Machine Vision

2024 Fall

Humans are good at perceiving the world from visual input alone. This comes so easily to us that we underestimate how difficult perception it is, and how hard it is for machines. Machine vision aims to enable machines to see as humans do. Despite the tremendous progress we have made, the problem is still far from clearly understood and reliably solved. This course will cover both the classic work based on geometry and physics, as well as the new ones based on low dimension model, nonconvexity, and deep learning.

Administrative

Instructor: Yuqian Zhang
Lecture Time: TBA
Office Hours: TBA

Textbooks and Resources

There is no required textbook, but students should refer to the following textbooks for reference.

- Computer Vision: Algorithms and Applications. Richard Szeliski (available online)
- High-Dimensional Data Analysis with Low-Dimensional Models: Principles, Computation, and Applications. John Wright, Yi Ma (available on Canvas, do not share out of class)

Prerequisite

Basic understanding of Fourier transformation, convex optimization

Tentative Syllabus (subject to change)

Part I: Visual Data Models

1. **introduction**
2. **image formation – geometry**

3. **image formation – color**
4. **basic image processing – convolution, Fourier, resizing**
5. **reconstruction 1: binocular stereo**
6. **reconstruction 2: epipolar geometry**
7. **correspondence 1: corner, blob**
8. **correspondence 2: SIFT**
9. **grouping: k-means and graph cut**
10. **illumination cone and subspace model, image manifold**

Part II: Learning Algorithms

1. **machine learning overview**
2. **convex optimization overview**
3. **sparse coding and recovery**
4. **sparse problems: low rank, rpca**
5. **nonconvex geometry: symmetry, eigenvector, low rank, dictionary learning**
6. **nonconvex algorithm: blind deconvolution, k-means, second algorithm**
7. **perceptron**
8. **neural network**
9. **transfer learning**
10. **graph-NN, GAN, deep learning**

Grading Policy

The course will be graded based on homework (20%), paper summaries (15%), course project (50%), and peer review (15%).

- **Homework**: There will be two assignments this semester, and must be done individually.
- **Paper summaries**: Select a paper of interest and submit a summary of the paper.

- **Project**: Feel free to work on any topic related to the course material, provided the project is well-executed. The project should be conducted in groups of 2-4. If you are doing the project by yourself, you will need approval from the instructor. The project itself will have three main deliverables: the project proposal (1-2 page), the final report and a project presentation.
- **Peer review**: Review projects of your classmates, ask insightful questions, give helpful suggestions.

Paper/Project Review Guideline

Read the paper: Its important to carefully read through the entire paper and to look up any related work and citations that will help you comprehensively evaluate it. Be sure to give yourself sufficient time for this step.

While reading, consider the following:

- Objective of the work: What is the goal of the paper? Is it to better address a known application or problem, draw attention to a new application or problem, or to introduce and/or explain a new theoretical finding? A combination of these? Different objectives will require different considerations as to potential value and impact.
- Strong points: is the submission clear, technically correct, experimentally rigorous, and reproducible, does it present novel findings (e.g. theoretically, algorithmically, etc.)?
- Weak points: is it weak in any of the aspects listed?

Be mindful of potential biases and try to be open-minded about the value and interest a paper can hold for the entire community, even if it may not be very interesting for you.

Answer three key questions for yourself, to make a recommendation to Accept or Reject:

- What is the specific question and/or problem tackled by the paper?
- Is the approach well motivated, including being well-placed in the literature?
- Does the paper support the claims? This includes determining if results, whether theoretical or empirical, are correct and if they are scientifically rigorous.

Write your initial review, organizing it as follows:

- Summarize what the paper claims to contribute. Be positive and generous.
- List the strong and weak points of the paper. Be as comprehensive as possible.
- Clearly state your recommendation (accept or reject) with one or two key reasons for this choice.
- Provide supporting arguments for your recommendation.

- Ask questions you would like answered by the authors to help you clarify your understanding of the paper and provide the additional evidence you need to be confident in your assessment.
- Provide additional feedback to improve the paper.

Deliverables

- **1st homework**: end of first month
- **2nd homework**: mid of second month
- **project proposal**: end of second month
- **paper summary**: beginning of third month
- **project presentation**: end of third month
- **project final report due**: end of the semester

Selected Papers

- Unsupervised Learning of Depth and Ego-motion from Video. Tinghui Zhou, Matthew Brown, Noah Snavely, David Lowe. CVPR, 2017.
- Generative Adversarial Nets. Goodfellow, Ian, et al. NeurIPS, 2014.
- Shape, Illumination and Reflectance from Shading. J. T. Barron and J. Malik. TPAMI, 2015.
- A comparative study of energy minimization methods for markov random fields with smoothness-based priors. Szeliski, R., Zabih, R., Scharstein, D., Veksler, O., Kolmogorov, V., Agarwala, A., Tappen, M. and Rother, C. In TPAMI, 2008.
- Semantic Image Segmentation with Deep Convolutional Nets and Fully Connected CRFs. Liang-Chieh Chen*, George Papandreou*, Iasonas Kokkinos, Kevin Murphy, and Alan L. Yuille. ICLR, 2015
- Lambertian Reflectance and Linear Subspaces. Ronen Basri, David Jacobs. TPAMI, 2003.