CS 1674: Intro to Computer Vision Introduction

Prof. Adriana Kovashka University of Pittsburgh January 11, 2022

Course Info

• Course website:

http://people.cs.pitt.edu/~kovashka/cs1674_sp22

- Instructor: Adriana Kovashka (kovashka@cs.pitt.edu)
- Office: Sennott Square 5325
- Zoom: <u>https://pitt.zoom.us/s/4168010698</u>
- **Class:** Tue/Thu, 9:30am-10:45am
- Office hours: Tue/Thu, 11am-12:30pm

About the Instructor





Born 1985 in Sofia, Bulgaria



Got BA in 2008 at Pomona College, CA (Computer Science & Media Studies)



Got PhD in 2014 at University of Texas at Austin (Computer Vision)

About the TA

- TA: Mesut Erhan Unal
- Office: SenSq 5404
- Zoom: TBD
- Office hours: TBD
 - Do this Doodle by the end of Friday: <u>https://doodle.com/poll/yfxdcgix8p9f3499</u>
 - Mark *all* times you're available during *any* week



Course Goals

- To learn the basics of low-level image analysis
- To learn about some classic and modern approaches to high-level computer vision tasks (e.g. object recognition)
- To get experience with vision techniques
- To learn/apply basic machine learning (a key component of modern computer vision)
- To think critically about vision approaches, and to see connections between works

Textbooks

- <u>Computer Vision: Algorithms and Applications</u> by Richard Szeliski, 2nd edition
- Visual Object Recognition by Kristen Grauman and Bastian Leibe
- More resources available on course webpage

• Your notes from class are your best study material, slides are *not* complete with notes

Programming Language

- We'll use Matlab
- It can be downloaded for free from MyPitt -> Software Downloads
- Please download version 2020b or later
- We'll do a short tutorial; ask TA if you need further help

Course Structure

- Lectures
- Homework
 - 9 assignments x 6% each = 54%
- Exams
 - First exam (20%)
 - Second exam (20%)
- Participation (6%)

Policies and Schedule

http://people.cs.pitt.edu/~kovashka/cs1674_sp22

Zoom Etiquette

- Please turn your camera on if you are joining class on Zoom
- Please send me email if you are unable to do so and explain the reason

Health and Well-being

- Follow university policies
- Wear your mask over your mouth and nose
- When talking to instructor or classmates, please keep **6ft distance**
- Stay home if sick, but let instructor know so I can record the lecture and share it with you
- Please reach out to the instructor if you want to talk about anything privately, and **don't hesitate** to ask for help

Warning #1

- This class is a lot of work
- I've opted for shorter, more manageable HW assignments, but there is a lot of them
- I expect you'd be spending 6-8 hours on homework each week

... But you get to understand algorithms and concepts in detail!

Warning #2

- Some parts will be hard and require that you pay close attention!
- Use instructor's and TA's office hours

• ... You will learn a lot!



H/T Kirk Pruhs

Questions?

Plan for Today

- Blitz introductions
- What is computer vision?
 - Why do we care?
 - What are the challenges?
 - What is recent research like?
- Overview of topics
- Review and tutorial
 - Linear algebra
 - Matlab

Blitz introductions (15 sec)

- What is your name?
- What one thing outside of school are you passionate about?
- What do you hope to get out of this class?

 Every time you speak, please remind me your name

Computer Vision

What is computer vision?



Done?

"We see with our brains, not with our eyes" (Oliver Sacks and others)

Kristen Grauman (adapted)

What is computer vision?

- Automatic understanding of images and video
 - Algorithms and representations to allow a machine to recognize objects, people, scenes, and activities
 - Algorithms to mine, search, and interact with visual data
 - Computing properties and navigating within the 3D world using visual data
 - Generating realistic synthetic visual data

Perception and interpretation



Kristen Grauman

Visual search, organization



Measurement

Structure from motion

Real-time stereo



input sequence Relating images feature matches Structure & Motion recovery 3D features and cameras **Dense Matching** dense depth maps 3D Model Building 3D surface model

Multi-view stereo for community photo collections



Goesele et al.



Pollefeys et al.

Generation



Karras et al., "Progressive Growing of GANs for Improved Quality, Stability, and Variation", ICLR 2018

Related disciplines



Vision and graphics



Inverse problems: analysis and synthesis.

Kristen Grauman

Why vision?

144k hours uploaded to YouTube daily4.5 mil photos uploaded to Flickr daily10 bil images indexed by Google

Images and video are everywhere!



Personal photo albums

Movies, news, sports







Surveillance and security Adapted from Lana Lazebnik



Medical and scientific images

Why vision?

- As image sources multiply, so do applications
 - Relieve humans of boring, easy tasks
 - Perception for robotics / autonomous agents
 - Organize and give access to visual content
 - Description of content for the visually impaired
 - Human-computer interaction
 - Fun applications (e.g. art styles to my photos)

Things that work well

Faces and digital cameras



Camera waits for everyone to smile to take a photo [Canon]



Setting camera focus via face detection

Face recognition



Optical character recognition



https://www.mathworks.com/matlabcentral/fileexchange/18169-optical-character-recognition-ocr

Accurate object detection in real time

	Pascal 2007 mAP	Speed	
DPM v5	33.7	.07 FPS	14 s/img
R-CNN	66.0	.05 FPS	20 s/img
Fast R-CNN	70.0	.5 FPS	2 s/img
Faster R-CNN	73.2	7 FPS	140 ms/img
YOLO	69.0	45 FPS	22 ms/img



Redmon et al., "You Only Look Once: Unified, Real-Time Object Detection", CVPR 2016



Our ability to detect objects has gone from 34 mAP in 2008 to 73 mAP at 7 FPS (frames per second) or 63 mAP at 45 FPS in 2016



YOLO: Real-Time Object Detection



You only look once (YOLO) is a system for detecting objects on the Pascal VOC 2012 dataset. It can detect the 20 Pascal object classes:

- person
- bird, cat, cow, dog, horse, sheep
- aeroplane, bicycle, boat, bus, car, motorbike, train
- bottle, chair, dining table, potted plant, sofa, tv/monitor

Exploring photo collections

Photo Tourism

Exploring photo collections in 3D

Microsoft







Snavely et al.

Linking to info with a mobile device



Yeh et al., MIT



MSR Lincoln



e.	Casino F	toyale	
Cine	eman: Rev	iews, Tra	iler
🐌 Film	blog.ch		
🔌 Ami	azon Mobil	9	
A Eba	y Mobile		
MSI	N Mobile N	lovies	
Goo	gle Mobile		
T Call	Kitag for	Ticket	
B Tell	a friend (by SMS)	
1 Hon	ne		
Search	for anoth	er movie	title
on aur	movie po	rtal:	
60317	h 1		

kooaba
Transferring art styles



















DeepArt.io – try it for yourself!

Image generation (faces)



Interactive systems



Shotton et al.



Video-based interfaces



Human joystick NewsBreaker Live YouTube Link



Assistive technology systems Camera Mouse Boston College

Vision for medical & neuroimages



Image guided surgery MIT AI Vision Group



fMRI data Golland et al.



Safety & security



Navigation, driver safety



Pedestrian detection MERL, Viola et al.



Monitoring pool (Poseidon)



Surveillance

Kristen Grauman

Things that need more work

From recent CVPR, ICCV, ECCV

CVPR = IEEE/CVF Conference on Computer Vision and Pattern Recognition ICCV = IEEE/CVF International Conference on Computer Vision ECCV = European Conference on Computer Vision

Recognition in novel modalities



Figure 6: Qualitative Results. YOLO running on sample artwork and natural images from the internet. It is mostly accurate although it does think one person is an airplane.

Redmon et al., CVPR 2016

Learning from weak supervision



Chen et al., CVPR 2017

Understanding activities and intents



Understanding how well activity performed



Pirsiavash et al., ECCV 2014

Imagining motion in static images



Decoding physics from video



Video generation













Understanding stories in film



Figure 1: Our MovieQA dataset contains 14,944 questions about 408 movies. It contains multiple sources of information: plots, subtitles, video clips, scripts, and DVS transcriptions. In this figure we show example QAs from *The Matrix* and localize them in the timeline.

Understanding stories in film





Automatic Understanding of Image and Video Advertisements



Zaeem Hussain, Mingda Zhang, Xiaozhong Zhang, Keren Ye, Christopher Thomas, Zuha Agha, Nathan Ong, Adriana Kovashka University of Pittsburgh



Understanding advertisements is more challenging than simply recognizing physical content from images, as ads employ a variety of strategies to persuade viewers.







We collect an advertisement dataset containing 64,832 images and 3,477 videos, each annotated by 3-5 human workers from Amazon Mechanical Turk.

Image	Topic	204,340	Strategy	20,000
	Sentiment	102,340	Symbol	64,131
	Q+A Pair	202,090	Slogan	11,130
Video	Торіс	17,345	Fun/Exciting	15,380
	Sentiment	17,345	English?	17,374
	Q+A Pair	17,345	Effective	16,721

Here are some sample annotations in our dataset.



New Caddy Maxi Life, infinitely bigge

What's being advertised in this image?

Cars, automobiles

What sentiments are provoked in the viewer?

Amused, Creative, Impressed, Youthful, Conscious

What strategies are used to persuade viewer?

Symbolism, Contrast, Straightforward, Transferred qualities

What should the viewer do, and why should they do this?

- I should buy Volkswagen because it can hold a big bear.
- I should buy VW SUV because it can fit anything and everything in it.
- I should buy this car because it can hold everything I need.

More information available at http://cs.pitt.edu/~kovashka/ads

Reasoning and acting: Embodied question answering



Computer vision is not solved

- Deep learning makes excellent use of massive data (labeled for the task of interest?)
 - But it's hard to understand how it does so, makes it hard to fix when it doesn't work well
 - It doesn't work well when massive data is not available and your task is different than tasks for which data is available
 - We can recognize objects with 97% accuracy but reasoning about relationships and intent is harder

Seeing Al

YouTube link



Microsoft Cognitive Services: Introducing the Seeing AI project

Obstacles?

MASSACHUSETTS INSTITUTE OF TECHNOLOGY

PROJECT MAC

Artificial Intelligence Group Vision Memo. No. 100. July 7, 1966

THE SUMMER VISION PROJECT

Seymour Papert

The summer vision project is an attempt to use our summer workers effectively in the construction of a significant part of a visual system. The particular task was chosen partly because it can be segmented into sub-problems which will allow individuals to work independently and yet participate in the construction of a system complex enough to be a real landmark in the development of "pattern recognition".

Read more about the history: Szeliski Sec. 1.2

Why is vision difficult?

- Ill-posed problem: real world much more complex than what we can measure in images
 - $-3D \rightarrow 2D$
 - Motion \rightarrow static
- Impossible to literally "invert" image formation process with limited information
 - Need information outside of this particular image to generalize what image portrays (e.g. to resolve occlusion)

What the computer gets







146 151 164 172 141 154 170 123 127 135 148 154 162 165 170 171 160 183 201 210 123 130 132 138 150 157 158 174 182 189 186 198 221 224 221 125 127 126 129 130 135 139 141 150 165 175 172 185 195 207 210 212 226 229 222 224 135 143 143 143 144 146 145 147 160 174 184 191 199 207 211 213 217 224 227 223 223 221 221 218 224 223 140 147 146 149 157 162 160 159 165 174 181 198 201 210 212 216 223 224 225 225 220 215 217 215

Why is this problematic?

Adapted from Kristen Grauman and Lana Lazebnik

Challenges: many nuisance parameters



Illumination



Occlusions



Object pose



Intra-class appearance



Clutter



Viewpoint

Think again about the pixels...

Challenges: intra-class variation



slide credit: Fei-Fei, Fergus & Torralba

Challenges: importance of context





slide credit: Fei-Fei, Fergus & Torralba

Challenges: Complexity

- Thousands to millions of pixels in an image
- 3,000-30,000 human recognizable object categories
- 30+ degrees of freedom in the pose of articulated objects (humans)
- Billions of images indexed by Google Image Search
- 1.424 billion smart camera phones sold in 2015
- About half of the cerebral cortex in primates is devoted to processing visual information [Felleman and van Essen 1991]

Challenges: Limited supervision



Less





More







Kristen Grauman

Challenges: Vision requires reasoning



What color are her eyes? What is the mustache made of?



How many slices of pizza are there? Is this a vegetarian pizza?



Is this person expecting company? What is just under the tree?



Does it appear to be rainy? Does this person have 20/20 vision?

Evolution of datasets

• Challenging problem \rightarrow active research area





PASCAL: 20 categories, 12k images



ImageNet: 22k categories, 14mil images



Microsoft COCO: 80 categories, 300k images

Some Visual Recognition Problems: Why are they challenging?



Recognition: What objects do you see?



Detection: Where are the cars?



Activity: What is this person doing?



Scene: Is this an indoor scene?



Instance: Which city? Which building?


Visual question answering: Why is there a carriage in the street?



What to do about variance? (albino koalas)

- You can't always do anything about it
- Basic assumption: training and test set are sampled from the same distribution
- Practically speaking: Make sure training set is representative
- Tricks (e.g. data preprocessing) make learning stable
- You can try to artificially expand variance in training set
- You can detect outliers, report confidence of prediction
- Methods exist that explicitly try to handle out-ofdistribution data, but it's an open problem

Problem with categorization (Borges' Animals)

"These ambiguities, redundancies and deficiencies recall those that Dr. Franz Kuhn attributes to a certain Chinese dictionary entitled The Celestial Emporium of Benevolent *Knowledge*. In its remote pages it is written that animals can be divided into (a) those belonging to the Emperor, (b) those that are embalmed, (c) those that are tame, (d) pigs, (e) sirens, (f) imaginary animals, (g) wild dogs, (h) those included in this classification, (i) those that are crazy-acting, (j) those that are uncountable, (k) those painted with the finest brush made of camel hair, (I) miscellaneous, (m) those which have just broken a vase, and (n) those which, from a distance, look like flies."

Overview of topics

Features and filters

512







Kristen Grauman

Features and filters

- Detecting distinctive and repeatable features
- Describing images with local statistics







Grouping and fitting



Parallelism



Symmetry



Continuity

Closure







[fig from Shi et al]

 Clustering, segmentation, fitting; what parts belong together?







Multiple views







Hartley and Zisserman



 Multi-view geometry, matching, stereo vision





Kristen Grauman

Image categorization

• Fine-grained recognition



Visipedia Project

Image categorization

• Material recognition





[Bell et al. CVPR 2015]

Image categorization

Image style recognition



HDR



Vintage



Macro



Noir



Minimal



Long Exposure



Hazy



Romantic

Flickr Style: 80K images covering 20 styles.



Baroque



Northern Renaissance



Impressionism



Abs. Expressionism



Roccoco



Cubism



Post-Impressionism



Color Field Painting

Wikipaintings: 85K images for 25 art genres.

[Karayev et al. BMVC 2014]

Visual recognition and SVMs











 Recognizing objects and categories, learning techniques

Convolutional neural networks (CNNs)

State-of-the-art on many recognition tasks



Recurrent neural networks (RNNs)

• Sequence processing, e.g. question answering



Wu et al., CVPR 2016

Motion and tracking

• Tracking objects, video analysis





Pose and actions

- Automatically annotating human pose (joints)
- Recognizing actions in first-person video





Unsupervised learning



Unsupervised learning







Doersch et al., ICCV 2015

Linear algebra review

See <u>http://cs229.stanford.edu/section/cs229-linalg.pdf</u> for more

What are images? (in Matlab)

- Matlab treats images as matrices of numbers
- To proceed, let's talk very briefly about how images are formed



Digital images





a b

FIGURE 2.17 (a) Continuos image projected onto a sensor array. (b) Result of image sampling and quantization.

- Sample the 2D space on a regular grid
- Quantize each sample (round to nearest integer)

Digital images

- Sample the 2D space on a regular grid
- **Quantize** each sample (round to nearest integer)
- What does quantizing signal look like?



2

• Image thus represented as a matrix of integer values.



62	79	23	119	120	105	4	0
10	10	9	62	12	78	34	0
10	58	197	46	46	0	0	48
176	135	5	188	191	68	0	49
2	1	1	29	26	37	0	77
0	89	144	147	187	102	62	208
255	252	0	166	123	62	0	31
166	63	127	17	1	0	99	30

Adapted from S. Seitz

2D

Digital color images



@ 2000 How Stuff Works

Slide credit: Kristen Grauman

Digital color images

Color images, RGB color space:

Split image into three channels









B Adapted from Kristen Grauman

Images in Matlab

- Color images represented as a matrix with multiple channels (=1 if grayscale) ٠
- Suppose we have a NxM RGB image called "im" ٠
 - im(1,1,1) = top-left pixel value in R-channel
 - im(y, x, b) = y pixels **down**, x pixels **to right** in the bth channel
 - im(N, M, 3) = bottom-right pixel in B-channel
- imread(filename) returns a uint8 image (values 0 to 255) ٠
 - Convert to double format with double or im2double

	column															
row								ī		•		R				
	0.92	0.93	0.94	0.97	0.62	0.37	0.85	0.97	0.93	0.92	0.99					
1	0.95	0.89	0.82	0.89	0.56	0.31	0.75	0.92	0.81	0.95	0.91			0		
	0.89	0.72	0.51	0.55	0.51	0.42	0.57	0.41	0.49	0.91	0.92	0 92	مم	l G		
	0.96	0.95	0.88	0.94	0.56	0.46	0.91	0.87	0.90	0.97	0.95	0.52	0.55			
	0.71	0.81	0.81	0.87	0.57	0.37	0.80	0.88	0.89	0.79	0.85	0.93	0.91			В
	0.49	0.62	0.60	0.58	0.50	0.60	0.58	0.50	0.61	0.45	0.33	0.91	0.92	0.92	0.99	
	0.86	0.84	0.74	0.58	0.51	0.39	0.73	0.92	0.91	0.49	0.74	0.97	0.95	0.95	0.91	
	0.96	0.67	0.54	0.85	0.48	0.37	0.88	0.90	0.94	0.82	0.93	0.79	0.85	0.91	0.92	
	0.69	0.49	0.56	0.66	0.43	0.42	0.77	0.73	0.71	0.90	0.99	0.45	0.33	0.97	0.95	
	0.79	0.73	0.90	0.67	0.33	0.61	0.69	0.79	0.73	0.93	0.97	0.49	0.74	0.79	0.85	
V	0.91	0.94	0.89	0.49	0.41	0.78	0.78	0.77	0.89	0.99	0.93	0.82	0.93	0.45	0.33	
	0131	0.5	0.05	0.45	0.50	0.00	0.45	0.72	0.77	0.75	0.7 1	0.90	0.99	0.49	0.74	
			0.79	0.73	0.90	0.67	0.33	0.61	0.69	0.79	0.73	0.93	0.97	0.82	0.93	
			0.91	0.94	0.89	0.49	0.41	0.78	0.78	0.77	0.89	0.99	0.93	0.90	0.99	
					0.79	0.73	0.90	0.67	0.33	0.61	0.69	0.79	0.73	0.93	0.97	
					0.01	0.75	0.90	0.07	0.33	0.78	0.78	0.75	0.75	0.95	0.97	
					0.91	0.94	0.09	0.49	0.41	0.78	0.78	0.77	Ac	lapted	from D	l erek Hoiem

Vectors and Matrices

- Vectors and matrices are just collections of ordered numbers that represent something: movements in space, scaling factors, word counts, movie ratings, pixel brightnesses, etc.
- We'll define some common uses and standard operations on them.

Vector

• A column vector $\mathbf{v} \in \mathbb{R}^{n imes 1}$ where

$$\mathbf{v} = \begin{bmatrix} v_1 \\ v_2 \\ \vdots \\ v_n \end{bmatrix}$$

• A row vector $\mathbf{v}^T \in \mathbb{R}^{1 \times n}$ where

$$\mathbf{v}^T = egin{bmatrix} v_1 & v_2 & \dots & v_n \end{bmatrix}$$
 T denotes the transpose operation

Vector

- You'll want to keep track of the orientation of your vectors when programming in MATLAB.
- You can transpose a vector V in MATLAB by writing V'.

Vectors have two main uses



- Vectors can represent an offset in 2D or 3D space
- Points are just vectors from the origin

- Data can also be treated as a vector
- Such vectors don't have a geometric interpretation, but calculations like "distance" still have value

Norms

• L1 norm

$$\left\|oldsymbol{x}
ight\|_{1}:=\sum_{i=1}^{n}\left|x_{i}
ight|$$

• L2 norm

$$\|oldsymbol{x}\|:=\sqrt{x_1^2+\cdots+x_n^2}$$

• L^p norm (for real numbers $p \ge 1$)

$$\left\|\mathbf{x}
ight\|_p := igg(\sum_{i=1}^n |x_i|^pigg)^{1/p}$$

Distances

• L1 (Manhattan) distance

$$d_1(\mathbf{p},\mathbf{q}) = \|\mathbf{p}-\mathbf{q}\|_1 = \sum_{i=1}^n |p_i-q_i|,$$

• L2 (Euclidean) distance

$$d(\mathbf{p},\mathbf{q})=\sqrt{\sum_{i=1}^n(q_i-p_i)^2}$$

$$d(p,q) = \sqrt{(p_1-q_1)^2 + (p_2-q_2)^2 + \dots + (p_i-q_i)^2 + \dots + (p_n-q_n)^2}$$

Example: Feature representation

- A vector representing measurable characteristics of a data sample we have
- E.g. a glass of juice can be represented via its color = {yellow=1, red=2, green=3, purple=4} and taste = {sweet=1, sour=2}
- A given glass *i* can be represented as a vector: *x_i* = [3 2] represents *sour, green* juice
- For D features, this defines a D-dimensional space where we can measure similarity between samples

Example: Feature representation



Matrix

• A matrix $A \in \mathbb{R}^{m \times n}$ is an array of numbers with size $m \downarrow$ by $n \rightarrow$, i.e. m rows and n columns.

$$\mathbf{A} = \begin{bmatrix} a_{11} & a_{12} & a_{13} & \dots & a_{1n} \\ a_{21} & a_{22} & a_{23} & \dots & a_{2n} \\ \vdots & & & \vdots \\ a_{m1} & a_{m2} & a_{m3} & \dots & a_{mn} \end{bmatrix}$$

• If m = n, we say that \mathbf{A} is square.

Matrix Operations

• Addition

$$\begin{bmatrix} a & b \\ c & d \end{bmatrix} + \begin{bmatrix} 1 & 2 \\ 3 & 4 \end{bmatrix} = \begin{bmatrix} a+1 & b+2 \\ c+3 & d+4 \end{bmatrix}$$

 Can only add a matrix with matching dimensions, or a scalar.

$$\begin{bmatrix} a & b \\ c & d \end{bmatrix} + 7 = \begin{bmatrix} a+7 & b+7 \\ c+7 & d+7 \end{bmatrix}$$

Scaling

$$\begin{bmatrix} a & b \\ c & d \end{bmatrix} \times 3 = \begin{bmatrix} 3a & 3b \\ 3c & 3d \end{bmatrix}$$
Matrix Multiplication

- Let X be an *axb* matrix, Y be an *bxc* matrix
- Then Z = X*Y is an *a*xc matrix
- Second dimension of first matrix, and first dimension of second matrix have to be the same, for matrix multiplication to be possible
- Practice: Let X be an 10x5 matrix. Let's factorize it into 3 matrices...

Matrix Multiplication



• Each entry in the result is (that row of A) dot product with (that column of B)

Matrix Multiplication

• Example:



 Each entry of the matrix product is made by taking the dot product of the corresponding row in the left matrix, with the corresponding column in the right one.

Inner Product

 Multiply corresponding entries of two vectors and add up the result

$$\mathbf{x}^T \mathbf{y} = \begin{bmatrix} x_1 & \dots & x_n \end{bmatrix} \begin{bmatrix} y_1 \\ \vdots \\ y_n \end{bmatrix} = \sum_{i=1}^n x_i y_i \quad (\text{scalar})$$

- x·y is also |x||y|Cos(angle between x and y)
- If B is a unit vector, then A·B gives the length of A which lies in the direction of B (projection)

(if B is unit-length hence norm is 1)

Different Types of Product

- *x*, *y* = column vectors (nx1)
- X, Y = matrices (mxn)
- *x*, *y* = scalars (1x1)
- $\mathbf{x}^T \mathbf{y} = \mathbf{x} \cdot \mathbf{y} = \text{inner product (1xn x nx1 = scalar)}$
- $\mathbf{x} \otimes \mathbf{y} = \mathbf{x} \mathbf{y}^T$ = outer product (nx1 x 1xn = matrix)
- **X** * **Y** = matrix product
- X.* Y = element-wise product

Matrix Operations

Transpose – flip matrix, so row 1 becomes column 1

$$\begin{bmatrix} 0 & 1 \\ 2 & 3 \\ 4 & 5 \end{bmatrix}^T = \begin{bmatrix} 0 & 2 & 4 \\ 1 & 3 & 5 \end{bmatrix}$$

• A useful identity:

$$(ABC)^T = C^T B^T A^T$$

Matrix Operations

• MATLAB example:

$$AX = B$$
$$A = \begin{bmatrix} 2 & 2 \\ 3 & 4 \end{bmatrix}, B = \begin{bmatrix} 1 \\ 1 \end{bmatrix}$$
$$>> \mathbf{x} = \mathbf{A} \setminus \mathbf{B}$$
$$\mathbf{x} = \mathbf{1.0000}$$
$$-0.5000$$

Matrix Operation Properties

Matrix addition is commutative and associative

$$-A+B = B+A$$

$$-A + (B + C) = (A + B) + C$$

Matrix multiplication is associative and distributive but *not* commutative

$$-A(B^*C) = (A^*B)C$$

$$-A(B+C) = A^*B + A^*C$$

− A*B != B*A

Special Matrices

- Identity matrix I
 - Square matrix, 1's along diagonal, 0's elsewhere
 - I [another matrix] = [that matrix]

$$\begin{bmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{bmatrix}$$

- Diagonal matrix
 - Square matrix with numbers along diagonal, 0's elsewhere
 - A diagonal [another matrix]
 scales the rows of that matrix

$$\begin{bmatrix} 3 & 0 & 0 \\ 0 & 7 & 0 \\ 0 & 0 & 2.5 \end{bmatrix}$$

Special Matrices

Symmetric matrix

$$\mathbf{A}^T = \mathbf{A}$$

 $\begin{bmatrix} 1 & 2 & 5 \\ 2 & 1 & 7 \\ 5 & 7 & 1 \end{bmatrix}$

Matlab

Matlab tutorial

http://www.cs.pitt.edu/~kovashka/cs1674_fa18/tutorial.m http://www.cs.pitt.edu/~kovashka/cs1674_fa18/myfunction.m http://www.cs.pitt.edu/~kovashka/cs1674_fa18/myotherfunction.m

Please cover whatever we don't finish at home.

Other tutorials and exercises

- <u>https://people.cs.pitt.edu/~milos/courses/cs2</u>
 <u>750/Tutorial/</u>
- <u>http://www.math.udel.edu/~braun/M349/Ma</u>
 <u>tlab_probs2.pdf</u>
- <u>http://www.facstaff.bucknell.edu/maneval/he</u> <u>lp211/basicexercises.html</u>
 - Do Problems 1-8, 12
 - Most also have solutions
 - Ask the TA if you have any problems