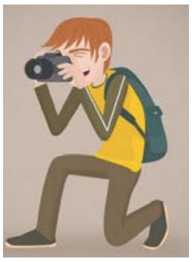

DIGITAL IMAGE PROCESSING

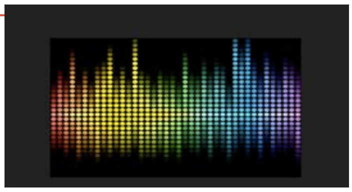

Lecture 1
Introduction

Tammy Riklin Raviv
Electrical and Computer Engineering
Ben-Gurion University of the Negev

Image understanding

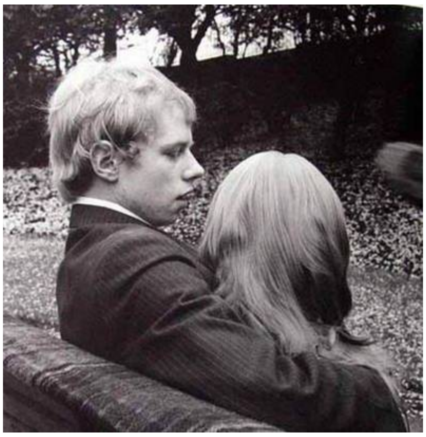

or misunderstanding ?

Introduction to Digital Image Processing

- Lecturer: Dr. Tammy Riklin Raviv
- Teaching assistants: Tal Ben-Haim and Ron Sofer
- No.: 361-1-4751
- Time: Wednesday 10:00-13:00
- Location: Virtual via zoom
- Prerequisites:
  - Signal Processing
  - Digital Signal Processing
  - Introduction to Stochastic Processes
  - Computational methods
- Course website:

http://www.ee.bgu.ac.il/~rrtammy/DIP/DIP

Course Objectives

- The primary objective of this course is to provide the students the necessary computational tools to:
  - Understand the main principles of image processing and computer vision.
  - Be familiar with different classical and commonly used algorithms and understand their mathematical foundation.
  - Implement (Matlab) and test commonly used image analysis algorithms.
  - Develop critical reading of computer vision and digital signal processing and analysis literature.
  - Plan, commit, present a system based on image processing and computer vision principles.

Course Resources

- Szeliski, Richard. Computer vision: algorithms and applications. Springer Science & Business Media, 2010.
- Gonzalez, Rafael C., and Richard E. Woods. "Image processing." Digital image processing 2 (2007).
- Forsyth, David A., and Jean Ponce. "A modern approach." Computer vision: a modern approach (2003).
- Duda, Richard O., Peter E. Hart, and David G. Stork. Pattern classification. John Wiley & Sons, 2012.
- Bishop, C. "Pattern Recognition and Machine Learning (Information Science and Statistics), 1st edn. 2006. corr. 2nd printing edn." Springer, New York(2007).

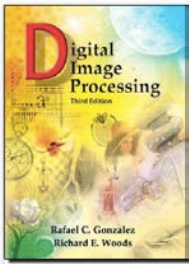

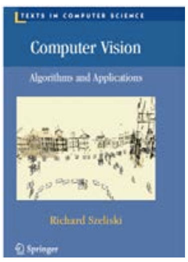

What should I do in order to succeed in the course?

- 6 Matlab assignments assignments, each 5% -> 30%
- 10 out of 13 – classwork assignments 5% (0.5% each)
- Bonus assignments
- Final project and presentation 65% Feb. 24
- Preparation meeting – mandatory

Preliminary report – 5%
Final meeting – 5%
Project presentation 15%
Final project report 40%

The instructor

- Tammy Riklin Raviv,
  Research interests:
  Signal processing: Biomedical Image Analysis, Computer Vision, Machine Learning
  Contact info:
  Telephone: 086428812
  Fax: 08-647 2949
  E-mail: rrtammy@ee.bgu.ac.il
  Office: 212/33
  Reception hours:
  please email
- Personal web page:

http://www.ee.bgu.ac.il/~rrtammy/

The Final Project

- A list of possible projects will be distributed in a few weeks
- Students can choose a subject out of this list or come up with their own ideas.
- It is the students responsibility to schedule two meetings with the teaching assistants to discuss the project of their choice.
- Students should work in teams of four (three).
- Students are expected to base their project on a scientific publication, make sure they understand it and are able to implement it using Matlab.
- All students should present their final project.

List of topics (Tentative)

Overview on digital image processing,
Visual Perception
What is an image?
Sampling, quantization,
Histogram processing
Color image processing
Edge detection
Frequency domain analysis, Fourier transform
Representation and compression: Hough transform, pyramids, quad trees, PCA, wavelates
Imaging geometry: Scaling, rotation, camera model, pose estimation

List of topics (tentative)

- Photometry, shape from shading
- Image segmentation
- Features and descriptors, SIFTs and Hogs
- Stereo and Motion
- Face detection

A graphical view on the syllabus

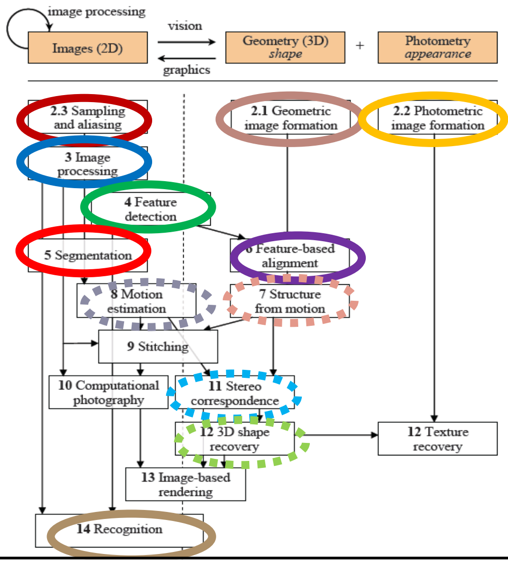

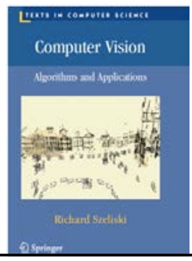

The Rest of Today's Class

- Brief Overview
- Human Vision and Visual Perception
- What is an Image?

The Rest of Today's Class

- Brief Overview
- Human Vision and Visual Perception
- What is an Image?

A brief overview: What is an Image?

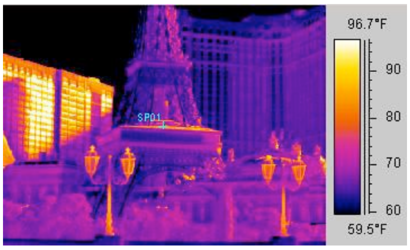

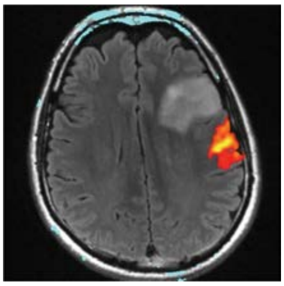

Brown's CV course

A Brief Overview: Image Formation

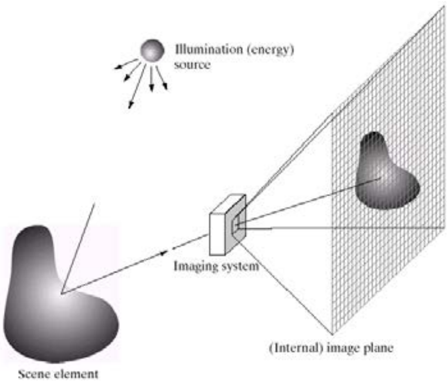

Camera Sensor

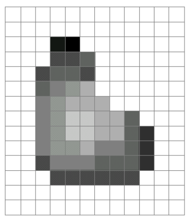

Output Image

Brown's CV course

A Brief Overview: Image Formation

See: Introduction to Medical Imaging

Magnetic Resonance Imaging

A Brief Overview: Image Formation

The main focus

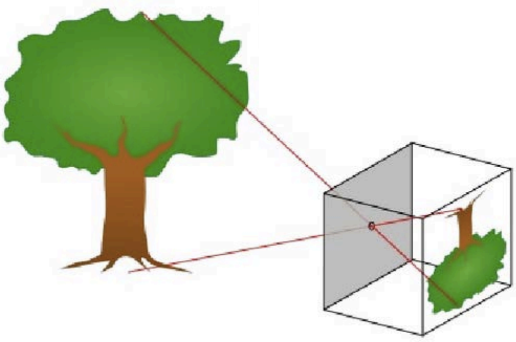

A Brief Overview: Image Formation

Sensor Array – Active Pixel Sensor

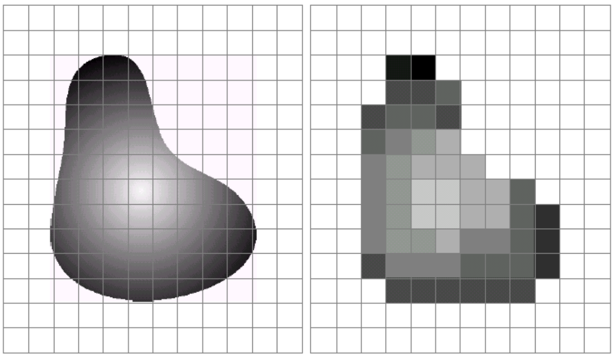

a b

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

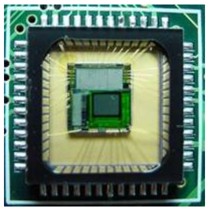

CMOS sensor

Complementary metal-oxide seminconductor

James Hays

A Brief Overview: Image Processing

Sampling & Quantization

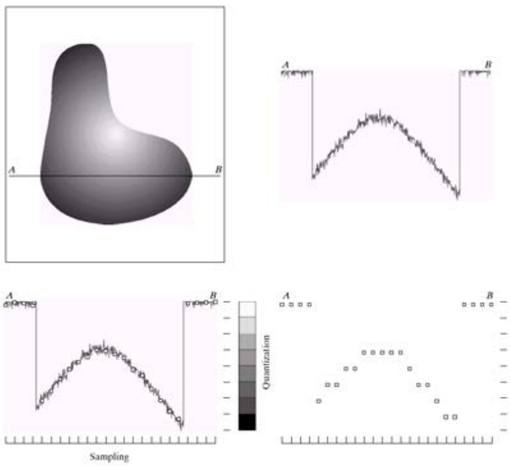

a b
c d

FIGURE 2.16 Generating a digital image. (a) Continuous image. (b) A scan line from *A* to *B* in the continuous image, used to illustrate the concepts of sampling and quantization. (c) Sampling and quantization. (d) Digital scan line.

Sampling rate determines the spatial resolution of the digitized image.
Quantization level determines the number of grey levels in the digitized image.

James Hays

A Brief Overview: Image Processing

Images as 2D Signals

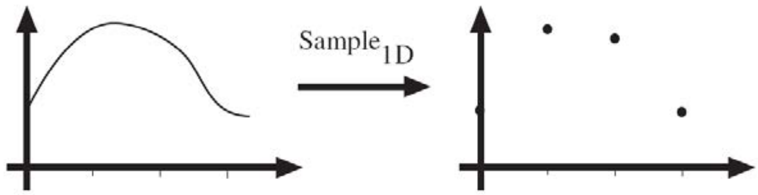

Sampling

A Brief Overview: Image Processing

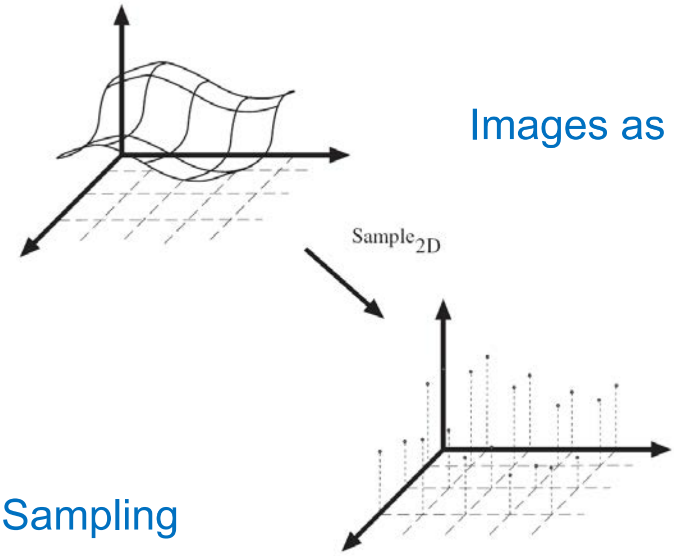

Images as 2D Signals

Sampling

A Brief Overview: Image Representation

Grayscale Digital Image

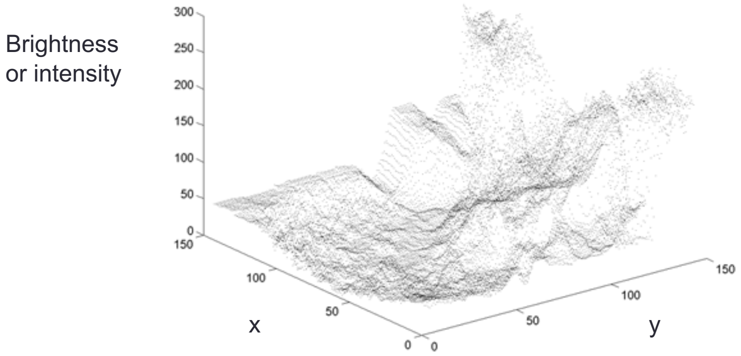

Danny Alexander

A Brief Overview: Light and Color

Danny Alexander

A Brief Overview: Illumination

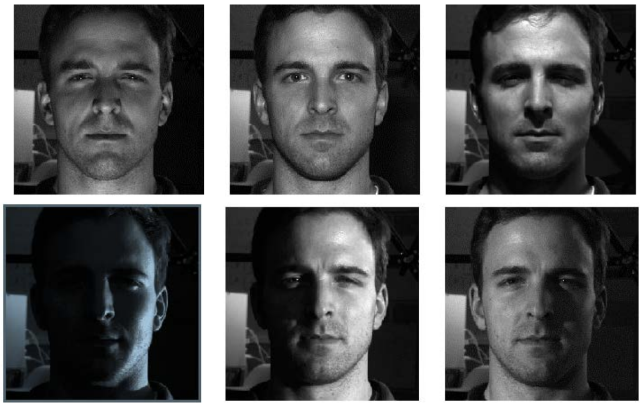

Shashua & Riklin Raviv, TPAMI 2001

Danny Alexander

A Brief Overview: Edge detection

Camera man

Danny Alexander

A Brief Overview: Frequency Analysis

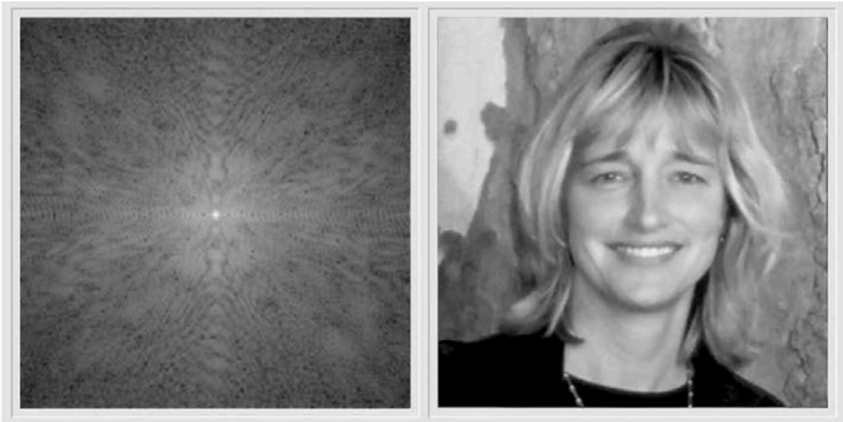

Fourier domain Image domain

A Brief Overview: Frequency Analysis

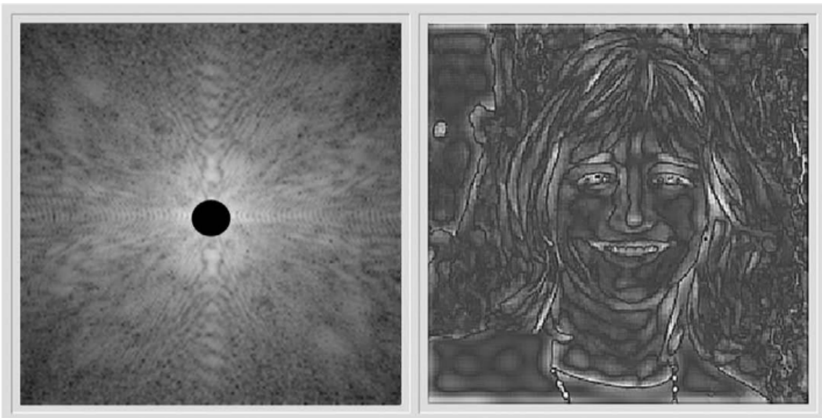

Fourier domain Image domain

A Brief Overview: Frequency Analysis

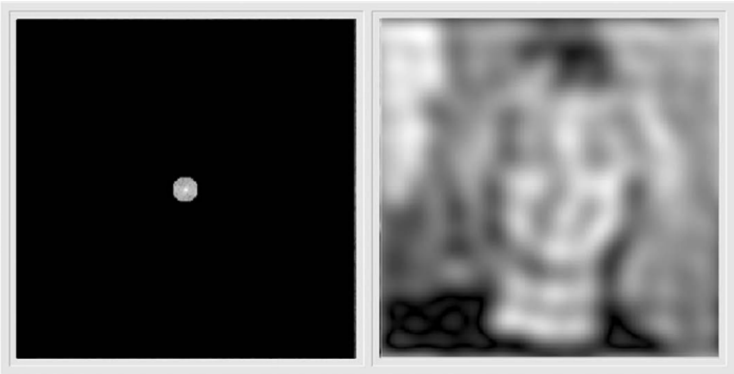

Fourier domain Image domain

A Brief Overview: Image Segmentation

Comaniciu, and Meer, TPAMI 2002

Agarwala et. al SIGGRAPH 2004

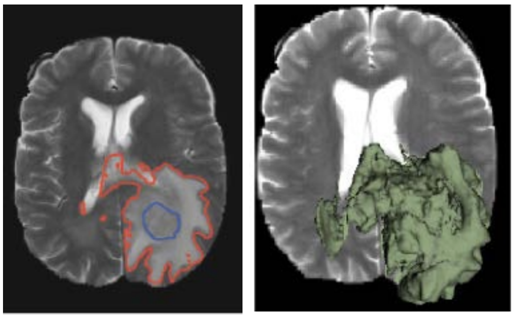

Riklin Raviv SSVM 2017

A Brief Overview: Image Segmentation

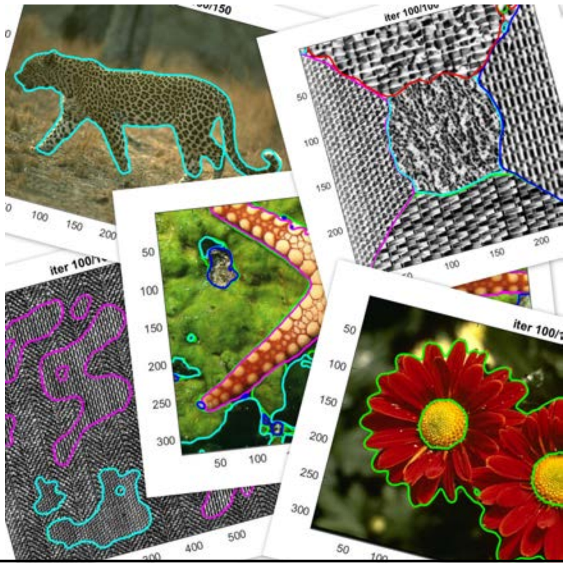

Texture segmentation

A Brief Overview: Image Segmentation

Prior based segmentation

Riklin Raviv et al, IJCV 2007

A Brief Overview: Image Segmentation

Symmetry based segmentation

Riklin Raviv et al, TPAMI 2009

A Brief Overview: Feature Detection

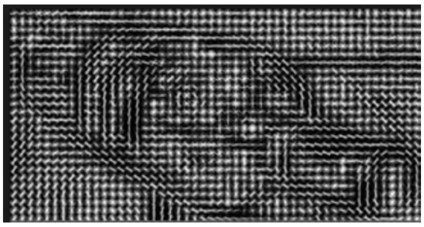

A Brief Overview: Correspondences

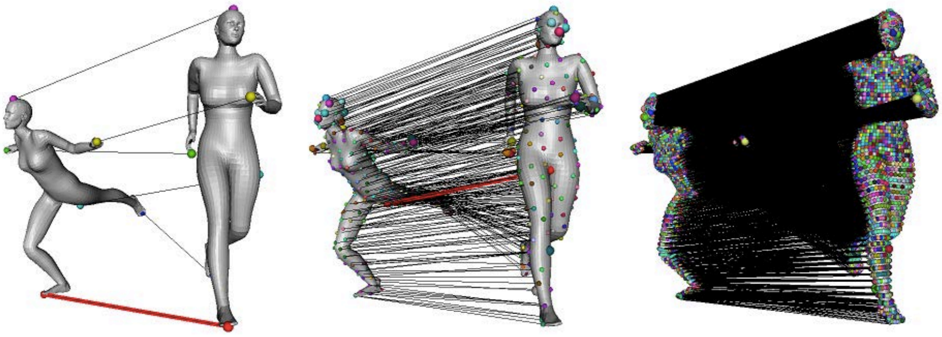

A Brief Overview: Correspondences

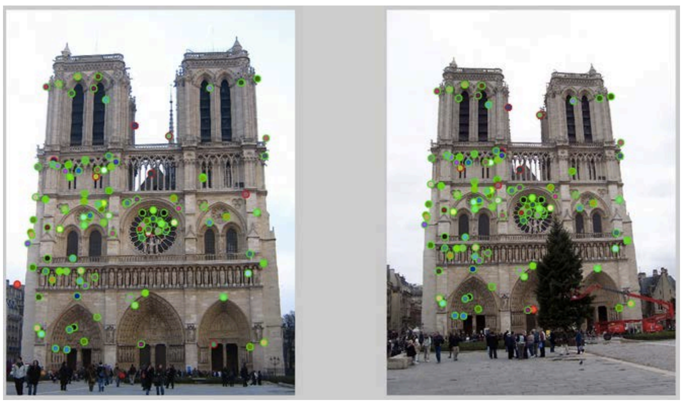

A Brief Overview: Stereo and 3D reconstruction

A Brief Overview: Object Detection and Recognition

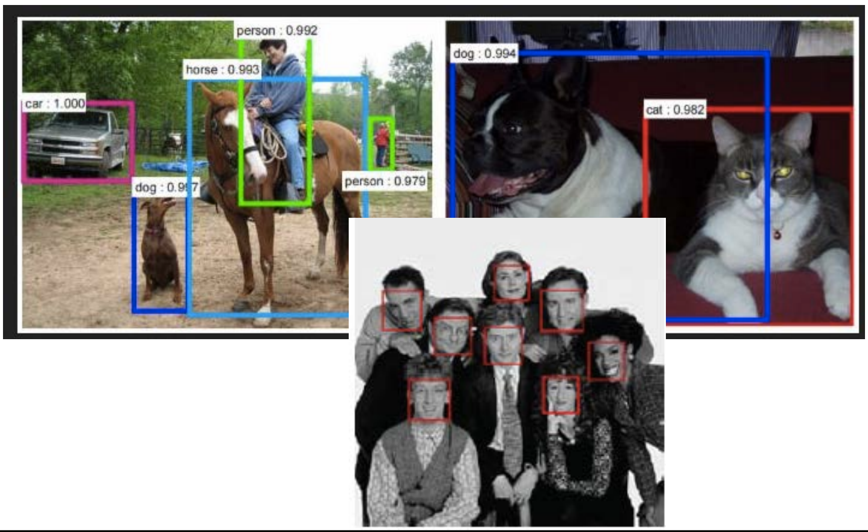

A Brief Overview: Motion

Optical Flow

A Brief Overview: Tracking

Hyun Tae Na Thesis

A Brief Overview: Tracking

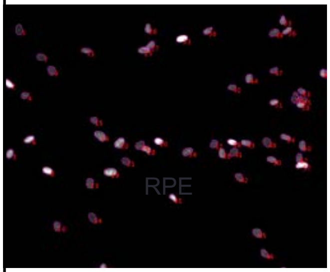

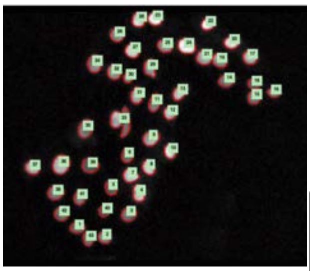

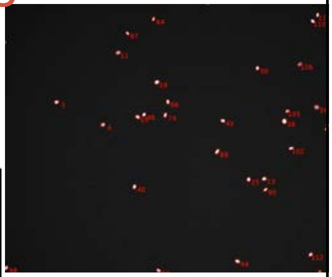

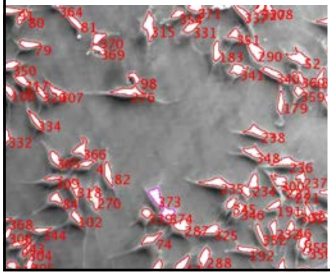

Arbelle's Thesis

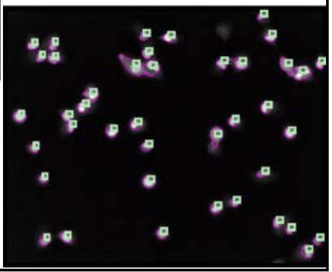

The Rest of Today's Class

- Brief Overview
- Human Vision and Visual Perception
- What is an Image?

Human Vision/ Visual Perception

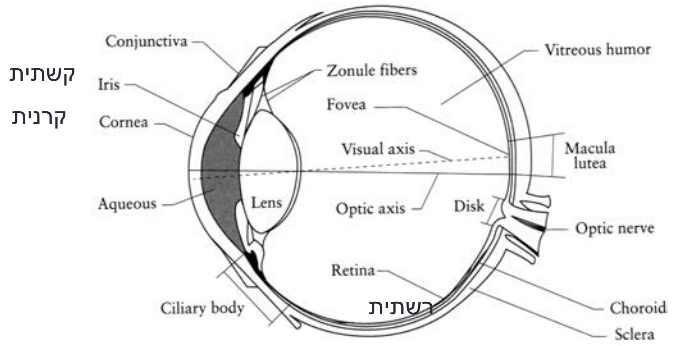

- The human eye is a camera
  - **Iris** - colored annulus with radial muscles
  - **Pupil** - the hole (aperture) whose size is controlled by the iris
  - What's the sensor?
    - photoreceptor cells (rods and cones) in the **retina**

Slide by Steve Seitz

Human Vision/ Visual Perception

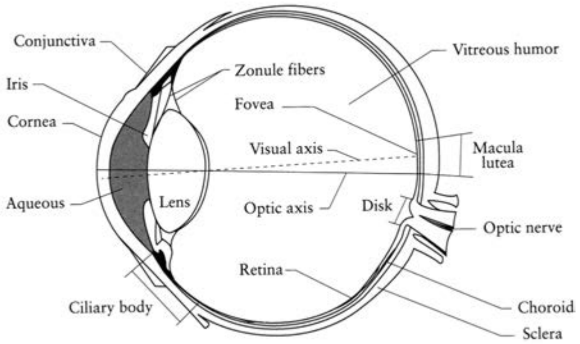

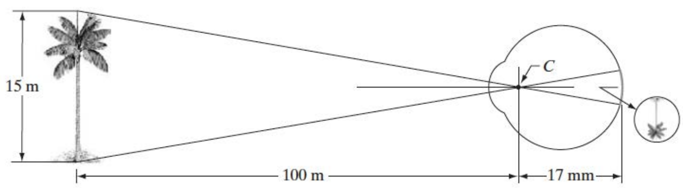

›y Steve Seitz

Human Vision: the Retina

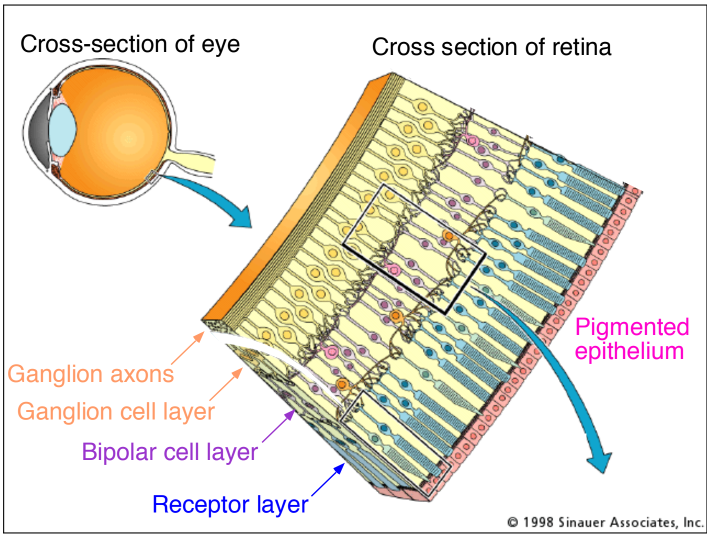

Human Vision: Retina up-close

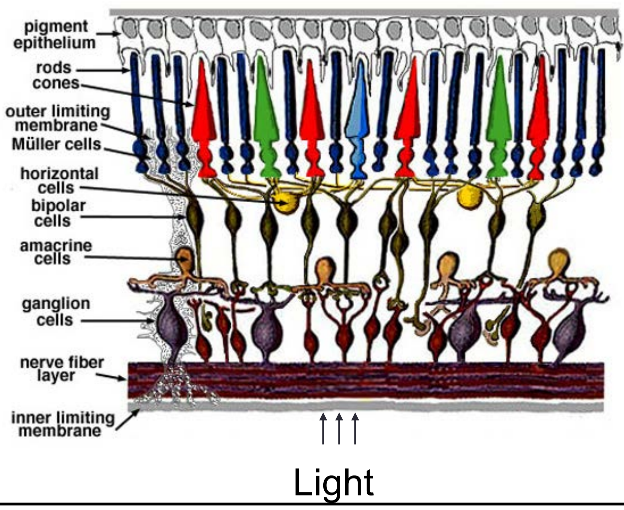

Human Vision: Retina up-close

Cones
cone-shaped
less sensitive
operate in high light
color vision

Rods
rod-shaped
highly sensitive
operate at night
gray-scale vision

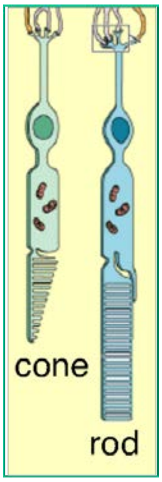

Two types of light-sensitive receptors

James Hays

Human Visual Perception

What can we learn from human visual perception?

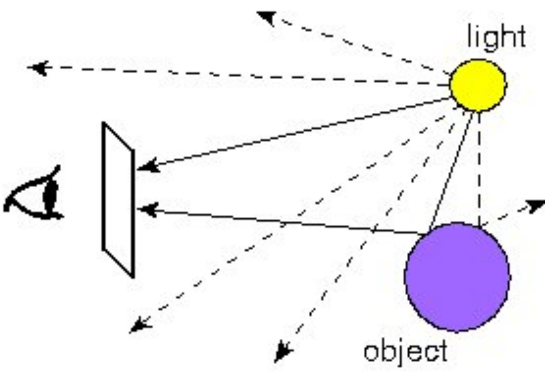

Human Perception

Muller-Lyer illusion

Ponzo Illusion

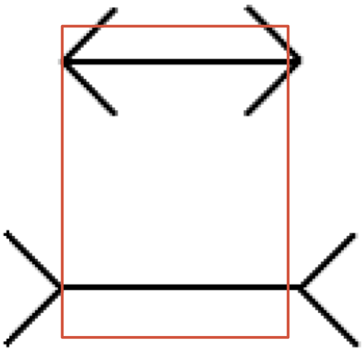

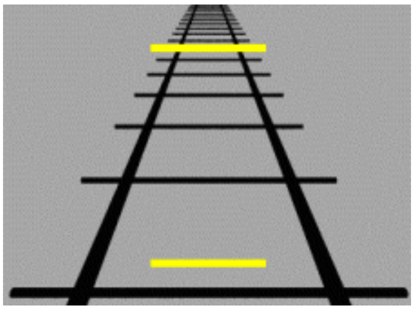

https://commons.wikimedia.org/w/index.php?curid=1211098

Human Perception

brightness constancy

https://plus.google.com/109794669788083578017/posts/YPZANXYcNFU

Human Perception

Kanizsa triangle

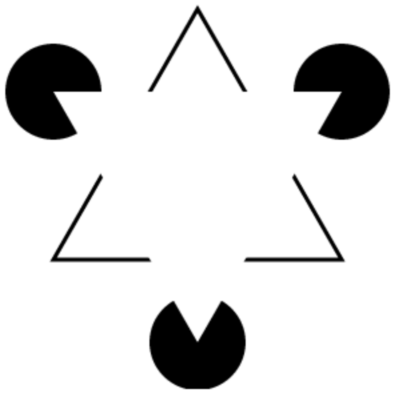

https://commons.wikimedia.org/wiki/File%3AKanizsa_triangle.svg

Human Perception

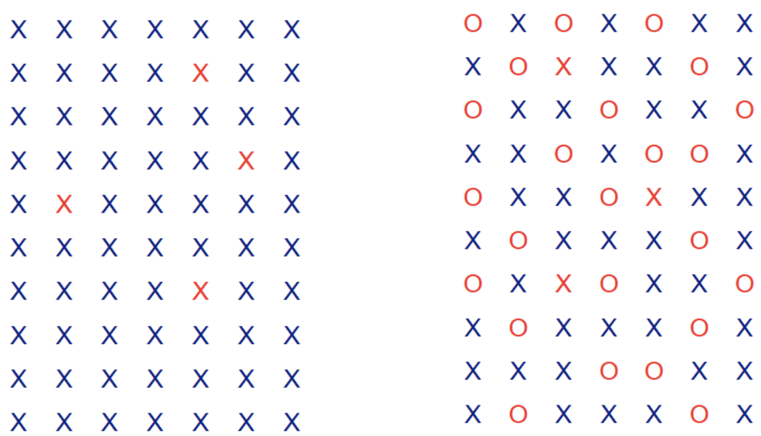

Count the red X

pop-out effect (Treisman 1985)

The Rest of Today's Class

- Brief Overview
- Human Vision and Visual Perception
- What is an Image?

What is an Image?

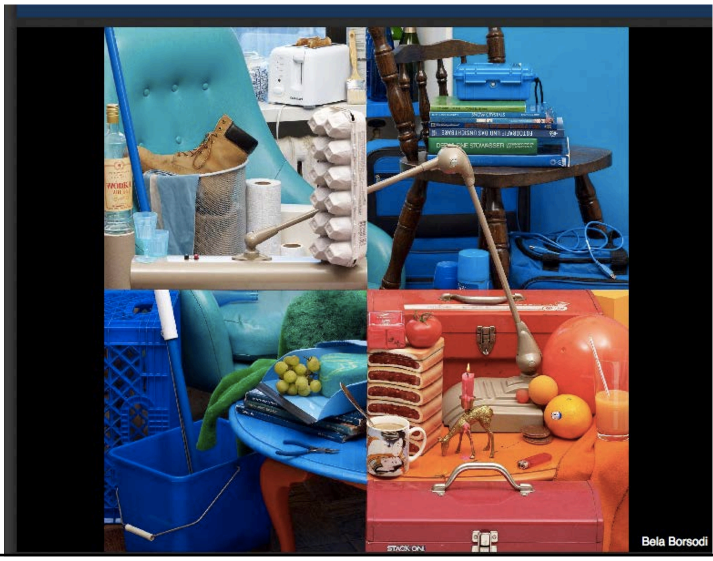

This Image is taken from Brown's Computer Vision Course

How would it look through the "computer's eyes" ?

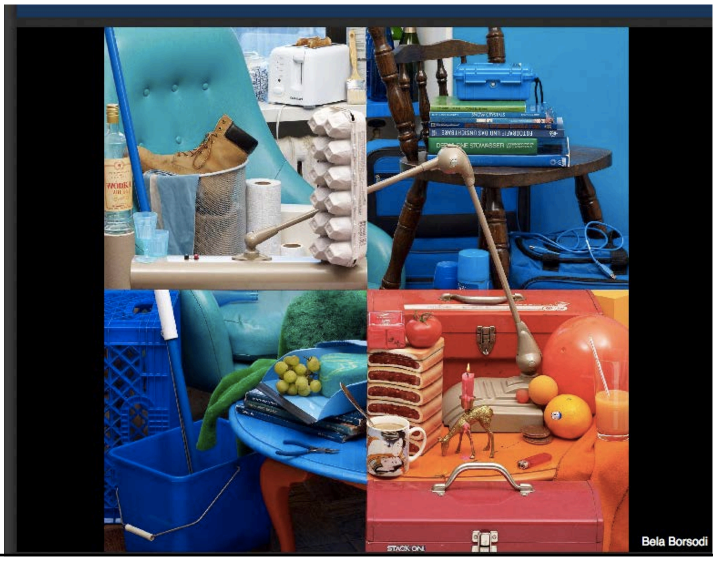

Why is this an image?

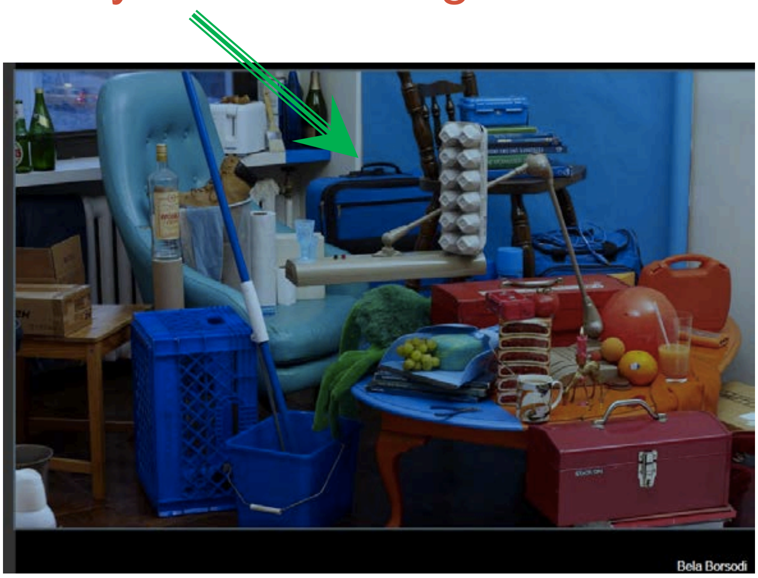

Hello Word !

```
>>
>> myFirstImage = imread('someImage.png');
>> whos
  Name                Size                 Bytes  Class

  ans                 526x764x3           1205592  uint8
  myFirstImage        526x764x3           1205592  uint8

>> imshow(myFirstImage);
```

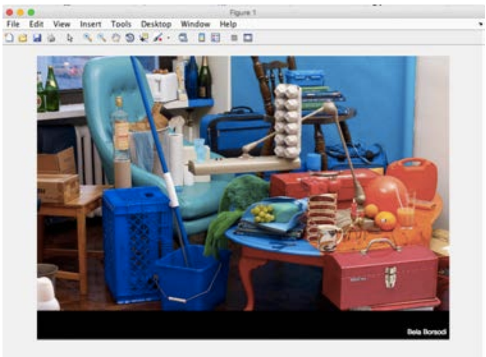

Hello Word !

```
>>
>> myFirstImage = imread('someImage.png');
>> whos
  Name                Size                 Bytes  Class

  ans                 526x764x3           1205592  uint8
  myFirstImage        526x764x3           1205592  uint8

>> imshow(myFirstImage);
>> I1 = myFirstImage(:,:,1);
>> figure;imshow(I1)
>>
```

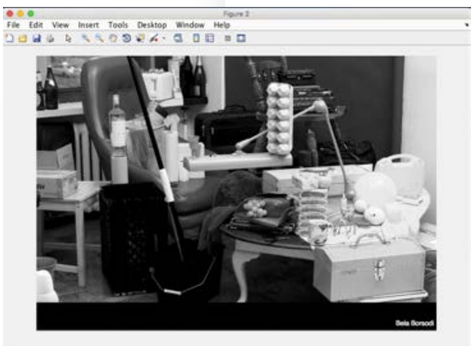

Hello Word !

```
>> I1 = myFirstImage(:,:,1);
>> figure;imshow(I1)
>> colorbar
>>
```

Hello World ☺ RGB

```
>> imshow(myFirstImage);
>> I1 = myFirstImage(:,:,1);
>> figure;imshow(I1)
>> I2 = myFirstImage(:,:,2);
>> figure;imshow(I2)
>> I3 = myFirstImage(:,:,3);
>> figure;imshow(I3)
```

Let's zoom in

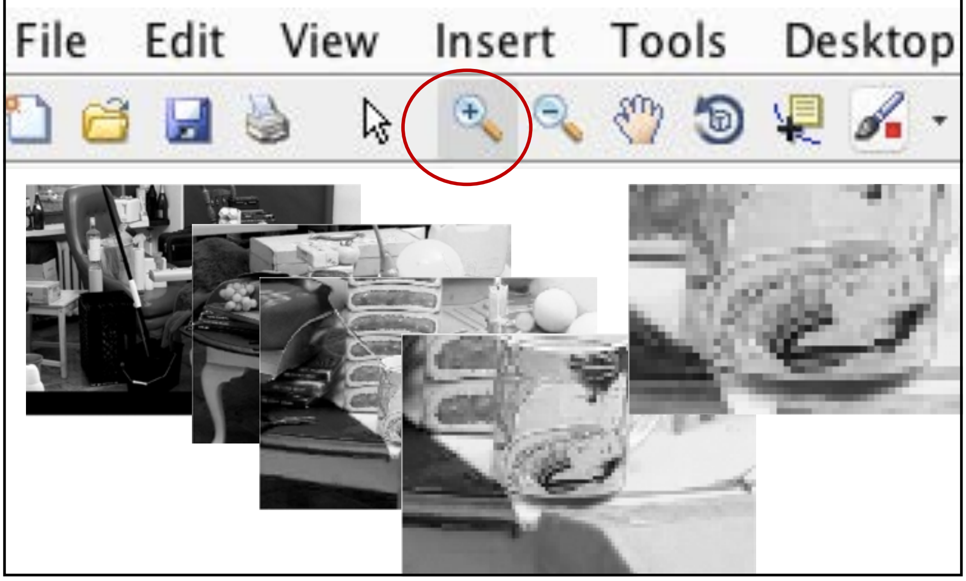

Let's Pixel

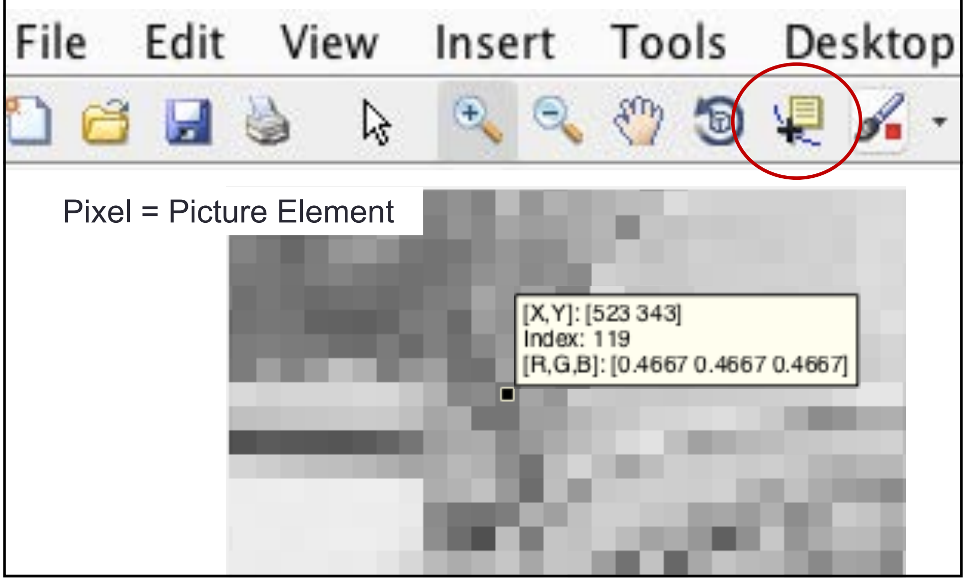

Pixel = Picture Element

Let's Pixel

220	218	214	224	224	179	223	208	230	237	210
211	217	191	215	192	187	221	213	232	223	231
212	219	188	229	181	189	221	212	208	230	235
217	214	220	198	190	194	205	187	221	236	222
202	210	175	154	187	221	200	206	244	229	226
173	196	218	137	201	226	195	241	241	238	242
201	207	210	199	180	183	207	245	247	249	245
208	206	211	199	170	184	216	252	250	244	219
204	177	211	211	189	178	206	244	238	217	224
207	137	203	212	193	151	207	227	221	230	236
203	159	207	209	208	143	209	221	231	236	235

Let's count

1 Pixel = 8 bits (UINT 8) = 1 Byte

```
>> whos myFirstImage
  Name              Size                Bytes    Class

  myFirstImage      526x764x3           1205592  uint8
```

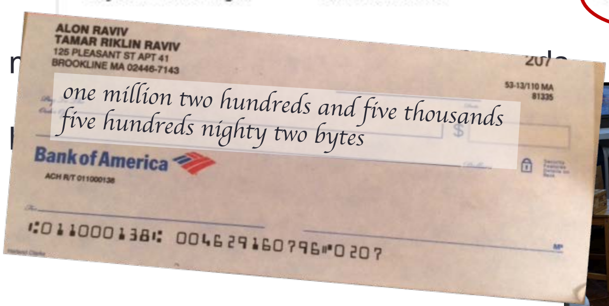

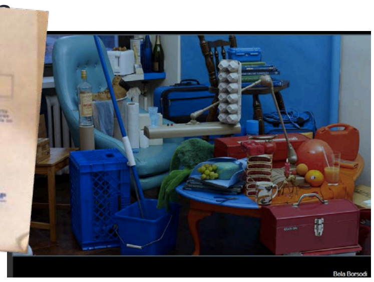

What is an Image?

An image I is a two dimensional (2D) function that maps the image domain Ω to $[0, 255]$

$I : \Omega \rightarrow [0, 255]$

or (for RGB)

$$I(\mathbf{x}) = I(x, y) = \vec{v},\ v \in [0, 255]^3$$

and the value of a single pixel is:

$I(\mathbf{x}) = I(x, y) = v,\ v \in [0, 255]$ Gray Level Pixel

$I(\mathbf{x}) = I(x, y) = (v_R, v_G, v_B)$ RGB pixel

Is this an image?

```
>>
>> I = uint8(255*rand(526,764));
>> I = uint8(255*rand(size(I1)));
>> figure ; imshow(I);
>>
```

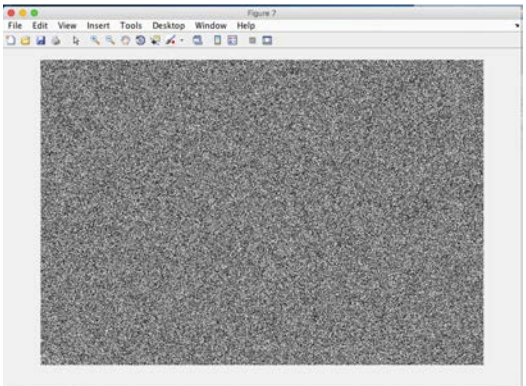

What's the difference?

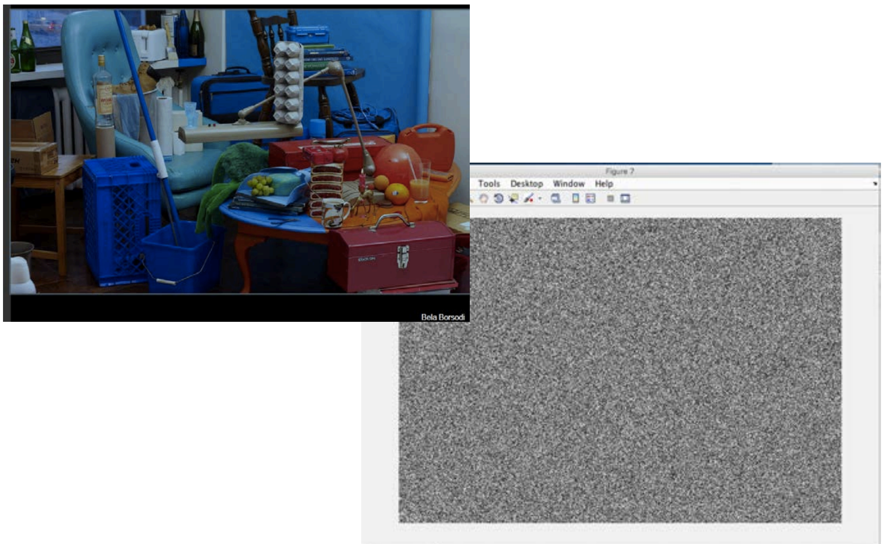

Building an histogram

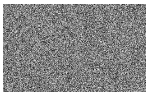

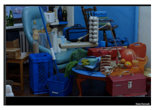

```
>>
>> I=double(I);
>> figure;hist(I(:),256);
>>
```

```
>>
>> I1 = double(I1);
>> figure;hist(I1(I1~=0),256)
>>
```

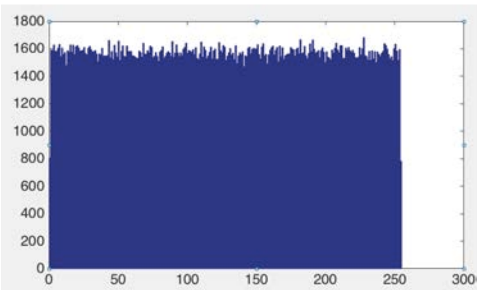

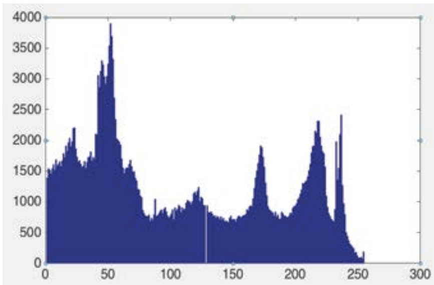

Can we measure the differences?

```
>>
>> I=double(I);
>> figure;hist(I(:),256);
>>
```

```
>> I1 = double(I1);
>> figure;hist(I1(I1~=0),256)
>>
```

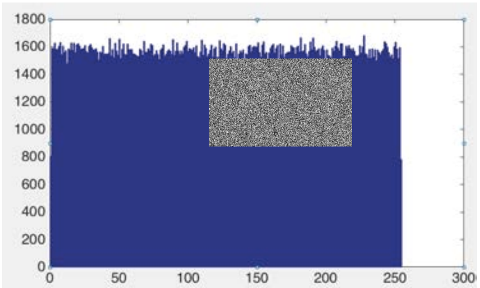

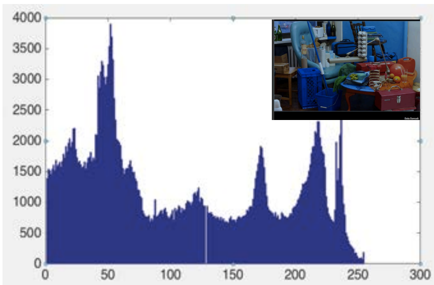

A bit on information theory

(in the context of images)

$\{0, 1\}$

Choose an arbitrary pixel in an image, can you guess its value?

Well, we can build an histogram and gamble on the value with the highest frequency

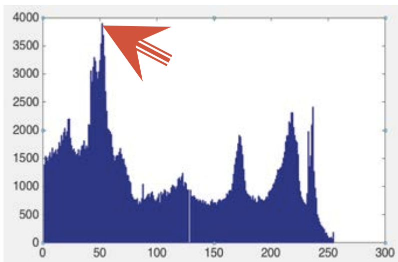

A bit on information theory

$\{0, 1\}$

Choose an arbitrary pixel in an image, can you guess its value?

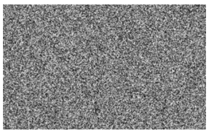

Well, we can build an histogram and gamble on the value with the highest frequency

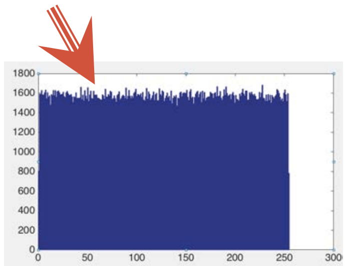

A bit on information theory

(in the context of images)

By normalizing an histogram, one can get the

probability p_i for the occurrence of the i-th value.

The Shannon entropy (measured in bits) is given by:

$$H = -\sum_i p_i log_2(p_i) \qquad i \in [0, 255]$$

where $-log_2(p_i)$ is the self-information,

which is the entropy contribution of an individual pixel.

Entropy of an Image

What does it mean ? Does it mean anything?

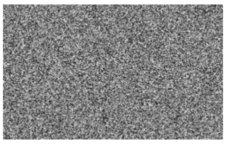

Entropy = 7.98

Entropy = 6.98

Entropy of an Image

What does it mean ? Does it mean anything?

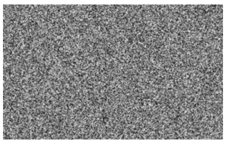

Entropy = 7.98

Entropy = 6.98

```
>>
>> Isame = uint8(100*ones(size(I1)));
>> figure; imshow(Isame)
```

But

Entropy = 0

Entropy of an Image

What does it mean ? Does it mean anything?

Entropy () = Entropy (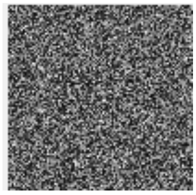)

Obtained by pixel permutation

Entropy of an Image

What does it mean ? Does it mean anything?

SmallI1

SmallI1im

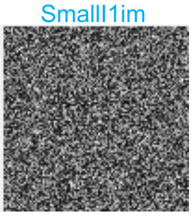

```
>>
>> myFirstImage = imread('someImage.png');
>> I1 = myFirstImage(:,:,1);
>> smallI1 = I1(1:100,1:100);
>> figure;imshow(smallI1)
>> randOrd = randperm(numel(smallI1));
>> permSI1 = smallI1(randOrd);
>> whos
  Name              Size               Bytes  Class

  I1              526x764              401864  uint8
  myFirstImage    526x764x3           1205592  uint8
  permSI1           1x10000             10000  uint8
  randOrd           1x10000             80000  double
  smallI1         100x100               10000  uint8

>> permSI1im = reshape(permSI1,size(smallI1));
>> figure;imshow(permSI1im)
```

Next-door neighbors

You have chosen a pixel and you know its value.

What can you say about the value of its next-door neighbor?

Next-door neighbors

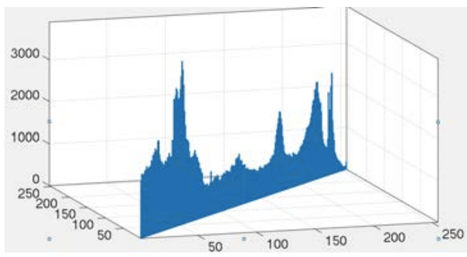


```
figure;histogram2(I1(I1~=0),I1(I1~=0),256)
```

Next-door neighbors

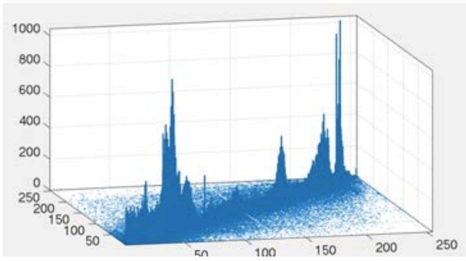


```
>> figure;histogram2(I1L(I1L~=0 & I1R~=0),I1R(I1L~=0 & I1R~=0),256)
```

Next-door neighbors

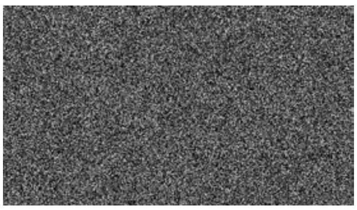

Random permutation of the red channel of myFirstImage

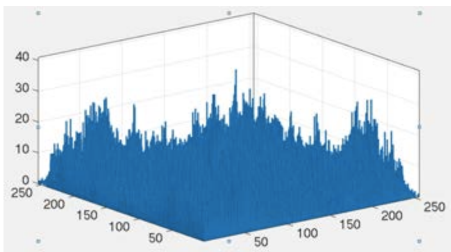


```
>> figure;histogram2(permI1L(permI1L~=0 & permI1R~=0),permI1R(permI1L~=0 & permI1R~=0),256)
```

Mutual Information

(a little bit more on information theory)

- The Mutual Information of two random variables is a measure of the variables' mutual dependence.
- The most common unit of measurement of mutual information is the bit.

Mutual Information

(a little bit more on information theory)

- The Mutual Information of two random variables is a measure of the variables' mutual dependence.

$$\mathcal{I}(J;K) = \sum_{j \in J} \sum_{k \in K} p(j,k) \log\left(\frac{p(j,k)}{p(j)p(k)}\right)$$

$p(j,k)$ is the joint probability function of J and K.
$p(j), p(k)$ are the marginal probability distribution functions of J and K (respectively).

Mini-assignment #1 (bonus)

- Read an image (any image)
- Present one of its RGB channels – I1
- Permute I1 and present it.
- Present the histogram of I1.
- Calculate its entropy
- Calculate the Mutual Information between I1 pixels and their respective left-neighbors
- Calculate the Mutual Information between the permutation image's pixels and their respective left-neighbors

Is it enough?

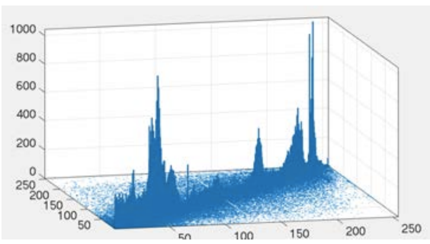

Is it enough?

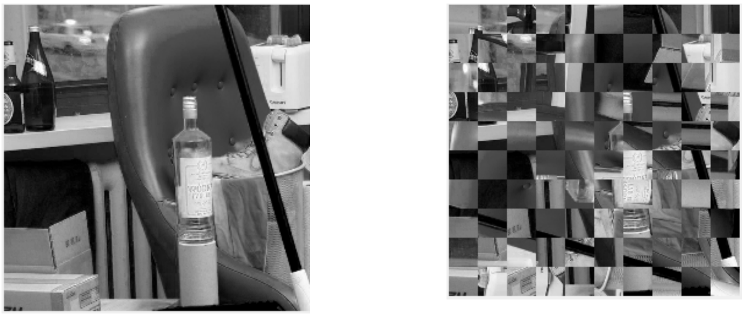

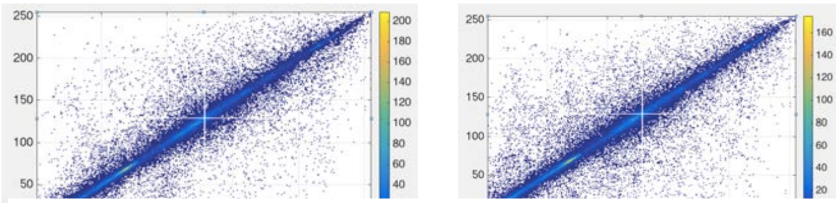


```
>> figure;histogram2(I1L(I1L~=0 & I1R~=0),I1R(I1L~=0 & I1R~=0),256,'DisplayStyle','tile')
>> colorbar
```

Is it enough?

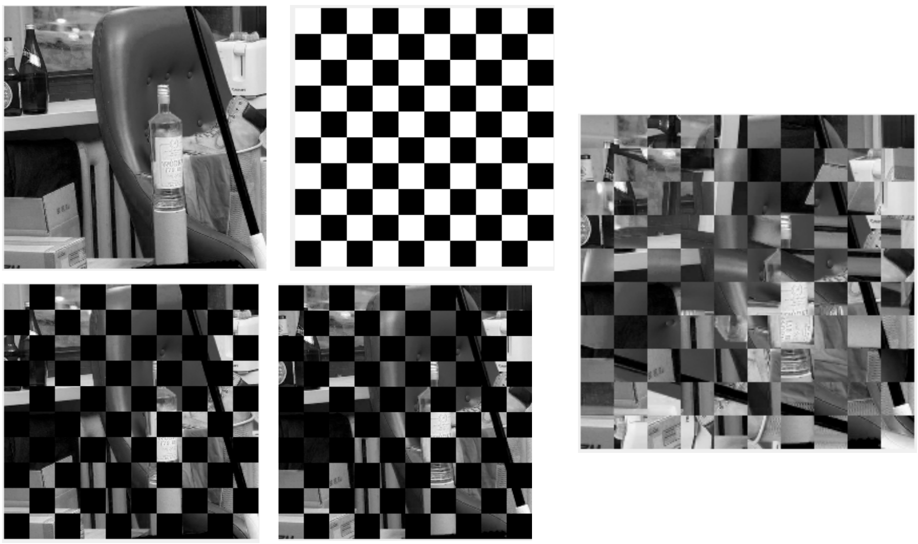

not yet there

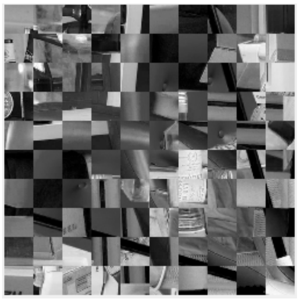

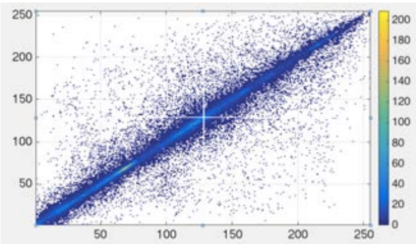

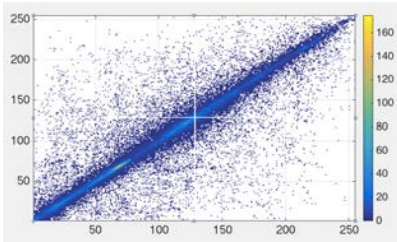

Next Class

Sampling

Quantization

Histogram Processing