Statistics 551, Winter 2022: Bayesian statistics

Instructor: Long Nguyen
Department of Statistics, Univ. of Michigan
Time: 2:30–4pm, 1230 USB

Instructors. Prof. Long Nguyen, xuanlong@umich.edu, OH: tbd.
The graduate student instructors (GSIs) for this course:

- Yangyi Lu, yylu@umich.edu, OH: tbd

Course description. This course provides basic concepts and several modern techniques of Bayesian modeling and computation. They include basic models, conjugate priors, and posterior computation, as well as techniques associated with complex models, such as hierarchical models, spatiotemporal models, and dynamical models. A substantial part of the course is devoted to computational algorithms based on Markov Chain Monte Carlo sampling for complex models. If time permits, we will also introduce advanced topics such as nonparametric Bayes, variational inference, and Hamiltonian Monte Carlo techniques. Foundational topics will be discussed when appropriate, although they are not our primary focus in this course; such topics may include decision theoretic characterization of Bayesian inference and its relation to frequentist methods, de Finetti-type theorems and the existence of priors, objective prior distributions, and Bayesian model selection.

Prerequisites. The prerequisites are previous coursework in linear algebra, multivariate calculus, and basic probability and statistics. Previous experience in numerical analysis and optimization would be helpful but is not required. Familiarity with R, Python, Matlab, or a related programming language will be necessary.

Structure/Evaluation. The course will meet twice a week and will follow a regular lecture format. There will be three to four homework assignments, due approximately one week after being passed out. The homework assignments will be solved by three-person teams; the team assignment will be randomly generated each time. There will be a final data analysis project, done in three-person teams. The project component consists of a short presentation and a project report.

Homework assignments account for %60, and project %40 of the final grade.

Course homepage. Please use Canvas for all announcements, homeworks, project information and data sets.

Textbook. Lecture notes and other relevant reading materials will be provided via Canvas. The following references will be useful:

- Peter Hoff's "A first course in Bayesian statistical methods".
- Christian Robert's "The Bayesian choice".
- Michael Jordan's "Introduction to probabilistic graphical models".
- C. Bishop's "Pattern recognition and machine learning".

Tentative outline.

1. Introduction and examples
2. Interpretation of probabilities and Bayes rules
3. One-parameter models
4. Monte Carlo approximation
5. The multivariate normal
6. Linear regression
7. Posterior computation via MCMC
8. Exponential families and conjugate priors
9. Group comparison and hierarchical modeling
10. Nonconjugate priors and Metropolis-Hasting algorithms
11. Linear and generalized linear mixed effects models
12. Bayesian classification
13. Unsupervised learning via latent variable models
14. Optional topics (time permitting):
    - nonparametric Bayesian methods
    - Hamiltonian MCMC methods
    - variational Bayesian inference