CIS 5930: Probabilistic Machine Learning

Spring 2025

Instructor:	Shibo Li	**Time:**	Mod & Wed 6:35pm - 7:50pm
Email:	shiboli@cs.fsu.edu	**Place:**	LOV 307

1 Course Information

Course Pages:

- https://cis5930.github.io/

Mode of Delivery

Lectures will meet in person at the specified time and location. Students should attend the lecture for the section they are enrolled in. This class is in-person only.

Office Hours

- Mon 3:30pm - 5:30pm, LOV 206A

Objective:

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. under-stand 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
3. 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 TensorFlow and/or PyTorch. The coding workload is not heavy, but requires mathematical derivations, especially in linear algebra and careful debugging.

Prerequisites:

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.

2 Getting Help

Take advantage of the instructor 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 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.

3 Course Materials

Textbooks:

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

Linear Algebra Resources

- Old and New Matrix Algebra Useful for Statistics from Thomas P. Minka
- The Matrix Cookbook from Kaare Brandt Petersen and Michael Syskind Pedersen
- Linear Algebra Review and Reference from Stanford
- The Linear Algebra chapter in the text book on deep learning by Goodfellow, Bengio and Courville.
- Linear Algebra lectures, MIT
- Review of some elements of linear algebra, by Fernando Paganini

Probability and Statistics Resources

- David Blei's review of probability
- Probability for data miners, slides by Andrew Moore
- Review of basic concepts in probability by Padhraic Smyth.
- A review of probability theory from Stanford
- The Probability and information theory chapter in Yoshua Bengio's upcoming book on deep learning.

4 Tentative Course Schedule

This is a tentative schedule for the course. Please see the weekly calendar on course website https://cis5930.github.io/for details and updates.

1. 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, Betadistribution, Gamma and inverse Gamma distributions, Dirichlet distribution, etc.– Maximum likelihood estimation (MLE), maximum a posterior estimation (MAP), predictive dis-tribution, type II MLE, empirical Bayes
   - Bayesian decision theory, Bayesian model selection
   - Exponential family and conjugate priors
2. Generalized linear models
   - Bayesian linear regression, logistic regression and probit regression
   - Multi-class logistic regression and ordinal regression
   - Generalized linear models in exponential family
3. 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
4. Approximate inference
   - Laplace approximation
   - Variational inference, EM algorithm, variational model evidence lower bound, mean-field, local-variational inference, convex conjugate, variational message passing
   - Expectation propagation, assumed density filtering
   - Markov Chain Monte-Carlo algorithms: Metropolis Hasting, Gibbs sampling, Hamiltonian Monte Carloe
5. Bayesian neural networks
   - Bayes back-propagation
   - Variational auto-encoder
   - Reparameterization tricks, concrete distribution – General adversarial neural networks
6. Probabilistic generative models
   - Diffusion models
   - Flow-based models
7. Nonparametric Bayes: Gaussian Process
   - Gaussian process regression and classification
   - Sparse Gaussian processes
   - Deep Gaussian processes

5 Grading Policy

Assignments - 90% for total of 5 assignments

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 TensorFlow. **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) that directly implements the required models/algorithms.**

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.

Assignments **must be electronically submitted through Canvas by midnight of the due date.** Instructions about submission will be given in each assignment. **Hand written versions or scans will *not* be accepted.**

Attendance - 10%

Attendance and participation are expected and REQUIRED to do well in the course. Attendance translates to staying current with the course. Attendance and class participation, measured through the in-class exercises, will count towards the In-Class Exercise component of the course grade.

Backup Copies of Your Programming Assignments

Always make multiple backup copies (on a flash drive, in the university's file space, in your private cloud space, private repository on Github ...) of your work!!! This is a course requirement, professional convention and good common sense. The teaching staff may under certain circumstances have to ask you to produce your backup file copies.

Regrade Requests

Requests for regrading should be ***within a week*** of grades being posted on Canvas by sending an email to the instructor. Only the file already submitted on Canvas. Students are not allowed to submit a newer version of their programs.

Letter Grade Mapping:

We do not curve your grade. Use table in the following to map your numerical final grade to letter grade.

Late Submission

All assignments should be submitted by the deadline. If the deadline is missed, the late submissions will have 10% penalty. In every subsequent 24 hours, the late submissions will loose another 10% credicts. For

A	>=90
A-	[75, 90)
B+	[60, 75)
B	<60

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 within 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.

6 Vital Policies

University Attendance Policy

Excused absences include documented illness, deaths in the family and other documented crises, call toactive military duty or jury duty, religious holy days, and official University activities. These absences will be accommodated in a way that does not arbitrarily penalize students who have a valid excuse. Consideration will also be given to students whose dependent children experience serious illness.

Academic Honor Policy

The Florida State University Academic Honor Policy outlines the University's expectations for the integrityof students' academic work, the procedures for resolving alleged violations of those expectations, and the rights and responsibilities of students and faculty members throughout the process. Students are responsible for reading the Academic Honor Policy and for living up to their pledge to "...be honest and truthfuland...[to] strive for personal and institutional integrity at Florida State University." (Florida State UniversityAcademic Honor Policy, found at https://fda.fsu.edu/academic-resources/academic-integrity-and-grievances/academic-honor-policy)

Academic Success

Your academic success is a top priority for Florida State University. University resources to help you succeed include tutoring centers, computer labs, counseling and health services, and services for designated groups, such as veterans and students with disabilities. The following information is not exhaustive, soplease check with your advisor or the Department of Student Support and Transitions to learn more.

Americans With Disabilities Act

Florida State University (FSU) values diversity and inclusion; we are committed to a climate of mutual respect and full participation. Our goal is to create learning environments that are usable, equitable, inclusive, and welcoming. FSU is committed to providing reasonable accommodations for all persons with disabilities in a manner that is consistent with academic standards of thecourse while empowering the student to meet integral requirements of the course.

Students with disabilities needing academic accommodation should:

1. register with and provide documentation to the Office of Accessibility Services; and

2. request a letter from the Office of Accessibility Services to be sent to the instructor indicatingthe need for accommodation and what type; and,

3. meet (in person, via phone, email, skype, zoom, etc...) with each instructor to whom a letter of accommodation was sent to review approved accommodations.

Please note that instructors are not allowed to provide classroom accommodations to a student until appropriate verification from the Office of Accessibility Services has been provided.

This syllabus and other class materials are available in alternative format upon request.

For the latest version of this statement and more information about services available to FSU students with disabilities, contact the:

```
Office of Accessibility Services
874 Traditions Way
108 Student Services Building
Florida State University
Tallahassee, FL 32306-4167
(850) 644-9566 (voice)
(850) 644-8504 (TDD)
oas@fsu.edu
https://dsst.fsu.edu/oas
```

Instructor's note on course exams: any requests for specific special exam arrangements due to aregistered disability must be brought to the course instructor at least two weeks prior to the examdate, or they will not be considered. In addition students must follow all rules and procedures setforth by the FSU OAS.

FSU Student Conduct Code

Students are expected to follow the FSU Student Conduct Code in all interactions and situationsat the university. See https://dos.fsu.edu/srr/conduct-codes/student-conduct-codes.

Confidential Campus Resources

Various centers and programs are available to assist students with navigating stressors thatmight impact academic success. These include the following:

- Victim Advocate Program

  ```
  University Center A
  Room 4100
  (850)644-7161 (Available 24/7/365)
  Office Hours:  M-F 8 am-5 pm
  https://dsst.fsu.edu/vap
  ```

- Counseling and Psychological Services

  ```
  Askew Student Life Center, 2nd Floor
  942 Learning Way
  (850)644-8255
  Office Hours:  M-F 8 am-5 pm
  https://counseling.fsu.edu/
  ```

- Counseling and Psychological Services

```
Health and Wellness Center
(850)644-6230
Office Hours:  M-F 8 am-5 pm
https://uhs.fsu.edu/
```

Syllabus Change Policy

Except for changes that substantially affect implementation of the evaluation (grading) statement, this syllabus is a guide for the course and is subject to change with appropriate notice.