EE P 596 Computer Vision: Classical and Deep Methods
Fall 2021

Overview:
Computer vision has made tremendous progress over the past decade in solving problems such as image classification, object detection, semantic segmentation, and 3D reconstruction. A major paradigm shift occurred in 2012, when the technique of deep learning began to replace hand-crafted features with automatically learned ones. The purpose of this class is to introduce students to both classic techniques (pre-2012) as well as modern ones (since 2012), along with fundamental concepts that underly both. The class is aimed at first-year (or later) graduate students who have little to no familiarity with image processing or computer vision.

Weekly meetings:
Tuesdays, 6:00pm – 9:50pm
- Lecture (6:00pm – 8:00pm, approximately)
- Lab (8:00pm – 9:50pm)

Instructor: Stan Birchfield, sbirchfield@nvidia.com (put "EEP596" in the subject line)
TA: Aniket Rege, aniketr@uw.edu
Office hours: TBD

Textbooks, Software:
- Stan Birchfield, *Image Processing and Analysis*, Cengage, 2018 (recommended).
- Python 3.x, NumPy, PyTorch

Grading:
Homework 40%
Project 40%
Quizzes 10%
Class participation 10%
No midterm; no final exam

Prerequisite skills:
- Linear Algebra
- Probability and Statistics
- Calculus (basic)
- Programming (in some imperative language)

Schedule:
We will have 10 classes, from Oct. 5 through Dec. 7. The last class (during final exam week, Dec. 13-17) will be project presentations.

Topics covered:
- Image formation and acquisition
- Color spaces and transforms
- Morphological operators, warping, and interpolation
- Filtering in spatial and frequency domains

- Model fitting and classification
- 3D vision and geometry of multiple views
- Convolutional neural networks
- Network architectures, loss functions, regularization, backpropagation
- Object detection
- Semantic segmentation

Projects:
The purpose of the projects is to allow students the opportunity to explore an application of the material to a problem that is of interest to them. Students may work individually or in groups of two. (With instructor permission and proper justification, groups of three may be allowed.) Students are expected to propose a topic of their own choosing and bring the project to completion. It is recommended that students scope their project accordingly, ensuring that it is neither too easy nor too ambitious. The final project grade will be determined by a brief oral presentation and written report.

Detailed schedule (tentative):
Week 1. 10/5. Introduction; Image transformations; Binary image processing. **(HW#1 due 10/10)**
Week 2. 10/12. Convolution; Spatial domain processing; Convolution revisited. **(HW#2 due 10/17)**
Week 3. 10/19. Deep neural networks; Convolutional networks; Classification. **(HW#3 due 10/24)**
Week 4. 10/26. Backpropagation; Optimization; Loss functions. **(HW#4 due 10/31)**
Week 5. 11/2. Backbone architectures; Batchnorm. **(HW#5 due 11/7)**
Week 6. 11/9. Frequency-domain processing; Edges and features. **(HW#6 due 11/15)**
Week 7. 11/16. Compression; Segmentation; Model fitting. **(HW#7 due 11/21)**
Week 8. 11/23. Color; Stereo; Motion. **(Project proposals due 11/28)**
Week 9. 11/30. Camera calibration; Geometry of multiple views. **(Project updates due 12/5)**
Week 10. 12/7. Object detection; Unsupervised learning. **(Project updates due 12/12)**
Week 11. 12/14. **Final presentations. (Project reports due 12/16)**