

L46: PROJECT SPECIFICATION

V1.7

Summary. The aim of the project is to provide students the opportunity to explore in detail machine learning systems related concepts encountered either in class, or during their own self-study. The selection of project topic and objectives are left open to each project group. (Although we offer some hints below and examples of projects on the website, in the event they are useful.) A project group can be comprised just of a single individual student, if they wish to work alone – or more conventionally a group of multiple students. It is recommended that students should select something about which they are intensely curious and passionate.

Project Types. Some representative categories of projects include the following:

- Replicating results observed in the literature or textbooks (*negative results are fine!*)
- Studying conventional wisdom discussed in textbooks or papers
- Building an application that leverages principles and algorithms seen in class (or that are related to topics discussed in the class)
- Comparing and contrasting approaches
- Exploring a novel ML system related idea or hypothesis (*note, novelty is not a requirement*)
- Investigating the capabilities of interesting hardware
- Performing a study to understand a (relevant) topic in greater depth
- Taking two methods often used in practice together (e.g., differential privacy and federated learning) and apply them together in a circumstance that often is not studied

Example Submissions. A select few projects from prior editions of this course are posted to the course website to act as examples to students. Please note, these tend to be good projects, but they also received a range of scores; they should not be treated as perfect projects that all were given the very best grades. They have flaws like any other project, that said – they all were well above a passing grade.

Submission. Submission of all material will be via Moodle. Each student will be assessed based on a written project report. This report should include the full names of the students who worked together (if the project is done in a group). Also included should be a description of how to access the project repository (repository described later in this document) as a separate document from the report itself. The report will be made available publicly (unless otherwise requested). A submission is a single file that is uploaded and will often be a zip file, or similar, containing the two mentioned documents and anything else necessary. Projects performed as a group need only be submitted by one of the students in the group.

Due Date. Submission Due: 16th January 2024, 12:00 noon

Report. The report is the primary method of assessment for the project (which accounts for 70% of the course grade) and is written by an individual student (if working alone) or the entire group (if working in a group). The report should seek to summarize and convey the outcomes and efforts behind the entire project.

Students should look towards typical scientific technical reports in the academic literature as examples to follow for style of writing and adopt the formatting instructions of the NeurIPS conference for the report (<https://nips.cc/Conferences/2020/PaperInformation/StyleFiles>). A maximum of 8 pages should be used with the space used for references not counting towards this limit.

If there is material you wish to include beyond this maximum page limit (e.g., additional experimental results, algorithm details, proofs, low-level dataset information) – you are welcome to include this in a report appendix section that is positioned at the end of the primary report after the report references. **Please note, using an appendix is not encouraged**, but if you feel there is necessary material to further support the main report content it can be included in this appendix. The appendix will not receive the same level of scrutiny as the primary text of the report and will mainly used to further clarify the primary report. The first 8 pages of the report must be self-contained, and not depend on the appendix to be understood.

The following is a non-exhaustive list of content the report should include (as applicable, not every list element will apply to all project types):

- A clearly articulated umbrella objective and purpose for the project, that is carefully broken-down into a series of inter-related sub-questions that build towards the objective. If necessary, more than just one hypothesis is included.
- Adequate (but not excessive) motivation and background information so that the reader has enough content to understand the reasons for, and importance of, the exploration and work performed. Crisp arguments motivate not only the objectives but also potential resulting findings and observations. The space given in the report to this type of content should not be excessive, and leave plenty of room for reporting and describing the technical work and thinking that forms of the core of the project.
- Where necessary, relationships are drawn with prior knowledge and the text in the report bridges between what is sought in the project and what is already known. The project is placed in context of prior and current on-going research or conventional practice/wisdom.
- A sound systematic methodology is employed. Where possible it builds on methods seen in the literature, or deviates from these when there is a reason to do so and discusses the reasons why. The steps within this method align well with the objectives of the project. If required, baseline methods that act as a point of comparison are offered and included in the discussion of results. Awareness of the impact of certain datasets or baselines (if applicable) not being included in a method are acknowledged.
- A description of the implementation developed and/or used within the project as part of the methodology. This should not be overly verbose but still explain sufficiently to indicate various frameworks and open-source software used. This content ties to the experiment discussion and helps explain how the experiments were performed.
- Experiment/analytical results, and how the results are reached, are clearly described. Careful consideration of how results are presented occurs, for example selecting a form that emphasizes the point being considered, or the alternative deciding to focus on, enabling results to be intuitive understood. Care is taken to enable, if reasonable effort applied, experimental repeatability.
- Convincing technical depth and solid fundamental understanding of techniques, algorithms and concepts used are provided. These are scoped based on relevance to the project. Limitations of these tools/techniques are appreciated and included in thinking and writing, relationships to well-known mainstream related techniques are acknowledged, as are very closely related work (if applicable).
- Implications of results/observations are clearly stated. An understanding of what claims *are* and what claims *are not* supported by individual and aggregate results is demonstrated.
- Overall conclusions and observations that wrap up the project, which is followed by appropriate references to the literature, datasets, implementations and online material.
- Care is taken in the presentation of the report and the experimental results presented. Example: figures are not just thrown into the report, but are tidily shown and sized appropriately (considering also font size).

Loose Marking Criteria. In addition to expectations implied by the report description above and each of the report content items (also see above), below is a list of issues that will also be considered during the marking process:

- Technical depth achieved and strength of understanding, few errors of understanding
- Convincing, well organized, clear **and succinct** technical writing
- Clear organization of purpose and aims, broken into smaller tractable units
- Degree of difficulty, **relative to the number of students contributing to the project**
- Systematic scientific methodology devised that corresponds to purpose and aims
- Results are well presented, connect strongly with the aims and implications are drawn
- Any purpose or aims not achieved, negative/surprising results discovered, key limitations are acknowledged and managed – appropriate implications and conclusions drawn
- Limitations and discussion show maturity; but more importantly entire work demonstrates awareness of how project fits with respect to class material (e.g., textbooks or similar where available), literature and recent developments being discussed in the community
- Project has a clear stolid motivation that justifies the effort applied

Repository. Each project should be performed using a shared repository to which the L46 team (lecturer and assisting TAs etc.) will be given access. This repository will contain experiment code and results, along with also a copy of the written report. The documentation of planning and decision making through the project is also required and must be stored within the repository as well. Please note, there is no need for this documentation to be overly detailed. It would be appreciated if additional documentation (e.g., README files and similar guidance) is provided in the repository to aid in understanding and navigation. It would also be helpful if you provide brief documentation regarding how to execute the code (even a portion of it) – if this applicable to your project. **Where needed, marks will be differentiated within groups using this repository as an input.**

Project Advice and Feedback. Starting from November 1st, drop-in sessions will be available until the end of the term each Wednesday at 11am. These will be with Filip Svoboda and email announcements provided when they are held in person (otherwise they will be offered via zoom <https://cam-ac-uk.zoom.us/j/5632858745?pwd=VEpONVdnYS9QNEUOUUpHM0dYLzBydz09>). A variety of project issues can be discussed, including project ideas, estimating GPU resource requirements, discussions about the project report – and so on. Similar sessions can be scheduled by appointment with Nic Lane.

FAQ

Do I need to perform novel research? This is absolutely not a requirement, it is completely fine to explore very well-known existing results, a paper you have read or even what has evolved to be current practice. At the same time this should not deter anyone who wishes to pursue the merits of a new idea they have had and are eager to explore.

Are there marks for implementation? There are not explicitly any marks assigned for the implementation. But the difficulty of implementation will be considered within the context of the overall difficulty of the project, **and if a tough implementation was necessary for the purpose of the project it will be rewarded.** That said, students should not feel obligated to re-invent the wheel and implement things that are already available unless, for example, again the purpose/goal of the project requires it – such as, a project that focuses on learning more about a certain model/algorithm/concept. When using existing code students should also be careful not to just blindly use libraries and toolkits, at all times they should develop a clear understanding of what this code is actually doing.

What if I get a negative result? This is not a problem, if you encounter some negative results you can use them but will need to properly frame them and explain why it is unexpected, while having also a solid methodological approach.

Do you prefer us to work as a group? Not at all, feel free to work in a group or work alone as an individual based on your preference. We will take into consideration the size of the group when judging appropriate project scope.

Are you going to judge my code? You will not be assessed for your style of coding or general code clarity and/or quality, do not be concerned about improving the presentation of your code before submitting. This is simply for us to have some level of knowledge of the implementation involved.

Can Part III students and MPhil students be in the same group? We request that these two cohorts of students not mix when forming groups. This is at the request of the administration team.