MIT CSAIL

6.8300/6.8301: Advances in Computer Vision

MIT COMPUTER VISION

Spring 2024

Project Overview & Team Formation

The final project is an opportunity for you to apply what you have learned in this class to a problem in computer vision that interests you. You can either pick one of the suggested project topics below (Option 1) or come up with your own project idea (Option 2).

Your project must investigate a scientific question. It is not enough to do a literature review, nor to directly reimplement a method without making any changes to it. A good project is one where you have scoped out a reasonable problem and demonstrated first experimental evidence that your proposal is viable. We expect at least 30 hours of work per team member on the project, and we are allocating time in the schedule to accommodate this (the workload will be consistent with problem set expectations during the equivalent time period). We expect two-person projects to be about twice the amount of work and content as one-person projects. 6.8300 projects will be graded to a higher standard than 6.8301 projects.

Although we recommend you work in teams, it is not required. If you do decide to form a team, please note that each team can have a maximum of 3 members (increased from last year) and you and your project partner must be registered in the same course. If you are registered in 6.8301, your project partner must additionally be in the same CI-M recitation section as you.

Further below you will find instructions regarding the project proposal, presentation, and report, and detailed grading rubrics.

Project Topic

Option 1: Choose one of the suggested project topics:

- ControlNet for Stable Diffusion
- Computer-Aided Diagnosis
- Stata Navigation
- 3D Printing Martian Rocks
- 3D Shape from 2D Surface Contours
- 3D Shape Reconstruction
- Reconstructing 3D Structure from Drawing & Sketches
- Imaging the Black Hole at the Center of the Milky Way
- Scene Segmentation
- Note: Visual Question Answering project removed due to most project attempts last year being over-ambitious in compute resources. Feel free to propose a well-scoped project on those lines.
- NeRF: Neural Radiance Fields
- An Automatic Metric for Surprisal
- Motion Magnification for Planets
- Gaze Coding for Developmental Research
- Video Super Resolution
- Visually Indicated Sounds

Option 2: Come up with your own project idea:

You could select a topic in computer vision that interests you and create your own project around it. A potential project could focus on an application or on creating new models or improving existing ones:

- Application: You could apply computer vision techniques to a specific application with your background and interest.
- Models: You could create new models or improve previous models or methods, then evaluate them systematically on standard image datasets to demonstrate their strengths and weaknesses.

Project Proposal Requirements

6.8301 Students:

The proposal is due Thursday, April 4 at 11:59 pm and must be submitted in PDF format once on Canvas for a CI-M grade and once on Gradescope for the TAs.

The proposal requires you to define the research space you'll investigate, and to describe the research you plan to do within that space. You will: identify a research gap; move from that problem/opportunity to a solution/goal; and come up with an approach for how that might be done. The technical staff will read your proposal for content, and you will receive more detailed feedback from your communication instructor and your peers to guide you in the revision and development process. The project proposal should be written and submitted as a group (not individually). Furthermore, independent of whether you choose Option 1 or 2 above, your proposal should answer the following questions:

- What is the problem/question that you will be investigating?

- Why is the problem interesting/important?
- What are the most relevant readings (2-4 papers)?
- What data will you use?
- What are the existing methods? Are their implementations available?
- What method or algorithm will you use, and why? What motivates your choice of this approach?
- What computing resources will you use to train and run your model(s)?
- How will you evaluate your results?
- Qualitatively, what kind of results do you expect (e.g. plots or figures)?
- Quantitatively, what kind of analysis (performance metrics, statistical tests etc.) will you use to evaluate and/or compare your results?
- If your approach is successful, what difference will this contribution make?

Formatting: The formatting of the proposal is up to you: it can take the form of sections separated by headings, each addressing one or several of the questions above, or it can be in a narrative format (i.e. without section headings). Whichever option you choose, your proposal should be written in full sentences, rather than bulleted lists. (Writing your proposal in full sentences allows you to prepare for the writing you'll do for the final report). The length of the proposal should be 1 single-spaced page, or 2 double-spaced pages, maximum.

Grading: The CIs grade will be on a 1-100 scale and will make up the CI-M portion of the project grade. You can find a rubric that describes how CI will assess the project here.

6.8300 Students:

The proposal is due Thursday, April 4 at 11:59 pm and must be submitted in PDF format on Gradescope for the TAs.

If you choose Option 1 (going with a suggested topic), then your or your team's project proposal only needs to specify which topic you picked (that's literally it). If you choose Option 2 (coming up with your own project idea), your or your team's proposal should answer the same questions and be of the same format and length as specified above for 6.8301 students. While we will not grade your proposal, we will provide feedback, in particular regarding the scope of your proposed project.

If you have a project partner, only one of you should submit the proposal on Canvas. In that case, please also list both partners' names and MIT Kerberos IDs (if you have one) in the proposal.

Project Report Requirements

The report should ideally be at least 3 and no more than 5 pages in length (excluding references). We will deduct points for excessively and unnecessarily long reports (a well-written concise report is better than a long and wordy one!). If you're working in a team, the expectation is that you'll likely need the full 5 pages for your report. Furthermore, if you're part of a team, you should write the report together but you must include a section that lists the individual contributions of each team member. We will apply a penalty if this section is missing for 2-person teams. The report should be structured like a research paper, starting with an abstract, followed by sections for Introduction, Related Work, Methodology, Experimental Results, Discussion, Conclusion, and ending with references. Some of the sections can be combined if you want (specifically, Introduction/Motivation & Related Work as well as Results & Discussion). We recommend you write the report in the CVPR format.

You should describe and evaluate what you did in your project, which may not necessarily be what you hoped to do originally. A small result described and evaluated well will earn more credit than an ambitious result where no aspect was done well. Be accurate in describing the problem you tried to solve. Explain in detail your approach, and specify any simplifications or assumptions you have made. Also demonstrate the limitations of your approach. When doesn't it work? Why? What steps would you have taken had you continued working on it? Make sure to add references to all related work you reviewed or used.

Submission: The report is due Tuesday, May 14 at 11:59 pm and must be submitted in PDF format. Late submissions will not be accepted. If you're part of a team, only one of you should submit the PDF. In that case, please list both partners' names and MIT Kerberos IDs (if you have one) at the top of the PDF and pay attention to the Gradescope team submission details.

Project Presentation Requirements

The presentation gives you a chance to provide us with a concise overview of the current state of your project and get us excited to your read your report! The presentation should highlight the key contributions your project is making and the insights you've gained while tackling a new research problem.

You should aim for your presentation to last 4 - 4.5 min and it should not exceed 5 min. After 4.5 min we will signal to you that you have 30 seconds left to wrap up your talk. At 5 min we will cut you off whether or not you have covered all your slides. If you try to continue with your presentation after we cut you off, we will have to apply a penalty. For teams of two, each student must present and you should aim to speak for approximately the same amount of time (~2 min each). Given the size of the class, unfortunately there will be no time for Q&A.

You will have to prepare your slides in Google Slides. We'll ask you to share with us a link to your slides so that at the time of the presentation, the grader can simply click on the link to pull up your slides. This way you can still make changes to your slides last minute (if need be). Details will be announced on Piazza.

Note: Whether you present at an earlier or later date will not affect your grade. It is perfectly fine if you do not have all (or any) of your results ready by the time you present. Make sure to describe the experiments you are planning on running and explain what you think your experiments will demonstrate.

For those who did not opt in for live presentations, the pre-recorded presentations are due the last day of live presentations, the same deadline as the final report: May 14th, 11:59 pm.

Project Grading Policy

The final project makes up 40% of the course grade. For 6.8300 students, the project is comprised of the report and the presentation, worth 30% and 10%, respectively. For 6.8301 students, the project is comprised of the report, the presentation, and the CI-M component, worth 22.5%, 7.5%, and 10%, respectively.

Report Grading Rubric: Please refer to this link. Presentation Grading Rubric (/100)

Note: Since presentations are holistic, we won't assign percentages to the individual sub-points.

- Clarity (/40)
  - Was your presentation easy to follow? Was it clear to someone who took 6.8300/1 but isn't familiar with the particulars of the project?
  - Were your slides clear, concise, and well organized? Did they get the message across?
  - Was your talk supported by meaningful visuals with a clear message? Were they easy to understand and legible?
  - Did you articulate your thoughts effectively? Did you speak clearly and at a good pace to ensure everyone can follow your talk?
  - Was the delivery of your presentation well organized? Did you cover all your slides within the time limit without rushing through them?
    - For teams of two: Was the coordination between joint speakers handled smoothly? Was the presentation divided up effectively and the presentation time well balanced between joint speakers?
- Content (/60)
  - Intro/Motivation & Related work: What are you trying to do? How is it done today, and what are the limits of current practice?
  - Methodology: What is your approach, how is it different from previous work, and why is it better?
  - Results & Discussion: How did you evaluate your approach or conduct your analysis? Which experiments did you run to support your claims? What other experiments will you run and why? What are the limitations of your approach or analysis? (As mentioned above, not having your results ready for the presentation is perfectly fine. In that case, please make sure to discuss the experiments you will run and what you think they will demonstrate.)