Computer Vision (CSE P576)

Staff

Prof. Steve Seitz
seitz@cs

T.A. Ricardo Martin
rmartin@cs@

T.A. Igor Mordatch
mordatch@cs

Web Page

- http://www.cs.washington.edu/education/courses/csep576/15sp/

Handouts

- signup sheet

Today

- Intros
- Computer vision overview
- Course overview
- Image processing
- Features

Readings for this week

- Book: Richard Szeliski, Computer Vision: Algorithms and Applications
- Ch 1.0, 3.1-3.2, 4.1

What is computer vision?

image recognition
image de-rasterization
self driving cars

What is computer vision?

Terminator 2

Every picture tells a story

Goal of computer vision is to write computer programs that can interpret images

What do computers see?

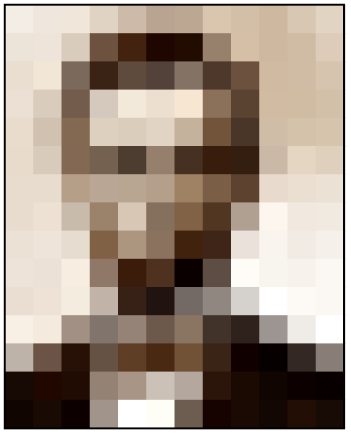

slide by Larry Zitnick

Can computers match (or beat) human vision?

Yes and no (but mostly no!)

- humans are much better at “hard” things
- computers can be better at “easy” things

Human perception has its shortcomings…

Sinha and Poggio, *Nature*, 1996

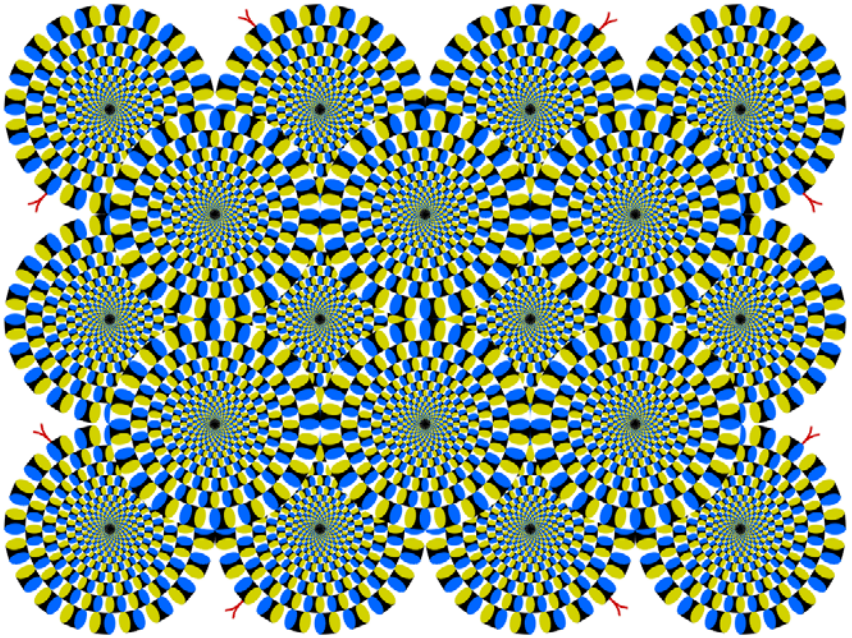

Copyright A.Kitaoka 2003

Current state of the art

The next slides show some examples of what current vision systems can do

3D Maps

Games

Microsoft's XBox Kinect

Face detection

Most digital cameras detect faces

Face recognition

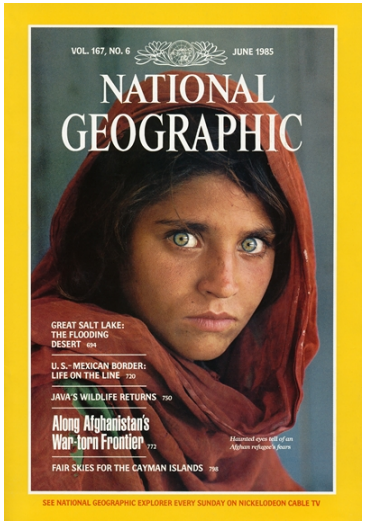

Who is she?

Vision-based biometrics

"*How the Afghan Girl was Identified by Her Iris Patterns*" Read the story

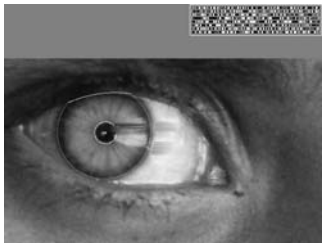

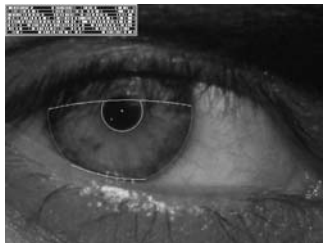

Object recognition

Google Goggles
Bing Vision

Special effects: shape capture

The Matrix movies, ESC Entertainment, XYZRGB, NRC

Sports

Sportvision first down line
Nice explanation on www.howstuffworks.com

Smart cars

Slide content courtesy of Amnon Shashua

Mobileye

- Vision systems currently in high-end BMW, GM, Volvo models

Self-driving cars

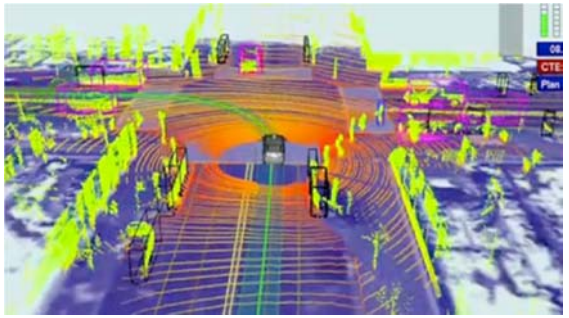

"Our self-driving cars have now traveled nearly 200,000 miles on public highways in California and Nevada, 100 percent safely. They have driven from San Francisco to Los Angeles and around Lake Tahoe, and have even descended crooked Lombard Street in San Francisco. They drive anywhere a car can legally drive."

- Sebastian Thrun, Google

Robotics

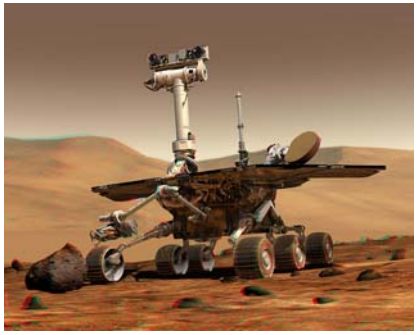

NASA's Mars Spirit Rover
http://en.wikipedia.org/wiki/Spirit_rover

http://www.robocup.org/

Virtual Reality

Google Cardboard

Microsoft Hololens

Current state of the art

You just saw examples of current systems.

- Many of these are less than 5 years old

This is a very active research area, and rapidly changing

- Many new apps in the next 5 years

To learn more about vision applications and companies

- David Lowe maintains an excellent overview of vision companies
  - http://www.cs.ubc.ca/spider/lowe/vision.html

This course

http://www.cs.washington.edu/education/courses/csep576/15sp/

Project 1: Features

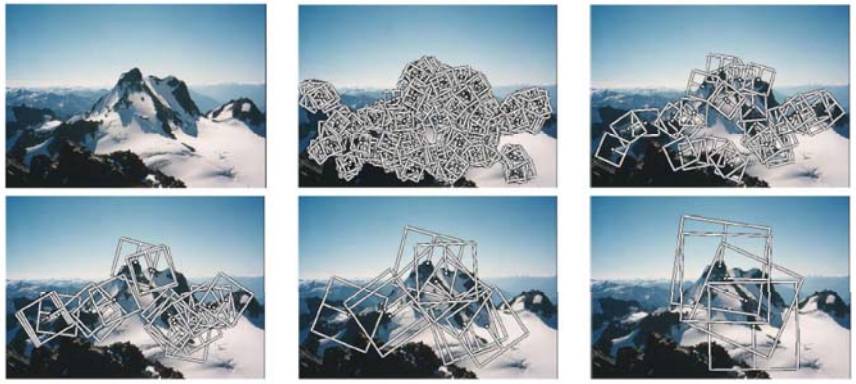

Figure 1. Multi-scale Oriented Patches (MOPS) extracted at five pyramid levels from one of the Matier images. The boxes show the feature orientation and the region from which the descriptor vector is sampled.

Project 2: Panoramas

Indri Atmosukarto, 576 08sp

Project 3: 3D Shape Reconstruction

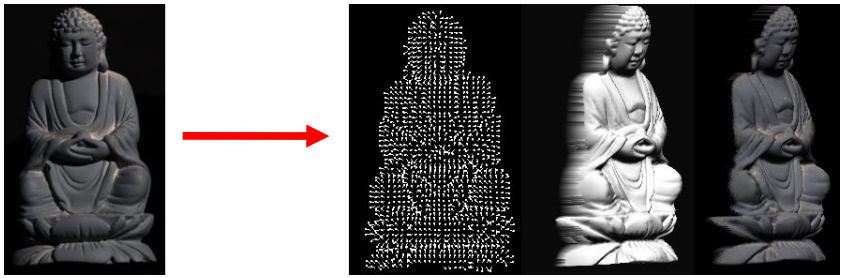

Project 4: Face Recognition

Grading

Programming Projects (100%)

- features
- panoramas
- 3D shape modeling
- face recognition

General Comments

Prerequisites—*these are essential*!

- Data structures
- A good working knowledge of C and C++ programming
- Linear algebra
- Vector calculus

Course does ***not*** assume prior imaging experience

- computer vision, image processing, graphics, etc.