Frequencies

# Agenda

- Texture
- Pyramid matching
- Frequencies

## Cost function

 $\min_{Z,D} C(Z, D, X) \quad \text{where} \quad C(Z, D, X) = \sum_{i} ||x_i - d_{z_i}||^2$ 

 $x_i$ : i<sup>th</sup> input vector (to be clustered)  $z_i \in \{1 \dots K\}$  : i<sup>th</sup> label

d<sub>k</sub>: k<sup>th</sup> dictionairy element (or mean)

## Coordinate descent optimization

 $\min_{Z,D} C(Z, D, X) \quad \text{where} \quad C(Z, D, X) = \sum_{i} ||x_i - d_{z_i}||^2$ 

1. 
$$\min_{Z} C(Z, D, X)$$
  
2.  $\min_{D} C(Z, D, X)$ 

Aside: what happens if we are just given a *matrix* of pairwise distances?

## Sparse reconstructions

 $\min_{D,Z} ||X - DZ||_F^2 \quad \text{subject to sparse constraints on Z}$ 

$$X = [x_1, \dots x_n]$$
$$D = [d_1, \dots, d_K]$$
$$Z = [z_1, \dots z_n]$$

K-means:  $z_i = [..., 0, 1, 0, ...]$ LO sparse-coding:  $||z_i||_0 \leq M$  (greedy algorithms known as "matching pursuit") L1 sparse-coding:  $||z_i||_1 \leq M$ 

(convex program)

# Dictionary learning

 $\min_{D,Z} ||X - DZ||_F^2$ 

Some folks claim that k-means should always be replaced by sparse coding (never hurts, sometimes better) even for bag-of-words

... k-means is far simpler, right?

#### "In between" k-means and sparse coding

 $\min_{D,Z} ||X - DZ||_F^2 \quad \text{subject to sparse constraints on Z}$ 

K-means:  $z_i = [\dots, 0, 1, 0, \dots]$ L0 sparse-coding:  $||z_i||_0 \le M$  (For M = 1, we can solve for Z in closed form)



Similar to k-means using cosine distance, but allows for negative matches

Histograms of Sparse Codes for Object Recognition, CVPR13

# Agenda

- Texture
- Pyramid matching
- Frequencies

## Recall texture pipeline



## Recall texture pipeline



mesh

#### Alternative to quantization: explicitly compute matchings between image pairs

Grauman & Darrell



#### Digression: alternative to quantization

Approximate matching with histogram similarity



## Aside: what's the "right" way to compare histograms? $D_r(x,y) = \left(\sum_i (x_i - y_i)^r\right)^{\frac{1}{r}}$







#### Aside: what's the "right" way to compare histograms?

euclidean (r=2)  
or manhattan (r=1)
$$D_r(x,y) = \left(\sum_i (x_i - y_i)^r\right)^{\frac{1}{r}}$$
chi-squared distance  
[derived from chi-squared text in statistics]
$$Chi(x,y) = \sum_i \frac{(x_i - y_i)^2}{x_i + y_i}$$
[log probability of seeing x under model y]
$$D_{KL}(x,y) = \sum_i x_i \log \frac{x_i}{y_i}$$

eu or m

chi-squ [derived from chi-

K-L

## Earth mover's distance

Cast as "transportation problem"



## Earth mover's distance

Bipartite network flow



 $\min_{f_{ij}} \sum_{ij} c_{ij} f_{ij} \quad s.t.$  $f_{ij} \ge 0$  $\sum_{i} f_{ij} = y_j$  $\sum_{j} f_{ij} = x_i$ 

$$EMD(\mathbf{x}, \mathbf{y}) = \sum_{ij} c_{ij} f_{ij}$$

https://www.cs.duke.edu/~tomasi/papers/rubner/rubnerTr98.pdf

# Similarity Kernals

[sometimes more intuitive to define than distance functions]

$$D(x,y) = K(x,x) + K(y,y) - 2K(x,y)$$
  
 $K_{lin}(x,y) = \sum_{i} x_{i}y_{i}$  [what's corresponding distance function?]

# Similarity Kernals

[sometimes more intuitive to define than distance functions]

$$D(x, y) = K(x, x) + K(y, y) - 2K(x, y)$$
$$K_{lin}(x, y) = \sum_{i} x_{i} y_{i}$$
$$K_{int}(x, y) = \sum_{i} \min(x_{i}, y_{i})$$

What happens if x,y are binary vectors?

# Similarity Kernals

[sometimes more intuitive to define than distance functions]

$$D(x, y) = K(x, x) + K(y, y) - 2K(x, y)$$
$$K_{lin}(x, y) = \sum_{i} x_{i}y_{i}$$
$$K_{int}(x, y) = \sum_{i} \min(x_{i}, y_{i})$$
$$K_{bat}(x, y) = \sum_{i} \sqrt{x_{i}y_{i}}$$

It turns out, we can compute transformations f(x) and f(y) such that L2 distance in transformed space corresponds to these kernals (allows use of linear predictors)

http://www.robots.ox.ac.uk/~vgg/software/homkermap/

### Histogram intersection kernal

$$\mathcal{I}(H(\mathbf{X}), X(\mathbf{Y})) = \sum_{k} \min(H(\mathbf{X}_{k}), X(\mathbf{Y}_{k}))$$

= 4



#### Back to correspondence matching



But what about bin effects (partial credit for near matches)?

#### Back to correspondence matching



Count matches obtained from larger bins

### Counting new matches

-listogram 
$$\mathcal{I}\left(H(\mathbf{X}),H(\mathbf{Y})\right) = \sum_{j=1}^{r}\min\left(H(\mathbf{X})_{j},H(\mathbf{Y})_{j}\right)$$



#### Giving partial credit for new matches



Weight new matches inversely porportional to bin size

 $\frac{1}{2^i}$ 

## Pyramid match kernel



- Weights inversely proportional to bin size
- Normalize kernel values to avoid favoring large sets

## Spatial Pyramid Matching

Quantize features into words, but build pyramid in space Nifty way to encode constraints like "eye" words lie near top of image



# Agenda

- Texture
- Pyramid matching
- Frequencies

## Motivation

Representing a signal as a linear combination of impulse functions



..allows us to characterize linear-shift invariant (LSI) systems by their impulse response

What about other basis functions besides impulses?

## Lets pick a smooth basis

to reflect fact that real-world signals tend to be slow-changing





sine functions (sinusoids)

#### Where we are headed...



#### Why sinusoidal basis functions?

$$(1*) = .1* (1*) + .3* (1*) + ...$$

#### 1. Efficient representation

Punchline: better compression (only need to encode low-frequency sinusoids - basis of jpeg)

-

#### 2. Sinusoids are *eigenfunctions* of LSI systems

$$\bigwedge \longrightarrow \mathbb{I} \longrightarrow \mathbb{I}^*$$

Punchline: LSI systems (or convolutions) can be implemented by ...

#### Implement LSI operations by multiplication



"Replace convolutions with multiplications"

## Background: sinusoids



$$f(t) = A\sin(2\pi ft + \phi) = A\sin(\omega x + \phi)$$

A: amplitude
φ: phase
f: frequency (cycles in t=1 sec)
ω:
λ:

# Background: complex numbers



(often "j" instead of "i")

Alternative parameterizations:  $z = x + iy = re^{i\theta}$   $e^{i\theta} = \cos \theta + i \sin \theta$  (Euler's formula) Re(z) = Im(z) = $z^* =$ 

# Background: complex numbers



What's the product of two complex numbers?

$$r_1 e^{j\theta_1} r_2 e^{j\theta_2} = r_1 r_2 e^{j(\theta_1 + \theta_2)}$$

What's the inverse of a complex number?

$$Inv(z) = \frac{1}{r}e^{-j\theta} = z^* \quad \text{for} \quad r = 1$$



Think of point travelling around unit circle on imaginary plane at  $\frac{\omega}{2\pi}$  cycles per sec
Proof: complex sinusoids are eigenfunctions of LSI systems

$$g[u] = \sum_{v} h[v]f[u-v], \quad f[u] = e^{j\omega u}$$

Proof: complex sinusoids are eigenfunctions of LSI systems

 $g[u] = \sum_{v} h[v]f[u - v], \quad f[u] = e^{j\omega u}$  $= e^{j\omega u} \sum_{v} h[v]e^{-j\omega v}$  $= e^{j\omega u}H(w) \quad \text{where} \quad H(w) = \sum_{v} h[v]e^{-j\omega v}$  $= f[u]H(\omega)$ 

If we hit a LSI system with a sinuoid, output is that sinusoid scaled by H(w) (a function of system's impulse response)

Proof: complex sinusoids are eigenfunctions of LSI systems

$$\begin{split} g[u] &= \sum_{v} h[v] f[u-v], \quad f[u] = e^{j\omega u} \\ &= e^{j\omega u} \sum_{v} h[v] e^{-j\omega v} \\ &= e^{j\omega u} H(w) \quad \text{where} \quad H(w) = \sum_{v} h[v] e^{-j\omega v} \\ &= f[u] H(\omega) \end{split}$$

Aside: Could this have worked for  $f[u] = z^u$  ?

The resulting H[z] is known as a z-transform: <u>https://en.wikipedia.org/wiki/Z-transform</u>

Evaluate below transformation at N different frequencies



$$F[k] = \frac{1}{N} \sum_{u=0}^{N-1} f[u] e^{\frac{-j2\pi ku}{N}}$$

$$Re(F[k]) = \frac{1}{N} \sum_{u=0}^{N-1} f[u] \cos(\frac{-2\pi ku}{N})$$



N-long sequence N-long sequence

$$F[k = 0] = \frac{1}{N} \sum_{u=0}^{N} f[u]$$
$$F[k = 1] = \frac{1}{N} \sum_{u=0}^{N} f[u] e^{\frac{-j2\pi u}{N}}$$

$$F[k = N/2] = \frac{1}{N} \sum_{u=0}^{N} f[u] e^{-j\pi u}$$

• • •

Comput F[k] by taking inner product of f[u] with a complex sinusoid  $g_k[u]$ What does the (real part of) this sinuoidal functions look like?

$$F[k = 0] = \frac{1}{N} \sum_{u=0}^{N} f[u]$$
$$F[k = 1] = \frac{1}{N} \sum_{u=0}^{N} f[u] e^{\frac{-j2\pi u}{N}}$$



$$F[k = N/2] = \frac{1}{N} \sum_{u=0}^{N} f[u] e^{-j\pi u}$$

• • •



$$F[k = 0] = \frac{1}{N} \sum_{u=0}^{N} f[u]$$
$$F[k = 1] = \frac{1}{N} \sum_{u=0}^{N} f[u] e^{\frac{-j2\pi u}{N}}$$

$$F[k = N/2] = \frac{1}{N} \sum_{u=0}^{N} f[u] e^{-j\pi u}$$

$$F[k] = \frac{1}{N} \sum_{u=0}^{N-1} f[u] e^{\frac{-j2\pi ku}{N}}$$

F[k=N] = ?

. . .

## Linear transformation



How expensive is the conversion?

## Complex matrix multiplication



## Complex orthonormality

$$U^* = [U_{ij}^*]^T \quad "$$
$$(U^*)U = I$$
$$( \bigcirc ) \bigcirc = I = I$$

"complex conjugate or adjoint"

Implies the inverse transform is obtained by transpose + flipping imaginary part

$$F = Uf$$
$$(U^*)F = f$$

#### Inverse DFT

$$f[u] = \sum_{k=0}^{N-1} F[k] e^{\frac{j2\pi ku}{N}} \quad \text{(IDFT)}$$
$$F[k] = \frac{1}{N} \sum_{u=0}^{N-1} f[u] e^{\frac{-j2\pi ku}{N}} \quad \text{(DFT)}$$



Sometimes the 1/N is moved or factored (really need sqrt(N) to make transform self-inverting)

## General fourier terminology

https://en.wikipedia.org/wiki/Fourier\_transform



#### $a_n \cos(nx) + b_n \sin(nx)$

Decompose any continuous periodic signal into a summation of sinusoids of increasing frequency

## General fourier terminology

• Continuous fourier transfrom

https://en.wikipedia.org/wiki/Fourier\_transform

• Discrete-time fourier transform

https://en.wikipedia.org/wiki/Discrete-time\_Fourier\_transform

(useful for analyzing samples of a continuous fn)

• Discrete fourier transform (DFT)

https://en.wikipedia.org/wiki/Discrete\_Fourier\_transform

## DFTs as a change of basis



How do we interpret complex weighting coefficients?

$$F[k] = a + bj = re^{j\theta}$$

$$re^{j\theta}e^{\frac{j2\pi ku}{N}} = re^{j(\frac{2\pi ku}{N} + \theta)}$$

$$Re\{\cdot\} = r\cos(\frac{2\pi ku}{N} + \theta)$$

u

Scaling coefficients allow us to shift sinusoids!

#### Visualizing DFTs



Visualize frequencies from [-N/2..N/2] instead of [0..N]

(Equivalent info because FTs are periodic)

## Proof: DFT of real signal

$$F[k] = \frac{1}{N} \sum_{u=0}^{N-1} f[u] e^{\frac{-j2\pi ku}{N}}$$

$$F^*[-k] =$$

## Properties of DFT

- Symmetric for real-valued f[u]:  $F[k] = F[-k]^*$  (we only really need N/2 values)
- Periodic: F[k+N] = F[k]. (we only need to encode N values; often use k = [-N/2...N/2])
- Eigenfunction property: the DFT of the convolution of two functions is the product of their DFTs (convolution in time domain is multiplication in frequency domain)
- DFT of a Gaussian with variance sigma is a another Gaussian with variance 1/sigma

# Crucial property: convolution is multiplication in frequency space



Extensions to 2D  

$$F[k,l] = \frac{1}{N^2} \sum_{u=v}^{N} \sum_{v=v}^{N} f[u,v] e^{-2j\pi(\frac{ku+lv}{N})}$$

 $e^{\pi i (ux+vy)}$ 

$$g[u, v] = \cos(-2\pi \frac{ku + lv}{N})$$

$$\begin{bmatrix} k \\ l \end{bmatrix}^T \begin{bmatrix} u \\ v \end{bmatrix} = ||a||||b||\cos\theta$$

magnitude of [k,1] gives the frequency  $e^{-\pi i(ux+vy)}$ 



u

## Visualizing DFTs in 2D

#### Arrange coefficients into an *image*



## Visualizing DFTs in 2D

#### What does basis image look like for [-k,-l]?





#### Extension to 2D



#### Extension to 2D



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

## Fourier analysis in images



http://sharp.bu.edu/~slehar/fourier/fourier.html#filtering

## Signals can be composed



http://sharp.bu.edu/~slehar/fourier/fourier.html#filtering More: http://www.cs.unm.edu/~brayer/vision/fourier.html

# Let's build up a signal by randomly adding in fourier transform coefficients






































#### Example 2D Fourier transform

#### Image with periodic structure





f(x,y)

|F(u,v)|

FT has peaks at spatial frequencies of repeated texture

#### **Example – Forensic application**



# Recall: convolution is multiplication in frequency space



# Recall: convolution is multiplication in frequency space



Implies that previous frequency manipulations can be implemented with convolutions!

#### Fast Fourier transform (FFT)

"The most important numerical algorithm of our lifetime"

Strang, Gilbert (May-June 1994). "Wavelets". American Scientist 82 (3): 253

Compute (I)DFTs in NlogN instead of N<sup>2</sup>(standard matrix multiplication)



Recursively compute N-element DFT with 2 N/2-element DFT

# Filtering

# Why does the Gaussian give a nice smooth image, but the square filter give edgy artifacts?

Gaussian



Box filter





# Recall...



# Another look: blurring



#### original

# Another look: blurring



#### smoothed (5x5 Gaussian)

# A "High-Pass" filter

#### smoothed – original





# Hybrid Images

**Gaussian Filter** 

#### A. Oliva, A. Torralba, P.G. Schyns, <u>"Hybrid Images,"</u> SIGGRAPH 2006



#### DaVinci and Hybrid Images



#### Humans have less resolution (fewer sensors) near periphery



low frequencies

middle frequencies

high frequencies

# Low-pass, Band-pass, High-pass filters

#### low-pass:



#### band-pass:







#### Efficient construction of a collection of bandpass images

#### Lowpass Images



Bandpass Images

#### Gabor filters



reason for a "quadrature pair"

### Phase and Magnitude

- Curious fact
  - all natural images have about the same magnitude transform
  - hence, phase seems to matter, but magnitude largely doesn't
- Demonstration
  - Take two pictures, swap the phase transforms, compute the inverse - what does the result look like?













Reconstruction vith zebra phase, heetah nagnitude



Reconstruction with cheetah phase, zebra magnitude



#### Phase and Magnitude

Image with cheetah phase (and zebra magnitude)



Computer Vision - A Modern Approach - Set: Pyramids and Texture - Slides by D.A. Forsyth



Image with zebra phase (and cheetah magnitude)



Optimize DFT for real images; we only need to encode half the values

Turns out we don't need complex exponentials, only cosine functions = > discrete cosine transforms (DCT)

$$X_{k} = \sum_{n=0}^{N-1} x_{n} \cos \left[ \frac{\pi}{N} \left( n + \frac{1}{2} \right) k \right] \qquad k = 0, \dots, N-1.$$

#### Apply DCT on 8x8 pixel blocks

#### low frequencies



#### Quantize

More coarsely for high frequencies (which also tend to have smaller values)

v

Many quantized high frequency values will be zero

#### Encode

• Can decode with inverse dct



#### Quantization table

	[16	11	10	16	24	40	51	61
Q =	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



89k



12k

#### Application: <u>"Style Transfer for Headshot Portraits" (SIGGRAPH '14)</u>





Example



# Step 2: multi-scale local transfer



Input



Example
# Step 2: multi-scale local transfer

Construct Laplacian stacks for the input and the example.



Example

# Step 2: multi-scale local transfer

1. Construct Laplacian stacks for the input and the example



 Local match at each scale

# Step 2: multiscale transfer of local statistics

#### 1. Construct Laplacian stacks for the input and the example



## Local energy S



Example Laplacian Local energy Gaussian kernel at this scale

#### At each scale: match local energy



Input energy



Example energy

#### At each scale: match local energy



Compute the gain map

## At each scale: match local energy



Compute the gain map

Modulate

#### https://www.youtube.com/watch?v=Hj5IGFzIubU