CS5560: Probabilistic Models for ML

Overview

This is an intermediate-level course in machine learning that exclusively studies (parametric) probabilistic models. The main topics covered are:

Basic: Exponential family based models, MLE/Bayesian parameter estimation, loss-based/Bayesian inference, Sampling techniques.

Advanced: Models with latent variables, Implicit generative models, Probabilistic graphical models.

The lectures and assignments focus on the technical aspects of the topics and provided sufficient details as desirable for an intermediate-level course.

Pre-requisites

The course assumes undergraduate level knowledge of probability theory,and multivariate calculus. It also assumes familiarity with machine learning models, in general. Further, expertise in statistics, and mathematical optimization, are desirable, but are not necessary pre-requisites.

Books

Books by Kevin P. Murphy are good for the basic topics and latent variables (`https://probml.github.io/pml-book/`). The book by Daphne Koller and Nir Friedman is a masterpiece on graphical models. Relevant papers will suffice for implicit models.

Evaluation Scheme

For regular students, there will be two Quizes and one end sem:

S.No.	Exam	Weightage	Date
1	Quiz-1	25%	11-Feb-2021 6:30pm-8pm
2	Quiz-2	25%	15-Mar-2021 6:30pm-8pm
3	EndSem	50%	5-May-2021 10am-1pm

For MDS students, there will be one quiz every weekend with equal weightage. Suppose 10 quizzes are conducted, best 8 of 10 are taken for final score.

Contact

Apart from regular 3 lecture hours, a 1hr tutorial will be conducted every week (online). Attending the tutorial is completely optional. In the tutorial, doubts from students shall be clarified etc.