CSC589 Introduction to Computer Vision
Lecture 10

Gaussian Pyramid

Bei Xiao

Man-made Scene

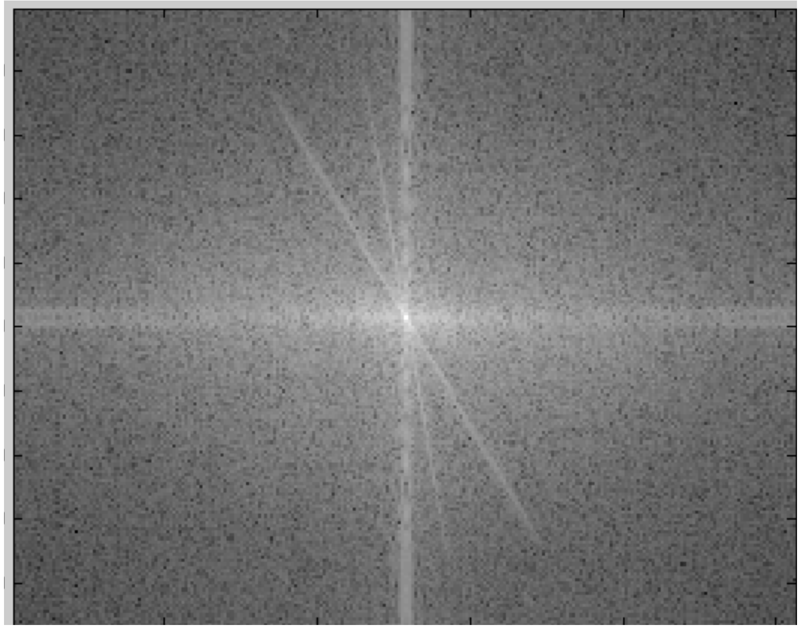

Can change spectrum, then reconstruct

Low and High Pass filtering

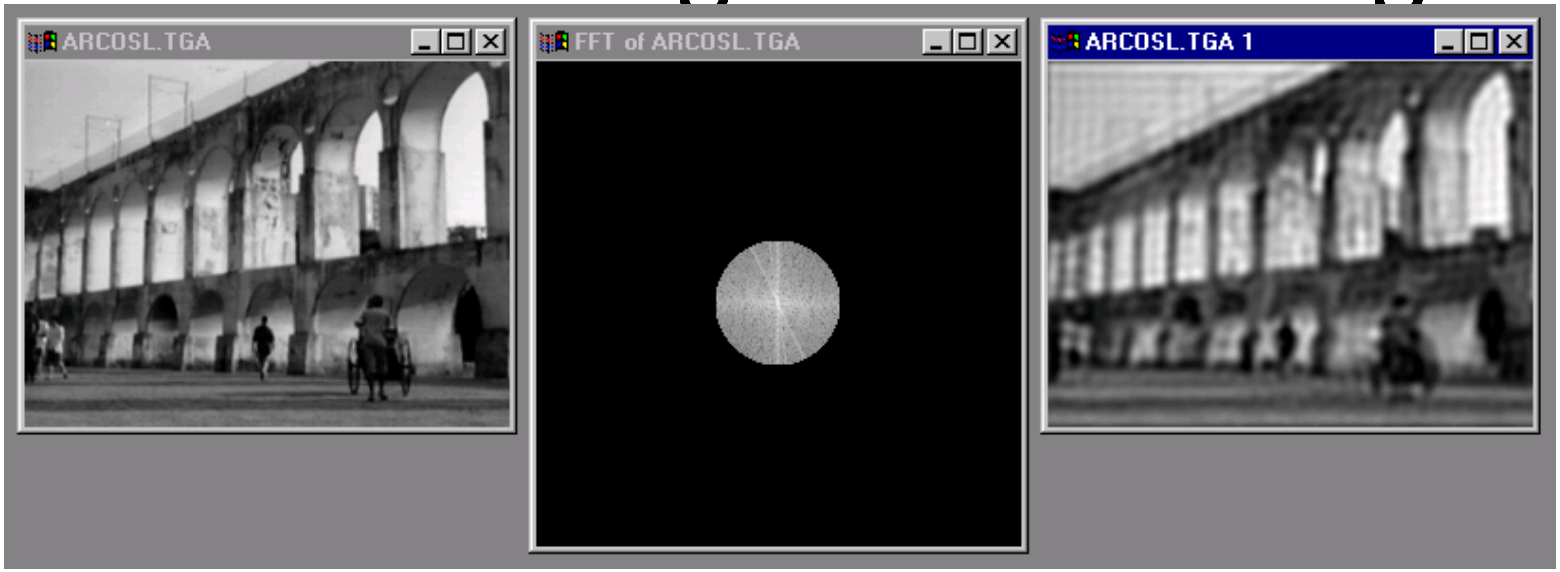

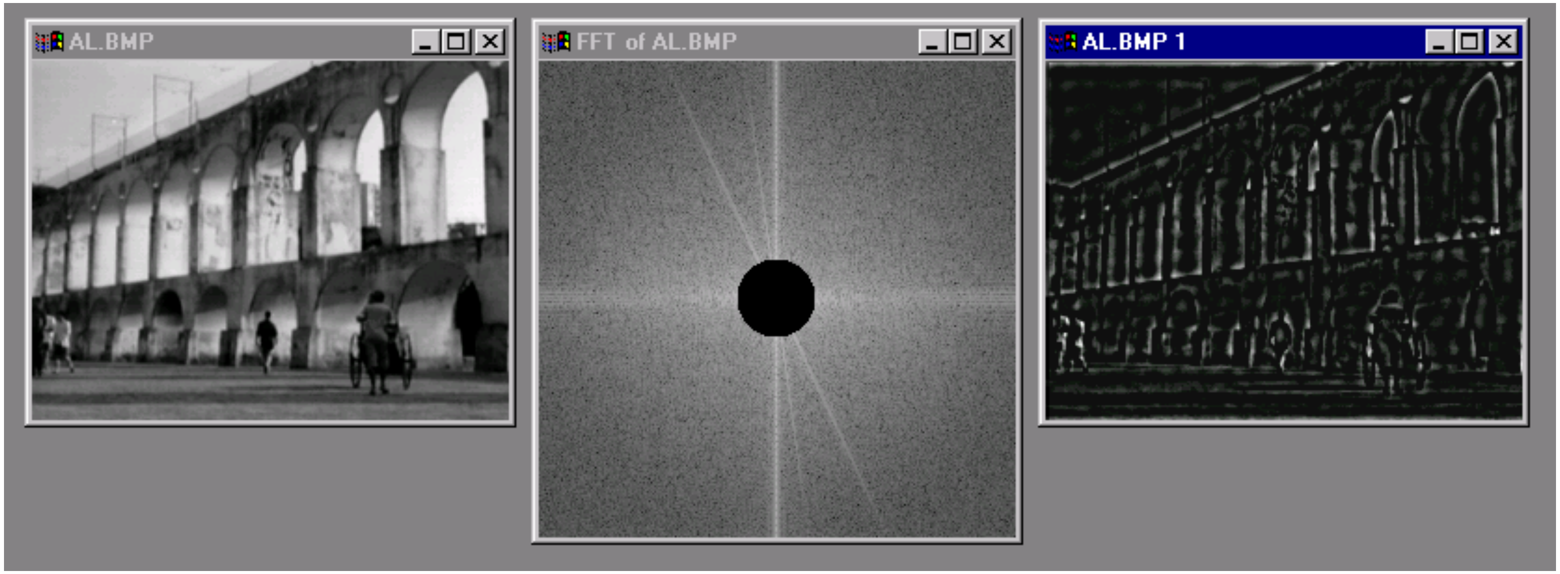

Today’s class

- Image Pyramids

Image representation

- Pixels: great for spatial resolution, poor access to frequency

- Fourier transform: great for frequency, not for spatial info

- Pyramids/filter banks: balance between spatial and frequency information

Fourier vs. Spatial Representation

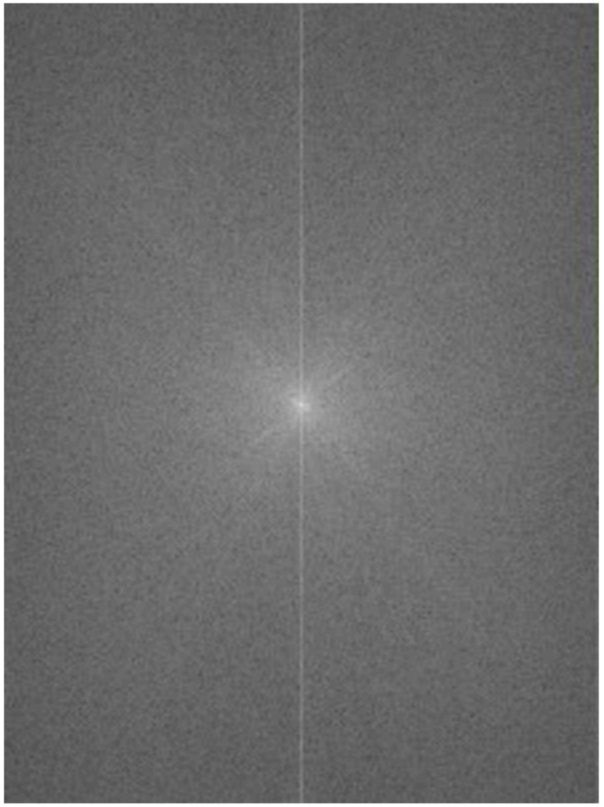

Image

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

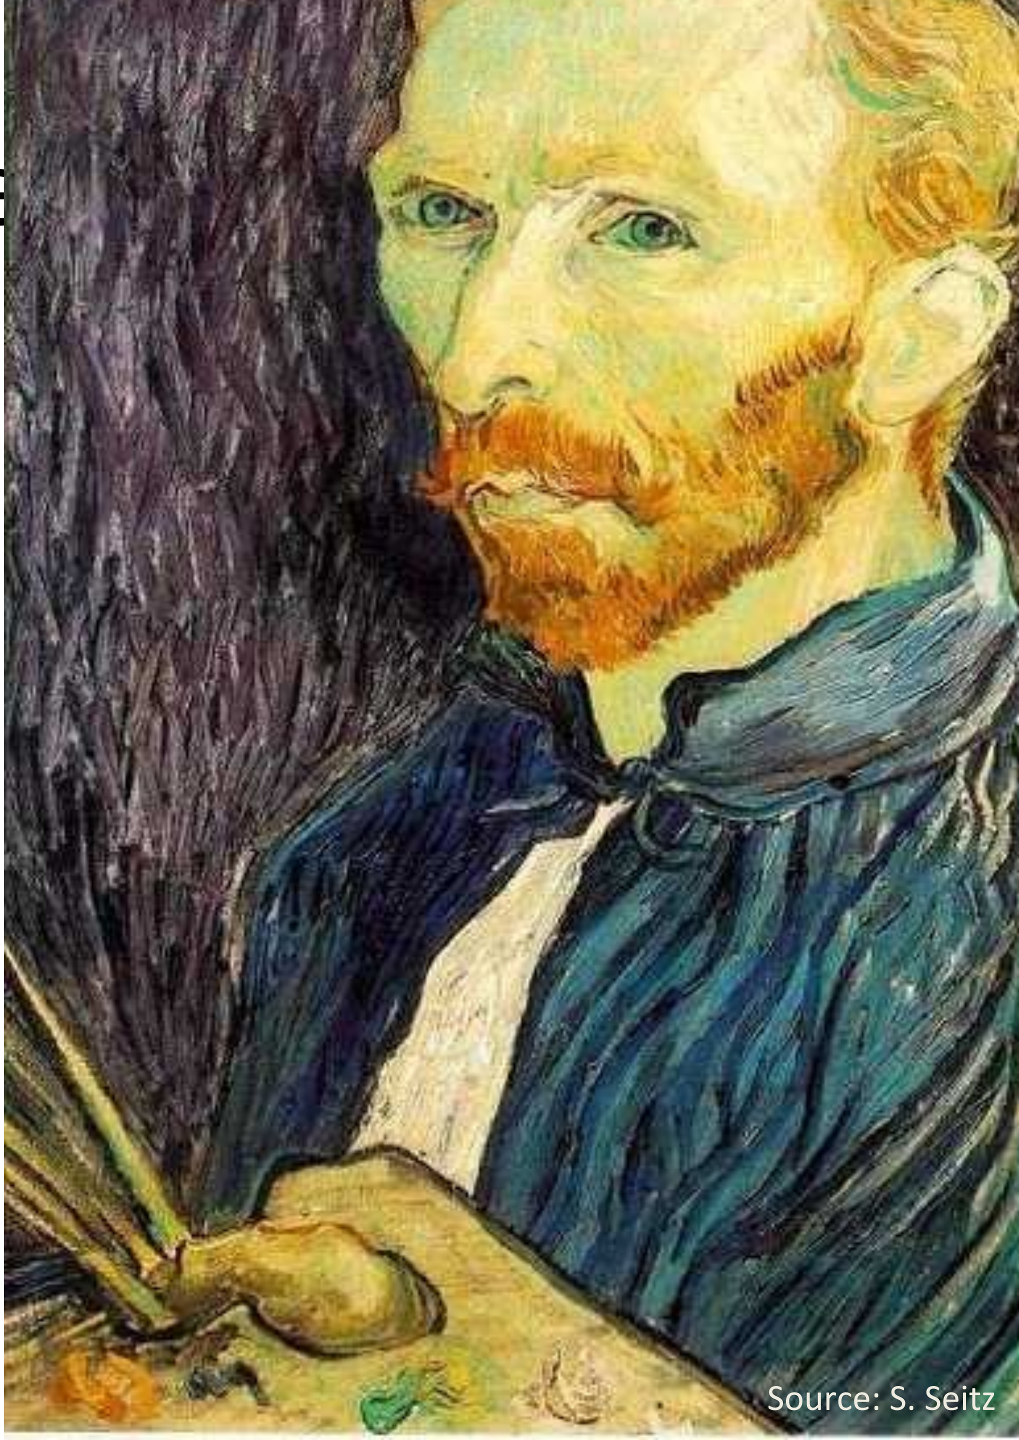

Image sub-sampling

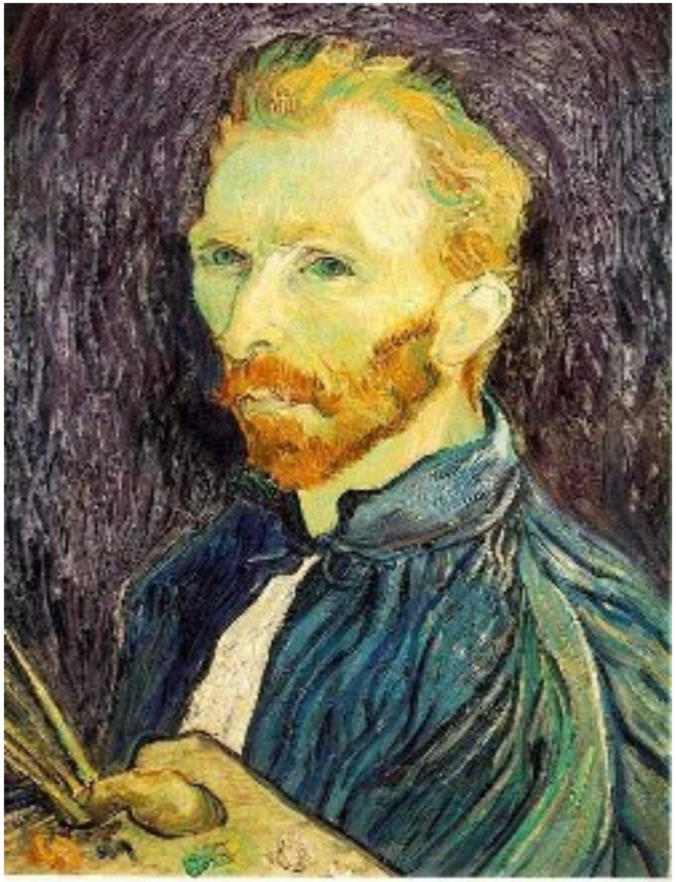

1/4

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

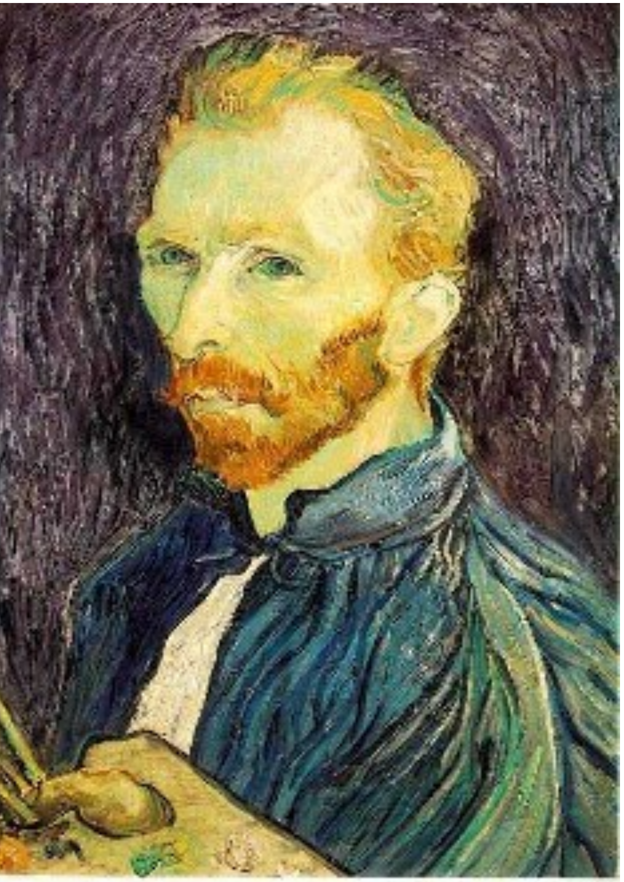

1/2

1/4 (2x zoom)

1/8 (4x zoom)

Why does this look so crufty?

Source: S. Seitz

Image sub-sampling

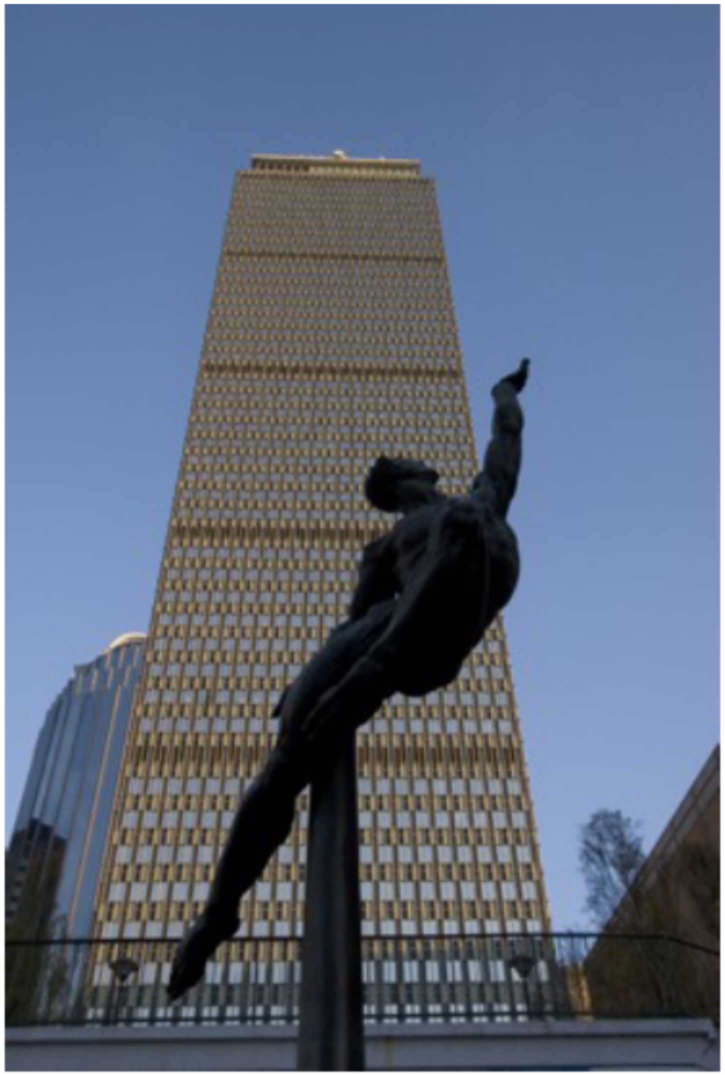

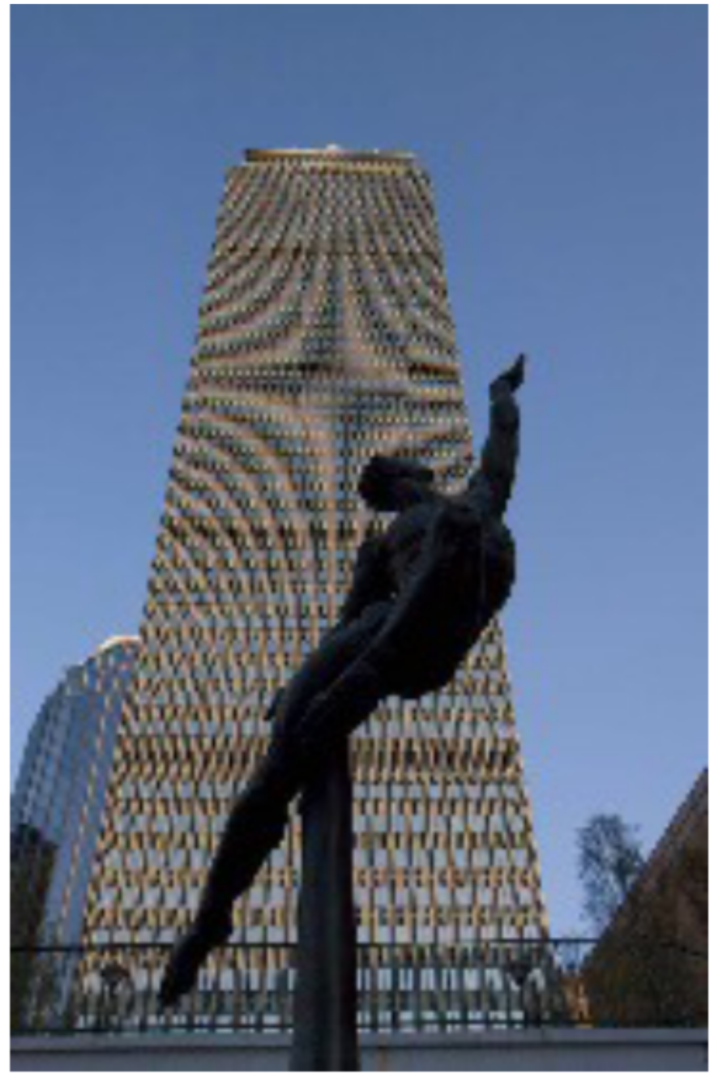

Source: F. Durand

Nyquist limit – 2D example

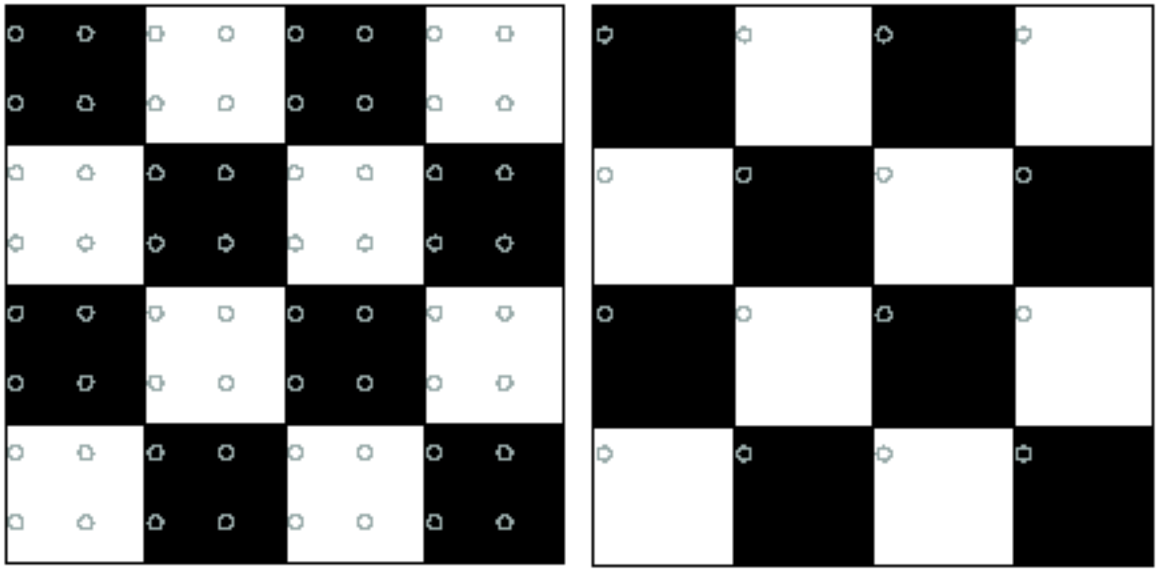

Good sampling

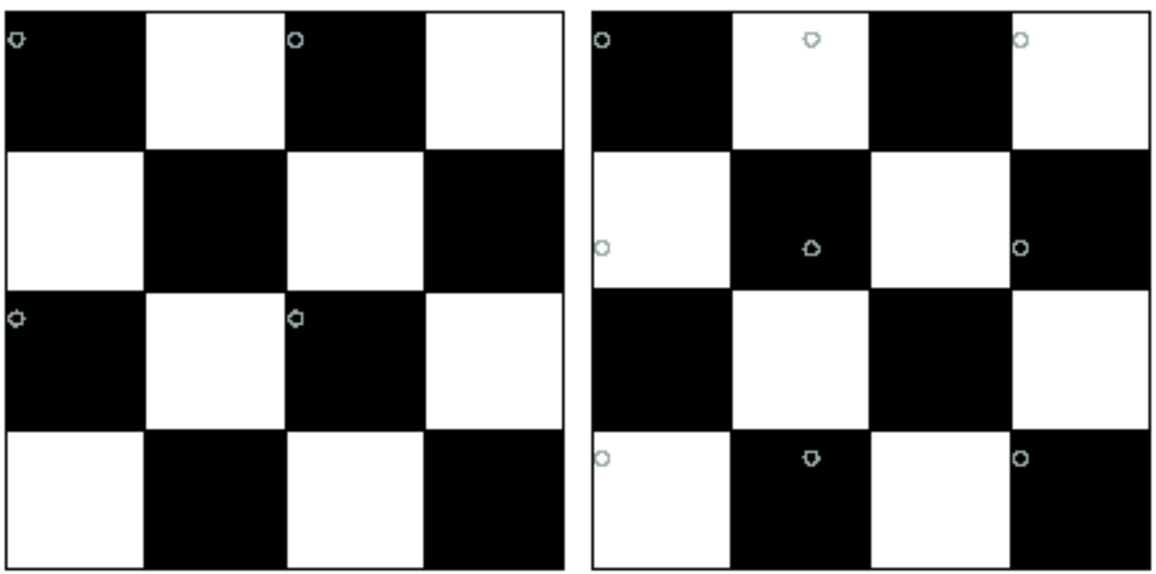

Bad sampling

Aliasing

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

- How can we fix this?

Gaussian pre-filtering

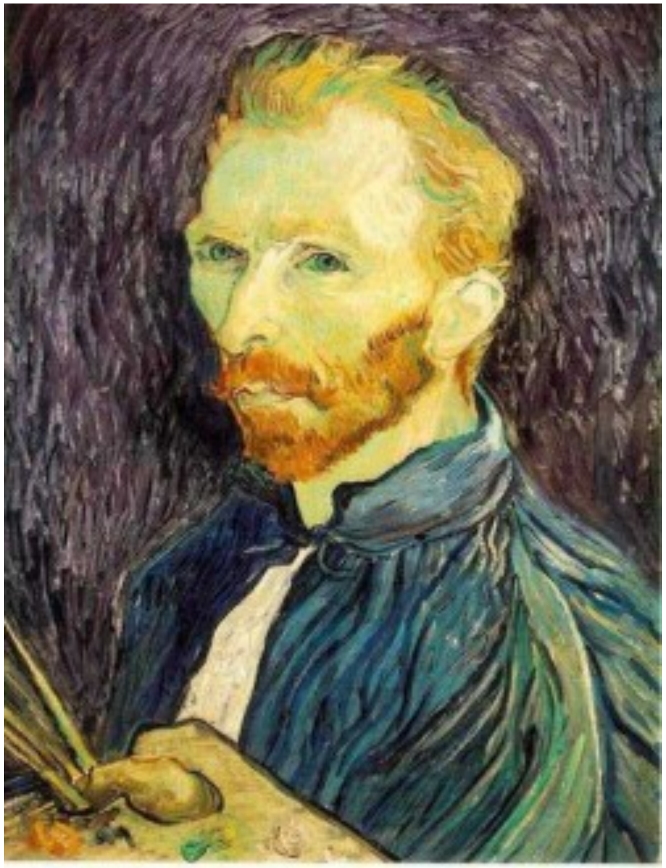

Gaussian 1/2

G 1/4

G 1/8

- Solution: filter the image, *then* subsample

Source: S. Seitz

Subsampling with Gaussian pre-filtering

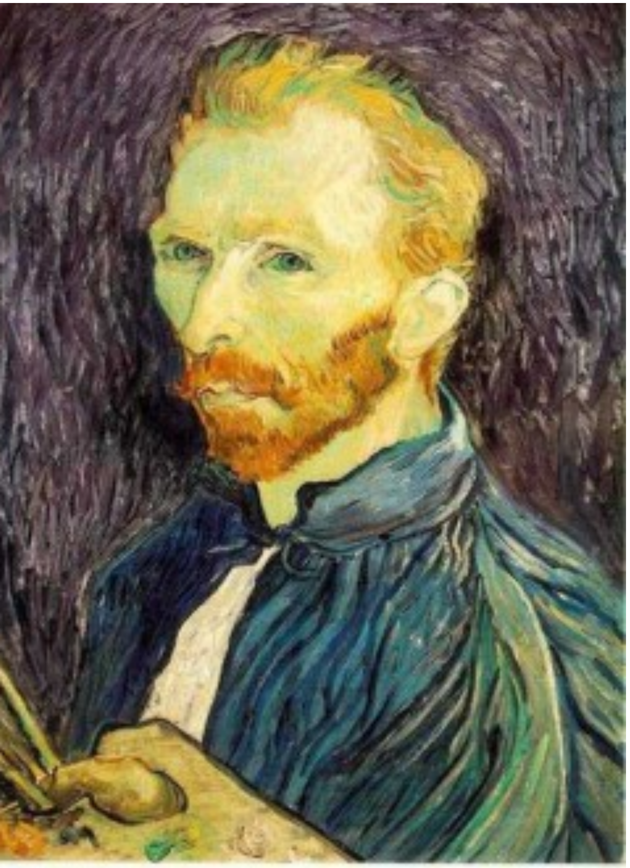

Gaussian 1/2

G 1/4

G 1/8

- Solution: filter the image, *then* subsample

Source: S. Seitz

Compare with...

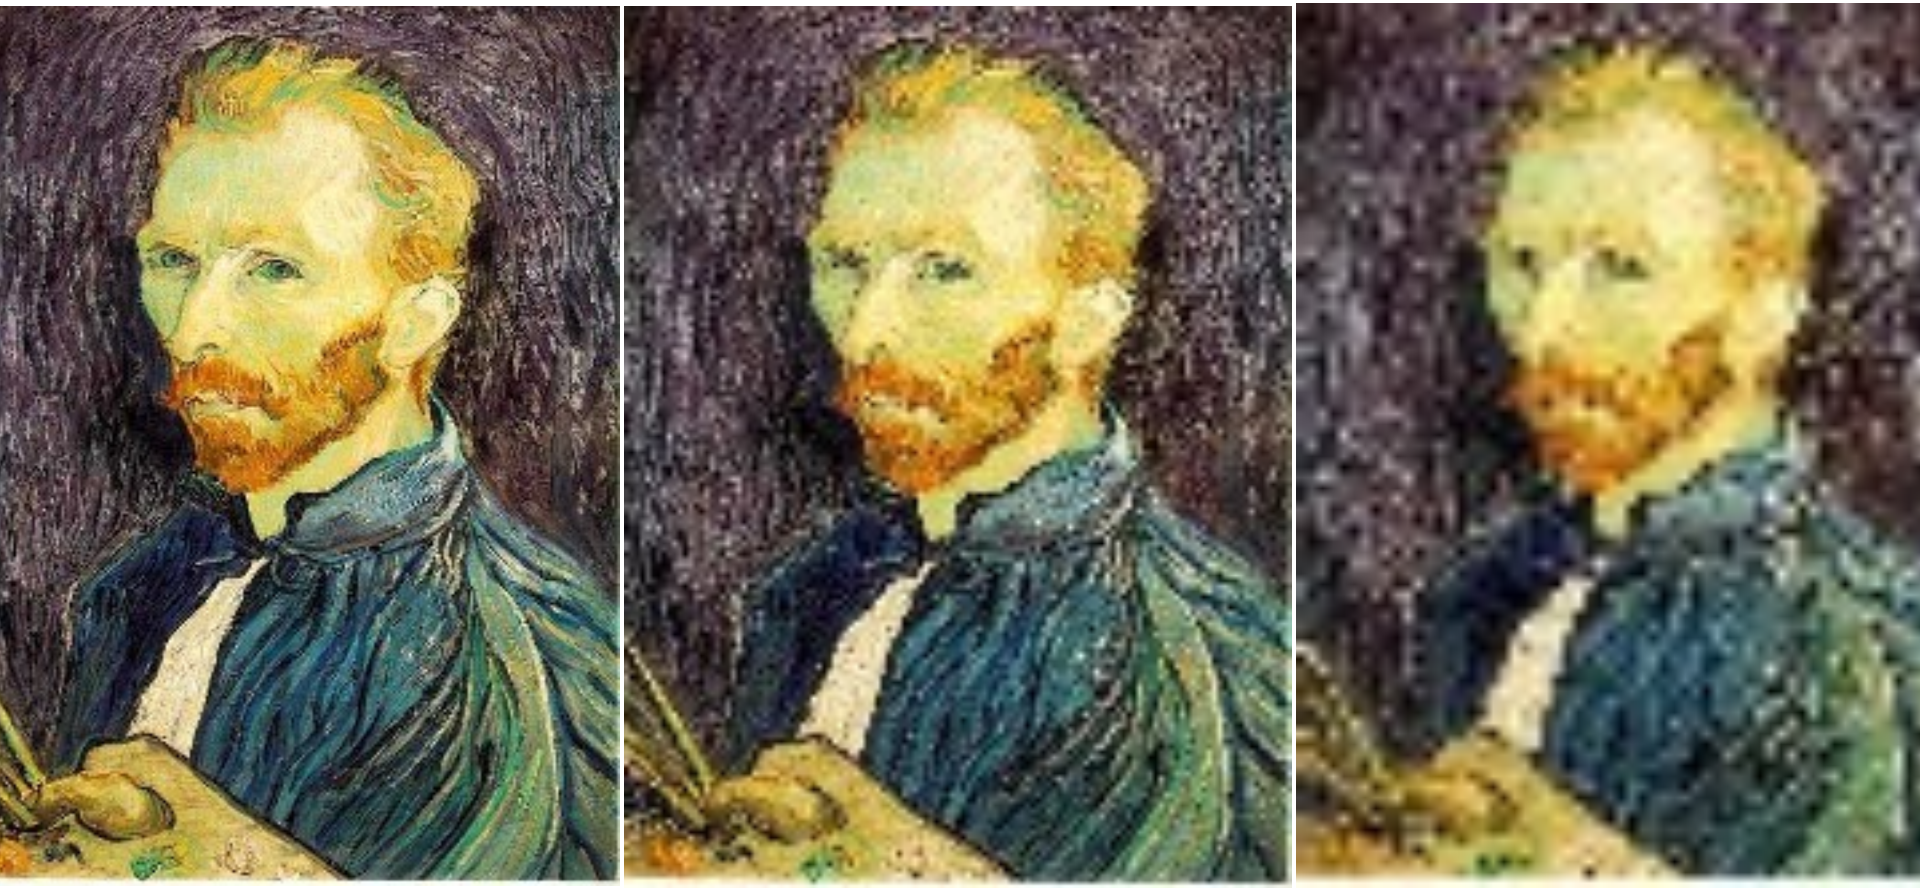

1/2 1/4 (2x zoom) 1/8 (4x zoom)

Source: S. Seitz

Gaussian pre-filtering

- Solution: filter the image, *then* subsample

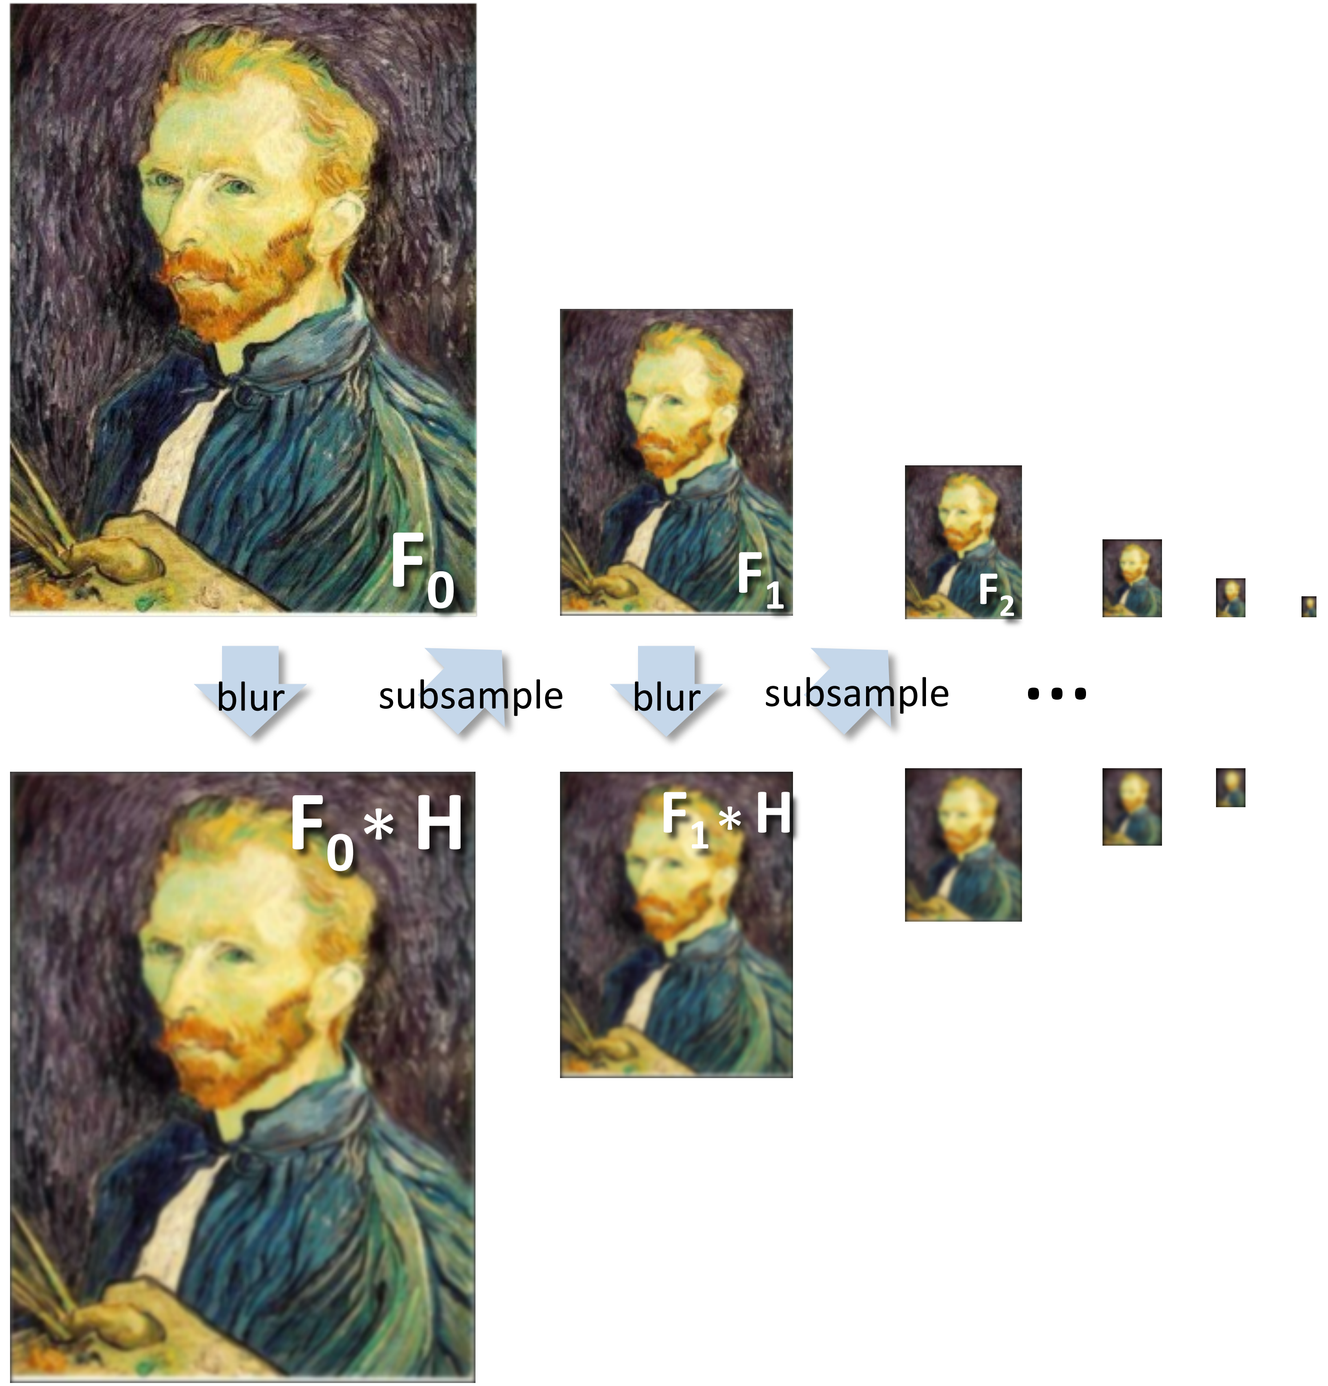

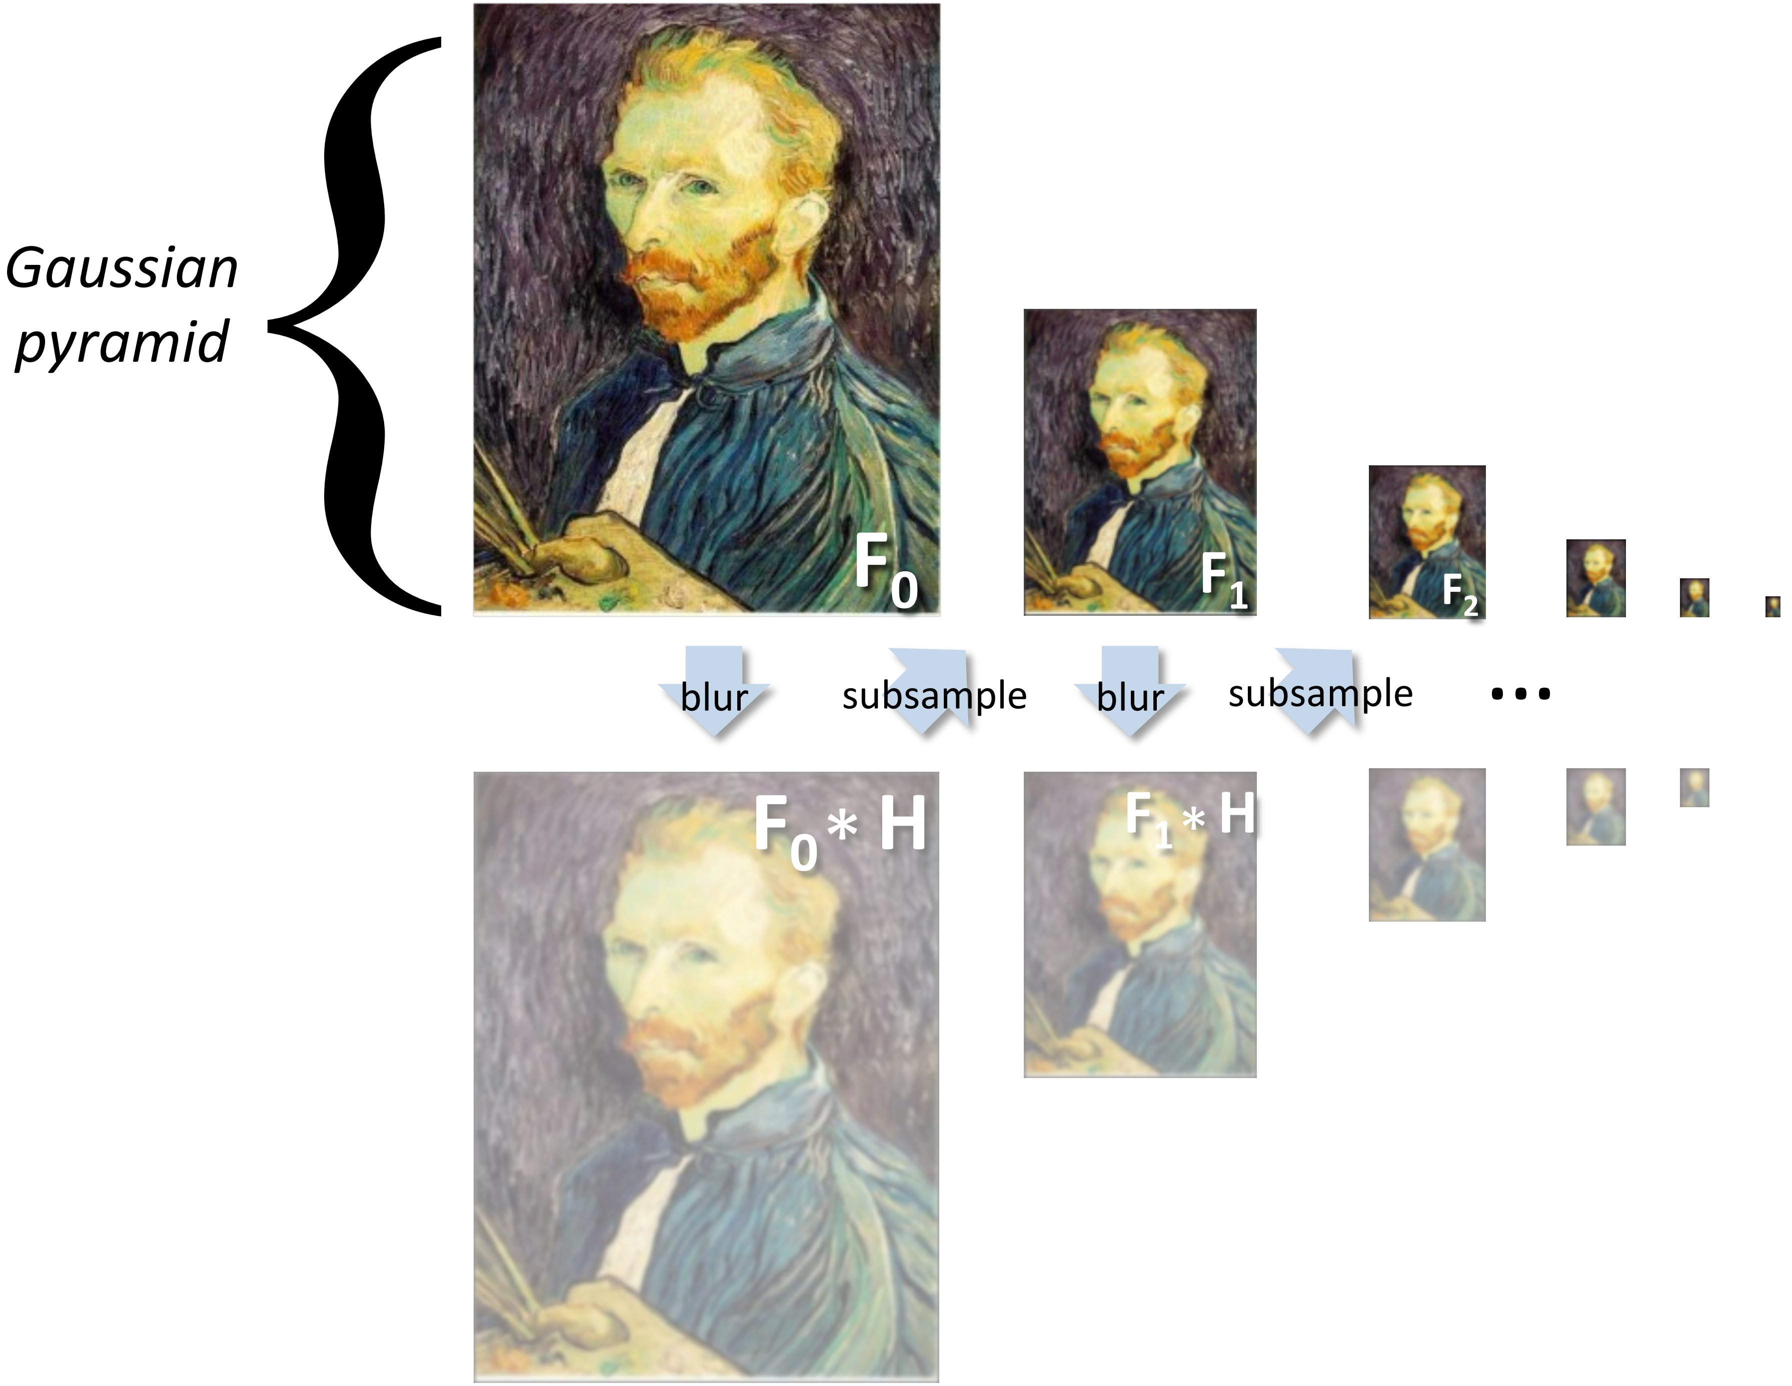
Gaussian pyramid
F0
F1
F2
blur
subsample
blur
subsample
...
F0 * H
F1 * H

Why is it called a pyramid

- Idea: Representation can be pictured as a "pyramid" of 3×3, 5×5, 9×9, …, $(2^n+1)(2^n+1)$ images.

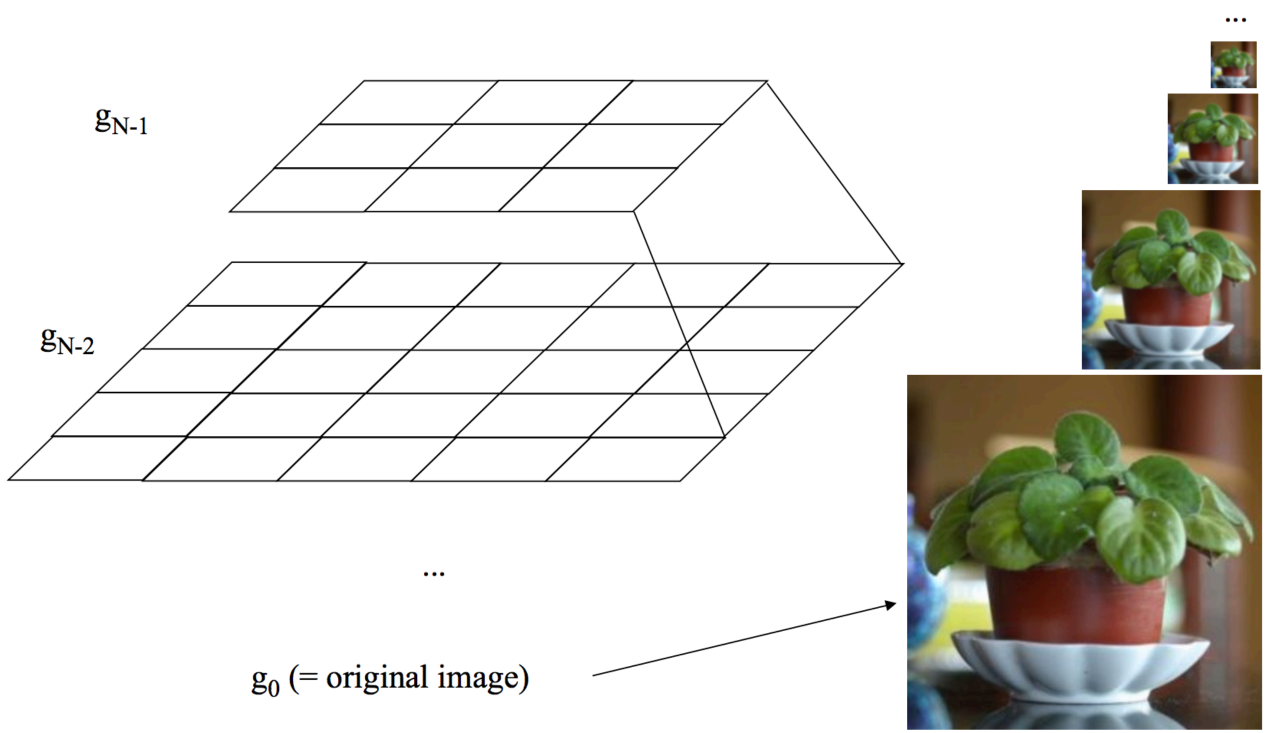

Gaussian pyramids
[Burt and Adelson, 1983]

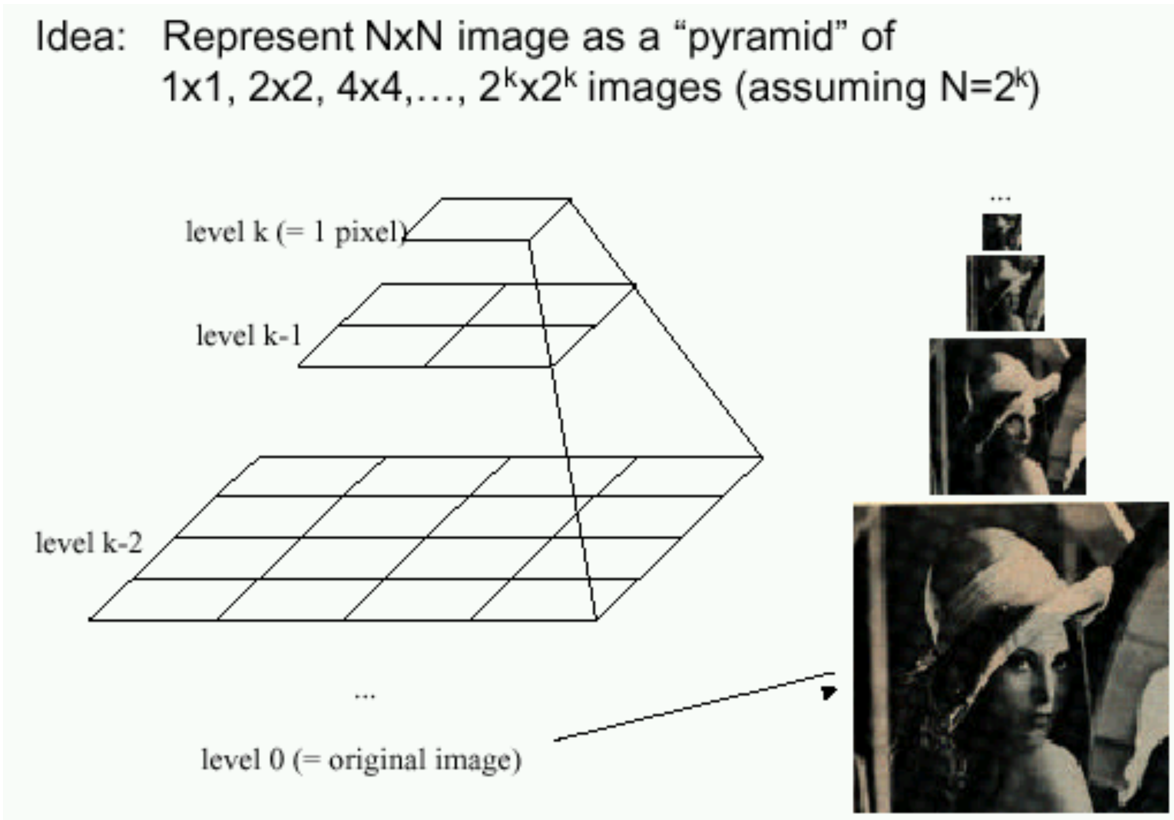

- In computer graphics, a *mip map* [Williams, 1983]
- A precursor to *wavelet transform*

Gaussian Pyramids have all sorts of applications in computer vision

Source: S. Seitz

Gaussian pyramids
[Burt and Adelson, 1983]

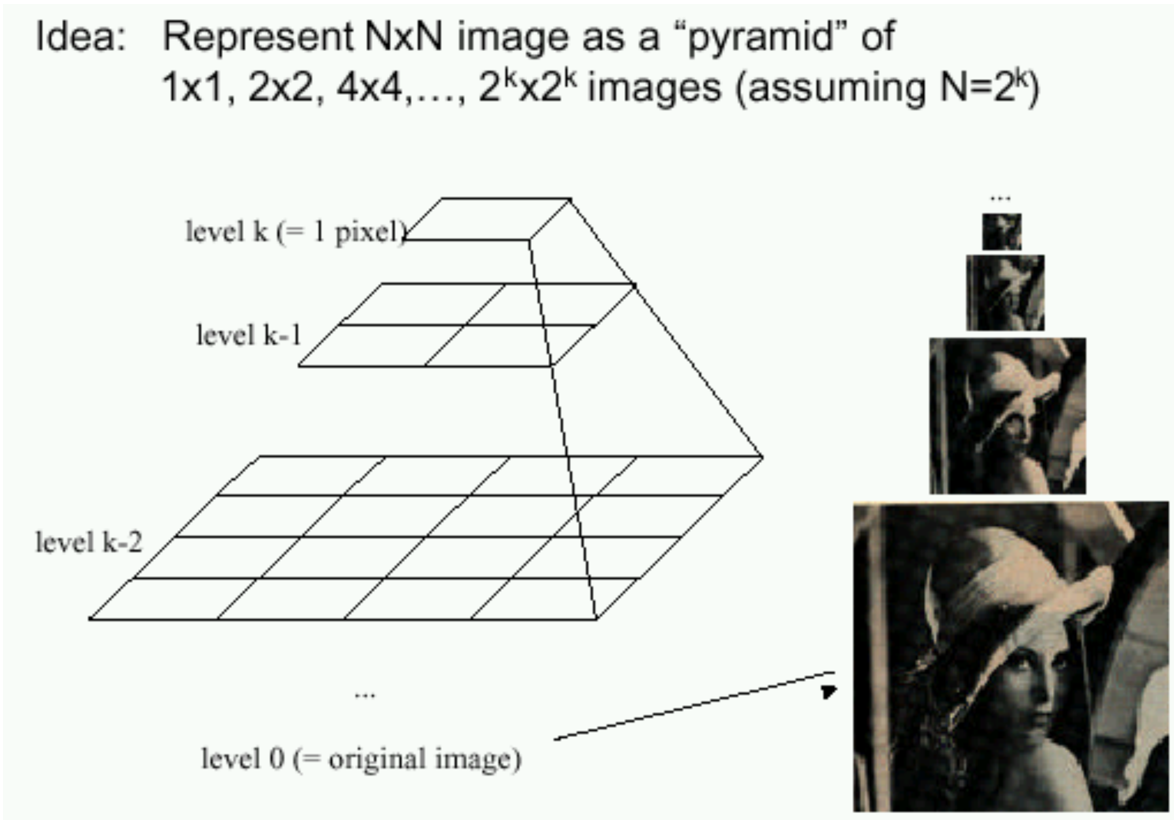

- How much space does a Gaussian pyramid take compared to the original image?

Source: S. Seitz

Gaussian Pyramid

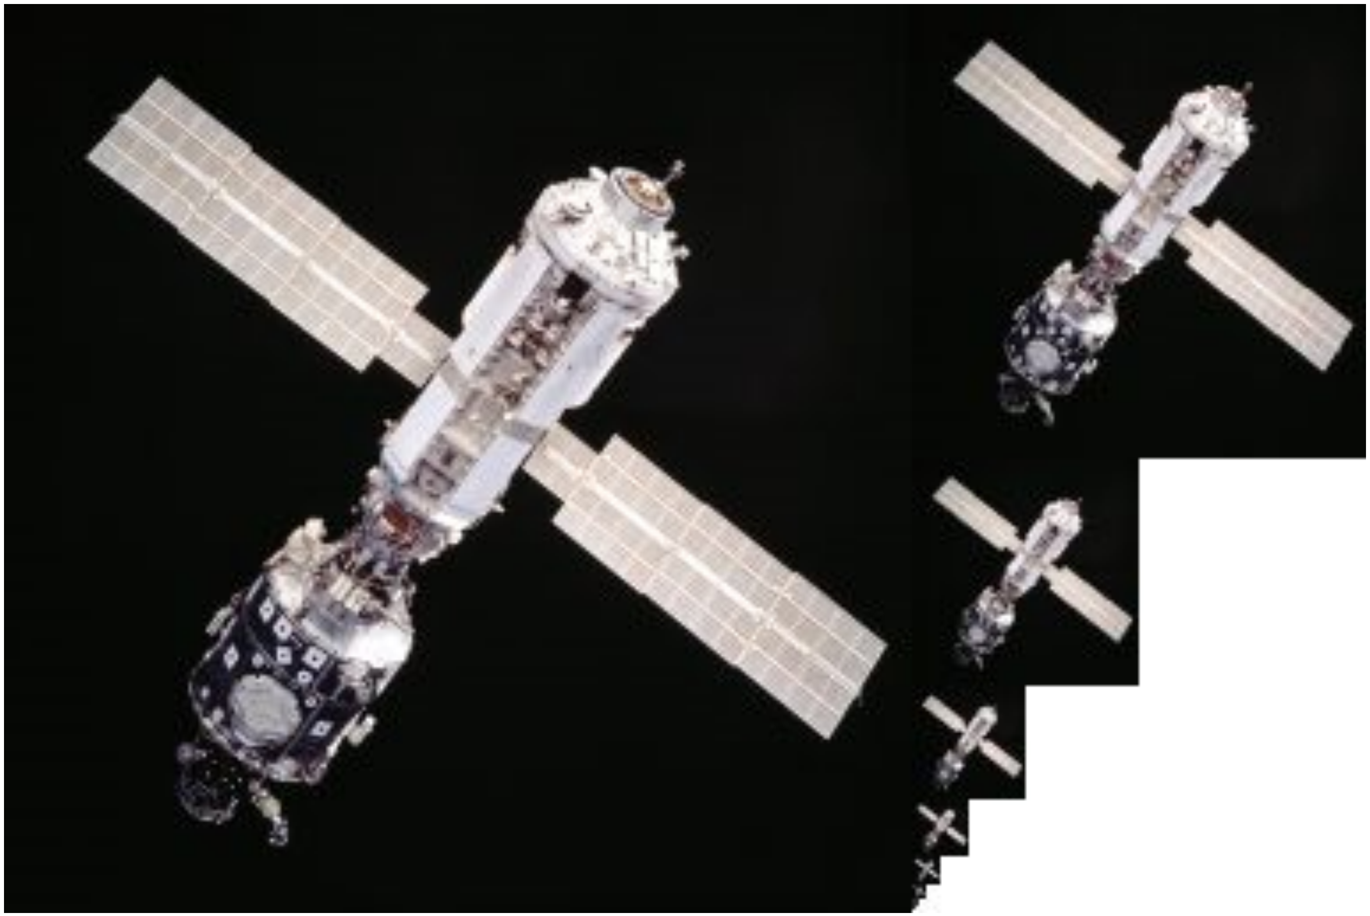

Gaussian pyramid

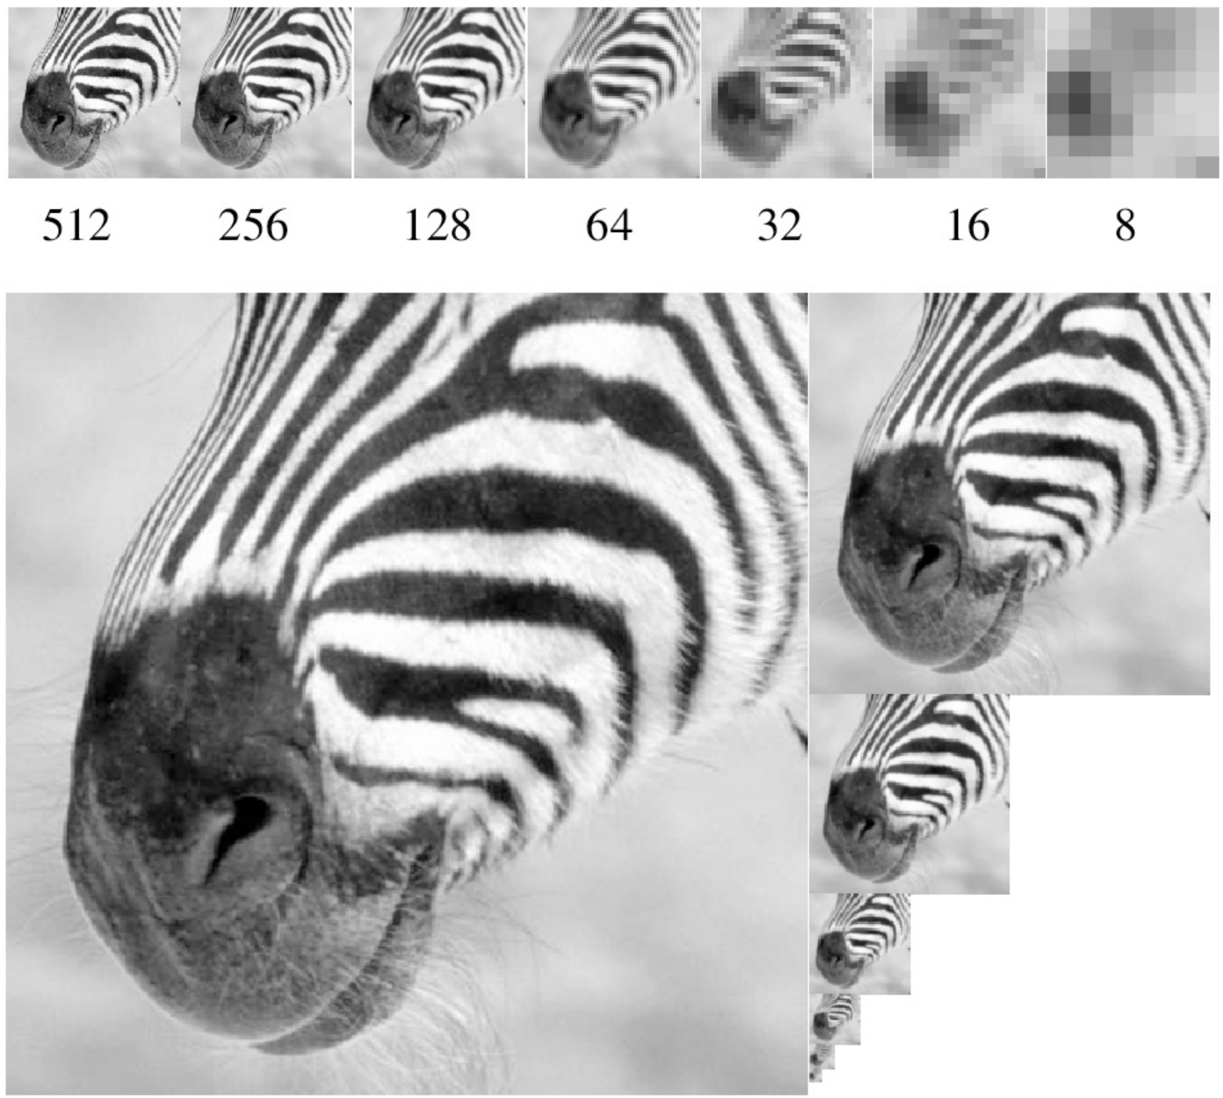

Source: Forsyth

Coarse-to-fine Image Registration

1. Compute Gaussian pyramid
2. Align with coarse pyramid
3. Successively align with finer pyramids
   - Search smaller range

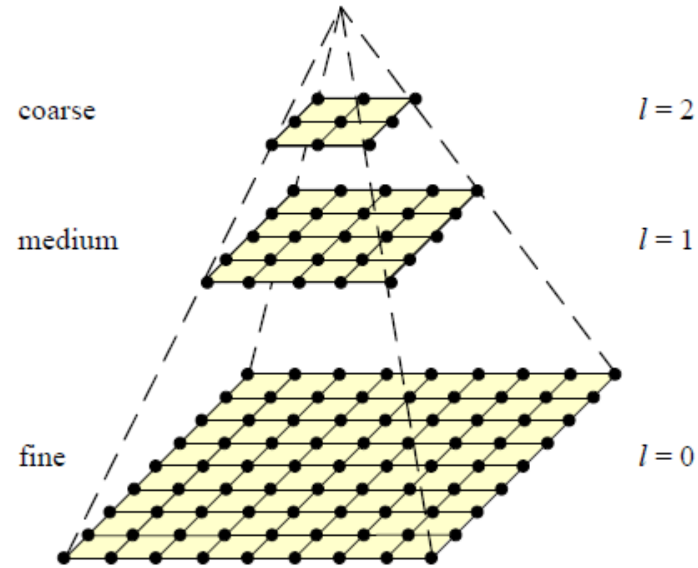

Why is this faster?

Are we guaranteed to get the same result?

Laplacian filter

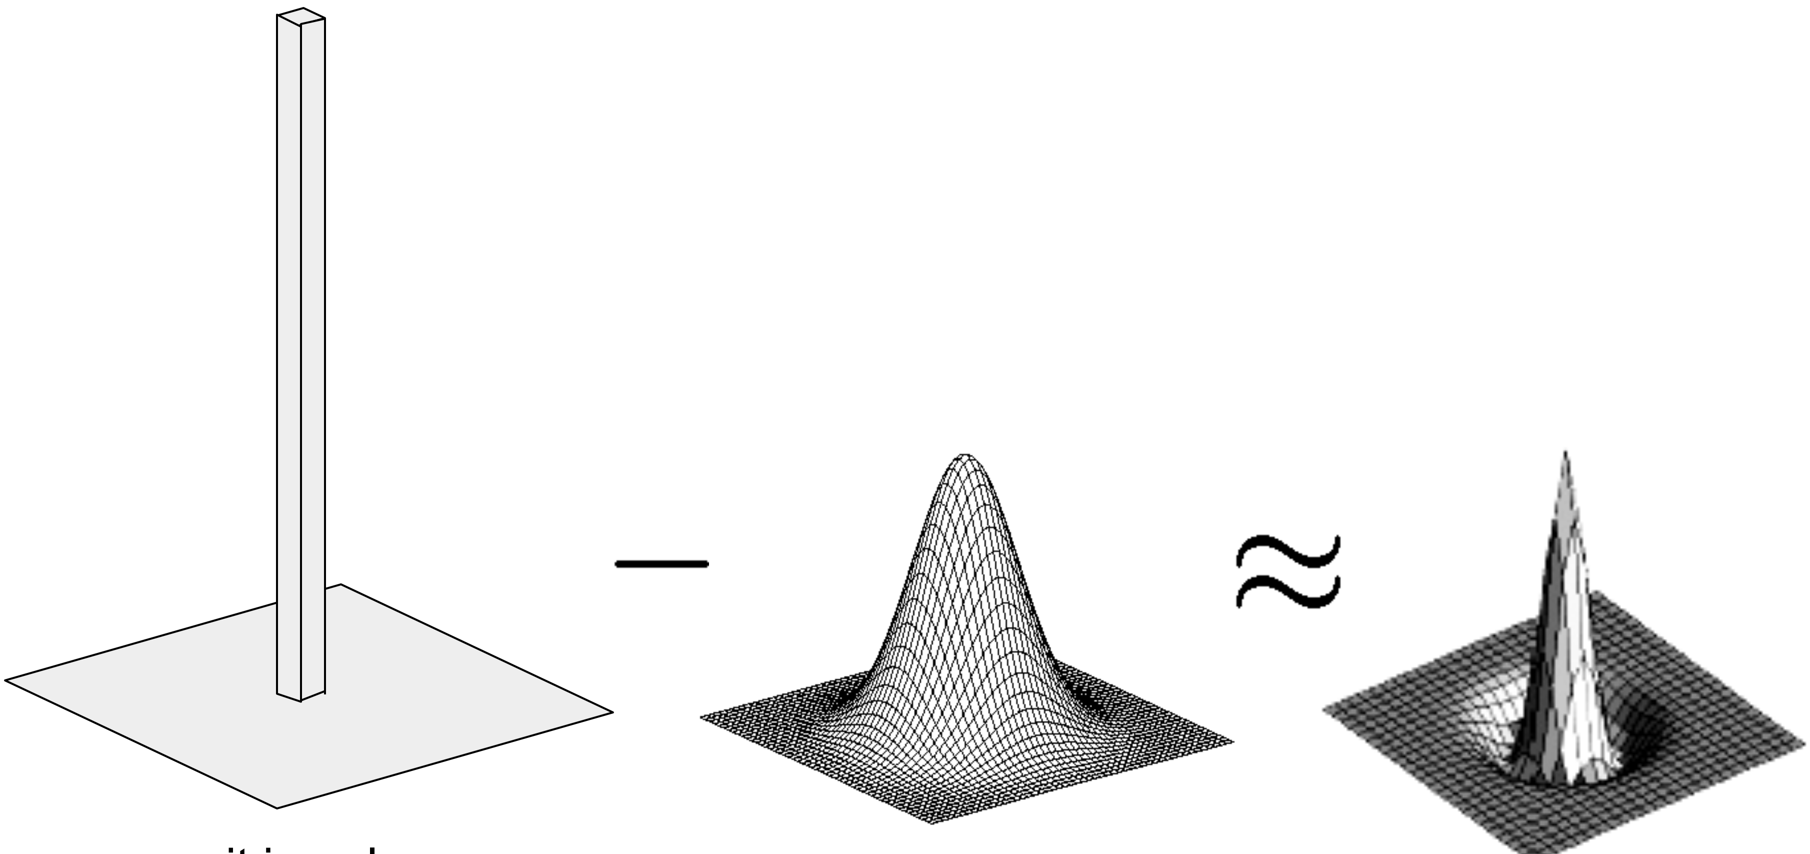

unit impulse

Gaussian

Laplacian of Gaussian

Source: Lazebnik

2D edge detection filters

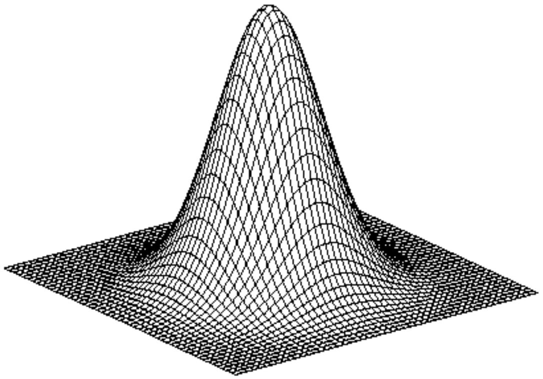

Gaussian

$$h_\sigma(u,v) = \frac{1}{2\pi\sigma^2} e^{-\frac{u^2+v^2}{2\sigma^2}}$$

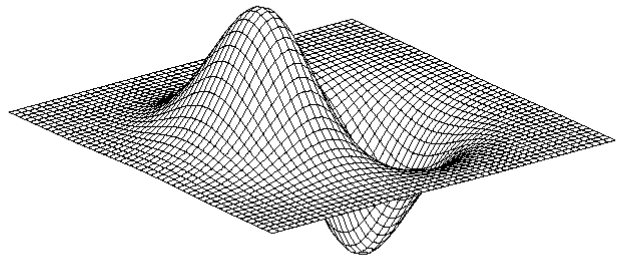

derivative of Gaussian

$$\frac{\partial}{\partial x} h_\sigma(u,v)$$

Laplacian of Gaussian

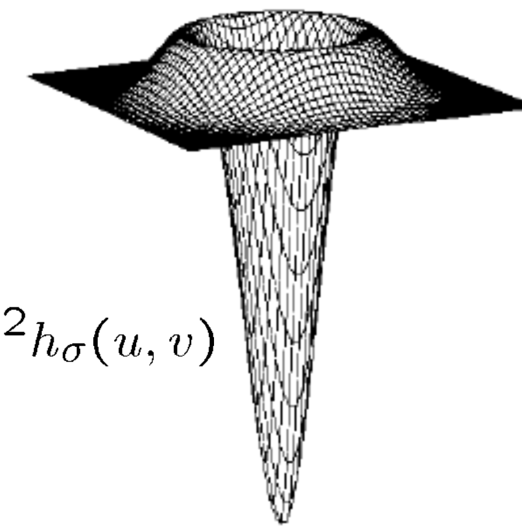

$$\nabla^2 h_\sigma(u,v)$$

∇^2 is the **Laplacian** operator:

$$\nabla^2 f = \frac{\partial^2 f}{\partial x^2} + \frac{\partial^2 f}{\partial y^2}$$

Computing Gaussian/Laplacian Pyramid

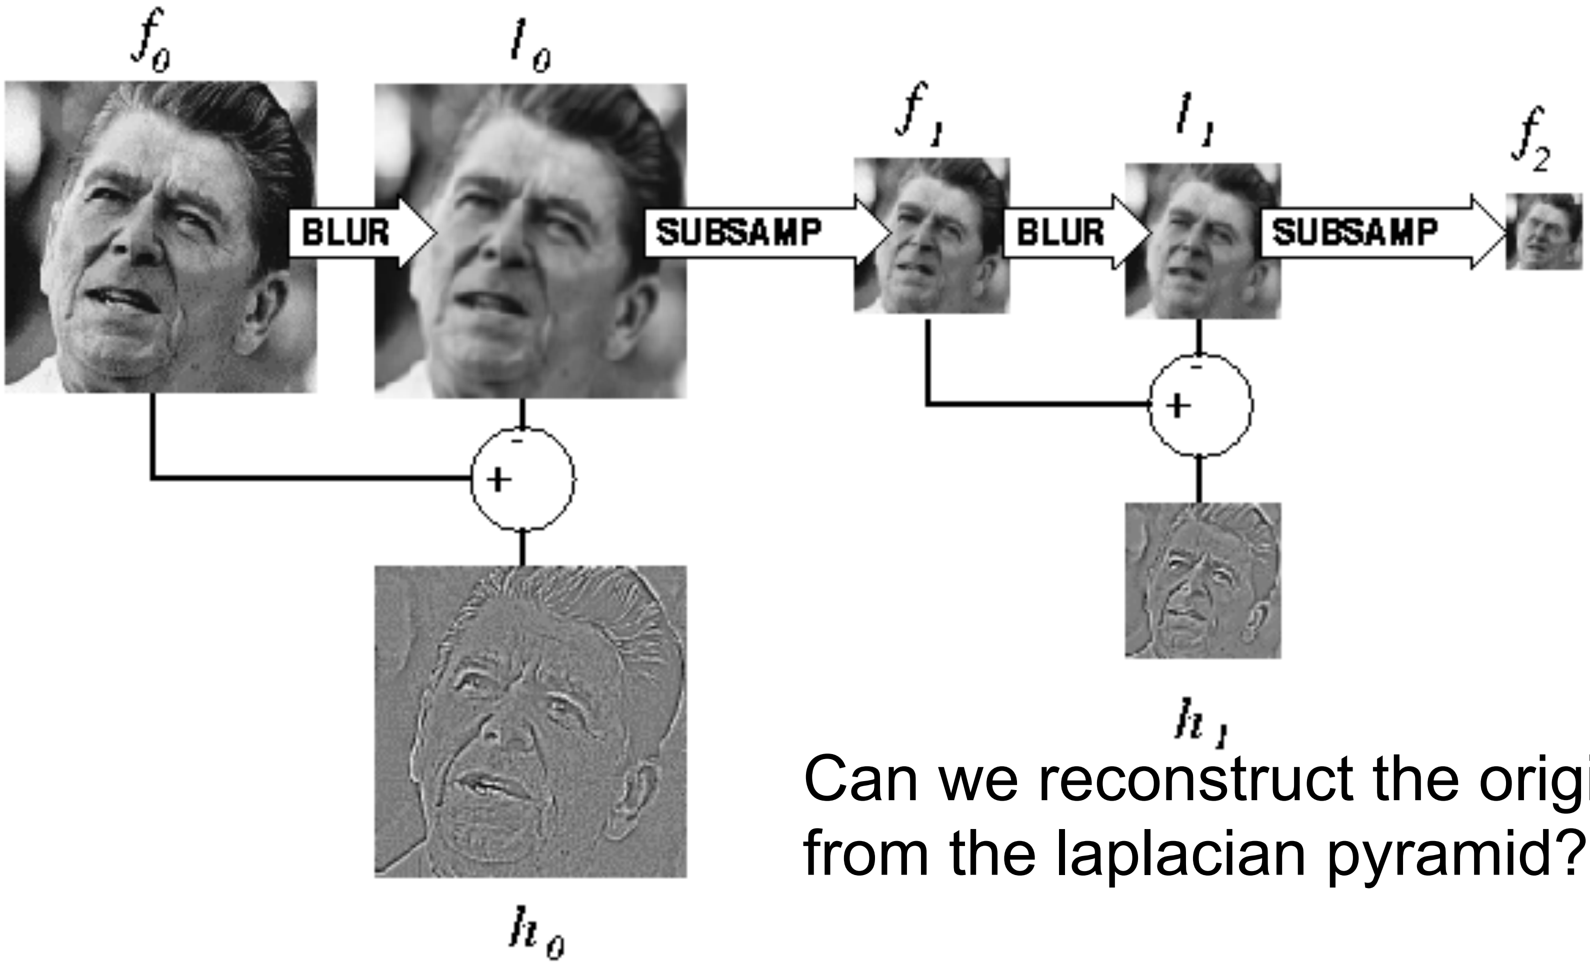

Can we reconstruct the original from the laplacian pyramid?

http://sepwww.stanford.edu/~morgan/texturematch/paper_html/node3.html

Can we reconstruct the original from the laplacian pyramid?

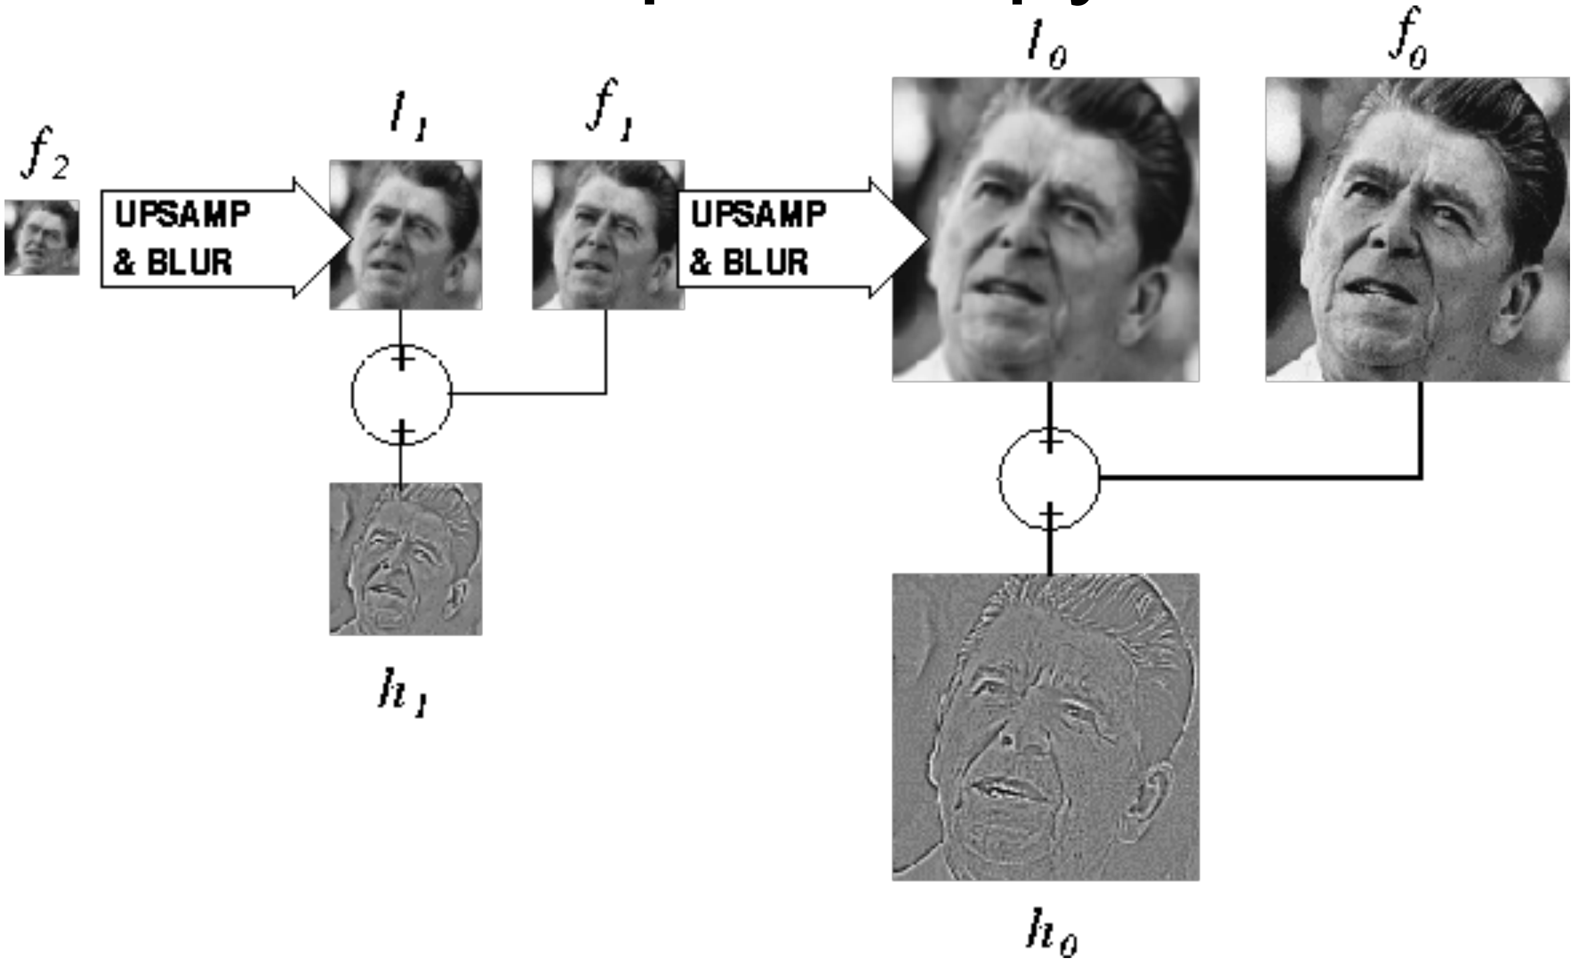

http://sepwww.stanford.edu/~morgan/texturematch/paper_html/node3.html

Laplacian Pyramid

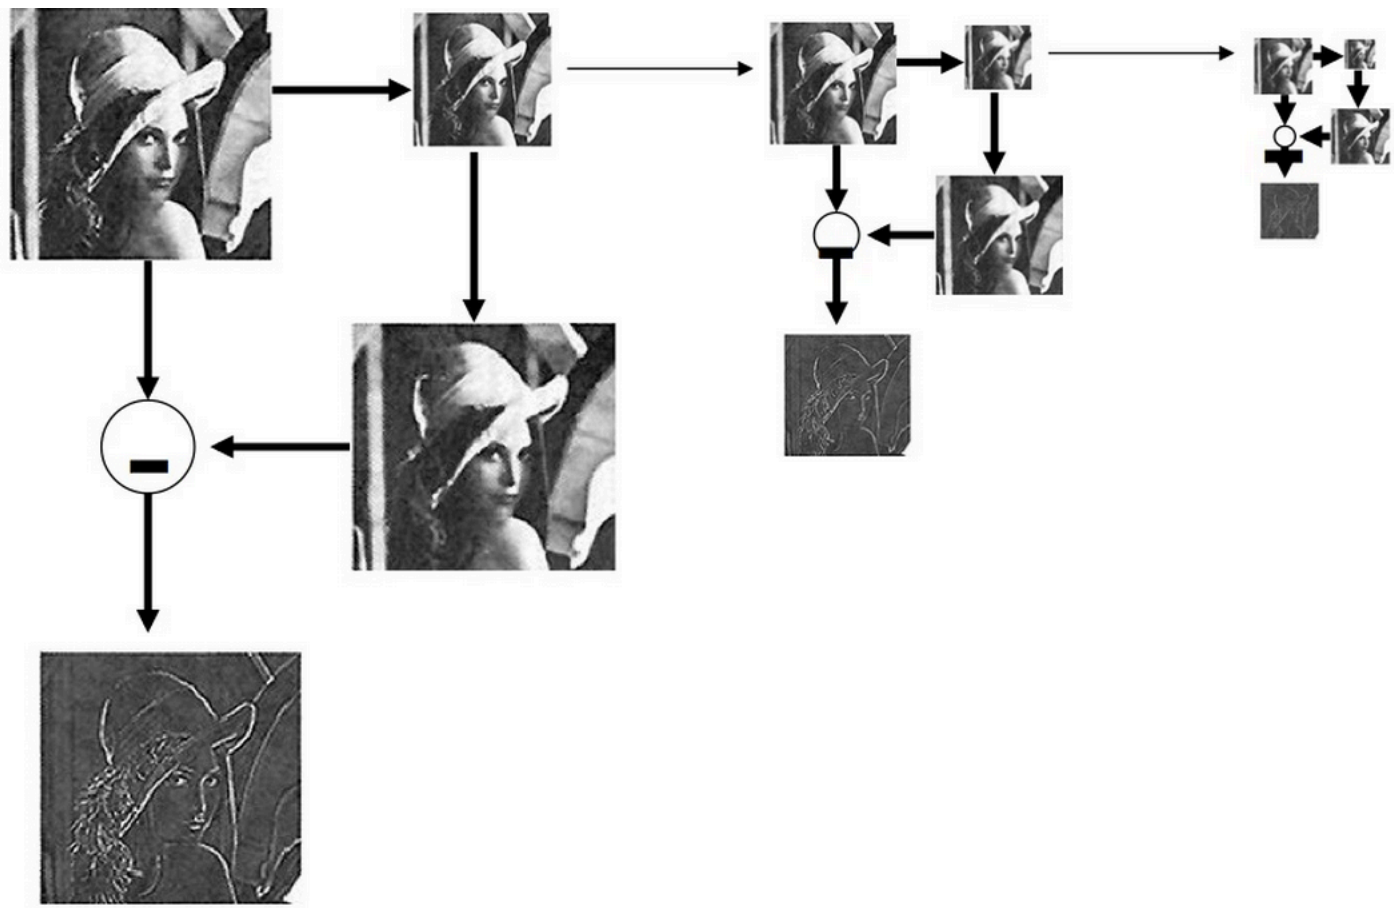

Laplace pyramid

Rather than store the smoothed images, store only the difference between the levels g_l and g_{l+1}

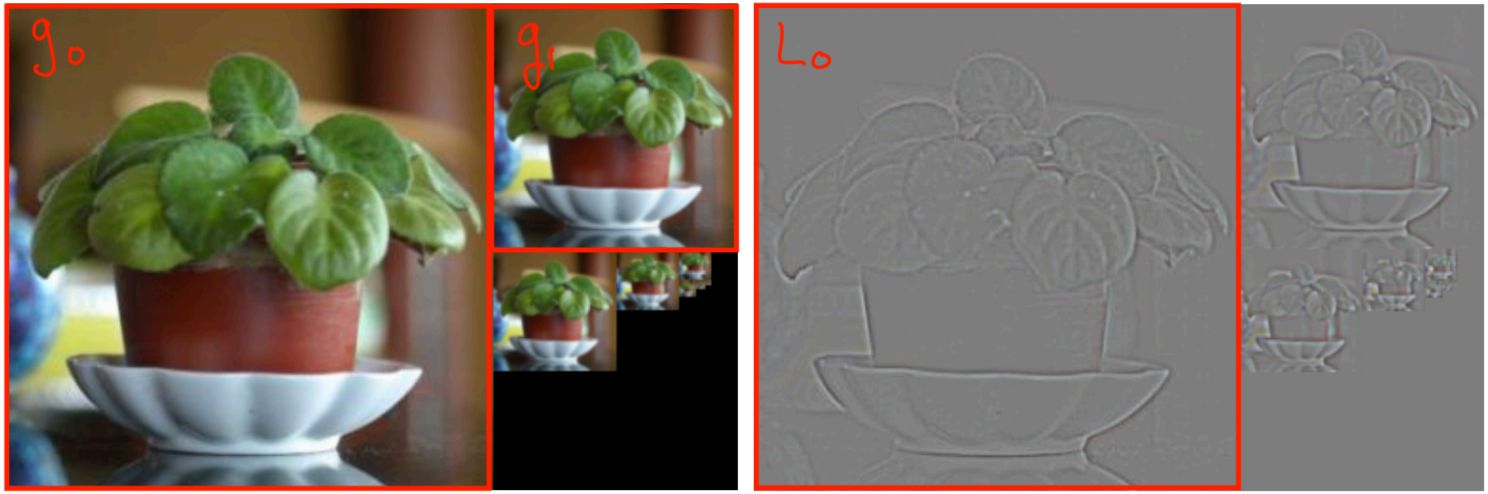

Laplacian pyramid

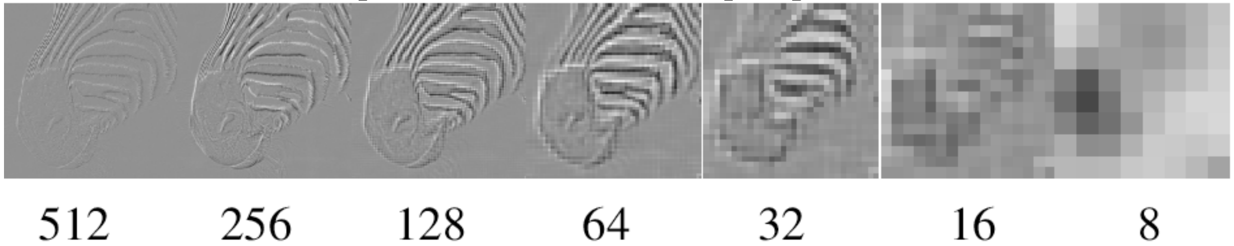

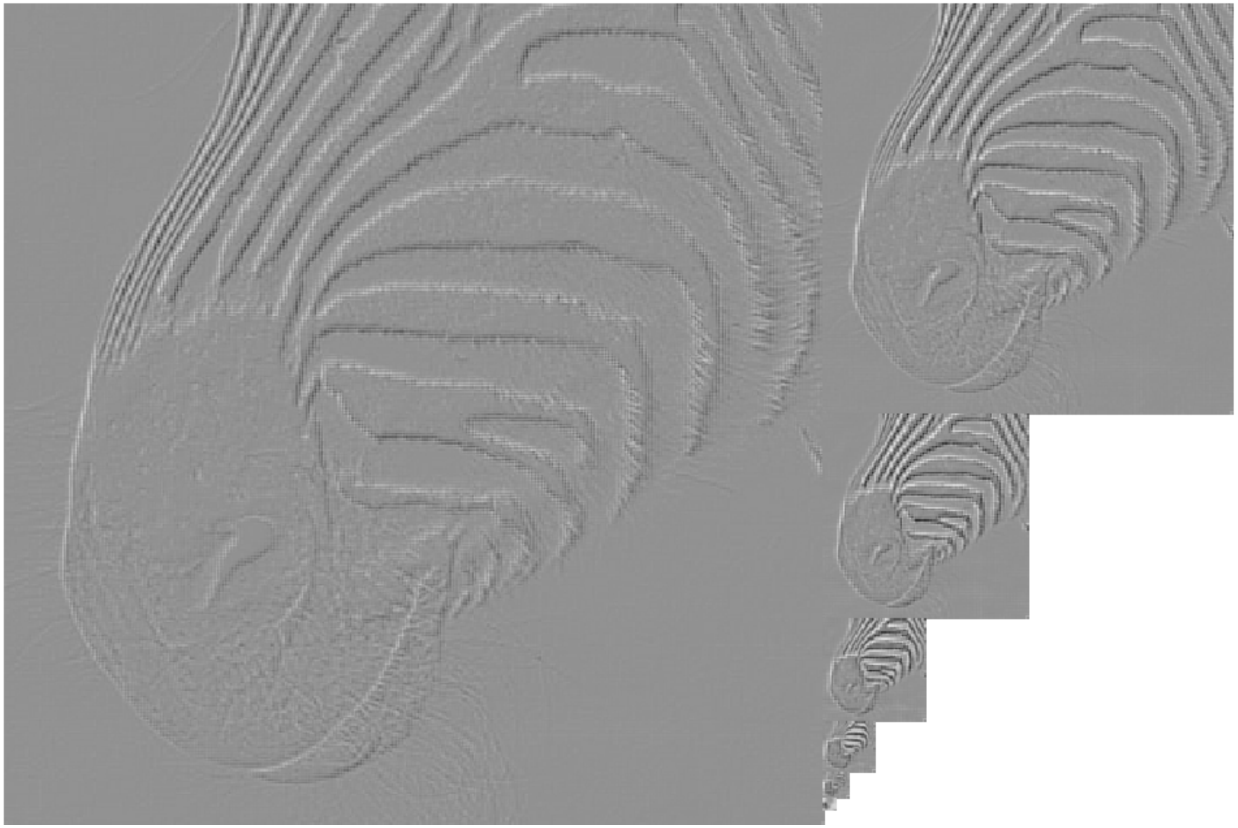

Source: Forsyth

How to build Pyramid in Python

A Gaussian pyramid is basically a series of increasingly decimated images, traditionally at down sampling rate r=2. At each level, the image is first blurred by convolving with a Gaussian-like filter to prevent aliasingin the downsampled image.

```
def decimate(image):
"""
Decimates at image with downsampling rate r=2.
"""
# Blur
image_blur = ndimage.filters.convolve(image,kernel, mode='constant')

# Downsample
return image_blur[::2, ::2]
```

How to build Pyramid in Python

Need a filter:

Binomial filter:

[1 4 6 4 1]

5-tap filter :

$$\frac{1}{256}\begin{pmatrix} 1 & 4 & 6 & 4 & 1 \\ 4 & 16 & 24 & 16 & 4 \\ 6 & 24 & 36 & 24 & 6 \\ 4 & 16 & 24 & 16 & 4 \\ 1 & 4 & 6 & 4 & 1 \end{pmatrix}$$

How to build Pyramid in Python

To build the Laplacian pyramid, we take each level of the Gaussian pyramid and subtract from it the next level interpolated to the same size.

```
def interpolate(image):
"""
Interpolates an image with upsampling rate r=2.
"""
image_up = np.zeros((2*image.shape[0], 2*image.shape[1]))

# Upsample
image_up[::2, ::2] = image

# Blur (we need to scale this up since the kernel has unit area)
# (The length and width are both doubled, so the area is quadrupled)

return ndimage.filters.convolve(image_up,4*kernel, mode='constant')
```

How to build Pyramid in Python

```
G = [image, ]
L = []

# Build the Gaussian pyramid to maximum depth
while image.shape[0] >= 2 and image.shape[1] >= 2:
    image = decimate(image)
    G.append(image)

# Build the Laplacian pyramid
for i in range(len(G) - 1):
    L.append(G[i] - interpolate(G[i + 1]))

return G[:-1], L
```

Hybrid Image

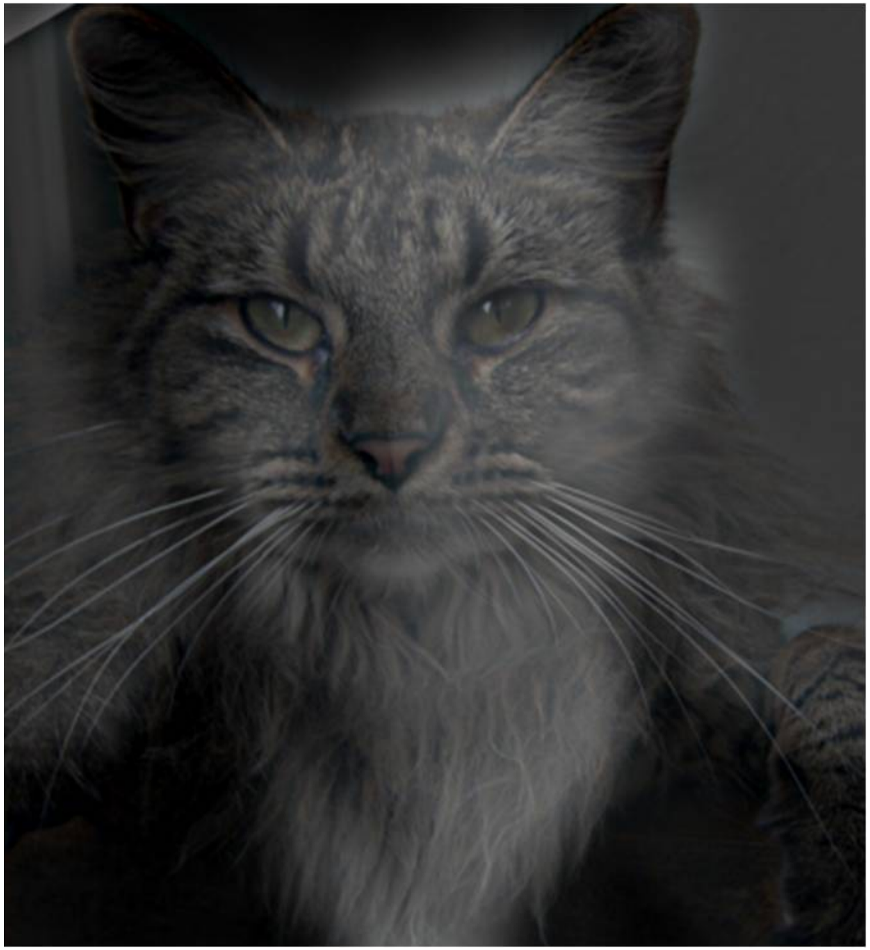

Hybrid Image in Laplacian Pyramid

High frequency → Low frequency

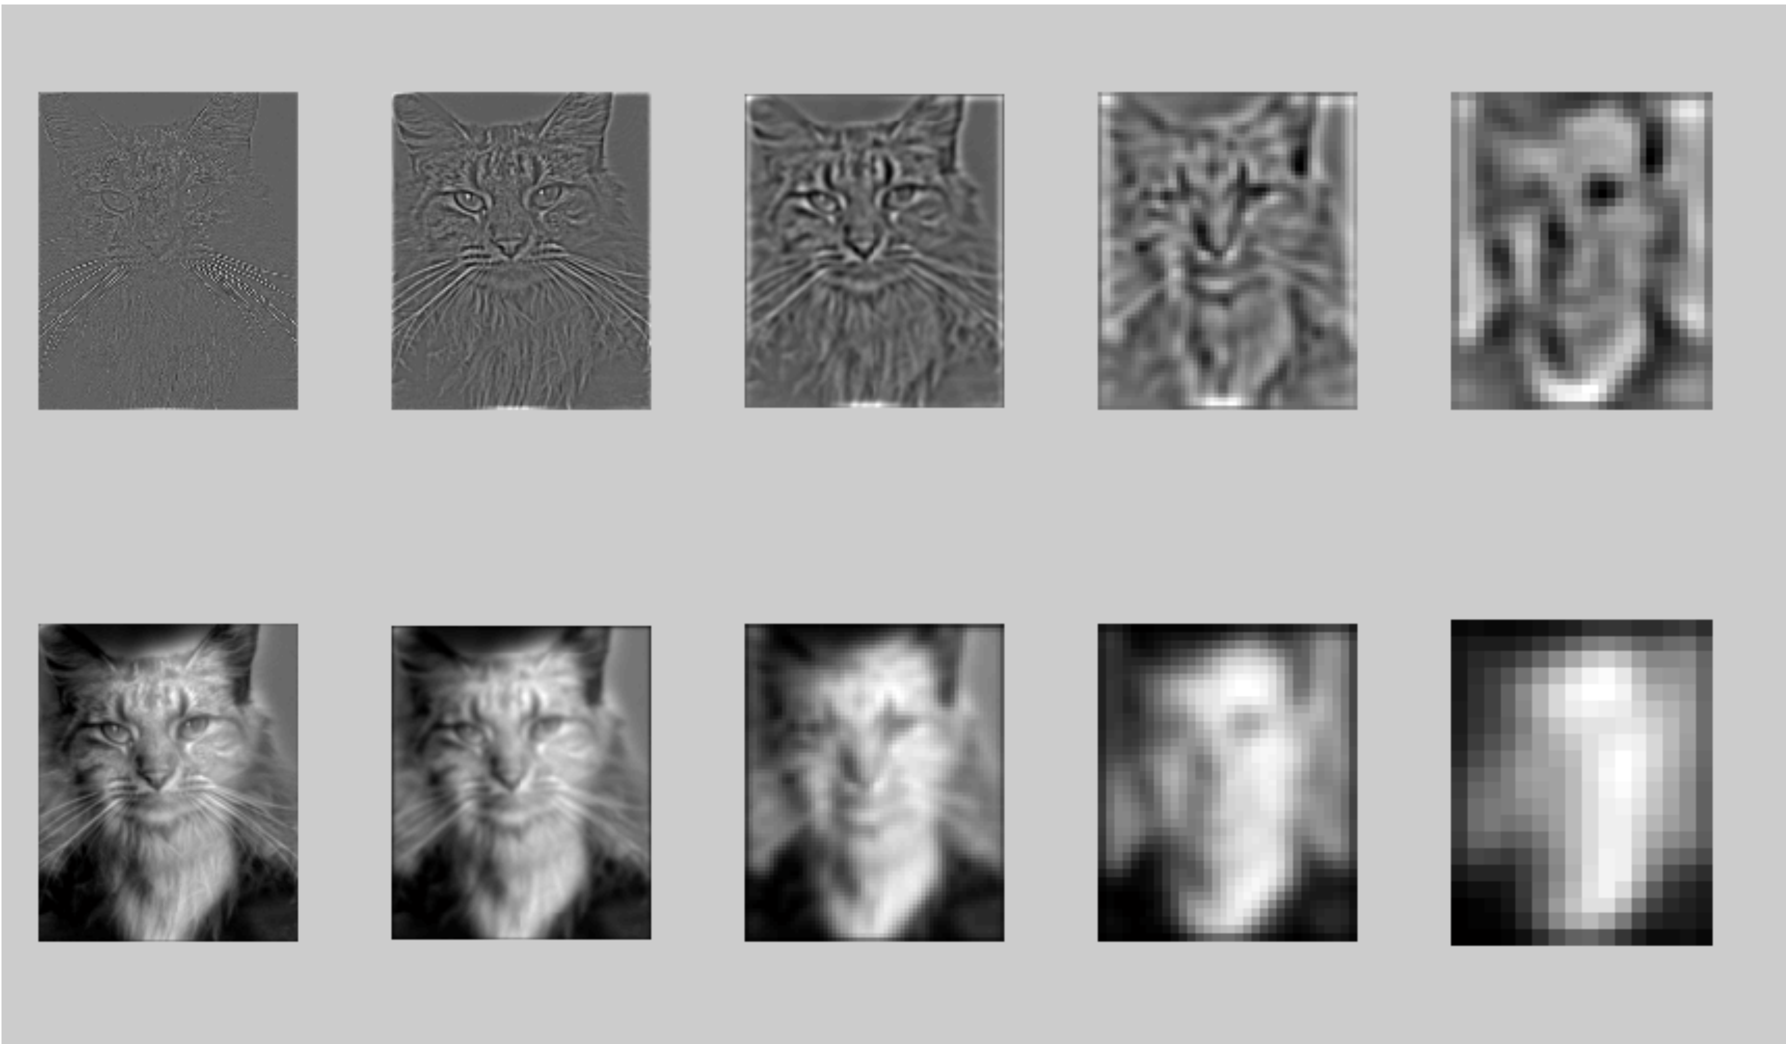

Major uses of image pyramids

- Compression
- Object detection
  - Scale search
  - Features
- Detecting stable interest points
- Registration
  - Course-to-fine

Take-home reading

- Original paper of Image Pyramids:
- http://persci.mit.edu/pub_pdfs/pyramid83.pdf
- Chapter 3.5 Szeliski
- Create a GitHub account and commit your homework code