# Lecture 16

Recognition, kNN and PCA

Ruta Desai, Chun-Liang Li



# Administrative

A4 should be out soon

- Due May 30th

Memorial day:

- No class on May 26th.
- We will cover some dimensionality reduction today (PCA)

Recitation this week:

- Will cover LDA for dimensionality reduction

Final exam:

- June 9th, Make up on June 6th
- T/F, MCQ, Short answers

#### Ruta Desai, Chun-Liang Li



## So far: Segmentation

cluster meaningful groups of pixels



#### Ruta Desai, Chun-Liang Li

#### Lecture 16 - 3

## So far: Segmentation



Segmentation using graph cuts





#### Ruta Desai, Chun-Liang Li

#### Lecture 16 - 4

# Today's agenda

- Introduction to recognition
- A object recognition pipeline
- Choosing the right features
- A training algorithm: KNN
- Testing an algorithm
- Challenges with kNN
- Dimensionality reduction: PCA

#### Ruta Desai, Chun-Liang Li

#### Lecture 16 - 5

# Today's agenda

- Introduction to recognition
- A object recognition pipeline
- Choosing the right features
- A training algorithm: KNN
- Testing an algorithm
- Challenges with kNN
- Dimensionality reduction: PCA

#### Ruta Desai, Chun-Liang Li

#### Lecture 16 - 6

## What do we mean by recognition?



Ruta Desai, Chun-Liang Li



## **Classification**: Does this image contain a building? [yes/no]



Ruta Desai, Chun-Liang Li



## **Classification:** Is this an beach?



#### Ruta Desai, Chun-Liang Li



## Applications: Image Search & Organizing photo collections



#### Ruta Desai, Chun-Liang Li



## **Detection:** Does this image contain a car? [where?]



Ruta Desai, Chun-Liang Li



## **Detection:** Which object does this image contain? [where?]



Ruta Desai, Chun-Liang Li

Lecture 16 - 12

## **Detection:** Accurate localization (segmentation)



Ruta Desai, Chun-Liang Li



## **Detection:** Estimating object semantic & geometric attributes



Ruta Desai, Chun-Liang Li

Lecture 16 - 14

## Levels of recognition: Category-level vs instance-level

Does this image contain the Chicago Macy's building?



Ruta Desai, Chun-Liang Li



## **Categorization vs Single instance recognition**

We have seen a form of single instance categorization already: Where is the crunchy nut?





#### Ruta Desai, Chun-Liang Li



## Applications of computer vision



Recognizing landmarks in mobile devices

#### Ruta Desai, Chun-Liang Li



## Activity recognition: What are these people doing?



Ruta Desai, Chun-Liang Li



## **Visual Recognition**

- Design algorithms that can:
  - $\circ \text{Classify}$  images or videos
  - $\circ \text{Detect}$  and localize objects
  - Estimate semantic and geometrical attributes
  - Classify human activities and events

## Why is this challenging?

Ruta Desai, Chun-Liang Li



How many object categories are there?



Ruta Desai, Chun-Liang Li

Lecture 16 - 20

## **Challenges**: viewpoint variation



Michelangelo 1475-1564

#### Ruta Desai, Chun-Liang Li



## **Challenges**: illumination



image credit: J. Koenderink

#### Ruta Desai, Chun-Liang Li

#### Lecture 16 - 22

#### Challenges: scale



#### Ruta Desai, Chun-Liang Li

Lecture 16 - 23

#### **Challenges**: deformation





#### Ruta Desai, Chun-Liang Li



### Challenges: occlusion



Magritte, 1957

#### Ruta Desai, Chun-Liang Li

Lecture 16 - 25

## Art Segway - <u>Magritte</u>





#### Ruta Desai, Chun-Liang Li

#### Lecture 16 - 26

## **Challenges**: background clutter



Kilmeny Niland. 1995

#### Ruta Desai, Chun-Liang Li

Lecture 16 - 27

### Challenges: intra-class variation



#### Ruta Desai, Chun-Liang Li

#### Lecture 16 - 28

# Today's agenda

- Introduction to recognition
- A object recognition pipeline
- Choosing the right features
- A training algorithm: KNN
- Testing an algorithm
- Challenges with kNN
- Dimensionality reduction: PCA

#### Ruta Desai, Chun-Liang Li Lecture 16 - 29



## Object recognition: a classification framework

• Apply a prediction function to a feature representation of the image to get the desired output:



Ruta Desai, Chun-Liang Li



## A simple pipeline - Training



#### Ruta Desai, Chun-Liang Li





#### Lecture 16 - 32













Lecture 16 - 35

# What we will learn today?

- Introduction to recognition
- A object recognition pipeline
- Choosing the right features
- A training algorithm: KNN
- Testing an algorithm
- Challenges with kNN
- Dimensionality reduction: PCA

#### Ruta Desai, Chun-Liang Li




Ruta Desai, Chun-Liang Li

#### Lecture 16 - 37

### May 21, 2025

	Invariances									
	Translation	Scale	<b>Rotation</b> (relative to camera plane)	<b>Rotation</b> (unconstrained)	Partial Occlusion	Illumination	Gaussian Noise			
RGB-histogram	?									

Ruta Desai, Chun-Liang Li



	Invariances									
	TranslationScaleRotation (relative to camera plane)Rotation (unconstrained)Partial OcclusionIlluminationGaussian Noise									
RGB-histogram										

✓ (global color counts don't change if the image shifts)

### Ruta Desai, Chun-Liang Li



	Invariances									
	Translation	Scale	<b>Rotation</b> (relative to camera plane)	<b>Rotation</b> (unconstrained)	Partial Occlusion	Illumination	Gaussian Noise			
RGB-histogram		?								

Ruta Desai, Chun-Liang Li



	Invariances									
	TranslationScaleRotation (relative to camera plane)Rotation (unconstrained)Partial OcclusionIlluminationGaussian Noise									
RGB-histogram		X								

✓ (if the *entire* image is uniformly scaled, the color distribution remains the same)

#### Ruta Desai, Chun-Liang Li



	Invariances									
	TranslationScaleRotation (relative to camera plane)Rotation (unconstrained)Partial OcclusionIllumination									
RGB-histogram	$\checkmark$	×	?	?						

Ruta Desai, Chun-Liang Li



	Invariances									
	Translation	Scale	<b>Rotation</b> (relative to camera plane)	<b>Rotation</b> (unconstrained)	Partial Occlusion	Illumination	Gaussian Noise			
RGB-histogram		×		X						

✓ (rotating the entire image does not change overall color distribution)

X (appearance/colors can change if out-of-plane rotation reveals different surfaces)

### Ruta Desai, Chun-Liang Li



	Invariances									
	Translation	Scale	<b>Rotation</b> (relative to camera plane)	<b>Rotation</b> (unconstrained)	Partial Occlusion	Illumination	Gaussian Noise			
RGB-histogram		×		×	?					

Ruta Desai, Chun-Liang Li



	Invariances									
	Translation	Scale	<b>Rotation</b> (relative to camera plane)	<b>Rotation</b> (unconstrained)	Partial Occlusion	Illumination	Gaussian Noise			
RGB-histogram		×		×	X					

**X** (removing part of the image can significantly alter color histogram)

Ruta Desai, Chun-Liang Li



	Invariances									
	Translation	Scale	<b>Rotation</b> (relative to camera plane)	<b>Rotation</b> (unconstrained)	Partial Occlusion	Illumination	Gaussian Noise			
RGB-histogram		×		×	×	?	?			

Ruta Desai, Chun-Liang Li



	Invariances									
	Translation	Scale	Rotation (relative to camera plane)	<b>Rotation</b> (unconstrained)	Partial Occlusion	Illumination	Gaussian Noise			
RGB-histogram	$\checkmark$	×		×	×	X	X			

**X** (shifts in illumination change color intensities/distribution)

**X** (noise directly alters pixel distribution)

### Ruta Desai, Chun-Liang Li



	Invariances									
	Translation	Scale	<b>Rotation</b> (relative to camera plane)	<b>Rotation</b> (unconstrained)	Partial Occlusion	Illumination	Gaussian Noise			
RGB-histogram		×		×	×	×	×			
HoG	?	?	?	?						

Ruta Desai, Chun-Liang Li



	Invariances									
	Translation	Scale	<b>Rotation</b> (relative to camera plane)	<b>Rotation</b> (unconstrained)	Partial Occlusion	Illumination	Gaussian Noise			
RGB-histogram		×		×	×	×	×			
HoG	X	X	X	X						

X (local bins move)

X (needs re-computation at multiple scales)

- X (oriented gradients are tied to an image grid)
- X (same reason as the ^)

#### Ruta Desai, Chun-Liang Li



	Invariances									
	Translation	Scale	<b>Rotation</b> (relative to camera plane)	<b>Rotation</b> (unconstrained)	Partial Occlusion	Illumination	Gaussian Noise			
RGB-histogram		×		×	×	×	×			
HoG	×	×	×	×	?	?	?			

Ruta Desai, Chun-Liang Li



	Invariances									
	Translation	Scale	<b>Rotation</b> (relative to camera plane)	<b>Rotation</b> (unconstrained)	Partial Occlusion	Illumination	Gaussian Noise			
RGB-histogram		×		×	×	×	×			
HoG	×	×	×	×	X		X			

- **X** (partial occlusion would result in no match)
- ✓ (gradients are more stable under monotonic intensity changes)
- X (gradient orientations can be disrupted by significant noise)

### Ruta Desai, Chun-Liang Li



	Invariances									
	Translation	Scale	<b>Rotation</b> (relative to camera plane)	<b>Rotation</b> (unconstrained)	Partial Occlusion	Illumination	Gaussian Noise			
RGB-histogram		×		×	×	×	×			
HoG	X	×	×	×	X		×			
SIFT	?	?	?	?						

Ruta Desai, Chun-Liang Li



	Invariances									
	Translation	Scale	<b>Rotation</b> (relative to camera plane)	<b>Rotation</b> (unconstrained)	Partial Occlusion	Illumination	Gaussian Noise			
RGB-histogram		×		×	×	×	×			
HoG	×	×	×	×	×		×			
SIFT				X						

- ✓ (keypoint-based, unaffected by shift)
- ✓ (built-in scale normalization)
- ✓ (SIFT normalizes orientation)
- X (local keypoints might disappear if the object rotates)

### Ruta Desai, Chun-Liang Li



	Invariances									
	Translation	Scale	<b>Rotation</b> (relative to camera plane)	<b>Rotation</b> (unconstrained)	Partial Occlusion	Illumination	Gaussian Noise			
RGB-histogram	$\checkmark$	×		×	×	×	×			
HoG	X	×	×	×	×		×			
SIFT	$\checkmark$		$\checkmark$	×	?	?	?			

#### Ruta Desai, Chun-Liang Li



	Invariances									
	Translation	Scale	<b>Rotation</b> (relative to camera plane)	<b>Rotation</b> (unconstrained)	Partial Occlusion	Illumination	Gaussian Noise			
RGB-histogram	$\checkmark$	×		×	×	×	×			
HoG	×	×	×	×	X		X			
SIFT	$\checkmark$			×						
Deep learning										

- ✓ (local keypoints can still match if some are visible)
- ✓ (gradient-based + normalization)
- ✓ (SIFT is relatively robust to moderate noise)

### Ruta Desai, Chun-Liang Li



	Invariances									
	Translation	Scale	<b>Rotation</b> (relative to camera plane)	<b>Rotation</b> (unconstrained)	Partial Occlusion	Illumination	Gaussian Noise			
RGB-histogram		×		×	×	×	×			
HoG	X	×	×	×	×		X			
SIFT	$\checkmark$	$\checkmark$		×	$\checkmark$		$\checkmark$			
Deep learning	?	?	?	?						

### Ruta Desai, Chun-Liang Li



	Invariances										
	Translation	Scale	<b>Rotation</b> (relative to camera plane)	<b>Rotation</b> (unconstrained)	Partial Occlusion	Illumination	Gaussian Noise				
RGB-histogram		×	$\checkmark$	×	×	×	×				
HoG	X	×	X	X	×		×				
SIFT			$\checkmark$	×							
Deep learning	usually	usually	usually	X							

Deep learning features are usually invariant to translation, scale, and planar rotation if the training data has these translations. It is not invariant to other rotations.
Aside: ImageNet has objects centered in the middle of images. So models trained on ImageNet are not translation or scale invariant.

#### Ruta Desai, Chun-Liang Li



	Invariances									
	Translation	Scale	<b>Rotation</b> (relative to camera plane)	<b>Rotation</b> (unconstrained)	Partial Occlusion	Illumination	Gaussian Noise			
RGB-histogram	$\checkmark$	×		×	×	×	×			
HoG	×	×	×	×	×		×			
SIFT	$\checkmark$	$\checkmark$		×	$\checkmark$					
Deep learning	usually	usually	usually	×	?	?	?			

Ruta Desai, Chun-Liang Li



	Invariances									
	Translation	Scale	<b>Rotation</b> (relative to camera plane)	<b>Rotation</b> (unconstrained)	Partial Occlusion	Illumination	Gaussian Noise			
RGB-histogram	$\checkmark$	×		X	×	×	×			
HoG	×	×	×	X	×		×			
SIFT	$\checkmark$	$\checkmark$		X	$\checkmark$					
Deep learning	usually	usually	usually	×	X					

- X (standard CNNs are not strictly occlusion-invariant; partial robustness depends on training)
- ✓ (can learn robustness if trained on varied lighting)
- ✓ (CNNs can learn to be noise-robust with proper training)

### Ruta Desai, Chun-Liang Li



# So, which features should we choose?

	Invariances									
	Translation	Scale	<b>Rotation</b> (relative to camera plane)	<b>Rotation</b> (unconstrained)	Partial Occlusion	Illumination	Gaussian Noise			
RGB-histogram	$\checkmark$	×		×	×	×	×			
HoG	×	×	×	×	×		×			
SIFT	$\checkmark$			×	$\checkmark$					
Deep learning	usually	usually	usually	×	×		$\checkmark$			

Ruta Desai, Chun-Liang Li



# What we will learn today?

- Introduction to recognition
- A simple Object Recognition pipeline
- Choosing the right features
- A training algorithm: KNN
- Testing an algorithm
- Challenges with kNN
- Dimensionality reduction: PCA

### Ruta Desai, Chun-Liang Li



Learning a classifier to map inputs to outputs



- Training: given a *training set* of labeled examples {(x<sub>1</sub>,y<sub>1</sub>), ..., (x<sub>N</sub>,y<sub>N</sub>)}, estimate the prediction function f by minimizing the prediction error on the training set
- Testing: apply f to a never before seen test example x and output the predicted value y = f(x)

### Ruta Desai, Chun-Liang Li



### A simple pipeline - Training Training Labels Training Images Image Learned Training Classifier Features 1997 IN 1998 Test Image Image Learned Prediction Classifier Features

### Ruta Desai, Chun-Liang Li

#### Lecture 16 - 63

### May 21, 2025

### Many classifiers to choose from

- K-nearest neighbor
- SVM
- Neural networks
- Naïve Bayes
- Bayesian network
- Logistic regression
- Randomized Forests
- Boosted Decision Trees
- RBMs
- Etc.

### Which is the best one?

### Ruta Desai, Chun-Liang Li



### An example training dataset



Training set (labels known)

Ruta Desai, Chun-Liang Li

Apples Pear Tomatos Cow Dog Horse For kNN classifier, training simply means to store all training data.



# A stored training set



Slide credit: L. Lazebnik

#### Ruta Desai, Chun-Liang Li



# During testing, we assign the label of the nearest neighbot in feature space



Slide credit: L. Lazebnik

#### Ruta Desai, Chun-Liang Li



# What we will learn today?

- Introduction to recognition
- A simple Object Recognition pipeline
- Choosing the right features
- A training algorithm: kNN
- Testing an algorithm
- Challenges with kNN
- Dimensionality reduction

### Ruta Desai, Chun-Liang Li



### A simple pipeline - Training



Ruta Desai, Chun-Liang Li

#### Lecture 16 - 69

### May 21, 2025

### Generalization



Training set (labels known)

Test set (labels unknown)

• How well does a learned model generalize from the data it was trained on to a new test set?

### Ruta Desai, Chun-Liang Li



### Intuition for Nearest Neighbor Classifier

Given a training dataset, simply store each image's features and their corresponding label.



#### Ruta Desai, Chun-Liang Li



### Intuition for Nearest Neighbor Classifier

Given a training dataset, simply store each image's features and their corresponding label.



#### Ruta Desai, Chun-Liang Li

#### Lecture 16 - 72

### May 21, 2025
### Nearest Neighbor Classifier

• Assign label of majority of K=3 nearest neighbors



#### Ruta Desai, Chun-Liang Li

#### Lecture 16 - 73

### Classification

- Assign input vector to one of many classes (categories)
- Geometric interpretation of classifiers: A classifier divides input space into *decision regions* separated by *decision boundaries*



#### Ruta Desai, Chun-Liang Li

Lecture 16 - 74

### Nearest Neighbor Classifier

• Assign label of nearest training data point to each test data point



from Duda et al.

Partitioning of feature space for two-category 2D and 3D data

#### Ruta Desai, Chun-Liang Li



### How do we find the nearest neighbors in feature space?

Distance measure (same as the ones from segmentation)

Euclidean:  

$$Dist(X^n, X^m) = \sqrt{\sum_{i=1}^{D} (X_i^n - X_i^m)^2}$$

Where X<sup>n</sup> and X<sup>m</sup> are the n-th and m-th data points



#### Ruta Desai, Chun-Liang Li



### 1-nearest neighbor

Distance measure (same as the ones from segmentation)

Euclidean:  

$$Dist(X^n, X^m) = \sqrt{\sum_{i=1}^{D} (X_i^n - X_i^m)^2}$$

Where X<sup>n</sup> and X<sup>m</sup> are the n-th and m-th data points



#### Ruta Desai, Chun-Liang Li

#### Lecture 16 - 77

### 3-nearest neighbor

Distance measure (same as the ones from segmentation)

Euclidean:  

$$Dist(X^n, X^m) = \sqrt{\sum_{i=1}^{D} (X_i^n - X_i^m)^2}$$

Where X<sup>n</sup> and X<sup>m</sup> are the n-th and m-th data points



#### Ruta Desai, Chun-Liang Li



### 5-nearest neighbor

Distance measure (same as the ones from segmentation)

Euclidean:  

$$Dist(X^n, X^m) = \sqrt{\sum_{i=1}^{D} (X_i^n - X_i^m)^2}$$

Where X<sup>n</sup> and X<sup>m</sup> are the n-th and m-th data points



#### Ruta Desai, Chun-Liang Li



# Choosing the right features is important but dataset-dependent



	Color	$D_x D_y$	Mag-Lap	PCA Masks	PCA Gray	Cont. Greedy	Cont. DynProg	Avg.
apple	57.56%	85.37%	80.24%	78.78%	88.29%	77.07%	76.34%	77.66%
pear	66.10%	90.00%	85.37%	99.51%	99.76%	90.73%	91.71%	89.03%
tomato	98.54%	94.63%	97.07%	67.80%	76.59%	70.73%	70.24%	82.23%
cow	86.59%	82.68%	94.39%	75.12%	62.44%	86.83%	86.34%	82.06%
dog	34.63%	62.44%	74.39%	72.20%	66.34%	81.95%	82.93%	67.84%
horse	32.68%	58.78%	70.98%	77.80%	77.32%	84.63%	84.63%	69.55%
cup	79.76%	66.10%	77.80%	96.10%	96.10%	99.76%	99.02%	87.81%
car	62.93%	98.29%	77.56%	100.0%	97.07%	99.51%	100.0%	90.77%
total	64.85%	79.79%	82.23%	83.41%	82.99%	86.40%	86.40%	80.87%

Dataset: ETH-80, by B. Leibe, 2003

May 21, 2025

#### Ruta Desai, Chun-Liang Li

### Results

					0				Ó
	Ò		6						
0	Ó	٢	Ó	٢		•			0
-7PI	<b>1</b>	1	1	<u>IS</u>	1	7	1		1.7
**	<b>%</b> ?-	92:	73	33	7933	Ħ	M	35	Yn:
M	77	1	37	77	1	27	To	1	1
Þ			<b>P</b>			V	<b>_</b>	9	
		<u>~</u>	-	<b>1</b>			<b>1</b>		<b>~</b>

Category	Primary feature(s)	Secondary reature(s)
apple	PCA Gray	Texture $D_x D_y$
pear	PCA Gray / Masks	
tomato	Color	Texture Mag-Lap
cow	Texture Mag-Lap	Contour / Color
dog	Contour	
horse	Contour	
cup	Contour	PCA Gray / Masks
car	PCA Masks / Contour	Texture $D_x D_y$

Dataset: ETH-80, by B. Leibe, 2003

May 21, 2025

#### Ruta Desai, Chun-Liang Li

### K-NN: a very useful algorithm

- Simple, a good one to try first
- Very flexible decision boundaries

#### Ruta Desai, Chun-Liang Li



# What we will learn today?

- Introduction to recognition
- A simple Object Recognition pipeline
- Choosing the right features
- A training algorithm: kNN
- Testing an algorithm
- Challenges with kNN
- Dimensionality reduction

#### Ruta Desai, Chun-Liang Li



- Choosing the value of k:
  - $\circ$  If too small, sensitive to noise points
  - If too large, neighborhood may include points from other classes



#### Ruta Desai, Chun-Liang Li



- Choosing the value of k:
  - $\circ$  If too small, sensitive to noise points
  - If too large, neighborhood may include points from other classes





#### Ruta Desai, Chun-Liang Li

#### Lecture 16 - 85

- Choosing the value of k:
  - $\circ$  If too small, sensitive to noise points

 If too large, neighborhood may include points from other classes

Solution: Cross validate



#### Ruta Desai, Chun-Liang Li

#### Lecture 16 - 86

All Data

Training data Test data

Ruta Desai, Chun-Liang Li



All Data	
Training data	Test data

	Fold 1	Fold 2	Fold 3	Fold 4	Fold 5
Split 1	Fold 1	Fold 2	Fold 3	Fold 4	Fold 5

#### Ruta Desai, Chun-Liang Li



	All Data							
		Г	Test data					
	Fold 1	Fold 2	Fold 3	Fold 4	Fold 5			
Split 1	Fold 1	Fold 2	Fold 3	Fold 4	Fold 5			
Split 2	Fold 1	Fold 2	Fold 3	Fold 4	Fold 5			

Ruta Desai, Chun-Liang Li





Ruta Desai, Chun-Liang Li



- Choosing the value of k:
  - If too small, sensitive to noise points
  - If too large, neighborhood may include points from other classes
  - Solution: cross validate!
- Curse of Dimensionality

#### Ruta Desai, Chun-Liang Li



### Curse of dimensionality

- As the dimensionality increases, the number of data points required for good performance increases exponentially.
- Let's say that for a model to perform well, we need at least 10 data points for each combination of feature values.

Need for Data Points with Increase in Dimensions



Ruta Desai, Chun-Liang Li

#### Lecture 16 - 92

#### • Choosing the value of k:

○ If too small, sensitive to noise points

 If too large, neighborhood may include points from other classes

• Solution: cross validate!

• Curse of Dimensionality

Solution: dimensionality reduction

#### Ruta Desai, Chun-Liang Li



# What we will learn today

- Introduction to recognition
- A simple Object Recognition pipeline
- Choosing the right features
- A training algorithm: kNN
- Testing an algorithm
- Challenges with kNN
- Dimensionality reduction

#### Ruta Desai, Chun-Liang Li



# Singular Value Decomposition (SVD)

### $\mathbf{U}\mathbf{\Sigma}\mathbf{V}^{\mathsf{T}} = \mathbf{A}$

• Where **U** and **V** are rotation matrices, and  $\boldsymbol{\Sigma}$  is a scaling matrix. For example:

$$\begin{bmatrix} U & \Sigma & V^T & A \\ -.40 & .916 \\ .916 & .40 \end{bmatrix} \times \begin{bmatrix} 5.39 & 0 \\ 0 & 3.154 \end{bmatrix} \times \begin{bmatrix} -.05 & .999 \\ .999 & .05 \end{bmatrix} = \begin{bmatrix} 3 & -2 \\ 1 & 5 \end{bmatrix}$$

Ruta Desai, Chun-Liang Li



# Singular Value Decomposition (SVD)

- Beyond 2x2 matrices:
  - In general, if A is m x n, then U will be m x m, Σ will be m x n, and V<sup>T</sup> will be n x n.
  - (Note the dimensions work out to produce *m* x *n* after multiplication)

$$\begin{bmatrix} U & \Sigma & V^T \\ -.39 & -.92 \\ -.92 & .39 \end{bmatrix} \times \begin{bmatrix} 9.51 & 0 & 0 \\ 0 & .77 & 0 \end{bmatrix} \times \begin{bmatrix} -.42 & -.57 & -.70 \\ .81 & .11 & -.58 \\ .41 & -.82 & .41 \end{bmatrix} = \begin{bmatrix} 1 & 2 & 3 \\ 4 & 5 & 6 \end{bmatrix}$$

#### Ruta Desai, Chun-Liang Li

Lecture 16 - 96

# Singular Value Decomposition (SVD)

- U and V are always rotation matrices.
  - Geometric rotation may not be an applicable concept, depending on the matrix. So we call them "unitary" matrices – each column is a unit vector.
- $\Sigma$  is a diagonal matrix
  - The number of nonzero entries = rank of **A**
  - $\circ$  The algorithm always sorts the entries high to low

$$\begin{bmatrix} U & \Sigma & V^T \\ -.39 & -.92 \\ -.92 & .39 \end{bmatrix} \times \begin{bmatrix} 9.51 & 0 & 0 \\ 0 & .77 & 0 \end{bmatrix} \times \begin{bmatrix} -.42 & -.57 & -.70 \\ .81 & .11 & -.58 \\ .41 & -.82 & .41 \end{bmatrix} = \begin{bmatrix} 1 & 2 & 3 \\ 4 & 5 & 6 \end{bmatrix}$$

#### Ruta Desai, Chun-Liang Li

#### Lecture 16 - 97

# **SVD** Applications

- We've discussed SVD in terms of geometric transformation matrices
- But SVD of an image matrix can also be very useful
- To understand this, we'll look at a less geometric interpretation of what SVD is doing



# What is SVD actually doing for images?

$$\begin{bmatrix} U & \Sigma & V^T \\ -.39 & -.92 \\ -.92 & .39 \end{bmatrix} \times \begin{bmatrix} 9.51 & 0 & 0 \\ 0 & .77 & 0 \end{bmatrix} \times \begin{bmatrix} -.42 & -.57 & -.70 \\ .81 & .11 & -.58 \\ .41 & -.82 & .41 \end{bmatrix} = \begin{bmatrix} 1 & 2 & 3 \\ 4 & 5 & 6 \end{bmatrix}$$

- Look at how the multiplication works out, left to right:
- Column 1 of **U** gets scaled by the first value from **Σ**.



#### Ruta Desai, Chun-Liang Li



## What is SVD actually doing for images?

$$\begin{bmatrix} U & \Sigma & V^T \\ -.39 & -.92 \\ -.92 & .39 \end{bmatrix} \times \begin{bmatrix} 9.51 & 0 & 0 \\ 0 & .77 & 0 \end{bmatrix} \times \begin{bmatrix} -.42 & -.57 & -.70 \\ .81 & .11 & -.58 \\ .41 & -.82 & .41 \end{bmatrix} = \begin{bmatrix} 1 & 2 & 3 \\ 4 & 5 & 6 \end{bmatrix}$$

- Look at how the multiplication works out, left to right:
- Column 1 of **U** gets scaled by the first value from  $\Sigma$ .

$$\begin{bmatrix} U\Sigma & V^T \\ -3.67 & -.71 & 0 \\ -8.8 & .30 & 0 \end{bmatrix} \times \begin{bmatrix} -.42 & -.57 & -.70 \\ .81 & .11 & -.58 \\ .41 & -.82 & .41 \end{bmatrix}$$

Ruta Desai, Chun-Liang Li



# What is SVD actually doing for images?

$$\begin{bmatrix} U & \Sigma & V^T \\ -.39 & -.92 \\ -.92 & .39 \end{bmatrix} \times \begin{bmatrix} 9.51 & 0 & 0 \\ 0 & .77 & 0 \end{bmatrix} \times \begin{bmatrix} -.42 & -.57 & -.70 \\ .81 & .11 & -.58 \\ .41 & -.82 & .41 \end{bmatrix} = \begin{bmatrix} 1 & 2 & 3 \\ 4 & 5 & 6 \end{bmatrix}$$

- Look at how the multiplication works out, left to right:
- Column 1 of **U** gets scaled by the first value from **Σ**.

$$\begin{bmatrix} -3.67 & -.71 & 0 \\ -8.8 & .30 & 0 \end{bmatrix} \times \begin{bmatrix} -.42 & -.57 & -.70 \\ .81 & .11 & -.58 \\ .41 & -.82 & .41 \end{bmatrix} \xrightarrow{A_{partial}} \begin{bmatrix} 1.6 & 2.1 & 2.6 \\ 3.8 & 5.0 & 6.2 \end{bmatrix}$$

 The resulting vector gets scaled by row 1 of V<sup>T</sup> to produce a contribution to the columns of A

#### Ruta Desai, Chun-Liang Li





Each product of (column i of U)·(value i from Σ)·(row i of V<sup>T</sup>) produces a component of the final A.

#### Ruta Desai, Chun-Liang Li





- We're building **A** as a linear combination of the columns of **U**
- Using all columns of **U**, we'll rebuild the original matrix perfectly
- But, in real-world data, often we can just use the first few columns of *U* and we'll get something close (e.g. the first *A*<sub>partial</sub>, above)

#### Ruta Desai, Chun-Liang Li





- We can call those first few columns of *U* the Principal Components of the data
- They show the major patterns that can be added to produce the columns of the original matrix
- The rows of V<sup>T</sup> show how the *principal components* are mixed to produce the columns of the matrix

#### Ruta Desai, Chun-Liang Li

#### Lecture 16 - 104

$$\begin{bmatrix} U & \Sigma & V^T \\ -.39 & -.92 \\ -.92 & .39 \end{bmatrix} \times \begin{bmatrix} 9.51 & 0 & 0 \\ 0 & .77 & 0 \end{bmatrix} \times \begin{bmatrix} -.42 & -.57 & -.70 \\ .81 & .11 & -.58 \\ .41 & -.82 & .41 \end{bmatrix} = \begin{bmatrix} 1 & 2 & 3 \\ 4 & 5 & 6 \end{bmatrix}$$

We can look at Σ to see that the first column has a large effect

while the second column has a much smaller effect in this example

#### Ruta Desai, Chun-Liang Li

#### Lecture 16 - 105

### Image compression

Ruta Desai, Chun-Liang Li



May 21, 2025

• For this image, using **only the first 16** of 300 principal components produces a recognizable reconstruction

Lecture 16 - 106

 Using the first 64 almost perfectly reconstructs the image

### Intuition behind PCA: high dimensional data usually lives in some lower dimensional space

**Covariance** between the two dimensions of features is high. Can we reduce the number of dimensions to just 1?



#### Ruta Desai, Chun-Liang Li

#### Lecture 16 - 107

### Geometric interpretation of PCA



Ruta Desai, Chun-Liang Li


### Geometric interpretation of PCA

- Let's say we have a set of 2D data points x. But we see that all the points lie on a line in 2D.
- So, 2 dimensions are redundant to express the data. We can express all the points with just one dimension.



#### Ruta Desai, Chun-Liang Li



## PCA: Principal Component Analysis

- Given a dataset of images, can we compressed them like we can compress a single image?
  - Yes, the trick is to look into the correlation between the dimensions of the image
  - $\circ$  The tool for doing this is called PCA

PCA can be used to compress image RGB pixel values or also be used to compress their features!

Ruta Desai, Chun-Liang Li



### Toy example to explain covariance

- What is covariance between dimensions?
- Let's say we have a dataset of students

  each student is represented with 3 dimensions
  x: number of hours studied for a class
  y: grades obtained in that class
  z: number of lectures attended

  covariance value between x and y is say: 104.53
  - what does this value mean?

#### Ruta Desai, Chun-Liang Li



### **Covariance interpretation**

- $\circ$  **x**: number of hours studied for a subject
- y: marks obtained in that subject
- covariance value between **x** and **y** is say: 104.53

 $\circ$  what does this value mean?



#### Ruta Desai, Chun-Liang Li

Lecture 16 - 112

## Visualizing this covariance matrix

- We can represent these covariance correlation numbers in a matrix
- e.g. for 3 dimensions:



- Diagonal is the **variances** of x, y and z
- cov(x,y) = cov(y,x) hence C is symmetrical about the diagonal
- N-dimensional data will result in NxN covariance matrix

#### Ruta Desai, Chun-Liang Li



### **Covariance interpretation**

- Exact value is not as important as it's sign.
- A **positive value** of covariance indicates both dimensions increase or decrease together e.g. as the number of hours studied increases, the marks in that subject increase.
- A **negative value** indicates while one increases the other decreases, or vice-versa e.g. active social life at PSU vs performance in CS dept.
- If covariance is zero: the two dimensions are independent of each other e.g. heights of students vs the marks obtained in a subject



• To relate this to PCA, we consider the <u>image (or feature) matrix</u>

$$X = \begin{bmatrix} 1 & 1 \\ x_1 & \dots & x_n \\ 1 & 1 \end{bmatrix}$$

 The sample mean of this dataset (or in plain english, the average image) is:

$$\mu = \frac{1}{n} \sum_{i} x_{i} = \frac{1}{n} \begin{bmatrix} 1 & 1 & 1 \\ x_{1} & \dots & x_{n} \\ 1 & 1 \end{bmatrix} = \frac{1}{n} X 1$$

Ruta Desai, Chun-Liang Li



- Center the data by subtracting the mean to each column of X
- The <u>centered dataset matrix</u> is

$$X_{c} = \begin{bmatrix} 1 & 1 \\ x_{1} & \dots & x_{n} \\ 1 & 1 \end{bmatrix} - \begin{bmatrix} 1 & 1 \\ \mu & \dots & \mu \\ 1 & 1 \end{bmatrix}$$



• The sample <u>covariance</u> matrix is

$$C = \frac{1}{n} \sum_{i} (x_i - \mu) (x_i - \mu)^T = \frac{1}{n} \sum_{i} x_i^c (x_i^c)^T$$

where  $x_i^{c}$  is the i<sup>th</sup> column of  $X_c$ 

• This can be written as



#### Ruta Desai, Chun-Liang Li

#### Lecture 16 - 117

• The matrix

$$\boldsymbol{X}_{c}^{T} = \begin{bmatrix} - & \boldsymbol{X}_{1}^{c} & - \\ \vdots & \\ - & \boldsymbol{X}_{n}^{c} & - \end{bmatrix}$$

is real (n x d). Assuming n>d it has SVD decomposition

$$X_c^T = U\Sigma V^T \qquad \qquad U^T U = I \qquad \qquad V^T V = I$$

and

$$C = \frac{1}{n} X_c X_c^T$$

Ruta Desai, Chun-Liang Li



### Calculating covariance matrix

$$C = \frac{1}{n} X_c X_c^T$$
$$= \frac{1}{n} U \Sigma V^T (U \Sigma V^T)^T$$
$$= \frac{1}{n} U \Sigma V^T V \Sigma U^T$$
$$= \frac{1}{n} U \Sigma^2 U^T$$

n

#### Ruta Desai, Chun-Liang Li



$$C = \frac{1}{n} U \Sigma^2 U^T$$

- Note that U is (d x d) and orthonormal, and Σ<sup>2</sup> is diagonal.
   This is just the eigenvalue decomposition of C
- This means that we can calculate the eigenvectors of C using the eigenvectors of X<sub>c</sub>
- It follows that
  - $\circ\,$  The eigenvectors of C are the columns of U
  - $\circ$  The eigenvalues of C are the diagonal entries of  $\Sigma^2$ :  $\lambda_i^2$

#### Ruta Desai, Chun-Liang Li



- In summary, computation of PCA by SVD
- Given X with one image (or feature) per column

Create the centered data matrix

$$\boldsymbol{X}_{c} = \begin{bmatrix} \boldsymbol{I} & \boldsymbol{I} \\ \boldsymbol{X}_{1} & \boldsymbol{I} & \boldsymbol{X}_{n} \\ \boldsymbol{I} & \boldsymbol{I} \end{bmatrix} - \begin{bmatrix} \boldsymbol{I} & \boldsymbol{I} \\ \boldsymbol{\mu} & \boldsymbol{I} & \boldsymbol{\mu} \\ \boldsymbol{I} & \boldsymbol{I} \end{bmatrix}$$

• Compute its SVD

Ruta Desai, Chun-Liang Li

$$X_c^T = U\Sigma V^T$$

Principal components of the covariance matrix C are columns of U



To compress an image dataset, pick the largest eigenvalues and their corresponding eigenvectors

Pick the eigenvectors that explain p% of the image data variability
 Can be done by plotting the ratio r<sub>k</sub> as a function of k



 $r_{k} = \sum_{i=1}^{k} \lambda_{i}^{2}$   $\sum_{i=1}^{n} \lambda_{i}^{2}$ 

 $\circ$  E.g. we need k=3 eigenvectors to cover 70% of the variability of this dataset

#### Ruta Desai, Chun-Liang Li



### What exactly is the covariance

- Variance and Covariance are a measure of the "**spread**" of a set of points around their center of mass (mean)
- Variance measure of the deviation from the mean for points in one dimension e.g. heights
- **Covariance** as a measure of how much each of the dimensions vary from the mean with respect to each other.
- Covariance is measured between 2 dimensions to see if there is a relationship between the 2 dimensions e.g. number of hours studied & marks obtained.
- The covariance between one dimension and itself is the variance

#### Ruta Desai, Chun-Liang Li



#### What happens with PCA during training?



Ruta Desai, Chun-Liang Li

Lecture 16 - 124



#### Ruta Desai, Chun-Liang Li

#### Lecture 16 - 125

## PCA during training

Let's say that we choose k top eigenvalues and their corresponding eigenvectors:  $[u_1, ..., u_k]$ 

Replace all image features x with:

$$\hat{x} = \begin{bmatrix} u_1^T x \\ u_2^T x \\ \dots \\ u_k^T x \end{bmatrix}$$

Ruta Desai, Chun-Liang Li





Ruta Desai, Chun-Liang Li

Lecture 16 - 127



Ruta Desai, Chun-Liang Li

#### Lecture 16 - 128

How PCA was originally used in vision: To identify celebrities using their faces

- An image is a point in a high dimensional space
  - $\circ$  In grayscale, an N x M image is a point in  $R^{\text{NM}}$
  - E.g. 100x100 images lives in a 10,000-dimensional space



#### Ruta Desai, Chun-Liang Li

#### Lecture 16 - 129

### SVD for symmetric matrices

• If A is a symmetric matrix, it can be decomposed as the following:

# • Compared to a traditional $A=\Phi\Sigma\Phi^T\quad \mathbf{V}^{\mathrm{T}}$ and is an orthogonal matrix.

#### Ruta Desai, Chun-Liang Li



### What we have learned today?

- Introduction to recognition
- A simple Object Recognition pipeline
- Choosing the right features
- A training algorithm: kNN
- Testing an algorithm
- Challenges with kNN
- Dimensionality reduction: PCA

#### Ruta Desai, Chun-Liang Li



## Next lecture

**Object detection** 

Ruta Desai, Chun-Liang Li



## Extra slides (out of scope)

for those of you curious about how SVD is calculated and what PCA is usually used for outside of computer vision

Ruta Desai, Chun-Liang Li



### **Principal Component Analysis**



- Remember, columns of **U** are the *Principal Components* of the data: the major patterns that can be added to produce the columns of the original matrix
- One use of this is to construct a matrix where each column is a separate data sample
- Run SVD on that matrix, and look at the first few columns of U to see patterns that are common among the columns
- This is called *Principal Component Analysis* (or PCA) of the data samples

#### Ruta Desai, Chun-Liang Li



### **Principal Component Analysis**



- Often, raw data samples have a lot of redundancy and patterns
- PCA can allow you to represent data samples as weights on the principal components, rather than using the original raw form of the data
- By representing each sample as just those weights, you can represent just the "meat" of what's different between samples.
- This minimal representation makes machine learning and other algorithms much more efficient

#### Ruta Desai, Chun-Liang Li



## How is SVD computed?

- For this class: tell PYTHON to do it. Use the result.
- But, if you're interested, one computer algorithm to do it makes use of Eigenvectors!

#### Ruta Desai, Chun-Liang Li



### **Eigenvector definition**

- Suppose we have a square matrix **A**. We can solve for vector x and scalar  $\lambda$  such that Ax=  $\lambda$ x
- In other words, find vectors where, if we transform them with **A**, the only effect is to scale them with no change in direction.
- These vectors are called eigenvectors (German for "self vector" of the matrix), and the scaling factors λ are called eigenvalues

Lecture 16 - 137

May 21, 2025

• An  $m \ge m$  matrix will have  $\le m$  eigenvectors where  $\lambda$  is nonzero

#### Ruta Desai, Chun-Liang Li

## Finding eigenvectors

- Computers can find an x such that  $Ax = \lambda x$  using this iterative algorithm:
  - $\circ$  X = random unit vector
  - while(x hasn't converged)
    - X = Ax
    - normalize x
- x will quickly converge to an eigenvector
- Some simple modifications will let this algorithm find all eigenvectors

#### Ruta Desai, Chun-Liang Li



## Finding SVD

- Eigenvectors are for square matrices, but SVD is for all matrices
- To do svd(A), computers can do this:
  - $\circ$  Take eigenvectors of AA<sup>T</sup> (matrix is always square).
    - These eigenvectors are the columns of **U**.
    - Square root of eigenvalues are the singular values (the entries of Σ).
  - $\circ$  Take eigenvectors of A<sup>T</sup>A (matrix is always square).
    - These eigenvectors are columns of **V** (or rows of **V**<sup>T</sup>)

Lecture 16 - 139

May 21, 2025

#### Ruta Desai, Chun-Liang Li

## Finding SVD

- Moral of the story: SVD is fast, even for large matrices
- It's useful for a lot of stuff
- There are also other algorithms to compute SVD or part of the SVD
  - Python's np.linalg.svd() command has options to efficiently compute only what you need, if performance becomes an issue

A detailed geometric explanation of SVD is here: http://www.ams.org/samplings/feature-column/fcarc-svd

#### Ruta Desai, Chun-Liang Li



### Bias-Variance Trade-off



- Models with too few parameters are inaccurate because of a large bias (not enough flexibility).
- Models with too many parameters are inaccurate because of a large variance (too much sensitivity to the sample).

#### Ruta Desai, Chun-Liang Li



### **Bias versus variance**

- Components of generalization error
  - Bias: how much the average model over all training sets differ from the true model?
    - Error due to inaccurate assumptions/simplifications made by the model
  - Variance: how much models estimated from different training sets differ from each other
- **Underfitting:** model is too "simple" to represent all the relevant class characteristics
  - $\circ\,$  High bias and low variance
  - $\circ\,$  High training error and high test error
- **Overfitting:** model is too "complex" and fits irrelevant characteristics (noise) in the data
  - $\circ$  Low bias and high variance
  - $\circ\,$  Low training error and high test error

#### Ruta Desai, Chun-Liang Li



### Bias versus variance trade off



Ruta Desai, Chun-Liang Li



### No Free Lunch Theorem



In a supervised learning setting, we can't tell which classifier will have best generalization

#### Ruta Desai, Chun-Liang Li


### Remember...

- No classifier is inherently better than any other: you need to make assumptions to generalize
- Three kinds of error
  Inherent: unavoidable
  - Bias: due to over-simplifications
  - Variance: due to inability to perfectly estimate parameters from limited data



### Ruta Desai, Chun-Liang Li

### Lecture 16 - 145

### May: 21:, 2025

### How to reduce variance?

- Choose a simpler classifier
- Regularize the parameters
- Get more training data

How do you reduce bias?

### Ruta Desai, Chun-Liang Li



# Last remarks about applying machine learning methods to object recognition

- There are machine learning algorithms to choose from
- Know your data:
  - o How much supervision do you have?
  - $\circ$  How many training examples can you afford?
  - How noisy?
- Know your goal (i.e. task):
  - $\circ$  Affects your choices of representation
  - Affects your choices of learning algorithms
  - Affects your choices of evaluation metrics
- Understand the math behind each machine learning algorithm under consideration!

### Ruta Desai, Chun-Liang Li



# PCA by SVD

- An alternative manner to compute the principal components, based on singular value decomposition
- Quick reminder: SVD
  - Any real n x m matrix (n>m) can be decomposed as

 $A = M\Pi N^T$ 

- Where M is an (n x m) column orthonormal matrix of left singular vectors (columns of M)
- $\circ~\Pi$  is an (m x m) diagonal matrix of singular values
- $\circ$  N<sup>T</sup> is an (m x m) row orthonormal matrix of right singular vectors (columns of N)

$$\mathbf{M}^{\mathsf{T}}\mathbf{M} = \mathbf{I} \qquad \mathbf{N}^{\mathsf{T}}\mathbf{N} = \mathbf{I}$$

### Ruta Desai, Chun-Liang Li

### Lecture 16 - 148

<u>May 21, 2025</u>

## **PCA Formulation**

• Basic idea:

 If the images (or their features) live in a subspace, it is going to look very flat when viewed from the full feature space, e.g.



Slide inspired by N. Vasconcelos

#### Ruta Desai, Chun-Liang Li



# **Alternative PCA Formulation**

- Assume images x is Gaussian with covariance Σ.
- Recall that a gaussian is defined with it's mean and variance:

 $\mathbf{X}~\sim~\mathcal{N}(oldsymbol{\mu},\,oldsymbol{\Sigma})$ 

• Recall that  $\mu$  and  $\Sigma$  of a gaussian are defined as:

$$\boldsymbol{\mu} = \mathrm{E}[\mathbf{X}]$$

$$oldsymbol{\Sigma} =: \mathrm{E}[(oldsymbol{X} - oldsymbol{\mu})^{\mathrm{T}}] = [\mathrm{Cov}[X_i, X_j]; 1 \leq i,j \leq k]$$



May 21, 2025

#### Ruta Desai, Chun-Liang Li

### Alternative PCA formulation

 Since gaussians are symmetric, it's covariance matrix is also a symmetric matrix. So we can express it as:

 $\circ \boldsymbol{\Sigma} = \boldsymbol{U} \boldsymbol{\Lambda} \boldsymbol{U}^{\mathsf{T}} = \boldsymbol{U} \boldsymbol{\Lambda}^{1/2} (\boldsymbol{U} \boldsymbol{\Lambda}^{1/2})^{\mathsf{T}}$ 

$$\mathbf{X} \sim \mathcal{N}(oldsymbol{\mu}, oldsymbol{\Sigma}) \iff \mathbf{X} \sim oldsymbol{\mu} + \mathbf{U} \mathbf{\Lambda}^{1/2} \mathcal{N}(0, \mathbf{I})$$

 $\iff \mathbf{X} \sim \boldsymbol{\mu} + \mathbf{U} \mathcal{N}(0, \boldsymbol{\Lambda}).$ 

May 21, 2025

Ruta Desai, Chun-Liang Li

# **Alternative PCA Formulation**

- If x is Gaussian with covariance  $\Sigma$ ,
  - Principal components φ<sub>i</sub> are the eigenvectors of Σ
    Principal lengths λ<sub>i</sub> are the eigenvalues of Σ
- by computing the eigenvalues we know the data is
  Not flat if λ<sub>1</sub> ≈ λ<sub>2</sub>
  Flat if λ<sub>1</sub> >> λ<sub>2</sub>

#### Ruta Desai, Chun-Liang Li



### Alternative PCA Algorithm (training)

▶ Given sample 
$$\mathcal{D} = {\mathbf{x}_1, ..., \mathbf{x}_n}, x_i \in \mathcal{R}^d$$

- compute sample mean:  $\hat{\mu} = \frac{1}{n} \sum_{i} (\mathbf{x}_i)$
- compute sample covariance:  $\hat{\Sigma} = \frac{1}{n} \sum_i (\mathbf{x}_i \hat{\mu}) (\mathbf{x}_i \hat{\mu})^T$
- $\bullet$  compute eigenvalues and eigenvectors of  $\widehat{\Sigma}$

$$\hat{\Sigma} = \Phi \Lambda \Phi^T, \ \Lambda = diag(\sigma_1^2, \dots, \sigma_n^2) \ \Phi^T \Phi = I$$

• order eigenvalues 
$$\sigma_1^2 > ... > \sigma_n^2$$

• if, for a certain k,  $\sigma_k \ll \sigma_1$  eliminate the eigenvalues and eigenvectors above k.

#### Ruta Desai, Chun-Liang Li

#### Lecture 16 - 153

### May 21, 2025

### Alternative PCA Algorithm (testing)

▶ Given principal components  $\phi_i, i \in 1, ..., k$  and a test sample  $\mathcal{T} = \{\mathbf{t}_1, ..., \mathbf{t}_n\}, t_i \in \mathcal{R}^d$ 

- subtract mean to each point  $\mathbf{t}'_i = \mathbf{t}_i \hat{\mu}$
- project onto eigenvector space  $\mathbf{y}_i = \mathbf{A}\mathbf{t}'_i$  where

$$\mathbf{A} = \begin{bmatrix} \phi_1^T \\ \vdots \\ \phi_k^T \end{bmatrix}$$

use \$\mathcal{T}' = {y\_1, ..., y\_n}\$ to estimate class conditional densities and do all further processing on \$\mathcal{y}\$.

### Ruta Desai, Chun-Liang Li

