

Course Description

Course Motivation: Broadly speaking, this is a course on decision making under uncertainty. Specifically, the course focuses on the study of stochastic linear systems where the dynamics have certain components of a stochastic nature. The course will cover both estimation and control of such systems. The first half of the course establishes the fundamentals of the estimation problem, culminating in the derivation of the fact that state estimation in linear systems is equivalent to projection onto a closed linear subspace generated by an observation process in a Hilbert space of random variables. This leads to the Kalman filter, which finds use in many applications ranging from aerospace to finance. The course will then cover the issues of stochastic optimal control (based on dynamic programming), control of Markov chains, and optimal stopping time problems (e.g., when is the optimal time to buy or sell a product?).

Calendar Description for ECE 686:

This is a course on continuous-parameter state estimation and control for stochastic linear systems. It is based on a single unifying theme, namely that state estimation in linear systems is equivalent to projection onto a closed linear subspace generated by an observation process in a Hilbert space of random variables. This formulation of state estimation leads to the innovations theorem of Kailath, and this in turn has a number of corollaries of considerable practical importance, such as the Kalman-Bucy filtering formulae and the Rauch-Tung-Striebel prediction formulae which are much used for example in problems of inertial guidance and control in aerospace, in stochastic optimal control, and (more recently) in econometrics.

Learning Outcomes

By the end of this course, students should be able to

- Apply probability theory to model stochastic systems
- Understand the Projection theorem and its application in estimation and filtering
- Develop estimators using least squares optimization
- Apply the Kalman filter for state estimation of linear systems
- Utilize dynamic programming to formulate optimal control problems
- Apply linear quadratic regulators (LQR) to control stochastic linear systems
- Understand Markov Chains and their corresponding optimal control problems

Instructor

Prof. Stephen L. Smith (stephen.smith@uwaterloo.ca).Office: E5 5112

- Please put “ECE686” at the front of your subject line for all course related email.

Tentative Course Schedule

The following is tentative list of topics that will be covered in the course.

1. Introduction and Mathematical Background
 - Motivation and overview
 - Review of probability theory and stochastic processes
2. Derivation of the Projection Theorem

- Metric and linear spaces
 - Hilbert spaces and the projection theorem for linear estimation
3. Estimation of Discrete-Time Linear Systems
 - Least-squares and recursive estimation
 - The discrete-time Kalman filter
 4. Stochastic Optimal Control
 - Dynamic programming and the principle of optimality
 - Optimal control with complete and partial observations
 - Linear quadratic regulators
 5. Control of Markov Chains
 - Overview of Markov chains
 - Markov policies and the cost of optimal policies
 6. Introduction to reinforcement learning

Text/Materials

There is no required textbook, notes will be provided each week with blanks that should be filled in during lecture. Parts of the course are based on notes by Prof. Andrew Heunis, and the following book:

1. P. R. Kumar and P. Varaiya, *Stochastic Systems: Estimation, Identification and Adaptive Control*, 1986.

The following textbooks may also be useful for additional information:

1. J. L. Speyer and W. H. Chung, *Stochastic Processes, Estimation and Control*, 2008
2. D. P. Bertsekas, *Dynamic Programming and Optimal Control*, 2005

Student Assessment

The course will consist of assignments and a project in which you will present a paper and implement their algorithm on filtering or control for stochastic systems, and a final exam.

- Assignments: 25%
- Project: 25% (proposal 5% and final presentation/code 20%)
- Final Exam: 50%

Assessment Details:

- Assignments will be submitted via crowdmark.
- Project details will be posted on LEARN. In the project you will select a research paper and/or topic that utilizes estimation or optimal control. You should implement the main algorithm/controller from the paper in MATLAB (or python) and have a running simulation that reproduces the main result. A proposal (worth 5%) will be due half-way through the term, and a the final presentation and code (worth 20%) will be due on the last day of term.
- The final exam is open book in the sense that you may consult your course notes and material posted in the course LEARN site and lecture videos. Use of any other resource (including file-sharing services such as chegg.com, coursehero.com, stack-exchange.com, etc.) is prohibited.

University Policy

Academic integrity: In order to maintain a culture of academic integrity, members of the University of Waterloo community are expected to promote honesty, trust, fairness, respect and responsibility. [Check the Office of Academic Integrity for more information.]

Grievance: A student who believes that a decision affecting some aspect of their university life has been unfair or unreasonable may have grounds for initiating a grievance. Read Policy 70, Student Petitions and Grievances, Section 4. When in doubt, please be certain to contact the department's administrative assistant who will provide further assistance.

Discipline: A student is expected to know what constitutes academic integrity to avoid committing an academic offence, and to take responsibility for their actions. [Check the Office of Academic Integrity for more information.] A student who is unsure whether an action constitutes an offence, or who needs help in learning how to avoid offences (e.g., plagiarism, cheating) or about "rules" for group work/collaboration should seek guidance from the course instructor, academic advisor, or the undergraduate associate dean. For information on categories of offences and types of penalties, students should refer to Policy 71, Student Discipline. For typical penalties, check Guidelines for the Assessment of Penalties.

Appeals: A decision made or penalty imposed under Policy 70, Student Petitions and Grievances (other than a petition) or Policy 71, Student Discipline may be appealed if there is a ground. A student who believes they have a ground for an appeal should refer to Policy 72, Student Appeals.

Note for students with disabilities: AccessAbility Services, located in Needles Hall, Room 1401, collaborates with all academic departments to arrange appropriate accommodations for students with disabilities without compromising the academic integrity of the curriculum. If you require academic accommodations to lessen the impact of your disability, please register with AccessAbility Services at the beginning of each academic term.

Turnitin.com: Text matching software (Turnitin) may be used to screen assignments in this course. Turnitin® is used to verify that all materials and sources in assignments are documented. Students' submissions are stored on a U.S. server, therefore students must be given an alternative (e.g., scaffolded assignment or annotated bibliography), if they are concerned about their privacy and/or security. Students will be given due notice, in the first week of the term and/or at the time assignment details are provided, about arrangements and alternatives for the use of Turnitin in this course.

It is the responsibility of the student to notify the instructor if they, in the first week of term or at the time assignment details are provided, wish to submit alternate assignment.