CS839: Differential Privacy and Learning February 23, 2026

Course Project Instructions

Instructor: Gavin Brown

Every student will complete a final project, representing 30% of their final grade. Each project will result in an in-class presentation and a written document. Students may work alone or in groups of two or three; I will expect more from larger groups.

The main goal of the project is for you to engage as a researcher with the course's subject or related areas. You will need to pick a topic, survey the relevant literature, formulate precise open questions, and build an understanding of possible paths to their resolution. By the end of the class, you should have built up the strong start of a real research project. Of course, in the process of building up this foundation you may be able to (partially) solve some or all of your questions. I expect that multiple course projects will eventually lead to published work, although you will not be graded on whether or not you have made novel research contributions.

Here are sketches of potential projects.

- Take a notable theory paper from the last few years, understand the problem it's solving and the new ideas it introduced. In your writeup, present key steps of the proof. Discuss remaining open questions and how the techniques could be extended or applied elsewhere.
- Take a recent work focused on implementation or empirical evaluation and try to replicate its results. Are there subtle obstacles or drawbacks? Is the method ready for use by practitioners?
- Many papers analyze new DP algorithms that are efficient/practical in theory, but the paper provides or code or only a limited empirical evaluation. Contributing a stronger empirical evaluation could be very valuable, and may also inform future theoretical analyses.

1 Timelines and Guidelines

Feel free to pick a topic that overlaps with your ongoing research. However, the work for this project must be new. You are not allowed to present work that is already done.

As with the rest of the class, you are permitted (and encouraged!) to use AI tools throughout the project, but you are responsible for everything you turn in.

1.1 Timelines

Proposal/Initial Discussion Before spring break (i.e., on or before **March 28**), each group needs to meet with me in-person to discuss their plans. At this point, you should have a clear idea of the work you plan to do. (You do not need to have made any progress, although the more you've done the more useful feedback I can provide.) Come by my office hours or, if you can't make them, contact me to set up a meeting.

Paper Draft Your group must submit a draft writeup by **April 21**. It's fine if this is very rough, but it should contain a start on all the basic sections.

Presentation The last three lectures (April 23, 28, and 30) will be student presentations.

Finished Paper Please submit the final version of your paper by **May 1**, the last day of classes.

1.2 Details on the Paper

Your group will submit a write-up of your review/results, which will be posted on the class website.[1] The main body of your paper should be roughly six pages, with no space restriction for references or appendices. You must typeset your work in LaTeXand use the ICML 2026 submission instructions.

Your paper should include the following sections:

- Abstract
- Introduction
- Related Work
- Results
- Future Directions

Aside from Results, each of these should look somewhat like a slimmed-down version of what you see in published work. Your Results section will depend heavily on your project; for example, the presentation of the core technical ideas of a theory paper.

1.3 Details on the Presentation

Each member of your group should participate. Plan for these to be around 20 minutes each plus time for questions. (Details may change once the groups are finalized.)

2 Some Open Questions/Directions

Some open questions of varying concreteness.

- What is the sample complexity of approximate DP-PAC learning? For pure DP, the answer is something called "probabilistic representation dimension." Without privacy, the VC dimension characterizes the sample complexity of PAC learning up to constants. Under approx DP, we have a major open question. We know a hypothesis class is DP-PAC learnable if and only if it has finite Littlestone dimension, but we lack a characterization of the sample complexity.
- Recent work of Raskhodnikova et al. [2021] (which arose from a course project!) introduced "DP sampling from distributions." This is an interesting problem with only a few papers on it. I will also note that there may be variations of the original question which are also interesting.
- The sample complexity of privately estimating the mean of $\mathcal{N}(\mu, 1)$ is known up to constants [Karwa and Vadhan, 2017]. What are the optimal constants for this task?
- The work of Steinke and Steinke [2025] gives an algorithm which is efficient in its use of queries but still has a large running time. Is there a computationally efficient version with similar accuracy guarantees?

[1] Similar to a dissertation, I can "embargo" any papers which contain partial results that you plan to continue working on and submit for publication. Just let me know.

- A number of recent papers, including Bombari et al. [2025], Lin et al. [2025], Dwork et al. [2025], use sophisticated tools from optimization to analyze variants of DP gradient descent for statistical problems, with a special focus on linear regression. It's natural to examine how these tools apply to other problems, for example ridge regression or logistic regression.
- Empirically, many of the best-performing algorithms for DP linear regression require carefully tuned hyperparameters. There are some methods that avoid some of these issues, but they are usually more sample-hungry. Are there practical adaptive methods that tune hyperparameters automatically?
- Can you improve the state-of-the-art in differentially private image classification?
- How does the noise added for privacy interact with training or fine-tuning neural networks? We always want to add "less" noise, but how should we think about the shape or directions? How might we build mechanisms to add better forms of noise? (There is lots of work on this, but much more to be done!)

3 Resources for Finding Projects

The annual workshop "Theory and Practice of Differential Privacy" has talks and posters covering a large fraction of the DP literature in the past year. Its website has videos from talks and pointers to many papers which appeared as posters (2023, 2024, 2025).

In particular, Google researcher Matthew Jagielski gave a talk at TPDP 2024 officially titled "Data and Privacy in Data Privacy" and unofficially titled "Some research directions I think are interesting and maybe you will, too" (video, notes). It presents a number of important and understudied directions.

There are lots of open problems about correlated noise for DP deep learning in Pillutla et al. [2025].

The conferences COLT and ALT have many theoretical papers on differential privacy. Most years, STOC and FOCS will have a handful of papers on DP. The main machine learning conferences also host many important DP papers, but their proceedings are (slightly) harder to search.

4 Topics Beyond Our Class Lectures

Our class has focused on a particular subset of the differential privacy literature. There is a lot we have not touched on! Here are some examples.

- Local/shuffle model
- Combining DP with cryptographic tools
- Computational differential privacy
- Continual learning
- Person-level (i.e., each individual may contribute multiple points)
- DP for graph data.
- Private learning with public data

Beyond differential privacy, here are some areas worthy of study:

- Attacks on privacy (e.g., membership inference or training-data reconstruction)
- Privacy auditing
- Machine unlearning
- Copyright
- Adaptive data analysis/replicable learning
- Understanding privacy in language and LLMs

5 A Few Interesting Recent Papers

A list of papers, without further comment.

- How to Make the Gradients Small Privately: Improved Rates for Differentially Private Non-Convex Optimization
- DPZero: Private Fine-Tuning of Language Models without Backpropagation
- Adapting to Linear Separable Subsets with Large-Margin in Differentially Private Learning
- Local Node Differential Privacy
- Keeping a Secret Requires a Good Memory: Space Lower-Bounds for Private Algorithms
- Skirting Additive Error Barriers for Private Turnstile Streams
- The Gaussian Mixing Mechanism: Renyi Differential Privacy via Gaussian Sketches
- A Unified Characterization of Private Learnability via Graph Theory
- Private Stochastic Convex Optimization with Heavy Tails: Near-Optimality from Simple Reductions
- Privacy of Noisy Stochastic Gradient Descent: More Iterations without More Privacy Loss
- Differentially Private Bilevel Optimization: Efficient Algorithms with Near-Optimal Rates
- Towards Separating Computational and Statistical Differential Privacy
- Private Everlasting Prediction
- On Differentially Private Linear Algebra
- Algorithmically Effective Differentially Private Synthetic Data

References

Simone Bombari, Inbar Seroussi, and Marco Mondelli. Better rates for private linear regression in the proportional regime via aggressive clipping. *arXiv preprint arXiv:2505.16329*, 2025.

Cynthia Dwork, Pranay Tankala, and Linjun Zhang. Differentially private learning beyond the classical dimensionality regime. In *Theory of Cryptography Conference*, pages 321–355. Springer, 2025.

Vishesh Karwa and Salil Vadhan. Finite sample differentially private confidence intervals. *arXiv preprint arXiv:1711.03908*, 2017.

Shurong Lin, Eric D Kolaczyk, Adam Smith, and Elliot Paquette. High-dimensional privacy-utility dynamics of noisy stochastic gradient descent on least squares. *arXiv preprint arXiv:2510.16687*, 2025.

Krishna Pillutla, Jalaj Upadhyay, Christopher A Choquette-Choo, Krishnamurthy Dvijotham, Arun Ganesh, Monika Henzinger, Jonathan Katz, Ryan McKenna, H Brendan McMahan, Keith Rush, et al. Correlated noise mechanisms for differentially private learning. *arXiv preprint arXiv:2506.08201*, 2025.

Sofya Raskhodnikova, Satchit Sivakumar, Adam Smith, and Marika Swanberg. Differentially private sampling from distributions. *Advances in Neural Information Processing Systems*, 34:28983–28994, 2021.

Günter F Steinke and Thomas Steinke. Privately estimating black-box statistics. *arXiv preprint arXiv:2510.00322*, 2025.