Image Interpolation

Image Processing

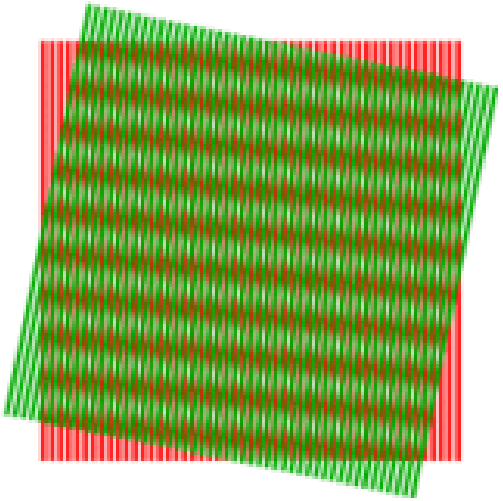

Brent M. Dingle, Ph.D.
Game Design and Development Program
Mathematics, Statistics and Computer Science
University of Wisconsin - Stout

2015

public domain image from wikipedia

Lecture Objectives

- Previously
  - What a Digital Image is
  - Acquisition of Digital Images
  - Human Perception of Digital Images
  - Digital Representation of Images
  - Various HTML5 and JavaScript Code
    - Pixel manipulation
    - Image Loading
    - Filtering
- Today
  - Image Interpolation
    - Sub-Sampling
      - Aliasing
      - Gaussian Blur Fix
    - Up-Sampling
      - Interpolation
  - Examples of Using Interpolation

Lecture Objectives

- Previously
  - What a Digital Image is
  - Acquisition of Digital Images
  - Human Perception of Digital Images
  - Digital Representation of Images
  - Various HTML5 and JavaScript Code
    - Pixel manipulation
    - Image Loading
    - Filtering
- Today
  - Image Interpolation
    - Sub-Sampling
      - Aliasing
      - Pre-Filtering
    - Up-Sampling
      - Interpolation
  - Examples of Using Interpolation

Image

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

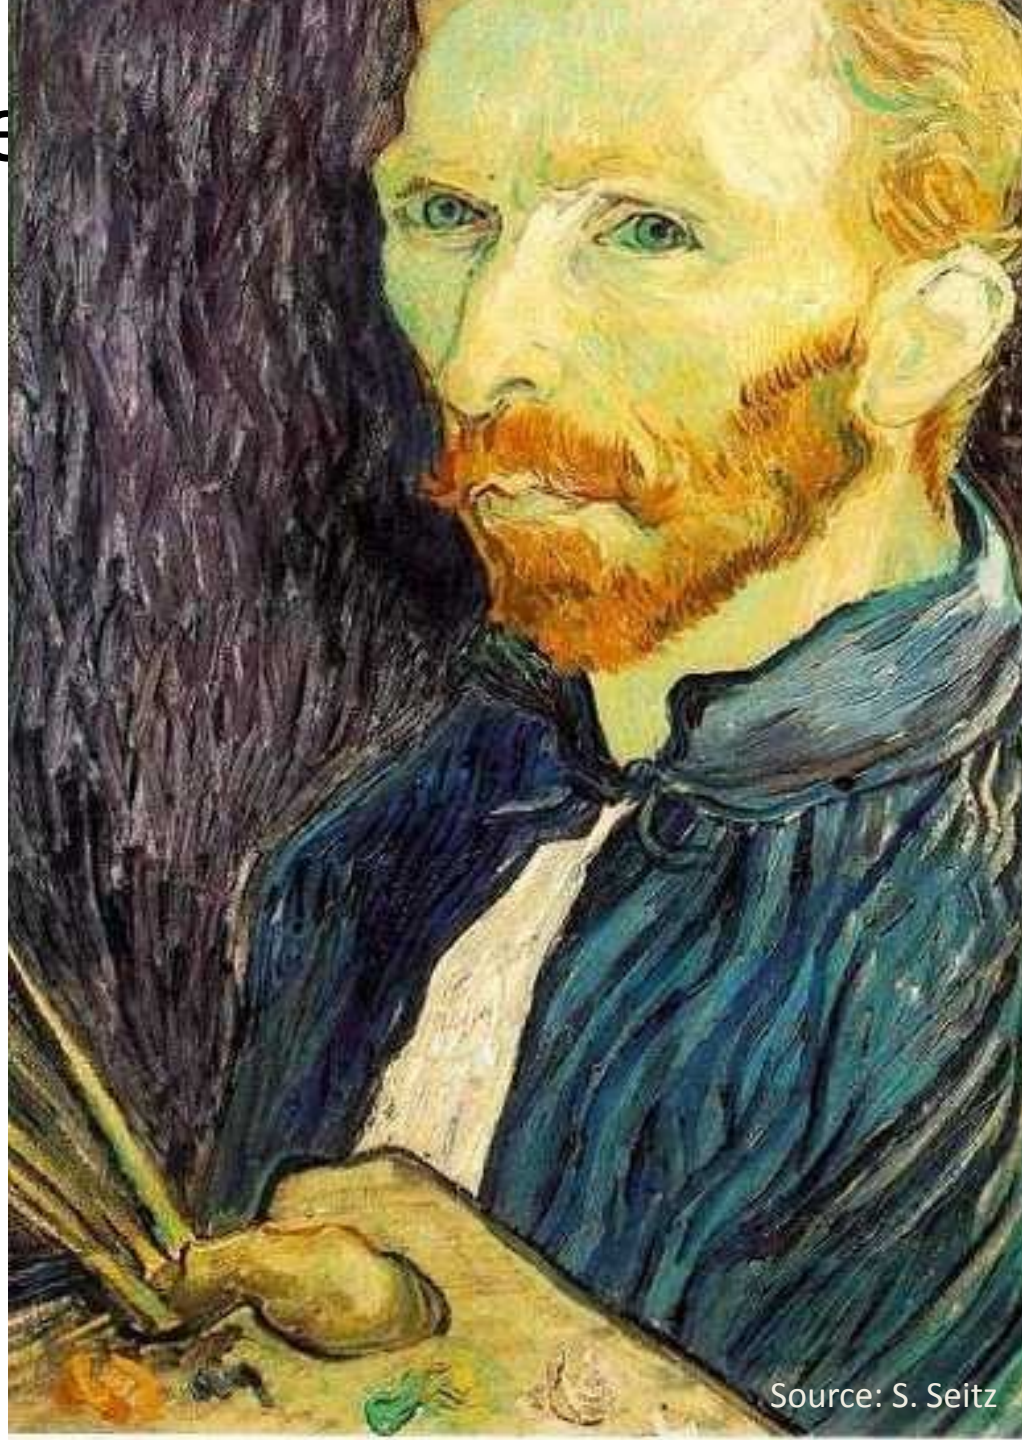

Slide from: Noah Snavely, 2012, Cornell

Image sub-sampling

Why do these reduced sizes look bad?

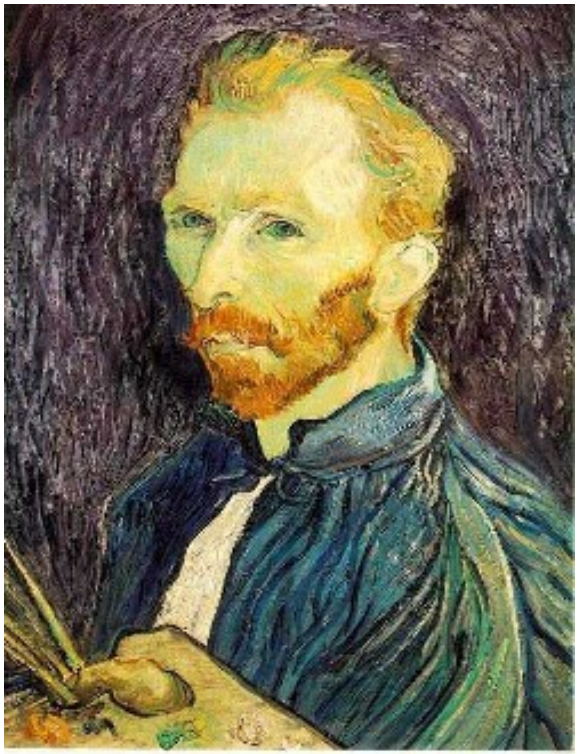

scaled to ½ size

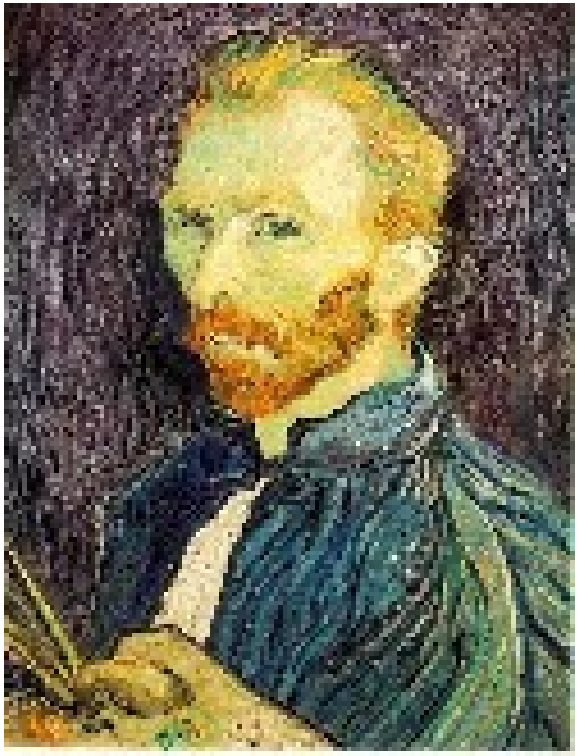

scaled to ¼ size
(2x zoom)

scaled to 1/8 size
(4x zoom)

Slide from: Noah Snavely, 2012, Cornell

Source: S. Seitz

Image sub-sampling

What happened here?

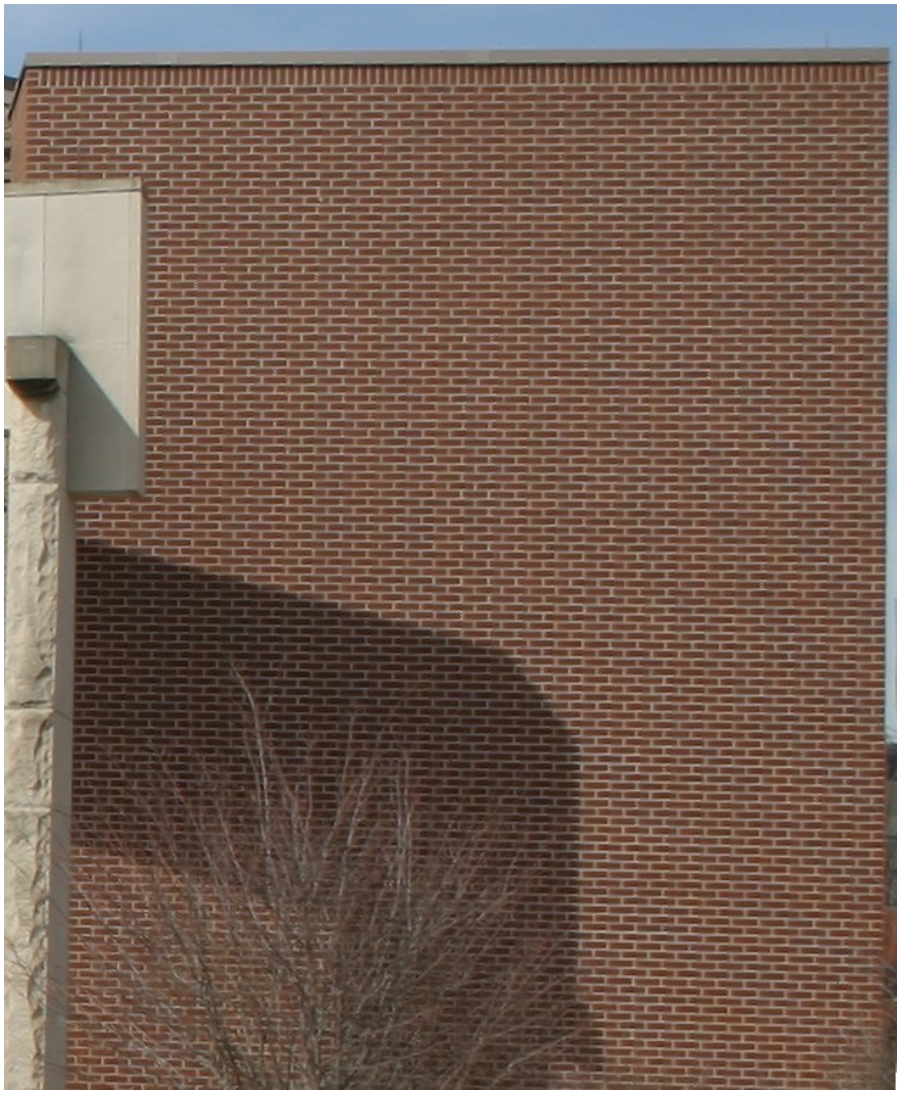

public domain images from wikipedia

Even worse for synthetic images

And what happened here?

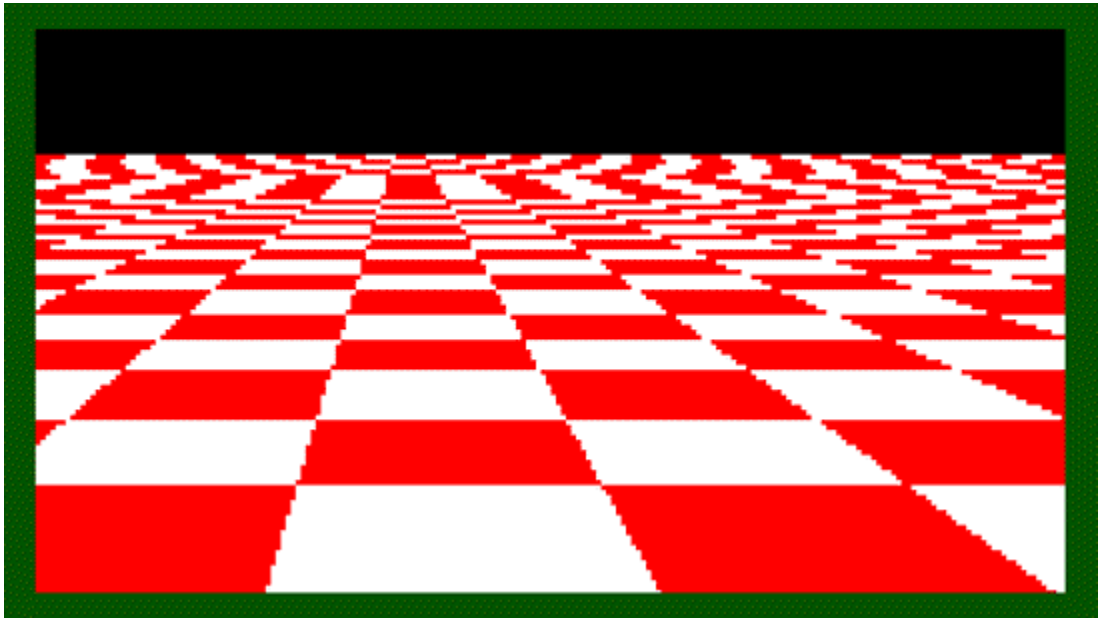

Source: L. Zhang

Aliasing

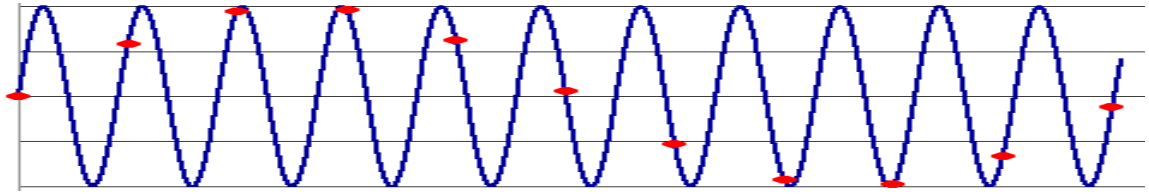

- Occurs when your sampling rate is not high enough to capture the amount of detail in your image
  - Can give you the **wrong signal/image—an *alias***
- To do sampling right,
  need to understand the structure of your signal/image
  - **Hello Fourier!**
- To avoid aliasing:
  - sampling rate ≥ 2 * max frequency in the image
    - said another way: ≥ two samples per cycle
  - This minimum sampling rate is called the **Nyquist rate**

Source: L. Zhang

Wagon-wheel effect

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):

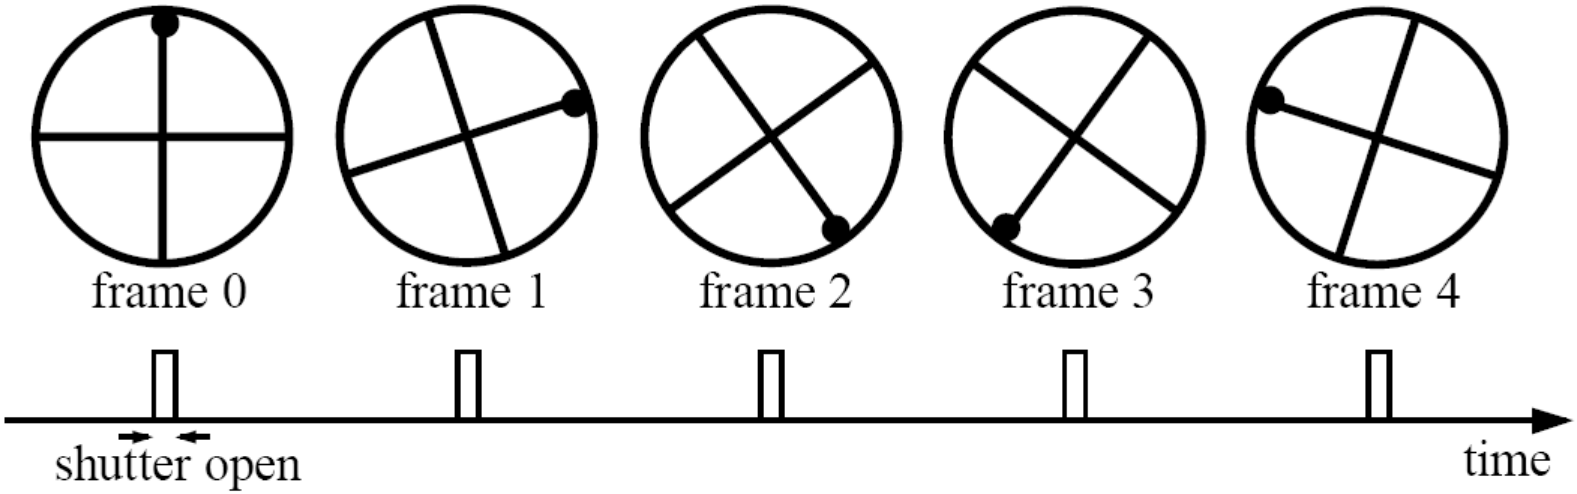

Without dot, wheel appears to be rotating slowly backwards! (counterclockwise)

(See http://www.michaelbach.de/ot/mot_wagonWheel/index.html)

Source: L. Zhang

Wagon Wheel Variations

Which way is the train going?

Which way does the middle dancer spin?
Focus on the left – middle matches
Focus on the right – middle matches

Images from: http://stanleykemp.blogspot.com/2015/06/lets-start-weekend-with-couple-of_6.html

Nyquist limit – 2D example

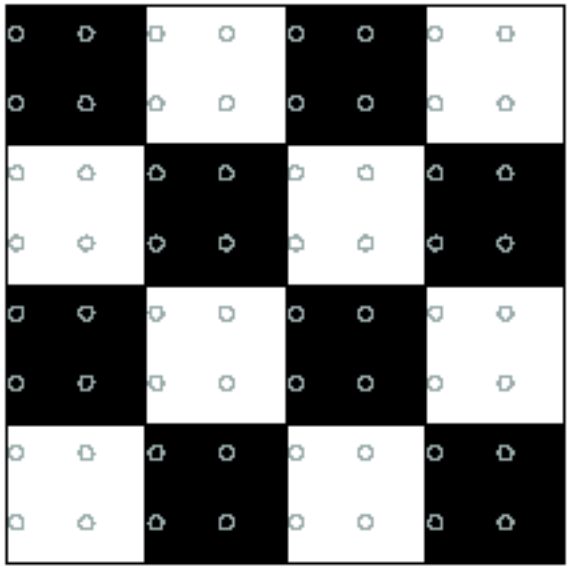

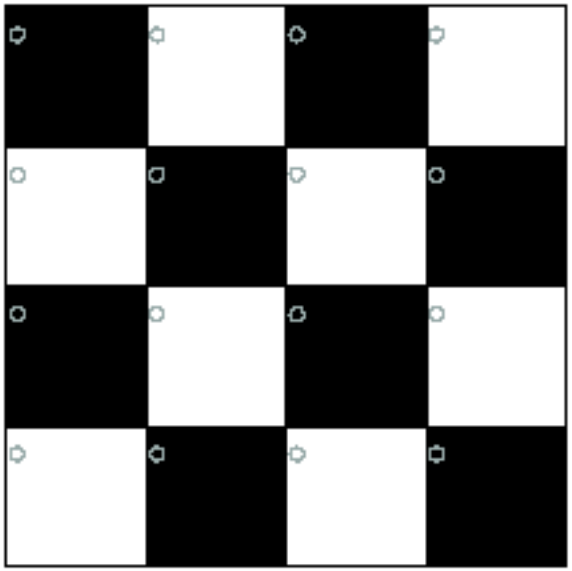

Good sampling

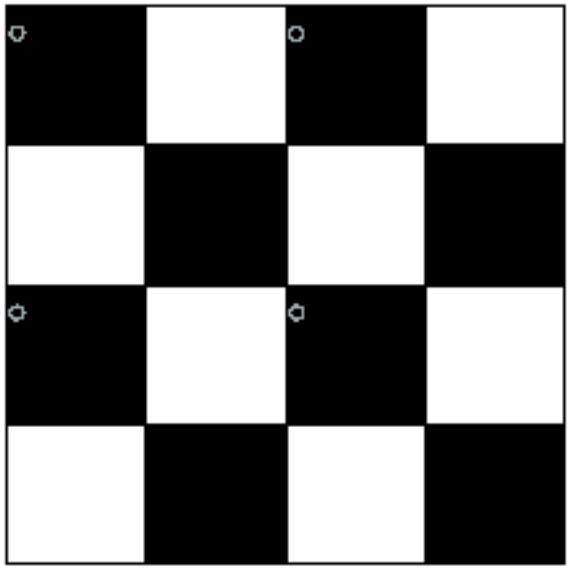

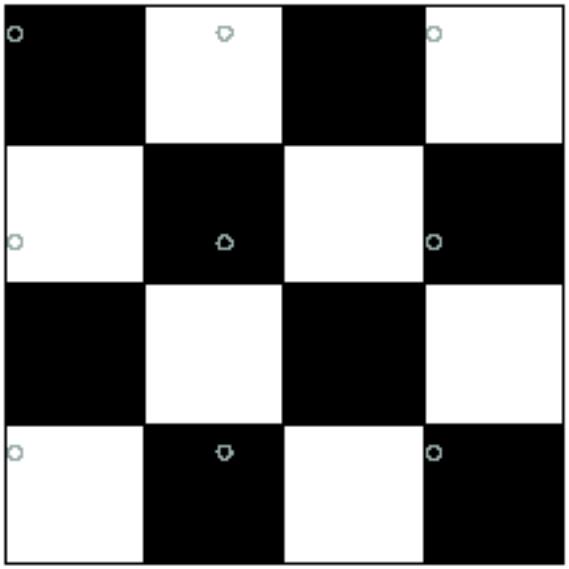

Bad sampling

sampling rate < 2*(max freq)

Aliasing

- When downsampling by a factor of two
  - Original image has frequencies that are too high

- How can we fix this?

Gaussian pre-filtering

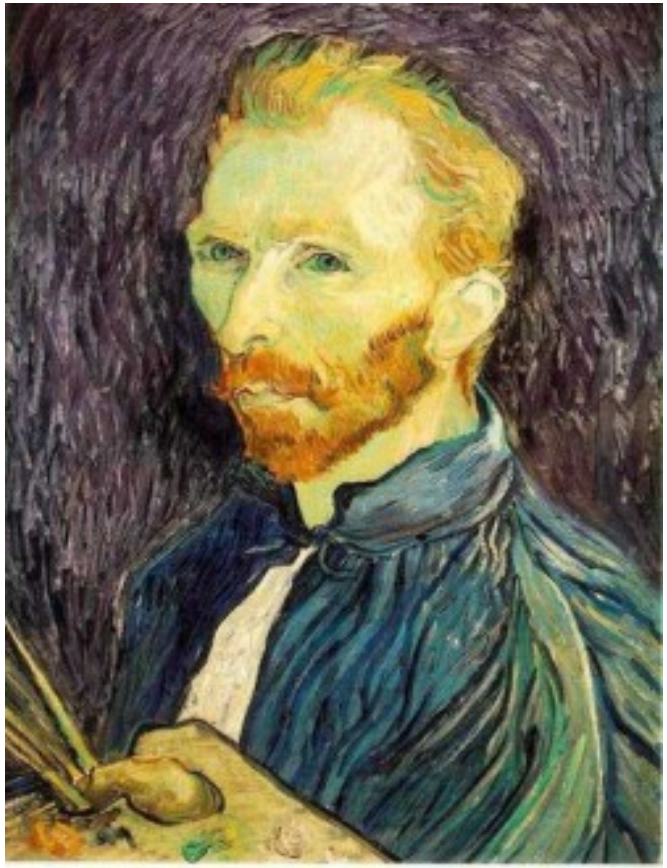

Gaussian 1/2

G 1/4

G 1/8

- Solution: filter the image, *then* subsample

Source: S. Seitz

Subsampling with Gaussian pre-filtering

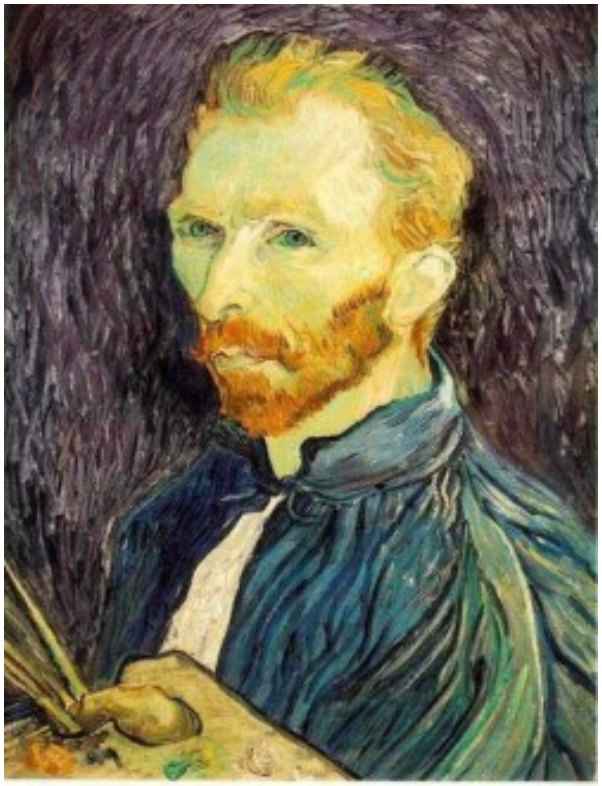

Gaussian 1/2

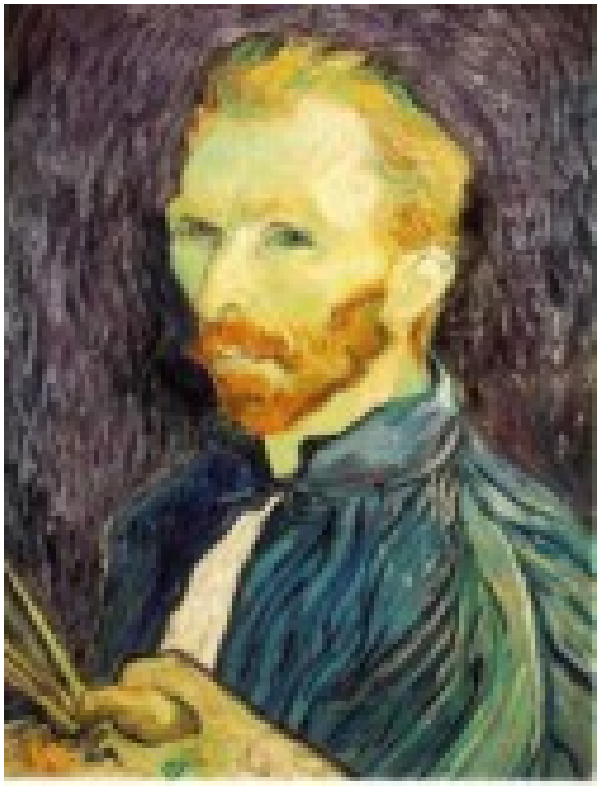

G 1/4

G 1/8

- Solution: filter the image, *then* subsample

Source: S. Seitz

Compare with...

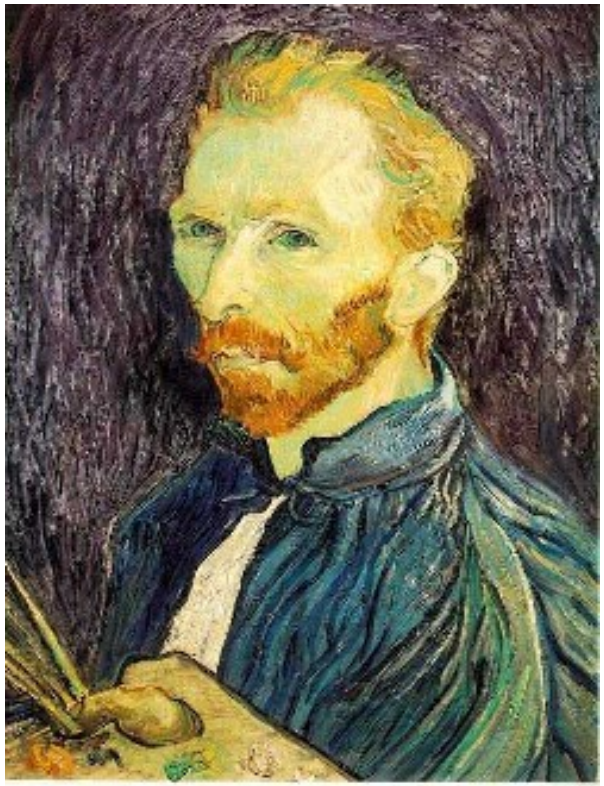

1/2

1/4 (2x zoom)

1/8 (4x zoom)

Source: S. Seitz

Gaussian pre-filtering

- Solution: filter the image, *then* subsample

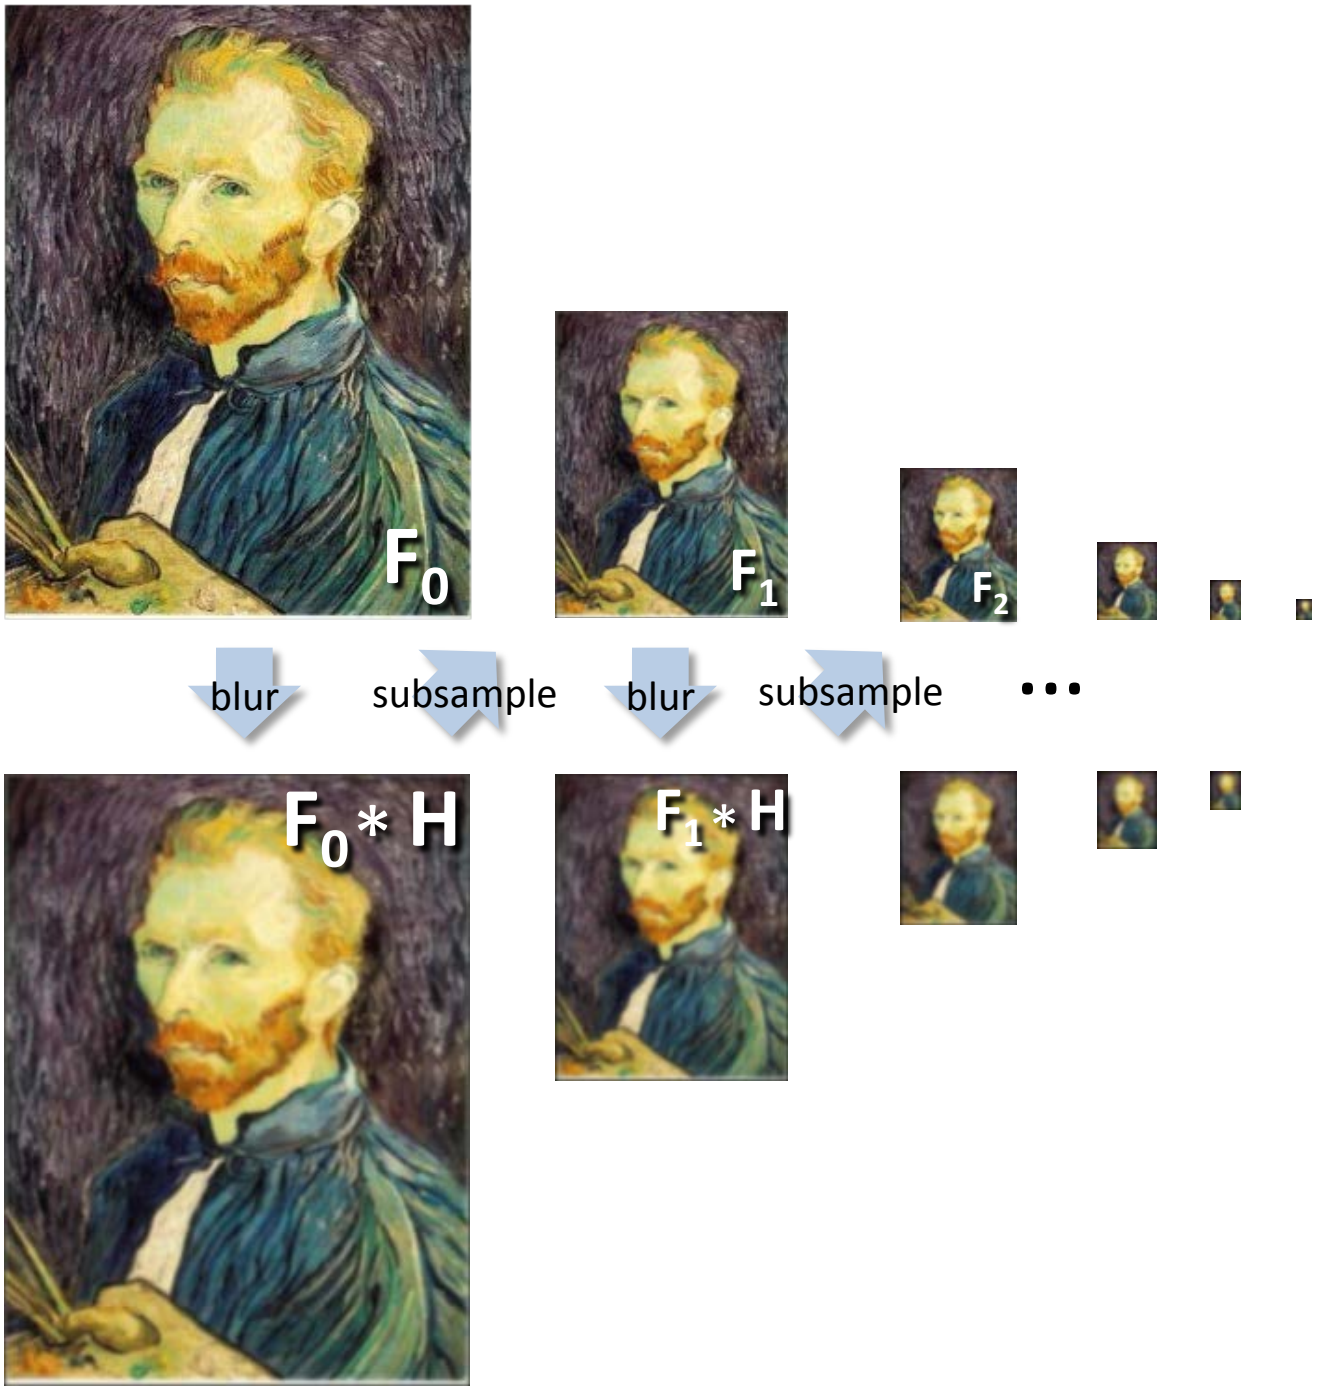

Gaussian pyramid

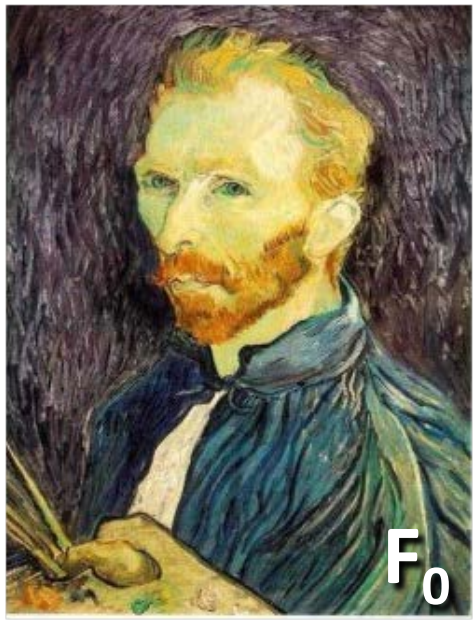

blur subsample blur subsample • • •

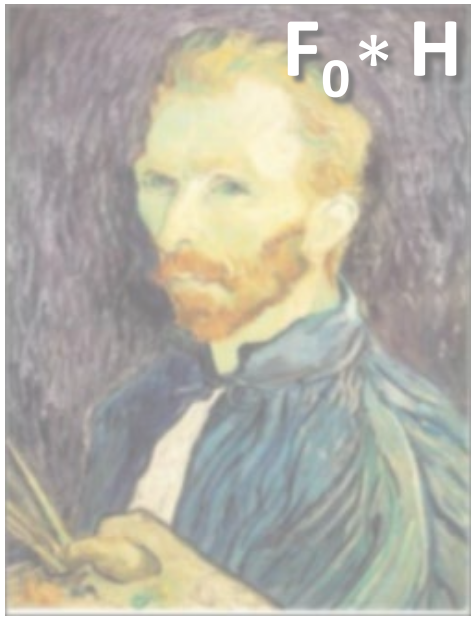

Gaussian pyramids [Burt and Adelson, 1983]

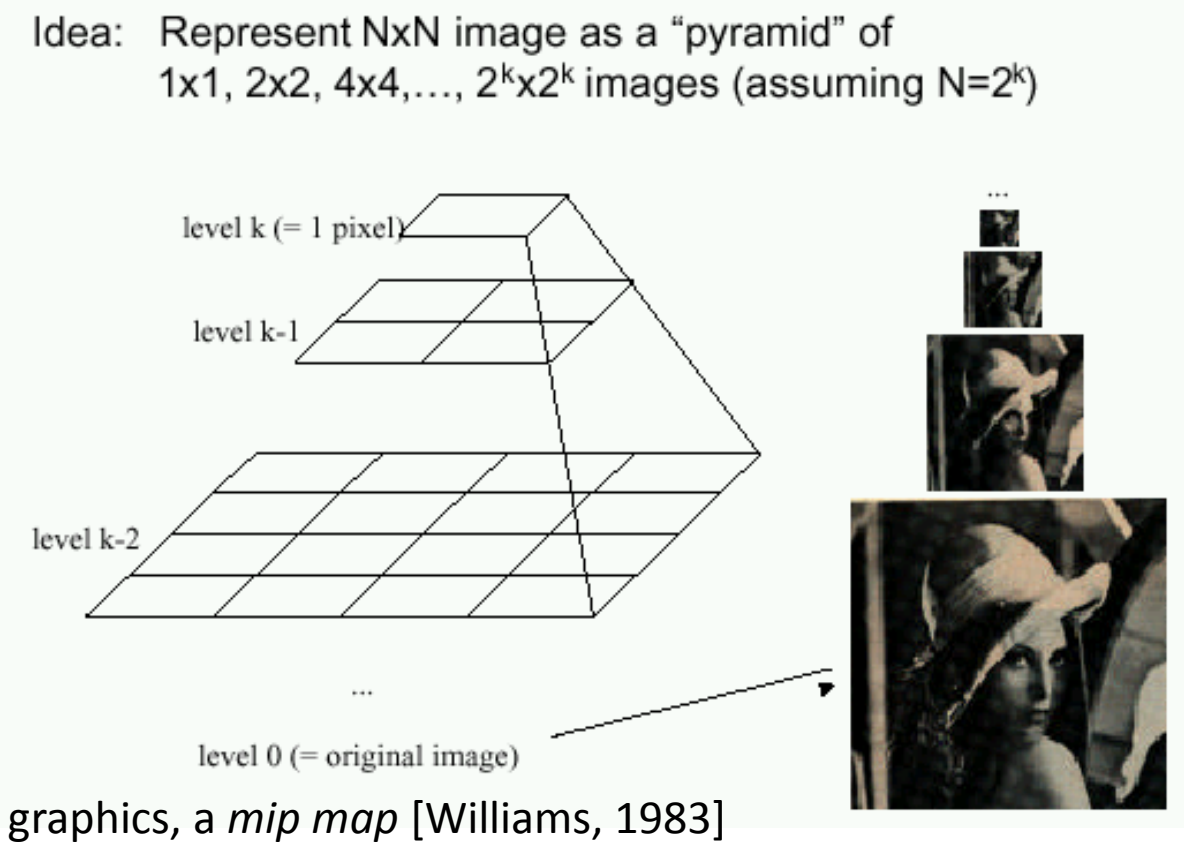

- In computer graphics, a *mip map* [Williams, 1983]
- A precursor to *wavelet transform*

Gaussian Pyramids have all sorts of applications in computer graphics, vision, imaging…

Source: S. Seitz

Gaussian pyramids [Burt and Adelson, 1983]

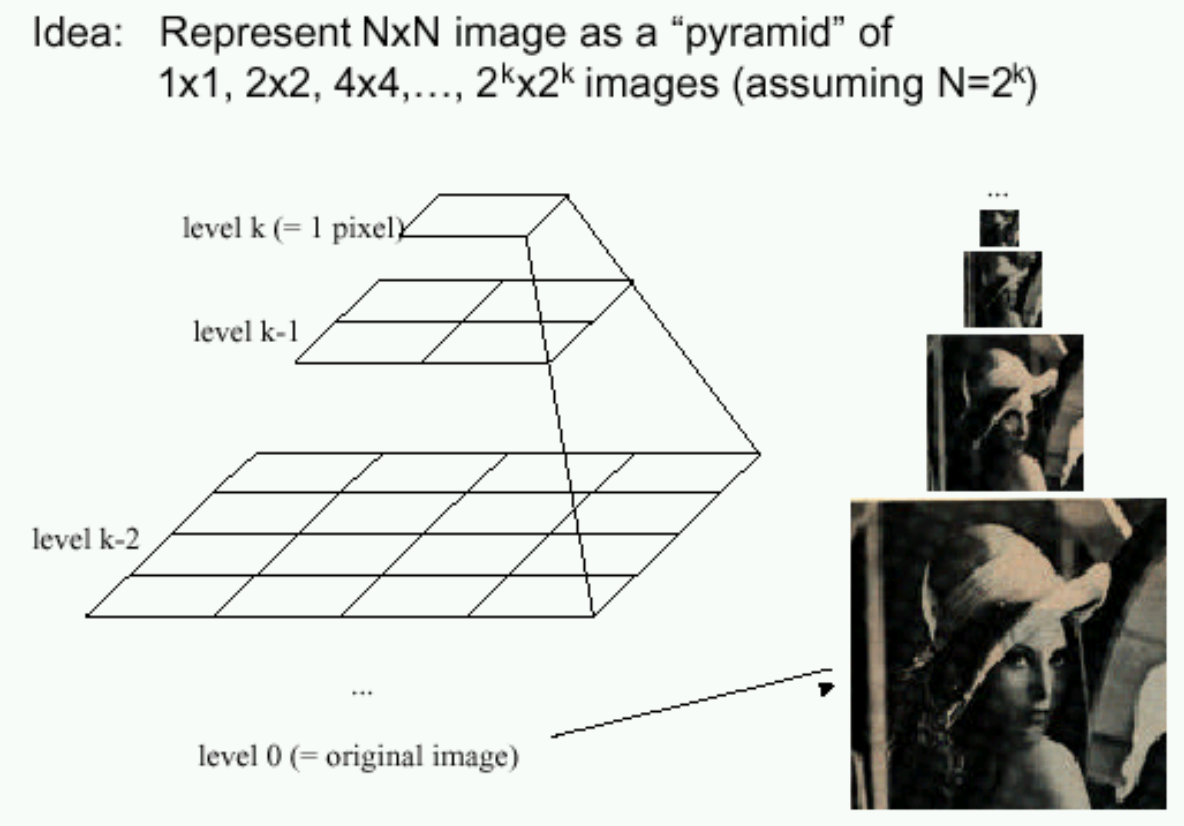

- How much space does a Gaussian pyramid take compared to the original image?

Source: S. Seitz

Gaussian Pyramid

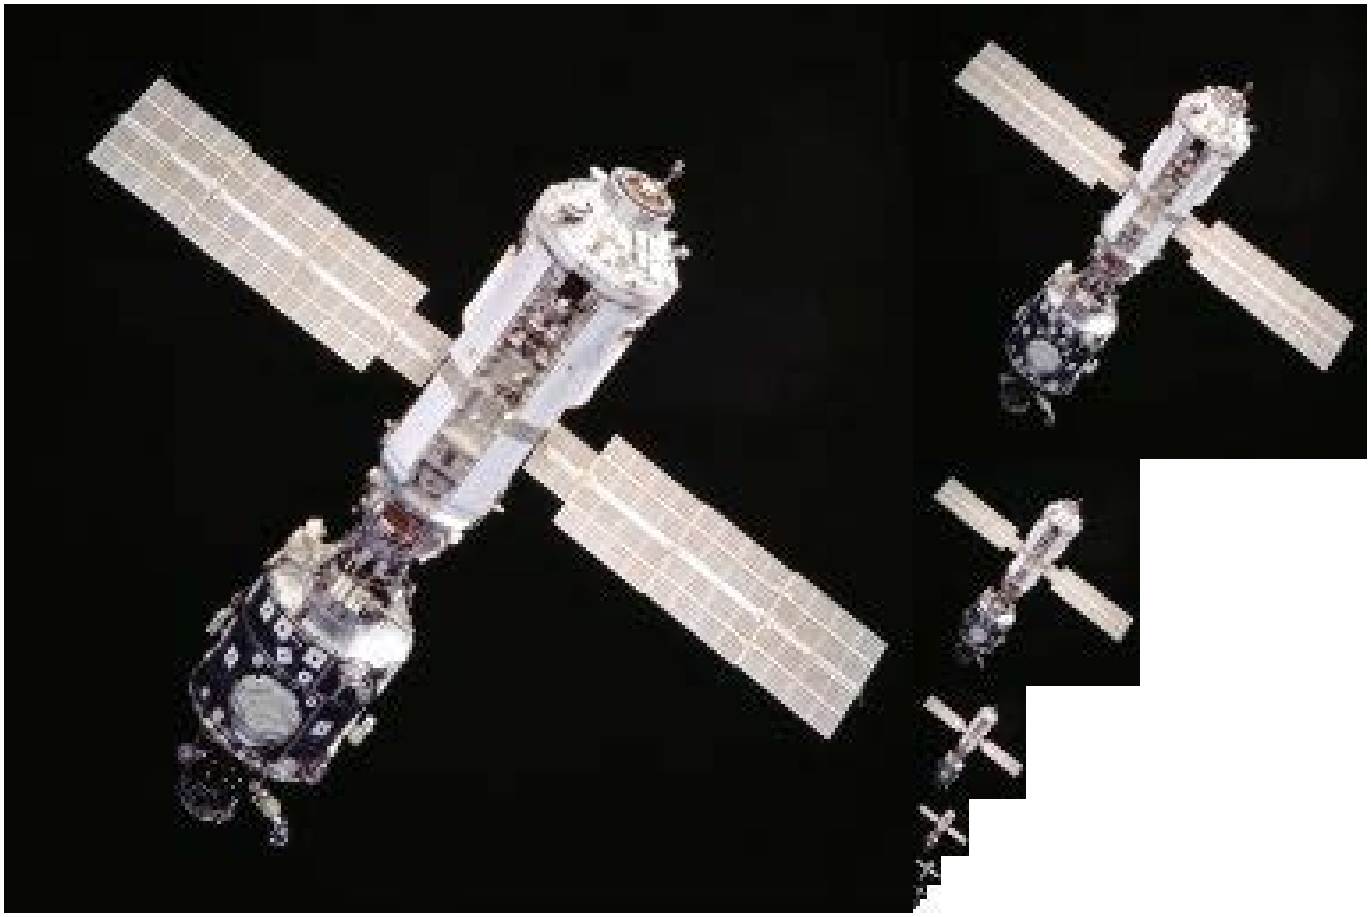

Questions so far?

- Questions on Sub-Sampling
  - Aliasing?
  - Gaussian pre-filter?

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")

Slide from: Noah Snavely, 2012, Cornell

Image interpolation

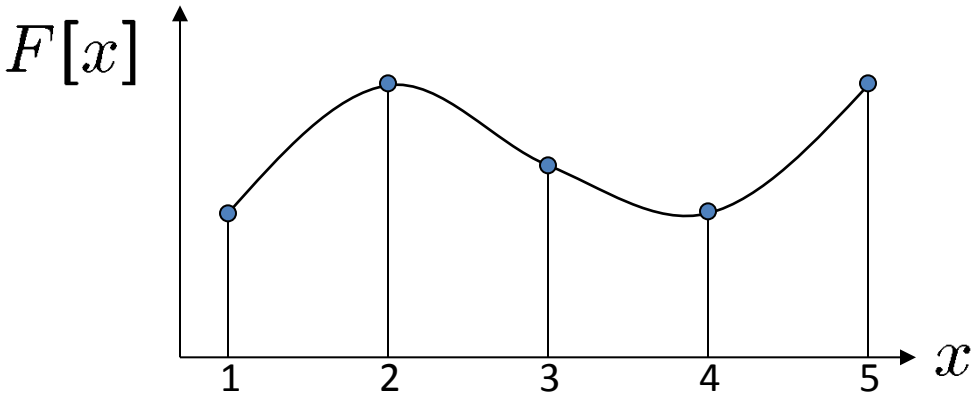

d = 1 in this example

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

Slide from: Noah Snavely, 2012, Cornell

Adapted from: S. Seitz

Image interpolation

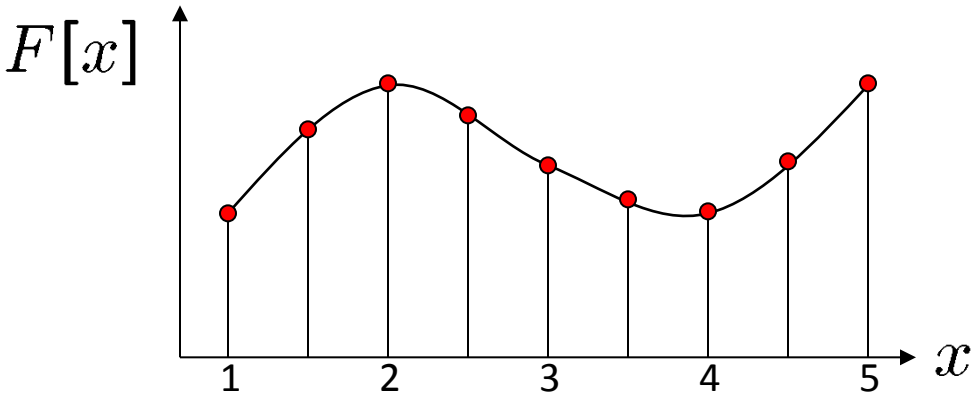

d = 1 in this example

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

Slide from: Noah Snavely, 2012, Cornell

Adapted from: S. Seitz

Image interpolation

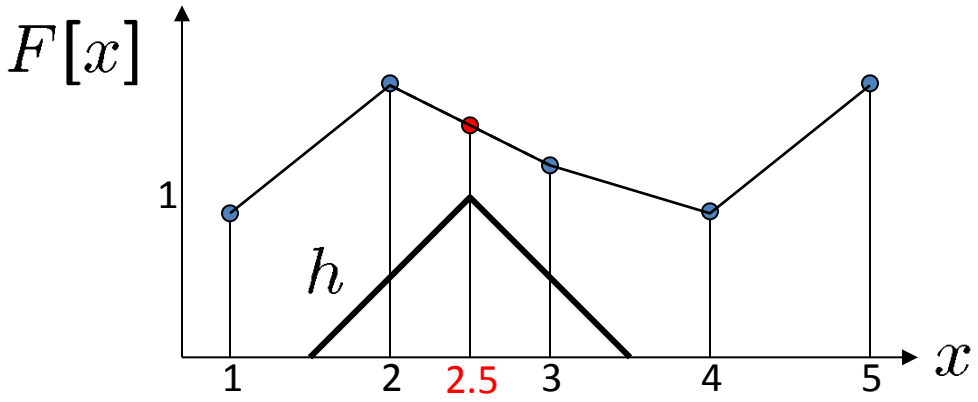

d = 1 in this example

- 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*

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

Slide from: Noah Snavely, 2012, Cornell

Adapted from: S. Seitz

Image interpolation

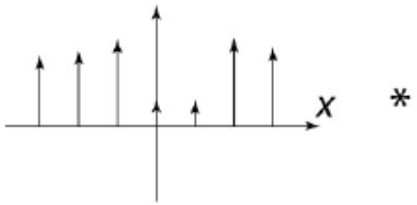

*

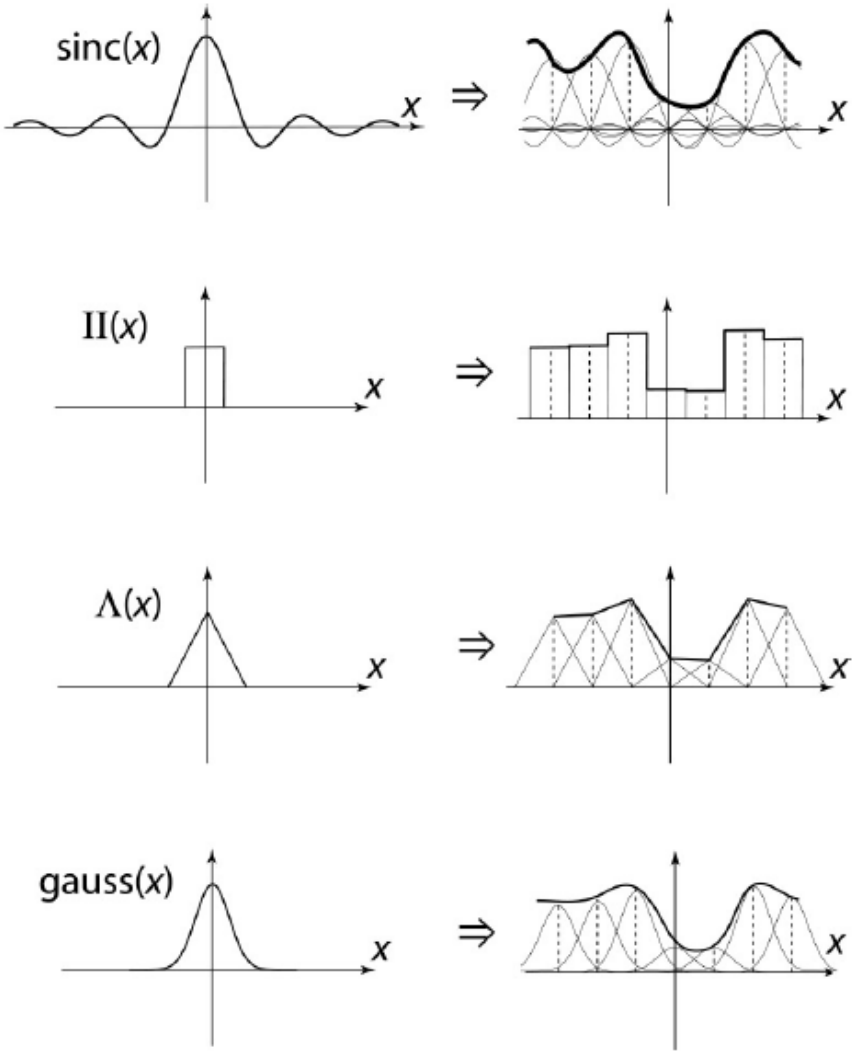

"Ideal" reconstruction

Nearest-neighbor interpolation

Linear interpolation

Gaussian reconstruction

Slide from: Noah Snavely, 2012, Cornell

Source: B. Curless

Reconstruction filters

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

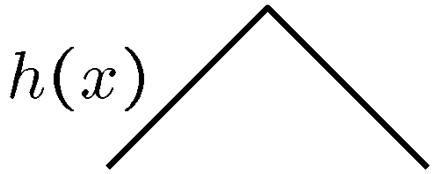

performs
linear interpolation

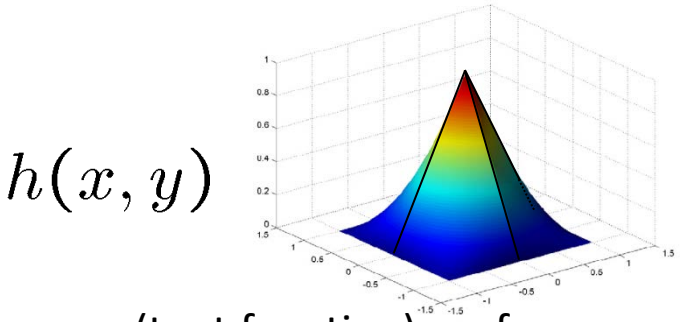

(tent function) performs
bilinear interpolation

3x3 Matrix for 2D tent filter might look like:

$$\frac{1}{16}\begin{bmatrix}1 & 2 & 1\\ 2 & 4 & 2\\ 1 & 2 & 1\end{bmatrix}$$

Derived from: Noah Snavely, 2012, Cornell

Reconstruction filters

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

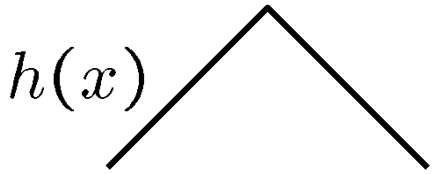

performs
linear interpolation

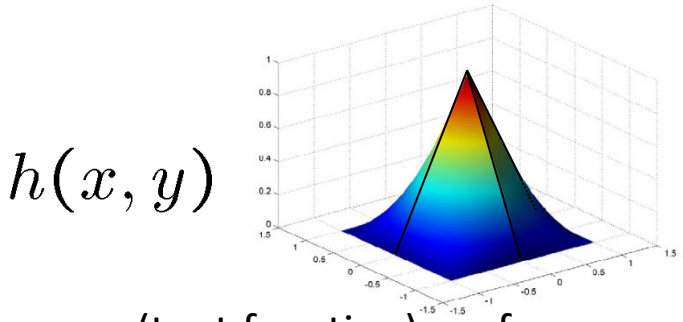

(tent function) performs
bilinear interpolation

More information, start at:

- http://en.wikipedia.org/wiki/Bilinear_interpolation

Better filters give better resampled images

- **Bicubic** is common choice

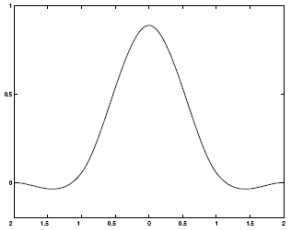

Cubic reconstruction filter

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

Slide from: Noah Snavely, 2012, Cornell

Image interpolation

Original image:  x 10

Nearest-neighbor interpolation

Bilinear interpolation

Bicubic interpolation

Slide from: Noah Snavely, 2012, Cornell

Image interpolation

Also used for *resampling*

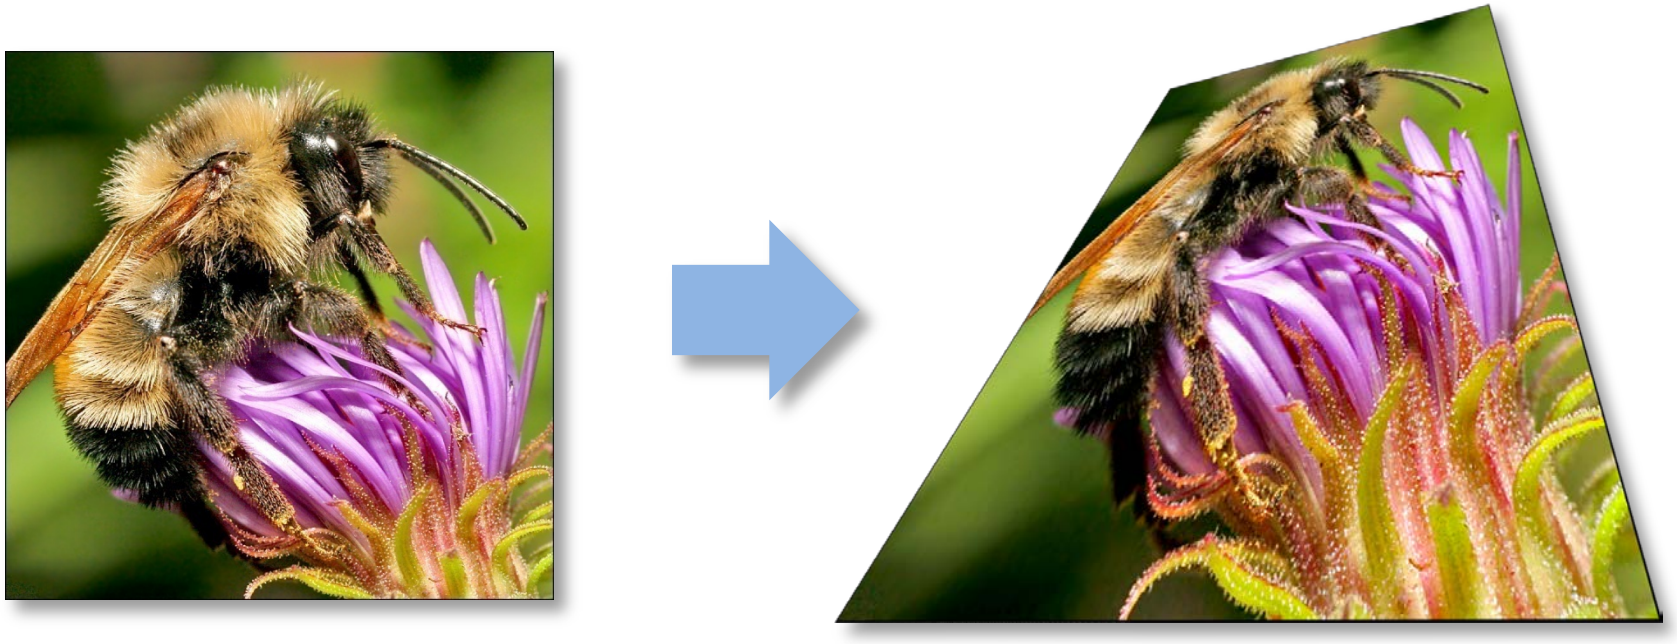

Slide from: Noah Snavely, 2012, Cornell

Examples of Interpolation Usage

- Image Bayer De-mosaicking
- Image Error Correction/Fill
- Image Reconstruction/Inpainting
- Merging/Mosaicking/Panoramic Splicing

Demosaicing Image Example

Ad-hoc CFA Interpolation

Advanced CFA Interpolation

Error Correction/Fill

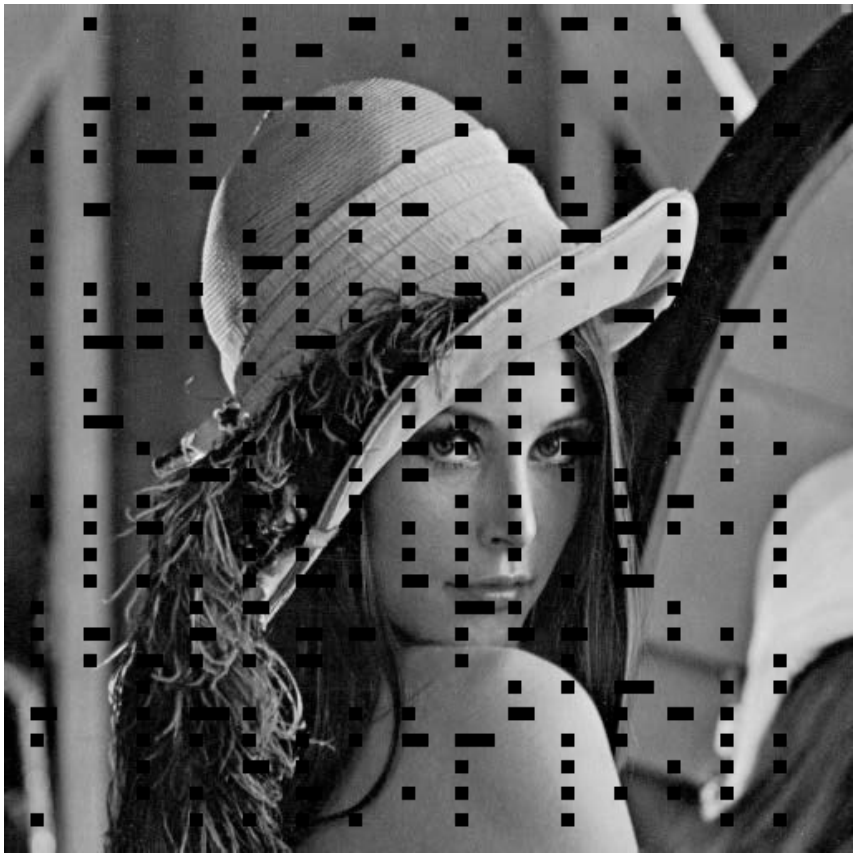

damaged

interpolated

Image Inpainting/Restoration

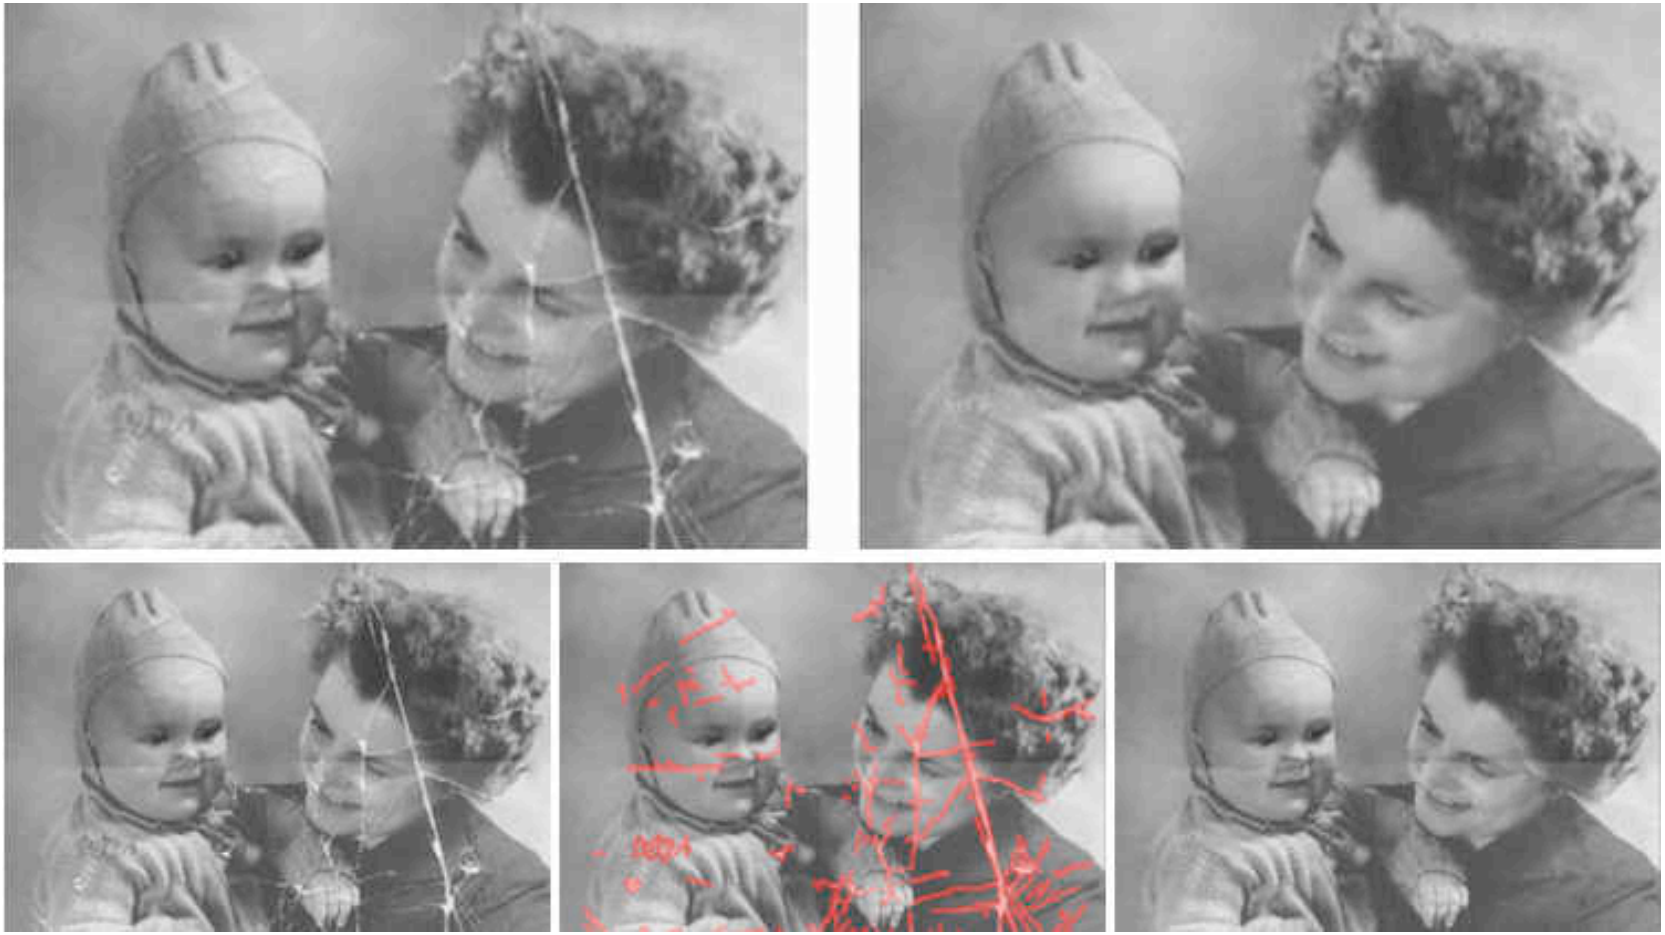

Image Mosaicing/Merging

Summary: Image Interpolation

- Image Interpolation
  - Fundamental tool in digital processing of images
    - bridges the continuous world and the discrete world
  - Wide range of application use

Challenge: Bayer De-Mosaicking

- Separate an image into its color planes (RGB)

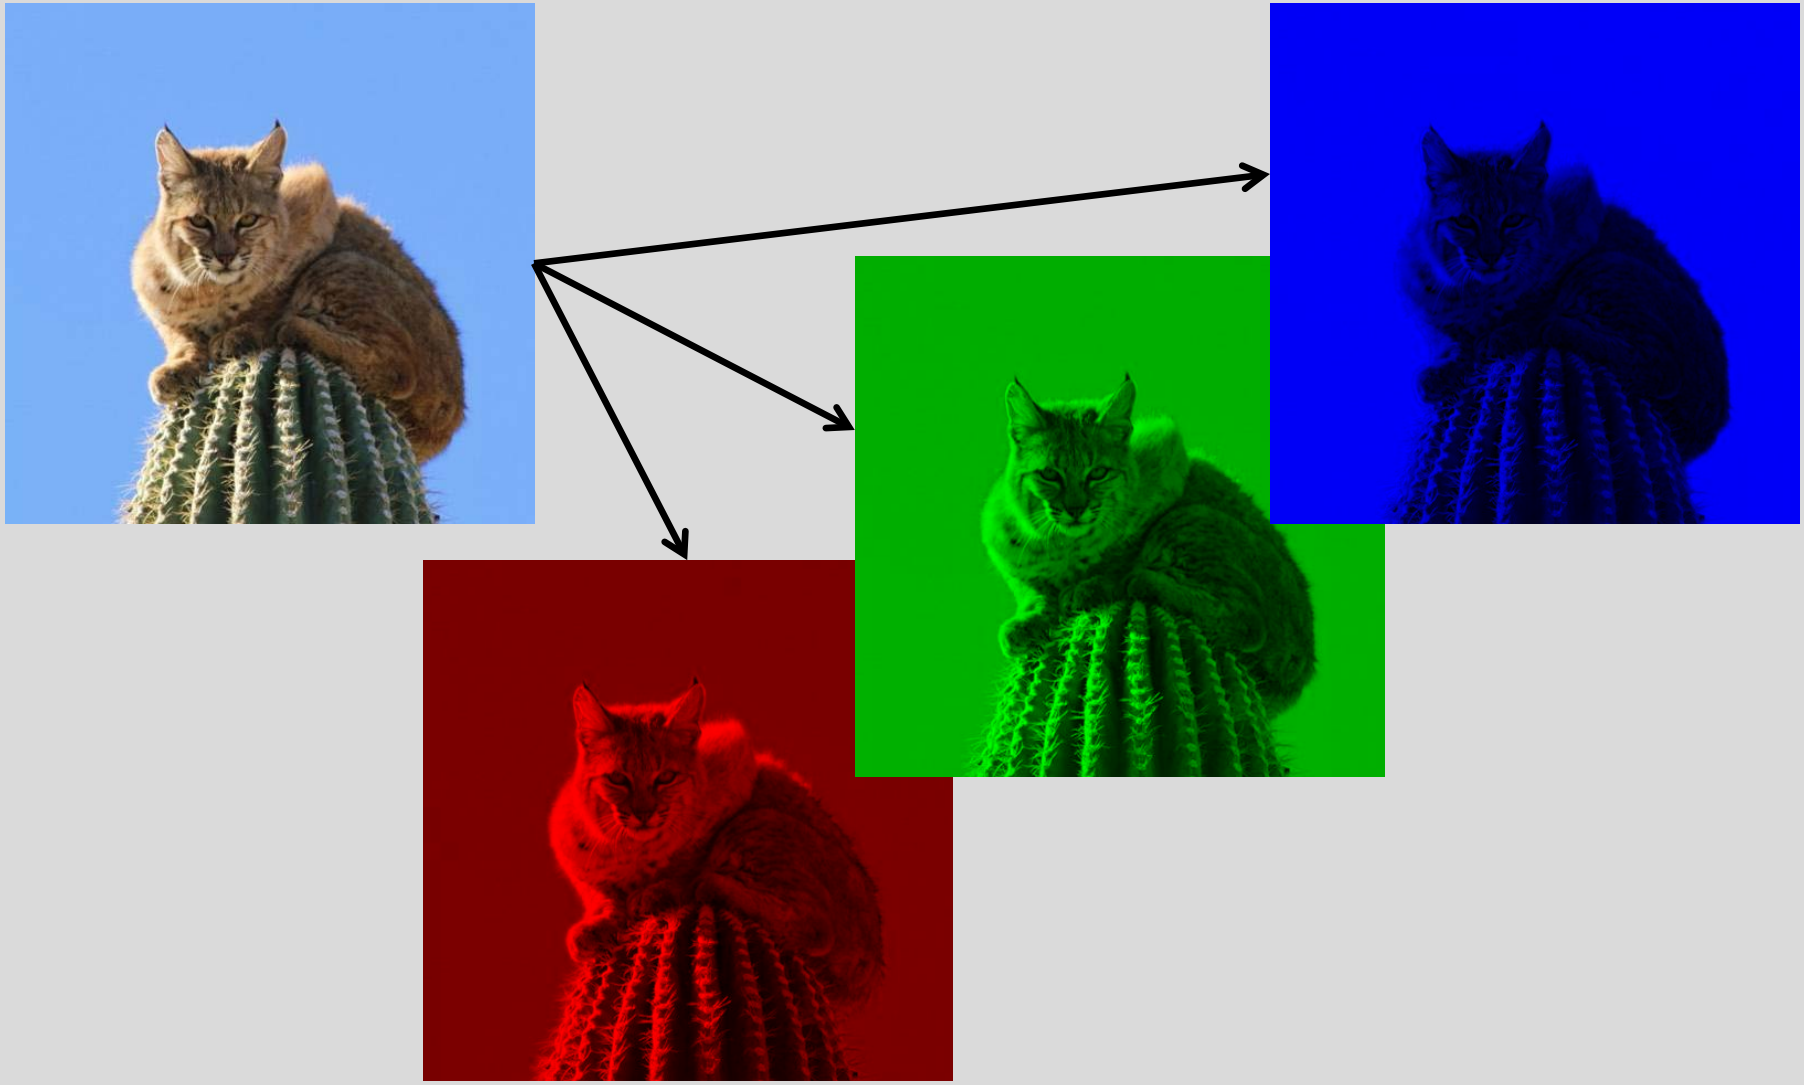

Challenge: Bayer De-Mosaicking

- Reduce panes to Bayer Pattern Equivalent

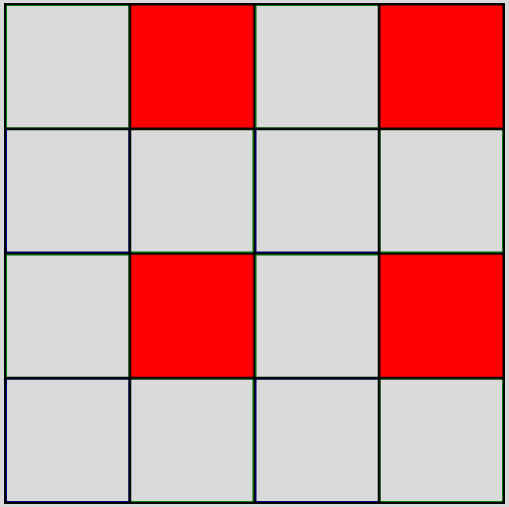

Challenge: Bayer De-Mosaicking

- Reduce panes to Bayer Pattern Equivalent

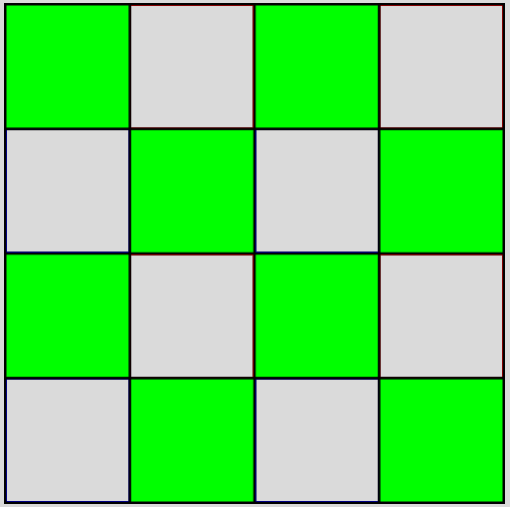

Challenge: Bayer De-Mosaicking

- Reduce panes to Bayer Pattern Equivalent

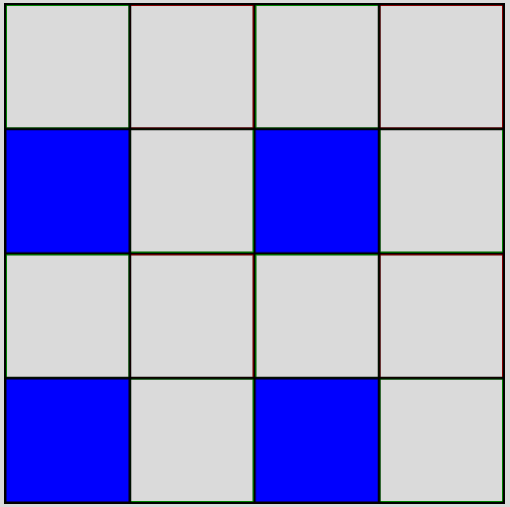

Challenge: Bayer De-Mosaicking

- Interpolate each pane and re-composite

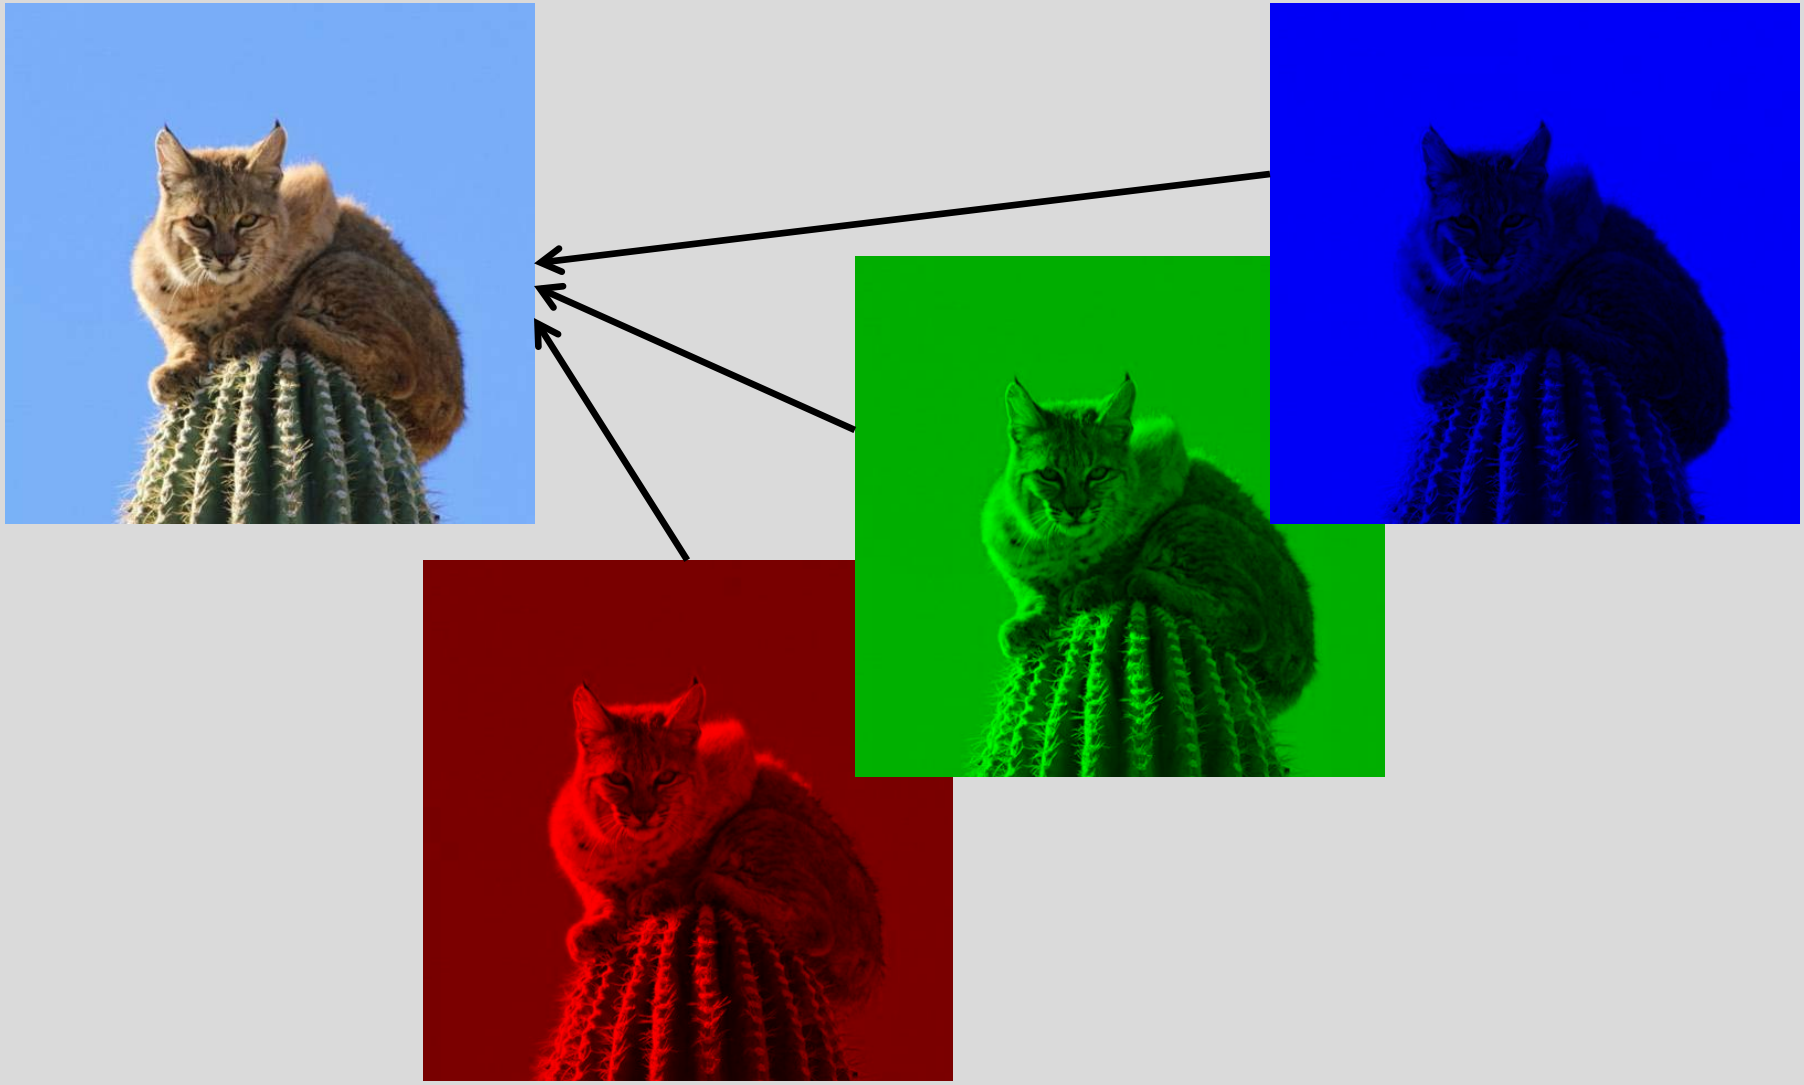

Questions?

- Beyond D2L
  - Examples and information can be found online at:
    - *http://docdingle.com/teaching/cs.html*

    - *Continue to more stuff as needed*

Extra Reference Stuff Follows

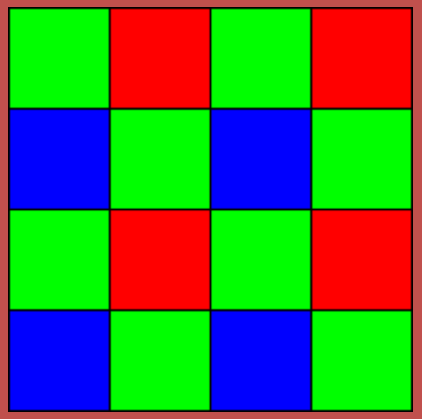

Credits

- Much of the content derived/based on slides for use with the book:
  - *Digital Image Processing,* Gonzalez and Woods

- Some layout and presentation style derived/based on presentations by
  - Donald House, Texas A&M University, 1999
  - Bernd Girod, Stanford University, 2007
  - Shreekanth Mandayam, Rowan University, 2009
  - Noah Snavely, Cornell University, 2012
  - Igor Aizenberg, TAMUT, 2013
  - Xin Li, WVU, 2014
  - George Wolberg, City College of New York, 2015
  - Yao Wang and Zhu Liu, NYU-Poly, 2015
  - Sinisa Todorovic, Oregon State, 2015