CS5560: Probabilistic Models for ML

Overview

This is a foundational course in machine learning with focus on probabilistic models. The course is intended for those who wish to pursue research in machine learning and related fields. Hence, the developments are technical and involve appropriate mathematical details. Broadly, the course includes studying basic probabilistic models that are often used in machine learning set-ups, and methods of estimating parameters of such models from data. The course ends with an introduction to a very rich class models known as probabilistic graphical models[1].

Pre-requisites

The course assumes basic engineering level knowledge of probability theory, and multivariate calculus. Familiarity with machine learning, statistics, and mathematical optimization, will aid better appreciation of the course, but are not necessary pre-requisites.

Syllabus and Text

The textbook for this course is "Machine Learning: A Probabilistic Perspective" By Kevin P. Murphy (MIT Press). The detailed syllabus is:

1. Maximum likelihood estimation (MLE) in Binomial, Multinomial, Gaussian, models in exponential family.
2. MAP, Bayesian estimation in Beta-Binomial, Dirichlet-Multinomial, Normal-Inverse-Gamma-Gaussian, conjugate prior based models in exponential family.
3. Learning with Supervised Models: Generative and discriminant models, Linear Regression, Logistic Regression, Gaussian Discriminant Analysis, Generalized linear models.

[1] In the version of the course taught to EMDS/MDS students this advanced topic on Graphical Models may not be covered.

4. Expectation Maximization (EM) based learning in Mixture models, Hidden Markov Model, Dirichlet processes (Clustering).

5. Introduction to directed (Bayes nets) and un-directed (Markov Random Fields) graphical models.

Evaluation Scheme

Date	Duration	Percentage
31-08-2018 (Fri)	60 min.	15%
23-10-2018 (Tue)	90 min.	35%
23-11-2018 (Fri)	180 min.	50%

In addition to the above, there will be weekly assignments. In case these are not submitted as per the instructions, two negative marks will be awarded (per default). Some assignments will carry bonus marks and will be used for grade-promotions during the final grading.

In case you are auditing this course, then 100% attendance is the only requirement for passing the course.

For EMDS/MDS students all instructions remain same, except that there will be only one exam (the one on the last row of the above table) carrying 100% marks.

Contact

You are welcome to drop-by my office (C-519) anytime during my regular office hours (8:30am-12:30pm;2:30pm-5:00pm). Though you are not required to take prior permission, in case you don't notify me, I may have stepped out for some meeting/break when you visit. You are highly encouraged to ask questions/clarifications either in lecture or when you visit my office.

All correspondences in this course (including assignment submissions) will be done via Google-Classroom (code will be announced in lecture). This forum can also be used for clarifications/questions.

References by Topic

Here is list of sections in your textbook (Murphy's book) for each topic covered in the lectures:

Method of Moments: Simple method with no elaborate reference in your textbook. Defined in equation (9.47).

Maximum Likelihood: Defined in equation (3.7). Some details in section 6.2.2.

Gaussian-MLE: Section 4.1.3.

Multinoulli-MLE: Again, easy! So not explicitly covered. Covered as special case in section 3.2.

Generative-LinearRegression: Section 4.3.

Discriminative-LinearRegression: Sections 7.1-7.3.

GDA (BayesClassifier): Section 4.2.

Logistic Regression: Sections 8.1-8.3.

Exponential Family: Sections 9.1-9.2.

GLMs: Section 9.3.

Bayesian Inference and MAP: Sections 3.2.2-3.2.4, 3.3.2-3.3.4, 3.4.2-3.4.4, 3.5.1.2, 4.6, 7.5.1, 7.6 (non-*), 8.4 (light reading), 9.2.5.

Hierarchical Bayes: Sections 5.5, 5.6.

CrossValidation: Section 7.10 in https://web.stanford.edu/ hastie/Papers/ESLII.pdf. Why it works is because of Lemma1 in www.kyb.mpg.de/publications/pdfs/pdf1436.pdf.

Model Selection (Bayesian): Sections 5.3, 5.5, 5.6.

Mixture Models: Section 11.2, 11.3.

EM for MixtureModels: Sections 11.4.1-11.4.2.

Markov Models: Section 17.2.

HMMs: Sections 17.3, 17.4, 17.5.

Bayes Nets: Entire chapter 10 except * sections. Additional reading: Chapter 3 except sections 3.3.3, 3.4.2, 3.4.3 in DaphneKoller's book.

Markov Nets: Sections 19.1-19.4. Additional reading: CHapter 4 except sections 4.4, 4.5, 4.6 in DaphneKoller's book.