

Advanced Topics in Machine Learning

COCS 89/189, Spring 2016

Instructor: Qiang Liu

Full description This is an advanced research topic course for students that have already taken COCS74/174 or an equivalent machine learning class, and have a strong math background. It will give an in-depth coverage of currently active research areas in machine learning. The course will connect to open research questions in machine learning, giving starting points for future work. Main topics include probabilistic graphical models, inference and learning, Markov chain Monte Carlo, deep networks, structured output prediction, advanced optimization methods, nonparametric and kernel methods, as well as related application areas. The students will read and present recently published papers, propose and complete a research-oriented term project. In addition, each student is responsible to scribe one lecture note in LaTeX. There is no midterm or final exam. The registration is by the instructor's permission only. Please send email to qliu@cs.dartmouth.edu. In addition, there will be an entrance exam in the first week.

Prerequisites COS74/174

Coursework The final course grade will be based 20% on in-class participation, 20% on the paper presentations, 20% on the scribe notes, and 40% on the term project.

Required Textbook

- *Kevin Murphy*, Machine learning: A Probabilistic Perspective

Optional Textbook

- *Daphne Koller, Nir Friedman*, Probabilistic Graphical Models: Principles and Techniques
- *David Barber* Bayesian Reasoning and Machine Learning
- *Sergios Theodoridis* Machine Learning: A Bayesian and Optimization Perspective

Time TuTh 4:00-5:50; x-period: M 5:00-5:50

Topics to be covered (tentative)

- Week1
 - L1: Intro to graphical models, learning and inference
 - L2: Variational inference
 - Reading: Murphy Section 19, 22

- Week 2
 - L1: Markov chain Monte Carlo
 - L2: Paper presentations
 - Reading: Murphy Section 24

- Week 3
 - L1: Learning in graphical models
 - L2: Paper presentations
 - Reading: Murphy Section 26

- Week 4
 - L1: Project proposal presentations
 - L2: Structured output prediction

- Week 5
 - L1: Deep generative networks
 - L2: Paper presentations
 - Reading: Murphy Section 28

- Week 6
 - L1: Neural networks, non-convex optimization
 - L2: Paper presentations
 - Reading: Murphy Section 16.5

- Week 9
 - L1: Deep reinforcement learning
 - L2: Paper presentations

- Week 7
 - L1: Nonparametric and kernel methods, reproducing kernel Hilbert spaces

L2: Project milestone presentations
Reading: Murphy Section 14

- Week 8
L1: Gaussian process, Bayesian optimization, Bayesian quadrature
L2: Paper presentations
Reading: Murphy Section 15
- Week 9
L1: Crowdsourcing in machine learning
L2: Paper presentations
- Week 10
Final project presentations