BIL 717
Image Processing

Feb. 8, 2016

Introduction

Erkut Erdem

Hacettepe University

Computer Vision Lab (HUCVL)

Today

- About me
- About you
- Introduction to Image Processing
- Course outline and logistics

Today

- About me
- About you
- Introduction to Image Processing
- Course outline and logistics

About me

- Asst. Prof. Erkut Erdem
  http://web.cs.hacettepe.edu.tr/~erkut/

Hacettepe University
Faculty Member
2010-now

Middle East Technical University
1997-2008
Ph.D., 2008
M.Sc., 2003
B.Sc., 2001

Ecole Nationale Supérieure des Télécommunications
Post-doctoral Researcher
2009-2010

University of California Los Angeles
Visiting Researcher
Oct. 2007 - Dec. 2007

About me

- Asst. Prof. Erkut Erdem
  http://web.cs.hacettepe.edu.tr/~erkut/

HACETTEPE UNIVERSITY
COMPUTER VISION LAB

http://vision.cs.hacettepe.edu.tr/

About my research

- My research centers on the areas of computer vision and machine learning.
- specifically interested in the role of context in visual processing.

- I try to incorporate different kinds of context (*spatial*, *temporal* and/or *cross-modal*) into all levels of visual processing from low to mid and high-level vision.

About my research

image smoothing

image colorization

visual saliency

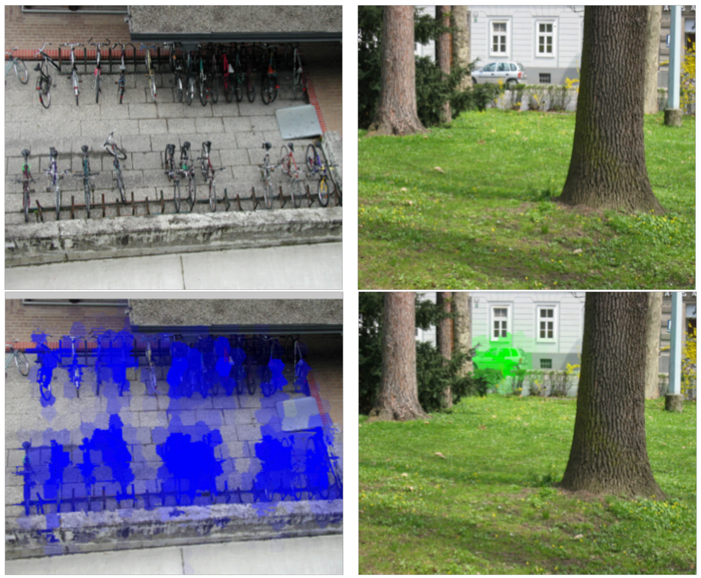

object segmentation

Today

- About me
- About you
- Introduction to Image Processing
- Course outline and logistics

Introduce yourselves

- Who are you?
- What do you know about image processing?
- Why you want to take BIL717?

Today

- About me
- About you
- Introduction to Image Processing
- Course outline and logistics

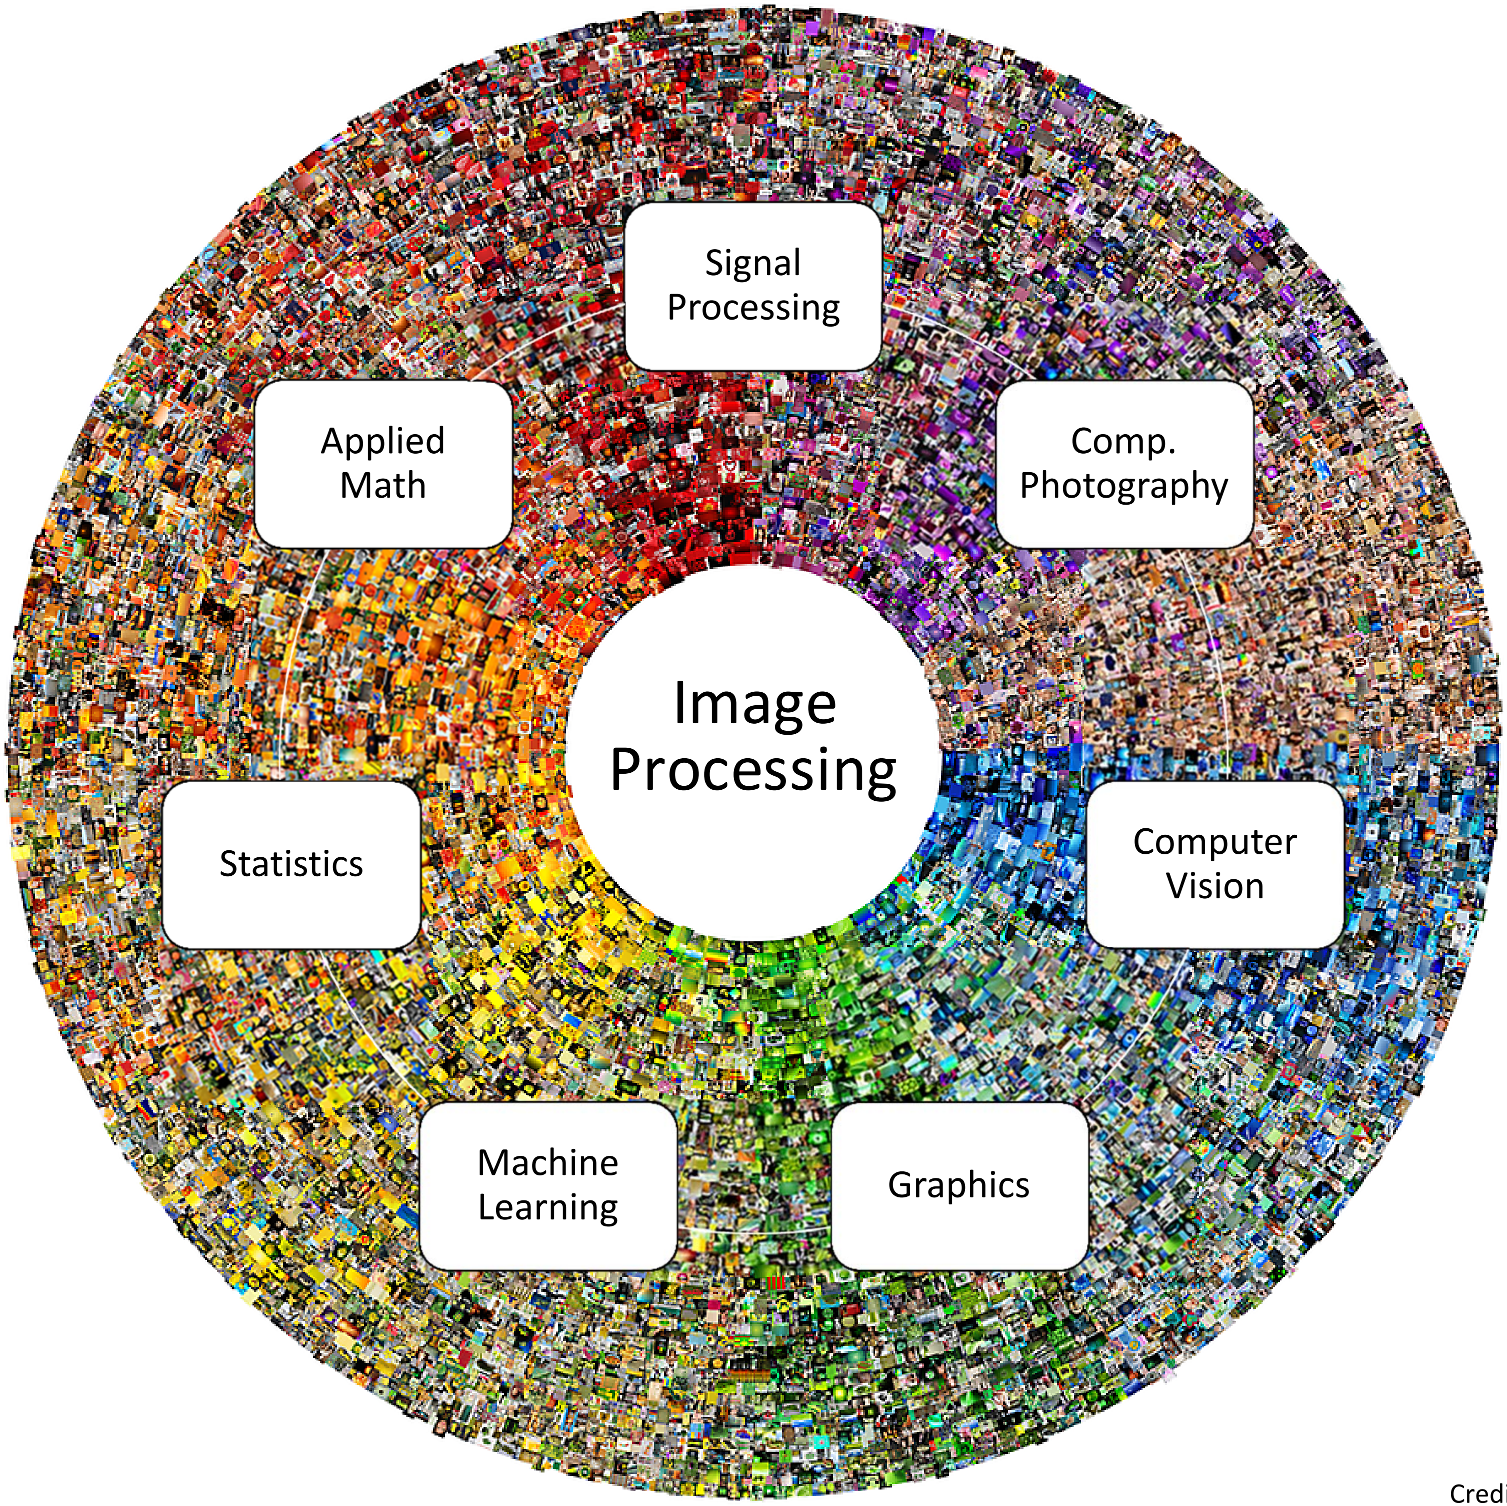

Credit: P. Milanfar

What does it mean, to see?

- *"The plain man's answer (and Aristotle's, too) would be, to know what is where by looking. In other words, vision is the process of discovering from images what is present in the world, and where it is."* David Marr, Vision, 1982

- Our brain is able to use an image as an input, and interpret it in terms of objects and scene structures.

What does Salvador Dali's *Study for the Dream Sequence in Spellbound (1945)* say about our visual perception?

What does Paul Signac's *Place des Lices (1893)* say about our visual perception?

Why does vision appear easy to humans?

- Our brains are specialized to do vision.
- Nearly half of the cortex in a human brain is devoted to doing vision (cf. motor control ~20-30%, language ~10-20%)
- *“Vision has evolved to convert the ill-posed problems into solvable ones by adding premises: assumptions about how the world we evolved in is, on average, put together”*
  Steven Pinker, How the Mind Works, 1997
- Gestalt Theory (Laws of Visual Perception), Max Wertheimer, 1912

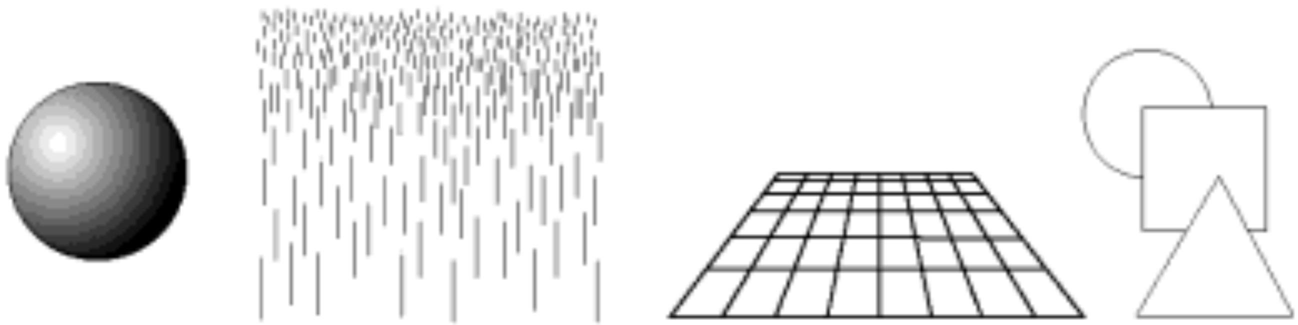

Figures: Steven Pinker, How the Mind Works, 1997

Computer Vision

- “Vision is a process that produces from images of the external world a description that is useful to the viewer and not cluttered with irrelevant information” ~David Marr

- The goal of Computer Vision:
  To develop artificial machine vision systems that make inferences related to the scene being viewed through the images acquired with digital cameras.

Marr's observation: Studying vision at three different levels

- Vision as an information processing task
  [David Marr, 1982]

- Three levels of understanding:

1. Computational theory
   - What is computed? Why it is computed?
2. Representation and Algorithm
   - How it is computed?
   - Input, Output, Transformation
3. Physical Realization
   - Hardware

Visual Modules and the Information Flow

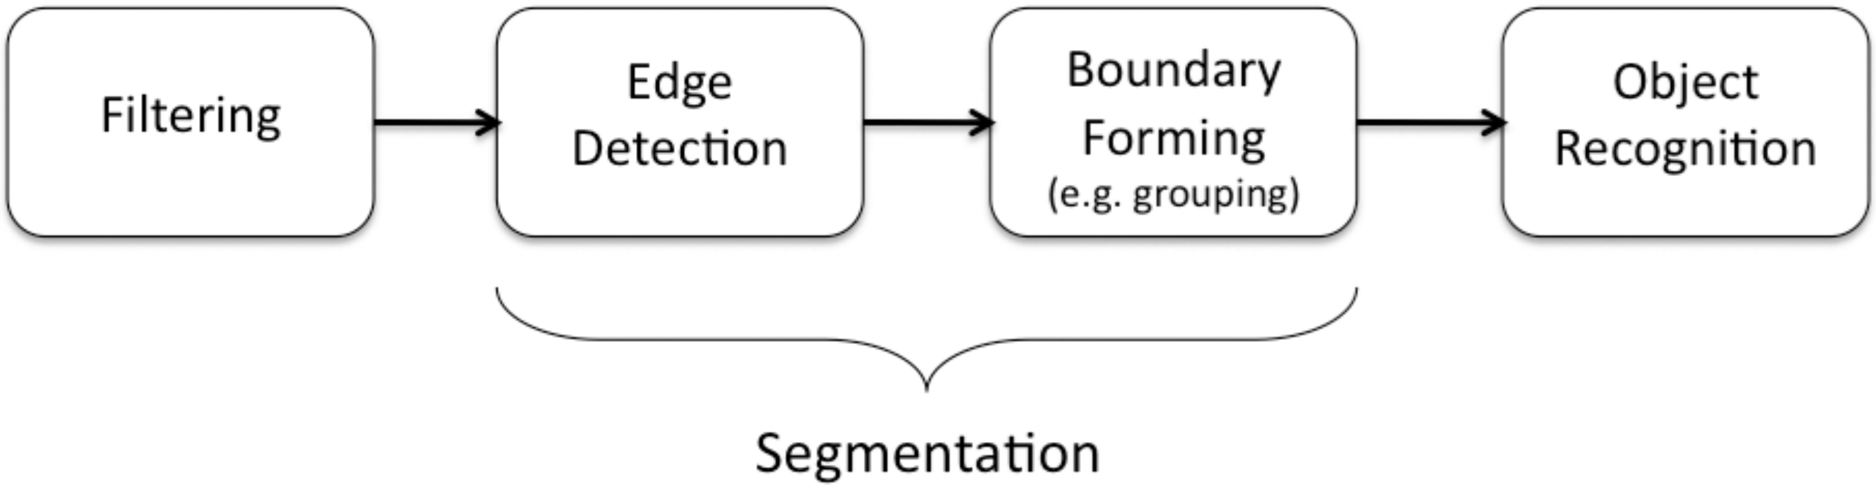

- Visual perception as a data-driven, bottom-up process (traditional view since D. Marr)
- Unidirectional information flow
- Simple low-level cues >> Complex abstract perceptual units

Visual Modules and the Information Flow

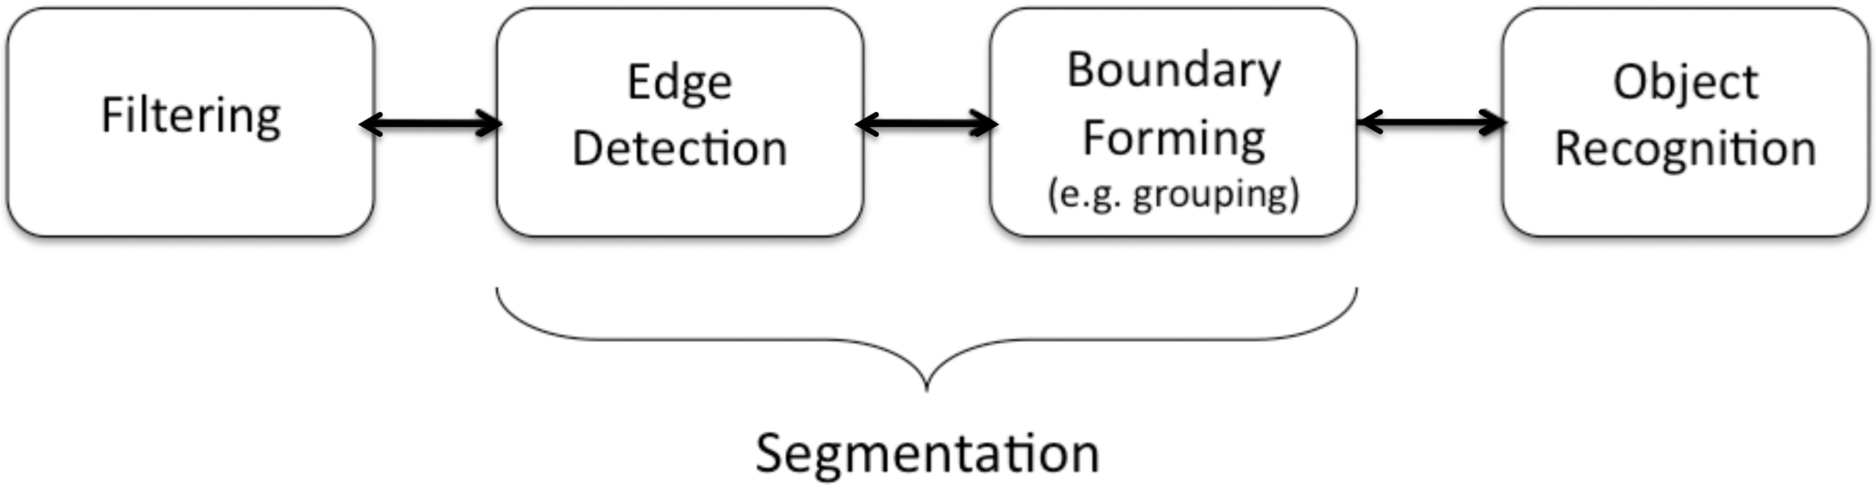

- Vision modules can be categorized into three groups according to their functionality:
  - Low-level vision: filtering out irrelevant image data
  - Mid-level vision: grouping pixels or boundary fragments together
  - High-level vision: complex cognitive processes

Visual Modules and the Information Flow

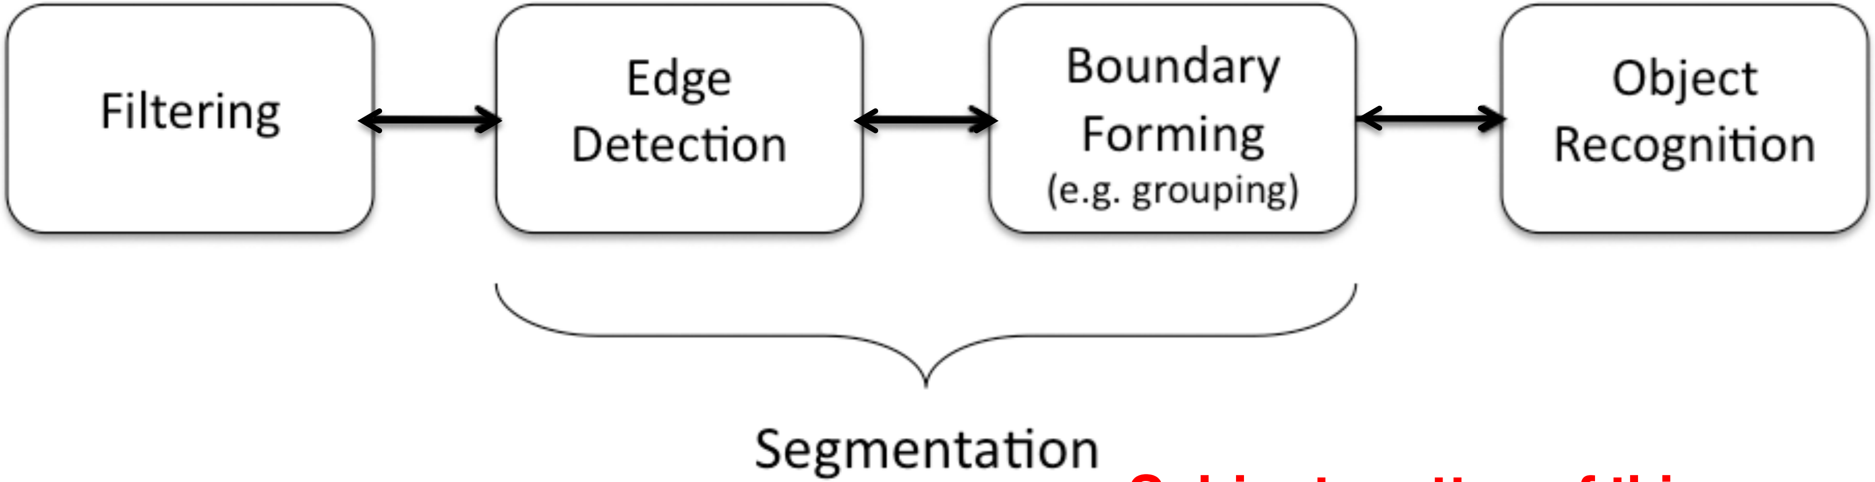

Subject matter of this course

- Vision modules can be categorized into three groups according to their functionality:
  - **Low-level vision: filtering out irrelevant image data**
  - **Mid-level vision: grouping pixels or boundary fragments together**
  - High-level vision: complex cognitive processes

Fundamentals of Image Processing

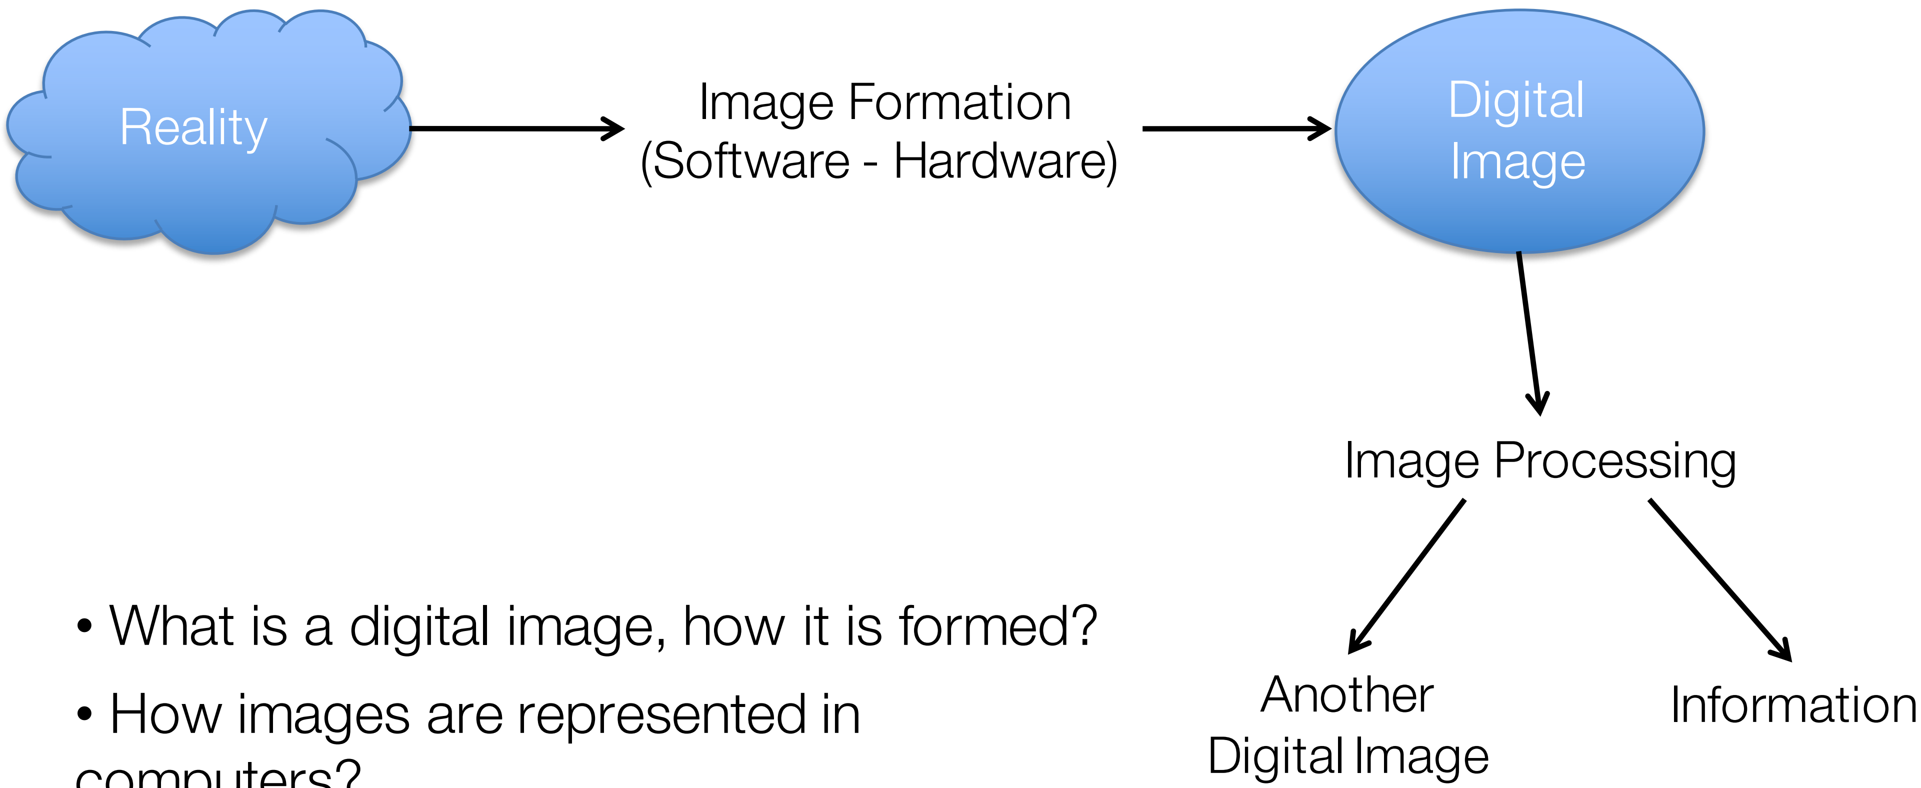

- What is a digital image, how it is formed?
- How images are represented in computers?
- Why we process images?
- How we process images?

Image Formation

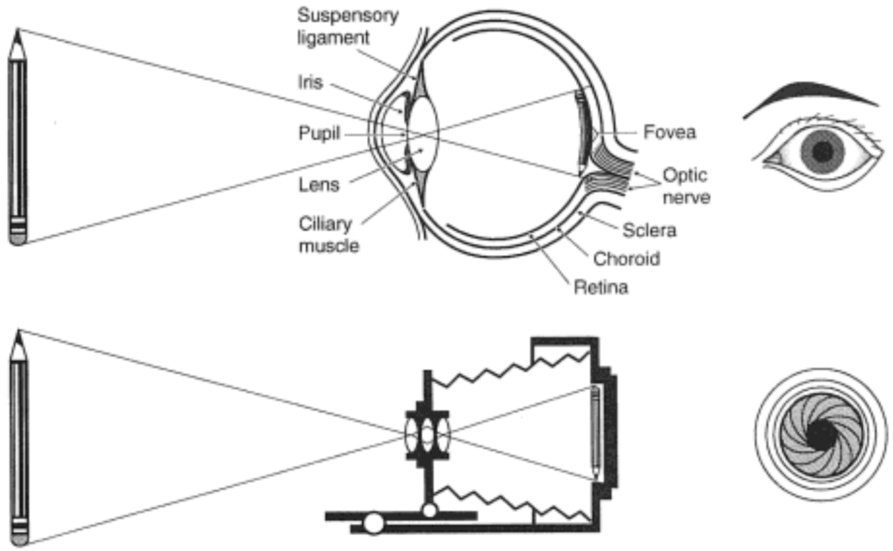

Three Dimensional World → Two Dimensional Image Space

- What is measured in an image location?
  - brightness
  - color

<< viewpoint
illumination conditions
local geometry
local material properties

Figures: Francis Crick, The Astonishing Hypothesis, 1995

Image Formation

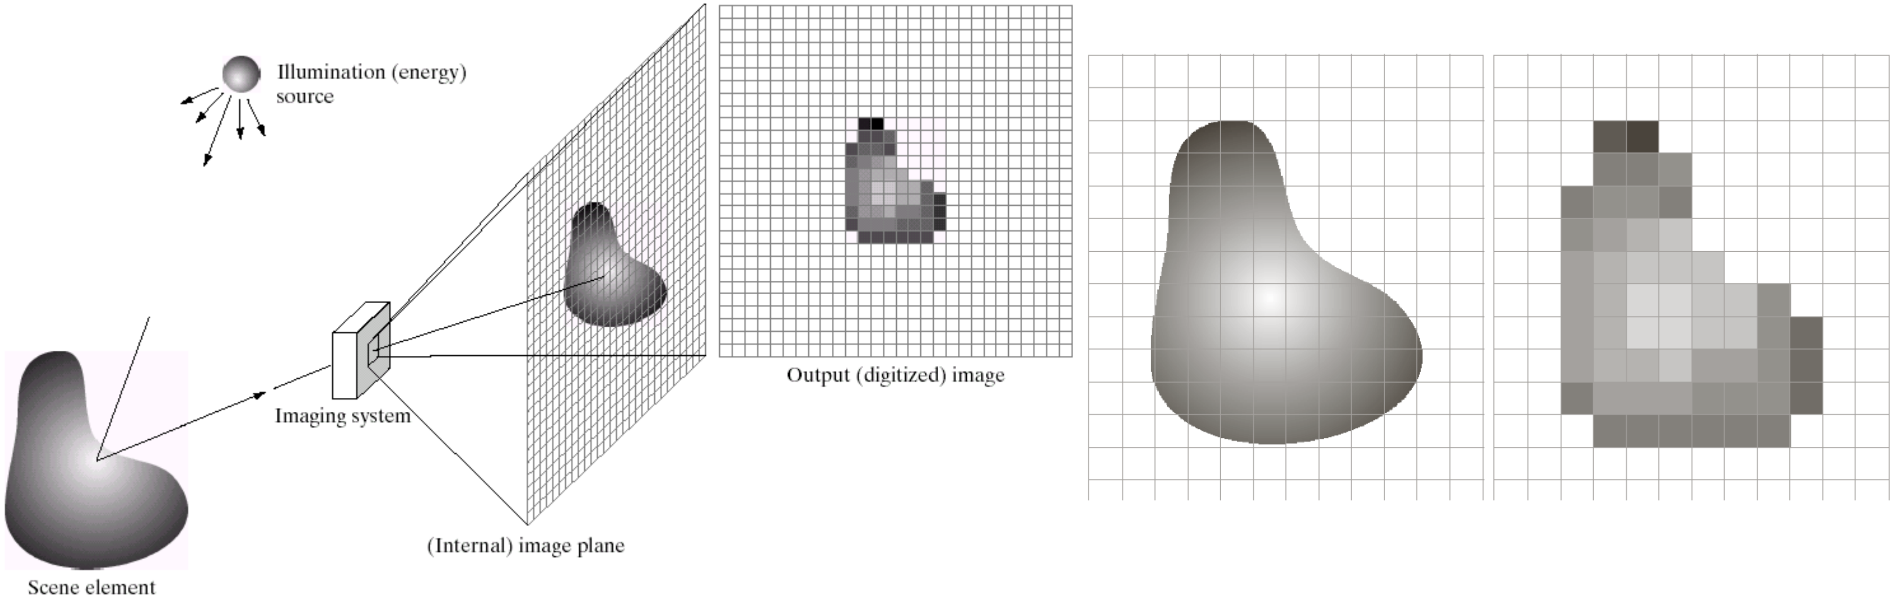

- Discretization
  - in image space - sampling
  - In image brightness - quantization

Figures: Gonzalez and Woods, Digital Image Processing, 3rd Edition, 2008

Image Representation

- **Digital image:** 2D discrete function f
- **Pixel**: Smallest element of an image f(x,y)

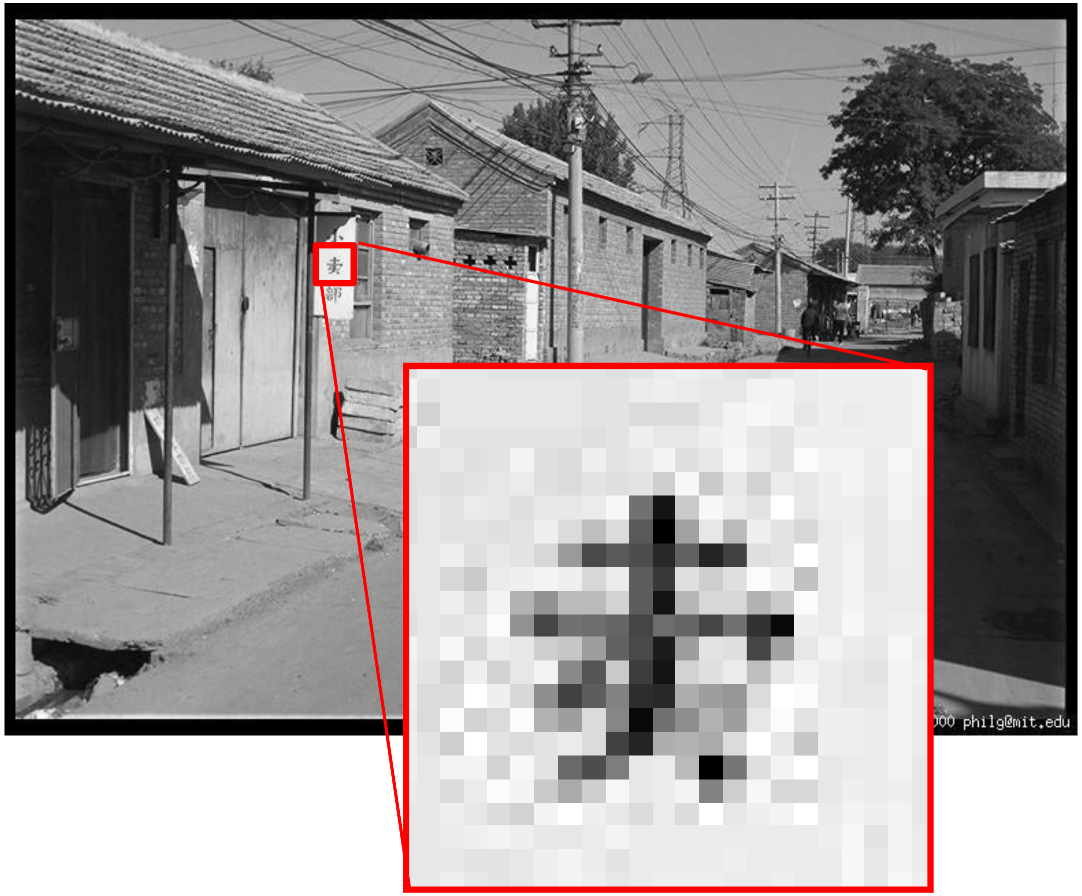

Figure: M. J. Black

Image Representation

- **Digital image:** 2D discrete function f
- **Pixel**: Smallest element of an image f(x,y)

0.92	0.93	0.94	0.97	0.62	0.37	0.85	0.97	0.93	0.92	0.99
0.95	0.89	0.82	0.89	0.56	0.31	0.75	0.92	0.81	0.95	0.91
0.89	0.72	0.51	0.55	0.51	0.42	0.57	0.41	0.49	0.91	0.92
0.96	0.95	0.88	0.94	0.56	0.46	0.91	0.87	0.90	0.97	0.95
0.71	0.81	0.81	0.87	0.57	0.37	0.80	0.88	0.89	0.79	0.85
0.49	0.62	0.60	0.58	0.50	0.60	0.58	0.50	0.61	0.45	0.33
0.86	0.84	0.74	0.58	0.51	0.39	0.73	0.92	0.91	0.49	0.74
0.96	0.67	0.54	0.85	0.48	0.37	0.88	0.90	0.94	0.82	0.93
0.69	0.49	0.56	0.66	0.43	0.42	0.77	0.73	0.71	0.90	0.99
0.79	0.73	0.90	0.67	0.33	0.61	0.69	0.79	0.73	0.93	0.97
0.91	0.94	0.89	0.49	0.41	0.78	0.78	0.77	0.89	0.99	0.93

Figure: M. J. Black

Sample Problems and Techniques

- Edge Detection
- Image Denoising
- Image Smoothing
- Image Deblurring
- Image Segmentation
- Visual Saliency
- Semantic Segmentation

- PDEs
- Variational models
- MRFs
- Graph Theory
- Sparse Coding
- Deep Learning

Image Filtering

- Filtering out the irrelevant information

$$f(x) = u(x) + n(x)$$

$f(x)$ ↓ observed image

$u(x)$ ↓ desired image

$n(x)$ ↓ irrelevant data

- Image denoising, image sharpening, image smoothing, image deblurring, etc.
- Edge detection

Edge Detection

Canny edge detector

- Edges: abrupt changes in the intensity
  - Uniformity of intensity or color
- Edges to object boundaries

Image Filtering

- **<u>Difficulty:</u>** Some of the irrelevant image information have characteristics similar to those of important image features

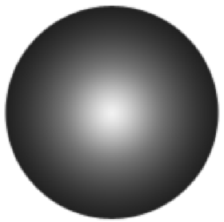

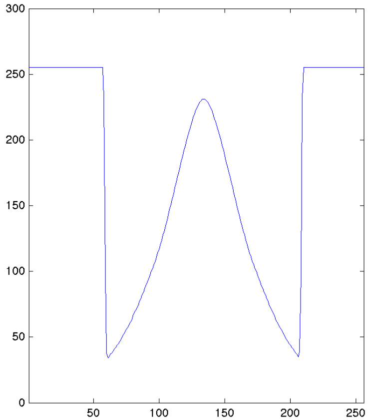

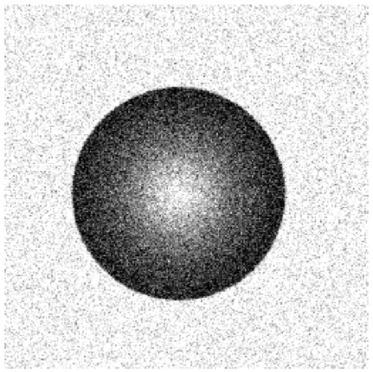

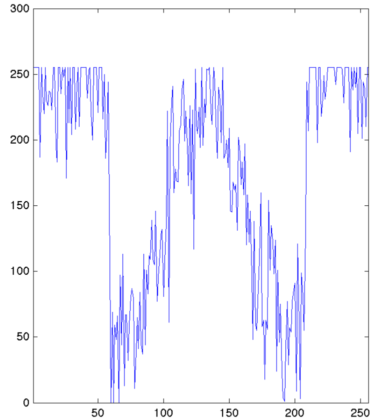

Image Smoothing - A Little Bit of History

- Gaussian Filtering / linear diffusion $$\frac{\partial u}{\partial t} = \nabla \cdot (\nabla u) = \nabla^2 u$$
  - the most widely used method

- mid 80's – unified formulations
  - methods that combine smoothing and edge detection
  - Geman & Geman'84, Blake & Zisserman'87, Mumford & Shah'89, Perona & Malik'90

Image Denoising

- Images are corrupted with 70% salt-and-pepper noise

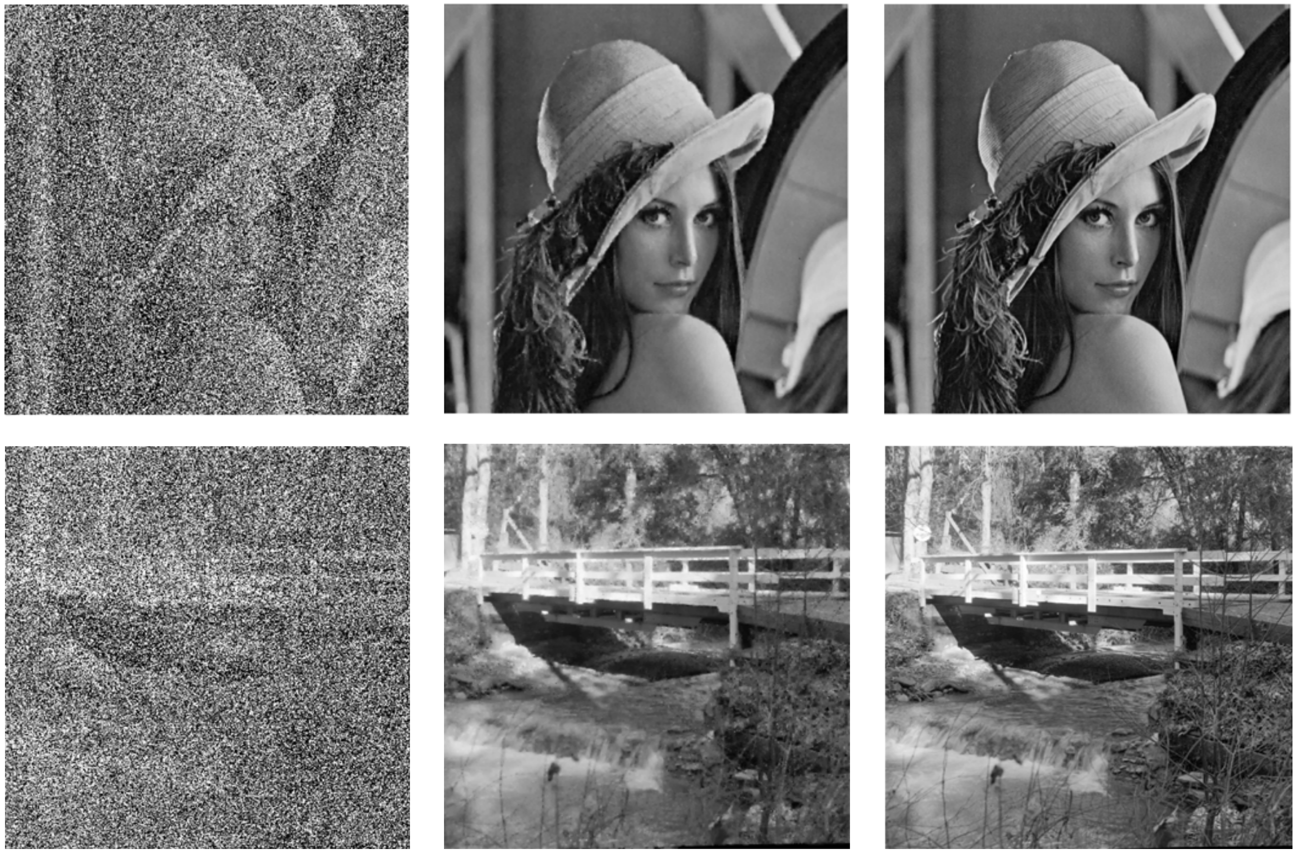

Noisy input Recovered image Original image

What do these examples demonstrate?

R. H. Chan, C.-W. Ho, and M. Nikolova, Salt-and-Pepper Noise Removal by Median-Type Noise Detectors and Detail-Preserving Regularization. IEEE TIP 2005

Non-local Means Denoising

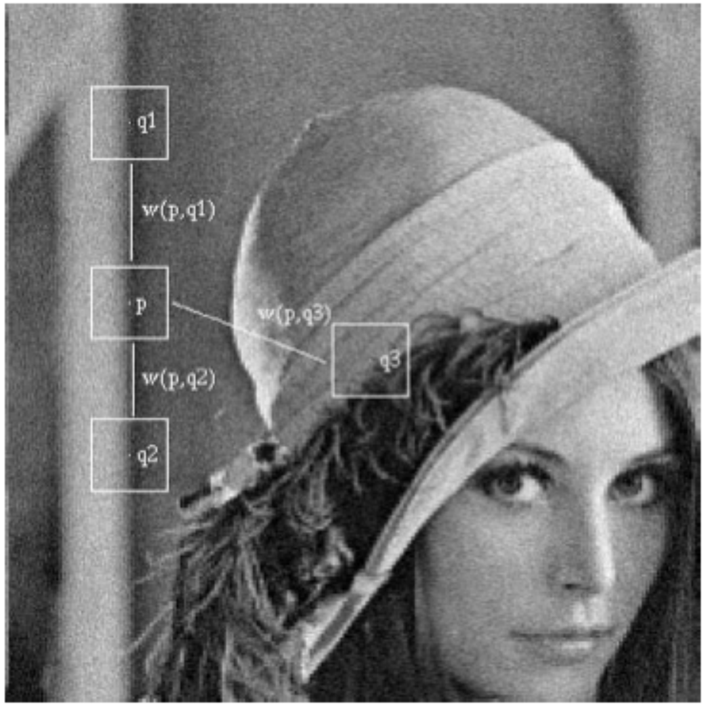

Figure 1. Scheme of NL-means strategy. Similar pixel neighborhoods give a large weight, w(p,q1) and w(p,q2), while much different neighborhoods give a small weight w(p,q3).

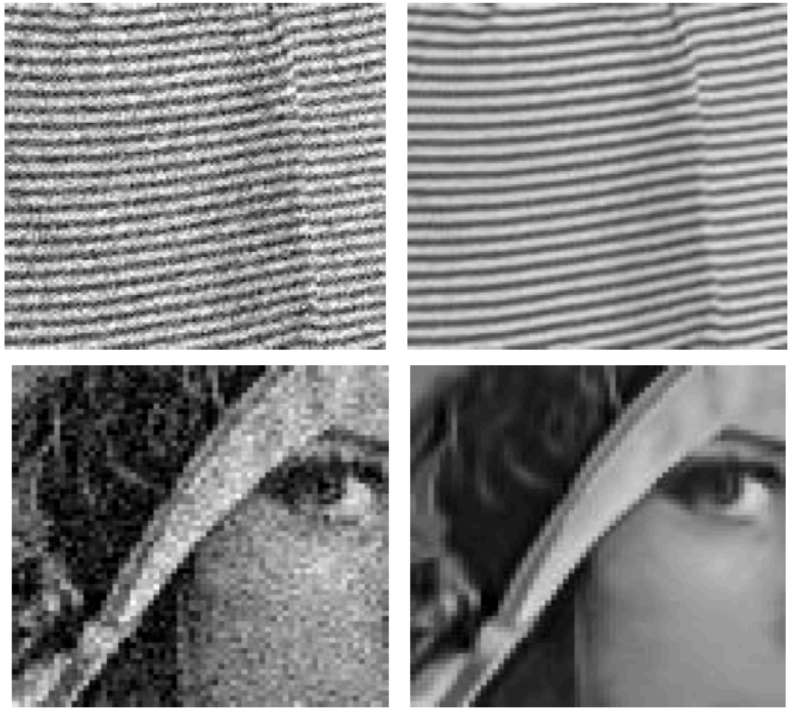

Preserve fine image details and texture during denoising

A. Buades, B. Coll, J. M. Morel, A non-local algorithm for image denoising, CVPR, 2005

Context-Guided Filtering

- Use local image context to steer filtering

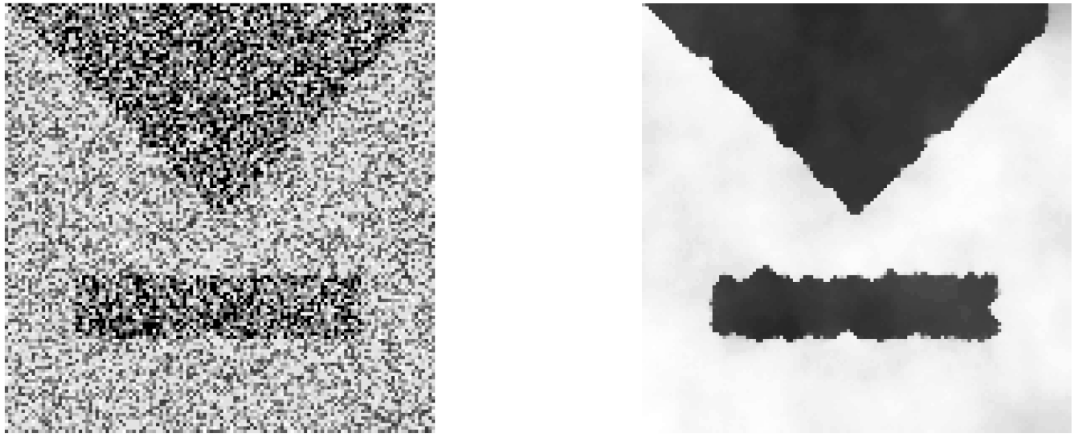

Preserve main image structures during filtering

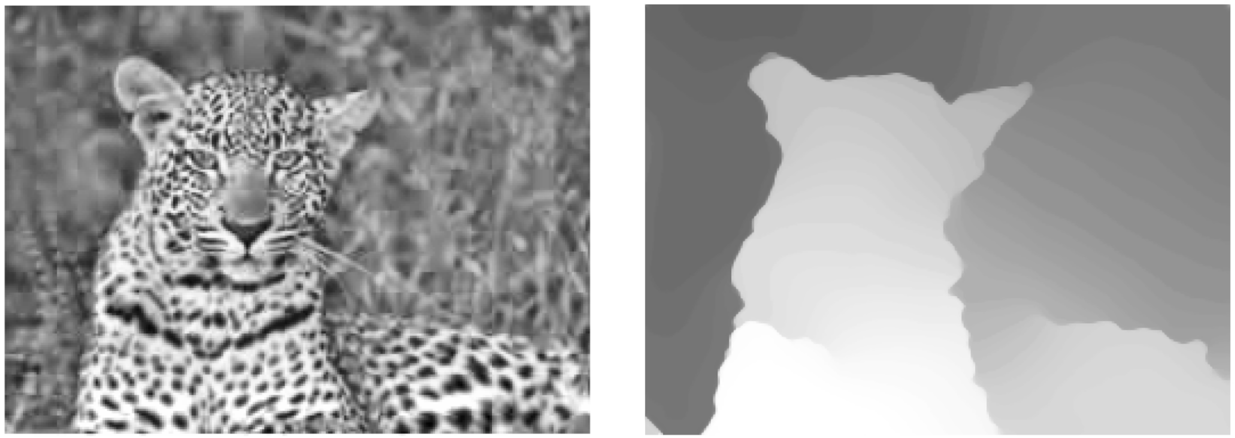

E. Erdem and S. Tari, Mumford-Shah Regularizer with Contextual Feedback, JMIV, 2009

Image Smoothing

L. Xu, C. Lu, Y. Xu, J. Jia, Image Smoothing via L0 Gradient Minimization, ACM Trans. Graphics 2011 (SIGGRAPH Asia 2011)

Image Smoothing

L. Karacan, E. Erdem, A. Erdem, Structure Preserving Image Smoothing via Region Covariances, ACM Trans. Graphics 2013 (SIGGRAPH Asia 2013)

Image Deblurring

- Remove blur and restore a sharp image

from a given blurred image

find its latent sharp image

Slide credit: Lee and Cho

Image Deblurring

- Remove blur and restore a sharp image

Input blurred image

Levin et al. CVPR 2010

Slide credit: Lee and Cho

Visual Saliency

- The problem of predicting where people look at images

The squares shows where the observers looked
in eye tracking experiments

E. Erdem and A. Erdem, Visual saliency estimation by nonlinearly integrating features using region covariances,
Journal of Vision 2013

Visual Saliency

- The problem of predicting where people look at images

E. Erdem and A. Erdem, Visual saliency estimation by nonlinearly integrating features using region covariances, Journal of Vision 2013

Image Retargetting

- automatically resize an image to arbitrary aspect ratios while preserving important image features

How we define the importance?

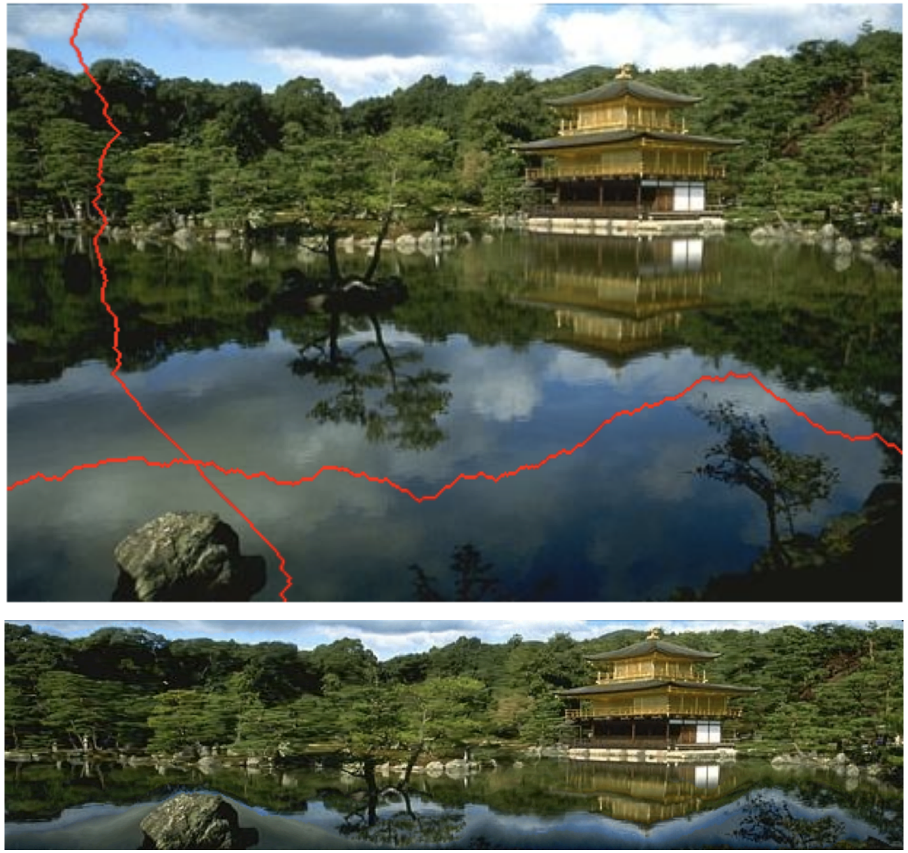

S. Avidan and A. Shamir, Seam Carving for Content-Aware Image Resizing, SIGGRAPH, 2007

Image retargeting by Seam Carving with different importance maps

Sparse Coding

- The problem of finding a small number of representative atoms from a dictionary which when combined with right weights represent a given signal.

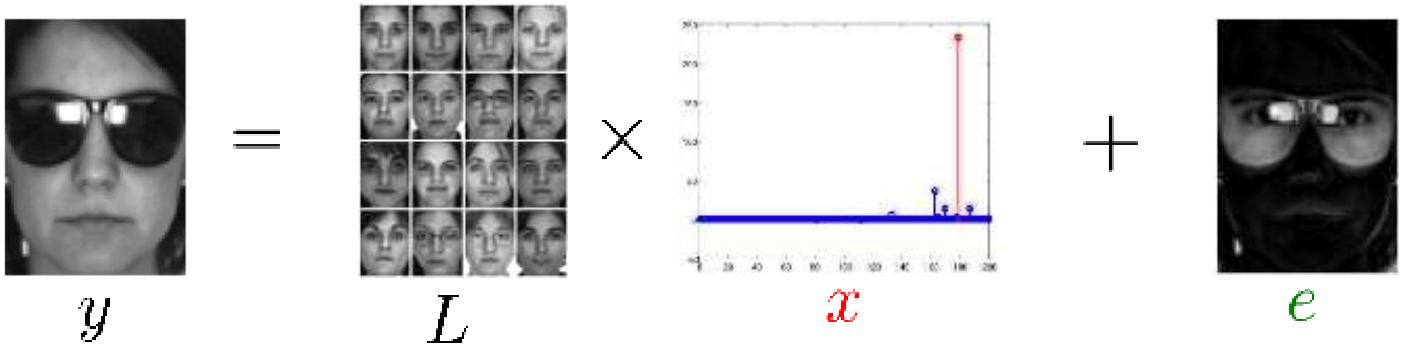

$$y = Lx + e$$

Credit: Yi Ma

Low-Rank Matrix Approximations

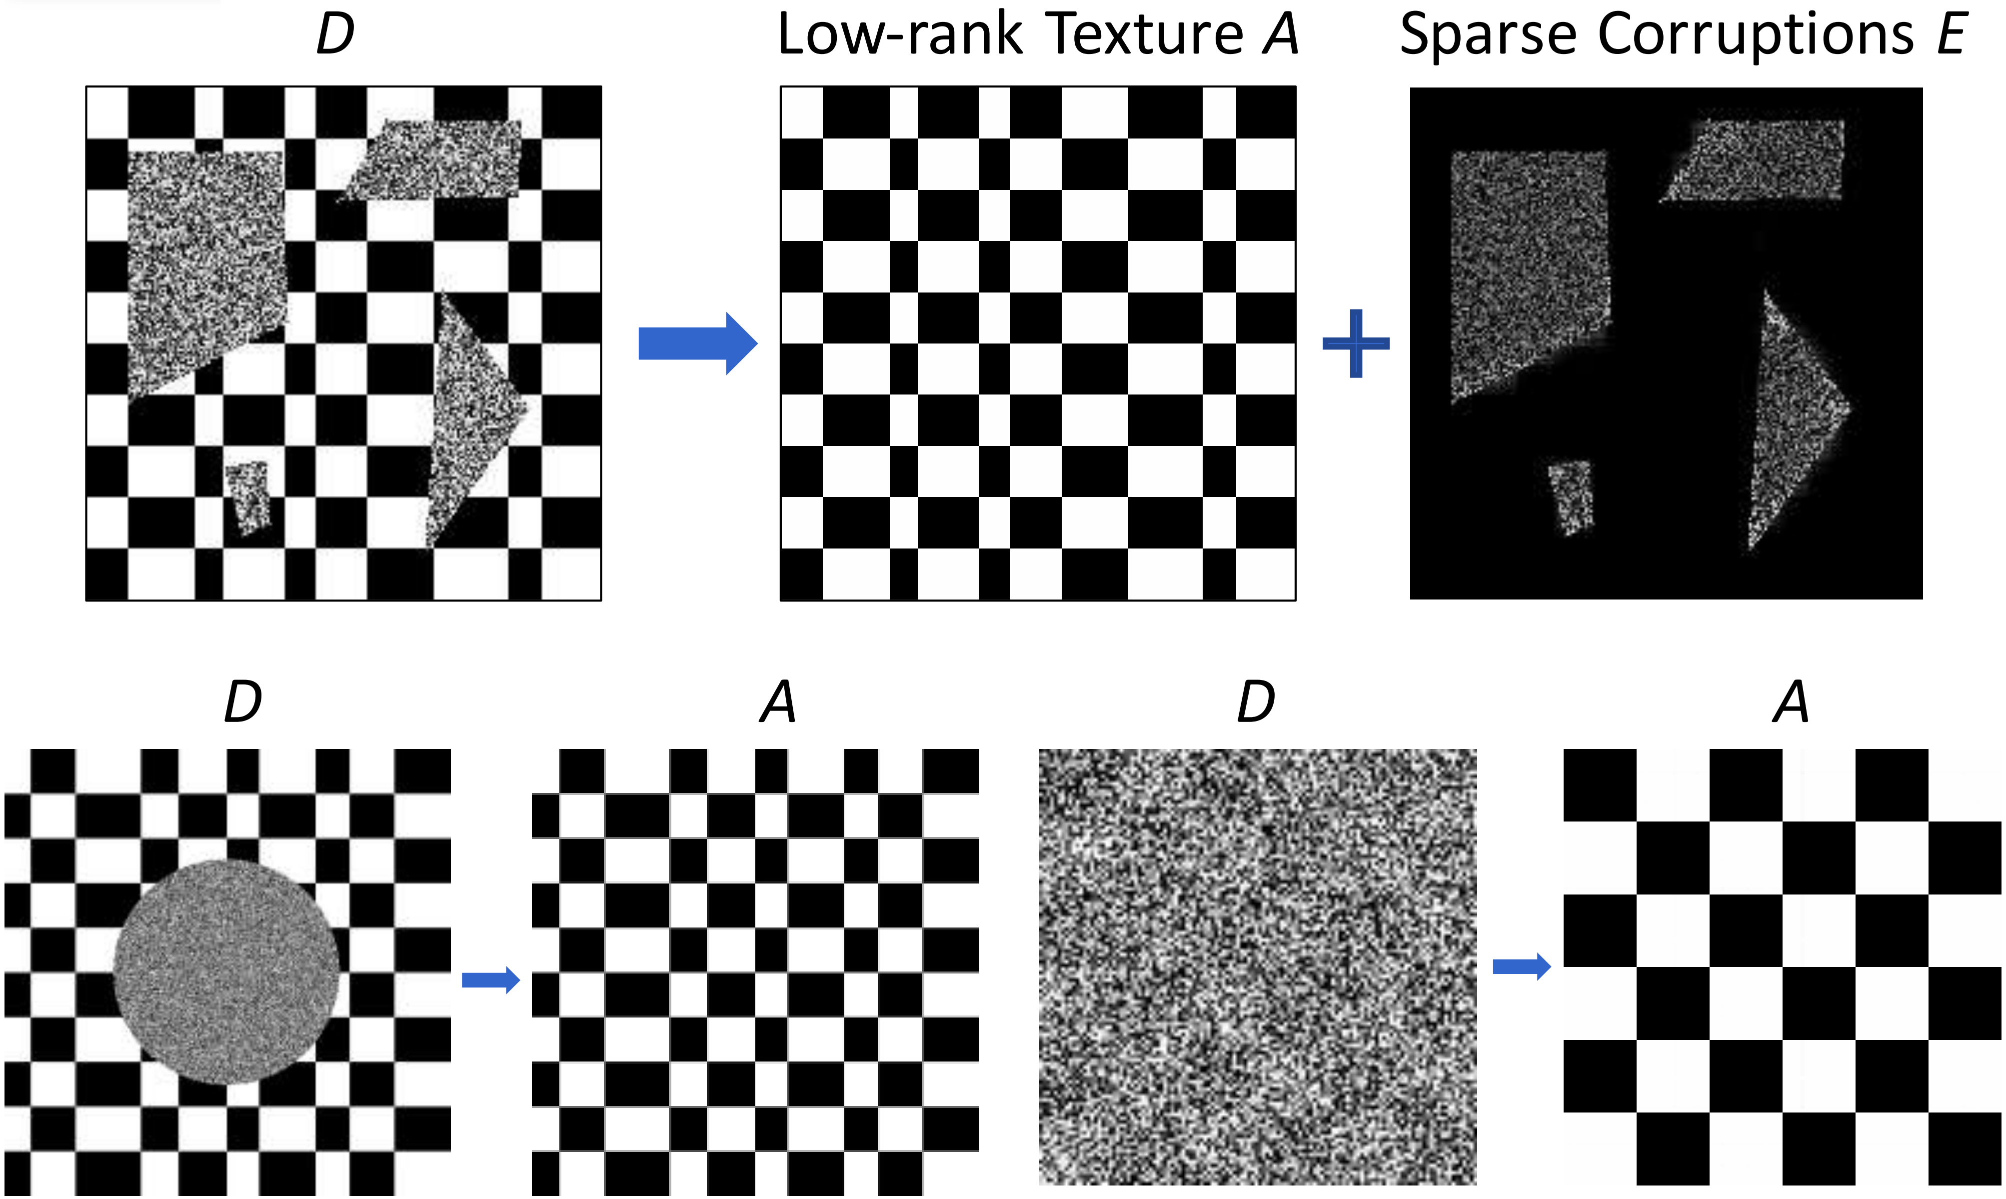

Credit: Yi Ma

Image Inpainting

- Reconstructing lost or deteriorated parts of images

What do these examples demonstrate?

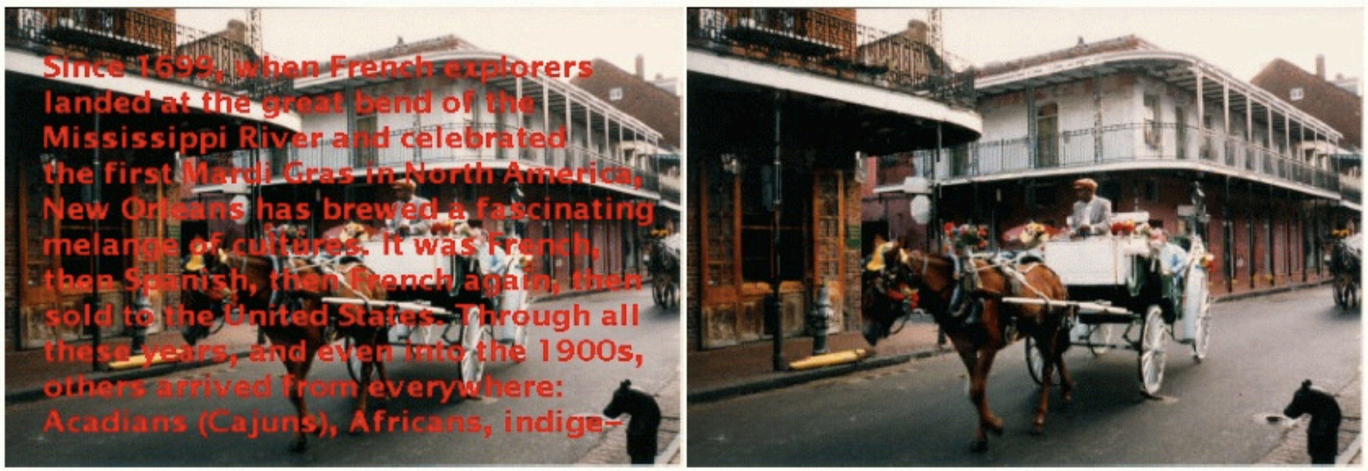

M. Bertalmio, G. Sapiro, V. Caselles and C. Ballester, Image Inpainting, SIGGRAPH, 2000

Image Segmentation

- Partition an image into meaningful regions that are likely to correspond to objects exist in the image

Grouping of pixels

according to what criteria?

high-level object specific knowledge matters!

Figures: A. Erdem

Image Segmentation

- Boundary-based segmentation
- Region-based segmentation
- Unified formulations

Snakes

- Curve Evolution - parametric curve formulation

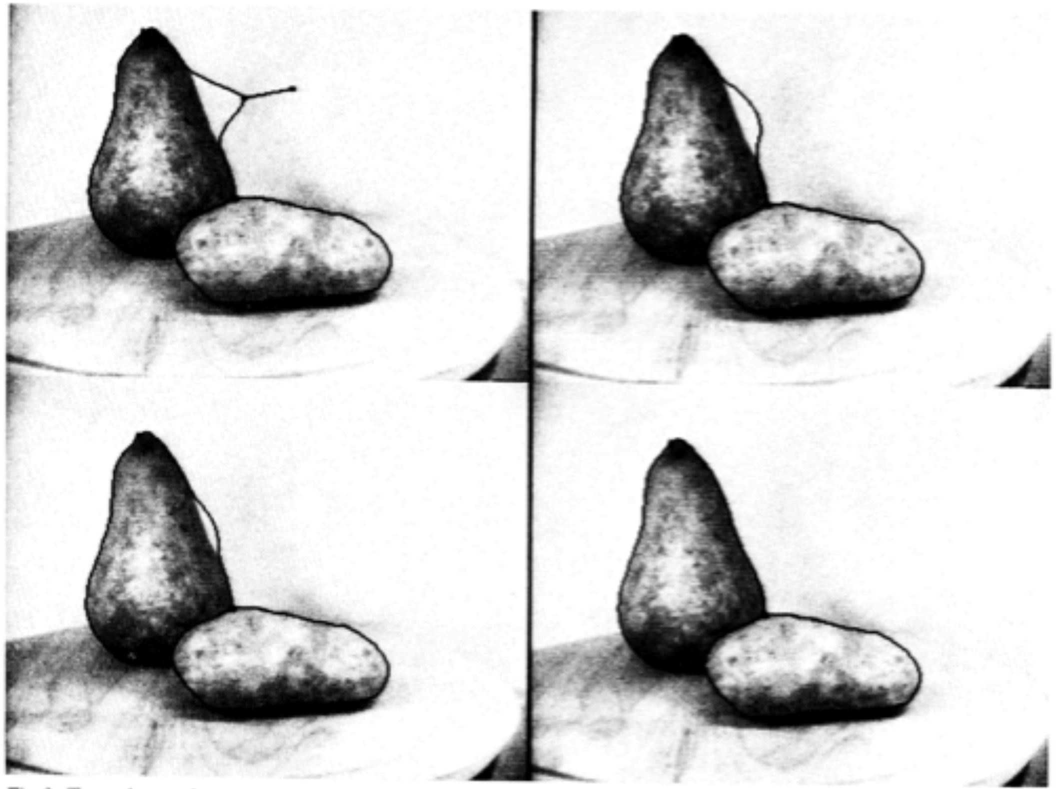

Fig. 3. Two edge snakes on a pear and potato. Upper-left: The user has pulled one of the snakes away from the edge of the pear. Others: After the user lets go, the snake snaps back to the edge of the pear.

M. Kass, A. Witkin, and D. Terzopoulos, Snakes: Active Contour Models, IJCV, 1988

Active Contours Without Edges

- Curve Evolution – a level-set based curve formulation

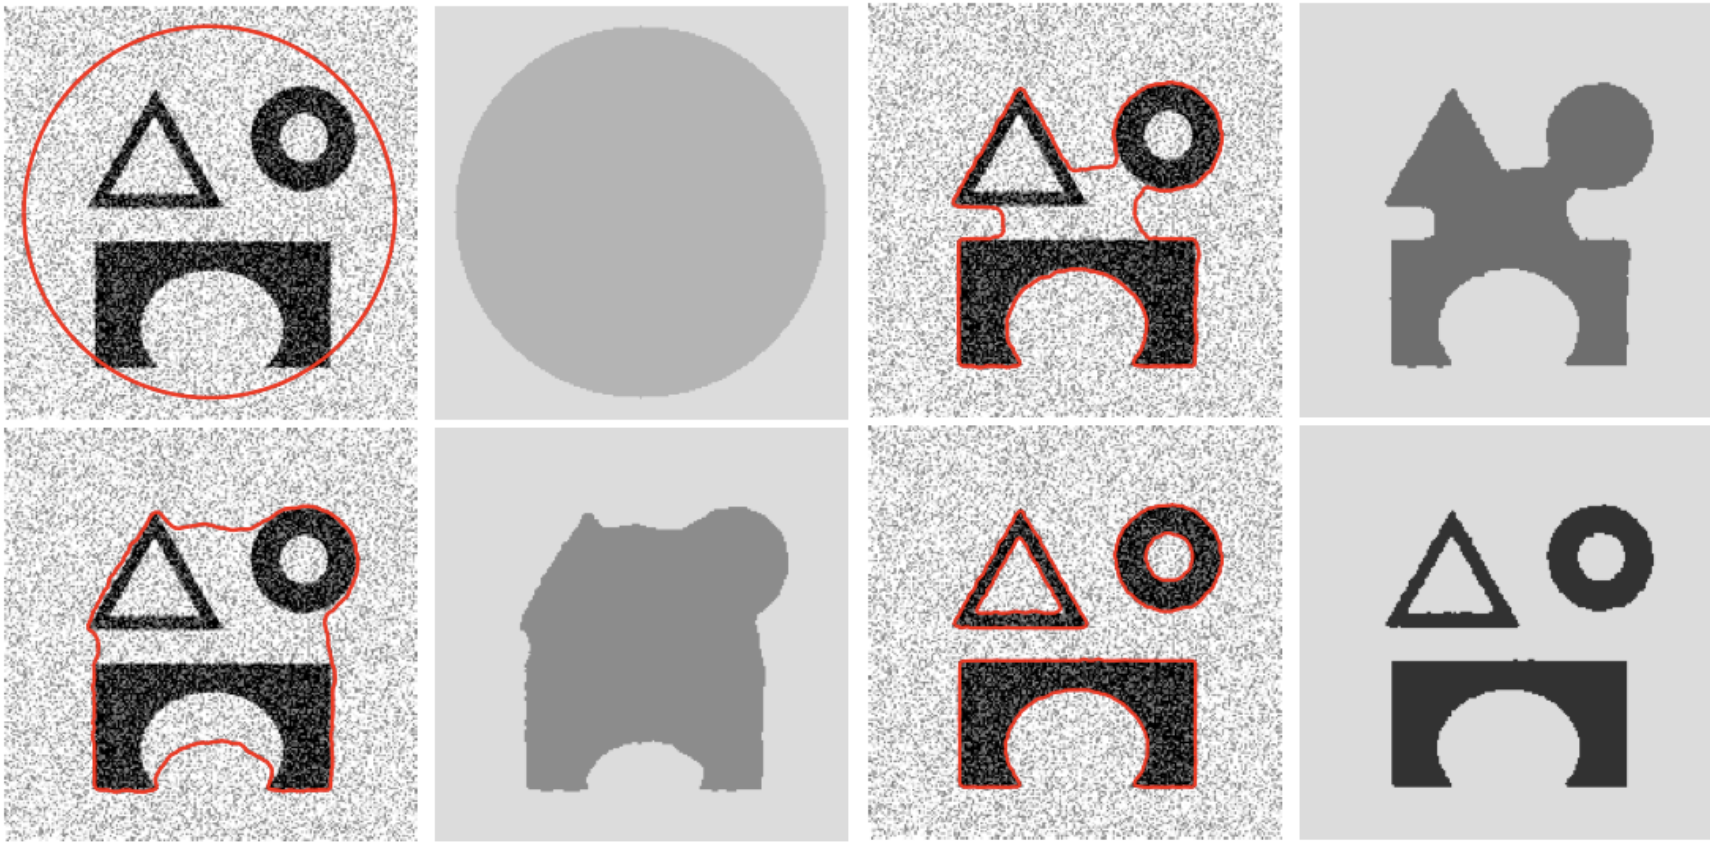

T. Chan and L. Vese. Active Contours Without Edges, IEEE Trans. Image Processing, 2001

Normalized Cuts

- A graph-theoretic formulation for segmentation

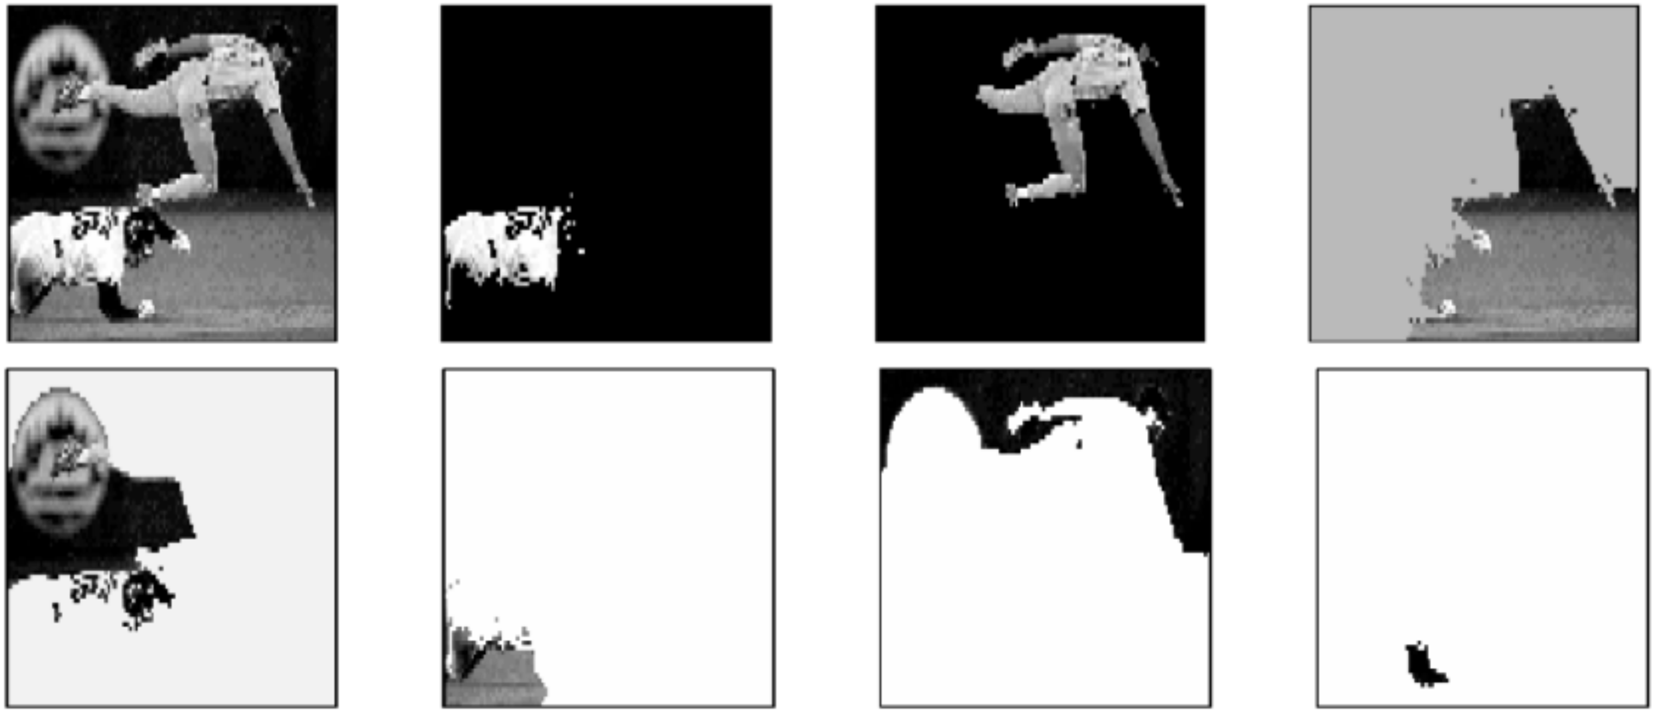

J. Shi and J. Malik, Normalized Cuts and Image Segmentation, IEEE Trans. Pattern Anal. Mach. Intel.

Normalized Cuts

Slide credit: S. Lazebnik

From contours to regions

- <u>State-of-the-art:</u> gPb-owt-ucm segmentation algorithm

P. Arbelaez, M. Maire, C. Fowlkes and J. Malik, Contour Detection and Hierarchical Image Segmentation, IEEE Trans Pattern Anal. Mach. Intell. 33(5):898-916, 2011

From contours to regions

- <u>State-of-the-art:</u> gPb-owt-ucm segmentation algorithm

P. Arbelaez, M. Maire, C. Fowlkes and J. Malik, Contour Detection and Hierarchical Image Segmentation, IEEE Trans Pattern Anal. Mach. Intell. 33(5):898-916, 2011

Prior-Shape Guided Segmentation

- Incorporate prior shape information into the segmentation process

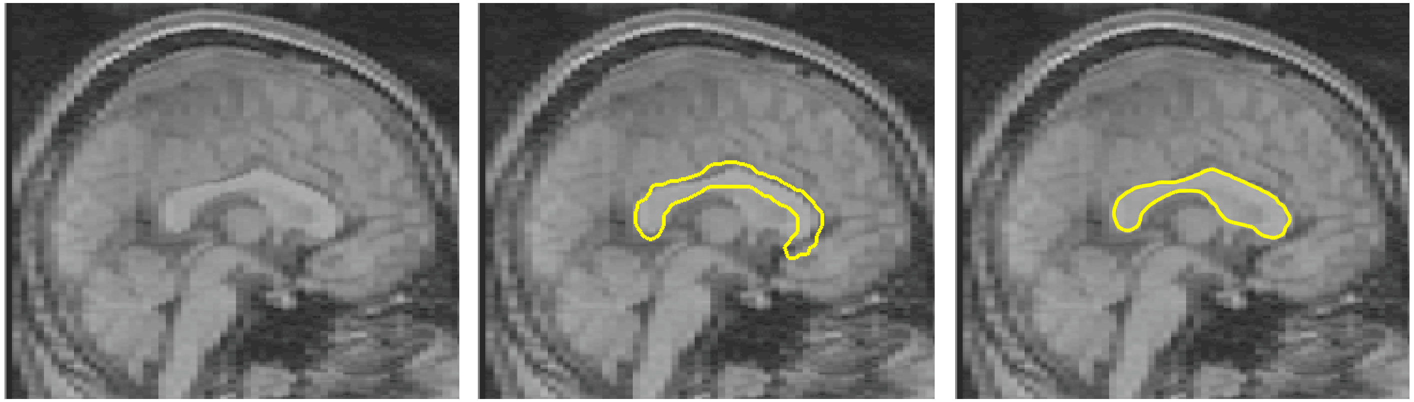

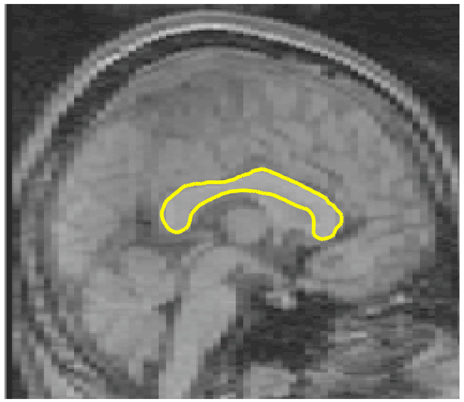

Our result

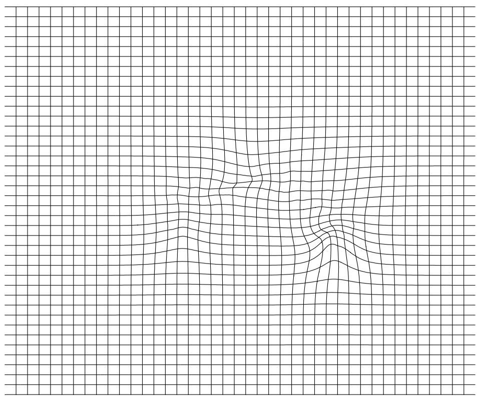

Deformation map

E. Erdem, S. Tari, and L. Vese, Segmentation Using The Edge Strength Function as a Shape Prior within a Local Deformation Model, ICIP 2009

Graphical Models in Vision

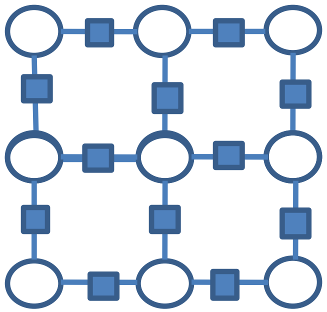

4-connected;
pairwise MRF

$$E(x) = \sum_{i,j \in N_4} \theta_{ij}(x_i,x_j)$$

Order 2

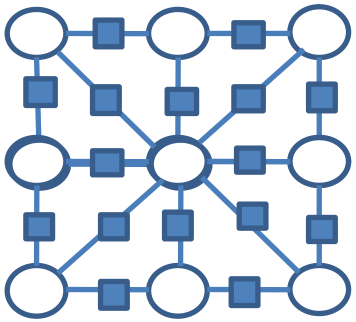

higher(8)-connected;
pairwise MRF

$$E(x) = \sum_{i,j \in N_8} \theta_{ij}(x_i,x_j)$$

Order 2

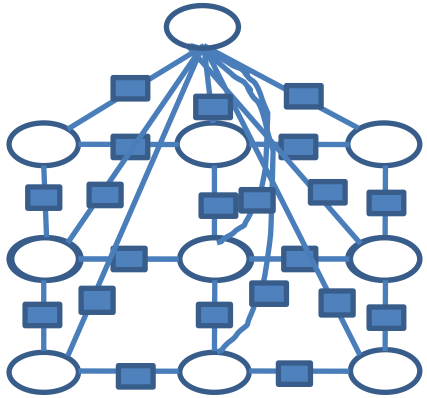

MRF with
global variables

$$E(x) = \sum_{i,j \in N_8} \theta_{ij}(x_i,x_j)$$

Order 2

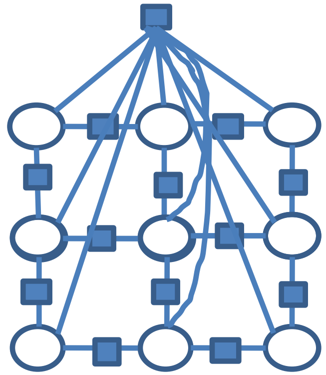

Higher-order MRF

$$E(x) = \sum_{i,j \in N_4} \theta_{ij}(x_i,x_j) + \theta(x_1,\ldots,x_n)$$

Order n

C. Rother

Semantic Segmentation

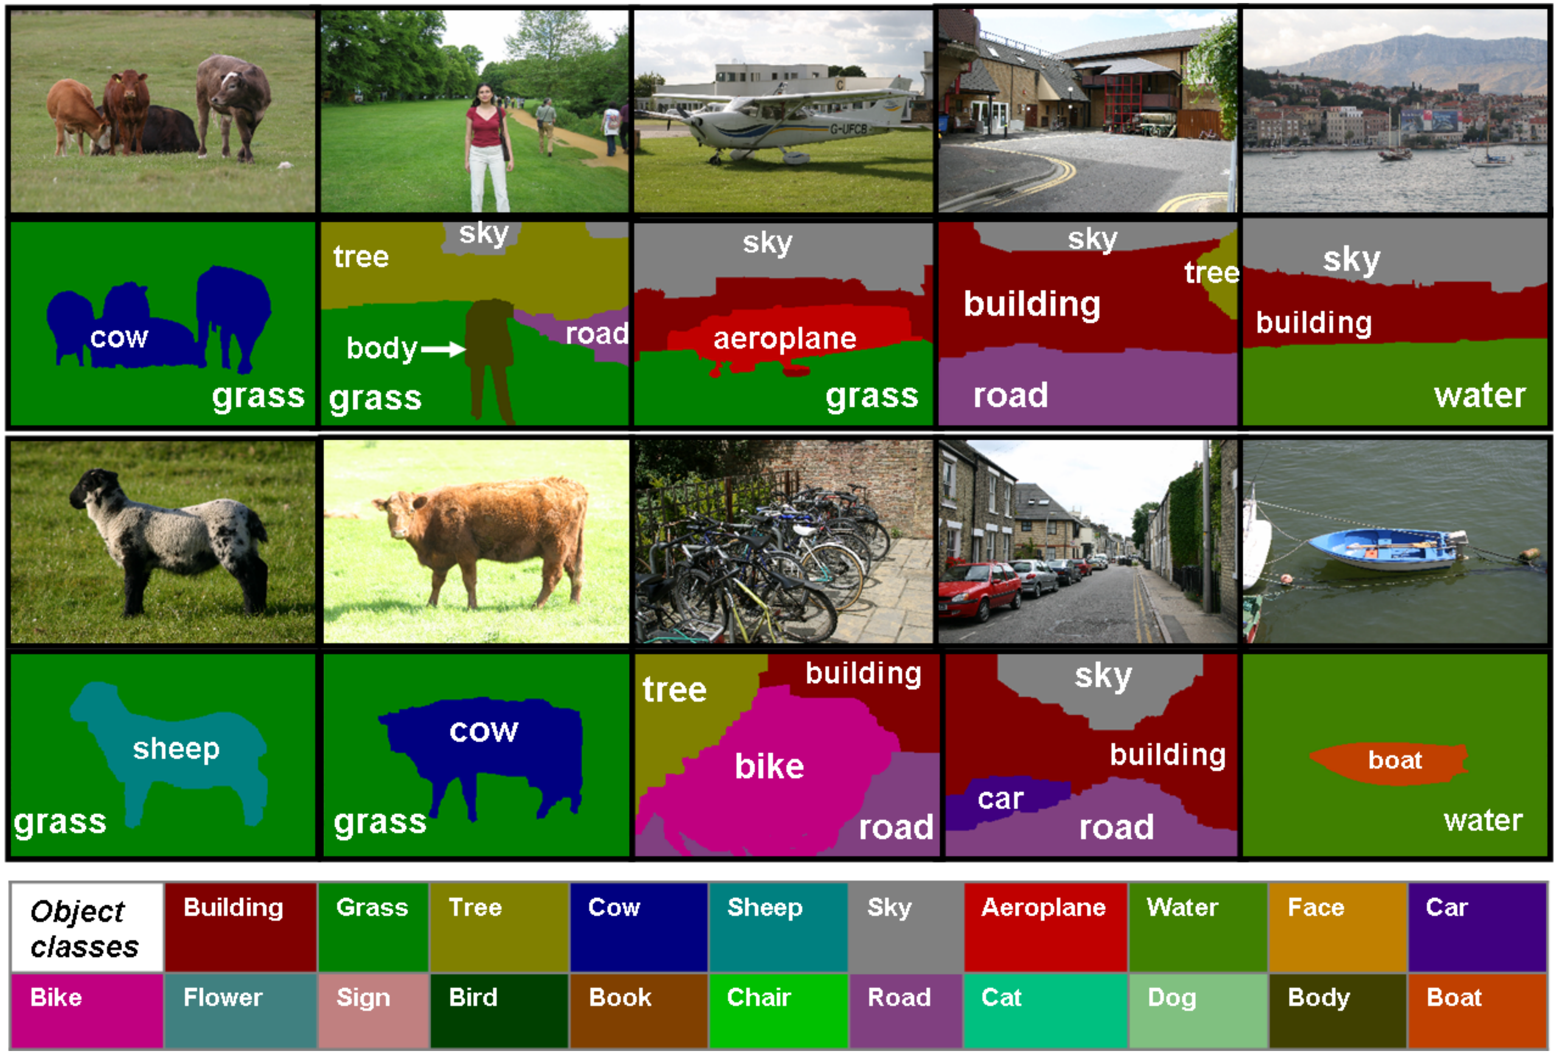

[TextonBoost; Shotton et al, ‘06]

C. Rother

Semantic Segmentation

- The problem of joint recognition and segmentation

Carreira et al., Semantic Segmentation with Second-Order Pooling , ECCV, 2012

Top-down Saliency

- Task-oriented models (e.g. searching for a target object from a specific category)

A. Kocak et al., Top down saliency estimation via superpixel-based discriminative dictionaries, BMVC 2014

Top-down Saliency

- Task-oriented models (e.g. searching for a target object from a specific category)

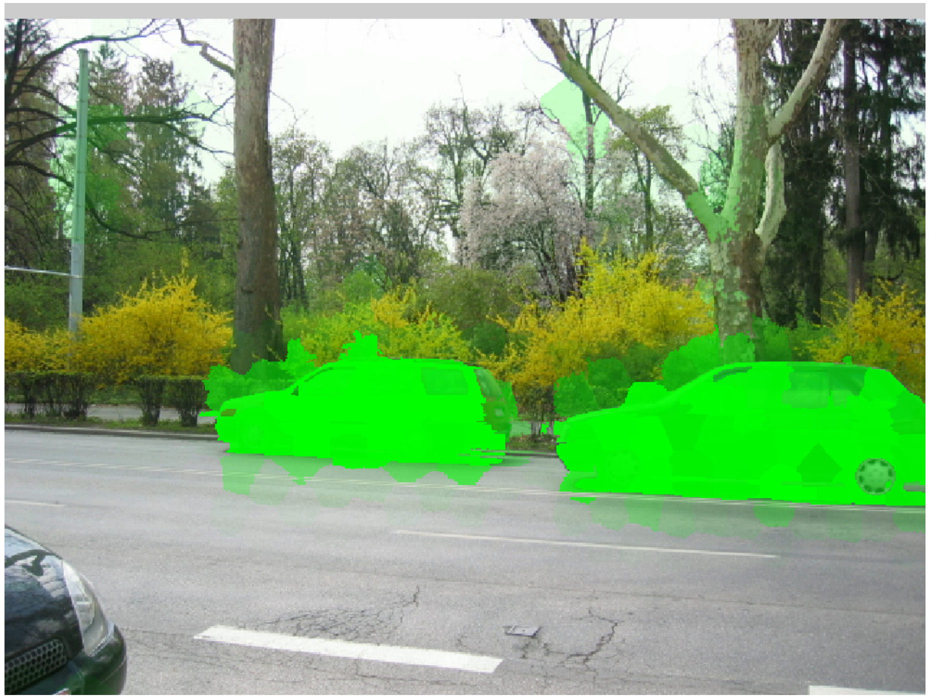

A. Kocak et al., Top down saliency estimation via superpixel-based discriminative dictionaries, BMVC 2014

Top-down Saliency

- Task-oriented models (e.g. searching for a target object from a specific category)

A. Kocak et al., Top down saliency estimation via superpixel-based discriminative dictionaries, BMVC 2014

Deep Learning

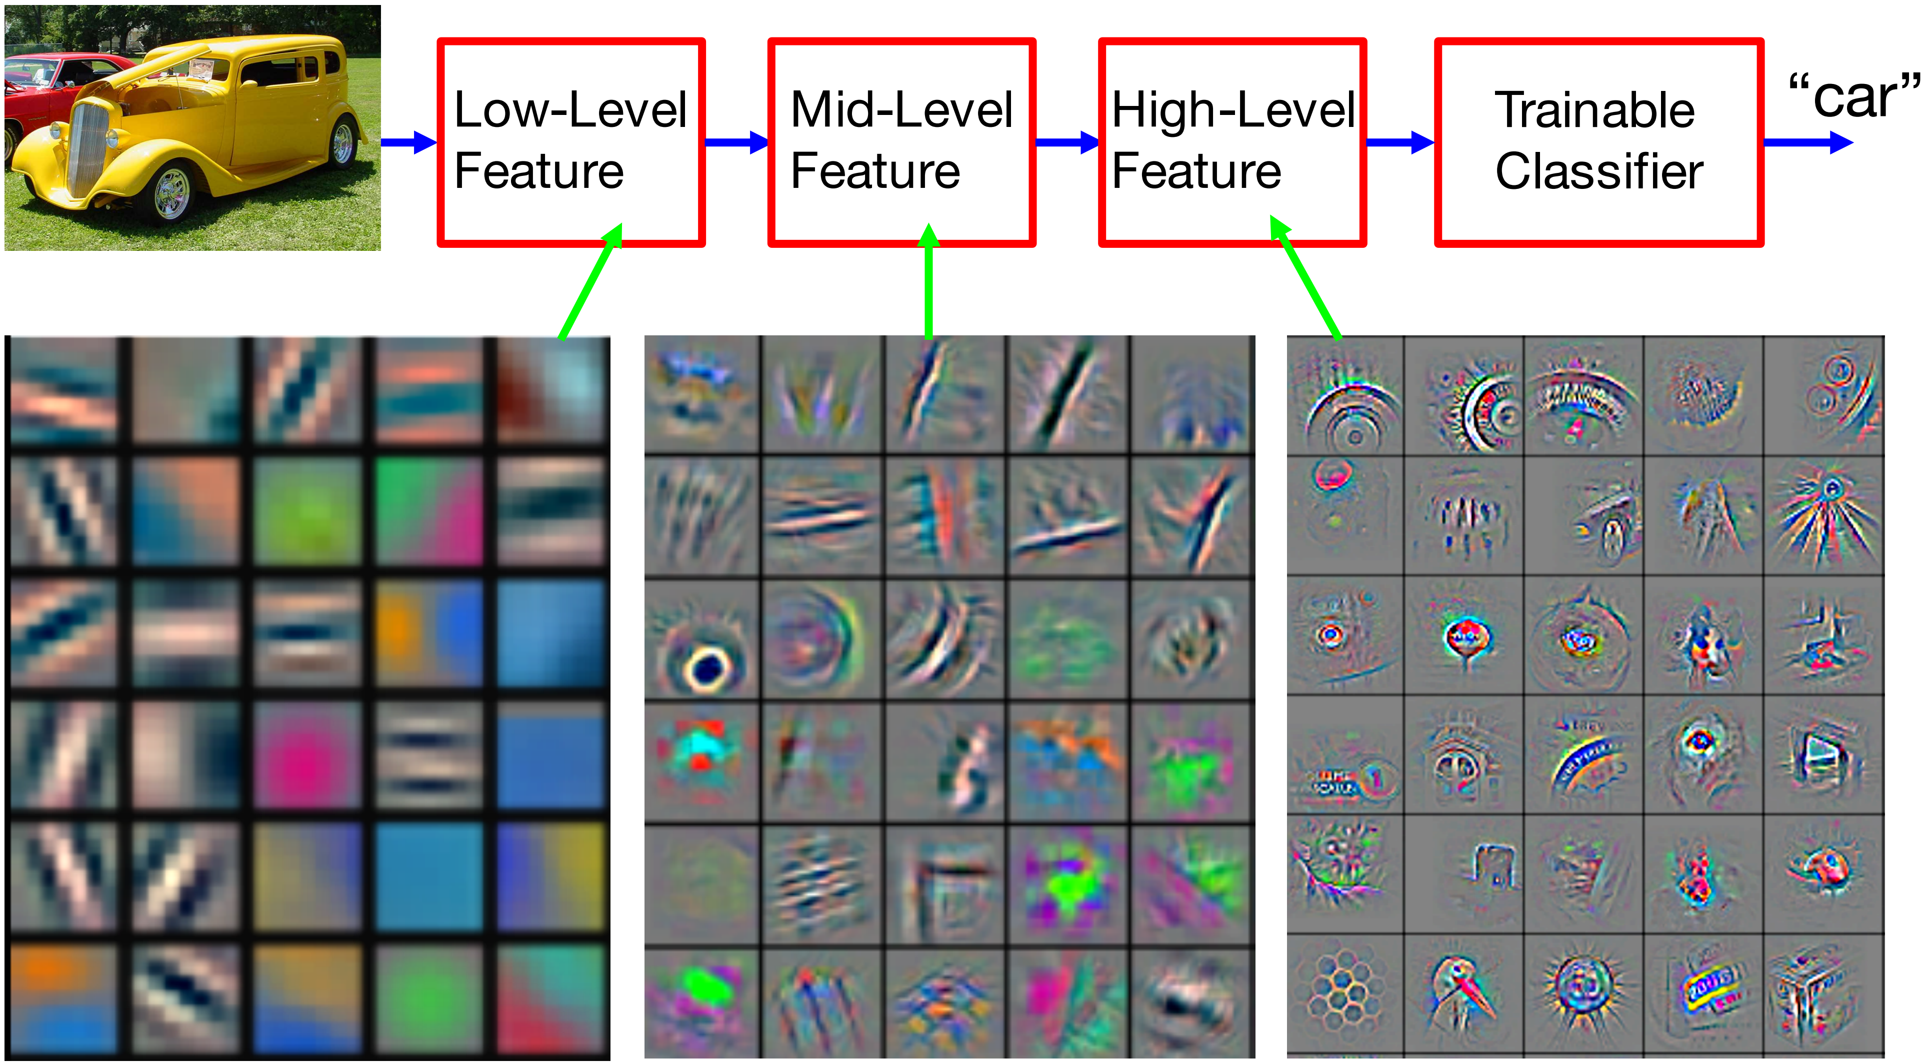

Feature visualization of convolutional net trained on ImageNet from [Zeiler & Fergus 2013]

Slide Credit: Marc'Aurelio Ranzato, Yann LeCun

Deep Learning

- [Krizhevsky et al. NIPS12]
  - 54 million parameters; 8 layers (5 conv, 3 fully-connected)
  - Trained on 1.4M images in ImageNet

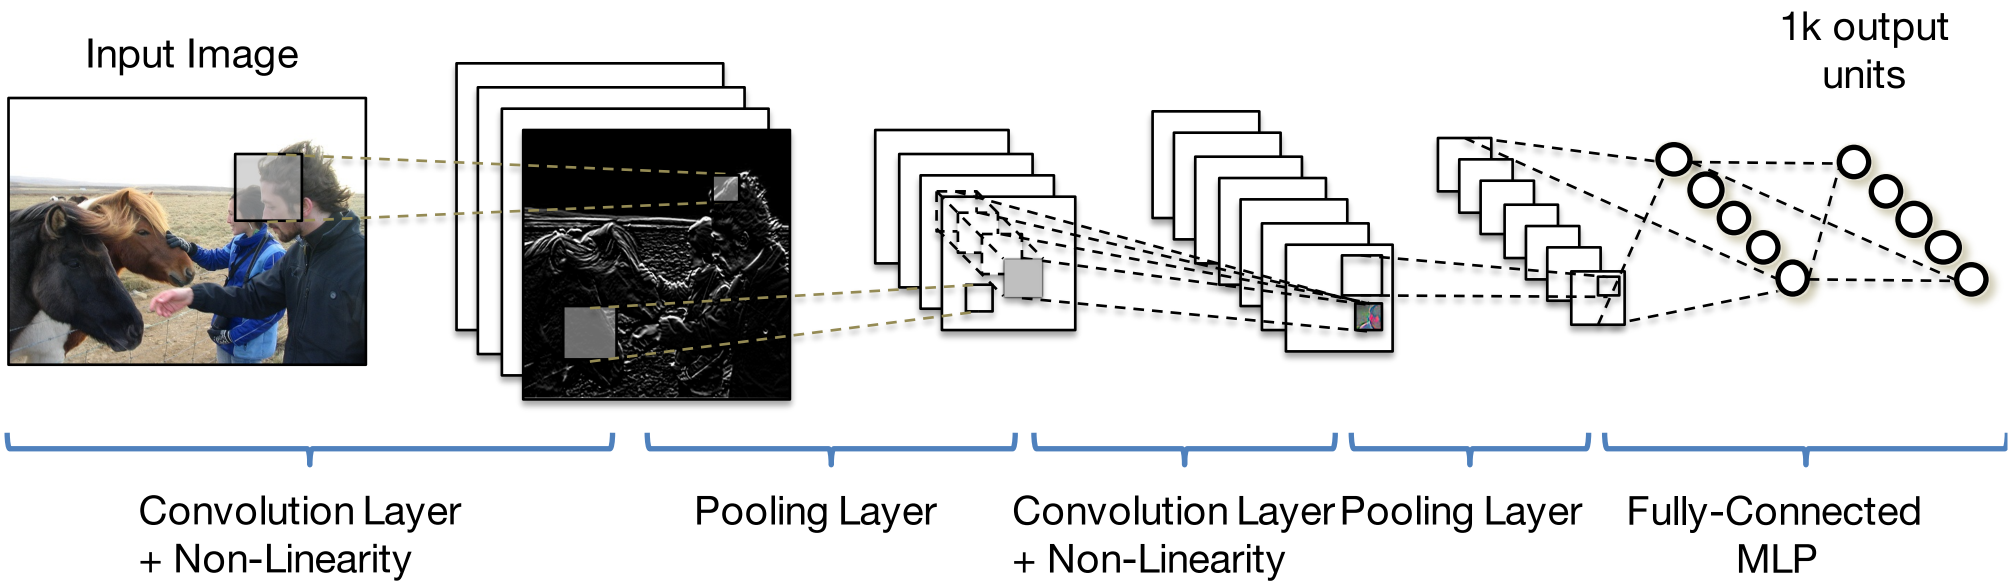

Slide Credit: Dhruv Batra

Deep Learning

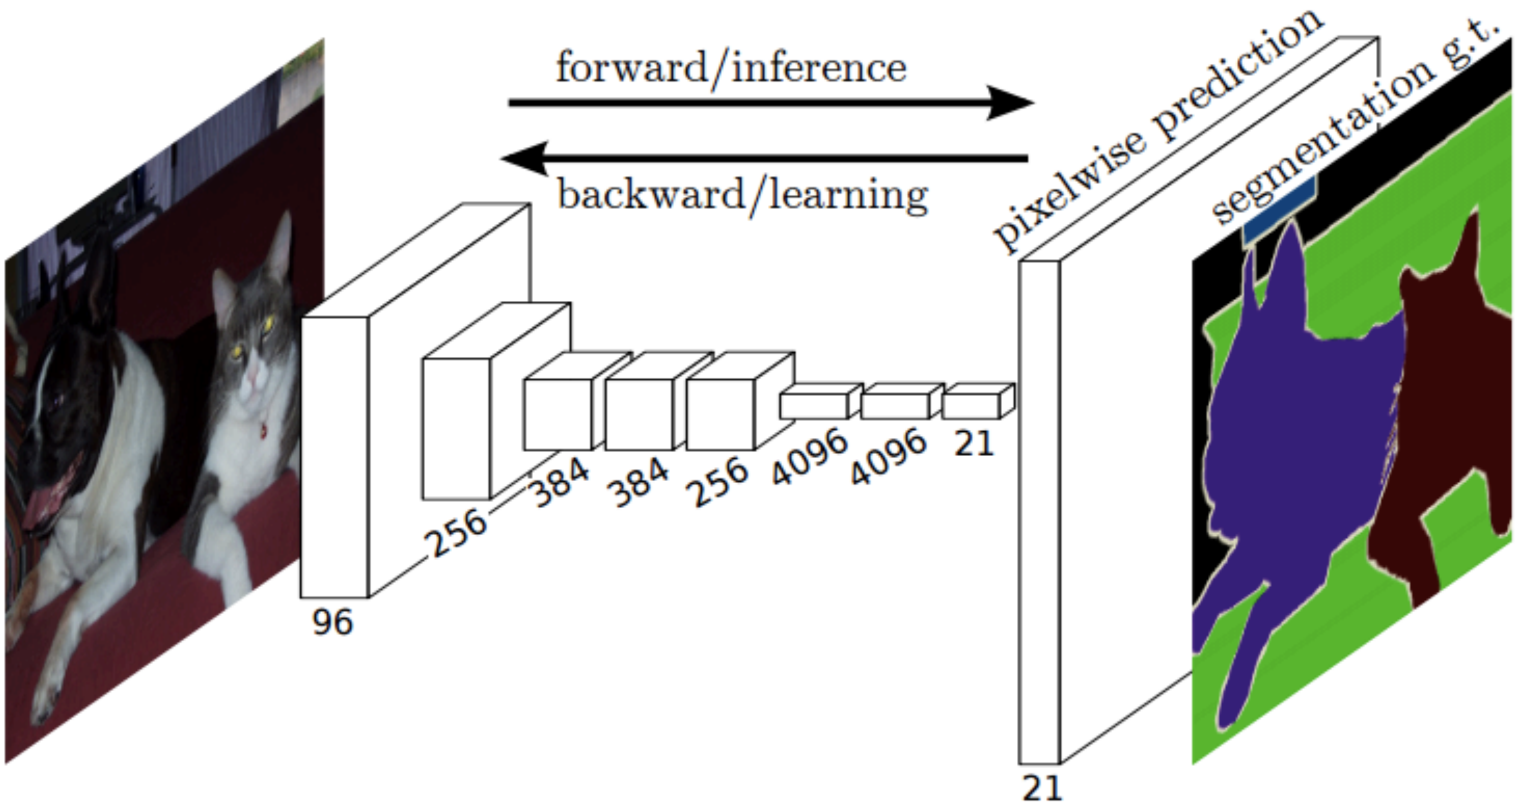

Fully Convolutional Networks for Semantic Segmentation[Long, Shelmer & Darrell 2015]

Today

- About me
- About you
- Introduction to Image Processing
- Course outline and logistics

Logistics

- Asst. Prof. Erkut ERDEM
- erkut@cs.hacettepe.edu.tr
- Office: 114

- Lectures: Monday, 13:30-16:30
- Office Hour: *By appointment.*

About BIL717

- This course provides a comprehensive overview of fundamental topics in image processing for graduate students.

- The goal of this course is to provide a deeper understanding of the state-of-the-art methods in image processing literature and to study their connections.

- The course makes the students gain knowledge and skills in key topics and provides them the ability to employ them in their advanced-level studies.

Communication

- The course webpage will be updated regularly throughout the semester with lecture notes, programming and reading assignments and
  important deadlines.
  http://web.cs.hacettepe.edu.tr/~erkut/bil717.s16

- All other communications will be carried out through Piazza. Please enroll it by following the link
  https://piazza.com/hacettepe.edu.tr/spring2016/bil717

Prerequisites

- Programming skills
  (C/C++, Matlab)
- Good math background
  (Calculus, Linear Algebra, Statistical Methods)
- A prior, introductory-level course in image processing is highly recommended.

Reference Books

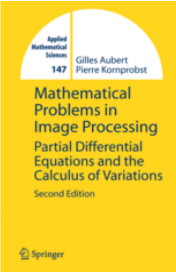

- Mathematical Problems in Image Processing: Partial Differential Equations and the Calculus of Variations, G. Aubert and P. Kornprobst,
2nd Edition, Springer-Verlag, 2006

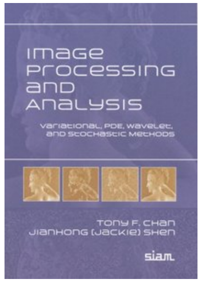

- Image Processing And Analysis: Variational, PDE, Wavelet, And Stochastic Methods,
T. Chan and J. Shen, Society for Industrial and Applied Mathematics, 2005

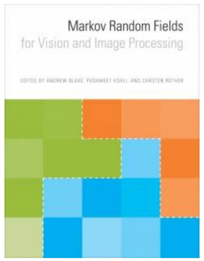

- Markov Random Fields For Vision And Image Processing, Edited by A. Blake, P. Kohli and
C. Rother, MIT Press, 2011

- Deep Learning,
Ian Goodfellow, Aaron Courville, and Yoshua Bengio, book in preparation for MIT Press

Reading Material

- Lecture notes and handouts
- Papers and journal articles

Related Conferences

- IEEE International Conference on Computer Vision (ICCV)
- European Conference on Computer Vision (ECCV)
- IEEE Conference on Computer Vision and Pattern Recognition (CVPR)
- IEEE Winter Conference on Applications of Computer Vision (WACV)
- British Machine Vision Conference (BMVC)
- ACM SIGGRAPH
- ACM SIGGRAPH Asia
- Advances in Neural Information Processing Systems (NIPS)
- IEEE International Conference on Pattern Recognition (ICPR)
- IEEE International Conference on Image Processing (ICIP)

Related Journals

- IEEE Transactions on Image Processing (IEEE TIP)
- IEEE Transactions on Pattern Analysis and Machine Intelligence (IEEE TPAMI)
- ACM Transactions on Graphics (TOG)
- International Journal of Computer Vision (IJCV)
- Computer Vision and Image Understanding (CVIU)
- Image and Vision Computing (IMAVIS)
- Pattern Recognition (PR)

Grading Policy

- 20% Quizzes
- 20% Programming Assignments
- 20% Paper presentations/Class participation
- 40% Project and final term paper

Paper presentations and Quizzes

- The students will be required to present at least one research paper either of their choice or from the suggested reading list.
- These papers should be read by every student as the quizzes about the presented papers will be given on the weeks of the presentations.
- The schedule for the presentations will be finalized on 15th of February.

Programming Assignments

- There will be three assignments related to the topics covered in the class.
- Each assignment will involve implementing an algorithm, carrying out a set of experiments to evaluate it, and writing up a report on the experimental results.
- All assignments have to be done individually, unless stated otherwise.

Project

- The aim of the project is to give the students some experience on conducting research.
- Students should work individually or groups in two.
- This project may involve
  - design of a novel approach and its experimental analysis,
  - an extension to a recent study (published after 2010) of non-trivial complexity and its experimental analysis,
  - an in-depth empirical evaluation and analysis of two or more related methods not covered in the class.

Project – Important Dates

- Project proposals: 29th of February
- Project progress reports: 4th of April
- Project progress presentations: 11th of April
- Project presentations: *will be announced!*
- Project final reports: 23rd of May

- Late submissions will be penalized!

Tentative Outline

- (1 week) Overview of Image Processing
- (1 week) Edge Detection, Linear Filtering
- (1 week) Image Segmentation, Boundary Detection
- (1 week) Nonlinear Filtering
- (1 week) Snakes, Variational Segmentation Models
- (2 weeks) Modern Image Filtering
- (1 week) Image Deblurring

Tentative Outline

- (1 week) Sparse Coding
- (1 week) Graphical Models
- (1 week) Semantic Segmentation
- (1 week) Visual Saliency
- (1 week) Deep Learning