ECE 4973: Lecture 10 Resampling

Samuel Cheng

Slide credits: Noah Snavely

Image Scaling

This image is too big to fit on the screen. How can we generate a half-sized version?

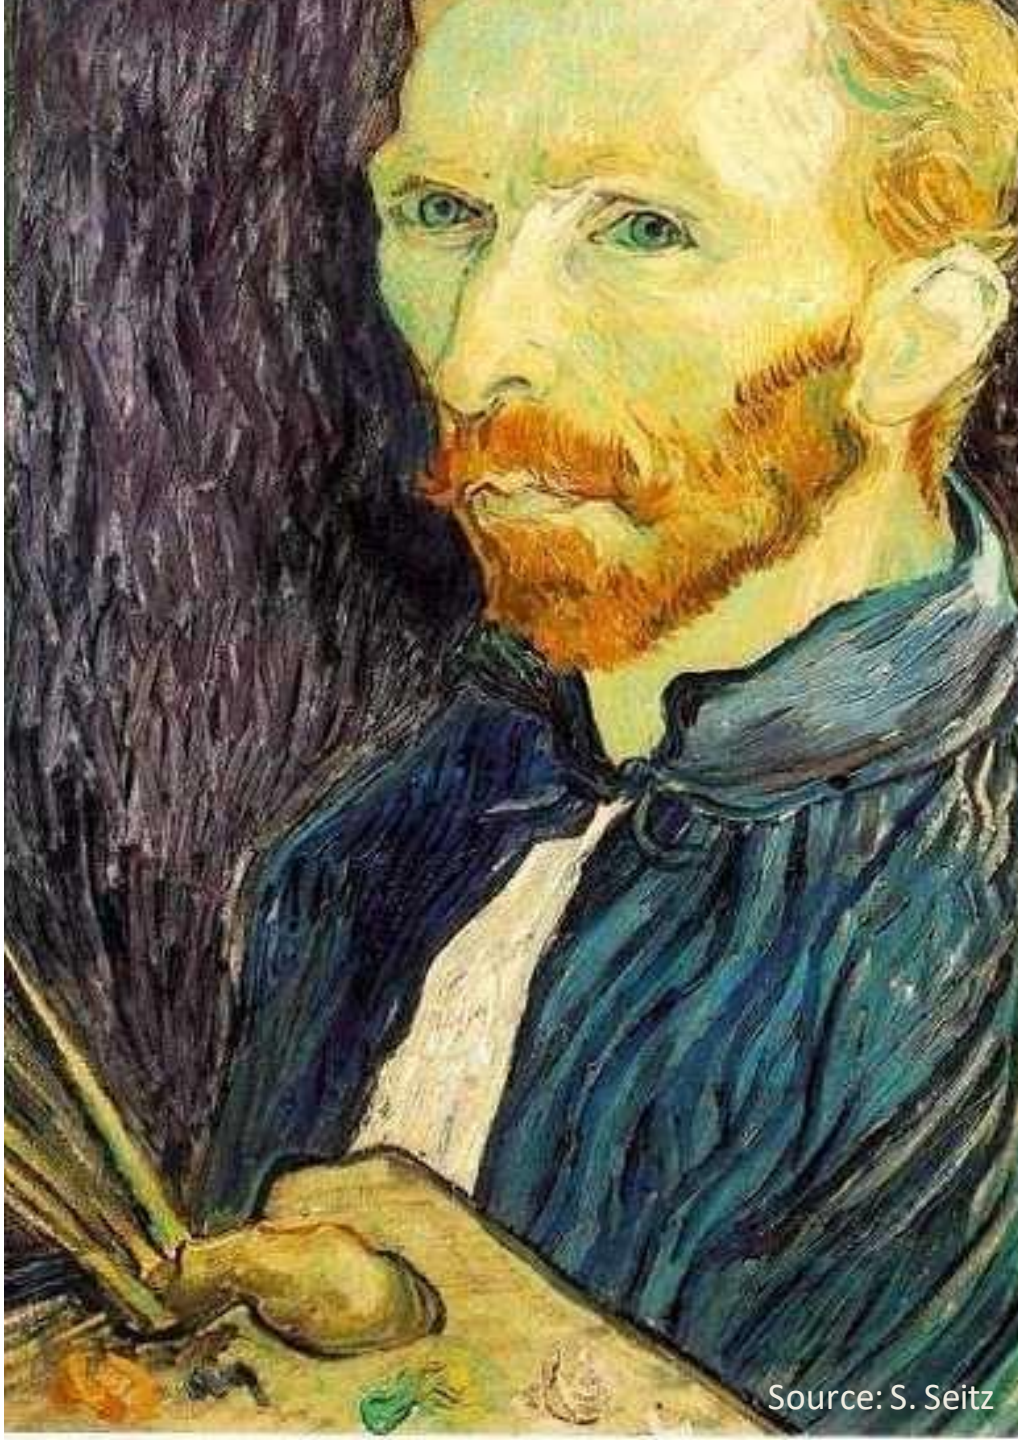

Image sub-sampling

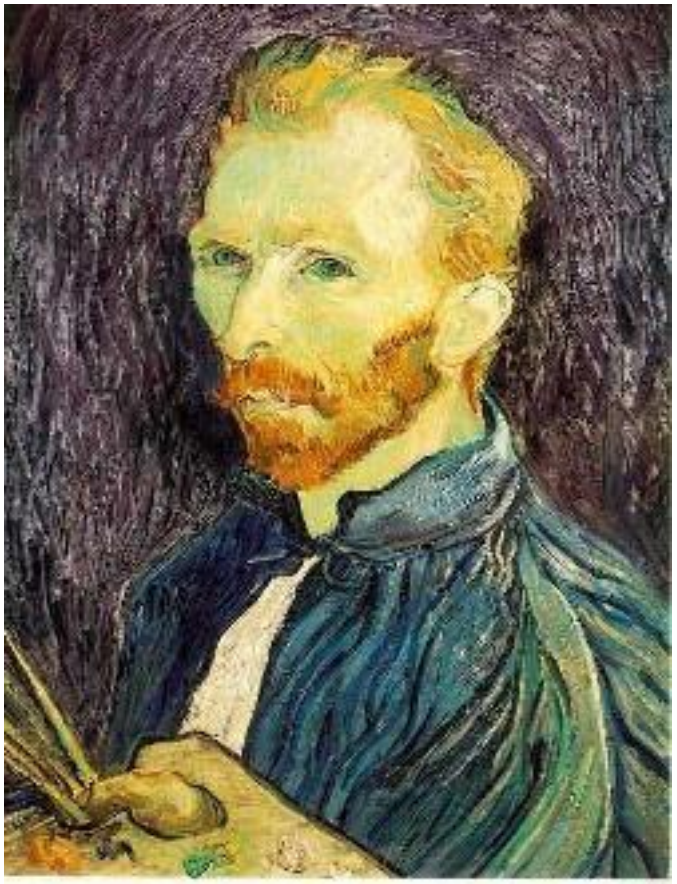

1/4

1/8

Throw away every other row and column to create a 1/2 size image

- called *image sub-sampling*

Source: S. Seitz

Image sub-sampling

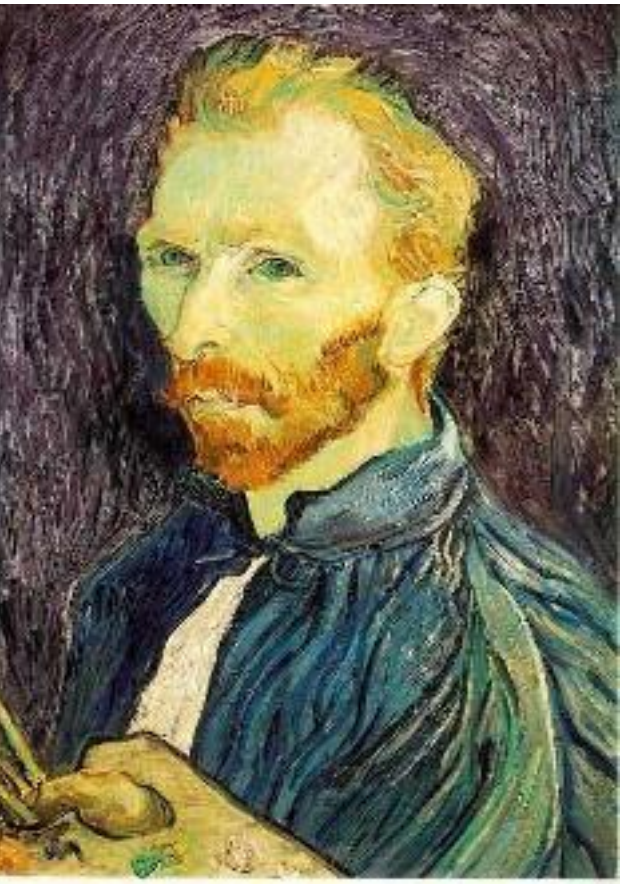

1/2

1/4 (2x zoom)

1/8 (4x zoom)

Why does this look so crufty?

Source: S. Seitz

Image sub-sampling

Source: F. Durand

Even worse for synthetic images

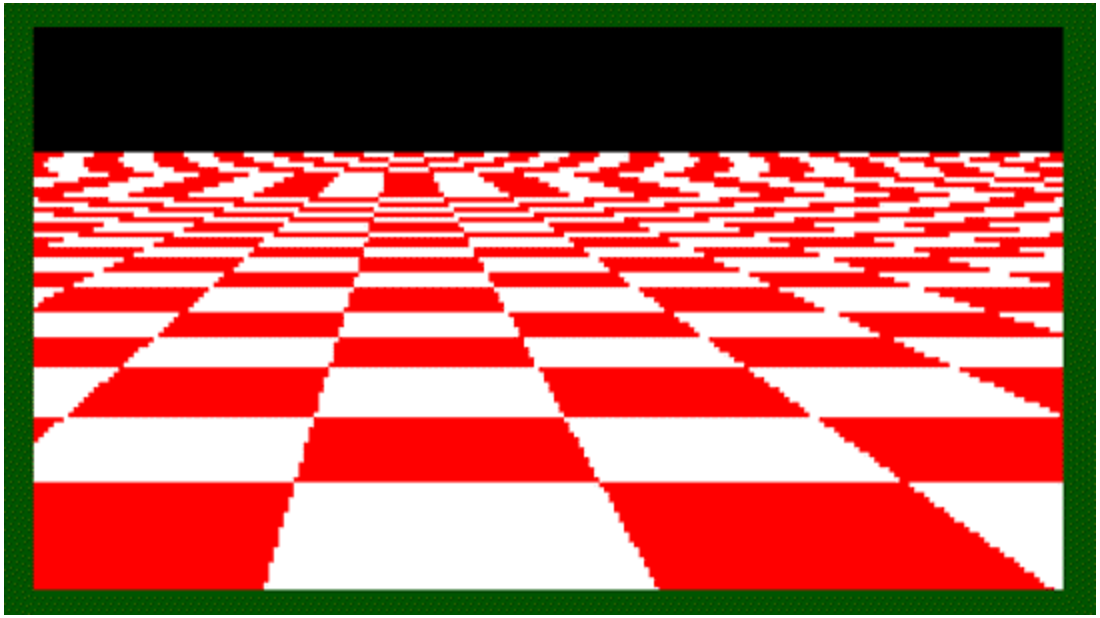

Source: L. Zhang

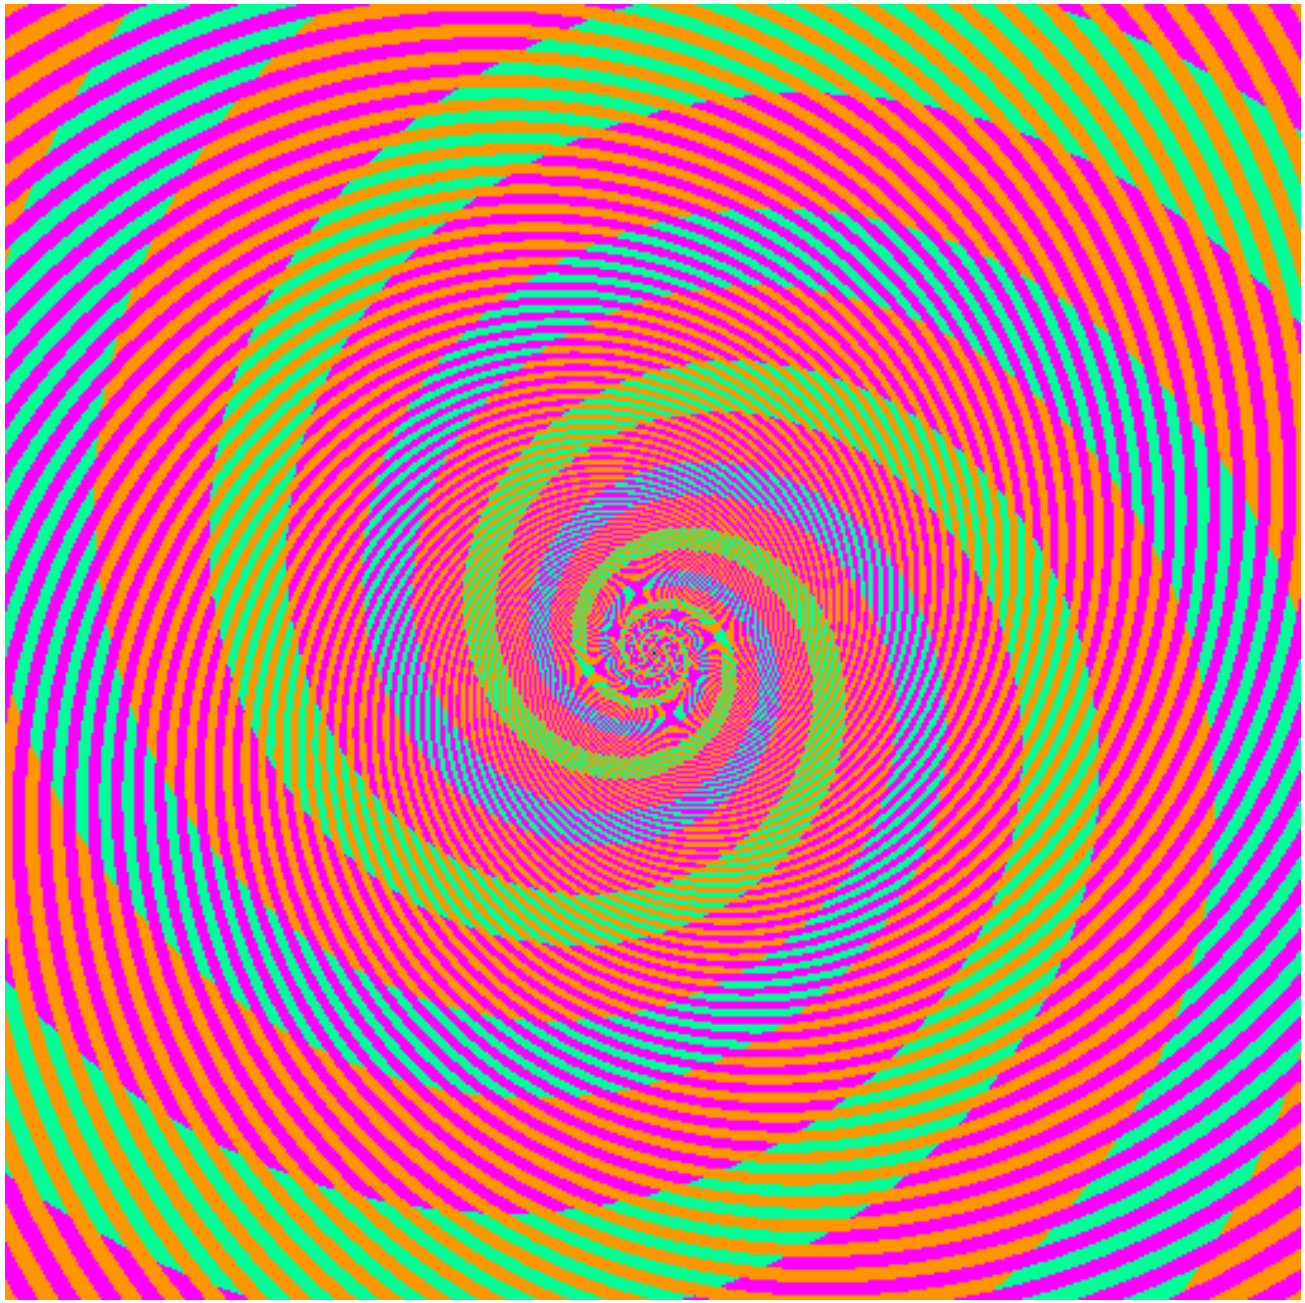

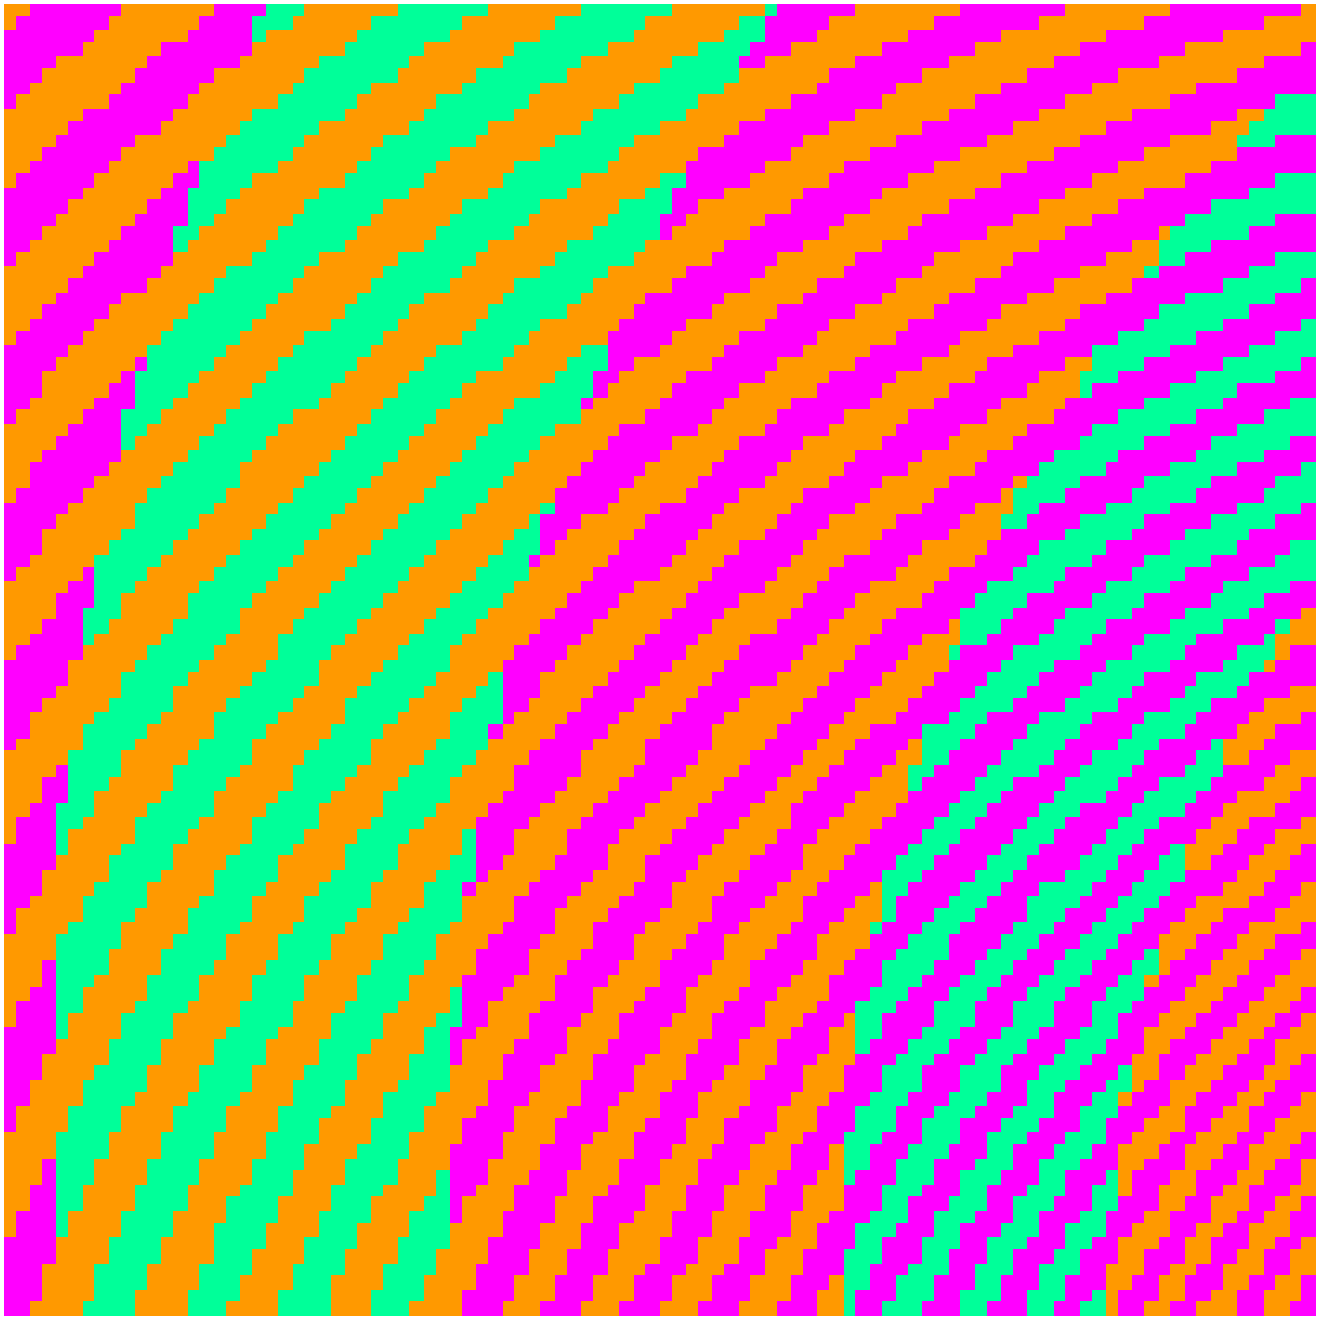

The blue and green colors are actually the same

http://blogs.discovermagazine.com/badastronomy/2009/06/24/the-blue-and-the-green/

256x256 128x128 64x64 32x32 16x16

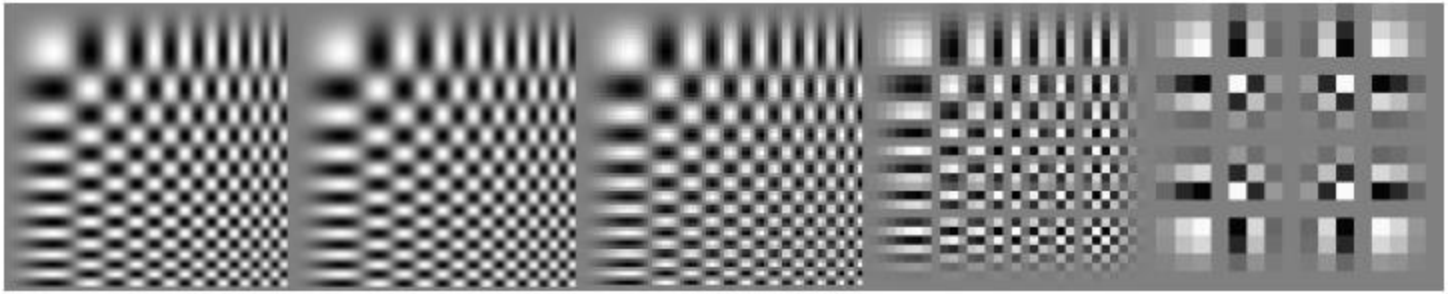

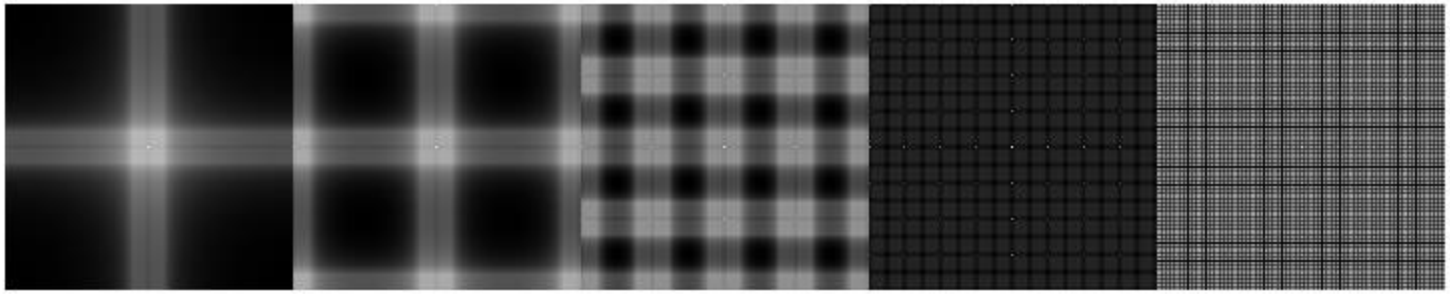

Interesting videos

Aliasing 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):

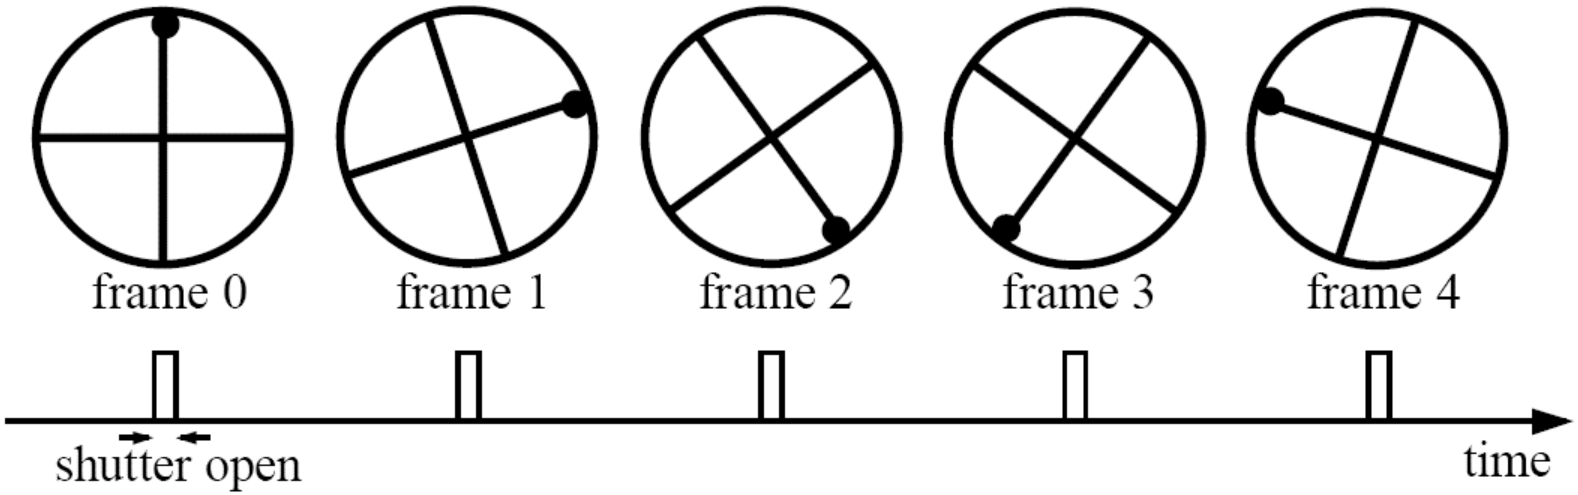

Without dot, wheel appears to be rotating slowly backwards! (counterclockwise)

Slide by Steve Seitz

Aliasing problem

- 1D example (sinewave):

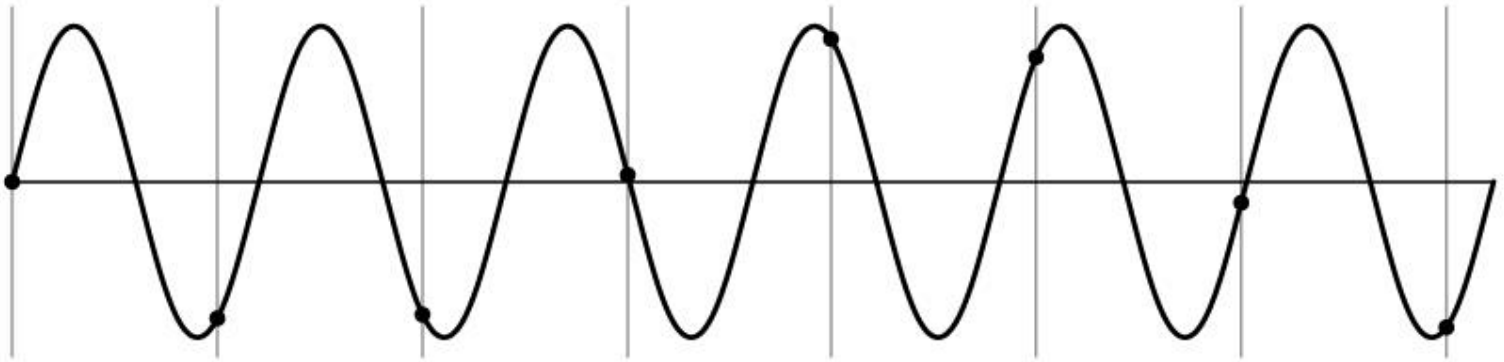

Source: S. Marschner

Aliasing problem

- 1D example (sinewave):

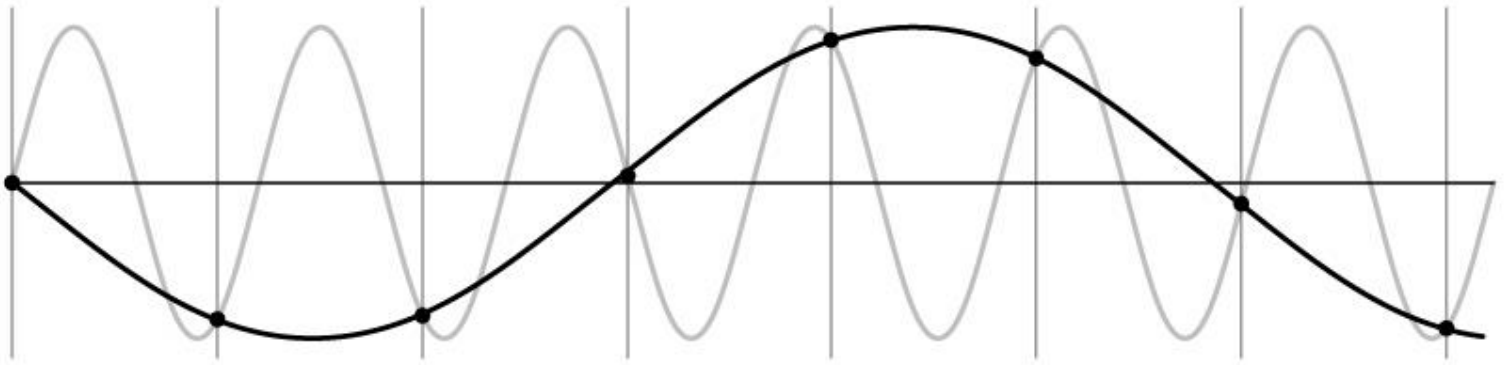

Source: S. Marschner

Nyquist-Shannon Sampling Theorem

- When sampling a signal at discrete intervals, the sampling frequency must be $\geq 2 \times f_{max}$
- f_{max} = max frequency

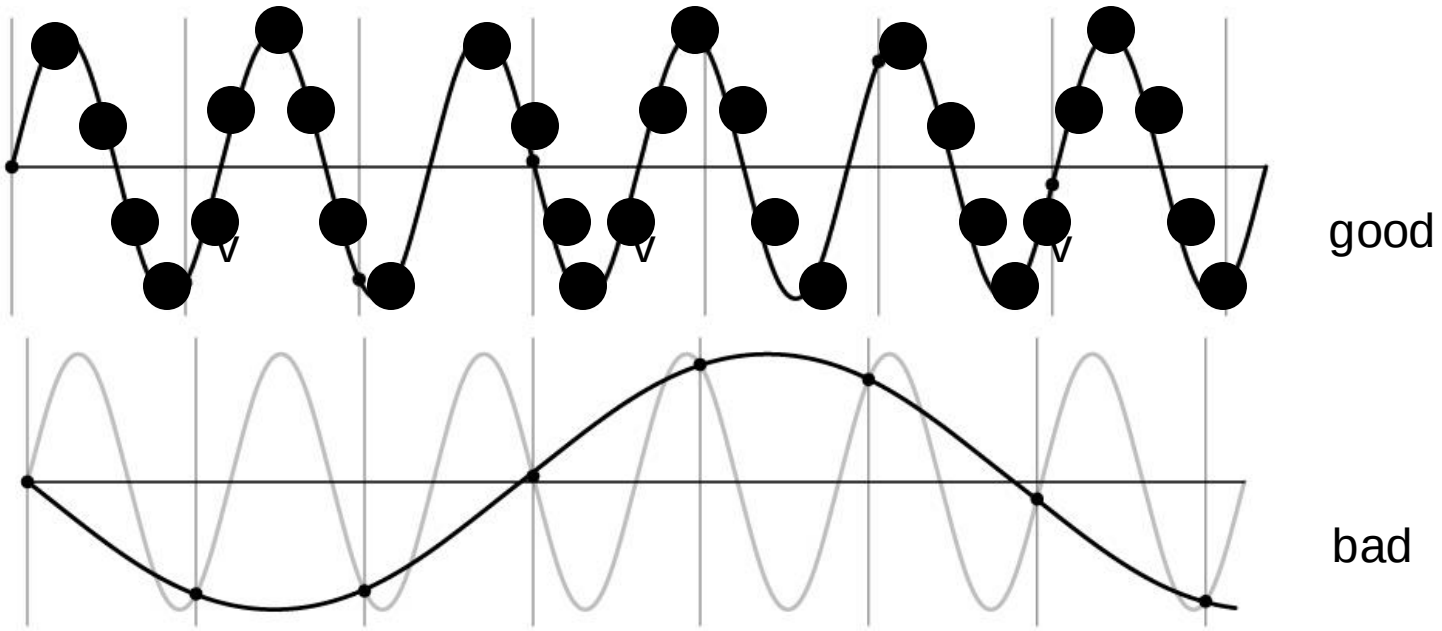

Nyquist limit – 2D example

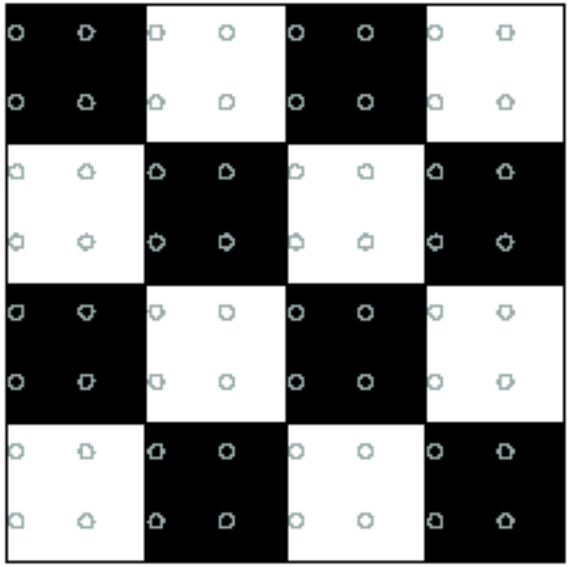

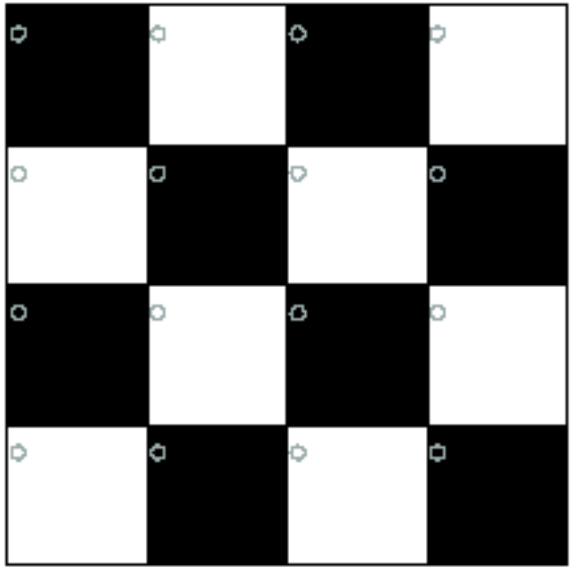

Good sampling

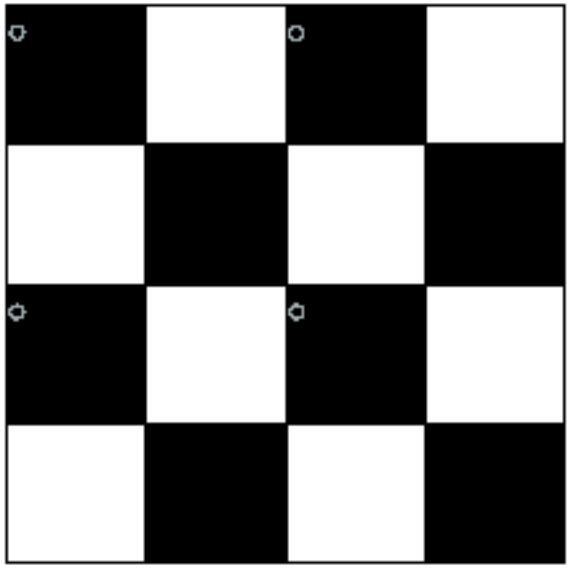

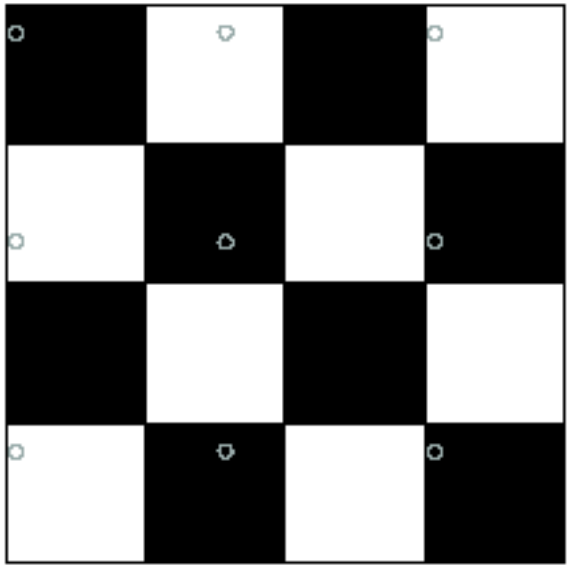

Bad sampling

FT under different sampling rates

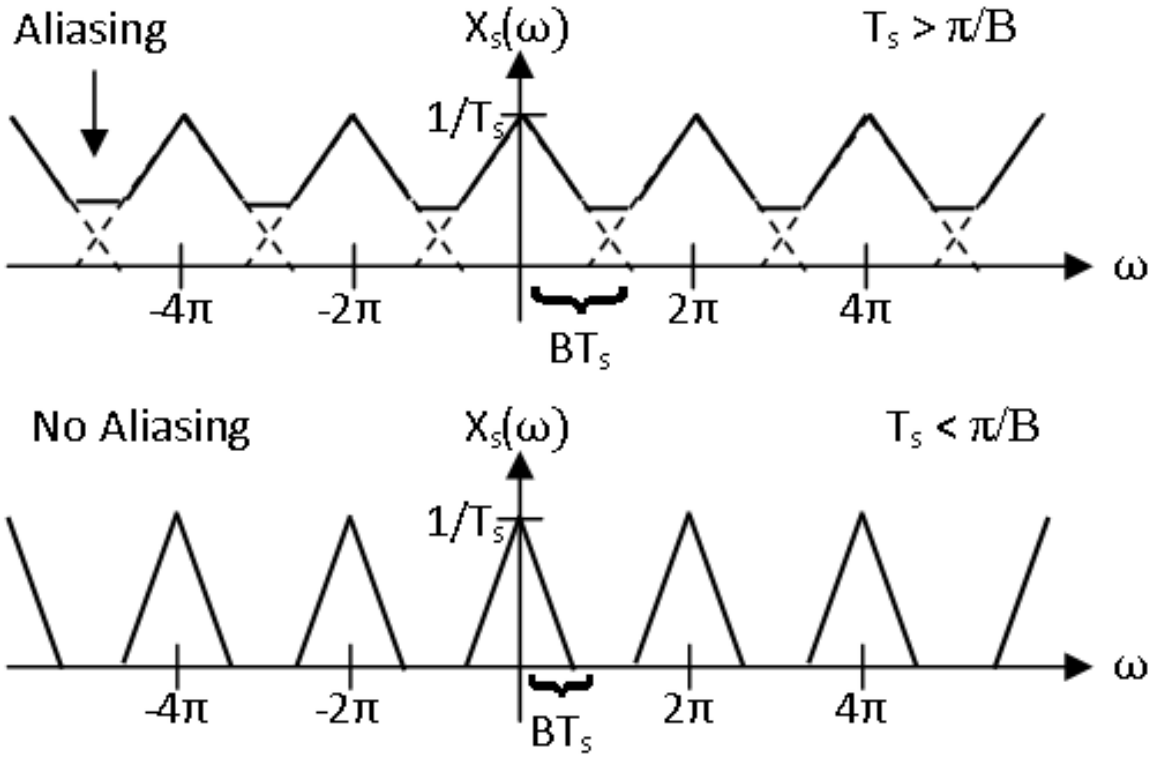

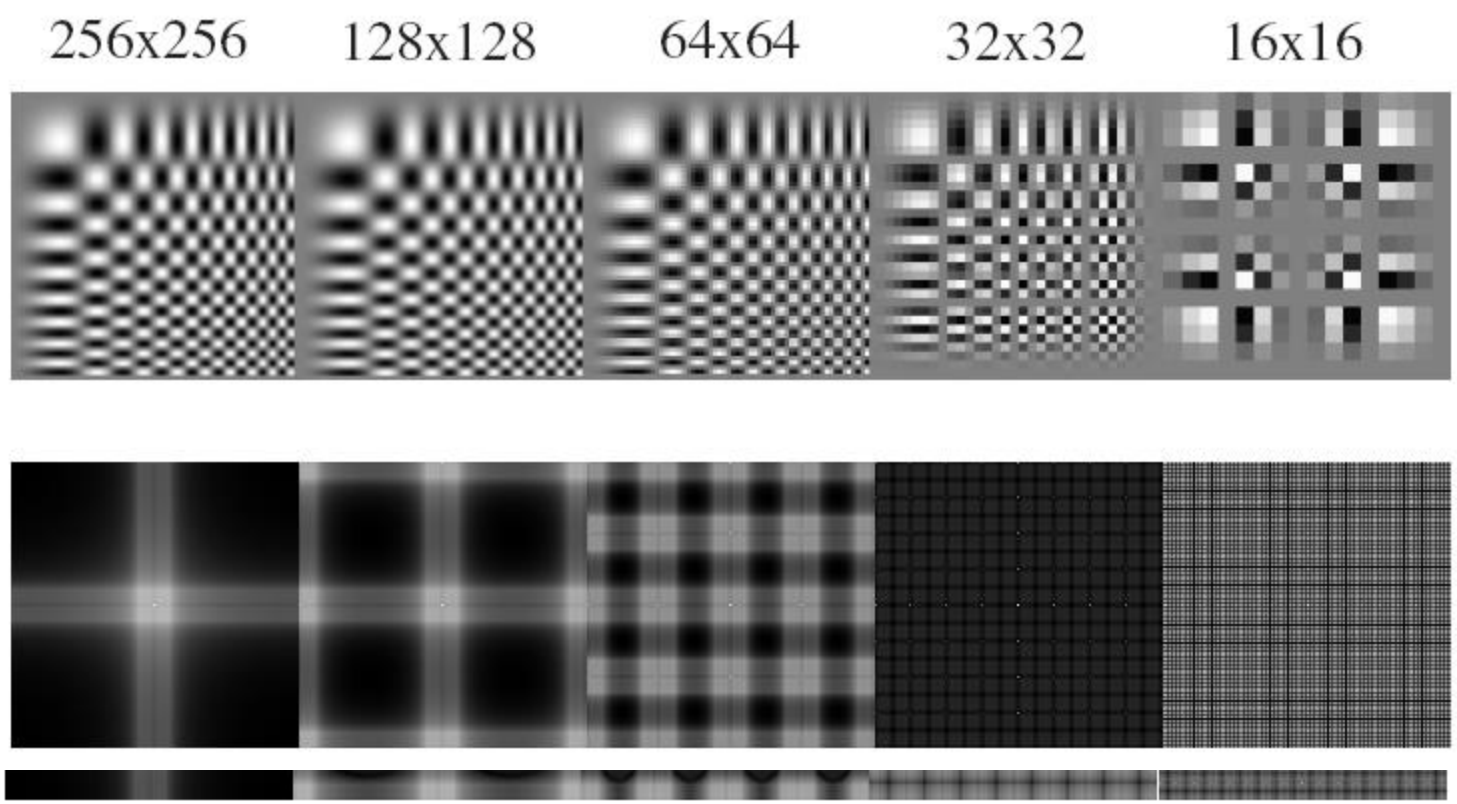

Forsyth and Ponce 2002

Gaussian pre-filtering

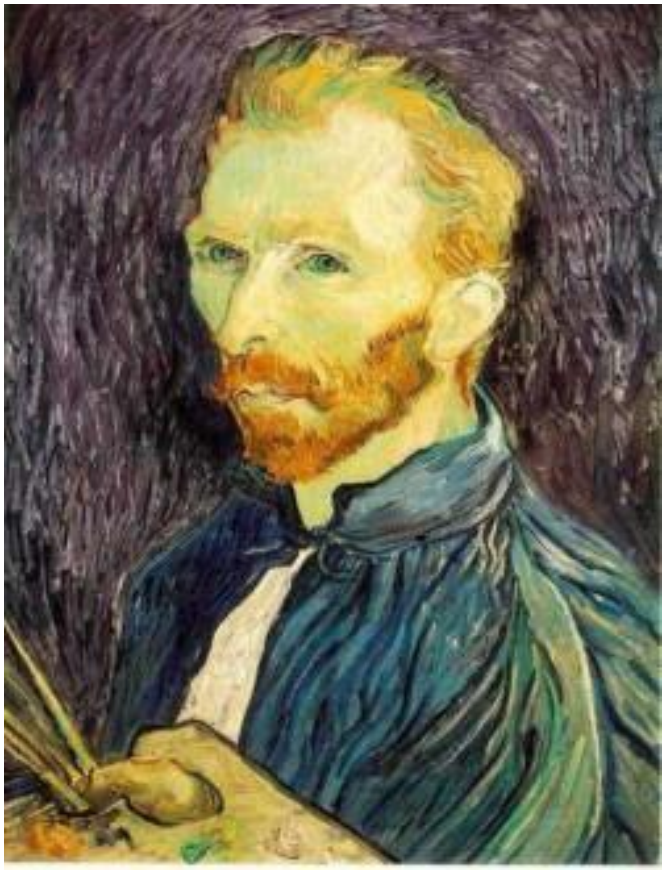

Gaussian 1/2

G 1/4

G 1/8

- Solution: filter the image, *then* subsample

Source: S. Seitz

Subsampling with Gaussian pre-filtering

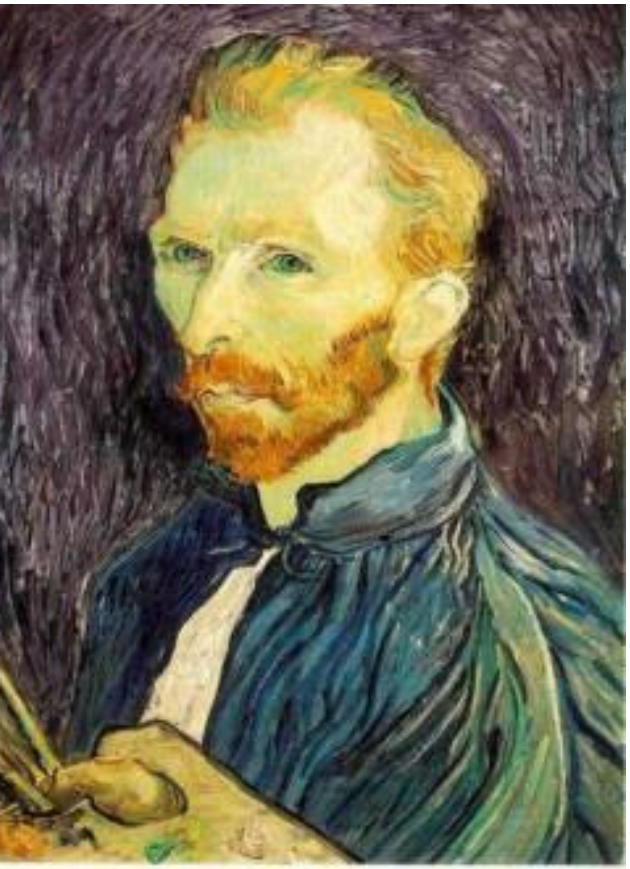

Gaussian 1/2

G 1/4

G 1/8

- Solution: filter the image, *then* subsample

Source: S. Seitz

Compare with...

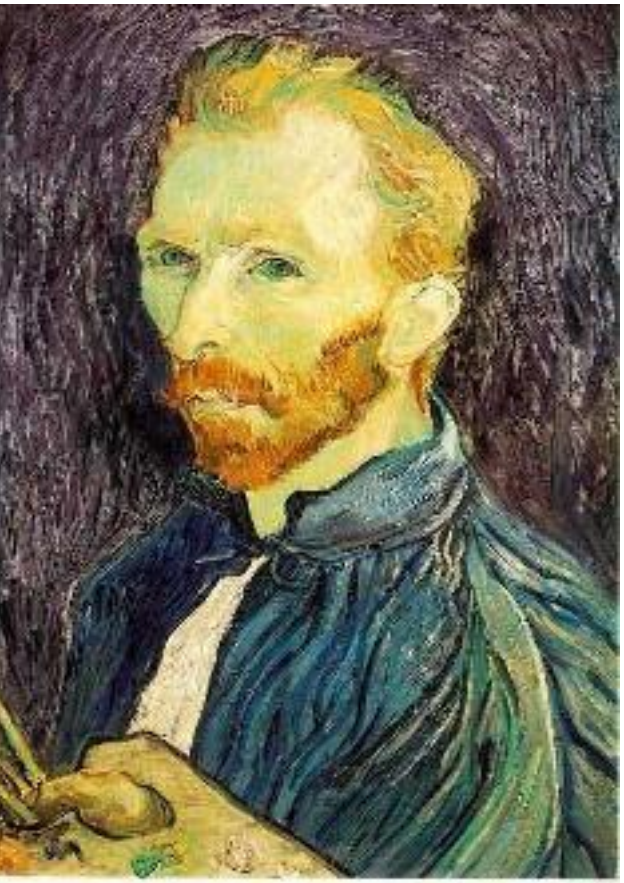

1/2

1/4 (2x zoom)

1/8 (4x zoom)

Source: S. Seitz

Upsampling

- This image is too small for this screen: 
- How can we make it 10 times as big?
- Simplest approach:
  repeat each row
  and column 10 times
- (“Nearest neighbor
  interpolation”)

Image interpolation

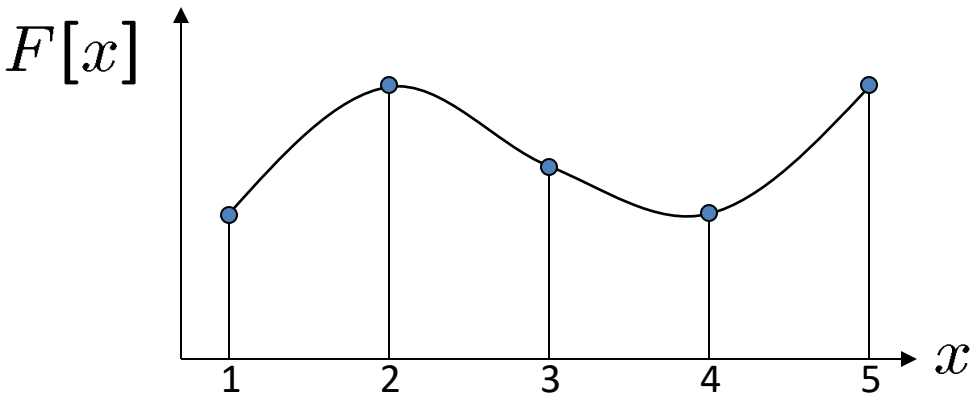

d = 1 in this example

Recall how a digital image is formed

$$F[x, y] = \text{quantize}\{f(xd, yd)\}$$

- It is a discrete point-sampling of a continuous function
- If we could somehow reconstruct the original function, any new image could be generated, at any resolution and scale

Adapted from: S. Seitz

Image interpolation

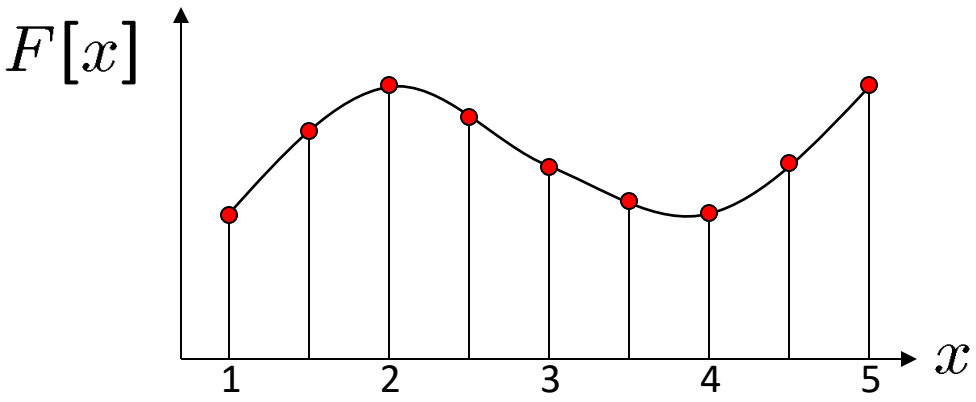

d = 1 in this example

Recall how a digital image is formed

$$F[x, y] = \text{quantize}\{f(xd, yd)\}$$

- It is a discrete point-sampling of a continuous function
- If we could somehow reconstruct the original function, any new image could be generated, at any resolution and scale

Adapted from: S. Seitz

Image interpolation

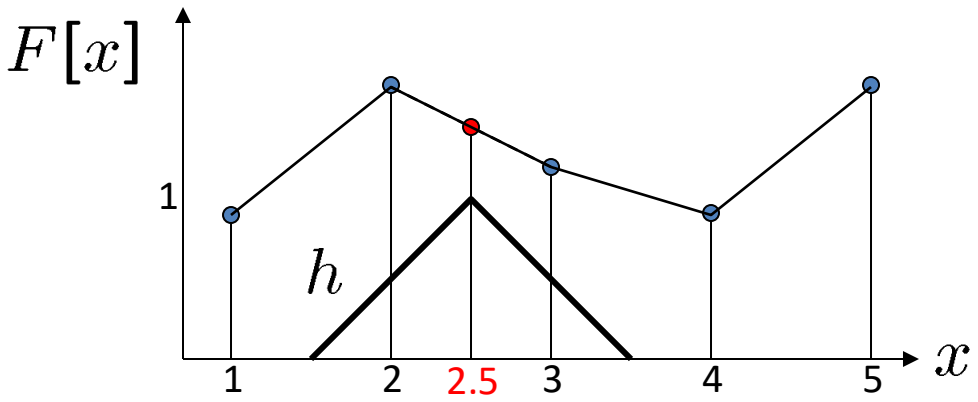

d = 1 in this example

- What if we don't know f ?
  - Guess an approximation: $\tilde{f}$
  - Can be done in a principled way: filtering
  - Convert F to a continuous function:

$$f_F(x) = F(\tfrac{x}{d}) \text{ when } \tfrac{x}{d} \text{ is an integer, 0 otherwise}$$

  - Reconstruct by convolution with a *reconstruction filter, h*

$$\tilde{f} = h * f_F$$

Adapted from: S. Seitz

Frequency representation

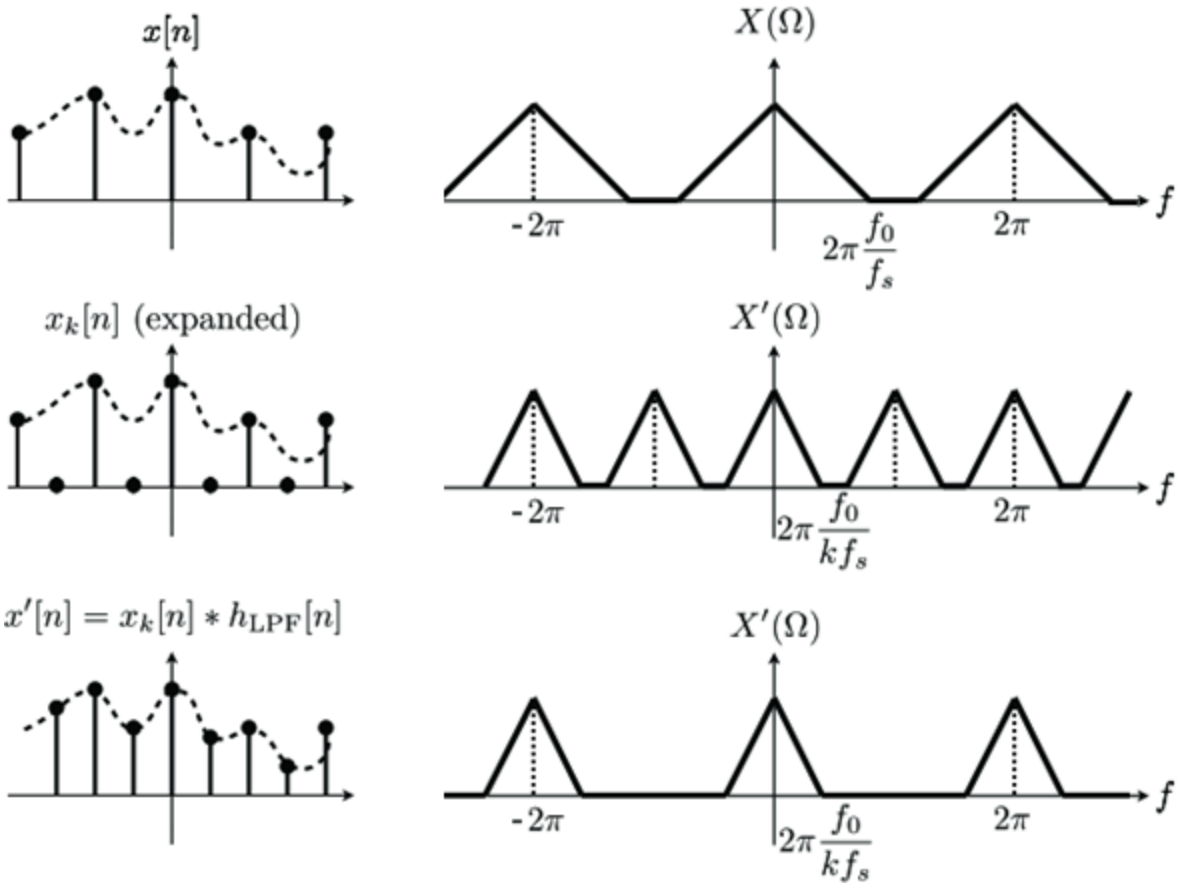

Image interpolation

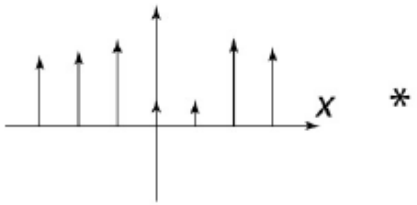

*

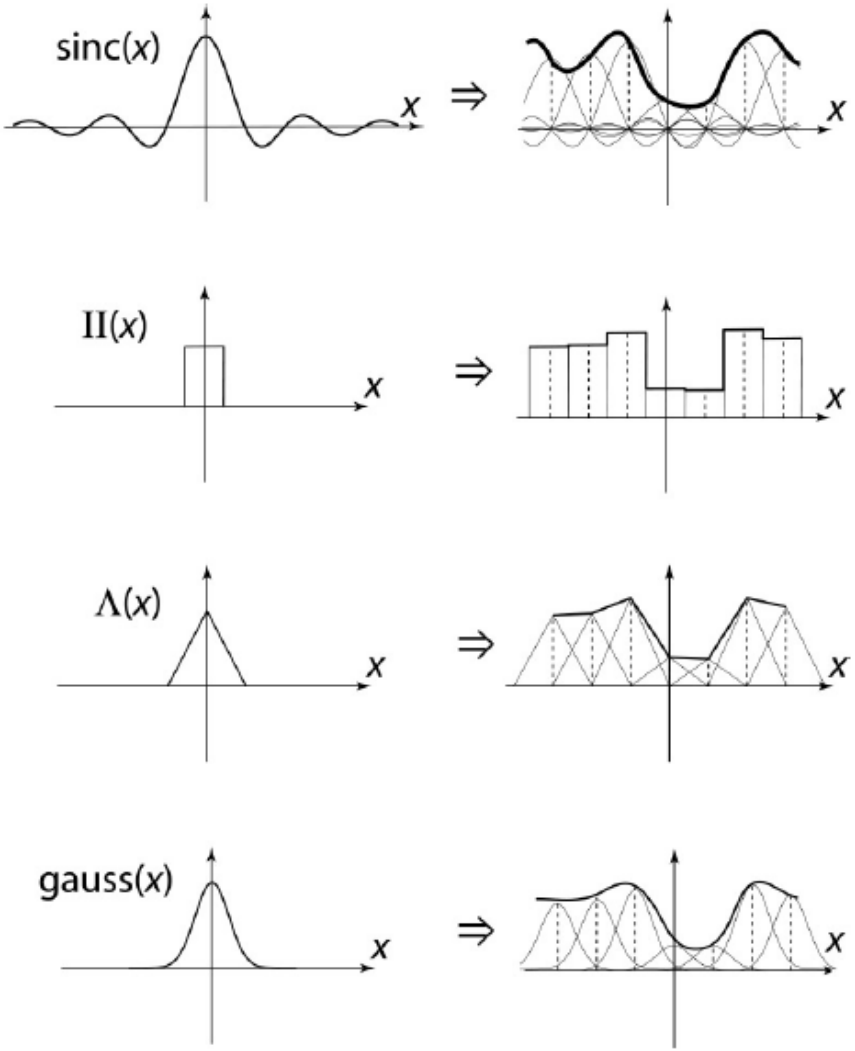

"Ideal" reconstruction

Nearest-neighbor interpolation

Linear interpolation

Gaussian reconstruction

Source: B. Curless

Reconstruction filters

- What does the 2D version of this hat function look like?

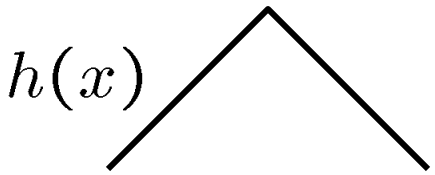

performs
linear interpolation

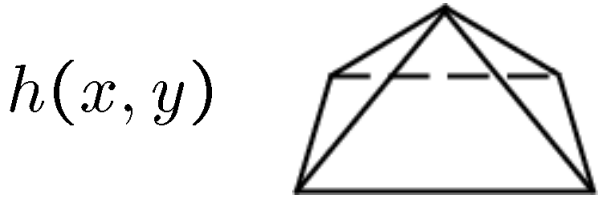

(tent function) performs
bilinear interpolation

Better filters give better resampled images

- **Bicubic** is common choice

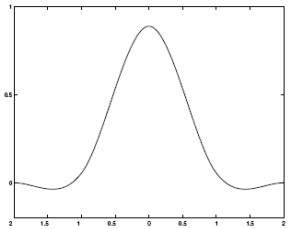

Cubic reconstruction filter

$$r(x) = \frac{1}{6}\begin{cases} (12-9B-6C)|x|^3 + (-18+12B+6C)|x|^2 + (6-2B) & |x|<1 \\ ((-B-6C)|x|^3 + (6B+30C)|x|^2 + (-12B-48C)|x| + (8B+24C) & 1\le|x|<2 \\ 0 & \textit{otherwise} \end{cases}$$

Image interpolation

Original image:  x 10

Nearest-neighbor interpolation

Bilinear interpolation

Bicubic interpolation

DSP Interpretation

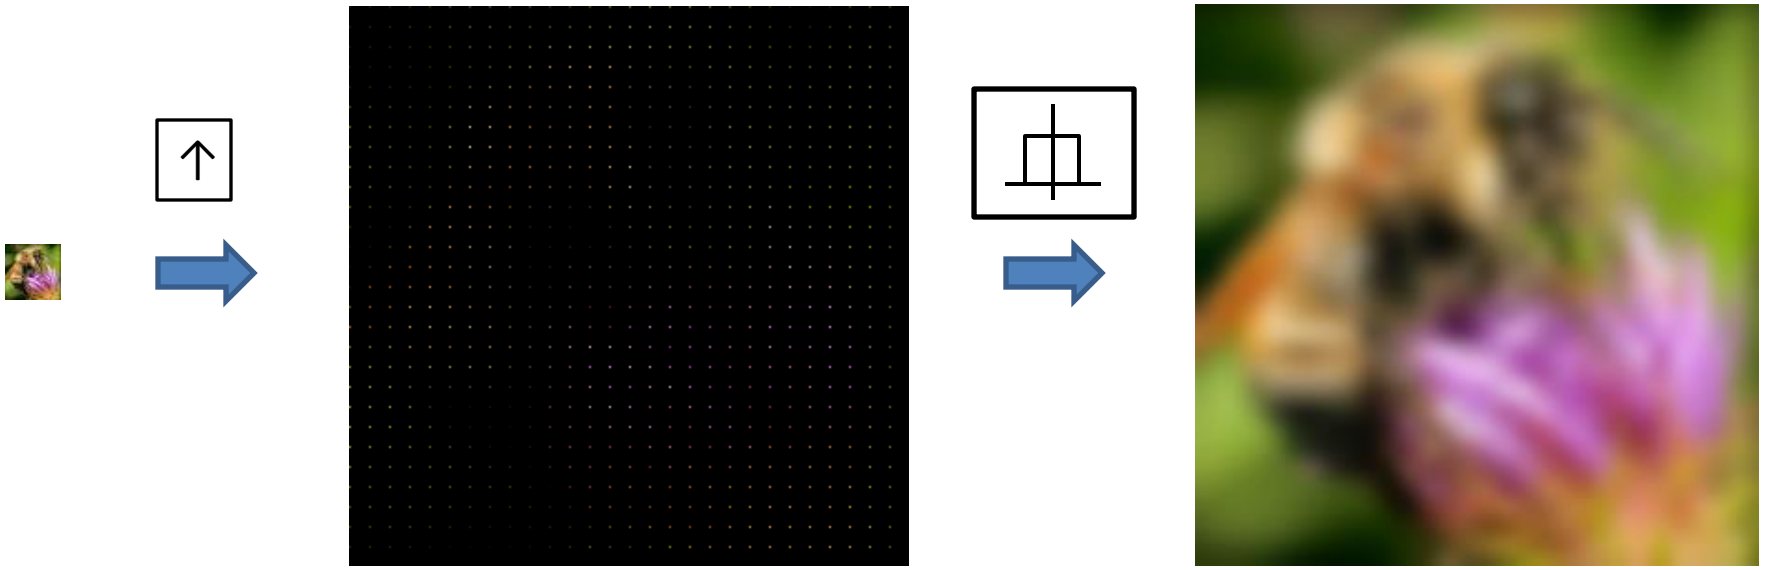

Image resampling

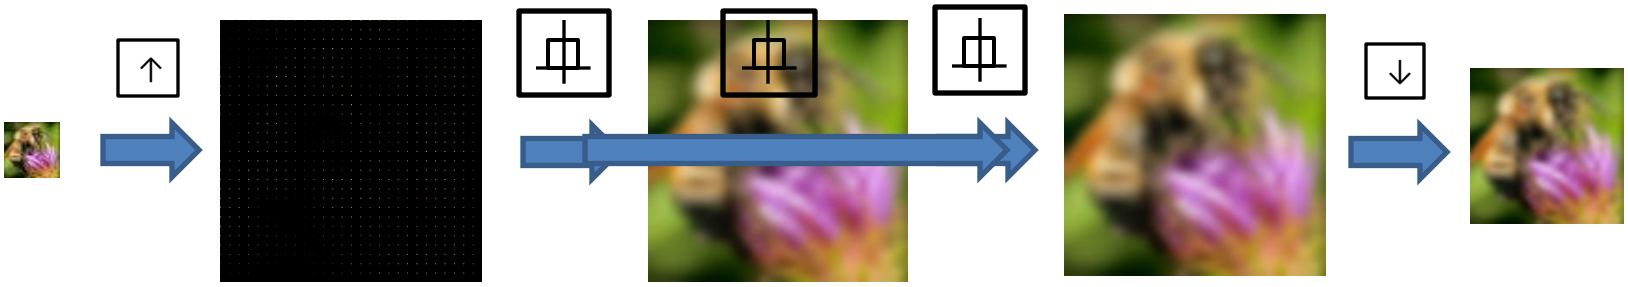

Upsampling

downsampling

Hybrid Image

Salvador Dali, 1976

Another example

- Who is (s)he?

Hybrid Images

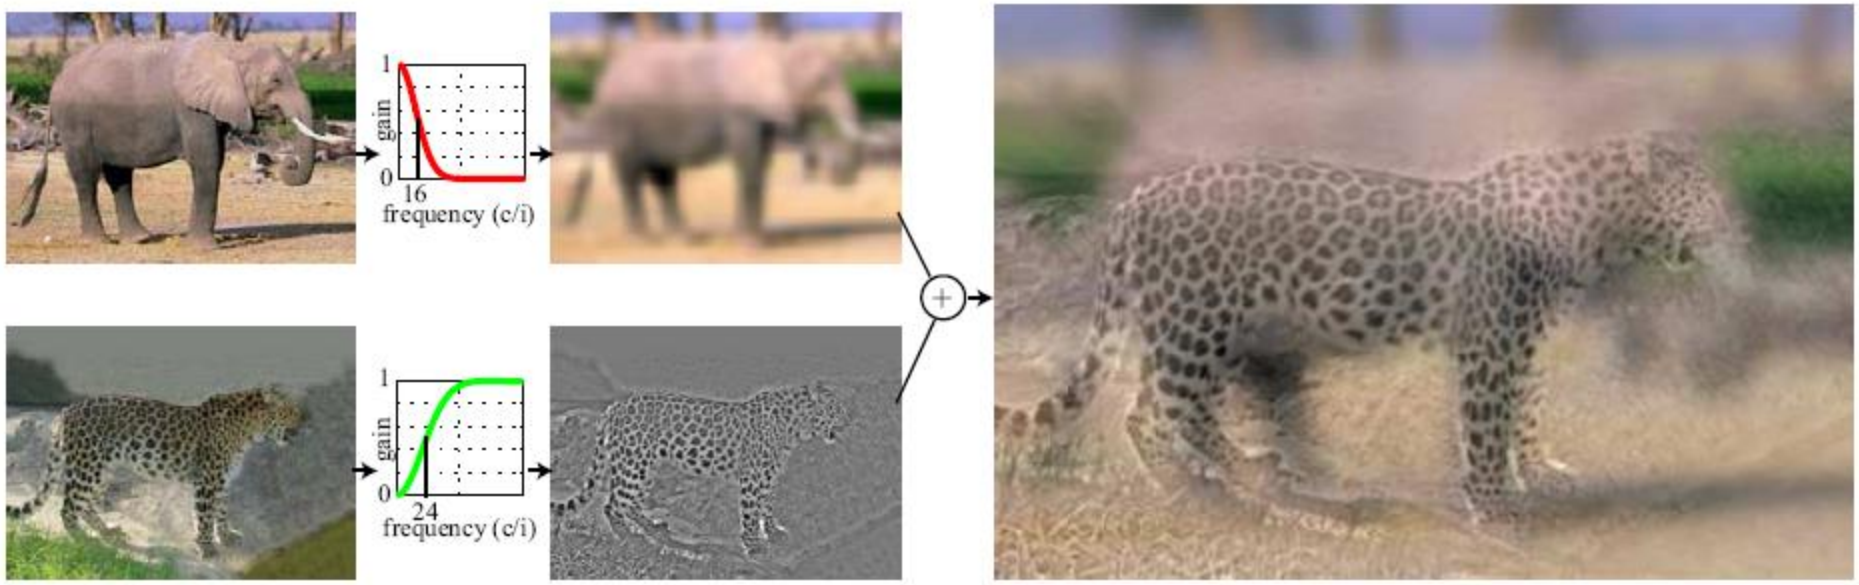

- A. Oliva, A. Torralba, P.G. Schyns, "Hybrid Images," SIGGRAPH 2006

Hybrid Image in FFT

Hybrid Image

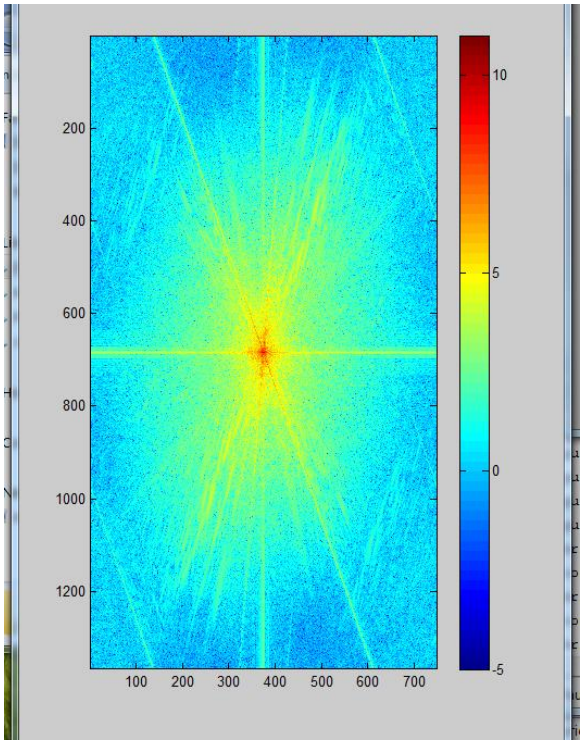

=

Low-passed Image

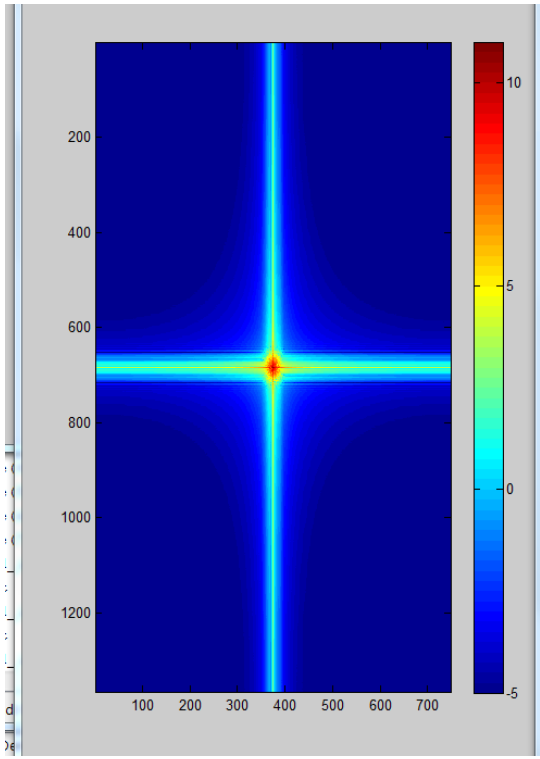

+

High-passed Image

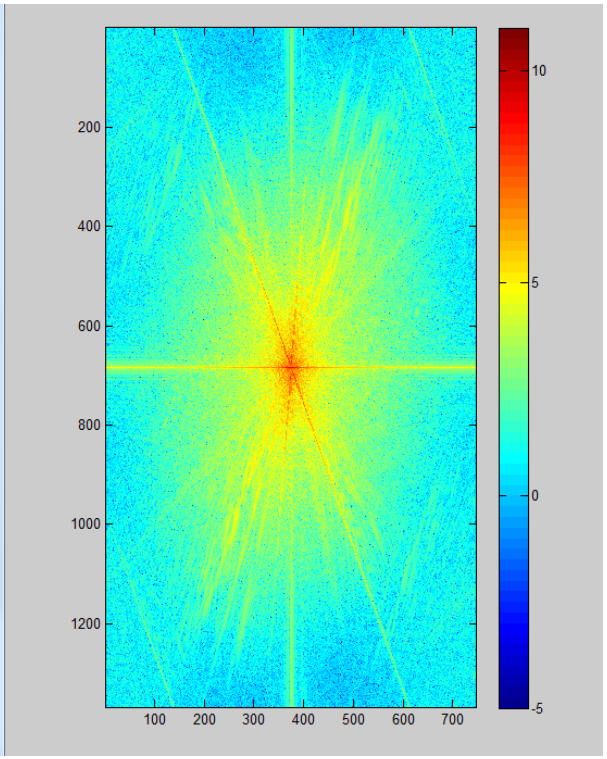

Why do we get different, distance-dependent interpretations of hybrid images?

?

Hays

Campbell-Robson contrast sensitivity curve

Perceptual cues in the mid-high frequencies dominate perception.

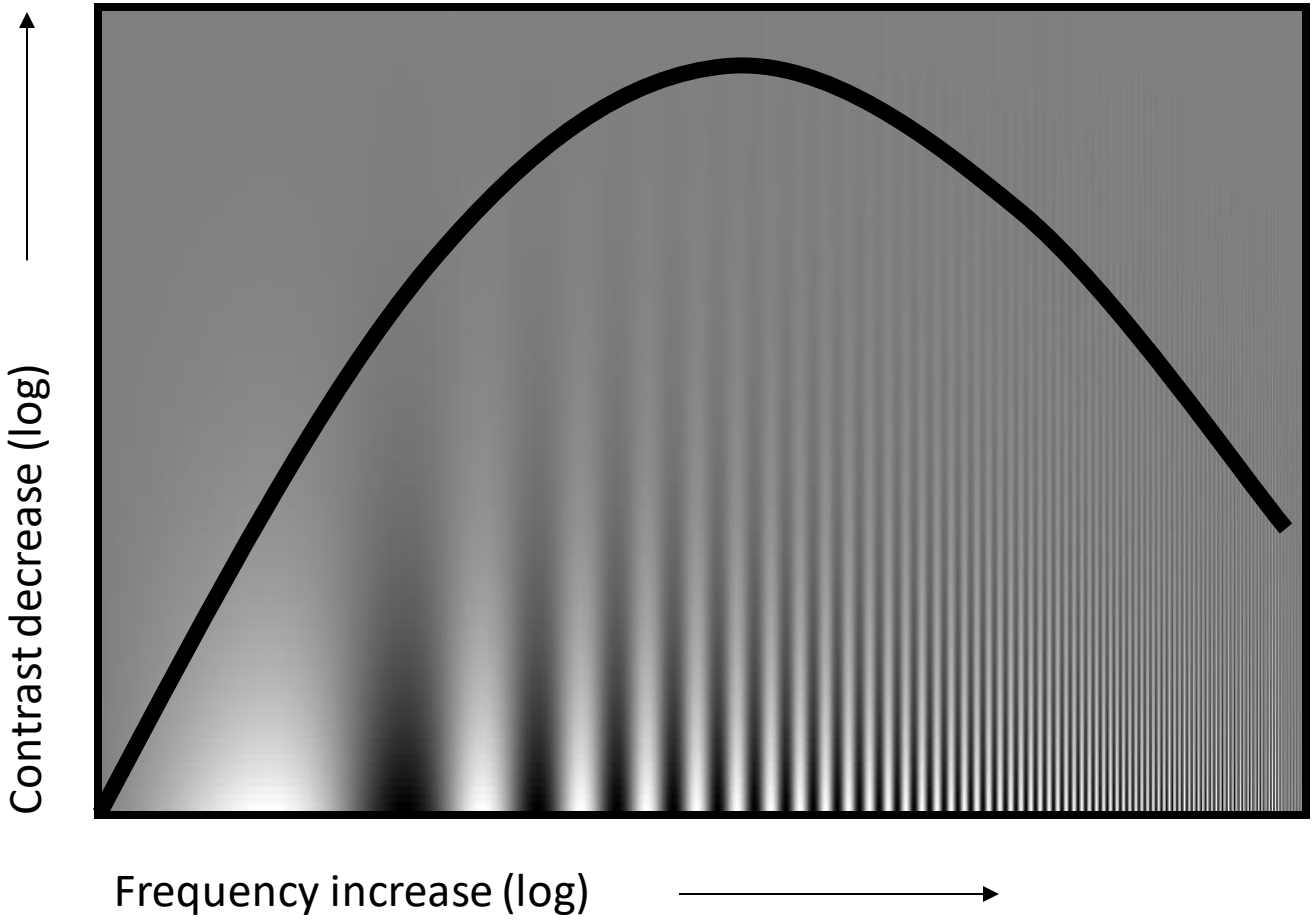

Questions?

Image Pyramids

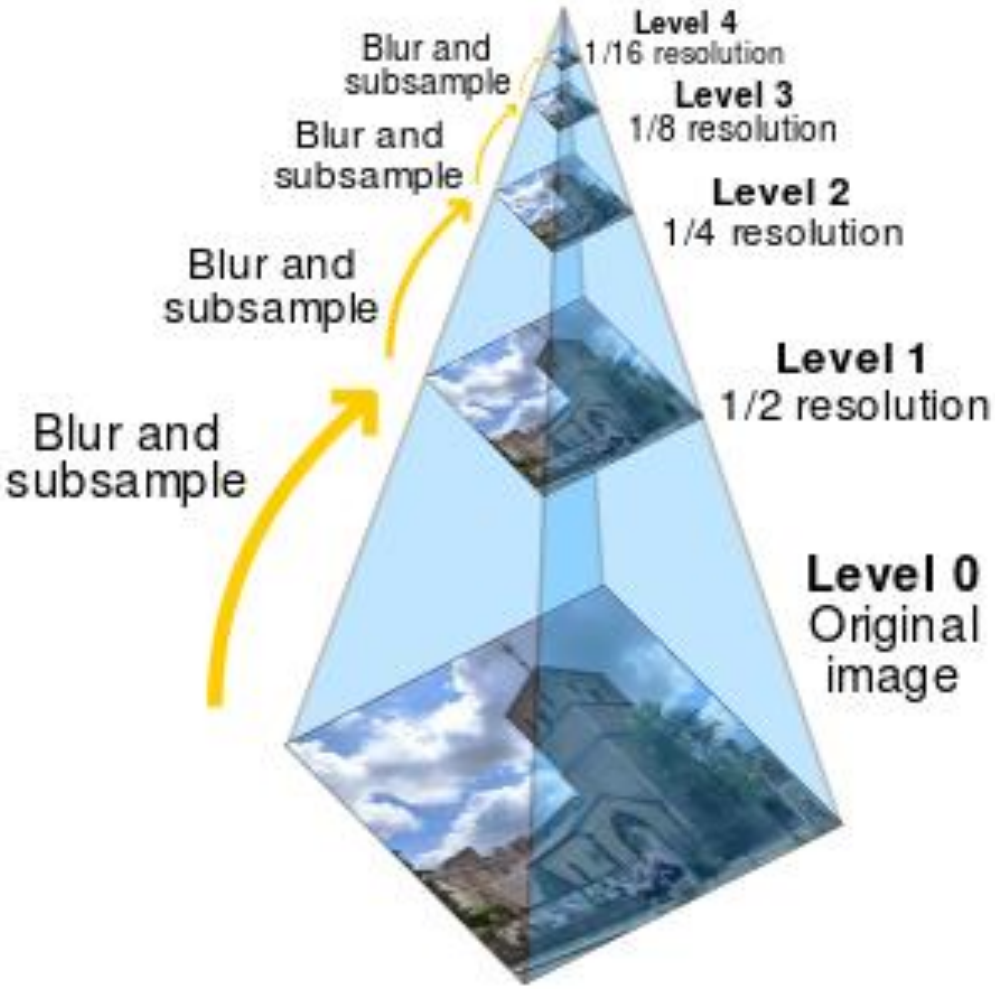

Project 1 function:
vis_hybrid_image.m

Wikipedia – Image Pyramids

Gaussian pyramid

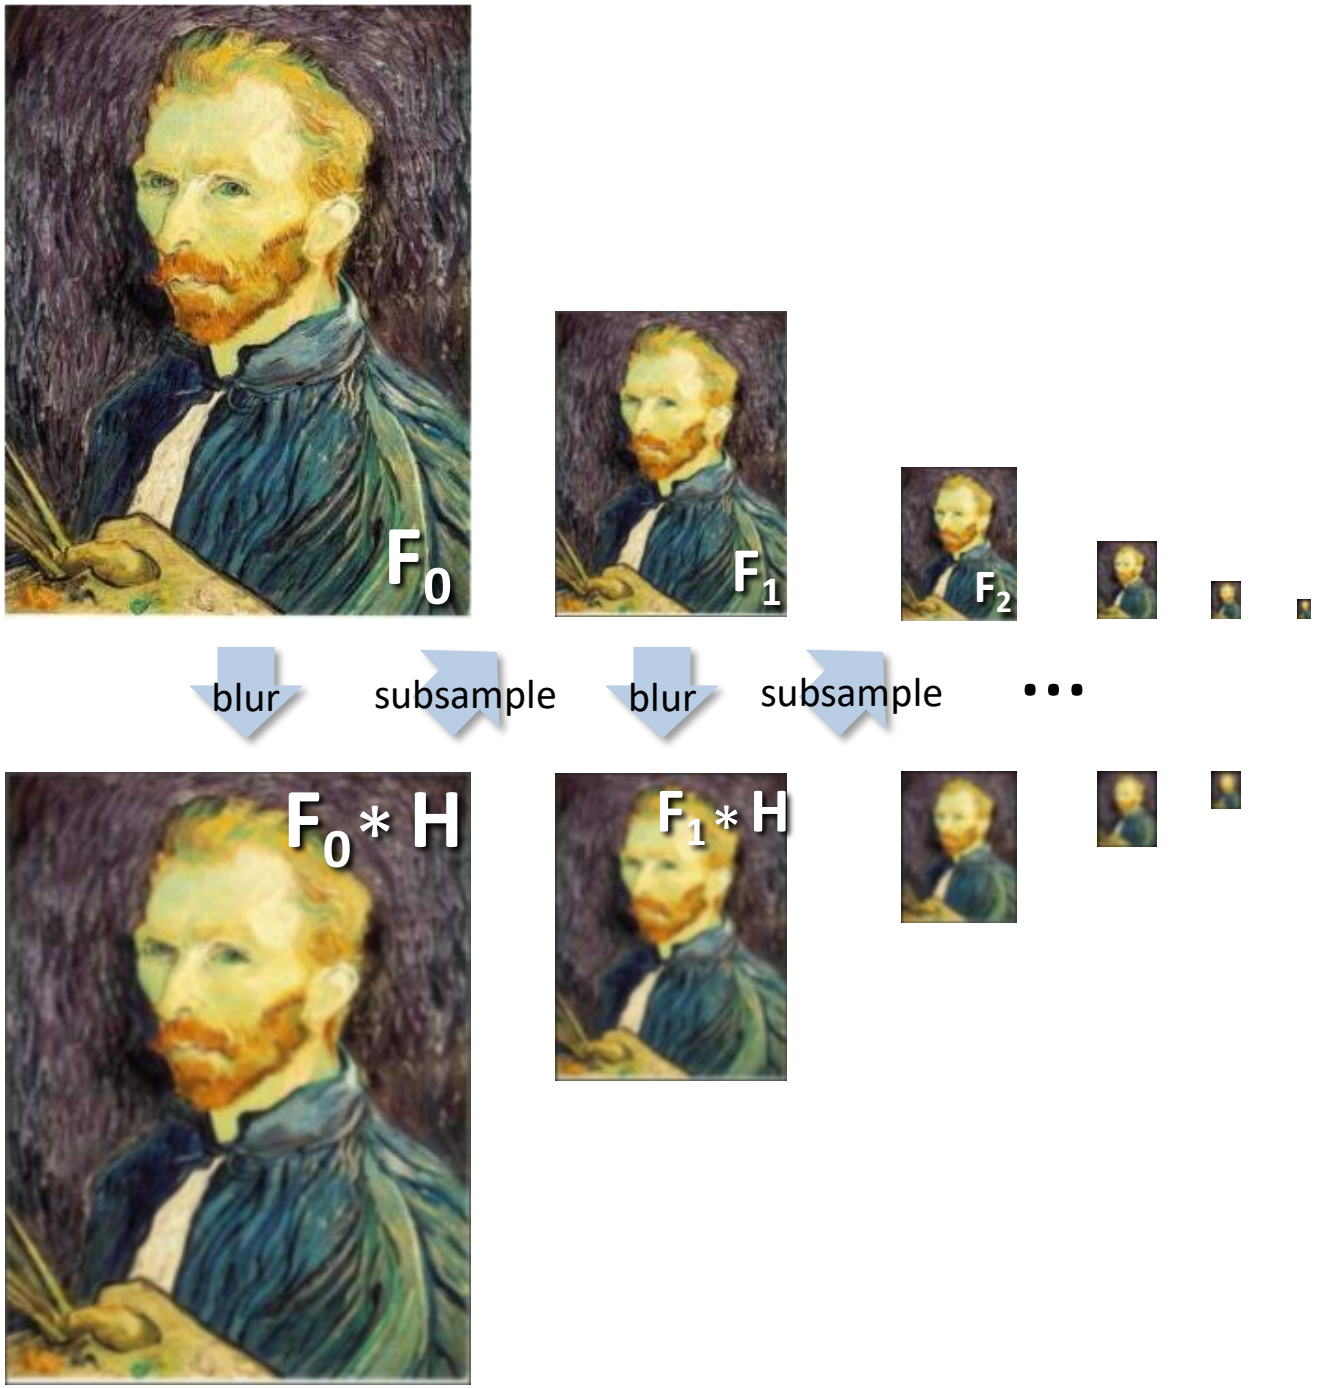

Gaussian pyramid

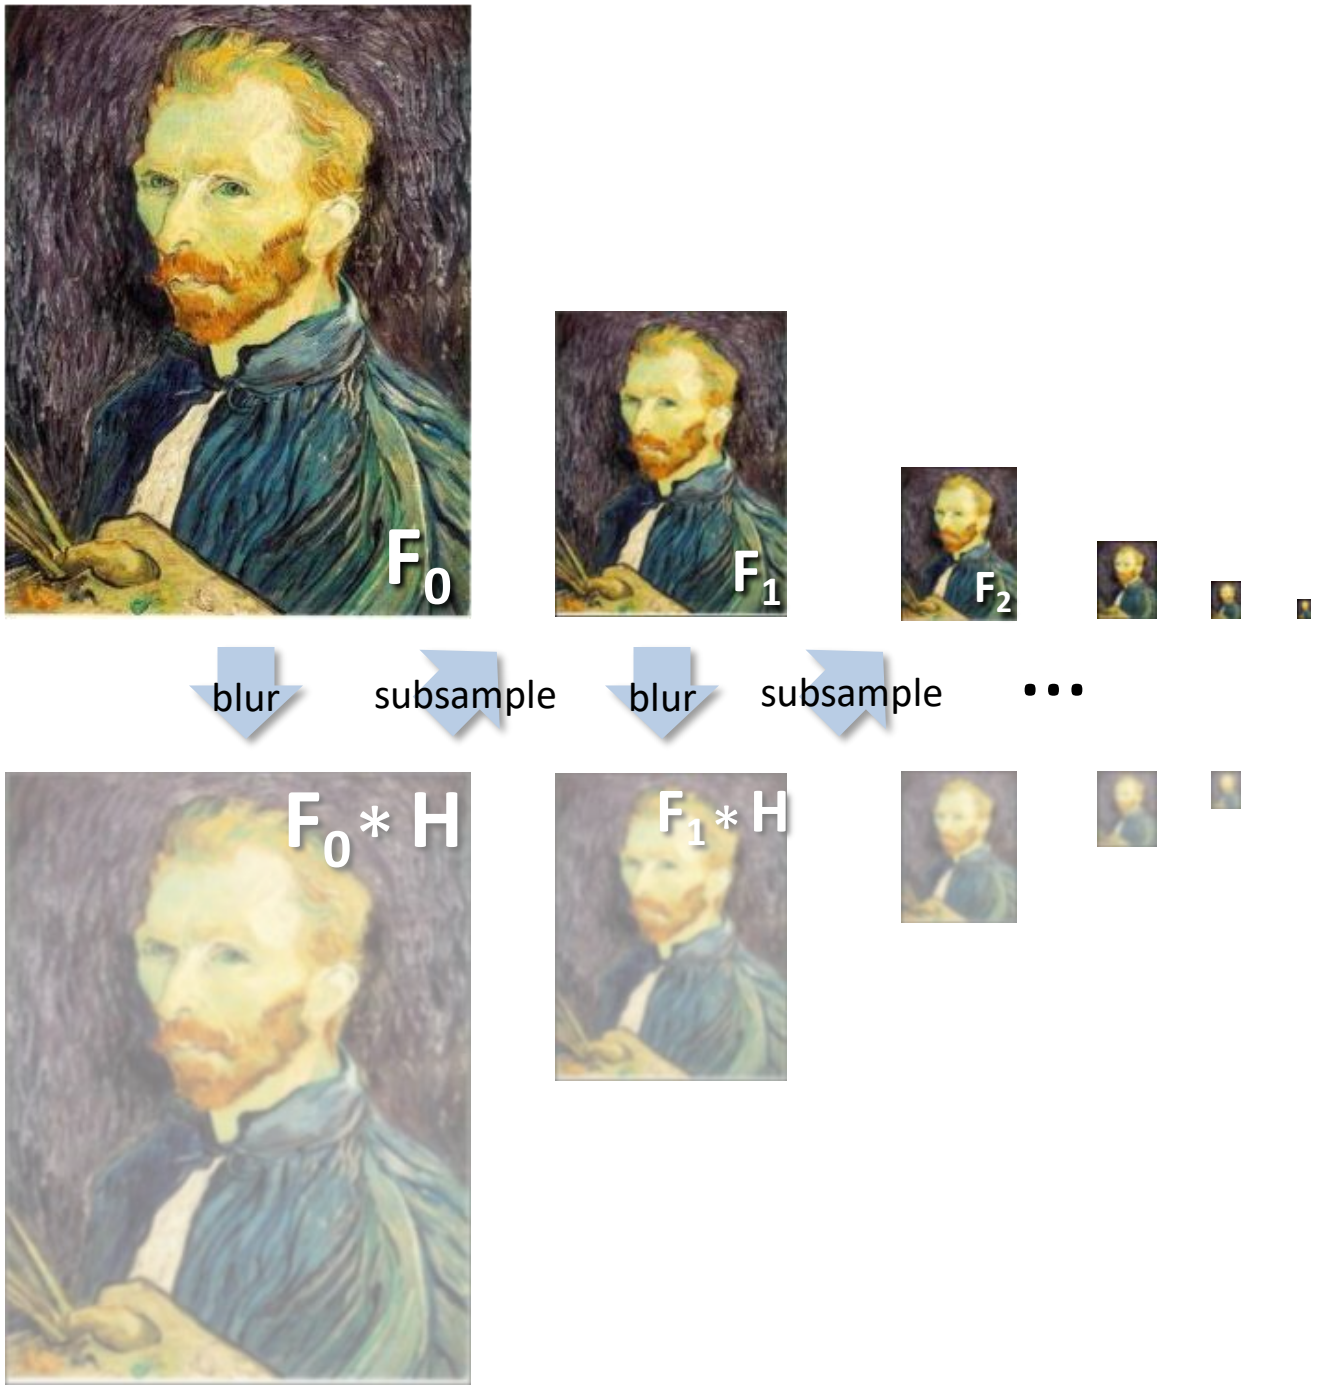

Laplace Pyramid

- Derive from Gaussian pyramid
  - G1=pydn(G0); G2=pydn(G1), …
  - One level of laplace pyramid is difference between approximated and original Gaussian pyramid levels
  - L0=G0-pyup(G1); L1 = G1-pyup(G2)

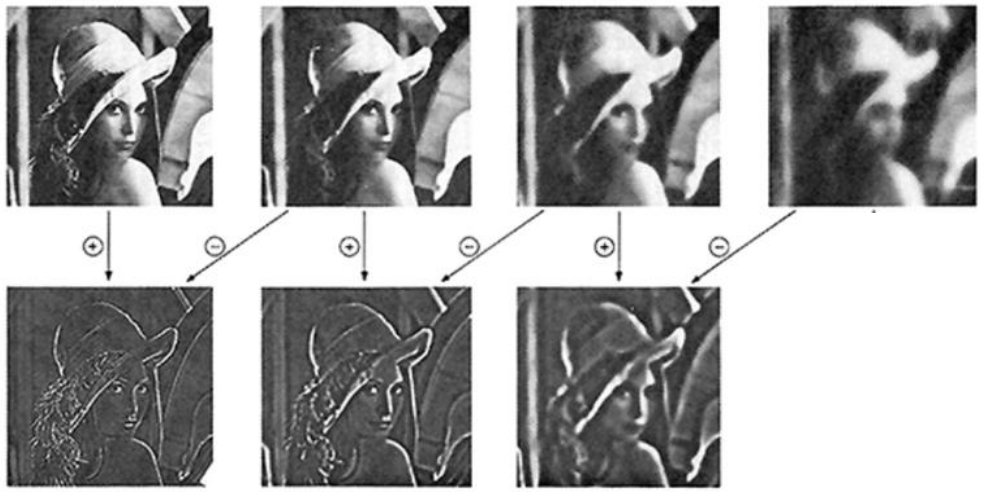

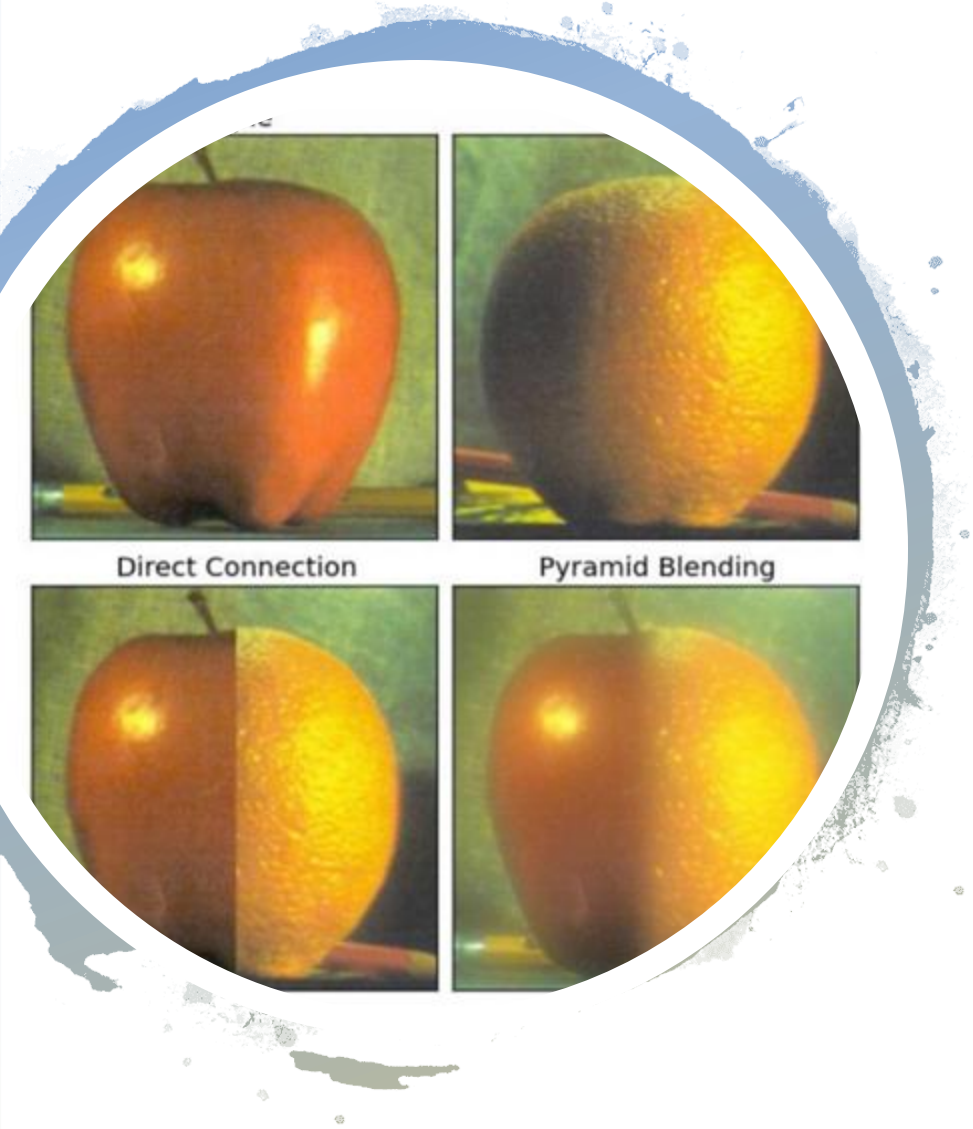

Image composting

- Generate L-pyramid of orange
- Generate L-pyramid of apple
- Combine two pyramids
  - For all levels, one half from one pyramid, the other half from another
- Reconstruct image from combine pyramid