Single-view Metrology + More with Pinhole Geometry

Computational Photography

Derek Hoiem, University of Illinois

Last Tuesday

- David Forsyth talked about physically grounded image editing, including inserting 3D objects into photographs

Last Thursday

- Amin Sadeghi talked transforming faces and PCA

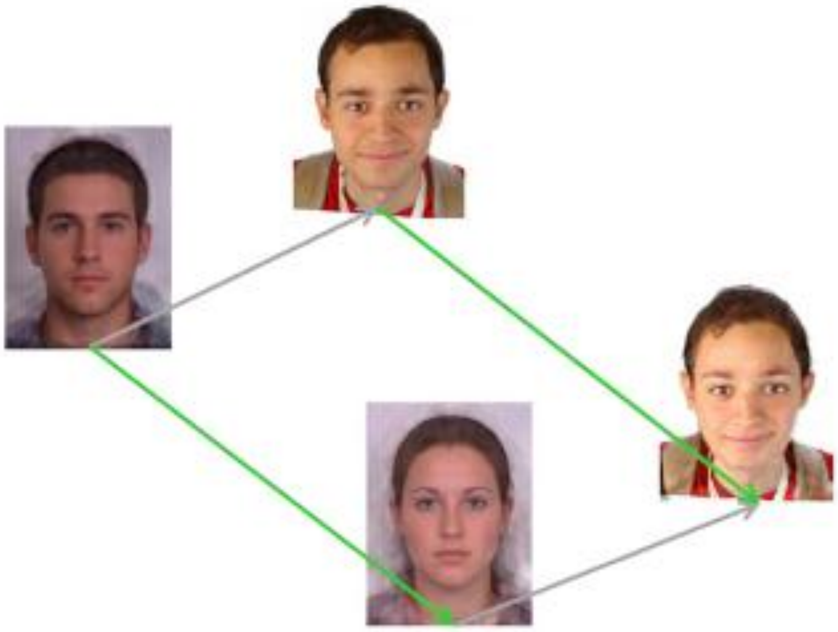

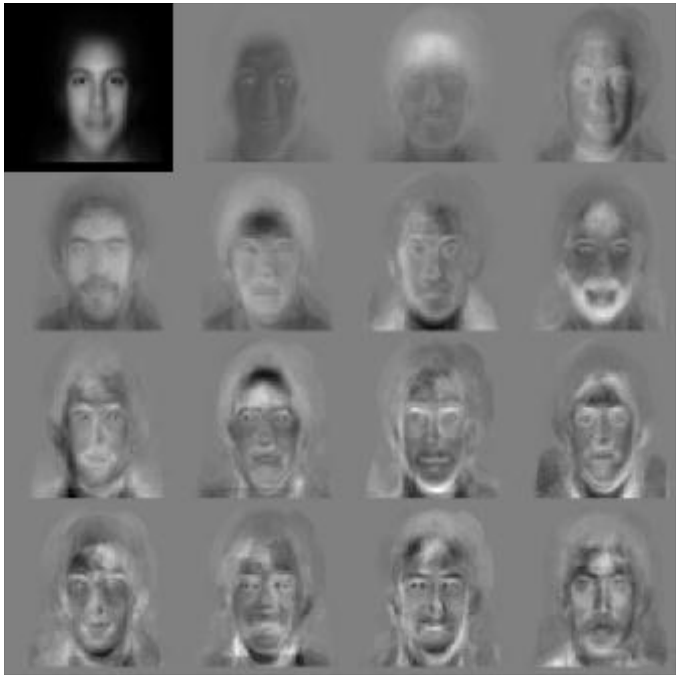

Project 2

- Class favorites: vote please
  - “Favorites” will be presented next Tues

Project 3

- Due Monday
  - http://courses.engr.illinois.edu/cs498dh3/projects/gradient/ComputationalPhotography_ProjectGradient.html
- Discussed in lecture 9 (9/25)

Project 3: Toy Problem

1. Preserve x-y gradients
2. Preserve intensity of one pixel

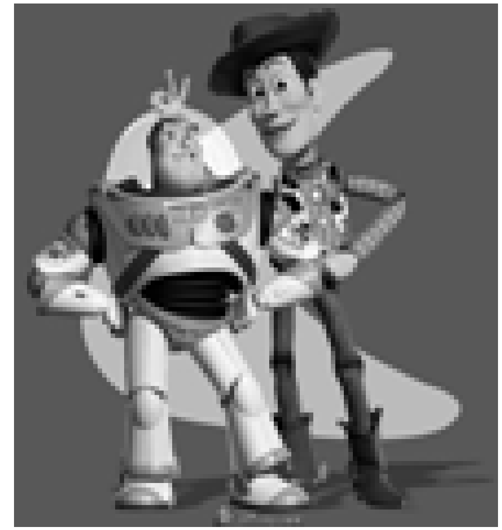

Source pixels: s

Variable pixels: v

1. minimize (v(x+1,y)-v(x,y) - (s(x+1,y)-s(x,y))^2
2. minimize (v(x,y+1)-v(x,y) - (s(x,y+1)-s(x,y))^2
3. minimize (v(1,1)-s(1,1))^2

Example

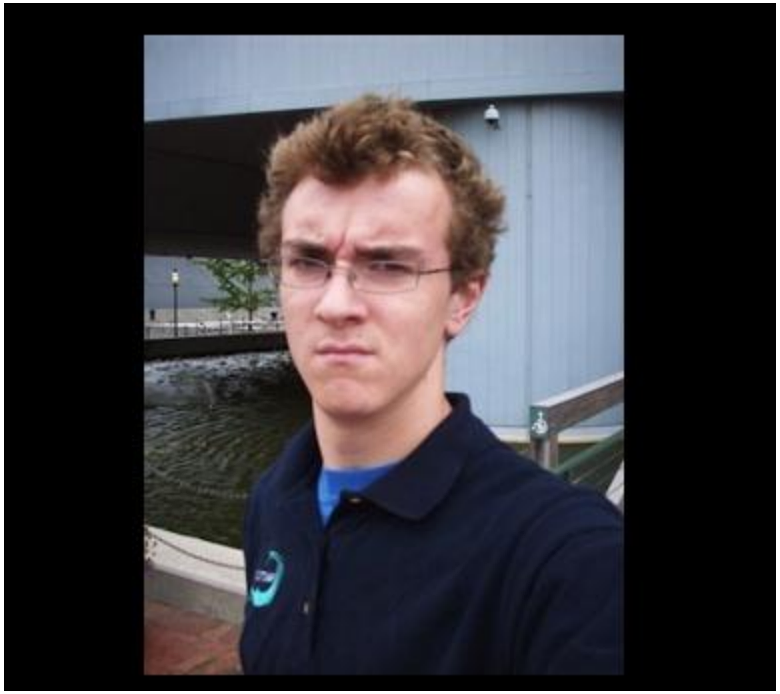

Gradient Visualization

Source: Evan Wallace

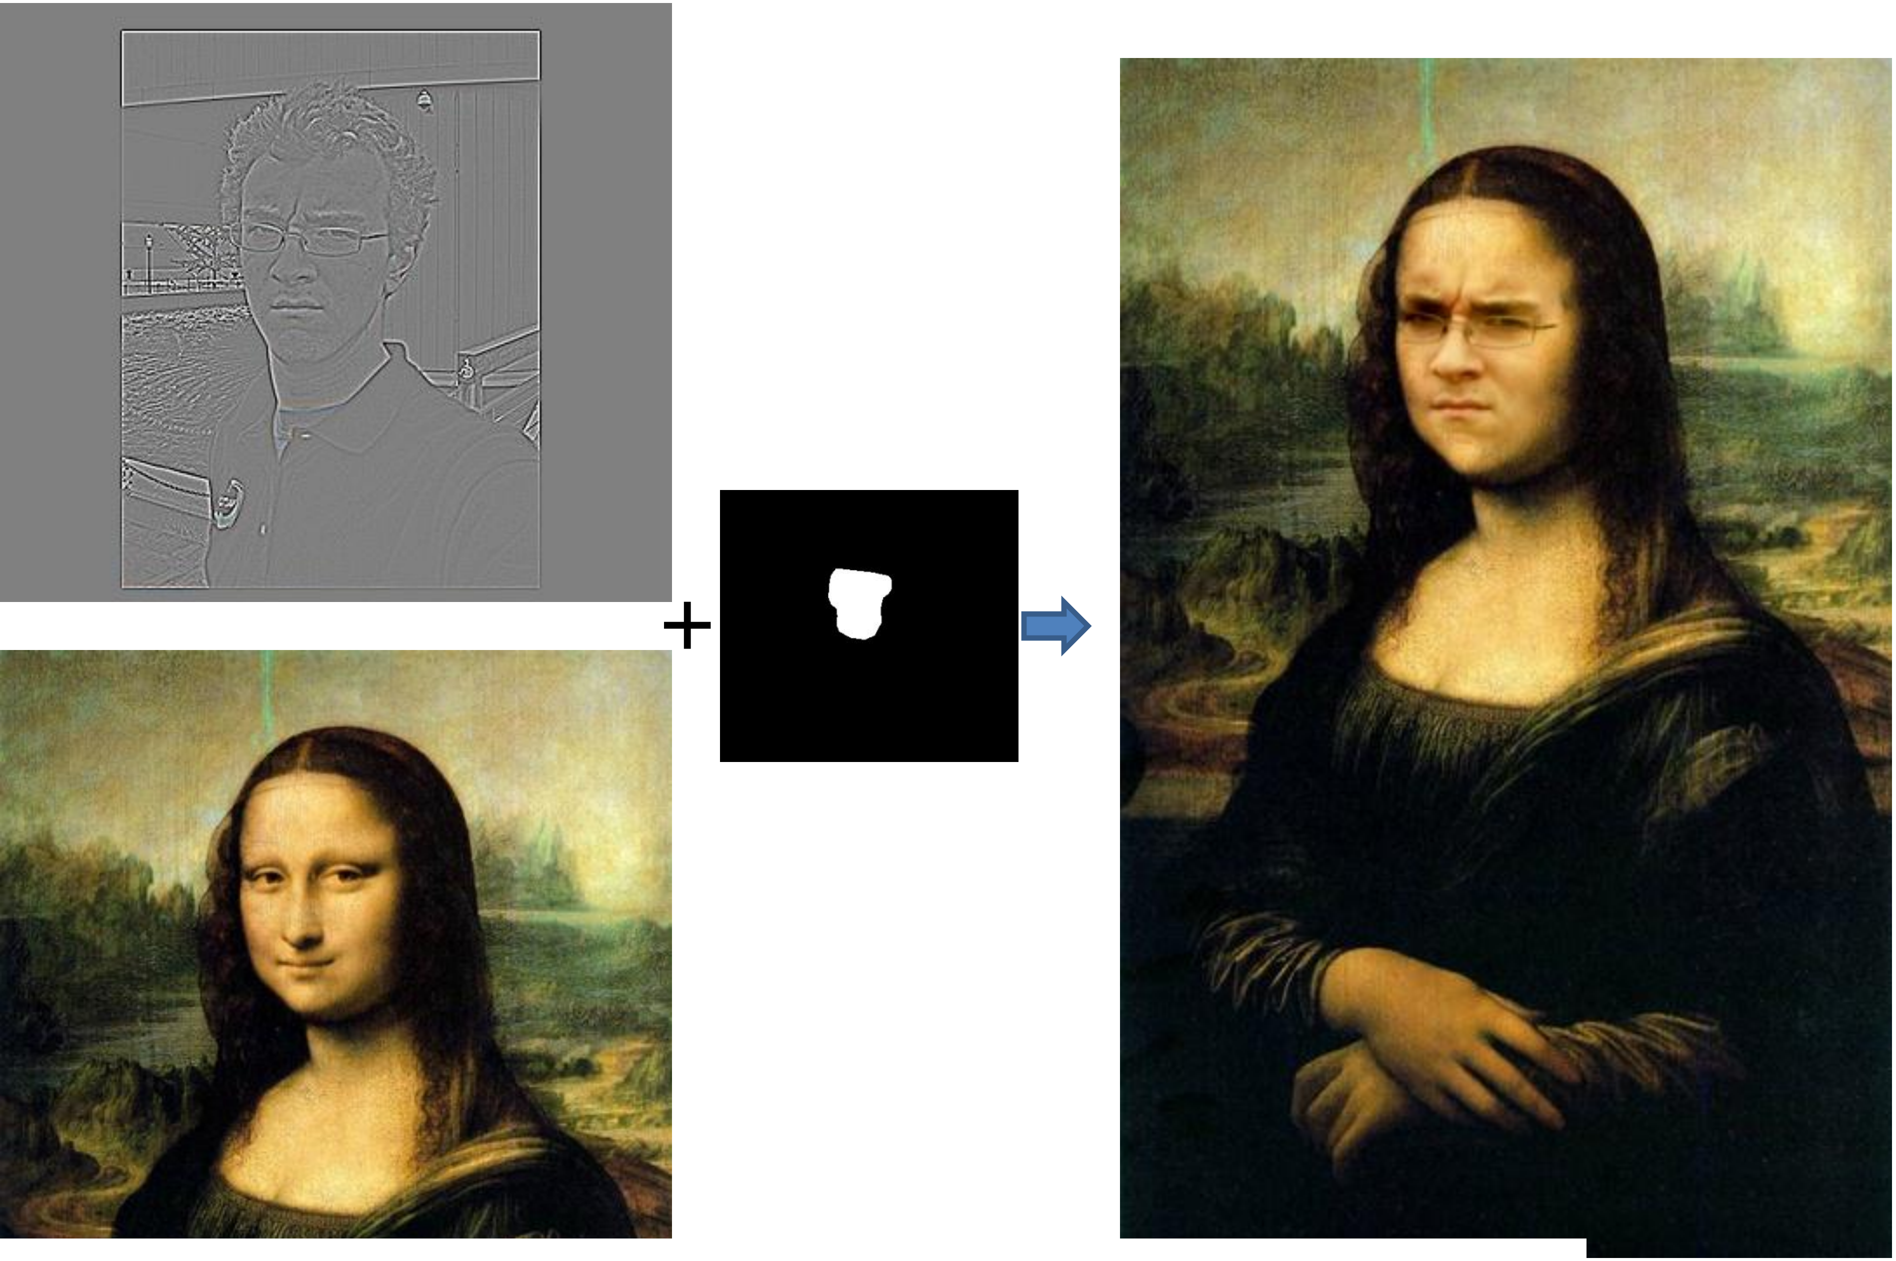

$$\mathbf{v} = \underset{\mathbf{v}}{\operatorname{argmin}} \sum_{i \in S, j \in N_i \cap S} ((v_i - v_j) - (s_i - s_j))^2 + \sum_{i \in S, j \in N_i \cap \neg S} ((v_i - t_j) - (s_i - s_j))^2$$

Source: Evan Wallace

Review: Pinhole Camera

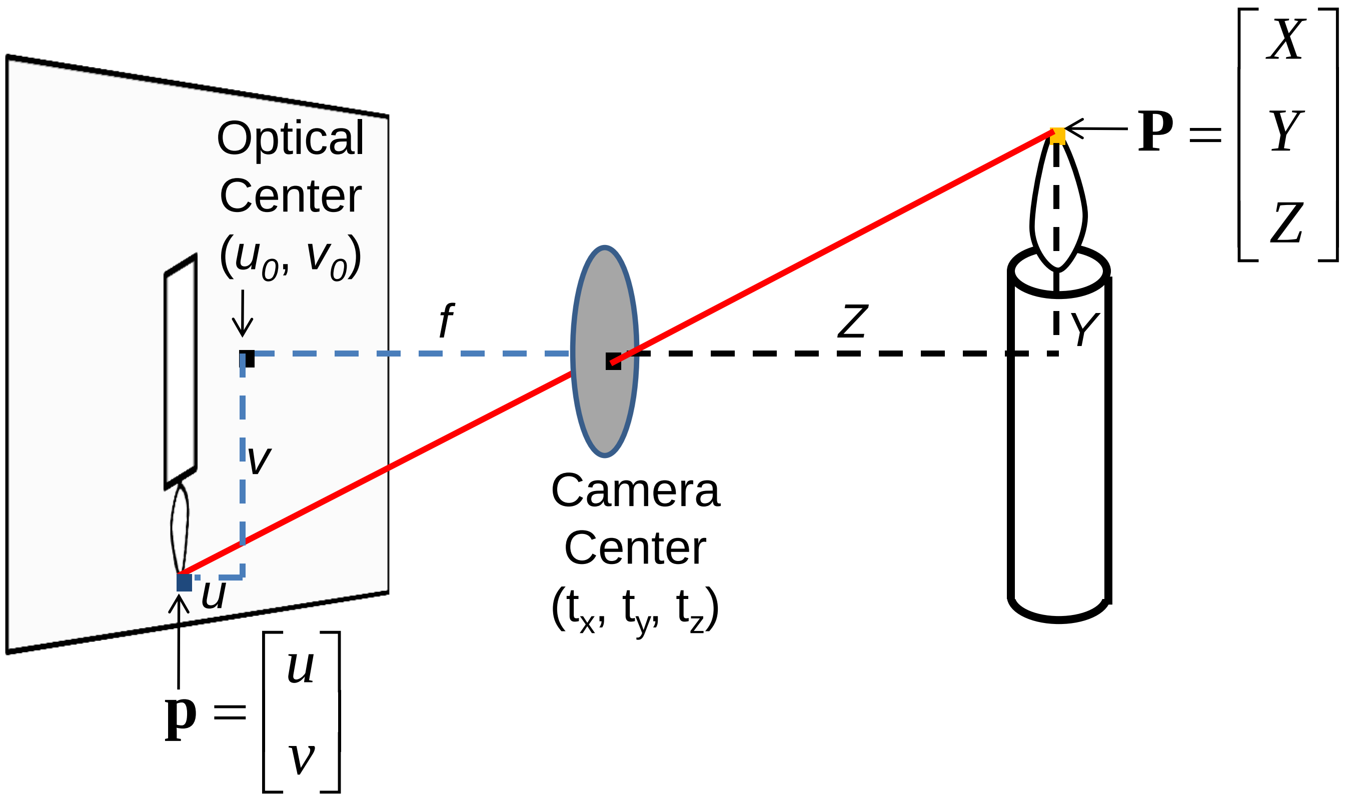

Review: Projection Matrix

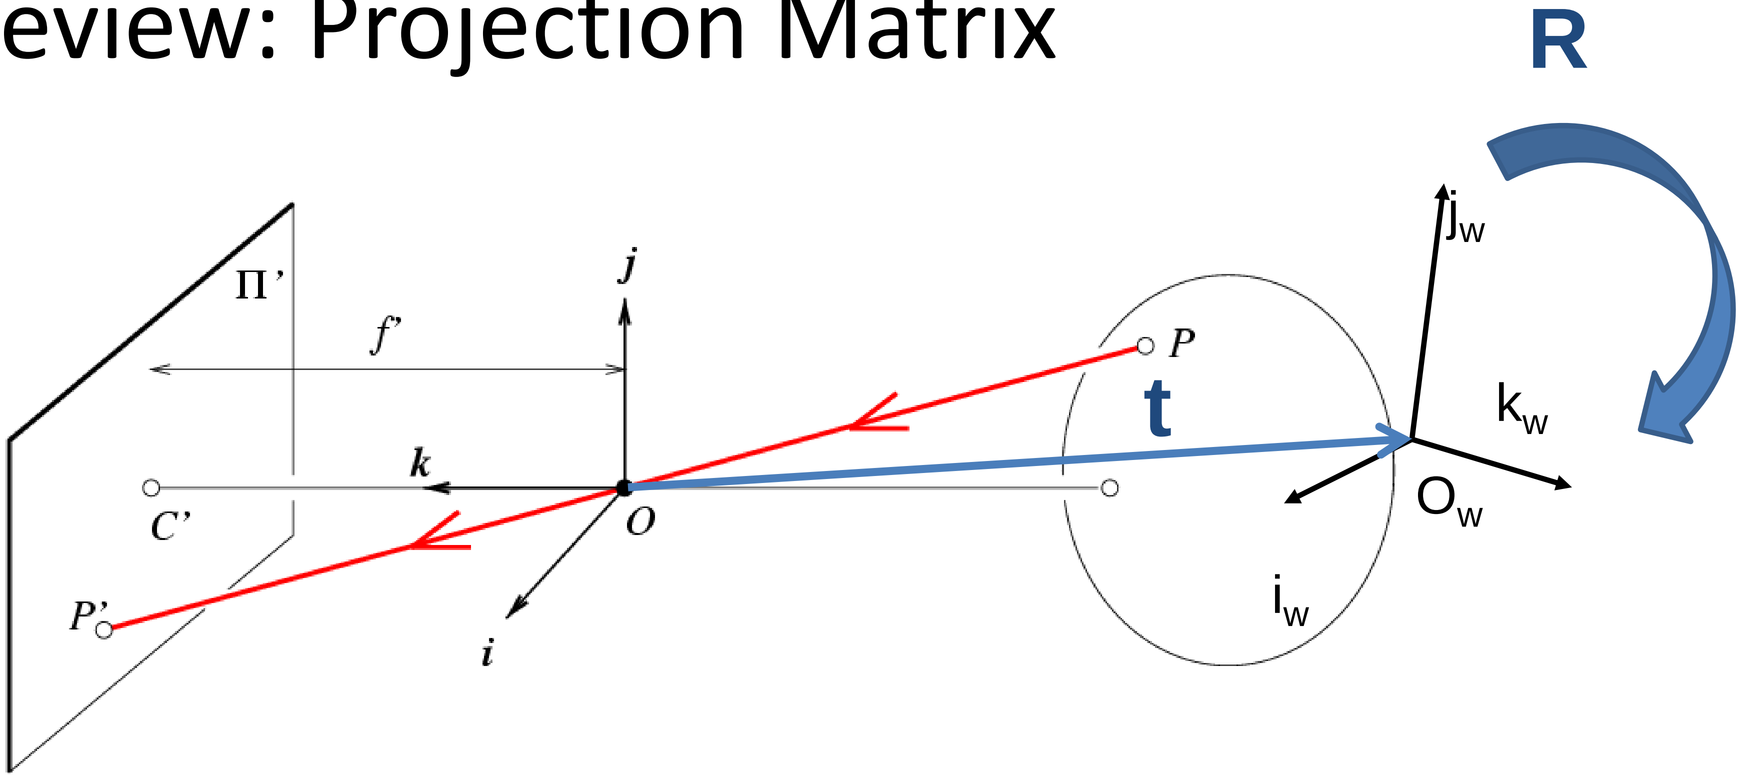

$$\mathbf{x} = \mathbf{K}\begin{bmatrix}\mathbf{R} & \mathbf{t}\end{bmatrix}\mathbf{X} \Rightarrow w\begin{bmatrix} u \\ v \\ 1 \end{bmatrix} = \begin{bmatrix} f & s & u_0 \\ 0 & \alpha f & v_0 \\ 0 & 0 & 1 \end{bmatrix}\begin{bmatrix} r_{11} & r_{12} & r_{13} & t_x \\ r_{21} & r_{22} & r_{23} & t_y \\ r_{31} & r_{32} & r_{33} & t_z \end{bmatrix}\begin{bmatrix} X \\ Y \\ Z \\ 1 \end{bmatrix}$$

Review: Vanishing Points

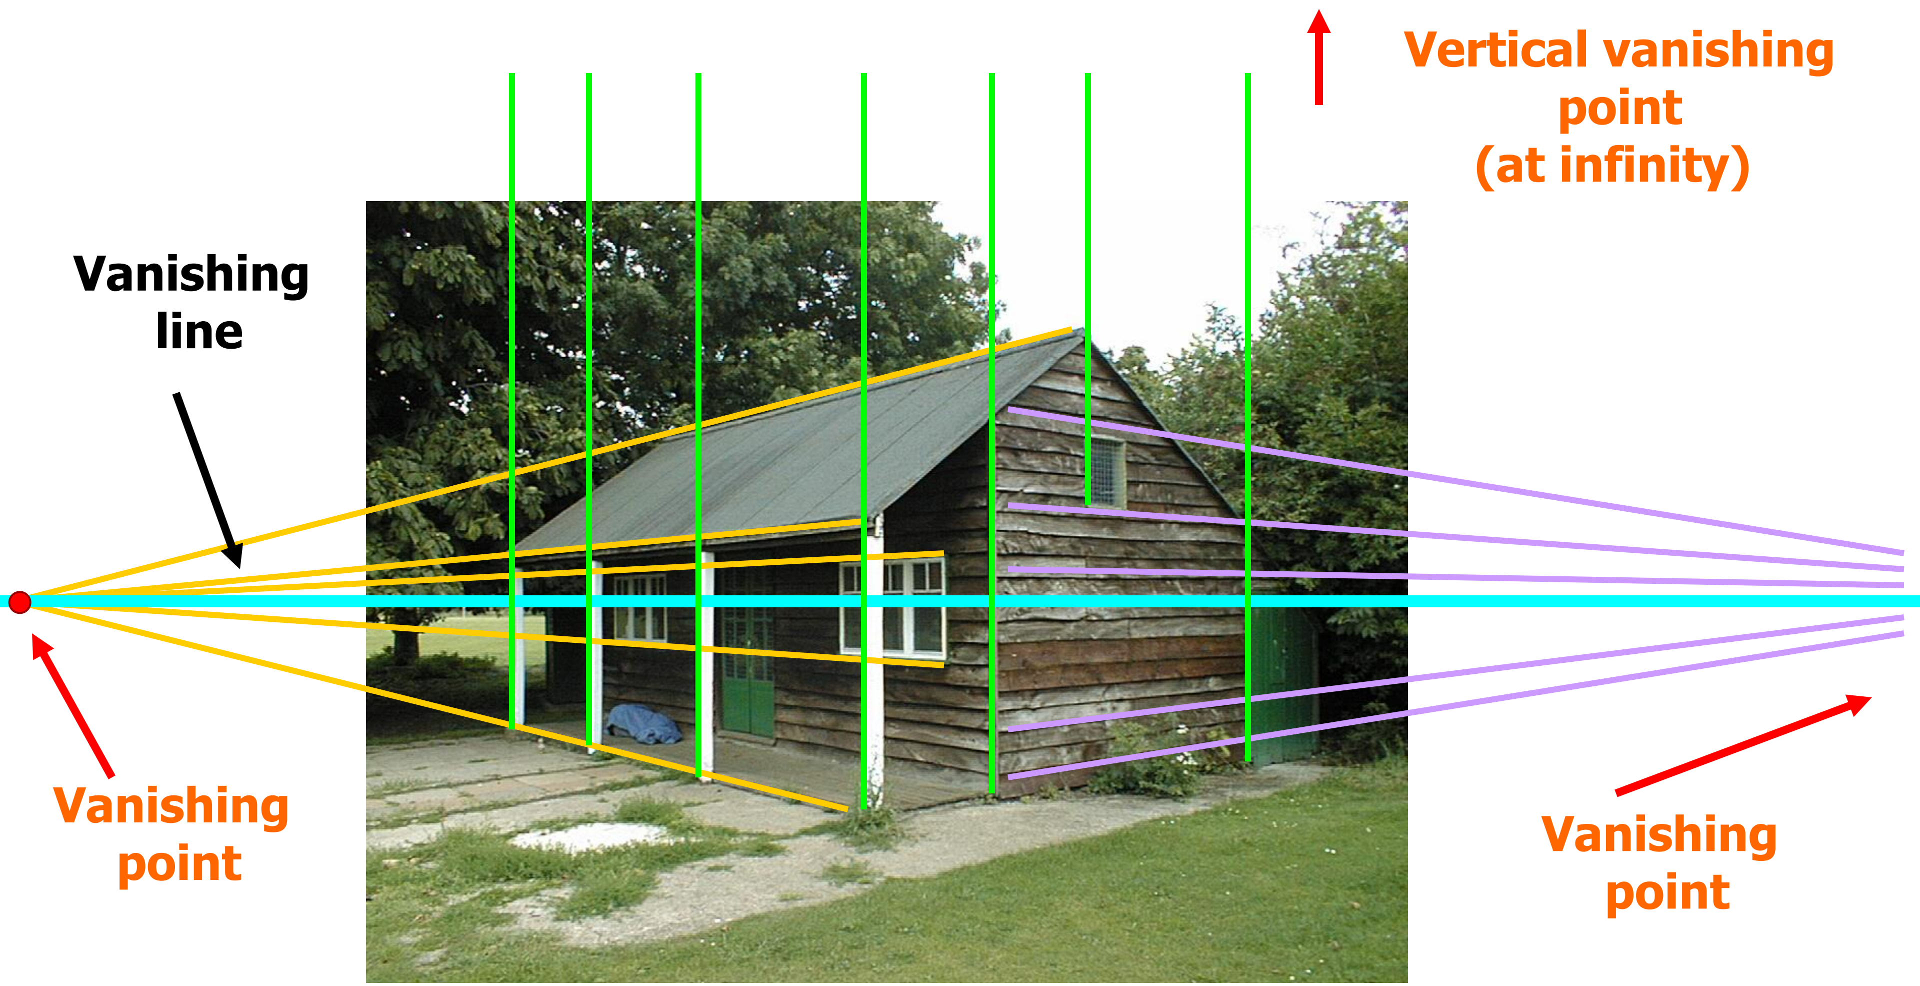

Slide from Efros, Photo from Criminisi

This class

- How can we calibrate the camera?
- How can we measure the size of objects in the world from an image?
- What about other camera properties: focal length, field of view, depth of field, aperture, f-number?
- How to do “focus stacking” to get a sharp picture of a nearby object
- How the “vertigo effect” works

How to calibrate the camera?

$$\mathbf{x} = \mathbf{K}\begin{bmatrix}\mathbf{R} & \mathbf{t}\end{bmatrix}\mathbf{X}$$

$$\begin{bmatrix} wu \\ wv \\ w \end{bmatrix} = \begin{bmatrix} * & * & * & * \\ * & * & * & * \\ * & * & * & * \end{bmatrix} \begin{bmatrix} X \\ Y \\ Z \\ 1 \end{bmatrix}$$

Calibrating the Camera

Method 1: Use an object (calibration grid) with known geometry

- Correspond image points to 3d points
- Get least squares solution (or non-linear solution)

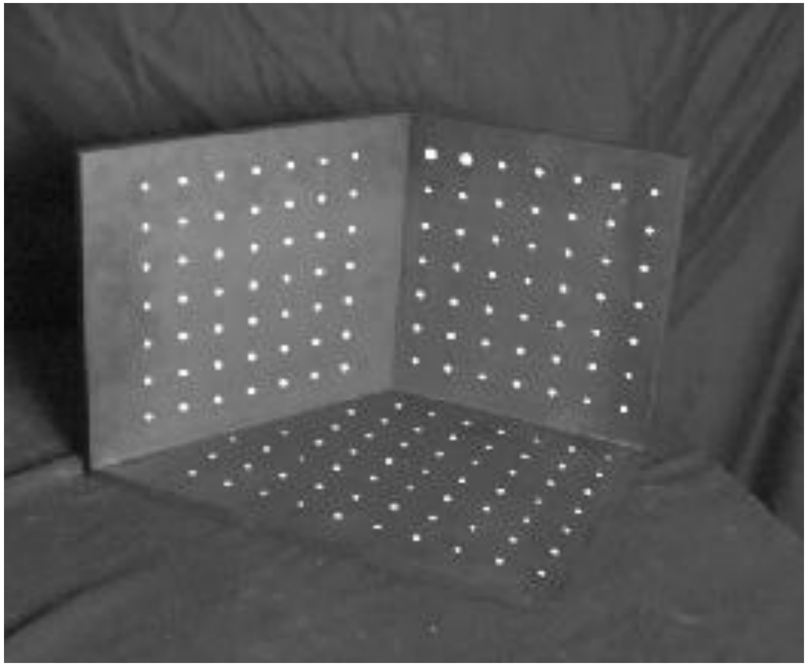

$$\begin{bmatrix} wu \\ wv \\ w \end{bmatrix} = \begin{bmatrix} m_{11} & m_{12} & m_{13} & m_{14} \\ m_{21} & m_{22} & m_{23} & m_{24} \\ m_{31} & m_{32} & m_{33} & m_{34} \end{bmatrix} \begin{bmatrix} X \\ Y \\ Z \\ 1 \end{bmatrix}$$

Calibrating the Camera

Method 2: Use vanishing points

– Find vanishing points corresponding to orthogonal directions

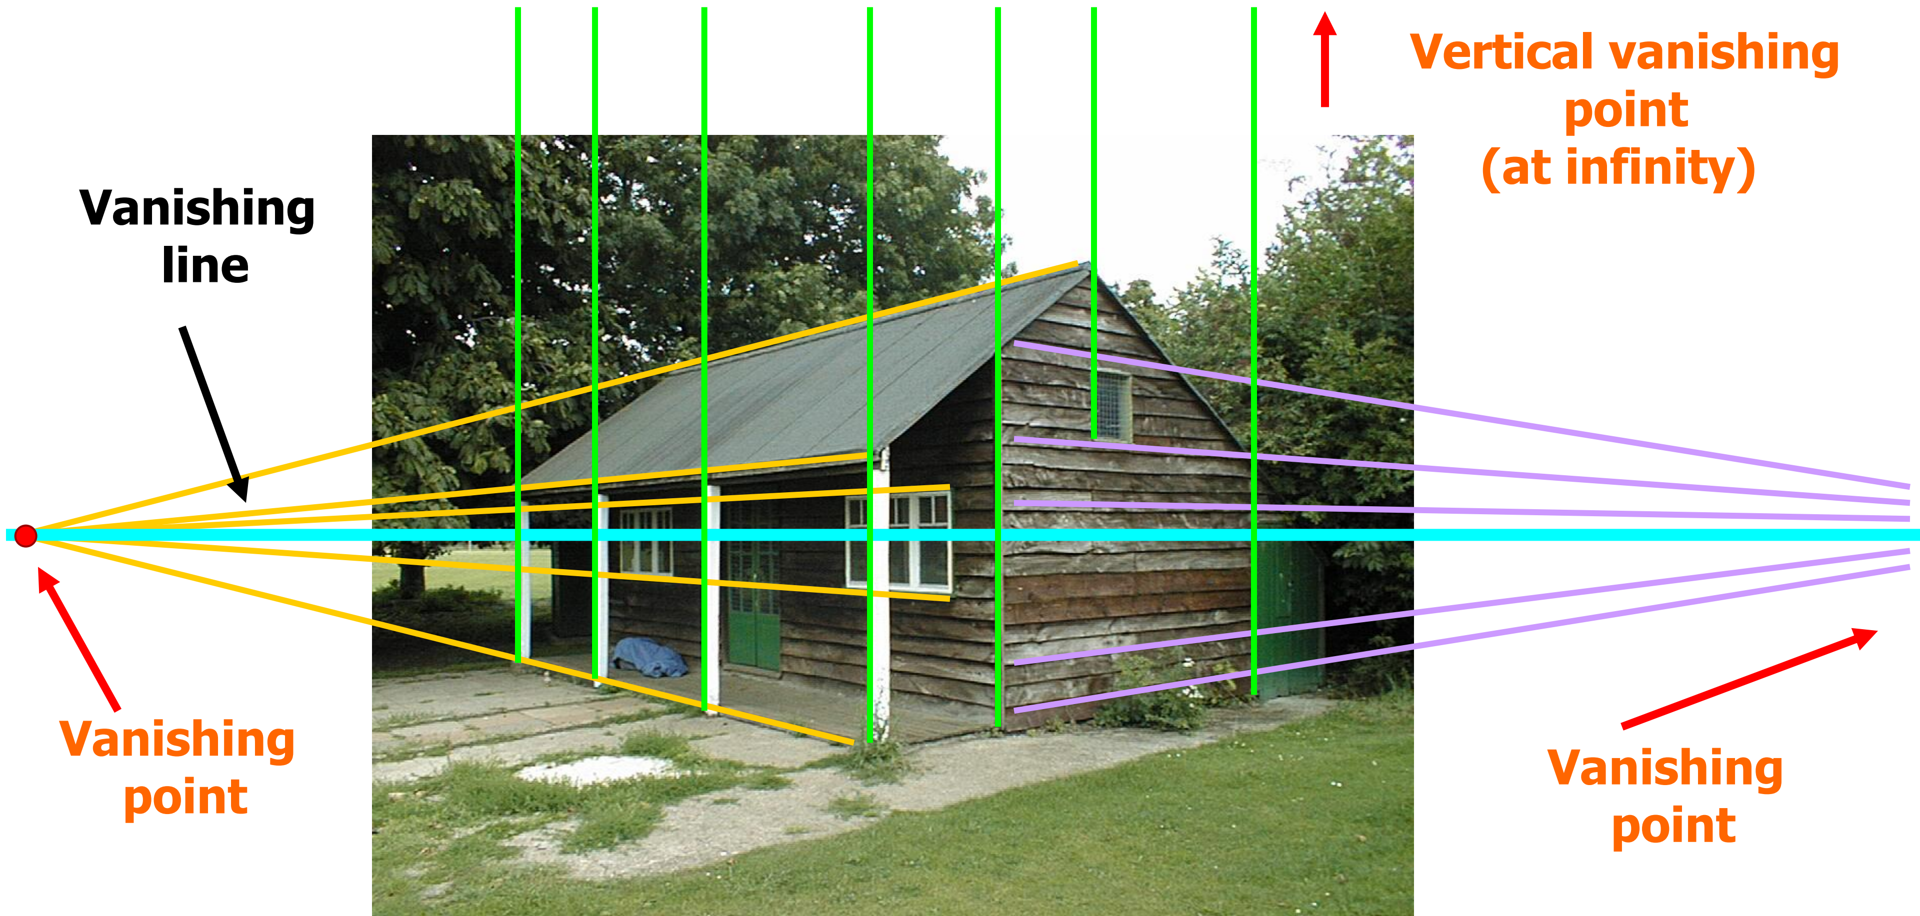

Take-home question

Suppose you have estimated three vanishing points corresponding to orthogonal directions. How can you recover the rotation matrix that is aligned with the 3D axes defined by these points?

- Assume that intrinsic matrix K has three parameters
- Remember, in homogeneous coordinates, we can write a 3d point at infinity as (X, Y, Z, 0)

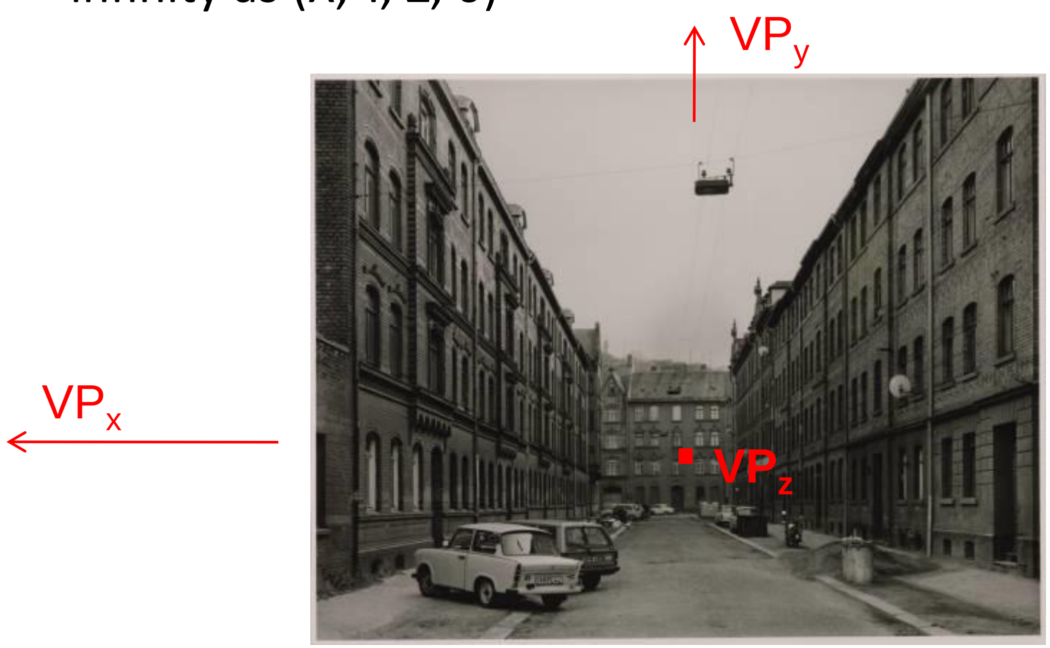

Photo from online Tate collection

How can we measure the size of 3D objects from an image?

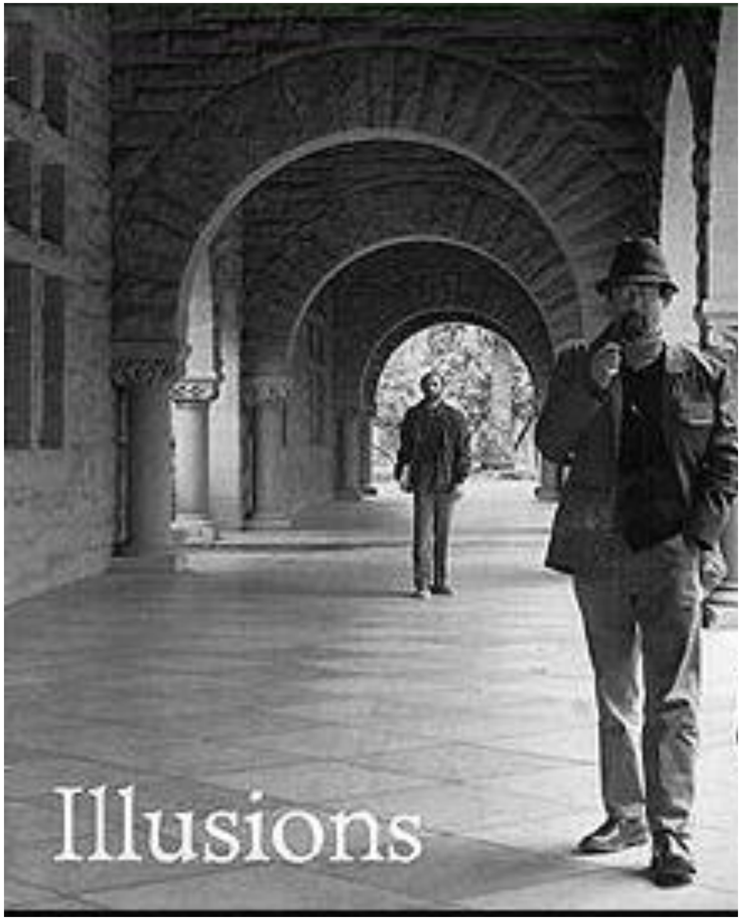

Slide by Steve Seitz

Slide by Steve Seitz

Perspective cues

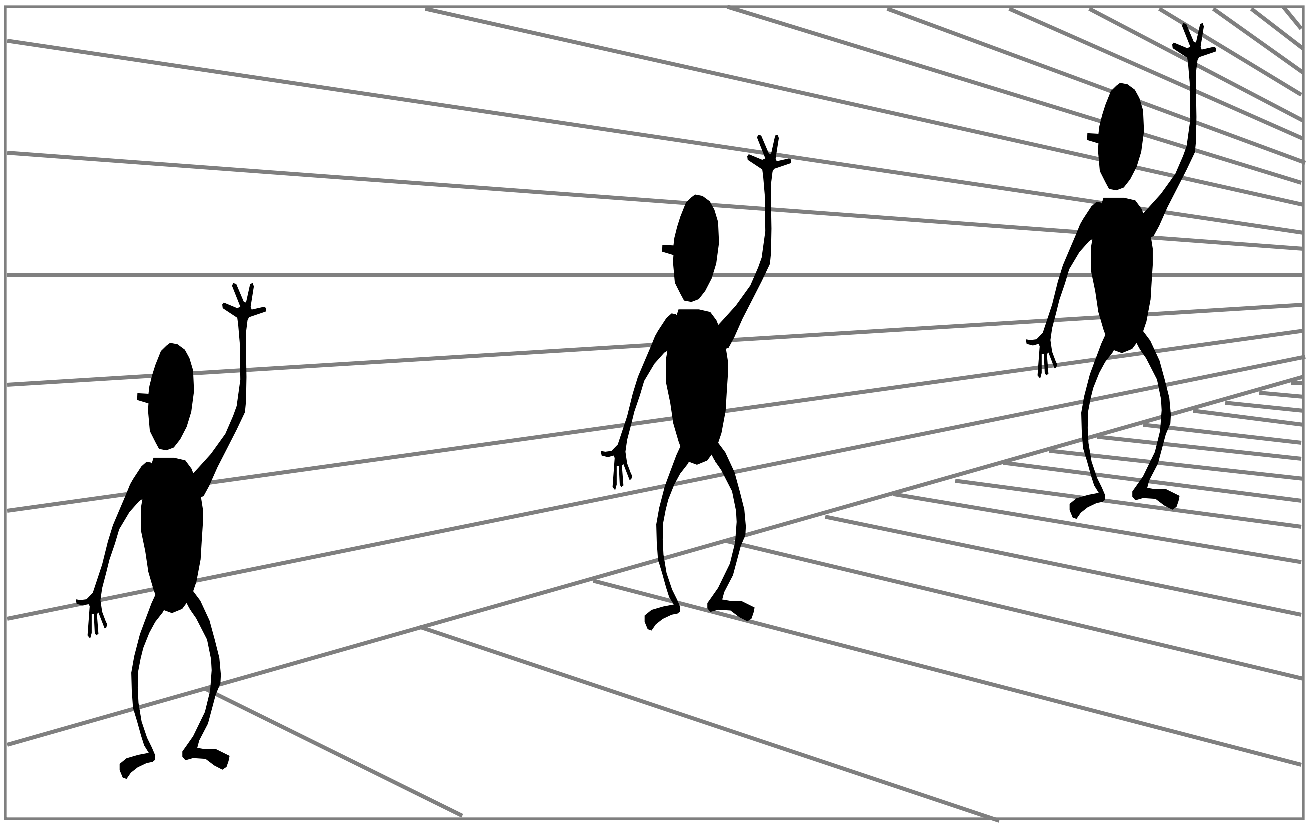

Slide by Steve Seitz

Perspective cues

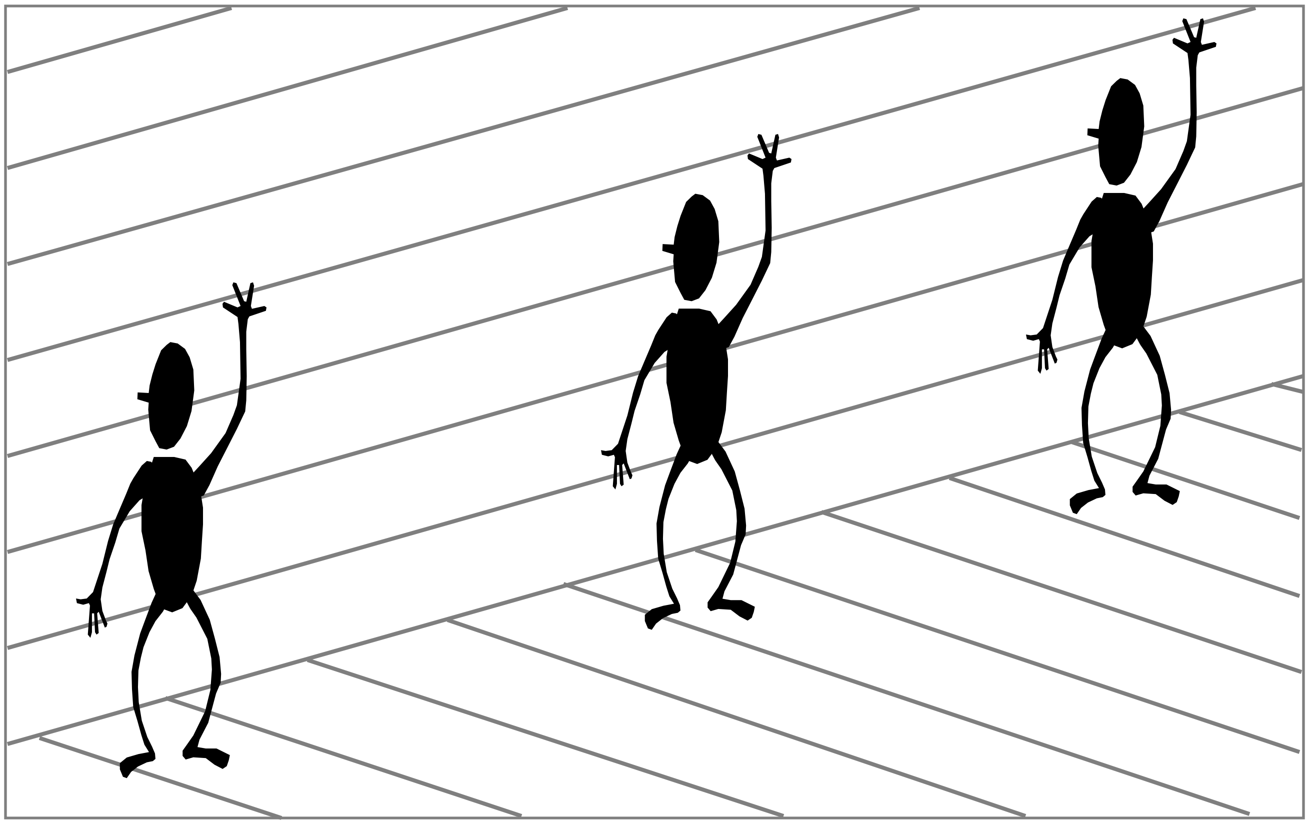

Perspective cues

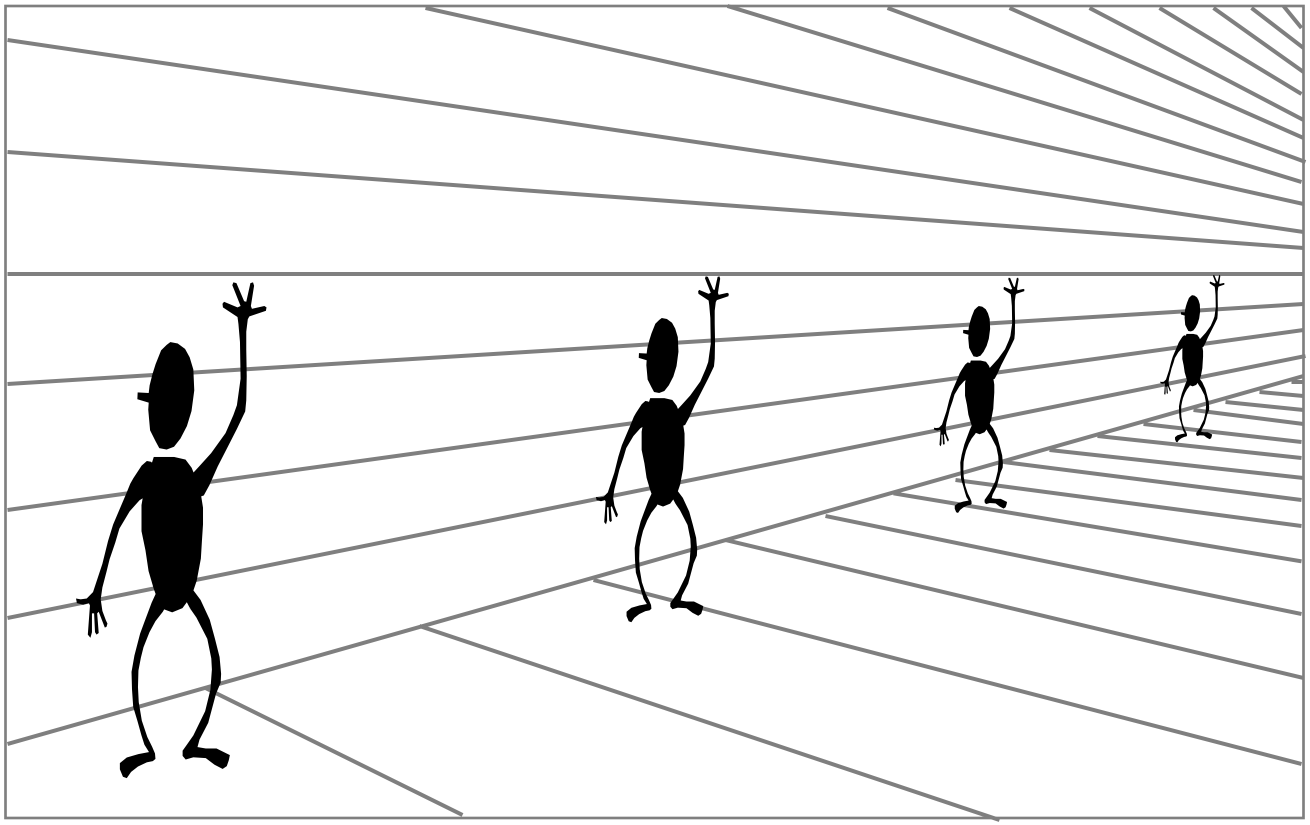

Ames Room

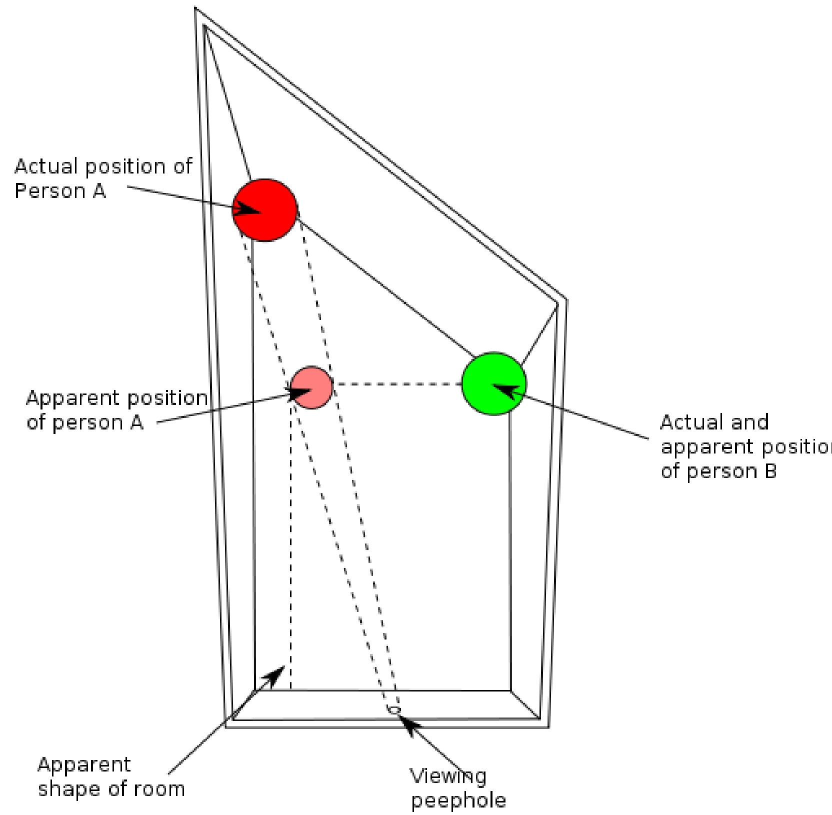
Actual position of
Person A
Apparent position
of person A
Actual and
apparent position
of person B
Apparent
shape of room
Viewing
peephole

Comparing heights

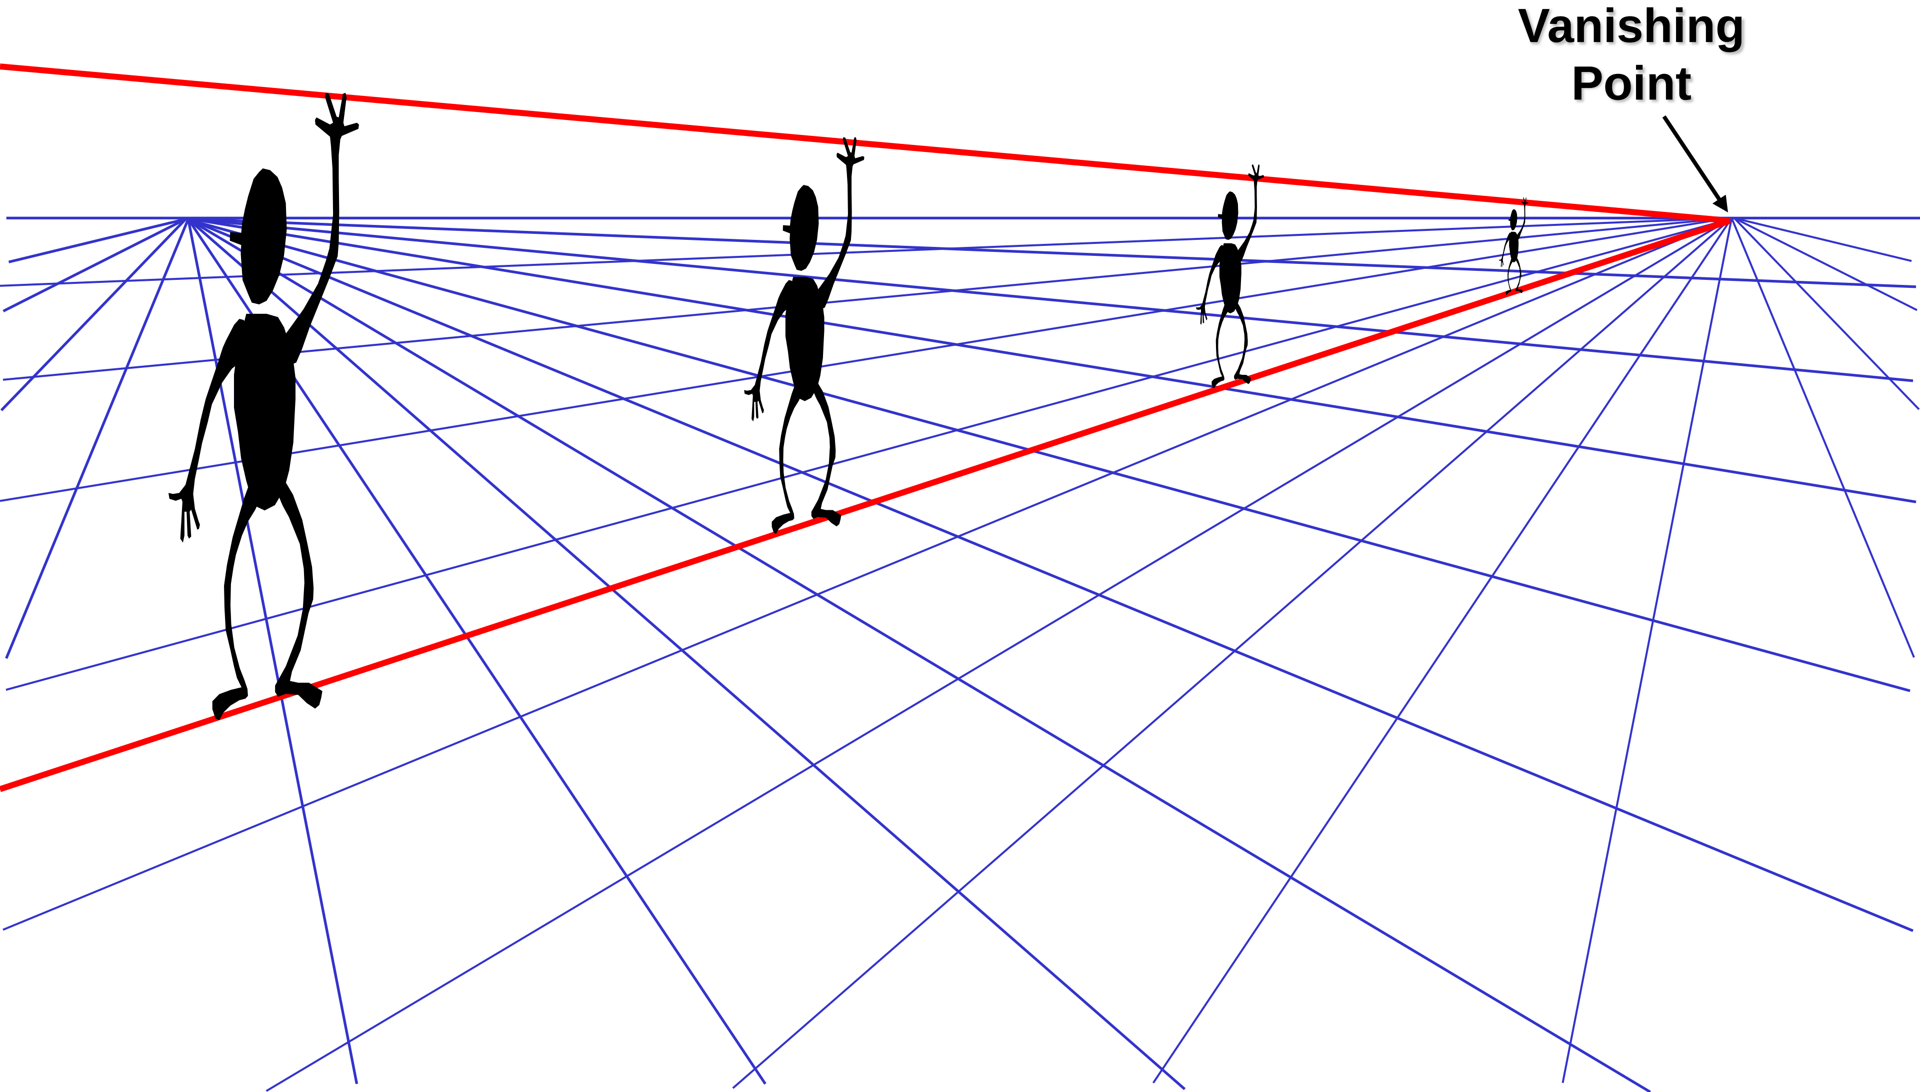

Slide by Steve Seitz

Measuring height

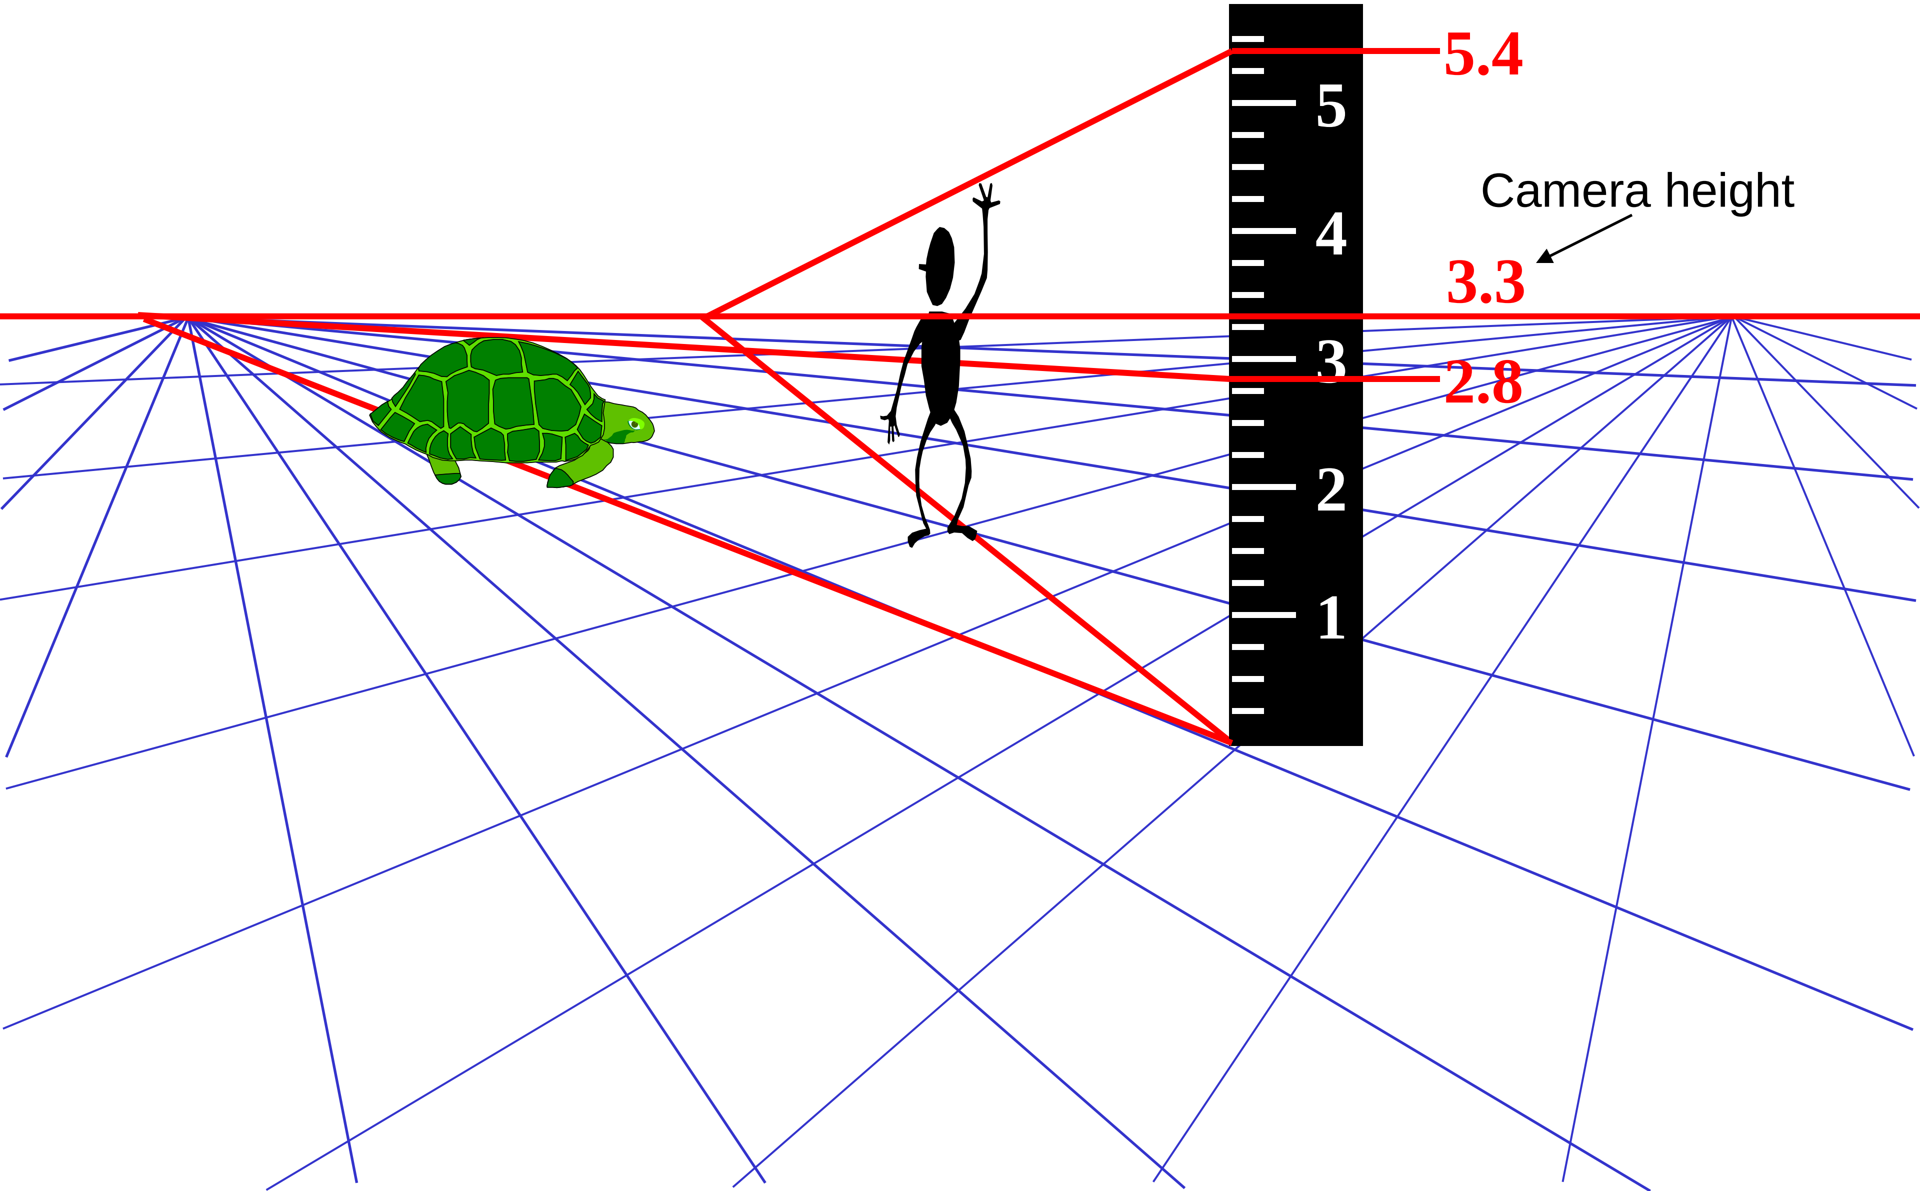

Which is higher – the camera or the parachute?

Measuring height without a giant ruler

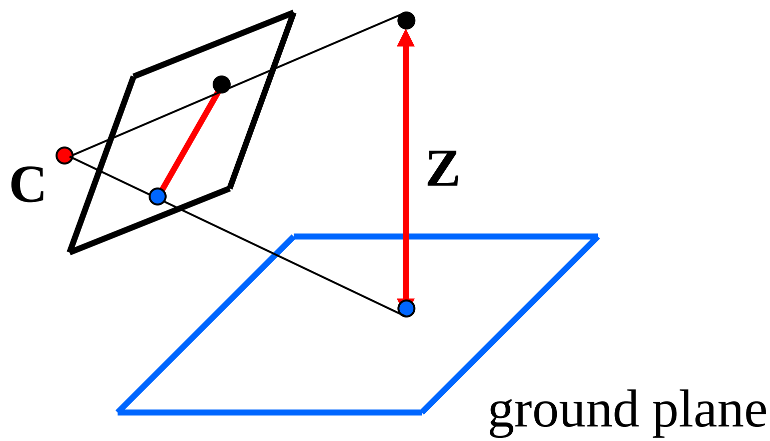

Compute Z from image measurements

- Need a reference object

Slide by Steve Seitz

The cross ratio

A Projective Invariant

- Something that does not change under projective transformations (including perspective projection)

The cross-ratio of 4 collinear points

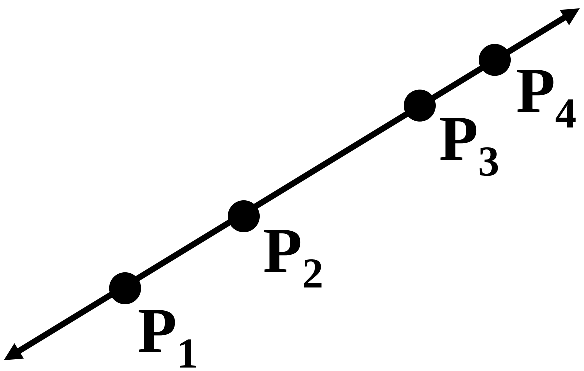

$$\frac{\|\mathbf{P}_3 - \mathbf{P}_1\| \, \|\mathbf{P}_4 - \mathbf{P}_2\|}{\|\mathbf{P}_3 - \mathbf{P}_2\| \, \|\mathbf{P}_4 - \mathbf{P}_1\|}$$

$$\mathbf{P}_i = \begin{bmatrix} X_i \\ Y_i \\ Z_i \\ 1 \end{bmatrix}$$

Can permute the point ordering

$$\frac{\|\mathbf{P}_1 - \mathbf{P}_3\| \, \|\mathbf{P}_4 - \mathbf{P}_2\|}{\|\mathbf{P}_1 - \mathbf{P}_2\| \, \|\mathbf{P}_4 - \mathbf{P}_3\|}$$

- 4! = 24 different orders (but only 6 distinct values)

This is the fundamental invariant of projective geometry

Measuring height

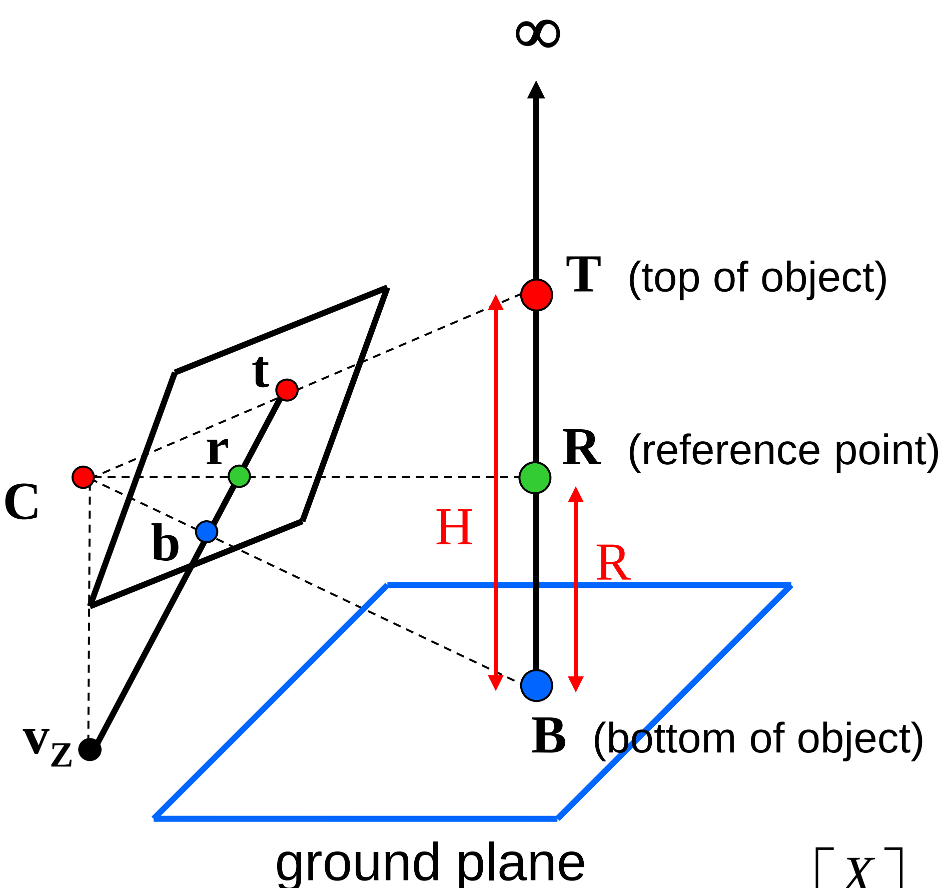

$$\frac{\|\mathbf{B}-\mathbf{T}\|\,\|\infty-\mathbf{R}\|}{\|\mathbf{B}-\mathbf{R}\|\,\|\infty-\mathbf{T}\|}=\frac{H}{R}$$

scene cross ratio

$$\frac{\|\mathbf{b}-\mathbf{t}\|\,\|\mathbf{v}_Z-\mathbf{r}\|}{\|\mathbf{b}-\mathbf{r}\|\,\|\mathbf{v}_Z-\mathbf{t}\|}=\frac{H}{R}$$

image cross ratio

scene points represented as $\mathbf{P}=\begin{bmatrix}X\\Y\\Z\\1\end{bmatrix}$ image points as $\mathbf{p}=\begin{bmatrix}x\\y\\1\end{bmatrix}$

Slide by Steve Seitz

Measuring height

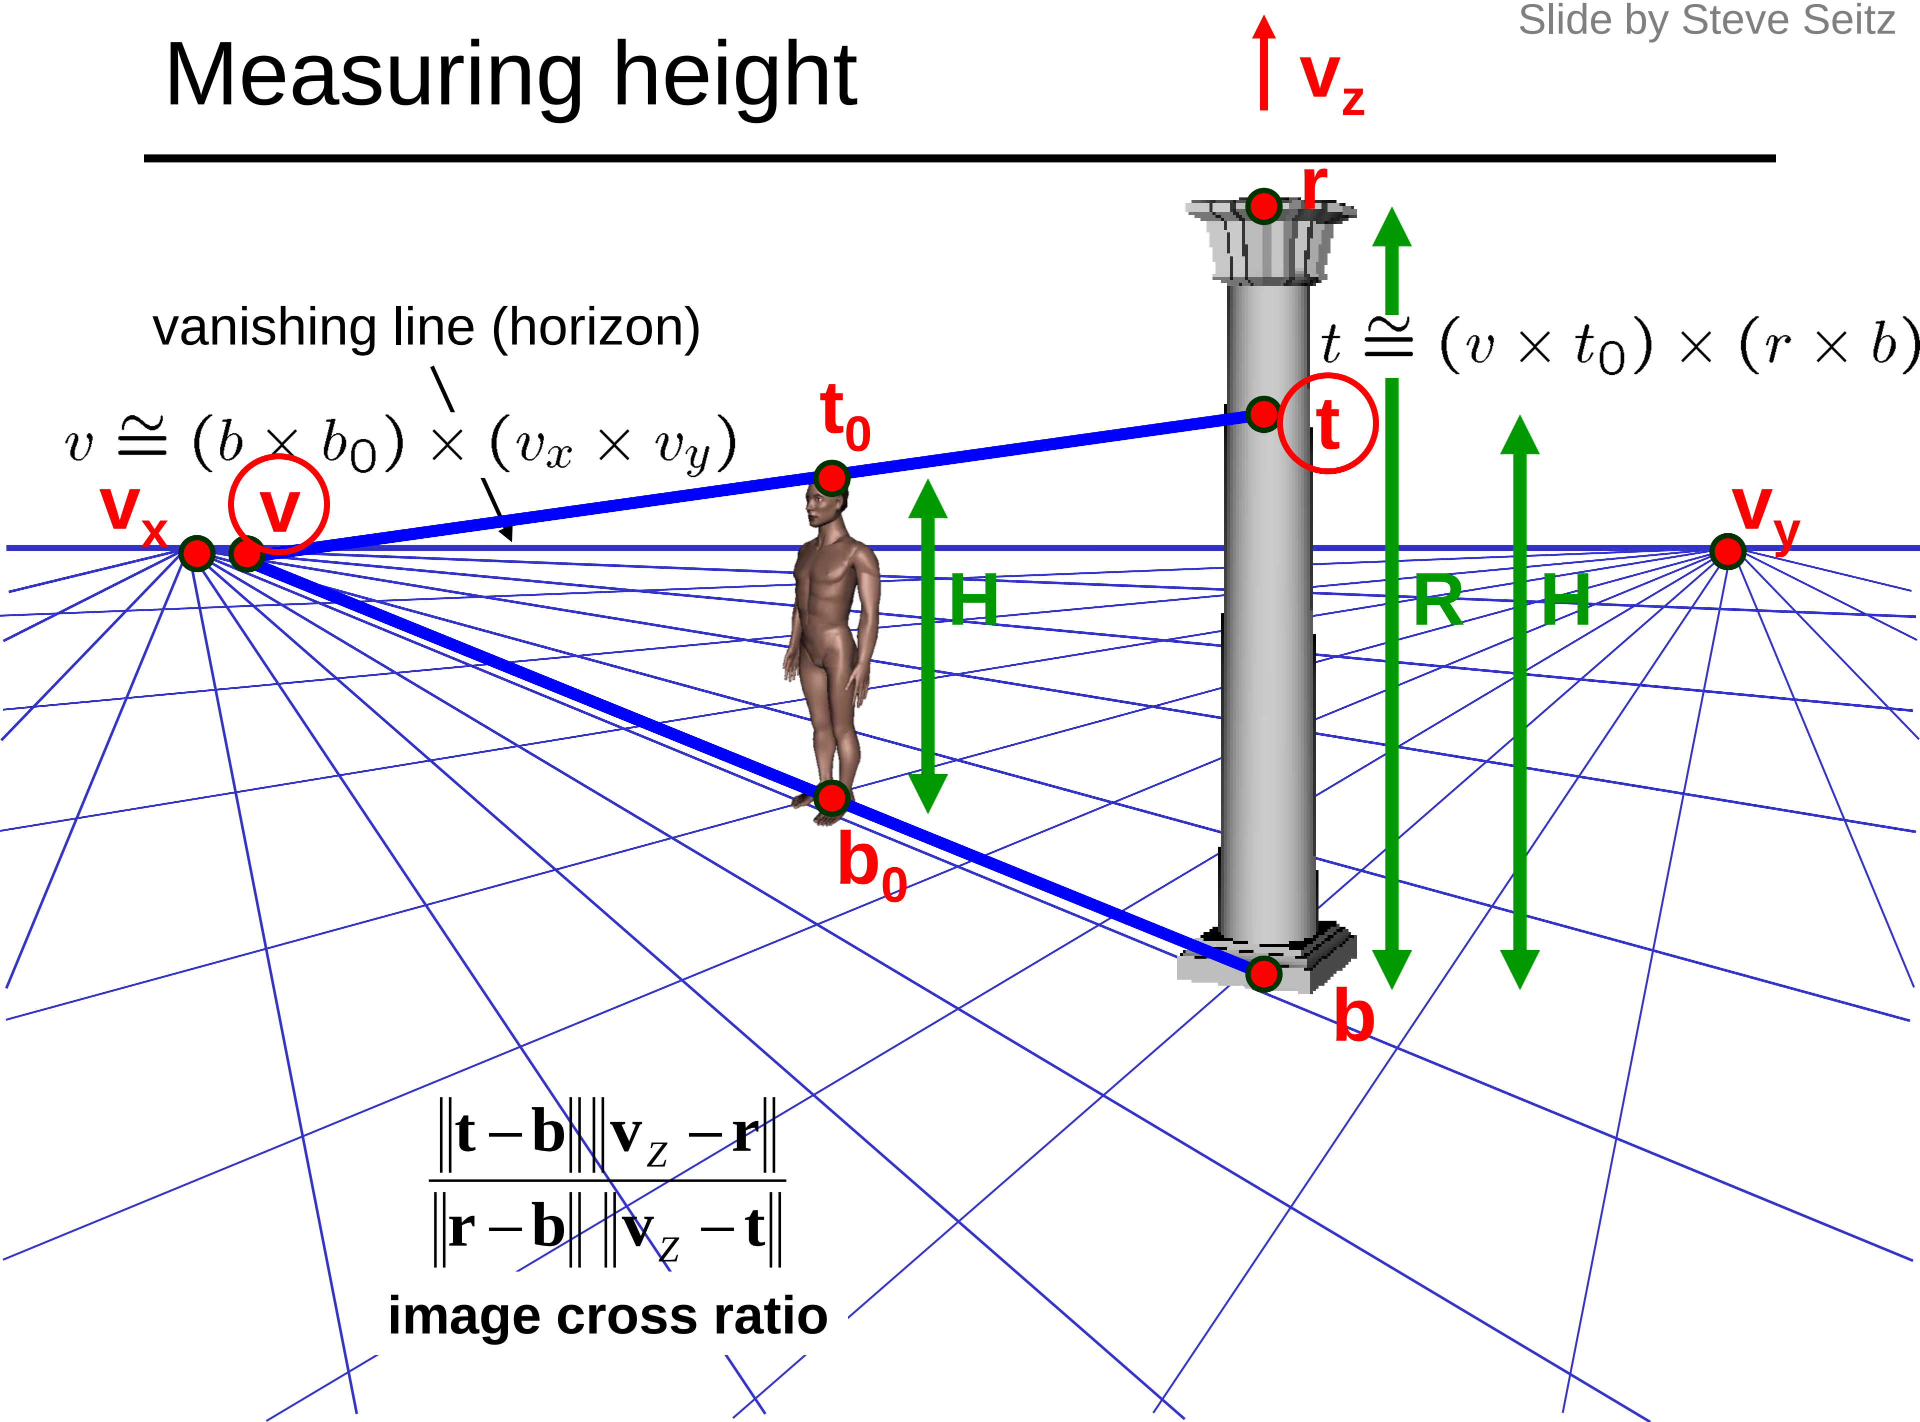

Measuring height

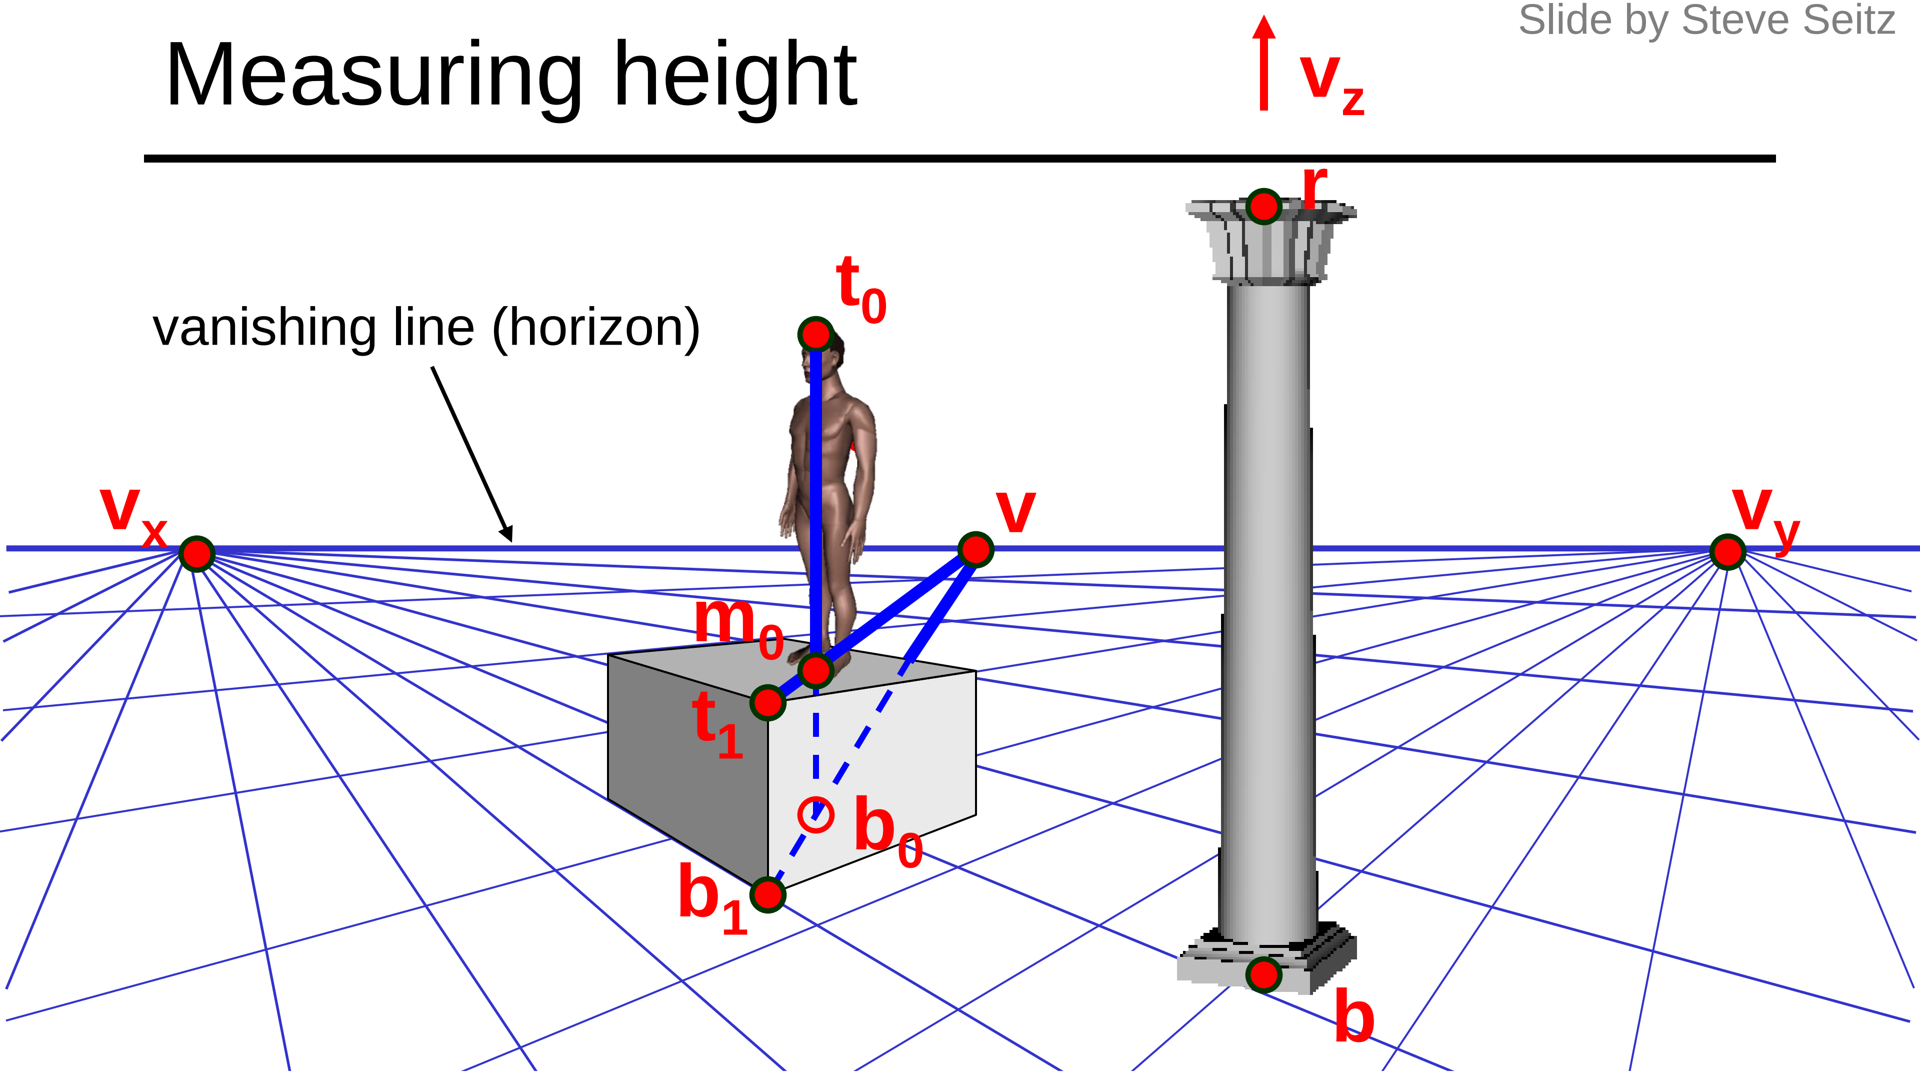

What if the point on the ground plane $\mathbf{b_0}$ is not known?

- Here the guy is standing on the box, height of box is known
- Use one side of the box to help find $\mathbf{b_0}$ as shown above

Take-home question

Assume that the camera height is 5 ft.

- What is the height of the man?
- What is the height of the building?

Beyond the pinhole: What about focus, aperture, DOF, FOV, etc?

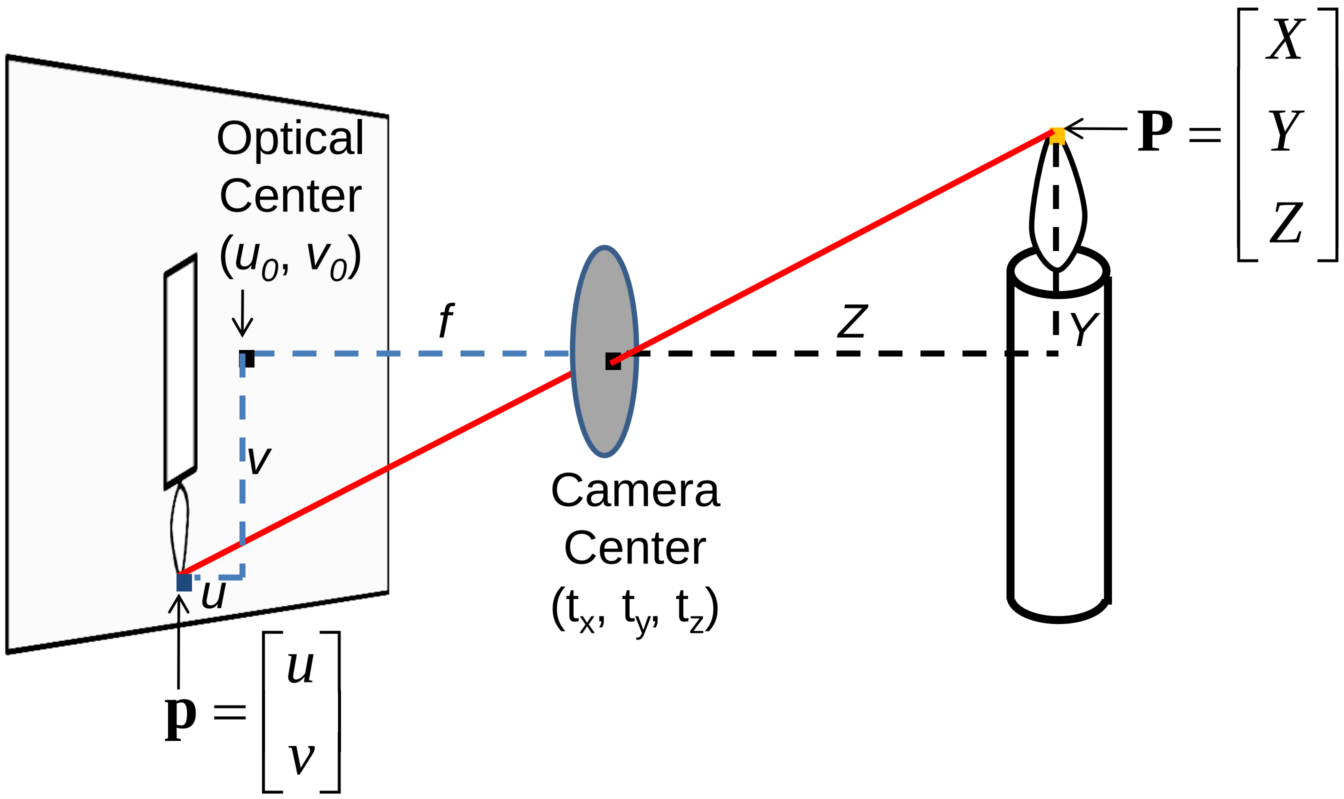

Adding a lens

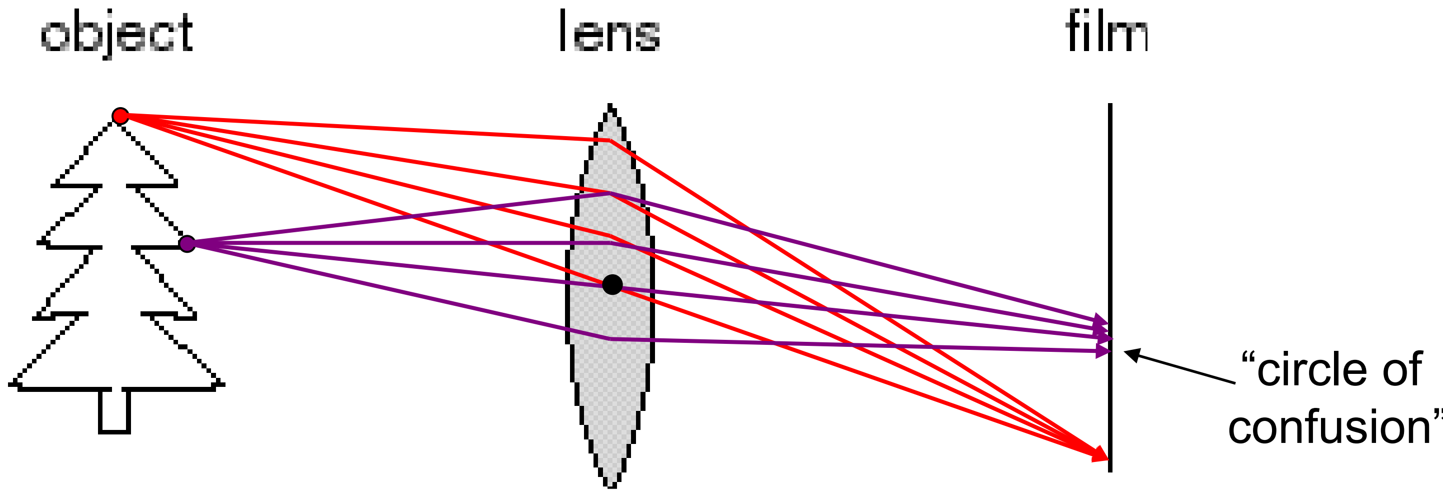

- A lens focuses light onto the film
  - There is a specific distance at which objects are "in focus"
    - other points project to a "circle of confusion" in the image
  - Changing the shape of the lens changes this distance

Focal length, aperture, depth of field

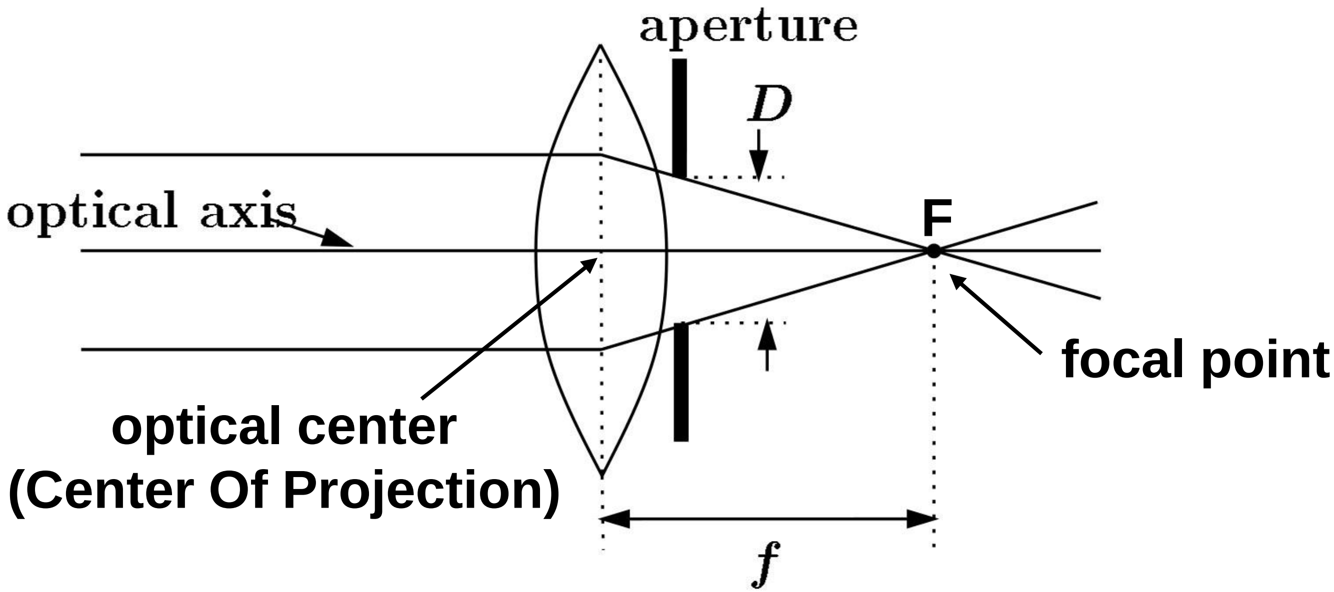

A lens focuses parallel rays onto a single focal point

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

Slide source: Seitz

The eye

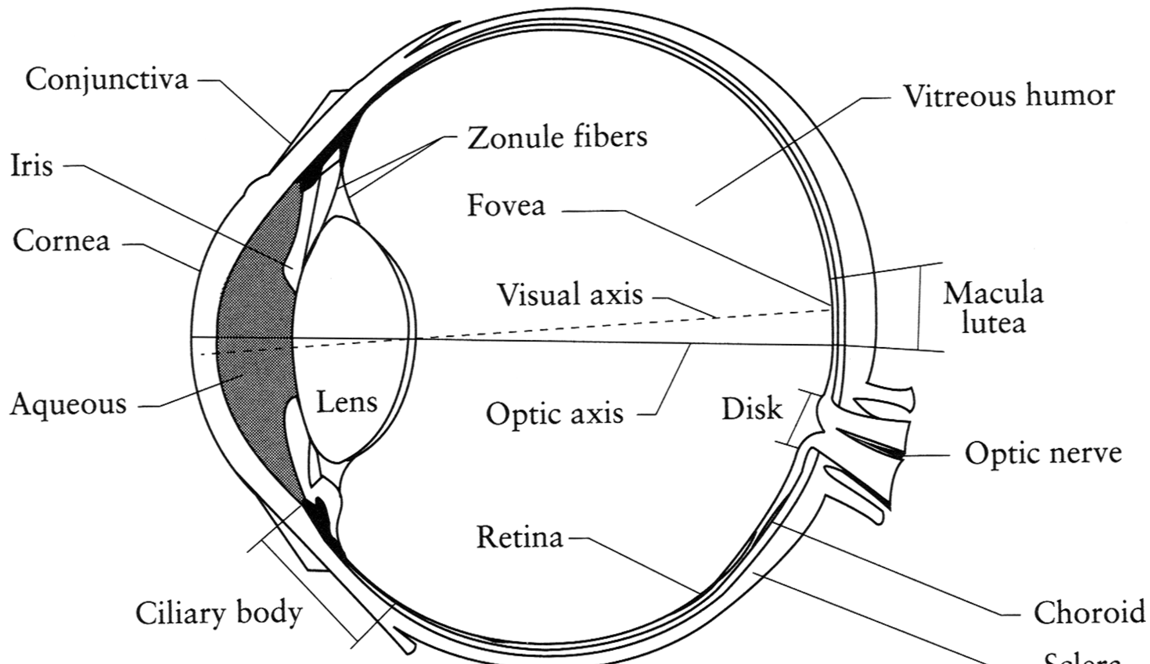

- The human eye is a camera
  - **Iris** - colored annulus with radial muscles
  - **Pupil** - the hole (aperture) whose size is controlled by the iris

Focus with lenses

Distance to object

Distance to sensor

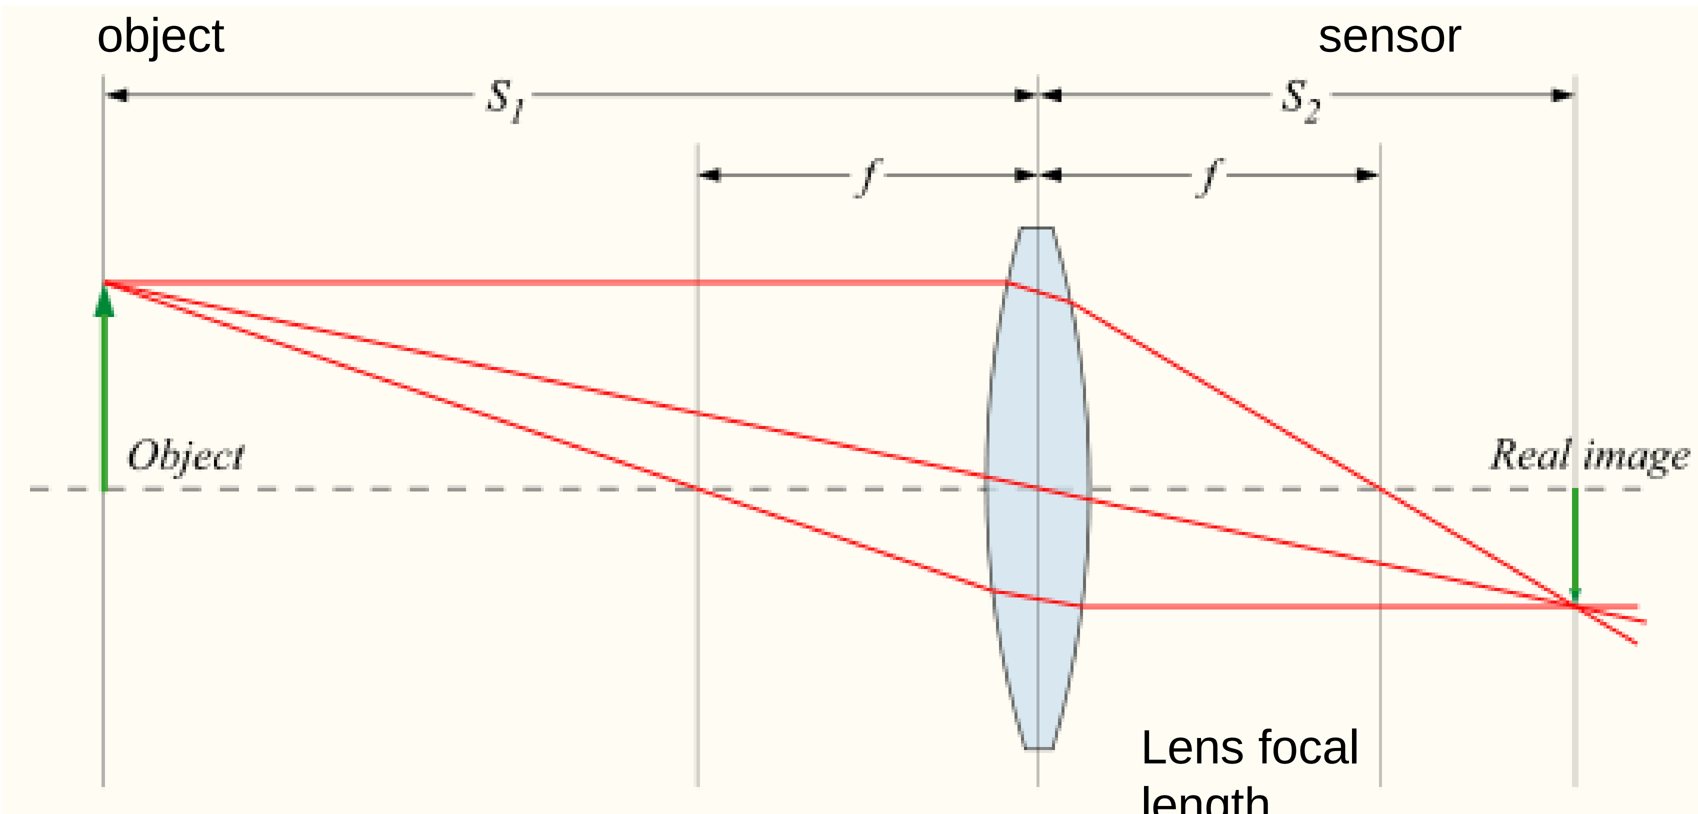

Lens focal length

Equation for objects in focus

$$\frac{1}{S_1} + \frac{1}{S_2} = \frac{1}{f}$$

Source: http://en.wikipedia.org/wiki/File:Lens3.svg

Slide source: Seitz

The aperture and depth of field

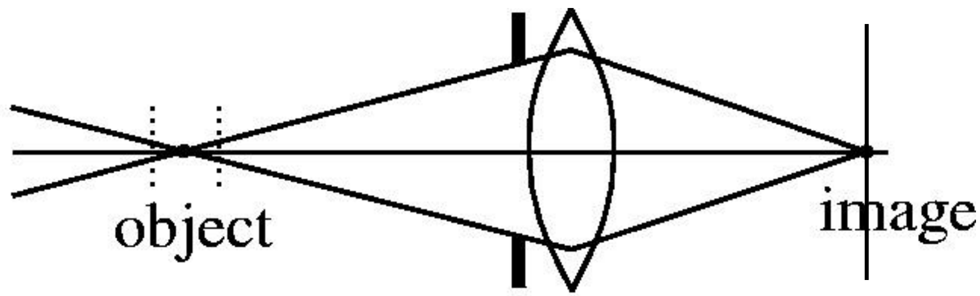

f / 5.6

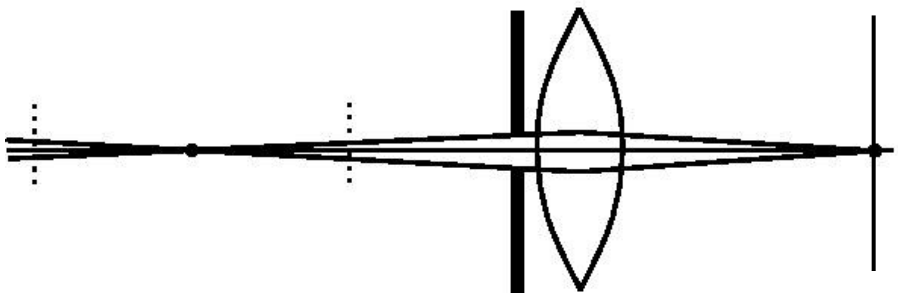

f / 32

Changing the aperture size or focusing distance affects depth of field

- f-number (f/#) =focal_length / aperture_diameter (e.g., f/16 means that the focal length is 16 times the diameter)
- When you change the f-number, you are changing the aperture

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

Varying the aperture

Large aperture = small DOF

Small aperture = large DOF

Shrinking the aperture

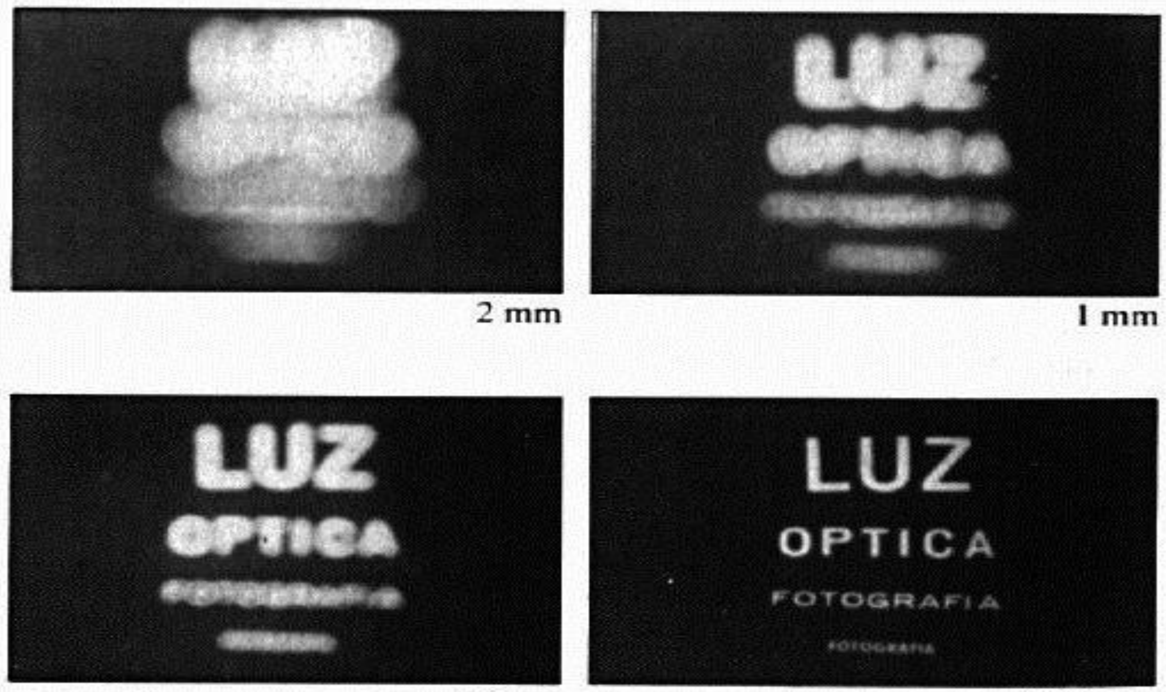

2 mm 1 mm 0.6mm 0.35 mm

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

Slide by Steve Seitz

Shrinking the aperture

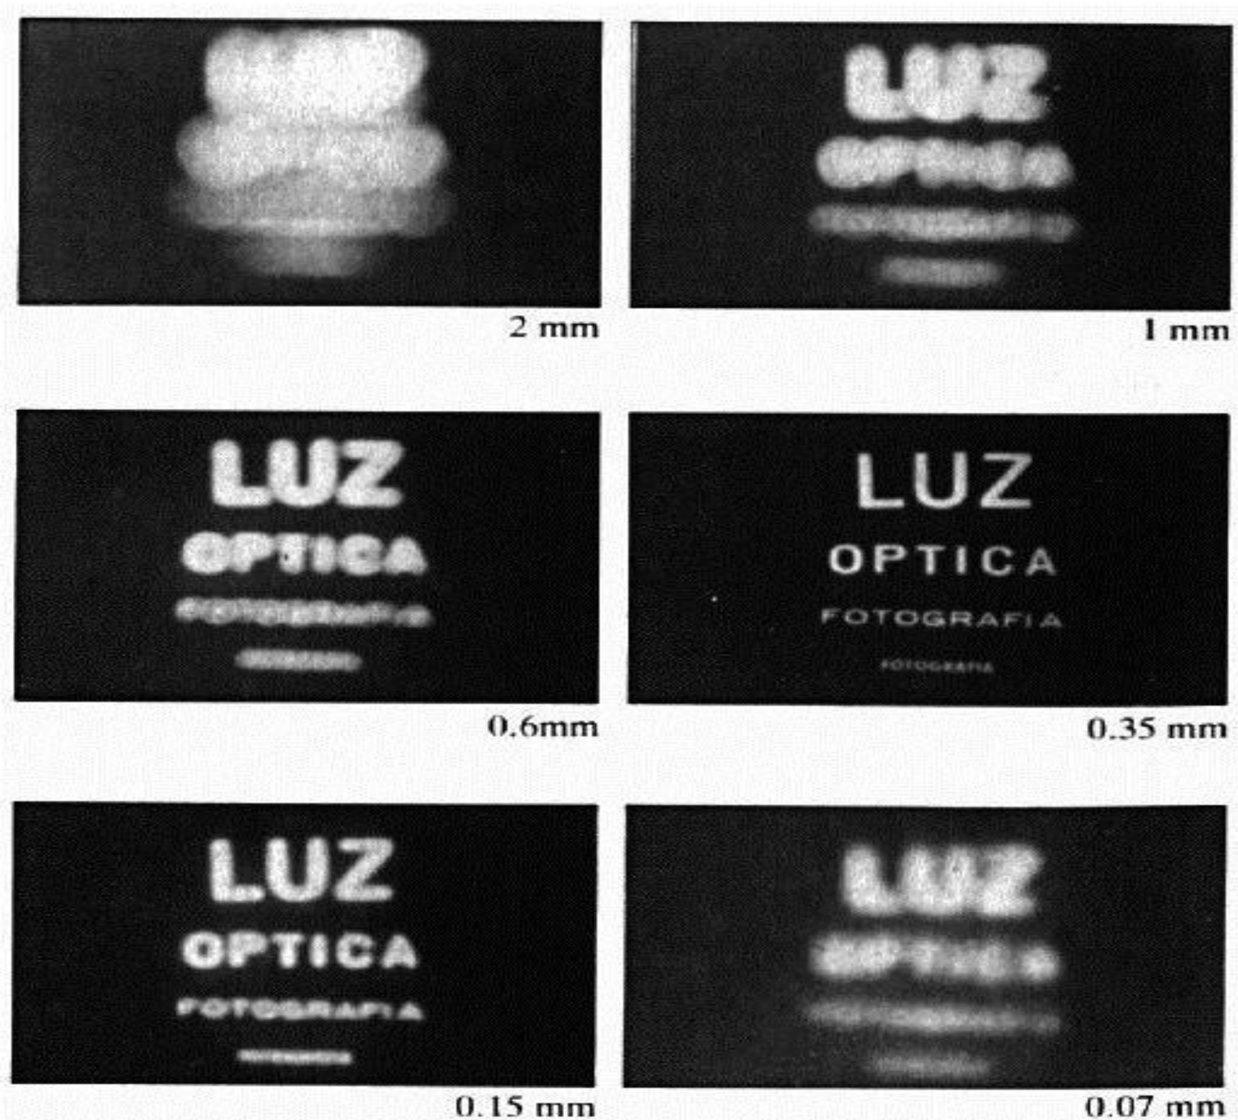

Slide by Steve Seitz

The Photographer's Great Compromise

What we want	How we get it	Cost
More spatial resolution	Increase focal length	Light, FOV
	Decrease focal length	DOF
Broader field of view		
	Decrease aperture	Light
More depth of field	Increase aperture	DOF
	Shorten exposure	Light
More temporal resolution	Lengthen exposure	Temporal Res
More light		

Difficulty in macro (close-up) photography

- For close objects, we have a small relative DOF
- Can only shrink aperture so far

How to get both bugs in focus?

Solution: Focus stacking

1. Take pictures with varying focal length



Example from
http://www.wonderfulphotos.com/articles/macro/focus_stacking/

Solution: Focus stacking

1. Take pictures with varying focal length
2. Combine

Focus stacking

http://www.wonderfulphotos.com/articles/macro/focus_stacking/

Focus stacking

How to combine?

Web answer: With software (Photoshop, CombineZM)

How to do it automatically?

Focus stacking

How to combine?

1. Align images (e.g., using corresponding points)
2. Two ideas
   a) Mask regions by hand and combine with pyramid blend
   b) Gradient domain fusion (mixed gradient) without masking

Automatic solution would make a very interesting final project

Recommended Reading:

http://www.digital-photography-school.com/an-introduction-to-focus-stacking

http://www.zen20934.zen.co.uk/photography/Workflow.htm#Focus%20Stacking

Relation between field of view and focal length

Field of view (angle width)

Film/Sensor Width

$$fov = 2\tan^{-1}\frac{d}{2f}$$

Focal length

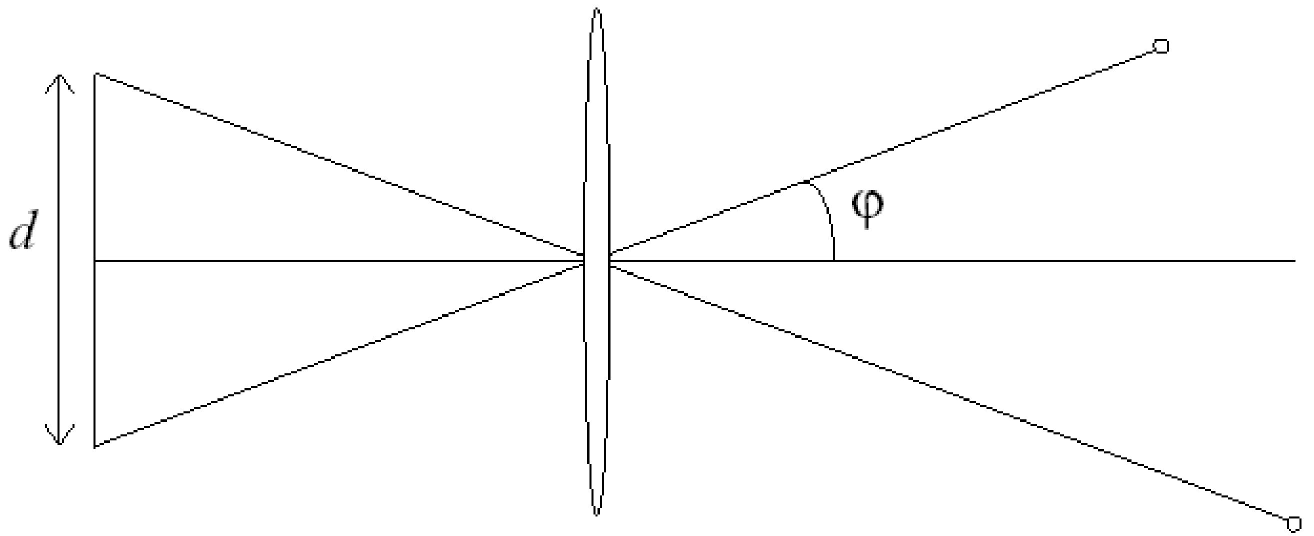

Dolly Zoom or “Vertigo Effect”

http://www.youtube.com/watch?v=NB4bikrNzMk

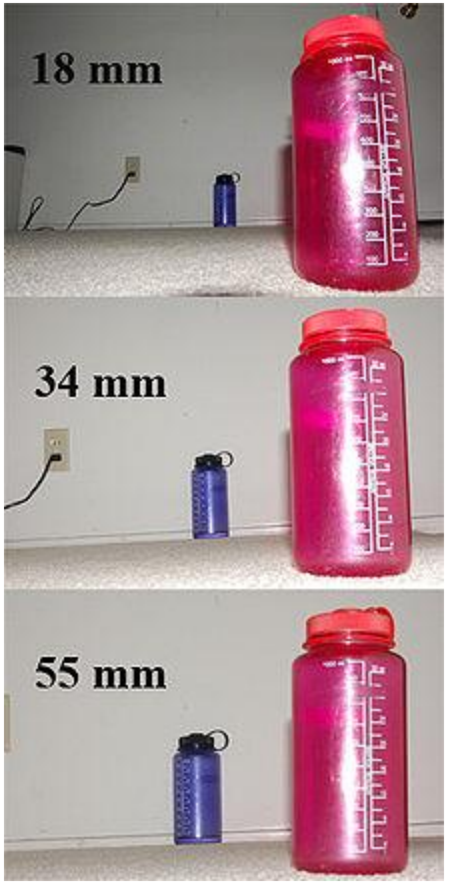

How is this done?

Zoom in while moving away

http://en.wikipedia.org/wiki/Focal_length

Dolly zoom (or "Vertigo effect")

Field of view (angle width)

Film/Sensor Width

$$fov = 2\tan^{-1}\frac{d}{2f}$$

Focal length

width of object

$$2\tan\frac{fov}{2} = \frac{width}{distance}$$

Distance between object and camera

Things to remember

- Can calibrate using grid or VP

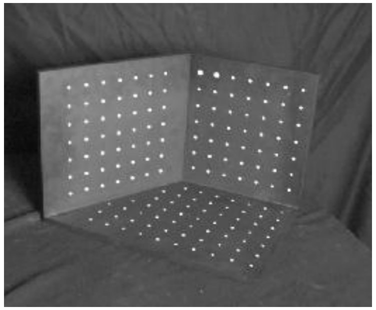

- Can measure relative sizes using VP

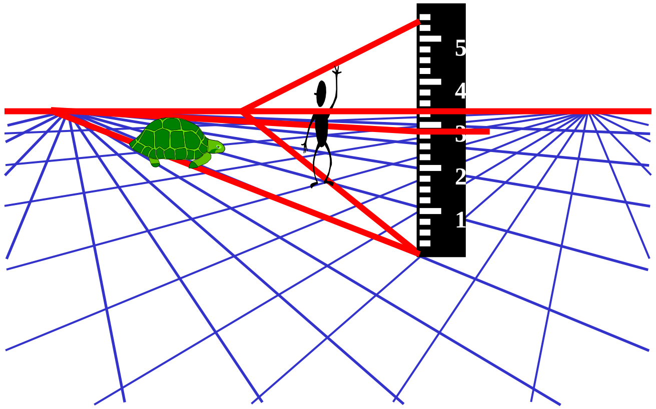

- Effects of focal length, aperture + tricks

Next class

- Go over take-home questions from today
- Single-view 3D Reconstruction