

Project Guidelines

Computer Vision (Spring 2012)

The goal of your project is to learn something new and, perhaps, to contribute something to the field of computer vision. Most projects should be done in groups of two or three. Larger projects could have a group of four, and in special cases groups of one are OK.

Proposal (*send description and meet by Mar 30*)

List the group members and a short (1/2 page) description of the goals of the project. If known, sketch a potential approach. You could switch later, but think it through as best you can now. As a guide to scope, you should expect to spend around 30-40 hours per person on the project. If you want to work by yourself, explain why it's not possible to work with a partner (e.g., because you're building on your own existing research). Otherwise, indicate what each group member will work on and how the project will be evaluated.

The proposal is not graded, but, if late, a 10% per day penalty will be assessed on the final project grade.

Deliverables: Send the proposal in an e-mail to Derek Hoiem *and* Ruiqi Guo by **March 29**. Then, meet with at least one of us to discuss by March 30. You must send the short written proposal and meet with one of us, for your proposal to be complete.

Paper (*due during final exam period*)

Use the style files here for your paper: <http://nips.cc/PaperInformation/StyleFiles>

The paper should be a minimum of 4 pages, and a maximum of 8 pages, including figures and references. Be sure to include an introduction that motivates the problem, discusses background material or related work, and summarizes the approach. The main section should provide details of the approach, experiments, and analysis. Finally, the conclusion section should summarize main conclusions, drawn from analysis and experiments.

Also, include a statement of individual contribution if there is more than one group member. If one person has done more or less than his/her fair share, please e-mail me separately. The project will be graded as a whole, but I may modulate individual grades based on this feedback.

Poster session (*during final exam period*)

Posters can be presented on poster-size paper or with stitched printouts of normal-sized pages. The posters should graphically illustrate the key ideas, approaches, and results of the project. I will split posters into two groups. Each poster will be presented for one hour, and each student will evaluate two other posters during the hour he/she is not presenting. If you would like to do a demo and need special equipment or setup, please let me know in advance.

Grading

Grades will be based on the quality of the project (originality, thoroughness, extent of analysis, etc.) and the clarity of the written report and poster presentation. Ideally, you will try something new or apply ideas from class to your domain or research. More will be expected of larger groups. It's great if you do something publishable, but you can also get full marks if you do something challenging or interesting that helps you learn something new.

Topic Ideas

Try something new or interesting: apply vision to a mobile robot, make a data-driven interface for image editing, organize home photos, do visual search in a home photo collection, reconstruct a 3D scene from multiple images, estimate material properties, etc.

Compare two or more approaches: Implement two approaches (e.g., for object recognition) and try to understand when one works better than the other. As part of this, you could try to create a better benchmark dataset (maybe a smaller prototype).

Some specific ideas:

- **Shadow detection:** Try to find cast shadows in outdoor images
- **Pedestrian detection:** Build a detector for standing/walking people
- **Multiview reconstruction:** Build a system to reconstruct an object or a scene from multiple images
- **Evaluation of object detection:** Run a state-of-the-art detector on a PASCAL VOC dataset and study (quantitatively and qualitatively) which factors make detection difficult.
- **Similar category differentiation:** Make a classifier that can tell the difference between dogs/cats or bicycles/motorbike, etc.
- **Material detection:** Try to classify materials on natural objects in images.
- **Action recognition:** Try to detect when somebody is performing a particular action in a video.
- **Tracking:** Try to track players and the ball in a sports video.
- **Photo organization:** Build a system that can organize your photos by the people in them.
- **Gender/age classification:** Given a face, try to predict the age and gender.
- **Fake or Real:** Try to predict whether an input image is natural or was generated by a computer.