CSE527: Introduction to Computer Vision

https://www3.cs.stonybrook.edu/~hling/Teaching/20S_527/index.html

- Instructor: Prof. Haibin Ling
- Spring 2020: Mon/Wed 5:30pm – 6:50pm, Engineering 143, West Campus
- Office hour: Tue 2pm-4pm, NCS 259
- Email: hling AT cs.stonybrook.edu

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. Topics in this course:

1. Image Formation
- Basic facts about light
- Anatomy of a camera
- Matting

2. Image Noise
- Modeling image noise
- Convolution
- Smoothing images
- Image pyramids

3. Image Features
- Point features, corners
- Edge features
- Scale
- Orientation

4. Model Fitting
- Lines, curves
- Hough transform
- Deformation
- RANSAC

5. Perspective Projection
- Homogeneous coordinates
- Image warping, mosaics

6. Multiple View Geometry
- Stereo viewing and reconstruction
- 3D range scanning

7. Object Recognition
- Object representation
- PCA for image patches
- Classifiers
- Object categories

8. Deep Learning
- Convolutional neural networks
- Architectures
- Applications

9. Deep Learning Practice
- Pre-training
- Data augmentation

10. Illumination
- Shading, shadows
- Reflectance properties

11. Deep Generative Models
- Autoencoders, VAEs
- Generative adversarial networks
- Application in computer vision

12. Segmentation
- Grouping, superpixels
- U-Net
- Semantic segmentation

13. Motion
- Motion capture
- Tracking in 2D and 3D
- Recurrent neural networks
- Action recognition

14. Advanced Topics
- Simultaneous localization and mapping
- Vision and graphics
- Medical image analysis

Intended Audience:
This course is intended for graduate students with interests in all areas of Visual Computing and Machine Learning, 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).

Grading:
There will be homeworks, two midterms and a final. Homeworks will be 60%, each midterm 10%, and the final 20%. Weights are approximate and subject to change. You are expected to do homeworks (5-6) by yourselves. Even if you discuss them with your classmates, you should turn in your own code and write-up. Do not share your code! There will be 4 free late dates for the semester. After that there will be 10% penalty per day.

- You can do a project instead of the final homework. 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 homework. Two people projects will be scaled accordingly.
- Midterm 1 date: Mar 2nd 2020
- Midterm 2 date: Apr 8th 2020
- Final date: May 13th 2020, 8:30pm-11pm. Projects due May 14th 2020.
- You can have one sheet of paper with notes in the exams.

Textbook:

- Computer Vision: Algorithms and Applications, by Richard Szeliski (2010) Main text, available online.
- Computer Vision: A Modern Approach, by David Forsyth and Jean Ponce (2012)
- Computer Vision: Models, Learning, and Inference, by S. Prince, Cambridge Univ. Press, 2012
- Introductory Techniques for 3-D Computer Vision by E. Trucco and A. Verri, (1998)
- Readings from these books and notes for all topics will be posted on blackboard

Academic misconduct policy:
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/

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.

Class discussion:
We will be using Piazza for class discussion. Rather than emailing questions to the teaching staff, I encourage you to post your questions on Piazza so everyone may benefit from the answers. The Piazza webpage of this class is https://piazza.com/stonybrook/spring2020/cse527/home.