BBM 413
Fundamentals of Image Processing

Erkut Erdem
Dept. of Computer Engineering
Hacettepe University

Color and Point Operations

Today

- Perception of color and light
- Color spaces
- Point operations
- Histogram processing

Today

- Perception of color and light
- Color spaces
- Point operations
- Histogram processing

Acknowledgement: This part is mostly adapted from the slides prepared by Steve Marschner, James Hays, Ali Farhadi and Anat Levin

What is color?

- Color is the result of interaction between physical light in the environment and our visual system
- Color is a psychological property of our visual experiences when we look at objects and lights, *not* a physical property of those objects or lights (S. Palmer, *Vision Science: Photons to Phenomenology*)

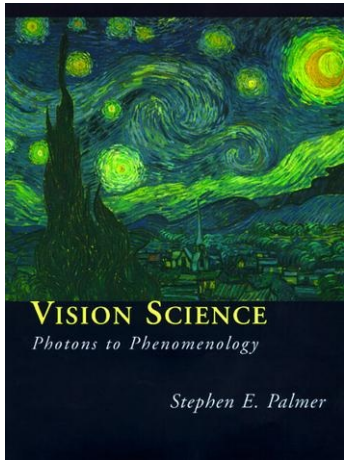

Slide credit: A. Efros

#thedress

- What is the color of the dress?

- blue and black
- white and gold
- blue and brown

- What #thedress tell about our color perception?

http://nyti.ms/186m3wE

#thedress

- Let's take averages

two pieces of the dress | averages | basic pattern

http://nyti.ms/186m3wE

#thedress

- The dress in the photograph

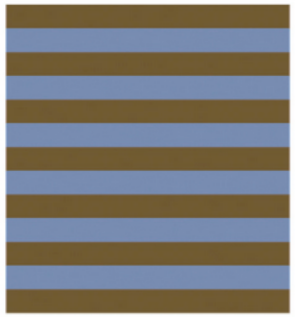

http://nyti.ms/186m3wE

#thedress

- Consider the dress is in shadow.

- Your brain remove the blue cast, and perceive it as white and gold.

http://nyti.ms/186m3wE

#thedress

- The dress in the photograph

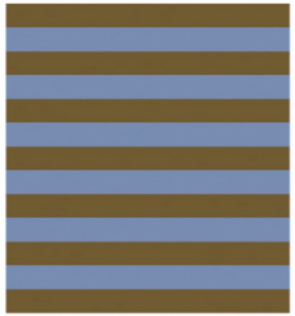

http://nyti.ms/186m3wE

#thedress

- Consider the dress is in bright light.

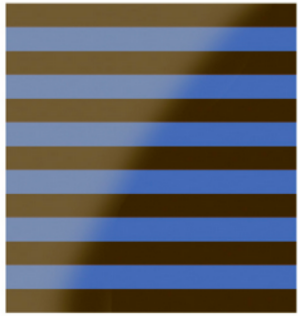

- Your brain perceive the dress as a darker blue and black

http://nyti.ms/186m3wE

#thedress

- Answer:

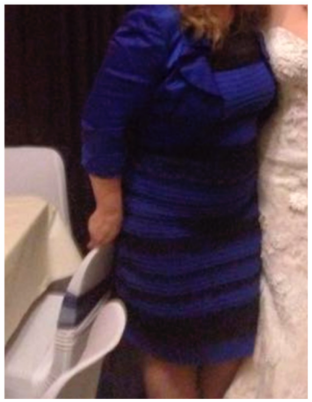

- The dress is actually blue and black.

http://nyti.ms/186m3wE

Brightness perception

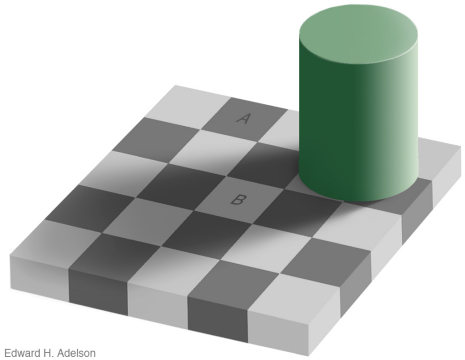

Edward Adelson

http://web.mit.edu/persci/people/adelson/illusions_demos.html

Brightness perception

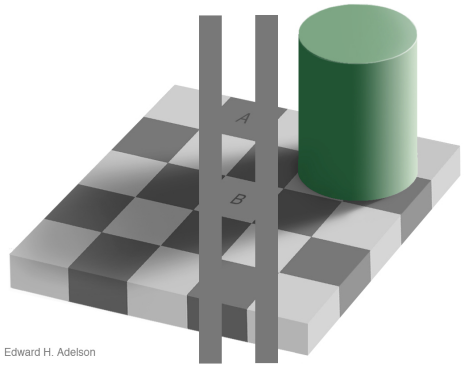

Edward Adelson

http://web.mit.edu/persci/people/adelson/illusions_demos.html

Brightness perception

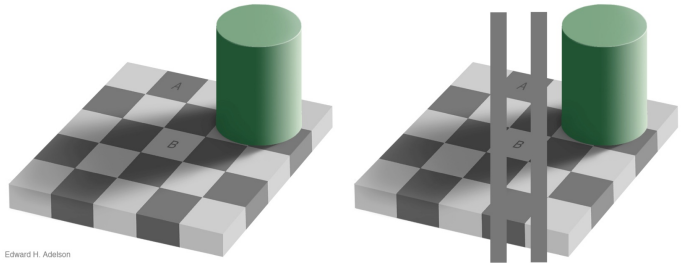

Edward Adelson

http://web.mit.edu/persci/people/adelson/illusions_demos.html

Brightness perception

Land's Experiment (1959)

- Cover all patches except a blue rectangle
- Make it look gray by changing illumination
- Uncover the other patches

Color Constancy

We filter out illumination variations

Slide credit: S. Narasimhan

Land's Experiment (1959)

- Cover all patches except a blue rectangle
- Make it look gray by changing illumination
- Uncover the other patches

Color Constancy

We filter out illumination variations

Slide credit: S. Narasimhan

Color Constancy in Gold Fish

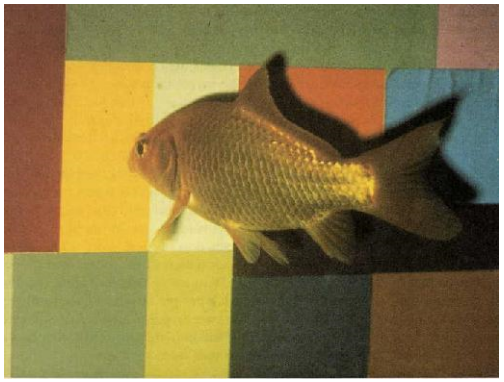

In David Ingle's experiment, a goldfish has been trained to swim to a patch of a given color for a reward—a piece of liver. It swims to the green patch regardless of the exact setting of the three projectors' intensities. The behavior is strikingly similar to the perceptual result in humans.

http://neuro.med.harvard.edu/site/dh/b45.htm

Slide credit: S. Narasimhan

Reading Assignment #2

- Watch Beau Lotto's TED talk on "Optical illusions show how we see" [link available on course webpage]

TED Ideas worth spreading

Image Brightness (Intensity)

irradiance: the power of electromagnetic radiation per unit area (radiative flux) incident on a surface.

reflectance: the fraction of incident electromagnetic power that is reflected at an interface

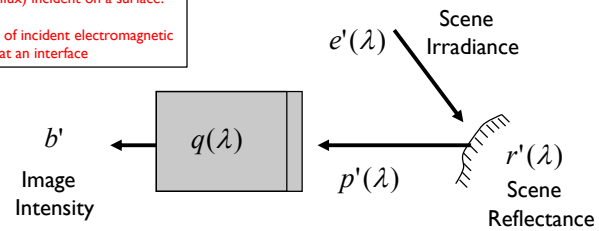

- Monochromatic Light : $(\lambda = \lambda_i)$

$$b'(x,y) = r'(x,y)\, e'(x,y) \qquad q(\lambda_i) = 1$$

NOTE: The analysis can be applied to COLORED LIGHT using FILTERS

Slide credit: S. Narasimhan

Recovering Lightness

- Image Intensity: $b'(x,y) = r'(x,y)\, e'(x,y)$

  Can we recover e' and r' from b' ?

  An illposed problem!

- Retinex theory, Land and McCann, 1971

- use constraints (or priors) on shading and reflectance
- employ additional information (multiple images, depth maps, etc.)

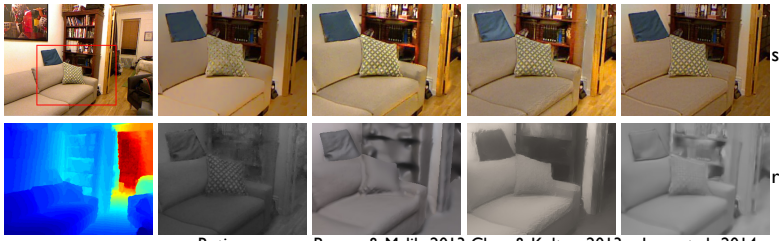

Retinex Barron & Malik, 2013 Chen & Koltun, 2013 Jeon et al., 2014

Slide credit: S. Narasimhan

Color and light

- **Color of light** arriving at camera depends on
  - Spectral reflectance of the surface light is leaving
  - Spectral radiance of light falling on that patch

- **Color perceived** depends on
  - Physics of light
  - Visual system receptors
  - Brain processing, environment

- Color is a phenomenon of human perception; it is **not** a universal property of light

Slide credit: K. Grauman, S. Marschner

Color

White light: composed of about equal energy in all wavelengths of the visible spectrum

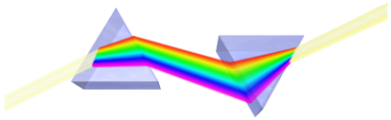

Newton 1665

Color

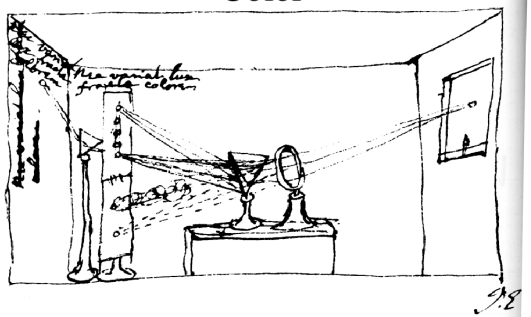

4.1 NEWTON'S SUMMARY DRAWING of his experiments with light. Using a point source of light and a prism, Newton separated sunlight into its fundamental components. By reconverging the rays, he also showed that the decomposition is reversible.

From Foundations of Vision, by Brian Wandell, Sinauer Assoc., 1995

Slide credit: B. Freeman, A. Torralba, K. Grauman

Electromagnetic spectrum

- Light is electromagnetic radiation
  - exists as oscillations of different frequency (or, wavelength)

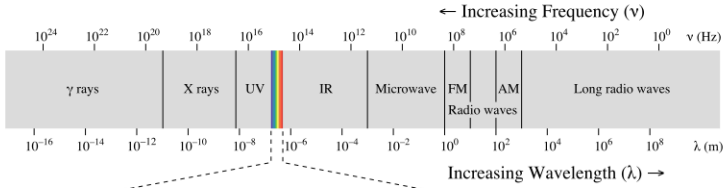

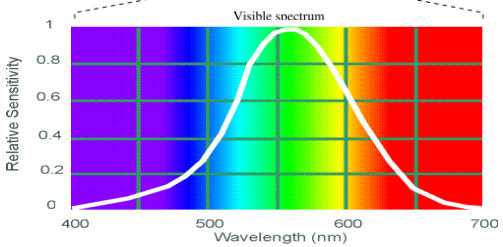

Human Luminance Sensitivity Function

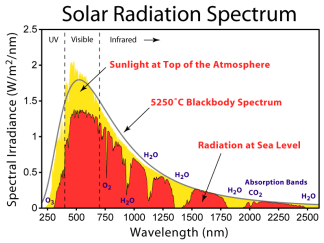

Slide credit: A. Efros

The Physics of light

Any source of light can be completely described physically by its spectrum: the amount of energy emitted (per time unit) at each wavelength 400 - 700 nm.

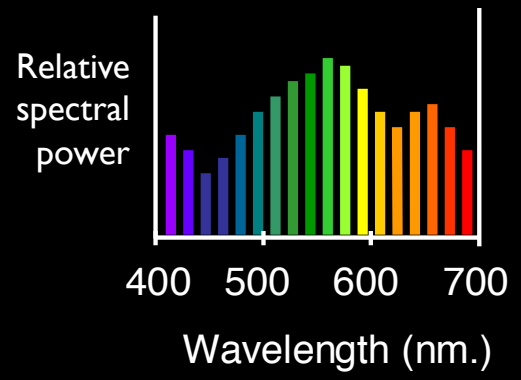

Slide credit: A. Efros

The Physics of light

Some examples of the spectra of light sources

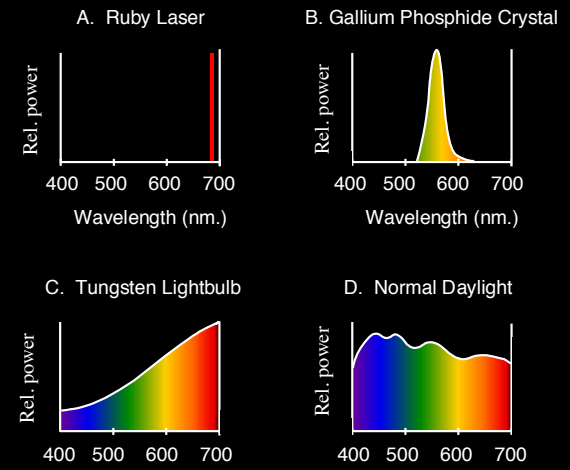

Slide credit: A. Efros

The Physics of light

Some examples of the reflectance spectra of surfaces

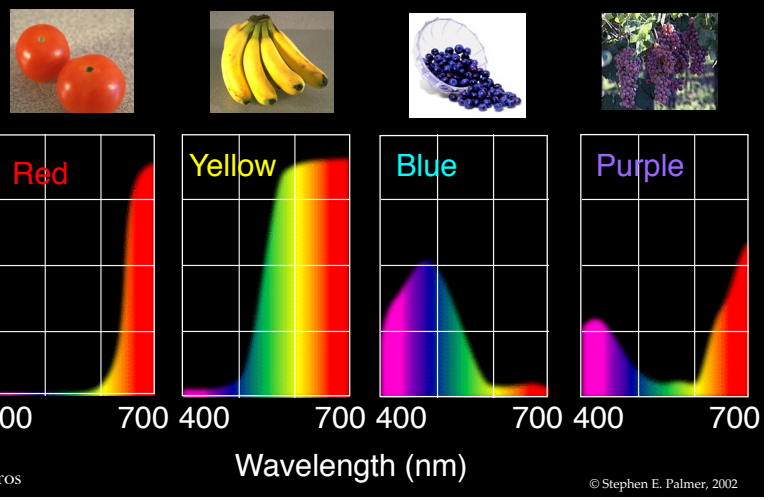

Slide credit: A. Efros

Image formation

- What determines the brightness of an image pixel?

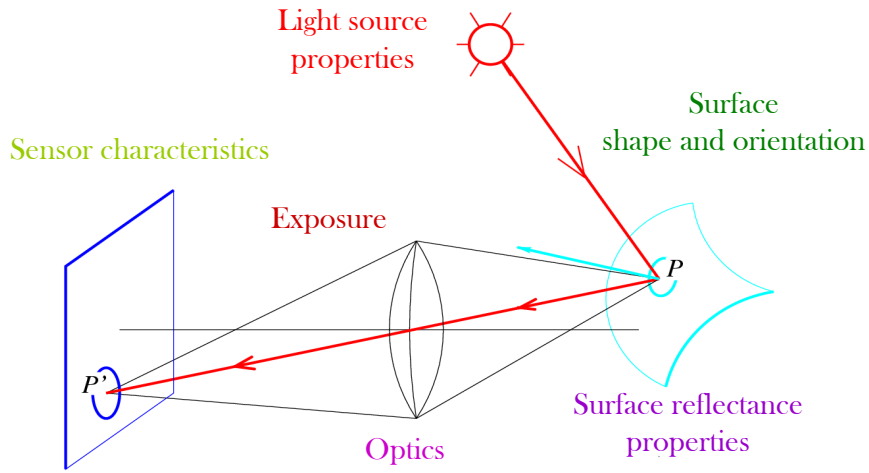

Slide credit: L. Fei-Fei

Color mixing

Cartoon spectra for color names:

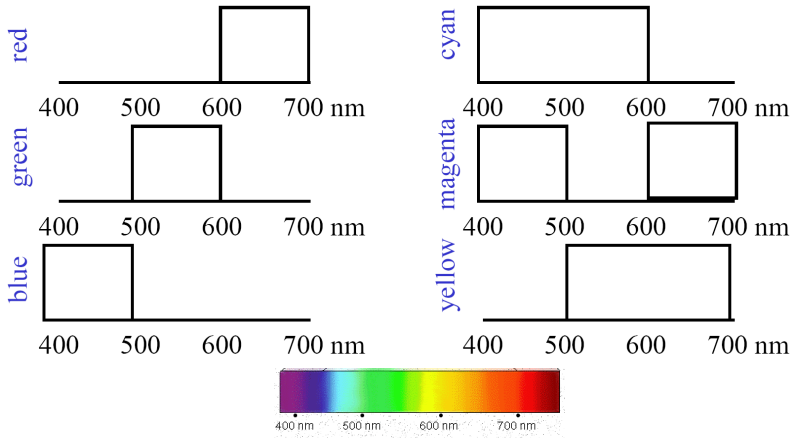

Credit: W. Freeman

Additive color mixing

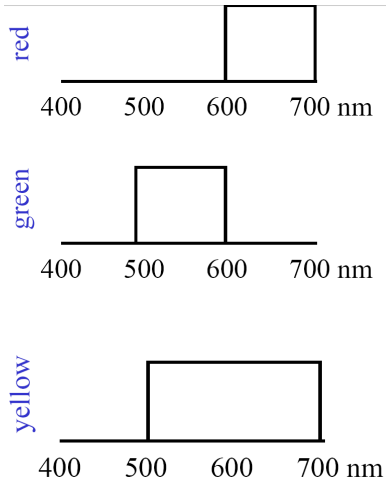

Colors combine by *adding* color spectra

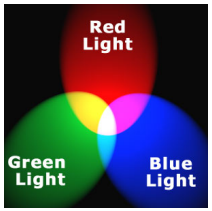

Light *adds* to black.

Credit: W. Freeman

Examples of additive color systems

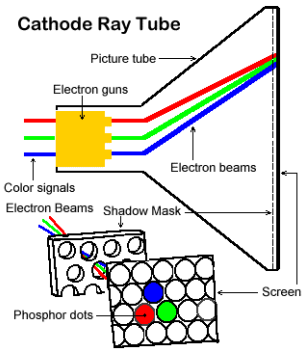

CRT phosphors

multiple projectors

http://www.jegsworks.com
http://www.crtprojectors.co.uk/

Slide credit: K. Grauman

Subtractive color mixing

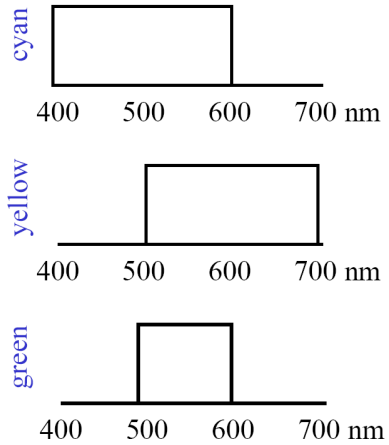

Colors combine by *multiplying* color spectra.

Pigments *remove* color from incident light (white).

Credit: W. Freeman

Examples of subtractive color systems

- Printing on paper
- Crayons
- Photographic film

Slide credit: K. Grauman

Interaction of light and surfaces

- Reflected color is the result of interaction of light source spectrum with surface reflectance

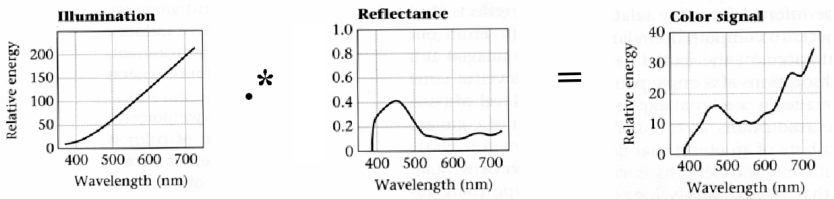

Slide credit: A. Efros

Reflection from colored surface

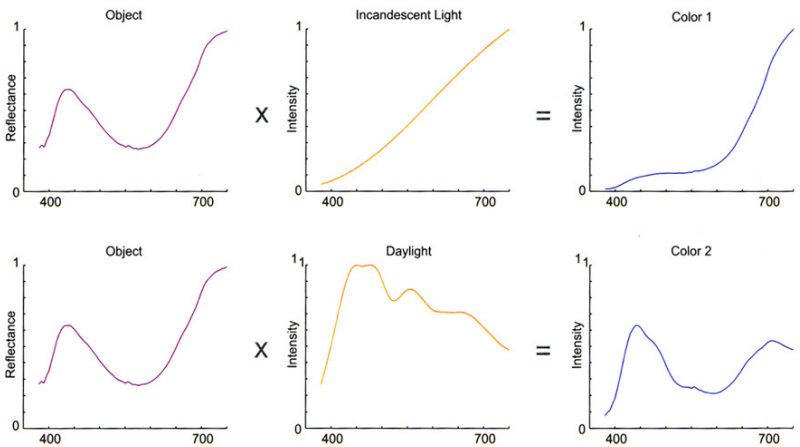

[Stone 2003]

Slide credit: S. Marschner

The Eye

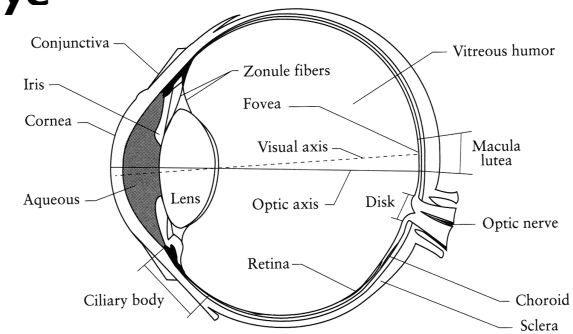

- **Iris** - colored annulus with radial muscles
- **Pupil** - the hole (aperture) whose size is controlled by the iris
- **Lens** - changes shape by using ciliary muscles (to focus on objects at different distances)
- **Retina** - photoreceptor cells

Slide credit: S. Seitz

The eye as a measurement device

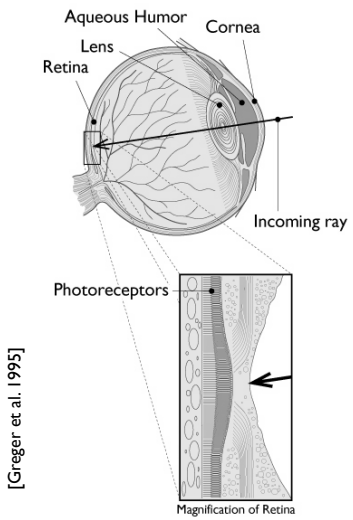

[Greger et al. 1995]

- We can model the low-level behavior of the eye by thinking of it as a light-measuring machine
  - its optics are much like a camera
  - its detection mechanism is also much like a camera
- Light is measured by the *photoreceptors* in the retina
  - they respond to visible light
  - different types respond to different wavelengths
- **The human eye is a camera!**

Slide credit: S. Marschner

Layers of the retina

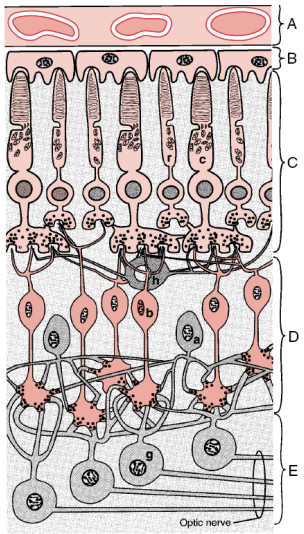

Slide credit: S. Ullman

Layers of the retina

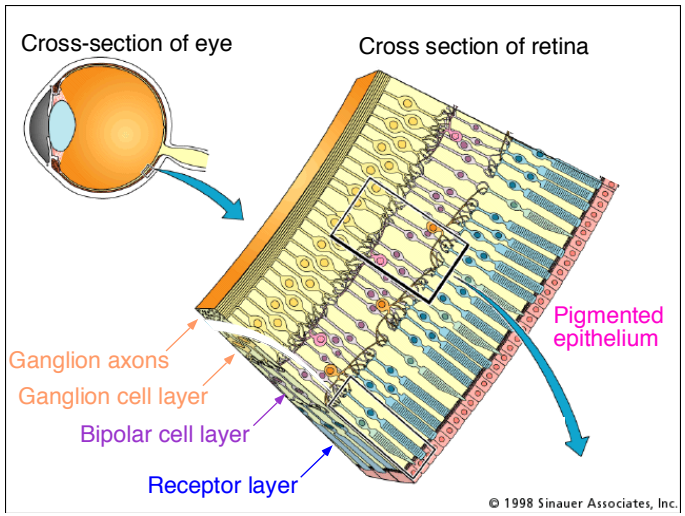

Slide credit: J. Hays

Receptors Density - Fovea

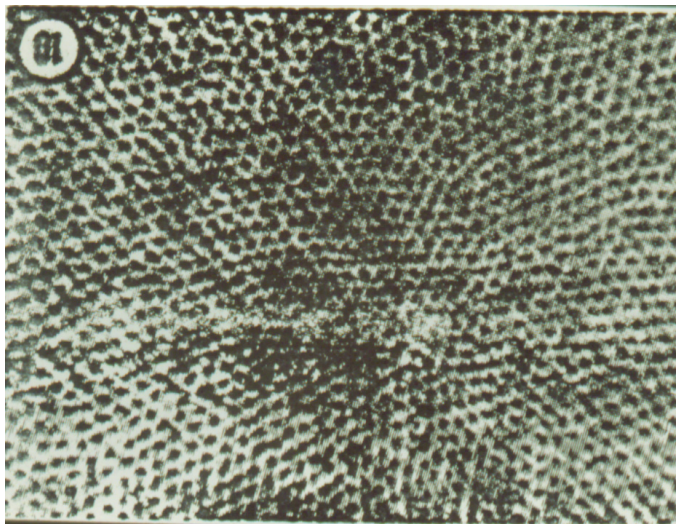

Slide credit: S. Ullman

Receptors Density - Fovea

64	66	76	85	99	100	101	101	106	112	117	118	105	77	57	50	51	43	52	55	62
65	69	76	84	97	89	93	107	121	121	121	122	125	101	71	43	45	41	52	52	68
66	72	78	83	91	86	91	102	108	104	106	113	136	118	86	43	49	47	60	55	64
73	79	83	85	94	93	90	83	79	79	85	92	124	124	108	62	58	43	57	57	64
78	84	86	86	69	71	68	68	86	108	115	109	117	135	139	93	73	37	49	58	70
75	75	73	77	75	80	62	84	90	94	98	102	102	110	114	100	80	58	51	51	51
77	72	73	83	84	91	80	77	71	70	73	80	80	87	99	103	93	67	53	50	51
74	66	69	88	98	101	95	65	56	55	55	60	64	70	93	114	112	82	56	47	53
64	59	66	86	108	103	98	54	52	57	54	54	67	77	103	124	125	96	64	46	53
56	57	66	83	112	108	104	59	55	60	59	60	78	94	115	125	121	98	68	43	46
56	58	66	80	114	121	117	85	71	67	69	76	87	101	116	117	112	94	68	43	46
61	57	61	77	111	125	119	114	98	87	87	94	97	102	111	113	108	90	65	43	44
63	52	54	73	103	117	107	126	119	108	103	104	106	103	108	115	112	91	65	48	42
66	63	58	63	94	115	120	108	102	104	106	108	105	108	107	105	105	97	72	47	41
68	65	58	61	86	108	115	106	102	103	103	104	98	99	97	97	103	101	81	57	43
72	68	62	64	78	102	111	105	101	101	101	103	99	98	96	97	104	104	86	63	48
74	71	64	64	69	93	104	99	94	93	96	101	99	101	102	103	108	106	90	69	53

Slide credit: S. Ullman

Human Photoreceptors

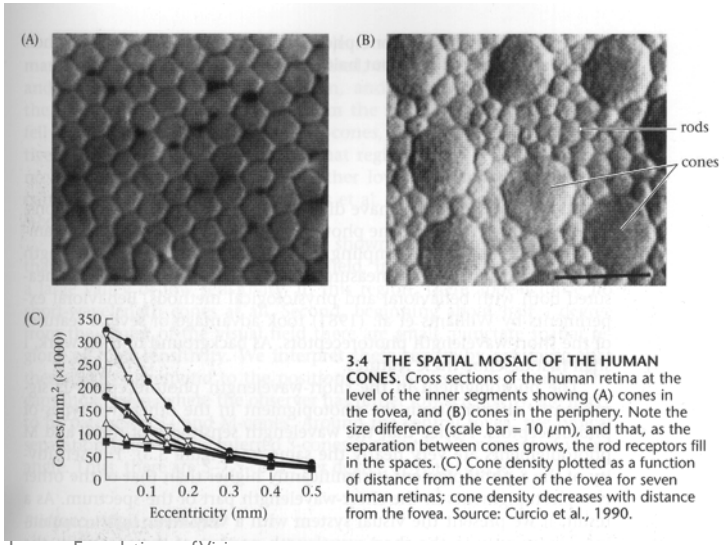

3.4 THE SPATIAL MOSAIC OF THE HUMAN CONES. Cross sections of the human retina at the level of the inner segments showing (A) cones in the fovea, and (B) cones in the periphery. Note the size difference (scale bar = 10 μm), and that, as the separation between cones grows, the rod receptors fill in the spaces. (C) Cone density plotted as a function of distance from the center of the fovea for seven human retinas; cone density decreases with distance from the fovea. Source: Curcio et al., 1990.

Images: Foundations of Vision, by Brian Wandell, Sinauer Assoc., 1995

Slide Credit: B. Freeman and A. Torralba

Human eye photoreceptor spectral sensitivities

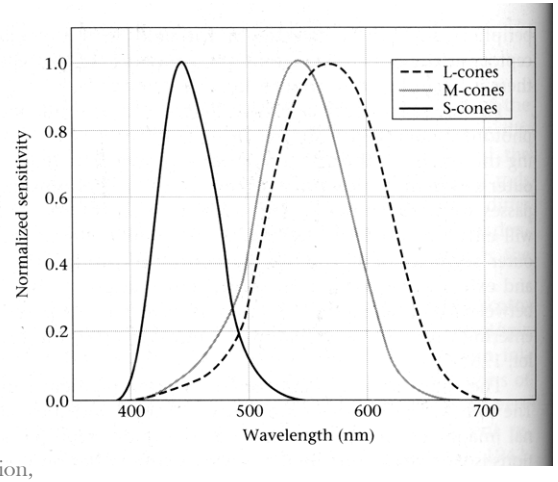

3.3 SPECTRAL SENSITIVITIES OF THE L-, M-, AND S-CONES in the human eye. The measurements are based on a light source at the cornea, so that the wavelength loss due to the cornea, lens, and other inert pigments of the eye plays a role in determining the sensitivity. Source: Stockman and MacLeod, 1993.

Images: Foundations of Vision, by Brian Wandell, Sinauer Assoc., 1995

Slide Credit: B. Freeman and A. Torralba

Two types of light-sensitive receptors

Cones

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

Rods

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

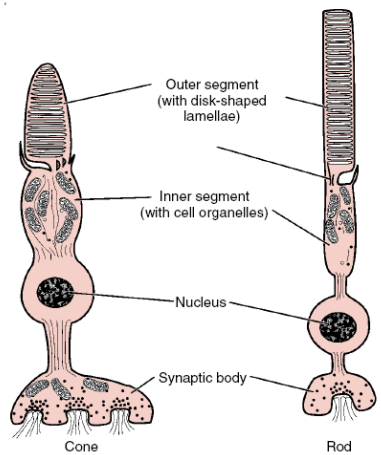

Images by Shimon Ullman

Slide credit: A. Efros

Rods and cones

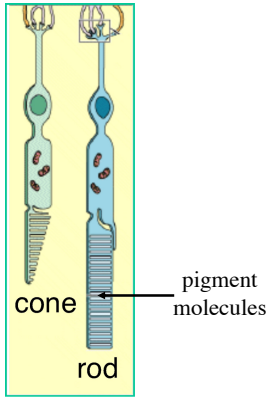

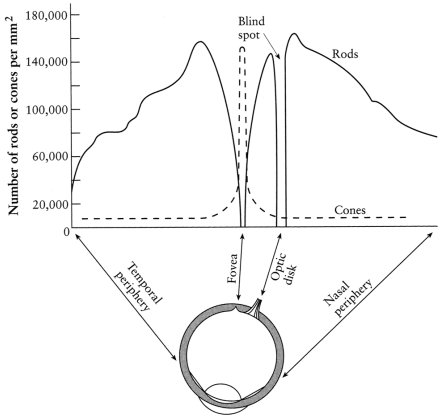

Rods are responsible for intensity, cones for color perception

Rods and cones are non-uniformly distributed on the retina

- Fovea - Small region (1 or 2°) at the center of the visual field containing the highest density of cones (and no rods)

Slide credit: S. Seitz

Physiology of Color Vision

Three kinds of cones:

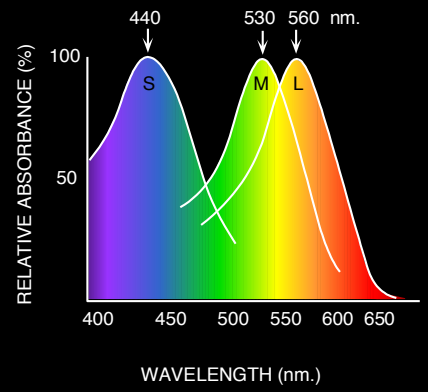

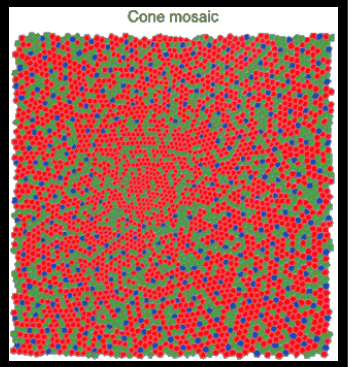

- Ratio of L to M to S cones: approx. 10:5:1
- Almost no S cones in the center of the fovea

Slide credit: A. Efros

Color perception

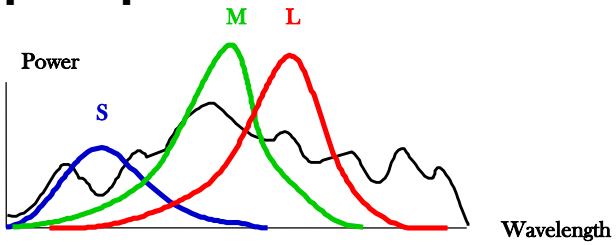

Rods and cones act as filters on the spectrum

- To get the output of a filter, multiply its response curve by the spectrum, integrate over all wavelengths
  - Each cone yields one number

Q: How can we represent an entire spectrum with 3 numbers?

A: We can't! Most of the information is lost.

- As a result, two different spectra may appear indistinguishable

Slide credit: S. Seitz

Today

- Perception of color and light
- Color spaces
- Point operations
- Histogram processing

Acknowledgement: This part is mostly adapted from the slides prepared by Steve Marschner, James Hays, Ali Farhadi and Anat Levin

Recall: Digital images

- Sample the 2D space on a regular grid
- Quantize each sample (round to nearest integer)
- Image thus represented as a matrix of integer values.

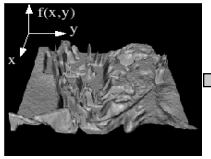

j →

i ↓

62	79	23	119	120	105	4	0
10	10	9	62	12	78	34	0
10	58	197	46	46	0	0	48
176	135	5	188	191	68	0	49
2	1	1	29	26	37	0	77
0	89	144	147	187	102	62	208
255	252	0	166	123	62	0	31
166	63	127	17	1	0	99	30

2D

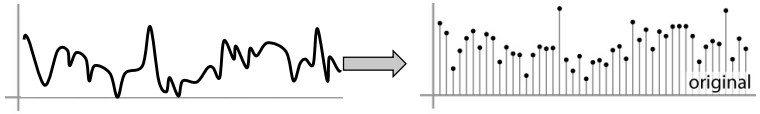

1D

Slide credit: K. Grauman, S. Seitz

Color Images: Bayer Grid

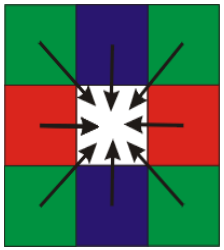

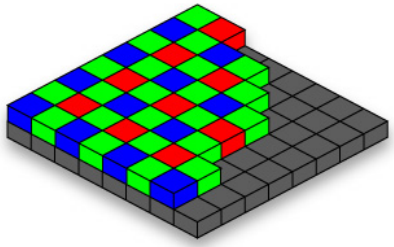

- Estimate RGB at 'G' cells from neighboring values

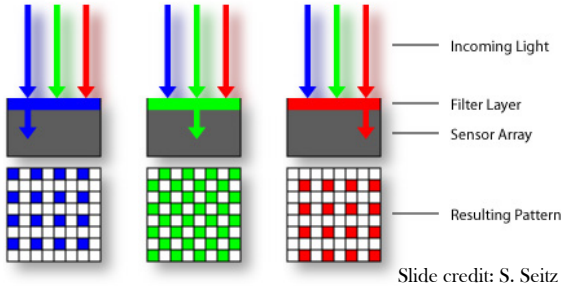

http://www.cooldictionary.com/words/Bayer-filter.wikipedia

Slide credit: S. Seitz

Digital color images

Color images, RGB color space

R

G

B

Slide credit: K. Grauman

Images

- Images represented as a matrix
- Suppose we have a NxM RGB image called "im"
  - `im(1,1,1)` = top-left pixel value in R-channel
  - `im(y,x,b)` = y pixels down, x pixels to right in the bth channel
  - `im(N,M,3)` = bottom-right pixel in B-channel

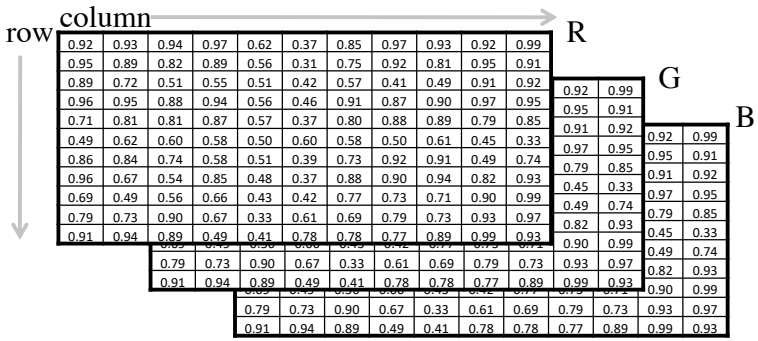

Slide credit: D. Hoiem

Color spaces

- How can we represent color?

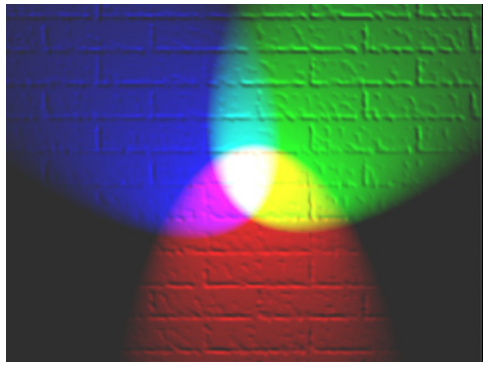

Image from http://en.wikipedia.org/wiki/File:RGB_illumination.jpg

Slide credit: D. Hoiem

Color spaces: RGB

- Single wavelength primaries
- makes a particular monitor RGB standard
- Good for devices (e.g., phosphors for monitor), but not for perception

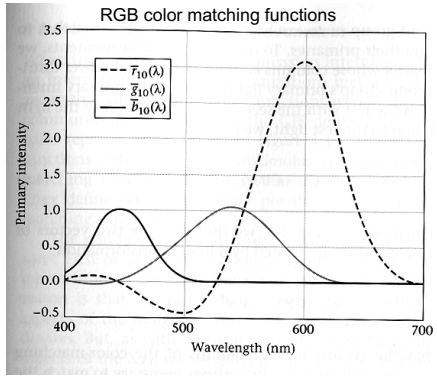

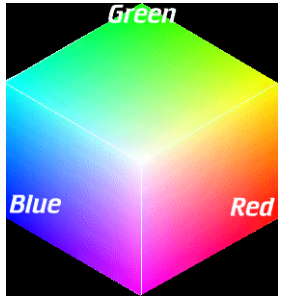

Slide credit: K. Grauman, S. Marschner

Color spaces: RGB

Default color space

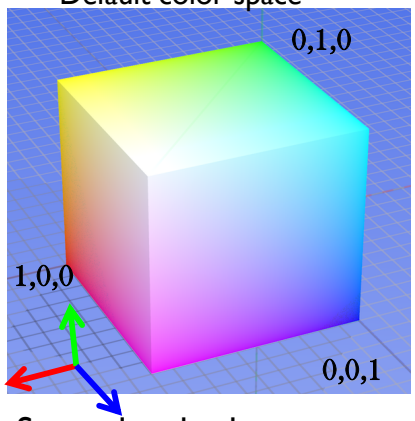

R
(G=0,B=0)

G
(R=0,B=0)

B
(R=0,G=0)

Some drawbacks
- Strongly correlated channels
- Non-perceptual

Image from: http://en.wikipedia.org/wiki/File:RGB_color_solid_cube.png

Slide credit: D. Hoiem

Color spaces: CIE XYZ

- Standardized by CIE (*Commission Internationale de l'Eclairage*, the standards organization for color science)
- Based on three "imaginary" primaries **X**, **Y**, and **Z**
  - imaginary = only realizable by spectra that are negative at some wavelengths
  - separates out luminance: **X**, **Z** have zero luminance, so *Y* tells you the luminance by itself

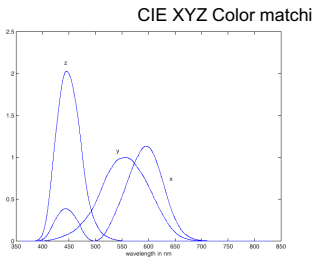

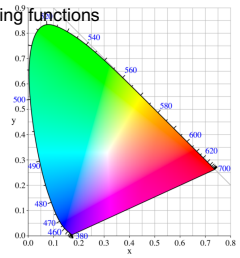

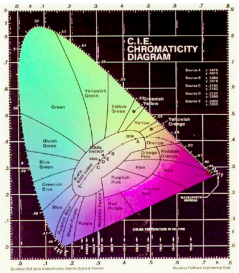

Slide credit: K. Grauman, S. Marschner

Color spaces: CIE XYZ

- Standardized by CIE (*Commission Internationale de l'Eclairage,* the standards organization for color science)
- Based on three "imaginary" primaries **X**, **Y**, and **Z**
  - imaginary = only realizable by spectra that are negative at some wavelengths
  - separates out luminance: **X**, **Z** have zero luminance, so *Y* tells you the luminance by itself

$$\begin{bmatrix} X \\ Y \\ Z \end{bmatrix} = \frac{1}{0.17697} \begin{bmatrix} 0.49 & 0.31 & 0.20 \\ 0.17697 & 0.81240 & 0.01063 \\ 0.00 & 0.01 & 0.99 \end{bmatrix} \begin{bmatrix} R \\ G \\ B \end{bmatrix}$$

Slide credit: K. Grauman, S. Marschner

Perceptually organized color spaces

- Artists often refer to colors as *tints*, *shades*, and *tones* of pure pigments
  - tint: mixture with white
  - shade: mixture with black
  - tones: mixture with black and white
  - gray: no color at all (aka. neutral)

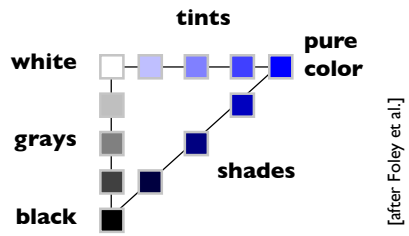

- This seems intuitive
  - tints and shades are inherently related to the pure color
    - "same" color but lighter, darker, paler, etc.

Slide credit: S. Marschner

Perceptual dimensions of color

- Hue
  - the "kind" of color, regardless of attributes
  - colorimetric correlate: dominant wavelength
  - artist's correlate: the chosen pigment color
- Saturation
  - the "colorfulness"
  - colorimetric correlate: purity
  - artist's correlate: fraction of paint from the colored tube
- Lightness (or value)
  - the overall amount of light
  - colorimetric correlate: luminance
  - artist's correlate: tints are lighter, shades are darker

Slide credit: S. Marschner

Color spaces: HSV

- **H**ue, **S**aturation, **V**alue
- Nonlinear – reflects topology of colors by coding **hue** as an angle
- Matlab: `hsv2rgb`, `rgb2hsv`.

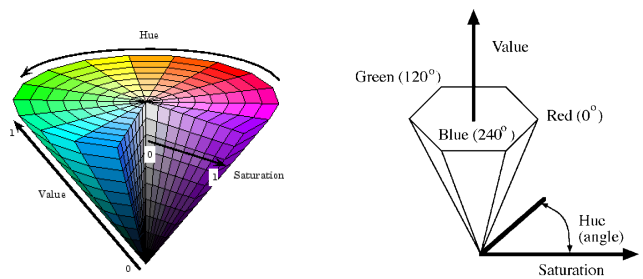

Image from mathworks.com

Slide credit: K. Grauman

Color spaces: HSV

- **H**ue, **S**aturation, **V**alue
- Nonlinear – reflects topology of colors by coding **hue** as an angle
- Matlab: `hsv2rgb`, `rgb2hsv`.

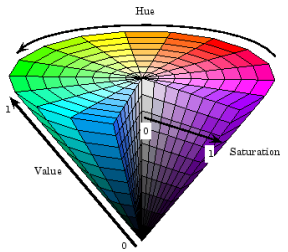

$$H = \begin{cases} \left(\dfrac{G' - B'}{MAX - MIN}\right)/6, & \text{if } R' = MAX \\ \left(2 + \dfrac{B' - R'}{MAX - MIN}\right)/6, & \text{if } G' = MAX \\ \left(4 + \dfrac{R' - G'}{MAX - MIN}\right)/6, & \text{if } B' = MAX \end{cases}$$

$$S = \frac{MAX - MIN}{MAX}$$

$$V = MAX$$

Image from mathworks.com

Slide credit: K. Grauman

Color spaces: HSV

Intuitive color space

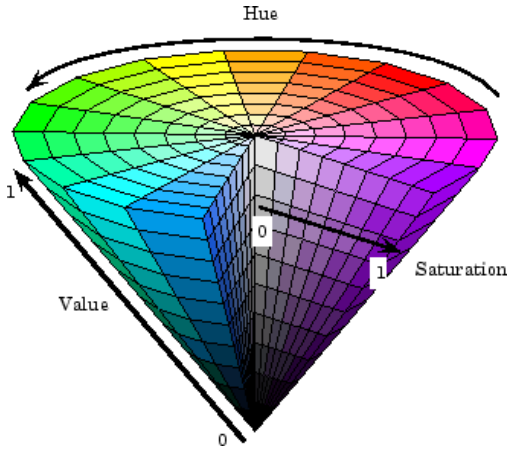

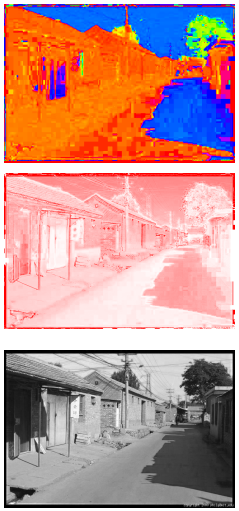

Slide credit: D. Hoiem

Color spaces: YCbCr

Fast to compute, good for compression, used by TV

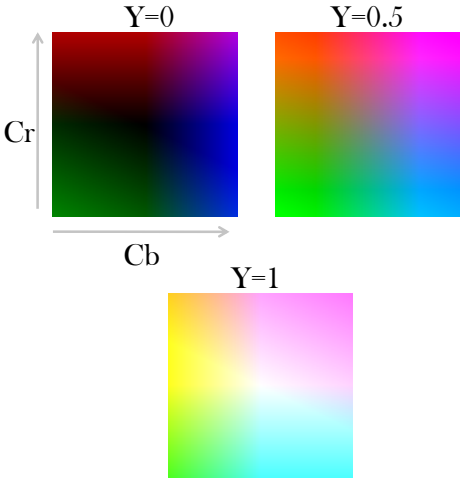

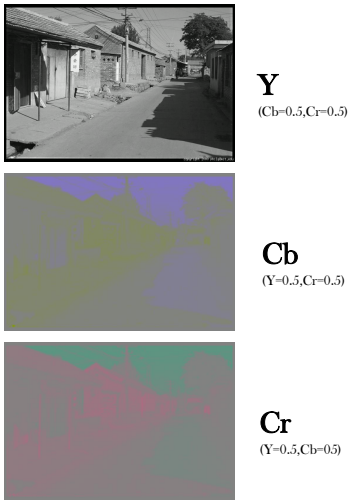

Slide credit: D. Hoiem

Color spaces: YCbCr

Fast to compute, good for compression, used by TV

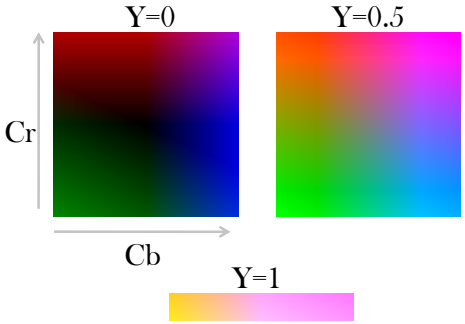

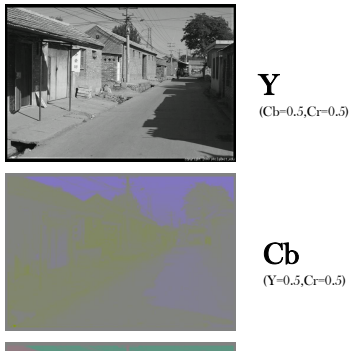

$$\begin{bmatrix} Y' \\ C_b \\ C_r \end{bmatrix} = \begin{bmatrix} 0.299 & 0.587 & 0.114 \\ -0.168736 & -0.331264 & 0.5 \\ 0.5 & -0.418688 & -0.081312 \end{bmatrix} \begin{bmatrix} R' \\ G' \\ B' \end{bmatrix} + \begin{bmatrix} 0 \\ 128 \\ 128 \end{bmatrix}$$

Slide credit: D. Hoiem

Distances in color space

- Are distances between points in a color space perceptually meaningful?

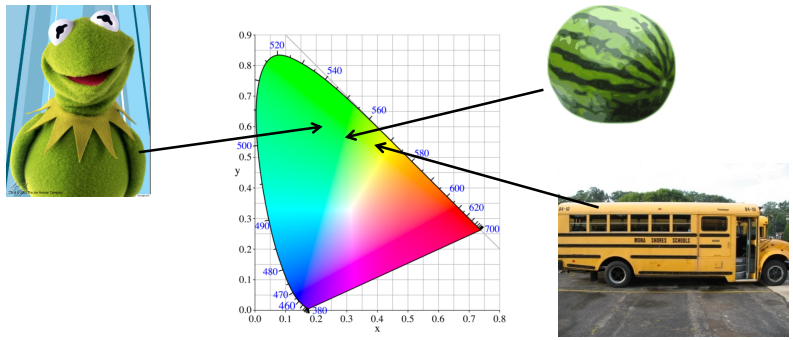

Slide credit: K. Grauman

Distances in color space

- Not necessarily: CIE XYZ is not a uniform color space, so magnitude of differences in coordinates are poor indicator of color "distance".

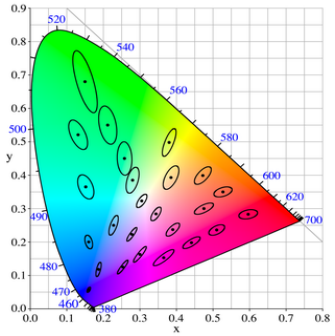

McAdam ellipses:
Just noticeable differences in color

Slide credit: K. Grauman

Uniform color spaces

- Attempt to correct this limitation by remapping color space so that just-noticeable differences are contained by circles → distances more perceptually meaningful.

- Examples:
  - CIE u'v'
  - CIE Lab

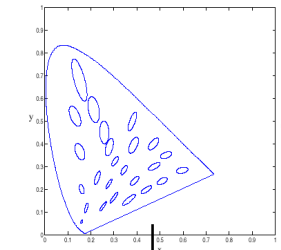

CIE XYZ

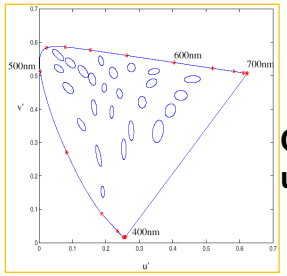

CIE u' v'

Slide credit: K. Grauman

Perceptually uniform spaces

- Two major spaces standardized by CIE
  - designed so that equal differences in coordinates produce equally visible differences in color
  - by remapping color space so that just-noticeable differences are contained by circles→ distances more perceptually meaningful.
  - LUV: earlier, simpler space; *L**, *u**, *v**
  - LAB: more complex but more uniform: *L**, *a**, *b**
  - both separate luminance from chromaticity
  - including a gamma-like nonlinear component is important

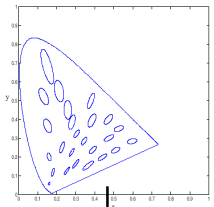

CIE XYZ

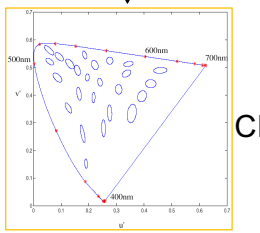

CIE u' v'

Slide credit: K. Grauman, S. Marschner

Color spaces: L*a*b*

"Perceptually uniform"* color space

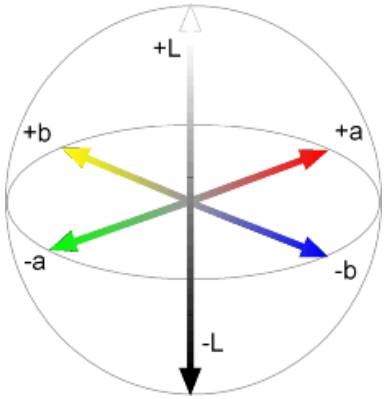

L
(a=0,b=0)

a
(L=65,b=0)

b
(L=65,a=0)

Slide credit: D. Hoiem

Color spaces: L*a*b*

"Perceptually uniform"* color space

$$L^* = 116 f\left(\frac{Y}{Y_n}\right)$$

$$f(t) = \begin{cases} t^{1/3} & t > \delta^3 \\ t/(3\delta^2) + 2\delta/3 & \text{else,} \end{cases}$$

$$a^* = 500\left[f\left(\frac{X}{X_n}\right) - f\left(\frac{Y}{Y_n}\right)\right]$$

$$b^* = 200\left[f\left(\frac{Y}{Y_n}\right) - f\left(\frac{Z}{Z_n}\right)\right]$$

(X_n, Y_n, Z_n): measured white point

L
(a=0,b=0)

a
(L=65,b=0)

b
(L=65,a=0)

Slide credit: D. Hoiem

Most information in intensity

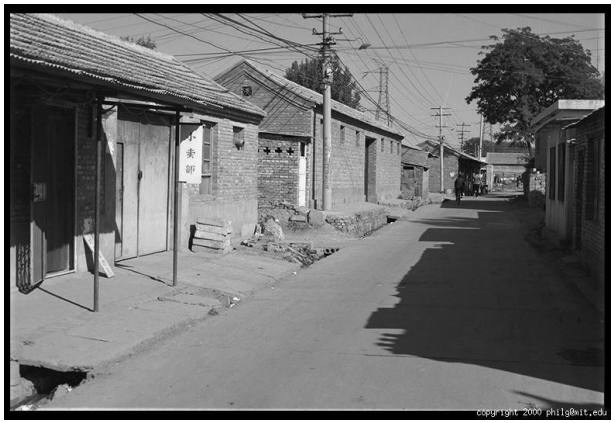

Only intensity shown – constant color

Slide credit: D. Hoiem

Most information in intensity

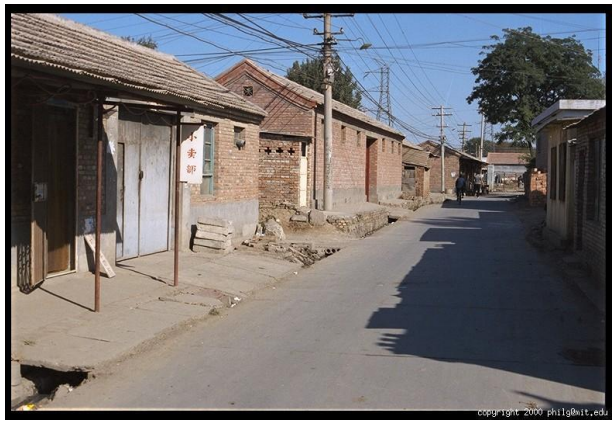

Original image

Slide credit: D. Hoiem

Today

- Perception of color and light
- Color spaces
- Point operations
- Histogram processing

Digital images

- Sample the 2D space on a regular grid
- Quantize each sample (round to nearest integer)

- Image thus represented as a matrix of integer values.

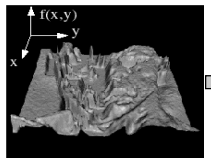

j

i

62	79	23	119	120	105	4	0
10	10	9	62	12	78	34	0
10	58	197	46	46	0	0	48
176	135	5	188	191	68	0	49
2	1	1	29	26	37	0	77
0	89	144	147	187	102	62	208
255	252	0	166	123	62	0	31
166	63	127	17	1	0	99	30

2D

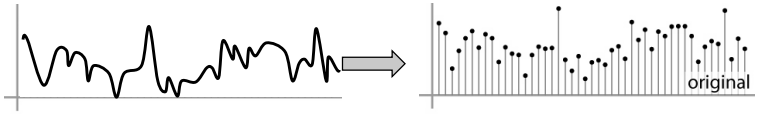

1D

Slide credit: K. Grauman, S. Seitz

Image Transformations

- $g(x,y)=T[f(x,y)]$

$g(x,y)$: output image $f(x,y)$: input image M: transformation function

1. Point operations: operations on single pixels
2. Spatial filtering: operations considering pixel neighborhoods
3. Global methods: operations considering whole image

Image Transformations

- $g(x,y)=T[f(x,y)]$

$g(x,y)$: output image $f(x,y)$: input image M: transformation function

1. Point operations: operations on single pixels
2. Spatial filtering: operations considering pixel neighborhoods
3. Global methods: operations considering whole image

$$g(x,y)=M(f(x,y))$$

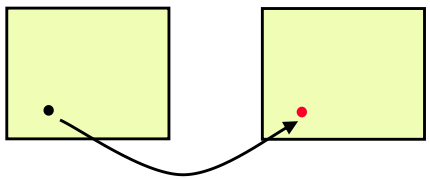

Image Transformations

- $g(x,y)=M[f(x,y)]$

$g(x,y)$: output image $f(x,y)$: input image M: transformation function

1. Point operations: operations on single pixels
2. <u>Spatial filtering: operations considering pixel neighborhoods</u>
3. Global methods: operations considering whole image

$$g(x,y)=M(\{f(i,j)\,|\,(i,j)\in N(x,y)\})$$

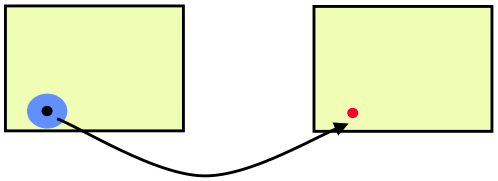

Point operations

- Smallest possible neighborhood is of size 1x1
- Process each point independently of the others
- Output image *g* depends only on the value of *f* at a single point (x,y)
- Map each pixel's value to a new value
- Transformation function *T* remaps the sample's value:

  $s = T(r)$

  where
  - *r* is the value at the point in question
  - *s* is the new value in the processed result
  - *T* is a *intensity transformation* function

Point operations

- Is mapping one color space to another (e.g. RGB2HSV) a point operation?
- Is image arithmetic a point operation?
- Is performing geometric transformations a point operation?
  - Rotation
  - Translation
  - Scale change
  - etc.

Sample intensity transformation functions

- Image negatives
- Log transformations
  - Compresses the dynamic range of images
- Power-law transformations
  - Gamma correction

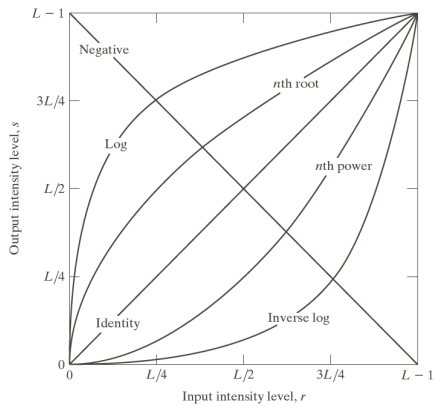

Point Processing Examples

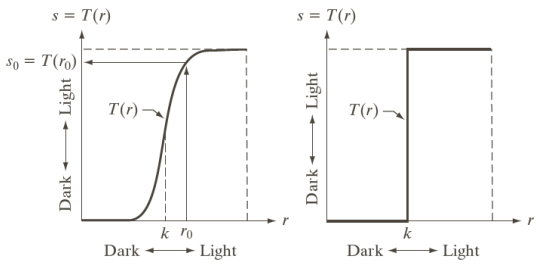

produces an image of higher contrast than the original by darkening the intensity levels below k and brightening intensities above k

produces a binary (two-intensity level) image

Image Mean

$$I_{av} = \frac{\sum_i \sum_j I(i,j)}{\sum_i \sum_j 1}$$

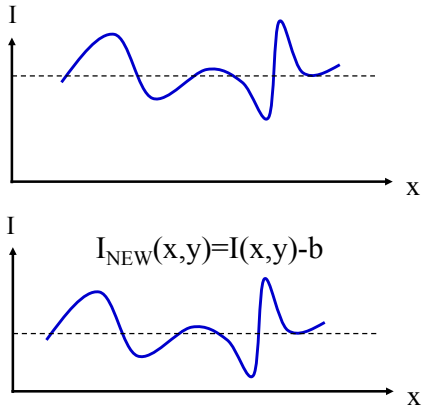

Slide credit: Y. Hel-Or

Image Mean

Changing the image mean

Slide credit: Y. Hel-Or

Image Negative

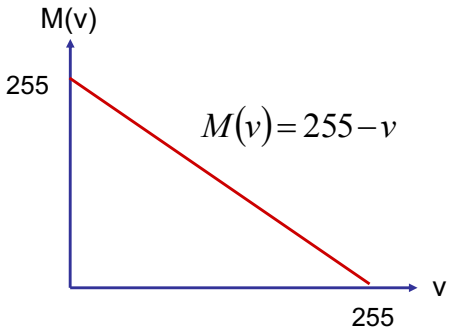

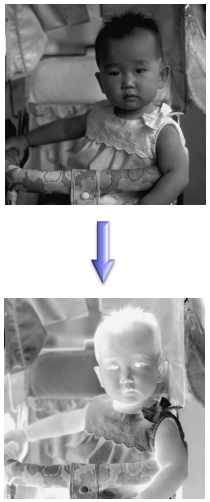

Slide credit: Y. Hel-Or

Dynamic range

- Dynamic range $R_d = I_{max} / I_{min}$, or $(I_{max} + k) / (I_{min} + k)$
  - determines the degree of image contrast that can be achieved
  - a major factor in image quality
- Ballpark values
  - Desktop display in typical conditions: 20:1
  - Photographic print: 30:1
  - High dynamic range display: 10,000:1

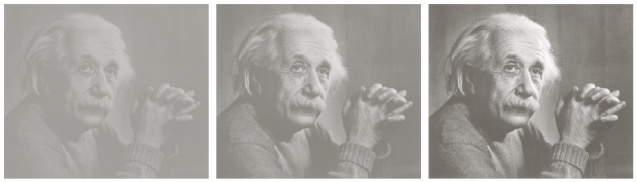

low contrast medium contrast high contrast

Slide credit: S. Marschner

Point Operations: Contrast stretching and Thresholding

- Contrast stretching: produces an image of higher contrast than the original

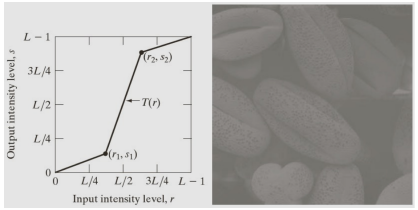

- Thresholding: produces a binary (two-intensity level) image

Point Operations: Contrast stretching and Thresholding

- Contrast stretching: produces an image of higher contrast than the original

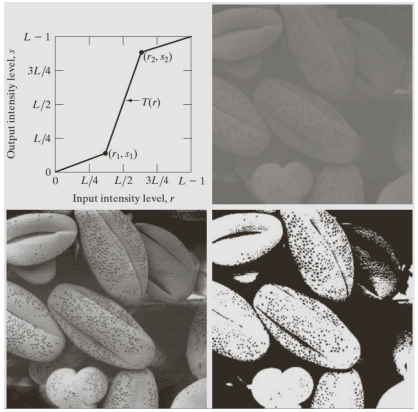

- Thresholding: produces a binary (two-intensity level) image

Histogram

- Histogram: a discrete function *h(r)* which counts the number of pixels in the image having intensity *r*
- If *h(r)* is normalized, it measures the probability of occurrence of intensity level *r* in an image

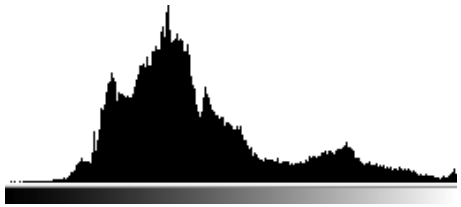

Point Operations

- What can you say about the image having the following histogram?

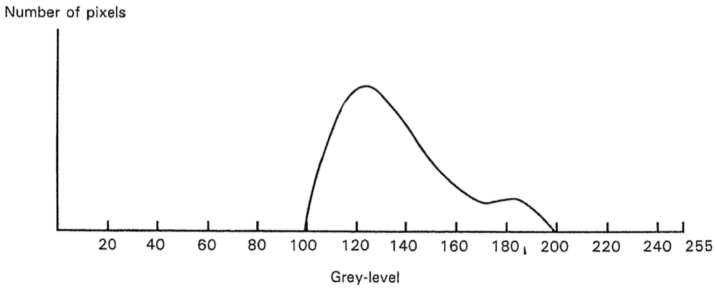

- A low contrast image
- How we can process the image so that it has a better visual quality?

Point Operations

- How we can process the image so that it has a better visual quality?

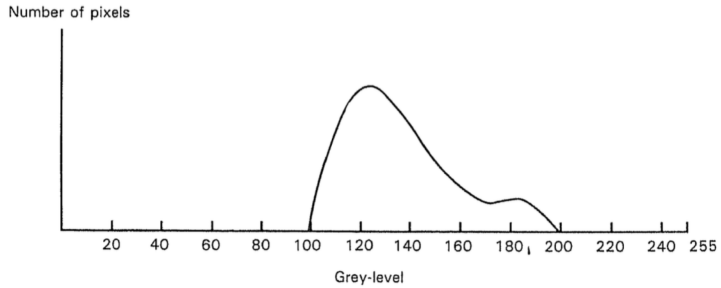

- Answer is contrast stretching!

Point Operations

- Let us devise an appropriate point operation.

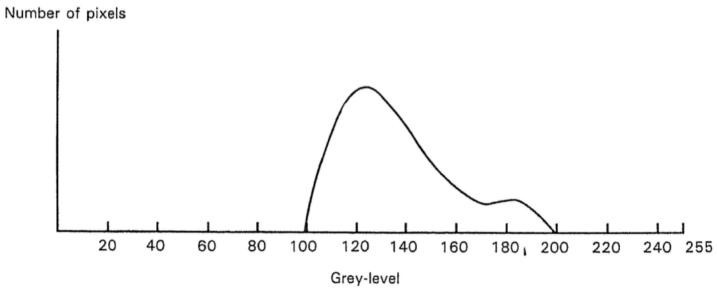

- Shift all values so that the observable pixel range starts at 0.

Point Operations

- Let us devise an appropriate point operation.

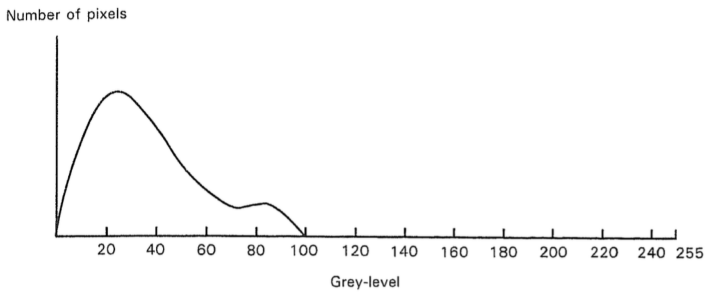

- Now, scale everything in the range 0-100 to 0-255.

Point Operations

- Let us devise an appropriate point operation.

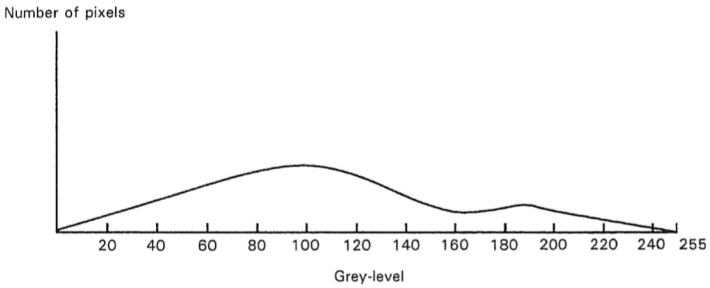

- What is the corresponding transformation function?
- T(r) = 2.55*(r-100)

Point Operations: Intensity-level Slicing

- highlights a certain range of intensities

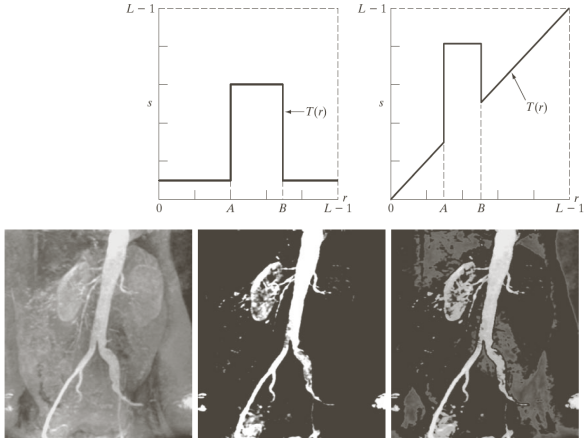

Point Operations: Intensity-level Slicing

- highlights a certain range of intensities

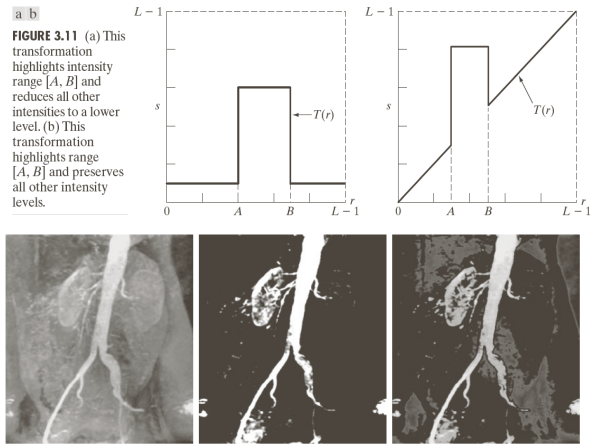

Intensity encoding in images

- Recall that the pixel values determine how bright that pixel is.
- Bigger numbers are (usually) brighter
- *Transfer function*: function that maps input pixel value to luminance of displayed image

$$I = f(n) \quad f : [0, N] \rightarrow [I_{\min}, I_{\max}]$$

- What determines this function?
  - physical constraints of device or medium
  - desired visual characteristics

adapted from: S. Marschner

What this projector does?

- Something like this:

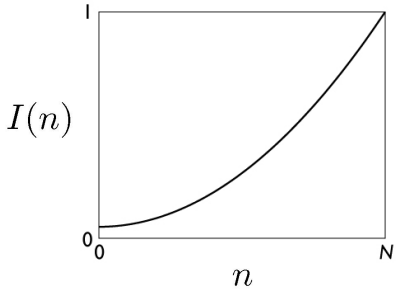

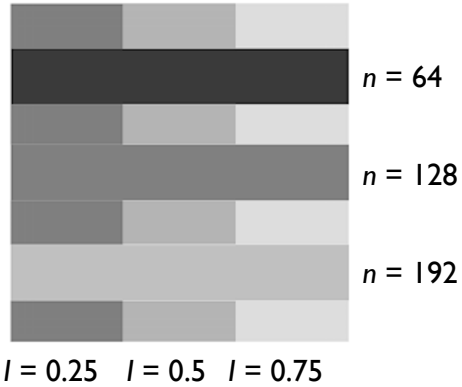

adapted from: S. Marschner

Constraints on transfer function

- Maximum displayable intensity, I_{max}
  - how much power can be channeled into a pixel?
    - LCD: backlight intensity, transmission efficiency (<10%)
    - projector: lamp power, efficiency of imager and optics
- Minimum displayable intensity, I_{min}
  - light emitted by the display in its "off" state
    - e.g. stray electron flux in CRT, polarizer quality in LCD
- Viewing flare, k: light reflected by the display
  - very important factor determining image contrast in practice
    - 5% of I_{max} is typical in a normal office environment [sRGB spec]
    - much effort to make very black CRT and LCD screens
    - all-black decor in movie theaters

Transfer function shape

- Desirable property: the change from one pixel value to the next highest pixel value should not produce a visible contrast
  - otherwise smooth areas of images will show visible bands
- What contrasts are visible?
  - rule of thumb: under good conditions we can notice a 2% change in intensity
  - therefore we generally need smaller quantization steps in the darker tones than in the lighter tones
  - most efficient quantization is logarithmic

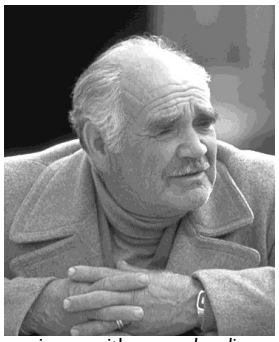

[Philip Greenspun]

an image with severe *banding*

Slide credit: S. Marschner

How many levels are needed?

- Depends on dynamic range
  - 2% steps are most efficient:

$$0 \mapsto I_{\min}; 1 \mapsto 1.02 I_{\min}; 2 \mapsto (1.02)^2 I_{\min}; \ldots$$

  - log 1.02 is about 1/120, so 120 steps per decade of dynamic range
    - 240 for desktop display
    - 360 to print to film
    - 480 to drive HDR display
- If we want to use linear quantization (equal steps)
  - one step must be < 2% (1/50) of I_{min}
  - need to get from ~0 to $I_{min} \cdot R_d$ so need about 50 R_d levels
    - 1500 for a print; 5000 for desktop display; 500,000 for HDR display
- Moral: 8 bits is just barely enough for low-end applications
  - but only if we are careful about quantization

Slide credit: S. Marschner

Intensity quantization in practice

- Option 1: linear quantization $I(n) = (n/N)\, I_{\max}$
  - pro: simple, convenient, amenable to arithmetic
  - con: requires more steps (wastes memory)
  - need 12 bits for any useful purpose; more than 16 for HDR
- Option 2: power-law quantization $I(n) = (n/N)^{\gamma}\, I_{\max}$
  - pro: fairly simple, approximates ideal exponential quantization
  - con: need to linearize before doing pixel arithmetic
  - con: need to agree on exponent
  - 8 bits are OK for many applications; 12 for more critical ones
- Option 2: floating-point quantization $I(x) = (x/w)\, I_{\max}$
  - pro: close to exponential; no parameters; amenable to arithmetic
  - con: definitely takes more than 8 bits
  - 16–bit "half precision" format is becoming popular

Slide credit: S. Marschner

Why gamma?

- Power-law quantization, or *gamma correction* is most popular
- Original reason: CRTs are like that
  - intensity on screen is proportional to (roughly) voltage^2
- Continuing reason: inertia + memory savings
  - inertia: gamma correction is close enough to logarithmic that there's no sense in changing
  - memory: gamma correction makes 8 bits per pixel an acceptable option

Slide credit: S. Marschner

Gamma quantization

$$I(n) = (n/N)^{\gamma}\, I_{\max}$$

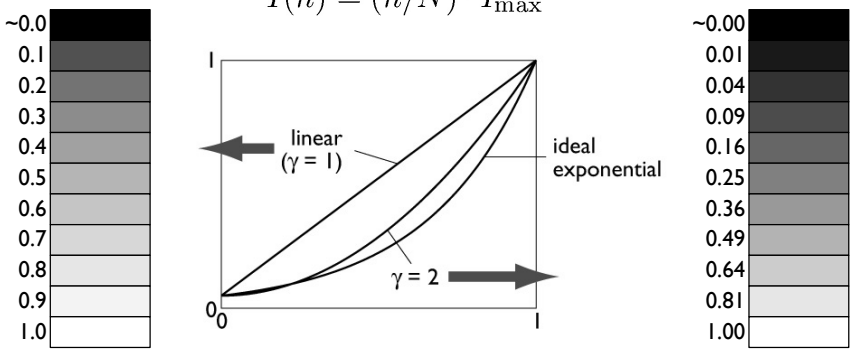

- Close enough to ideal perceptually uniform exponential

Slide credit: S. Marschner

Gamma correction

- Sometimes (often, in graphics) we have computed intensities *a* that we want to display linearly
- In the case of an ideal monitor with zero black level,

$$I(n) = (n/N)^{\gamma}$$

(where N = $2^n - 1$ in *n* bits). Solving for *n*:

$$n = N a^{\frac{1}{\gamma}}$$

- This is the "gamma correction" recipe that has to be applied when computed values are converted to 8 bits for output
  - failing to do this (implicitly assuming gamma = 1) results in dark, oversaturated images

Slide credit: S. Marschner

Gamma correction

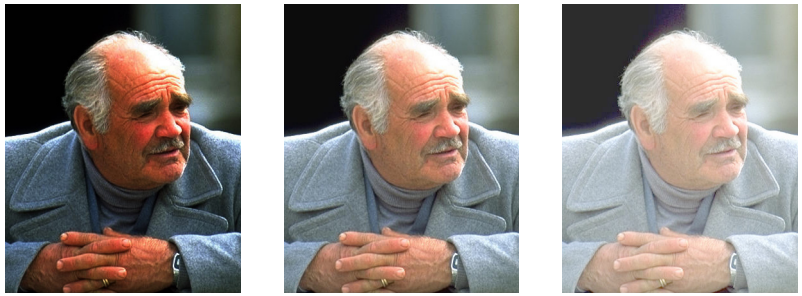

corrected for γ lower than display

OK

corrected for γ higher than display

[Philip Greenspun]

Slide credit: S. Marschner

Instagram Filters

- How do they make those Instagram filters?

"It's really a combination of a bunch of different methods. In some cases we draw on top of images, in others we do pixel math. It really depends on the effect we're going for." --- Kevin Systrom, co-founder of Instagram

Source: C. Dyer

Example Instagram Steps

1. Perform an independent RGB color point transformation on the original image to increase contrast or make a color cast

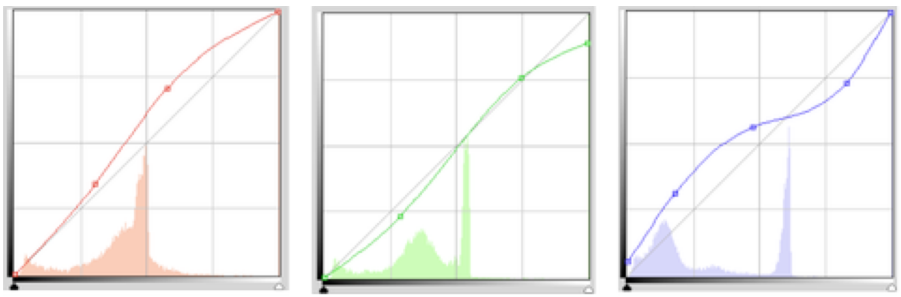

R G B

Source: C. Dyer

Example Instagram Steps

2. Overlay a circle background image to create a vignette effect

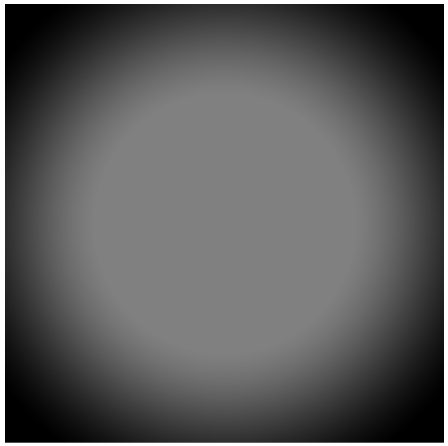

Source: C. Dyer

Example Instagram Steps

3. Overlay a background image as decorative grain

Source: C. Dyer

Example Instagram Steps

4. Add a border or frame

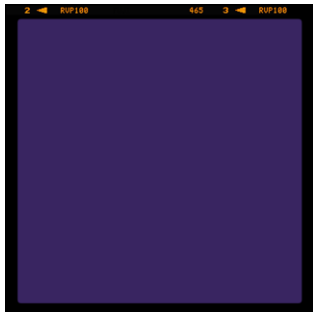

Source: C. Dyer

Result

Javascript library for creating Instagram-like effects, see: http://alexmic.net/filtrr/

Source: C. Dyer

Today

- Perception of color and light
- Color spaces
- Point operations
- Histogram processing

Histogram

- Histogram: a discrete function *h(r)* which counts the number of pixels in the image having intensity *r*
- If *h(r)* is normalized, it measures the probability of occurrence of intensity level *r* in an image

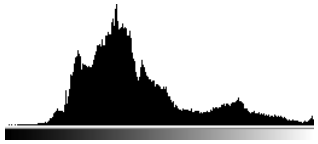

- What histograms say about images? A descriptor for visual information
- What they don't?
  - No spatial information

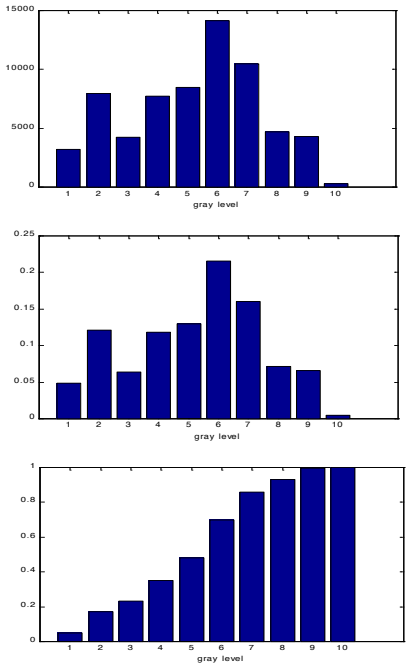

Histogram

Normalized Histogram

Cumulative Histogram

Slide credit: Y. Hel-Or

Images and histograms

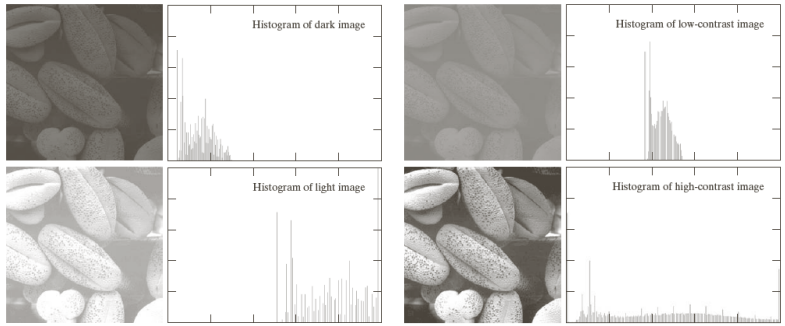

- How do histograms change when
  - we adjust brightnesss? shifts the histogram horizontally
  - we adjust constrast? stretches or shrinks the histogram horizontally

Image Representations: Histograms

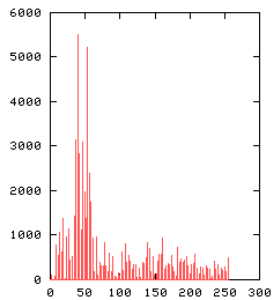

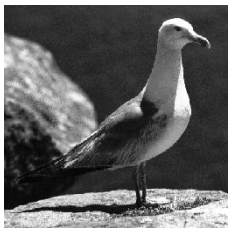

Global histogram

- Represent distribution of features
  - Color, texture, depth, …

Image credit: D. Kauchak

Image Representations: Histograms

Histogram: Probability or count of data in each bin

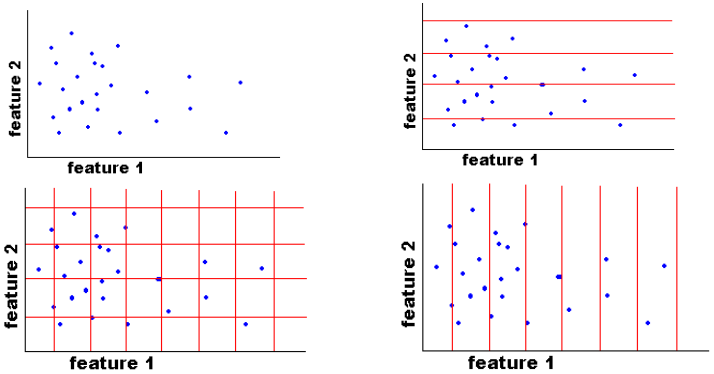

Joint histogram

- Requires lots of data
- Loss of resolution to avoid empty bins

Marginal histogram

- Requires independent features
- More data/bin than joint histogram

Image credit: D. Kauchak

Histograms: Implementation issues

- Quantization
  - Grids: fast but applicable only with few dimensions

Few Bins
Need less data
Coarser representation

Many Bins
Need more data
Finer representation

- Matching
  - Histogram intersection or Euclidean may be faster
  - Chi-squared often works better
  - Earth mover's distance is good for when nearby bins represent similar values

Slide credit: J. Hays

What kind of things do we compute histograms of?

- Color

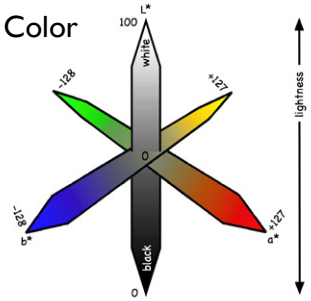

L*a*b* color space

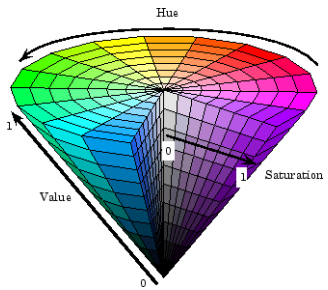

HSV color space

- Texture (filter banks over regions – later on)

Slide credit: J. Hays

What kind of things do we compute histograms of?

- Histograms of oriented gradients (later on)

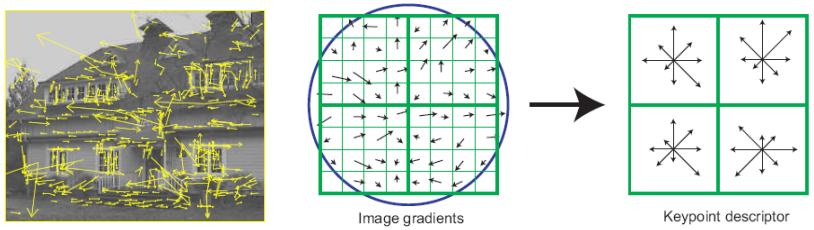

SIFT – Lowe IJCV 2004

Slide credit: J. Hays

Examples

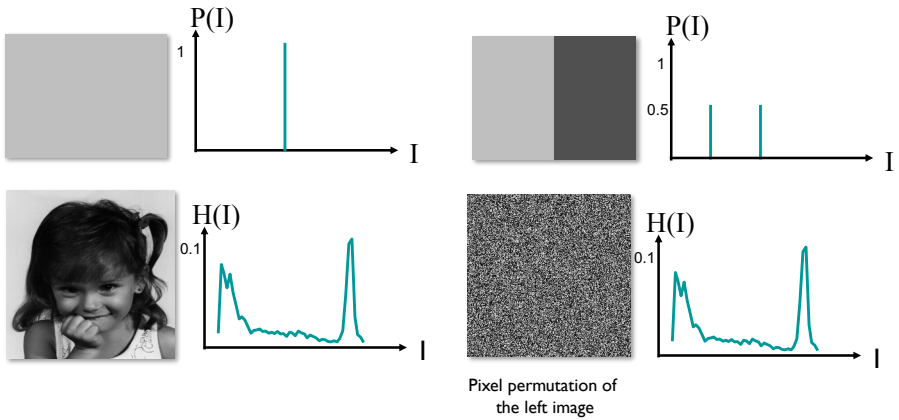

- The image histogram does not fully represent the image

Slide credit: Y. Hel-Or

Examples

Original image

Decreasing contrast

Increasing average

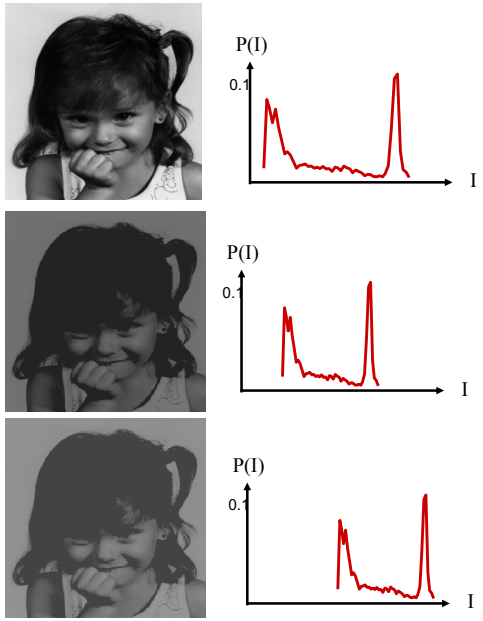

Slide credit: Y. Hel-Or

Image Statistics

- The image mean: $E\{I\} = \frac{1}{N}\sum_{i,j} I(i,j) = \frac{1}{N}\sum_k k\,H(k) = \sum_k k\,P(k)$
- Generally: $E\{g(k)\} = \sum_k g(k)P(k)$
- The image s.t.d. : $\sigma(I) = \sqrt{E\{(I - E\{I\})^2\}} = \sqrt{E(I^2) - E^2(I)}$

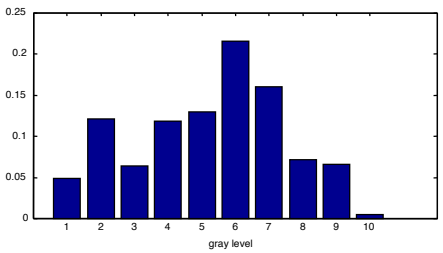

$where \quad E\{I^2\} = \sum_k k^2 P(k)$

Slide credit: Y. Hel-Or

Image Entropy

$$Entropy(I) = -\sum_k P(k)\log P(k)$$

- The image entropy specifies the uncertainty in the image values.
- Measures the averaged amount of information required to encode the image values.

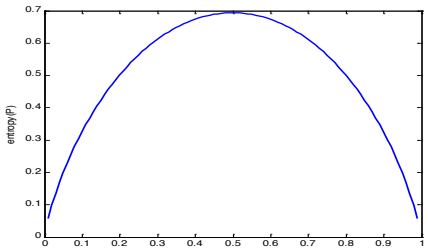

Entropy of a 2 values variable

Slide credit: Y. Hel-Or

Image Entropy

- An infrequent event provides more information than a frequent event
- Entropy is a measure of histogram dispersion

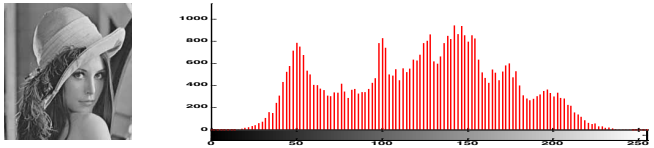

entropy=7.4635

entropy=0

Slide credit: Y. Hel-Or

Adaptive Histogram

- In many cases histograms are needed for local areas in an image
- Examples:
  - Pattern detection
  - adaptive enhancement
  - adaptive thresholding
  - tracking

Slide credit: Y. Hel-Or

Histogram Usage

- Digitizing parameters
- Measuring image properties:
  - Average
  - Variance
  - Entropy
  - Contrast
  - Area (for a given gray-level range)
- Threshold selection
- Image distance
- Image Enhancement
  - Histogram equalization
  - Histogram stretching
  - Histogram matching

Slide credit: Y. Hel-Or

Example: Auto-Focus

- In some optical equipment (e.g. slide projectors) inappropriate lens position creates a blurred ("out-of-focus") image
- We would like to automatically adjust the lens
- How can we measure the amount of blurring?

Slide credit: Y. Hel-Or

Example: Auto-Focus

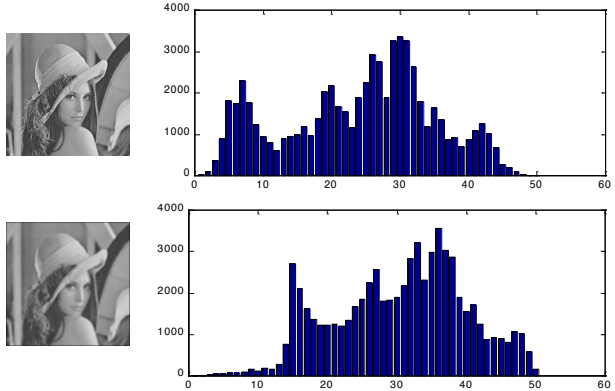

- Image mean is not affected by blurring
- Image s.t.d. (entropy) is decreased by blurring
- Algorithm: Adjust lens according the changes in the histogram s.t.d.

Slide credit: Y. Hel-Or

Recall: Thresholding

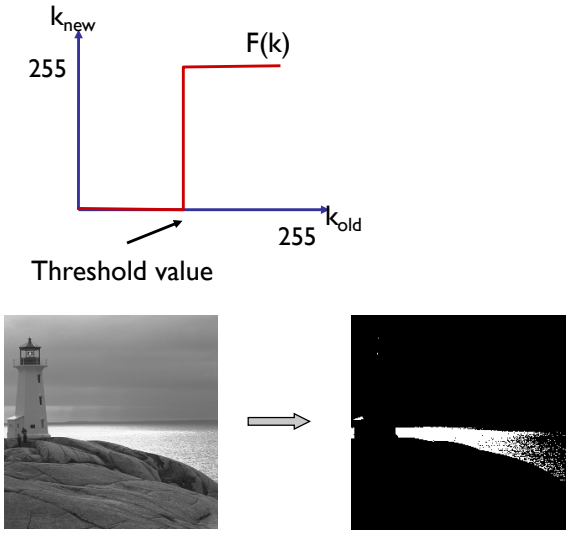

Slide credit: Y. Hel-Or

Threshold Selection

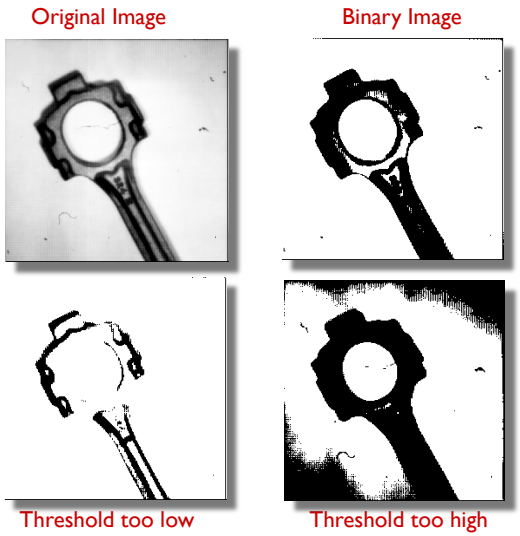

Slide credit: Y. Hel-Or

Segmentation using Thresholding

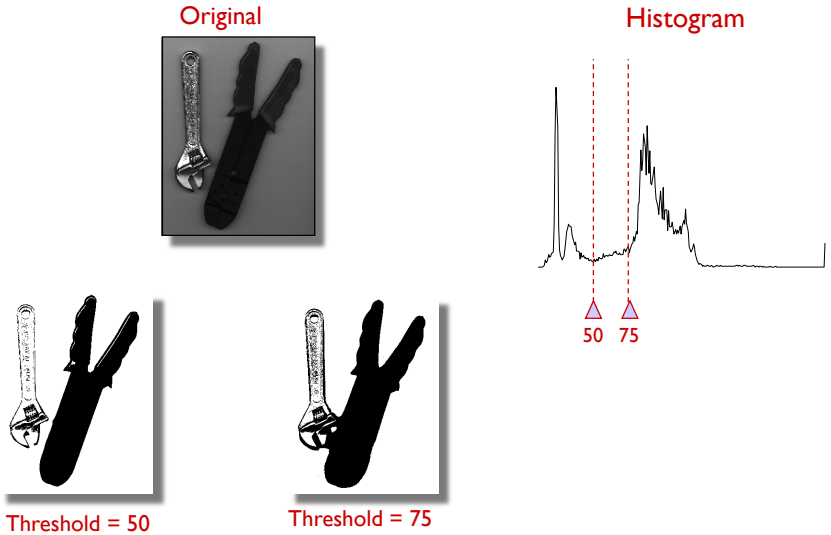

Slide credit: Y. Hel-Or

Segmentation using Thresholding

Original

Histogram

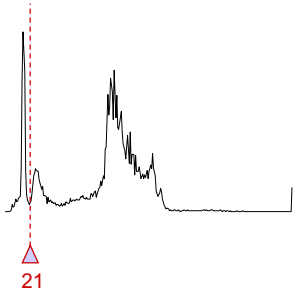

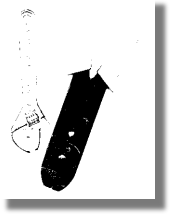

Threshold = 21

Slide credit: Y. Hel-Or

Adaptive Thresholding

- Thresholding is space variant.
- How can we choose the the local threshold values?

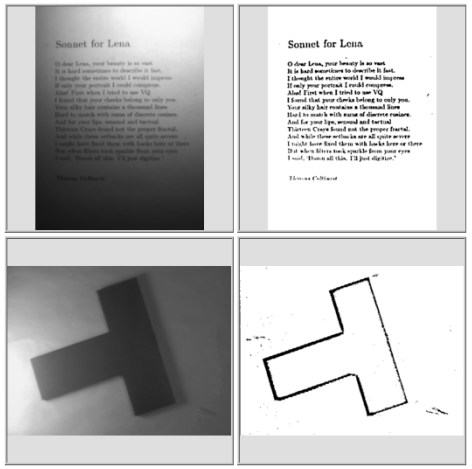

Slide credit: Y. Hel-Or

Histogram based image distance

- **Problem**: Given two images A and B whose (normalized) histogram are P_A and P_B define the distance D(A,B) between the images.

- Example Usage:
  - Tracking
  - Image retrieval
  - Registration
  - Detection
  - Many more ...

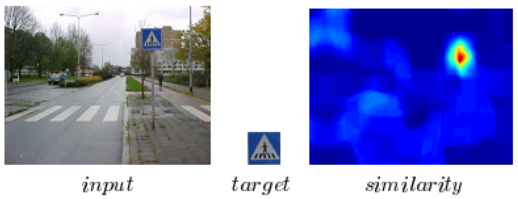

Porikli 05

Slide credit: Y. Hel-Or

Option 1: Minkowski Distance

$$D_p(A,B) = \left[\sum_k \left| P_A(k) - P_B(k) \right|^p \right]^{1/p}$$

- **Problem**: distance may not reflect the perceived dissimilarity:

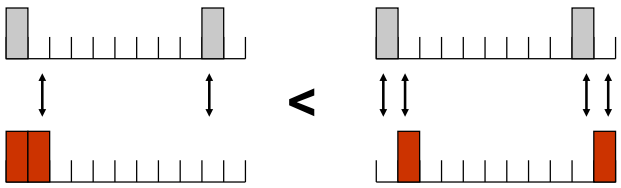

Slide credit: Y. Hel-Or

Option 2: Kullback-Leibler (KL) Distance

$$D_{KL}(A \| B) = -\sum_k P_A(k) \log \frac{P_A(k)}{P_B(k)}$$

- Measures the amount of added information needed to encode image A based on the histogram of image B.
- Non-symmetric: $D_{KL}(A,B) \neq D_{KL}(B,A)$
- Suffers from the same drawback of the Minkowski distance.

Slide credit: Y. Hel-Or

Option 3: The Earth Mover Distance (EMD)

- Suggested by Rubner & Tomasi 98
- Defines as the minimum amount of "work" needed to transform histogram H_A towards H_B
- The term d_{ij} defines the "ground distance" between gray-levels i and j.
- The term $F=\{f_{ij}\}$ is an admissible flow from $H_A(i)$ to $H_B(j)$

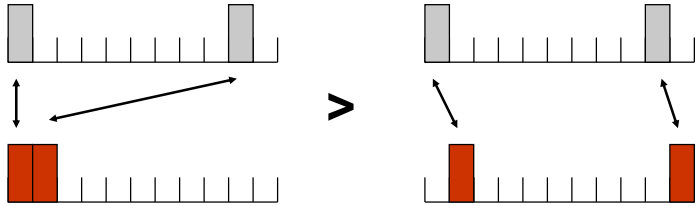

Slide credit: Y. Hel-Or

Option 3: The Earth Mover Distance (EMD)

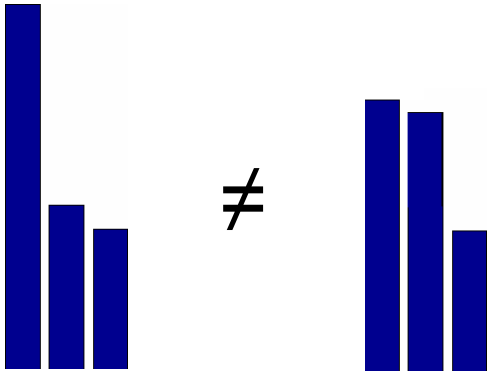

Slide credit: P. Barnum

Option 3: The Earth Mover Distance (EMD)

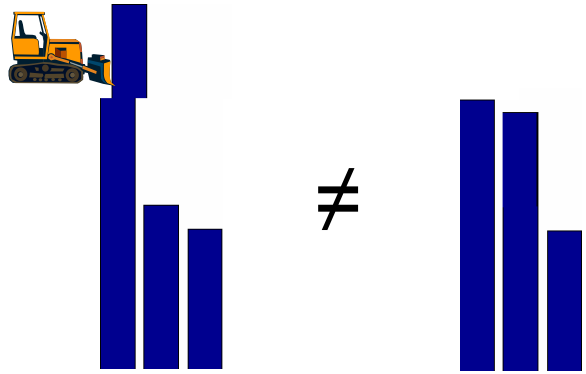

From: Pete Barnum

Slide credit: P. Barnum

Option 3: The Earth Mover Distance (EMD)

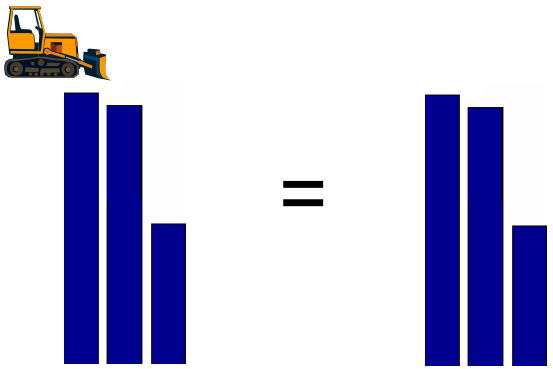

Slide credit: P. Barnum

Option 3: The Earth Mover Distance (EMD)

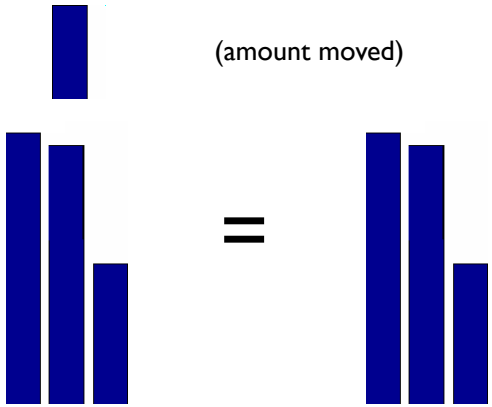

Slide credit: P. Barnum

Option 3: The Earth Mover Distance (EMD)

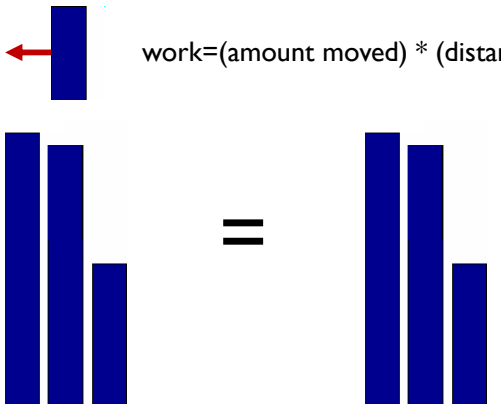

Slide credit: P. Barnum

Option 3: The Earth Mover Distance (EMD)

$$D_{EMD}(A,B) = \min_F \sum_i \sum_j f_{ij} \cdot d_{ij}$$

$$s.t. \quad f_{ij} \geq 0 \; ; \quad P_B(k) = \sum_i f_{ik} \quad ; \quad P_A(k) \geq \sum_i f_{ki}$$

- Constraints:
  - Move earth only from A to B
  - After move P_A will be equal to P_B
  - Cannot send more "earth" than there is
- Can be solved using Linear Programming
- Can be applied in high dim. histograms (color).

Slide credit: Y. Hel-Or

Special case: EMD in 1D

- Define C_A and C_B as the cumulative histograms of image A and B respectively:

$$D_{EMD}(A,B) = \sum_k \left| C_A(k) - C_B(k) \right|$$

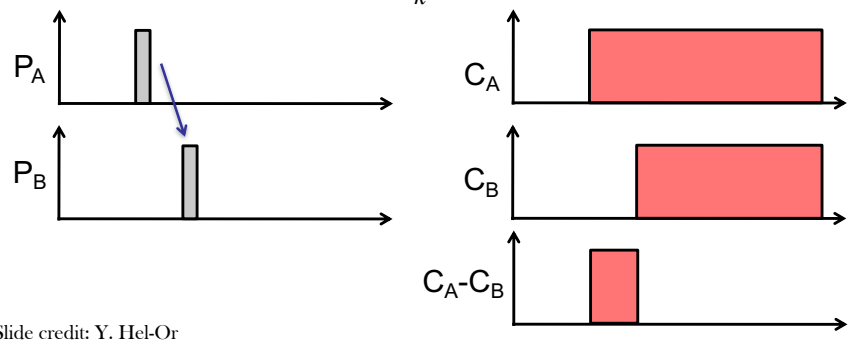

Slide credit: Y. Hel-Or

Histogram equalization

- A good quality image has a nearly uniform distribution of intensity levels. Why?
- Every intensity level is equally likely to occur in an image
- *Histogram equalization:* Transform an image so that it has a uniform distribution
  - create a lookup table defining the transformation

Histogram as a probability density function

- Recall that a normalized histogram measures the probability of occurrence of an intensity level *r* in an image
- We can normalize a histogram by dividing the intensity counts by the area

$$p(r) = \frac{h(r)}{Area}$$

Histogram equalization: Continuous domain

- Define a transformation function of the form

$$s = T(r) = (L-1) \underbrace{\int_0^r p(w)\,dw}_{\text{cumulative distribution function}}$$

where
- *r* is the input intensity level
- *s* is the output intensity level
- *p* is the normalized histogram of the input signal
- *L* is the desired number of intensity levels

(Continuous) output signal has a uniform distribution!

Histogram equalization: Discrete domain

- Define the following transformation function for an MxN image

$$s_k = T(r_k) = (L-1)\sum_{j=0}^{k}\frac{n_j}{MN} = \frac{(L-1)}{MN}\sum_{j=0}^{k} n_j$$

for $k = 0,\ldots,L-1$

where

- r_k is the input intensity level
- s_k is the output intensity level
- n_j is the number of pixels having intensity value j in the input image
- L is the number of intensity levels

(Discrete) output signal has a nearly uniform distribution!

Histogram equalization

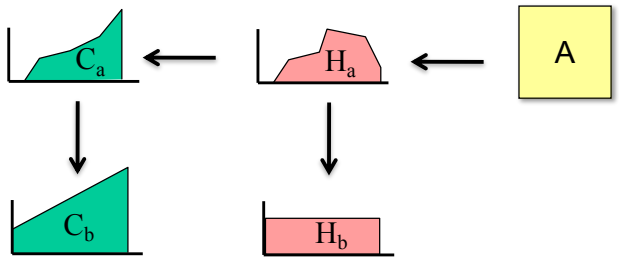

- Define: $C_b(v) = v * \frac{(\#\, pixels)}{\#\, grayValues}$
- Assign: $v_b = C_b^{-1}(C_a(v_a)) = M(v_a)$

Slide credit: Y. Hel-Or

Histogram equalization

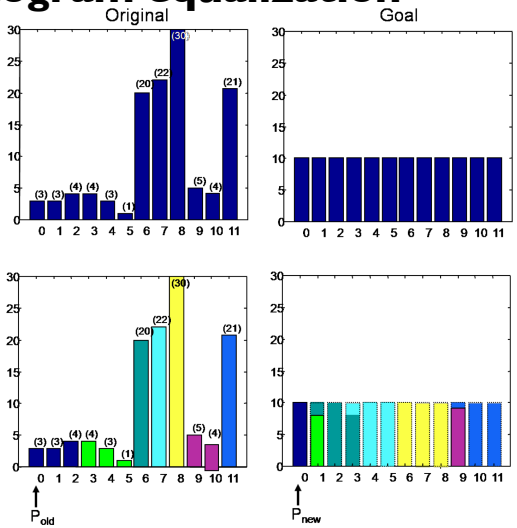

0 1 2 3 4 5 6 7 8 9 10 11 old
0 0 0 1 1 1 3 5 8 9 9 11 new

Slide credit: Y. Hel-Or

Histogram equalization

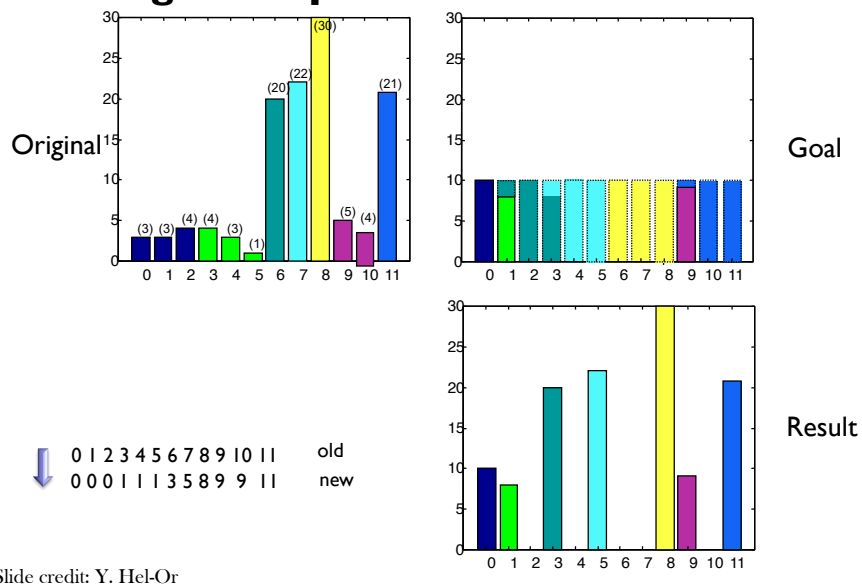

0 1 2 3 4 5 6 7 8 9 10 11 old
0 0 0 1 1 1 3 5 8 9 9 11 new

Slide credit: Y. Hel-Or

Histogram equalization examples

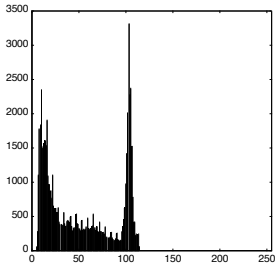

Original

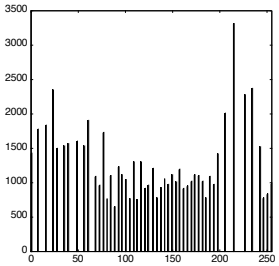

Equalized

Slide credit: Y. Hel-Or

Histogram equalization examples

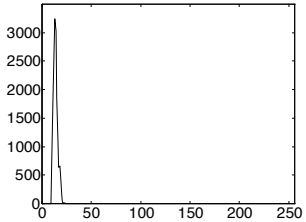

Original

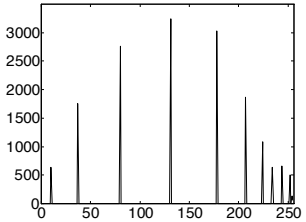

Equalized

Slide credit: Y. Hel-Or

Histogram equalization examples

Original

Equalized

Slide credit: Y. Hel-Or

Histogram equalization examples

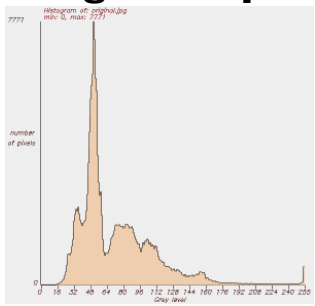

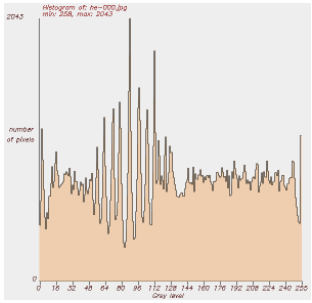

Slide credit: C. Dyer

Histogram equalization examples

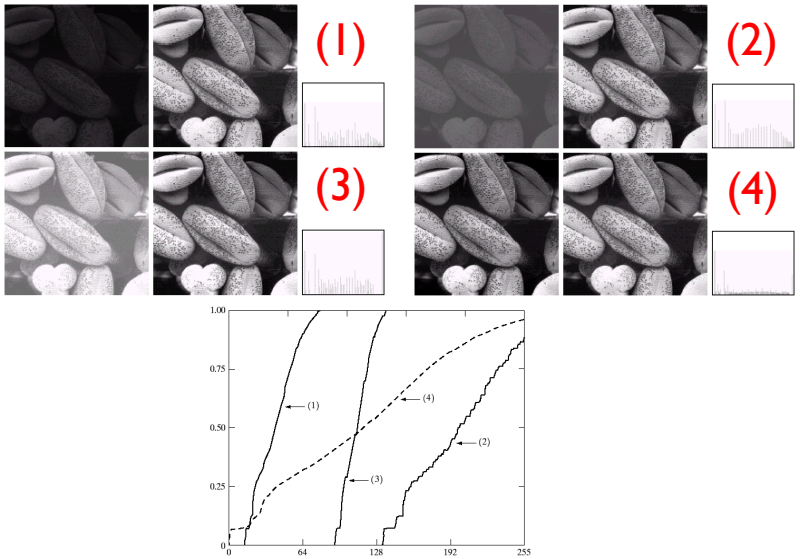

Histogram Specification

- Given an input image f and a specific histogram $p_2(r)$, transform the image so that it has the specified histogram

source target mapped source

- How to perform histogram specification?
- Histogram equalization produces a (nearly) uniform output histogram
- Use histogram equalization as an intermediate step

Image credit: Y. Hel-Or

Histogram Specification

1. Equalize the histogram of the input image

$$T_1(r) = (L-1)\int_0^r p_1(w)\,dw$$

2. Histogram equalize the desired output histogram

$$T_2(r) = (L-1)\int_0^r p_2(w)\,dw$$

3. Histogram specification can be carried out by the following point operation:

$$s = T(r) = T_2^{-1}(T_1(r))$$

Histogram Specification

- In cases where corresponding colors between images are not "consistent", this mapping may fail:

Images from: S. Kagarlitsky, M.Sc. thesis 2010.

Slide credit: Y. Hel-Or

Histogram Specification: Discussion

- Histogram matching produces the optimal **monotonic** mapping so that the resulting histogram will be as **close** as possible to the target histogram.
- This does not necessarily imply similar images.

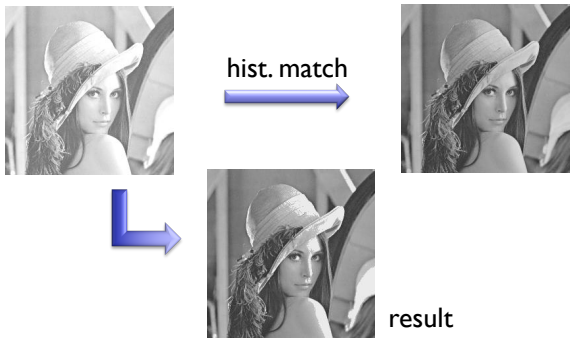

Slide credit: Y. Hel-Or

Next week

- Spatial filtering