AER1517 Control for Robotics

Syllabus and Course Information, Winter 2020

Course Overview

This course presents optimal, adaptive and learning control principles from the perspective of robotics applications. It covers the foundations of optimal control and derives practical control algorithms for robotics. Working from the Hamilton-Jacobi-Bellman formulation, optimal control methods for aerial and ground robots are developed. Real-world challenges such as disturbances, state estimation errors and model errors are addressed, and adaptive and reinforcement learning approaches are derived to address these challenges. Course assignments involve, for example, the simulated control of an aerial vehicle, with aerodynamic models and wind disturbances.

Learning Objectives

At the end of the course, you will be able to:

- derive optimal controller equations for mobile robots,
- implement optimal controllers in practice and analyze their properties,
- understand state-of-the-art approaches to learning-based control and reinforcement learning,
- design learning-based controllers to cope with non-idealities such as model errors.

Instructor

- Prof. Angela Schoellig, University of Toronto Institute for Aerospace Studies (UTIAS), Room 187, e-mail: schoellig@utias.utoronto.ca, phone: 416-667-7518, web: http://schoellig.name

Teaching Assistants

- Siqi Zhou, UTIAS Room 191, e-mail: siqi.zhou@robotics.utias.utoronto.ca

Schedule

- Lecture: Mon 14h00–17h00, UTIAS Lecture Hall
- Office hours: after class on Mondays (please come to Prof. Schoellig right after class) and by appointment (please send an e-mail to Prof. Schoellig).

Website

We will use Quercus for course administration including announcements: http://q.utoronto.ca.

Reference Material

Recommended reference texts include:

- Bertsekas, *Dynamic Programming and Optimal Control,* Vol. I and II
- Stengel, *Optimal Control and Estimation*
- Sutton, Barto, *Reinforcement Learning: An Introduction*
- Bertsekas, *Reinforcement Learning and Optimal Control*
- Borrelli, Bemporad, Morari, P*redictive Control for Linear and Hybrid Systems*
- Ljung, *System Identification: Theory for User (2nd ed.)*

Lecture notes and slides will be made available on Quercus after each lecture, including links to references and relevant literature.

Course Topics and Syllabus

A tentative outline of the topics covered in class is provided below.

Lecture	Date	Title
Module 1: Optimal Control		
1	Jan 13	Introduction and Motivation
2	Jan 20	*Discrete-Time Optimal Control:* Principle of Optimality, Dynamic Programming, Solving the Bellman Equation
3	Jan 27	*Continuous-Time Optimal Control:* Hamilton-Jacobi-Bellman (HJB) Equation, Pontryagin's Minimum Principle
4	Feb 3	*Iterative Algorithms:* Linear Quadratic Regulator, Iterative Linear Quadratic Controller, Differential Dynamic Programming
5	Feb 10	*Model Predictive Control (MPC):* Introduction, Feasibility, Stability
	Feb 17	➢ *Reading Week* ,
Module 2: Reinforcement Learning		
6	Feb 24	*Introduction to Reinforcement Learning:* Markov Decision Processes, State Value Function, Action Value Function, Policy Evaluation
7	Mar 2	*Model-Based Reinforcement Learning:* Policy Iteration, Value Iteration
8	Mar 9	*Model-Free (or Sample-Based) Reinforcement Learning:* Monte Carlo, Q Learning, Experience Replay
9	Mar 16	*Path Integral Optimal Control*
10	Mar 23	*Representation Learning:* Policy Gradient, Unsupervised and Self-Supervised Learning, Deep Reinforcement Learning
11	Mar 30	*Learning System Models:* System Identification, Regression, Gaussian Processes
12	Apr 6	Summary and Outlook

Assignments

There will be four assignments throughout this course. Three assignments will involve writing Matlab code to implement various control algorithms and test them in simulation. The programming assignments will be released at least 2 weeks prior to the deadlines and will be due on the dates noted below at midnight.

The fourth assignment will be a presentation in class. For this assignment, you can work in *groups of up to two students*. The main goal of the presentation assignment is to enable the inclusion of more recent methods into the course. You will review a specific topic or paper and present results of the method applied to a sample robotics problem. Presentations will be 15 minutes long, starting in the fourth week. Presentation slots will be assigned based on topic selected, so as to align with the lecture of the week.

The due dates for the programming assignments are as follows:

Assignment	Due Date
1	Monday, Feb 24
2	Monday, Mar 23
3	Monday, Apr 6

Project

The project will provide a chance to apply the methods taught in this course to a problem of your choosing within your own field. A two-page proposal is required after 6 weeks of the course. The mark will be based on a demonstration video and a final report. The deadlines for the project are as follows:

Component	Due Date
Proposal	Monday, Mar 2
Demonstration Video	Monday, Apr 27
Final Report	Monday, Apr 27

Evaluation

The grading scheme is: three programming assignments ($3 \times 10\% = 30\%$), one presentation assignment (15%), and one project (15% video $+ 40\%$ report $= 55\%$).