BBM 413 Fundamentals of Image Processing

Oct. 10, 2012

Erkut Erdem
Dept. of Computer Engineering
Hacettepe University

Image Formation and the Digital Camera

Acknowledgement: The course slides are mostly adapted from the slides prepared by Steve Marschner and Anat Levin

What is an image?

- A photographic print
- A photographic negative?
- This projection screen
- Some numbers in RAM?

An image is:

- A 2D distribution of intensity or color
- A function defined on a two-dimensional plane

$$I : \mathbb{R}^2 \rightarrow \ldots$$

- Note: no mention of pixels yet
- To process images, must:
  - obtain images—capture the scenes via hardware
  - represent images—encode them numerically

Image Formation

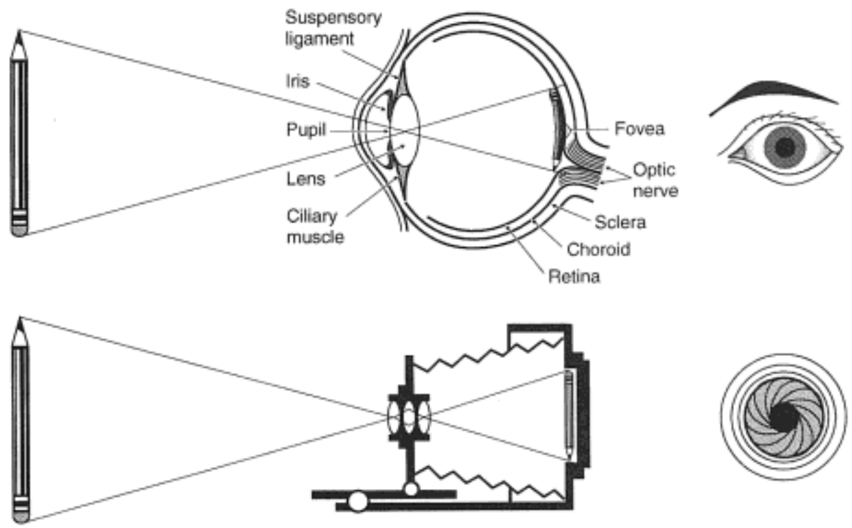

Three Dimensional World → Two Dimensional Image Space

- What is measured in an image location?
  - brightness
  - color

<< viewpoint
illumination conditions
local geometry
local material properties

Figures: Francis Crick, The Astonishing Hypothesis,

Image Formation

Why is there no image on a white piece of paper?
-It receives light from all directions

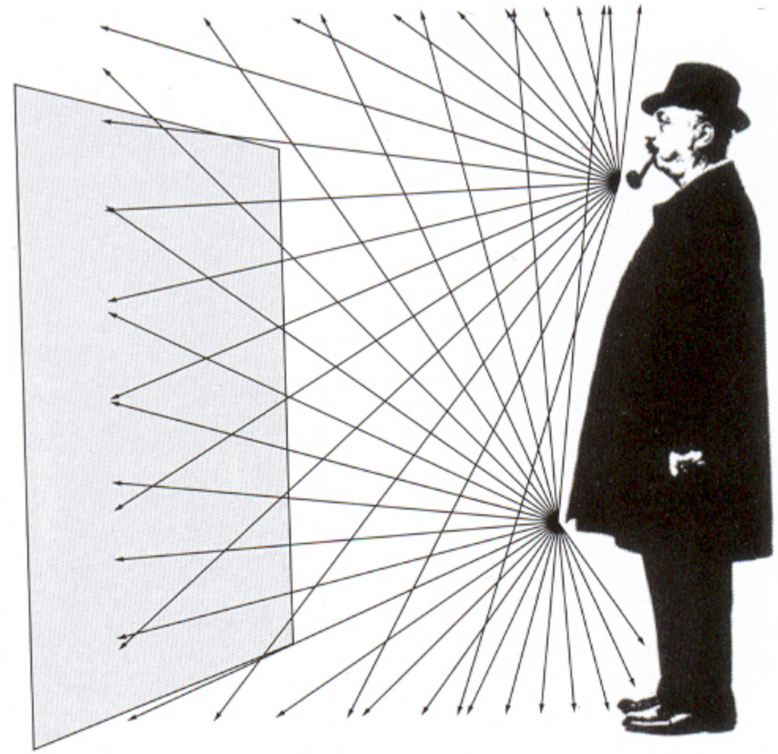

From Photography, London et al.

Pinhole

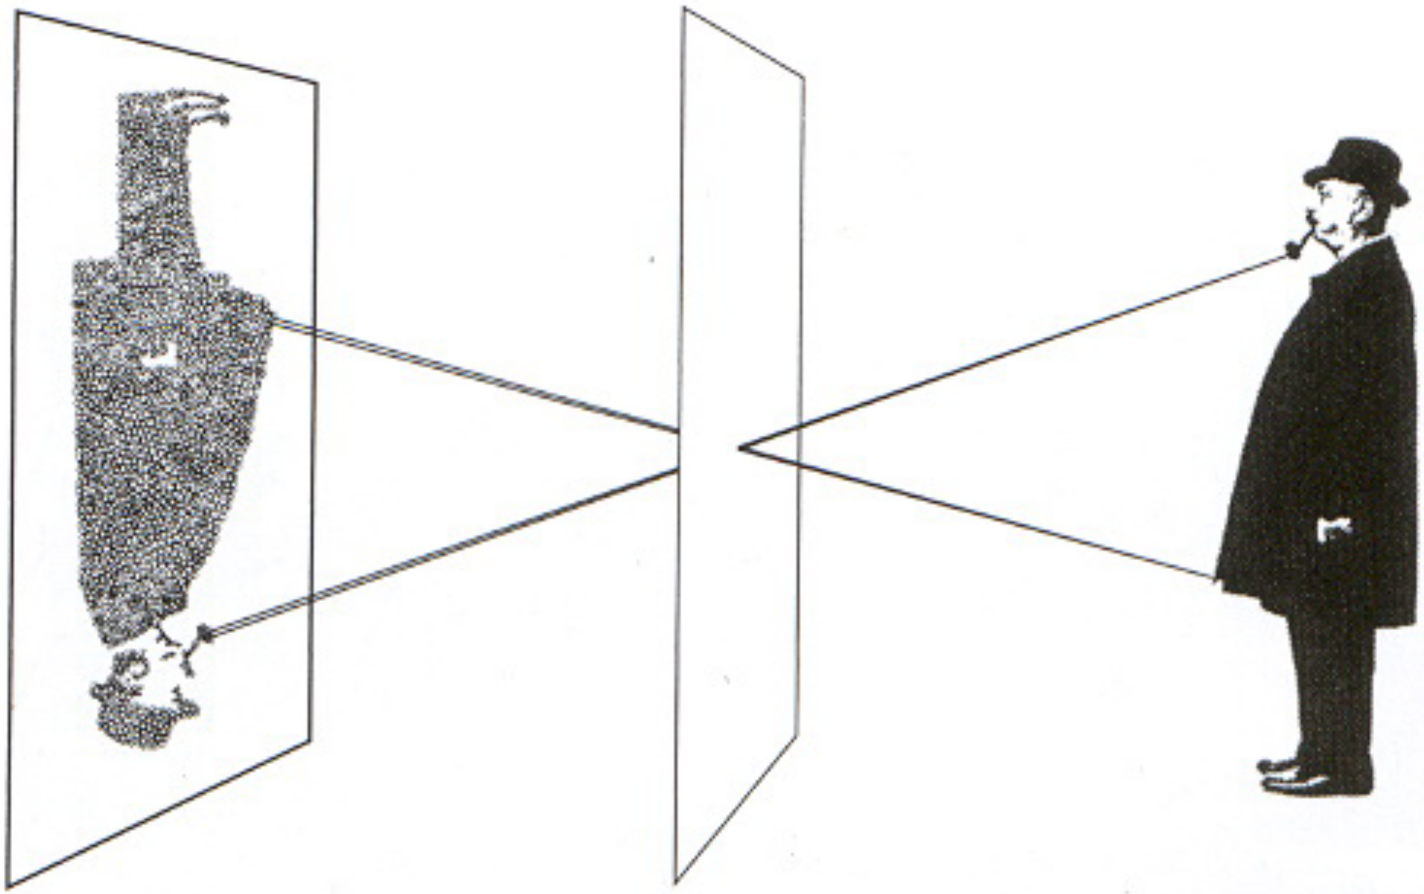

From Photography, London et al.

Pinhole

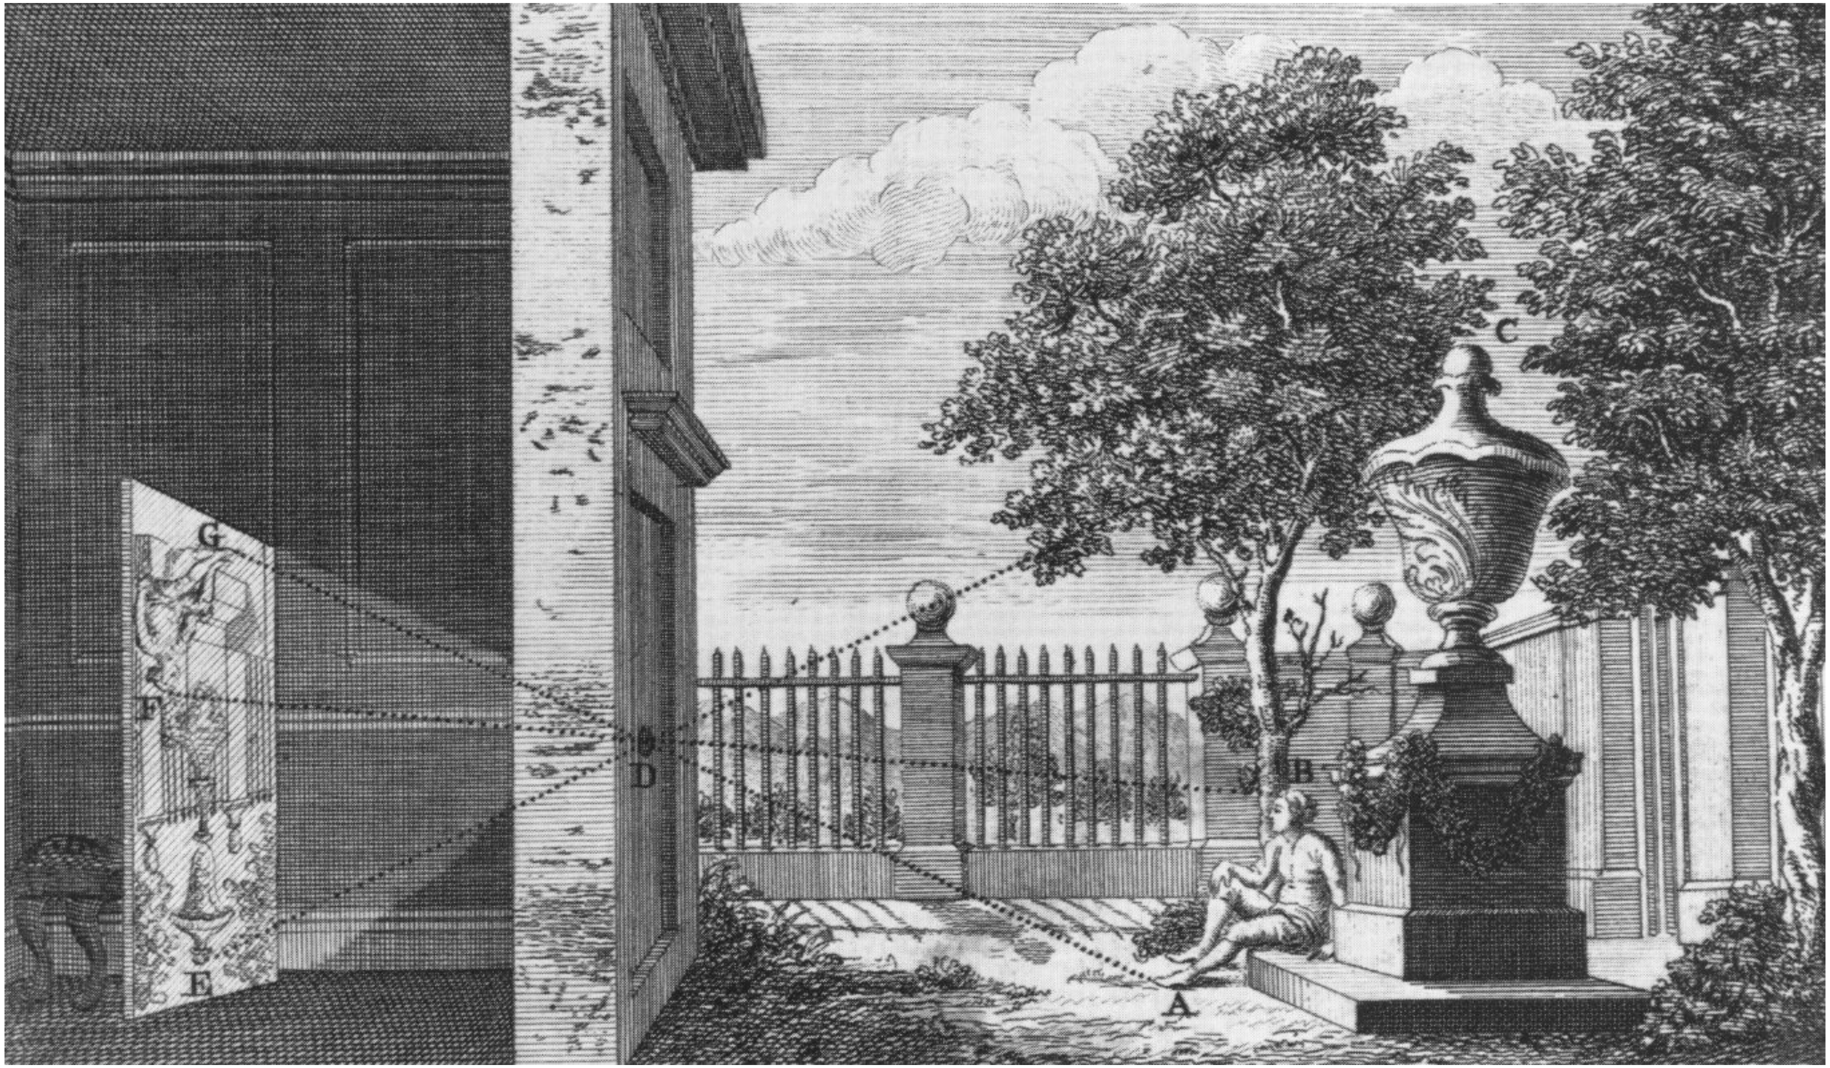

Pinhole Size?

Photograph made with small pinhole

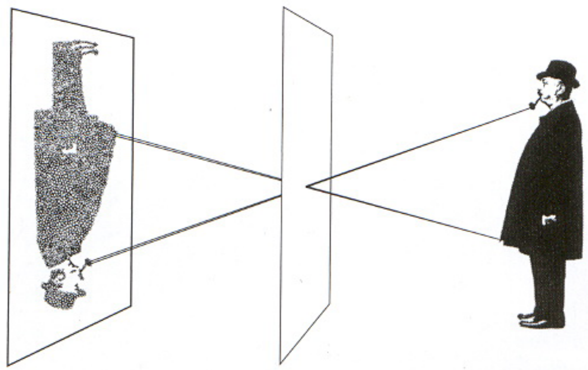

Small pinhole-
sharp but hard to
collect enough light

Photograph made with larger pinhole

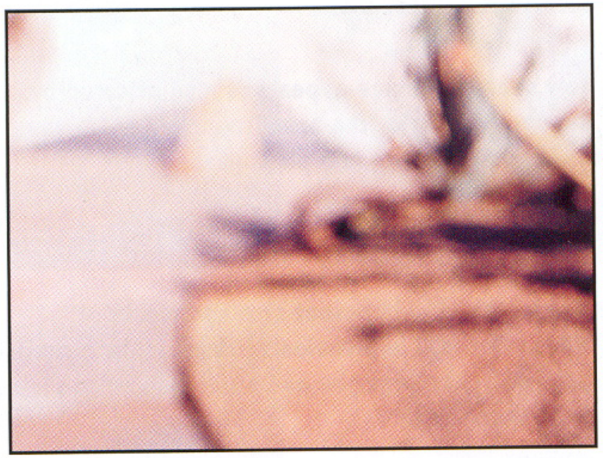

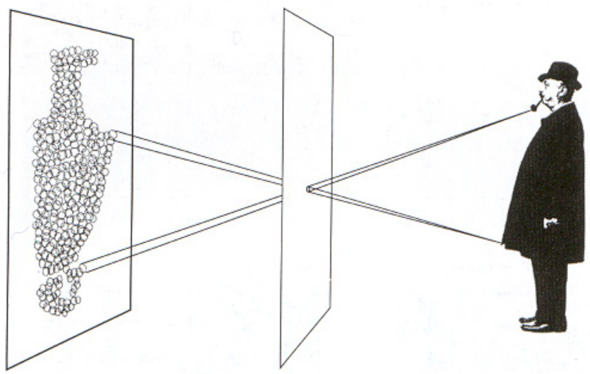

Larger pinhole-
Blur

From Photography, London et al.

Pinhole Size

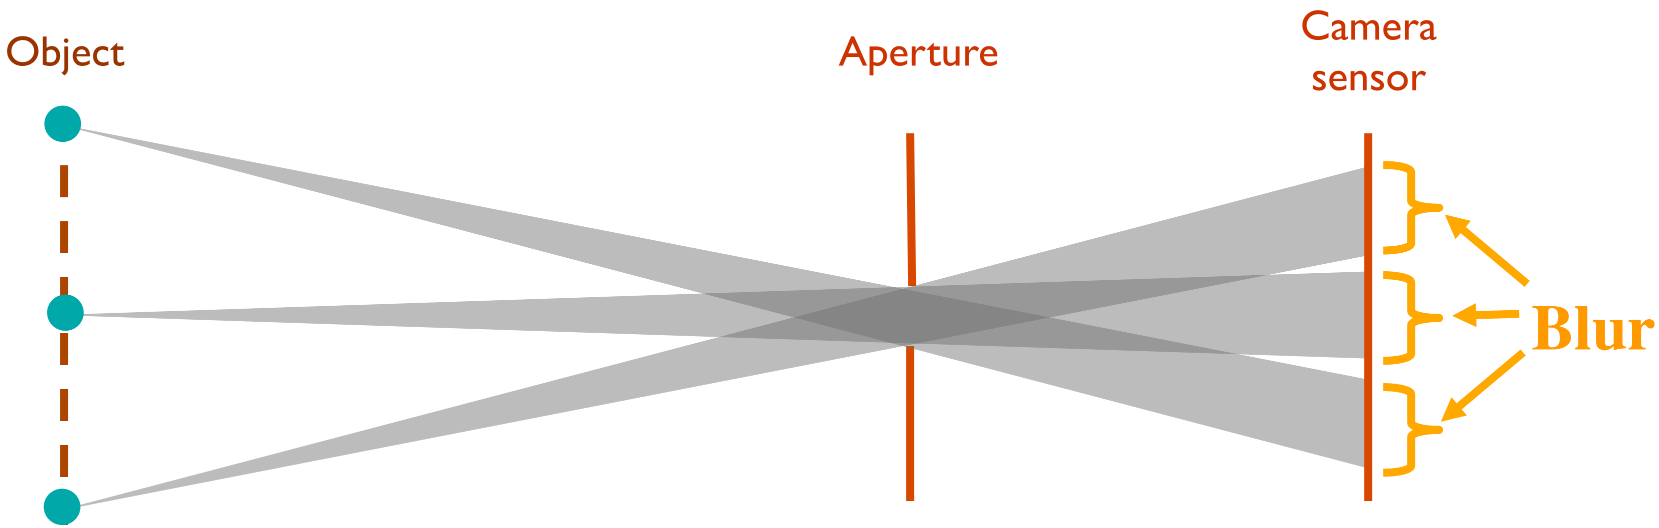

small hole => sharp, but don't collect enough light (noise)

Solution: refraction!

From Photography, London et al.

Lenses

- gather more light!
- But need to be focused

Photograph made with small pinhole

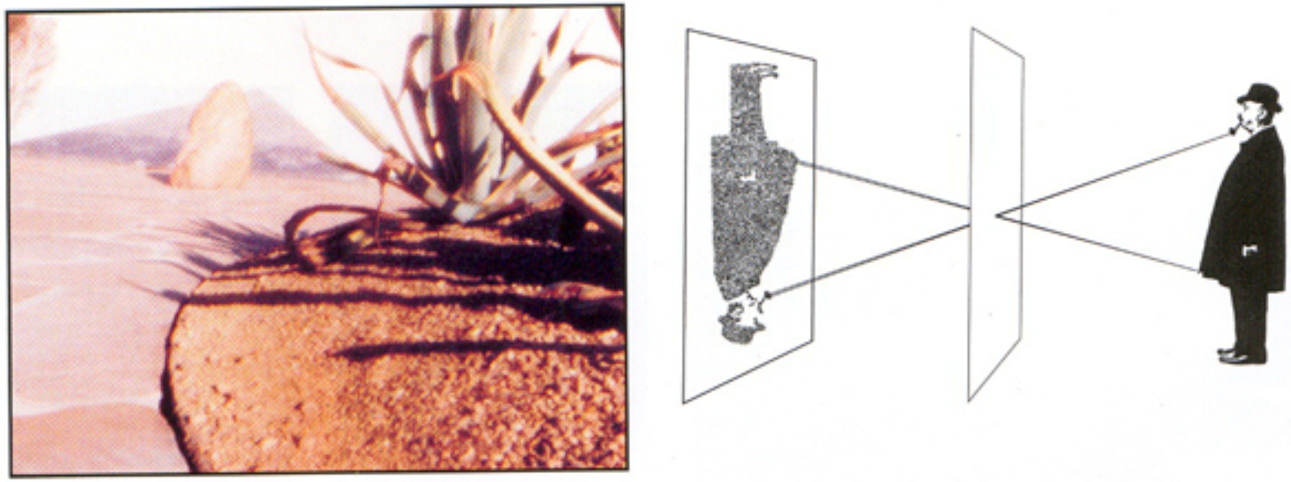

To make this picture, the lens of a camera was replaced with a thin metal disk pierced by a tiny pinhole, *equivalent in size to an aperture of f/182. Only a few rays of light from each point on the subject got through the tiny opening, producing a soft but acceptably clear photograph. Because of the small size of the pinhole, the exposure had to be 6 sec long.*

Photograph made with lens

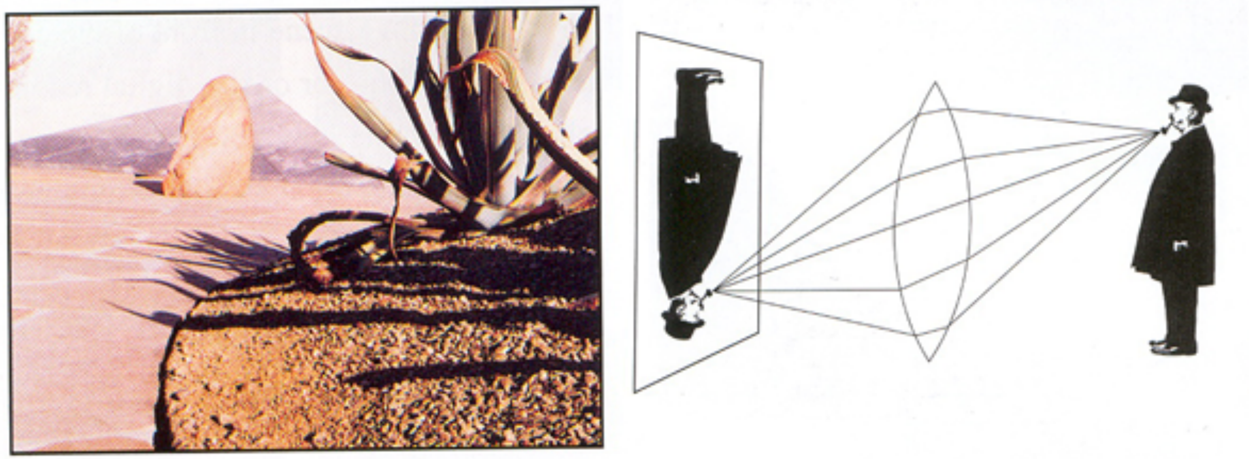

This time, using a simple convex lens with an f/16 aperture, *the scene appeared sharper than the one taken with the smaller pinhole, and the exposure time was much shorter, only 1/100 sec.*

The lens opening was much bigger than the pinhole, letting in far more light, but it focused the rays from each point on the subject precisely so that they were sharp on the film.

From Photography, London et al.

A lens is focused at a single depth

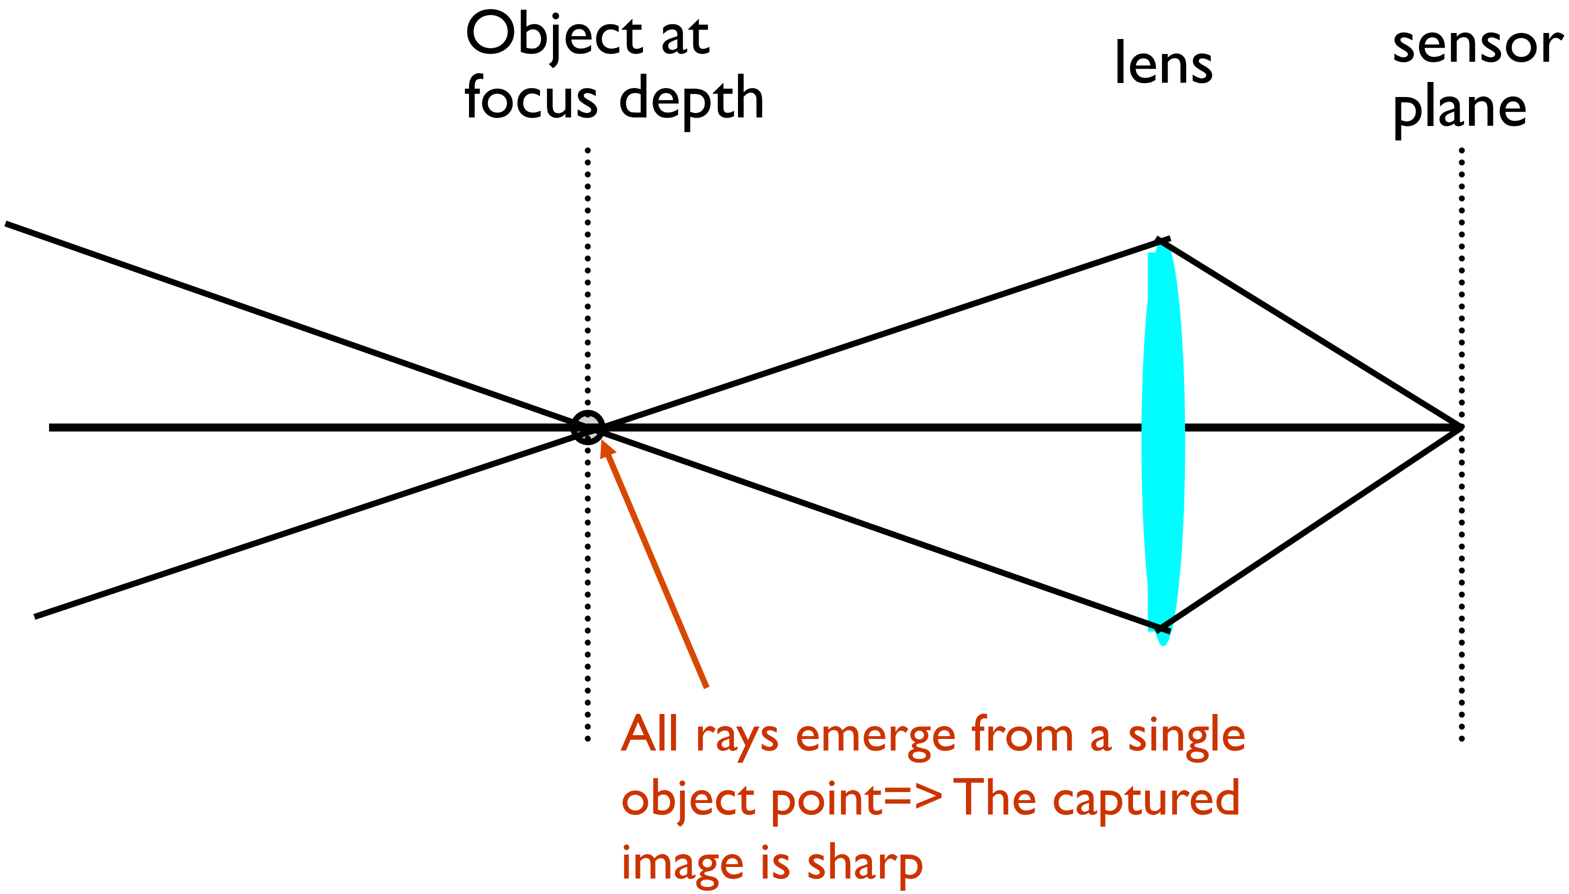

A lens is focused at a single depth

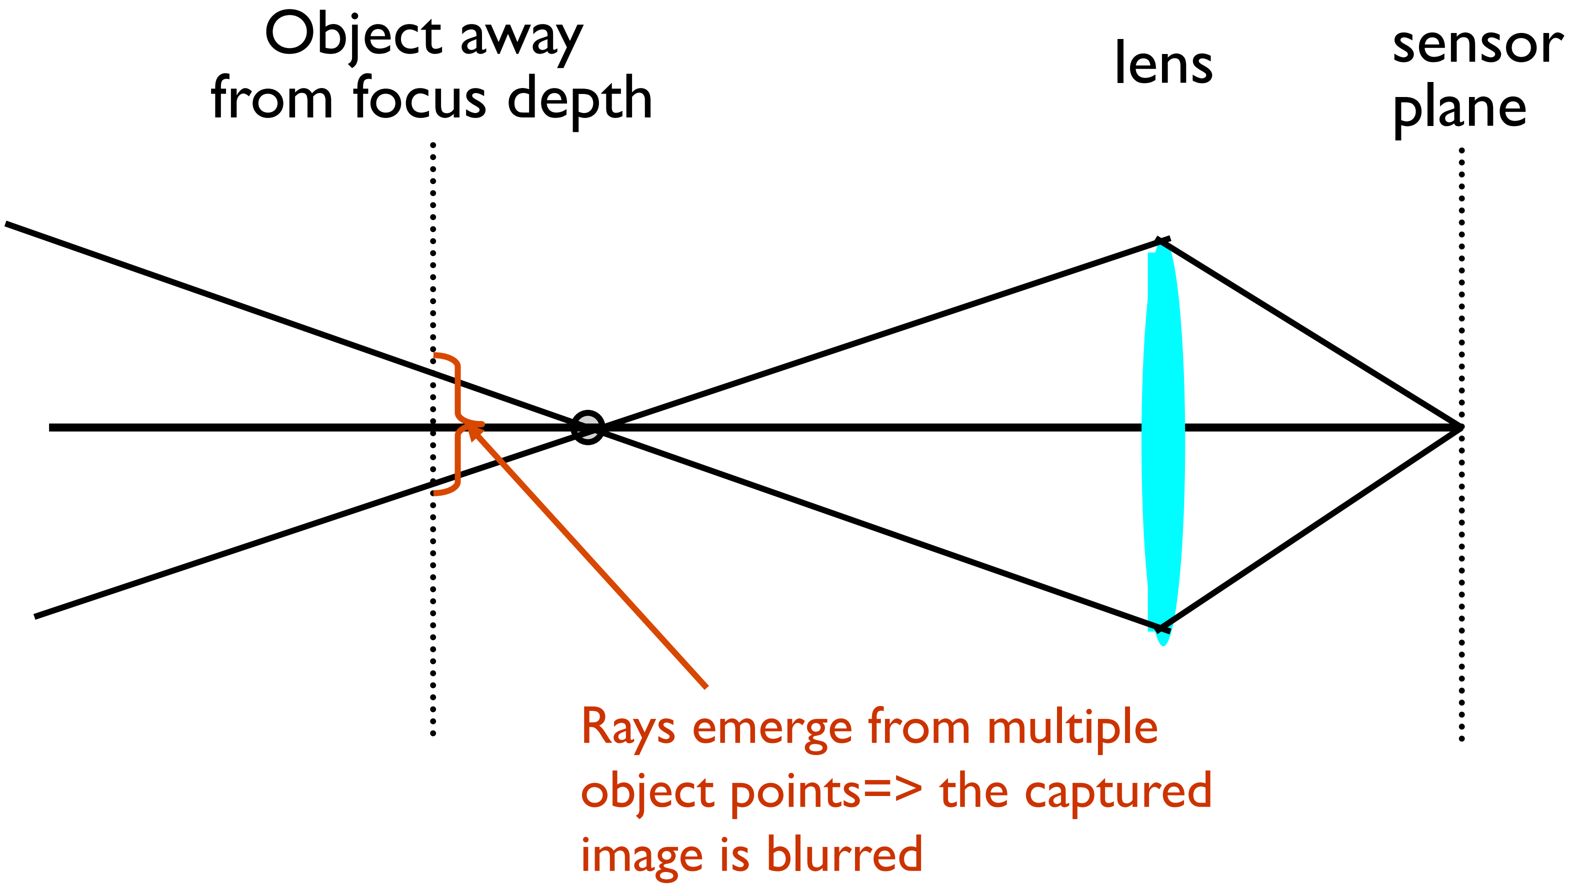

A lens is focused at a single depth

Aperture

- Diameter of the lens opening (controlled by diaphragm)
- Expressed as a fraction of focal length, in f-number
  - f/2.0 on a 50mm means that the aperture is 25mm
  - f/2.0 on a 100mm means that the aperture is 50mm
- Disconcerting: small f number = big aperture
- What happens to the area of the aperture when going from f/2.0 to f/4.0?
- Typical f numbers are
  f/2.0, f/2.8, f/4, f/5.6, f/8, f/11, f/16, f/22, f/32
  - See the pattern?

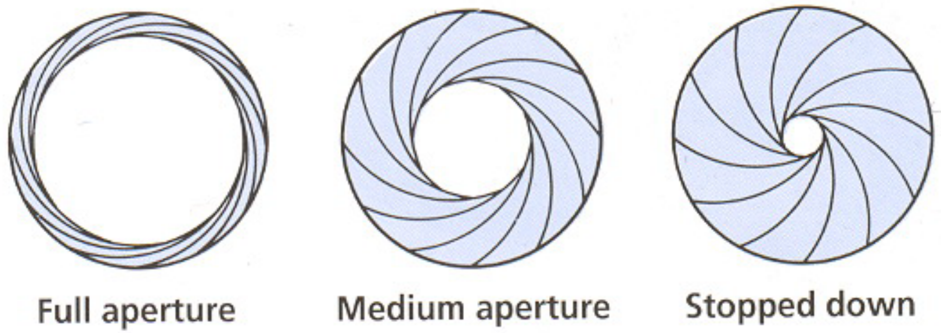

Main effect of aperture

- Depth of field

Large aperture opening

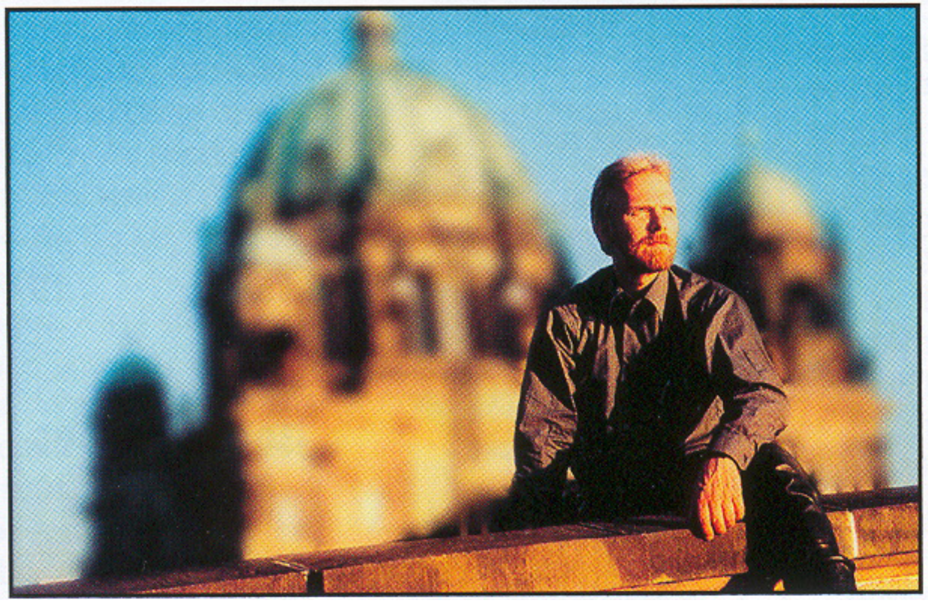

Small aperture opening

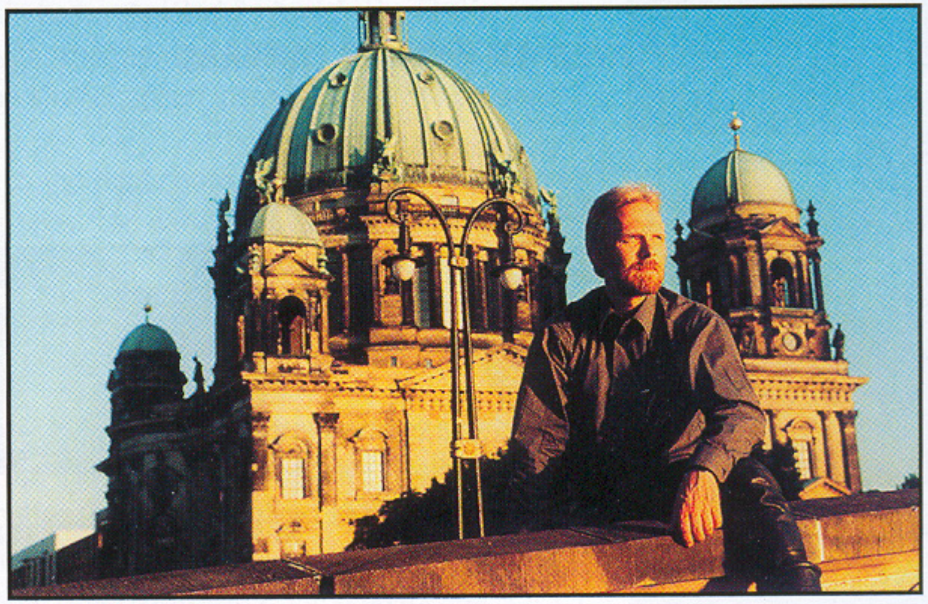

From Photography, London et al.

Depth of field

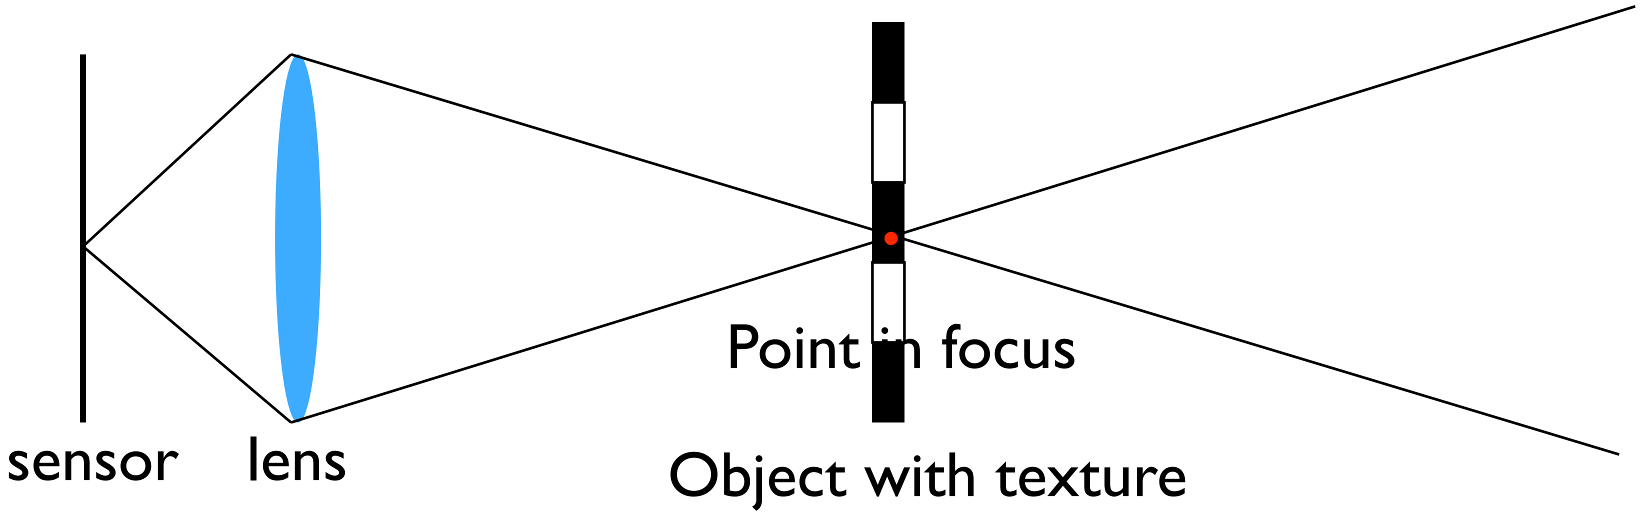

Image of object in focus- sharp (all rays hitting a single sensor point emerge from a single point on the object)

Depth of field

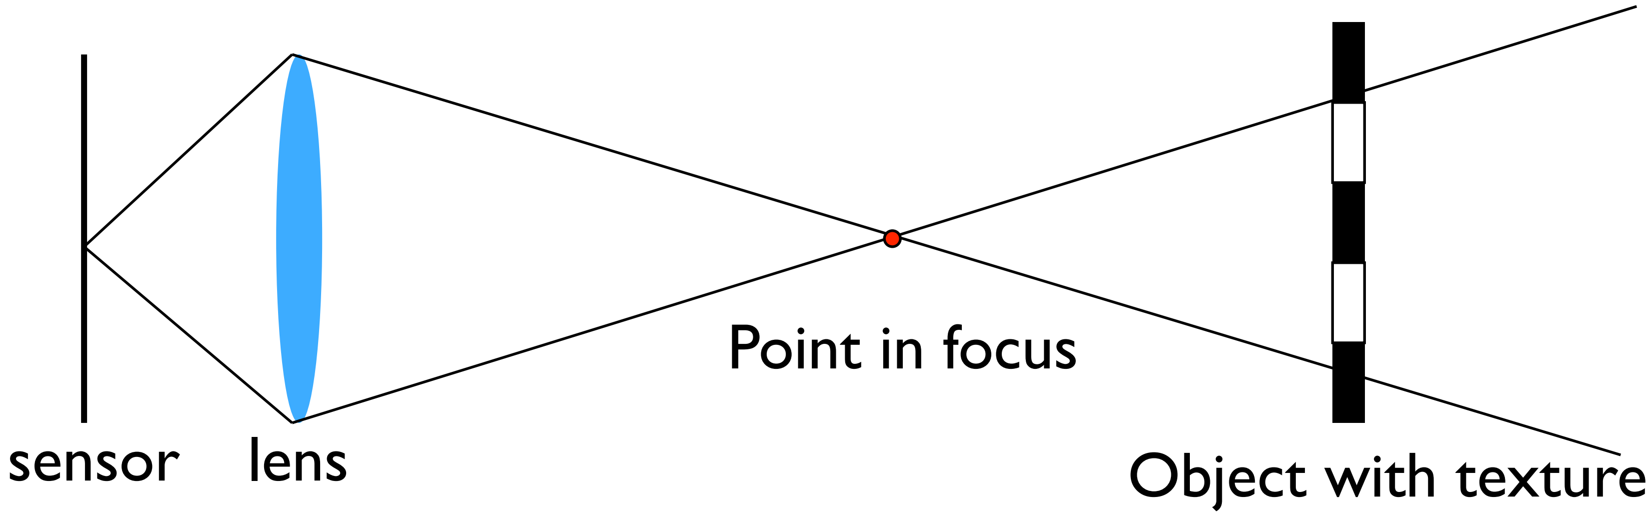

Image of object in focus- sharp (all rays hitting a single sensor point emerge from a single point on the object)

Image of an object away from focus depth- blurred (rays hitting a single sensor point emerge from multiple points on the object)

Depth of field

- We allow for some tolerance

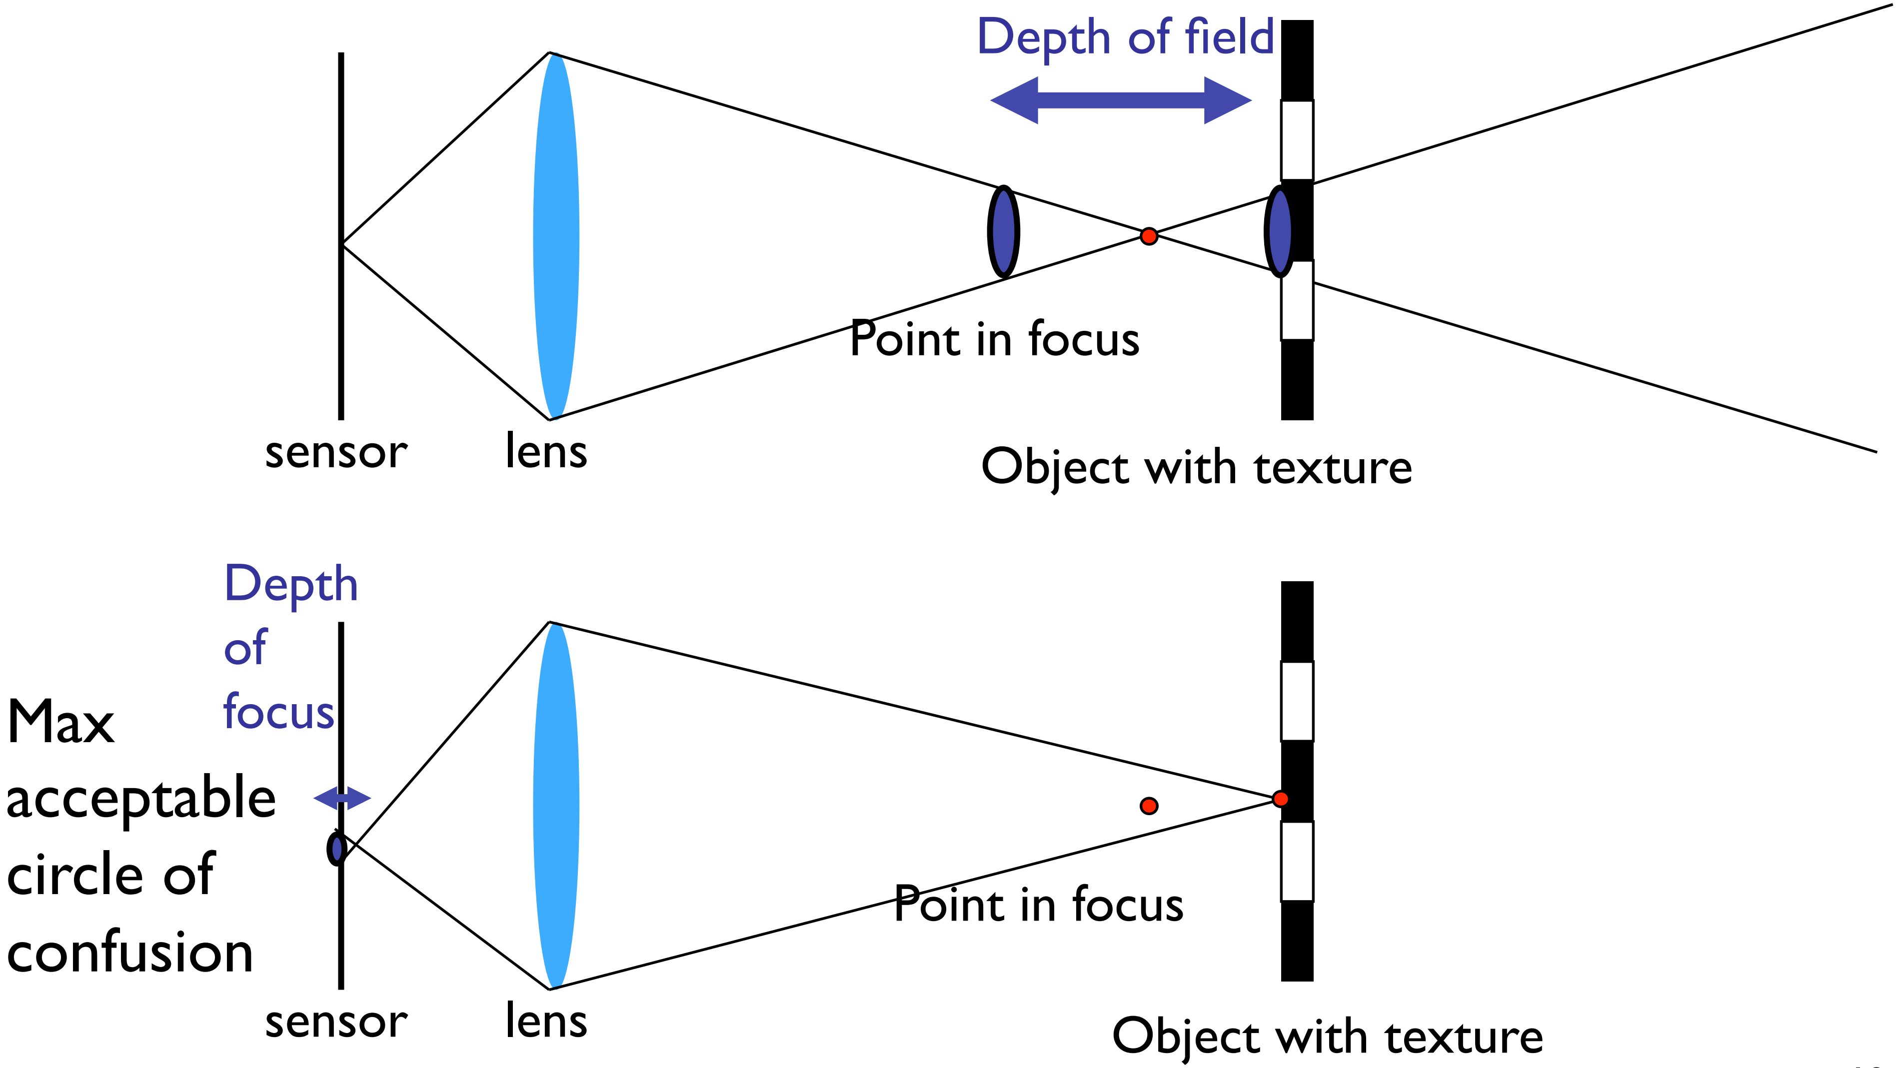

Depth of Field

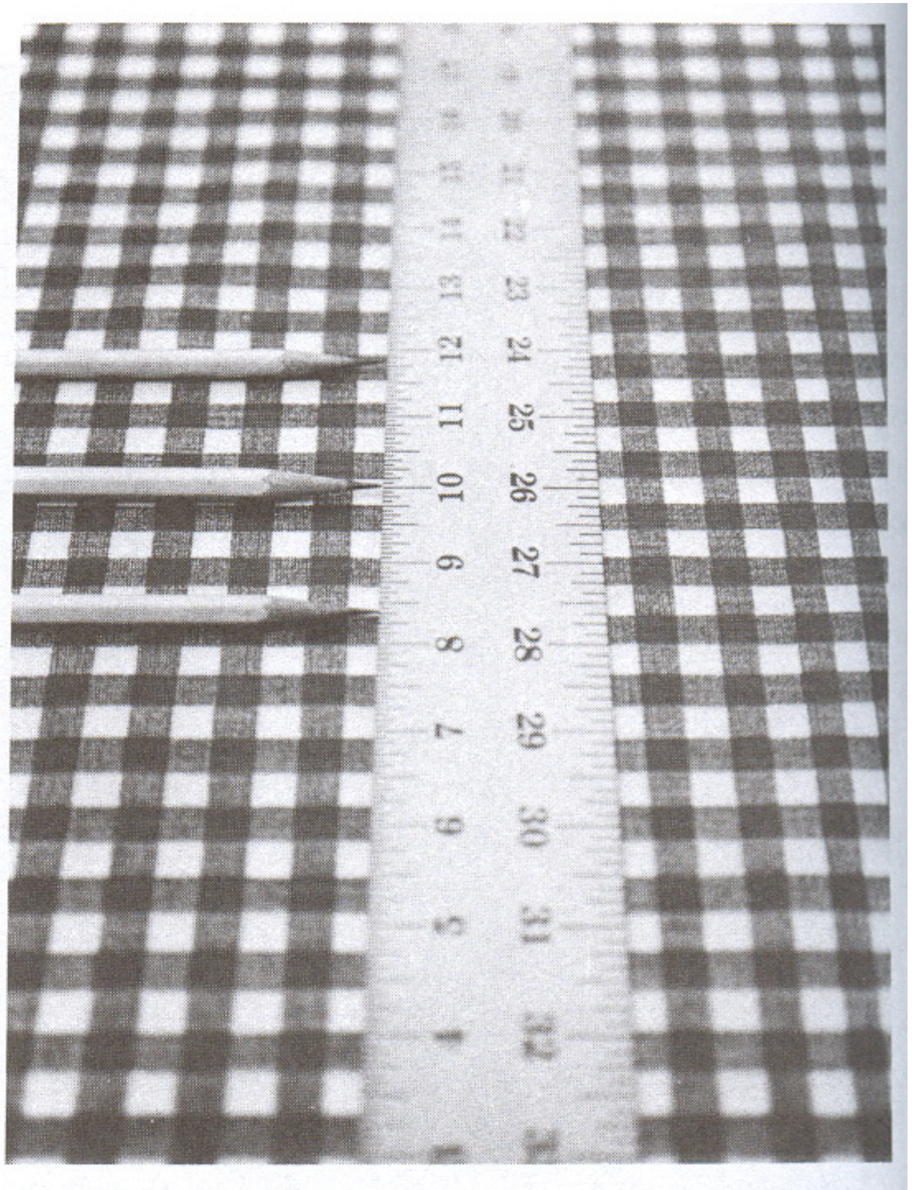

Depth of Field

	Portrait	Landscape	
Shallow Depth of Field			Large Aperture 
Large Depth of Field			Small Aperture 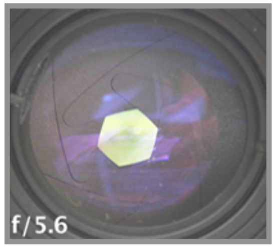

http://photographertips.net

Exposure

- Get the right amount of light to sensor/film
- Two main parameters:
  - Shutter speed
  - Aperture (area of lens)

Shutter speed

- Controls how long the film/sensor is exposed
- Pretty much linear effect on exposure
- Usually in fraction of a second:
  - 1/30, 1/60, 1/125, 1/250, 1/500
  - Get the pattern ?
- On a normal lens, normal humans can hand-hold down to 1/60
  - In general, the rule of thumb says that the limit is the inverse of focal length, e.g. 1/500 for a 500mm

Shutter speed

Short exposure- dark

medium exposure

long exposure- saturation

Shutter speed

Short exposure after contrast adjustment- noise

medium exposure

long exposure- saturation

Main effect of slower shutter speed

- Motion blur – camera shake

Image taken with a tripod

Image taken with a hand held camera

Main effect of slower shutter speed

- Motion blur- scene motion

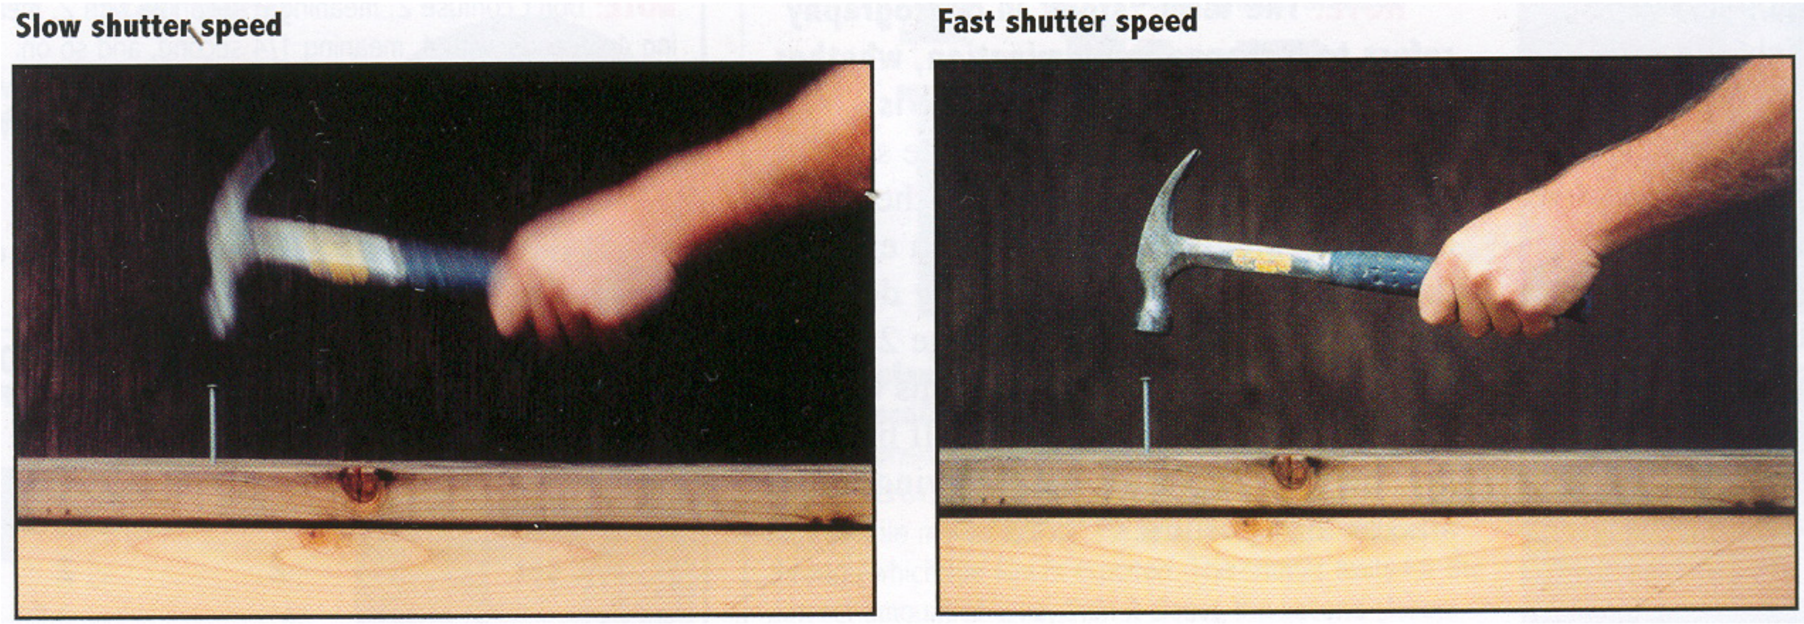

From Photography, London et al.

Representative display technologies

Computer displays

- Raster CRT display
- LCD display

Printers

- Laser printer
- Inkjet printer

Effect of shutter speed

- Freezing motion

Walking people

1/125

Running people

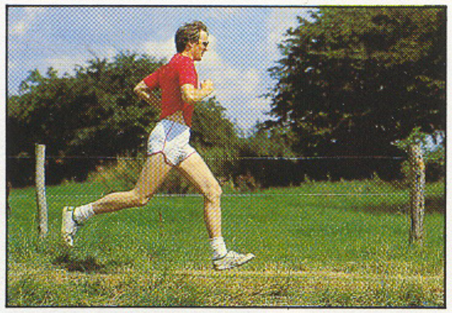

1/250

Car

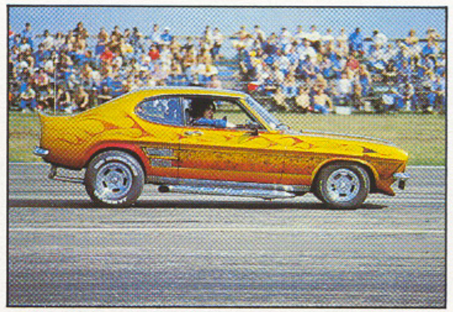

1/500

Fast train

1/1000

Cathode ray tube

- First widely used electronic display
  - developed for TV in the 1920s–1930s

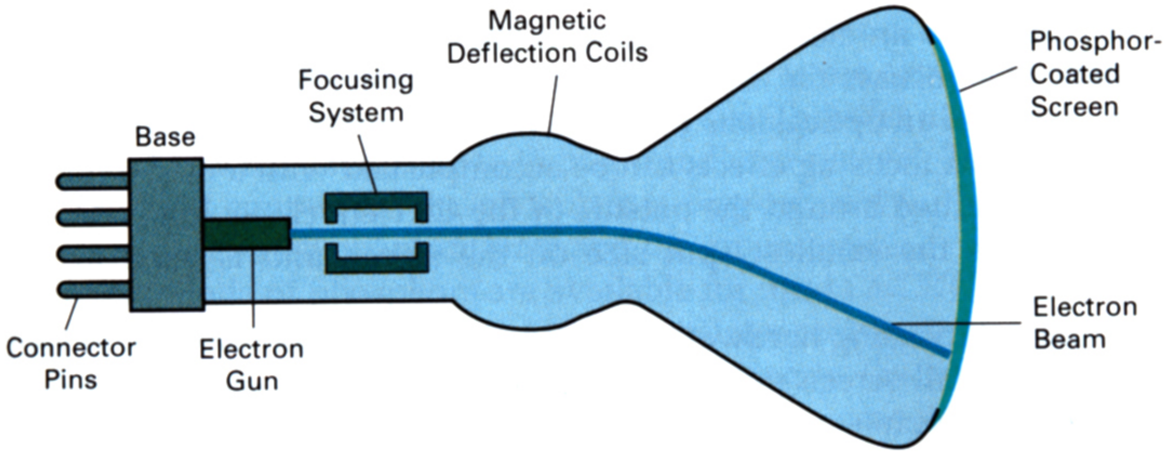

Raster CRT display

- Scan pattern fixed in display hardware
- Intensity modulated to produce image
- Originally for TV
  - (continuous analog signal)
- For computer, intensity determined by contents of *framebuffer*

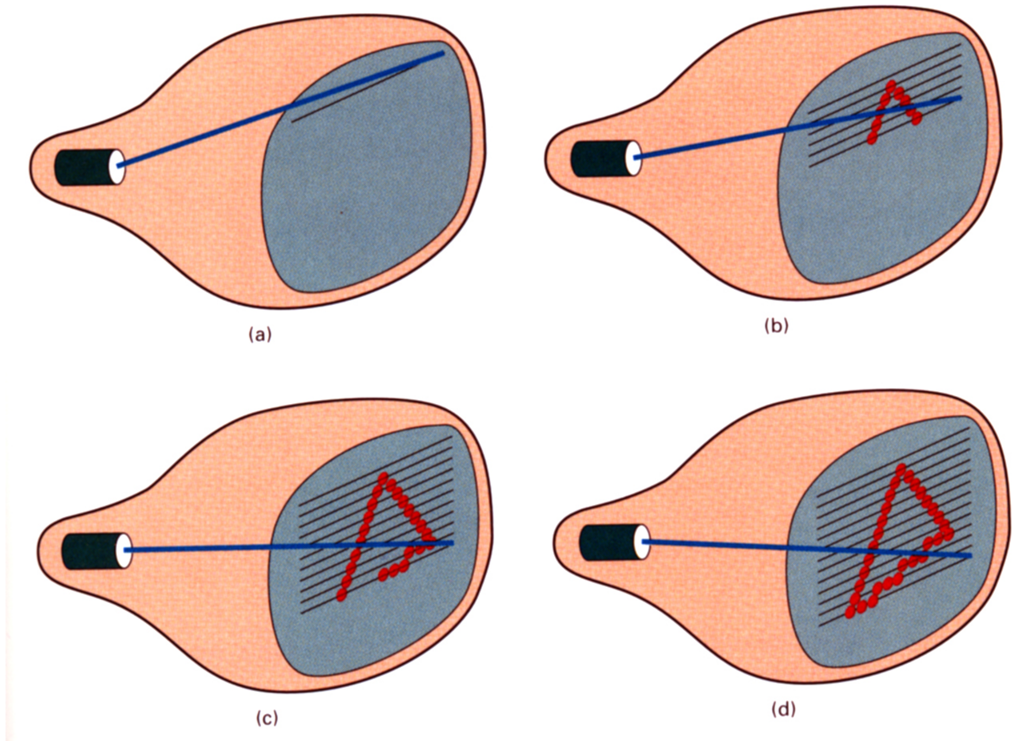

LCD flat panel or projection display

- Principle: block or transmit light by twisting its polarization
- Intermediate intensity levels possible by partial twist
- Fundamentally raster technology
- Fixed format

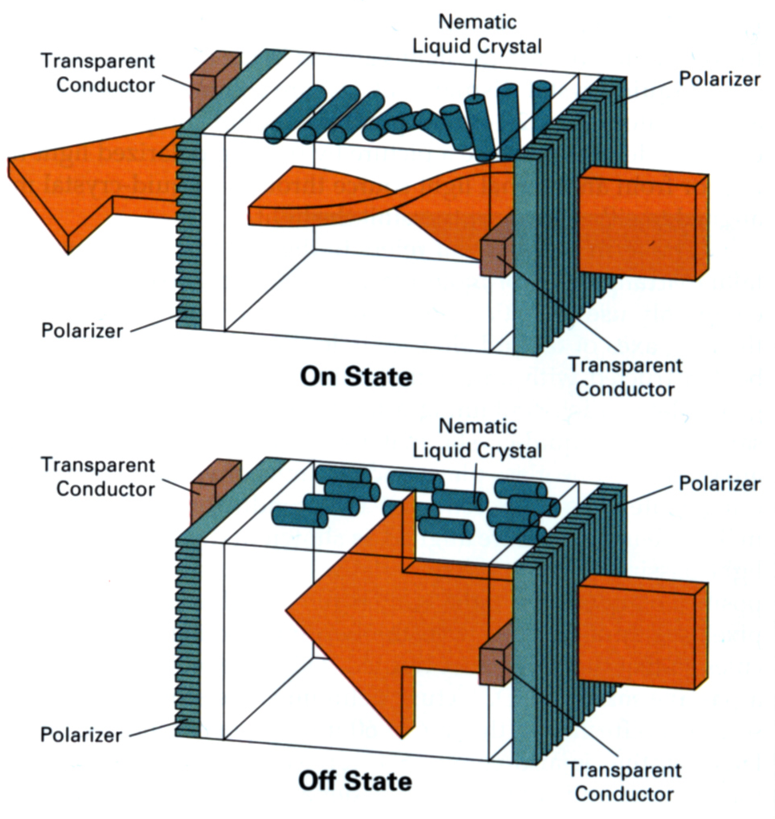

Raster display system

- Screen image defined by a 2D array in RAM
  - for CRT, read out and convert to analog in sync with scan
- In most systems today, it's in a separate memory
- The memory area that maps to the screen is called the *frame buffer*

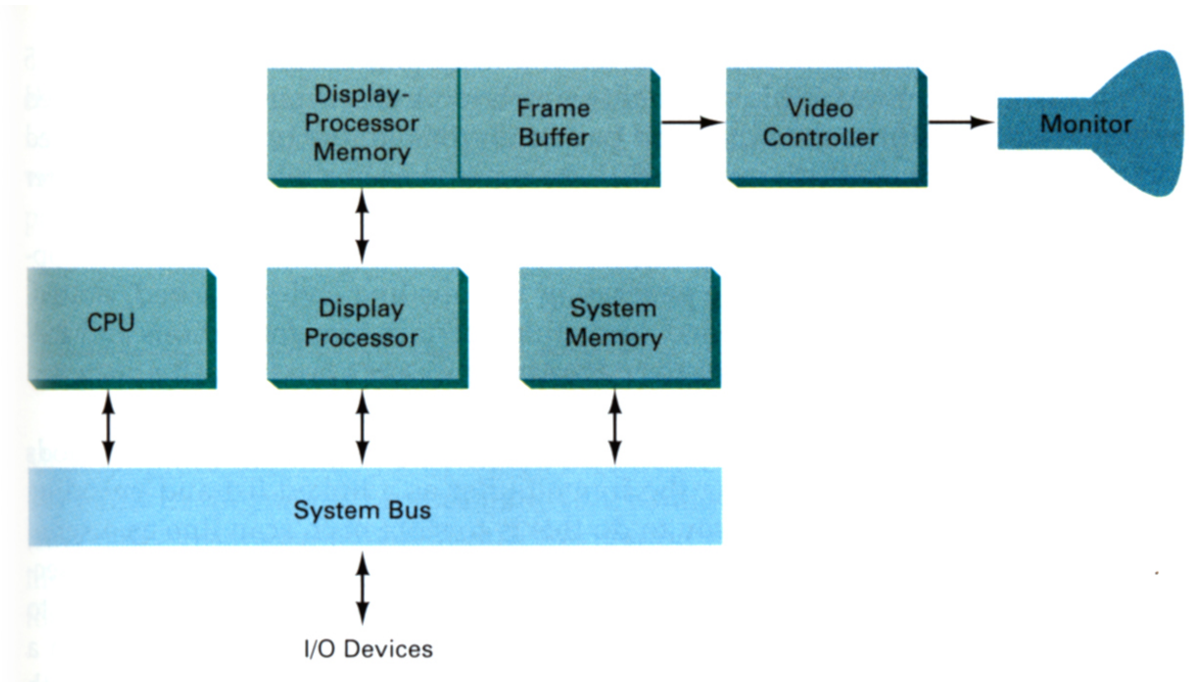

Color displays

- Operating principle: humans are trichromatic
  - match any color with blend of three
  - therefore, problem reduces to producing 3 images and blending
- Additive color
  - blend images by sum
  - e.g. overlapping projection
  - e.g. unresolved dots
  - R, G, B make good primaries

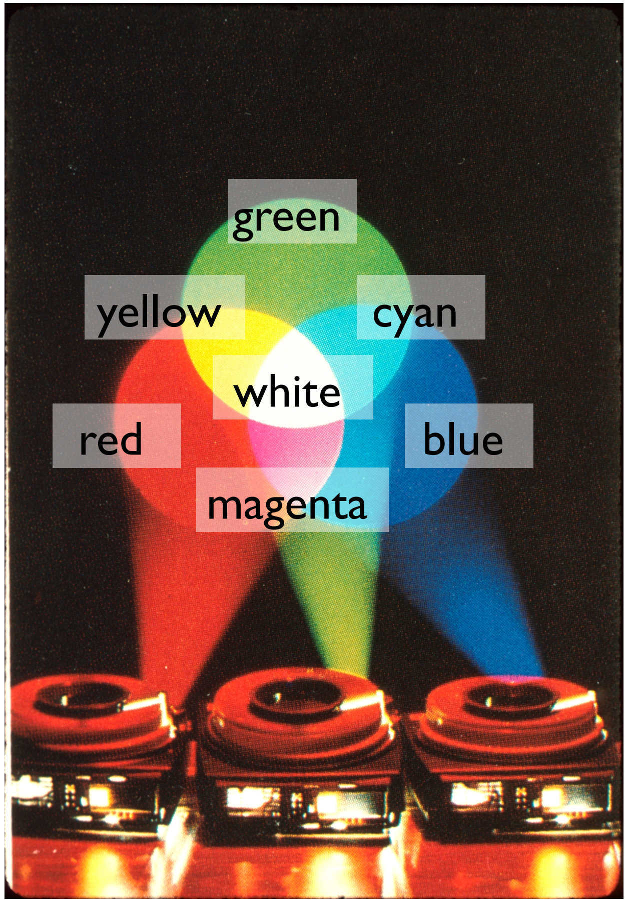

[source unknown]

Color displays

- CRT: phosphor dot pattern to produce finely interleaved color images

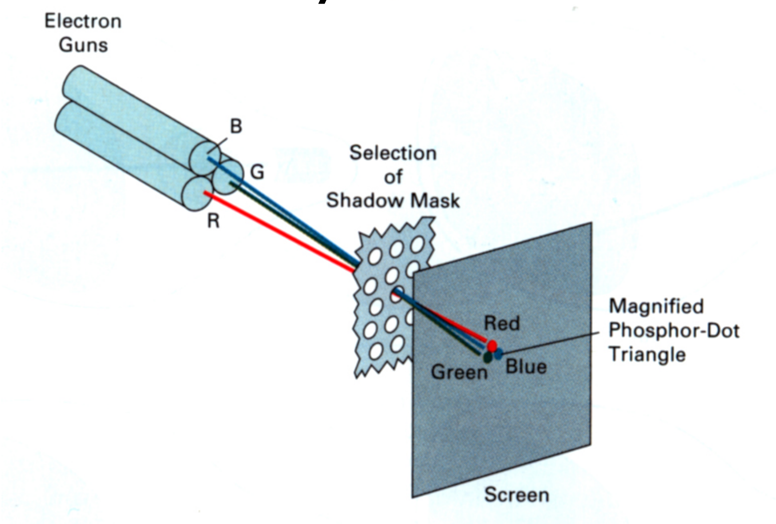

- LCD: interleaved R,G,B pixels

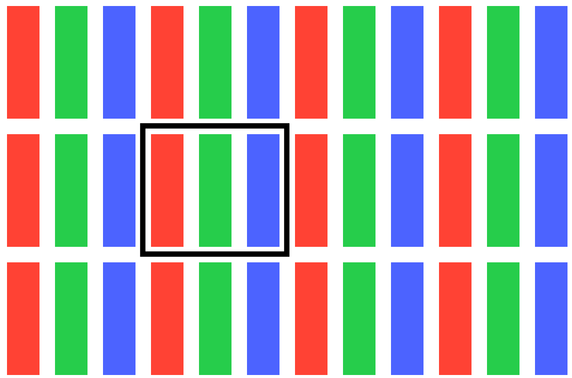

Laser printer

- Xerographic process
- Like a photocopier but with laser-scanned raster as source image
- Key characteristics
  - image is binary
  - resolution is high
  - very small, isolated dots are not possible

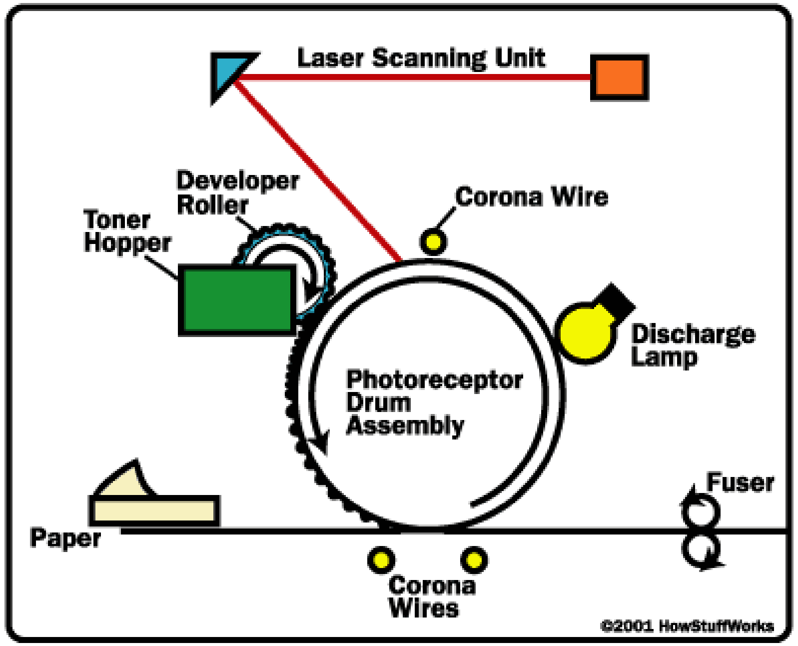

[howstuffworks.com]

Inkjet printer

- Liquid ink sprayed in small drops
  - very small—measured in picoliters
- Head with many jets scans across paper
- Key characteristics:
  - image is binary (drop or no drop; no partial drops)
  - isolated dots are reproduced well

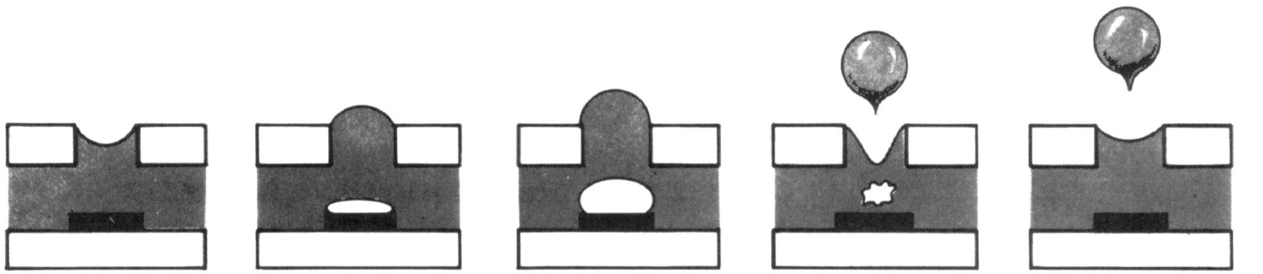

[source unknown]

Digital camera

- A raster input device
- Image sensor contains 2D array of photosensors

[CS 417 Spring 2002]

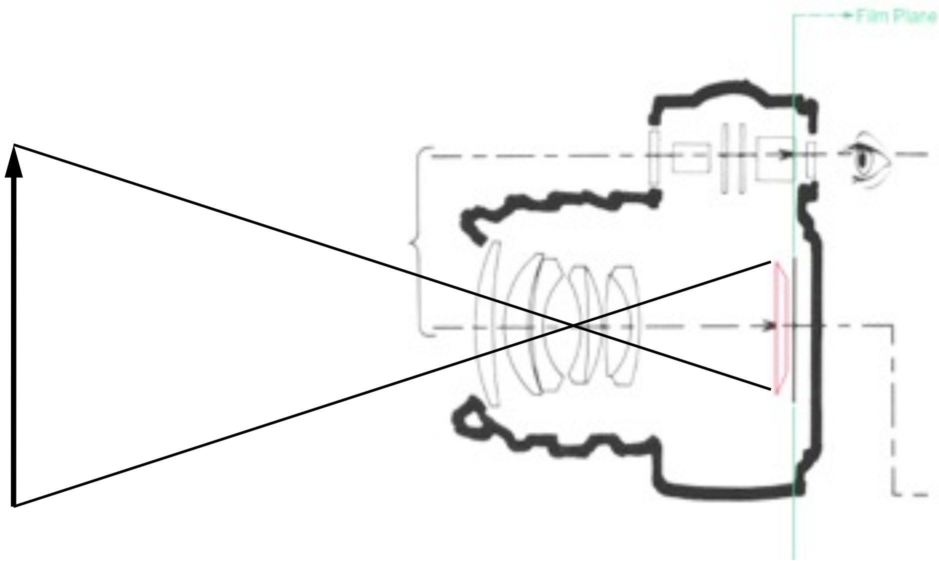

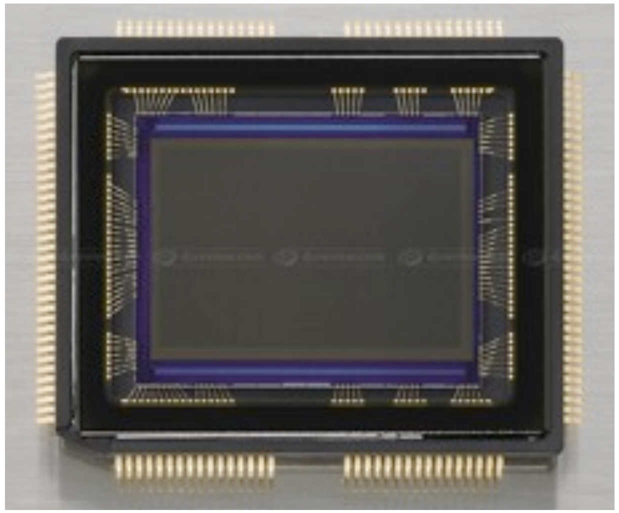

[dpreview.com]

Digital camera

- Color typically captured using color mosaic

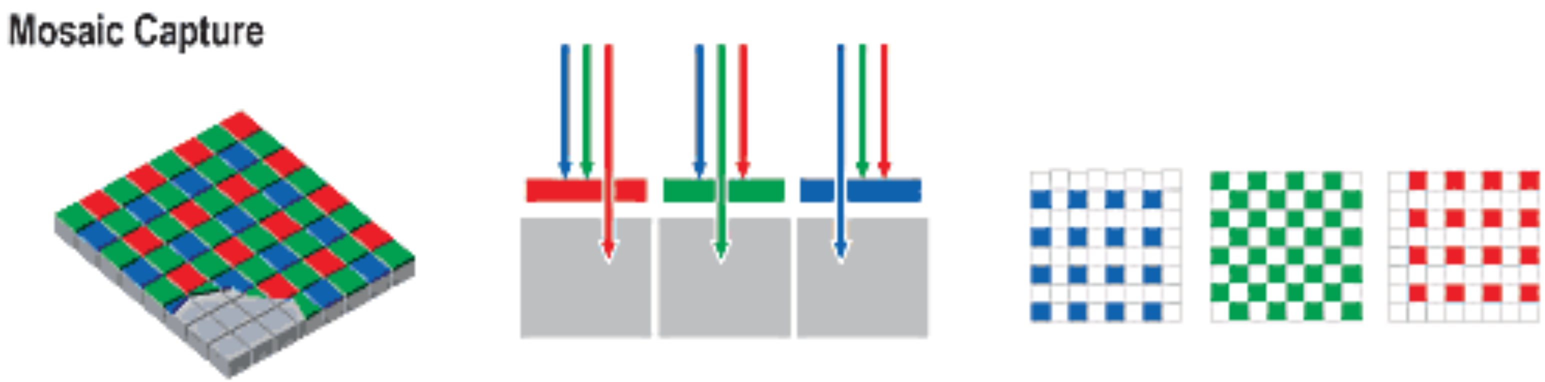

[Foveon]

Raster image representation

- All these devices suggest 2D arrays of numbers
- Big advantage: represent arbitrary images
  - approximate arbitrary functions with increasing resolution
  - works because memory is cheap (brute force approach!)

[Philip Greenspun]

Meaning of a raster image

- Meaning of a given array is a function on 2D
- Define meaning of array = result of output device?
  - that is, piecewise constant for LCD, blurry for CRT
  - but: we don't have just one output device
  - but: want to define images we can't display (e.g. too big)
- Abstracting from device, problem is reconstruction
  - image is a sampled representation
  - pixel means "this is the intensity around here"
    - LCD: intensity is constant over square regions
    - CRT: intensity varies smoothly across pixel grid
  - will discuss specifics of reconstruction later

Image Representation

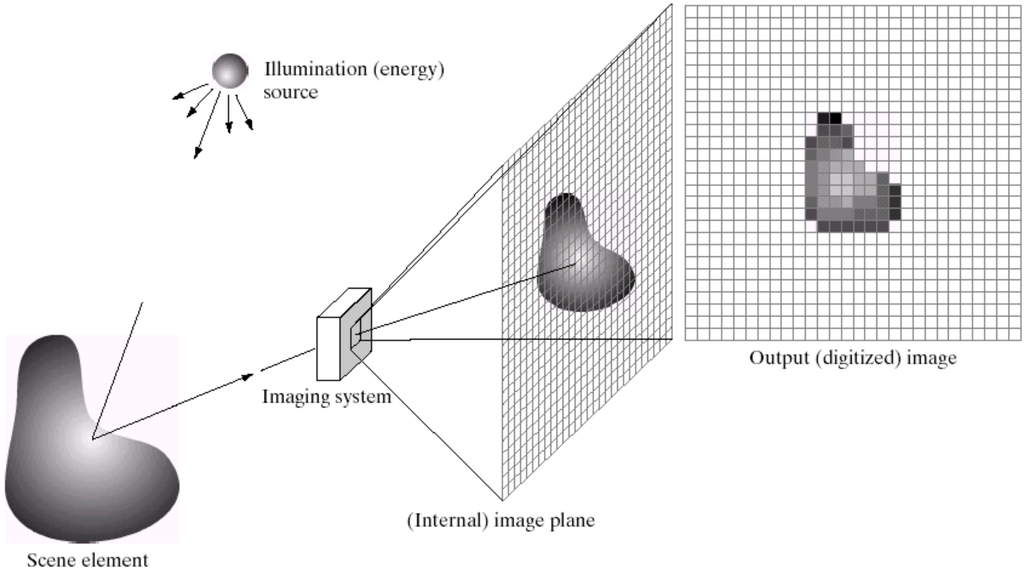

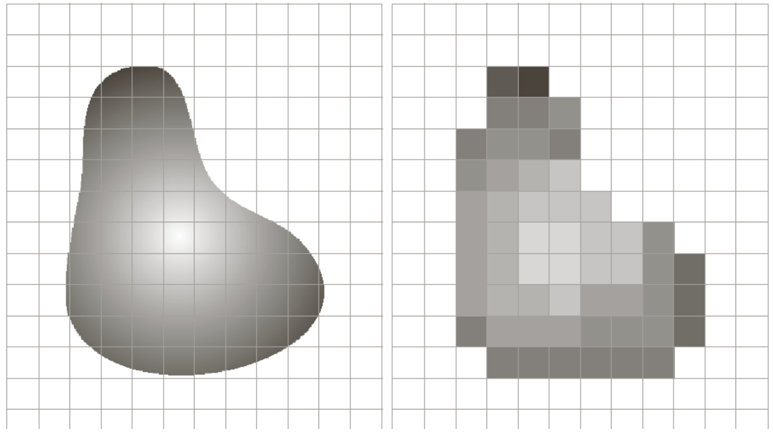

- Discretization
  - in image space - sampling
  - In image brightness - quantization

Figures: Gonzalez and Woods, Digital Image Processing, 3rd Edition, 2008

Image Representation

- **Digital image:** 2D discrete function f
- **Pixel**: Smallest element of an image f(x,y)

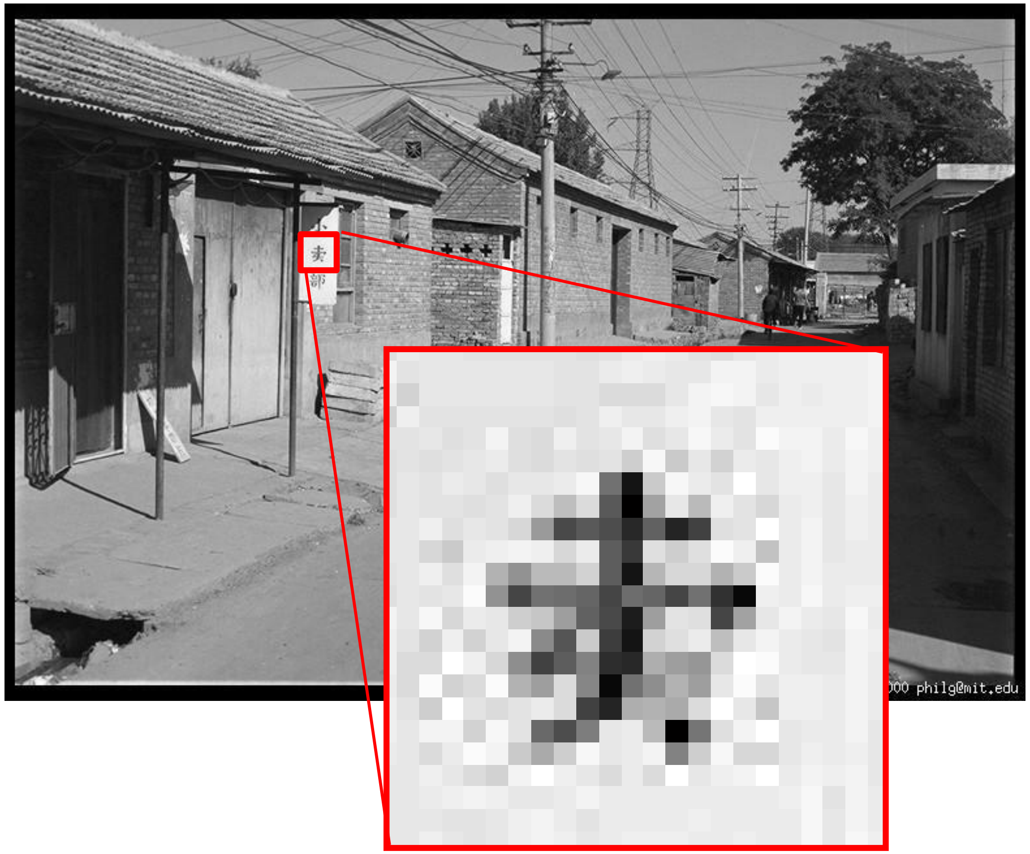

Figure: M. J. Black

Image Representation

- **Digital image:** 2D discrete function f
- **Pixel**: Smallest element of an image f(x,y)

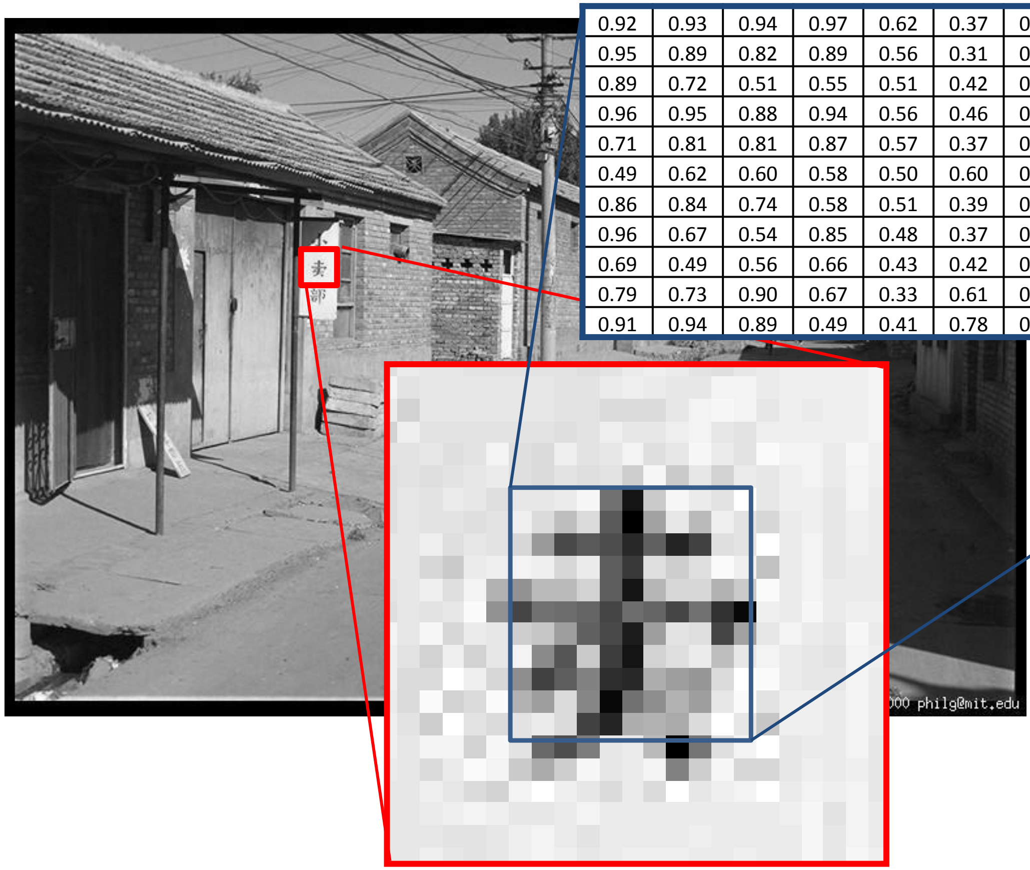

0.92	0.93	0.94	0.97	0.62	0.37	0.85	0.97	0.93	0.92	0.99
0.95	0.89	0.82	0.89	0.56	0.31	0.75	0.92	0.81	0.95	0.91
0.89	0.72	0.51	0.55	0.51	0.42	0.57	0.41	0.49	0.91	0.92
0.96	0.95	0.88	0.94	0.56	0.46	0.91	0.87	0.90	0.97	0.95
0.71	0.81	0.81	0.87	0.57	0.37	0.80	0.88	0.89	0.79	0.85
0.49	0.62	0.60	0.58	0.50	0.60	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

Figure: M. J. Black

Datatypes for raster images

- Bitmaps: `boolean` per pixel (1 bpp): $I : \mathbb{R}^2 \rightarrow \{0, 1\}$
  - interp. = black and white; e.g. fax
- Grayscale: integer per pixel: $I : \mathbb{R}^2 \rightarrow [0, 1]$
  - interp. = shades of gray; e.g. black-and-white print
  - precision: usually `byte` (8 bpp); sometimes 10, 12, or 16 bpp
- Color: 3 integers per pixel: $I : \mathbb{R}^2 \rightarrow [0, 1]^3$
  - interp. = full range of displayable color; e.g. color print
  - precision: usually `byte[3]` (24 bpp)
  - sometimes 16 (5+6+5) or 30 or 36 or 48 bpp
  - indexed color: a fading idea

Datatypes for raster images

- Floating point: $I : \mathbb{R}^2 \rightarrow \mathbb{R}_+$ or $I : \mathbb{R}^2 \rightarrow \mathbb{R}_+^3$
  - more abstract, because no output device has infinite range
  - provides *high dynamic range* (HDR)
  - represent real scenes independent of display
  - becoming the standard intermediate format in graphics processors
- Clipping and white point
  - common to compute FP, then convert to integer
  - full range of values may not "fit" in display's output range
  - simplest solution: choose a maximum value, scale so that value becomes full intensity (2^n–1 in an n-bit integer image)

Intensity encoding in images

- What do the numbers in images (pixel values) mean?
  - they determine how bright that pixel is
  - bigger numbers are (usually) brighter

Datatypes for raster images

- For color or grayscale, sometimes add *alpha* channel
  - describes transparency of images

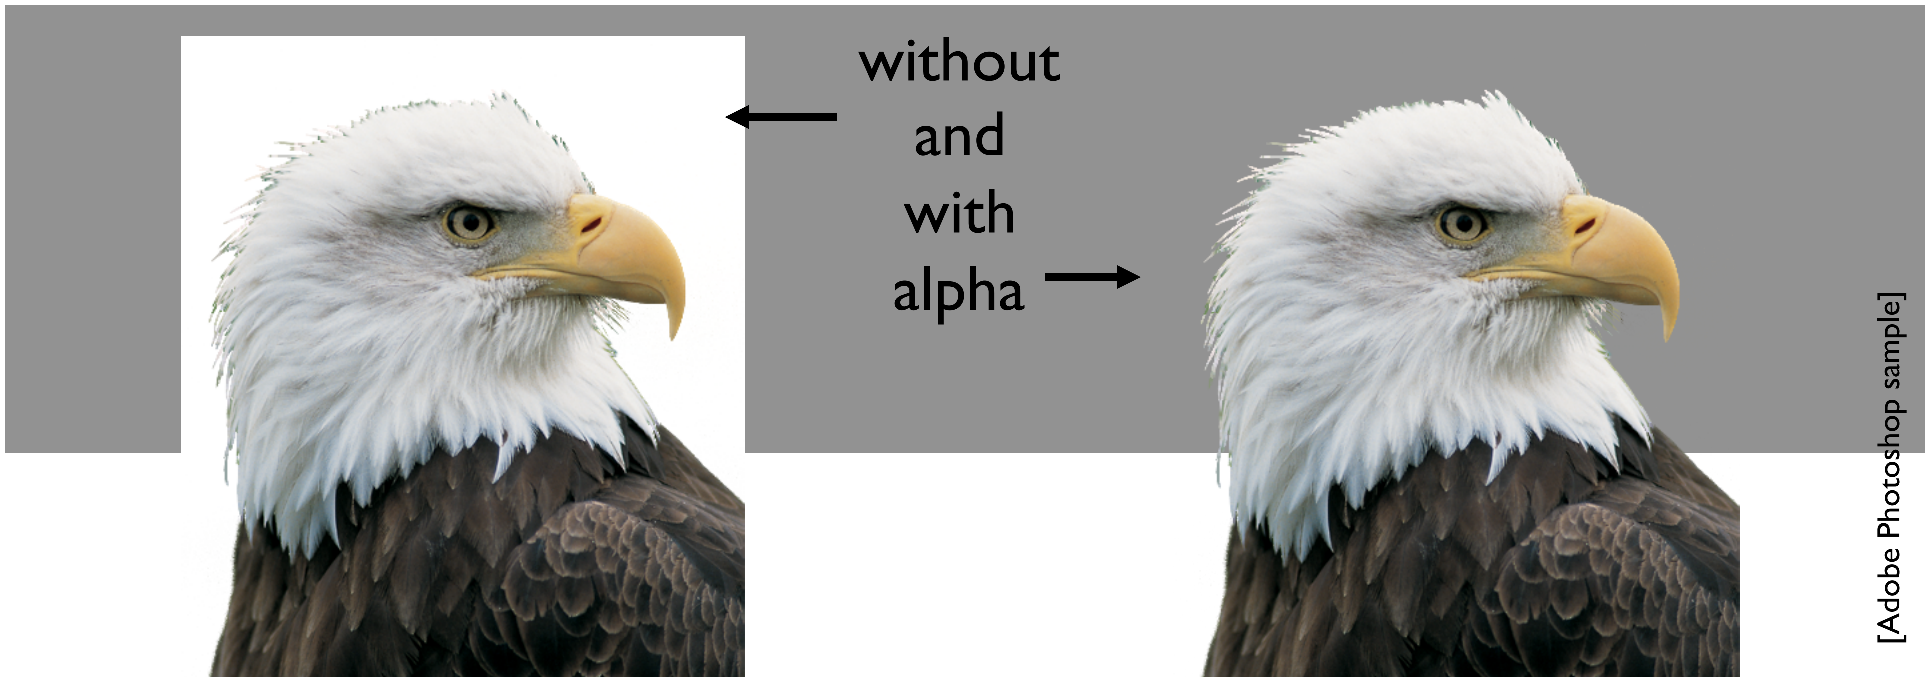

[Adobe Photoshop sample]

Storage requirements for images

- 1024x1024 image (1 megapixel)
  - bitmap: 128KB
  - grayscale 8bpp: 1MB
  - grayscale 16bpp: 2MB
  - color 24bpp: 3MB
  - floating-point HDR color: 12MB

Converting pixel formats

- Color to gray
  - could take one channel (blue, say)
    - leads to odd choices of gray value
  - combination of channels is better
    - but different colors contribute differently to lightness
    - which is lighter, full blue or full green?
    - good choice: gray = 0.2 R + 0.7 G + 0.1 B
    - more on this in color, later on

Same pixel values.

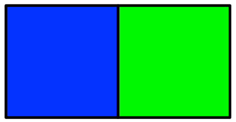

Same luminance?

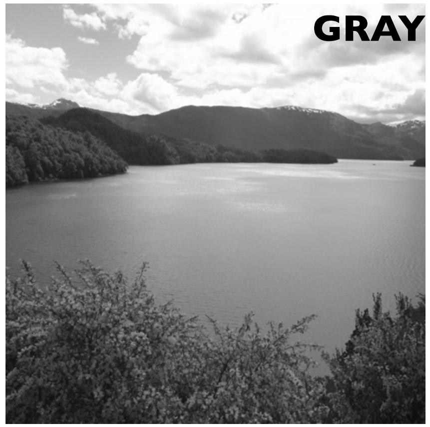

Converting pixel precision

- Up is easy; down loses information—be careful

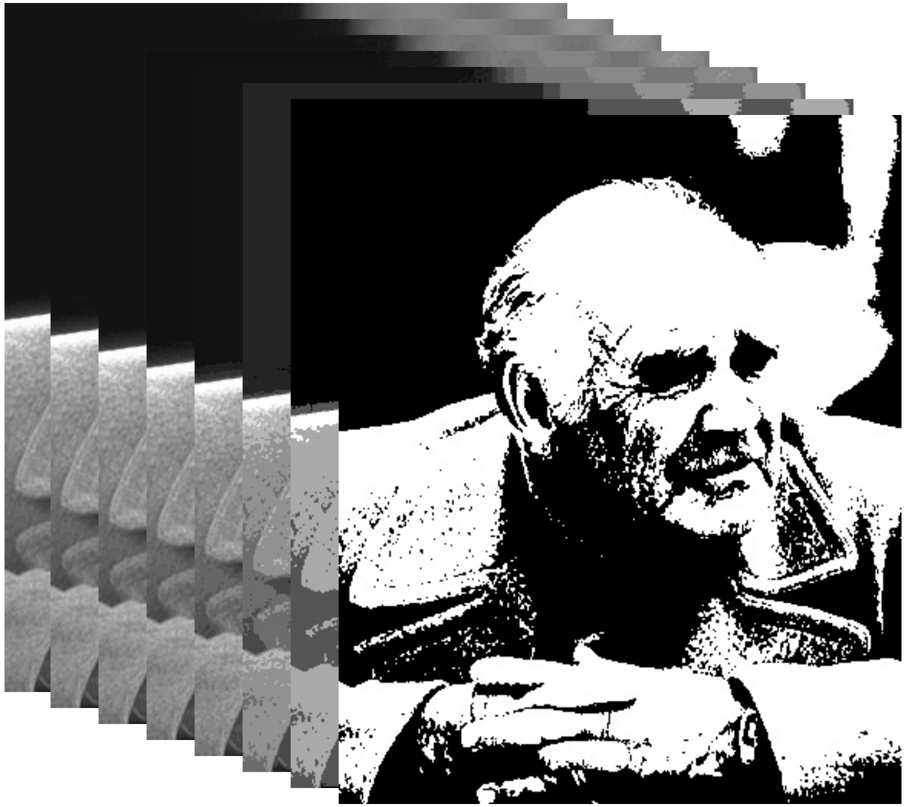

1 bpp (2 grays)

[photo: Philip Greenspun]

Dithering

- When decreasing bpp, we quantize
- Make choices consistently: banding
- Instead, be inconsistent—dither
  - turn on some pixels but not others in gray regions
  - a way of trading spatial for tonal resolution
  - choose pattern based on output device
  - laser, offset: clumped dots required (halftone)
  - inkjet, screen: dispersed dots can be used

Dithering methods

- Ordered dither
  - based on traditional, optically produced halftones
  - produces larger dots
- Diffusion dither
  - takes advantage of devices that can reproduce isolated dots
  - the modern winner for desktop printing

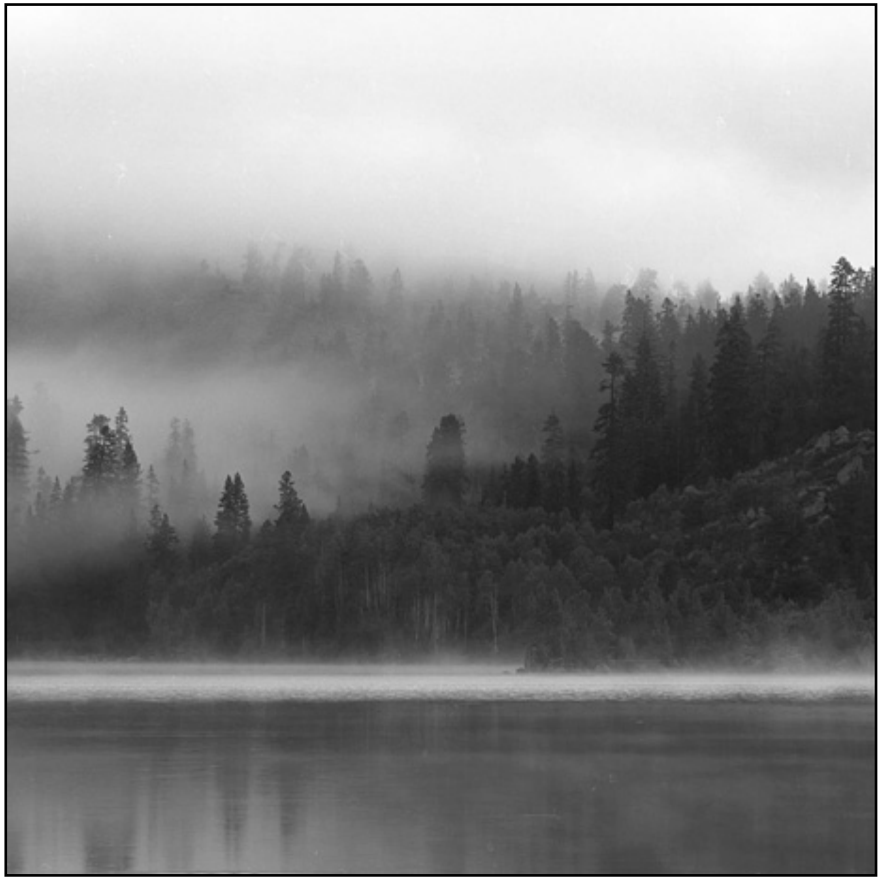

[Philip Greenspun]

Ordered Dither example

- Produces regular grid of compact dots

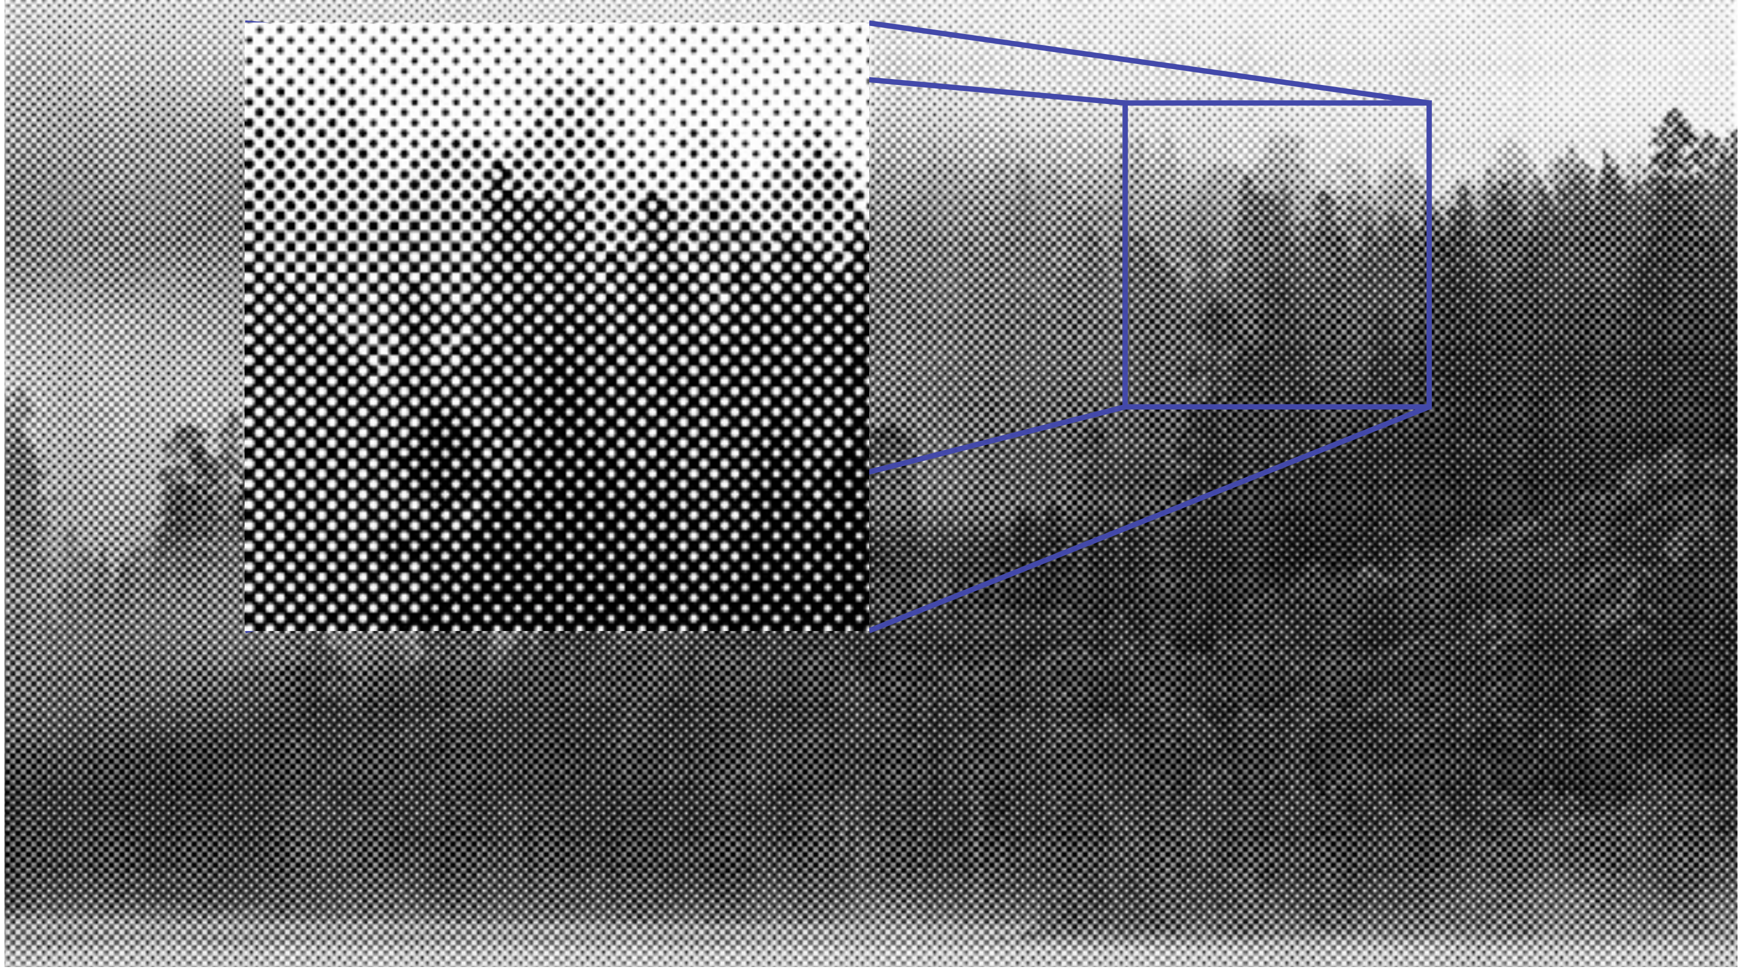

[photo: Philip Greenspun]

Diffusion dither

- Produces scattered dots with the right local density

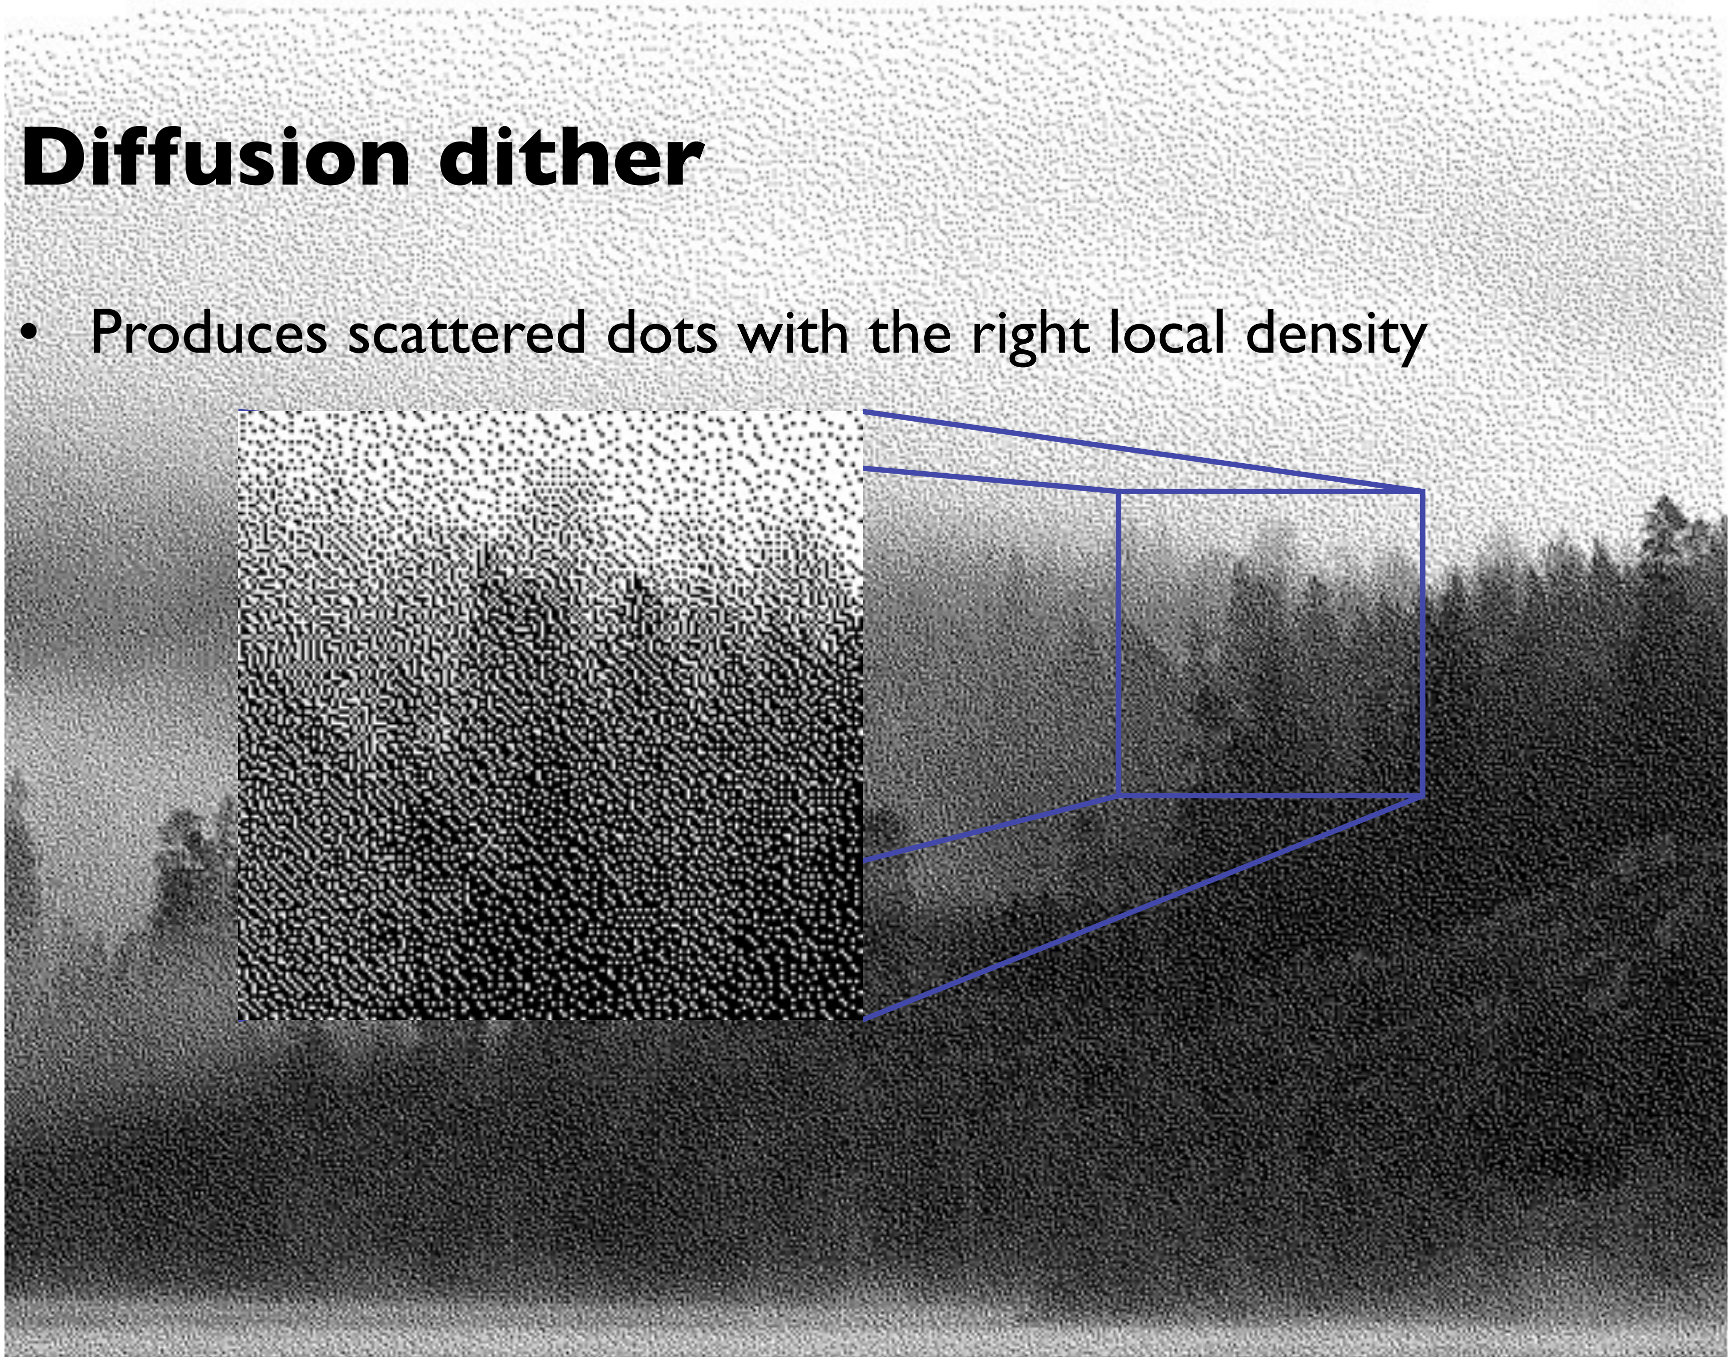

[photo: Philip Greenspun]