Sampling and pixels

CS 178, Spring 2009

Begun on Tuesday, 4/14/09, finished on Thursday, 4/16. Additional slide (#24) added on 4/21.

Marc Levoy
Computer Science Department
Stanford University

Announcements (from whiteboard)

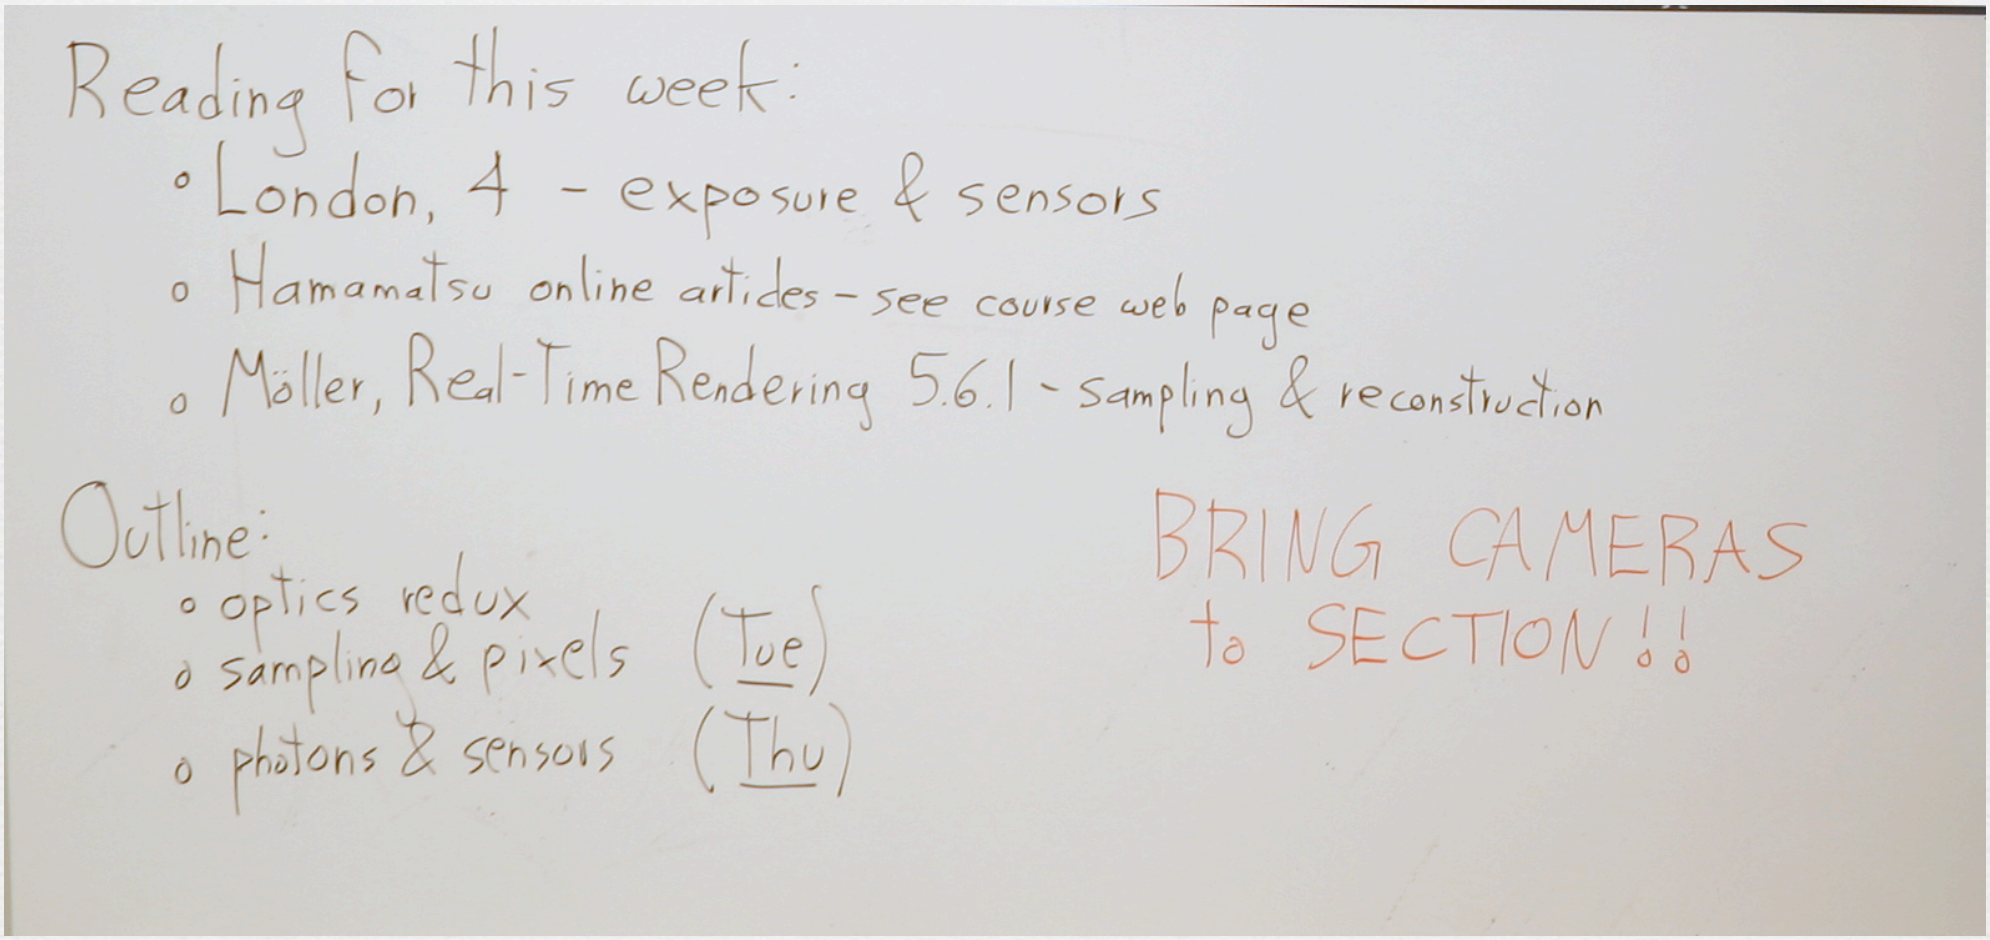

Outline

- image output technologies
  - displays
  - printing
- human perception
  - human spatial sensitivity
    - How many pixels?
  - human intensity discrimination
    - How many bits per pixel?
- sampling and reconstruction
  - including aliasing and antialiasing

I skipped human intensity discrimination. I'll cover it when I talk about exposure metering, noise, and ISO.

Liquid crystal displays

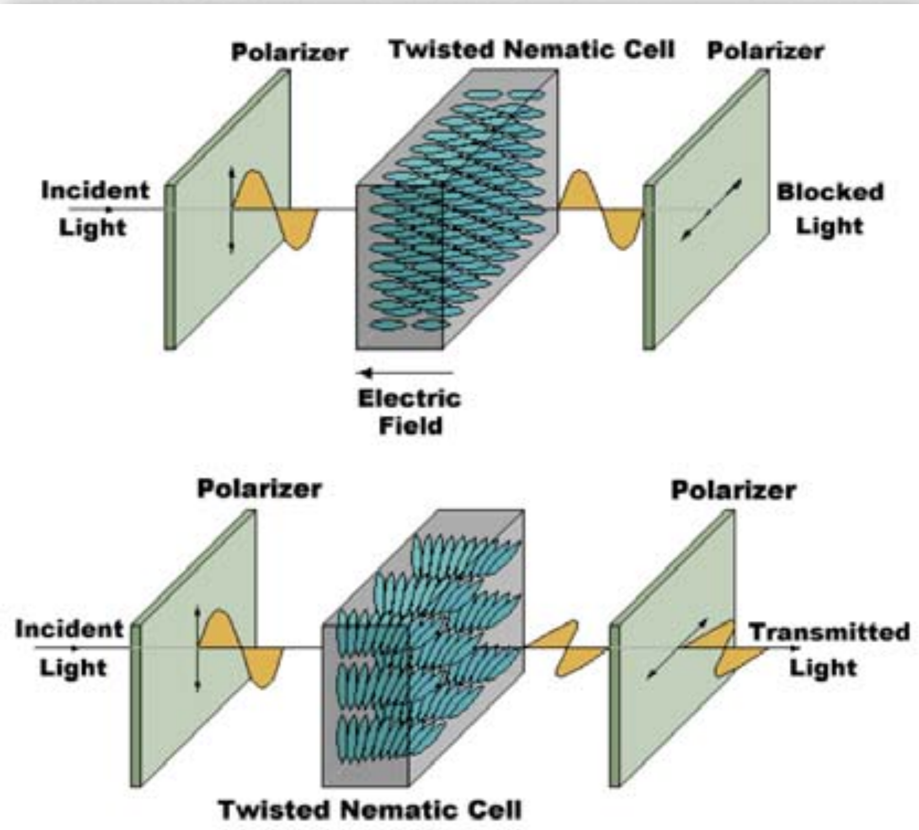

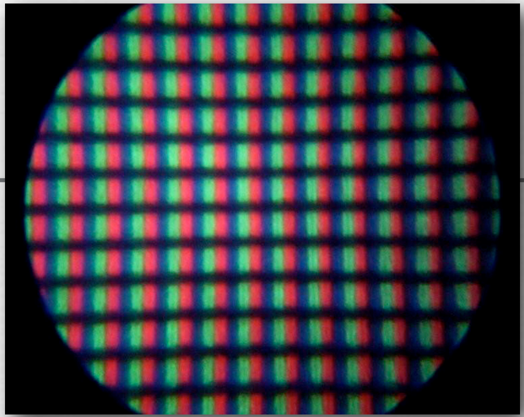

(Katie Dektar)

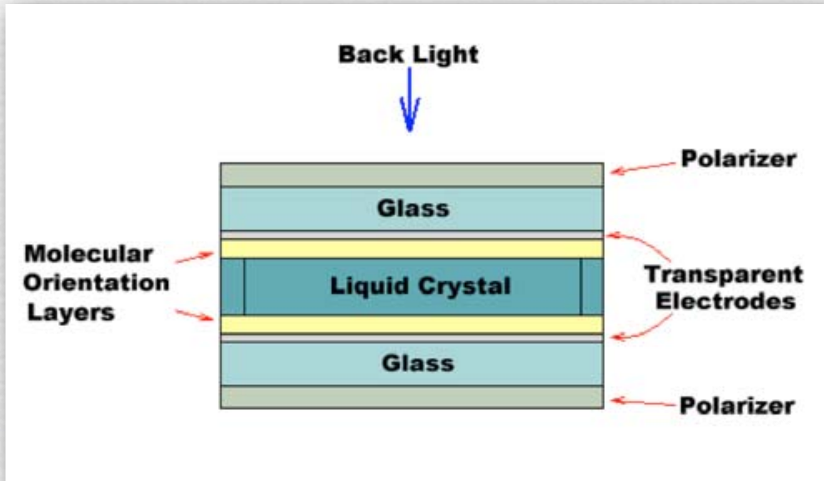

- cell on (top): LC aligned, light blocked → black
- cell off (bottom): LC twists light 90°, light transmitted → white
- color produced using a superimposed array of R,G,B filters

Dynamic micro-mirror device (DMD)

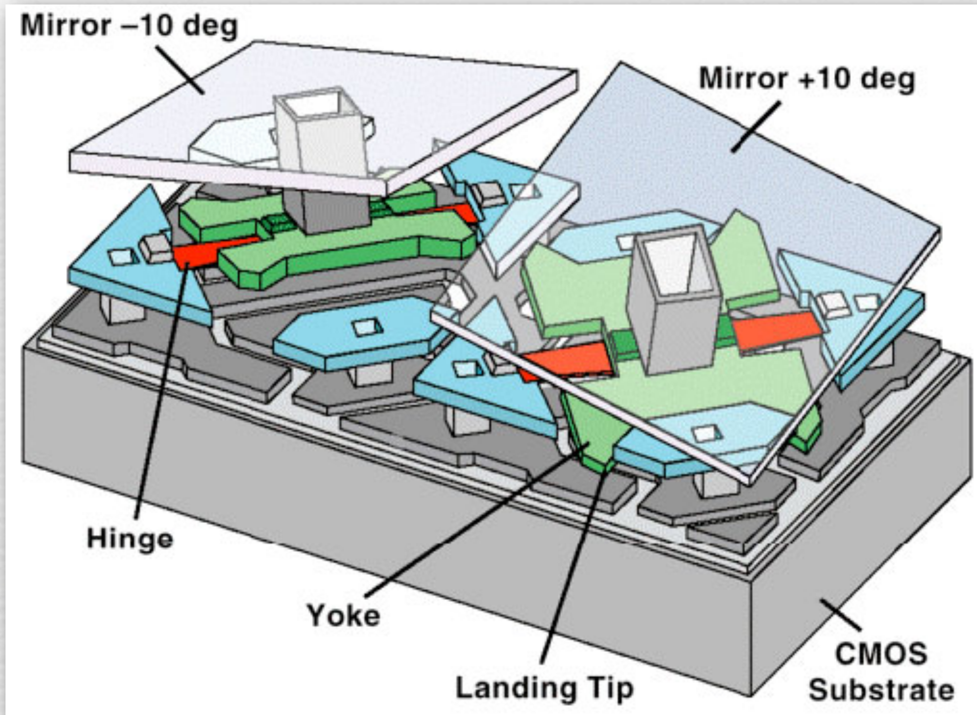

<u>Other acronyms:</u>

DLP = digital light processing

SLM = spatial light modulator

- each mirror is 7-16μ on a side
  - created using MEMS technology
- mirrors vibrate at KHz rates
  - fraction of time spent facing optics determines gray level
- color produced using spinning color wheel or LEDs

Organic LEDs (OLEDs)

(http://www.artlebedev.com/everything/optimus)

- requires no backlight
- low power, thin, flexible (bendable)
- lousy color quality

Electronic ink

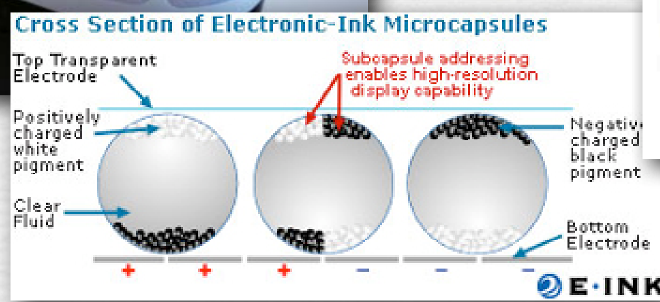

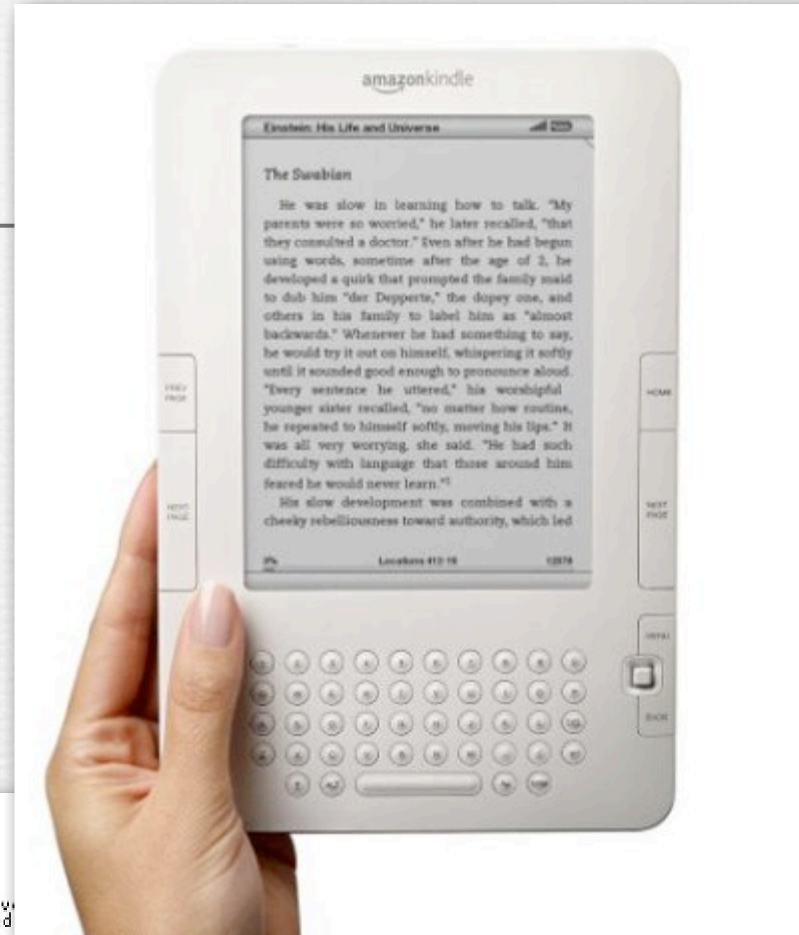

Amazon's Kindle 2

- black (-) & white (+) particles in a clear fluid
  - position in fluid controlled by a transparent grid of electrodes
- flexible, requires no power except when changing images
- changing images is slow

Other image output technologies

- display
  - CRT = cathode ray tube (obsolete)
  - LCOS = liquid crystal on silicon - like LCD but reflective
  - plasma panel - emissive
- printing
  - dye sublimation
  - inkjet
  - offset

Human field of view

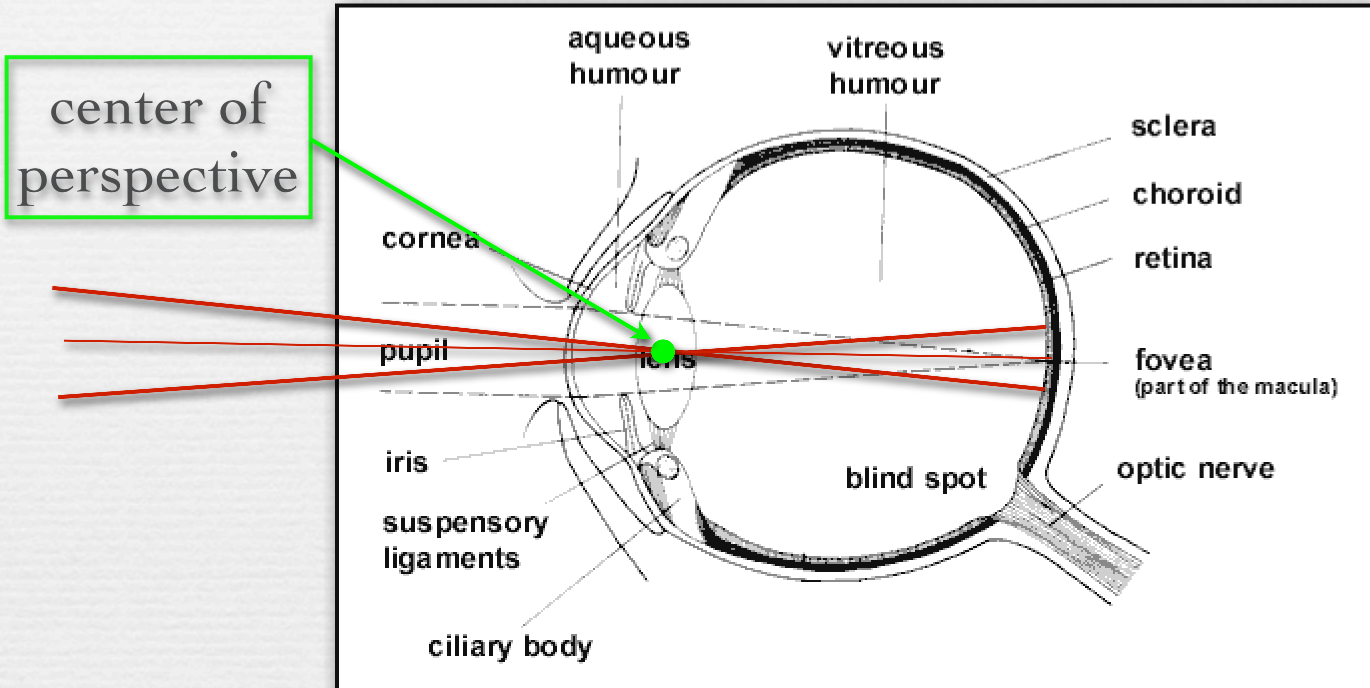

(schools.net.au)

- center of ray bundle is nominal line of sight
- a range of such lines constitute a field of view
  - the proper measure (from “natural perspective”) is retinal arc
  - maximum human FOV is 180°, with 140° of binocular overlap
  - human fovea (area of highest acuity) is 1°

Human retinal arc

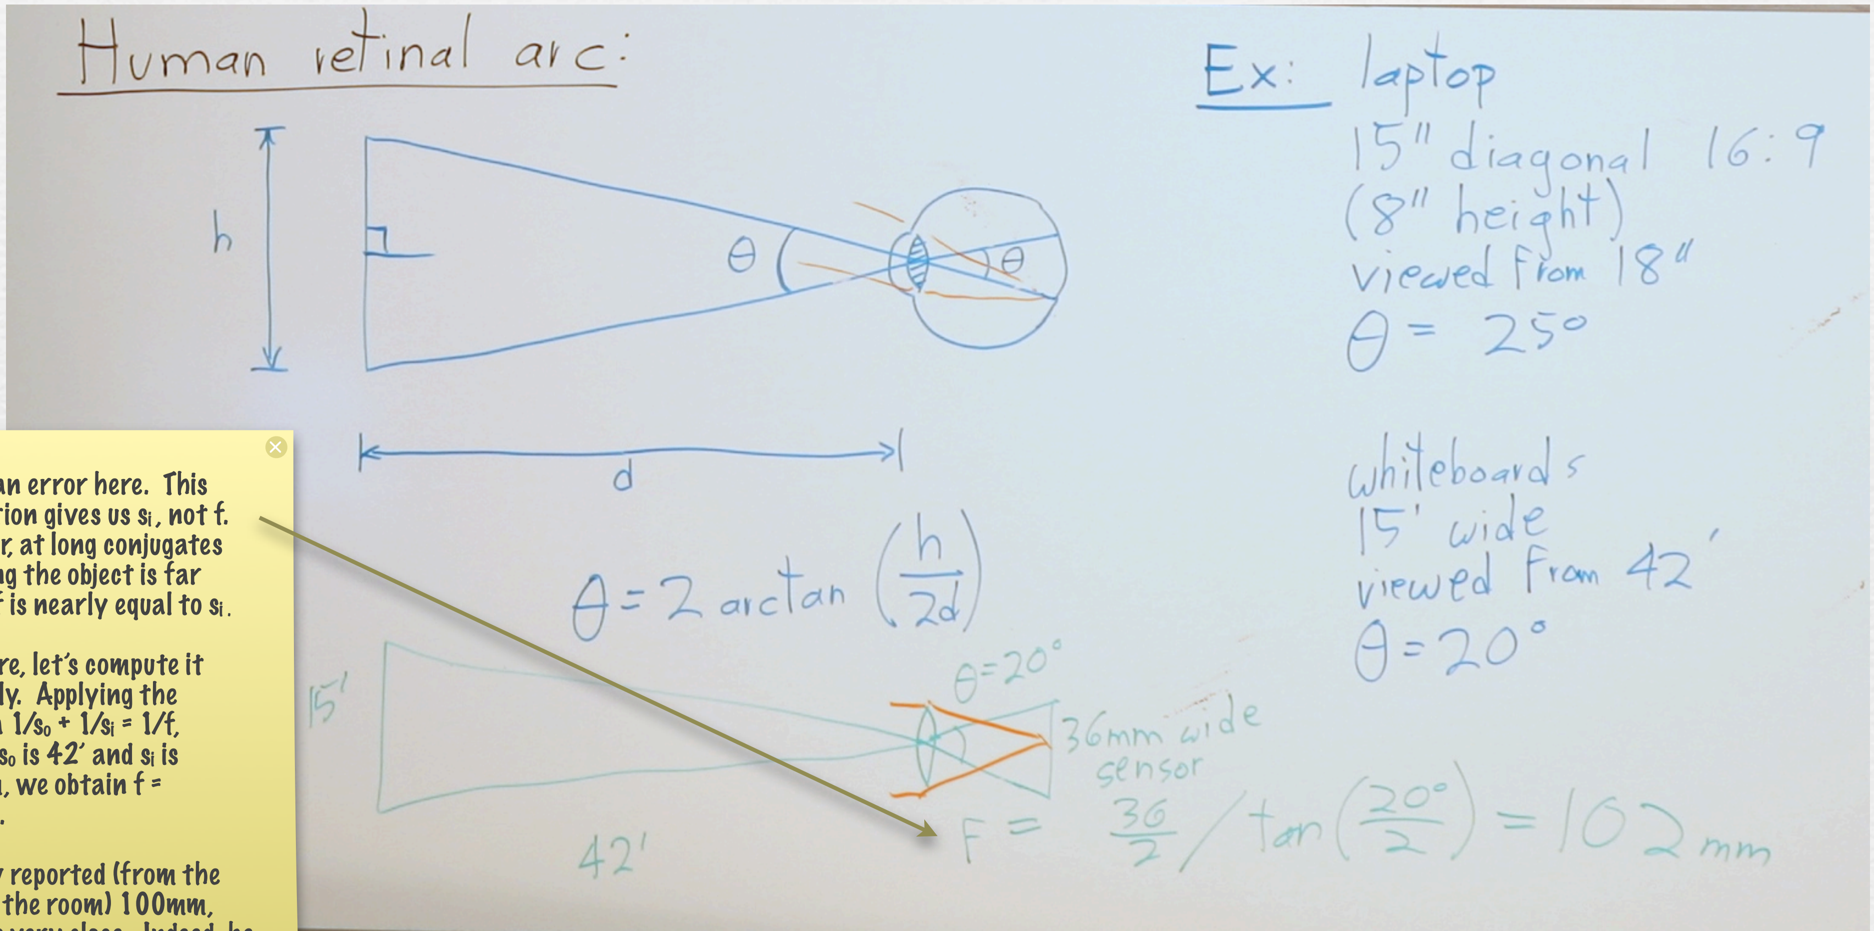

I made an error here. This calculation gives us s_i, not f. However, at long conjugates (meaning the object is far away), f is nearly equal to s_i.

To be sure, let's compute it correctly. Applying the formula $1/s_o + 1/s_i = 1/f$, where s_o is 42' and s_i is 102mm, we obtain f = 101mm.

Andrew reported (from the back of the room) 100mm, which is very close. Indeed, he was deliberately set a bit wide to make sure he didn't clip the whiteboards.

Spatial frequency of a medium

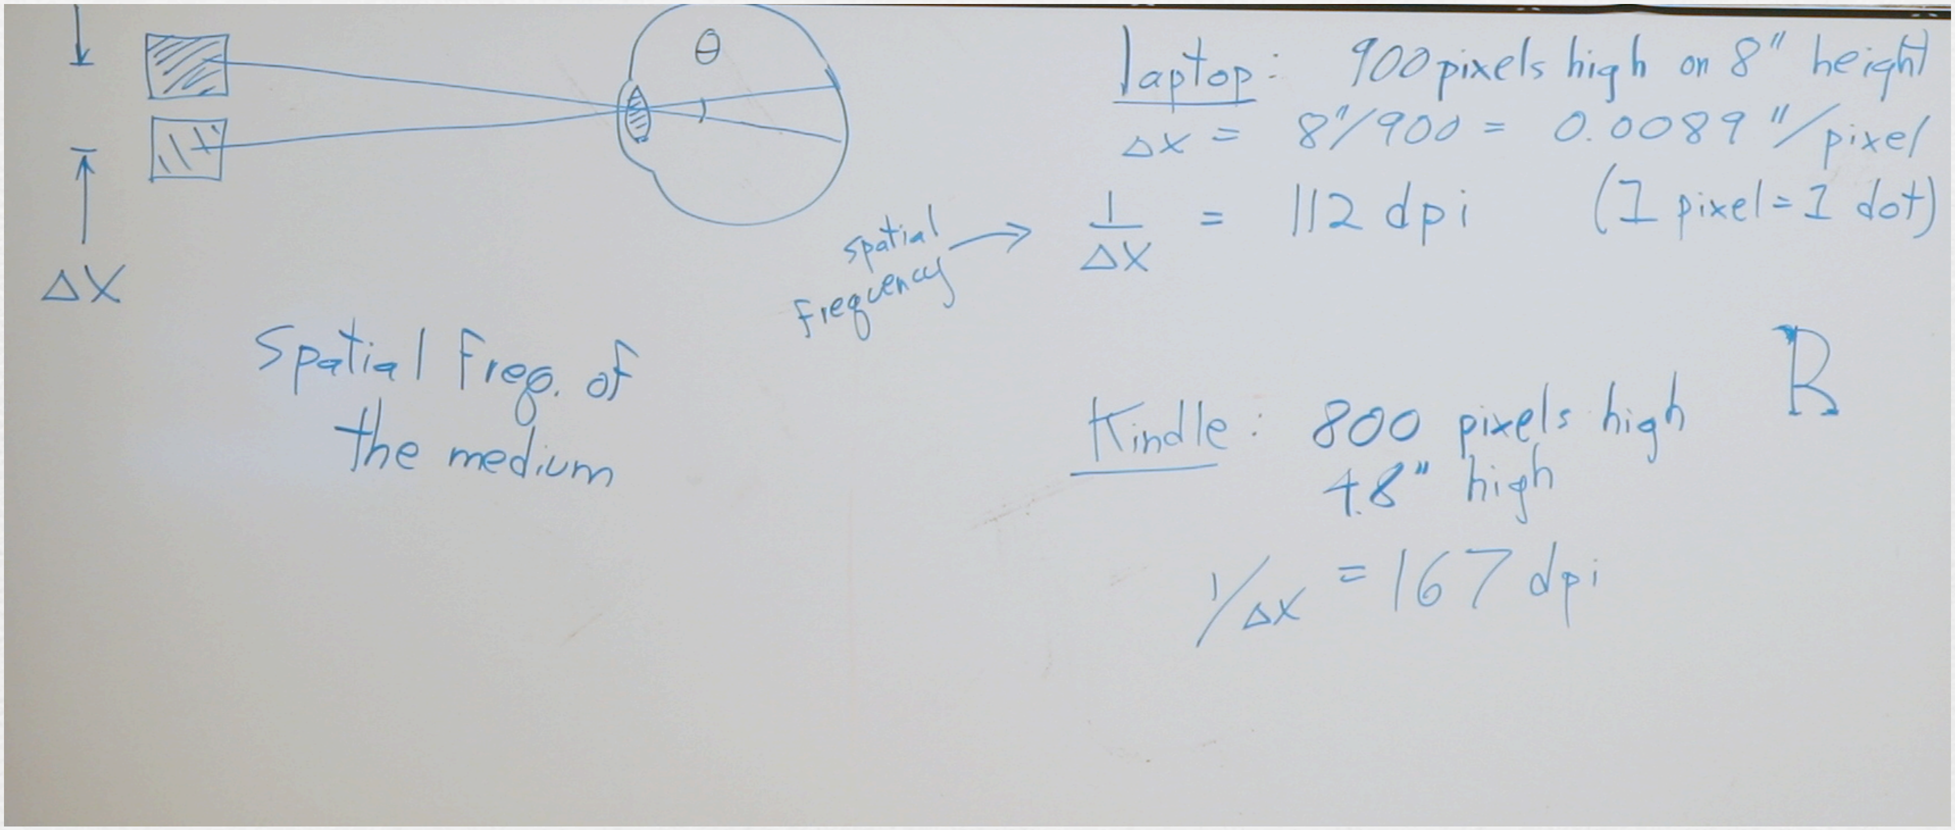

Spatial frequency on the retina

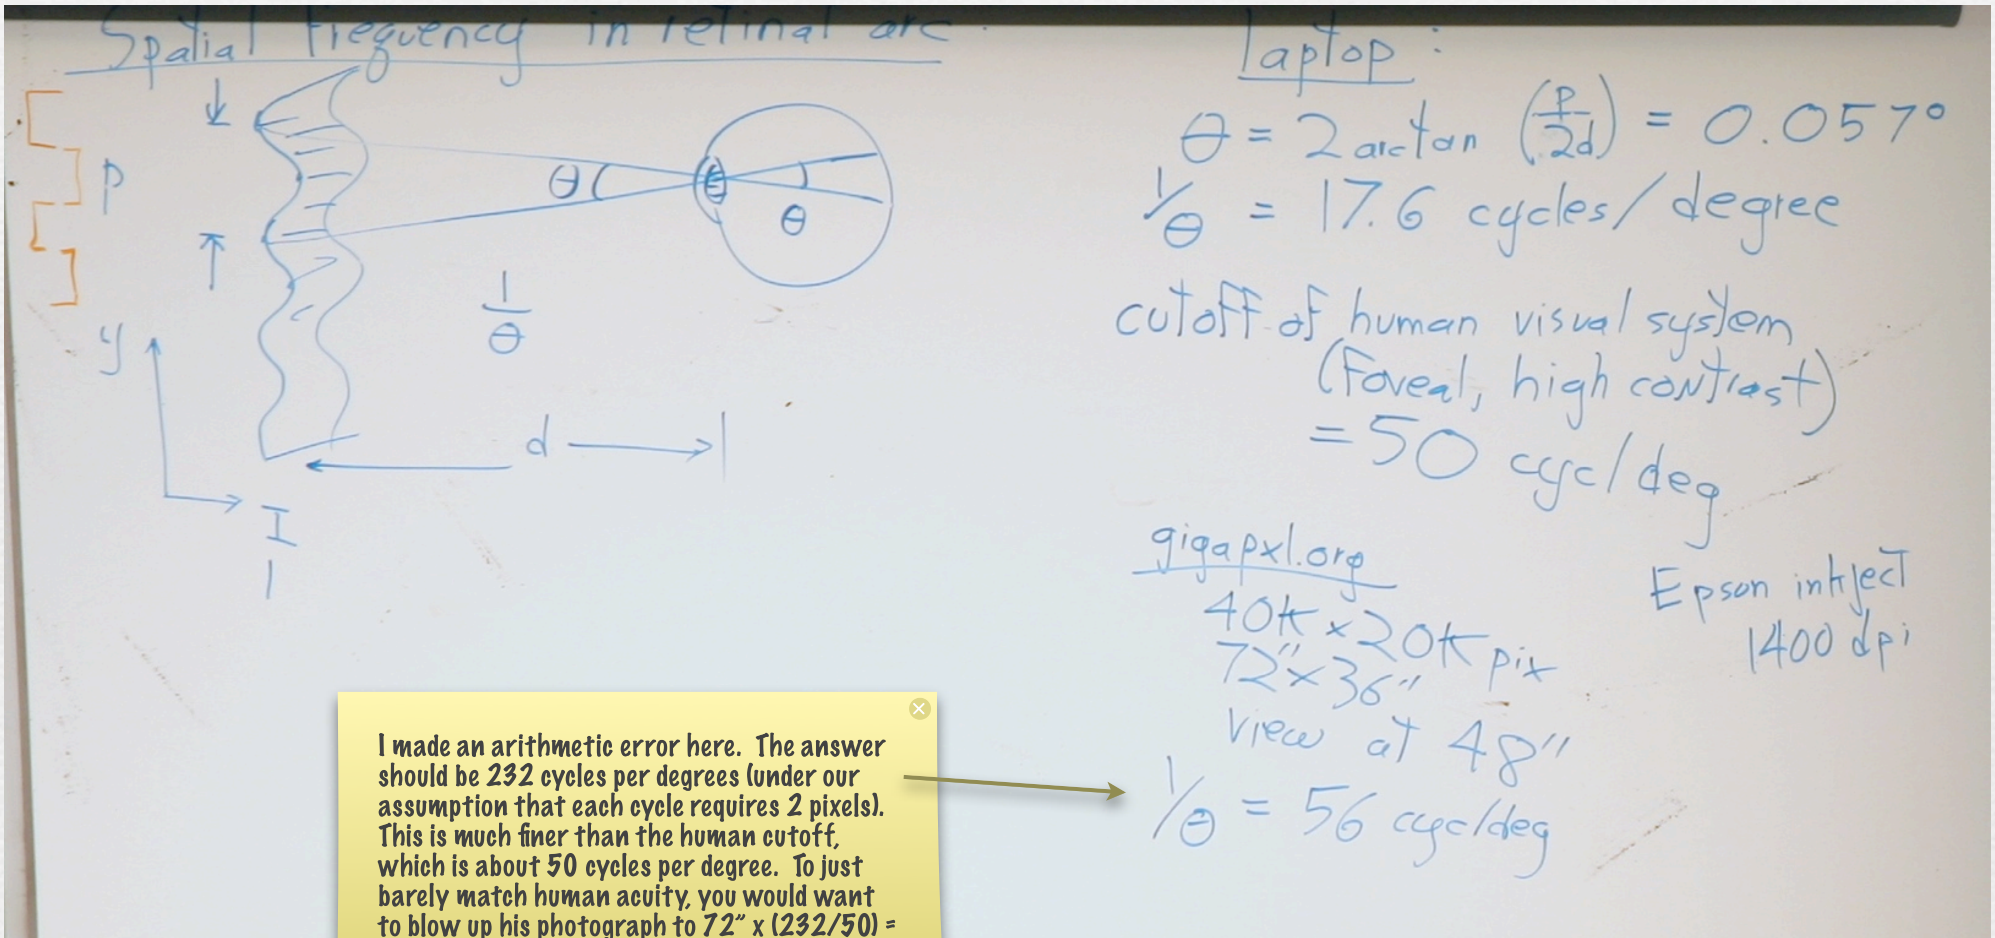

I made an arithmetic error here. The answer should be 232 cycles per degrees (under our assumption that each cycle requires 2 pixels). This is much finer than the human cutoff, which is about 50 cycles per degree. To just barely match human acuity, you would want to blow up his photograph to 72" x (232/50) = 27 feet wide! (assuming you still view it from 4' away)

Human spatial sensitivity

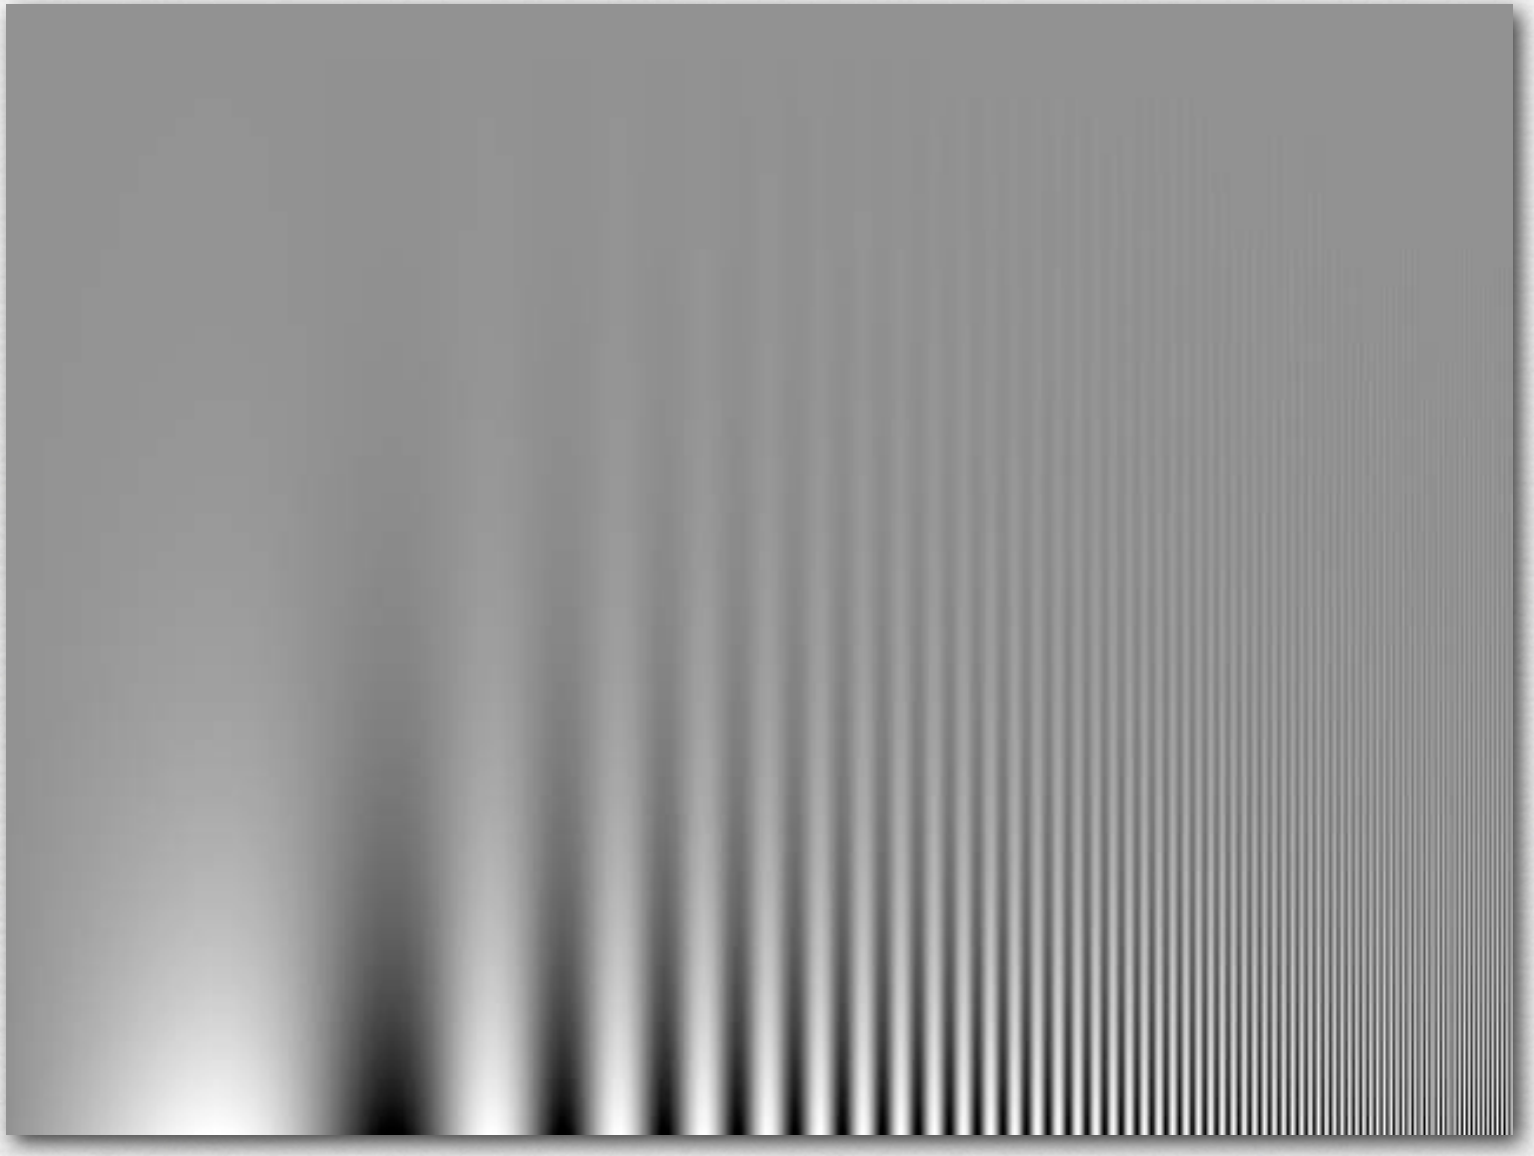

Human spatial sensitivity

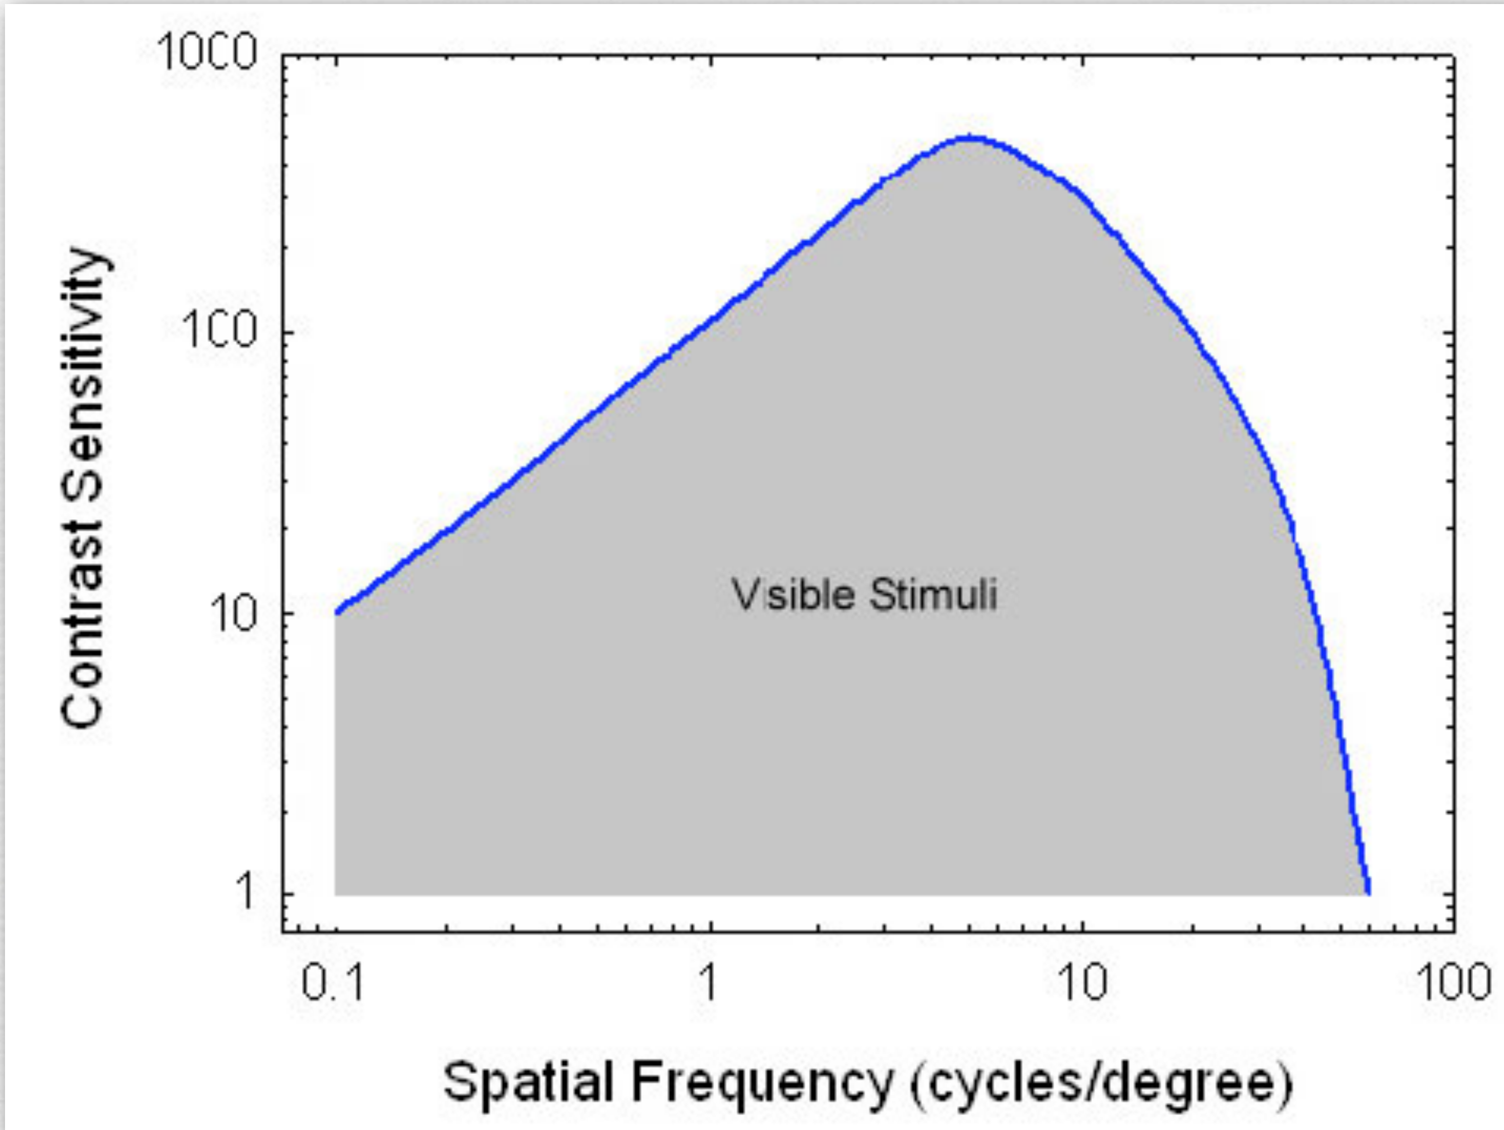

(Graham Flint)

Balboa Park, San Diego

(original is 40K × 20K pixels, Gates Hall print is *72" × 36"*)

The human retina

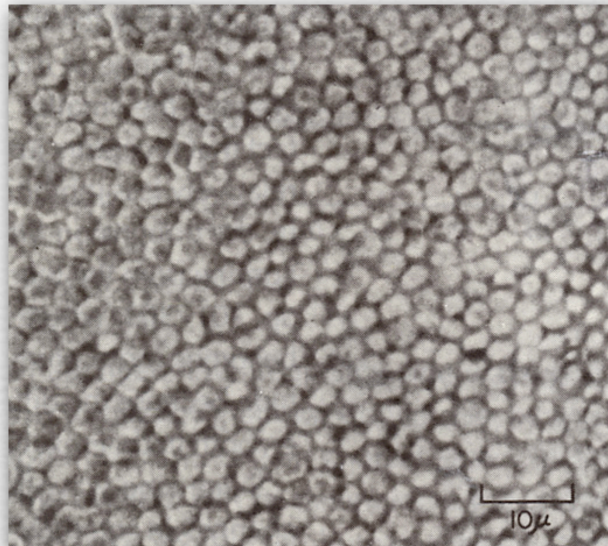

(Cornsweet)

- retina contains 100M rods and 5M cones (3 types: L,M,S)
  - size of cones is 1μ (foveal), 10μ (peripheral)
  - expressed as retinal arc: 30 arcsecs (L,M), 600 arcsecs (S)
- eyes are diffraction limited
  - pupil is 4mm
  - circle of confusion on retina is 6μ (bigger than a foveal cone)

Vernier acuity

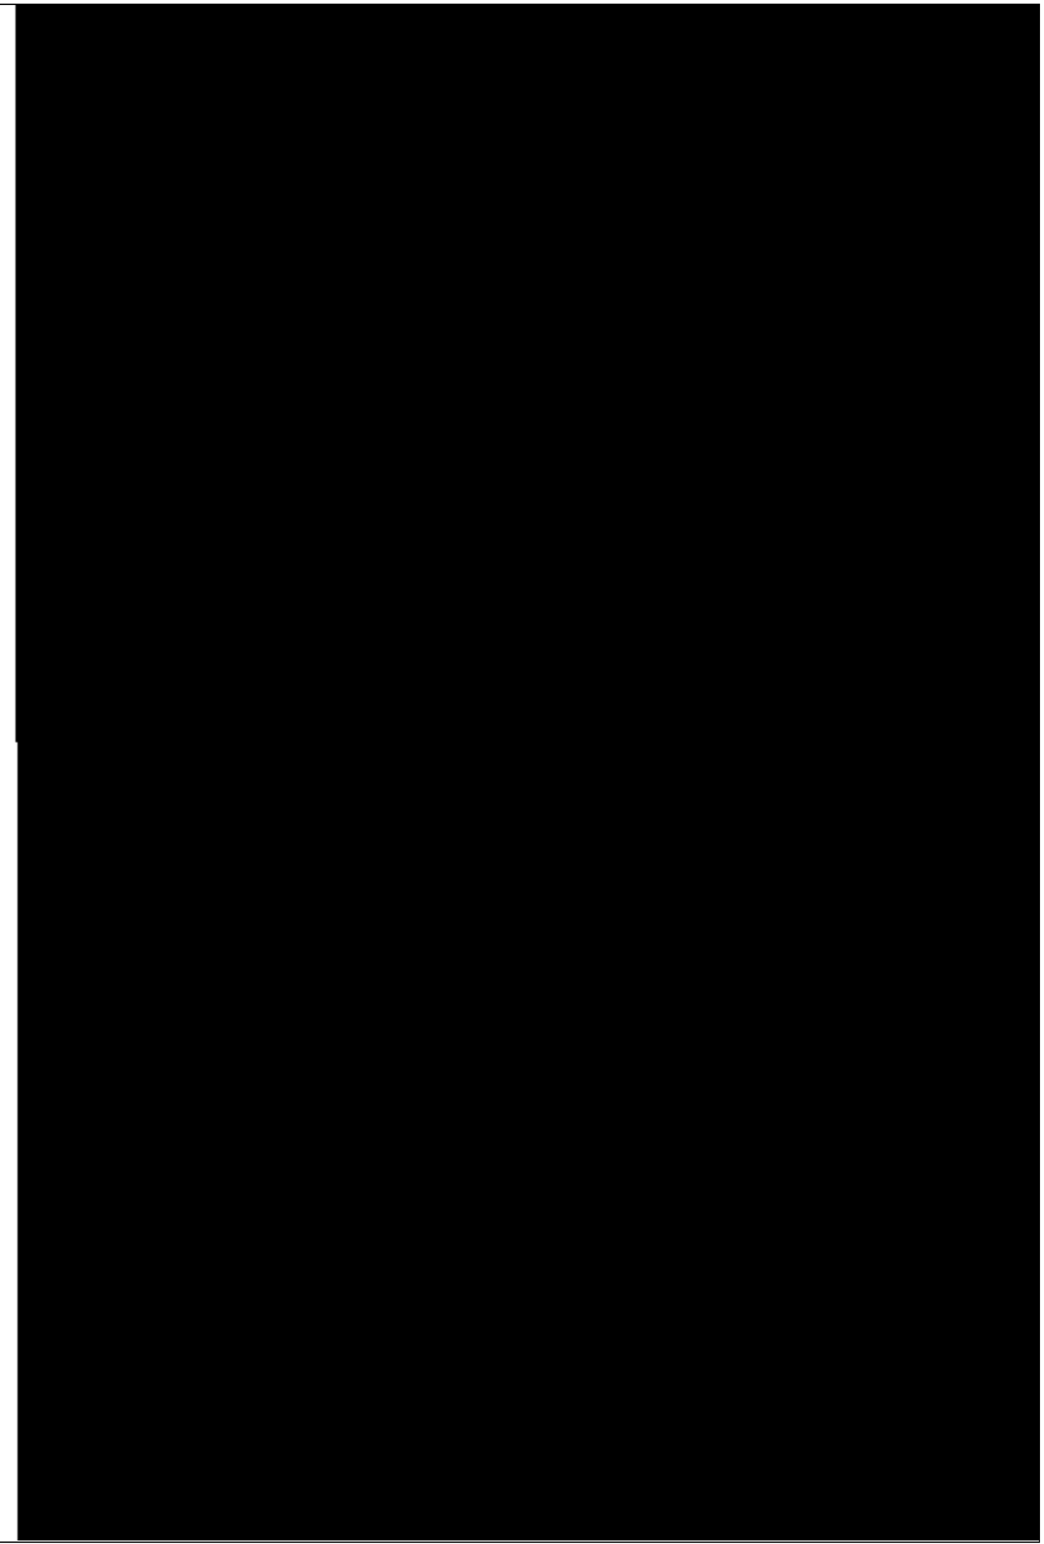

Vernier acuity

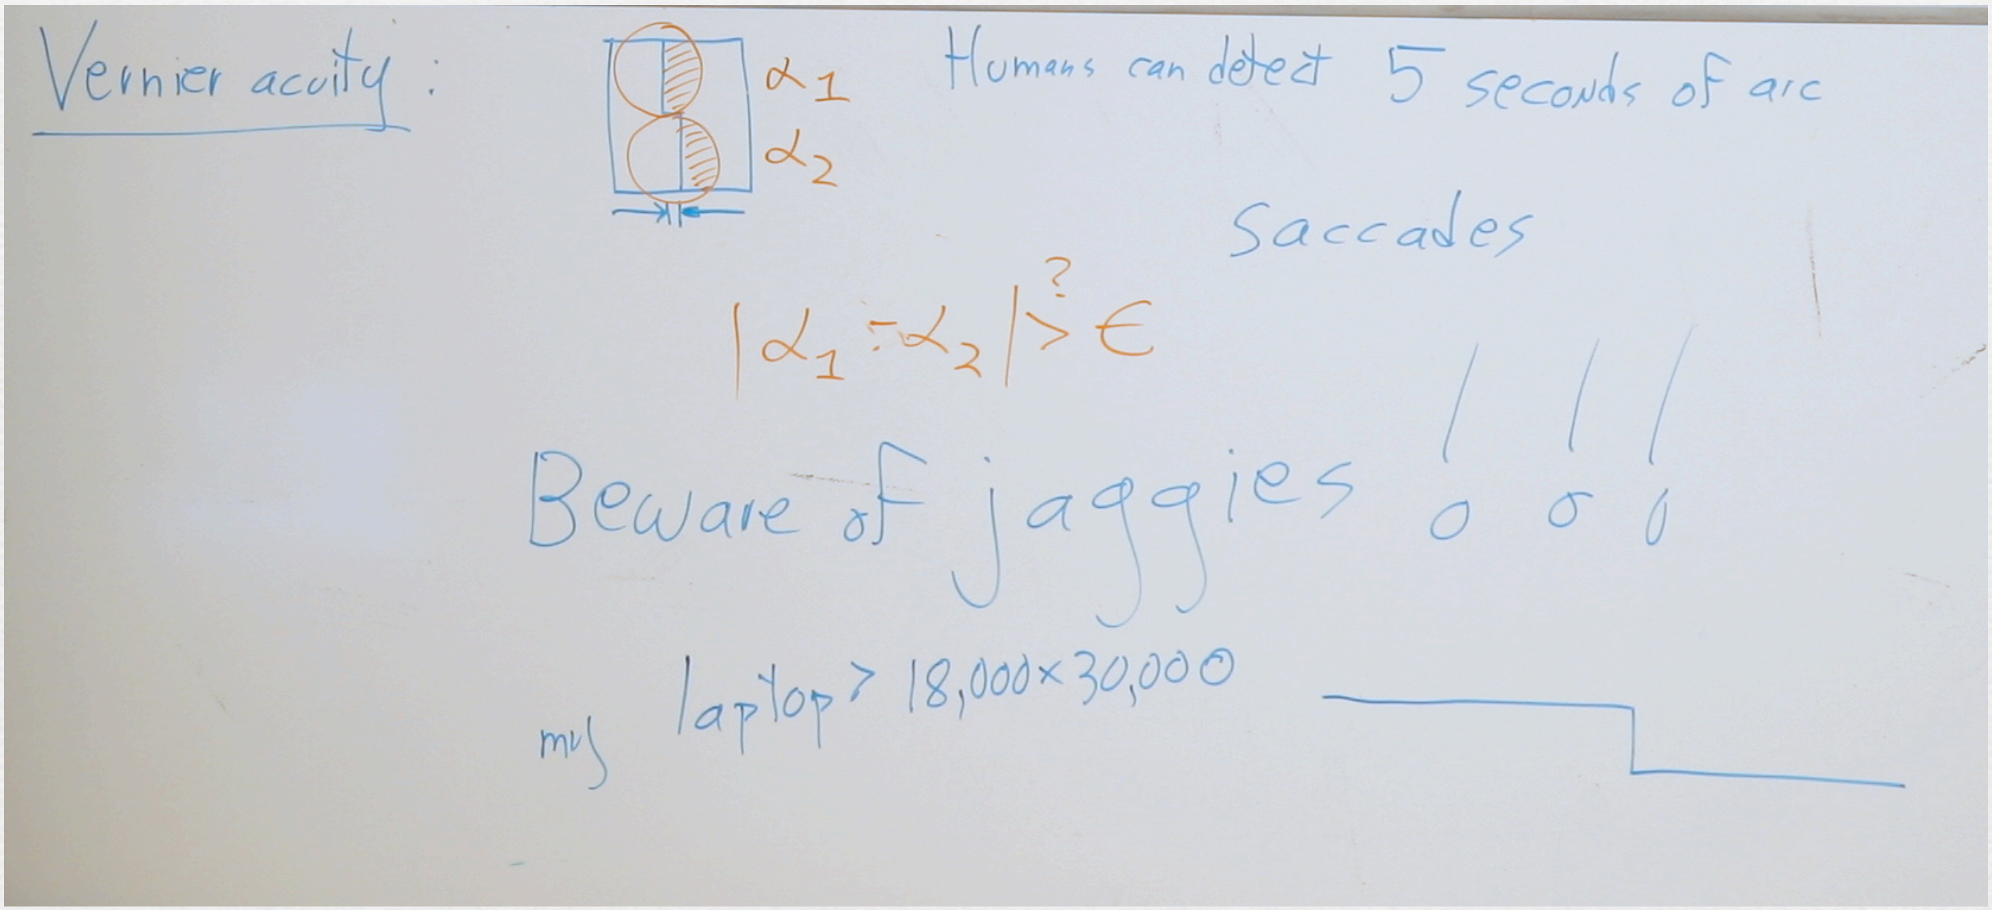

Sampling and reconstruction

- the pipeline

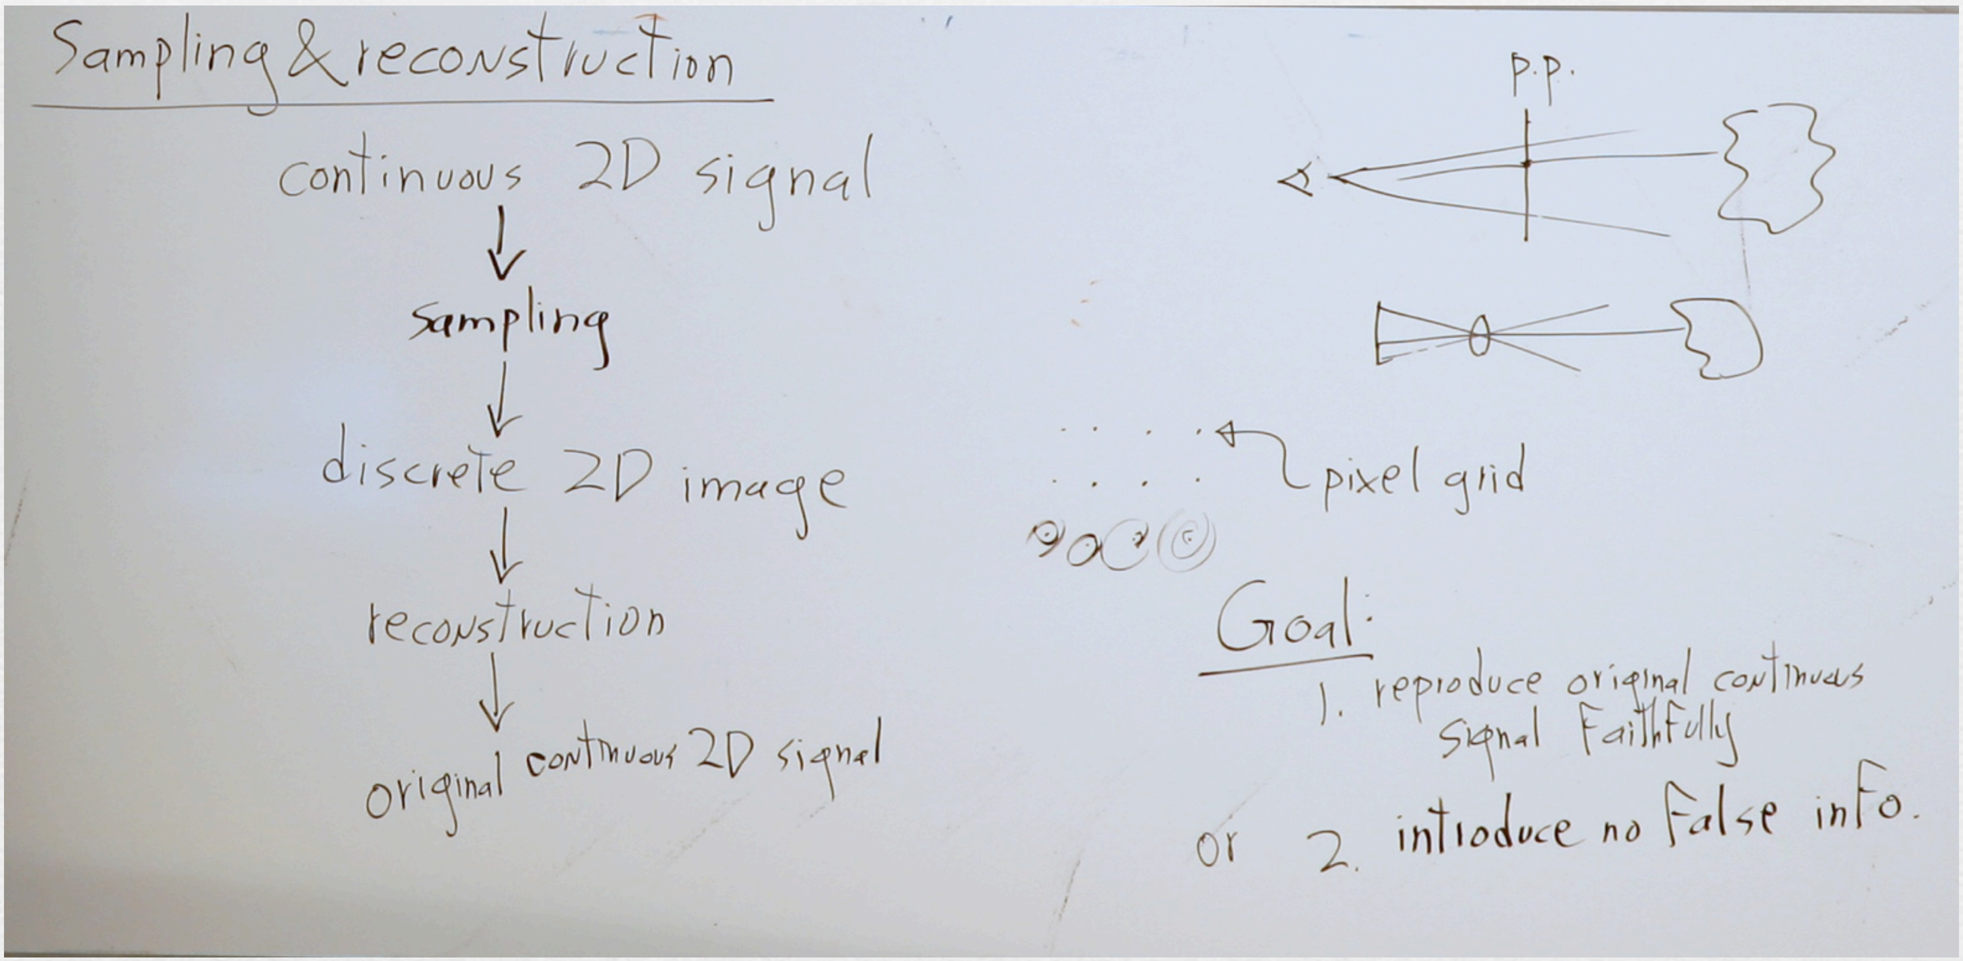

Sampling and reconstruction

- insufficiently fine sampling, producing *aliasing* (high frequencies masquerading as low frequencies)

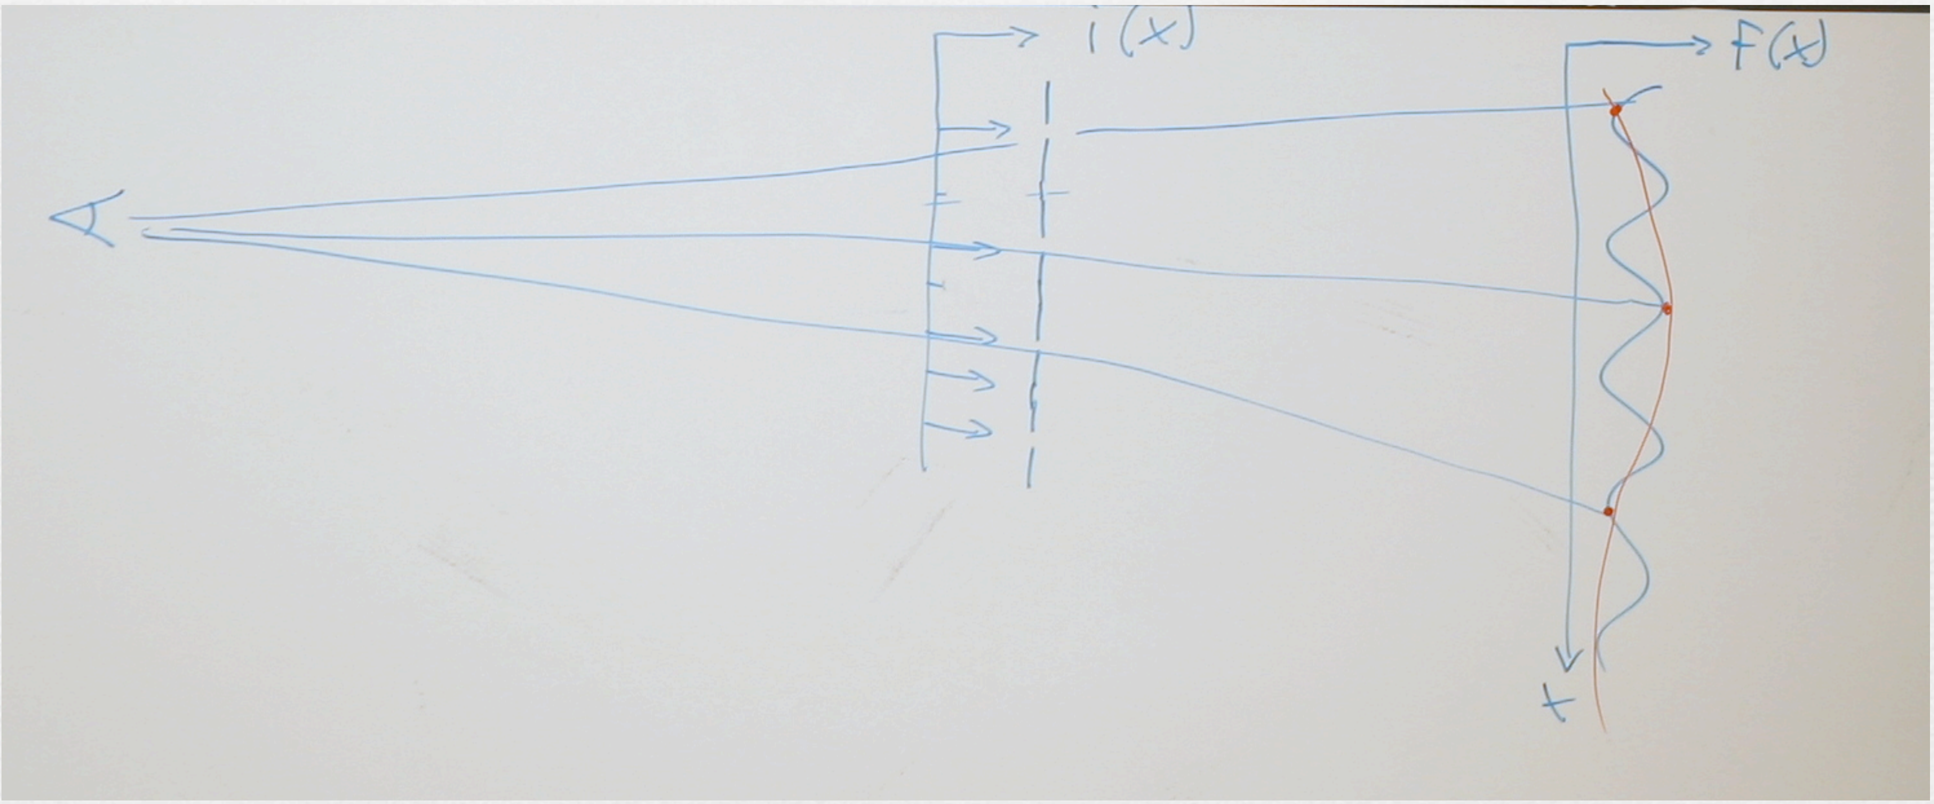

Sampling and reconstruction

- solution #1: raise the sampling rate

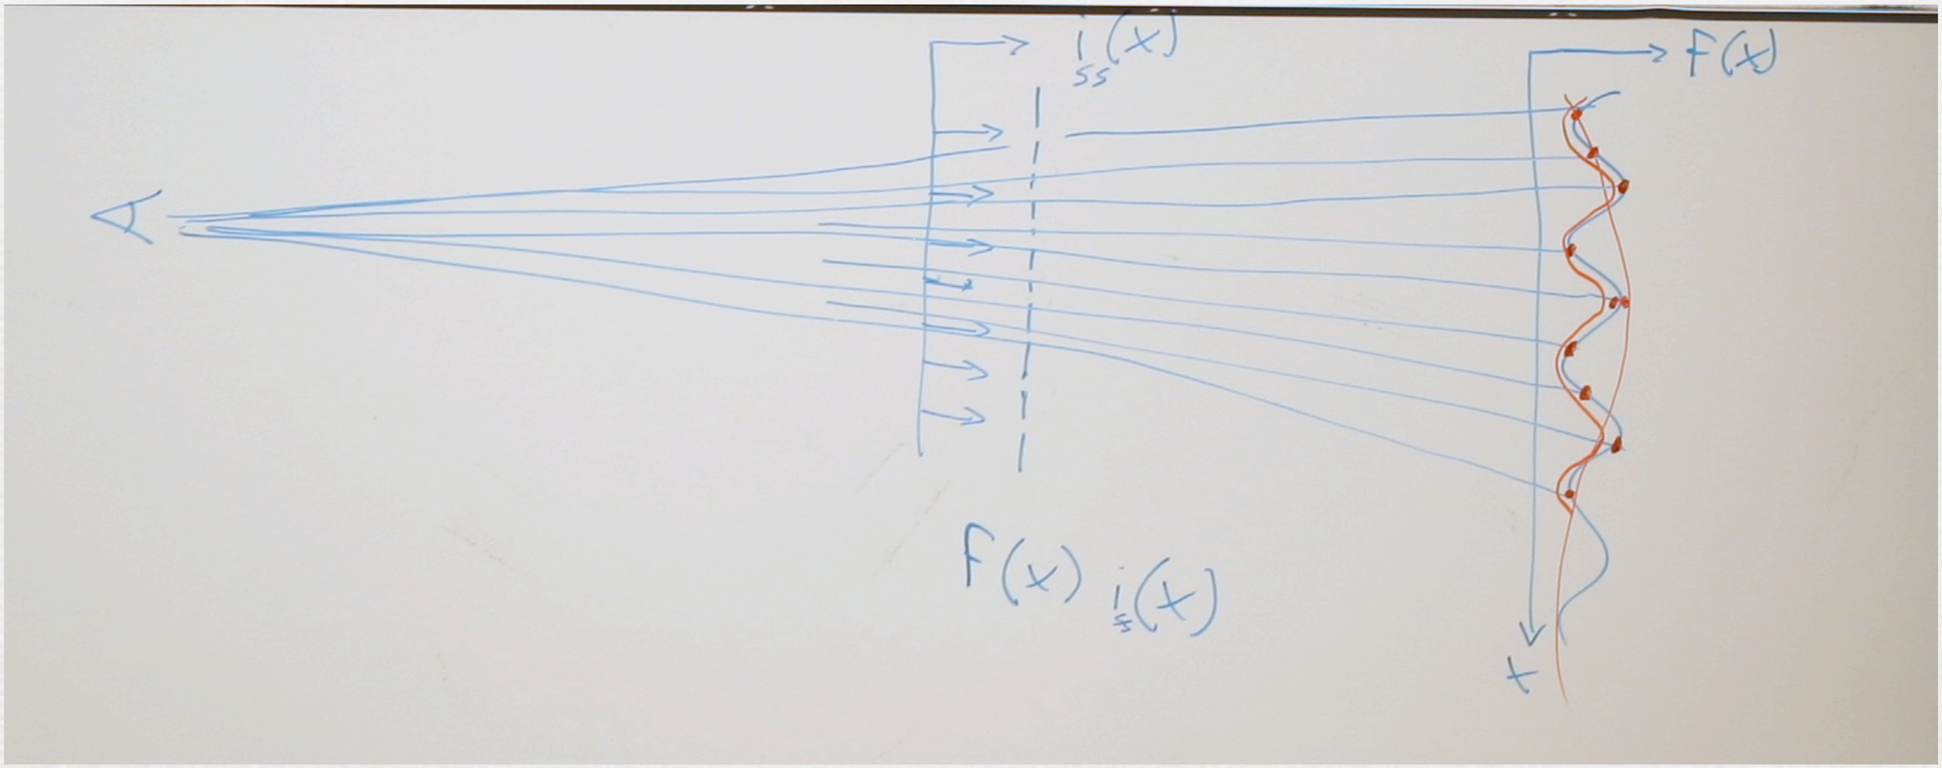

Sampling and reconstruction

- solution #2: filter the signal <u>before</u> sampling, a.k.a. *pre-filtering*, performed using a spatial convolution

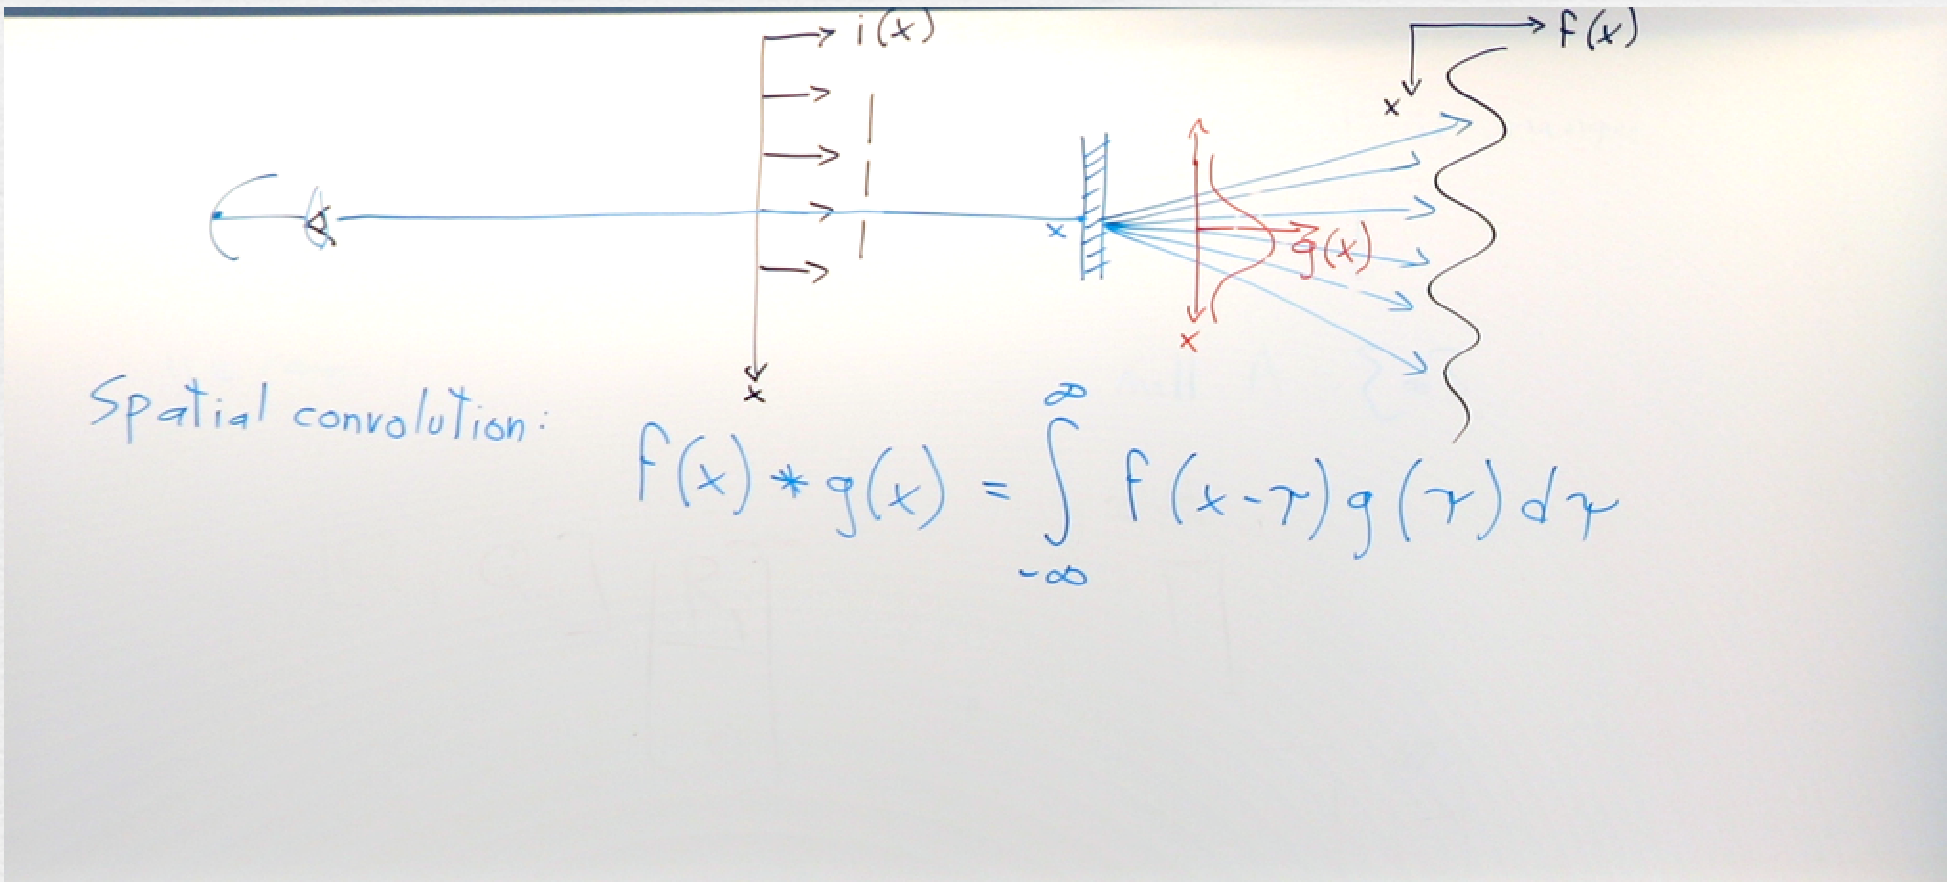

A graphical construction for how convolution works

- on the left is continuous convolution (an integral); on the right is discrete convolution (a sum)
  - a blur filter in Photoshop performs a discrete convolution
- at bottom right is the Rect function - a.k.a. box filter
  - a square pixel in a digital camera performs a (continuous) convolution of the focused image with a (2D) Rect function, thereby pre-filtering it to avoid aliasing when converted to pixels

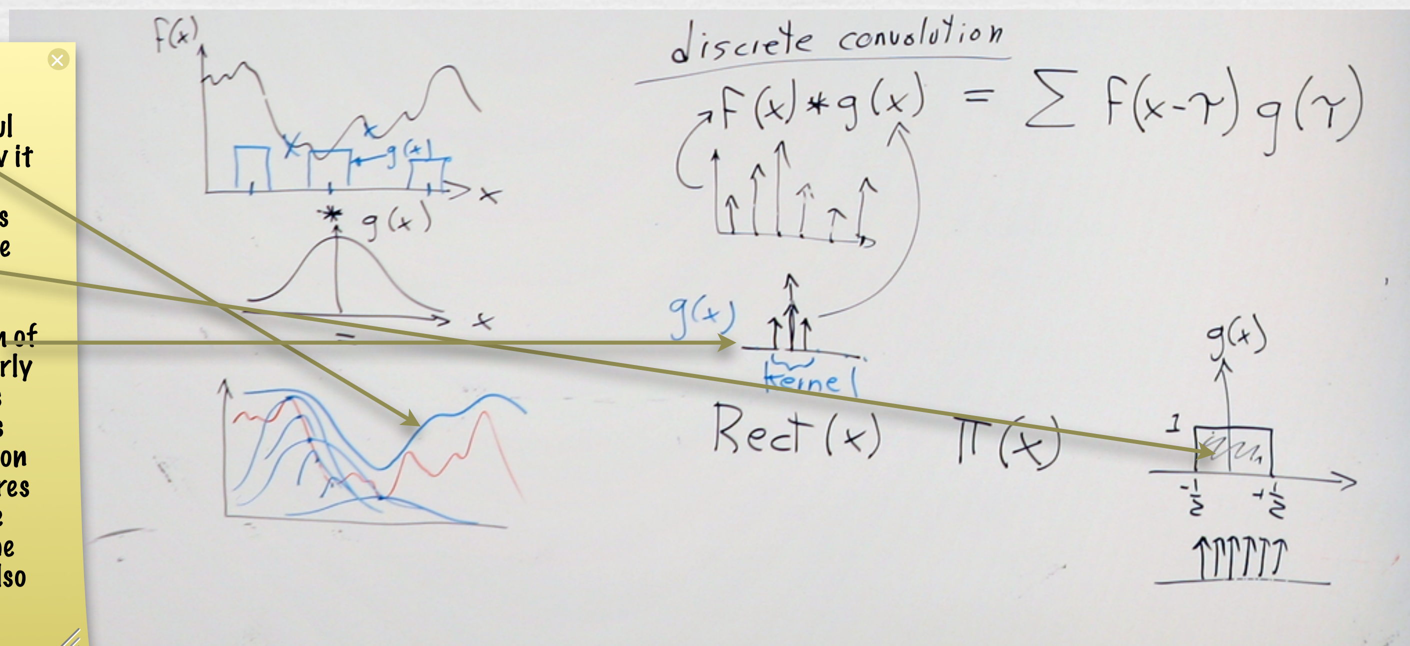

In this graphical construction for convolution, I should have been more careful when drawing the smoothed result to draw it passing *through* the original signal, not above it. For continuous convolution, this is accomplished by dividing the integral by the area under the filter function g(x) (or r(x)), which by design should be 1.0. For discrete convolution, this means dividing by the sum of the non-zero taps in the filter, which similarly should sum to 1.0 by design. This division is called normalization. When convolution is used for reconstruction, i.e. for interpolation between samples, this normalization ensures that the black (or white) area between the samples being interpolated doesn't make the result look too dark (or washed out). See also next slide, added 4/21/09.

Normalization during convolution

- with appropriate normalization, the smoothed result of convolution passes more or less through the original function values

As noted in class on 4/21/09, in order to make the area of g(x) equal to 1, it may be necessary to make its maximum value less than 1. An example, for the Rect function, is shown in green at lower right. In this case, the vertically scaled copies (red curves) will not reach as high as the original function values. However, if the area of g(x) is 1, then the result of the convolution integral is guaranteed to pass more or less through the original function values.

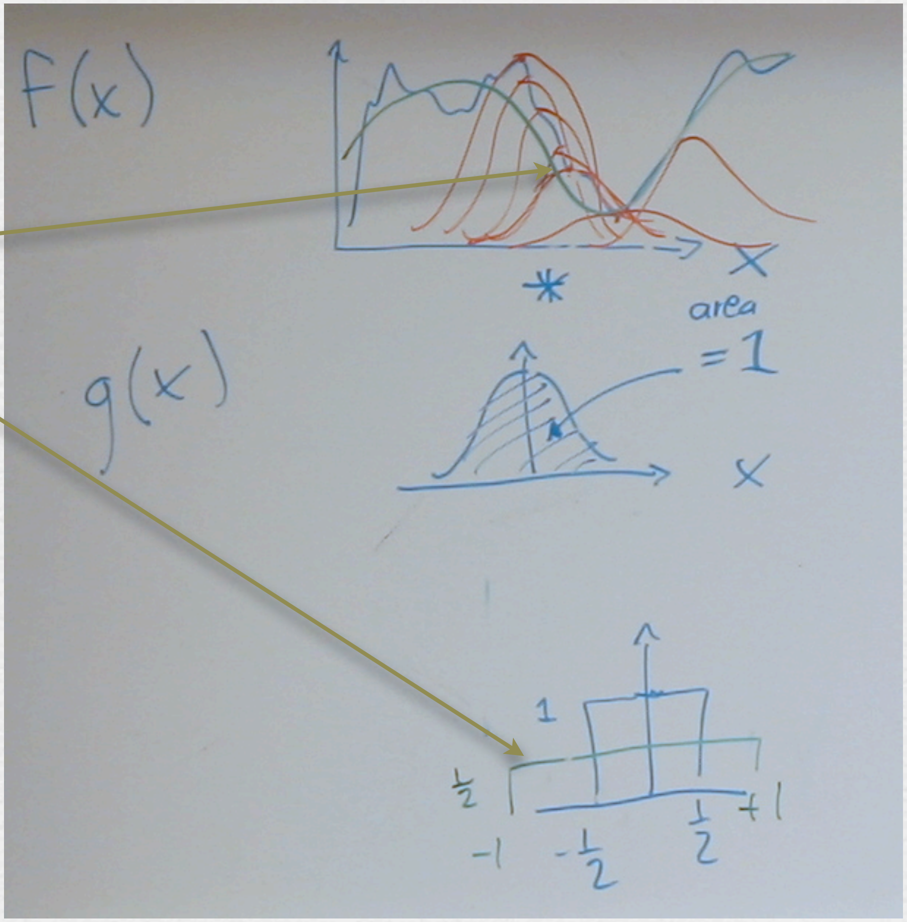

Sampling and reconstruction (cont.)

- reconstruction (a.k.a. *post-filtering*) by performing another spatial convolution, typically with a filter of a different size
  - its task is to interpolate between the discrete samples
- in computer graphics and digital photography, we don't worry about reconstruction; we hope the display or printer does it
  - blurring in the human eye helps with the reconstruction task

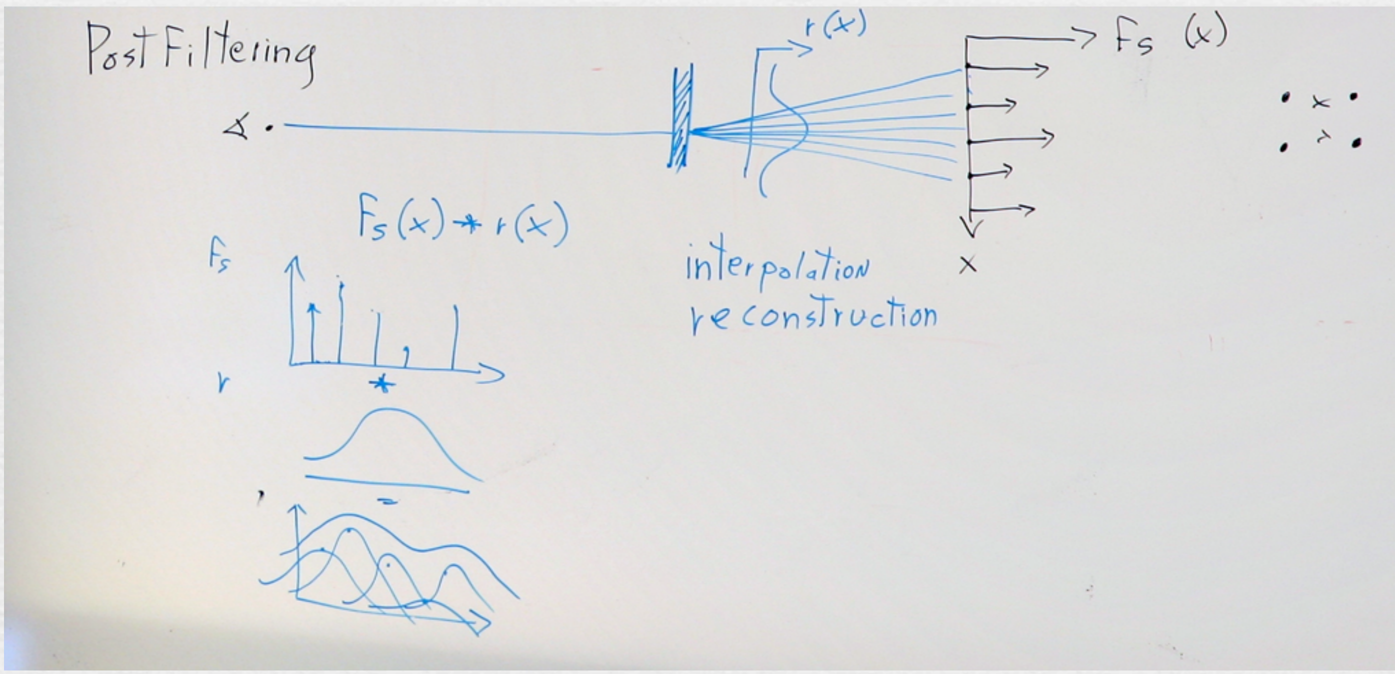

Sampling and reconstruction (cont.)

- summary of the pipeline
  - aliasing error is high frequencies masquerading as low frequencies due to insufficient sampling or prefiltering
  - reconstruction error is a residual of the sampling pattern due to insufficient postfiltering, a.k.a. rastering in the case of an image
  - no amount of postfiltering can remove aliasing error; once moire is in the image, it's too late to remove it

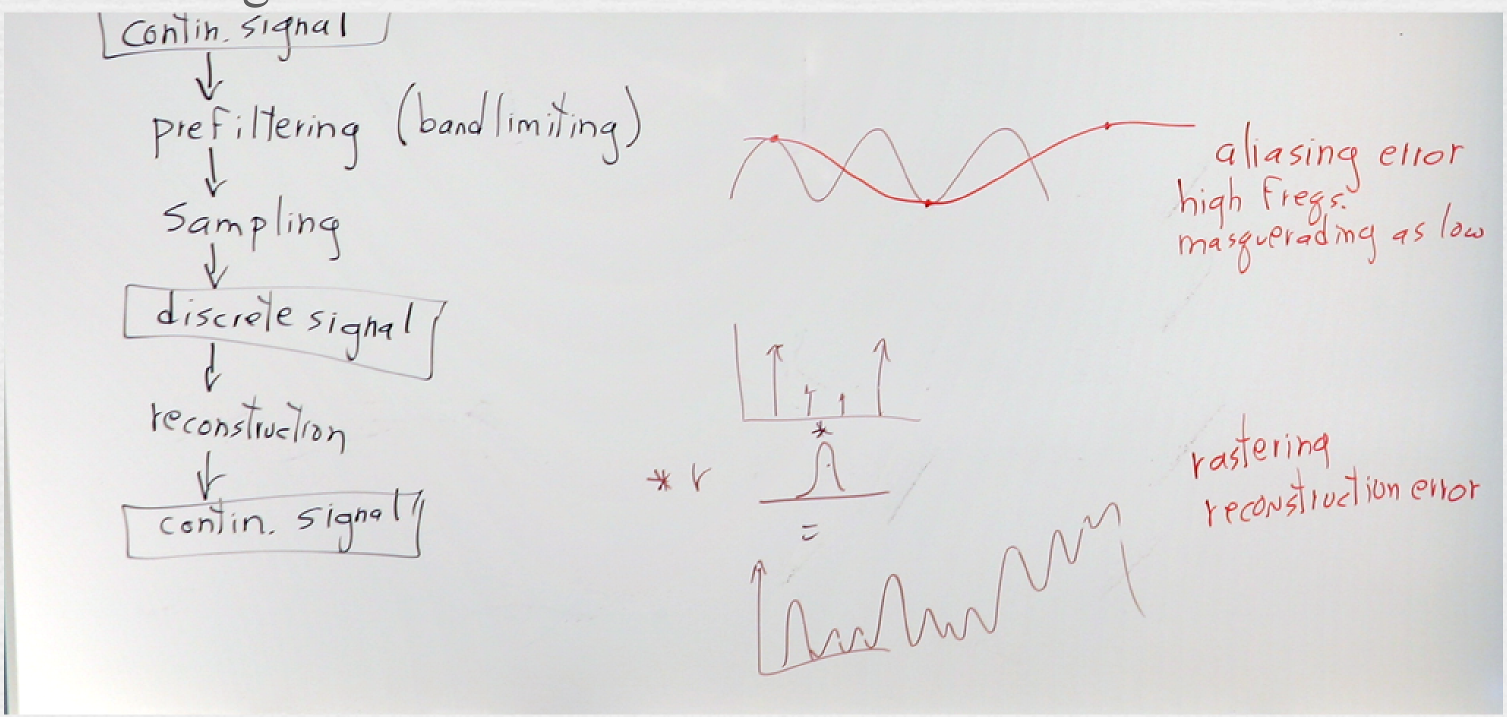

Example of spatial aliasing: jaggies in computer graphics

- an edge (e.g. of a colored triangle) contains (infinitely) high frequencies; merely sampling it produces jaggies
  - the rhythm of these jaggies is aliasing - a false low frequency
- to remove the jaggies, blur the edge (in software), i.e. draw it with pixels of intermediate colors
  - feathering a Selection in Photoshop will do this for you
- no amount of blurring a jaggy image will remove the jaggies
  - that's attempting to use postfiltering to remove aliasing

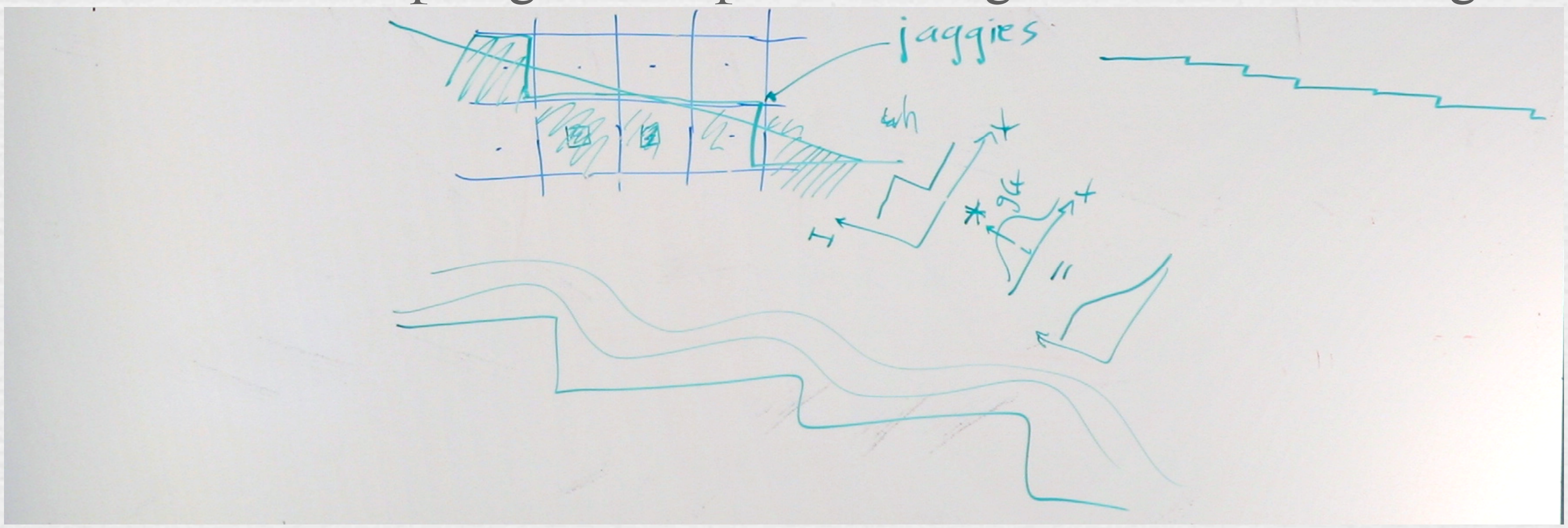

Example of temporal aliasing: wagon wheels spinning backwards in movies

- old movie cameras had longer intervals when the shutter was closed between frames than modern movie cameras
  - each exposure was like convolving the time-varying scene with a (1D) Rect in time, but the Rect wasn't as long as the interval between frames, so some action was missed
  - this lack of an adequate prefilter caused aliasing, as shown below

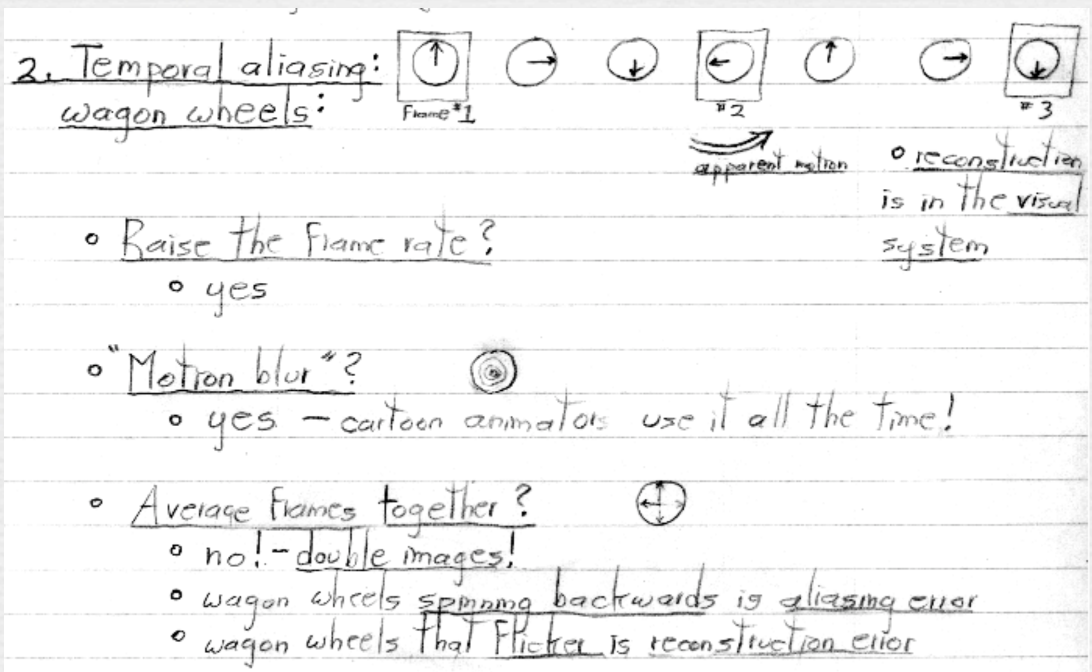

(from CS 248 notes at http://graphics.stanford.edu/courses/cs248-08/samp/samp1.html, because the whiteboard wasn't photographed at the right moment)

Slide credits

- Pat Hanrahan

- Cornsweet, T.N., *Visual Perception,* Kluwer Academic Press, 1970.