NEURL-GA 3042.005 -- Spring, 2014
(cross-listed as: MATH-GA 2855, CSCI-GA 2715, PSYCH-GA 3405.005)

Representation and Analysis of Visual Images

Eero Simoncelli

00 - Introduction

Logistics

- Tuesdays, 10-12, Meyer 851
- web: http://www.cns.nyu.edu/~eero/imrep-course/
- Prerequisites: linAlg, vectorCalc, linear systems / Fourier, probability, matlab programming.
- Grade based on homework sets
- Today's handouts: course description, poll

Photographic Images

Diverse specialized structures:

- edges/lines/contours
- shadows/highlights
- smooth regions
- textured regions

Photographic Images

67	223	103	231	200	147	166	176
205	148	19	242	100	15	187	47
110	141	61	126	62	60	166	94
233	37	32	125	103	90	115	160
47	218	47	86	25	210	140	200
68	159	61	230	34	4	76	21
37	90	107	95	241	11	191	238
35	131	13	28	245	43	48	199

Edges?

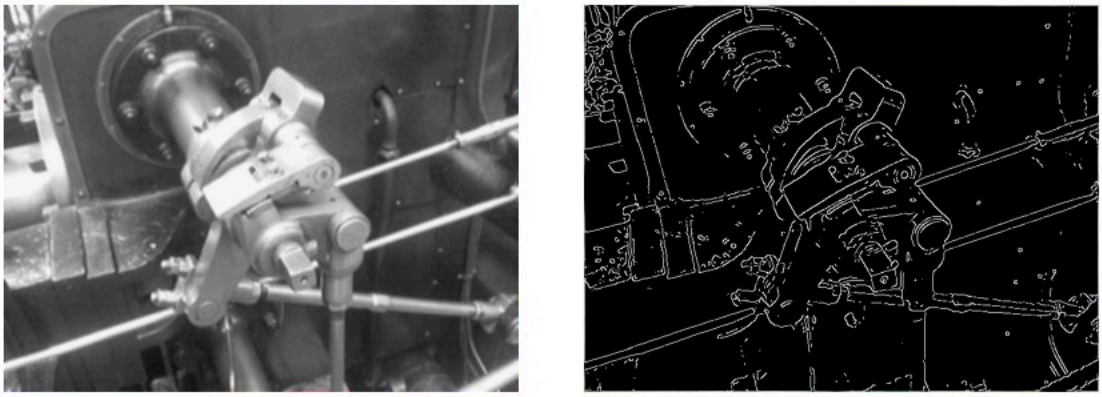

Canny edge detector (1986)

Edges?

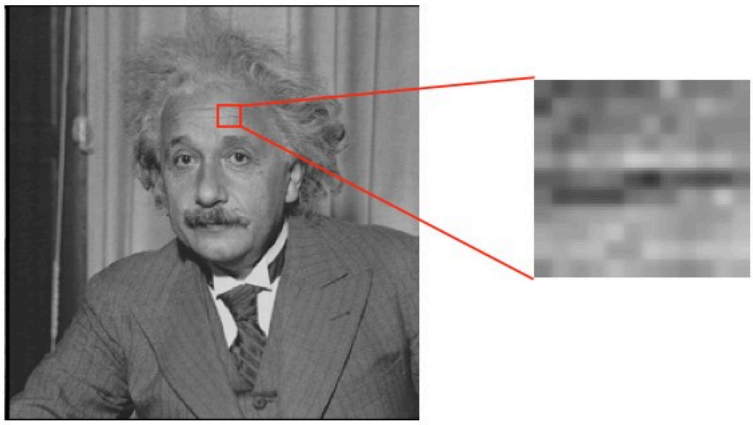

- Step edges are rare (lighting, junctions, texture, noise)
- One scale's texture is another scale's edge
- Need seamless transitions from isolated features to dense textures

What makes a good representation?

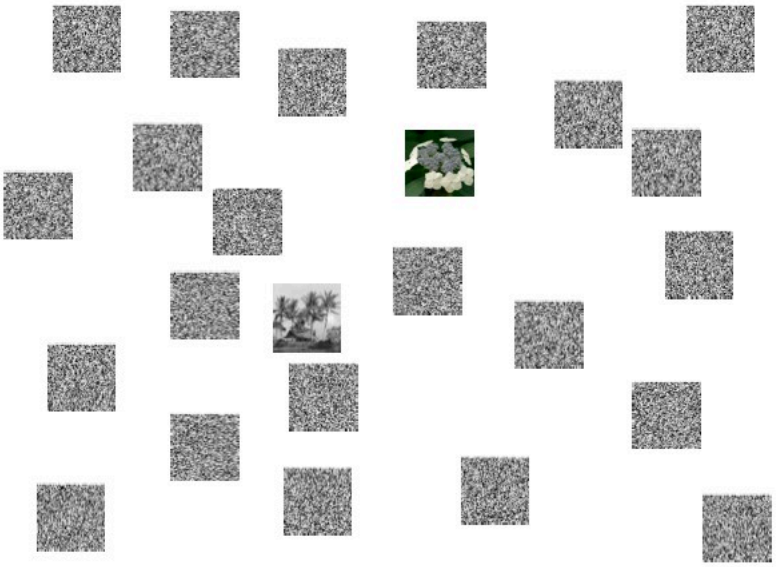

The distribution of natural images is complicated. Perhaps it is something like beer foam, which is mostly empty but contains a thin mesh-work of fluid which fills the space and occupies almost no volume.

- Ruderman, 1996

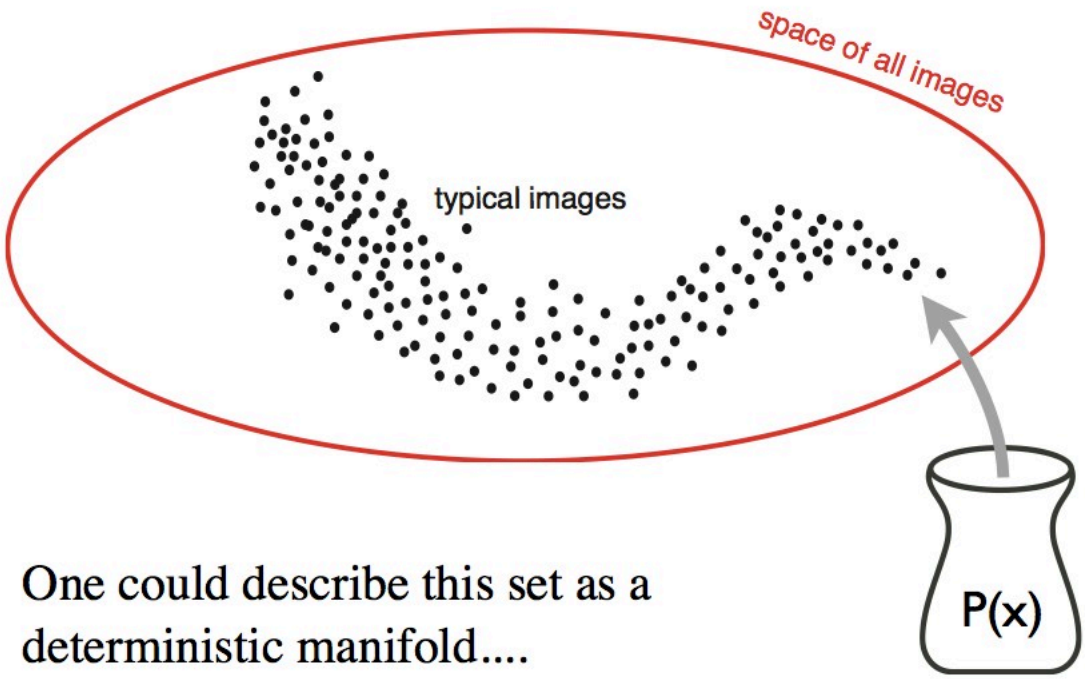

One could describe this set as a deterministic manifold....

But seems more natural to use probability

Density models

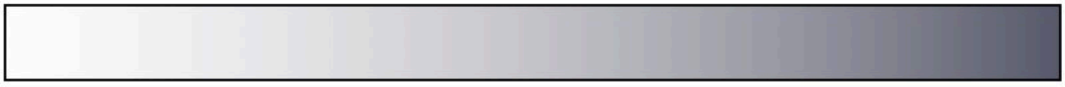

nonparametric

historical trend
(technology driven)

parametric/
constrained

build a histogram from lots of observations...

use "natural constraints" (geometry/photometry of image formation, computation, maxEnt)

Some big numbers:

seconds since big bang: 10^{17}

atoms in the visible universe: 10^{80}

10-bit images, 1000x1000: $10^{3,000,000}$

we can't enumerate them => we need models

What makes a good analysis?

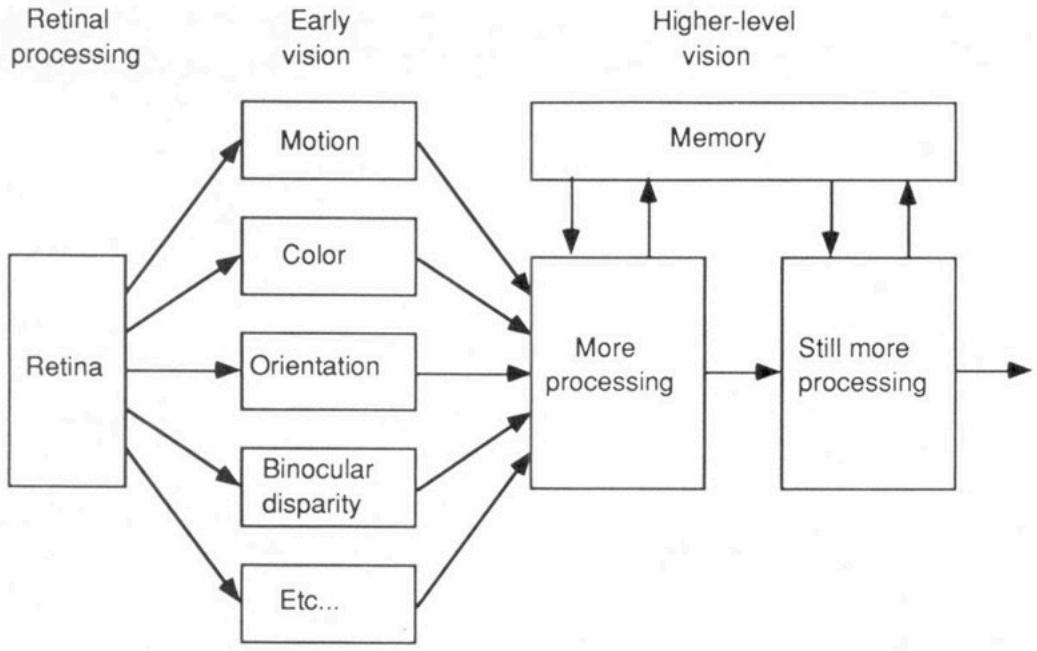

[Adelson & Bergen, 1990]

General methodology

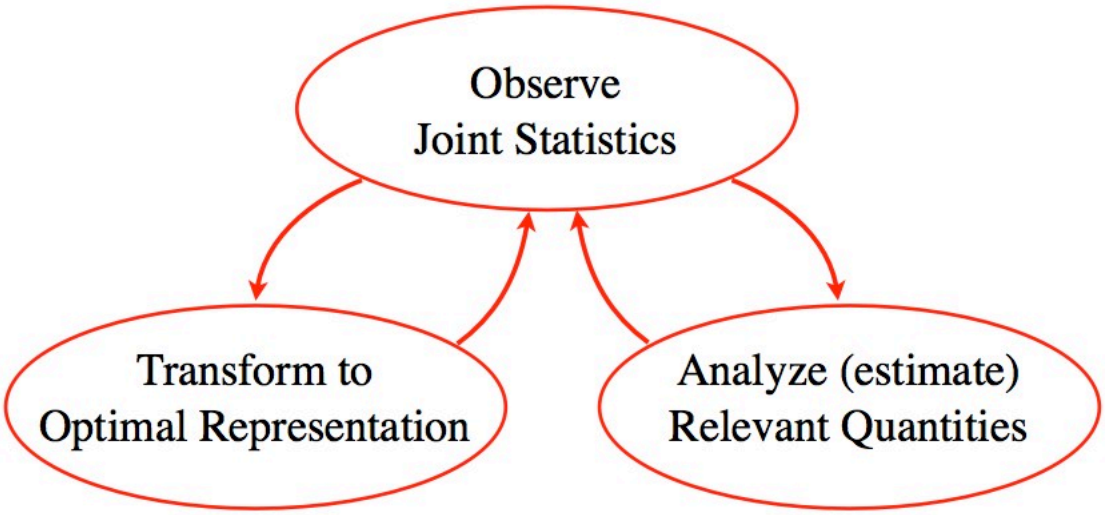

Image Analysis/Representation

- Engineering: compression, denoising, restoration, enhancement/modification, synthesis, manipulation
- Science: optimality principles for neurobiology (evolution, development, learning, adaptation)

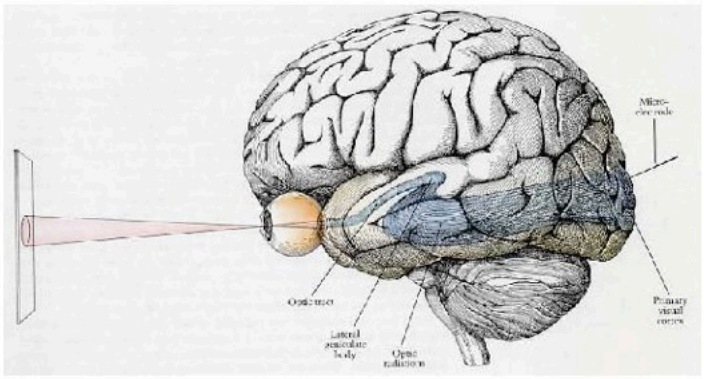

[figure: Hubel]

the course is *not* primarily about

- photography
- computer vision
- machine learning
- biological vision

... but covers fundamentals that underly all of these

Course priorities/emphasis

- conceptual/geometric
- mathematical/statistical
- algorithmic

Basic tools:
  convolution, Fourier, sampling, boundaries in 2D
  rotation-invariance, derivs, orientation tensor, steerability, simple texture
  multi-scale representations, wavelets

Color estimation:
  Demosaicing (color/space joint models)
  Bayes color constancy

Motion estimation
  differential optic flow
  coarse-to-fine estimation

Statistically optimal representation:
  Compression, rate-distortion
  Local stats: PCA, ICA, RG,
  Contextual statistics: GSM, orientation, phase,

Image estimation (denoising)
  Classical (Wiener)
  Shrinkage, thresholding, coring
  Contextual modeling
  MRFs

Texture representation:

Image Segmentation
  min-cuts
  EM and mixture models

Questions?