CAP 6412: Advanced Computer Vision

Introduction

- Seminar-Style Class
- We will all take turns presenting papers
- Goals:
  - Survey “latest and greatest” research
  - Learn to read papers
  - Learn to present

Administrivia

- Office Hours:
- T-Th – 11AM to 12:30PM
- ENG3-230
- Course Web-page: http://www.cs.ucf.edu/~mtappen/cap6412

Expectations

- Each student will be required to present one paper
- All others will be required to turn in a 1-2 page report on the paper
- Due at beginning of class
  - If you have genuine reasons to miss class, let me know

Grading

- Reports - 30%
- Presentation - 30%
- Implementations(2) - 40% -
- Two implementation projects will be assigned over the course of the semester.
  - You will create a basic implementation of a paper

Reports

- Summarize Contributions
- Describe Strengths
- Describe Weaknesses
- What's Next?
  - What interesting research ideas stem from this paper
  - Is there a cool idea that you could build on?

Presentations

- You will probably only have to do one.
- Please spend time on it.
- What's the main idea?
- What's the cool idea?
- Is there some subtle, interesting math?
- What would break this algorithm?

Tips on Reading (Borrowed From Dr. Shah)

- First read once
  - Skip related work
  - Understand problem
  - Do not try to understand details
- Read once more – understand details
  - Try to extract sub-problems
  - Think about how to implement them
  - Read as many
  - late to other Vision Problems you are aware of
- Understand how the sub-problems are combined together.

The Papers

- My Interests
  - Machine Learning
  - Statistical Modeling
  - Machine Vision
- Most of the papers involve a statistical model of images or scenes
- A few learning –only papers
- Three Main Themes:
  - Low-Level Vision
  - Object/Scene Recognition
  - Activity Recognition
    - Could use suggestions for papers on this topic!

Next Time

- I'll present on graphical models
- From a Book Chapter
- No Report Due
- Please sign-up today or tomorrow