

# Numerical Linear Algebra

EECS 442 – David Fouhey

Fall 2019, University of Michigan

[http://web.eecs.umich.edu/~fouhey/teaching/EECS442\\_W19/](http://web.eecs.umich.edu/~fouhey/teaching/EECS442_W19/)

# Administrivia

- HW 1 out – due in two weeks
- Follow submission format (wrong format = 0)
- The homeworks are **not** fill-in-the-blank. This is harder to do but mirrors life
- If it's ambiguous: make a decision, document what you think and why in your homework, and move on
- Highly encouraged to work together. See piazza
- Please check syllabus for what's allowed. I guarantee checking the syllabus thoroughly will help boost your grade.

# This Week – 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

~~**M** | USA MATHEMATICS  
UNIVERSITY OF MICHIGAN~~

# This Week – Goal

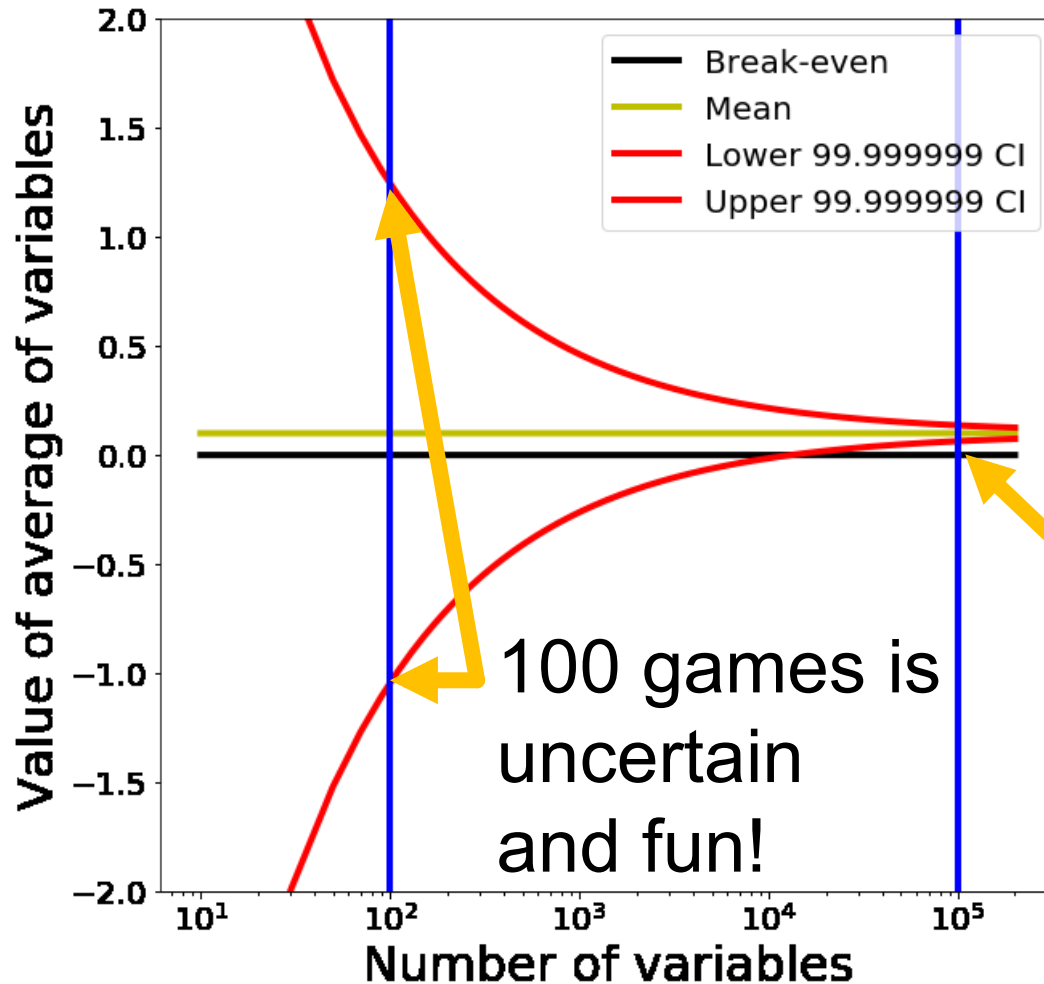
- Not a “Linear algebra in two lectures” – that’s impossible.
- *Some of this you should know!*
- Aimed at reviving your knowledge and plugging any gaps
- Aimed at giving you intuitions

# Adding Numbers

- **$1 + 1 = ?$**
- Suppose  $x_i$  is normally distributed with mean  $\mu$  and standard deviation  $\sigma$  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  $\mu$  and standard deviation  $\frac{\sigma}{\sqrt{N}}$ .*

# Free Drinks in Vegas

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



100K games is guaranteed profit: 99.999999% lowest value is \$0.064. \$0.01 for drinks \$0.054 for profits

# Let's Make It Big

- Suppose I average 50M normally distributed numbers (mean: 31, standard deviation: 1)
- For instance: have predicted and actual depth for 200 480x640 images and want to know the average error ( $|\text{predicted} - \text{actual}|$ )

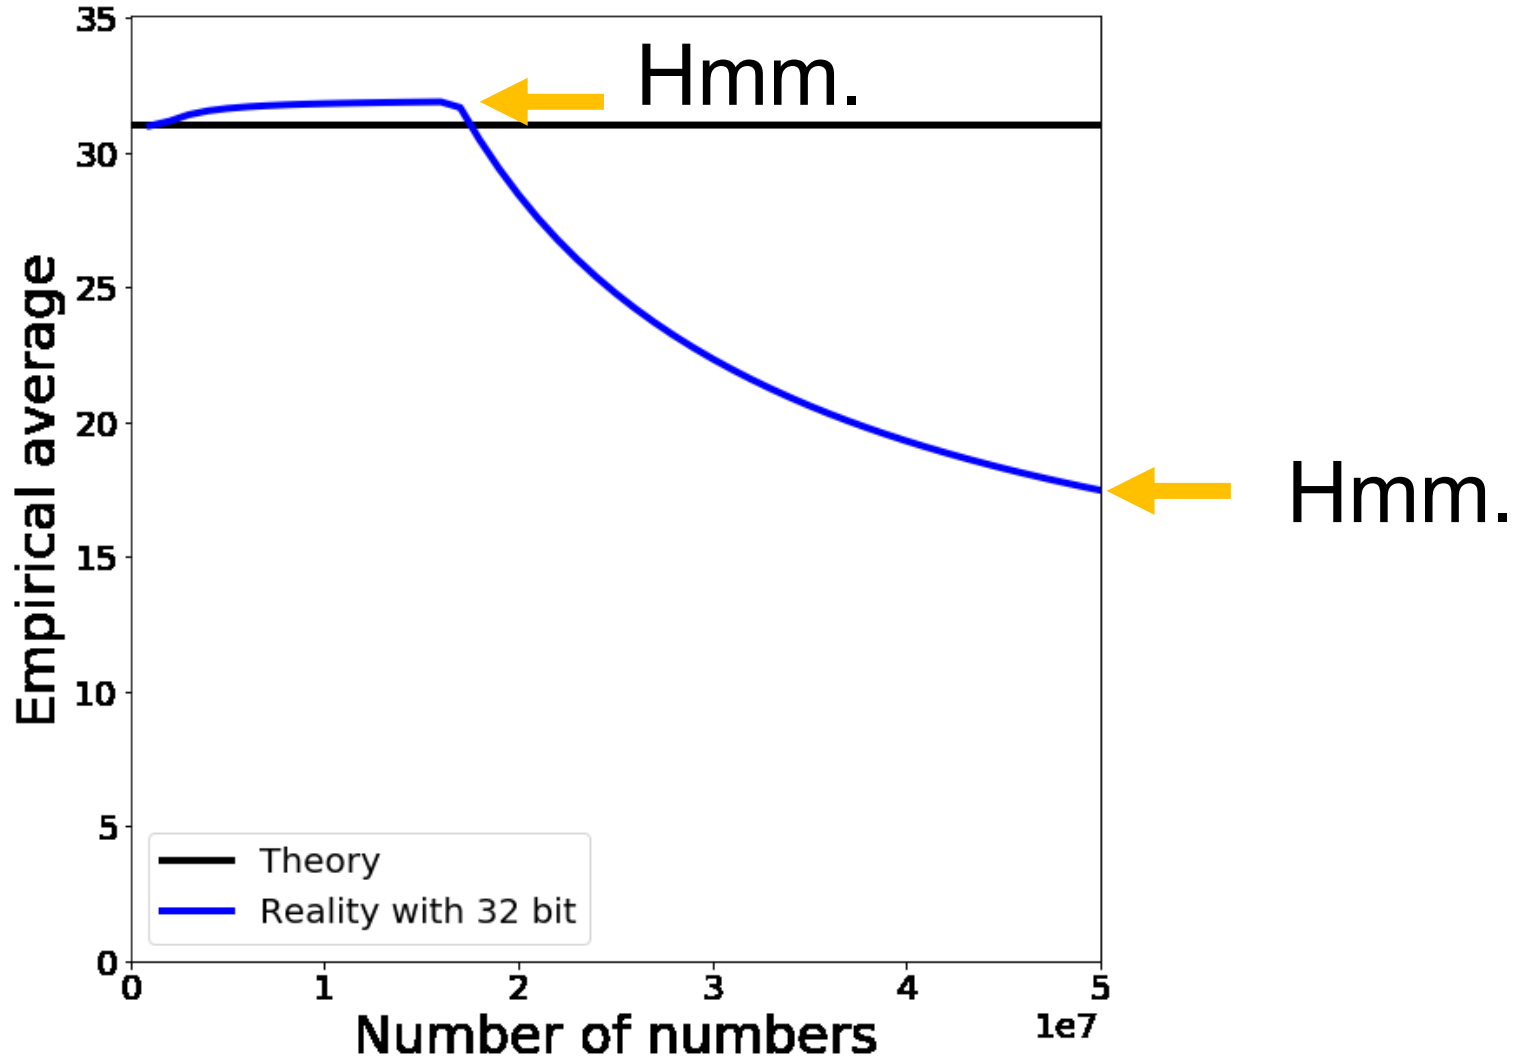
```
numerator = 0
for x in xs:
    numerator += x
return numerator / len(xs)
```

# 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's a Number?

$$\begin{array}{cccccccc} 2^7 & 2^6 & 2^5 & 2^4 & 2^3 & 2^2 & 2^1 & 2^0 & & \\ 1 & 0 & 1 & 1 & 1 & 0 & 0 & 1 & 185 & \\ 128 & + & 32 & + & 16 & + & 8 & + & 1 & = & 185 \end{array}$$

# 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	105
<hr/>									
1	0	0	1	0	0	0	1	0	34
<hr/>									
↓	↓								
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

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

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

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

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

$$32 + 0 \quad =$$

$$32 + 30 \quad =$$

$$32$$

$$62$$

*Ok – you have to multiply before dividing*

# Some Gotchas

**Should be:  
9x4=36**

math 32 + (9 x 40) / 10 = 68

uint8 32 + (9 x 40) / 10 = 42

## Overflow

32 + 9 x 40 / 10 =

32 + 104 / 10 =

32 + 10 =

42

## Why 104?

9 x 40 = 360

360 % 256 = 104

# What's 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

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

**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

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; Note: fraction = significant = mantissa; exponents of all ones or all zeros are special numbers



# Floating Point

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

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

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

# Revisiting Adding 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

# Revisiting Adding 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

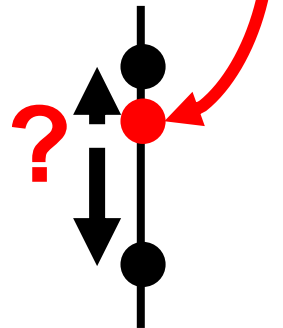
$-2^2 \times 1.00 = -4$

1

1 0 0 1

0 0 1

$-2^2 \times 1.125 = -4.5$



# Revisiting Adding 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	$-2^2 \times 1.00 = -4$

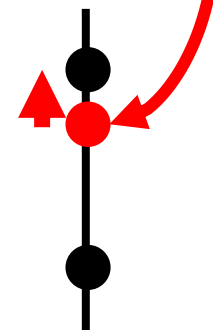
---

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

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

For a and b, these can happen

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



# Revisiting Adding Numbers

## IEEE 754 Single Precision / Single

8 bits

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

23 bits

$\approx 7$  decimal digits



## IEEE 754 Double Precision / Double

11 bits

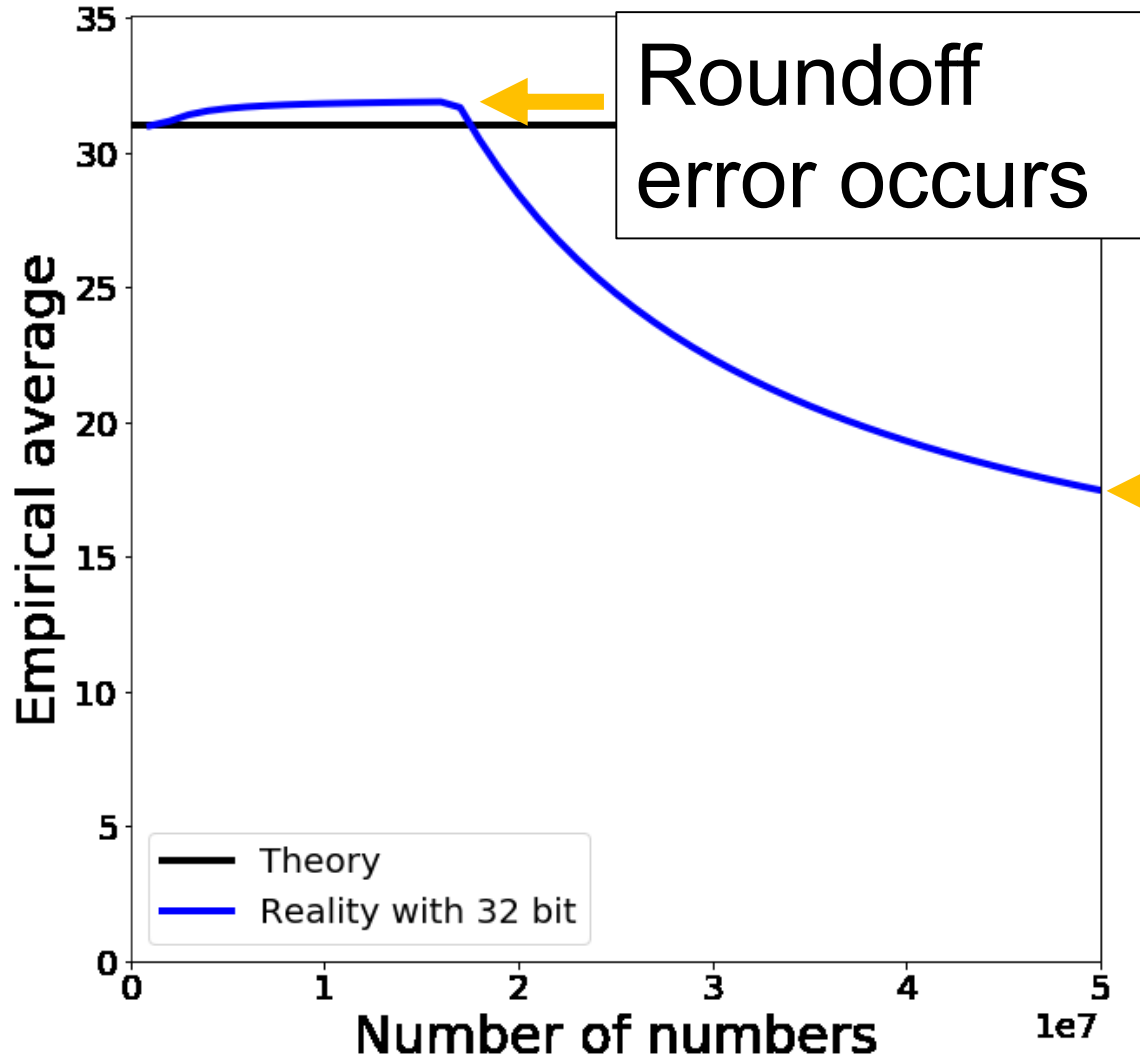
$$2^{1023} \approx 10^{308}$$

52 bits

$\approx 15$  decimal digits



# Trying it Out



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



# Take-homes

- 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

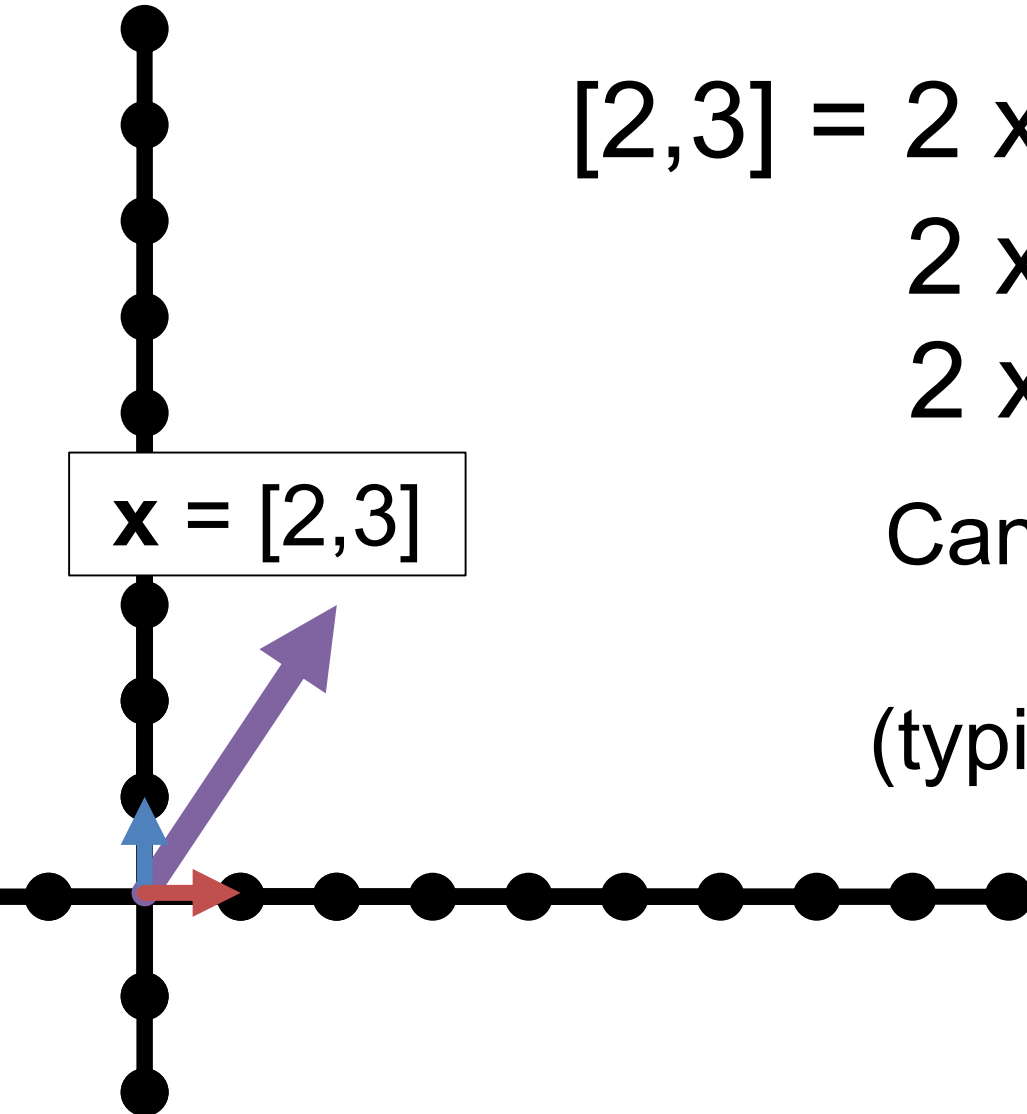
# Vectors

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

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

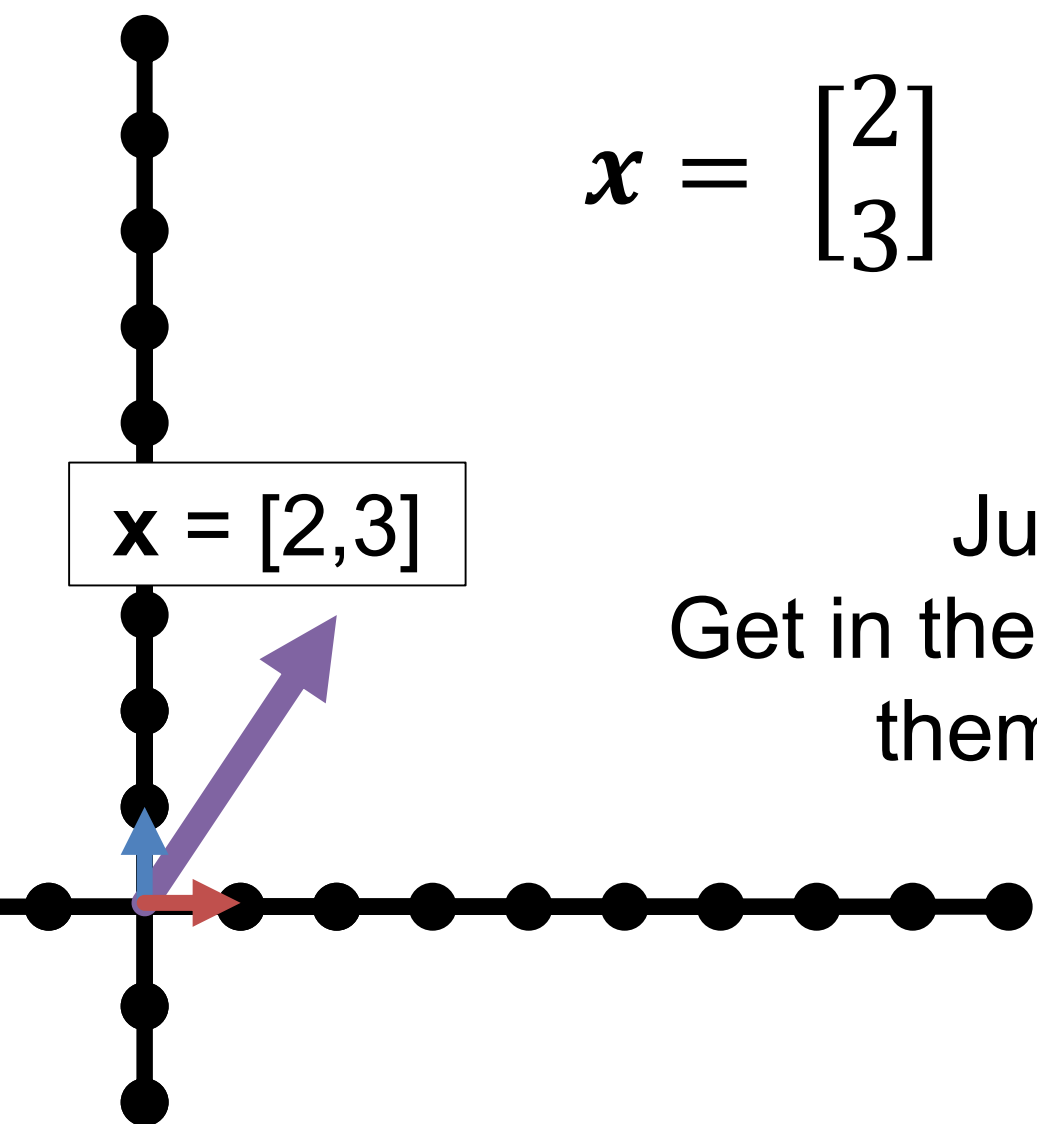
$$2 \times e_1 + 3 \times e_2$$

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



# Vectors

$$\mathbf{x} = \begin{bmatrix} 2 \\ 3 \end{bmatrix} \quad \begin{array}{l} x_1 = 2 \\ x_2 = 3 \end{array}$$

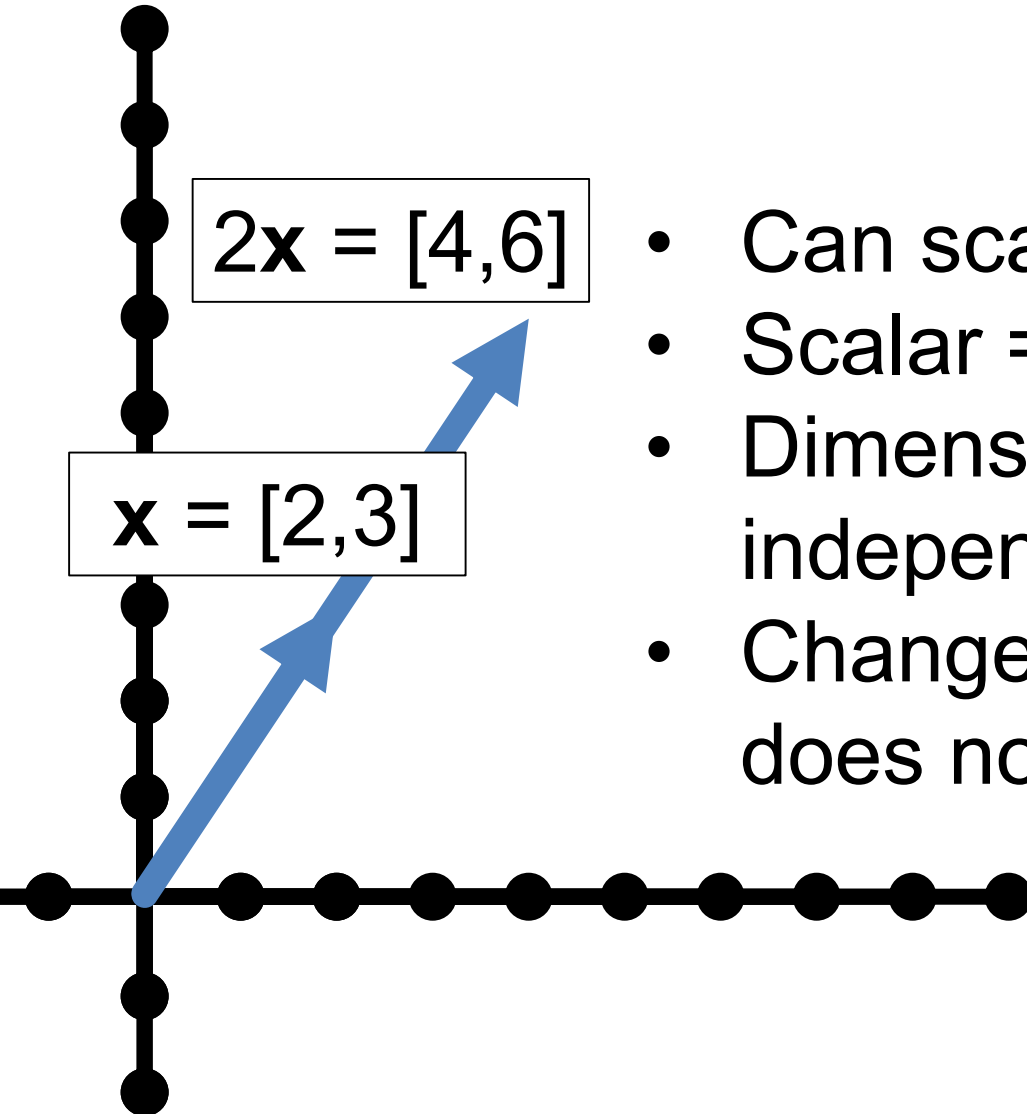


$\mathbf{x} = [2, 3]$

Just an array!

Get in the habit of thinking of them as columns.

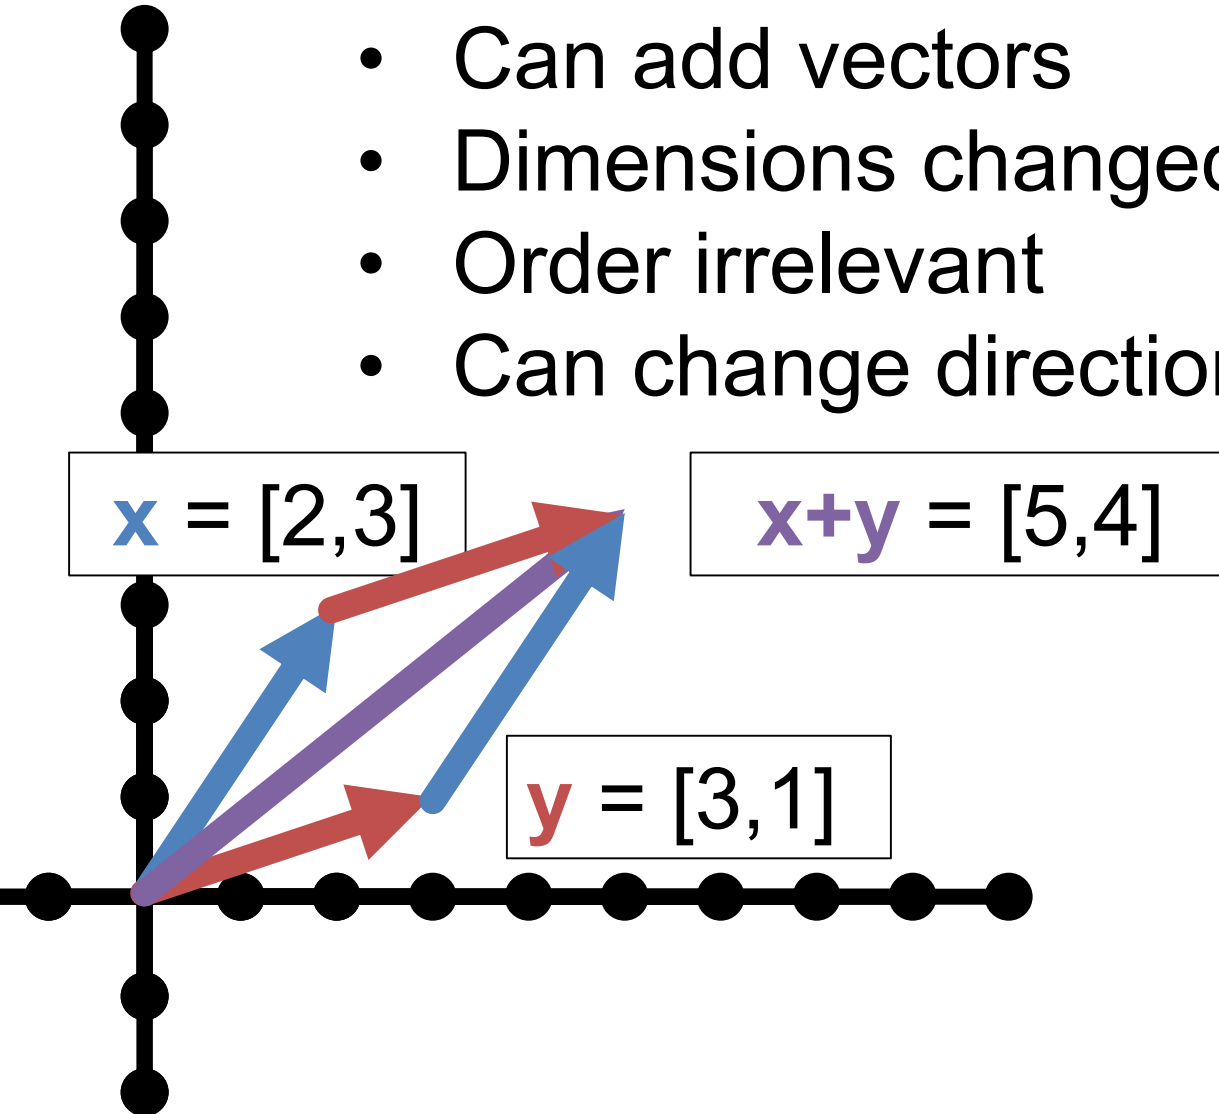
# Scaling Vectors



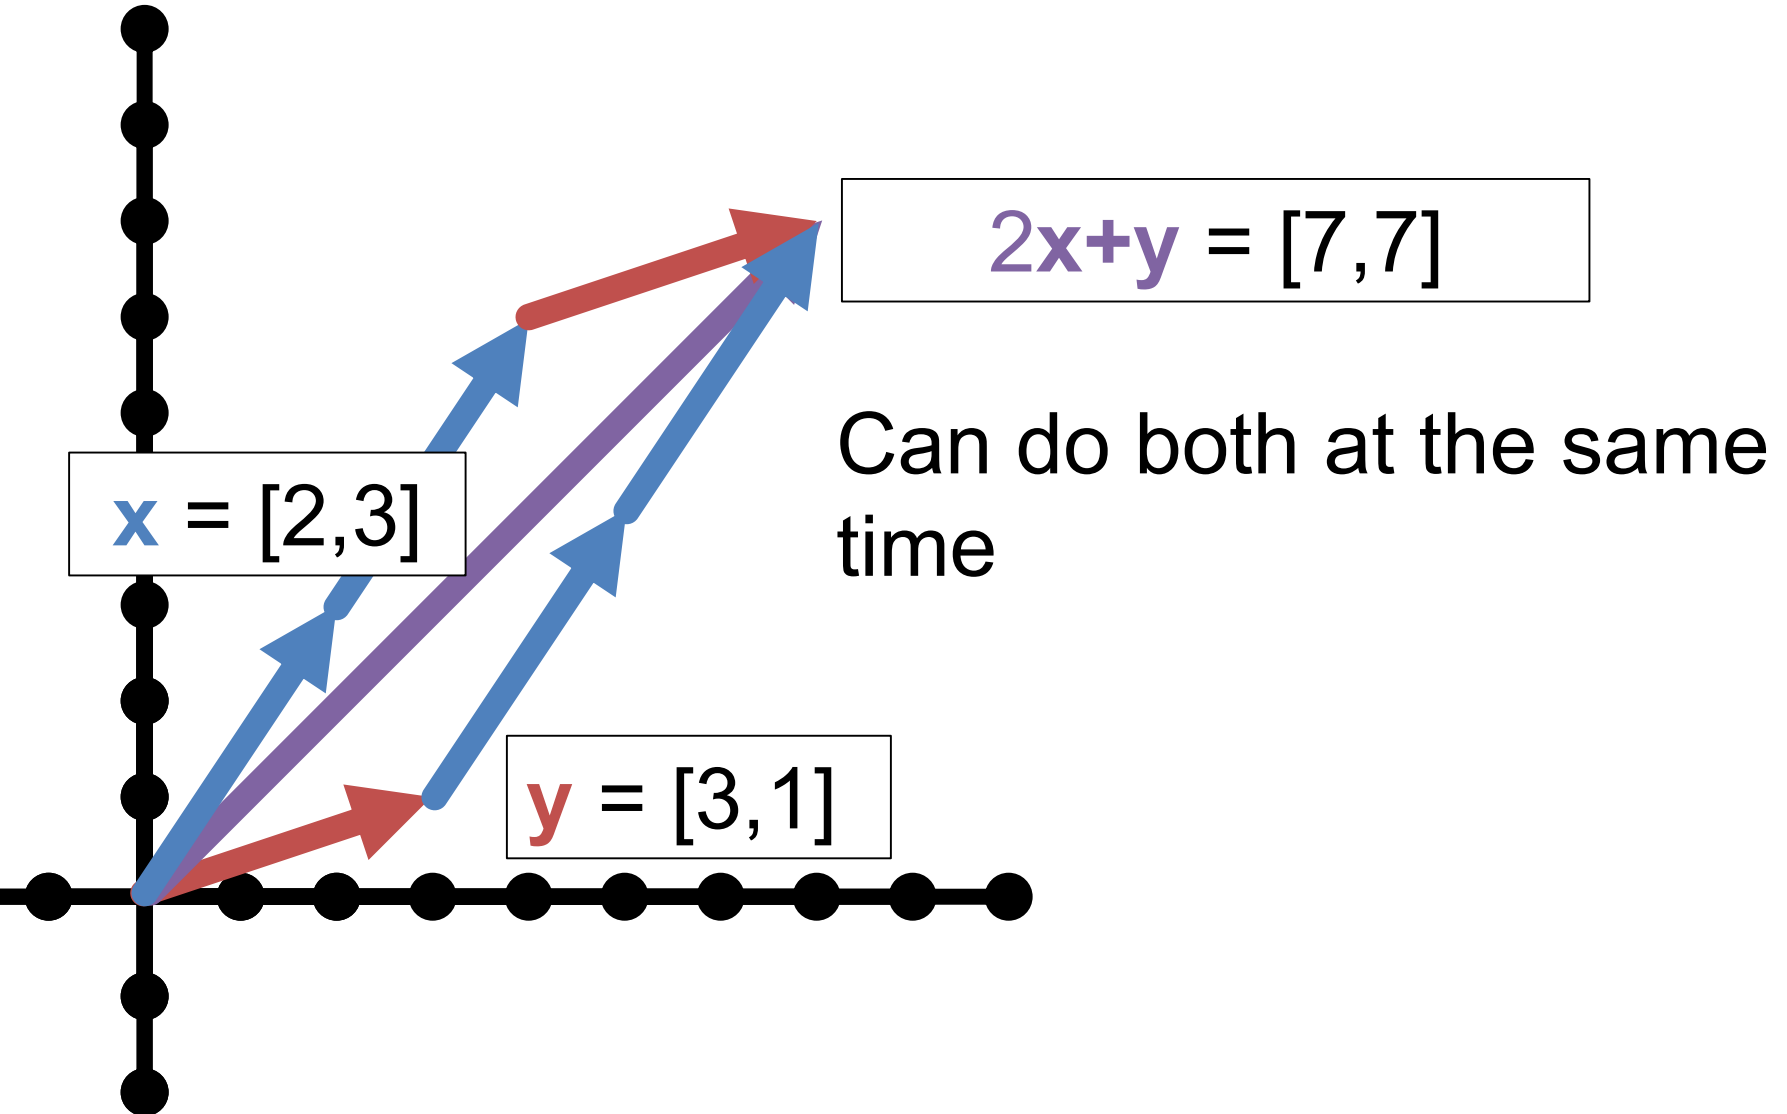
- Can scale vector by a *scalar*
- Scalar = single number
- Dimensions changed independently
- Changes *magnitude / length*, does not change *direction*.

# Adding Vectors

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



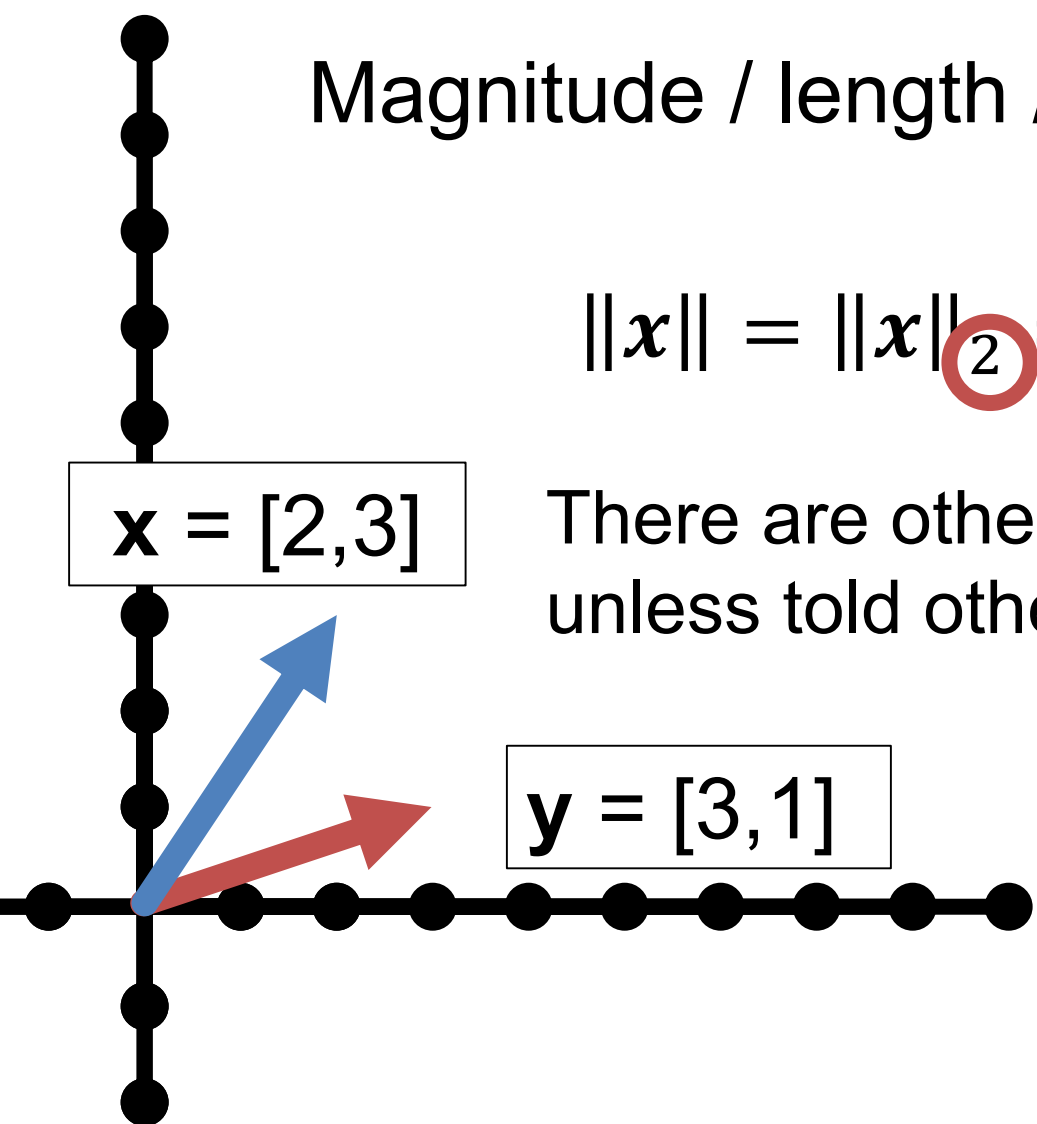
# Scaling and Adding



# Measuring Length

Magnitude / length / (L2) norm of vector

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



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

There are other norms; assume L2 unless told otherwise

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

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

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

**Why?**

# Normalizing a Vector

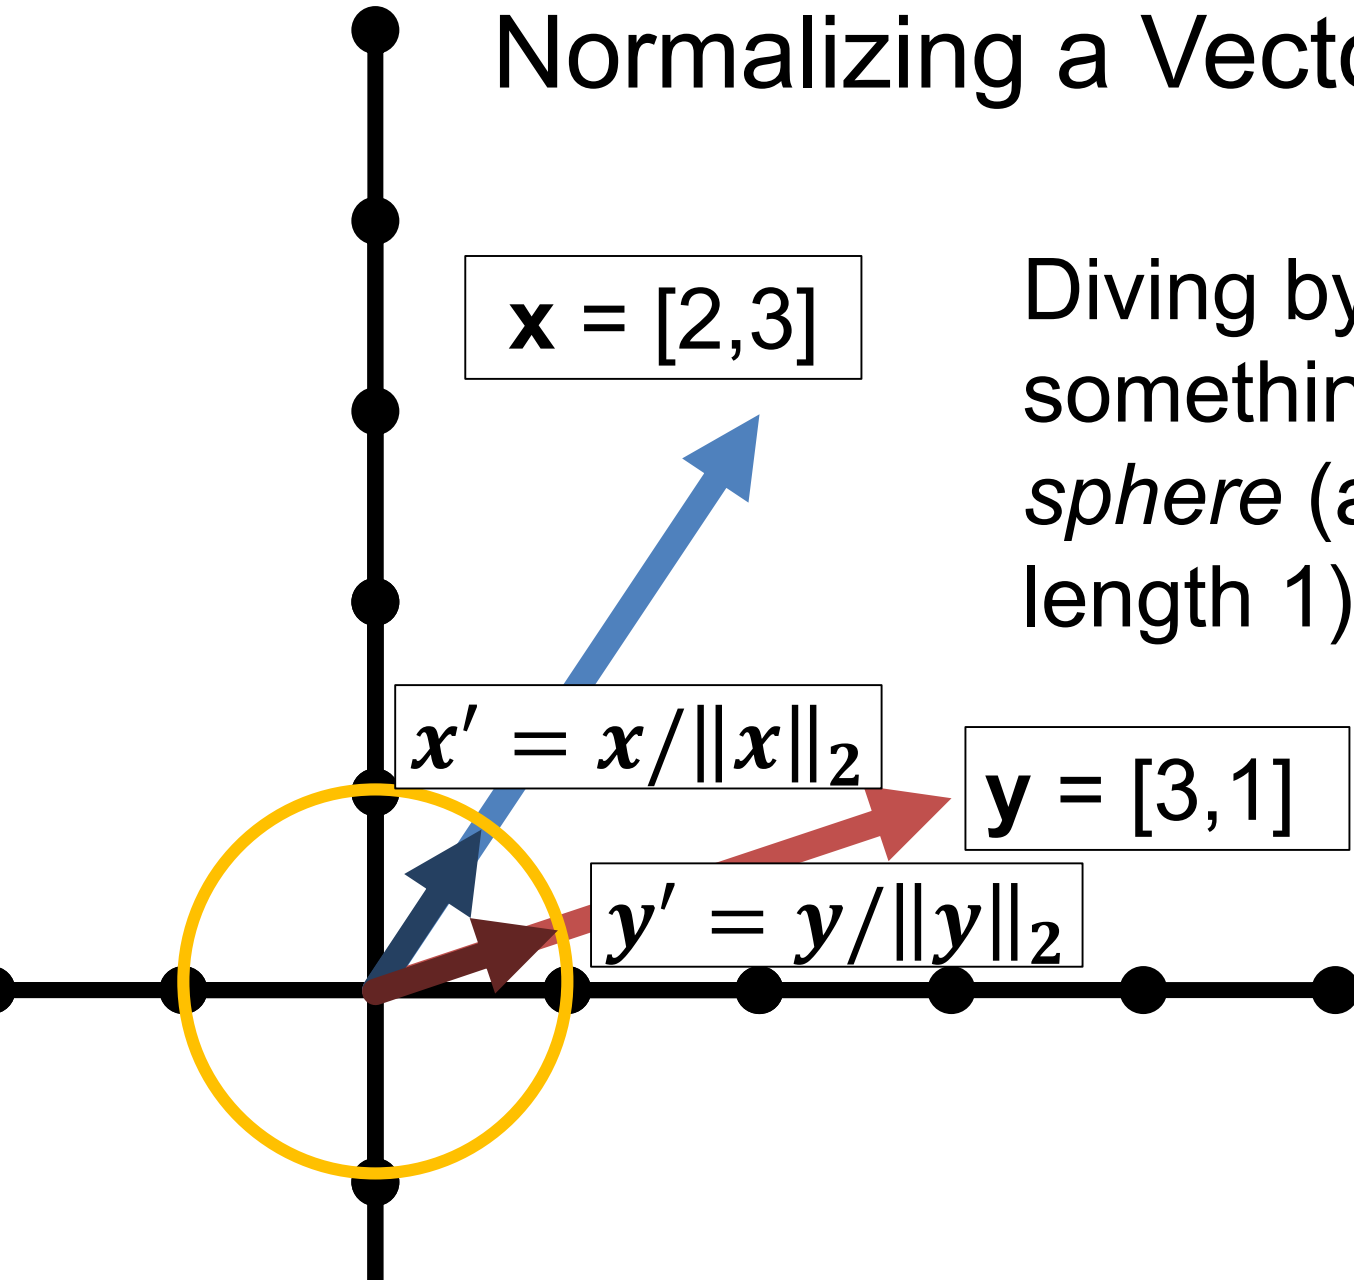
$$\mathbf{x} = [2, 3]$$

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

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

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

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



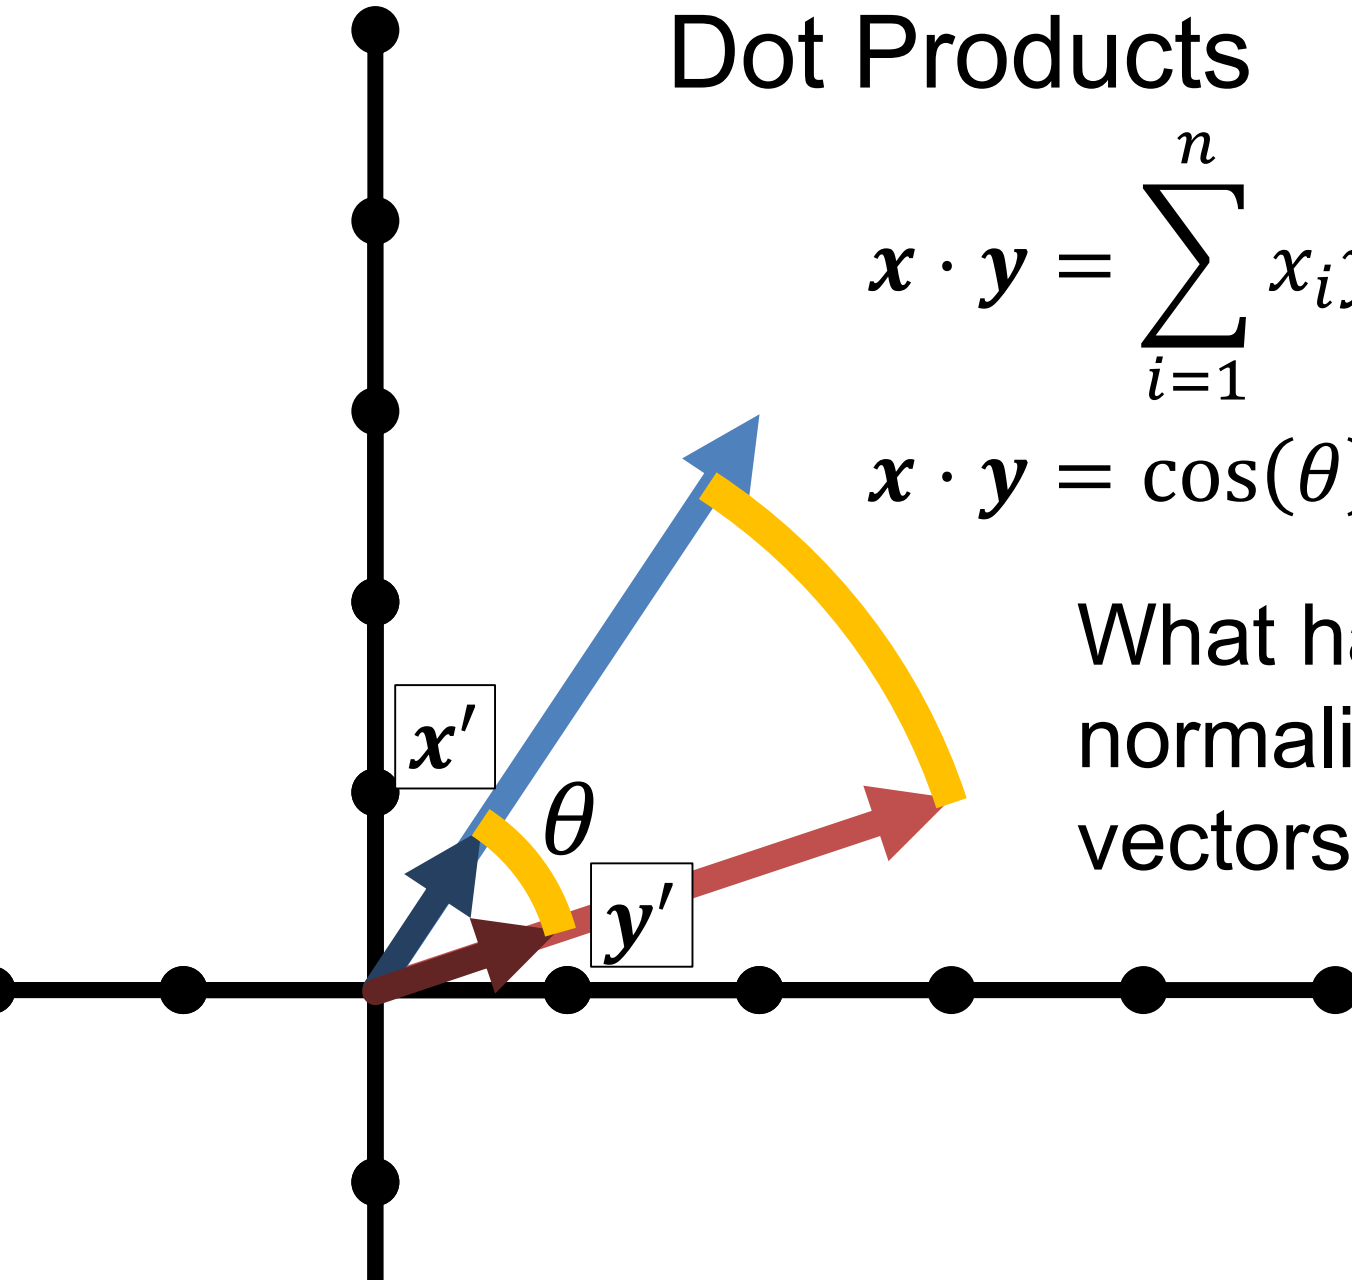


# Dot Products

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

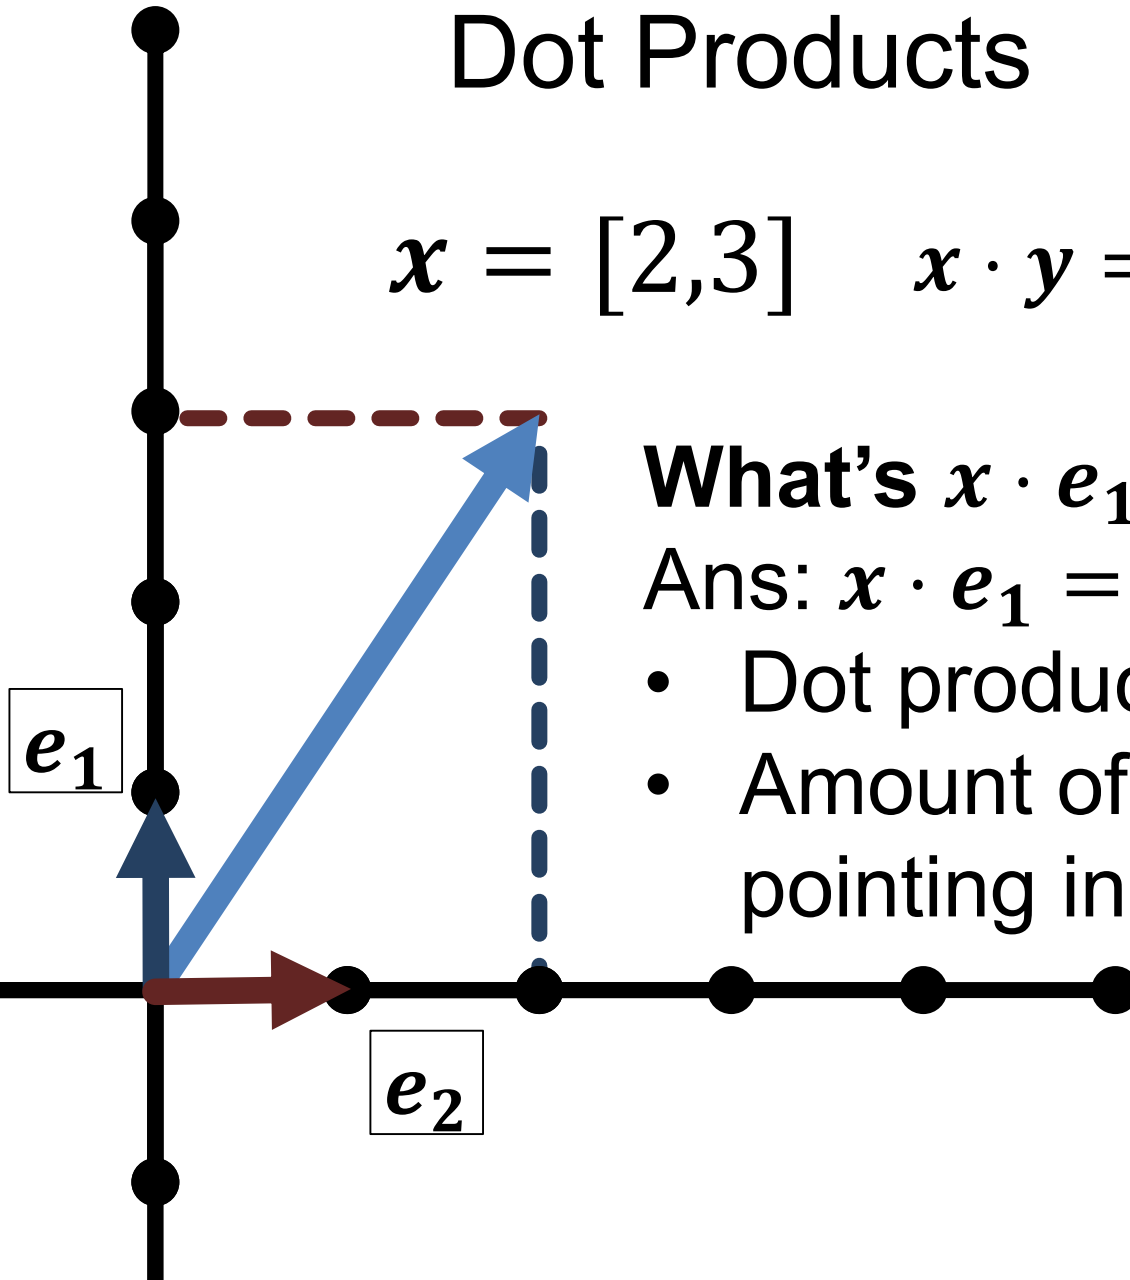
$$\mathbf{x} \cdot \mathbf{y} = \cos(\theta) \|\mathbf{x}\| \|\mathbf{y}\|$$

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:  $\mathbf{x} \cdot \mathbf{e}_1 = 1$  ;  $\mathbf{x} \cdot \mathbf{e}_2 = 3$

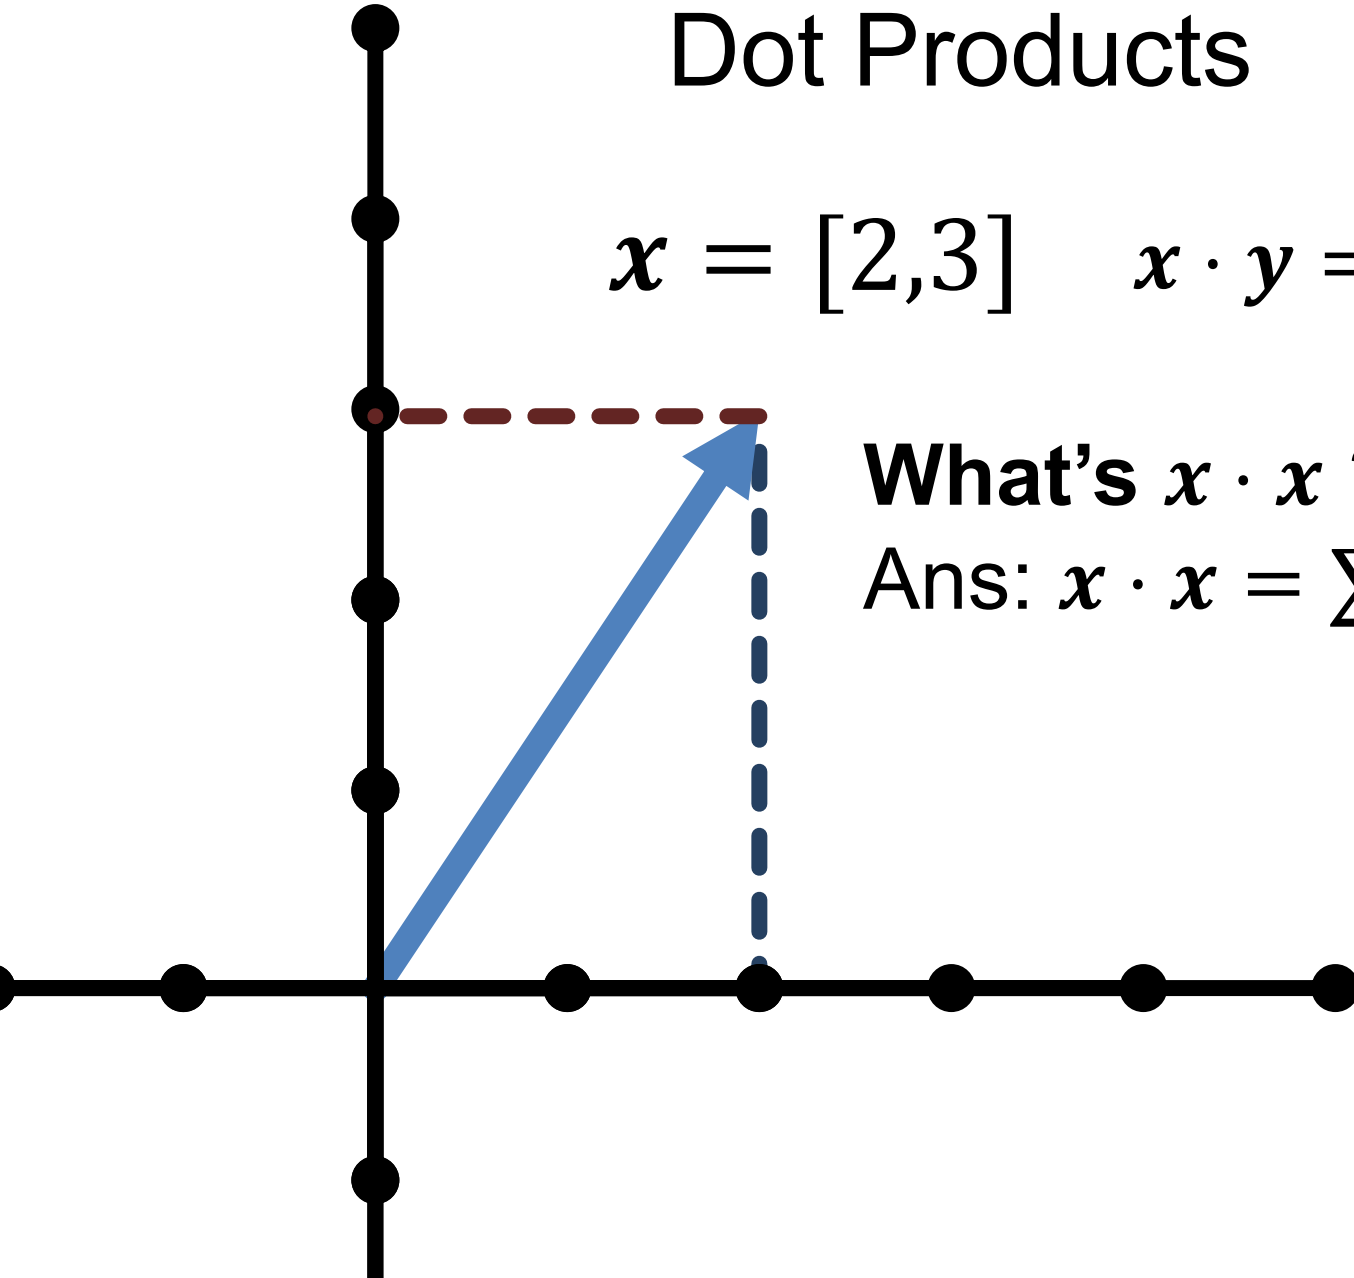
- Dot product is projection
- Amount of  $\mathbf{x}$  that's also pointing in direction of  $\mathbf{y}$

# Dot Products

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

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

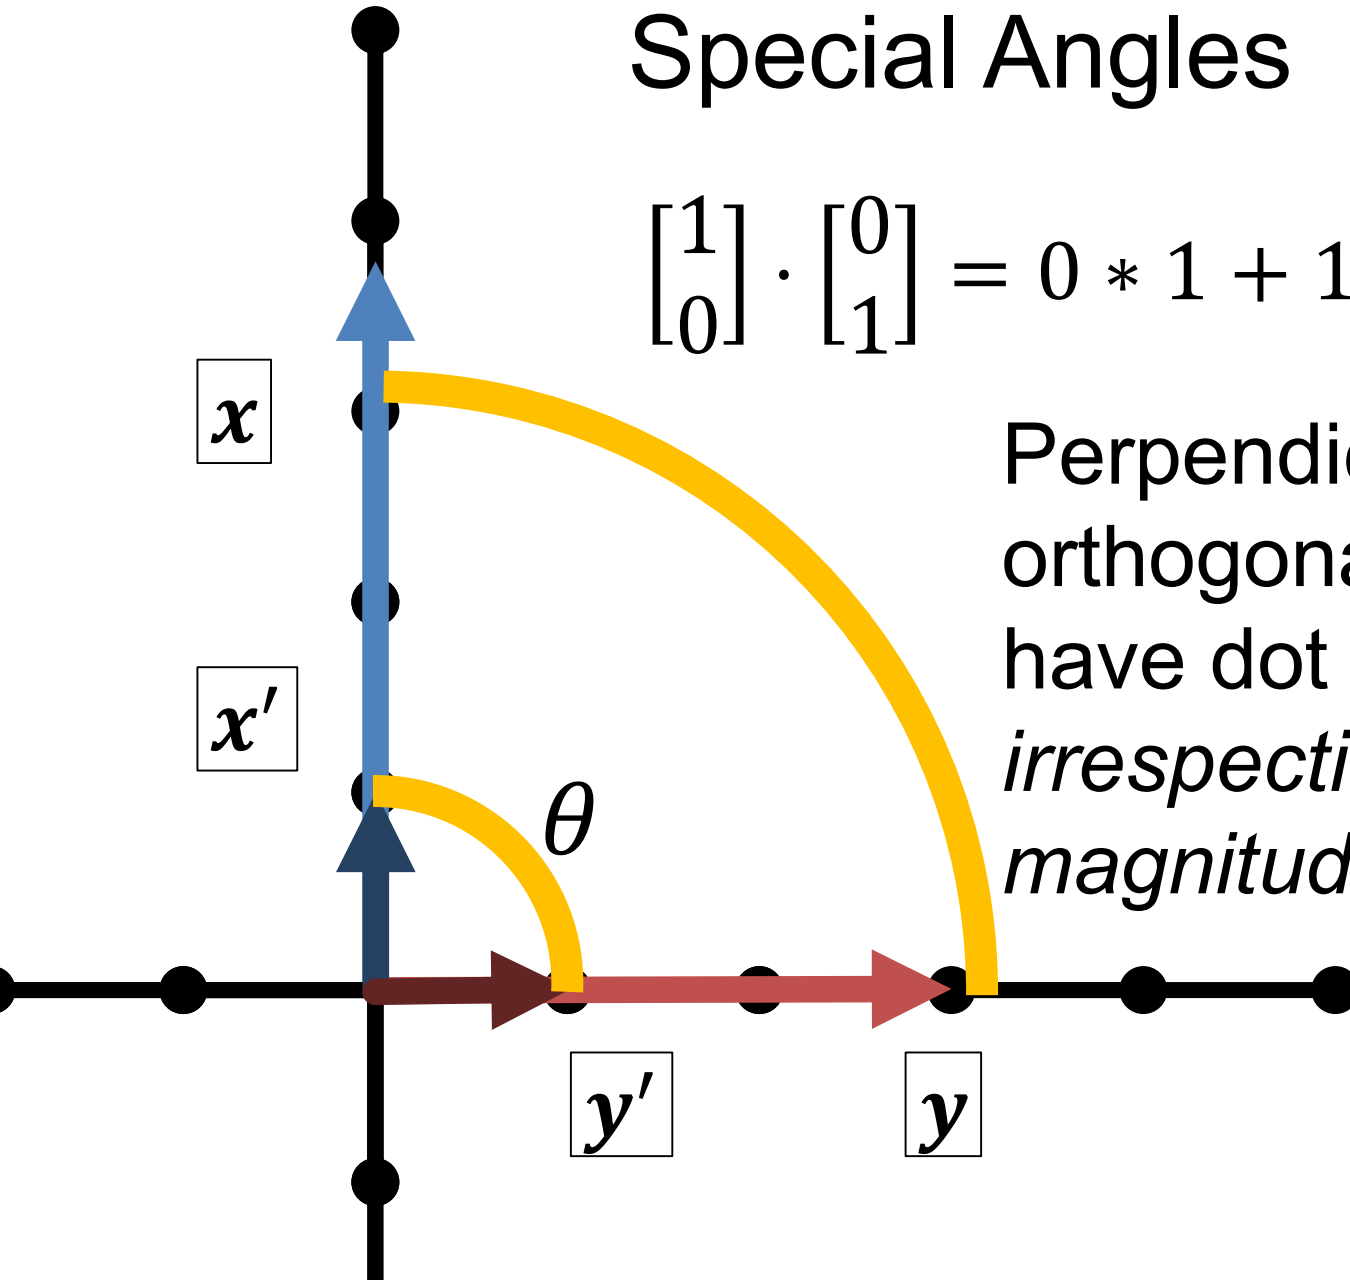
Ans:  $\mathbf{x} \cdot \mathbf{x} = \sum x_i x_i = \|\mathbf{x}\|_2^2$



# Special Angles

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

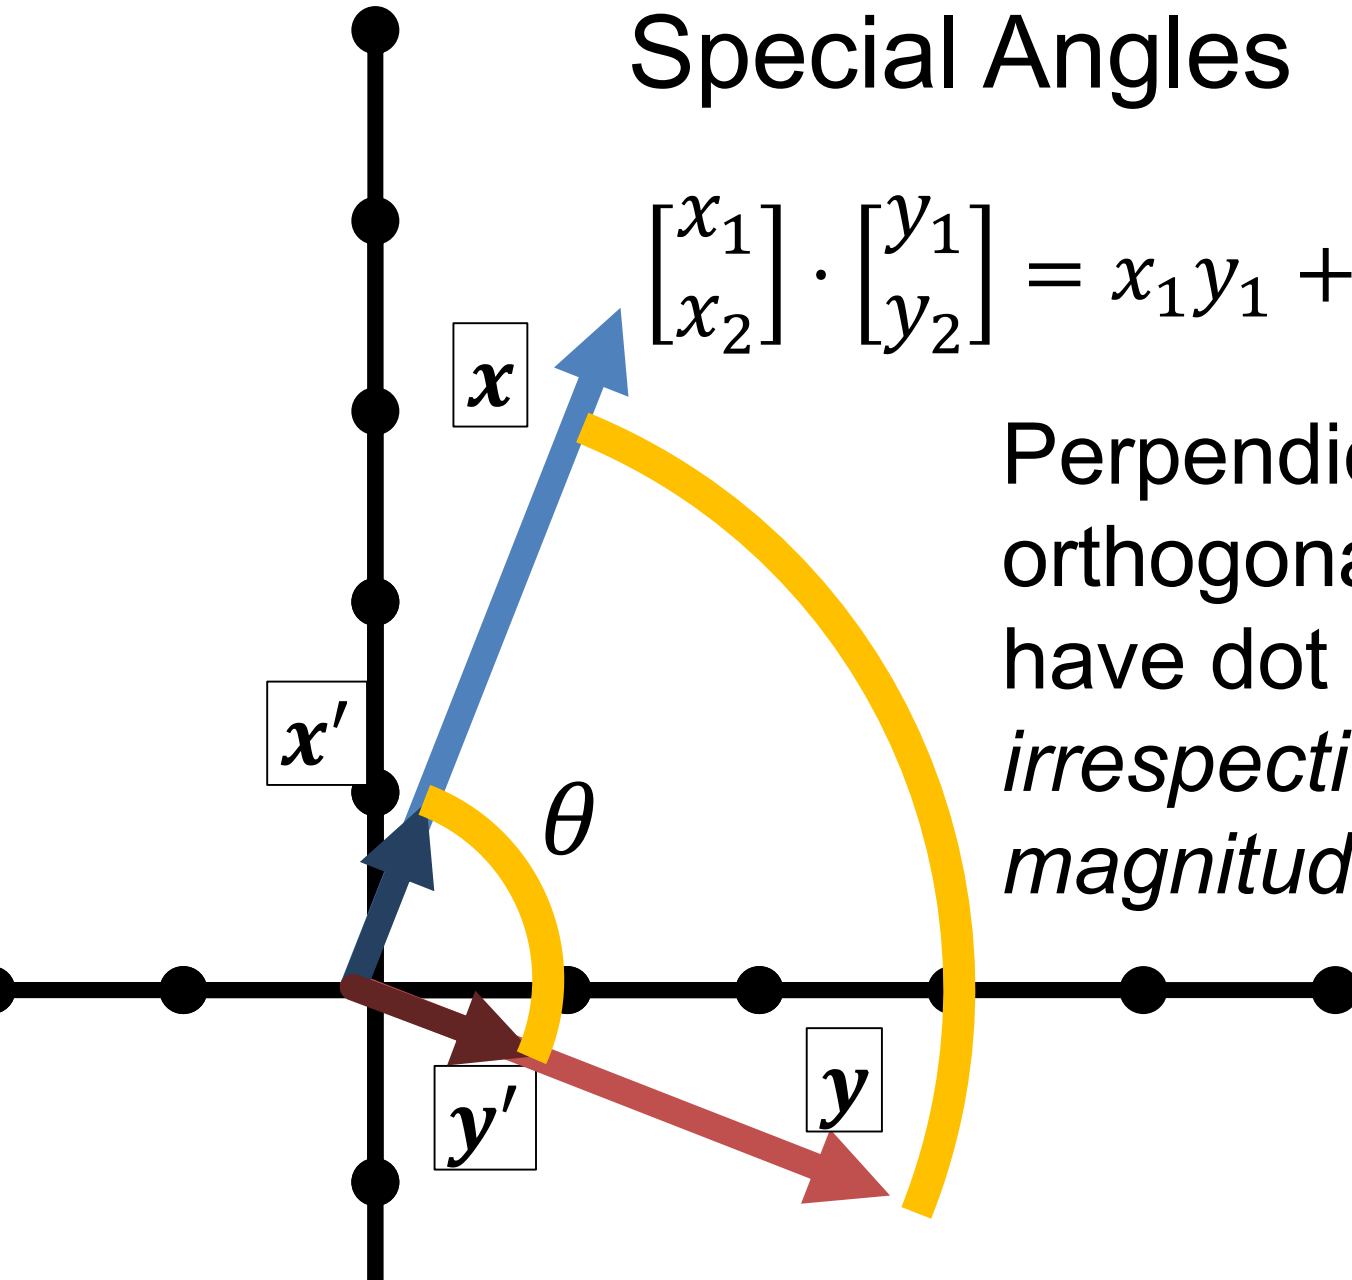
Perpendicular /  
orthogonal vectors  
have dot product 0  
*irrespective of their  
magnitude*



# Special Angles

$$\begin{bmatrix} x_1 \\ x_2 \end{bmatrix} \cdot \begin{bmatrix} y_1 \\ y_2 \end{bmatrix} = x_1 y_1 + x_2 y_2 = 0$$

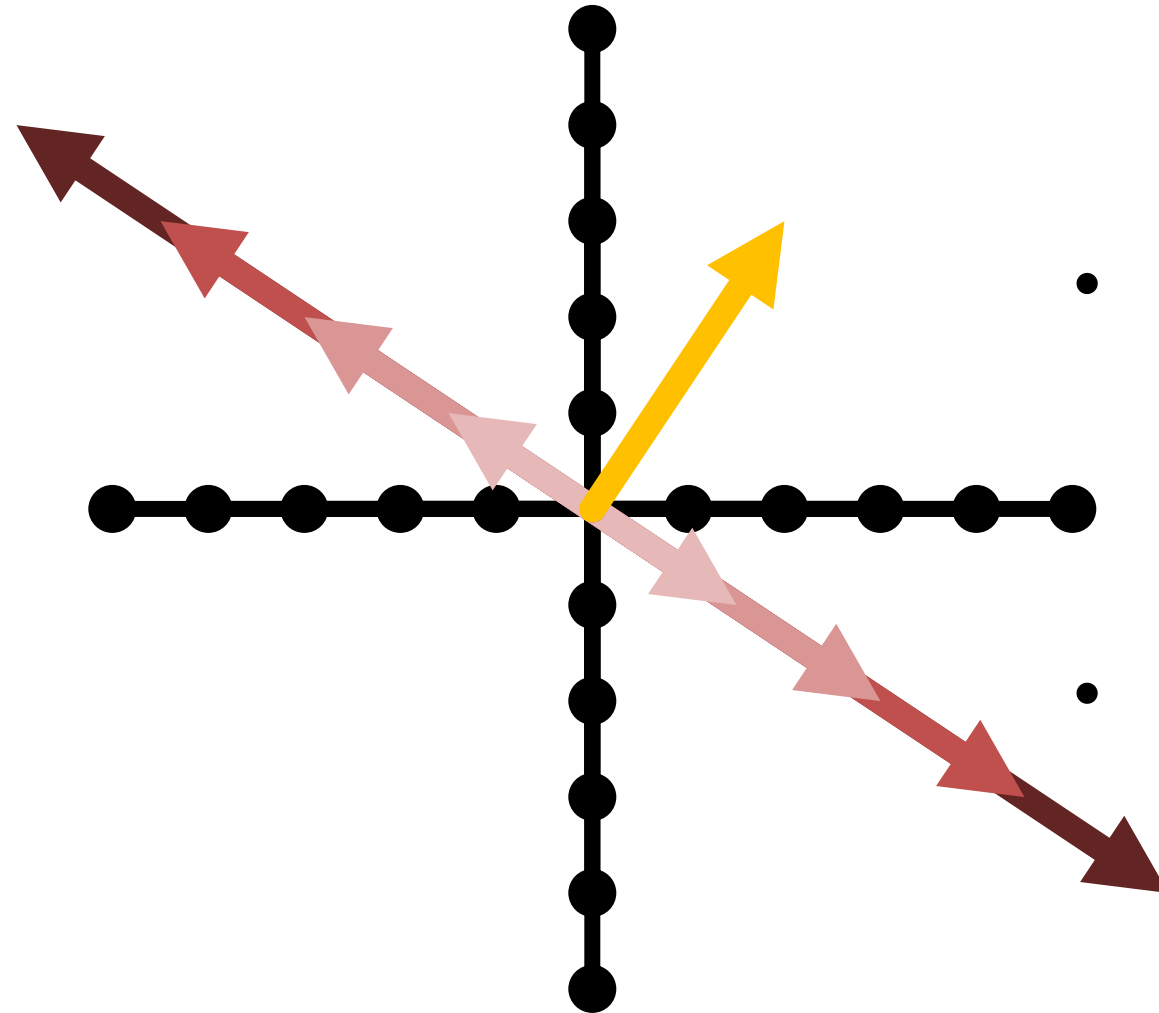
Perpendicular /  
orthogonal vectors  
have dot product 0  
*irrespective of their  
magnitude*



# Orthogonal Vectors

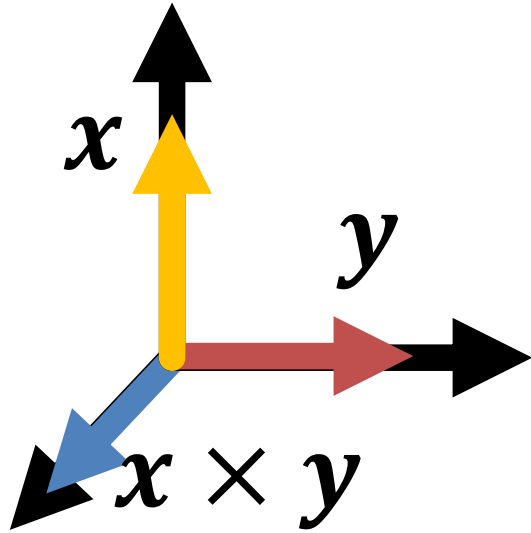
$$x = [2, 3]$$

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

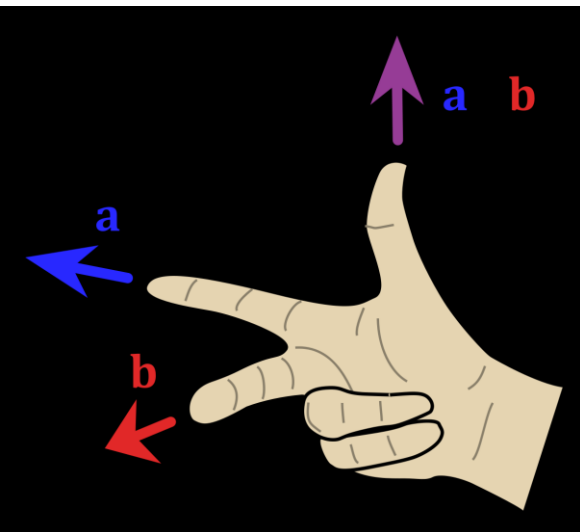




# Cross Product



- 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!





# 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

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

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

**a** (scalar)  
lowercase  
undecorated

**a** (vector)  
lowercase  
**bold or arrow**

**A** (matrix)  
uppercase  
**bold**

# Matrices

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  $2 \times 3$ )

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

# 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

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

$$\begin{array}{c} \leftarrow \boxed{3} \rightarrow \\ \left[ \begin{array}{c} y_1 \\ y_2 \end{array} \right] = \left[ \begin{array}{c} \mathbf{u}_1 \\ \mathbf{u}_2 \end{array} \right] \mathbf{x} \begin{array}{c} \updownarrow \\ \boxed{3} \\ \updownarrow \end{array} \end{array}$$

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

*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  $(\mathbf{AB})_{mp}$

$$\mathbf{AB} = \begin{bmatrix} \text{---} a_1^T \text{---} \\ \vdots \\ \text{---} a_m^T \text{---} \end{bmatrix} \begin{bmatrix} \downarrow b_1 & \dots & \downarrow b_p \\ a_1^T b_1 & \dots & a_1^T b_p \\ \vdots & \ddots & \vdots \\ a_m^T b_1 & \dots & a_m^T b_p \end{bmatrix}$$

$$AB_{ij} = a_i^T b_j$$

# 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  $\mathbf{D}$

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

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

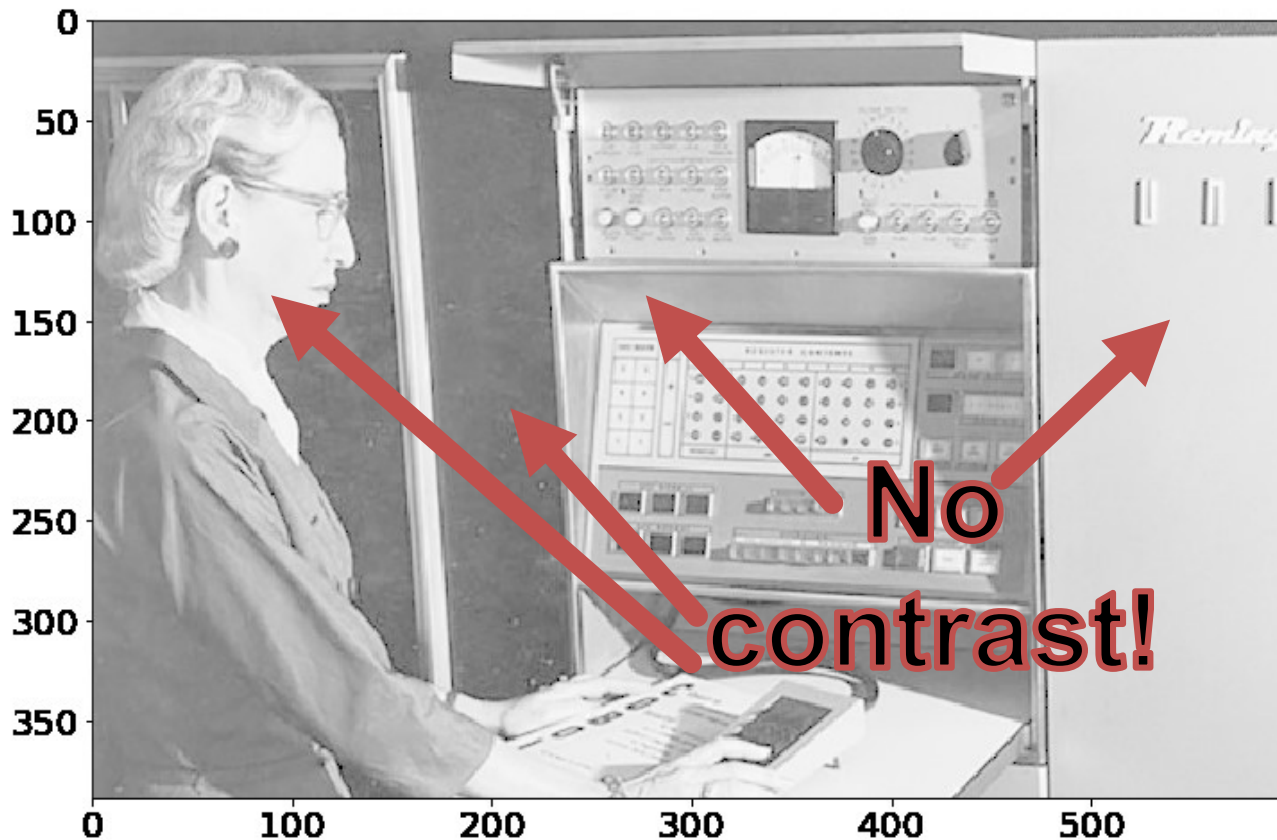
# Two Uses for Matrices

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 (**Ax**)
  - *Typical/Atypical*: reverse of above

# Images as Matrices

Suppose someone hands you this matrix.

**What's wrong with it?**

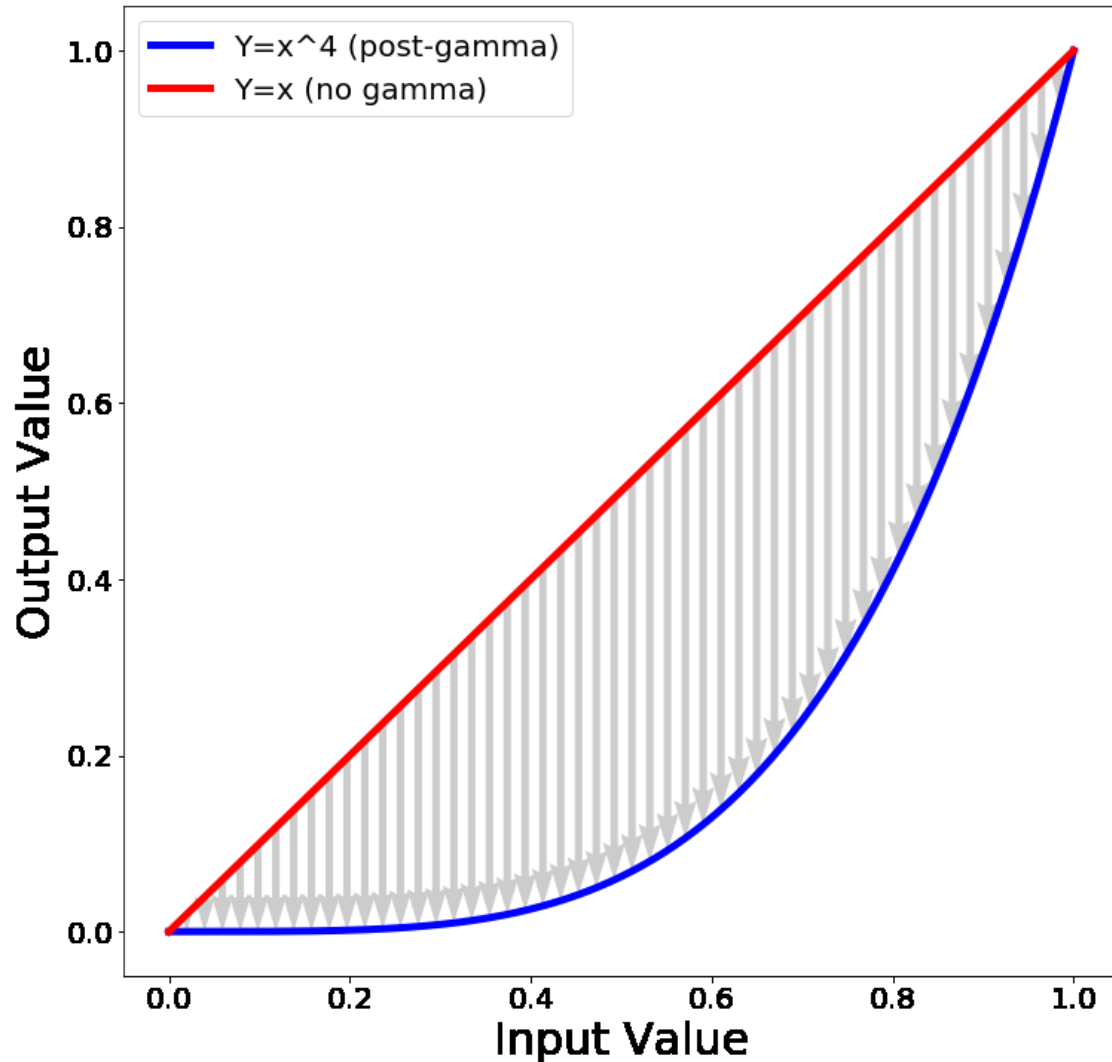


# Contrast – Gamma curve

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

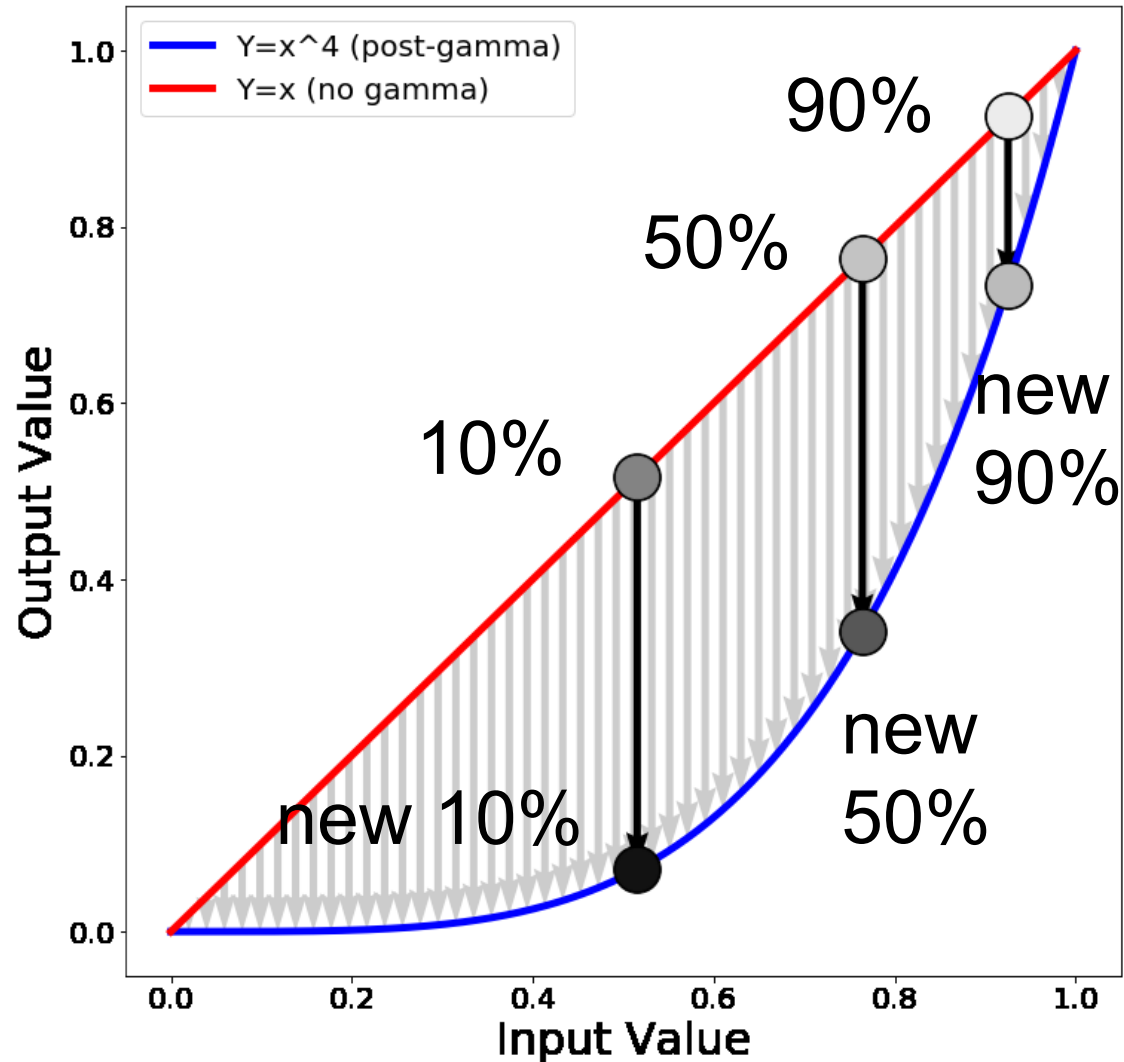
$$\text{pixelvalue}^\gamma$$

The quantity  $\gamma$  controls how much contrast gets added



# Contrast – Gamma curve

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



# Implementation

Python+Numpy (right way):

```
imNew = im**0.25
```

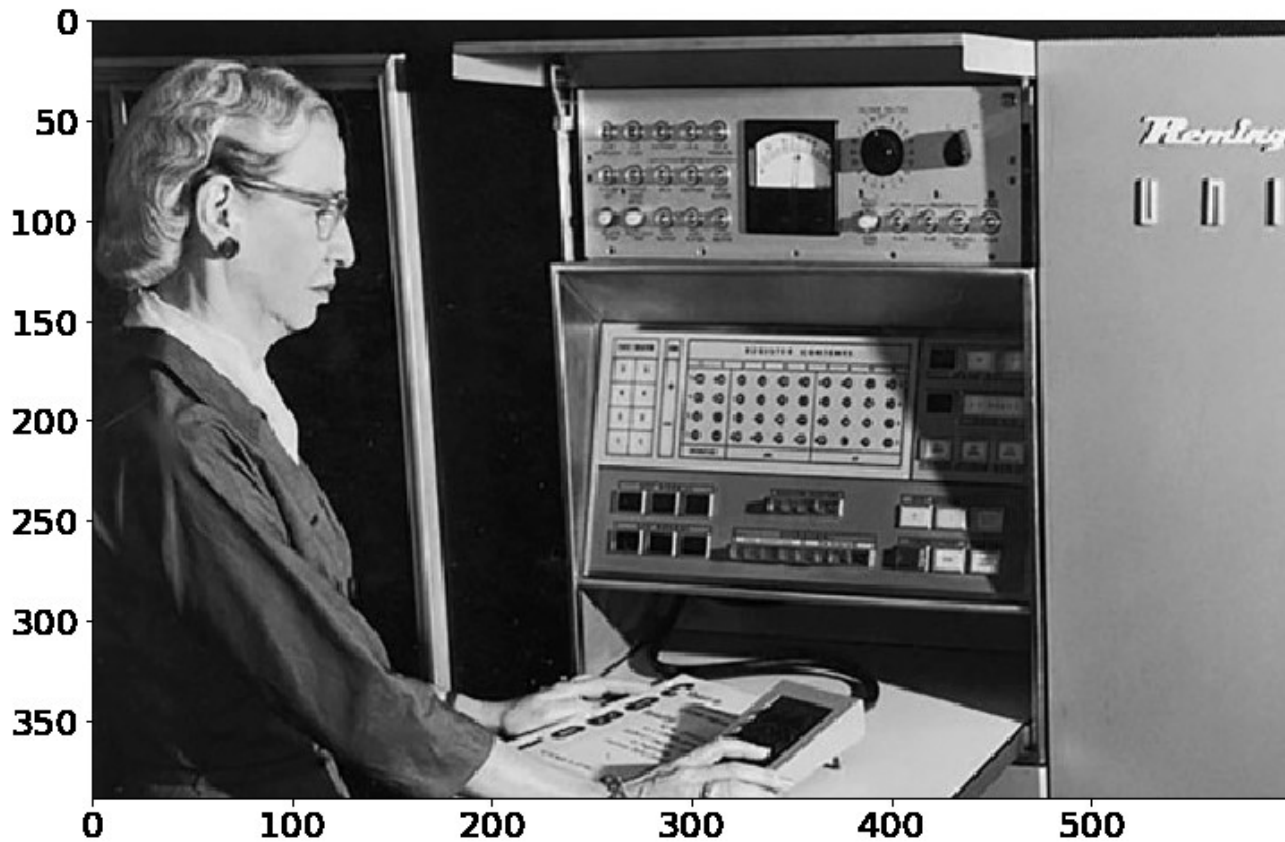
Python+Numpy (slow way – **why?**):

```
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
```



# Results

Phew! Much Better.



# Element-wise Operations

Element-wise power – beware notation

$$(\mathbf{A}^p)_{ij} = A_{ij}^p$$

“Hadamard Product” / Element-wise multiplication

$$(\mathbf{A} \odot \mathbf{B})_{ij} = A_{ij} * B_{ij}$$

Element-wise division

$$(\mathbf{A}/\mathbf{B})_{ij} = \frac{A_{ij}}{B_{ij}}$$

# Sums Across Axes

Suppose have  
Nx2 matrix **A**

$$\mathbf{A} = \begin{bmatrix} x_1 & y_1 \\ \vdots & \vdots \\ x_n & y_n \end{bmatrix}$$

ND col. vec.

$$\Sigma(\mathbf{A}, 1) = \begin{bmatrix} x_1 + y_1 \\ \vdots \\ x_n + y_n \end{bmatrix}$$

2D row vec

$$\Sigma(\mathbf{A}, 0) = \left[ \sum_{i=1}^n x_i \quad , \quad \sum_{i=1}^n y_i \right]$$

*Note – libraries distinguish between N-D column vector and Nx1 matrix.*

# Vectorizing Example

- Suppose I represent each image as a 128-dimensional vector
- I want to compute all the pairwise distances between  $\{\mathbf{x}_1, \dots, \mathbf{x}_N\}$  and  $\{\mathbf{y}_1, \dots, \mathbf{y}_M\}$  so I can find, for every  $\mathbf{x}_i$  the nearest  $\mathbf{y}_j$
- Identity:  $\|\mathbf{x} - \mathbf{y}\|^2 = \|\mathbf{x}\|^2 + \|\mathbf{y}\|^2 - 2\mathbf{x}^T \mathbf{y}$
- Or:  $\|\mathbf{x} - \mathbf{y}\| = (\|\mathbf{x}\|^2 + \|\mathbf{y}\|^2 - 2\mathbf{x}^T \mathbf{y})^{1/2}$

# Vectorizing Example

$$\mathbf{X} = \begin{bmatrix} - & \mathbf{x}_1 & - \\ & \vdots & \\ - & \mathbf{x}_N & - \end{bmatrix} \quad \mathbf{Y} = \begin{bmatrix} - & \mathbf{y}_1 & - \\ & \vdots & \\ - & \mathbf{y}_M & - \end{bmatrix} \quad \mathbf{Y}^T = \begin{bmatrix} | & & | \\ \mathbf{y}_1 & \cdots & \mathbf{y}_M \\ | & & | \end{bmatrix}$$

Compute a  $N \times 1$   
vector of norms  
(can also do  $M \times 1$ )

$$\Sigma(\mathbf{X}^2, \mathbf{1}) = \begin{bmatrix} \|\mathbf{x}_1\|^2 \\ \vdots \\ \|\mathbf{x}_N\|^2 \end{bmatrix}$$

Compute a  $N \times M$   
matrix of dot products

$$(\mathbf{X}\mathbf{Y}^T)_{ij} = \mathbf{x}_i^T \mathbf{y}_j$$

# Vectorizing Example

$$\mathbf{D} = \left( \Sigma(\mathbf{X}^2, \mathbf{1}) + \Sigma(\mathbf{Y}^2, \mathbf{1})^T - 2\mathbf{X}\mathbf{Y}^T \right)^{1/2}$$

$$\begin{bmatrix} \|\mathbf{x}_1\|^2 \\ \vdots \\ \|\mathbf{x}_N\|^2 \end{bmatrix} + [\|\mathbf{y}_1\|^2 \quad \cdots \quad \|\mathbf{y}_M\|^2]$$

$$\begin{bmatrix} \|\mathbf{x}_1\|^2 + \|\mathbf{y}_1\|^2 & \cdots & \|\mathbf{x}_1\|^2 + \|\mathbf{y}_M\|^2 \\ \vdots & \ddots & \vdots \\ \|\mathbf{x}_N\|^2 + \|\mathbf{y}_1\|^2 & \cdots & \|\mathbf{x}_N\|^2 + \|\mathbf{y}_M\|^2 \end{bmatrix}$$

**Why?**

$$(\Sigma(\mathbf{X}^2, \mathbf{1}) + \Sigma(\mathbf{Y}^2, \mathbf{1})^T)_{ij} = \|\mathbf{x}_i\|^2 + \|\mathbf{y}_j\|^2$$

# Vectorizing Example

$$\mathbf{D} = \left( \Sigma(\mathbf{X}^2, 1) + \Sigma(\mathbf{Y}^2, 1)^T - 2\mathbf{X}\mathbf{Y}^T \right)^{1/2}$$

$$\mathbf{D}_{ij} = \|\mathbf{x}_i\|^2 + \|\mathbf{y}_j\|^2 + 2\mathbf{x}_i^T \mathbf{y}_j$$

Numpy code:

```
XNorm = np.sum(X**2, axis=1, keepdims=True)
```

```
YNorm = np.sum(Y**2, axis=1, keepdims=True)
```

```
D = (XNorm + YNorm.T - 2*np.dot(X, Y.T)) ** 0.5
```

---

\*May have to make sure this is at least 0  
(sometimes roundoff issues happen)

# Does it Make a Difference?

Computing pairwise distances between 300 and 400 128-dimensional vectors

1. for x in X, for y in Y, using native python: 9s
2. for x in X, for y in Y, using numpy to compute distance: 0.8s
3. vectorized: 0.0045s (~2000x faster than 1, 175x faster than 2)

*Expressing things in primitives that are optimized is usually faster*



# Linear Independence

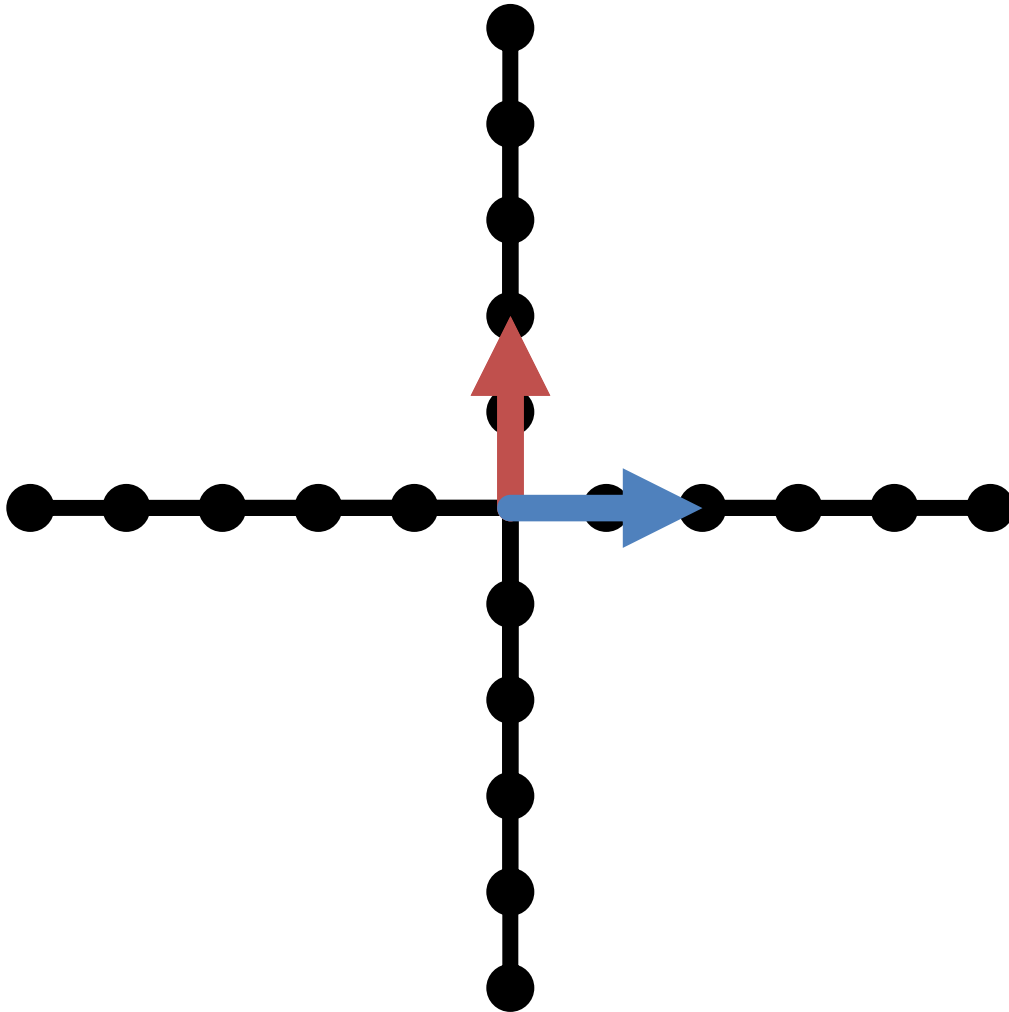
A set of vectors is linearly independent if you can't write one as a linear combination of the others.

$$\text{Suppose: } \mathbf{a} = \begin{bmatrix} 0 \\ 0 \\ 2 \end{bmatrix} \quad \mathbf{b} = \begin{bmatrix} 0 \\ 6 \\ 0 \end{bmatrix} \quad \mathbf{c} = \begin{bmatrix} 5 \\ 0 \\ 0 \end{bmatrix}$$

$$\mathbf{x} = \begin{bmatrix} 0 \\ 0 \\ 4 \end{bmatrix} = 2\mathbf{a} \quad \mathbf{y} = \begin{bmatrix} 0 \\ -2 \\ 1 \end{bmatrix} = \frac{1}{2}\mathbf{a} - \frac{1}{3}\mathbf{b}$$

- Is the set  $\{\mathbf{a}, \mathbf{b}, \mathbf{c}\}$  linearly independent?
- Is the set  $\{\mathbf{a}, \mathbf{b}, \mathbf{x}\}$  linearly independent?
- Max # of independent 3D vectors?

# Span



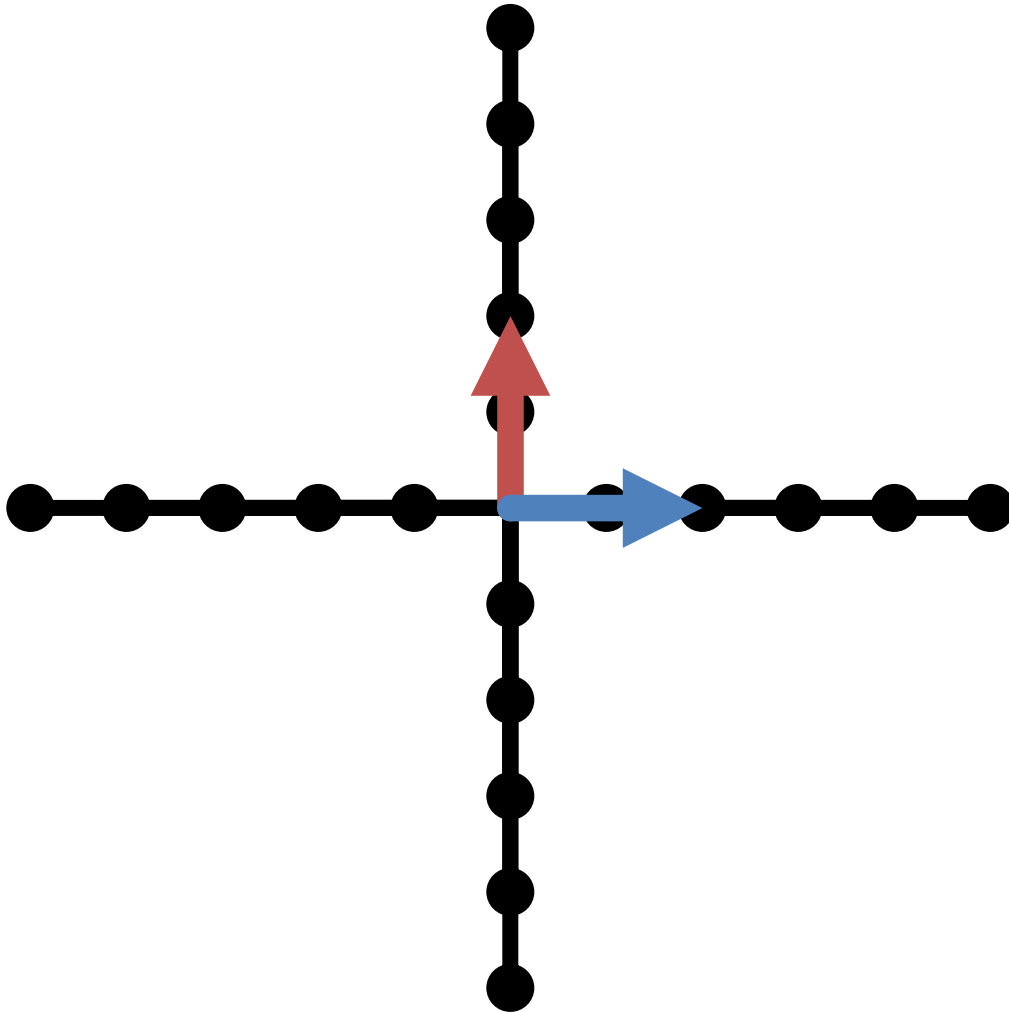
Span: all linear combinations of a set of vectors

Span( $\{\uparrow\}$ ) =  
Span( $\{[0,1]\}$ ) = ?

All vertical lines through origin =  
 $\{\lambda[0,1]: \lambda \in R\}$

Is **blue** in **{red}**'s span?

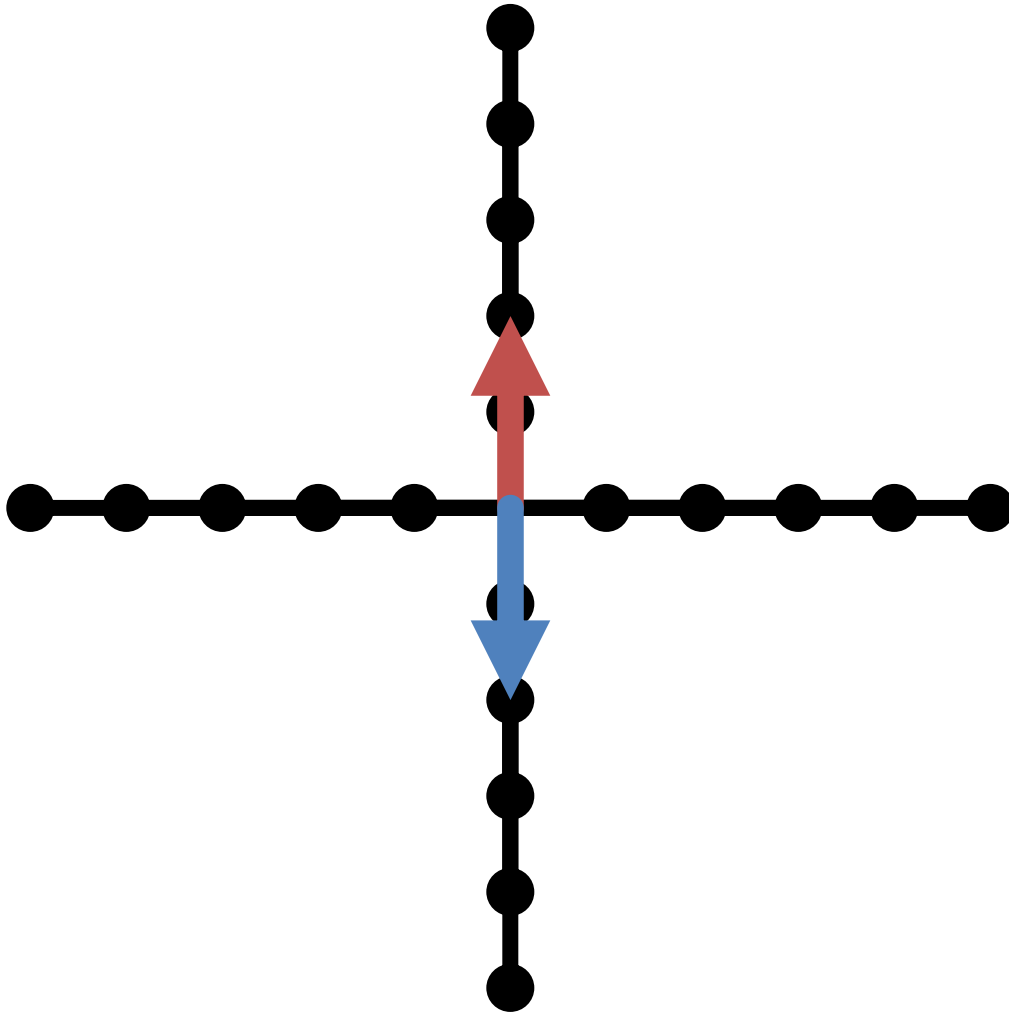
# Span



Span: all linear combinations of a set of vectors

$$\text{Span}(\{\uparrow, \rightarrow\}) = ?$$

# Span



Span: all linear combinations of a set of vectors

$$\text{Span}(\{\uparrow, \downarrow\}) = ?$$

# Matrix-Vector Product

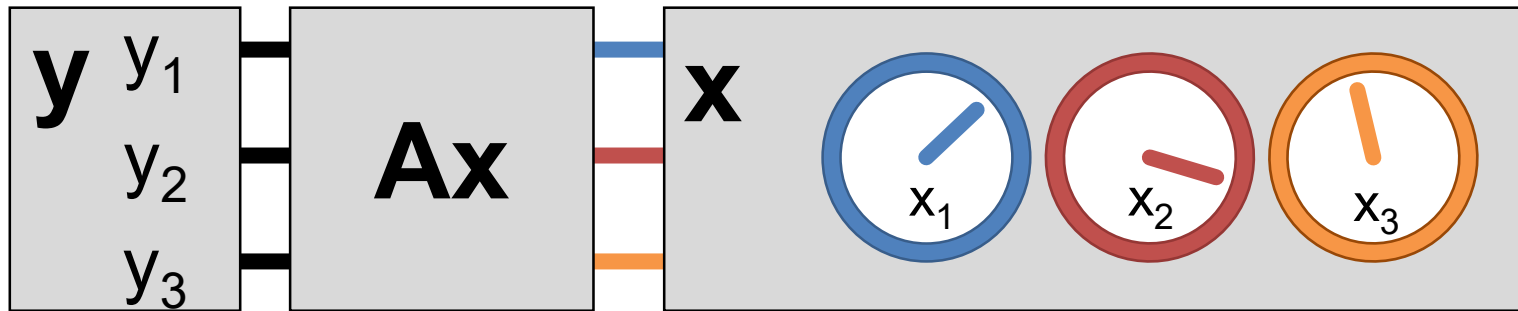
$$\mathbf{Ax} = \begin{bmatrix} | & & | \\ \mathbf{c}_1 & \cdots & \mathbf{c}_n \\ | & & | \end{bmatrix} \mathbf{x}$$

Right-multiplying  $\mathbf{A}$  by  $\mathbf{x}$   
mixes columns of  $\mathbf{A}$   
according to entries of  $\mathbf{x}$

- The output space of  $f(\mathbf{x}) = \mathbf{Ax}$  is constrained to be the *span* of the columns of  $\mathbf{A}$ .
- Can't output things you can't construct out of your columns

# An Intuition

$$\mathbf{y} = \mathbf{A}\mathbf{x} = \begin{bmatrix} | & | & | \\ \mathbf{c}_1 & \mathbf{c}_2 & \mathbf{c}_n \\ | & | & | \end{bmatrix} \begin{bmatrix} x_1 \\ x_2 \\ x_3 \end{bmatrix}$$



- $\mathbf{x}$  – knobs on machine (e.g., fuel, brakes)
- $\mathbf{y}$  – state of the world (e.g., where you are)
- $\mathbf{A}$  – machine (e.g., your car)

# Linear Independence

Suppose the columns of 3x3 matrix **A** are *not* linearly independent ( $c_1, \alpha c_1, c_2$  for instance)

$$\mathbf{y} = \mathbf{Ax} = \begin{bmatrix} | & | & | \\ \mathbf{c}_1 & \alpha\mathbf{c}_1 & \mathbf{c}_2 \\ | & | & | \end{bmatrix} \begin{bmatrix} x_1 \\ x_2 \\ x_3 \end{bmatrix}$$

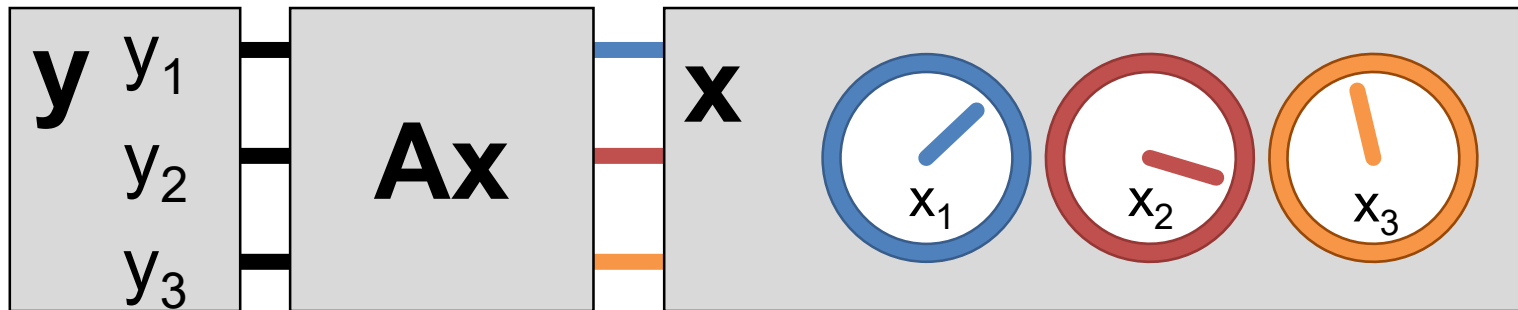
$$\mathbf{y} = x_1\mathbf{c}_1 + \alpha x_2\mathbf{c}_1 + x_3\mathbf{c}_2$$

$$\mathbf{y} = (x_1 + \alpha x_2)\mathbf{c}_1 + x_3\mathbf{c}_2$$

# Linear Independence Intuition

Knobs of  $\mathbf{x}$  are redundant. Even if  $\mathbf{y}$  has 3 outputs, you can only control it in two directions

$$\mathbf{y} = (x_1 + \alpha x_2) \mathbf{c}_1 + x_3 \mathbf{c}_2$$





# Linear Independence

Recall:  $A\mathbf{x} = (x_1 + \alpha x_2)\mathbf{c}_1 + x_3\mathbf{c}_2$

$$\mathbf{y} = \mathbf{A} \begin{bmatrix} x_1 + \beta \\ x_2 - \beta/\alpha \\ x_3 \end{bmatrix} = \left( \cancel{x_1 + \beta} + \alpha x_2 - \alpha \cancel{\frac{\beta}{\alpha}} \right) \mathbf{c}_1 + x_3 \mathbf{c}_2$$

- Can write  $\mathbf{y}$  an infinite number of ways by adding  $\beta$  to  $\mathbf{x}_1$  and subtracting  $\frac{\beta}{\alpha}$  from  $\mathbf{x}_2$
- Or, given a vector  $\mathbf{y}$  there's not a unique vector  $\mathbf{x}$  s.t.  $\mathbf{y} = \mathbf{A}\mathbf{x}$
- Not all  $\mathbf{y}$  have a corresponding  $\mathbf{x}$  s.t.  $\mathbf{y} = \mathbf{A}\mathbf{x}$

# Linear Independence

$$A\mathbf{x} = (x_1 + \alpha x_2)\mathbf{c}_1 + x_3\mathbf{c}_2$$

$$\mathbf{y} = A \begin{bmatrix} \beta \\ -\beta/\alpha \\ 0 \end{bmatrix} = \left( \beta - \alpha \frac{\beta}{\alpha} \right) \mathbf{c}_1 + 0\mathbf{c}_2$$

- What else can we cancel out?
- An infinite number of non-zero vectors  $\mathbf{x}$  can map to a zero-vector  $\mathbf{y}$
- Called the right null-space of  $A$ .

# Rank

- Rank of a  $n \times n$  matrix **A** – number of linearly independent columns (**or rows**) of A / the dimension of the span of the columns
- Matrices with *full rank* ( $n \times n$ , rank  $n$ ) behave nicely: can be inverted, span the full output space, are one-to-one.
- Matrices with *full rank* are machines where every knob is useful and every output state can be made by the machine

# Inverses

- Given  $\mathbf{y} = \mathbf{A}\mathbf{x}$ ,  $\mathbf{y}$  is a linear combination of columns of  $\mathbf{A}$  proportional to  $\mathbf{x}$ . If  $\mathbf{A}$  is full-rank, we should be able to invert this mapping.
- Given some  $\mathbf{y}$  (output) and  $\mathbf{A}$ , what  $\mathbf{x}$  (inputs) produced it?
- $\mathbf{x} = \mathbf{A}^{-1}\mathbf{y}$
- Note: if you don't need to compute it, **never ever compute it**. Solving for  $\mathbf{x}$  is much faster and stable than obtaining  $\mathbf{A}^{-1}$ .

# Symmetric Matrices

- Symmetric:  $A^T = A$  or  $A_{ij} = A_{ji}$
- Have **lots** of special properties

$$\begin{bmatrix} a_{11} & a_{12} & a_{13} \\ a_{21} & a_{22} & a_{23} \\ a_{31} & a_{32} & a_{33} \end{bmatrix}$$

Any matrix of the form  $A = \mathbf{X}^T \mathbf{X}$  is symmetric.

Quick check:

$$A^T = (\mathbf{X}^T \mathbf{X})^T$$
$$A^T = \mathbf{X}^T (\mathbf{X}^T)^T$$
$$A^T = \mathbf{X}^T \mathbf{X}$$

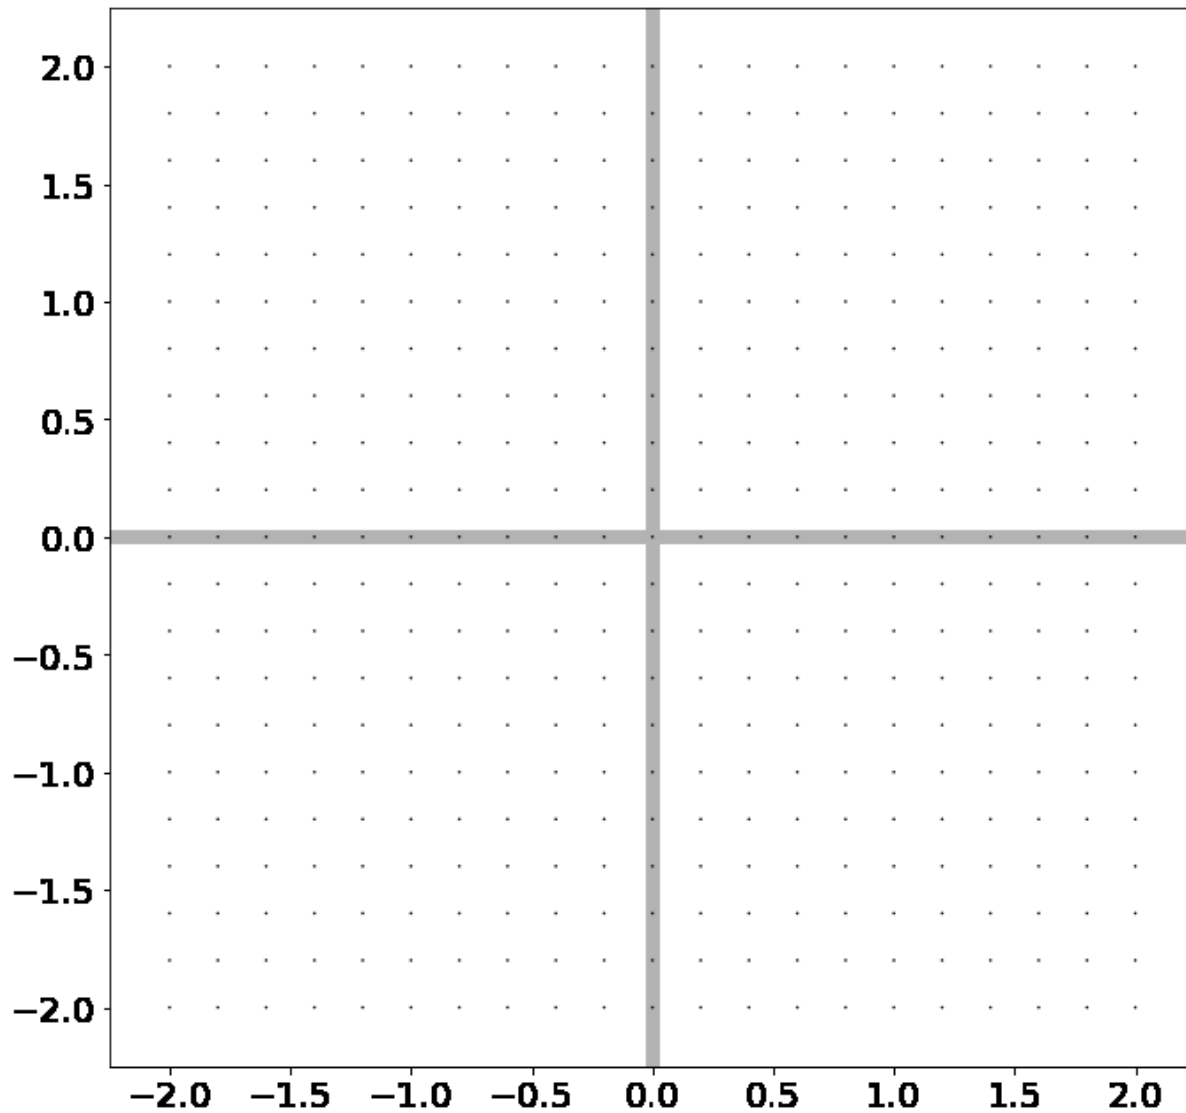
# Special Matrices – Rotations

$$\begin{bmatrix} r_{11} & r_{12} & r_{13} \\ r_{21} & r_{22} & r_{23} \\ r_{31} & r_{32} & r_{33} \end{bmatrix}$$

- Rotation matrices  $R$  rotate vectors and ***do not change vector L2 norms*** ( $\|R\mathbf{x}\|_2 = \|\mathbf{x}\|_2$ )
- Every row/column is unit norm
- Every row is linearly independent
- Transpose is inverse  $RR^T = R^T R = I$
- Determinant is 1 (otherwise it's also a coordinate flip/reflection), eigenvalues are 1

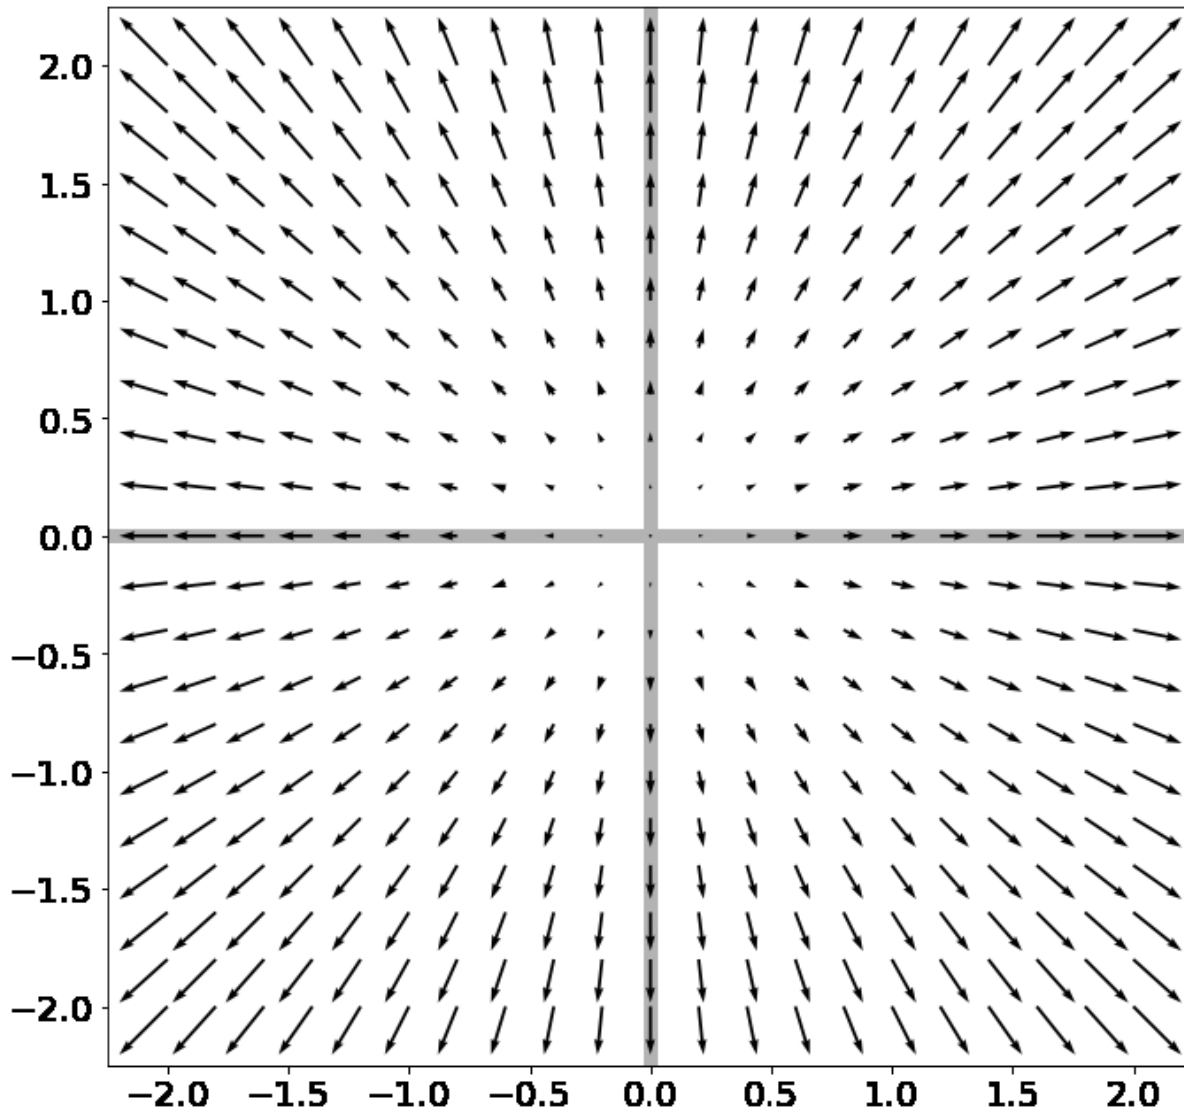
# Eigensystems

- An eigenvector  $\mathbf{v}_i$  and eigenvalue  $\lambda_i$  of a matrix  $A$  satisfy  $A\mathbf{v}_i = \lambda_i\mathbf{v}_i$  ( $A\mathbf{v}_i$  is scaled by  $\lambda_i$ )
- Vectors and values are always paired and typically you assume  $\|\mathbf{v}_i\|^2 = 1$
- Biggest eigenvalue of  $A$  gives bounds on how much  $f(\mathbf{x}) = A\mathbf{x}$  stretches a vector  $\mathbf{x}$ .
- Hints of what people really mean:
  - “Largest eigenvector” = vector w/ largest value
  - Spectral just means there’s eigenvectors



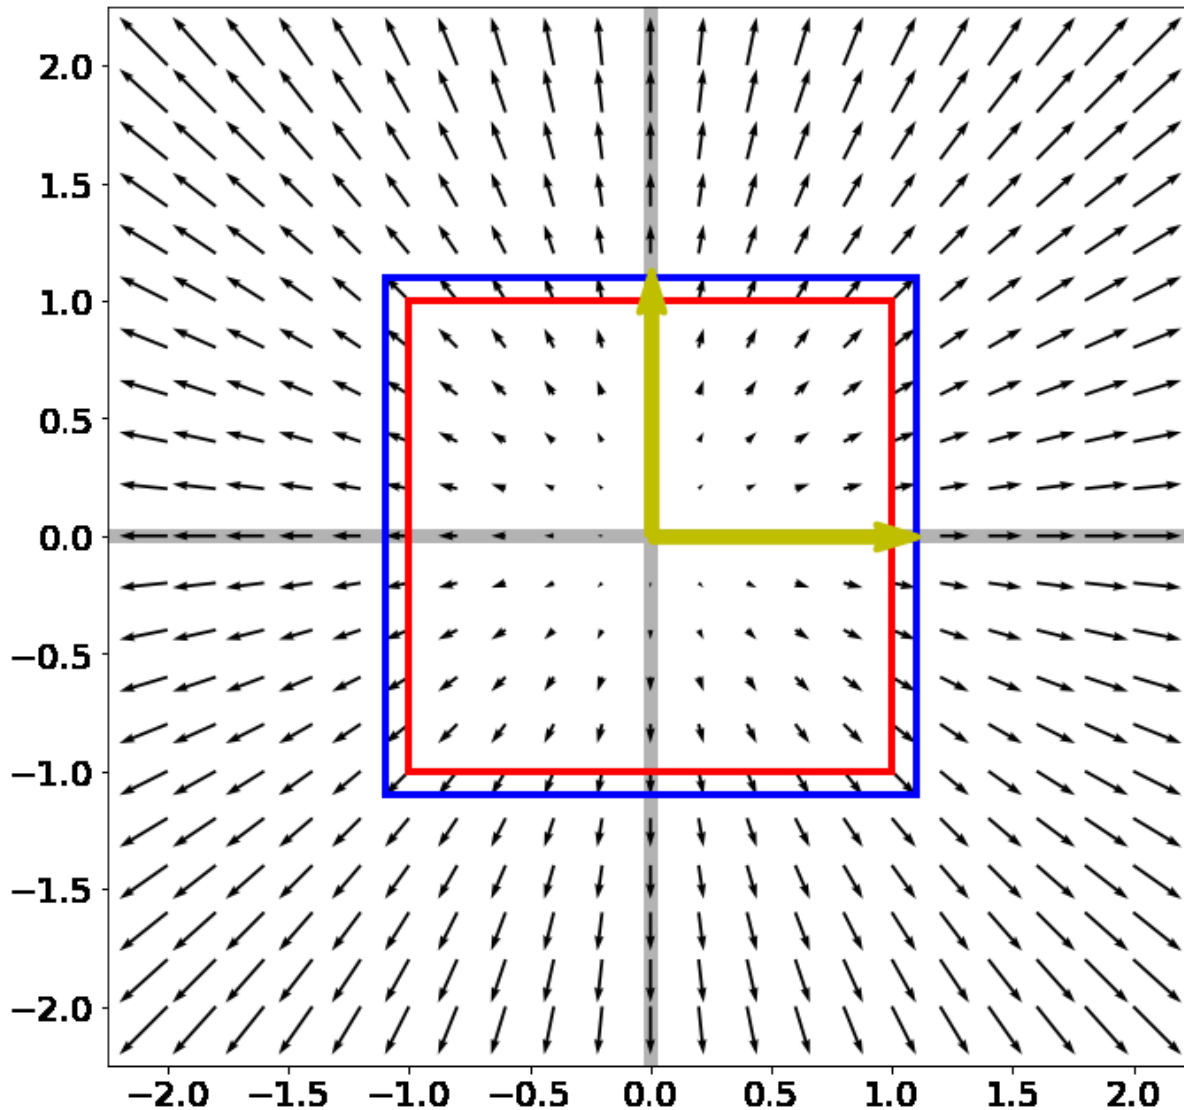
Suppose I have points in a grid





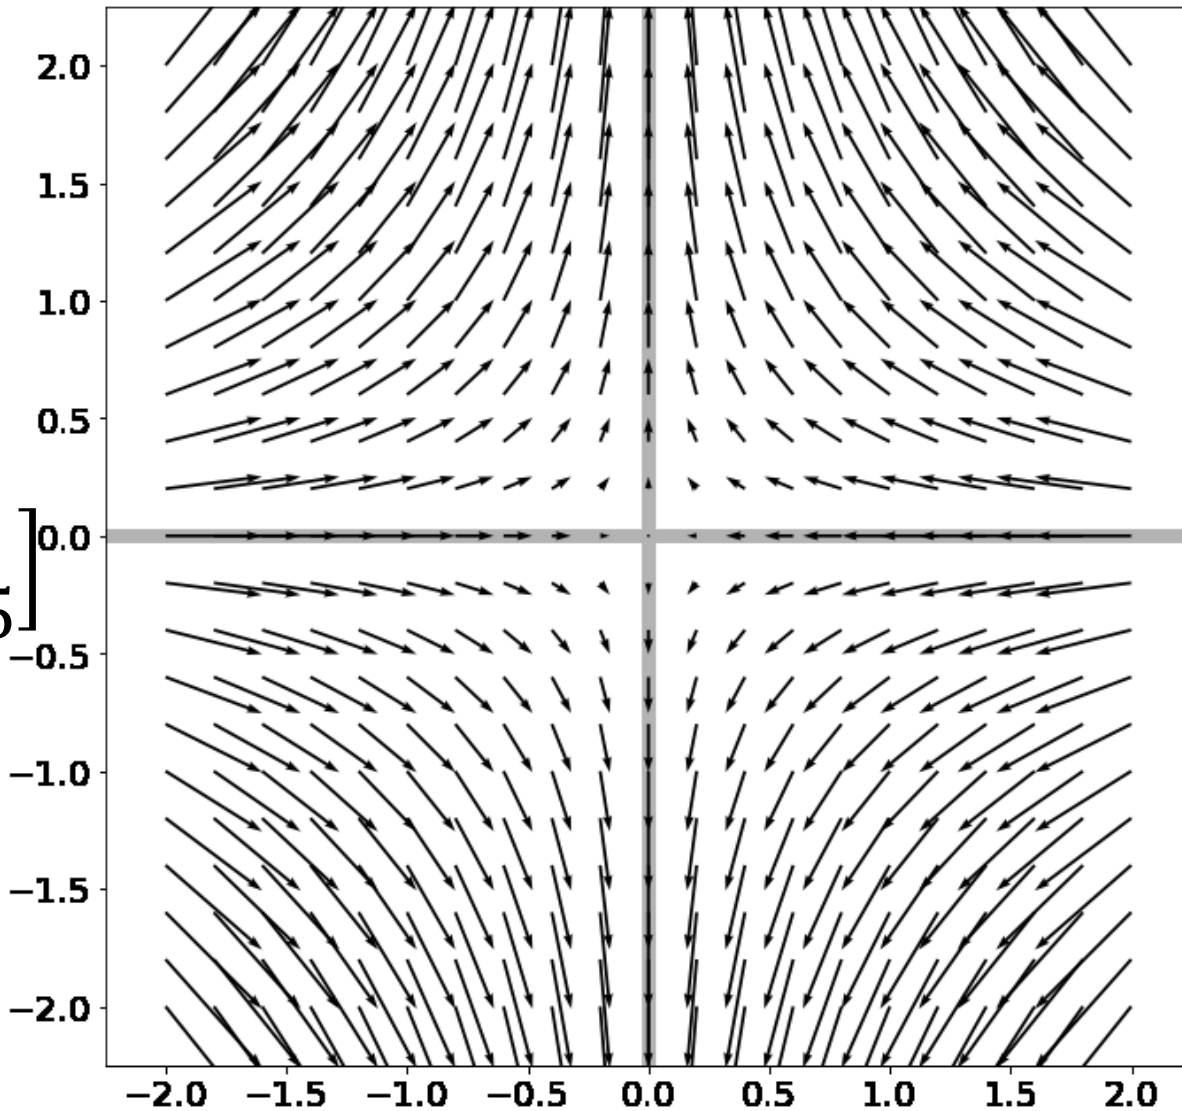
Now I apply  $f(\mathbf{x}) = \mathbf{A}\mathbf{x}$  to these points  
 Pointy-end:  $\mathbf{A}\mathbf{x}$  . Non-Pointy-End:  $\mathbf{x}$

$$\mathbf{A} = \begin{bmatrix} 1.1 & 0 \\ 0 & 1.1 \end{bmatrix}$$



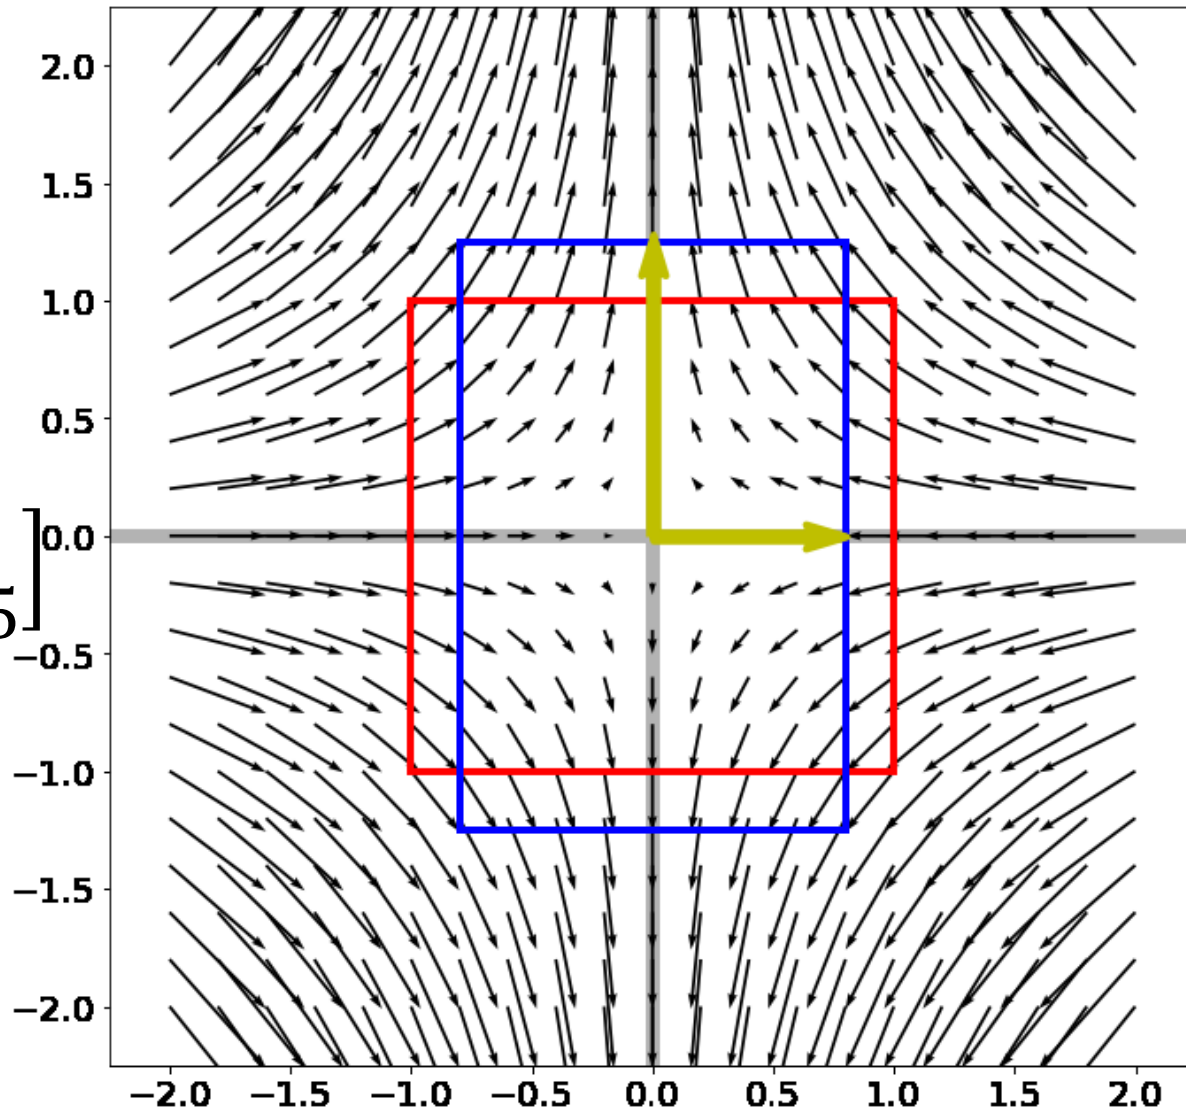
Red box – unit square, Blue box – after  $f(\mathbf{x}) = \mathbf{A}\mathbf{x}$ .  
**What are the yellow lines and why?**

$$\mathbf{A} = \begin{bmatrix} 0.8 & 0 \\ 0 & 1.25 \end{bmatrix}$$

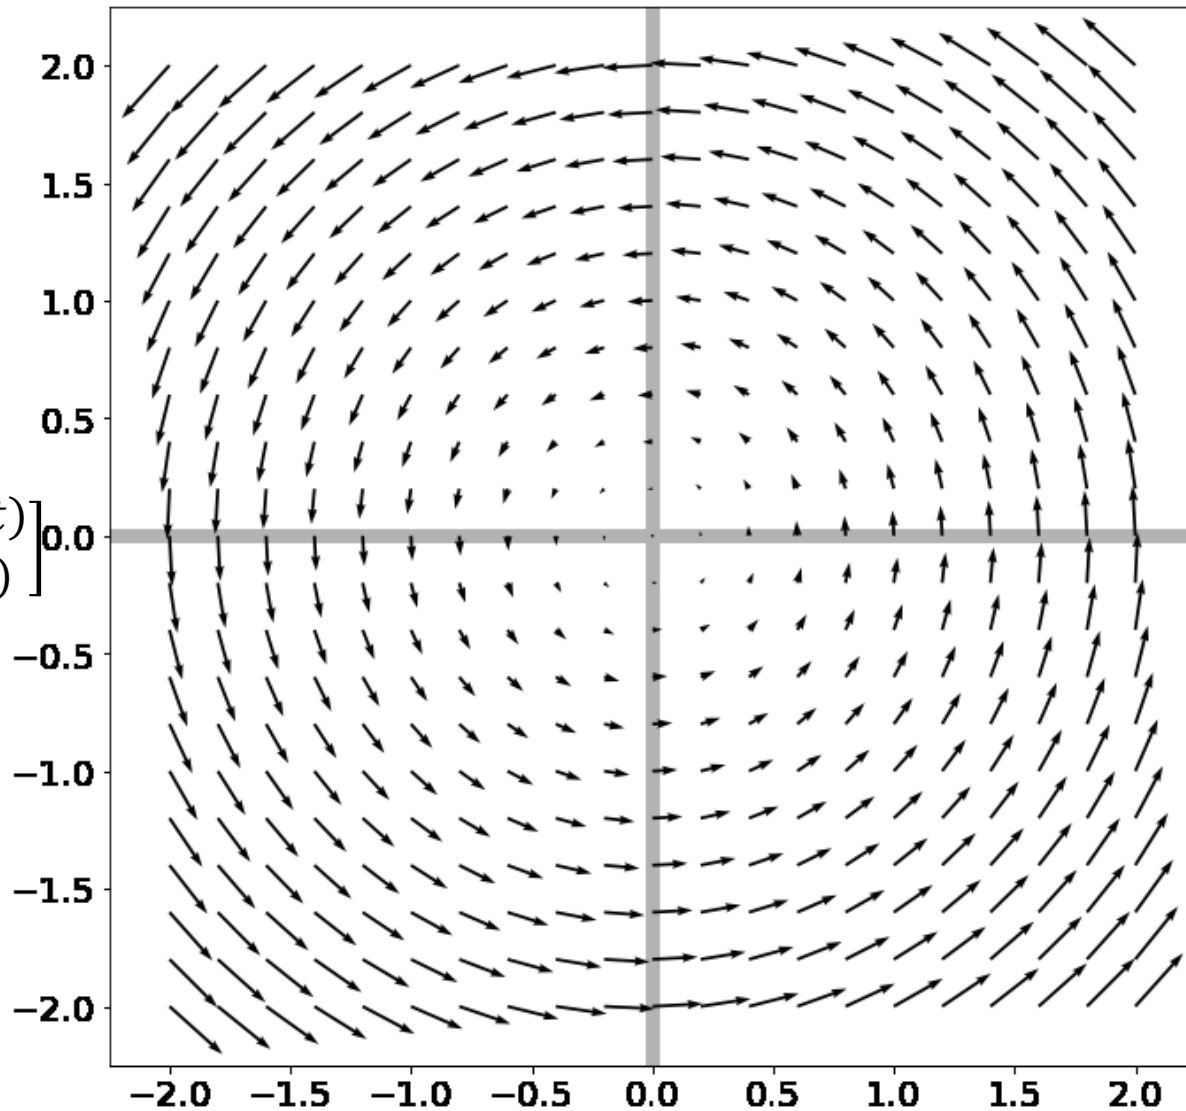


Now I apply  $f(\mathbf{x}) = \mathbf{A}\mathbf{x}$  to these points  
Pointy-end:  $\mathbf{A}\mathbf{x}$  . Non-Pointy-End:  $\mathbf{x}$

$$\mathbf{A} = \begin{bmatrix} 0.8 & 0 \\ 0 & 1.25 \end{bmatrix}$$



Red box – unit square, Blue box – after  $f(\mathbf{x}) = \mathbf{A}\mathbf{x}$ .  
**What are the yellow lines and why?**



$$\mathbf{A} = \begin{bmatrix} \cos(t) & -\sin(t) \\ \sin(t) & \cos(t) \end{bmatrix}$$

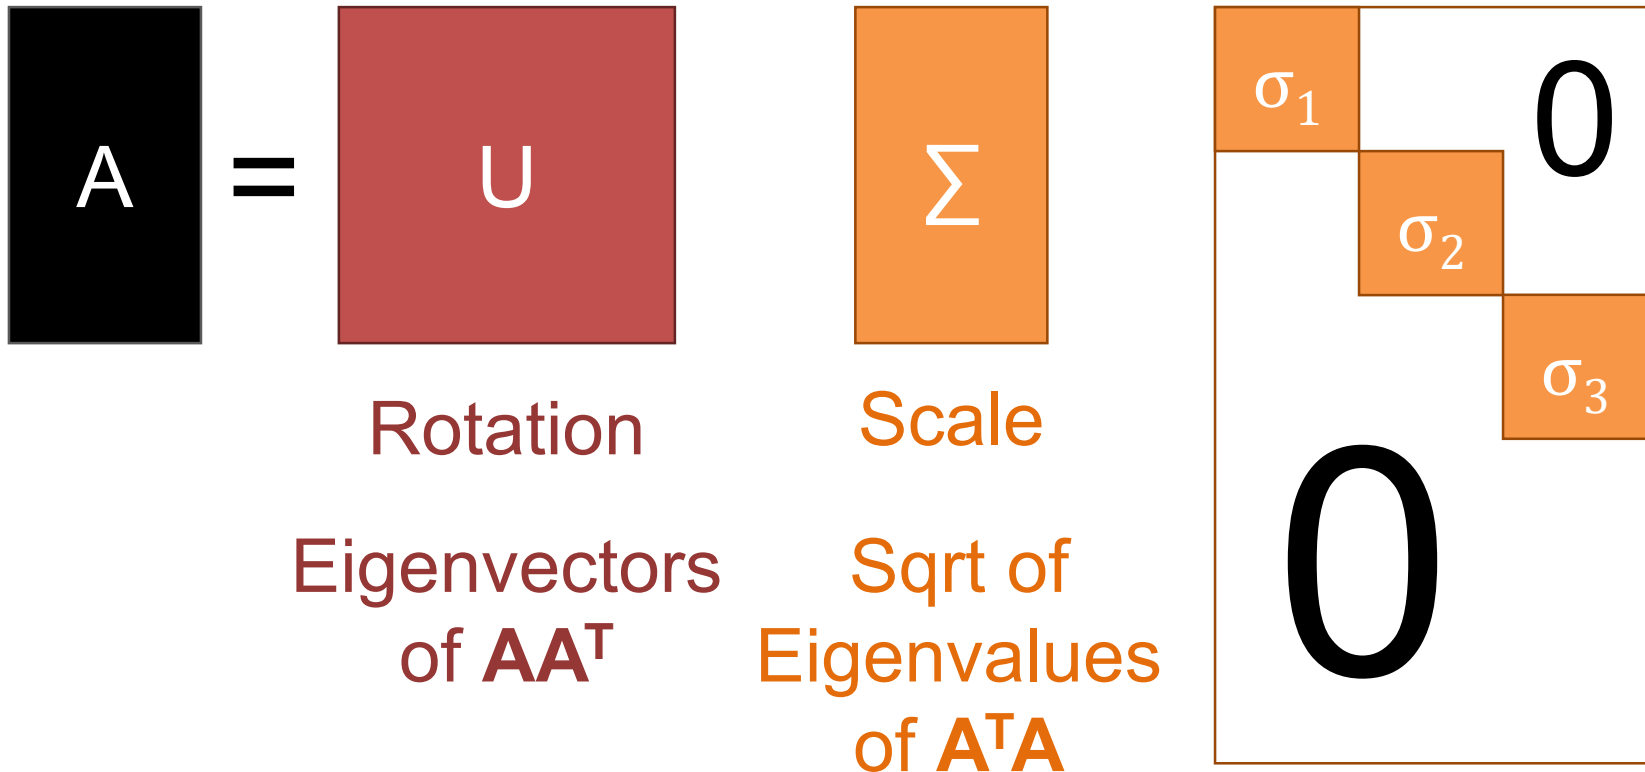
Red box – unit square, Blue box – after  $f(\mathbf{x}) = \mathbf{A}\mathbf{x}$ .  
**Can we draw any yellow lines?**

# Eigenvectors of Symmetric Matrices

- Always  $n$  mutually orthogonal eigenvectors with  $n$  (not necessarily) distinct eigenvalues
- For symmetric  $A$ , the eigenvector with the largest eigenvalue maximizes  $\frac{x^T A x}{x^T x}$   
(smallest/min)
- So for unit vectors (where  $x^T x = 1$ ), that eigenvector maximizes  $x^T A x$
- A surprisingly large number of optimization problems rely on (max/min)imizing this

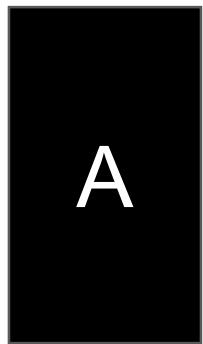
# The Singular Value Decomposition

Can **always** write a  $m \times n$  matrix  $\mathbf{A}$  as:  $\mathbf{A} = \mathbf{U}\mathbf{\Sigma}\mathbf{V}^T$

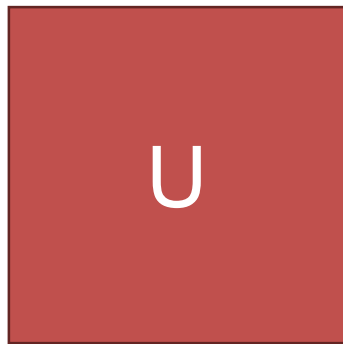


# The Singular Value Decomposition

Can **always** write a  $m \times n$  matrix  $\mathbf{A}$  as:  $\mathbf{A} = \mathbf{U}\mathbf{\Sigma}\mathbf{V}^T$



=



Rotation

Eigenvectors  
of  $\mathbf{A}\mathbf{A}^T$



Scale

Sqrt of  
Eigenvalues  
of  $\mathbf{A}^T\mathbf{A}$



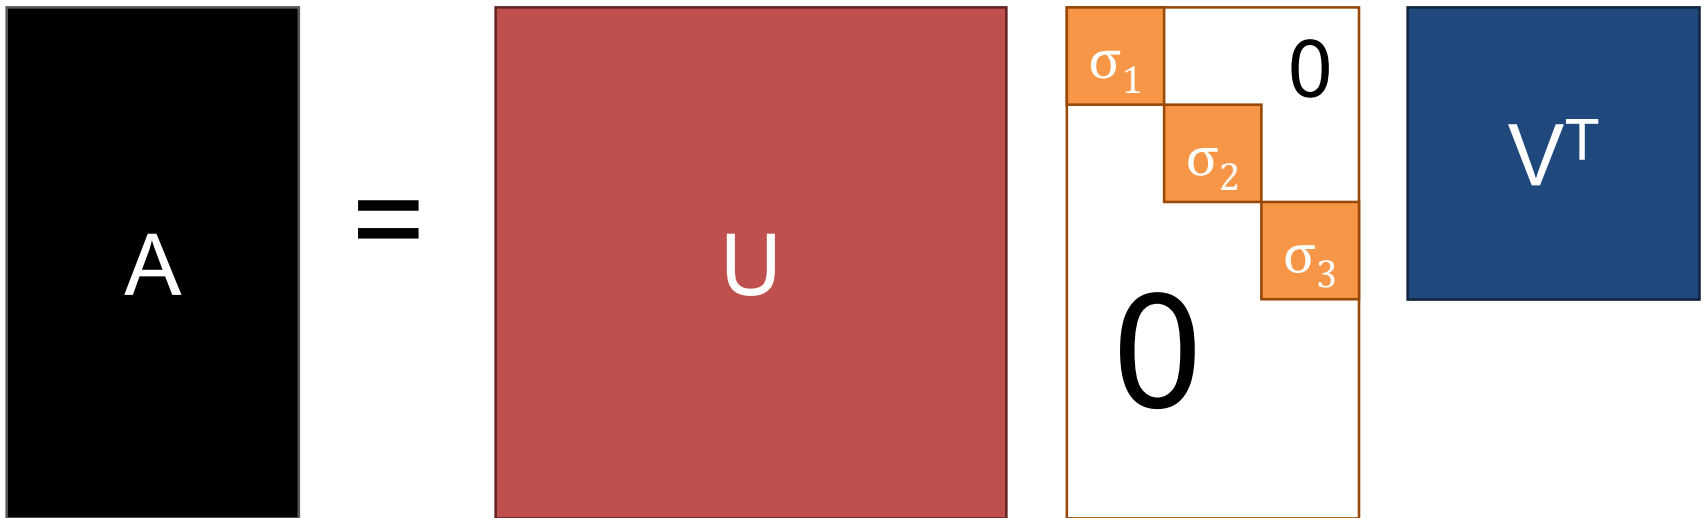
Rotation

Eigenvectors  
of  $\mathbf{A}^T\mathbf{A}$



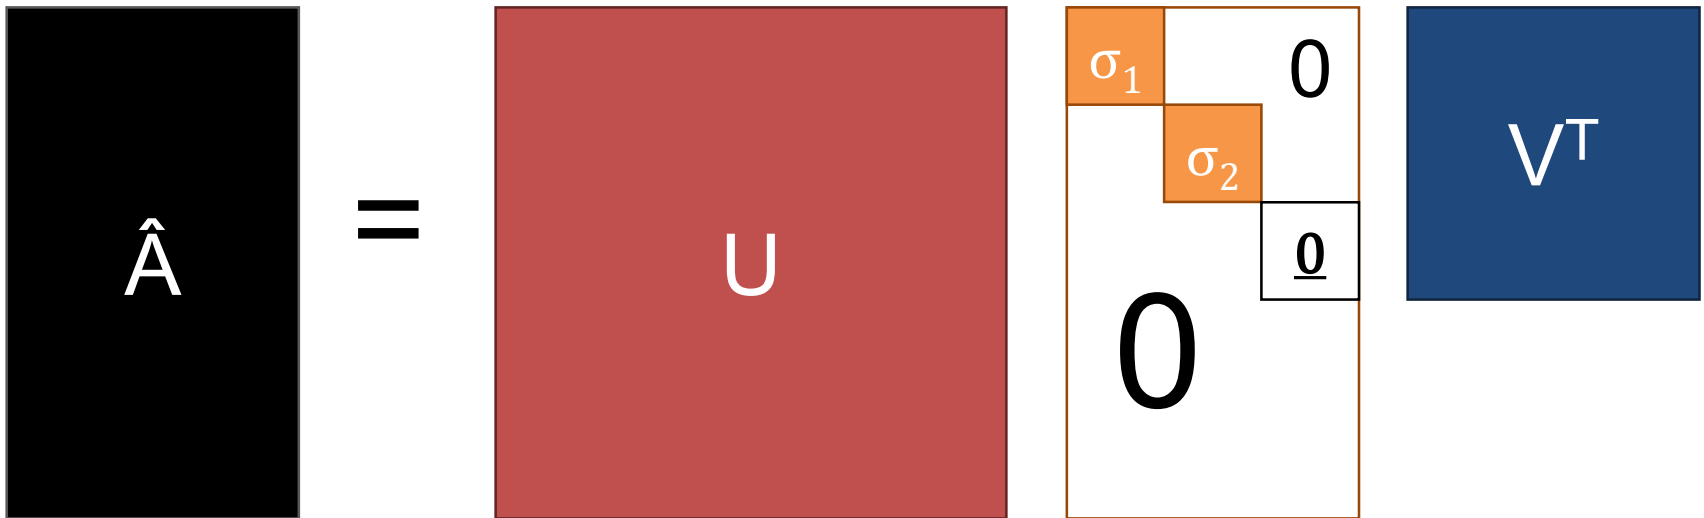
# Singular Value Decomposition

- *Every* matrix is a rotation, scaling, and rotation
- Number of non-zero singular values = rank / number of linearly independent vectors
- “Closest” matrix to  $\mathbf{A}$  with a lower rank



# Singular Value Decomposition

- *Every* matrix is a rotation, scaling, and rotation
- Number of non-zero singular values = rank / number of linearly independent vectors
- “Closest” matrix to  $\mathbf{A}$  with a lower rank

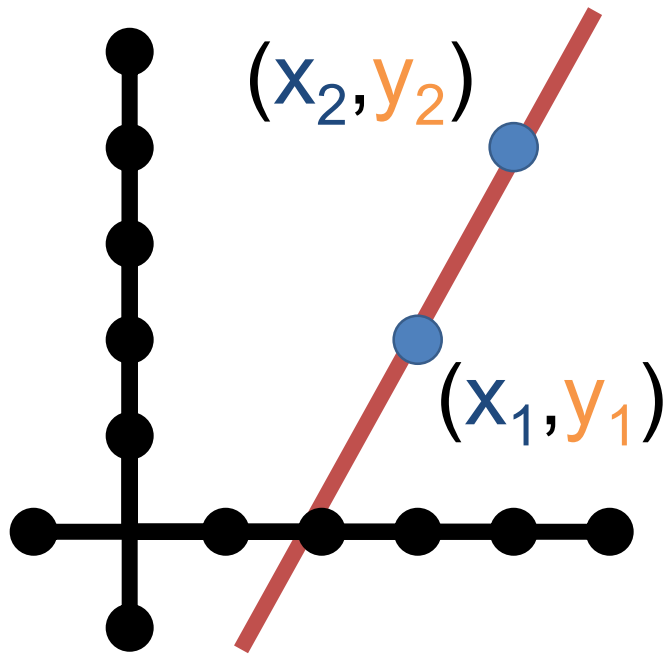


# Singular Value Decomposition

- *Every* matrix is a rotation, scaling, and rotation
- Number of non-zero singular values = rank / number of linearly independent vectors
- “Closest” matrix to  $\mathbf{A}$  with a lower rank
- Secretly behind basically many things you do with matrices



# Solving Least-Squares



Start with two points  $(x_i, y_i)$

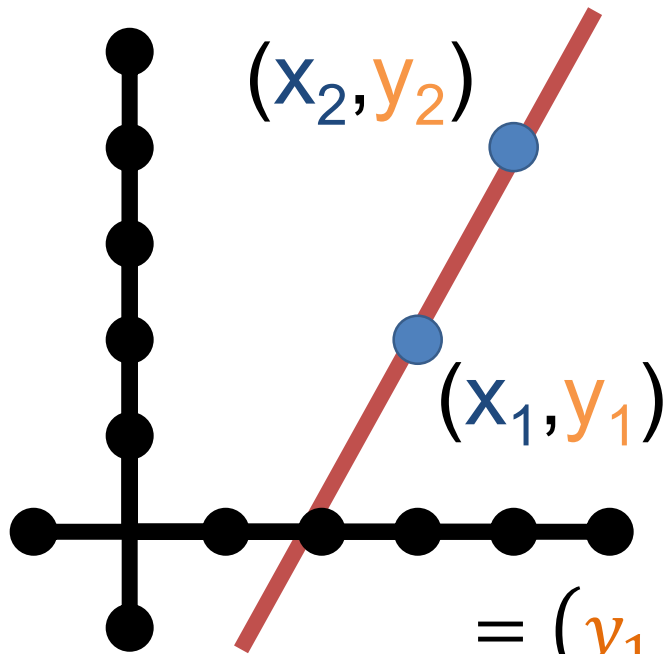
$$y = Av$$

$$\begin{bmatrix} y_1 \\ y_2 \end{bmatrix} = \begin{bmatrix} x_1 & 1 \\ x_2 & 1 \end{bmatrix} \begin{bmatrix} m \\ b \end{bmatrix}$$

$$\begin{bmatrix} y_1 \\ y_2 \end{bmatrix} = \begin{bmatrix} mx_1 + b \\ mx_2 + b \end{bmatrix}$$

We know how to solve this – invert  $A$  and find  $v$  (i.e.,  $(m, b)$  that fits points)

# Solving Least-Squares



Start with two points  $(x_i, y_i)$

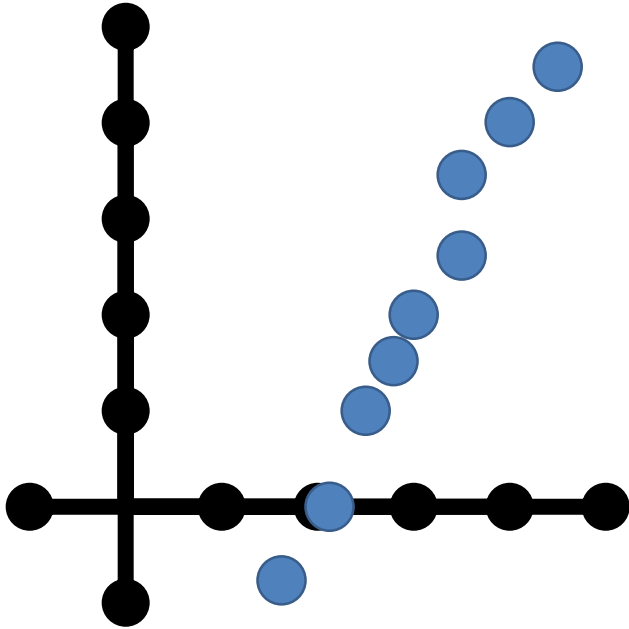
$$y = Av$$

$$\begin{bmatrix} y_1 \\ y_2 \end{bmatrix} = \begin{bmatrix} x_1 & 1 \\ x_2 & 1 \end{bmatrix} \begin{bmatrix} m \\ b \end{bmatrix}$$

$$\begin{aligned} \|y - Av\|^2 &= \left\| \begin{bmatrix} y_1 \\ y_2 \end{bmatrix} - \begin{bmatrix} mx_1 + b \\ mx_2 + b \end{bmatrix} \right\|^2 \\ &= (y_1 - (mx_1 + b))^2 + (y_2 - (mx_2 + b))^2 \end{aligned}$$

The sum of squared differences between **the actual value of  $y$**  and **what the model says  $y$  should be.**

# Solving Least-Squares



Suppose there are  $n > 2$  points

$$\mathbf{y} = \mathbf{A}\mathbf{v}$$

$$\begin{bmatrix} y_1 \\ \vdots \\ y_N \end{bmatrix} = \begin{bmatrix} x_1 & 1 \\ \vdots & \vdots \\ x_N & 1 \end{bmatrix} \begin{bmatrix} m \\ b \end{bmatrix}$$

Compute  $\|\mathbf{y} - \mathbf{A}\mathbf{x}\|^2$  again

$$\|\mathbf{y} - \mathbf{A}\mathbf{v}\|^2 = \sum_{i=1}^n (y_i - (mx_i + b))^2$$

# Solving Least-Squares

Given  $\mathbf{y}$ ,  $\mathbf{A}$ , and  $\mathbf{v}$  with  $\mathbf{y} = \mathbf{A}\mathbf{v}$  overdetermined  
( $\mathbf{A}$  tall / more equations than unknowns)

We want to minimize  $\|\mathbf{y} - \mathbf{A}\mathbf{v}\|^2$ , or find:

$$\boxed{\arg \min_{\mathbf{v}} \|\mathbf{y} - \mathbf{A}\mathbf{v}\|^2}$$

*(The value of  $x$  that makes  
the expression smallest)*

$$\text{Solution satisfies } (\mathbf{A}^T \mathbf{A}) \mathbf{v}^* = \mathbf{A}^T \mathbf{y}$$

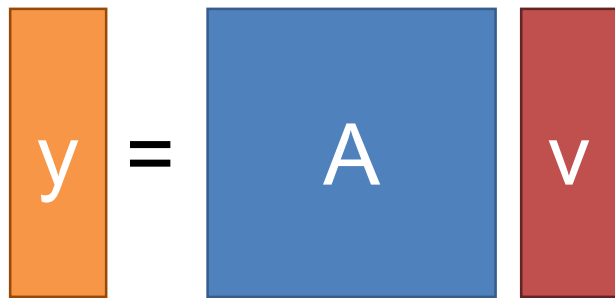
or

$$\mathbf{v}^* = (\mathbf{A}^T \mathbf{A})^{-1} \mathbf{A}^T \mathbf{y}$$

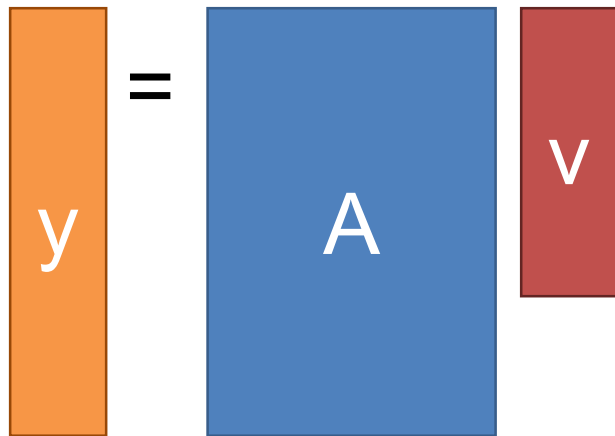
*(Don't actually compute the inverse!)*

# When is Least-Squares Possible?

Given  $\mathbf{y}$ ,  $\mathbf{A}$ , and  $\mathbf{v}$ . Want  $\mathbf{y} = \mathbf{A}\mathbf{v}$



Want  $n$  outputs, have  $n$  knobs to fiddle with, every knob is useful if  $\mathbf{A}$  is full rank.

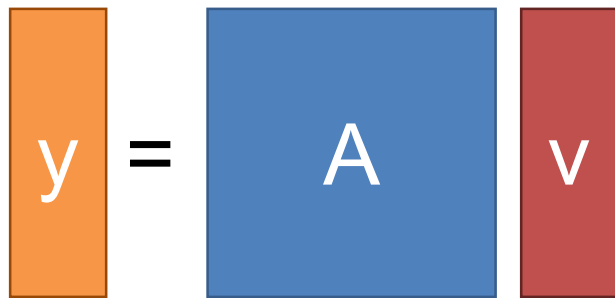


$\mathbf{A}$ : rows (outputs)  $>$  columns (knobs). Thus can't get precise output you want (not enough knobs). So settle for "closest" knob setting.



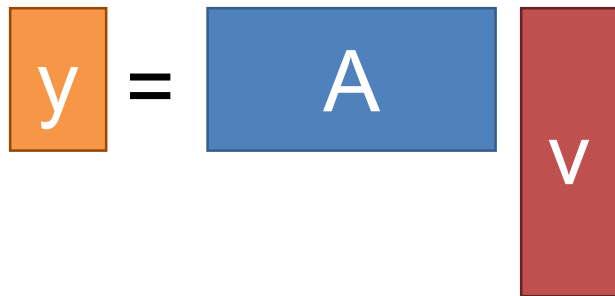
# When is Least-Squares Possible?

Given  $\mathbf{y}$ ,  $\mathbf{A}$ , and  $\mathbf{v}$ . Want  $\mathbf{y} = \mathbf{A}\mathbf{v}$



A diagram representing the equation  $\mathbf{y} = \mathbf{A}\mathbf{v}$ . On the left is a vertical orange rectangle containing the letter 'y'. To its right is an equals sign. Further right is a blue square containing the letter 'A'. To the right of 'A' is a vertical red rectangle containing the letter 'v'.

Want  $n$  outputs, have  $n$  knobs to fiddle with, every knob is useful if  $\mathbf{A}$  is full rank.

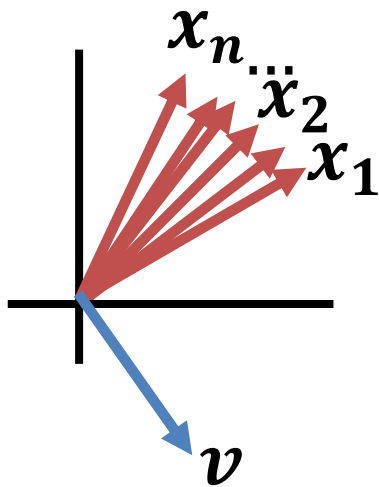


A diagram representing the equation  $\mathbf{y} = \mathbf{A}\mathbf{v}$ . On the left is a vertical orange rectangle containing the letter 'y'. To its right is an equals sign. Further right is a blue horizontal rectangle containing the letter 'A'. To the right of 'A' is a vertical red rectangle containing the letter 'v'.

$\mathbf{A}$ : columns (knobs)  $>$  rows (outputs). Thus, any output can be expressed in infinite ways.

# Homogeneous Least-Squares

Given a set of unit vectors (aka directions)  $x_1, \dots, x_n$  and I want vector  $v$  that is as orthogonal to all the  $x_i$  as possible (for some definition of orthogonal)



Stack  $x_i$  into  $A$ , compute  $Av$

$$Av = \begin{bmatrix} - & x_1^T & - \\ & \vdots & \\ - & x_n^T & - \end{bmatrix} v = \begin{bmatrix} x_1^T v \\ \vdots \\ x_n^T v \end{bmatrix} \begin{matrix} 0 \text{ if} \\ \text{orthog} \end{matrix}$$

$$\text{Compute } \|Av\|^2 = \sum_i (x_i^T v)^2$$

Sum of how orthog.  $v$  is to each  $x$

# Homogeneous Least-Squares

- A lot of times, given a matrix  $\mathbf{A}$  we want to find the  $\mathbf{v}$  that minimizes  $\|\mathbf{A}\mathbf{v}\|^2$ .
- I.e., want  $\mathbf{v}^* = \arg \min_{\mathbf{v}} \|\mathbf{A}\mathbf{v}\|_2^2$
- What's a trivial solution?
- Set  $\mathbf{v} = \mathbf{0} \rightarrow \mathbf{A}\mathbf{v} = \mathbf{0}$
- Exclude this by forcing  $\mathbf{v}$  to have unit norm

# Homogeneous Least-Squares

Let's look at  $\|\mathbf{A}\mathbf{v}\|_2^2$

$$\|\mathbf{A}\mathbf{v}\|_2^2 = \quad \text{Rewrite as dot product}$$

$$\|\mathbf{A}\mathbf{v}\|_2^2 = (\mathbf{A}\mathbf{v})^T (\mathbf{A}\mathbf{v}) \quad \text{Distribute transpose}$$

$$\|\mathbf{A}\mathbf{v}\|_2^2 = \mathbf{v}^T \mathbf{A}^T \mathbf{A} \mathbf{v} = \mathbf{v}^T (\mathbf{A}^T \mathbf{A}) \mathbf{v}$$

We want the vector minimizing this quadratic form

Where have we seen this?

# Homogeneous Least-Squares

Ubiquitous tool in vision:

$$\arg \min_{\|v\|^2=1} \|Av\|^2$$

- (1) “Smallest”\* eigenvector of  $A^T A$   
(2) “Smallest” right singular vector of  $A$

For min → max, switch smallest → largest

\*Note:  $A^T A$  is positive semi-definite so it has all non-negative eigenvalues

# Derivatives

Remember derivatives?

Derivative: rate at which a function  $f(x)$  changes at a point as well as the direction that increases the function

Given quadratic function  $f(x)$

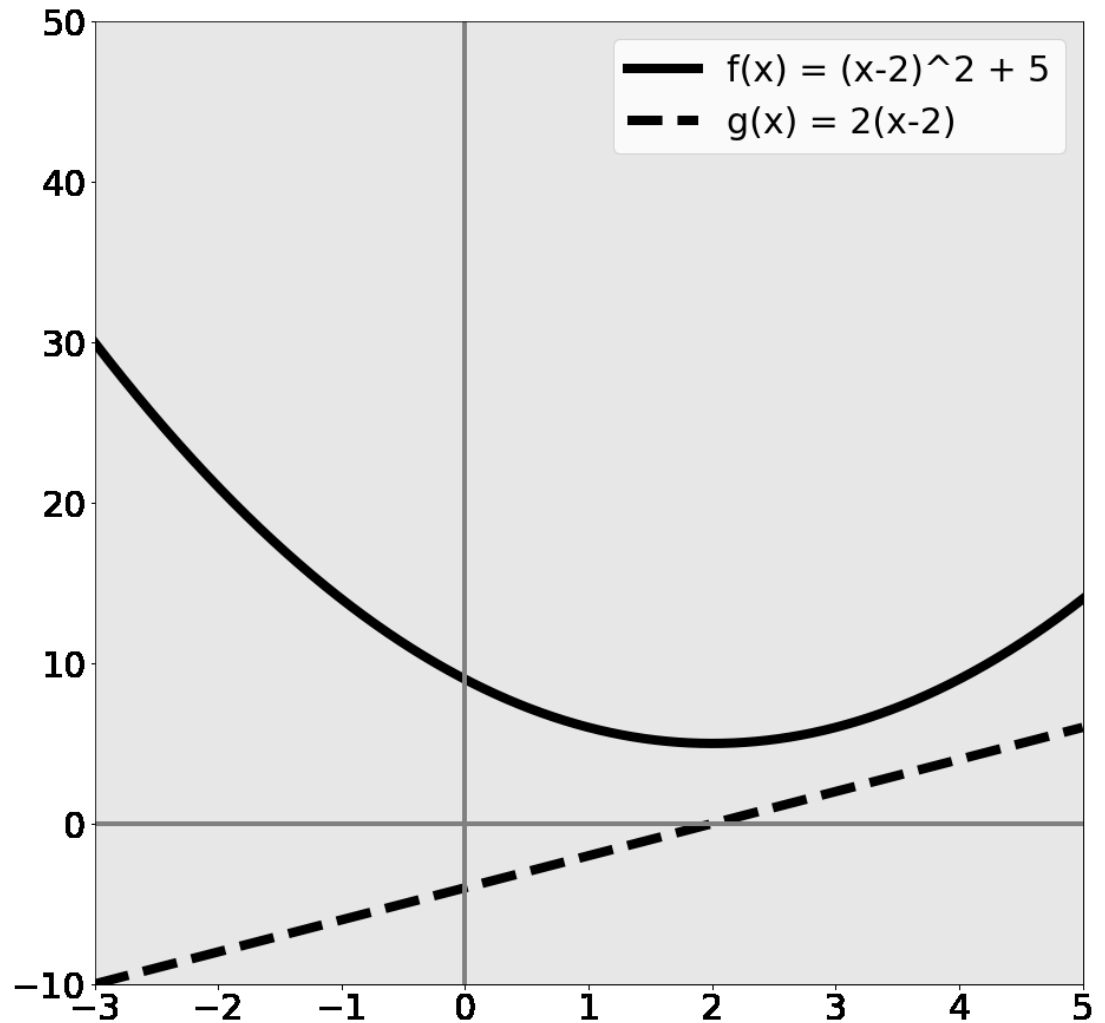
$$f(x, y) = (x - 2)^2 + 5$$

$f(x)$  is function

$$g(x) = f'(x)$$

aka

$$g(x) = \frac{d}{dx} f(x)$$



Given quadratic function  $f(x)$

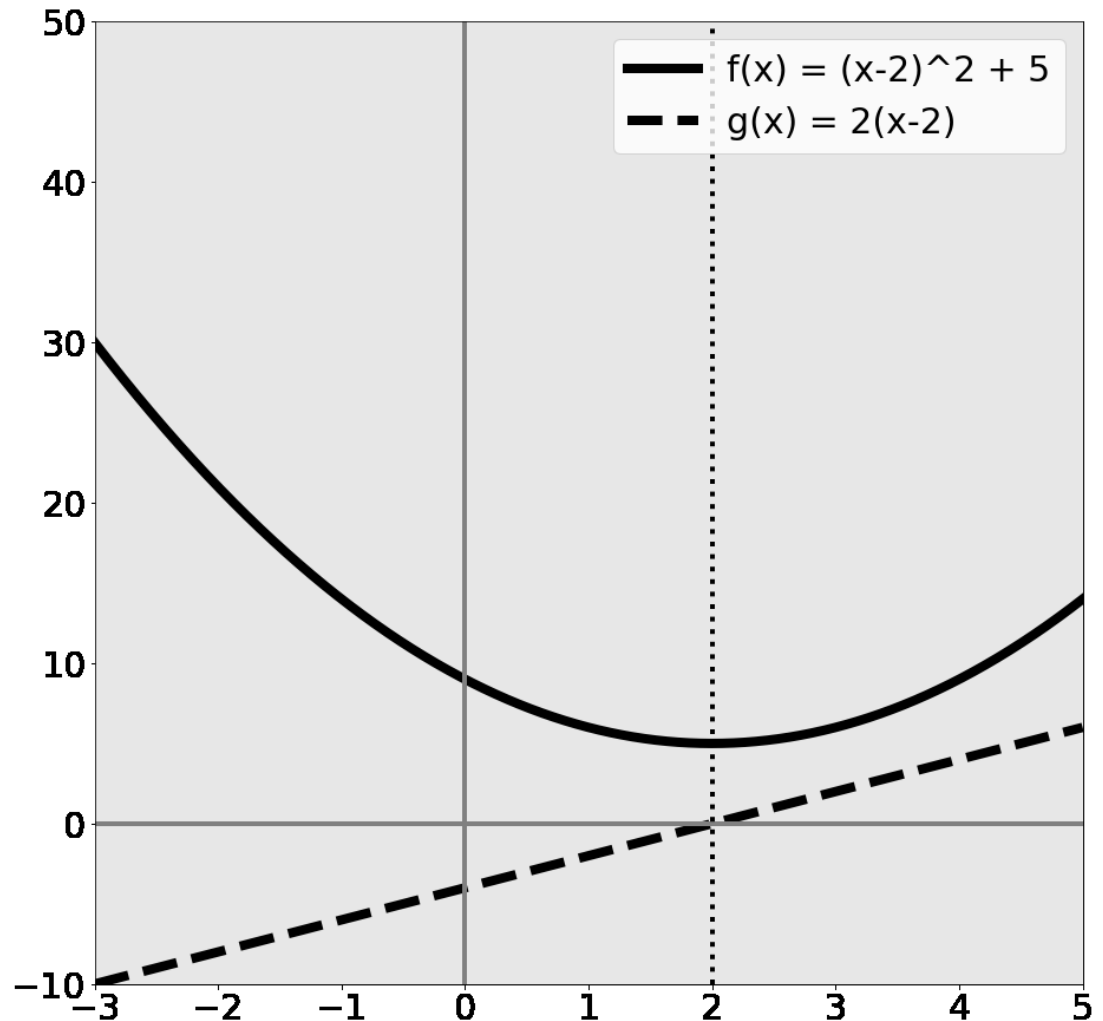
$$f(x, y) = (x - 2)^2 + 5$$

**What's special  
about  $x=2$ ?**

$f(x)$  minim. at 2  
 $g(x) = 0$  at 2

$a = \text{minimum of } f \rightarrow$   
 $g(a) = 0$

Reverse is **not true**





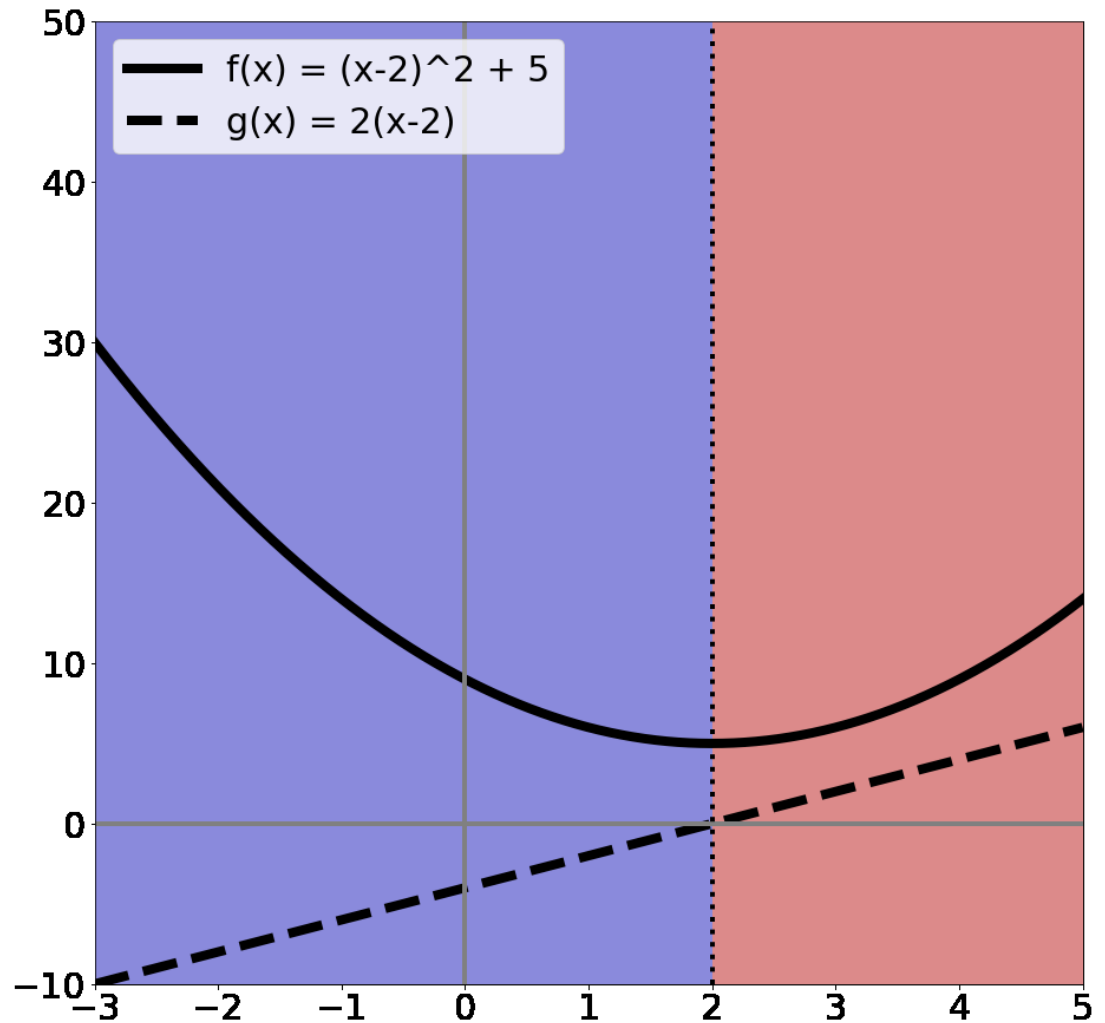
# Rates of change

$$f(x, y) = (x - 2)^2 + 5$$

Suppose I want to increase  $f(x)$  by changing  $x$ :

Blue area: move left  
Red area: move right

Derivative tells you direction of ascent and rate



# What Calculus Should I Know

- Really need intuition
- Need chain rule
- Rest you should look up / use a computer algebra system / use a cookbook
- Partial derivatives (and that's it from multivariable calculus)

# Partial Derivatives

- Pretend other variables are constant, take a derivative. That's it.
- Make our function a function of two variables

$$f(x) = (x - 2)^2 + 5$$

$$\frac{\partial}{\partial x} f(x) = 2(x - 2) * 1 = 2(x - 2)$$

$$f_2(x, y) = (x - 2)^2 + 5 + (y + 1)^2$$

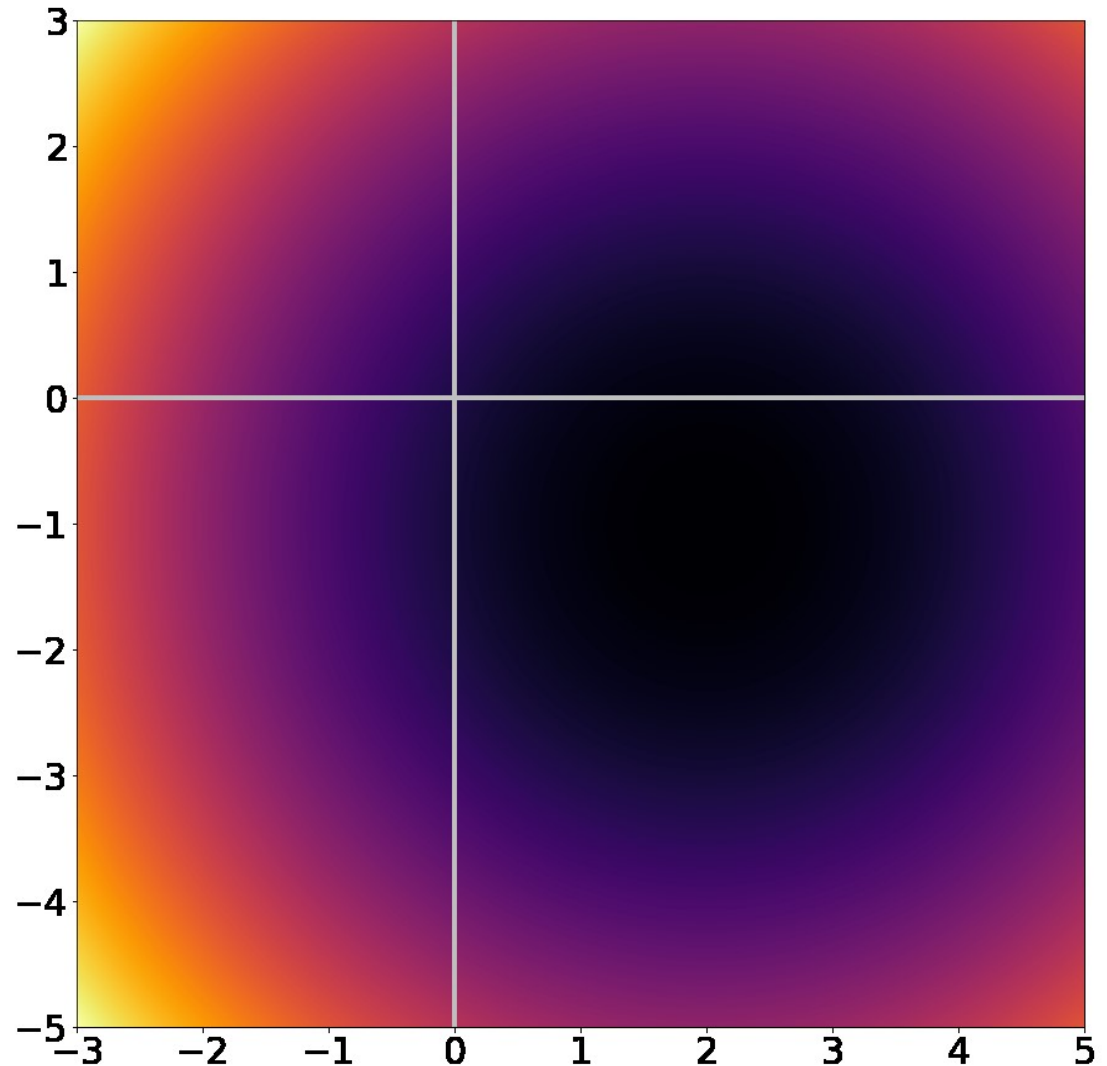
$$\frac{\partial}{\partial x} f_2(x) = 2(x - 2)$$

Pretend it's  
constant  $\rightarrow$   
derivative = 0

# Zooming Out

$$f_2(x, y) = (x - 2)^2 + 5 + (y + 1)^2$$

Dark =  $f(x, y)$  low  
Bright =  $f(x, y)$  high



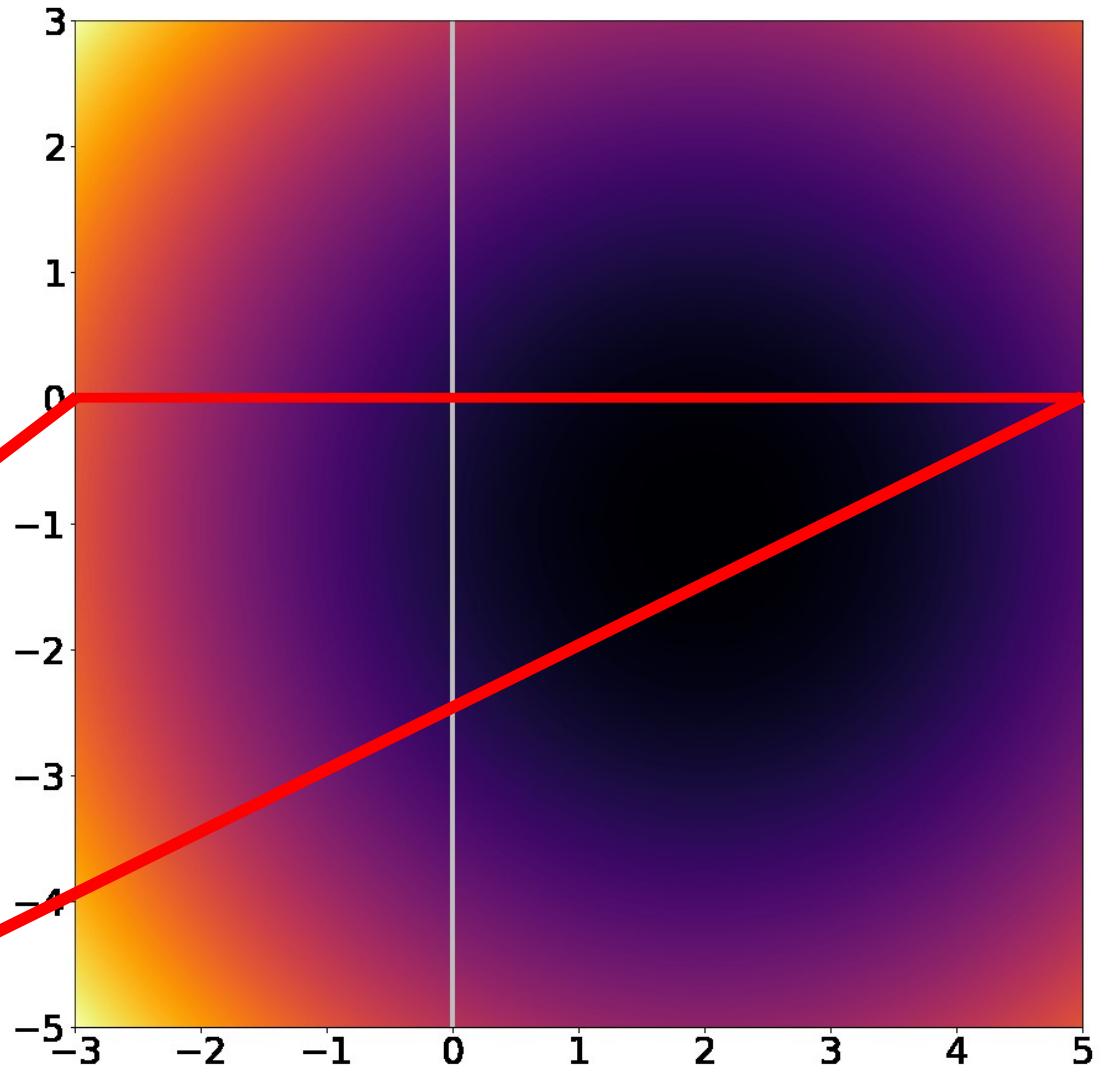
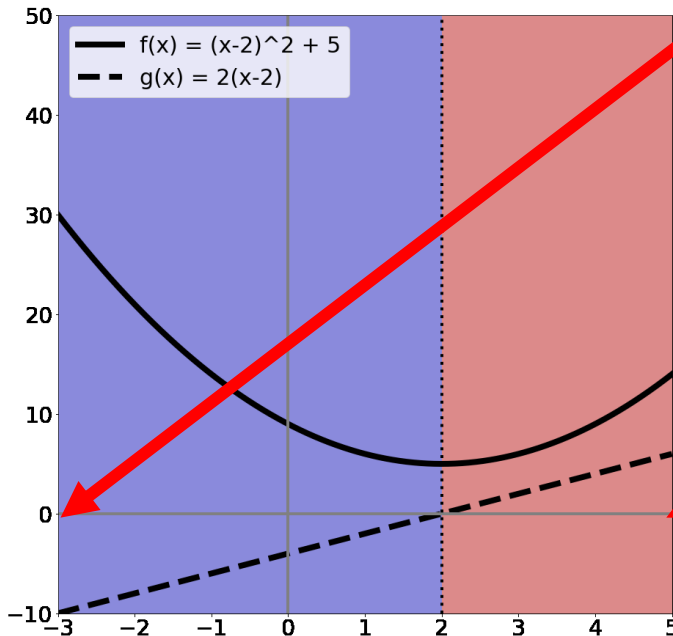
# Taking a slice of

$$f_2(x, y) = (x - 2)^2 + 5 + (y + 1)^2$$

Slice of  $y=0$  is the  
function from before:

$$f(x) = (x - 2)^2 + 5$$

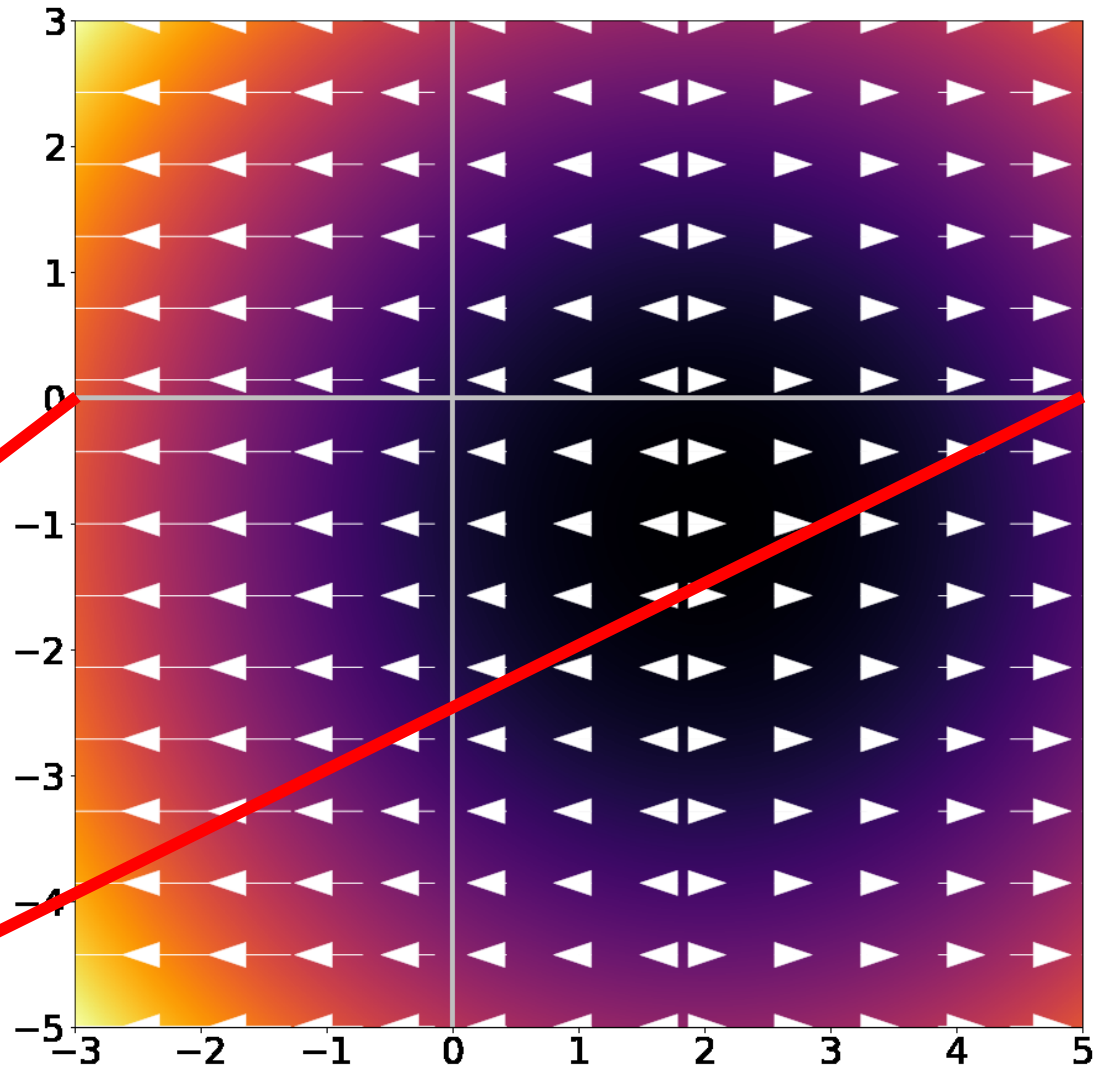
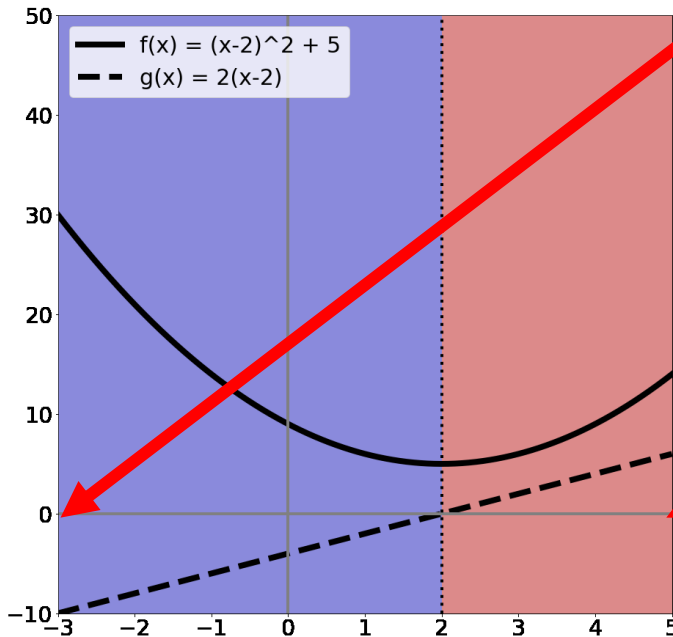
$$f'(x) = 2(x - 2)$$



# Taking a slice of

$$f_2(x, y) = (x - 2)^2 + 5 + (y + 1)^2$$

$\frac{\partial}{\partial x} f_2(x, y)$  is rate of change & direction in x dimension



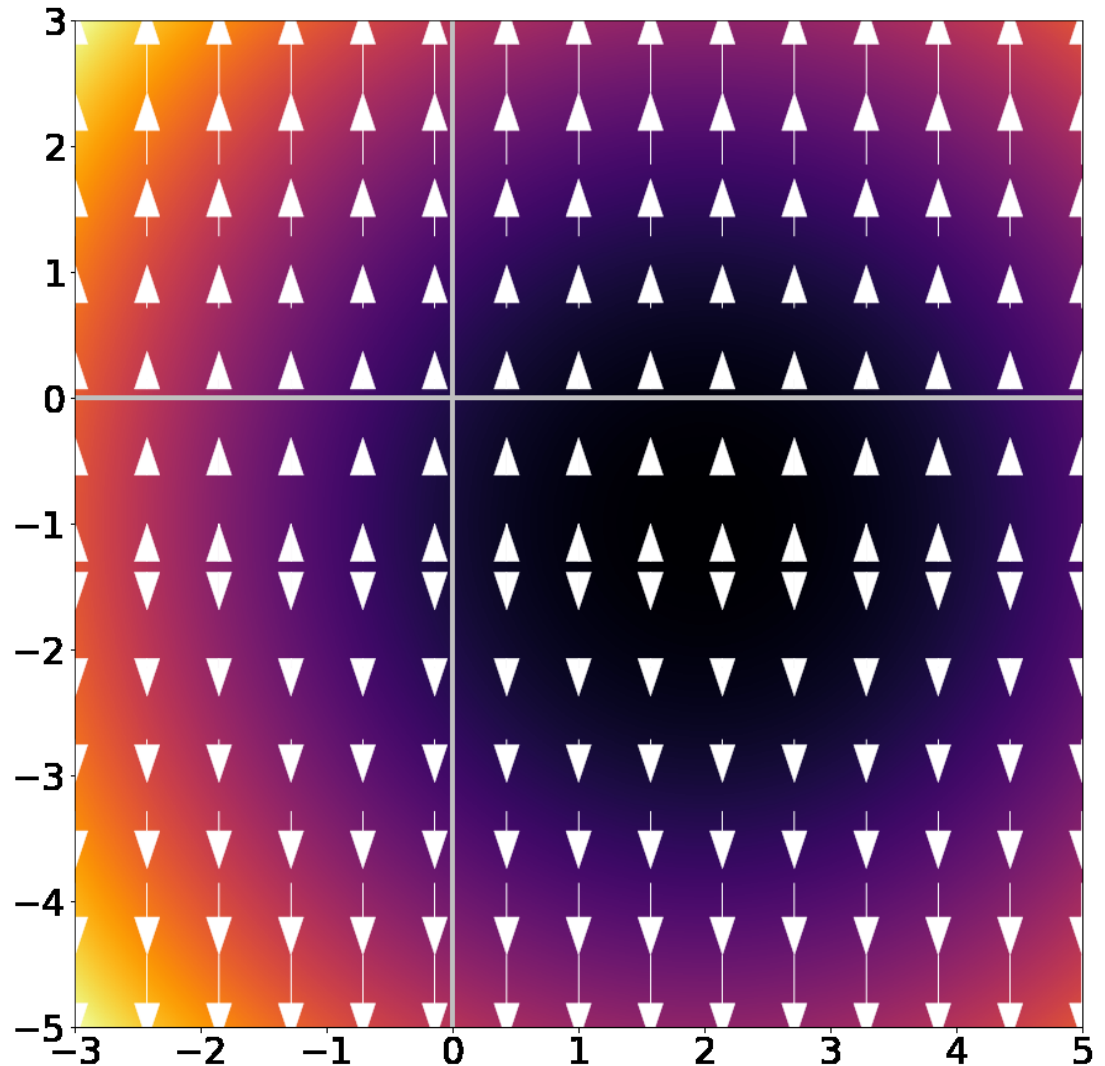
# Zooming Out

$$f_2(x, y) = (x - 2)^2 + 5 + (y + 1)^2$$

$\frac{\partial}{\partial y} f_2(x, y)$  is

$$2(y + 1)$$

and is the rate of  
change & direction in  
y dimension



# Zooming Out

$$f_2(x, y) = (x - 2)^2 + 5 + (y + 1)^2$$

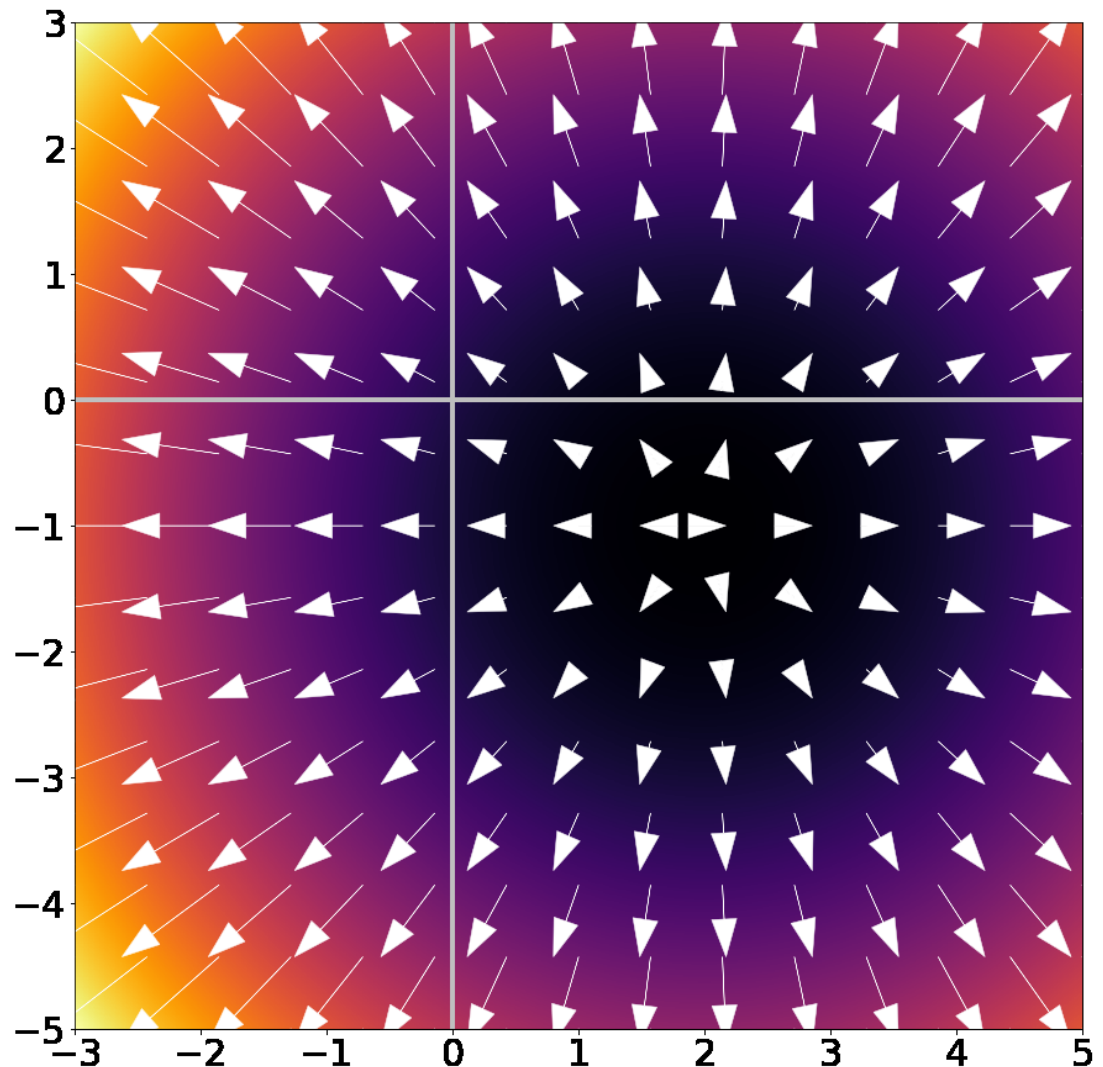
## Gradient/Jacobian:

Making a vector of

$$\nabla_f = \left[ \frac{\partial f}{\partial x}, \frac{\partial f}{\partial y} \right]$$

gives rate and  
direction of change.

Arrows point OUT of  
minimum / basin.





# What Should I Know?

- Gradients are simply partial derivatives per-dimension: if  $\mathbf{x}$  in  $f(\mathbf{x})$  has  $n$  dimensions,  $\nabla_f(\mathbf{x})$  has  $n$  dimensions
- Gradients point in direction of ascent and tell the rate of ascent
- If  $\mathbf{a}$  is minimum of  $f(\mathbf{x}) \rightarrow \nabla_f(\mathbf{a}) = \mathbf{0}$
- Reverse is not true, especially in high-dimensional spaces

