

BBM 413

Fundamentals of

Image Processing

Erkut Erdem
Dept. of Computer Engineering
Hacettepe University

Image Formation and Color

Acknowledgement: The course slides are mostly adapted from the slides prepared by Steve Marschner, James Hays, Ali Farhadi and Anat Levin

Today

- Image formation
- Digital images
- Perception of color and light
- Color spaces

Today

- **Image formation**
- Digital images
- Perception of color and light
- Color spaces


An image is:

- A 2D distribution of intensity or color
- A function defined on a two-dimensional plane

$$I : \mathbb{R}^2 \rightarrow \dots$$

- Note: no mention of pixels yet
- To process images, must:
 - obtain images—capture the scenes via hardware
 - represent images—encode them numerically

Image Formation




Three Dimensional
World



Two Dimensional
Image Space

- What is measured in an image location?
 - brightness
 - color
- viewpoint
- << illumination conditions
- local geometry
- local material properties

Image Formation



Three Dimensional World \longrightarrow Two Dimensional Image Space

- What is measured in an image location?
 - brightness
 - color
- viewpoint
 << illumination conditions
 surface properties
 (local geometry and local material properties)⁶

Slide source: Seitz




CoolOpticalIllusions.com



CoolOpticalIllusions.com


A photon's life choices

- Absorption
- Diffusion
- Reflection
- Transparency
- Refraction
- Fluorescence
- Subsurface scattering
- Phosphorescence
- Interreflection




A photon's life choices

- **Absorption**
- Diffusion
- Reflection
- Transparency
- Refraction
- Fluorescence
- Subsurface scattering
- Phosphorescence
- Interreflection




A photon's life choices

- Absorption
- **Diffuse Reflection**
- Reflection
- Transparency
- Refraction
- Fluorescence
- Subsurface scattering
- Phosphorescence
- Interreflection




A photon's life choices

- Absorption
- Diffusion
- **Specular Reflection**
- Transparency
- Refraction
- Fluorescence
- Subsurface scattering
- Phosphorescence
- Interreflection




A photon's life choices

- Absorption
- Diffusion
- Reflection
- **Transparency**
- Refraction
- Fluorescence
- Subsurface scattering
- Phosphorescence
- Interreflection




A photon's life choices

- Absorption
- Diffusion
- Reflection
- Transparency
- **Refraction**
- Fluorescence
- Subsurface scattering
- Phosphorescence
- Interreflection




A photon's life choices

- Absorption
- Diffusion
- Reflection
- Transparency
- Refraction
- **Fluorescence**
- Subsurface scattering
- Phosphorescence
- Interreflection




A photon's life choices

- Absorption
- Diffusion
- Reflection
- Transparency
- Refraction
- Fluorescence
- **Subsurface scattering**
- Phosphorescence
- Interreflection



A photon's life choices

- Absorption
- Diffusion
- Reflection
- Transparency
- Refraction
- Fluorescence
- Subsurface scattering
- **Phosphorescence**
- Interreflection



A photon's life choices

- Absorption
- Diffusion
- Reflection
- Transparency
- Refraction
- Fluorescence
- Subsurface scattering
- Phosphorescence
- **Interreflection**





Image Formation

Images cannot exist without light!


Why is there no image on a white piece of paper?

It receives light from all directions




From Photography, London et al.

Image Formation



- Let's design a camera
 - Idea I: put a piece of film in front of an object
 - Do we get a reasonable image?


Pinhole camera



A pinhole projects all rays through
a common center of projection.


From Photography, London et al.

Pinhole camera




- Add a barrier to block off most of the rays
 - This reduces blurring
 - The opening is known as the **aperture**
 - How does this transform the image?

Camera Obscura



- Basic principle known to Mozi (470-390 BC), Aristotle (384-322 BC)
- Drawing aid for artists: described by Leonardo da Vinci (1452-1519)


Camera Obscura



Abelardo Morell:
Through the Looking Glass

<https://www.youtube.com/watch?v=DTa9wLkaizQ> 25


Camera Obscura



- The first camera
 - How does the aperture size affect the image?



Pinhole Size?

Photograph made with small pinhole




Small pinhole-
sharp but hard to
collect enough light

Photograph made with larger pinhole



Larger pinhole-
Blur


Pinhole Size



small hole => sharp, but doesn't collect enough light (noise)

larger hole => easy to collect enough light, but blur occurs

Pinhole Size



2 mm



1 mm







0.6mm



0.35 mm

- Why not make the aperture as small as possible?
 - Less light gets through
 - Diffraction effects...

Solution: light refraction!




Lenses

- gather more light!
- But need to be focused

Photograph made with small pinhole



To make this picture, the lens of a camera was replaced with a thin metal disk pierced by a tiny pinhole, equivalent in size to an aperture of f/182. Only a few rays of light from each point on the




subject got through the tiny opening, producing a soft but acceptably clear photograph. Because of the small size of the pinhole, the exposure had to be 6 sec long.

Photograph made with lens




This time, using a simple convex lens with an f/16 aperture, the scene appeared sharper than the one taken with the smaller pinhole, and the exposure time was much shorter, only 1/100 sec.




The lens opening was much bigger than the pinhole, letting in far more light, but it focused the rays from each point on the subject precisely so that they were sharp on the film.

Adding a lens




- A lens focuses light onto the film
 - There is a specific distance at which objects are “in focus”
 - other points project to a “circle of confusion” in the image
 - Changing the shape of the lens changes this distance

Lenses



- A lens focuses parallel rays onto a single focal point
 - focal point at a distance f beyond the plane of the lens
 - f is a function of the shape and index of refraction of the lens
 - Aperture of diameter D restricts the range of rays
 - aperture may be on either side of the lens
 - Lenses are typically spherical (easier to produce)
 - Real cameras use many lenses together (to correct for aberrations)

Thin lenses



- Thin lens equation: $\frac{1}{d_o} + \frac{1}{d_i} = \frac{1}{f}$
 - Any object point satisfying this equation is in focus
 - What is the shape of the focus region?
 - How can we change the focus region?


A lens is focused at a single depth

$$\frac{1}{z_o} + \frac{1}{z_i} = \frac{1}{f}$$


z_o : distance to the (focused) object

z_i : distance behind the lens at which the image is formed

f : focal length



A lens is focused at a single depth



Rays emerge from multiple object points
(circle of confusion) => the captured image is blurred

A lens is focused at a single depth




$$\frac{1}{z_o} + \frac{1}{z_i} = \frac{1}{f}$$



Thin lens assumption

The thin lens assumption assumes the lens has no thickness, but this isn't true...




By adding more elements to the lens, the distance at which a scene is in focus can be made roughly planar.

Projection


- Mapping from the world (3d) to an image (2d)
 - Can we have a 1-to-1 mapping?
 - How many possible mappings are there?
-
- An optical system defines a particular projection.
 - Two examples:
 1. Perspective projection (how we see “normally”)
 2. Orthographic projection (e.g., telephoto lenses)

Modeling projection



- The coordinate system
 - We will use the pin-hole model as an approximation
 - Put the optical center (Center Of Projection) at the origin
 - Put the image plane (Projection Plane) in front of the COP
 - The camera looks down the negative z axis
 - we need this if we want right-handed-coordinates

Modeling projection



- Projection equations
 - Compute intersection with PP of ray from (x, y, z) to COP
 - Derived using similar triangles

$$(x, y, z) \rightarrow \left(-d \frac{x}{z}, -d \frac{y}{z}, -d \right)$$

- We get the projection by throwing out the last coordinate:

$$(x, y, z) \rightarrow \left(-d \frac{x}{z}, -d \frac{y}{z} \right)$$

Homogeneous coordinates

- Is this a linear transformation?
 - no—division by z is nonlinear

Trick: add one more coordinate:

$$(x, y) \Rightarrow \begin{bmatrix} x \\ y \\ 1 \end{bmatrix}$$

homogeneous image
coordinates

$$(x, y, z) \Rightarrow \begin{bmatrix} x \\ y \\ z \\ 1 \end{bmatrix}$$

homogeneous scene
coordinates

Converting *from* homogeneous coordinates

$$\begin{bmatrix} x \\ y \\ w \end{bmatrix} \Rightarrow (x/w, y/w)$$

$$\begin{bmatrix} x \\ y \\ z \\ w \end{bmatrix} \Rightarrow (x/w, y/w, z/w)$$

Perspective Projection

- Projection is a matrix multiply using homogeneous coordinates:

$$\begin{bmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & -1/d & 0 \end{bmatrix} \begin{bmatrix} x \\ y \\ z \\ 1 \end{bmatrix} = \begin{bmatrix} x \\ y \\ -z/d \end{bmatrix} \Rightarrow \left(-d\frac{x}{z}, -d\frac{y}{z} \right)$$

divide by third coordinate

This is known as **perspective projection**

- The matrix is the **projection matrix**

Perspective Projection Example

I. Object point at (10, 6, 4), d=2

$$\begin{bmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & -1/d & 0 \end{bmatrix} \begin{bmatrix} x \\ y \\ z \\ 1 \end{bmatrix} = \begin{bmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & -1/2 & 0 \end{bmatrix} \begin{bmatrix} 10 \\ 6 \\ 4 \\ 1 \end{bmatrix} = \begin{bmatrix} 10 \\ 6 \\ -2 \end{bmatrix}$$
$$\Rightarrow x' = -5, y' = -3$$

2. Object point at (25, 15, 10)


$$\begin{bmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & -1/d & 0 \end{bmatrix} \begin{bmatrix} x \\ y \\ z \\ 1 \end{bmatrix} = \begin{bmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & -1/2 & 0 \end{bmatrix} \begin{bmatrix} 25 \\ 15 \\ 10 \\ 1 \end{bmatrix} = \begin{bmatrix} 25 \\ 15 \\ -5 \end{bmatrix}$$
$$\Rightarrow x' = -5, y' = -3$$

Perspective projection is not 1-to-1!

Perspective Projection



- preserves lines (collinearity), cross ratio
- does not always preserve parallel lines.
- Lines parallel to projection plane remain parallel.
- Lines not parallel to projection plane converge to a single point on the horizon called the vanishing point.




Perspective Projection

- What happens when $d \rightarrow \infty$?

$$\begin{bmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & -1/d & 0 \end{bmatrix} \begin{bmatrix} x \\ y \\ z \\ 1 \end{bmatrix} = \begin{bmatrix} x \\ y \\ -z/d \end{bmatrix} \Rightarrow \left(-d\frac{x}{z}, -d\frac{y}{z} \right)$$



Orthographic projection

- Special case of perspective projection
 - Distance from the COP to the PP is infinite



- Good approximation for telephoto optics
- Also called “parallel projection”:
$$(x, y, z) \rightarrow (x, y)$$

Orthographic projection




- preserves ratios, but not angles.
- parallel lines remain parallel.
- loses depth information.

Orthographic (“telecentric”) lenses





Navitar telecentric zoom lens



<http://www.lhup.edu/~dsimanek/3d/telecent.htm>


Depth of field



- Changing the aperture size affects depth of field
 - A smaller aperture increases the range in which the object is approximately in focus

Aperture


- Diameter of the lens opening (controlled by diaphragm)
- Controls depth of field
- Expressed as a fraction of focal length, in f-number
 - f/2.0 on a 50mm means that the aperture is 25mm
 - f/2.0 on a 100mm means that the aperture is 50mm
- Disconcerting: small f number = big aperture
- What happens to the area of the aperture when going from f/2.0 to f/4.0?
- Typical f numbers are f/2.0, f/2.8, f/4, f/5.6, f/8, f/11, f/16, f/22, f/32




Main effect of aperture

- **Depth of field:** Allowable depth variation in the scene that limits the circle of confusion to a tolerable number

Large aperture opening



Small aperture opening



Depth of field




Image of object in focus- sharp (all rays hitting a single sensor point emerge from a single point on the object)

Depth of field






Image of object in focus- sharp (all rays hitting a single sensor point emerge from a single point on the object)


Image of an object away from focus depth- blurred (rays hitting a single sensor point emerge from multiple points on the object)

Depth of field

- We allow for some tolerance



Depth of Field



Depth of Field

Shallow Depth
of Field

Portrait



Large Depth
of Field




Landscape



Large
Aperture



Small
Aperture




<http://photographertips.net>

Exposure

- Exposure: How much light falls on sensor
- Get the right amount of light to sensor/film
- Main parameters:
 - Shutter speed: How long sensor is exposed to light
 - Aperture (area of lens): How much light can pass through from the lens
 - Sensitivity: How much light is needed by the sensor
 - Lighting conditions

Field of View (Zoom, focal length)




From London and Upton

Shutter speed

- Controls how long the film/sensor is exposed, i.e. the amount of light reaching the sensor
- Pretty much linear effect on exposure
- Usually in fraction of a second:
 - 1/30, 1/60, 1/125, 1/250, 1/500
 - Get the pattern ?
- Faster shutter (e.g. 1/500th sec) = less light
- Slower shutter (e.g. 1/30th sec) = more light
- On a normal lens, normal humans can hand-hold down to 1/60
 - In general, the rule of thumb says that the limit is the inverse of focal length, e.g. 1/500 for a 500mm

Shutter speed



Short exposure- dark



medium exposure



long exposure- saturation

Shutter speed



Short exposure after
contrast adjustment-
noise



medium exposure



long exposure- saturation


Main effect of slower shutter speed

- For dynamic scenes, the shutter speed also determines the amount of *motion blur* in the resulting picture.
- Camera shake

Image taken with a tripod




Image taken with a hand held camera



Main effect of slower shutter speed

- For dynamic scenes, the shutter speed also determines the amount of *motion blur* in the resulting picture.
- Scene motion



From Photography, London et al.


Effect of Shutter Speed

- Freezing motion


Walking people



Running people



Car



Fast train



1/125

1/250

1/500

1/1000

Today

- Image formation
- **Digital images**
- Perception of color and light
- Color spaces

Digital camera





- A digital camera replaces film with a sensor array
 - Each cell in the array is light-sensitive diode that converts photons to electrons
 - Two common types
 - Charge Coupled Device (CCD)
 - CMOS
 - <http://electronics.howstuffworks.com/digital-camera.htm>

Digital camera


- Color typically captured using color mosaic
- Demosaicing

Mosaic Capture



[Foveon]

Digital camera




How do digital cameras work?

<https://www.youtube.com/watch?v=lc0czeUjrGE>


69

Sensor Array



a b

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



CMOS sensor

Issues with digital cameras

- Noise
 - big difference between consumer vs. SLR-style cameras
 - low light is where you most notice noise
- Compression
 - creates artifacts except in uncompressed formats (tiff, raw)
- Color
 - color fringing artifacts from Bayer patterns
- Blooming
 - charge overflowing into neighboring pixels
- In-camera processing
 - oversharpening can produce halos
- Interlaced vs. progressive scan video
 - even/odd rows from different exposures
- Are more megapixels better?
 - requires higher quality lens
 - noise issues
- Stabilization
 - compensate for camera shake (mechanical vs. electronic)

More info online, e.g.,

<http://electronics.howstuffworks.com/digital-camera.htm>

<http://www.dpreview.com/>

Sampling and Quantization







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.

Image Representation



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

Image Representation

- **Digital image:** 2D discrete function f
- **Pixel:** Smallest element of an image $f(x,y)$




Figure: M. J. Black

Image Representation

- **Digital image:** 2D discrete function f
- **Pixel:** Smallest element of an image $f(x,y)$




Figure: M. J. Black

Datatypes for raster images

- Bitmaps: boolean per pixel (1 bpp): $I : \mathbb{R}^2 \rightarrow \{0, 1\}$
 - interp. = black and white; e.g. fax
- Grayscale: integer per pixel: $I : \mathbb{R}^2 \rightarrow [0, 1]$
 - interp. = shades of gray; e.g. black-and-white print
 - precision: usually byte (8 bpp); sometimes 10, 12, or 16 bpp
- Color: 3 integers per pixel: $I : \mathbb{R}^2 \rightarrow [0, 1]^3$
 - interp. = full range of displayable color; e.g. color print
 - precision: usually byte[3] (24 bpp)
 - sometimes 16 (5+6+5) or 30 or 36 or 48 bpp
 - indexed color: a fading idea

Datatypes for raster images


- Floating point: $I : \mathbb{R}^2 \rightarrow \mathbb{R}_+$ or $I : \mathbb{R}^2 \rightarrow \mathbb{R}_+^3$
 - more abstract, because no output device has infinite range
 - provides *high dynamic range* (HDR)
 - represent real scenes independent of display
 - becoming the standard intermediate format in graphics processors
- Clipping and white point
 - common to compute FP, then convert to integer
 - full range of values may not “fit” in display’s output range
 - simplest solution: choose a maximum value, scale so that value becomes full intensity ($2^n - 1$ in an n -bit integer image)

Intensity encoding in images

- What do the numbers in images (pixel values) mean?
 - they determine how bright that pixel is
 - bigger numbers are (usually) brighter

Datatypes for raster images

- For color or grayscale, sometimes add *alpha* channel
 - describes transparency of images



[Adobe Photoshop sample]

Storage requirements for images

- 1024x1024 image (1 megapixel)
 - bitmap: 128KB
 - grayscale 8bpp: 1MB
 - grayscale 16bpp: 2MB
 - color 24bpp: 3MB
 - floating-point HDR color: 12MB

Converting pixel formats

- Color to gray
 - could take one channel (blue, say)
 - leads to odd choices of gray value
 - combination of channels is better
 - but different colors contribute differently to lightness
 - which is lighter, full blue or full green?
 - good choice: $\text{gray} = 0.2 \text{ R} + 0.7 \text{ G} + 0.1 \text{ B}$
 - more on this in color, later on



COLOR




BLUE ONLY

Same pixel values.



Same luminance?



GRAY

Converting pixel precision

- Up is easy; down loses information—be careful



1 bpp (2 grays)

[photo: Philip Greenspun]

Today

- Image formation
- Digital images
- Perception of color and light
- Color spaces

Why does a visual system need color?




Why does a visual system need color? (an incomplete list...)

- To tell what food is edible.
- To distinguish material changes from shading changes.
- To group parts of one object together in a scene.
- To find people's skin.
- Check whether a person's appearance looks normal/healthy.

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?



#thedress


- Let's take averages



two pieces
of the dress



averages




basic pattern




#thedress

- The dress in the photograph



#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



#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.


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.

Image Brightness (Intensity)



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

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

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

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.)




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

Color


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



Color



Newton 1665




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.




The Physics of light

Some examples of the spectra of light sources


A. Ruby Laser




B. Gallium Phosphide Crystal



C. Tungsten Lightbulb



D. Normal Daylight



The Physics of light

Some examples of the reflectance spectra of surfaces








Image formation

- What determines the brightness of an image pixel?






Color mixing

Cartoon spectra for color names:




Credit: W. Freeman

Additive color mixing




Colors combine by
adding color spectra



Light *adds* to black.

Examples of additive color systems






CRT phosphors




multiple projectors

Subtractive color mixing



Colors combine by
multiplying color
spectra.



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

Examples of subtractive color systems



- Printing on paper
- Crayons
- Photographic film




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]


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


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!**

Layers of the retina



Slide credit: S. Ullman

Layers of the retina



Slide credit: J. Hays

Wait, the blood vessels are in front of the photoreceptors??




Why we have blind spots - and how to see the blood vessels inside your own eye! https://www.youtube.com/watch?v=L_W-IXqoxHA

Slide credit: J. Hays

Eye Movements

- Saccades
 - Can be consciously controlled. Related to perceptual attention.
 - 200ms to initiation, 20 to 200ms to carry out. Large amplitude.
- Microsaccades
 - Involuntary. Smaller amplitude. Especially evident during prolonged fixation. Function debated.
- Ocular microtremor (OMT)
 - involuntary. high frequency (up to 80Hz), small amplitude.
- Smooth pursuit – tracking an object

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


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 = $10 \mu\text{m}$), 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.

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

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)


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

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

Today

- Image formation
- Digital images
- Perception of color and light
- **Color spaces**

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

Color Images: Bayer Grid



- Estimate RGB at 'G' cells from neighboring values



[http://www.cooldictionary.com/
words/Bayer-filter.wikipedia](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 in Matlab

- Images represented as a matrix
- Suppose we have a NxM RGB image called “im”
 - $\text{im}(1, 1, 1)$ = top-left pixel value in R-channel
 - $\text{im}(y, x, b)$ = y pixels down, x pixels to right in the bth channel
 - $\text{im}(N, M, 3)$ = bottom-right pixel in B-channel
- **`imread(filename)`** returns a uint8 image (values 0 to 255)
 - Convert to double format (values 0 to 1) with `im2double`

		column												
												R	G	B
row		0.92	0.93	0.94	0.97	0.62	0.37	0.85	0.97	0.93	0.92	0.99		
		0.95	0.89	0.82	0.89	0.56	0.31	0.75	0.92	0.81	0.95	0.91		
1		0.89	0.72	0.51	0.55	0.51	0.42	0.57	0.41	0.49	0.91	0.92		
2		0.96	0.95	0.88	0.94	0.56	0.46	0.91	0.87	0.90	0.97	0.95		
3		0.71	0.81	0.81	0.87	0.57	0.37	0.80	0.88	0.89	0.79	0.85		
4		0.49	0.62	0.60	0.58	0.50	0.60	0.58	0.50	0.61	0.45	0.33		
5		0.86	0.84	0.74	0.58	0.51	0.39	0.73	0.92	0.91	0.49	0.74		
6		0.96	0.67	0.54	0.85	0.48	0.37	0.88	0.90	0.94	0.82	0.93		
7		0.69	0.49	0.56	0.66	0.43	0.42	0.77	0.73	0.71	0.90	0.99		
8		0.79	0.73	0.90	0.67	0.33	0.61	0.69	0.79	0.73	0.93	0.97		
9		0.91	0.94	0.89	0.49	0.41	0.78	0.78	0.77	0.89	0.99	0.93		
10		0.65	0.45	0.56	0.66	0.45	0.42	0.77	0.77	0.75	0.71	0.90	0.99	
11		0.79	0.73	0.90	0.67	0.33	0.61	0.69	0.79	0.73	0.93	0.97		
12		0.91	0.94	0.89	0.49	0.41	0.78	0.78	0.77	0.89	0.99	0.93		
13		0.65	0.45	0.56	0.66	0.45	0.42	0.77	0.77	0.75	0.71	0.90	0.99	
14		0.79	0.73	0.90	0.67	0.33	0.61	0.69	0.79	0.73	0.93	0.97		
15		0.91	0.94	0.89	0.49	0.41	0.78	0.78	0.77	0.89	0.99	0.93		
16		0.65	0.45	0.56	0.66	0.45	0.42	0.77	0.77	0.75	0.71	0.90	0.99	
17		0.79	0.73	0.90	0.67	0.33	0.61	0.69	0.79	0.73	0.93	0.97		
18		0.91	0.94	0.89	0.49	0.41	0.78	0.78	0.77	0.89	0.99	0.93		

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



Some drawbacks

- Strongly correlated channels
- Non-perceptual



R
(G=0,B=0)






G
(R=0,B=0)



B
(R=0,G=0)

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

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

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.



[after Foley et al.]

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

Color spaces: HSV

- **Hue, Saturation, Value**
- 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

- **Hue, Saturation, Value**
- Nonlinear – reflects topology of colors by coding **hue** as an angle
- Matlab: `hsv2rgb`, `rgb2hsv`.





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

Color spaces: HSV




Intuitive color space



H
($S=1, V=1$)



S
($H=1, V=1$)




V
($H=1, S=0$)

Slide credit: D. Hoiem

Color spaces: YCbCr

Fast to compute, good for compression, used by TV



Y
($Cb=0.5, Cr=0.5$)




Cb
($Y=0.5, Cr=0.5$)



Cr
($Y=0.5, Cb=0.5$)

Color spaces: YCbCr

Fast to compute, good for compression, used by TV



Y
(Cb=0.5,Cr=0.5)




Cb
(Y=0.5,Cr=0.5)

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

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




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'

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




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

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

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

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

(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




copyright 2000 philg@mit.edu

Only intensity shown – constant color

Slide credit: D. Hoiem

Most information in intensity




copyright 2000 philg@mit.edu

Original image

Slide credit: D. Hoiem

Back to grayscale intensity



Slide credit: D. Hoiem

Today

- Image formation
- Digital images
- Perception of color and light
- Color spaces

Next week

- Point operations
- Histogram processing