

CIS 625, Theory of Machine Learning

Final Project Guidelines

Prof. Michael Kearns

This document Dec. 7, 2020; project proposals due by email Dec. 14; project Write-ups due Dec. 31

For the final projects, you are free to work in teams of up to *four* people, but all team members should contribute significantly to both the work and the write-up. You are also free to work alone. The amount and quality of a team's work should be commensurate with its size.

As has been mentioned during lectures, there is a great deal of flexibility around the nature of your projects. The sample categories below are suggestive and not exhaustive. Regardless of the type of project you pursue, the one firm requirement is that your write-up have at least a section that *clearly relates your project content directly to course topics*. This could be in the form of direct connections to course models or results, a critique of course models or results, an extension to course topics, an experimental comparison to course topics, etc.

One thing to be avoided is “double-dipping” with other course projects and requirements — e.g. if you're taking a related course this semester, or are doing your own ML research, please don't try submitting that work for this course project. It's fine if there is some overlap, but in such cases you should clearly declare what the other work is, and clearly indicate the original contribution for this course project.

Below are some sample project categories. If you want to do something that doesn't seem to fit any of these, just go ahead and propose it and I'll provide feedback.

1. **Theory Research.** Here you could either pose a problem in one of the models we examined that (hopefully) has not already been considered/solved in the literature, and try to solve or make progress on it. You could also consider natural extensions of course models and explore what you can prove. In this type of project, it's great if you “solve” your problem(s), but that's not a requirement — if you pick a natural and interesting project and demonstrate that you've thought hard about it, by documenting your ideas and partial progress, that's fine.
2. **Literature Survey.** You can pick a few papers from the scientific ML literature, read and understand them deeply, then write an original syn-

thesis and comparison of them. Imagine that you're writing a document that educates others.

3. **Solve Additional Problems.** From either K&V or other ML theory texts, choose a challenging set of problems and solve them carefully as in the course problem sets.
4. **Experimental ML Work.** It's also fine if you want to do an experimental/applied ML project; just remember to clearly relate it to course themes in some meaningful way.