Computer Vision Course
Lecture 02

Image Formation
Light and Color

Ceyhun Burak Akgül, PhD
cba-research.com

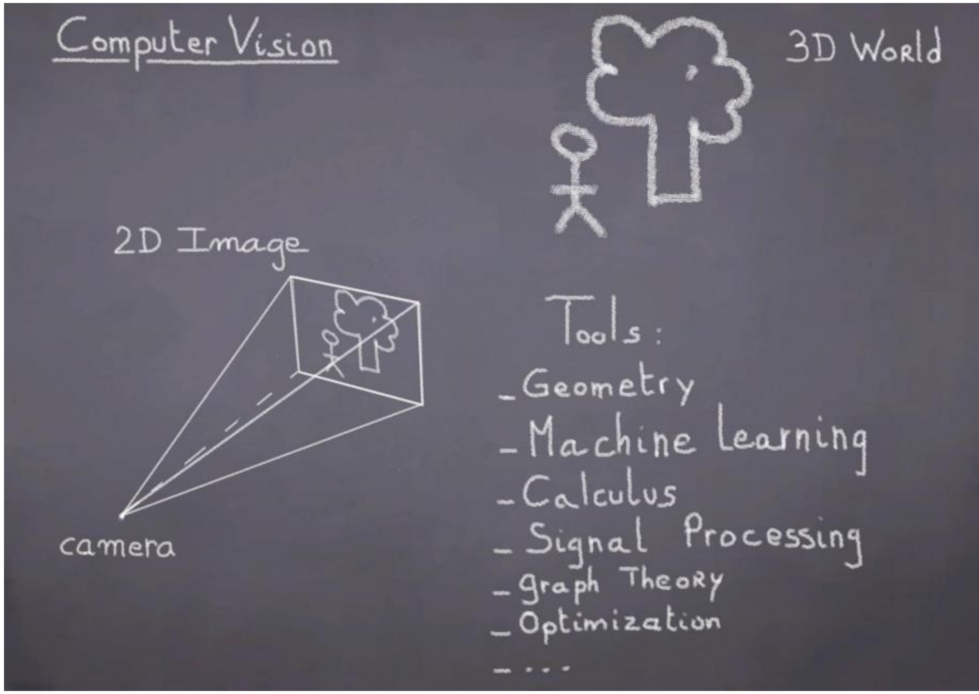

Photo credit: Olivier Teboul
vision.mas.ecp.fr/Personnel/teboul

Spring 2015
Last updated 04/03/2015

These slides have been adapted from James Hays's 2014 Computer Vision course slides at Brown University.

Course Outline

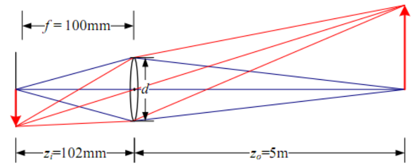

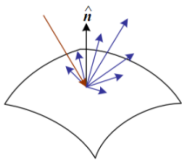

Image Formation and Processing

Light, Shape and Color

The Pin-hole Camera Model, The Digital Camera

Linear filtering, Filter banks, Multiresolution

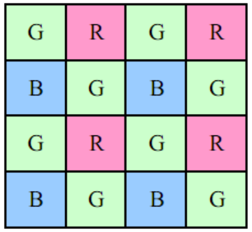

Feature Detection and Matching

Edge Detection, Interest Points: Corners and Blobs

Local Image Descriptors

Feature Matching and Hough Transform

Multiple Views and Motion

Geometric Transformations, Camera Calibration

Feature Tracking , Stereo Vision

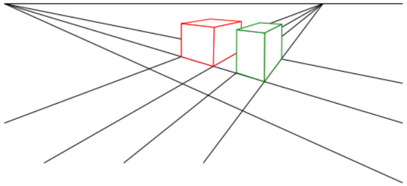

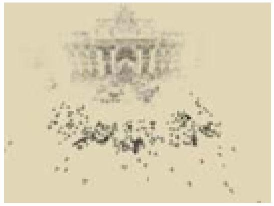

Segmentation and Grouping

Segmentation by Clustering, Region Merging and Growing

Advanced Methods Overview: Active Contours, Level-Sets, Graph-Theoretic Methods

Detection and Recognition

Problems and Architectures Overview

Statistical Classifiers, Bag-of-Words Model, Detection by Sliding Windows

A Very Primitive Image

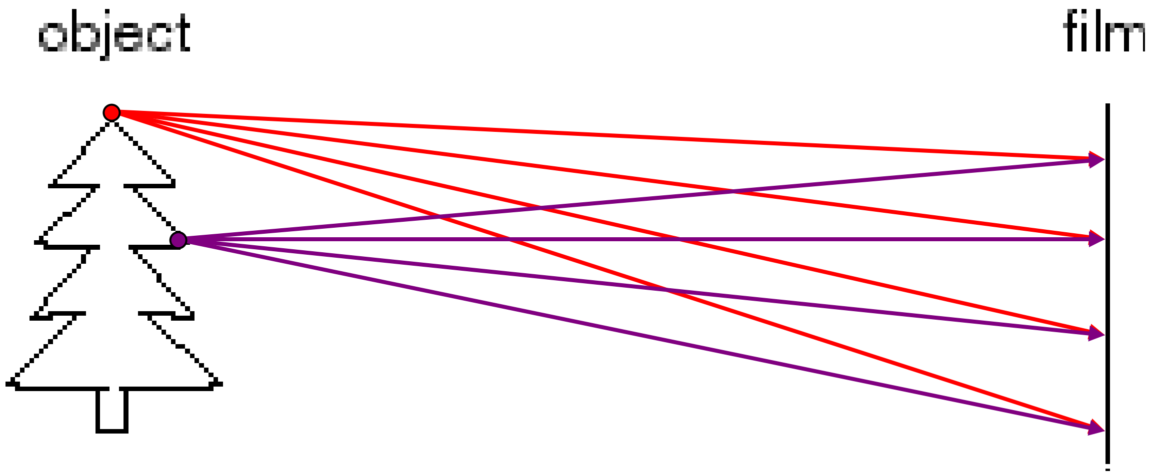

Let's design a camera

- Idea 1: put a piece of film in front of an object
- Do we get a reasonable image?

Slide source: Seitz

Pinhole Camera

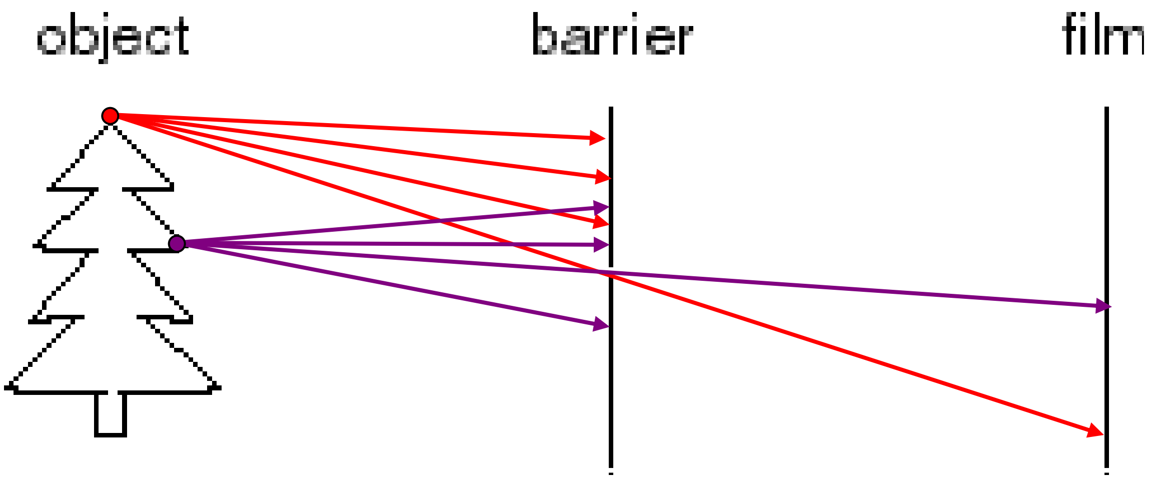

Idea 2: add a barrier to block off most of the rays

- This reduces blurring
- The opening known as the **aperture**

Slide source: Seitz

Pinhole Camera

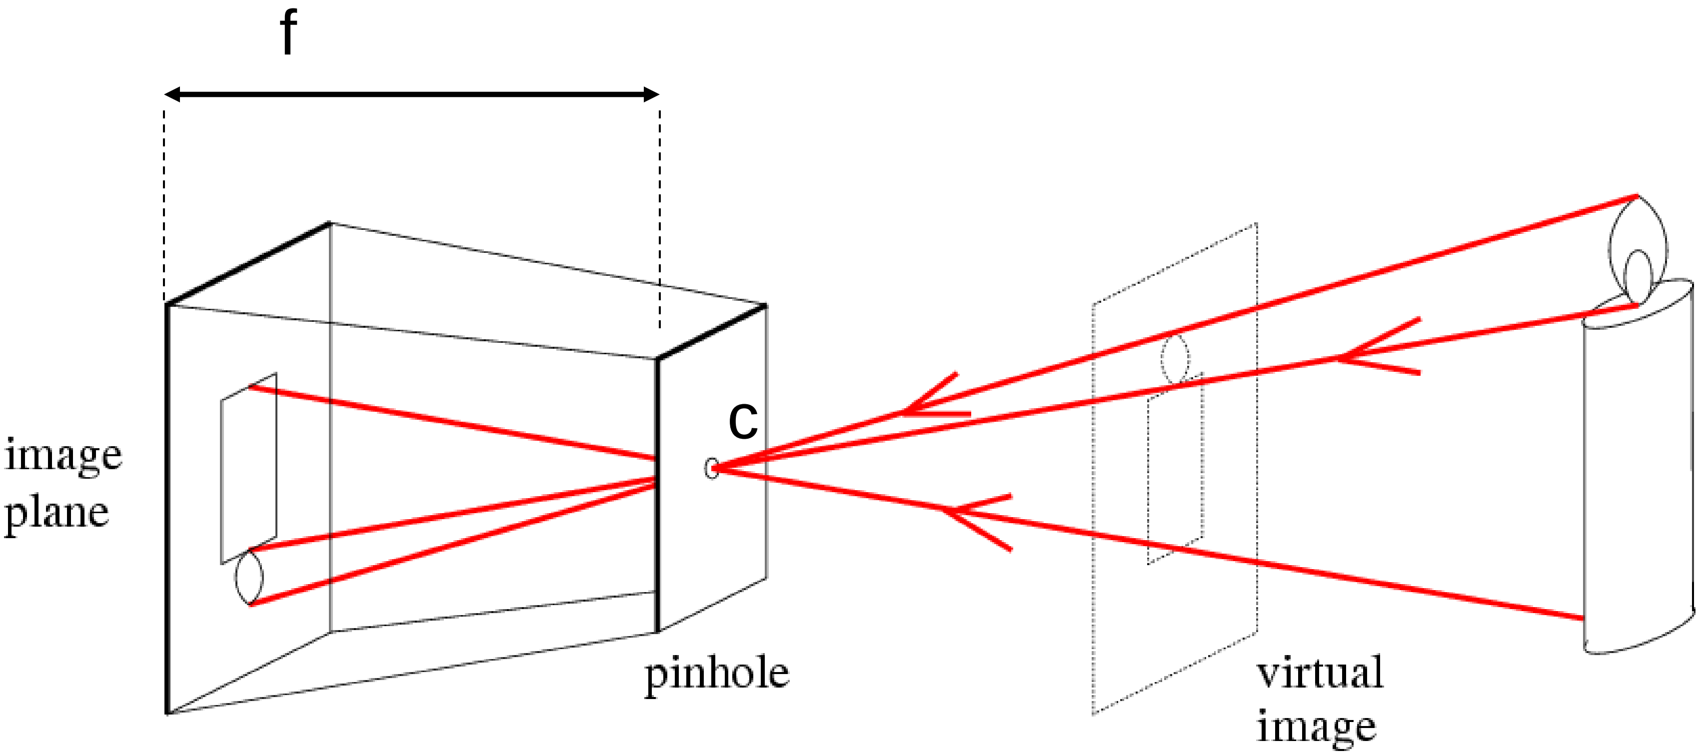

f = focal length
c = center of the camera

Figure from Forsyth

Mapping 3D World to 2D Plane

What is lost?

- Length

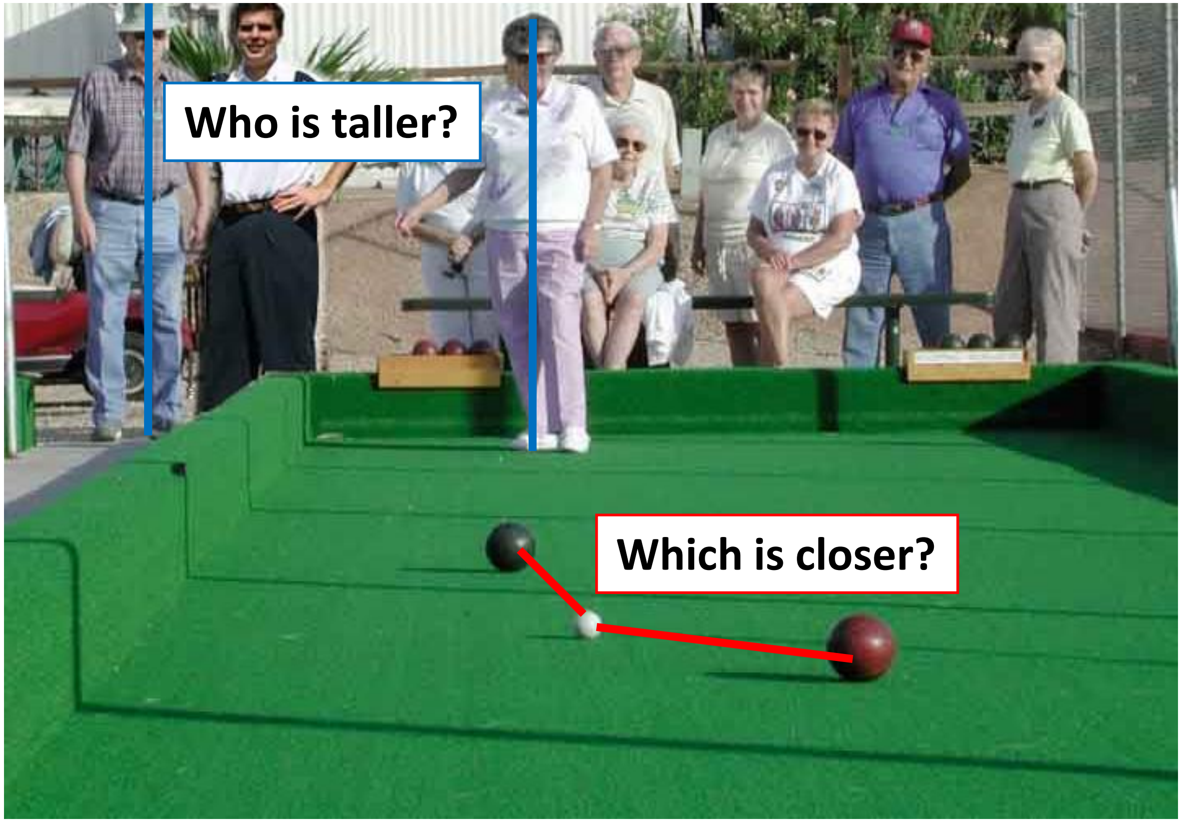

Length is not preserved

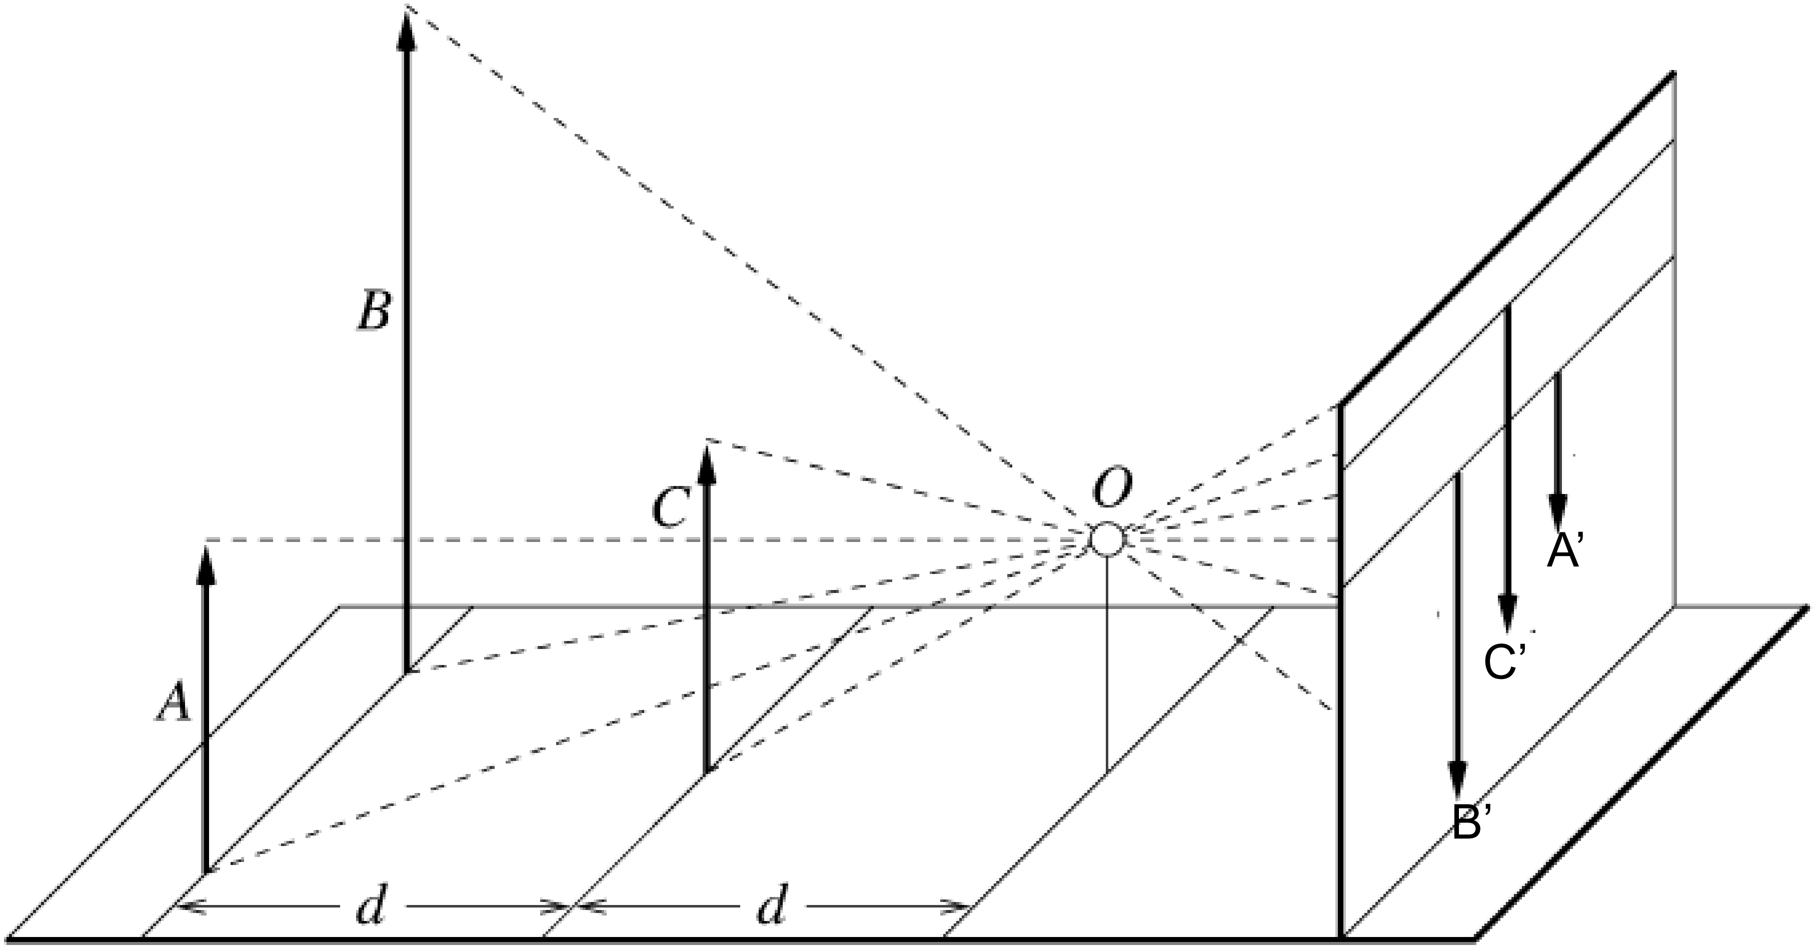

Mapping 3D World to 2D Plane

What is lost?

- Length
- Angle

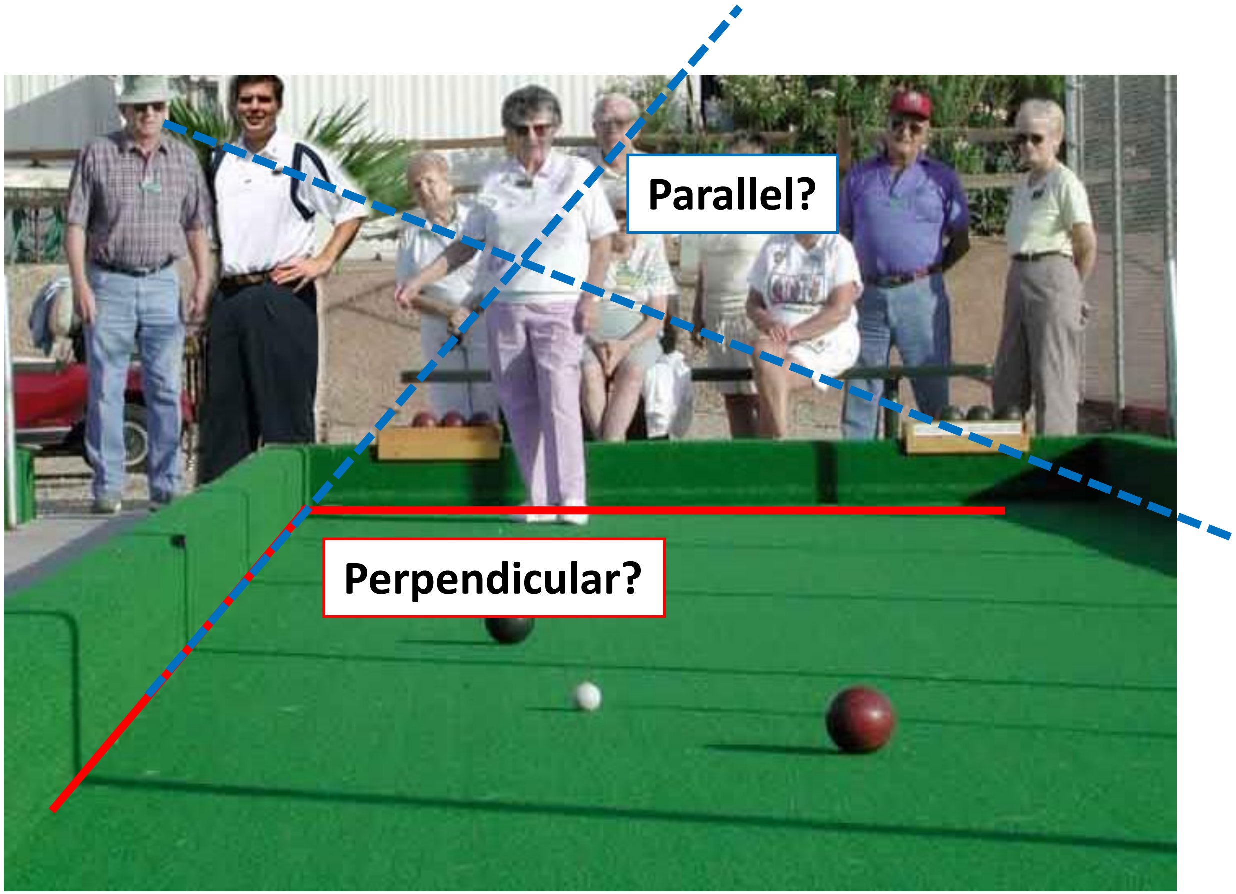

Mapping 3D World to 2D Plane

What is preserved?

- Straight lines are still straight

Vanishing Points and Lines

Parallel lines in the world intersect in the image at a “vanishing point”

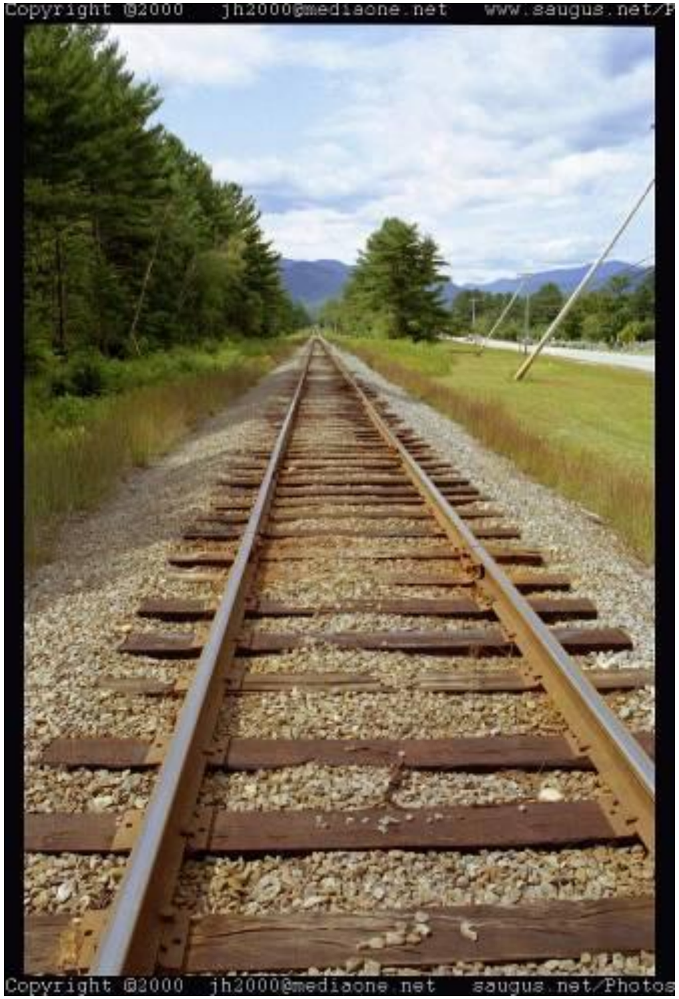

Vanishing Points and Lines

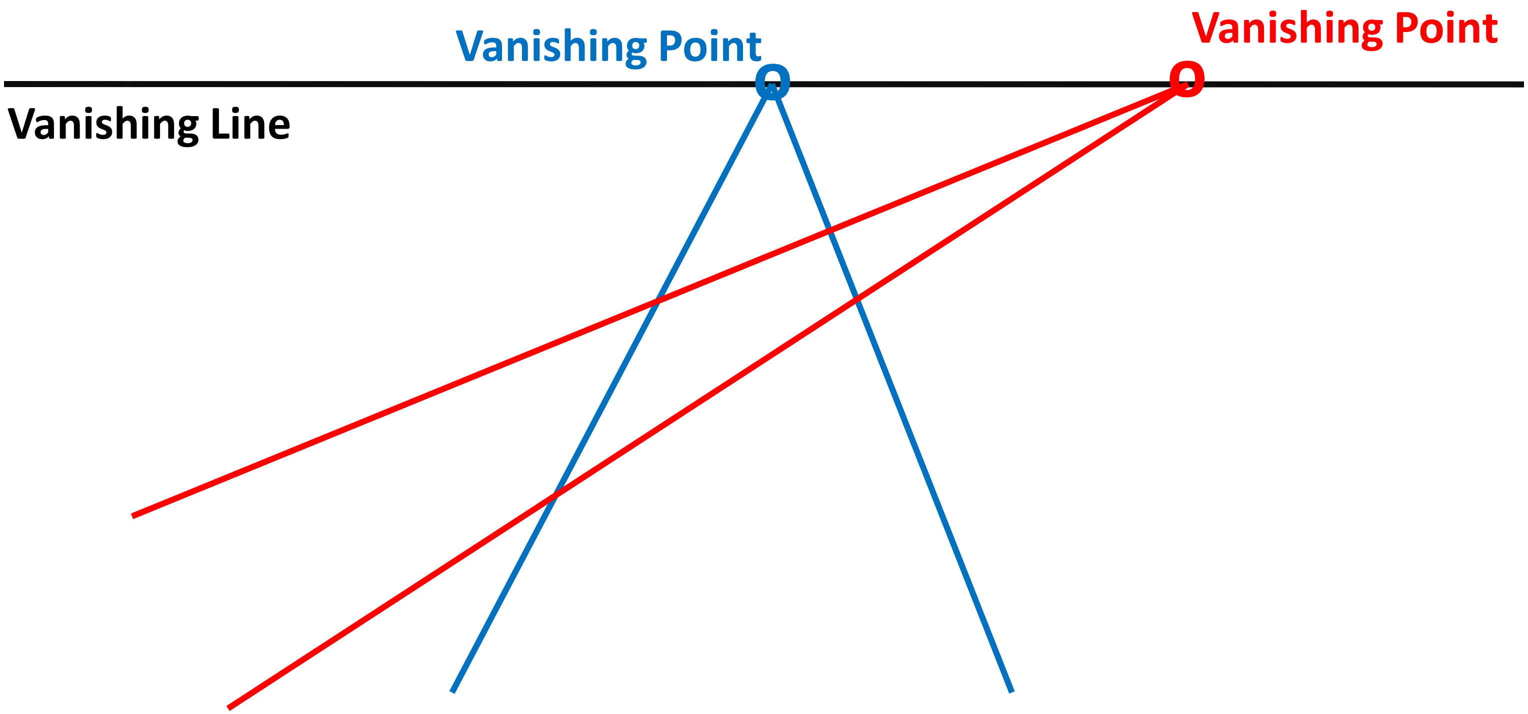

Vanishing Points and Lines

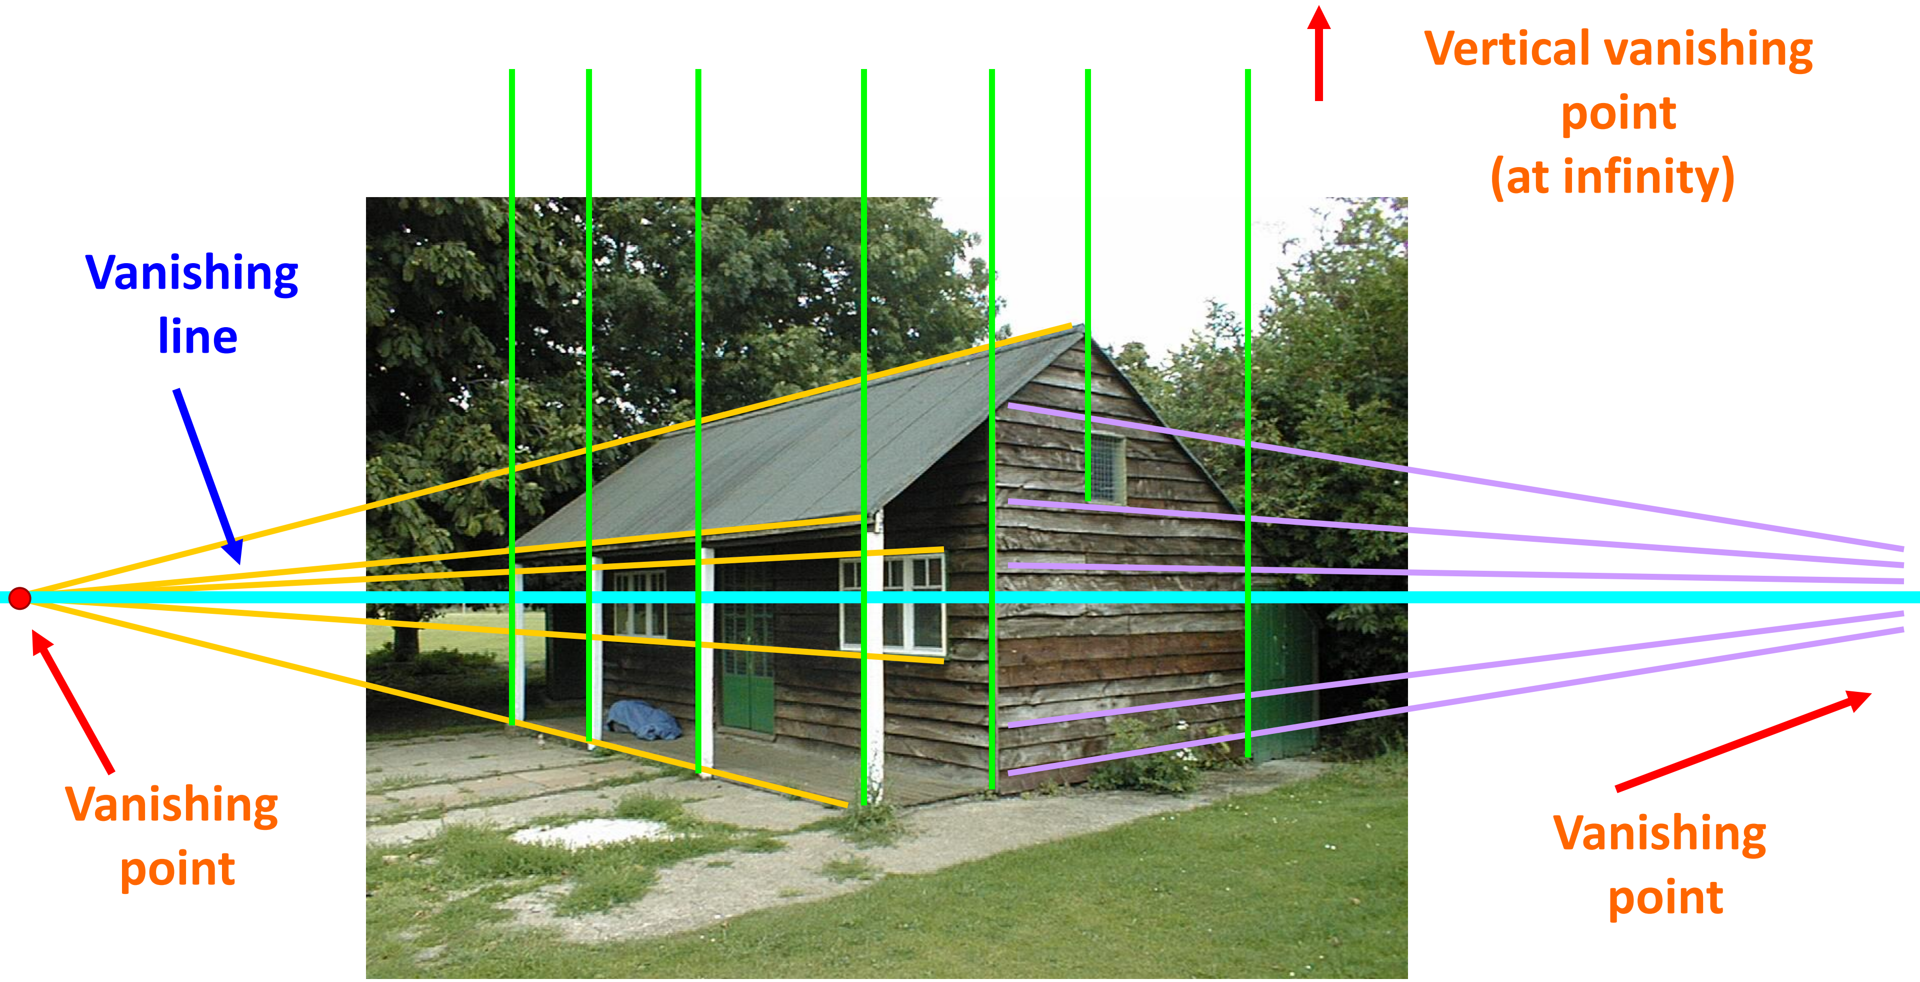

Slide from Efros, Photo from Criminisi

Mapping 3D World to 2D Plane

Projection = World coordinates → Image coordinates

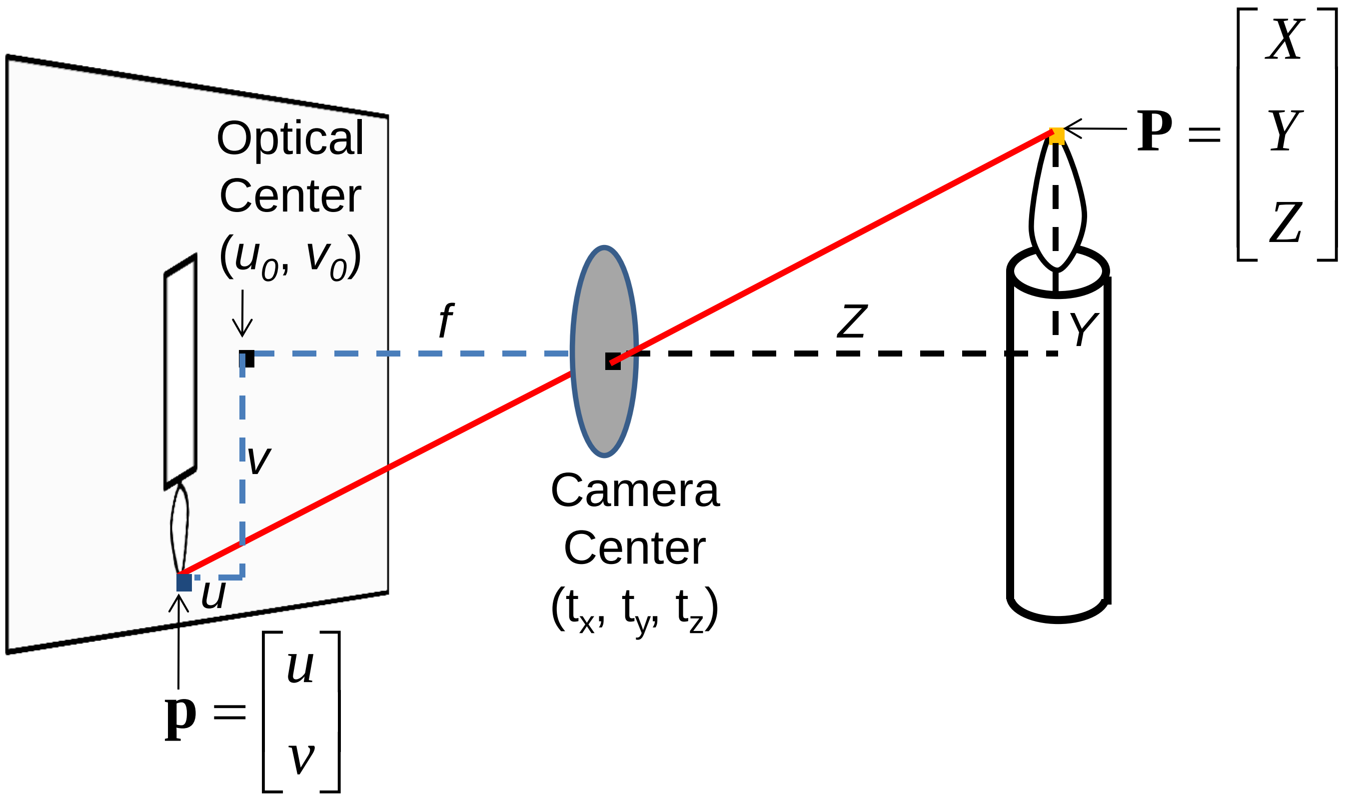

Adding a Lens

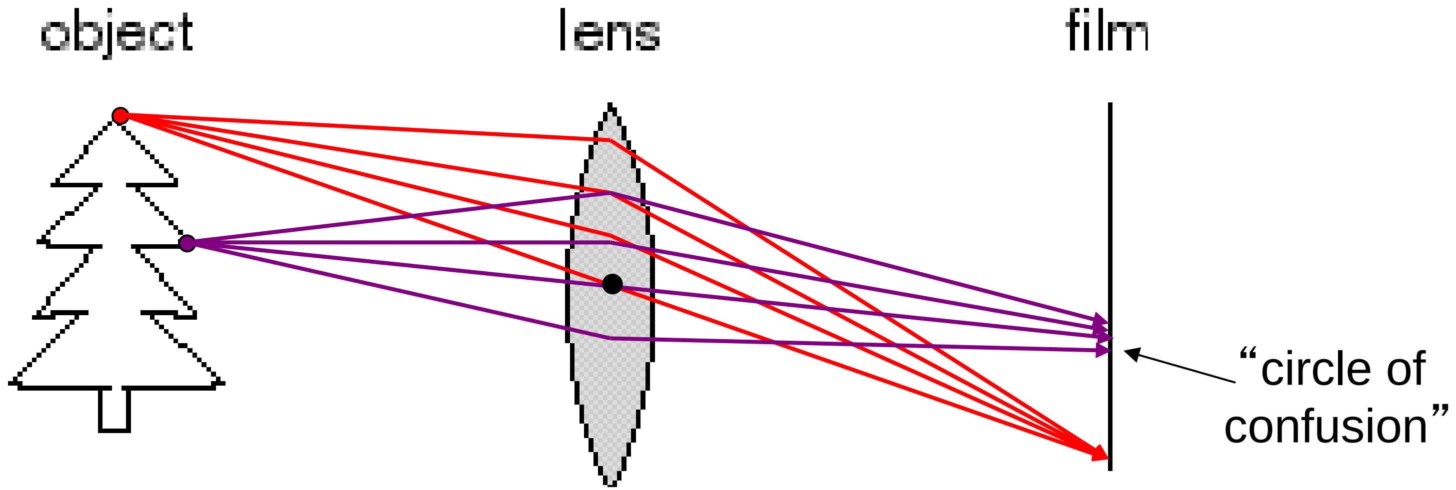

A lens focuses light onto the film

- 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

Basic Characteristics of a Lens

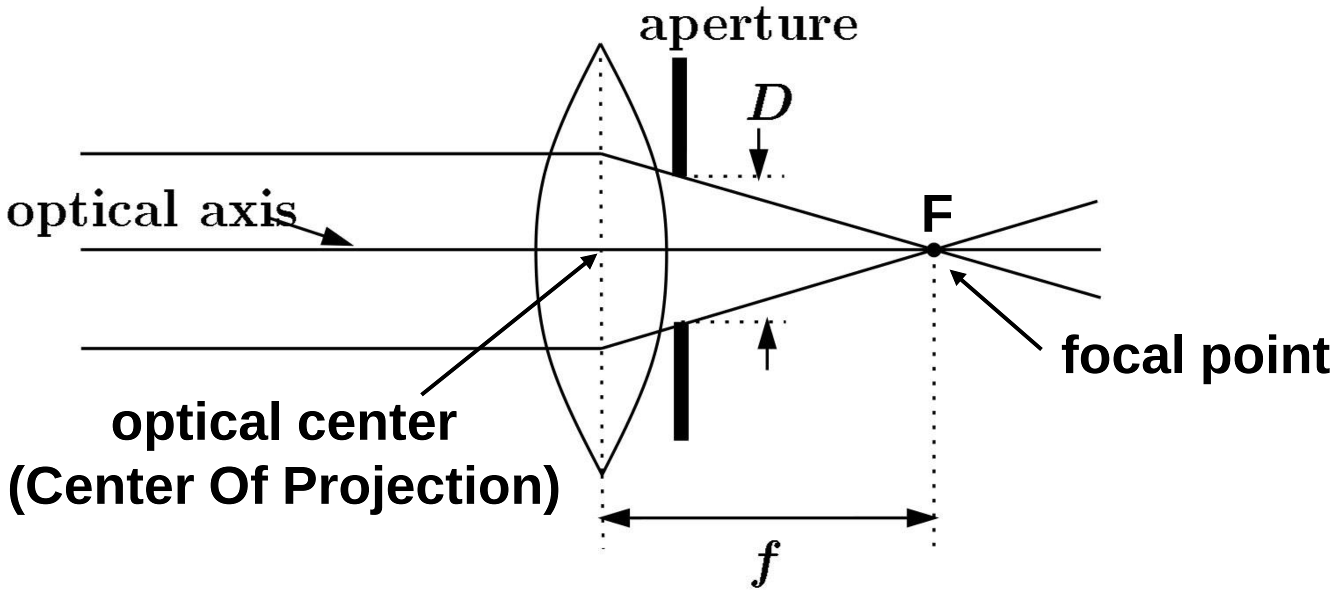

A lens focuses parallel rays onto a single focal point

- Focal point at a distance *f* beyond the plane of the lens
- Aperture of diameter D restricts the range of rays

Slide source: Seitz

Depth of Field

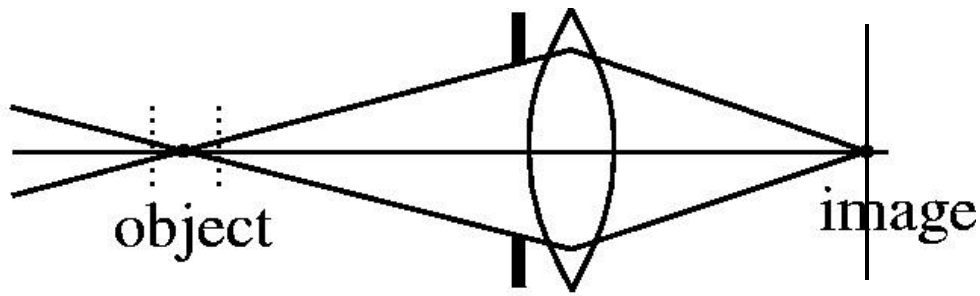

f / 5.6

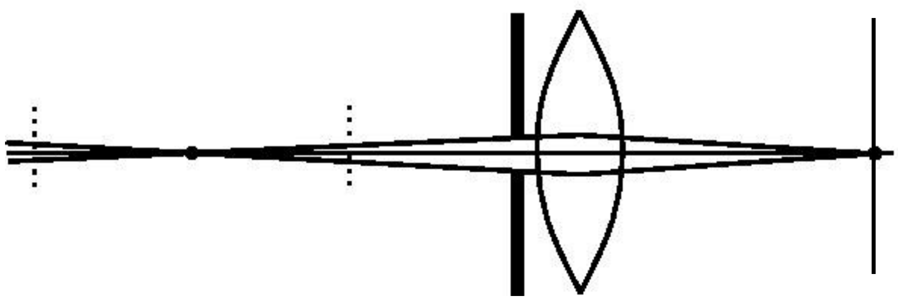

f / 32

Changing the aperture size or focal length affects depth of field

Image Formation

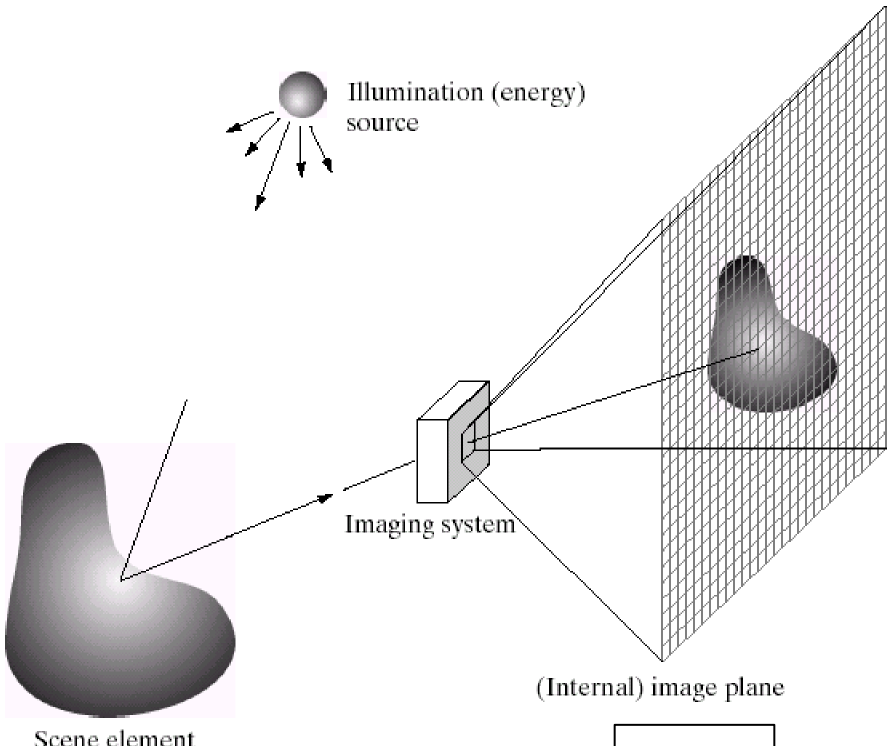

Film

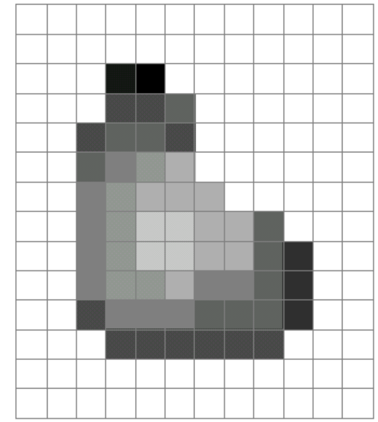

Digital Camera

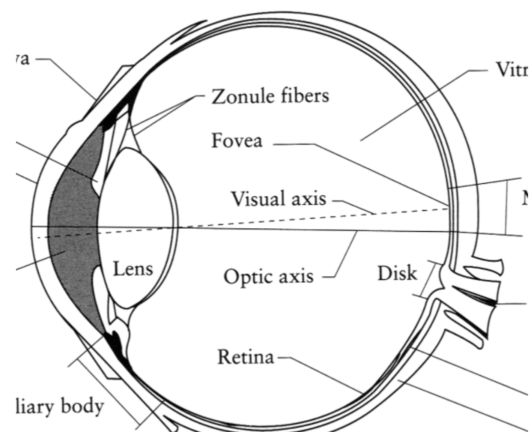

The Eye

What Happens to a Traveling Photon?

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

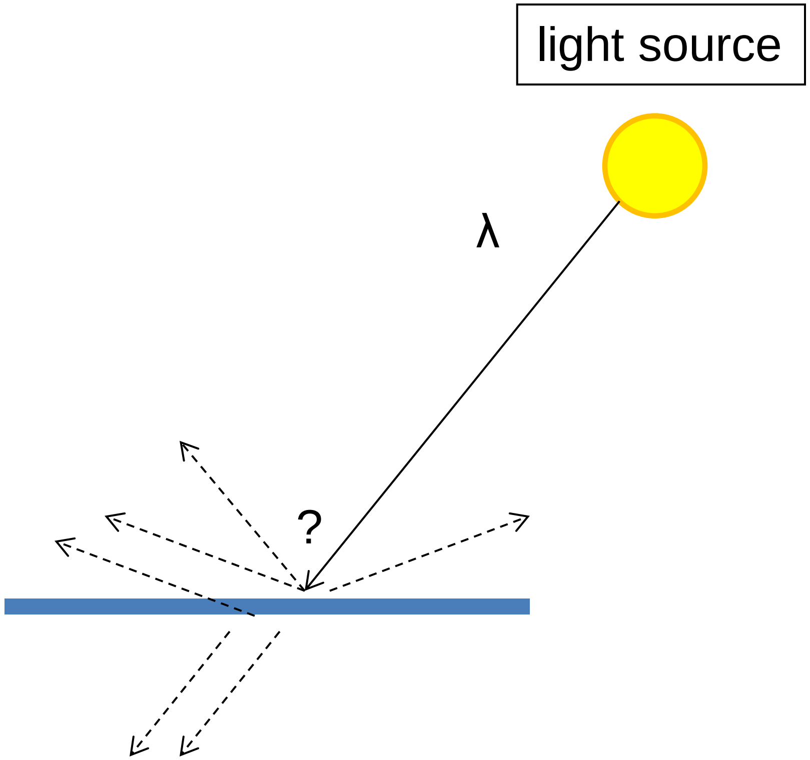

What Happens to a Traveling Photon?

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

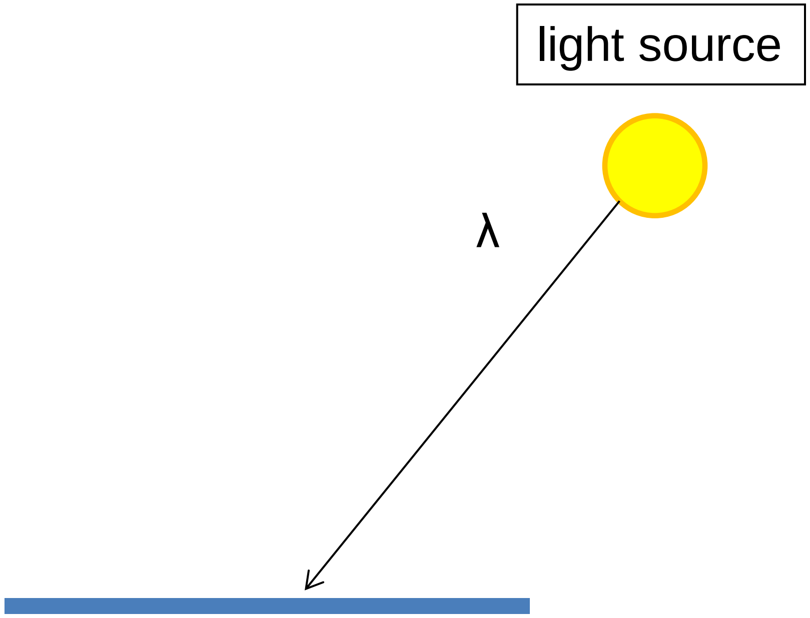

What Happens to a Traveling Photon?

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

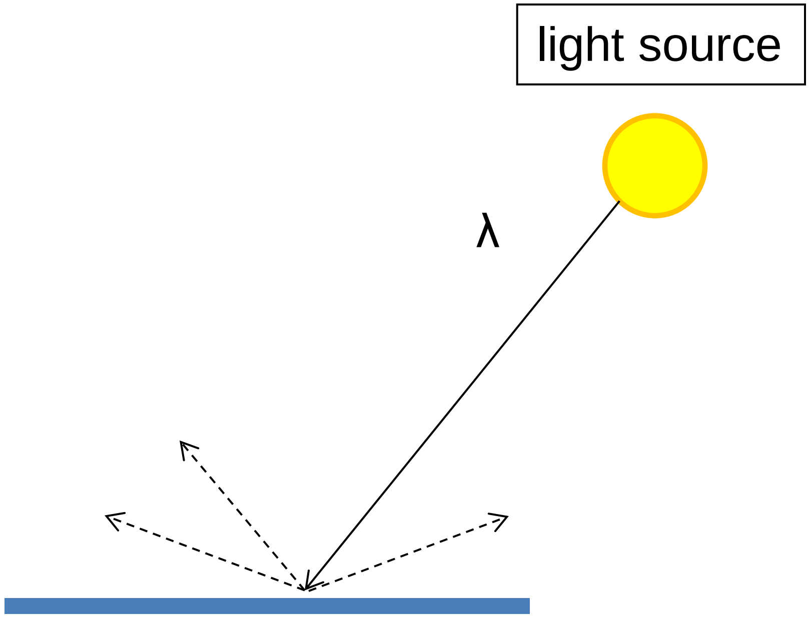

What Happens to a Traveling Photon?

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

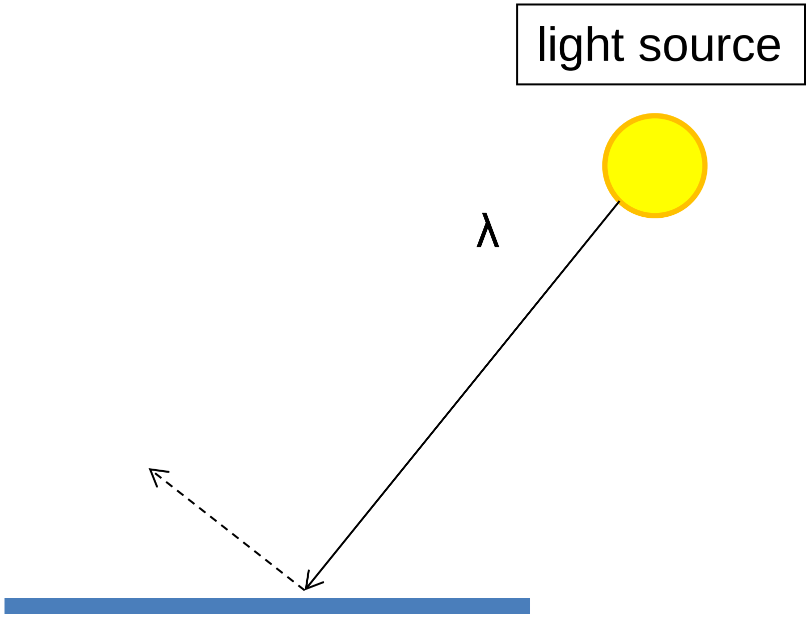

What Happens to a Traveling Photon?

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

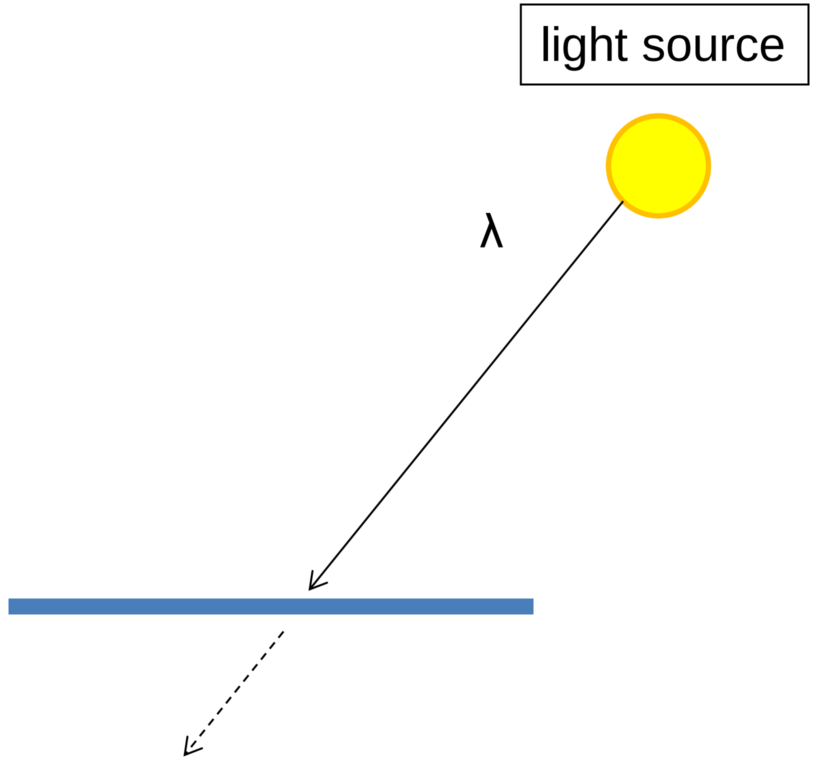

What Happens to a Traveling Photon?

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

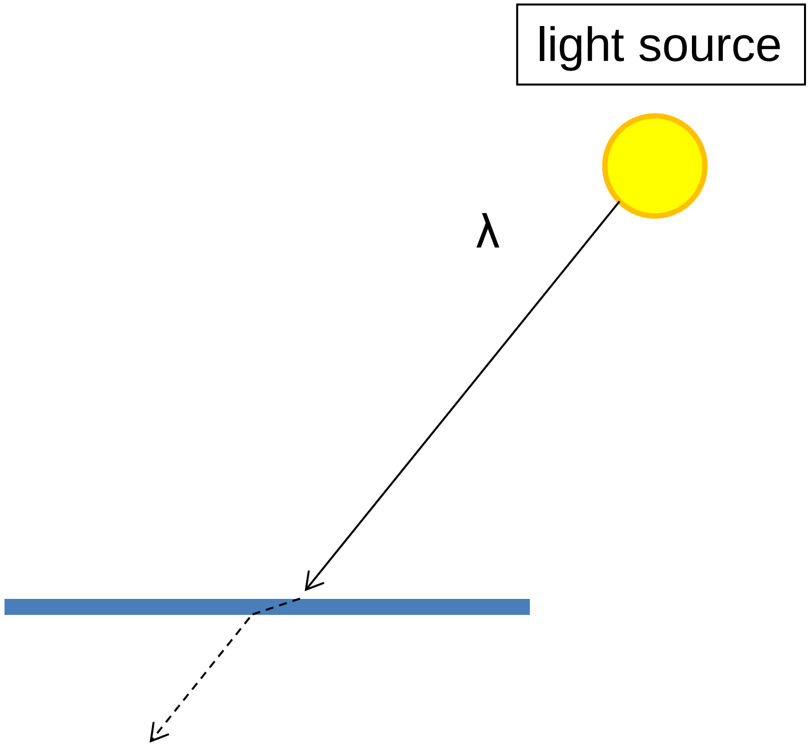

What Happens to a Traveling Photon?

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

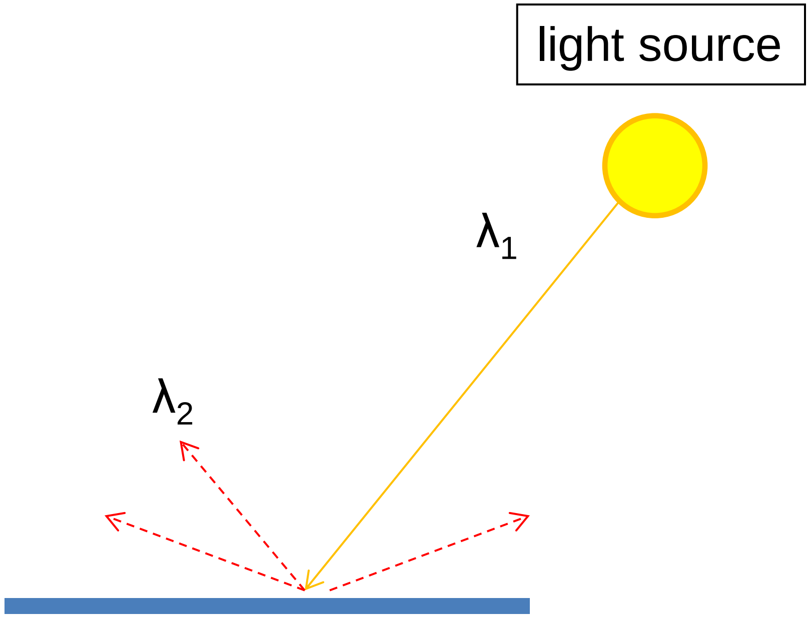

What Happens to a Traveling Photon?

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

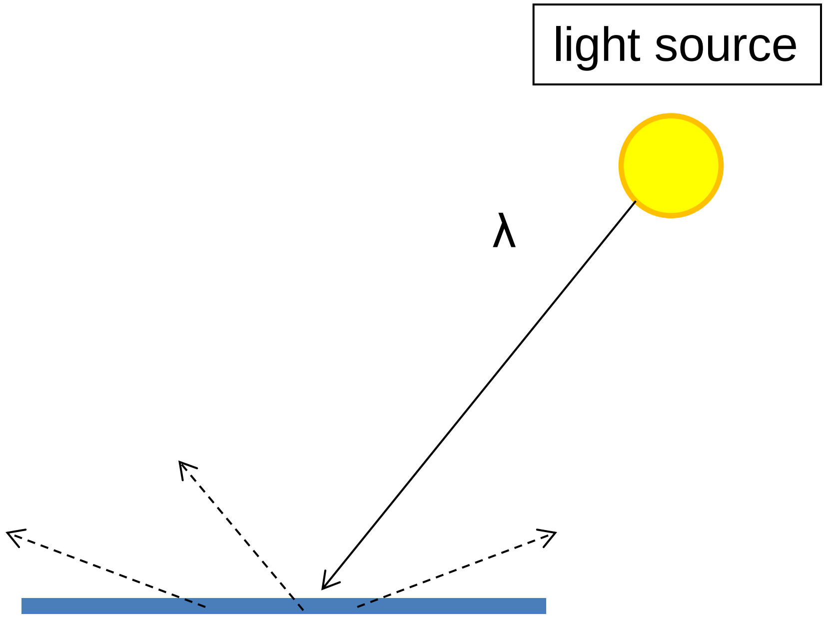

What Happens to a Traveling Photon?

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

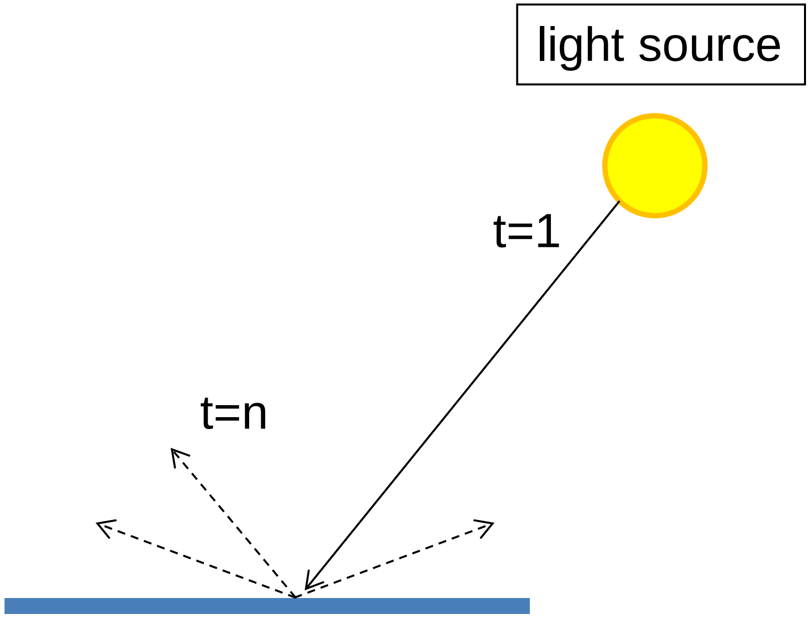

What Happens to a Traveling Photon?

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

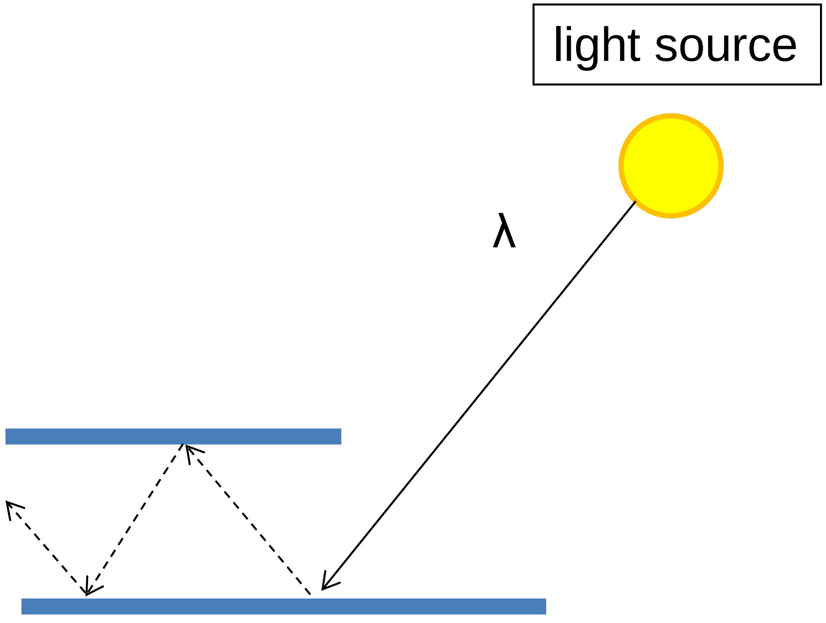

(Specular Interreflection)

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) and CMOS

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

Slide by Steve Seitz

Sensor Array

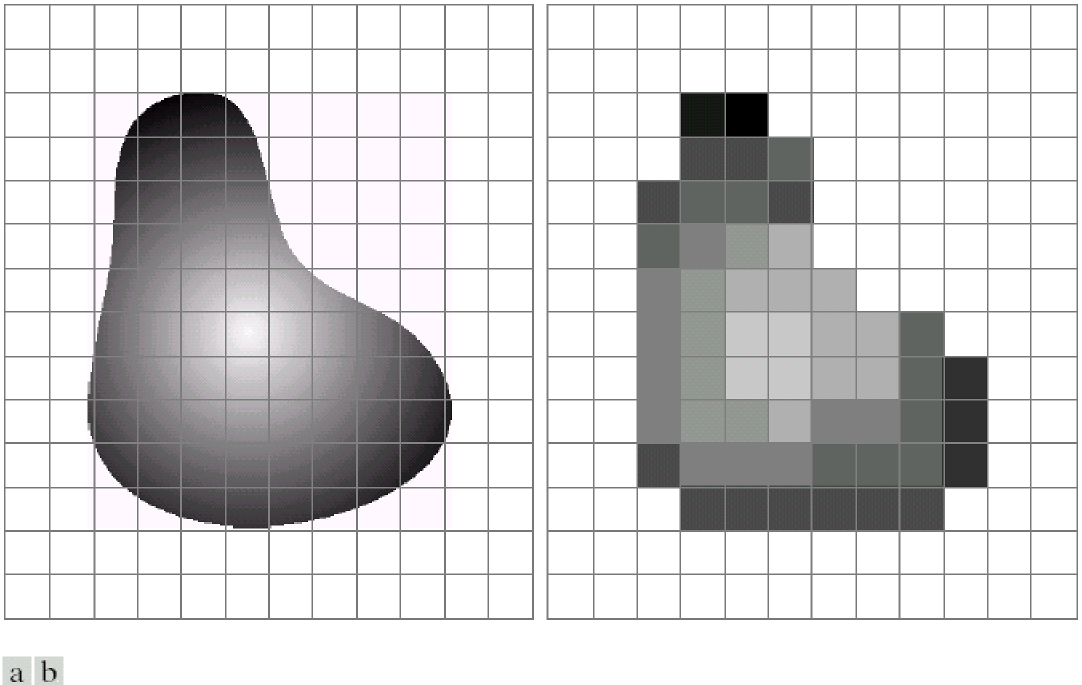

FIGURE 2.17 (a) Continuos image projected onto a sensor array. (b) Result of image sampling and quantization.

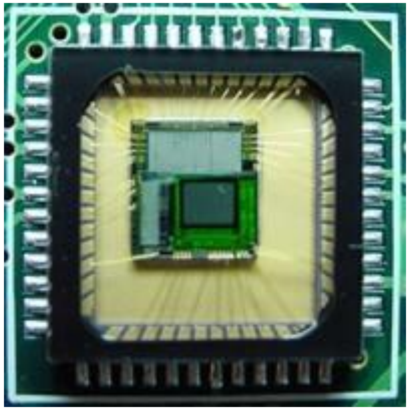

CMOS sensor

Sampling and Quantization

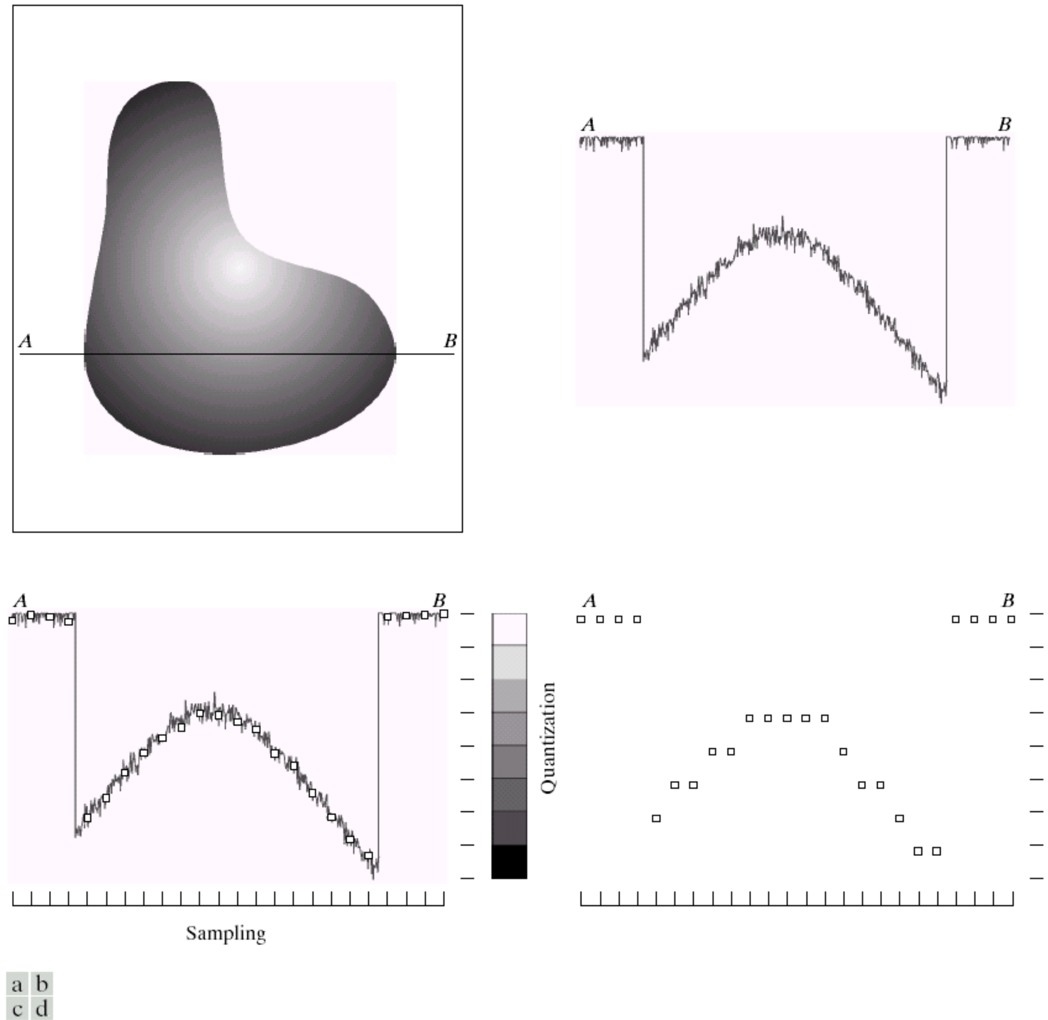

FIGURE 2.16 Generating a digital image. (a) Continuous image. (b) A scan line from A to B in the continuous image, used to illustrate the concepts of sampling and quantization. (c) Sampling and quantization. (d) Digital scan line.

The Eye

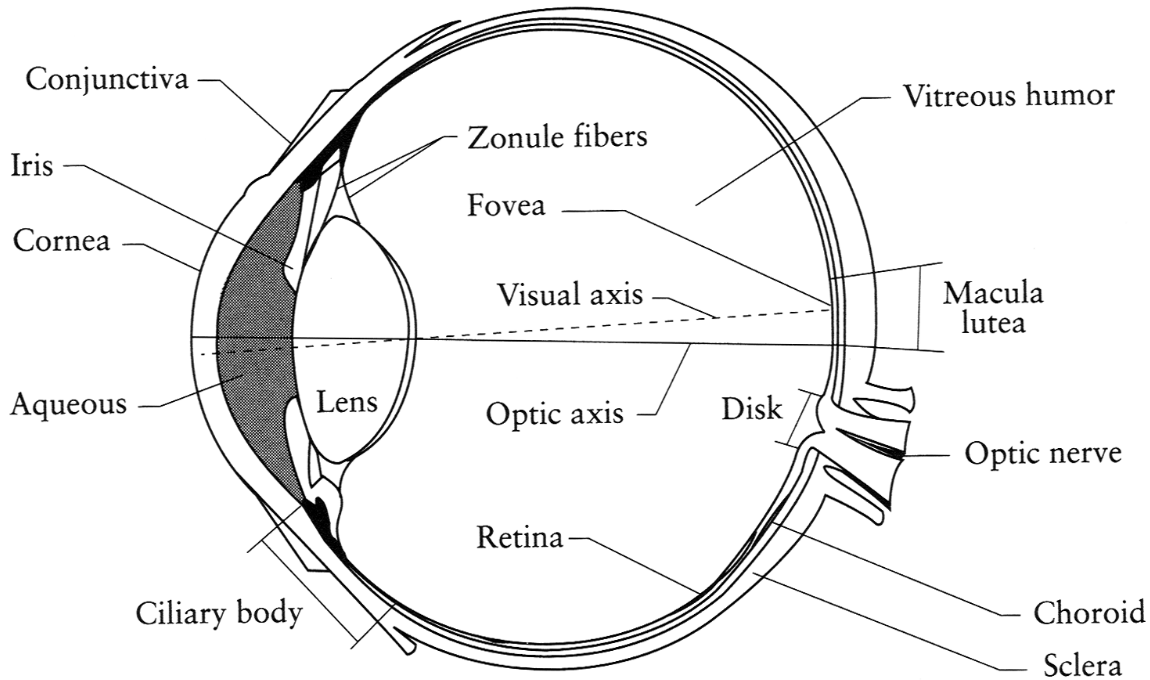

The human eye is a camera!

- **Iris** – Colored annulus with radial muscles
- **Pupil** – The hole (aperture) whose size is controlled by the iris
- **Retina: Cones and Rodes** – The «Film»

Slide by Steve Seitz

The Retina

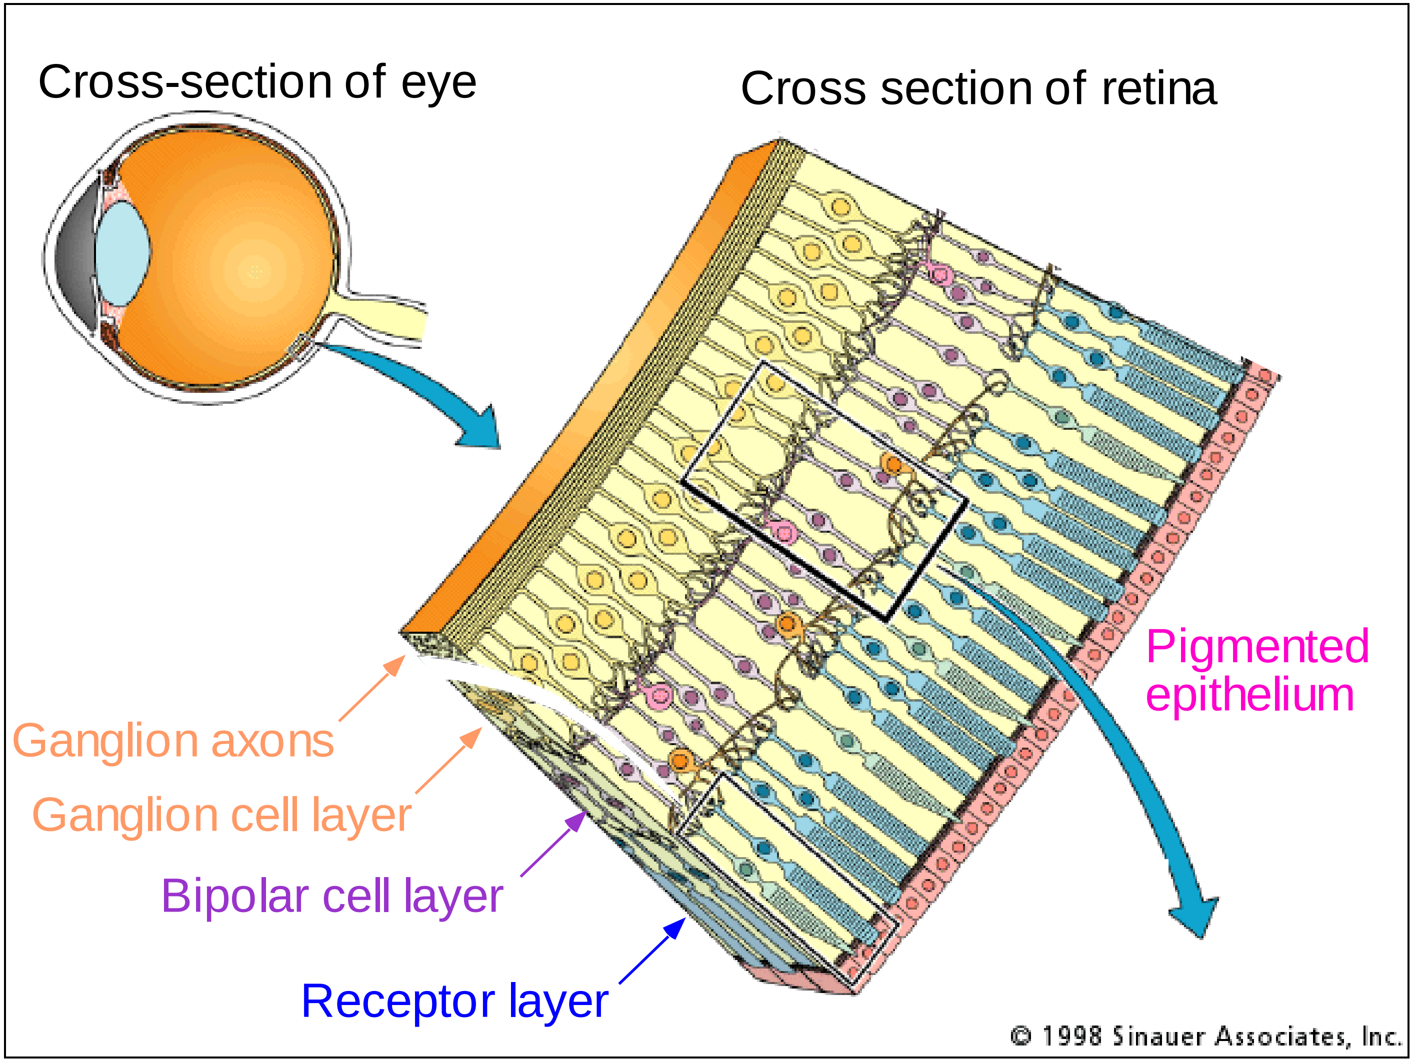

Cross-section of eye
Cross section of retina
Ganglion axons
Ganglion cell layer
Bipolar cell layer
Receptor layer
Pigmented epithelium
© 1998 Sinauer Associates, Inc.

Two Types of Light-Sensitive Receptors

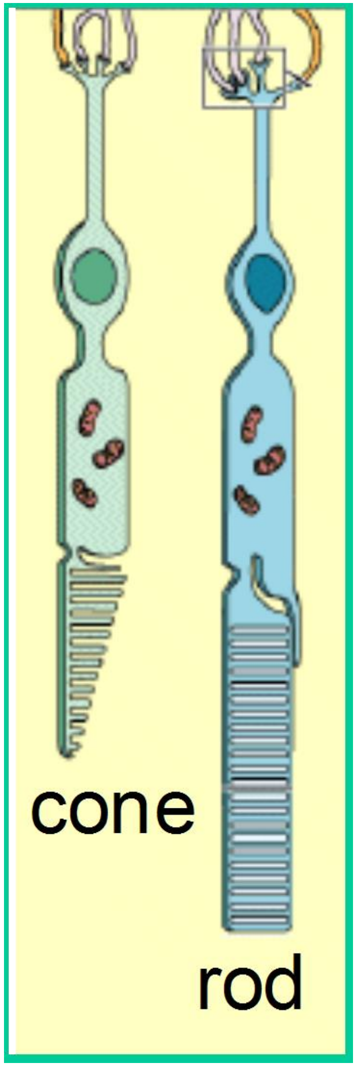

Cones

cone-shaped
less sensitive
operate in high light
color vision

Rods

rod-shaped
highly sensitive
operate at night
gray-scale vision (shape)

Electromagnetic vs. Visible Spectrum

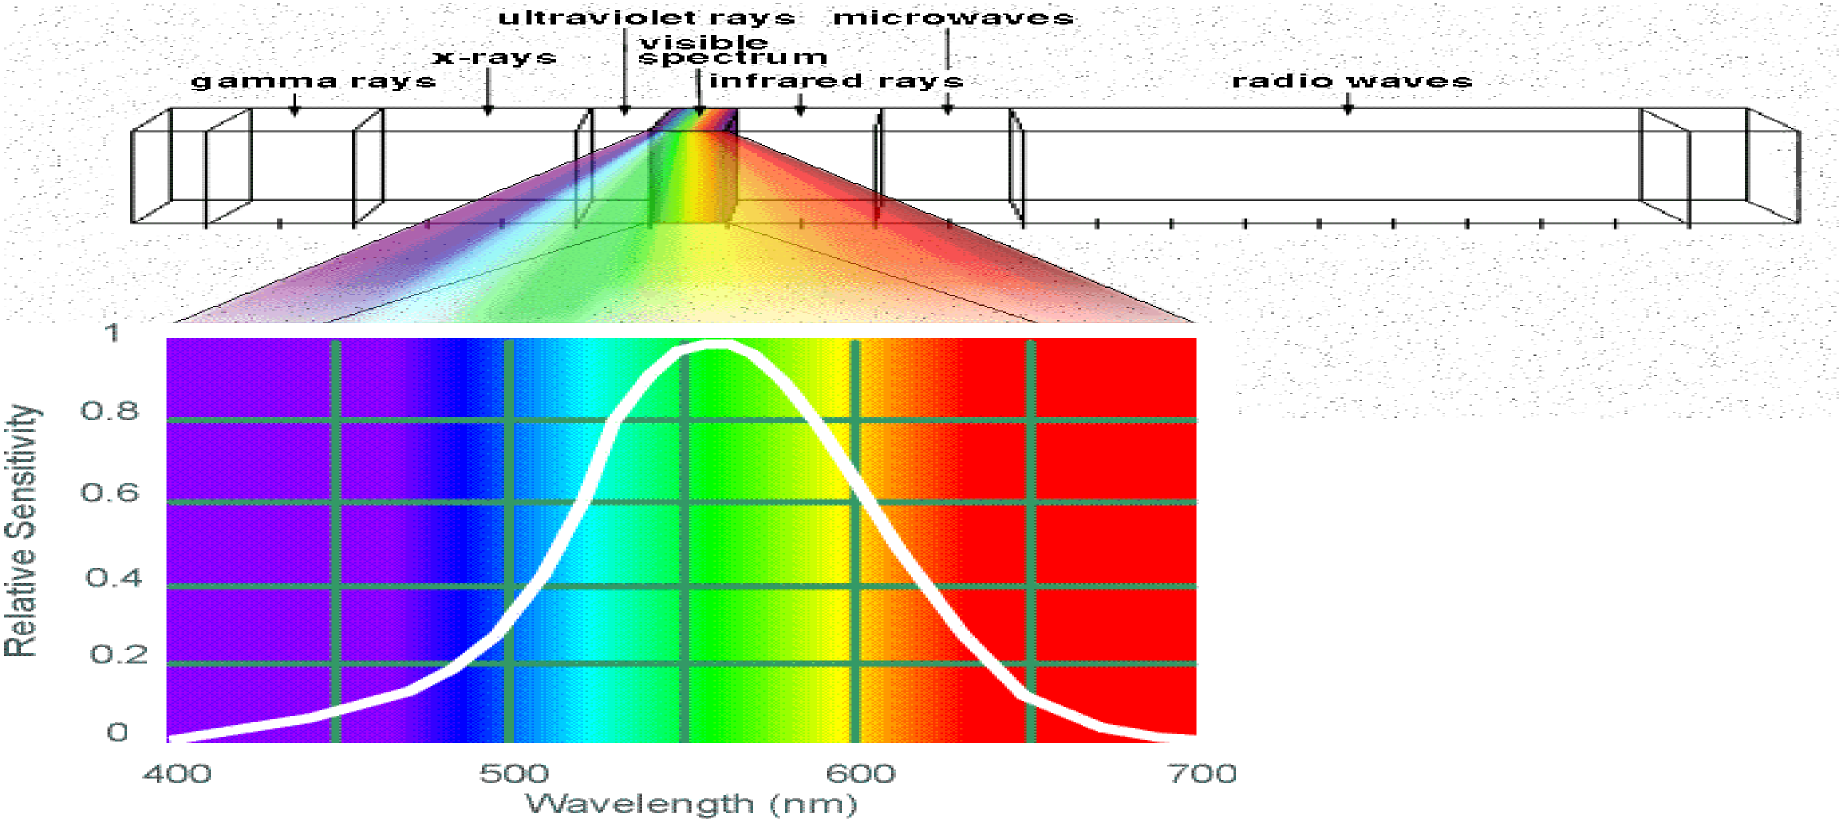

Human Luminance Sensitivity Function

Light Source Spectra

Some examples of the spectra of **light sources**

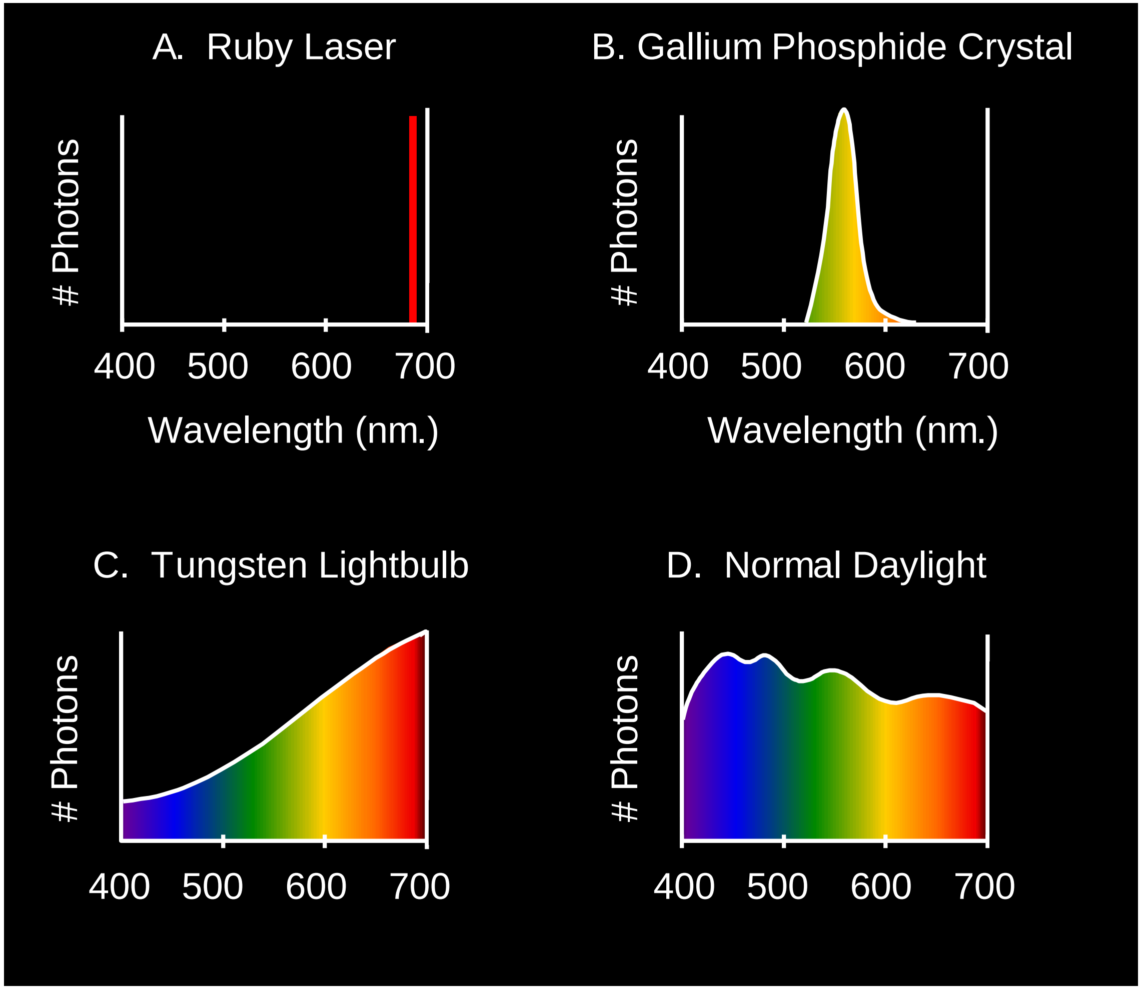

Reflectance Spectra

Some examples of the **reflectance** spectra of **surfaces**

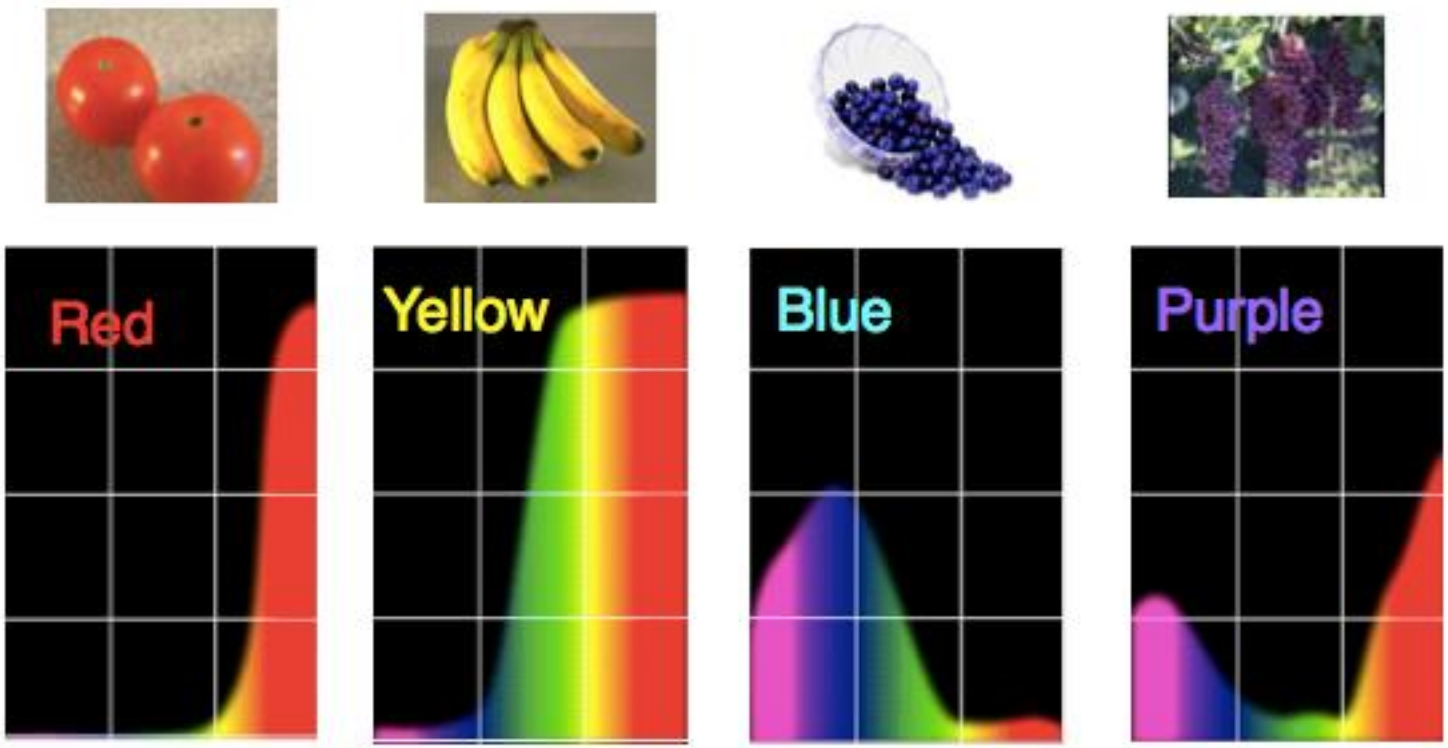

Physics vs. Perception

There is no simple functional description for the perceived color of all lights under all viewing conditions, but there is…

A helpful constraint:
Consider only physical spectra with normal distributions

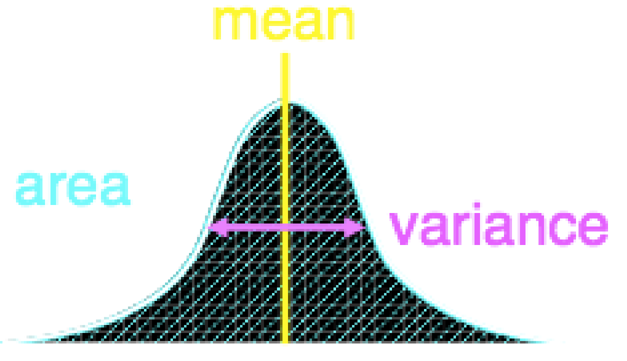

Physics vs. Perception

Mean ⟷ Hue

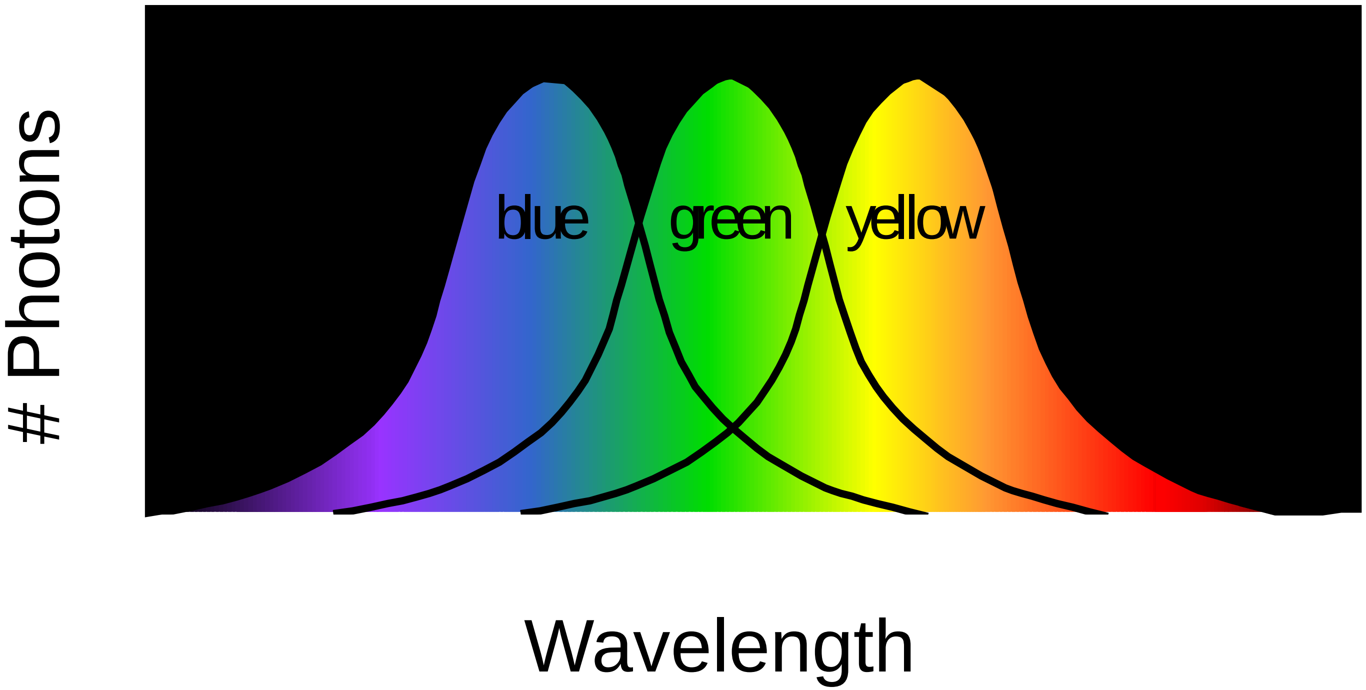

Physics vs. Perception

Variance ⟷ Saturation

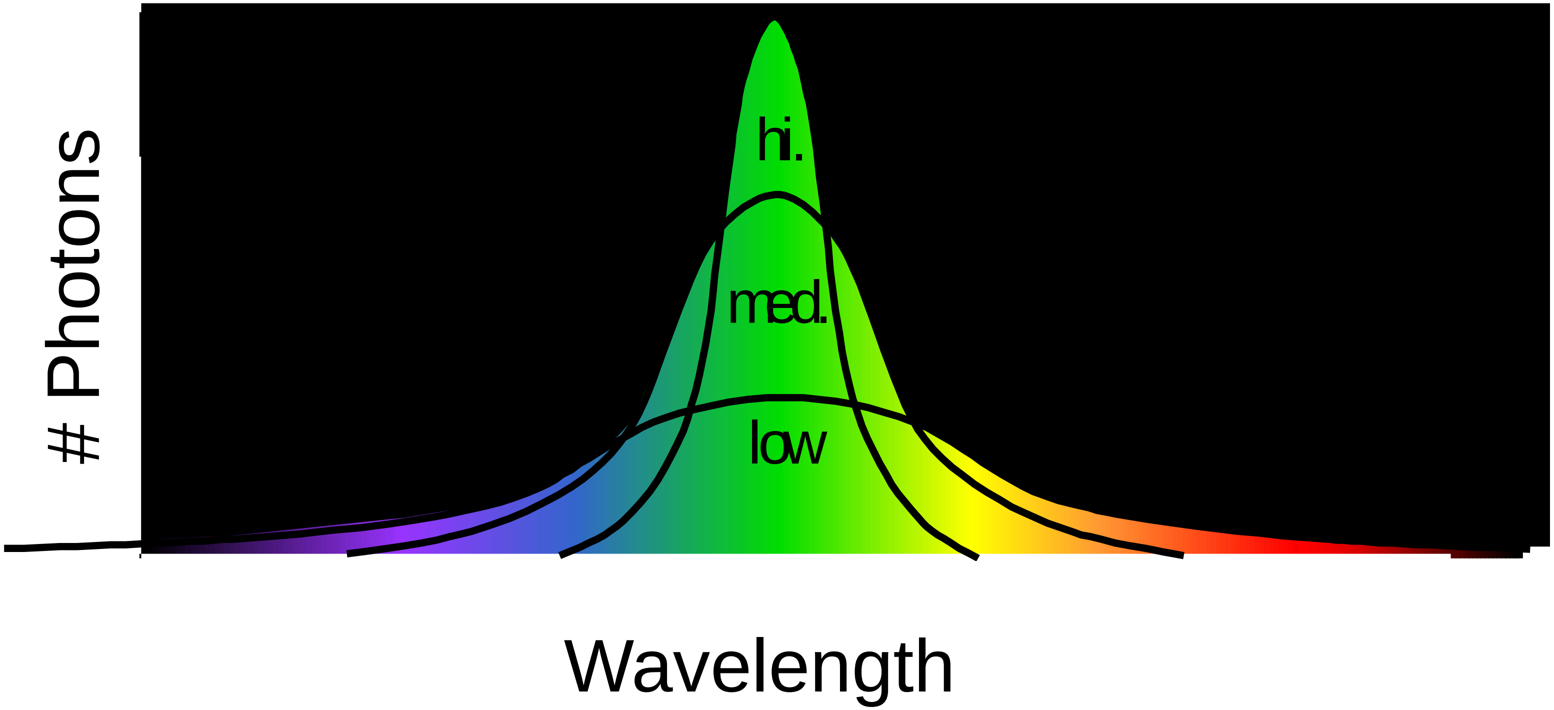

Physics vs. Perception

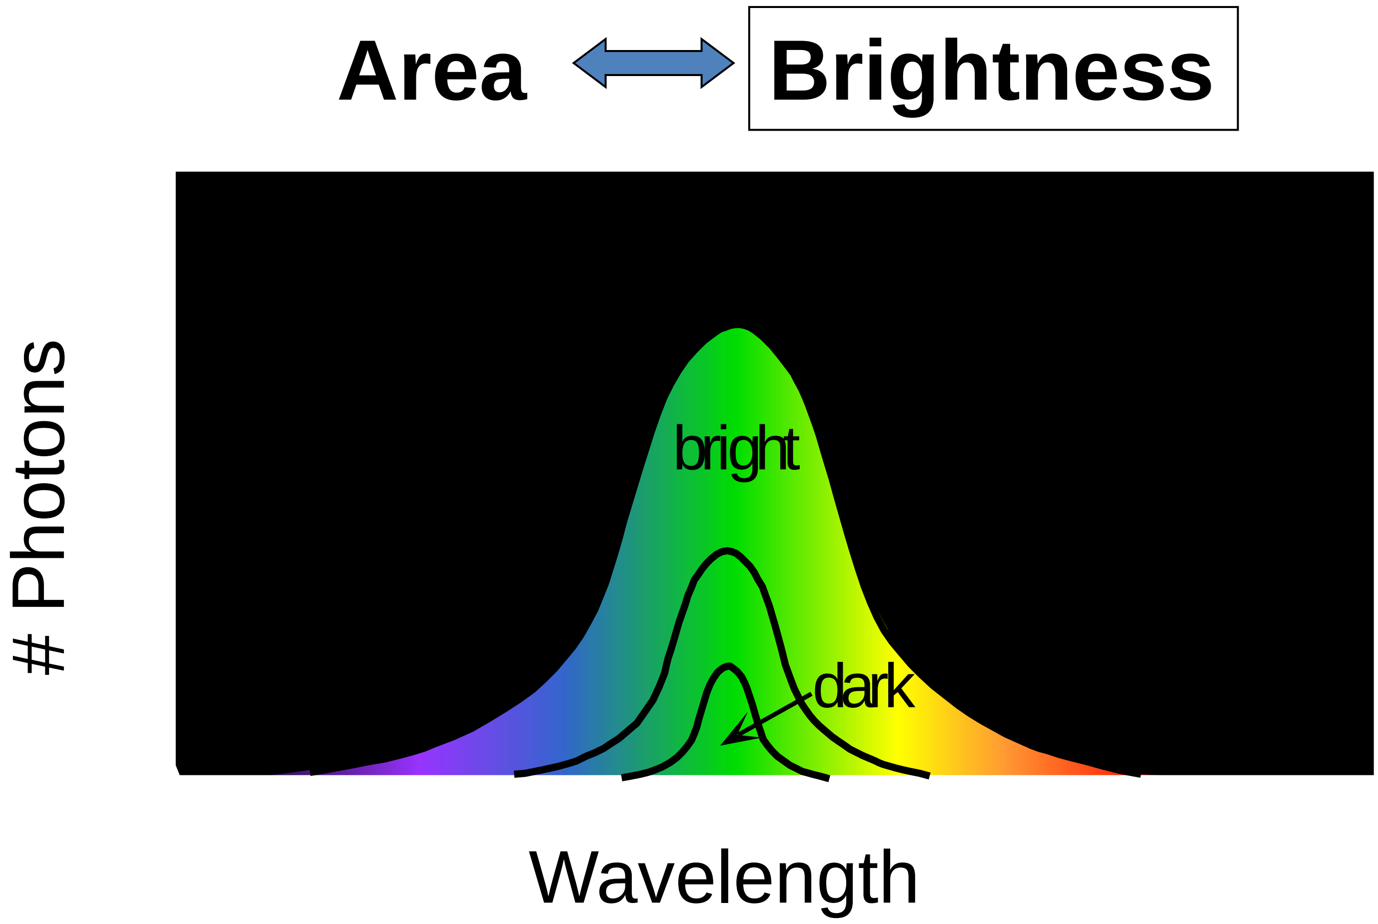

Physiology of Color Vision

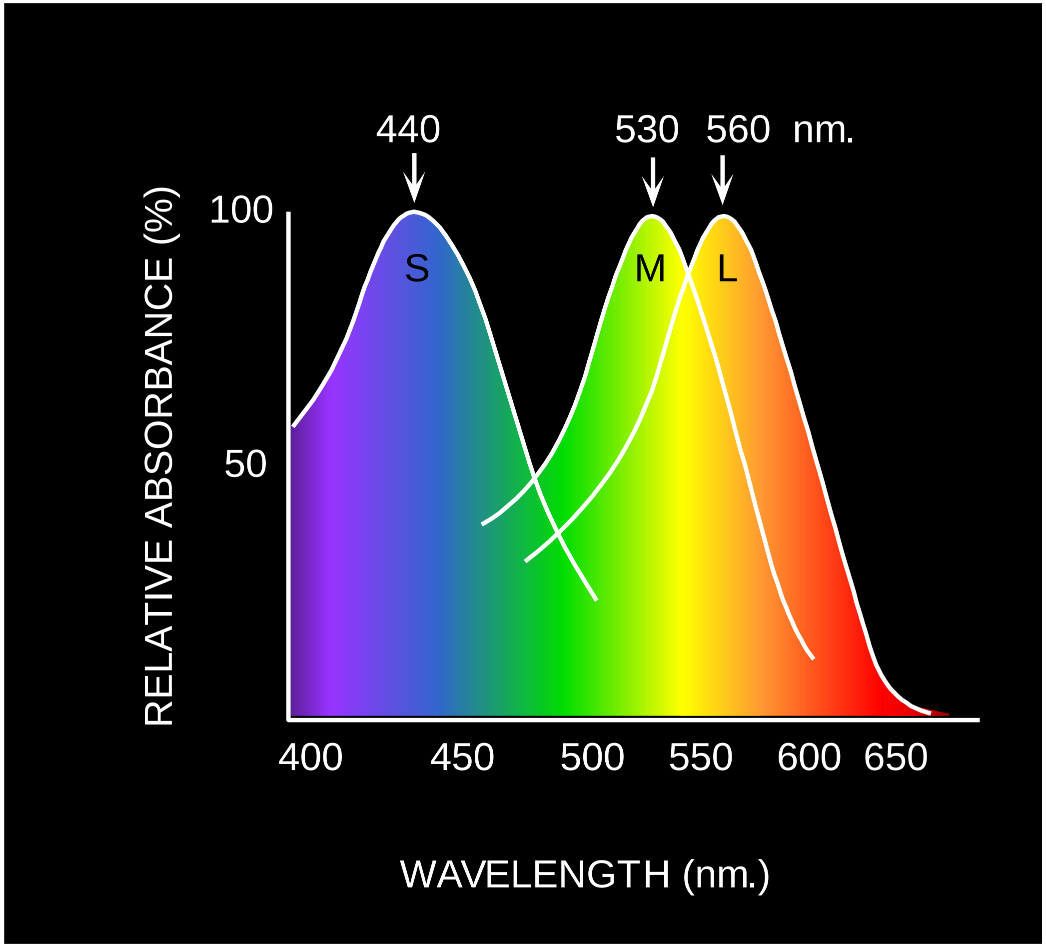

Tetrachromatism

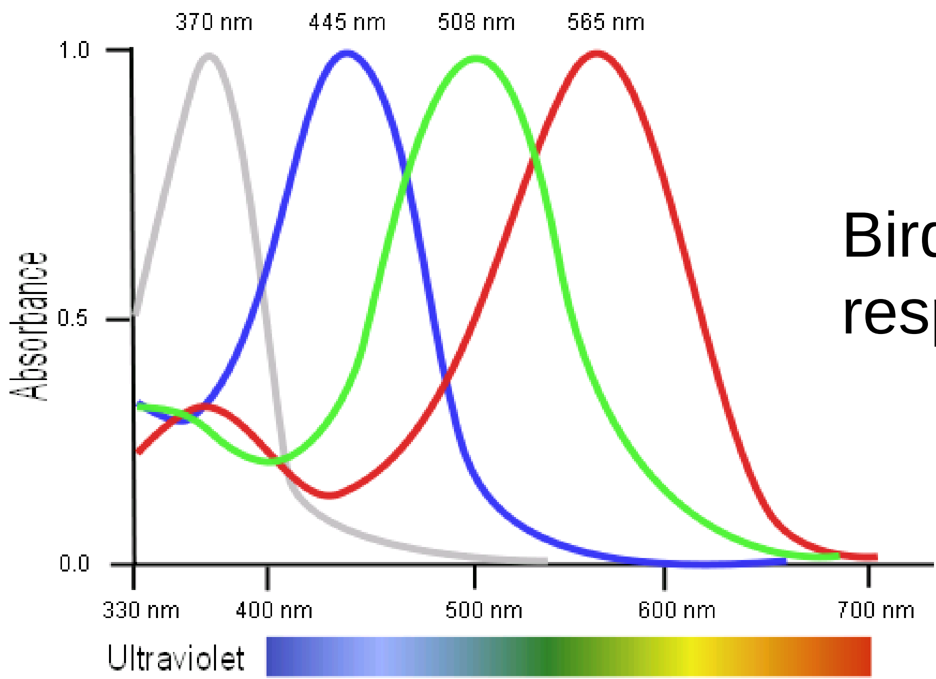

Bird cone responses

- Most birds, and many other animals, have cones for ultraviolet.
- Some humans, mostly female, have slight tetrachromatism.

Practical Color Sensing: Bayer Grid

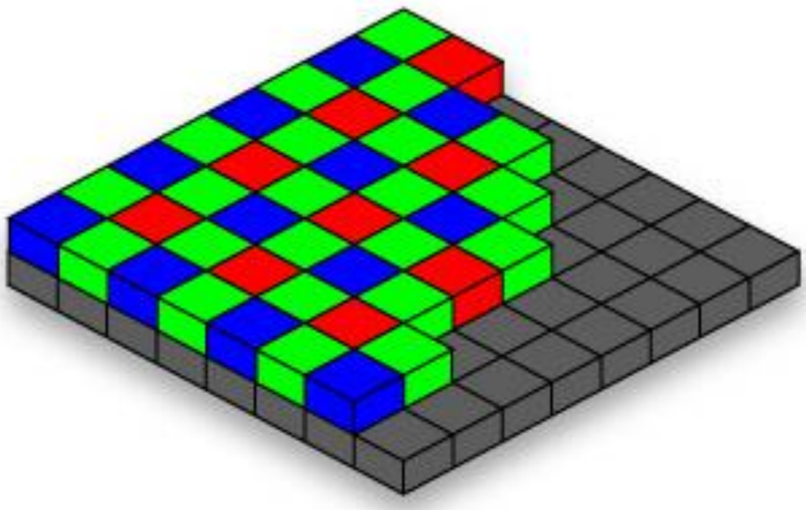

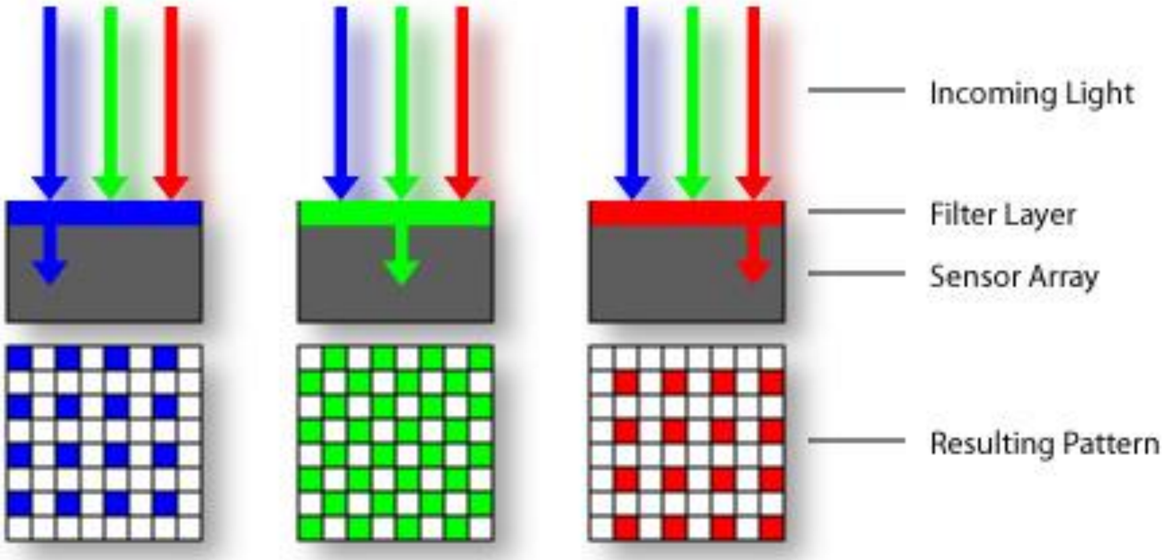

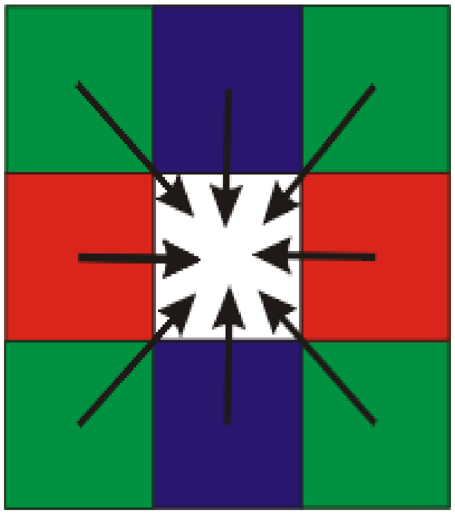

Estimate RGB at ‘G’ cells from neighboring values

Slide by Steve Seitz

Color Image

Images in Matlab

- 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 b^{th} channel
  - 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

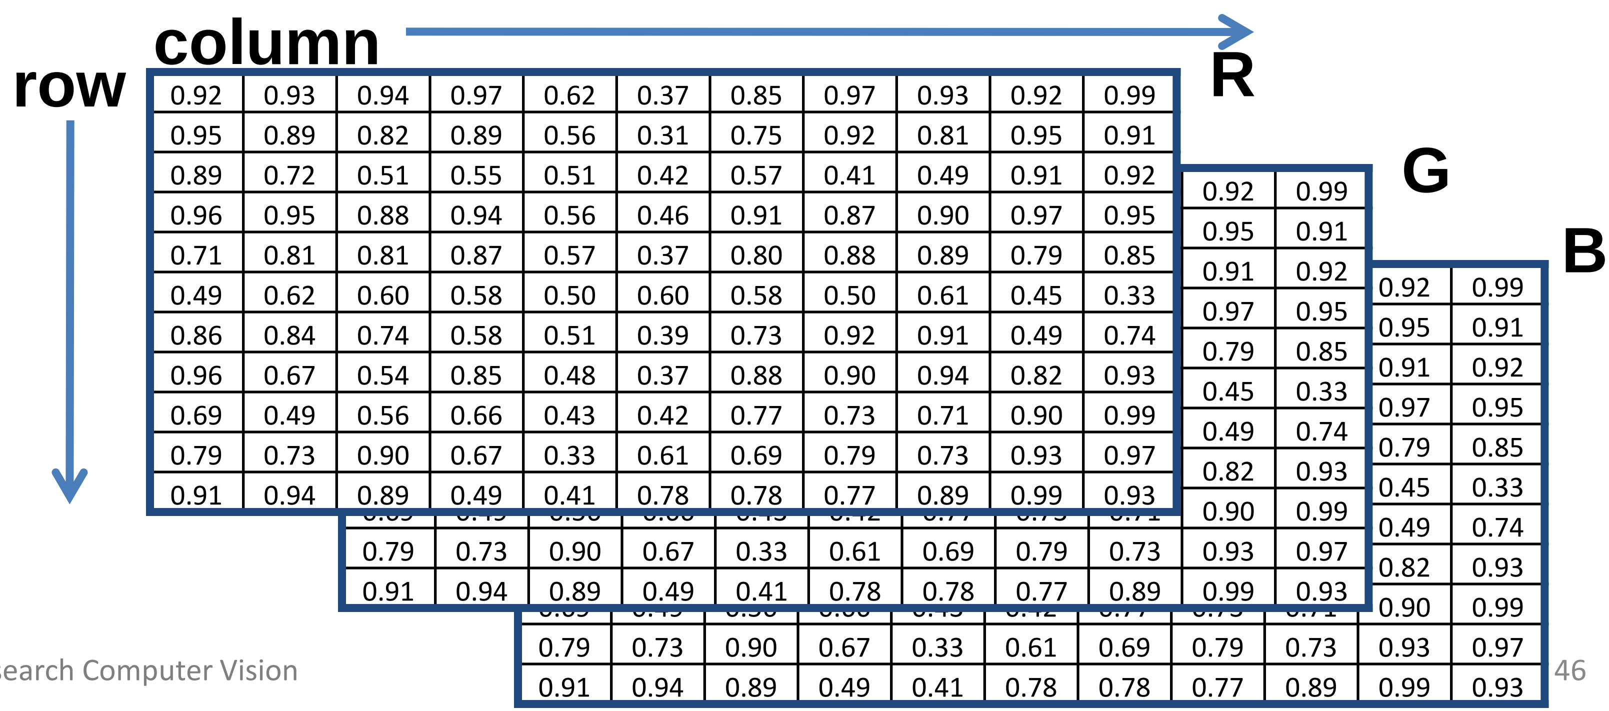

0.92	0.93	0.94	0.97	0.62	0.37	0.85	0.97	0.93	0.92	0.99
0.95	0.89	0.82	0.89	0.56	0.31	0.75	0.92	0.81	0.95	0.91
0.89	0.72	0.51	0.55	0.51	0.42	0.57	0.41	0.49	0.91	0.92
0.96	0.95	0.88	0.94	0.56	0.46	0.91	0.87	0.90	0.97	0.95
0.71	0.81	0.81	0.87	0.57	0.37	0.80	0.88	0.89	0.79	0.85
0.49	0.62	0.60	0.58	0.50	0.60	0.58	0.50	0.61	0.45	0.33
0.86	0.84	0.74	0.58	0.51	0.39	0.73	0.92	0.91	0.49	0.74
0.96	0.67	0.54	0.85	0.48	0.37	0.88	0.90	0.94	0.82	0.93
0.69	0.49	0.56	0.66	0.43	0.42	0.77	0.73	0.71	0.90	0.99
0.79	0.73	0.90	0.67	0.33	0.61	0.69	0.79	0.73	0.93	0.97
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?

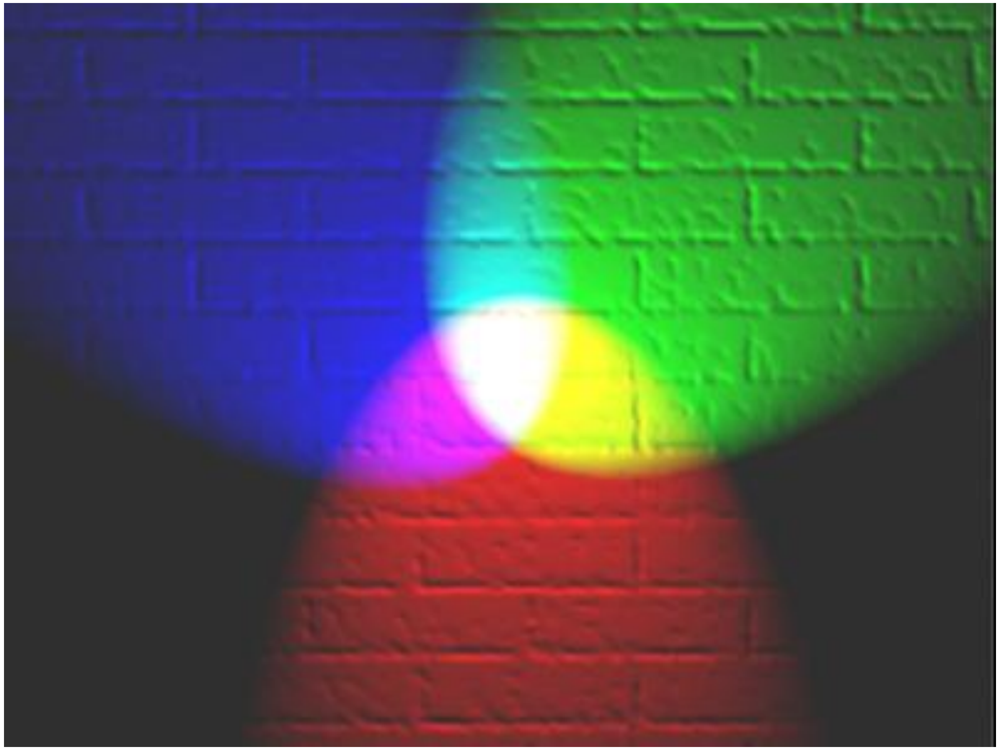

http://en.wikipedia.org/wiki/File:RGB_illumination.jpg

Color Spaces: RGB

Default color space

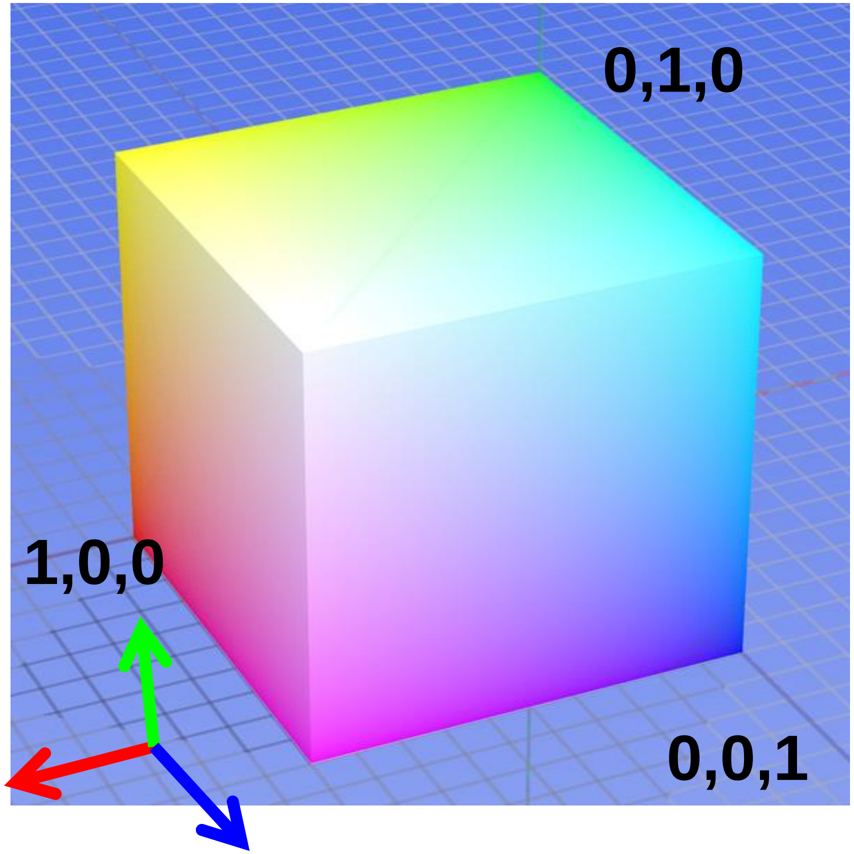

R

(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

Color spaces: HSV

Intuitive color space

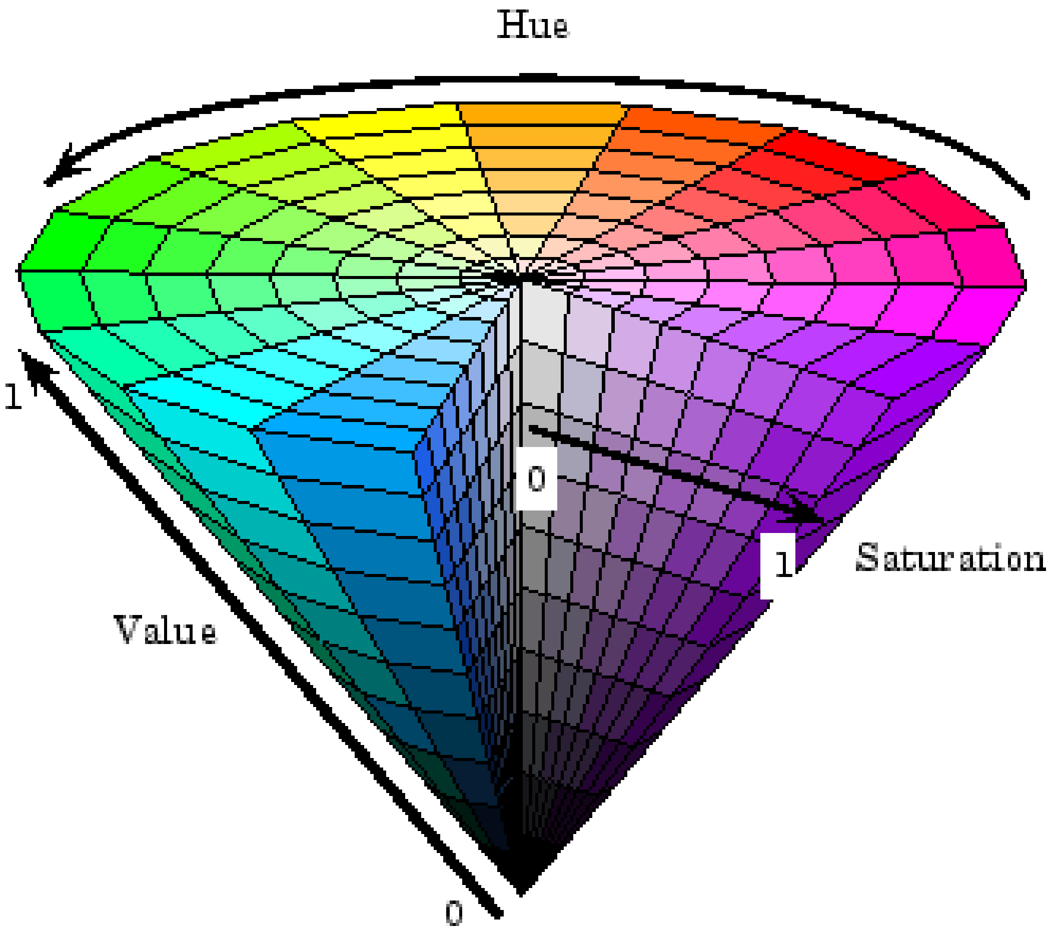

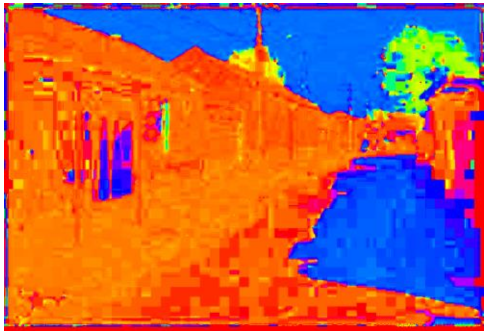

H
(S=1,V=1)

S
(H=1,V=1)

V
(H=1,S=0)

Color spaces: YCbCr

Fast to compute,
good for compression, used by TV

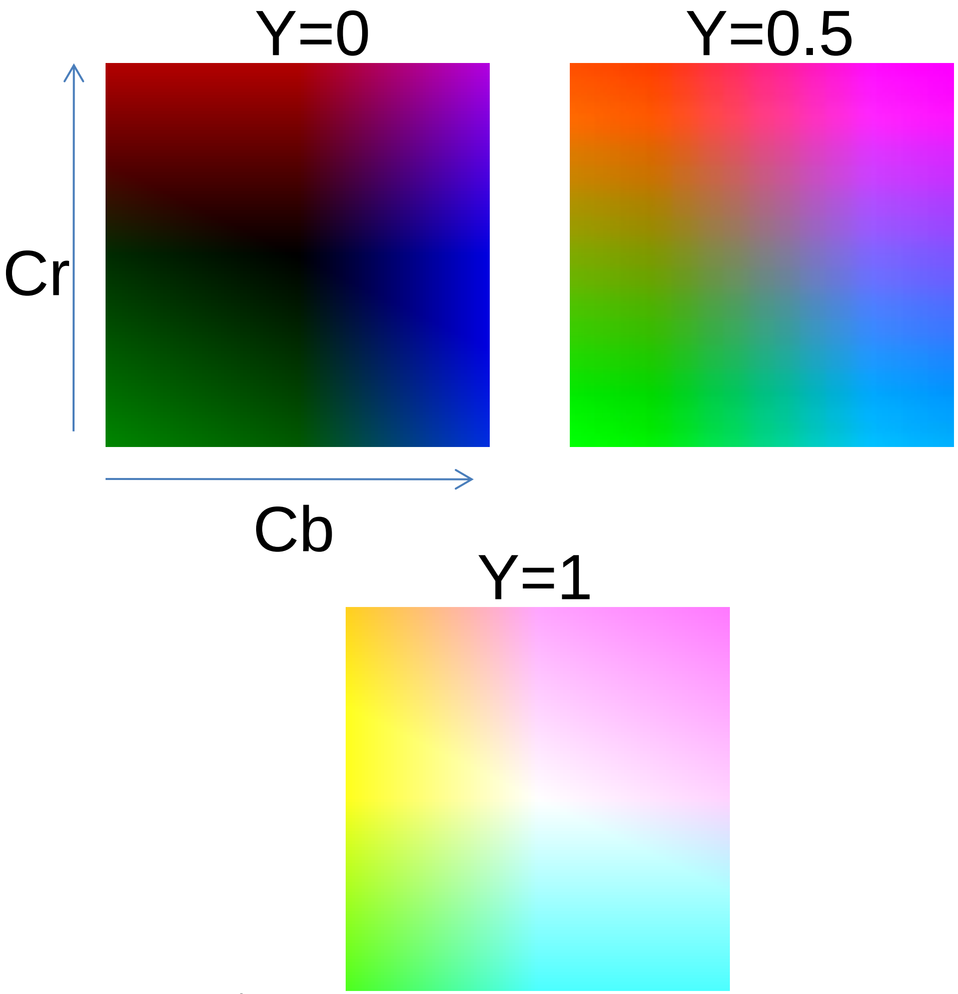

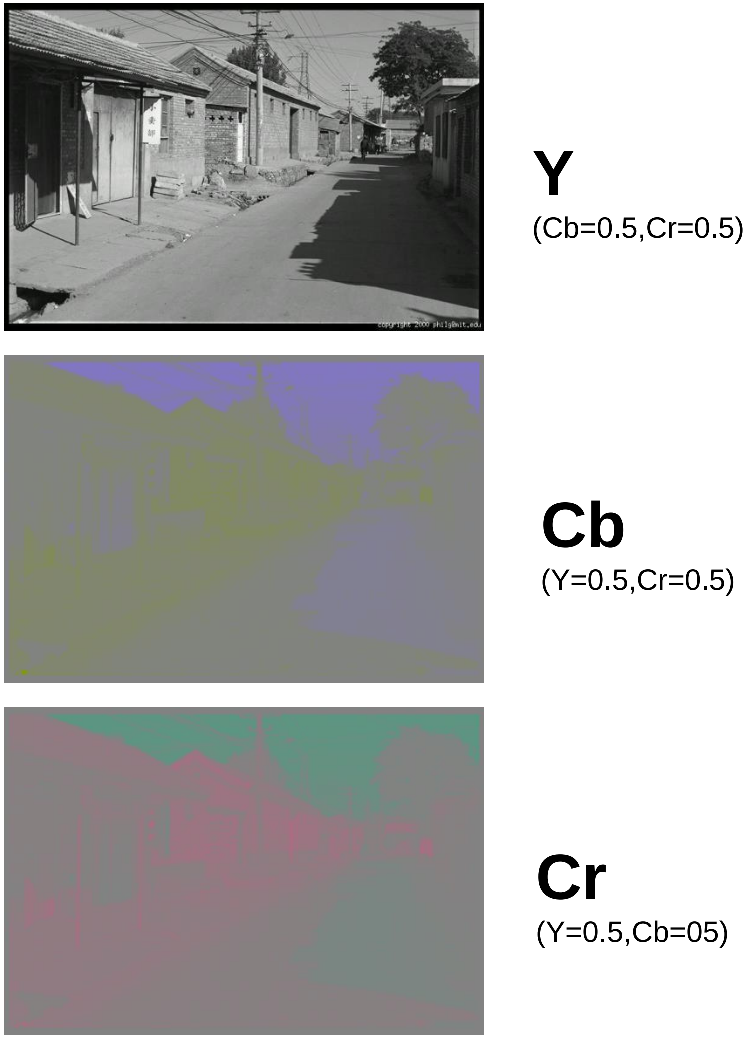

Color spaces: L*a*b*

"Perceptually uniform"* color space

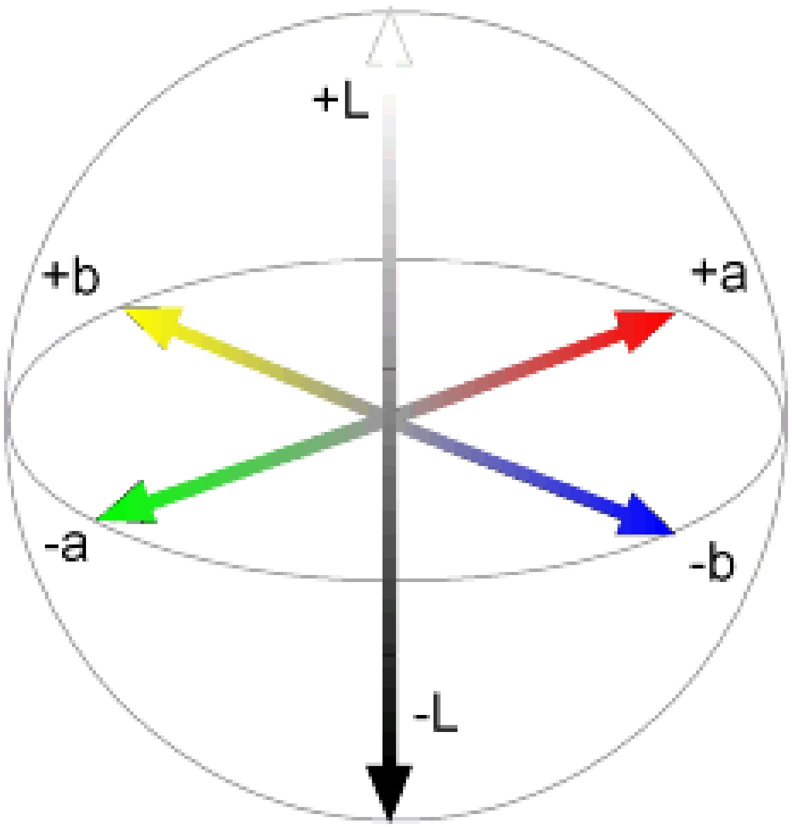

L
(a=0,b=0)

a
(L=65,b=0)

b
(L=65,a=0)

Luminance or Chrominance?

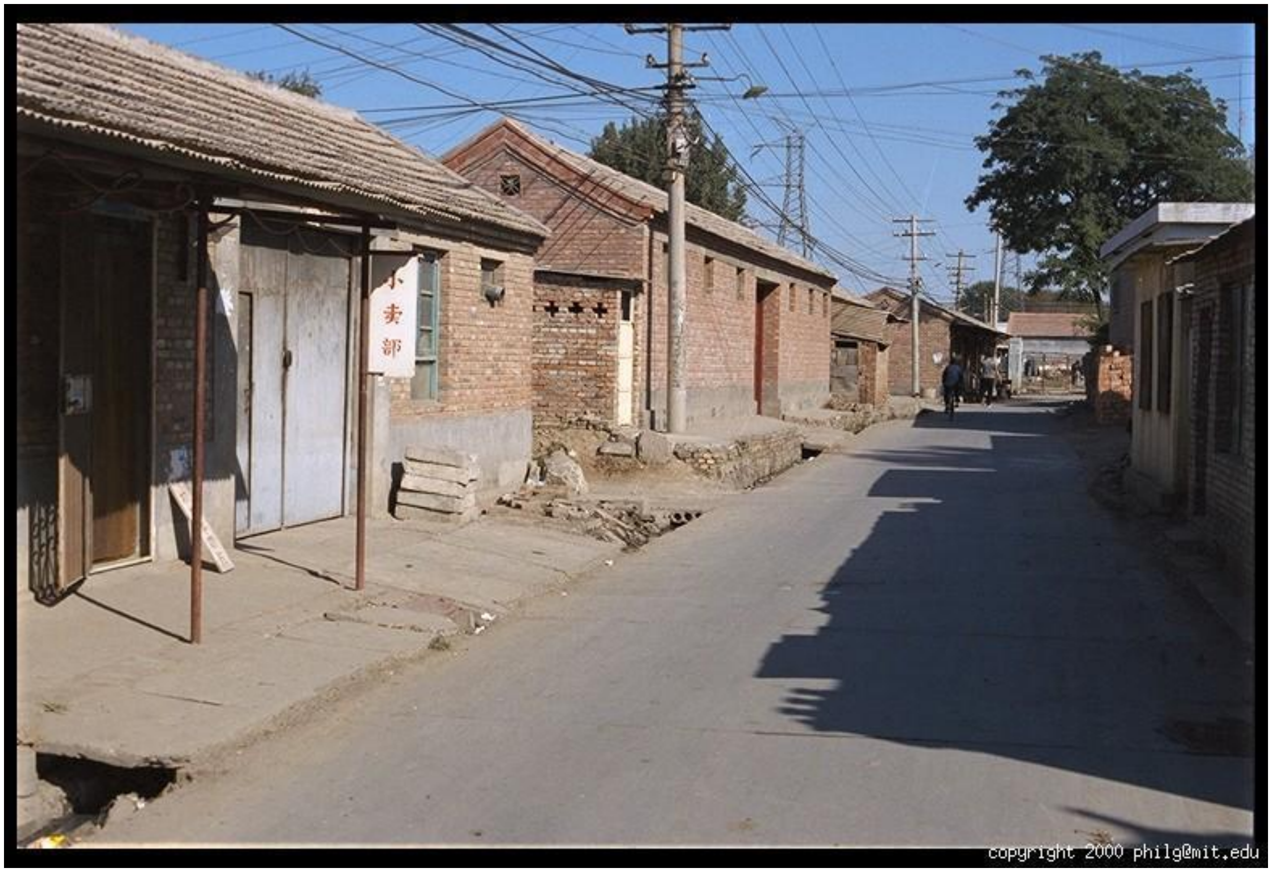

Only Color (Chrominance)

Only Intensity (Luminance)

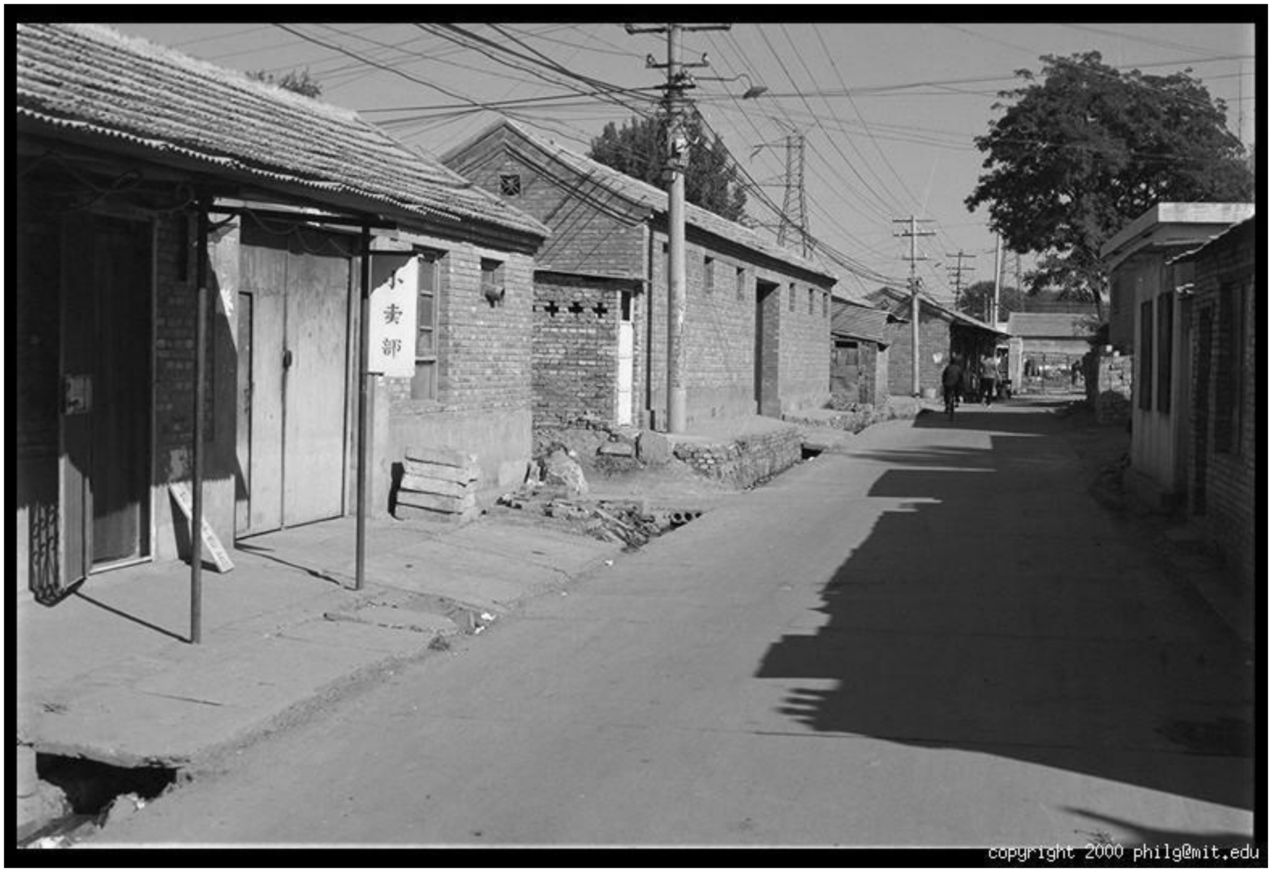

Back to Grayscale

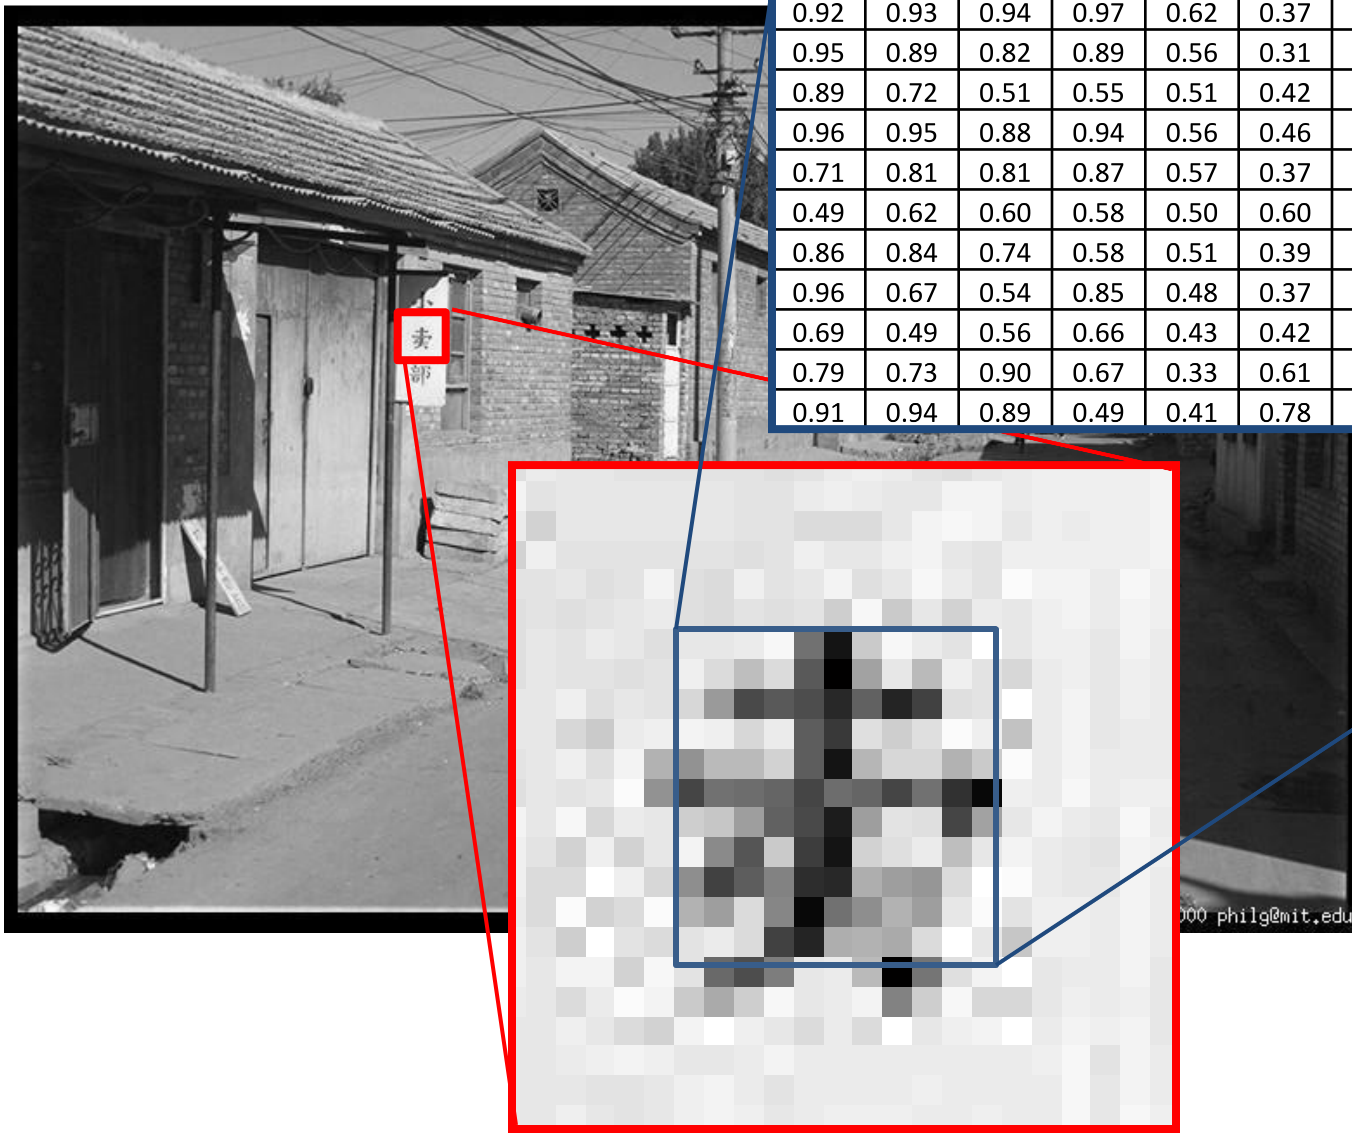

0.92	0.93	0.94	0.97	0.62	0.37	0.85	0.97	0.93	0.92	0.99
0.95	0.89	0.82	0.89	0.56	0.31	0.75	0.92	0.81	0.95	0.91
0.89	0.72	0.51	0.55	0.51	0.42	0.57	0.41	0.49	0.91	0.92
0.96	0.95	0.88	0.94	0.56	0.46	0.91	0.87	0.90	0.97	0.95
0.71	0.81	0.81	0.87	0.57	0.37	0.80	0.88	0.89	0.79	0.85
0.49	0.62	0.60	0.58	0.50	0.60	0.58	0.50	0.61	0.45	0.33
0.86	0.84	0.74	0.58	0.51	0.39	0.73	0.92	0.91	0.49	0.74
0.96	0.67	0.54	0.85	0.48	0.37	0.88	0.90	0.94	0.82	0.93
0.69	0.49	0.56	0.66	0.43	0.42	0.77	0.73	0.71	0.90	0.99
0.79	0.73	0.90	0.67	0.33	0.61	0.69	0.79	0.73	0.93	0.97
0.91	0.94	0.89	0.49	0.41	0.78	0.78	0.77	0.89	0.99	0.93

Course Outline

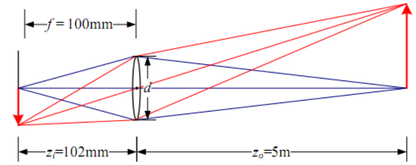

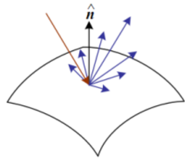

Image Formation and Processing

Light, Shape and Color

The Pin-hole Camera Model, The Digital Camera

Linear filtering, Filter banks, Multiresolution

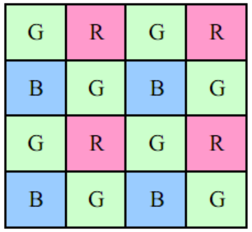

Feature Detection and Matching

Edge Detection, Interest Points: Corners and Blobs

Local Image Descriptors

Feature Matching and Hough Transform

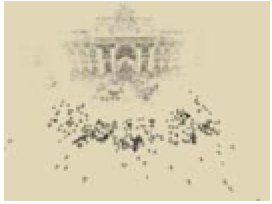

Multiple Views and Motion

Geometric Transformations, Camera Calibration

Feature Tracking , Stereo Vision

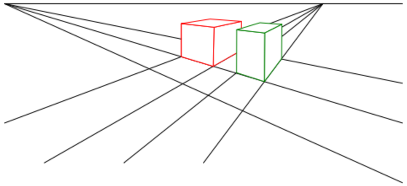

Segmentation and Grouping

Segmentation by Clustering, Region Merging and Growing

Advanced Methods Overview: Active Contours, Level-Sets, Graph-Theoretic Methods

Detection and Recognition

Problems and Architectures Overview

Statistical Classifiers, Bag-of-Words Model, Detection by Sliding Windows

Resources

Books

R. Szeliski, Computer Vision: Algorithms and Applications, 2010 – *available online*

D. A. Forsyth and J. Ponce, Computer Vision: A Modern Approach, 2003

L. G. Shapiro and G. C. Stockman, Computer Vision, 2001

Web

CVonline: The Evolving, Distributed, Non-Proprietary, On-Line Compendium of Computer Vision

http://homepages.inf.ed.ac.uk/rbf/CVonline/

Dictionary of Computer Vision and Image Processing

http://homepages.inf.ed.ac.uk/rbf/CVDICT/

Computer Vision Online

http://www.computervisiononline.com/

Programming

Development environments/languages: Matlab, Python and C/C++

Toolboxes and APIs: OpenCV, VLFeat Matlab Toolbox, Piotr's Computer Vision Matlab Toolbox, EasyCamCalib Software, FLANN, Point Cloud Library PCL, LibSVM, Camera Calibration Toolbox for Matlab