

Machine Learning Projects - Fall 2022

How should a project look like

The main purpose of the project is to apply machine learning techniques that we covered in class to a problem of your choice. In the end, your project **should include something interesting** that has ideally not been done before in that particular context. Good ideas for the project may include:

- something that you need to get done for your master thesis
- applying machine learning techniques to an atypical data set
- testing effects of parameter changes, preprocessing techniques, or different algorithms
- testing how model size influences prediction time and accuracy for a specific problem
- developing methods for using unlabeled data in the context of supervised learning
- using unsupervised learning techniques to nicely visualize a complex data set

The project does not need to lead to a successful application, but in case it is unsuccessful, you need to include a detailed analysis of why that happened and suggestions for future work. A typical machine learning project is a series of several trial-error-analysis development cycles, you need to perform at least one such cycle.

There are three possible forms of the output:

- summary in the form of a conference paper (5 - 15 pages, containing introduction, methods, results and discussion of results) + a pointer to a GitHub/GitLab/whatever repository with all code and data sets
- interactive web application (there should be short summary of your methods and results on the webpage) + 1 page summary of methods
- get at least 20th place in a serious Kaggle competition (by serious we mean there is money involved and the competitions is not about Santa) + 1 page summary of methods

The projects are due 2 days before your final exam or **January 30**, whichever comes first. Your project report should be sent via classroom.

Project proposals

Project proposals are not required, but well prepared project proposal may earn you up to 5 bonus points and a constructive feedback. It is highly recommended to submit a project proposal. Proposals are due **January 8**.

Good project proposals will include at least the following:

- what problem are you trying to solve / formulation as a machine learning problem
- which data sets you are planning to use; are these data sets large enough to accomplish your goals?
- you may include proposed methods that you plan to use, but this is not required.