USC Viterbi
School of Engineering

CSCI 699: Probabilistic and Generative Models

Units: 4.0
Spring 2025 — Fridays — 2:00-5:20PM

Location: [TBD]

Instructor: Willie Neiswanger
Office: Powell Hall of Engineering (PHE), Room 216
Office Hours: Friday 11-1pm and 5:30-6:30pm
Contact Info: Email to neiswang@usc.edu, please include "[class]" in the subject line.

Catalogue Description

Probabilistic and generative models, including approximate inference algorithms (MCMC, variational inference), deep generative models (autoregressive, score-matching, diffusion, and flow-based models), and model-based sequential decision making.

Course Description

This course focuses on the foundational principles of probability and its central role within modern machine learning and generative modeling. With probability increasingly driving advancements in AI, this course will explore its applications across a range of topics. From approximate inference algorithms, to generative models via large-scale self-supervised learning, to probabilistic-model based decision-making methods, you will gain a deeper understanding of how these methods shape contemporary AI research.

Learning Objectives

Through this course, you will be introduced to core topics in probabilistic machine learning, including probabilistic graphical models and approximate inference algorithms (e.g., MCMC and variational inference), deep generative models (e.g., autoregressive, score-matching, diffusion, and flow-based approaches), and methods for probabilistic model-based sequential decision-making (e.g., Bayesian optimization, and information-based experimental design). By the end of this course, you will have gained familiarity with both cutting-edge research and the historical foundations of the field.

Recommended Preparation

This course is designed for students currently pursuing research, or who wish to pursue research, in probabilistic machine learning or deep generative models. It will be beneficial to have familiarity with machine learning (at the level of CSCI 567), algorithms (at the level of CSCI 570), and probability (at the level of MATH 505a). Students are expected to be comfortable with reading and presenting modern machine learning conference papers.

Course Notes

Grading type: letter grades. Copies of lecture slides and other class information will be posted on the course website.

Technological Proficiency and Hardware/Software Required

There are no specific software requirements for this course.

This class will periodically plan to use portions of lectures for in-class (practical) “lab sessions”, to provide more hands-on experience with the theoretical concepts that we will learn. These sessions will involve running code and using implementations of probabilistic and generative models. It is therefore encouraged (though not required) that students bring a laptop with them to each class, so that they can follow along and participate in these practical lab sessions. These sessions will also be helpful for implementation skills that may be used in the course project running over the course of the semester. Additionally, please see the following information about the USC Computing Center Laptop Loaner Program (link).

Required Readings and Supplementary Materials

There are no required readings, supplementary materials, or textbook in this class.

Optional Readings and Supplementary Materials

The following resources are useful for many of the topics covered in this class:

1. Kevin Murphy, “Machine Learning: A Probabilistic Perspective”, 2012 (link).
2. Kevin Murphy, “Probabilistic Machine Learning: Advanced Topics”, 2023 (link).
3. Chris Bishop, “Pattern Recognition and Machine Learning”, 2006 (link).
4. Chris Bishop, “Deep Learning - Foundations and Concepts”, 2024 (link).
5. Stefano Ermon, Deep Generative Models, Course Notes (link).

6. Stefano Ermon, Probabilistic Graphical Models, Course Notes (link).

Description of Assignments and How They Will Be Assessed

A course project will give hands-on experience in probabilistic and generative models. It will be worth a substantial portion of the grade, and consist of four main assignments, detailed below:

1. Project group pitch: Students will form small groups. Each group will come up with a project idea, make a few slides, and share their idea with the class for feedback.
2. Midway report: Each group will be responsible for writing a short (4 page) midway report for their project, which focuses on a literature review, implementation plan, and any initial experiments.
3. Project group presentation: Each group will give a presentation to the class on their project.
4. Final report: Each group will be responsible for writing a 10-page report detailing their full project and all implementation and experimental results.

In addition to the course project, each student will give a short (20-30 minute) presentation on a paper or small group of papers (encouraged to be from literature review of their project). Presentations will take place in class.

There will be two short homework assignments given over the course of this class, focusing on methods of probabilistic or deep generative models, and/or approximate inference algorithms.

Participation

There will be credit given for participation in this class (see grading breakdown below). To recieve full credit, students will need to demonstrate ~15 minutes of participation each week. Students will have the opportunity to participate in discussion both in-class (during paper/project presentations), as well as in online project discussions.

Grading Breakdown

Assessment Tool (assignments)	**% of Grade**
1. Course project.	55 %
1.a Project group pitch	10 %
1.b Midway report	10 %
1.c Project group presentation	10 %
1.d Final report	25 %
2. Paper Presentation	15 %
3. Homework	15 %
4. Participation and Discussion	15 %
TOTAL	**100 %**

Assignment Submission Policy

Assignments will primarily be submitted in class.

Course-Specific Policies

Students are allowed to use their laptops in the classroom (as it is encouraged for in-class practical sessions).

Attendance

Attendance is encouraged, and will be a great opportunity to satisfy the participation and discussion requirement of this class.

Academic Integrity for this Class

Unless otherwise noted, this course will follow the expectations for academic integrity as stated in the USC Student Handbook. The general USC guidelines on Academic Integrity and Course Content Distribution are provided later in this syllabus.

Please ask the instructor [and/or TA(s)] if you are unsure about what constitutes unauthorized assistance on an exam or assignment, or what information requires citation and/or attribution.

Class Recordings and Course Content Distribution: You may not record this class without the express permission of the instructor and all other students in the class. Distribution of any notes, recordings, exams, or other materials from a university class or lectures — other than for individual or class group study — is prohibited without the express permission of the instructor; violations will be considered an intentional act to facilitate or enable academic dishonesty and reported to the university.

Use of Generative AI in this Course

Policy on the use of large language models (LLMs), and other AI generative models:

I expect that you may use AI (e.g., ChatGPT and image generation tools) in this class. Learning to use AI is an emerging skill, and I welcome the opportunity to meet with you to provide guidance with these tools during office hours or after class. Keep in mind the following:

- AI tools are permitted to help you brainstorm topics or revise work you have already written.
- If you provide minimum-effort prompts, you will get low-quality results. You will need to refine your prompts to get good outcomes. This will take work.
- Proceed with caution when using AI tools and do not assume the information provided is accurate or trustworthy If it gives you a number or fact, assume it is incorrect unless you either know the correct answer or can verify its accuracy with another source. You will be responsible for any errors or omissions provided by the tool. It works best for topics you understand.
- AI is a tool, but one that you need to acknowledge using. Please include a paragraph at the end of any assignment that uses AI explaining how (and why) you used AI. Failure to do so is a violation of academic integrity policies.
- Be thoughtful about when AI is useful. Consider its appropriateness for each assignment or circumstance. The use of AI tools requires attribution. You are expected to clearly attribute any material generated by the tool used.

Course Evaluations

Course evaluation occurs at the end of the semester university-wide, and they are an important review of students' experience in the class, and for the instructor to receive useful feedback. We will discuss and carry out course evaluations during a class session.

Course Schedule

A course calendar including a schedule and list of deliverables is given below. Note that, for each student, the paper presentation is due on the day they are scheduled to give their presentation.

	Topics/Daily Activities	Deliverables
Week 1	Introduction to probabilistic machine learning. a. Applications in deep generative modeling and decision making. b. Review of core concepts: sampling, (approximate) inference, learning.	
Week 2	Probability fundamentals and review. a. Probabilistic modeling preliminaries: useful definitions, terminology, and formulas. b. Where these fundamentals appear in modern machine learning models.	
Week 3	Probabilistic modeling and PGMs. a. Types of probabilistic graphical models (PGMs), plate notation. b. Overview on learning and approximate inference in PGMS.	Project pitch due.
Week 4	Learning in PGMs: classic algorithms. a. Expectation Maximization (EM) algorithm. b. Forward-backward, Viterbi, Baum–Welch, and belief propagation algorithms.	
Week 5	Approximate Bayesian inference. a. Markov-chain Monte Carlo (MCMC). b. Variational inference (VI).	
Week 6	Probabilistic (deep) generative models. a. Paradigms of generative modeling. b. Overview of model types: autoregressive, diffusion/score-based, flow-based.	HW 1 due.
Week 7	Autoregressive generative models. a. Large language models (LLMs) from a probabilistic perspective. b. Decoding in LLMs.	
Week 8	VAEs and GANs. a. Variational autoencoders (VAEs). b. Generative adversarial networks (GANs).	
Week 9	Diffusion models. a. Score-based and diffusion generative models. b. Recent advances in diffusion modeling.	Midway report due.
Week 10	Flow-based models. a. Normalizing flows. b. Continuous normalizing flow, neural ODEs. c. Score-matching.	
Week 11	Predictive uncertainty quantification. a. Gaussian processes (GPs), Bayesian neural networks (BNNs), and more. b. Ensemble models and dropout.	HW 2 due.
Week 12	Active learning. a. Sequential decision making via probabilistic models. b. Bayesian optimization and optimal experimental design. c. Acquisition functions.	
Week 13	Generative models in decision making. a. Diffusion for planning and generative design.	
Week 14	Project presentations	Group project presentation due.
Week 15	Project presentations	
FINAL	Final Project Report	Refer to the final exam schedule in the USC *Schedule of Classes* at classes.usc.edu.

Academic Integrity

The University of Southern California is foremost a learning community committed to fostering successful scholars and researchers dedicated to the pursuit of knowledge and the transmission of ideas. Academic misconduct — which includes any act of dishonesty in the production or submission of academic work (either in draft or final form) — is in contrast to the university's mission to educate students through a broad array of academic, professional, and extracurricular programs.

This course will follow the expectations for academic integrity as stated in the USC Student Handbook. All students are expected to submit assignments that are their own original work and prepared specifically for this course and section in this academic term. You may not submit work written by others or "recycle" work prepared for other courses without obtaining written permission from the instructor(s). Students suspected of engaging in academic misconduct will be reported to the Office of Academic Integrity.

Other violations of academic misconduct include, but are not limited to, cheating, plagiarism, fabrication (e.g., falsifying data), knowingly assisting others in acts of academic dishonesty, and any act that gains or is intended to gain an unfair academic advantage.

Academic dishonesty has a far-reaching impact and is considered a serious offense against the university. Violations will result in a grade penalty, such as a failing grade on the assignment or in the course, and disciplinary action from the university itself, such as suspension or even expulsion.

For more information about academic integrity see the student handbook or the Office of Academic Integrity's website, and university policies on Research and Scholarship Misconduct.

Please ask your instructor if you are unsure what constitutes unauthorized assistance on an exam or assignment or what information requires citation and/or attribution.

Course Content Distribution and Synchronous Session Recordings Policies

USC has policies that prohibit recording and distribution of any synchronous and asynchronous course content outside of the learning environment.

Recording a university class without the express permission of the instructor and announcement to the class, or unless conducted pursuant to an Office of Student Accessibility Services (OSAS) accommodation. Recording can inhibit free discussion in the future, and thus infringe on the academic freedom of other students as well as the instructor. (Living our Unifying Values: The USC Student Handbook, page 13).

Distribution or use of notes, recordings, exams, or other intellectual property, based on university classes or lectures without the express permission of the instructor for purposes other than individual or group study. This includes but is not limited to providing materials for distribution by services publishing course materials. This restriction on unauthorized use also applies to all information, which had been distributed to students or in any way had been displayed for use in relation to the class, whether obtained in class, via email, on the internet, or via any other media. Distributing course material without the instructor's permission will be presumed to be an intentional act to facilitate or enable academic dishonestly and is strictly prohibited. (Living our Unifying Values: The USC Student Handbook, page 13).

Statement on University Academic and Support Systems

Students and Disability Accommodations:
USC welcomes students with disabilities into all of the University's educational programs. The Office of Student Accessibility Services (OSAS) is responsible for the determination of appropriate accommodations for students who encounter disability-related barriers. Once a student has completed the OSAS process (registration, initial appointment, and submitted documentation) and accommodations are determined to be reasonable and appropriate, a Letter of Accommodation (LOA) will be available to generate for each course. The LOA must be given to each course instructor by the student and followed up with a discussion. This should be done as early in the semester as possible as accommodations are not retroactive. More information can be found at osas.usc.edu. You may contact OSAS at (213) 740-0776 or via email at osasfrontdesk@usc.edu.

Student Financial Aid and Satisfactory Academic Progress:
To be eligible for certain kinds of financial aid, students are required to maintain Satisfactory Academic Progress (SAP) toward their degree objectives. Visit the Financial Aid Office webpage for undergraduate- and graduate-level SAP eligibility requirements and the appeals process.

Support Systems:
Counseling and Mental Health - (213) 740-9355 – 24/7 on call
Free and confidential mental health treatment for students, including short-term psychotherapy, group counseling, stress fitness workshops, and crisis intervention.

988 Suicide and Crisis Lifeline - 988 for both calls and text messages – 24/7 on call
The 988 Suicide and Crisis Lifeline (formerly known as the National Suicide Prevention Lifeline) provides free and confidential emotional support to people in suicidal crisis or emotional distress 24 hours a day, 7 days a week, across the United States. The Lifeline consists of a national network of over 200 local crisis centers, combining custom local care and resources with national standards and best practices. The new, shorter phone number makes it easier for people to remember and access mental health crisis services (though the previous 1 (800) 273-8255 number will continue to function indefinitely) and represents a continued commitment to those in crisis.

Relationship and Sexual Violence Prevention Services (RSVP) - (213) 740-9355(WELL) – 24/7 on call
Free and confidential therapy services, workshops, and training for situations related to gender- and power-based harm (including sexual assault, intimate partner violence, and stalking).

Office for Equity, Equal Opportunity, and Title IX (EEO-TIX) - (213) 740-5086
Information about how to get help or help someone affected by harassment or discrimination, rights of protected classes, reporting options, and additional resources for students, faculty, staff, visitors, and applicants.

Reporting Incidents of Bias or Harassment - (213) 740-2500
Avenue to report incidents of bias, hate crimes, and microaggressions to the Office for Equity, Equal Opportunity, and Title for appropriate investigation, supportive measures, and response.

The Office of Student Accessibility Services (OSAS) - (213) 740-0776
OSAS ensures equal access for students with disabilities through providing academic accommodations and auxiliary aids in accordance with federal laws and university policy.

USC Campus Support and Intervention - (213) 740-0411
Assists students and families in resolving complex personal, financial, and academic issues adversely affecting their success as a student.

Diversity, Equity and Inclusion - (213) 740-2101
Information on events, programs and training, the Provost's Diversity and Inclusion Council, Diversity Liaisons for each academic school, chronology, participation, and various resources for students.

USC Emergency - UPC: (213) 740-4321, HSC: (323) 442-1000 – 24/7 on call
Emergency assistance and avenue to report a crime. Latest updates regarding safety, including ways in which instruction will be continued if an officially declared emergency makes travel to campus infeasible.

USC Department of Public Safety - UPC: (213) 740-6000, HSC: (323) 442-1200 – 24/7 on call
Non-emergency assistance or information.

Office of the Ombuds - (213) 821-9556 (UPC) / (323-442-0382 (HSC)
A safe and confidential place to share your USC-related issues with a University Ombuds who will work with you to explore options or paths to manage your concern.

Occupational Therapy Faculty Practice - (323) 442-2850 or otfp@med.usc.edu
Confidential Lifestyle Redesign services for USC students to support health-promoting habits and routines that enhance quality of life and academic performance.