Image Processing

Overview

Images

Pixel Filters

Neighborhood Filters

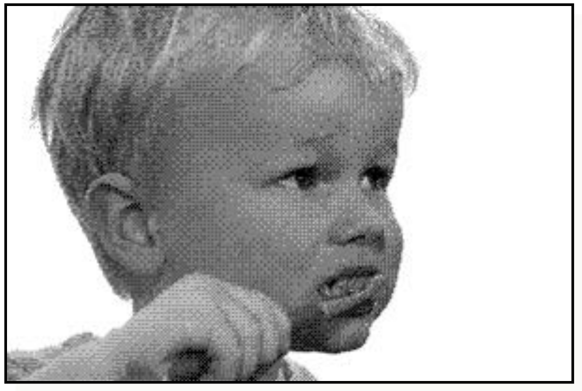

Dithering

Image as a Function

- We can think of an **image** as a function, *f*,
- f: $R^2 \rightarrow R$
  - $f(x, y)$ gives the **intensity** at position (x, y)
  - Realistically, we expect the image only to be defined over a rectangle, with a finite range:
    - f: $[a,b] \times [c,d] \rightarrow [0,1]$
- A color image is just three functions pasted together. We can write this as a "vector-valued" function:

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

Image as a Function

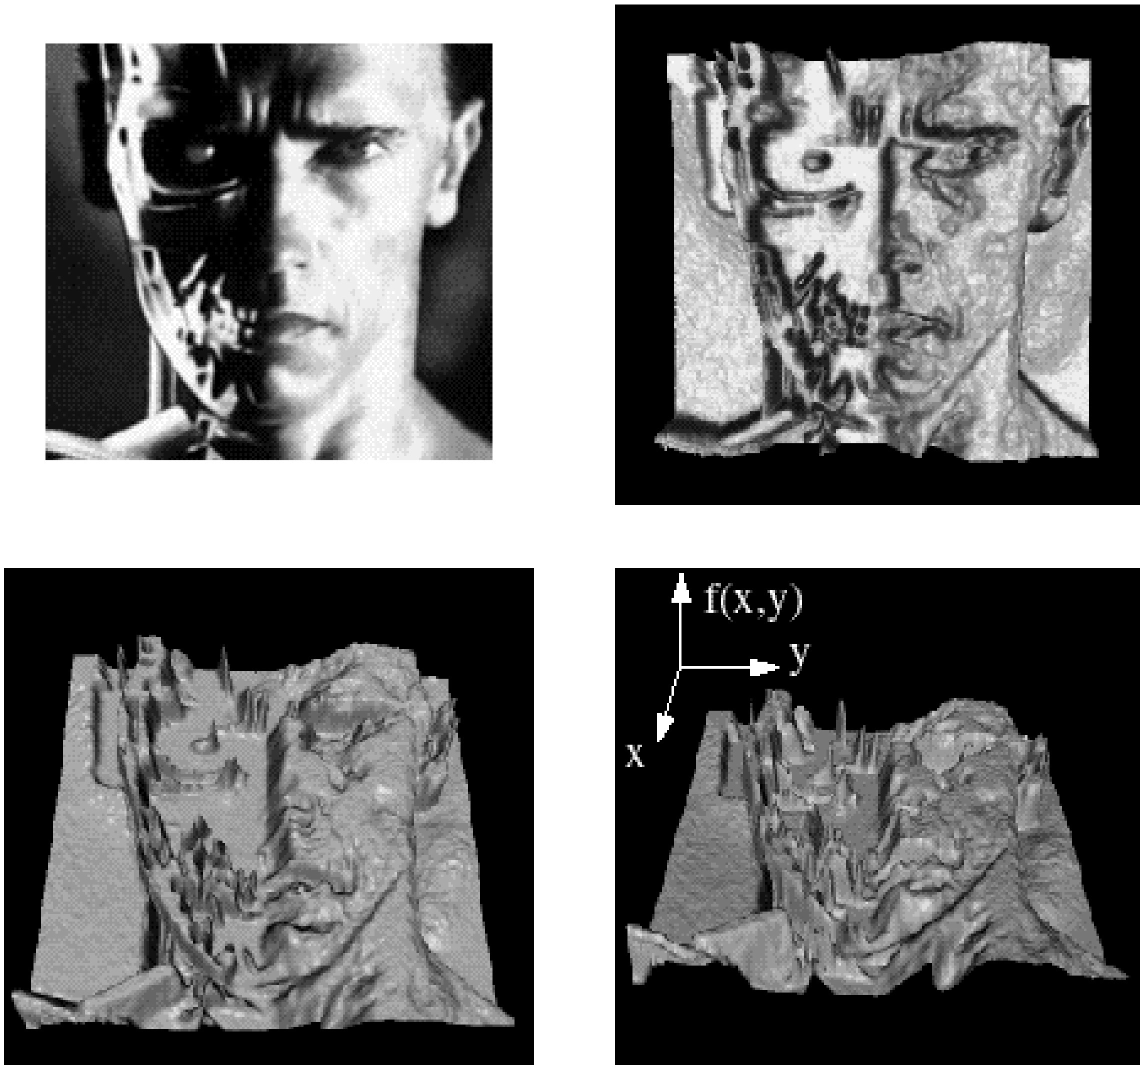

Image Processing

- Define a new image *g* in terms of an existing image *f*
  - We can transform either the domain or the range of *f*
- Range transformation:

$$g(x, y) = t(f(x, y))$$

What kinds of operations can this perform?

Image Processing

- Some operations preserve the range but change the domain of *f* :

$$g(x, y) = f(t_x(x, y), t_y(x, y))$$

What kinds of operations can this perform?

- Still other operations operate on both the domain and the range of f .

Point Operations

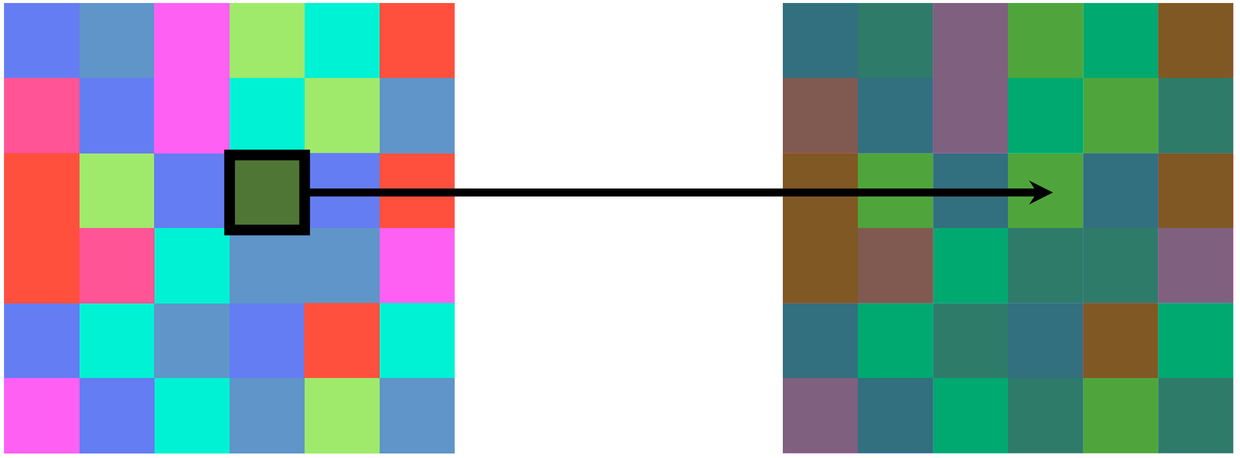

Point Processing

Original

Darken

Lower Contrast

Nonlinear Lower Contrast

Invert

Lighten

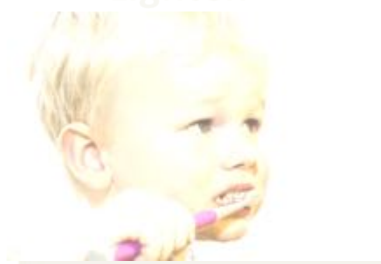

Raise Contrast

Nonlinear Raise Contrast

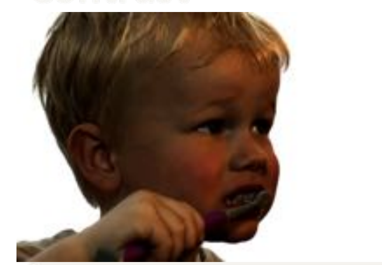

Point Processing

Original

`x`

Darken

`x - 128`

Lower Contrast

`x / 2`

Nonlinear Lower Contrast

`((x / 255.0) ^ 0.33) * 255.0`

Invert

`255 - x`

Lighten

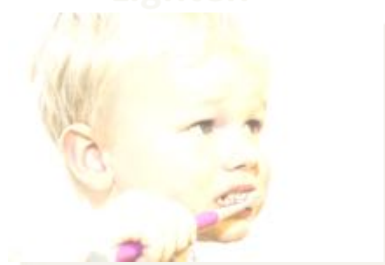

`x + 128`

Raise Contrast

`x * 2`

Nonlinear Raise Contrast

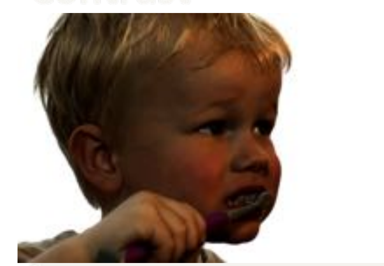

`((x / 255.0) ^2) * 255.0`

Gamma correction

Monitors have a intensity to voltage response curve which is roughly a 2.5 power function

Send v → actually display a pixel which has intensity equal to $v^{2.5}$

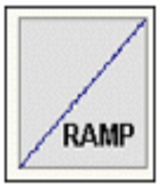

Graph of Input

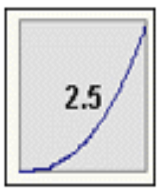

Graph of Output L = V ^ 2.5

$\Gamma = 1.0; f(v) = v$

$\Gamma = 2.5; f(v) = v^{1/2.5} = v^{0.4}$

Neighborhood Operations

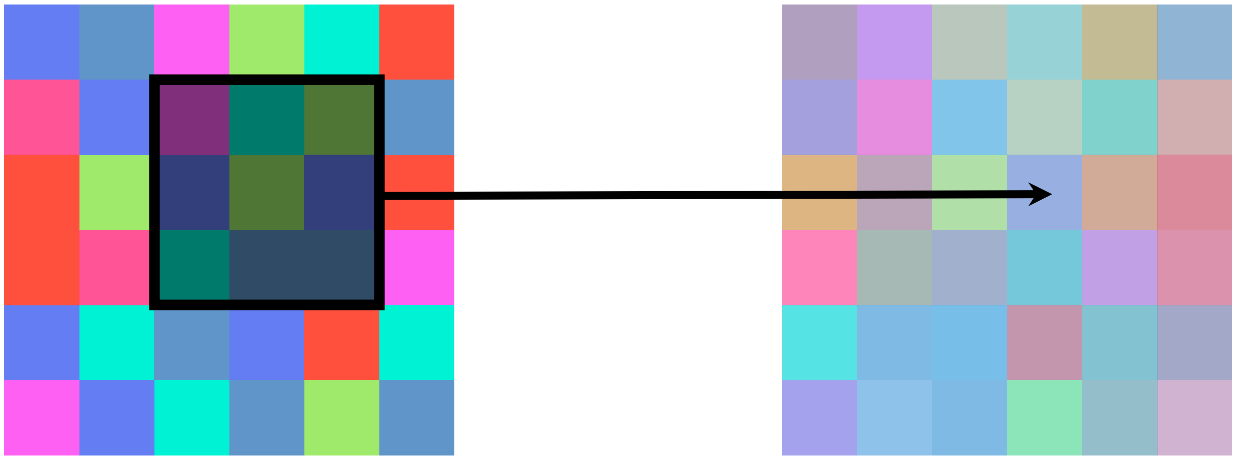

Convolution

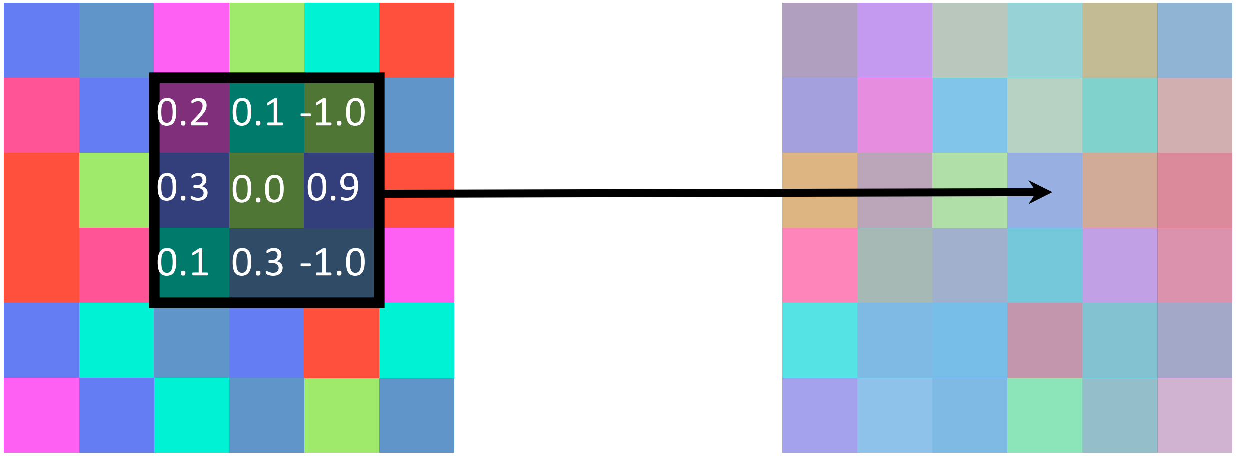

$$F = \begin{bmatrix} 0.2 & 0.1 & -1.0 \\ 0.3 & 0.0 & 0.9 \\ 0.1 & 0.3 & -1.0 \end{bmatrix}$$

$$I' = F * I$$

Properties of Convolution

- Commutative

$$a * b = b * a$$

- Associative

$$(a * b) * c = a * (b * c)$$

- Cascade system

$$f \longrightarrow \boxed{h_1} \longrightarrow \boxed{h_2} \longrightarrow g$$

$$= f \longrightarrow \boxed{h_1 * h_2} \longrightarrow g$$

$$= f \longrightarrow \boxed{h_2 * h_1} \longrightarrow g$$

Convolution

Convolution is linear and shift invariant

$$g(x) = \int_{-\infty}^{\infty} f(\tau)h(x-\tau)d\tau \qquad g = f * h$$

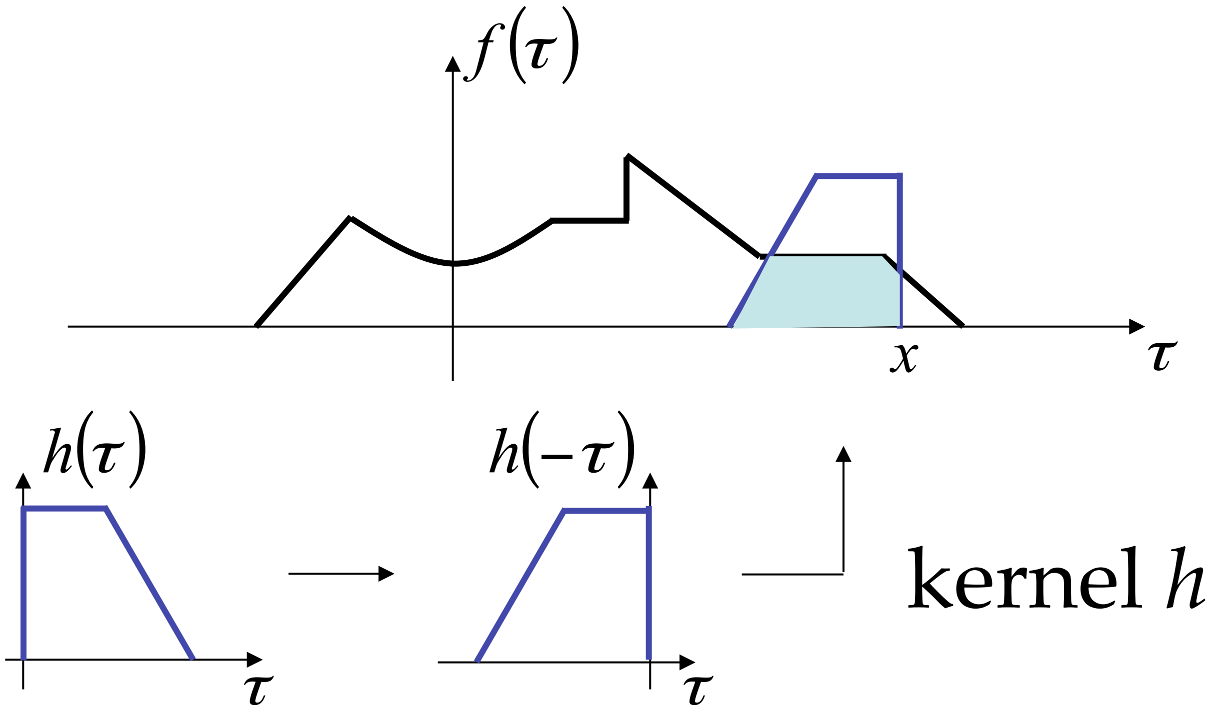

Convolution - Example

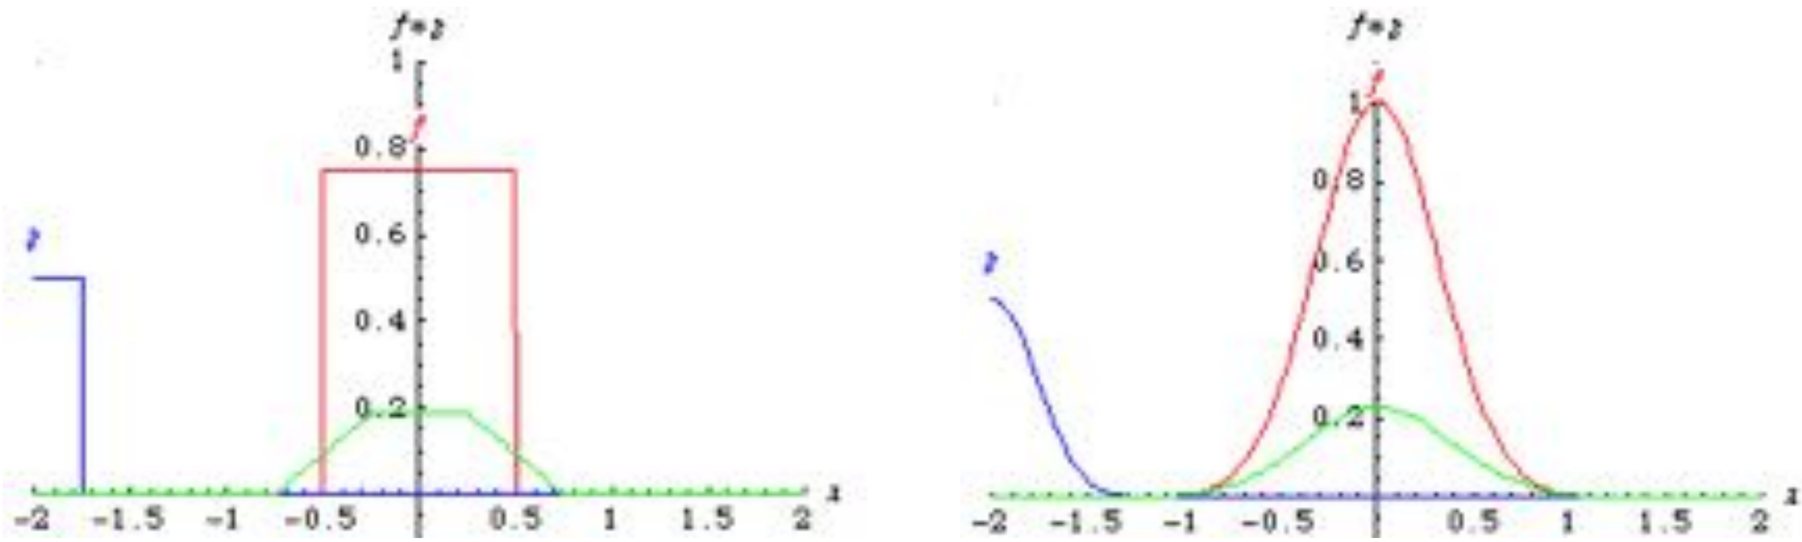

f

g

$f * g$

Eric Weinstein's Math World

Convolution - Example

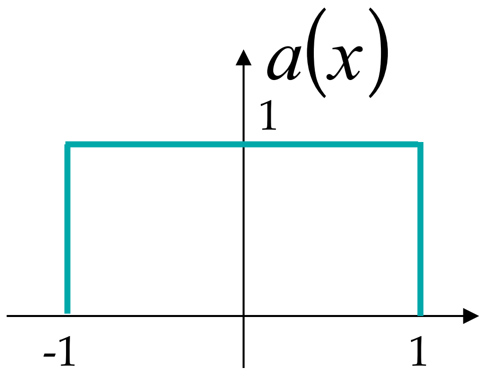

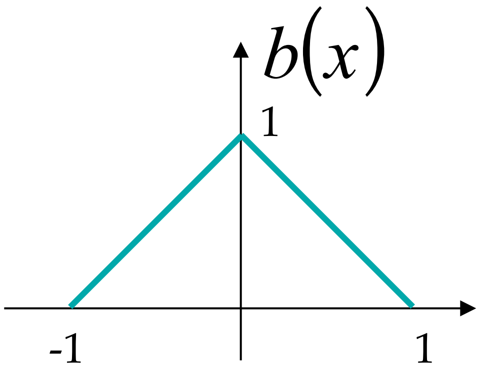

$$\downarrow \; c = a * b$$

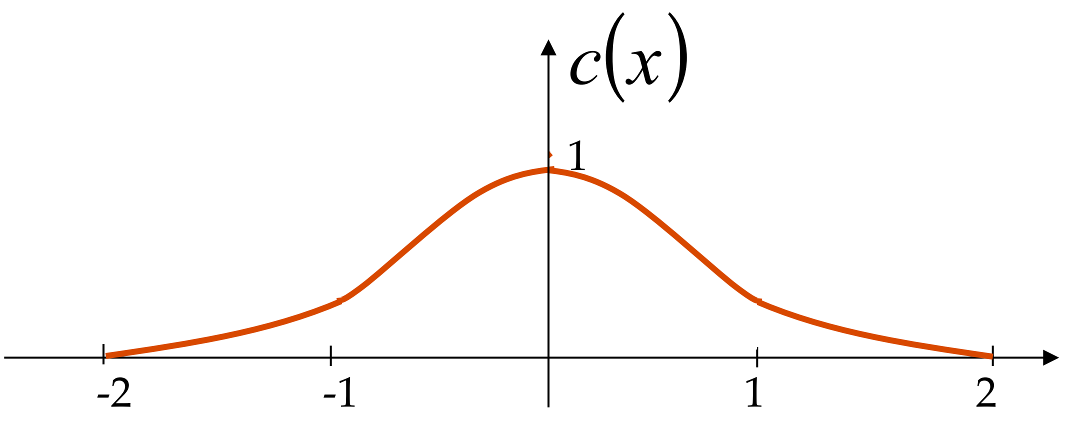

Point Spread Function

$$\text{scene} \longrightarrow \boxed{\text{Optical System}} \longrightarrow \text{image}$$

- Ideally, the optical system should be a Dirac delta function.

- However, optical systems are never ideal.

$$\underset{\text{point source}}{\delta(x)} \longrightarrow \boxed{\text{Optical System}} \longrightarrow \underset{\text{point spread function}}{PSF(x)}$$

- Point spread function of Human Eyes

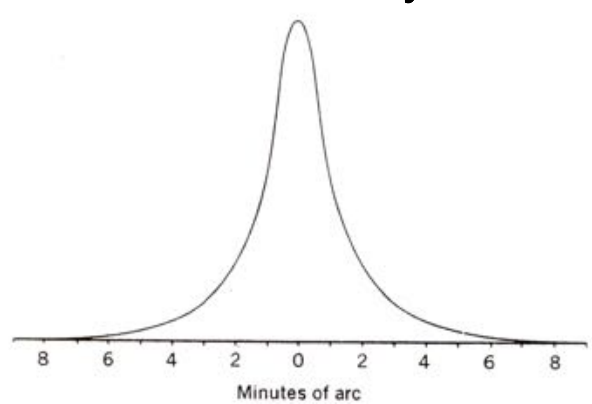

Point Spread Function

normal vision

myopia

hyperopia

astigmatism

Images by Richmond Eye Ass

Original Image

Blurred Image

Gaussian Smoothing

by Charles Allen Gillbert

by Harmon & Julesz

http://www.michaelbach.de/ot/cog_blureffects

Gaussian Smoothing

http://www.michaelbach.de/ot/cog_blureffects

Original Image

Sharpened Image

Sharpened Image

Noise

Blurred Noise

Median Filter

- Smoothing is averaging
  (a) Blurs edges
  (b) Sensitive to outliers

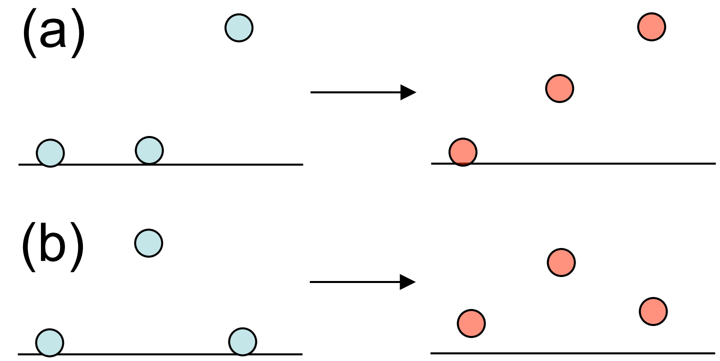

- Median filtering
  - Sort $N^2 - 1$ values around the pixel
  - Select middle value (median)

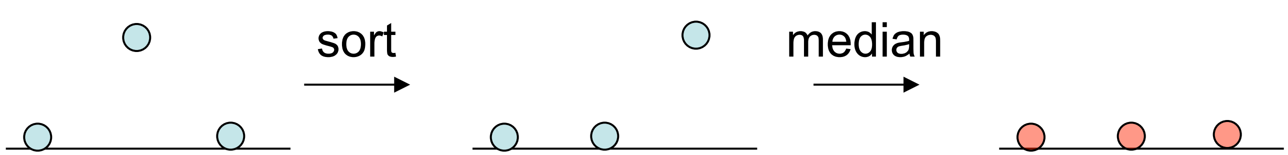

  - Non-linear (Cannot be implemented with convolution)

Noise

Blurred Noise

Median Filter

Can this be described as a convolution?

Original Image

Example: Noise Reduction

Image with noise

Median filter (5x5)

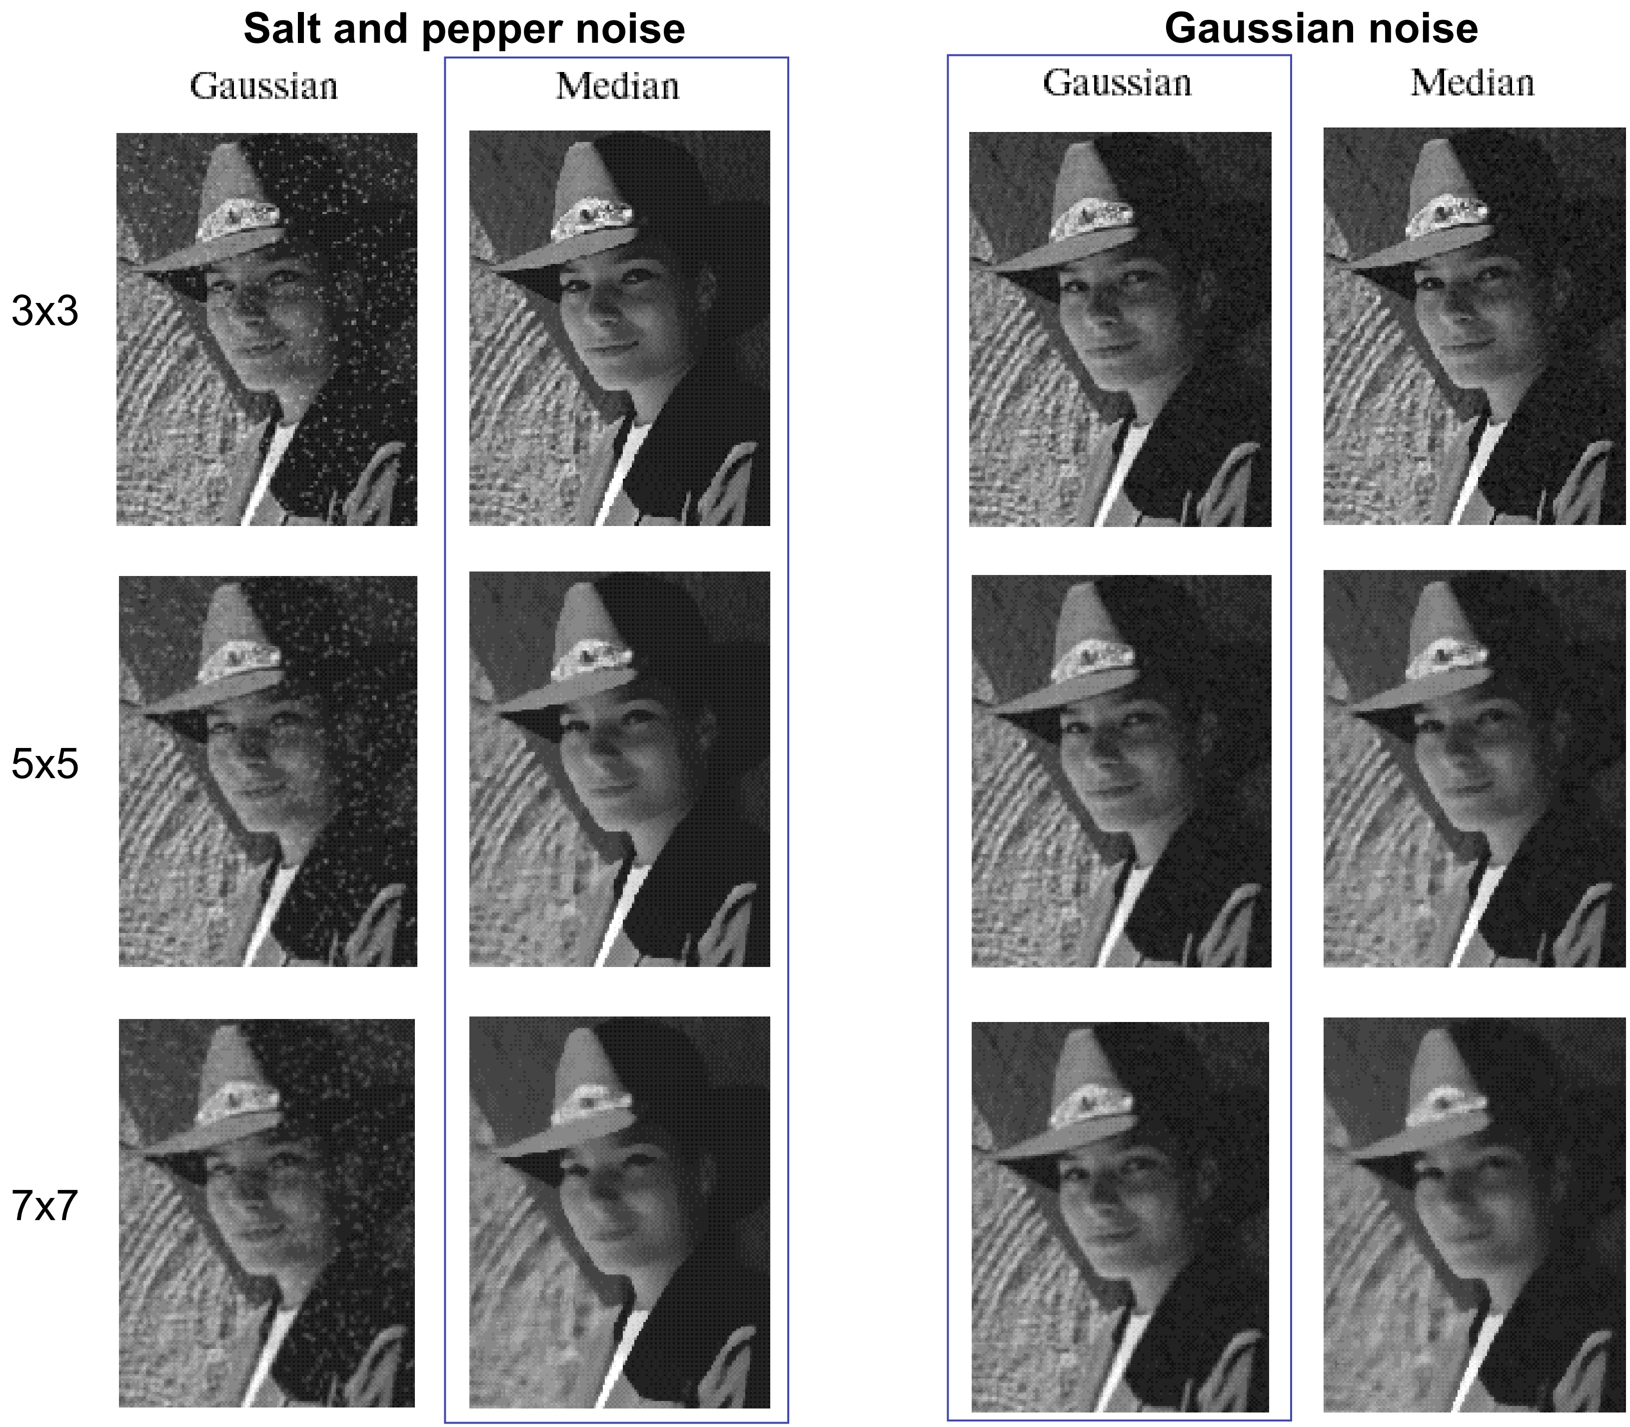
Salt and pepper noise
Gaussian noise
Gaussian
Median
Gaussian
Median
3x3
5x5
7x7

Example: Noise Reduction

Original image

Image with noise

Median filter (5x5)

Original Image

X-Edge Detection

Y-Edge Detection

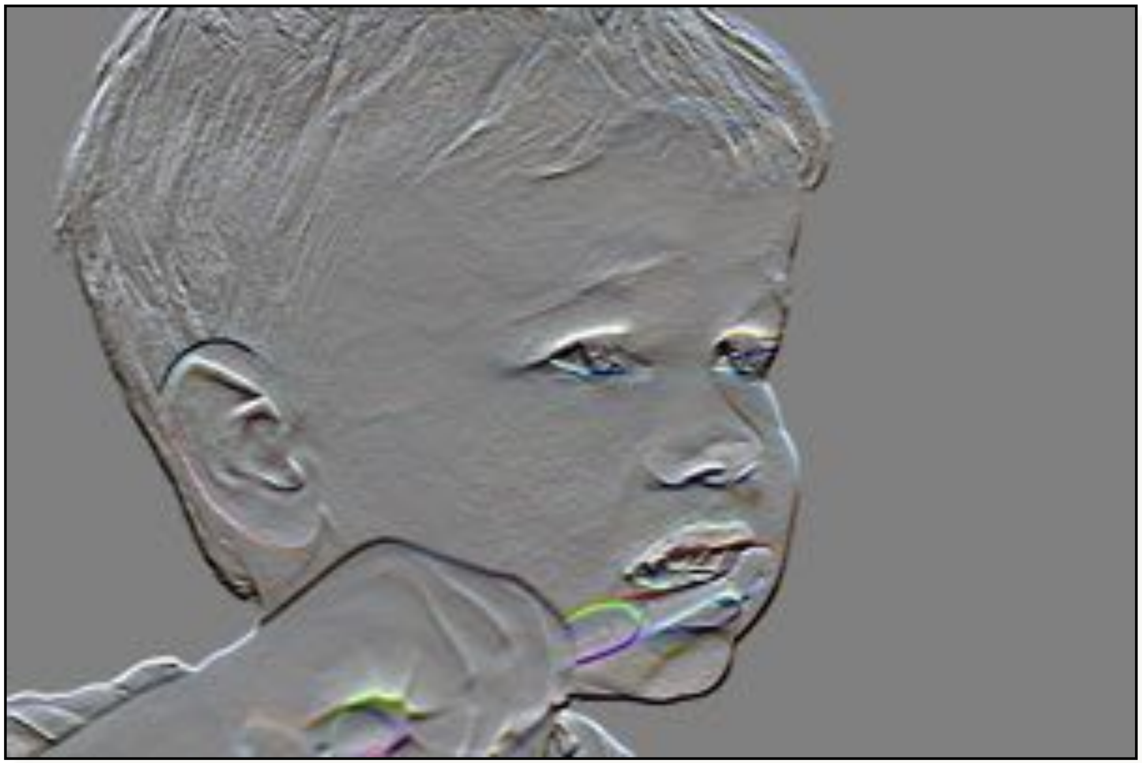

General Edge Detection

Can this be described as a convolution?

Image Processing

- Some operations preserve the range but change the domain of *f* :

$$g(x, y) = f(t_x(x, y), t_y(x, y))$$

What kinds of operations can this perform?

- Still other operations operate on both the domain and the range of f .

Image Scaling

This image is too big to fit on the screen. How can we reduce it?

How to generate a half-sized version?

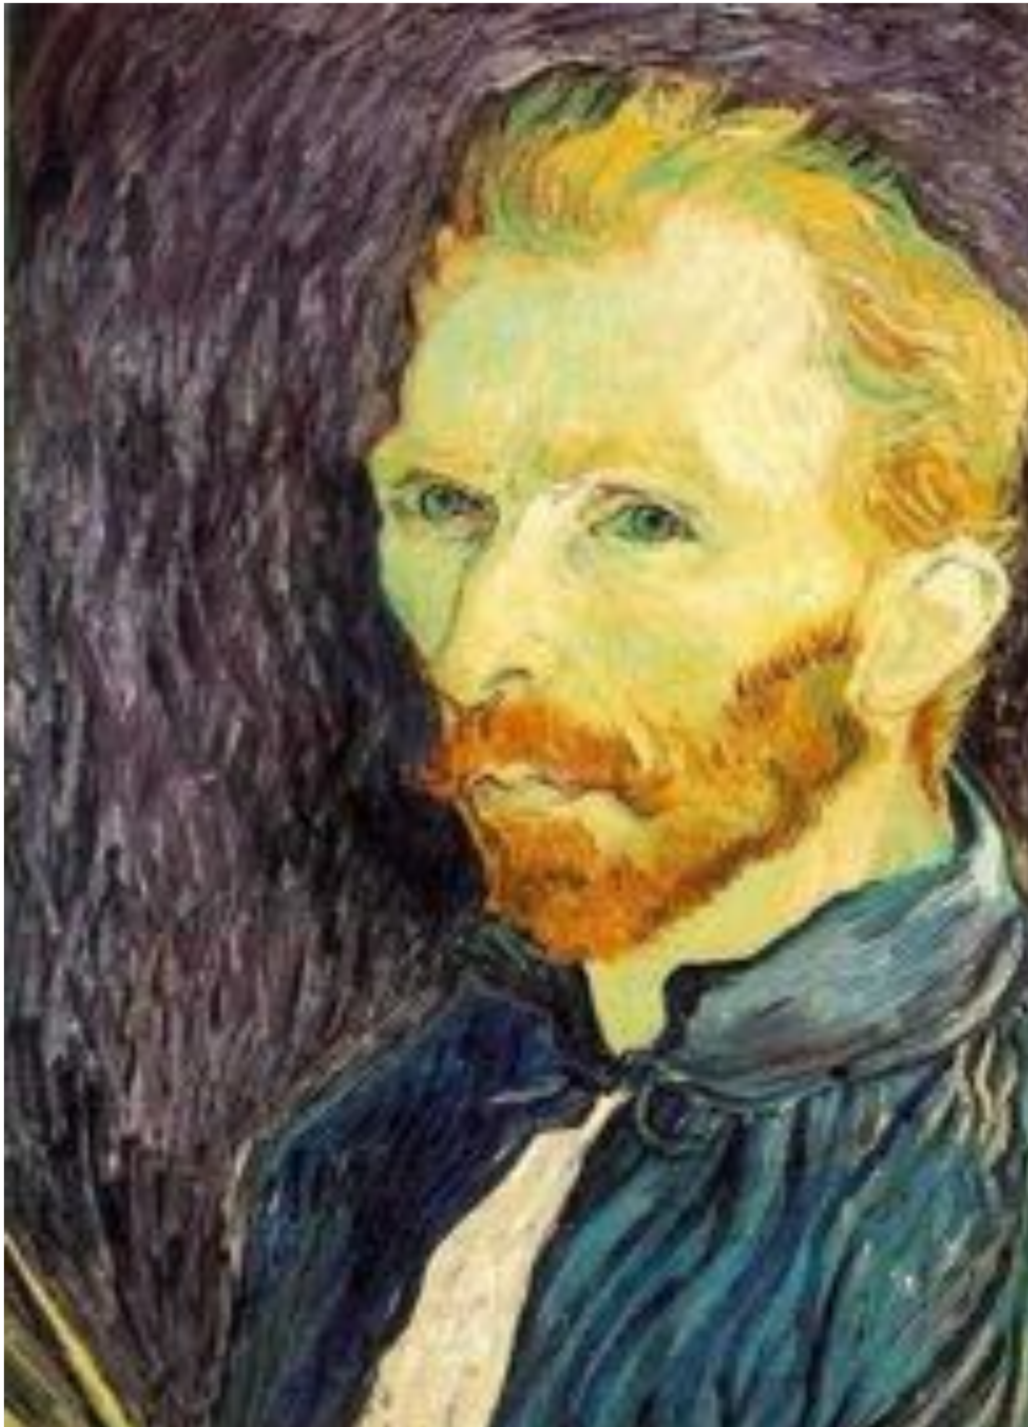

Image Sub-Sampling

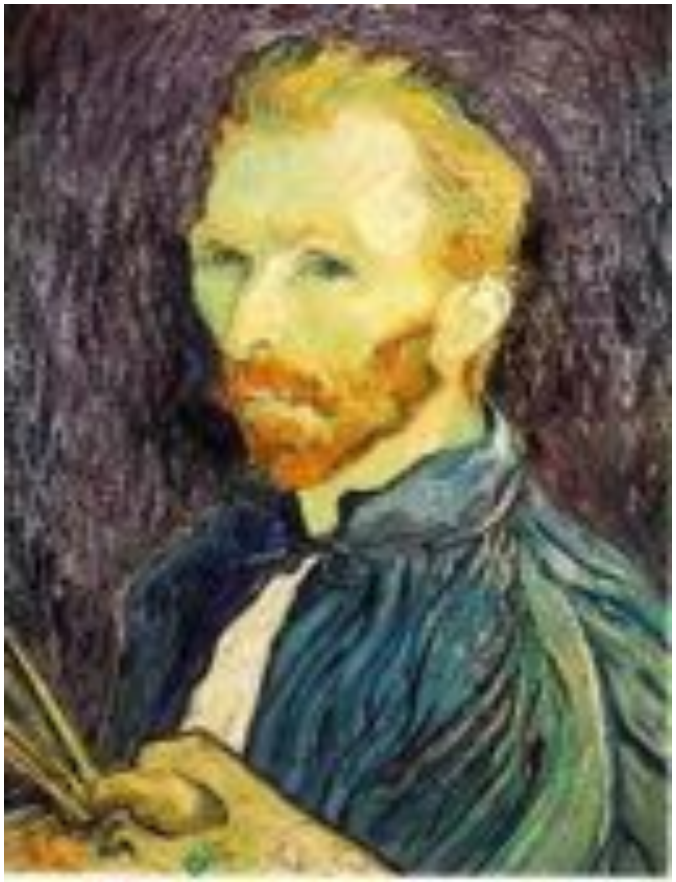

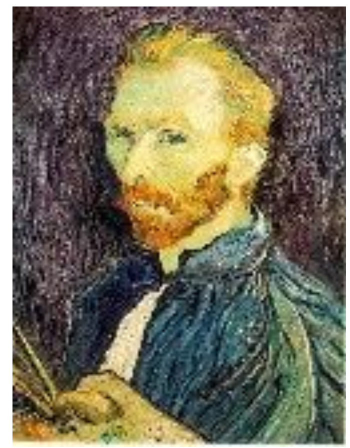

1/4

1/8

Throw away every other row and column to create a 1/2 size image
- called *image sub-sampling*

Image Sub-Sampling

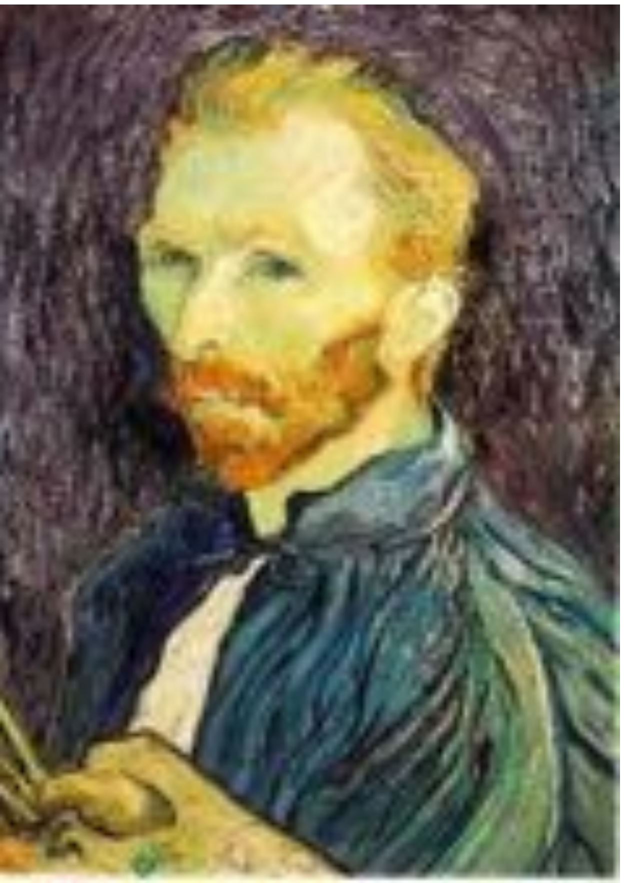

1/2

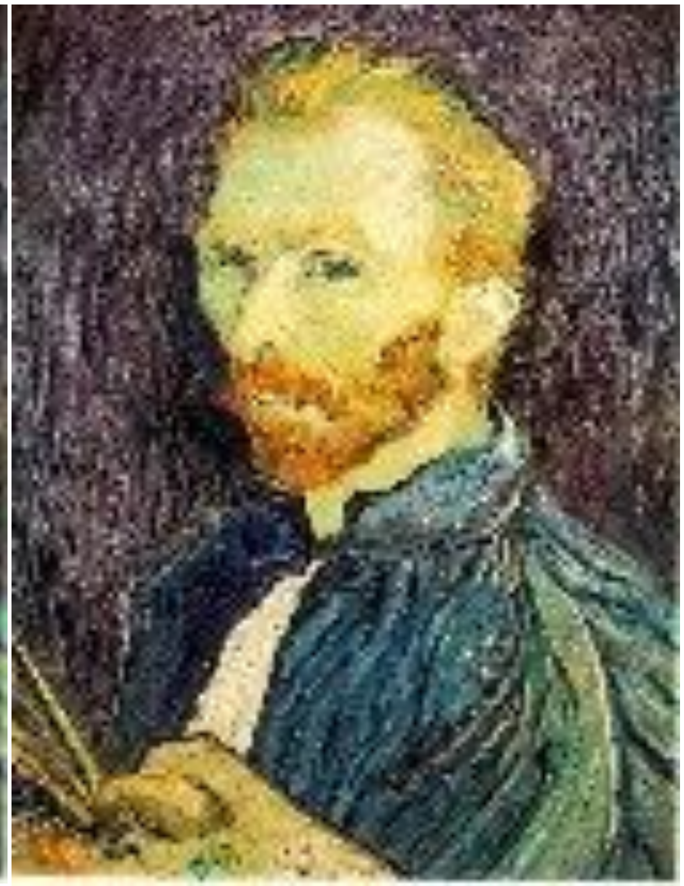

1/4 (2x zoom)

1/8 (4x zoom)

Good and Bad Sampling

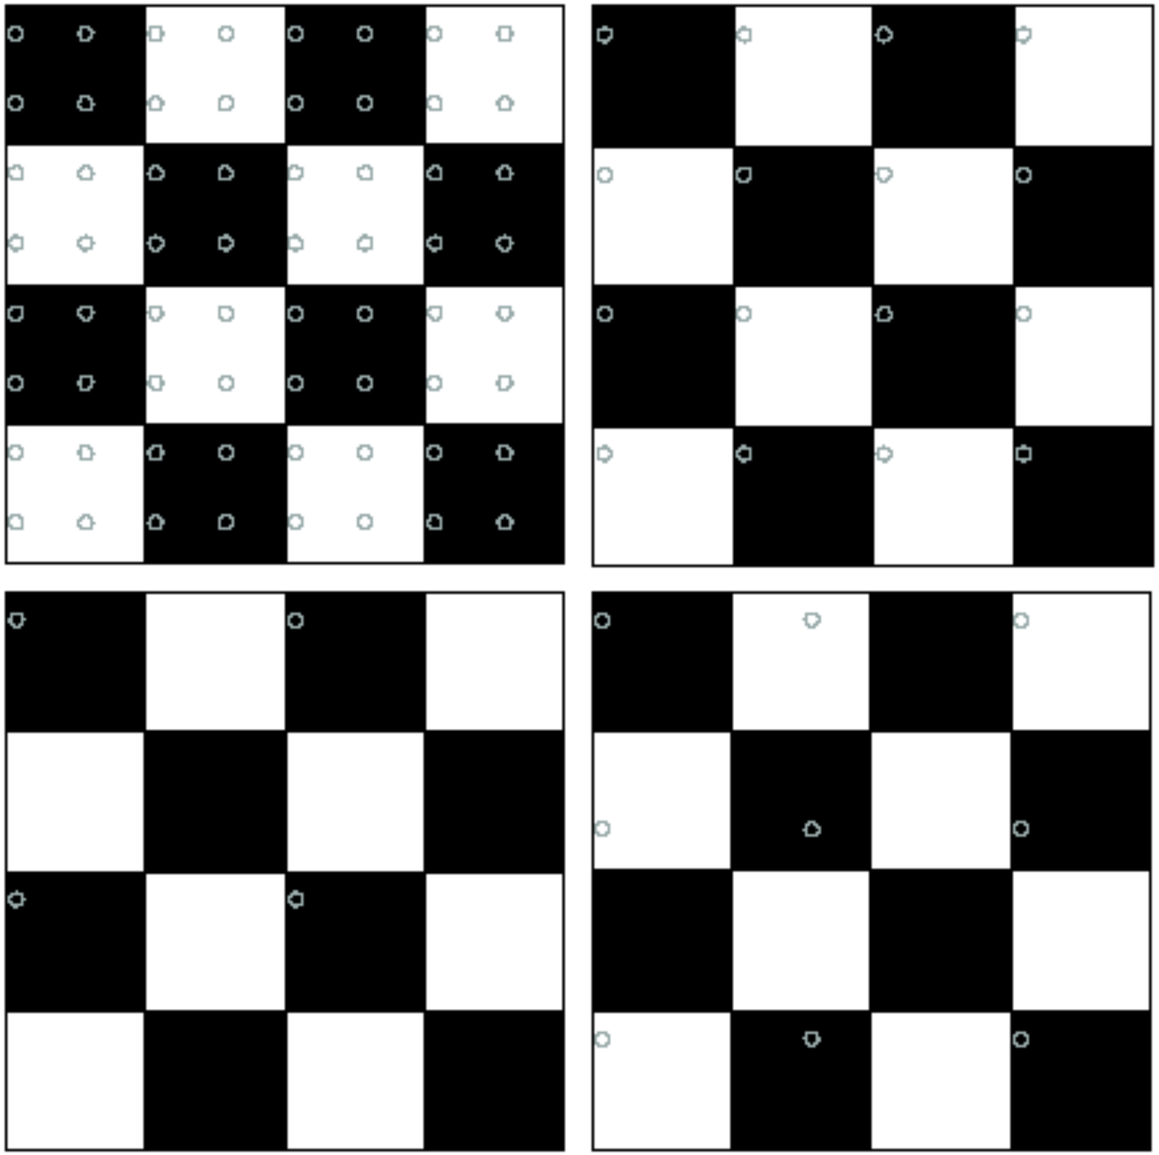

Good sampling:
- Sample often or,
- Sample wisely

Bad sampling:
- see aliasing in action!

Aliasing

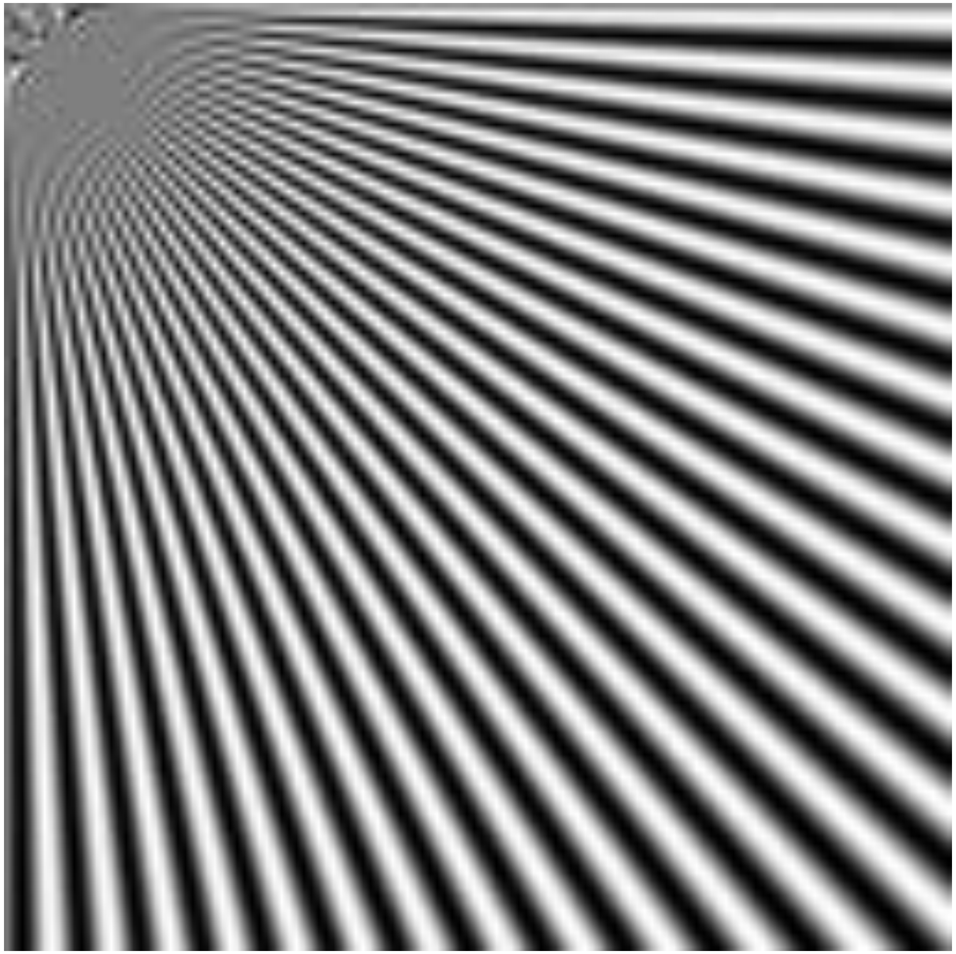

Alias: n., an assumed name

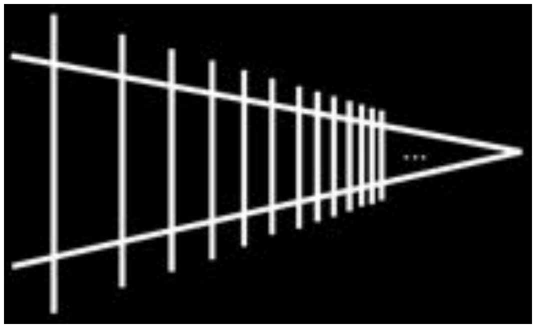

Picket fence receding into the distance will produce aliasing…

WHY?

Input signal:

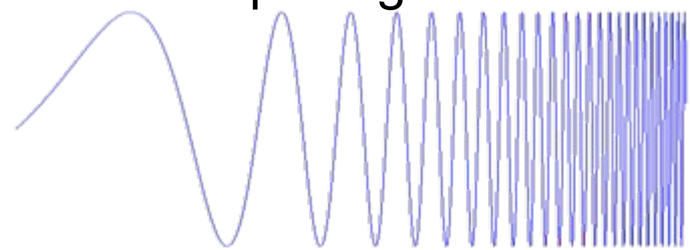

Matlab output:

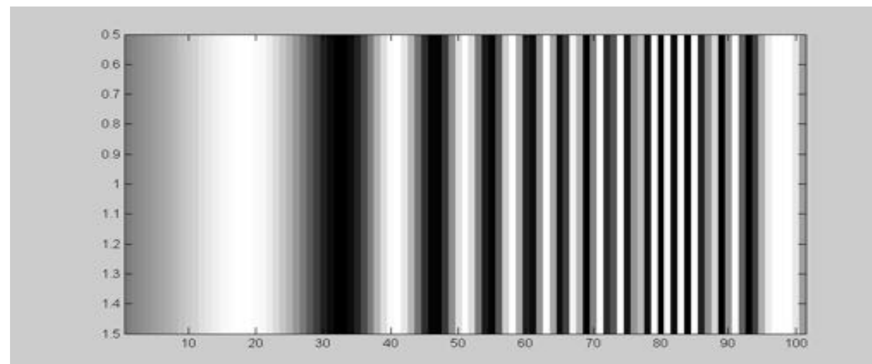

x = 0:.05:5; imagesc(sin((2.^x).*x))

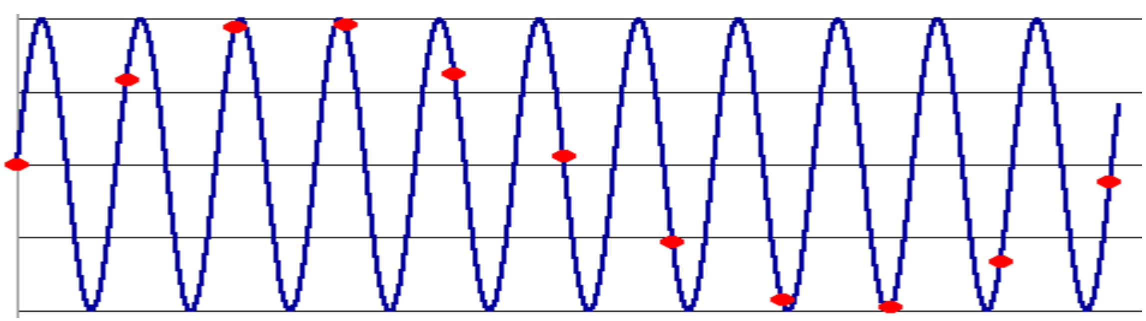

Alias!

Not enough samples

Really bad in video

Imagine a spoked wheel moving to the right (rotating clockwise). Mark wheel with dot so we can see what's happening.

If camera shutter is only open for a fraction of a frame time (frame time = 1/30 sec. for video, 1/24 sec. for film):

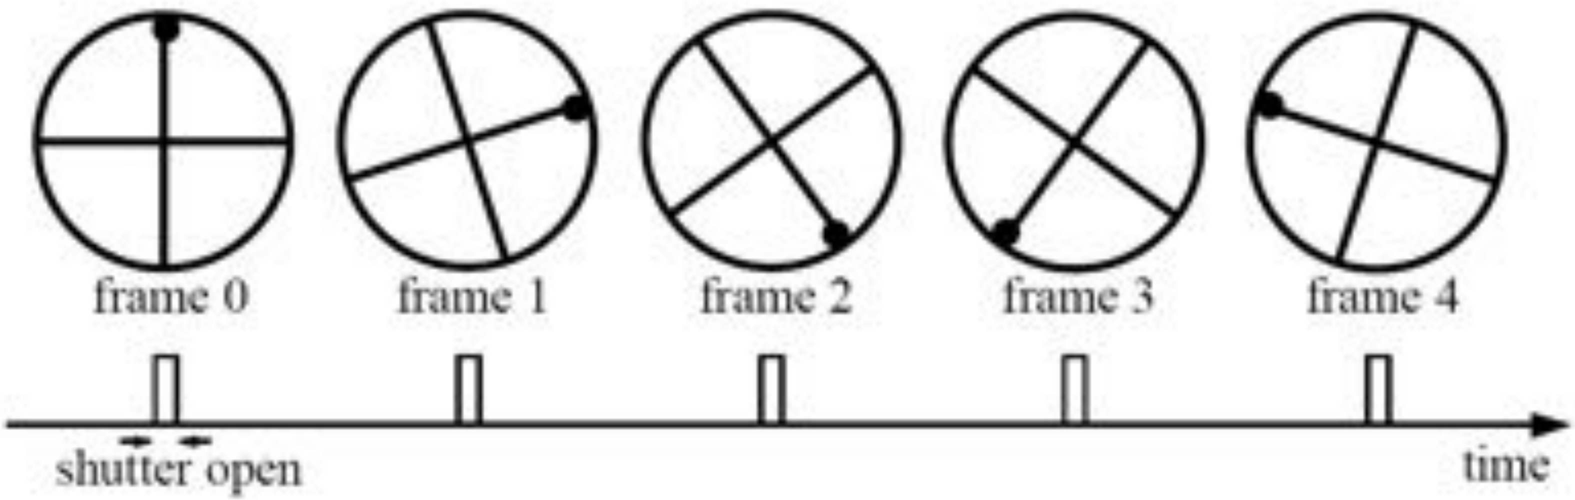

Without dot, wheel appears to be rotating slowly backwards! (counterclockwise)

Sub-Sampling with Gaussian Pre-Filtering

Gaussian 1/2

G 1/4

G 1/8

- Solution: filter the image, *then* subsample
  - Filter size should double for each ½ size reduction. Why?

Sub-Sampling with Gaussian Pre-Filtering

Gaussian 1/2

G 1/4

G 1/8

Compare with...

1/2

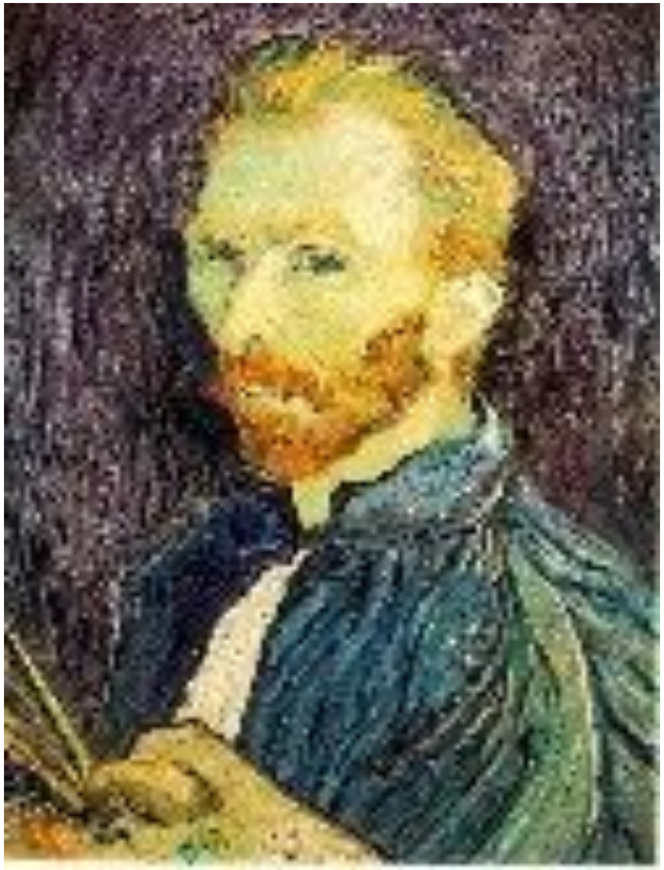

1/4 (2x zoom)

1/8 (4x zoom)

Canon D60 (w/ anti-alias filter)

Sigma SD9 (w/o anti-alias filter)

From Rick Matthews website, images by Dave Etchells

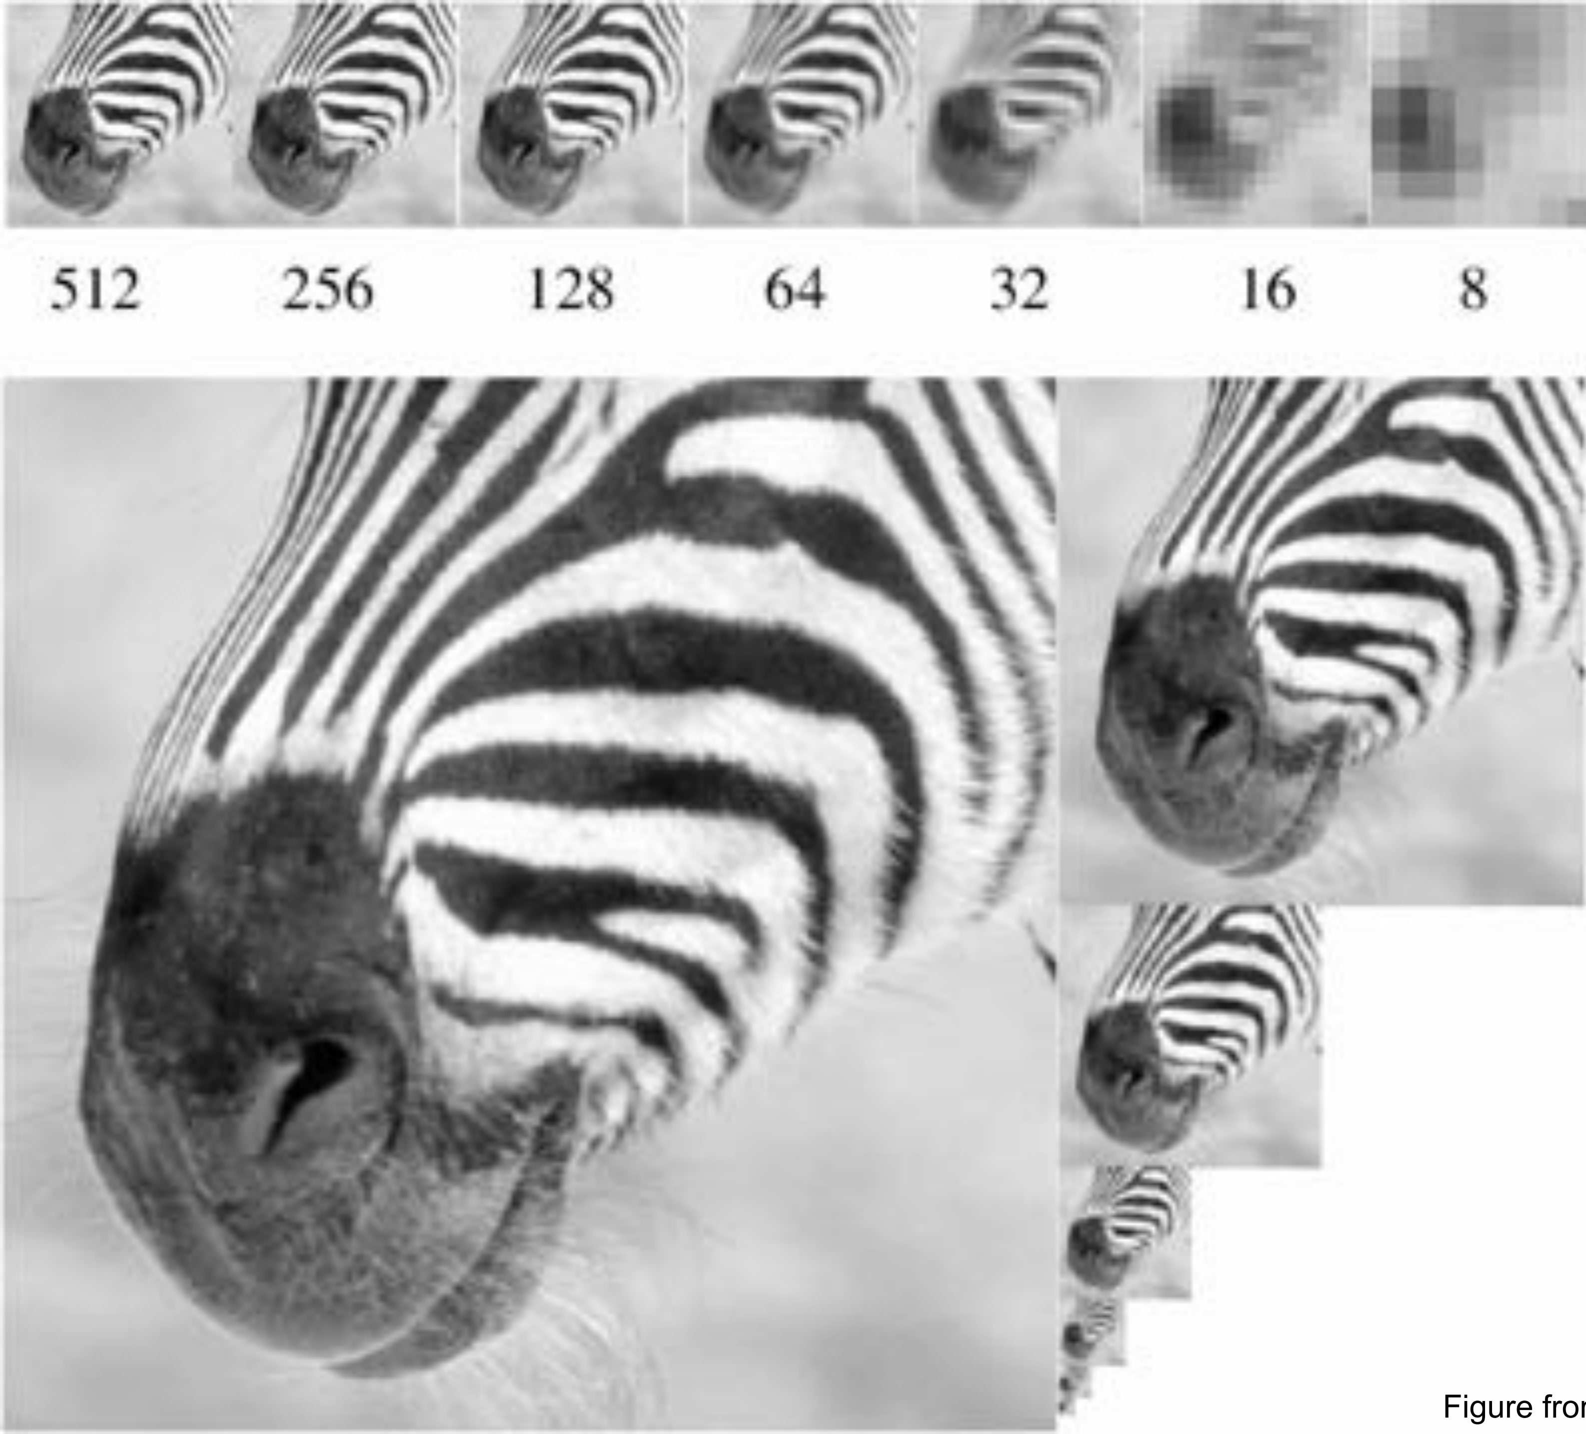

Figure from David Forsyth

Original Image

Warped Image

Warped Image

+

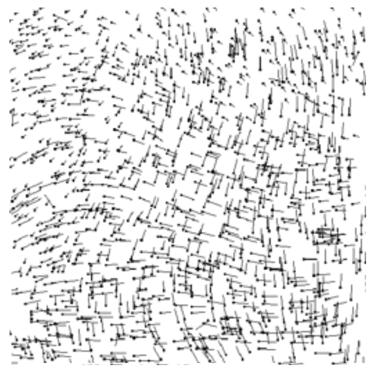

=

orig

vector field

warped

how?

Advection (just like a fluid)

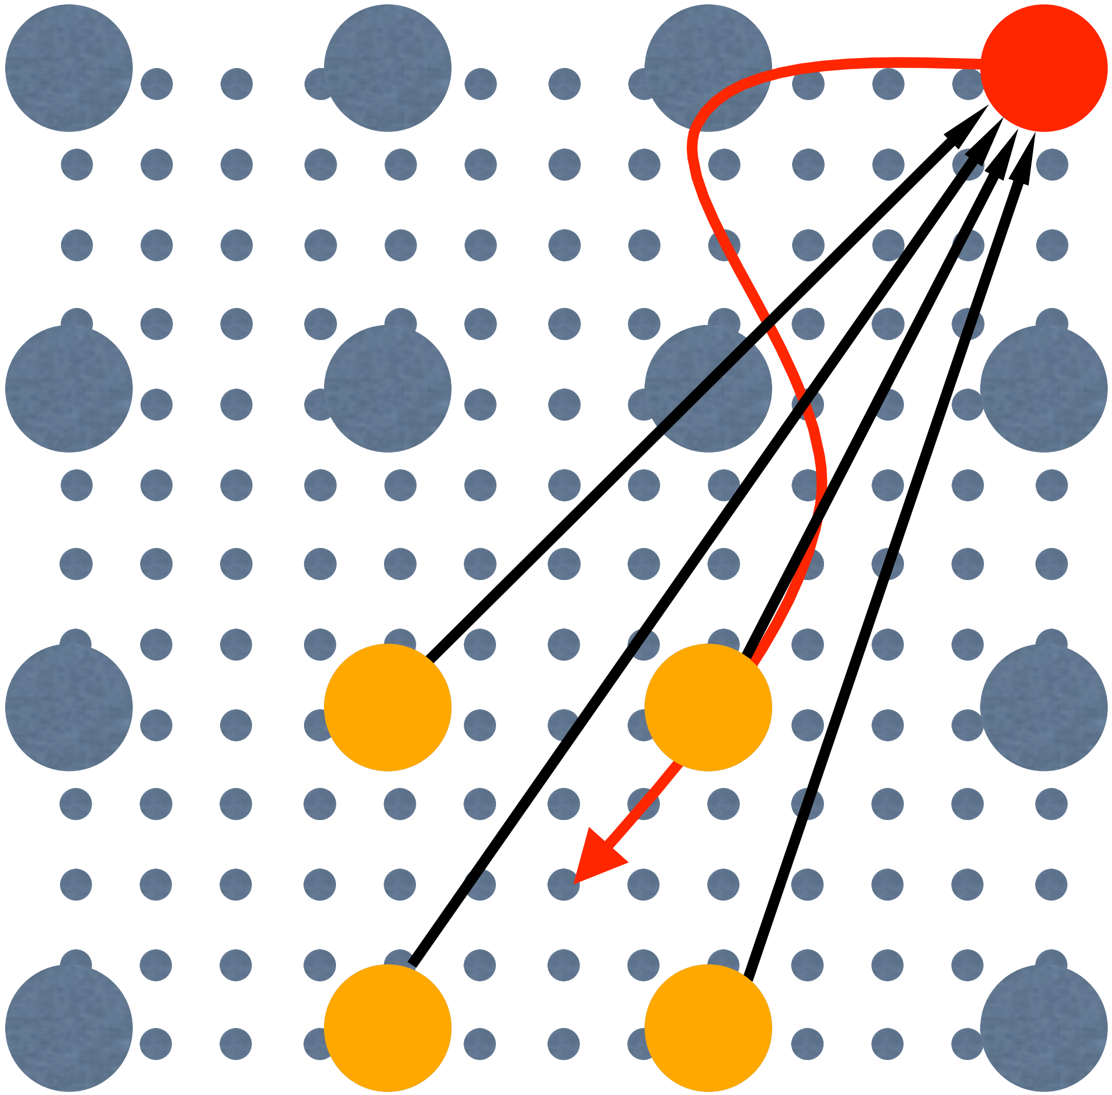

Check this out

Tailored Displays to Compensate for Visual Aberrations

Vitor F. Pamplona[1], Manuel M. Oliveira[1], Daniel G. Aliaga[2], Ramesh Raskar[3]

ACM SIGGRAPH 2012. ACM Trans. on Graph. 31(4). Pages 87:1-12

[1] UFRGS Instituto de Informática [2] Purdue University [3] MIT Media Lab

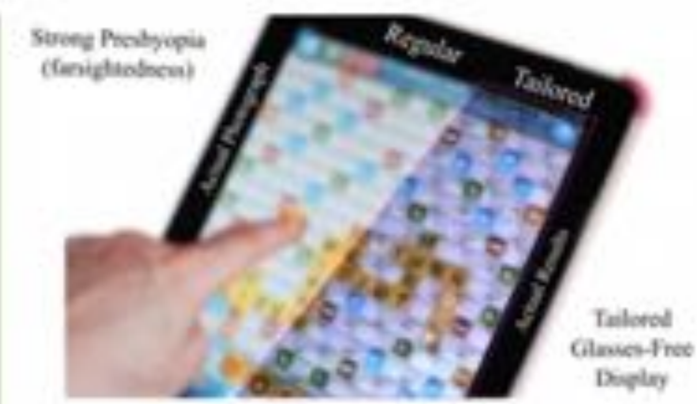

Can we create a display that adapts itself to improve one's eyesight? Top figure compares the view of a 2.5-diopter farsighted individual in regular and tailored displays. We use currently available inexpensive technologies to warp light fields to compensate for refractive errors and scattering sites in the eye.

http://tailoreddisplays.com/

http://eyenetra.com/index.html

Published in *Computer Graphics*, **25**(4), July 1991, pp. 319-328.
(ACM SIGGRAPH '91 Conference Proceedings, Las Vegas, Nevada, July 1991.)

Artificial Evolution for Computer Graphics

Karl Sims

Thinking Machines Corporation

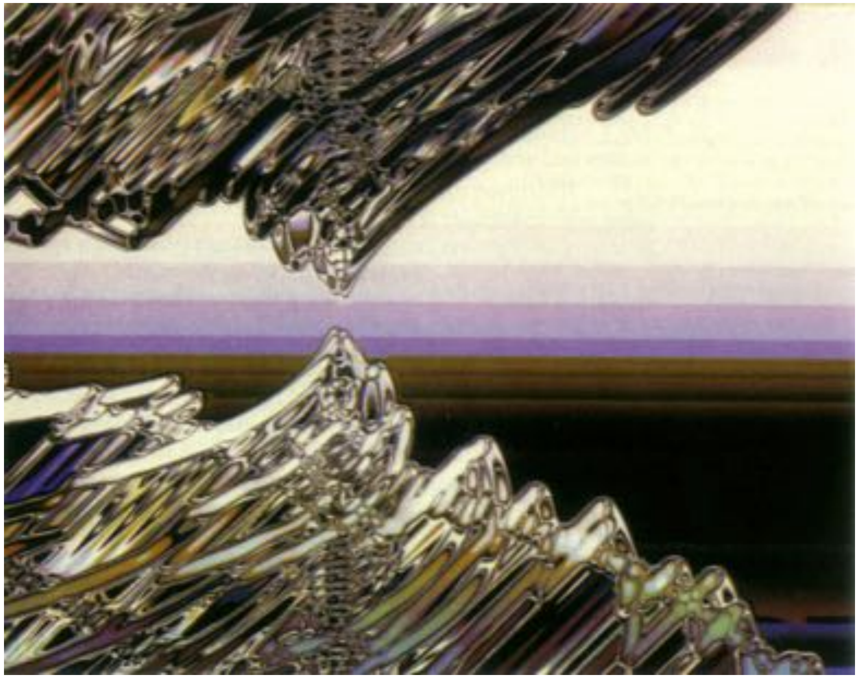

http://www.karlsims.com/genetic-images.html

http://www.karlsims.com/papers/siggraph91.html

High Dynamic Range Imaging and Tone Mapping

Paul Debevec's SIGGRAPH Course

Dynamic Range in the Real World

Office interior

Indirect light from window

$1/60^{th}$ sec shutter

f/5.6 aperture

0 ND filters

0dB gain

Sony VX2000 video camera

Dynamic Range in the Real World

Outside in the shade

1/1000th sec shutter

f/5.6 aperture

0 ND filters

0dB gain

16 times the light as inside

Dynamic Range in the Real World

Outside in the sun

1/1000th sec shutter

f/11 aperture

0 ND filters

0dB gain

64 times the light as inside

Dynamic Range in the Real World

Straight at the sun

1/10,000th sec shutter

f/11 aperture

13 stops ND filters

0dB gain

5,000,000 times the light as inside

Dynamic Range in the Real World

Very dim room

1/4th sec shutter

f/1.6 aperture

0 stops ND filters

18dB gain

1/1500th the light than inside

Dynamic Range in the Real World

The real world is
high dynamic range.

The Problem of Dynamic Range

- Dynamic Range: Range of brightness values measurable with a camera

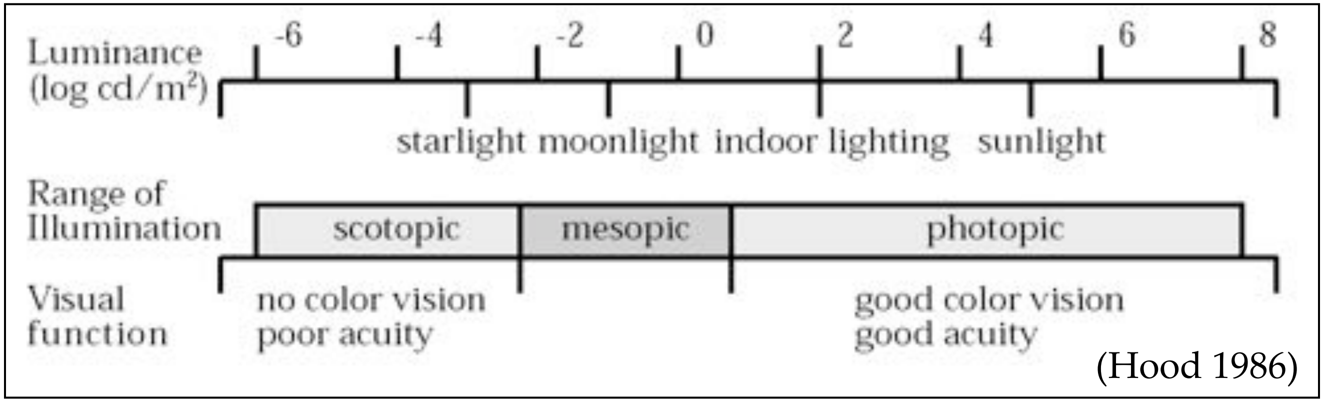

- Today's Cameras: Limited Dynamic Range

High Exposure Image

Low Exposure Image

- We need about 5-10 million values to store all brightnesses around us.
- But, typical 8-bit cameras provide only 256 values!!

High Dynamic Range Imaging

- Capture a lot of images with different exposure settings.
- Apply radiometric calibration to each camera.
- Combine the calibrated images (for example, using averaging weighted by exposures).

(Mitsunaga)

(Debevec)

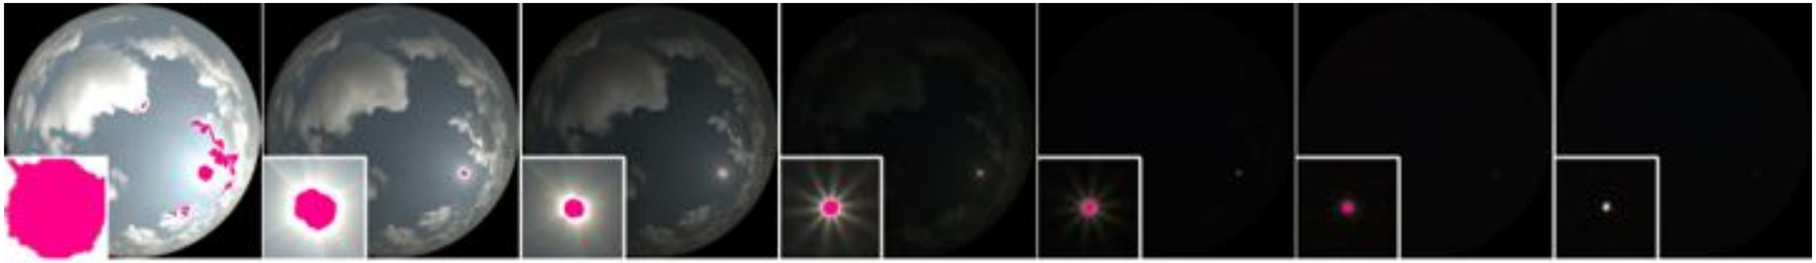

f/4, 1s f/4, $\frac{1}{4}$s f/4, $\frac{1}{32}$s f/16, $\frac{1}{16}$s f/16, $\frac{1}{125}$s f/16, $\frac{1}{1000}$s f/16, $\frac{1}{8000}$s

Images taken with a fish-eye lens of the sky show the wide range of brightnesses.

Relationship between Scene and Image Brightness

- Before light hits the image plane:

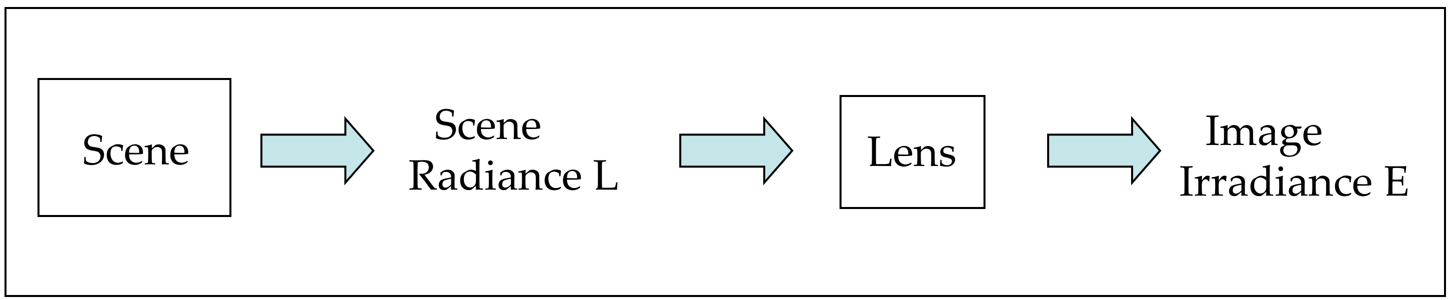

Linear Mapping!

- After light hits the image plane:

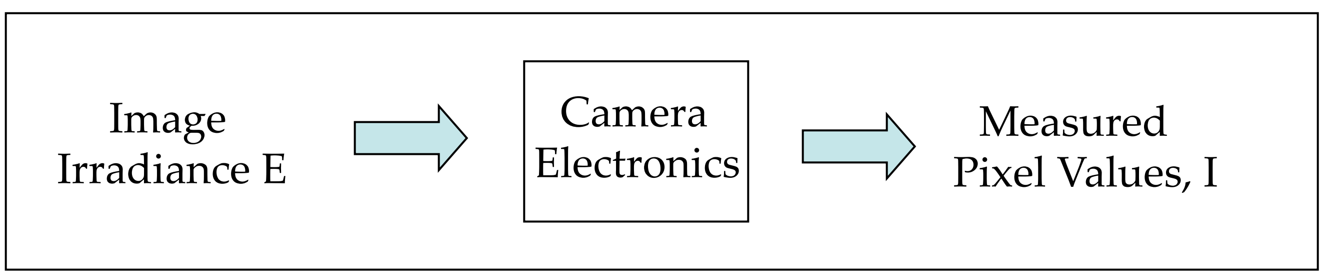

Non-linear Mapping!

Can we go from measured pixel value, I, to scene radiance, L?

Radiometric Calibration

- Important preprocessing step for many vision and graphics algorithms such as photometric stereo, invariants, de-weathering, inverse rendering, image based rendering, etc.

$$g^{-1} : I \rightarrow E$$

- Use a color chart with precisely known reflectances.

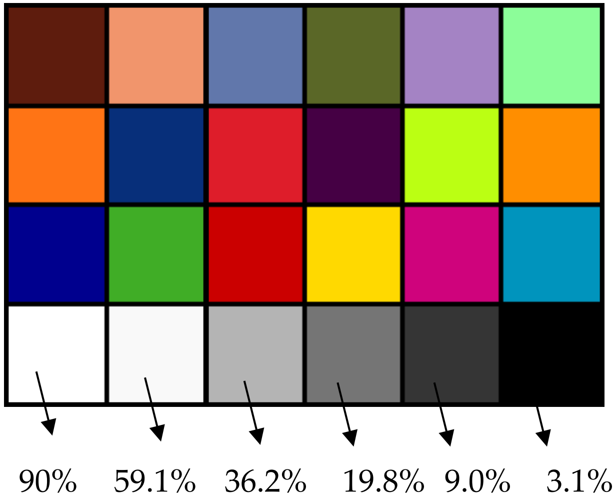

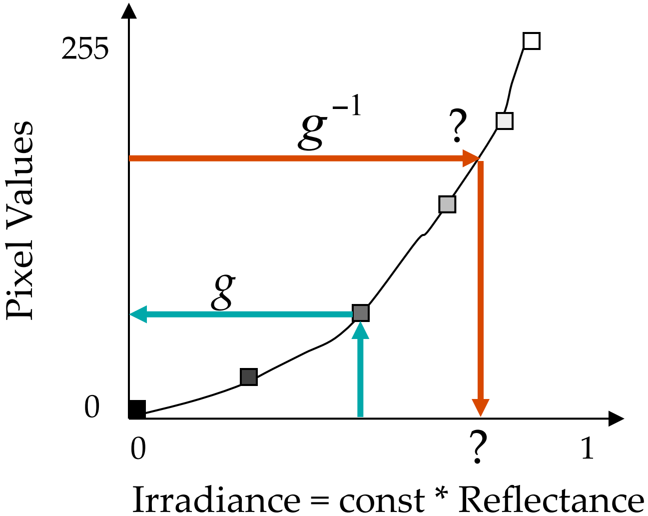

- Use more camera exposures to fill up the curve.
- Method assumes constant lighting on all patches and works best when source is far away (example sunlight).
- Unique inverse exists because g is monotonic and smooth for all cameras.

Ways to vary exposure

- Shutter Speed
- F/stop (aperture, iris)
- Neutral Density (ND) Filters
- Gain / ISO / Film Speed

RADIANCE Format

Greg Ward's "Real Pixels" format

32 bits / pixel

Red Green Blue Exponent

(145, 215, 87, 149) =

(145, 215, 87) * 2^(149-128) =

(1190000, 1760000, 713000)

(145, 215, 87, 103) =

(145, 215, 87) * 2^(103-128) =

(0.00000432, 0.00000641, 0.00000259)

Ward, Greg, "Real Pixels," in Graphics Gems IV, edited by James Arvo, Academic Press, 1994

8-bit Images (TIF, BMP, TGA, JPG, etc.)

- Useful for representing images to be output on a computer monitor or printer
- Less useful for representing images for film
- Inadequate for representing HDR images
- Usually nonlinearly encoded with a gamma curve, i.e.
- Amount of light = $(\text{pixel value})^{2.2}$

High-Dynamic Range Photography

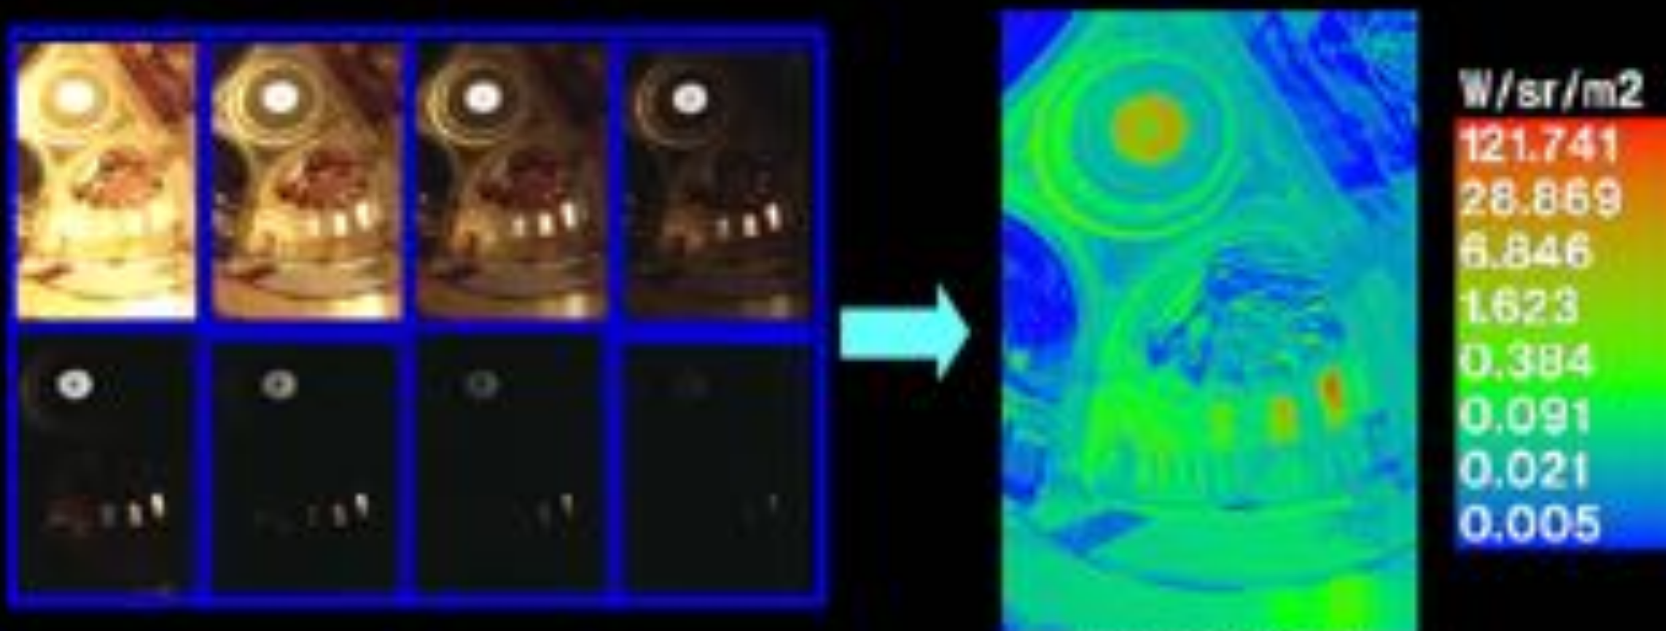

300,000 : 1

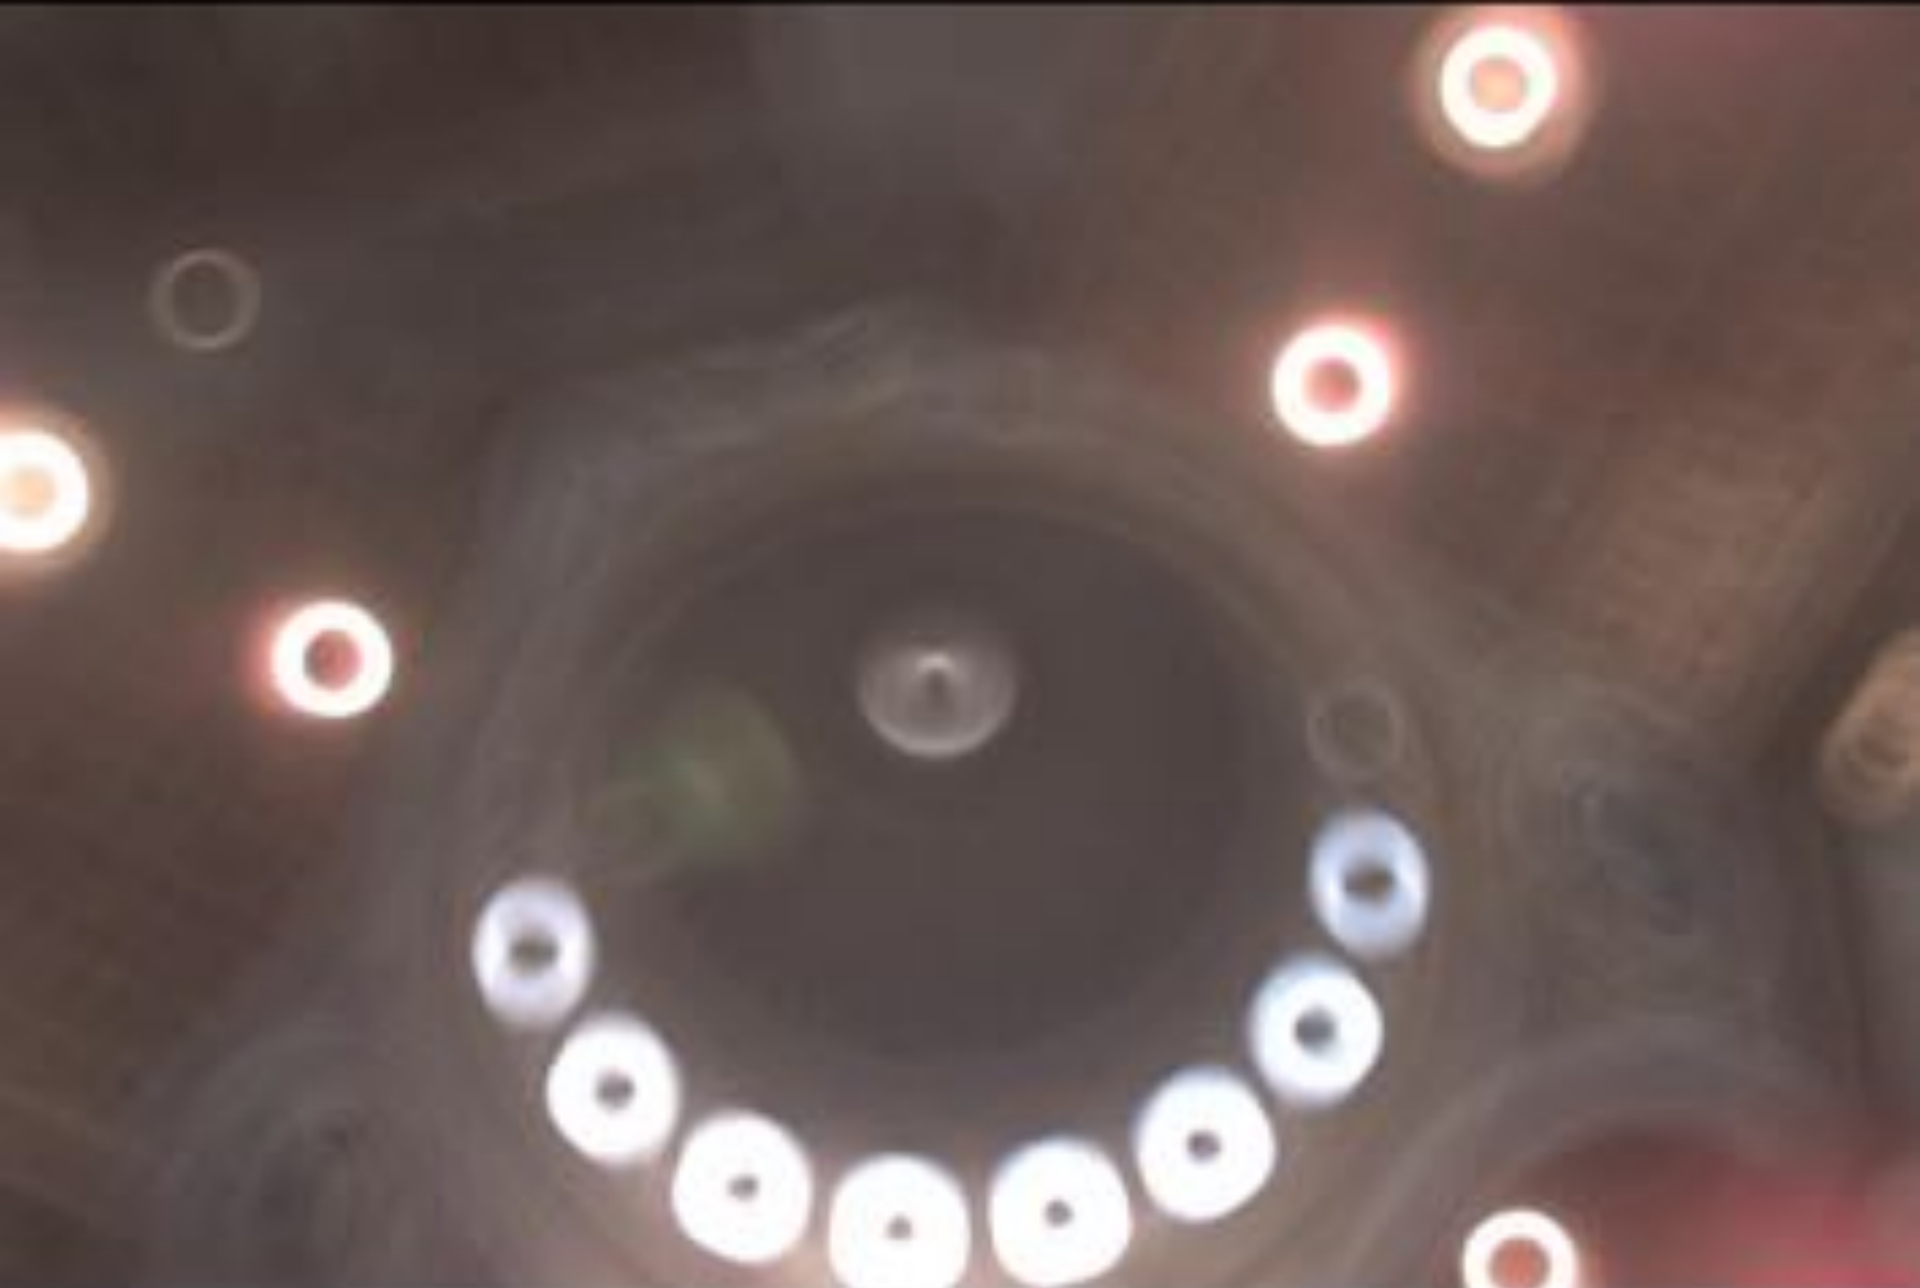

[Greg Ward]

HDR Tone-mapping

Linear tone-mapping

Non-linear tone-mapping

[Greg Ward]

Veiling Glare Simulation


H D R S h o p
High Dynamic Range Image Processing and Manipulation
www.debevec.org/HDRShop
Introduction | Tutorials | Reference | Plugins | FAQ | Download/Licensing | WWW Links | Mailing List


Capturing Real-World
Illumination

Mirrored Ball - Records light in all directions

Brightest regions are saturated

Intensity and color information lost

HDR Image of a Mirrored Ball

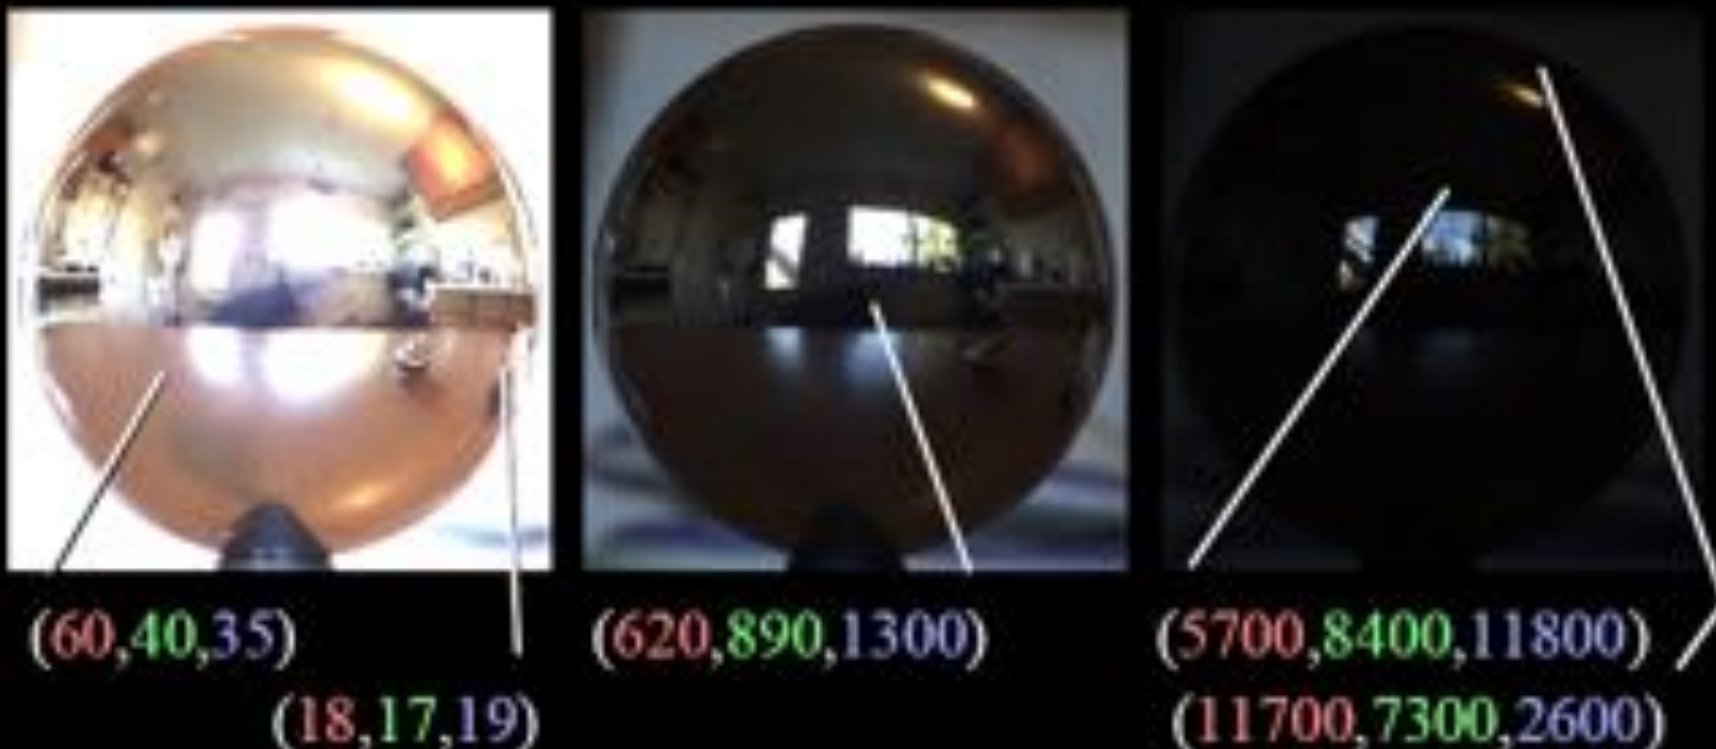

Assembled from ten digital images,

Δt = 1/4 to 1/10000 sec

Sources of Mirrored Balls

2-inch chrome balls < $20 ea.

King Bearing, Inc. / Applied Industrial Technologies

(many locations nationally, check www.bigbook.com)

6-12 inch large gazing balls

- Baker's Lawn Ornaments
  570 Berlin Plank Road
  Somerset, PA 15501-2413
  814-445-7028
- www.amazon.com - $8

Types of Omnidirectional Images

Latitude/Longitude

Cube Map

Illuminating Objects using Measurements of Real Light

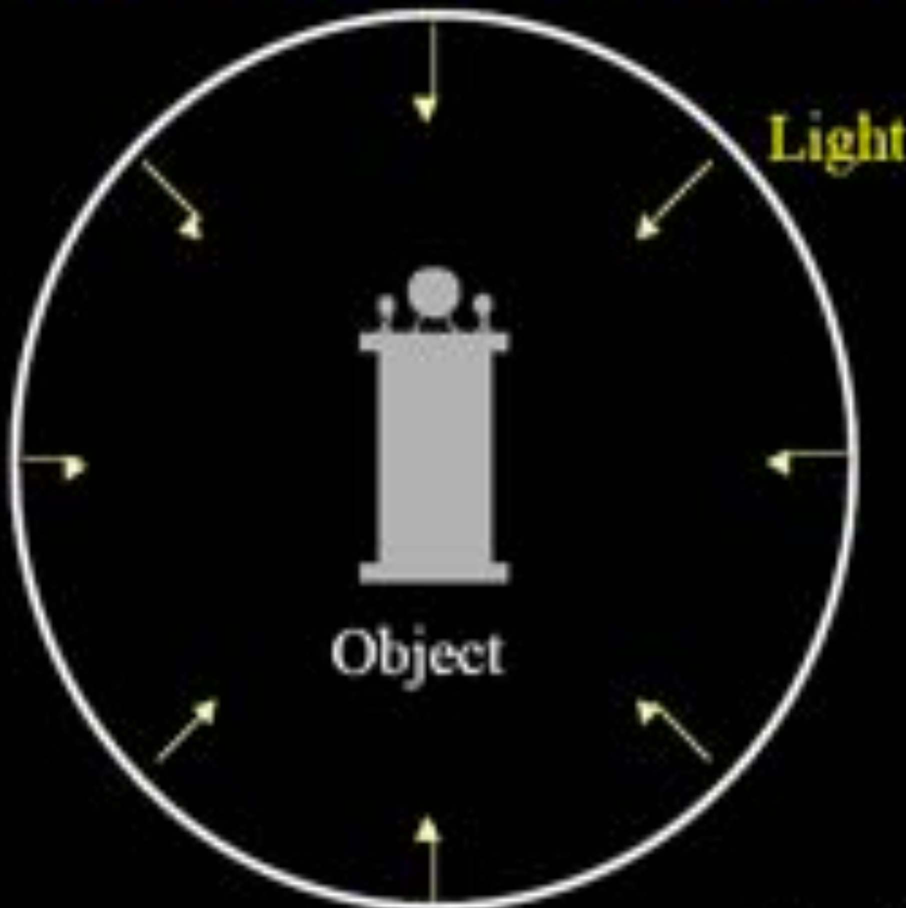

Environment assigned "glow" material property in Greg Larson's RADIANCE system.

http://radsite.lbl.gov/radiance/

See also: Larson and Shakespeare, "Rendering with Radiance", 1998


5. Illuminating Synthetic
Objects with Real Light

Comparison: Radiance map versus single image

Making

Rendering with Natural Light

SIGGRAPH 98 Electronic Theater

Light Probe Images

Eucalyptus Grove
UC Berkeley

Uffizi Gallery
Florence

St. Peter's Basilica
Rome

Grace Cathedral
San Francisco

Light Probe Image Gallery:
www.debevec.org/Probes

Photo Clip Art

Jean-François Lalonde Derek Hoiem Alexei A. Efros
Carnegie Mellon University

Carsten Rother John Winn Antonio Criminisi
Microsoft Research Cambridge

Figure 1: Starting with a present day photograph of the famous Abbey Road in London (left), a person using our system was easily able to make the scene much more lively. There are 4 extra objects in the middle image, and 17 extra in the right image. Can you spot them all?

http://graphics.cs.cmu.edu/projects/photoclipart/

http://vision.gel.ulaval.ca/~jflalonde/#research