

CSE 327: Fundamentals of Computer Vision

<https://www3.cs.stonybrook.edu/~hling/Teaching/24F-327/F24-327.html>

Fall 2024

Basic Course Information

- **Instructor:** Haibin Ling (haibin.ling@ stonybrook.edu)
- **Time:** Mon/Wed 3:30pm - 4:50pm
- **Office hour:** Wed 9:30pm-11:30am or by appointment
- **Room:** Frey Hall 205
- **TA:** TBA

Course Description

The aims of this course are to provide an understanding of the fundamentals of Computer Vision and to give a glimpse in the state-of-the-art, at a moment when the field is achieving "critical mass" and has significant commercial applications. Apart from basic theory we will look at applications of Computer Vision in Robotics, Graphics and Medicine.

Lecture Topics (roughly ordered by lecturing date)

Note: contents may be up to change to reflect the fast evolution of the frontier of computer vision.

- 1. Image Formation**
 - Basic facts about light
 - Camera anatomy, matting
- 2. Image Noise**
 - Modeling image noise
 - Convolution, image smoothing, pyramid
- 3. Image Features**
 - Points, corners, edges
 - Scale and orientation
- 4. Model Fitting**
 - Lines, curves, Hough transform
 - Deformation, RANSAC
- 5. Perspective Projection**
 - Homogeneous coordinates
 - Image warping, mosaics
- 6. Multiview Geometry**
 - Stereo viewing and reconstruction
 - 3D range scanning
- 7. Motion**
 - Motion Capture
 - Tracking in 2D and 3D
- 8. Illumination**
 - Shading, shadows
- 9. Segmentation**
 - Reflectance properties
 - Grouping, superpixels
 - U-Net, semantic segmentation
- 10. Object Recognition**
 - Object representation
 - PCA for image patches
 - Classifiers
 - Category classification
- 11. Deep Learning Practice**
 - Convolutional neural networks
 - Architectures and applications
 - Pre-training
 - Data augmentation
- 12. Object Detection**
 - Face detection and pedestrian detection
 - Generic object detection
 - Deep learning-based solution
- 13. Advanced Topics (topics vary)**
 - Computer vision for science
 - Simultaneous localization and mapping

Intended Audience

This course is intended for undergraduate students with interests in all areas of Visual Computing, such as Computer Vision, Computer Graphics, Visualization, Biomedical Imaging, Robotics, Virtual Reality, Computational Geometry, Optimization, Deep Learning, HCI. Prerequisites include a foundation in Linear Algebra and Calculus, and the ability to program. We will be programming in Python (OpenCV, NumPy, SciKit).

Textbook and References

- **Computer Vision: Algorithms and Applications** (Texts in Computer Science), by Richard Szeliski, Springer; 1st Edition. ISBN-10: 1848829345, ISBN-13: 978-1848829343, <http://szeliski.org/Book/>
- **Deep Learning**, Goodfellow and Yoshua Bengio, Aaron Courville, 2016, MIT press. <http://www.deeplearningbook.org/>
- **Mastering OpenCV with Practical Computer Vision Projects**, by Daniel Lélis Baggio, Shervin Emami, David Millán Escrivá, Khvedchenia Ievgen, Naureen Mahmood, Jason Saragih, Roy Shilkrot, Packt Publishing, ISBN-10: 1849517827, ISBN-13: 978-1849517829
- **Vision: A Computational Investigation into the Human Representation and Processing of Visual Information**, by David Marr, The MIT Press, ISBN-10: 0262514621, ISBN-13: 978-0262514620.
- **Papers** assigned in the class.

Grading

- There will be 5-6 homeworks, a final, a midterm and 4-5 10-min quizzes.
 - Homeworks will be 40%, quizzes 10%, the midterm 25%, and the final 25%.
 - Weights are approximate and subject to change.
- There will be 3 free late days for the semester. After that there will be 10% penalty per day.
- You are expected to do homeworks by yourself. Even if you discuss them with your classmates, you should turn in your own code and write-up. Do not share your code! Final projects can be done by one or two people. Two people projects will be scaled accordingly.
- You can do a project instead of the final exam. Projects will be done in up-to 2 people teams, and will require a significant programming and documentation effort. This will probably be much more work than doing the final exam. Two people projects will be scaled accordingly.
- You can have one sheet of paper with notes in the midterm and final.

Don't Cheat!

Cheating on anything will be dealt with as academic misconduct and handled accordingly. I will not spend a lot of time trying to decide if you actually cheated. If I think cheating might have occurred, then evidence will be forwarded to the University's Academic Judiciary and they will decide. If cheating has occurred, an F grade will be awarded. Discussion of assignments is acceptable, but you must do your own work. Near duplicate assignments will be considered cheating unless the assignment was restrictive enough to justify such similarities in independent work. Just think of it that way: Cheating impedes learning and having fun. The labs are meant to give you an opportunity to really understand the class material. If you don't do the lab yourself, you are likely to fail the exams. Please also note that opportunity makes thieves: It is your responsibility to protect your work and to ensure that it is not turned in by anyone else. No excuses! The University has a relevant policy:

"Each student must pursue his or her academic goals honestly and be personally accountable for all submitted work. Representing another person's work as your own is always wrong. Any suspected instance academic dishonesty will be reported to the Academic Judiciary. For more comprehensive information on academic integrity, including categories of academic dishonesty, refer to the academic judiciary website at <http://www.stonybrook.edu/uaa/academicjudiciary/>."

Anything that you hand in, whether it is a written problem or a computer program, must be written by you in your own words. If you base your code on any other solution written by someone else or you use an AI to generate any part of your code, you are cheating.

Disability Note

If you have a physical, psychological, medical or learning disability that may impact on your ability to carry out assigned course work, I would urge that you contact the staff in the Disabled Student Services office (DSS), Room 133 Humanities, 632-6748/TDD. DSS will review your concerns and determine, with you, what accommodations are necessary and appropriate. All information and documentation of disability is confidential.