CSE 493S: Advanced Topics in Machine Learning

Project Proposal Guidelines

Autumn 2025

Due: Tuesday, October 21, 2025 at 11:59 PM

Overview

The project is designed to give you an opportunity to engage with current machine learning research and contribute useful knowledge to the community. Your team (1–4 students) will choose a direction and develop a project based on recent research. The proposal is worth **10% of your project grade**.

Teams: Your team should be composed of 1–4 people. We encourage diverse teams—working with people whose perspectives and backgrounds are different from yours gives you new opportunities to learn. Diverse teams are stronger teams! Use the Team Finder thread on Ed Discussion to find teammates. If you are having trouble finding a team, let us know (email `cse493s-staff@cs.washington.edu`) and we will help you.

Project Directions

Your project should follow one of these directions:

- **Replication of recent work:** Reproduce results from a recent paper, potentially extending to new settings
- **Fair comparisons for algorithms:** Fairly compare attacks/defenses (backdoor attacks, data poisoning), or compare uncertainty quantification methods (conformal prediction, multicalibration, proper scoring rules)
- **Summarizing a line of theoretical work**, for instance:
  - Explain an empirical phenomenon (e.g., generalization despite large hypothesis class, nonconvex optimization success, mode connectivity, lottery ticket hypothesis)
  - Develop a new algorithm (e.g., extensions to SGD, online learning algorithms)
  - Survey calibration guarantees (temperature scaling, histogram binning, platt scaling, multicalibration)
- **Original research**, such as:
  - Proposing and evaluating a new idea on top of an existing codebase
  - Trying out new attacks on neural networks or designing new defenses
  - Memory/compute-efficient fine-tuning methods (zeroth-order optimization)

  - Understanding low-dimensional landscape of fine-tuning LLMs
  - Testing scaling properties of published methods
  - Applying online learning algorithms (multiplicative weights, OCO) to new domains
  - Developing new uncertainty quantification methods for distribution shift
- **Your own research project**, if relevant

Proposal Requirements (1 page)

Your proposal should be **one page** and include:

1. **Team information:** Names, NetIDs, undergraduate/graduate status for each member
2. **Citation(s):** A BibTeX citation of the paper(s) or research direction you plan to base your project on, *with a URL*
3. **Goals and expected outcomes of your project:** What specific results do you aim to produce? What are the concrete deliverables (e.g., "We will replicate Figure 3 from [Paper] and extend it to dataset X")?
4. **Concrete description of the first milestone** to be achieved by Nov 6 (whether it is reproducing a baseline, running an eval, reproducing a result in a referenced paper, surveying the related theoretical work, etc.). *We want to make sure you know where to start.*
5. **For empirical projects:**
   - A short description of whether and how you can access the data needed for the project
   - If applicable: the size of the dataset you want to use and the size of the model you want to use (e.g., "ImageNet-1K with ResNet-50" or "GPT-2 small on WikiText-103")
   - Whether you will use existing code (*in that case, provide a link to the code*) or implement yourself
   - A discussion of the feasibility of the computation you will need to do (essentially, an argument that the project will be feasible given your computational resources)
6. **For theoretical projects:**
   - Your conjecture of what you want to show
   - Some ideas on what techniques you plan to try out
   - What related theoretical results you will build on
7. **For empirical + theory projects:** Some combination of the above

Note: Estimations do not have to be exact, but you will get more useful feedback if you include specific details. You can be flexible on what you decide to include in the proposal, as long as you have all the basic information needed.

Important Information

- **Submission:** Upload as PDF to Gradescope. One submission per team with all members listed.
- **No late submissions:** Late submissions receive zero points.
- **Grade distribution:** Proposal 10%, Milestone 25%, Final Report + Presentation 65%
- **Final report:** Must not exceed 8 pages (excluding references). Code should be in a public GitHub repository.

Getting Help

- Attend office hours to discuss project ideas before the deadline
- Use the Team Finder thread on Ed Discussion to find teammates
- Post questions about project ideas on Ed Discussion
- Contact `cse493s-staff@cs.washington.edu` with questions