Lecture 1:

Course Intro:
Welcome to Computer Graphics!

Computer Graphics
CMU 15-462/662

Hi!

Nancy Pollard

Daniel

Denise

Shilin

Eve

Alan

TODAY

- Two main objectives
  - Quick overview of computer graphics
  - Launch right in with some of the math we need for this course

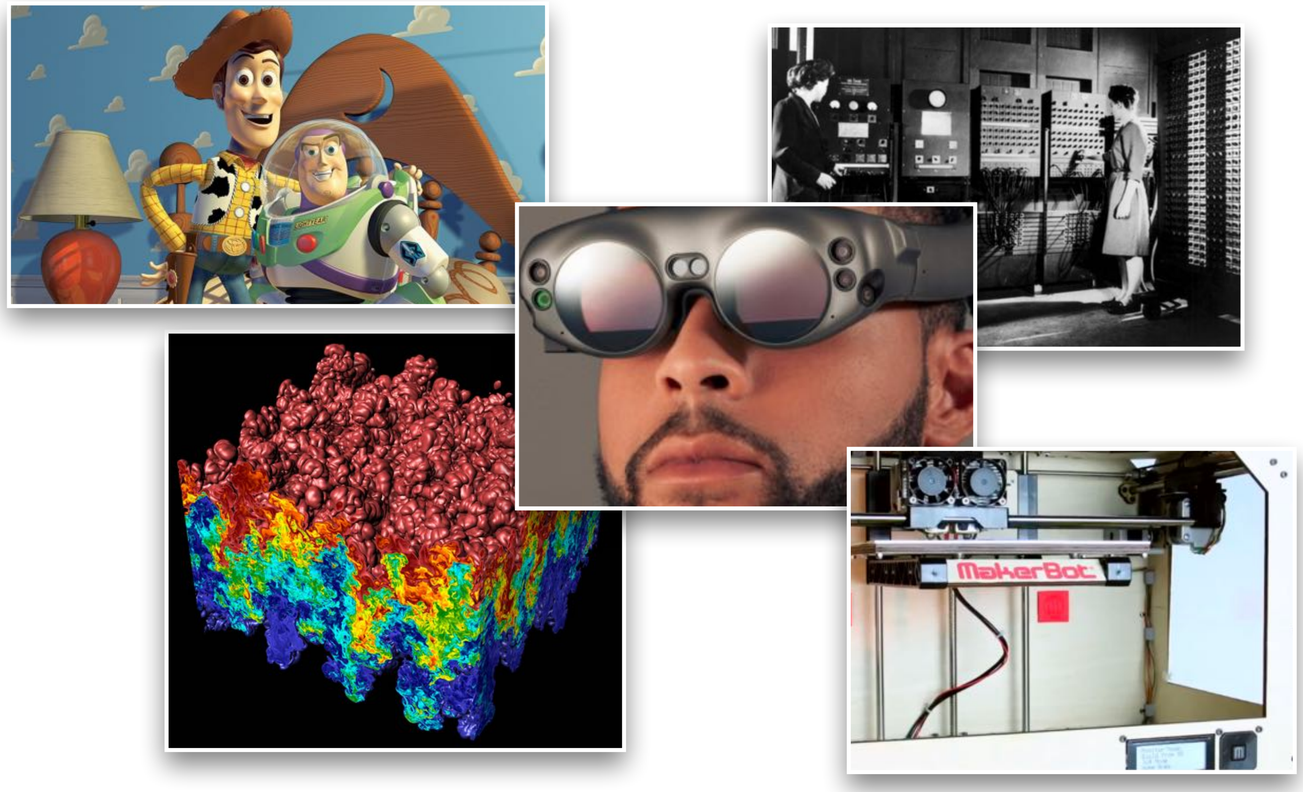

Linear Algebra for Computer Graphics

Q: What is computer graphics?

Probably an image like this comes to mind:

Or this:

2021 virtual reality headset: 2 x 2160 x 2160 @ 90Hz => 2.3GB/s

So what is computer graphics?

com•put•er graph•ics /kəmˈpyo͞odər ˈgrafiks/ *n.*
The use of computers to synthesize visual information.

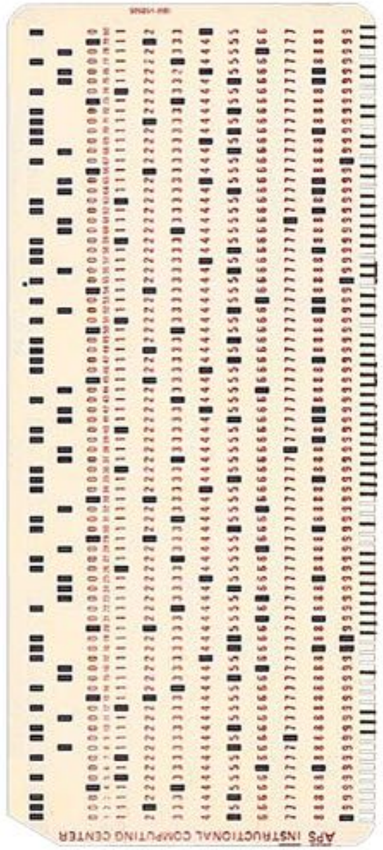

digital information

computation

visual information

Why visual information?

About 30% of brain dedicated to visual processing…

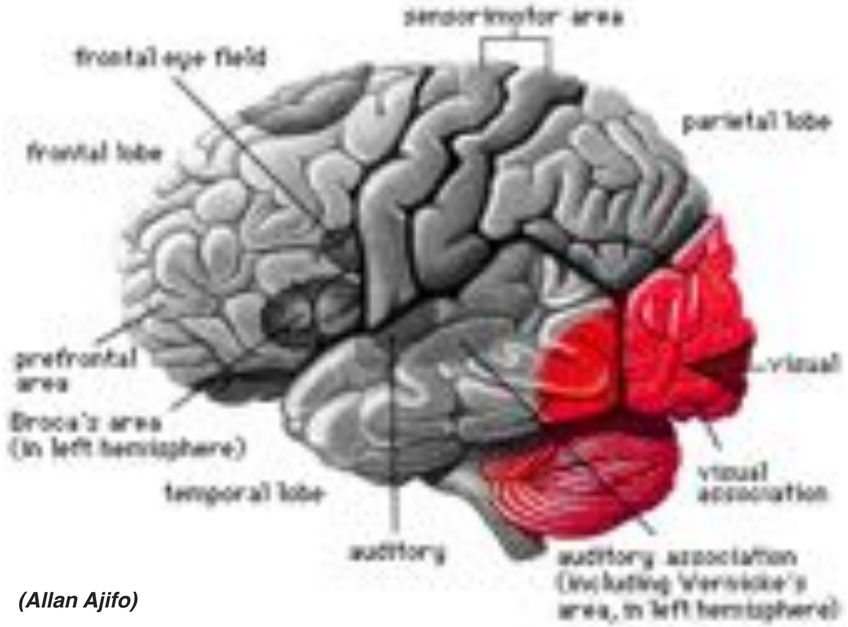

(Allan Ajifo)

(Petar Milošević)

…eyes are highest-bandwidth port into the head!

So what is computer graphics?

com•put•er graph•ics /kəmˈpyoodərˈgrafiks/ *n.*
The use of computers to synthesize visual information.

Why only visual?

visual information

digital information

computation

Graphics has evolved a lot since its early days… no longer just about turning on pixels!

Turning digital information into sensory stimuli

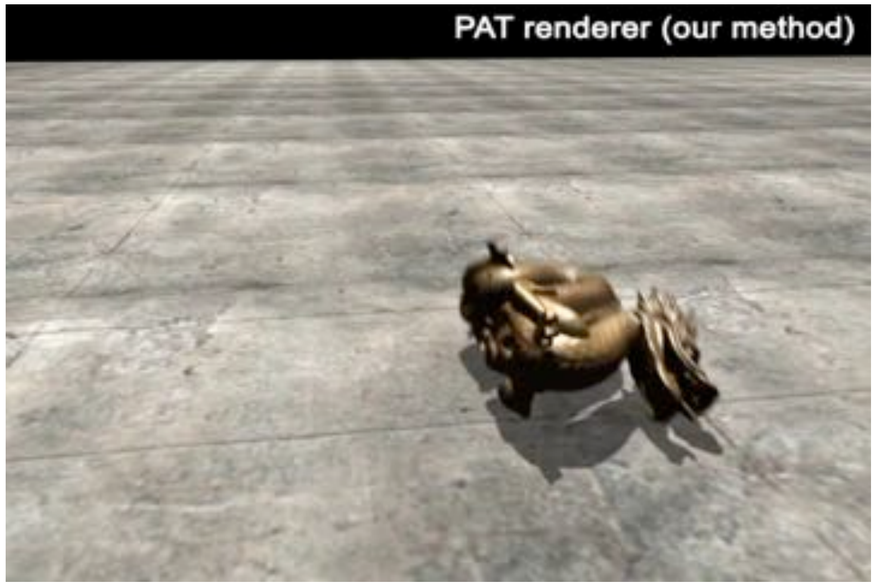

(sound)

(touch)

com•put•er graph•ics /kəmˈpyo͞odər ˈgrafiks/ *n.*
The use of computers to synthesize and manipulate **sensory** information.

(…What about taste? Smell?!)

Turning digital information into physical matter

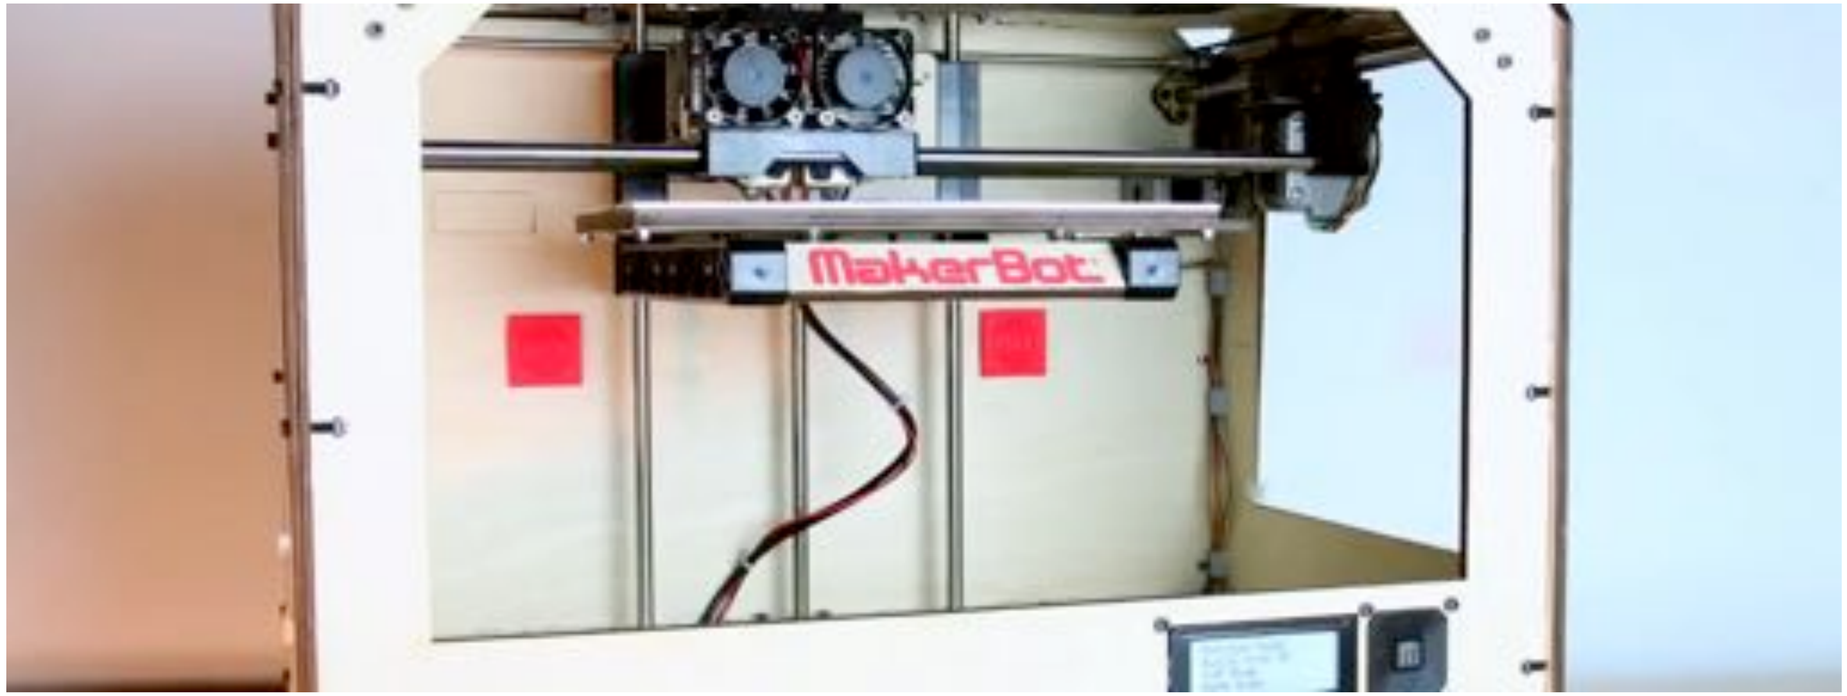

Definition of Graphics, Revisited

com•put•er graph•ics /kəmˈpyo͞odər ˈgrafiks/ *n*.
The use of computation to turn digital information into **sensory stimuli**.

Even this definition is too narrow…

SIGGRAPH 2022 Technical Papers Trailer

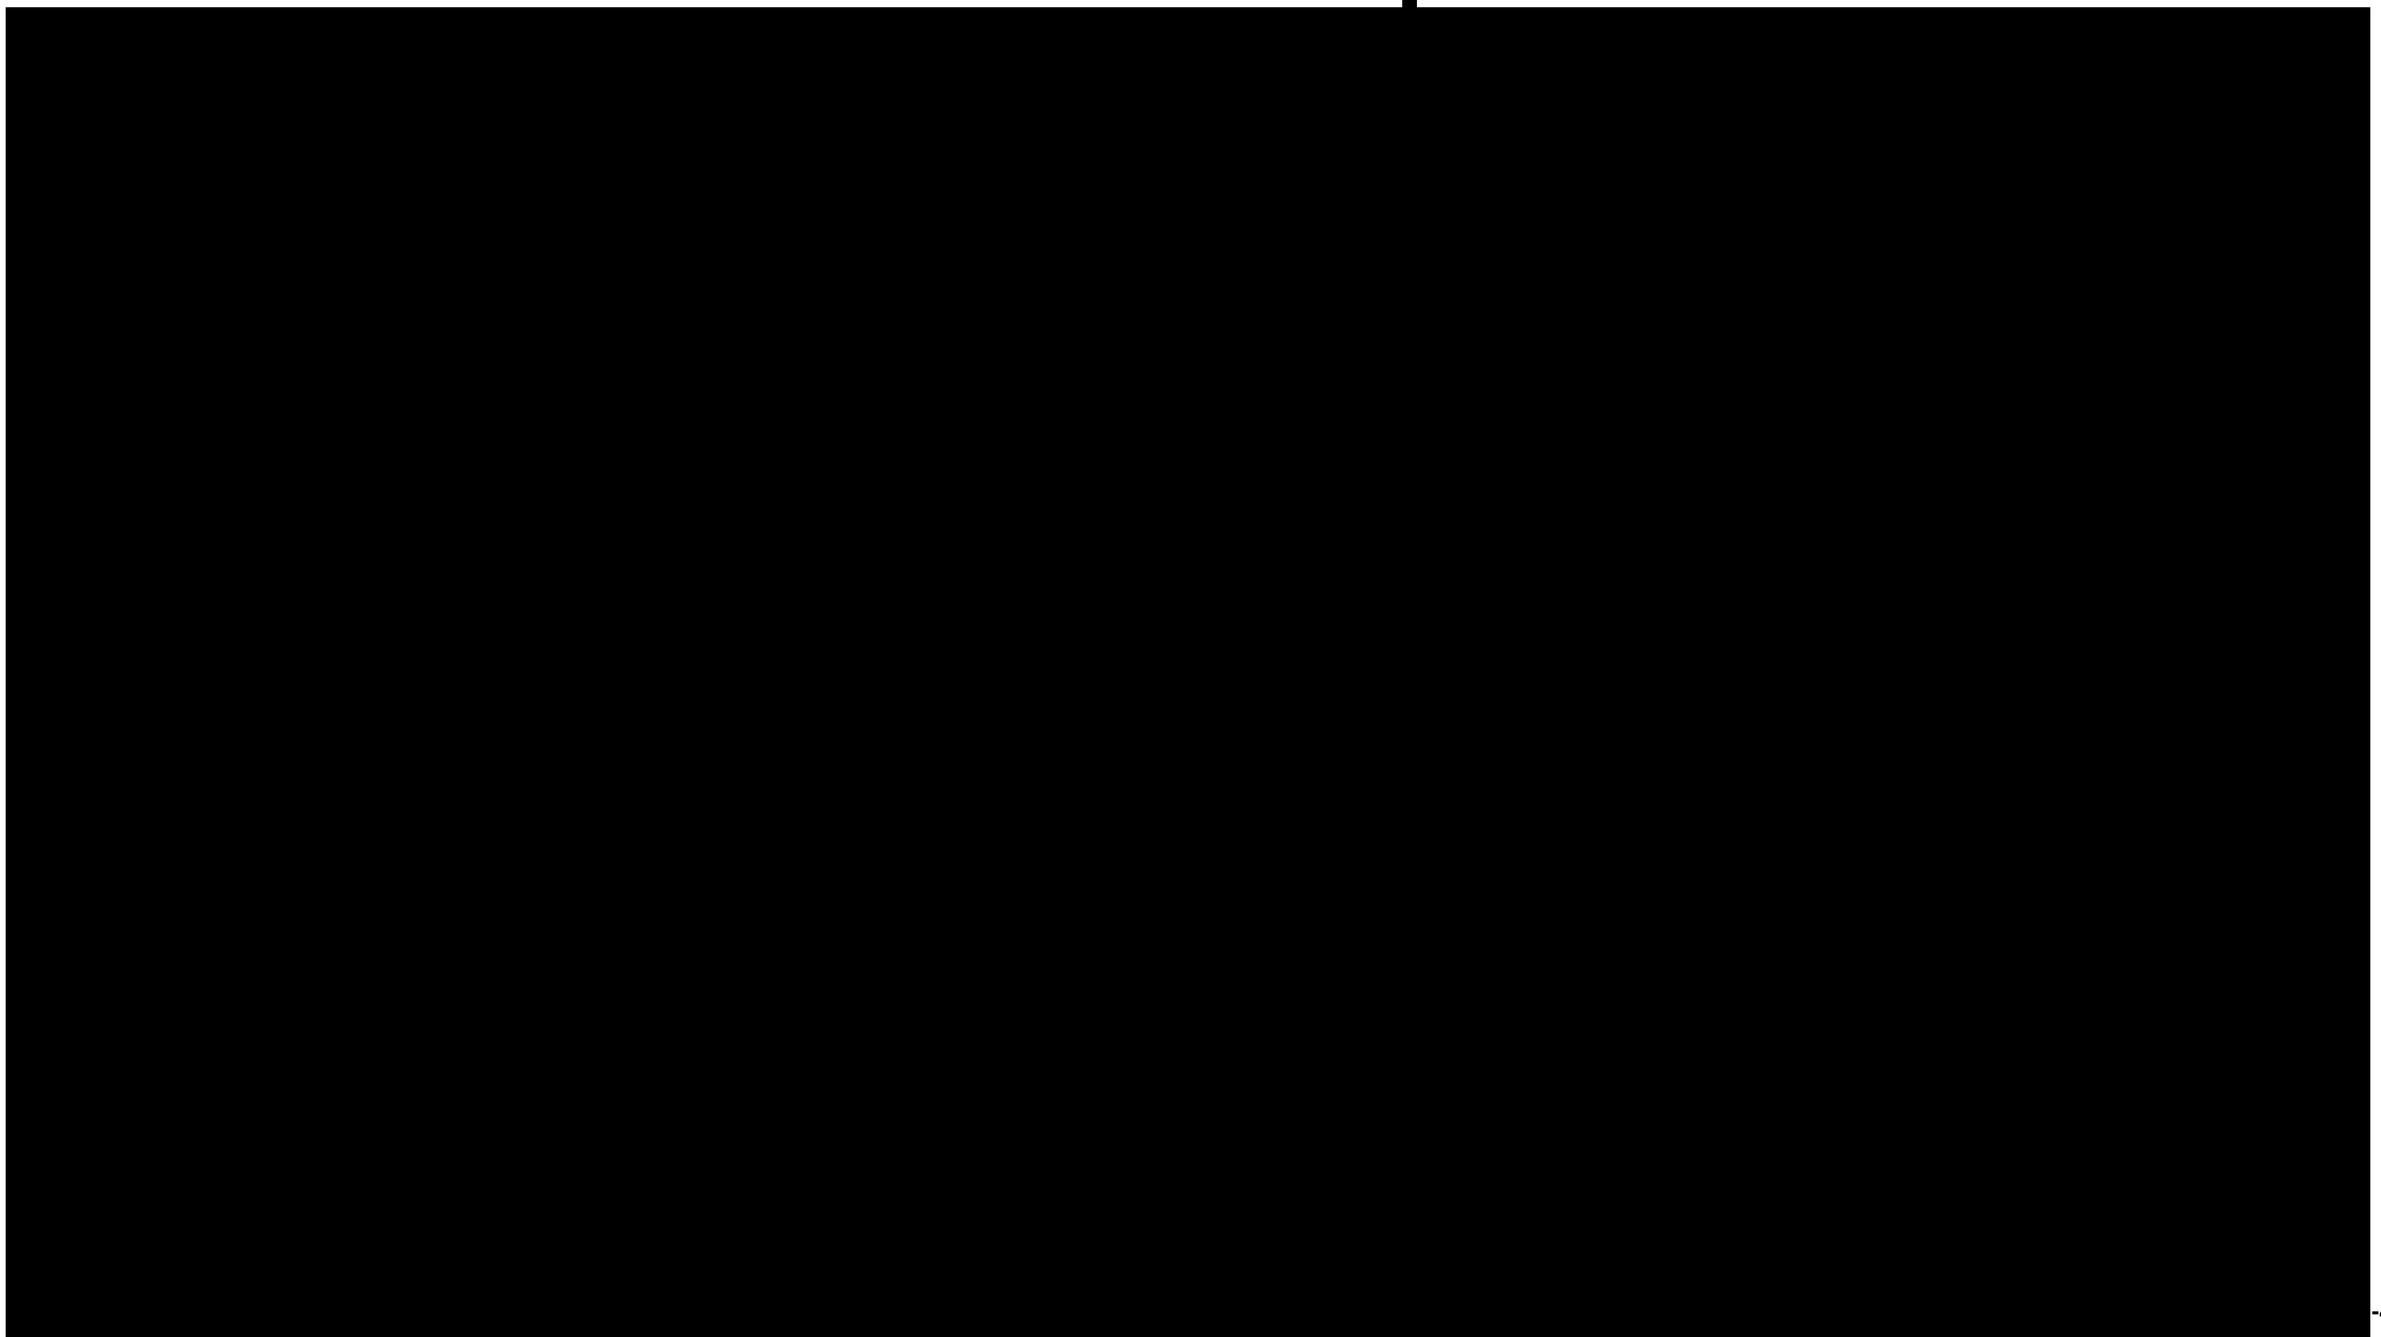

Computer graphics is everywhere!

Entertainment (movies, games)

Entertainment

- Not just cartoons!

Art and design

Industrial design

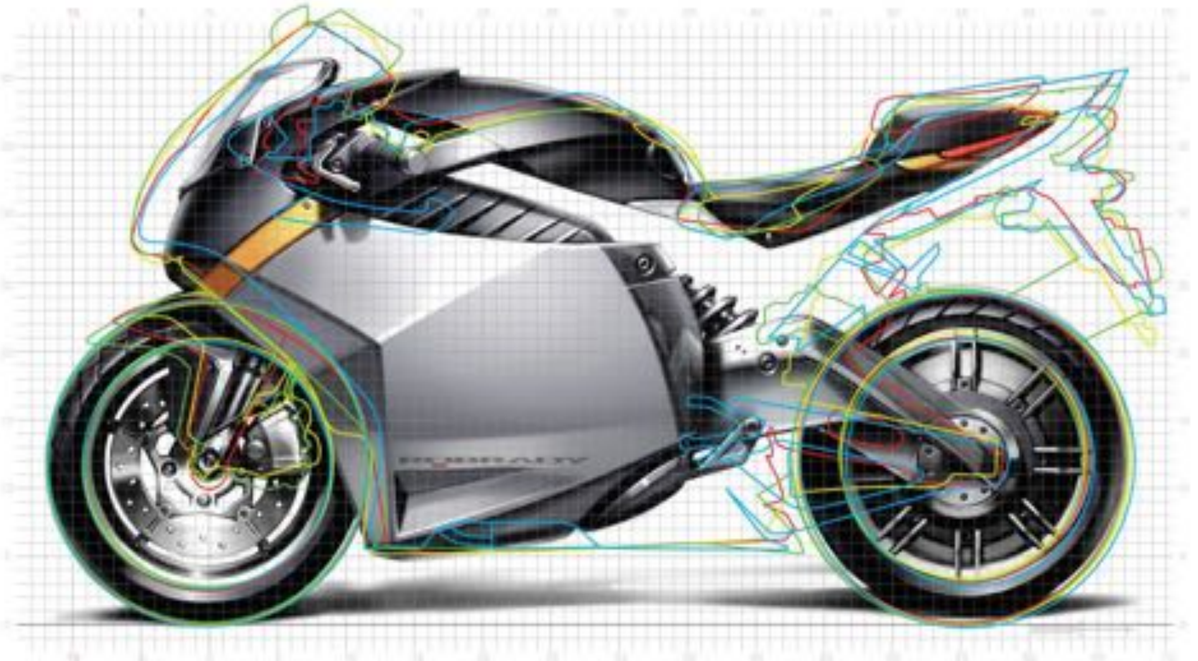

Computer aided engineering (CAE)

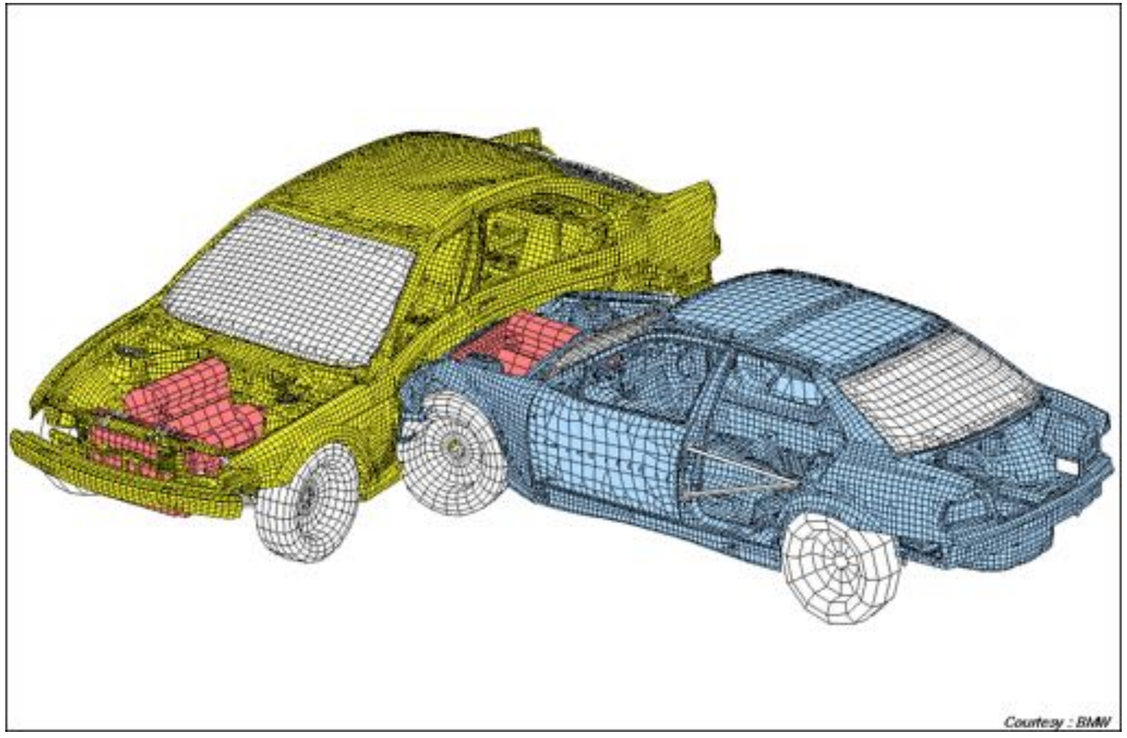

Architecture

Scientific/mathematical visualization

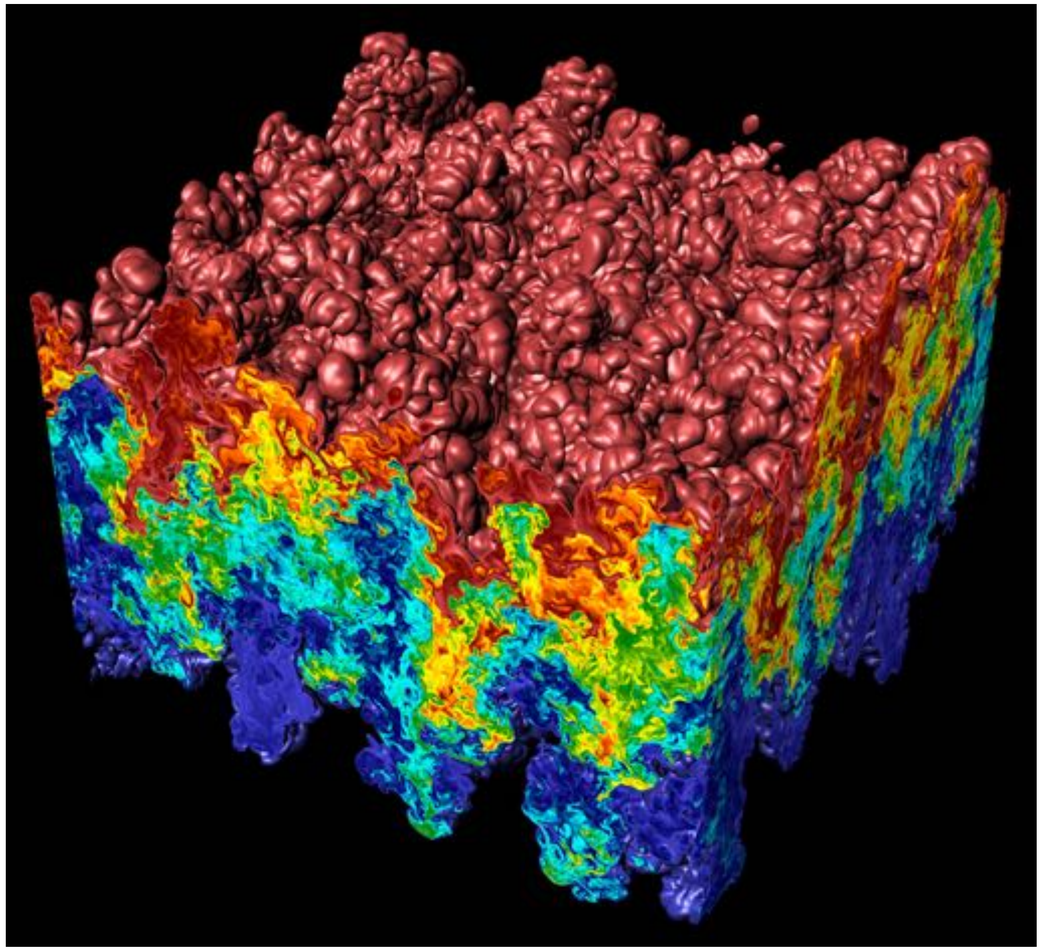

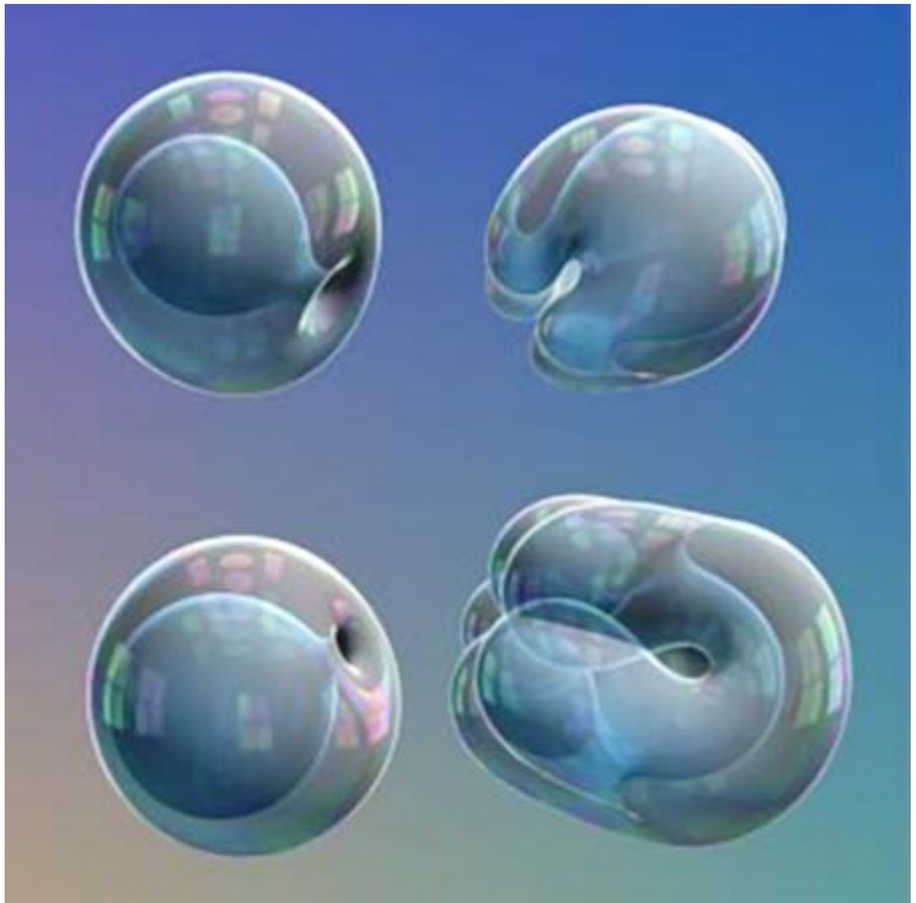

Medical/anatomical visualization

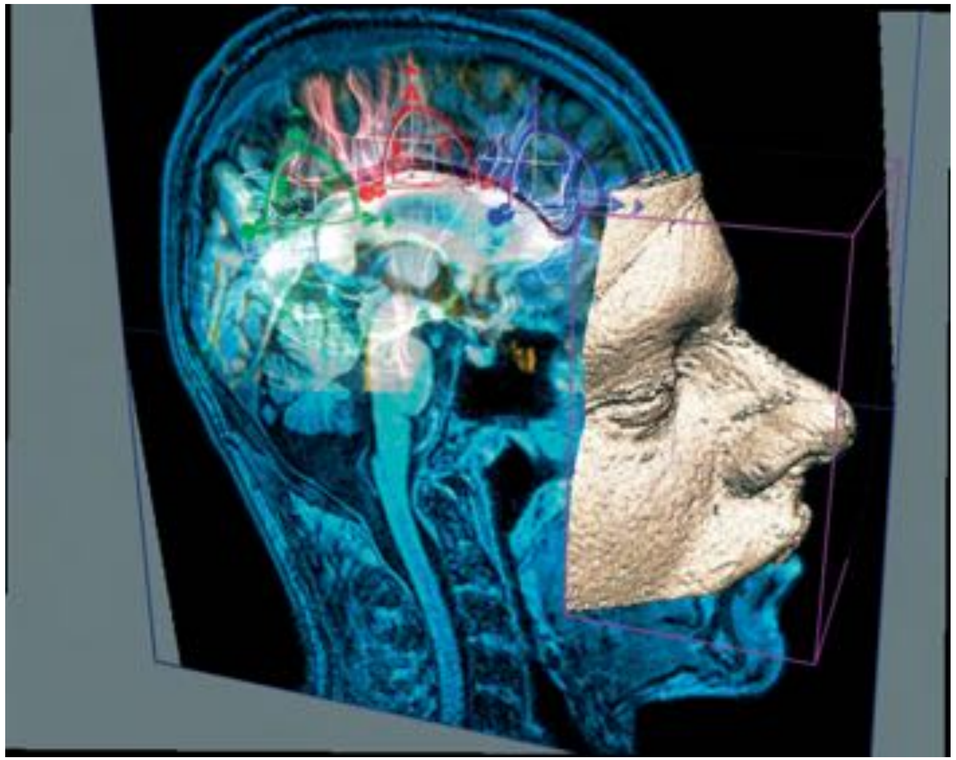

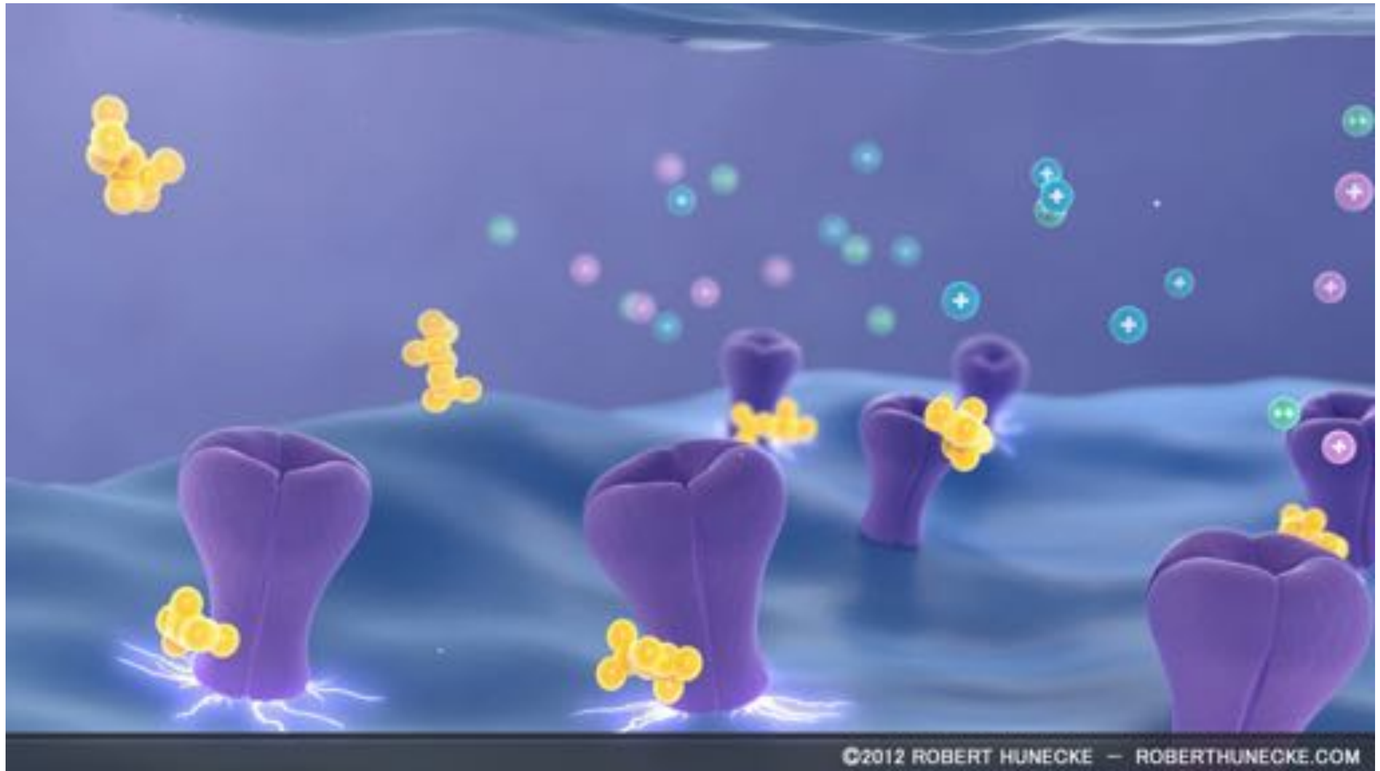

Navigation

Communication

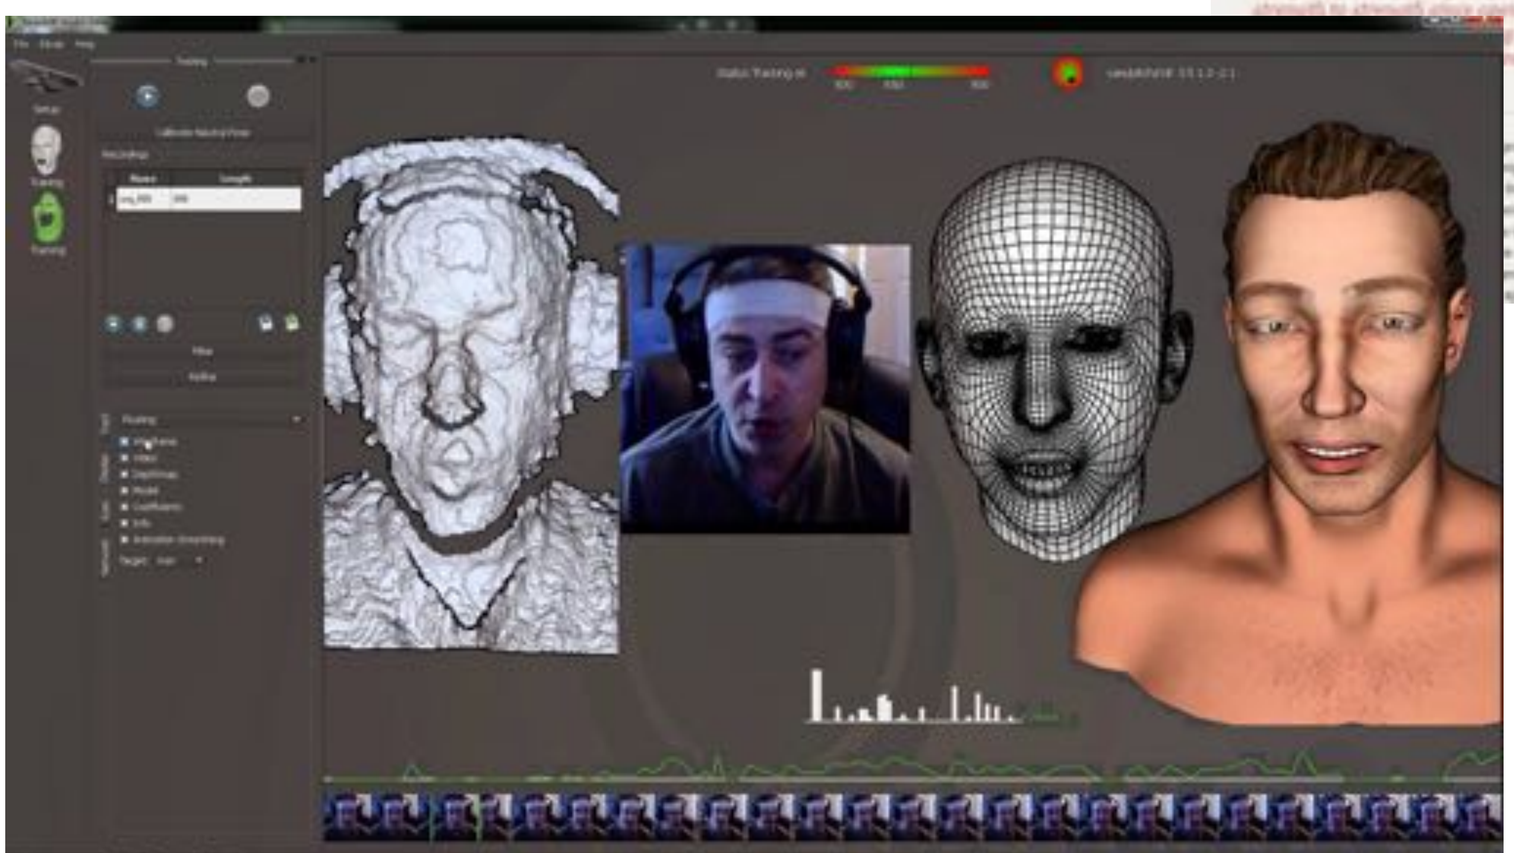

Foundations of computer graphics

- **All these applications demand sophisticated theory & systems**
- **Theory**
  - **basic representations** (how do you digitally encode shape, motion?)
  - **sampling & aliasing** (how do you acquire & reproduce a signal?)
  - **numerical methods** (how do you manipulate signals numerically?)
  - **radiometry & light transport** (how does light behave?)
  - **perception** (how does this all relate to humans?)
  - **…**
- **Systems**
  - **parallel, heterogeneous processing**
  - **graphics-specific programming languages**
  - **…**

Learn by making/doing!

- Build up “Scotty3D” package for modeling/rendering/animation

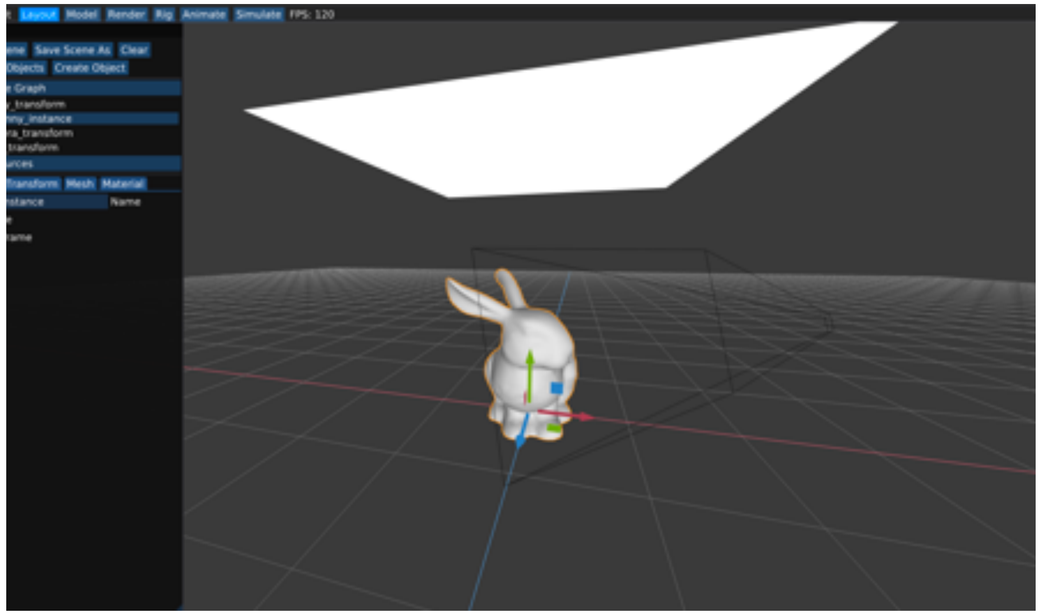

Broken up into four major assignments…

Assignment 1: Rasterization

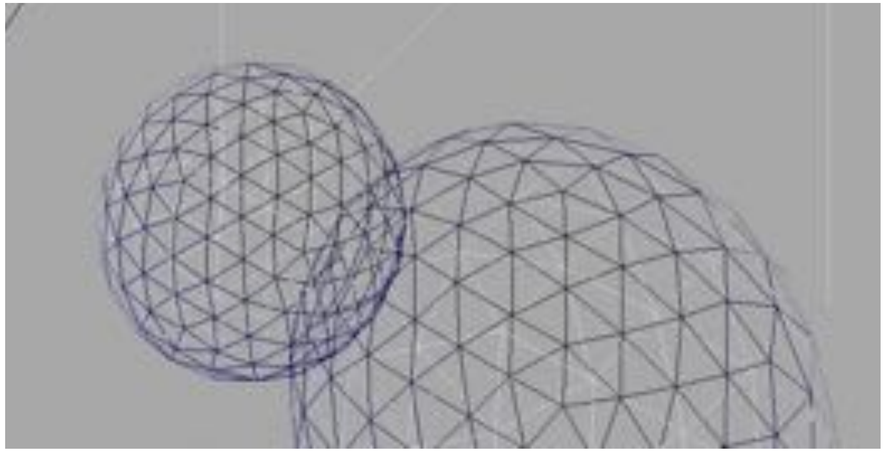

Motivation: 3D without a GPU!

Assignment 2: Geometric Modeling

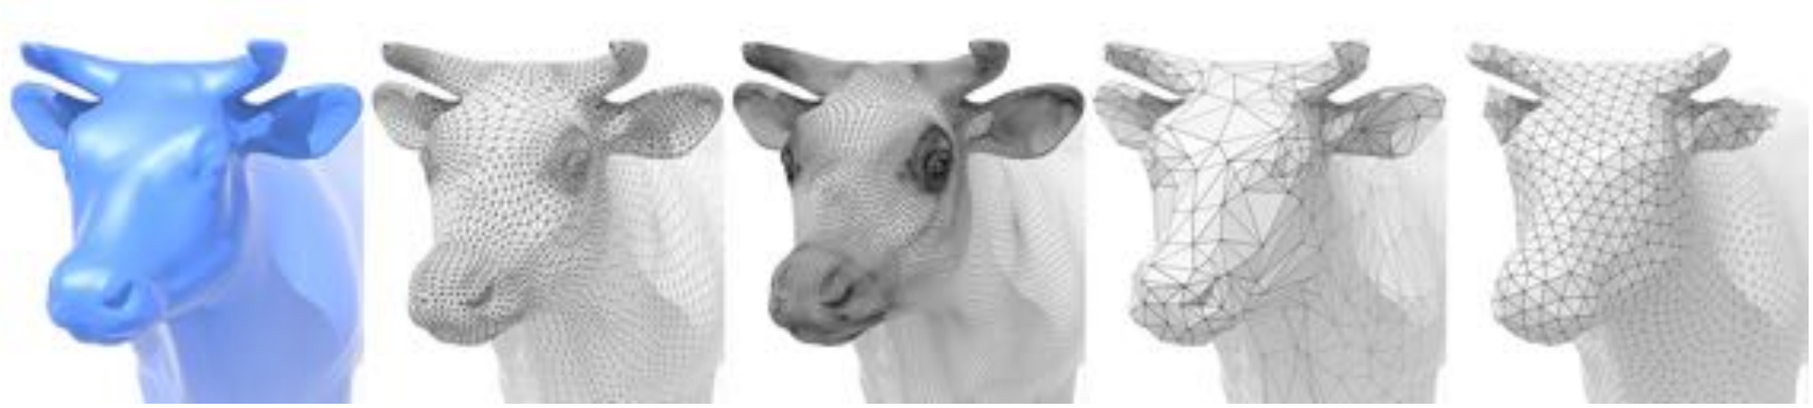

Motivation: create models like these!

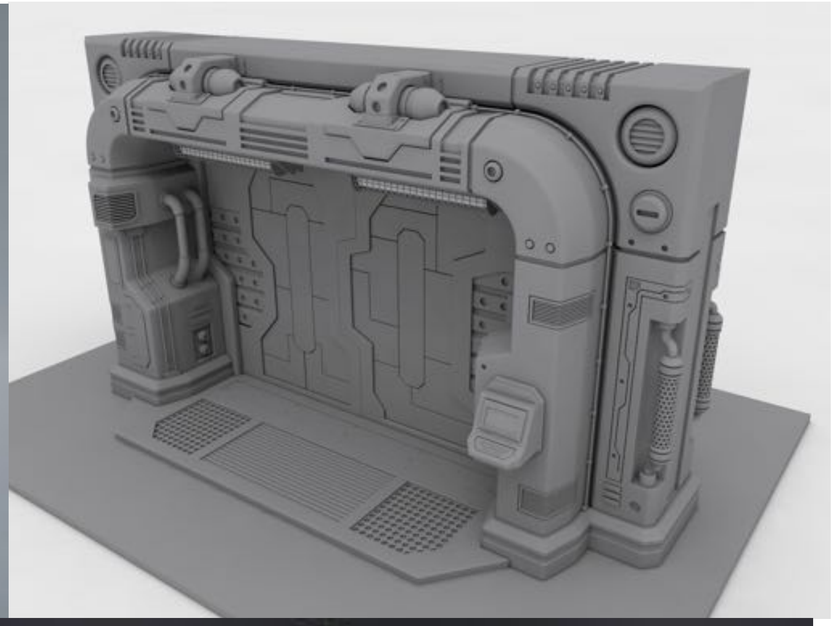

[sources: Richard Yot, 3D-Ace, contrafibbularities,3ddd.ru]

Assignment 3: Photorealistic Rendering

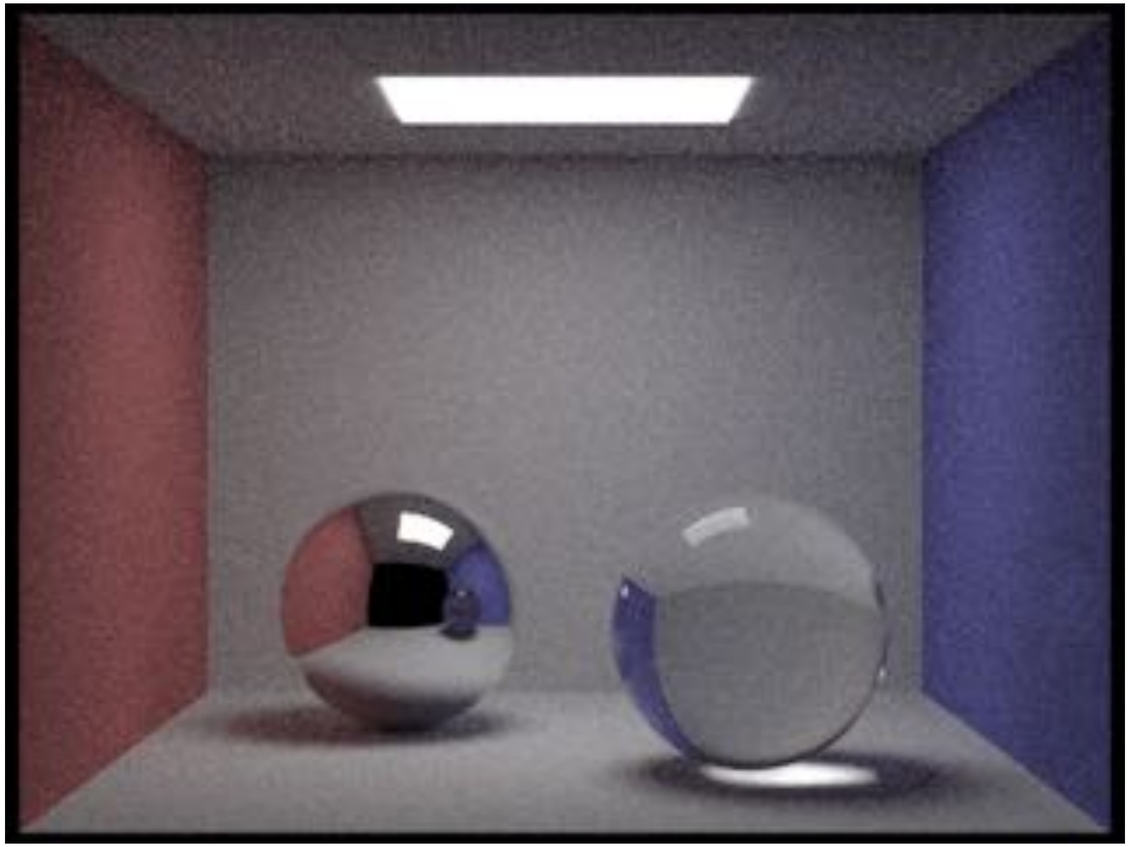

Motivation: render images like these!

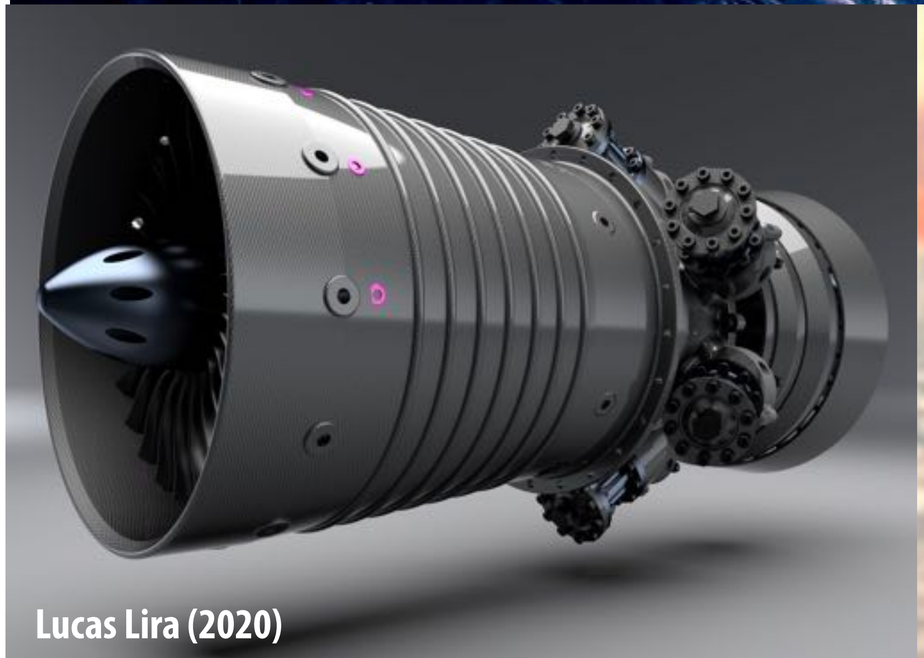

Assignment 4: Animation

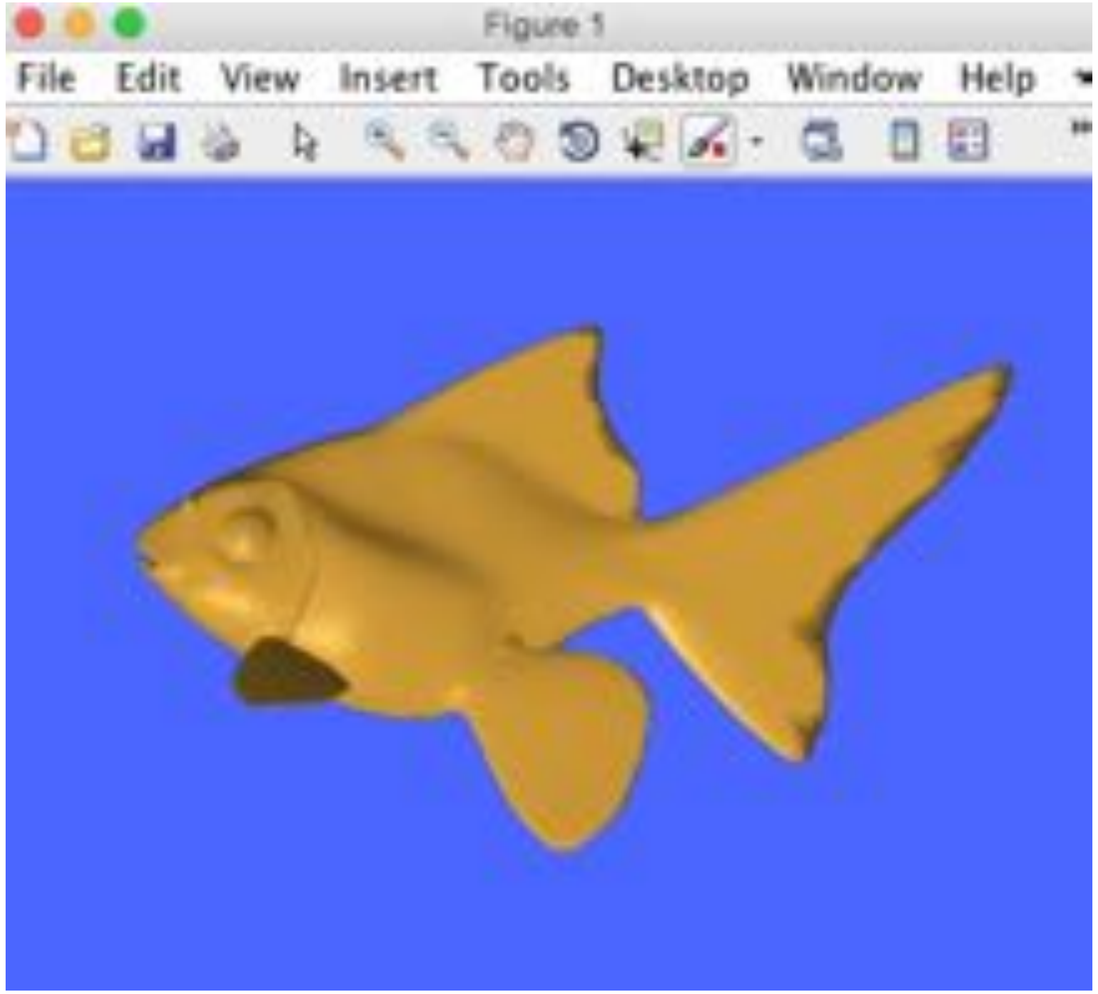

(cribbed from Alec Jacobson)

Motivation: make animations like these!

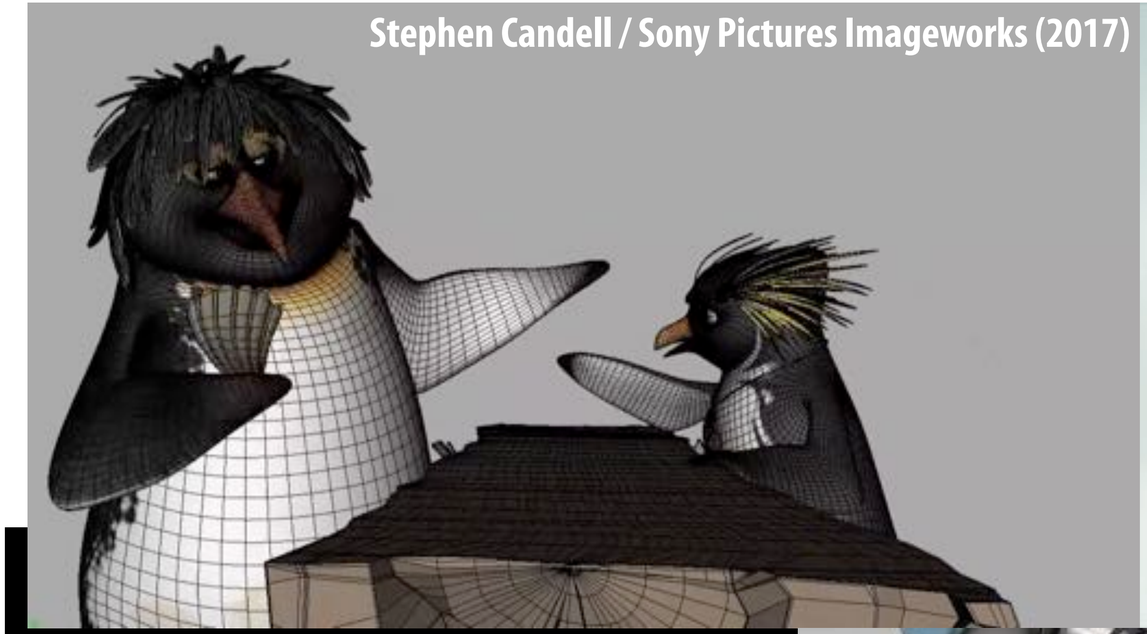

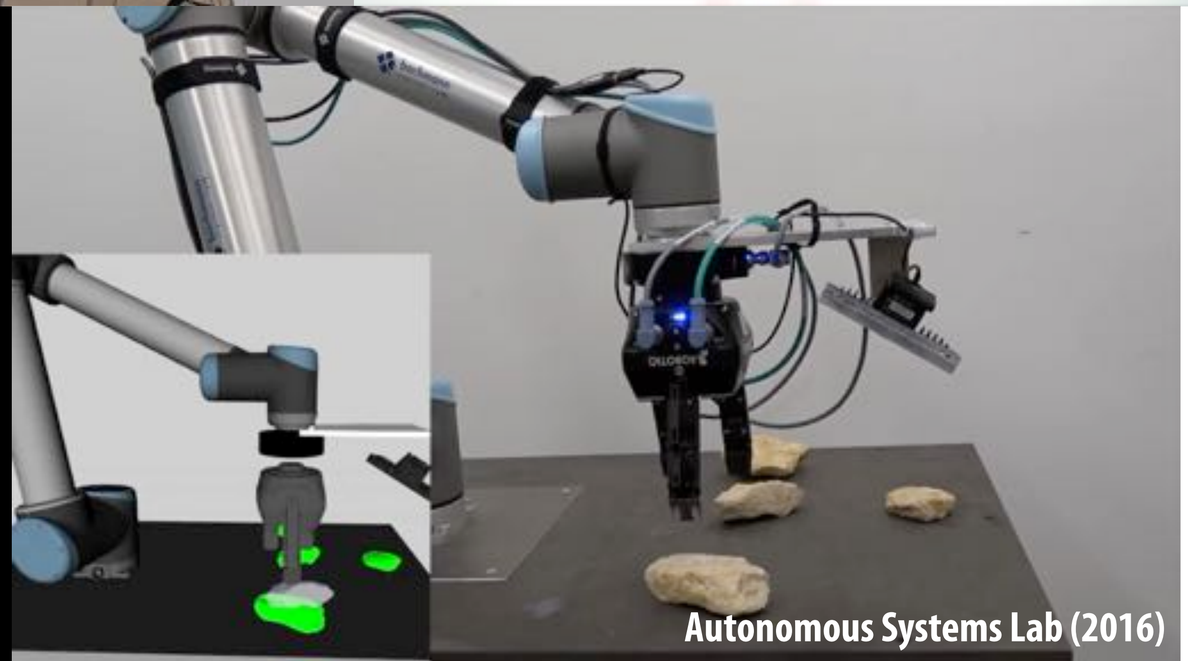

Let’s get started!

Math (P)Review Part I: Linear Algebra

Computer Graphics
CMU 15-462/15-662

Homework 0.0 (Due Monday!)

- Exercises will be a bit harder / more detailed than what you will do for the rest of the class.
- Goal is to help you build strength for the upcoming journey.
- We are here to help!

1 Linear Algebra

1.1 Basic Vector Operations

Exercise 1. *Letting* $\mathbf{u} := (4,3)$, $\mathbf{v} := (4,3)$, $a := 7$ *and* $b := 7$, *calculate the following quantities:*

(a) $\mathbf{u} + \mathbf{v}$

(b) $b\mathbf{u}$

(c) $a\mathbf{u} - b\mathbf{v}$

Exercise 2. *Letting* $\mathbf{u} := (8,2,7)$ *and* $\mathbf{v} := (8,7,3)$, *calculate the following quantities:*

1. $\mathbf{u} - \mathbf{v}$
2. $\mathbf{u} + 6\mathbf{v}$

Exercise 3. *So far we have been working with vectors in* $\mathbb{R}^2$ *and* $\mathbb{R}^3$, *but it is important to remember that other objects, like functions, also behave like vectors in the sense that we can add them, subtract them, multiply them by scalars,* etc. *Calculate the following quantities for the two polynomials* $p(x) := 8x^2 + 2x + 7$ *and* $q(x) := 8x^2 + 7x + 3$, *and evaluate the result at the point* $x = 7$*:*

1. $p(x) - q(x)$
2. $p(x) + 6q(x)$

p.s.
Also see the introductory Mini Homework designed to get you oriented in the class

Private posts on Piazza

- Don't make homework questions private on Piazza
  - other students can't benefit from your question
  - we have to answer the same question repeatedly
  - post as anonymous instead
  - official policy: we won't answer until you make it public :-)
- Private questions should be for "personal stuff" only
  - your grade in the class
  - making up an exam
  - etc.

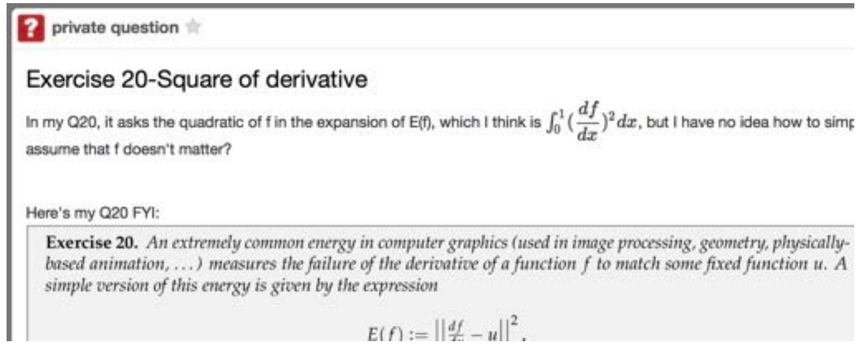

Linear Algebra in Computer Graphics

- **Why is linear algebra important for computer graphics?**
  - **Effective bridge between geometry, physics, etc., and computation.**
  - **In many areas of graphics, once you can express the solution to a problem in terms of linear algebra, you're essentially done: now ask the computer to solve Ax=b.**
  - **Fast numerical linear algebra has really made modern computer graphics possible (image processing, physically-based animation, geometry processing…)**

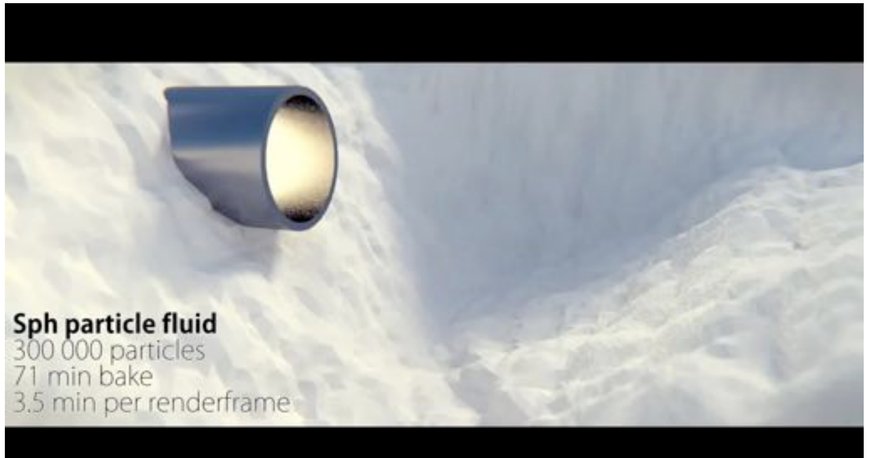

Vectors - Intuition

- First things first: what is a vector?
- Intuitively, a vector is a little arrow:

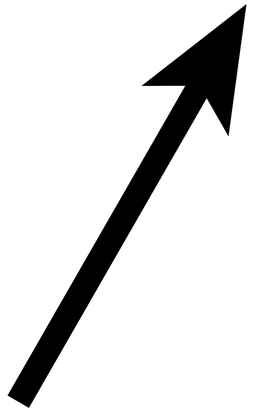

A vector.

- In computer graphics, we work with many types of data that may not look like little arrows (polynomials, images, radiance...). But they still behave like vectors. So, this little arrow is still often a useful mental model.

Vectors - What Can We Measure?

- What information does a vector encode?
- Fundamentally, just direction and magnitude*:

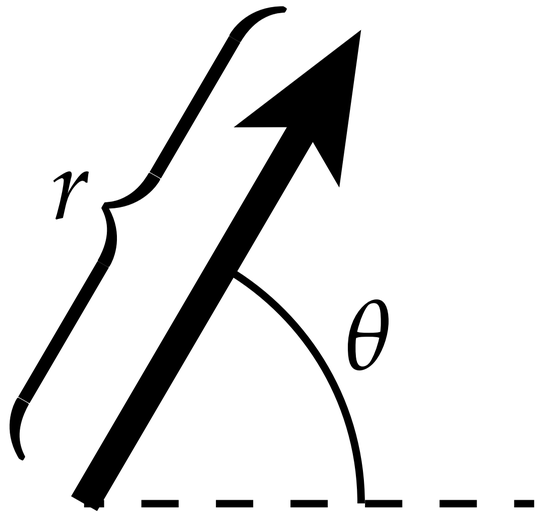

- For instance, a vector in 2D can be encoded by a length and an angle relative to some fixed direction (“polar coordinates”).
- (Side note: are these values the same in any coordinate system?)
- How else might we encode a vector?

*Traditionally, a vector does not include a “basepoint”; a vector with a basepoint is sometimes called a tangent vector.

Vector in Cartesian Coordinates

- Can also measure components of a vector with respect to some chosen coordinate system:

René Descartes, Est. 1596

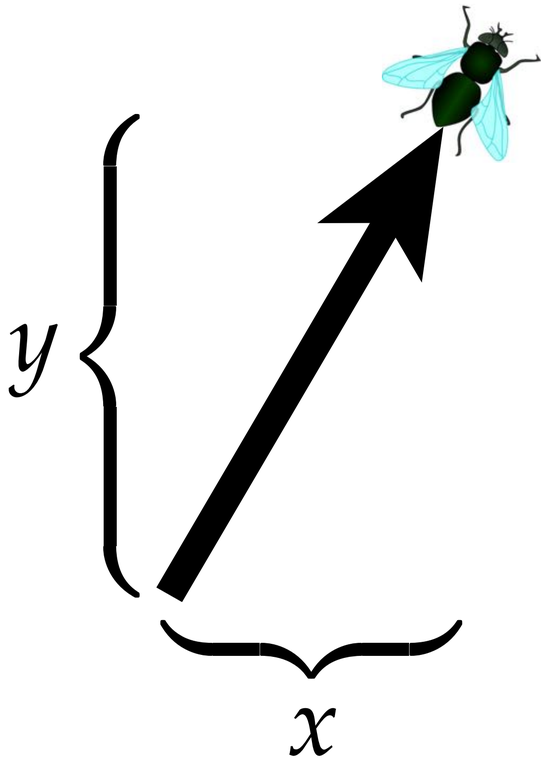

- **WARNING**: Can't directly compare coordinates in different systems! (Also shouldn't compare (r,θ) to (x,y).)

What Can We Do with a Vector?

- Two basic operations. First, we can add them "end to end":

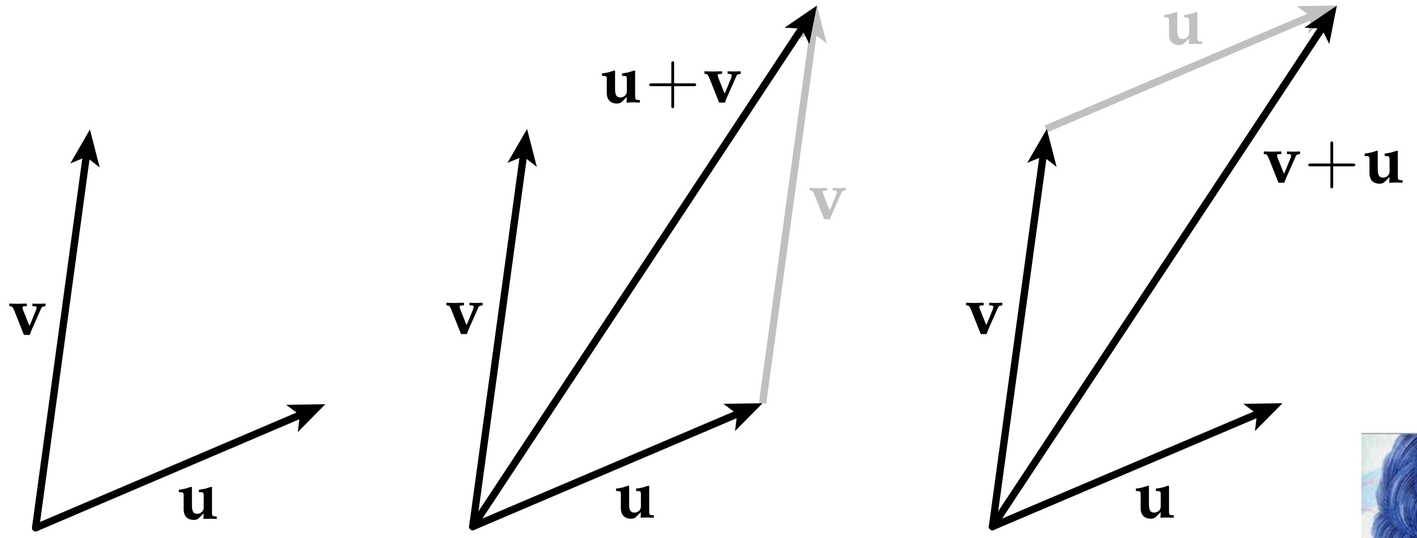

Niels Henrik Abel

- What if we walk along v first, then u?
- Actually, it doesn't seem to matter: u+v = v+u
- Language: vector addition is "commutative" or "abelian"

What Else Can We Do with a Vector?

- Other basic operation? Scaling:

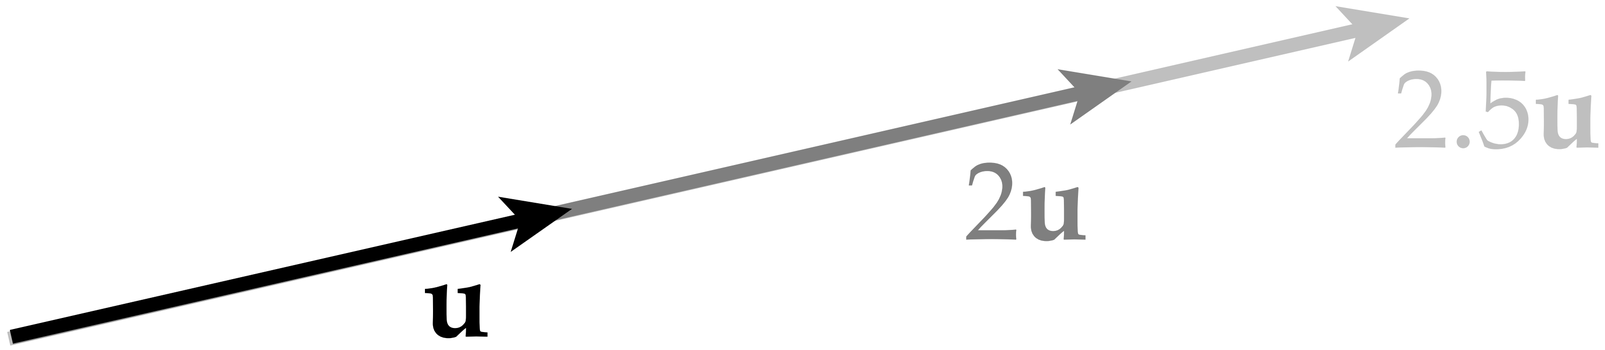

- In general, can multiply any vector u by a number or "scalar" a to get a new vector a**u**.
- Multiplication behaves the way we would expect, based on the geometric behavior of scaling "little arrows." E.g.,

$$a(b\mathbf{u}) = (ab)\mathbf{u}$$

Interaction of Addition & Scaling

- What if we try to add two scaled vectors? Or scale two vectors that have been added together?

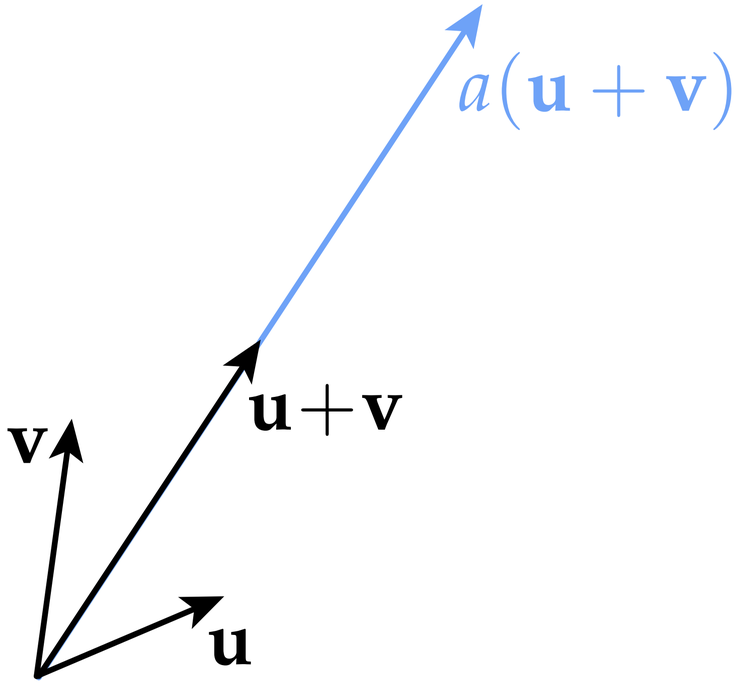

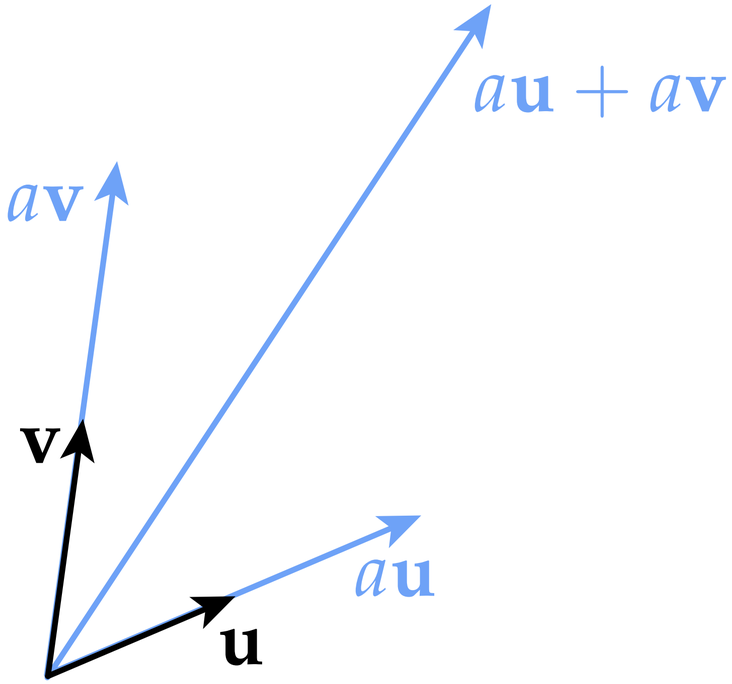

- Interesting—seems we get the same result either way:

$$a(\mathbf{u}+\mathbf{v}) = a\mathbf{u}+a\mathbf{v}$$

Vector Space—Formal Definition

- **If we keep playing around vectors, eventually we come up with a complete set of "rules" that vectors seem to obey:**

  For all vectors $\mathbf{u}, \mathbf{v}, \mathbf{w}$ and scalars a, b:

  - $\mathbf{u} + \mathbf{v} = \mathbf{v} + \mathbf{u}$
  - $\mathbf{u} + (\mathbf{v} + \mathbf{w}) = (\mathbf{u} + \mathbf{v}) + \mathbf{w}$
  - There exists a *zero vector* "$\mathbf{0}$" such that $\mathbf{v} + \mathbf{0} = \mathbf{0} + \mathbf{v} = \mathbf{v}$
  - For every $\mathbf{v}$ there is a vector "$-\mathbf{v}$" such that $\mathbf{v} + (-\mathbf{v}) = \mathbf{0}$
  - $1\mathbf{v} = \mathbf{v}$
  - $a(b\mathbf{v}) = (ab)\mathbf{v}$
  - $a(\mathbf{u} + \mathbf{v}) = a\mathbf{u} + a\mathbf{v}$
  - $(a + b)\mathbf{v} = a\mathbf{v} + b\mathbf{v}$

- **These rules did not "fall out of the sky!" Each one comes from the geometric behavior of "little arrows." (Can you draw a picture for each one?)**

- **Any collection of objects satisfying all of this properties is a vector space (even if they don't look like little arrows!)**

Euclidean Vector Space

- Most common example: Euclidean n-dimensional space
- Typically denoted by $\mathbb{R}^n$, meaning "n real numbers"
- E.g., (1.23, 4.56, π/2) is a point in $\mathbb{R}^3$
- Why such a common example?
  - Looks a lot like the space we live in!
  - That's what we can easily encode on a computer (a list of floating-point numbers).

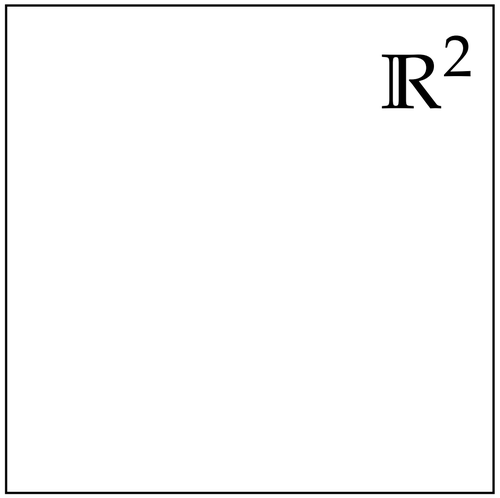

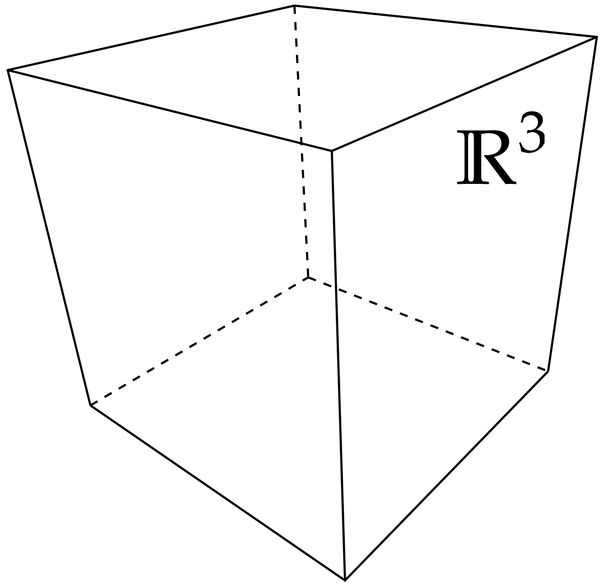

Functions as Vectors

- Another very important example of vector spaces in computer graphics are spaces of functions.
- Why? Because many of the objects we want to work with in graphics are functions! (Images, radiance from a light source, surfaces, modal vibrations, ...)

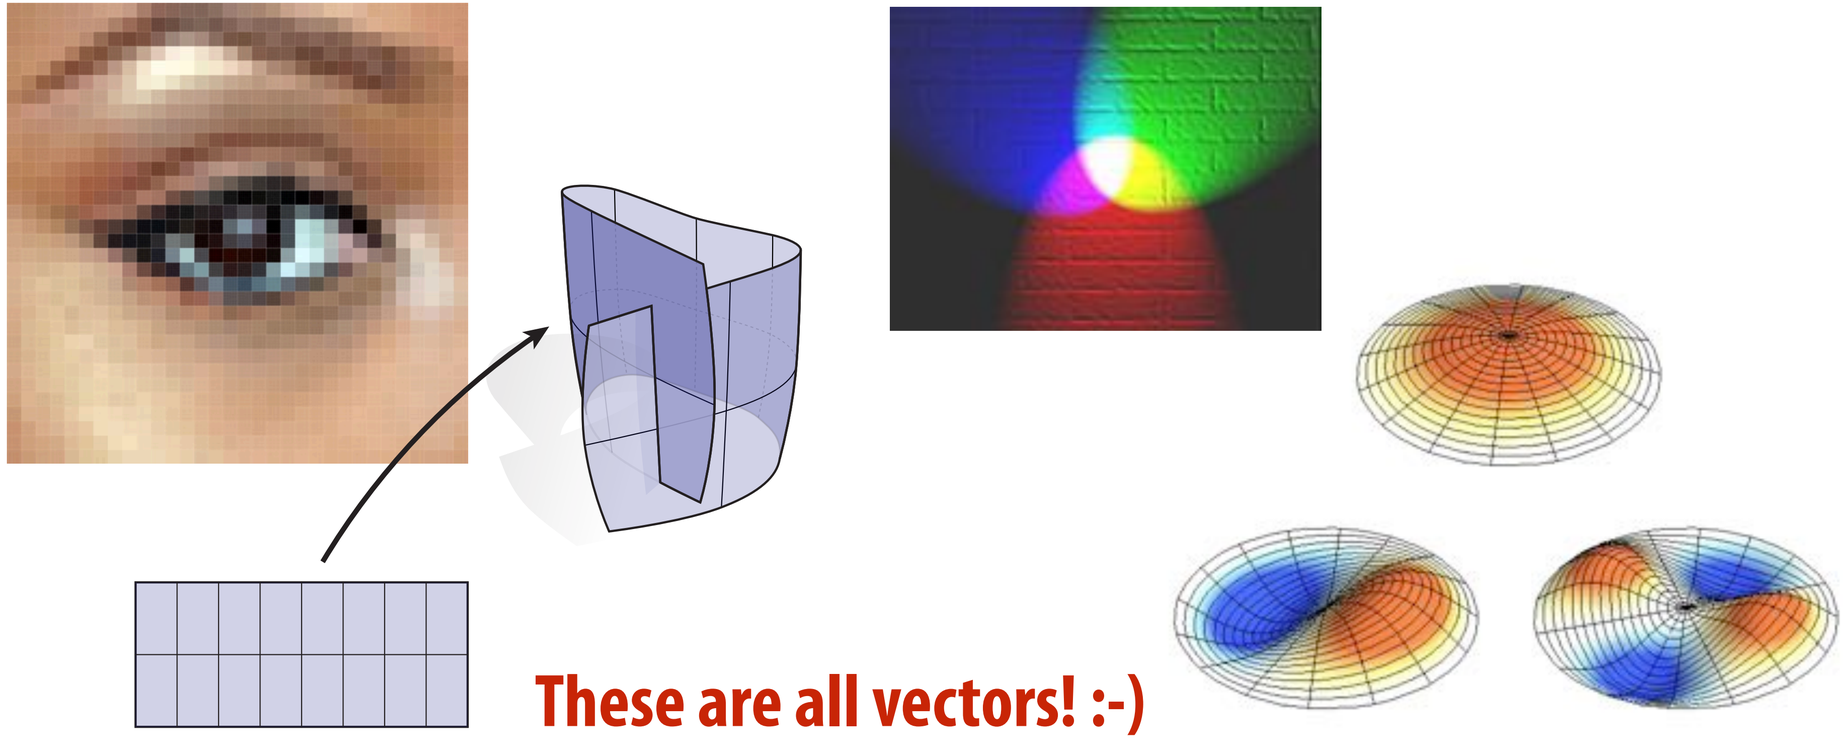

These are all vectors! :-)

Functions as Vectors

- Do functions exhibit the same behavior as "little arrows?"
- Well, we can certainly add two functions:

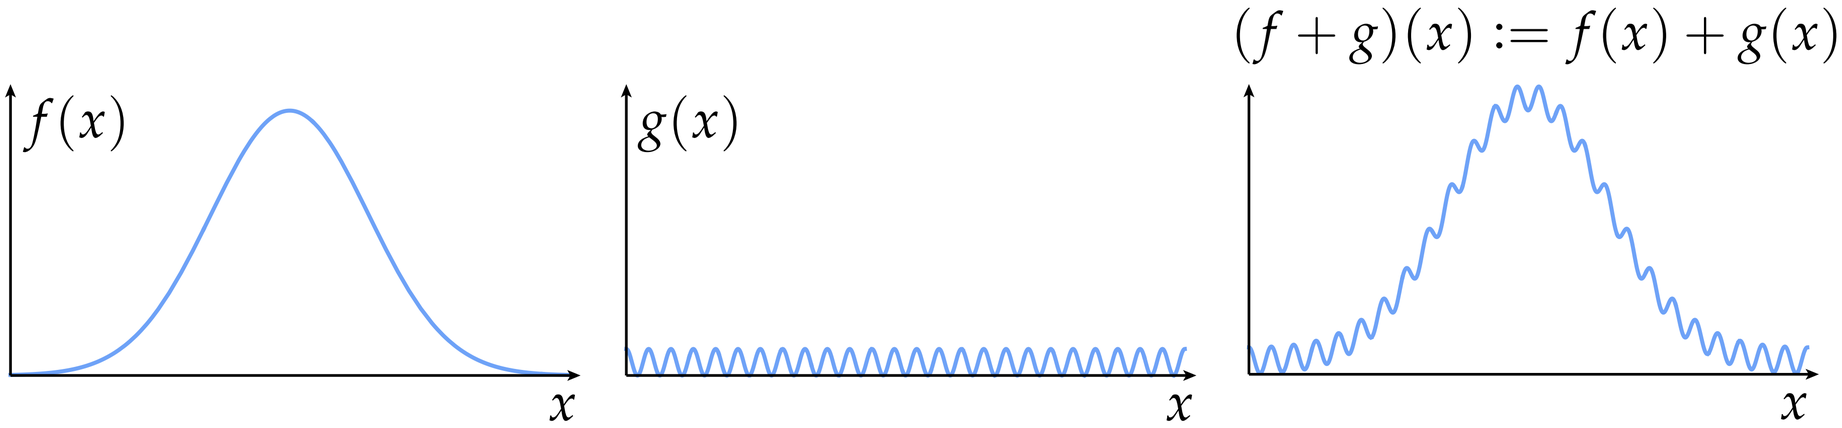

- We can also scale a function:

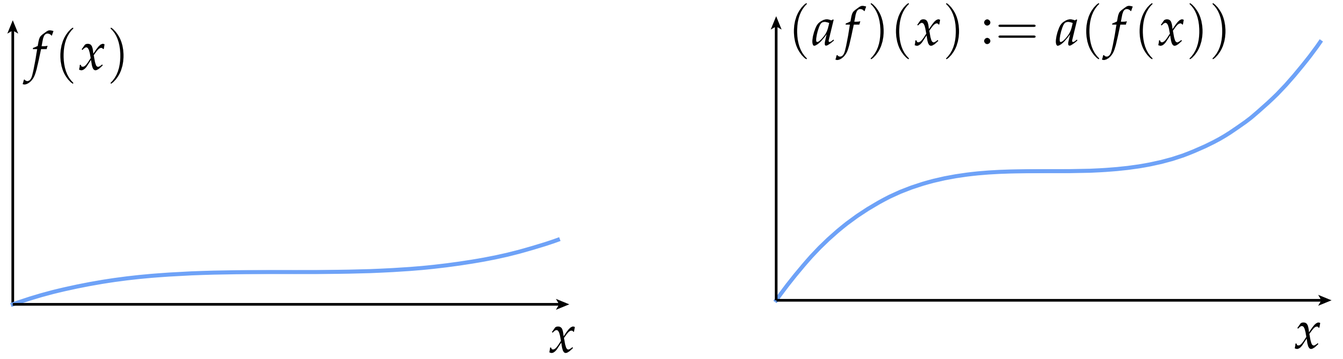

Functions as Vectors

- **What about the rest of these properties?**

  For all vectors $\mathbf{u}, \mathbf{v}, \mathbf{w}$ and scalars a, b:

  - $\mathbf{u} + \mathbf{v} = \mathbf{v} + \mathbf{u}$
  - $\mathbf{u} + (\mathbf{v} + \mathbf{w}) = (\mathbf{u} + \mathbf{v}) + \mathbf{w}$
  - There exists a *zero vector* "$\mathbf{0}$" such that $\mathbf{v} + \mathbf{0} = \mathbf{0} + \mathbf{v} = \mathbf{v}$
  - For every $\mathbf{v}$ there is a vector "$-\mathbf{v}$" such that $\mathbf{v} + (-\mathbf{v}) = \mathbf{0}$
  - $1\mathbf{v} = \mathbf{v}$
  - $a(b\mathbf{v}) = (ab)\mathbf{v}$
  - $a(\mathbf{u} + \mathbf{v}) = a\mathbf{u} + b\mathbf{v}$
  - $(a + b)\mathbf{v} = a\mathbf{v} + b\mathbf{v}$

- **Try it out at home!**
- **E.g., the "zero vector" is the function equal to zero for all x.**
- **Short answer: yes, functions are vectors! (Even if they don't look like "little arrows".)**

Vectors in Coordinates

- So far, we've only drawn our vector operations via pictures.
- How do we actually compute with vectors?
- Return to our coordinate representation:

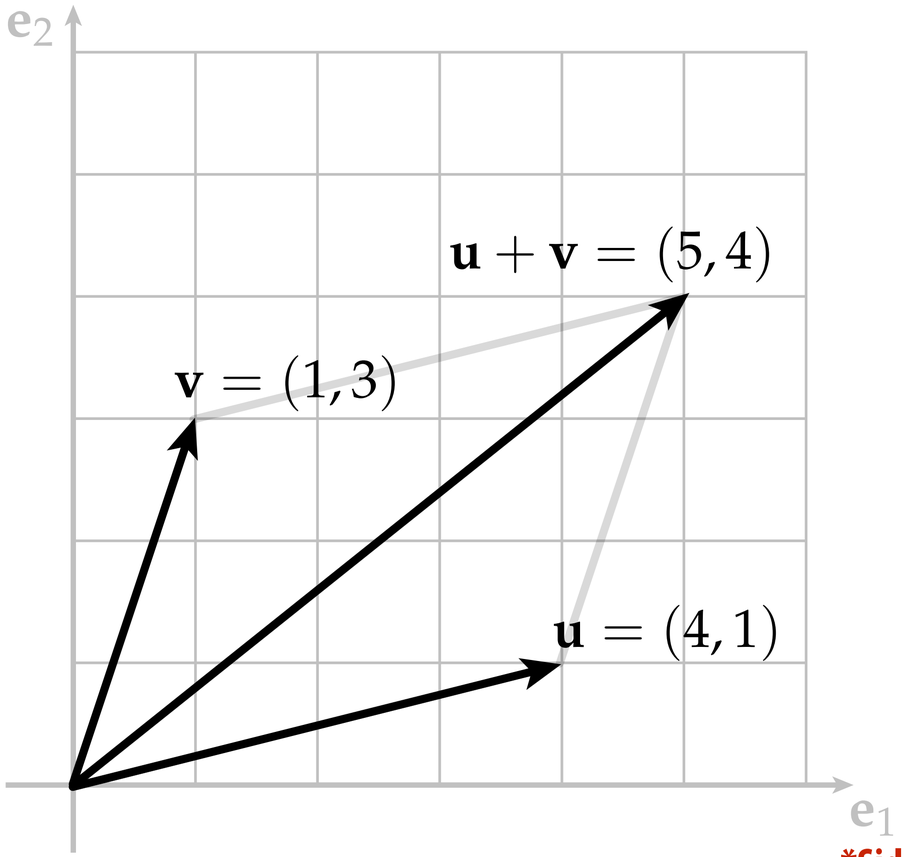

$$\mathbf{u}+\mathbf{v}$$
$$=(1,3)+(4,1)$$
$$=(1+4,3+1)$$
$$=(5,4)$$

*Side note: does it make sense to add vectors encoded as (r,θ)?

Ok, so we came up with some
rule for adding pairs of numbers.

How can we check that it faithfully encodes
geometric behavior of “little arrows?”

From Geometry to Algebra

- **Just check that it agrees with our list of rules that we know (from reasoning geometrically) "little arrows" must obey:**

For all vectors $\mathbf{u}, \mathbf{v}, \mathbf{w}$ and scalars a, b:

- $\mathbf{u} + \mathbf{v} = \mathbf{v} + \mathbf{u}$
- $\mathbf{u} + (\mathbf{v} + \mathbf{w}) = (\mathbf{u} + \mathbf{v}) + \mathbf{w}$
- There exists a *zero vector* "$\mathbf{0}$" such that $\mathbf{v} + \mathbf{0} = \mathbf{0} + \mathbf{v} = \mathbf{v}$
- For every $\mathbf{v}$ there is a vector "$-\mathbf{v}$" such that $\mathbf{v} + (-\mathbf{v}) = \mathbf{0}$
- $1\mathbf{v} = \mathbf{v}$
- $a(b\mathbf{v}) = (ab)\mathbf{v}$
- $a(\mathbf{u} + \mathbf{v}) = a\mathbf{u} + a\mathbf{v}$
- $(a + b)\mathbf{v} = a\mathbf{v} + b\mathbf{v}$

- **For instance, for any two vectors $\mathbf{u} := (u_1, u_2)$ and $\mathbf{v} := (v_1, v_2)$ we have**

$$\mathbf{u} + \mathbf{v} = (u_1, u_2) + (v_1, v_2) = (u_1 + v_1, u_2 + v_2) =$$
$$(v_1 + u_1, v_2 + u_2) = (v_1, v_2) + (u_1, u_2) = \mathbf{v} + \mathbf{u}.$$

Turning geometric observations into algebraic rules is convenient for symbolic manipulation & numerical computation.

But you should never blindly accept a rule given by authority.

Always ask: where does this rule come from?
What does it mean geometrically?
(Can you draw a picture?)

Scaling Vectors in Coordinates

- We'd also like to be able to scale vectors using coordinates.
- Any ideas?

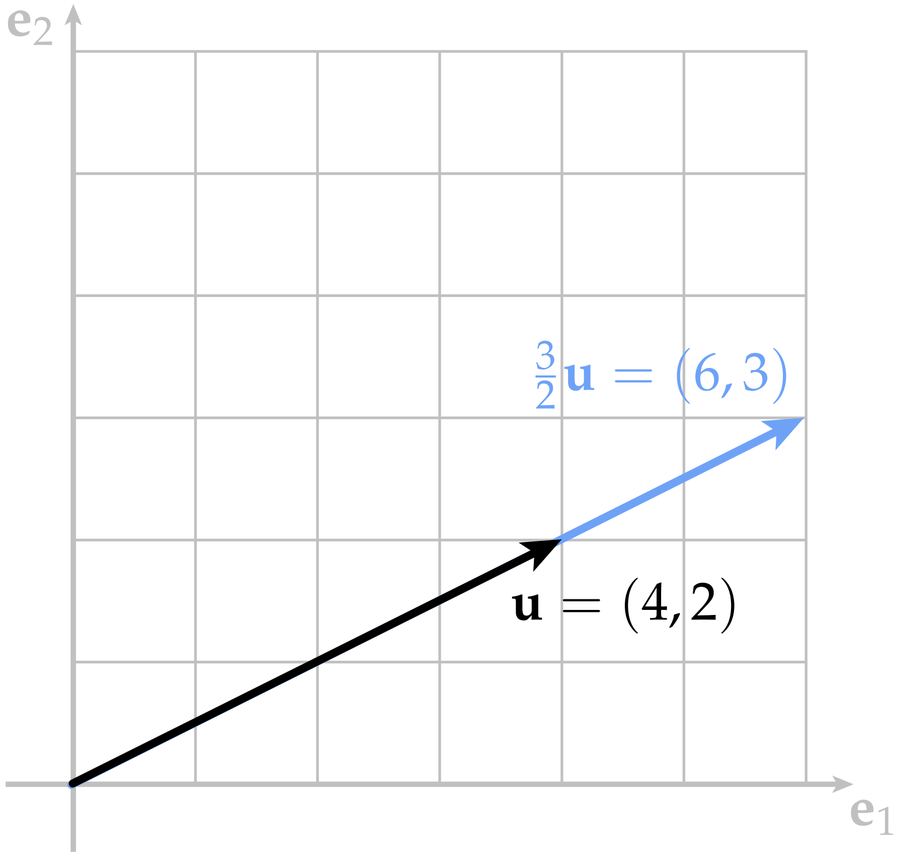

$$\frac{3}{2}\mathbf{u}$$
$$= \frac{3}{2}(4,2)$$
$$= (4 \cdot 3/2, 2 \cdot 3/2)$$
$$= (12/2, 6/2)$$
$$= (6,3)$$

(From here, check the rest of the properties...)

Computing the Midpoint

- As we start to combine vector operations, we build up operations needed for computer graphics.
- E.g., how would I compute the midpoint m of a = (3,4) and b=(7,2)?

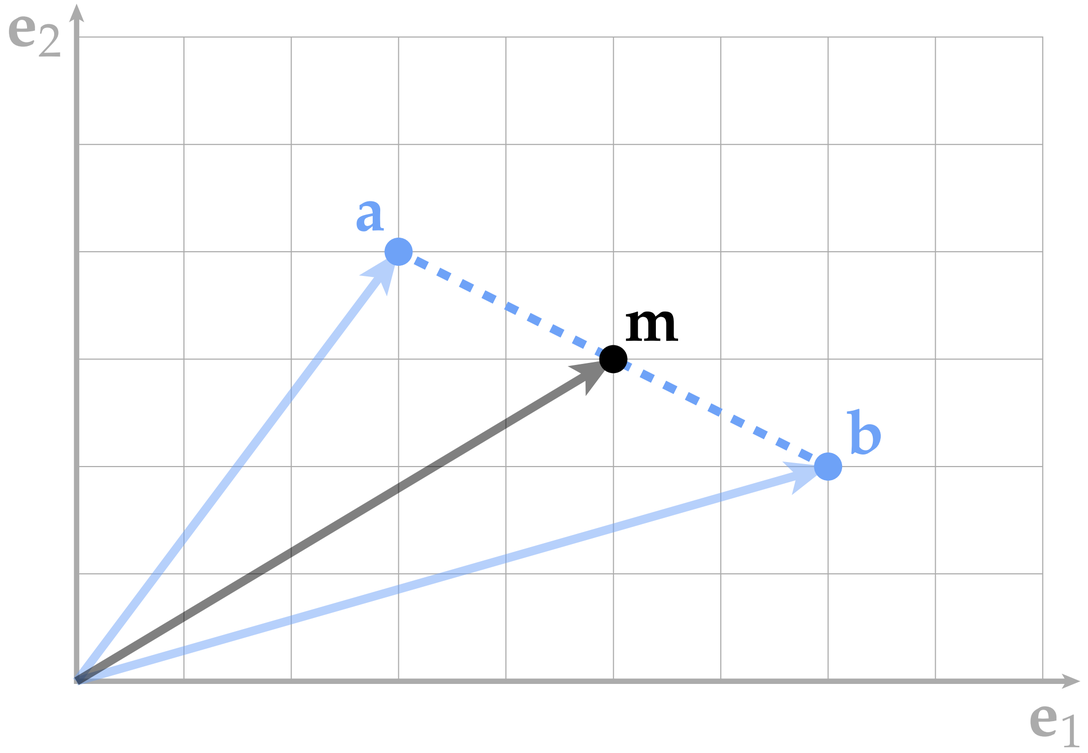

$$\mathbf{m} = \tfrac{1}{2}(\mathbf{a} + \mathbf{b})$$

$$= \tfrac{1}{2}((3,4) + (7,2))$$

$$= \tfrac{1}{2}(10,6)$$

$$= (5,3)$$

Measuring Vectors

- Earlier we asked, “what information does a vector encode?”
- (A: Orientation and magnitude.)
- How do we actually measure these quantities?

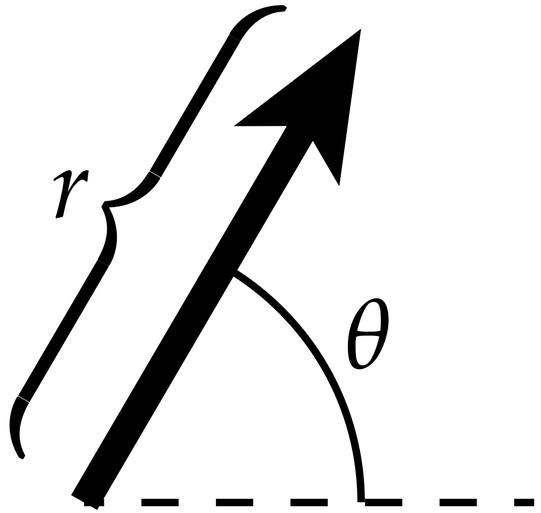

Norm of a Vector

- Let's start with magnitude—for a given vector v, we want to assign it a number |v| called its length or magnitude or norm.
- Intuitively, the norm should capture how "big" the vector is.

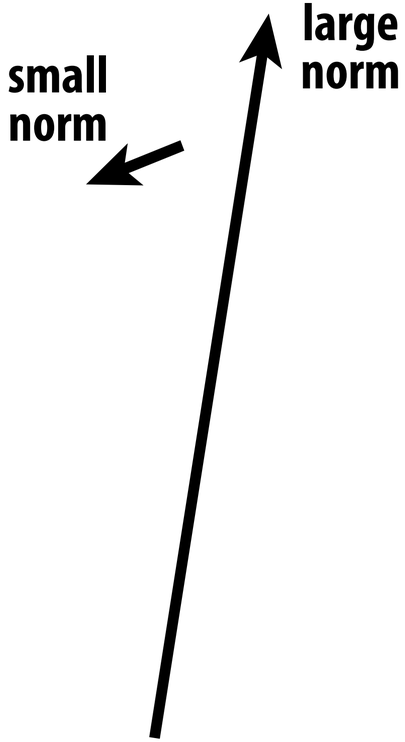

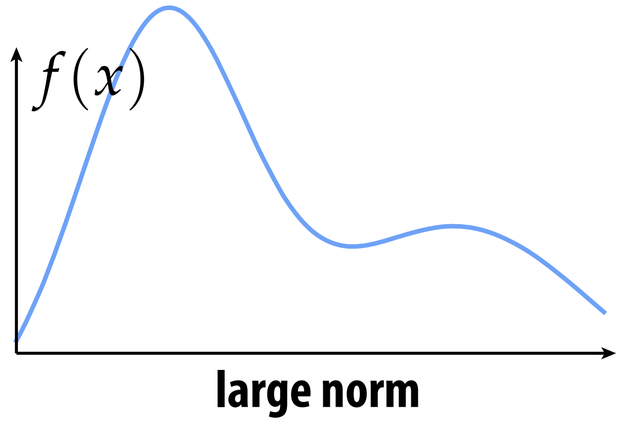

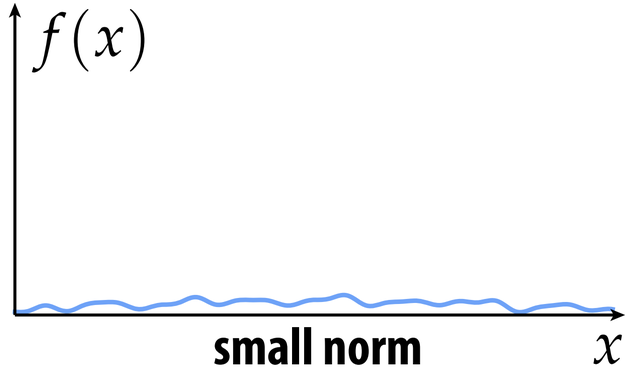

small norm

large norm

Natural Properties of Length—Positivity

- What properties might you expect the norm (or length) of a vector to satisfy?
- For one thing, it probably shouldn't be negative!

$$|\mathbf{u}| \geq 0$$

- And probably it should be zero only for the zero vector:

$$|\mathbf{u}| = 0 \quad \Longleftrightarrow \quad \mathbf{u} = \mathbf{0}$$

Natural Properties of Length, Continued

- Also, if we scale a vector by a factor c, its norm (i.e., length) really should scale by the same amount:

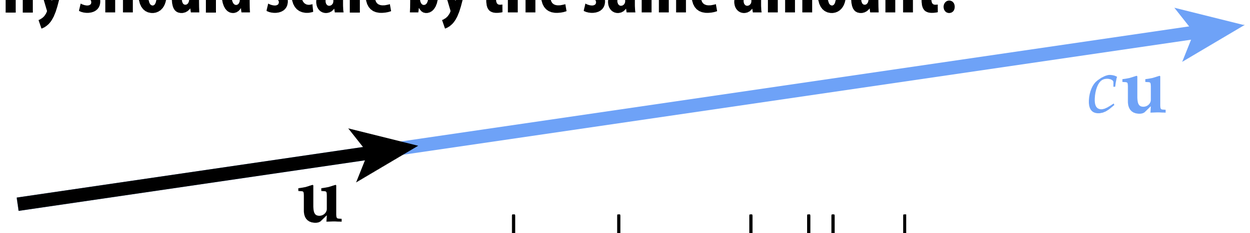

$$|c\mathbf{u}| = |c||\mathbf{u}|$$

- Finally, we know that the shortest path between two points is always along a straight line:

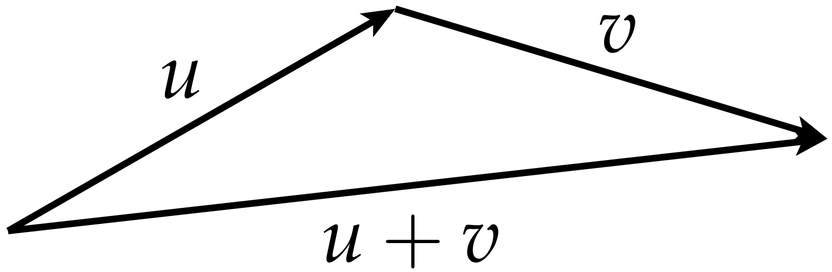

$$|\mathbf{u}| + |\mathbf{v}| \geq |\mathbf{u} + \mathbf{v}|$$

- (This final property is sometimes called the "triangle inequality," since the diagram looks like a triangle.)

Norm—Formal Definition

- A norm is any function that assigns a number to each vector and satisfies the following properties for all vectors u, v, and all scalars a:

  - $|\mathbf{v}| \geq 0$
  - $|\mathbf{v}| = 0 \quad \Longleftrightarrow \quad \mathbf{v} = \mathbf{0}$
  - $|a\mathbf{v}| = |a||\mathbf{v}|$
  - $|\mathbf{u}| + |\mathbf{v}| \geq |\mathbf{u} + \mathbf{v}|$

- But you don't have to take my word for it—for each rule, you now have a concrete geometric picture explaining why this "rule" is there.

Euclidean Norm in Cartesian Coordinates

- **A standard norm is the so-called Euclidean norm of n-vectors:**

$$|\mathbf{u}| = |(u_1, \ldots, u_n)| := \sqrt{\sum_{i=1}^{n} u_i^2}$$

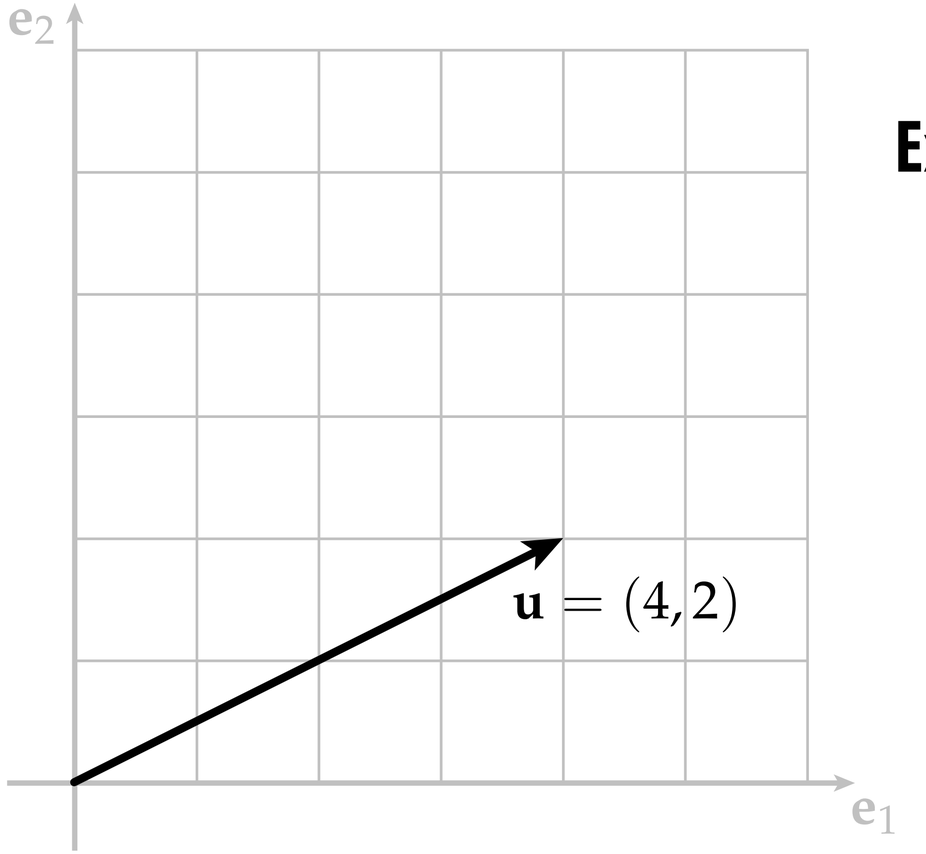

Example: $\mathbf{u} = (4,2)$

$$|\mathbf{u}| = \sqrt{4^2 + 2^2}$$

$$= 2\sqrt{5}$$

Q: Does this formula satisfy all the natural, geometric properties of a norm?

L^2 Norm of Functions

- Less familiar idea, but same basic intuition: the so-called L^2 norm measures the total magnitude of a function.
- Consider real-valued functions on the unit interval [0,1] whose square has a well-defined integral. The L^2 norm is defined as:

$$||f|| := \sqrt{\int_0^1 f(x)^2 \, dx}$$

- Not too different from the Euclidean norm: we just replaced a sum with an integral (which is kind of like a sum...).

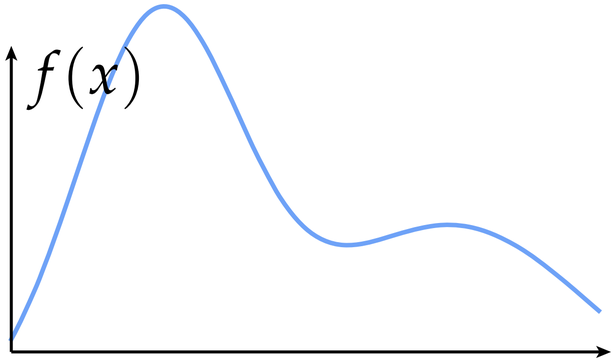

large L^2 norm

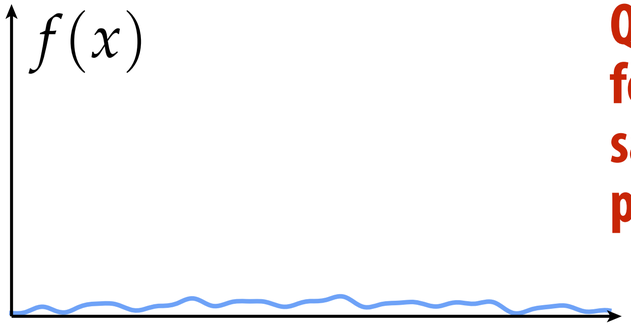

small L^2 norm

Q: Careful—does the formula above exactly satisfy all our desired properties for a norm?

L^2 Norm of Functions—Example

- Consider the function $f(x) := x\sqrt{3}$, defined over the unit interval [0,1].
- Q: What is its L^2 norm?
- A:

$$||f|| := \sqrt{\int_0^1 f(x)^2 \, dx}$$

$$||f||^2 = \int_0^1 3x^2 \, dx = \Big[x^3 \Big]_0^1 = 1^3 - 0^3 = 1$$

$$\Rightarrow ||f|| = \sqrt{1} = 1.$$

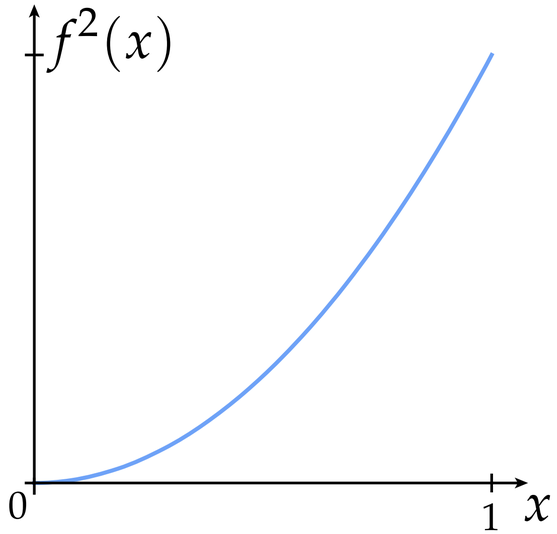

For clarity we will use $||\bullet||$ for the norm of a function, and $|\bullet|$ for the norm of a vector in R^n.

P.S. Most integrals in graphics are not calculated this way (at least not for more challenging functions, or functions described by data). Later on we'll talk a lot about numerical integration.

Inner Product—Motivation

- What else can we measure? In addition to magnitude, we said that vectors have orientation. Just as norm measured length, inner product measures how well vectors "line up."

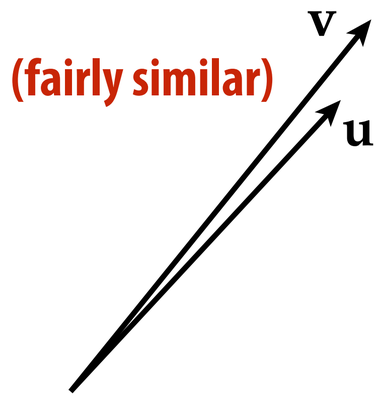

(fairly similar)

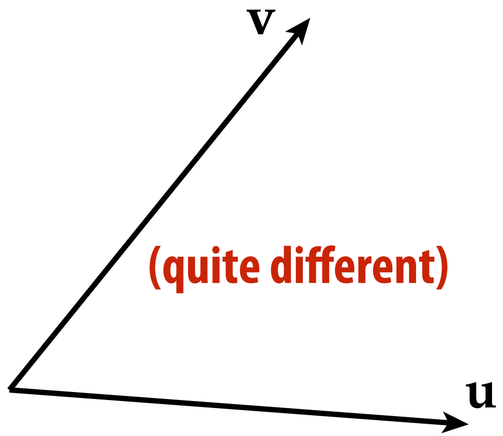

(quite different)

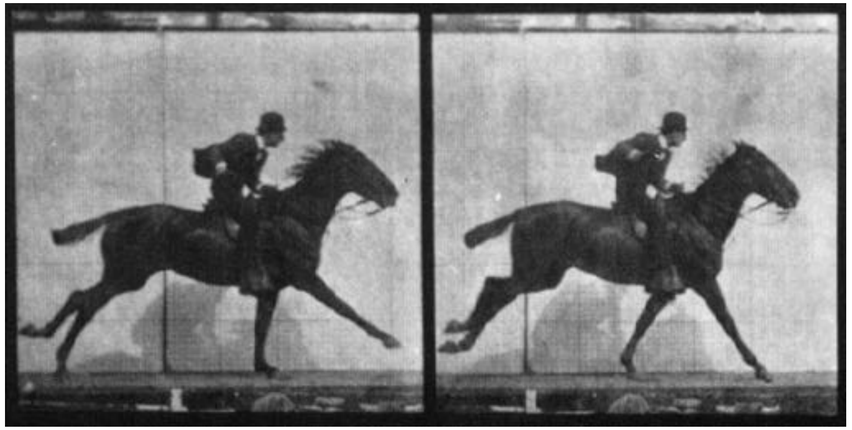

(fairly similar)

(quite different!)

Inner Product—Symmetry

- Will write inner product (also sometimes called the scalar product or dot product) using the notation <u,v> (some folks also write it as u•v).
- When measuring the alignment of two vectors u,v, what are some natural properties you might expect?
- One "obvious" property: order shouldn't matter, since u is just as well-aligned with v as v is with u:

$$\langle \mathbf{u}, \mathbf{v} \rangle = \langle \mathbf{v}, \mathbf{u} \rangle$$

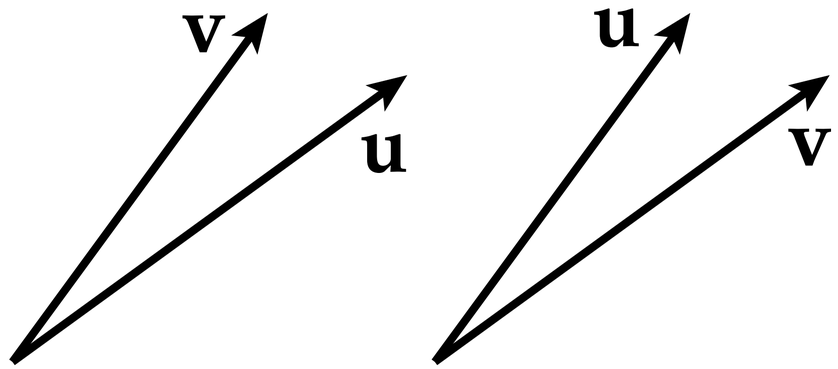

- Moreover, simply re-naming the vectors should have no effect on how well-aligned they are!

Inner Product—Projection & Scaling

- For unit vectors |u|=|v|=1, an inner product measures the extent of one vector along the direction of the other:

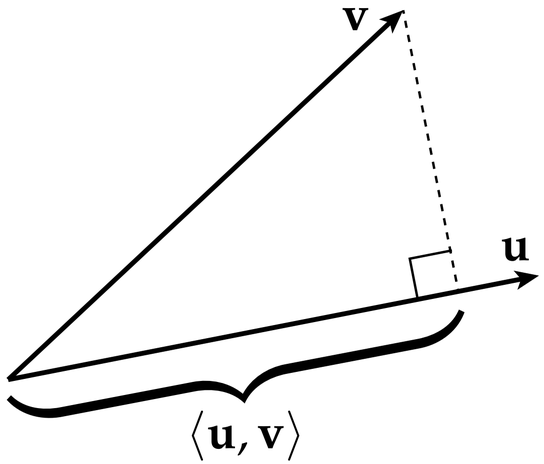

Q: Is this property symmetric? I.e., is the length of v along u the same as the length of u along v?

- If we scale either of the vectors, the inner product also scales:

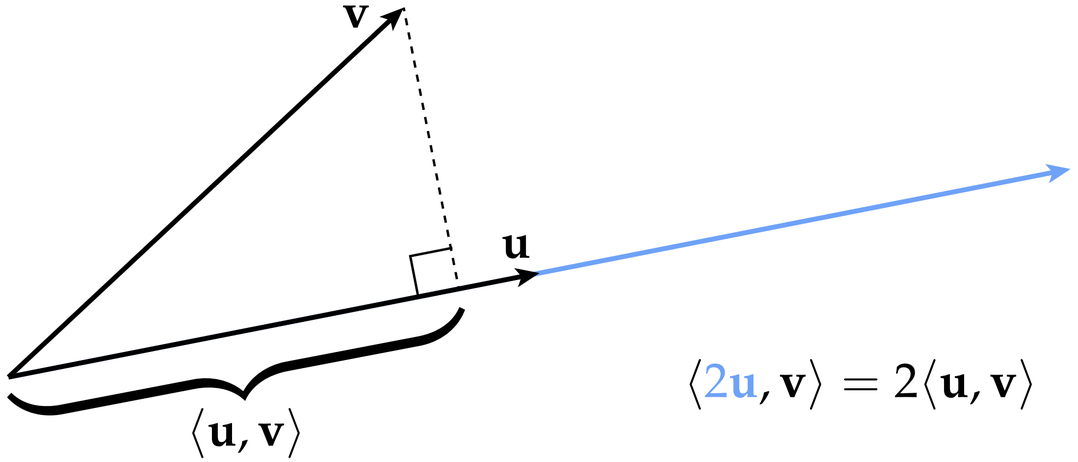

$$\langle 2\mathbf{u}, \mathbf{v} \rangle = 2\langle \mathbf{u}, \mathbf{v} \rangle$$

Inner Product—Positivity

- Also, a vector should always be aligned with itself, which we can express by saying that the inner product of a vector with itself should be positive (or at least, non-negative):

$$\langle \mathbf{u}, \mathbf{u} \rangle \geq 0$$

- In fact, if we continue to think of the inner product of a vector as the length of one vector along another then for unit-length vectors we must have

$$\langle \mathbf{u}, \mathbf{u} \rangle = 1$$

- Q: In general, then, what must <u,u> be equal to?
- A: Letting $\hat{\mathbf{u}} := \mathbf{u} / |\mathbf{u}|$, we have

$$\langle \mathbf{u}, \mathbf{u} \rangle = \langle |\mathbf{u}| \hat{\mathbf{u}}, |\mathbf{u}| \hat{\mathbf{u}} \rangle = |\mathbf{u}|^2 \langle \hat{\mathbf{u}}, \hat{\mathbf{u}} \rangle = |\mathbf{u}|^2 \cdot 1 = |\mathbf{u}|^2$$

Inner Product—Formal Definition

- An inner product is any function that assigns to any two vectors u,v a number <u,v> satisfying the following properties:
  - $\langle \mathbf{u}, \mathbf{v} \rangle = \langle \mathbf{v}, \mathbf{u} \rangle$
  - $\langle \mathbf{u}, \mathbf{u} \rangle \geq 0$
  - $\langle \mathbf{u}, \mathbf{u} \rangle = 0 \quad \iff \quad \mathbf{u} = \mathbf{0}$
  - $\langle a\mathbf{u}, \mathbf{v} \rangle = a\langle \mathbf{u}, \mathbf{v} \rangle$
  - $\langle \mathbf{u} + \mathbf{v}, \mathbf{w} \rangle = \langle \mathbf{u}, \mathbf{w} \rangle + \langle \mathbf{v}, \mathbf{w} \rangle$
- Q: Which of these properties didn't we talk about? Can you argue that they make sense geometrically?

Inner Product in Cartesian Coordinates

- **A standard inner product is the so-called Euclidean inner product, which operates on a pair of n-vectors:**

$$\langle \mathbf{u}, \mathbf{v} \rangle = \langle (u_1, \ldots, u_n), (v_1, \ldots, v_n) \rangle := \sum_{i=1}^{n} u_i v_i$$

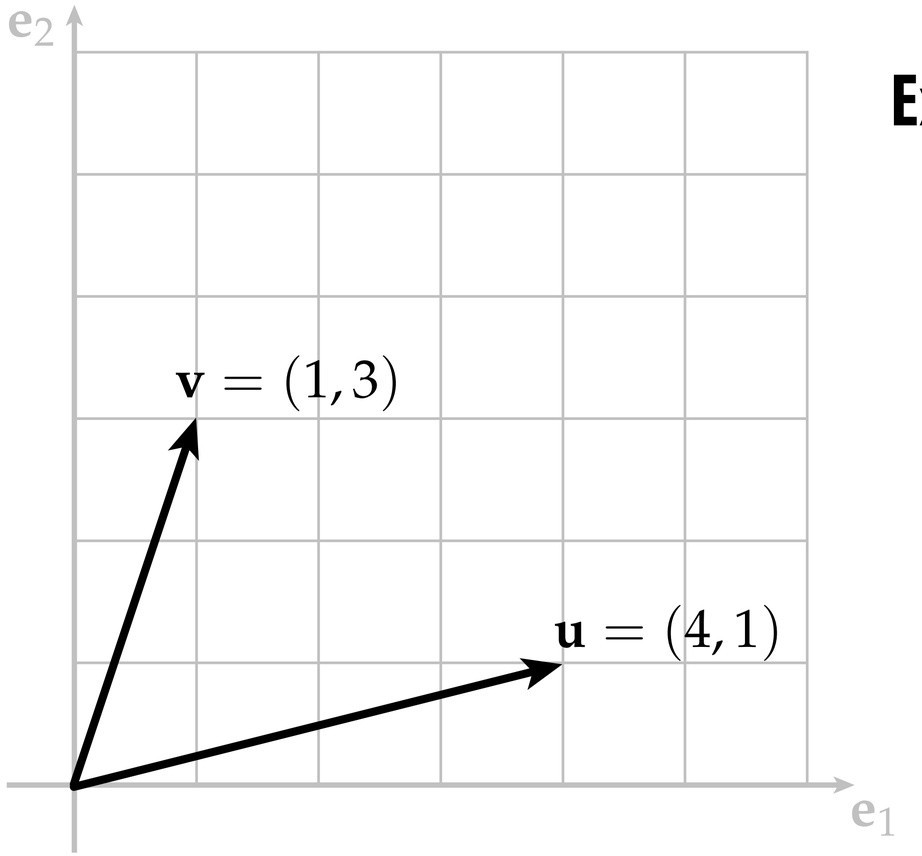

Example:

$$\mathbf{u} = (4,1)$$
$$\mathbf{v} = (1,3)$$

$$\langle \mathbf{u}, \mathbf{v} \rangle = 4 \cdot 1 + 1 \cdot 3 = 7$$

L^2 Inner Product of Functions—Example

- Just like we had a norm for functions, we can also define an inner product that measures how well two functions "line up".
- E.g., for square-integrable functions on the unit interval:

$$\langle\langle f, g \rangle\rangle := \int_0^1 f(x) g(x) \, dx$$

Example: $f(x) := x^2, \quad g(x) := (1-x)^2$

$$\langle\langle f, g \rangle\rangle = \int_0^1 x^2 (1-x)^2 \, dx = \cdots = \frac{1}{30}$$

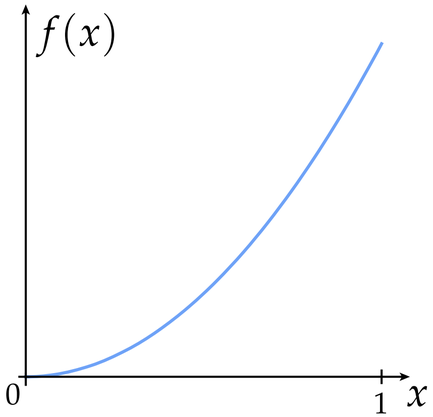

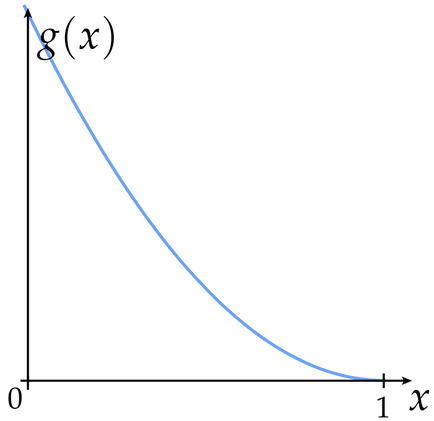

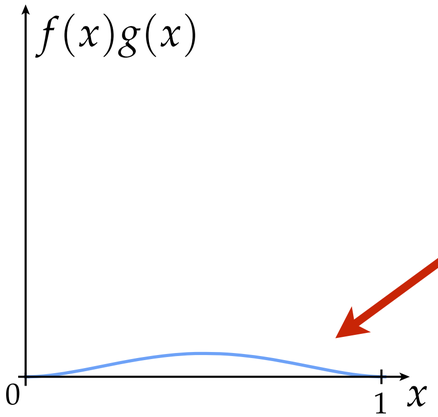

small number; functions don't "line up" much!

Measuring Images, Other Signals?

- Many ways to measure “how big” a signal is (norm) or “how well-aligned” two signals are (inner product).
- Choice of inner product depends on application.
- For instance, suppose we want images with “interesting stuff”
- Might try measuring norm of derivative (captures edges):

(dimmer) LARGER

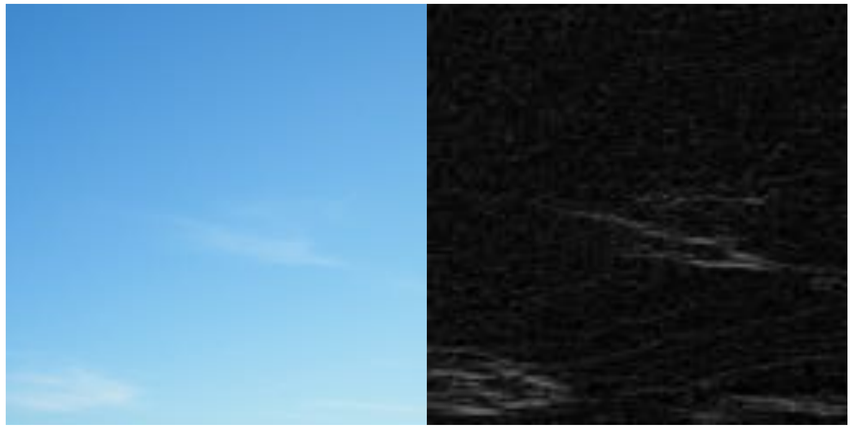

(brighter) SMALLER

Linear Maps

- At the beginning, said linear algebra was study of vector spaces and linear maps between them.
- Have a pretty good handle on vector (and inner product) spaces.
- But what's a linear map? And why is it useful for graphics?
- We'll get to the 1st question in a moment. As for the 2nd question, a few reasons:
  - Computationally, easy to solve systems of linear equations.
  - Basic transformations (rotation, translation, scaling) can be expressed as linear maps. (Will see this in a later lecture!)
  - All maps can be approximated as linear maps over a short distance/short time. (Taylor's theorem). This approximation is used all over geometry, animation, rendering, image processing...

Linear Maps—Geometric Definition

- What is a linear map?
- Especially in graphics, can think about them visually.
- Example:

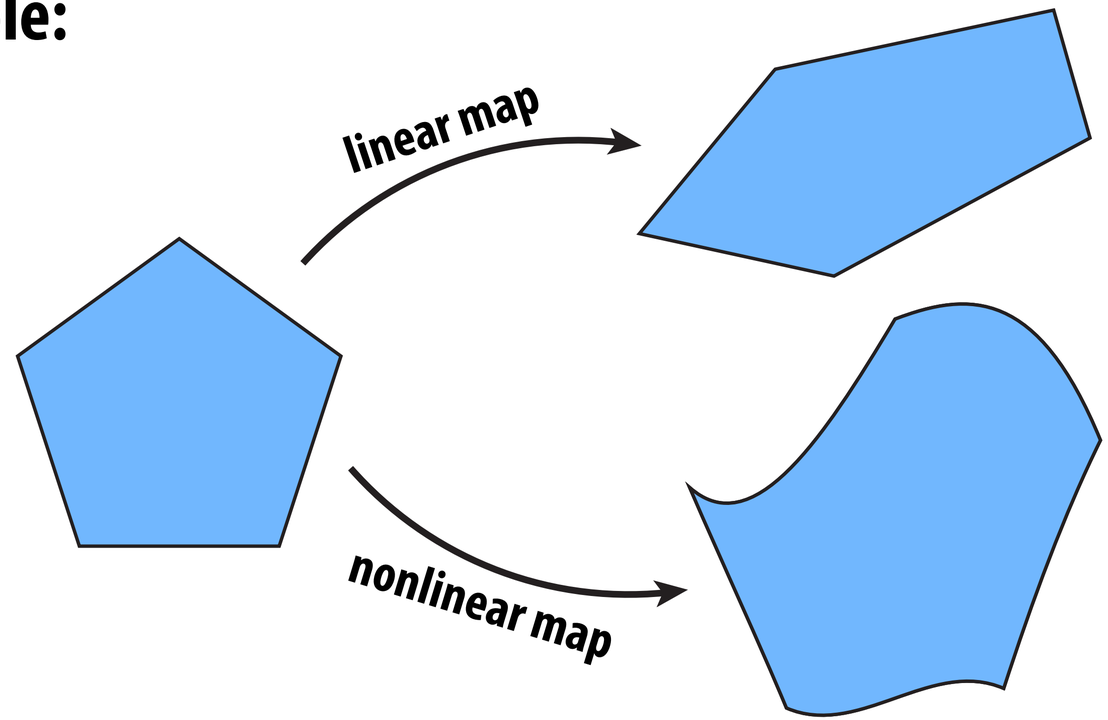

Key idea: linear maps take lines to lines*

Linear Maps—Geometric Definition

- What is a linear map?
- Especially in graphics, can think about them visually.
- Example:

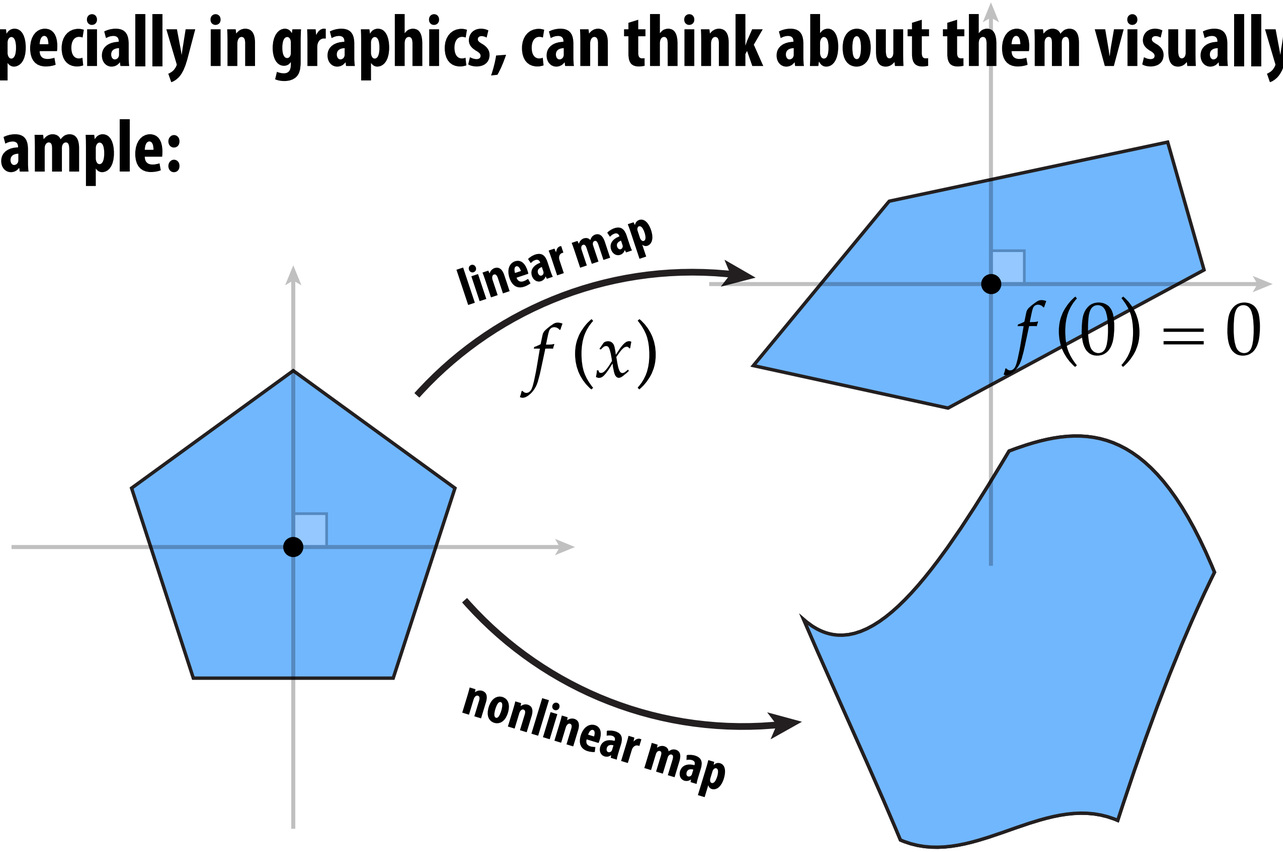

Key idea: linear maps take lines to lines*

*...while keeping the origin fixed.

Linear Maps—Algebraic Definition

- A map f is linear if it maps vectors to vectors, and if for all vectors u,v and scalars a we have

$$f(\mathbf{u}+\mathbf{v}) = f(\mathbf{u}) + f(\mathbf{v})$$
$$f(a\mathbf{u}) = af(\mathbf{u})$$

- In other words: if it doesn't matter whether we add the vectors and then apply the map, or apply the map and then add the vectors (and likewise for scaling):

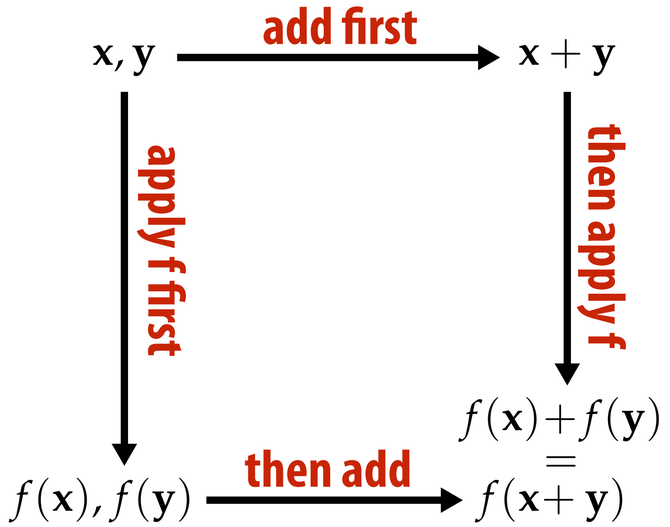

Linear Maps in Coordinates

- For maps between $\mathbb{R}^m$ and $\mathbb{R}^n$ (e.g., a map from 2D to 3D), we can give an even more explicit definition.
- A map is linear if it can be expressed as

$$f(u_1, \ldots, u_m) = \sum_{i=1}^{m} u_i \mathbf{a}_i$$

- In other words, if it is a linear combination of a fixed set of vectors $\mathbf{a}_i$:

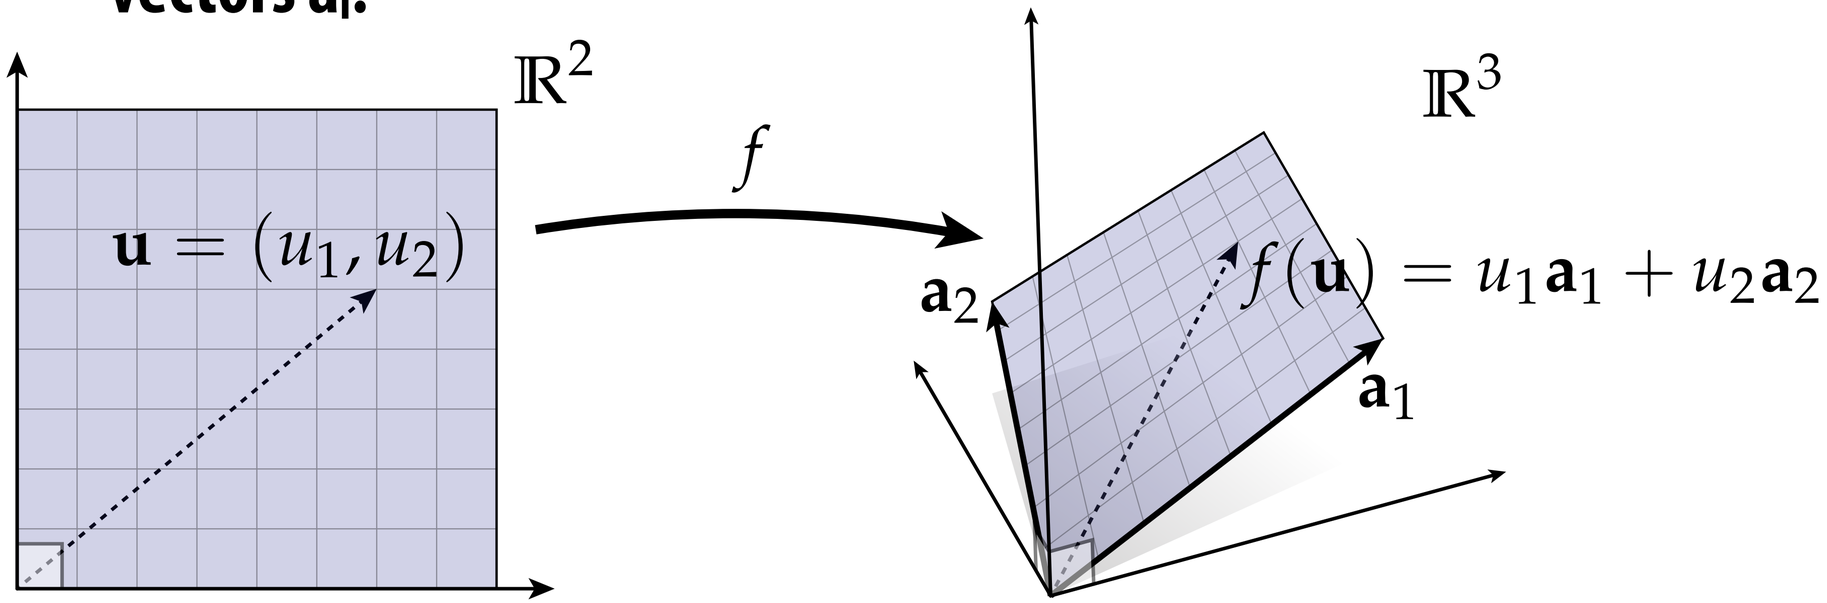

Q: Is $f(x) := ax + b$ a linear function?

Linear vs. Affine Maps

- No! But it's easy to be fooled, since the graph looks like a line:

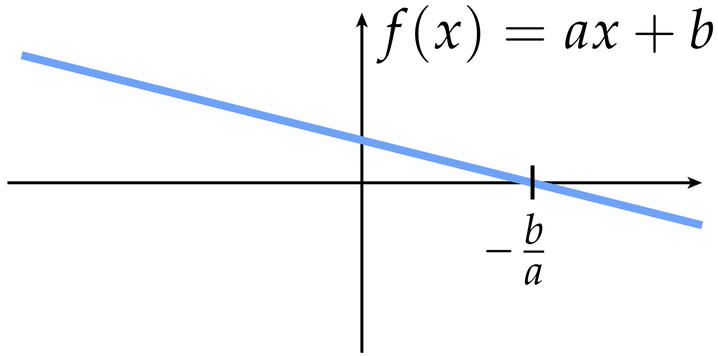

- However, it's not a line through the origin, i.e., f(0) ≠ 0.
- Another way to see it's not linear? Doesn't preserve sums:

$$f(x_1 + x_2) = a(x_1 + x_2) + b = ax_1 + ax_2 + b$$
$$f(x_1) + f(x_2) = (ax_1 + b) + (ax_2 + b) = ax_1 + ax_2 + 2b$$

- This function is called an AFFINE function (not a LINEAR one).
- Later we'll see an important computer graphics magic trick: turn affine functions (e.g., translation) into linear ones via homogeneous coordinates.

More interesting question:

Q: Is $f(u) := \int_0^1 u(x)\, dx$ **a linear map?**

(Think about it—it will be part of your homework!)

Span

Q: Geometrically, what is the span of two vectors u, v?

A: The span is the set of all vectors that can be written as a linear combination of u and v, i.e., vectors of the form

$$a\mathbf{u} + b\mathbf{v}$$

for any two numbers a, b**.**

More generally: $\operatorname{span}(\mathbf{u}_1, \ldots, \mathbf{u}_k) = \left\{ \mathbf{x} \in V \;\middle|\; \mathbf{x} = \sum_{i=1}^{k} a_i \mathbf{u}_i,\ a_1, \ldots, a_k \in \mathbb{R} \right\}$

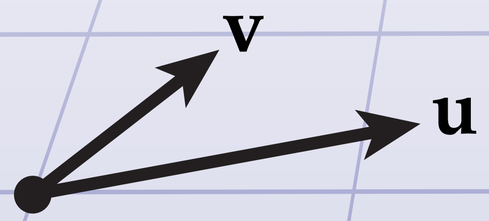

Span & Linear Maps

- Just a bit of language—can connect "span" and "linear map":
- "The image of any linear map is the span of some collection of vectors."

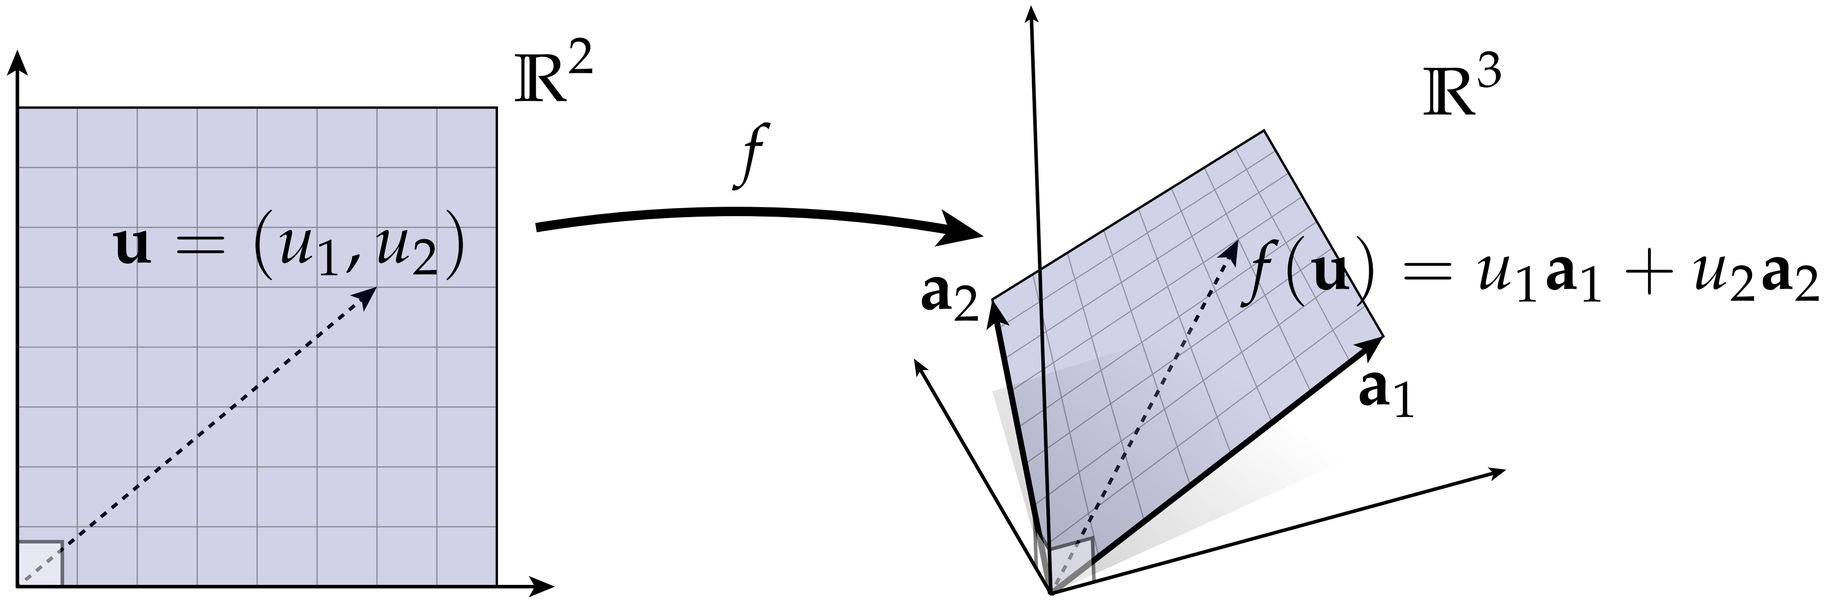

Q: What's the image of a function?

Basis

- Span is also closely related to the idea of a basis.
- In particular, if we have exactly n vectors e_1, ..., e_n such that

$$\text{span}(\mathbf{e}_1, \ldots, \mathbf{e}_n) = \mathbb{R}^n$$

- Then we say that these vectors are a basis for R^n.
- Note: many different choices of basis!
- Q: Which of the following are bases for the 2D plane (n=2)?

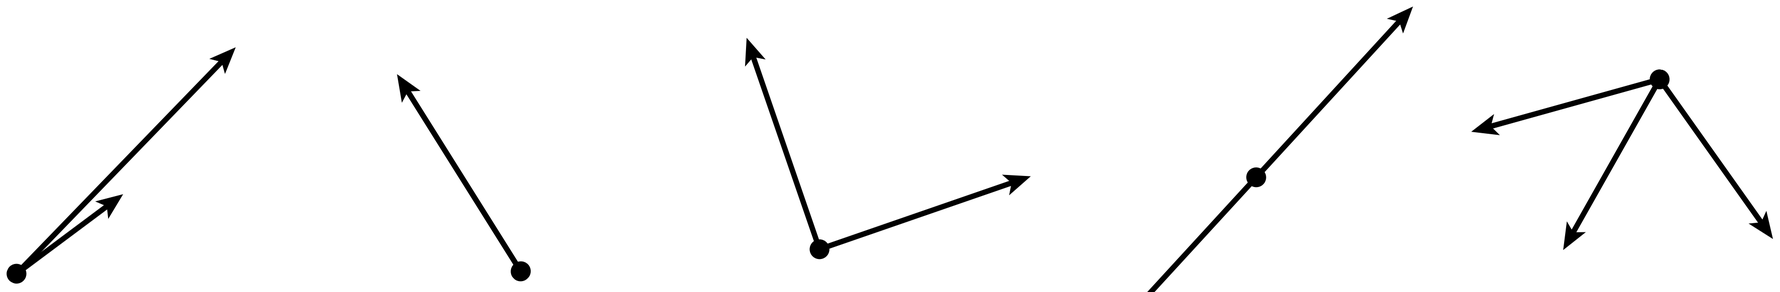

(A) (B) (C) (D) (E)

Orthonormal Basis

- Most often, it is convenient to have to basis vectors that are (i) unit length and (ii) mutually orthogonal.
- In other words, if $e_1, \ldots, e_n$ are our basis vectors then

$$\langle \mathbf{e}_i, \mathbf{e}_j \rangle = \begin{cases} 1, & i = j \\ 0, & \text{otherwise.} \end{cases}$$

- This way, the geometric meaning of the sum $u_1^2 + \ldots + u_n^2$ is maintained: it is the length of the vector u.

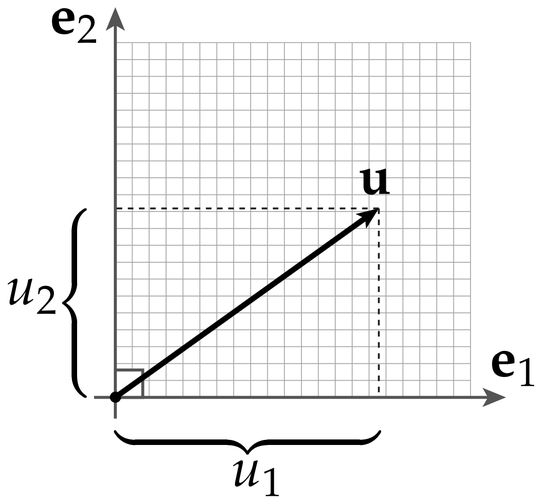

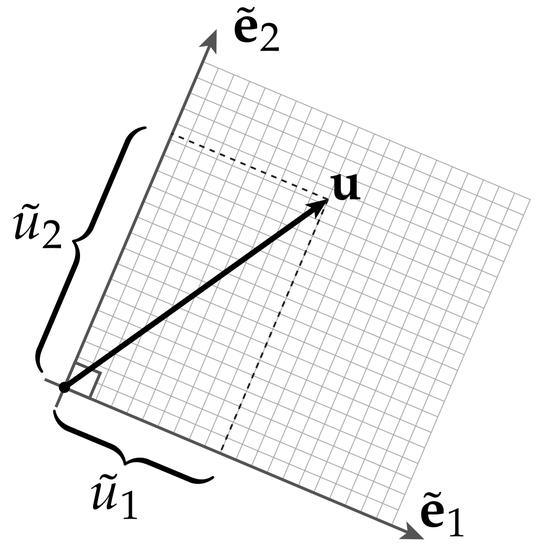

Common bug: projecting a vector onto a basis that is NOT orthonormal while continuing to use standard norm / inner product.

Gram-Schmidt

- Given a collection of basis vectors $a_1, \ldots a_n$, how do we find an orthonormal basis $e_1, \ldots e_n$,?
- Gram-Schmidt algorithm:
  - normalize the first vector (i.e., divide by its length)
  - subtract any component of the 1st vector from the 2nd one
  - normalize the 2nd one
  - repeat, removing components of first k vectors from vector k+1

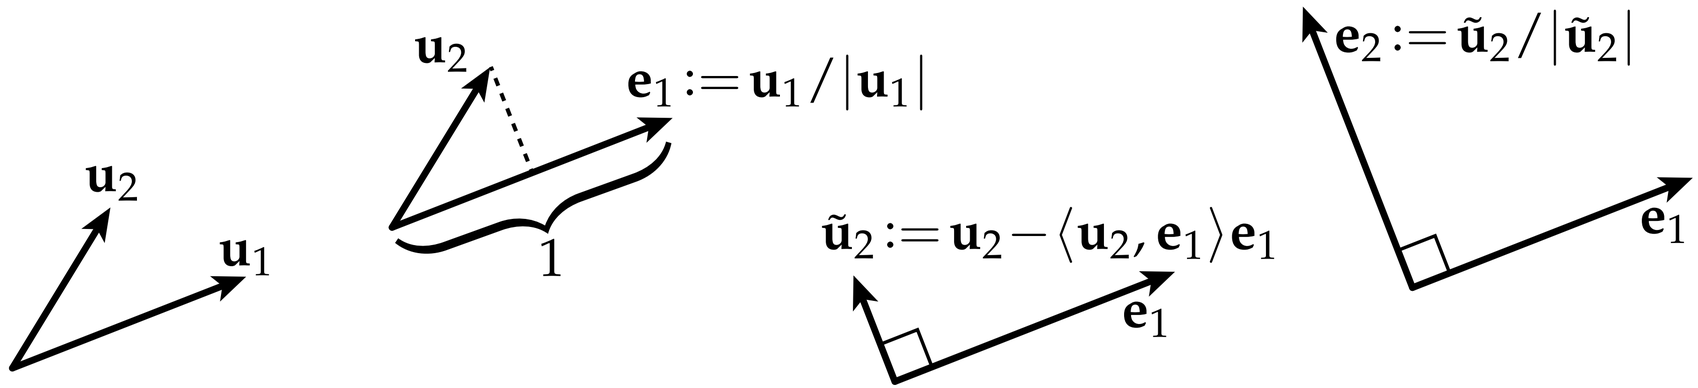

*WARNING: for large number of vectors / nearly parallel vectors, not the best algorithm...

Fourier Transform

- Functions are also vectors. Do they have an orthonormal basis?
- Yes! This is the basic idea behind the Fourier transform.
- Simple example: functions that repeat at intervals of 2π:

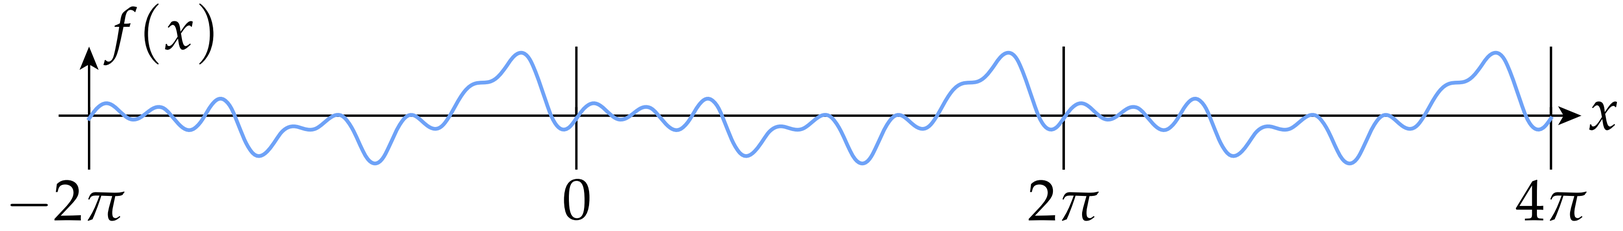

- Can project onto basis of sinusoids: $\cos(nx), \sin(mx),\ m,n \in \mathbb{N}$
  - really just a linear map from one basis to another
  - fundamental building block for many graphics algorithms

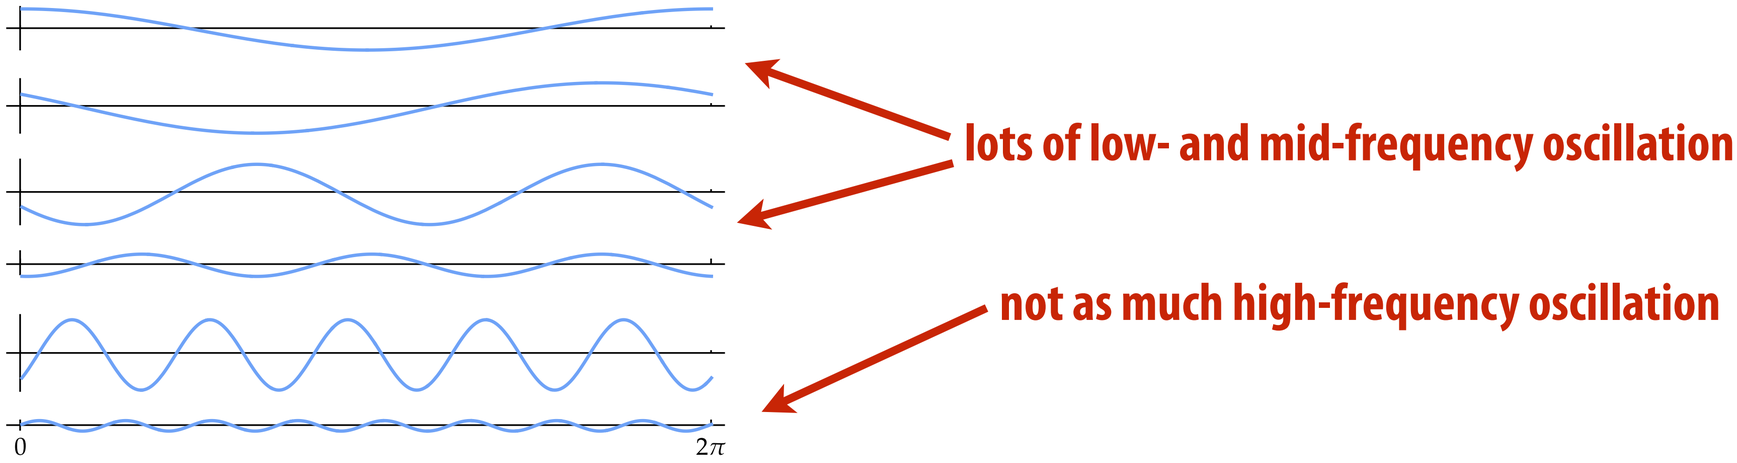

Frequency Decomposition of Signals

- More generally, this idea of projecting a signal onto different "frequencies" is known as Fourier decomposition
- Can be applied to all sorts of signals; basic tool used across, image processing, rendering, geometry, physical simulation…
- Will have plenty more to say as course goes on!

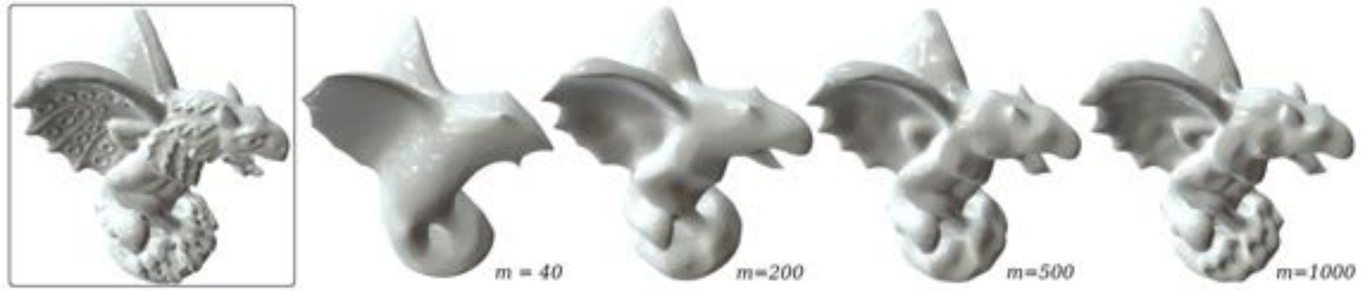

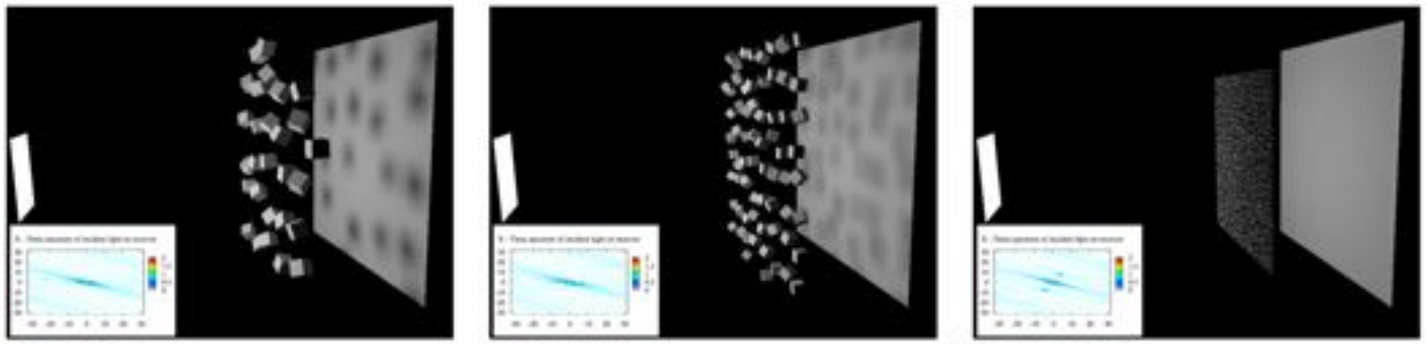

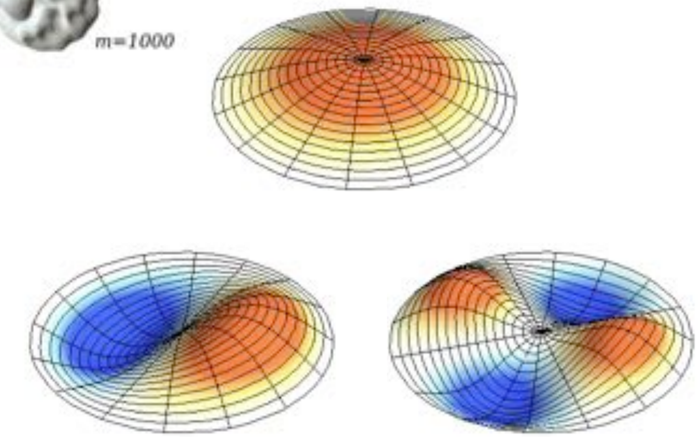

System of Linear Equations

- A system of linear equations is exactly what it sounds like: a bunch of equations where left-hand side is a linear function, right hand side is constant. E.g.,

$$\begin{aligned} x + 2y &= 3 \\ 4x + 5y &= 6 \end{aligned}$$

- Unknown values are sometimes called "degrees of freedom" (DOFs); equations sometimes called "constraints"
- Goal: solve for DOFs that simultaneously satisfy constraints:

$$x = 3 - 2y$$

$$4(3 - 2y) + 5y = 6$$

$$\boxed{\begin{aligned} y &= 2 \\ x &= -1 \end{aligned}}$$

What does solving a linear system mean?

Linear System, Visualized

- Of course, a linear system can be used to represent many different practical tasks (simulation, processing, etc.).
- But for any linear system, there are some good mental models to visualize:

Find the point where two lines meet:

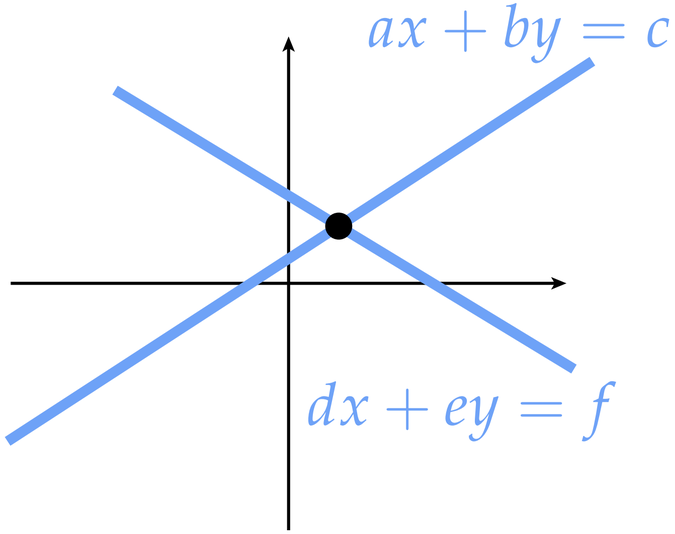

GIVEN a point b, FIND the point x that maps to it:

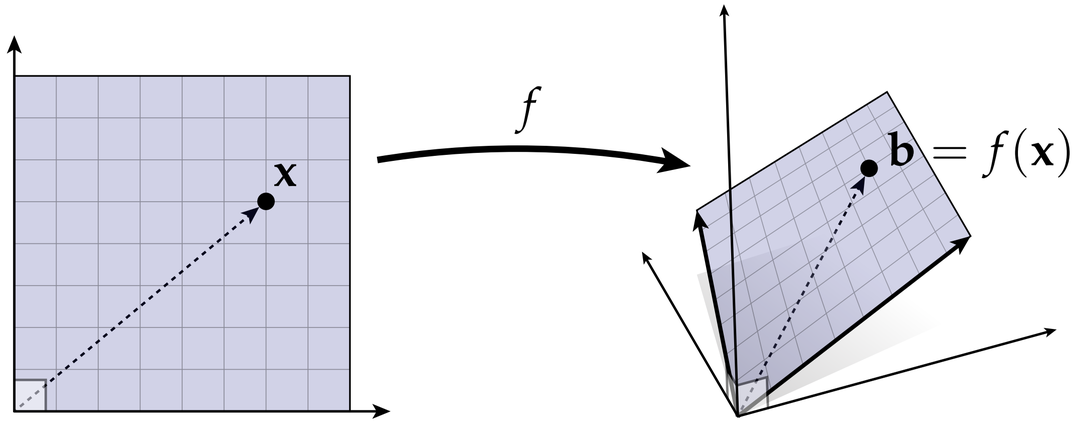

Uniqueness, Existence of Solutions

- Of course, not all linear systems can be solved! (And even those that can be solved may not have a unique solution.)

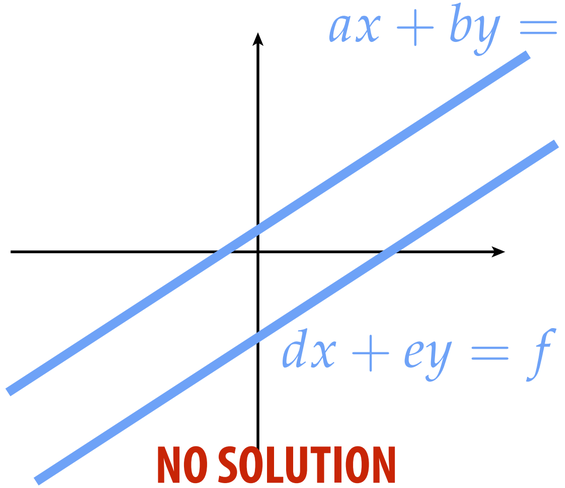

NO SOLUTION
equations disagree

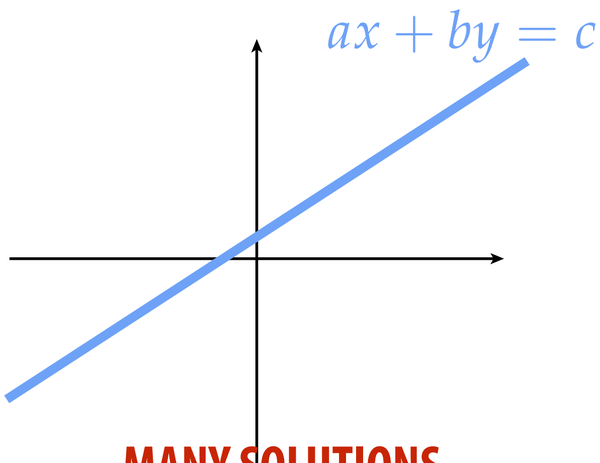

MANY SOLUTIONS
any point along the line

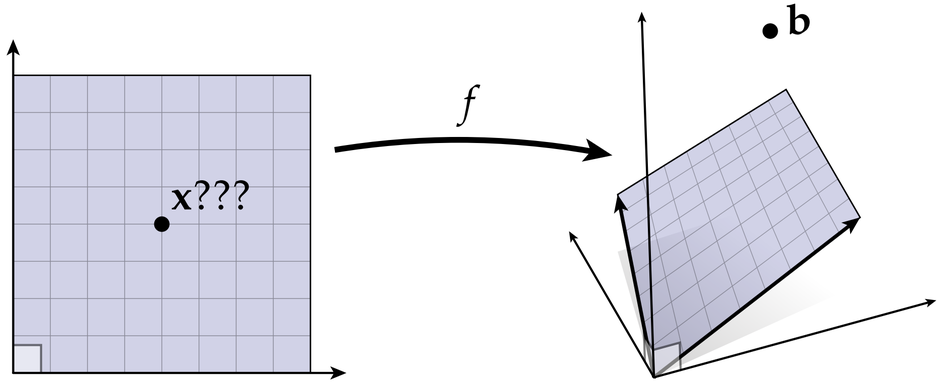

NO SOLUTION
not every point can be "reached" by f

Wait, what about matrices?!

Matrices in Linear Algebra

- Linear algebra often taught from the perspective of matrices, i.e., pushing around little blocks of numbers.
- But linear algebra is not fundamentally about matrices.
- As you've just seen, you can understand almost all the basic concepts without ever touching a matrix!
- Likewise, matrices can interfere with understanding / lead to confusion, since the same object (a block of numbers) is used to represent many different things (linear map, quadratic form, ...) in many different bases.
- Still, VERY useful!
  - symbolic manipulation
  - numerical computation

$$\begin{bmatrix} 1 & 7 & 3 \\ 4 & 9 & 2 \\ 0 & 1 & 1 \end{bmatrix}$$

What does this thing mean/ encode/do/represent?

Representing Linear Maps via Matrices

- Key example: suppose I have a linear map

$$f(\mathbf{u}) = u_1\mathbf{a}_1 + u_2\mathbf{a}_2$$

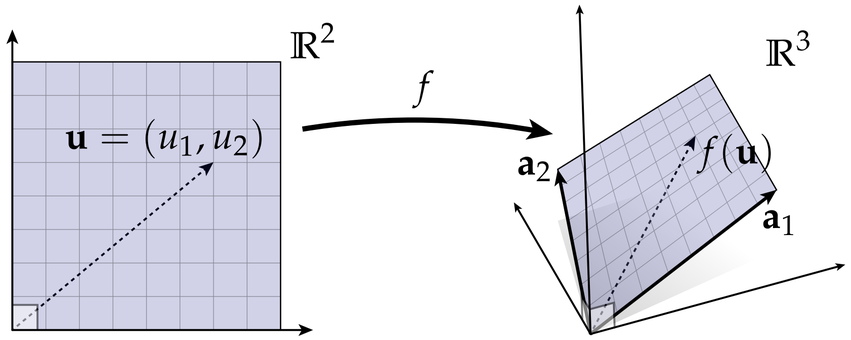

- How do I encode as a matrix?
- Easy: "a" vectors become matrix columns:

$$A := \begin{bmatrix} a_{1,x} & a_{2,x} \\ a_{1,y} & a_{2,y} \\ a_{1,z} & a_{2,z} \end{bmatrix}$$

- Now, matrix-vector multiply recovers original map:

$$\begin{bmatrix} a_{1,x} & a_{2,x} \\ a_{1,y} & a_{2,y} \\ a_{1,z} & a_{2,z} \end{bmatrix} \begin{bmatrix} u_1 \\ u_2 \end{bmatrix} = \begin{bmatrix} a_{1,x}u_1 + a_{2,x}u_2 \\ a_{1,y}u_1 + a_{2,y}u_2 \\ a_{1,z}u_1 + a_{2,x}u_2 \end{bmatrix} = u_1\mathbf{a}_1 + u_2\mathbf{a}_2$$

Don’t worry: if you love matrices, there will be plenty of them in your homework!

Next time: Math (P)Review Part II
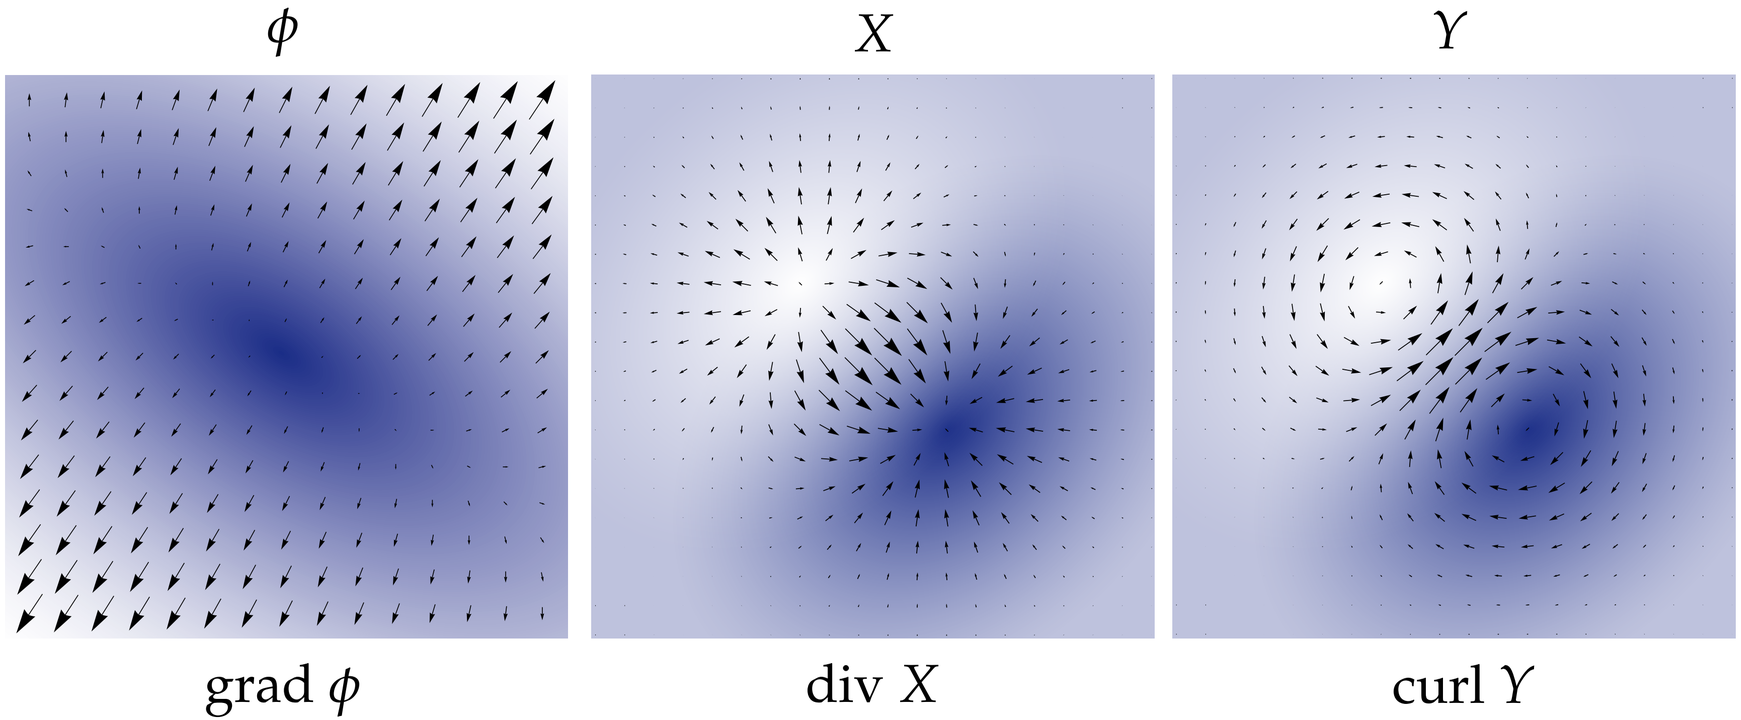
- Vector calculus
- Eigenvalue problems
- Complex numbers