Image Processing and Computer Vision Introduction

Amartya Kundu Durjoy

Lecturer

CSE Department (UGV)

What is an image?

- We can think of an **image** as a function, *f*, from R^2 to R:
  - $f(x, y)$ gives the **intensity** at position (x, y)
  - Realistically, we expect the image only to be defined over a rectangle, with a finite range:
    - f: $[a,b]$x$[c,d]$ → $[0,1]$
    - 0 → black; 1 → white; in-between → gray
- A color image is just three functions pasted together. We can write this as a "vector-valued" function:

$$f(x,y) = \begin{bmatrix} r(x,y) \\ g(x,y) \\ b(x,y) \end{bmatrix}$$

Analog Image

- An image can be understood as a 2D light intensity function f(x,y) where:
  - x and y are spatial coordinates
  - The value of f at any point (x, y) is proportional to the brightness or gray value of the image at that point
- Cannot be stored as such on a digital computer.

Digital Image

Recall two ways of visualizing an image

Intensity pattern

2d array of numbers

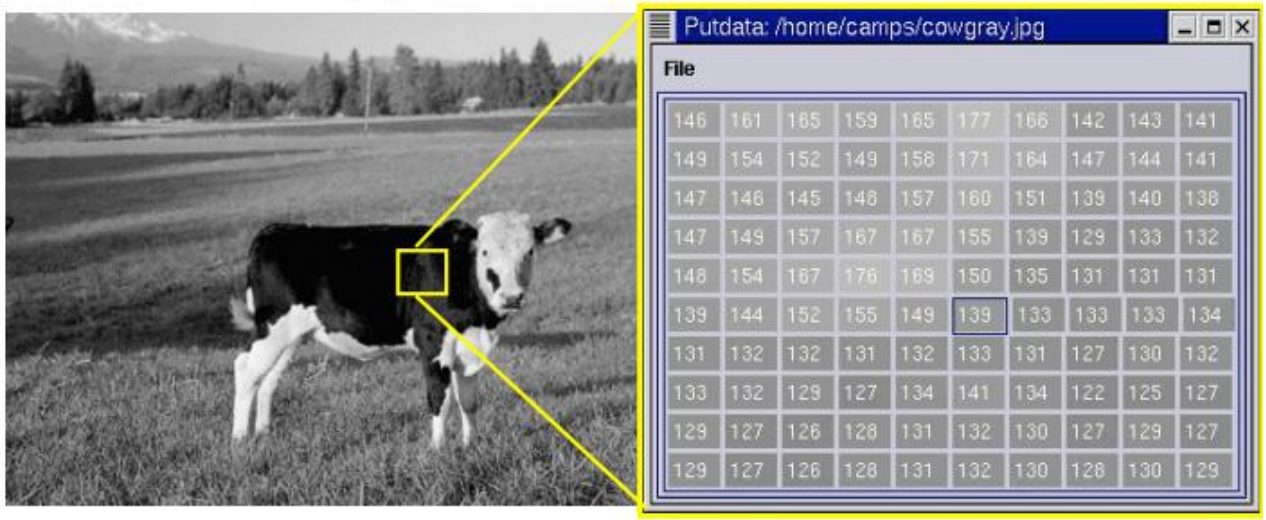

We "see it" at this level

Computer works at this level

- A digitized image is one in which:
  - Spatial and grayscale values have been made discrete.
  - Intensities measured across a **regularly spaced grid** in x and y directions are sampled to
    - 8 bits (256 values) per point for black and white,
    - 3x8 bits per point for color images.
  - Stored as a 2D arrays of gray-level values. The array elements are called pixels and identified by their x, y coordinates.

Image Representation

- **Discrete representation of images**
  - we'll carve up image into a rectangular grid of pixels $P[x,y]$
  - each pixel p will store an intensity value in [0 1]
  - 0 → black; 1 → white; in-between → gray
  - Image size mxn → (mn) pixels

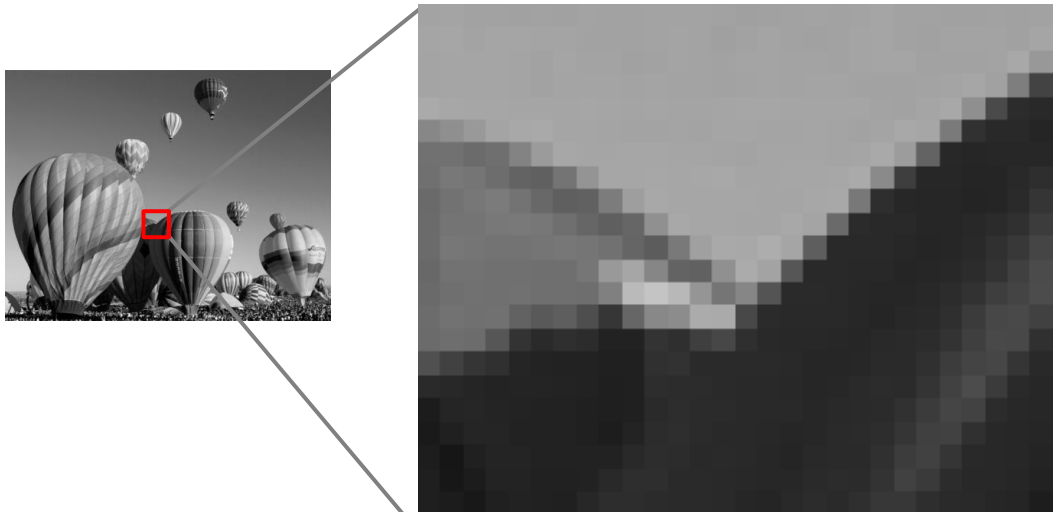

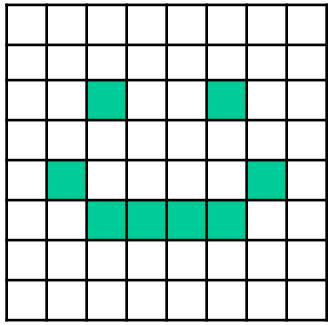

Color Image

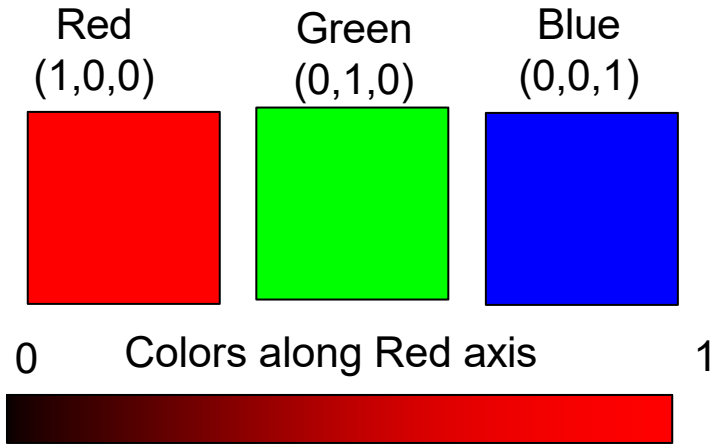

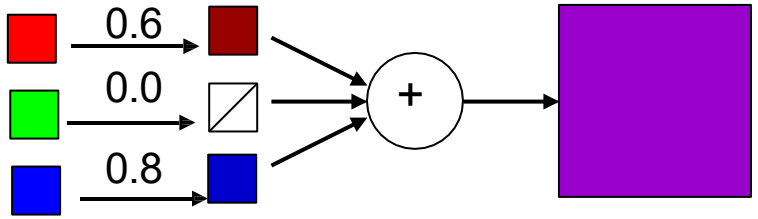

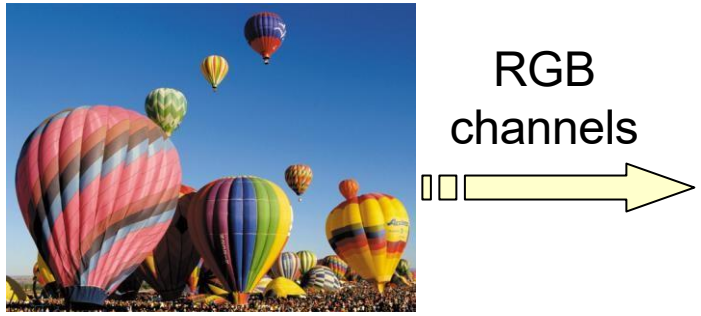

RGB channels

What is a Digital Image?

•A **digital image** is a representation of a two-dimensional image as a finite set of digital values, called picture elements or pixels

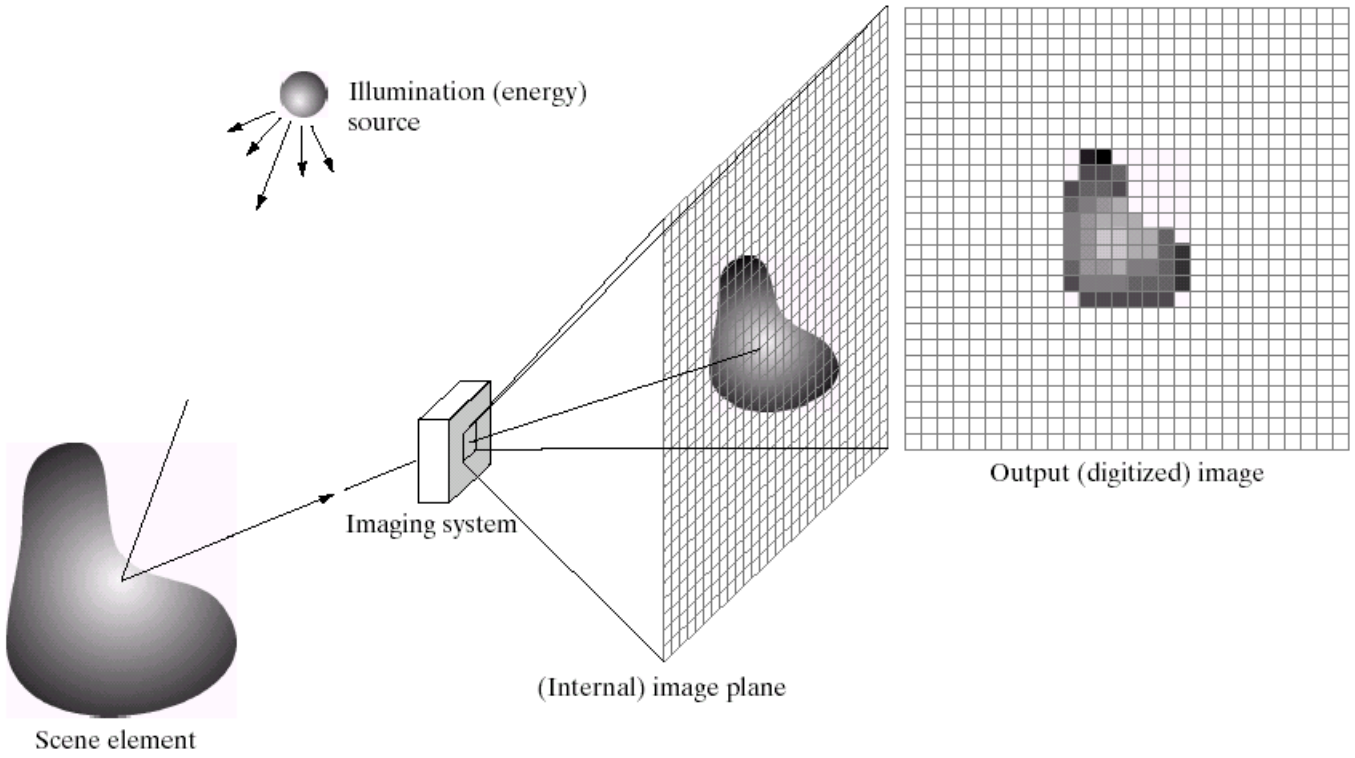

What is a Digital Image? (cont…)

- Pixel values typically represent gray levels, colors, heights, opacities etc
- **Remember** *digitization* implies that a digital image is an *approximation* of a real scene

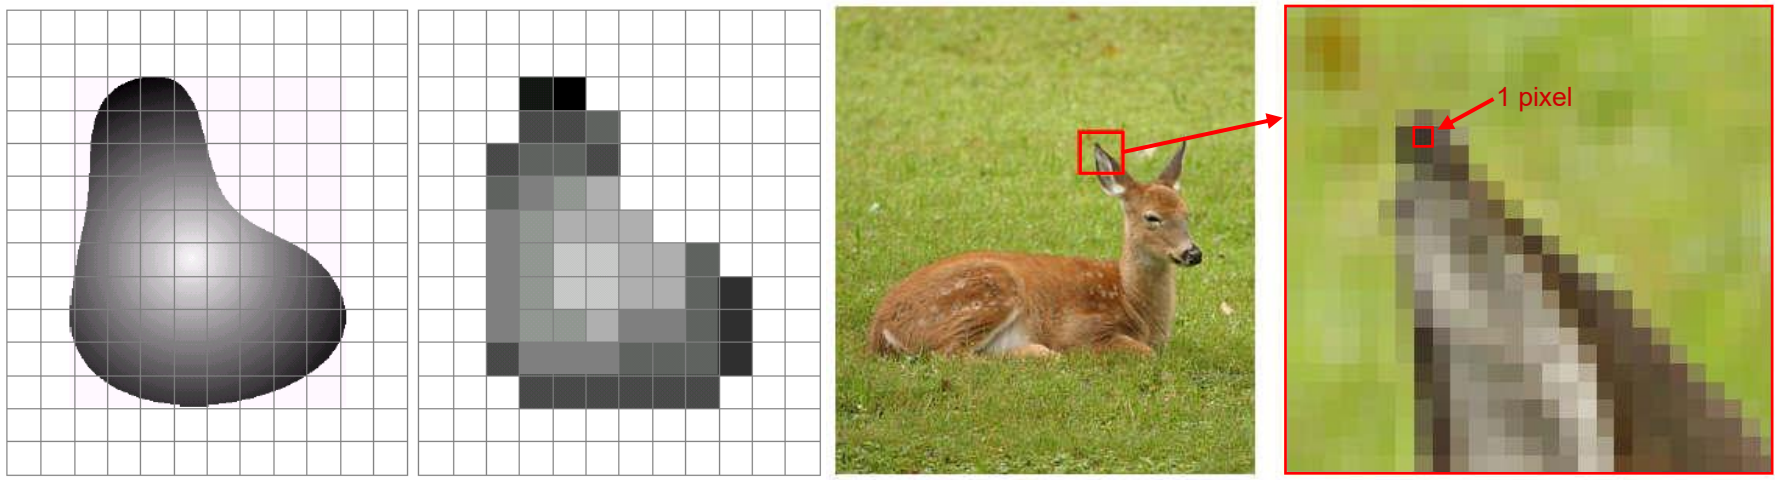

What is a Digital Image? (cont…)

- Common image formats include:
- 1 sample per point (B&W or Grayscale)
- 3 samples per point (Red, Green, and Blue)
- 4 samples per point (Red, Green, Blue, and “Alpha”, a.k.a. Opacity)

- For most of this course we will focus on grey-scale images

Images as Surfaces

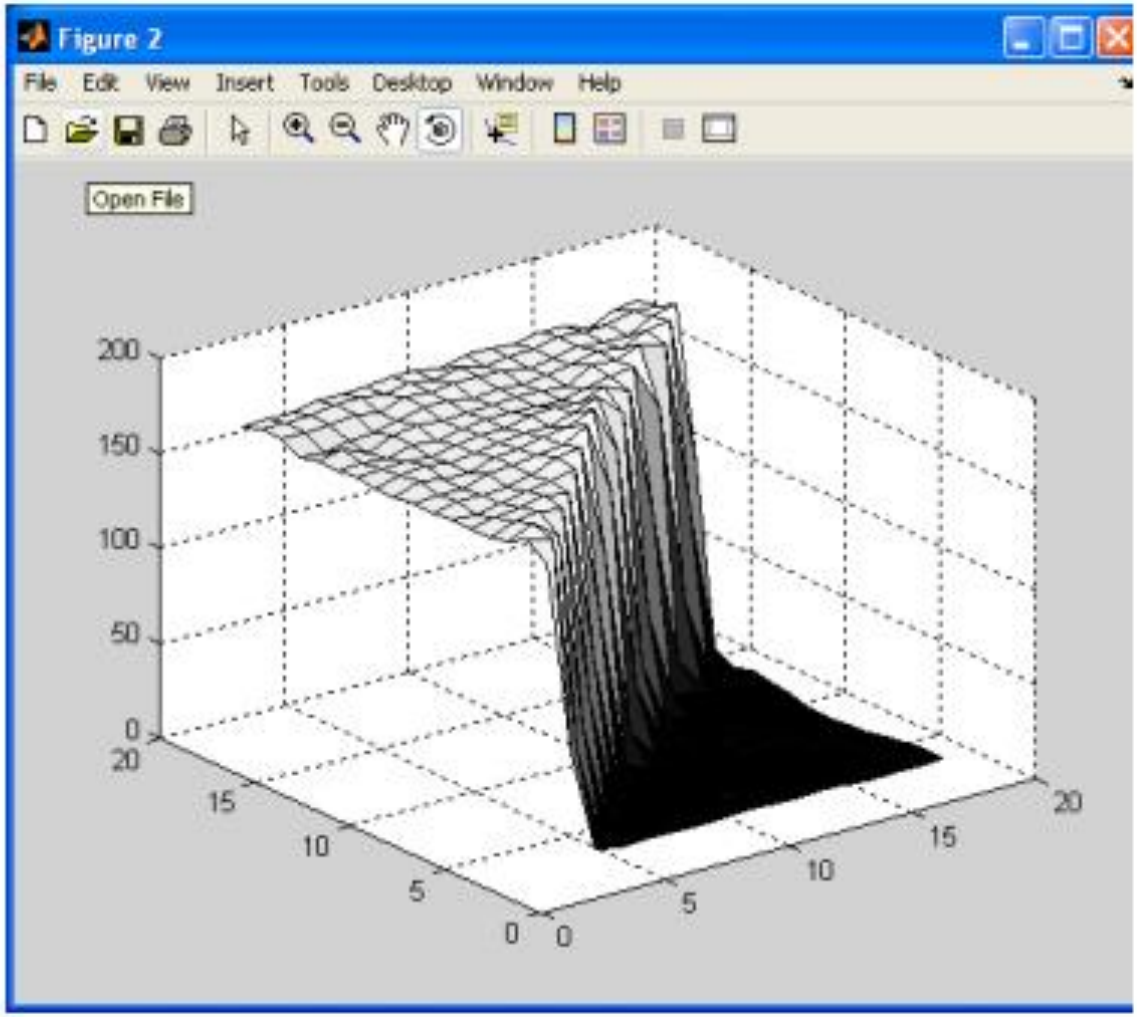

What is Digital Image Processing?

- Digital image processing focuses on two major tasks
  - Improvement of pictorial information for human interpretation
  - Processing of image data for storage, transmission and representation for autonomous machine perception
- Some argument about where image processing ends and fields such as image analysis and computer vision start

Image Processing

Hubble telescope – image restoration example:

- A defective mirror made many of Hubble's images useless.

- Image restoration techniques were used to improve image quality before fixing the problem.

Image Processing

Image Compression

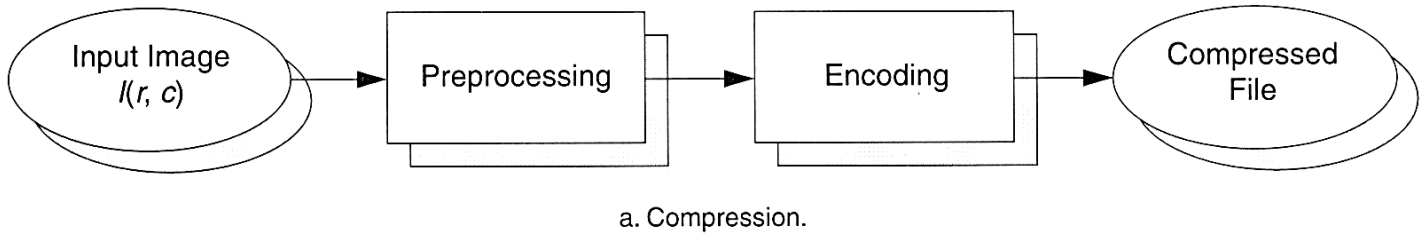

a. Compression.

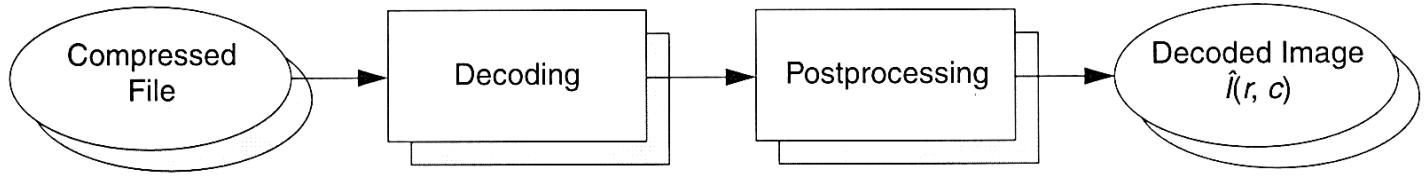

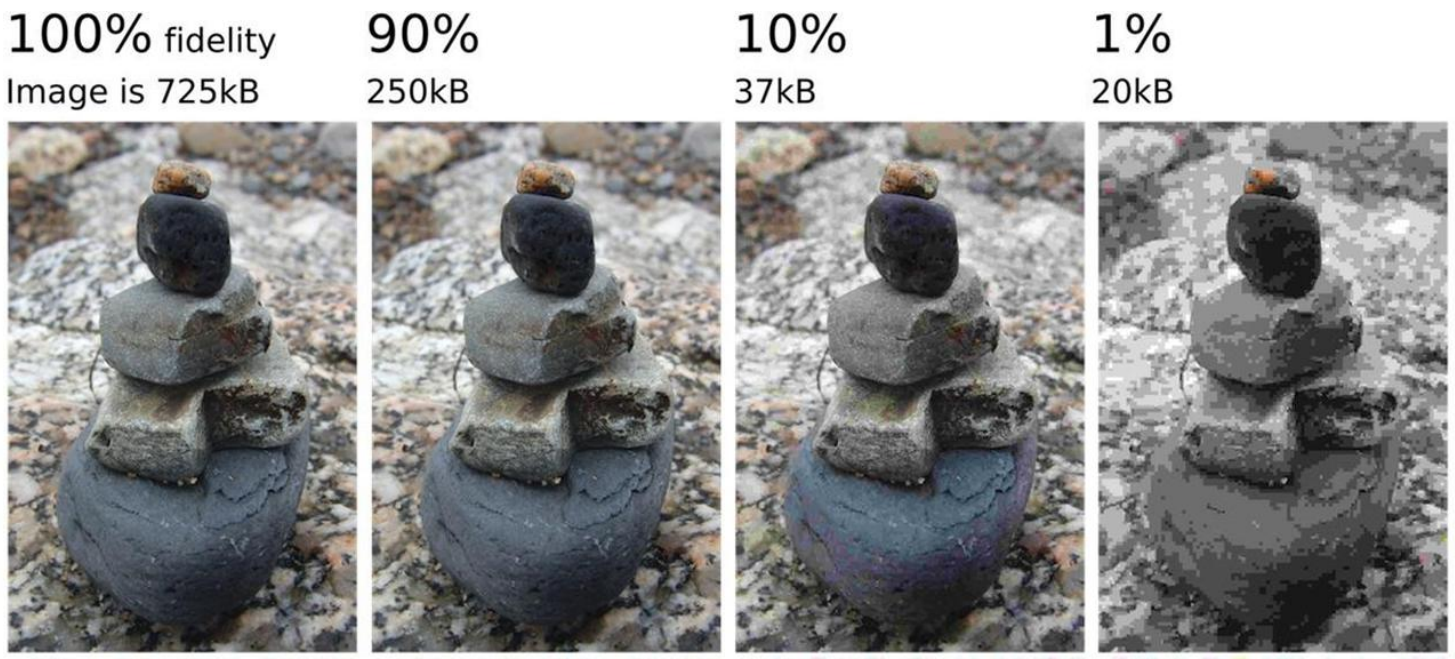

Computer Vision

- Make computers understand images and video.
  - Computing properties of the 3D world from visual data *(measurement)*
  - Algorithms and representations to allow a machine to recognize objects, people, scenes, and activities. *(perception and interpretation)*

What kind of scene?

Where are the cars?

How far is the building?

…

What is Computer Vision?

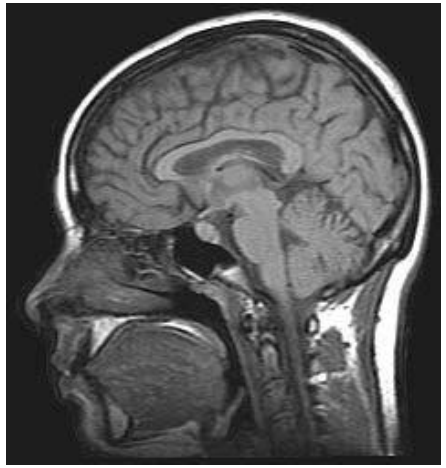

- **Computer vision** is the science and technology of machines that see.
- Concerned with the theory for building artificial systems that obtain information from images.
- The image data can take **many forms**, such as a **video sequence**, **depth images**, views from multiple cameras, or multi-dimensional data from a **medical scanner**

Examples

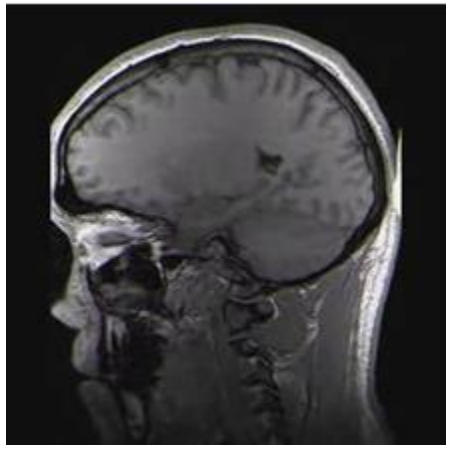

1. Brain MRI

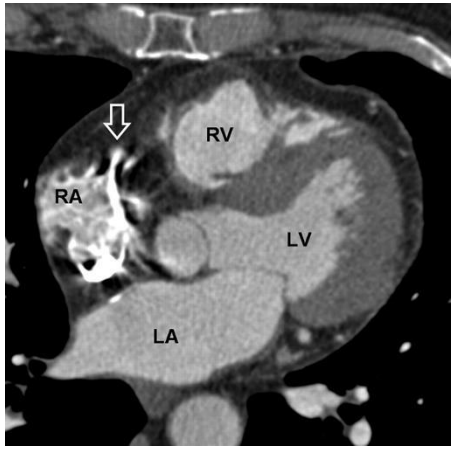

2. Cardiac CT

3. Fetus Ultrasound

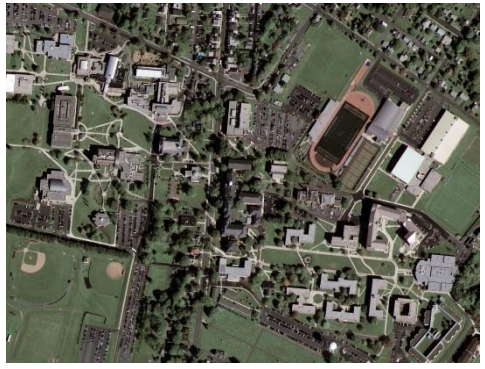

4. Satellite image

5. IR image

1 and 3. http://en.wikipedia.org
2. http://radiology.rsna.org
4. http://emap-int.com
5. http://www.imaging1.com

DIP to CV

•The continuum from image processing to computer vision can be broken up into low-, mid- and high-level processes

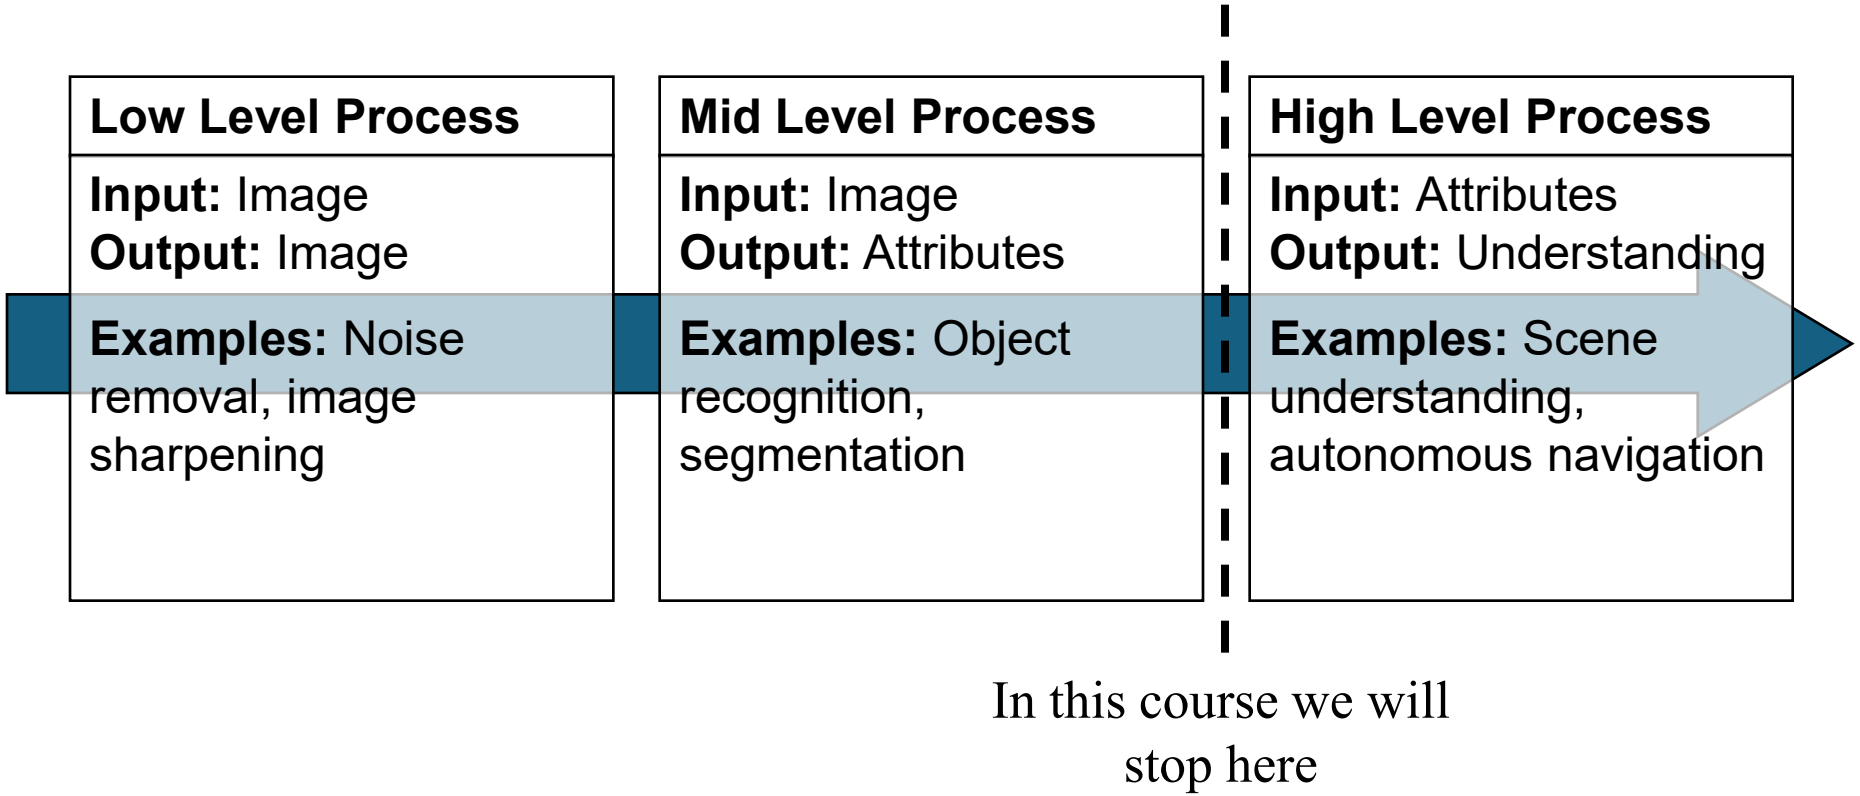

Image Processing vs Computer Vision

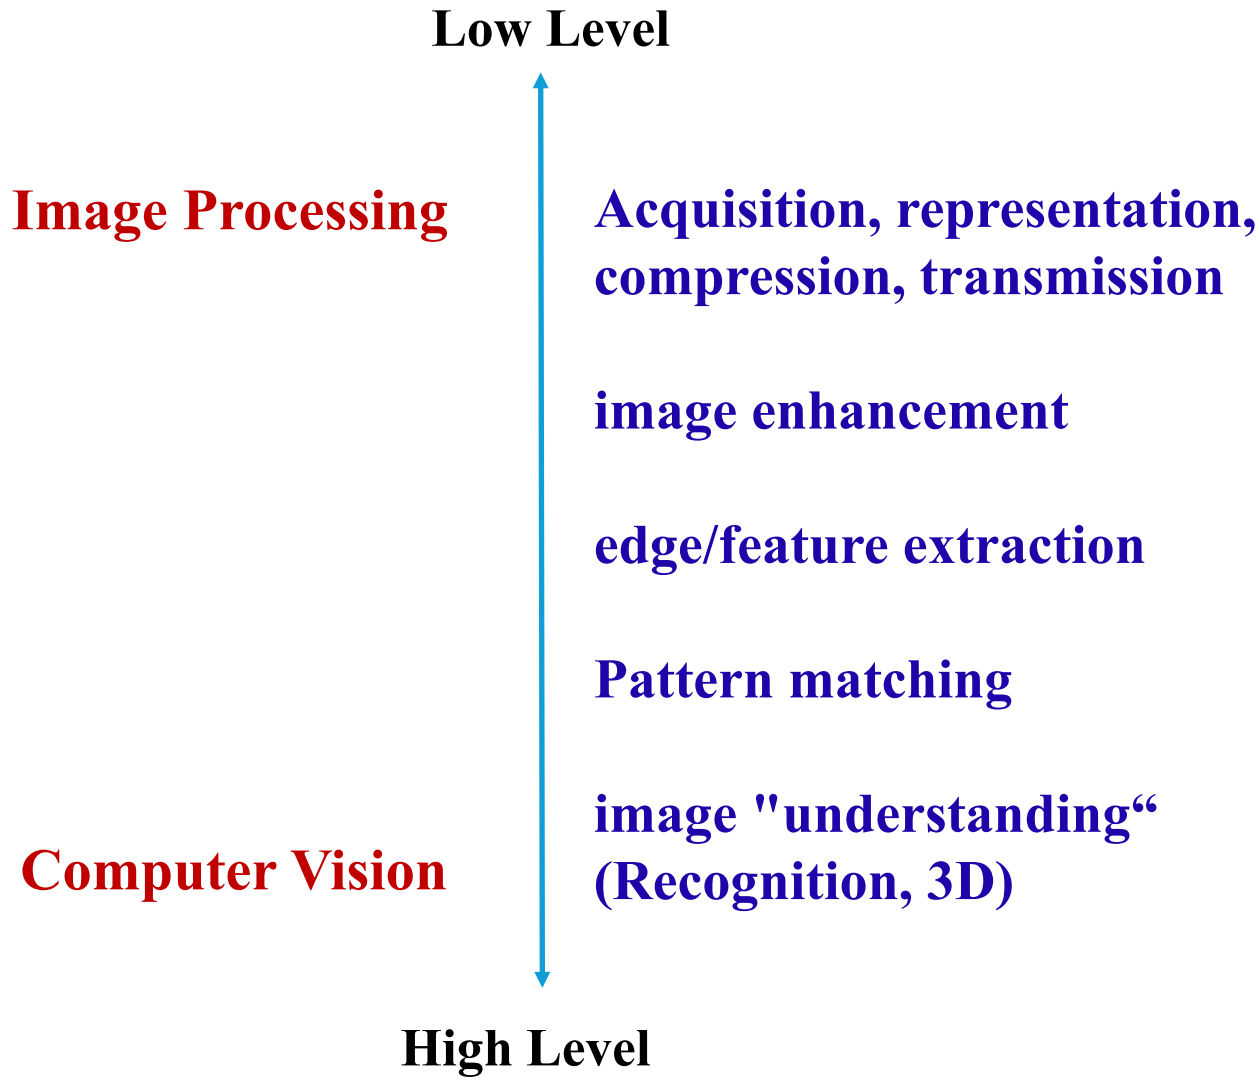

Image Processing → Image Analysis

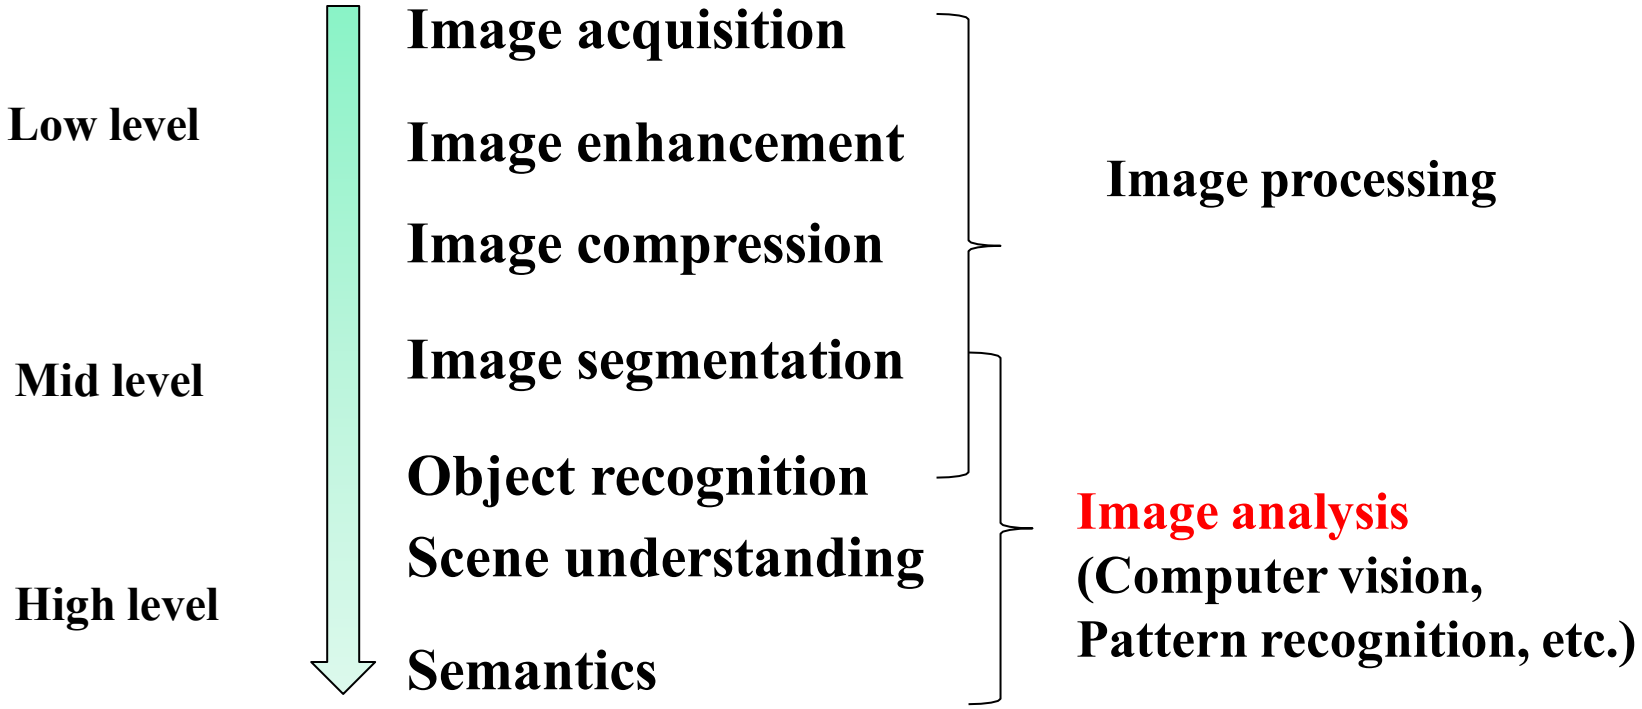

Why Computer Vision is Hard?

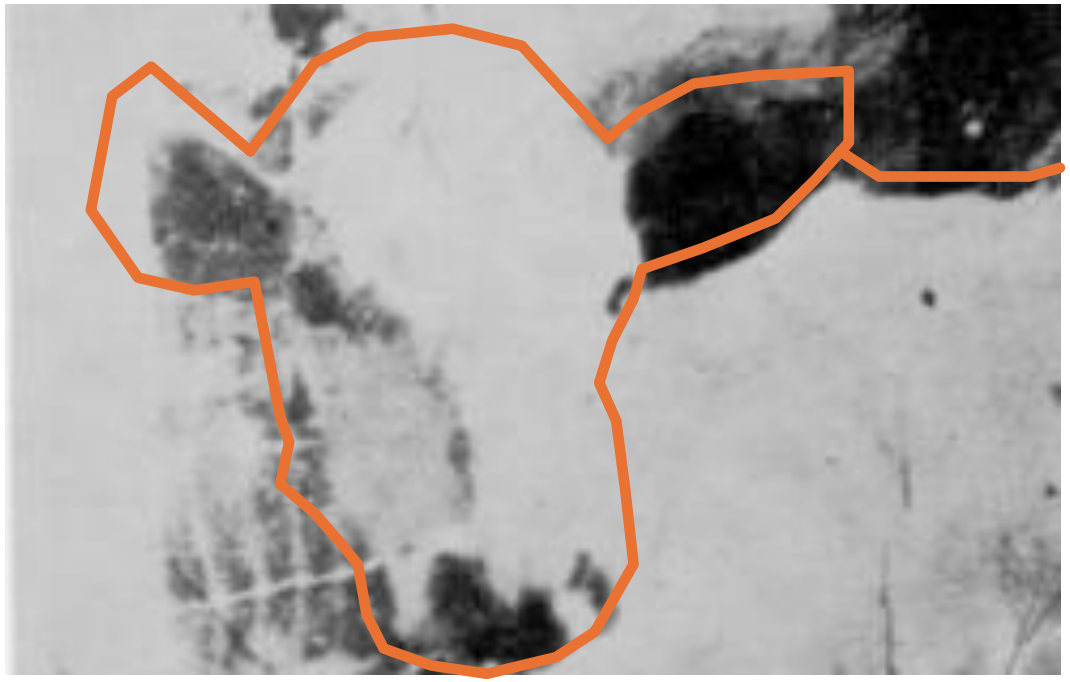

Vision is really hard

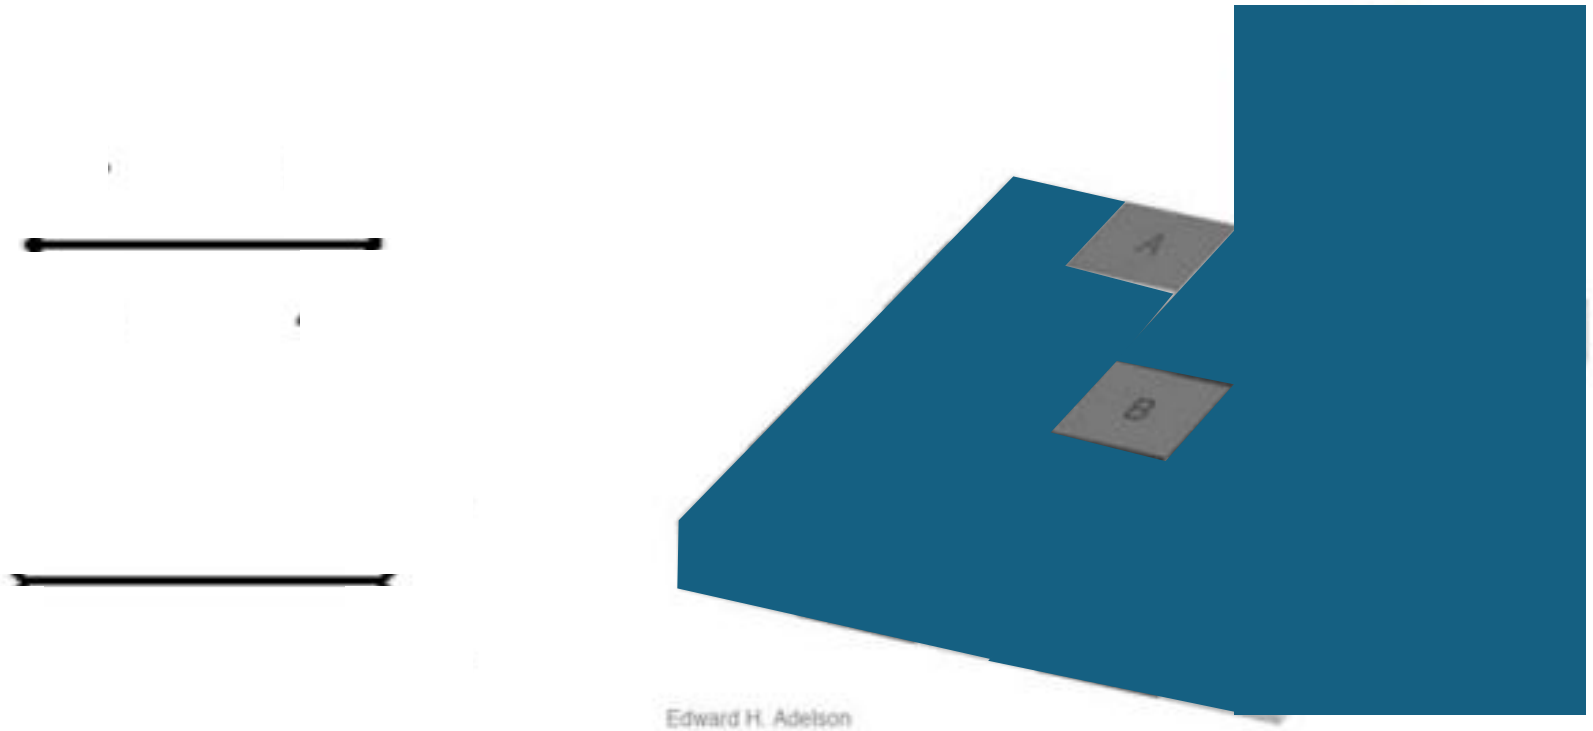

(a) (b)

Perceived Intensity is Not a Simple Function of the Actual Intensity (1)

a
b
c

FIGURE 2.7
Illustration of the Mach band effect. Perceived intensity is not a simple function of actual intensity.

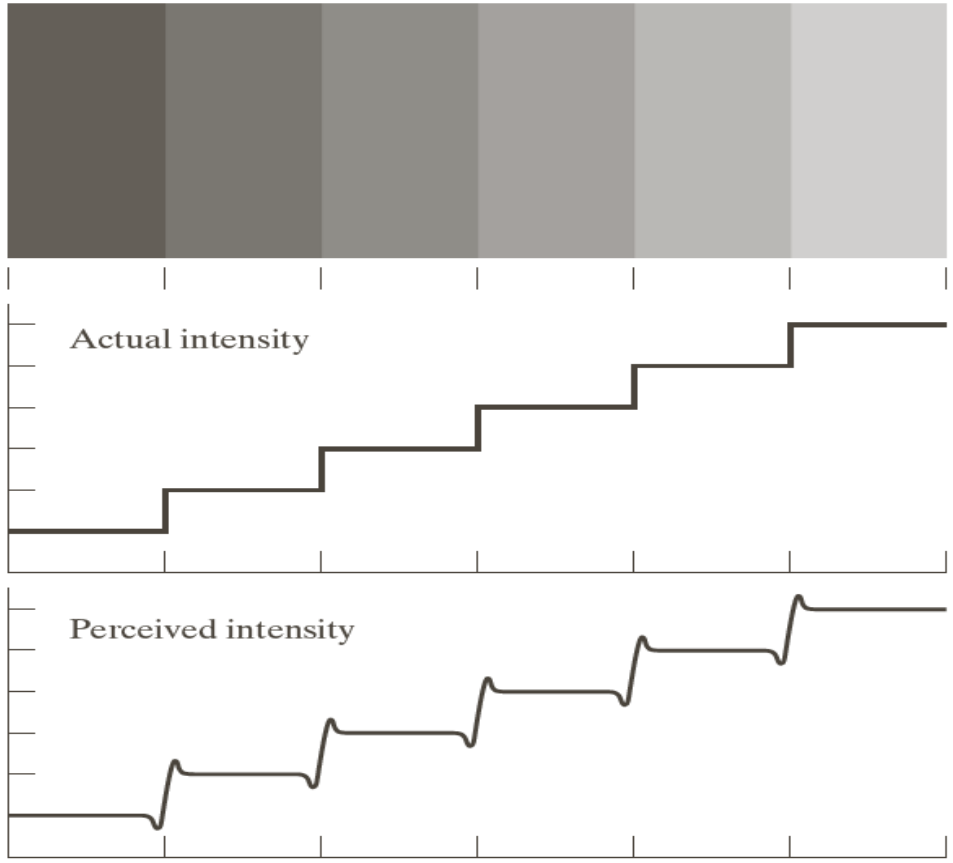

Perceived Intensity is Not a Simple Function of the Actual Intensity – Simultaneous Contrast

a b c

FIGURE 2.8 Examples of simultaneous contrast. All the inner squares have the same intensity, but they appear progressively darker as the background becomes lighter.

Optical Illusions: Complexity of Human Vision

a b
c d

FIGURE 2.9 Some well-known optical illusions.

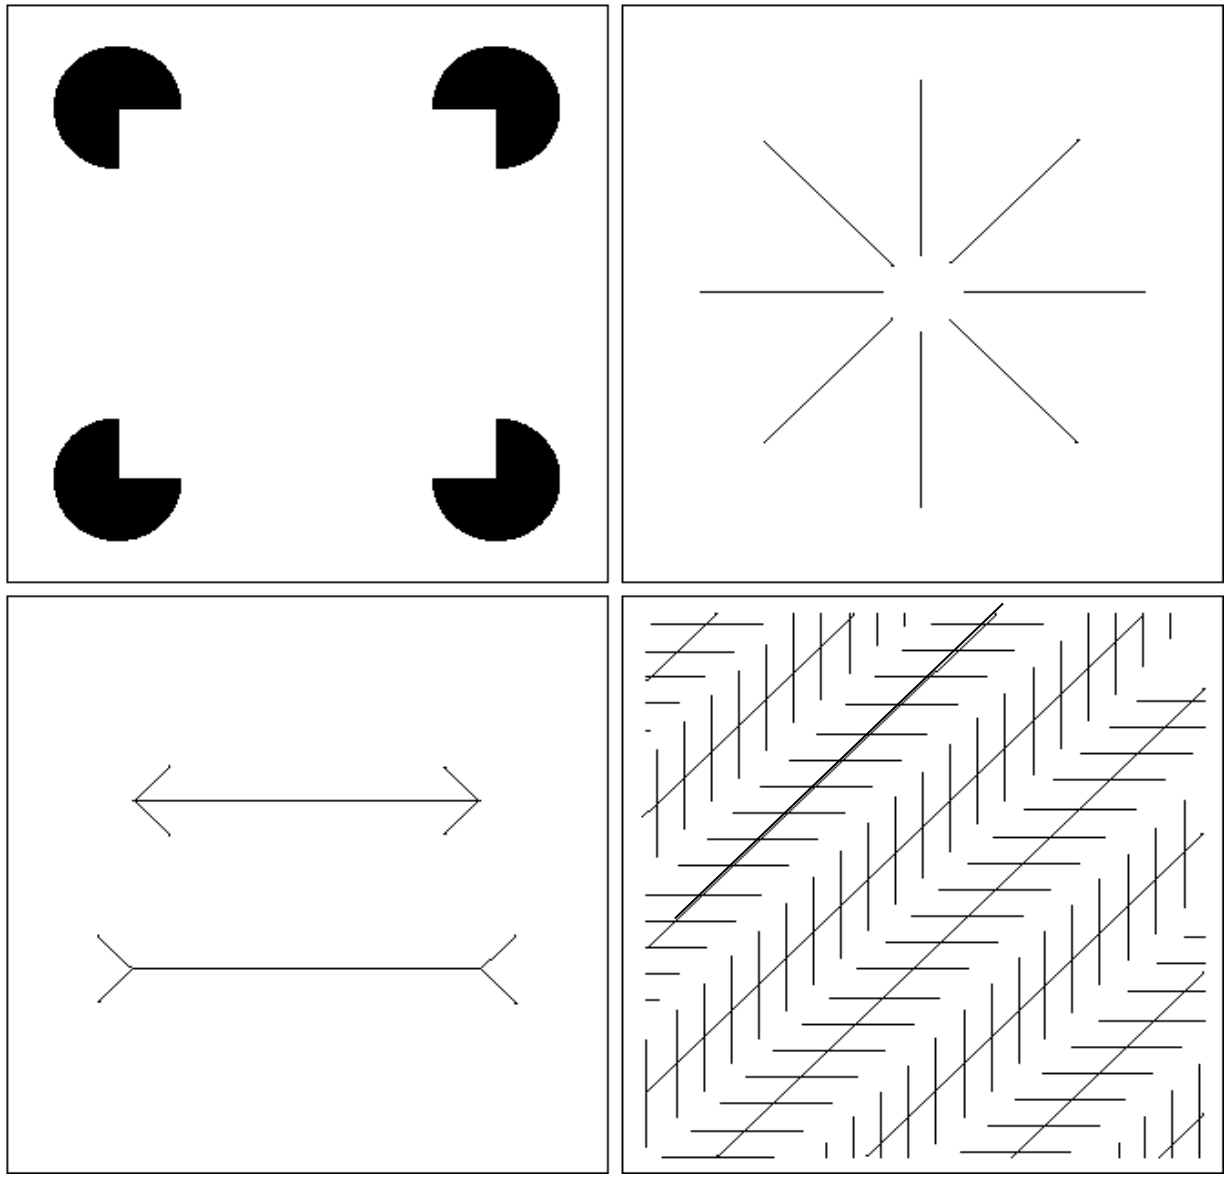

More Optical Illusions

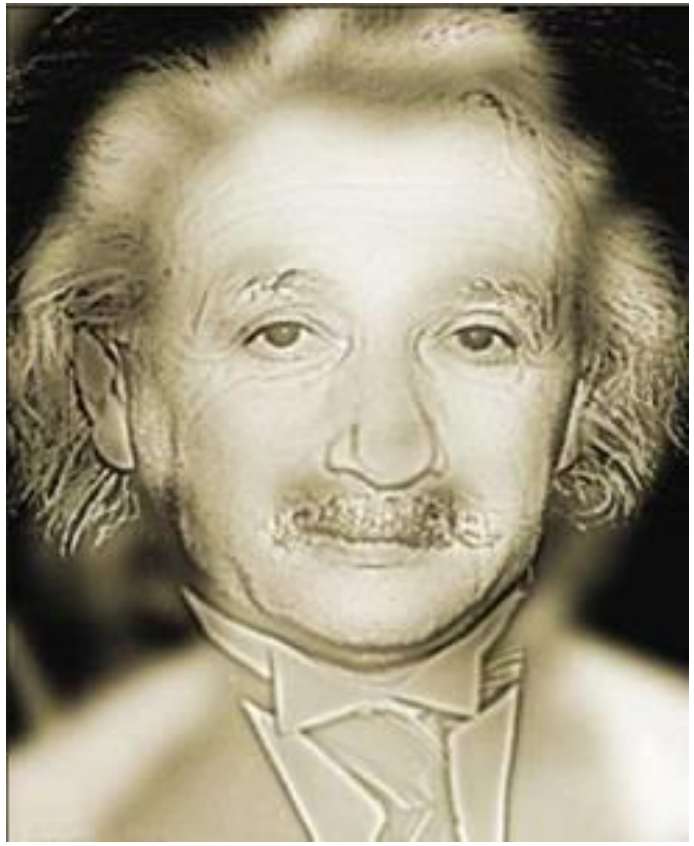

http://www.123opticalillusions.com/

http://brainden.com/optical-illusions.htm

Vision is Challenging

- Inverse problems
- Apriori-knowledge is required
- Complexity is extensive
- Non-local operations

Vision is really hard

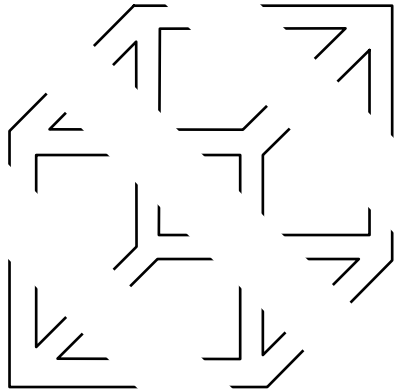

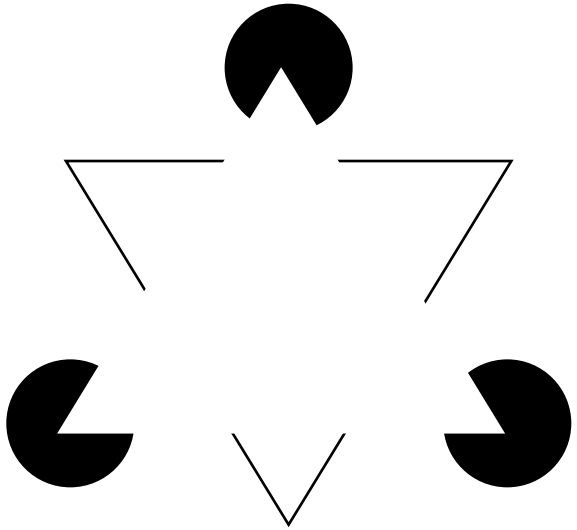

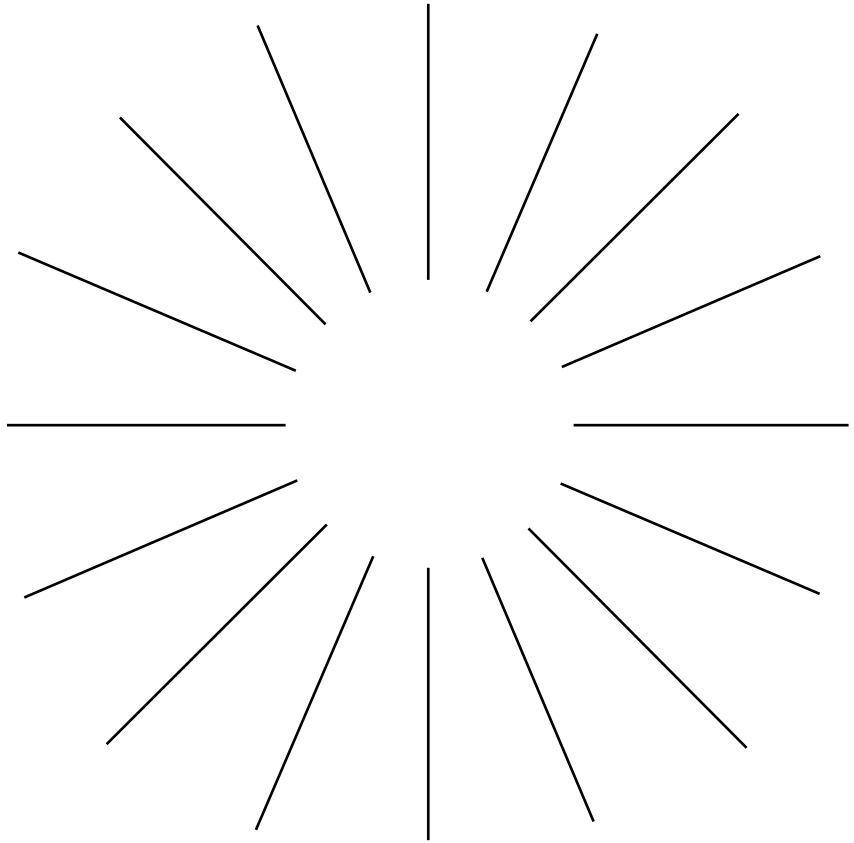

Vision is really hard

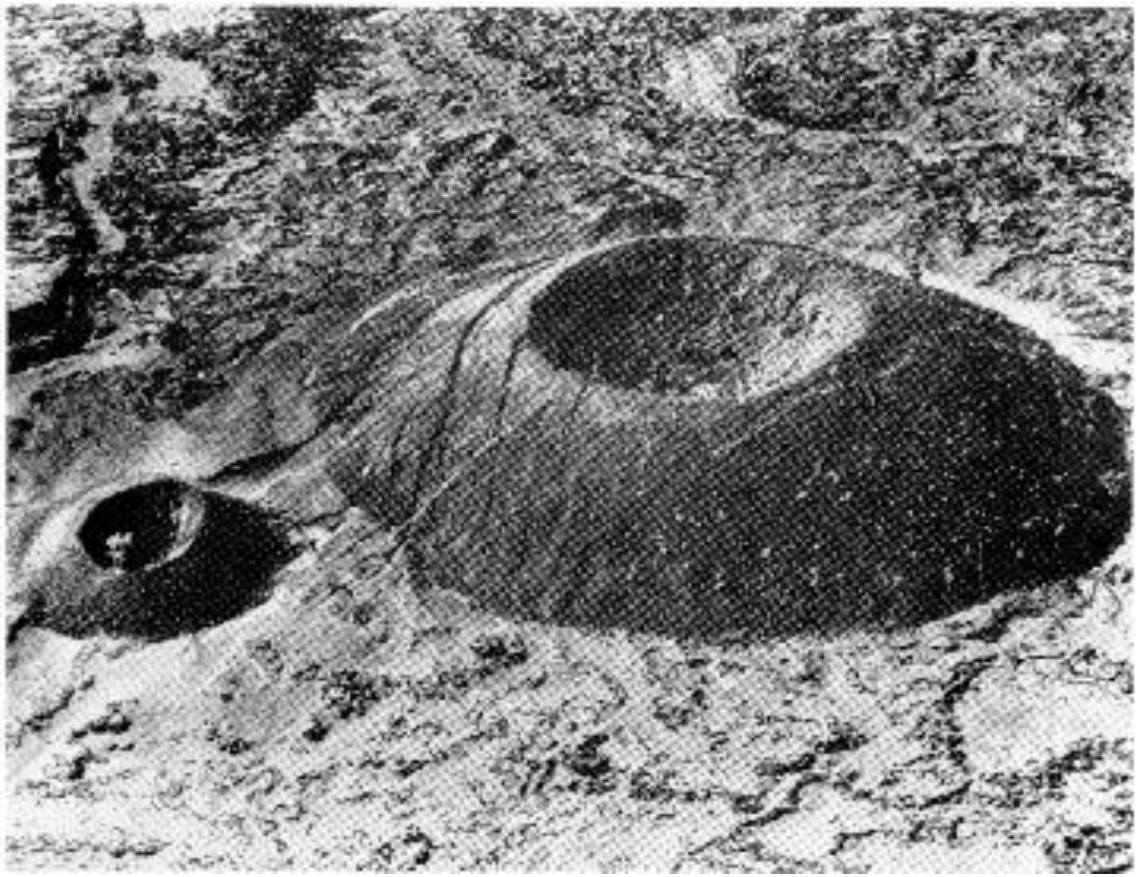

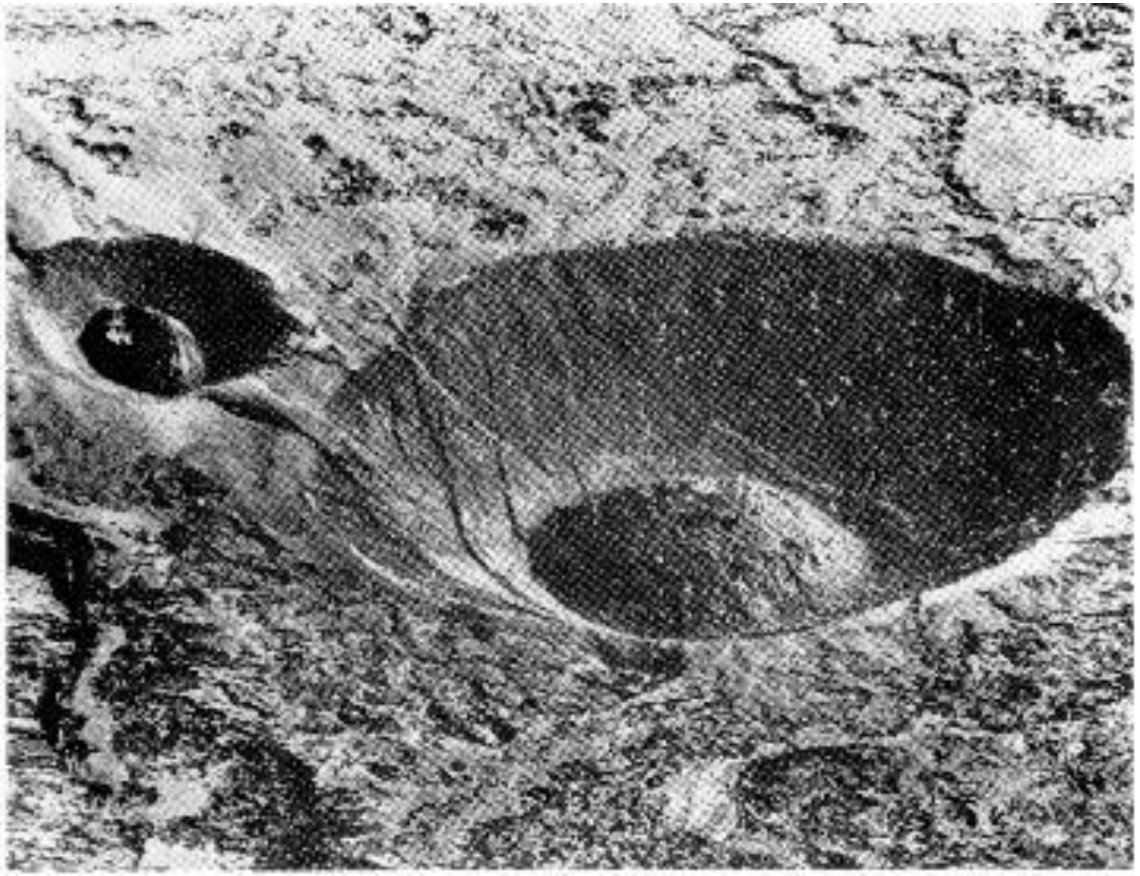

Components of a computer vision system

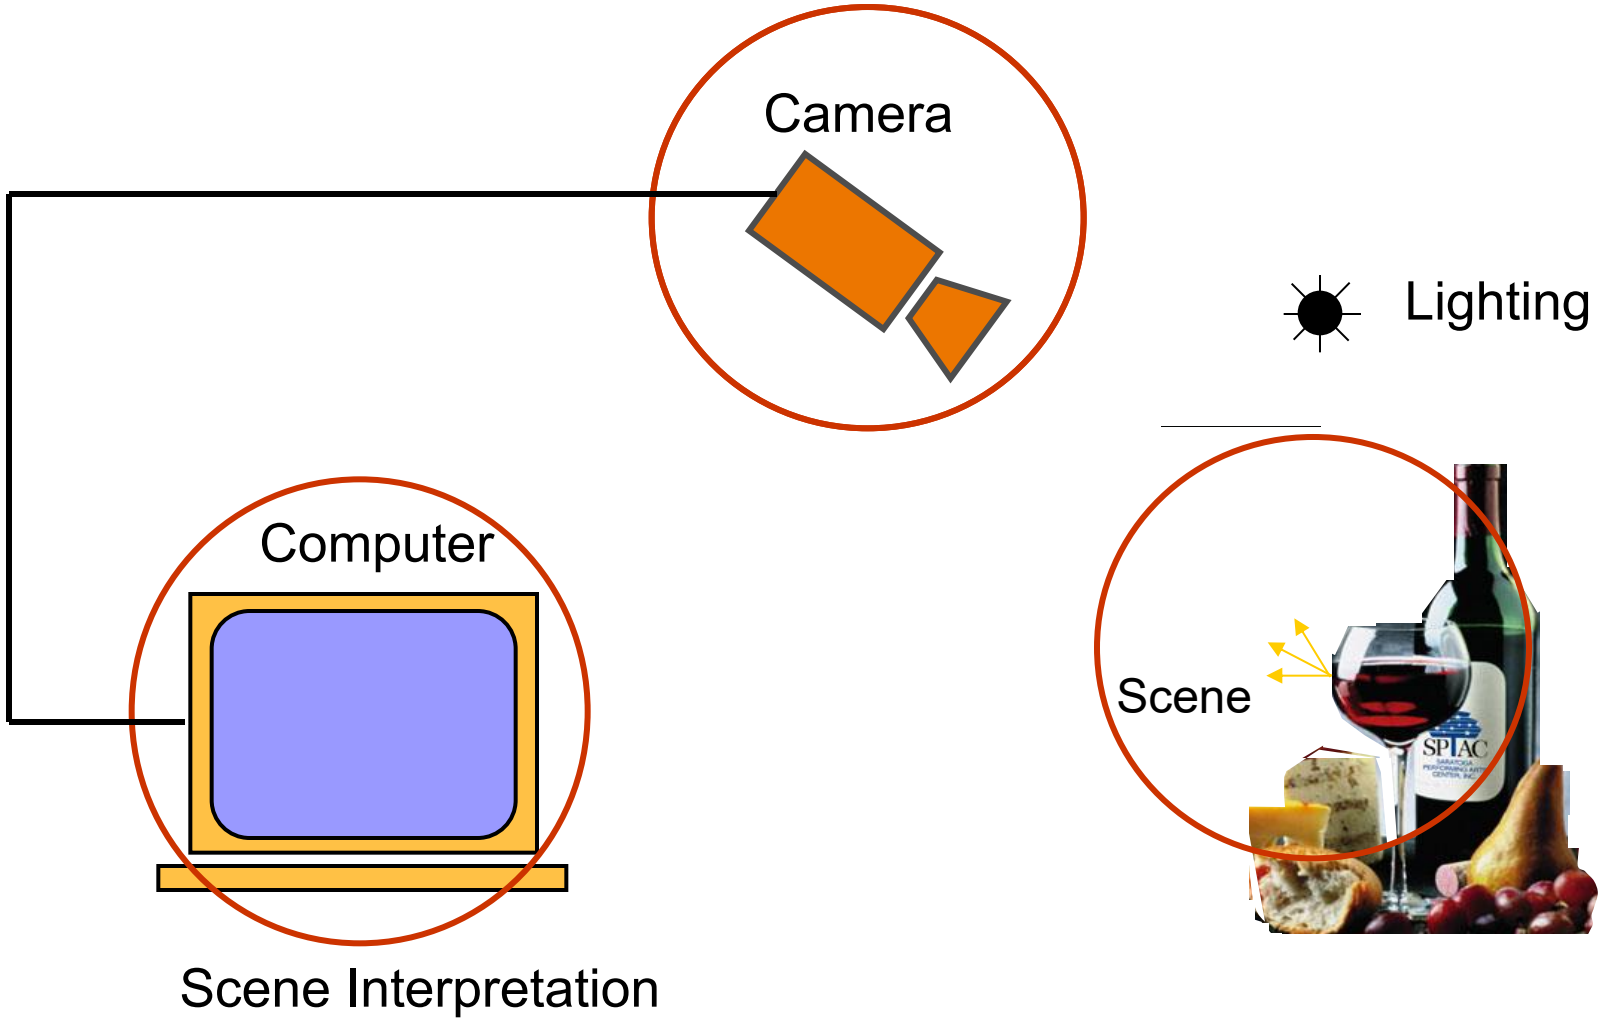

Related disciplines

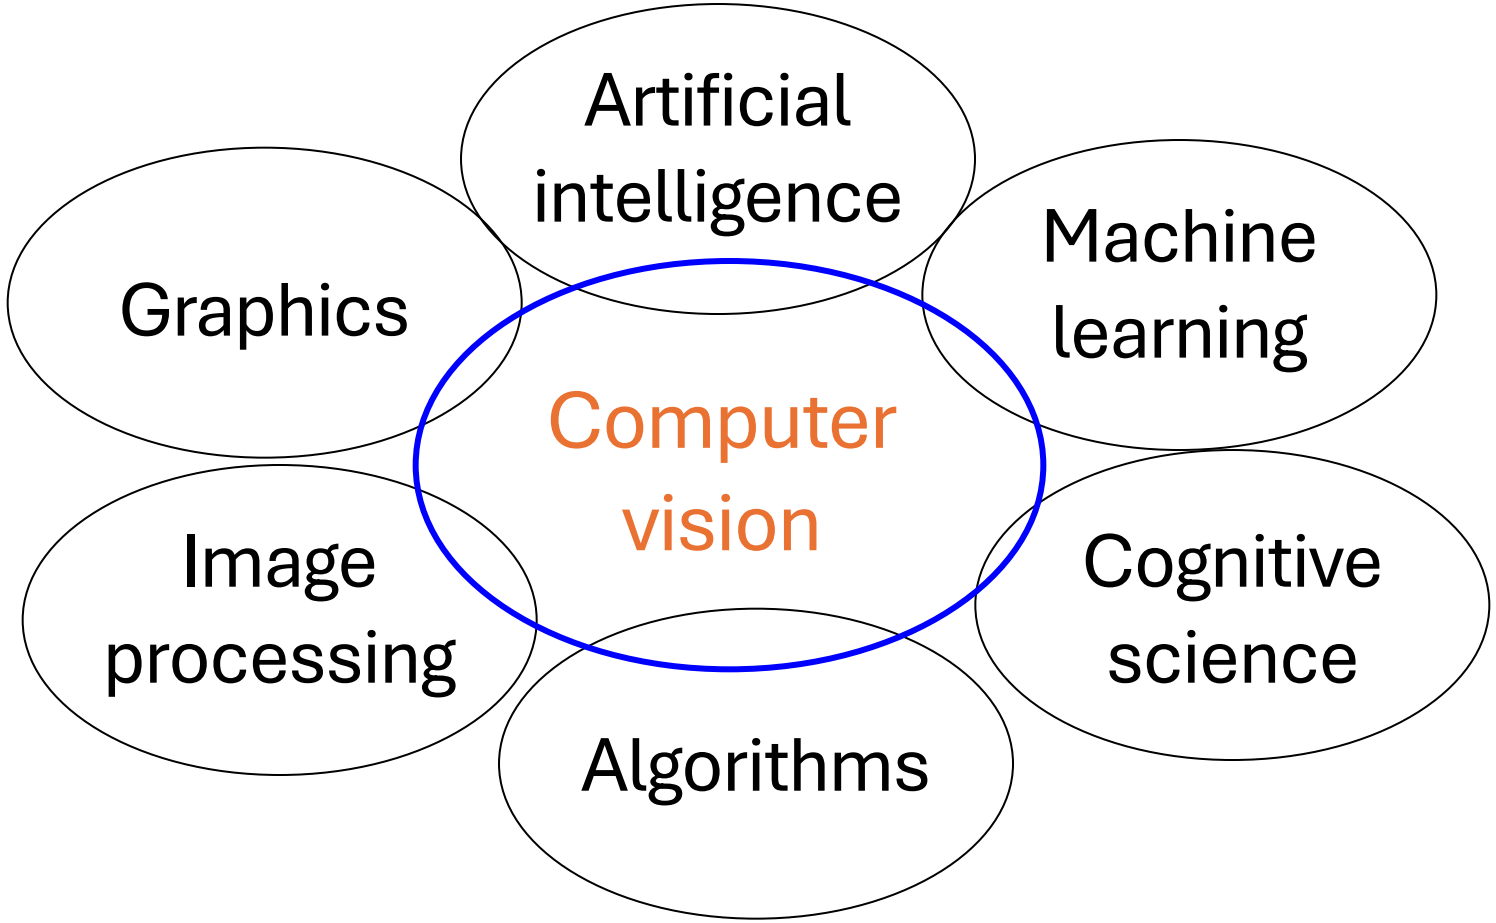

Vision and graphics

- Inverse problems: analysis and synthesis.

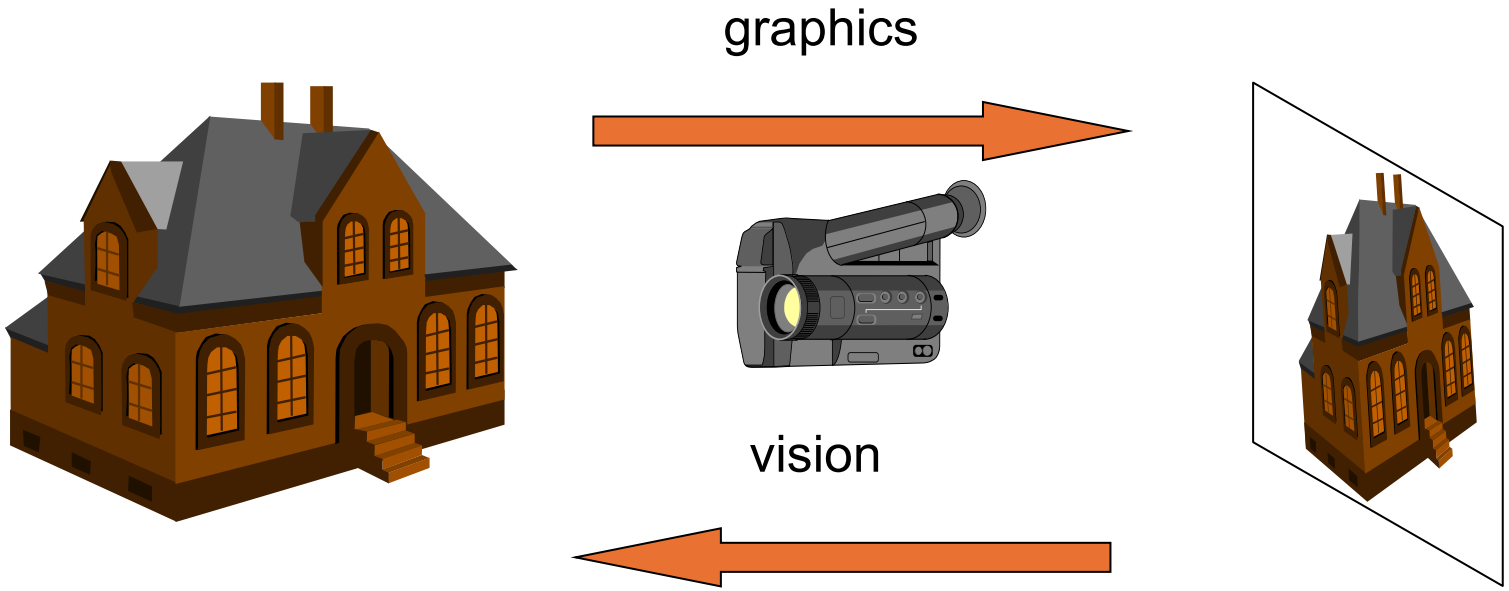

What is computer vision? (2D->3D)

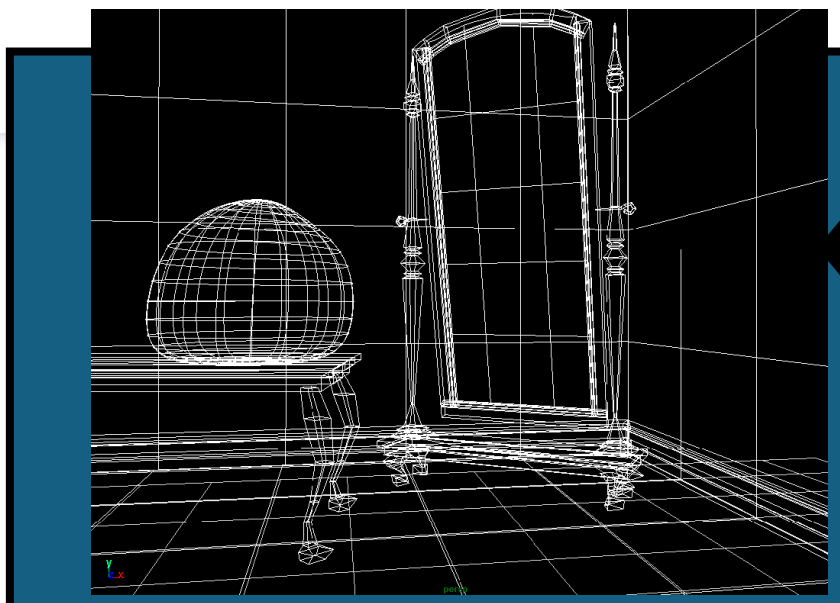

3D geometry

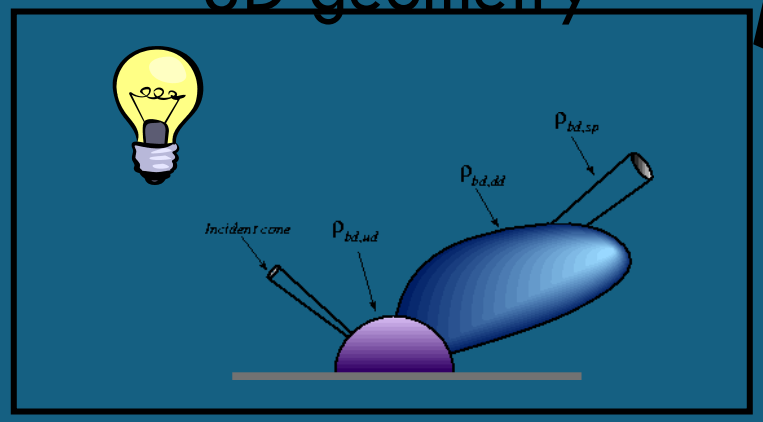

physics

Estimation

What is computer graphics? (3D->2D)

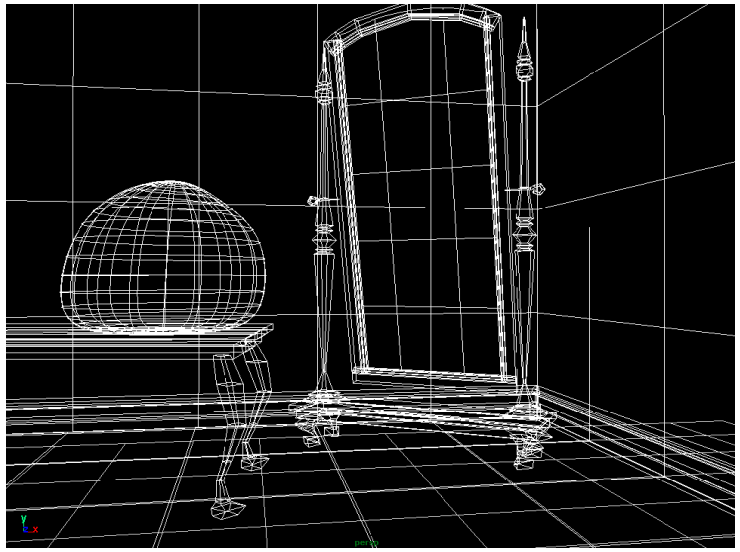

3D geometry

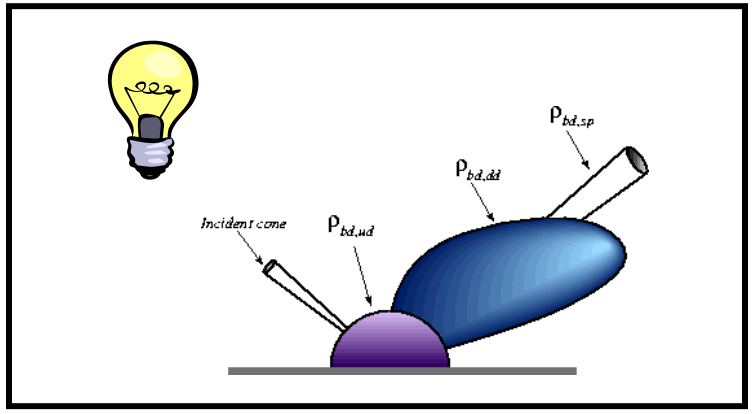

physics

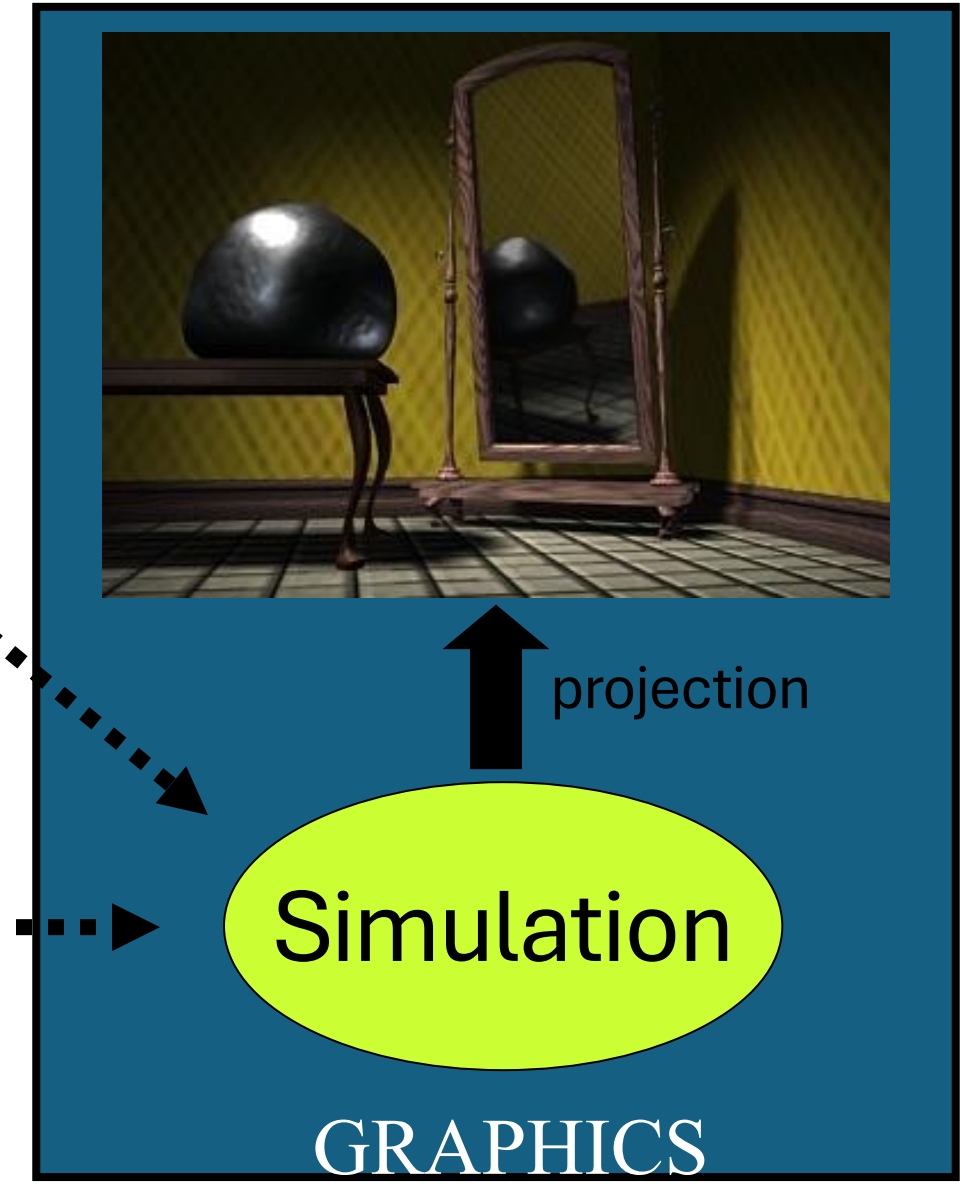

Computer Graphics

Projection, shading, lighting models

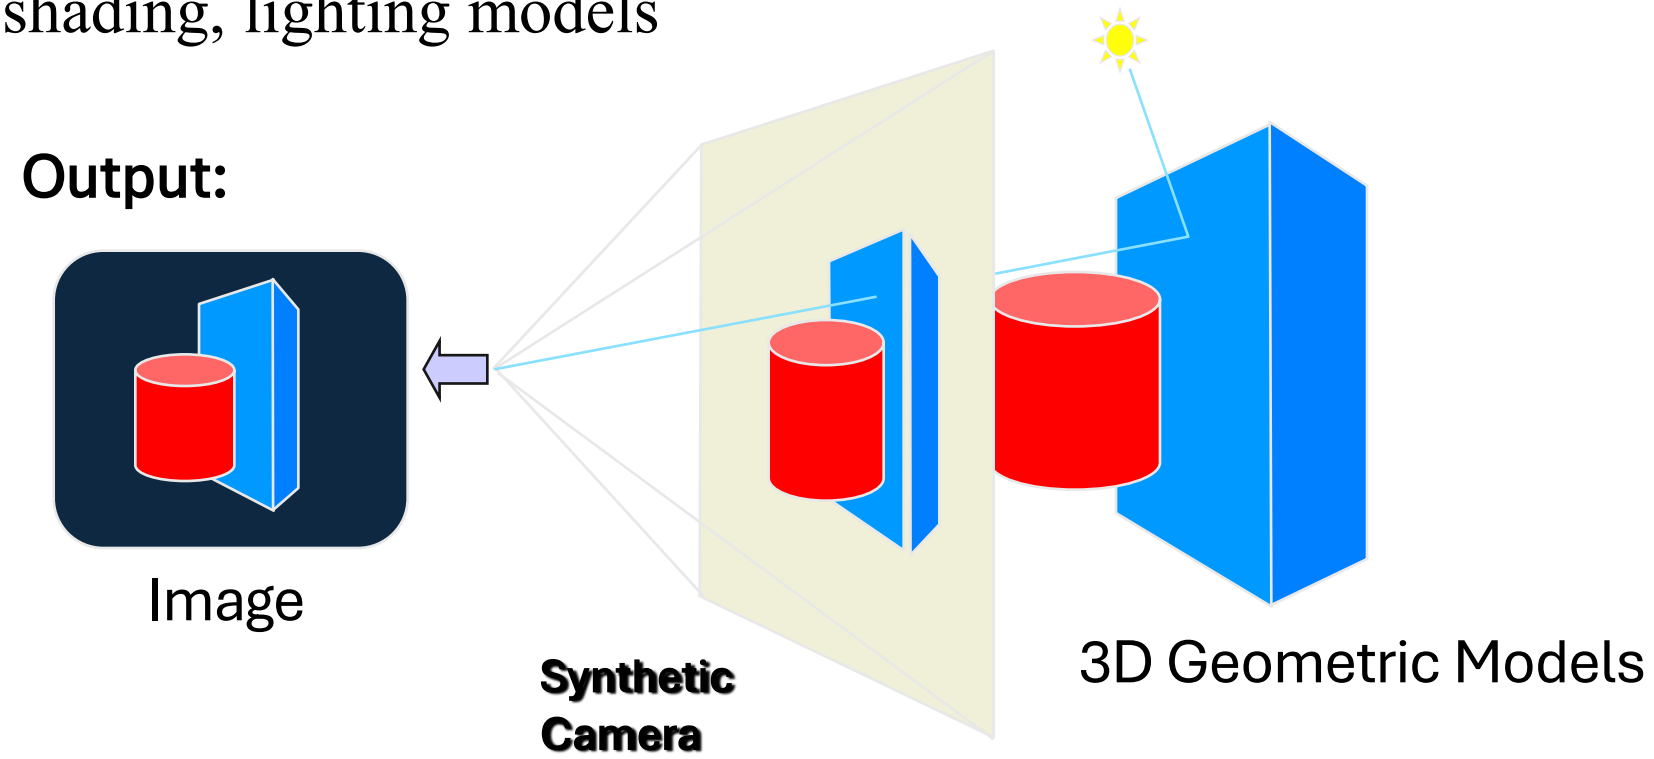

Why vision matters?

- Images and video are everywhere!

Personal photo albums

Movies, news, sports

Google Image Search™ Picasa™ flickr™ webshots picsearch™ YouTube Broadcast Yourself™

Surveillance and security

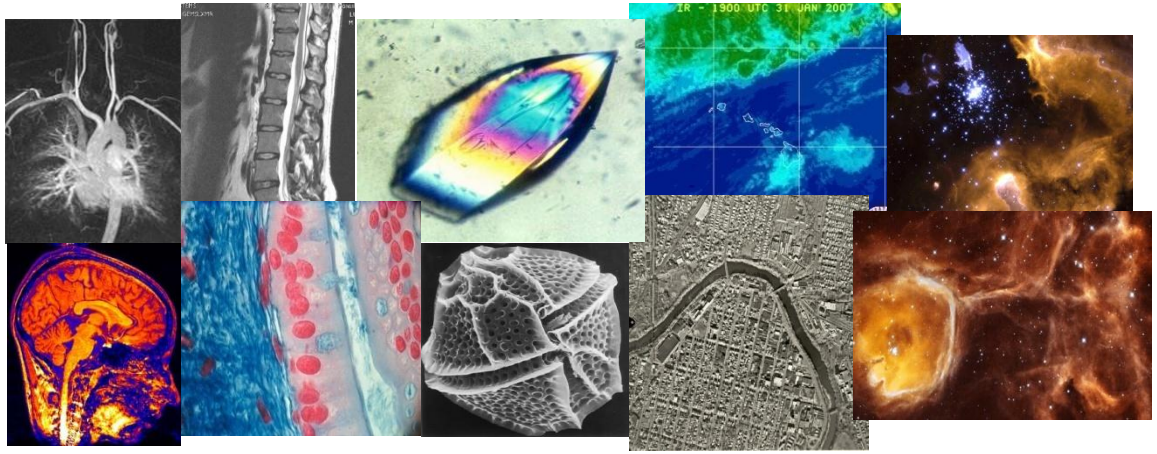

Medical and scientific images

Slide credit; L. Lazebnik

Image Processing and Computer Vision Applications

- Visual inspection/quality control
- Surveillance and security
- Autonomous vehicles
- Space applications
- Medical imaging
- Digital photography and 3D modeling
- Games and much more ……

Image Processing and Computer Vision Applications

Safety

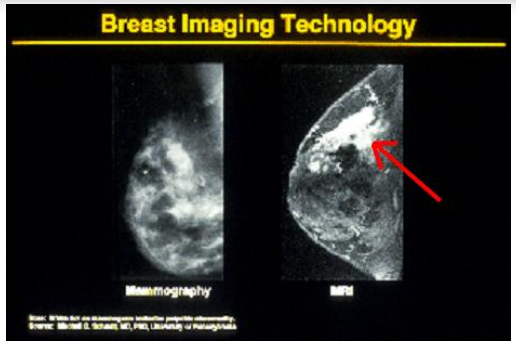

Health

Security

Comfort

Fun

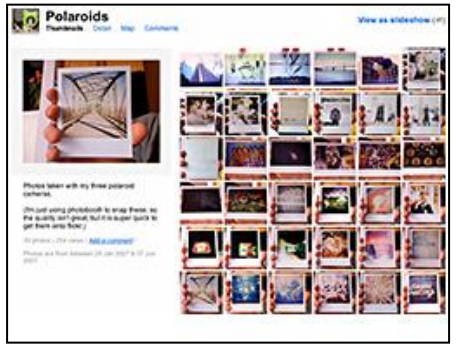

Access

Industry and Applications

- Automobile driver assistance
  - Lane departure warning
  - Adaptive cruise control
  - Obstacle warning
- Digital Photography
  - Image Enhancement
  - Compression
  - Color manipulation
  - Image editing
  - Digital cameras
- Sports analysis
  - sports refereeing and commentary
  - 3D visualization and tracking sports actions

MobilEye system

Industry and Applications

- Film and Video
  - Editing
  - Special effects
- Image Database
  - Content based image retrieval
  - visual search of products
  - Face recognition
- Industrial Automation and Inspection
  - vision-guided robotics
  - Inspection systems
- Medical and Biomedical
  - Surgical assistance
  - Sensor fusion
  - Vision based diagnosis
- Astronomy
  - Astronomical Image Enhancement
  - Chemical/Spectral Analysis

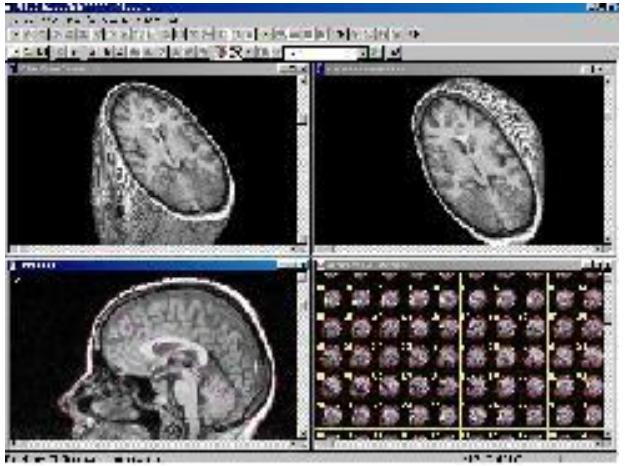

Industry and Applications

- Arial Photography
  - Image Enhancement
  - Missile Guidance
  - Geological Mapping
- Robotics
  - Autonomous Vehicles
- Security and Safety
  - Biometry verification (face, iris)
  - Surveillance (fences, swimming pools)
- Military
  - Tracking and localizing
  - Detection
  - Missile guidance
- Traffic and Road Monitoring
  - Traffic monitoring
  - Adaptive traffic lights

Cruise Missiles

Key Processes in Image Analysis

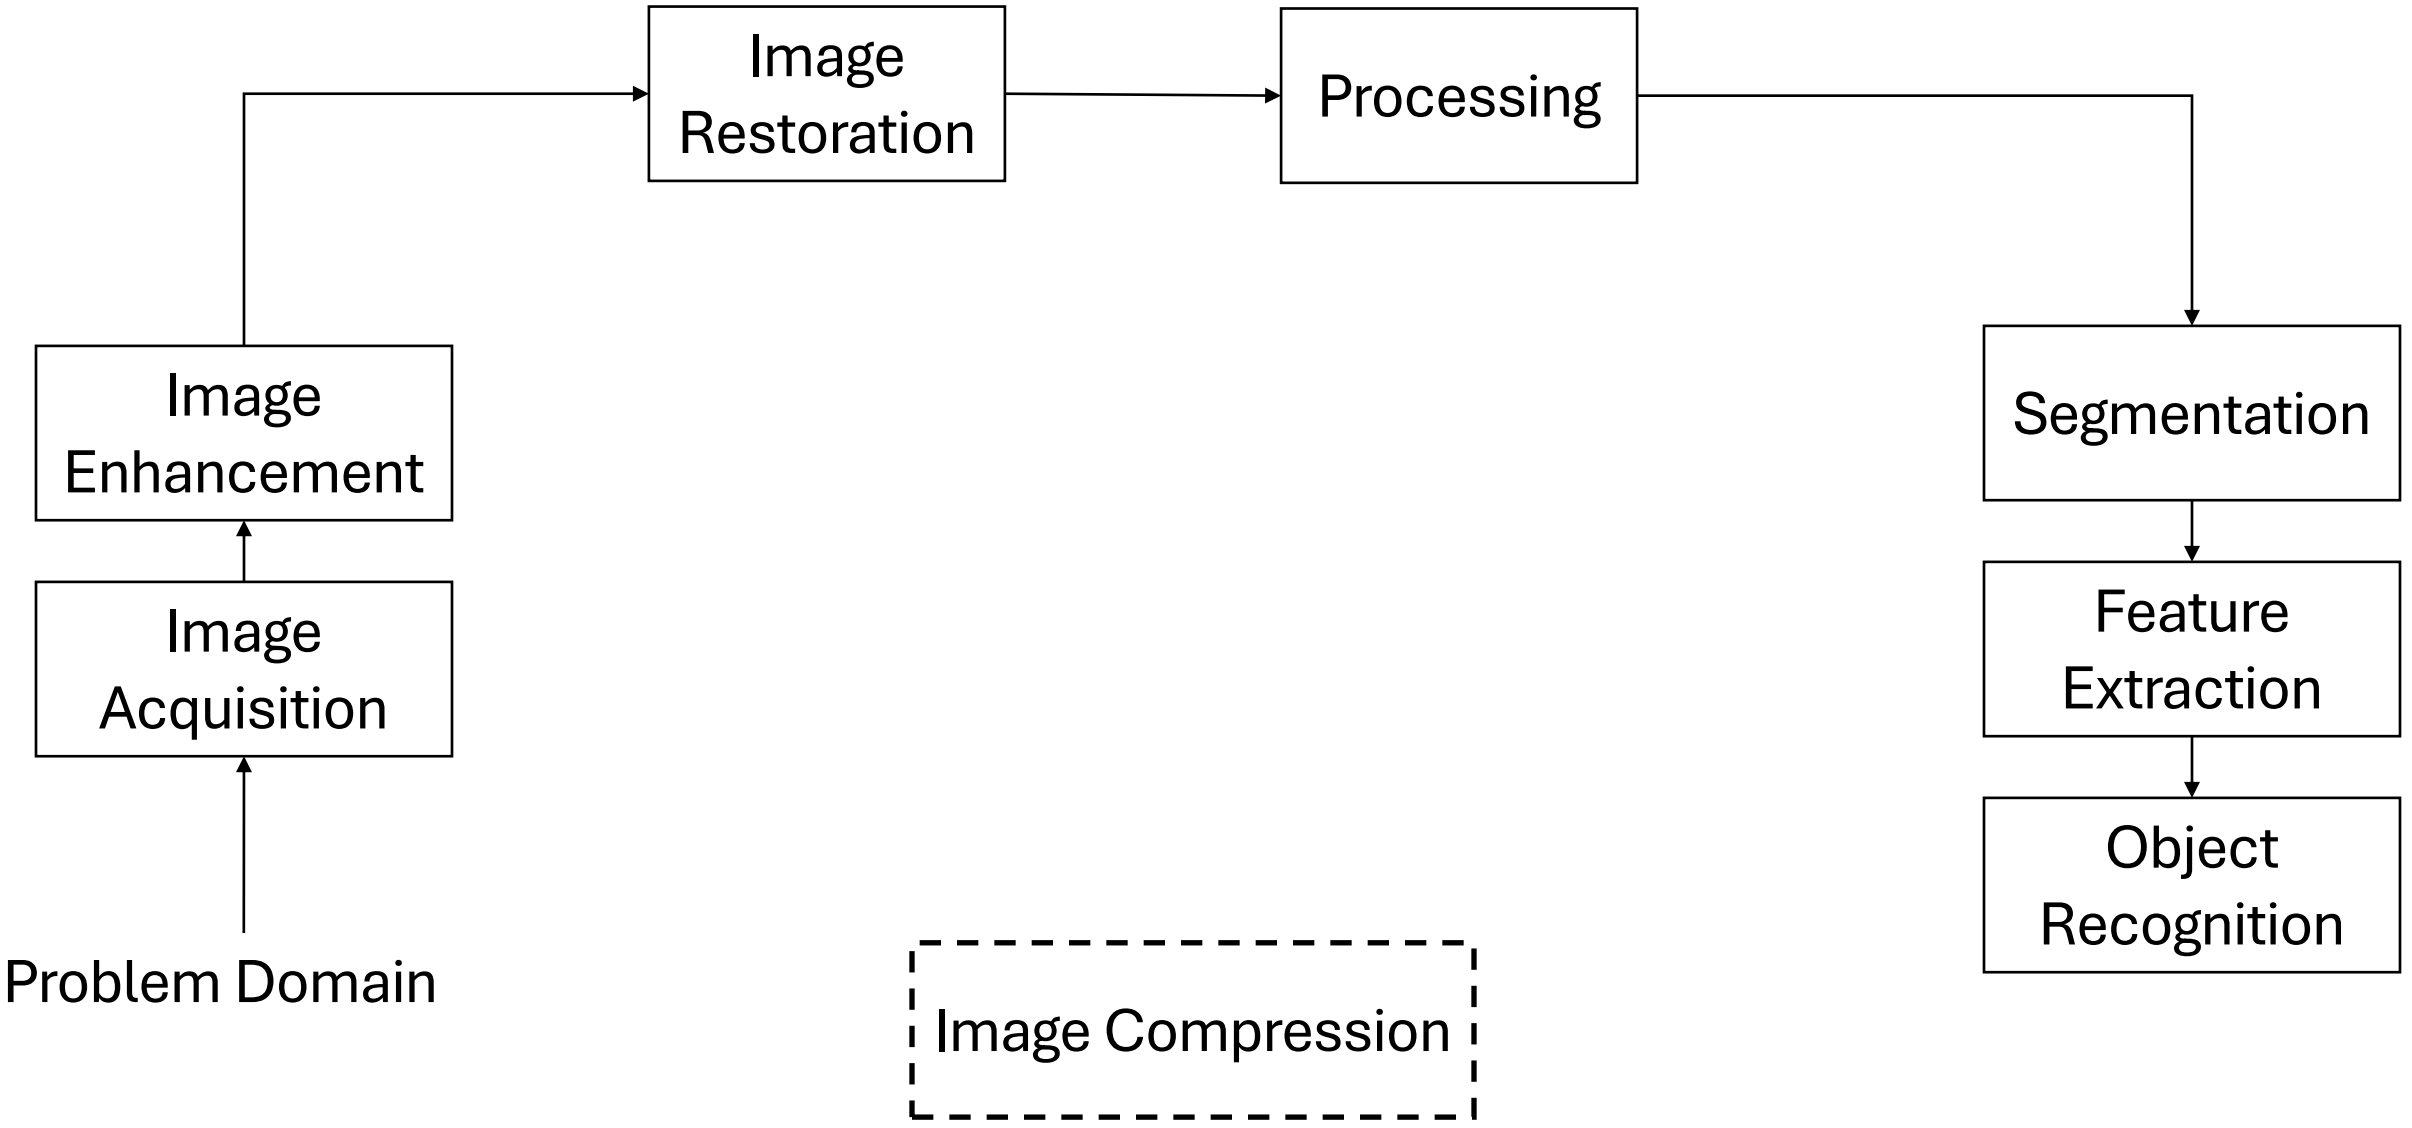

Image Acquisition

Image Restoration

Processing

Image Enhancement

Image Acquisition

Problem Domain

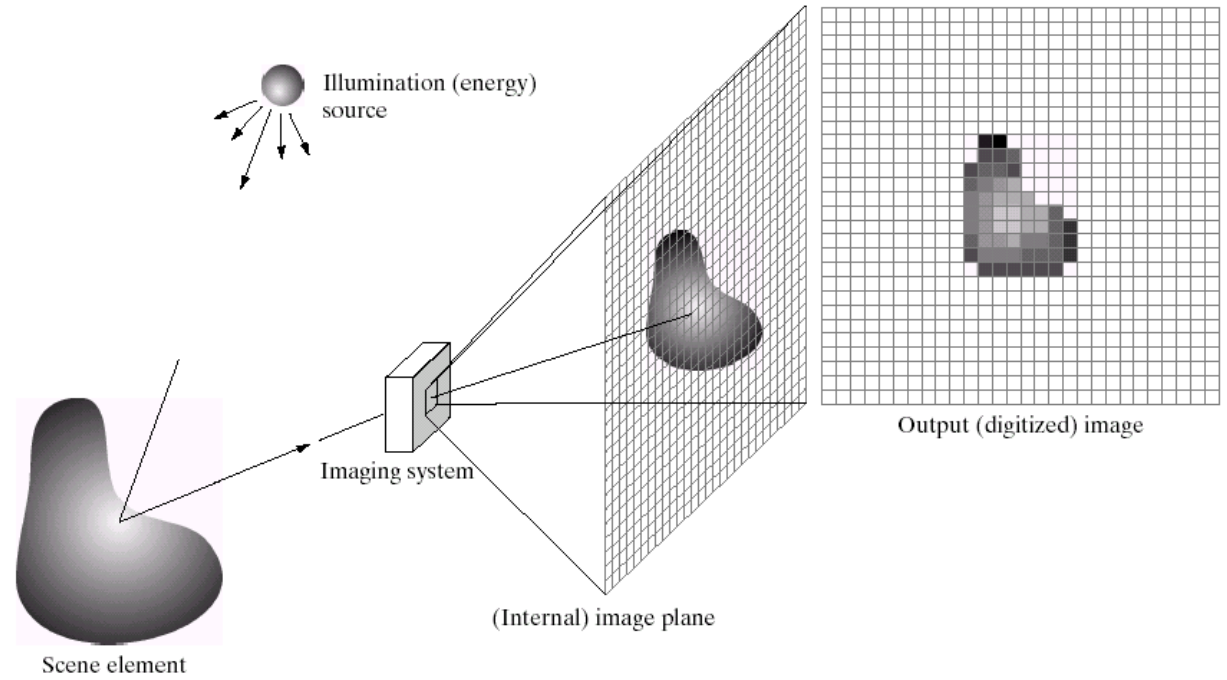

Segmentation

Feature Extraction

Object Recognition

Image Compression

Image Enhancement

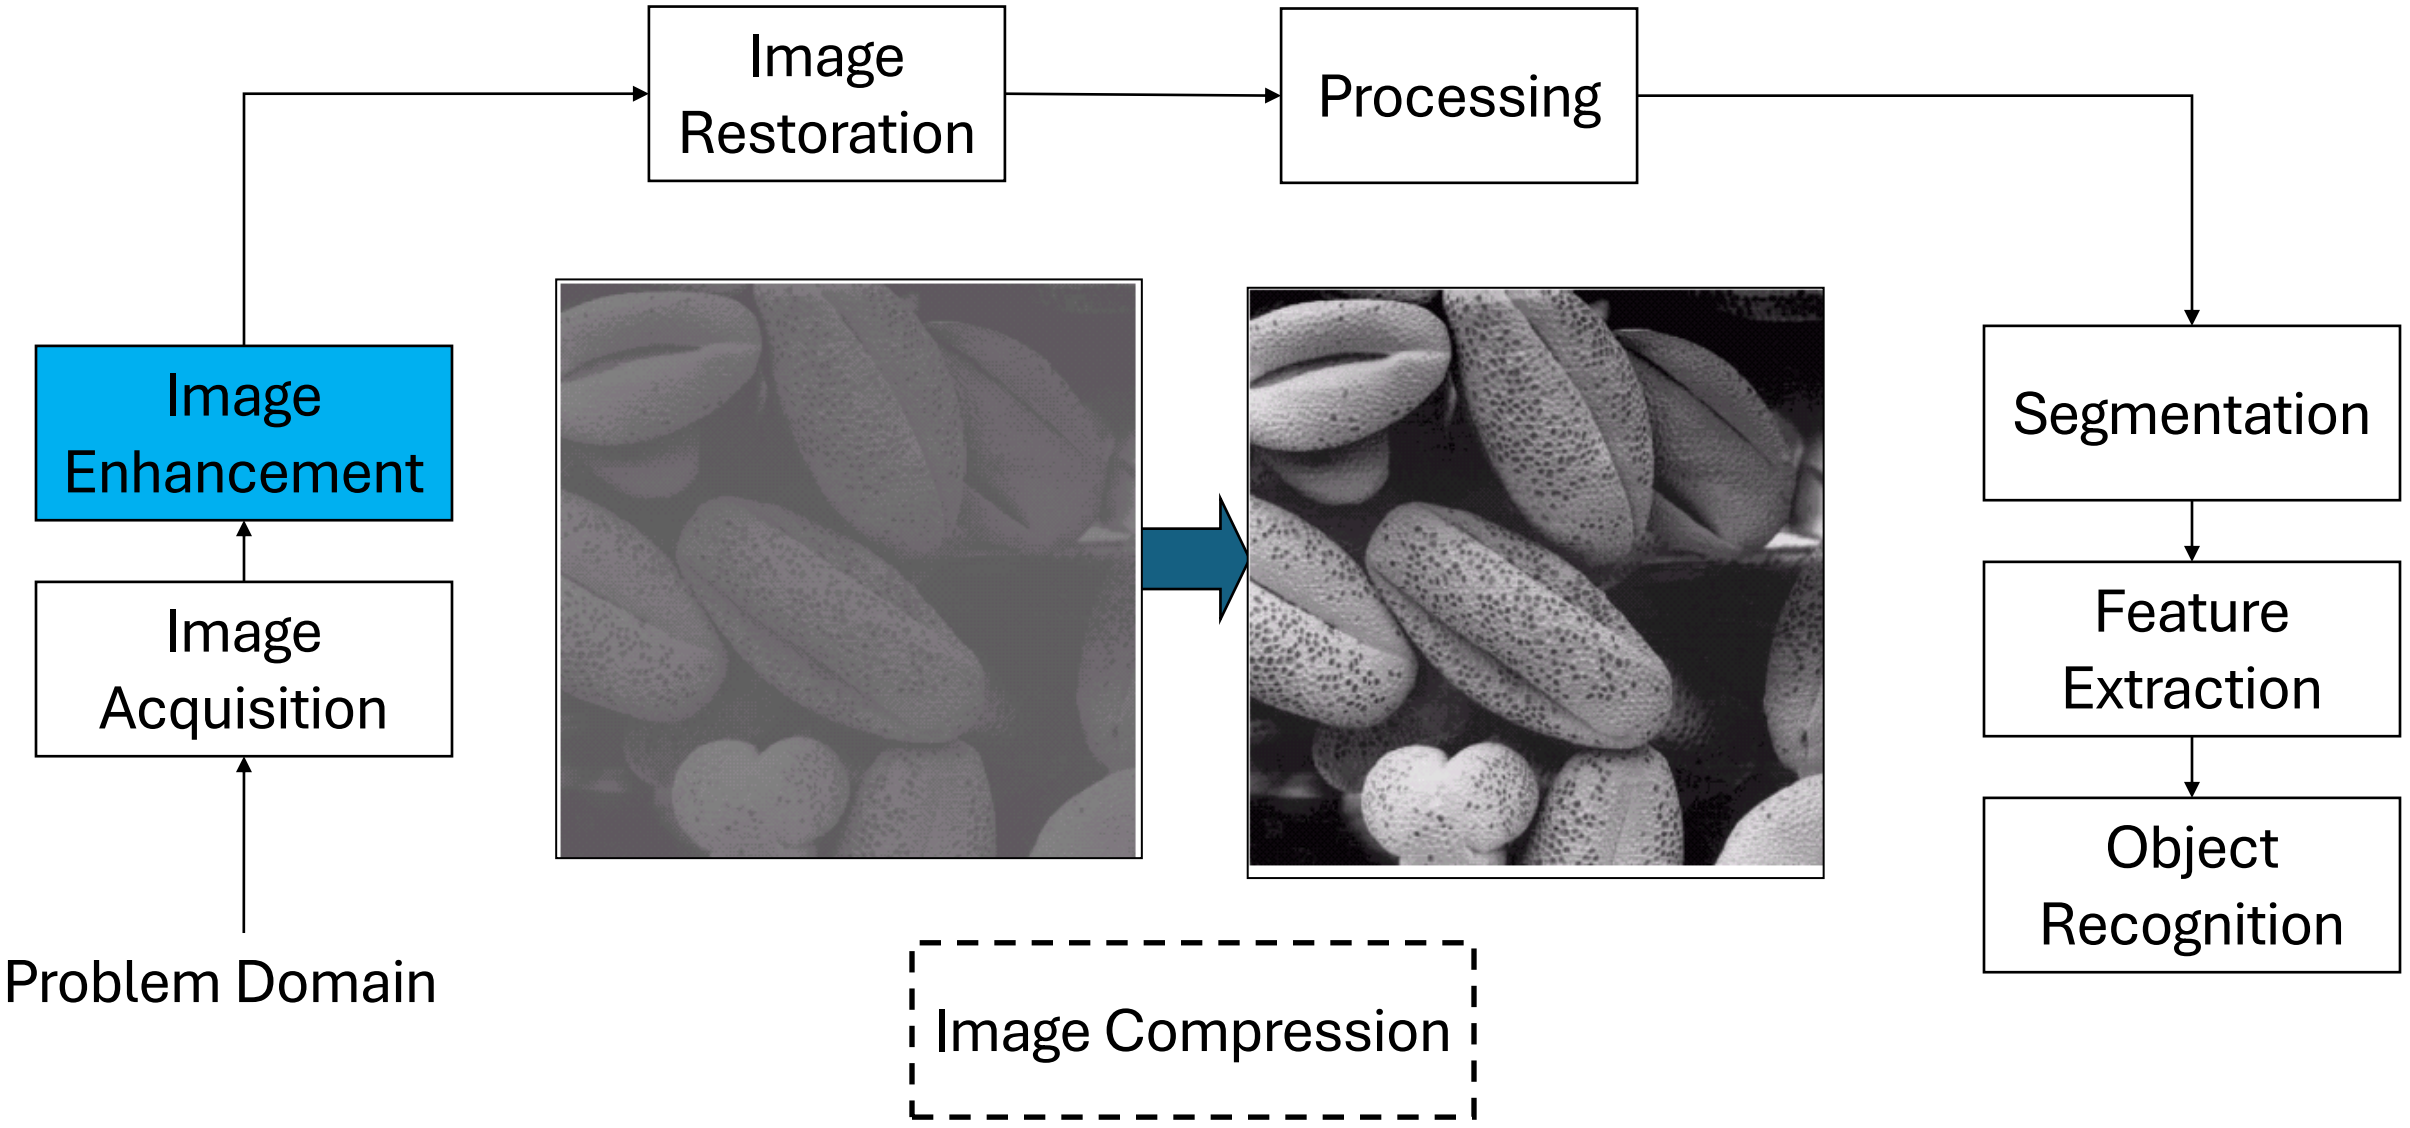

Image Restoration

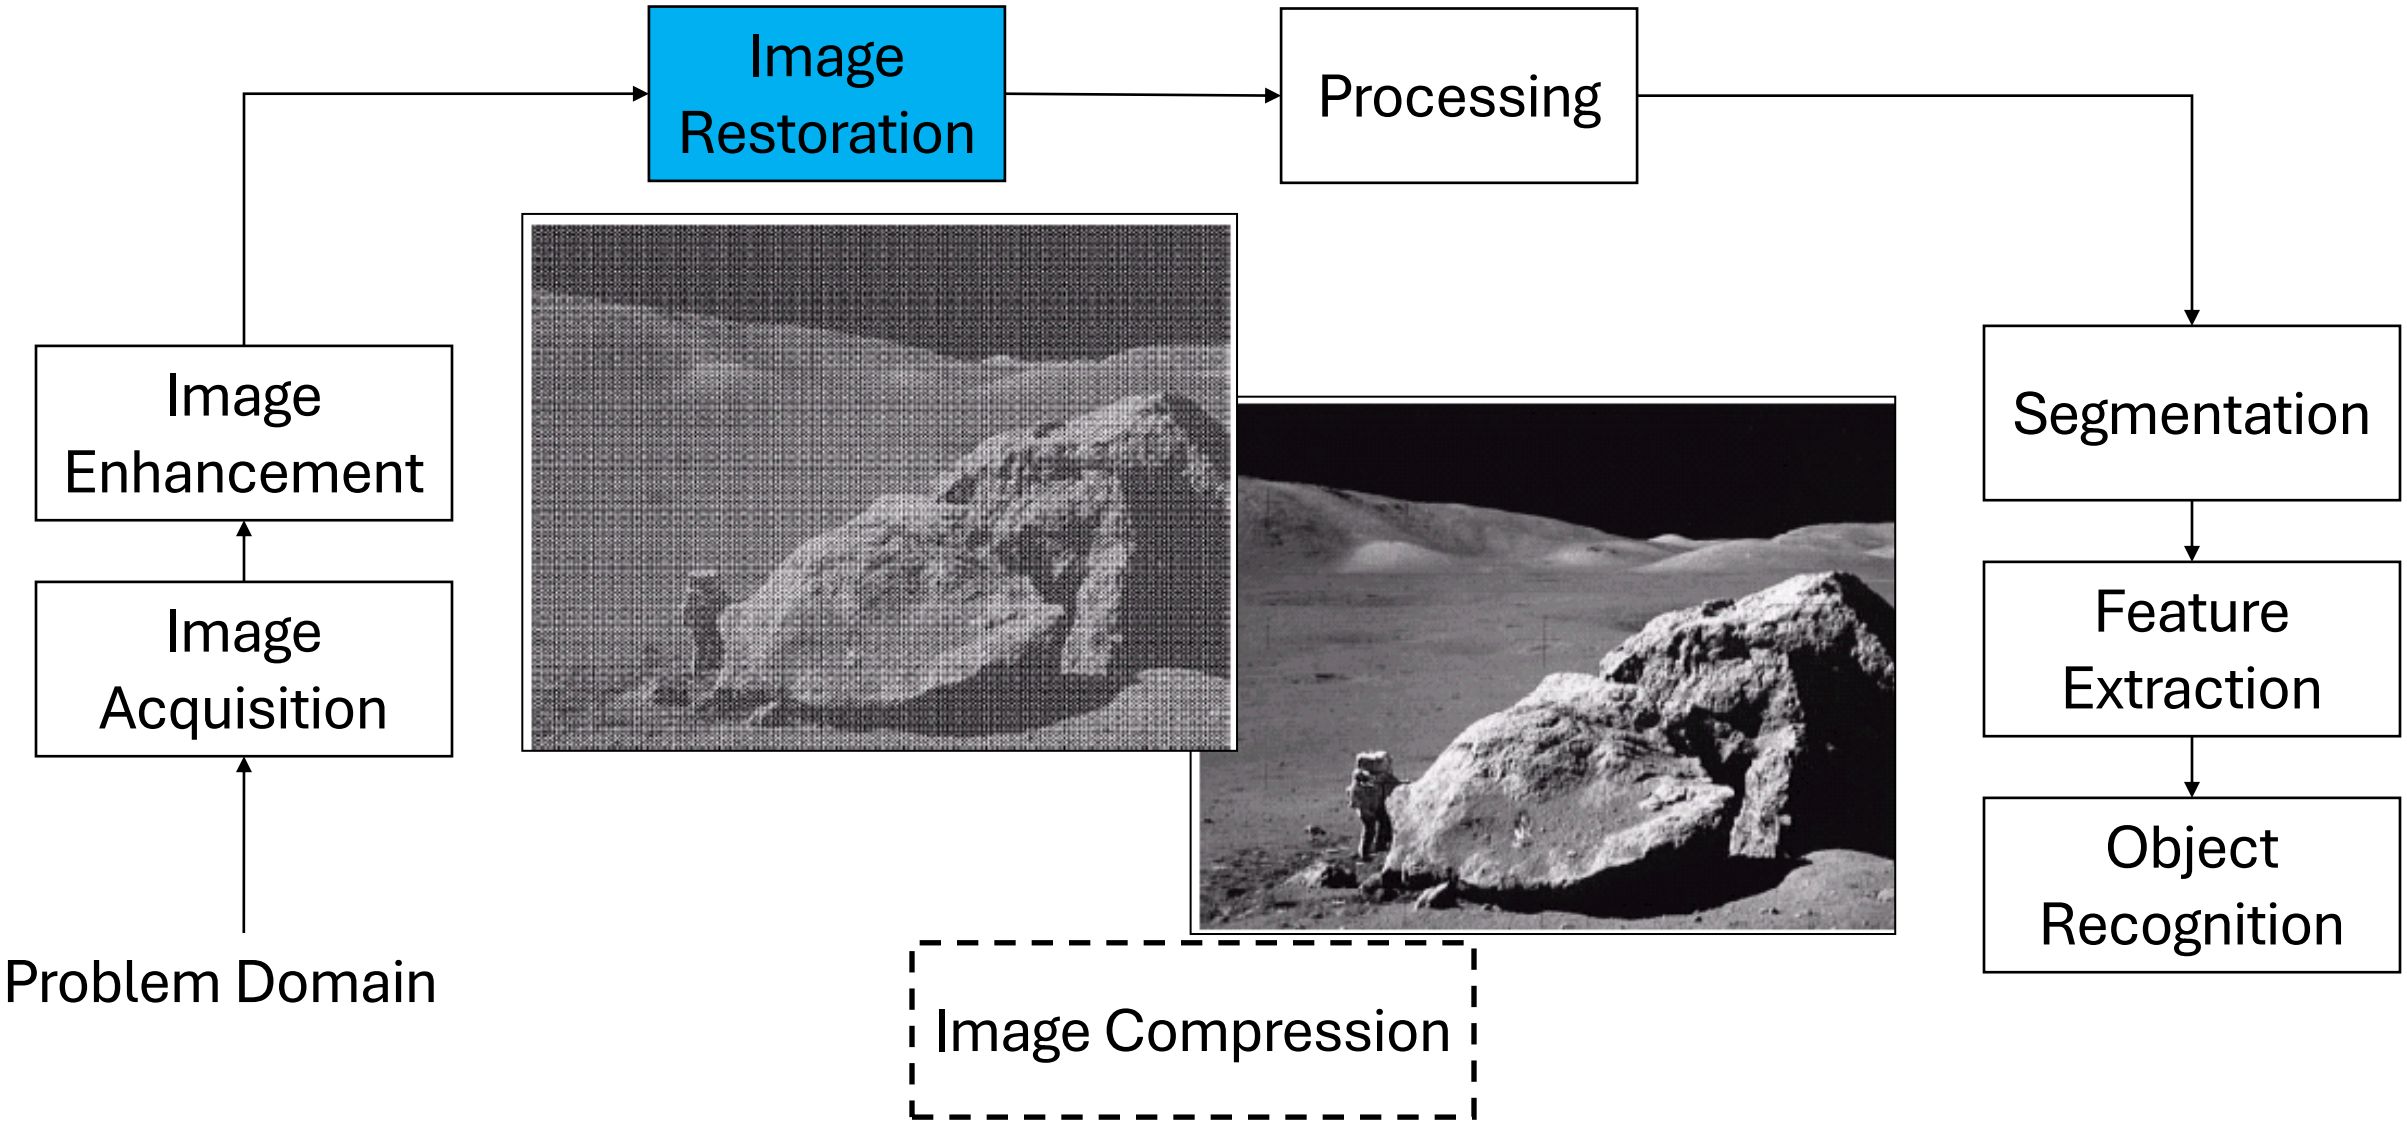

Processing

Problem Domain

Image Acquisition

Image Enhancement

Image Restoration

Processing

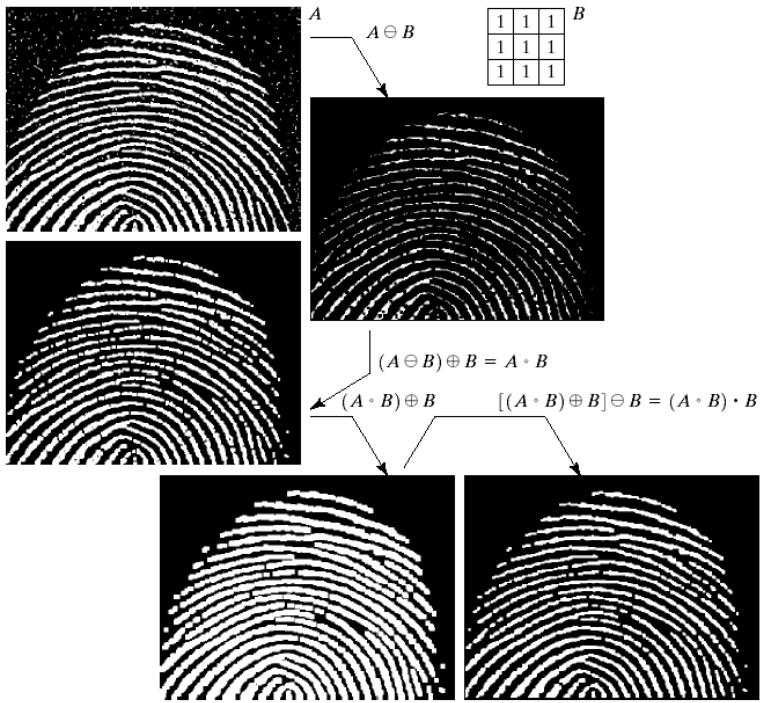

Image Compression

Segmentation

Feature Extraction

Object Recognition

Segmentation

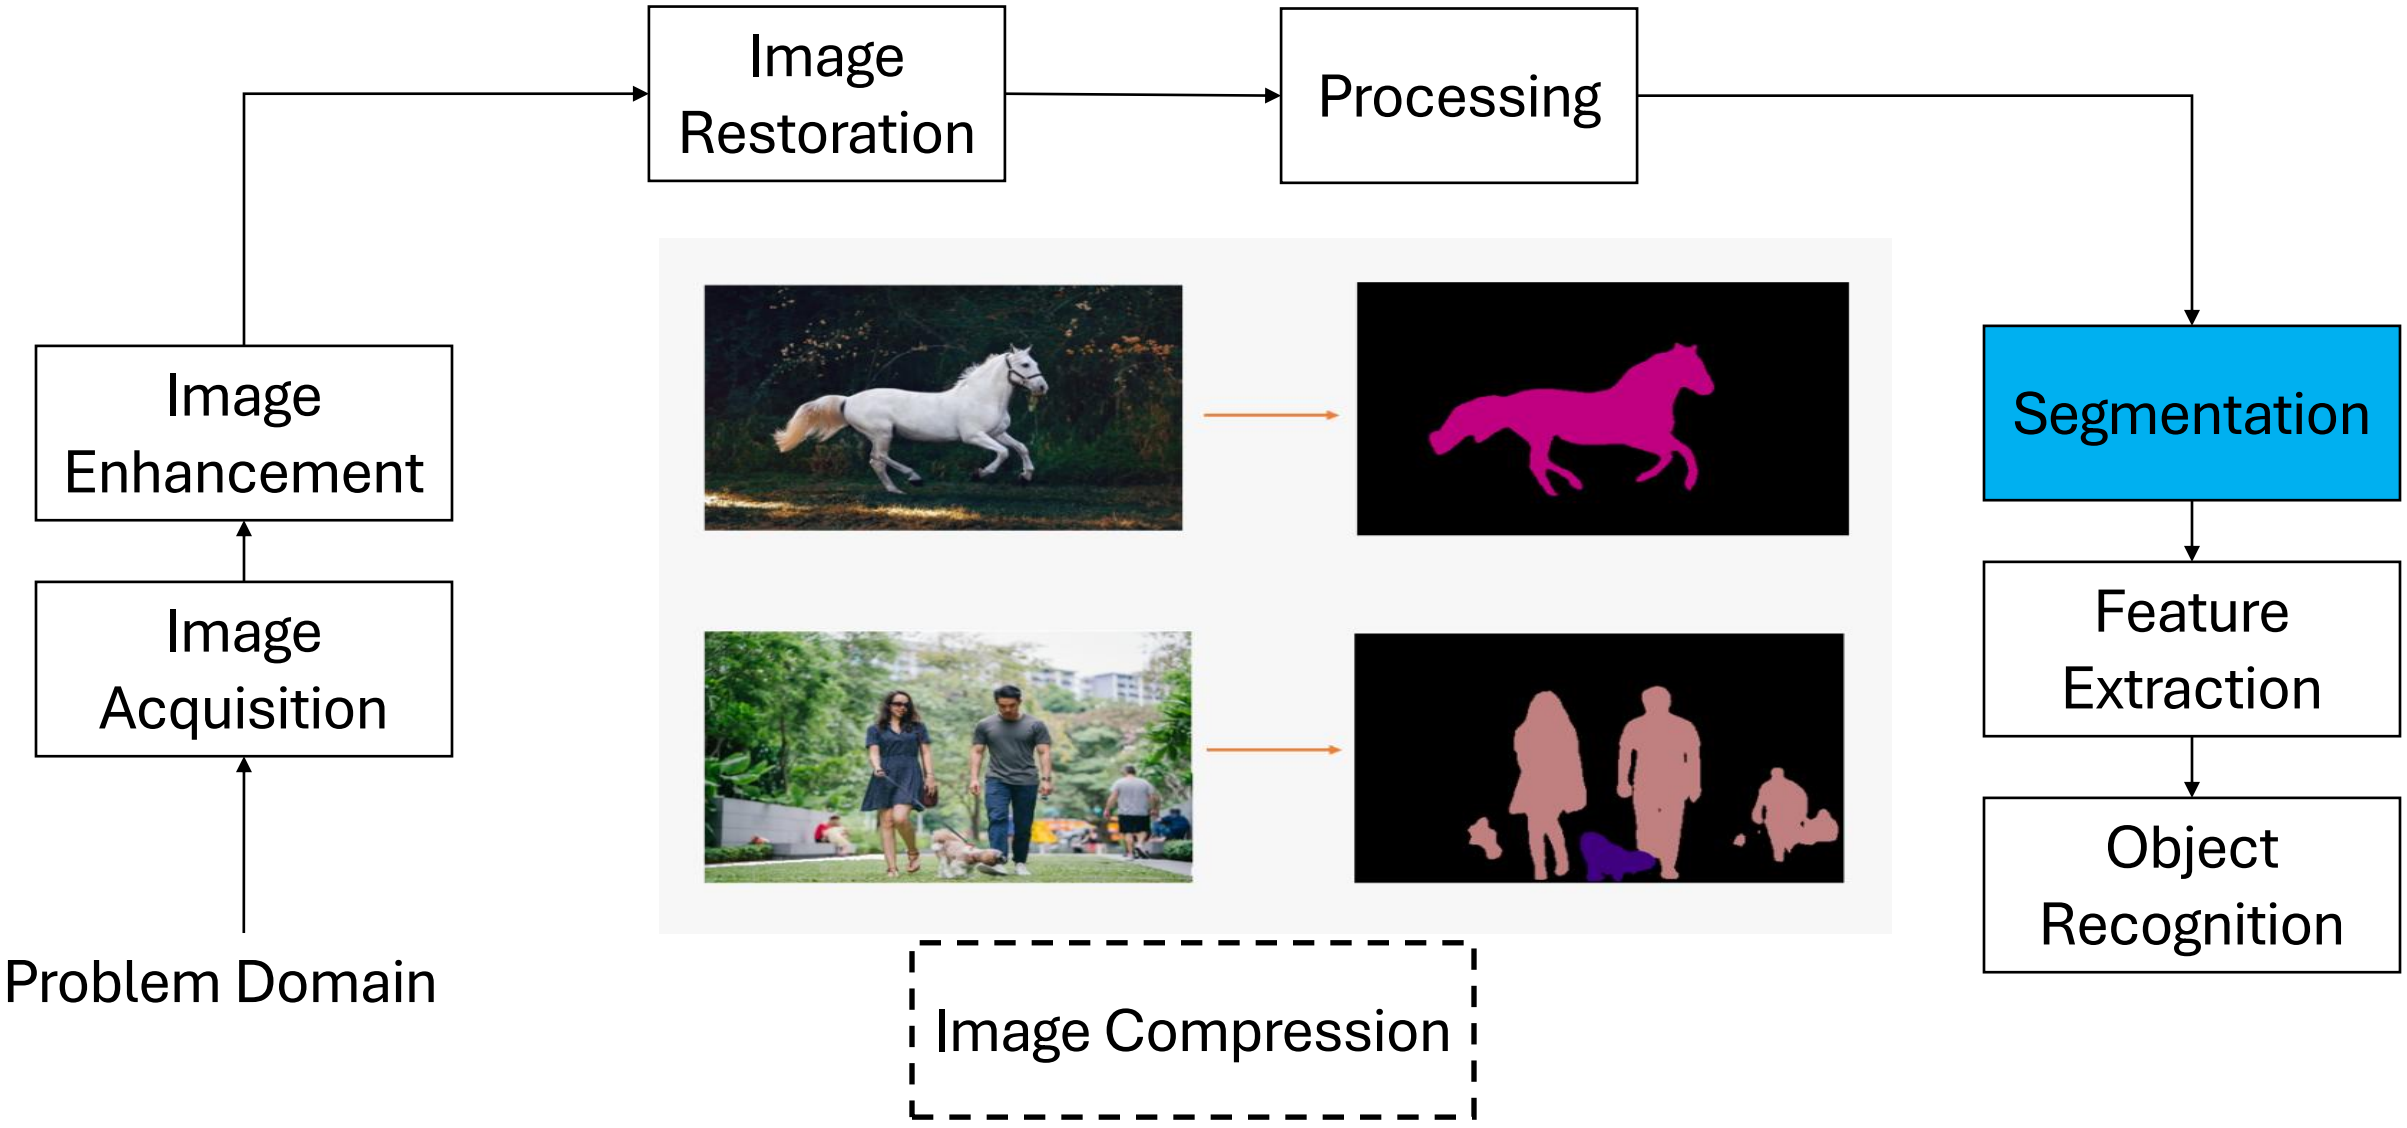

Representation & Description

Image Restoration

Processing

Segmentation

Feature Extraction

Object Recognition

Image Enhancement

Image Acquisition

Problem Domain

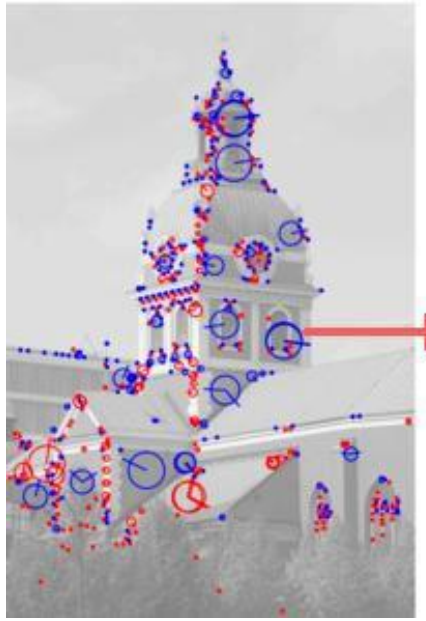

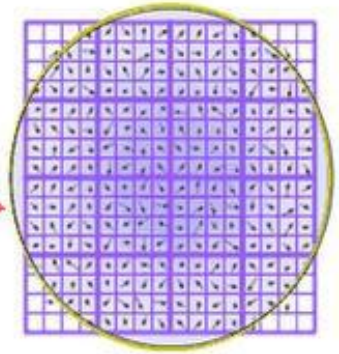

Image gradients

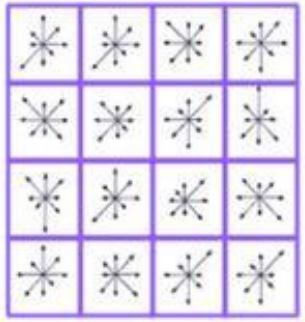

Keypoint descriptor

Image Compression

Object Recognition

Problem Domain

Image Acquisition

Image Enhancement

Image Restoration

Processing

Segmentation

Feature Extraction

Object Recognition

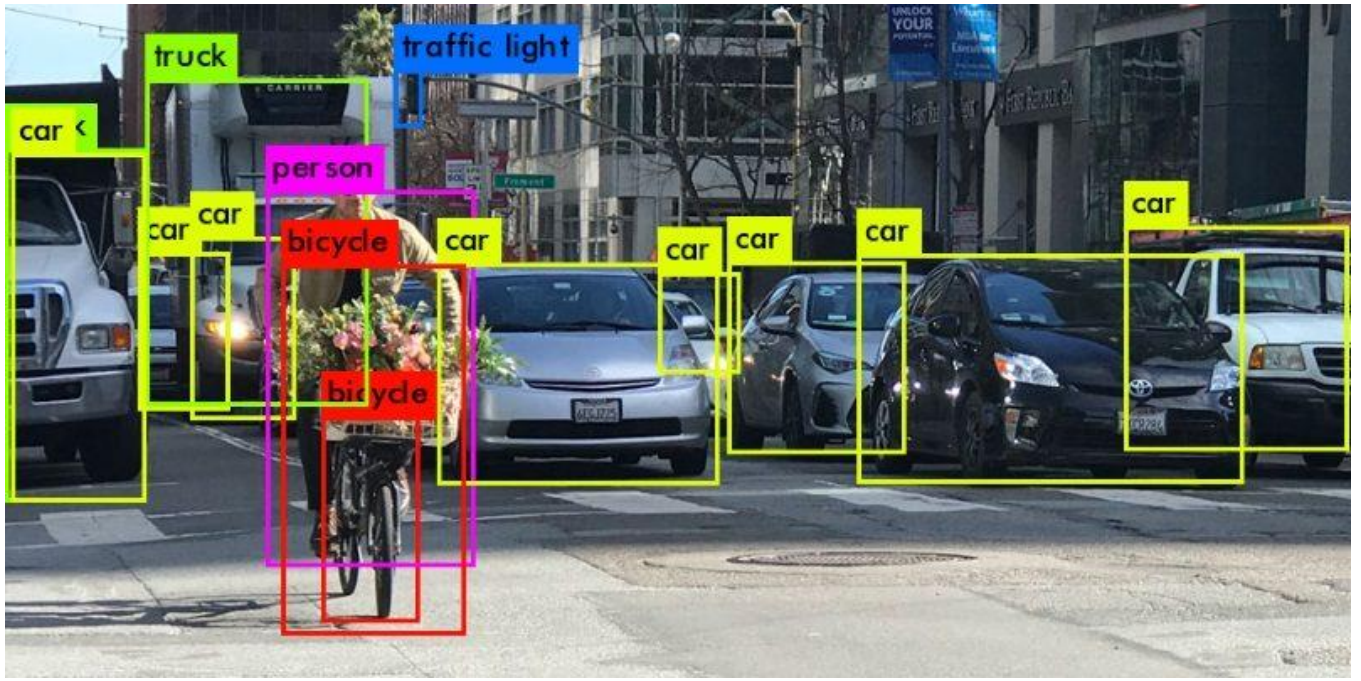

Image Compression

Image Compression

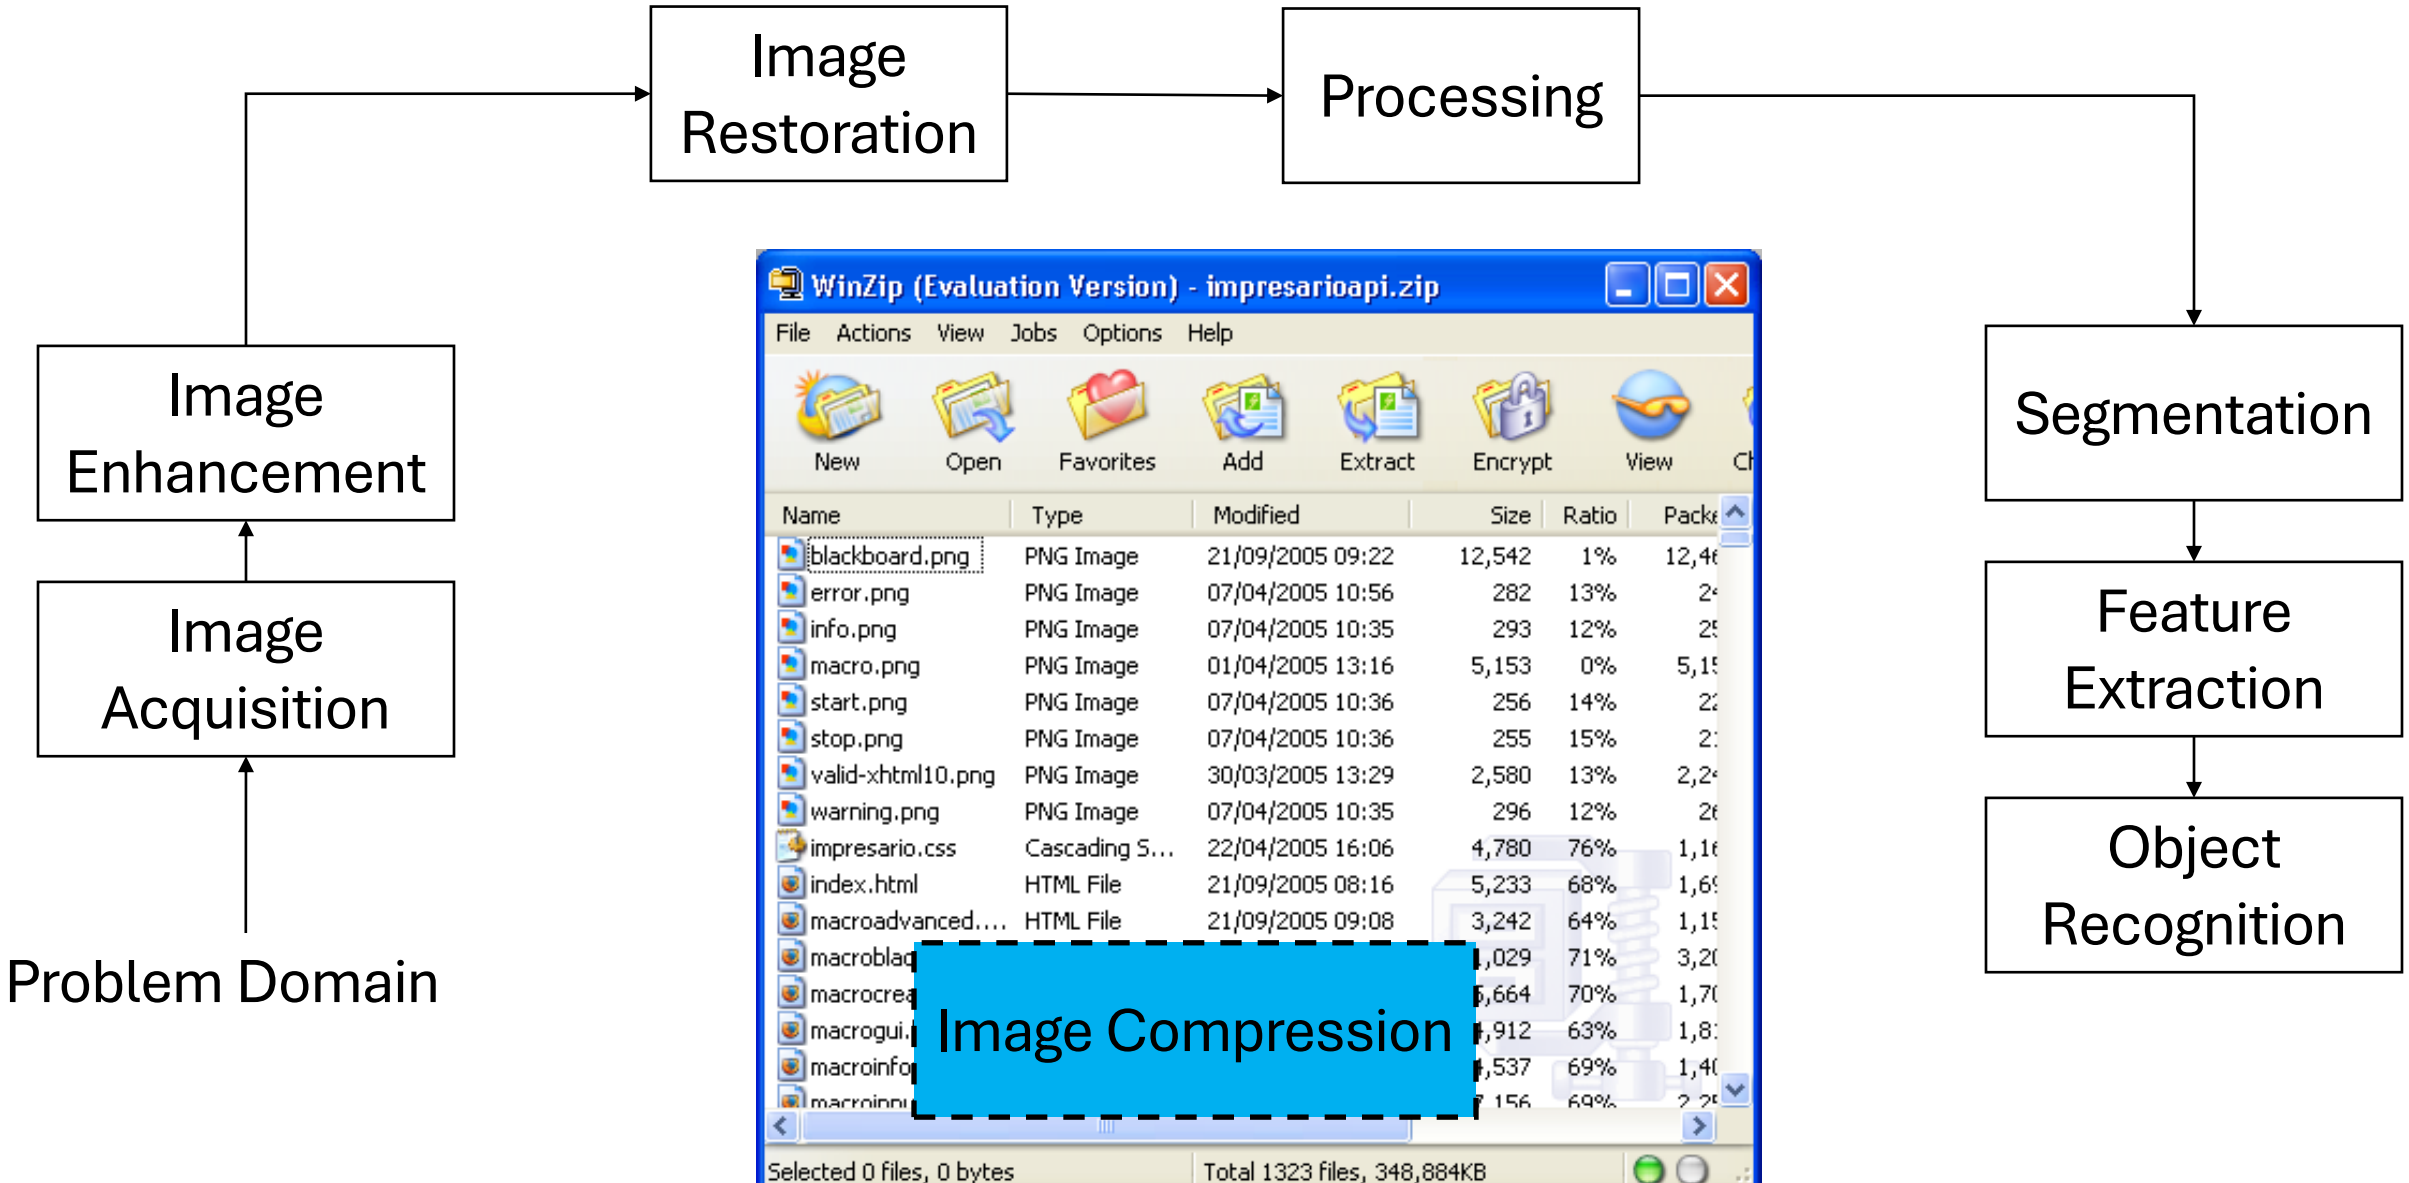

Challenges for Vision Algorithms

viewpoint variation

Michelangelo 1475-1564

Illumination

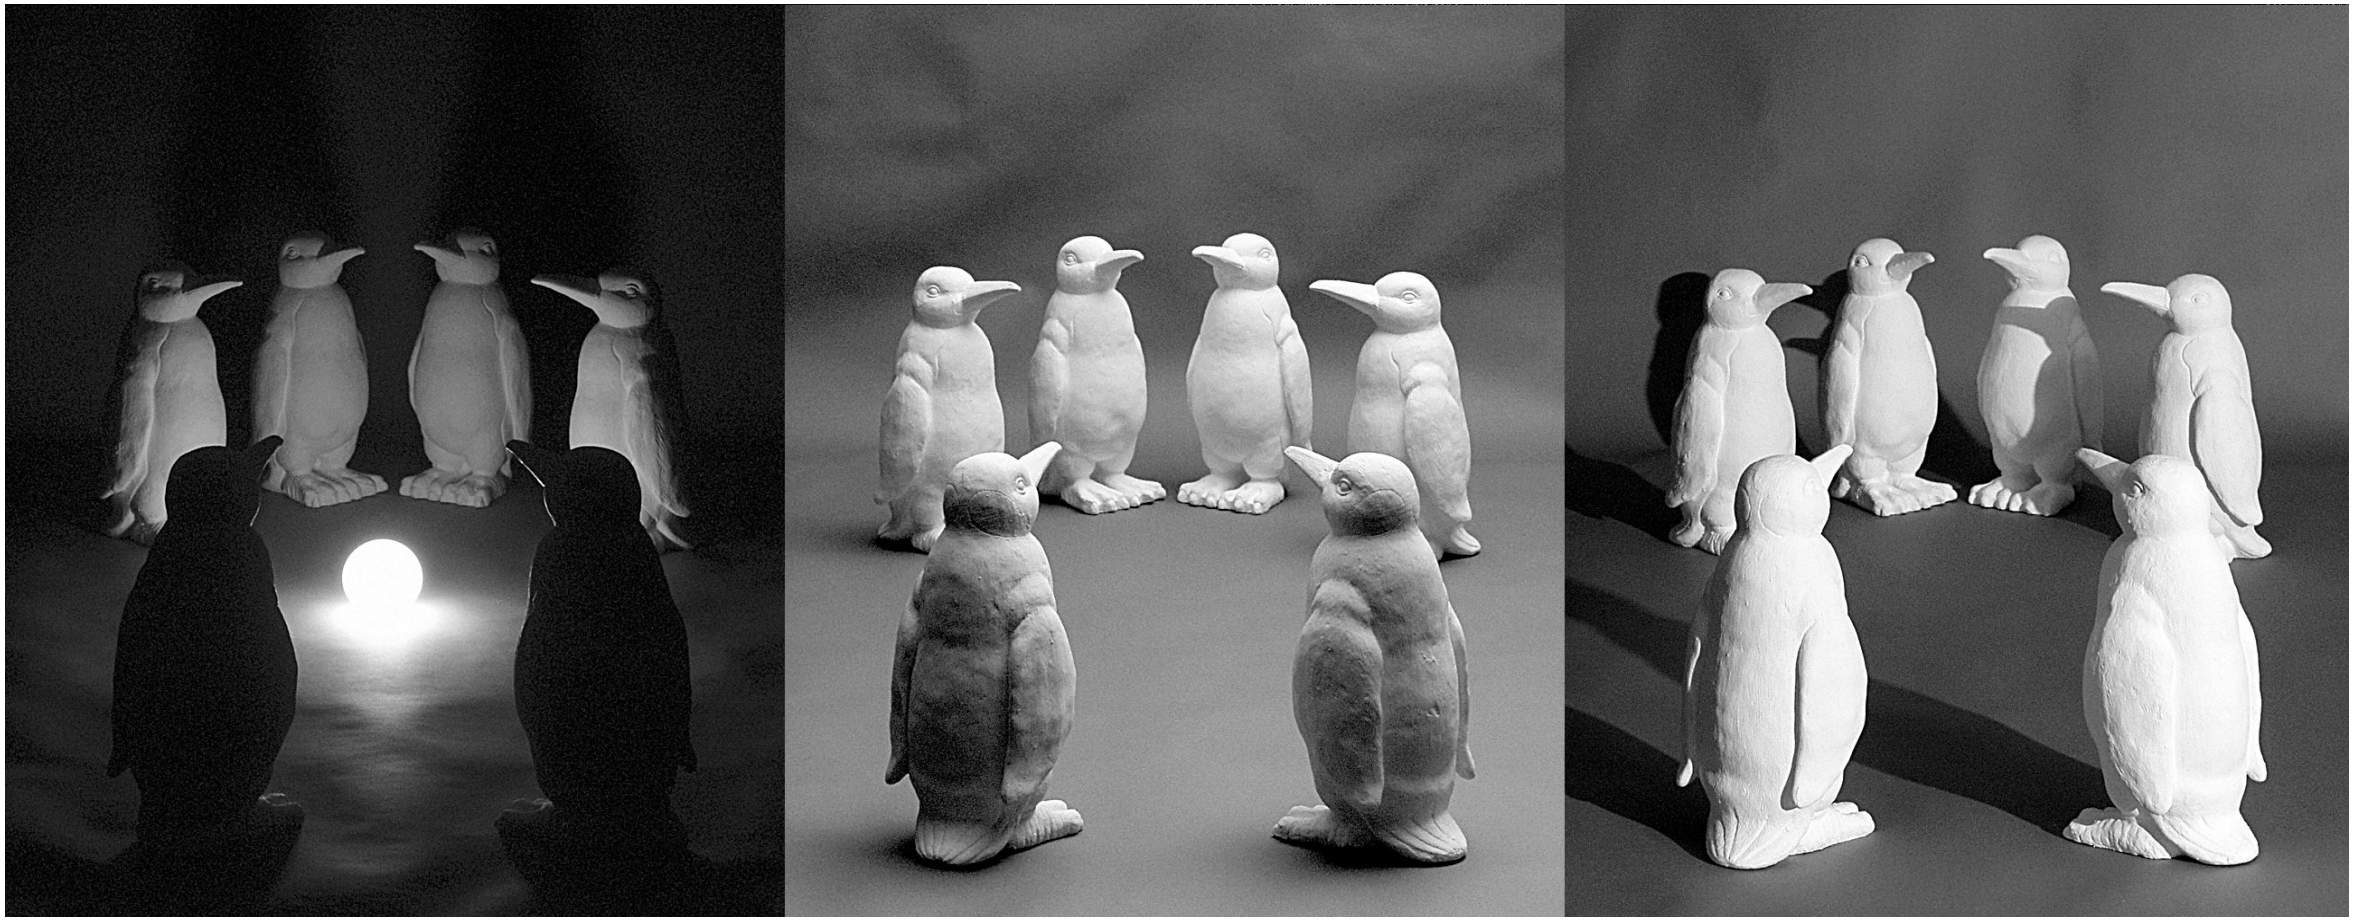

Illumination




Scale

Deformation

Occlusion

slide credit: Fei-Fei, Fergus & Torralba



Background Clutter

Background Clutter

Object intra-class variation

Local ambiguity

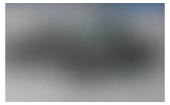

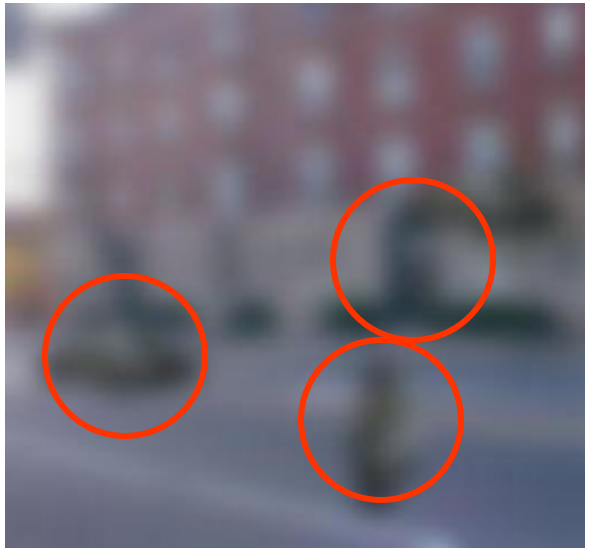

Challenges or opportunities?

- Images are confusing, but they also reveal the structure of the world through numerous cues
- Our job is to interpret the cues!

Bottom line

- Perception is an inherently ambiguous problem
  - Many different 3D scenes could have given rise to a particular 2D picture

- Possible solutions
  - Bring in more constraints (or more images)
  - Use prior knowledge about the structure of the world
- Need both exact measurements and statistical inference!

Some more Applications of IP and CV

Image Enhancement

Contrast stretching

Deblurring

Image Enhancement

Denoising

Inpainting

Industry and Applications

- Face detection
  - Almost all digital cameras now detect faces

Machine vision

Automated visual inspection

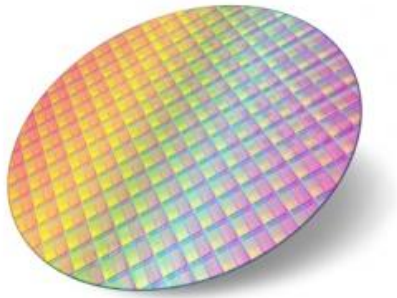

Face detection

- Many new digital cameras now detect face
  - Canon, Sony, Fuji, …

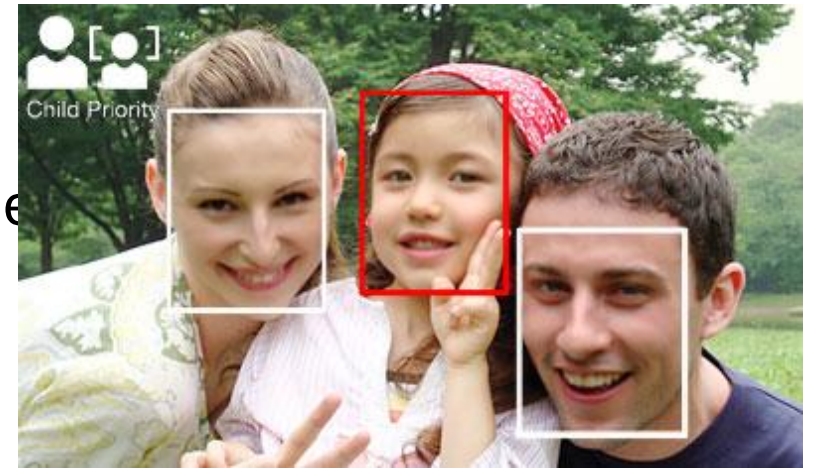

Age recognition

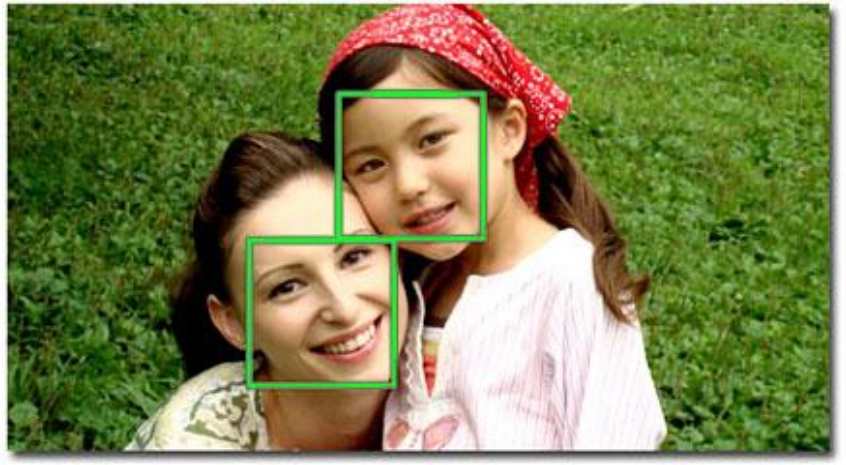

Smile recognition

Smile detection?

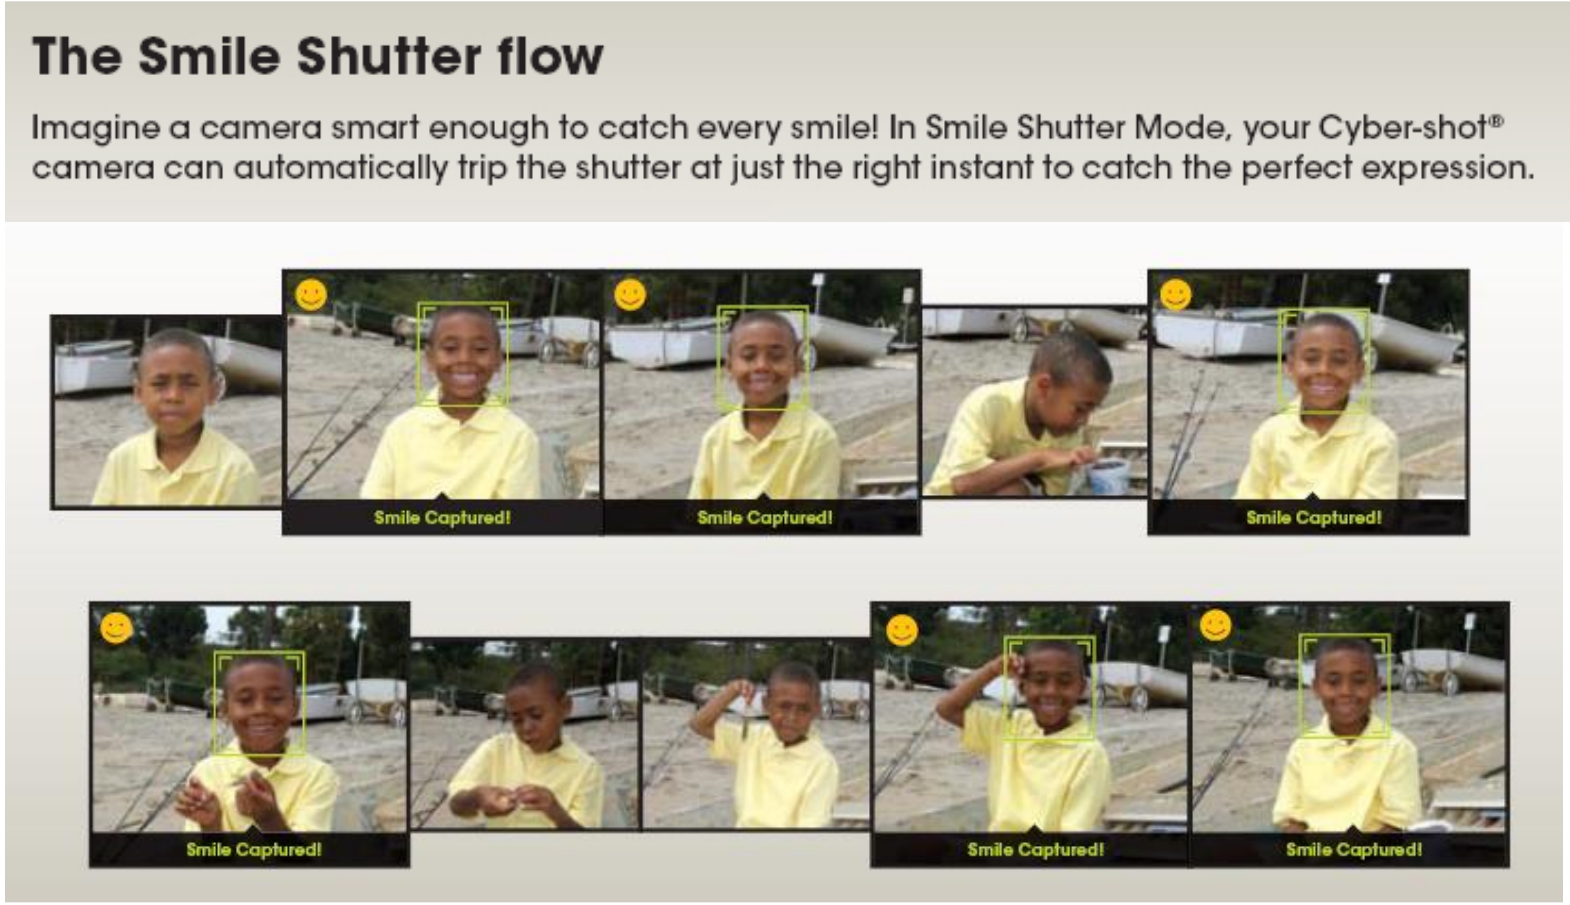

Sony Cyber-shot® T70 Digital Still Camera

Face makeovers

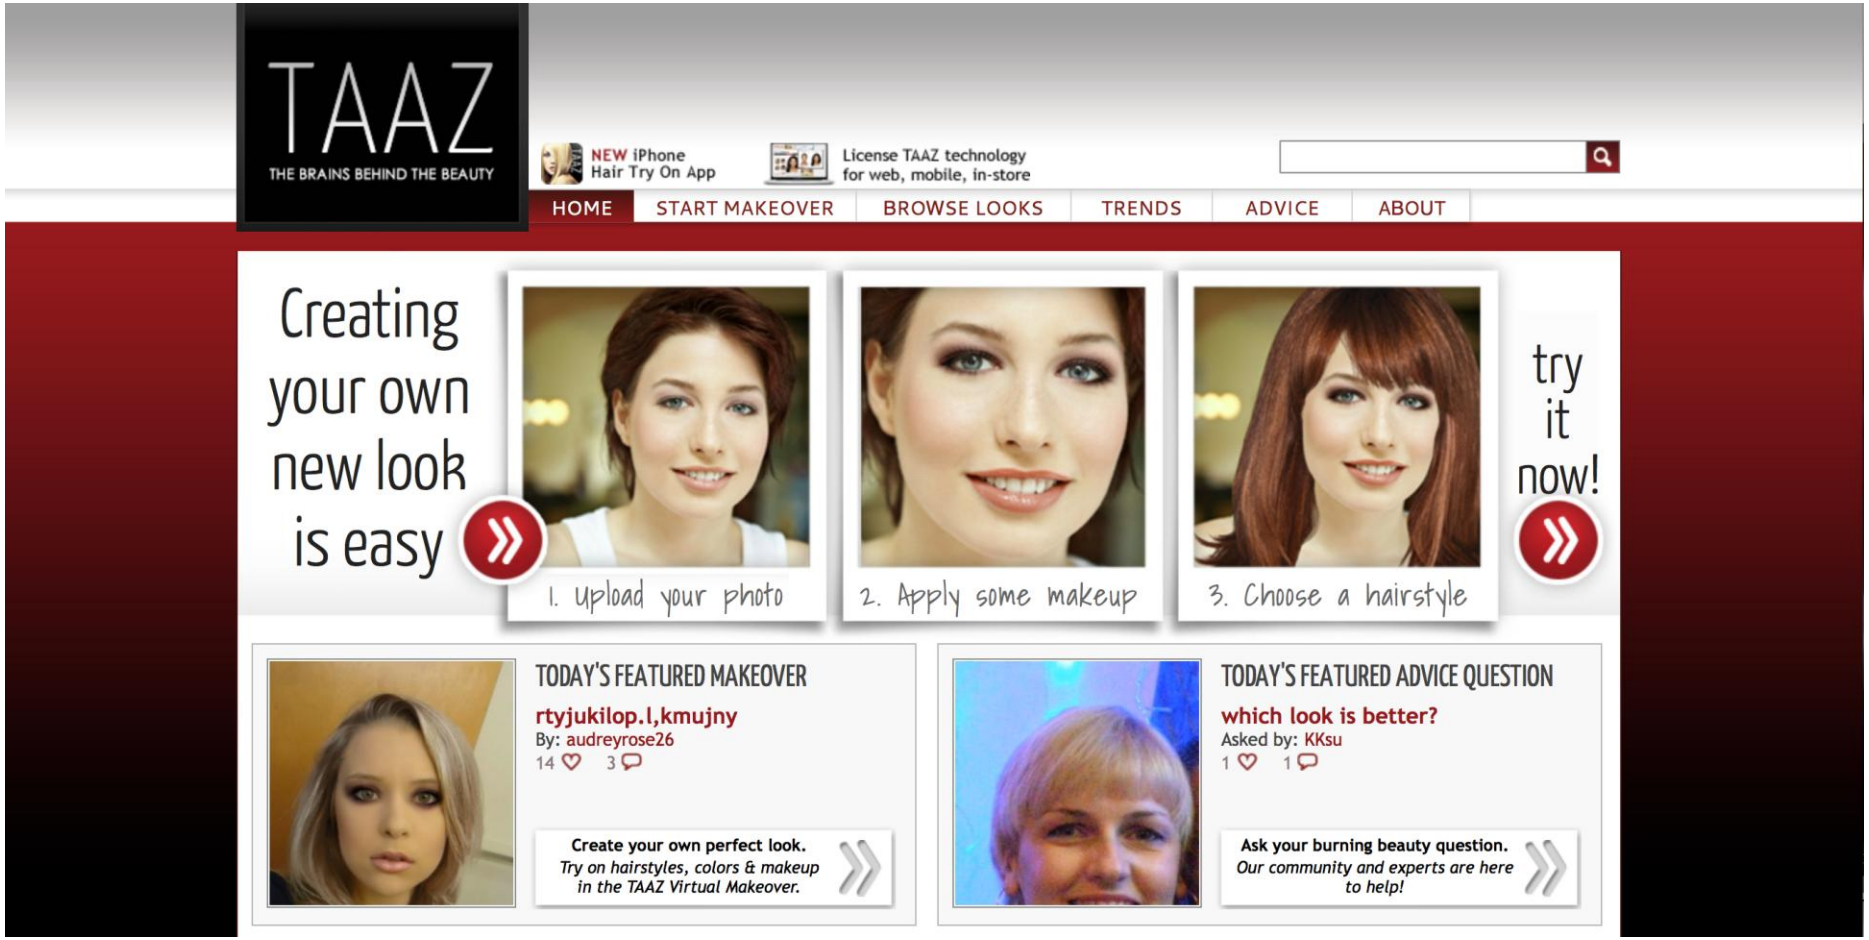

Industry and Applications

- Face recognition

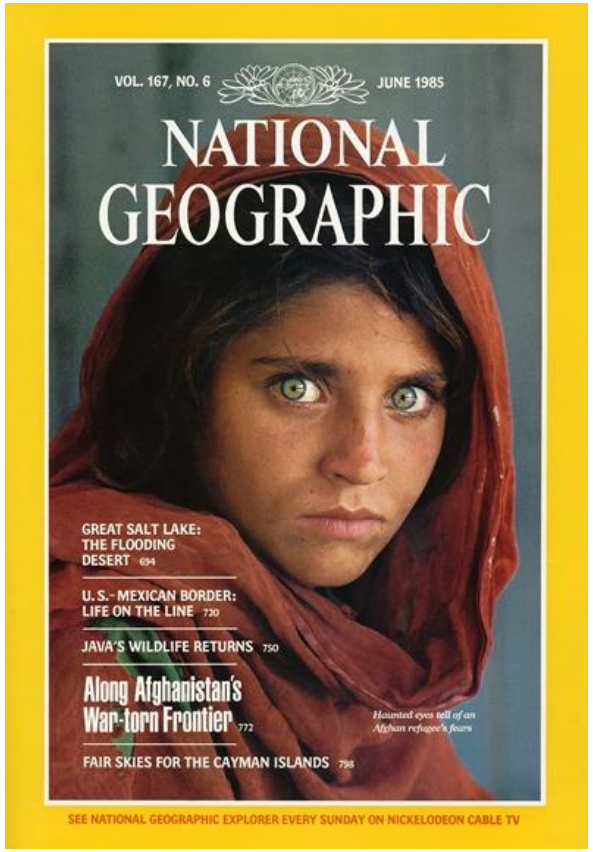

Who is she?

"*How the Afghan Girl was Identified by Her Iris Patterns*" Read the story

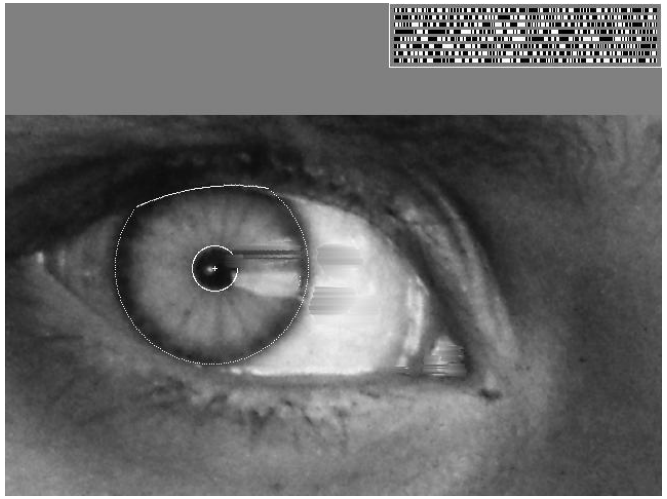

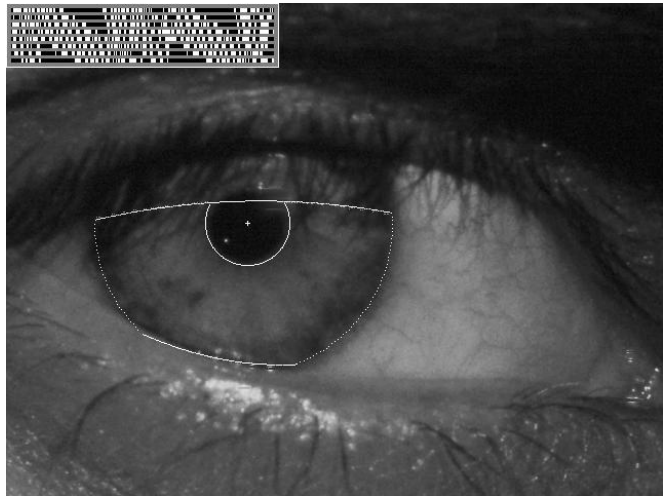

Face Recognition

http://www.face-rec.org/

Challenge: appearance changes

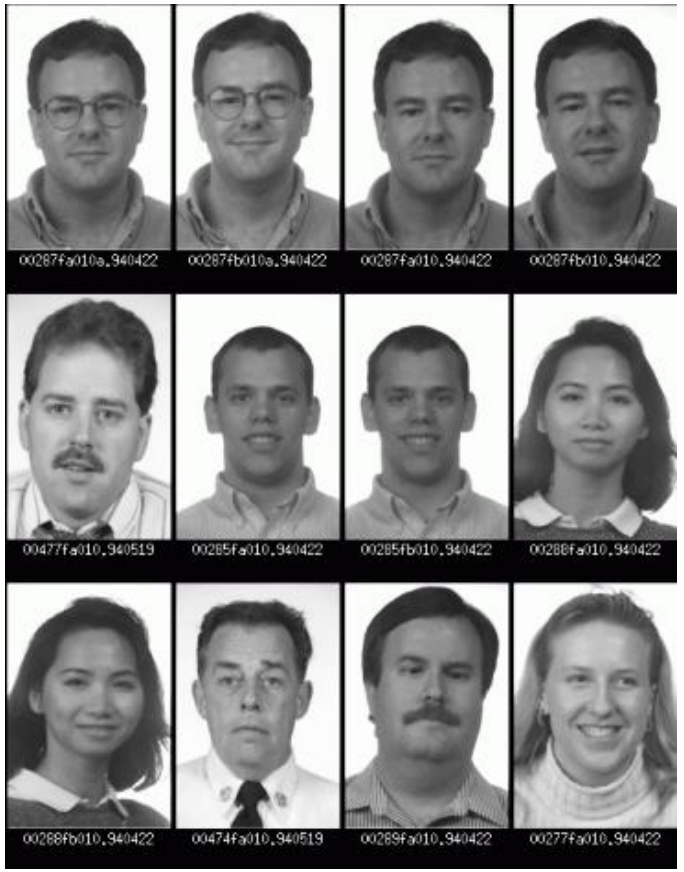

Gender Classification

- Useful for collecting demographic data but also boosting face recognition performance!
- Related applications: race classification, age classification.

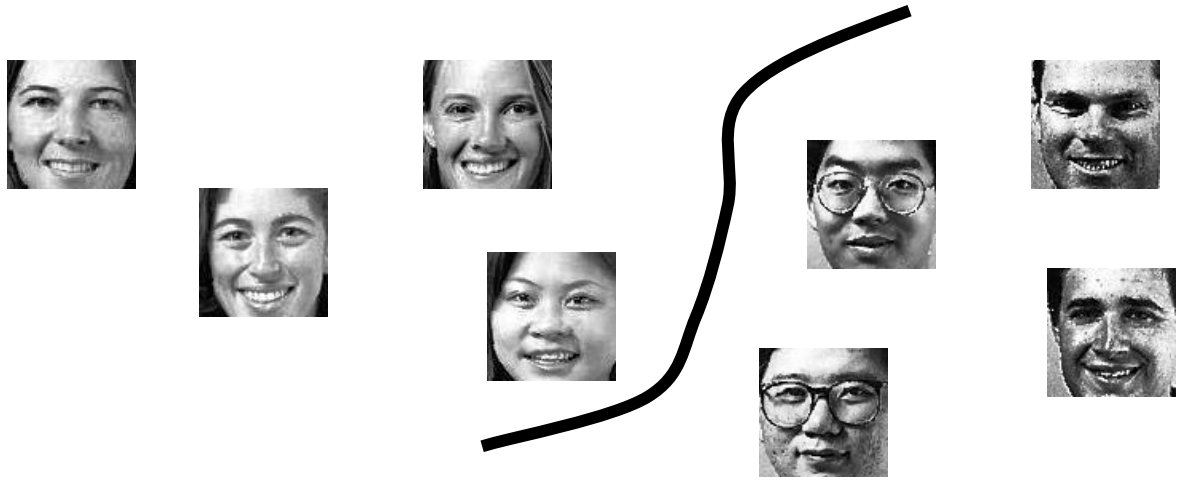

Key challenge: choose features that encode gender information but not identity information!

Facial Expression Recognition

http://www.youtube.com/watch?v=M1WgnisIyPQ&feature=related

Login without a password…

Fingerprint scanners on many new laptops, other devices

Face recognition systems now beginning to appear more widely
http://www.sensiblevision.com/

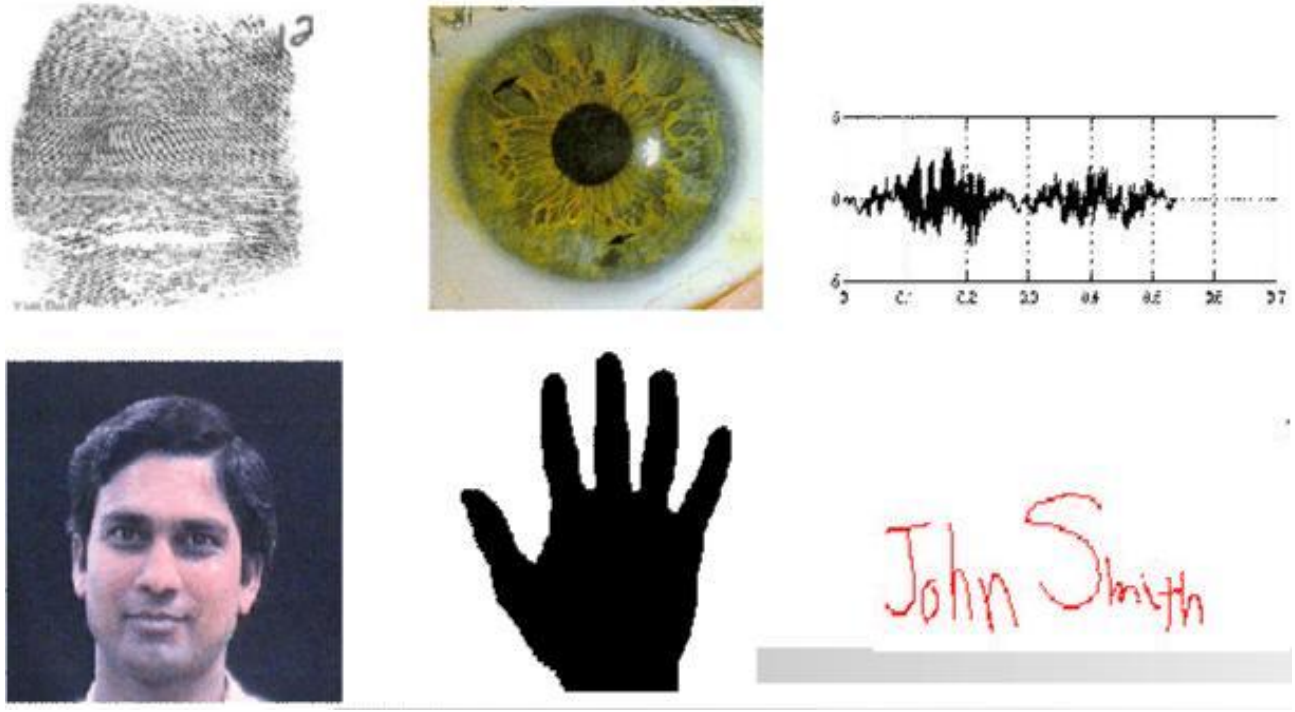

Authentication Using Biometrics

Fingerprint Recognition

minutiae

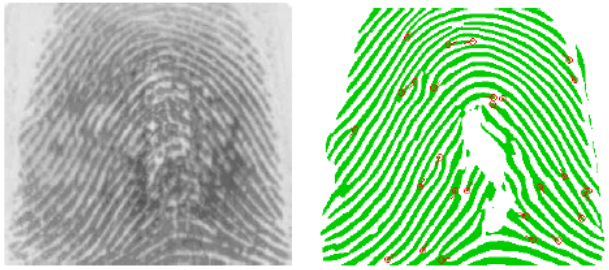

Challenge: small overlapping area

input

matching

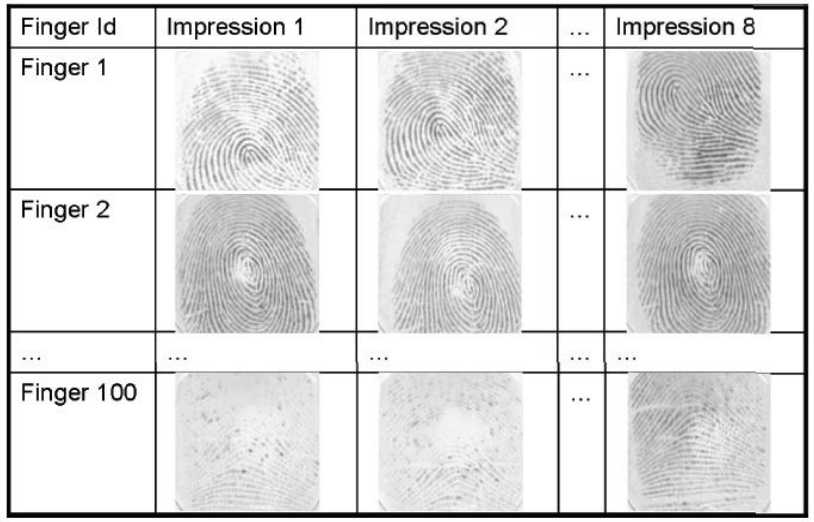

ID

Object Recognition

Toshiba Tech IS-910T 2013

DataLogic LaneHawk LH4000 2012

Special effects: shape capture

The Matrix movies, ESC Entertainment, XYZRGB, NRC

Special effects: motion capture

Pirates of the Carribean, Industrial Light and Magic
Click here for interactive demo

Automatic Panorama Stitching

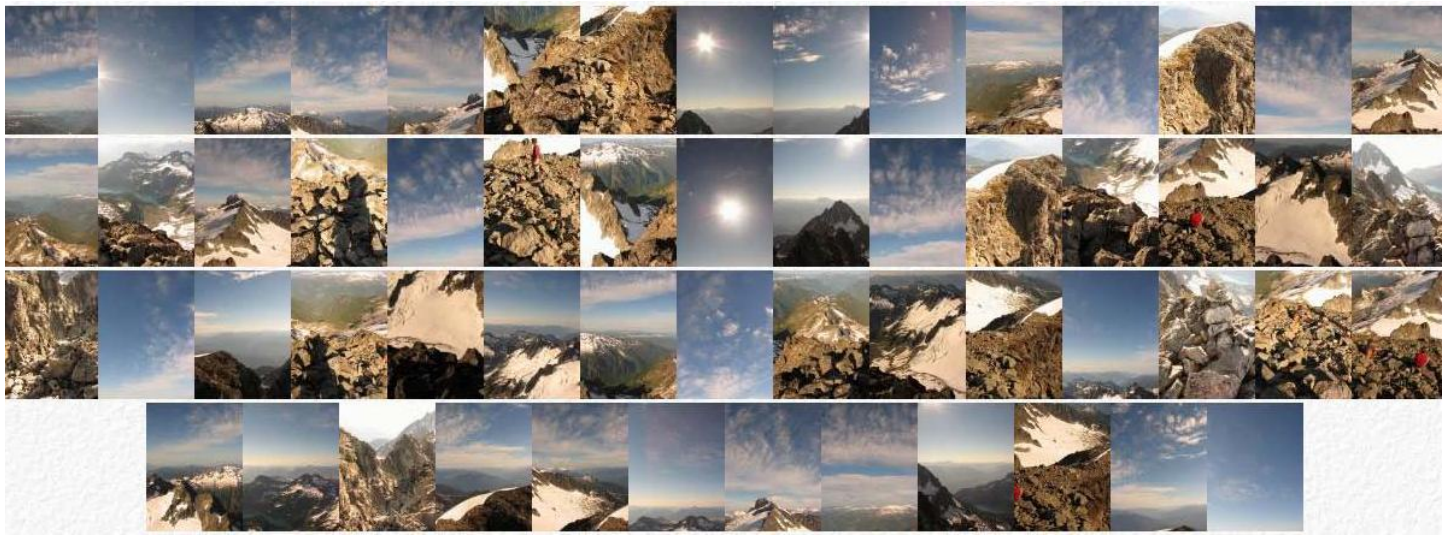

3D Reconstruction from internet photo collections

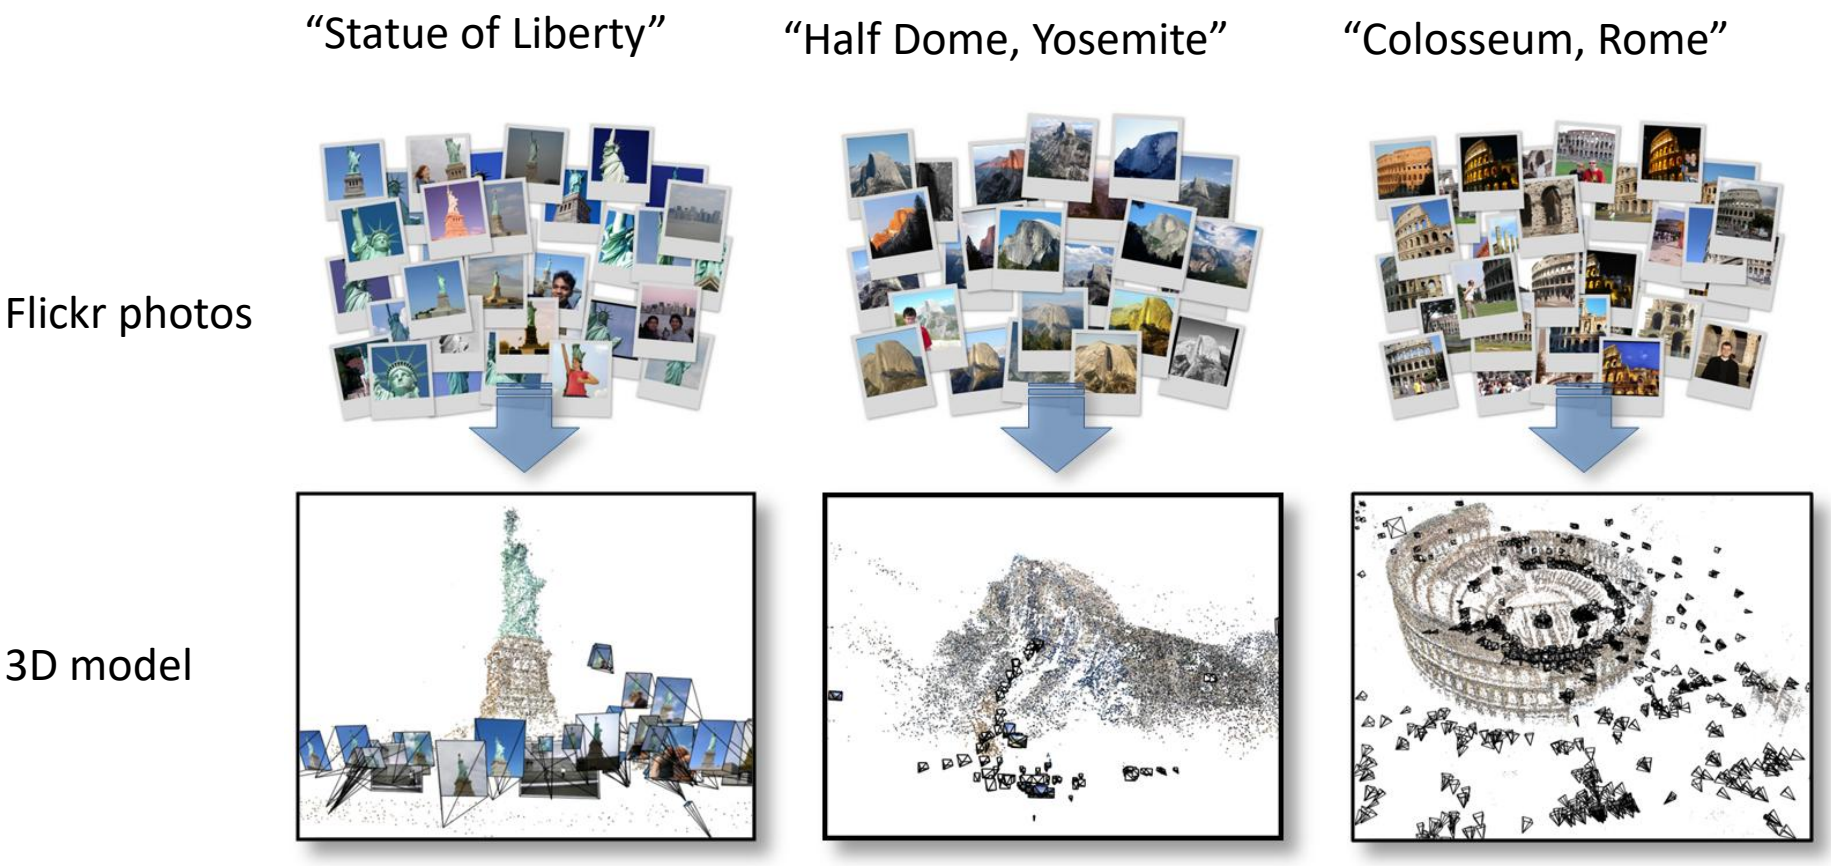

see "Building Rome in a day" project at U. Washington http://grail.cs.washington.edu/rome/

3D from thousands of images

Building Rome in a Day: Agarwal et al. 2009

Sports

Sportvision first down line
Nice explanation on www.howstuffworks.com

BMW 5 series

BMW night vision

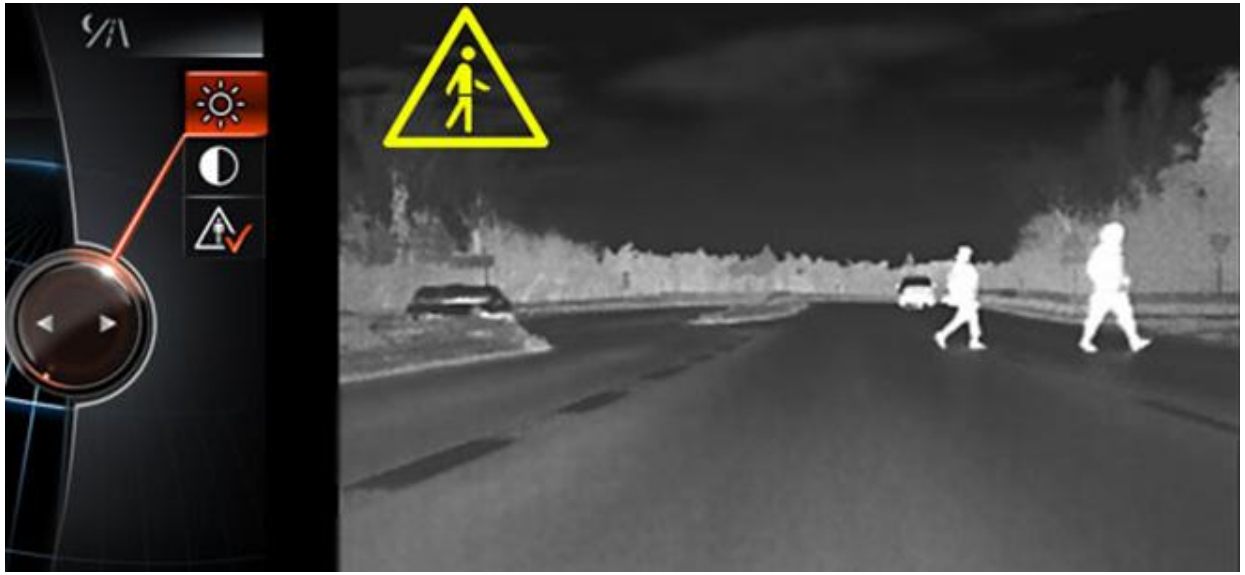

Games and Assistive Technologies

Nintendo Wii has camera-based IR tracking built in. See Lee's work at CMU on clever tricks on using it to create a multi-touch display!

Kinect

Assistive technologies

Virtual Fitting

2015

Interactive Games: Kinect

- Object Recognition: http://www.youtube.com/watch?feature=iv&v=fQ59dXOo63o
- Mario: http://www.youtube.com/watch?v=8CTJL5lUjHg
- 3D: http://www.youtube.com/watch?v=7QrnwoO1-8A
- Robot: http://www.youtube.com/watch?v=w8BmgtMKFbY

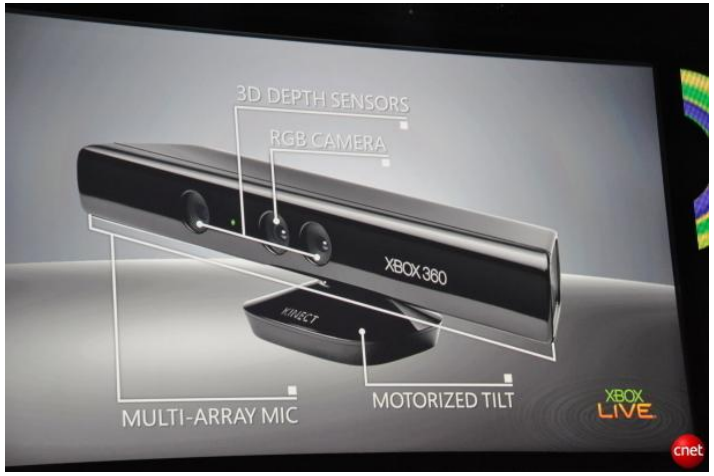

Vision in space

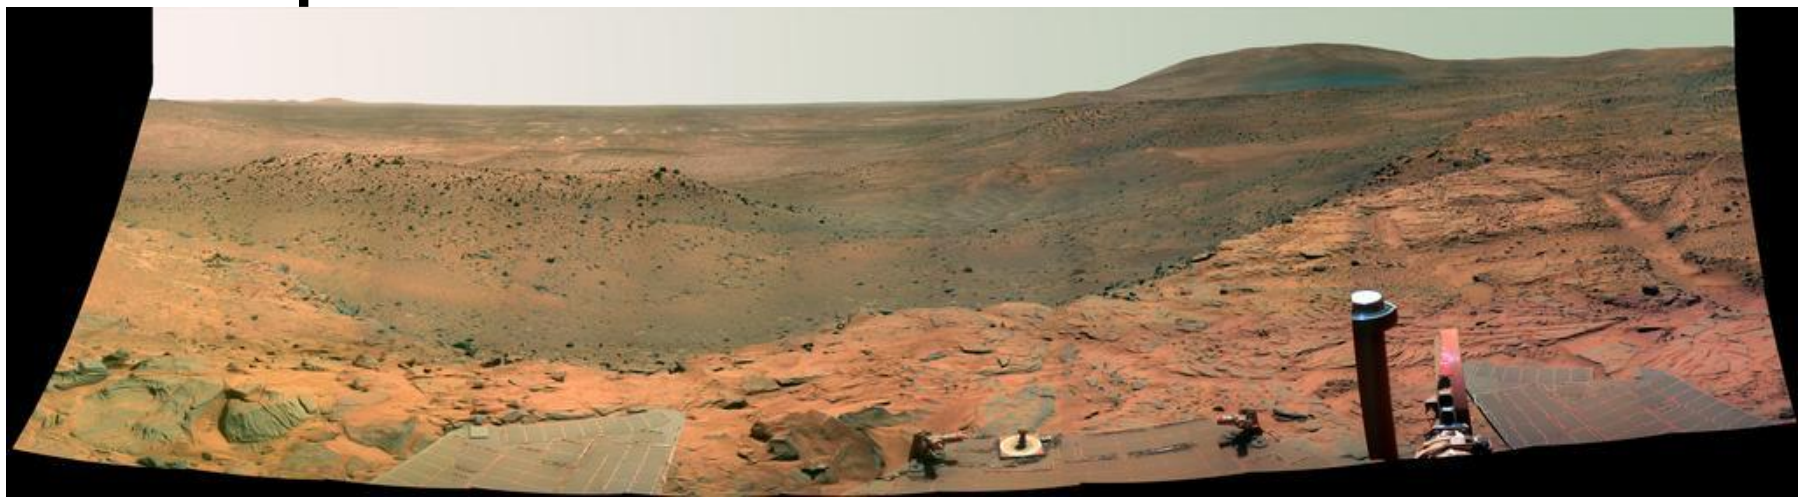

NASA'S Mars Exploration Rover Spirit captured this westward view from atop a low plateau where Spirit spent the closing months of 2007.

Vision systems (JPL) used for several tasks

- Panorama stitching
- 3D terrain modeling
- Obstacle detection, position tracking
- For more, read “Computer Vision on Mars” by Matthies et al.

Vision in Space

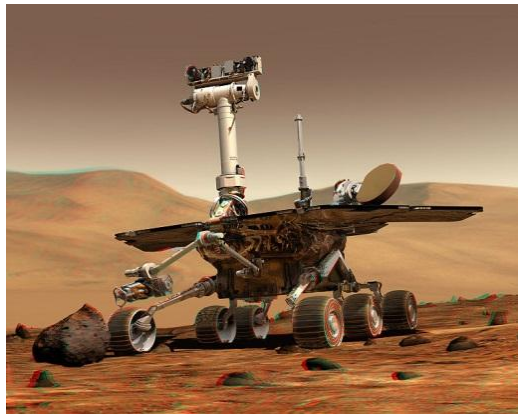

NASA'S Mars Exploration Rover Spirit

- Vision systems used for several tasks
  - Obstacle detection
  - Position tracking
  - 3D terrain modeling

For more info, read "Computer Vision on Mars" by Matthies et al. International Journal of Computer Vision, 2007.

Crater Detection

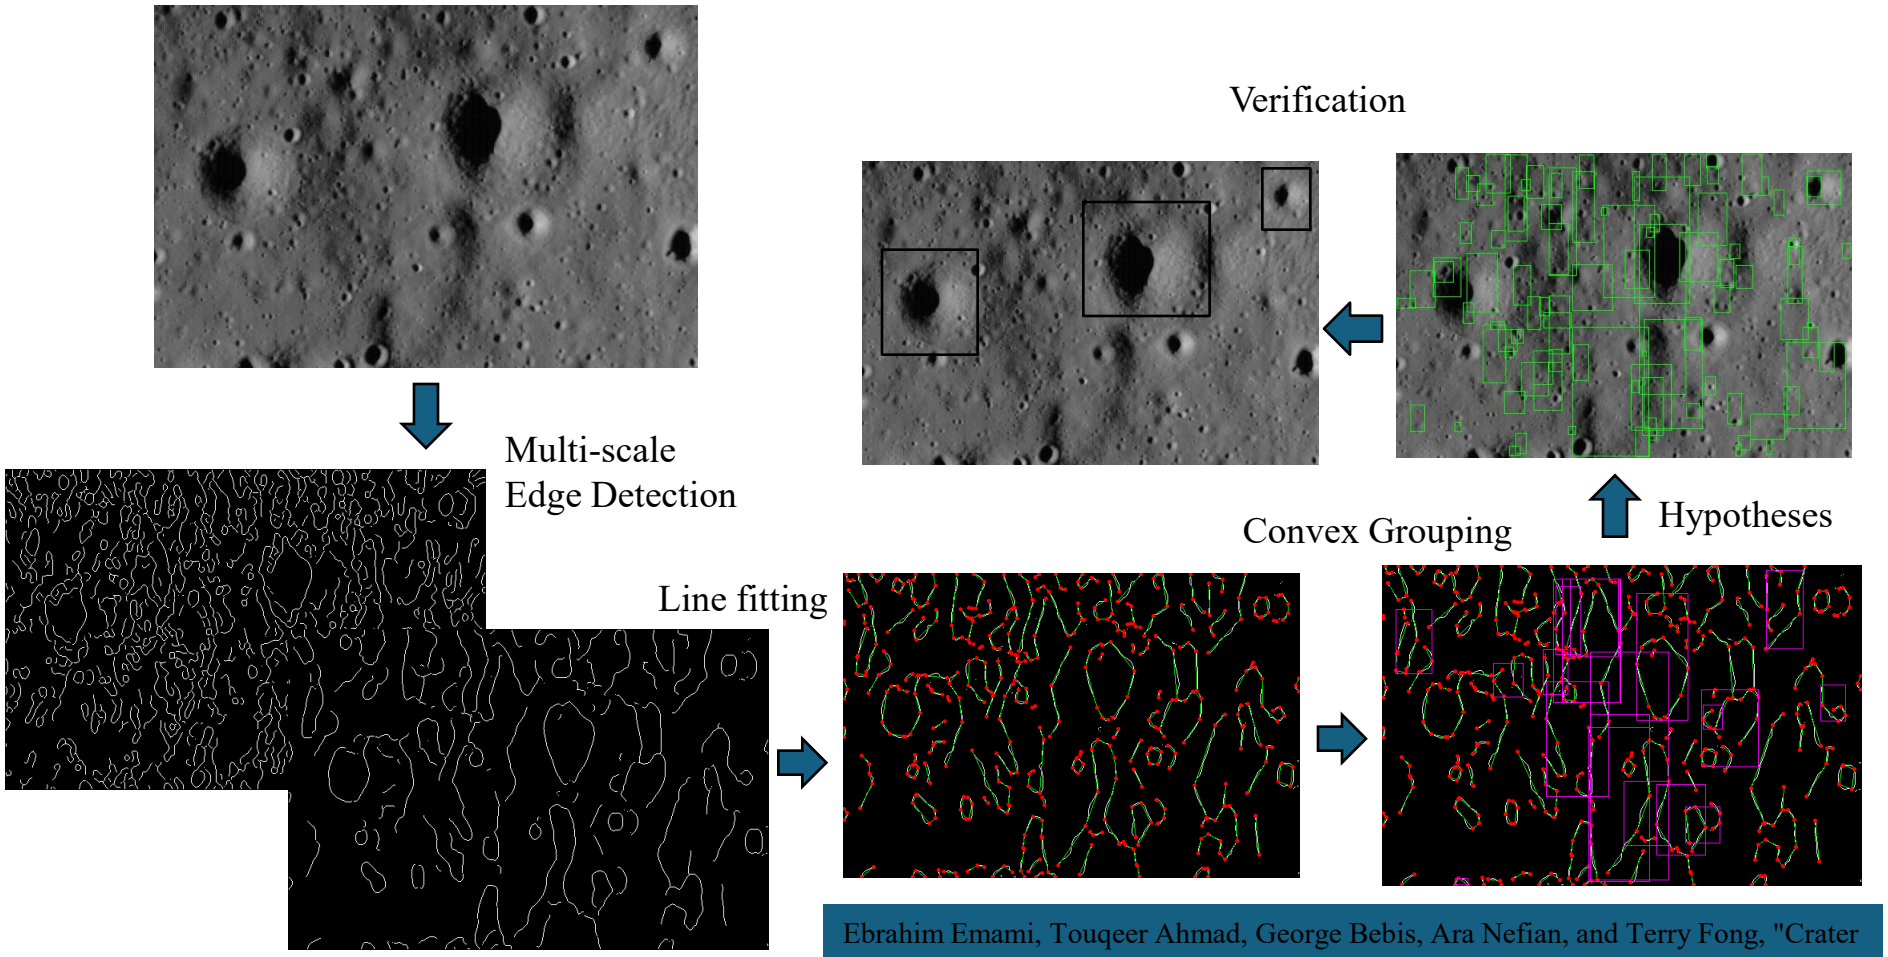

Ebrahim Emami, Touqeer Ahmad, George Bebis, Ara Nefian, and Terry Fong, "Crater Detection Using Unsupervised Algorithms and Convolutional Neural Networks", **IEEE Transactions on Geoscience and Remote Sensing**, vol. 57, no. 8, 2019.

Robotics

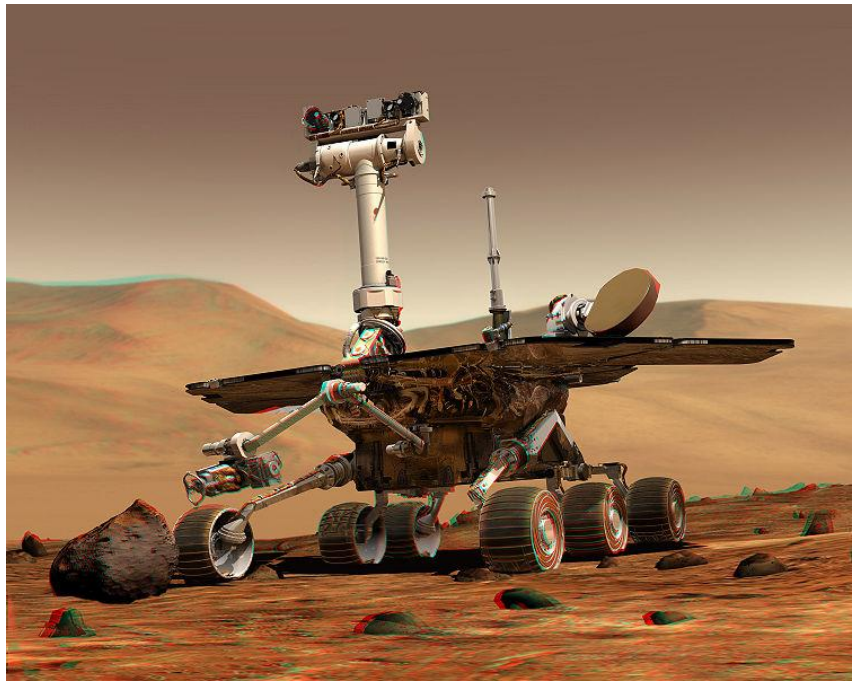

NASA's Mars Spirit Rover
http://en.wikipedia.org/wiki/Spirit_rover

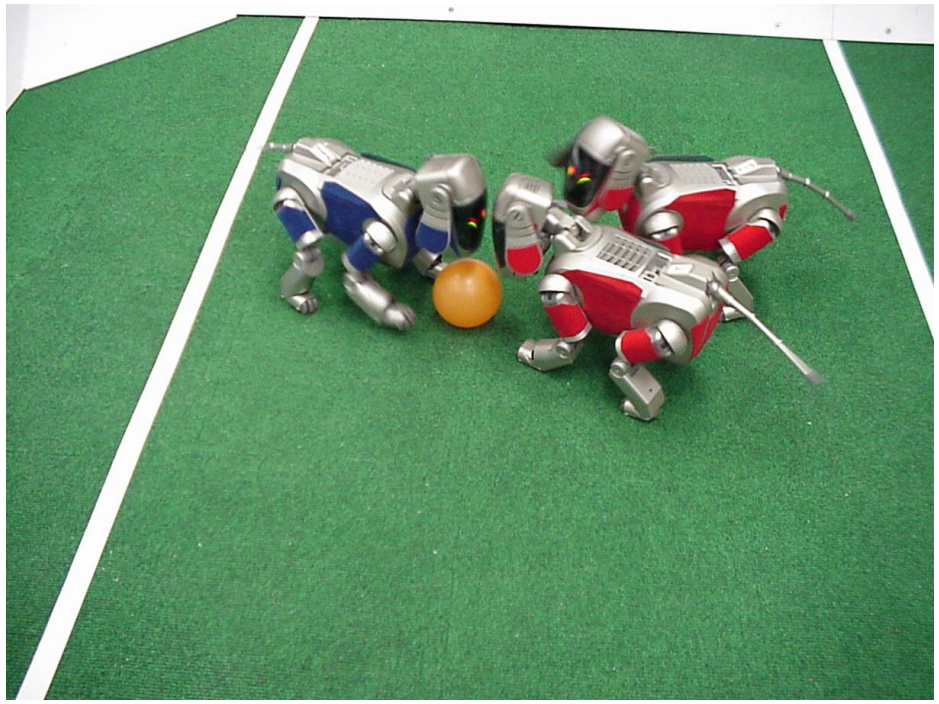

http://www.robocup.org/

Medical imaging

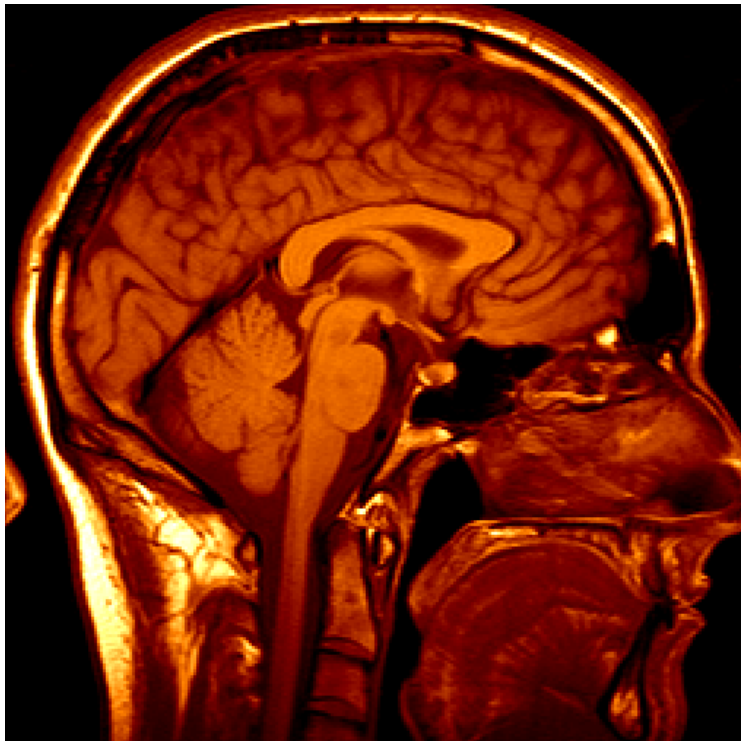

3D imaging
MRI, CT

Image guided surgery
Grimson et al., MIT

Image Operations

- Geometric Operations
- Point Operations
- Spatial Operations
- Global Operations (Freq. domain)
- Multi-Resolution Operations

Geometric Operations

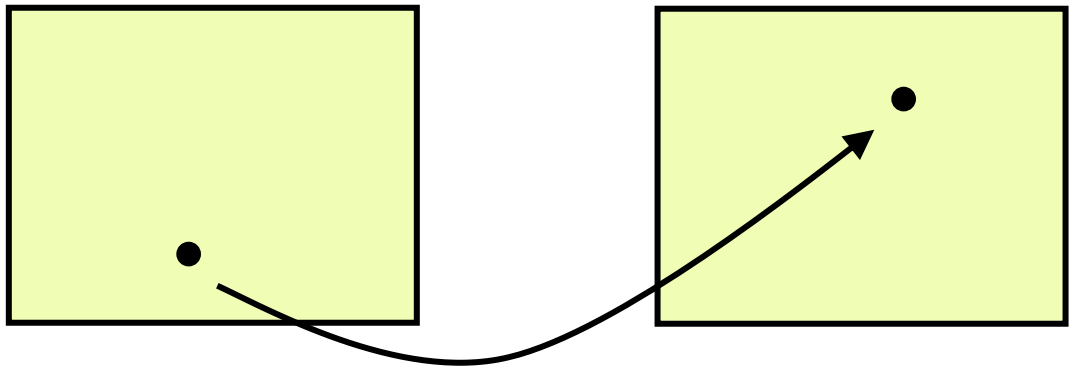

Point Operations

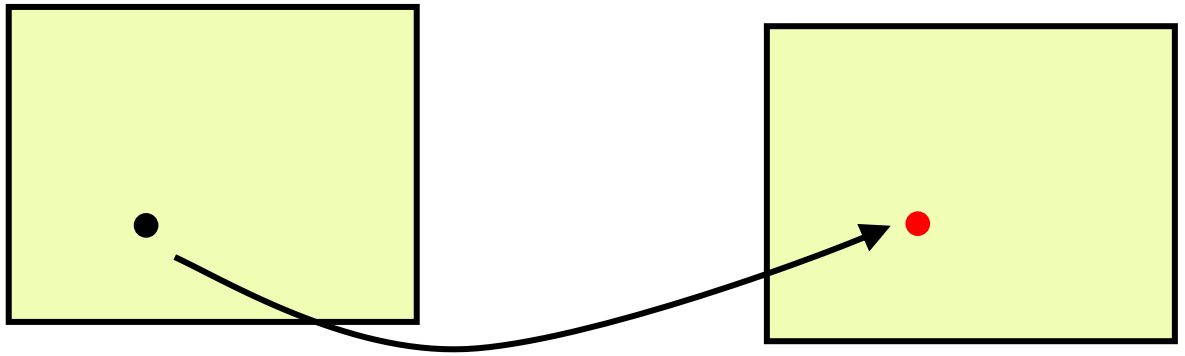

Geometric and Point Operations

Spatial Operations

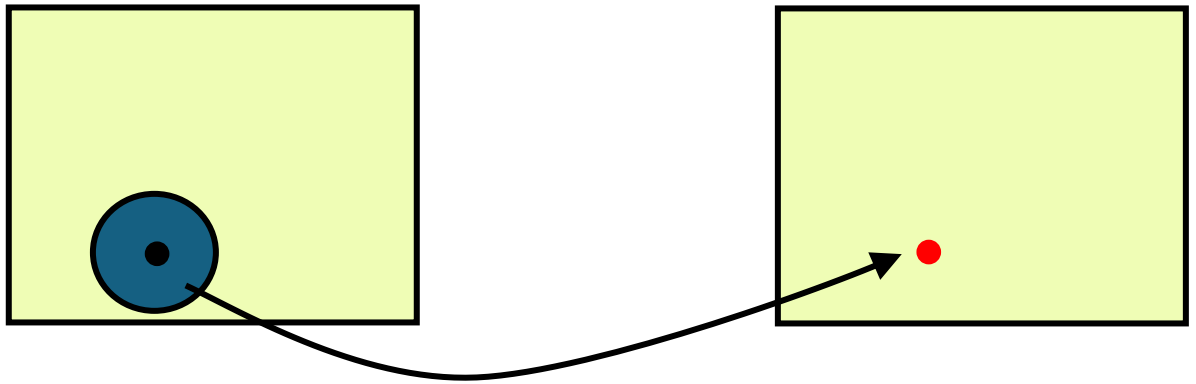

Global Operations

Global Operations

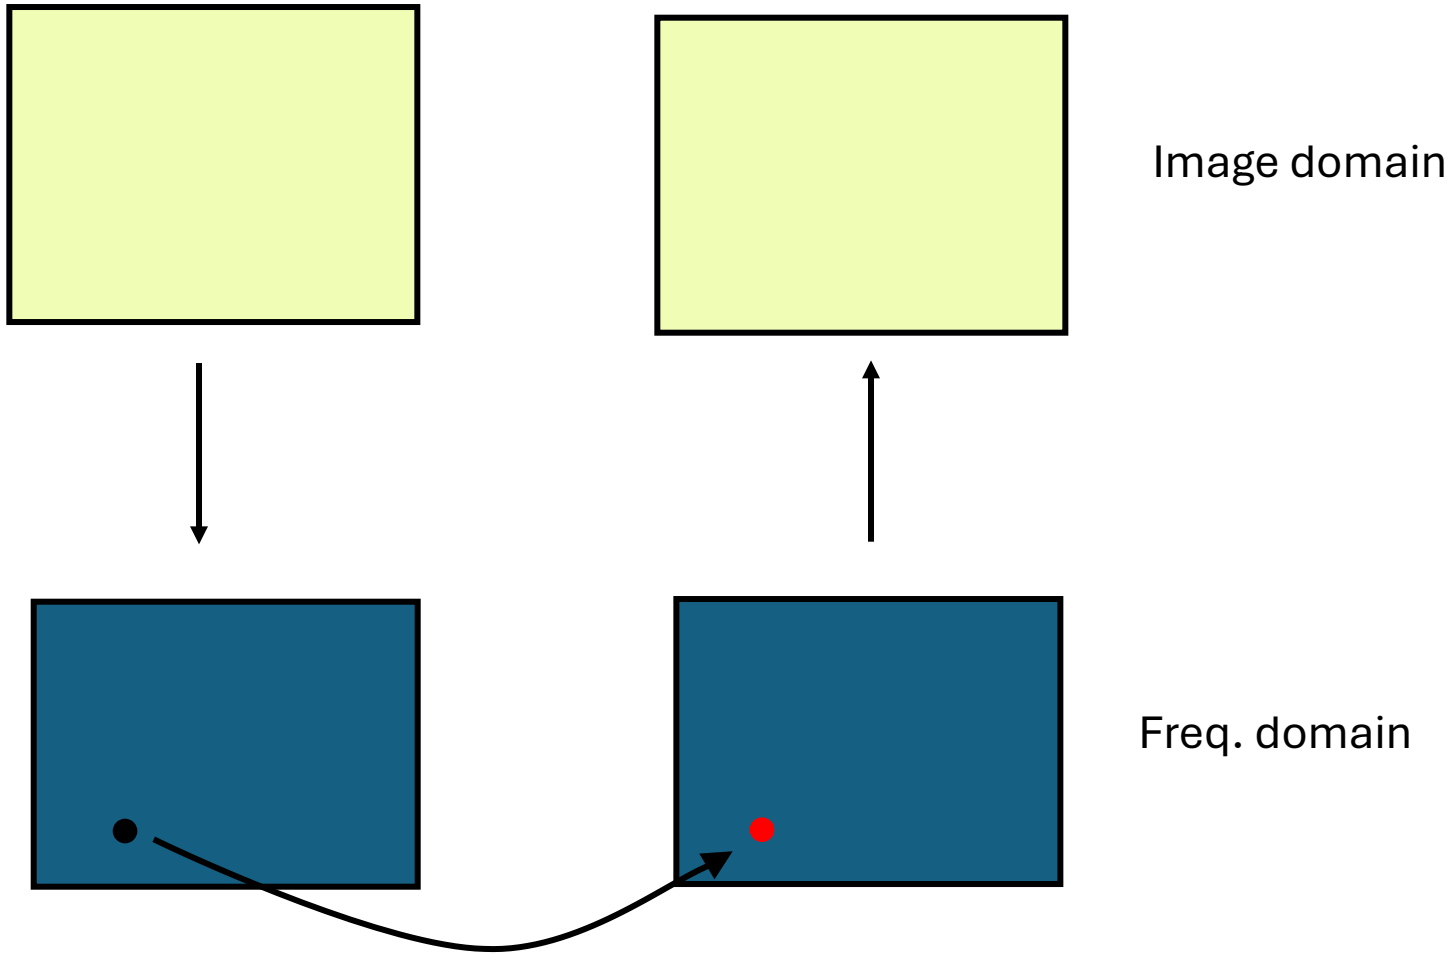

Multi-Resolution

Low resolution

High resolution

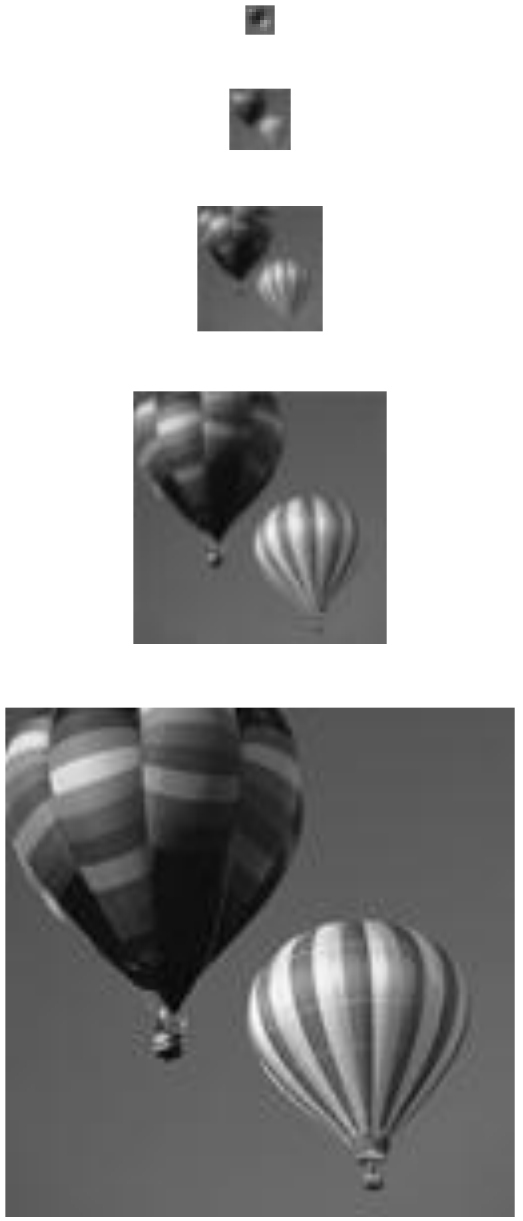

Image Digitization

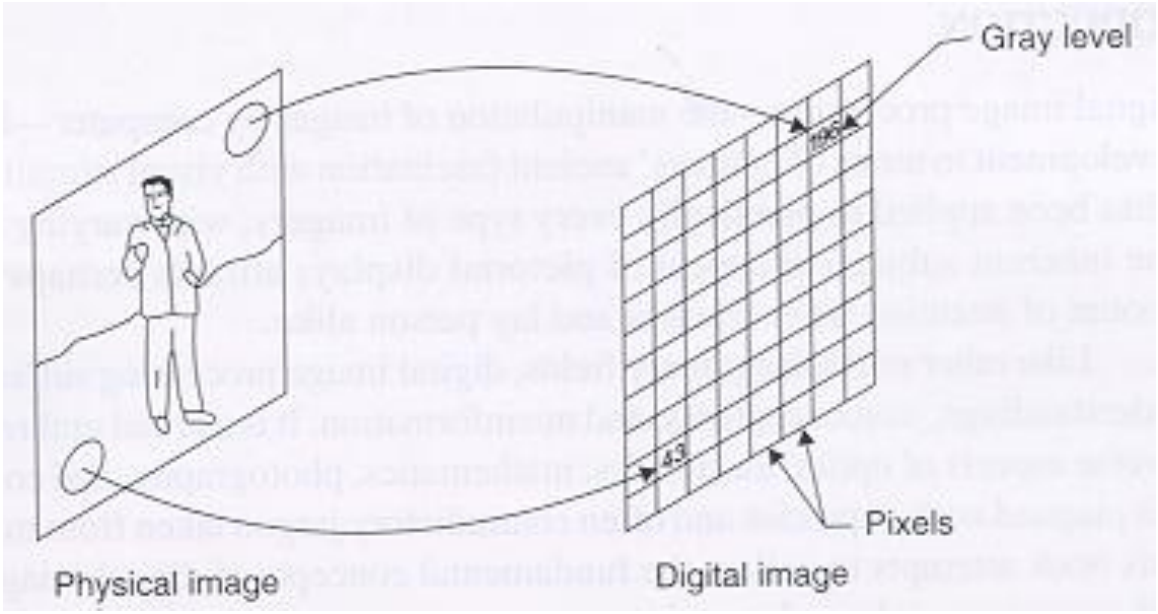

- **Sampling** is measuring the value of an image at a finite number of points (i.e., CCD array)
- **Quantization** is the representation of the measured value at the sampled point by an integer (i.e., frame grabber)

Image Digitization (cont'd)

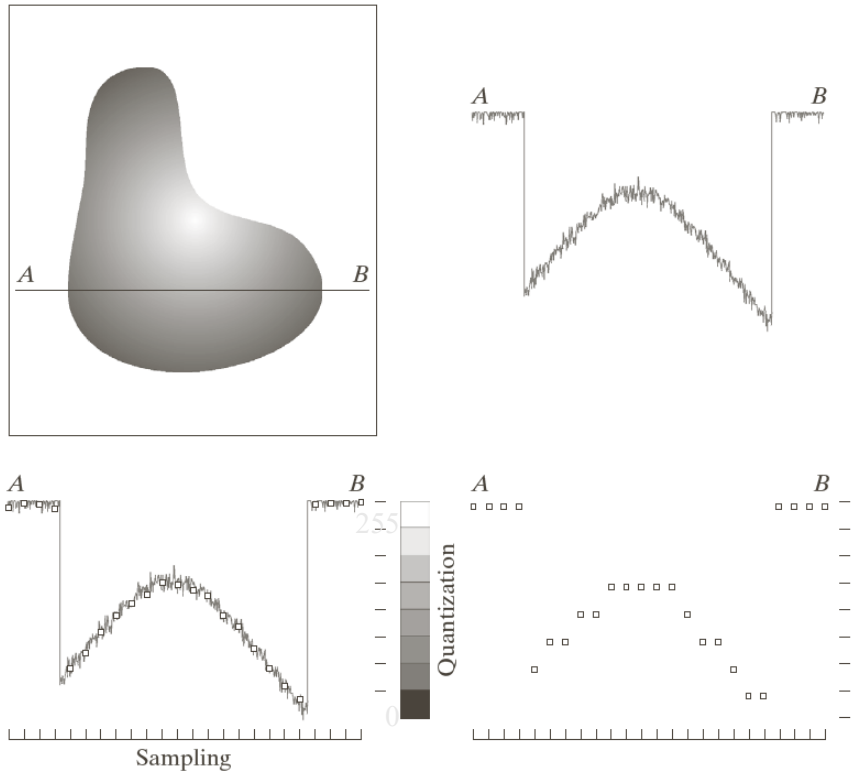

Sampling Quantization

Physical Image Digital Image

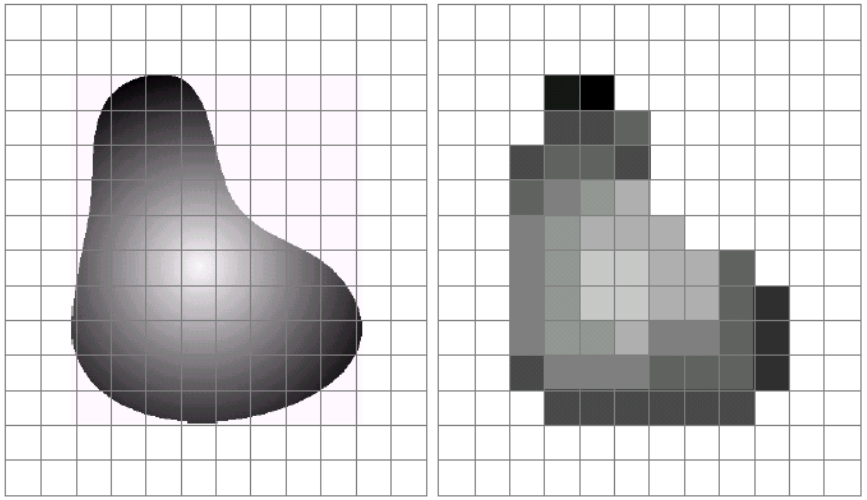

a b

FIGURE 2.17 (a) Continuos image projected onto a sensor array. (b) Result of image sampling and quantization.

Effect of Image Sampling

original image

sub-sampled by a factor of 2

sub-sampled by a factor of 4

sub-sampled by a factor of 8

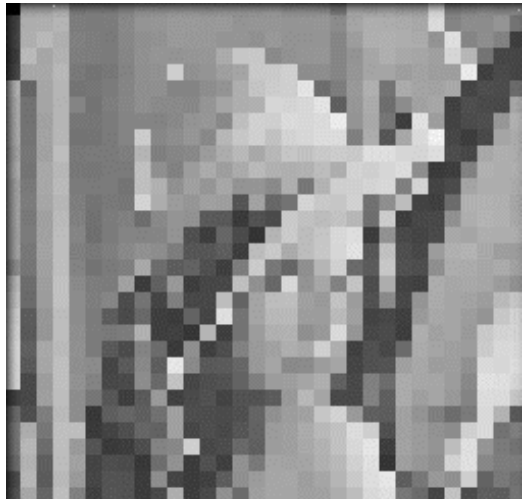

Note: images have been resized for comparison purposes

Effect of Image Quantization

256 gray levels (8 bits/pixel)

32 gray levels (5 bits/pixel)

16 gray levels (4 bits/pixel)

8 gray levels (3 bits/pixel)

4 gray levels (2 bits/pixel)

2 gray levels (1 bit/pixel)

What skills you need ?

- Strong programming skills (i.e., C, C++, Python, Matlab)
- Good knowledge of Data Structures and Algorithms
- Good skills in analyzing algorithm performance (i.e., time and memory requirements).
- Strong background in mathematics, especially in:
  - Linear Algebra
  - Probabilities and Statistics
  - Numerical Analysis
  - Geometry
  - Calculus

Textbook

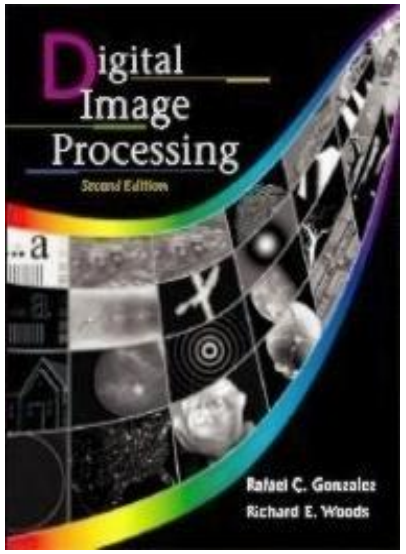

- Digital Image Processing

  Rafael C. Gonzalez & Richard E. Woods,

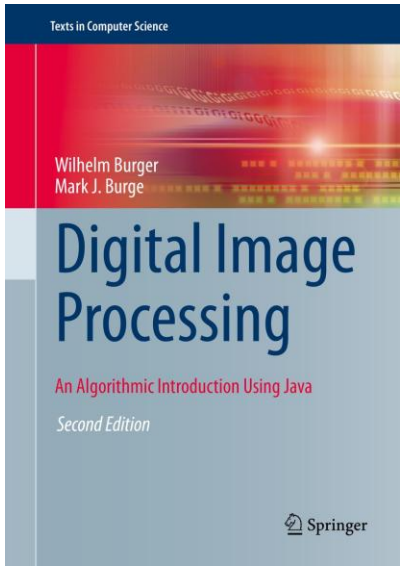

- Digital Image Processing

  S Esakkirajan T Veerakumar , S **Jayaraman**

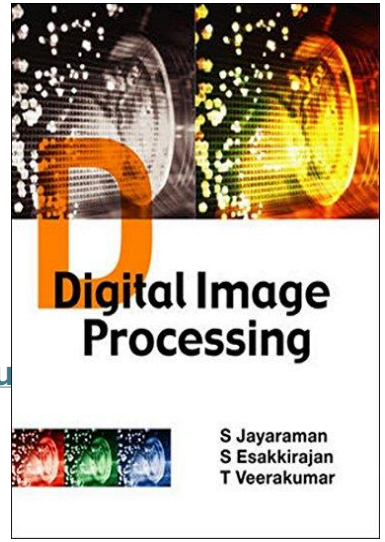

https://books.google.co.in/books?id=JeDGn6Wmf1kC&printsec=frontcover&source=gbs_ge_summary_r&cad=0#v=onepage&q&f=false

- Digital Image Processing An Algorithmic Introduction Using Java

  Wilhelm Burger, Mark J. Burge

Secondary Text

- Concise Computer Vision

  Reinhard Klette

- Fundamentals of Digital Image Processing

  Chris Solomon, Toby Breckon

.

Computer Vision: Algorithms and Applications

© 2010 Richard Szeliski, Microsoft Research

http://szeliski.org/Book/