Wrap-up

Class evaluation*s* – please take them!

- CMU's Faculty Course Evaluations (FCE): https://cmu.smartevals.com/

- 16-385 end-of-semester survey (see Piazza for link)

- Please take both, super helpful for developing future offerings of the class.

- Thanks in advance!

Course overview

1. Image processing. ← Lectures 1 – 6
   See also 18-793: Image and Video Processing

2. Geometry-based vision. ← Lectures 7 – 12
   See also 16-822: Geometry-based Methods in Vision

3. Semantic vision. ← Lectures 13 – 17
   See also 16-824: Vision Learning and Recognition
   See also 10-703: Deep Reinforcement Learning

4. Dealing with motion. ← Lectures 18 – 20
   See also 16-831: Statistical Techniques in Robotics
   See also 16-833: Robot Localization and Mapping

5. Physics-based vision. ← Lectures 21 – 25
   See also 16-823: Physics-based Methods in Vision
   See also 15-462: Computer Graphics
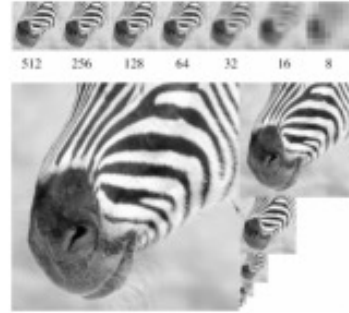   See also 15-463: Computational Photography

Image processing

Image filtering
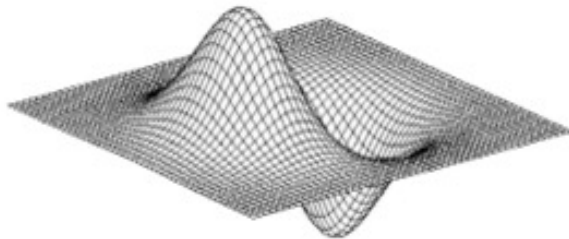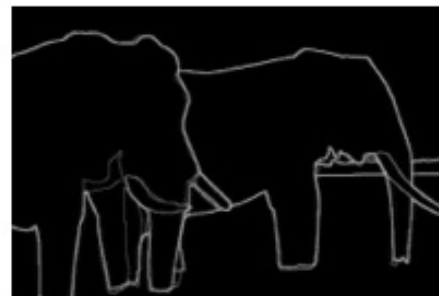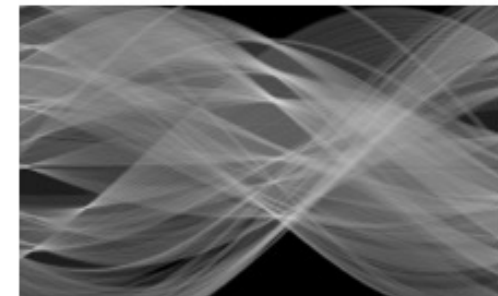
image pyramids

Fourier filtering
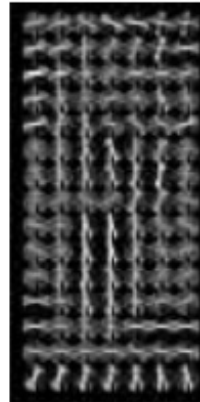
Image gradients

Boundaries

Hough Transform

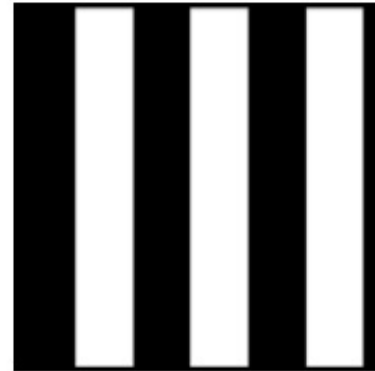# Image features

Corner detection   Multi-scale detection
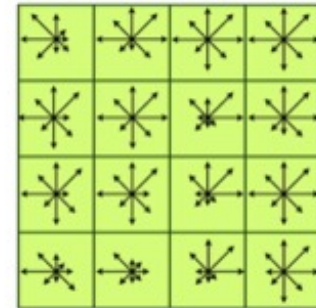

Haar-like   HOG   SURF   SIFT
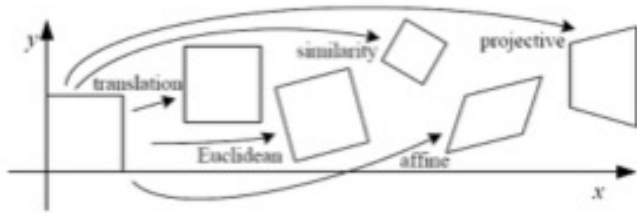
2D alignment

Figure 1: Basic set of 2D planar transformations

2D Transforms

DLT

RANSAC

Homography

Camera and multi-view geometry
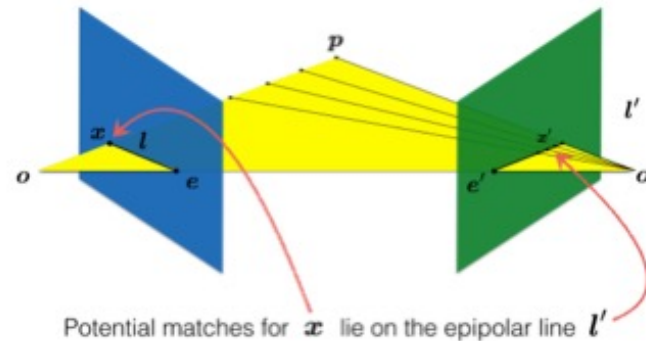
$$x = \mathbf{PX}$$

camera matrix

$$\mathbf{P}$$

pose estimation

$$\mathbf{X}$$

triangulation

$$\mathbf{F}$$

fundamental matrix

Potential matches for x lie on the epipolar line l'
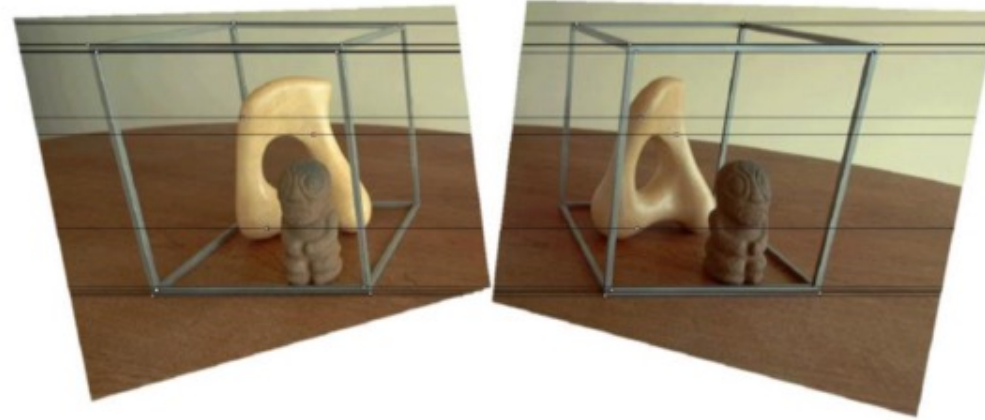
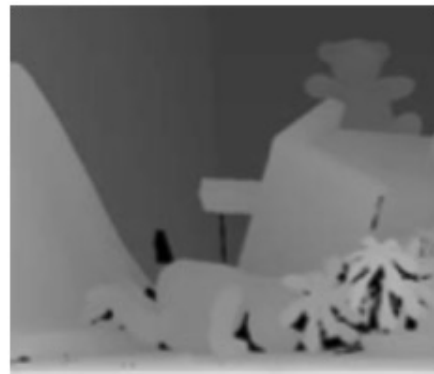epipolar geometry

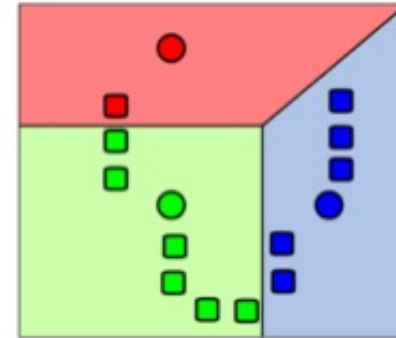Reconstruction

Stereo
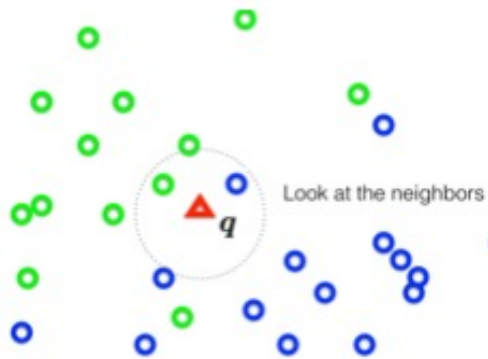

Stereo Rectification
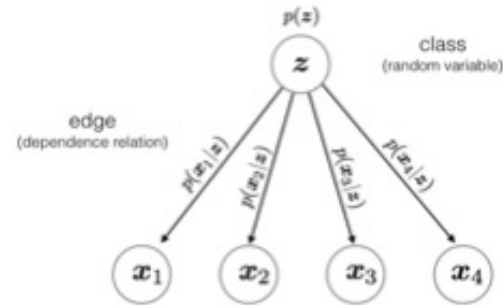

Block matching

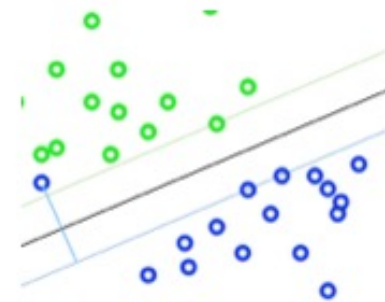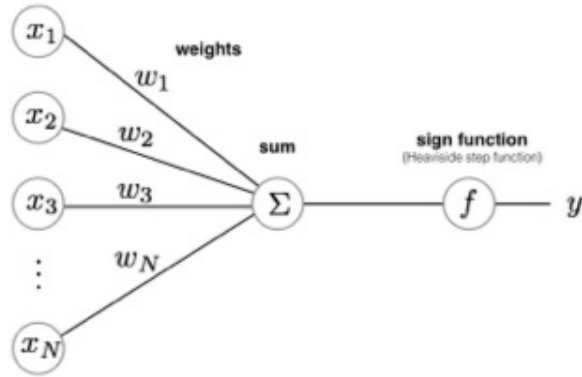Energy minimization

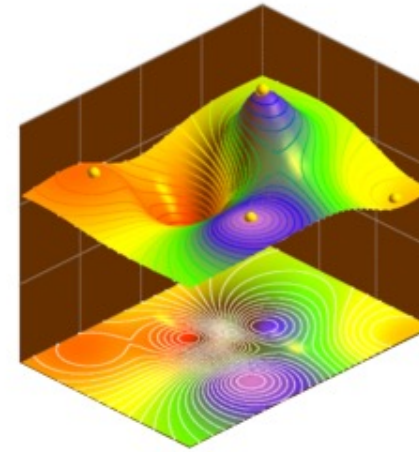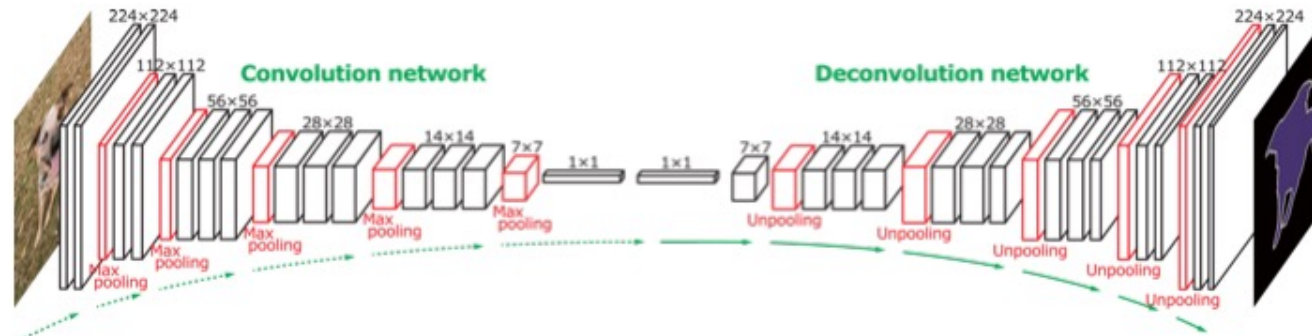# Object recognition

Bag-of-words

K-means

Nearest Neighbor

Naive Bayes

SVM

Neural networks

Perceptron

Gradient Decent

Convolutional Neural Networks

Optical flow and alignment

$$
\begin{bmatrix}
I_x(\boldsymbol{p}_1) & I_y(\boldsymbol{p}_1) \\
I_x(\boldsymbol{p}_2) & I_y(\boldsymbol{p}_2) \\
\vdots & \vdots \\
I_x(\boldsymbol{p}_{25}) & I_y(\boldsymbol{p}_{25})
\end{bmatrix}
\begin{bmatrix}
u \\
v
\end{bmatrix}
= -
\begin{bmatrix}
I_t(\boldsymbol{p}_1) \\
I_t(\boldsymbol{p}_2) \\
\vdots \\
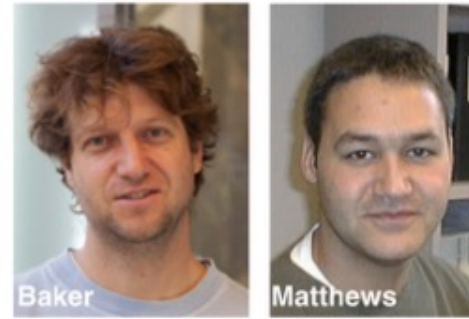I_t(\boldsymbol{p}_{25})
\end{bmatrix}
$$

Constant Flow

$$
\min_{\boldsymbol{u}, \boldsymbol{v}} \sum_{ij} \left\{ E_d(i,j) + \lambda E_s(i,j) \right\}
$$

Horn-Schunck

Lucas-Kanade
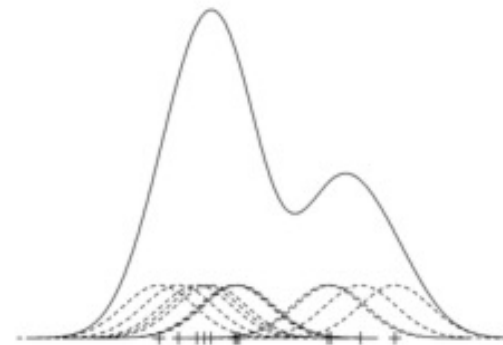(Forward additive)

Baker-Matthews
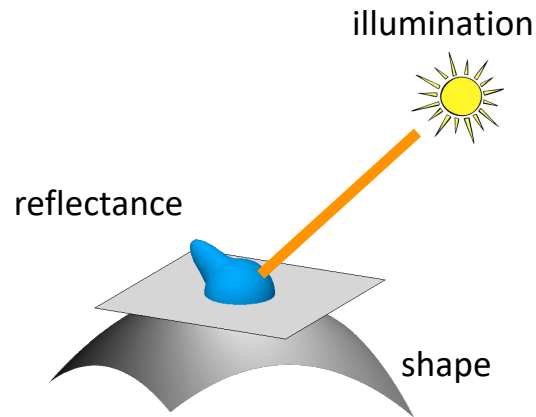(Inverse Compositional)

Tracking in videos

KLT

Mean shift

Image formation and physics
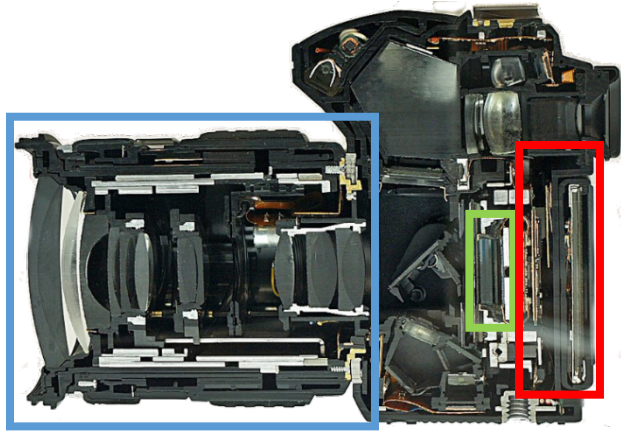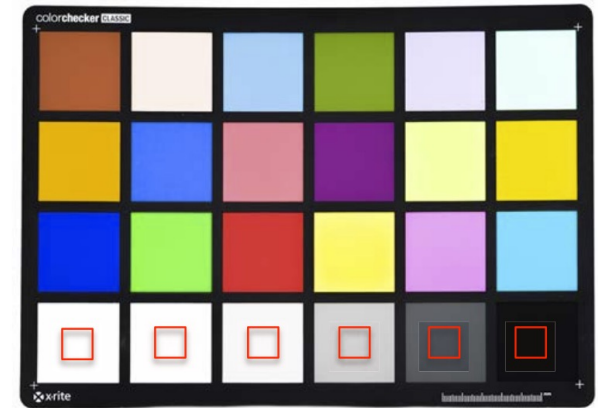

Radiometry and image formation

Photometric stereo

Image processing pipeline

Radiometric and color calibration

Things you should know how to do

1. Detect lines (circles, shapes) in an image.

2. Perform automatic image warping and basic AR.

3. Reconstruct 3D scene structure from two images.

4. Do photometric stereo and render simple images.

5. Recognize objects using a bag-of-words model.

6. Recognize objects using deep CNNs.

7. Track objects in video.

Questions?

Do you plan on taking any other vision courses?

Which part of the class did you like the most?

Which part of the class did you like the least?

Any topics you wanted to learn more about?

Any topics you wanted to learn less about?

Which was your favorite programming assignment?

Which was your least favorite programming assignment?

Did you find the take-home quizzes useful?

How does course workload compare to other classes?