UC Berkeley CS 280: Advanced Topics in Computer Vision

Instructor: Alexei Efros
GSIs: Lisa Dunlap
Suzie Petryk

Spring 2024
MO WE 12:30 – 2 pm
1102 Berkeley Way West

Website: https://cs280-berkeley.github.io/
Course email: cs280ucberkeley@gmail.com

Computer vision seeks to develop algorithms that replicate one of the most amazing capabilities of the human brain: inferring properties of the external world purely from light reaching the eyes and using this information to control actions in the world. We can determine how far away objects are, how they are oriented with respect to us, and in relationship to various other objects. We reliably guess their colors and textures, and we can recognize their semantic categories such as chairs and tables. We can segment out particular objects and track them over time, such as a basketball player weaving through the court. We can use the information extracted from images or video to manipulate objects in the world and navigate in environments while avoiding obstacles.

This course studies the concepts and algorithms behind the remarkable success of modern computer vision. We will build up from fundamentals and cover aspects of 2D vision, 3D vision, 4D vision, vision and action. A tentative list of topics includes:

1. geometry of image formation
2. radiometry of image formation, light fields, color vision
3. foundations of machine learning: neural networks, backpropagation
4. image processing
5. linear filters: low-pass, band-pass, spatiotemporal
6. sampling and multiscale image representations
7. neural architectures for vision: CNNs, recurrent networks, transformers
8. probabilistic models of images
9. generative image models and representation learning
10. challenges in learning-based vision: dataset bias
11. stereo correspondence and multiview geometry
12. understanding motion: optical flow, learning-based methods
13. understanding vision with language
14. failures of computer vision and where we should go next

Textbooks

The course requires the textbook *Foundations of Computer Vision* by Antonio Torralba, Phillip Isola, and Bill Freeman. Note that this is not publicly available yet; we will update you. All course materials will be posted on bCourses.

The following two books may additionally be useful as a reference:

1. David Forsyth and Jean Ponce: *Computer Vision : A Modern Approach*
2. Richard Szeliski: *Computer Vision: Algorithms and Applications* (available free online)

Grading

1. 35% homework: 4 assignments, first due in 3 weeks and each of the rest due in 2 weeks.
2. 35% exams: Includes in-class quizzes.
3. 30% final project: Presentations during the week of 5/6/2024.

Email Rules

1. Please use Ed Discussion for all communications, private or public, as much as possible. Others will benefit from answers and discussions on public questions.
2. Please email cs280ucberkeley@gmail.com only if the Ed channel does not work for you.
3. Please do not email us, instructor or GSIs, directly. We will not reply to any emails except at this official address.

Ed Discussion Rules

1. We encourage the use of Ed for discussions on conceptual and technical questions among classmates. There is no anonymity. Please be respectful to your classmates.
2. Ed will mostly be monitored by the GSIs, who would consult instructors whenever necessary.
3. Please post your questions for lectures, homework, and logistics in their dedicated folders.
4. Check Ed for already posted questions before posting a new one. Unnecessarily clogging up the forum makes it less usable for everybody.

Office Hours

Office hours will be posted on the course website.

Computational Resources

1. This course involves Python programming; you are free to use alternative packages.
2. For neural network assignments, we will use PyTorch http://pytorch.org/. Knowledge of python will be assumed. GPU resources and usage instructions will be given in the assignments.
3. CS280 students can use their existing EECS Windows accounts in EECS instructional labs, and they can request new accounts (for non-majors) or additional access to Instructional resources by following the instructions about 'named' accounts in
   http://inst.eecs.berkeley.edu/connecting.html#accounts

   They can logon remotely and run it on some of our servers:
   http://inst.eecs.berkeley.edu/connecting.html#labs