

CS510: Image Computation

Bruce A. Draper
draper@cs.colostate.edu
(970) 491-7873
Room 442, CSB

www.cs.colostate.edu/~cs510

The logo for Colorado State University, featuring a green wavy banner with yellow lines and the text "Colorado State University" in gold.

Colorado State University

Course Theme

- Stage 1 : Text-based Computers **Done**
 - Manual data entry, command-line interface
- Stage 2 : Image Output (graphics on screens) **Ending**
 - Gaming computers (e.g. Xbox)
 - Window-based operating systems
 - Applications: Maps, YouTube, etc.
- Stage 3 : Image Input (cameras everywhere) **Beginning!**
 - Gaming (e.g. Kinect)
 - Security (e.g. Face recognition,)
 - Ubiquitous Computing (e.g. Self-driving cars)

Course Goals

- Prepare students for graduate level research in computer vision
 - Note: no one on the current CSU faculty does (or advises) Ph.D.-level research in graphics.
- Prepare students to integrate vision into new applications (Stage 3!)

Prerequisites

- You are assumed to know:
 - 1) Geometric primitives: points, lines, vectors ...
 - 2) Homogeneous coordinates
 - 3) 3D transformations & their compositions
 - 4) Perspective Projection
 - 5) Geometry of camera models
 - 6) Lighting models & material properties
- Note that this material will not be reviewed: it is assumed that you know it and can build on it.

Course Outline

1. Images & Image Matching

A. Image Transformations

- Geometric Transformations
- Photometric Transformations

B. Template matching

- Pearson's correlation
- Linear correlation
- Convolution

C. Fourier Transforms

- 1D & 2D Discrete Fourier Transforms
- Image space interpretation
- Correlation & convolution in frequency domain

Course Outline (II)

1. Images & Image Matching (cont).

D. Scale

- Scale-space theory
- Image pyramids
- Super-resolution

E. Correlation filters

- MACH filters
- ASEF/MOSSE filters

F. Principal Components Analysis

- Covariance minimization & compression
- Gaussian process model
- Subspace projection

Course Outline (III)

2. Image Features & Feature Matching

A. Correspondence-free approaches

A. Local Features

- Edges (Sobel, Canny)
- Corners (Harris)
- DoG (Lowe)

B. Feature Descriptors

- SIFT
- HoG
- Color spaces & color histograms

C. Matching

- Bag of Words
- Deformable Shape Models

Course Outline (IV)

2. Image Features & Feature Matching (cont).

A. Correspondence-based Approaches

- RANSAC & Geometric Hashing
- Active Shape Models
- Pose estimation
- SLAM
 - Simultaneous Localization and Mapping

Textbooks

- No single textbook covers these topics.
- Nonetheless,
 - We will mostly use Rick Szeliski's text
 - <http://szeliski.org/book>
 - free & online
 - We will also use CVOnline
 - <http://homepages.inf.ed.ac.uk/rbf/CVonline/>
 - Free & online, but less well edited
 - Simon Prince's book is also available (but...)
 - <http://www.computervisionmodels.com/>
 - I will sometimes reference Shapiro & Stockman's text
 - Outdated, but good on basics
 - Unfortunately, not on line, so I must summarize...

Requirements

- 1) Four programming assignments (75%)
 - Each section has an “easy” assignment, then a hard one
 - In C, C++, or Python
 - Java OpenCV interface is still incomplete
 - Use of OpenCV recommended
 - Qt OK for GUIs.
 - Any other library : ask me first!
 - Variations on assignments are encouraged -- but check with me first!

Requirements (II)

- Midterm (10%)
- Final exam (10%)
 - Non cumulative
- In class participation (5%)
 - Subjective on my part
 - Discussions, questions, etc.
 - So be pro-active! Get involved!

Class Slides

- Class slides (lectures) go on the class web page.
 - As soon as the current bug is fixed.
- The class web page is a general source of useful information, including links to interesting computer graphics and/or computer vision web sites:
 - <http://www.cs.colostate.edu/~cs510>

Discussion

- What can a computer use a camera for?
- Examples:
 - iPhoto
 - Google Image Search
 - Star Walk (constellation finder)
 - Augmented reality
 - Subway finder, etc.
 - Driver assist
 - Lane changes
 - Obstacle avoidance
 - Automatic parking
 - Security / HCI
 - Is the person sitting at the computer the person who is logged in?

One last thing...

- Send me an email from the account you read most often, with the subject line “Spring CS510”
 - I will use this to make a real class mailing list
 - (Aries web emails are too often not real)