Lecture 2: Image formation and capture

COS 429: Computer Vision

PRINCETON UNIVERSITY

Slides adapted from: Szymon Rusinkiewicz, Jia Deng

Pinhole camera: overview

Let's design a camera

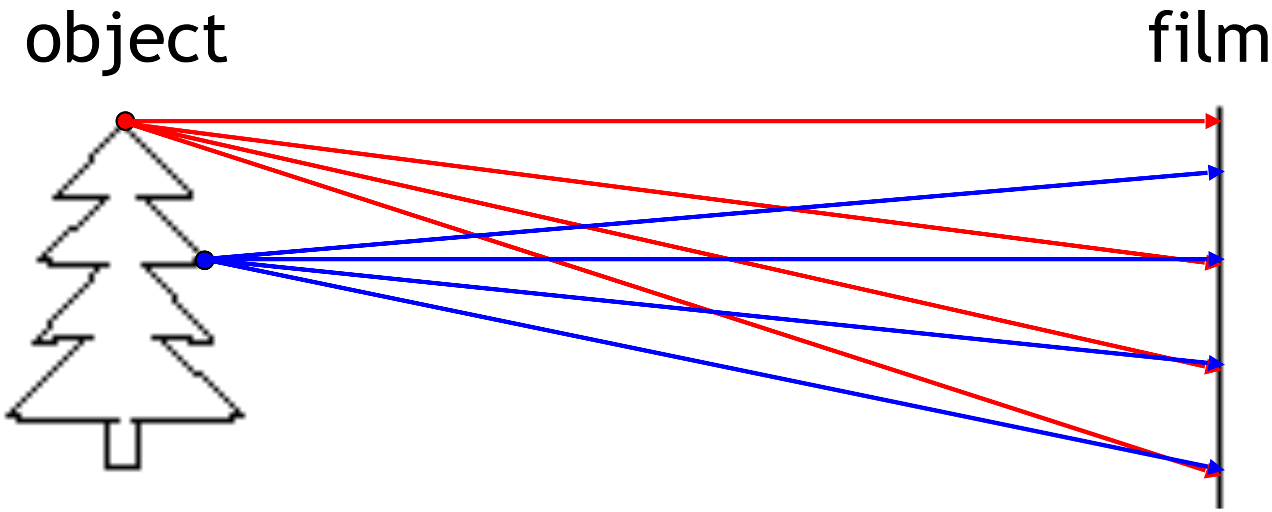

Idea 1: put a piece of film in front of an object

Do we get a reasonable image?

Slide by Steve Seitz

Let's design a camera

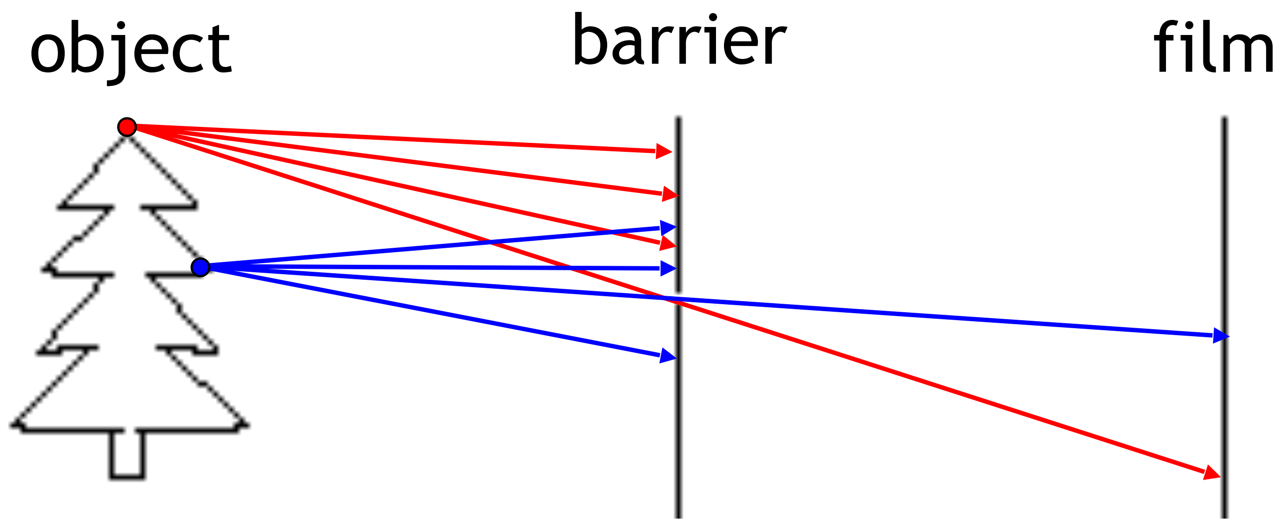

Add a barrier to block off most of the rays

Slide by Steve Seitz

Pinhole camera

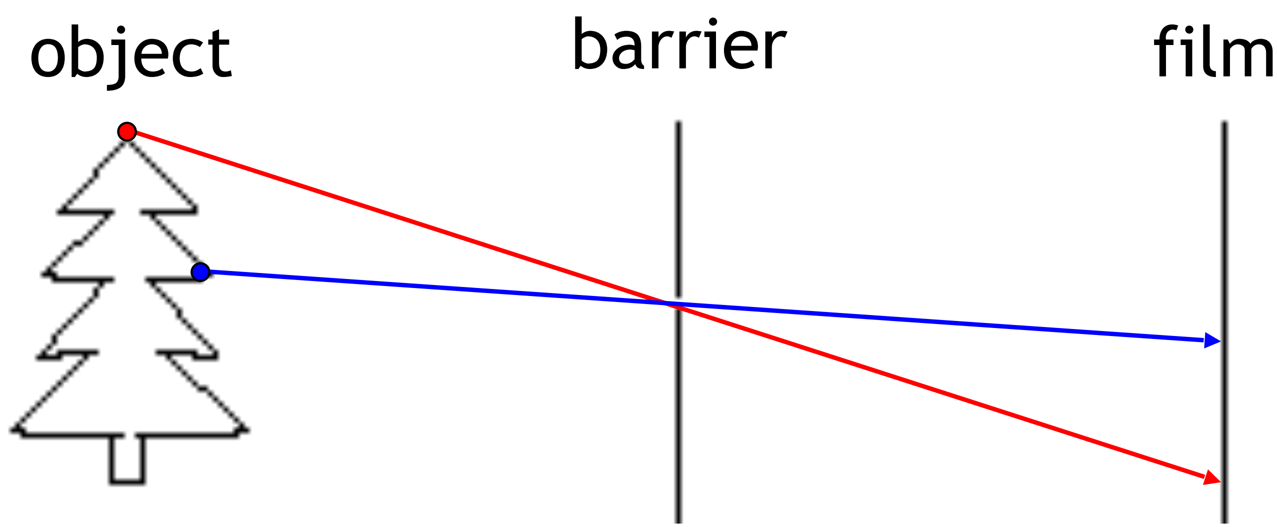

- Captures **pencil of rays** – all rays through a single point: **aperture, center of projection, optical center, focal point, camera center**
- The image is formed on the **image plane**

Slide by Steve Seitz

Camera obscura (Latin for “Dark Chamber”)

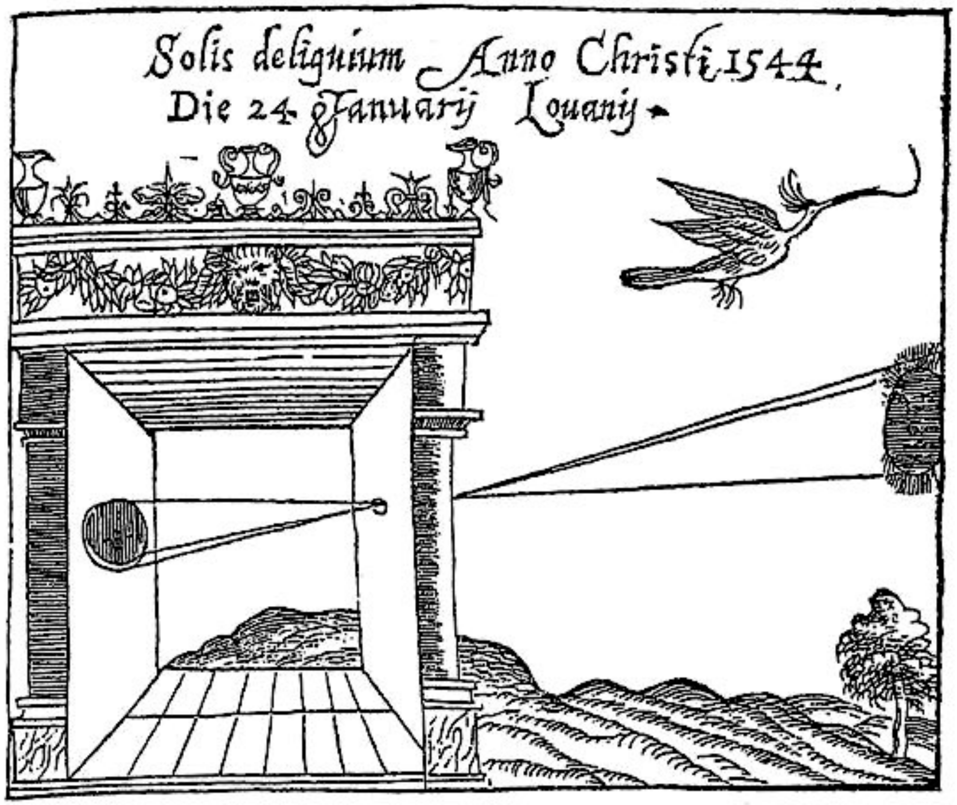

Gemma Frisius, 1558

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

Source: A. Efros

Turning a room into a camera obscura

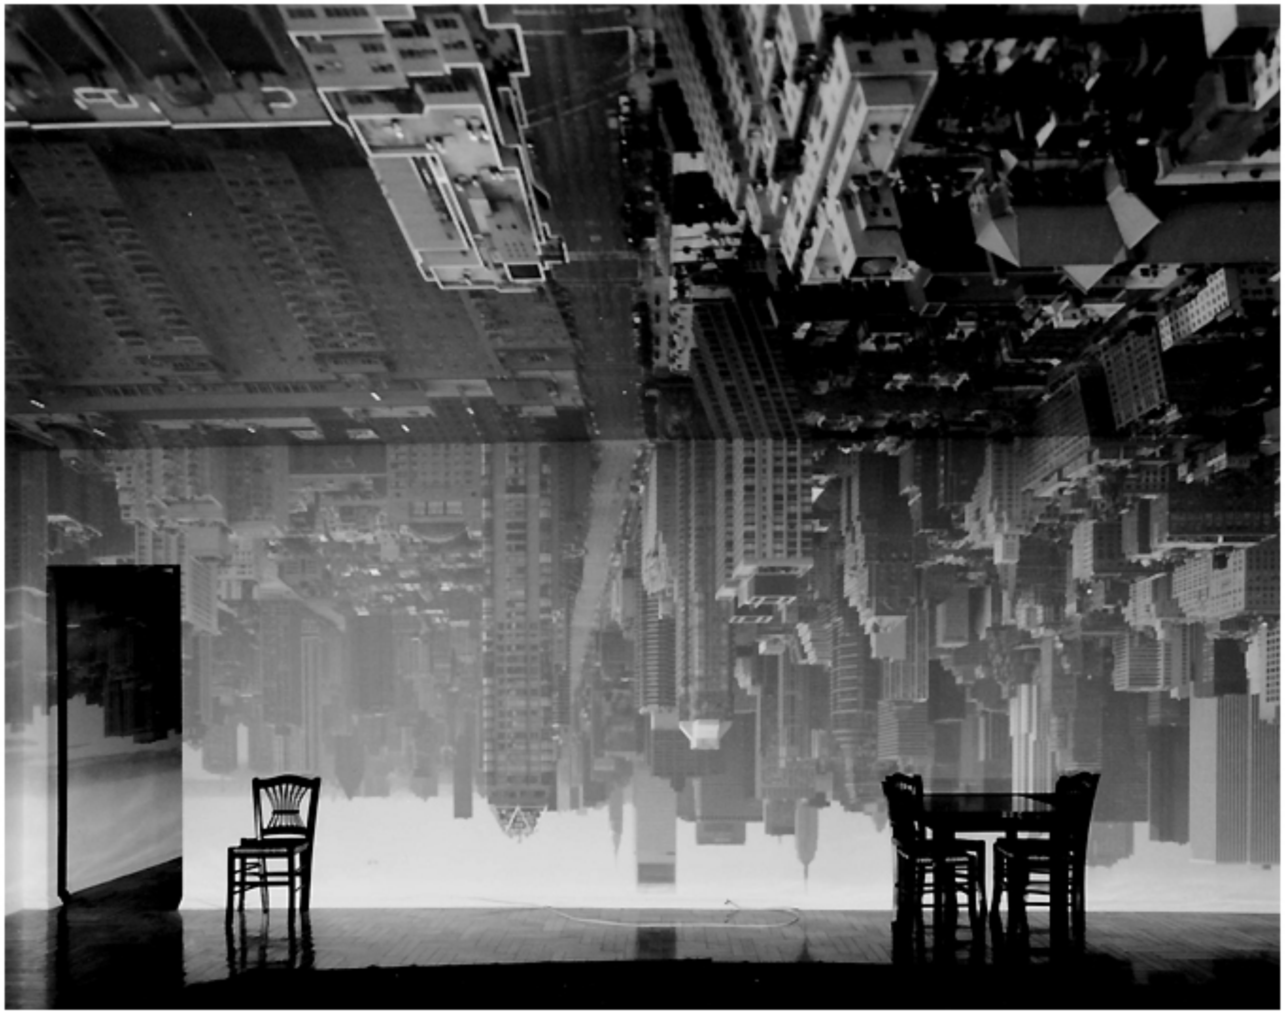

From *Grand Images Through a Tiny Opening*, **Photo District News,** February 2005

http://www.abelardomorell.net/project/camera-obscura/

Turning a room into a camera obscura

Hotel room, contrast enhanced

View from the window

Accidental pinholes produce images that are unnoticed or misinterpreted as shadows

A. Torralba and W. Freeman, Accidental Pinhole and Pinspeck Cameras, CVPR 2012

Pinhole camera

- Each point on image plane illuminated by light from one direction

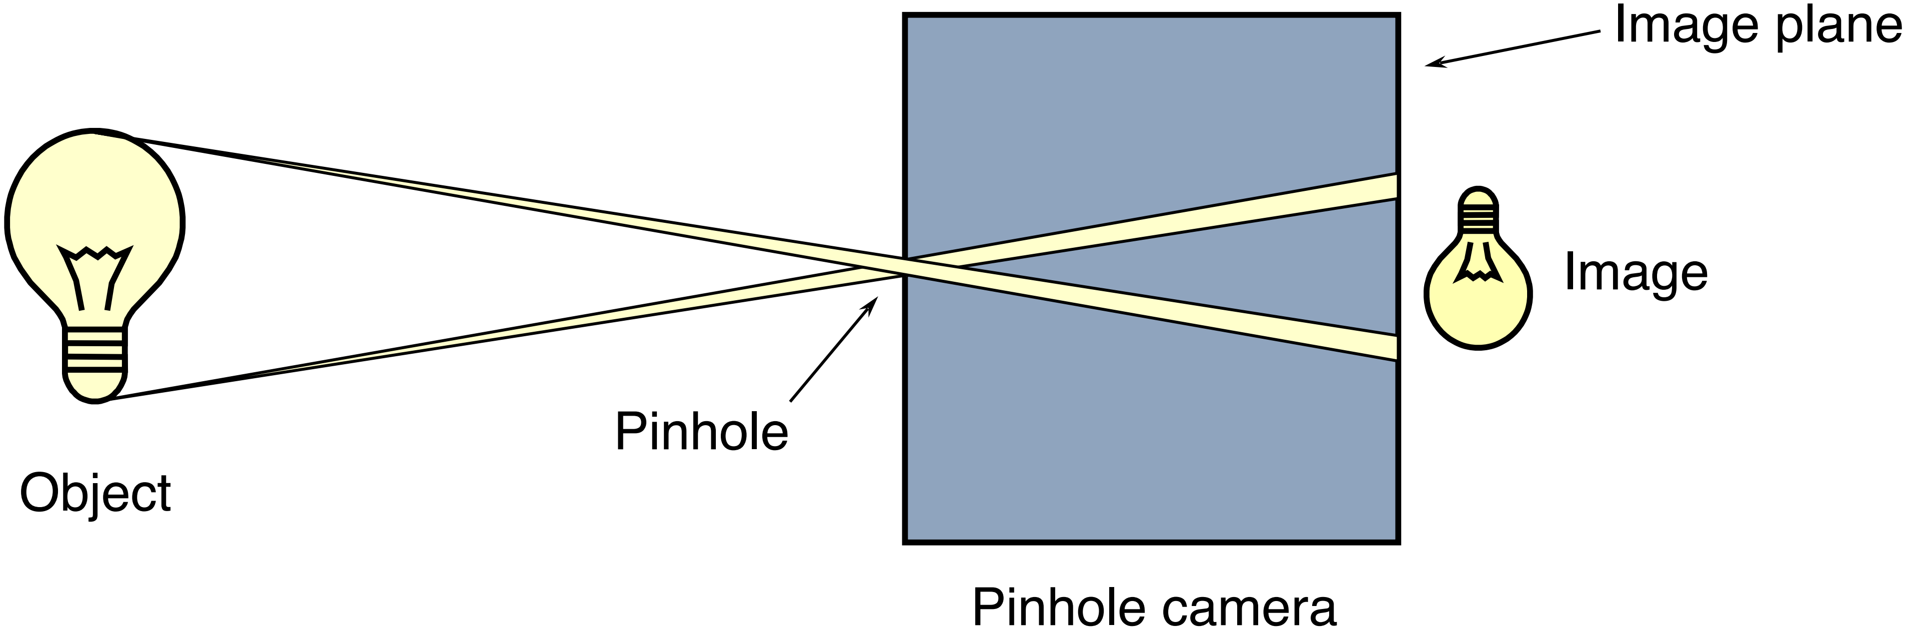

- Joseph Nicéphore Niépce: first recording onto pewter plate coated with bitumen (1826)

Dimensionality reduction: from 3D to 2D

3D world

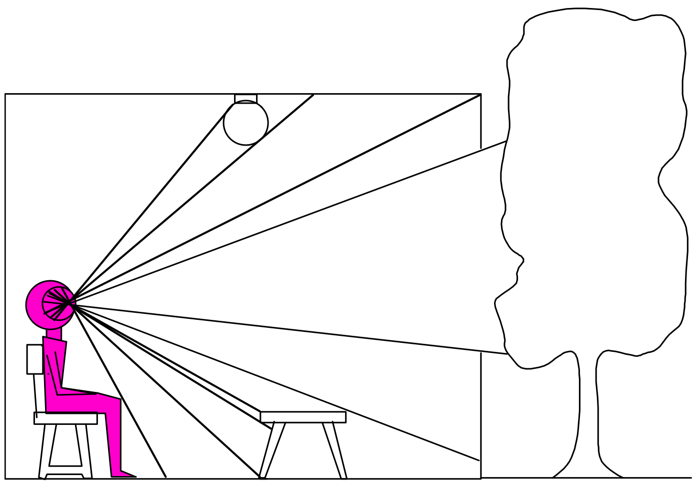

Point of observation

2D image

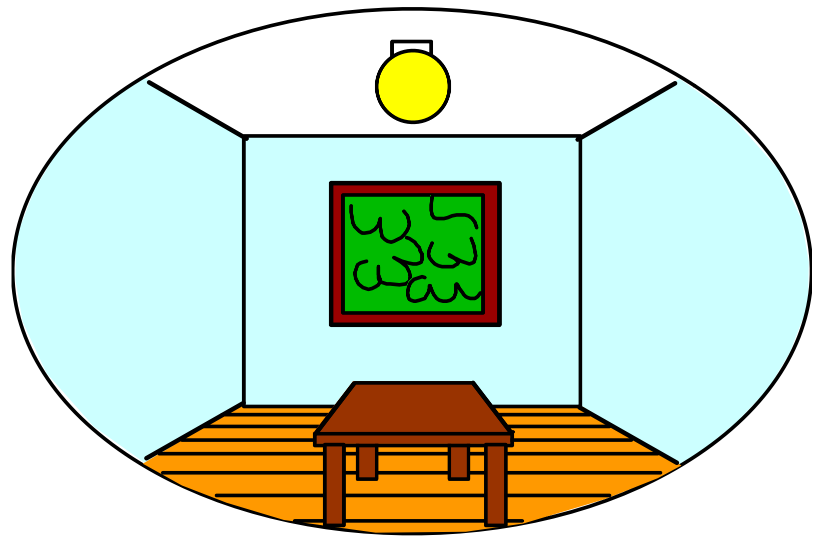

What properties of the world are preserved?

- Straight lines, incidence

What properties are not preserved?

- Angles, lengths

Slide by A. Efros

Nearer Objects Appear Bigger

size in image ~ 1/*distance*

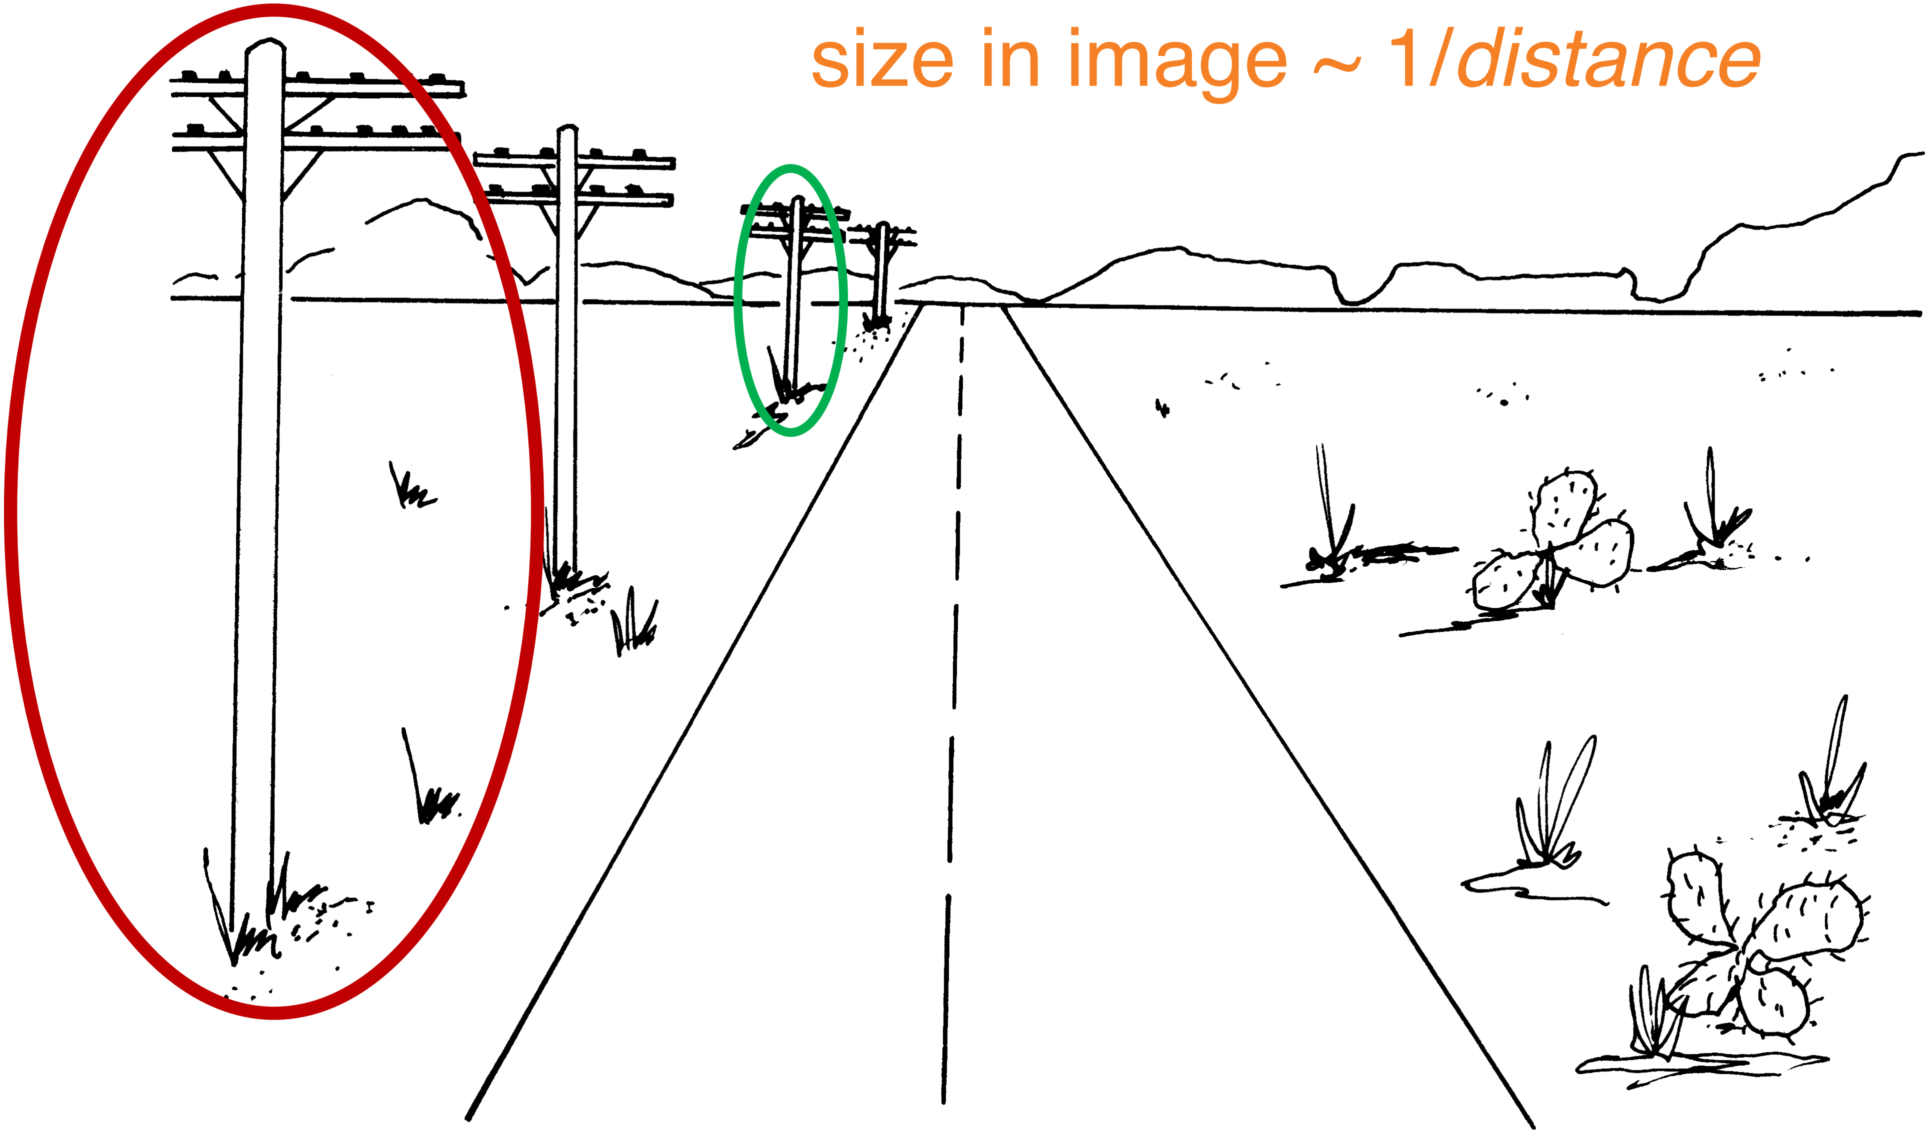

Perspective Projection Phenomena…

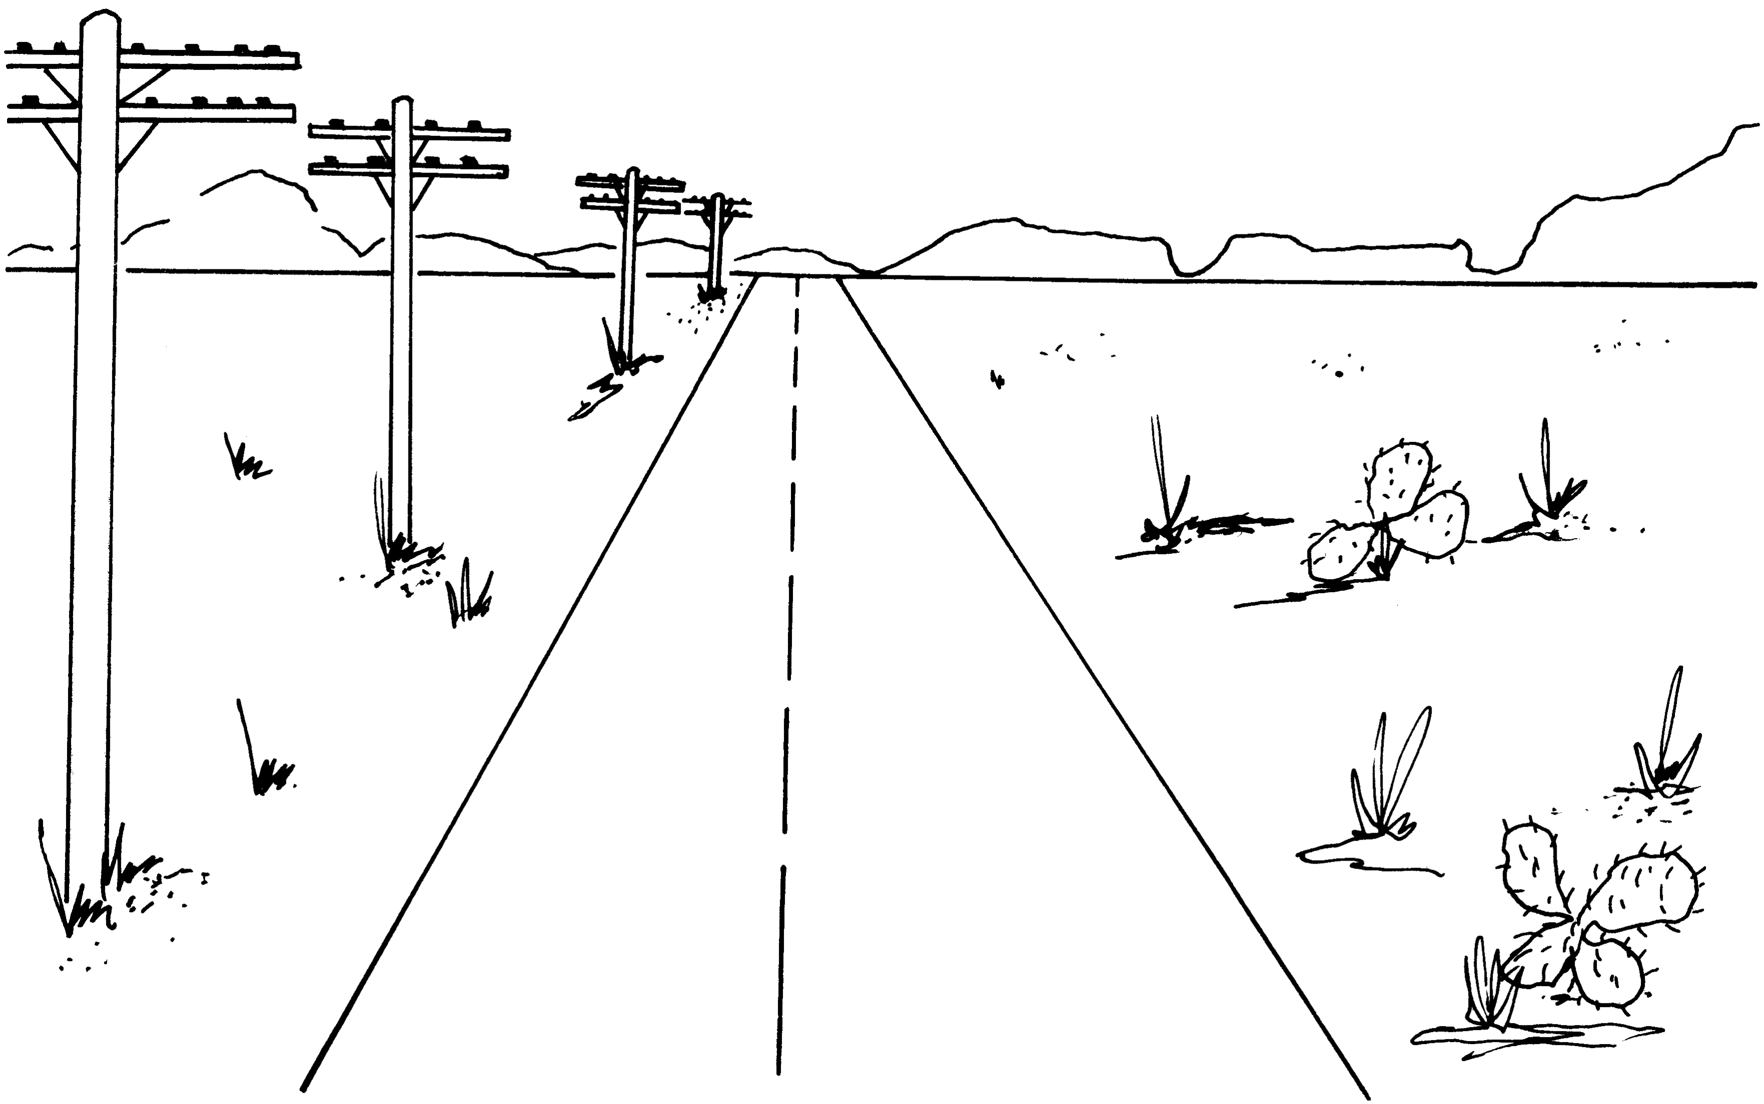

Straight Lines Remain Straight

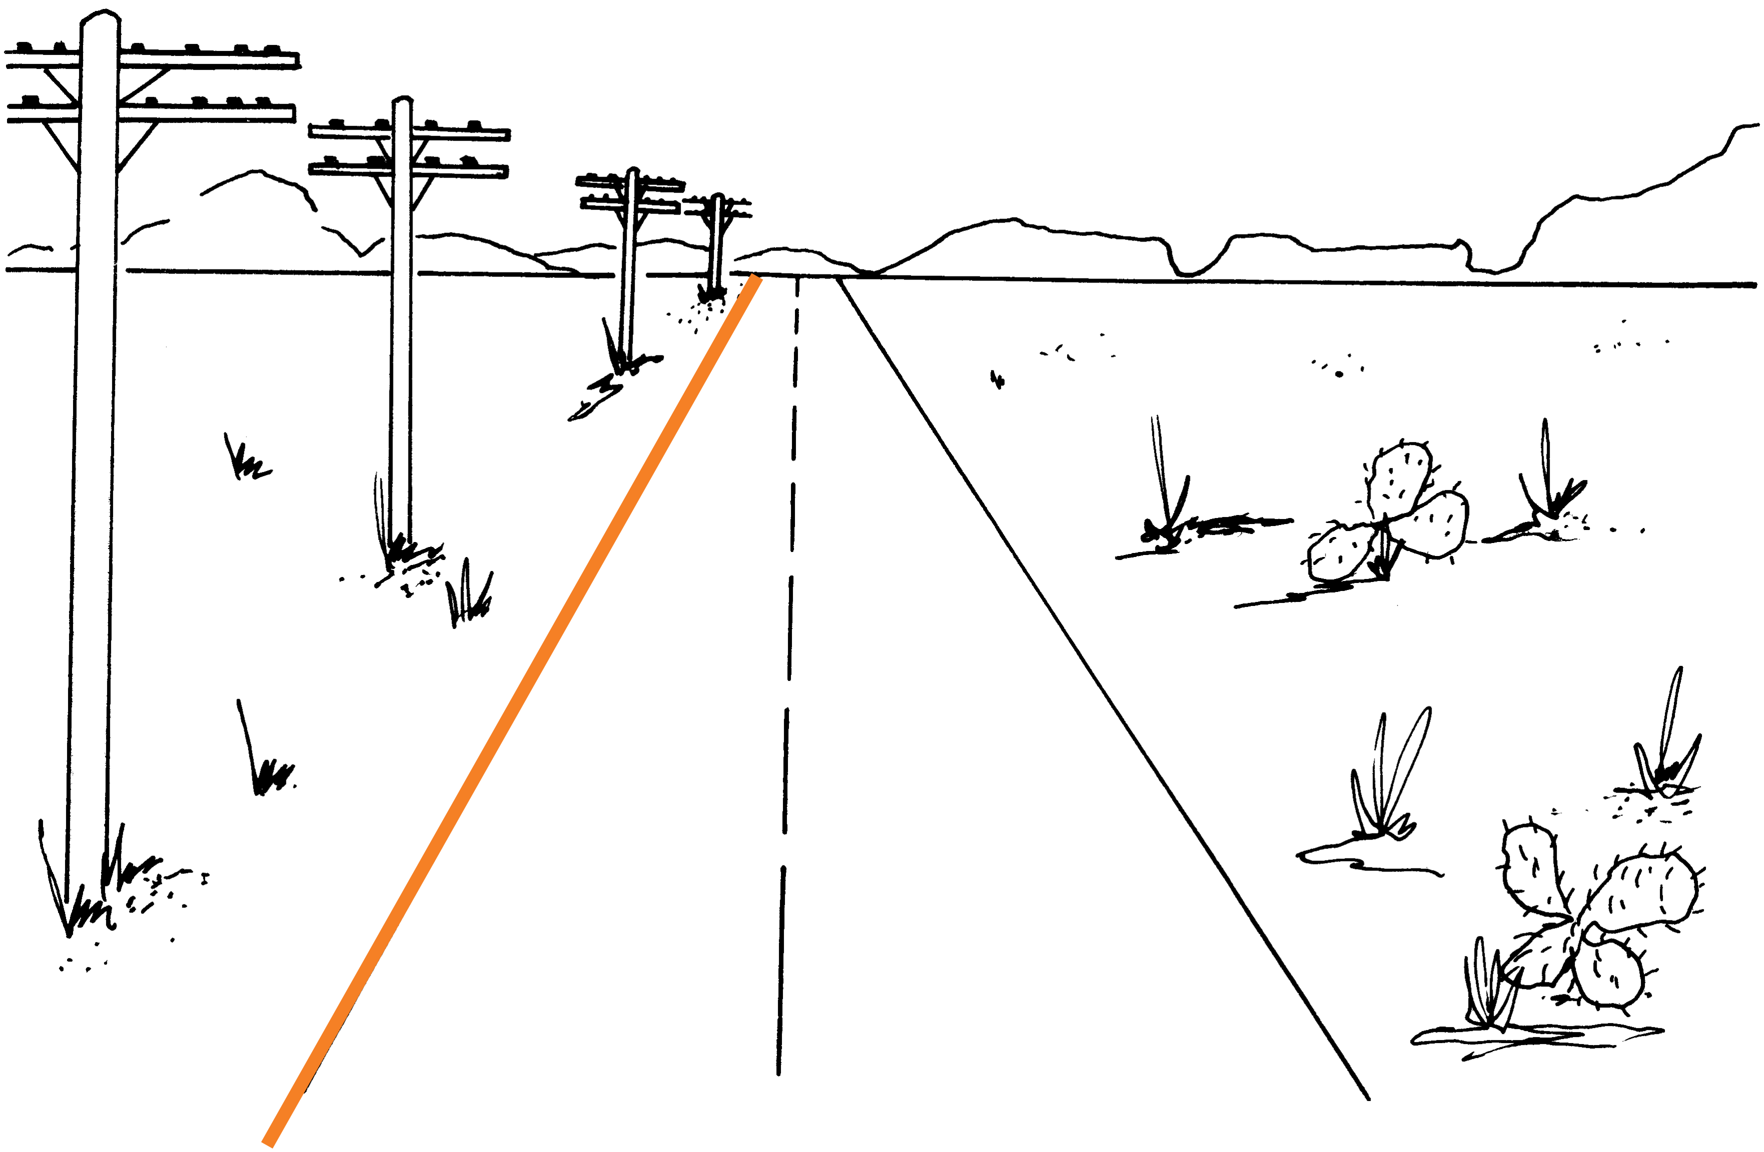

Parallel Lines Converge at Vanishing Points

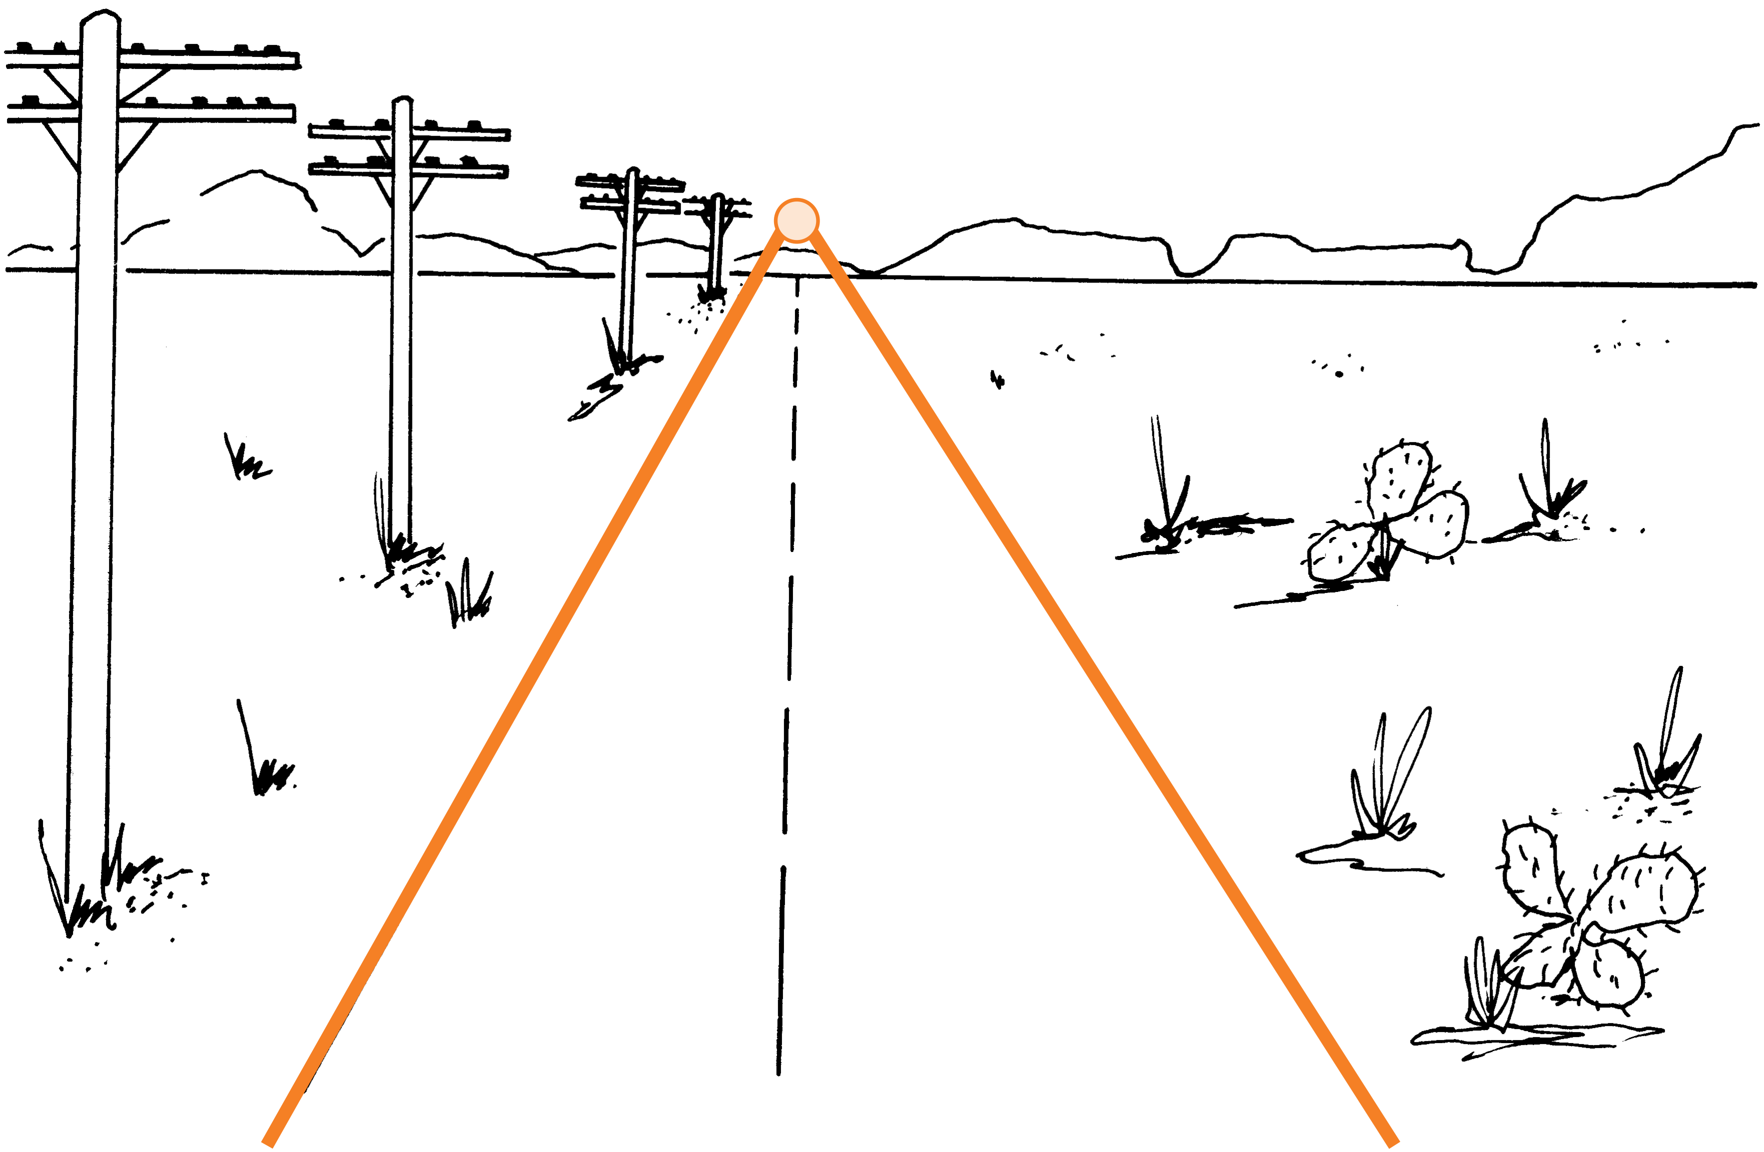

Parallel Lines Converge at Vanishing Points

Each family of parallel lines has its own vanishing point

Pinhole camera: projection of a point

Modeling projection

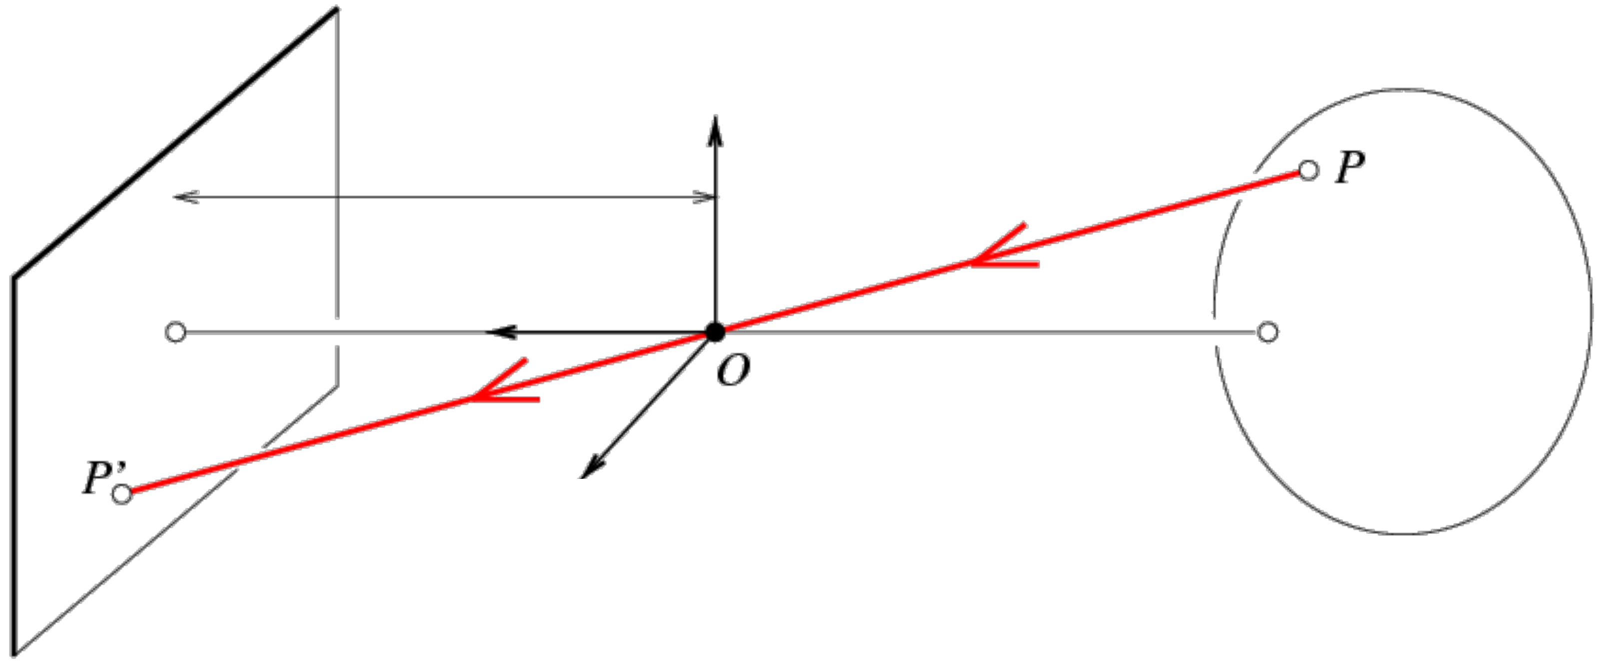

- To compute the projection *P'* of a scene point *P*, form the **visual ray** connecting *P* to the camera center *O* and find where it intersects the image plane
  - All scene points that lie on this visual ray have the same projection in the image
  - Are there scene points for which this projection is undefined?

The coordinate system

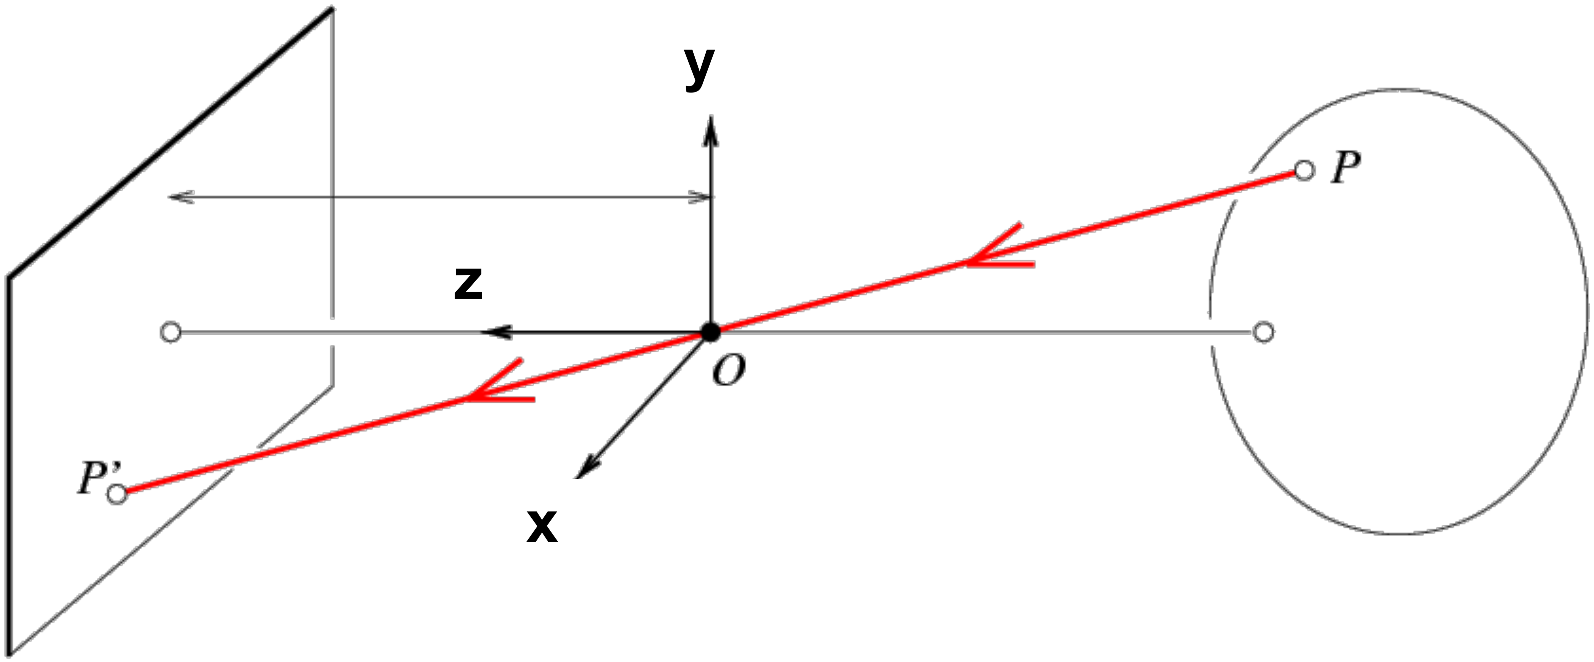

- The optical center (***O***) is at the origin
- The image plane is parallel to xy-plane or perpendicular to the z-axis, which is the *optical axis*

Projection equations

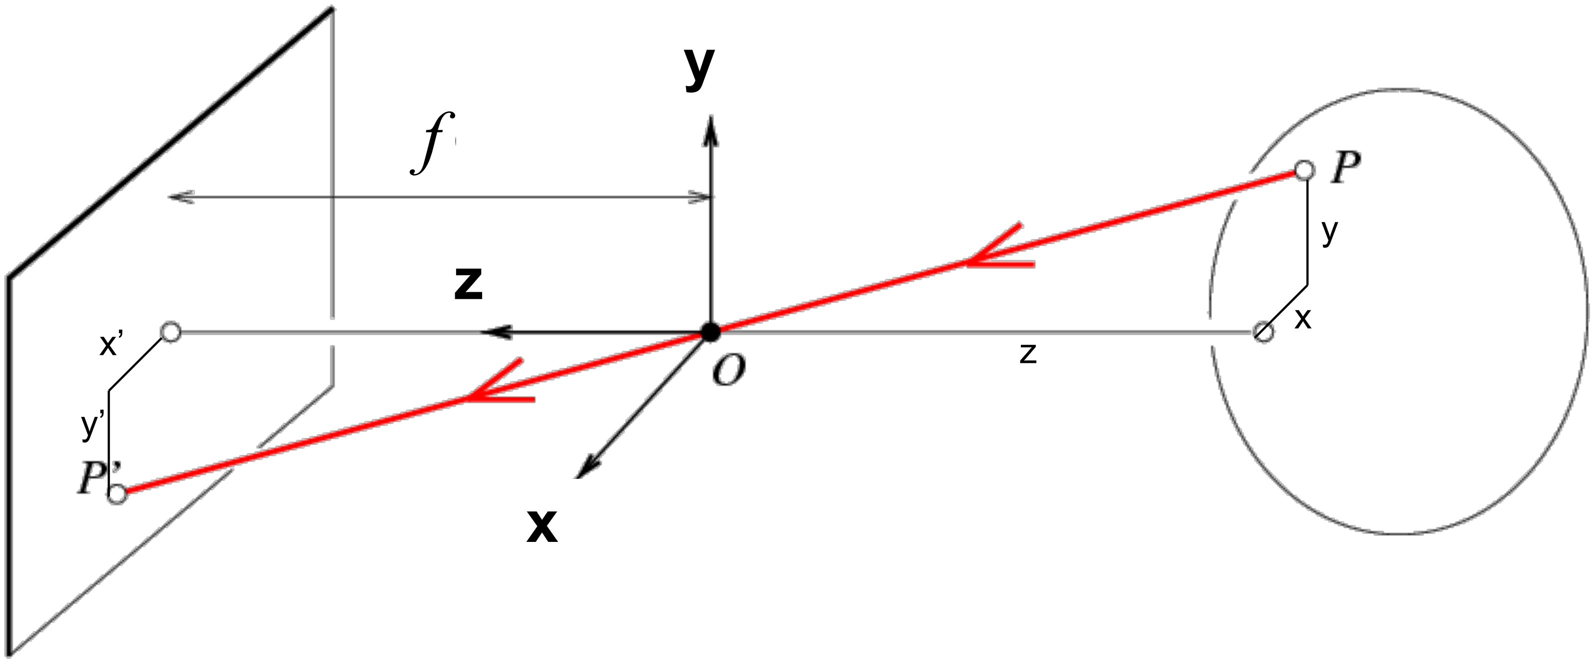

Projection equations

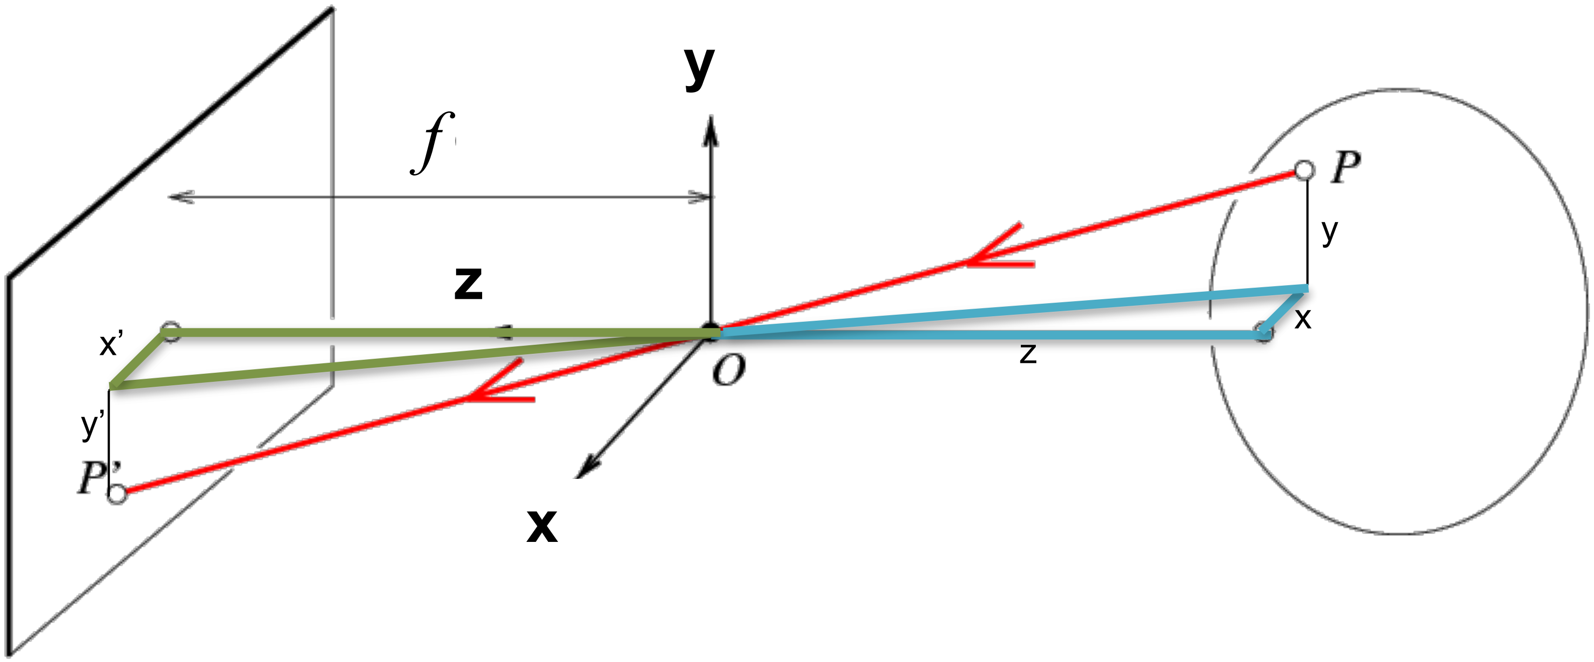

- Using similar triangles: $\frac{x}{z} = \frac{x'}{f}$

Projection equations

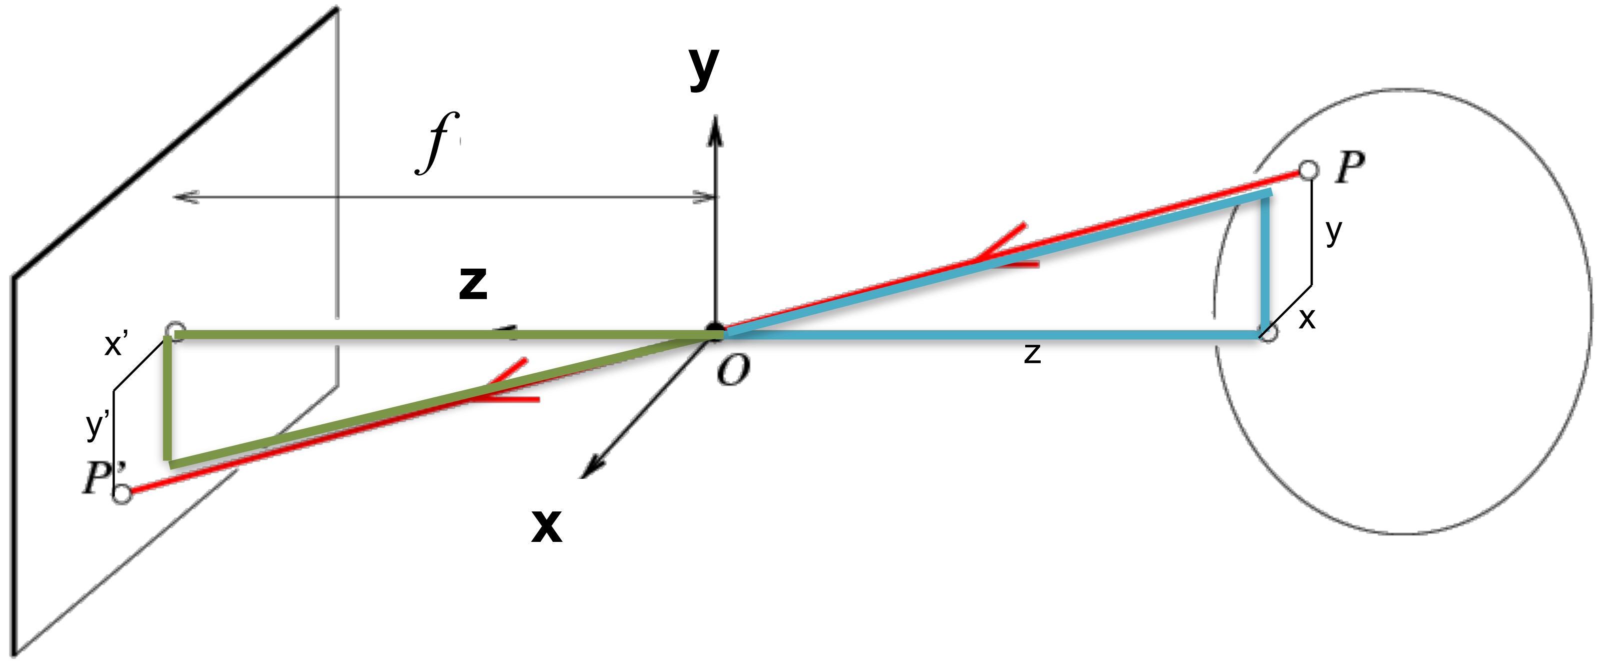

- Using similar triangles: $\frac{x}{z} = \frac{x'}{f}$ $\frac{y}{z} = \frac{y'}{f}$

Projection equations

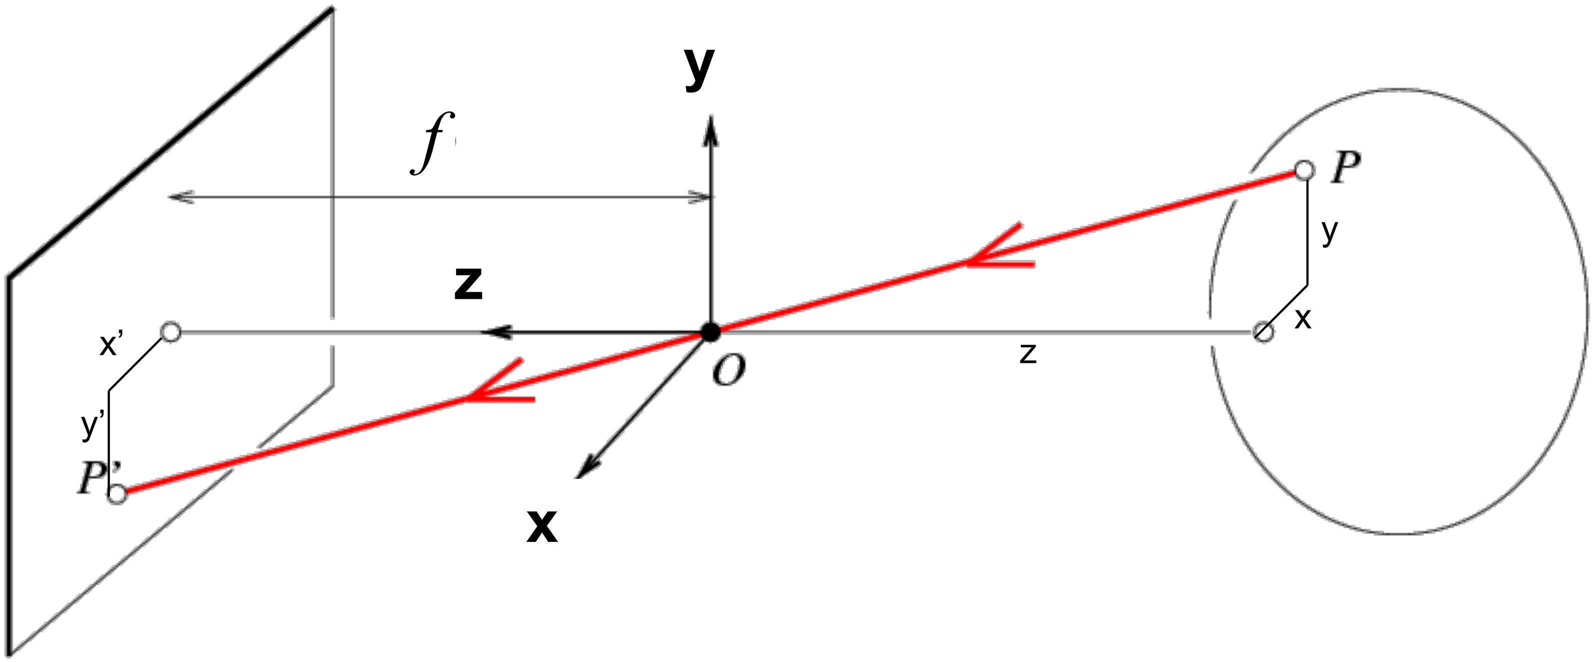

- Using similar triangles: $\frac{x}{z} = \frac{x'}{f} \qquad \frac{y}{z} = \frac{y'}{f}$

- Thus: $(x, y, z) \rightarrow (f\frac{x}{z}, f\frac{y}{z})$

Fronto-parallel planes

- What happens to the projection of a pattern on a plane parallel to the image plane?
  - All points on that plane are at a fixed *depth z*
  - The pattern gets scaled by a factor of *f / z*, but angles and ratios of lengths/areas are preserved

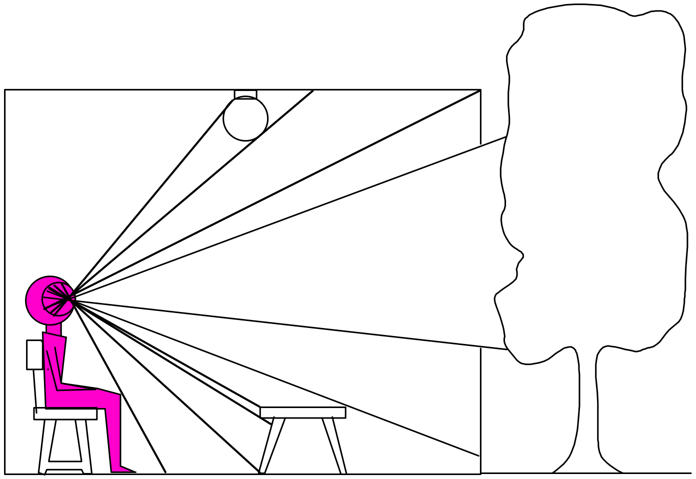

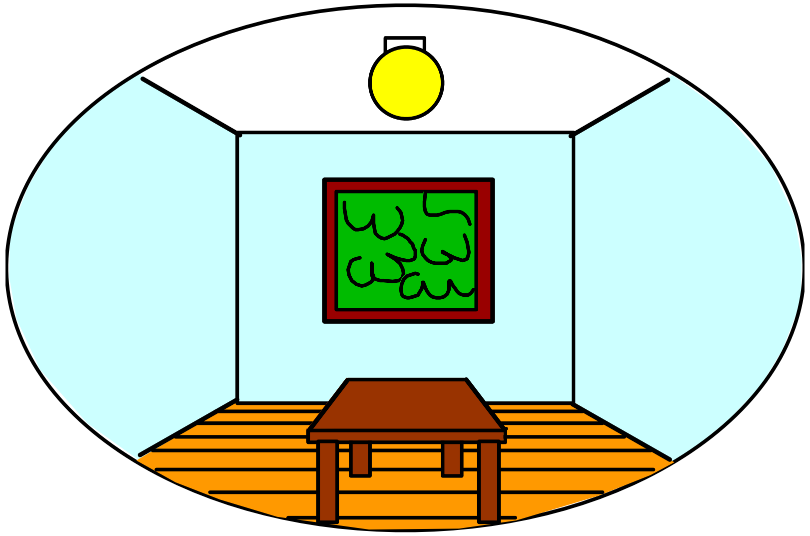

$$(x, y, z) \rightarrow (f\frac{x}{z}, f\frac{y}{z})$$

Pinhole camera: projection of a line

Vanishing points

- All parallel lines converge to a *vanishing point*
  - Each direction in space is associated with its own vanishing point
  - Exception: *directions parallel to the image plane*

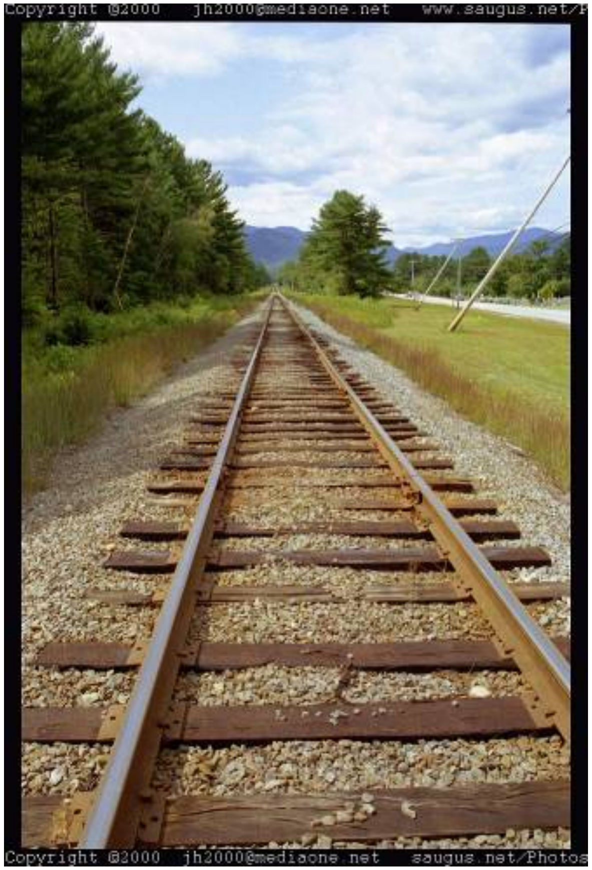

Constructing the vanishing point of a line

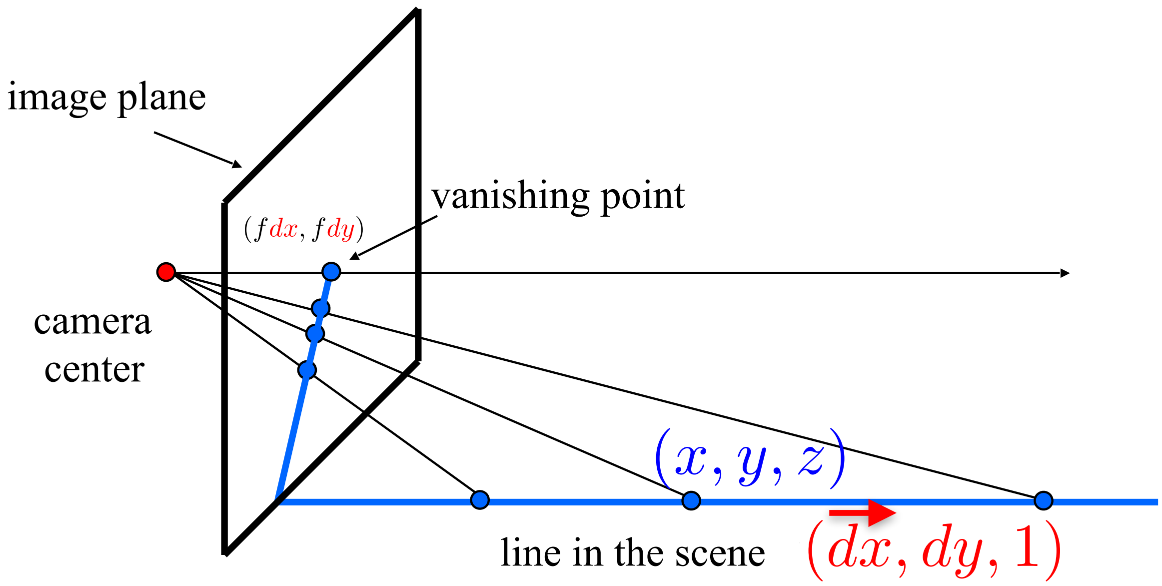

- **Claim**: Each direction in space is associated with one vanishing point
  - Any point on the line: $(x, y, z) + \alpha(dx, dy, 1)$
  - This point is projected to: $\left(f\frac{x+\alpha dx}{z+\alpha}, f\frac{y+\alpha dy}{z+\alpha}\right)$
  - The limit as $\alpha \rightarrow \inf:\ (f dx, f dy)$
  - Thus the vanishing point is independent of the location (x, y, z) and uniquely determined by the direction $(dx, dy, 1)$

Constructing the vanishing point of a line

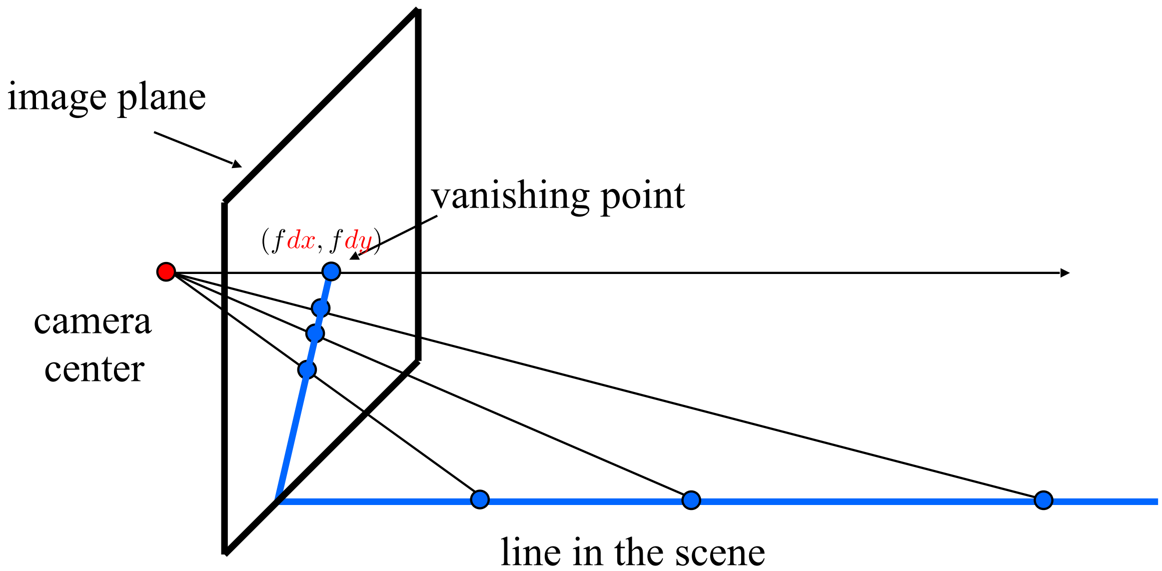

- To find the vanishing point, shoot a ray from camera center along the same direction. Find the intersection with the image plane.
- How does the vanishing point move if the camera is moved without rotation?

Source: Steve Seitz

Perspective cues

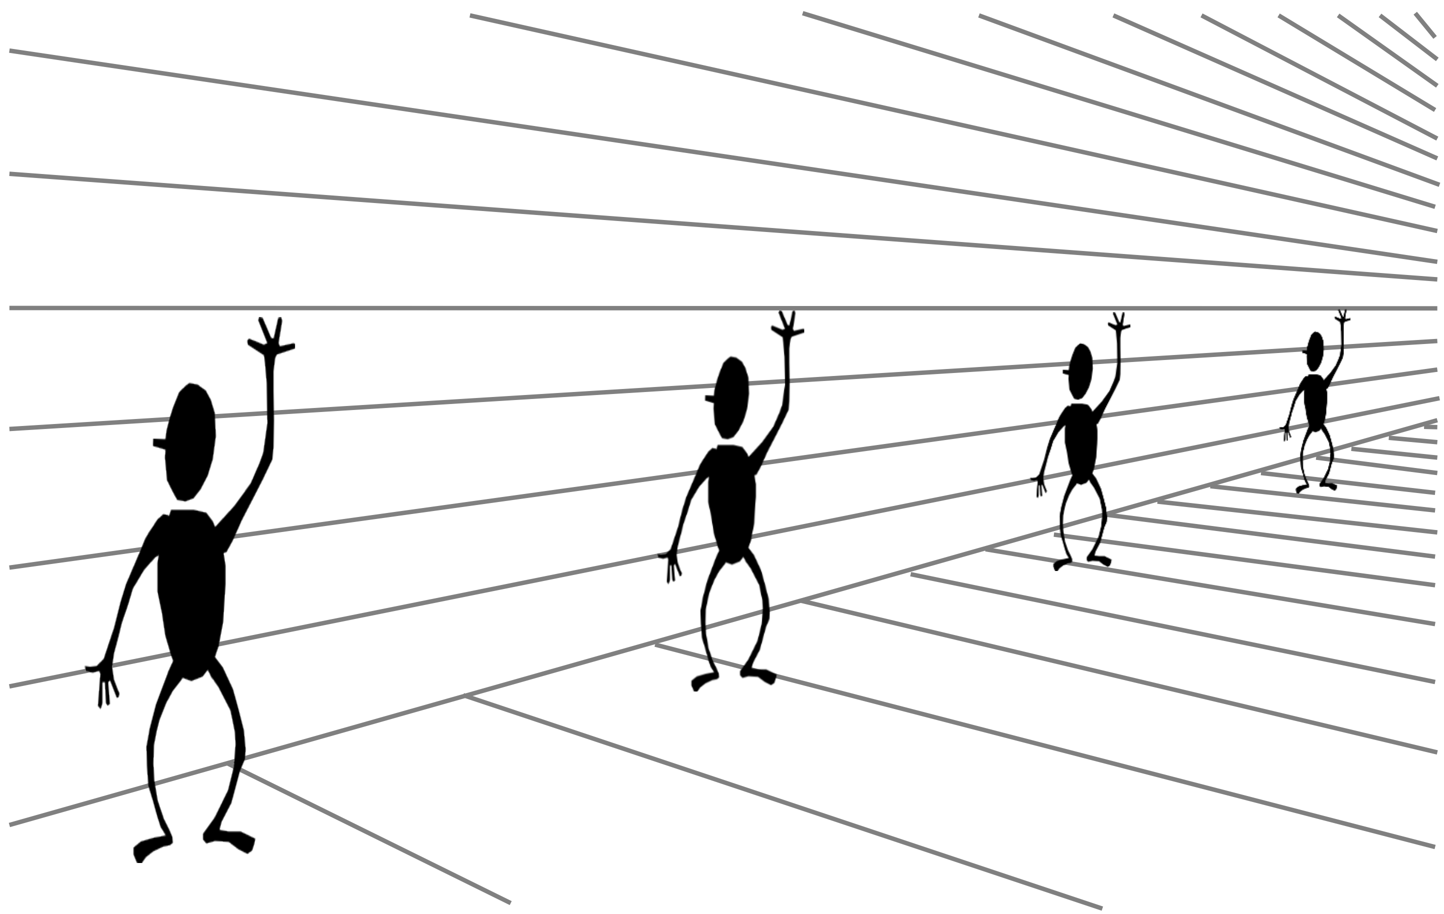

Slide by Steve Seitz

Perspective cues

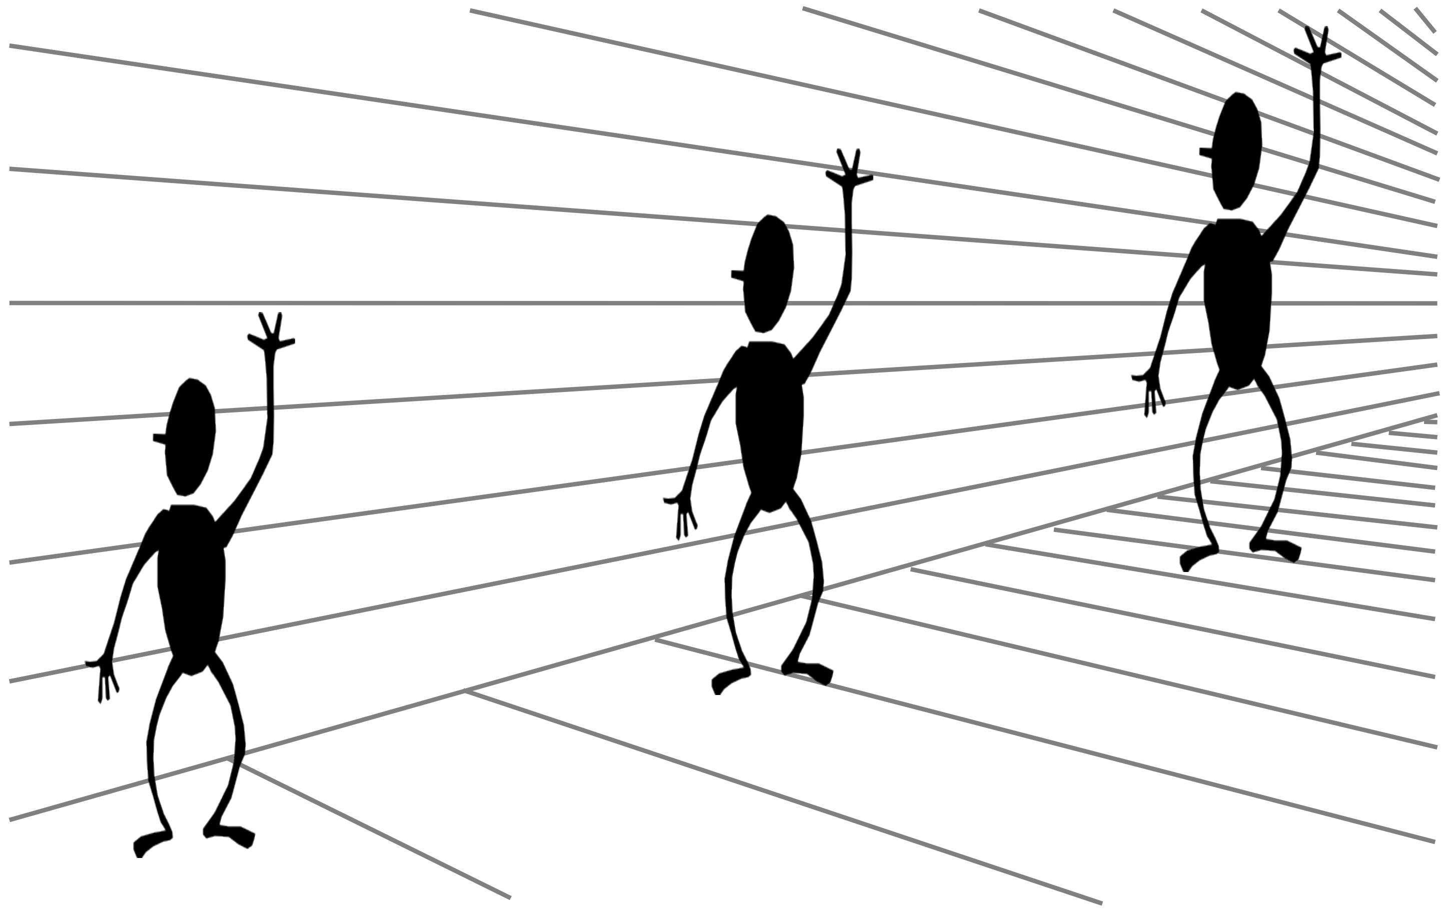

Slide by Steve Seitz

Perspective cues

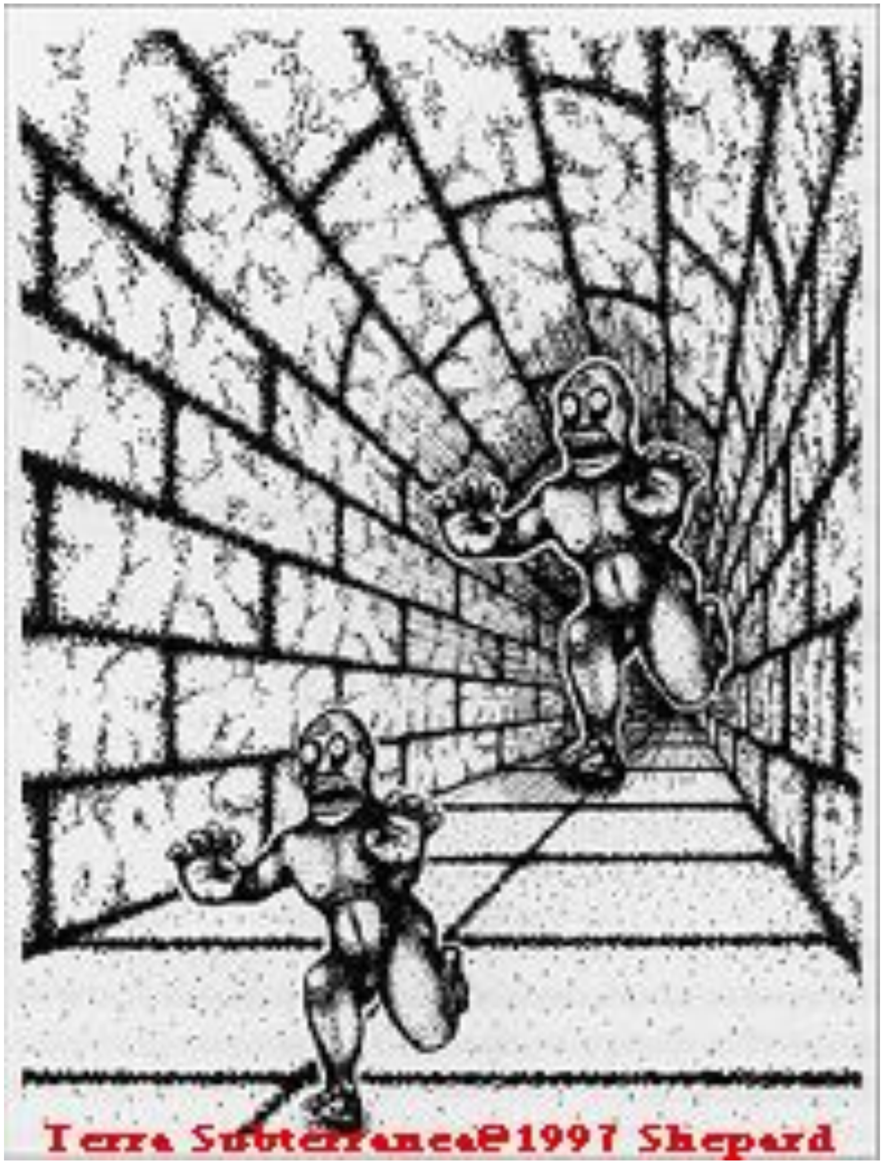

Pinhole camera: projection beyond a single line

Vanishing lines of planes

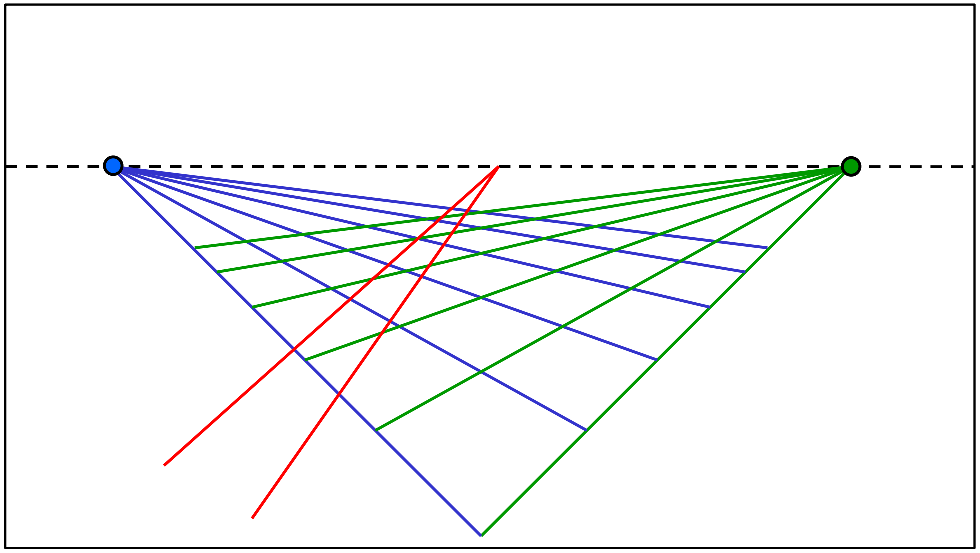

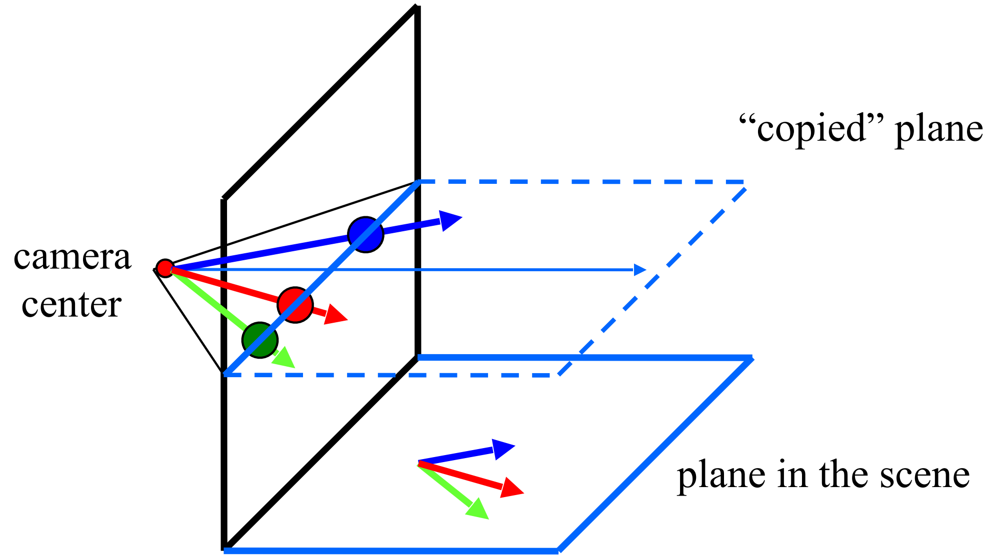

- Vanishing points *of co-planar directions* form a *vanishing line* (exercise for fun: show this algebraically)
- *Horizon*: vanishing line of the ground plane
  - All points at the same height as the camera project to the horizon
  - Points higher (resp. lower) than the camera project above (resp. below) the horizon
  - Provides way of comparing height of objects

Adapted from Steve Seitz

Comparing heights

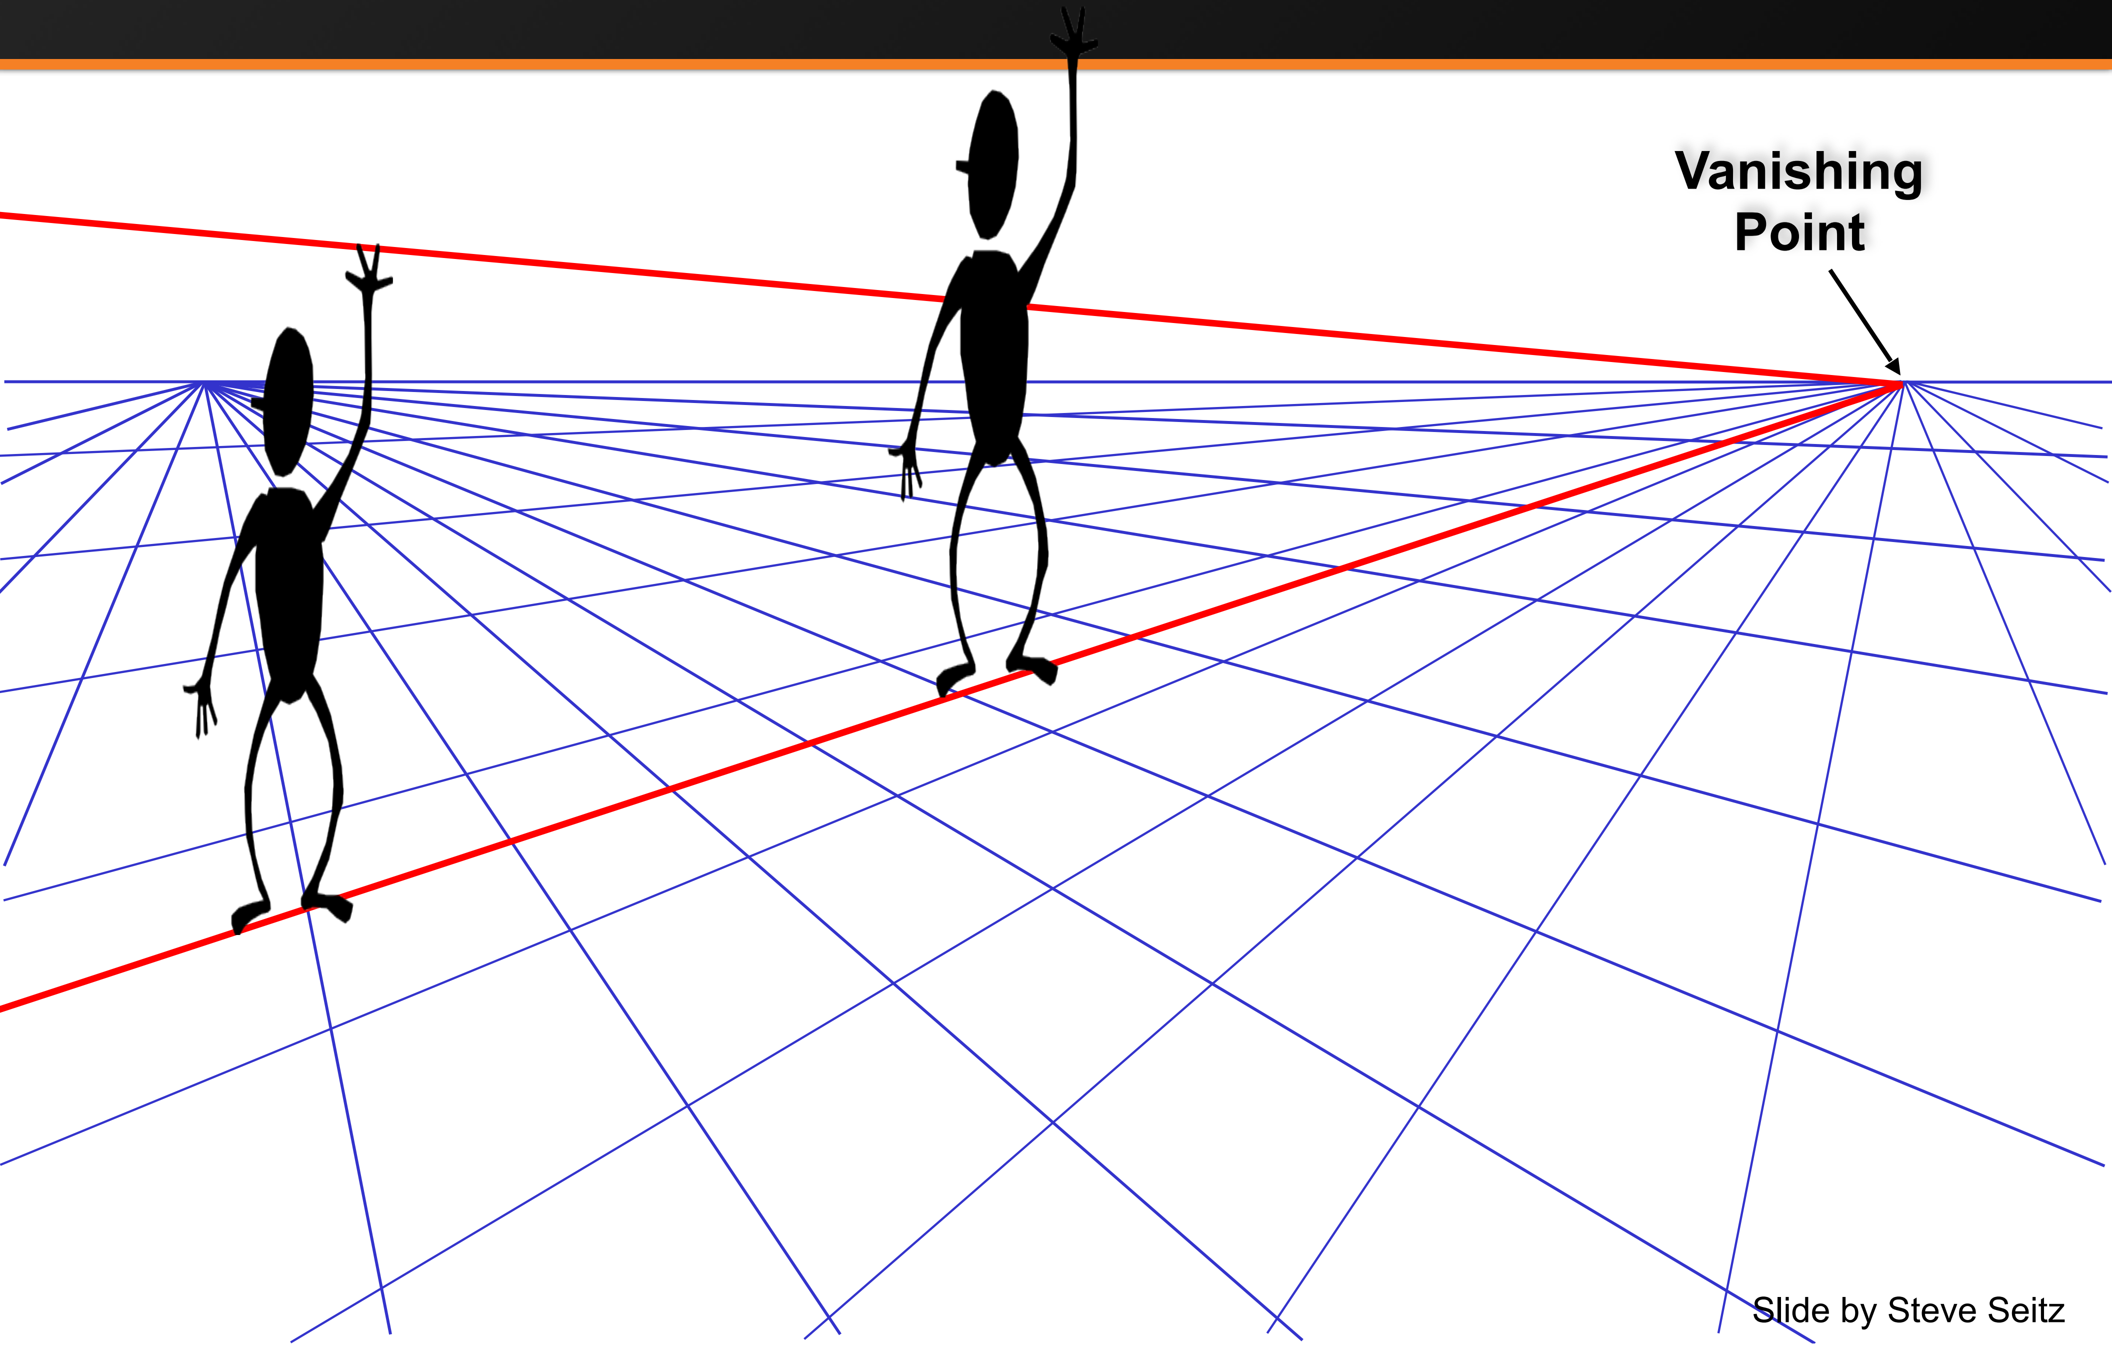

Slide by Steve Seitz

Measuring height

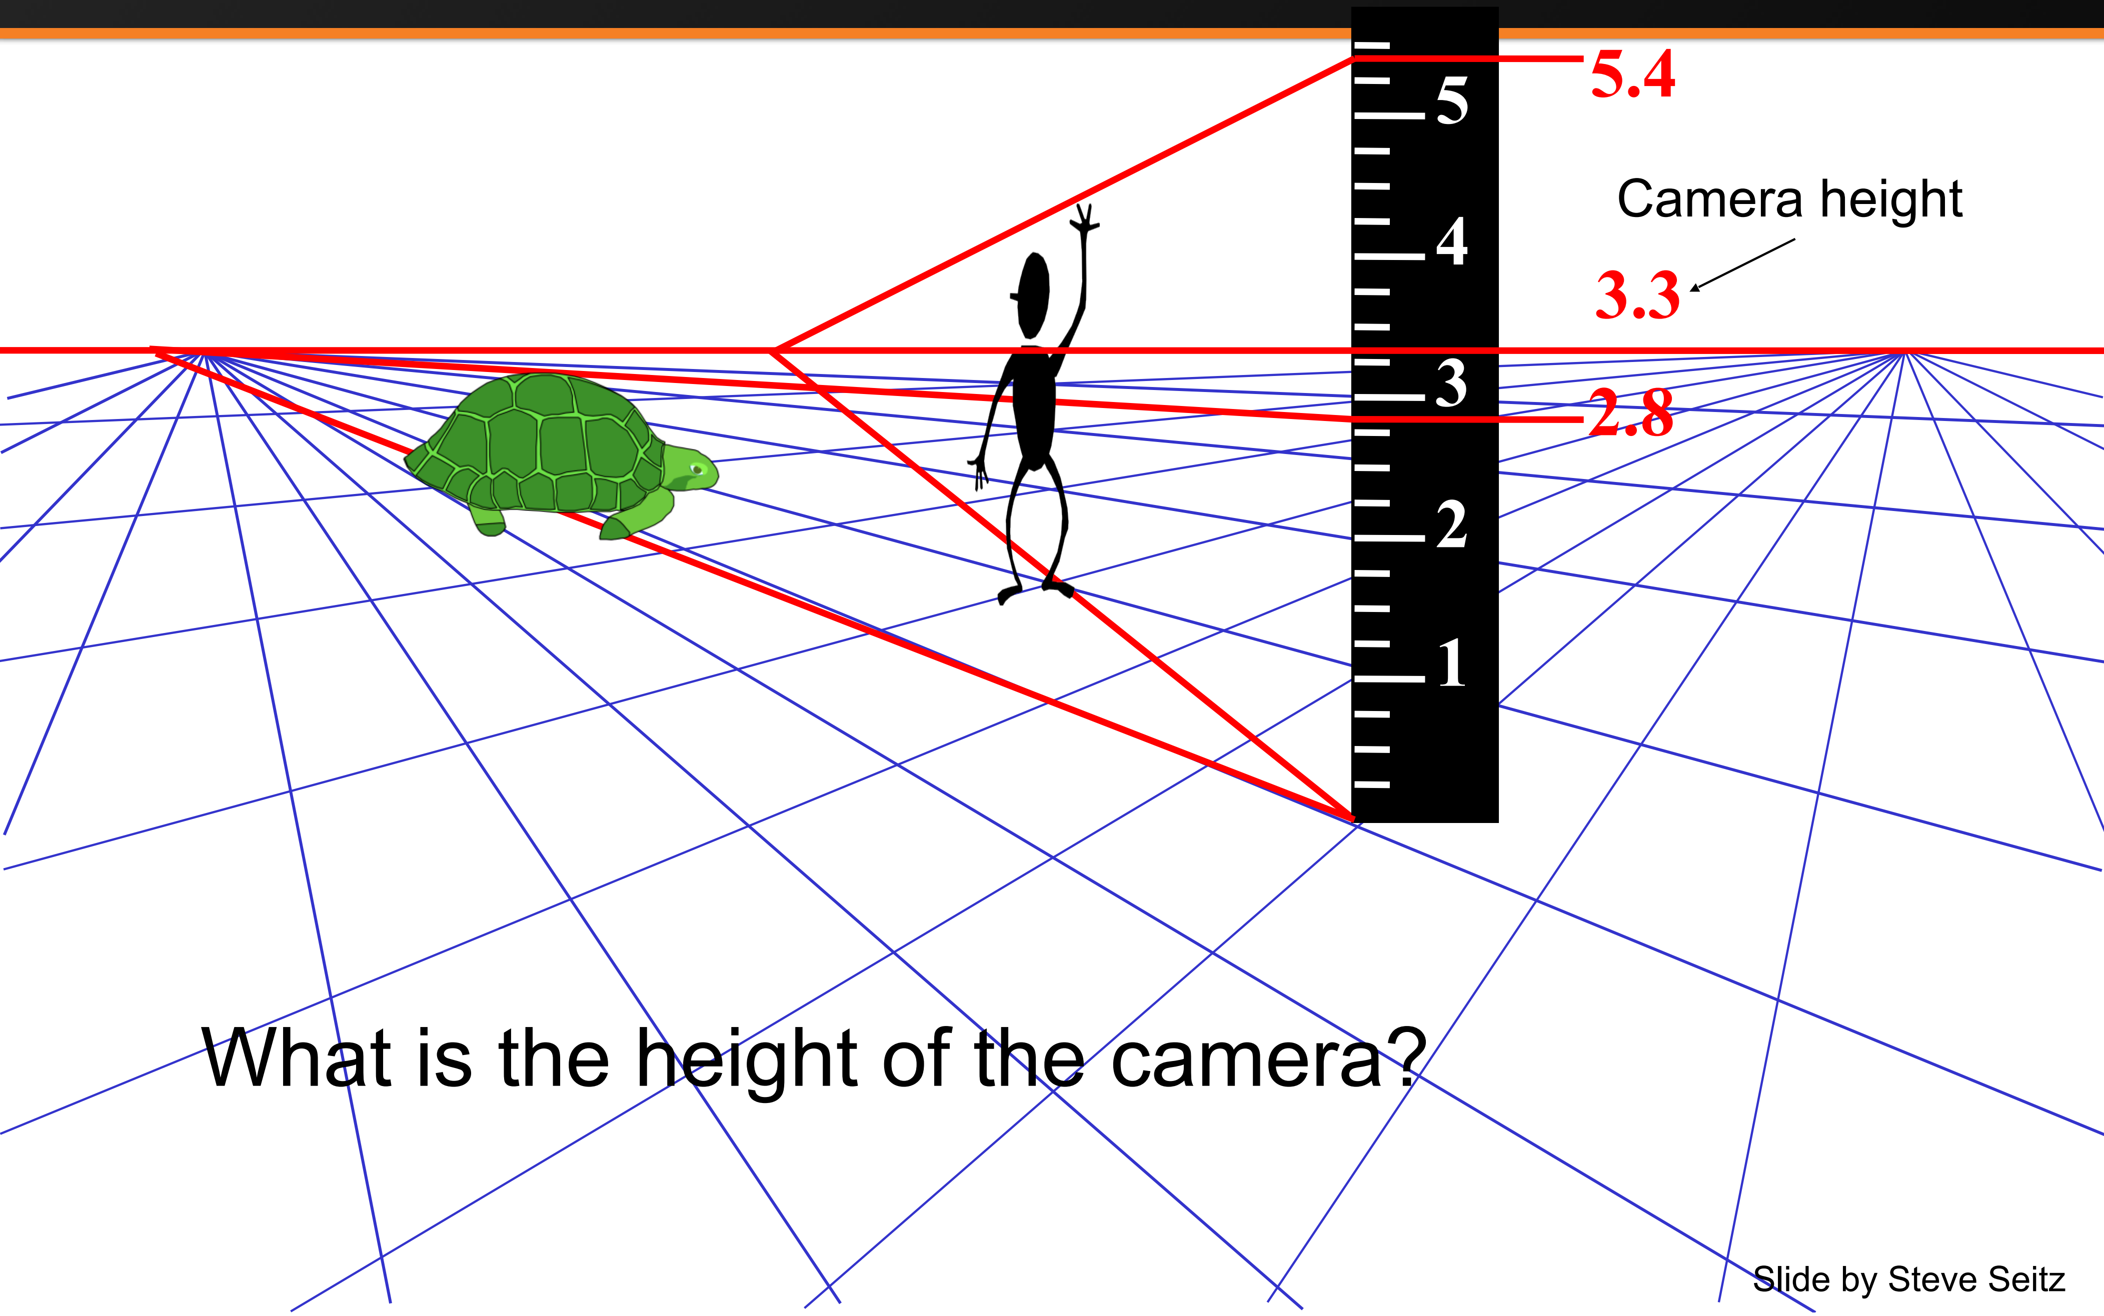

What is the height of the camera?

Slide by Steve Seitz

Quiz 1

Which is higher? The parachutist or the person taking the picture?

Image source: S. Seitz

Quiz 2

How does the location of the horizon change if the person taking the picture raise the camera but keep its orientation?

Image source: S. Seitz

Perspective distortion

- Are the widths of the projected columns equal?
  - The exterior columns are wider
  - This is not an optical illusion, and is not due to lens flaws
  - Phenomenon pointed out by Da Vinci

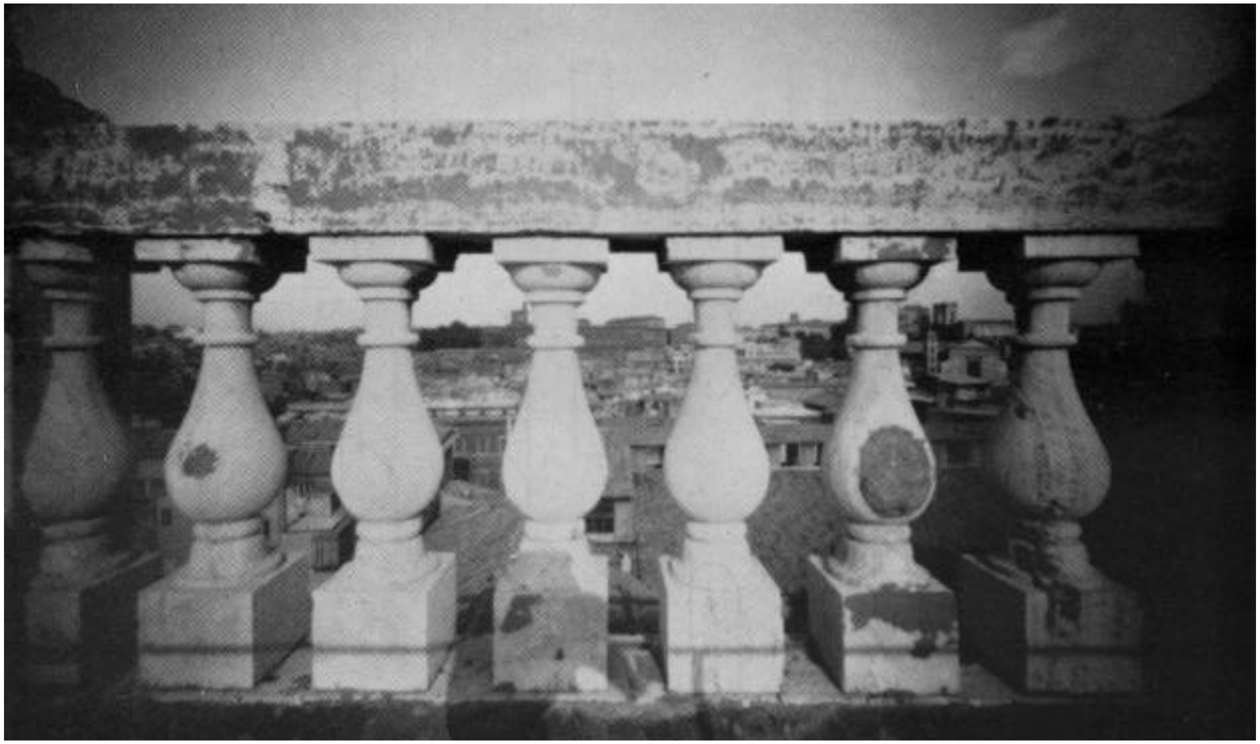

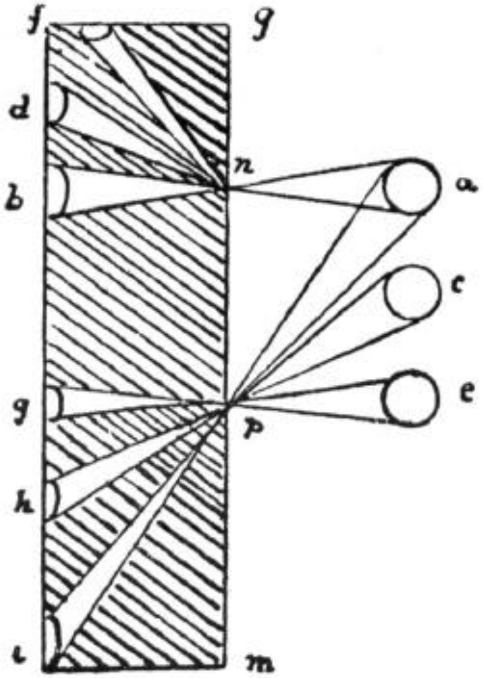

Source: F. Durand

Perspective distortion

- What is the shape of the projection of a sphere?

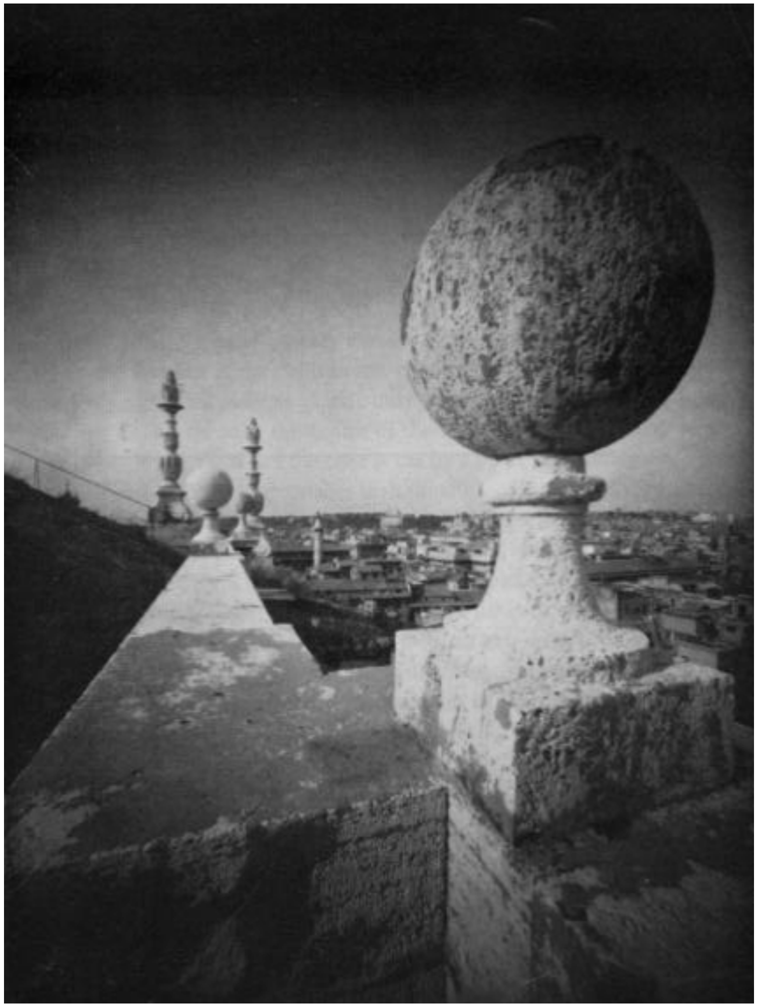

Image source: F. Durand

Perspective distortion

- What is the shape of the projection of a sphere?

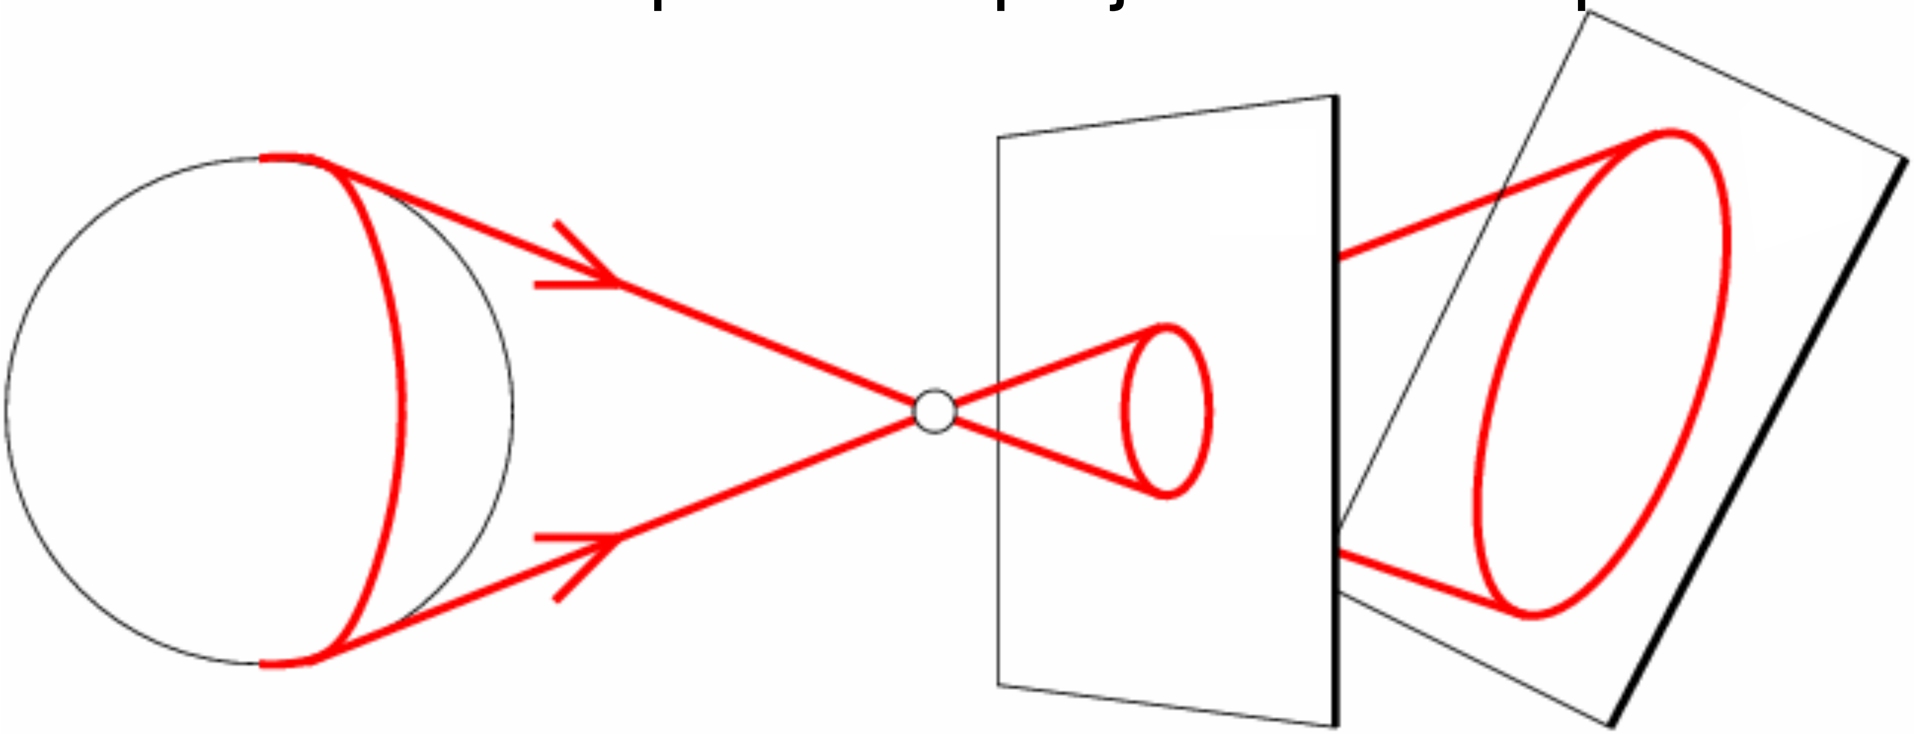

Perspective distortion: People

Building a real camera

Pinhole camera

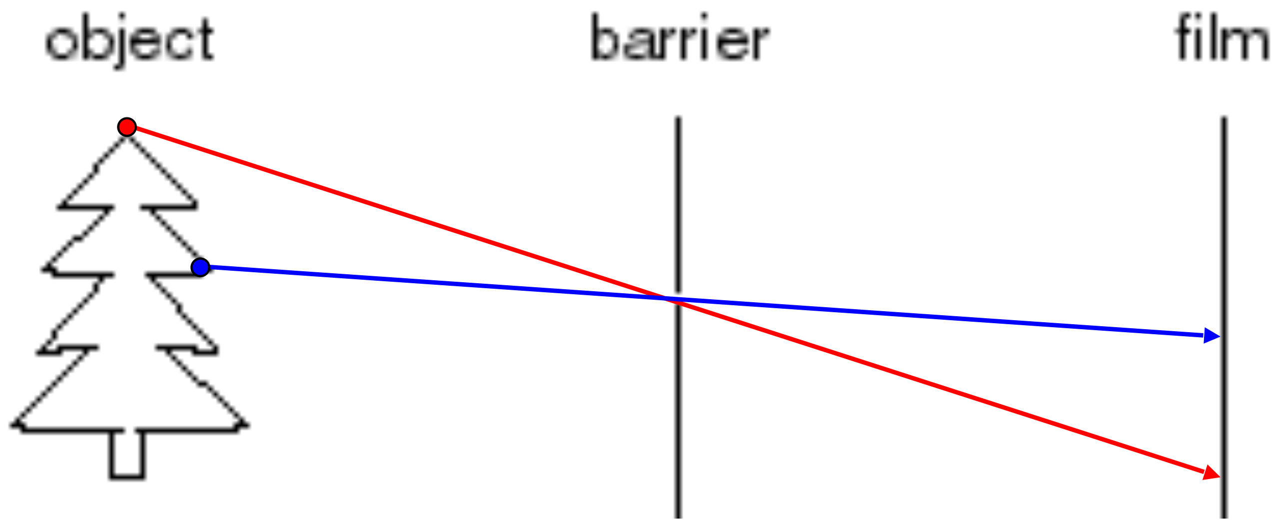

Slide by Steve Seitz

Home-made pinhole camera

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

Slide by A. Efros

Aperture

- Controls amount of light

Shrinking the aperture

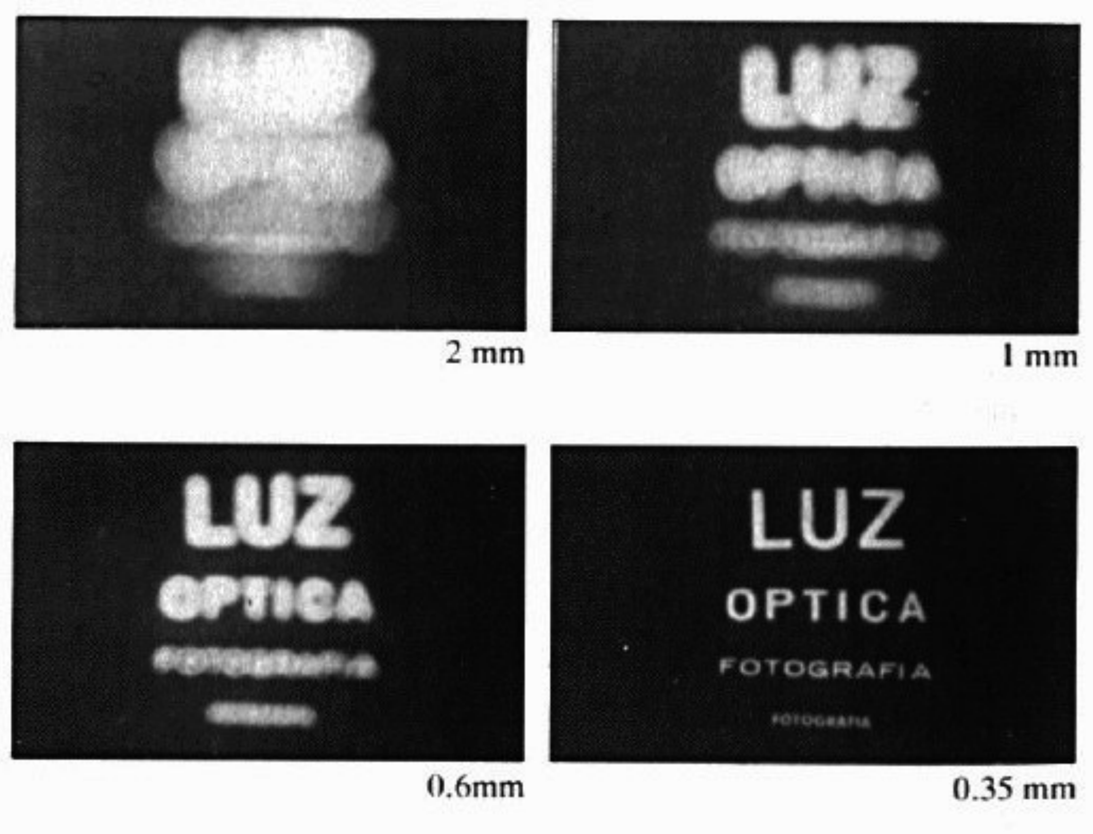

Why not make the aperture as small as possible?

- Less light gets through
- Diffraction effects…

Slide by Steve Seitz

Shrinking the aperture

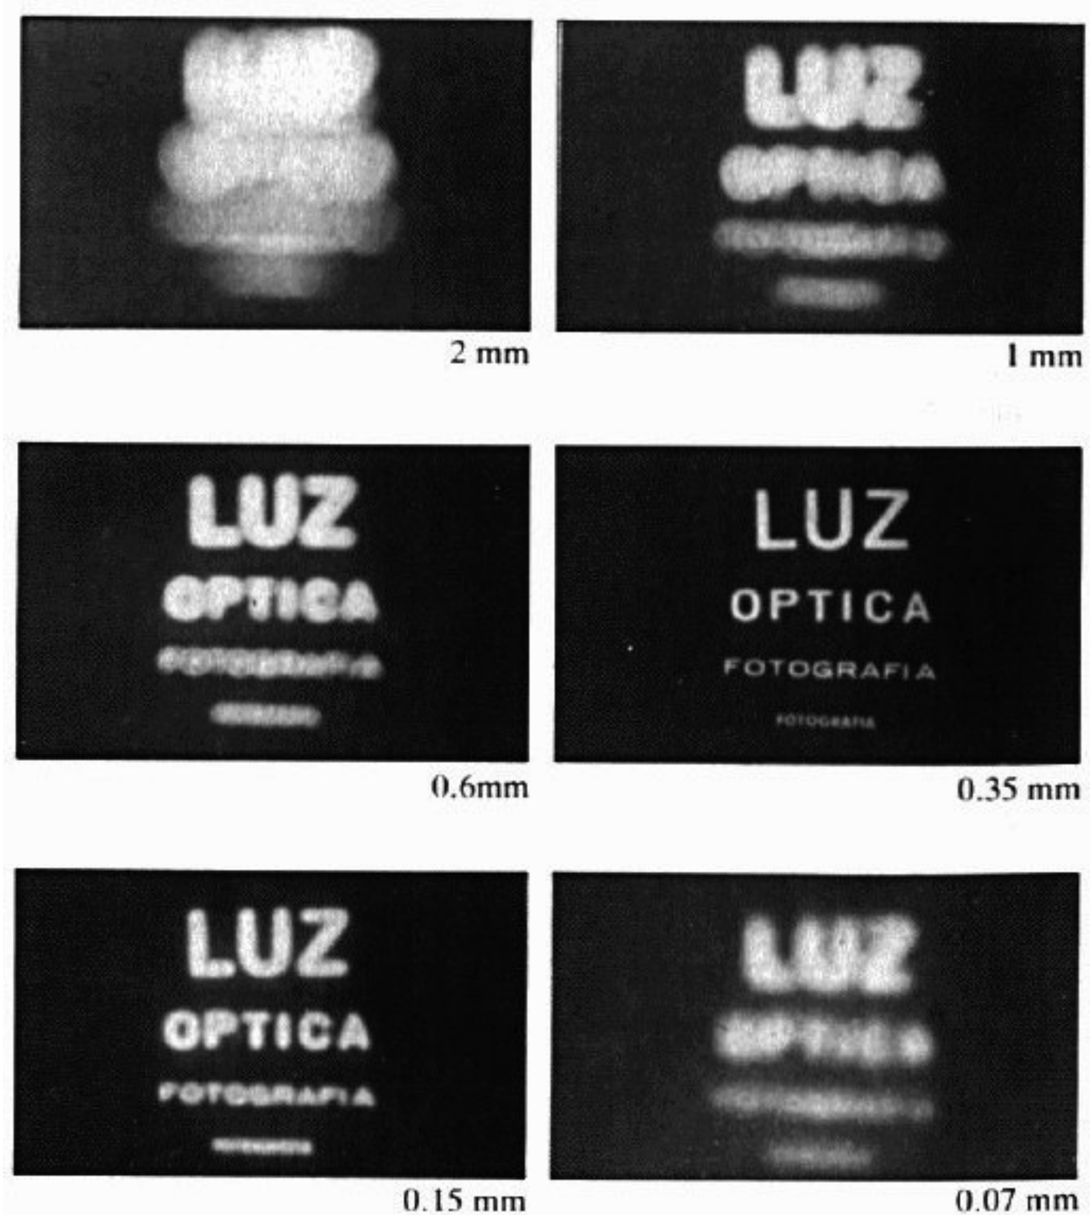

Slide by Steve Seitz

Adding a lens

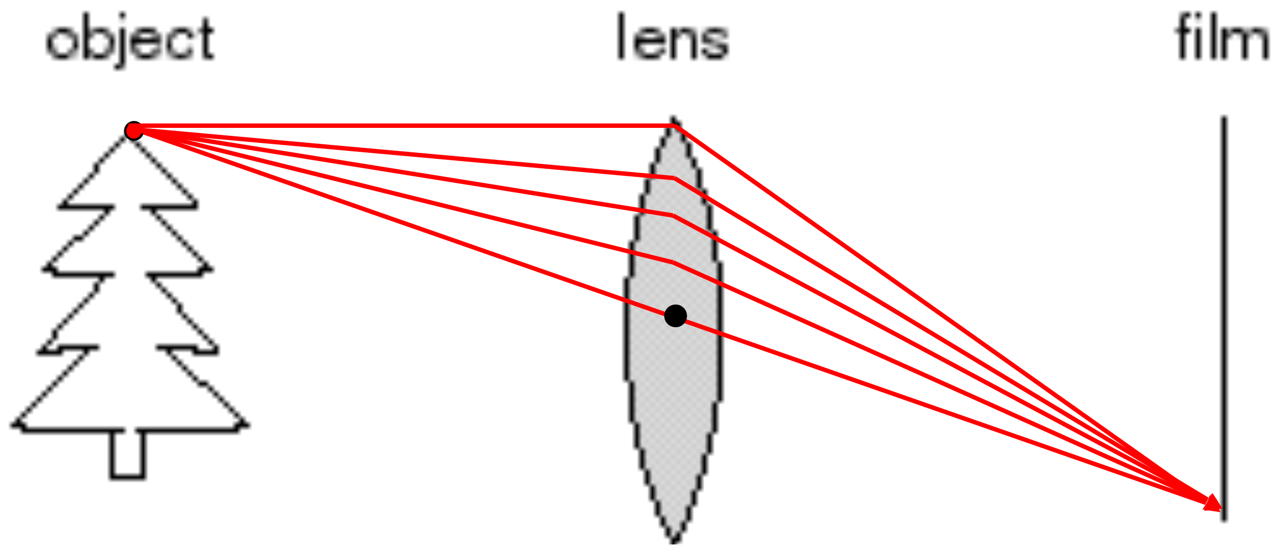

A lens focuses light onto the film

- Thin lens model:
  – Rays passing through the center are not deviated (pinhole projection model still holds)

Slide by Steve Seitz

Adding a lens

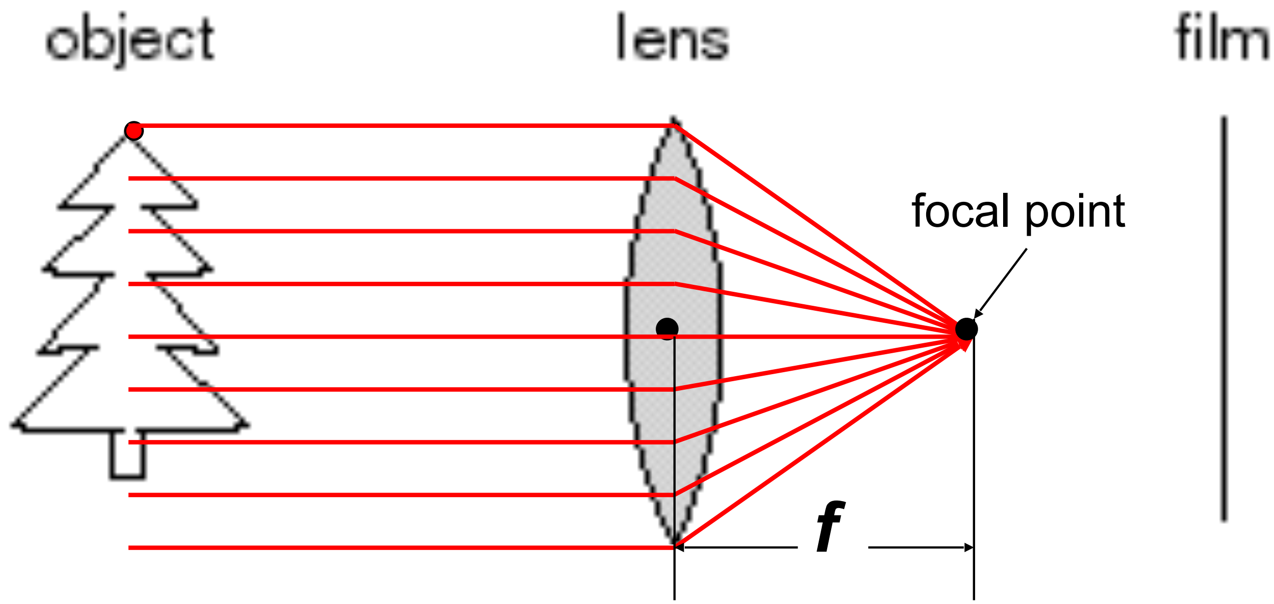

A lens focuses light onto the film

- Thin lens model:
  - Rays passing through the center are not deviated (pinhole projection model still holds)
  - All parallel rays (along the principal axis) converge to one point on a plane located at the *focal length f*

Slide by Steve Seitz

Adding a lens

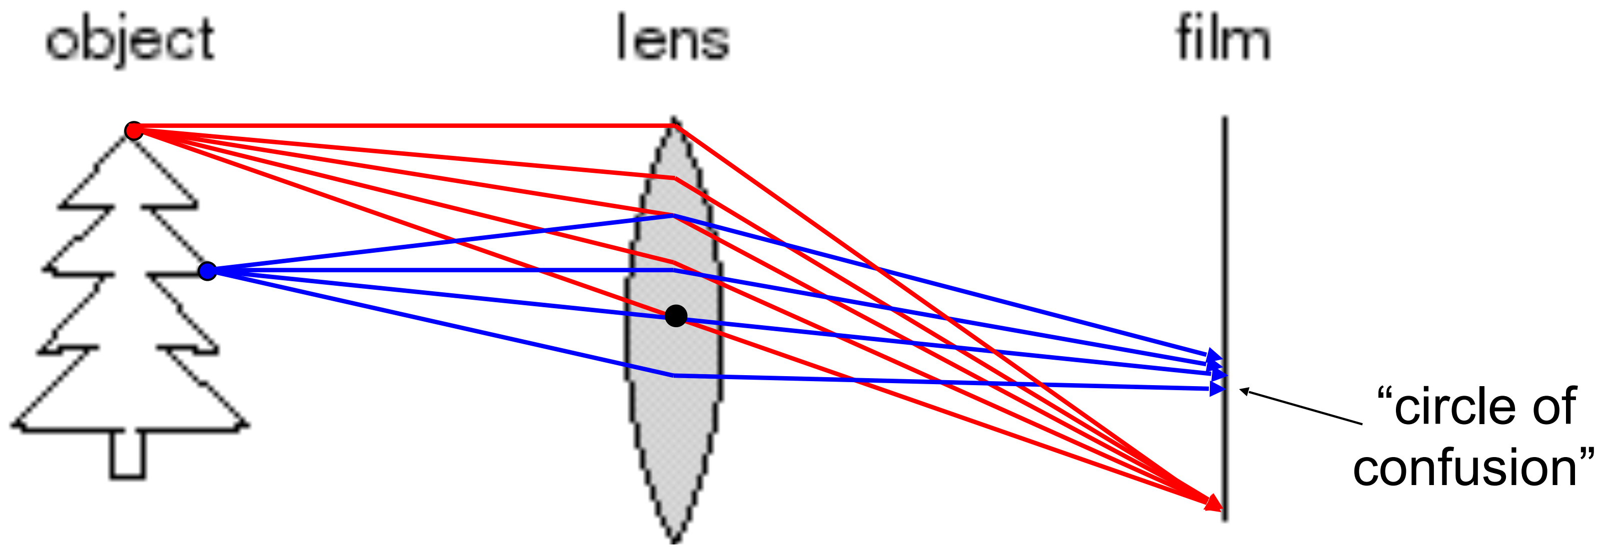

A lens focuses light onto the film

- There is a specific distance at which objects are "in focus"
  - Other points project to a "circle of confusion" in the image

Slide by Steve Seitz

Thin lens formula

- What is the relation between the focal length (f), the distance of the object from the optical center (D), and the distance at which the object will be in focus (D')?

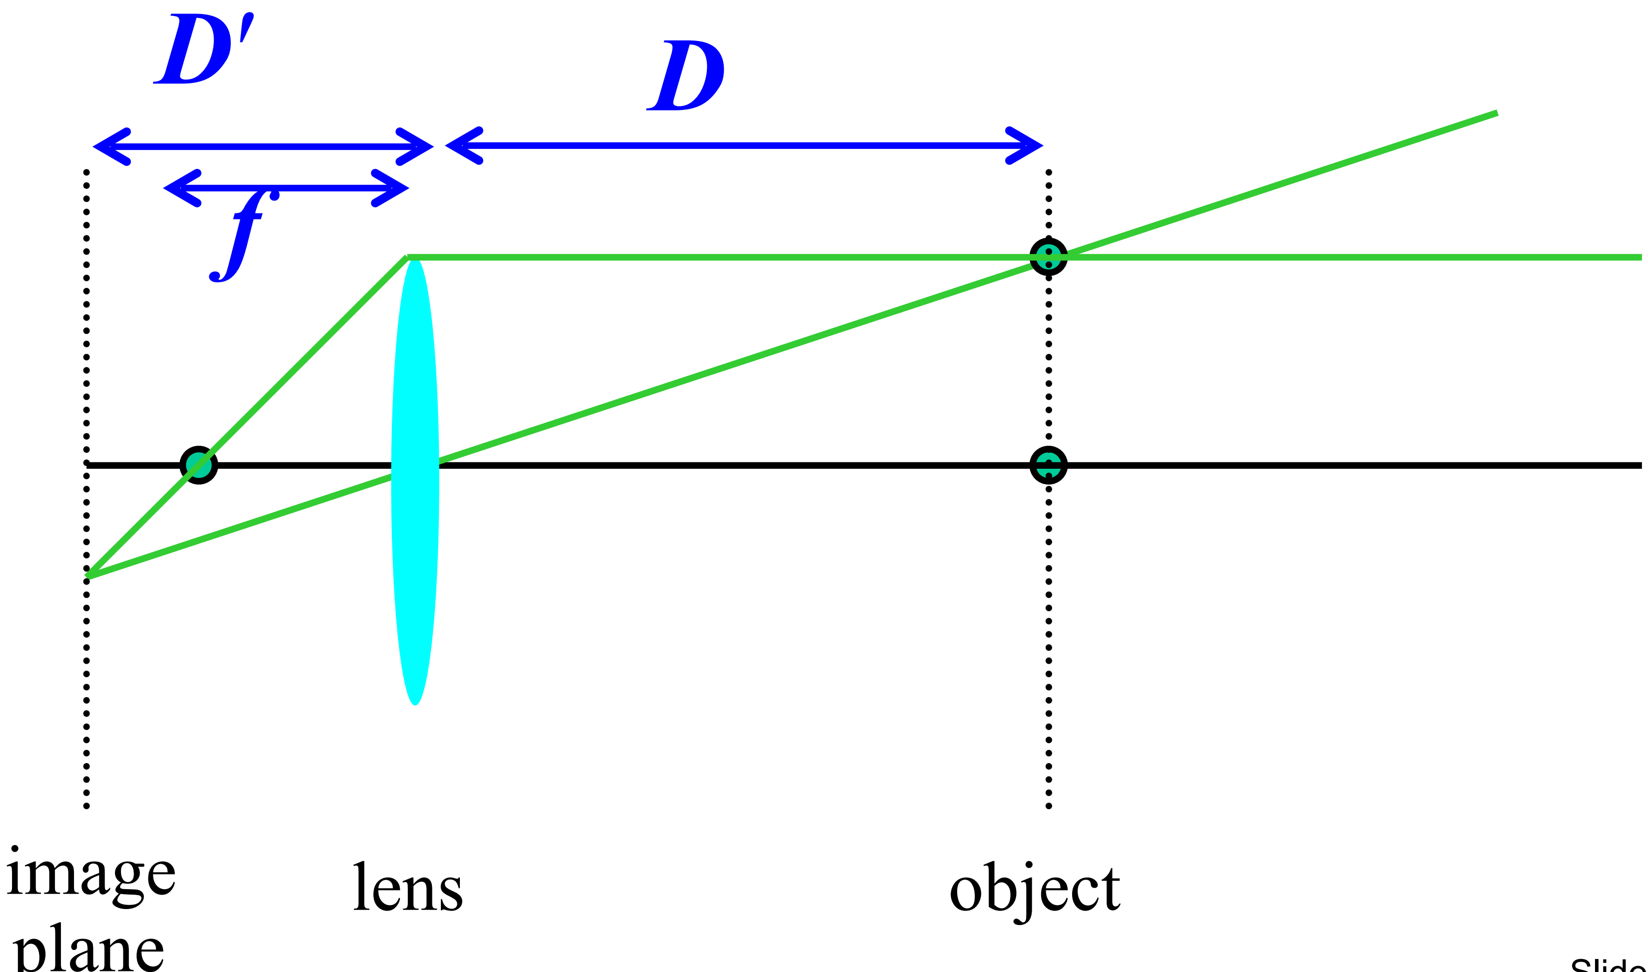

Slide by Frédo Durand

Thin lens formula

Similar triangles everywhere!

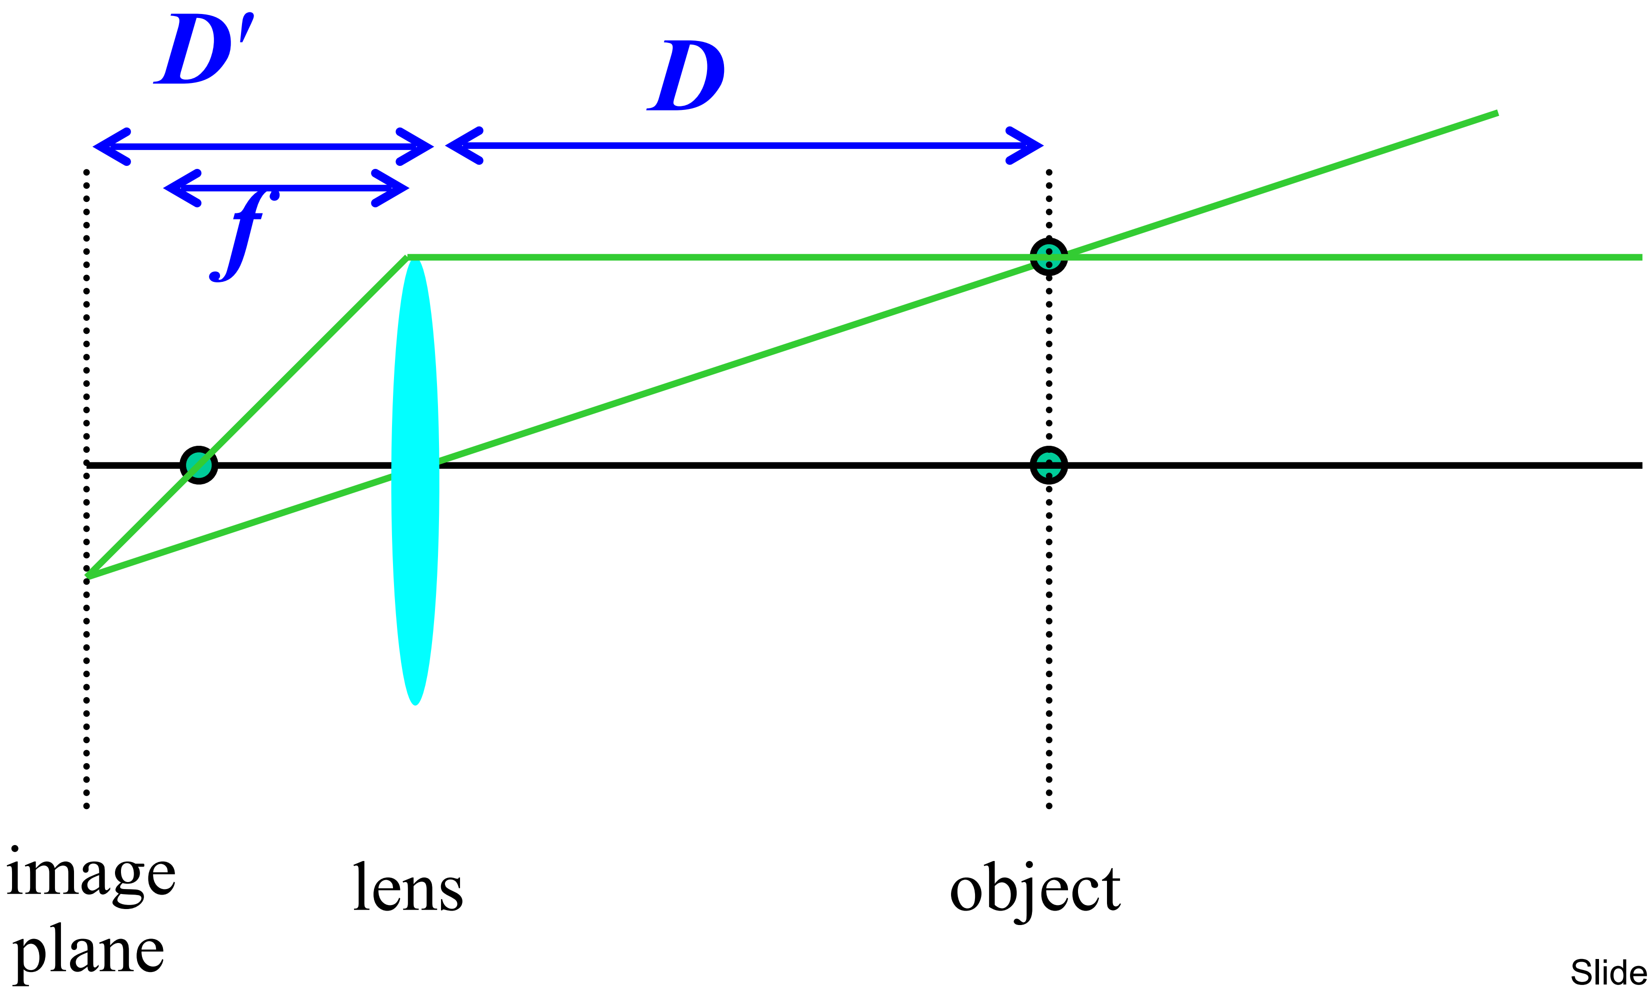

Slide by Frédo Durand

Thin lens formula

Similar triangles everywhere!

$$y / y' = D / D'$$

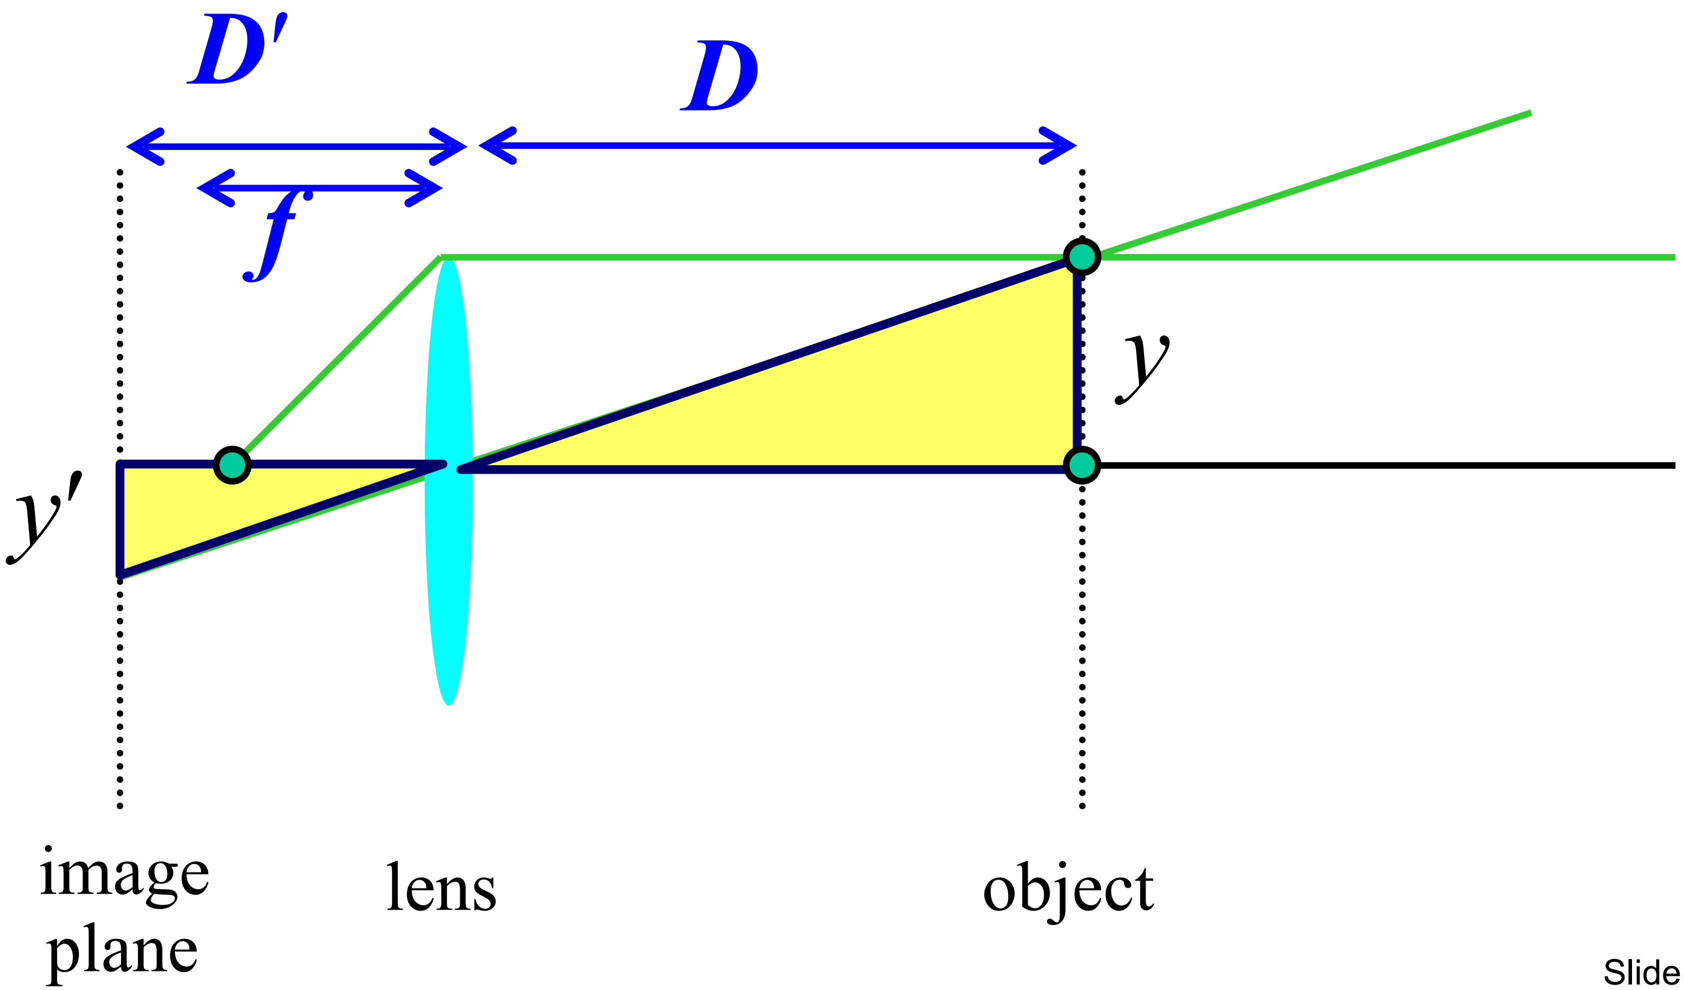

Slide by Frédo Durand

Thin lens formula

Similar triangles everywhere!

$$y / y' = D / D'$$

$$y / y' = f / (D' - f)$$

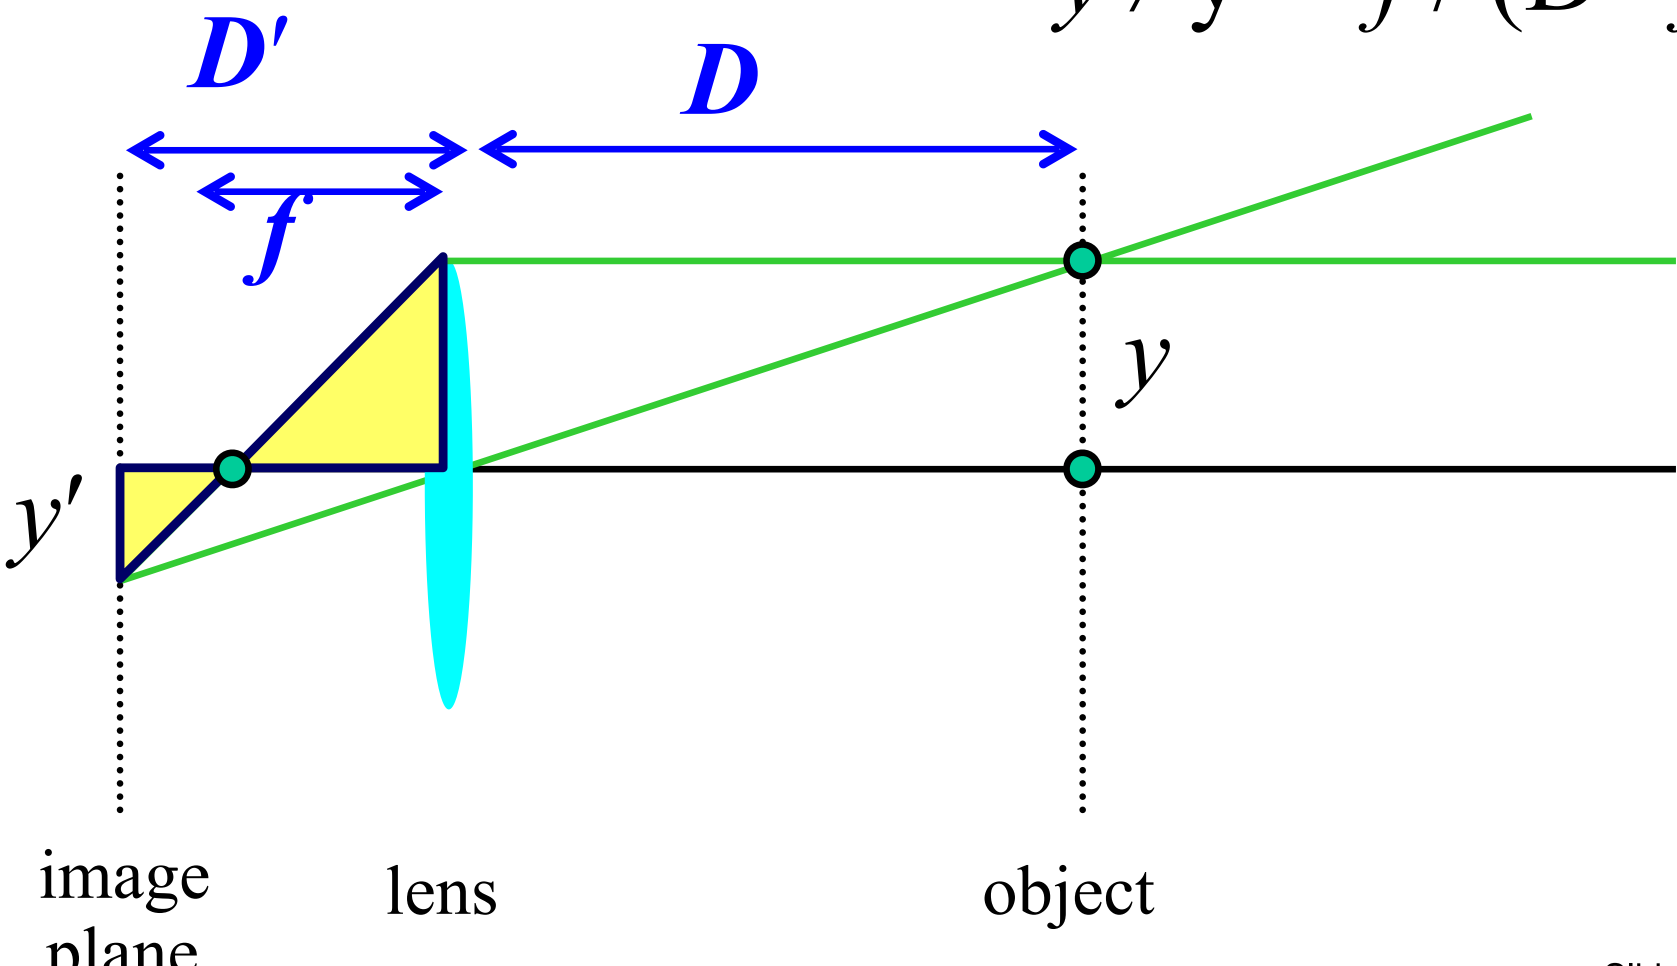

Slide by Frédo Durand

Thin lens formula

$$\frac{1}{D'} + \frac{1}{D} = \frac{1}{f}$$

Any point satisfying the thin lens equation is in focus.

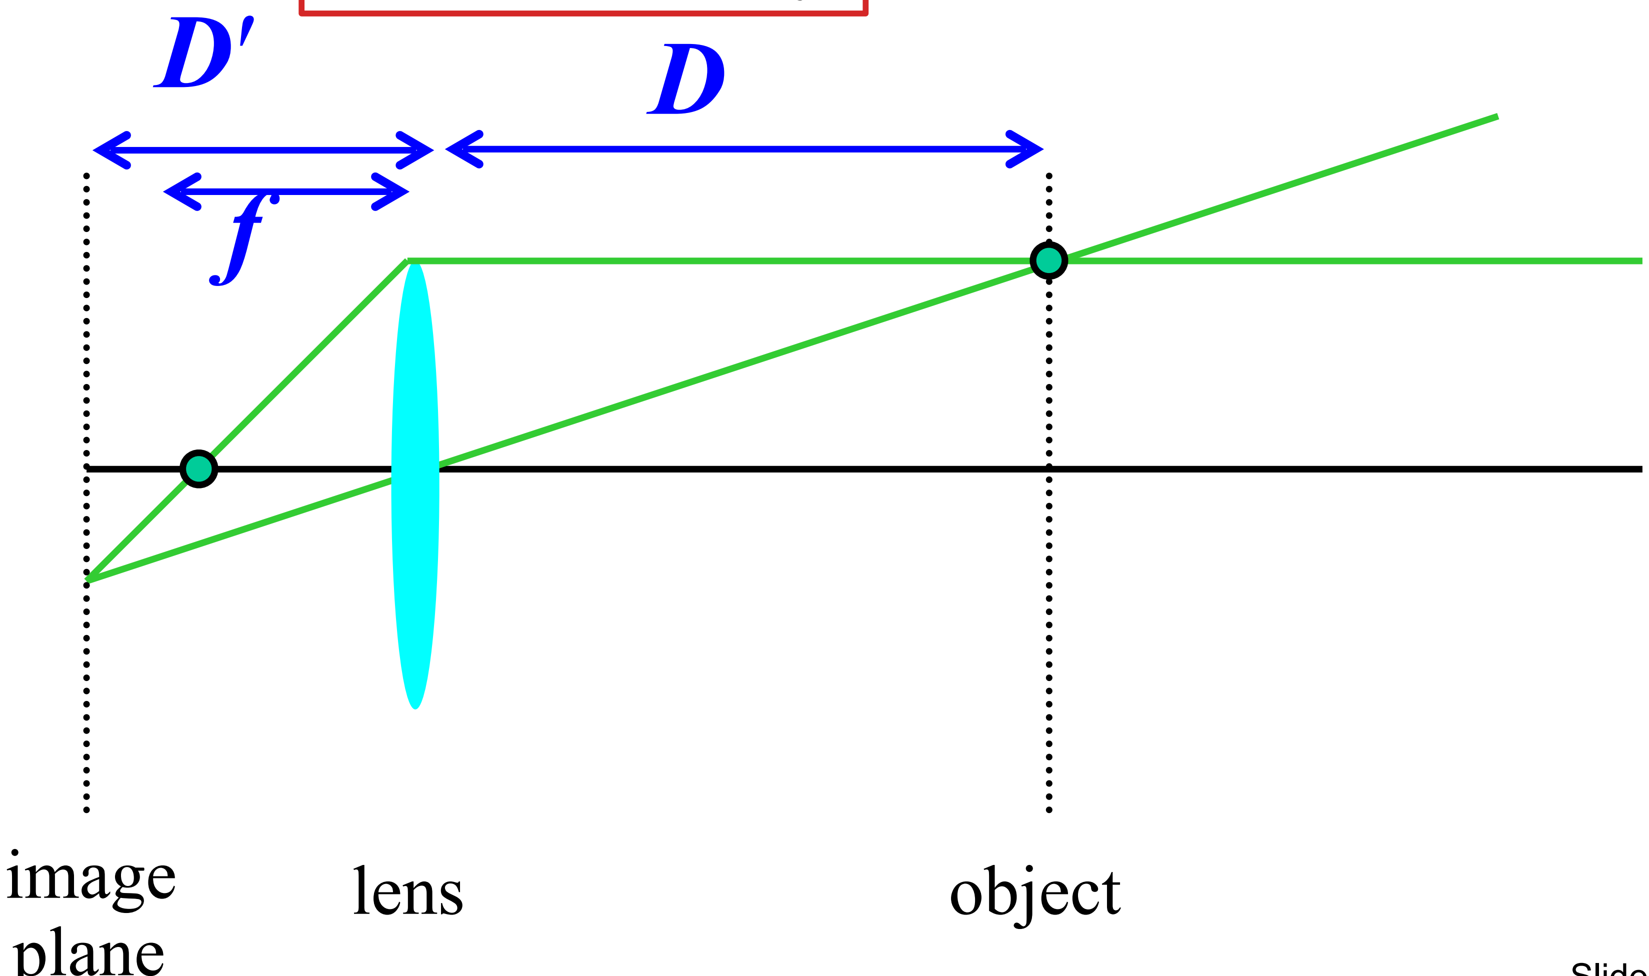

Slide by Frédo Durand

Focus and Depth of Field

- For a given *D*, “perfect” focus at only one *D’*
- In practice, OK for some range of depths
  - Circle of confusion smaller than a pixel

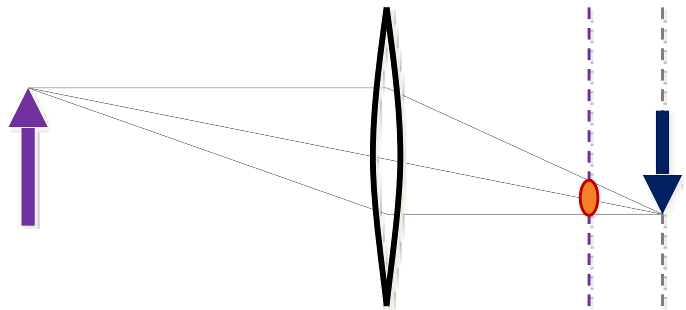

Depth of Field

http://www.cambridgeincolour.com/tutorials/depth-of-field.htm

Slide by A. Efros

Controlling depth of field

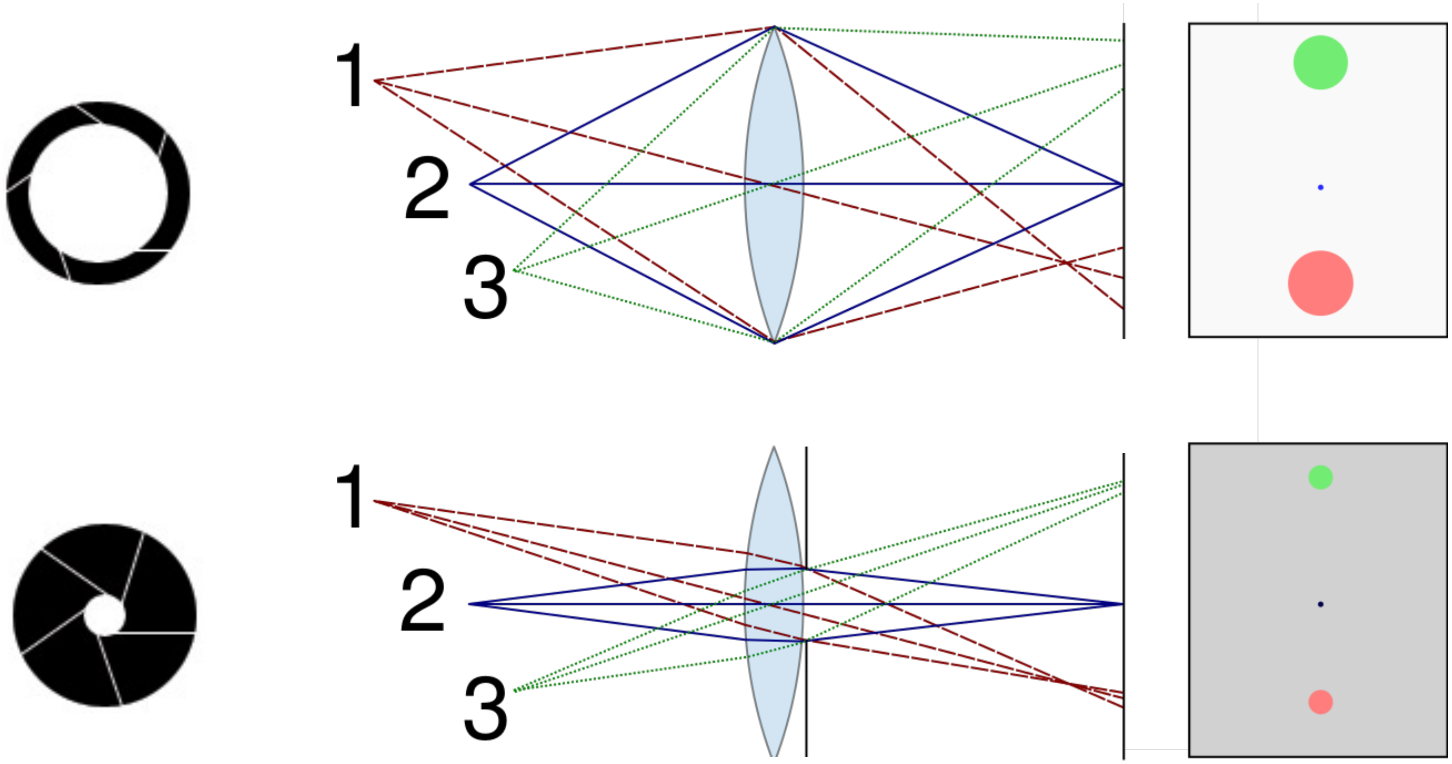

Changing the aperture size affects depth of field

- A smaller aperture increases the range in which the object is approximately in focus
- But small aperture reduces amount of light – need to increase exposure

http://en.wikipedia.org/wiki/File:Depth_of_field_illustration.svg

Slide by S. Lazebnik

Varying the aperture

Large aperture = small DOF

Small aperture = large DOF

Slide by A. Efros

Aperture

- Aperture typically given as "*f*-number"
- What is *f*/4?
  - Aperture *diameter* is ¼ the focal length
- One "*f*-stop" equals change of *f*-number by √2
  - Equals change in aperture *area* by factor of 2
  - Equals change in amount of light by factor of 2
  - Example: *f*/2 → *f*/2.8 → *f*/4 (each one halves the amount of light)

Building a real camera: field of view

Field of View

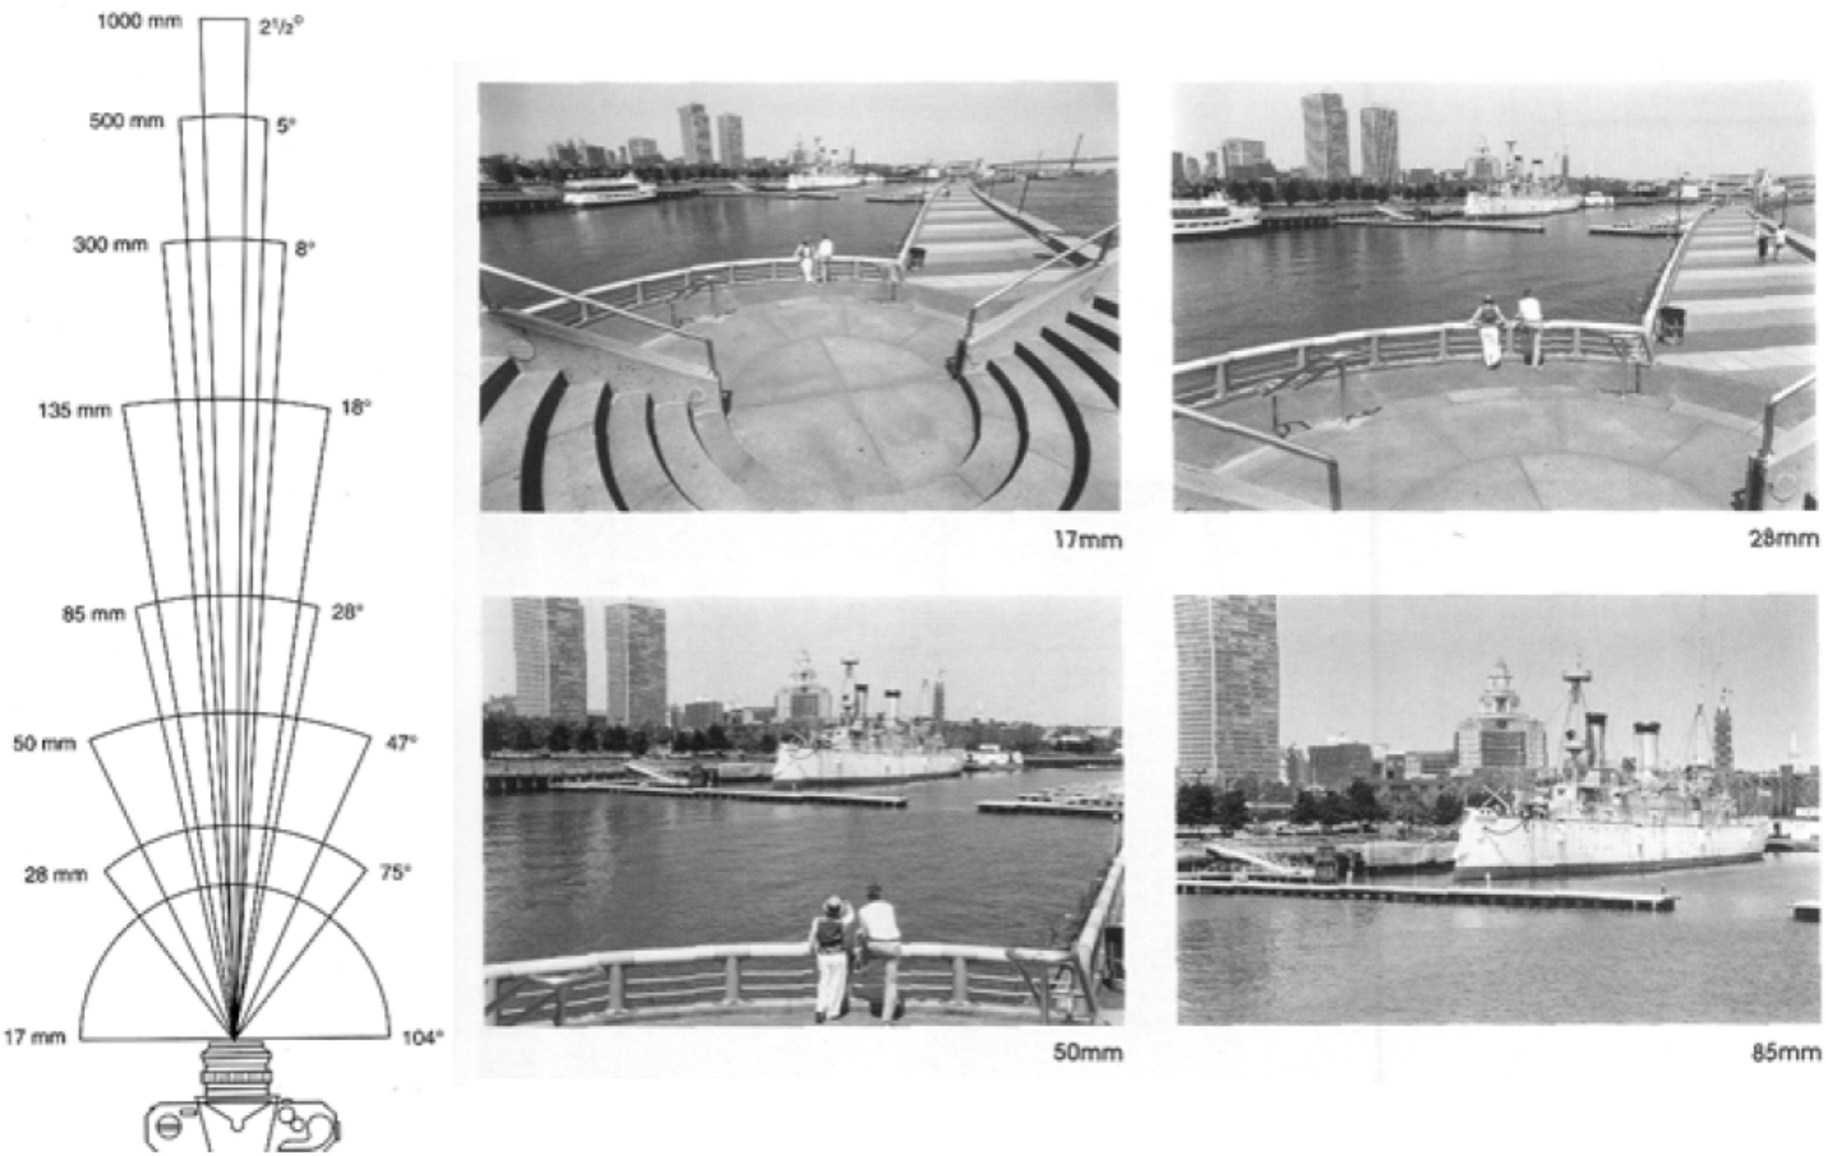

Slide by A. Efros

Field of View

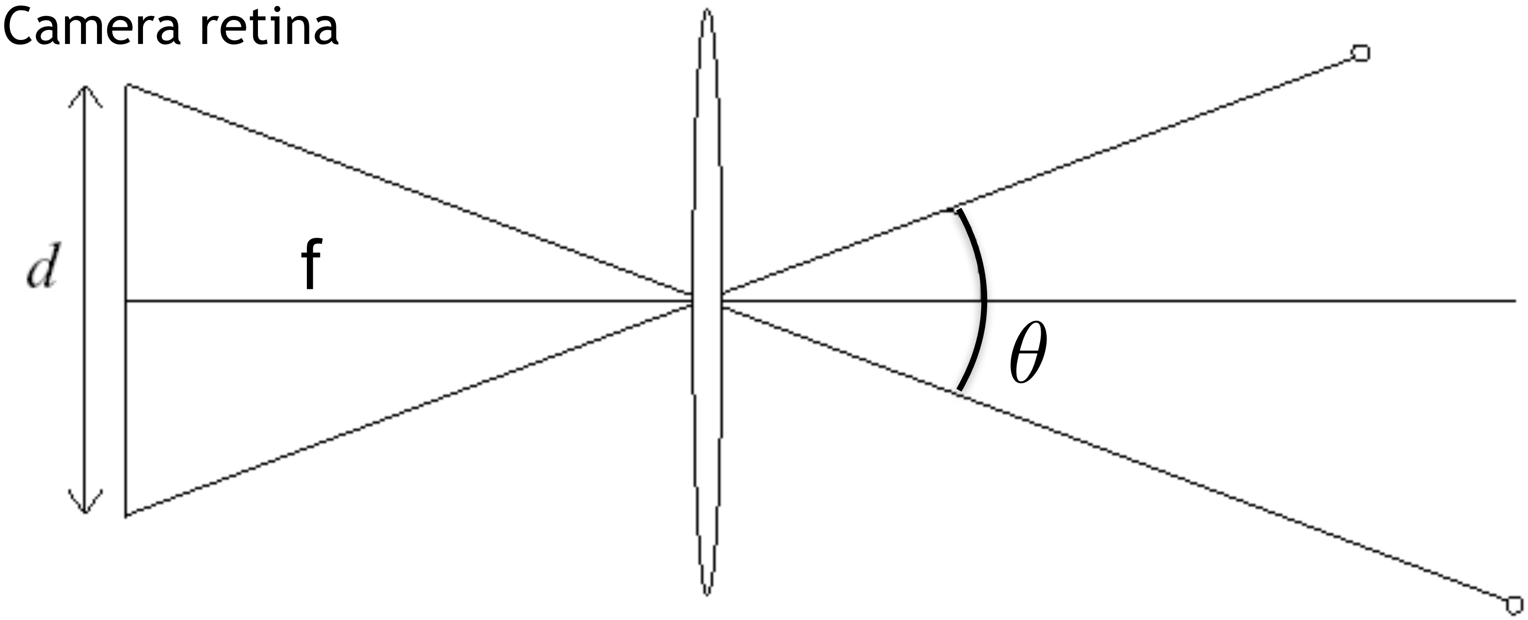

FOV depends on focal length and size of the camera retina

$$\theta = 2\tan^{-1}\left(\frac{d}{2f}\right) \approx \frac{d}{f}$$

Larger focal length = smaller FOV

Slide by A. Efros

Field of View / Focal Length

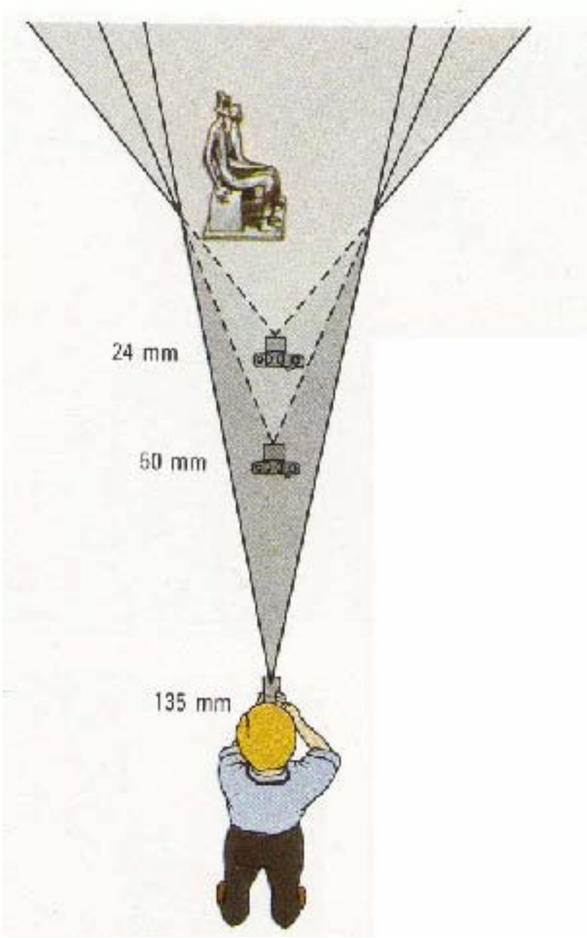

Large FOV, small *f*
Camera close to car

Small FOV, large *f*
Camera far from the car

Sources: A. Efros, F. Durand

A bit about digital cameras

Digital camera

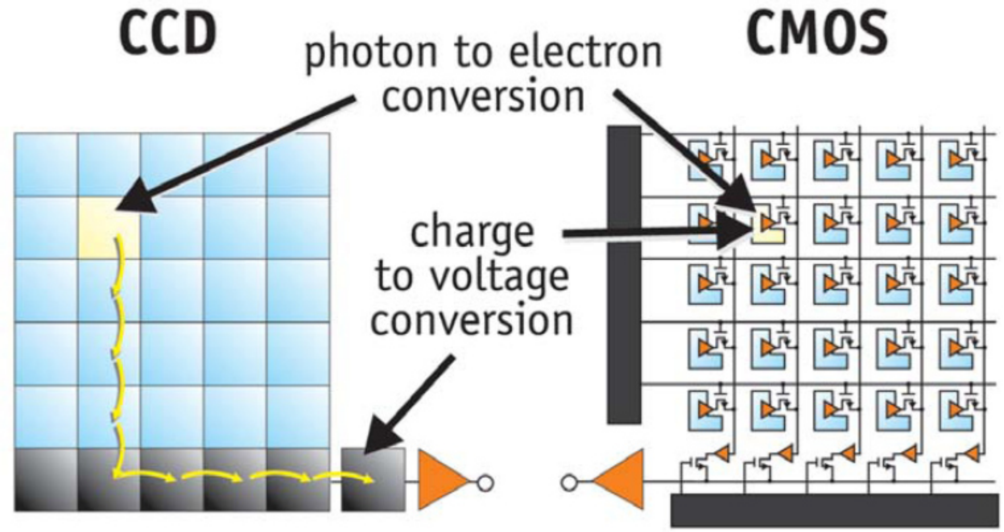

CCDs move photogenerated charge from pixel to pixel and convert it to voltage at an output node. CMOS imagers convert charge to voltage inside each pixel.

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)
  - **Complementary metal oxide semiconductor** (CMOS)
- http://electronics.howstuffworks.com/digital-camera.htm

Slide by Steve Seitz

Color sensing: Color filter array

Bayer grid

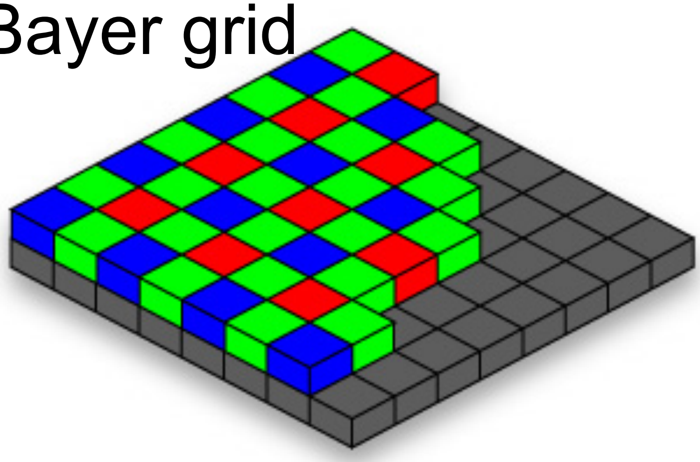

Estimate missing components from neighboring values (**demosaicing**)

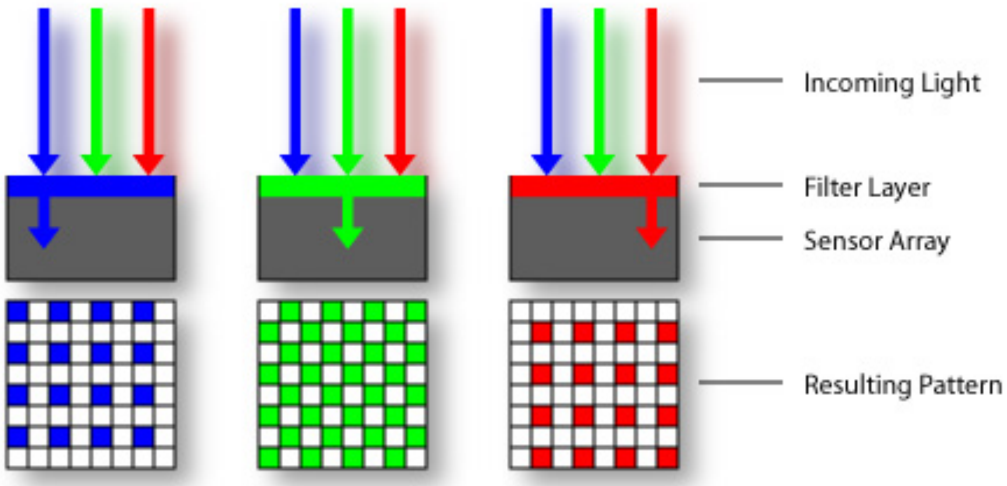

Why more green?

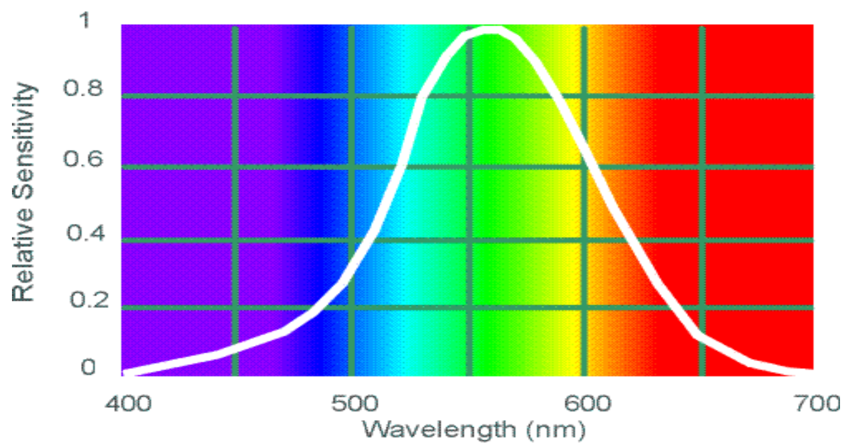

Human Luminance Sensitivity Function

Source: Steve Seitz

Next class: convolution and filtering

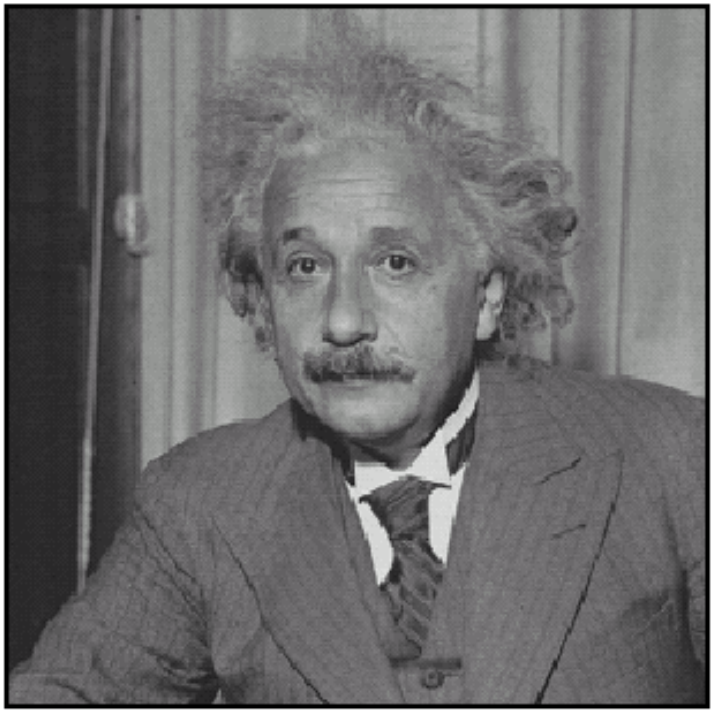

before

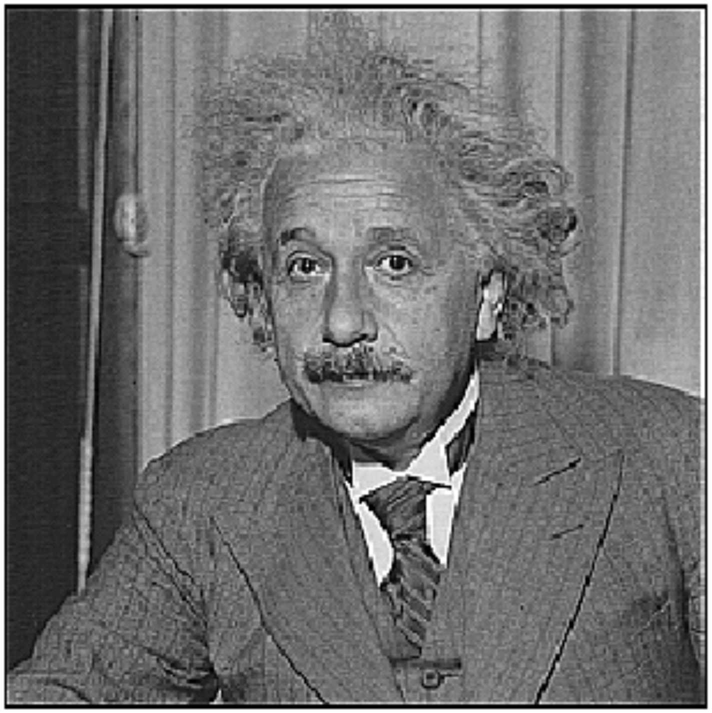

after