More Image Manipulation

Prof. Graeme Bailey

http://cs1114.cs.cornell.edu

(notes modified from Noah Snavely, Spring 2009)

Cornell University
Computer Science

Last time: image transformations

2D Linear Transformations

- Can be represented with a 2D matrix

$$T = \begin{bmatrix} a & b \\ c & d \end{bmatrix}$$

- And applied to a point using matrix multiplication

$$\begin{bmatrix} a & b \\ c & d \end{bmatrix} \begin{bmatrix} x \\ y \end{bmatrix} = \begin{bmatrix} ax + by \\ cx + dy \end{bmatrix}$$

Forward mapping

- Lots of problems came up trying to enlarge a image…

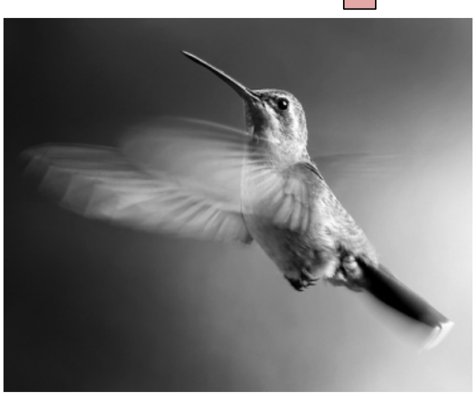

How do we fix this?

- Answer: do the opposite!

1. Create an output image
2. For each pixel in the output image, find the corresponding pixel in the input image*
3. Give that output pixel the same color

- *Requires that we invert the mapping

Inverse mapping

- How do we invert the mapping?

- With linear transformations T, we invert T

$$T = \begin{bmatrix} a & b \\ c & d \end{bmatrix} \quad T^{-1}T = \begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix}$$

$$T^{-1} = \frac{1}{ad - bc} \begin{bmatrix} d & -b \\ -c & a \end{bmatrix}$$

Inverse mapping

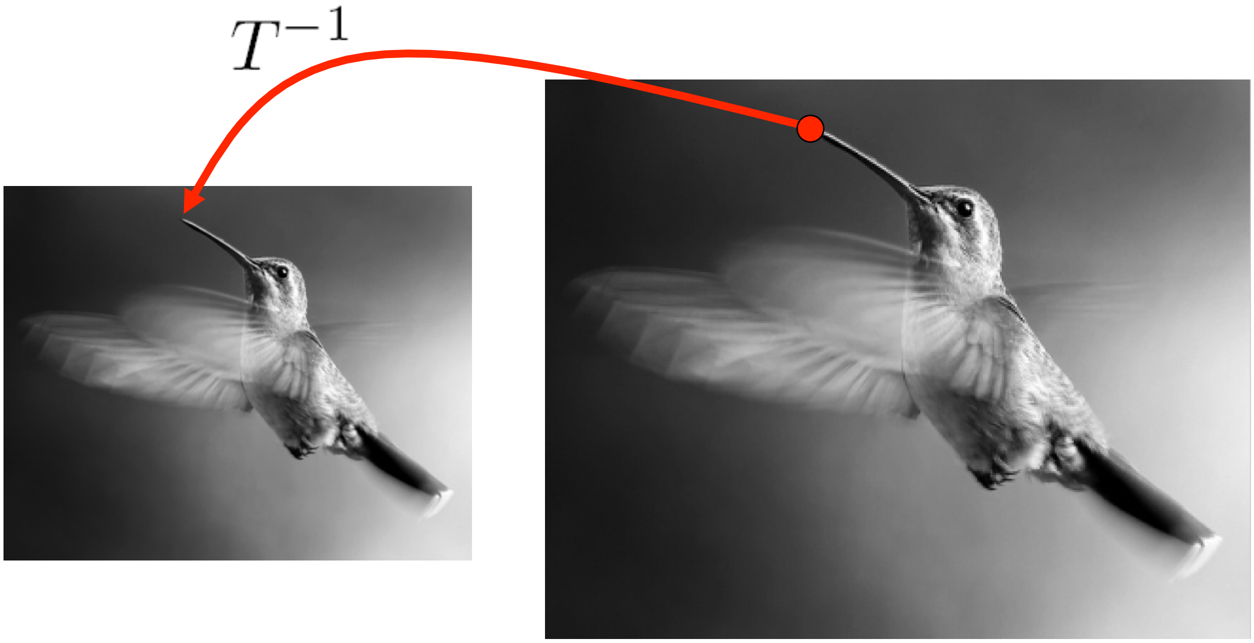

Resampling

- Suppose we scale image by 2
- T = [2 0 ; 0 2]
- inv(T) = ?
- Pixel (5,5) in img_out should be colored with pixel (2.5, 2.5) in img_in
- How do we find the intensity at (2.5, 2.5)?

Downsampling

- Suppose we scale image by 0.25 – a different problem with some surprising issues …

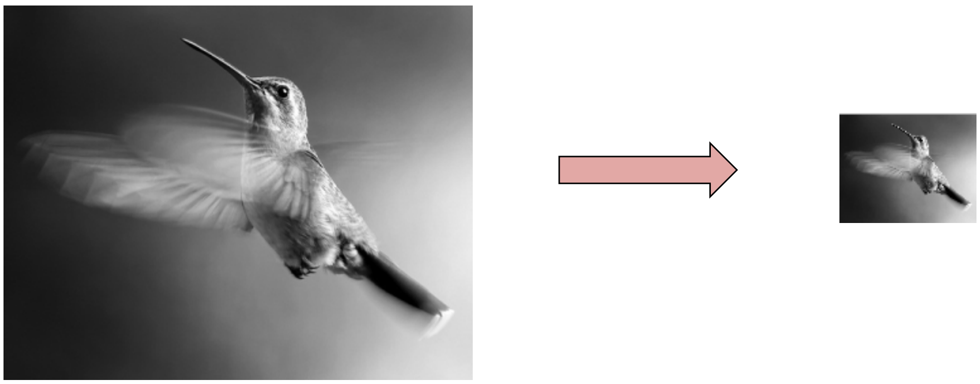

Downsampling

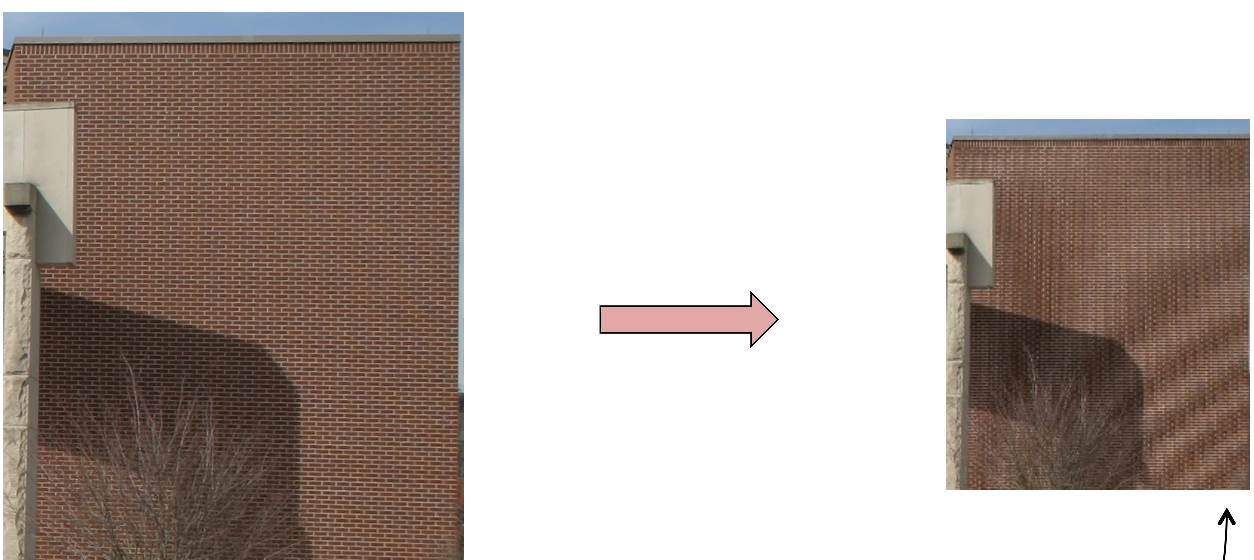

Which pixels were picked from the larger image to choose colour values?

What's going on?

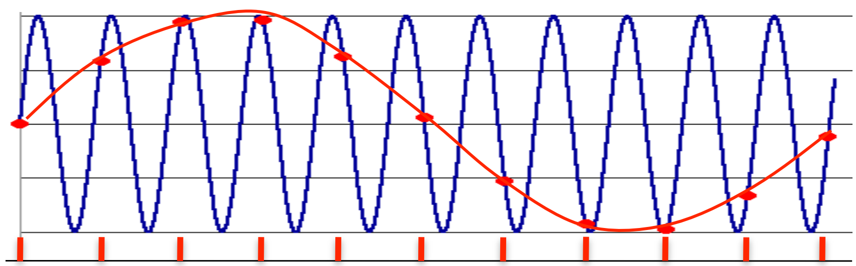

- **Aliasing** can arise when you sample a continuous signal or image
- Occurs when the sampling rate is not high enough to capture the detail in the image
- Can give you the wrong signal/image—an alias

2D example

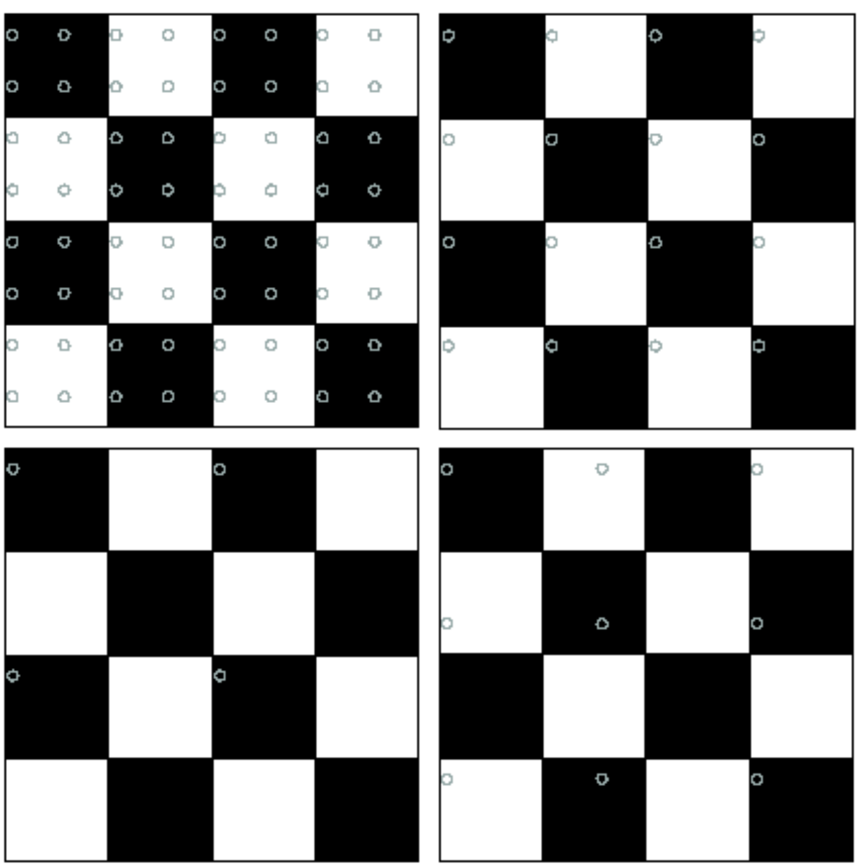

Good sampling

Bad sampling

Examples of aliasing

- Wagon wheel effect

- Moiré patterns

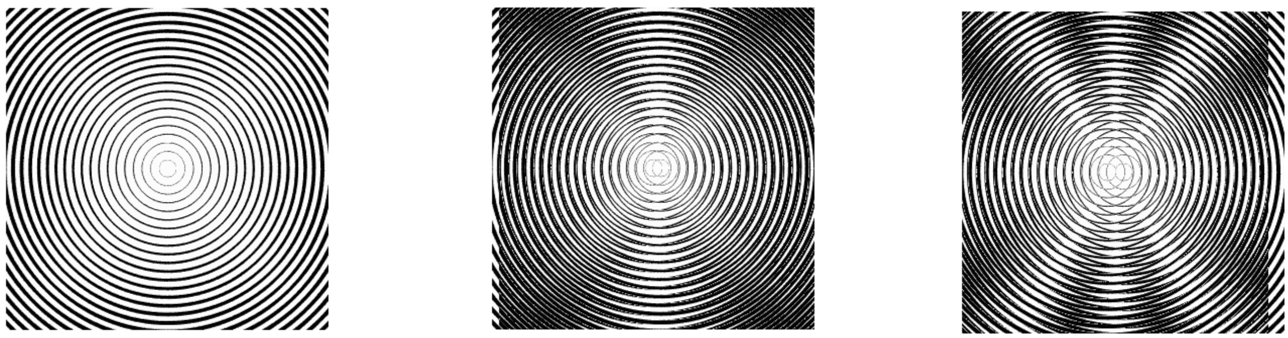

Image credit: Steve Seitz

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

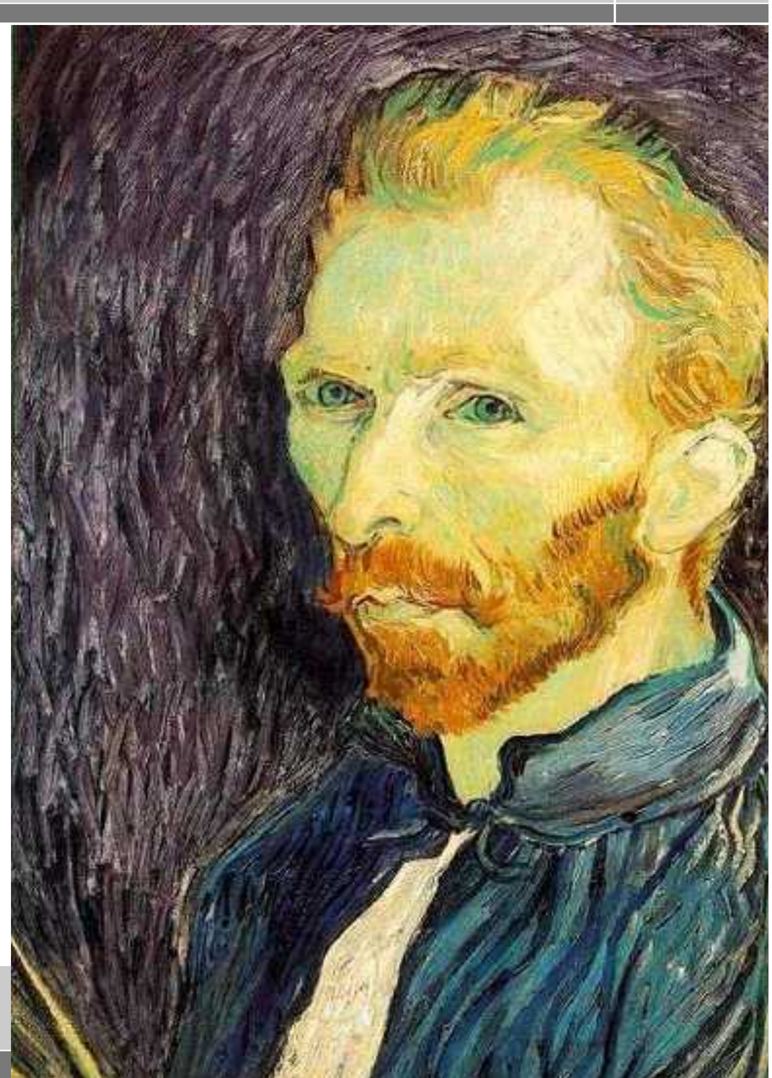

Slide credits: Steve Seitz

Image sub-sampling

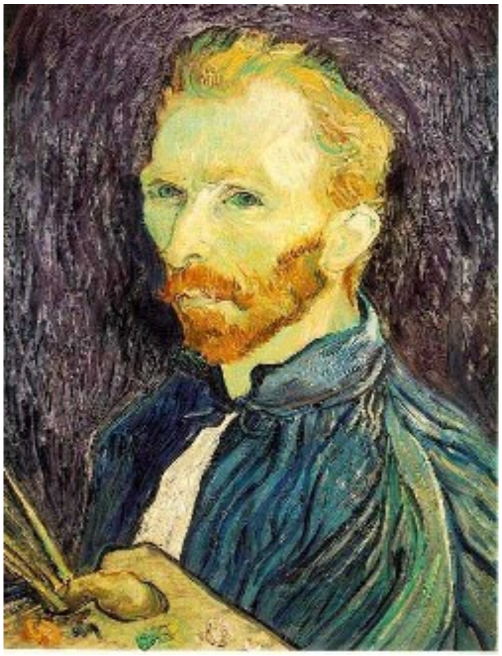

1/4

1/8

Current approach: throw away every other row and column (subsample)

Image sub-sampling

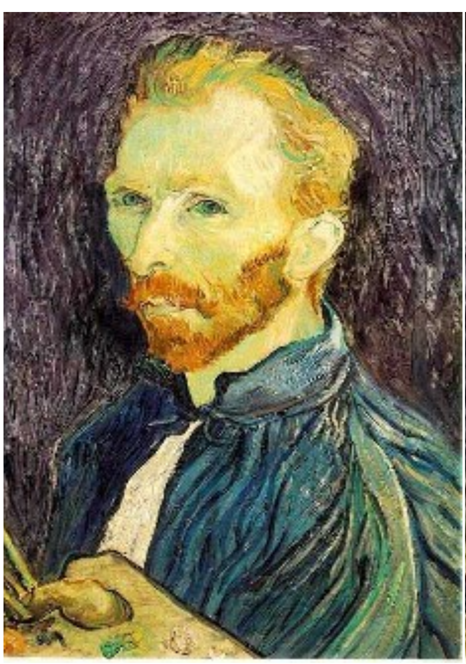

•*1/2*

•*1/4* *(2x zoom)*

•*1/8* *(4x zoom)*

Image sub-sampling

- What’s really going on?

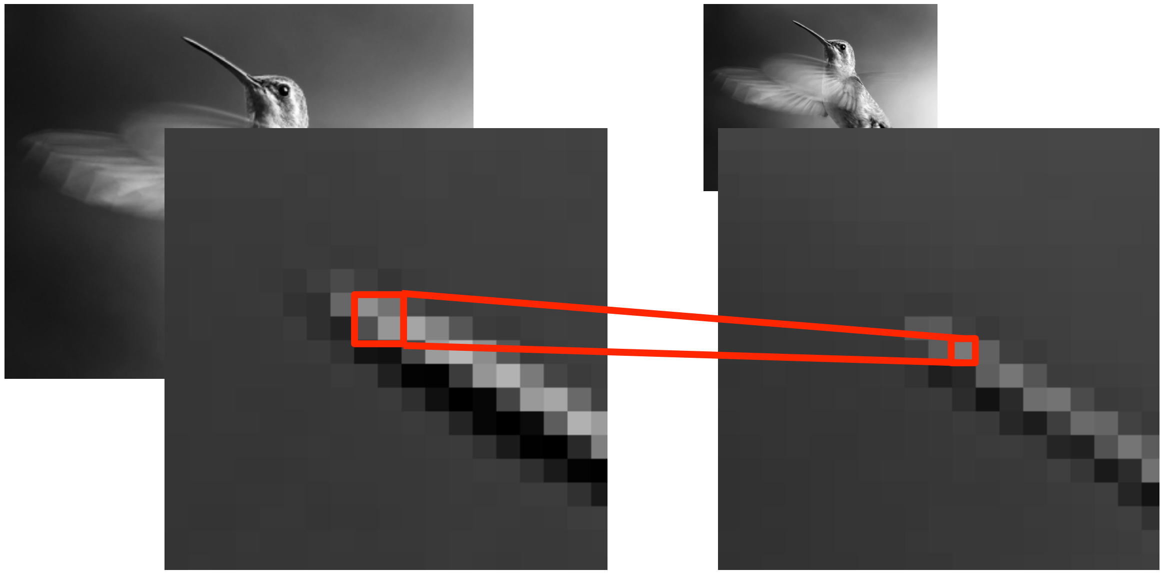

Which one to pick, or should we somehow merge all values?

Subsampling with pre-filtering

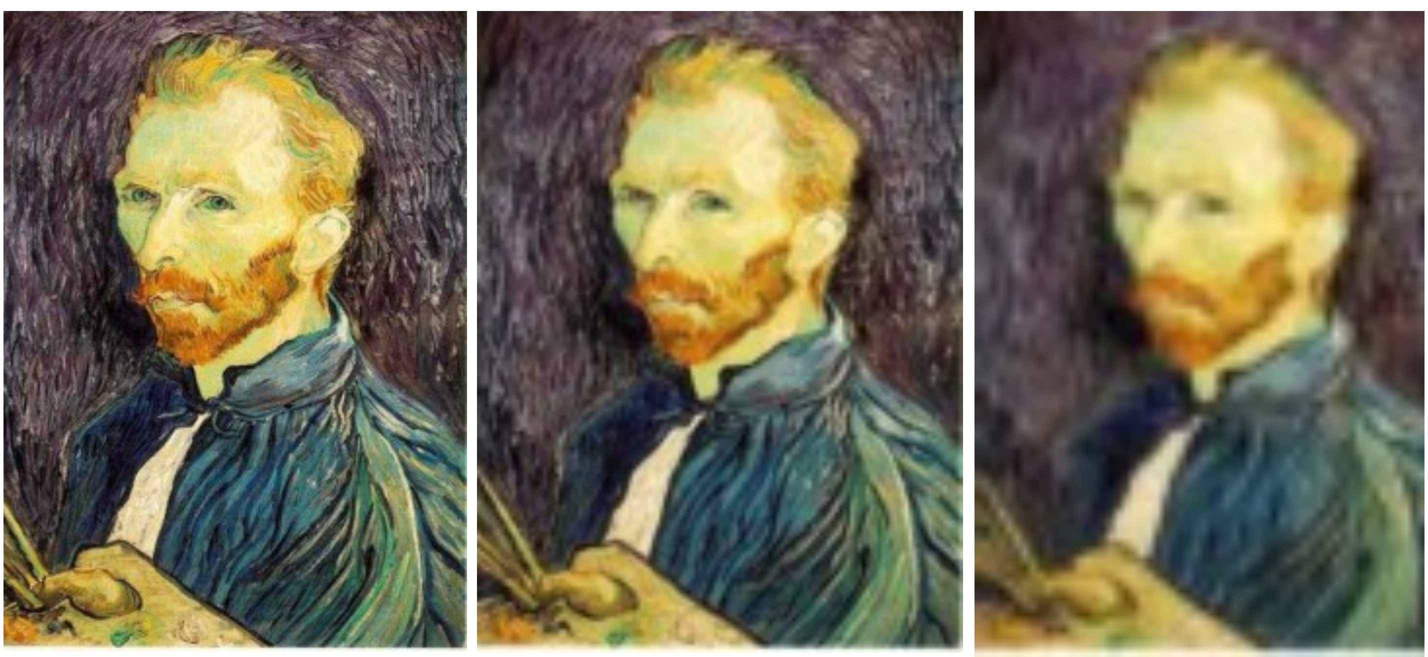

Average 2x2 *Average 4x4* *Average 8x8*

- Solution: blur the image, then subsample
  - Filter size should double for each ½ size reduction.

Subsampling with pre-filtering

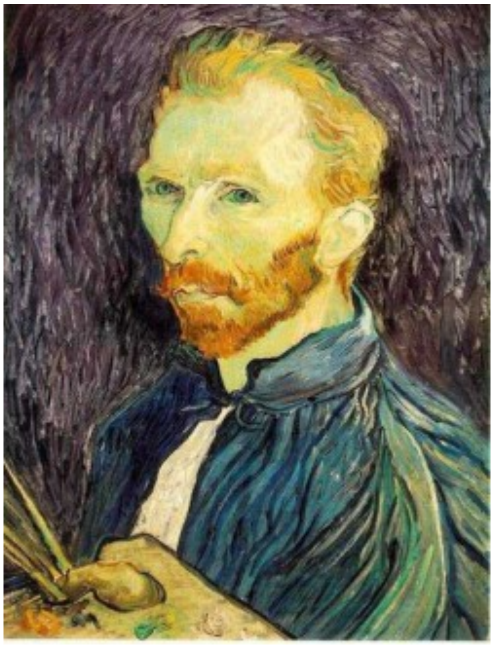

Average 2x2

Average 4x4

Average 8x8

- Solution: blur the image, then subsample
  - Filter size should double for each ½ size reduction.

Compare with

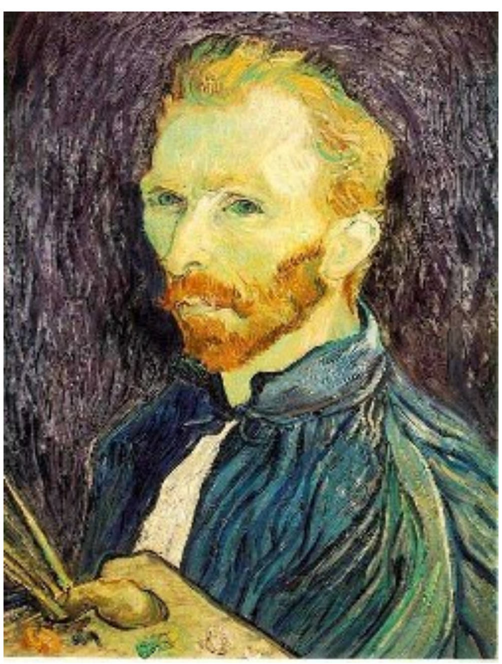

1/4

1/8

Weighted averaging

- Filtering, aka *convolution* (a *very* big topic!)
- Take one image, the *kernel* (usually small) to act as a (weighted) averager, then slide it over another image (usually big)
- At each point/pixel, multiply the kernel times the image, and add up the results
- This is the new value of the image … example:

0	1/5	0
1/5	1/5	1/5
0	1/5	0

Sometimes we want many resolutions

Idea: Represent NxN image as a “pyramid” of $1x1, 2x2, 4x4, \ldots, 2^k x 2^k$ images (assuming $N=2^k$)

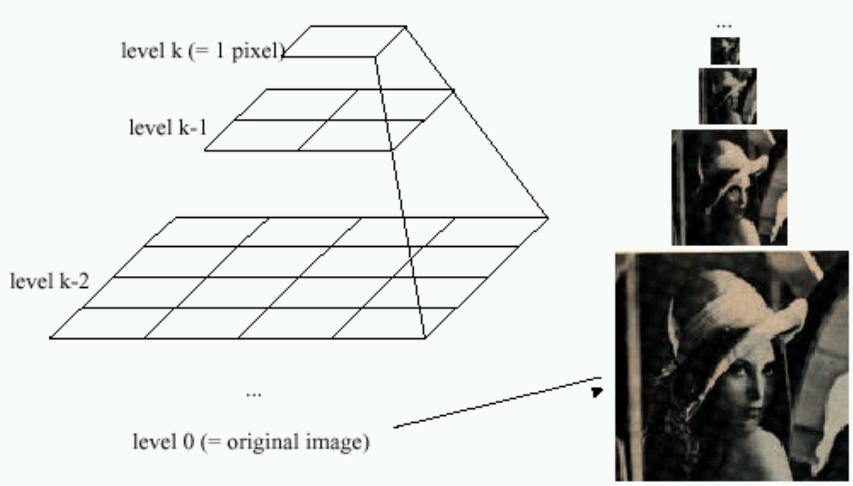

- Known as a **Gaussian Pyramid** [Burt and Adelson, 1983]
  - In computer graphics, a mip map [Williams, 1983]
  - A precursor to wavelet transform

More image transformations

- Default rotation is around the point (0, 0) – the upper-left corner of the image

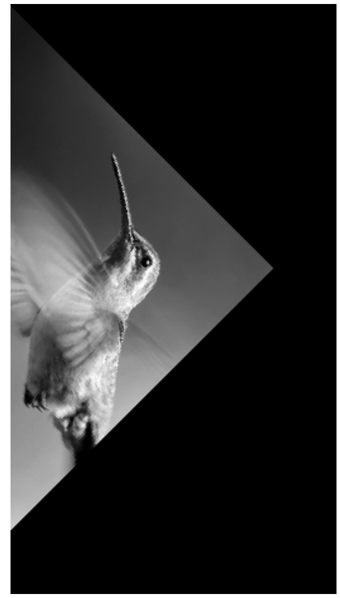

- This probably isn't really what we want…

Translation

- We really want to rotate around the *center* of the image
- Nice trick: move the center of the image to the origin, apply the default rotation, then move the center back
- (Moving an image is called "translation")
- But technically, translation isn't a linear function - check the two definition properties for $\mathbf{f}(\mathbf{v}) = \mathbf{v} + \mathbf{w}$ (for some constant vector $\mathbf{w}$)…
  - $\mathbf{f}(\mathbf{v} + \mathbf{u}) \;=?\; \mathbf{f}(\mathbf{v}) + \mathbf{f}(\mathbf{u})$ $\quad LHS = \mathbf{v} + \mathbf{u} + \mathbf{w} \quad RHS = \mathbf{v} + \mathbf{u} + 2\mathbf{w}$
  - $\mathbf{f}(a\mathbf{v}) \;=?\; a\,\mathbf{f}(\mathbf{v})$ $\quad LHS = a\mathbf{v} + \mathbf{w} \quad RHS = a\mathbf{v} + 2a\mathbf{w}$
- Formally we call this an *affine* linear function