Welcome to CSE481 (ML Capstone)!

Spring 2026

Pang Wei Koh

What is this class about?

- Intensive 10-week course on doing independent ML research in groups of three students
- Starting from project proposal and finishing with a full report
- Minimal lectures; focus is on small-group interactions and presentations

What is research?

- Pushing the boundaries of what's currently possible
- Understanding things that are not currently known

- Corollary: It should not be obvious if what you're proposing will actually work!

- Research is a community endeavor: read papers, build on open-source code

Why this class?

- Research unlocks many career paths, especially in AI

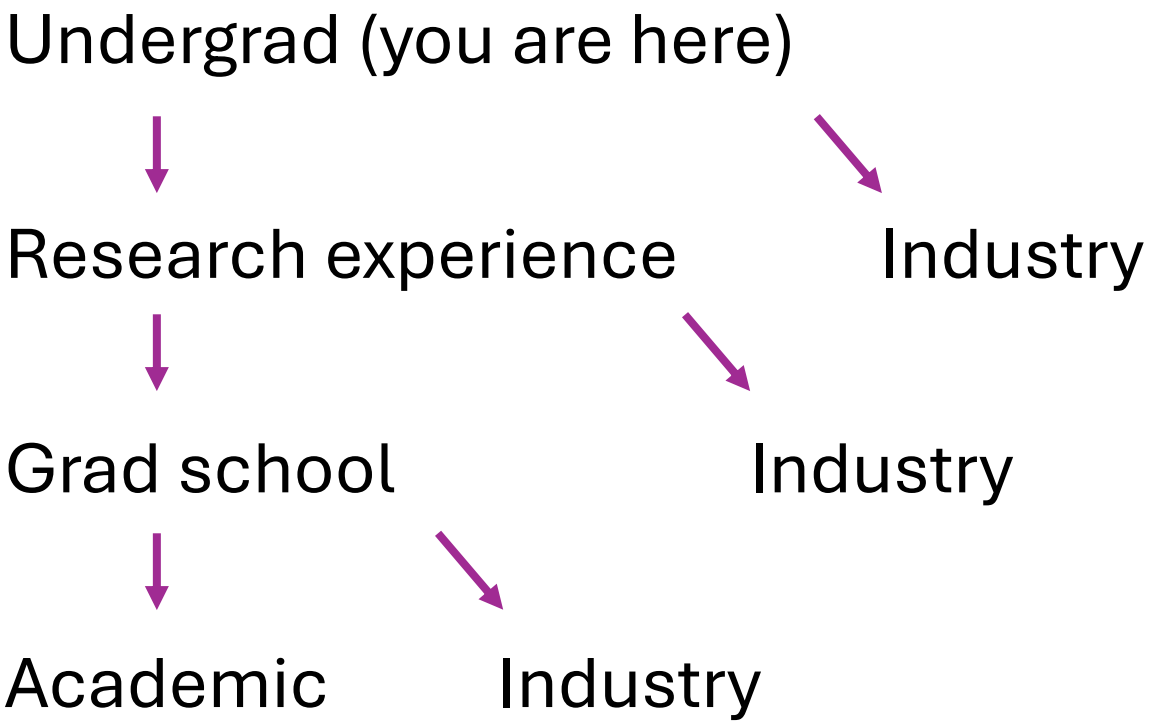

Why this class?

- Research experience is hard to obtain because labs are often full
- Our goal: Get a substantial research project under your belt that you can use to open other doors
  - We'll provide scaffolding + many opportunities for small-group interactions and class presentations/discussions
- Also, let you make a more informed choice about whether you'd like to keep doing research

Who is this class for?

- The final product is your project, so what you get is what you put in
- Please make sure you are committed and able to put in the time

- Prereqs: ML (CSE446/546) + one follow-up class
- You should be able to read papers in your chosen field (with effort) and build upon common libraries and tools

The course project

- Has to involve ML research at its core!
- It can involve other components: systems engineering, UI/UX, etc.
- Can be theoretical, empirical, and/or application-focused
- ML research is broadly construed
  - E.g., Studying AI-human interaction
- Scope it down
  - You don’t have to aim to publish by the end of the quarter
  - Resources are limited
  - Spend time reproducing existing work and implementing simple baselines

Good research is surprising

- There should be a surprise somewhere
  - In the problem formulation?
  - In the methods?
  - In the results?
- It should make readers learn something new and change how they think

Project FAQ

Must I be in a team of 3 people? Yes (mostly).

Let us know if you need help finding a team.

Project FAQ

Can I use a project that I've already been working on, and/or a project that I'm working on with another lab outside this class? No (mostly).

- Your course project must be started from scratch
- You're welcome to get external advice and feedback
- All coding and writing must be done by you and your team members

Course structure & grading

- Regular progress reports / blog posts, posted on Ed
  - 9 regular reports at 4% each, mostly in teams
  - + Project proposal: 9%
  - + Final report: 15%
- Four 5-min class presentations (5% each)
- Class and project participation (20% each)
- No curving; everyone can do well
  - Good idea but experiments didn't work out: Ok
  - Tried to do too much and didn't complete: Not ok
  - Each team member should pull their weight

Course structure & grading

- No final exam / poster presentation
- Upload your final 5-min presentation to Youtube

Classes

- Tuesday classes are always mandatory and in-person only

- Thursdays, first 2 weeks: Optional, focus on forming teams and developing project proposals
- Thursdays, remainder: In-class working time + mandatory scheduled discussions with course staff

Compute

Apply for free cloud compute:

- **Google Cloud:** $300 in free credits for new customers.
- **Google Colab Pro:** Available for higher education.
- **Microsoft Azure:** $100 of credits for students.

You may use other external compute resources that they have access to (e.g., from other research labs).

We will distribute a small amount of additional Google Cloud credits at the end of the second week

LLM policy

All submitted work must reflect the student's or team's own thinking and analysis. External tools, including software and computational resources, may be used as part of the research process. If students use automated tools (for example, for coding assistance, drafting, or analysis), they must:

- Be transparent about how these tools were used.
- Take responsibility for verifying and interpreting all outputs.
- Ensure that the final work reflects their own intellectual contribution. Teams will discuss their weekly reports with the course staff. We expect all students to know all details about their projects, especially around their own contributions, regardless of whether AI tools were used. Not following these guidelines will be treated as a violation of academic integrity and will have to be reported externally.

Any other questions?

- Feel free to contact us via Ed at any time (public or private posts)
- https://edstem.org/us/courses/96256/
- https://www.gradescope.com/courses/1266737/

- Read website for more info

Step 1: Teams and project proposals

- Today: Intros
- This Thursday (4/2): Optional class to form teams and discuss projects with course staff. Preliminary teams due 11:59pm. (Don’t panic!)
- Next Tuesday (4/7): Discuss project ideas with course staff and each other. Teams and 3 project ideas due 11:59pm.
- Next Thursday (4/9): Optional class to discuss projects with course staff. 3 fleshed out project ideas due 11:59pm.

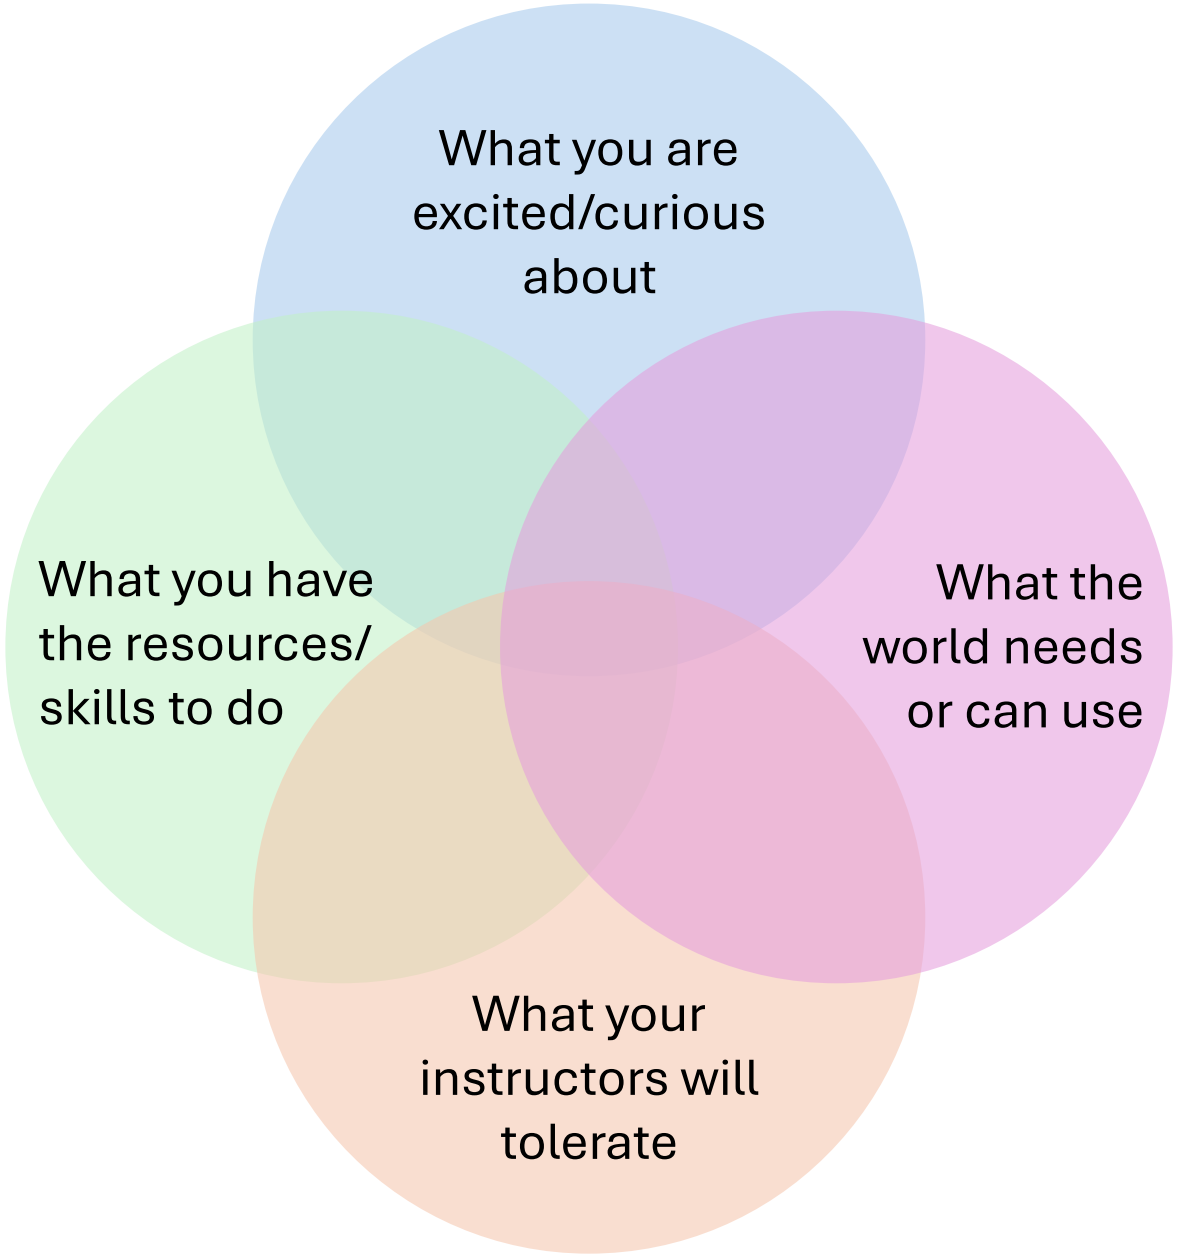
What you are excited/curious about
What you have the resources/ skills to do
What the world needs or can use
What your instructors will tolerate

Intros!

- https://docs.google.com/presentation/d/1e3CQqttl4TjiLk7at6Db2nYJLib2OgAhjjl2umm20y8/edit?usp=sharing