

Image formation

CS 178, Spring 2012

Begun April 3, finished April 5.



Marc Levoy
Computer Science Department
Stanford University

Outline


- ◆ perspective
 - natural versus linear perspective
 - vanishing points
 - ◆ image formation
 - pinhole cameras
 - lenses
-
- ◆ exposure
 - shutter speed
 - aperture
 - ISO
 - ◆ choosing a camera

The laws of perspective

- ◆ common assumptions
 1. Light leaving an object travels in straight lines.
 2. These lines converge to a point at the eye.

- ◆ natural perspective (Euclid, 3rd c. B.C.)
 - 3a. More distant objects subtend smaller visual angles.

The laws of perspective



- ◆ natural perspective (Euclid, 3rd c. B.C.)
 - 3a. More distant objects subtend smaller visual angles.

Roman wall paintings



from Villa Publius Fannius Synistor,
Boscotrecase, Pompeii (c. 40 B.C.)



Still life with peaches, from
Herculaneum (before 79 A.D.)

The laws of perspective



- ◆ common assumptions
 1. Light leaving an object travels in straight lines.
 2. These lines converge to a point at the eye.

- ◆ natural perspective (Euclid, 3rd c. B.C.)
 - 3a. More distant objects subtend smaller visual angles.

- ◆ linear perspective (Filippo Brunelleschi, 1413)
 - 3b. A perspective image is formed by the intersection of these lines with a “picture plane” (the canvas).

The laws of perspective

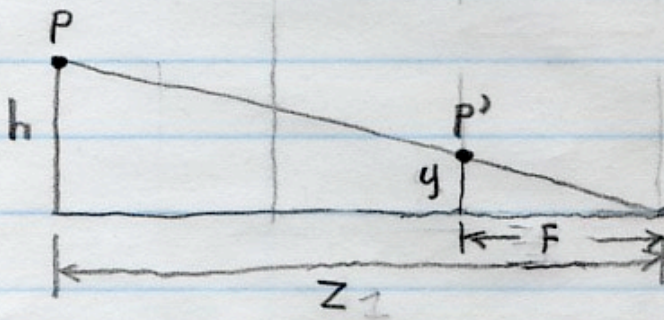
$$\frac{y_2}{y_1} \neq \frac{\theta_2}{\theta_1}$$



- ◆ natural perspective (Euclid, 3rd c. B.C.)
3a. More distant objects subtend smaller visual angles.
- ◆ linear perspective (Filippo Brunelleschi, 1413)
3b. A perspective image is formed by the intersection of these lines with a “picture plane” (the canvas).

Projection onto picture plane (contents of whiteboard)

I accidentally forgot to show this derivation in class today. It shows how you can use similar triangles to compute where in a planar canvas that lies f units from the eye one should place the projection P' of a point P in the scene that lies Z units from the eye. This simple calculation effectively implements linear perspective.



• How do we compute y , so we can place point P on p.p. as P' ?

By similar triangles, $\frac{h}{Z} = \frac{y}{F}$, so $y = \frac{h \cdot F}{Z}$.


(Perspective requires division by z !)

- ◆ the division by z means that the size of an object in a photograph is inversely proportional to its distance from the camera




Filippo Brunelleschi,
dome of the cathedral,
Florence (1419)

The problem of drawing pavimento





Giovanni de Paolo, Birth of St. John the Baptist (1420)

Alberti's method (1435)



(Cole)






(Cole)


Piero della Francesca, The Flagellation (c.1460)

Vanishing points


Q. How many vanishing points can there be in a perspective drawing?



1-point



2-point

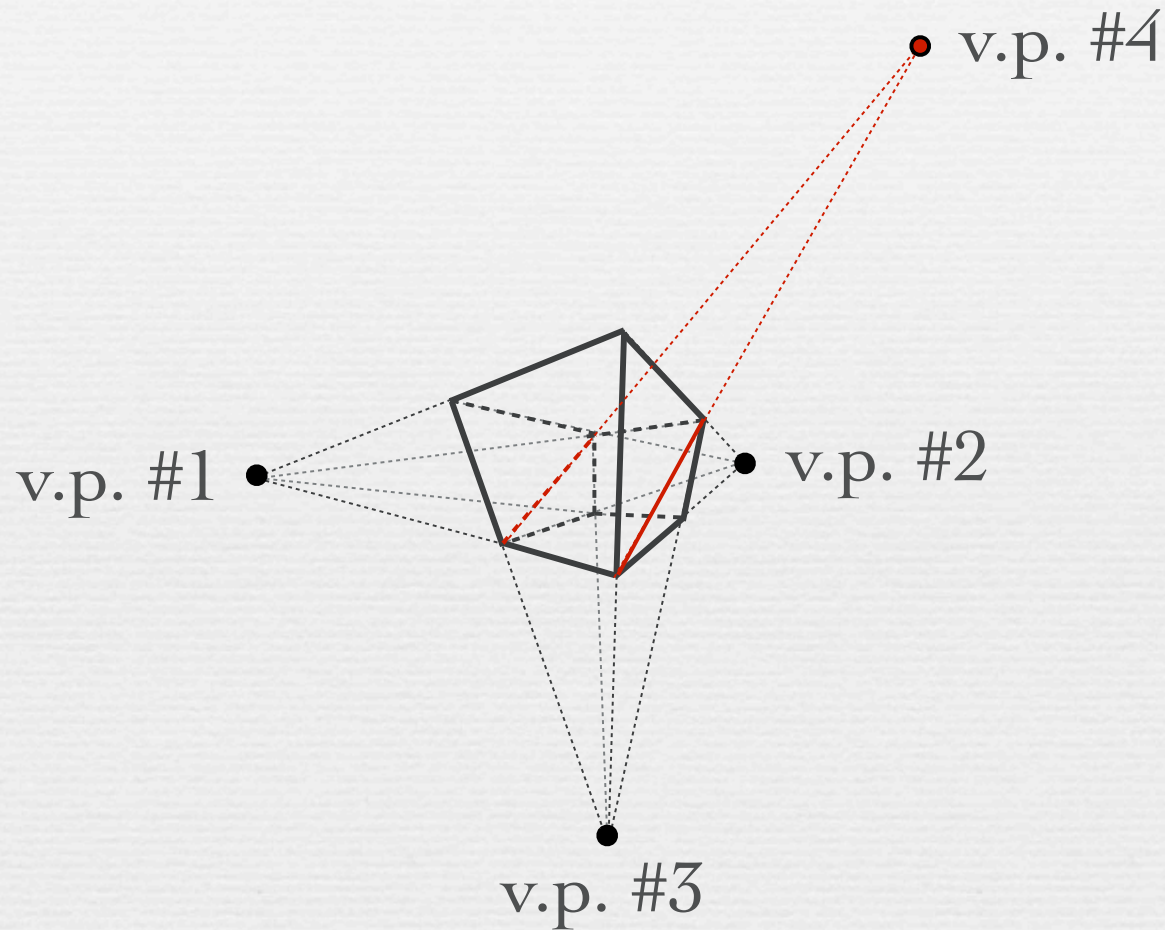


3-point

(D'Amelio)

© Marc Levoy


Example of a 4th vanishing point



- ◆ each direction of parallel lines will converge to a unique vanishing point

Q. Should the distant ends of a long facade be drawn smaller than its center in a perspective drawing?

Imagine the figure in slide #8 as the plan view of a bay in the arcade, where h is the distance on the ground between two columns. To figure out what happens to faraway columns in a linear perspective drawing, imagine replacing h with $10 \times h$ on that figure. From similar triangles, and the resulting formula in slide #8, it's clear that the proper placement of this faraway column in the perspective drawing will be $10 \times y$. Thus, the spacing between columns in the drawing will be constant, assuming the picture plane for the drawing is parallel to the facade containing the columns. (I stumbled over this point in class, but eventually got it right.)



?

- ◆ no, in linear perspective straight lines remain straight
- ◆ lines parallel to the picture plane do not converge
- ◆ they appear smaller when you view the drawing, due to natural perspective (angles subtended at eye)

Recap

- ◆ natural perspective
 - visual angle subtended by a feature in the world
- ◆ linear perspective
 - intersections of lines of sight with a picture plane
 - the correct way to make a drawing on a flat surface
- ◆ vanishing points
 - one per direction of line in the scene
 - lines parallel to the picture plane do not converge


Questions?

Single lens reflex camera (SLR)



Nikon F4
(film camera)


Why not use sensors without optics?



(London)

- ◆ each point on sensor would record the integral of light arriving from every point on subject
- ◆ all sensor points would record similar colors

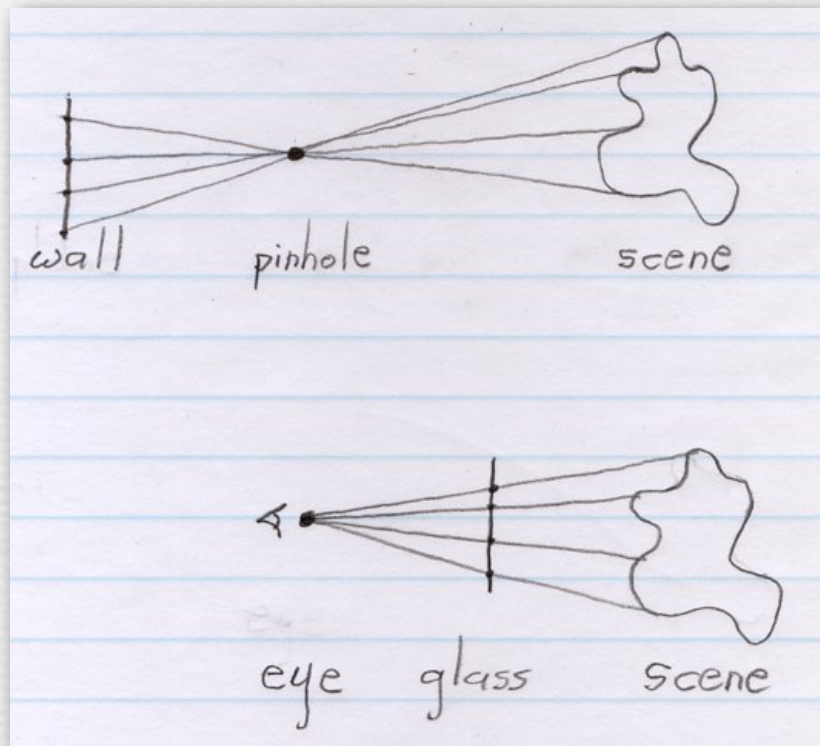
Pinhole camera (a.k.a. *camera obscura*)




- ◆ linear perspective with viewpoint at pinhole
- ◆ tilting the picture plane changes the number and location of vanishing points

Equivalence of Dürer's glass and *camera obscura* (contents of whiteboard)

camera obscura




Dürer's glass



- ◆ both devices compute 2D planar geometric projections,
i.e. projections along straight lines through a point and onto a plane
 - the images differ only in scale (and a reflection around the origin)

Pinhole photography

- ◆ no distortion
 - straight lines remain straight
- ◆ infinite depth of field
 - everything is in focus



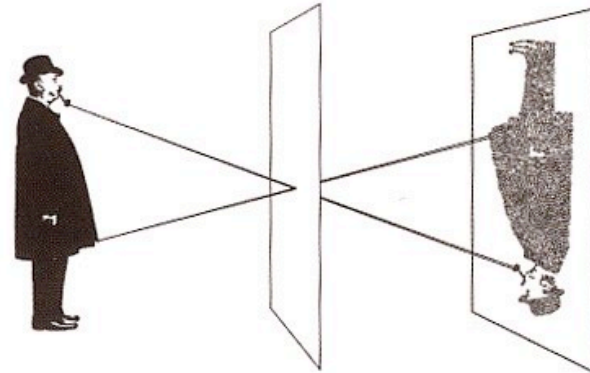
(Bami Adedoyin)




© Marc Levoy

Effect of pinhole size

Photograph made with small pinhole

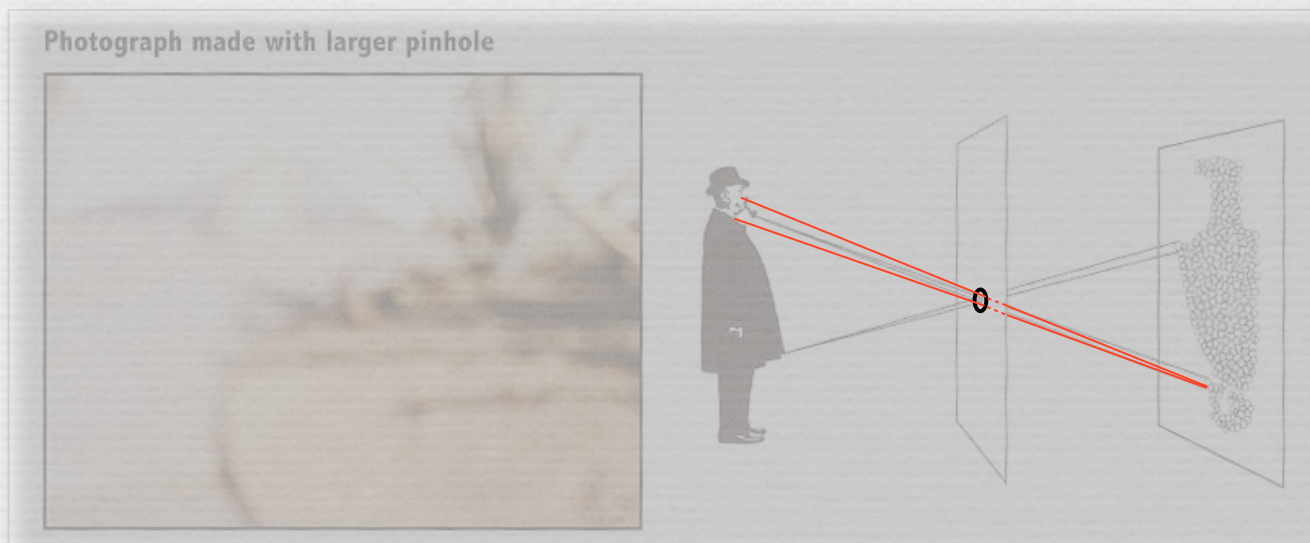



Photograph made with larger pinhole



(London)


Effect of pinhole size




(London)

Replacing the pinhole with a lens

Photograph made with small pinhole




Photograph made with lens




(London)

Replacing the pinhole with a lens

Photograph made with small pinhole



- ◆ a photographic camera produces the same 2D geometric projection as a *camera obscura*
 - a lens replaces the pinhole, and film or digital sensor replaces the wall
 - rotating the camera around the center of the lens adds or removes vanishing points




(London)

Geometrical optics


- ◆ parallel rays converge to a point located at focal length f from lens



- ◆ rays going through center of lens are not deviated
 - hence same perspective as pinhole



Gauss' ray tracing construction




- ◆ rays coming from points on a plane parallel to the lens are focused on another plane parallel to the lens

Changing the focus distance


- ◆ to focus on objects at different distances, move sensor relative to lens
- ◆ in a handheld camera, one actually moves the lens, not the sensor

by convention, the “focus distance” is on the object side of the lens



Changing the focal length

- ◆ weaker lenses have longer focal lengths
- ◆ to stay in focus, move the sensor further back
- ◆ focused image of tree is located slightly beyond the focal length

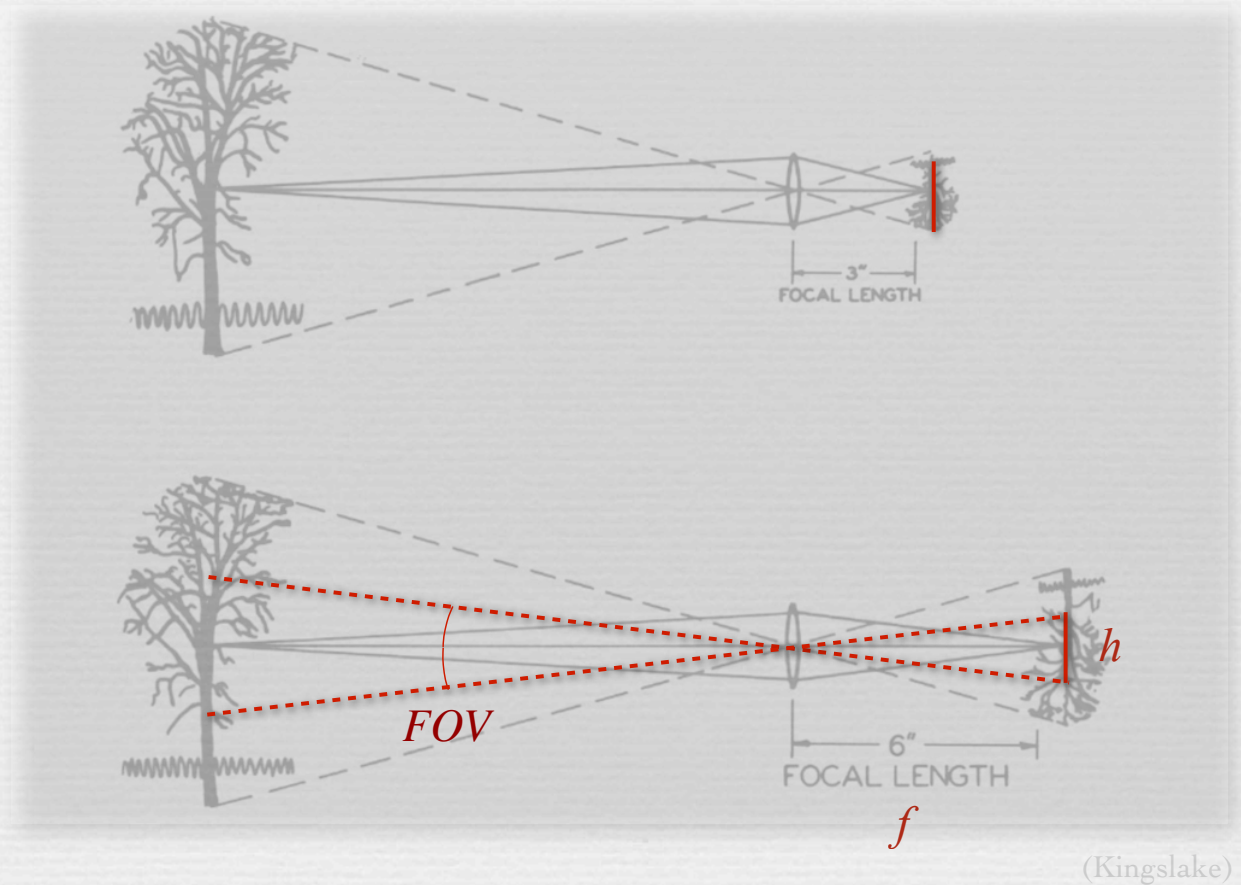


(Kingslake)

The tree would be in focus at the lens focal length only if it were infinitely far away.


Changing the focal length

- ◆ if the sensor size is constant, the field of view becomes smaller



$$FOV = 2 \arctan (h / 2f)$$

Focal length and field of view



17mm



28mm



50mm




85mm


(London)

FOV measured diagonally on a 35mm full-frame camera (24 × 36mm)


Focal length and field of view



135mm



300mm



500mm

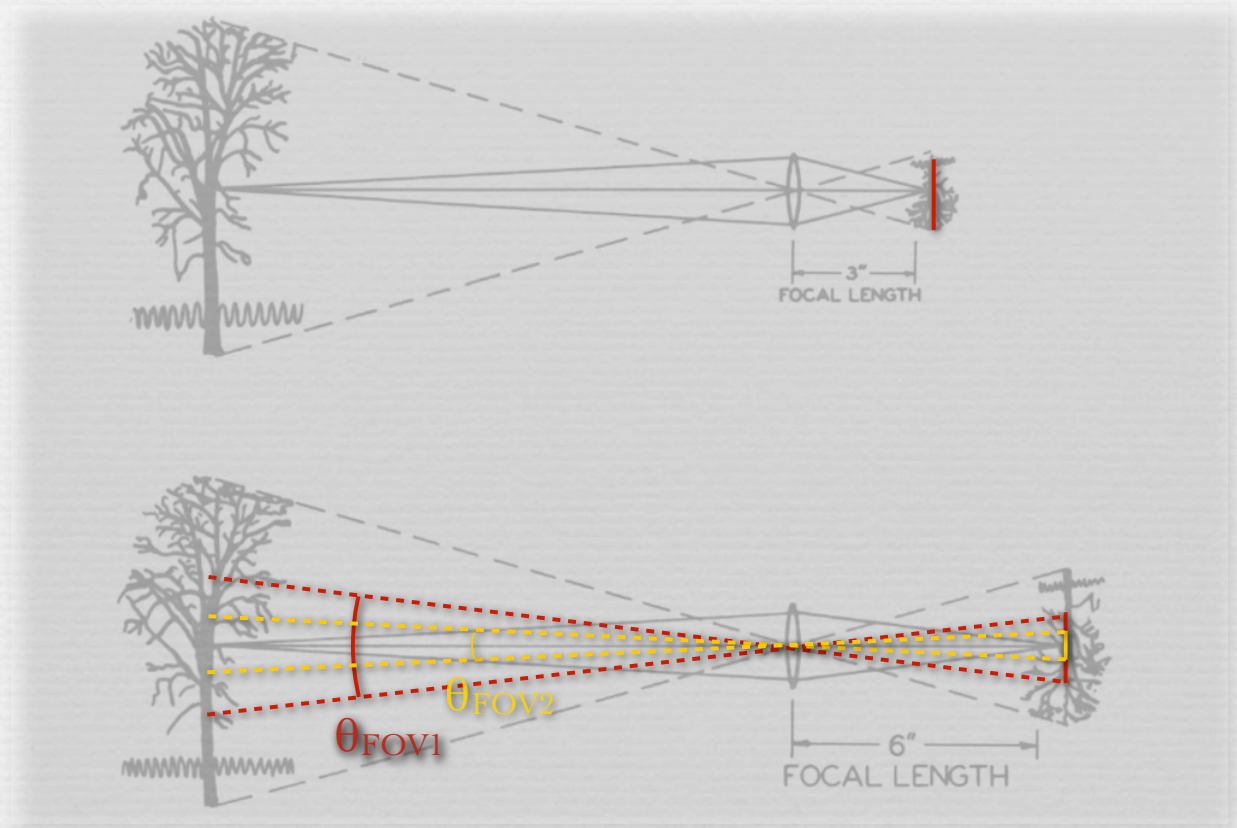


(London)

FOV measured diagonally on a 35mm full-frame camera (24 × 36mm)

Changing the sensor size

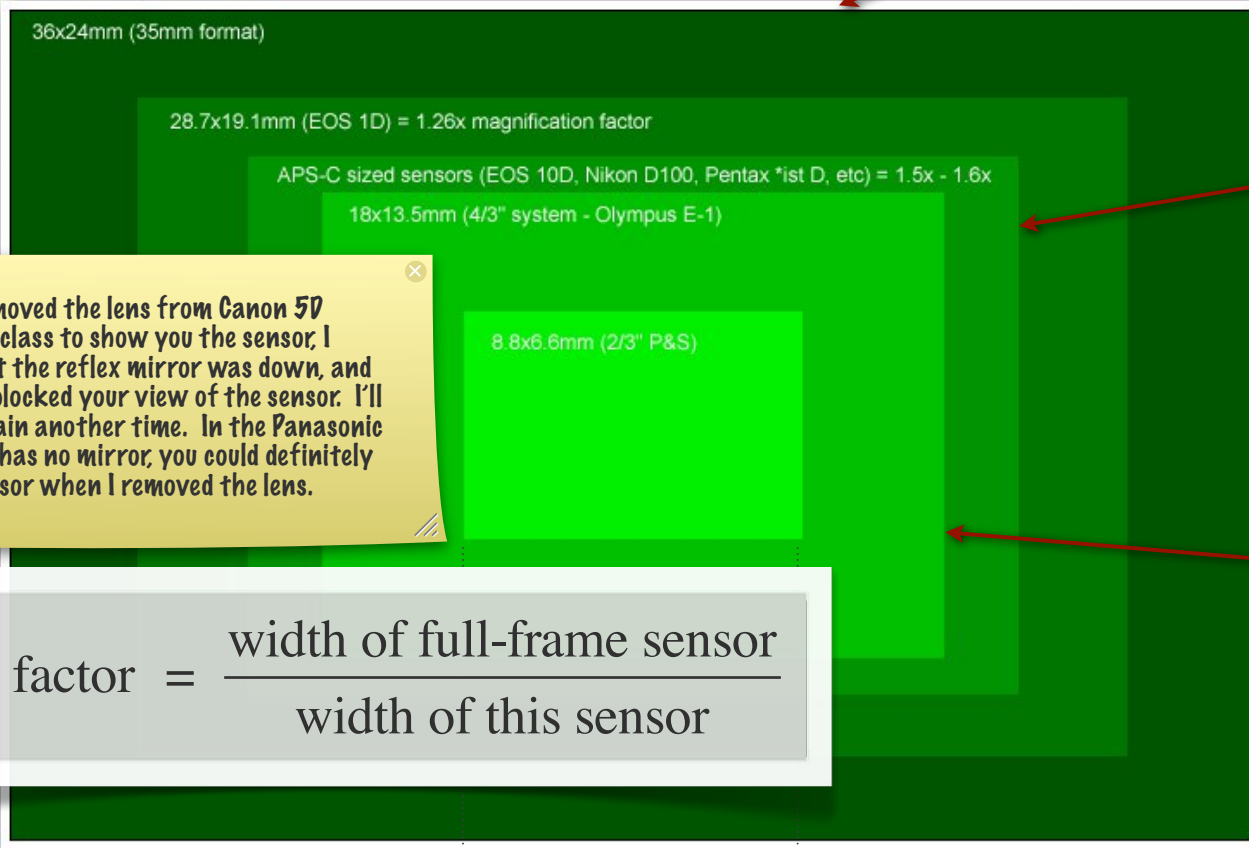
- ◆ if the sensor size is smaller, the field of view is smaller too
- ◆ smaller sensors either have fewer pixels, or smaller pixels, which are noisier



(Kingslake)

Sensor sizes

“full frame”
Canon 5D Mark II
(24mm × 36mm)



“APS-C”
Nikon D40
(15.5mm × 23.7mm)
(~1.5× crop factor)

“micro 4/3”
Panasonic GF1
(13mm × 17.3mm)
(~2× crop factor)

“point-and-shoot”
Canon A590
(5.75mm × 4.31mm)
(~6× crop factor)

When I removed the lens from Canon 5D Mark II in class to show you the sensor, I forgot that the reflex mirror was down, and therefore blocked your view of the sensor. I'll show it again another time. In the Panasonic GF1, which has no mirror, you could definitely see the sensor when I removed the lens.

$$\text{crop factor} = \frac{\text{width of full-frame sensor}}{\text{width of this sensor}}$$

iPhone 4S
is 4.5mm ×
3.4mm





Paris, 2009 (Panasonic GF1 + Leica 90mm)

Changing the focal length versus changing the viewpoint

(Kingslake)



(a)

wide-angle



(b)




(c)

telephoto and
moved back

- ◆ changing the focal length lets us move back from a subject, while maintaining its size on the image
- ◆ but moving back changes perspective relationships

Changing the focal length versus changing the viewpoint

- ◆ moving back while changing the focal length lets you keep objects at one depth the same size
- ◆ in cinematography, this is called the dolly zoom, or “Vertigo effect”, after Alfred Hitchcock’s movie



Effect of focal length on portraits

- ◆ standard “portrait lens” is 85mm



wide angle



standard



telephoto

Recap

- ◆ pinhole cameras compute correct linear perspectives
 - but dark
- ◆ lenses gather more light
 - but only one plane of scene is in focus
 - focus by moving the sensor or lens
- ◆ focal length determines field of view
 - from wide angle to telephoto
 - depends on sensor size


more in the lens lectures next week

Questions?

Exposure


- ◆ $H = E \times T$
- ◆ exposure = irradiance \times time
- ◆ irradiance (E)
 - amount of light falling on a unit area of sensor per second
 - controlled by aperture
- ◆ exposure time (T)
 - in seconds
 - controlled by shutter

Single lens reflex camera



Nikon F4
(film camera)


Shutters




- ◆ quiet
- ◆ slow
(max 1/500s)
- ◆ out of focus
- ◆ need one
per lens






Shutters



- ◆ quiet
- ◆ slow
(max 1/500s)
- ◆ out of focus
- ◆ need one
per lens



(London)



- ◆ loud
- ◆ fast
(max 1/4000)
- ◆ in focus
- ◆ distorts motion



Jacques-Henri Lartigue, Grand Prix (1912)

Shutter speed

- ◆ controls how long the sensor is exposed to light
- ◆ linear effect on exposure until sensor saturates
- ◆ denoted in fractions of a second:
 - 1/2000, 1/1000, ..., 1/250, 1/125, 1/60, ..., 15, 30, B(ulb)
- ◆ normal humans can hand-hold down to 1/60 second
 - *rule of thumb*: shortest exposure = $1 / f$
 - e.g. 1/180 second for a 180mm lens

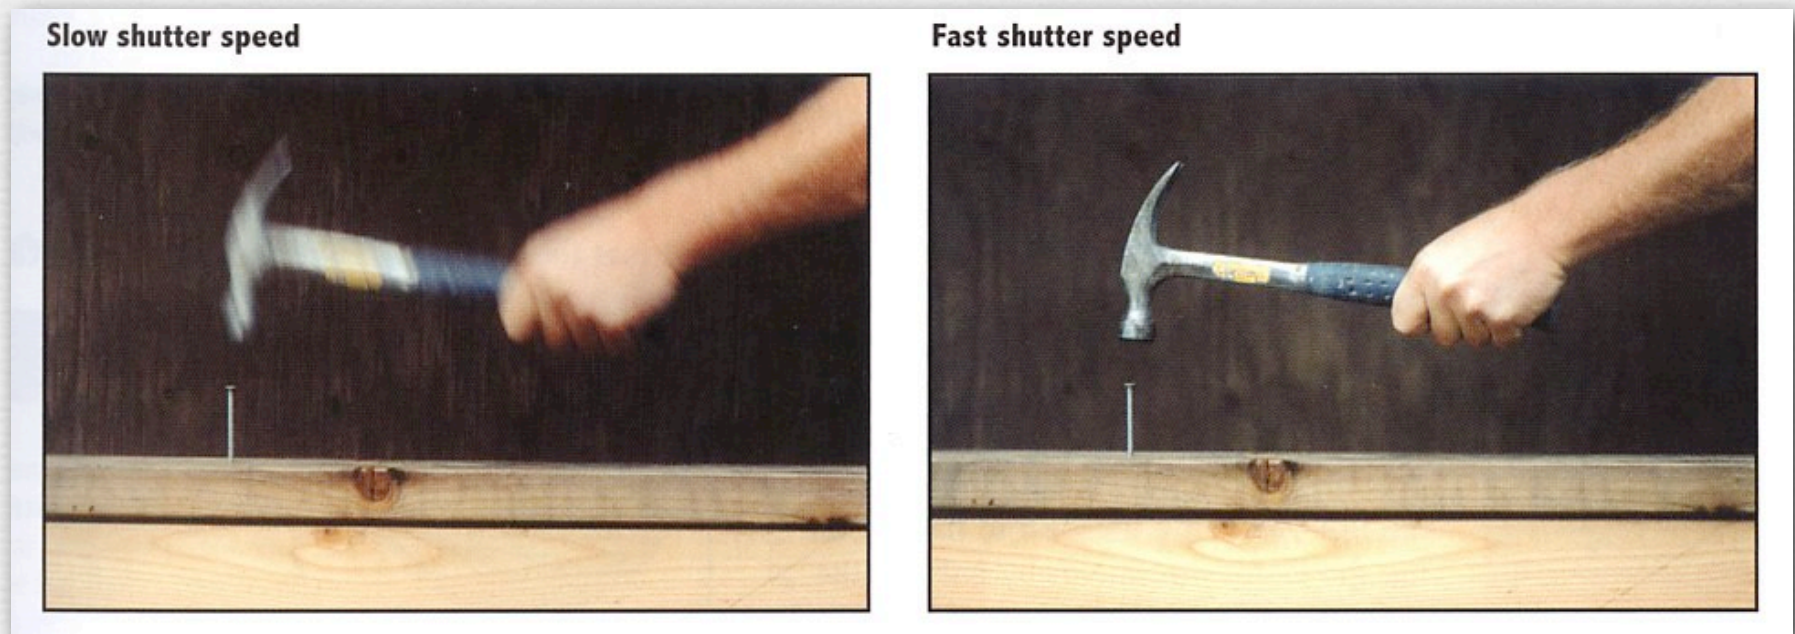
using 35mm
equivalent
focal length

GF1 (2× crop) + Leica 90mm



Main side-effect of shutter speed

- ◆ motion blur
- ◆ doubling exposure time doubles motion blur



(London)

Useful shutter speeds



1/25 sec (lucky!)

1/40 sec



Useful shutter speeds



1/250 sec

1/125 sec



Useful shutter speeds

1/2500 sec



1/800 sec



How fast is a volleyball spike?

Women's
volleyball

(Canon 1DIII,
1/800 second)




- ◆ work out shutter speed from motion in #pixels / second
- ◆ 5 pixels in 1/800sec = 4000 pixels/sec \Rightarrow 1/4000 sec !

focal plane
shutter distortion

Women's volleyball

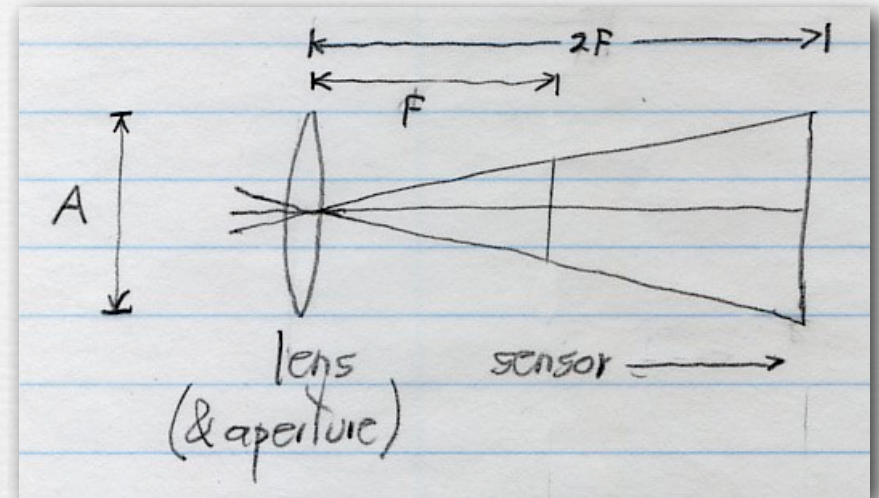

(Canon 1DIII,
1/800 second)



Aperture

- ◆ irradiance on sensor is proportional to
 - square of aperture diameter A
 - inverse square of distance to sensor (\sim focal length f)

Irradiance on sensor (contents of whiteboard)



- ◆ As the diameter A of the aperture doubles, its area (hence the light that can get through it) increases by $4\times$ (first drawing).
- ◆ Think of the lens as a collection of pinholes, each having a fixed angular field of view (cone in 2nd drawing) determined by the lens design.
- ◆ A certain amount of light gets through each pinhole. By conservation of energy, that light will fall on whatever sensor is placed in its path.
- ◆ If the distance to the sensor is doubled, the area intersecting the cone increases by $4\times$, so the light falling per unit area decreases by $4\times$.

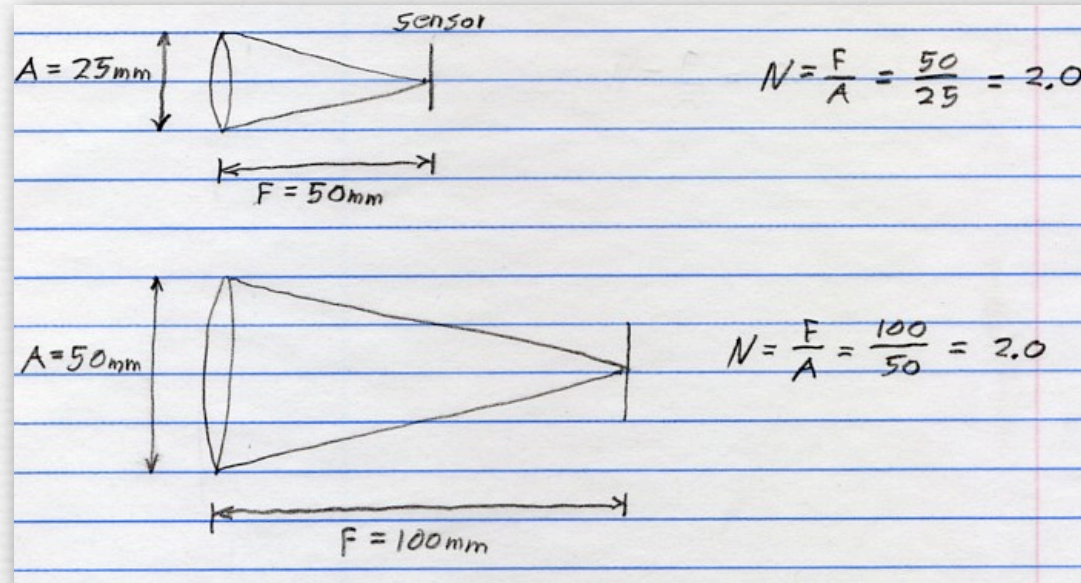
Aperture

- ◆ irradiance on sensor is proportional to
 - square of aperture diameter A
 - inverse square of distance to sensor (\sim focal length f)
- ◆ so that aperture values give irradiance regardless of lens, *aperture number* N is defined relative to focal length

$$N = \frac{f}{A}$$

- $f/2.0$ on a 50mm lens means the aperture is 25mm
- $f/2.0$ on a 100mm lens means the aperture is 50mm
- \therefore low F-number (N) on long telephotos require fat lenses

Example F-number calculations (contents of whiteboard)



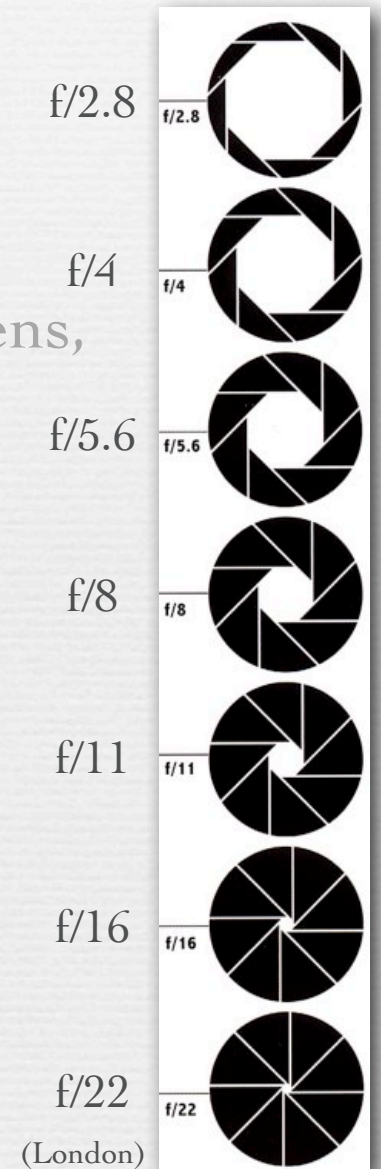
- ◆ A relative aperture size (called F-number or just N) of 2 is sometimes written f/2, reflecting the fact that the absolute aperture (A) can be computed by dividing focal length (f) by the relative aperture (N).
- ◆ As this drawing shows, doubling both the absolute aperture diameter (A) and the focal length (f) cancel; leaving the same relative aperture size (N). In this example, both lenses are f/2.

Aperture

- ◆ irradiance on sensor is proportional to
 - square of aperture diameter A
 - inverse square of distance to sensor (\sim focal length f)
- ◆ so that aperture values give irradiance regardless of lens, *aperture number* N is defined relative to focal length

$$N = \frac{f}{A}$$


- $f/2.0$ on a 50mm lens means the aperture is 25mm
- $f/2.0$ on a 100mm lens means the aperture is 50mm
- \therefore low F-number (N) on long zooms require fat lenses
- ◆ doubling N reduces A by $2\times$, hence light by $4\times$
 - going from $f/2.0$ to $f/4.0$ cuts light by $4\times$
 - to cut light by $2\times$, increase N by $\sqrt{2}$



Main side-effect of aperture

- ◆ depth of field
- ◆ doubling N (two f/stops) doubles depth of field

Large aperture opening



Small aperture opening

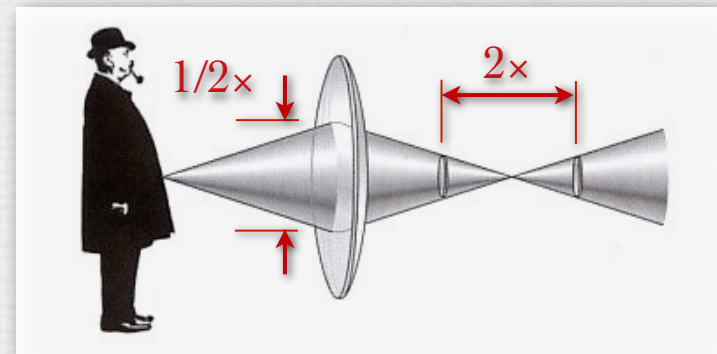
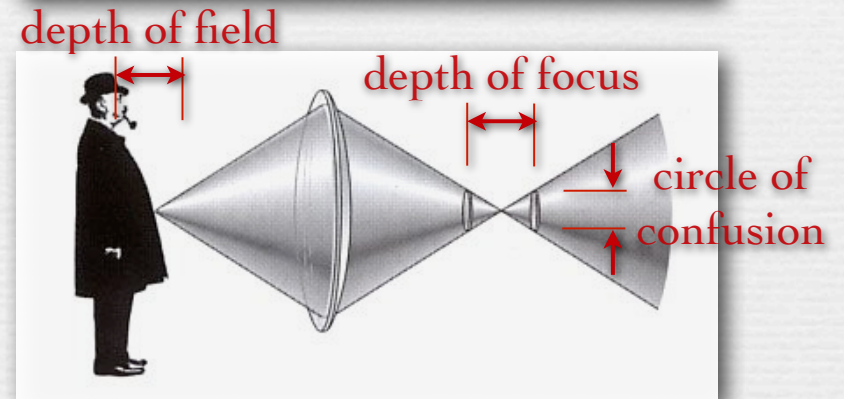
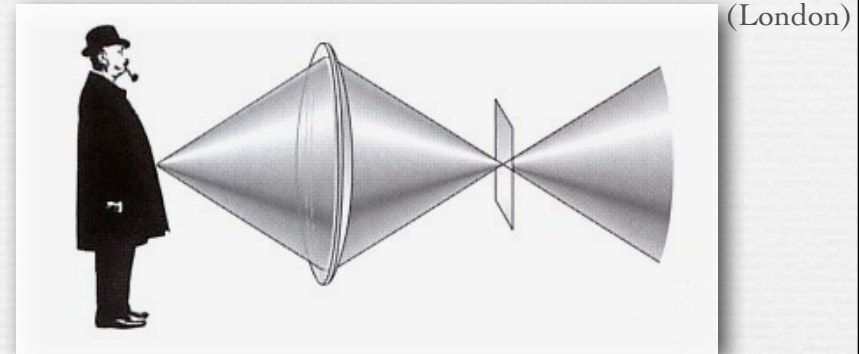


(London)


Depth of field (briefly)

This figure isn't quite right;
we'll fix it next week


- ◆ a point in the scene is focused at a point on the sensor
- ◆ if we move the sensor in depth, the point becomes blurred
- ◆ if it blurs too much, it exceeds our allowable *circle of confusion*
- ◆ the zone in which it's sharp enough is called the *depth of focus*
- ◆ this corresponds in the scene to a *depth of field*
- ◆ halving the aperture diameter doubles the depth of field



Trading off motion blur and depth of field



(London)




Questions?


Sensitivity (ISO)

- ◆ third variable for exposure
- ◆ film: trade sensitivity for grain
- ◆ digital: trade sensitivity for noise
 - multiply signal before analog-to-digital conversion
 - linear effect (200 ISO needs half the light as 100 ISO)

more in noise lecture



Tradeoffs affecting brightness



(FLASH DEMO)

<http://graphics.stanford.edu/courses/cs178/applets/exposure.html>

Slide credits

- ◆ Steve Marschner
- ◆ Fredo Durand
- ◆ Eddy Talvala

- ◆ Cole, A., *Perspective*, Dorling Kindersley, 1992.
- ◆ Kemp, M., *The Science of Art*, Yale University Press, 1990.
- ◆ Hecht, E., *Optics* (4th ed.), Pearson / Addison-Wesley, 2002.
- ◆ Renner, E., *Pinhole Photography* (2nd ed.), Focal Press, 2000.
- ◆ London, Stone, and Upton, *Photography* (9th ed.), Prentice Hall, 2008.
- ◆ D'Amelio, J., *Perspective Drawing Handbook*, Tudor Press, 1964.
- ◆ Dubery, F., Willats, J., *Perspective and other drawing systems*, Van Nostrand Reinhold, 1972.
- ◆ Kingslake, R. *Optics in Photography*, SPIE Press, 1992.
- ◆ <http://dpreview.com>