Digital Image Processing

Hongkai Xiong

熊红凯

http://ivm.sjtu.edu.cn

电子工程系

上海交通大学

22 Feb. 2016

About Me

Hongkai Xiong, distinguished professor

- Office: 5-419, the 5-th building of dianxin building group
- Email: xionghongkai@sjtu.edu.cn
- Web-page: http://ivm.sjtu.edu.cn

About TA

Yuehan Xiong, phd candidate

- Email: xiongyuehan@sjtu.edu.cn

Xing Gao, phd candidate

- Email: william-g@sjtu.edu.cn

About The Class

- Requirements and Grading:

  - Homework and attendance: 20%
  - Projects (2+1): 20%+20%
  - Final Exam (close-book): 40%

About The Class

Text book and reference:

- R.C. Gonzalez and R.E Woods, **Digital Image Processing**， Third Edition , Publishing House of Electronic Industry， 2010
- **数字图像处理，**第三版中文版， R.C. Gonzalez and R.E Woods, 阮秋琦、阮智宇等译 , 电子工业出版社

About The Class

- Programming resources:
  - Matlab
  - OpenCV library (c/c++)
- Website: http://ivm.sjtu.edu.cn/dip.html

What you will learn

Part I Digital Image Fundamentals

- Human visual perception
- Image sensing and Acquisition
- Some basic knowledge

What you will learn

Part II Low-level processing

- Intensity Transformations and Image Filtering
- Image Restoration and Reconstruction
- Wavelets and Multiresolution Processing
- Image Compression

What you will learn

Part III high-level processing

- Image Segmentation
- Morphological Image Processing
- Representation and Recognition

A bit more about us

图像-视频-多媒体通信实验室 IVM Laboratory

Research Topic:

- Computer Vision:
  - Image classification – ImageNet *international challenge*
  - 3-D reconstruction
  - Activity identification

A bit more about us

图像-视频-多媒体通信实验室 IVM Laboratory

Research Topic:

- Machine Learning and Deep Learning:
  - Multitask learning
- Computational Photography
  - Light-field camera
- Biomedical Image Processing
  - Gene sequence compression

Welcome you genius guys to join us!

Course Contents(16weeks, 48hours)

- Course Review
- Part 1

  Digital Image Fundamentals
- Part 2

  Low Level Digital Image Processing
- Part 3

  High Level Digital Image Processing

Part I

Introduction

- History and examples of fields that use DIP

Digital Image Fundamentals

- Visual Perception
- Light and the Electromagnetic Spectrum
- Sensing and Acquisition
- Sampling and Quantization
- Image Quality Assessment

Color Image Processing

- Color Fundamentals & Color Models
- Pseudocolor Image Processing & Full-Color Image Processing
- Color Transformations, Smoothing and Sharpening

Part II

Image Filtering
- Image Filtering in Spatial Domain
- Image Filtering in Frequency Domain
- Image Enhancement

Image Restoration and Reconstruction
- Image Restoration
- Image Reconstruction

Wavelets and Multiresolution Processing
- Multi-resolution Expansions
- Wavelet Transforms

Image Compression
- Fundamentals of Image Compression
- Basic Compression Methods
- Image Compression Standards

Part III

Image Segmentation

- Fundamentals
- Point, Line, and Edge Detection
- Thresholding
- Region-Based Segmentation

Morphological Image Processing

- Preliminaries
- Erosion and Dilation
- Opening and Closing
- Some Basic Morphological Algorithms

This lecture will cover

- Why Do We Process Images?
- History of Digital Image Processing
- Fields that Use Digital Image Processing
- Key Stages in Digital Image Processing
- Something Cool

Why Do We Process Images?

- Acquire an image
  - Correct aperture and color balance
  - Reconstruct image from projections
- Prepare for display or printing
  - Adjust image size
  - Halftoning
- Facilitate picture storage and transmission
  - Efficiently store an image in a digital camera
  - Send an image from Mars to Earth
- Enhance and restore images
  - Remove scratches from an old movie
  - Improve visibility of tumor in a radiograph
- Extract information from images
  - Read the ZIP code on a letter
  - Measure water pollution from aerial images

History of Digital Image Processing

Early 1920s: One of the first applications of digital imaging was in the news-paper industry

- The Bartlane cable picture transmission service
- Images were transferred by submarine cable between London and New York
- Pictures were coded for cable transfer and reconstructed at the receiving end on a telegraph printer

Early digital image

History of Digital Image Processing

Mid to late 1920s: Improvements to the Bartlane system resulted in higher quality images

- New reproduction processes based on photographic techniques
- Increased number of tones in reproduced images

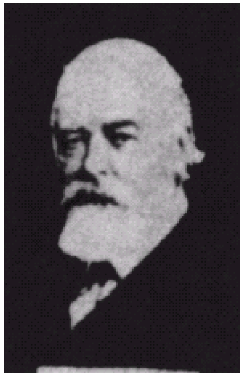

Improved digital image

Early 15 tone digital image

History of Digital Image Processing

1960s: Improvements in computing technology and the onset of the space race led to a surge of work in digital image processing

- **1964:** Computers were used to improve the quality of images of the moon taken by the *Ranger 7* probe
- Such techniques were used in other space missions including the Apollo landings

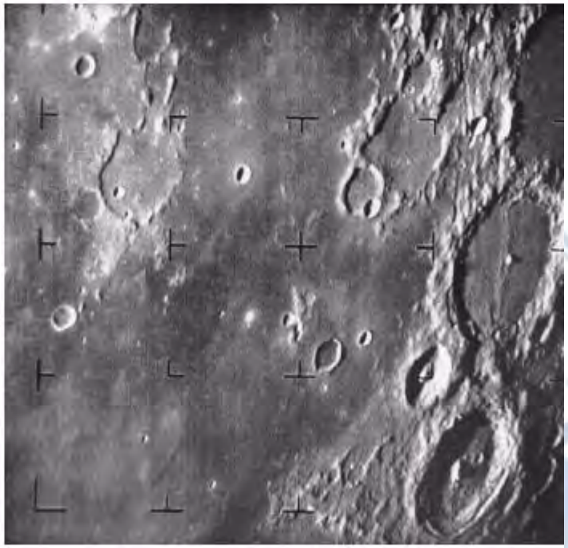

A picture of the moon taken by the Ranger 7 probe minutes before landing

History of Digital Image Processing

1970s: Digital image processing begins to be used in medical applications

- **1979:** Sir Godfrey N. Hounsfield & Prof. Allan M. Cormack shared the Nobel Prize in medicine for the invention of tomography, the technology behind Computerised Axial Tomography (CAT) scans

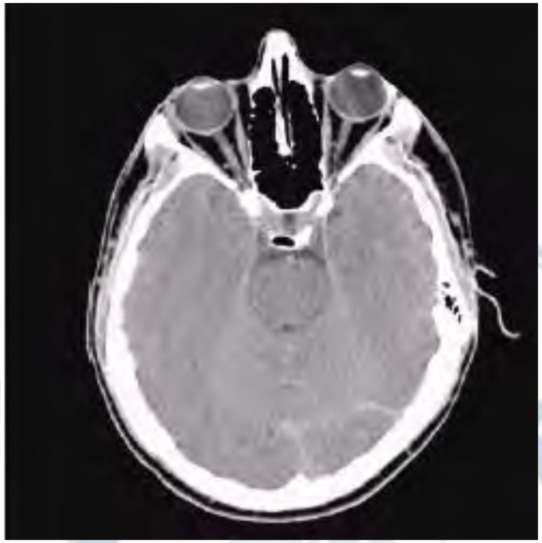

Typical head slice CAT image

History of Digital Image Processing

1980s - Today: The use of digital image processing techniques has exploded and they are now used for all kinds of tasks in a broad range of areas

- Image enhancement/restoration
- Artistic effects
- Medical visualisation
- Industrial inspection
- Law enforcement
- Human computer interfaces

Fields that Use Digital Image Processing

- Energy of one photon

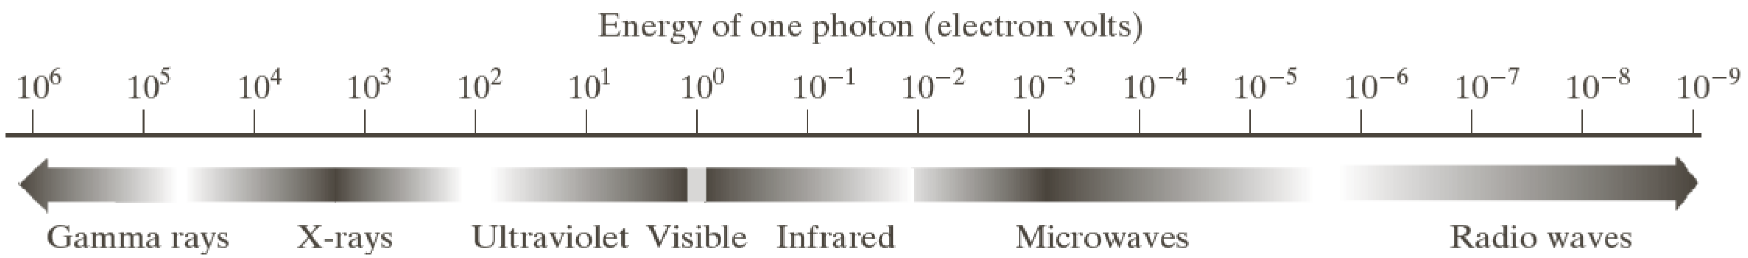

FIGURE 1.5 The electromagnetic spectrum arranged according to energy per photon.

- Image from Invisible light
  - γ- ray imaging
  - X- ray imaging
  - Imaging in the ultraviolet band
  - Imaging in the infrared band
  - Imaging in the microwave band
  - Imaging in the radio band

Fields that Use Digital Image Processing

- Examples of gamma-ray imaging

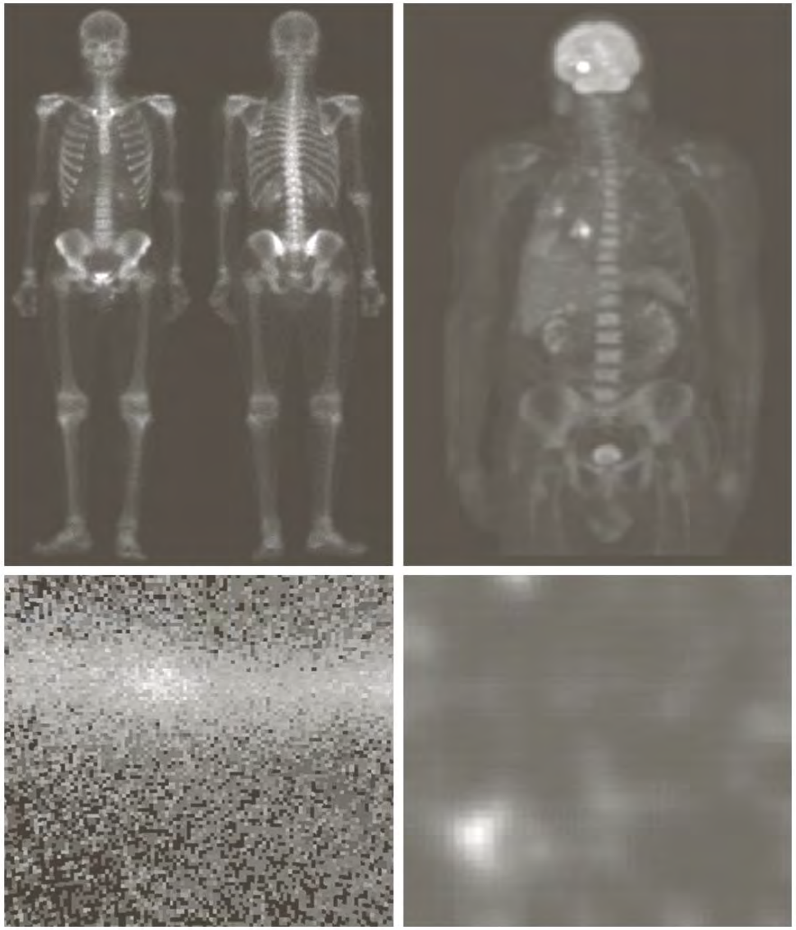

a	b
c	d

FIGURE 1.6 Examples of gamma-ray imaging. (a) Bone scan. (b) PET image. (c) Cygnus Loop. (d) Gamma radiation (bright spot) from a reactor valve. (Images courtesy of (a) G.E. Medical Systems, (b) Dr. Michael E. Casey, CTI PET Systems, (c) NASA, (d) Professors Zhong He and David K. Wehe, University of Michigan.)

Fields that Use Digital Image Processing

- Examples of X-ray imaging

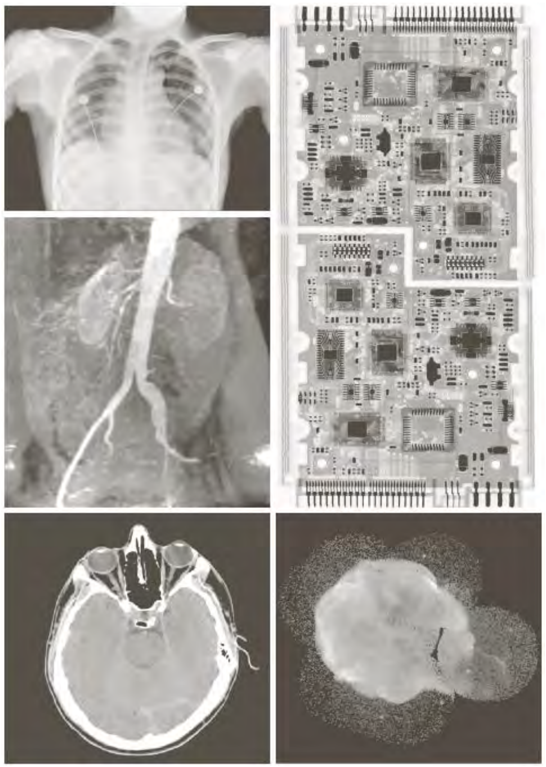

a	d
b	
c	e

FIGURE 1.7 Examples of X-ray imaging. (a) Chest X-ray. (b) Aortic angiogram. (c) Head CT. (d) Circuit boards. (e) Cygnus Loop. (Images courtesy of (a) and (c) Dr. David R. Pickens, Dept. of Radiology & Radiological Sciences, Vanderbilt University Medical Center; (b) Dr. Thomas R. Gest, Division of Anatomical Sciences, University of Michigan Medical School; (d) Mr. Joseph E. Pascente, Lixi, Inc.; and (e) NASA.)

The First X-ray Photo
Wilhelm Röntgen
(1845～1923)

Fields that Use Digital Image Processing

- Examples of ultraviolet imaging

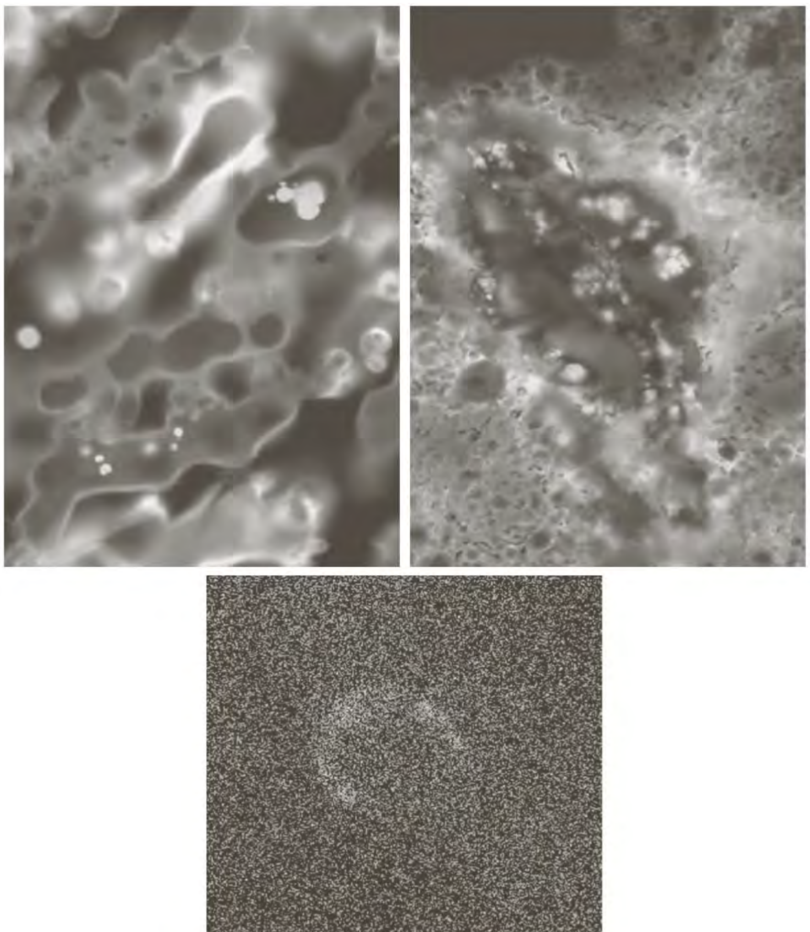

a b
c

FIGURE 1.8 Examples of ultraviolet imaging. (a) Normal corn. (b) Smut corn. (c) Cygnus Loop. (Images courtesy of (a) and (b) Dr. Michael W. Davidson, Florida State University, (c) NASA.)

Fields that Use Digital Image Processing

- Examples of light microscopy imaging

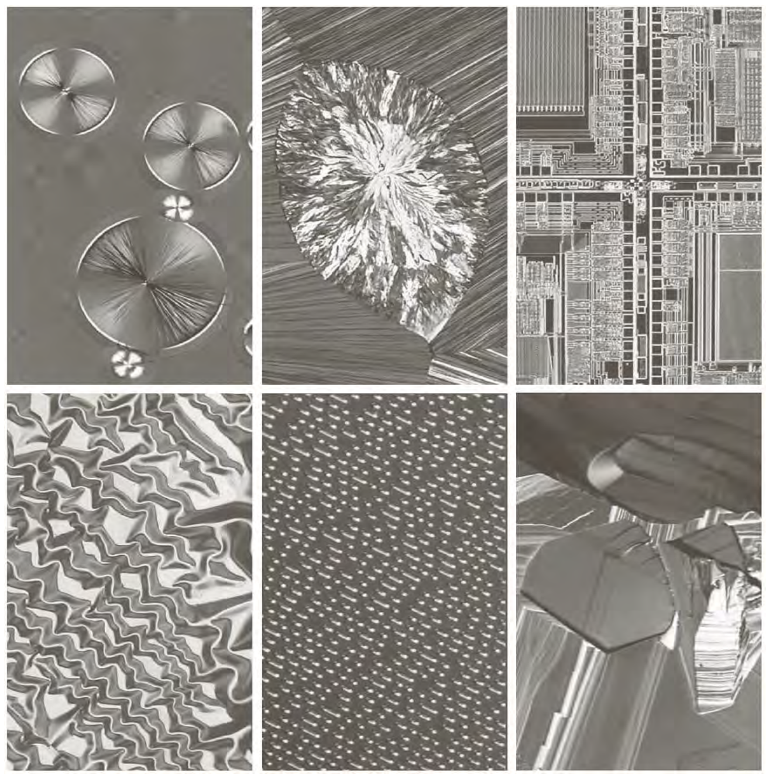

a b c
d e f

FIGURE 1.9 Examples of light microscopy images. (a) Taxol (anticancer agent), magnified 250×. (b) Cholesterol—40×. (c) Microprocessor—60×. (d) Nickel oxide thin film—600×. (e) Surface of audio CD—1750×. (f) Organic superconductor—450×. (Images courtesy of Dr. Michael W. Davidson, Florida State University.)

Fields that Use Digital Image Processing

- LANDSAT satellite images of the Washington, D.C. area

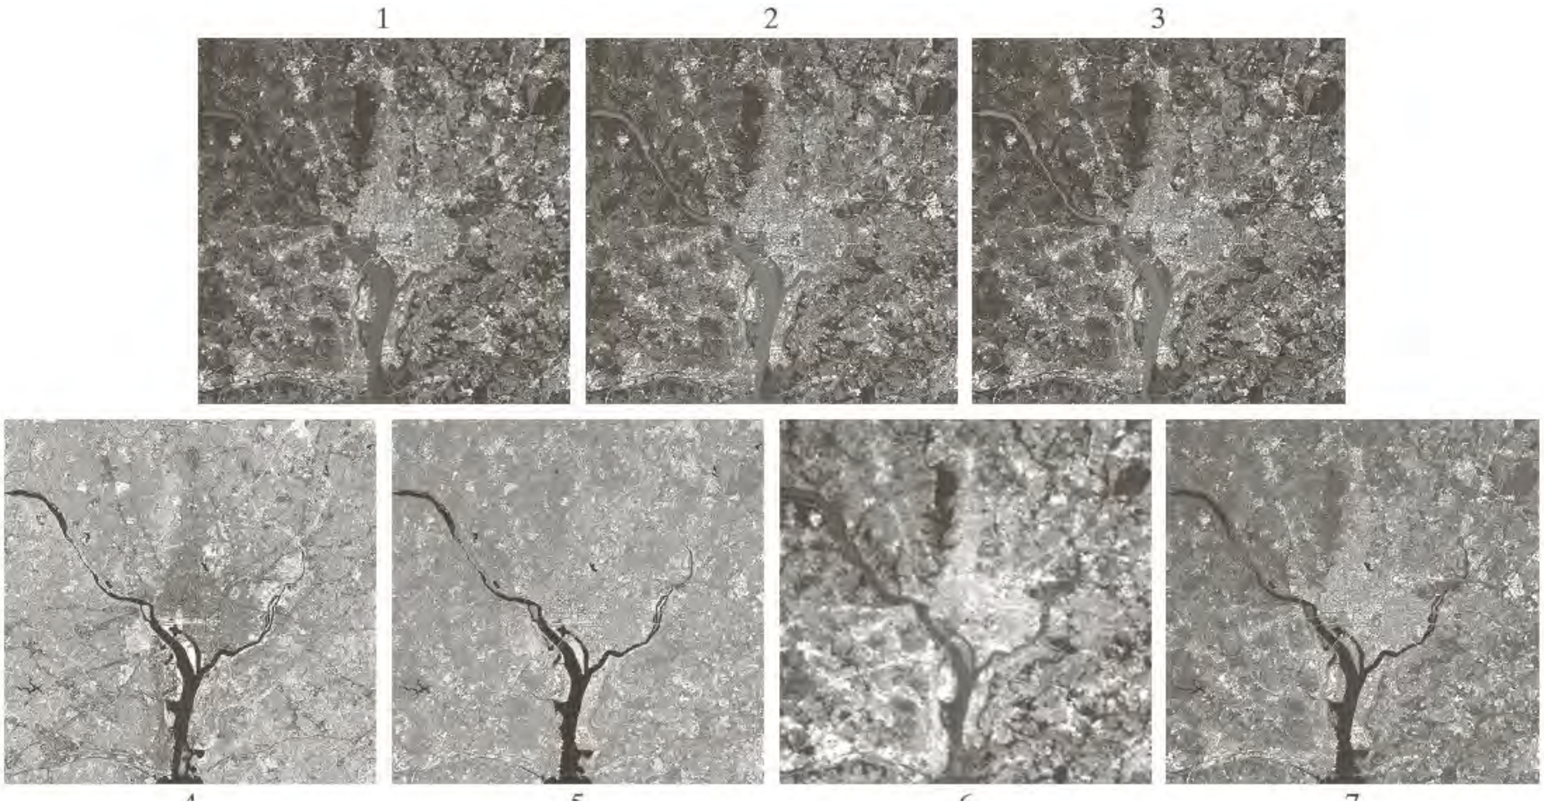

FIGURE 1.10 LANDSAT satellite images of the Washington, D.C. area. The numbers refer to the thematic bands in Table 1.1. (Images courtesy of NASA.)

Fields that Use Digital Image Processing

- Satellite image of Hurricane

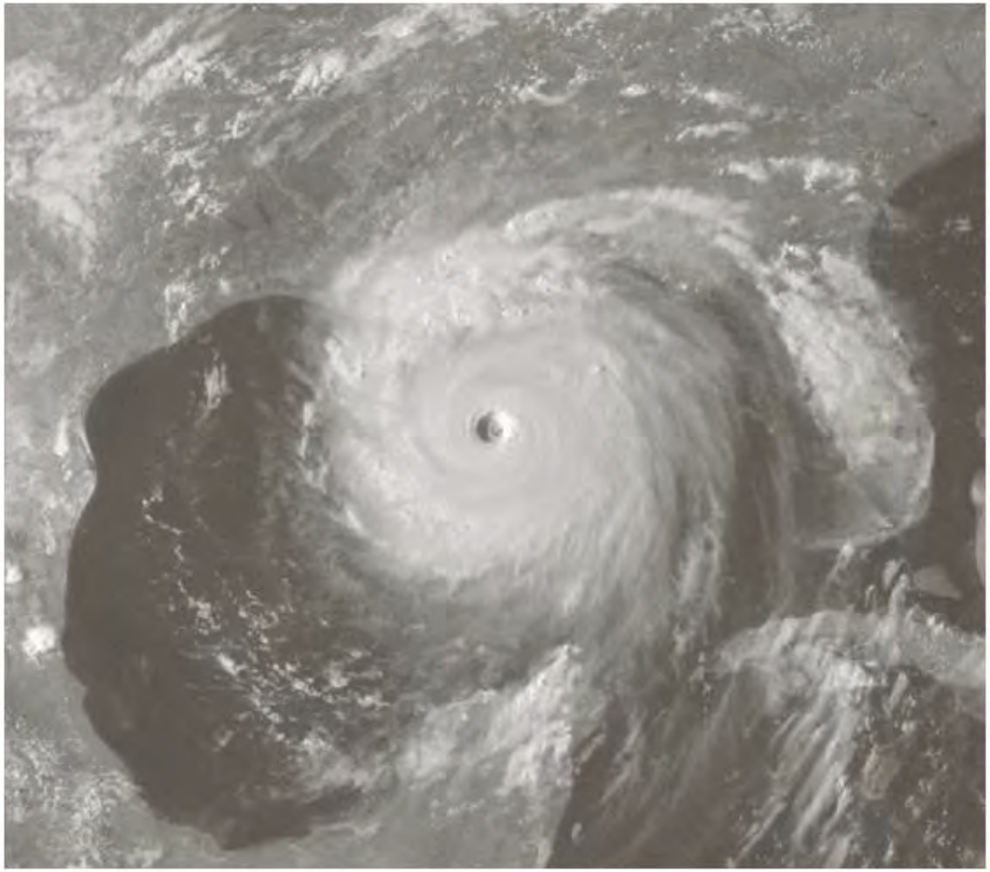

FIGURE 1.11 Satellite image of Hurricane Katrina taken on August 29, 2005. (Courtesy of NOAA.)

Fields that Use Digital Image Processing

- Infrared satellite images of the Americans.

FIGURE 1.12 Infrared satellite images of the Americas. The small gray map is provided for reference. (Courtesy of NOAA.)

Fields that Use Digital Image Processing

shanghai

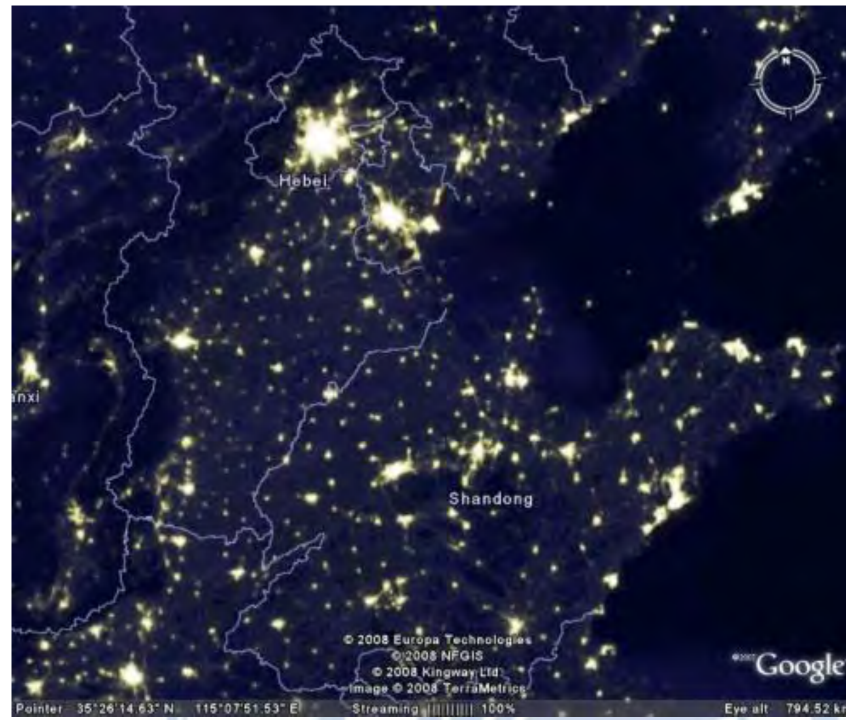

beijing

To see the level of development from brightness

Fields that Use Digital Image Processing

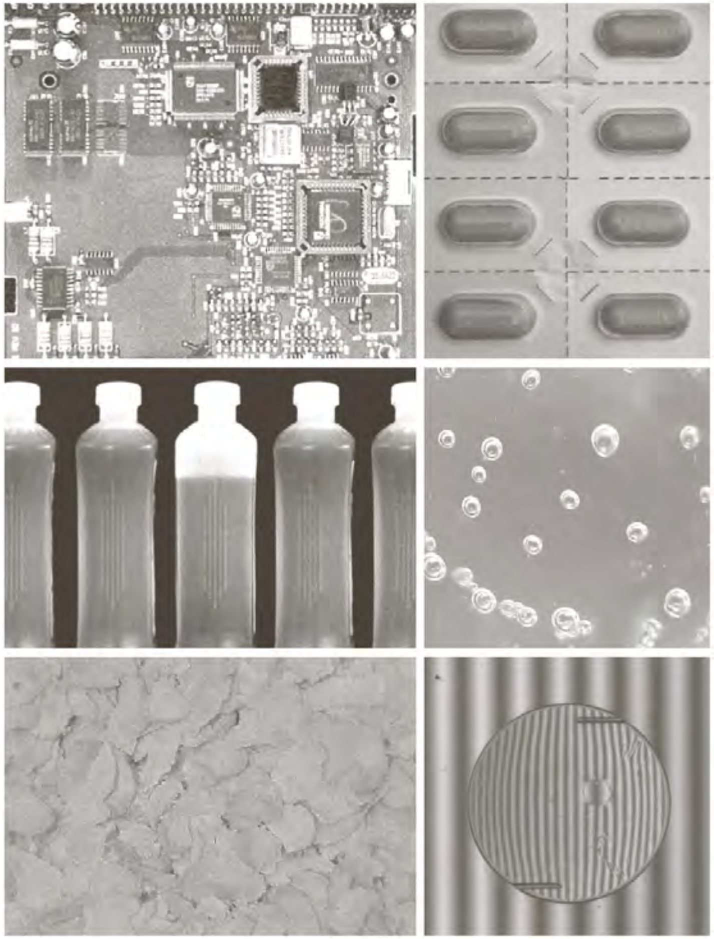

a	b
c	d
e	f

FIGURE 1.14 Some examples of manufactured goods often checked using digital image processing. (a) A circuit board controller. (b) Packaged pills. (c) Bottles. (d) Air bubbles in a clear-plastic product. (e) Cereal. (f) Image of intraocular implant. (Fig. (f) courtesy of Mr. Pete Sites, Perceptics Corporation.)

Fields that Use Digital Image Processing

a b
c
d

FIGURE 1.15 Some additional examples of imaging in the visual spectrum. (a) Thumb print. (b) Paper currency. (c) and (d) Automated license plate reading. (Figure (a) courtesy of the National Institute of Standards and Technology. Figures (c) and (d) courtesy of Dr. Juan Herrera, Perceptics Corporation.)

Fields that Use Digital Image Processing

FIGURE 1.16 Spaceborne radar image of mountains in southeast Tibet. (Courtesy of NASA.)

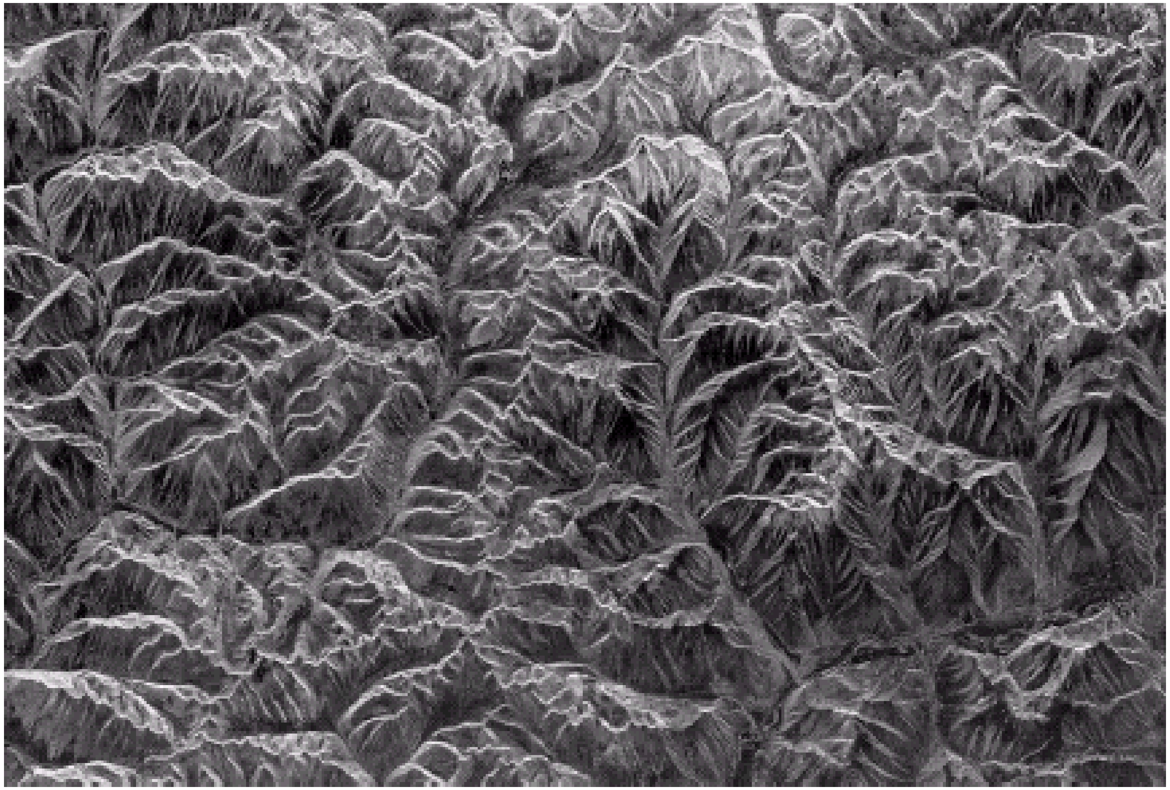

Fields that Use Digital Image Processing

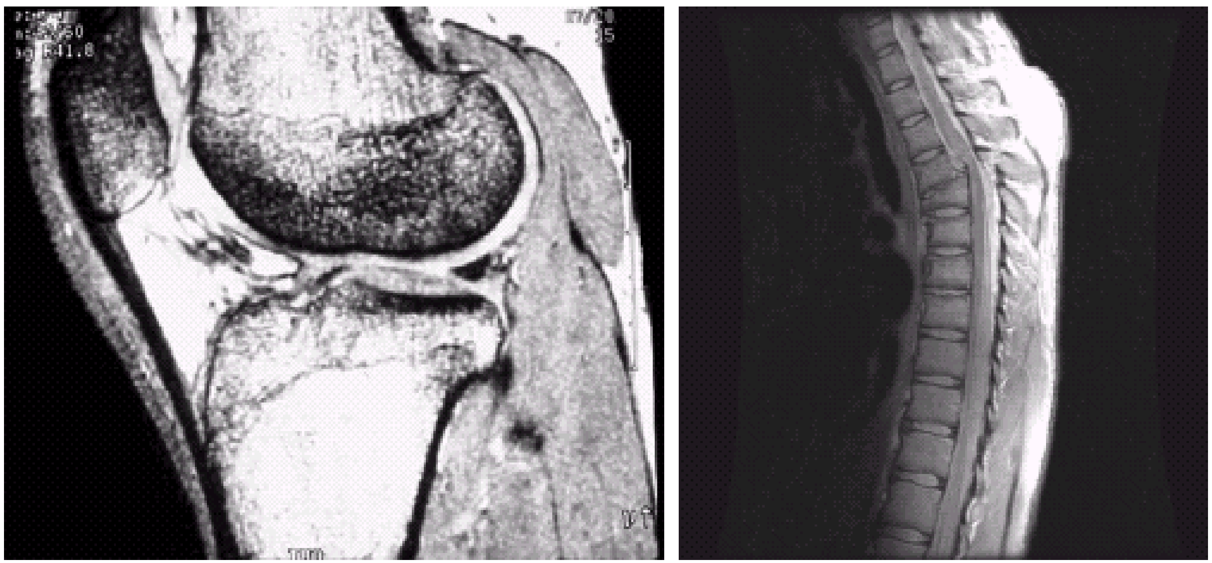

a b

FIGURE 1.17 MRI images of a human (a) knee, and (b) spine. (Image (a) courtesy of Dr. Thomas R. Gest, Division of Anatomical Sciences, University of Michigan Medical School, and (b) Dr. David R. Pickens, Department of Radiology and Radiological Sciences, Vanderbilt University Medical Center.)

Fields that Use Digital Image Processing

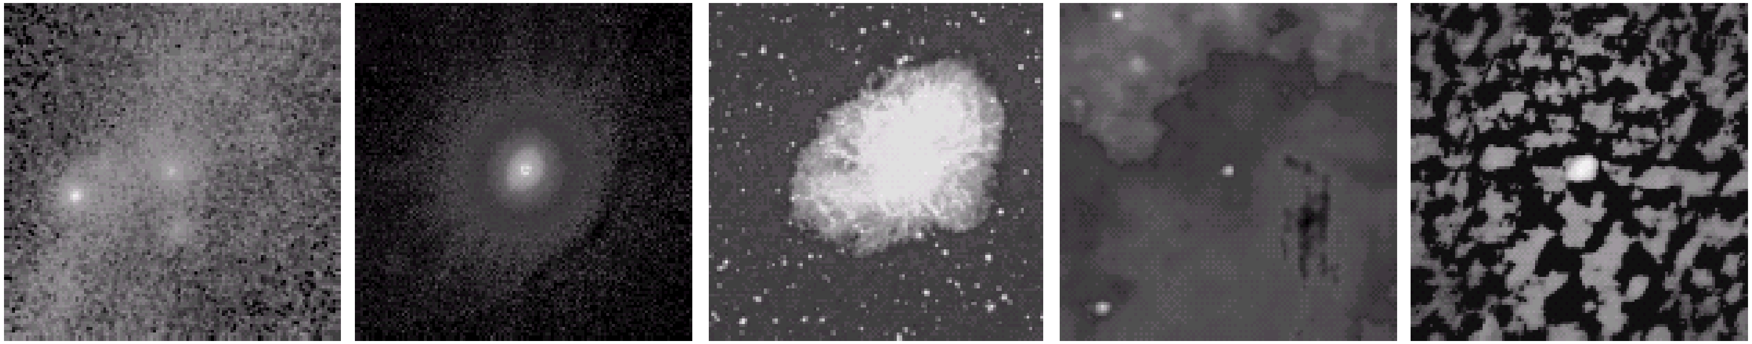

FIGURE 1.18 Images of the Crab Pulsar (in the center of images) covering the electromagnetic spectrum. (Courtesy of NASA.)

Fields that Use Digital Image Processing

Moving images (Video)

- Movie: 24 frames/second
- TV: 25 frames/second
- Gray scale image: $f_k(m, n)$
- Color image:
  $R_k(m, n)$, $G_k(m, n)$, $B_k(m, n)$

Fields that Use Digital Image Processing

Digital Image Processing

- Low-level processing
  - Pixel level (image → image)
  - This course only discuss low-level processing
  - Difficulties:
    - Real time
    - Adjacent region

Fields that Use Digital Image Processing

Image Compression

Compression at 0.5 bit per pixel by means of JPEG and JPEG2000

Fields that Use Digital Image Processing

Image Transform

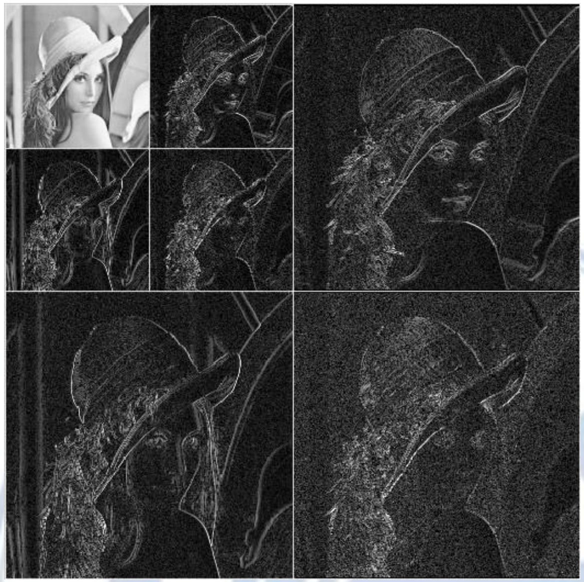

2-D wavelet transform

Fields that Use Digital Image Processing

Restoration of image from Hubble Space Telescope

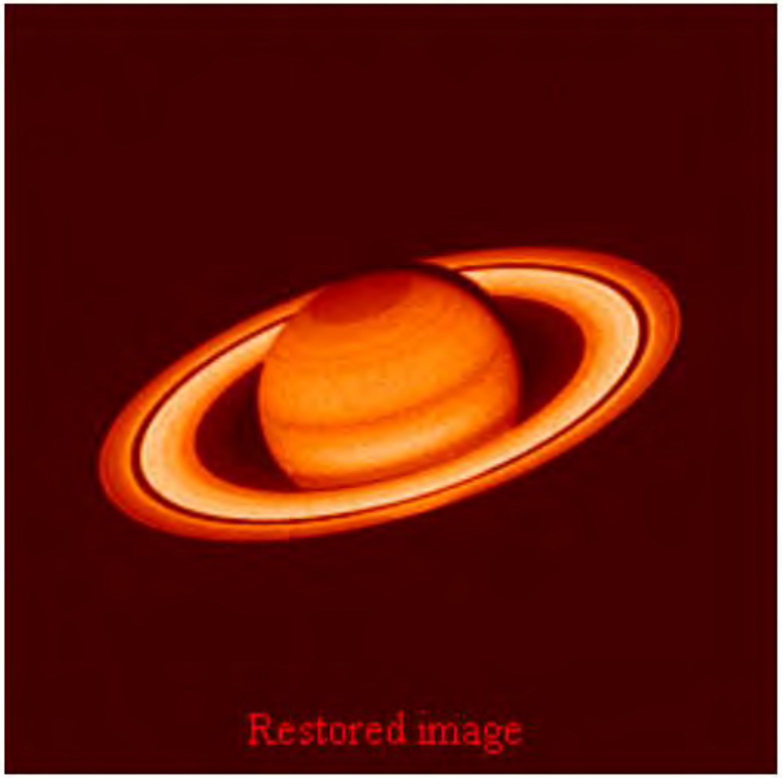

Fields that Use Digital Image Processing

Image Denosing

Noisy image$_p$eppers512rgb, PSNR: 20.176 dB (sigma: 25)

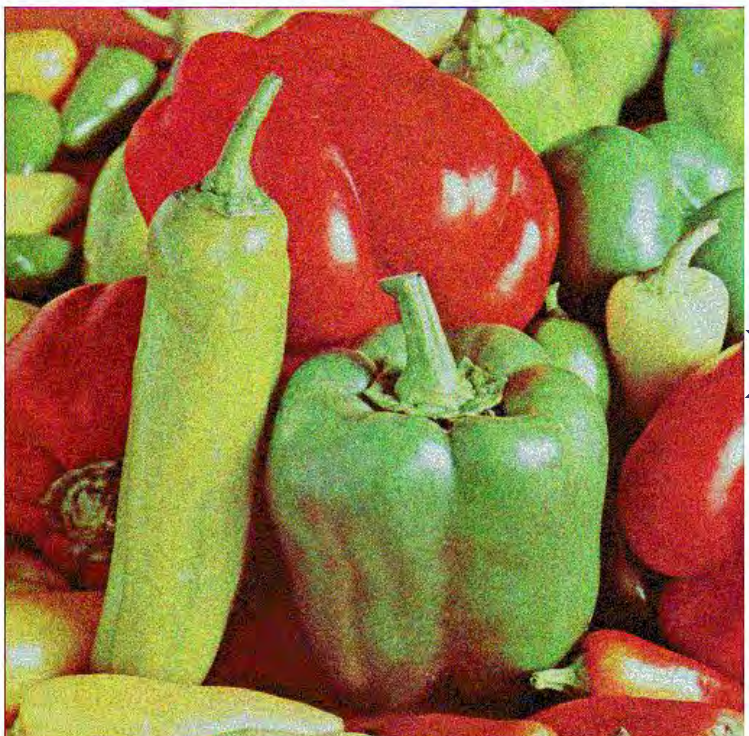

Denoised image$_p$eppers512rgb, PSNR: 31.199 dB

"Image Denoising by Sparse 3D Transform-Domain Collaborative Filtering"

Fields that Use Digital Image Processing

Noise reduction

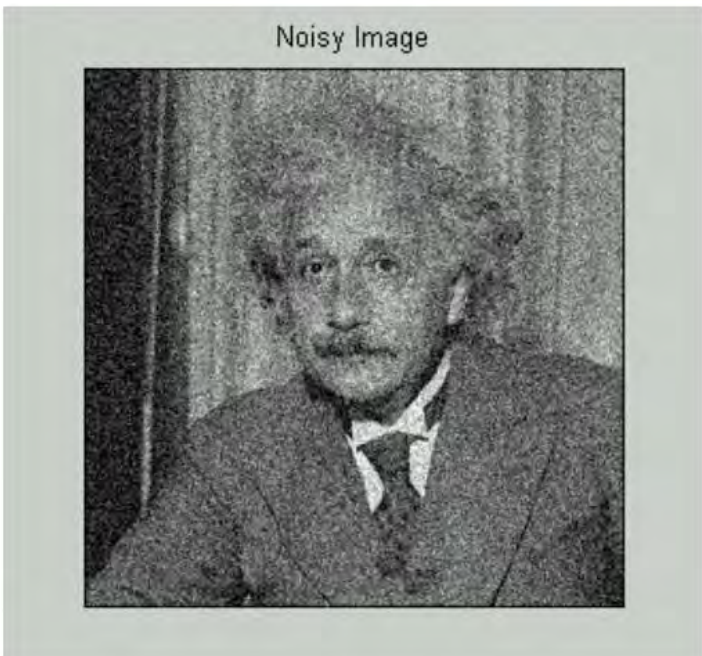

Degraded image

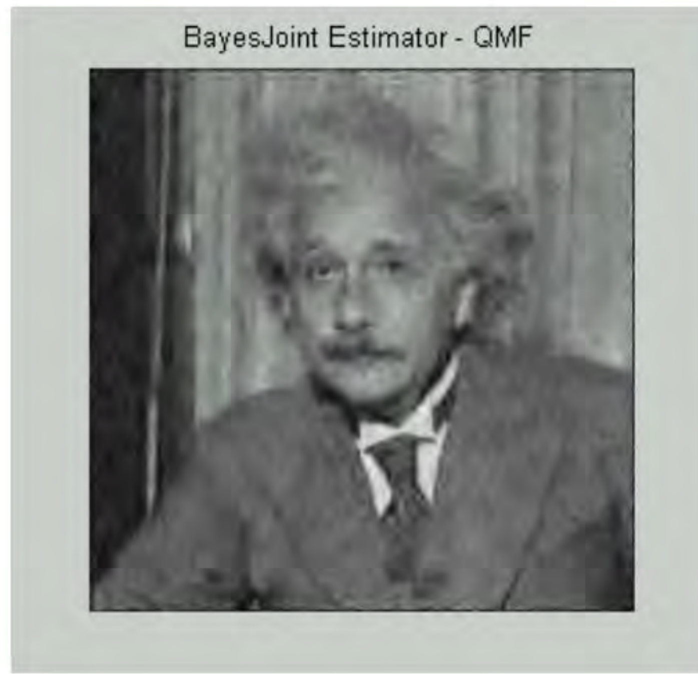

Noise-reduced image

Fields that Use Digital Image Processing

Video Denosing

“Video Denoising by Sparse 3D Transform-Domain Collaborative Filtering"

Low-level processing

original image

Canny

edge image

Middle-level processing

edge image

ORT

data structure

circular arcs and line segments

Middle-level processing

original color image

K-means clustering followed by connected component analysis

data structure

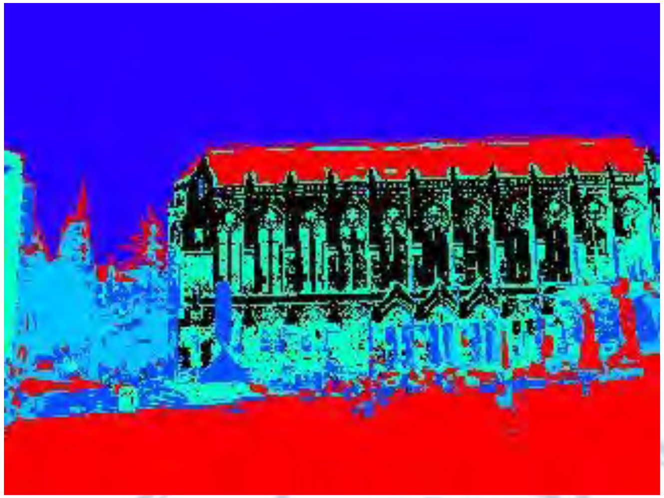

regions of homogeneous color

Low-level to high-level processing

low-level →

edge image

↓ middle-level

consistent line clusters

← high-level

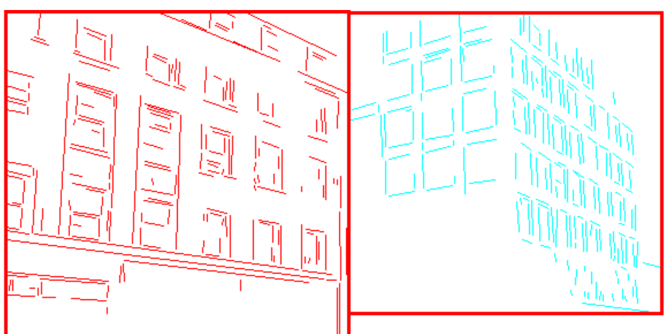

Fields that Use Digital Image Processing

Middle-level & High-level processing

- Image $\longrightarrow$ features/attributes, features $\longrightarrow$ recognition
- Image Analysis, Image Recognition, Image Comprehension
- Pattern Recognition, Computer Vision
- Difficulty
  - Computer has no intelligence
  - Machine Learning!!

Fields that Use Digital Image Processing

Cell Segmentation (2D)

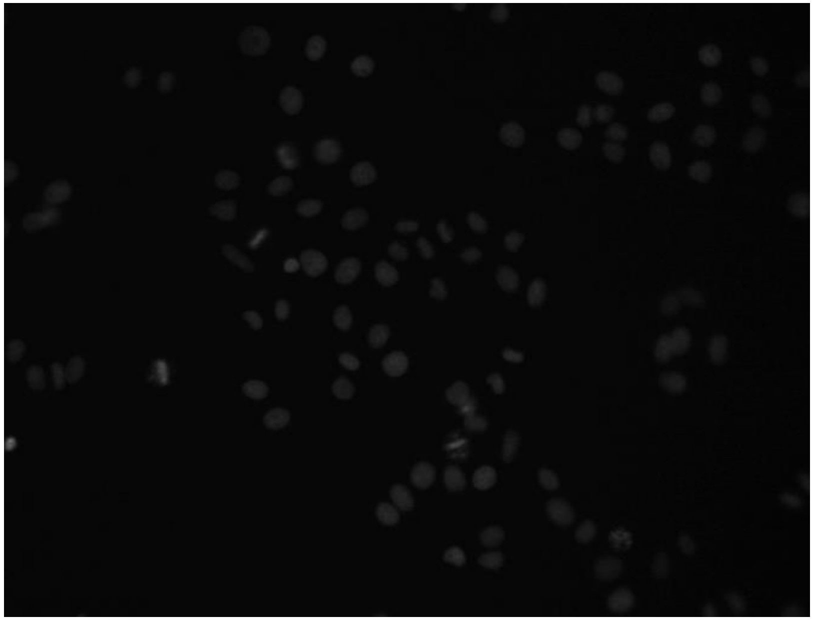

Original Image

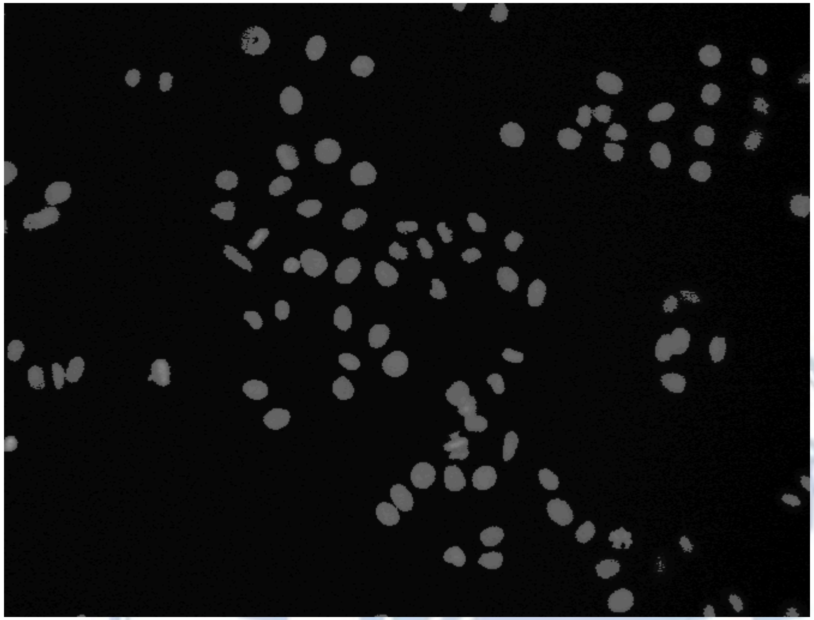

Segment Result

Fields that Use Digital Image Processing

Cell Segmentation (3D)

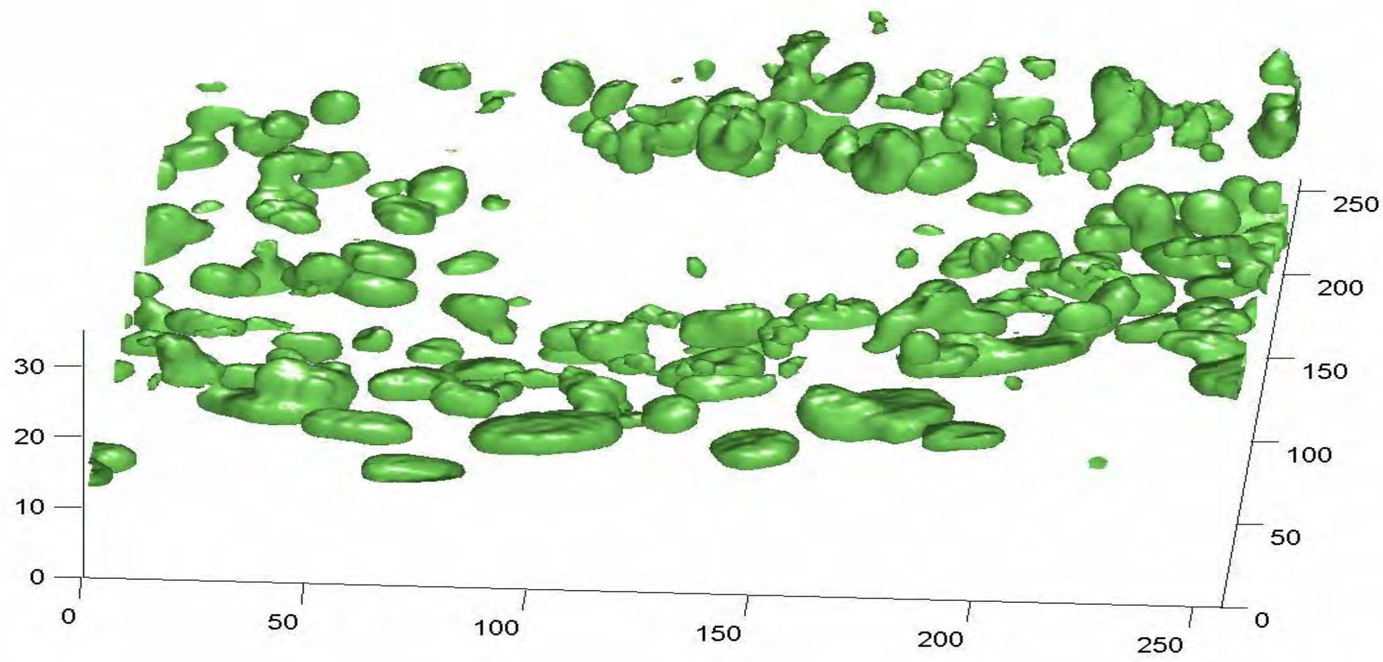

Fields that Use Digital Image Processing

Matching Result (2D)

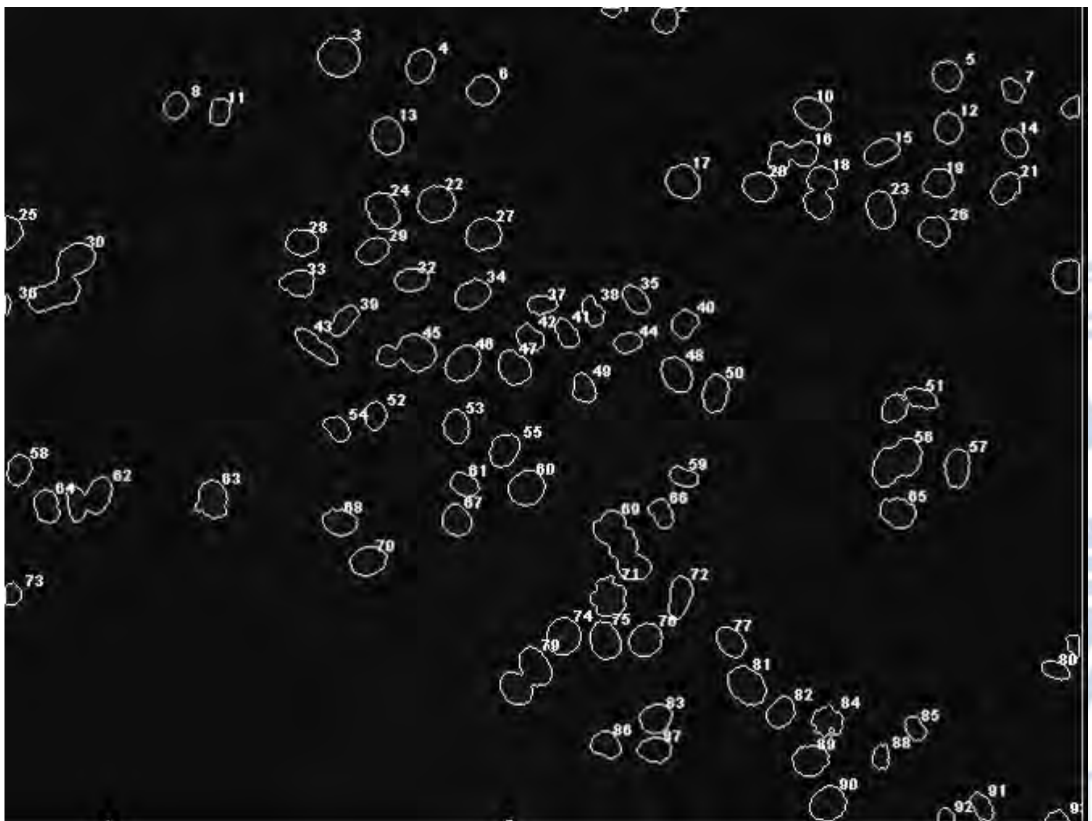

Fields that Use Digital Image Processing

Matching Result (3D)

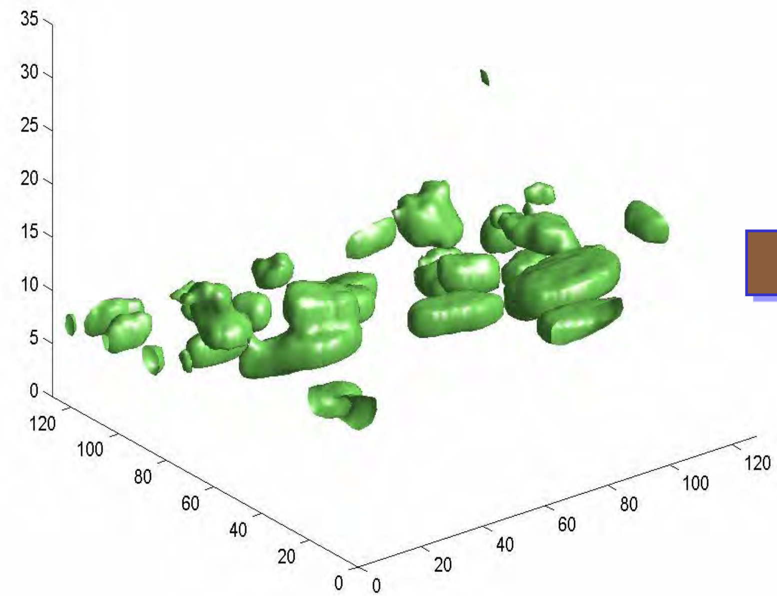

Segment Result

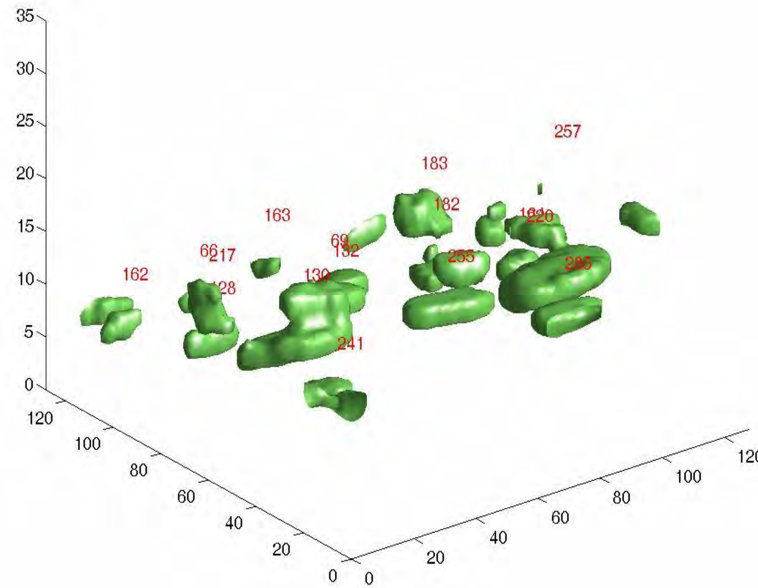

Matching Result

Fields that Use Digital Image Processing

Edge Detection

gx^2+gy^2

$gx^2+gy^2 > T$

Fields that Use Digital Image Processing

Color-Based Segmentation

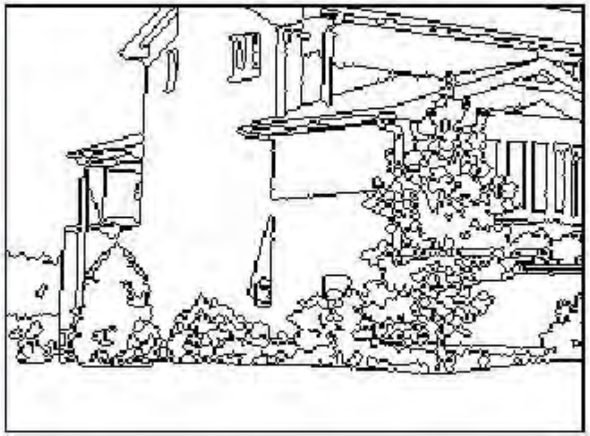

Fields that Use Digital Image Processing

Erosion

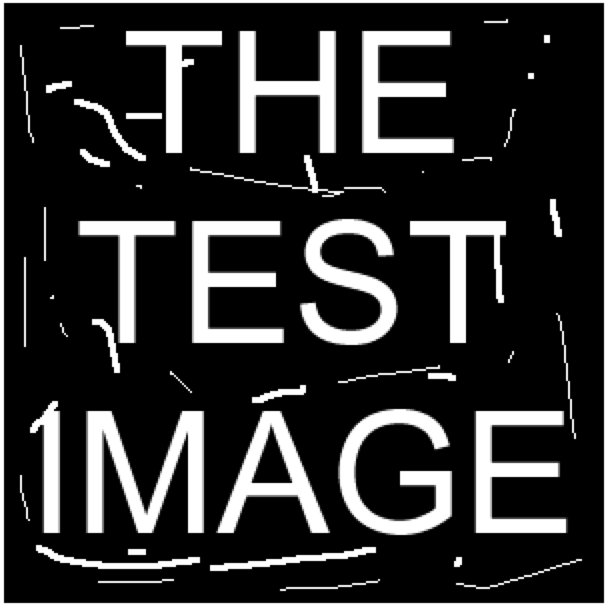

Original image

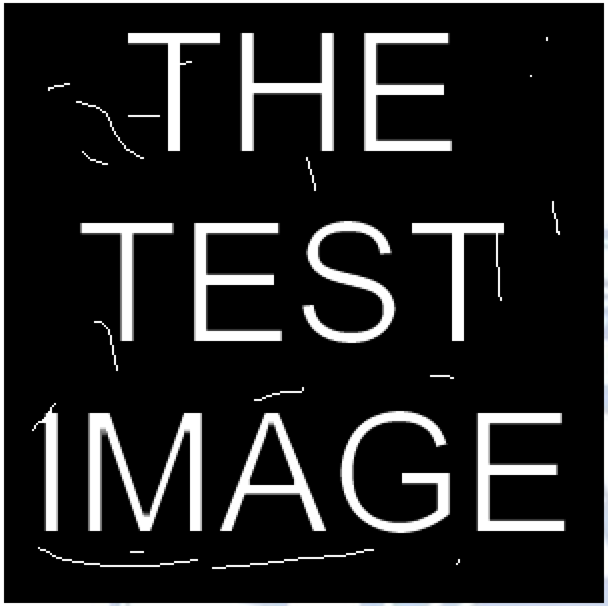

Eroded image

Fields that Use Digital Image Processing

Erosion

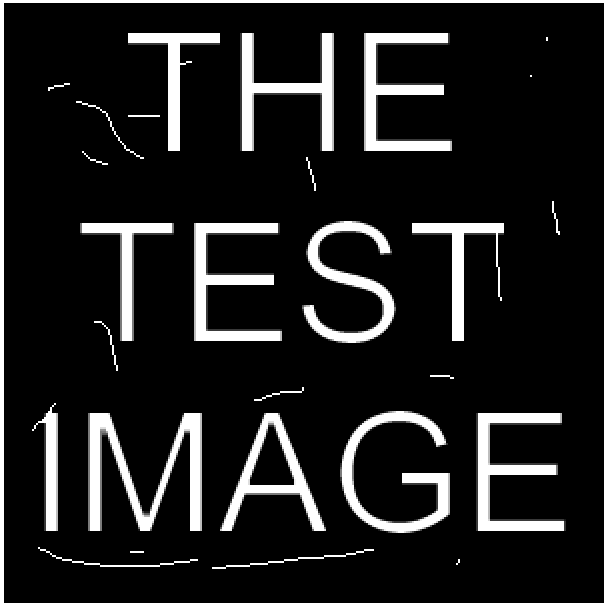

Eroded once

Eroded twice

Fields that Use Digital Image Processing

Vision-based biometrics

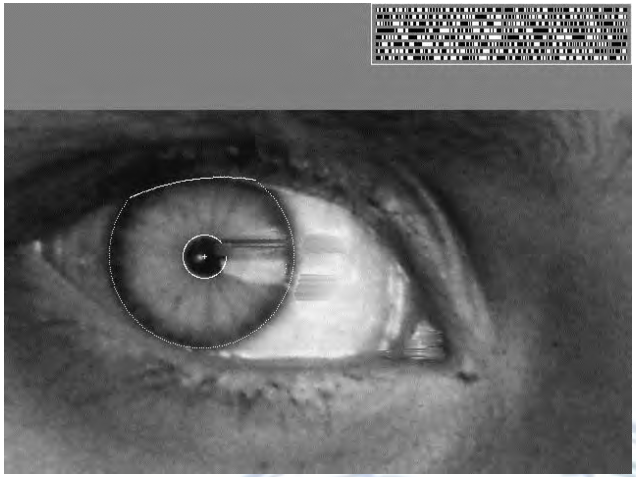

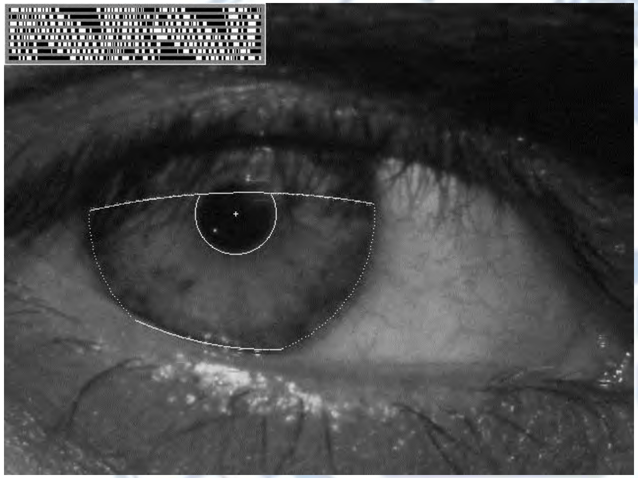

The Afghan Girl Identified by Her Iris Patterns

Fields that Use Digital Image Processing

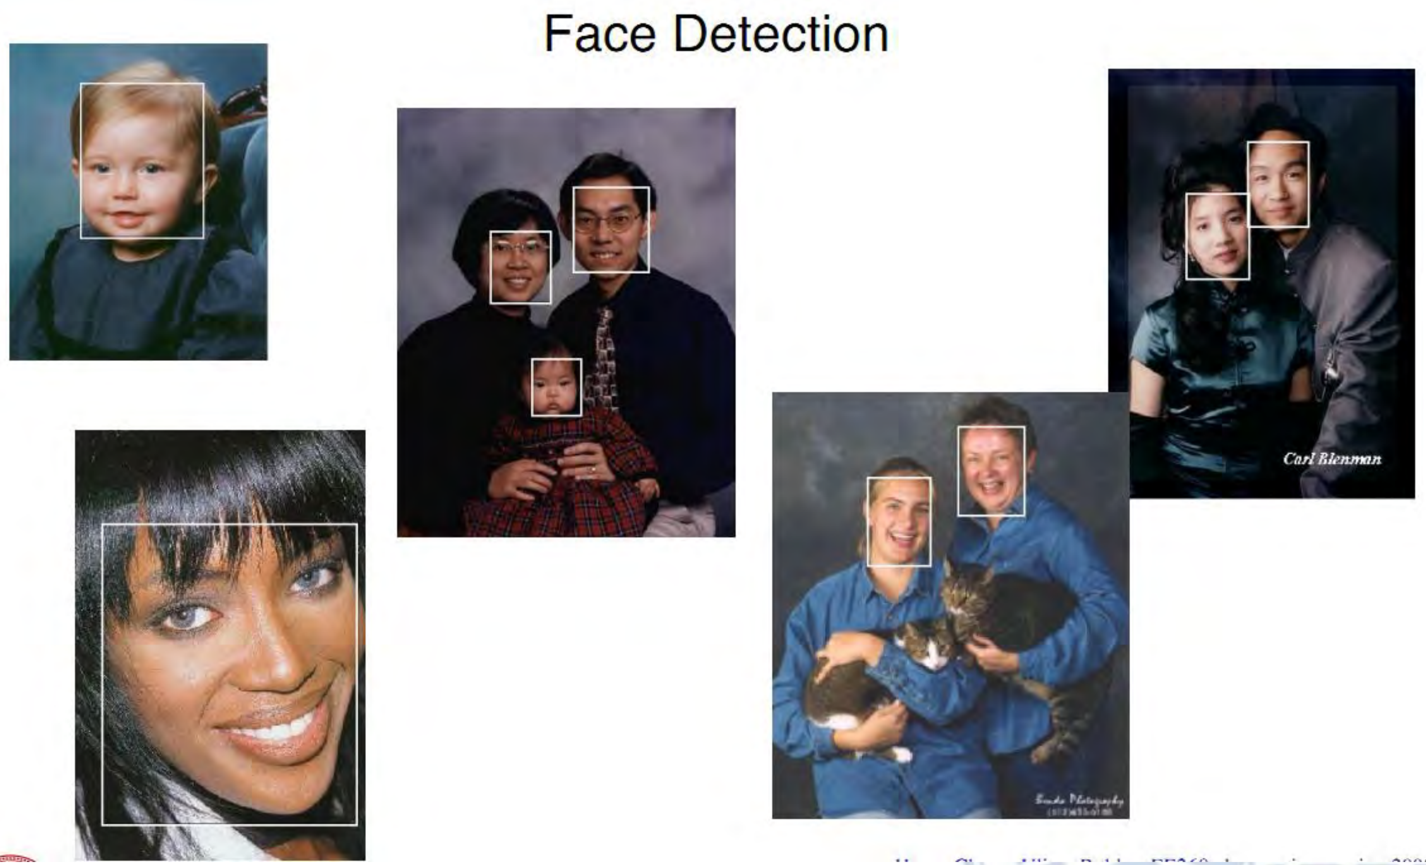

Fields that Use Digital Image Processing

Fields that Use Digital Image Processing

Surveillance and tracking

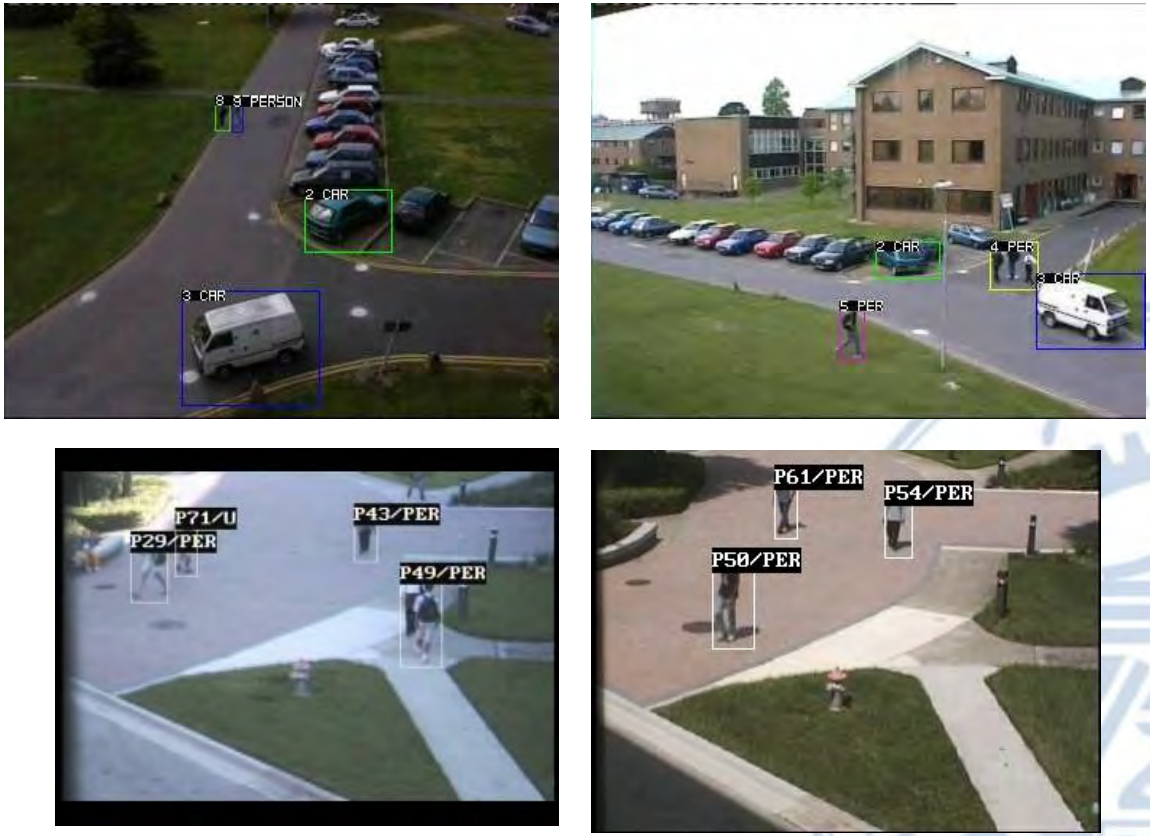

Fields that Use Digital Image Processing

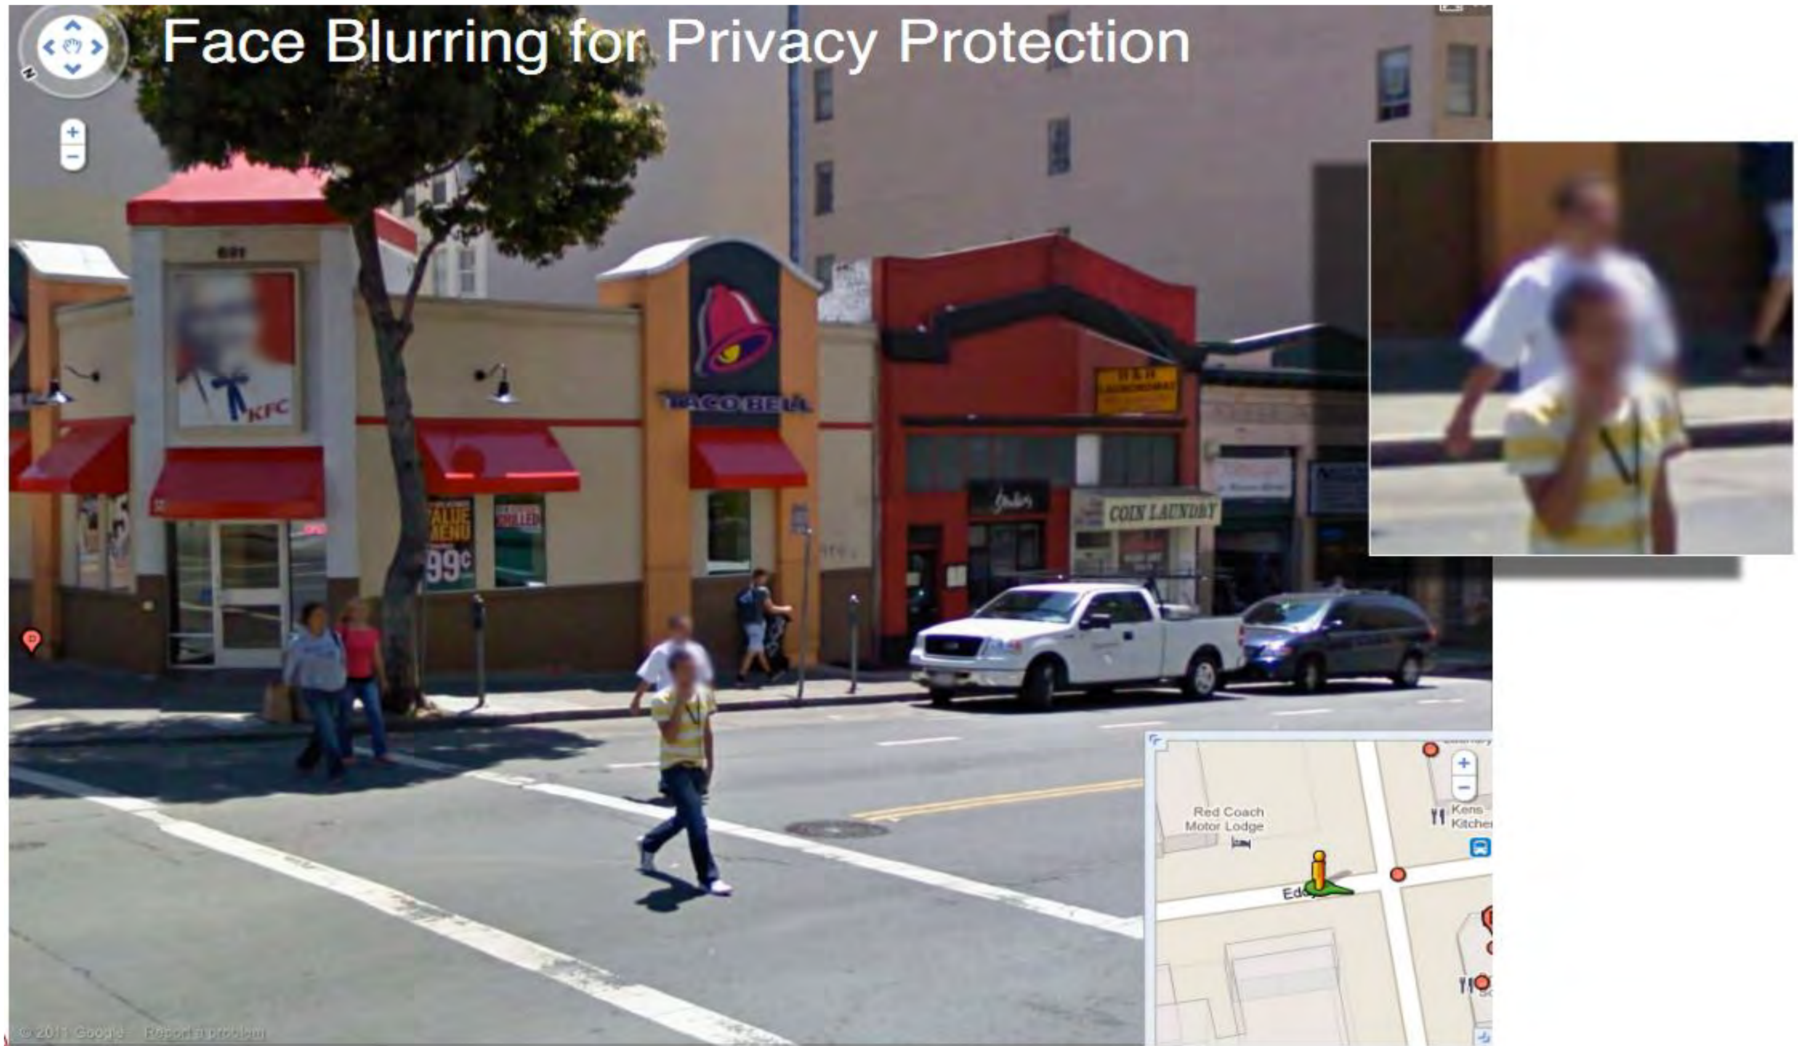

Fields that Use Digital Image Processing

Augmented reality

Fields that Use Digital Image Processing

Content-based retrieval

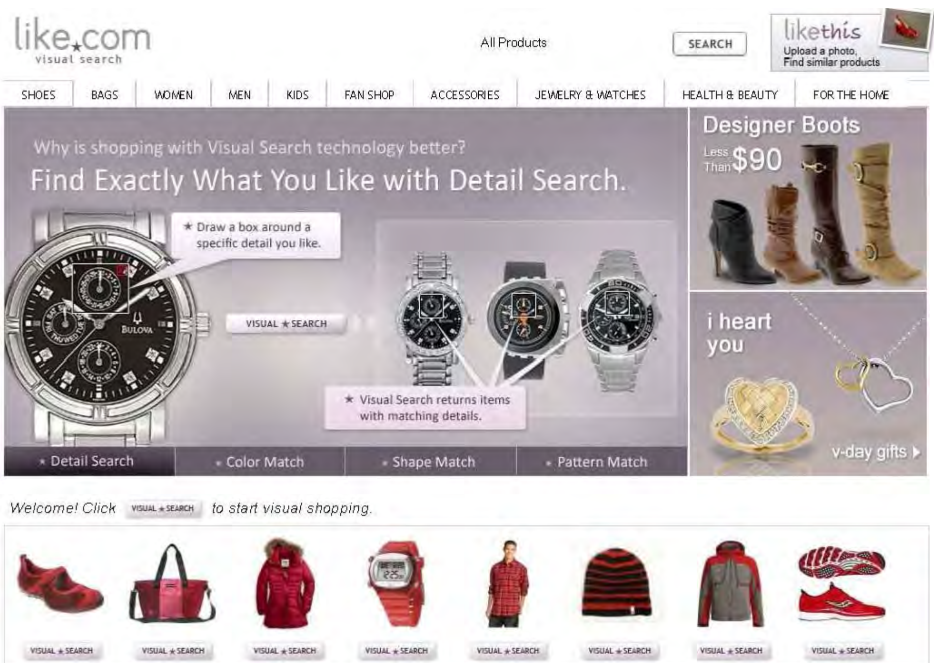

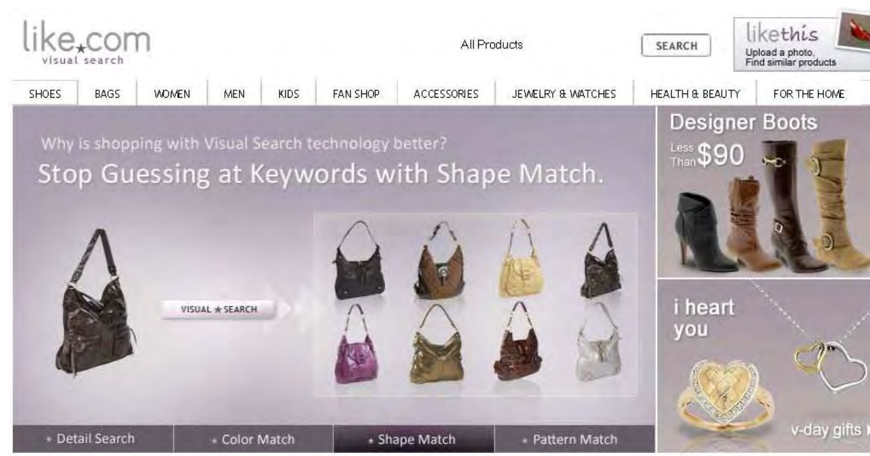

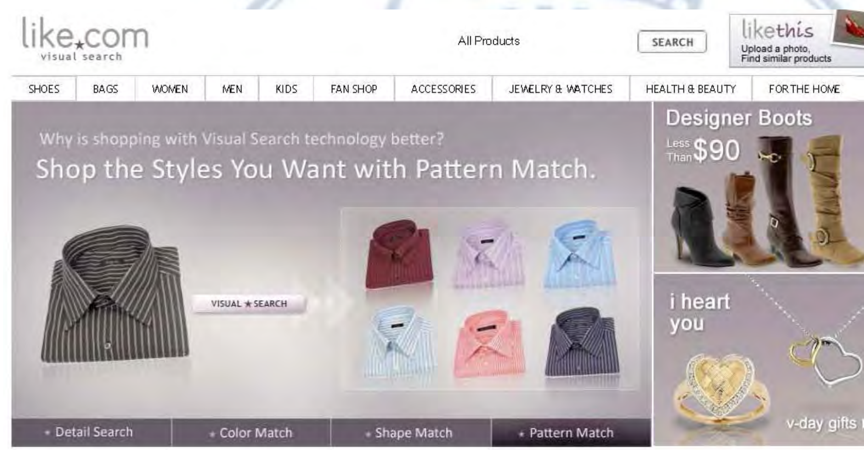

Online shopping catalog search

Fields that Use Digital Image Processing

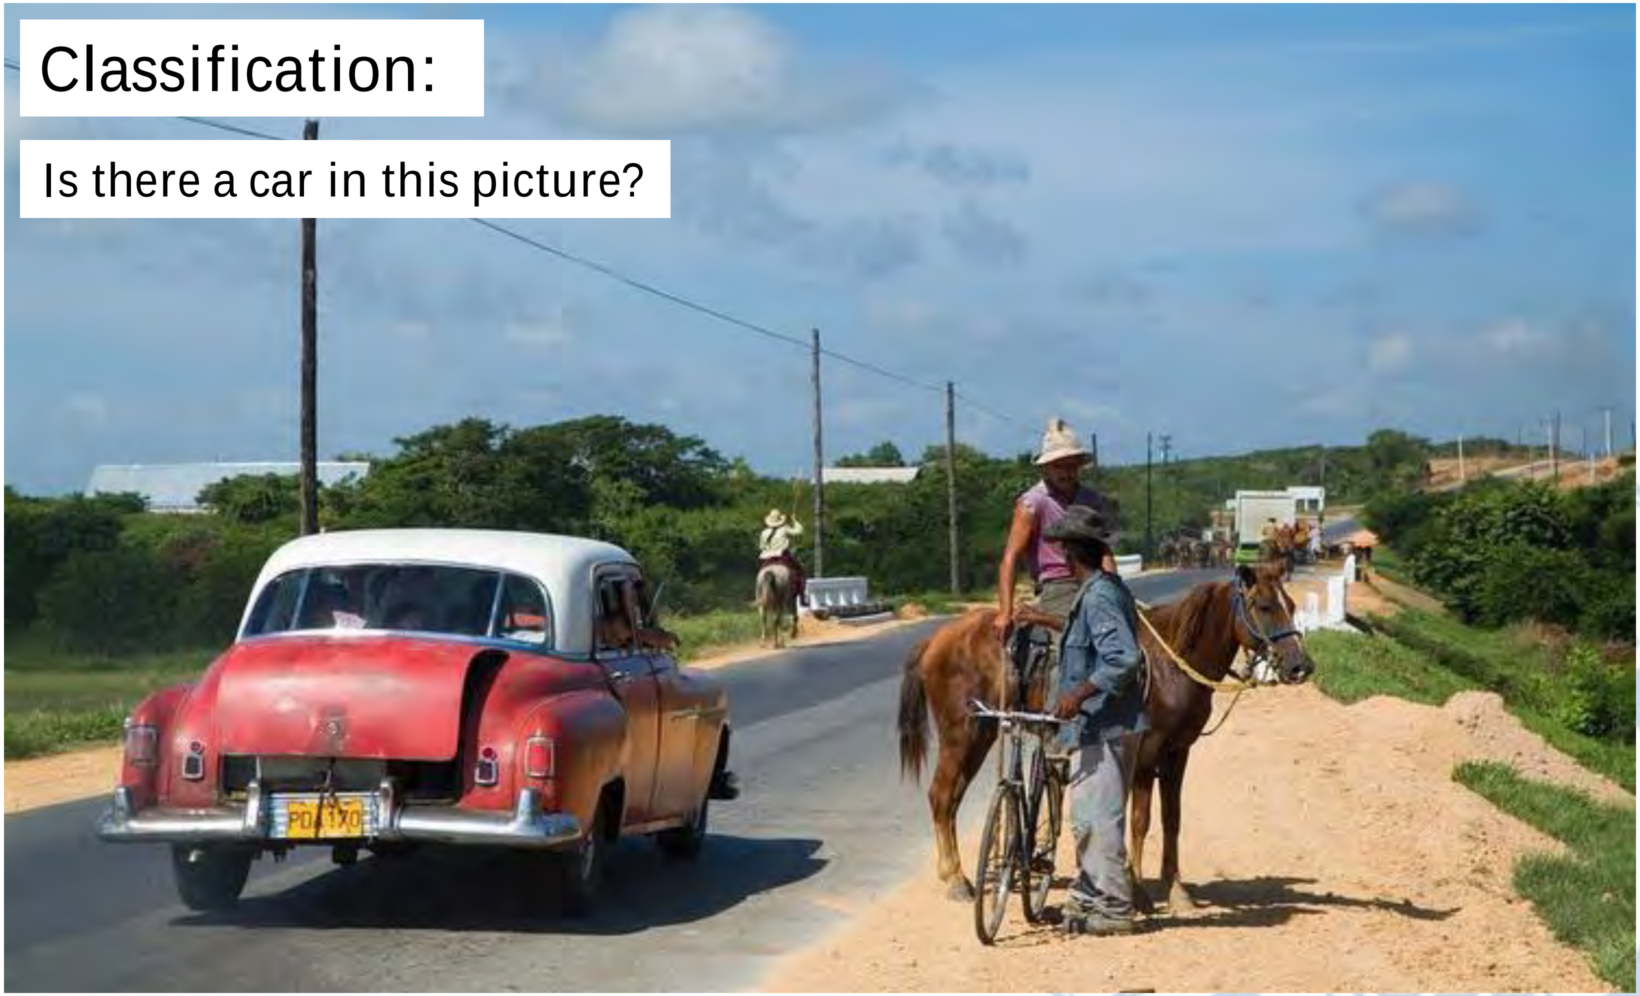

Fields that Use Digital Image Processing

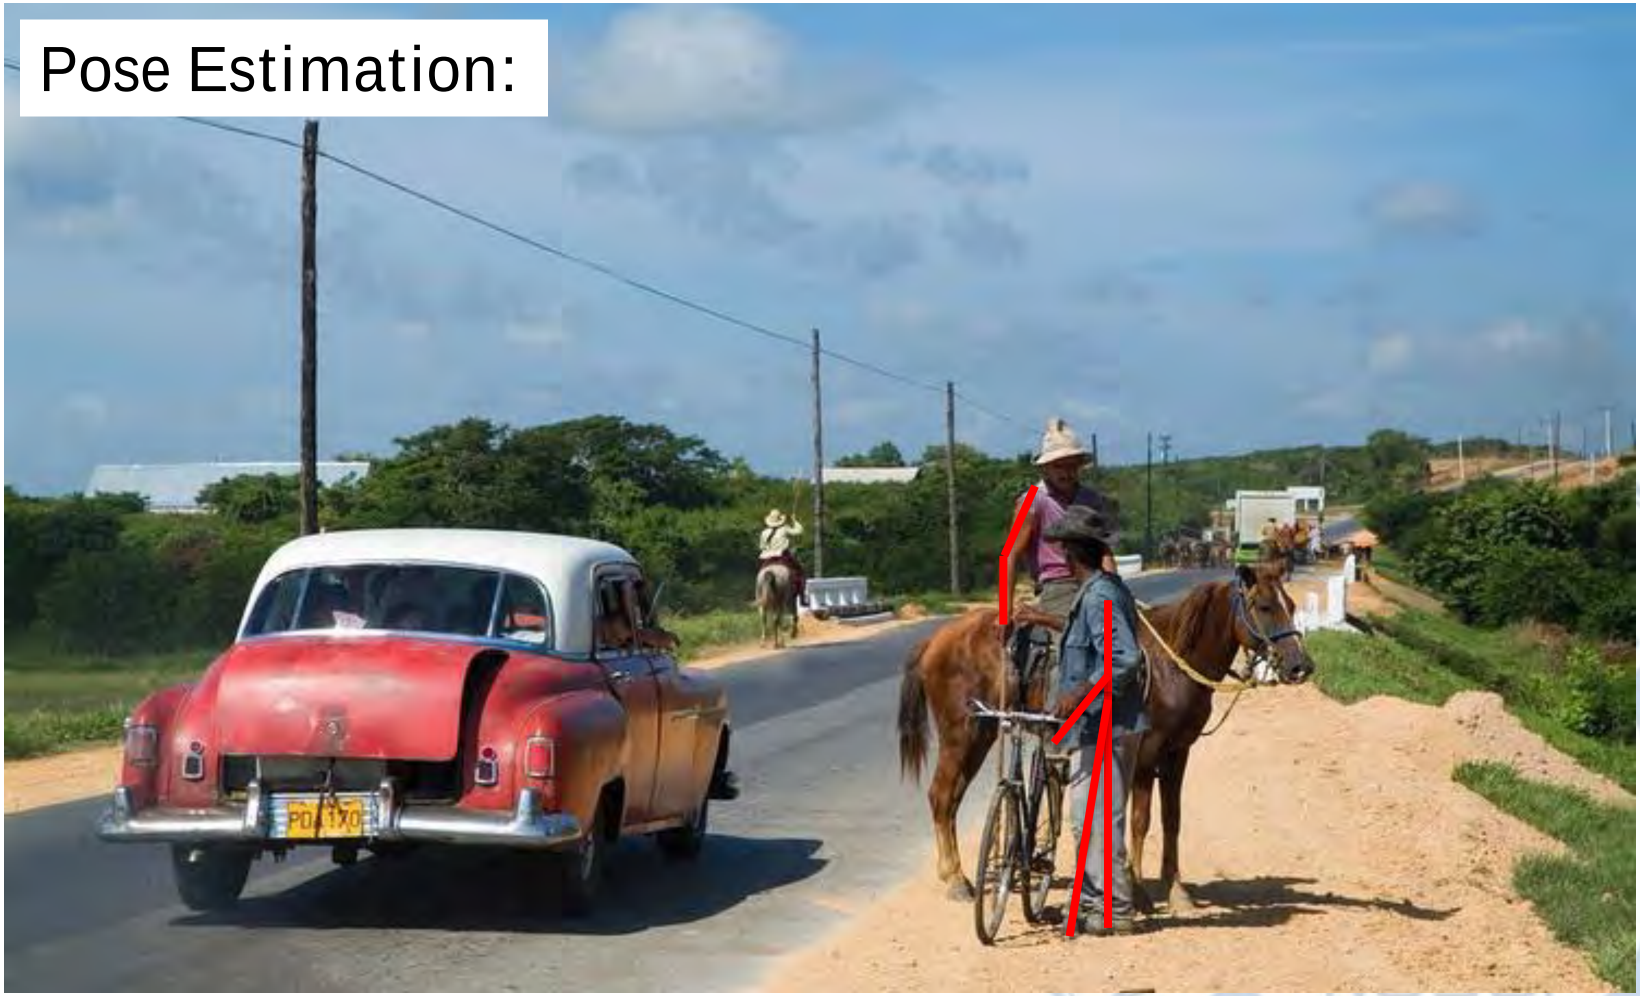

Fields that Use Digital Image Processing

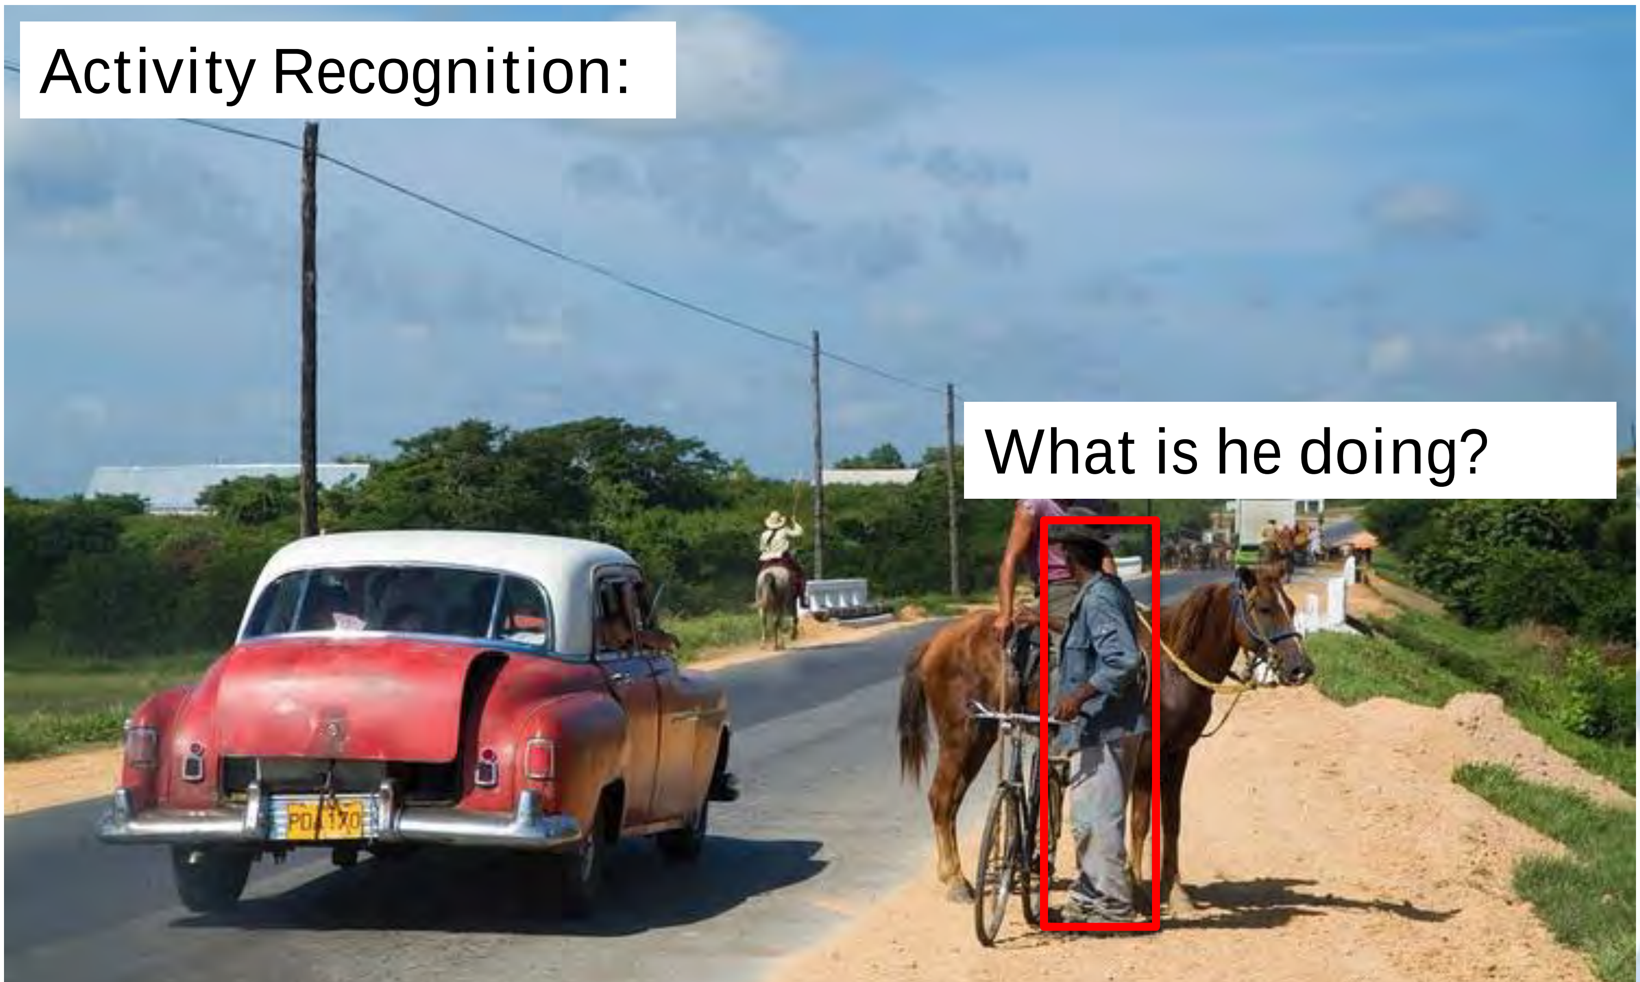

Fields that Use Digital Image Processing

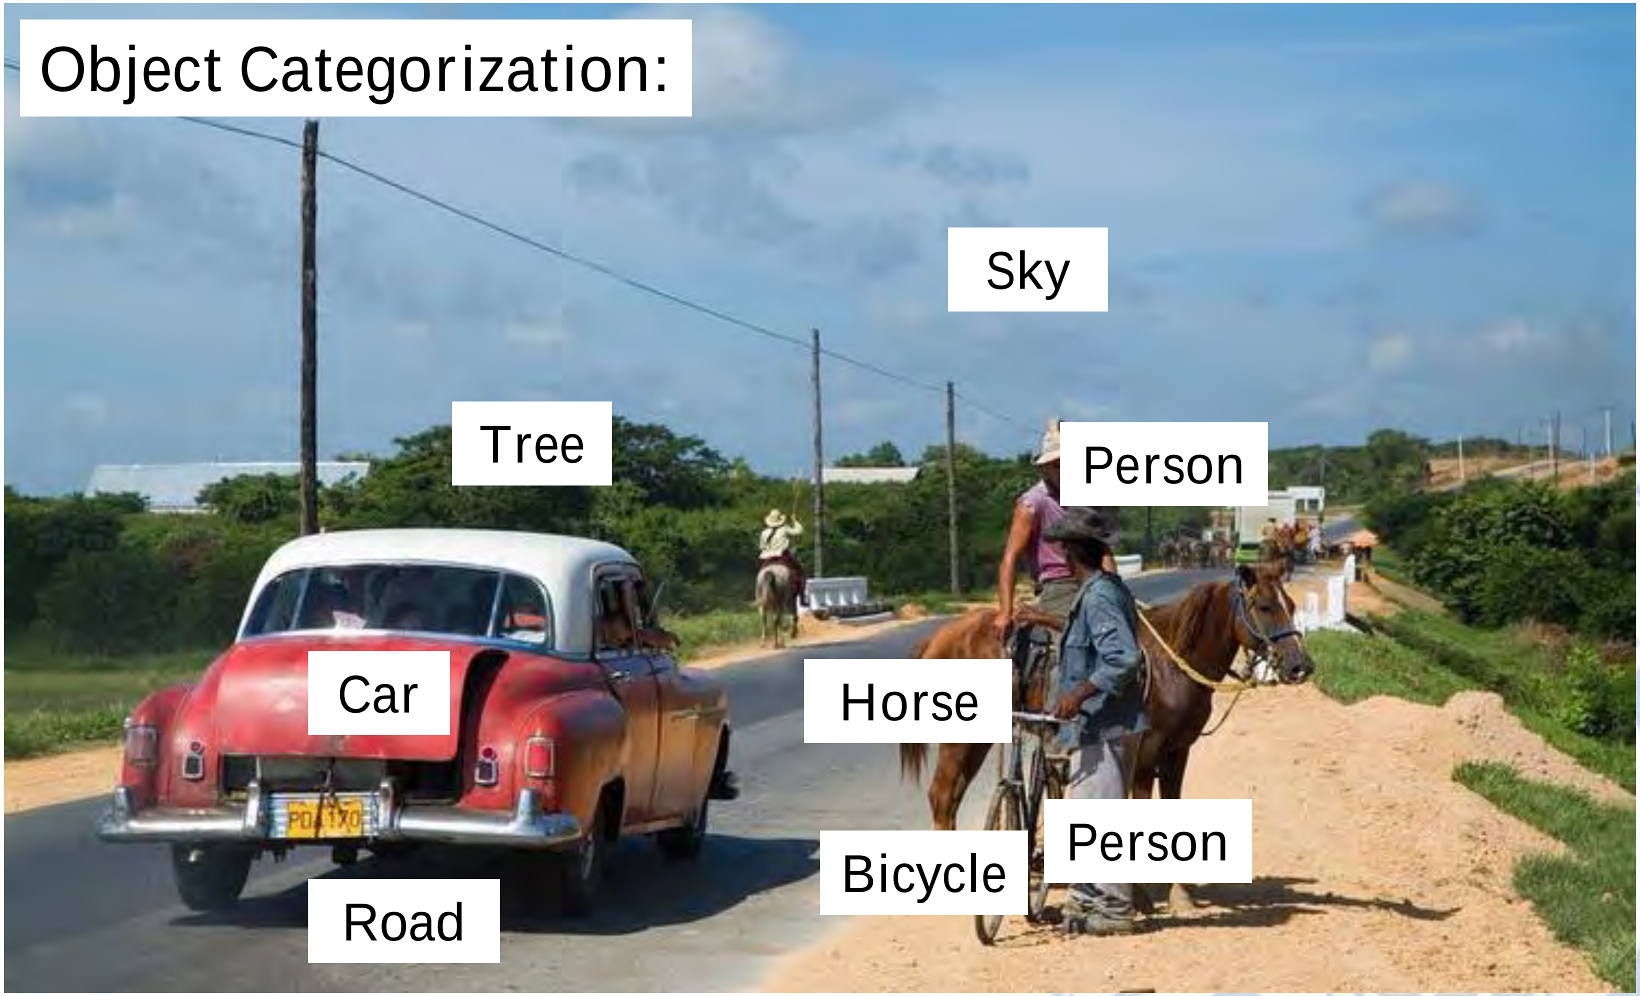

Fields that Use Digital Image Processing

Public security
- Video surveillance system
- Human face recognition & tracking
- Fingerprint enhancement & recognition

Traffic
- Car license plate recognition
- Vehicle recognition
- Electronic police

Universe exploration
- Airship
- Moon exploration

Telemetry
- Weather forecast
- Mineral resources detection

Fields that Use Digital Image Processing

National Defense
- Pilotless aircraft
- Cruise missile

Biomedicine
- CT
- MRI

Other
- Mobile phone
- Digital camera
- Digital recorder
- VOD
- MSN
- ...

Key Stages in Digital Image Processing

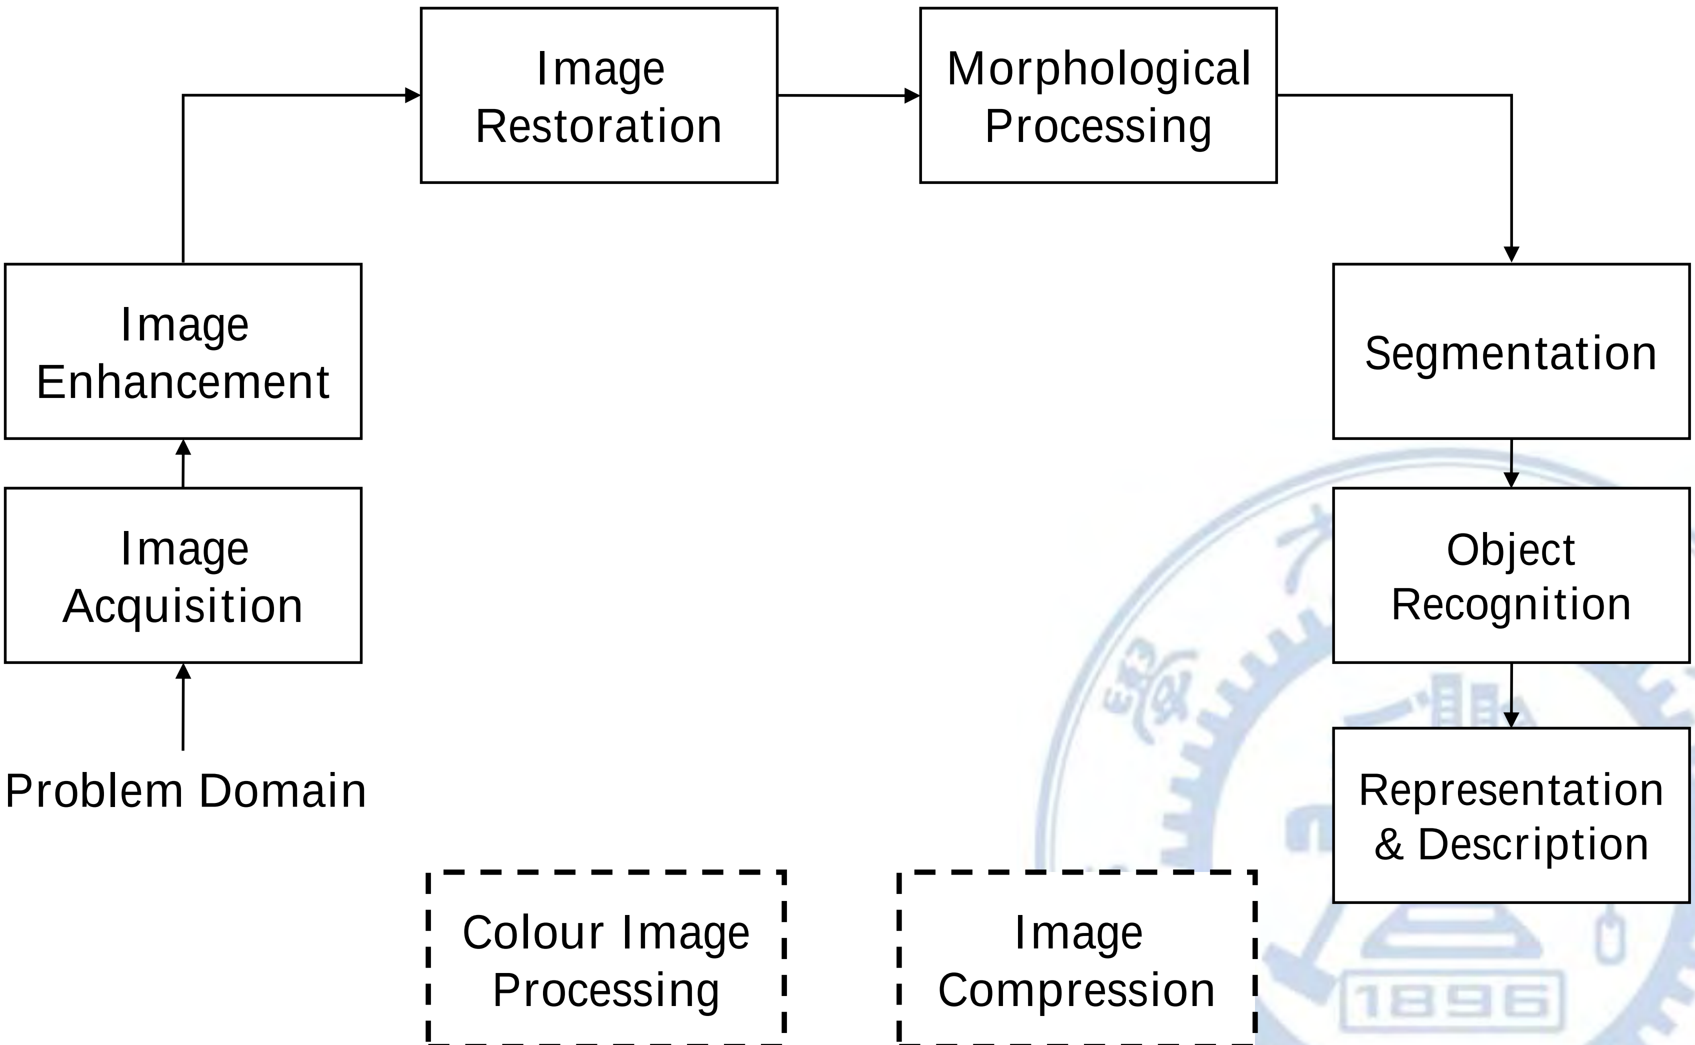

Key Stages in Digital Image Processing: Image Aquisition

Image Restoration

Morphological Processing

Image Enhancement

Segmentation

Image Acquisition

Object Recognition

Representation & Description

Problem Domain

Colour Image Processing

Image Compression

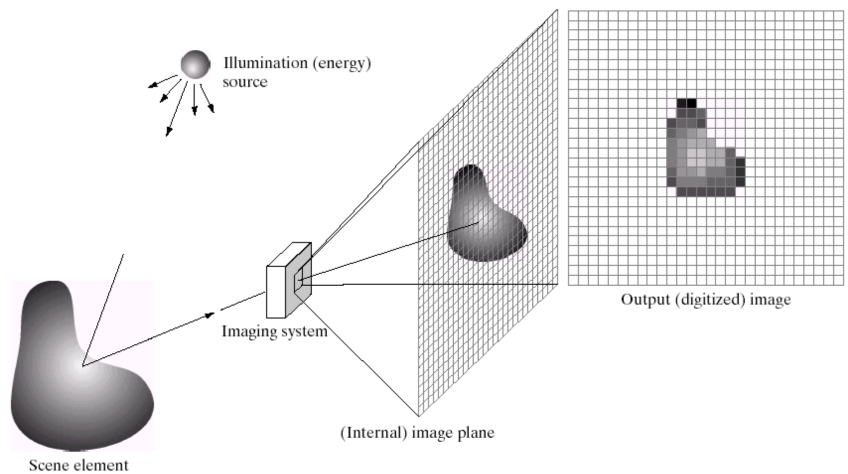

Key Stages in Digital Image Processing: Image Enhancement

Image Restoration

Morphological Processing

Image Enhancement

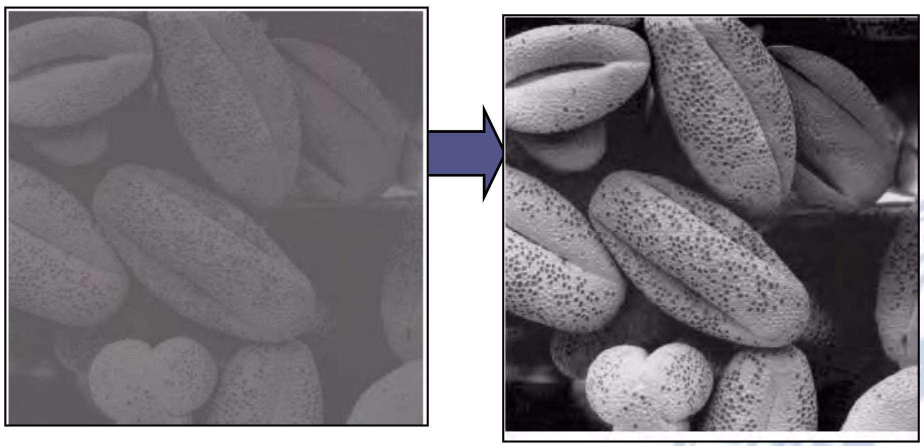

Segmentation

Image Acquisition

Object Recognition

Problem Domain

Representation & Description

Colour Image Processing

Image Compression

Key Stages in Digital Image Processing: Image Restoration

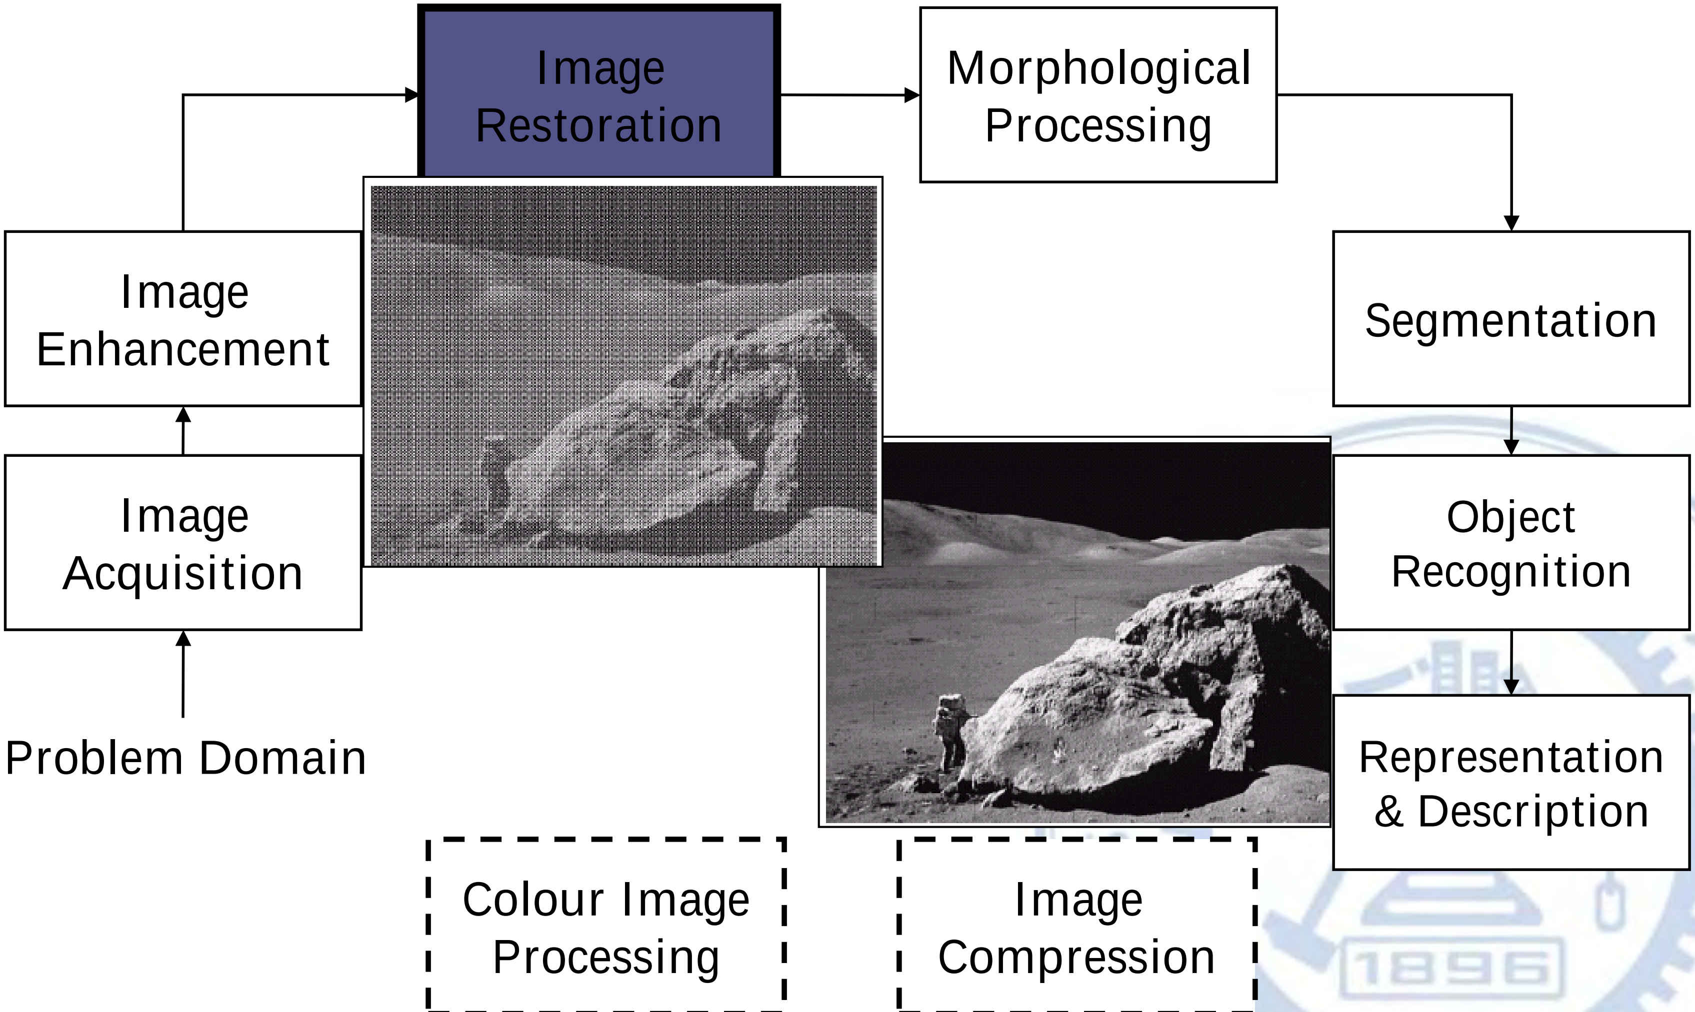

Key Stages in Digital Image Processing: Morphological Processing

Image Acquisition → Image Enhancement → Image Restoration → **Morphological Processing** → Segmentation → Object Recognition → Representation & Description

Problem Domain

Colour Image Processing

Image Compression

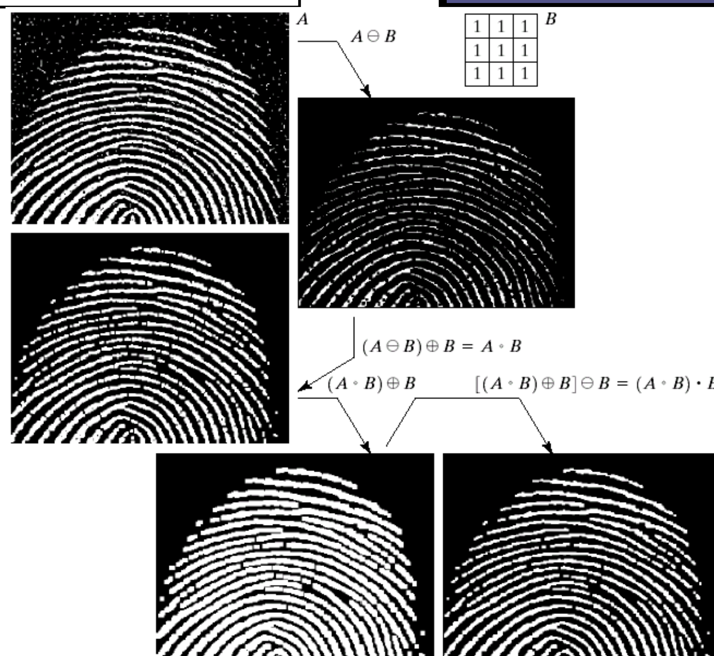

Key Stages in Digital Image Processing: Segmentation

Problem Domain → Image Acquisition → Image Enhancement → Image Restoration → Morphological Processing → **Segmentation** → Object Recognition → Representation & Description

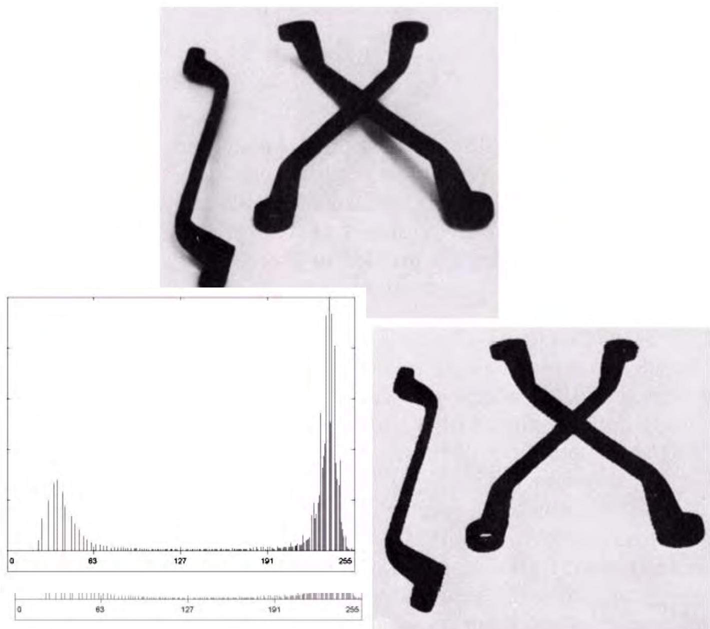

Colour Image Processing

Image Compression

Key Stages in Digital Image Processing: Object Recognition

Image Restoration

Morphological Processing

Image Enhancement

Segmentation

Image Acquisition

Object Recognition

Representation & Description

Problem Domain

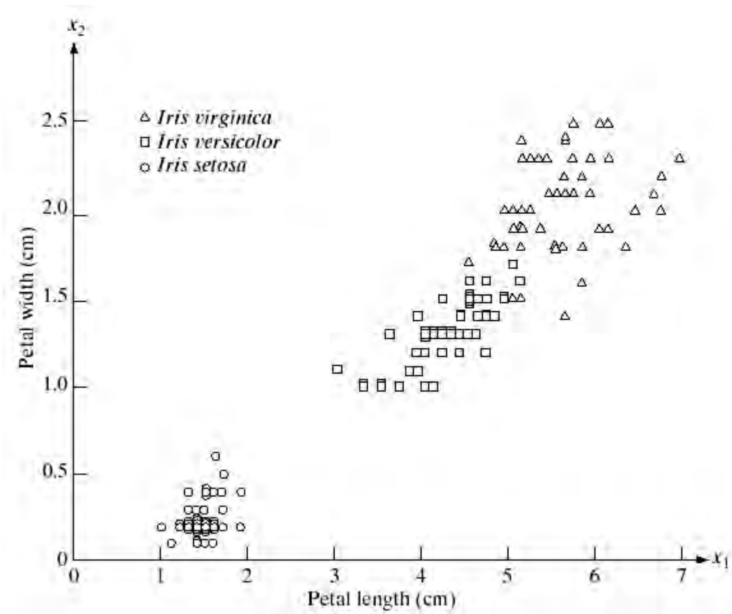

Colour Image Processing

Image Compression

Key Stages in Digital Image Processing: Representation & Description

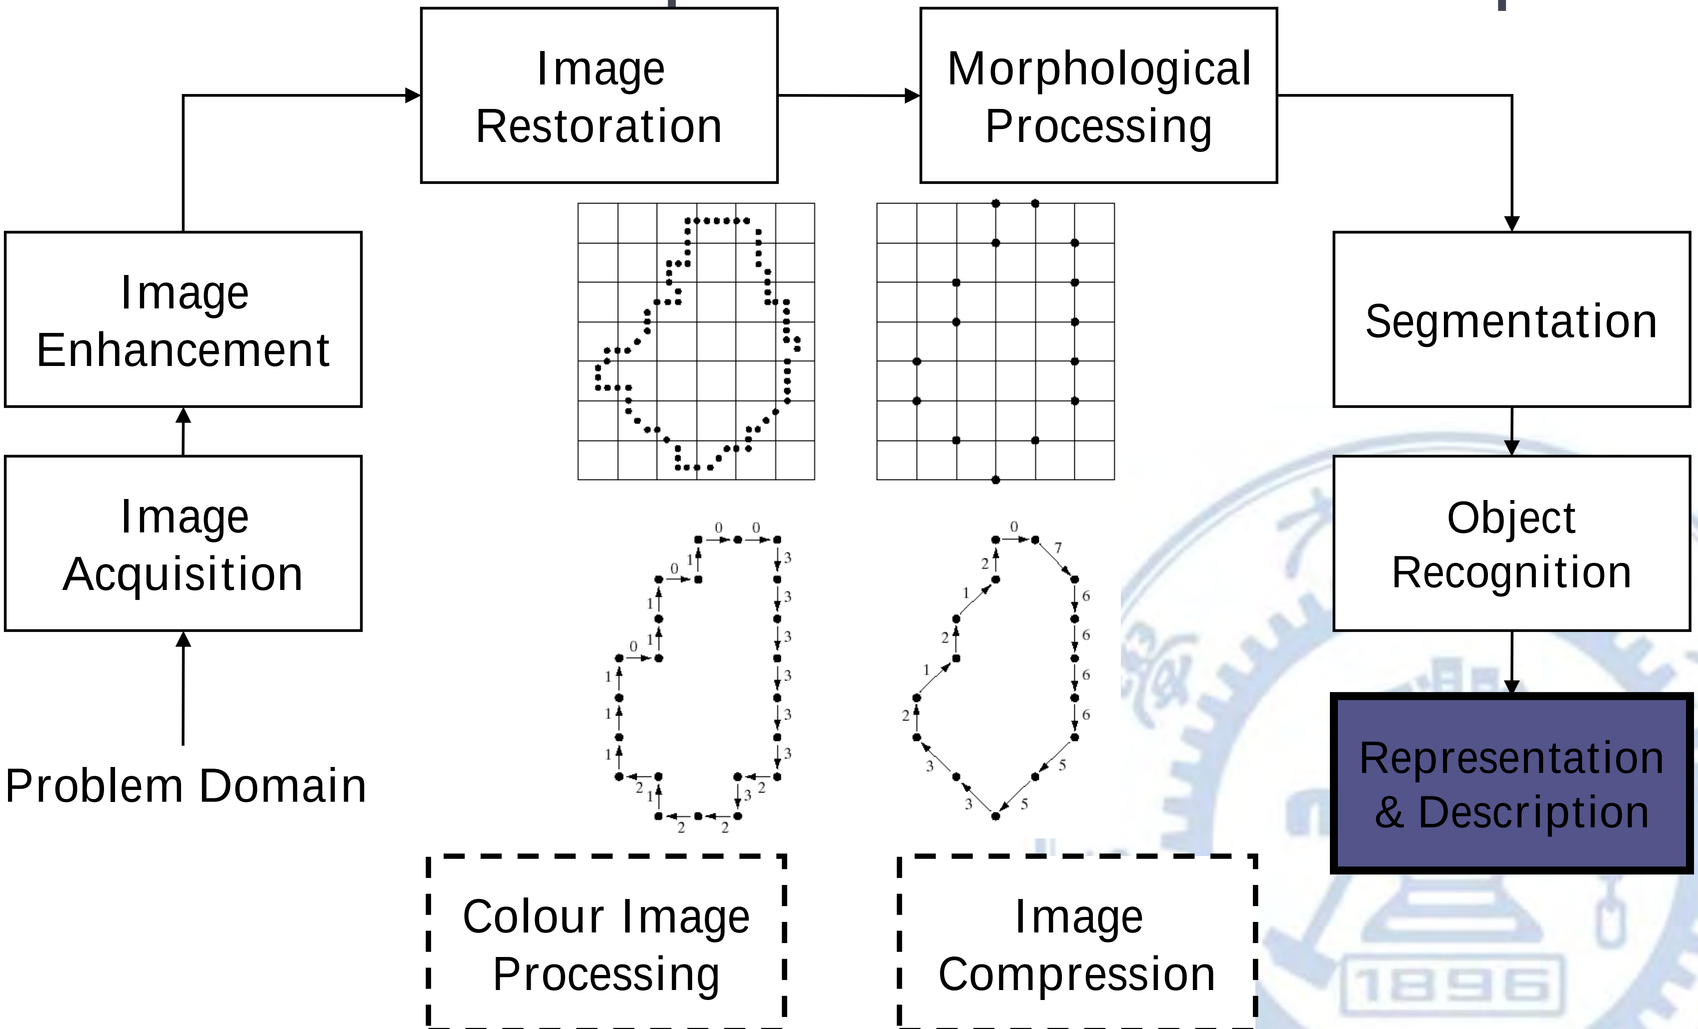

Key Stages in Digital Image Processing: Image Compression

Image Restoration

Morphological Processing

Image Enhancement

Segmentation

Image Acquisition

Object Recognition

Representation & Description

Problem Domain

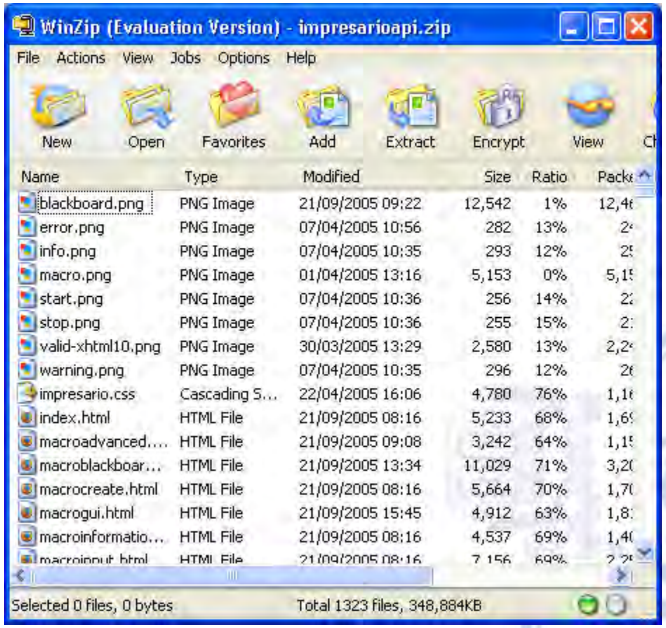

Colour Image Processing

Image Compression

Key Stages in Digital Image Processing: Colour Image Processing

Image Restoration

Morphological Processing

Image Enhancement

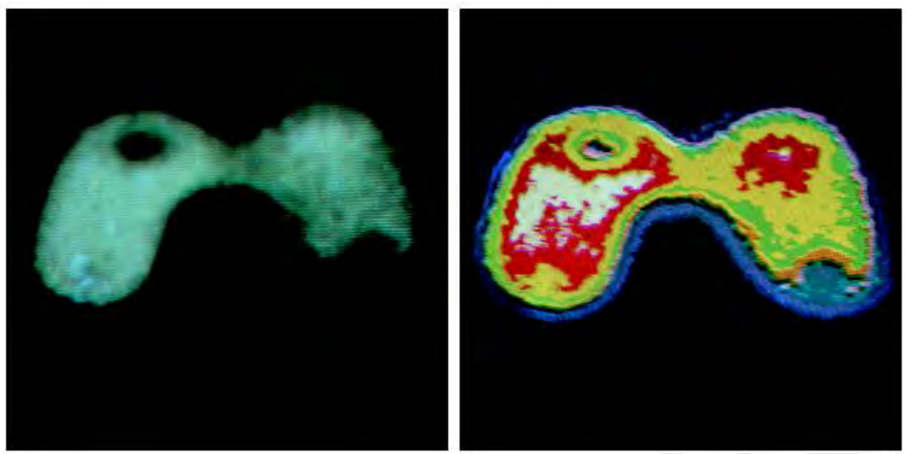

Segmentation

Image Acquisition

Object Recognition

Representation & Description

Problem Domain

Colour Image Processing

Image Compression

Something Cool!!!

Camera rotations with homographies (Single View)

St.Petersburg
photo by A. Tikhonov

Virtual camera rotations

Something Cool!!!

Stereo

- Input Images:

Something Cool!!!

- User select edges and corners

Something Cool!!!

Camera Position and Orientation

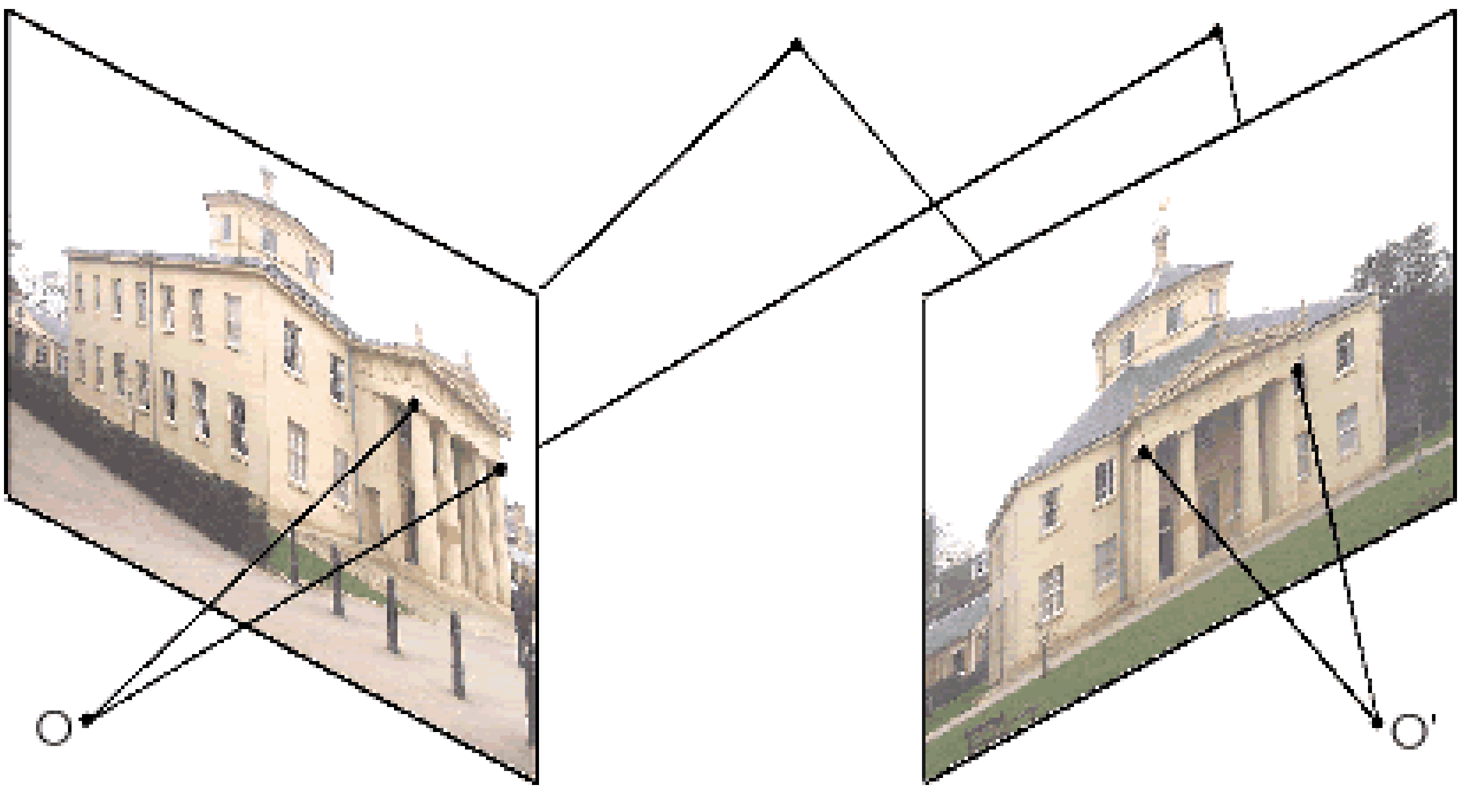

Something Cool!!!

Compute 3D textured triangles

Something Cool!!!

Panoramas

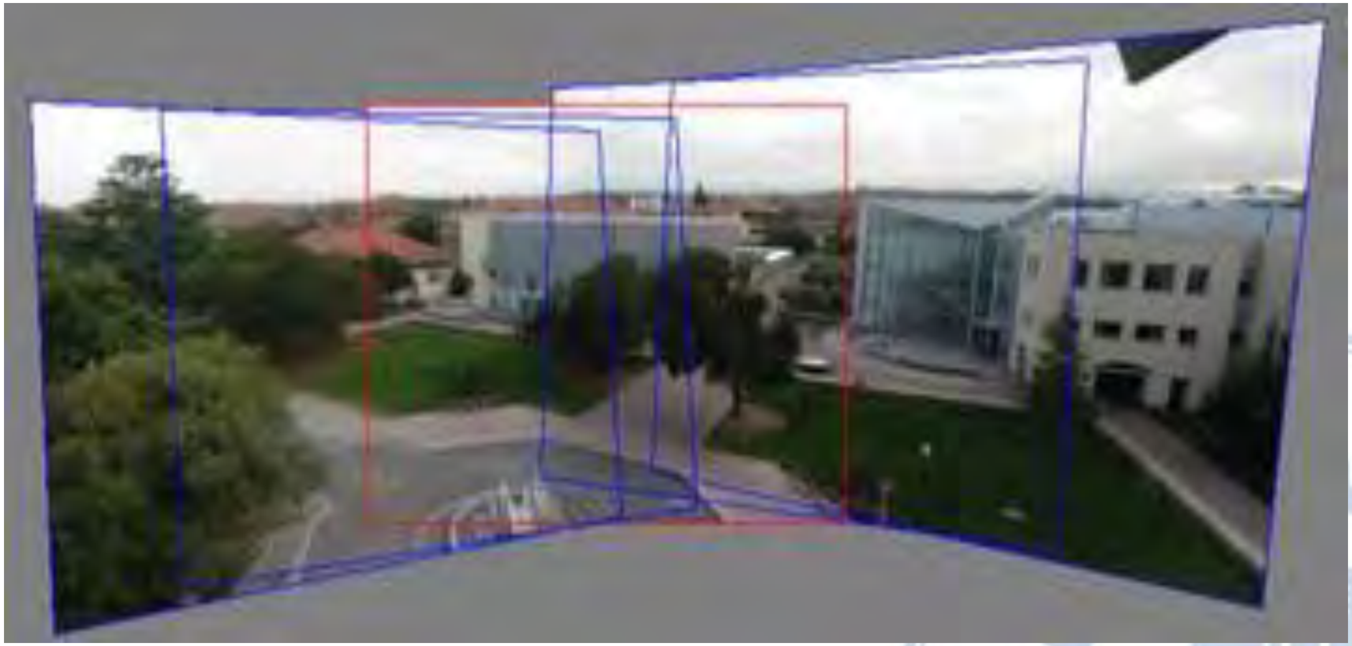

1. Pick one image (red)
2. Warp the other images towards it (usually, one by one)
3. blend

3D Applications

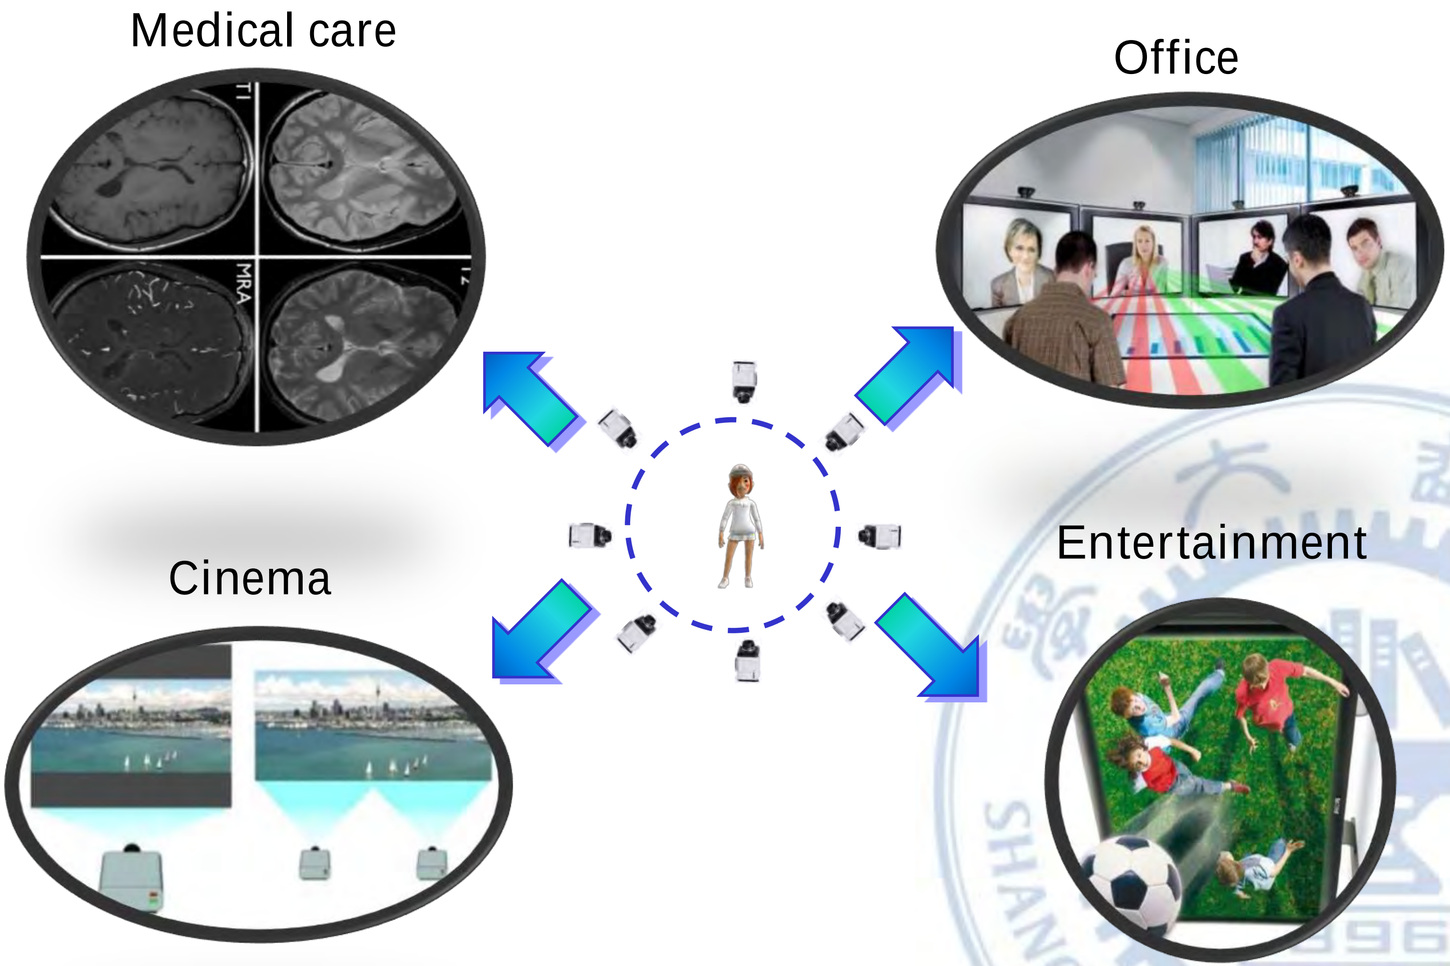

3D Video

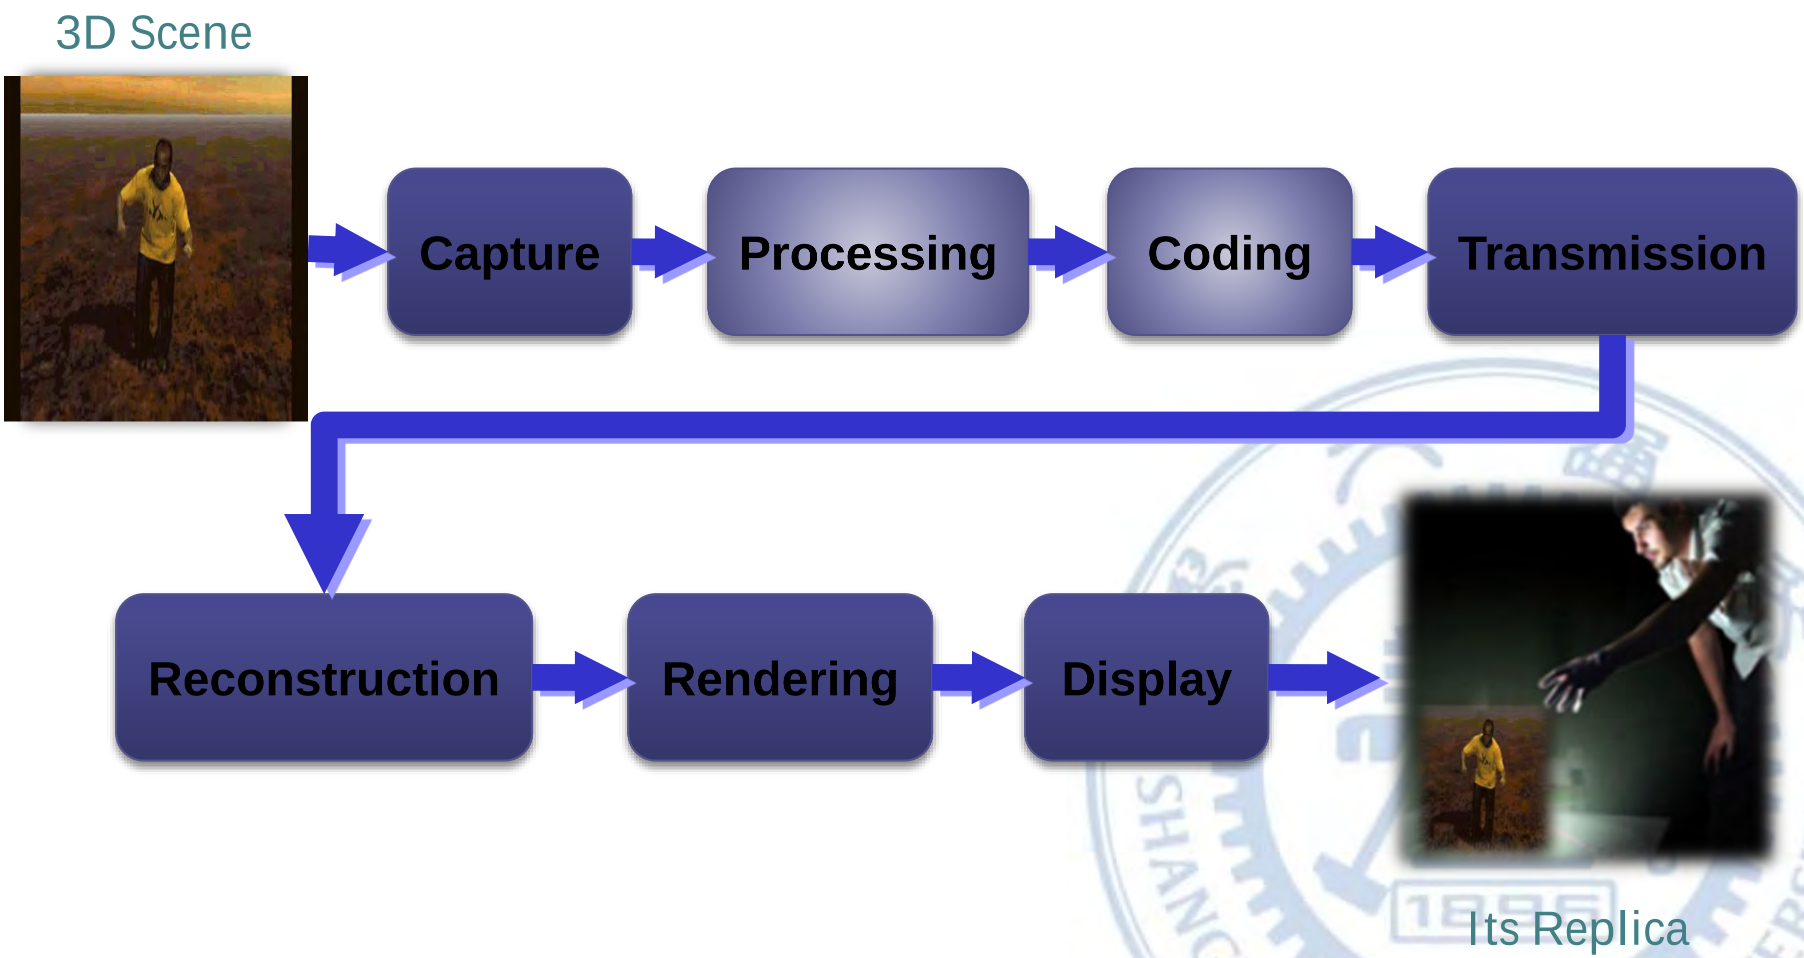
3D Scene
Capture
Processing
Coding
Transmission
Reconstruction
Rendering
Display
Its Replica

3D Data Capture

CT / MRI scanner

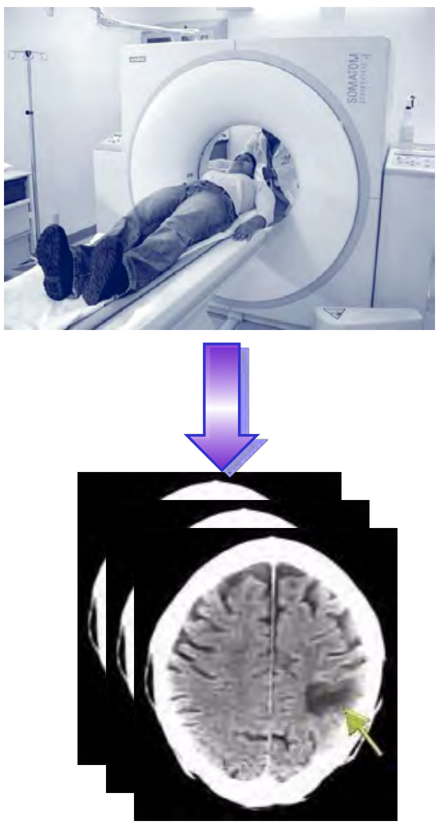

Multi-view

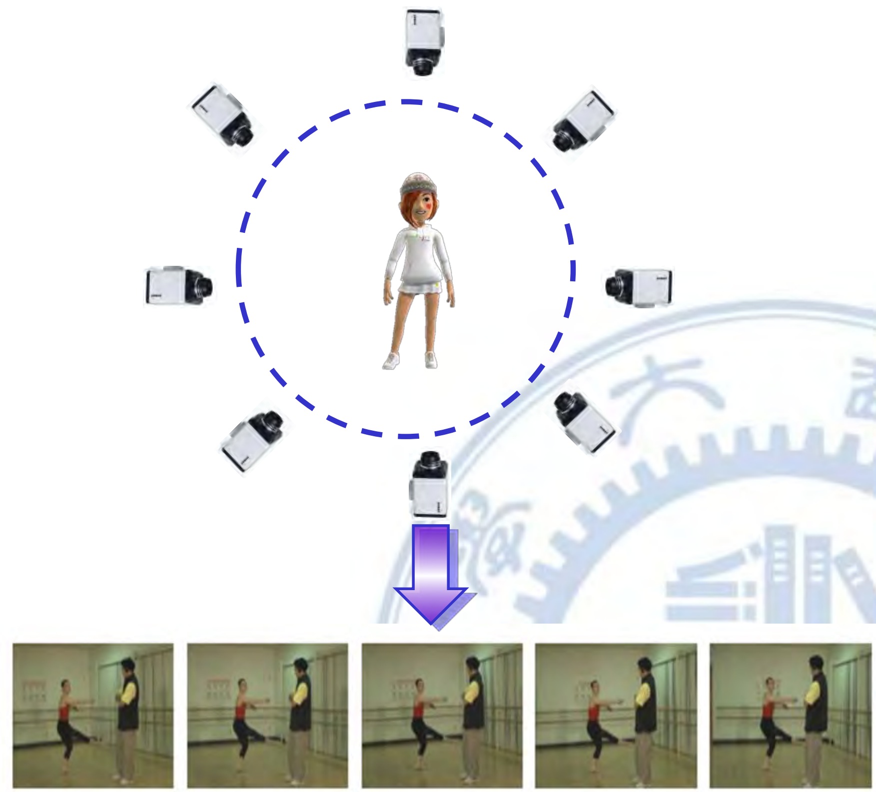

3D Capture Technique in Avatar

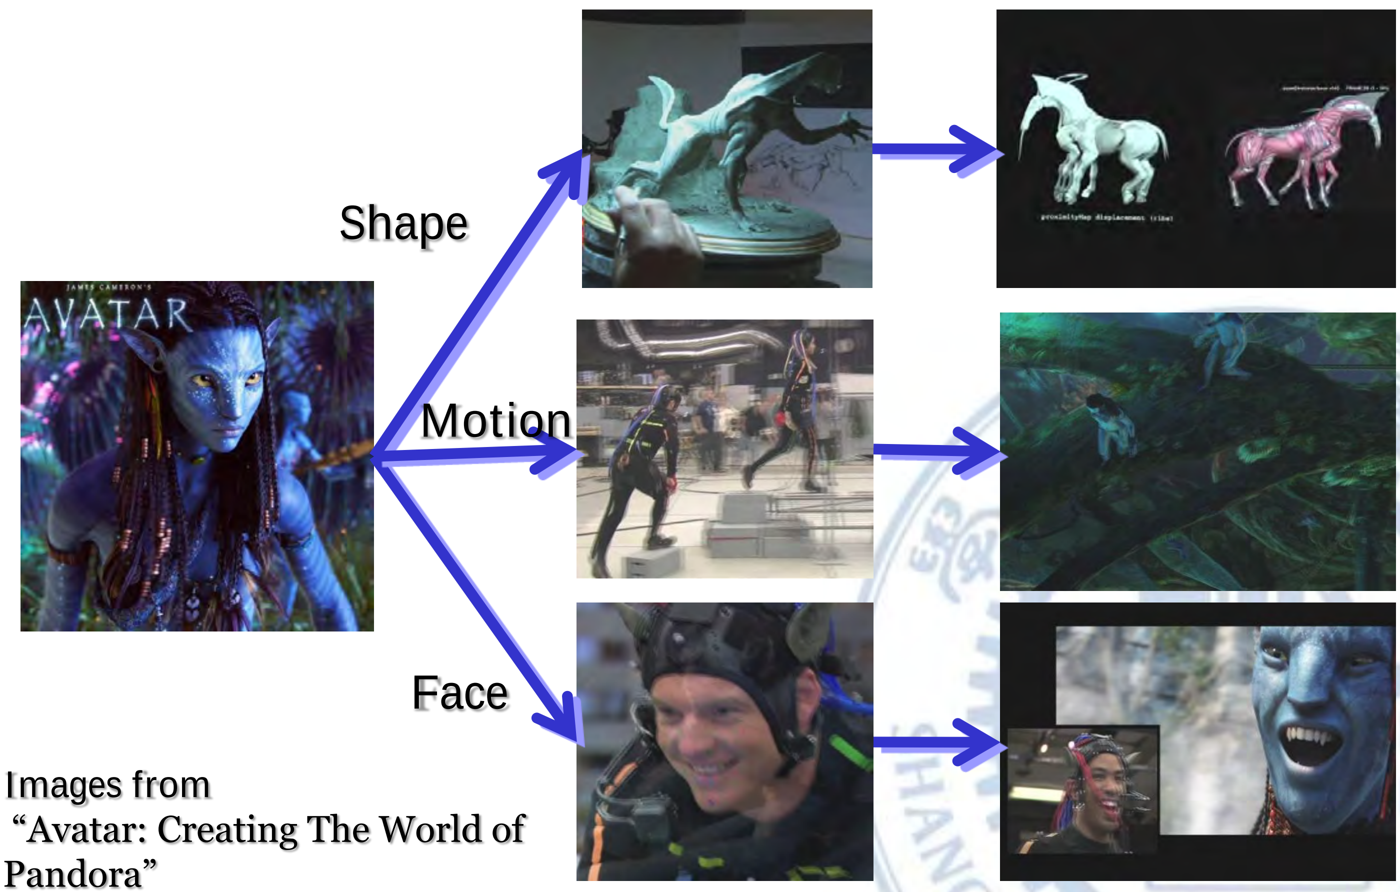

Images from
"Avatar: Creating The World of Pandora"

3D Surface Reconstruction

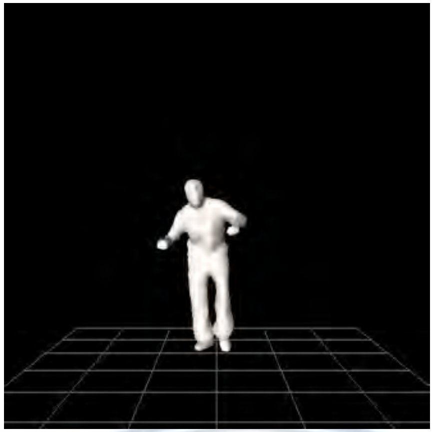

Surface reconstruction
Using Visual-Hull and geometric constraints

Automatic 3D reconstruction from internet photo collections

"Statue of Liberty" "Half Dome, Yosemite" "Colosseum, Rome"

Flickr photos

3D model

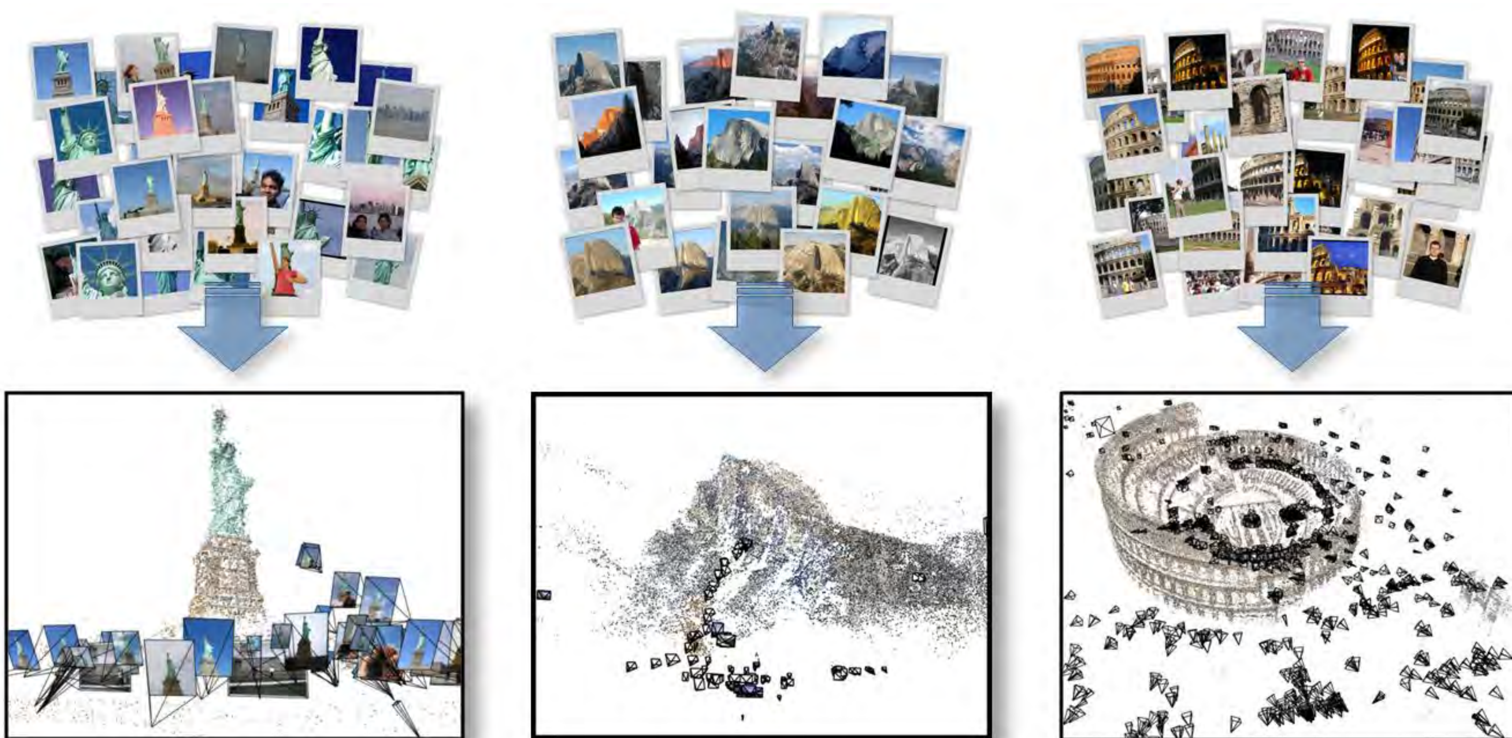

Seam carving

Seam carving (also known as image retargeting, content-aware image resizing, content-aware scaling, liquid resizing, or liquid rescaling), is an algorithm for image resizing.

It functions by establishing a number of *seams* (paths of least importance) in an image and automatically removes seams to reduce image size or inserts seams to extend it. Seam carving also allows manually defining areas in which pixels may not be modified, and features the ability to remove whole objects from photographs. The purpose of the algorithm is to display images without distortion on various media (cell phones, PDAs) using document standards, like HTML, that already support dynamic changes in page layout and text, but not images.

Seam Carving

Seam Carving

Seam Carving

Seam Carving

Seam Carving

Seam Carving

Seam Carving

Seam Carving

Simple object removal: the user marks a region for removal (green), and possibly a region to protect (red), on the original image (see inset in left image). On the right image, consecutive vertical seam were removed until no 'green' pixels were left.

Seam Carving

Find the missing shoe!

Object removal: In this example, in addition to removing the object (one shoe), the image was enlarged back to its original size. Note that this example would be difficult to accomplish using in-painting or texture synthesis.

Software Recommended

GIMP is an acronym for GNU Image Manipulation Program. It is a freely distributed program for such tasks as photo retouching, image composition and image authoring.

It has many capabilities. It can be used as a simple paint program, an expert quality photo retouching program, an online batch processing system, a mass production image renderer, an image format converter, etc.

GIMP is expandable and extensible. It is designed to be augmented with plug-ins and extensions to do just about anything. The advanced scripting interface allows everything from the simplest task to the most complex image manipulation procedures to be easily scripted.

GIMP is written and developed under X11 on UNIX platforms. But basically the same code also runs on MS Windows and Mac OS X.

GIMP

Project Main Page
- http://www.gimp.org/

- A repository of extensions for GIMP, the FREE and Open Source image manipulation program.
  - http://registry.gimp.org/
- Example
  - Liquid Rescale
  - http://liquidrescale.wikidot.com/en:examples

Liquid Rescale

- Calculate the weight/density/energy of each pixel
- Generate a list of seams

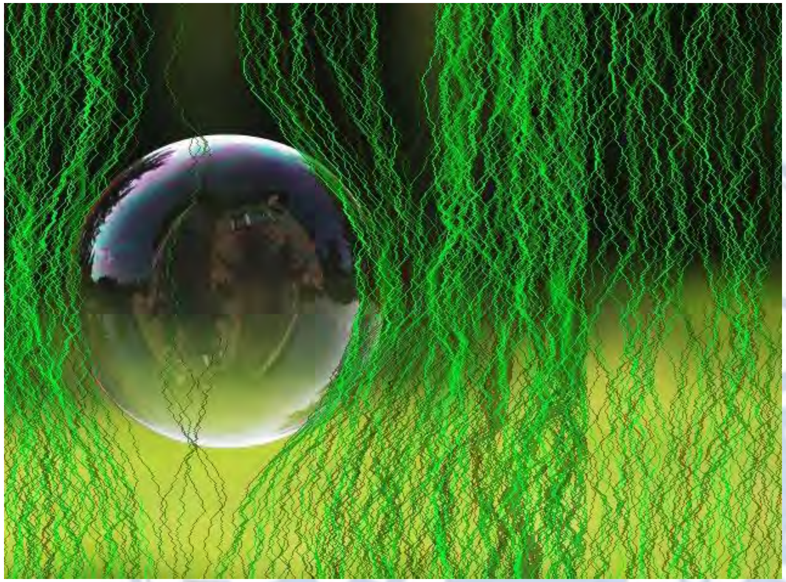

Liquid Rescale

- Calculate the weight/density/energy of each pixel
- Generate a list of seams

Why is computer vision difficult?

- What do computers see?

47	49	51	47	41	41	41	38	42	54	66	66	58	56	53	48	43	43	45	47	50	47	47	47
45	44	39	38	37	48	67	95	138	151	156	157	165	157	125	79	36	38	47	48	48	43	38	36
43	35	31	45	64	109	155	179	178	160	142	132	146	187	195	170	133	86	45	46	51	41	36	32
33	24	24	47	88	149	135	136	160	170	166	135	111	153	169	169	109	113	86	57	49	46	40	36
22	19	22	47	122	131	99	120	204	199	185	150	119	152	159	173	110	80	83	82	63	58	45	42
22	20	24	60	114	108	123	191	215	212	198	169	156	169	168	172	151	115	91	77	82	59	53	53
20	19	29	86	127	87	169	223	219	218	212	182	178	190	194	185	169	108	88	85	74	55	52	51
20	20	26	131	138	129	214	228	224	222	221	206	207	208	203	193	177	136	88	87	72	54	44	42
24	23	28	130	125	152	226	224	222	223	217	218	214	201	185	168	164	114	70	39	45	47	39	34
29	26	25	104	92	123	220	226	230	228	218	213	210	193	152	118	136	97	50	26	39	41	36	33
26	24	25	66	95	140	222	223	228	225	218	208	205	181	140	97	101	121	71	35	78	51	40	37
26	30	24	51	149	179	224	221	218	215	205	204	210	191	140	108	107	127	112	43	46	42	39	40
27	34	30	23	142	198	210	226	233	220	205	204	222	210	175	154	134	125	137	51	54	55	44	34
26	32	29	18	124	197	178	174	140	113	182	183	174	112	98	74	34	69	126	54	53	78	59	41
30	27	26	19	114	197	207	138	73	43	167	191	49	29	139	66	33	76	92	60	85	50	42	40
26	25	23	18	91	198	220	221	184	133	210	214	40	112	210	129	120	105	81	62	60	28	22	30
23	19	16	13	53	201	211	227	220	227	226	216	75	72	196	190	130	58	62	58	32	21	24	26
18	14	12	11	13	93	198	220	226	209	219	218	121	34	148	170	53	37	50	25	17	17	23	24
17	15	14	13	15	25	177	203	189	151	223	219	139	59	33	78	30	39	45	26	22	21	16	38
12	14	17	13	15	11	125	201	149	194	223	203	67	19	15	22	33	43	55	37	29	28	31	68
10	13	14	11	16	15	58	196	170	193	213	175	123	34	19	48	37	93	35	32	30	38	93	118
17	19	19	20	31	35	30	145	191	201	215	182	134	47	66	89	45	196	45	16	52	98	141	149
25	28	34	34	28	32	20	105	216	215	213	187	168	130	73	26	148	195	34	12	21	76	121	123
31	36	30	26	29	42	20	77	220	215	221	213	185	131	37	117	201	85	56	11	16	10	22	38
24	20	21	40	43	42	24	106	190	235	212	188	134	85	138	178	45	89	40	13	19	13	19	21

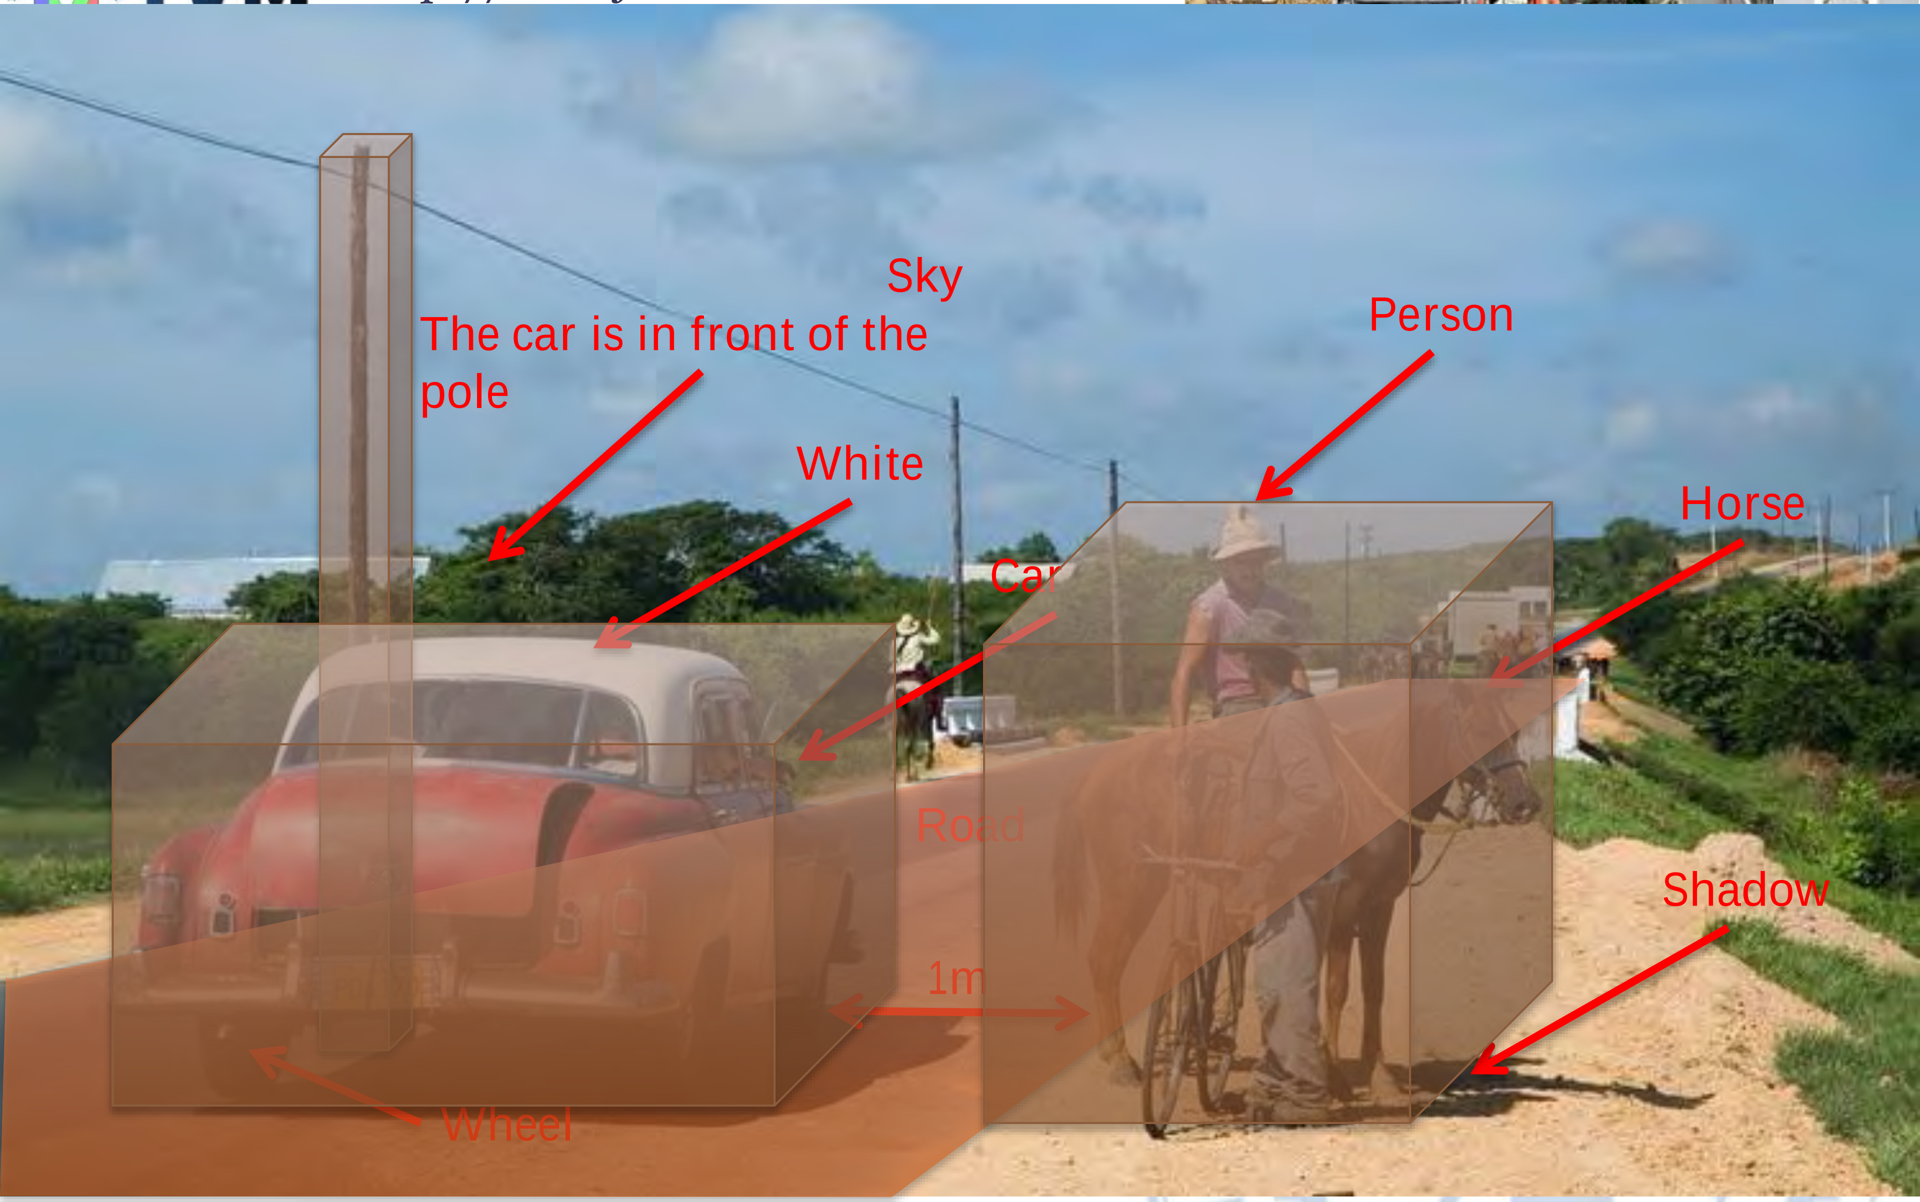
Sky
The car is in front of the pole
Person
White
Horse
Car
Road
Shadow
1m
Wheel

Visual Cues

People use information from various visual cues for recognition (e.g., color, shape, texture etc.)

- How important is each visual cue?
- How do we combine information from various visual cues?

Color Cues

Texture Cues

Shape Cues

Grouping Cues

Similarity (color, texture, proximity)

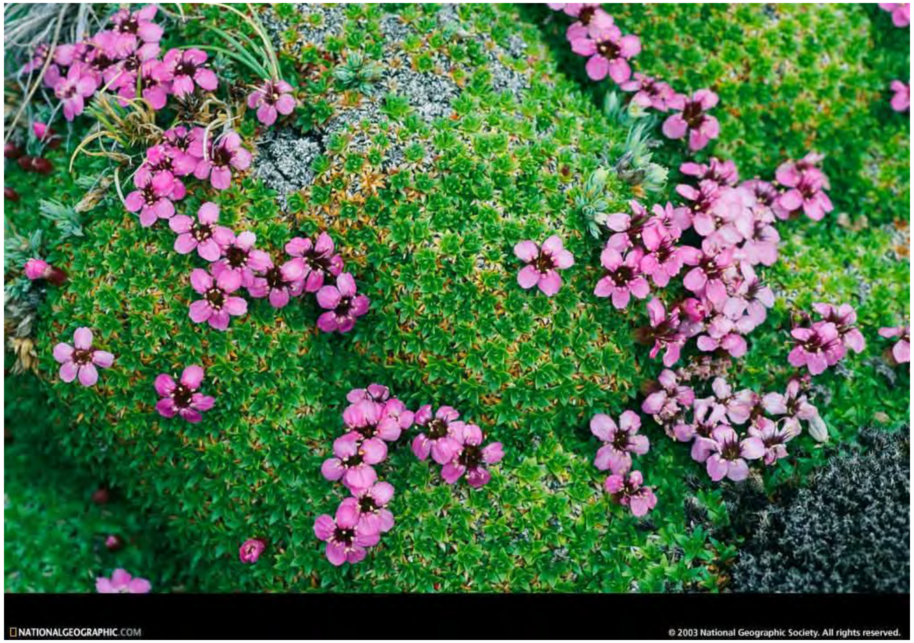

Depth Cues


NATIONALGEOGRAPHIC.COM
© 2003 National Geographic Society. All rights reserved.

Shading Cues

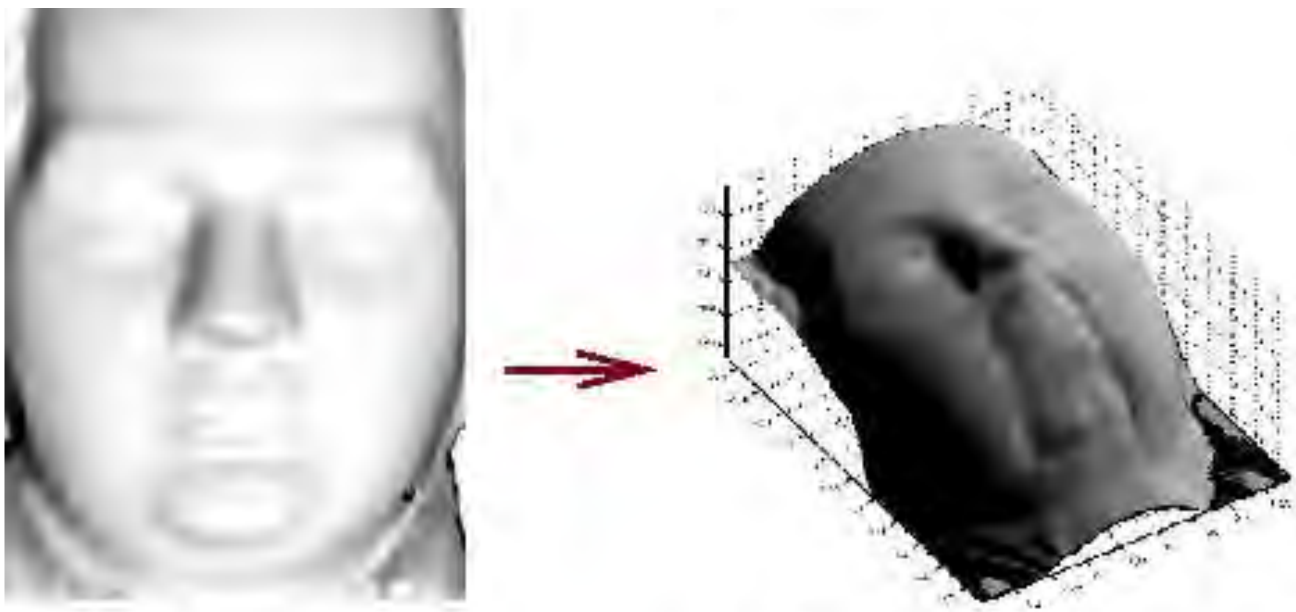

a) Image

b) 3D surface reconstructed from the single image a)

Source: J. Koenderink

Learning representations/features

The traditional model of pattern recognition (since the late 50's)

- Fixed/engineered features (or fixed kernel) + trainable classifier

hand-crafted Feature Extractor → "Simple" Trainable Classifier →

End-to-end learning / Feature learning / Deep learning

- Trainable features (or kernel) + trainable classifier

Trainable Feature Extractor → Trainable Classifier →

Deep Learning: Learning hierarchical representations

It's deep if it has more than one stage of non-linear feature transformation.

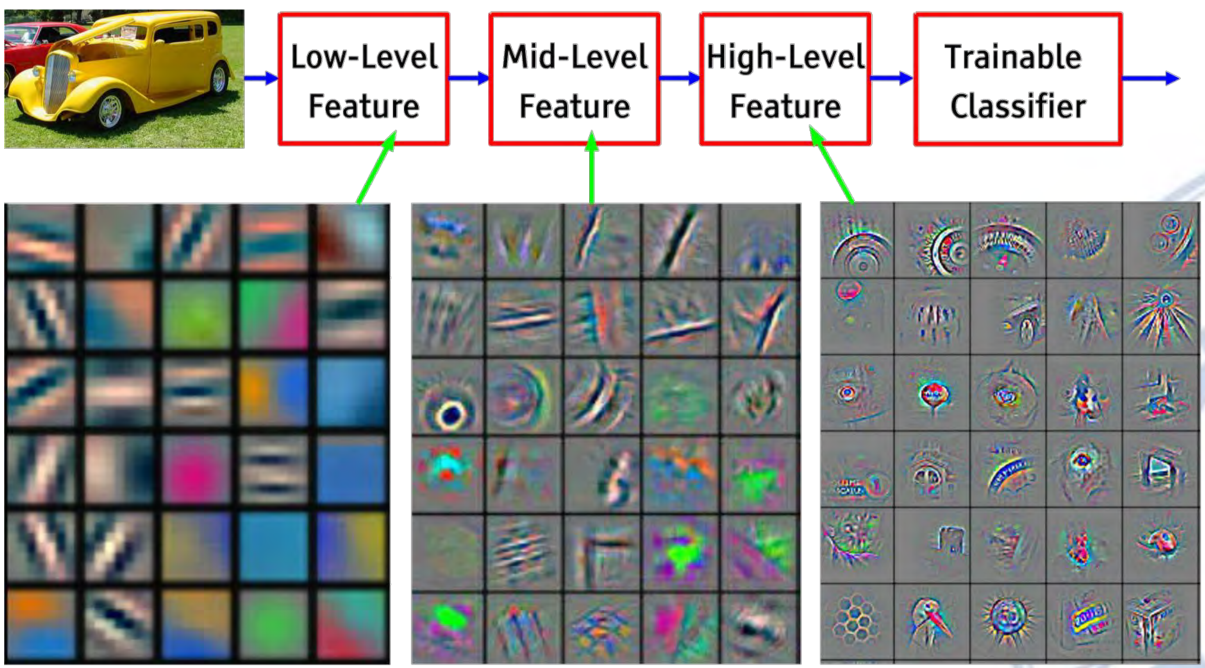

Feature visualization of convolutional net trained on ImageNet from [Zeiler & Fergus 2013]

Why Deep Learning?

- How does the cortex learn perception?

The Mammalian Visual Cortex is Hierarchical

- The ventral (recognition) pathway in the visual cortex has multiple stages
- Retina-LGN- V1 - V2 - V4 - PIT - AIT ....
- Lots of intermediate representations

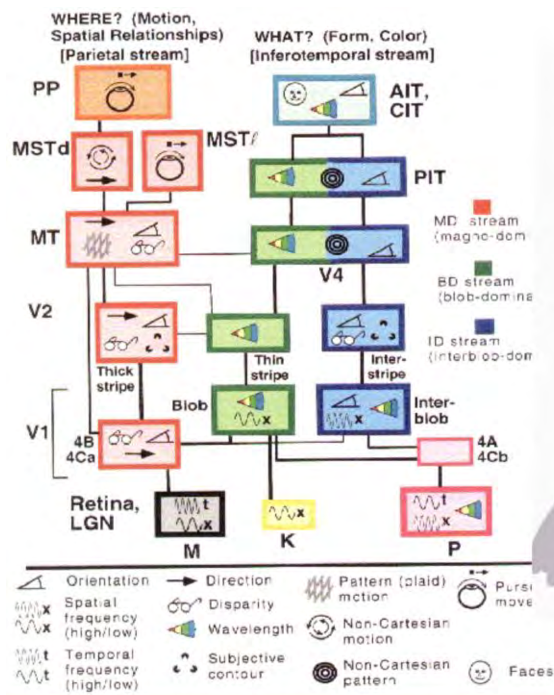

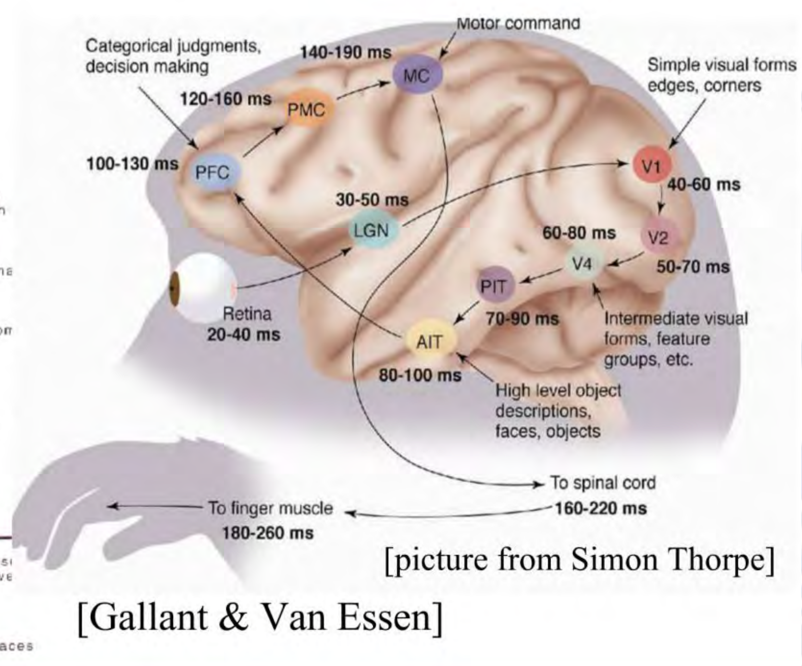

[picture from Simon Thorpe]

[Gallant & Van Essen]

Deep Learning: CNN ILSVRC Architecture

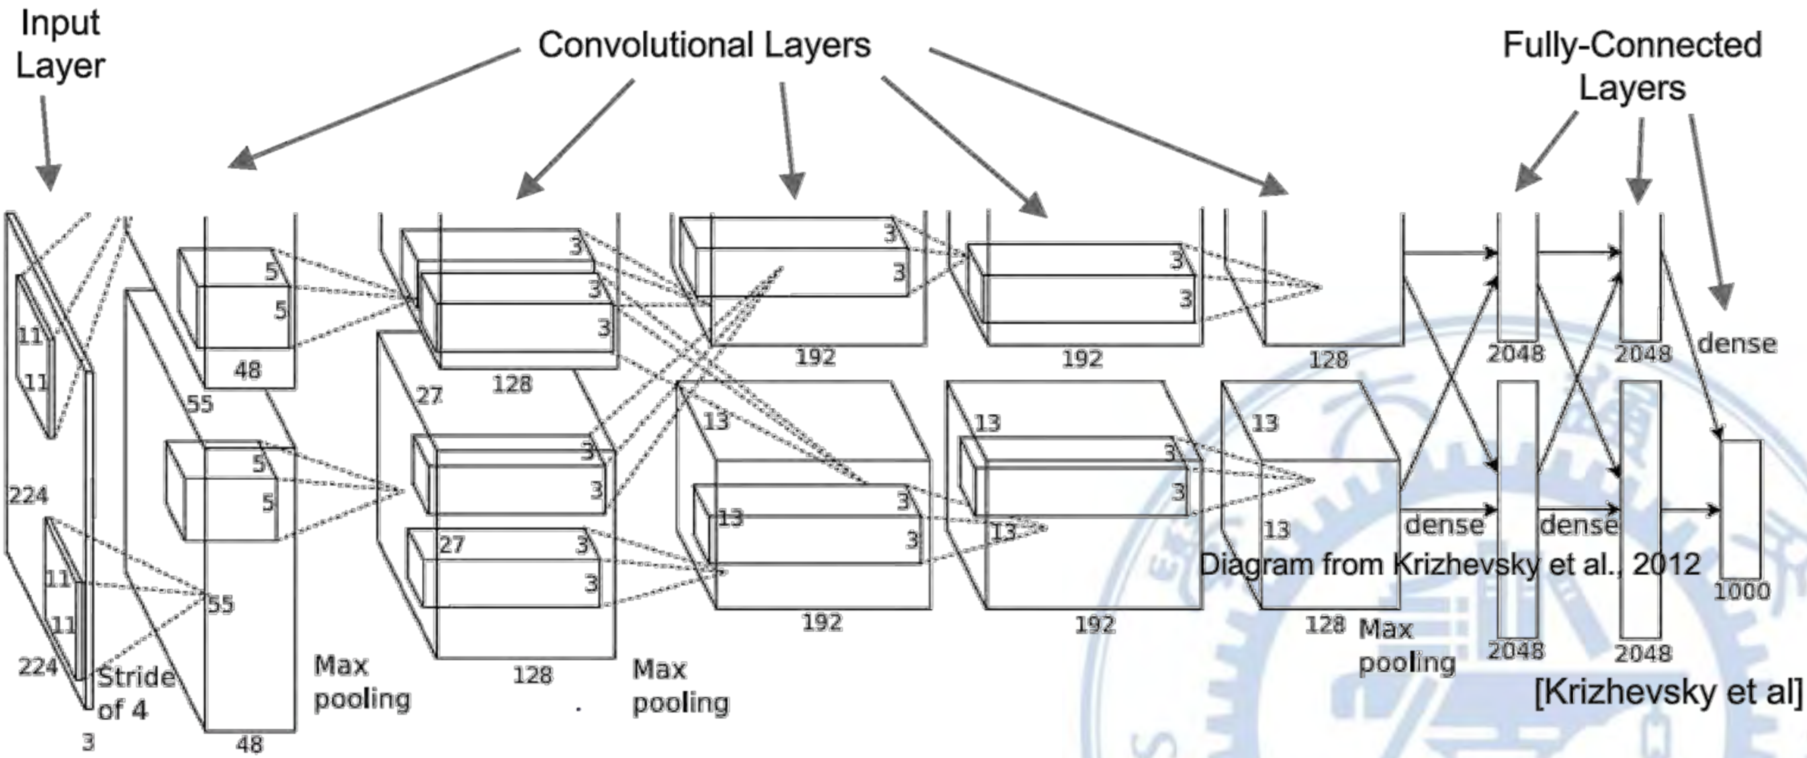

Convolve-Quantize-Pool → [Convolve-Quantize-Pool] → [[Convolve-Quantize-Pool]] → …

Deep Learning for Object Detection

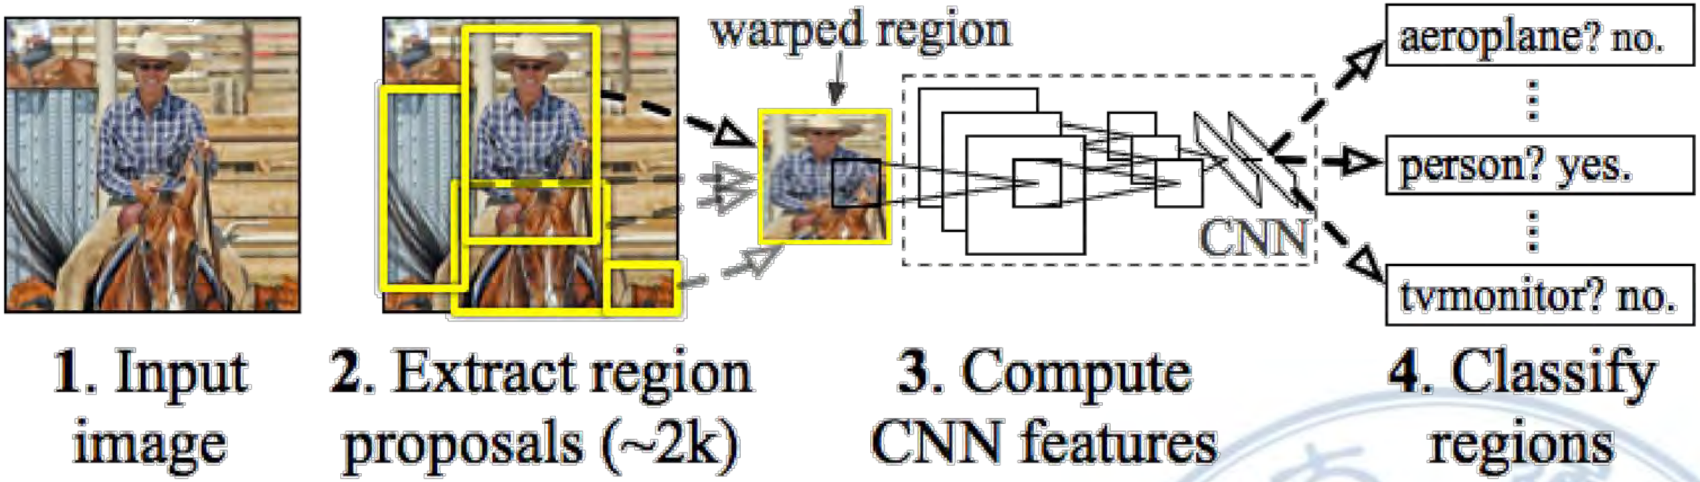

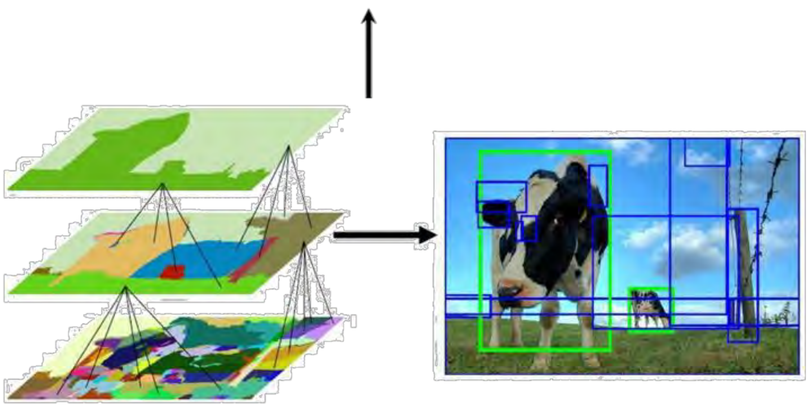

Top bicycle FPs (AP 62.5%)

Caffe: Open Sourcing Deep Learning

Convolutional Architecture for Fast Feature Extraction

- Seamless switching between CPU and GPU
- Fast computation (2.5ms / image with GPU)
- Full training and testing capability
- Reference ImageNet model available

A framework to support multiple applications:

Classification

Embedding

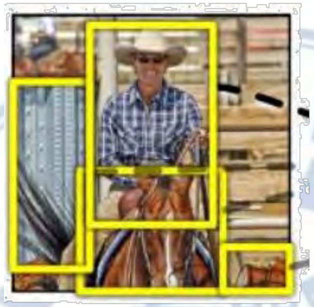

Detection

Main Page

- http://www.berkeleyvision.org/

- You will learn a basic set of image-based techniques
  - All quite simple
  - Most can be done "at home"
- You have your digital camera
- You have your imagination

Go off and explore!

Thank You!