Lecture 17

Final Review

All good things must...

Course Grading

Programming Projects (80%)

- Image scissors (20%) **- DONE!**
- Panoramas (20%) **- DONE!**
- Content-based image retrieval (20%) **- DONE!**
- Face recognition & detection (20%) – Due tomorrow (11:59pm)

Final (20%)

Everything you wanted to know about the final

but were afraid to ask…

The Final Exam

- When & Where: Thursday March 19 10:30am-12:20pm here (EE 037)
- Closed book, closed notes except for one 8 ½'' x 11'' sheet of notes
- Format: Similar to sample exam on class website (under Calendar)
  - Five questions
  - Short answer questions as well as longer problems (see sample exam for examples)

How do I ace the final?

Option A: Memorize this

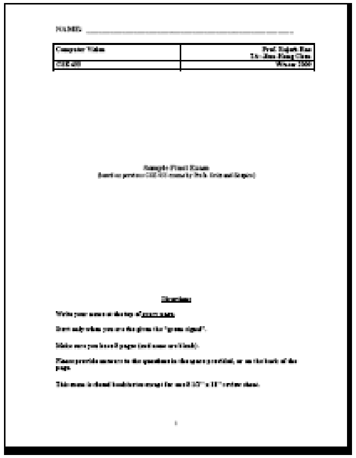
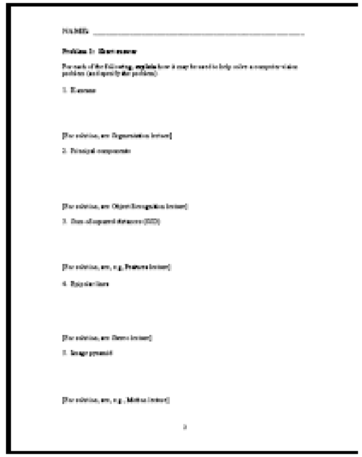
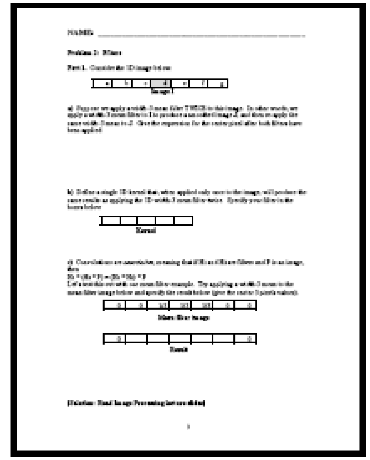
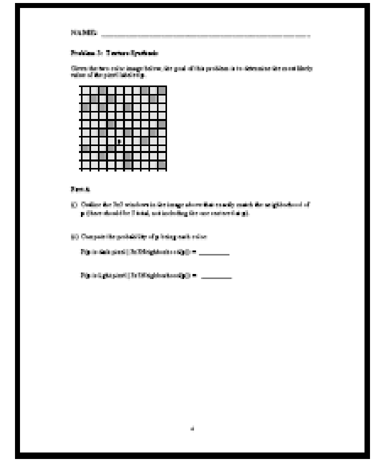
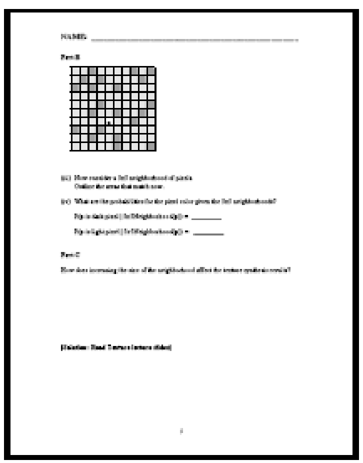
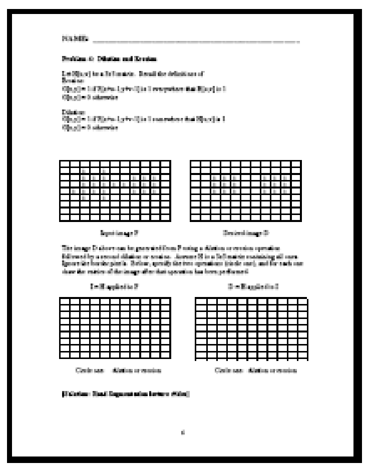
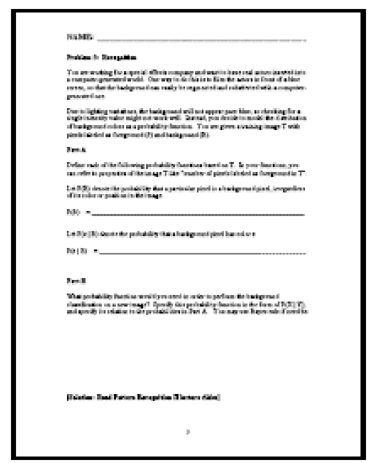
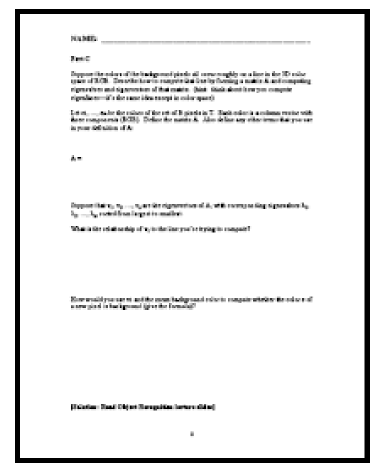

Option B

- Read the lecture slides
  - Only topics covered in the lecture slides will be the subject of questions in the final
- Consult the accompanying readings (on Calendar page) and projects for details where necessary
- Try the sample final

Course review and highlights

In the beginning…

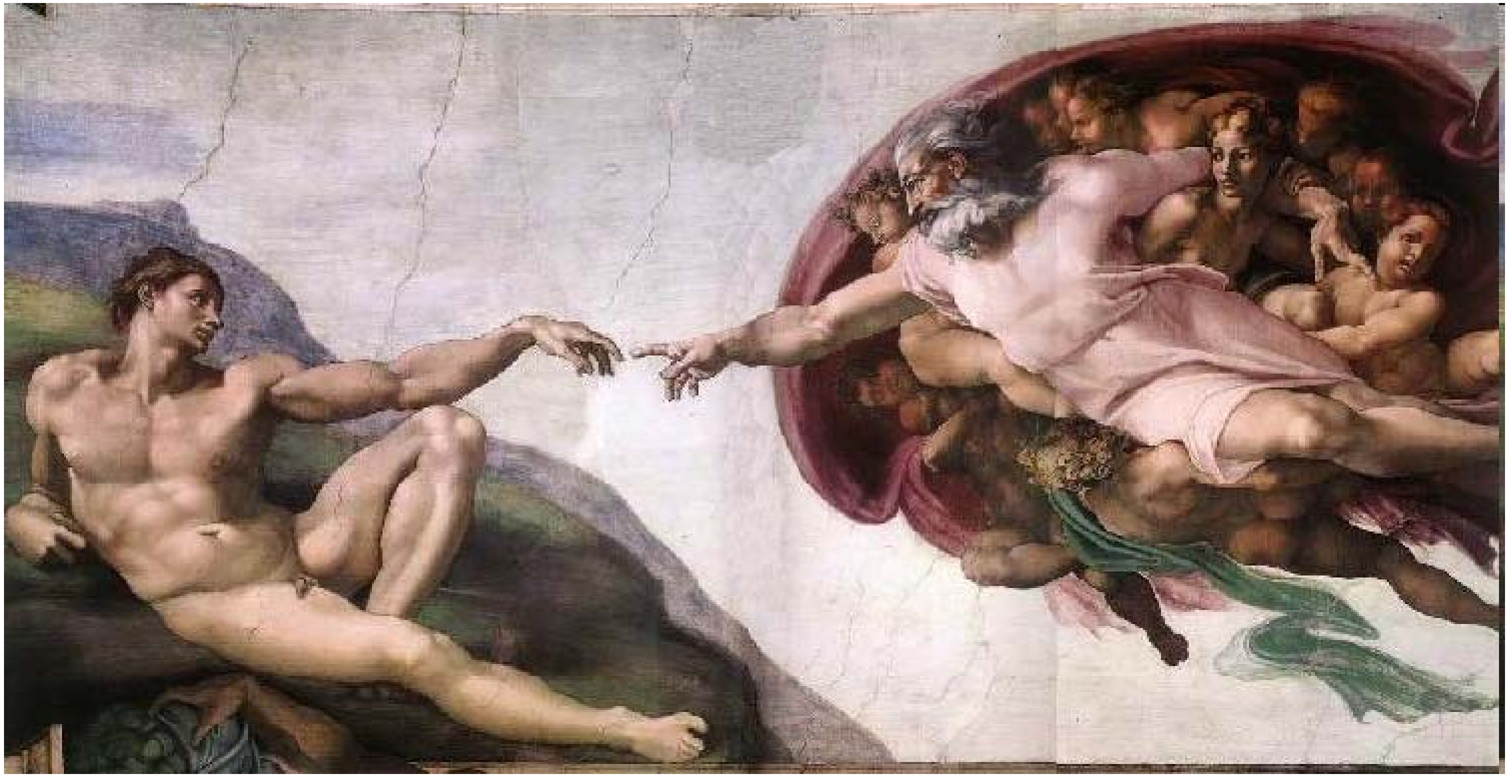

Image Processing and Filtering

- Image as a function
- Image formation: illumination, reflectance
- Domain transformations
  - Homogenous coordinates
  - Affine transformations
- Range transformations:

  Filtering
  - Cross-correlation, template matching
  - Convolution
  - Filter kernel
  - Mean filtering, Gaussian filtering
  - Median filtering

Highlight

Noise reduction through mean filtering

Input image

Salt and pepper noise

Filtered Images

3 x 3

5 x 5

7 x 7

Kernel size

Image Sampling and Pyramids

- Sub-sampling artifacts
- Aliasing and Nyquist rate
- Gaussian pre-filtering
- Gaussian pyramid
- Image resampling and interpolation
  - Linear and bilinear

Highlight

Gaussian Pyramid

Blur the image (low pass filter) the image, *then* subsample

- Blur using Gaussian filter

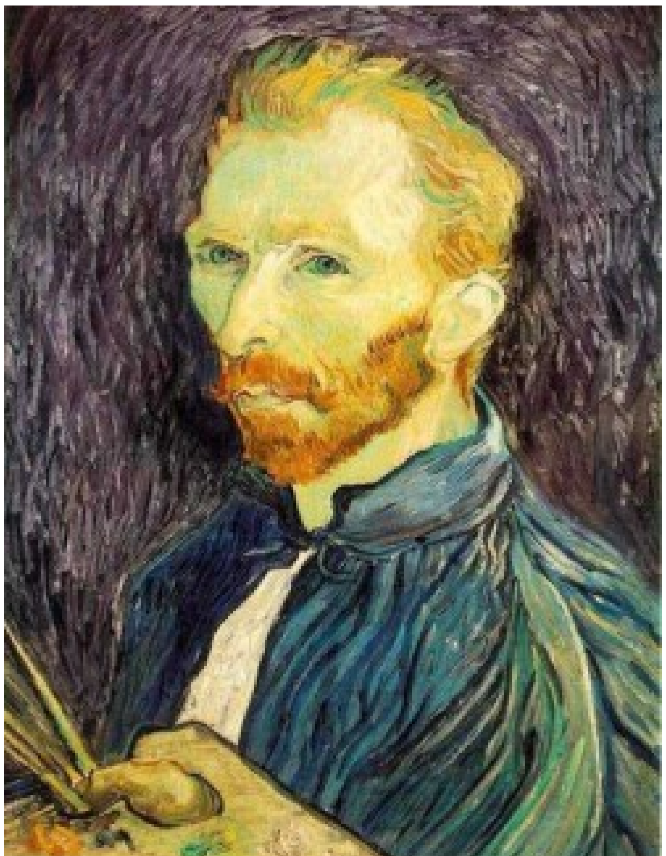

Gaussian 1/2

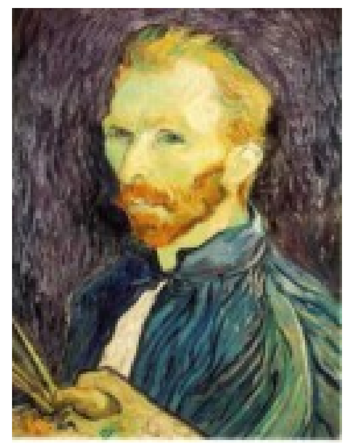

G 1/4

G 1/8

Edge Detection

- Image gradient
  - Magnitude and direction
  - Discrete gradient, Sobel operator
- Derivative theorem of convolution
- Edge detection filters
  - Derivative of Gaussian
  - Laplacian, Laplacian of Gaussian (LoG)
- Canny edge detector (5 steps)

Highlight

Edge detection using LoGs

original image (Lena)

LoG followed by zero crossing detection

Intelligent Scissors

- Problem being solved?
  - Semi-automatic segmentation
- Computing lowest cost paths from seed to all pixels
- Image as a graph
  - Cost of links computed from cross-correlation filter
  - Dijkstra’s shortest path algorithm
- Creating composite images

Highlight

Your creativity at work!

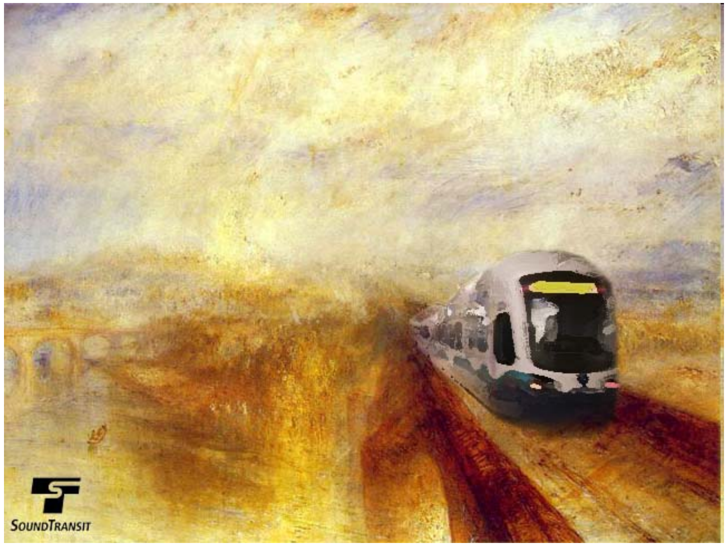

Cameras and Projection

- Pinhole camera
  - Aperture and effects of its size
- Lenses
  - Thin lens equation
  - Magnification by a lens
- Modeling projection
  - COP, PP
  - Projection equations (derive using similar triangles)
  - Homogenous coordinates
  - Perspective projection using homogenous coordinates
  - Vanishing points

Highlight

Perspective Projection

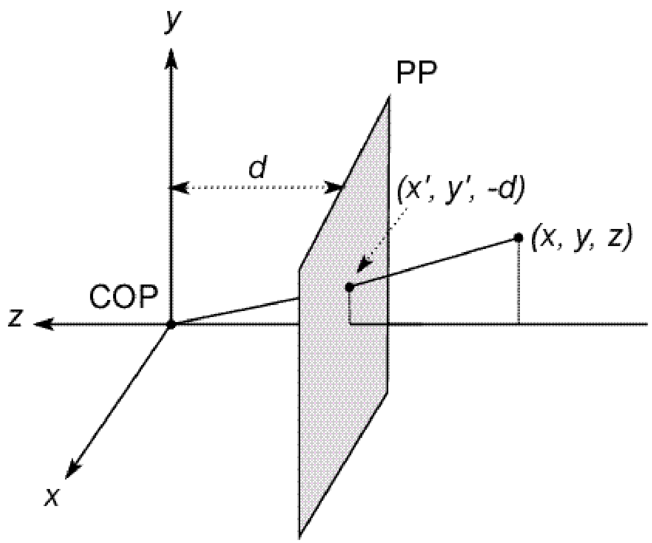

Other Highlights

Your work!

Cameras and Projection (cont.)

- Simplified projection models
  - Weak perspective
  - Orthographic
- Radial distortion modeling
- Extrinsic vs. intrinsic camera parameters
- World coordinates to pixel coordinates transform
- Camera calibration

Features and Image Matching

- Feature Detection
  - Finding interest points

$$E(u,v) \approx [u \;\; v] \underbrace{\left(\sum_{(x,y)\in W} \begin{bmatrix} I_x^2 & I_x I_y \\ I_x I_y & I_y^2 \end{bmatrix} \right)}_{H} \begin{bmatrix} u \\ v \end{bmatrix}$$

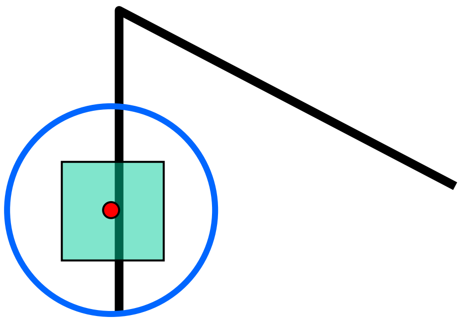

  - Eigenvalues and eigenvectors of H
  - Choose points where smaller eigenvalue is a local maximum as interest points
  - Rotation and brightness invariance
- Harris corner detector
- Scale invariance using a pyramid

Highlight

Harris detector example

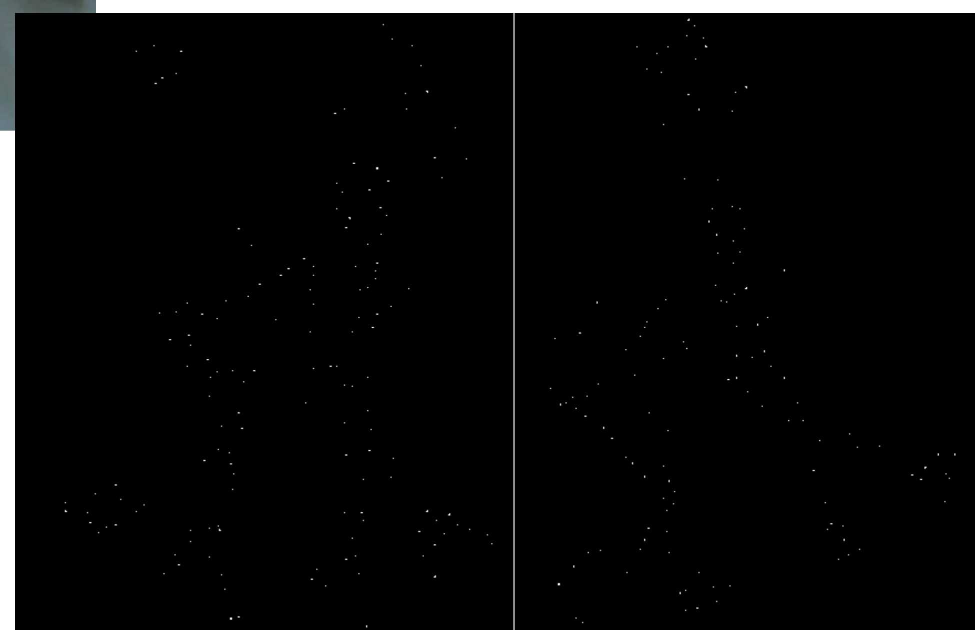

Other Highlights

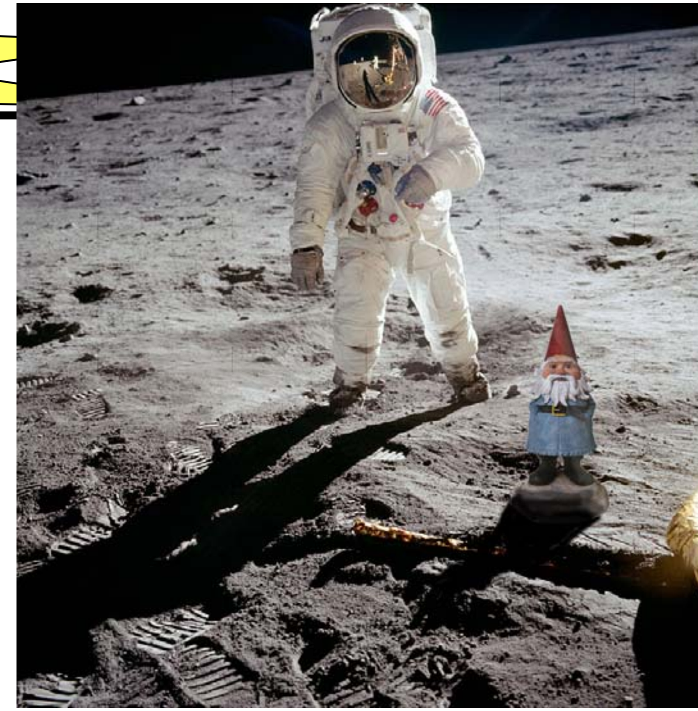

Your work!

Features and Image Matching

- Feature Descriptors
  - MOPS
  - SIFT (just the basic idea of orientation histogram)
- Feature matching
  - SSD
  - Ratio of SSDs
- True positives and false positives
- ROC curve
  - How is it computed?
  - How is it used?

Image Stitching and Panoramic Mosaics

- Go back and review the steps you used in your project 2 (Panoroma)
  - Warp to spherical coordinates
  - Feature extraction and matching
    - Note: The correct definition for the original Harris detector is the one in the Feature detection lecture
  - Aligning using RANSAC
  - Correcting for drift and blending

Highlight

Your Panoramas!

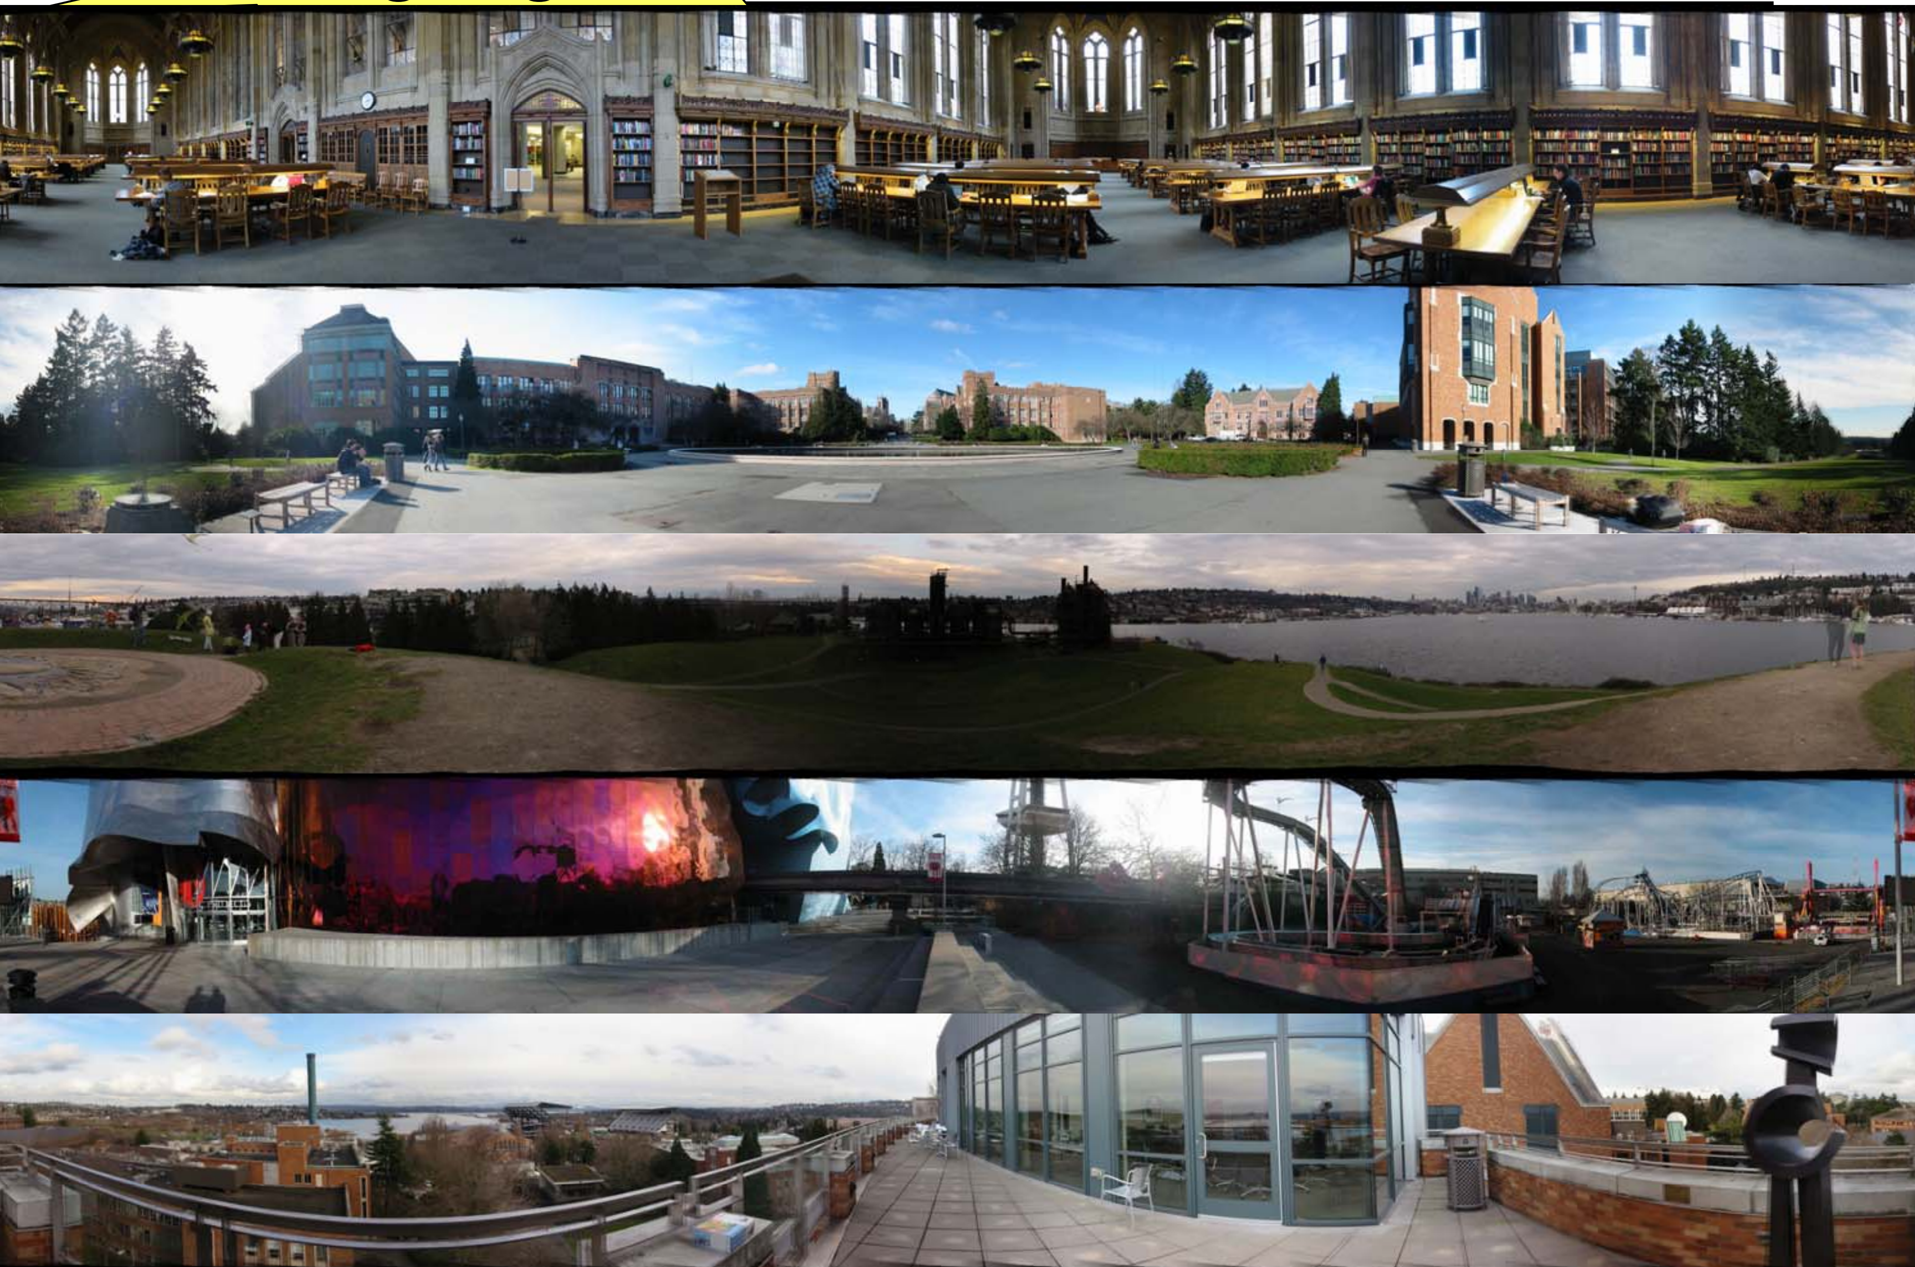

Human Vision

- Visible spectrum
- Retina
  - Rods, cones
  - Space-time-color filtering in the retina
- Eye movements
- Visual pathway
  - Simple, complex receptive fields
  - Orientation, direction, depth selectivity
  - Hierarchical organization
  - “What” and “Where”
  - Face selective cells

Highlight

"Jennifer Aniston cell"

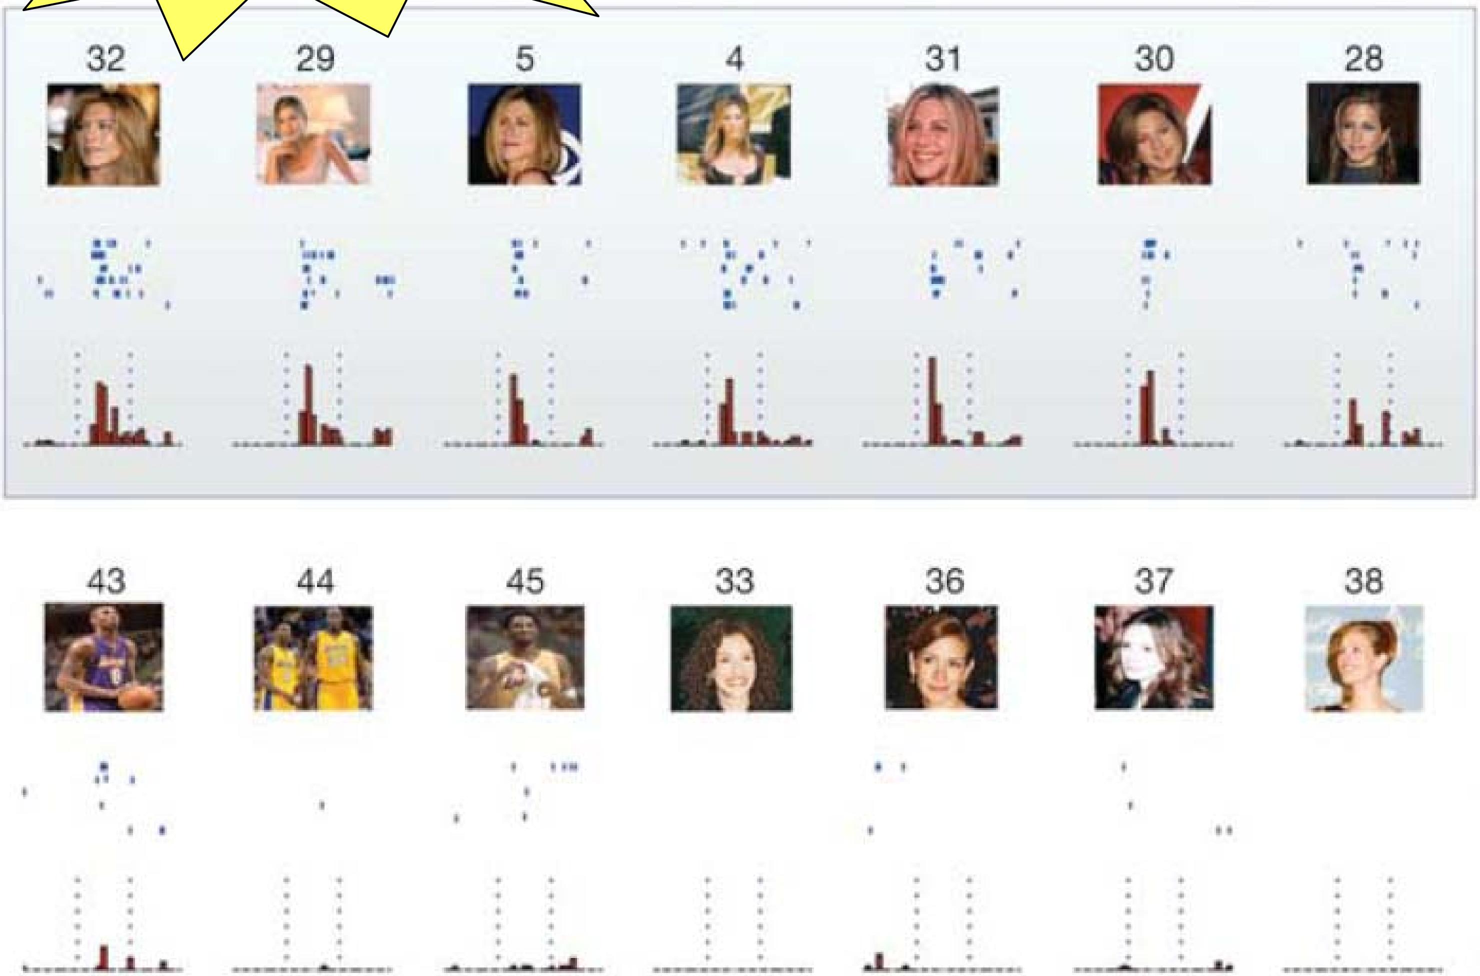

Cell in medial temporal lobe responded selectively to different images of JA
Other cells were found that responded to Bill Clinton, Halle Berry…
(Quiroga et al., 2005)

Human vision relies heavily on context

Sinha and Poggio, *Nature*, 1996

Human vision relies heavily on context

Sinha and Poggio, *Nature*, 1996

Change Blindness

Something big is changing in this scene – what is it?

http://www.psych.ubc.ca/~rensink/flicker/download/

Attention is needed to perceive changes in scenes

Other Highlights

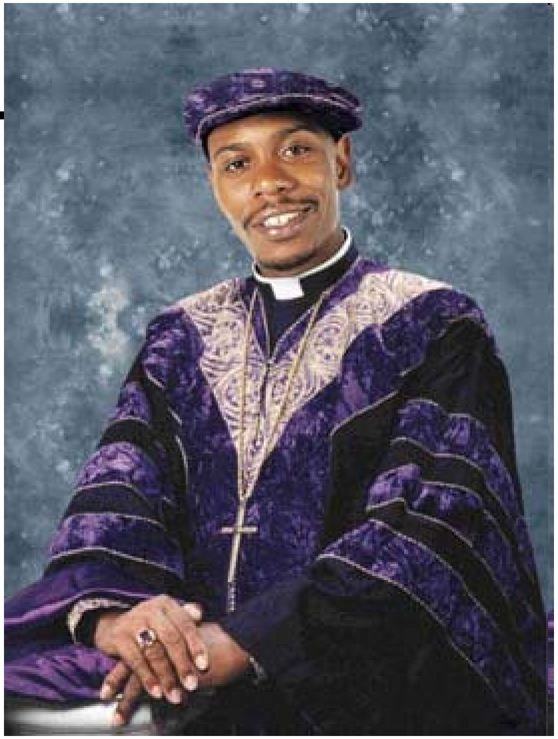

Your work!

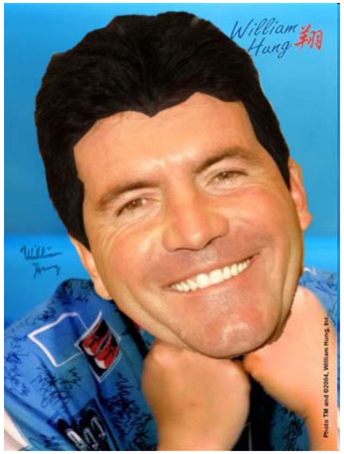

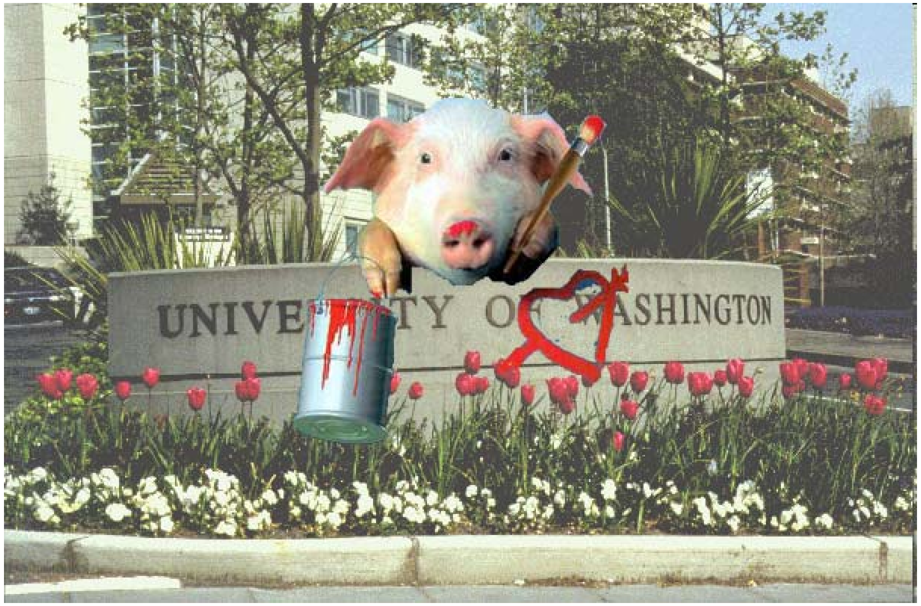

Pattern Recognition and Learning

- Decision tree learning
- Classification using nearest class mean
- Nearest neighbors and K-NN
- Perceptrons
- Multilayer Sigmoidal networks
  - Backpropagation learning (no need to derive)
- Max margin classification and SVMs
  - No need to derive but know the basic idea and intuition
  - Kernel trick
- Probabilistic (Bayesian) approach (skin classification example)
  - Histogram-based vs. parametric (Gaussian) PDFs
  - Bayes rule: prior, likelihood, and posterior
  - MAP vs. ML estimation

Color

- Color perception
- Color coding
  - RGB, color triangle, HSI, YIQ
- Color histograms for image matching
  - Histogram intersection
- Color constant color indexing

Texture

- Texture analysis
  - Edgeness, histograms, Laws texture energy features
  - Autocorrelation and Fourier power spectrum
- Texture synthesis
  - Markov random fields (MRFs)
  - Synthesizing one pixel at time (read Efros & Leung paper)
  - Block-based synthesis (read Efros & Freeman paper for basic idea)
  - Other methods: Graphcut textures, image analogies (know the basic idea, details not necessary)

Highlight

Texture transfer

+

parmesan

=

+

rice

=

Highlight

Image Analogies: Artistic Filters

A

A’

B

B’

Highlight

Texture-by-numbers

A

A’

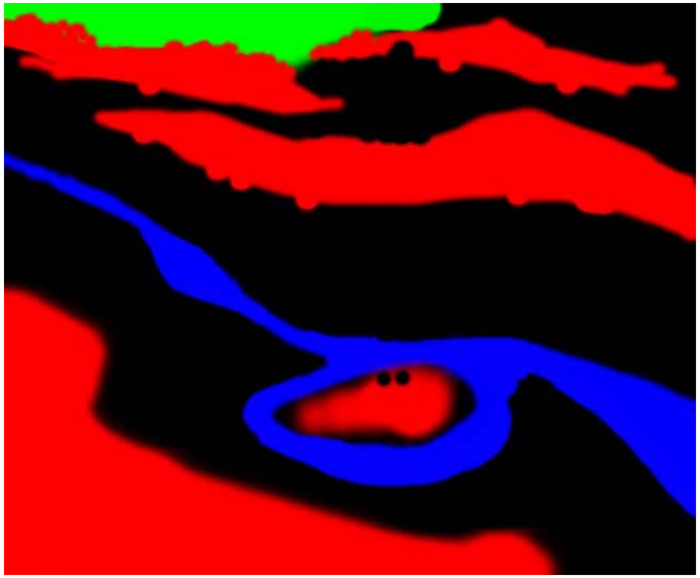

B

B’

Segmentation

- Histogram-based segmentation
- K-means clustering
  - EM algorithm
- Morphological operators
  - Dilation, erosion, opening, closing
- Graph-cut based segmentation
  - Normalized cuts (read Shi and Malik paper for basic idea and problem formulation, no need to know all the details)

Highlight

Graph-cut based segmentation

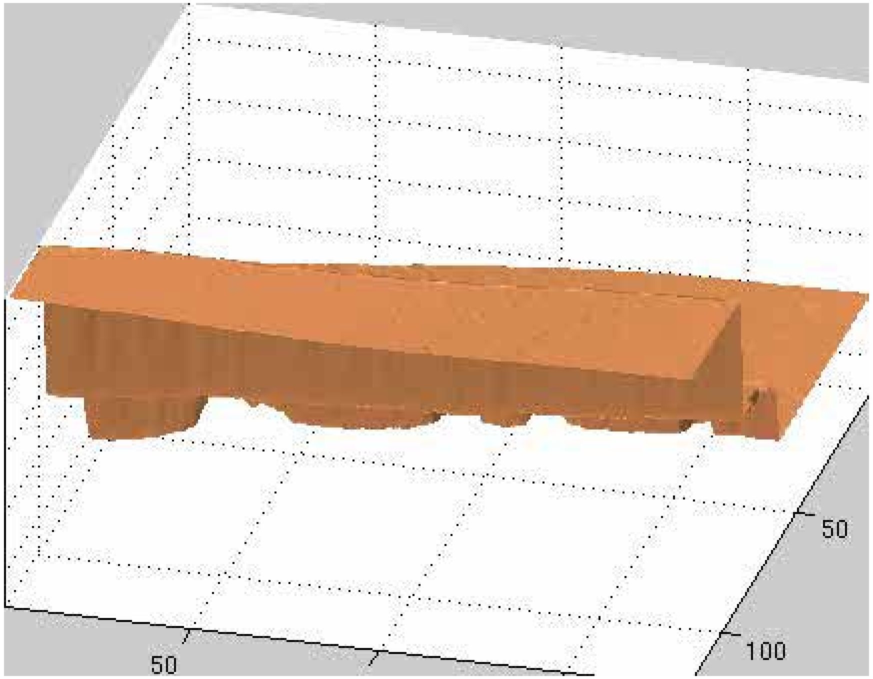

Other Highlights

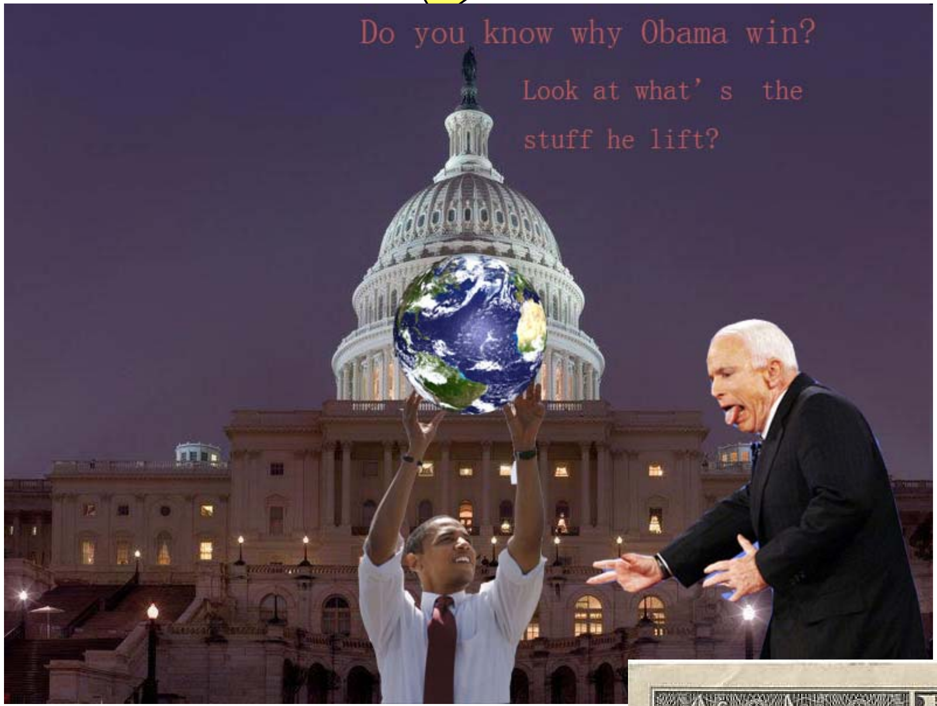

Your work!

Object Recognition

- Principal Component Analysis (PCA) for Recognition
  - Linear subspaces and dimensionality reduction
  - Eigenvectors and eigenvalues of data covariance matrix
  - Eigenfaces, reconstruction using eigenfaces
  - Recognition and detection with eigenfaces
- Object class recognition by parts (Fergus et al. paper)
  - Constellation model
  - Know the general idea: Generative probabilistic model that includes Location, Scale, and Appearance of Parts; EM for learning
  - Browse Fergus et al. paper (no need to know the details or math derivations, just understand their basic generative model)

Highlight

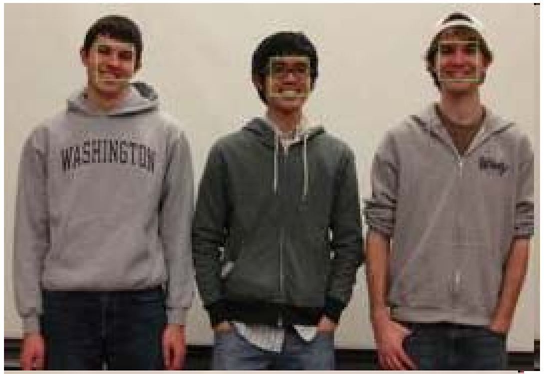

Your eigenfaces results here!

Motion

- Optical flow problem
  - Brightness constancy and small motion assumptions
  - Aperture problem
  - Lucas-Kanade technique
  - Using higher-order terms
  - Iterative L-K technique
- Coarse-to-fine optical flow using a pyramid
- Applications: Structure from motion, MPEG
- Feature tracking
  - Example application: Rotoscoping (know the basic idea, no need to know the details)

Highlights

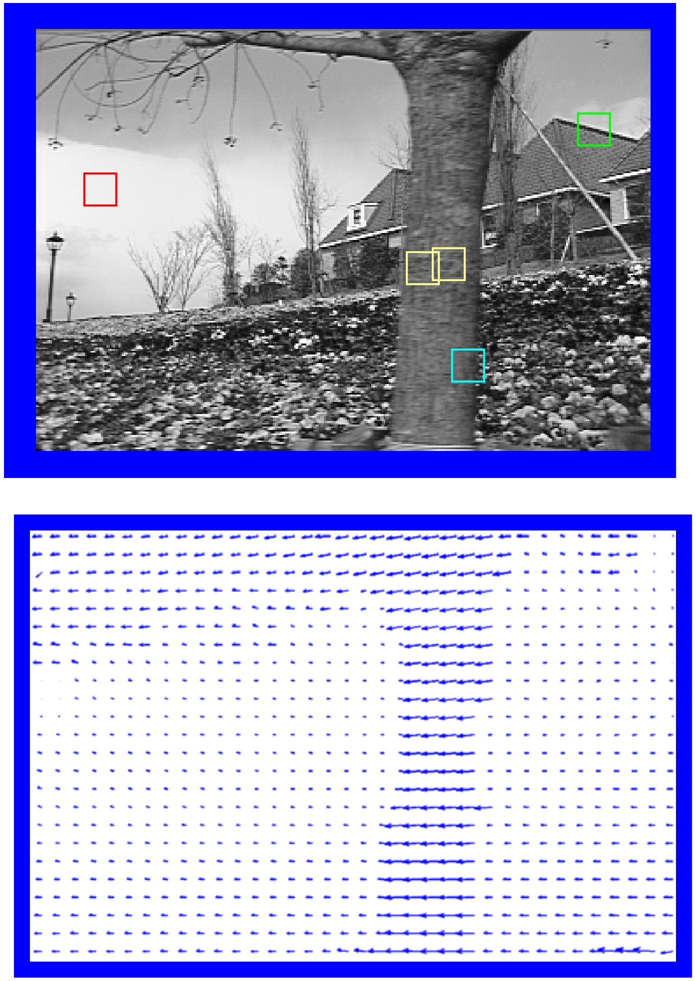

Optical flow

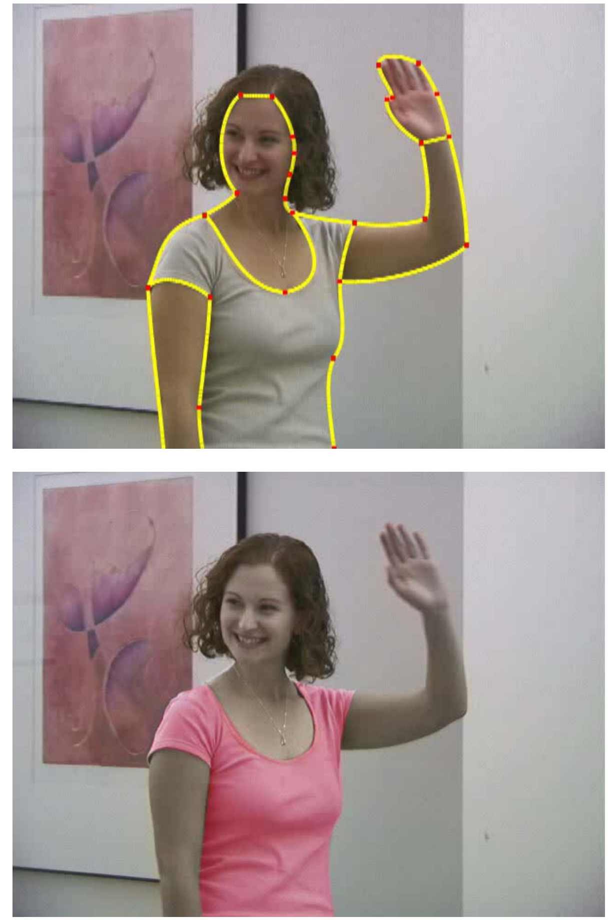

Feature tracking and Rotoscoping

Other Highlights

Your work!

Stereo

- Disparity and perception of depth
  - Stereoscopes, anaglyphs
- 3D from stereo: Triangulation method
  - Depth from disparity (derive from projection equations)
- Correspondence problem
  - Epipolar constraint
  - Stereo image rectification (know what it is, no need to know details from Loop and Zhang paper)
  - Window-based matching using SSD
  - Stereo correspondence as energy minimization
    - Know the two costs, no need to know details of actual algorithm
  - Sources of errors
- 3D reconstruction using structured light

Highlights

Video View Interpolation

(from “The Matrix”)

Applications

Content-Based Image Retrieval (CBIR)

- Browse through Prof. Shapiro’s lecture slides
- Know how features and object recognition techniques can be applied to CBIR
- Understand the use of key images to speed up retrieval

Tactile Graphics

- What is tactile graphics?
- How can computer vision techniques be used to solve problems in tactile graphics?
- Browse through Prof. Ladner’s lecture slides

Other Highlights

Other Highlights

Other Highlights

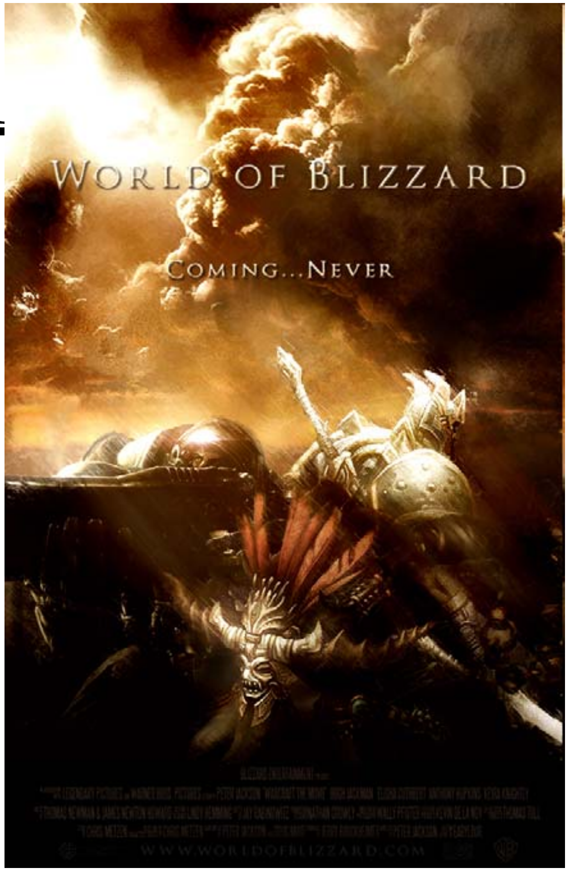

Other Highlights

Other Highlights

Good luck on the final exam!