

EE P 596 Computer Vision: Classical and Deep Methods Fall 2025

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:

Thursdays, 6:30pm – 9:50pm, ECE 125

Instructor: Stan Birchfield, sbirchfield@nvidia.com (put “EEP596” in the subject line)

TA: Ryan Chiu-Kam Chu, Zhihao Zhou

Office hours: TBD

Textbooks, Software (recommended):

- Stan Birchfield, *Image Processing and Analysis*, Cengage, 2018.
- Rick Szeliski, *Computer Vision*, Springer, 2010.
<https://szeliski.org/Book/1stEdition.htm>
- Jeremy Howard and Sylvain Gugger, *Deep Learning for Coders*
<https://github.com/fastai/fastbook>
They also have a free online course: <https://course.fast.ai/>
- E. Stevens, L. Antiga, T. Viehmann, *Deep Learning with PyTorch*, Manning, 2020.
<https://pytorch.org/assets/deep-learning/Deep-Learning-with-PyTorch.pdf>
- I. Goodfellow, Y. Bengio, A. Courville, *Deep Learning*, MIT, 2016.
<https://www.deeplearningbook.org/>
- M. A. Nielsen, *Neural Networks and Deep Learning*, Determination Press, 2015.
<http://neuralnetworksanddeeplearning.com/>
- Python 3.x, NumPy, PyTorch

Grading:

Homework 50%

Project 30%

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 regular classes, plus project presentations (during final exam week).

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 function optimization, backpropagation
- Object detection, semantic segmentation
- Transformers, unsupervised learning, foundation models

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.

Collaborating:

Students are encouraged to work together and help each other. But you should turn in your own work. You should be able to explain anything you turn in; otherwise, you are bypassing the learning process. The same goes for using LLMs: you are free to use them to aid the learning process, but you should not turn in anything that you cannot explain.

Detailed schedule (tentative):

Week 1. 9/25. Introduction; Image transformations; Binary image processing; Python. **(HW#1 due 10/5)**

Week 2. 10/2. Smoothing; Differentiating; Convolution. **(HW#2 due 10/12)**

Week 3. 10/9. Classification; Deep learning; Convolutional neural networks. **(HW#3 due 10/19)**

--- No in-person class on 10/16. Instructor will be out of town that week. ---

Week 4. 10/16. Frequency-domain processing; Edges; Features. **(HW#4 due 10/26)**

--- No in-person class on 10/16. Instructor will be out of town that week. ---

Week 5. 10/23. Optimization; Backpropagation; Batch norm. **(HW#5 and project proposals due 11/2)**

Week 6. 10/30. Semantic segmentation. Backbones. Network architectures. **(HW#6 due 11/9)**

Week 7. 11/6. Color; Stereo; Motion. **(HW#7 and project updates due 11/16)**

Week 8. 11/13. Projective geometry; Geometry of multiple views; NeRF.

Week 9. 11/20. Object detection; Model fitting; Unsupervised learning.

--- [Thanksgiving break] ---

Week 10. 12/4. Transformers; LLMs and VLMs; Foundation models.

Week 11. 12/11. **Project presentations. (Project reports due 12/13)**

All assignments are due at 11:59pm, with an automatic 8-hour grace period.