

Course Proposal: CSC 262, Computer Vision

Department of Computer Science

November 7, 2013

Background

Computer Vision has been offered twice as a four-credit special topics course, CSC 295. Now that it has reached the limit in that regard, we propose to regularize the course as an alternate-year offering. The regularization is neutral with regard to staffing, insofar as it has been considered part of Jerod Weinman's regular instruction load since his arrival at Grinnell.

CSC 262 is an elective, open to both majors and non-majors, and satisfies a requirement for advanced laboratory work in the neuroscience concentration. The interdisciplinary-themed course develops programming, general laboratory, and communication skills (oral and written) in addition to bridging concepts from computing, psychology, physics, mathematics, and philosophy.

Catalog Description

How can computers understand images? This course covers principles in computational vision and their relationship to human visual perception. Topics may include geometry of image formation, image filtering and representation, texture analysis, 3-D reconstruction from stereo and motion, image segmentation, object detection, and recognition. Students implement and evaluate methods on real images in laboratory exercises and an independent project.

Syllabus Description

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 twofold. 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? In summary, what is the nature of information processing for visual data?

To answer such a question, computer vision draws on a surprisingly large and diverse set of fields beyond computer science, including psychology, neuroscience, mathematics, photogrammetry, optics, physics, signal processing, pattern recognition, artificial intelligence, and even philosophy. In this course, we will glean what we need from each of these in a truly liberal (as in "liberal arts") fashion, examining computational approaches to problems in low-level visual processing and 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, appreciating, and writing* about **limitations** of models and algorithms

Textbooks

As there is no universally suitable textbook for an undergraduate course in computer vision, the sources are varied. Our primary source will be the following textbook.

Richard Szeliski, *Computer Vision: Algorithms and Applications*, Springer, 2011.

The author hosts a free online preprint version (used for the prior two iterations of the special topics course). For topics that are not treated well by Szeliski, chapters from the following serve as useful additional references:

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.

Previously, we have casually used excerpts from these two books. However, as we regularize the course, we should formally obtain the publishers' permission and pay the appropriate licensing fees.

Topics

Image formation: Lenses, geometry, and sensors	(Week 1)
Local image operators: Gain and bias, gamma correction, and histogram equalization	(Week 2)
Image filtering: Noise, convolution, correlation, edge detection, and scale selection	(Weeks 2–3)
Image representation: Fourier transform, image pyramids, and wavelets	(Weeks 3–5)
Object detection/recognition: Keypoint detection, local descriptors, and feature matching	(Weeks 5–7)
3-D reconstruction: Stereo epipolar geometry, disparity estimation, and structure from motion	(Weeks 8–9)
Image segmentation and grouping: Top-down vs. bottom-up, background subtraction, and clustering	(Week 10)
Color: Color temperature, color spectrum, human perception, and color spaces	(Week 11)
Texture modeling: Image statistics, representation, analysis, and synthesis	(Week 12)
Current research: Recent award-winning primary literature	(Weeks 13–14)
Independent project work: In-class consultations and oral presentations	(Weeks 9–14)

Course Activities

This class meets three times per week for fifty minutes. In-class activities are evenly divided between discussions clarifying concepts from textbook readings, and directed, workshop-style lab exercises that range from empirically validating, testing, and modifying existing implementations of algorithms, to creating their own implementations for solving relatively simple problems.

In the ninth week, students will also propose a final group project incorporating several aspects of techniques they have learned (or will learn soon in the term). An ideal project should involve

- an interesting application,
- possibility for evaluation on synthetic and actual data (either an established standard or student-generated),
- clear milestones, and
- a small amount of mentored, but self-directed, knowledge acquisition

Students are assessed on their class participation, written lab reports, two mid-term exams, and their project reports (written and oral).