

BSB 663
Image Processing
Feb. 26, 2013

Erkut Erdem
Hacettepe University

Introduction

Instructor and Course Schedule

- Dr. Erkut ERDEM
- erkut@cs.hacettepe.edu.tr
- Office: 114
- Tel: 297 7500 / 149

- Lectures (BSB663): Tuesday, 18:15-21:00@D7

About BSB663

- A course about the fundamentals of image processing.
- The aim of this course is to provide an introduction to students who wish to specialize in interrelated disciplines like image processing, computer vision.
- The students will also be expected to gain hand-on experience via a set of programming assignments.

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/bsb663.s13>

- All other communications will be carried out through Piazza. Please enroll it by following the link <https://piazza.com/hacettepe.edu.tr/spring2013/bsb663>


Prerequisites

- Programming skills
(Matlab)
- Good math background
(Calculus, Linear Algebra, Statistical Methods)
- Students are not expected to have any prior knowledge of image processing techniques.


Reading Material

- Lecture notes and handouts
- Papers and journal articles

Textbooks



- Computer Vision: Algorithms and Applications, Richard Szeliski, Springer, 2010



- Digital Image Processing, R. C. Gonzalez, R. E. Woods, 3rd Edition, Prentice Hall, 2008

Grading Policy

- BSB663 Image Processing
 - Class participation, 5%
 - Reading Assignments, 5%
 - Programming Assignments, 20%
 - Midterm Exam, 30%
 - Final Exam, 40%

Programming Assignments

- 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.

Important Dates


- Programming Assignment 1 26 March 2013
- Programming Assignment 2 4 April 2013
- Programming Assignment 3 14 May 2013

- Midterm exam 4 April 2013
- Final exam To be announced later

Tentative Outline

- (1 week) Image formation and the digital camera
- (1 week) Color perception and color spaces
- (1 week) Point operations
- (1 week) Spatial filtering
- (1 week) Fourier Transform
- (1 week) Image pyramids and wavelets
- (1 week) Gradients, edges, contours
- (1 week) Image smoothing
- (2 weeks) Image segmentation
- (1 week) Advanced topics

Image Processing



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?

We see a two dimensional image


But, we perceive depth information

light reflected
on the retina



converging lines



shadows of the eye



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



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.

MASSACHUSETTS INSTITUTE OF TECHNOLOGY
PROJECT MAC

Artificial Intelligence Group
Vision Memo. No. 100.


July 7, 1966

THE SUMMER VISION PROJECT


Seymour Papert

The summer vision project is an attempt to use our summer workers effectively in the construction of a significant part of a visual system. The particular task was chosen partly because it can be segmented into sub-problems which will allow individuals to work independently and yet participate in the construction of a system complex enough to be a real landmark in the development of "pattern recognition".


Origins of computer vision




(a) Original picture.



(b) Differentiated picture.



(c) Line drawing.



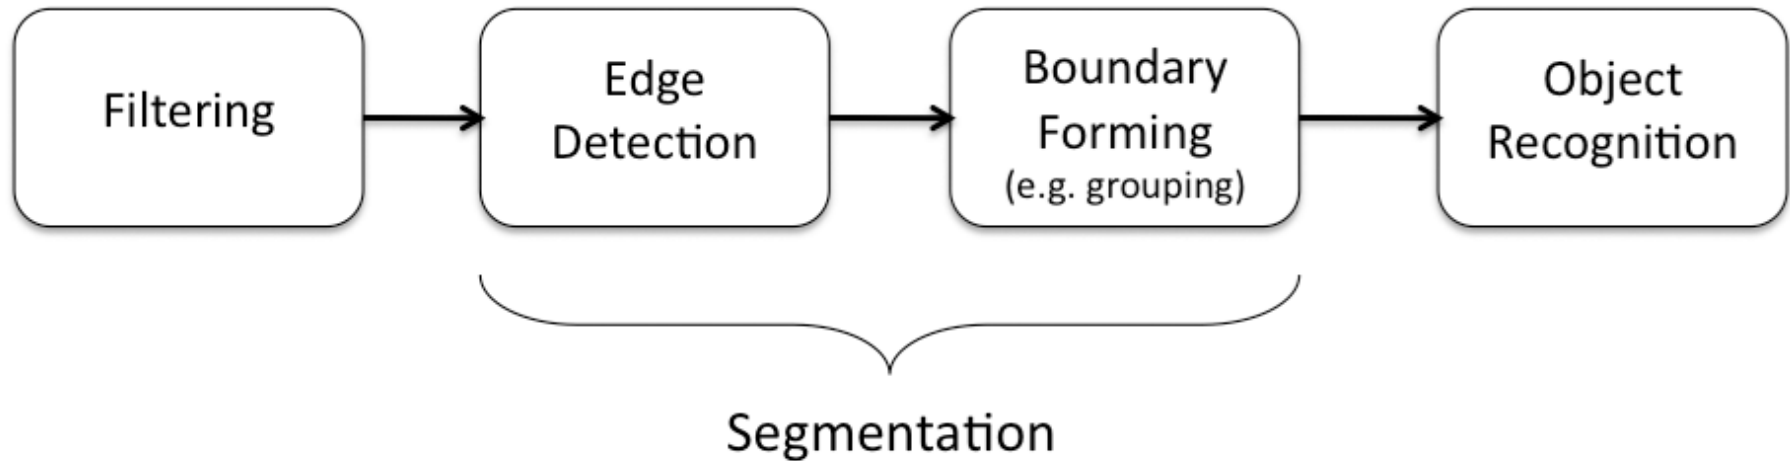
(d) Rotated view.

L. G. Roberts, *Machine Perception of Three Dimensional Solids*, Ph.D. thesis, MIT Department of Electrical Engineering, 1963.

Marr's observation: Studying vision at 3 levels

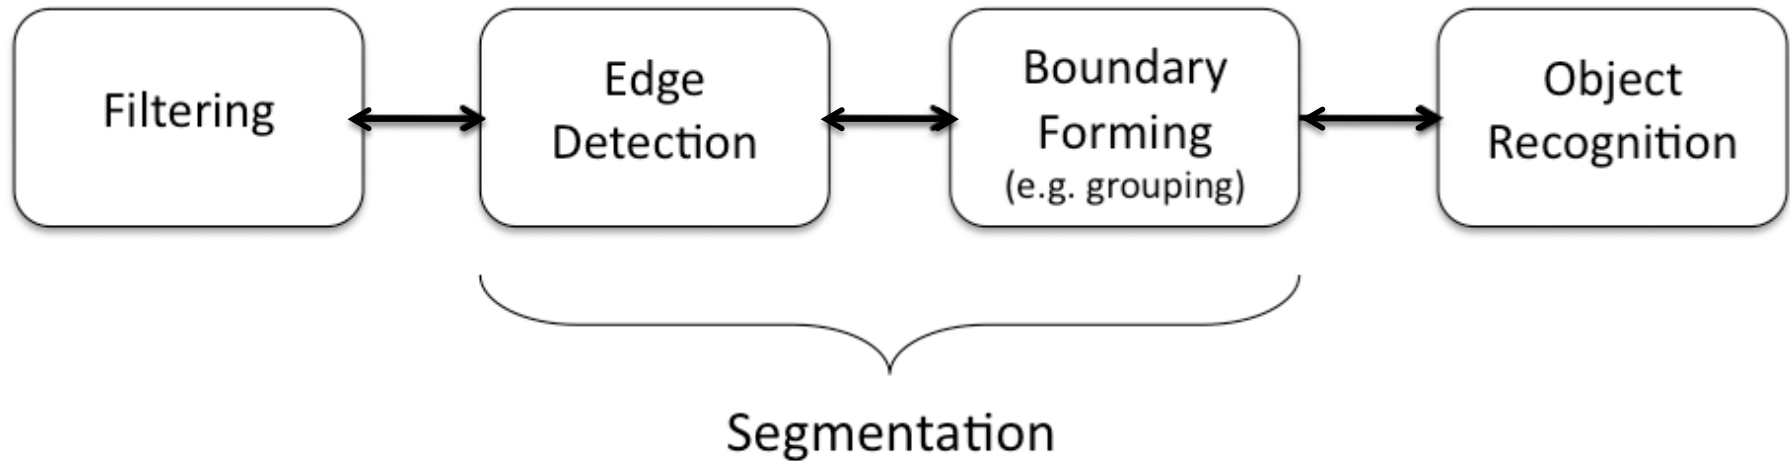
- Vision as an information processing task [David Marr, 1982]
- Three levels of understanding:
 1. Computational theory
 - What is computed? Why it is computed?
 2. Representation and Algorithm
 - How it is computed?
 - Input, Output, Transformation
 3. Physical Realization
 - Hardware

Visual Modules and the Information Flow



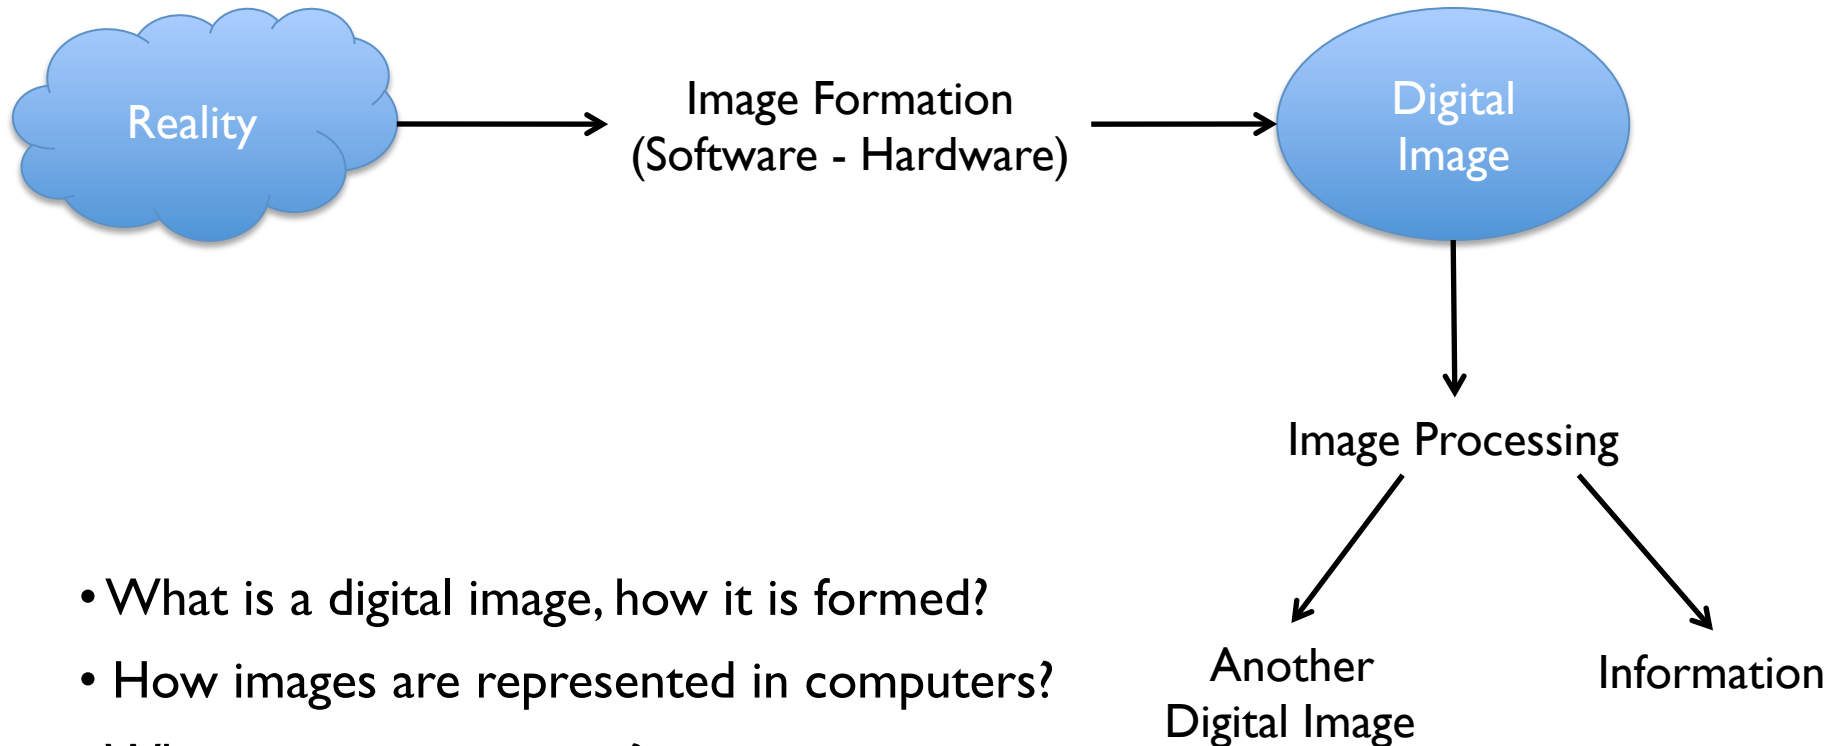
- Visual perception as a data-driven, bottom-up process (traditional view since D. Marr)
- Unidirectional information flow
- Simple low-level cues >> Complex abstract perceptual units

Visual Modules and the Information Flow



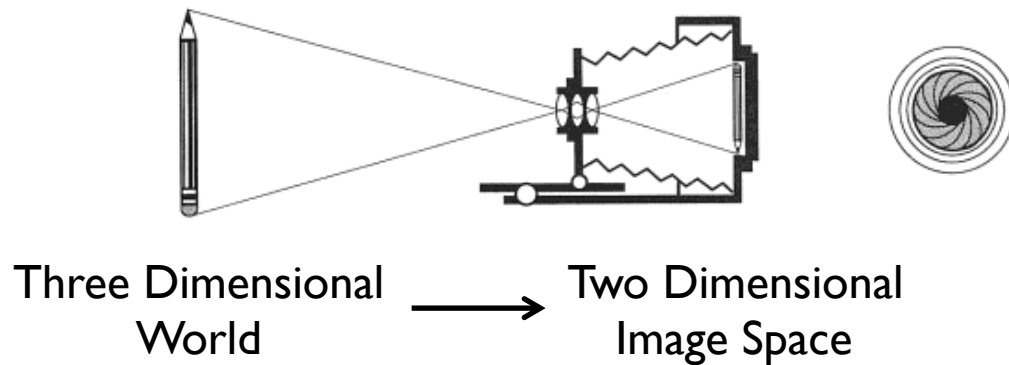
- Vision modules can be categorized into three groups according to their functionality:
 - Low-level vision: filtering out irrelevant image data
 - Mid-level vision: grouping pixels or boundary fragments together
 - High-level vision: complex cognitive processes

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

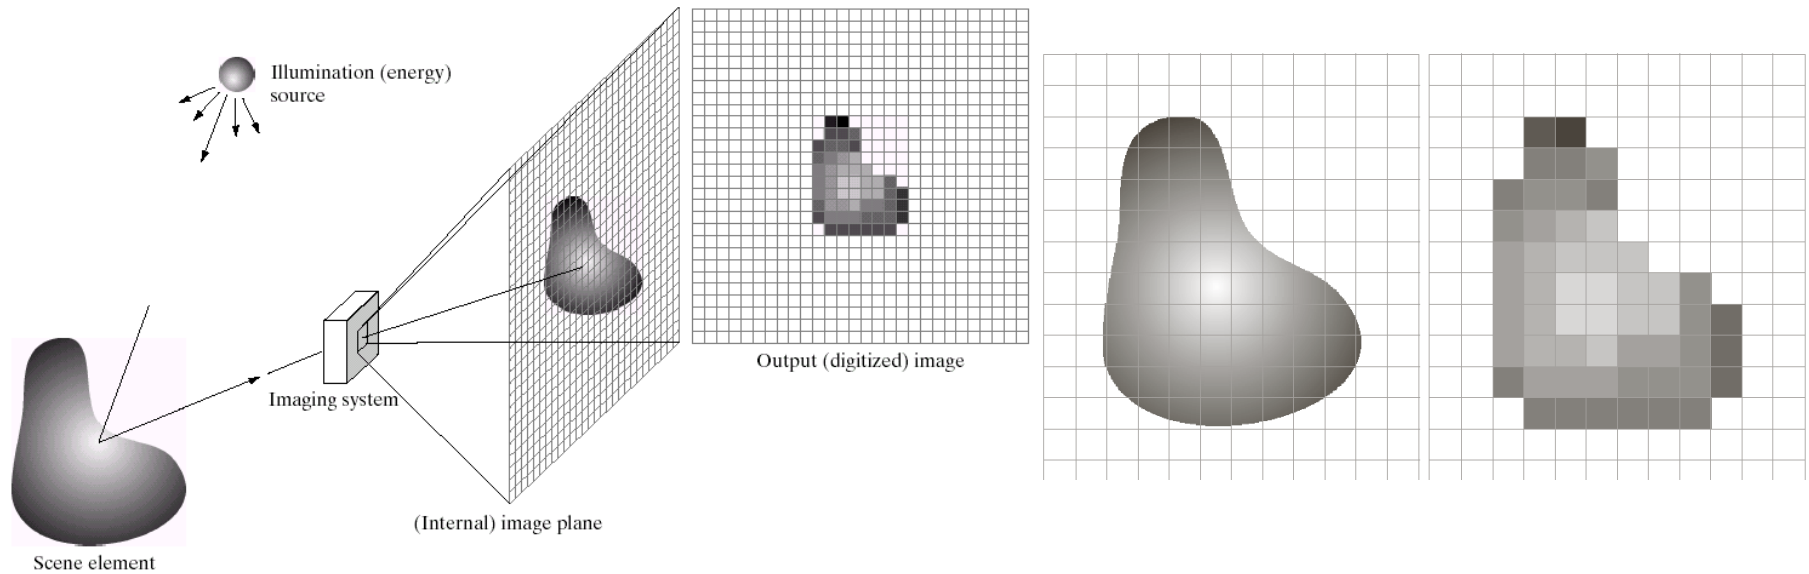


- What is measured in an image location?

- brightness
- color

<< viewpoint
illumination conditions
local geometry
local material properties

Image Formation



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

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 $f(x,y)$





Figure: M. J. Black

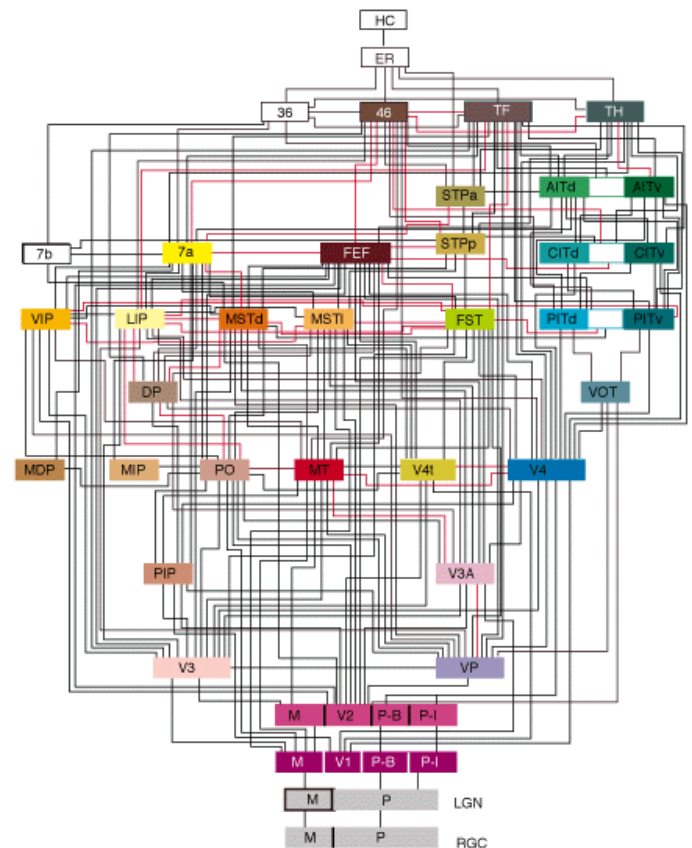
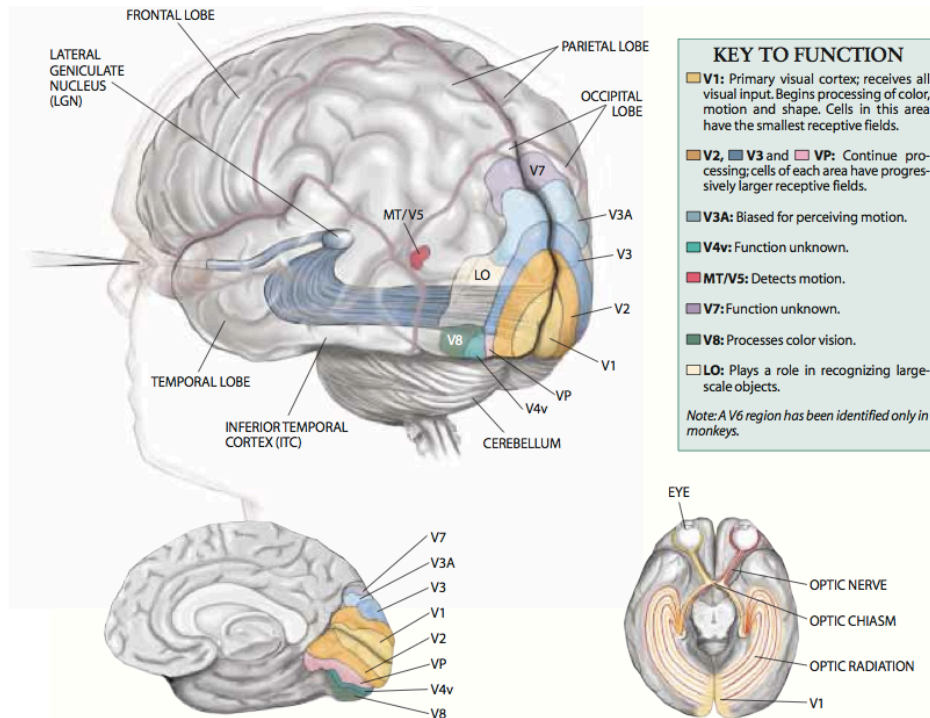
Human Eye



- Two types of receptor cells in retina:
 - Cone Receptor cells: 6-7 million → function in bright light, color sensitive, fine detail
 - Rod receptor cells: 75-150 million → function in dim light, color insensitive, coarse detail
- A recent discovery: Photosensitive retinal ganglion cells → sensitive to blue light

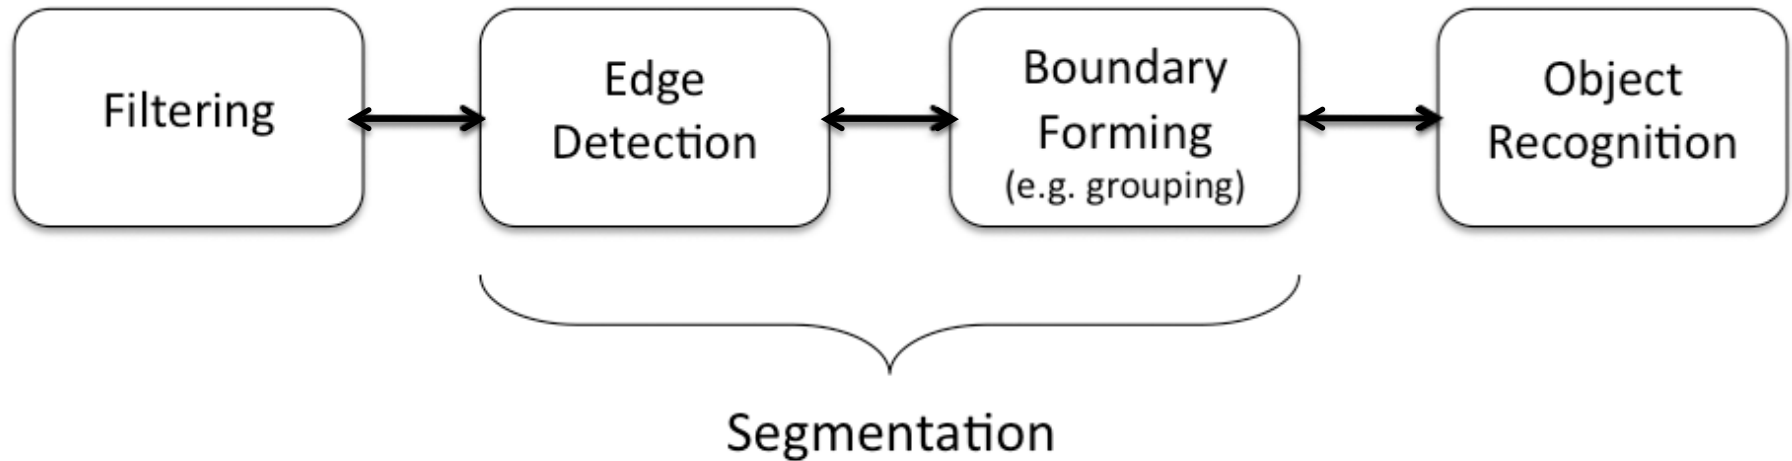
Hierarchy of Visual Areas

- There are many different neural connections between different visual areas.



Figures: Nikos K. Logothetis, Vision: A Window on Consciousness, SciAm, Nov 1999F (on the left)
 Felleman & van Essen, 1991 (on the right)

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

Sample Problems

- Edge Detection
- Image Denoising
- Image Smoothing
- Image Segmentation
- Image Registration
- Image Inpainting
- Image Retargeting
- ...

Image Filtering

- Filtering out the irrelevant information

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

↓ ↓ ↓

observed desired irrelevant
image image 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

- **Difficulty:** 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
 - the most widely used method



$$\frac{\partial u}{\partial t} = \nabla \cdot (\nabla u) = \nabla^2 u$$




- 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



What do these examples demonstrate?



Noisy input

Recovered image

Original image

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

Context-Guided Filtering

- Use local image context to steer filtering



Preserve main image structures during filtering



Image Smoothing



L. Xu, C. Lu, Y. Xu, J. Jia, Image Smoothing via L0 Gradient Minimization, SIGGRAPH ASIA 2011

Image Smoothing



L. Karacan, E. Erdem, A. Erdem, work in progress

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!

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.

Snakes


- Curve Evolution - parametric curve formulation



Non-rigid,
deformable objects
can change their
shape over time,
e.g. lips, hands...

Normalized Cuts

- A graph-theoretic formulation for segmentation




Normalized Cuts



slide credit: S. Lazebnik


From contours to regions

- State-of-the-art: gPb-owt-ucm segmentation algorithm




From contours to regions

- State-of-the-art: gPb-owt-ucm segmentation algorithm



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

Registration

- Estimate a transformation function between
 - two images
 - two point sets
 - two shapes
 - ...

Registration

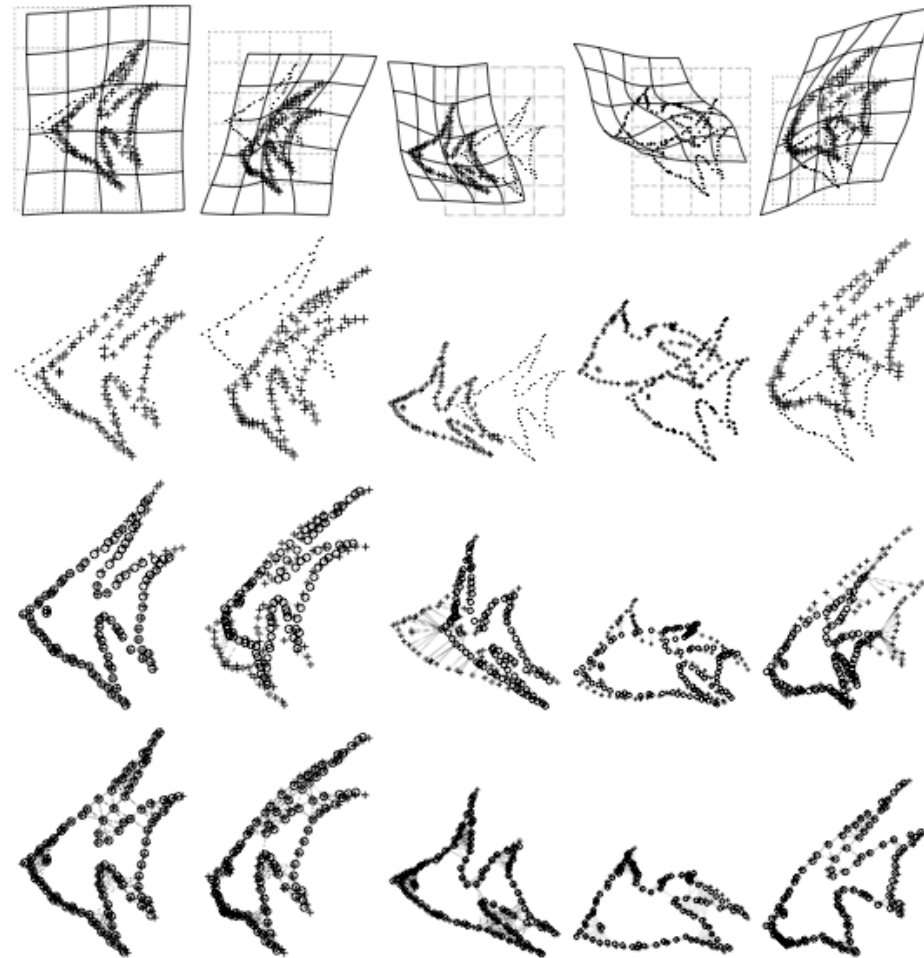


Fig. 5. Experiments on deformation. Each column represent one example. From left to right, increasing degree of deformation. *Top row*: warped template. *Second row*: template and target (same as the warped template). *Third row*: ICP results. *Bottom row*: RPM results.

Image Registration






Fig. 2. An example of a geodesic between images (original images taken from the Olivetti face database). The three intermediate images are generated by the optimization algorithm.



A tumor progressively appearing on a brain



Tumor: Reference image, registered target and deformation

(top) Alain Trounev and Laurent Younes, *Metamorphoses Through Lie Group Action*, *Found. Comput. Math.*, 2005
(bottom) M. I. Miller and L. Younes, *Group Actions, Homeomorphisms, and Matching: A General Framework*, *IJCV*, 2001

Image Inpainting

- Reconstructing lost or deteriorated parts of images



What do these examples demonstrate?





Since 1699, when French explorers landed at the great bend of the Mississippi River and celebrated the first Mardi Gras in North America, New Orleans has brewed a fascinating melange of cultures. It was French, then Spanish, then French again, then sold to the United States. Through all these years, and even into the 1900s, others arrived from everywhere: Acadians (Cajuns), Africans, indige-



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

Input

Seam Carving

GBVS

sigLab

sigRGB

our map

