

We launch into our consideration of smooth and discrete geometry by studying *curves*: one-dimensional pieces of geometry. We owe our inspiration for starting this way to Grinspun and Seedorf, and in particular many of the ideas in this chapter also appear in their discussion [10].

Curves exemplify key considerations that appear when studying geometry. We will start our discussion by asking a deceptively simple question: What *is* a curve? After establishing a careful definition, we will see how the geometry of a curve can be expressed through measurements like curvature, a key building block to defining the curvature of higher-dimensional structures in future chapters. Beyond considering the classical theory of curve geometry, however, we also will develop a parallel story for discrete curves that can be stored on a computer. This will lead to our first “no free lunch” observation about discrete geometry: It will be impossible to have a discrete curve representation that preserves all the nice properties of curves we see in smooth geometry.

In this chapter, we will attempt to keep our mathematical arguments fairly precise and formal. We will relax this approach in future chapters to avoid having to explain the details of both smooth and discrete differential geometry, but our intention here is to show the flavor of what it takes to construct a careful geometric theory.

Starting with curves is no doubt a classical approach, but it is also worth noting where this theory diverges from the broader philosophy of geometry. The shape of a curve is completely determined by the space *around* the curve: A (curvy) measuring tape glued to a curve cannot distinguish between a knotted-up curve and a straight line with the same length. When we move to higher dimensions in Chapter [REF](#), measuring tapes will become more useful: Many species of curvature can be measured without leaving a surface.

3.1 WHAT IS A CURVE?

We all have a reasonable intuition for curves as geometric objects; roughly, curves are abstractions of coils of string embedded in space. But formalizing this notion mathematically reveals some significant subtleties. Basic ideas from calculus can lead us down astray when defining the notion of being a curve:

Example 3.1 (A curve is not a function). *Calculus might suggest that a reasonable definition of a curve is as a continuous function $\gamma(t) : (a, b) \rightarrow \mathbb{R}^n$. While this abstraction is often useful for calculations, it is not a suitable definition. The simplest counterexample is the function $\gamma(t) \equiv (0, 0, 0)$, which clearly does not trace out an interesting geometric object.*

Even if we exclude point-like features, defining curves as functions remains problematic. For instance, the following two functions define the same curve in \mathbb{R}^2 , a straight line:

$$\begin{aligned} \gamma_1(t) &= (t, 2t) \\ \gamma_2(t) &= \begin{cases} (t, 2t) & \text{if } t \leq 1 \\ (2(t - \frac{1}{2}), 4(t - \frac{1}{2})) & \text{if } t > 1. \end{cases} \end{aligned}$$

Notice γ_1 and γ_2 are different functions of t but define the same curve in \mathbb{R}^3 .

It is also worth noting that the same terminology in calculus can take on different significance in the world of shapes:

Example 3.2 (Non-differentiable curve). *Consider the trace of the function*

$$\gamma(t) := (t^2, t^3), \tag{3.1}$$




Figure 5: The cusp in this plot is a geometric feature, even though it is the trace of an infinitely-differentiable function γ defined in equation (3.1).

DRAW ME!

Figure 6: Illustration of Definition 3.1.

plotted in Figure 5. As a function of t , $\gamma(t)$ is smooth. But as a geometric object, $\gamma(t)$ contains a cusp at $t = 0$. Roughly, if $\gamma(t)$ gives the position of a car driving along the trace of γ as a function of time t , at time $t = 0$ the car grinds to a halt, turns its steering wheel, and moves off in a completely different direction. Quantities like curvature are ill-defined at this position, calling into question whether we would consider the trace of $\gamma(t)$ a “smooth curve” even though γ is a smooth function.

These counterexamples suggest a simpler starting point:

A curve is a set of points.

This observation may appear obvious, but in some sense it is counter to the starting points we might try from calculus class. While a *parameterization* $\gamma(t) : (a, b) \rightarrow \mathbb{R}^n$ may be convenient for doing calculations, the facts that (1) there are multiple ways to parameterize the same curve and (2) some functions $\gamma(t)$ do not correspond to curves make it hard to build a definition out of parameterizations.

One reasonable definition for what it means to be a curve is below:

Definition 3.1 (Differentiable curve in \mathbb{R}^n). A differentiable curve in \mathbb{R}^n is a set of points $C \subset \mathbb{R}^n$ with the property that for every $\mathbf{p} \in C$ there exists an open neighborhood $U \subseteq \mathbb{R}^n$ containing \mathbf{p} and a smooth function $\gamma_{\mathbf{p}} : (a, b) \rightarrow C \cap U$, so that

$$C \cap U = \{\gamma_{\mathbf{p}}(t) : t \in (a, b)\}$$

and $\gamma'_{\mathbf{p}}(t) \neq \mathbf{0}$ for all $t \in (a, b)$. The function $\gamma_{\mathbf{p}}$ is known as a local parameterization of C at \mathbf{p} .

This definition is a mouthful, but in reality it is a careful encoding of our intuition for what it means to be a curve. First, we have defined C to be a set of points in \mathbb{R}^n rather than defining C to be a function; from the start, this makes it clear that C is a geometric object. The remainder of the definition, illustrated in Figure 6, codifies the following steps:

- Take any point \mathbf{p} on the curve C .
- Take a magnifying glass and zoom close to \mathbf{p} by choosing a neighborhood U around \mathbf{p} .

DRAW ME!

Figure 7: Objects that are not considered curves by our definition.

- Take scissors and cut out the portion of the curve C inside the neighborhood U .
- Glue a piece of string $\gamma_{\mathbf{p}}$ to cover $C \cap U$. The condition $\gamma'_{\mathbf{p}}(0) \neq 0$ makes sure the string does not have a kink, as in Figure 5.

Our definition excludes a few cases shown in Figure 7, including curves with cusps and curves that self-intersect.

Example 3.3 (Circle). *The unit circle*

$$S^1 := \{\mathbf{p} \in \mathbb{R}^2 : \|\mathbf{p}\|_2 = 1\}$$

is a differentiable curve. It is straightforward to show that the function

$$\gamma(t) = (\cos t, \sin t)$$

parameterizes S^1 completely. Alternatively, we could parameterize S^1 using four functions of $t \in (-1, 1)$:

$$\begin{aligned}\gamma_1(t) &= (t, \sqrt{1-t^2}) \\ \gamma_2(t) &= (t, -\sqrt{1-t^2}) \\ \gamma_3(t) &= (-\sqrt{1-t^2}, t) \\ \gamma_4(t) &= (\sqrt{1-t^2}, t)\end{aligned}$$

These four functions cover the top, bottom, left, and right halves of C . They individually only cover overlapping pieces of S^1 , but together they show that the entire circle admits local parameterization.

Definition 3.1 represents our first example of the intricate interplay between calculus and geometry. On the one hand, the fact that curves are geometric objects dictates that we represent curves as sets $C \subseteq \mathbb{R}^n$. But since we eventually want to compute derivatives along C , we have to link it back to calculus through use of a local parameterization.

We pause to note a related definition common in the world of curve geometry:

Definition 3.2 (Parameterized curve). *A parameterized curve in \mathbb{R}^n is the image (or trace) of a differentiable function $\gamma : (a, b) \rightarrow \mathbb{R}^n$, where $-\infty \leq a < b \leq \infty$ and $\gamma'(t) \neq \mathbf{0}$ for all $t \in (a, b)$.*

Exercise 3.1. relates Definitions 3.1 and 3.2. In effect, this exercise allows us to define properties of parameterized curves and have them apply to differentiable curves as well. Intuitively, the function γ gives the position of a car driving along the curve as a function of time $t \in (a, b)$; the relevant geometric object is the shape of the smoke left out of the tailpipe of the car, rather than the velocity with which the car drives along the curve.

3.2 CAN YOU MEASURE VELOCITY?

Our eventual goal is to define geometric measurements near points $\mathbf{p} \in C$. A common approach to this task is to take a local parameterization $\gamma_{\mathbf{p}}(t)$ and define the measurement in terms of $\gamma_{\mathbf{p}}$ and its derivatives. For this to make sense, however, we must double check that if we took a *different* parameterization at \mathbf{p} , the measurement would yield the same value: Otherwise it is not a measurement about C but rather about the function $\gamma_{\mathbf{p}}$.

Example 3.4 (Velocity and reparameterization). *Suppose we have a parameterized curve C traced by $\gamma(t) : (a, b) \rightarrow \mathbb{R}^3$. We can define the velocity of γ at t as the first derivative $\mathbf{v}(t) = \gamma'(t)$. While it is*

perfectly legitimate to claim that $\mathbf{v}(t)$ is a property of the function $\gamma(t)$, it is **not** reasonable to think of $\mathbf{v}(t)$ as a property of the point $\mathbf{p} := \gamma(t) \in C$ on the curve without referring to γ , as we show below.

Suppose $\phi(t) : (\tilde{a}, \tilde{b}) \rightarrow (a, b)$ is a smooth, bijective function. Geometrically, $\tilde{\gamma}(t) := \gamma \circ \phi(t) = \gamma(\phi(t))$ traces out the same curve; it is known as a reparameterization of C . But by the chain rule

$$\tilde{\gamma}'(t) = \frac{d}{dt}\gamma(\phi(t)) = \mathbf{v}(\phi(t))\phi'(t).$$

In other words, the velocity of γ and the velocity of $\tilde{\gamma}$ differ by a factor of ϕ' . This is a mathematical reflection of the fact that different cars can drive along the same curve C with different velocities and accelerations.

As a concrete example, recall the parameterization of S^1 given by

$$\gamma(t) = (\cos t, \sin t).$$

Taking $\phi(t) = 2t$ yields a second parameterization

$$\tilde{\gamma}(t) = (\cos 2t, \sin 2t).$$

The velocities of these two functions are

$$\begin{aligned}\mathbf{v}(t) &= (-\sin t, \cos t) \\ \tilde{\mathbf{v}}(t) &= (-2\sin 2t, 2\cos 2t).\end{aligned}$$

At $t = 0$, we have $\gamma(0) = \tilde{\gamma}(0) = (1, 0)$. But,

$$\mathbf{v}(0) = (0, 1) \neq (0, 2) = \tilde{\mathbf{v}}(0).$$

Hence, velocity is not a property of the circle S^1 itself, since different parameterizations of S^1 led to different velocities at $\mathbf{p} = (0, 1)$.

3.3 ARC LENGTH

We have set the stage for our first geometric measurement along a curve:

Definition 3.3 (Arc length). *The arc length of a parameterized curve C traced by $\gamma : (a, b) \rightarrow \mathbb{R}^n$ is given by*

$$L[C] := \int_a^b \|\gamma'(t)\|_2 dt. \quad (3.2)$$

To make sure that our definition of L makes sense, we need to check that it is invariant to reparameterization, i.e. that it is a function of the geometry of $C \subset \mathbb{R}^n$ rather than the particular choice of γ . To do so, take $\tilde{\gamma} : (\tilde{a}, \tilde{b}) \rightarrow \mathbb{R}^n$ to be a second parameterization of C . Applying the intuition that parameterizations are just different ways to walk along the same curve—or more formally by applying Exercise 3.2.—we can convince ourselves that there exists a bijective function $\phi(\tilde{t}) : (\tilde{a}, \tilde{b}) \rightarrow (a, b)$ such that $\tilde{\gamma}(\tilde{t}) = \gamma(\phi(\tilde{t}))$ with $|\phi'(\tilde{t})| > 0$. Abusing notation slightly by using γ as an argument to $L[\cdot]$ we have:

$$\begin{aligned}L[\tilde{\gamma}] &= \int_{\tilde{a}}^{\tilde{b}} \|\tilde{\gamma}'(\tilde{t})\|_2 d\tilde{t} \text{ by definition of } L \\ &= \int_{\tilde{a}}^{\tilde{b}} \|[\gamma \circ \phi]'(\tilde{t})\|_2 d\tilde{t} \text{ by definition of } \tilde{\gamma} \\ &= \int_{\tilde{a}}^{\tilde{b}} \|\gamma'(\phi(\tilde{t}))\|_2 |\phi'(\tilde{t})| d\tilde{t} \text{ by the chain rule} \\ &= \int_a^b \|\gamma'(t)\|_2 dt \text{ by taking } t := \phi(\tilde{t}) \\ &= L[\gamma].\end{aligned}$$

DRAW ME!

Figure 8: Illustration of Proposition 3.1.

DRAW ME!

Figure 9: Up to a change of base point, there exist two parameterizations of a curve by arc length; these are its two orientations.

This chain of equalities justifies our use of notation $L[C]$ rather than $L[\gamma]$ to denote length, but it does not justify that the integral (3.2) has anything to do with the length of C . Although we will not prove it here, a slightly more involved argument justifies the definition above:

Proposition 3.1 ([15], Proposition 1.6). *If C is parameterized by γ , then for every $\varepsilon > 0$ there exists $\delta > 0$ such that*

$$\max_i |t_i - t_{i-1}| < \delta \implies \left| L[C] - \sum_i \|\gamma(t_i) - \gamma(t_{i-1})\|_2 \right| < \varepsilon.$$

Here, P is any partition $P := \{t_0 = a < t_1 < \dots < t_n = b\}$.

This proposition indicates that the length of C is approximated by the sum of the lengths of a set of line segments along C , illustrated in Figure 8.

3.4 PARAMETERIZATION BY ARC LENGTH

Take a time $t_0 \in (a, b)$, and define the *arc length function* $S : (a, b) \rightarrow \mathbb{R}_+$ as the integral

$$S(t) := \int_{t_0}^t \|\gamma'(t)\|_2 dt. \tag{3.3}$$

This function gives the distance traveled along the curve C by the parameterization γ in the time interval from t_0 to t .

By the Fundamental Theorem of Calculus, the derivative of (3.3) is

$$S'(t) = \|\gamma'(t)\|_2 > 0. \tag{3.4}$$

This shows that S is a *strictly increasing* function, and in particular it admits an inverse ϕ such that $t = \phi \circ S(t)$. Take $\tilde{\gamma} := \gamma \circ \phi$; this reparameterization of γ is known as the *reparameterization by arc length*. The reparameterization $\tilde{\gamma}$ admits many special properties, most notably the following:

$$\begin{aligned} \tilde{\gamma}'(s) &= \frac{d}{ds} \gamma(\phi(s)) \text{ by definition of } \tilde{\gamma} \\ &= \gamma'(\phi(s)) \cdot \phi'(s) \text{ by the chain rule} \\ &= \frac{\gamma'(\phi(s))}{\|\gamma'(\phi(s))\|_2} \text{ since } \phi = S^{-1} \text{ and by (3.4)} \\ \implies \|\tilde{\gamma}'(s)\|_2 &= 1. \end{aligned} \tag{3.5}$$

In other words, the derivative of $\tilde{\gamma}$ has unit length!

Parameterization by arc length is a fairly intuitive construction. Given a distance s , the point $\tilde{\gamma}(s)$ is distance s away from $\tilde{\gamma}(0)$. Or, as a function of s , $\tilde{\gamma}(s)$ gives the position of a car driving along C with constant speed, as verified by the fact that $\|\tilde{\gamma}'(s)\|_2 \equiv 1$ for all s .

For ease of notation, from here on we will always use s to denote arc length, and we will use t to denote a more general parameter. For example, the notation $\gamma(s)$ denotes that γ is *parameterized by arc length*, and hence $\|\gamma'(s)\|_2 \equiv 1$. Up to a constant shift in s (corresponding to a change in base point), there exist two parameterizations by arc length, as illustrated in Figure 9; these correspond

DRAW ME!

Figure 10: Frenet frame in two dimensions, including unit tangent \mathbf{T} , unit normal \mathbf{N} , and turning angle θ .

to driving along the curve C forward or backward. We will call these two directions *orientations* of the curve C , and we will denote an *oriented curve* as one with a single orientation chosen.

Any measurement written in terms of the function $\gamma(s)$ and its derivatives is a property of the local geometry of the oriented curve C at the point $\mathbf{p} := \gamma(s)$; this follows from uniqueness of the arc length parameterization. In a sense, we have gotten rid of acceleration *along* the curve, e.g. a car pushing its gas pedal harder and harder while driving along C , by choosing a parameterization with constant speed. This gives us an easy way to make *purely geometric* measurements about C using constructions from calculus, as we do below.

3.5 TWO-DIMENSIONAL CURVES

Take $\gamma(s) : (a, b) \rightarrow \mathbb{R}^2$ to be a curve in the plane, and denote its *unit tangent* as

$$\mathbf{T}(s) := \gamma'(s).$$

Recall from (3.5) that $\|\mathbf{T}(s)\|_2 = 1$. Hence, we can find an angle $\theta(s)$ so that

$$\mathbf{T}(s) = \cos \theta(s) \mathbf{e}_1 + \sin \theta(s) \mathbf{e}_2. \quad (3.6)$$

Differentiating, we find

$$\begin{aligned} \mathbf{T}'(s) &= \theta'(s) [-\sin \theta(s) \mathbf{e}_1 + \cos \theta(s) \mathbf{e}_2] \\ &= \kappa(s) \mathbf{N}(s), \end{aligned} \quad (3.7)$$

where

$$\kappa(s) := \theta'(s), \quad \mathbf{N}(s) := J\mathbf{T}(s), \quad \text{and} \quad J := \begin{pmatrix} 0 & -1 \\ 1 & 0 \end{pmatrix}. \quad (3.8)$$

Differentiating once more shows

$$\begin{aligned} \mathbf{N}'(s) &= J\mathbf{T}'(s) \text{ by (3.8)} \\ &= \kappa(s)J\mathbf{N}(s) \text{ by (3.7)} \\ &= -\kappa(s)\mathbf{T}(s) \text{ since } J^2 = -I_{2 \times 2} \text{ and } \mathbf{N} = J\mathbf{T}. \end{aligned} \quad (3.9)$$

Combining (3.7) and (3.9) gives the *Frenet equations* for our curve:

$$\frac{d}{ds} \begin{pmatrix} \mathbf{T}(s) \\ \mathbf{N}(s) \end{pmatrix} = \begin{pmatrix} 0 & \kappa(s) \\ -\kappa(s) & 0 \end{pmatrix} \begin{pmatrix} \mathbf{T}(s) \\ \mathbf{N}(s) \end{pmatrix}. \quad (3.10)$$

Figure 10 illustrates the construction above. Given an oriented curve, we constructed a *frame* of two orthonormal vectors $\{\mathbf{T}, \mathbf{N}\}$ at every point. Each of these vectors and our derived quantities can be described in terms of shape:

- The *unit tangent* vector $\mathbf{T}(s)$ grazes the curve at $\gamma(s)$.
- The *unit normal* vector $\mathbf{N}(s)$ points to the left side of the curve perpendicularly $\mathbf{T}(s)$.
- The *turning angle* $\theta(s)$ measures the angle of $\mathbf{T}(s)$ from the x axis \mathbf{e}_1 .

- The *signed curvature* $\kappa(s)$ measures the change in turning angle as a function of arc length. Since $\kappa(s) = \theta'(s)$, if $\gamma(s)$ denotes the position of a car driving along C with unit speed, then $\kappa(s)$ is the rate at which we must turn the steering wheel of the car to stay driving along γ . Alternatively, since $\gamma''(s) = \mathbf{T}'(s) = \kappa(s)\mathbf{N}(s)$, we can interpret $\kappa(s)$ as the *force* felt by the driver of the car driving with constant speed along the curve; as expected, this force is zero when the path of the car is straight, since arc length parameterization prevents the driver from accelerating in the tangent direction \mathbf{T} .

The key geometric measurement about the oriented curve is its curvature $\kappa(s)$, associated with a point $\gamma(s)$.

The Frenet equations (3.10) are reflective of a certain philosophy taken in differential geometry. If we take a curve and rotate or shift it, its geometry has not changed; in a sense, reference to Cartesian (x, y) coordinates is against the idea that all interesting geometry can be measured using rulers and compasses. But without coordinates, it is difficult to understand measurements like the derivatives of \mathbf{T} and \mathbf{N} . Hence, we choose a coordinate system *intrinsic* to the curve C itself: The Frenet equations (3.10) write the derivatives of \mathbf{T} and \mathbf{N} in the $\{\mathbf{T}, \mathbf{N}\}$ basis!

The curvature $\kappa(s)$ is not just an interesting measurement one can make along an oriented two-dimensional curve, but rather is all the information you need to describe the shape of a curve in 2D. This is formalized by the following proposition:

Proposition 3.2 (Fundamental Theorem of the Local Theory of Plane Curves). *Given a differentiable function $\kappa(s) : (a, b) \rightarrow \mathbb{R}$, there exists a curve $\gamma(s) : (a, b) \rightarrow \mathbb{R}^2$ with curvature $\kappa(s)$. If two curves have the same curvature $\kappa(s)$, they are the same up to rotation and translation.*

Proof. Take $s_0 \in (a, b)$, and define

$$\theta(s) := \int_{s_0}^s \kappa(s) ds.$$

Then, define

$$\gamma(s) := \left[\int_{s_0}^s \cos \theta(s) ds \right] \mathbf{e}_1 + \left[\int_{s_0}^s \sin \theta(s) ds \right] \mathbf{e}_2.$$

It is straightforward to verify that $\gamma(s)$ has curvature $\kappa(s)$.

Now, suppose a second curve $\tilde{\gamma}(s)$ also has curvature $\kappa(s)$. Notice that $\tilde{\gamma}$ and γ both satisfy the same ODE

$$\frac{d}{ds} \begin{pmatrix} \gamma(s) \\ \mathbf{T}(s) \\ \mathbf{N}(s) \end{pmatrix} = \begin{pmatrix} 0 & 1 & 0 \\ 0 & 0 & \kappa(s) \\ 0 & -\kappa(s) & 0 \end{pmatrix} \begin{pmatrix} \gamma(s) \\ \mathbf{T}(s) \\ \mathbf{N}(s) \end{pmatrix},$$

just with different initial conditions. Translating and rotating $\gamma, \tilde{\gamma}$ to align at $s = 0$ shows they are the same through uniqueness of solutions to ODE. \square

The local theory explored above, the function κ also encodes a *topological invariant*. Define a *closed curve* to be a curve whose endpoints matched, i.e., a curve that can be parameterized by a function $\gamma(t)$ with $\gamma(a) = \gamma(b)$ and $\mathbf{T}(a) = \mathbf{T}(b)$. Then, we make the following definition:

Definition 3.4 (Winding number). *The winding number of a closed curve $\gamma(s) : [a, b] \rightarrow \mathbb{R}^2$ is given by*

$$W[\gamma] := \frac{1}{2\pi} \int_a^b \kappa(s) ds, \quad (3.11)$$

where s is arc length.

We can prove that if we smoothly deform γ its winding number does not change via the following proposition:

Theorem 3.1 (Winding number theorem). *$W[\gamma]$ is an integer, and smoothly deforming γ does not affect its winding number.*

DRAW ME!

Figure 11: First variation formula for two-dimensional arc length.

Proof. By definition, we can write $W[\gamma] = \theta(b) - \theta(a)$. Since $\mathbf{T}(a) = \mathbf{T}(b)$, we have $\cos\theta(a) = \cos\theta(b)$ and $\sin\theta(a) = \sin\theta(b)$, which together imply that $\theta(a) - \theta(b)$ is an integer multiple of 2π , as needed. Smoothly deforming γ should smoothly change $W[\gamma]$ since it is defined via the integral (3.11); since it remains integer, it must be constant. \square

Intuition for the winding number is given in Figure REF; roughly, it counts the number of times the tangent vector to γ rotates about its base while traversing the length of the curve. The only way to change $W[\gamma]$ is to introduce a kink, as shown in Figure REF.

3.5.1 Advanced Topic: First Variation Formula

We briefly outline an alternative way to motivate the curvature $\kappa(s)$ of a two-dimensional curve. Suppose $a < b$ and take $\gamma(s) : [a, b] \rightarrow \mathbb{R}^2$ to be a parameterized plane curve with endpoints $\gamma(a)$ and $\gamma(b)$. Now, take $\mathbf{v}(t) : (a, b) \rightarrow \mathbb{R}^2$ to be a *vector field* along $\gamma(t)$, illustrated in Figure 11; for simplicity we will assume $\mathbf{v}(a) = \mathbf{v}(b) = \mathbf{0}$. We can define a one-parameter family of curves

$$\gamma_r(t) := \gamma(t) + r\mathbf{v}(t).$$

Here, $\gamma_0(t) \equiv \gamma(t)$, and as r changes we slide γ along the displacement specified by $\mathbf{v}(t)$.

We can think of the length of γ_r as a function of r given by

$$\ell(r) := \int_a^b \|\gamma'(t) + r\mathbf{v}'(t)\|_2 dt.$$

This function gives the length of $\gamma_r(t)$ from $t = a$ to $t = b$. We can compute the derivative of this function at $r = 0$ as follows:

$$\begin{aligned} \ell'(r) &= \frac{d}{dr} \int_a^b \|\gamma'(t) + r\mathbf{v}'(t)\|_2 dt \\ &= \int_a^b \frac{d}{dr} \|\gamma'(t) + r\mathbf{v}'(t)\|_2 dt \text{ by differentiating under the integral sign} \\ &= \int_a^b \frac{(\gamma'(t) + r\mathbf{v}'(t)) \cdot \mathbf{v}'(t)}{\|\gamma'(t) + r\mathbf{v}'(t)\|_2} dt \text{ by exercise 3.4.} \\ \implies \ell'(0) &= \int_a^b \frac{\gamma'(t)}{\|\gamma'(t)\|_2} \cdot \mathbf{v}'(t) dt \text{ by substituting } r = 0 \\ &= \int_a^b \mathbf{T}(t) \cdot \mathbf{v}'(t) dt \text{ by definition of the unit tangent} \\ &= - \int_a^b \mathbf{T}'(t) \cdot \mathbf{v}(t) dt \text{ by integration by parts, since } \mathbf{v}(a) = \mathbf{v}(b) = \mathbf{0} \\ &= - \int_a^b \kappa(t)\mathbf{N}(t) \cdot \mathbf{v}(t) dt \text{ by definition of } \kappa, \mathbf{N}. \end{aligned} \tag{3.12}$$

Now we make a conceptual leap. Suppose we have a function $f : \mathbb{R}^n \rightarrow \mathbb{R}$, and take $\mathbf{x} \in \mathbb{R}^n$ and $\mathbf{v} \in \mathbb{R}^n$. Then, one way to understand the gradient is through the relationship

$$\left. \frac{d}{dr} f(\mathbf{x} + r\mathbf{v}) \right|_{r=0} = \nabla f(\mathbf{x}) \cdot \mathbf{v}. \tag{3.13}$$

Take a look at our derivation above of $\ell'(0)$. Think of arc length L as a function from curves $\gamma(t)$ to real numbers; γ replaces \mathbf{x} in our finite-dimensional story about $f(\mathbf{x})$. Then, we can think if $\mathbf{v}(t)$ as

DRAW ME!

Figure 12: Curve-shortening flow.

a perturbation, similar to the vector \mathbf{v} . In this light, by comparison to (3.13), the expression (3.12) takes on different meaning, showing that the negative curvature normal $-\kappa\mathbf{N}$ is in some sense the “gradient” of the arc length functional L from curves to real numbers; in differential geometry language, we call the curvature-weighted normal the *first variation* of arc length.

As an example application, recall that *gradient descent* on a function $f(\mathbf{x})$ attempts to modify \mathbf{x} locally by moving along the gradient of f :

$$\mathbf{x}' = -\nabla f(\mathbf{x}).$$

This ODE can be considered an optimization routine that attempts to find the \mathbf{x} that minimizes f locally. Similarly, for a curve γ with fixed endpoints, *curve-shortening flow* moves the points of γ along $\kappa\mathbf{N}$, reducing the arc length until a straight line is reached; an example is illustrated in Figure 12.

3.6 CURVES IN \mathbb{R}^3

We conclude our consideration of smooth curves by examining parameterized curves embedded in \mathbb{R}^3 , formalized as $\gamma(s) : (a, b) \rightarrow \mathbb{R}^3$. This leads to one more familiar object, the *Frenet frame*, which appears in many applications of smooth and discrete differential geometry.

We begin with a short lemma:

Proposition 3.3. *Suppose $\mathbf{v}(t) : (a, b) \rightarrow \mathbb{R}^n$ is differentiable with $\|\mathbf{v}(t)\|_2 \equiv 1$ for all $t \in (a, b)$. Then, $\mathbf{v}'(t) \cdot \mathbf{v}(t) \equiv 0$ for all t , that is, $\mathbf{v}'(t) \perp \mathbf{v}(t)$. Similarly, if $\mathbf{v}(t) \cdot \mathbf{w}(t) \equiv 0$ for all $t \in (a, b)$, then $\mathbf{v}'(t) \cdot \mathbf{w}(t) = -\mathbf{v}(t) \cdot \mathbf{w}'(t)$.*

Proof. This is a direct consequence of the product rule:

$$0 = \frac{d}{dt} 1 = \frac{d}{dt} \|\mathbf{v}(t)\|_2^2 = \frac{d}{dt} \mathbf{v}(t) \cdot \mathbf{v}(t) = 2\mathbf{v}'(t) \cdot \mathbf{v}(t).$$

The second relationship is derived similarly. □

Suppose $\gamma(s)$ is parameterized by arc length with $\gamma''(s) \neq \mathbf{0}$ for all $s \in (a, b)$. We define the unit tangent again to be $\mathbf{T}(s) := \gamma'(s)$ and have $\|\mathbf{T}(s)\|_2 = 1$ for all $s \in (a, b)$. Hence, by the proposition above, there exists a unit vector $\mathbf{N}(s)$ and a scaling factor $\kappa(s)$ such that $\mathbf{T}'(s) = \kappa(s)\mathbf{N}(s)$; so far our story agrees with the two-dimensional picture in §3.5.

Since we are in \mathbb{R}^3 , however, given two orthogonal vectors $\mathbf{T}(s)$ and $\mathbf{N}(s)$, we can define a third *binormal* vector $\mathbf{B}(s)$ via

$$\mathbf{B}(s) := \mathbf{T}(s) \times \mathbf{N}(s). \tag{3.14}$$

The vectors $\{\mathbf{T}, \mathbf{N}, \mathbf{B}\}$ span \mathbb{R}^3 and define a convenient basis for the geometry of the curve γ near $\gamma(s)$. To imitate the relationships in (3.10), we need to compute some derivatives. Since $\{\mathbf{T}, \mathbf{N}, \mathbf{B}\}$ form an orthonormal basis, for any $\mathbf{v} \in \mathbb{R}^3$ we can write

$$\mathbf{v} = (\mathbf{v} \cdot \mathbf{T})\mathbf{T} + (\mathbf{v} \cdot \mathbf{N})\mathbf{N} + (\mathbf{v} \cdot \mathbf{B})\mathbf{B}.$$

Hence,

$$\begin{aligned} \mathbf{N}'(s) &= (\mathbf{N}'(s) \cdot \mathbf{T}(s))\mathbf{T}(s) + (\mathbf{N}'(s) \cdot \mathbf{B}(s))\mathbf{B}(s) \text{ since } \mathbf{N}' \perp \mathbf{N} \text{ by Proposition 3.3} \\ &= -(\mathbf{N}(s) \cdot \mathbf{T}'(s))\mathbf{T}(s) + (\mathbf{N}'(s) \cdot \mathbf{B}(s))\mathbf{B}(s) \text{ by the second half of Proposition 3.3} \\ &= -\kappa(s)\mathbf{T}(s) + \tau(s)\mathbf{B}(s) \text{ by definition of curvature } \kappa(s) \text{ and a new definition below.} \end{aligned}$$

DRAW ME!

Figure 13: Frenet frame in 3D.

Here, we introduce a second quantity, the *torsion* of the curve γ at s :

$$\tau(s) := \mathbf{N}'(s) \cdot \mathbf{B}(s).$$

Completing our suite of derivatives,

$$\begin{aligned} \mathbf{B}'(s) &= (\mathbf{B}'(s) \cdot \mathbf{T}(s))\mathbf{T}(s) + (\mathbf{B}'(s) \cdot \mathbf{N}(s))\mathbf{N}(s) \text{ since } \mathbf{B}' \perp \mathbf{B} \text{ by Proposition 3.3} \\ &= (-\mathbf{B}(s) \cdot \mathbf{T}'(s))\mathbf{T}(s) + (-\mathbf{B}(s) \cdot \mathbf{N}'(s))\mathbf{N}(s) \text{ by the second half of Proposition 3.3} \\ &= -\tau(s)\mathbf{N}(s) \text{ by definition of } \tau. \end{aligned}$$

We now have the three-dimensional equivalent of (3.10):

$$\frac{d}{ds} \begin{pmatrix} \mathbf{T}(s) \\ \mathbf{N}(s) \\ \mathbf{B}(s) \end{pmatrix} = \begin{pmatrix} 0 & \kappa(s) & 0 \\ -\kappa(s) & 0 & \tau(s) \\ 0 & -\tau(s) & 0 \end{pmatrix} \begin{pmatrix} \mathbf{T}(s) \\ \mathbf{N}(s) \\ \mathbf{B}(s) \end{pmatrix}. \quad (3.15)$$

The intuition for this construction is drawn in Figure 13. As an extension of 2D curvature, the curvature $\kappa(s)$ of a curve in three dimensions measures in-plane bending, while the torsion $\tau(s)$ measures how much the curve spirals out of the plane.

Our construction, however, has some drawbacks that do not appear in two dimensions. In particular, recall we assumed $\gamma''(s) \neq \mathbf{0}$, in other words, that the curve is not a straight line. When the curve straightens out, there is no clear way to distinguish \mathbf{N} from \mathbf{B} , the two directions orthogonal to the unit tangent \mathbf{T} . Even when $\gamma''(s) \neq \mathbf{0}$, however, there is a more subtle issue: The sign of $\kappa(s)$ is less clear. Indeed, we could flip the direction of \mathbf{N} as well as the sign of κ and maintain the relationships above. By convention, we often take $\kappa(s) > 0$ locally.

Degeneracies of the Frenet frame aside, we can derive an analogous result to Proposition 3.2: Given the curvature function $\kappa(s)$ and torsion function $\tau(s)$, we know a curve in \mathbb{R}^3 up to rigid motion, and furthermore any pair $(\kappa(s), \tau(s))$ can be used to construct a curve. We omit the proof, although it does not use machinery more complicated than the proof of Proposition 3.2.

3.7 REPRESENTING CURVES ON A COMPUTER

There is a *huge* space of curves embedded in \mathbb{R}^n , from fractals to Bézier splines. Since our hard drives only can store finite numbers of bits, we are forced to adapt the theory of curves to the finite-dimensional regime. For this reason, we have to build a computational theory of curve processing from the ground up, again starting with how to define curves and then defining measurements like length and curvature.

The question of how to store a curve on a computer is strongly application-dependent and has been studied for many decades. A few options are listed below and illustrated in Figure [REF](#):

- Perhaps the simplest representation to store is simply a list of points $\{\mathbf{x}_k\}_k \subset \mathbb{R}^n$ sampled from the curve, in no particular order. This representation obviously makes it challenging to do even simple computations along the curve, although it is sufficient for display purposes.
- An only slightly more complex representation is a *poly-line*, or ordered sequence of points $\mathbf{x}_1, \mathbf{x}_2, \dots \in \mathbb{R}^n$. This representation at least encodes some notion of a local neighborhood, using points with adjacent indices. We can think of poly-lines as collections of line segments connecting pairs $(\mathbf{x}_k, \mathbf{x}_{k+1})$.

- Considerable effort in the theory of computer-aided geometric design and other disciplines goes into the development of *splines* and other curves that have piecewise smoothness properties. For instance, a common curve representation for vector graphics is the piecewise-cubic curve $\gamma(t)$, whose coordinate functions are cubic in different ranges $[t_k, t_{k+1}]$. Well-known conditions on the coefficients of the cubic functions in the different segments can guarantee continuity of the different derivatives of γ .

Each of these representations—in particular the latter two—has advantages and disadvantages. Poly-lines are extremely simple to store and render. Splines and their relatives are represented by differentiable functions for which the smooth computations in the previous sections are relevant.

Rather than applications in rendering, graphic design, time series, or animation, here our focus is on shape analysis. For this task, elegant differentiable curve representations can be difficult to work with:

Example 3.5 (Cubic curve). *A cubic curve has component functions that are cubic polynomials in t :*

$$\gamma(t) = \sum_k \mathbf{e}_k \left[a_0^k + a_1^k t + a_2^k t^2 + a_3^k t^3 \right].$$

No closed-form formula is known for the arc length of $\gamma(t)$ for $t \in [a, b]$ in terms of the a 's [2].

Even though spaces of curves like piecewise cubics are useful for many applications, these rich spaces still do not span the space of *all* curves in \mathbb{R}^3 . In terms of approximation power, in the most basic sense even poly-lines will suffice: By sampling enough t values from a curve $\gamma(t)$, the poly-line connecting these samples can approximate $\gamma(t)$ arbitrarily well. Since poly-lines are the simplest discrete representations of curves with local structure but still can approximate any smooth curve well, we will focus on developing a discrete theory of curves operating on these representations.

3.8 DISCRETE CURVATURE IN TWO DIMENSIONS

As our running example of discrete geometry for the rest of this chapter, we will consider the possibility of measuring geometric quantities along piecewise-linear curves, or poly-lines. Our goal is to illustrate the balance of engineering and mathematics that goes into discrete differential geometry: We will have multiple options for measuring curvature, and no single one of them will preserve *all* the favorable properties we enjoy in the smooth case. Hence, we must choose a species of curvature that makes sense for a given application or task.

To begin, suppose we have a poly-line represented using a chain of points $\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_p \in \mathbb{R}^2$. Our goal is to measure (or approximate) the curvature of our poly-line at each of the \mathbf{x}_i 's; we will ignore \mathbf{x}_i 's near the boundary. We explore three approaches below:

3.8.1 True Curvature

The most basic approach to determining the curvature of a poly-line might be to consider the poly-line as its own geometric object, consisting of line segments $\mathbf{x}_k \rightarrow \mathbf{x}_{k+1}$ connected at angular joints. Unfortunately, this leads to a fairly uninteresting differential structure: $\kappa \equiv 0$ along the interiors of the line segments, joined together by sharp joints at which $\kappa = \infty$. But, we might be able to make sense of this scenario by thinking of κ as a *measure*—so, for example, the joints become Dirac δ -functions roughly encoding “impulses” of curvature. While we will not explore this approach explicitly here, the turning angle discretization in §3.8.3 can be interpreted in such a fashion.

3.8.2 Approximate Approach

The most straightforward approach to curvature approximation suggested by classical numerical analysis or even statistics is to use *curve-fitting*, as illustrated in Figure [REF](#). To approximate the curvature at \mathbf{x}_k , for some $\ell > 0$ we can sample points along the curve $\mathbf{x}_{k-\ell}, \mathbf{x}_{k-\ell+1}, \dots, \mathbf{x}_k, \dots, \mathbf{x}_{k+\ell-1}, \mathbf{x}_{k+\ell}$ and find the best-fitting $\gamma(t)$ in some parametric family to this set of points by solving a least-squares problem; for instance, perhaps we fit a parabolic arc to every triplet of points. If we design γ so that $\gamma(0) \approx \mathbf{x}_i$, then we can use $\kappa(0)$ as an approximation of the curvature at \mathbf{x}_i .

This approach is often described as a “sliding window” or “moving stencil.” Here, our window is the set of $2\ell + 1$ points surrounding \mathbf{x}_i , which are used to approximate the curvature at \mathbf{x}_i . The value of p has a strong bearing on the quality of such an approximation. Small p makes sense from the perspective that curvature is a local derivative measurement, but in this case slight noise in the \mathbf{x}_i 's leads to noisy curvature estimates. As we increase p , our curvature estimates smooth out but can be prone to oversmoothing: In theory only a tiny local neighborhood of t should affect $\kappa(t)$.

It may be possible to prove certain mathematical properties of these curvature approximations. We will not focus on them here, but for careful choices of the class of functions for γ , we may be able to show that our estimate of κ *converges in the limit of refinement*, illustrated in Figure [REF](#). That is, as the spacing of the \mathbf{x}_i 's shrinks to zero, that is, as we sample an underlying smooth curve more and more densely, the approximated values of κ converge to the true curvature of the curve.

3.8.3 Turning Angle

Although convergence in the limit of refinement is a classical property reminiscent of Taylor series from elementary calculus, it neglects a key aspect of typical application scenarios: We may not have access to an arbitrarily dense sampling of a curve. Essentially the only information we have about a curve might be a sequence of points \mathbf{x}_i , gathered e.g. from the outside world, and our task is to measure curvature in this finitely-sampled regime without the option of collecting more data. We cannot, for example, use adaptive refinement to improve approximation when a curve locally fits the data poorly.

With this motivation in mind, recent methods in computational differential geometry have focused on the *discrete* behavior of measurements, before the spacing of the sample points is taken to zero. The resulting theory of discrete differential geometry aims to *redefine* notions like curvature directly on finitely-sampled data in a way that preserves features from the smooth case exactly. Of course, we also aim for our definitions of quantities like discrete curvature to be reasonable approximations of the smooth measurements.

As an initial clue toward structure-preserving discretization of curvature, recall the *exterior angle theorem* from Euclidean geometry, illustrated in Figure [REF](#). If we walk around the vertices of a polygon and sum up the exterior angles, we get 2π . A slight generalization of this classical result looks even more familiar: If the polygon is allowed to self-intersect, adding up the exterior angles gets 2π times the number of times the curve winds around itself.

The exterior angle theorem strongly resembles the Winding Number Theorem, Theorem [3.1](#). Furthermore, the exterior angle is constructed identically to the turning angle $\theta(s)$ in [\(3.6\)](#). Hence, it is completely natural to use the exterior angle at a vertex \mathbf{x}_k to discretize curvature.

Define θ_k to be the turning angle at vertex \mathbf{x}_k . As shown in Figure [REF](#), it does not quite make sense to discretize curvature as $\kappa_k := \theta_k$. In particular, this approximation does not behave similarly to the curvature of a smooth curve; as we refine our approximation of a smooth curve, all the θ_k 's approach zero!

To get around this issue, define Γ_k to be the *Voronoi region* of \mathbf{x}_k , shown in Figure [REF](#), which consists of the two line segments

$$\frac{\mathbf{x}_{k-1} + \mathbf{x}_k}{2} \rightarrow \mathbf{x}_k \rightarrow \frac{\mathbf{x}_k + \mathbf{x}_{k+1}}{2}.$$

In other words, Γ_k consists of two half-segments adjacent to \mathbf{x}_k . Then, we observe

$$\theta_i \approx \int_{\Gamma_k} \kappa(s) ds, \quad (3.16)$$

that is, the exterior angle resembles the *total curvature* along the segment Γ_k . This explains why θ_k vanishes as the lengths of the segments goes to zero, and it yields a *discrete winding number theorem*:

$$W[\gamma] = \frac{1}{2\pi} \sum_k \int_{\Gamma_k} \kappa ds = k.$$

This is our first example of *structure preservation* in a discrete computation: We found a version of curvature on a discrete curve that preserves the winding number theorem exactly.

Integrated curvature and pointwise curvature are of course not the same, and for many applications we may wish to approximate κ_k directly. Under the assumption that curvature is uniformly distributed along Γ_k , a reasonable estimator might be

$$\kappa_k \approx \frac{\theta_k}{L[\Gamma_k]}.$$

While integrated curvature makes complete sense for a poly-line, we might regard this formula as more of an approximation; it is not clear that pointwise curvature is reasonable or natural to measure when a poly-line is piecewise flat. This theme will appear and re-appear in our discussion: On discrete structures, many quantities from differential geometry only can be measured as integrated quantities.

3.8.4 Variation of Arc Length

In the previous section, we started with a structure from discrete geometry—the exterior angle theorem—which resembled a smooth structure—the winding number theorem, and we leveraged this connection to motivate a definition of curvature on a discrete structure. A natural question to ask is whether all roads lead to the same destination: If we choose a different connection between smooth and discrete geometry, do we get the same definition of curvature along a polyline?

An alternative starting point for adapting curvature to poly-lines might be to start from the first variation of arc length, equation (3.12). Treating our poly-line as a collection of line segments leads to a reasonable definition of arc length:

$$L[\{\mathbf{x}_k\}_{k=1}^p] := \sum_{k=1}^{p-1} \|\mathbf{x}_{k+1} - \mathbf{x}_k\|_2.$$

We can think of arc length in this case as a function $L[\mathbf{x}_1, \dots, \mathbf{x}_p]$ taking in an ordered collection of p points and giving the sum of the lengths of the segments between them.

For $k \in \{2, \dots, p-1\}$, we can differentiate to find

$$\nabla_{\mathbf{x}_k} L = \frac{\mathbf{x}_k - \mathbf{x}_{k-1}}{\|\mathbf{x}_k - \mathbf{x}_{k-1}\|_2} + \frac{\mathbf{x}_k - \mathbf{x}_{k+1}}{\|\mathbf{x}_k - \mathbf{x}_{k+1}\|_2}. \quad (3.17)$$

Examples of this vector are shown in Figure [REF](#). Notice it roughly points normal to the polyline at \mathbf{x}_i , and that its length shrinks to zero as the turning angle vanishes.

As derived in Exercise 3.6., however, the length of this vector at \mathbf{x}_k is

$$\|\nabla_{\mathbf{x}_k} L\|_2 = 2 \sin \frac{\theta}{2}.$$

Why is this problematic? From (3.12), we could reasonably use this length as a definition of the integrated curvature over Γ_k , the two half-segments centered at \mathbf{x}_k . But if we do this—thus preserving the property that curvature gives the gradient of arc length suggested in (3.12)—we *cannot* preserve the Winding Number Theorem in our discretization as in §3.8.3, since the latter requires taking integrated curvature to be θ . That is, there is no free lunch: Preserving one property of smooth curves in our discretization can actually contradict preserving another.

From a convergence perspective, however, we are still safe. In particular, for small θ we have $2 \sin \frac{\theta}{2} \approx \theta$ to first order. So, both of these discrete notions of curvature converge to the same measurement in the limit of refinement toward a smooth curve.

3.9 DISCRETE FRENET FRAMES IN THREE DIMENSIONS

JS: Omit for now, refer to slides for a short summary

3.10 CASE STUDY: DISCRETE ELASTIC RODS

JS: Omit for now, refer to [3]

3.11 EXERCISES

- 3.1. Show that every differentiable curve is a parameterized curve.
- 3.2. Suppose C is a parameterized curve, and take $\gamma : (a, b) \rightarrow C$ and $\tilde{\gamma} : (\tilde{a}, \tilde{b}) \rightarrow C$ to be two parameterizations of C . Show that there exists a bijective, smooth $\phi : (\tilde{a}, \tilde{b}) \rightarrow (a, b)$ such that $\tilde{\gamma}(t) = \gamma(\phi(t))$ and $|\phi'(t)| > 0$.
- 3.3. **JS: Something justifying radius of curvature**
- 3.4. **JS: derivative of the norm of a vector**
- 3.5. **JS: Frenet frames in \mathbb{R}^n with $n > 3$.**
- 3.6. Show that the length of the gradient vector $\nabla_{\mathbf{x}_k} L$ in (3.17) is $2 \sin \frac{\theta}{2}$, where θ is the exterior angle at vertex \mathbf{x}_k .
- 3.7. **JS: Smoke ring flow $\dot{\gamma} = \kappa B$**