

Synopsis: We explore some fundamentals of image processing and understanding, considering several image representations for extracting information and survey algorithms for solving a variety of problems, such as image de-noising, panorama creation, 3-D reconstruction, segmentation, and object recognition.

CSC295 Computer Vision Spring 2010

| |
|-------------------------------|
| MWF 3:15-4:05 pm Science 3819 |
|-------------------------------|

Instructor: Jerod Weinman
Office: Noyce 3825
Phone: x9812
E-mail: [weinman]

Office hours:

| | |
|-----------|--------------|
| Monday | 1:00-2:00 PM |
| Tuesday | 3:30-4:30 PM |
| Wednesday | 4:15-5:15 PM |
| Friday | 1:30-2:30 PM |

or by appointment.

Course web page: <http://www.cs.grinnell.edu/~weinman/courses/CSC295/2010S>

1 Overview

How can we make computers understand images? What information is contained in an image and how can a meaningful representation be extracted? The problems of computer vision are two-fold. First, how can we take a two-dimensional image and reconstruct the three-dimensional world it came from? Second, how can we understand and recognize objects from those three-dimensional worlds?

Computer vision draws on a surprisingly large and diverse set of other fields, including psychology, neuroscience, mathematics, photogrammetry, optics, physics, signal processing, pattern recognition, and artificial intelligence. We will try to find a way to glean what we need from each of these in a truly liberal (as in “liberal arts”) fashion.

In this course, we will examine computational approaches to problems in low-level visual processing as well as some applications in higher-level vision.

Our major objectives for this course include:

- Understanding the fundamental problems in computer vision
- Learning both new and classical approaches to vision problems
- Finding and appreciating limitations and caveats of algorithms

Why take it? Computers can do many things much better than humans, but this area is one where humans still excel. If you are interested in an exciting field with many open problems and practical applications, this course is for you.

What do I need to know? This course assumes you have experience with programming in both the functional *and* imperative paradigms. You should also be fluent in some basic computational ideas like encapsulation and code re-use. Some mathematical maturity, especially familiarity with multivariate calculus and linear algebra, will be helpful. We will review these briefly as necessary. If linear algebra is a distant memory, I strongly urge you to study Eero Simoncelli’s “A Geometric Review of Linear Algebra” <http://www.cns.nyu.edu/~eero/NOTES/geomLinAlg.pdf> right away.

2 Texts

Due to discrepancy between topics or details appropriate to an undergraduate course, we have no single, official textbook. Instead, we will use selected readings from a variety of texts, which will be available on the course web page, PioneerWeb, and E-Reserves:

David Forsyth and Jean Ponce, *Computer Vision: A Modern Approach*, Prentice Hall, 2003.

Emanuele Trucco and Alessandro Verri, *Introductory Techniques for 3-D Computer Vision*, Prentice Hall, 2008.

Richard Szeliski, *Computer Vision: Algorithms and Applications*, 2010.

<http://research.microsoft.com/en-us/um/people/szeliski/Book/>

The latter text is a work-in-progress by a senior researcher at Microsoft. He has graciously made drafts of the manuscript available on the internet for review by folks like us. Throughout the semester, please send me errors and comments for the author via email; I will collect and distribute these to him so that he can improve the book.

Our laboratory exercises will be done in MATLAB. This software is available on the MathLAN, but if you wish to have MATLAB on your own personal computer, a student version may be purchased from www.mathworks.com/store, www.journeyed.com, or www.academicssuperstore.com.

3 Class Attendance

Class meetings will involve a mix of discussions, labs, and lectures. In short: *You are expected to attend and actively participate in class. I am expected to make class worth attending.*

I know that sometimes “things happen.” Therefore, you will be granted **one** unexcused absence from class without penalty. However, this is a collaborative, lab-based course, so your presence is integral to your learning. Thus, the following will be deducted from your participation grade in the case of absences:

| | |
|--------------------|-----|
| 2-3 absences | 20% |
| 4-5 absences | 40% |
| 6 or more absences | 80% |

If you wish me to acknowledge your absence as excused, you must either:

1. Notify me in writing (email is acceptable) at least 7 days in advance to make arrangements for your absence, *or*
2. Ensure I have received documentation of your absence’s circumstances *post hoc* from Health Services or Student Affairs.

Our meetings benefit from your contributions. If you do miss a class, you must first talk to a classmate about any material that you may have missed. *After* that, you may follow up with the instructor about any further questions or concerns.

4 Schedule of Topics

The following is an *approximate* schedule of topics to be discussed during the course. See the web page schedule for daily details.

| Week | Topic | Week | Topic |
|------|-----------------------|------|---------------------------|
| 1 | Image Formation | 8 | Feature Matching |
| 2 | Image Processing | 9 | 3-D Structure from Motion |
| 3 | Edge Detection | 10 | Stereo Geometry |
| 4 | Fourier Transforms | 11 | Texture and Color |
| 5 | Pyramids and Wavelets | 12 | Segmentation |
| 6 | Feature Detection | 13 | Project Work |
| 7 | Feature Descriptors | 14 | Recognition |

5 Activities

Under a normal 16 credit load, I expect that you will spend at least 40 hours per week on your studies (class time, homework, and studying).¹ Thus, you should plan to spend a *minimum* of 10 hours/week on work for this course (but hopefully more). With class time clocking in at $2\frac{1}{2}$ hours, you’ll have $7\frac{1}{2}$ hours/week left for the following:

¹This is a minimum recommendation for achieving “satisfactory” (i.e., C-level) results. “Good” or “excellent” results may require a greater investment of time.

5.1 Reading

Our class meetings will be heavily discussion and lab-based, and this will require a significant amount of preparation on your part. You should check the class schedule for updates and read any material that has been assigned *before* coming to class. Reading the textbook entails the following:

Overview You should quickly skim through the reading once to get an overview of the material to be covered, paying particular attention to subject headings and topic introductions. This first “reading” can (and should) be very quick. (*Expected time: 5 to 10 minutes.*)

In-Depth Next, read the material closely. Try to understand what individual steps of algorithms or mathematical equations are accomplishing. Not everything will make sense at this point, but hopefully many things will. (*Expected time: 40 to 50 minutes.*)

Final Notes After carefully reading the material, mentally review and try making a few notes to yourself about what you think are the most important concepts being covered, as well as any questions you have. Bring your questions to class. I will ask for them. (*Expected time: 5 to 10 minutes.*)

Most readings for this class contain mathematics or algorithms that require a moderate amount of study. While I realize not everyone learns best by reading, you are asked to make your best effort and come to class with any questions you may have. Then we can proceed with discussion, examples, and exercises that enhance and clarify the material in class.

5.2 Participation

Because much of our work in this course involves collaboration and discussion, you will be evaluated on your participation.

Participating in class involves:

- being present in class (physically and mentally)
- coming to class on time
- coming to class prepared
- asking questions when appropriate
- making positive contributions to class discussion by volunteering and when called upon
- staying on task during collaborative exercises, and
- working effectively with your group.

Students who regularly meet these criteria can expect to earn 90% (i.e., an A-) for their participation grade. I will reward students who regularly provide significant insights or guide discussion in productive ways with a higher participation score. Students who fail to participate regularly or who participate in counterproductive ways (e.g., by dominating the conversation or making inappropriate comments) can expect to earn a lower score.

One unexcused absence will have no effect on your participation score. (See the Attendance Policy, Section 3 above.)

5.3 Laboratory Exercises and Write-Ups

Over the course of the term, approximately half of our class meetings will involve laboratory exercises. The “hands-on” labs will provide an instructor-assisted setting for you concretely engage in the material introduced in class and readings. Labs will be posted in advance of the date they are begun in class; since our in-class lab time is limited, it is very important for you to come prepared by reading the lab assignment beforehand, bringing any questions to resolve at the beginning of class.

You will be expected to produce a formal write-up of your results, including images and figures as appropriate, a description of your method(s), and providing some conclusions. You will also submit your code. Please see the course web page (and individual labs) for further details on what you are expected to submit. Although we will begin the labs in class, it is possible that you may not finish the technical material in class and unlikely you will finish the write-up.

You will typically complete these labs in pairs or groups, oftentimes assigned by me on a rotating basis. Though you may start the lab as a group, I know that some times it is difficult to synchronize schedules on short time frames. Thus, you may elect to complete your write-ups individually. In that case, you must clearly identify the contributions of your original partner(s) in the interest of academic honesty. Conversely, everyone whose name appears on a submitted group lab report has the responsibility to ensure everyone fully understands the submission.

Discussion of concepts and approaches with other classmates is encouraged. Debugging programs can be difficult and is often helped by explaining your code to someone else (which will also frequently help you to explain the bug to *yourself*). All such contributions by others (not in your group) must be properly attributed in your report. Furthermore, all the work submitted (code, experimental data, write-ups, etc.) must be your own. Code provided by the instructor should be attributed if it is included, but no other code or written work (from *any* source) may be shared with others *or* copied for your own use.

You are highly encouraged to use the PioneerWeb Discussion Board for questions related to the course. If a post is related to an assignment, it must adhere to the standards of collaboration for that particular assignment. If you are debugging, one good rule of thumb is to not post code. (I often find that when I am in the process of explaining some error by writing up a post or email, I typically find its source.)

5.4 Exams

As opportunities for you to demonstrate your image interpretation prowess and grasp of vision principles there will be two mid-semester exams, one before spring break and the other during the scheduled final exam period. Although exam 2 is not intended to be cumulative, certain topics build on each other naturally and cannot be excluded.

| | |
|--------|------------------------|
| Exam 1 | Friday, March 12 |
| Exam 2 | Tuesday, May 18 (9 am) |

Do not make airline reservations that will conflict with your exam schedule.

5.5 Independent Project

In the second half of the semester, you will propose and deploy an independent project that uses and builds on the principles we are learning in the class. You will have several weeks to work on the project and will be expected to conduct a small amount of mentored, but self-directed learning. Several example possibilities will be provided to get your imagination running. You will propose a project including a description of how a successful deployment will be produced, as well as two benchmarks. You will present the results of your work in a final report and a presentation on the last day of class. Further details will be given after spring break.

6 Grading

My goal is for everyone taking this course to be able to demonstrate familiarity and fluency with the course concepts. I would be very happy if you all met the goals above and received “A”s. The following weighting will provide a basis for evaluation.

| | |
|----------------------|-----|
| Laboratory Write-ups | 55% |
| Participation | 10% |
| Exams | 20% |
| Project | 15% |

7 Academic Honesty

You, as students, are members of the academic community. Both the College and I expect the highest standards of academic honesty. (See the Grinnell College Student Handbook, e.g.,

<http://www.grinnell.edu/offices/studentaffairs/shb/section3/academichonesty>).

Among other things, this means clearly distinguishing between work that is your own, and work that should be attributed to others. It is expected that the collaboration policies given in this syllabus and on particular assignments will be followed. Furthermore, any program results or output must be faithfully recorded, not forged. (A thoughtful explanation of unexpected behavior can often be a worthwhile submission and is *much* better than the alternative.)

In your homework assignments, you must give specific attribution for any assistance you receive. For example, one possible acknowledgment format is “[Person X] helped me to do [thing Y] by [explaining Z].”

As an instructor, I will meet my obligation to bring any work suspected to be in violation of the College’s Academic Honesty Policy to the attention of the Committee on Academic Standing, after which there is no recourse with me.

8 Deadlines

Assignments are due at the *beginning* of class on the specified date. Work submitted more than five minutes after the beginning of class will be considered late. Assignments due on days for which you have a prior excused absence must still be submitted by the deadline.

A late penalty of 33.33% will be deducted at each subsequent class meeting. Thus, you have at most two additional meetings to submit your work.

Exception: Deadlines for MathLAN computer-based assignments will automatically be extended by at least one class period if MathLAN is down for an unscheduled period of 3 or more hours during the week preceding the assignment due date.

9 Contacting Me

Please come by during my office hours to discuss the course content, get any extra assistance, or just talk about how the course is going. Note that if multiple students have similar questions or issues, we may work together as a group. If you cannot attend a scheduled office hour, you may also email me to schedule an appointment; please include 3-4 possible meeting times so that I can pick one that works for me, too.

I enjoy getting to know my students, but I prefer to reserve office hours for academic matters. If you would like to have a more informal conversation, I would be delighted to accept an invitation to lunch.

Email is also a reliable way to contact me, but please allow 24 hours for a response (except on weekends, when I do not regularly read email). You may also call me in my office (x9812).

10 Accommodations

If you have any disability that requires accommodations, please meet with me right away so that we can work together to find accommodations that meet your learning needs. You will also need to provide documentation of your disability to the Dean for Student Academic Support and Advising, Joyce Stern, located on the 3rd floor of the Rosenfield Center (x3702).

With thanks to Janet Davis for the “Reading Suggestions” and other key policies.