

STCS 6701: Probabilistic Models and Machine Learning

David M. Blei
Columbia University

September 6, 2025

Description. *Probabilistic Models and Machine Learning* is a PhD-level course about how to design and use probability models. We study their mathematical properties, algorithms for computing with them, and applications to real problems. We study both the foundations and modern methods in this field. Our goals are to understand probabilistic modeling, to begin research that makes contributions to this field, and to develop good practices for building and applying probabilistic models.

Prerequisites. The prerequisites are: knowledge of basic probability and statistics, calculus, and some optimization; comfort writing software to analyze data; familiarity with a good programming language for statistics and machine learning, such as R or Python.

Requirements and Grades. The requirements are weekly reader reports, several homework assignments, and a final project.

- *Reader reports.* Each week, write what you thought about the reading. These papers can be up to one page; they can be as short as one paragraph. (Short is good.) They are required, but not graded. They must be handed in during class. Late papers are not accepted.
- *Homework.* The homework involves problems, programming, and data analysis. There will be three assignments. Each student has five late days to use for late homework.
- *Final project.* The main requirement is the final project. Most projects use and develop probabilistic machine learning to analyze real-world data. Some projects develop novel theoretical research in probabilistic machine learning. Ideally, your project connects to your doctoral research.

A project report is due at the end of the semester. There will also be some intermediate assignments to check your progress. We grade your project on both content and writing quality.

Please prepare all written work using the LaTeX template we provide.

Your grade decomposes as follows:

- Final Project: 60%
- Homeworks: 30%
- Completion of reader reports: 10%

Independent of this breakdown, you must hand in all homework, at least 10 reader reports, and a final project. Without satisfying these requirements, you may not receive a passing grade.

In all of the assignments, good writing is important. Three good books about writing are [Strunk and White \(1979\)](#), [Williams \(1981\)](#), and [Francis-Noel and Turner \(2011\)](#).

Syllabus

Below are the subjects we cover and in what order. (It may change.)

The Basics of Probabilistic Machine Learning

1. Introduction and the ingredients of probabilistic models
2. Basic distributions, conjugate priors, and the exchangeable data model
3. Conditional models: Linear and logistic regression, stochastic MAP estimation

Motifs and Algorithms in Probabilistic Models

4. Bayesian mixtures and the Gibbs sampler
5. Mixed-membership, topic models, and variational inference
6. Matrix factorization, recommendation systems, and efficient MAP

Generalizations and Advanced Ideas

7. Exponential families, conjugacy, and generalized linear models
8. Deep learning, diffusion models, and large language models
9. Black box variational inference

Further Topics

10. Model criticism and revision
11. Probabilistic machine learning and applied causality
12. Summary (and wiggle room)

References

Francis-Noel, T. and Turner, M. (2011). *Clear and Simple as the Truth: Writing Classic Prose, Second Edition*. Princeton University Press.

Strunk, W. and White, E. (1979). *Elements of Style*. Longman Press.

Williams, J. (1981). *Style: Towards Clarity and Grace*. University of Chicago Press.