CS 6190: Probabilistic Machine Learning

Spring 2024

Instructor:	Shandian Zhe	**Time:**	Tue & Thu 03:40PM-05:00PM
Email:	zhe@cs.utah.edu	**Place:**	WEB 2230

Course Page:

- `https://www.cs.utah.edu/~zhe/teach/cs6190.html`

Teaching Mentee:

- Su-wei Yang (u1429034@utah.edu)

Office Hours:

- Instructor: Tue/Thu 12:30-1:30pm, MEB 3466
- Su-wei Yang: Mon/Wed 10:00-11:00am, Friday 9:00-11:00am (MEB 3115)

Description:
The course introduces basic knowledge of probabilistic modeling and learning. Topics cover fundamental concepts of Bayesian statistics, probabilistic graphic models, generalized linear models, approximate inference (including variational inference, message passing and Markov-Chain Monte-Carlo), Bayesian (deep) neural networks, Gaussian process regression, etc. After taking this class, we expect that you will (1) *understand the principles and paradigms of probabilistic learning*, (2) *be able to explore relevant literature, exploit existing and/or create new probabilistic modeling/learning tools for your own research interests*, and (2) *be well prepared to dive into the cutting-edge research in probabilistic machine learning.*

Warning: This course is math intensive and requires a certain level of programming capabilities (with Matlab, R or Python). Python components may require PyTorch/JAX. The coding workload is not heavy, but requires mathematical derivations, especially in linear algebra and careful debugging.

Books: The major reference textbook for this course is *Pattern Recognition and Machine Learning* by Christopher Bishop, Springer, 2007. The book is free online. While the lecture slides will cover all the content, the students are encouraged to read through the corresponding chapters. There can be a few topics not covered by the reference book. For these topics, we will provide extra reading materials. In addition, we list several books to further extend the depth and breadth of the topics we will discuss in the class.

- Kevin Patrick Murphy, Machine Learning: a Probabilistic Perspective. MIT Press, 2012.
- David J.C. MacKay, Information Theory, Inference, and Learning Algorithms. Cambridge University Press, 2003.
- Larry Wasserman, All of Statistics: A Concise Course in Statistical Inference. Springer, 2004.
- Sidney I. Resnick, A Probability Path, Springer, 2014.
- Daphne Koller and Nir Friedman, Probabilistic Graphical Models: Principles and Techniques, MIT Press, 2009

Prerequisites: "C-" or better in (CS 3500 AND (MATH 2270 OR MATH 2250) AND (CS 3130 OR ECE 3530)) AND (Full Major status in Computer Science OR Computer Engineering). Corequisites: "C-" or better in (CS 4150 OR CS 3100). Basically, we assume that you

- know basics of calculus and statistics,
- are familiar with **linear algebra**, know vector/matrix derivatives,
- have algorithmic design and programming skills.

Getting Help: Take advantage of the instructor and TA office hours (posted on course web page). We will work hard to be accessible to students. Please send us emails if you need to meet outside of office hours. Don't be shy if you don't understand something: come to office hours, send emails, or speak up in class! Students are encouraged to use a discussion group for additional questions outside of class and office hours. The class will rely on the **Canvas discussion group**. Feel free to post questions regarding any questions related to class: homeworks, schedule, material covered in class. Also feel free to answer questions, the instructors and TAs will also actively be answering questions. But, do not post **potential homework answers**. Such posts will be immediately removed, and not answered. All important announcements will be made through the discussion group, there is otherwise no class mailing list.

Tentative Course Outline:

- Basic concepts and Bayesian statistics
  - Probability space, random variables, CDF, PDF, expectation, variance, independence, etc.
  - Probability distributions, e.g., (Multivariate) Gaussian distribution, student t-distribution, Beta distribution, Gamma and inverse Gamma distributions, Dirichlet distribution, etc.
  - Maximum likelihood estimation (MLE), maximum a posterior estimation (MAP), predictive distribution, type II MLE, empirical Bayes
  - Bayesian decision theory, Bayesian model selection
  - Noninformative priors, de Finetti's theorem, Bayesian philosophy
  - Exponential family and conjugate priors
- Generalized linear models
  - Bayesian linear regression, logistic regression and probit regression
  - Multi-class logistic regression and ordinal regression
  - Generalized linear models in exponential family
- Probabilistic graphical models
  - Bayes networks and Markov random fields
  - Conditional independence, Markov blanket, Bayes ball algorithm
  - Factor-graphs
  - Message passing, sum-product, and max-sum algorithms
- Approximate inference
  - Laplace approximation
  - Variational inference, EM algorithm, variational model evidence lower bound, mean-field, local variational inference, convex conjugate, variational message passing

- Markov Chain Monte-Carlo algorithms: Metropolis Hasting, Gibbs sampling, Hamiltonian Monte Carlo

- Bayesian neural networks
  - Bayes back-propagation
  - Variational auto-encoder
  - Reparameterization tricks, concrete distribution
  - Diffusion Modeling
- Gaussian process
  - Gaussian process regression and classification
  - Sparse Gaussian processes
  - Deep Gaussian processes

Grading Policy:

- Homeworks (60%): 6 homeworks. The homework assignment can include both analytical problems and programming assignments. The programming assignments usually request you to implement particular probabilistic learning models and test them in real-world/synthetic datasets. Each homework might have a different number of points, depending on the workload. *You can only use MATLAB, Python or R for the programming portion of the assignments or projects. Other programming languages are NOT accepted.* Some programming tasks may ask you to use PyTorch or Jax. That means, you have to use Python for those tasks. In your program assignments, you are free to use existing libraries (e.g., Numpy and Scipy in Python) to finish linear algebra computation and optimization (unless the homework says it is not allowed). *Your are never allowed to call APIs (e.g., scikit-learn/PyMC3/PyRO) that directly implements the required models/algorithms.*
- Course project (40%): Please use probabilistic learning techniques to address one research or practical problem of your own interest. You are requested to create a Github repository and submit mid-term and final reports (3-6 pages, font size less than or equal to 11). For details, please refer to the course web page `https://www.cs.utah.edu/~zhe/teach/cs6190-proj`
- Final exam (0%): There is NO final exam.

Letter Grade Mapping: We do not curve your grade.

Table 1: CS6190

91-100	A	74-78	B	59-63	C	44-49	D
85-90	A-	69-73	B-	54-58	C-	39-43	D-
79-84	B+	64-68	C+	49-54	D+	0-38	E

We will round up your score to the closest integer values. For example, 90.5-90.9 will be treated as 91, and 90.1-90.4 will be treated as 90.

Late Policy: Each assignment must be turned in through Canvas by the designated deadline, usually 11:59pm, to receive the full credit. If the deadline is missed, the late submission will have 10% penalty.

Then late submissions in every subsequent 24 hours will lose another 10% credit. For example, a 10 points assignment will have 2 points penalty, if it is submitted 30 hours late. However, if the assignment is not turned in until the other assignment have been graded and returned or 48 hours after the deadline, 0 grade will be given.

Assignments will be posted far enough ahead of time that I will not be able to make exceptions if a student falls ill. The exception will be prolonged illness accompanied by a doctor's note.

If you believe there is an error in grading homeworks, you may request a regrading within **one week** of receiving your grade. Requests must be made by email to instructor, explaining clearly why you think your solution is correct.

Homework Submission: We only accept homeworks written with LaTex. It is something that everyone should know for research and writing scientific documents. This linked directory (`http://www.cs.utah.edu/~jeffp/teaching/latex/`) contains a sample .tex file, as well as what its .pdf compiled outcome looks like. It also has a figure .pdf to show how to include figures. Overleaf (`https://www.overleaf.com/project`) is an excellent web editor for Latex documents. We encourage everyone to use overleaf. We will release the template tex file along with each homework assignment for convenient editing.

Academic Policy: This course follows School of Computing (SoC) Polices/Guidelines (`https://handbook.cs.utah.edu/`). Please read it carefully. If a student is caught cheating on homework or project, they will receive a failing grade for the course. For a detailed description of the university policy on cheating, please see the University of Utah Student Code: `https://regulations.utah.edu/academics/6-410.php`.

Collaboration policy: For assignments, students can discuss answers with anyone, but must write their own code, proofs and write-ups. If you collaborated with another student such that you expect the answers may look similar, you must explain explicitly in the homework submission to what extent you have collaborated. Students' homework submissions appearing too similar to others and without collaboration explanation will receive 0 grades.

Students with Disabilities

Please let me know at your earliest convenience. The University of Utah seeks to provide equal access to its programs, services, and activities for people with disabilities. If you need accommodations in this class, reasonable prior notice needs to be given to the Center for Disability Services, 162 Olpin Union Building, 581-5020 (V/TDD). CDS will work with you and the instructor to make arrangements for accommodations.

Safety

The University of Utah values the safety of all campus community members. To report suspicious activity or to request a courtesy escort, call campus police at 801-585-COPS (801-585-2677). You will receive important emergency alerts and safety messages regarding campus safety via text message. For more information regarding safety and to view available training resources, including helpful videos, visit safeu.utah.edu