Final Project

Modern Computer Vision (MCV) - Winter 2026-27

Groups of 4-5 | Proposal: TBD | Report, code & video: TBD

What This Is

A small, self-contained piece of **original research** in computer vision. The point is to find a question that is genuinely open, however narrow, and answer it: propose something, build it, test it honestly, and report what you found. Scope should be modest. We are far more interested in a clear, well-executed idea than in an ambitious one that does not come together. Aim for work you would be comfortable showing to someone in the field, not for a grand result.

The topic is entirely yours to choose, as long as it lives in computer vision. Theory counts if it brings something new. Any tools, models, libraries, or pretrained checkpoints are fair game.

If you already have a direction in mind, run with it. If you would like a starting point, a few directions that tend to work well:

- Take a recent paper you find interesting and push on it: an improvement, an extension, a case it did not handle, a claim worth testing.
- Apply existing methods to a new kind of data, or in a setting where they have not been tried before.
- Build on research you are already involved in, in your lab or elsewhere.

Proposal (1 page, deadline TBD)

One page. Submit early if you can: if you do, we will make an effort to send feedback before the deadline. Cover:

- **Goal.** What you are trying to do, in one or two sentences.
- **Method.** How you plan to do it.
- **Prior work & your delta.** What has been done before, and what is new in your approach.
- **Resources.** What you will need, mostly compute (GPUs, memory, time).
- **Challenges & feasibility.** What could go wrong, and why this fits in the time you have.

Deliverables (deadline TBD, after the exam period)

- **Report**, 2-3 pages, in LaTeX. Standard structure: abstract, introduction (including related work), method, results, discussion. Include figures; a good figure often explains more than a paragraph.
- **Public GitHub repository** with your code, so results can be reproduced.
- **Video, 5 minutes.** Choose any format you like, including a recorded talk. We encourage an AI-assisted explanatory video in the style of 3Blue1Brown, or any other style. Its quality is an important part of the grade, so prepare and refine it carefully. No live talk is required.

Groups

- Groups must be **4-5 people**. We cannot accept groups smaller than 4.
- We prefer **5 over 4**, and we ask groups with room to take in students who have not found a group.
- We recommend **mixing graduate and undergraduate students**.
- A shared sheet will be provided to help form and join groups.

Source Disclosure

For the final project, students are free to use any sources, including AI and existing works, as long as all sources are disclosed and all external code, methods or models are cited.

Proposal: TBD | Final report, code & video: TBD, after the exam period