598

Autonomy Against the Odds: Stochastic Control for Motion and Mission Planning

4 credit hours

This syllabus is not an exhaustive description of all details of the course. Students are free to contact the instructor with any additional questions or concerns at any time.

Instructor

Melkior Ornik (mornik@illinois.edu)
Office: 319H Talbot
Office hours: whenever
Contact: most easily by e-mail, and then we can arrange in person

Lecture Times and Organization

While all sections of this syllabus are subject to change given the fluid COVID-19 situation, this section is particularly fluid. The current plan is to teach in person twice a week, in **3018 Campus Instructional Facility on Tuesdays and Thursdays, 1-2:20pm**. An already known exception will be the first week of classes (January 17-21), when the classes will be on Zoom following the university's COVID-19 guidance; additional exceptions are unknown, but likely (see Tentative Course Outline below). All exceptions and online lecture locations, as well as all other announcements, will be communicated through the course Canvas page **– all students must ensure they have access to the Canvas page**. Because of technical issues outside of the instructor's control, the lectures this semester will **not** be video-taped, but basic lecture slides will be posted on Canvas.

Temporary or permanent adjustments to this plan, or to the schedule of classes, are possible and likely. Attendance, while recommended because of a lack of alternative lecture delivery options, is *not* mandatory – students are welcome to adjust their course experience to their learning style, as long as doing so does not disturb learning styles of the others.

Course Description

The overarching goal of control theory is arguably design of control strategies to satisfy standard specifications such as system stability, reference tracking, safety, or reward optimality. Initial control courses perform such design while assuming that everything about the system – its environment, dynamics, and objectives – is entirely known. Such an assumption is clearly incorrect in many domains, and methods such as robust and adaptive control seek to design control strategies for systems in the presence of bounded disturbances or unknown parameters, but still under the assumption of *structural* knowledge about the system.

The future playing fields of autonomy, however, face significantly more substantial lack of knowledge than bounded disturbances and unknown parameters: systems on Earth will operate in complex, possibly changing environments with allied and adversarial human agents, while extraterrestrial systems, by their nature, operate in previously unexplored environments, often without certain structural knowledge about their behavior prior to the mission start.

This course will present one common ingredient in modeling and control design for systems whose dynamics or environments are partially unknown, dependent on human interaction, or possibly too complex for standard methods of control theory. It relies on constructing a *probability function* that describes the outcome of performing a particular action; in other words, while we may not know the exact parameters of a coin being tossed, stating that there is a 50% chance that it will fall on its head is often a reasonable approximation of the truth.

Systems whose behavior depends on a probability function are driven by *stochastic control processes*. These will be the main subject of our course. However, performing even the simplest classical control tasks in such a framework – as opposed to the standard differential or difference equations – requires the development of new theories of dynamical systems, optimal control, reachability, and adaptation. This course will serve as a *gentle* first introduction to that endeavor.

The technical focus of the course will be on the Markovian framework of finite-state, finite-action stochastic control. It will begin by reminding the students of the basics of *formal probability and measure theory*, and proceed by introducing *Markov chains* and *Markov decision processes*. These will provide an appropriate ecosystem to be able to define and construct *optimal control policies*, making use of the methods of *dynamic programming*.

Moving on from the basics of Markov chains and decision processes, we will then begin adding further realistic features possibly faced by an autonomous system:

- for systems where the whole system state cannot be determined at all times, we will consider the notions of *partial observability* and *belief spaces*,
- for systems where not only are the system dynamics stochastic, but even the probability function is partially unknown, we will introduce the notion of *robustness* in Markov decision processes, as well as the *exploration-exploitation tradeoff* and the *value of information* that enable learning about the unknown system during its mission,
- as a special case of lack of knowledge, where the system dynamics are known, but the system costs and objectives are not, we will briefly introduce the framework of *bandits*,
- in environments where there are multiple interacting agents, we will consider *multi-agent planning,*
- finally, recognizing that autonomous systems often face objectives significantly more complex than simple reachability or reward maximization, we will note the tools of *product spaces* and *automata* to describe such objectives and derive the appropriate control laws.

While – or exactly because – the course will introduce many topics in stochastic control, it will often trade-off depth for breadth. By the end of the course, students will gain familiarity with the fundamental notions and features of stochastic control processes, as well as some of the methods for control design with additional system limitations. The course will, naturally, *not* provide an in-depth study on these features, each of which is likely deserving of another course on its own.

Assignments and Grading

The deliverables for the course will consist of 5 homework assignments, an in-person (or remotely proctored, as circumstances arise) midterm, and a final project. The weights for the deliverables will be distributed as follows:

Homework 1: 10%
Homework 2: 10%
Midterm: 20%
Homework 3: 10%
Homework 4: 10%
Homework 5: 10%
Final project: 30%

Additional extra credit may be offered during the semester, but should not be counted on.

While students are encouraged to consult additional literature, all deliverables will largely follow the material covered in the lectures.

The final grades for the course will be calculated by the following formula: A-/A/A+ = 90-100, B-/B/B+ = 80-89.99, C- /C/C+ = 70-79.99, D-/D/D+ = 60-69.99, F = 0-59.99, where the "-" modifier will be assigned to those grades with the unit digit 0-1 (e.g., 91.87 = A-) and "+" modifier to those grades with the unit digit 8-9 (e.g., 78.02 = C+). The grades will not be rounded up, rounded down, nor "curved".

Form and Submission of Deliverables

Written submissions of the homework assignments and the final project will be due at **noon** (Central time) of the deadline dates indicated by the instructor (see tentative course outline for predicted dates). Late submission of a particular deliverable, if not agreed with the instructor, will be penalized at the rate of 15% of the weight of the deliverable per day (prorated for the actual delay time; e.g., a 2-hour delay incurs a penalty of 1.25% of the total value of the deliverable).

Assignments should be submitted to the course Canvas page. Students are responsible for timely submission of the assignments. If there are issues with the page, students are welcome to submit the assignments to the instructor by e-mail.

The midterm will be in person or, for those students who will need to take it remotely, in a timed classroom-like setting. It will be 80 minutes in length; the students will be able to use any books, computers, or notes during any of the homeworks, the midterm, and the final project. Students will also be able to communicate with each other while working on any deliverables *except for the midterm*, but will naturally need to write their own answers and possibly answer additional questions from the instructor, which might cause a change in the deliverable mark.

The homework assignments and, particularly, the midterm will largely follow the material covered in the lectures. The final project will require students to produce work related to a topic in stochastic control, based on (a) literature review, (b) students' independent research, or (c) a novel application of previous concepts to an aerospace application.

Emergencies, injuries, mental health issues do happen; when faced with obstacles to proceeding in class along the regular plan – and such obstacles will almost certainly arise, given the current state of flux of the society – students should contact the instructor for any modifications to the deliverable schedules.

Prerequisites and Literature

There is no required text for the course. With possible small exceptions intended for independent study, all new topics required for success in the course will be discussed during the lectures. This course is a fairly advanced control class; while there are no official prerequisites, the students *should* have taken AE 504, AE 555, ECE 515, or another course that introduced them to the mathematical machinery of control theory (which usually goes beyond the undergraduate-level introductory control courses).

The course material will often – especially initially – follow *Markov Decision Processes: Discrete Stochastic Dynamic Programming* by Martin L. Puterman, which can be found in the university library in an online version. While this textbook will be useful for the course, the material covered in the course will not cover everything in the book, and there will be elements of the course not found in the book. Later elements of the course will be more closely related to material in reinforcement learning and formal methods, including *Reinforcement Learning: An Introduction* by Richard S. Sutton and Andrew G. Barto, as well as *Formal Methods for Discrete-Time Dynamical Systems* by Calin Belta, Boyan Yordanov, and Ebru Aydin Gol. The course deliverables will not explicitly refer to any book at any point. *Students are not required to purchase any textbooks or other materials.*

Academic Integrity

Students are welcome to use any literature and work together on their homework assignments, as well as help each other on the final projects (although projects will remain individual). They are, however, required to:

1) write submissions on their own,
2) mention any peers they were working and any online or offline sources that they used, and
3) respond to any subsequent questions on the material posed by the instructor, whether in person or over e-mail. The answers to the instructor's questions may affect the assigned mark.

Because of the nature of the midterm, students are expected to work alone on it, while making use of literature. Students are required to familiarize themselves with the University's Academic Integrity Policy and Procedure, available at http://studentcode.illinois.edu/article1/part4/1-401/, and abide by that policy in full.

Accommodations

To obtain disability-related academic adjustments and/or auxiliary aids, students that require special accommodations must contact the instructor and the Disability Resources and Educational Services (DRES) as soon as possible. Students are welcome to contact the instructor at any time with any accommodation-related needs. To contact DRES, visit 1207 S. Oak St., Champaign, call 217-333-4603 (V/TTY), e-mail a message to disability@illinois.edu, or visit https://www.disability.illinois.edu. If a student is concerned that they have a disability-related condition that is impacting their academic progress, there are academic screening appointments available that can help diagnosis a previously undiagnosed disability. These may be accessed by visiting the DRES website.

Illinois law requires the University to reasonably accommodate its students' religious beliefs, observances, and practices in regard to admissions, class attendance, and the scheduling of examinations and work requirements. If there is a conflict between course deadlines and any religious observances, students should notify their instructor as soon as they realize the conflict.

Privacy and Reporting

The University of Illinois is committed to combating sexual misconduct. Faculty and staff members are required to report any instances of sexual misconduct to the University's Title IX Office. In turn, an individual with the Title IX Office will provide information about rights and options, including accommodations, support services, the campus disciplinary process, and law enforcement options. A list of the designated University employees who, as counselors, confidential advisors, and medical professionals, do not have this reporting responsibility and can maintain confidentiality, can be found here: https://wecare.illinois.edu/resources/students. Other information about resources and reporting is available here: https://wecare.illinois.edu.

Any student who has suppressed their directory information pursuant to Family Educational Rights and Privacy Act (FERPA) should self-identify to the instructor to ensure protection of the privacy of their attendance in this course. See https://registrar.illinois.edu/academic-records/ferpa/ for more information on FERPA.

Campus Emergency Plan

The university's emergency response recommendations can be found at the following website: http://police.illinois.edu/emergency-preparedness/. Students should review this website and the appropriate campus building floor plans website within the first 10 days of class: http://police.illinois.edu/emergency-preparedness/building-emergency-action-plans/.

Modifications to the Syllabus

The instructor reserves the right to modify any and all parts of this syllabus throughout the semester. All modifications will be made solely in the interest of time scheduling, beneficial adaptation to changing public health circumstances, accurate measurement of the students' success, and improvement of the students' educational outcomes. Any modifications will be transparently communicated to the students.

Tentative Course Outline

Due to the delays in course scheduling, every element of this section is preliminary.

Week	Topic	Day 1	Day 2
1: Jan 17-21	Introduction to stochastic control; probability functions	Online	Online
2: Jan 24-28	Brief introduction to measure theory and probability theory	Regular class **Homework 1 given**	No class
3: Jan 31 – Feb 4	Markov chains; stationary distributions	Regular class **Homework 1 deadline**	Regular class
4: Feb 7-11	Markov decision processes	Regular class	Regular class **Homework 2 given**
5: Feb 14-18	Dynamic programming; optimal control	Regular class	Regular class **Homework 2 deadline**
6: Feb 21-25		No class **Midterm during the week**	
7: Feb 28 – Mar 4	Partial observability; belief spaces	Regular class	Regular class **Homework 3 given**
8: Mar 7-11	Robust optimal control	Regular class	Regular class **Homework 3 deadline**
--: Mar 14-18		March break	
9: Mar 21-25	Exploration vs. exploitation, value of information	Regular class	Regular class
10: Mar 28 – Apr 1	Bandits	Regular class	Regular class **Homework 4 given**
11: Apr 4-8	Multi-agent planning	Regular class	Regular class **Homework 4 deadline**
12: Apr 11-15	Temporal logic	Regular class	Regular class
13: Apr 18-22	Automata and product spaces	Regular class	Regular class **Homework 5 given**
14: Apr 25-29	Aerospace examples	Regular class	Regular class **Homework 5 deadline**
15: May 2-4	Project discussion	Regular class	No class
--	**Final project deadline: May 9**		