CS6670: Computer Vision

Noah Snavely

Lecture 5: Cameras and Projection

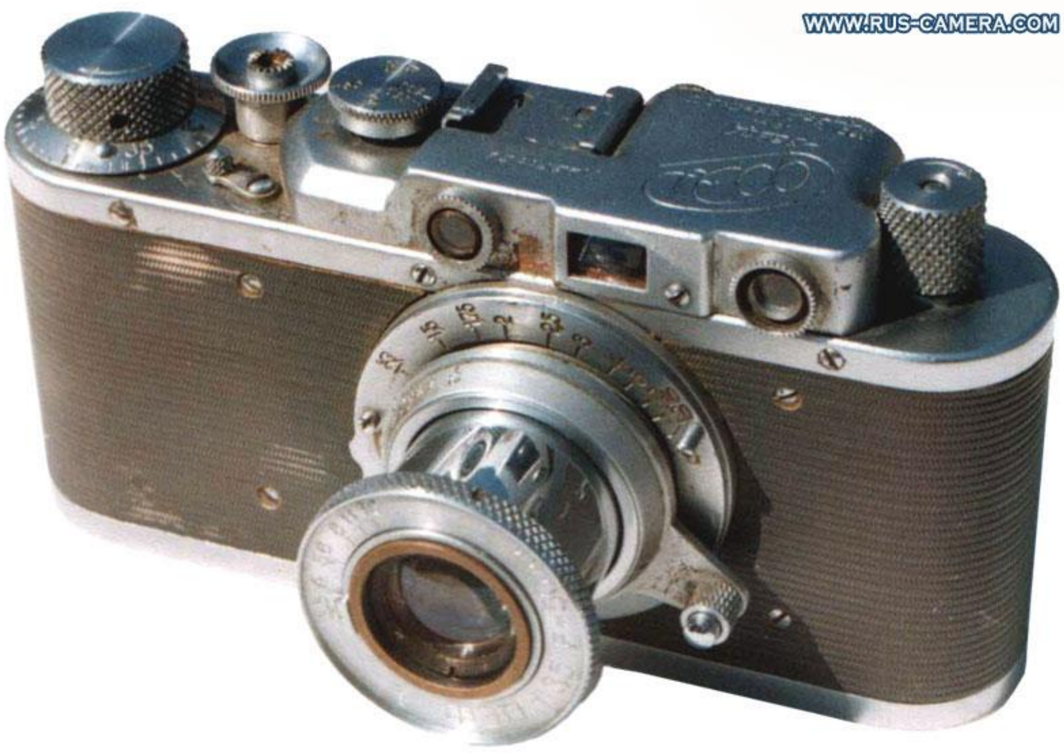

Reading

- Szeliski 2.1.3-2.1.6

Announcements

- Project 1 assigned, see projects page:
  - http://www.cs.cornell.edu/courses/cs6670/2011sp/projects/projects.html

- Quiz 1 on Wednesday

Image formation

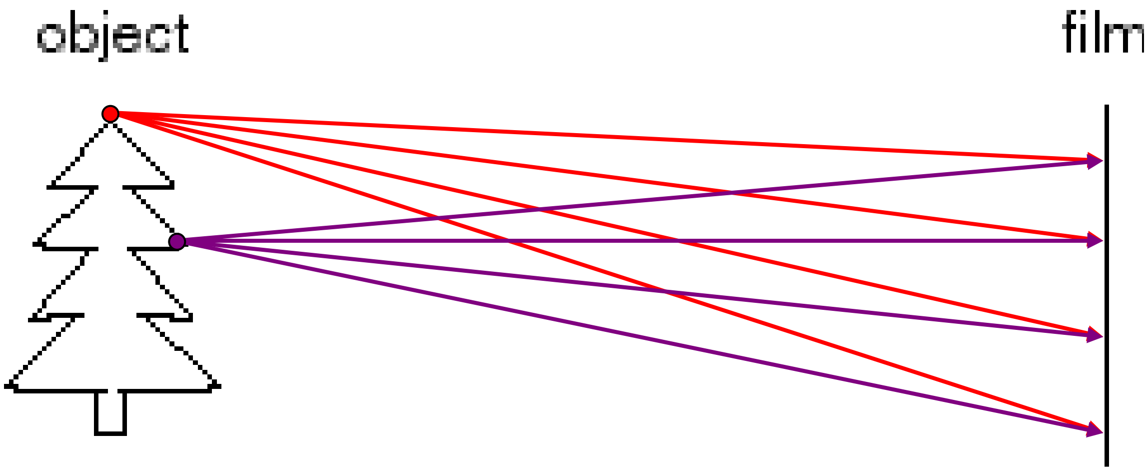

- Let's design a camera
  - Idea 1: put a piece of film in front of an object
  - Do we get a reasonable image?

Pinhole camera

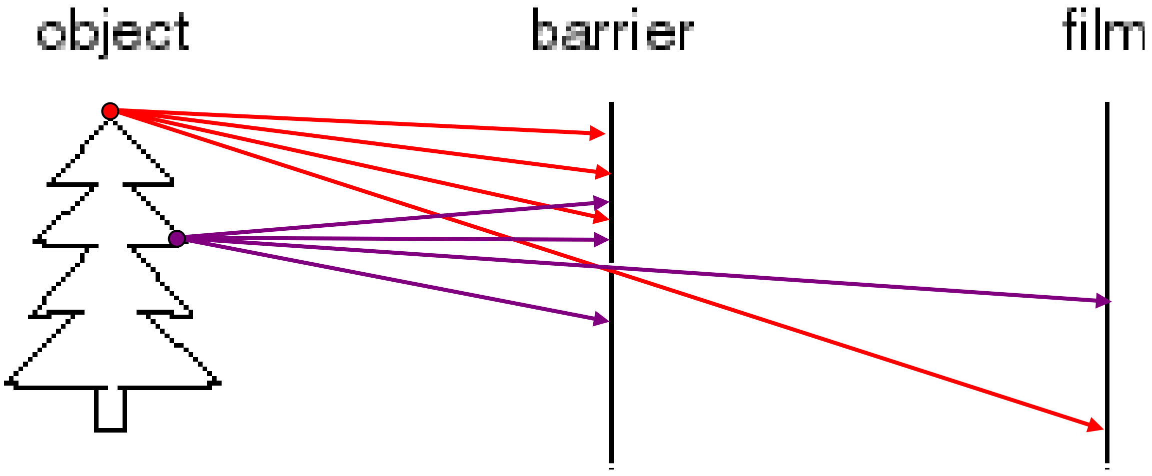

- Add a barrier to block off most of the rays
  - This reduces blurring
  - The opening known as the **aperture**
  - How does this transform the image?

Camera Obscura

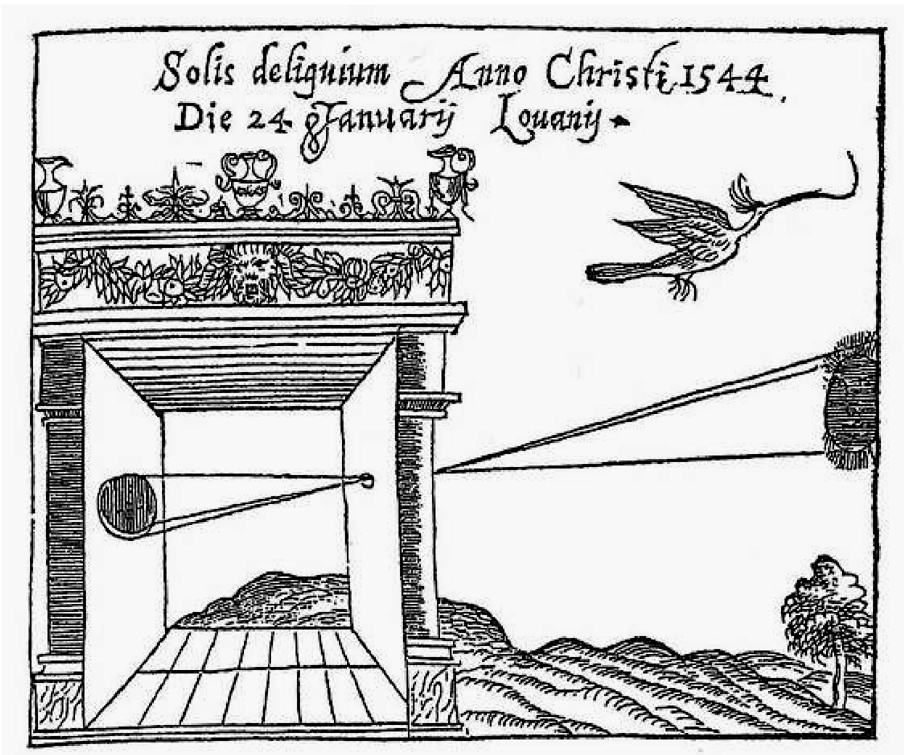

Gemma Frisius, 1558

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

Source: A. Efros

Camera Obscura

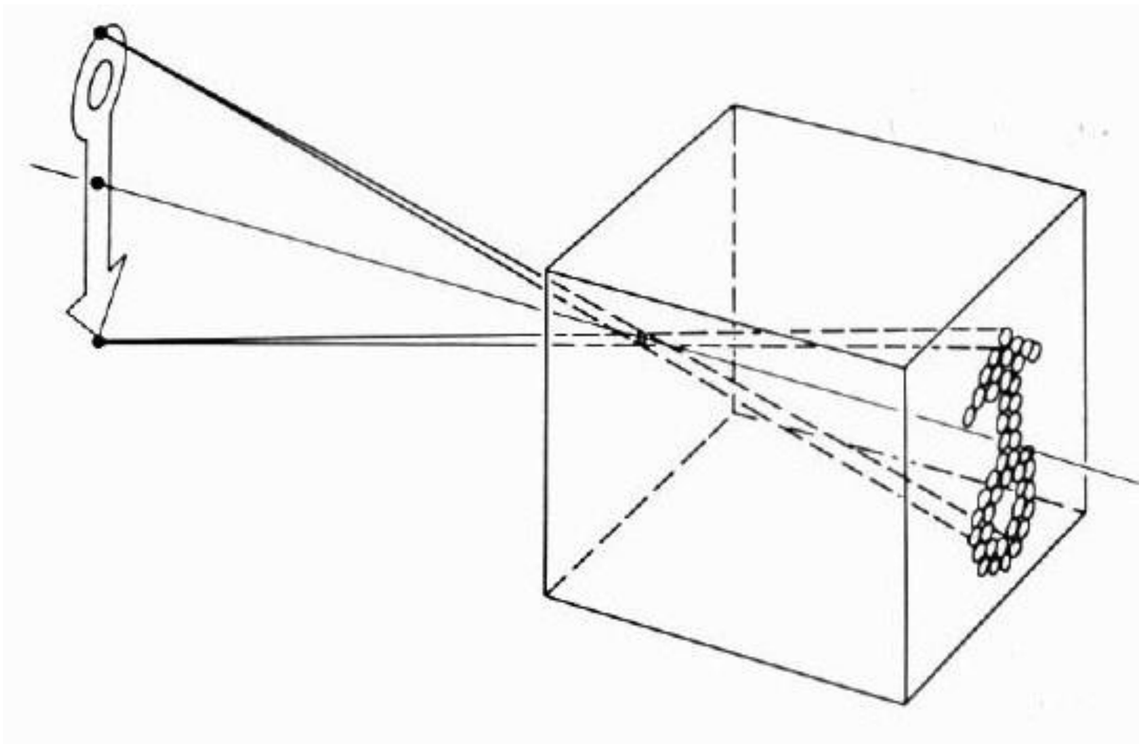

Home-made pinhole camera

Why so blurry?

http://www.debevec.org/Pinhole/

Slide by A. Efros

Pinhole photography

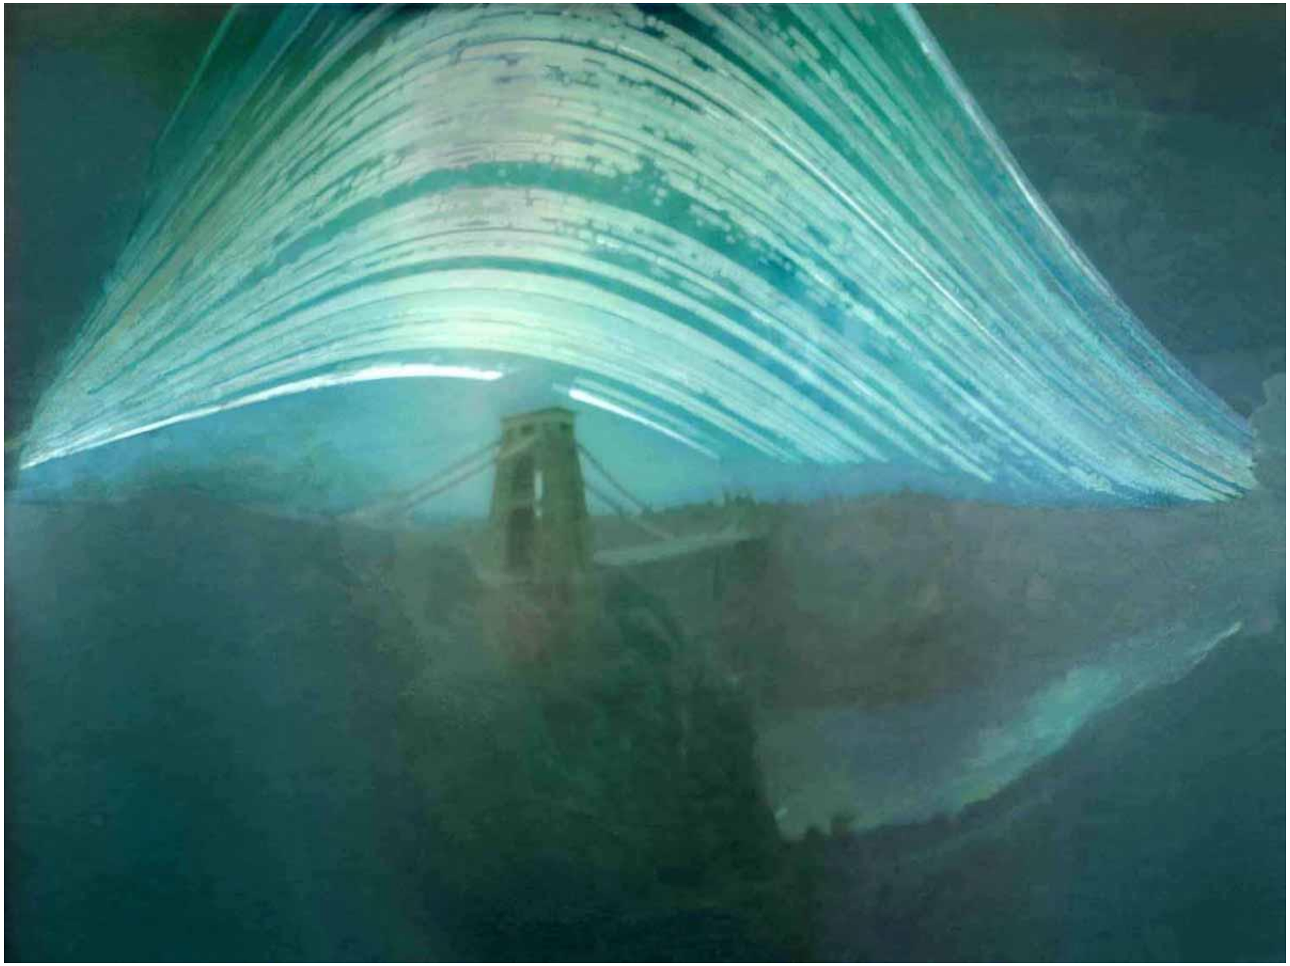

Justin Quinnell, The Clifton Suspension Bridge. December 17th 2007 - June 21st 2008
***6-month* exposure**

Shrinking the aperture

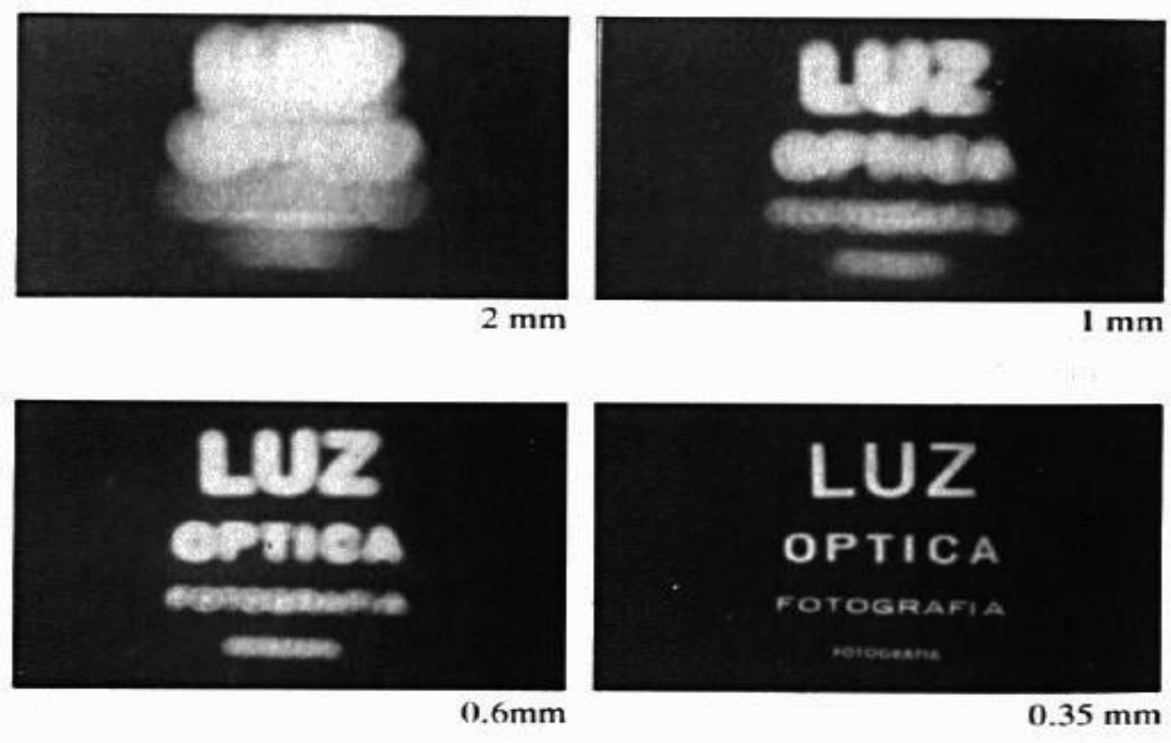

- Why not make the aperture as small as possible?
  - Less light gets through
  - *Diffraction* effects...

Shrinking the aperture

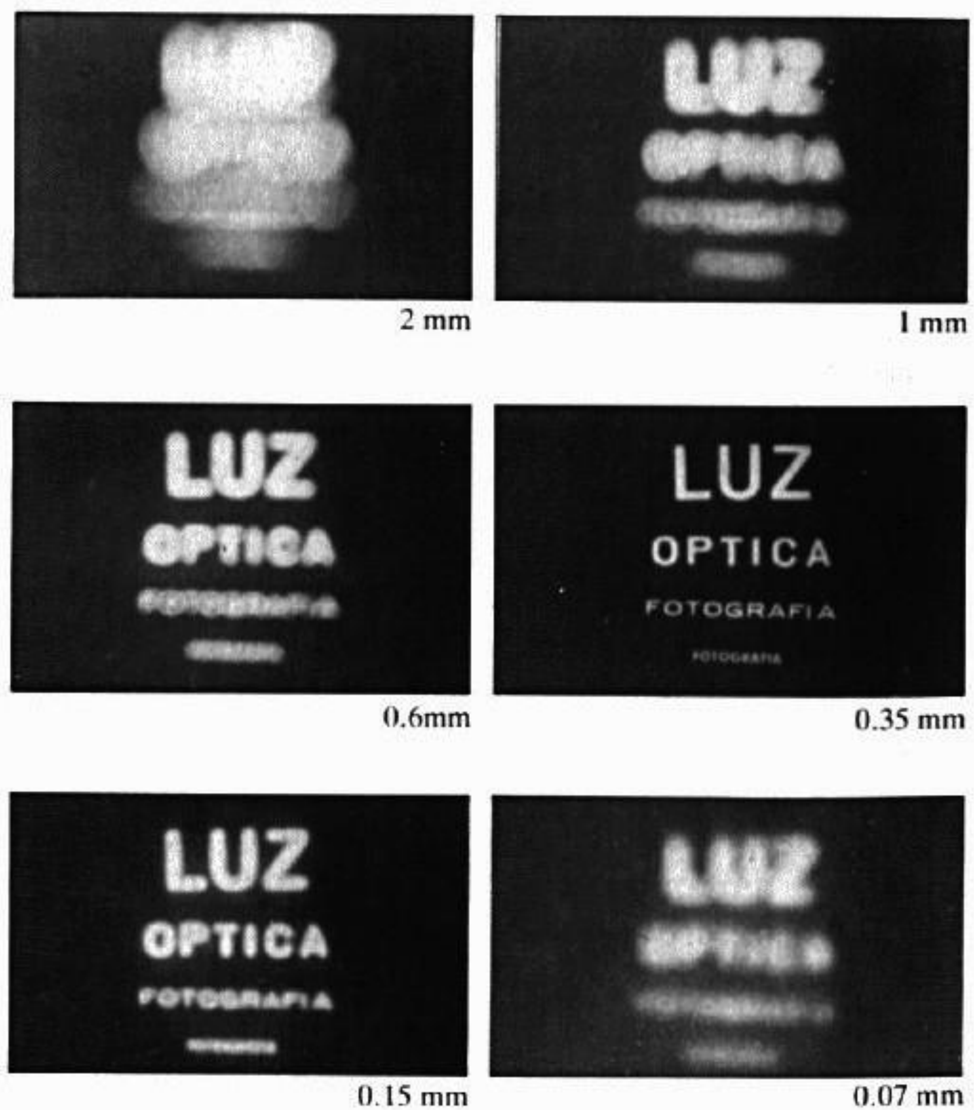

2 mm　1 mm

0.6mm　0.35 mm

0.15 mm　0.07 mm

Adding a lens

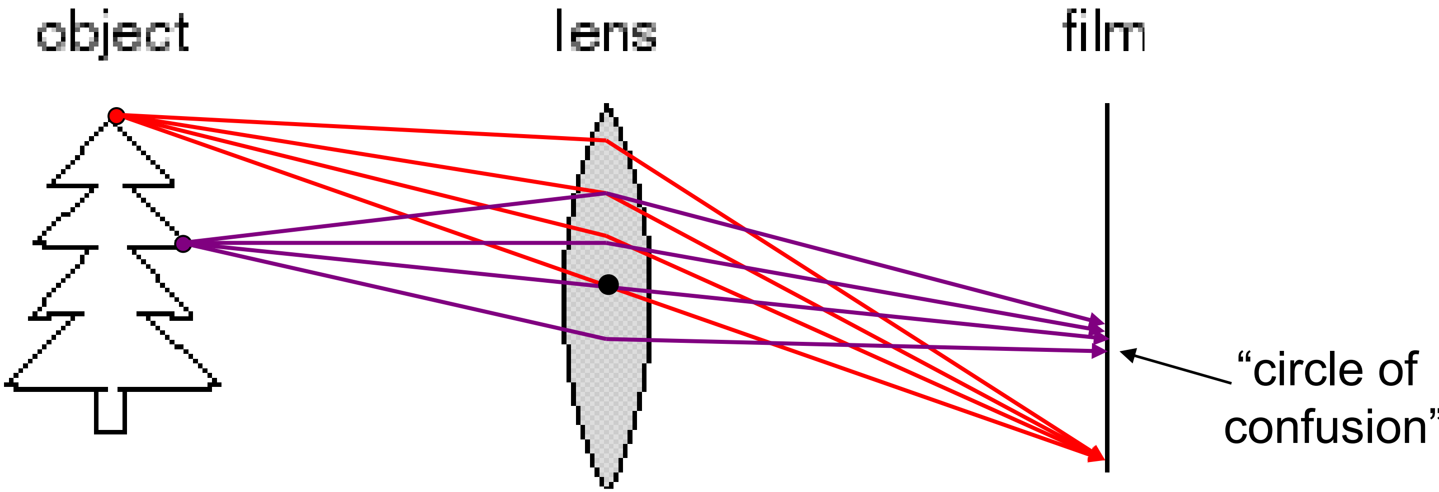

- A lens focuses light onto the film
  - There is a specific distance at which objects are "in focus"
    - other points project to a "circle of confusion" in the image
  - Changing the shape of the lens changes this distance

Lenses

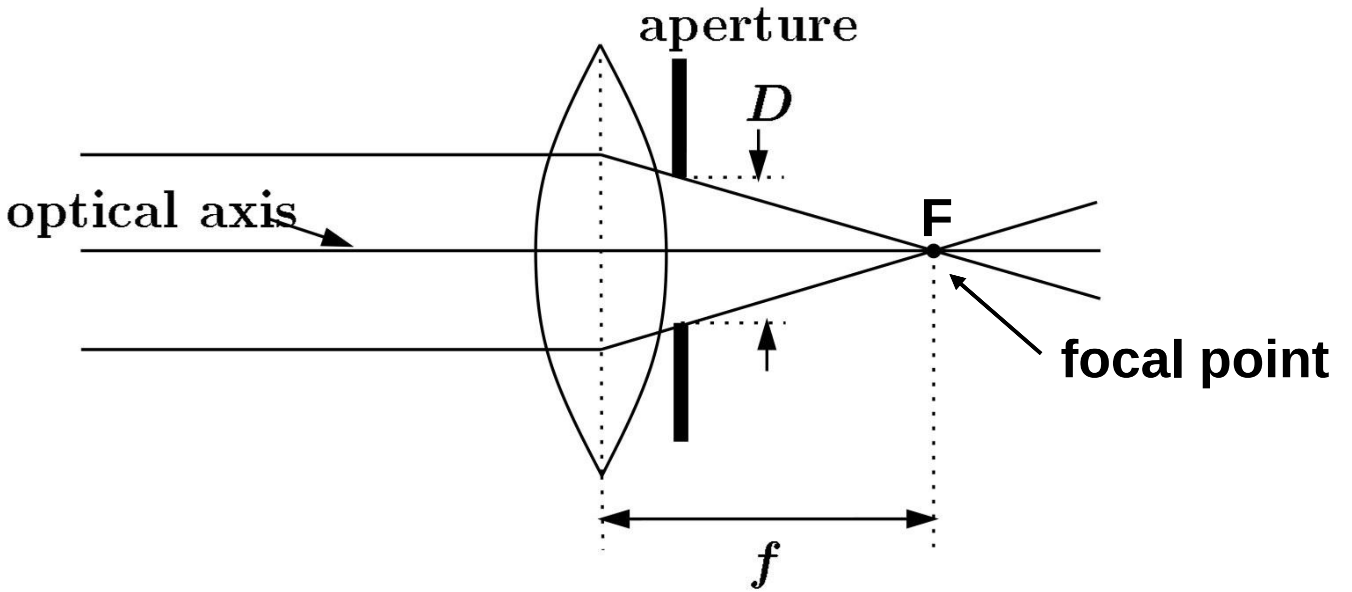

- A lens focuses parallel rays onto a single focal point
  - focal point at a distance *f* beyond the plane of the lens (the *focal length*)
    - *f* is a function of the shape and index of refraction of the lens
  - Aperture restricts the range of rays
    - aperture may be on either side of the lens
  - Lenses are typically spherical (easier to produce)

Thin lenses

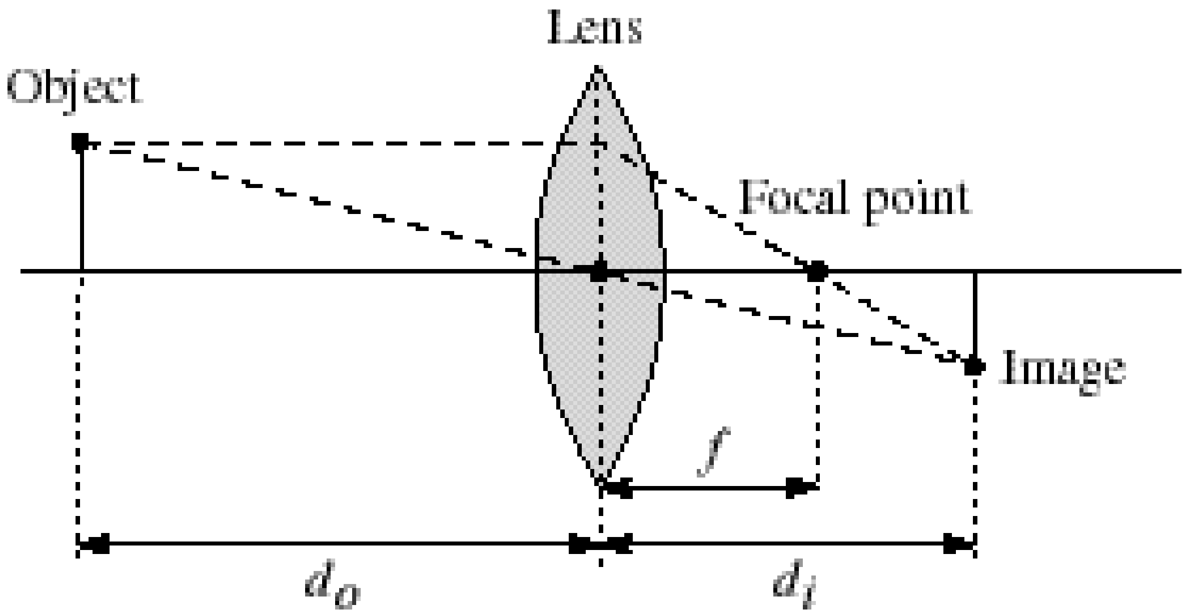

- Thin lens equation: $$\frac{1}{d_o} + \frac{1}{d_i} = \frac{1}{f}$$
  - Any object point satisfying this equation is in focus
  - What is the shape of the focus region?
  - How can we change the focus region?
  - Thin lens applet: http://www.phy.ntnu.edu.tw/java/Lens/lens_e.html (by Fu-Kwun Hwang)

Depth of Field

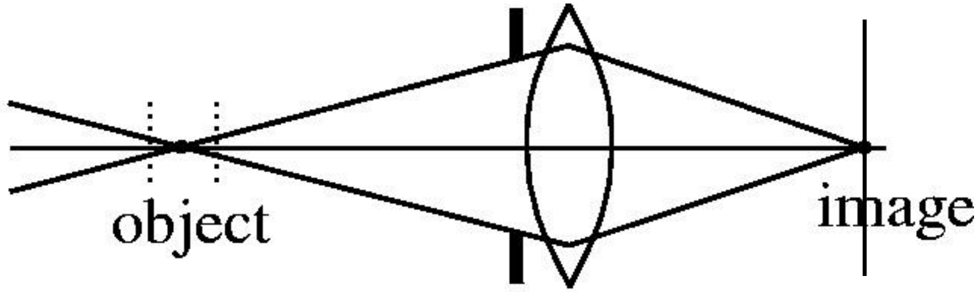

f / 5.6

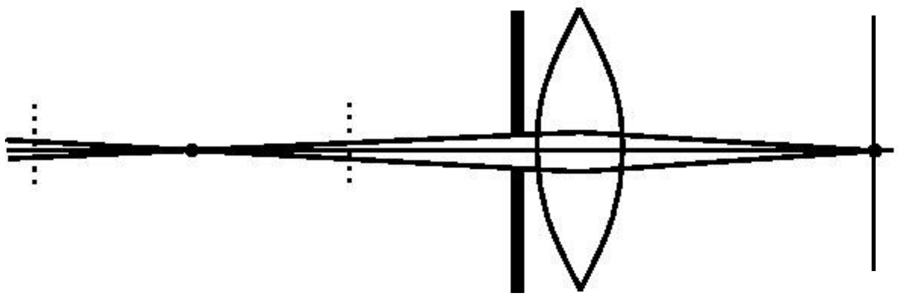

f / 32

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

Flower images from Wikipedia http://en.wikipedia.org/wiki/Depth_of_field

Depth of Field

Large aperture opening

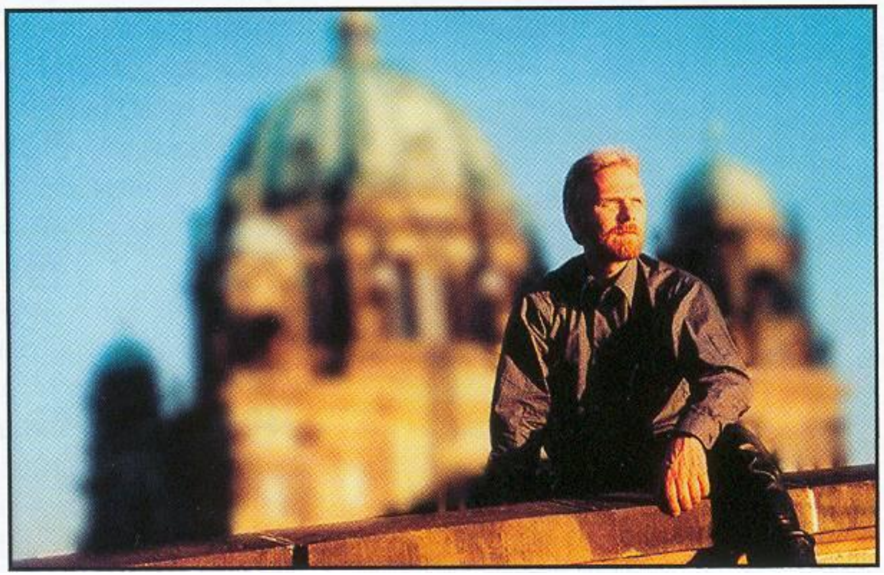

Small aperture opening

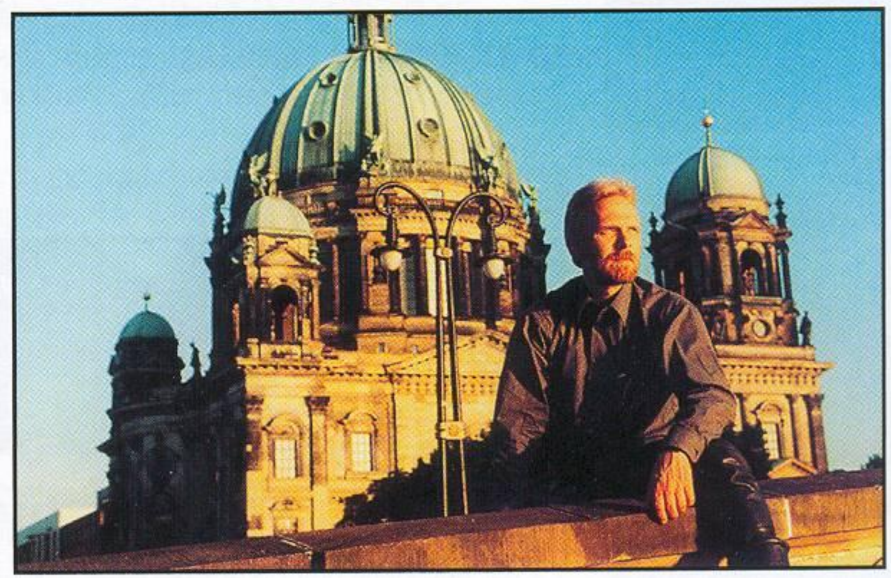

The eye

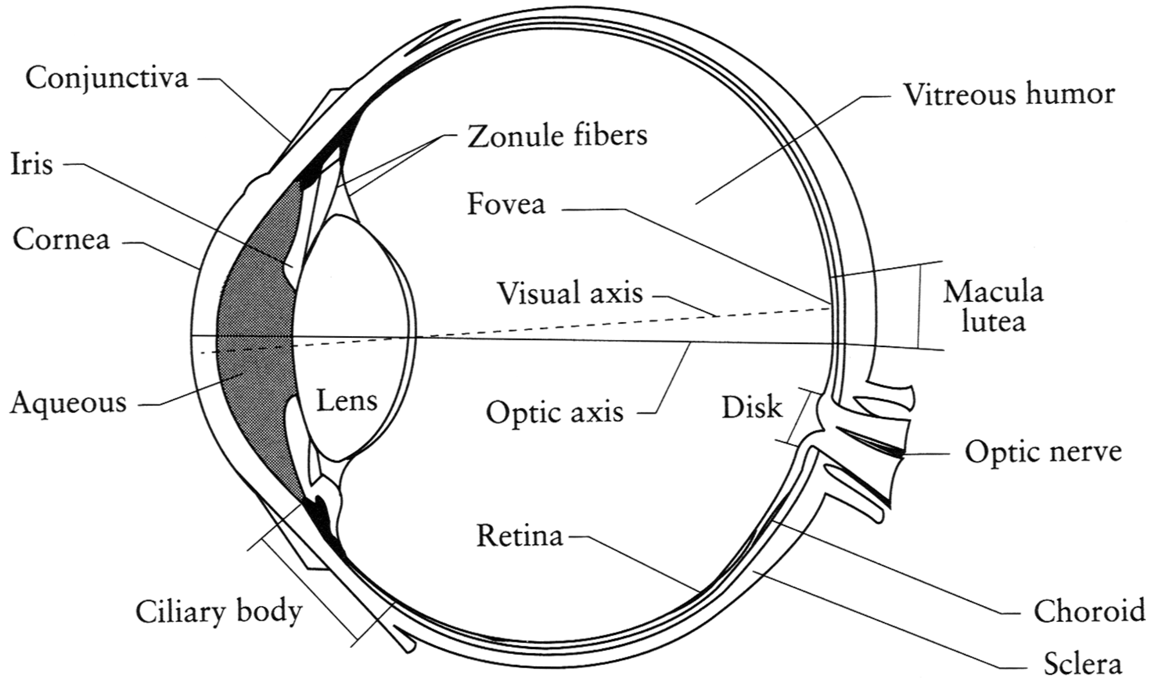

- The human eye is a camera
  - **Iris** - colored annulus with radial muscles
  - **Pupil** - the hole (aperture) whose size is controlled by the iris
  - What's the "film"?
    - photoreceptor cells (rods and cones) in the **retina**

Before Film

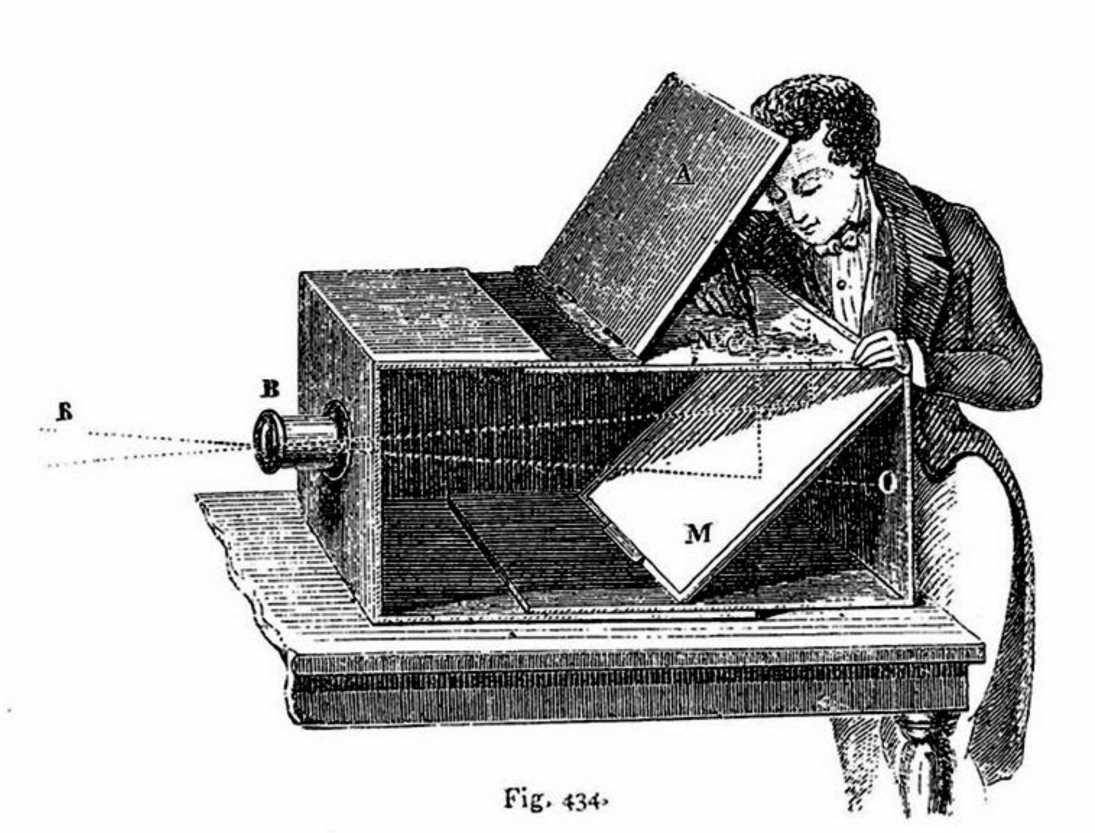

Lens Based Camera Obscura, 1568

Srinivasa Narasimhan's slide

Film camera

Still Life, Louis Jaques Mande Daguerre, 1837

Srinivasa Narasimhan's slide

Silicon Image Detector

Silicon Image Detector, 1970

Digital camera

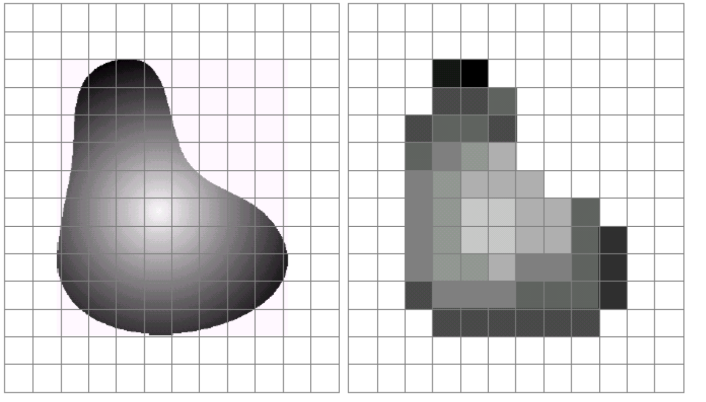

a b

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

- A digital camera replaces film with a sensor array
  - Each cell in the array is a **C**harge **C**oupled **D**evice
    - light-sensitive diode that converts photons to electrons
    - other variants exist: CMOS is becoming more popular
    - http://electronics.howstuffworks.com/digital-camera.htm

Color

- So far, we've talked about grayscale images
- What about color?
- Most digital images are comprised of three color channels – red, green, and, blue – which combine to create most of the colors we can see

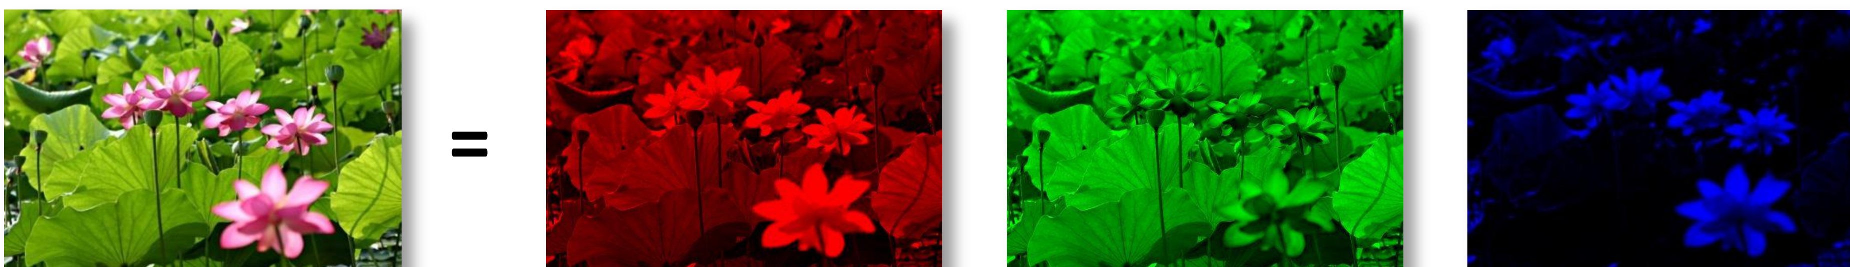

- Why are there three?

Color perception

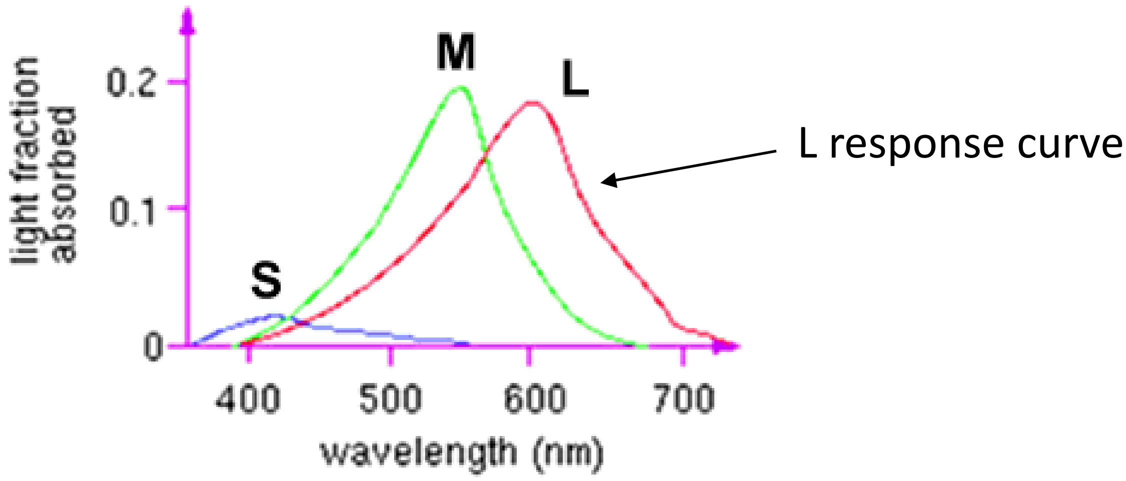

- Three types of cones
  - Each is sensitive in a different region of the spectrum
    - but regions overlap
    - Short (S) corresponds to blue
    - Medium (M) corresponds to green
    - Long (L) corresponds to red
  - Different sensitivities: we are more sensitive to green than red
    - varies from person to person (and with age)
  - Colorblindness—deficiency in at least one type of cone

Field sequential

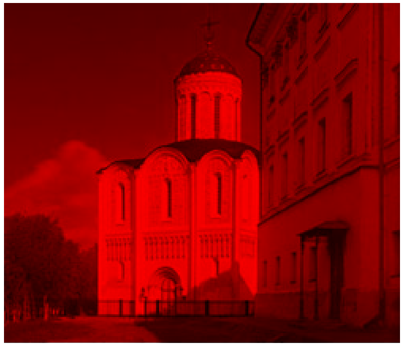

YungYu Chuang's slide

Field sequential

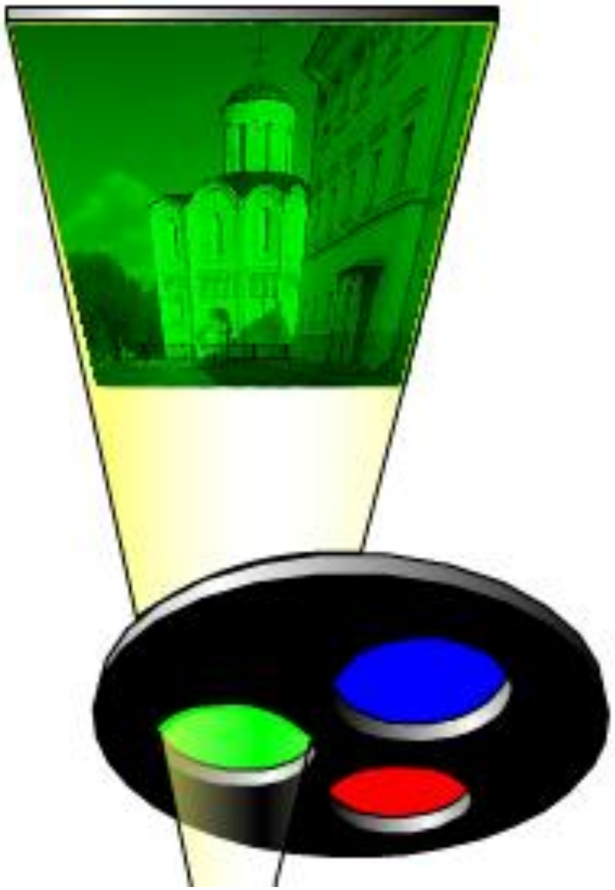

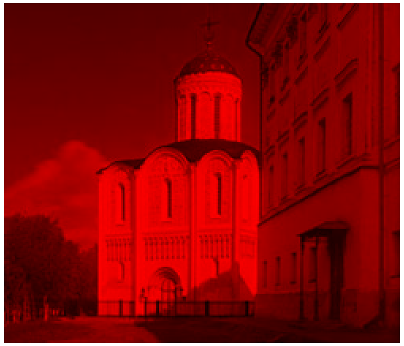

YungYu Chuang's slide

Field sequential

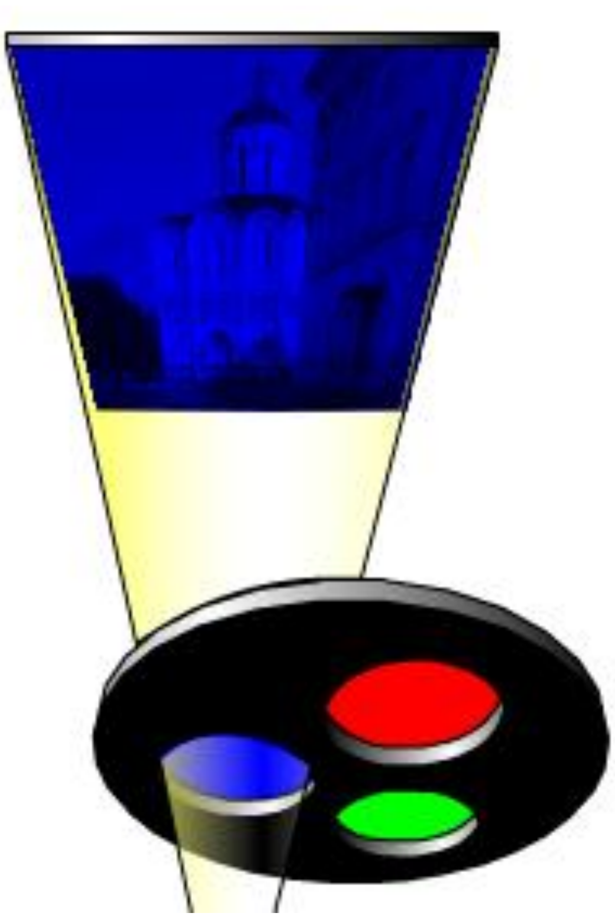

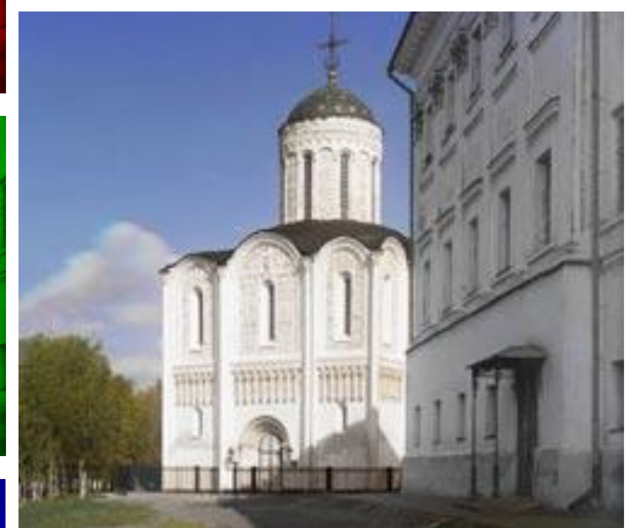

YungYu Chuang's slide

Prokudin-Gorskii (early 1900’s)

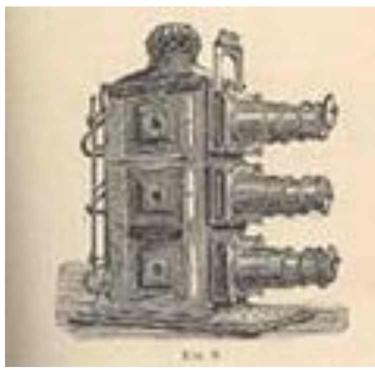

Lantern projector

http://www.loc.gov/exhibits/empire/

YungYu Chuang’s slide

Prokudin-Gorskii (early 1990's)

YungYu Chuang's slide

Color sensing in camera: Prism

- Requires three chips and precise alignment
- More expensive

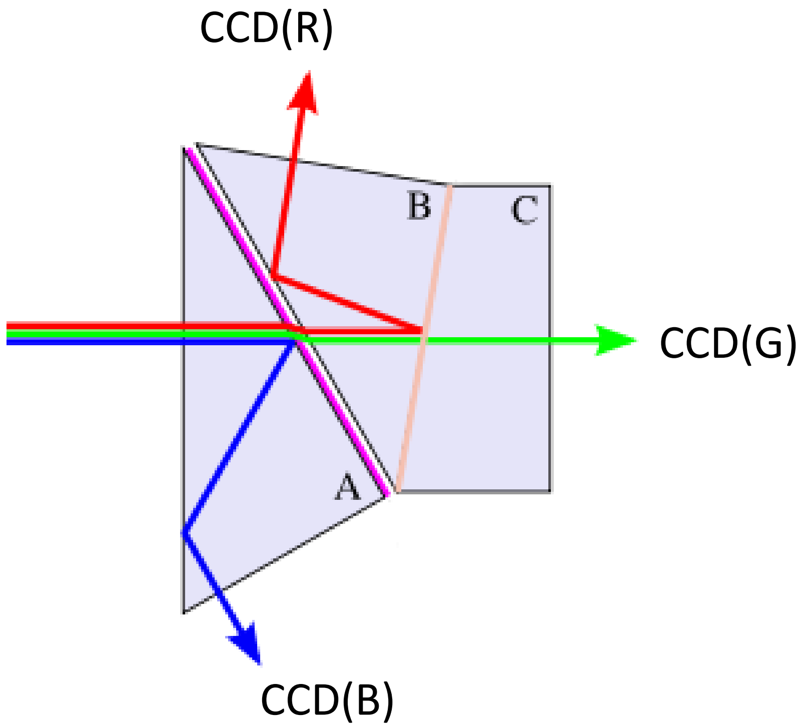

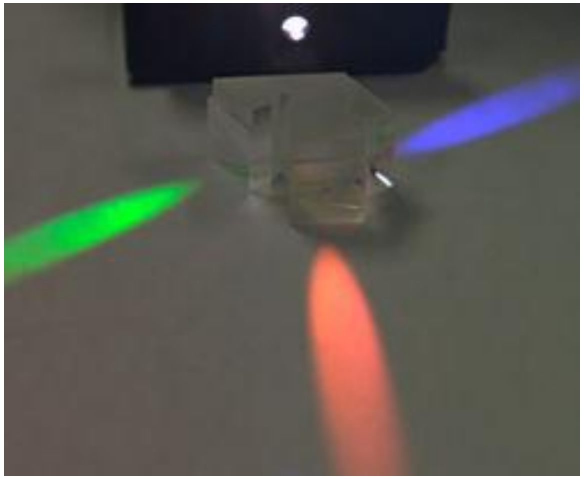

Color filter array

Bayer grid

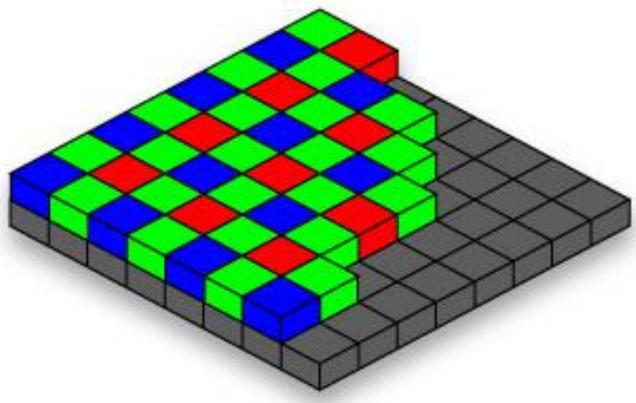

Estimate missing components from neighboring values (demosaicing)

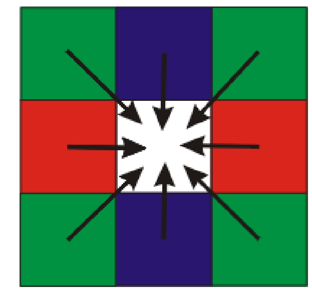

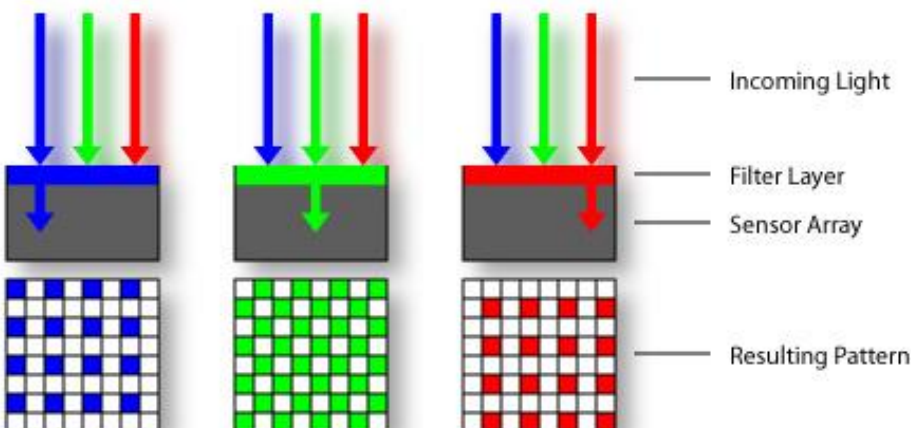

Why more green?

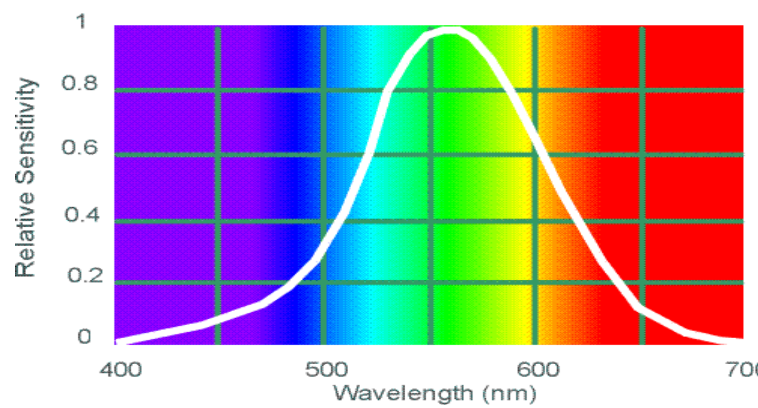

Human Luminance Sensitivity Function

Source: Steve Seitz

Bayer's pattern

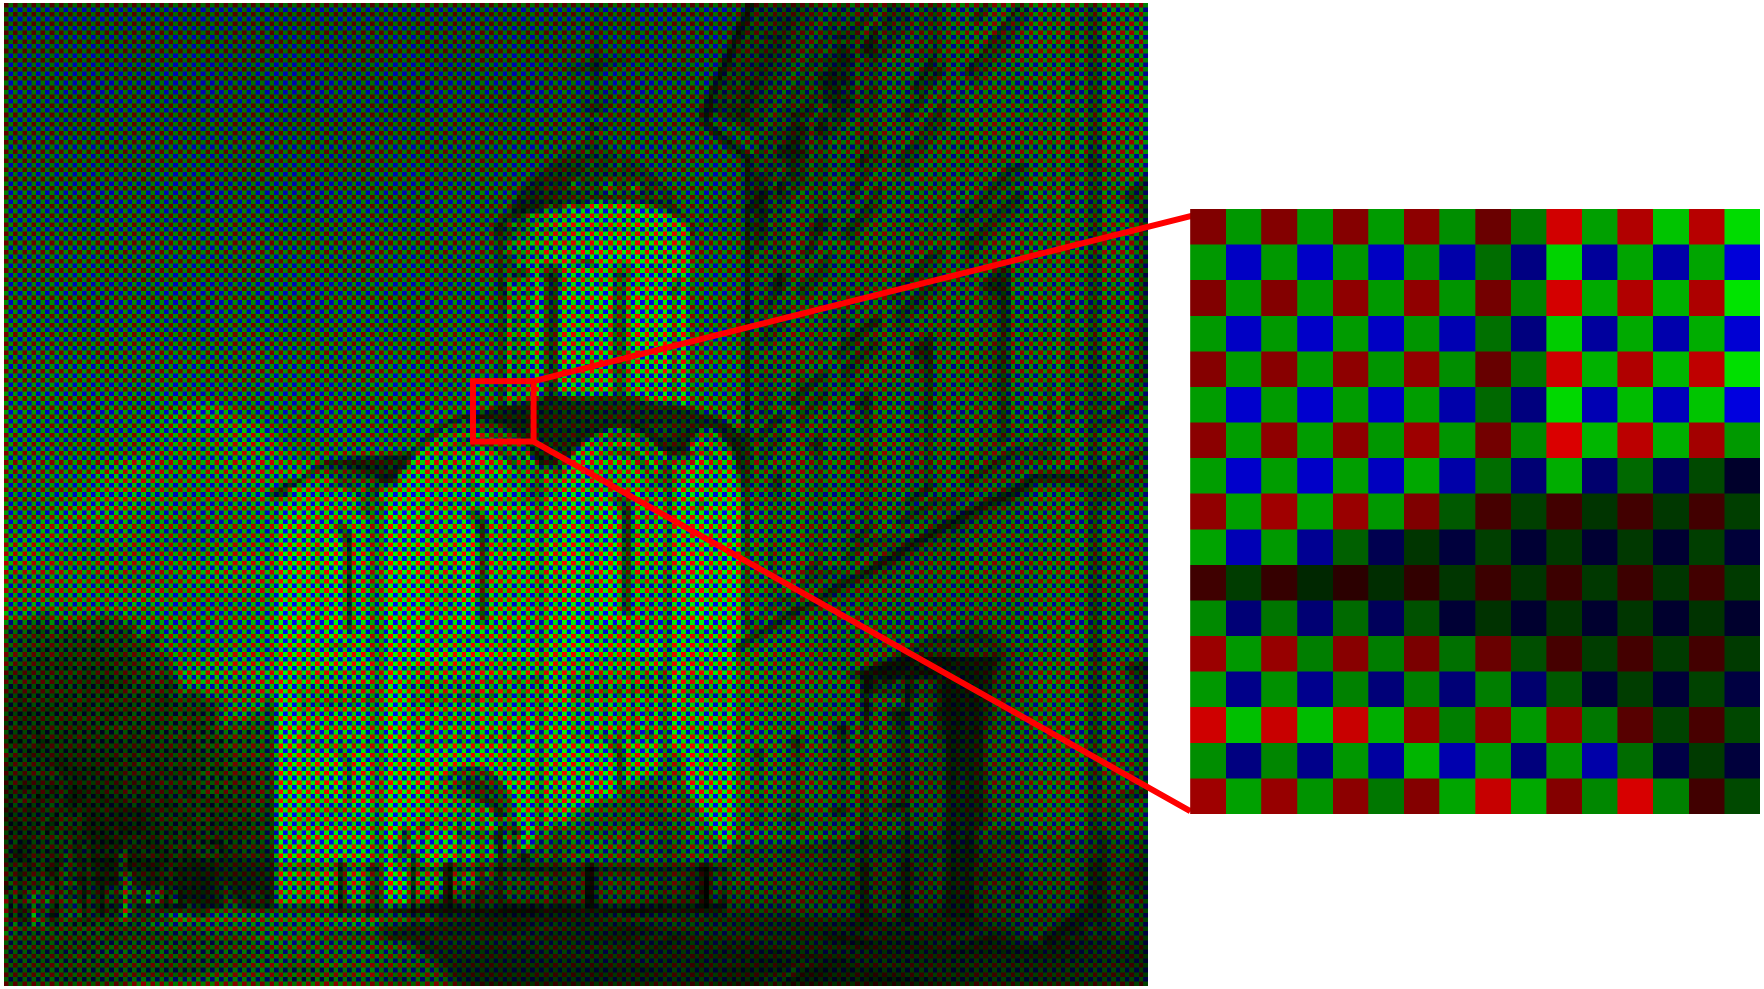

YungYu Chuang's slide

Color filter array

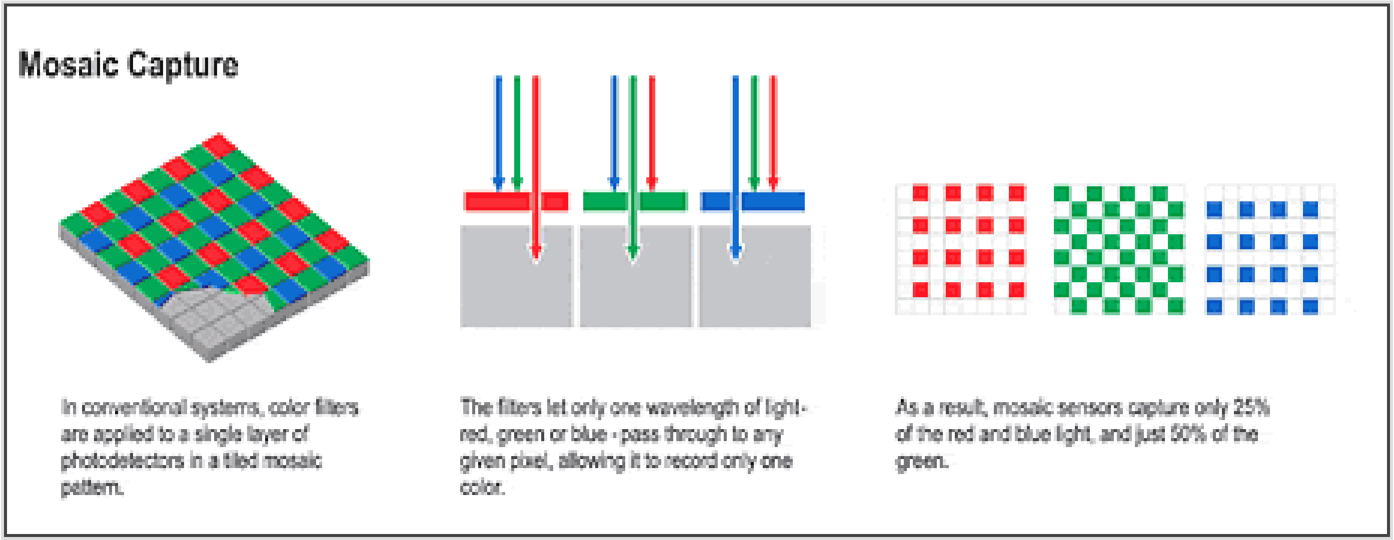

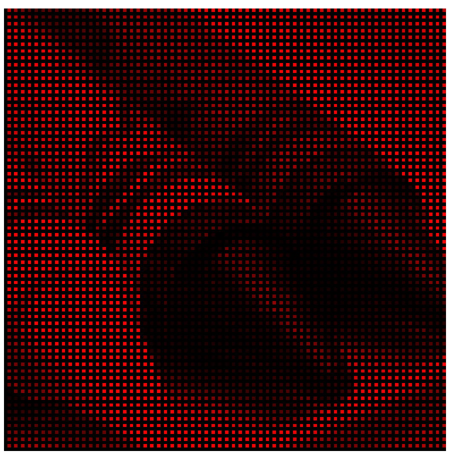

red

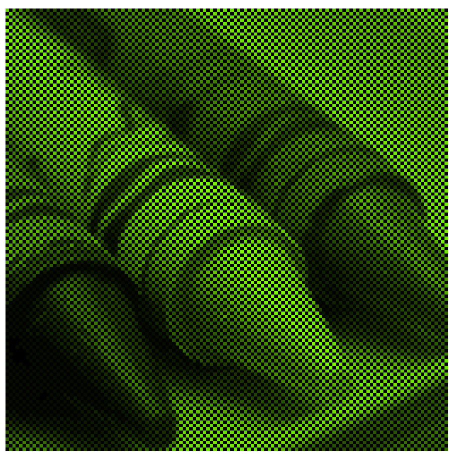

green

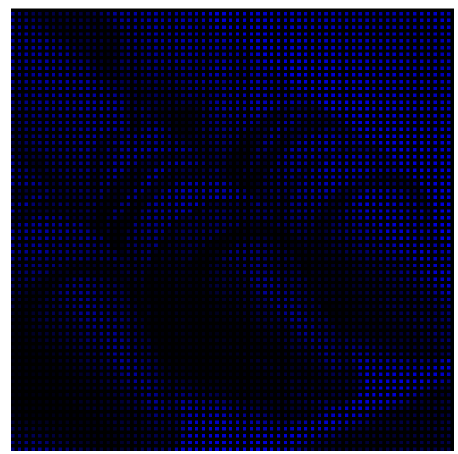

blue

output

YungYu Chuang's slide

Color images

- We'll treat color images as a vector-valued function:

$$f(x, y) = \begin{bmatrix} r(x, y) \\ g(x, y) \\ b(x, y) \end{bmatrix}$$

- We'll often convert to grayscale
  (e.g., 0.3 * r + 0.59 * g + 0.11 * b)

Dynamic range

- What is the range of light intensity that a camera can capture?
  - Called *dynamic range*
  - Digital cameras have difficulty capturing both high intensities and low intensities in the same image

Light response is nonlinear

- Our visual system has a large *dynamic range*
  - We can resolve both light and dark things at the same time
  - One mechanism for achieving this is that we sense light intensity on a *logarithmic scale*
    - an exponential intensity ramp will be seen as a linear ramp
  - Another mechanism is *adaptation*
    - rods and cones adapt to be more sensitive in low light, less sensitive in bright light.

Visual dynamic range

Background	Luminance (candelas per square meter)
Horizon sky	
Moonless overcast night	0.00003
Moonless clear night	0.0003
Moonlit overcast night	0.003
Moonlit clear night	0.03
Deep twilight	0.3
Twilight	3
Very dark day	30
Overcast day	300
Clear day	3,000
Day with sunlit clouds	30,000
Daylight fog	
Dull	300–1,000
Typical	1,000–3,000
Bright	3,000–16,000
Ground	
Overcast day	30–100
Sunny day	300
Snow in full sunlight	16,000

FIGURE 1.13
Luminance of everyday backgrounds. *Source:* Data from Rea, ed., *Lighting Handbook 1984 Reference and Application*, fig. 3-44, p. 3-24.

Dynamic range

- Our total dynamic range is high ($\sim10^9$)
- Our dynamic range *at a given time* is still pretty high ($\sim10^4$)
- A camera's dynamic range for a given exposure is relatively low ($2^8 = 256$ tonal values, range of about $\sim10^3$)

High dynamic range imaging

Historical context

- **Pinhole model:** Mozi (470-390 BC), Aristotle (384-322 BC)
- **Principles of optics (including lenses):** Alhacen (965-1039)
- **Camera obscura:** Leonardo da Vinci (1452-1519), Johann Zahn (1631-1707)
- **First photo:** Joseph Nicephore Niepce (1822)
- **Daguerréotypes** (1839)
- **Photographic film** (Eastman, 1889)
- **Cinema** (Lumière Brothers, 1895)
- **Color Photography** (Lumière Brothers, 1908)
- **Television** (Baird, Farnsworth, Zworykin, 1920s)
- **First consumer camera with CCD**: Sony Mavica (1981)
- **First fully digital camera:** Kodak DCS100 (1990)

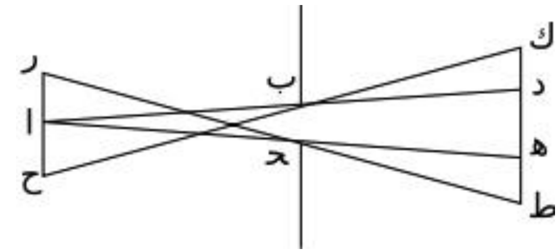

Alhacen’s notes

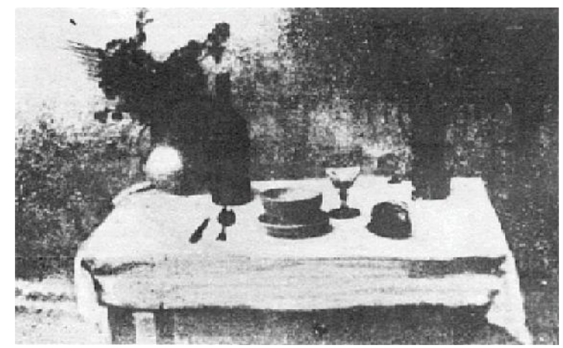

Niepce, “La Table Servie,” 1822

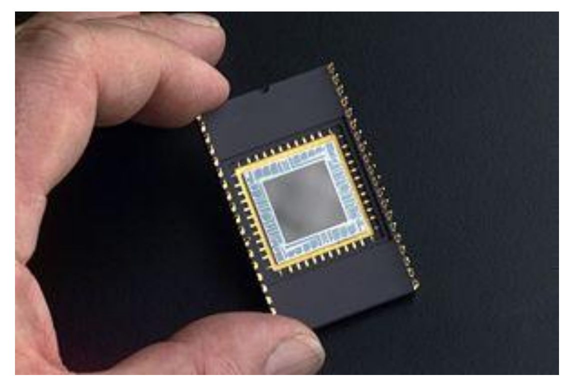

CCD chip

Questions?

- 3-minute break

Projection

Projection

Müller-Lyer Illusion

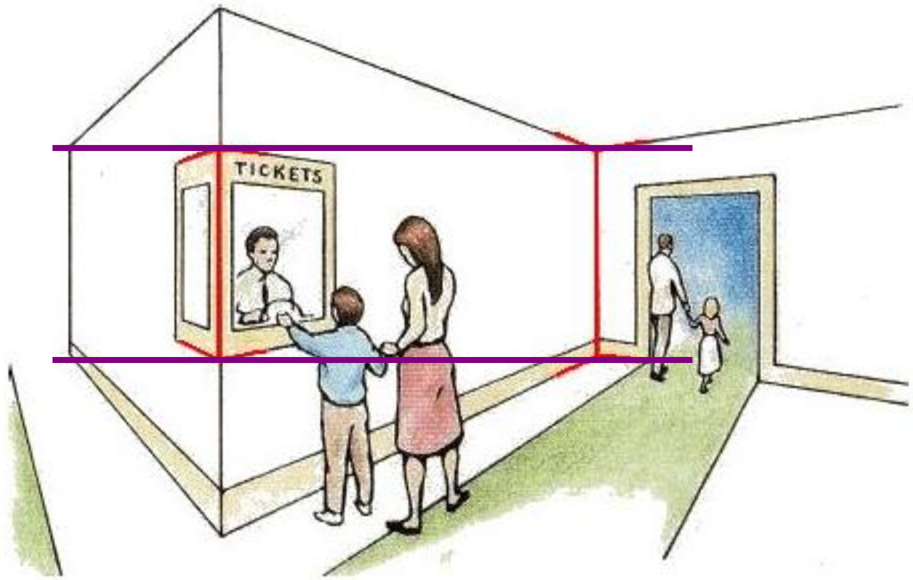

http://www.michaelbach.de/ot/sze_muelue/index.html

Modeling projection

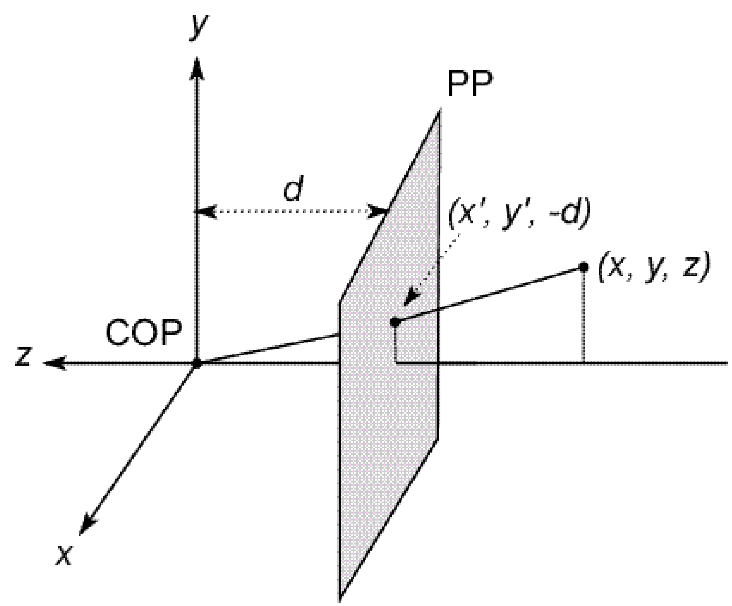

- The coordinate system
  - We will use the pinhole model as an approximation
  - Put the optical center (**C**enter **O**f **P**rojection) at the origin
  - Put the image plane (**P**rojection **P**lane) *in front* of the COP
    - Why?
  - The camera looks down the *negative* z axis
    - we need this if we want right-handed-coordinates

Modeling projection

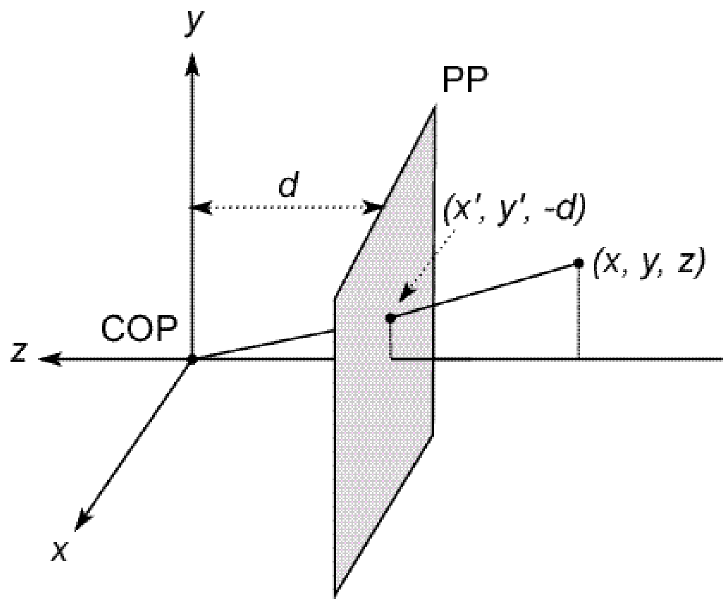

- Projection equations
  - Compute intersection with PP of ray from (x,y,z) to COP
  - Derived using similar triangles (on board)

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

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

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

Modeling projection

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

Homogeneous coordinates to the rescue!

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

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:

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

divide by third coordinate

This is known as **perspective projection**

- The matrix is the **projection matrix**

- (Can also represent as a 4x4 matrix – OpenGL does something like this)

Perspective Projection

- How does scaling the projection matrix change the transformation?

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

$$\begin{bmatrix} -d & 0 & 0 & 0 \\ 0 & -d & 0 & 0 \\ 0 & 0 & 1 & 0 \end{bmatrix} \begin{bmatrix} x \\ y \\ z \\ 1 \end{bmatrix} = \begin{bmatrix} -dx \\ -dy \\ z \end{bmatrix} \Rightarrow (-d\frac{x}{z}, \quad -d\frac{y}{z})$$

Orthographic projection

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

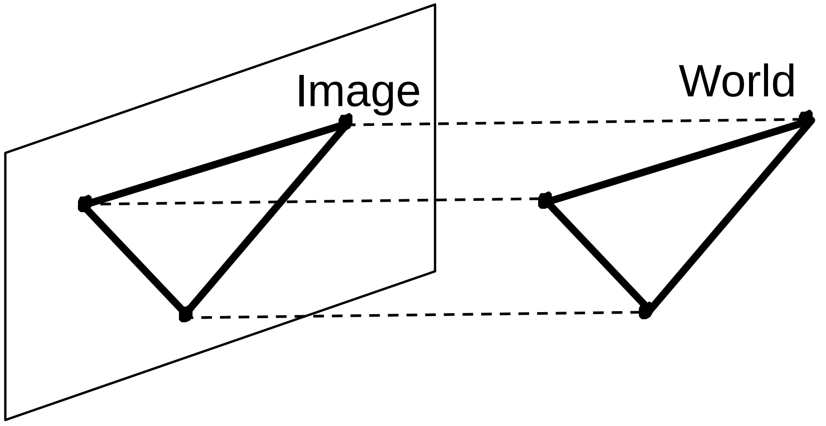

$$\begin{bmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} x \\ y \\ z \\ 1 \end{bmatrix} = \begin{bmatrix} x \\ y \\ 1 \end{bmatrix} \Rightarrow (x, y)$$

Orthographic projection

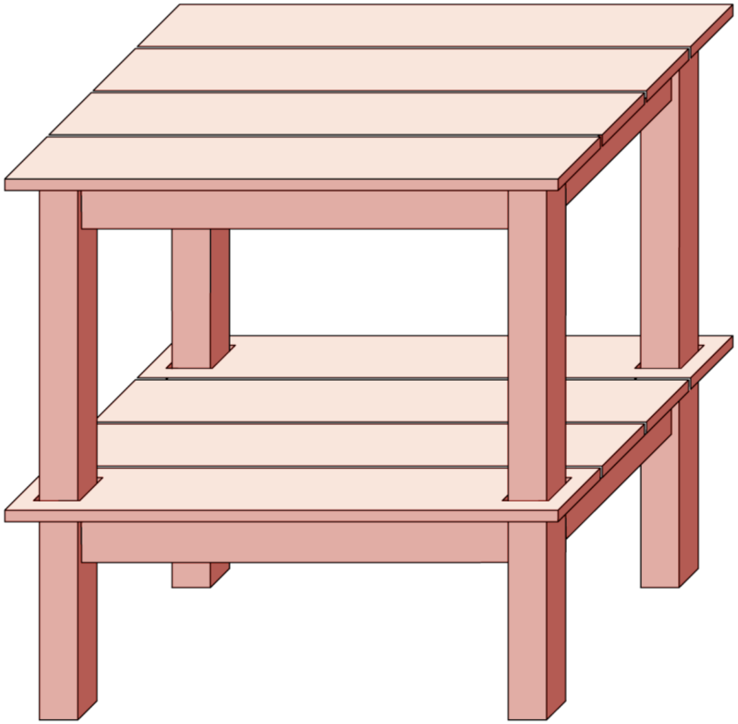

Perspective projection

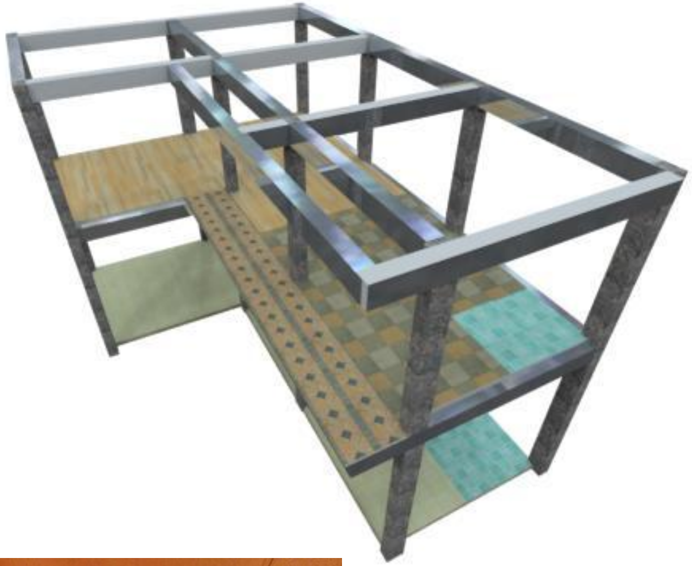

Perspective distortion

- What does a sphere project to?

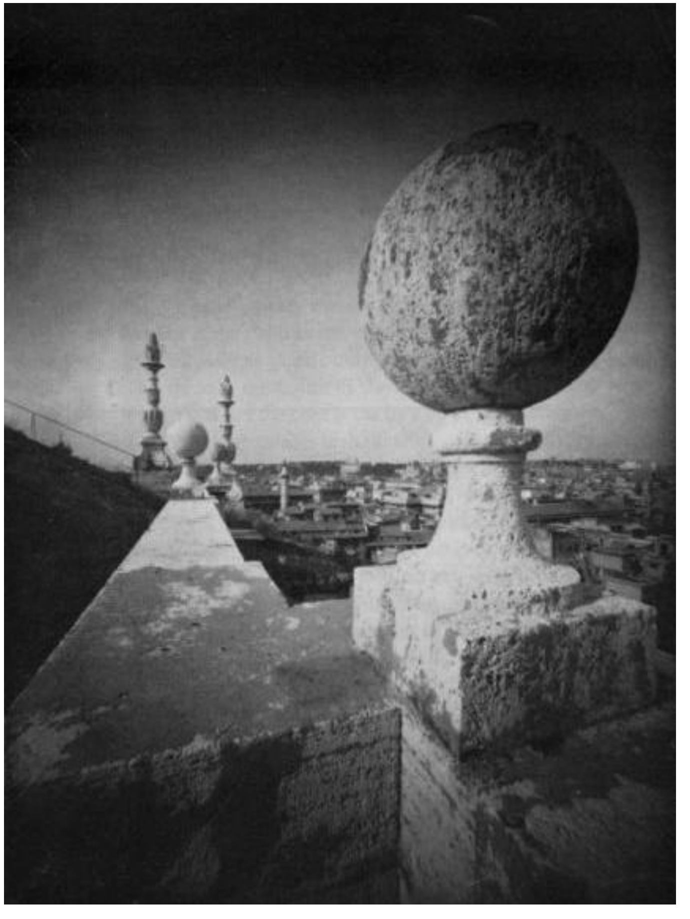

Image source: F. Durand

Perspective distortion

- The exterior columns appear bigger
- The distortion is not due to lens flaws
- Problem pointed out by Da Vinci

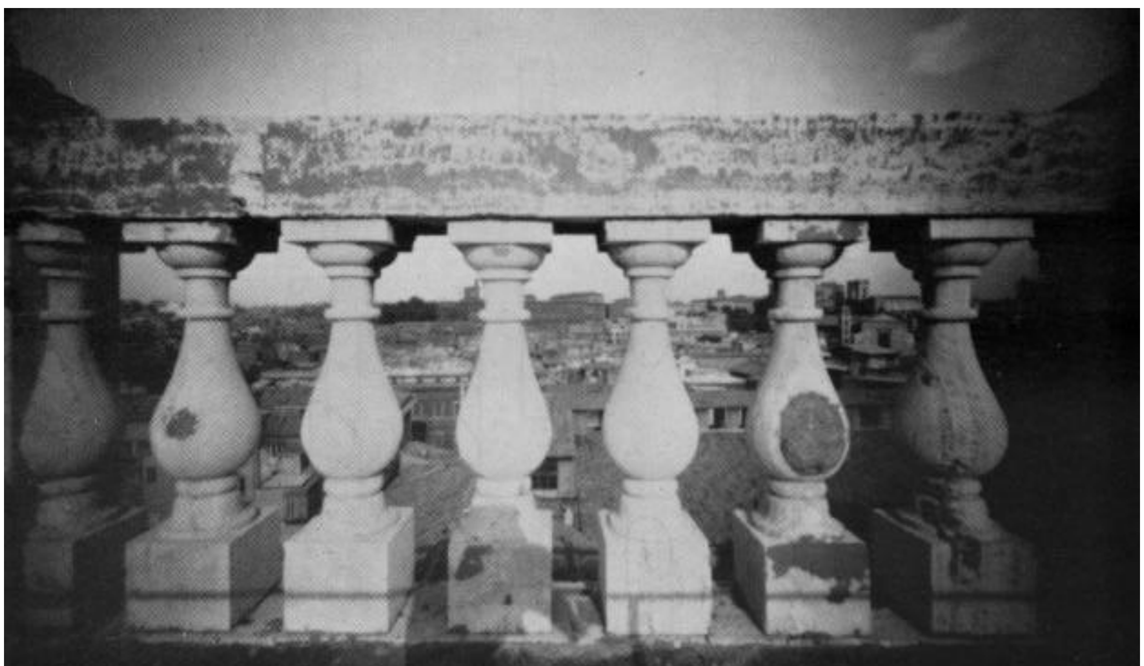

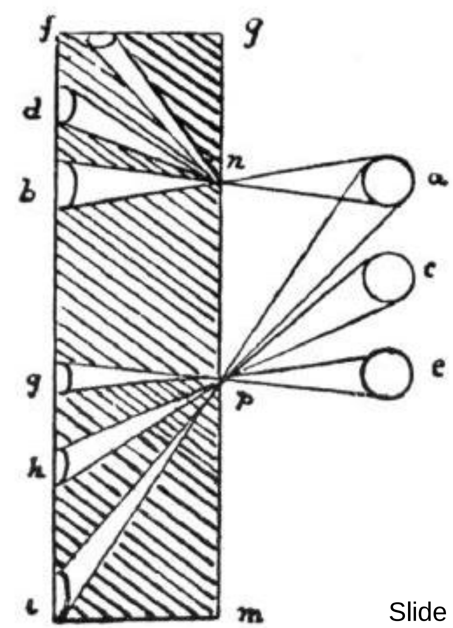

Slide by F. Durand

Perspective distortion: People

Distortion

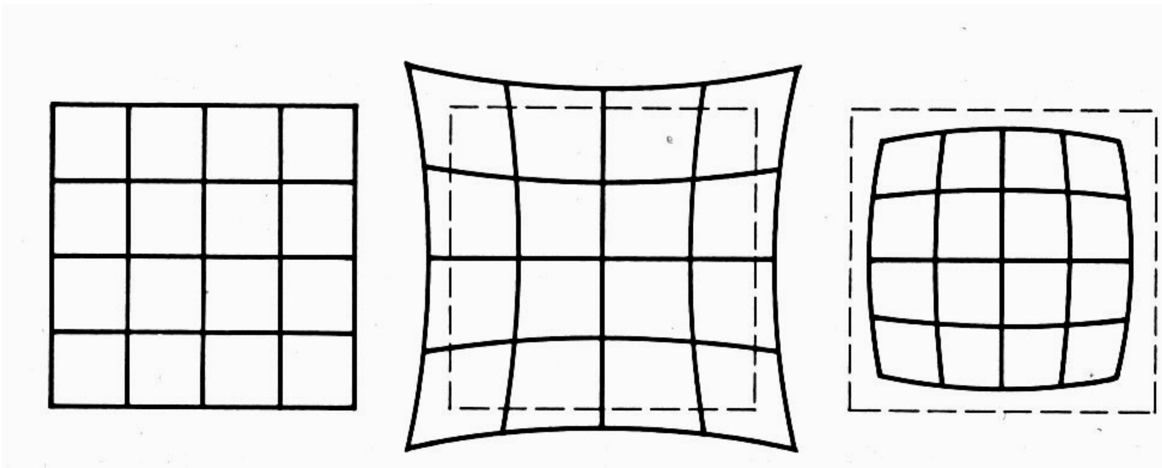

No distortion Pin cushion Barrel

- Radial distortion of the image
  - Caused by imperfect lenses
  - Deviations are most noticeable for rays that pass through the edge of the lens

Correcting radial distortion

from Helmut Dersch

Distortion

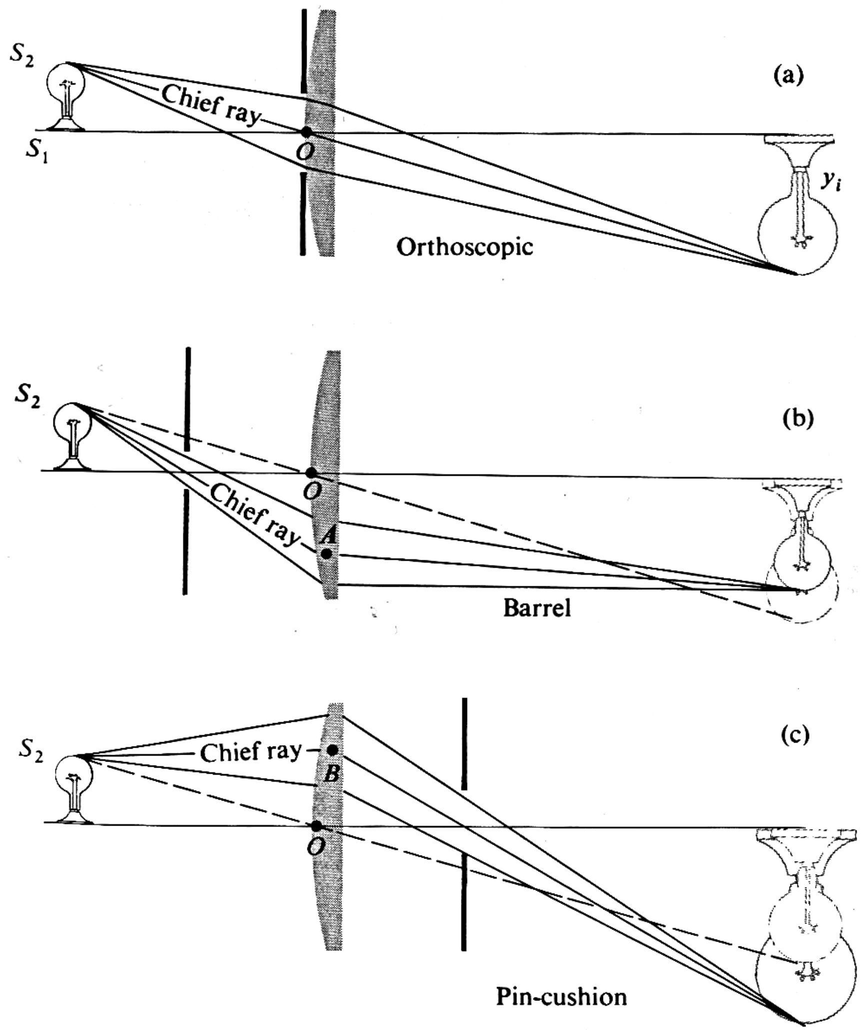

S_2
Chief ray
S_1
O
(a)
y_i
Orthoscopic
S_2
(b)
O
Chief ray
A
Barrel
S_2
Chief ray
B
O
(c)
Pin-cushion

Modeling distortion

Project $(\widehat{x}, \widehat{y}, \widehat{z})$ to "normalized" image coordinates

$$x'_n = \widehat{x}/\widehat{z}$$
$$y'_n = \widehat{y}/\widehat{z}$$

Apply radial distortion

$$r^2 = {x'_n}^2 + {y'_n}^2$$
$$x'_d = x'_n(1 + \kappa_1 r^2 + \kappa_2 r^4)$$
$$y'_d = y'_n(1 + \kappa_1 r^2 + \kappa_2 r^4)$$

Apply focal length translate image center

$$x' = f x'_d + x_c$$
$$y' = f y'_d + y_c$$

- To model lens distortion
  - Use above projection operation instead of standard projection matrix multiplication

Other types of projection

- Lots of intriguing variants…
- (I’ll just mention a few fun ones)

360 degree field of view…

- Basic approach
  - Take a photo of a parabolic mirror with an orthographic lens (Nayar)
  - Or buy one a lens from a variety of omnicam manufacturers…
    - See http://www.cis.upenn.edu/~kostas/omni.html

Tilt-shift

http://www.northlight-images.co.uk/article_pages/tilt_and_shift_ts-e.html

Titlt-shift images from Olivo Barbieri
and Photoshop imitations

Dimensionality Reduction Machine (3D to 2D)

3D world

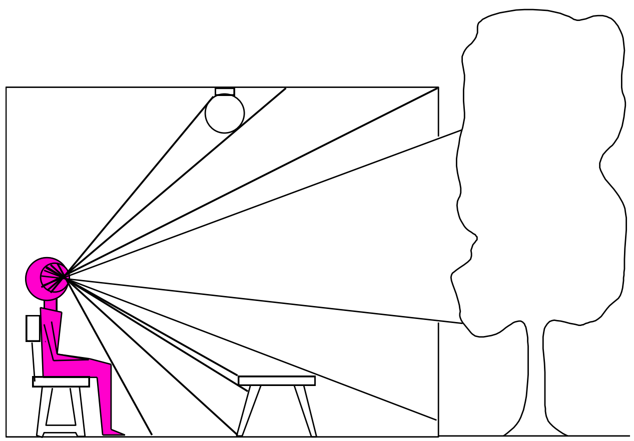

Point of observation

2D image

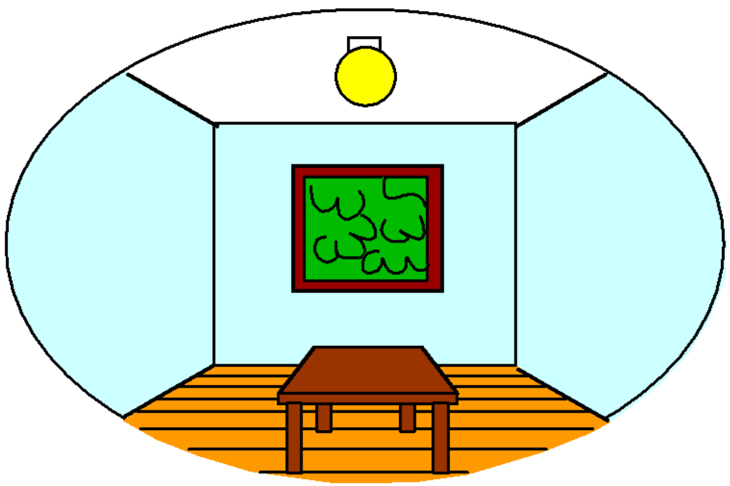

What have we lost?

- Angles
- Distances (lengths)

Slide by A. Efros

Projection properties

- Many-to-one: any points along same ray map to same point in image
- Points → points
- Lines → lines (collinearity is preserved)
  - But line through focal point projects to a point
- Planes → planes (or half-planes)
  - But plane through focal point projects to line

Projection properties

- Parallel lines converge at a vanishing point
  - Each direction in space has its own vanishing point
  - But parallels parallel to the image plane remain parallel

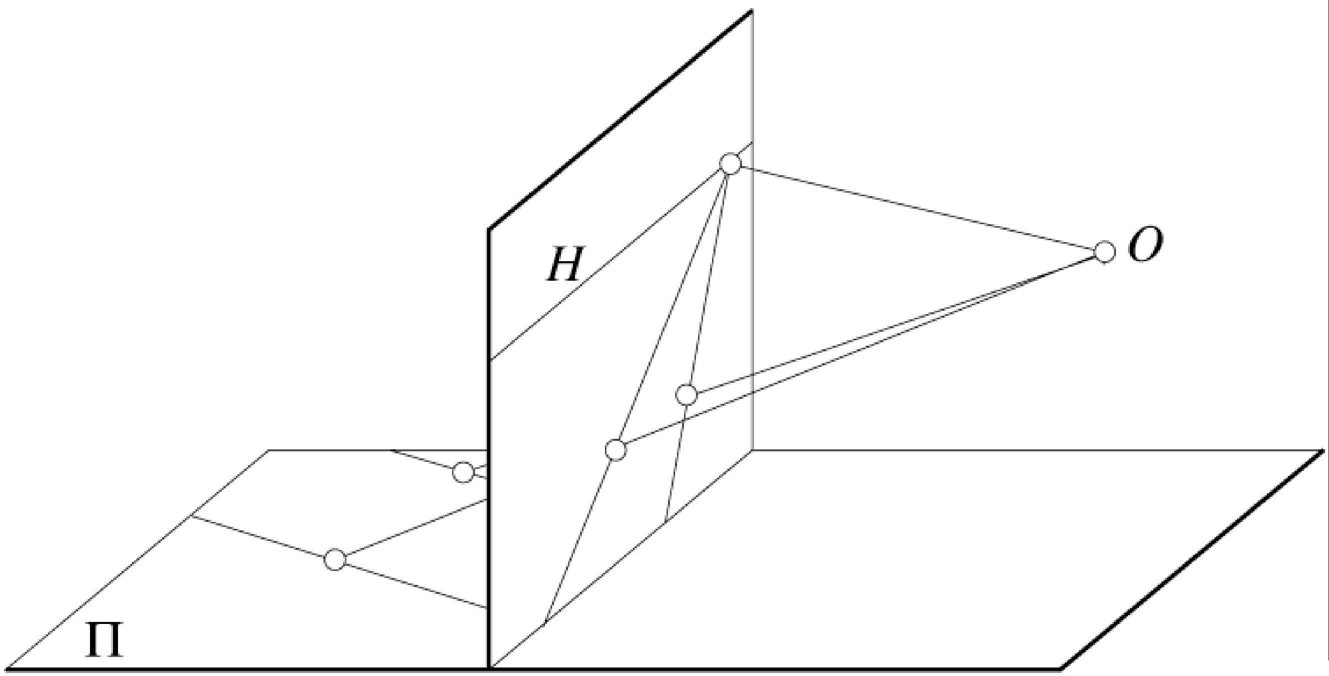

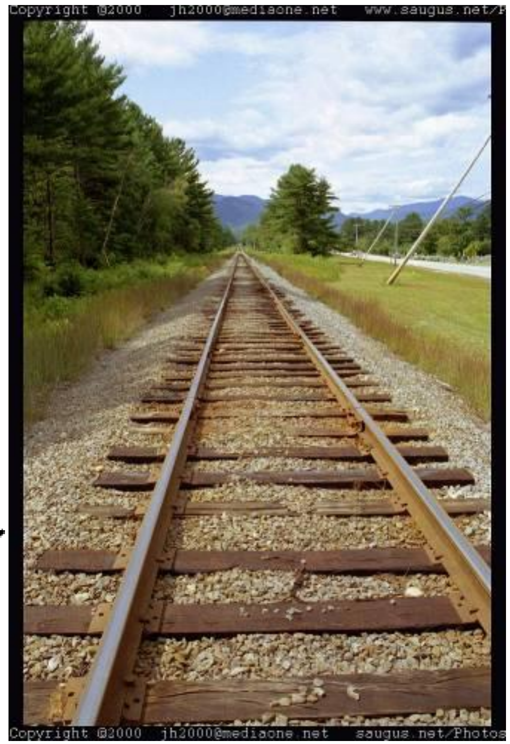

Camera parameters

- How can we model the geometry of a camera?

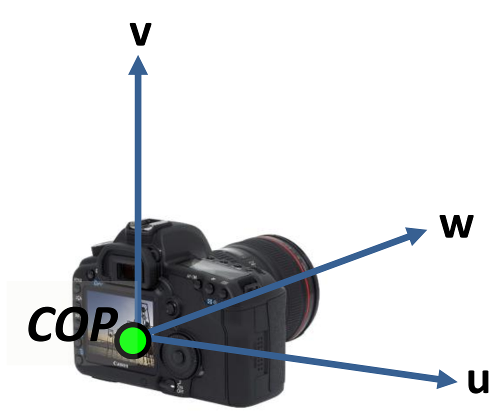

Camera

“The World”

Two important coordinate systems:

1. *World* coordinate system
2. *Camera* coordinate system

Camera parameters

- To project a point (*x*,*y*,*z*) in *world* coordinates into a camera
- First transform (*x*,*y*,*z*) into *camera* coordinates
- Need to know
  - Camera position (in world coordinates)
  - Camera orientation (in world coordinates)
- The project into the image plane
  - Need to know camera *intrinsics*

Camera parameters

A camera is described by several parameters

- Translation T of the optical center from the origin of world coords
- Rotation R of the image plane
- focal length f, principle point (x'_c, y'_c), pixel size (s_x, s_y)
- blue parameters are called "extrinsics," red are "intrinsics"

Projection equation

$$\mathbf{x} = \begin{bmatrix} sx \\ sy \\ s \end{bmatrix} = \begin{bmatrix} * & * & * & * \\ * & * & * & * \\ * & * & * & * \end{bmatrix} \begin{bmatrix} X \\ Y \\ Z \\ 1 \end{bmatrix} = \mathbf{\Pi X}$$

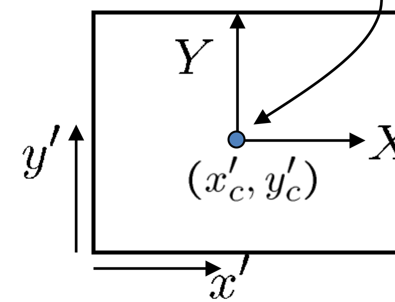

- The projection matrix models the cumulative effect of all parameters
- Useful to decompose into a series of operations

identity matrix

$$\mathbf{\Pi} = \begin{bmatrix} -fs_x & 0 & x'_c \\ 0 & -fs_y & y'_c \\ 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 \end{bmatrix} \begin{bmatrix} \mathbf{R}_{3x3} & \mathbf{0}_{3x1} \\ \mathbf{0}_{1x3} & 1 \end{bmatrix} \begin{bmatrix} \mathbf{I}_{3x3} & \mathbf{T}_{3x1} \\ \mathbf{0}_{1x3} & 1 \end{bmatrix}$$

intrinsics — projection — rotation — translation

- The definitions of these parameters are **not** completely standardized
  - especially intrinsics—varies from one book to another

Extrinsics

- How do we get the camera to "canonical form"?
  - (Center of projection at the origin, x-axis points right, y-axis points up, z-axis points backwards)

Step 1: Translate by -**c**

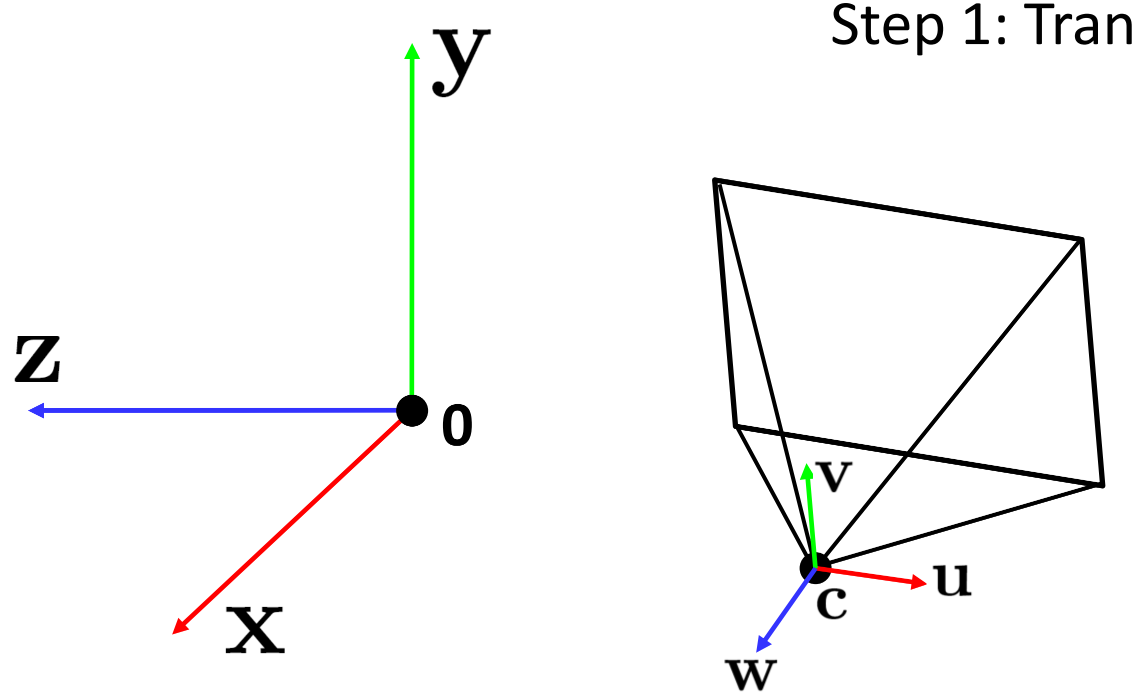

Extrinsics

- How do we get the camera to “canonical form”?
  - (Center of projection at the origin, x-axis points right, y-axis points up, z-axis points backwards)

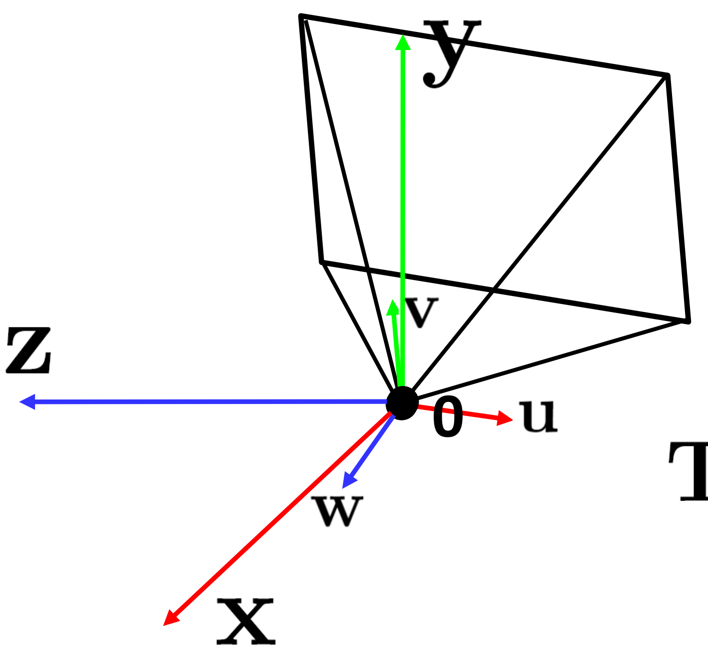

Step 1: Translate by -**c**

How do we represent translation as a matrix multiplication?

$$\mathbf{T} = \begin{bmatrix} \mathbf{I}_{3\times 3} & -\mathbf{c} \\ 0 \quad 0 \quad 0 & 1 \end{bmatrix}$$

Extrinsics

- How do we get the camera to "canonical form"?
  - (Center of projection at the origin, x-axis points right, y-axis points up, z-axis points backwards)

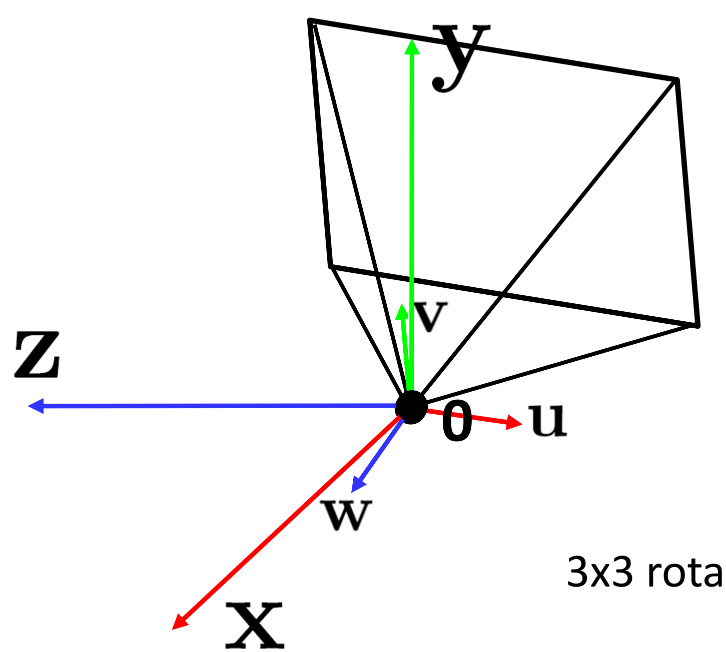

Step 1: Translate by -**c**

Step 2: Rotate by **R**

$$\mathbf{R} = \begin{bmatrix} \mathbf{u}^T \\ \mathbf{v}^T \\ \mathbf{w}^T \end{bmatrix}$$

3x3 rotation matrix

Extrinsics

- How do we get the camera to “canonical form”?
  - (Center of projection at the origin, x-axis points right, y-axis points up, z-axis points backwards)

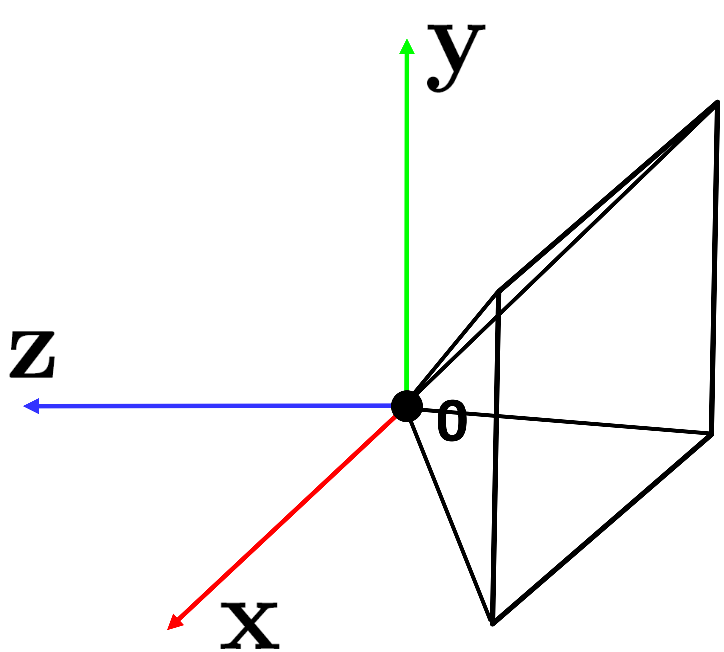

Step 1: Translate by -**c**

Step 2: Rotate by **R**

$$\mathbf{R} = \begin{bmatrix} \mathbf{u}^T \\ \mathbf{v}^T \\ \mathbf{w}^T \end{bmatrix}$$

Perspective projection

$$\underbrace{\begin{bmatrix} -f & 0 & 0 \\ 0 & -f & 0 \\ 0 & 0 & 1 \end{bmatrix}}_{\mathbf{K}} \begin{bmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 \end{bmatrix}$$

$\mathbf{K}$ (intrinsics) (converts from 3D rays in camera coordinate system to pixel coordinates)

in general, $\mathbf{K} = \begin{bmatrix} -f & s & c_x \\ 0 & -\alpha f & c_y \\ 0 & 0 & 1 \end{bmatrix}$ (upper triangular matrix)

α : **aspect ratio** (1 unless pixels are not square)

s : **skew** (0 unless pixels are shaped like rhombi/parallelograms)

(c_x, c_y) : **principal point** ((0,0) unless optical axis doesn’t intersect projection plane at origin)

Focal length

- Can think of as “zoom”

24mm

50mm

200mm

800mm

- Also related to *field of view*

Projection matrix

$$\mathbf{\Pi} = \underbrace{\mathbf{K}}_{\text{intrinsics}} \underbrace{\underbrace{\begin{bmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 \end{bmatrix}}_{\text{projection}} \underbrace{\begin{bmatrix} \mathbf{R} & \begin{matrix} 0 \\ 0 \\ 0 \end{matrix} \\ \begin{matrix} 0 & 0 & 0 \end{matrix} & 1 \end{bmatrix}}_{\text{rotation}} \underbrace{\begin{bmatrix} \mathbf{I}_{3\times 3} & -\mathbf{c} \\ \begin{matrix} 0 & 0 & 0 \end{matrix} & 1 \end{bmatrix}}_{\text{translation}}}$$

$$\left[\ \mathbf{R}\ \middle|\ \underbrace{-\mathbf{Rc}}_{\text{(}\mathbf{t}\text{ in book's notation)}}\ \right]$$

$$\mathbf{\Pi} = \mathbf{K}\left[\ \mathbf{R}\ \middle|\ -\mathbf{Rc}\ \right]$$

Projection matrix

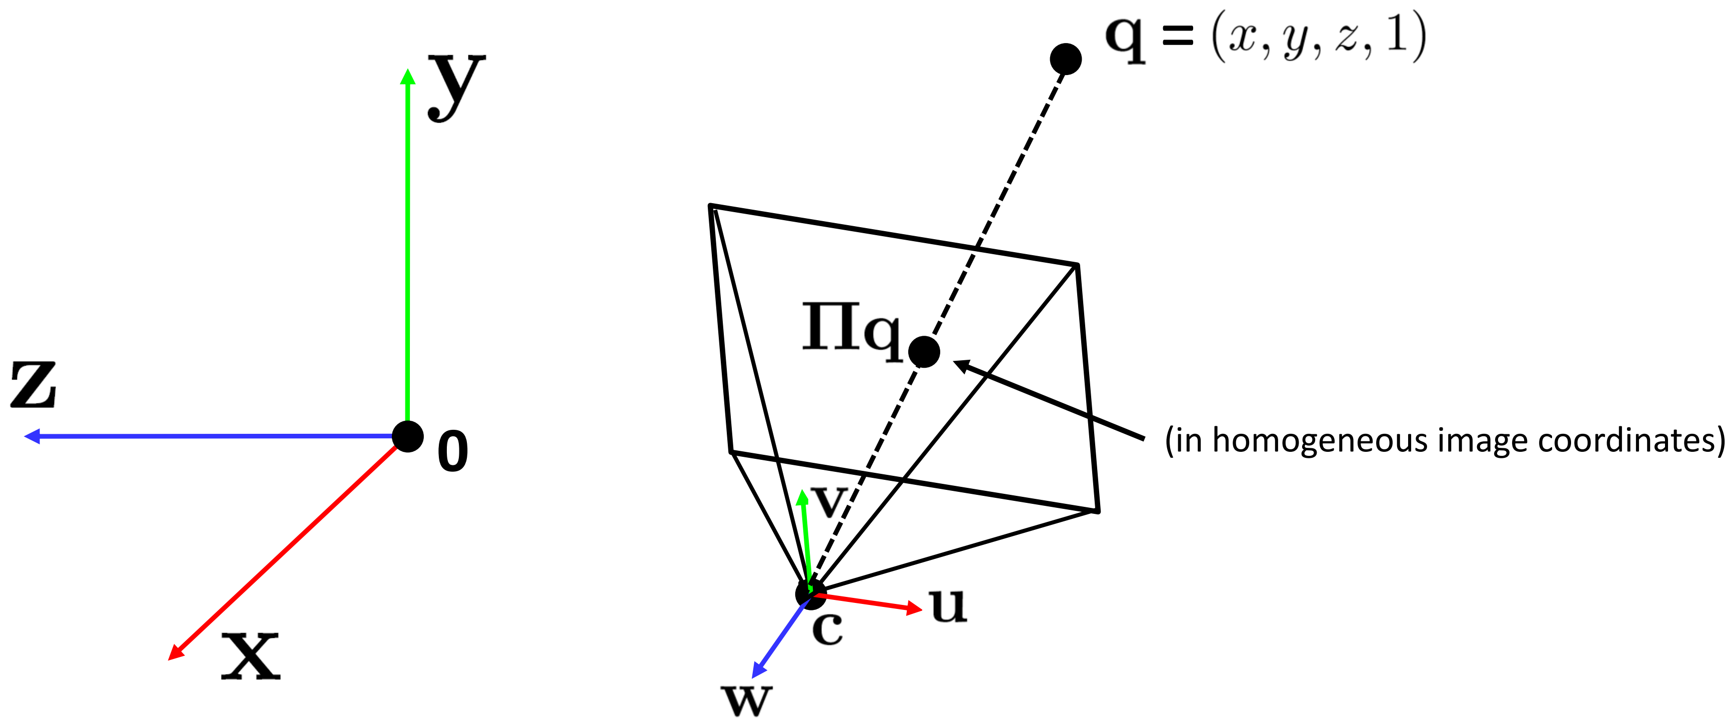

Questions?