ECSE 551 -- Machine Learning for Engineers
Department of Electrical and Computer Engineering
McGill University

Instructor

Prof. Narges Armanfard
Room: MC 623
Email: narges.armanfard@mcgill.ca
Office Hours: 1 hour per week

Lectures

3 hours per week. For more details, see McGill Class Schedule.
Number of credits: 4.

Tutorials

2 hours per week. For more details, see McGill Class Schedule.
Focus: Implementation of machine learning systems in Python (Numpy, Pandas, Scikit-Learn, Pytorch).

Teaching Assistants

TAs name and contact info are posted on the course homepage on myCourses.
Office Hours: 1-2 hours per week.

Communication

The course homepage on myCourses is used as the primary means of communication. Course materials, notes, etc. are distributed via the course homepage. Please consult the course homepage regularly.

Course topics and Schedule (tentative)

Introduction to machine learning (Lectures 1-2); Linear regression (Lectures 3-5); Linear Classification (6-8); Evaluation, Bias, Variance, Regularization (Lectures 9-10); Decision trees (Lecture 11); Feature construction, Dimension reduction (Lecture 12); Instance-based learning (Lecture 13-15); Ensemble methods (Lecture 16); Neural networks and deep learning (Lectures 18-22); Unsupervised learning (Lectures 23-24).

Pre(co)requisite

Corequisite: ECSE 443 or ECSE 543 or MATH 247.
Prerequisite(s): COMP 250 and ECSE 205 or MATH 323.
Restrictions: Not open to students who have taken COMP 551.

Evaluation (tentative)

10% quizzes,
45% assignments (projects),
40% midterm,
5% in-class activities.

Assignments

There are 3-4 projects. Projects are designed to be completed in groups of 2-3 students. Students can find their partner(s) in the Discussions section of the course homepage on myCourses. There is Kaggle competition for some of the projects.

Textbook

Shai Shalev-Shwartz and Shai Ben-David. Understanding Machine Learning: From Theory to Algorithms. Cambridge University Press. 2014.

Trevor Hastie, Robert Tibshirani and Jerome Friedman. The Elements of Statistical Learning: Data Mining, Inference, and Prediction, 2nd Edition. Springer. 2009.

Christopher Bishop. Pattern Recognition and Machine Learning. Springer. 2007.

Ian Goodfellow, Yoshua Bengio and Aaron Courville. Deep Learning. The MIT Press.