

Lecture 5: Math Review I

Administrative

HW0 due Wednesday 1/29
(1 week from yesterday)

HW1 out yesterday,
due Wednesday 2/5
(3 weeks from yesterday)

Floating Point Arithmetic

This Lecture and Next: Math

Two goals for the next two classes:

- Math with computers \neq Math
- Practical math you need to know but may not have been taught



~~MSA MATHEMATICS
UNIVERSITY OF MICHIGAN~~

This Lecture and Next: Goal


- Not a “Linear algebra in two lectures” – that’s impossible.
- *Some of this you should know!*
- Aimed at reviving your knowledge and plugging any gaps
- Aimed at giving you intuitions

Adding Numbers

- **$1 + 1 = ?$**
- Suppose x_i is normally distributed with mean μ and standard deviation σ for $1 \leq i \leq N$
- **How is the average, or $\hat{\mu} = \frac{1}{N} \sum_{i=1}^N x_i$, distributed (qualitatively), in terms of variance?**
- *The Free Drinks in Vegas Theorem: $\hat{\mu}$ has mean μ and standard deviation $\frac{\sigma}{\sqrt{N}}$.*

Free Drinks in Vegas

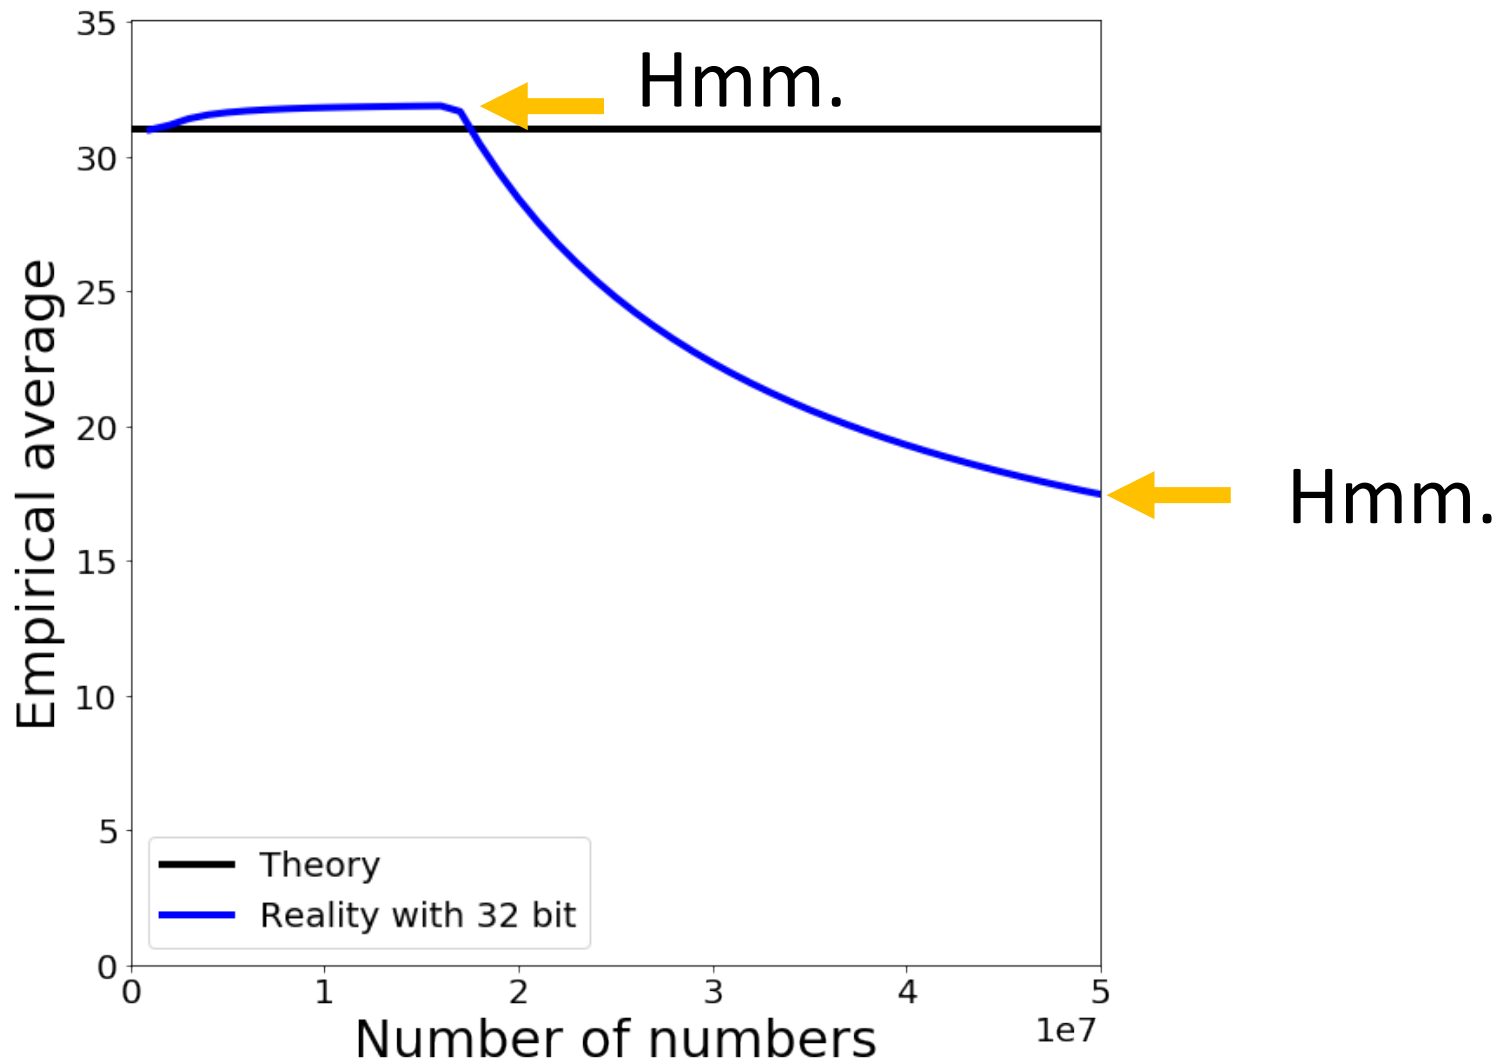
Each game/variable has mean \$0.10, std \$2



Let's make it big

- **What should happen qualitatively?**
- Theory says that the average is distributed with mean 31 and standard deviation $\frac{1}{\sqrt{50M}} \approx (10^{-5})$
- **What will happen?**
- Reality: 17.47



Trying it out



What is a number?

2^7	2^6	2^5	2^4	2^3	2^2	2^1	2^0	
1	0	1	1	1	0	0	1	185
$128 + 32 + 16 + 8 + 1 =$								185

Adding two numbers

	2^8	2^7	2^6	2^5	2^4	2^3	2^2	2^1	2^0	
		1	0	1	1	1	0	0	1	185
	+	0	1	1	0	1	0	0	1	105
<hr/>										
1		0	0	1	0	0	0	1	0	34
										
Carry Flag	Result									

“Integers” on a computer are integers modulo 2^k

Some Gotchas

$$32 + (3 / 4) \times 40 = 32$$

$$32 + (3 \times 40) / 4 = 62$$

Why?

Underflow

No Underflow

$$32 + (3 / 4) \times 40 =$$

$$32 + 0 \quad \times 40 =$$

$$32 + 0 \quad =$$

$$32$$

$$32 + (3 \times 40) / 4 =$$

$$32 + 120 \quad / 4 =$$

$$32 + 30 \quad =$$

$$62$$

Ok – you have to multiply before dividing

Some Gotchas

math $32 + (9 \times 40) / 10 = 68$

**Should be:
 $9 \times 4 = 36$**

uint8 $32 + (9 \times 40) / 10 = 42$

Overflow

$32 + 9 \times 40 / 10 =$

$32 + 104 / 10 =$

$32 + 10 =$

42

Why 104?

$9 \times 40 = 360$

$360 \% 256 = 104$

What is a number?

2^7	2^6	2^5	2^4	2^3	2^2	2^1	2^0	
1	0	1	1	1	0	0	1	185

How can we do fractions?

2^5	2^4	2^3	2^2	2^1	2^0	2^{-1}	2^{-2}	
1	0	1	1	1	0	0	1	45.25

45 0.25

Fixed-Point Arithmetic

$$\begin{array}{cccccccc} 2^5 & 2^4 & 2^3 & 2^2 & 2^1 & 2^0 & 2^{-1} & 2^{-2} \\ 1 & 0 & 1 & 1 & 1 & 0 & 0 & 1 \end{array} \quad 45.25$$

What's the largest number we can represent?

63.75 – Why?

How precisely can we measure at 63?

0.25

How precisely can we measure at 0?

0.25

Fine for many purposes but for science, seems silly

Floating Point Numbers

Sign (S)

Exponent (E)

Fraction (F)

1

0 1 1 1

0 0 1

1

7

1

-1

$2^{7-7} = 2^0 = 1$

$1 + 1/8 = 1.125$

$$(-1^S) (2^{E+bias}) \left(1 + \frac{F}{2^3} \right)$$

Bias: allows exponent to be negative (bias = -127 for float32)

Note: fraction = significant = mantissa;

exponents of all ones or all zeros are special numbers

Floating Point Numbers

Sign	Exponent	Fraction	
1	0 1 1 1	0 0 0	$-2^0 \times 1.00 = -1$
		0 0 1	$-2^0 \times 1.125 = -1.125$
		0 1 0	$-2^0 \times 1.25 = -1.25$
		...	
		1 1 0	$-2^0 \times 1.75 = -1.75$
		1 1 1	$-2^0 \times 1.875 = -1.875$

-1

$7-7=0$
 $*(-bias)*$

Floating Point Numbers

Sign	Exponent	Fraction	
1	1 0 0 1	0 0 0	$-2^2 \times 1.00 = -4$
		0 0 1	$-2^2 \times 1.125 = -4.5$
		0 1 0	$-2^2 \times 1.25 = -5$
		...	
		1 1 0	$-2^2 \times 1.75 = -7$
		1 1 1	$-2^2 \times 1.875 = -7.5$

-1

9-7=2
(-bias)

Floating Point Numbers

Sign	Exponent	Fraction	
1	0 1 1 1	0 0 0	$-2^0 \times 1.00 = -1$
		0 0 1	$-2^0 \times 1.125 = -1.125$
1	1 0 0 1	0 0 0	$-2^2 \times 1.00 = -4$
		0 0 1	$-2^2 \times 1.125 = -4.5$

Gap between numbers is *relative*, not absolute

Adding Floating Point Numbers

	Sign	Exponent	Fraction	
	1	0 1 1 0	0 0 0	$-2^{-1} \times 1.00 = -0.5$
+	1	1 0 0 1	0 0 0	$-2^2 \times 1.00 = -4$
<hr/>				
	1	1 0 0 1	0 0 1	$-2^2 \times 1.125 = -4.5$


Actual implementation is complex

Adding Floating Point Numbers

	Sign	Exponent	Fraction	
	1	0 1 0 0	0 0 0	$-2^{-3} \times 1.00 = -0.125$
+	1	1 0 0 1	0 0 0	$-2^2 \times 1.00 = -4$

$-2^2 \times 1.03125 = -4.125$

1	1 0 0 1	0 0 0	$-2^2 \times 1.00 = -4$
1	1 0 0 1	0 0 1	$-2^2 \times 1.125 = -4.5$

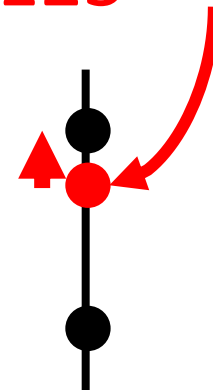


Adding Floating Point Numbers

	Sign	Exponent	Fraction	
	1	0 1 0 0	0 0 0	$-2^{-3} \times 1.00 = -0.125$
+	1	1 0 0 1	0 0 0	$-2^2 \times 1.00 = -4$

$-2^2 \times 1.03125 = -4.125$

1	1 0 0 1	0 0 0	$-2^2 \times 1.00 = -4$
---	---------	-------	-------------------------



For a and b, these can happen

$$a + b = a \quad a + b - a \neq b$$

Real Floating Point Numbers

IEEE 754 Single Precision / Single / float32

8 bits

$$2^{127} \approx 10^{38}$$

23 bits

≈ 7 decimal digits



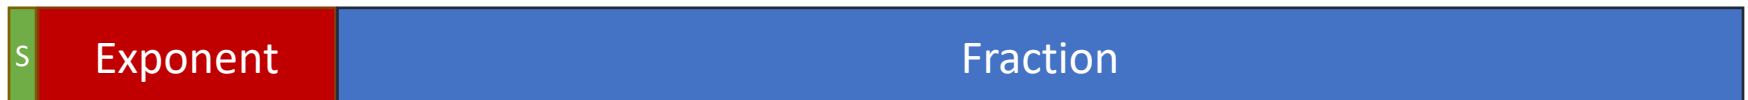
IEEE 754 Double Precision / Double / float64

11 bits

$$2^{1023} \approx 10^{308}$$

52 bits

≈ 15 decimal digits



Real Floating Point Numbers

IEEE 754 Half Precision / Half / float16

5 bits
 $2^{32} \approx 10^9$

10 bits
 ≈ 3 decimal digits



Brain Floating Point / bfloat16


8 bits
 $2^{127} \approx 10^{38}$

7 bits
 ≈ 2 decimal digits



Same range as FP32, but reduced precision

Trying it out




$a+b=a \rightarrow$
numerator is
stuck,
denominator
isn't

Things to Remember

- Computer numbers aren't math numbers
- Overflow, accidental zeros, roundoff error, and basic equalities are almost certainly incorrect for some values
- Floating point defaults and numpy try to protect you.
- Generally safe to use a double and use built-in-functions in numpy (not necessarily others!)
- Spooky behavior = look for numerical issues

Vectors

Vectors



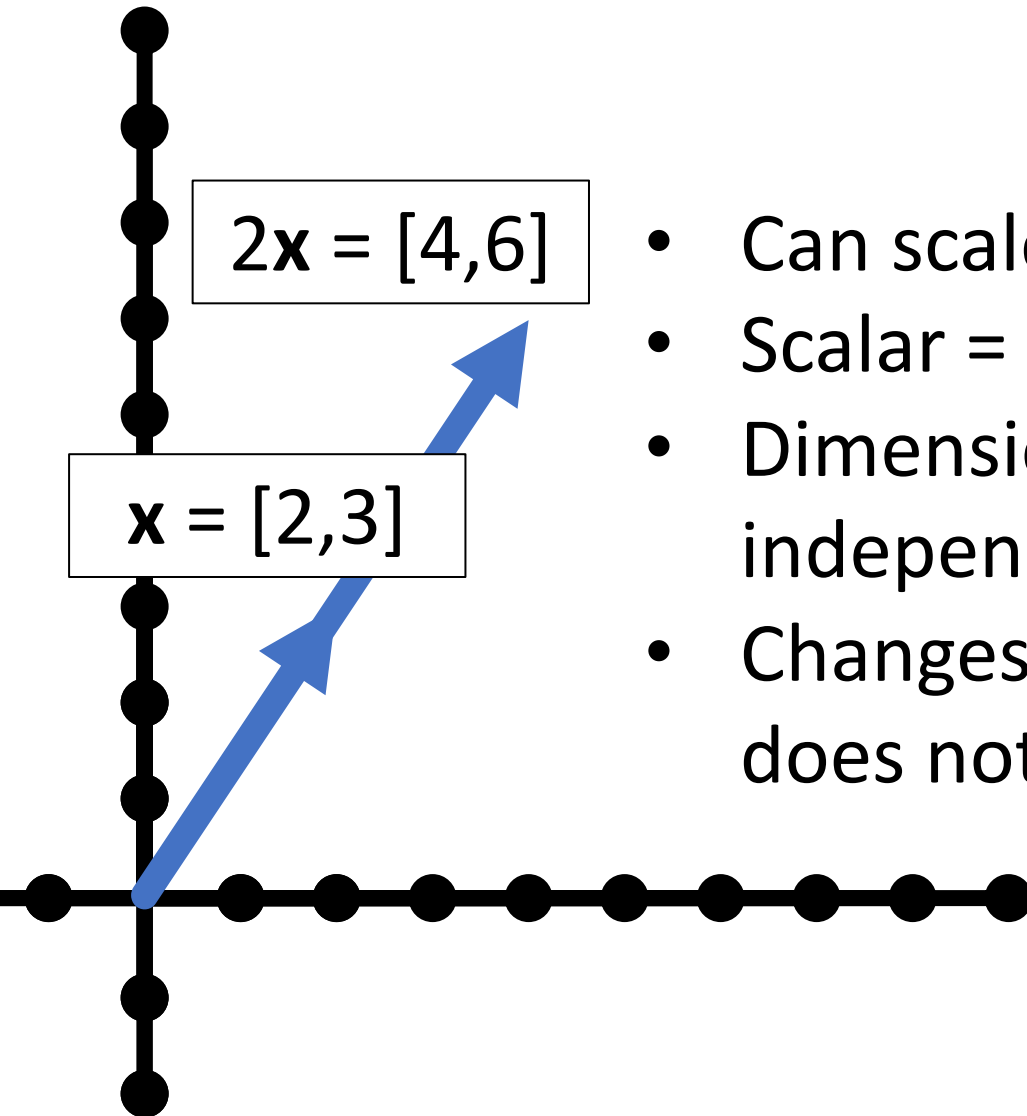
$$[2,3] = 2 \times [1,0] + 3 \times [0,1]$$

$$2 \times \rightarrow + 3 \times \uparrow$$

$$2 \times e_1 + 3 \times e_2$$

Can be arbitrary # of
dimensions
(typically denoted \mathbb{R}^n)


Scaling Vectors




- Can scale vector by a *scalar*
- Scalar = single number
- Dimensions changed independently
- Changes *magnitude / length*, does not change *direction*.

Adding Vectors

- Can add vectors
- Dimensions changed independently
- Order irrelevant
- Can change direction and magnitude



Scaling and Adding



Measuring Length


Magnitude / length / (L2) norm of vector

$$\|\mathbf{x}\| = \|\mathbf{x}\|_2 = \left(\sum_i^n x_i^2 \right)^{1/2}$$

There are other norms; assume L2 unless told otherwise

$$\|\mathbf{x}\|_2 = \sqrt{13}$$
$$\|\mathbf{y}\|_2 = \sqrt{10}$$

Why?



Normalizing a Vector


$$\mathbf{x} = [2, 3]$$

Dividing by norm gives something on the *unit sphere* (all vectors with length 1)

$$\mathbf{x}' = \mathbf{x} / \|\mathbf{x}\|_2$$

$$\mathbf{y} = [3, 1]$$

$$\mathbf{y}' = \mathbf{y} / \|\mathbf{y}\|_2$$

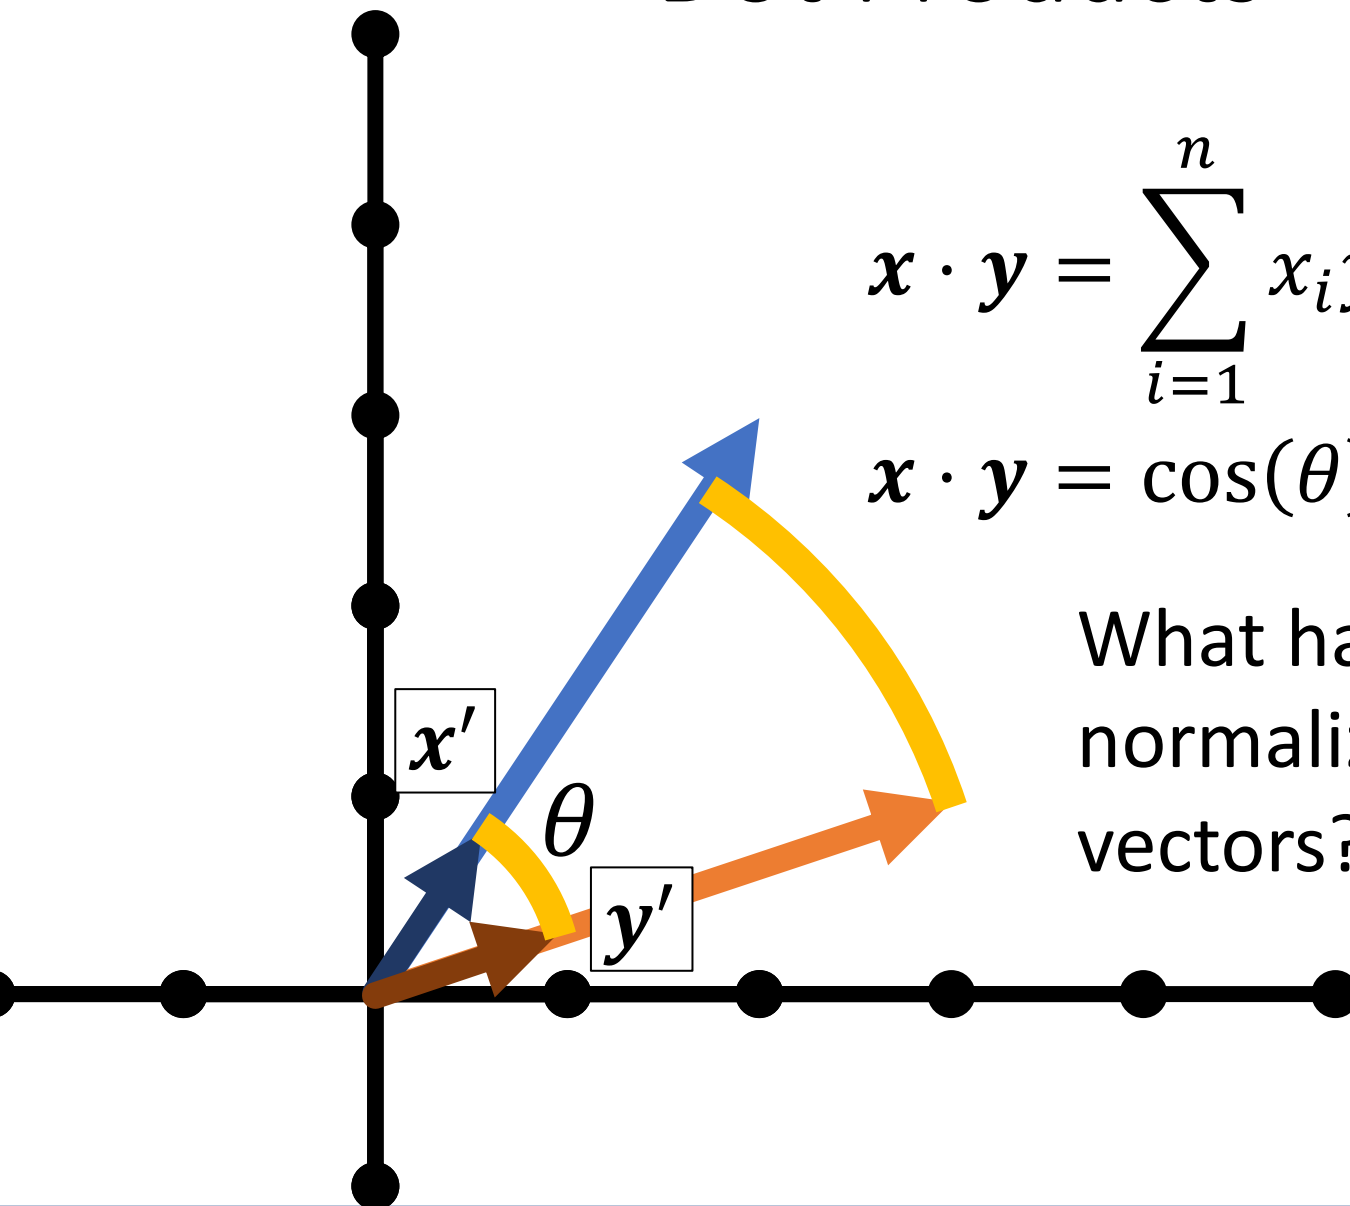


Dot Products

$$\mathbf{x} \cdot \mathbf{y} = \sum_{i=1}^n x_i y_i = \mathbf{x}^T \mathbf{y}$$

$$\mathbf{x} \cdot \mathbf{y} = \cos(\theta) \|\mathbf{x}\| \|\mathbf{y}\|$$

What happens with
normalized / unit
vectors?




Dot Products

$$\mathbf{x} = [2, 3] \quad \mathbf{x} \cdot \mathbf{y} = \sum_i^n x_i y_i$$

What's $\mathbf{x} \cdot \mathbf{e}_1$, $\mathbf{x} \cdot \mathbf{e}_2$?

Ans: $\mathbf{x} \cdot \mathbf{e}_1 = 2$; $\mathbf{x} \cdot \mathbf{e}_2 = 3$

- Dot product is projection
- Amount of \mathbf{x} that's also pointing in direction of \mathbf{y}




Dot Products

$$\mathbf{x} = [2,3] \quad \mathbf{x} \cdot \mathbf{y} = \sum_i^n x_i y_i$$

What's $\mathbf{x} \cdot \mathbf{x}$?


Ans: $\mathbf{x} \cdot \mathbf{x} = \sum x_i x_i = \|\mathbf{x}\|_2^2$



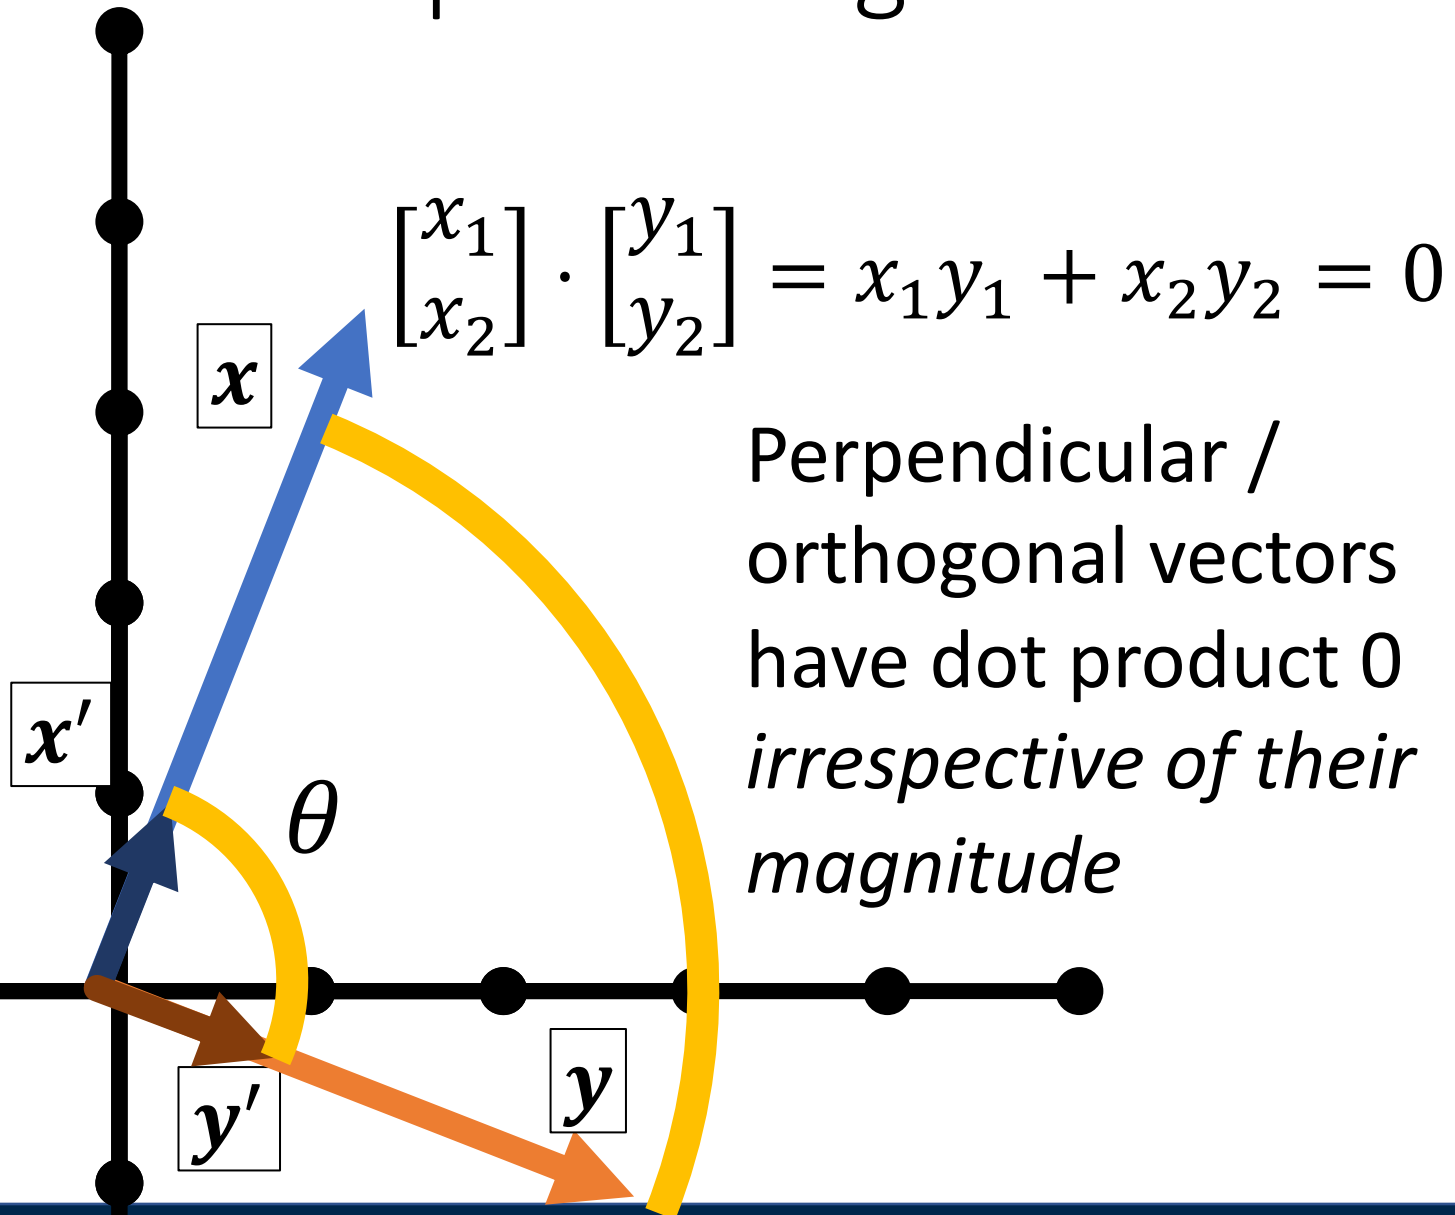
Special Angles

$$\begin{bmatrix} 1 \\ 0 \end{bmatrix} \cdot \begin{bmatrix} 0 \\ 1 \end{bmatrix} = 0 * 1 + 1 * 0 = 0$$


Perpendicular /
orthogonal vectors
have dot product 0
*irrespective of their
magnitude*



Special Angles




Orthogonal Vectors



$$x = [2, 3]$$

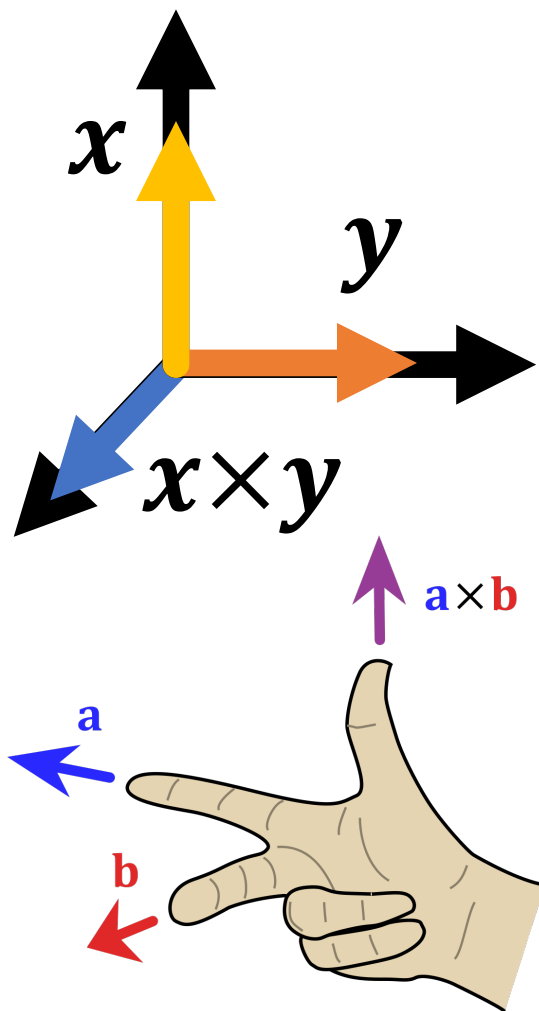
- Geometrically, what's the set of vectors that are orthogonal to x ?
- A line $[3, -2]$

Orthogonal Vectors



- What's the set of vectors that are orthogonal to $x = [5,0,0]$?
- A plane/2D space of vectors/any vector $[0, a, b]$
- What's the set of vectors that are orthogonal to x and $y = [0,5,0]$?
- A line/1D space of vectors/any vector $[0,0, b]$
- Ambiguity in *sign and magnitude*

Cross Product



- Set $\{z: z \cdot x = 0, z \cdot y = 0\}$ has an ambiguity in sign and magnitude
- Cross product $x \times y$ is: (1) orthogonal to x, y (2) has sign given by right hand rule and (3) has magnitude given by area of parallelogram of x and y
- **Important:** if x and y are the same direction or either is $\mathbf{0}$, then $x \times y = \mathbf{0}$
- Only in 3D!
- (See wedge product for $D \neq 3$)

Image credit: Wikipedia.org

Operations You Should Know

- Scale (vector, scalar \rightarrow vector)
- Add (vector, vector \rightarrow vector)
- Magnitude (vector \rightarrow scalar)
- Dot product (vector, vector \rightarrow scalar)
 - Dot products are projection / angles
- Cross product (vector, vector \rightarrow vector)
 - Vectors facing same direction have cross product 0
- You can **never** mix vectors of different sizes

Matrices

Matrices

Horizontally concatenate n , m -dim column vectors and you get a $m \times n$ matrix A (here 2×3)

$$A = [\mathbf{v}_1, \dots, \mathbf{v}_n] = \begin{bmatrix} v_{1_1} & v_{2_1} & v_{3_1} \\ v_{1_2} & v_{2_2} & v_{3_2} \end{bmatrix}$$

a (scalar)
lowercase
undecorated

a (vector)
lowercase
bold or arrow

A (matrix)
uppercase
bold

Matrices

Horizontally concatenate n , m -dim column vectors and you get a $m \times n$ matrix A (here 2×3)

$$A = [\mathbf{v}_1, \dots, \mathbf{v}_n] = \begin{bmatrix} v_{1_1} & v_{2_1} & v_{3_1} \\ v_{1_2} & v_{2_2} & v_{3_2} \end{bmatrix}$$

Watch out: In math, it's common to treat D -dim vector as a $D \times 1$ matrix (column vector);
In numpy these are different things

Matrices

Transpose: flip rows / columns

$$\begin{bmatrix} a \\ b \\ c \end{bmatrix}^T = [a \quad b \quad c] \quad (3 \times 1)^T = 1 \times 3$$

Vertically concatenate m , n -dim row vectors and you get a $m \times n$ matrix A (here 2×3)

$$A = \begin{bmatrix} \mathbf{u}_1^T \\ \vdots \\ \mathbf{u}_n^T \end{bmatrix} = \begin{bmatrix} u_{1_1} & u_{1_2} & u_{1_3} \\ u_{2_1} & u_{2_2} & u_{2_3} \end{bmatrix}$$

Matrix-vector Product

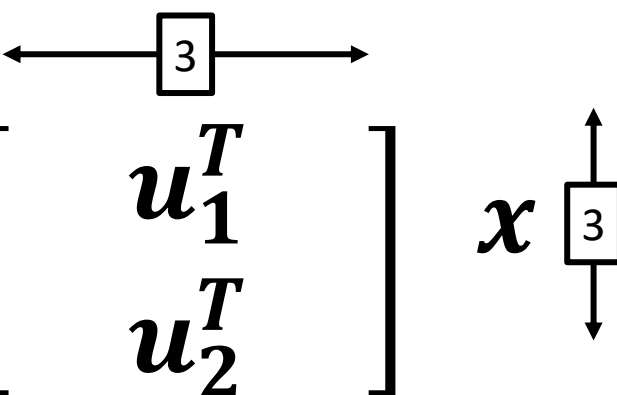
$$\mathbf{y}_{2 \times 1} = \mathbf{A}_{2 \times 3} \mathbf{x}_{3 \times 1}$$
$$\begin{bmatrix} y_1 \\ y_2 \end{bmatrix} = \begin{bmatrix} \mathbf{v}_1 & \mathbf{v}_2 & \mathbf{v}_3 \end{bmatrix} \begin{bmatrix} x_1 \\ x_2 \\ x_3 \end{bmatrix}$$

$$\mathbf{y} = x_1 \mathbf{v}_1 + x_2 \mathbf{v}_2 + x_3 \mathbf{v}_3$$

Linear combination of columns of \mathbf{A}

Matrix-vector Product

$$\mathbf{y}_{2 \times 1} = \mathbf{A}_{2 \times 3} \mathbf{x}_{3 \times 1}$$

$$\begin{bmatrix} y_1 \\ y_2 \end{bmatrix} = \begin{bmatrix} \mathbf{u}_1^T \\ \mathbf{u}_2^T \end{bmatrix} \mathbf{x}$$


$$y_1 = \mathbf{u}_1^T \mathbf{x} \quad y_2 = \mathbf{u}_2^T \mathbf{x}$$

Dot product between rows of \mathbf{A} and \mathbf{x}

Matrix Multiplication

Generally: \mathbf{A}_{mn} and \mathbf{B}_{np} yield product $(\mathbf{AB})_{mp}$

$$\mathbf{AB} = \begin{bmatrix} - & \mathbf{a}_1^T & - \\ & \vdots & \\ - & \mathbf{a}_m^T & - \end{bmatrix} \begin{bmatrix} | & & | \\ \mathbf{b}_1 & \cdots & \mathbf{b}_p \\ | & & | \end{bmatrix}$$

Yes – in \mathbf{A} , I'm referring to the rows, and in \mathbf{B} , I'm referring to the columns

Matrix Multiplication

Generally: \mathbf{A}_{mn} and \mathbf{B}_{np} yield product $(\mathbf{AB})_{mp}$

$$\mathbf{AB}_{ij} = \mathbf{a}_i^T \mathbf{b}_j$$
$$\mathbf{AB} = \begin{bmatrix} \mathbf{a}_1^T \\ \vdots \\ \mathbf{a}_m^T \end{bmatrix} \begin{bmatrix} \mathbf{b}_1 & \cdots & \mathbf{b}_p \end{bmatrix} = \begin{bmatrix} \mathbf{a}_1^T \mathbf{b}_1 & \cdots & \mathbf{a}_1^T \mathbf{b}_p \\ \vdots & \ddots & \vdots \\ \mathbf{a}_m^T \mathbf{b}_1 & \cdots & \mathbf{a}_m^T \mathbf{b}_p \end{bmatrix}$$

Matrix Multiplication

- Dimensions must match
- Dimensions must match
- Dimensions must match
- (Yes, it's associative): $ABx = (A)(Bx) = (AB)x$
- (*No it's not commutative*): $ABx \neq (BA)x \neq (BxA)$

Operations they don't teach

You Probably Saw Matrix Addition

$$\begin{bmatrix} a & b \\ c & d \end{bmatrix} + \begin{bmatrix} e & f \\ g & h \end{bmatrix} = \begin{bmatrix} a + e & b + f \\ c + g & d + h \end{bmatrix}$$

What is this? FYI: e is a scalar

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

Broadcasting

If you want to be pedantic and proper, you expand e by multiplying a matrix of 1s (denoted $\mathbf{1}$)

$$\begin{aligned} \begin{bmatrix} a & b \\ c & d \end{bmatrix} + e &= \begin{bmatrix} a & b \\ c & d \end{bmatrix} + \mathbf{1}_{2 \times 2} e \\ &= \begin{bmatrix} a & b \\ c & d \end{bmatrix} + \begin{bmatrix} e & e \\ e & e \end{bmatrix} \end{aligned}$$

Many smart matrix libraries do this automatically. This is the source of many bugs.

Broadcasting Example

Given: a $n \times 2$ matrix \mathbf{P} and a 2D column vector \mathbf{v} , Want:
 $n \times 2$ difference matrix \mathbf{D}

$$\mathbf{P} = \begin{bmatrix} x_1 & y_1 \\ \vdots & \vdots \\ x_n & y_n \end{bmatrix} \quad \mathbf{v} = \begin{bmatrix} a \\ b \end{bmatrix} \quad \mathbf{D} = \begin{bmatrix} x_1 - a & y_1 - b \\ \vdots & \vdots \\ x_n - a & y_n - b \end{bmatrix}$$

$$\mathbf{P} - \mathbf{v}^T = \begin{bmatrix} x_1 & y_1 \\ \vdots & \vdots \\ x_n & y_n \end{bmatrix} - \begin{bmatrix} a & b \\ \vdots & \vdots \\ a & b \end{bmatrix} \quad \begin{array}{l} \text{Blue stuff is} \\ \text{assumed /} \\ \text{broadcast} \end{array}$$


Two uses for Matrices

1. Storing things in a rectangular array (images, maps)
 - *Typical operations*: element-wise operations, convolution (which we'll cover next)
 - *Atypical operations*: almost anything you learned in a math linear algebra class
2. A linear operator that maps vectors to another space (\mathbf{Ax})
 - *Typical/Atypical*: reverse of above

Images as Matrices

Suppose someone hands you this matrix.

What's wrong with it?




Contrast: Gamma Curve

Typical way to change the contrast is to apply a nonlinear correction


$$\text{pixelvalue}^\gamma$$

The quantity γ controls how much contrast gets added




Contrast: Gamma Curve

Now the darkest regions (10th pctile) are **much** darker than the moderately dark regions (50th pctile).




Contrast: Gamma Correction



Contrast: Gamma Correction

Phew! Much Better.



Implementation

Python+Numpy (right way):

```
imNew = im**4
```

Python+Numpy (slow way – **why?**):

```
imNew = np.zeros(im.shape)
for y in range(im.shape[0]):
    for x in range(im.shape[1]):
        imNew[y,x] = im[y,x]**expFactor
```

Next Time:
Vectorization, Tensors,
Linear Algebra