

CS 839: PRESENTATION INFORMATION

Instructions: This document contains some instructions and useful information for the student presentations for the Fall '24 offering of CS 839: Foundation Models. The goal for the class presentation is to read a high-quality paper that covers our set of topics and to present it to the class. Some potential paper ideas are listed below, but if you have any questions about what is appropriate, please contact me!

Goals. We have several goals for these presentations.

- Familiarize yourself and the class with major research results in our area. This means we must understand the problem setting, results, the core mechanism that makes these results work, limitations, and more. This will also help you in your own paper reading; it is a key skill for any researcher,
- Develop an understanding of how results fit into the overall universe of machine learning work and research. In other words, we want to grasp the context of this work,
- Build skill with identifying potential future directions: extensions of the work, open cases and related open problems, applications to new areas,
- Work on our presentation and talk skills—also very useful for your own research presentations.

Deliverables. There are two items that must be turned in:

1. Presentation proposal: a short document (0.5 pages in length is fine) with the paper that you're going to present, the reason why you chose it, and how it fits this course. This is due Nov. 7. on our spreadsheet (linked on the course website); the earlier you sign up, the better.
2. Post-presentation: submit your notes or slides, and include answers to questions you received in class. This is due Dec. 9th, but you may submit it as early as you wish.

Team Sizes. Your group sizes should be 4-6 (but there is some flexibility). We will figure out based on the number of teams how many slots there will be and how long these slots will be.

1 Frequently Asked Questions

- What format is OK for the presentation? A: You can use slides or any other means of presenting. If you use slides, please avoid purely copying pages of the paper into the slides. You should use the format that is most conducive to explaining the results, as above. Useful: outlines/roadmaps in the presentation, providing motivation and useful connections to other work, giving examples and intuition, and coming up for air after diving into technical segments. This should be useful practice for a long talk, like the one you give for your defense.
- How long should the presentation be? A: It will depend on the number of teams we end up having, but you can expect somewhere between 20 and 50 minutes. We will have some time for questions for each presentation.
- What are we being graded on? A: We'll be using the goals above and the level of effort/sophistication that went into your presentation. You should be able to explain the key ideas clearly, identify limitations when asked, etc. If there are any questions on whether your planned presentation meets our bar, please ask any time.
- Can we present on multiple papers? A: This might be OK if you are comparing several techniques. Feel free to ask me about it. However, the number of papers should be low enough to be able to get in-depth in at least some of them. Don't view this as a survey.

- Do I have to attend other students' presentations? A: Yes; we will consider this to be part of your grade as well. In addition, you should participate by asking questions. It is one of our goals to make these presentations as interactive as possible; treat them as a discussion more than a lecture.
- Can I coordinate my presentation with my project? A: Yes; this is actually ideal. However, you are not required to. If you want to do a different line of work for your project, that is also all right. You can similarly switch teams if you wish to.
- What should I assume the class already knows? A: Anything we've worked on during the lectures or that is standard knowledge is fine. If your topic is more sophisticated than that, include the background that's needed for your classmates to understand.
- Do I need a separate report to submit along with the presentation? A: No, you can submit slides or the notes that you used to give the presentation. The only additional content you should include is the answers to questions that were asked (if they were not already present in your presentation).

2 Sample Papers

Here are a few examples—but you are not limited to these.

- Dao and Gu, "Transformers are SSMs: Generalized Models and Efficient Algorithms Through Structured State Space Duality".
- Carlini et al, "Stealing part of a production language model".
- Dettmers et al, "QLoRA: Efficient Finetuning of Quantized LLMs".
- Kondratyuk et al, "VideoPoet: A Large Language Model for Zero-Shot Video Generation".
- Liu et al, "Deja vu: Contextual sparsity for efficient llms at inference time".
- Solatorio et al, "REaLTabFormer: Generating Realistic Relational and Tabular Data using Transformers".
- Guo et al, "How Close is ChatGPT to Human Experts? Comparison Corpus, Evaluation, and Detection".
- Zheng et al, "Less is More: Focus Attention for Efficient DETR".
- Alayrac et al, "Flamingo: a Visual Language Model for Few-Shot Learning".
- Tian et al, "JoMA: Demystifying Multilayer Transformers via Joint Dynamics of MLP and Attention".
- Park et al, "How Do Vision Transformers Work?".
- Ahn et al, "Do As I Can, Not As I Say: Grounding Language in Robotic Affordances".