


# Pyramid Blending, Templates, NL Filters



CS180: Intro to Comp. Vision and Comp. Photo  
Alexei Efros, UC Berkeley, Fall 2024

# Low Pass vs. High Pass filtering

---

Image




Smoothed



-

Details



=




# Application: Hybrid Images

A. Oliva, A. Torralba, P.G. Schyns,  
["Hybrid Images,"](#) SIGGRAPH 2006

Gaussian Filter




Laplacian Filter




unit impulse

-



Gaussian


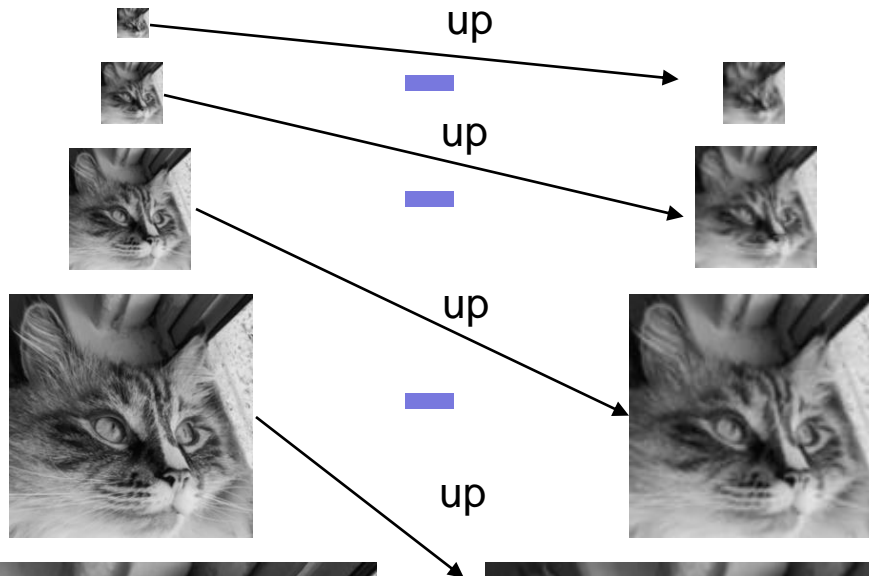
≈





Laplacian of Gaussian

# Band-pass filtering in spatial domain

Gaussian Pyramid  
(low-pass images) :



Laplacian Pyramid  
(sub-band images)



=

=



=

=

# As a stack

---

## Gaussian Pyramid (low-pass images)

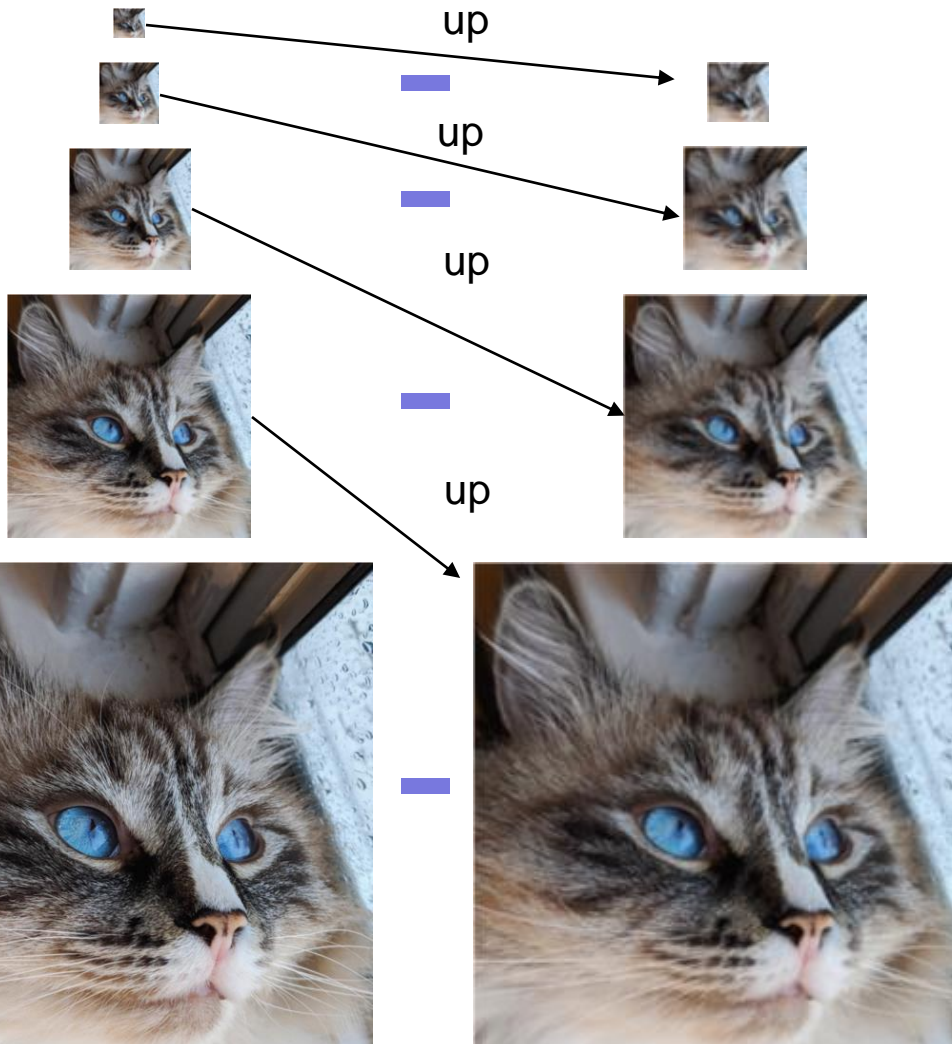


## Laplacian Pyramid (sub-band images)


Created from Gaussian pyramid by subtraction

# Band-pass filtering in spatial domain

Gaussian Pyramid  
(low-pass images) :

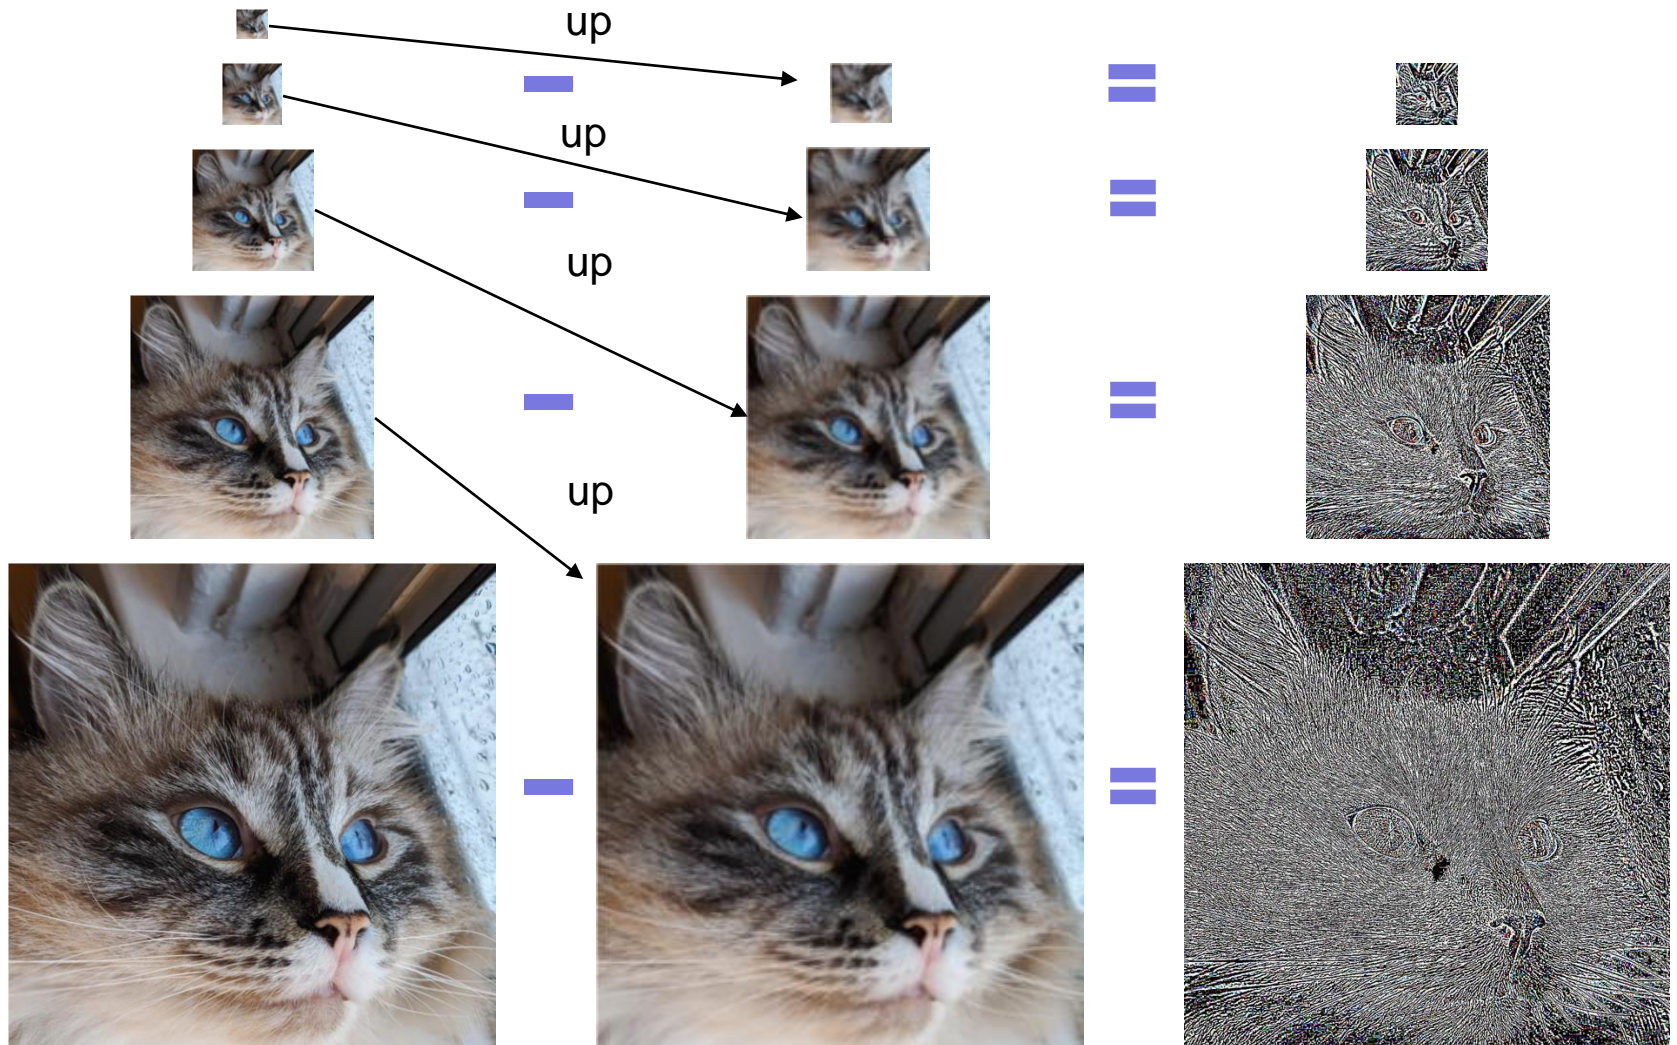


Laplacian Pyramid  
(sub-band images)



# Band-pass filtering in spatial domain



## Gaussian Pyramid (low-pass images)



# Band-pass filtering in spatial domain

---

## Gaussian Pyramid as a stack




## Laplacian Pyramid (sub-band images)

Created from Gaussian pyramid by subtraction



# Collapsing Laplacian Pyramid

---

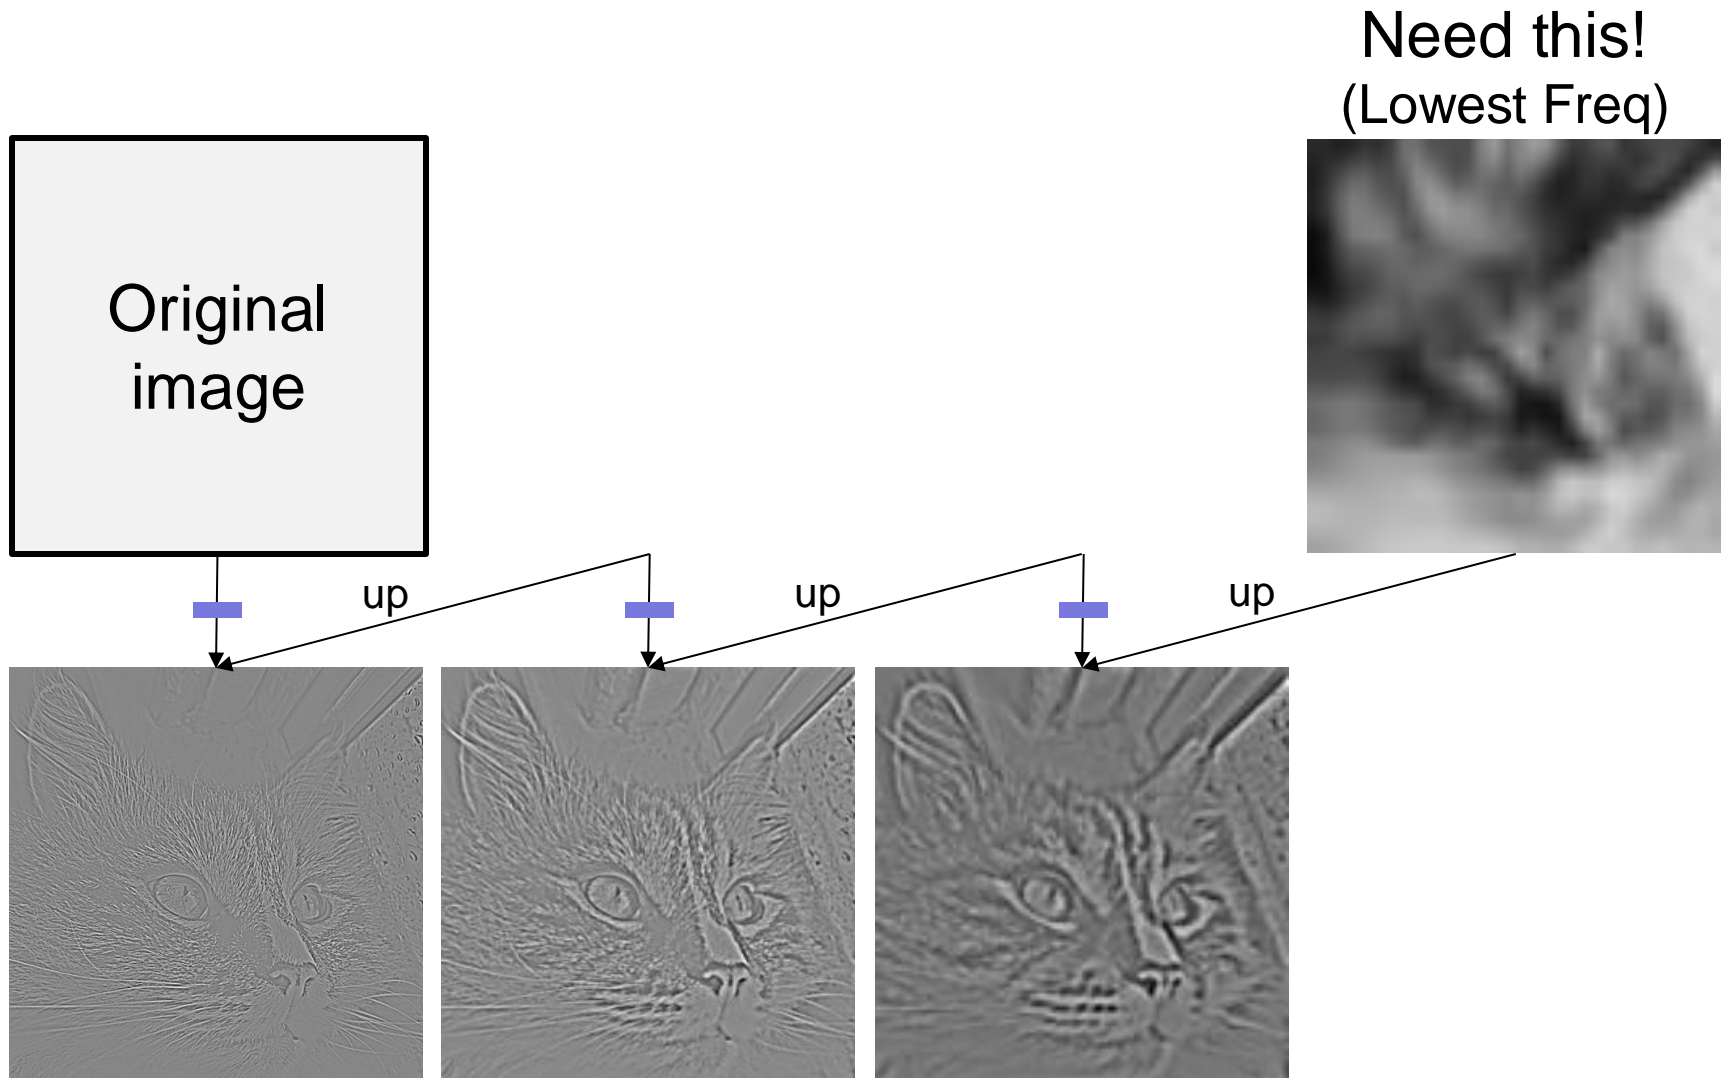


Laplacian Pyramid (sub-band images)

Created from Gaussian pyramid by subtraction

# Collapsing Laplacian Pyramid

---



How can we reconstruct (collapse) this pyramid into the original image?


# Da Vinci and The Laplacian Pyramid

---




# Da Vinci and The Laplacian Pyramid

---



coarse components  
(peripheral vision)



medium components  
(near peripheral vision)






fine details  
(central vision)

Leonardo playing with peripheral vision



# Blending

---



# Alpha Blending / Feathering





---



$$I_{\text{blend}} = \alpha I_{\text{left}} + (1-\alpha) I_{\text{right}}$$





# Affect of Window Size

---



# Affect of Window Size


---





# Good Window Size

---



“Optimal” Window: smooth but not ghosted

# What is the Optimal Window?


---

To avoid seams

- window = size of largest prominent feature

To avoid ghosting

- window  $\leq 2 \times$  size of smallest prominent feature




Natural to cast this in the *Fourier domain*

- largest frequency  $\leq 2 \times$  size of smallest frequency
- image frequency content should occupy one “octave” (power of two)




**FFT**  
→



# What if the Frequency Spread is Wide

---




## Use a band-pass (Laplacian) Pyramid!


- Split image into set of band-pass images (one octave of frequencies each)
- Blend each level of the pyramid separately
- Collapse the pyramid!

# Band-pass Pyramid Blending


---



Left pyramid






blend



Right pyramid


# Pyramid Blending (Burt and Adelson)

---




Burt and Adelson (1983), A Multiresolution Spline With Application to Image Mosaics


laplacian  
level  
4



(c)




(g)




(k)


laplacian  
level  
2



(b)




(f)




(j)


laplacian  
level  
0



(a)



(e)



(i)

left pyramid

right pyramid

blended pyramid

# Image Blending with mask


---



$I^A$



$I^B$



$m$






















$I$

$$I = m * I^A + (1 - m) * I^B$$

# Image Blending with mask

---



$$l_k = l_k^A * m_k + l_k^B * (1 - m_k)$$



# Result

---



# Blending Regions

---



# Image Blending with the Laplacian Pyramid

---

Build Laplacian pyramid for both images: LA, LB

Build Gaussian pyramid for mask: G

Build a combined Laplacian pyramid L

Collapse L to obtain the blended image

532

IEEE TRANSACTIONS ON COMMUNICATIONS, VOL. COM-31, NO. 4, APRIL 1983


## The Laplacian Pyramid as a Compact Image Code

PETER J. BURT, MEMBER, IEEE, AND EDWARD H. ADELSON



# Horror Photo

---



© david dmartin (Boston College)

# Results from this class (fall 2005)

---



© Chris Cameron

# Simplification: Two-band Blending

---

Brown & Lowe, 2003

- Only use two bands -- high freq. and low freq. – without downsampling
- Blends low freq. smoothly
- Blend high freq. with no smoothing: use binary alpha



# 2-band “Laplacian Stack” Blending

---



Low frequency ( $\lambda > 2$  pixels)



High frequency ( $\lambda < 2$  pixels)

# Linear Blending





# 2-band Blending




# Review: Smoothing vs. derivative filters

---



## Smoothing filters

- Gaussian: remove “high-frequency” components; “low-pass” filter
- Can the values of a smoothing filter be negative?
- What should the values sum to?
  - **One**: constant regions are not affected by the filter



## Derivative filters

- Derivatives of Gaussian
- Can the values of a derivative filter be negative?
- What should the values sum to?
  - **Zero**: no response in constant regions
- High absolute value at points of high contrast




# Template matching

---

Goal: find  in image

Main challenge: What is a good similarity or distance measure between two patches?

- Correlation
- Zero-mean correlation
- Sum Square Difference
- Normalized Cross Correlation



# Matching with filters


---

Goal: find  in image


Method 0: filter the image with eye patch

$$h[m,n] = \sum_{k,l} g[k,l] f[m+k,n+l]$$

f = image  
g = filter



Input



Filtered Image

What went wrong?

# Matching with filters

---


Goal: find  in image

f = image  
g = filter


Method 1: filter the image with zero-mean eye

$$h[m, n] = \sum_{k, l} (g[k, l] - \bar{g})(f[m + k, n + l])$$


← mean of g



Input



Filtered Image (scaled)



Thresholded Image


# Matching with filters

---

Goal: find  in image

Method 2: SSD (L2)


$$h[m,n] = \sum_{k,l} (g[k,l] - f[m+k,n+l])^2$$



Input



1- sqrt(SSD)



Thresholded Image

# Matching with filters

---

Can SSD be implemented with linear filters?

$$h[m,n] = \sum_{k,l} (g[k,l] - f[m+k,n+l])^2$$

# Matching with filters


---

Goal: find  in image

What's the potential  
downside of SSD?

Method 2: SSD

$$h[m,n] = \sum_{k,l} (g[k,l] - f[m+k,n+l])^2$$



Input



1- sqrt(SSD)



# Matching with filters

---

Goal: find  in image

Method 3: Normalized cross-correlation

$$h[m,n] = \frac{\sum_{k,l} (g[k,l] - \bar{g})(f[m+k,n+l] - \bar{f}_{m,n})}{\left( \sum_{k,l} (g[k,l] - \bar{g})^2 \sum_{k,l} (f[m+k,n+l] - \bar{f}_{m,n})^2 \right)^{0.5}}$$

mean template                      mean image patch


↓    ↓

# Matching with filters

---

Goal: find  in image


Method 3: Normalized cross-correlation



Input



Normalized X-Correlation




Thresholded Image

# Matching with filters

---

Goal: find  in image


Method 3: Normalized cross-correlation



Input



Normalized X-Correlation



Thresholded Image

# Q: What is the best method to use?

A: Depends


Zero-mean filter: fastest but not a great matcher

SSD: next fastest, sensitive to overall intensity


Normalized cross-correlation: slowest, invariant to local average intensity and contrast

# Denoising

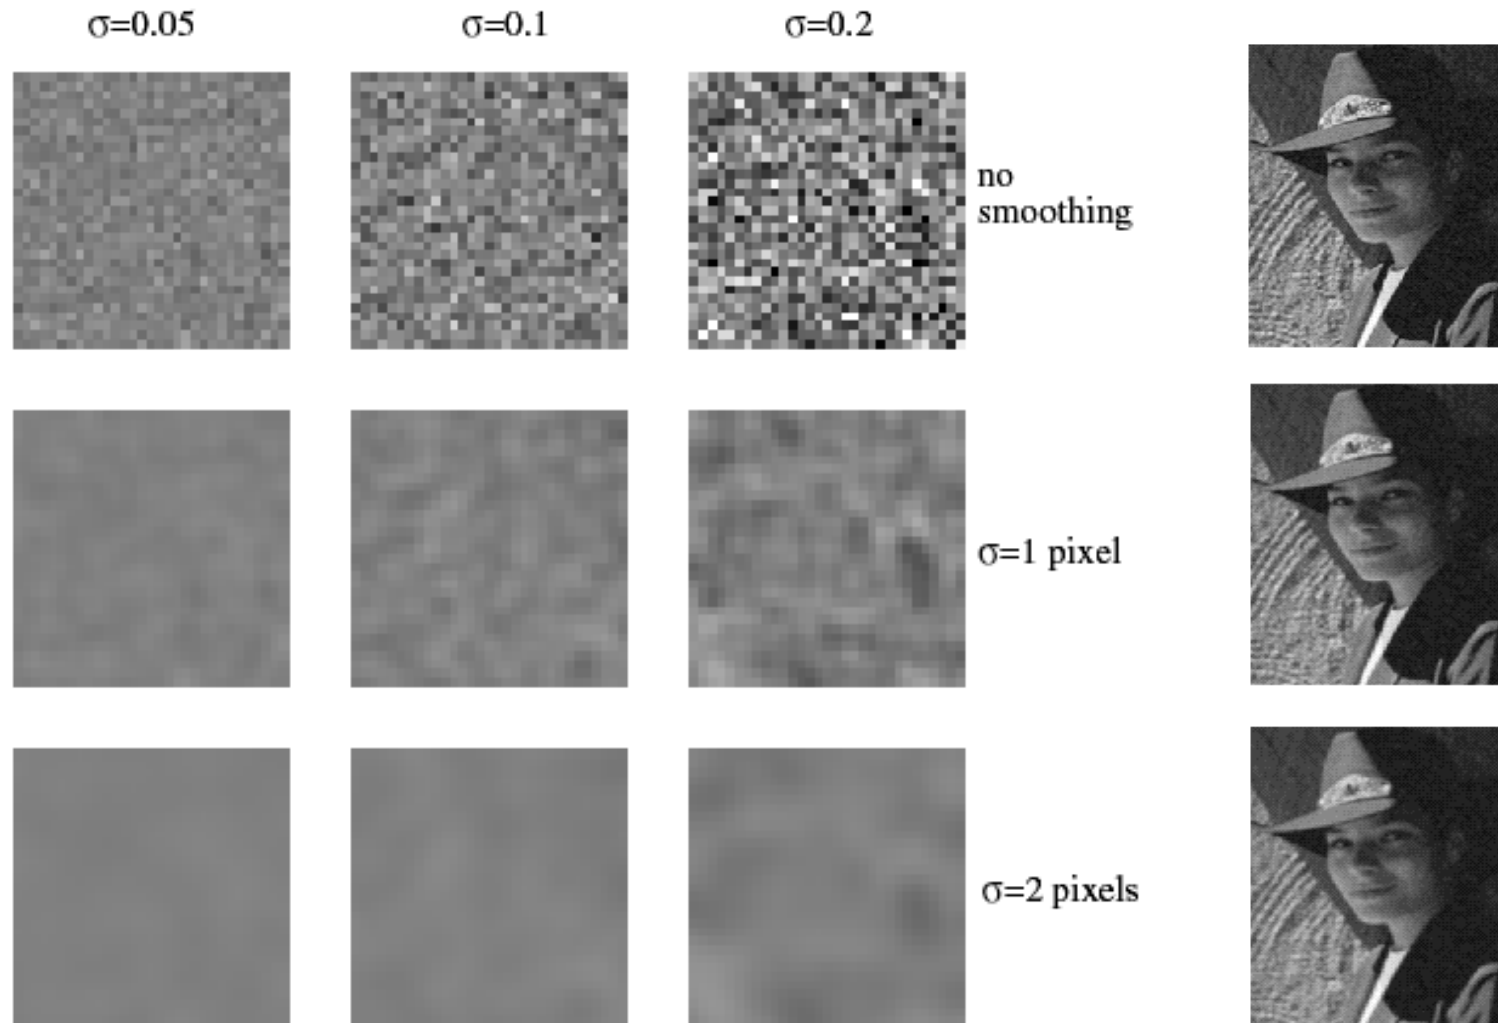
---



Additive Gaussian Noise



# Reducing Gaussian noise



Smoothing with larger standard deviations suppresses noise, but also blurs the image

# Reducing salt-and-pepper noise by Gaussian smoothing

---

3x3



5x5




7x7



# Alternative idea: Median filtering

---

A **median filter** operates over a window by selecting the median intensity in the window



- Is median filtering linear?




# Median filter

---

What advantage does median filtering have over Gaussian filtering?




- Robustness to outliers

filters have width 5 :



# Median filter

Salt-and-pepper noise    Median filtered



`medfilt2(image, [h w])`

# Median vs. Gaussian filtering

---

3x3

5x5

7x7

Gaussian



Median



# Side note: Image Compression

---




89k

# Lossless Compression (e.g. Huffman coding)

---


Input image:



Pixel code:

color	freq.	bit code
	14	0
	6	10
	3	110
	2	111

Pixel histogram:



Compressed image:


0 110 110 0 0

0 10 110 111 0

...


# Lossless Compression not enough

---




# Lossy Image Compression (JPEG)

---



cut up into 8x8 blocks





Block-based Discrete Cosine Transform (DCT)

# Using DCT in JPEG

---

The first coefficient  $B(0,0)$  is the DC component, the average intensity

The top-left coeffs represent low frequencies, the bottom right – high frequencies





# Image compression using DCT

## Quantize

- More coarsely for high frequencies (tend to have smaller values anyway)
- Many quantized high frequency values will be zero

## Encode

- Can decode with inverse dct

## Filter responses

$$G = \begin{matrix} & & & \xrightarrow{u} & & & & & \\ \begin{matrix} \downarrow v \\ \end{matrix} & \begin{bmatrix} -415.38 & -30.19 & -61.20 & 27.24 & 56.13 & -20.10 & -2.39 & 0.46 \\ 4.47 & -21.86 & -60.76 & 10.25 & 13.15 & -7.09 & -8.54 & 4.88 \\ -46.83 & 7.37 & 77.13 & -24.56 & -28.91 & 9.93 & 5.42 & -5.65 \\ -48.53 & 12.07 & 34.10 & -14.76 & -10.24 & 6.30 & 1.83 & 1.95 \\ 12.12 & -6.55 & -13.20 & -3.95 & -1.88 & 1.75 & -2.79 & 3.14 \\ -7.73 & 2.91 & 2.38 & -5.94 & -2.38 & 0.94 & 4.30 & 1.85 \\ -1.03 & 0.18 & 0.42 & -2.42 & -0.88 & -3.02 & 4.12 & -0.66 \\ -0.17 & 0.14 & -1.07 & -4.19 & -1.17 & -0.10 & 0.50 & 1.68 \end{bmatrix} \end{matrix}$$



## Quantized values

$$B = \begin{bmatrix} -26 & -3 & -6 & 2 & 2 & -1 & 0 & 0 \\ 0 & -2 & -4 & 1 & 1 & 0 & 0 & 0 \\ -3 & 1 & 5 & -1 & -1 & 0 & 0 & 0 \\ -3 & 1 & 2 & -1 & 0 & 0 & 0 & 0 \\ 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \end{bmatrix}$$

## Quantization table

$$Q = \begin{bmatrix} 16 & 11 & 10 & 16 & 24 & 40 & 51 & 61 \\ 12 & 12 & 14 & 19 & 26 & 58 & 60 & 55 \\ 14 & 13 & 16 & 24 & 40 & 57 & 69 & 56 \\ 14 & 17 & 22 & 29 & 51 & 87 & 80 & 62 \\ 18 & 22 & 37 & 56 & 68 & 109 & 103 & 77 \\ 24 & 35 & 55 & 64 & 81 & 104 & 113 & 92 \\ 49 & 64 & 78 & 87 & 103 & 121 & 120 & 101 \\ 72 & 92 & 95 & 98 & 112 & 100 & 103 & 99 \end{bmatrix}$$

# JPEG Compression Summary

---

Subsample color by factor of 2

- People have bad resolution for color

Split into blocks (8x8, typically), subtract 128

For each block

- a. Compute DCT coefficients
- b. Coarsely quantize
  - Many high frequency components will become zero
- c. Encode (e.g., with Huffman coding)

Spatial dimension of color channels are reduced by 2 (lecture 2)!

<http://en.wikipedia.org/wiki/YCbCr>

<http://en.wikipedia.org/wiki/JPEG>

# JPEG compression comparison

---



89k



12k