

CAP 6412 Advanced Computer Vision

Website:

<http://www.cs.ucf.edu/~bgong/CAP6412.html>

Jan 12, 2016

Introductions

- Instructor: Boqing Gong

bgong@crcv.ucf.edu

- Office hours: Tuesday 4:30—5:30pm at HEC 214

Email --- the best way to reach me

- Put **[CAP6412]** in subject line
- Summarize message in subject line
- Ex: [CAP6412] Meeting request: Thursday (Jan 14) 4:30pm?

Today


- Course overview
- Logistics, requirements

What is computer vision?

- Computer programs to recognize images
- Understand image sequences or videos
- Visual scenes around us
- Imitate human capabilities to understand scenes
- Perception, visual perceiving of the 3D world

What is computer vision?

- Duplicate the capabilities of human vision on machines
- “Acquire, process, analyze, and understand images, videos, 3D data, and other types of high-dimensional data of the real (visual) world” --- Wikipedia.org



Research on computer vision

- Topics

Scene understanding

Medical image analysis

Human-centric CV

Human action recognition,
Face

3D Scene reconstruction

Document/character recognition

Low-level computer vision

- Methods, tools, models

Graphical models (HMM, GMM, etc.), Neural networks, OpenCV

- Related disciplines

- Graphics, Machine learning, HCI, Image processing, Medicine, Statistics

Research on computer vision

- Topics

3D modeling and reconstruction

Stereo, structure from motion

Biomedical image analysis

Computational photography

Face and gestures

Low-level vision (edge, color, etc.)

Motion and tracking

Optimization

Statistical learning, interactive learning

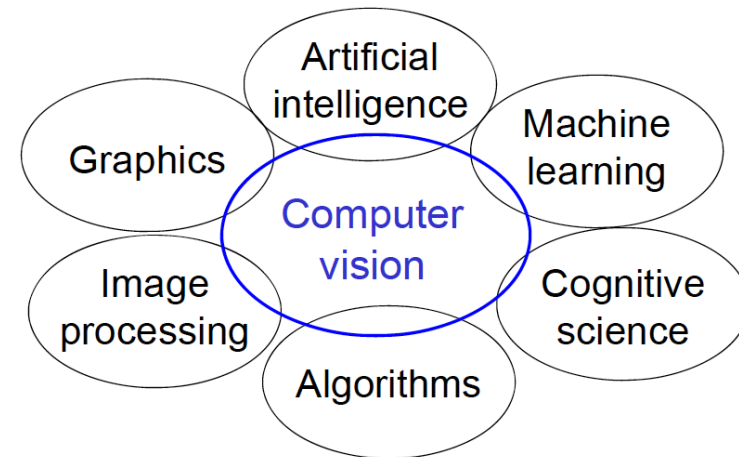
Recognition (of objects, attributes, scenes)

Segmentation

Video: action, tracking, surveillance

Vision for X

- Related disciplines



Disseminating computer vision research

- Conferences

- Journals, transactions

Disseminating computer vision research

- Conferences
 - CVPR, ICCV, ECCV, NIPS
 - BMVC, WACV, ACCV, etc.
- Submission, **double-blind** review, rebuttal, decision
- More noisy reviews
- Publication in about 6 months
- Journals, transactions
 - IEEE T-PAMI, IJCV
 - IEEE T-IP, T-MM, CVIU, etc.
- Submission, **single-blind** review, response, review, response,, decision
- Publication in >1 years

*arXiv.org: e-prints, **self-archived***

Digesting computer vision research

- #1 course objective

- Familiar with frontier problems, research groups, and researchers
- Understand state-of-the-art techniques and tools
- Familiar with computer vision datasets (benchmarks)
- Understand evaluation methods and metrics
- Improve critical analysis capabilities

How to read research papers? (1)

- Ask questions before you start

How to read research papers? (1)

- Ask questions before you start
 - What are the motivations for this paper?
 - What problem is being solved?
 - What is the proposed solution?
 - What experiments are designed to test the solution?
 - What are the evaluation methods and metrics
 - What are the contributions?
 - What are the future directions?
 - How is the paper related to what you previously knew?
 - How is the paper related to other works?

How to read research papers? (2)

- Repeat: decreasing marginal gain until convergence?
- Get the general idea first
 - What problem is being solved?
 - What are the main steps in the solution?
 - What is the high-level idea of the solution?
- Draw analogies to what you know well
- Read in order:
 - Abstract, conclusion, problem statement if it exists, intro, approach, experiments
- Divide and conquer
- How to read research papers by Dr. Mubarak Shah: <http://crcv.ucf.edu/people/faculty/HowToRead.html>
- The three-pass approach by Dr. Srinivasan Keshav: <http://ccr.sigcomm.org/online/files/p83-keshavA.pdf>

How to review research papers? (1)

- Write a review on the paper --- #2 course objective
- Summarize a take-home message
- Recall the main strengths of the paper
 - Novelty
 - Technical correctness
 - Clarity
 - Experimental evaluation

How to review research papers? (2)

- Write a review on the paper --- #2 course objective
 - Summarize a take-home message
 - Recall the main strengths of the paper
 - Be picky: weakness of the paper
 - Lack of novelty: comparing to prior work
 - Lack of clarity: language, organization, presentation
 - Technical errors: rare
 - Mismatched experiment design
 - Insufficient experiments
 - Unfair comparison with other methods
 - **Justify your comment**

How to review research papers? (3)

- Write a review on the paper --- **#2 course objective**
- Summarize a take-home message
- Recall the main strengths of the paper
- Be picky: weakness of the paper
- Overall rating (adapted from NIPS reviewer instructions)
 - **0:** Top 10% of the papers I have read, an excellent paper, a strong accept.
 - I will fight for acceptance. I will consider not reviewing papers for XX if this is rejected.
 - **1:** Top 50% of accepted NIPS papers, a very good paper, a clear accept.
 - I vote and argue for acceptance.
 - **2:** Good paper, accept.
 - I vote for acceptance, although would not be upset if it were rejected.
 - **3:** Marginally above the acceptance threshold.
 - I tend to vote for accepting it, but leaving it out of the program would be no great loss.
 - **4:** Marginally below the acceptance threshold.
 - I tend to vote for rejecting it, but having it in the program would not be that bad.
 - **5:** An OK paper, but not good enough. A rejection.
 - I vote for rejecting it, although would not be upset if it were accepted.

How to review research papers? (4)

- Write a review on the paper --- #2 course objective
 - Summarize a take-home message
 - Recall main strengths of the paper
 - Be picky: weaknesses of the paper
 - Overall rating (0, 1--5)
 - Explain your rating: how you weigh the strengths and weaknesses
 - Brainstorm future directions
 - ??

How to review research papers? (4)

- Write a review on the paper --- #2 course objective
 - Summarize a take-home message
 - Recall main strengths of the paper
 - Be picky: weaknesses of the paper
 - Overall rating (0, 1--5)
 - Explain your rating: how you weigh the strengths and weaknesses
 - Brainstorm future directions
 - New solutions
 - New solutions extending or inspired by the paper's solution
 - Open problems
 - Other problems that could benefit from the paper

How to present research papers (in class)? (1)

- Make good presentations --- **#3 course objective**
- Know your audience: fellow graduate students with good background
- Be aware of your time: Tues, Thurs, 3pm—4:15pm
- Adapt the presentation goal: explain and discuss the paper
- *Assume no one in the class has read the paper before*

How to present research papers? (2)

- Make good presentations --- **#3 course objective**
- Title, authors (full name), authors' institutes, your name and email
- Motivation of the research (1—2 slides)
- Problem statement
 - What is being solved?
 - Why is it an important problem?
- Main contributions of the paper
 - Studied a new and important problem
 - Proposed a novel approach
 - Improved or extended existing methods
 - Compared several popular methods
 - Explored a variety of use cases (many datasets of different kinds)
 - Presented new theories
 - Introduced new methodologies or tools to computer vision

How to present research papers? (3)

- Make good presentations --- **#3 course objective**
 - Title, authors (full name), authors' institutes, your name and email
 - Motivation of the research (1—2 slides)
 - Problem statement (1—2 slides)
 - Main contributions of the paper
 - Approach outline (1 slide)
 - Details of the proposed approach
 - Experiments
 - Data, features, baselines, evaluation metrics, results
 - Related work (1—3 slides)
 - Conclusion: take-home message (1—2 slides)

How to present research papers? (4)

- Make good presentations --- #3 course objective
- Title, authors (full name), authors' institutes, your name and email
- Motivation of the research (1—2 slides)
- Problem statement (1—2 slides)
- Main contributions of the paper
- Approach outline (1 slide)
- Details of the proposed approach
- Experiments
- Related work (1—3 slides)
- Conclusion: take-home message (1—2 slides)
- Strengths & weaknesses of the paper (1—2 slides)
- Overall rating & why (how you weigh the strengths and weaknesses) (1 slide)
- Future directions (1—3 slides)

Syllabus (1)

- **Overarching objective**

- Introduce the frontiers and trends in computer vision
- Prepare students for computer vision research

- **Through taking this course**

- Gain in-depth understanding of the state-of-the-arts in computer vision
- Write high-quality paper reviews
- Improve presentation skills
- Sharpen programming and critical analysis skills
- **Identify new research directions**

Syllabus (2)

- **Topics**

- Convolutional neural networks: basics and applications
- Vision and language: image captioning, visual question answering, etc.
- Representations: low-level, middle-level (attributes, parts), CNN layers, etc.
- Video: action, tracking, surveillance, etc.
- Statistical models and learning
- Image retrieval and matching
- Human-centered computer vision: face, pose, pedestrian, ego-centric, feedback, humans in the loop, active learning, etc.
- Visual saliency: saliency, gaze, object proposal, etc.
- For more topics, see **Syllabus** and **Papers** on course homepage:
<http://www.cs.ucf.edu/~bgong/CAP6412.html>

Submit up to three preferred topics by Jan 14th 2016 (Thurs), 1pm (UCF Financial Aid)

Syllabus (3)

- **Grading policy**

Reports (writing paper reviews): 25%

- Two reviews per week (**due resp. Tuesday, Thursday, 3pm**)
- Due on Jan 14th 2016 (Thurs), 1pm: up to three preferred topics

Paper Presentations: 20%

- 1—2 presentations every class (volunteers?)

Discussion and Attendance: 20%

Programming Projects: 25%

- Potential workshop & conference submissions
- Project presentation: 10%

Final project presentation: 1:00pm—3:50pm, April 28th.

Letter grades: 90—100: A, 80—89: B, 70—79: C, 60—69: D, below 60: F

Syllabus (4)

- Late homework policy
 - Three late days in total for all reports and projects
 - Counting at the granularity of 12 hours
 - No additional late days

Syllabus (5)

- **Important dates**

Jan 11th, 2016: Classes begin

Jan 14th, 2016: Last day to drop and request full refund

Jan 15th, 2016: Add deadline on myUCF

March 23rd, 2016: Withdrawal deadline

April 26th, 2016: Classes end; Last day to remove incomplete

April 28th, 2016: Final project presentation

See <http://calendar.ucf.edu/2016/spring> for any updates

Syllabus (6)

- **Statement on Academic Integrity:**

The UCF Golden Rule (<http://goldenrule.sdes.ucf.edu/>) will be observed in the class. Plagiarism and Cheating of any kind on an examination, quiz, or assignment will result at least in an "F" for that assignment (and may, depending on the severity of the case, lead to an "F" for the entire course) and may be subject to appropriate referral to the Office of Student Conduct for further action. I will assume for this course that you will adhere to the academic creed of this University and will maintain the highest standards of academic integrity. In other words, don't cheat by giving answers to others or taking them from anyone else. I will also adhere to the highest standards of academic integrity, so please do not ask me to change (or expect me to change) your grade illegitimately or to bend or break rules for one person that will not apply to everyone.

Questions?