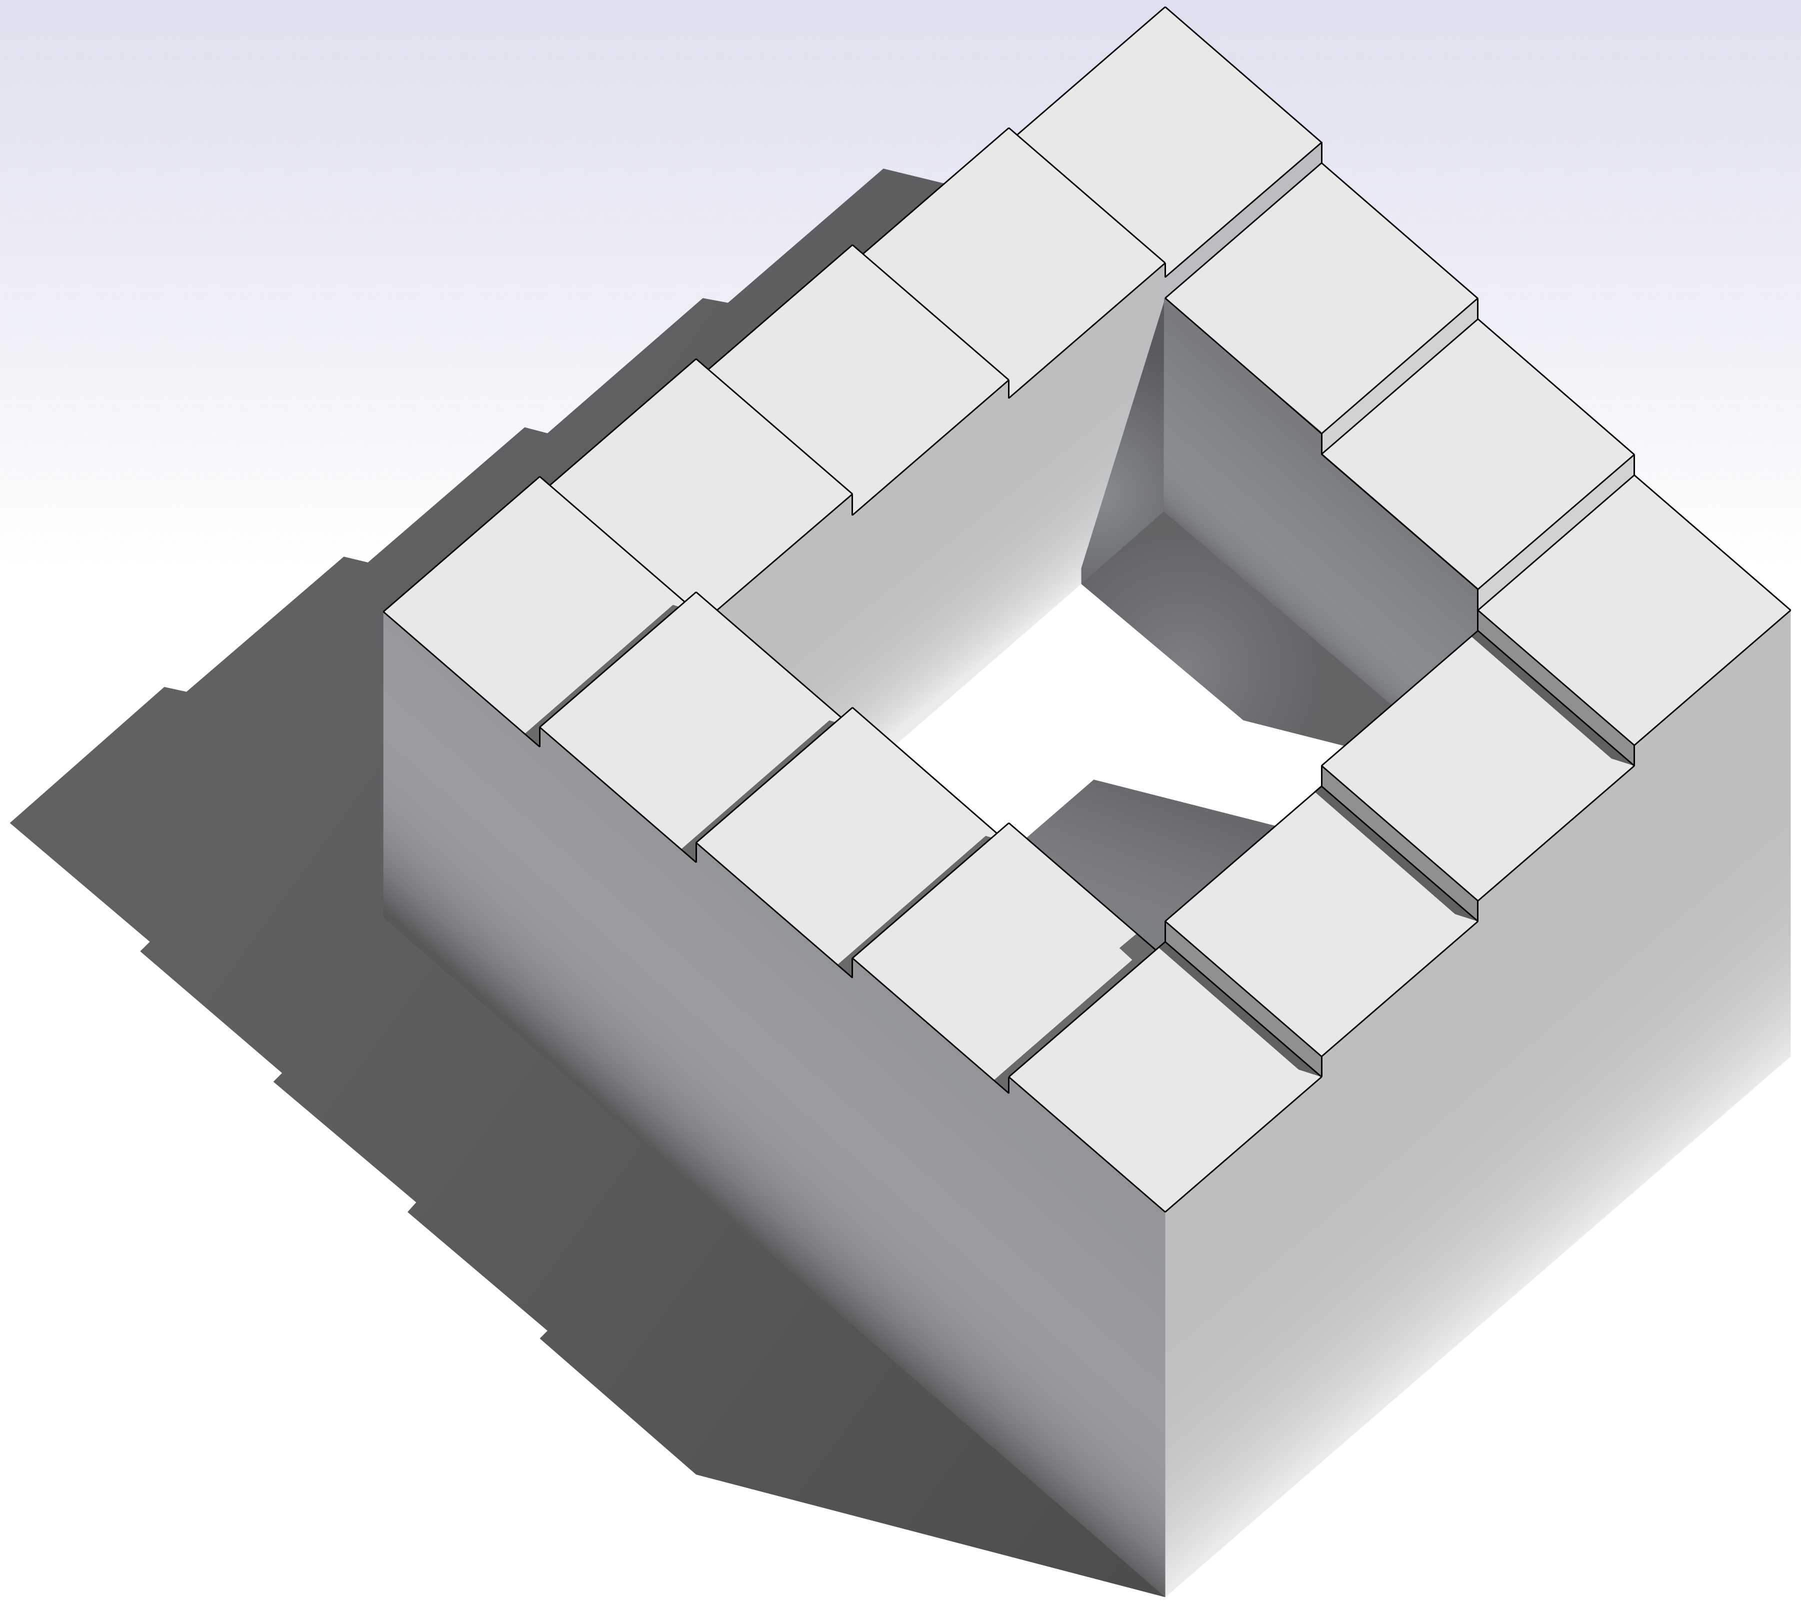

Discrete Differential Geometry: An Applied Introduction

Keenan Crane • CMU 15-458/858

Lecture 10:
Smooth Curves

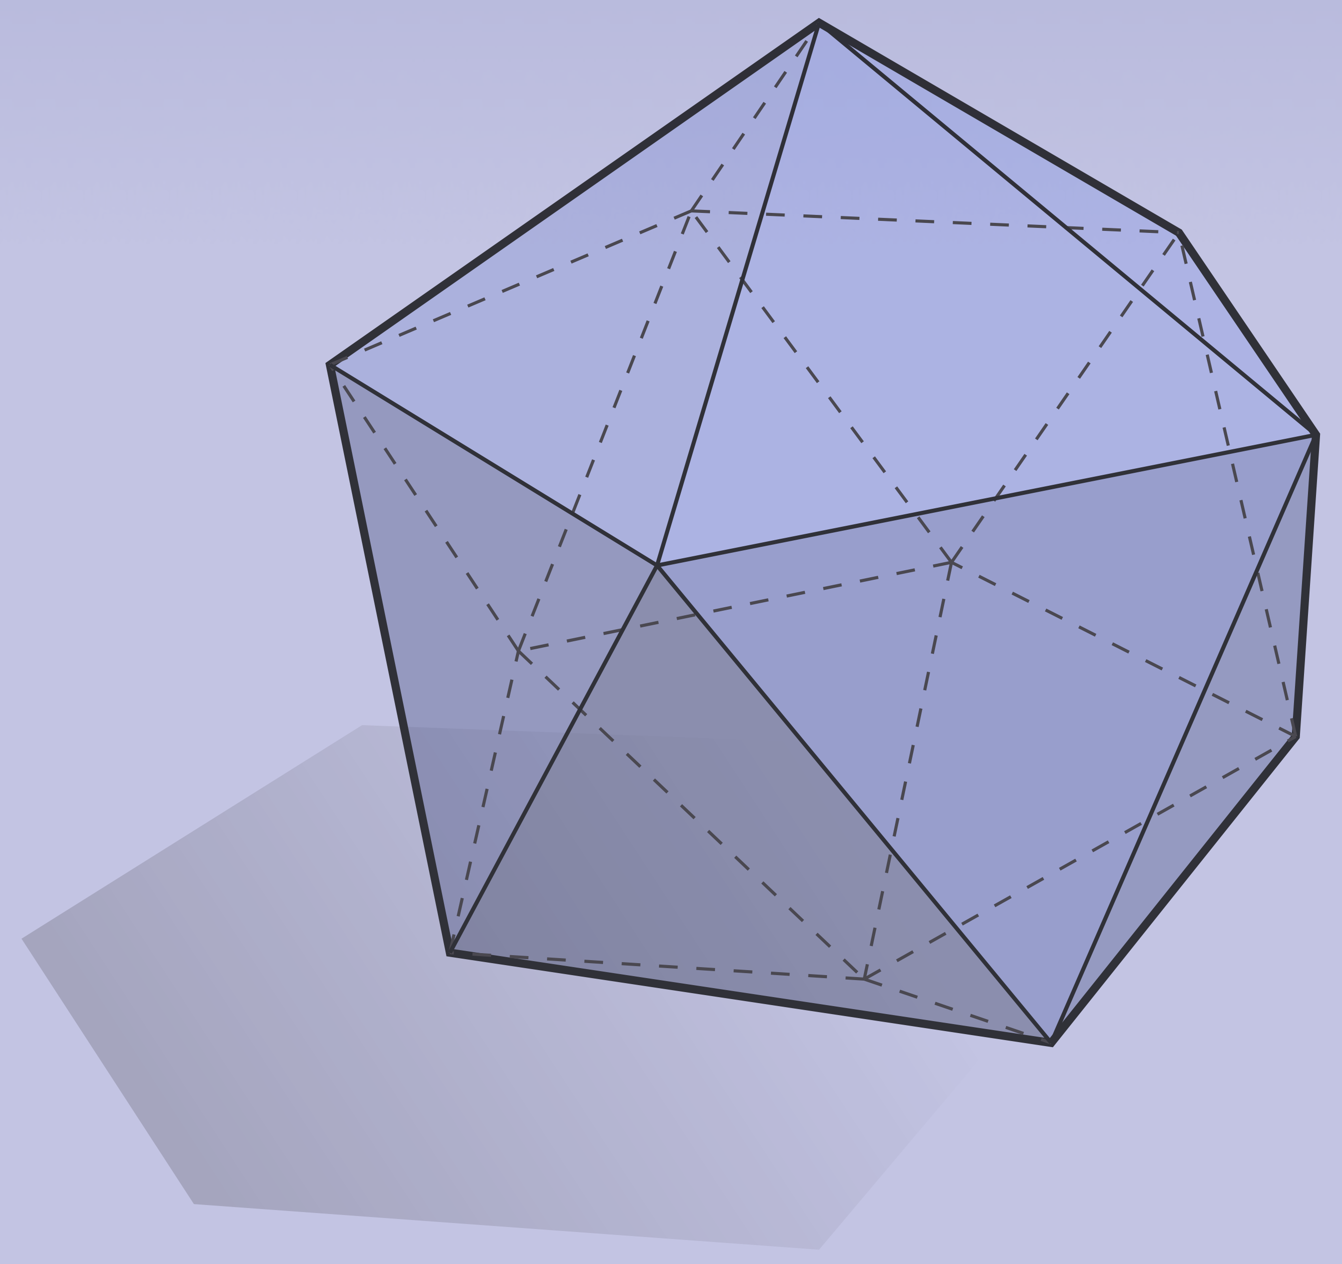

Discrete Differential Geometry:
An Applied Introduction

Keenan Crane • CMU 15-458/858

Curves, Surfaces, and Volumes

- In general, differential geometry studies *n*-dimensional manifolds; we'll focus mostly on low dimensions: curves (*n*=1), surfaces (*n*=2), and volumes (*n*=3)
- Why? Geometry we encounter in "every day life" (Common in applications!)
- Low-dimensional manifolds are not baby stuff! :-)
  - ***n*=1:** unknot recognition (open as of July 2017)
  - ***n*=2:** Willmore conjecture (2012 for genus 1)
  - ***n*=3:** Geometrization conjecture (2003, $1 million)

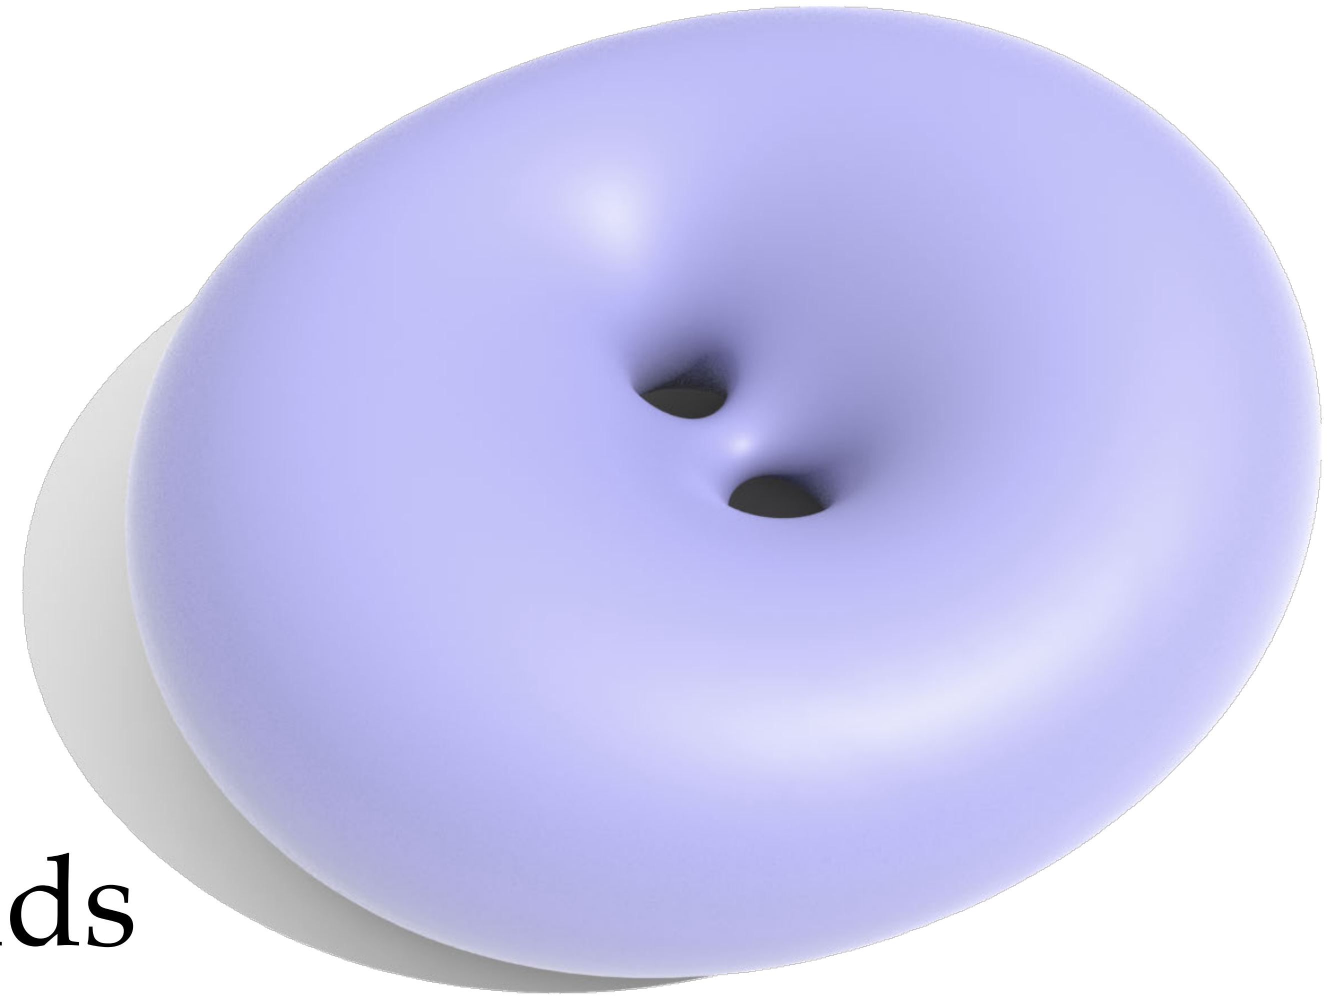

- Serious intuition gained by studying low-dimensional manifolds
- Conversely, problems involving very high-dimensional manifolds (e.g., statistics/ machine learning) involve less "deep" geometry than you might imagine!
  - *fiber bundles, Lie groups, curvature flows, spinors, symplectic structure, ...*
- Moreover... curves and surfaces are beautiful! (And sometimes boring for large *n*…)

Curves & Surfaces

- Much of the geometry we encounter in life well-described by *curves* and *surfaces**

(Curves)

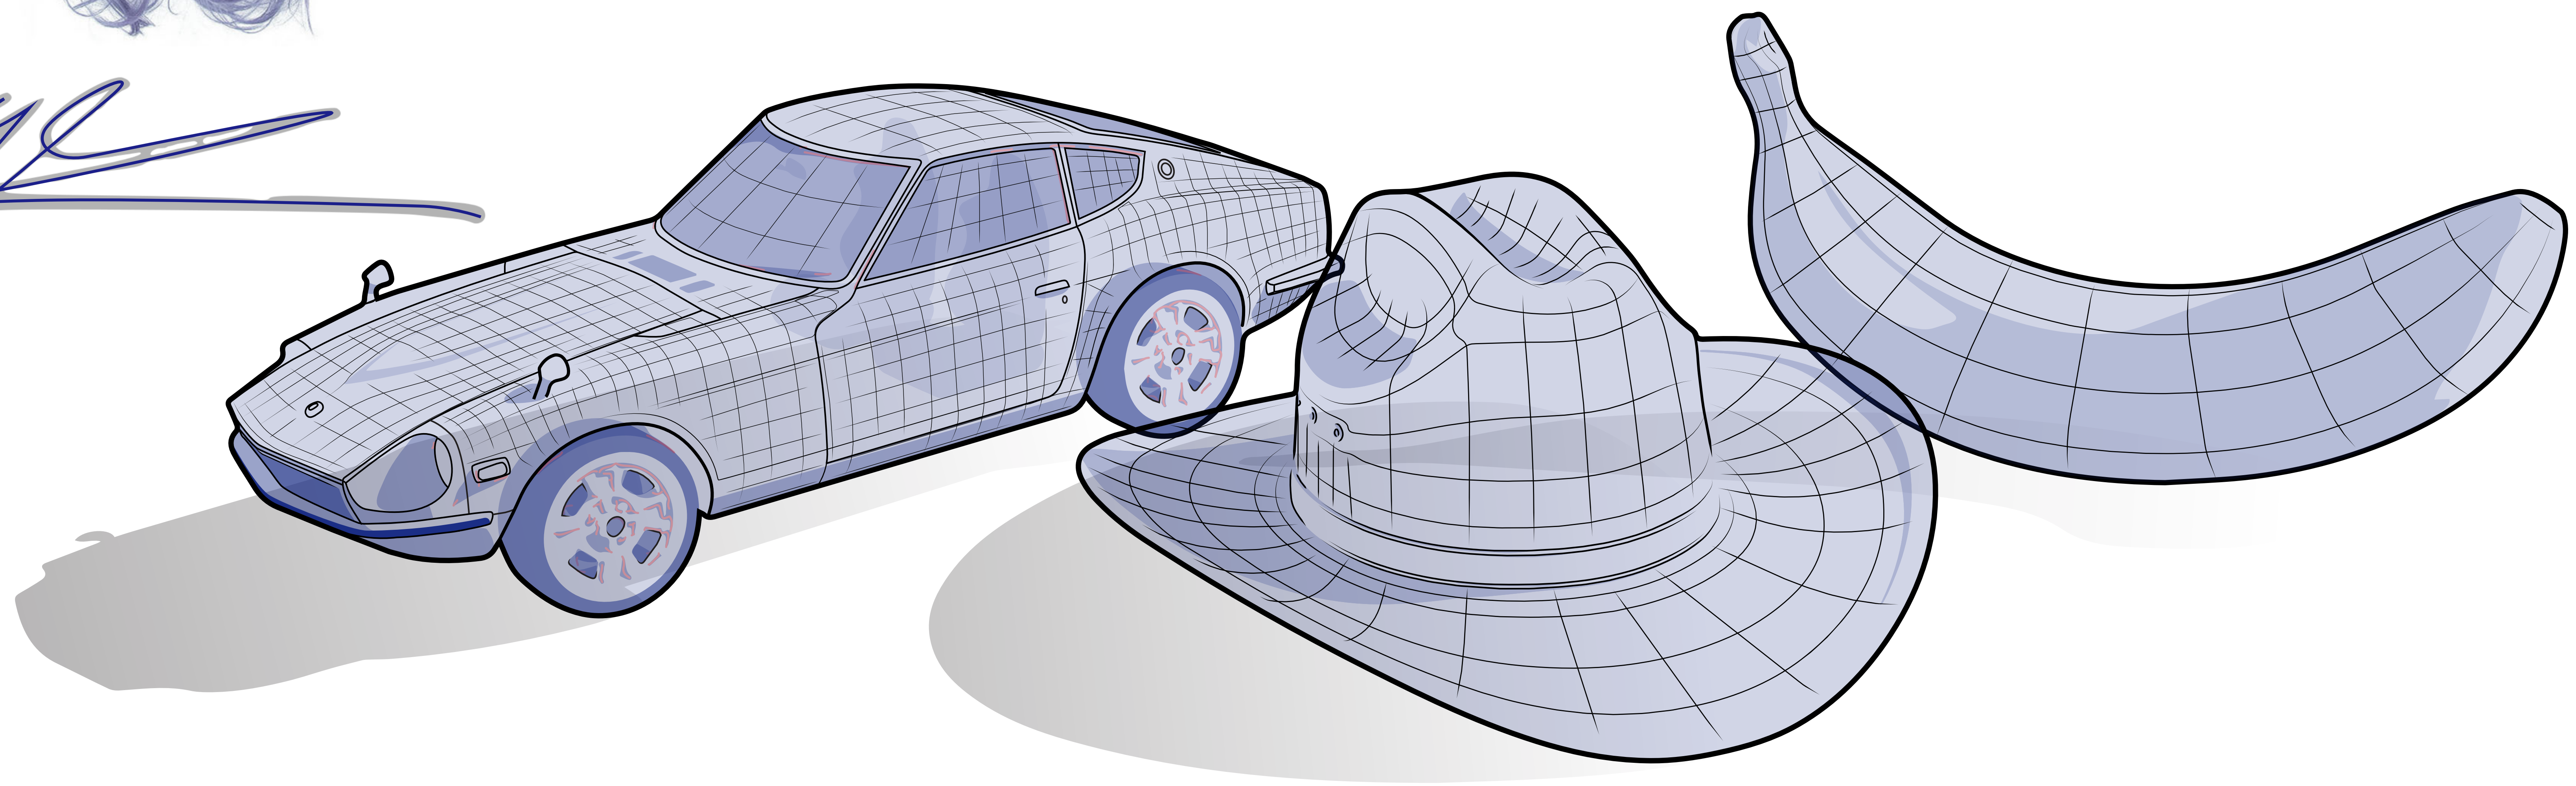

(Surfaces)

*Or solids… but the boundary of a solid is a surface!

Smooth Descriptions of Curves & Surfaces

- Many ways to express the geometry of a curve or surface:
  - height function over tangent plane
  - local parameterization
  - Christoffel symbols — coordinates/indices
  - **differential forms** — “coordinate free”
  - moving frames — change in *adapted frame*
  - Riemann surfaces (*local*); Quaternionic functions (*global*)
- People can get very religious about these different “dialects”... best to be multilingual!
- We'll dive deep into one description (**differential forms**) and touch on others

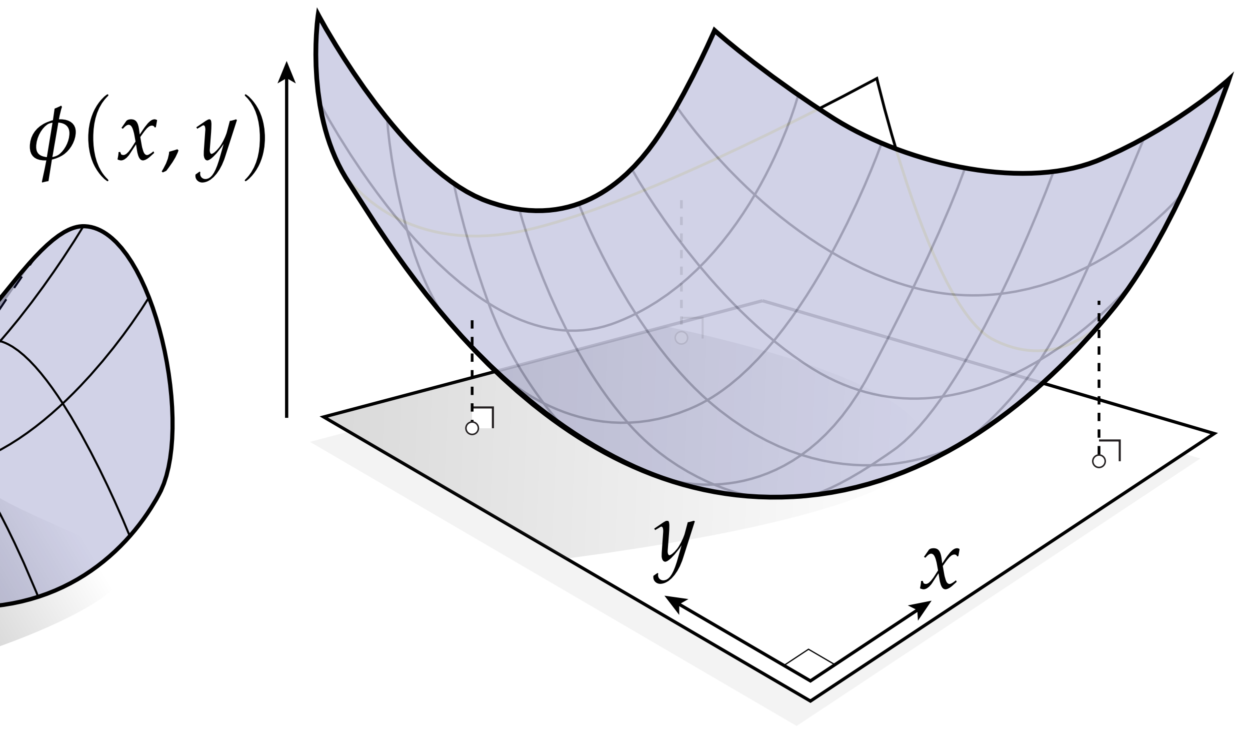

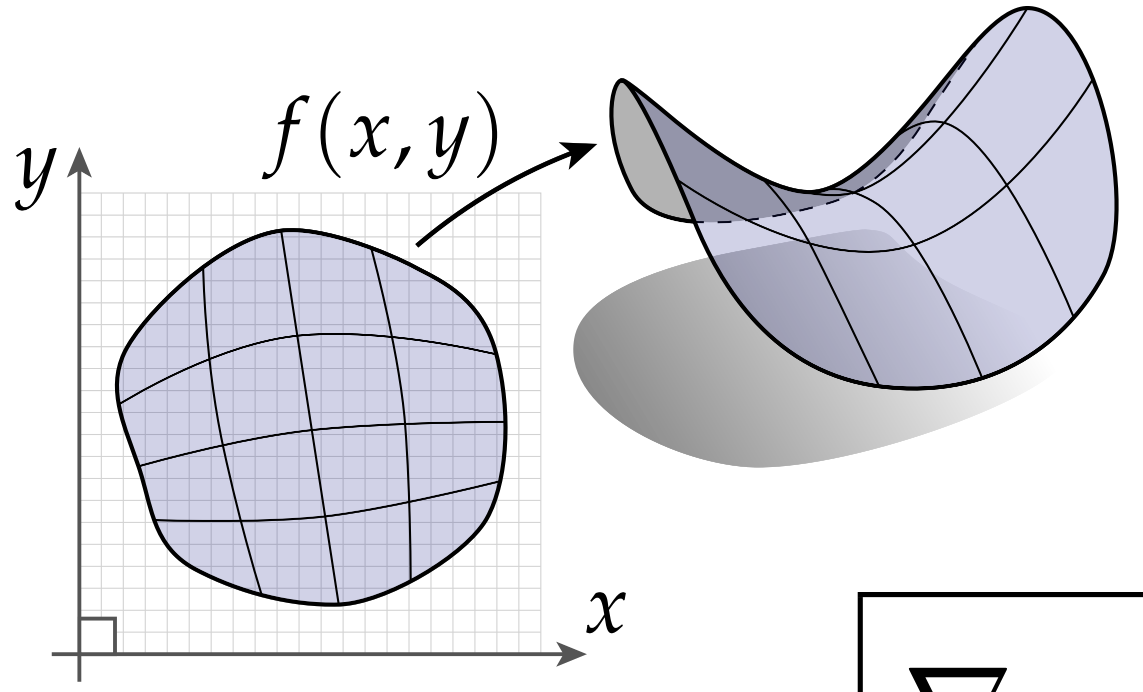

$$\nabla_i e_j = \Gamma_{ij}^k e_k$$

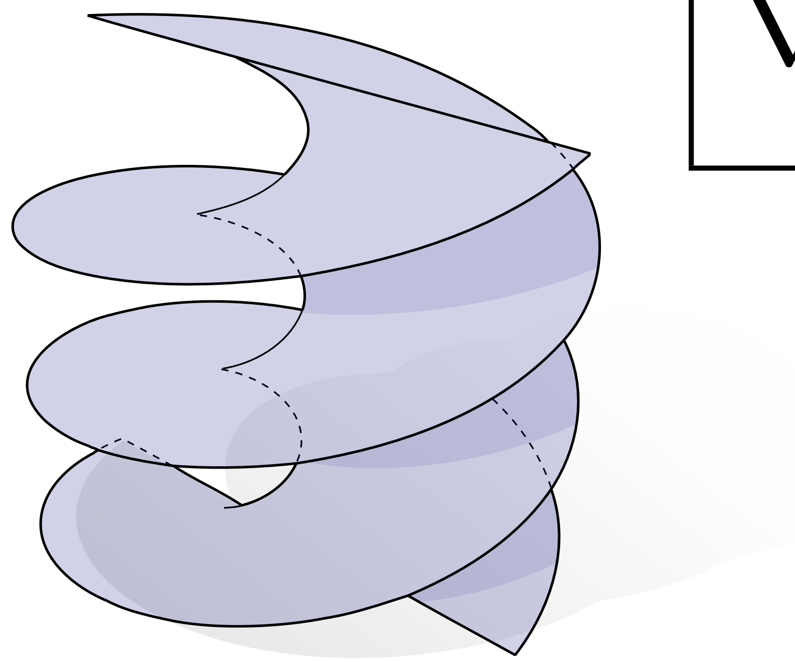

$$\begin{aligned} d\sigma &= \omega \wedge \sigma \\ d\omega &= \omega \wedge \omega \end{aligned}$$

Discrete Descriptions of Curves & Surfaces

- Also *many* ways to discretize a surface
- For instance:
  - **implicit** — *e.g.*, zero set of scalar function on a grid
    - good for changing topology, high accuracy
    - expensive to store/adaptivity is harder
    - hard to solve sophisticated equations *on* surface
  - **explicit** — *e.g.*, polygonal surface mesh
    - changing topology, high-order continuity is harder
    - cheaper to store / adaptivity is much easier
    - more mature tools for equations *on* surfaces
- Don't be "religious"; use the right tool for the job!

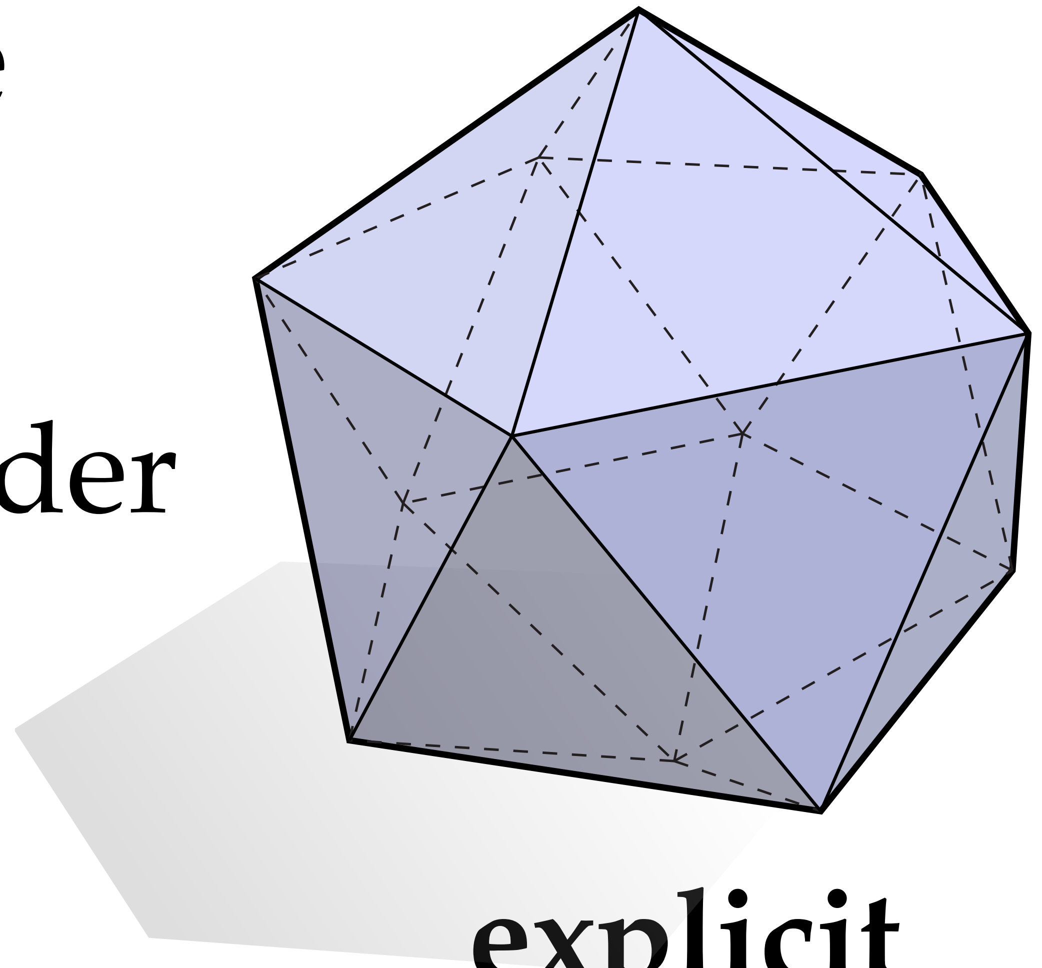

explicit

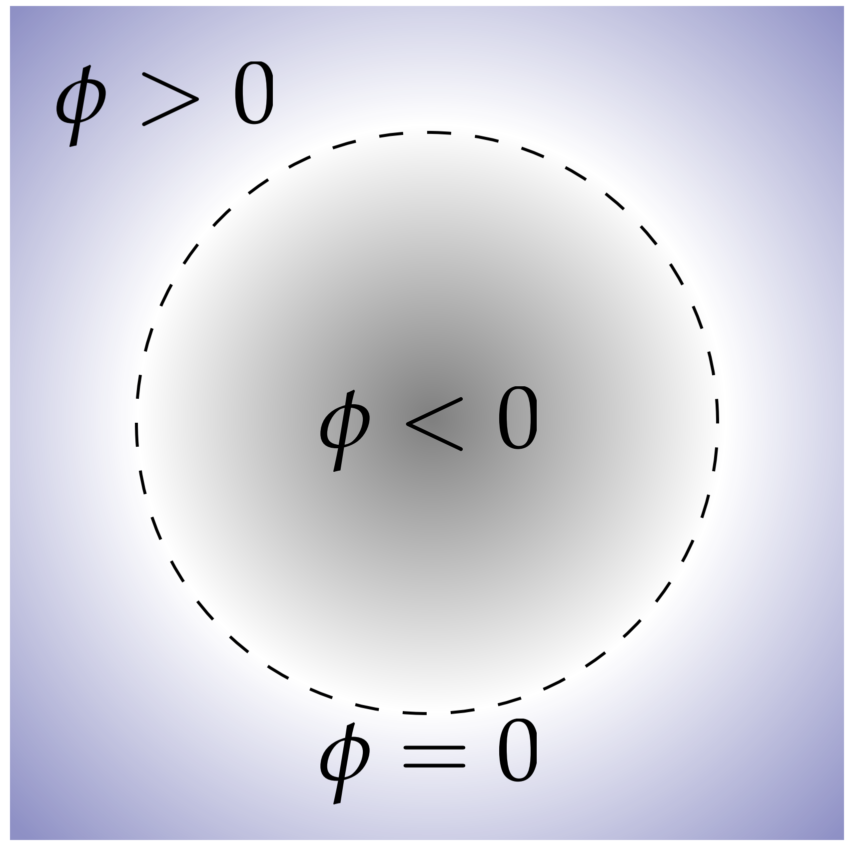

implicit

Curves & Surfaces—Overview

- **Goal:** understand curves & surfaces from complementary smooth and discrete points of view.

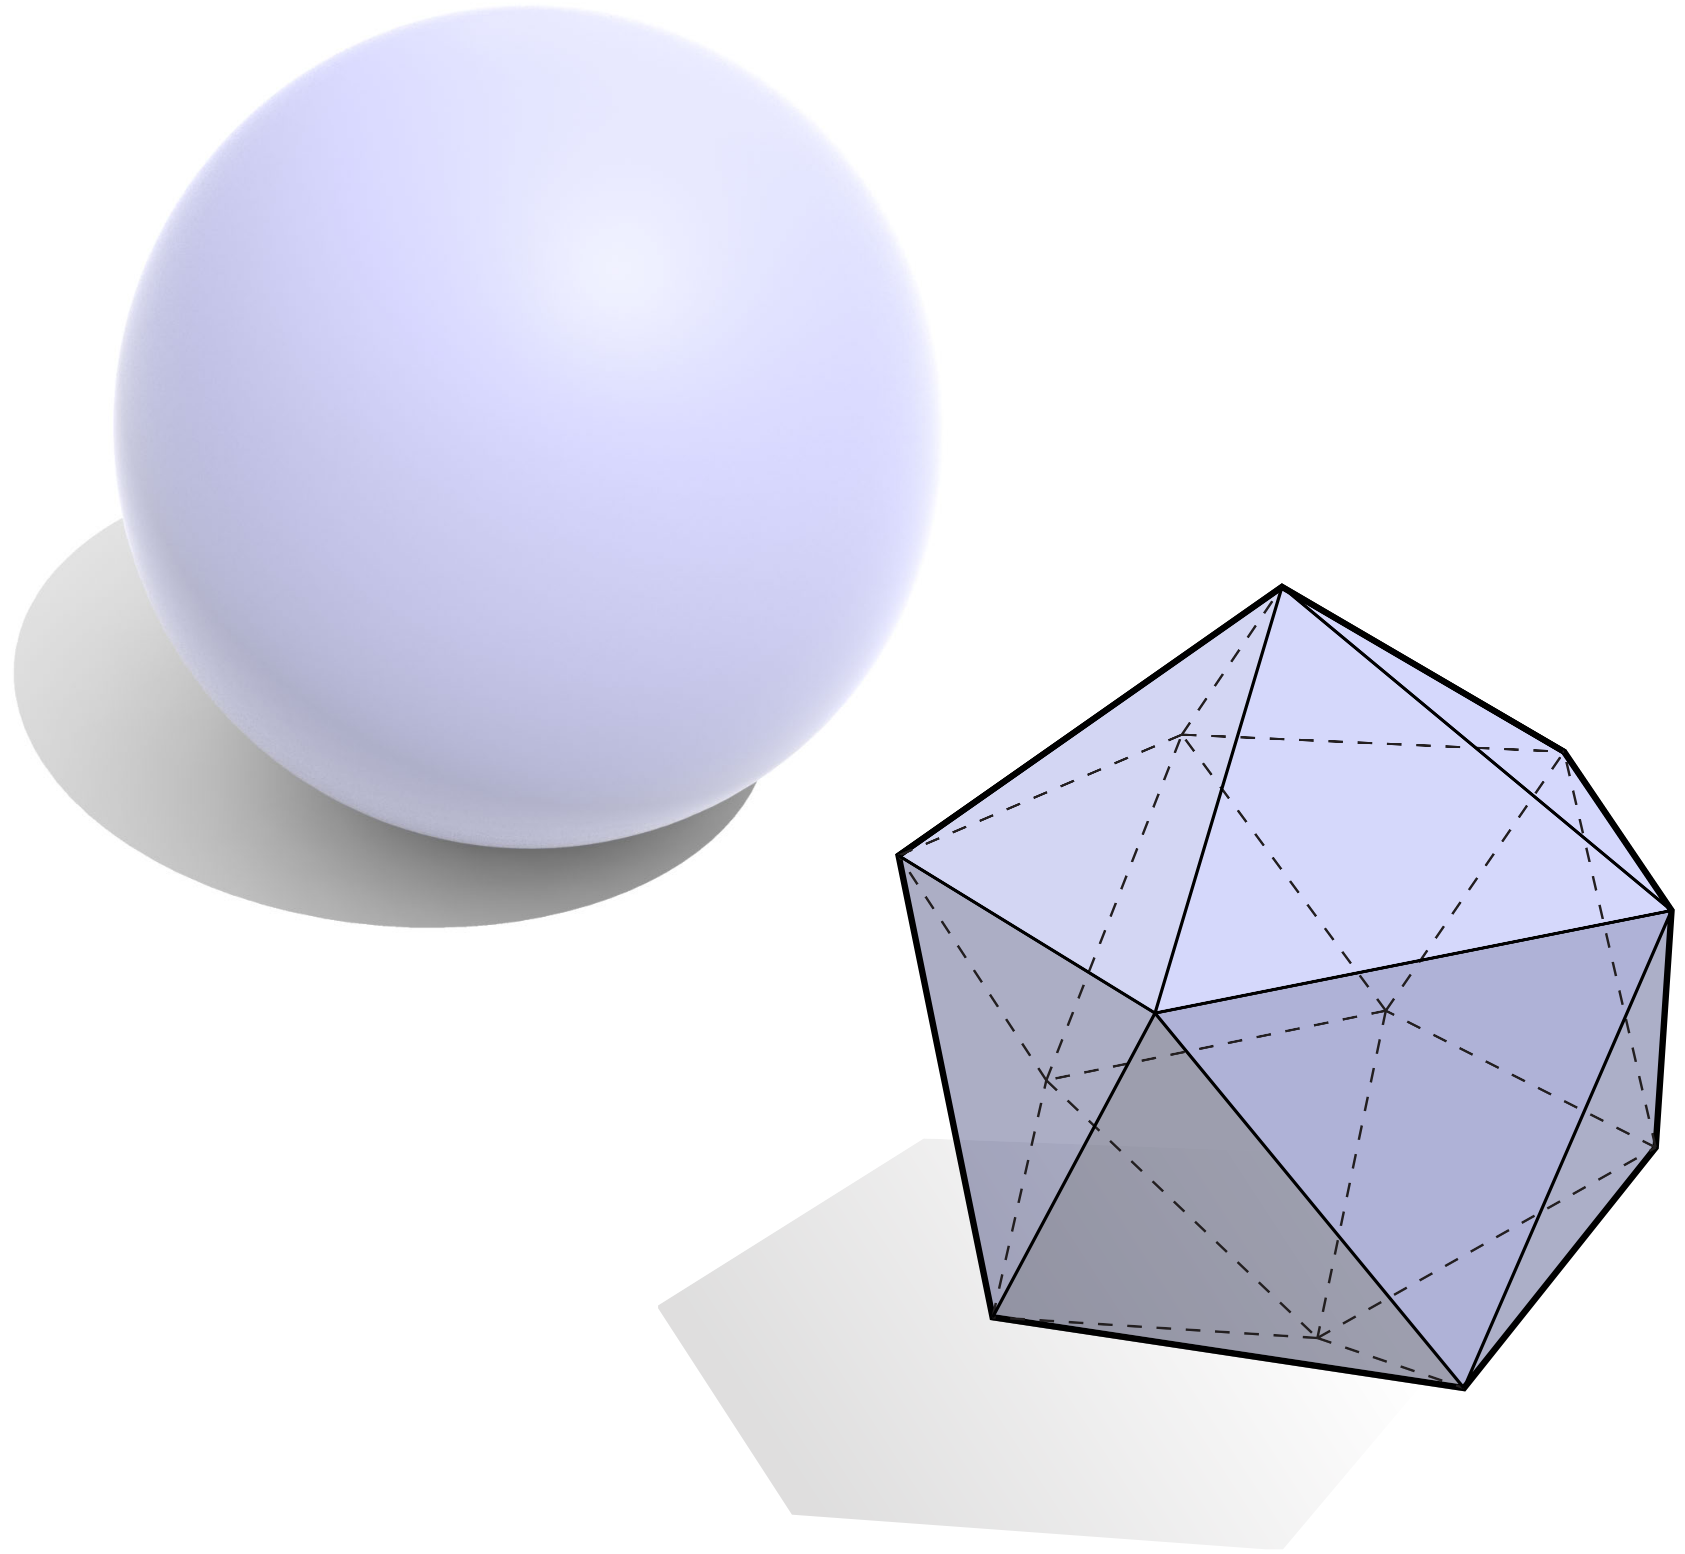

- **<u>Smooth setting:</u>**
  - express geometry via differential forms
  - will first need to think about *vector-valued* forms
- **<u>Discrete setting:</u>**
  - use explicit mesh as domain
  - express geometry via discrete differential forms
- **Payoff:** will become very easy to switch back & forth between smooth setting (*scribbling in a notebook*) and discrete setting (*running algorithms on real data!*)

Vector Valued Differential Forms

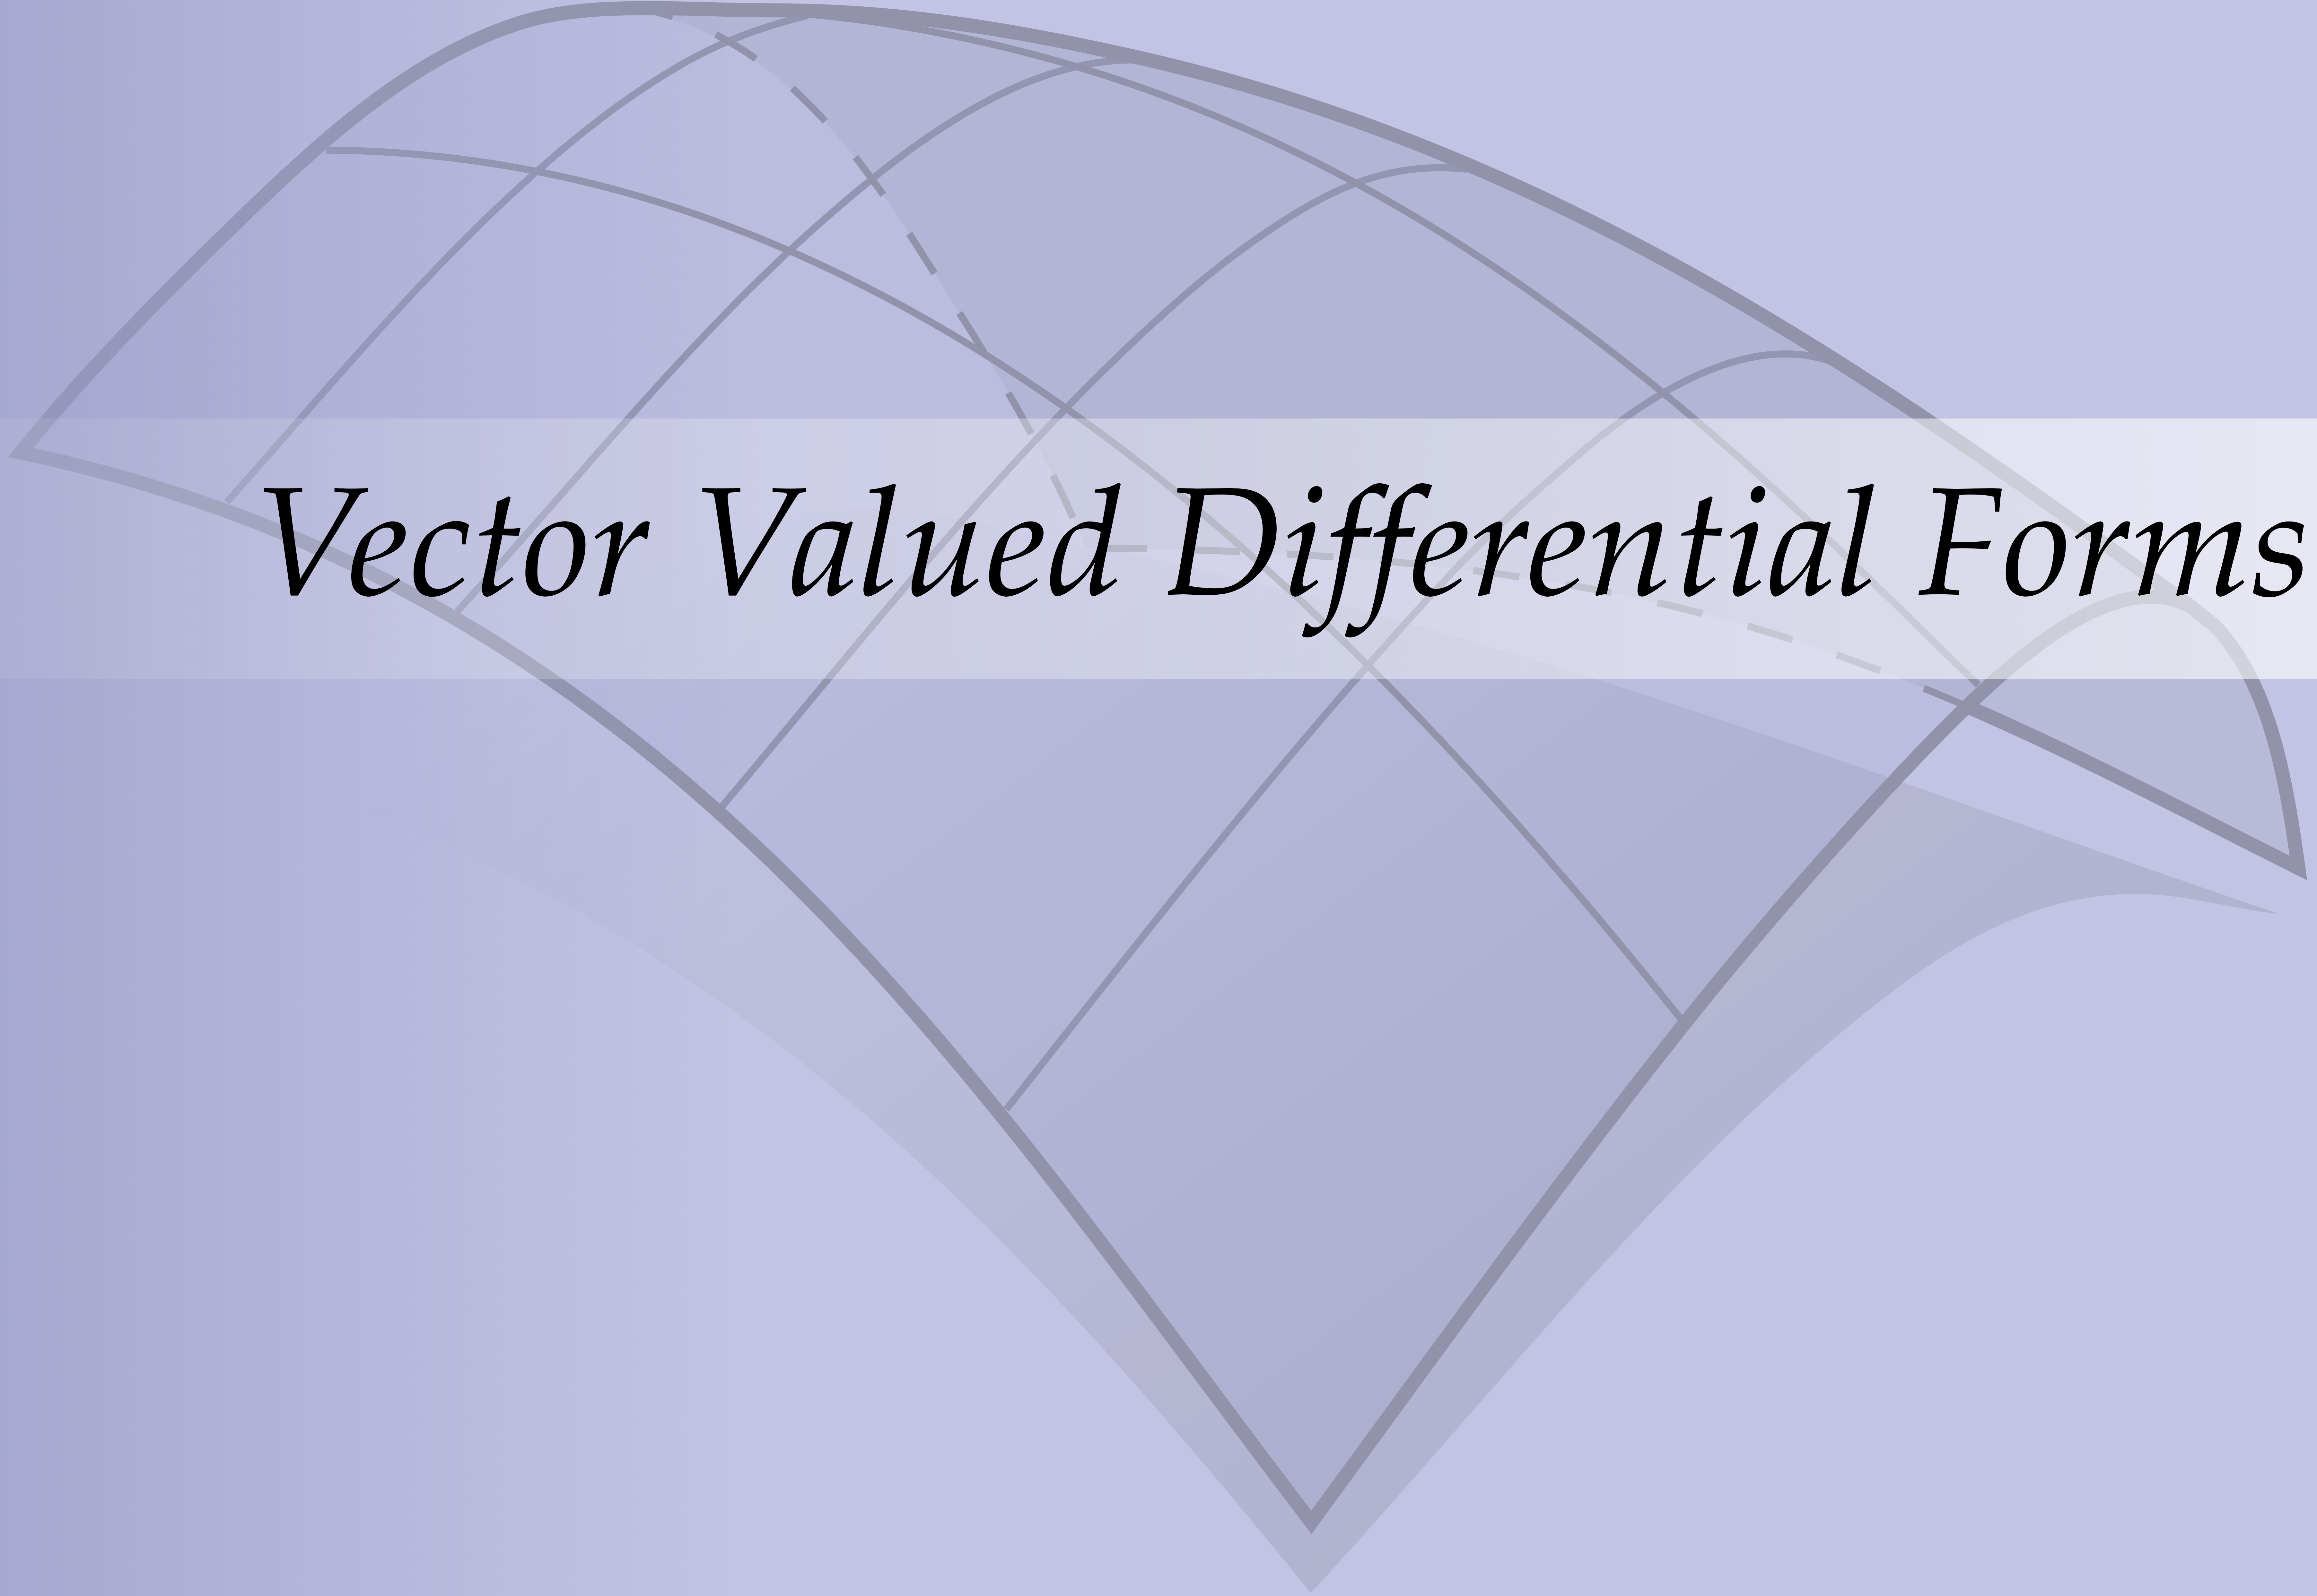

Vector Valued k-Forms

- So far, we've defined a *k*-form as a linear map from *k* vectors to a real number
- For working with curves and surfaces in R^n, it will be essential to generalize this definition to *vector-valued k-forms*.
- In particular, a **vector-valued *k*-form** is a multi-linear map from *k* vectors in a vector space *V* to some other vector space *U* (not necessarily *U*=*V*)
  - So far, for instance, all of our forms have been *R*-valued *k*-forms on R^n *(V=*R^n*,U=R)*
  - A R^3-valued 2-form on R^2 would instead be a multilinear, fully-antisymmetric map from a pair of vectors *u*,*v* in R^2 to a single vector in R^3:

$$\alpha : \mathbb{R}^2 \times \mathbb{R}^2 \to \mathbb{R}^3 \qquad \alpha(u,v) = -\alpha(v,u)$$

$$\alpha(au + bv, w) = a\alpha(u,w) + b\alpha(v,w), \quad \forall u,v,w \in \mathbb{R}^2, a,b \in \mathbb{R}$$

Q: What kind of object is a R^2-valued 0-form on R^2?

Vector-Valued k-forms—Example

Consider for instance the following R^3-valued 1-form on R^2:

$$\alpha := \begin{bmatrix} 1 \\ 2 \\ 3 \end{bmatrix} e^1 + \begin{bmatrix} 4 \\ 5 \\ 6 \end{bmatrix} e^2$$

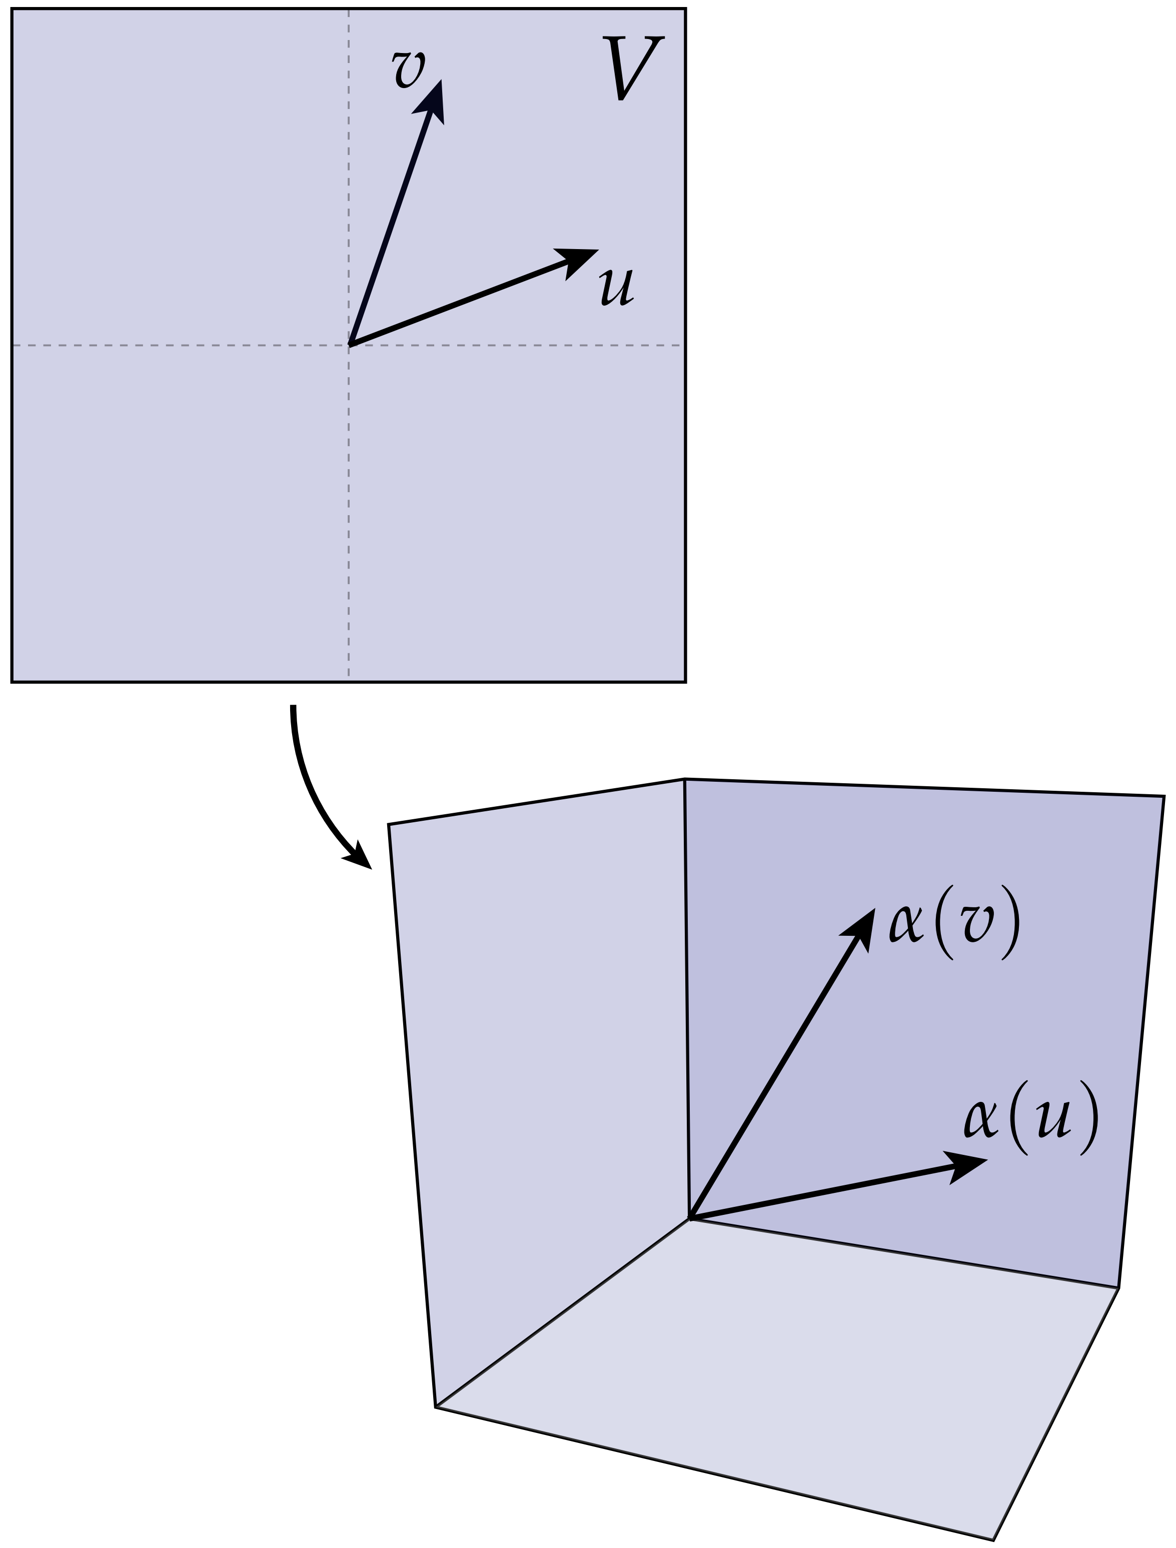

Q: What do we get if we evaluate this 1-form on the vector

$$u := e_1 - e_2$$

A: Evaluation is not much different from real-valued forms:

$$\alpha(u) = \begin{bmatrix} 1 \\ 2 \\ 3 \end{bmatrix} \underbrace{e^1(e_1 - e_2)}_{1} + \begin{bmatrix} 4 \\ 5 \\ 6 \end{bmatrix} \underbrace{e^2(e_1 - e_2)}_{-1} = \begin{bmatrix} 1 \\ 2 \\ 3 \end{bmatrix} - \begin{bmatrix} 4 \\ 5 \\ 6 \end{bmatrix} = \begin{bmatrix} -3 \\ -3 \\ -3 \end{bmatrix}$$

Key idea: coefficients just have a different type

Wedge Product of Vector-Valued k-Forms

- Most important change is how we evaluate wedge product for vector-valued forms.
- Consider for instance a pair of R^3-valued 1-forms:

$$\alpha, \beta : V \to \mathbb{R}^3$$

- To evaluate their wedge product on a pair of vectors *u*,*v* we would normally write:

$$(\alpha \wedge \beta)(u, v) = \alpha(u)\beta(v) - \alpha(v)\beta(u)$$

- If α and β were *real*-valued, then $\alpha(u)$, $\beta(v)$, $\alpha(v)$, $\beta(u)$, would just be *real numbers*, so we could just multiply the two pairs and take their difference.
- But what does it mean to take the "product" of two vectors from R^3?
- Many possibilities (*e.g.*, dot product), but if we want result to be an R^3-valued 2-form, the product we choose must produce another 3-vector!

Wedge Product of R^3-Valued k-Forms

- Most common case for our study of surfaces:
  - k-forms are R^3-valued
  - use **cross product** to multiply 3-vectors

$$\alpha, \beta : V \to \mathbb{R}^3$$

$$\alpha \wedge \beta : V \times V \to \mathbb{R}^3$$

$$\boxed{(\alpha \wedge \beta)(u,v) := \alpha(u) \times \beta(v) - \alpha(v) \times \beta(u)}$$

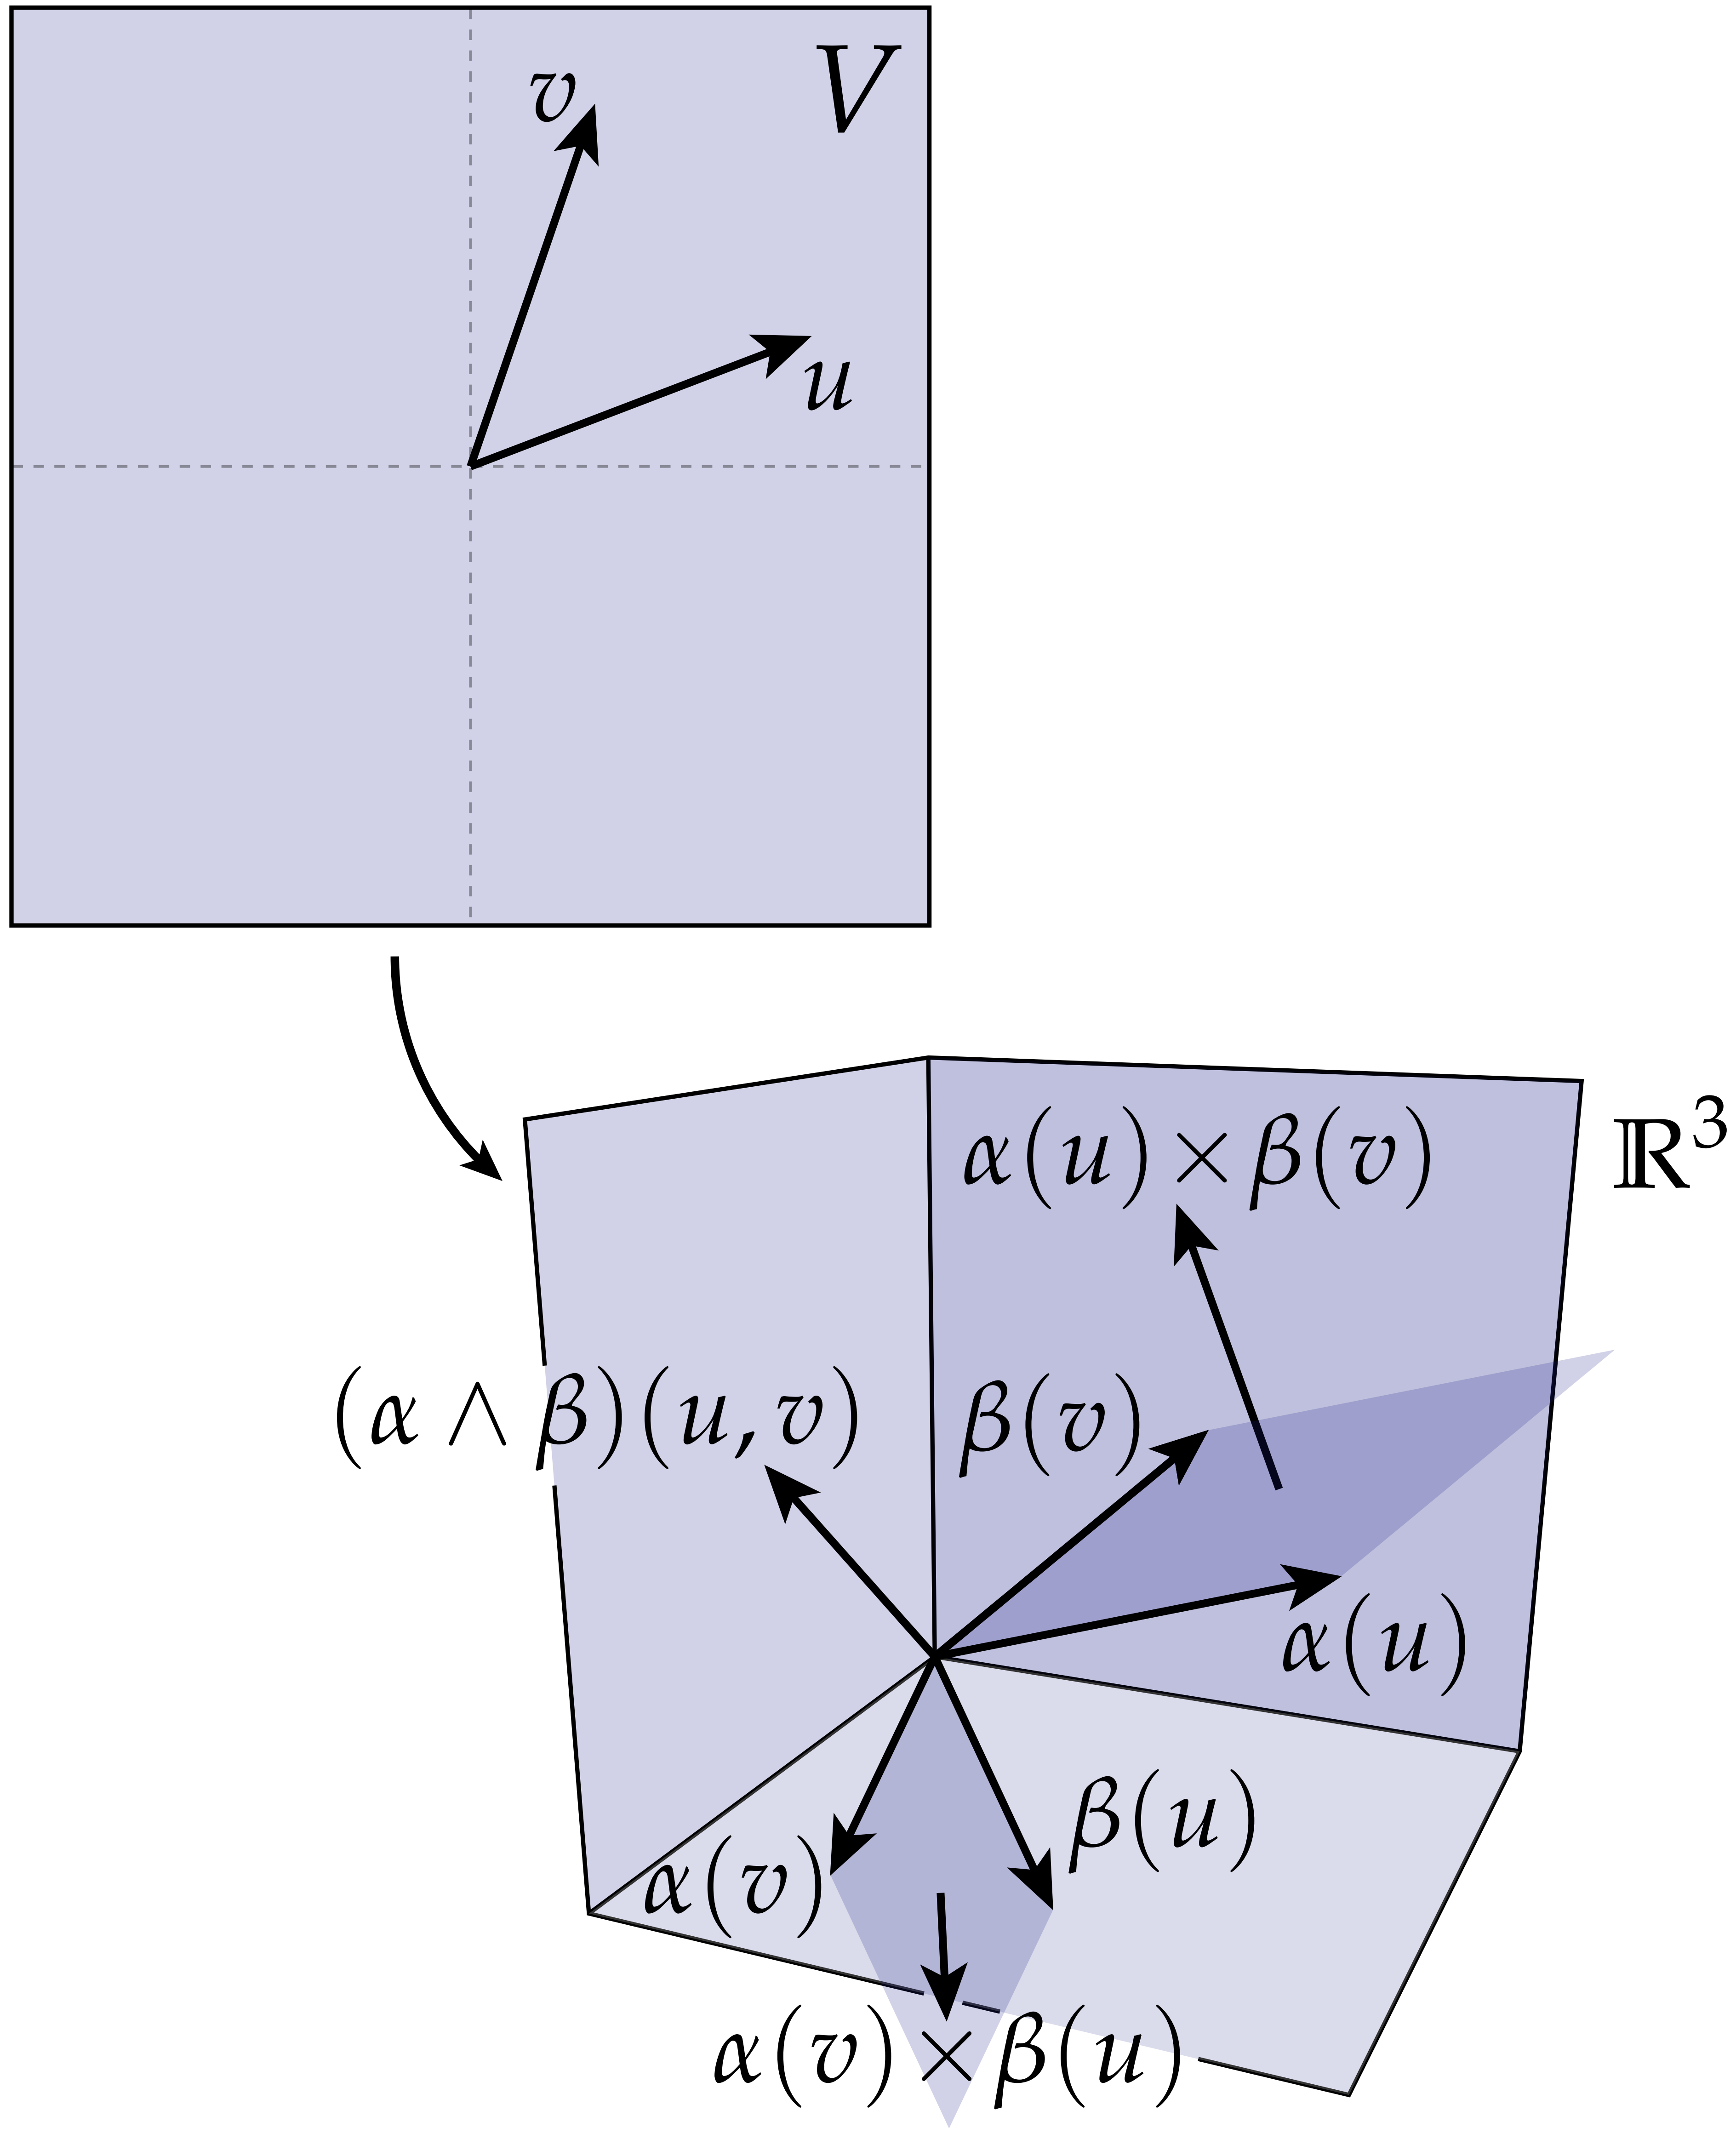

R^3-valued 1-forms: Antisymmetry & Symmetry

With real-valued forms, we observed antisymmetry in both the wedge product of 1-forms as well as the application of the 2-form to a pair of vectors, *i.e.*,

$$(\alpha \wedge \beta)(u, v) = -(\alpha \wedge \beta)(v, u)$$
$$(\beta \wedge \alpha)(u, v) = -(\alpha \wedge \beta)(u, v)$$

What happens w/ R^3-valued 1-forms? Since cross product is antisymmetric, we get

$$\begin{aligned}(\alpha \wedge \beta)(v, u) &= \alpha(v) \times \beta(u) - \alpha(u) \times \beta(v) \\ &= -(\alpha(u) \times \beta(v) - \alpha(v) \times \beta(u))\end{aligned}$$

$$\Rightarrow \boxed{(\alpha \wedge \beta)(u, v) = -(\alpha \wedge \beta)(v, u)}$$

(no change)

$$\begin{aligned}(\beta \wedge \alpha)(u, v) &= \beta(u) \times \alpha(v) - \beta(v) \times \alpha(u) \\ &= \alpha(u) \times \beta(v) - \alpha(v) \times \beta(u) \\ &= (\alpha \wedge \beta)(u, v)\end{aligned}$$

$$\Rightarrow \boxed{\alpha \wedge \beta = \beta \wedge \alpha}$$

(big change!)

R^3-valued 1-forms: Self-Wedge

Likewise, we saw that wedging a real-valued 1-form with itself yields zero:

$$\alpha \wedge \alpha = 0$$

Q: What was the *geometric* interpretation?

A: Parallelogram spanned by two copies of the same vector has zero area!

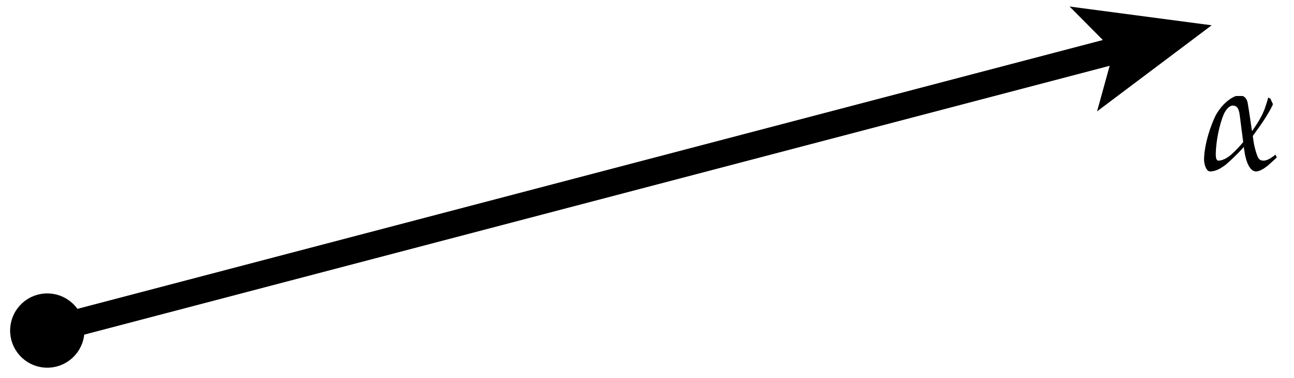

…But, no longer true with (R^3, ×)-valued 1-forms:

$$(\alpha \wedge \alpha)(u, v) = \alpha(u) \times \alpha(v) - \alpha(v) \times \alpha(u) = 2\alpha(u) \times \alpha(v) \neq 0$$

Geometric meaning will become clearer as we work with surfaces.

Vector-Valued Differential k-Forms

- Just as we distinguished between a *k-form* (value at a single point) and a *differential k-form* (value at ever point in space), we will also say that a *vector-valued differential k-form* is a vector-valued k-form at each point of space.

- Just like any differential form, a vector-valued differential *k*-form gets evaluated on k vector fields $X_1, \ldots, X_k$.

- **<u>Example:</u>** an R^3-valued differential 1-form on R^2 (with coordinates u,v):

$$\alpha := \begin{bmatrix} 0 \\ 1 \\ 0 \end{bmatrix} du + \begin{bmatrix} 0 \\ 0 \\ 1 \end{bmatrix} dv$$

Q: What does this 1-form do to any given vector field X on the plane?

A: It simply "copies" it to the *yz*-plane in 3D.

Exterior Derivative on Vector-Valued Forms

Unlike the wedge product, not much changes with the exterior derivative. For instance, if we have an R^n-valued k-form we can simply imagine we have n real-valued k-forms and differentiate as usual.

Example.

Consider an $\mathbb{R}^2$-valued differential 0-form $\phi_{(x,y)} := \begin{bmatrix} x^2 \\ xy \end{bmatrix}$

Then $d\phi = \frac{\partial \phi}{\partial x} dx + \frac{\partial \phi}{\partial y} dy = \begin{bmatrix} 2x \\ y \end{bmatrix} dx + \begin{bmatrix} 0 \\ x \end{bmatrix} dy$

Example.

Consider an $\mathbb{R}^2$-valued differential 1-form $\alpha_{(x,y)} := \begin{bmatrix} x^2 \\ xy \end{bmatrix} dx + \begin{bmatrix} xy \\ y^2 \end{bmatrix} dy$

Then $d\alpha = \left(\begin{bmatrix} 2x \\ y \end{bmatrix} dx + \begin{bmatrix} 0 \\ x \end{bmatrix} dy \right) \wedge dx + \left(\begin{bmatrix} y \\ 0 \end{bmatrix} dx + \begin{bmatrix} x \\ 2y \end{bmatrix} dy \right) \wedge dy = \begin{bmatrix} y \\ -x \end{bmatrix} dx \wedge dy$

Planar Curves

Parameterized Plane Curve

- A **parameterized plane curve** is a map* taking each point in an interval $[0,L]$ of the real line to some point in the plane $\mathbb{R}^2$:

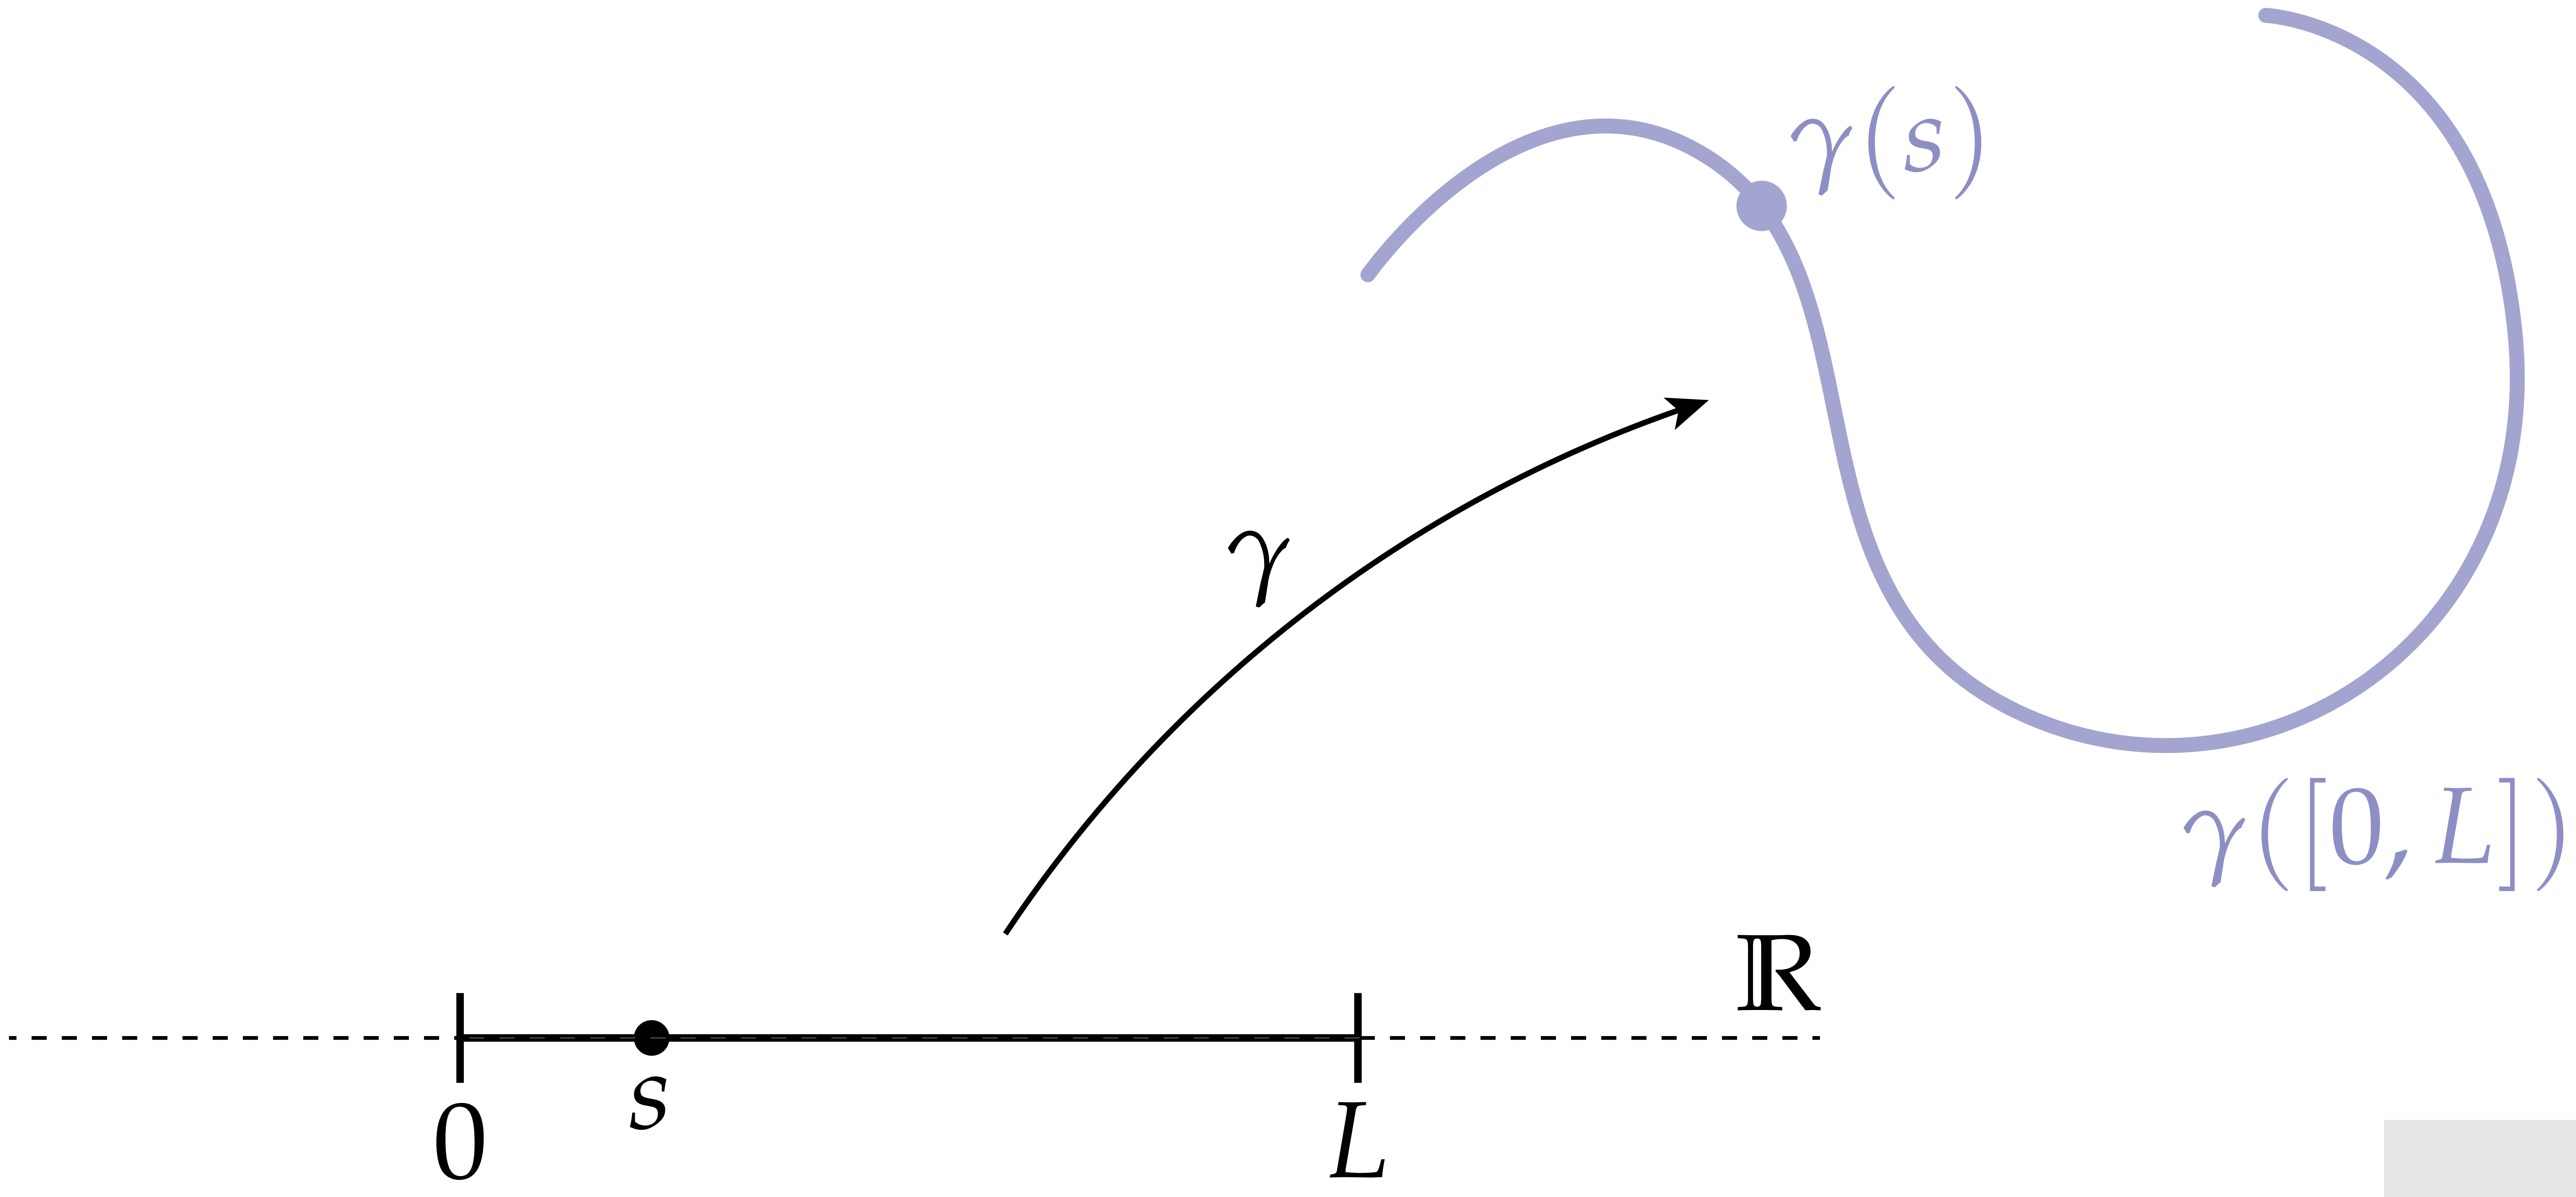

*Continuous, differentiable, smooth…

$$\gamma : [0,L] \to \mathbb{R}^2$$

Curves in the Plane—Example

- As an example, we can express a circle as a parameterized curve γ:

$$\gamma : [0, 2\pi) \to \mathbb{R}^2;\ s \mapsto (\cos(s), \sin(s))$$

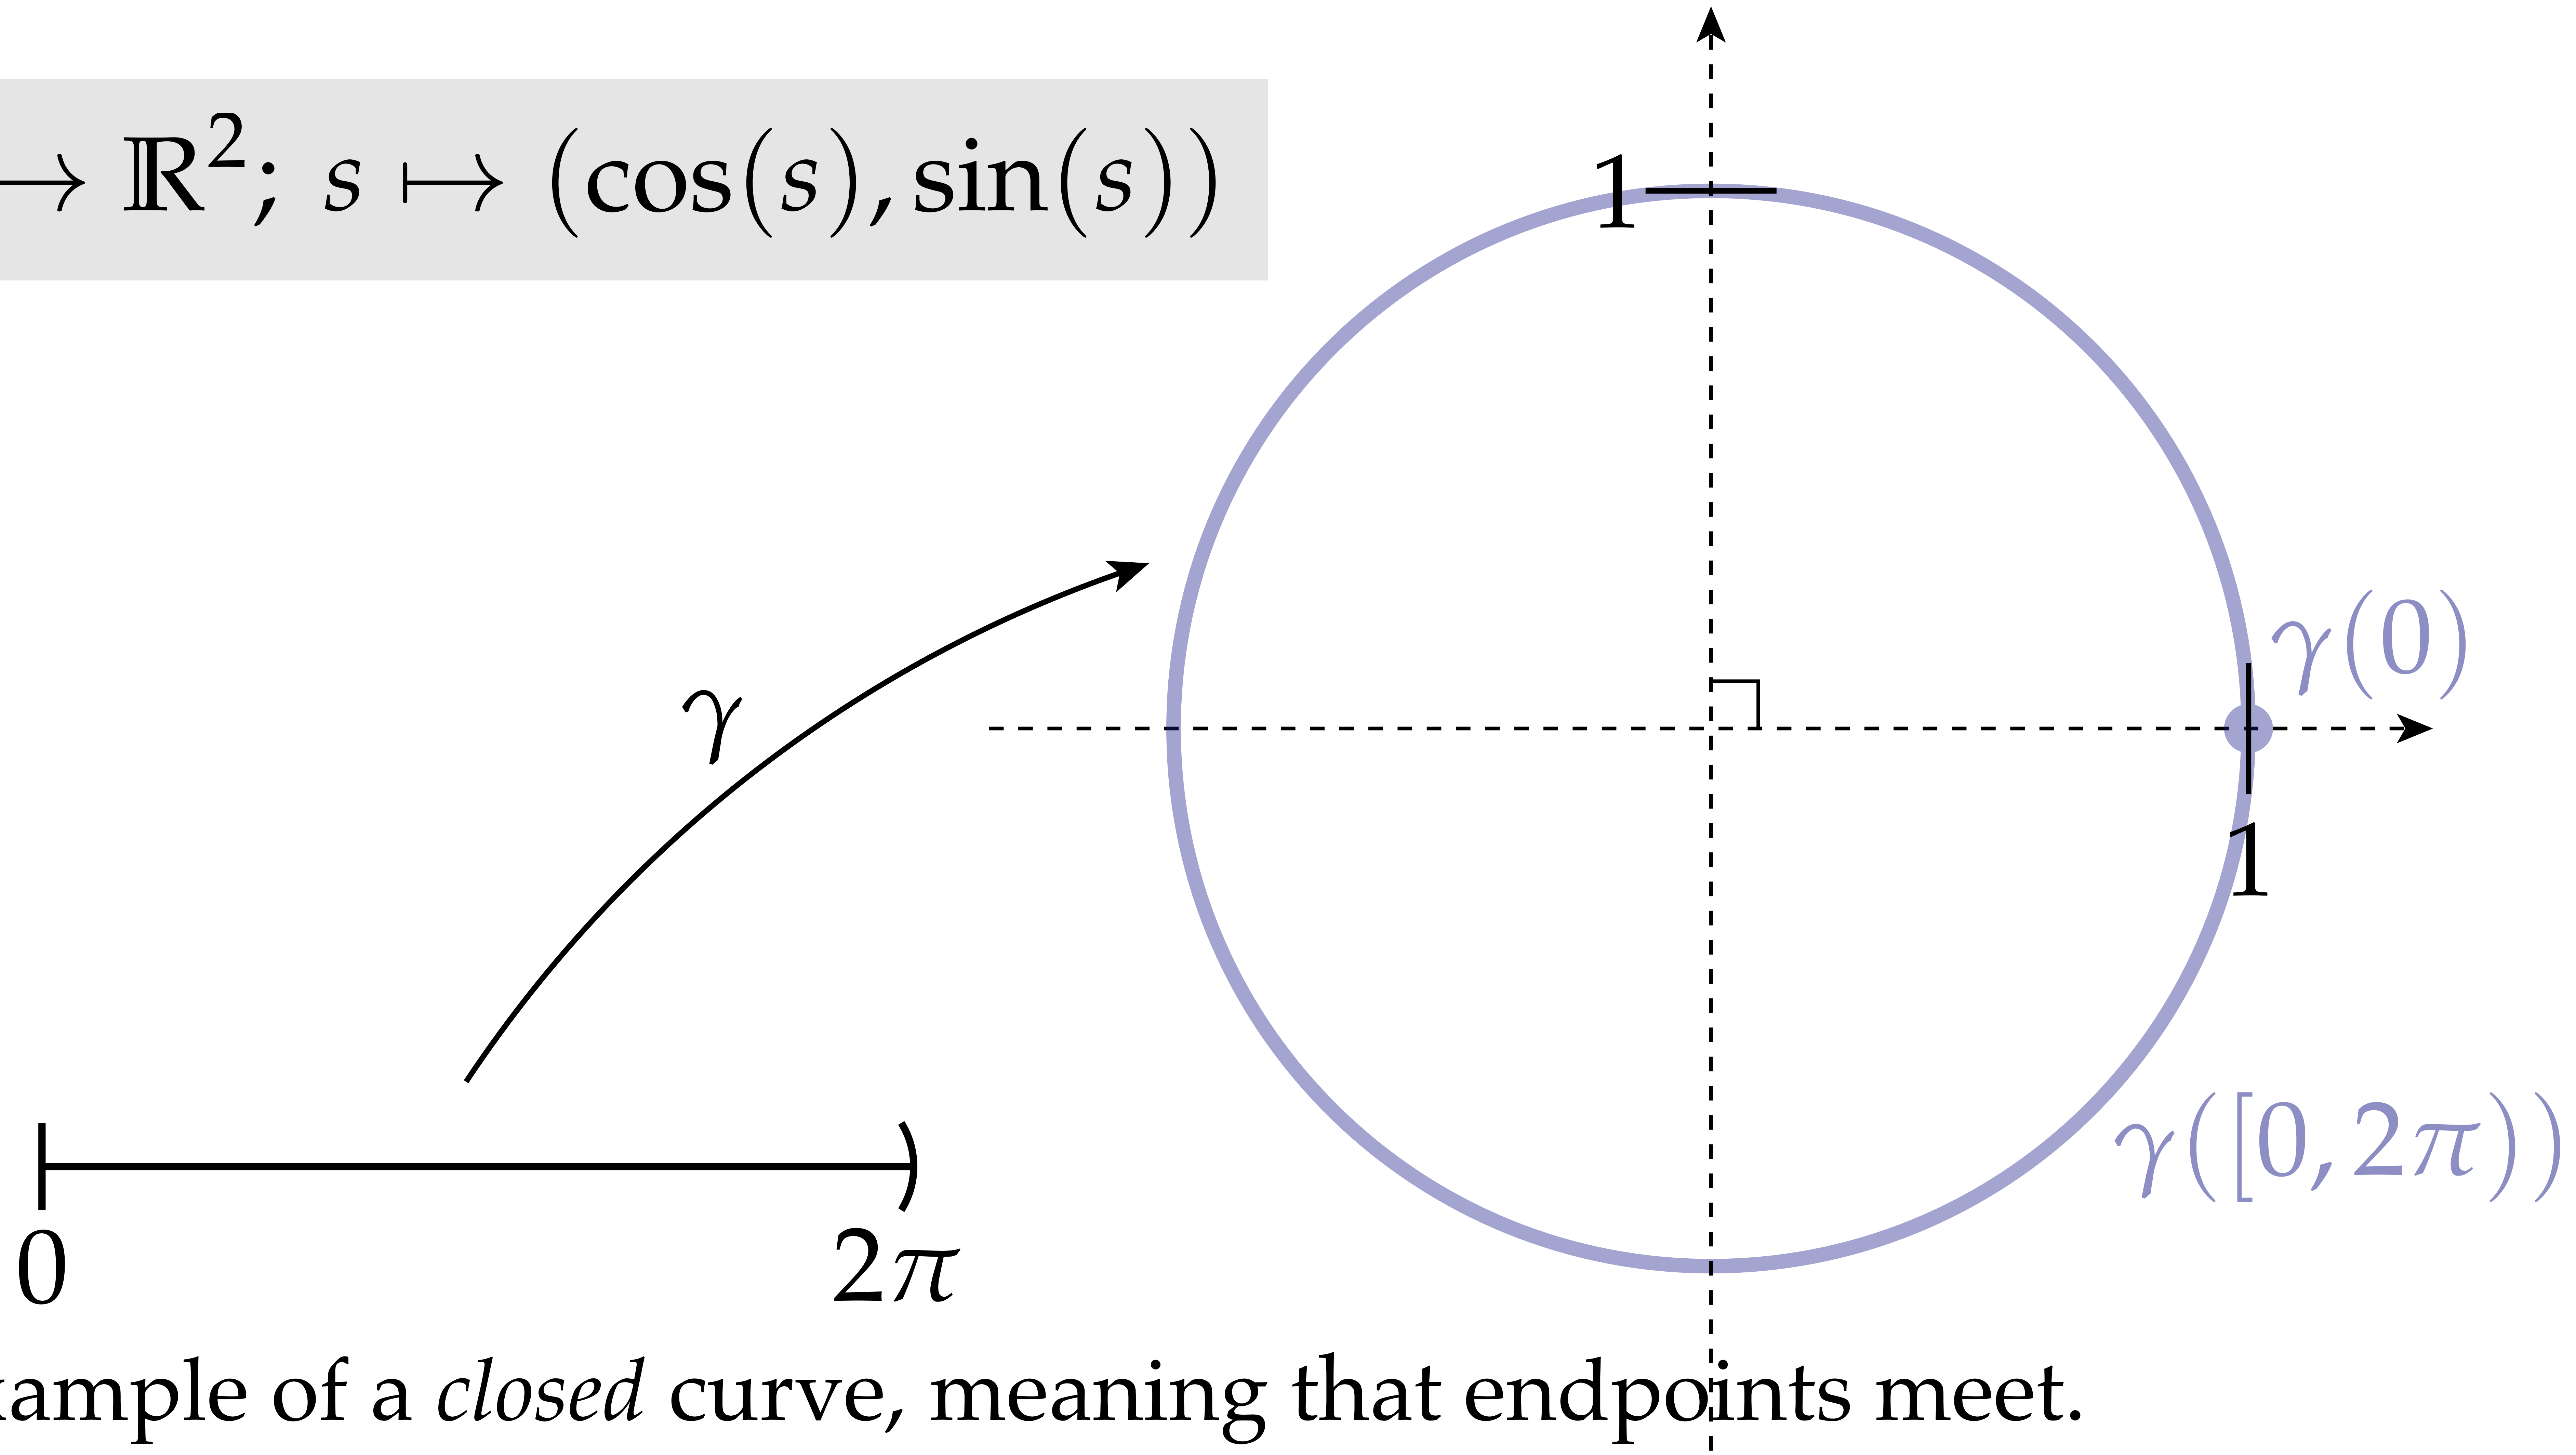

The circle is an example of a *closed* curve, meaning that endpoints meet.

Differential of a Curve

- If we think of a parameterized curve as an R^2-valued 0-form on an interval of the real line, then the *differential* (or exterior derivative) says how vectors on the domain get mapped into the plane:

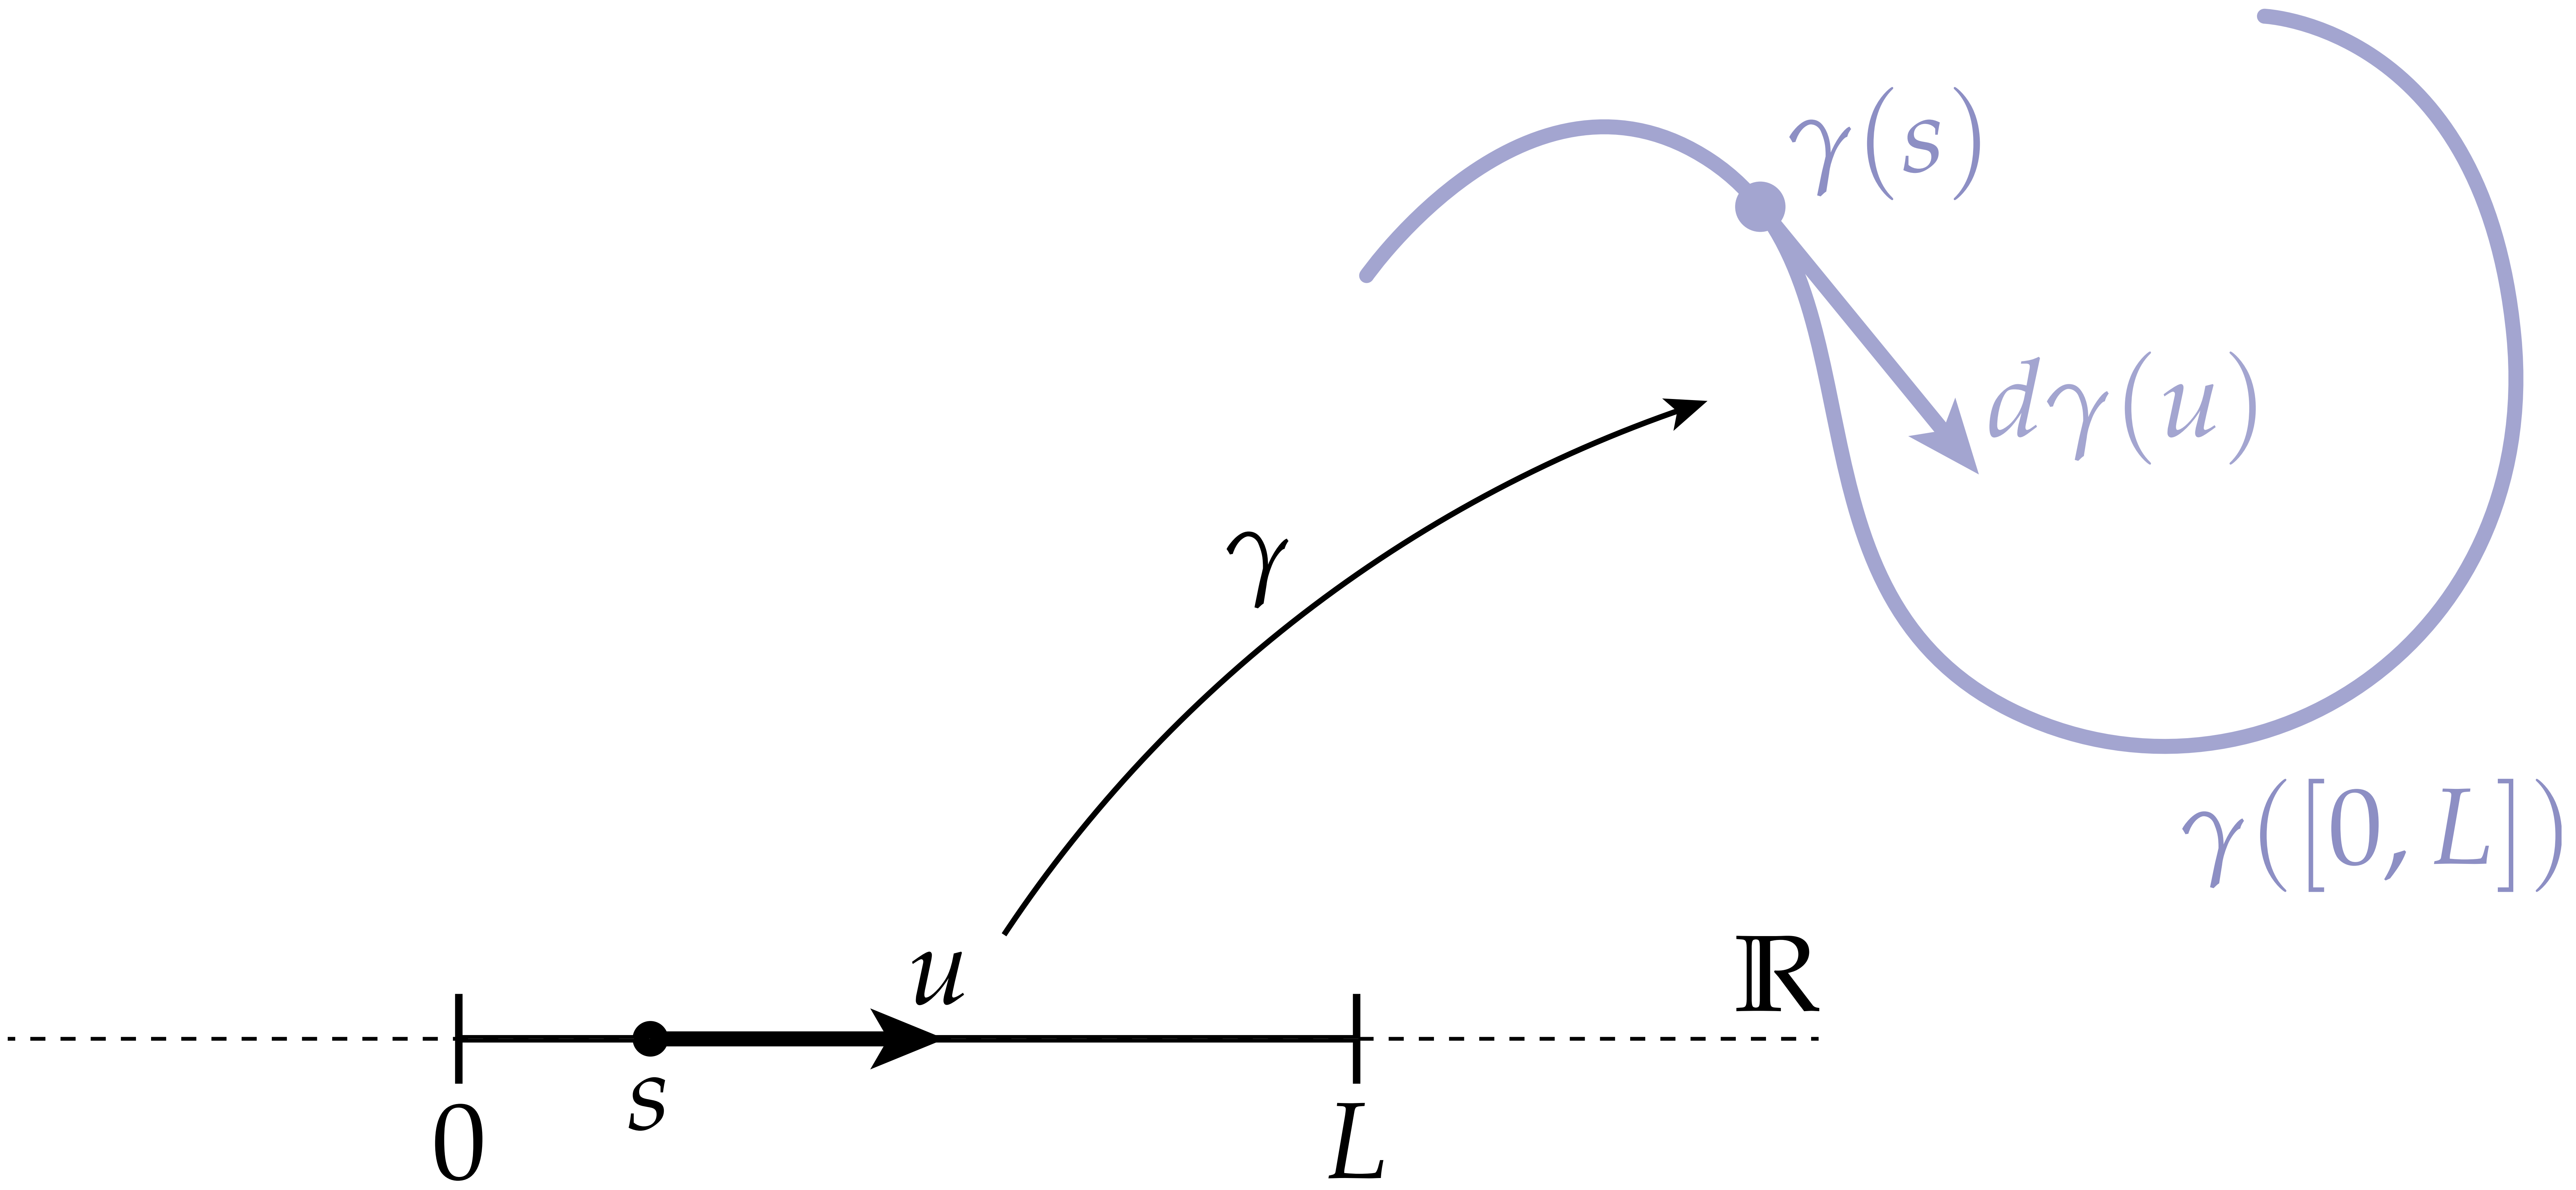

Tangent of a Curve

- Informally, a vector is *tangent* to a curve if it "just barely grazes" the curve.

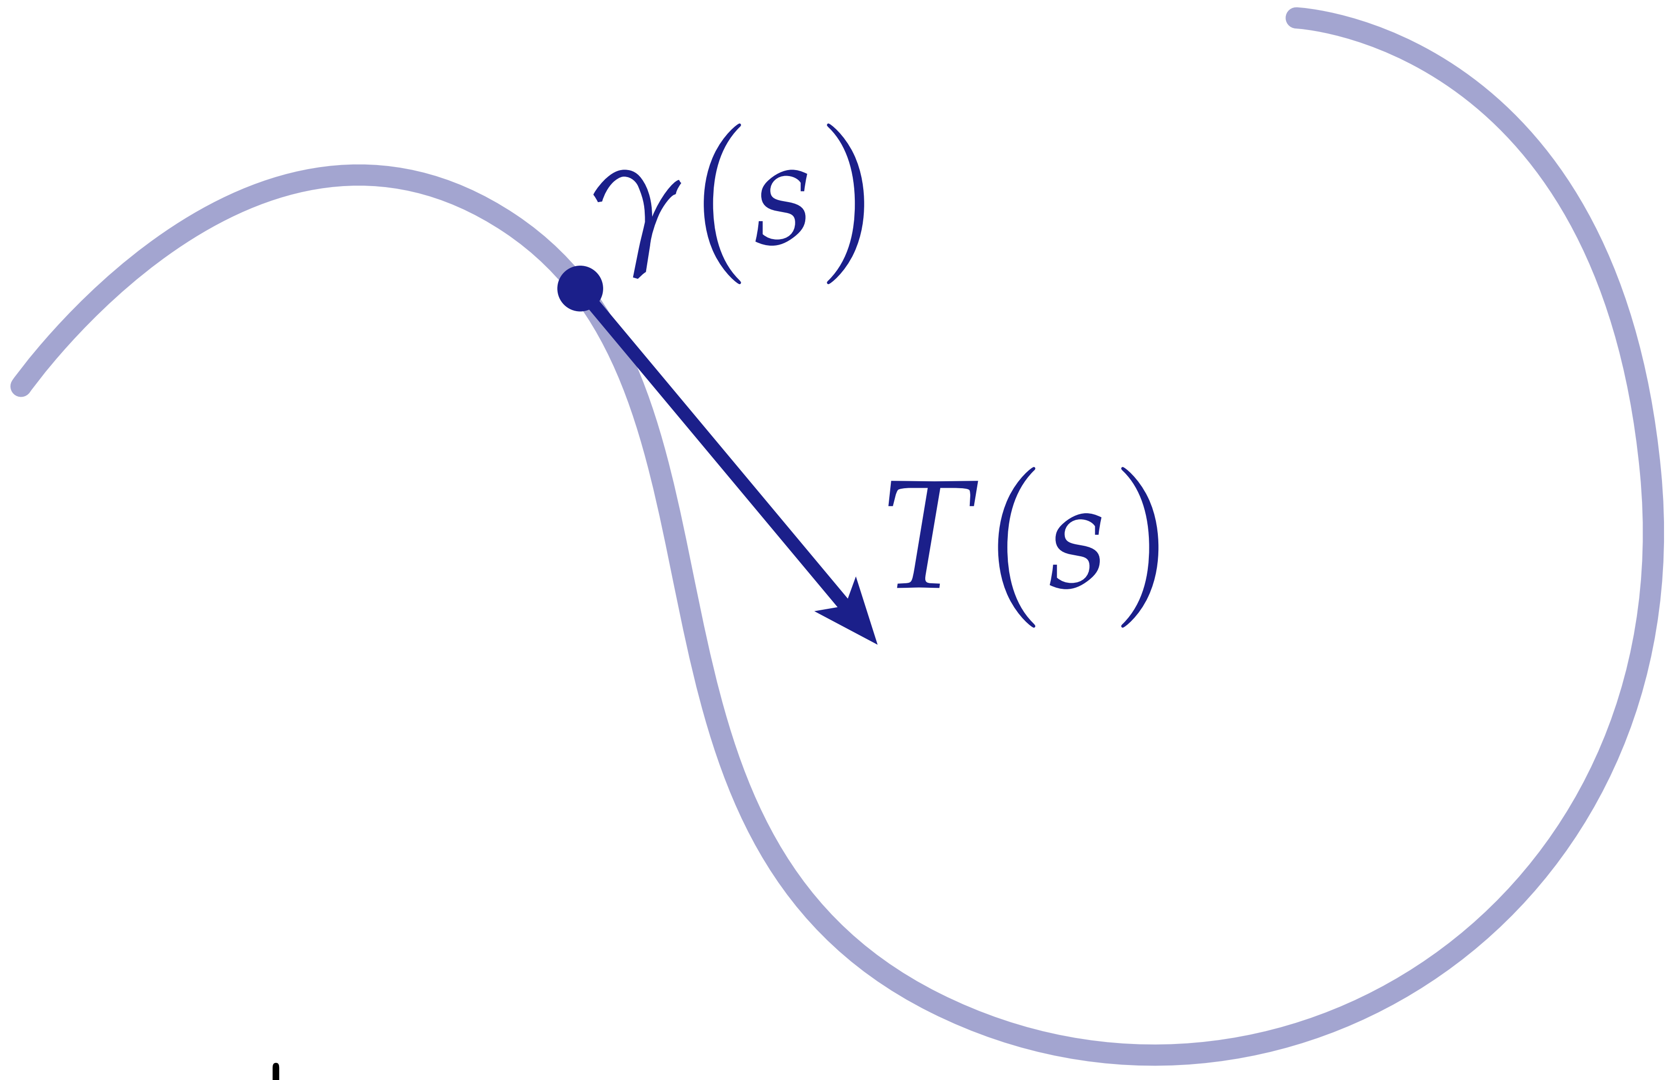

- More formally, the **unit tangent** (or just **tangent**) of a regular curve is the map obtained by normalizing its first derivative:

$$T(s) := \tfrac{d}{ds}\gamma(s) \Big/ \left|\tfrac{d}{ds}\gamma(s)\right| = d\gamma(\tfrac{d}{ds}) \Big/ \left|d\gamma(\tfrac{d}{ds})\right|$$

- If the derivative already has unit length, then we say the curve is **arc-length parameterized** and can write the tangent as just

$$T(s) := \tfrac{d}{ds}\gamma(s) = d\gamma(\tfrac{d}{ds})$$

Tangent of a Curve—Example

- Let's compute the unit tangent of a circle:

$$\gamma : [0, 2\pi) \to \mathbb{R}^2;\ s \mapsto (\cos(s), \sin(s))$$

$$d\gamma = (-\sin(s), \cos(s))ds$$

$$d\gamma(\tfrac{\partial}{\partial s}) = (-\sin(s), \cos(s))$$

$$\cos^2(s) + \sin^2(s) = 1$$

$$\Rightarrow T = d\gamma(\tfrac{d}{ds})$$

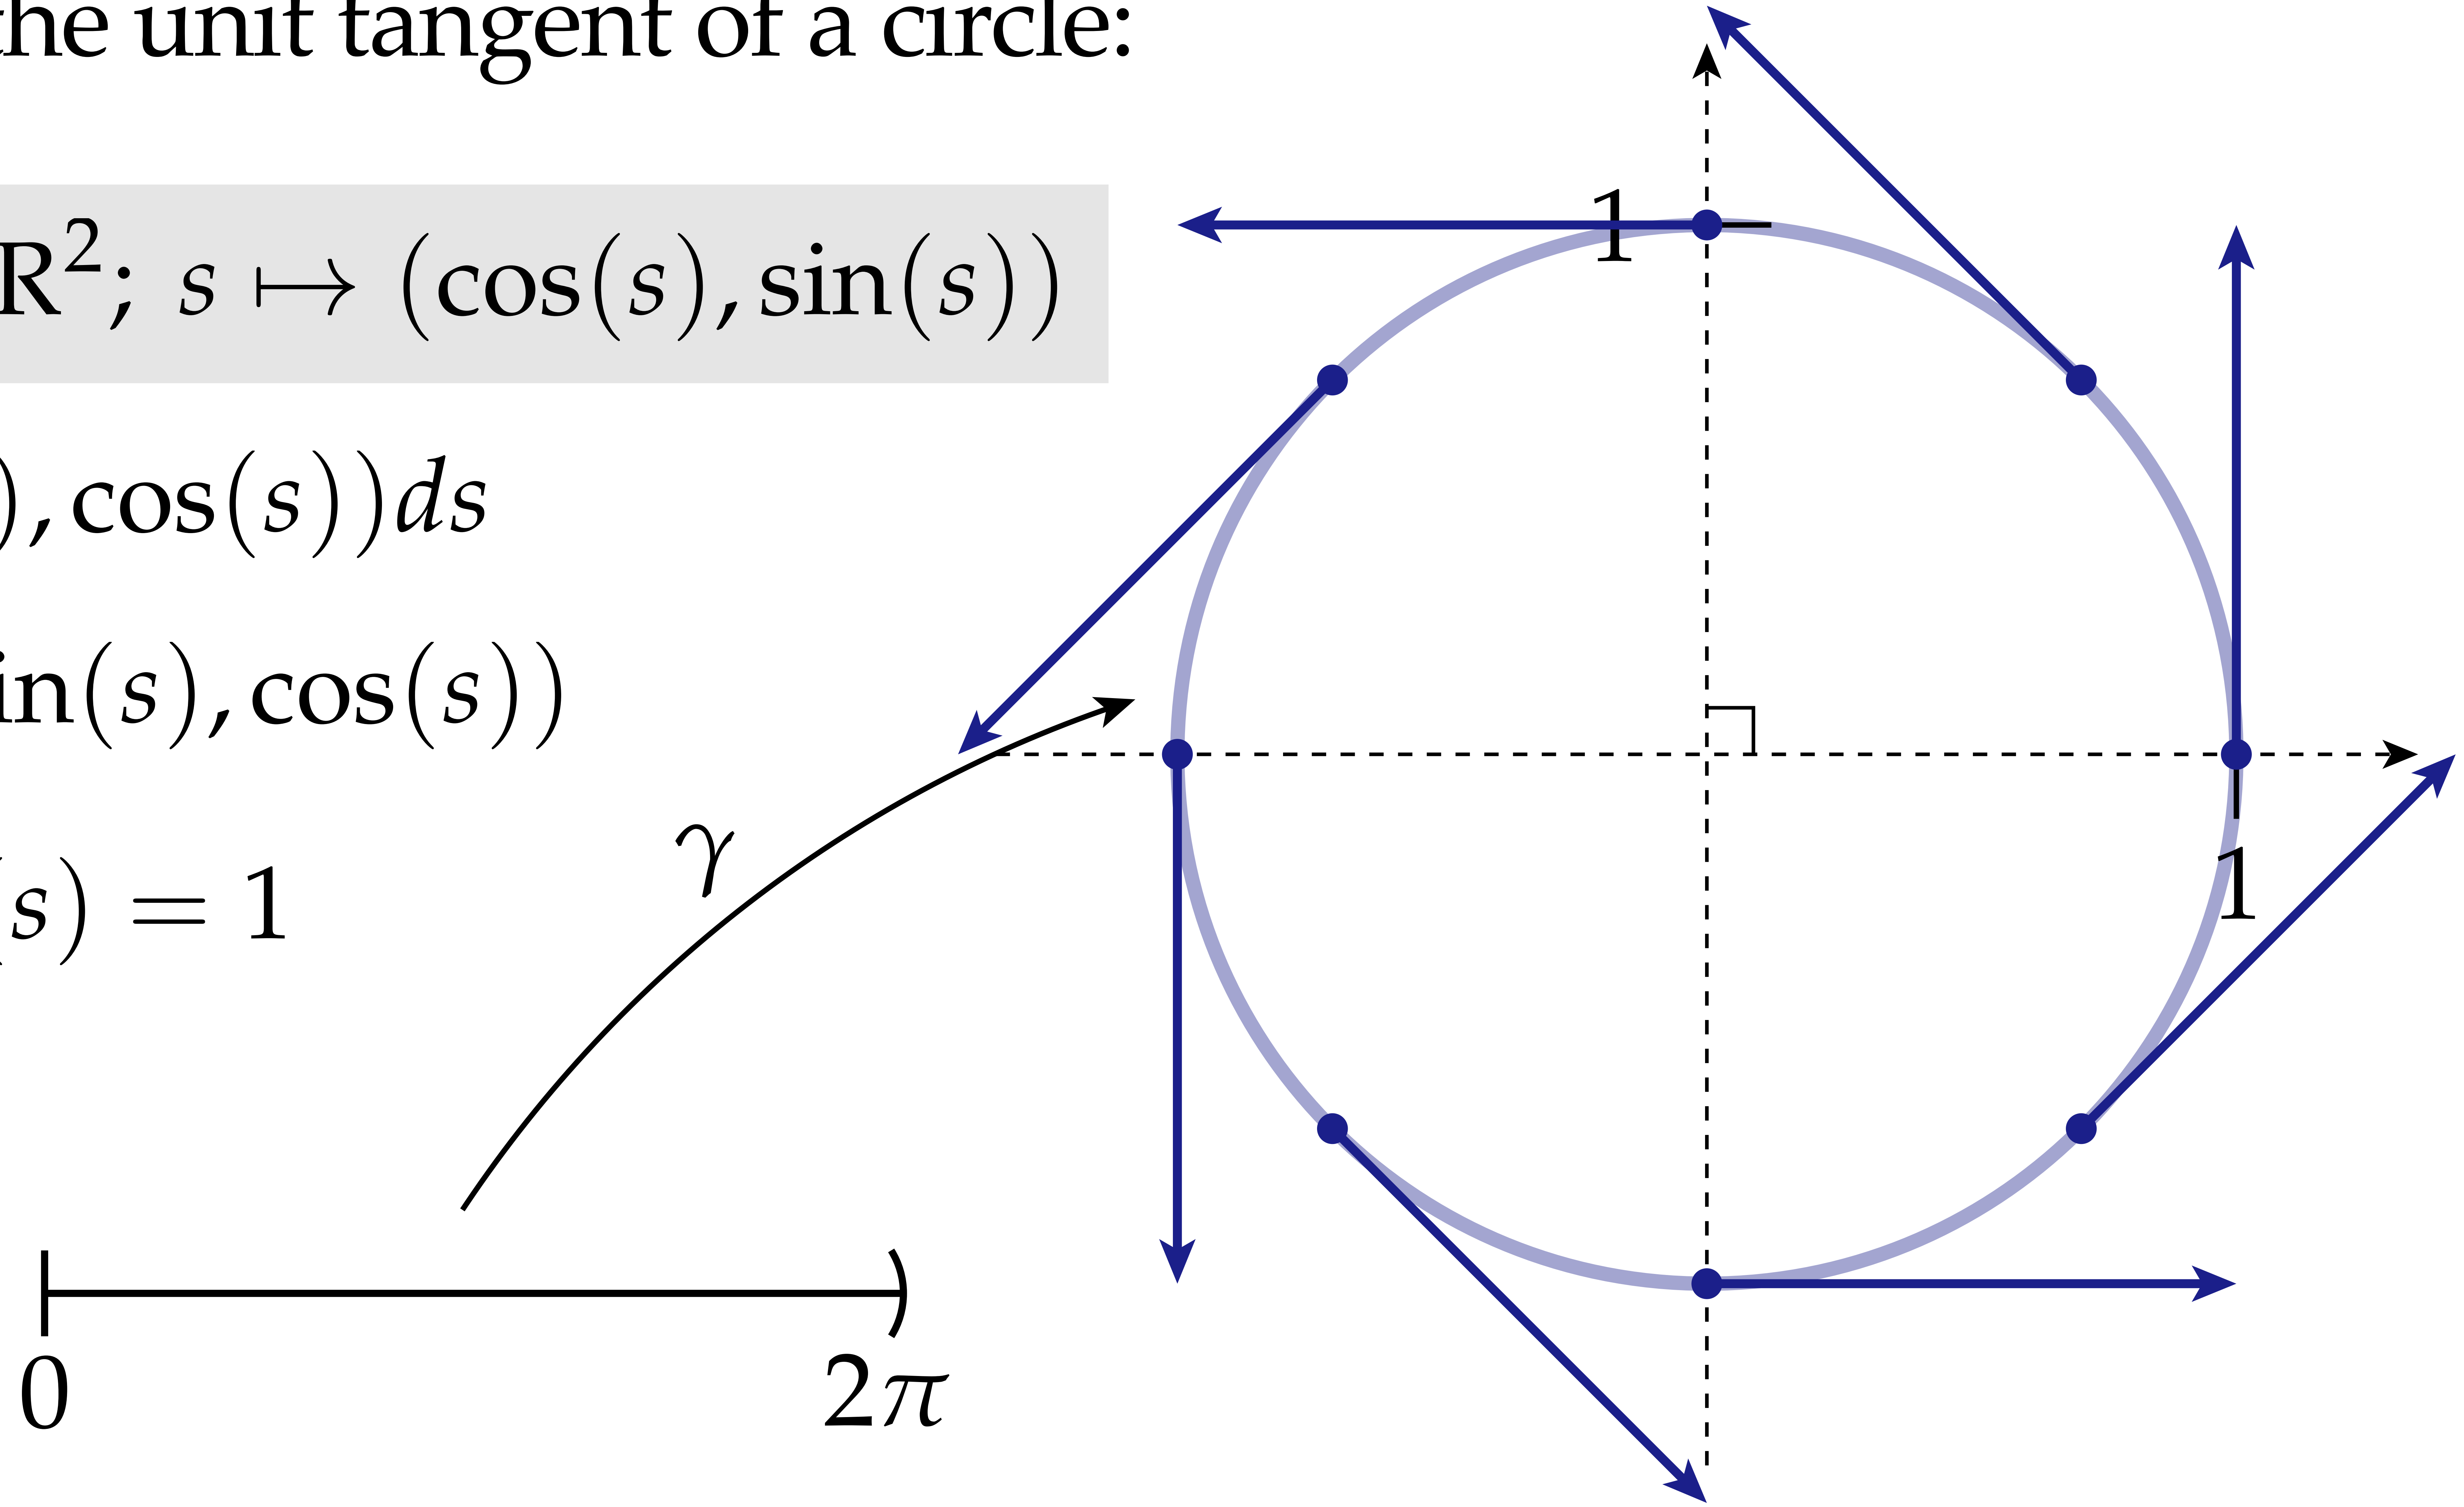

Reparameterization of a Curve

- We can *reparameterize* a curve $\gamma : \mathbb{R} \supset I \to \mathbb{R}$ by composing it with a bijection $\eta : I \to I$ to obtain a new parameterized curve

$$\tilde{\gamma}(s) := \gamma(\eta(s))$$

- The *image* of the new curve is the same, even though the map itself changes. For example:

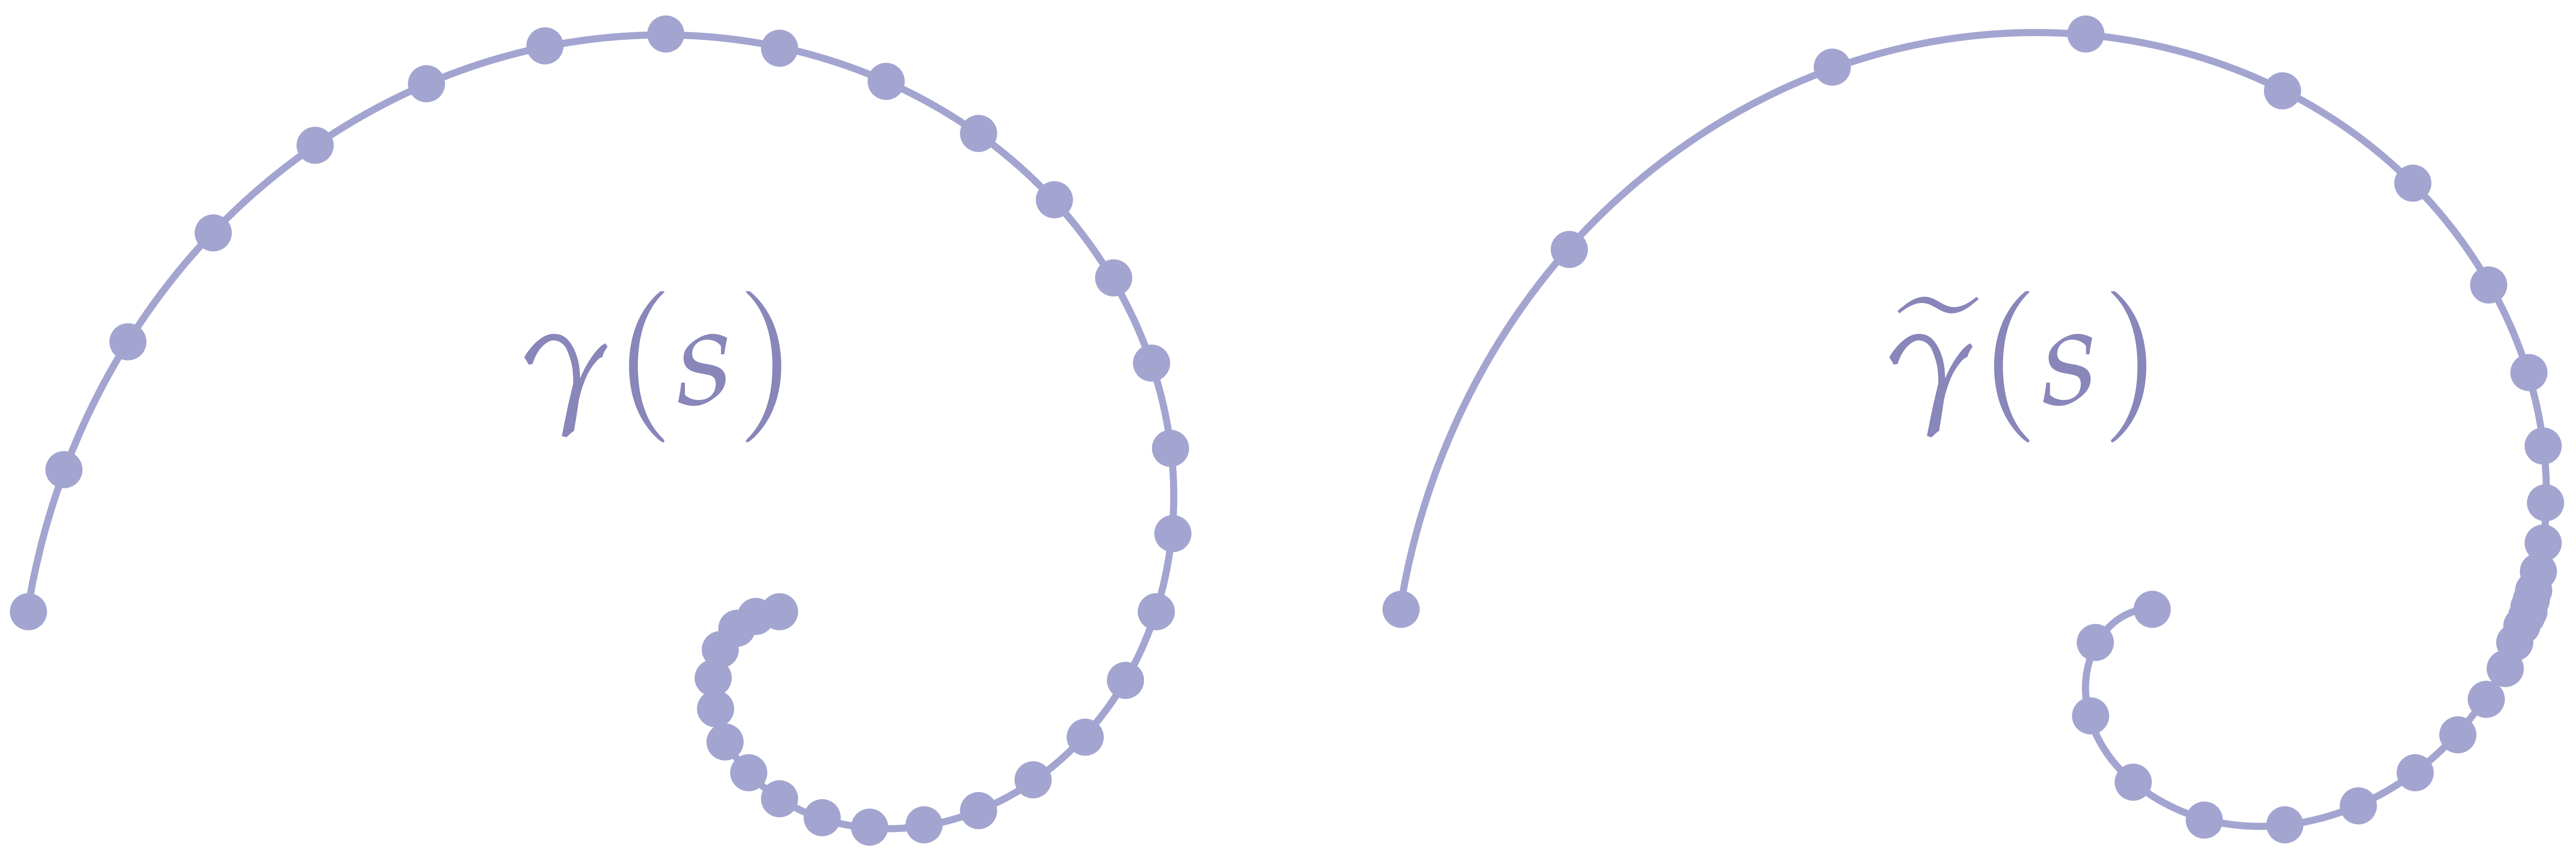

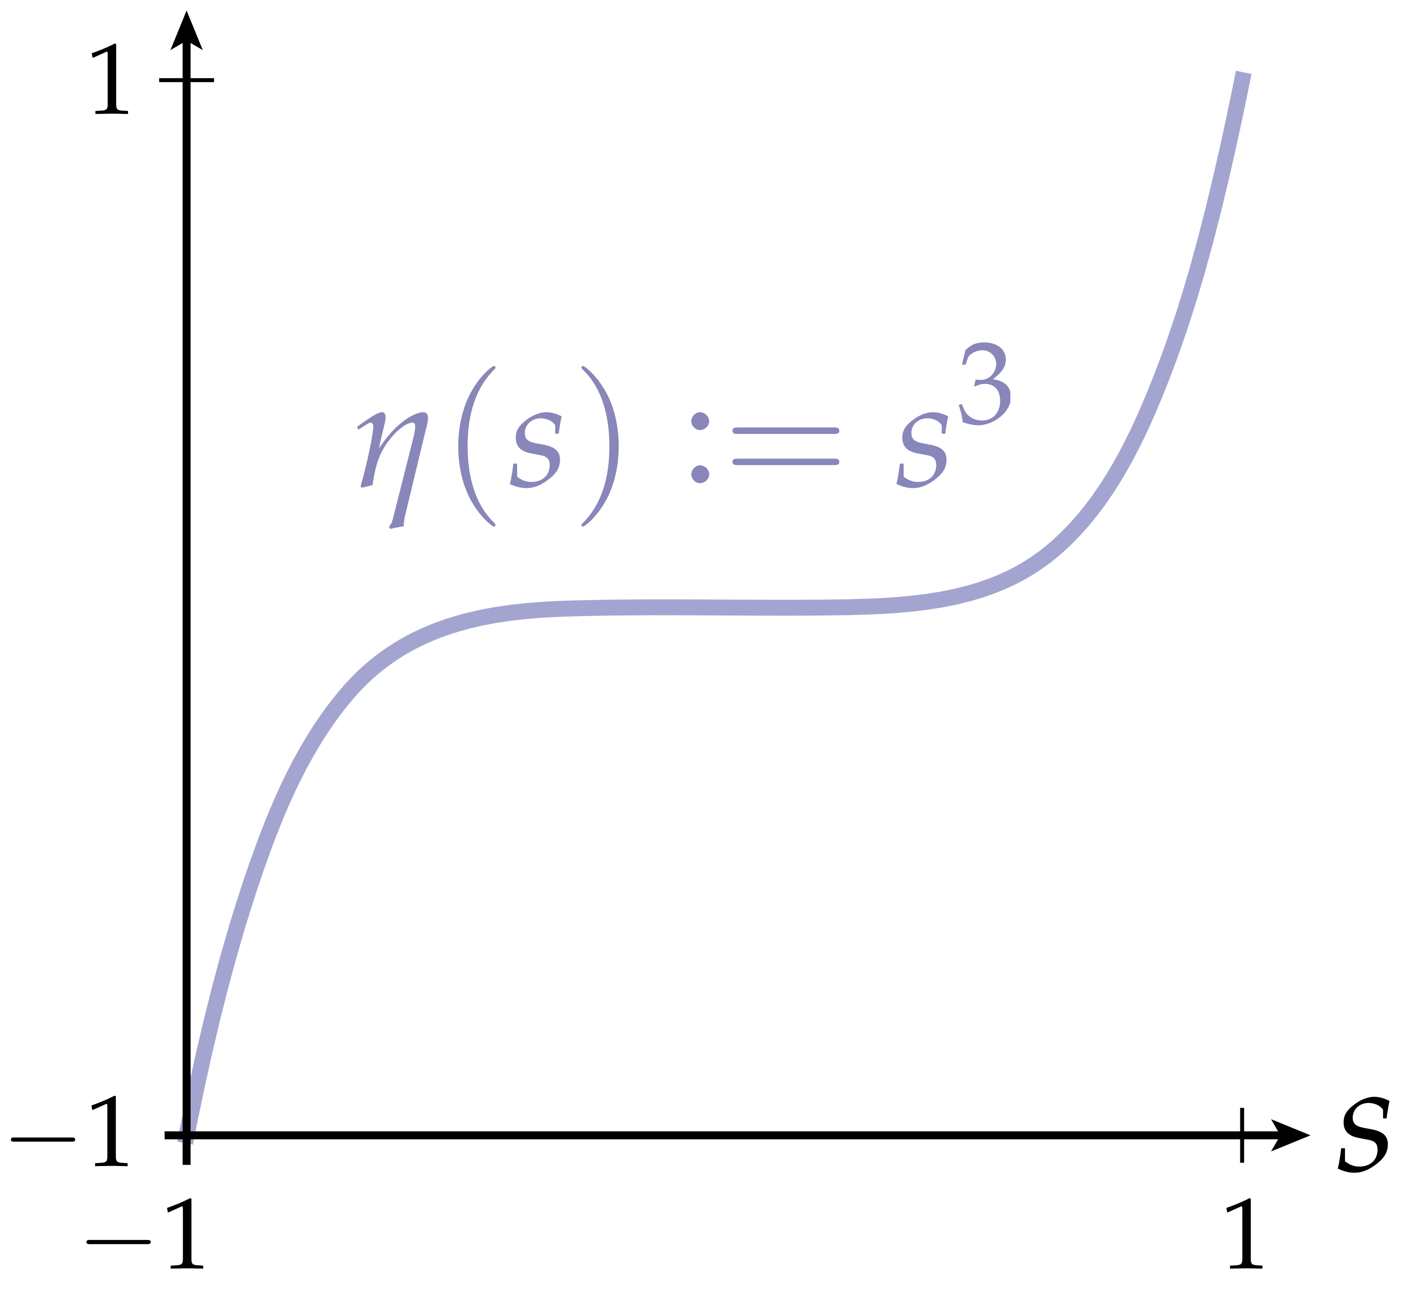

$\gamma(s) := (1 + s)(\cos(\pi s), \sin(\pi s))$

Regular Curve / Immersion

- A parameterized curve is *regular* (or *immersed*) if the differential is nonzero everywhere, *i.e.*, if the curve "never slows to zero"
- Without this condition, a parameterized curve may look non-smooth but actually be differentiable everywhere, or look smooth but fail to have well-defined tangents.

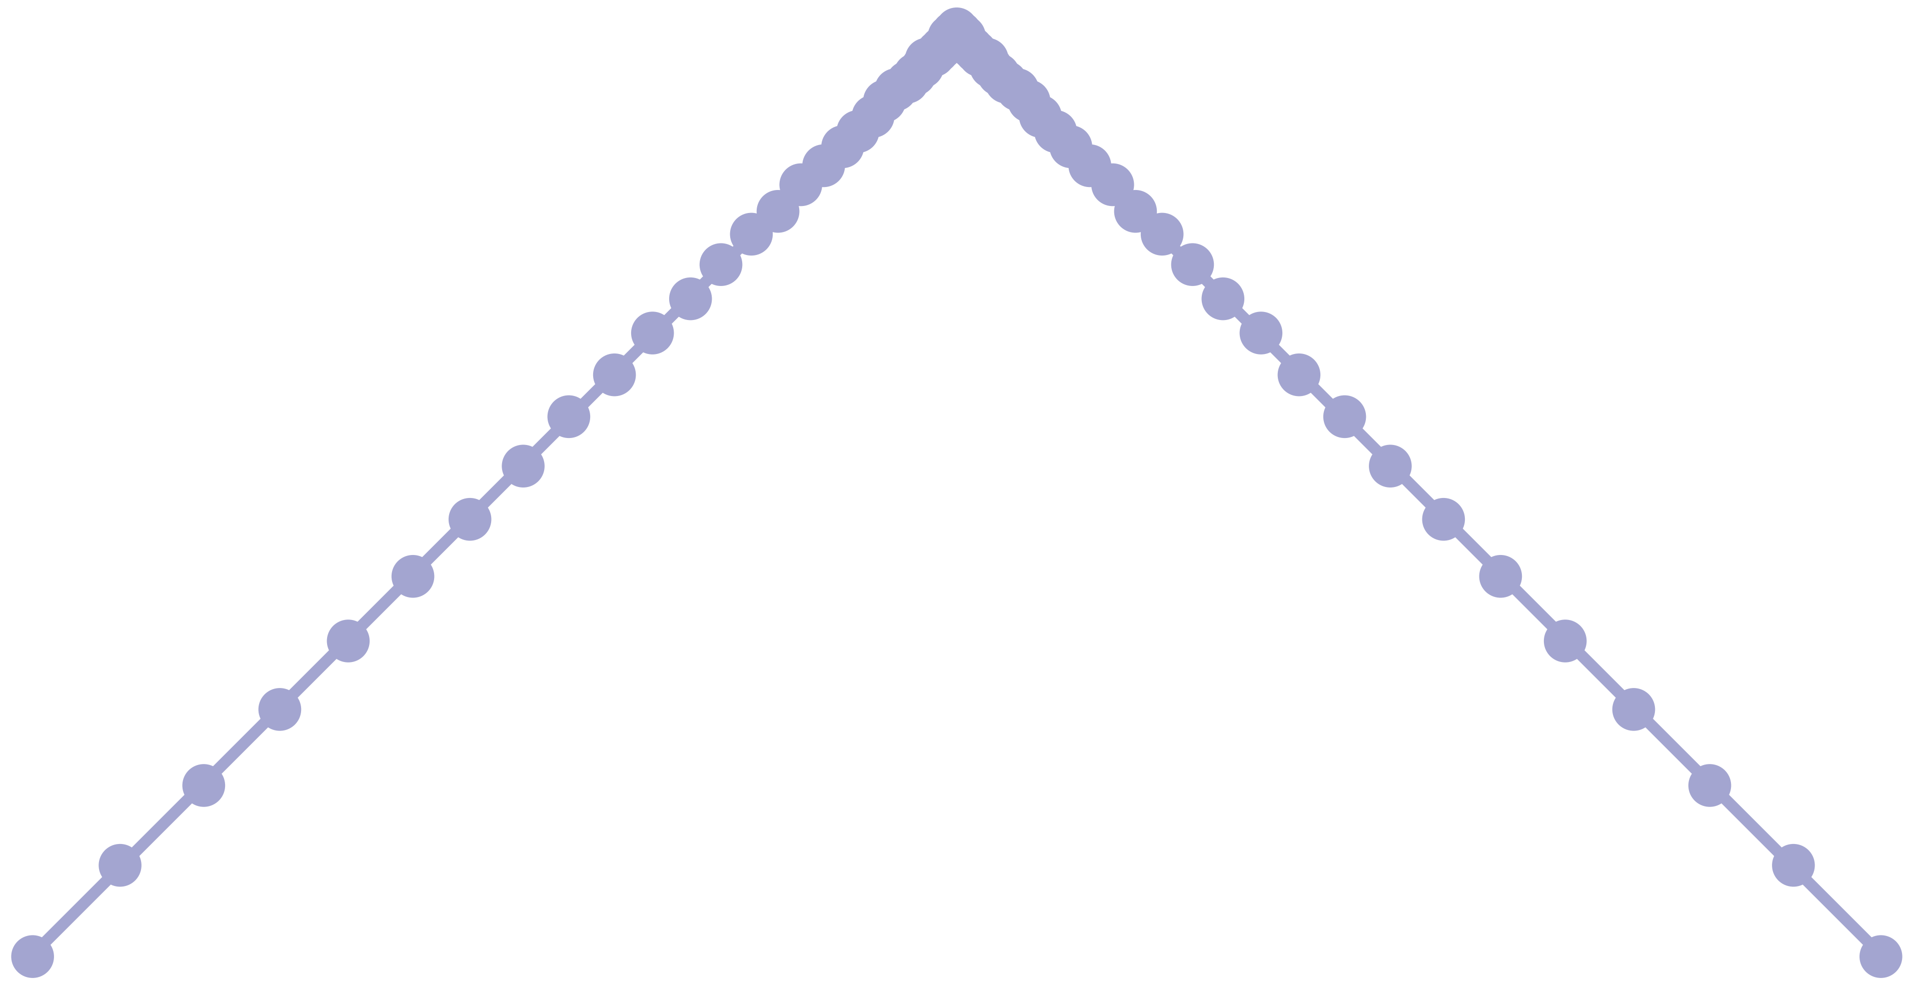

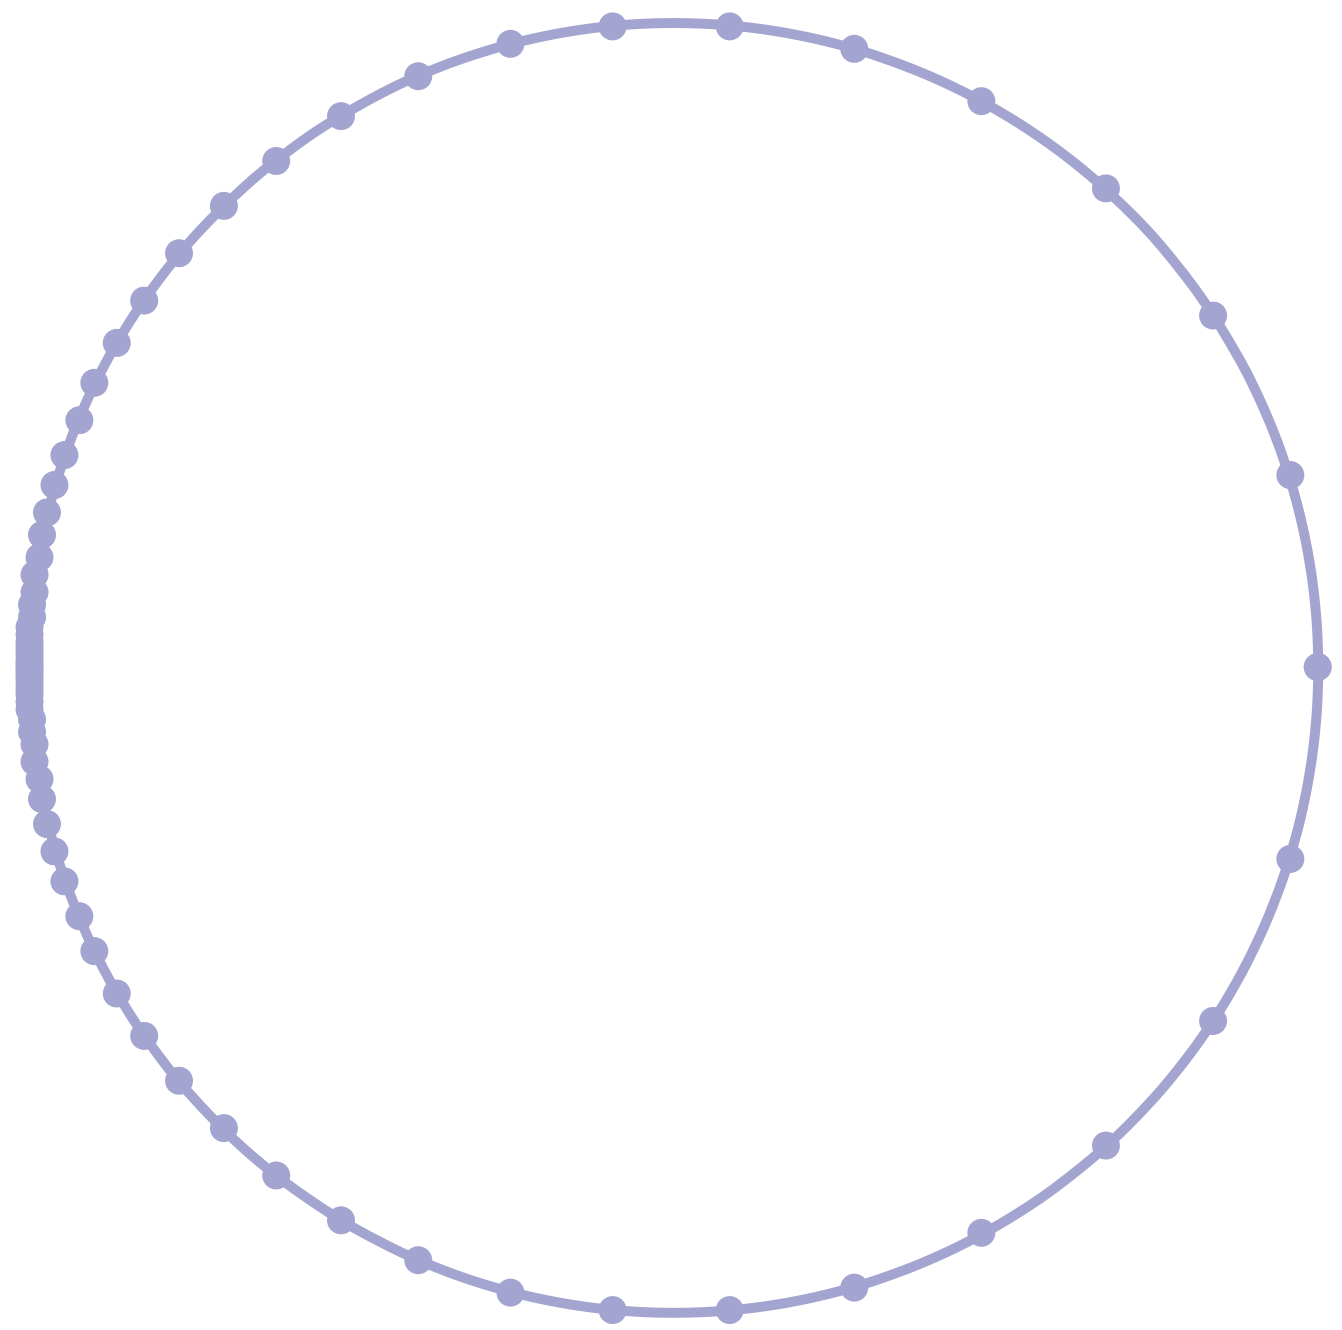

Irregular Curve—Example

- Consider the reparameterization of a piecewise linear curve:

$$\eta(s) := s^3 \qquad \gamma(s) := \begin{cases} (s,s) & s \leq 0 \\ (s,-s) & s > 0 \end{cases} \qquad \tilde{\gamma}(s) = \begin{cases} (s^3,s^3) & s \leq 0, \\ (s^3,-s^3) & s > 0 \end{cases}$$

- Even though the reparameterized curve has a continuous first derivative, this derivative goes to zero at $s = 0$:

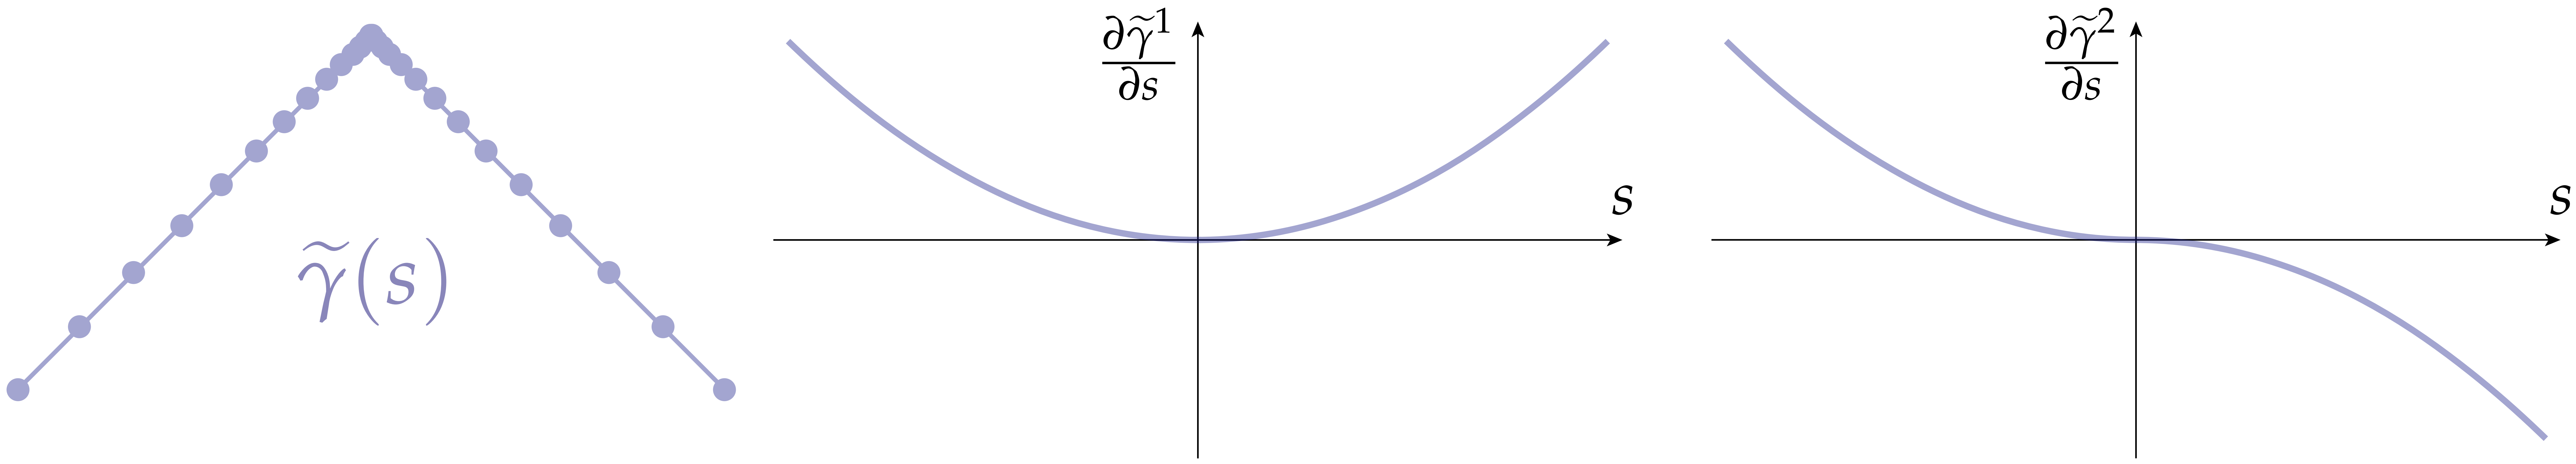

- Hence, (still) can't define tangent at zero.

Embedded Curve

- Roughly speaking, an *embedded* curve does not cross itself
- More precisely, a curve is embedded if it is a continuous and bijective map from its domain to its image, and the inverse map is also continuous
- **Q:** What's an example of a continuous injective curve that is not embedded?
- **A:** A half-open interval mapped to a circle (inverse is not continuous)

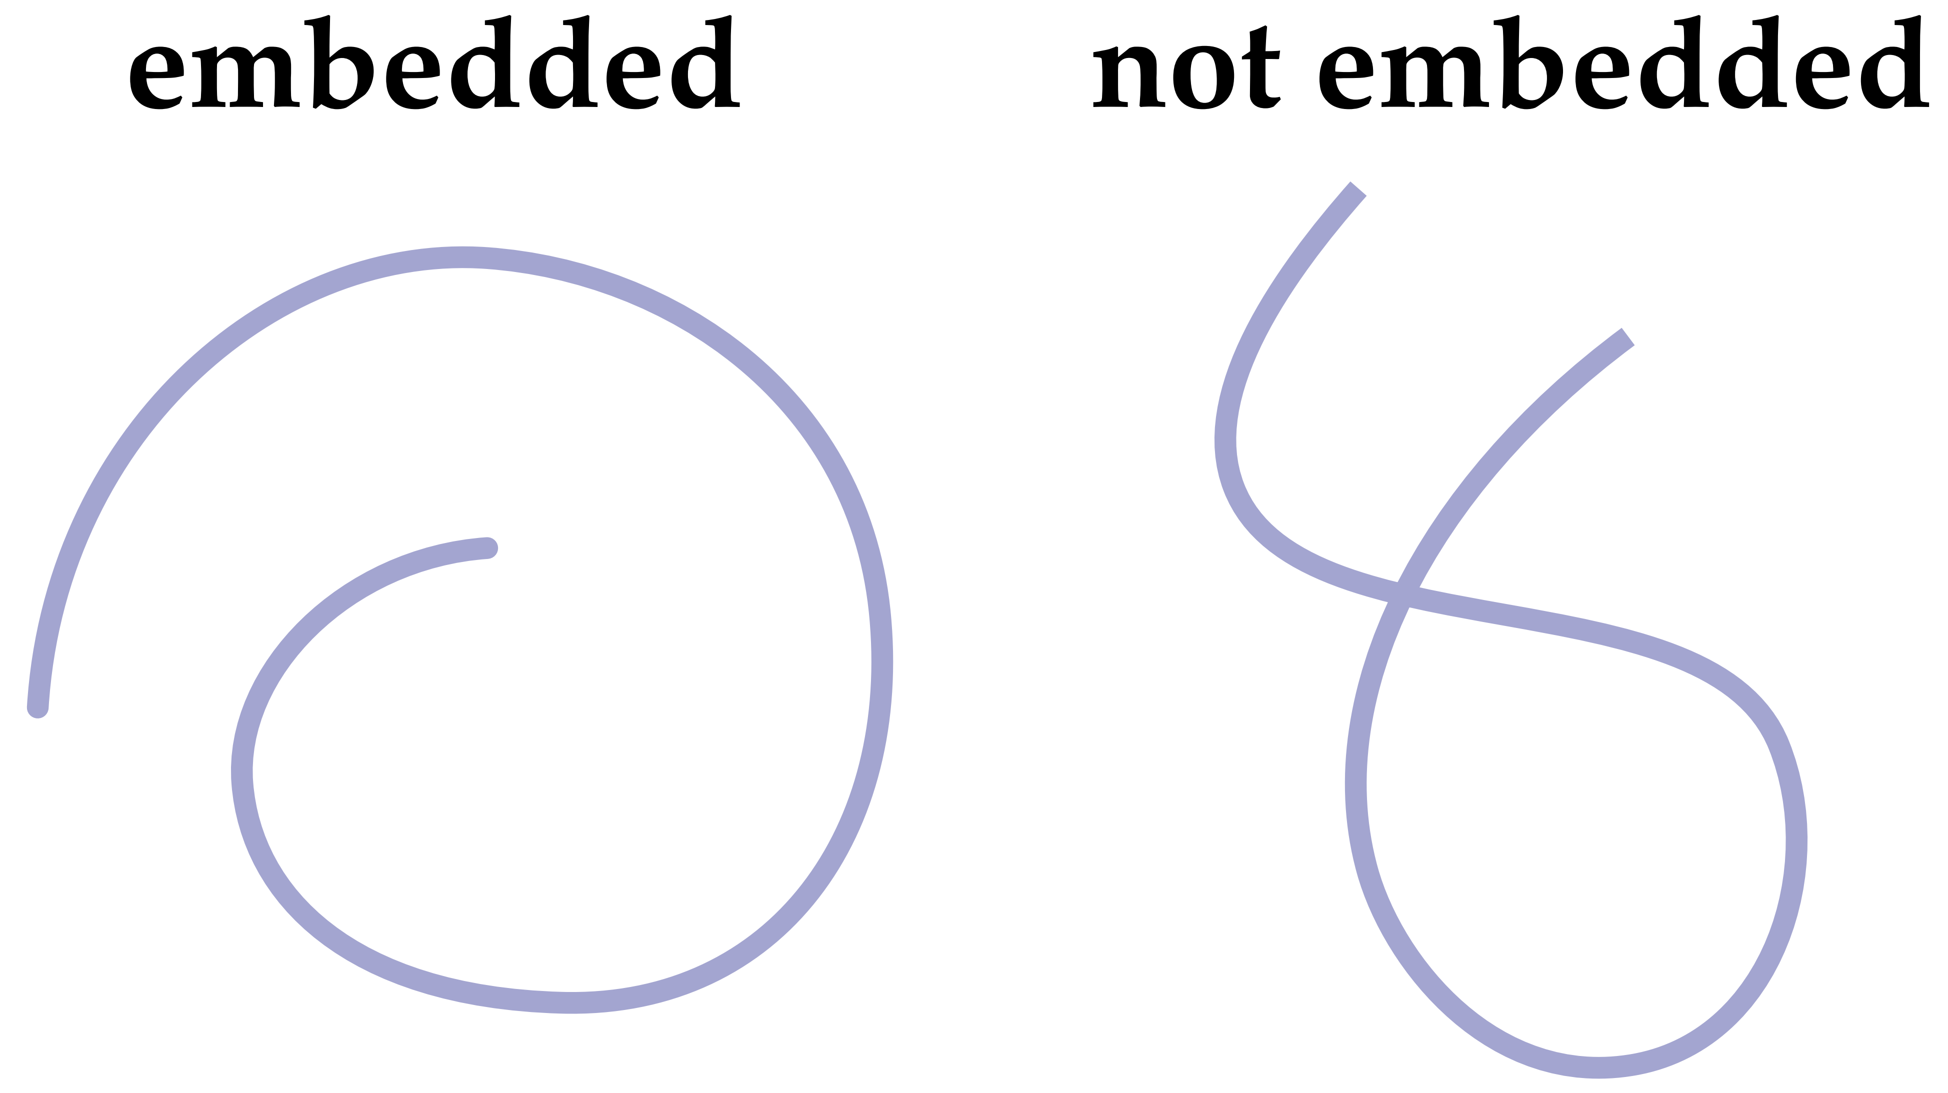

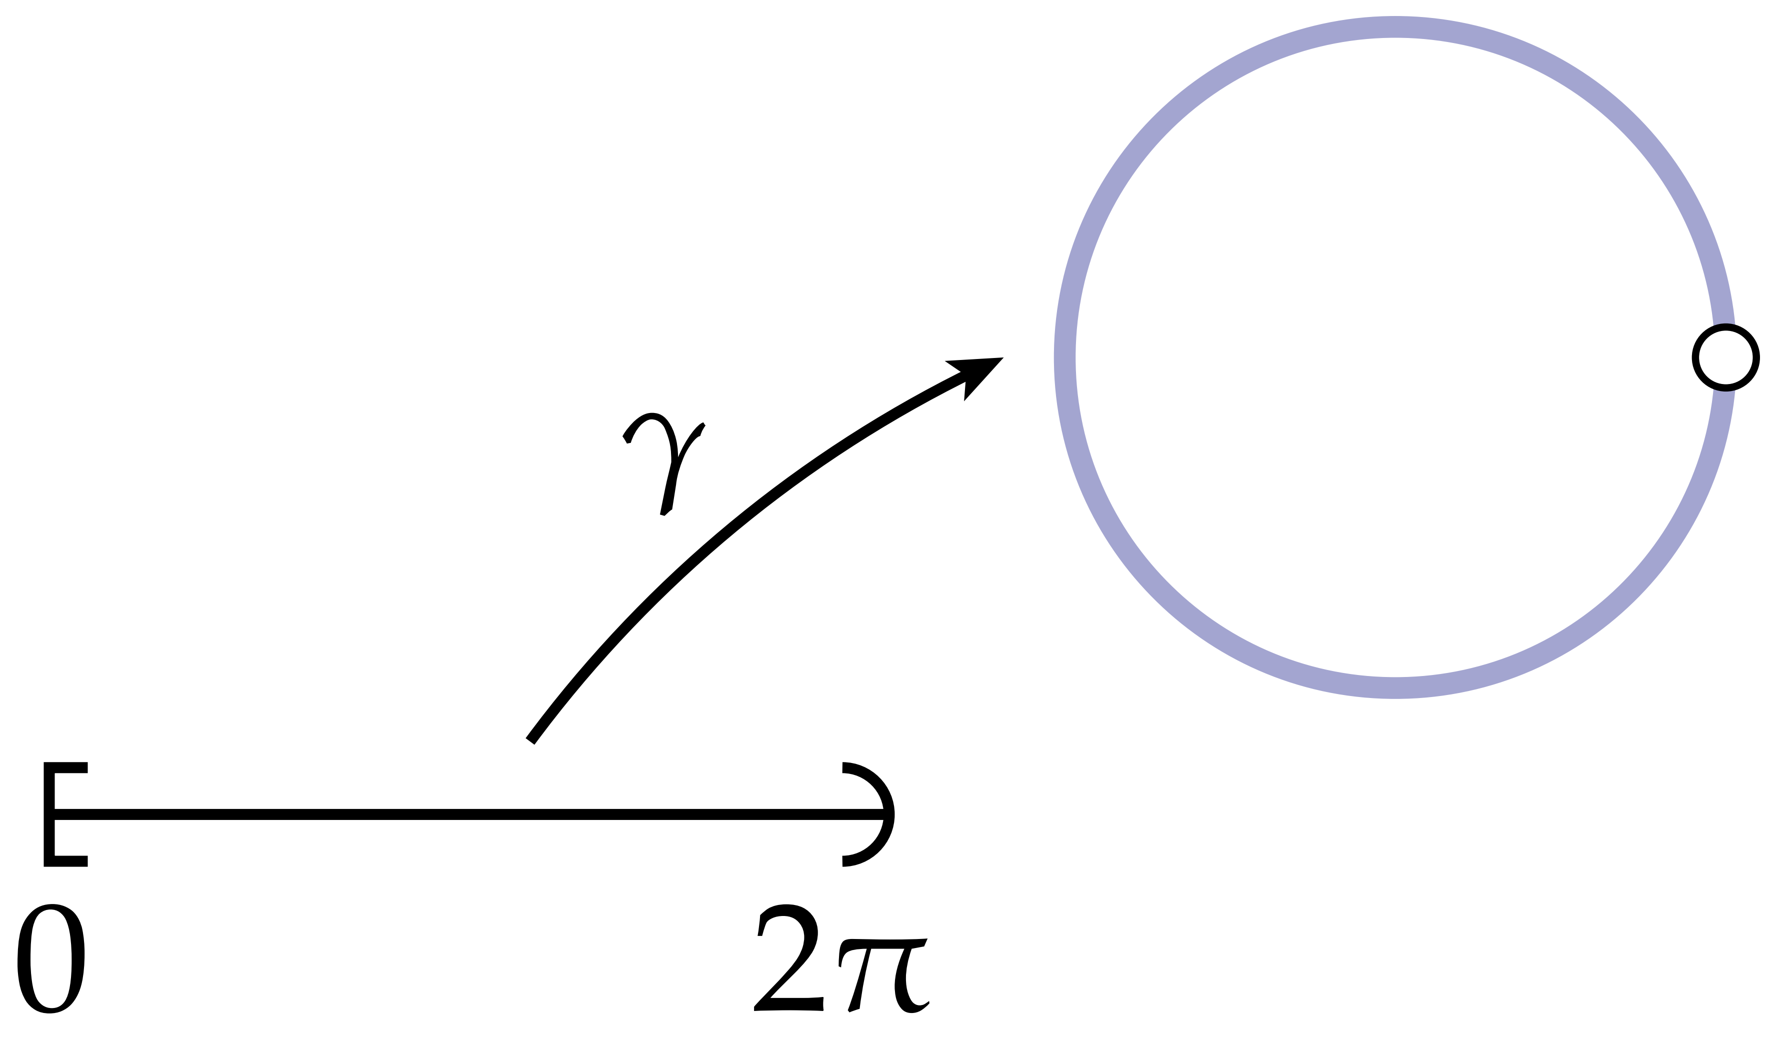

Normal of a Curve

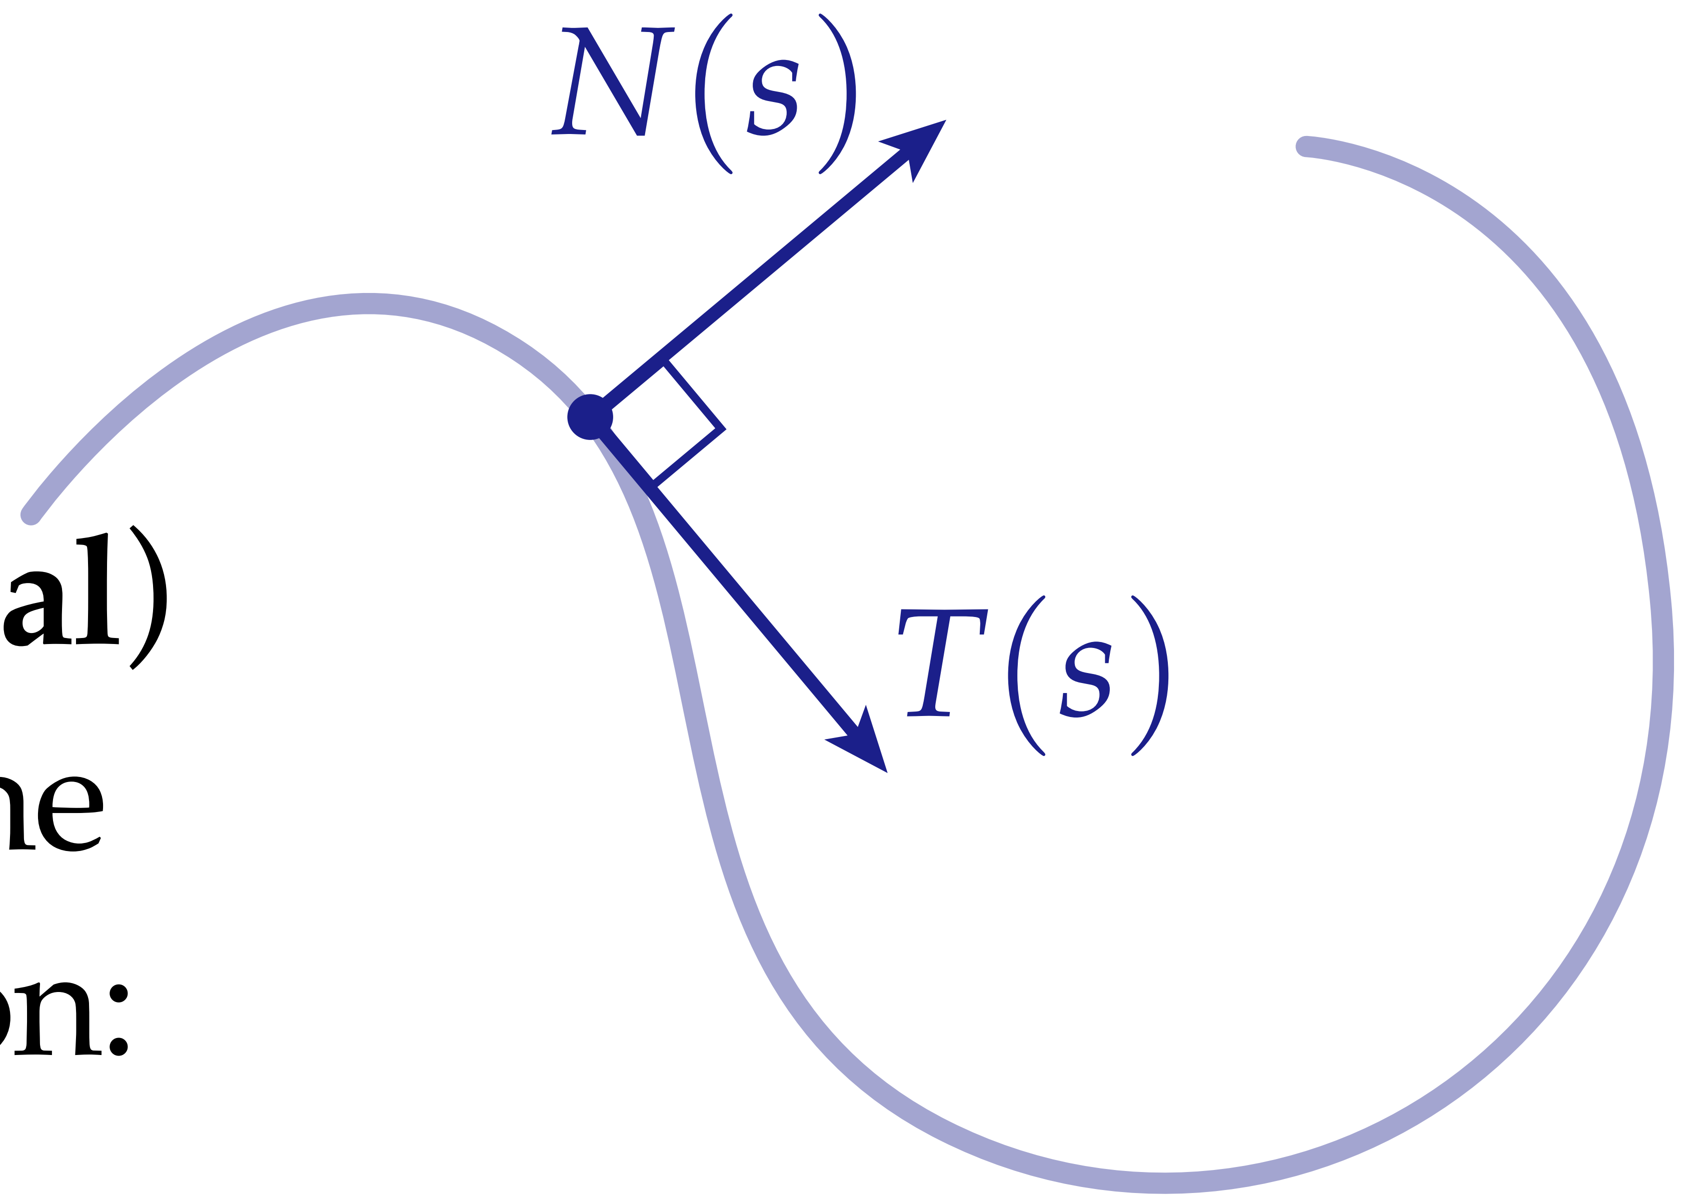

- Informally, a vector is *normal* to a curve if it “sticks straight out” of the curve.
- More formally, the **unit normal** (or just **normal**) can be expressed as a quarter-rotation $\mathcal{J}$ of the unit tangent in the counter-clockwise direction:

$$N(s) := \mathcal{J}T(s)$$

- In coordinates (x,y), a quarter-turn can be achieved by* simply exchanging x and y, and then negating y:

$$(x, y) \overset{\mathcal{J}}{\mapsto} (-y, x)$$

*Why does this work?

Normal of a Curve—Example

- Let's compute the unit normal of a circle:

$$\gamma : [0, 2\pi) \to \mathbb{R}^2;\ s \mapsto (\cos(s), \sin(s))$$

$$T(s) = (-\sin(s), \cos(s))$$

$$N(s) = \mathcal{J}T(s) = (-\cos(s), -\sin(s))$$

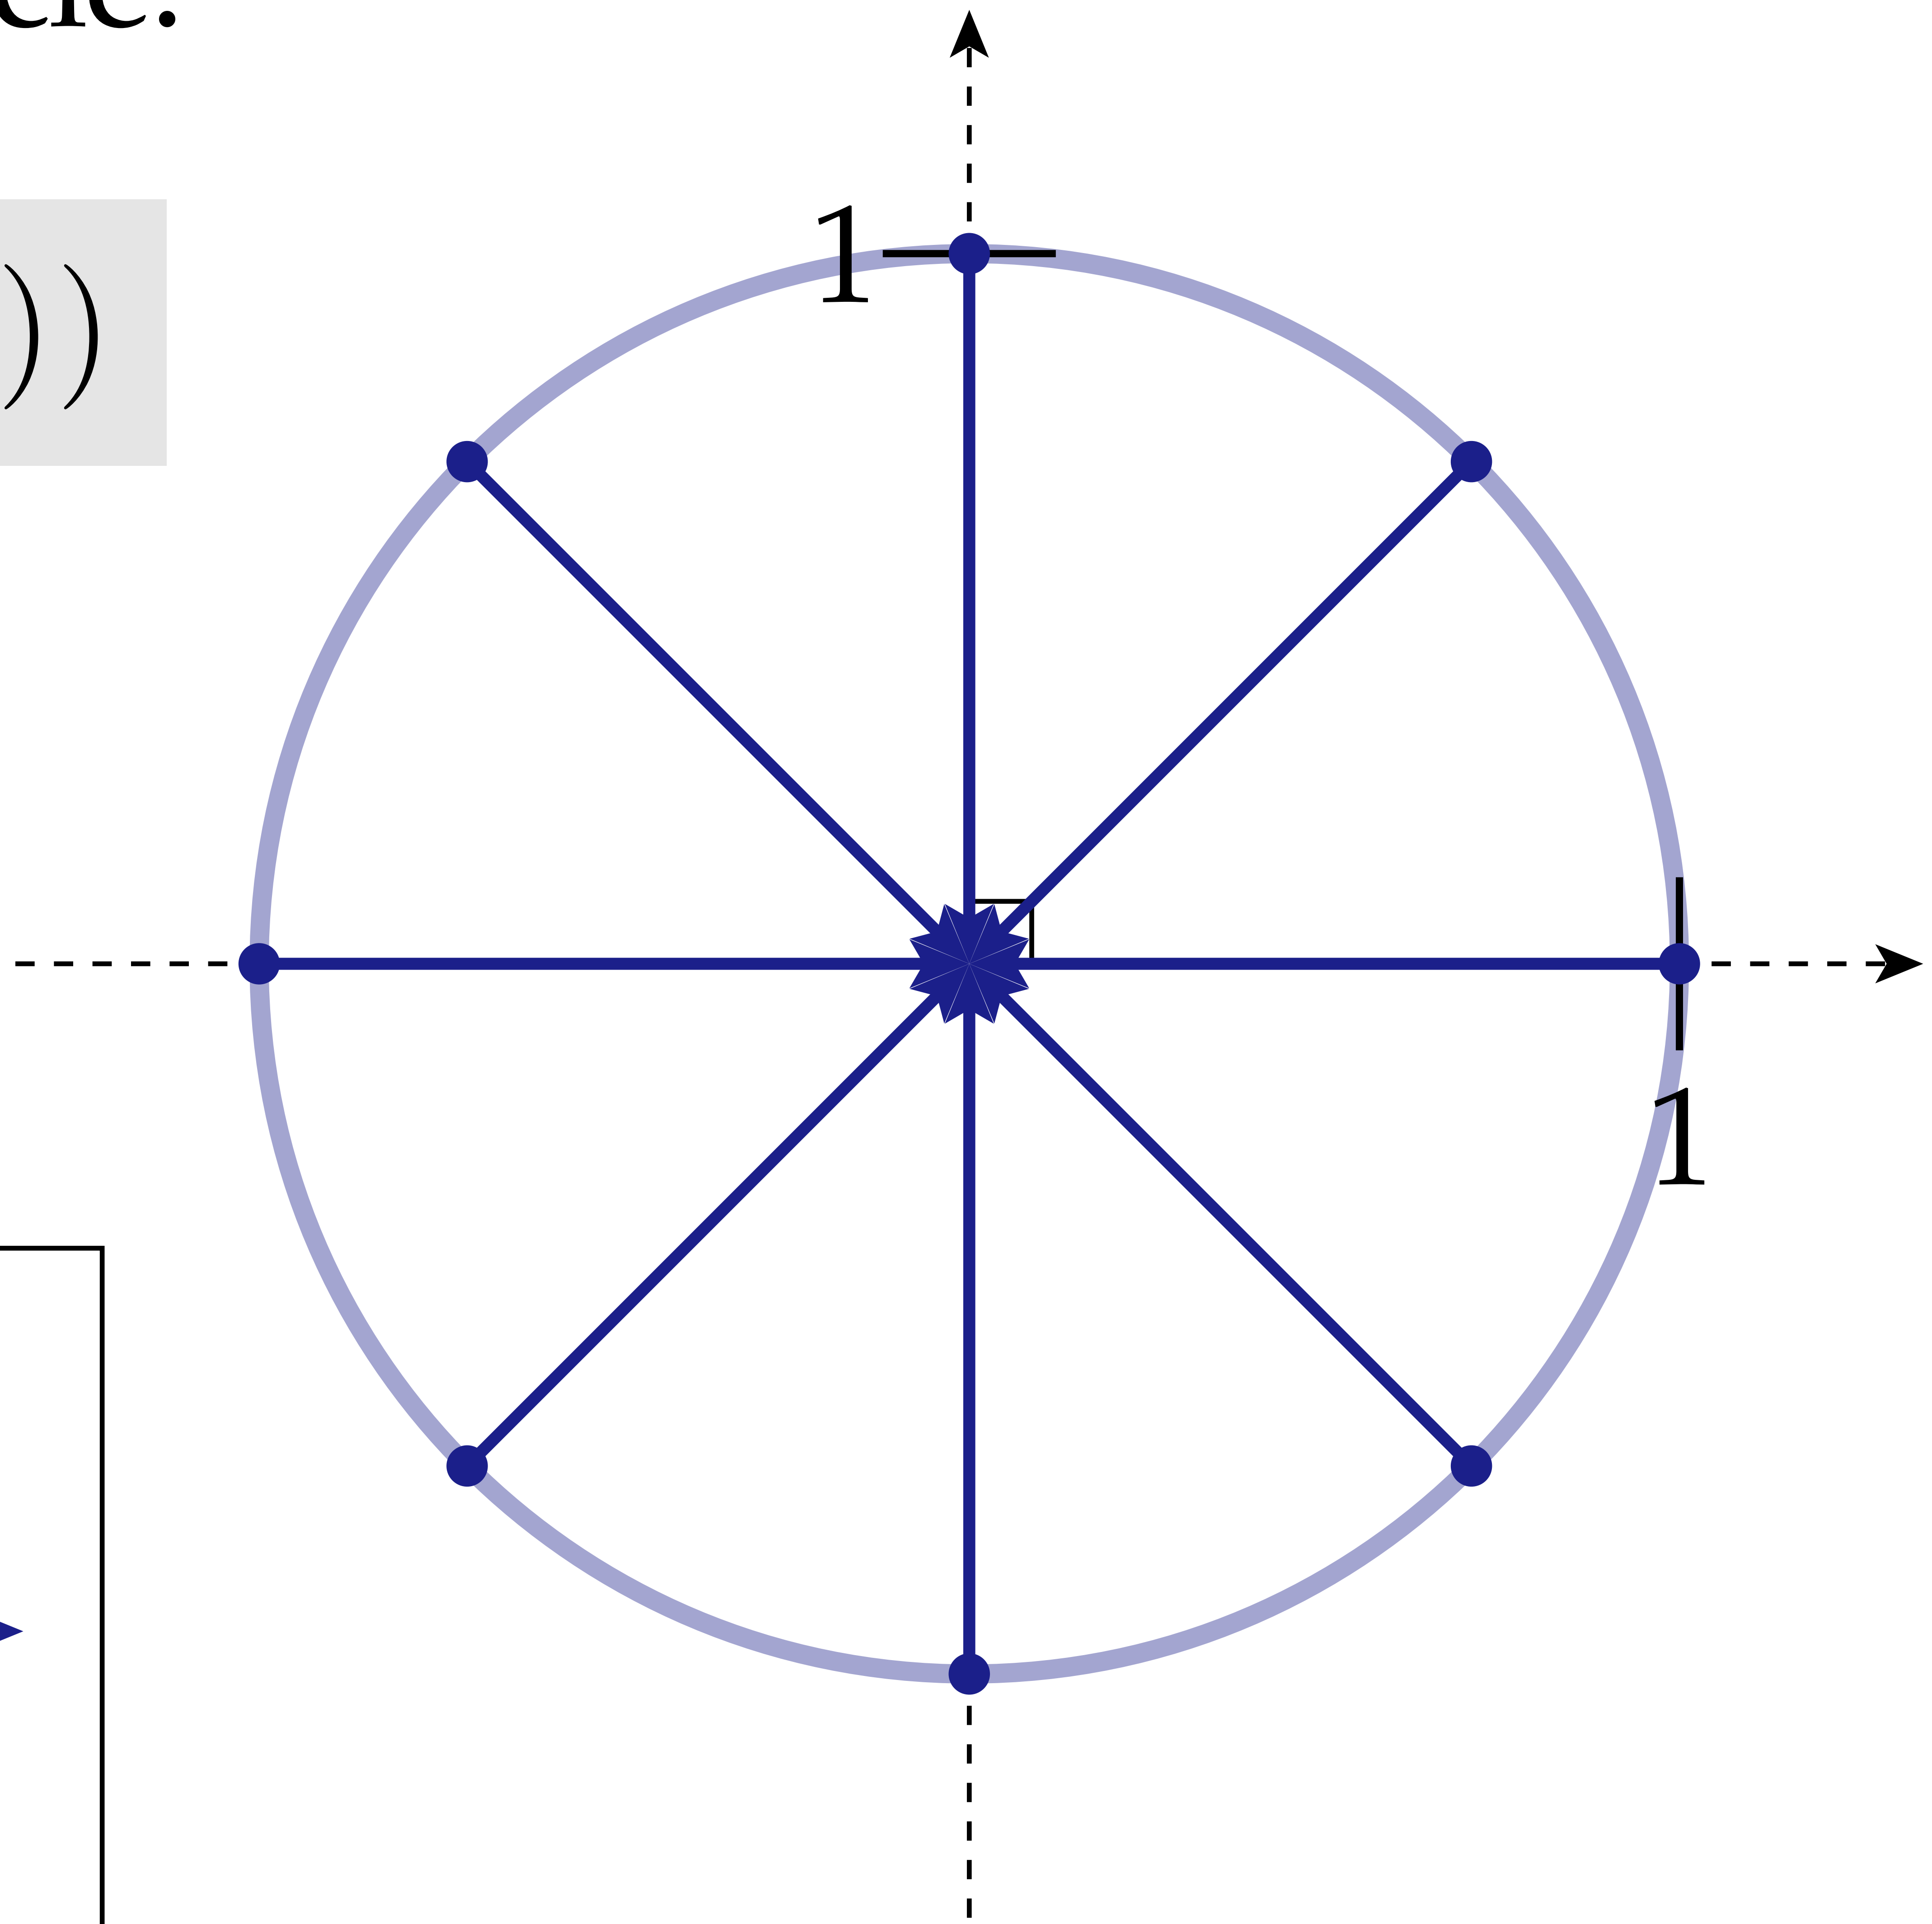

Note: could also adopt the convention $N = -\mathcal{J}T$. (Just remain consistent!)

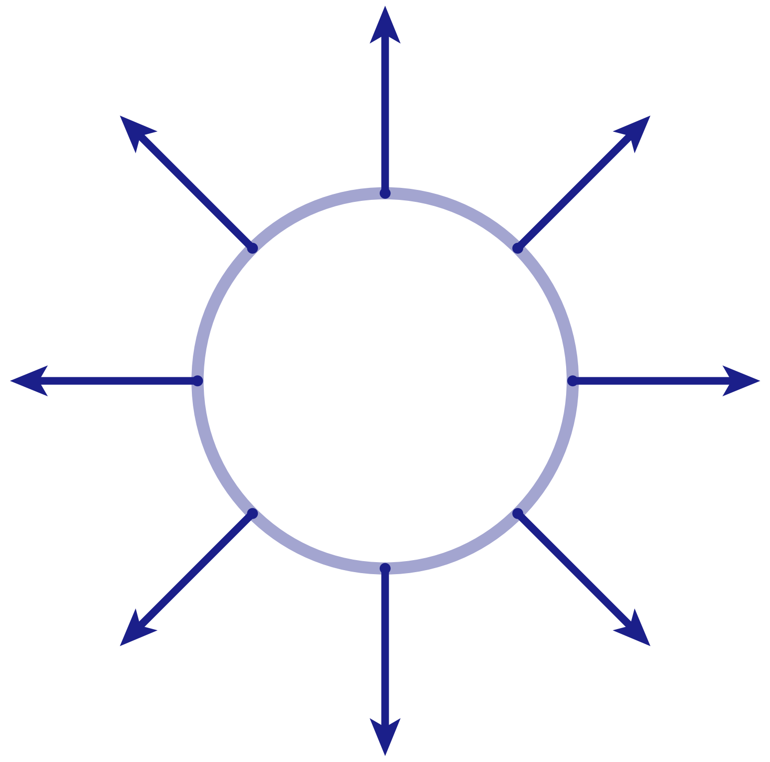

Curvature of a Plane Curve

- Informally, curvature describes "how much a curve bends"
- More formally, the **curvature** of an arc-length parameterized plane curve can be expressed as the rate of change in the tangent

$$\begin{aligned}\kappa(s) &:= \langle N(s), \tfrac{d}{ds}T(s)\rangle \\ &= \langle N(s), \tfrac{d^2}{ds^2}\gamma(s)\rangle\end{aligned}$$

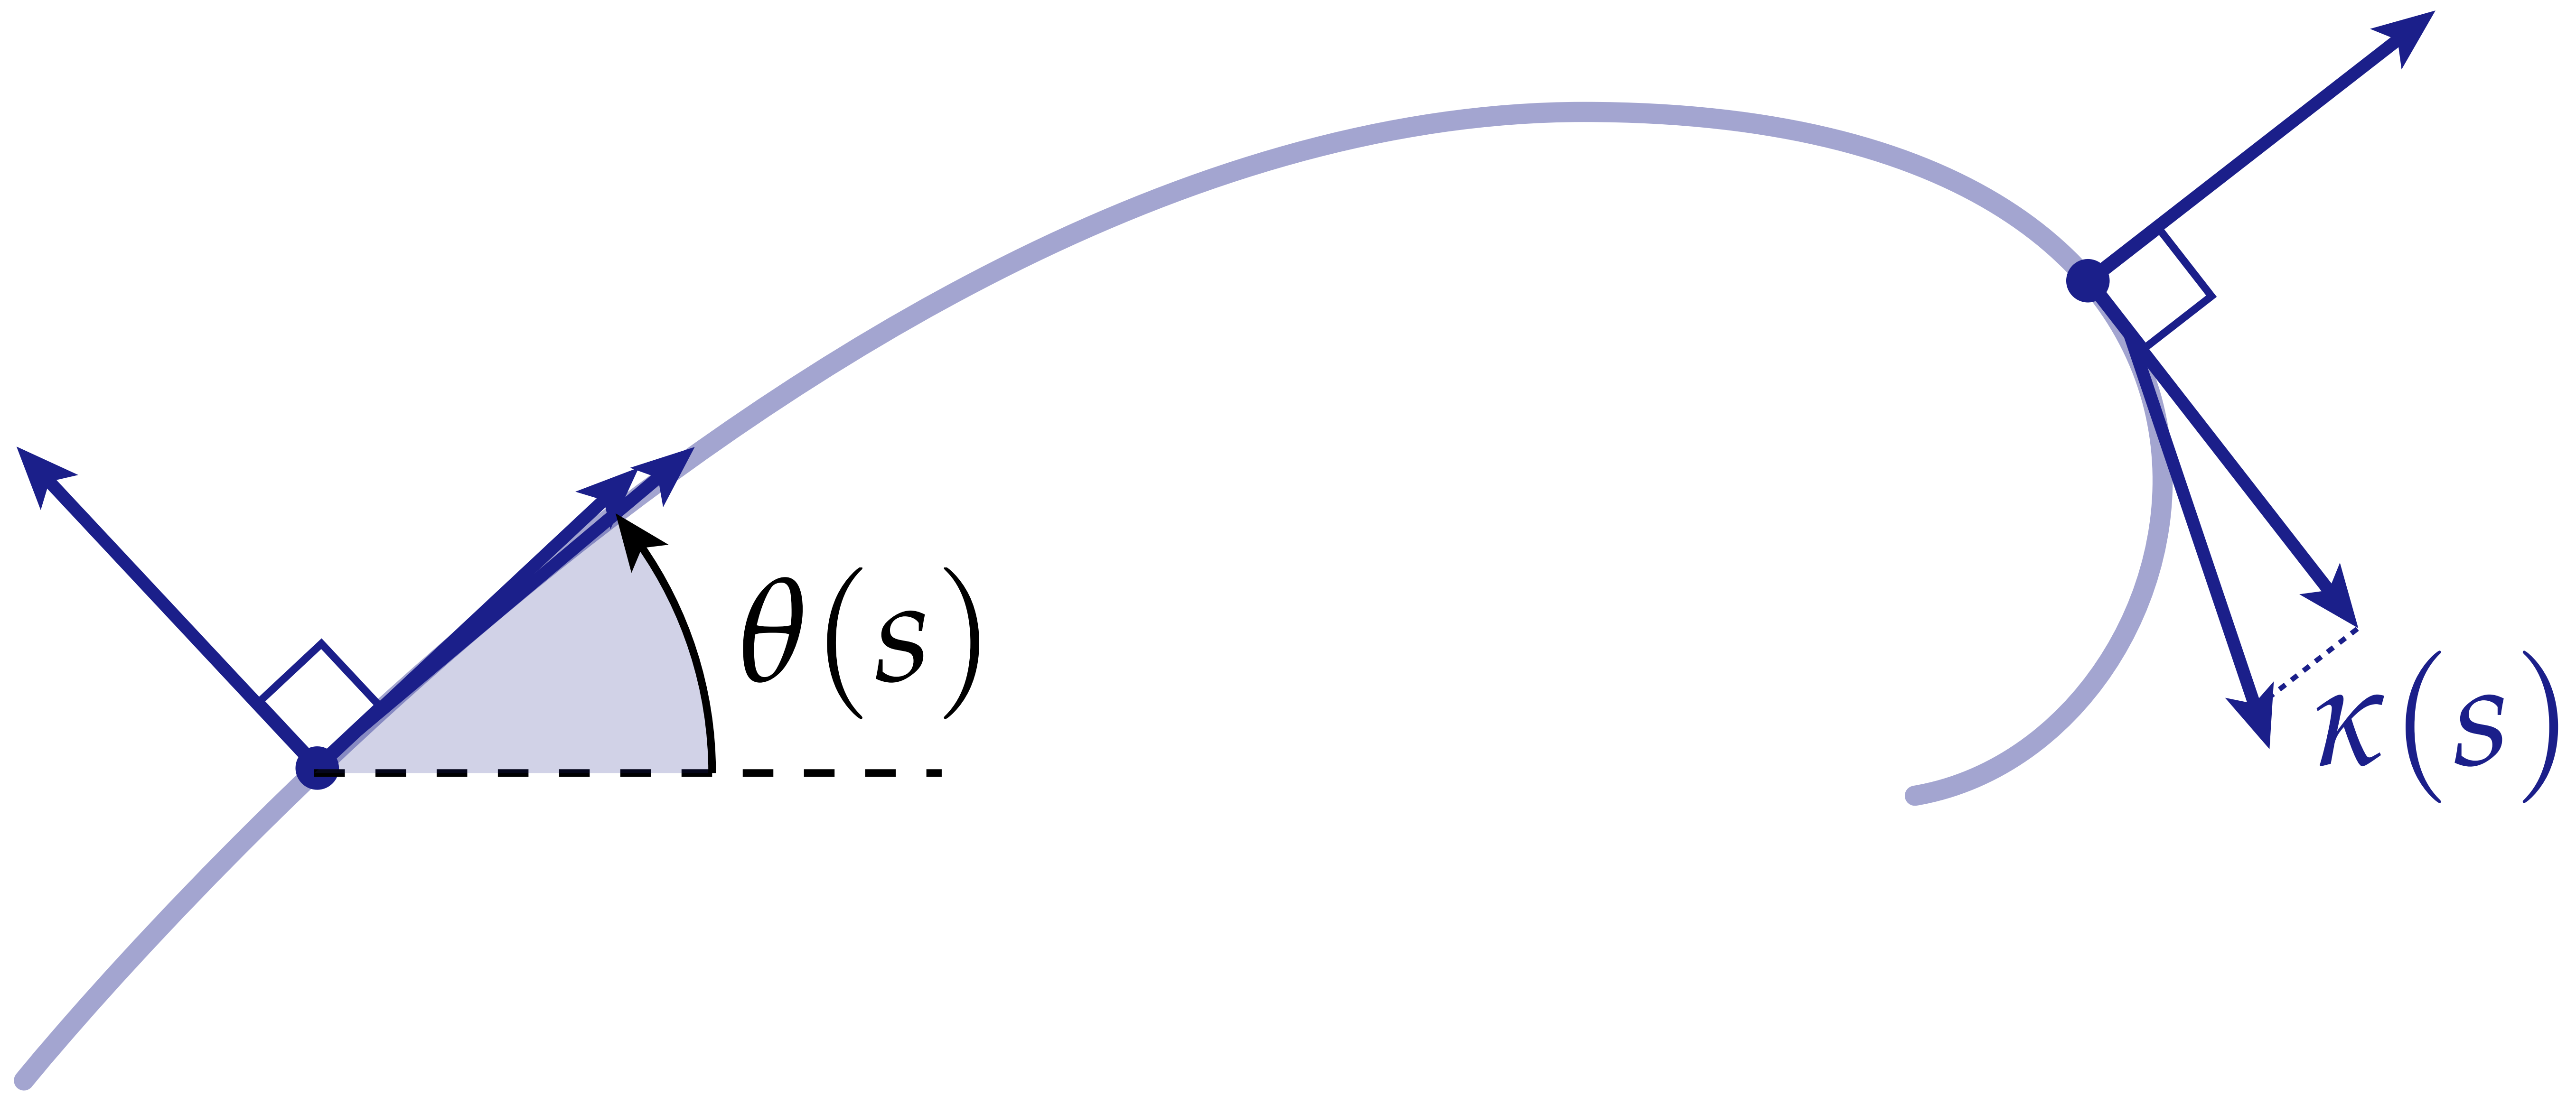

Equivalently:

$$\kappa(s) = \frac{d}{ds}\theta(s)$$

Here the angle brackets denote the usual dot product, i.e., $\langle (a,b),(x,y)\rangle := ax + by$.

Fundamental Theorem of Plane Curves

Fact. Up to rigid motions, an arc-length parameterized plane curve is uniquely determined by its curvature.

Q: Given only the curvature function, how can we recover the curve?

A: Just "invert" the two relationships $\frac{d}{ds}\theta = \kappa$, $\frac{d}{ds}\gamma = T$

First integrate curvature to get angle: $\theta(s) := \int_0^s \kappa(t)\, dt$

Then evaluate unit tangents: $T(s) := (\cos(\theta), \sin(\theta))$

Finally, integrate tangents to get curve: $\gamma(s) := \int_0^s T(t)\, dt$

Q: What about the rigid motion? Will this work for *closed* curves?

Turning and Winding Numbers

- For a closed regular curve in the plane…
- The **turning number** k is the number of counter-clockwise turns made by the tangent
- The **winding number** n is the number of times the curve goes around a particular point p
  - can also be viewed as the total *signed* length of the projection of the curve onto a unit-length circle around p

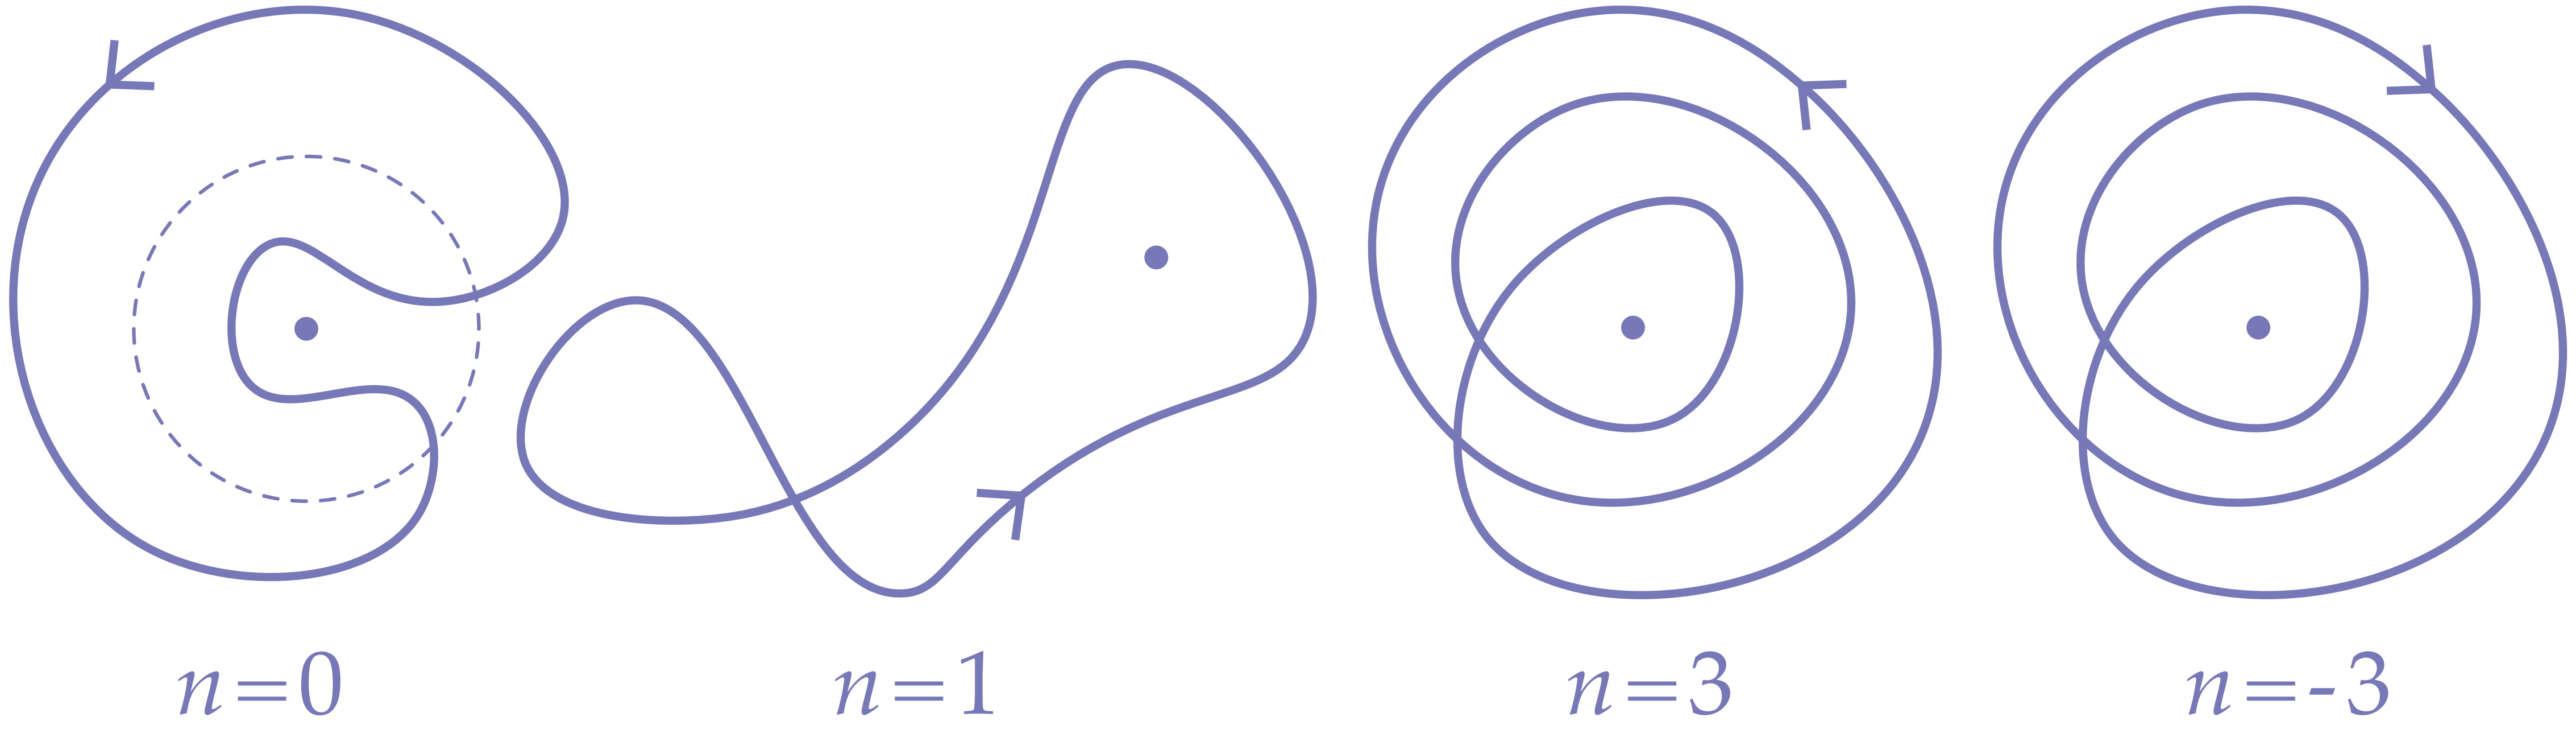

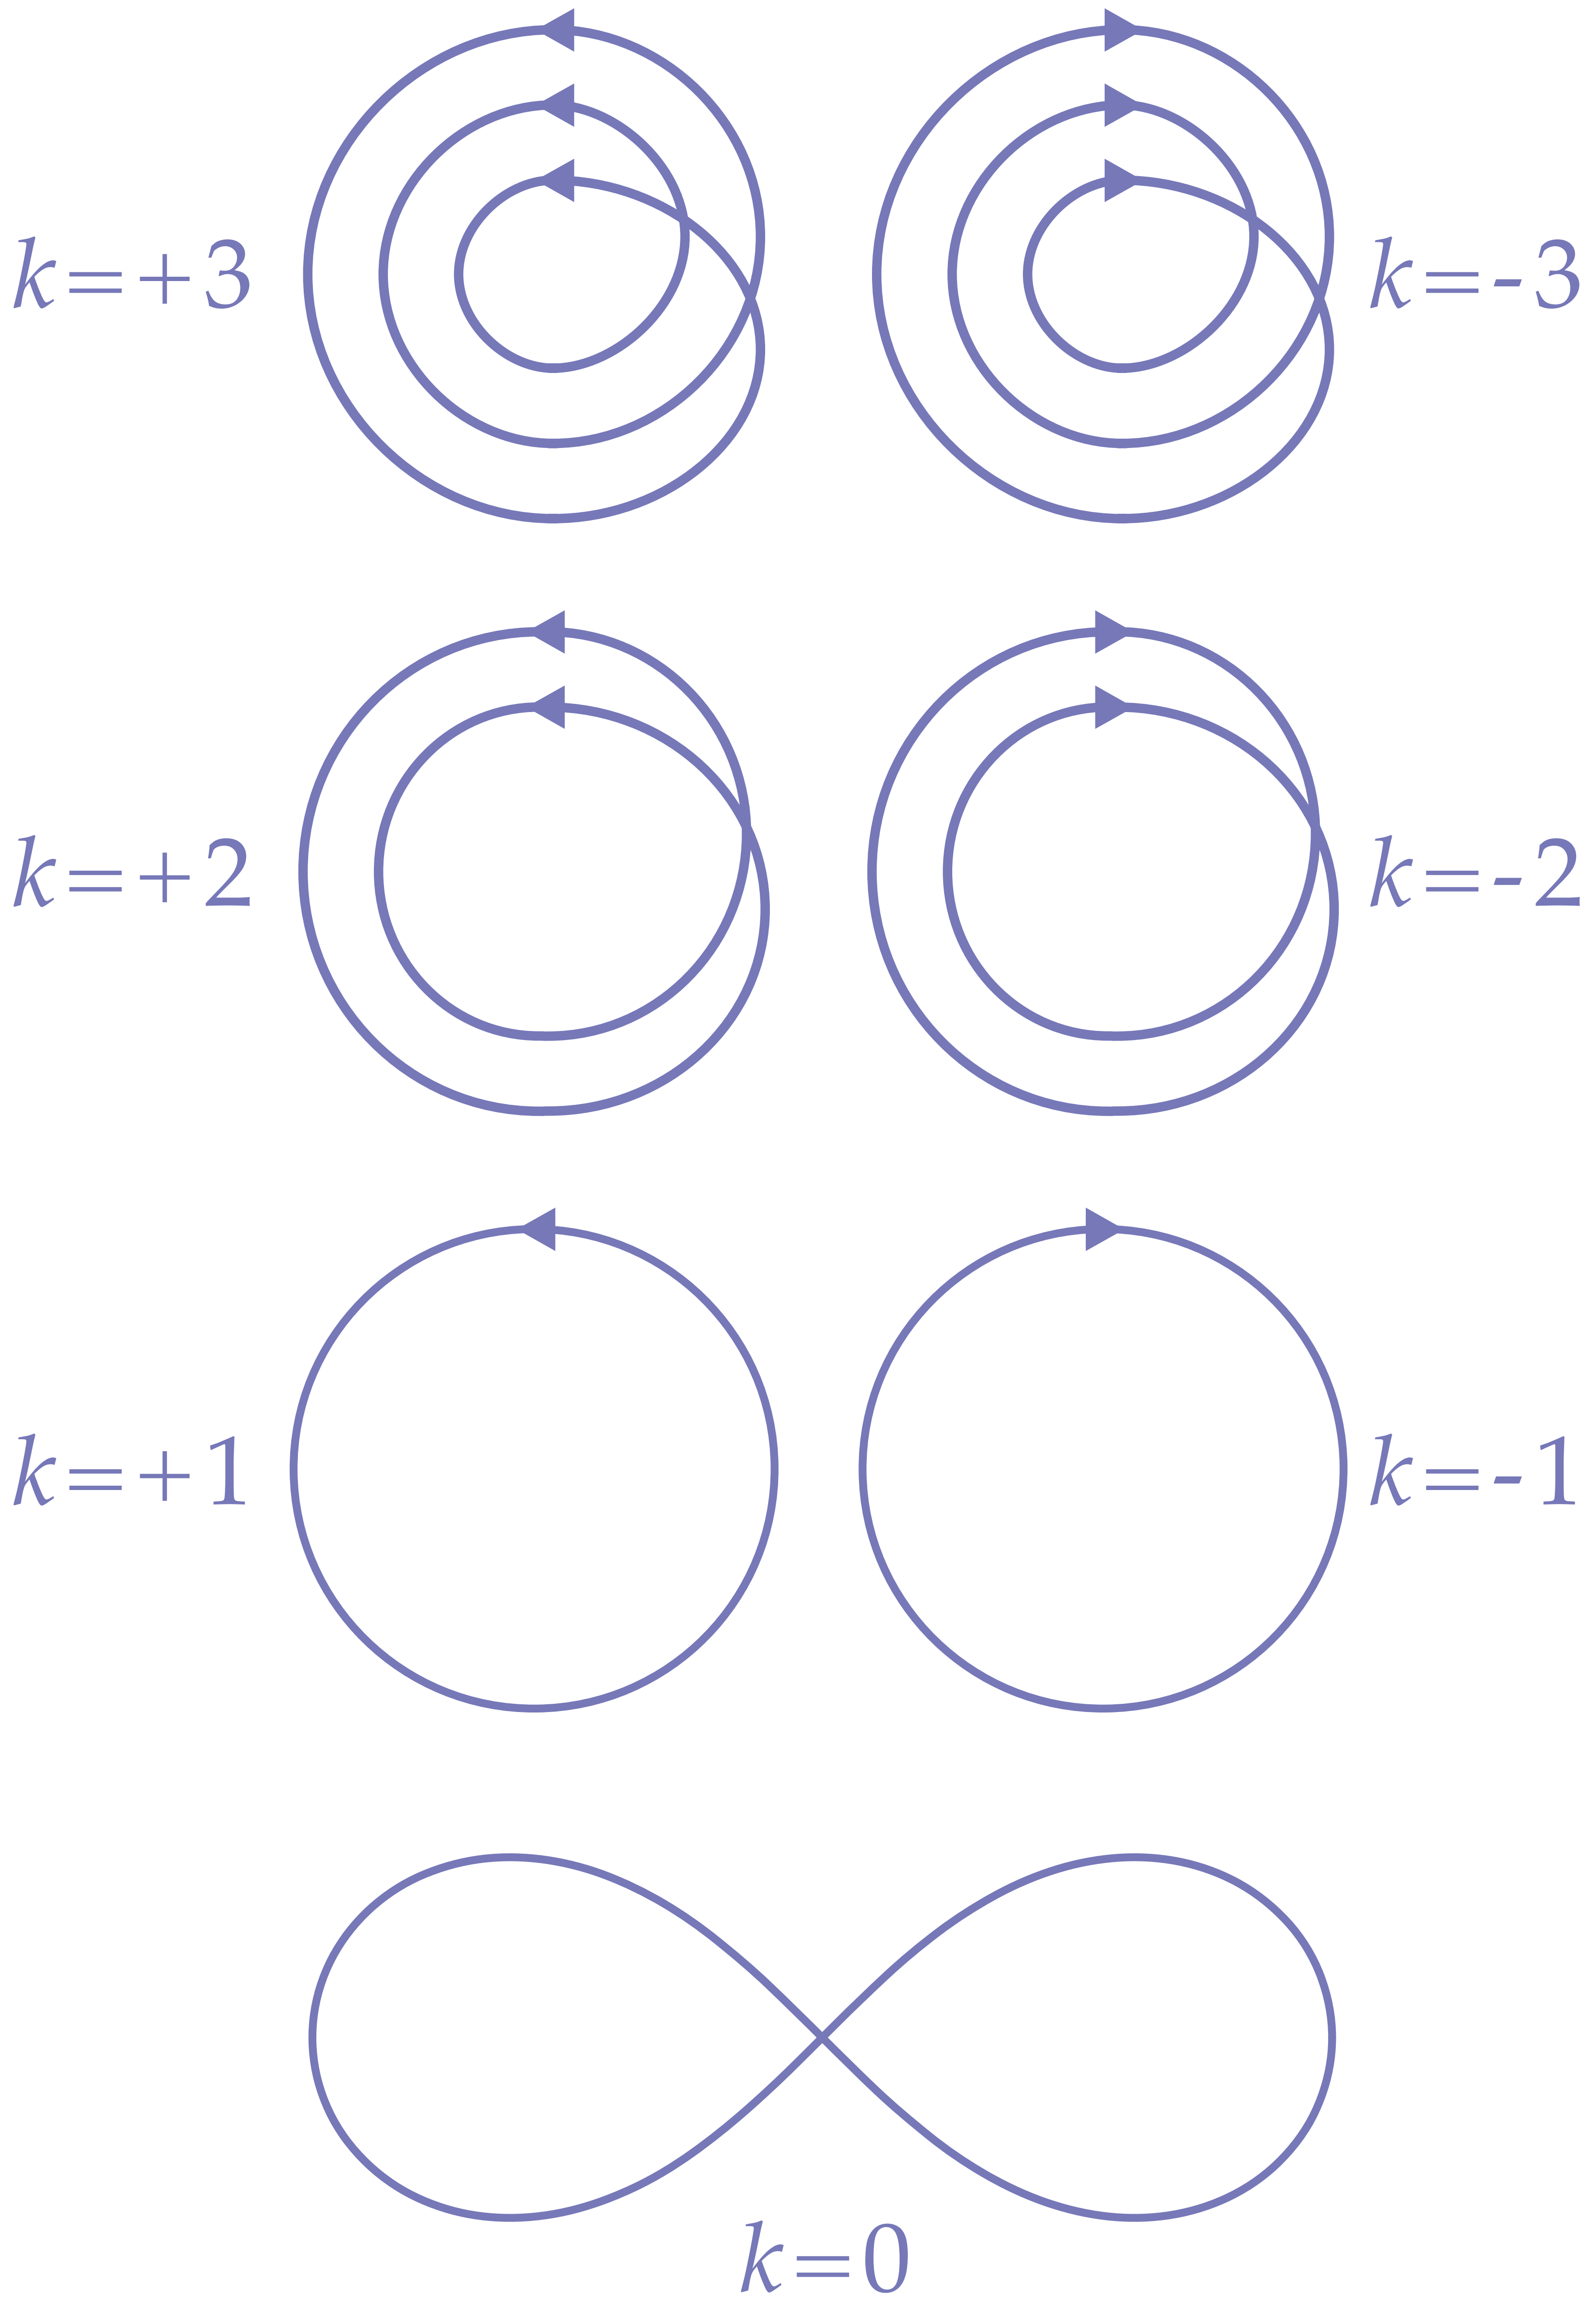

Whitney-Graustein Theorem

- (Whitney-Graustein) Two curves have the same *turning number k* if and only if they are related by *regular homotopy*, i.e., if one can continuously “deform” into the other while remaining regular (immersed).

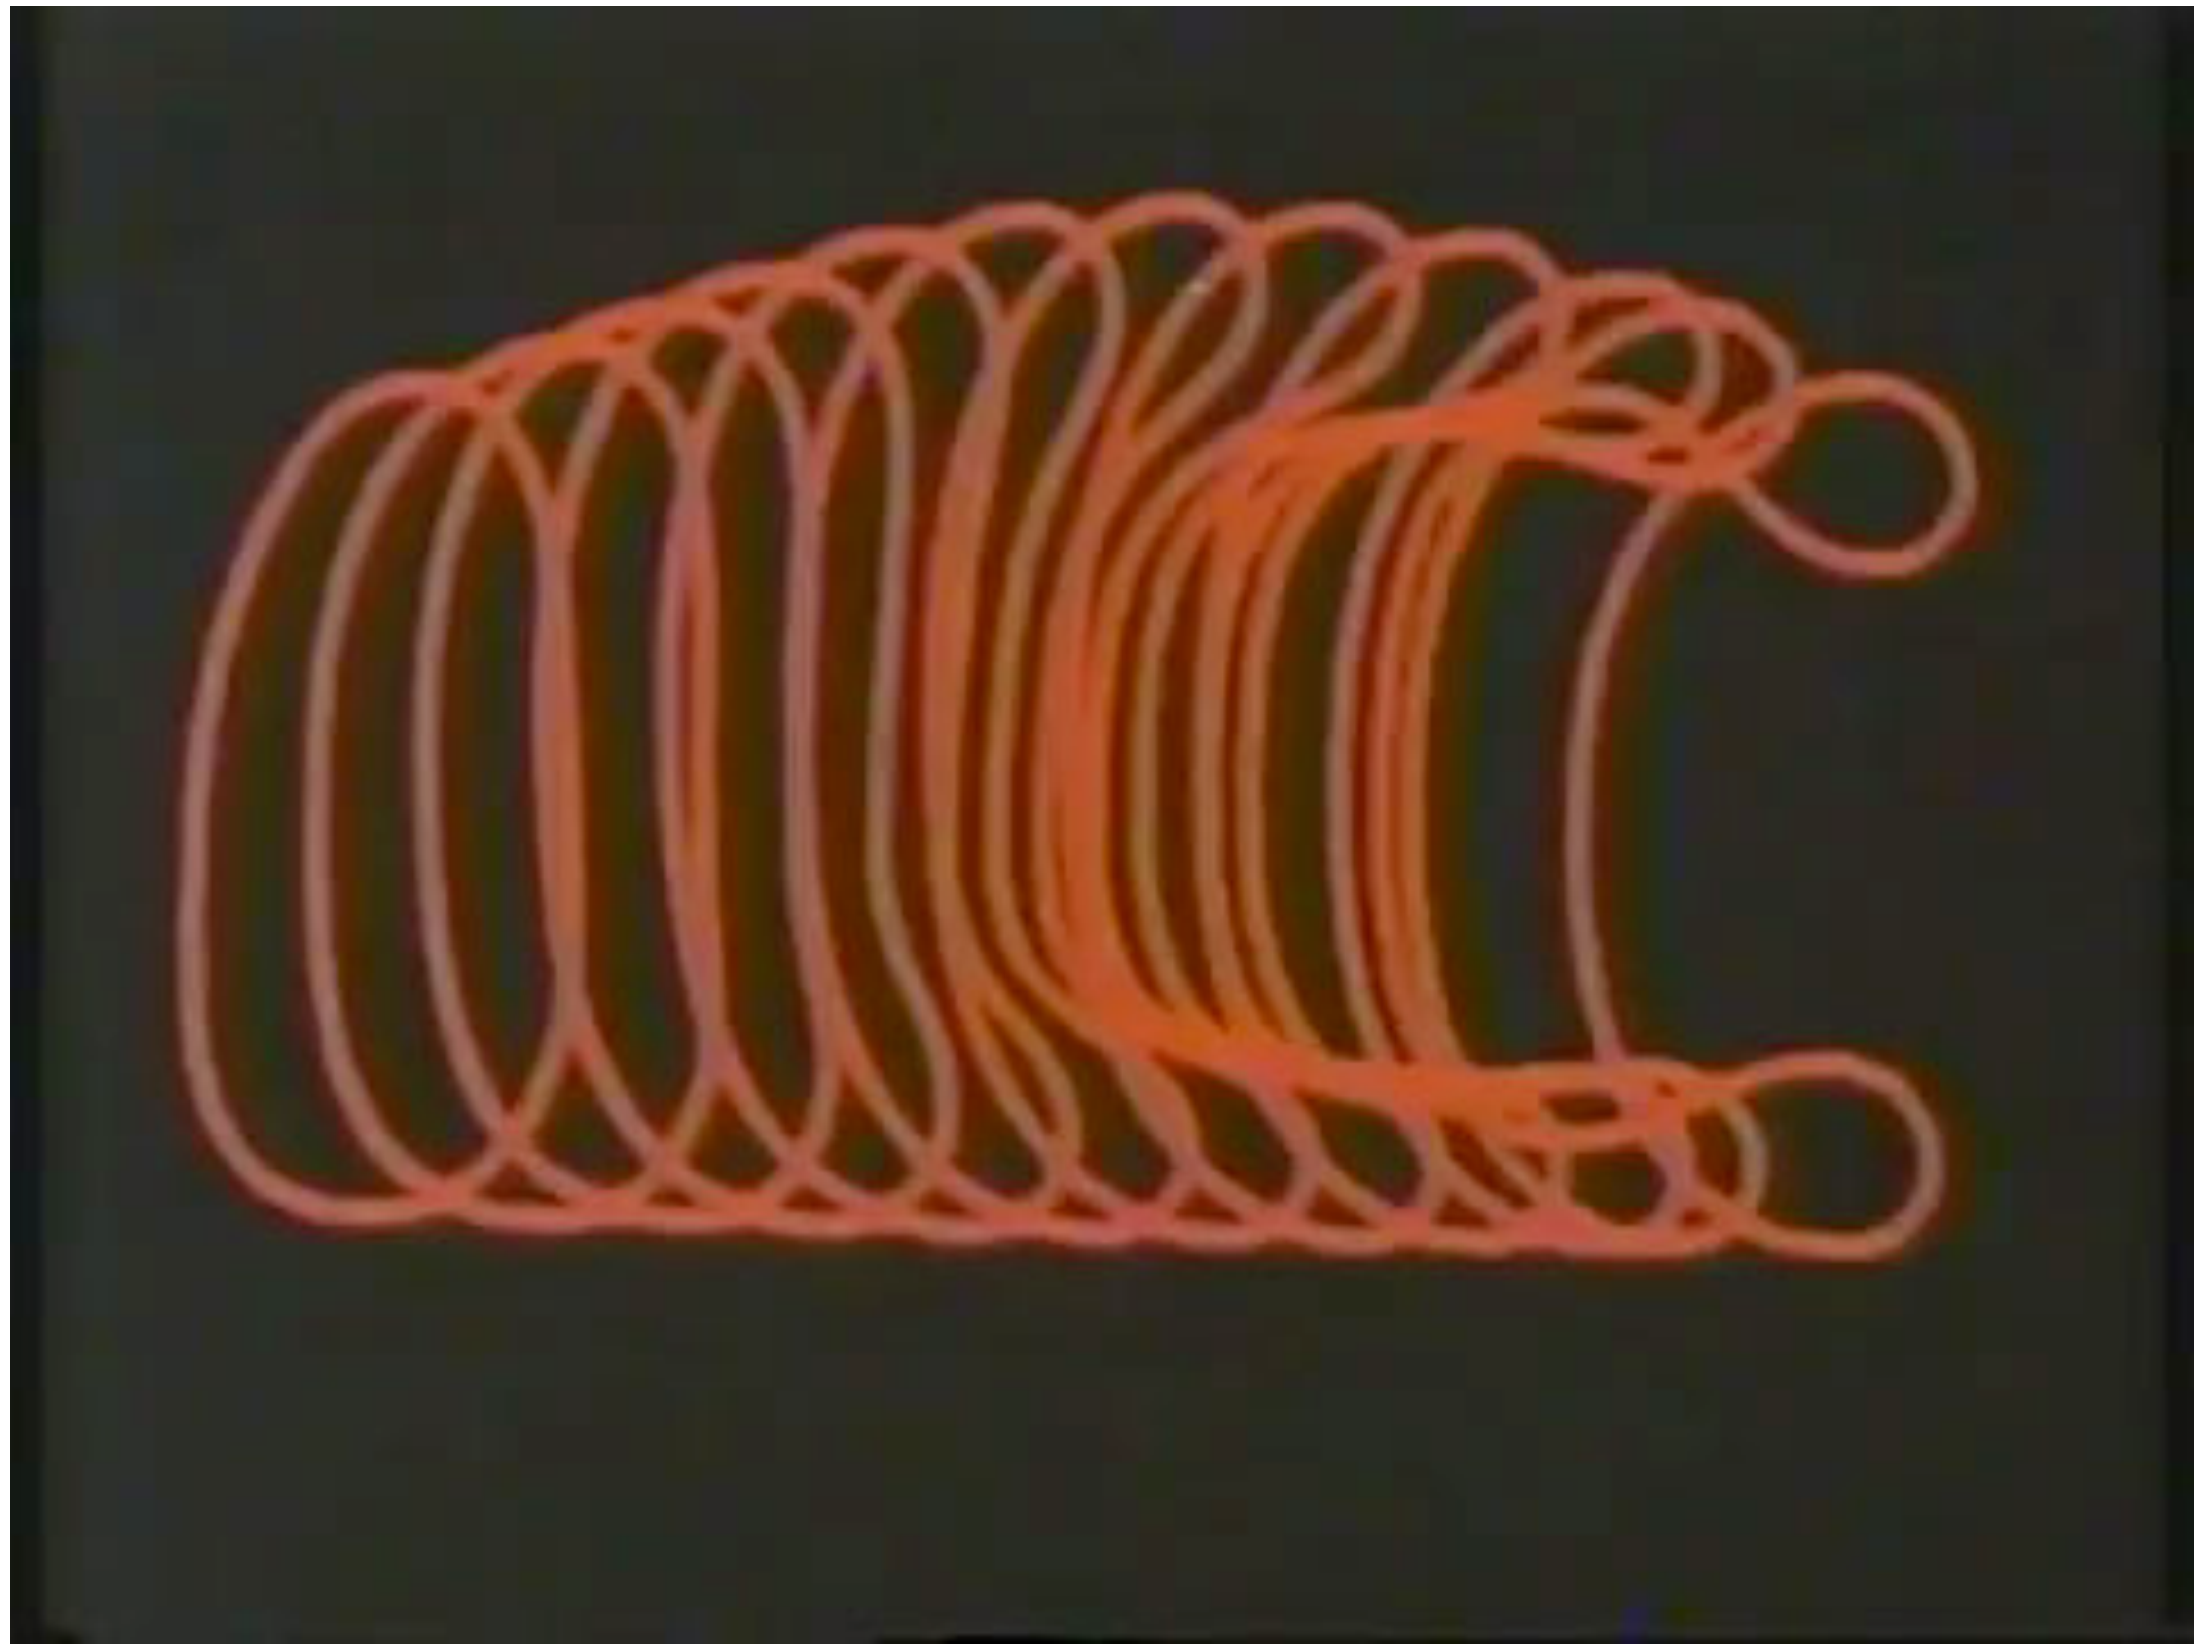

“Regular Homotopies in the Plane” — https://youtu.be/fKFH3c7b57s

Application: Generalized Winding Numbers

- For messy, "real world" data (instead of perfect closed curves) can still measure notion of how much a curve, surface, etc., "wraps around" a point
- Just sum up signed projected lengths (or areas)
- Fractional winding number gives good indication of which points are inside/outside
- Useful for a wide variety of practical tasks: extracting "watertight" mesh, tetrahedral meshing, constructive solid geometry (booleans), …

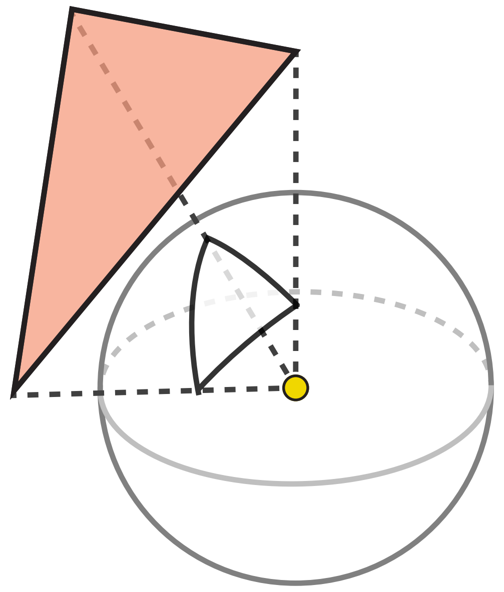

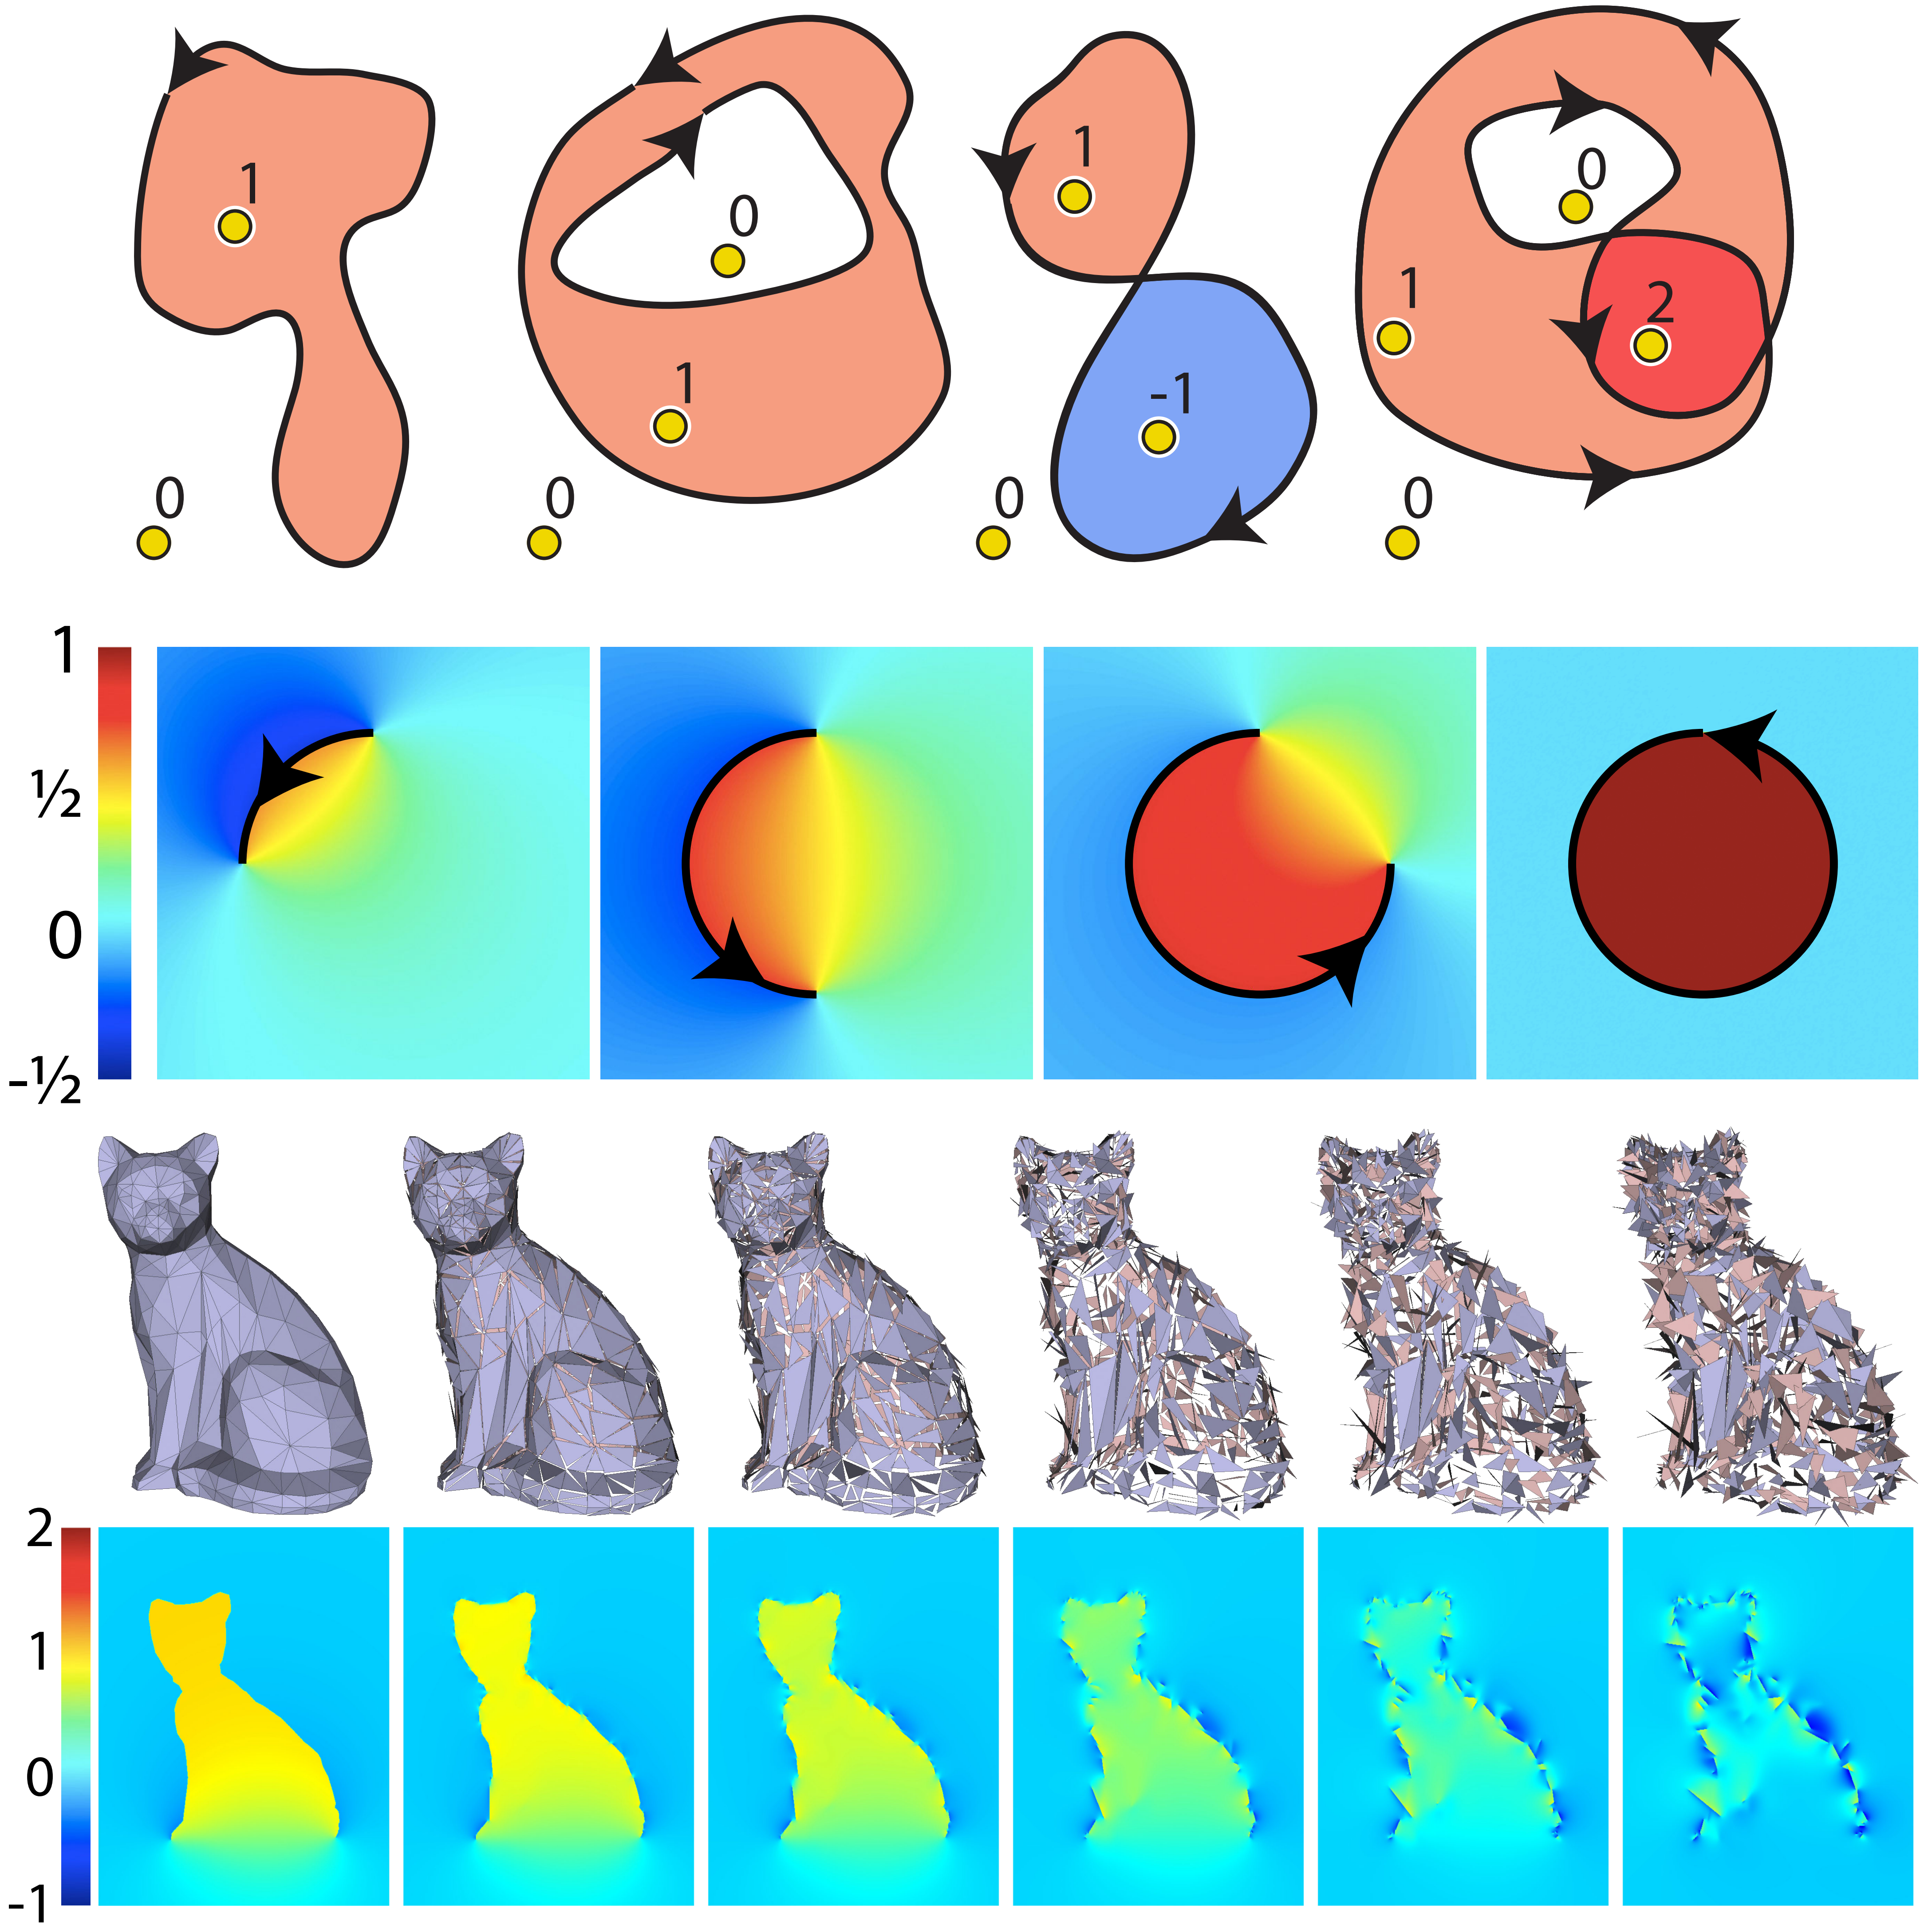

Jacobson et al, *"Robust Inside-Outside Segmentation using Generalized Winding Numbers"* (2013)

Space Curves

Parameterized Space Curve

- A **parameterized space curve** is a map* taking each point in an interval $[0,L]$ of the real line to some point in $\mathbb{R}^3$

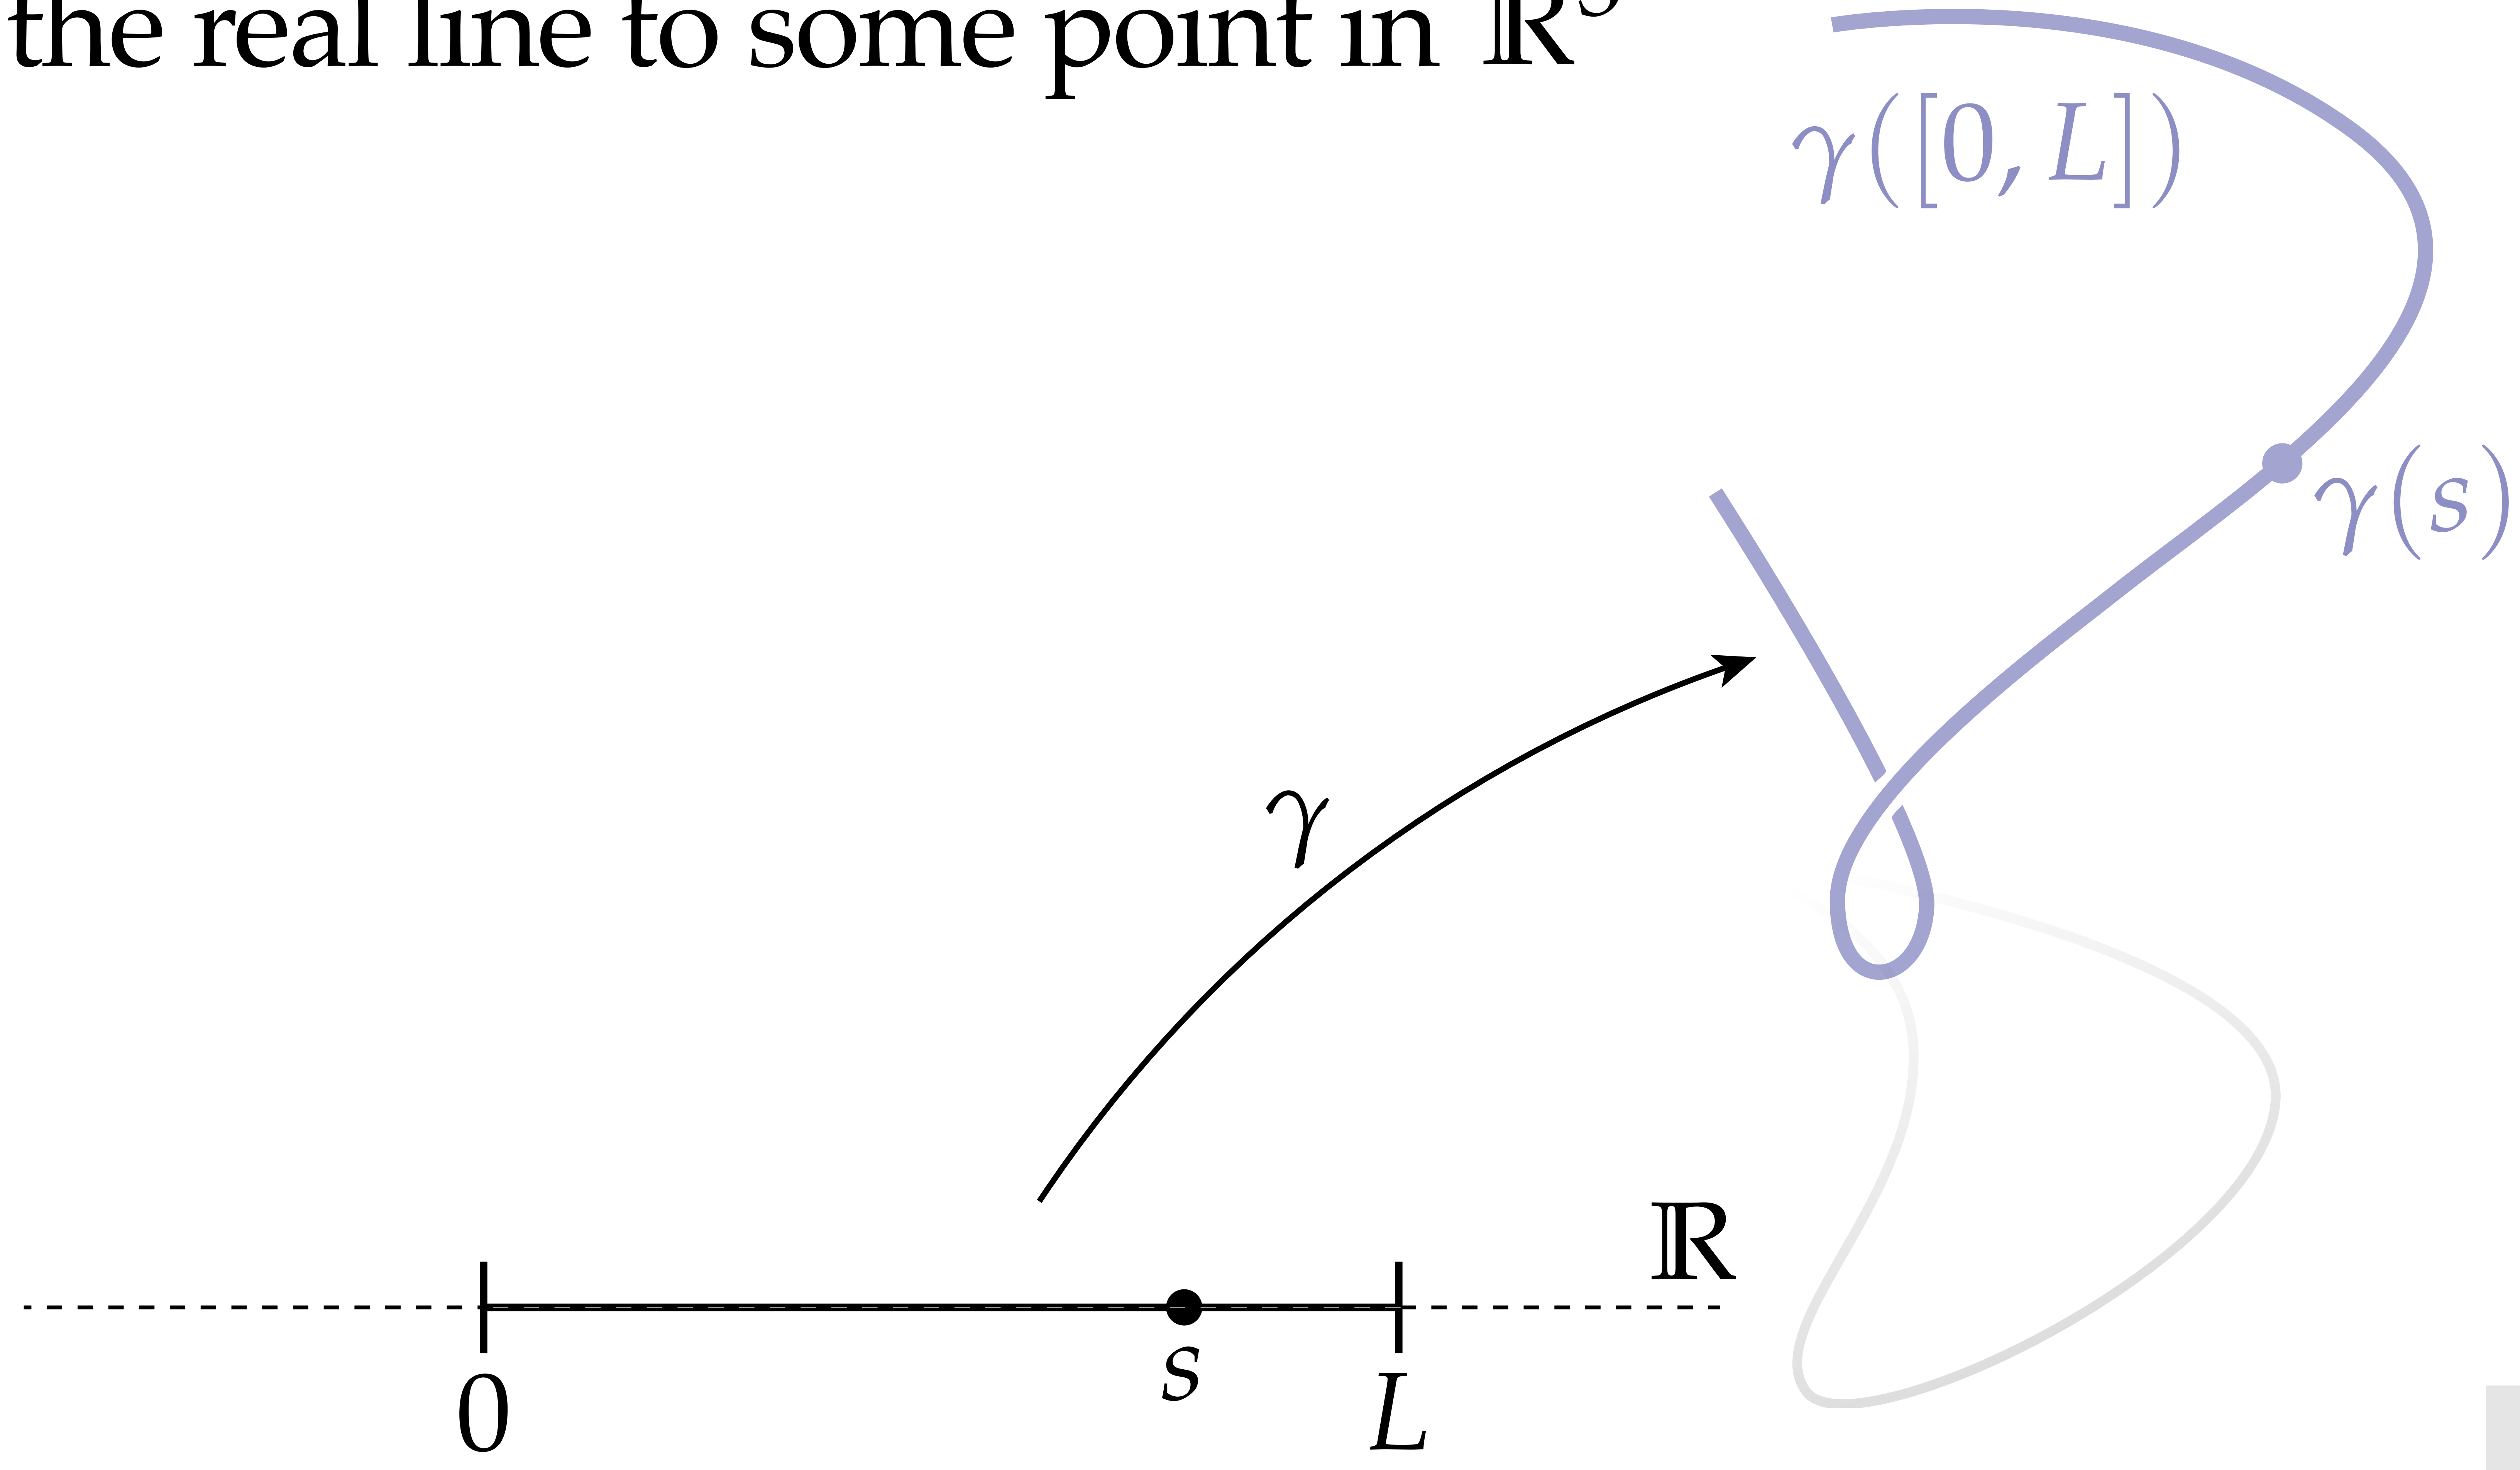

$\gamma : [0,L] \to \mathbb{R}^3$

*Continuous, differentiable, smooth…

Pushforward of Vectors on a Space Curve

Suppose we apply the differential of a parameterized space curve to a vector field X on its domain:

$$\gamma := (x, y, z), \;\; x, y, z : [0, L] \to \mathbb{R}$$

$$X := a\tfrac{\partial}{\partial s}, \;\; a : [0, L] \to \mathbb{R}$$

$$d\gamma = (\tfrac{\partial x}{\partial s}, \tfrac{\partial y}{\partial s}, \tfrac{\partial z}{\partial s})ds$$

$$d\gamma(X) = a(\tfrac{\partial x}{\partial s}, \tfrac{\partial y}{\partial s}, \tfrac{\partial z}{\partial s})$$

Q: What's the *geometric* meaning?

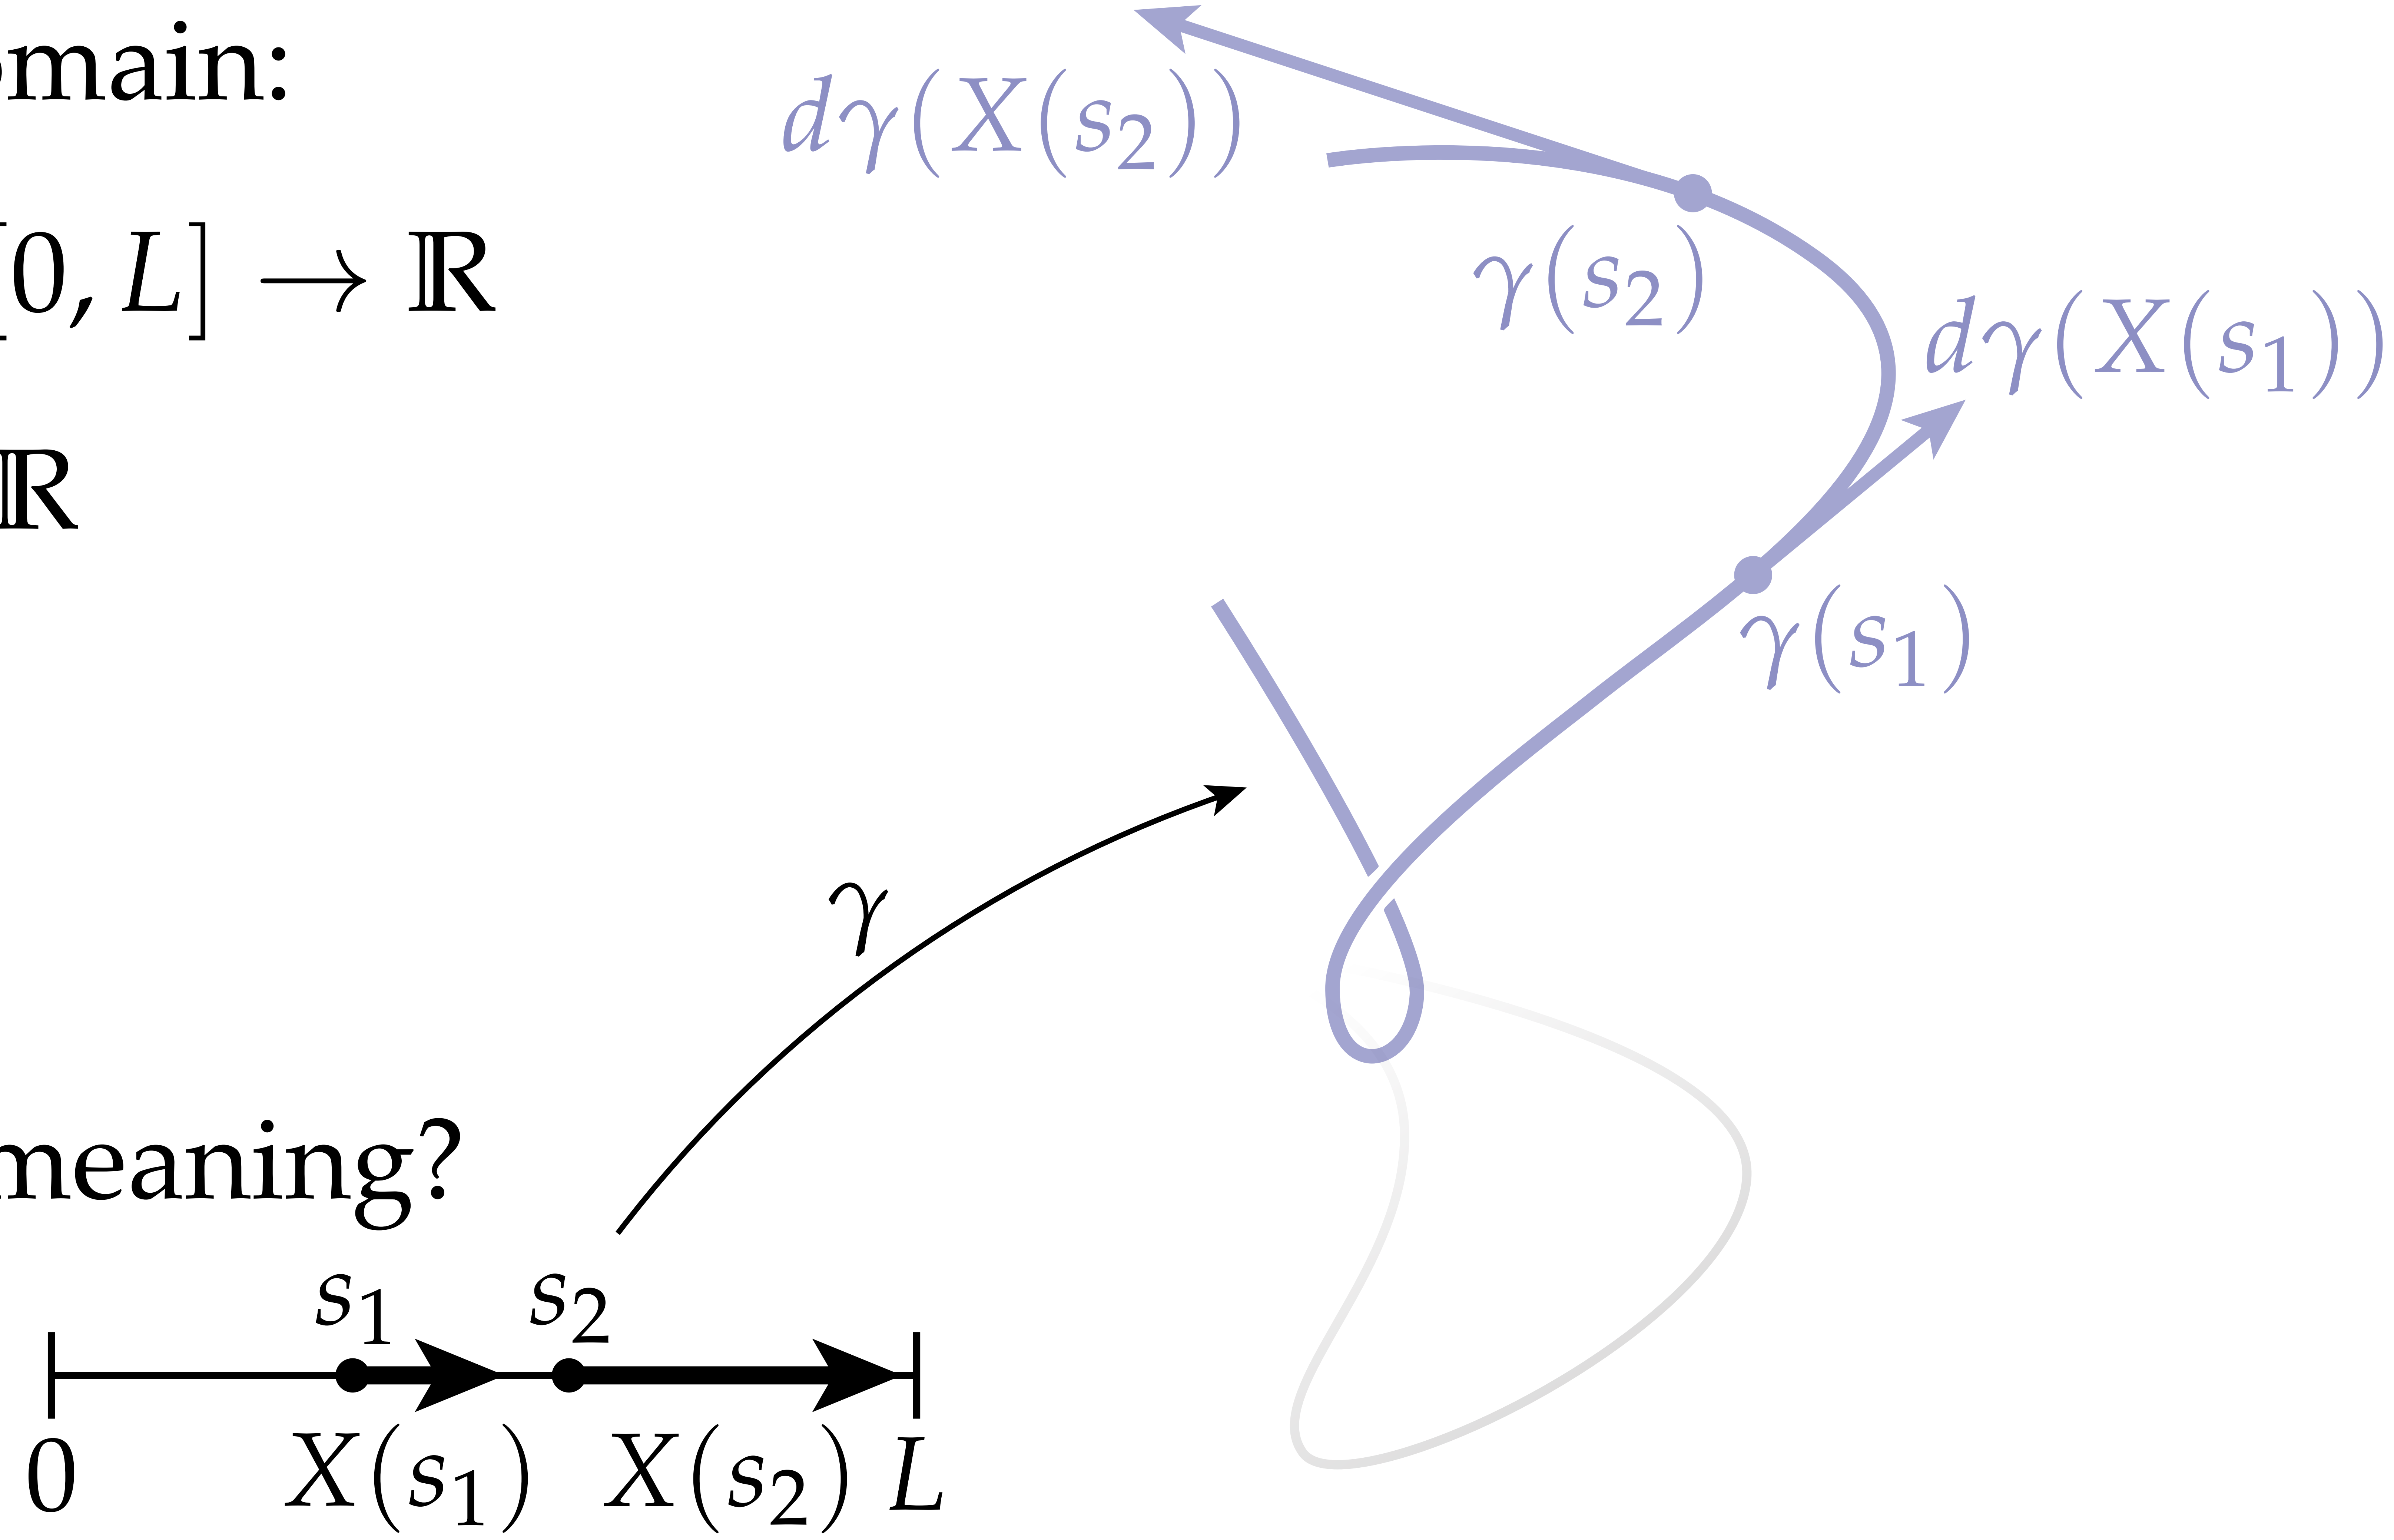

Parameterized Space Curve

- A **parameterized space curve** is a map* taking each point in an interval $[0,L]$ of the real line to some point in $\mathbb{R}^3$
- Its *differential* takes vectors on $\mathbb{R}$ to vectors in $\mathbb{R}^3$

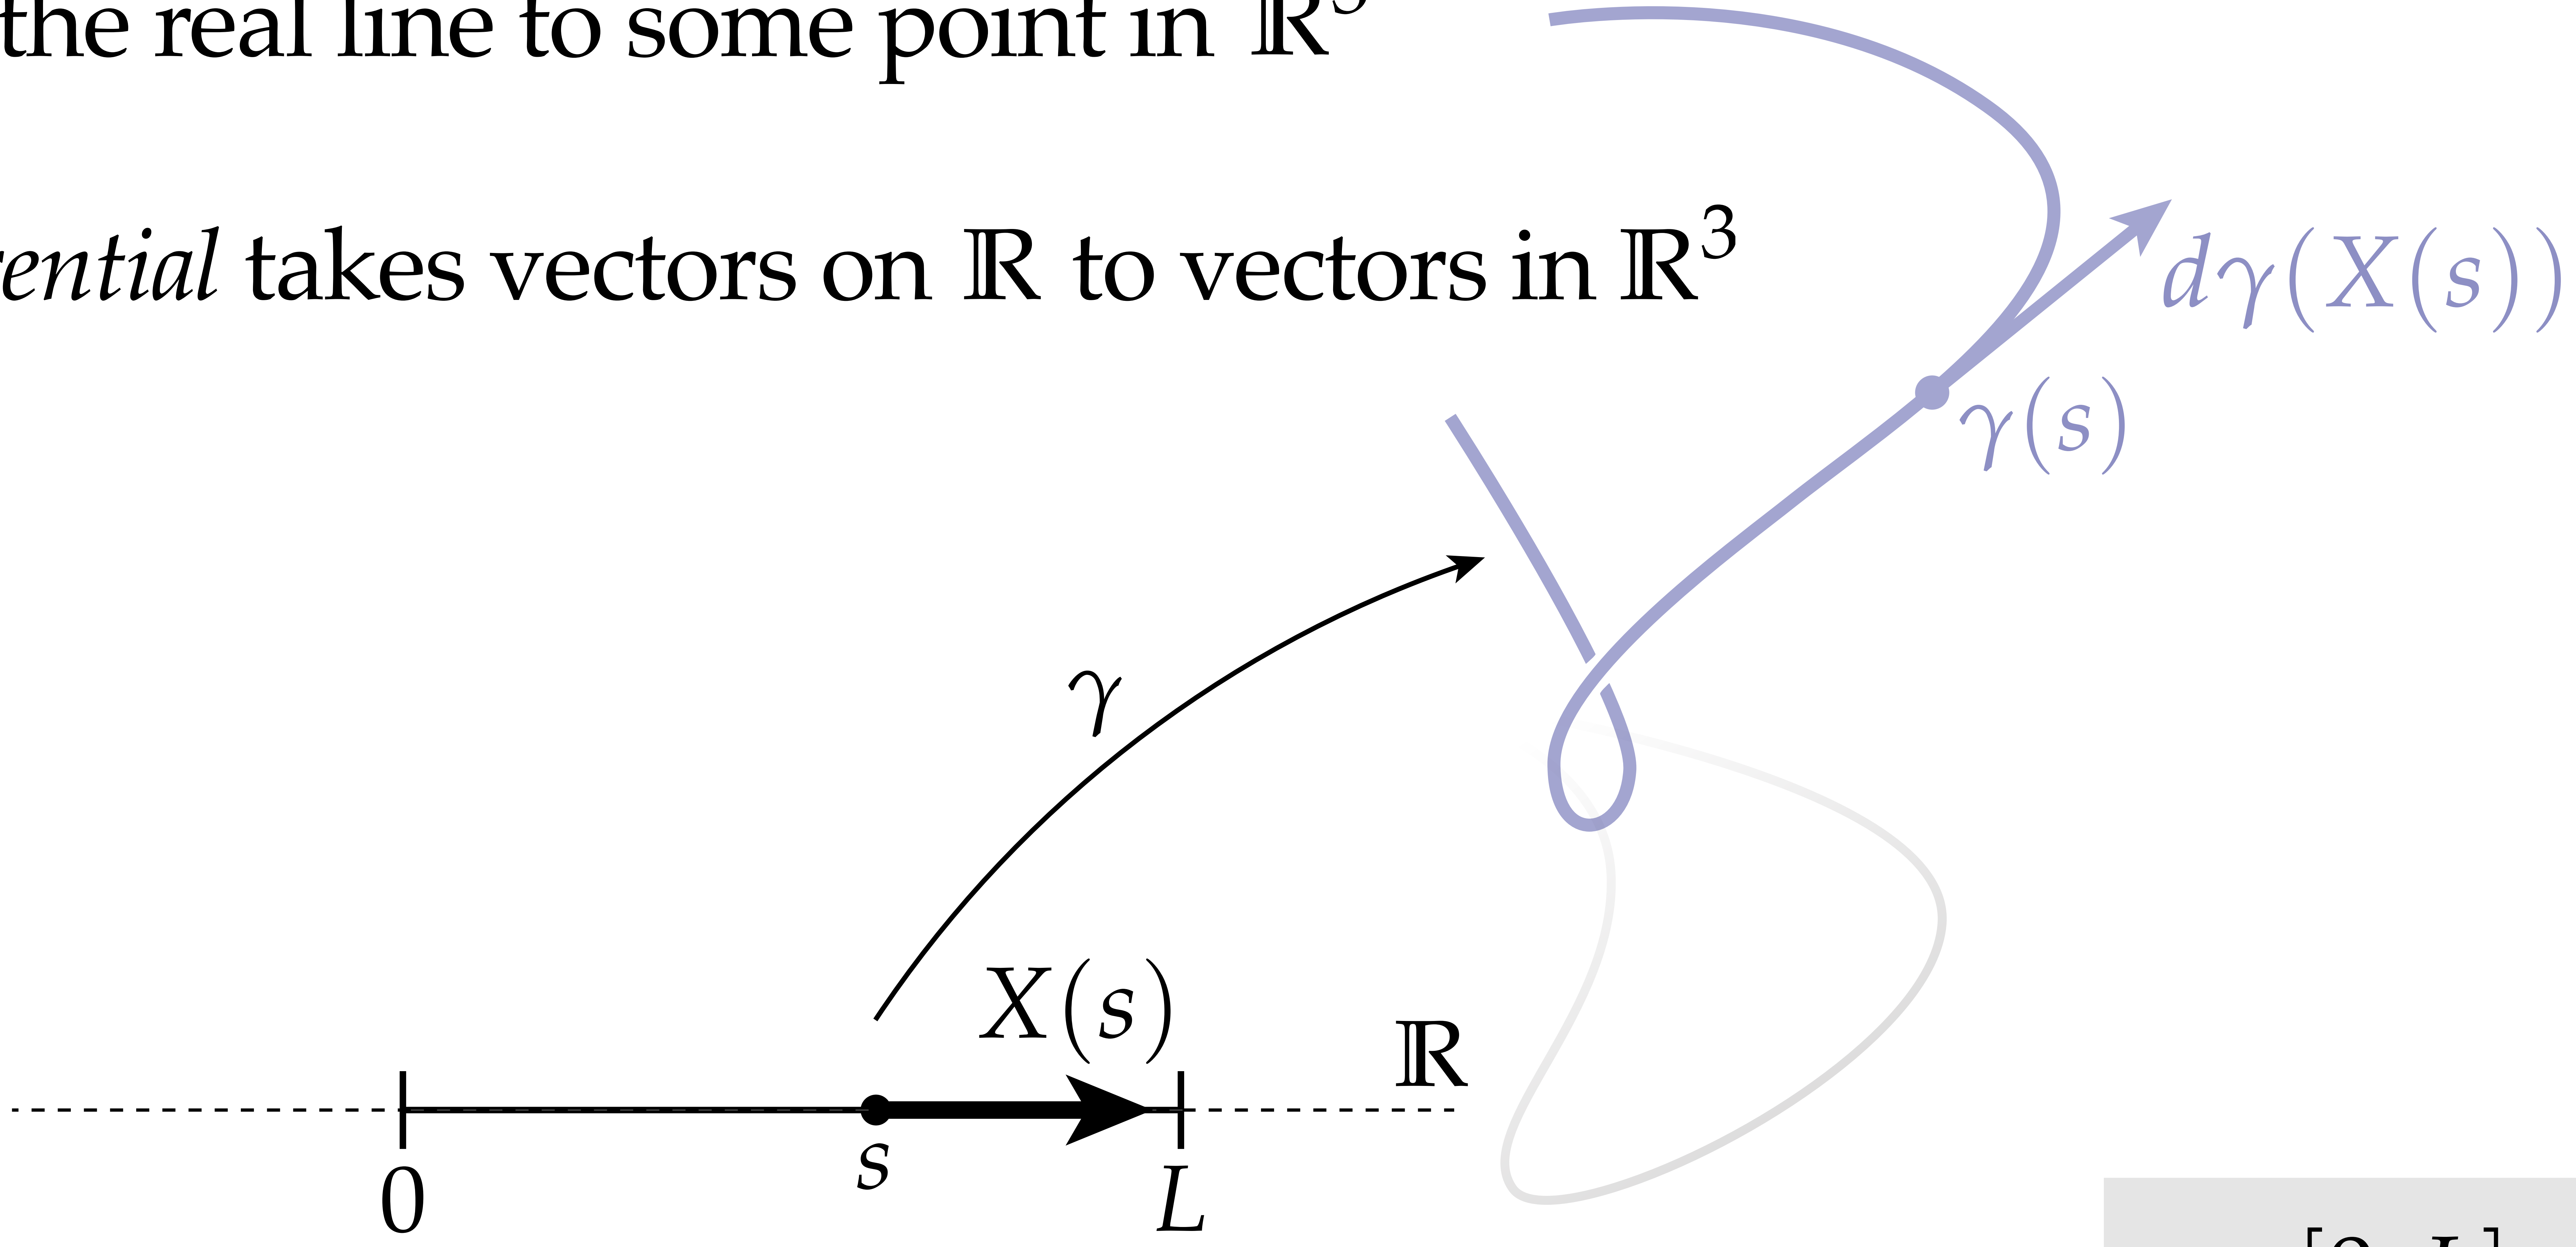

$\gamma : [0, L] \to \mathbb{R}^3$

*Continuous, differentiable, smooth…

Curvature and Torsion of a Space Curve

- For a plane curve, *curvature* captured the notion of "bending"
- For a space curve we also have *torsion*, which captures "twisting"

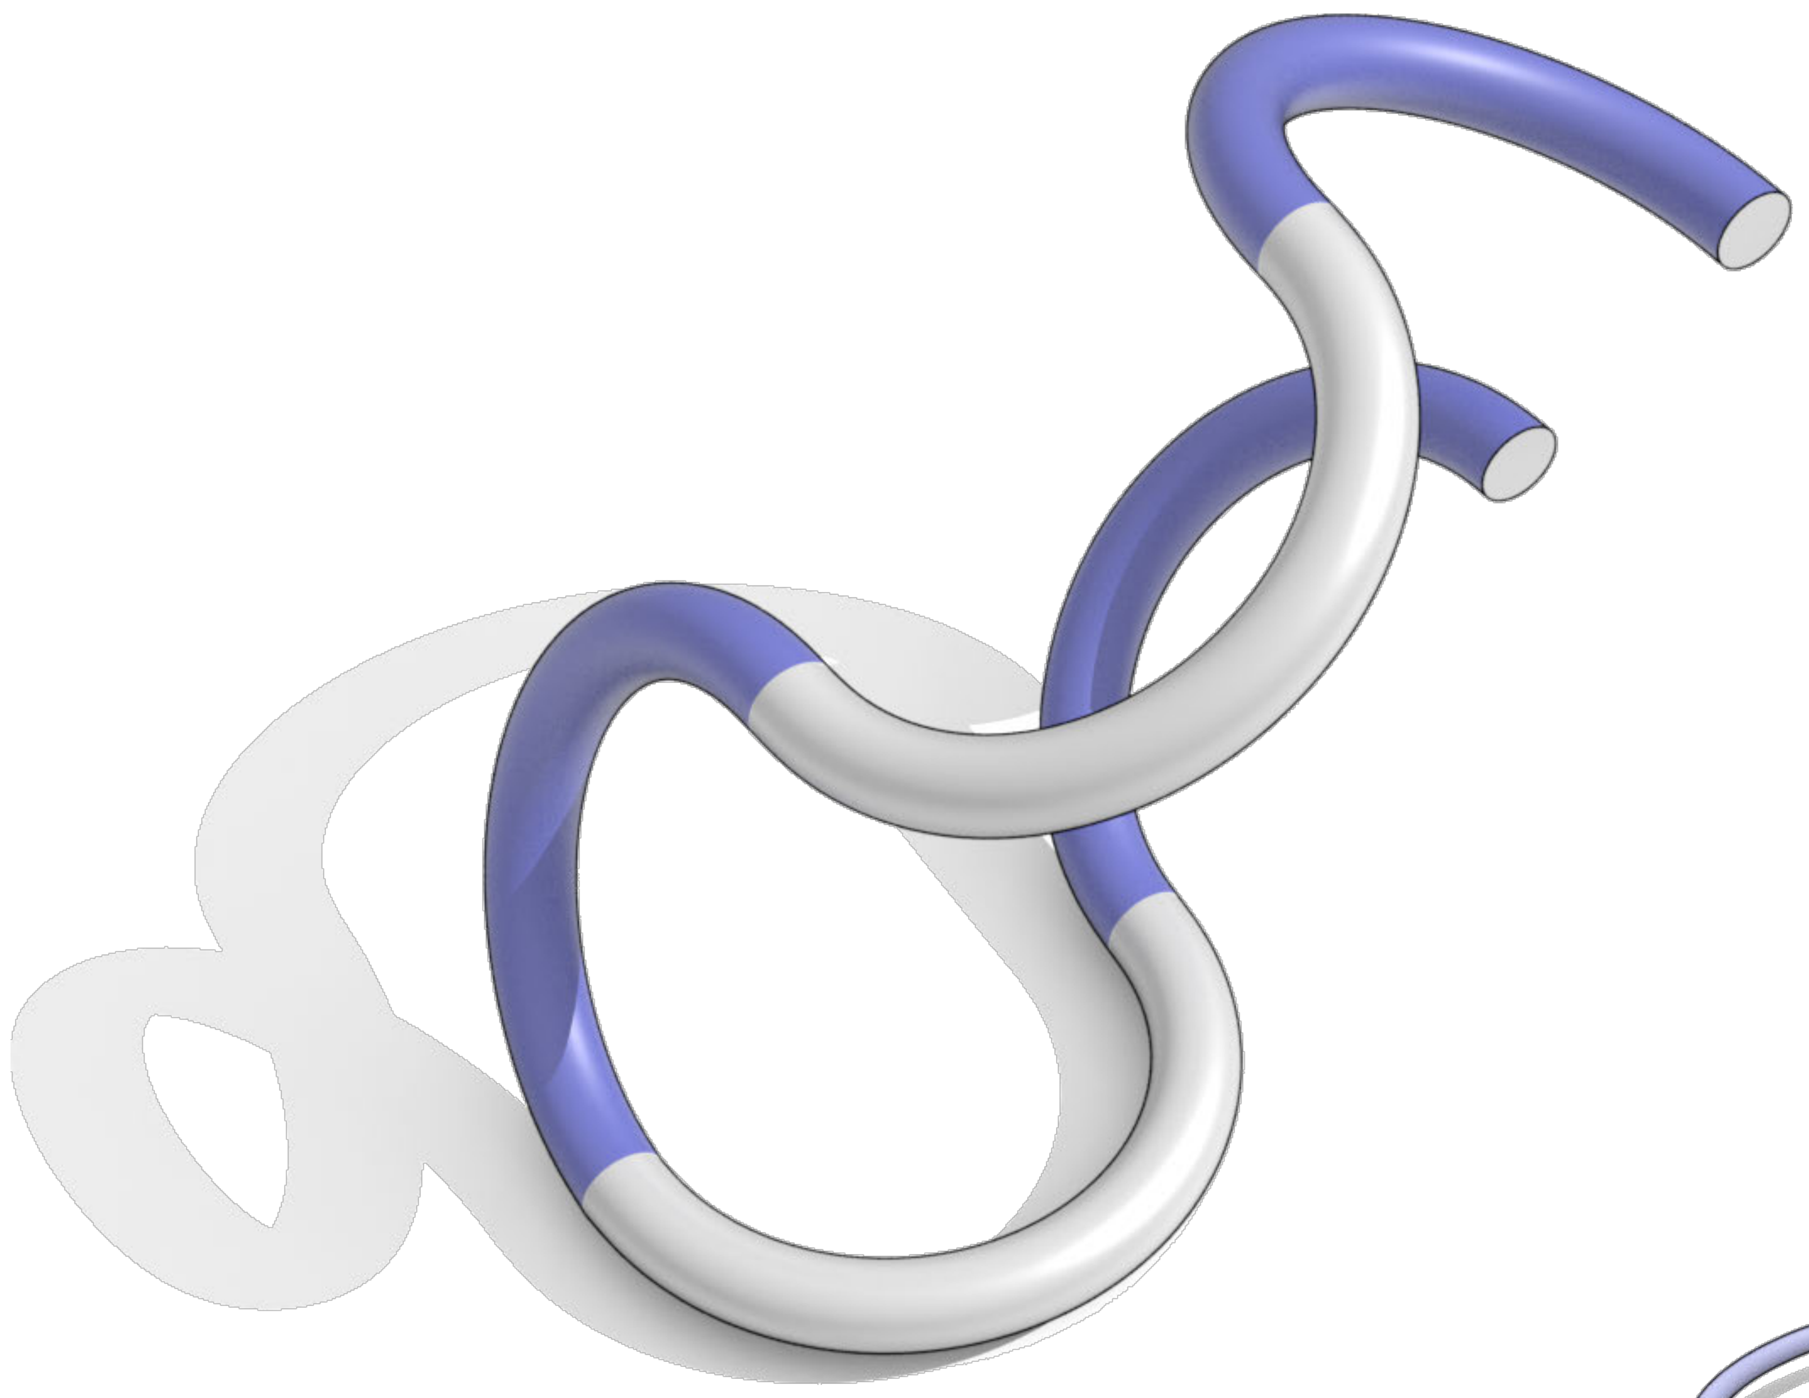

Intuition: torsion is "out of plane bending"

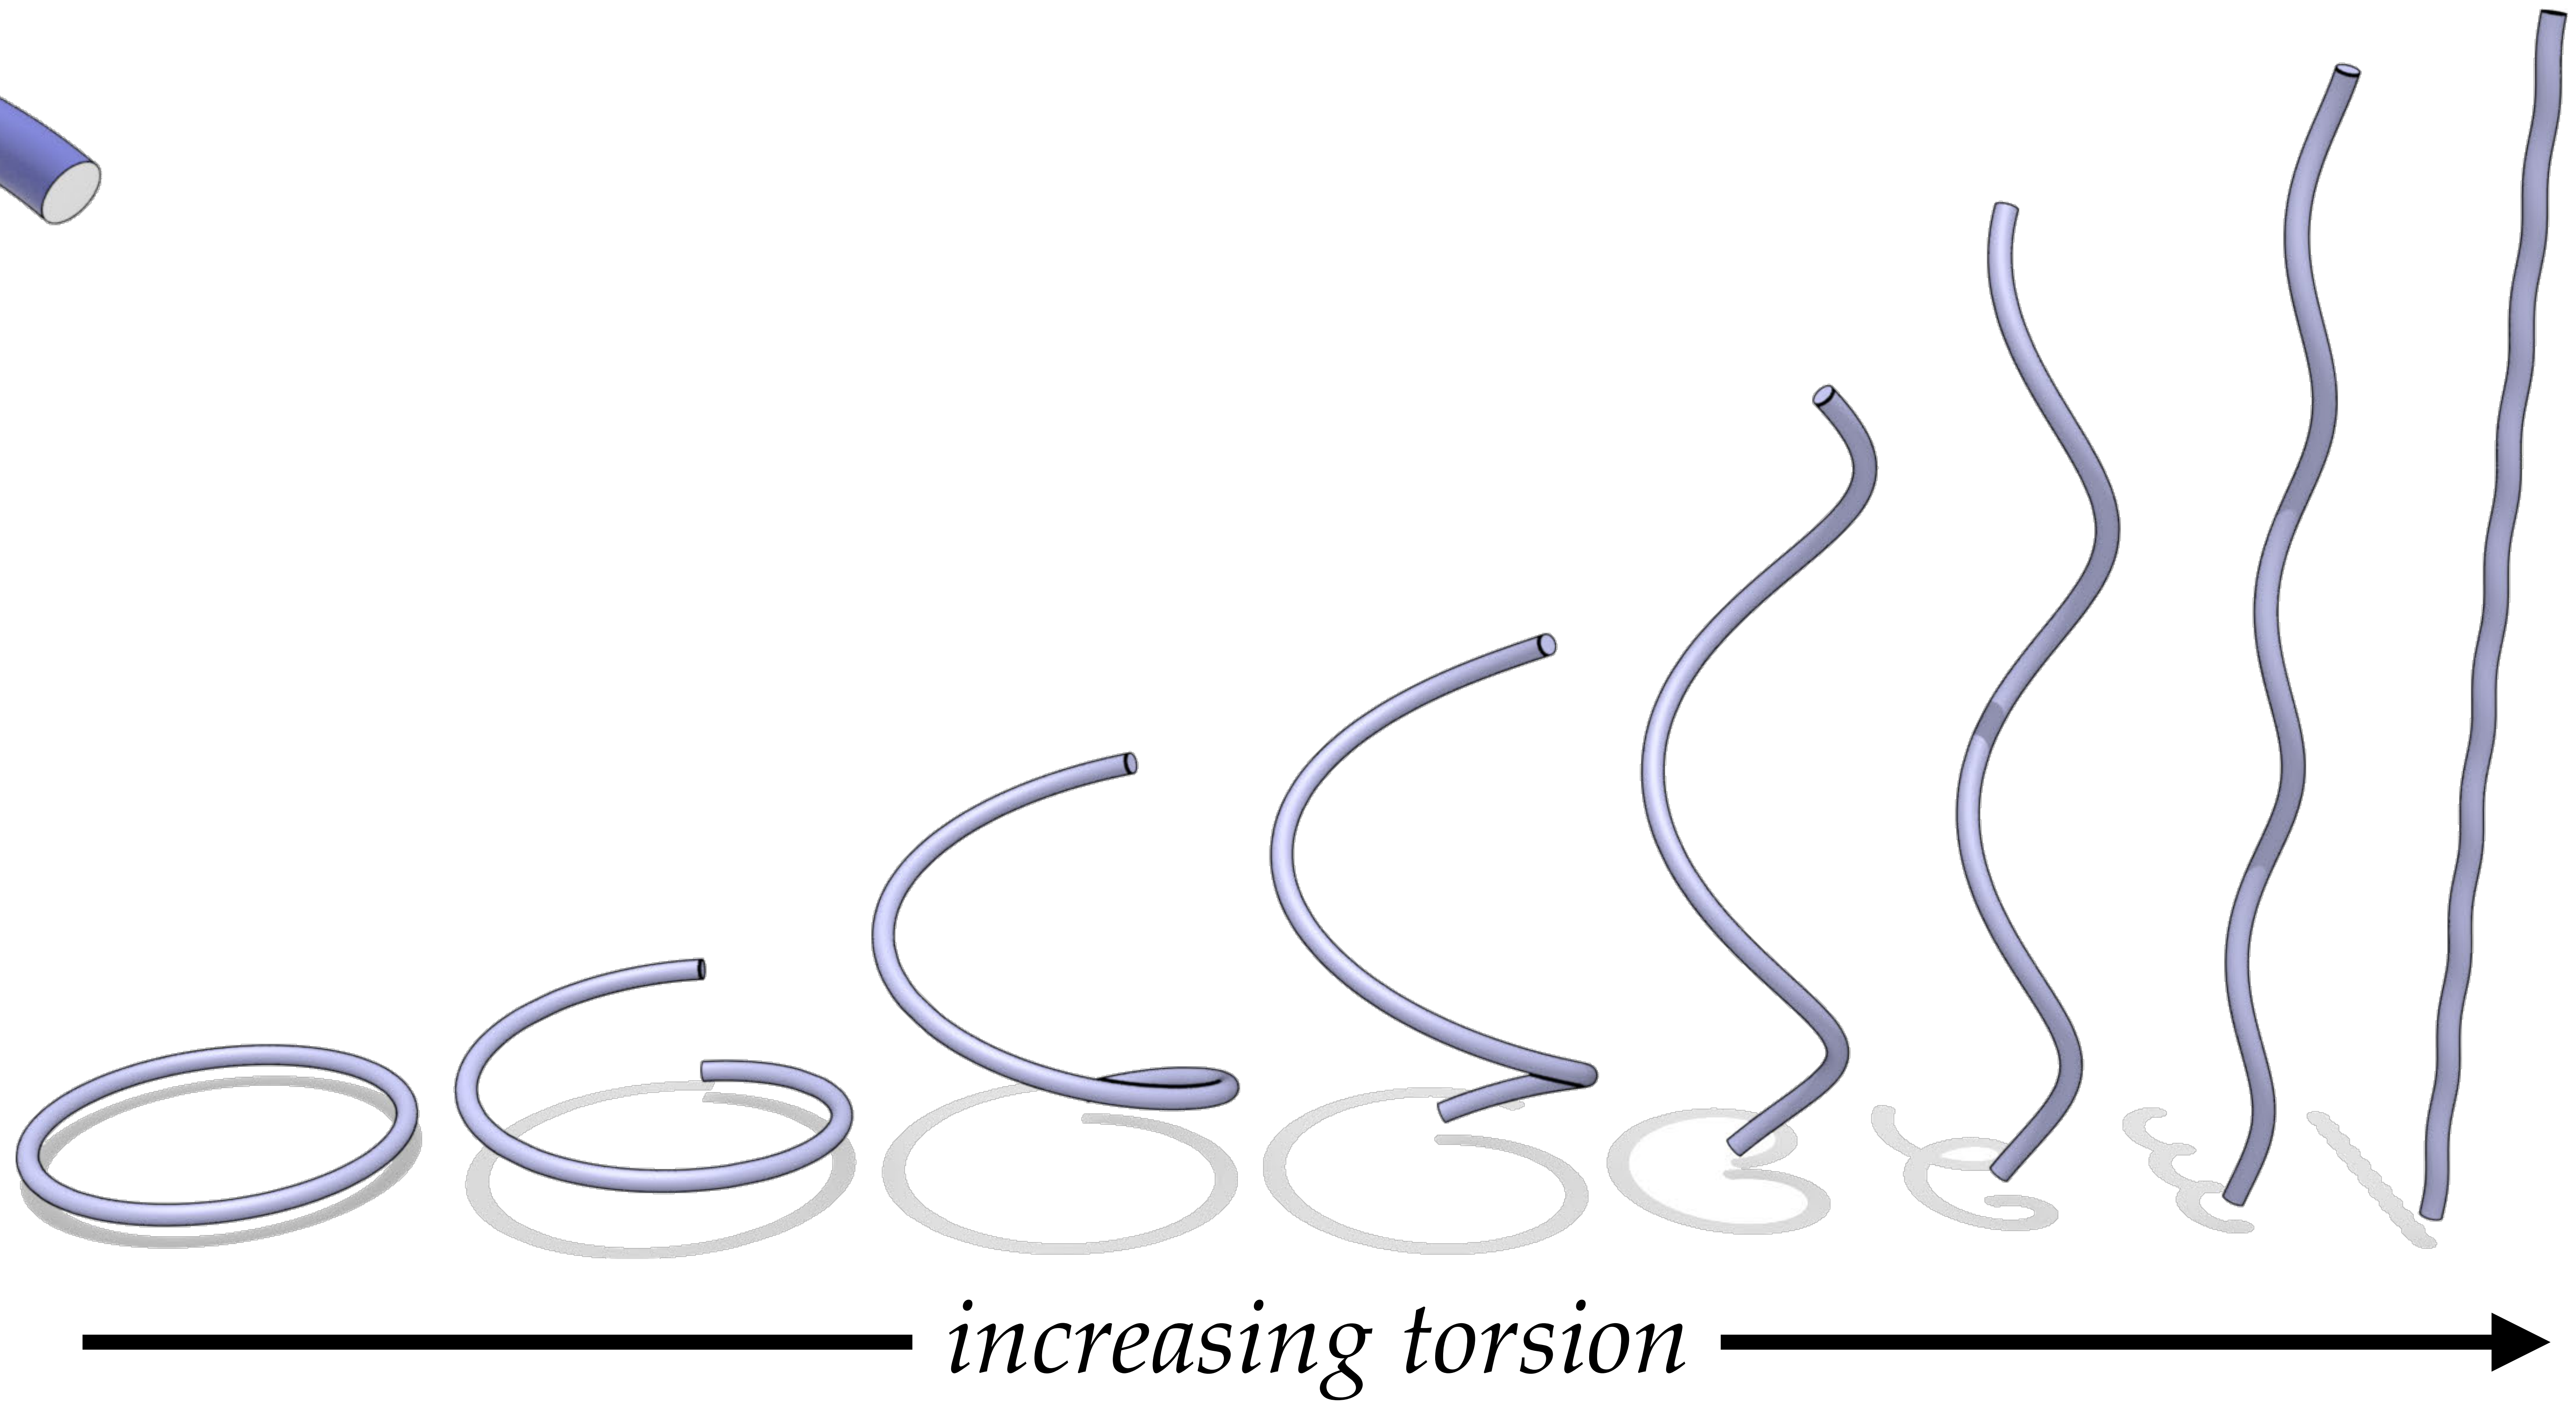

Frenet Frame

- Each point of a space curve has a natural coordinate frame called the *Frenet frame*, which depends only on the local geometry

- As in the plane, the tangent T is found by differentiating the curve, and differentiating the tangent yields the curvature times the normal N

- The binormal B then completes an orthonormal basis with T and N

$$T(s) := \tfrac{d}{ds}\gamma(s)$$
$$N(s) := \tfrac{d}{ds}T/|\tfrac{d}{ds}T|$$
$$B(s) := T(s) \times N(s)$$

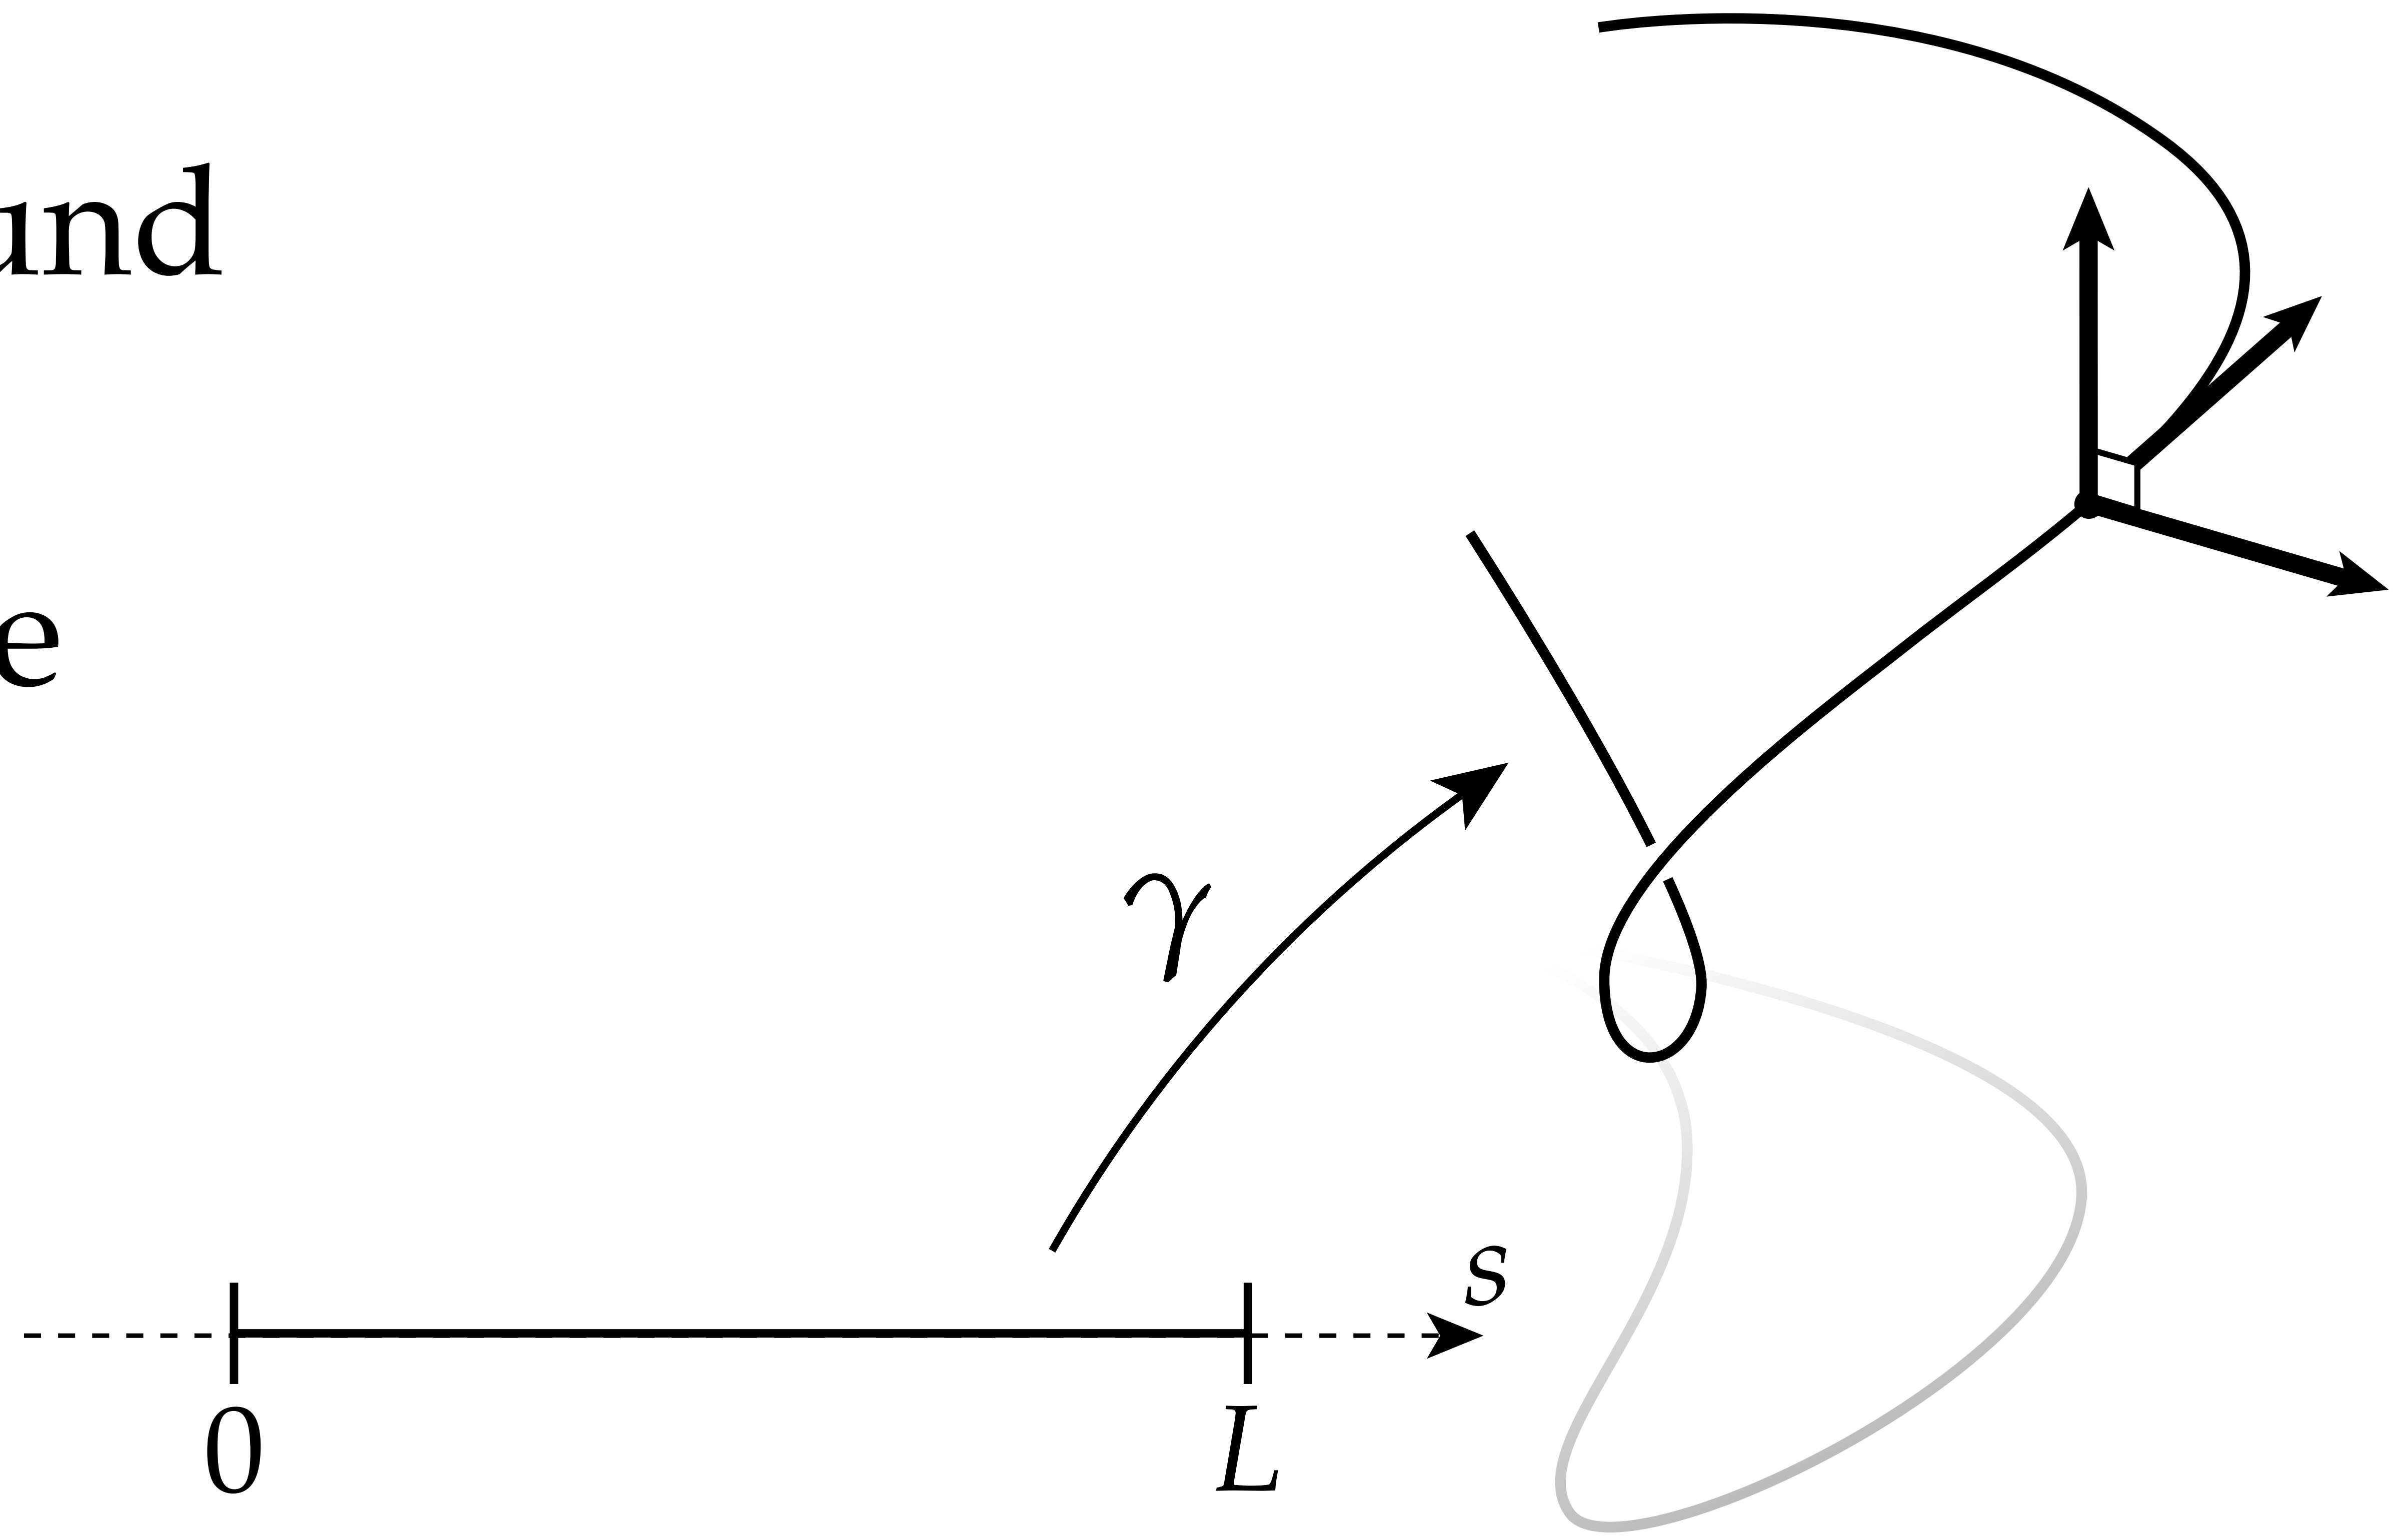

Frenet-Serret Equation

- Curvature κ and torsion τ can be defined in terms of the change in the Frenet frame as we move along the curve:

$$\frac{d}{ds}\begin{bmatrix} T \\ N \\ B \end{bmatrix} = \begin{bmatrix} 0 & -\kappa & 0 \\ \kappa & 0 & -\tau \\ 0 & \tau & 0 \end{bmatrix}\begin{bmatrix} T \\ N \\ B \end{bmatrix}$$

- Most importantly, change in the tangent describes bending (*curvature*); change in binormal describes twisting (*torsion*)

$$\kappa = -\langle N, \tfrac{d}{ds}T\rangle$$
$$\tau = \quad \langle N, \tfrac{d}{ds}B\rangle$$

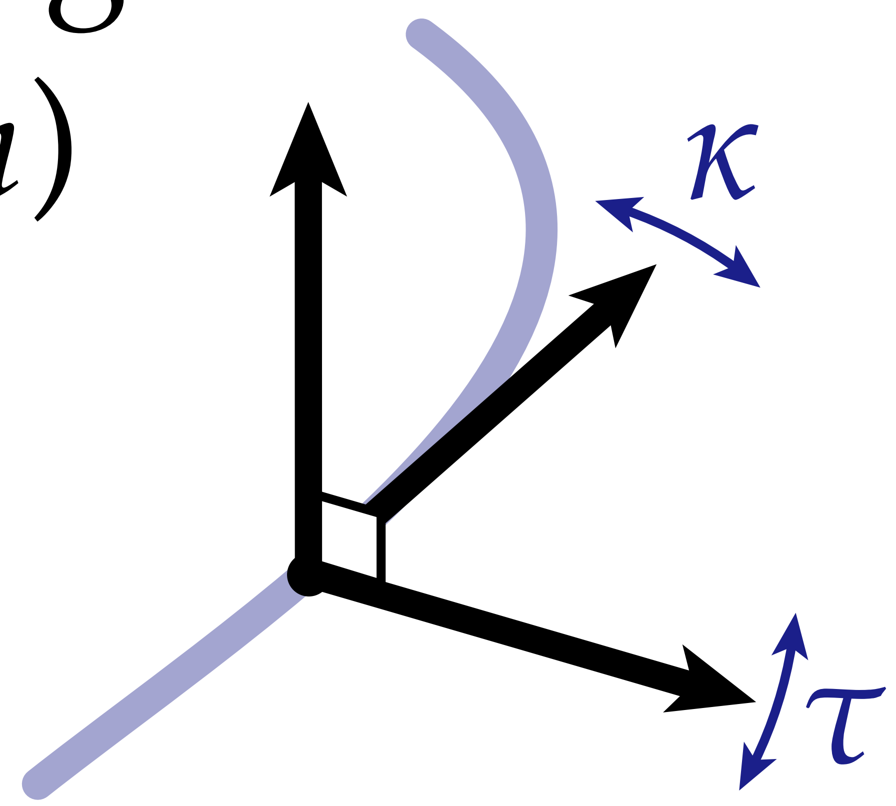

Example — Helix

- Let's compute the Frenet frame, curvature, and torsion for a *helix**

$$\gamma(s) := (a\cos(s), a\sin(s), bs)$$

$$\frac{d}{ds}\gamma(s) = (-a\sin(s), a\cos(s), b)$$

$$\left|\frac{d}{ds}\gamma\right| = \sqrt{a^2+b^2} = 1$$

$$\Rightarrow T(s) = \frac{d}{ds}\gamma(s)$$

$$\frac{d}{ds}T(s) = -a(\cos(s), \sin(s), 0)$$

$$\Rightarrow \kappa(s) = -a,\ N(s) = (\cos(s), \sin(s), 0)$$

$$B(s) = T(s) \times N(s) = (-b\sin(s), b\cos(s), -a)$$

$$\frac{d}{ds}B(s) = -b(\cos(s), \sin(s), 0)$$

$$\Rightarrow \tau(s) = -b$$

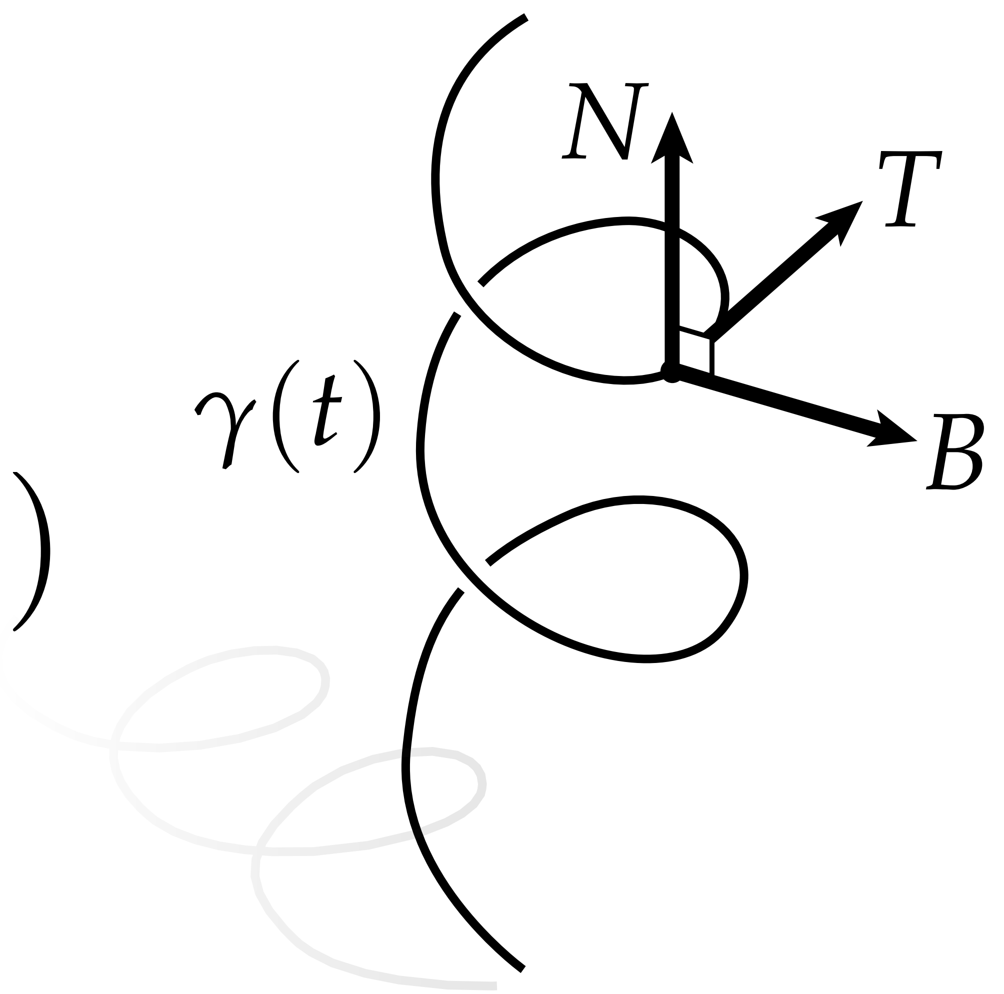

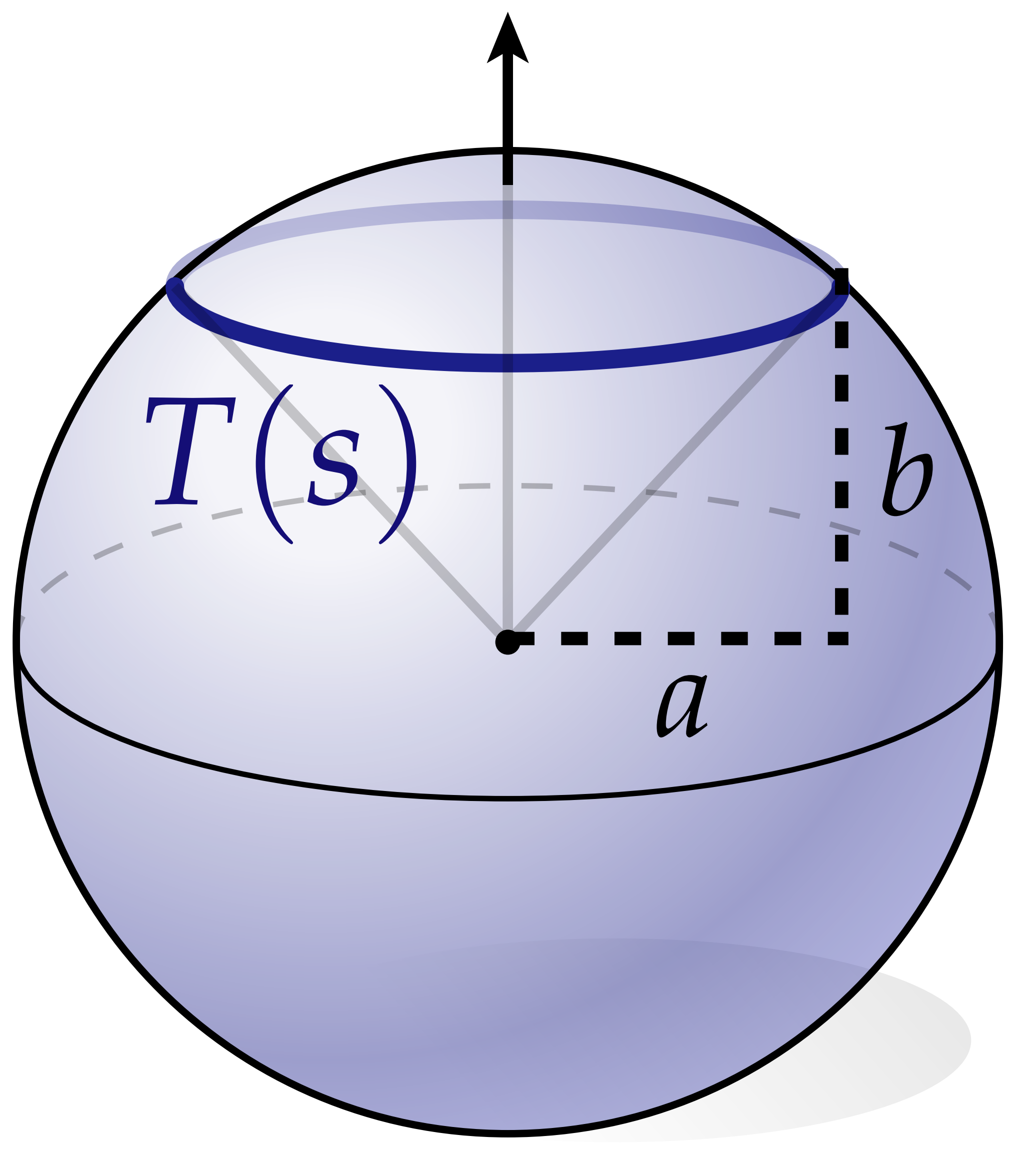

*For simplicity, let's pick a,b such that $a^2 + b^2 = 1$.

Fundamental Theorem of Space Curves

- The *fundamental theorem of space curves* tells us we can also go the other way: given the curvature and torsion of an arc-length parameterized space curve, we can recover the curve itself
- In 2D we just had to integrate a single ODE; here we integrate a system of three ODEs—namely, Frenet-Serret!

$$\frac{d}{ds}\begin{bmatrix} T \\ N \\ B \end{bmatrix} = \begin{bmatrix} 0 & -\kappa & 0 \\ \kappa & 0 & -\tau \\ 0 & \tau & 0 \end{bmatrix}\begin{bmatrix} T \\ N \\ B \end{bmatrix}$$

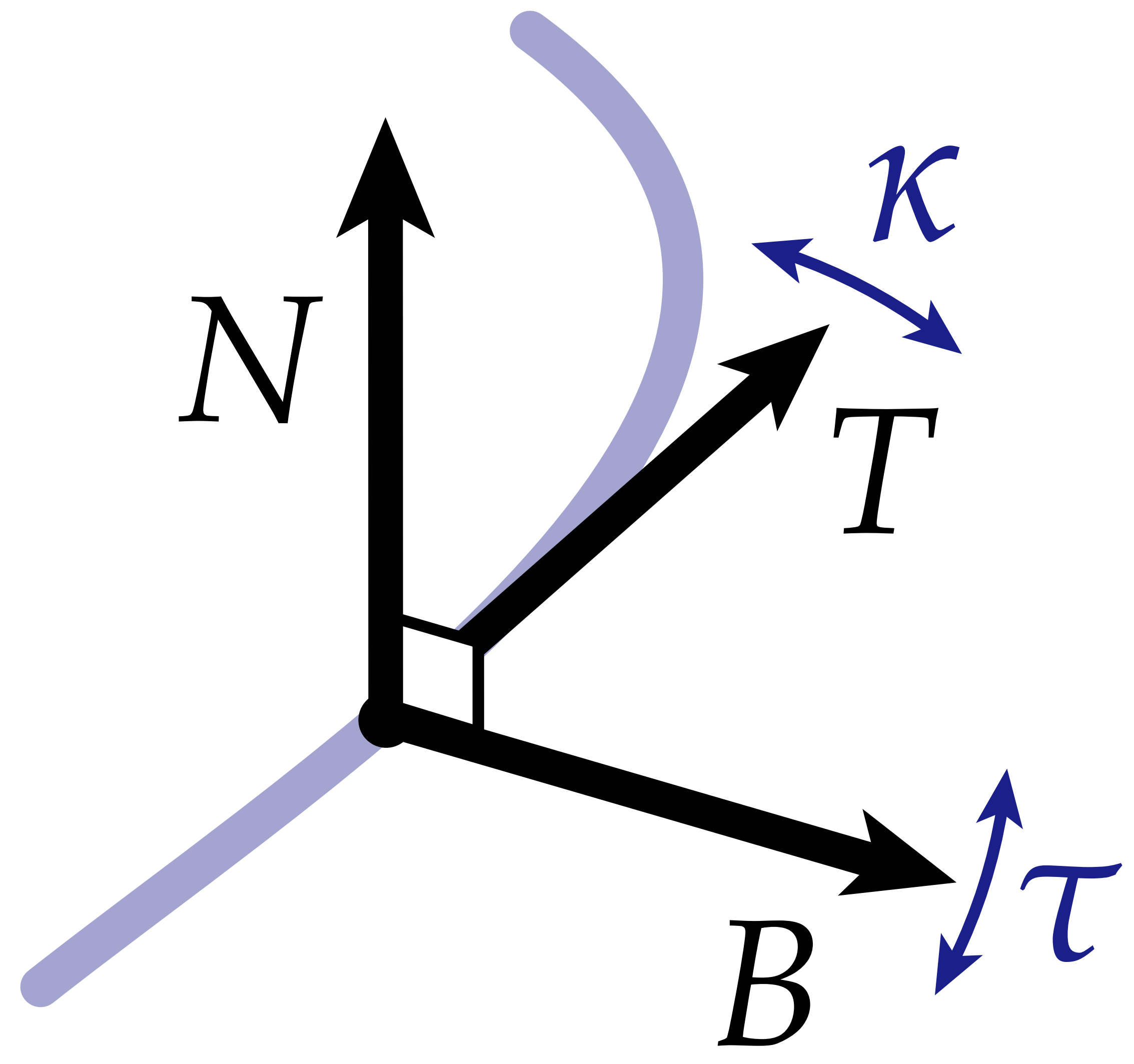

Adapted Frames on Curves

- **Q:** If our curve has a straight piece, is the Frenet frame well-defined?
- **A:** No, we don't have a clear normal/binormal (since, *e.g.*, $dT/ds = 0$)
- However, there are many ways to choose an *adapted frame*
- Any orthonormal frame including T
- *E.g.*, *least-twisting* frame (Bishop)
  - Unlike Frenet, *global* rather than *local*
- First example of *moving frames*
- (Will see more later for surfaces…)

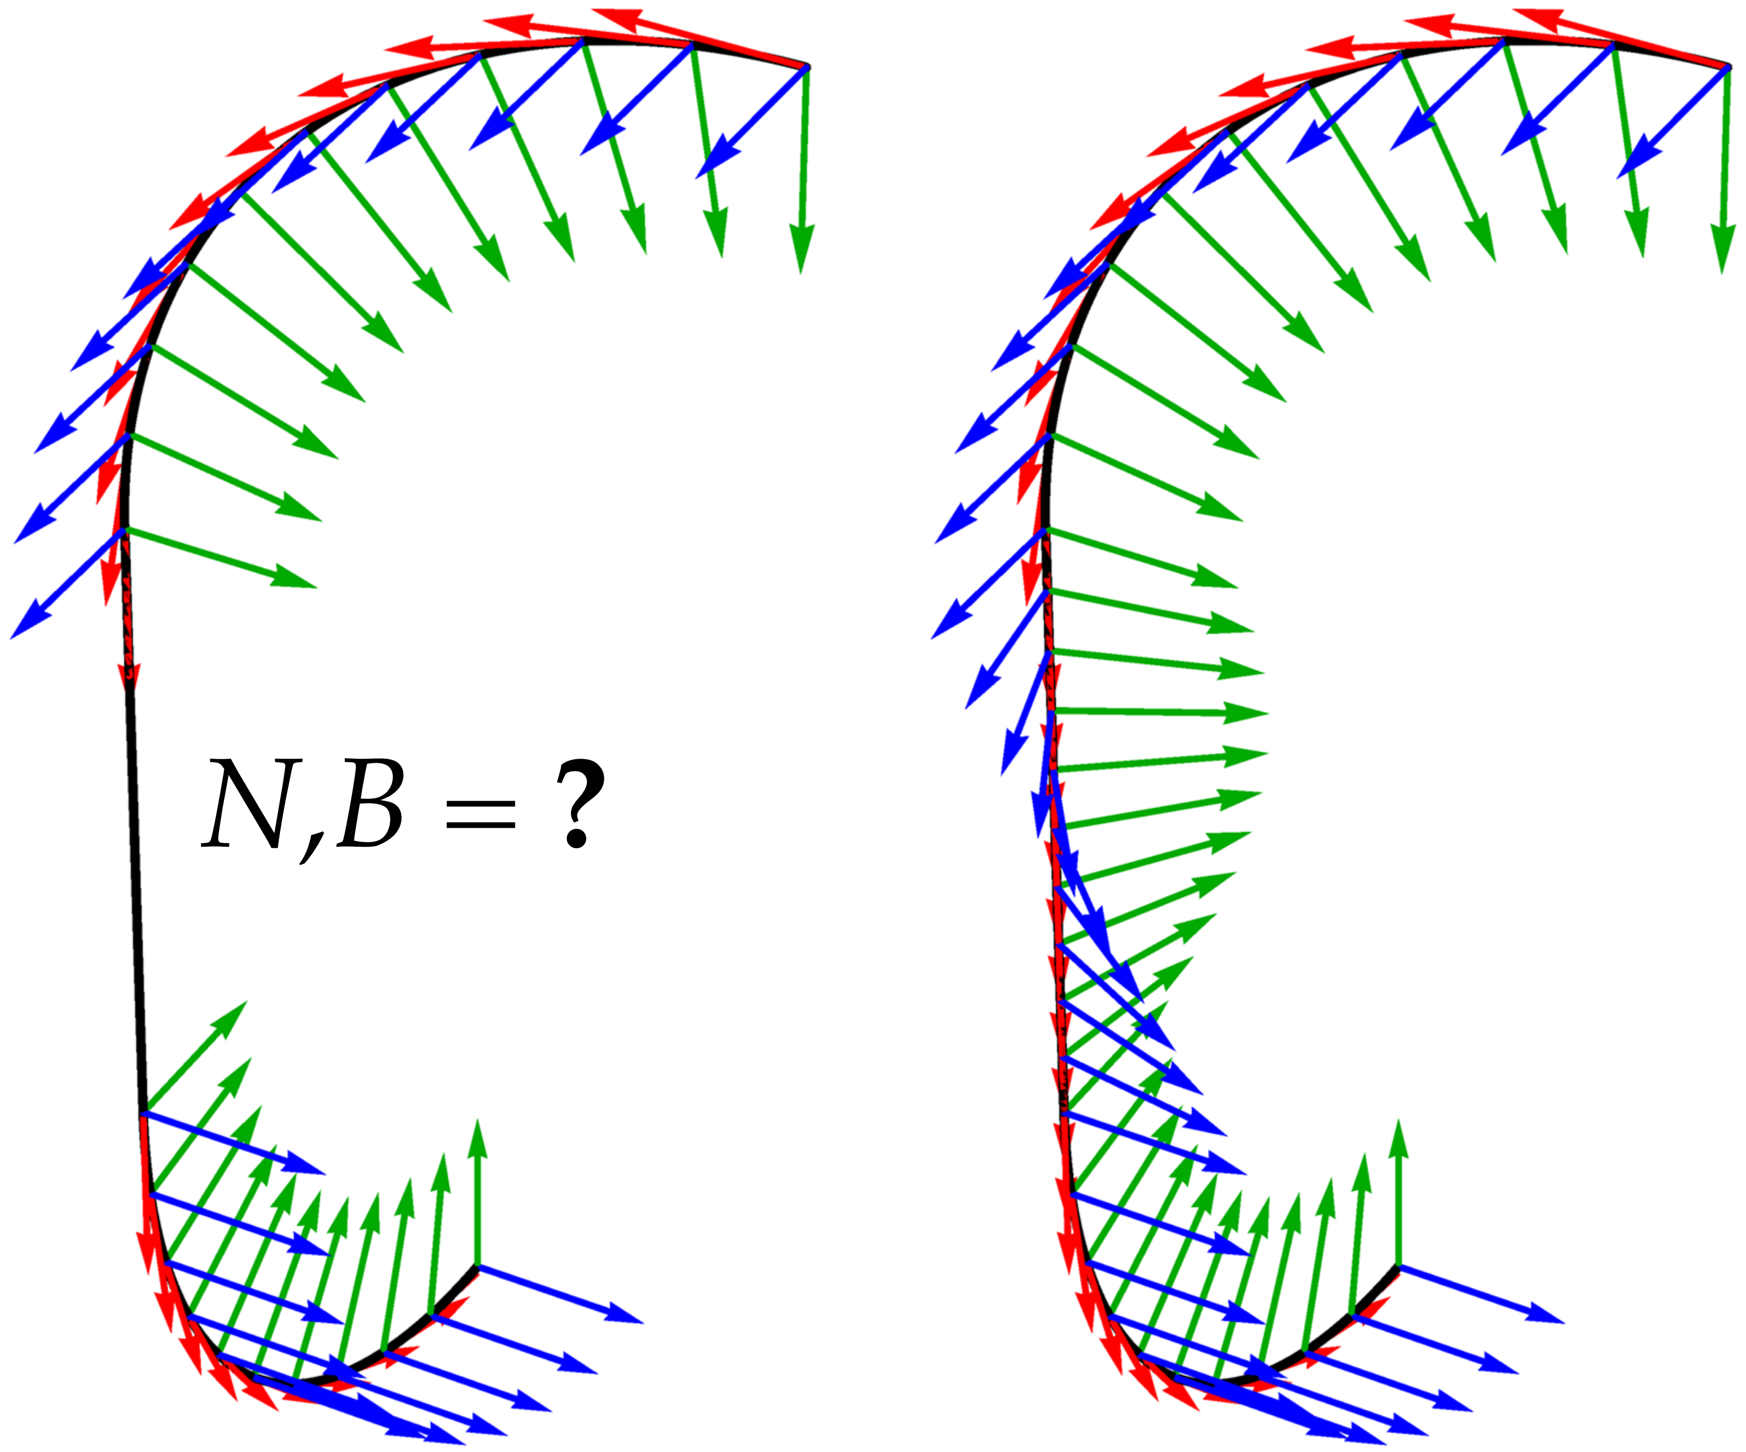

Thanks!

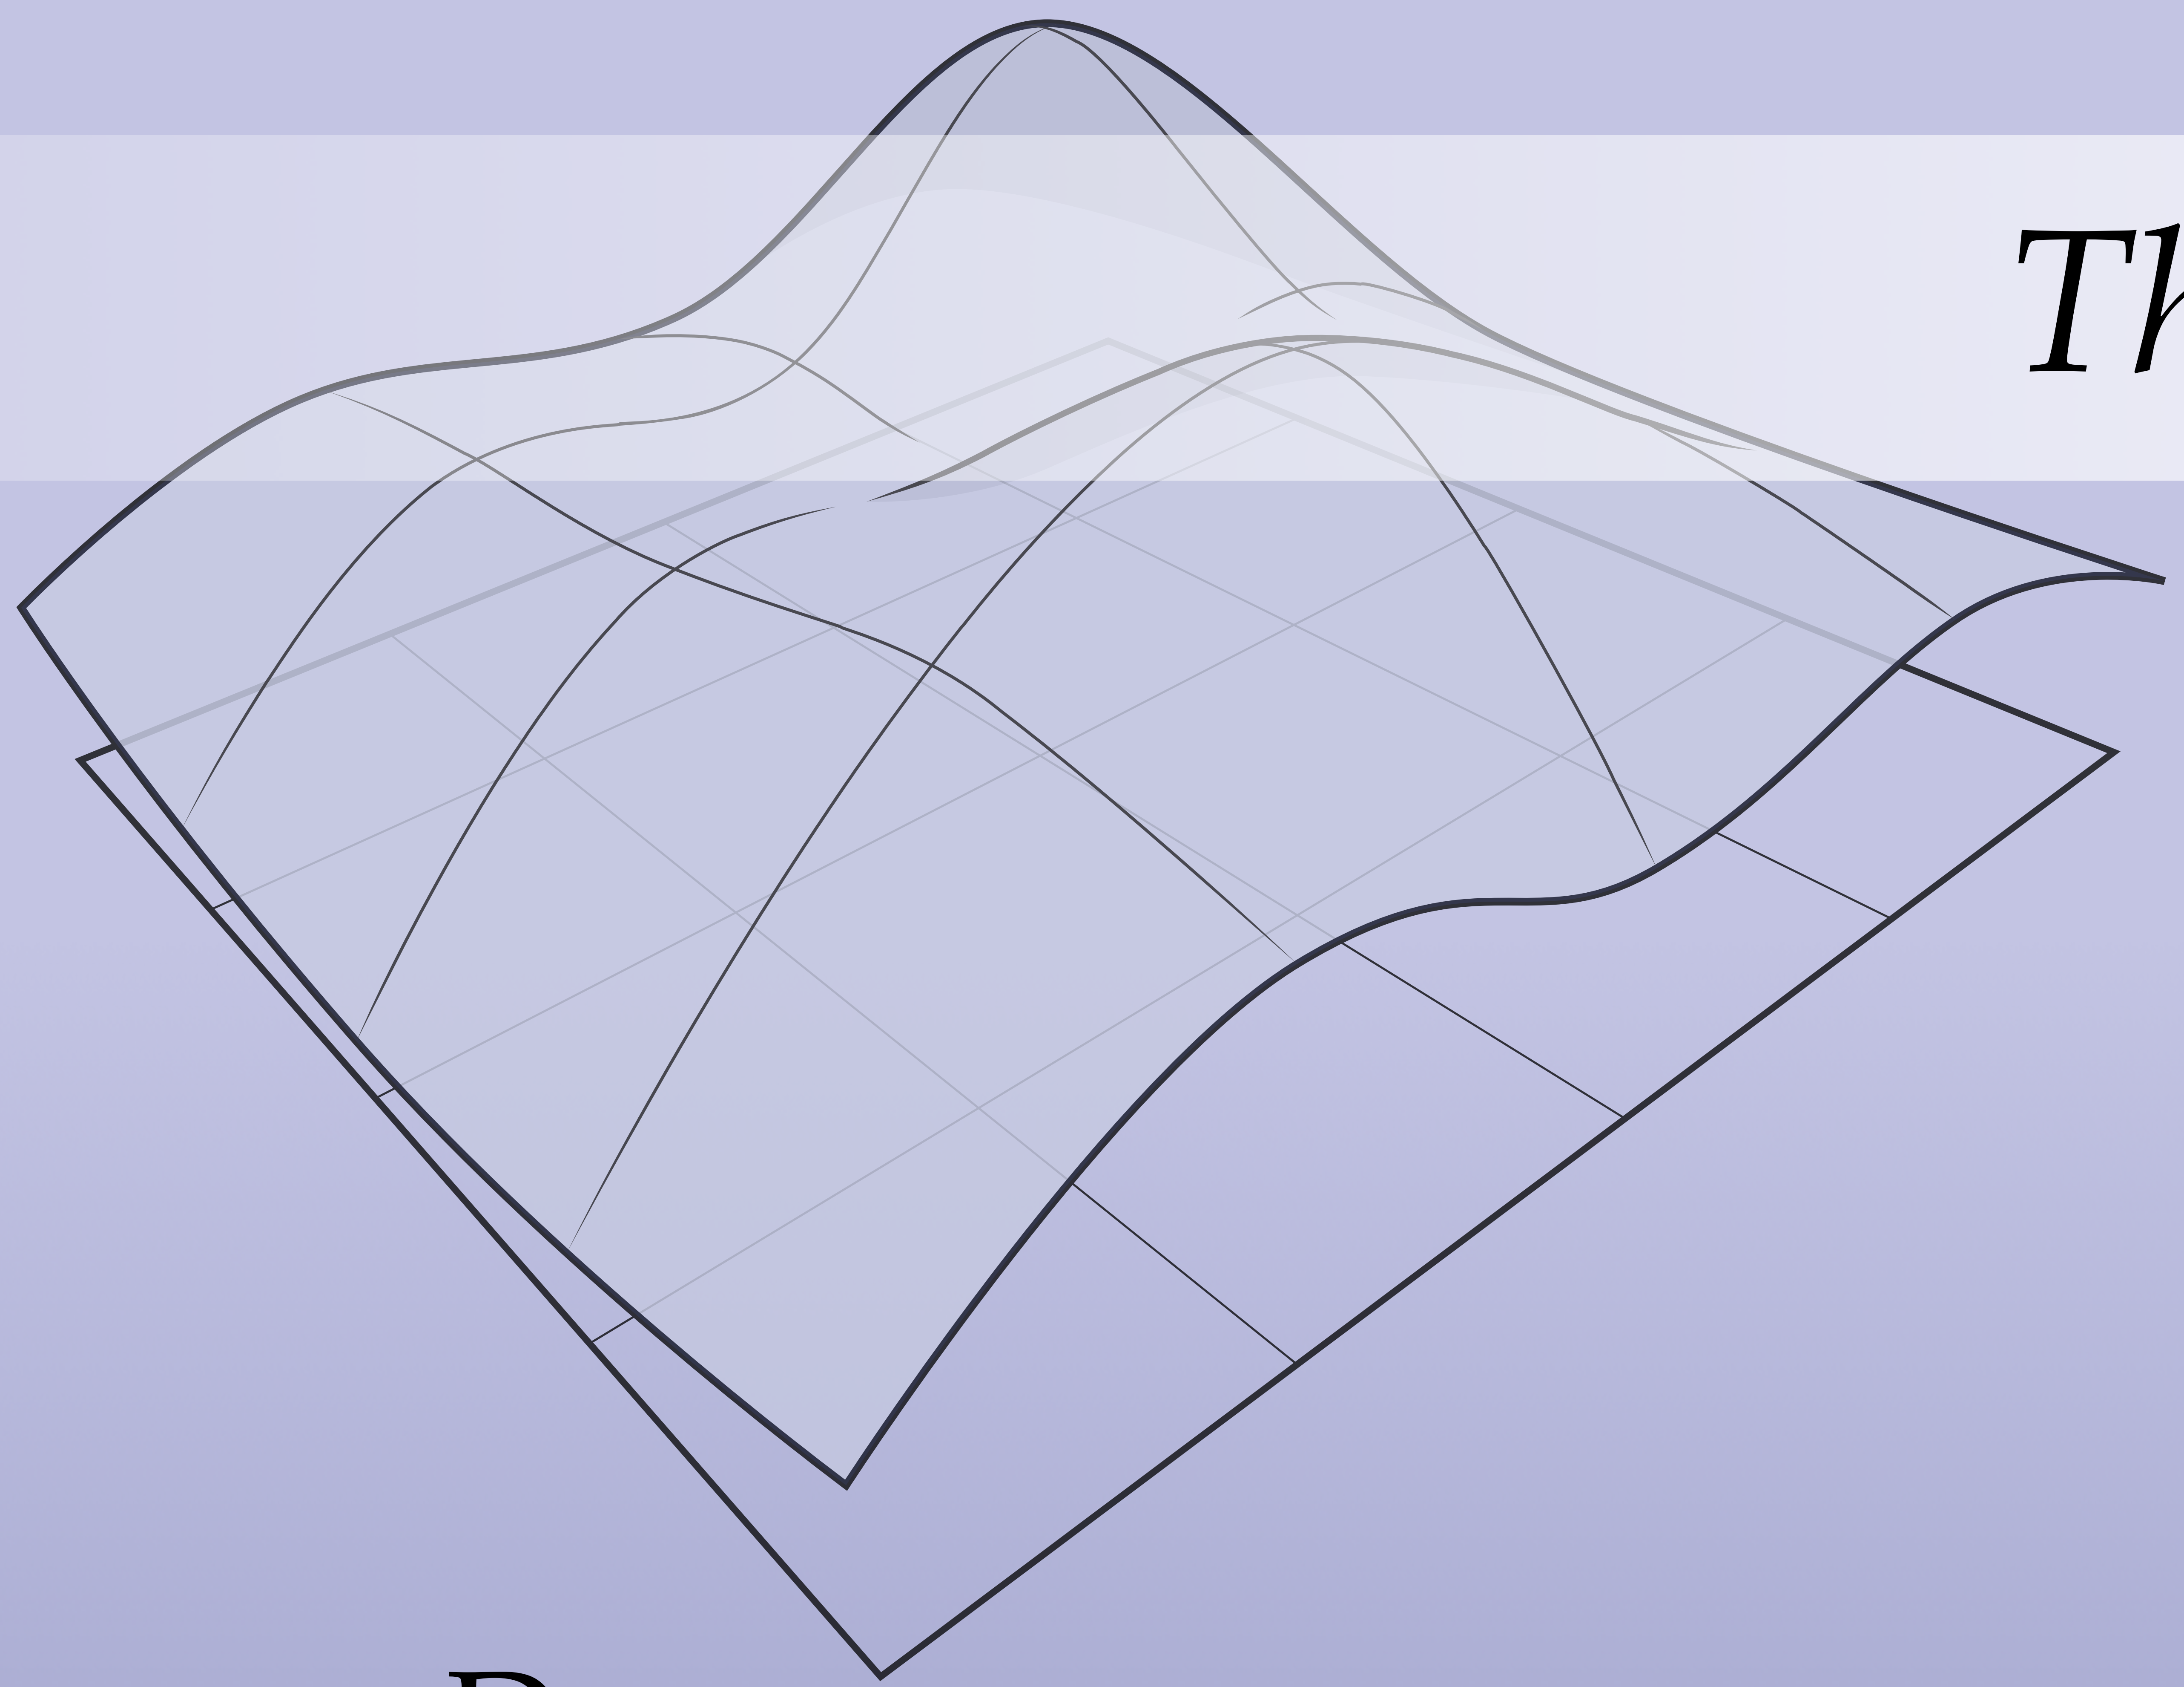

DISCRETE DIFFERENTIAL GEOMETRY: AN APPLIED INTRODUCTION

Keenan Crane • CMU 15-458/858