Image Pyramids and Applications



Computer Vision Jia-Bin Huang, Virginia Tech

Administrative stuffs

- HW 1 posted, due 11:59 PM Sept 19
- Anonymous feedback from students
 - Repeat students' questions and answers
 - Turn on some light in the classroom
 - Post frequently asked questions for HWs

Previous class: Image Filtering

- Sometimes it makes sense to think of images and filtering in the frequency domain
 - Fourier analysis
- Can be faster to filter using FFT for large images (N logN vs. N² for auto-correlation)
- Images are mostly smooth
 - Basis for compression
- Remember to low-pass before sampling







Spatial domain



Frequency domain

Fourier Transform

Teases away fast vs. slow changes in the image.



Image as a sum of basis images

Slide credit: A Efros

Extension to 2D Fourier domain with complex amplitude: a+jb Discrete Fourier Transform 13

in Matlab, check out: imagesc(log(abs(fftshift(fft2(im)))));

Phase vs. Magnitude



Today's class

- Template matching
- Image Pyramids
- Compression
- Introduction to HW1

Template matching

- Goal: find 💽 in image
- Main challenge: What is a good similarity or distance measure between two patches? D(,,)
 - Correlation
 - Zero-mean correlation
 - Sum Square Difference
 - Normalized Cross Correlation



- Goal: find 💽 in image
- Method 0: filter the image with eye patch

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

$$f = image$$

$$g = filter$$
What went wrong?

Filtered Image

Input

- Goal: find 💽 in image
- Method 1: filter the image with zero-mean eye

$$h[m,n] = \sum_{k,l} (g[k,l] - \overline{g}) \underbrace{(f[m+k,n+l])}_{\text{mean of template g}}$$



Input



Filtered Image (scaled)



Thresholded Image

- Goal: find 💽 in image
- Method 2: Sum of squared differences (SSD)

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







Input

1- sqrt(SSD)

Thresholded Image

Can SSD be implemented with linear filters? $h[m,n] = \sum_{k,l} (g[k,l] - f[m+k,n+l])^2$

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

Constant Filtering with g Filtering with box filter

- Goal: find 💽 in image
- Method 2: SSD

What's the potential downside of SSD?

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



Input

1- sqrt(SSD)

- Goal: find 💽 in image
- Method 3: Normalized cross-correlation

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

Matlab: normxcorr2 (template, im)

- Goal: find 💽 in image
- Method 3: Normalized cross-correlation



Input

Normalized X-Correlation

Thresholded Image

- 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

Q: What if we want to find larger or smaller eyes?

A: Image Pyramid

Review of Sampling



Gaussian pyramid



512 256 128 64 32 16 8



Source: Forsyth

Template Matching with Image Pyramids

Input: Image, Template

- 1. Match template at current scale
- 2. Downsample image
 - In practice, scale step of 1.1 to 1.2
- 3. Repeat 1-2 until image is very small
- 4. Take responses above some threshold, perhaps with non-maxima suppression



Source: Lazebnik

Laplacian pyramid





Creating the Gaussian/Laplacian Pyramid



- Use same filter for smoothing in each step (e.g., Gaussian with $\sigma = 2$)
- Downsample/upsample with "nearest" interpolation

Hybrid Image in Laplacian Pyramid

High frequency \rightarrow Low frequency





Reconstructing image from Laplacian pyramid

Image = L₁ + Smooth(Upsample(G₂))



- Use same filter for smoothing as in desconstruction
- Upsample with "nearest" interpolation
- Reconstruction will be lossless

Major uses of image pyramids

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







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?



Applications: Pyramid Blending







Applications: Pyramid Blending







Pyramid Blending

- At low frequencies, blend slowly
- At high frequencies, blend quickly





Right pyramid

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

Compression

How is it that a 4MP image (12000KB) can be compressed to 400KB without a noticeable change?

Lossy Image Compression (JPEG)



Block-based Discrete Cosine Transform (DCT)

Slides: Efros

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 (which also tend to have smaller values)
 - Many quantized high frequency values will be zero
- Encode
 - Can decode with inverse dct

Filter	r respo	nses		$\stackrel{u}{\longrightarrow}$					
G =	$\begin{bmatrix} -415.38 \\ 4.47 \\ -46.83 \\ -48.53 \\ 12.12 \\ -7.73 \\ -1.03 \end{bmatrix}$	$\begin{array}{r} -30.19 \\ -21.86 \\ 7.37 \\ 12.07 \\ -6.55 \\ 2.91 \\ 0.18 \end{array}$	$\begin{array}{r} -61.20 \\ -60.76 \\ 77.13 \\ 34.10 \\ -13.20 \\ 2.38 \\ 0.42 \end{array}$	$\begin{array}{r} 27.24 \\ 10.25 \\ -24.56 \\ -14.76 \\ -3.95 \\ -5.94 \\ -2.42 \end{array}$	$56.13 \\ 13.15 \\ -28.91 \\ -10.24 \\ -1.88 \\ -2.38 \\ -0.88$	$\begin{array}{r} -20.10 \\ -7.09 \\ 9.93 \\ 6.30 \\ 1.75 \\ 0.94 \\ -3.02 \end{array}$	$\begin{array}{r} -2.39 \\ -8.54 \\ 5.42 \\ 1.83 \\ -2.79 \\ 4.30 \\ 4.12 \end{array}$	0.46 4.88 -5.65 1.95 3.14 1.85 -0.66	$\bigg \downarrow^v$
Quantized values									
	В	$= \begin{bmatrix} -26\\ -3\\ -3\\ -3\\ 0\\ 0\\ 0\\ 0\\ 0\\ 0\\ 0\\ 0\\ 0\\ 0\\ 0\\ 0\\ 0\\$	$ \begin{array}{cccccccccccccccccccccccccccccccccccc$	$\begin{array}{ccccc} -6 & 2 \\ -4 & 1 \\ 5 & -1 \\ 2 & -1 \\ 0 & 0 \\ 0 & 0 \\ 0 & 0 \\ 0 & 0 \end{array}$	$\begin{array}{cccc} 2 & -1 \\ 1 & 0 \\ -1 & 0 \\ 0 & 0 \\ 0 & 0 \\ 0 & 0 \\ 0 & 0 \\ 0 & 0 \end{array}$	0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0			

Quantization table

Q =	16 12 14 14 18	11 12 13 17 22	10 14 16 22 37	16 19 24 29 56	24 26 40 51 68	40 58 57 87 109	51 60 69 80 103	61 55 56 62 77
	14	17	22	24 20	40 51	87	80	62
Q =	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

JPEG Compression Summary

- 1. Convert image to YCrCb
- 2. Subsample color by factor of 2
 - People have bad resolution for color
- 3. Split into blocks (8x8, typically), subtract 128
- 4. For each block
 - a. Compute DCT coefficients
 - b. Coarsely quantize
 - Many high frequency components will become zero
 - c. Encode (e.g., with Huffman coding)

http://en.wikipedia.org/wiki/YCbCr http://en.wikipedia.org/wiki/JPEG

Lossless compression (PNG)

- 1. Predict that a pixel's value based on its upper-left neighborhood
- 2. Store difference of predicted and actual value
- 3. Pkzip it (DEFLATE algorithm)

С	В	D	
Α	Х		

Three views of image filtering

- Image filters in spatial domain
 - Filter is a mathematical operation on values of each patch
 - Smoothing, sharpening, measuring texture
- Image filters in the frequency domain
 - Filtering is a way to modify the frequencies of images
 - Denoising, sampling, image compression
- Templates and Image Pyramids
 - Filtering is a way to match a template to the image
 - Detection, coarse-to-fine registration

HW 1 – Hybrid Image



• Hybrid image =

Low-Freq(Image A) + Hi-Freq(Image B)

HW 1 – Image Pyramid



HW 1 – Edge Detection



Derivative of Gaussian filters



Things to remember

- Template matching (SSD or Normxcorr2)
 - SSD can be done with linear filters, is sensitive to overall intensity
- Gaussian pyramid
 - Coarse-to-fine search, multi-scale detection
- Laplacian pyramid
 - More compact image representation
 - Can be used for compositing in graphics
- Compression
 - In JPEG, coarsely quantize high frequencies









Thank you

- See you this Thursday
- Next class:
 - Edge detection

