Manifolds are spaces that "locally look like" $\mathbb{R}^n$ in various ways. Up to some technical assumptions, if we locally look like $\mathbb{R}^n$'s topological structure, we are a *topological* manifold. Differentiable structure? A *smooth* manifold. Inner product space structure? A *Riemannian* manifold. Before unraveling what *locally* and *looks like* means, let us first explore what kinds of standard structures $\mathbb{R}^n$ is equipped with.

First, $\mathbb{R}^n$ is a (real) vector space. In particular, we can add and scale (by real numbers) elements of $\mathbb{R}^n$. We have the following properties (for all $\mathbf{u}, \mathbf{v}, \mathbf{w} \in \mathbb{R}^n$ and $s, t \in \mathbb{R}$):

- Closure under addition: $\mathbf{v} + \mathbf{w} \in \mathbb{R}^n$ and closure under scalar multiplication: $s\,\mathbf{v} \in \mathbb{R}^n$
- Associativity of addition: $(\mathbf{u} + \mathbf{v}) + \mathbf{w} = \mathbf{u} + (\mathbf{v} + \mathbf{w})$ and compatibility of multiplication of scalars and scalar multiplication: $(st)\mathbf{v} = s(t\mathbf{v})$
- Additive identity: $\mathbf{0} + \mathbf{v} = \mathbf{v} = \mathbf{v} + \mathbf{0}$ and multiplicative identity: $1\,\mathbf{v} = \mathbf{v}$
- Additive inverses: $\mathbf{v} + (-\mathbf{v}) = \mathbf{0} = (-\mathbf{v}) + \mathbf{v}$
- Commutativity: $\mathbf{v} + \mathbf{w} = \mathbf{w} + \mathbf{v}$ and somewhat unofficially: $\mathbf{v}\,s = s\,\mathbf{v}$
- Distributivity: $s(\mathbf{v} + \mathbf{w}) = s\,\mathbf{v} + s\,\mathbf{w}$ and $(s + t)\mathbf{v} = s\,\mathbf{v} + t\,\mathbf{v}$

This vector space structure allows us to do *linear algebra* stuff, but it does not really help us do real geometry or analysis. However, in addition to $\mathbb{R}^n$'s vector space structure, it has *inner product space* structure. In particular, $\mathbb{R}^n$ is equipped with the *dot product*:

$$(v_1, v_2, \dots v_n) \bullet (w_1, w_2, \dots, w_n) = v_1 w_1 + v_2 w_2 + \cdots + v_n w_n$$

This is an example of an inner product. In particular, we have (for all $\mathbf{u}, \mathbf{v}, \mathbf{w} \in \mathbb{R}^n$ and $s \in \mathbb{R}$):

- Bilinearity: $(\mathbf{u} + \mathbf{v}) \bullet \mathbf{w} = \mathbf{u} \bullet \mathbf{w} + \mathbf{v} \bullet \mathbf{w}$, $\mathbf{u} \bullet (\mathbf{v} + \mathbf{w}) = \mathbf{u} \bullet \mathbf{v} + \mathbf{u} \bullet \mathbf{w}$, and $(s\,\mathbf{v}) \bullet \mathbf{w} = s(\mathbf{v} \bullet \mathbf{w}) = \mathbf{v} \bullet (s\,\mathbf{w})$
- Symmetry: $\mathbf{v} \bullet \mathbf{w} = \mathbf{w} \bullet \mathbf{v}$
- Positive definiteness: $\mathbf{v} \bullet \mathbf{v} \geq 0$ and $\mathbf{v} \bullet \mathbf{v} = 0$ only if $\mathbf{v} = \mathbf{0}$

An inner product is quite powerful. It allows us to define concepts such as length and angle. In particular, because of positive definiteness, we can define (for all $\mathbf{v} = (v_1, v_2, \dots v_n) \in \mathbb{R}^n$):

$$\|\mathbf{v}\| = \sqrt{\mathbf{v} \bullet \mathbf{v}} = \sqrt{v_1^2 + v_2^2 + \cdots + v_n^2}$$

to be the *length* (= *magnitude* = *norm*) of $\mathbf{v}$. In general, a vector space equipped with a mapping into the reals such that (for all vectors $\mathbf{v}, \mathbf{w}$, and scalars $s \in \mathbb{R}$):

- Positive definiteness: $\|\mathbf{v}\| \geq 0$ and $\|\mathbf{v}\| = 0$ only if $\mathbf{v} = \mathbf{0}$
- Scalar multiplication compatibility: $\|s\,\mathbf{v}\| = |s|\,\|\mathbf{v}\|$
- Triangle inequality: $\|\mathbf{v} + \mathbf{w}\| \leq \|\mathbf{v}\| + \|\mathbf{w}\|$

is called a *normed space*. It is easy to see (except for the triangle inequality) that an inner product space is also a normed space when we define length as we did above. We prove the following rather important result which in turn lets us establish the triangle inequality with ease:

Theorem (Cauchy-Schwarz Inequality): For all $\mathbf{v}, \mathbf{w} \in \mathbb{R}^n$, $|\mathbf{v} \bullet \mathbf{w}| \leq \|\mathbf{v}\| \, \|\mathbf{w}\|$.

Proof: Let x be an arbitrary real number. Consider $0 \leq \|x\mathbf{v} + \mathbf{w}\|^2 = (x\mathbf{v} + \mathbf{w}) \bullet (x\mathbf{v} + \mathbf{w}) = x^2(\mathbf{v} \bullet \mathbf{v}) + x(\mathbf{v} \bullet \mathbf{w}) + x(\mathbf{w} \bullet \mathbf{v}) + \mathbf{w} \bullet \mathbf{w}$ where we used bilinearity to pull apart the dot product. Next, we can use symmetry to put the middle terms together. Therefore, we have a quadratic in x:

$$y = \|\mathbf{v}\|^2 x^2 + 2(\mathbf{v} \bullet \mathbf{w})x + \|\mathbf{w}\|^2 \quad (= \|x\mathbf{v} + \mathbf{w}\|^2 \geq 0)$$

Being non-negative, our quadratic cannot have distinct real roots. Therefore, the discriminant (for $Ax^2 + Bx + C$ the discriminant is $B^2 - 4AC$) must be non-positive: $(2\,\mathbf{v} \bullet \mathbf{w})^2 - 4\|\mathbf{v}\|^2\|\mathbf{w}\|^2 \leq 0$ so that $4(\mathbf{v} \bullet \mathbf{w})^2 \leq 4(\|\mathbf{v}\| \cdot \|\mathbf{w}\|)^2$. Thus $(\mathbf{v} \bullet \mathbf{w})^2 \leq (\|\mathbf{v}\| \cdot \|\mathbf{w}\|)^2$ so $|\mathbf{v} \bullet \mathbf{w}| \leq \|\mathbf{v}\| \, \|\mathbf{w}\|$. ■

The triangle inequality easily follows from Cauchy-Schwarz:

$$\|\mathbf{v} + \mathbf{w}\|^2 = \mathbf{v} \bullet \mathbf{v} + \mathbf{v} \bullet \mathbf{w} + \mathbf{w} \bullet \mathbf{v} + \mathbf{w} \bullet \mathbf{w} = \|\mathbf{v}\|^2 + 2(\mathbf{v} \bullet \mathbf{w}) + \|\mathbf{w}\|^2$$

$$\leq \|\mathbf{v}\|^2 + 2|\mathbf{v} \bullet \mathbf{w}| + \|\mathbf{w}\|^2 \leq \|\mathbf{v}\|^2 + 2\|\mathbf{v}\|\|\mathbf{w}\| + \|\mathbf{w}\|^2 = (\|\mathbf{v}\| + \|\mathbf{w}\|)^2$$

where we used $\mathbf{v} \bullet \mathbf{w} \leq |\mathbf{v} \bullet \mathbf{w}|$ (i.e., $x \leq |x|$) and then the Cauchy-Schwarz inequality. Square-rooting yields: $\|\mathbf{v} + \mathbf{w}\| \leq \|\mathbf{v}\| + \|\mathbf{w}\|$. ■

The Cauchy-Schwarz inequality implies that for nonzero vectors $\mathbf{v}$ and $\mathbf{w}$, we have $\dfrac{|\mathbf{v} \bullet \mathbf{w}|}{\|\mathbf{v}\| \, \|\mathbf{w}\|} \leq 1$. This means that $-1 \leq \dfrac{\mathbf{v} \bullet \mathbf{w}}{\|\mathbf{v}\| \, \|\mathbf{w}\|} \leq 1$ so that $\theta = \arccos\left(\dfrac{\mathbf{v} \bullet \mathbf{w}}{\|\mathbf{v}\| \, \|\mathbf{w}\|}\right)$ makes sense. In other words, we can define the notion of *angle between* vectors (leaving the angle between the zero vector and other vectors is undefined) using our inner product. In the reverse direction, a standard argument using the law of cosines shows that this "definition" of angle between vectors matches with our usual geometric definition of angle. In general, even when we involve the zero vector, we have our familiar identity: $\mathbf{v} \bullet \mathbf{w} = \|\mathbf{v}\| \, \|\mathbf{w}\| \cos(\theta)$ where θ is the angle between $\mathbf{v}$ and $\mathbf{w}$.

Next, normed space structure leads to metric space structure. Be careful: in manifold theory, a *metric* is a smoothly varying inner product assigned to tangent spaces (whatever that means). On the other hand, in our immediate setting, a *metric* refers something different:

Definition: A set X equipped with a mapping $d : X \times X \to \mathbb{R}$ (called a *metric* or a distance function) such that (for all $x, y, z \in X$):

- Positivity: $d(x, y) \geq 0$ and $d(x, y) = 0$ if and only if $x = y$
- Symmetry: $d(x, y) = d(y, x)$
- Triangle inequality: $d(x, z) \leq d(x, y) + d(y, z)$

is called a **metric space**.

Using our norm (i.e., vector length), we can turn $\mathbb{R}^n$ into a metric space. In particular, we have the familiar distance function (for $\mathbf{x} = (x_1, \ldots, x_n)$ and $\mathbf{y} = (y_1, \ldots, y_n)$):

$$d(\mathbf{x}, \mathbf{y}) = \|\mathbf{x} - \mathbf{y}\| = \sqrt{(x_1 - y_1)^2 + (x_2 - y_2)^2 + \cdots + (x_n - y_n)^2}$$

Proof: Notice $d(\mathbf{v}, \mathbf{w}) = \|\mathbf{v} - \mathbf{w}\| \geq 0$. This is 0 if and only if $\mathbf{v} - \mathbf{w} = \mathbf{0}$ (i.e., $\mathbf{v} = \mathbf{w}$). Symmetry is easy: $d(\mathbf{v}, \mathbf{w}) = \|\mathbf{v} - \mathbf{w}\| = \| - (\mathbf{w} - \mathbf{v})\| = | - 1| \cdot \|\mathbf{w} - \mathbf{v}\| = d(\mathbf{w}, \mathbf{v})$. The triangle inequality follows from its norm version: $d(\mathbf{u}, \mathbf{w}) = \|\mathbf{u} - \mathbf{w}\| = \|\mathbf{u} - \mathbf{v} + \mathbf{v} - \mathbf{w}\| \leq \|\mathbf{u} - \mathbf{v}\| + \|\mathbf{v} - \mathbf{w}\| = d(\mathbf{u}, \mathbf{v}) + d(\mathbf{v}, \mathbf{w})$. ■

Note: Working in $\mathbb{R}^2$, we have $d((x_1, x_2), (y_1, y_2)) = \sqrt{(x_1 - y_1)^2 + (x_2 - y_2)^2}$ and in $\mathbb{R} = \mathbb{R}^1$, we have $d(x, y) = \sqrt{(x - y)^2} = |x - y|$.

Completeness: A metric space space is called *complete* if every Cauchy sequence converges.[1] A normed space whose corresponding metric structure is complete is called a *Banach space*. An inner product space whose corresponding metric structure is complete is called a *Hilbert space*. One usually wants (at a minimum) Banach structure to do calculus stuff. We note that $\mathbb{R}^n$ is complete (i.e., it is a Hilbert space and thus Banach space too).

With distances in hand, we have a notion of *closeness*. This lets us define *topological* structure.

Definition: A set X equipped with a collection of subsets $\mathcal{T}$ (called a *topology* on X) such that:

- The empty set $\emptyset$ and X itself belong to $\mathcal{T}$
- Finite intersections belong: If A and B belong to $\mathcal{T}$ then so does $A \cap B$
- Arbitrary unions belong: Suppose A_i is in $\mathcal{T}$ (for all i in some index set I). Then their union $\bigcup_{i \in I} A_i = \{x \in X \mid x \in A_i \text{ for some } i \in I\}$ also belongs to $\mathcal{T}$

is called a **topological space**. The subsets of X belonging to $\mathcal{T}$ are called *open sets*.

Let $\mathbf{x}_0 \in \mathbb{R}^n$ and ϵ be a positive real number. Define $B_\epsilon(\mathbf{x}_0) = \{\mathbf{x} \in \mathbb{R}^n \mid d(\mathbf{x}, \mathbf{x}_0) < \epsilon\}$ (i.e., the set of all points less than ϵ distance to $\mathbf{x}_0$) to be the *open ball* centered at $\mathbf{x}_0$ of radius ϵ. Let $U \subseteq \mathbb{R}^n$. Then U is *open* if and only if for every $\mathbf{x}_0 \in U$ there is some $\epsilon > 0$ such that $B_\epsilon(\mathbf{x}_0) \subseteq U$. In other words, U is open if and only if given any point in U, the points *arbitrarily close by* also belong to U. It is not hard to show that this defines a topology on $\mathbb{R}^n$. In fact, this is how one can define a topology on any metric space.

Proof: The empty is vacuously open.[2] Using $\epsilon = 1$ (or really any positive number), one sees that $B_\epsilon(\mathbf{x}_0) \subseteq \mathbb{R}^n$ for all $\mathbf{x}_0 \in \mathbb{R}^n$ so $\mathbb{R}^n$ is open. Suppose A and B are open sets. Consider $\mathbf{x}_0 \in A \cap B$. Then since $\mathbf{x}_0$ belongs to both A and B and since these sets are open, $B_1 = B_{\epsilon_1}(\mathbf{x}_0) \subseteq A$ and $B_2 = B_{\epsilon_2}(\mathbf{x}_0) \subseteq B$ for some $\epsilon_1, \epsilon_2 > 0$. Notice that if $\epsilon = \min\{\epsilon_1, \epsilon_2\}$, then $B_\epsilon(\mathbf{x}_0)$ is contained in both (in fact equal to one of) B_1 and B_2. Thus $B_\epsilon(\mathbf{x}_0)$ is contained in both A and B, so $B_\epsilon(\mathbf{x}_0) \subseteq A \cap B$. Therefore, $A \cap B$ is open. Finally, if $\mathbf{x}_0$ belongs to a union of open sets, then it must belong to at least one of them, say U. Since U is open, there must be an open ball about $\mathbf{x}_0$ contained in that set, say B. Thus since $B \subseteq U$, B is contained in the union of the collection of open sets that include U. Therefore, the union of open sets is open. ■

Topology is a weird beast. We can still do large amounts of analysis in this context. We can even do some kinds of abstract geometry. But topology does not know about things like boundedness, distance, angle, and volume. However, topology does allow us to formulate a notion of *locality*. We will say a property holds *locally* at $\mathbf{x}_0$ if it holds on some open set containing $\mathbf{x}_0$. In $\mathbb{R}^n$, this is equivalent to saying that there is some $\epsilon > 0$ such that this property holds for all points within ϵ distance of $\mathbf{x}_0$.

In General: Every inner product space gives us a compatible normed space structure. Every normed space gives us a compatible metric space structure. Every metric space gives us a compatible topological space structure. None of this reverses (in general). Inner product spaces have many layers of rich structure. In such a space, linear algebra, geometry, analysis, and topology all have something to tell us.[3]

[1] What does this mean? Briefly, a *Cauchy sequence* is a list: $x_1, x_2, \ldots$ such that the further down the list we go, the closer the x_i's get to each other: For every $\epsilon > 0$ there is some $N > 0$ such that for all $i, j \geq N$ we have $d(x_i, x_j) < \epsilon$. A *convergent sequence* is a list such that the further down the list we go, the closer the x_i's get to some limit ℓ: For every $\epsilon > 0$ there is some $N > 0$ such that for all $i \geq N$ we have $d(x_i, \ell) < \epsilon$. Note that convergent sequences are always Cauchy. However, the converse may fail if our space is *incomplete*. For example, the sequence 3, 3.1, 3.14, 3.141, … is Cauchy in $\mathbb{Q}$, but not convergent since its limit is irrational (i.e., $\pi \notin \mathbb{Q}$). So $\mathbb{Q}$ is incomplete whereas $\mathbb{R}$ is complete – this is almost how $\mathbb{R}$ gets derived from $\mathbb{Q}$!

[2] A pun yet true.

[3] Functional analysis begins with a study of these structures (inner product, normed, and metric spaces). Also, infinite dimensional Hilbert spaces play a big role in quantum mechanics.

Differentiability: We will deal with differentiability more carefully later. Essentially a function f between subsets of $\mathbb{R}^n$ and $\mathbb{R}^m$ is *differentiable* at $\mathbf{x}_0$ if f can be *well approximated* by a linear function *near* $\mathbf{x}_0$. This means there is some linear transformation $J_{\mathbf{x}_0}$ from $\mathbb{R}^n$ to $\mathbb{R}^m$ such that

$$\lim_{\mathbf{x}\to\mathbf{x}_0} \frac{f(\mathbf{x}) - [f(\mathbf{x}_0) + J_{\mathbf{x}_0}(\mathbf{x}-\mathbf{x}_0)]}{\|\mathbf{x}-\mathbf{x}_0\|} = \mathbf{0}$$

where the limit is defined in the "usual" way.[4] This linear transformation $J_{\mathbf{x}_0}$ is called the *Jacobian* of f (evaluated at $\mathbf{x}_0$). The Jacobian is our (first) derivative. More on this later.

Looks Like / Mappings: When mapping between structured spaces, we would like our functions to carry over (i.e., *preserve*) some of that structure. Generally structure preserving maps are called homomorphisms (or morphisms). An invertible (i.e., one-to-one and onto) morphism whose inverse is also a morphism is called an *isomorphism.*[5]

When dealing with vector spaces, we call our morphisms *linear transformations.* In particular, a function T between vector spaces is linear if for any input vectors $\mathbf{v}, \mathbf{w}$ and scalar s, we have

$$T(\mathbf{v}+\mathbf{w}) = T(\mathbf{v}) + T(\mathbf{w}) \quad \text{and} \quad T(s\,\mathbf{v}) = s\,T(\mathbf{v})$$

If T is a linear map from a vector space to itself, we say T is a *linear operator.* For example, the derivative is a linear operator on the space of smooth functions.

In the context of inner product spaces, an invertible linear transformation that preserves inner products is an isomorphism. For the dot product, this looks like: $T(\mathbf{v}) \bullet T(\mathbf{w}) = \mathbf{v} \bullet \mathbf{w}$. Such a mapping would be *angle* and *length* preserving. For isomorphisms of normed spaces, we want invertible linear maps that preserve lengths: $\|T(\mathbf{v})\| = \|\mathbf{v}\|$. For metric spaces, isomorphisms are invertible maps (not necessarily linear) that preserves distances:

$$d(T(\mathbf{x}), T(\mathbf{y})) = d(\mathbf{x}, \mathbf{y})$$

For inner product, normed, and metric spaces we call isomorphisms *isometries.*

Topological spaces again are kind of weird. A morphism between topological spaces is called a *continuous* function. It turns out that the "right" way to preserve structure is to ask that inverse images of open sets are open. An invertible continuous map with a continuous inverse is called a *homeomorphism.* When working with $\mathbb{R}^n$, we can be a bit more direct.

Given a function f defined on a subset of $\mathbb{R}^n$ mapping into $\mathbb{R}^m$, we say f is *continuous* at $\mathbf{x}_0$ (some point in our domain) if given any $\epsilon > 0$ there is some $\delta > 0$ such that $f(B_\delta(\mathbf{x}_0) \cap (\text{domain of } f)) \subset B_\epsilon(f(\mathbf{x}_0))$. In other words, for every $\epsilon > 0$ there is some $\delta > 0$ such that whenever $\mathbf{x}$ is in the domain of f and $d(\mathbf{x}, \mathbf{x}_0) < \delta$ we have $d(f(\mathbf{x}), f(\mathbf{x}_0)) < \epsilon$ (i.e., the traditional ϵ-δ definition of continuity). Finally, a function is continuous if it is continuous at every point in its domain.

When working with structures where we have a concept of differentiability, we say that an invertible differentiable map with differentiable inverse is called a *diffeomorphism.*

We also have local versions of various concepts. For example, if f is an invertible continuous function from an open set containing $\mathbf{x}_0$ onto an open set and f's inverse is continuous, then f is a *local homeomorphism* (at $\mathbf{x}_0$). If you replace "continuous" with "differentiable", f becomes a *local diffeomorphism.*

---

[4] We have $\lim_{\mathbf{x}\to\mathbf{x}_0} g(\mathbf{x}) = \mathbf{L}$ if for every $\epsilon > 0$ there is some $\delta > 0$ such that for any $\mathbf{x}$ in the domain of g with $0 < d(\mathbf{x}, \mathbf{x}_0) < \delta$ we have $d(g(\mathbf{x}), \mathbf{L}) < \epsilon$. One usually only considers limits as $\mathbf{x}$ goes to $\mathbf{x}_0$ if g is defined on points *arbitrarily close to* but distinct from $\mathbf{x}_0$ (i.e., $\mathbf{x}_0$ is a *limit point* or *accumulation point* of the domain of g).

[5] With some structures, an invertible morphism's inverse will automatically preserve structure (i.e., the inverse of a morphism is a morphism). This is true for vector, normed, and inner product spaces. However, this can fail in the context of topological and metric spaces – there are invertible continuous maps with discontinuous inverses! For example, $f : [0, 2\pi) \to S^1$ (a half open interval of real numbers mapped to the unit circle) defined by $f(\theta) = (\cos(\theta), \sin(\theta))$ is continuous and invertible but its inverse fails to be continuous at $(1, 0) \in S^1$ (the angle jumps there).