The Levi-Civita Connection

Kirk Bonney

1 Project Overview

The aim of this project is to investigate and understand the Levi-Civita connection. We motivate this exploration in two ways. The first is by viewing the Levi-Civita as a necessary component for the generalization of curvature to manifolds. The other will focus on a simpler geometric concept - the straight line. We will demonstrate two different methods of generalizing the straight line to manifolds; it will turn out that the Levi-Civita connection is the natural way to reconcile these two points of view. Finally, we will end by stating the Fundamental Theorem of Riemannian Geometry, the theorem guaranteeing the existence and uniqueness of the Levi-Civita connection.

2 Curvature

To motivate the Levi-Civita connection and its relation to curvature, let's begin by reminding ourselves of some of the more basic perspectives on curvature. A definition of curvature given in a course on multi-variable calculus might look like the following:

Given a smooth, regular curve $\vec{\gamma}(t) : \mathbb{R} \to \mathbb{R}^3$ (i.e. one such that $\vec{\gamma}'(t)$ is continuous and $\vec{\gamma}'(t) \neq 0$) we define curvature to be

$$\kappa = \left\| \frac{\vec{T}'(t)}{\vec{\gamma}'(t)} \right\|$$

Where $\vec{T}$ denotes the unit tangent vector. Typically we assume that $|\vec{\gamma}'(t)| = 1$ since we can divide by $\|\vec{\gamma}'(t)\|$, a quantity we required to be non-zero. With this assumption in mind, it follows that $\vec{T} = \vec{\gamma}'$. Hence, the formula above can be reduced to

$$\kappa = \|\gamma''(t)\| .$$

At a particular point p on a smooth plane curve, we can find circles which are tangent to the curve at p. If we further require that the circles have the same velocity and acceleration vectors as the curve at p, we have a unique choice called the osculating circle (fun fact: osculating is a synonym for kissing). One can generalize this idea to space curves and osculating spheres with the appropriate details. This point of view encapsulates the historical approach to curvature taken by Oresme and Cauchy. We can relate this to the derivative approach to curvature in the following equation

$$\kappa = \frac{1}{R}$$

where κ is the curvature at a point, and R is the radius of the osculating circle at a point. It is worth noting that we can define an orientation of curvature to give us a notion of signed curvature. The idea being that curving 'inwards' is positive, and curving 'away' is negative.

In the grand scheme of smooth manifolds, we have kept things relatively simple here. We have only concerned ourselves with one-dimensional manifolds embedded in low-dimensional Euclidean space. What we gain from these tight restrictions is a very simple description of curvature: a single positive real number associated to each point on the curve. A natural question to ask about this definition of curvature is whether or not it is 'intrinsic' to the curve - that is, does a curve know about its curvature without reference to an embedding. The answer is no for any 1-dimensional manifold. To contrast this, we'll take a look at the curvature of surfaces. There we will see that one step up in dimension will already make things more complicated, however, thanks to Gauss, we will be able to define an embedding-independent (intrinsic) notion of curvature.

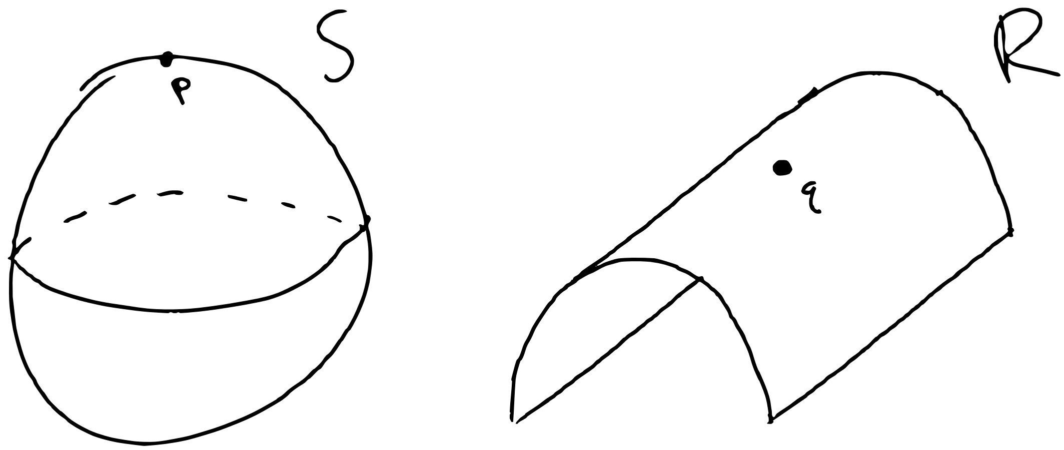

To get ourselves started, we'll consider two surfaces, S and R embedded in $\mathbb{R}^3$. For now focus on the 2-sphere S. We would like to come up with some notion of curvature on S at the distinguished point p. Since we already have curvature for curves established, it would be nice if we could leverage that theory here. So we would like to think about curves in S passing through the point p. We can obtain a family of such curves by considering planes containing the unit normal vector at p to S. The intersection of one of these planes with S gives us a curve in $\mathbb{R}^3$, something we can calculate the the curvature for. The case of the sphere is convenient since the curvature of each of these curves is the same (the osculating sphere for each curve is just S). It would be fair to define the curvature of S at p to be the shared curvature of the curves we defined. However, what we actually do is something a bit different. Examining R will enlighten us to why.

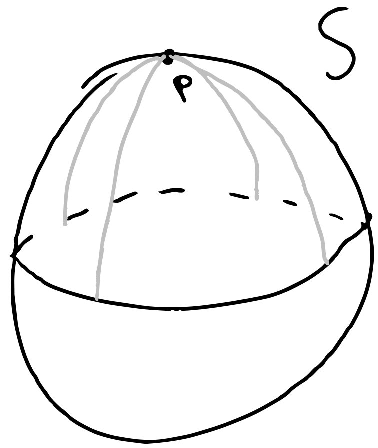

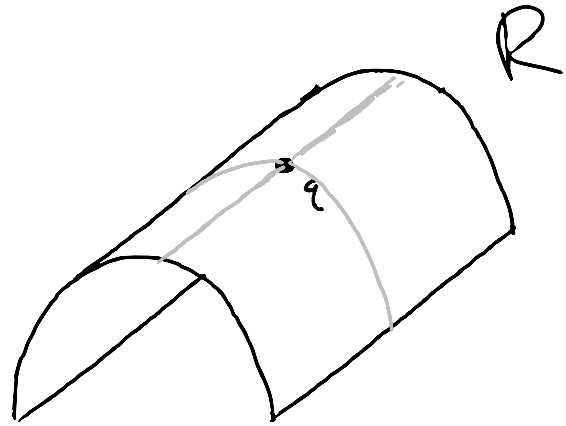

We proceed with the same strategy for the curvature of R at q. Those with keen visualization ability will quickly note that as we rotate our intersecting plane about the normal vector, we get drastically different curves. Along one plane we obtain a straight line, whereas the perpendicular plane gives us a parabola. These will certainly give us different curvatures, so we can't proceed as we did with S. Since the curvatures vary, it would be natural to think about the minimum and maximum of the set of curvatures. This is certainly well-defined for a fixed embedding, and the pair of values are known as the principal curvatures of R at q. However, this is not quite enough to give us something intrinsic.

The right thing to define is due to Gauss and his *Theorema Egregium* which demonstrates that the product of the principal curvatures is an intrinsic property of a 2-dimensional manifold. This value is called the Gaussian curvature. Consider the fact that R may be diffeomorphically deformed into a flat surface. Then the intrinsic curvature of the plane and R should be identical. R will have 0 either as a min or max of its curvatures at q, depending on the orientation chosen to sign the curvature, since the straight line is one of the possible curves. Therefor, the Gaussian curvature of R at q must be 0. This is the case for the plane, as its only curves obtained from normal, intersecting planes are lines.

In the following sections, we will take a tour through a collection of concepts from Riemannian geometry so that we may state and understand the Levi-Civita connection. From this perspective, one can define sectional curvature for a Riemannian manifold. Gaussian curvature depends only on the Riemannian metric, which in our examples were endowed on our manifolds by the embedding.

3 A Starting Point for Geometry

Our goal is to start making sense of geometric notions on a general manifold, i.e. one which is not embedded in $\mathbb{R}^n$. However, a manifold alone doesn't have much more shape to it than the underlying topological space. The atlas guarantees smoothness, but the manifold is essentially formless. To convince yourself that an ordinary manifold really can't have an inherent notion of geometry, recall that half of S_2 is diffeomorphic to an open set of R^2, but the two manifolds have very different geometries. For example, triangles on a sphere have a total angle sum greater than π.

One of the most straightforward ways to give our manifolds shape is with a Riemannian metric. This is how we turn $\mathbb{R}^n$ into the familiar realm of Euclidean geometry. Recall the usual Euclidean inner product is given by

$$\langle x, y\rangle = \sum x_i y_i$$

from which we can derive the Euclidean metric $d(x, y) = \sqrt{\langle x - y, x - y\rangle}$. The tangent spaces of $\mathbb{R}^n$ inherit this inner product since we can identify them with $\mathbb{R}^n$. However, this luxury does not extend to manifolds in general. This problem leads us to define a Riemannian metric.

Definition 3.1 (Lee)**.** A **Riemannian metric** on a manifold M is a smooth, covariant 2-tensor field g, whose value g_p at each $p \in M$ is an inner product on T_pM.

For example, we can define a Riemannian metric on $\mathbb{R}^n$ by setting $g_p = \langle\cdot,\cdot\rangle$. Once we equip a manifold with a Riemannian metric (from now on, we'll simply call it a metric), we immediately obtain a notion of distance and angles for our tangent spaces - a great start for geometry on manifolds. Conveniently, it is a fact that all manifolds can admit a metric. As we will shortly see in an example, there is a straightforward way to form a metric on an embedded manifold. We can combine this method with the Whitney embedding to convince ourselves of the aforementioned fact, though there are ways to prove it without needing such a powerful theorem.

Example 3.1. Given an embedded manifold $f : M \to \mathbb{R}^n$, we can pull back the metric $r(\cdot,\cdot)$ on $\mathbb{R}^n$ to the manifold, endowing M with the metric

$$g_p(x, y) = (f^*r)_p(x, y) = r_{f(p)}(df(x), df(y)).$$

In the particular case of $\iota : S^1 \to \mathbb{R}^2$, the inherited metric is really just the restriction of the metric on $\mathbb{R}^2$ to T_pS^1 at each point p. In symbols,

$$g_p(x, y) = (\iota^*r)_p(x, y) = r_{\iota(p)}(d\iota(x), d\iota(y)) = r_p(x, y).$$

This example generalizes in a straightforward way to an immersion $f : M \to N$ when N is already a Riemannian manifold. We can build new metrics from old ones using products and submersions as well, but instead of going into details on these, let's look at an explicit way we can pull geometry out of Riemannian manifold.

Definition 3.2. [3] Let (M, g) be a Riemannian manifold. For a smooth curve $\gamma : [a, b] \to M$ we define its length

$$L_g(\gamma) = \int_a^b \sqrt{\langle\gamma'(t), \gamma'(t)\rangle_g}\, dt.$$

We can then define the g-distance on M between two points $x, y \in M$ to be the greatest lower bound of the set $\{L_g(\gamma) \mid \gamma(a) = x,\ \gamma(b) = y\}$.

What does this look like on $\mathbb{R}^n$ with the usual metric? Hopefully, there aren't any surprises.

Example 3.2. Let $\gamma : [a,b] \to \mathbb{R}^2$ be a smooth curve and write $\gamma(t) = (\gamma_1(t), \gamma_2(t))$. Then

$$\begin{aligned} L_g(\gamma) &= \int_a^b \sqrt{\langle \gamma'(t), \gamma'(t) \rangle}_g dt \\ &= \int_a^b \sqrt{\langle (\gamma_1'(t), \gamma_2'(t)), (\gamma_1'(t), \gamma_2'(t)) \rangle}_g dt \\ &= \int_a^b \sqrt{\gamma_1'(t)^2 + \gamma_2'(t)^2}_g dt. \end{aligned}$$

This is exactly how we define the length of a curve in calculus, so we've given ourselves a proper generalization of length. With this in hand, we can generalize lines to Riemannian manifolds by requiring a line between two points to be a curve with length equal to the g-distance between the points. We will explore this concept in further detail in the next section and answer questions like 'is such a curve unique?'

4 Generalizing the straight line without a metric

In this section we will approach the problem of generalizing the straight line of Euclidean space from a different point of view. We will not start with a Riemannian manifold, but instead use a covariant derivative to allow us to generalize the zero-acceleration property of $ax + b$. Covariant derivatives are obtained from an affine connection, and we will give rigorous definitions for both. However, let's first justify why we need a covariant derivative: what does it give us, and why can't we proceed without it?

Recall the discussion of elementary curvature in the first section. There we indicated that curvature in $\mathbb{R}^n$ can be realized as the magnitude of an acceleration vector. If we want to generalize this formulation of curvature, we certainly need a way to obtain acceleration vectors on an arbitrary manifold. Let's try.

As usual, a first attempt at naive generalization follows from trying what works in $\mathbb{R}^n$ on an arbitrary manifold. Given a curve $\gamma(t)$, what we do in $\mathbb{R}^n$ to obtain $\gamma''(t)$ is to simply differentiate $\gamma'(t)$. Up until manifold theory, a student might think of $\gamma'(t)$ as another function with values in $\mathbb{R}^n$, so really it is no problem to differentiate it. However, for our purposes, it will be better to view $\gamma(t)$ as a vector field along a curve, $\gamma'(t) : [a,b] \to TM$. In Euclidean space we derive one vector field with respect to another in the following way

$$(\nabla_X Y)(p) := \lim_{t \to 0} \frac{Y(p + tX(p)) - Y(p)}{t}.$$

This works because of our natural identification of each tangent space with $\mathbb{R}^n$, so we can make sense of the quantities $p + tX(p)$ and $Y(p + tX(p)) - Y(p)$, despite their terms living in different worlds. However, we don't have such an identification in general, so these quantities are going to cause us trouble. The first can be remedied using flows similar to what one does in the case of the Lie derivative. In other words, we can write $p + tX(p)$ as $\phi_t(p)$ where ϕ_t is the local flow of X. When building the Lie derivative, the second problem is also solved with flows, but we are not going to do that. To see why, recall from section 3 that we can't have a reliable notion of geometry on a manifold since diffeomorphisms won't preserve them.

So if we try to build a geometry without introducing something new, we are doomed to fail. Furthermore, the Lie derivative can't be thought of as a proper generalization of the ∇ defined above in terms of basic properties. As we'll see in the coming definition, ∇ is C^∞-linear in the first argument, but this is not true for the Lie derivative.

The moral is that we need something else to get our geometry off the ground. A connection turns out to be the framework for the task. As we'll see, a connection lets us directly define a ∇ that works like the one for $\mathbb{R}^n$ in other manifolds without mucking about in technical details too much (unless you want to: see connection coefficients).

Definition 4.1. [4]

Let $\pi : E \to M$ be a smooth vector bundle over a smooth manifold M. We write $\mathfrak{X}(M)$ to denote the set of vector fields on M and $\Gamma(M)$ to denote the space of smooth sections of E. A **connection in E** is a map

$$\nabla : \mathfrak{X}(M) \times \Gamma(E) \to \Gamma(E)$$

written $(X, Y) \mapsto \nabla_X Y$, satisfying the following properties:

1. $\nabla_X Y$ is linear over $C^\infty(M)$ in X : for $f_1, f_2 \in C^\infty(M)$ and $X_1, X_2 \in \mathfrak{X}(M)$,

$$\nabla_{f_1 X_1 + f_2 X_2} Y = f_1 \nabla_{X_1} Y + f_2 \nabla_{X_2} Y$$

2. $\nabla_X Y$ is linear over $\mathbb{R}$ in Y : for $a_1, a_2 \in \mathbb{R}$ and $Y_1, Y_2 \in \Gamma(E)$,

$$\nabla_X (a_1 Y_1 + a_2 Y_2) = a_1 \nabla_X Y_1 + a_2 \nabla_X Y_2$$

3. ∇ satisfies the following product rule: for $f \in C^\infty(M)$,

$$\nabla_X (fY) = f \nabla_X Y + (Xf) Y$$

Sometimes ∇ is referred to as a **covariant derivative**, emphasizing its role as a differential operator. This definition is slightly more general than what we actually need. We will be interested in connections where $E = TM$, which we call **affine connections**. In particular, an affine connection is a map

$$\mathfrak{X}(M) \times \mathfrak{X}(M) \to \mathfrak{X}(M)$$

which satisfies the properties in the above definition. It is good to acknowledge that this is a very top-down approach to mathematics. We know how we would like ∇ to behave since we are modeling it off of the directional derivative of vector fields in $\mathbb{R}^n$, so we select a handful of key properties from the case in $\mathbb{R}^n$ and declare that any map satisfying them is a candidate for an analogous operator on our manifold. Indeed, you can check that properties 1-3 are satisfied by the affine connection defined earlier for $\mathbb{R}^n$.

The way to 'get your boots on the ground' with an affine connection is via connection coefficients.

Definition 4.2. The **connection coefficients** of an affine connection with respect to a chart ϕ with tangent vectors $\partial_1, \ldots, \partial_n$ are defined by

$$\nabla_{\partial_i}\partial_j = \Gamma_{ij}^k \partial_k.$$

The connection coefficients capture how the bases of the tangent space interact with each other, so it shouldn't be hard to believe that they fully determine the connection on the chart. In practice, connection coefficients are frequently used to prove results and to get your hands on explicit examples.

Example 4.1. Recall the definition of ∇ given for $\mathbb{R}^n$. Then

$$\nabla_{\partial_i}\partial_j = \lim_{t\to\infty} \frac{\partial_j(p + t\partial_i(p)) - \partial_j(p)}{t} = \frac{\partial_j - \partial_j}{t} = 0$$

So the usual geometry of $\mathbb{R}^n$ is characterized by global connection coefficients of 0.

To see that there exists at least one connection on any manifold, one can use a partition of unity argument to construct a connection induced by the standard connection on $\mathbb{R}^n$. The idea being that we can locally obtain a connection from $\mathbb{R}^n$ at all points of M, and we can use a partition of unity to patch these together into a full connection. However, a general manifold might have many possible affine connections, but we will see later that a Riemannian metric gives us a way to choose one uniquely. Before we get there, we will work with affine connections a bit more and see how we obtain a notion of a straight line for a manifold with an affine connection. To get started, let's make rigorous the concept of a vector field along a curve.

Definition 4.3 (cite Ben Andrews)**.** Let $\gamma : I \to M$ be a smooth curve. A **vector field along** γ is a smooth map $V : I \to TM$ with $V(t) \in T_{\gamma(t)}M$ for each t. The set of all smooth vector fields along γ is denoted $\mathfrak{X}_\gamma(M)$.

An immediate example is $\gamma'(t)$. We defined an affine connection as a map of vector fields, but we can use it to define an operator on vector fields along curves as the following proposition shows.

Proposition 4.1. *[1] Let M be a smooth manifold and ∇ an affine connection on M. Then for any smooth curve $\gamma : I \to M$ there is a natural covariant derivative along γ, denoted ∇_t, which takes a smooth vector field along γ to another, while satisfying*

1. $\nabla_t(V + W) = \nabla_t V + \nabla_t W$ *for all* V *and* W *in* $\mathfrak{X}\gamma(M)$;
2. $\nabla_t(fV) = f'V + f\nabla_t V$ *for all* $f \in C^\infty(I)$ *and* $V \in \mathfrak{X}_\gamma(M)$;
3. *If* V *is given by the restriction of a vector field* $\tilde{V} \in \mathfrak{X}(M)$ *to* γ*, then* $\nabla_t V = \nabla_{\dot\gamma}\tilde{V}$*, where* $\dot\gamma$ *is the tangent vector to* γ*.*

The first two properties give us a form of linearity and a product rule, inherited from ∇. The third property is more interesting. A fact implicitly used in the statement of property three is that $\nabla_X Y(p)$ only relies on the value of X at p, and the values that Y takes on in a neighborhood of p. This is why it makes sense to write $\nabla_v Y$ for v a vector rather than a vector field. In fact, sometimes affine connections are defined in this way, as a map with domain $TM \times \mathfrak{X}(M)$. So then, property three is strengthening this fact to $\nabla_X Y(p)$ depending only on the value of X at p and the values Y takes on the neighborhood of some curve through p.

Recall that the problem we began this section with was to generalize the straight line of $\mathbb{R}^n$ via the property of zero acceleration. Now that we can make sense of ∇ for curves, rather than whole fields, we properly define what acceleration actually looks like in a manifold with connection.

Definition 4.4. [4] Let $\gamma : I \to M$ be a smooth curve on a manifold M with connection ∇. Define the **acceleration of** γ to be the vector field $\nabla_t \gamma'$ along γ.

Now we are equipped to say that straight line for M with connection ∇ is a curve γ with acceleration zero. This is captured in the notion of parallel transport.

Definition 4.5. Let $\gamma : I \to M$ be a smooth curve. A vector field $V \in \mathfrak{X}_\gamma(M)$ is said to be **parallel** along γ if $\nabla_t V = 0$.

In other words, a straight line in M is a curve γ such that $\nabla_t \gamma'$ is parallel along γ. As promised, an affine connection on a general manifold gives us an alternative approach to generalizing the straight lines. We call these curves **geodesics**. Since we have two formulations of geodesics, we want to know how to make them agree with each other. The Fundamental Theorem of Riemannian Geometry will answer this question.

5 Levi-Civita

We have now established two different approaches to geometry on a smooth manifold. The first involved a symmetric 2-tensor called a Riemannianc metric which gives us a smoothly varying inner-product for each T_pM, and the second involved a choice of how to differentiate a vector field with respect to another, which we called an affine connection. The Levi-Civita connection is the connection that unites these two views, and its existence is the content of the following theorem.

Theorem 5.1. *[4](The Fundamental Theorem of Riemannian Geometry).*
Let (M, g) be a Riemannian manifold. There exists a unique connection ∇ on TM that is compatible with g and symmetric. It is called the Levi-Civita connection of g (or also, when g is positive definite, the Riemannian connection).

To fully understand this theorem, we need to break down what is meant by compatibility and symmetry. We once again look to $\mathbb{R}^n$ for guidance on how to structure our generalized theory. It is a fact of the standard connection and Riemannian metric on $\mathbb{R}^n$ that

$$\nabla_X \langle Y, Z \rangle = \langle \nabla_X Y, Z \rangle + \langle Y, \nabla_X Z \rangle.$$

So *nabla* and $\langle \cdot, \cdot \rangle$ satisfy a product rule. We generalize this with the following definition.

Definition 5.1. Given a Riemannian manifold (M, g), we say a connection ∇ on M is **compatible with the metric on M** if for all $X, Y \in \mathfrak{X}(M)$, and every $v \in T_xM$,

$$v(g(X,Y)) = g(\nabla_v X, Y) + g(X, \nabla_v Y).$$

It is true that a Riemannian manifold can have many compatible connections, so we need further constraints for the uniqueness of the theorem to hold. For this we turn to the concept of torsion.

Definition 5.2. Given a manifold M with connection ∇, we define the **torsion tensor** T by

$$T(X,Y) = \nabla_X Y - \nabla_Y X - [X,Y]$$

We can understand symmetry in terms of the torsion tensor.

Definition 5.3. A connection ∇ on M is said to be **symmetric** if for all $X, Y \in \mathcal{X}(M)$,

$$\nabla_X Y - \nabla_Y X = [X,Y].$$

In other words, a symmetric connection is one whose torsion tensor vanishes. With these definitions in hand, we can prove the uniqueness part of the theorem.

Proof. [1] Let (M, g) be a Riemannian manifold and let ∇ be a symmetric connection which is compatible with g. The symmetry condition gives us

$$\begin{aligned} \nabla_X Y - \nabla_Y X &= [X,Y]; \\ \nabla_Y Z - \nabla_Z Y &= [Y,Z]; \\ \nabla_Z X - \nabla_X Z &= [Z,X]. \end{aligned}$$

The compatibility condition gives us

$$\begin{aligned} g(\nabla_X Y, Z) + g(Y, \nabla_X Z) &= Xg(Y,Z); \\ g(\nabla_Y Z, X) + g(Z, \nabla_Y X) &= Yg(Z,X); \\ g(\nabla_Z X, Y) + g(X, \nabla_Z Y) &= Zg(X,Y). \end{aligned}$$

Taking the sum of the first two of the second set of equations and subtracting the third, we obtain

$$\begin{aligned} &g(\nabla_X Y, Z) + g(Y, \nabla_X Z) + g(\nabla_Y Z, X) + g(Z, \nabla_Y X) \\ &- g(\nabla_Z X, Y) - g(X, \nabla_Z Y) = Xg(Y,Z) + Yg(Z,X) - Zg(X,Y). \end{aligned}$$

Then we use the symmetry and linearity of g to arrange the left hand side as

$$g(X, \nabla_Y Z - \nabla_Z Y) + g(Y, \nabla_X Z - \nabla_Z X) + g(Z, \nabla_Y X + \nabla_X Y).$$

Next, we substitute the first set of equations

$$g(X, [Y,Z]) + g(Y, -[Z,X]) + 2g(Z, \nabla_Y X) + g(Z, [X,Y]).$$

Finally, we can write

$$g(Z, \nabla_Y X) = (Xg(Y,Z) + Yg(Z,X) - Zg(X,Y) \\ + g(X,[Z,Y]) + g(Y,[Z,X]) + g(Z,[Y,X]))/2.$$

The punchline is that the g inner-product of $\nabla_Y X$ with any vector field Z is written in terms of g, which fully determines ∇. Thus if such a ∇ exists, it is unique.

We won't prove existence here, but the basic idea is to define a connection by the last equation above and show that it satisfies all the necessary definitions. □

It is a consequence of the Hopf-Rinow theorem that if (M, g) is a complete metric space, then a geodesic (defined in terms of the Levi-Civita connection) between any two points always exists and agrees with the distance function (defined in terms of the metric). That is, for any $p, q \in M$ there exists a $\gamma : [0, 1] \to M$ so that $\gamma(0) = p$ and $\gamma(1) = q$, $\gamma_t \gamma'$ along γ is 0, and the length of γ is minimal among all unit-speed curves from p to q. So we have found a way to reconcile our two points of view on the straight line!

To conclude our journey, we will define the Riemann curvature tensor to demonstrate an application of the Levi-Civita connection.

Definition 5.4. Let (M, g) be a Riemannian manifold. We define the **Riemann curvature tensor** $R(X, Y, Z)$ by

$$R(X,Y,Z) = \nabla_X \nabla_Y Z - \nabla_Y \nabla_X Z - \nabla_{[X,Y]}$$

where ∇ is the Levi-Civita connection for g.

One can show that the map is multilinear over $C^\infty(M)$, so it really is a $(1, 3)$-tensor. The details of this tensor field may as well be the subject of a whole other project, so we'll simply offer an intuitive interpretation of what is going on. The basic idea is that R captures the deviation of the geometry of M from the standard geometry on $\mathbb{R}^n$. One can imagine this in terms of parallel transporting a vector through $\mathbb{R}^3$ on a plane versus on 2-sphere along a closed path. On the plane, no matter what closed path is taken, one should expect the vector to remain unchanged when it returns from its journey. However, there are closed paths (a triangle works) on the sphere which do not have this property. This sort of difference can be derived from the Riemann curvature tensor.

In the case of any 1-dimensional manifold, R will be identically zero due to certain symmetries of the tensor - which is exactly what we should expect since 1-dimensional manifolds all look like $\mathbb{R}$ locally. In the 2-dimensional case, one can use R to generalize both the Gaussian and sectional curvatures which we defined at the beginning of the paper. Past this, there are many other kinds of notions of curvature one can extract from R: Ricci, Scalar, and Weyl to name a few. Curvature at this level of generality is not easy to describe, hence the need many different tools to capture the various properties of curvature.

Conclusion

We have ended our discussion with a brief glimpse at abstract curvature. To access these ideas in a rigorous way, one needs to understand the fundamentals of Riemannian geometry and the build-up to the Levi-Civita connection. There are a lot of details to pick up along the way, many of which were omitted from this paper. Hopefully the survey of Euclidean curvature in one and two dimensions helped motivate the quest to generalize curvature to any smooth manifold. The Levi-Civita connection is an important step towards this generalization, as it equips a Riemannian manifold with a connection that meshes with the already established metric structure, paving the way for a large flurry of curvature-related tensors to be defined.

References

[1] Ben Andrews. Lectures on differential geometry. `https://maths-people.anu.edu.au/~andrews/DG/`. Accessed: 4/01/2021.

[2] Sigurdur Helgason. *Differential Geometry, Lie Groups, and Symmetric Spaces*. Pure and applied mathematics. Elsevier Science & Technology, Burlington, 1978.

[3] Serge Lang. *Differential and Riemannian Manifolds*, volume 160 of *Graduate Texts in Mathematics*. Springer, 2012.

[4] John M. Lee. *Introduction to Riemannian manifolds*. Graduate texts in mathematics, 176. Springer, Cham, Switzerland, second edition. edition, 2018.

[5] J. Milnor, M. SPIVAK, and R. WELLS. *Morse Theory. (AM-51), Volume 51*. Princeton University Press, 2016.

[6] A. N Pressley. *Elementary Differential Geometry*. Springer undergraduate mathematics series. Springer London, Limited, London, 2010.

[7] Andre Wibisono. Levi-civita connection. `http://awibisono.github.io/2016/09/19/levi-civita.html`. Accessed: 4/01/2021.