

CPSC 440/550 Machine Learning (Jan-Apr 2026)

Project Proposal – due Sunday March 29 at **11:59pm**

As discussed in the syllabus, 40% of the final grade for the course is from:

1. CPSC 440: the best of the final exam or the research project.
2. CPSC 550: the average of the final exam and the research project.

If you're a 440 student and confident you want to only do the final, you don't need to hand in a proposal.

If you choose to do the research project, 10% of the corresponding part of the final grade is for this proposal. That is, if you're a 440 student doing only the research project, the research project proposal will count for 4% of your final course grade, and the research project writeup will be 36% of your final grade. If you're a 550 student, this proposal will be worth 2% of your final grade.

Projects are **strongly encouraged** to be done in [groups of 2 or 3](#). If you really want to do it alone, that's allowed, but **expectations will not be decreased** so it'll be harder, and it will probably be a worse learning experience for you as well – ML research is almost always collaborative. It's fine for 440 students and 550 students to work together.

These proposals are short and lightly graded: if you put some reasonable effort into thinking about the project, you'll probably get full proposal marks. The point is so that you don't wait to the last second to find a project group or think about what you want to do, and so we can check that the scope and topic of the project is suitable for the course.

If you'd like [early feedback](#) (before the official deadline at the end of March), please submit a proposal on Gradescope, then make a private Piazza post directing me to it. (You're also welcome to ask me about general ideas less formally, in person or on Piazza.) Otherwise, even if you submit it way in advance I probably won't look at it until the deadline.

It's okay with me for your project to overlap with a project in another simultaneous course, as long as you check with the other instructor. In that case, make sure you do an appropriate writeup of the work for each course. Having some overlap with your ongoing thesis research/similar is also fine (this is a good idea if you think of one that makes sense!), but make sure this is a relatively discrete project with a clear scope and will be self-contained enough that the course staff can read and understand it.

The proposal is not necessarily “binding” – research projects very often shift from the idea you started with, and it's even okay to do something totally unrelated to what you proposed. If so, though, you probably want to check in with me or a TA that the new project is in scope: if you hand in a final project we think is very out of scope for the course without having checked with us, you'll get a bad grade on it. If your change in ideas also corresponds to a change in teammates and if anyone is unhappy with how that worked out, talk to me.

The proposal can be done in any format you'd like; default L^AT_EX style with `\usepackage{fullpage}` is fine.

The proposal should be a [maximum of 2 pages](#); it's totally fine to be shorter if you can describe the plan concisely. The proposal should be written for the instructor and the TAs, so you don't need to introduce ML background that's covered in the course or that you would reasonably expect graduate students working in ML to know, but you should introduce any required non-ML background that someone in a CS/stat/ECE grad program might not know.

There is quite a bit of flexibility in terms of the type of project you can do, since there are many ways that people can make valuable contributions to research. The final deliverable will be a written report consisting of at most 6 pages (in the provided L^AT_EX format), with unlimited additional space for references and possible

appendices (which, as with e.g. NeurIPS reviewers, you shouldn't count on the graders reading). That report should emphasize (at least) one particular "contribution": what has doing this project added to the world?

The three main questions your project proposal needs to answer are:

1. What problem you are focusing on?
2. What do you plan to do?
3. What will the "contribution" be?

For the course project, negative results are acceptable (and often unavoidable). In that case, the paper should probably include something like "here's why we thought this thing we tried would work in this setting, here's us convincing you that it didn't, and here's our best understanding at why we think it failed."

Here are some standard project "templates"; you might want to roughly follow one of these templates, but **it's fine if your project mixes and matches between these project types, or does something else entirely.**

Some of the examples below include topics not covered in the course (like random forests): there is flexibility in the topic, but it should be closely related to ML, and ideally use tools or ideas that are covered in this course.

- **Application bake-off:** you pick a specific application (from your research, personal interests, maybe from Kaggle) or a small number of related applications, and try out a bunch of techniques (e.g., random forests vs. logistic regression vs. generative classifiers). In this case, the contribution could be showing that some methods work better than others for this specific application, and hopefully some idea why – or your contribution could be that everything works equally well/badly. Again, make sure this includes some ideas covered in this course, not just stuff from 340. Note that a project that exclusively takes a pre-existing dataset, runs standard methods on it, and reports the result is a less ambitious project; see the note at the end.
- **New application:** you pick an application where where people aren't using ML, and you test out whether ML methods are effective for the task. In this case, the contribution would be knowing whether ML is suitable for the task (and perhaps how to prepare the data and constructed features).
- **Scaling up:** you pick a specific machine learning technique, and you try to figure out how to make it run faster or on larger datasets. (For example: how do we apply kernel methods when n is very large?) Your improvements might be a new approximation, a distributed version, a smarter implementation, or so on. In this case, the contribution would be the new technique and an evaluation of its performance, or could be a comparison of different ways to address the problem.
- **Improving performance:** you pick a specific machine learning technique, and try to extend it in some way to improve its performance (for example, trying to efficiently use non-linearities within graphical models). In this case, the contribution would be the new technique and an evaluation of its performance.
- **Generalization to new setting:** you pick a specific machine learning technique, and try to extend it to a new setting (for example, making a graphical model version of random forests). In this case, the contribution could be the new technique and an evaluation of its performance, or could be a comparison of different ways to address the problem.
- **Perspective/survey paper:** you pick a specific topic in ML, read a large number of papers on the topic, then write a report summarizing what has been done on the topic, how they relate to one another (*not* just a paragraph summarizing the abstract of each paper), and what you think are the most promising directions of future work. In this case, the contribution would be your summary of the relationships between the existing works, and your insights about where the field is going.

- **Coding project:** you pick a specific method or set of methods (like independent component analysis), and build an implementation of them. In this case, the contribution could be the implementation itself or a comparison of different ways to solve the problem.
- A **reproducibility report** of a recent paper, as in the ML Reproducibility Challenge: the official challenge dates don't line up well with this course, but you can draw inspiration from the challenge and past submissions there. This is related to the "coding project," but with somewhat different aims.
- **Theory:** you pick a theoretical topic (like the variance of cross-validation or the convergence of stochastic gradient in non-smooth and non-convex setting), read what has been done about it, and try to prove a new result (usually by relaxing existing assumptions or adding new assumptions). The contribution could be a new analysis of an existing method, or why some approaches to analyzing the method will not work.

Any one of these project "types" is enough to get a reasonable project grade. Some are naturally more ambitious than others, though; grading will take into account the ambition of your scope as well as how well you execute it. For instance, a project that does an application bake-off with standard methods and doesn't do anything "wrong" but also doesn't do anything super-impressive in terms of analysis, implementation, etc will likely get somewhere perhaps in the high B range, although a very insightful or particularly thorough bake-off that's written up well could absolutely get an A+. On the other hand, a project that seems like it's on track for publication at a standard ML venue will be likely to get a high grade even if not everything is executed perfectly yet.

Teamwork for the project In addition to the project writeup and any code/etc that you submit, for projects done in groups, you will also be asked to each, separately write a brief summary of your view of what each teammate contributed to the project, including yourself, and whether you feel that the total amount of work was split in a way so that the whole team getting the same grade feels fair. Details will be with the project submission instructions. If, as you start working on the project, you feel there are issues with your group's functioning, you can talk to me about it (on a private Piazza post or in person).