Course Information (draft)

CS 57800 – Statistical Machine Learning

Time/Location: MWF 1:30 PM - 2:20 PM at LWSN B155

Instructor: Ruizhe Zhang
Office hours: TBD

Teaching Assistants: Shreya Nasa, Yucheng Zhang
Office hours: TBD

Prerequisites: Probability (vectors and matrices, expectation, conditioning, common distributions), Linear algebra (eigenvalues, positive semidefinite matrices), Multivariate calculus, and comfort with proofs.
Undergraduates interested in taking the course should contact the instructor for permission.

Webpage:

1 Course Description

This graduate course develops the mathematical foundations of statistical machine learning, with statistical reasoning and probabilistic inference as recurring themes. Modern AI systems can often produce predictions or implement standard learning algorithms automatically. The harder questions are what those predictions mean, whether they are supported by the available data, which assumptions and approximations they rely on, and when they may fail. Accordingly, this offering is *theory-focused*: rather than presenting a catalog of methods, it develops a mathematical account of four questions that recur throughout modern machine learning:

- **Generalization:** What conclusions about an unknown population are justified by finitely many samples, and what cannot be learned without additional assumptions?
- **Inference:** How can predictions, latent structure, and marginal probabilities be computed or approximated in structured probabilistic models?
- **Generation:** How can we sample from a complicated learned distribution, and how do estimation, mixing, and discretization errors affect the result?
- **Reliability:** How are learning guarantees affected by approximation, distribution shift, privacy constraints, objective misspecification, or adversarial modification?

These questions organize the course into six connected units: statistical learning theory; unsupervised and latent-variable learning; probabilistic inference and statistical physics; diffusion-model theory; modern neural-network theory; and trustworthy machine learning. Students will also read selected research papers to see how these ideas are applied, tested, and extended in current research.

2 Course Materials

We will not be following any particular textbook, but the following resources could be helpful for or complementary to various parts of the course:

- Shai Shalev-Shwartz and Shai Ben-David, *Understanding Machine Learning: From Theory to Algorithms*;
- Mehryar Mohri, Afshin Rostamizadeh, and Ameet Talwalkar, *Foundations of Machine Learning*;
- Martin Wainwright and Michael Jordan, *Graphical Models, Exponential Families, and Variational Inference*;
- Larry Wasserman, *Statistical Methods for Machine Learning*;
- Mark Davenport, *Statistical Machine Learning*;
- Ankur Moitra, *Algorithmic Aspects of Machine Learning*;
- Varun Kanade, *Computational Learning Theory*;
- John Duchi, *Statistics and Information Theory*;
- Marc Mézard and Andrea Montanari, *Information, Physics, and Computation*;
- Yuansi Chen, *Computational and Statistical Aspects of Diffusion Models*;
- David Levin, Yuval Peres, and Elizabeth Wilmer, *Markov Chains and Mixing Times*;
- Joel A. Tropp. *An Introduction to Matrix Concentration Inequalities*;
- Stephen Boyd and Lieven Vandenberghe, *Convex Optimization*;
- Kaare Brandt Petersen and Michael Syskind Pedersen, *The Matrix Cookbook*

3 Grading

The grading breakdown is as follows:

- Problem sets (30%)
- Midterm exams (15%$\times$2)
- Final project (40%)

Problem sets They will be assigned once every week or two, for a total of 8 problem sets. All submissions must be in PDF format. You are allowed to discuss problems with other students, but the final submission must be written individually. Each problem set will be worth 4 points, and your total points will be capped at 30.

Midterm exams There will be two in-class midterm exams; each will be worth 15 points. You can expect them both to be "pen and paper" exams, with roughly 3 response questions on them.

Attendance Regular attendance at lectures is expected. If you have any unavoidable conflict (illness, travel, etc.), you should try to let me know in advance whenever possible.

Final project. MOST important part of this course. Students are encouraged to work as teams of two or three. Projects by one individual are discouraged. Each project must investigate a precise question that is not answered directly by the cited literature and make a focused, technically substantive, and independently verifiable contribution. Students should first establish an appropriate existing result or baseline and then go beyond it, for example by proving a nontrivial extension or counterexample; designing controlled experiments that distinguish between plausible hypotheses; resolving a reproducibility discrepancy; or characterizing both a failure regime and its underlying mechanism. Merely reproducing a published result, applying an existing method or software package to a new dataset, conducting an undirected parameter sweep, or reporting that an attempted approach did not work is not sufficient. Negative or inconclusive findings may receive full credit when the investigation is well designed, rules out important alternative explanations, and supports a precise scientific conclusion. The contribution need not be publication-ready, but it should teach a technically competent reader something that is not already immediate from the cited literature.

The project has three graded components: a short proposal, a final report with an accompanying technical artifact, and a presentation with individual questions. Each team will also complete a brief, ungraded mid-semester check-in with the instructor or a teaching assistant.

Generative AI may be used as a research tool. Students remain responsible for verifying and understanding every claim, citation, proof, computation, figure, and piece of code they submit. The final report must contain a brief statement describing how AI tools were used and how consequential outputs were checked. Students may be asked to explain any part of the submitted work during the presentation.

4 Tentative List of Topics

The detailed calendar will be maintained on the course webpage. The following sequence is tentative and may be adjusted to match the pace and preparation of the class.

- **Statistical learning problems.** Population and empirical risk; Bayes classification; conditional-mean regression; bias–variance; least squares and logistic regression.

- **Concentration and PAC learning.** Hoeffding's inequality; finite-class uniform convergence; empirical risk minimization; realizable and agnostic PAC learning.
- **VC dimension and the fundamental theorem.** Shattering and growth functions; the Sauer–Shelah lemma; VC generalization bounds; the fundamental theorem of statistical learning; halfspaces.
- **Margins and linear classification.** The perceptron mistake bound; maximum-margin classification; implicit bias of logistic regression.
- **Clustering.** The k-means objective; Lloyd's algorithm; local optima and initialization.
- **Method of moments for latent-variable models.** Topic models; second- and third-order moments; Jennrich's algorithm; identifiability and perturbation.
- **Exponential families and energy-based models.** The log-partition function; maximum entropy; the Gibbs variational principle; Ising and Boltzmann models.
- **Graphical models and exact inference.** Factorization and conditional independence; variable elimination; the sum–product algorithm; exact belief propagation on trees.
- **Variational inference and mean-field methods.** The ELBO; expectation–maximization; mean-field approximation; coordinate-ascent variational inference.
- **Statistical physics of learning.** Free energy and phase transitions; Bethe approximation and loopy belief propagation; Hopfield networks; the replica method and order parameters.
- **Markov-chain Monte Carlo and mixing.** Detailed balance; Metropolis–Hastings and Gibbs sampling; coupling; conductance and slow mixing.
- **Langevin dynamics and score estimation.** Langevin dynamics; the unadjusted Langevin algorithm; discretization bias; score matching and denoising score matching.
- **Diffusion models.** Forward and reverse diffusion; Tweedie's formula; score, noise, and denoiser parameterizations; reverse-time dynamics.
- **Neural-network training dynamics.** Gradient flow; implicit bias; the neural tangent kernel; lazy training and feature learning.
- **Attention and in-context learning.** Self-attention; permutation equivariance; in-context learning for linear regression.
- **Trustworthy machine learning.** Differential privacy; preference learning and alignment; watermark detection and robustness for generative models.

5 Course Policies[1]

Late Policy Problem sets may be submitted up to 2 days late, with a 10% deduction in credit. Late submissions of the final project report will not be accepted under any circumstances.

Academic Dishonesty Purdue prohibits "dishonesty in connection with any University activity. Cheating, plagiarism, or knowingly furnishing false information to the University are examples of dishonesty." Furthermore, the University Senate has stipulated that "he commitment of acts of cheating, lying, and deceit in any of their diverse forms (such as the use of substitutes for taking examinations, the use of illegal cribs, plagiarism, and copying during exams) is dishonest and must not be tolerated. Moreover, knowingly to aid and abet, directly or indirectly, other parties in committing dishonest acts is in itself dishonest." While it is all right to discuss homework in this class with other students in general terms, do not copy another student's homework or let anyone copy your homework. These discussions should be appropriately acknowledged. When two identical homeworks are found, both are sent to the Dean of Students to determine who copied from whom. Penalties for academic dishonesty range from a 0 grade on one assignment to a failing grade in the class and even expulsion from the University. A hearing at the Dean of Students' Office can ruin your whole day.

Violent Behavior Policy Purdue University is committed to providing a safe and secure campus environment for members of the university community. Purdue strives to create an educational environment for students and a work environment for employees that promote educational and career goals. Violent behavior impedes such goals. Therefore, violent behavior is prohibited in or on any University facility or while participating in any university activity.

Students with Disabilities Purdue University is required to respond to the needs of the students with disabilities as outlined in both the Rehabilitation Act of 1973 and the Americans with Disabilities Act of 1990 through the provision of auxiliary aids and services that allow a student with a disability to fully access and participate in the programs, services, and activities at Purdue University. The ODOS Testing Center is an excellent place for students with disabilities to take exams for this class. Please tell the instructor at least one week in advance if you wish to take the exams there. If you have a disability that requires other special accommodation, please tell the instructor early in the semester. It is the student's responsibility to notify the Disability Resource Center of an impairment/condition that may require accommodations and/or classroom modifications.

Emergencies In the event of a major campus emergency (such as a tornado, earthquake, flu epidemic or terrorist attack), course requirements, deadlines and grading percentages are subject to changes that may be necessitated by a revised semester calendar or other circumstances beyond the instructor's control. Relevant changes to this course will be posted onto the course website

[1] Most of the following material is based on the policy followed by Prof. Samuel Wagstaff in his cryptography courses.

or can be obtained by contacting the instructor via email or phone. You should read your Purdue email frequently.

Nondiscrimination Purdue University is committed to maintaining a community which recognizes and values the inherent worth and dignity of every person; fosters tolerance, sensitivity, understanding, and mutual respect among its members; and encourages each individual to strive to reach his or her own potential. In pursuit of its goal of academic excellence, the University seeks to develop and nurture diversity. The University believes that diversity among its many members strengthens the institution, stimulates creativity, promotes the exchange of ideas, and enriches campus life. Purdue University prohibits discrimination against any member of the University community on the basis of race, religion, color, sex, age, national origin or ancestry, marital status, parental status, sexual orientation, disability or status as a veteran. The University will conduct its programs, services and activities consistent with applicable federal, state and local laws, regulations and orders and in conformance with the procedures and limitations as set forth in Executive Memorandum No. D-1, which provides specific contractual rights and remedies. The instructor agrees completely with all Purdue policies mentioned in this document.

Privacy The Federal Educational Records Privacy Act (FERPA) protects information about students, such as grades. If you apply for a job and wish to use the instructor as a reference, you should tell the instructor beforehand. Otherwise, the instructor cannot say anything about you to a prospective employer who might call. The instructor is happy to provide references and to write letters of recommendation for his students as needed.

Disclaimer.

This syllabus and grading scheme are subject to change.