Can we do something cool with edges already?

S. Avidan and A. Shamir

*Seam Carving for Content-Aware Image Resizing*

SIGGRAPH 2007

Paper: `http://www.win.tue.nl/~wstahw/edu/2IV05/seamcarving.pdf`

# Simple Application: Seam Carving

- Content-aware resizing

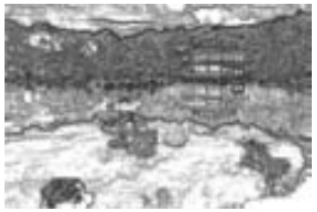

- Find path from top to bottom row with minimum gradient energy
- Remove (or replicate) those pixels

# Simple Application: Seam Carving

- Content-aware resizing

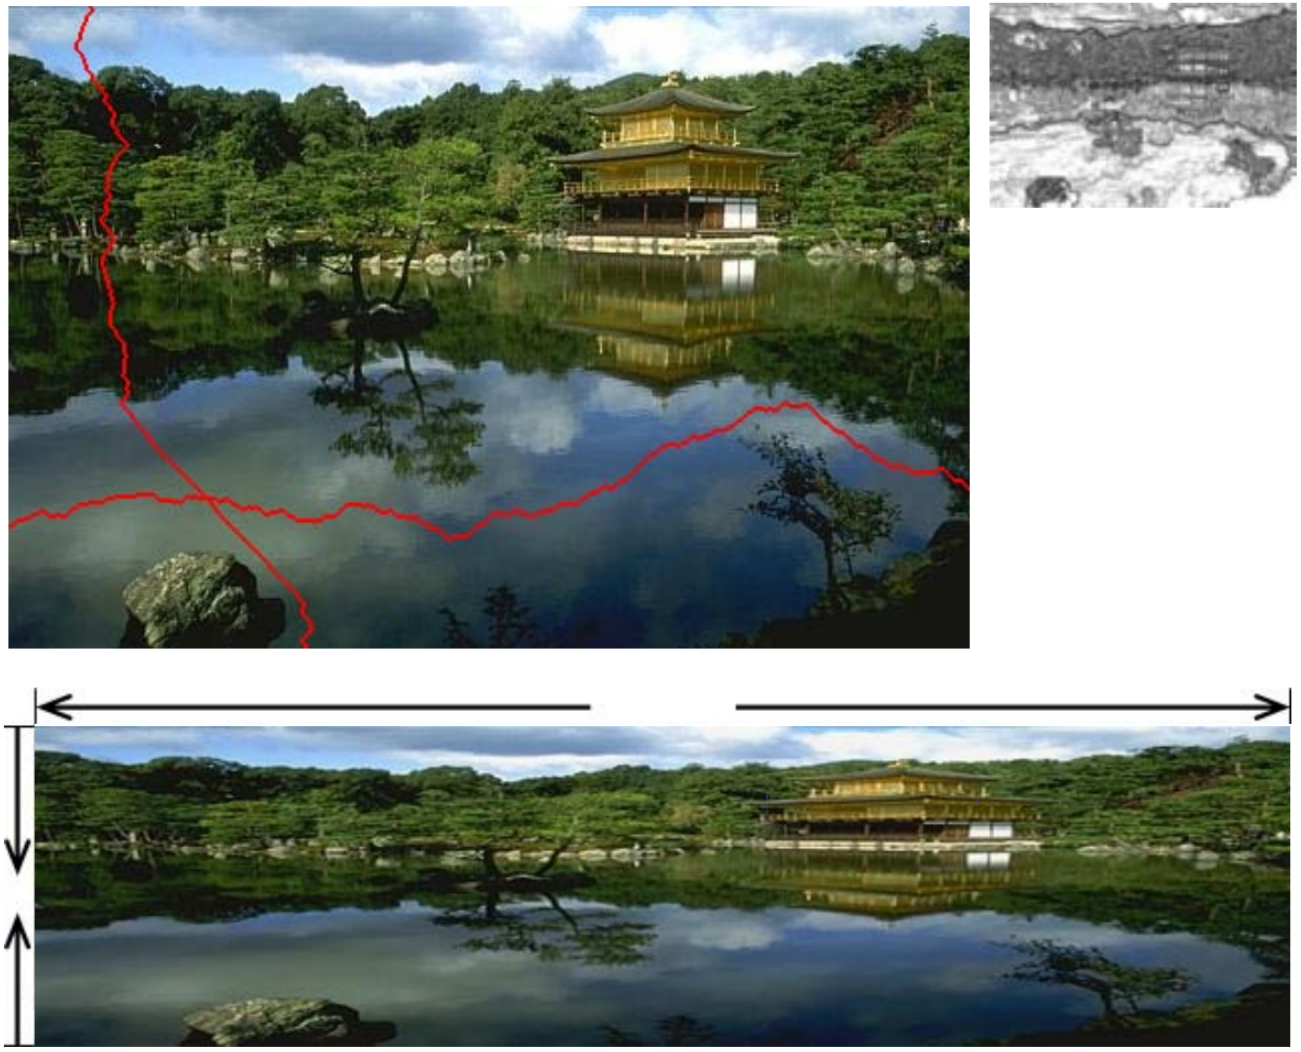

# Simple Application: Seam Carving

- Content-aware resizing

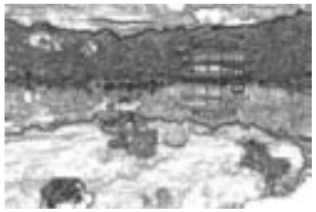

# Seam Carving

- A vertical seam **s** is a list of column indices, one for each row, where each subsequent column differs by no more than one slot.
- Let $G$ denote the image gradient magnitude. Optimal 8-connected path:

$$\mathbf{s}^* = \text{argmin}_{\mathbf{s}} E(\mathbf{s}) = \text{argmin}_{\mathbf{s}} \sum_{i=1}^{n} G(s_i)$$

- Can be computed via dynamic programming
- Compute the cumulative minimum energy for all possible connected seams at each entry $(i, j)$:

$$M(i,j) = G(i,j) + \min\left(M(i-1,j-1), M(i-1,j), M(i-1,j+1)\right)$$

- Backtrack from min value in last row of M to pull out optimal seam path.

# Seam Carving – Examples



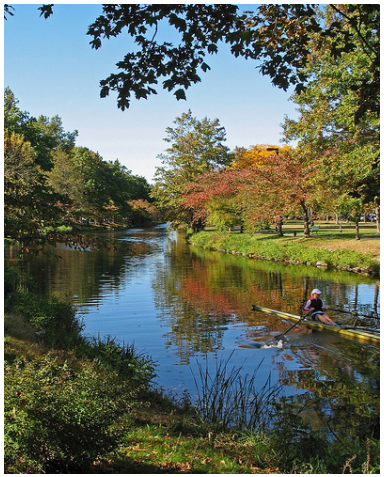


- Implement seam carving for 5% extra credit on first assignment

# Image Pyramids

# Finding Waldo

- Let's revisit the problem of finding Waldo
- This time he is on the road

image

template (filter)

# Finding Waldo

- He comes closer but our filter doesn't know that
- How can we find Waldo?

image

template (filter)

# Idea: Re-size Image

- Re-scale the image multiple times! Do correlation on every size!

template (filter)

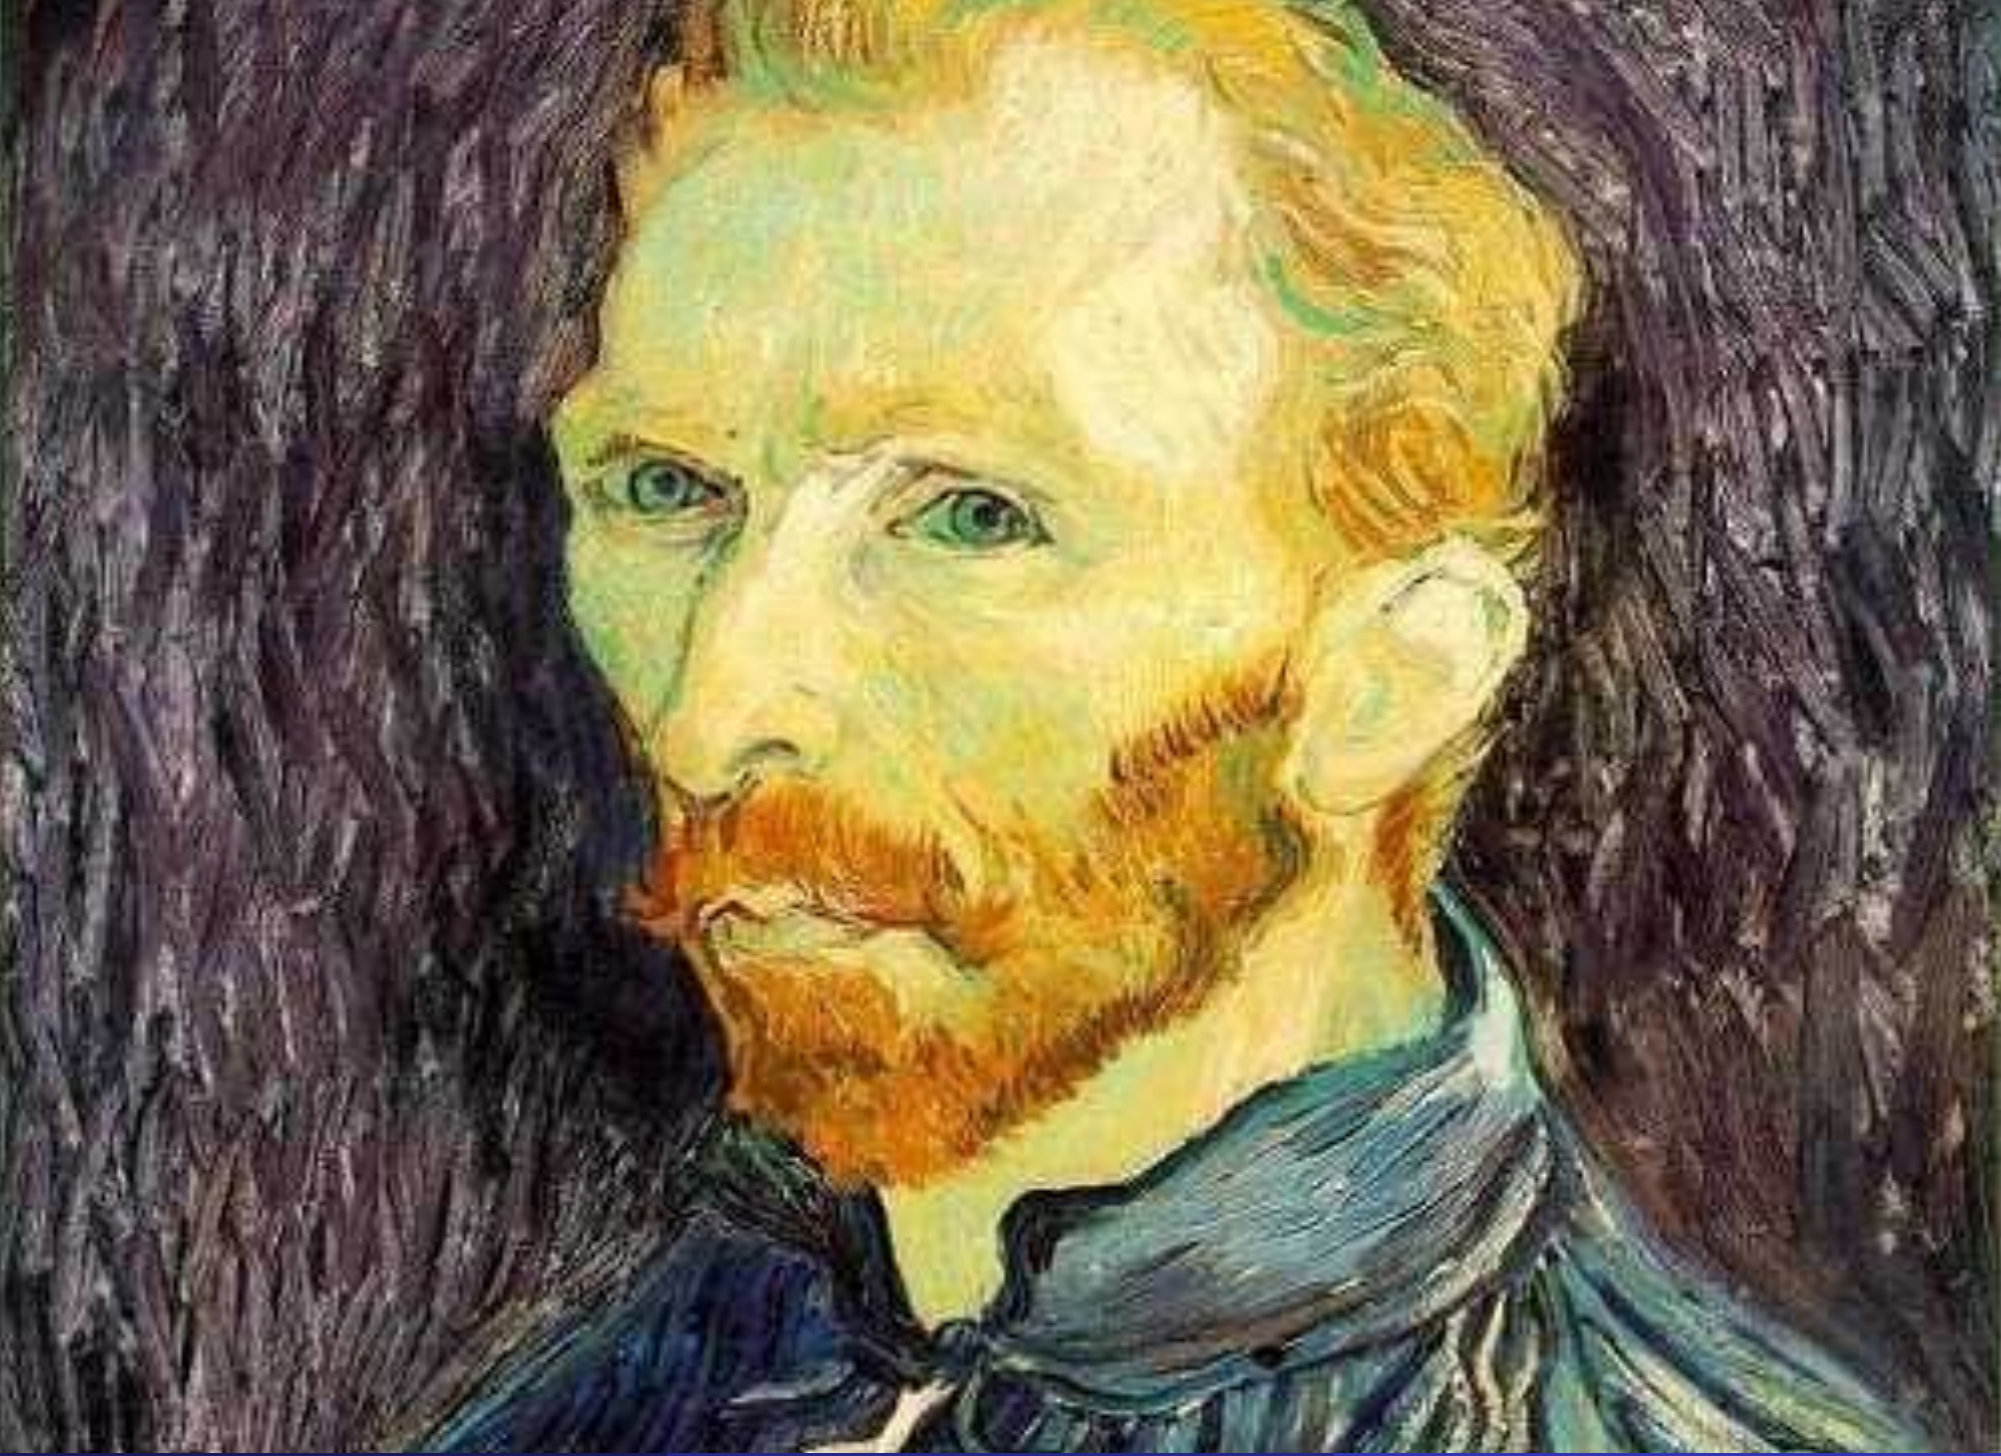

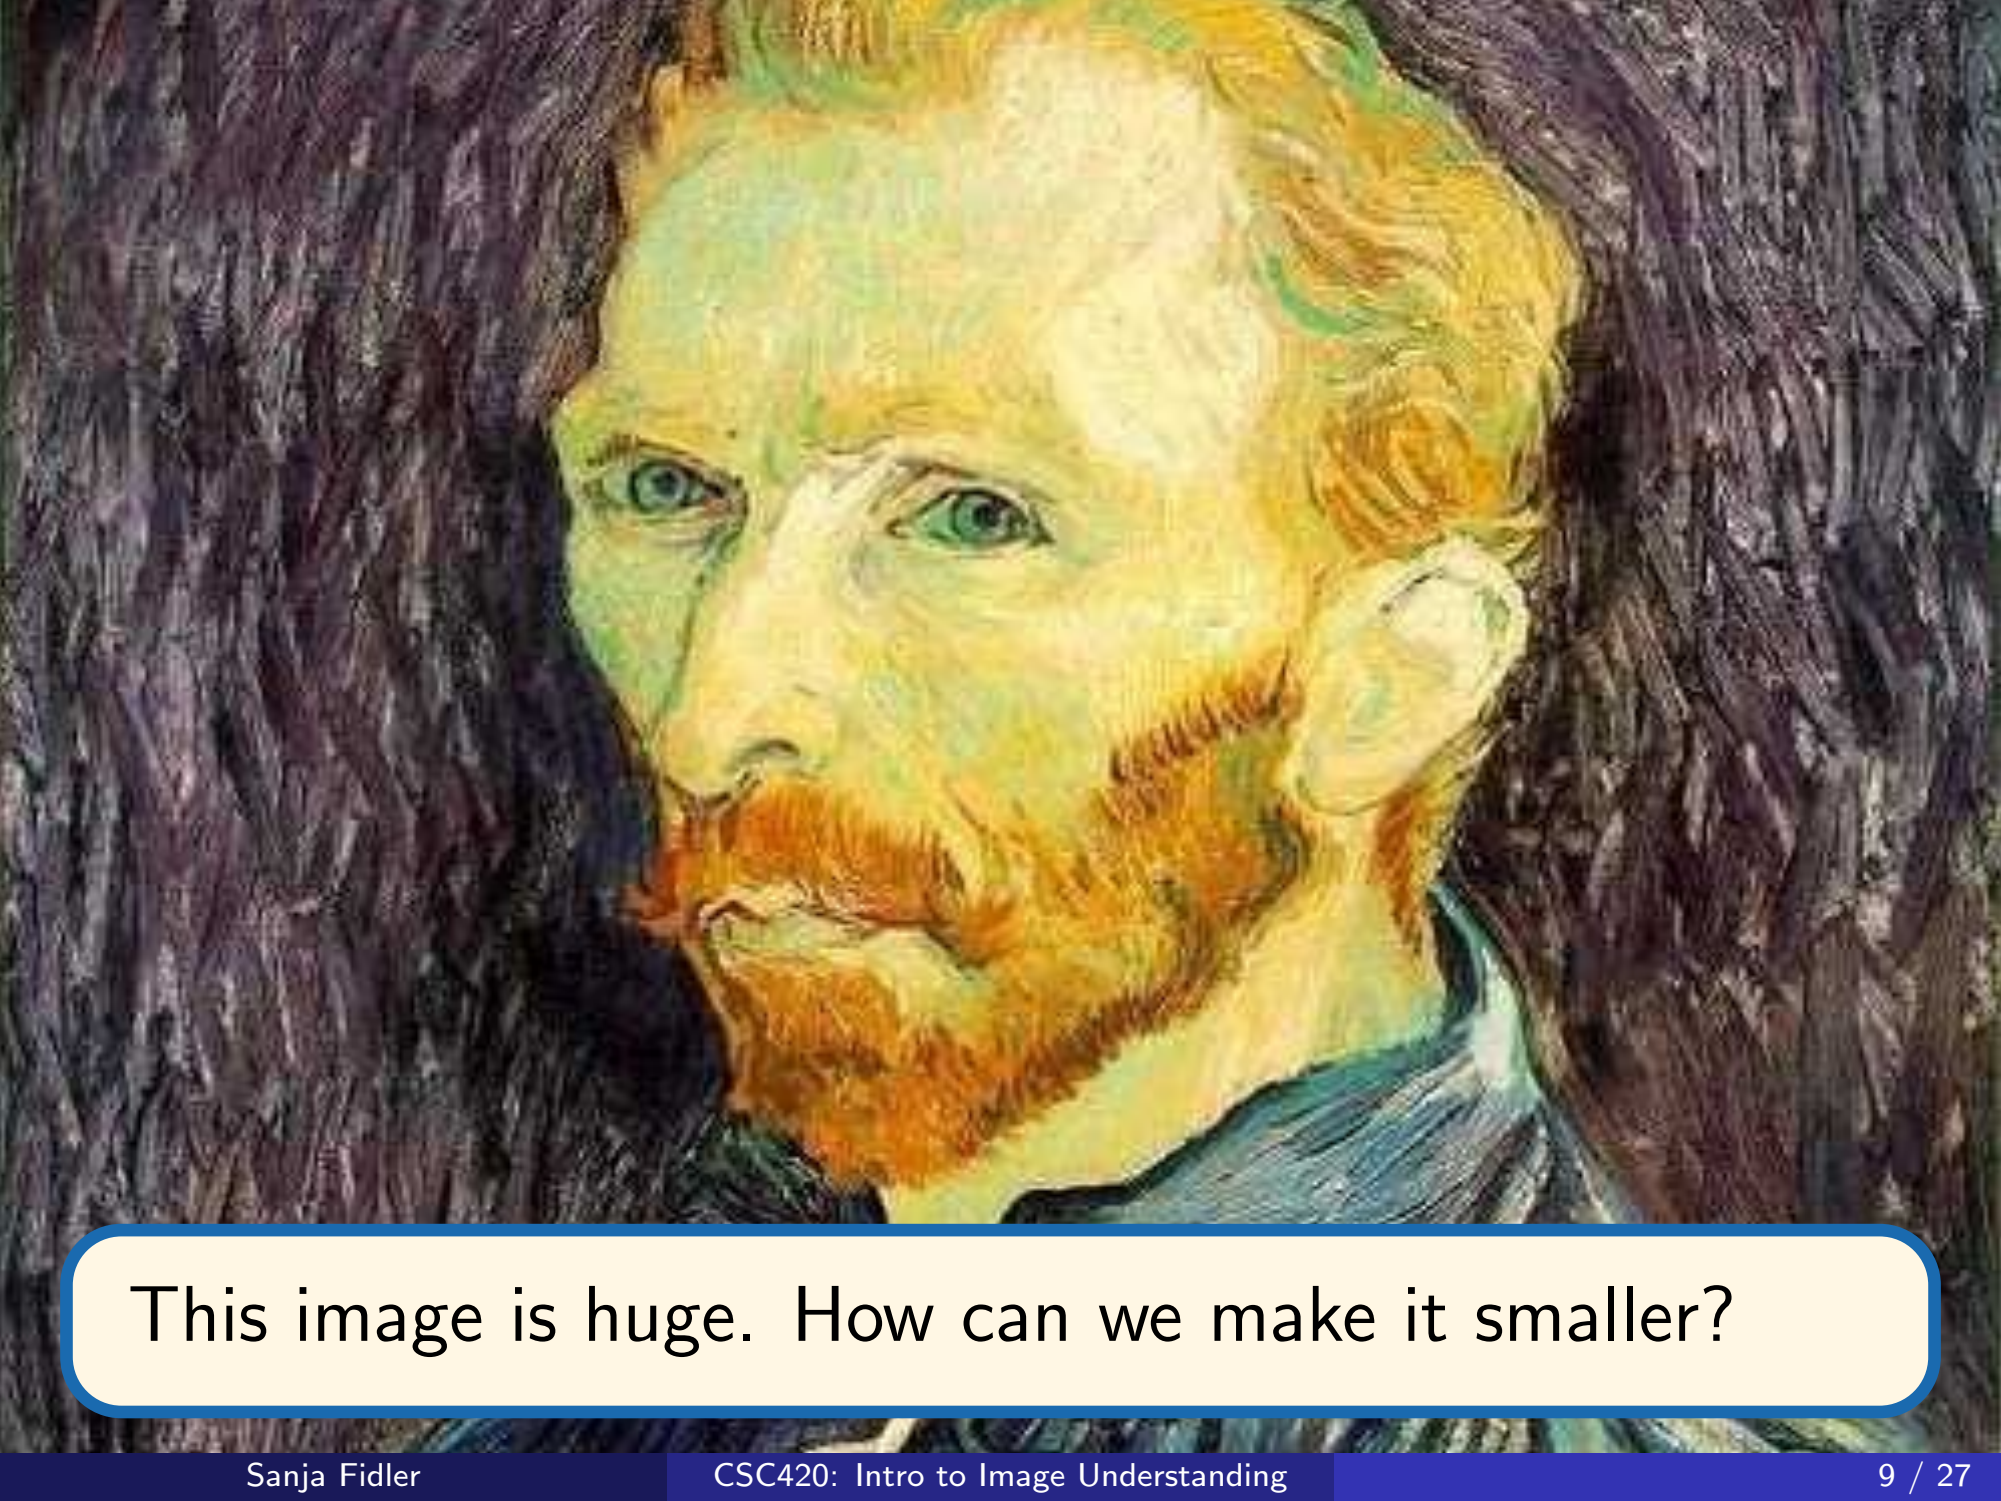

This image is huge. How can we make it smaller?

# Image Sub-Sampling

- **Idea**: Throw away every other row and column to create a 1/2 size image

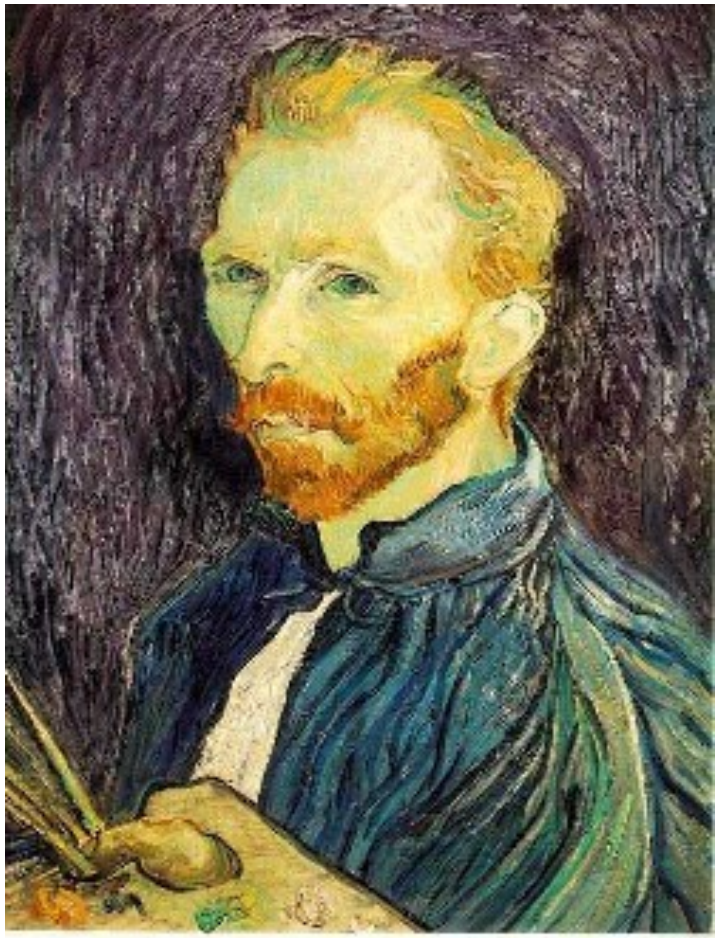

1/4

1/8

[Source: S. Seitz]

# Image Sub-Sampling

- Why does this look so crufty?

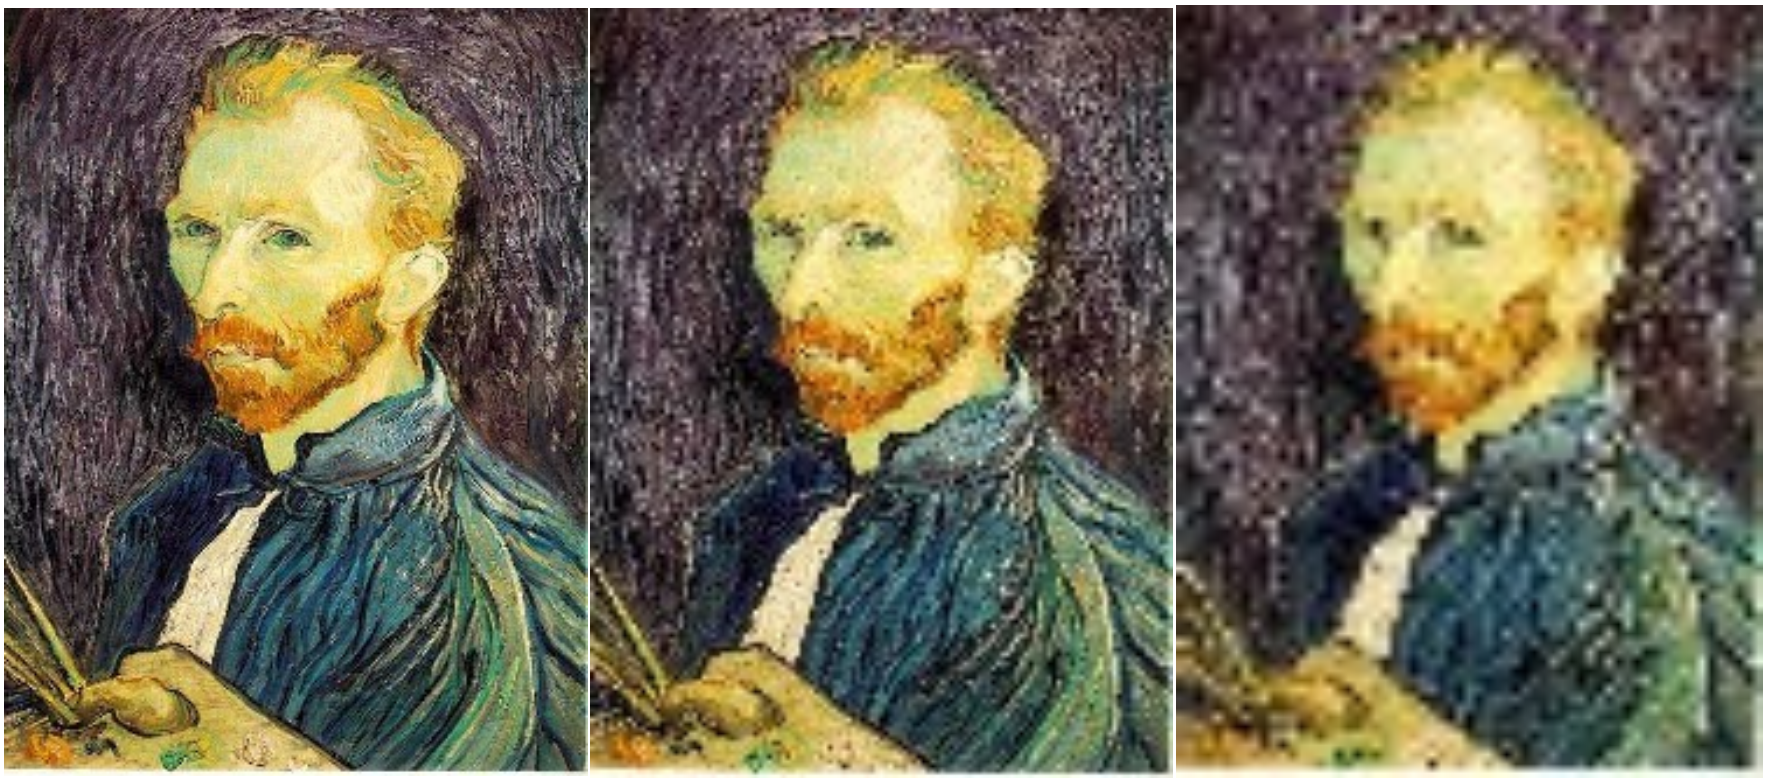

1/2 1/4 (2x zoom) 1/8 (4x zoom)

[Source: S. Seitz]

# Image Sub-Sampling

[Source: F. Durand]

# Even worse for synthetic images

- I want to resize my image by factor 2
- And I take every other column and every other row (1st, 3rd, 5th, etc)

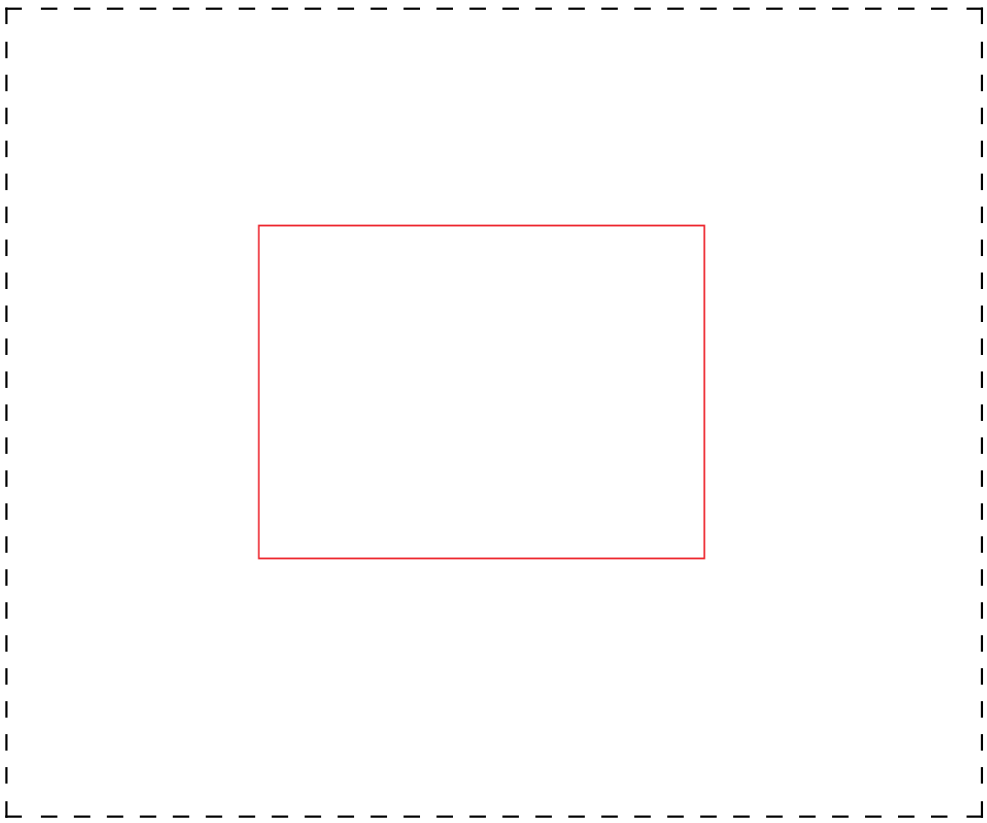

Figure: Dashed line denotes the border of the image (it's not part of the image)

# Even worse for synthetic images

- I want to resize my image by factor 2
- And I take every other column and every other row (1st, 3rd, 5th, etc)
- Where is the rectangle!

Figure: Dashed line denotes the border of the image (it's not part of the image)

# Even worse for synthetic images

- What's in the image?
- Now I want to resize my image by half in the width direction
- And I take every other column (1st, 3rd, 5th, etc)

# Even worse for synthetic images

- What's in the image?
- Now I want to resize my image by half in the width direction
- And I take every other column (1st, 3rd, 5th, etc)

# Even worse for synthetic images

- What's in the image?
- Now I want to resize my image by half in the width direction
- And I take every other column (1st, 3rd, 5th, etc)
- Where is the chicken!

# Even worse for synthetic images

- What's happening?

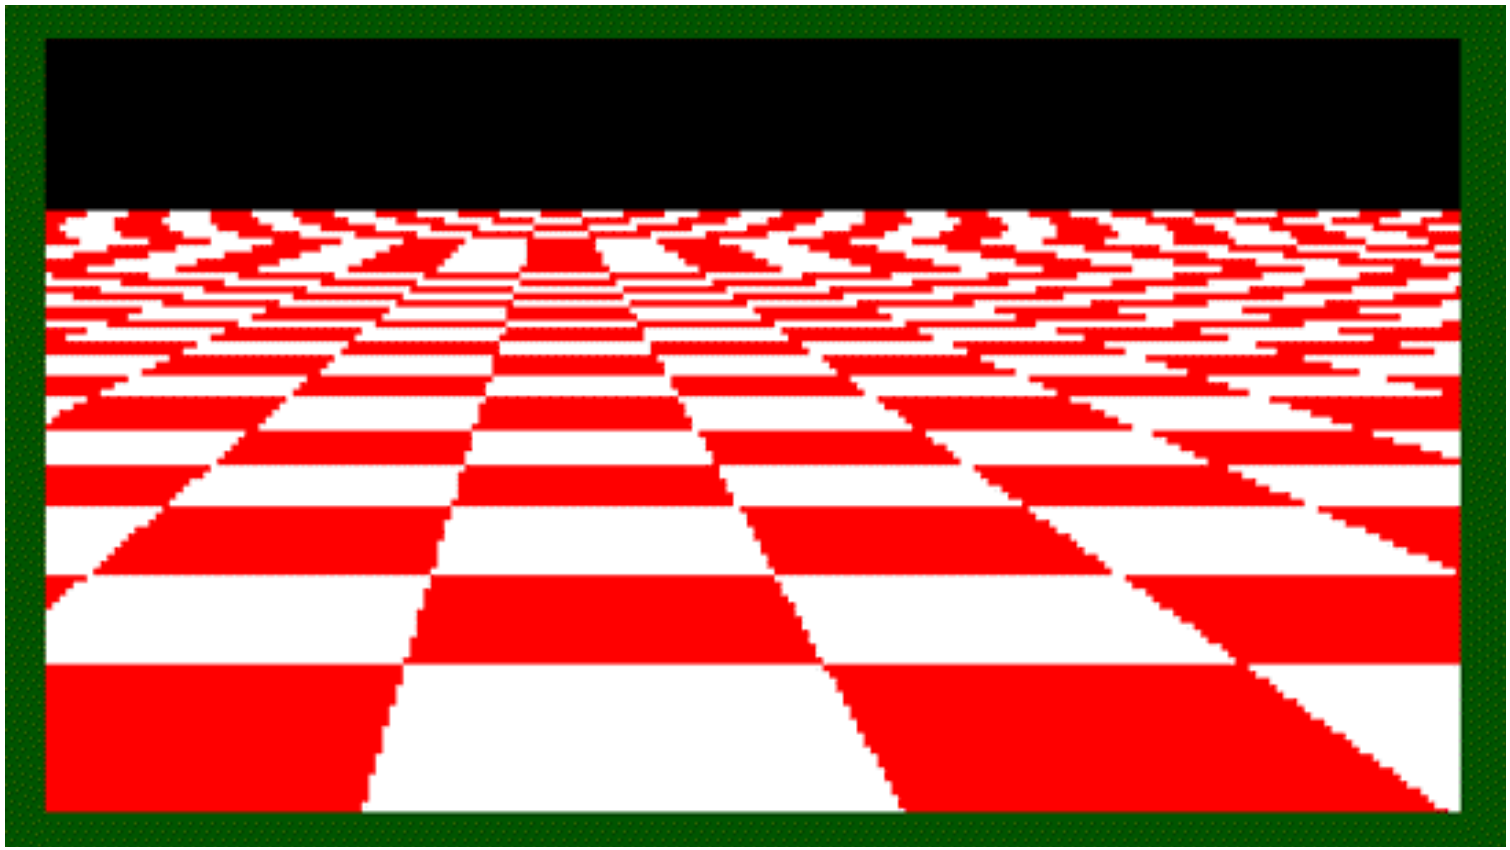

[Source: L. Zhang]

# Aliasing

- Occurs when your sampling rate is not high enough to capture the amount of detail in your image

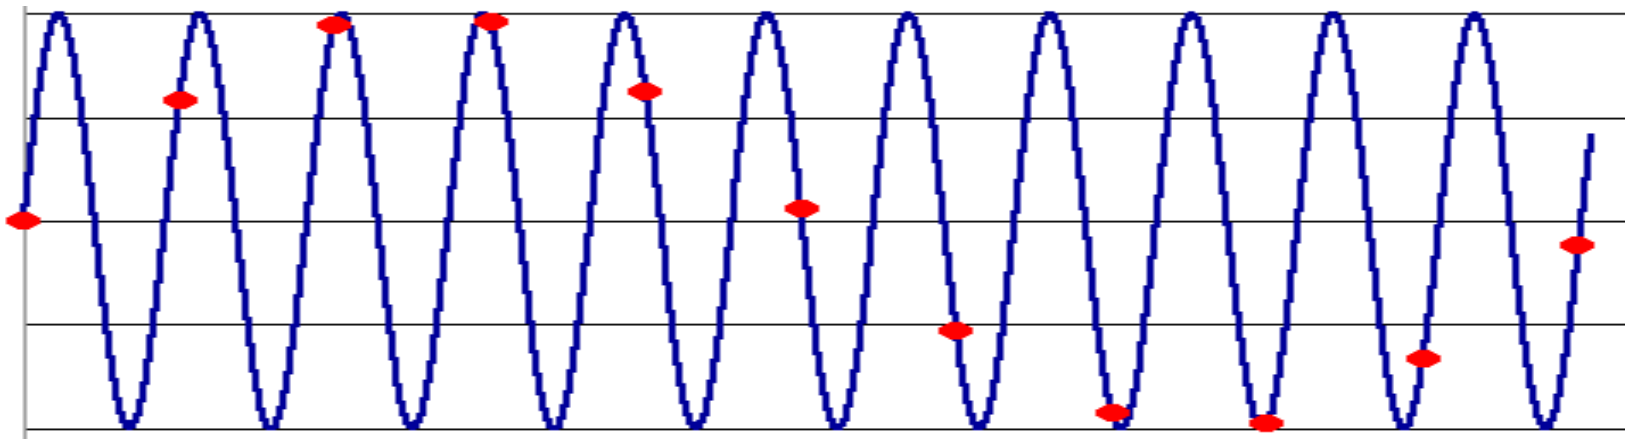

- To do sampling right, need to understand the structure of your signal/image

[Source: R. Urtasun]

# Aliasing

- Occurs when your sampling rate is not high enough to capture the amount of detail in your image

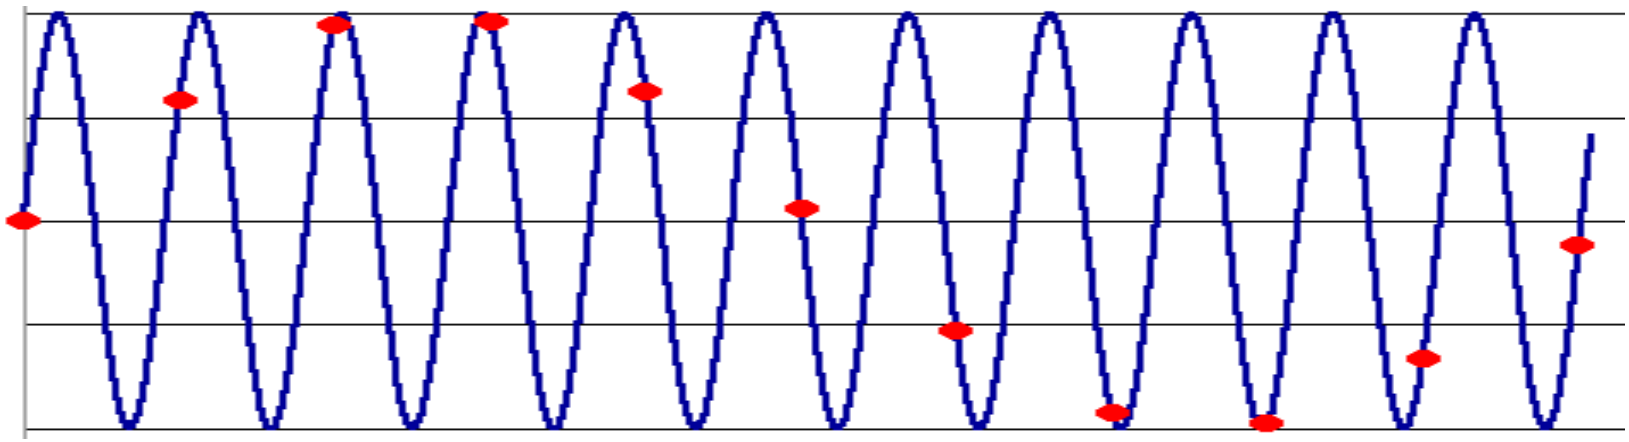

- To do sampling right, need to understand the structure of your signal/image
- The minimum sampling rate is called the **Nyquist rate**

[Source: R. Urtasun]

# Aliasing

- Occurs when your sampling rate is not high enough to capture the amount of detail in your image

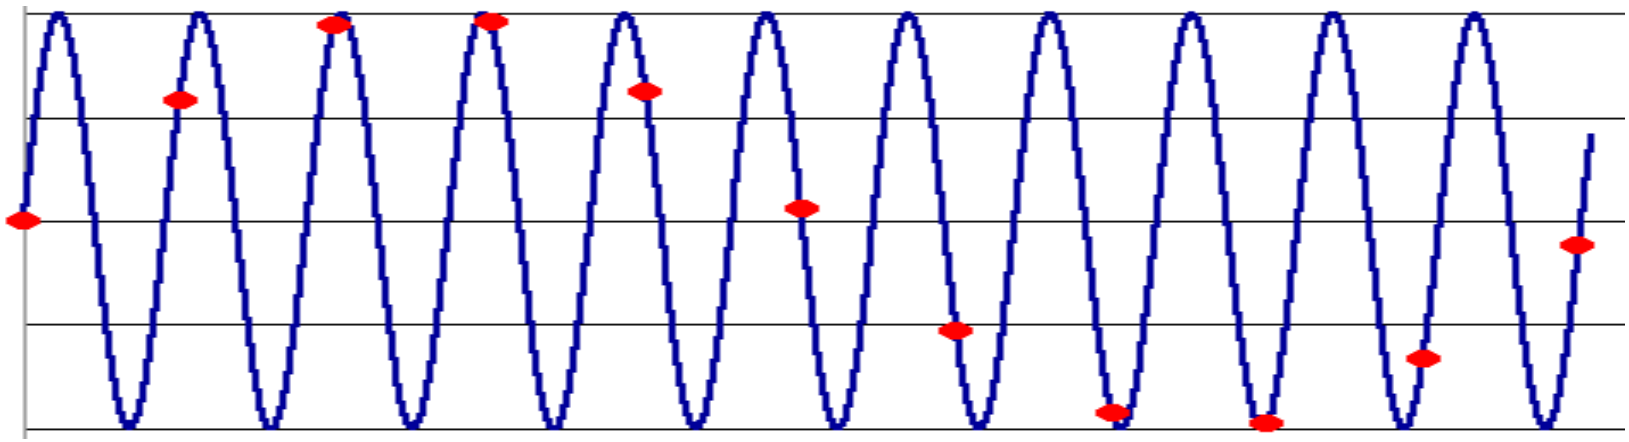

- To do sampling right, need to understand the structure of your signal/image
- The minimum sampling rate is called the **Nyquist rate**

[Source: R. Urtasun]

# Mr. Nyquist

- Harry Nyquist says that one should look at the frequencies of the signal.
- One should find the highest frequency (via Fourier Transform)
- To sample properly you need to sample with at least twice that frequency
- For those interested: `http://en.wikipedia.org/wiki/Nyquist%E2%80%93Shannon_sampling_theorem`

- He looks like a smart guy, we'll just believe him

# 2D example

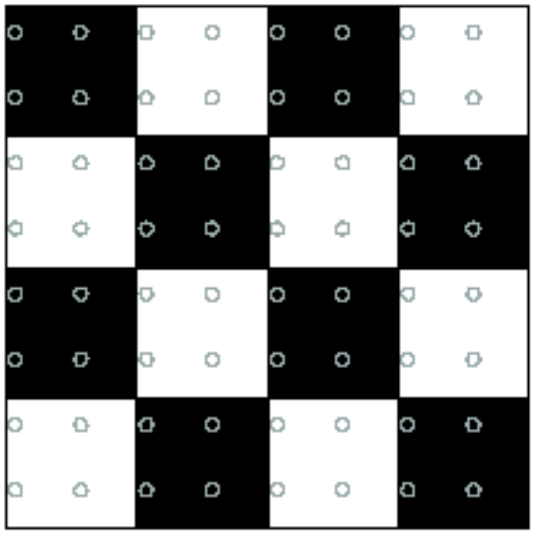

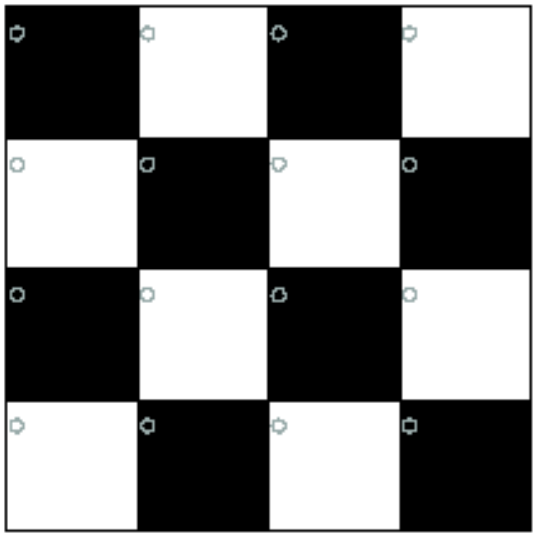

Good sampling

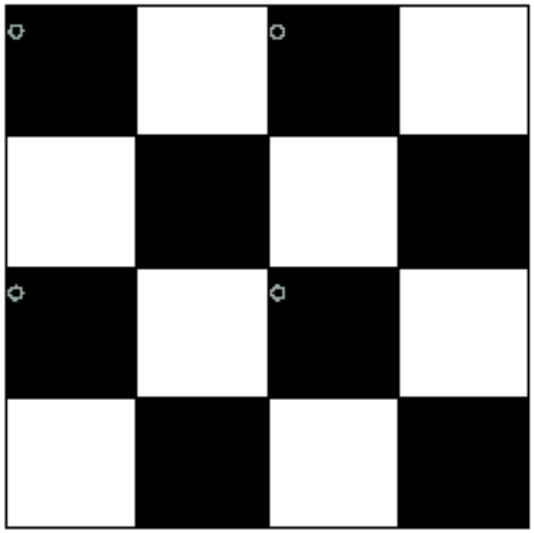

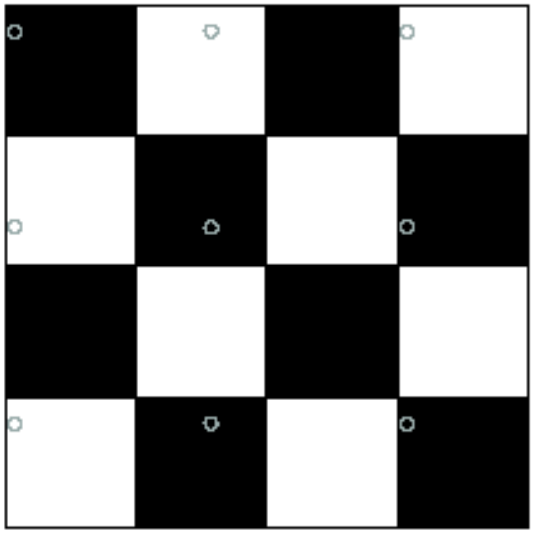

Bad sampling

[Source: N. Snavely]

# Going back to Downsampling ...

- When downsampling by a factor of two, the original image has frequencies that are too high
- High frequencies are caused by sharp edges
- How can we fix this?

[Adopted from: R. Urtasun]

# Going back to Downsampling ...

- When downsampling by a factor of two, the original image has frequencies that are too high
- High frequencies are caused by sharp edges
- How can we fix this?

[Adopted from: R. Urtasun]

# Gaussian pre-filtering

- Solution: Blur the image via Gaussian, then subsample. Very simple!

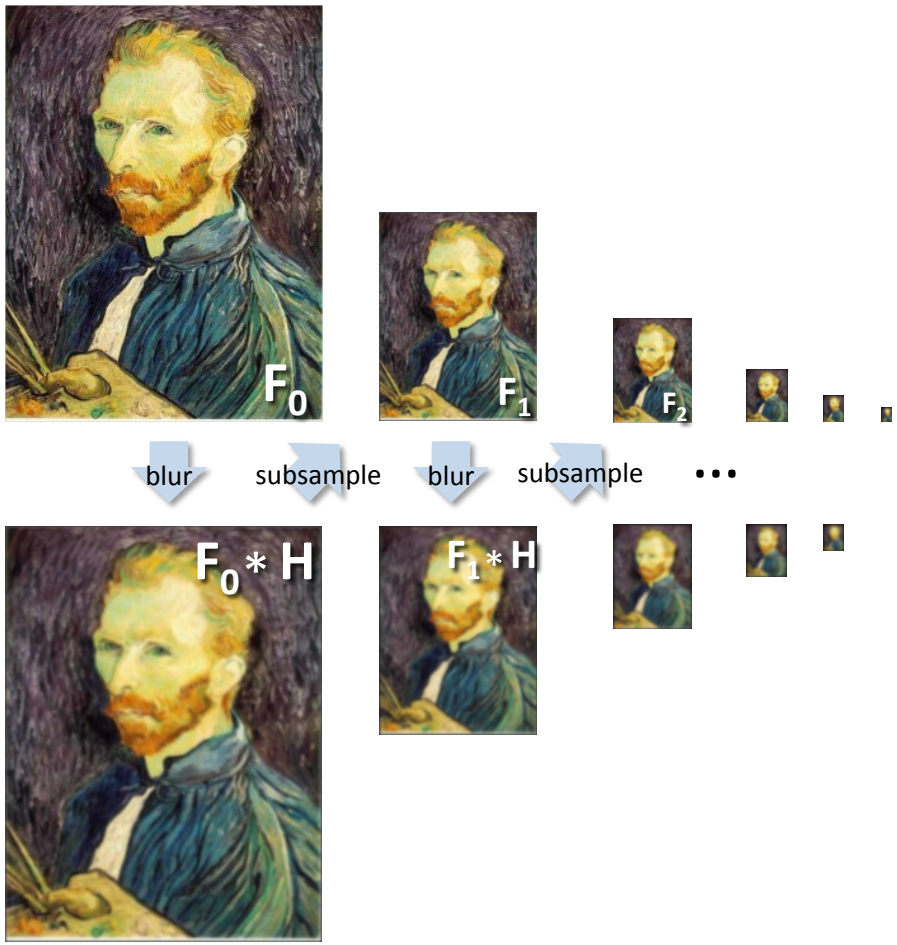


[Source: N. Snavely]

# Subsampling with Gaussian pre-filtering

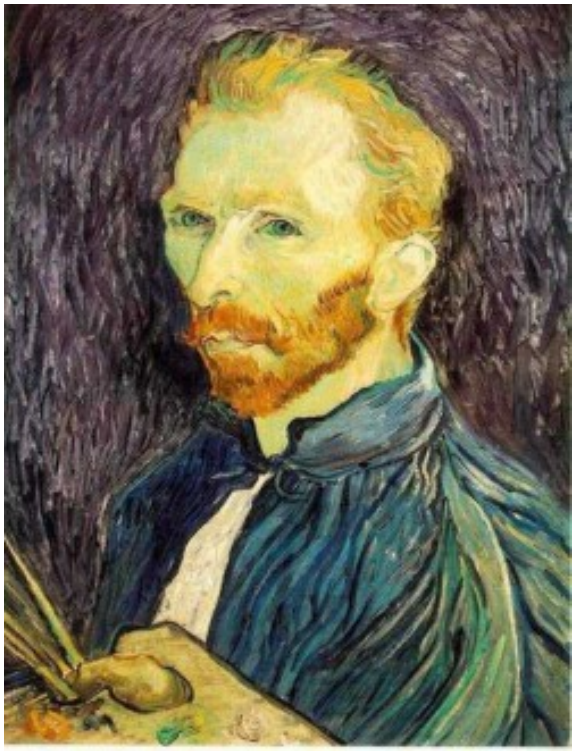

Gaussian 1/2

G 1/4

G 1/8

[Source: S. Seitz]

# Compare with ...

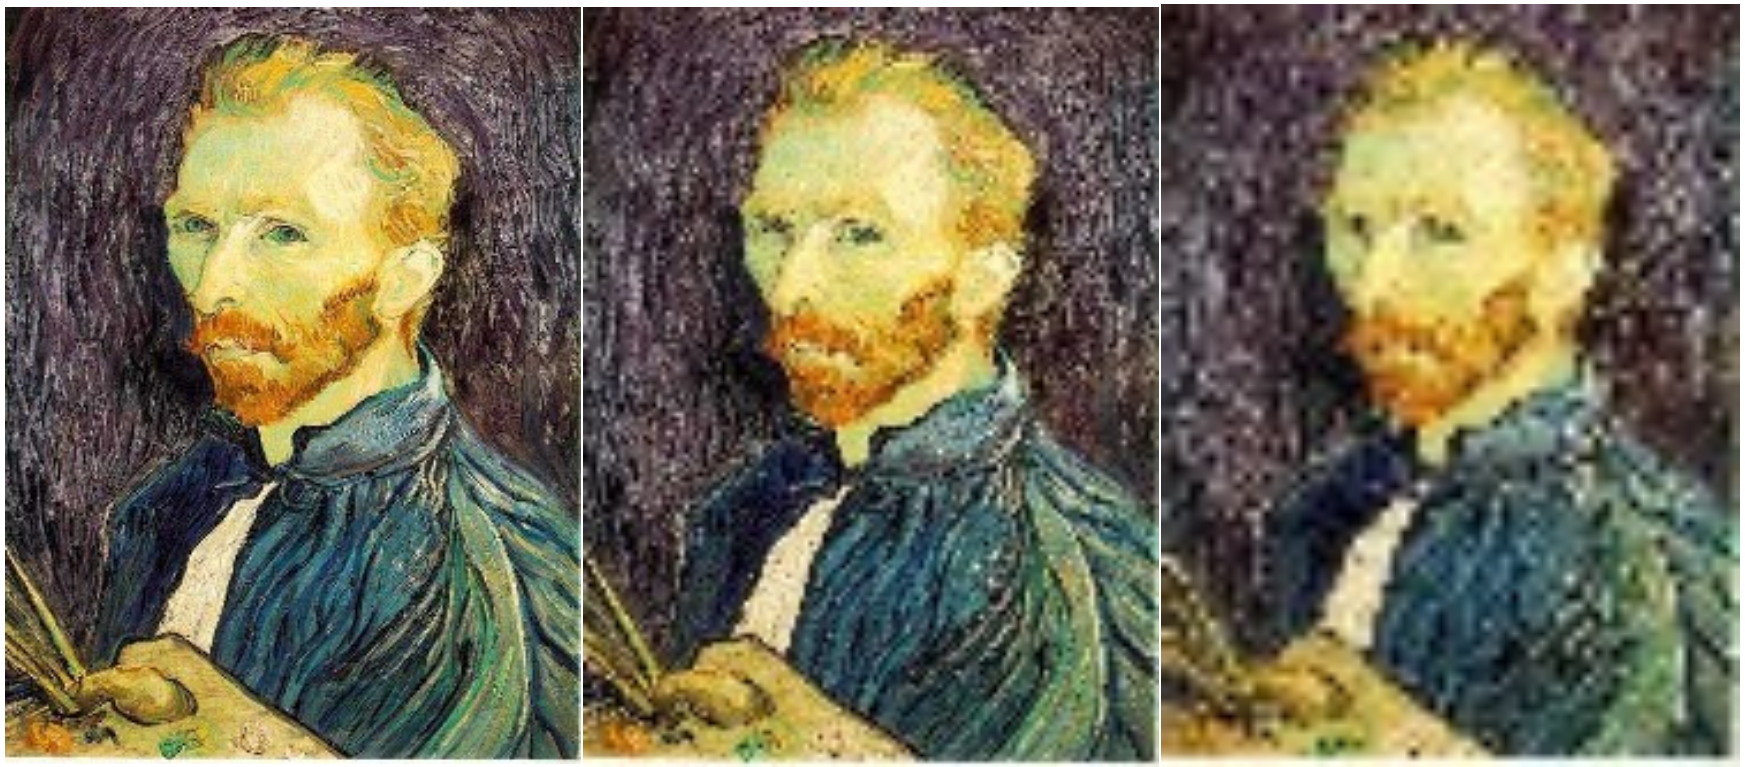

1/2 1/4 (2x zoom) 1/8 (4x zoom)

[Source: S. Seitz]

# Where is the Rectangle?

- My image

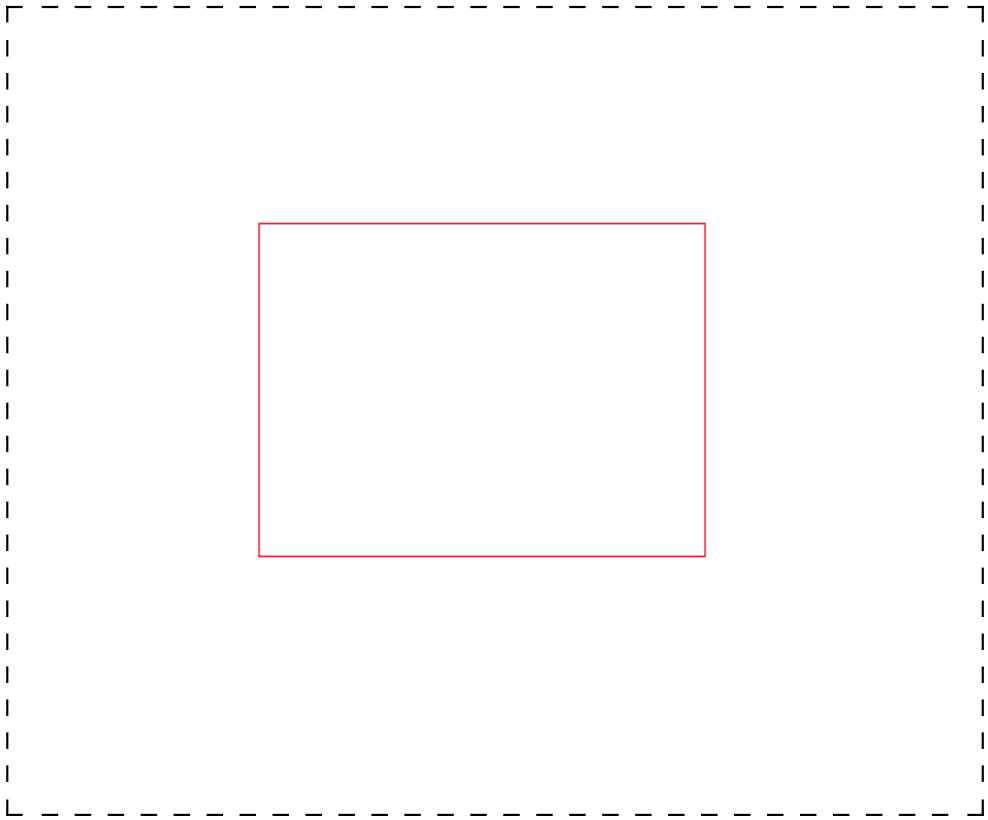

Figure: Dashed line denotes the border of the image (it's not part of the image)

# Where is the Rectangle?

- My image
- Let's blur

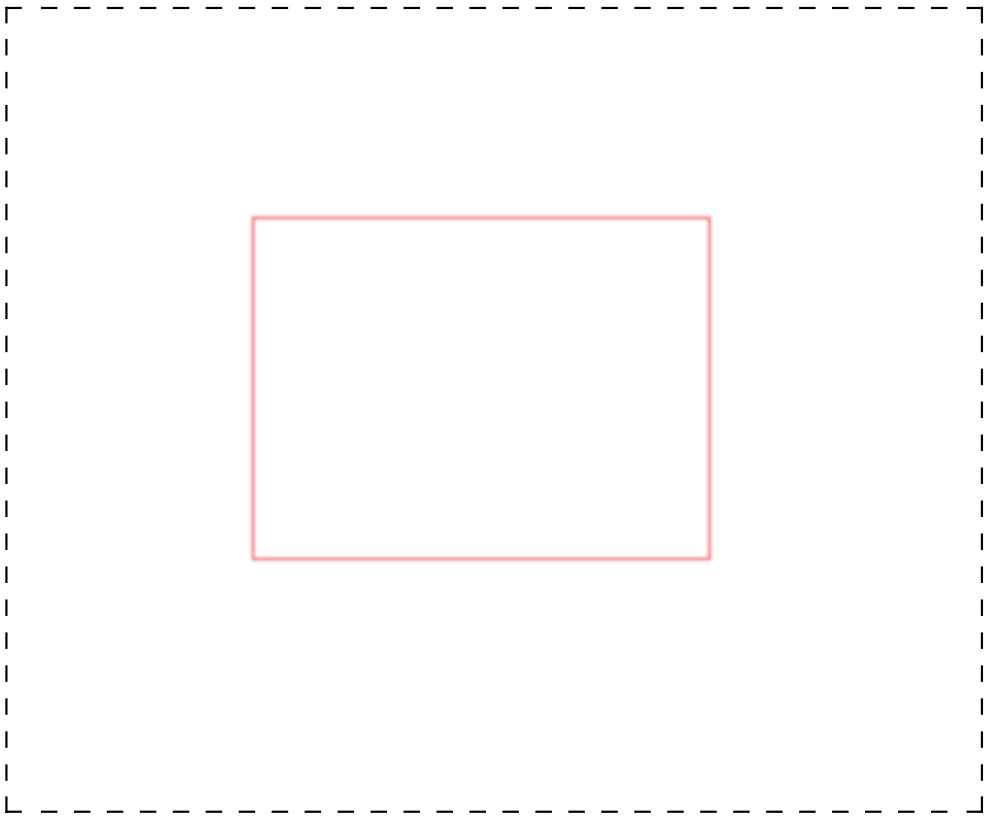

Figure: Dashed line denotes the border of the image (it's not part of the image)

# Where is the Rectangle?

- My image
- Let's blur
- And now take every other row and column

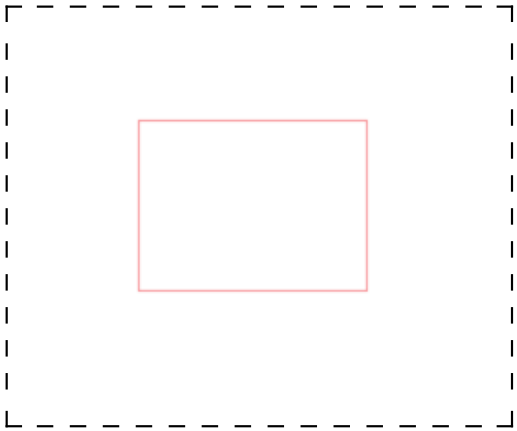

Figure: Dashed line denotes the border of the image (it's not part of the image)

# Where is the Chicken?

- My image

# Where is the Chicken?

- My image
- Let's blur

# Where is the Chicken?

- My image
- Let's blur
- And now take every other column

# Gaussian Pyramids [Burt and Adelson, 1983]

- A sequence of images created with Gaussian blurring and downsampling is called a **Gaussian Pyramid**
- In computer graphics, a *mip map* [Williams, 1983]

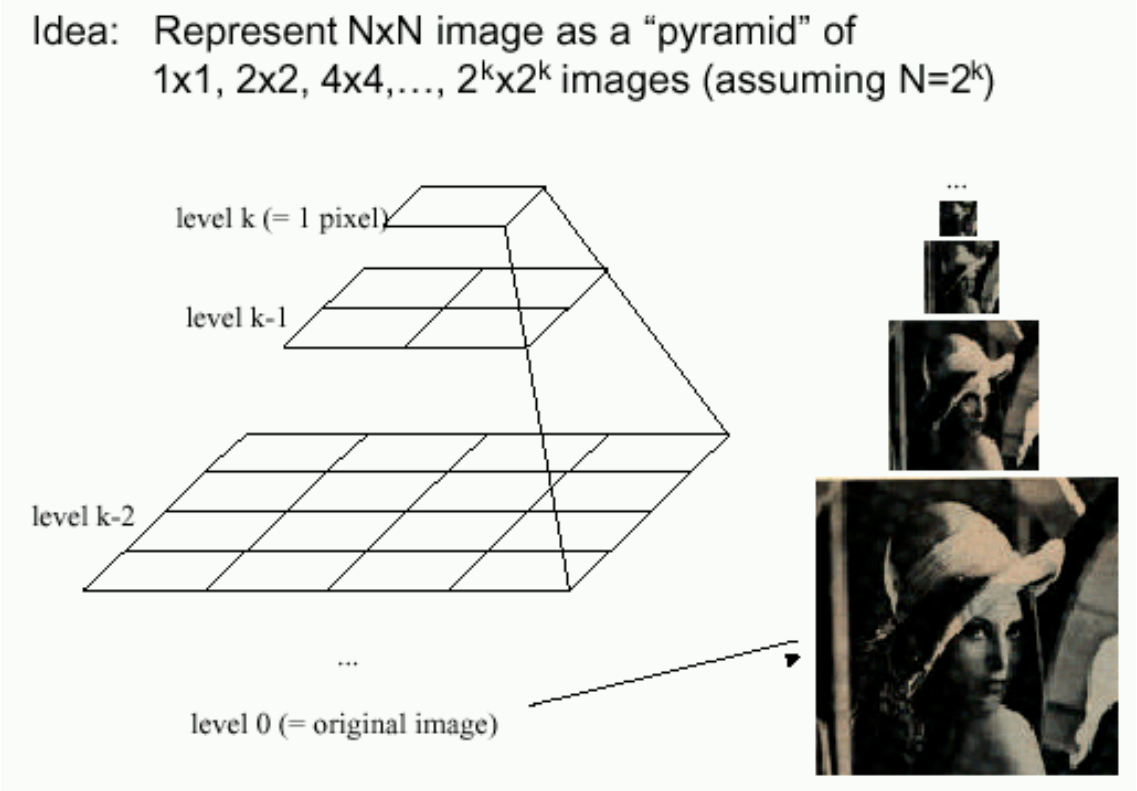


- How much space does a Gaussian pyramid take compared to original image?

[Source: S. Seitz]

# Gaussian Pyramids [Burt and Adelson, 1983]

- A sequence of images created with Gaussian blurring and downsampling is called a **Gaussian Pyramid**
- In computer graphics, a *mip map* [Williams, 1983]

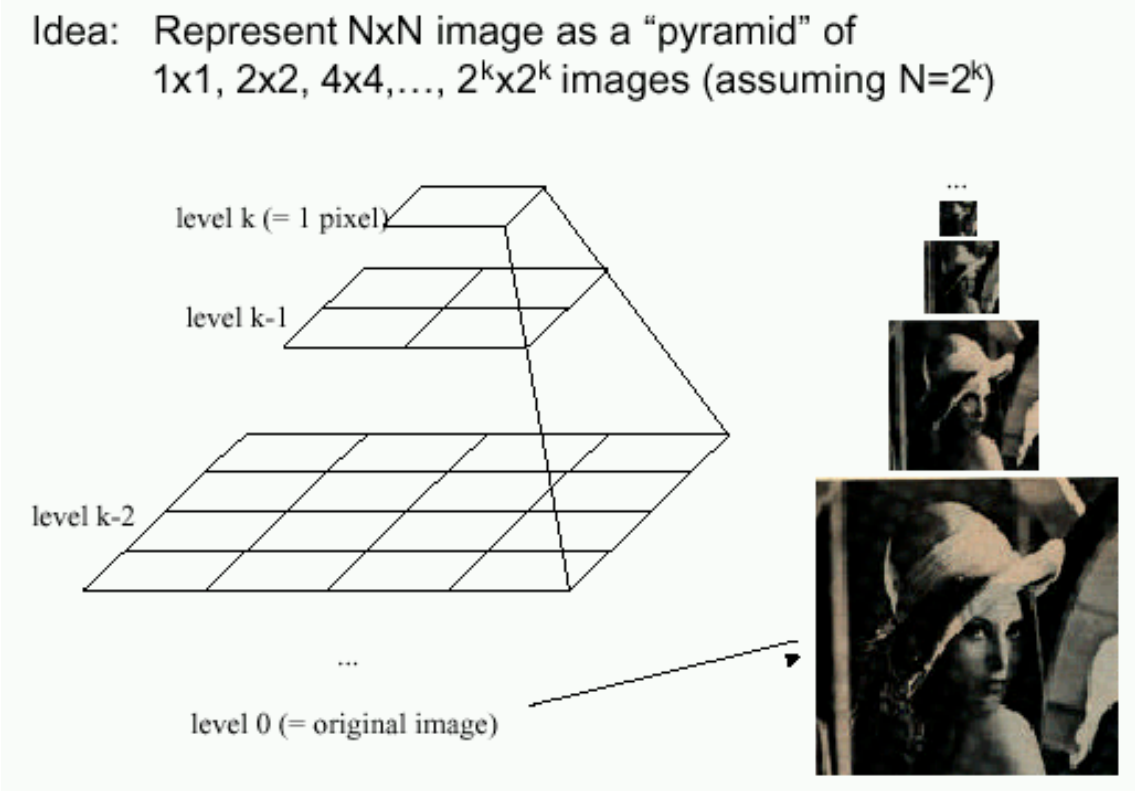


- How much space does a Gaussian pyramid take compared to original image?

[Source: S. Seitz]

# Example of Gaussian Pyramid

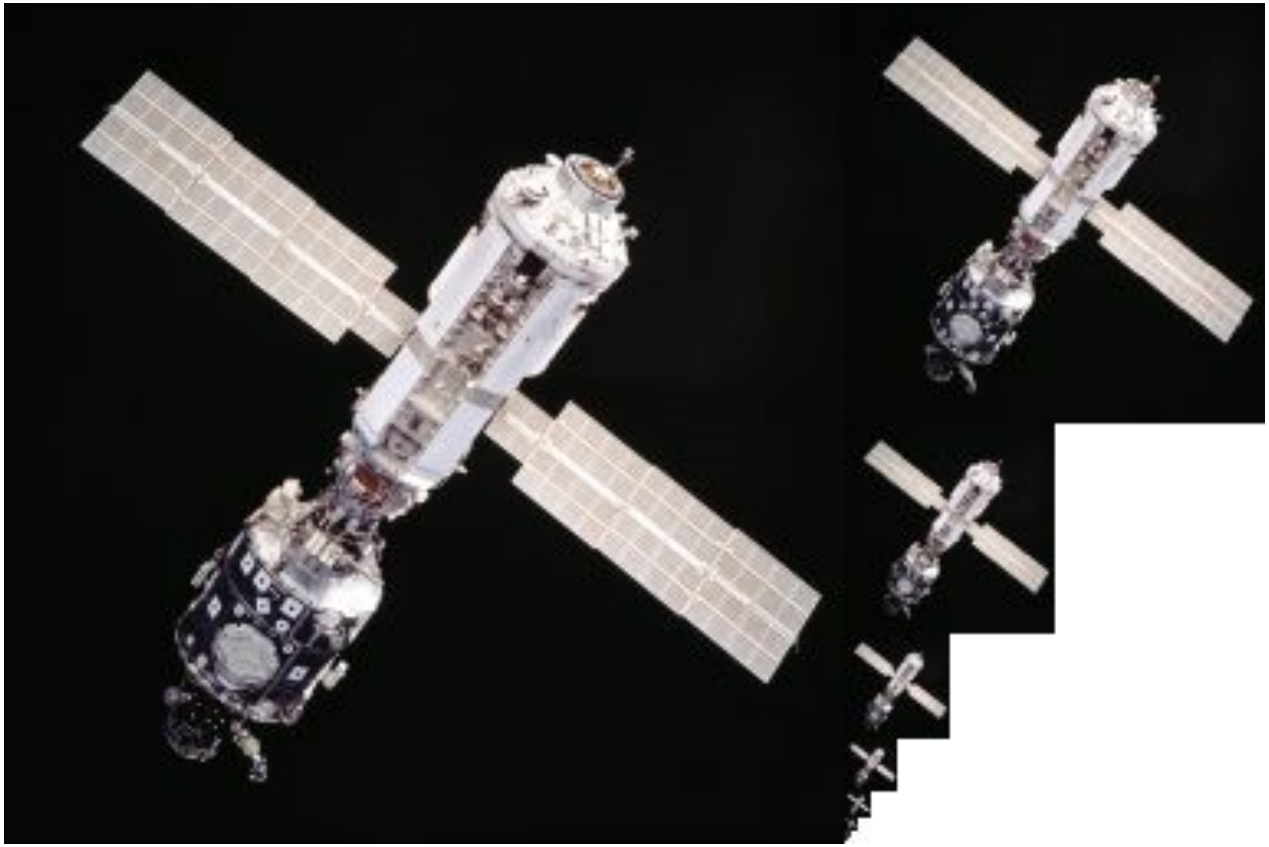

[Source: N. Snavely]

# Summary – Stuff You Should Know

- To down-scale an image: blur it with a small Gaussian (e.g., $\sigma = 1.4$) and downsample
- To up-scale an image: interpolation (we won't cover this in class)
- Gaussian pyramid: Blur with Gaussian filter, downsample result by factor 2, blur it with the Gaussian, downsample by 2...

**Matlab functions**:

- FSPECIAL: creates a Gaussian filter with specified $\sigma$
- IMFILTER: convolve image with the filter
- I(1:2:END, 1:2:END): takes every second row and column
- IMRESIZE: Matlab's function for resizing the image