

STA 380: Bayesian Methods for Machine Learning

Spring 2018, Thursday 1:00 - 4:00 PM, Room: CBA 6.420

Instructor: Mingyuan Zhou, Ph.D., Assistant Professor of Statistics

Office: CBA 6.458 (east side of the building that faces the entrance of Gregory Gym)

Phone: 512-232-6763

Email: mingyuan.zhou@mcombs.utexas.edu

Website: <http://mingyuanzhou.github.io/>

Office Hours: Thursday 4:00-5:30 PM

You are welcome to come by my office at other times. To make sure that I will be there then, you may first call my office or send me an email.

Course Description:

This Ph.D.-level course will discuss Bayesian statistical methods and their applications to machine learning. We will learn not only how to construct hierarchical probabilistic models for various types of data, but also how to perform Bayesian computation via both Markov chain Monte Carlo and variational Bayesian inference. We will study a number of representative Bayesian hierarchical models that are useful for data analysis problems in econometrics, risk analysis, finance, decision making, marketing, information systems, and physical and social sciences. Example topics include discrete choice models, latent variable models, count data analysis, data augmentation, Bayesian nonparametrics, scalable inference for big data, and deep neural networks. Example applications include missing-data imputation, nonlinear classification, survival analysis, text analysis, collaborative filtering, recommendation systems, market basket analysis, and social network analysis. Some recent progresses in Bayesian deep learning, including deep hierarchical models, variational auto-encoder, generative adversarial networks, and implicit models, will also be discussed.

There is no official prerequisite for the class, but it is expected that a student either has already had or is willing to work hard along the semester to build a solid foundation in probability and statistics.

Materials:

- Three textbooks are recommended but not required:
 - [1] Pattern Recognition and Machine Learning, by Christopher Bishop.
 - [2] Machine Learning: a Probabilistic Perspective, by Kevin Murphy.
 - [3] Deep Learning, by Ian Goodfellow, Yoshua Bengio, and Aaron Courville.
- Software:

Python, R, or Matlab

Tentative Course Schedule:

This schedule represents my current plans and objectives. As we go through the semester, those plans may need to change to enhance the class learning opportunity. Such changes, communicated clearly, are not unusual and should be expected.

Week 1 January 18: Likelihood, conjugate priors, posterior, Bayes' rule

Week 2 January 25: Markov chain Monte Carlo and variational Bayesian inference

Week 3 February 1: Probit regression, logistic regression, discrete choice models

Week 4 February 8: Bayesian sparse regression and classification

Week 5 February 15: Bayesian dictionary learning and sparse coding, matrix and tensor factorization, collaborative filtering, recommendation systems

Week 6 February 22: Topic models, mixed-membership modeling, blocked and collapsed Gibbs sampling

Week 7 March 1: Poisson factor analysis, count data analysis, data augmentation

Week 8 March 8: Scalable Bayesian inference for big data, stochastic variational inference, stochastic gradient MCMC

Week 9 March 22: Bayesian models for network analysis, stochastic blockmodel, edge partition model

Week 10 March 29: Bayesian nonparametrics (Poisson process, gamma process, Dirichlet process, Chinese restaurant process, beta process, Indian buffet process)

Week 11 April 5: Bayesian deep learning

Week 12 April 12: Bayesian deep learning

Week 13 April 19: Project presentations

Week 14 April 26: Project presentations

Week 15 May 3: Project presentations