Mitchell Faulk
June 22, 2014

A Note on Smooth Manifolds and Smooth Maps

1 Introduction

The notion of a smooth manifold is fundamental to modern geometry. In this note, we motivate and explore this notion along with the notion of a smooth map between manifolds.

2 Manifolds

We begin by recalling the notion of a manifold. Remember that the idea behind an n-dimensional manifold is that it is a topological space that looks locally like $\mathbb{R}^n$. More precisely, we have the following definition.

Definition 1. An **n-dimensional topological manifold** is a second-countable, Hausdorff space which is locally homeomorphic to the Euclidean space $\mathbb{R}^n$.

Here, locally homeomorphic to $\mathbb{R}^n$ means that every point of the manifold has an open neighborhood which is homeomorphic to some open subset of $\mathbb{R}^n$ (using the standard topology induced by the Euclidean norm). The additional requirements of being second-countable and Hausdorff simply stipulate that a manifold has "enough" open sets.

Another way of formulating the local condition in the definition of a manifold is the following. Let X be a topological space. An *n-dimensional coordinate chart* for X is a pair (U, ϕ) where U is an open subset of X and $\phi : U \to \mathbb{R}^n$ is a continuous map taking U homeomorphically onto its image $\phi(U)$ in $\mathbb{R}^n$. An *atlas of n-dimensional charts* for X is a collection $\{(U_\alpha, \phi_\alpha)\}$ of n-dimensional charts such that the open sets U_α form a cover for X, that is, such that $\cup_\alpha U_\alpha = X$. The following corollary is immediate from the definitions above.

Corollary 1. *Let X be a second-countable, Hausdorff topological space. Then X is an n-dimensional manifold if and only if there is an atlas of n-dimensional charts for X.*

This corollary brings up a subtle point about manifolds: there can be many different atlases for a given n-dimensional manifold X. In fact, there will usually be many such atlases. This begs the obvious question: does the manifold structure depend on the choice of atlas? In general, the answer is yes, so we try always to specify which atlas is being used whenever we are working with manifolds. However, it is somewhat common in the literature and when speaking about manifolds to neglect the choice of atlas, especially when dealing with standard examples for which there are standard choices of atlases.

3 Smooth manifolds

In the setting of Euclidean vector spaces, it makes sense to talk about differentiable functions between such spaces. Moreover, the notion of differentiability is a local notion, in the sense that the derivative of a function at a point depends only on the information of the function near that point. Because of these two observations, we could hope that the notion of differentiability could be extended to the setting of manifolds, since manifolds are locally the same as Euclidean spaces. Indeed, this turns out to be the case.

To understand how we can extend the notion of differentiability, let us recall what differentiability means in a setting where we understand the term, namely, the setting of finite-dimensional Euclidean spaces. There are more general settings than this one, but for our purposes, we are content to work in finite dimensions.

Recall that we say a map $f : \mathbb{R}^n \to \mathbb{R}^m$ is *differentiable at the point* $x \in \mathbb{R}^n$ if there is a linear map $Df_x : \mathbb{R}^n \to \mathbb{R}^m$ which approximates f at the point x in the sense that

$$\lim_{y \to 0} \frac{\|f(x+y) - f(x) - Df_x(y)\|}{\|y\|} = 0$$

If f is differentiable at every point of its domain, then we call f a *differentiable* map. We would also say that f is of class C^1. In such a case, the assignment $x \mapsto Df_x$ gives a map from $\mathbb{R}^n$ to the space $\mathcal{L}(\mathbb{R}^n, \mathbb{R}^m)$ of linear maps from $\mathbb{R}^n$ to $\mathbb{R}^m$. If this mapping is itself differentiable, then we say that f belongs to the class C^2. We define C^k inductively.

There is a more explicit description of C^k maps which is also useful. Working with f and x as above, let u be a vector in $\mathbb{R}^n$. The *directional derivative of f at x with respect to u*, written $f'(x; u)$, is the value of the limit

$$f'(x; u) = \lim_{h \to 0} \frac{f(x + hu) - f(x)}{h}$$

provided the limit exists. When u is the standard basis vector e_i, we write $f'(x, e_i) = D_i f(x)$ and we call this vector the *i-th partial derivative of f at x*. The assignment $x \mapsto D_i f(x)$ defines a function from $\mathbb{R}^n$ to $\mathbb{R}^m$. If for each i satisfying $1 \le i \le n$, the function $D_i f$ is a continuous map, then we say that f belongs to the class C^1. More generally, for a positive integer k, we say that f belongs to the class C^k if each function of the form

$$D_{i_1} \cdots D_{i_r} f$$

exists and is continuous, where i_t are positive integers satisfying $1 \le i_t \le n$ and $i_1 + \cdots + i_r = k$. The set of C^k maps on a subset U of $\mathbb{R}^n$ with values in $\mathbb{R}^m$ is denoted $C^k(U, \mathbb{R}^m)$. Note that we have the inclusions

$$C^1(U, \mathbb{R}^m) \subset C^2(U, \mathbb{R}^m) \subset \cdots \subset C^k(U, \mathbb{R}^m) \subset \cdots$$

Definition 2. A function $f : \mathbb{R}^n \to \mathbb{R}^m$ is **smooth** if f is of class C^k for each positive integer k.

Now let us try to extend this notion of differentiability to the setting of continuous maps $f : V \to \mathbb{R}^m$ from some subset V of a manifold X. For a point x of V, there is a neighborhood U containing x, which is mapped by some homeomorphism $\phi : U \to \mathbb{R}^n$ onto some subset $\phi(U)$ of $\mathbb{R}^n$. Let y denote the value $\phi(x)$ in $\mathbb{R}^n$. Now, note that the composition

$$f \circ \phi^{-1} : \mathbb{R}^n \longrightarrow \mathbb{R}^m$$

is a continuous map between two Euclidean spaces! This means that it makes sense to talk about differentiability of this composition near y. What we would like to say is that f is differentiable at x precisely when this composition is differentiable at y. However, note that such a definition would depend upon a choice of homeomorphism ϕ. So we would be obliged to include this clause in our definition, producing a definition such as the following: The function $f : U \to \mathbb{R}^m$ is *differentiable at the point x with respect to the chart* (U, ϕ) if the composition $f \circ \phi^{-1}$ is differentiable at the point $\phi(x)$. We similarly obtain notions of C^k with respect to the chart (U, ϕ) and smooth with respect to the chart (U, ϕ).

It would be easy enough to eliminate this clause from the definition of differentiability (resp. C^k, smooth, etc.): we simply declare f to be differentiable at x if f is differentiable at x with respect to all charts containing x. In this way, we can extend the notion of differentiability to a function $f : U \to \mathbb{R}^m$ from some subset U of a manifold X.

There is a minor problem in this definition: it would be tedious to verify. Indeed, we would need to check differentiability with respect to *all* such charts. It would be better if we could construct things so that we would only need to check differentiability with respect to one chart, and somehow differentiability with respect to other charts would follow. This motivates the notion of a C^k manifold.

Before introducing the notion of a C^k manifold, let us investigate how differentiability with respect to one chart could imply differentiability with respect to another. Suppose that $f : U \to \mathbb{R}^m$ is a continuous map that is C^k at x with respect to the chart (U, ϕ). Let $(\widetilde{U}, \widetilde{\phi})$ be another chart containing x. Note that we have the following equality of functions

$$f \circ \widetilde{\phi}^{-1} = (f \circ \phi^{-1}) \circ (\phi \circ \widetilde{\phi}^{-1})$$

wherever makes sense, that is, at the intersection of their domains, which happens to be $\widetilde{\phi}(U \cap \widetilde{U})$. We are already given that the function $f \circ \phi^{-1}$ is C^k at the point $\phi(x)$. Using the chain rule and the above equality of functions, we know that the assertion that the function $f \circ \widetilde{\phi}^{-1}$ is C^k at the point $\widetilde{\phi}(x)$ is equivalent to the assertion that the function $\phi \circ \widetilde{\phi}^{-1}$ is C^k at the point $\widetilde{\phi}(x)$.

The preceding paragraph implies the following. If we know that

- the function $f : U \to \mathbb{R}^m$ is C^k at x with respect to the chart (U, ϕ) and

- a chart $(\widetilde{U}, \widetilde{\phi})$ containing x satisfies the property that the map $\phi \circ \widetilde{\phi}^{-1}$ is C^k,

then we also know that the function f is C^k at x with respect to the chart $(\widetilde{U}, \widetilde{\phi})$. This motivates the following definitions.

Definition 3. Let X be an n-dimensional manifold. Two charts (U, ϕ) and (V, ψ) for X are said to be C^k **compatible** if either $U \cap V = \varnothing$ or the **transition map** $\phi \circ \psi^{-1} : \psi(U \cap V) \to \phi(U \cap V)$ is a C^k map. An atlas is said to be a C^k **atlas** if each chart is C^k compatible with every other chart. A **smooth atlas** is one which is C^k for each positive integer k. A manifold X together with a C^k (resp. smooth) atlas of charts is called a C^k **(resp. smooth) manifold**.

Because of the discussion before the previous definition, we also have the following definitions.

Definition 4. A function $f : V \to \mathbb{R}^m$ from an open subset V of a C^k manifold X into Euclidean space $\mathbb{R}^m$ is said to be C^k **differentiable at the point** x **in** V if there is a chart (U, ϕ) containing x such that the composition $f \circ \phi^{-1}$ is C^k differentiable at the point $\phi(x)$. The function f is said to be of class C^k if it is a C^k map at each point of its domain. Smooth functions on open subsets of smooth manifolds are defined similarly.

These definitions are well-defined because the previous discussion shows that for a C^k manifold, the notion of C^k differentiability is independent of the choice of chart.

For a map $f : X \to Y$ between two C^k manifolds (possibly of different dimension), we can also introduce a notion of C^k differentiability. Again, because of the C^k structure, one can show that the notion of C^k differentiability depends only on the behavior of f in one pair of charts (one from the atlas for X and one from the atlas for Y), not on the behavior of f in all such pairs. We leave this fact as an exercise for the interested reader.

Definition 5. Let X and Y be two C^k manifolds, possibly of different dimensions. A continuous map $f : W \to Y$ from some open subset W of X into Y is said to be C^k **differentiable at the point** $x \in W$ if there is a chart (U, ϕ) containing x and a chart (V, ψ) containing $f(x)$ such that the composition

$$\psi \circ f \circ \phi^{-1} : \phi(f^{-1}(V) \cap U) \longrightarrow \psi(V \cap f(U))$$

is a C^k differentiable map at $\phi(x)$. We say that f is **of class** C^k if f is C^k differentiable at all points of its domain. The notion of a smooth map between smooth manifolds is defined similarly.

Because we often want to be able to differentiate as many times as we'd like, we usually only consider the setting of smooth manifolds, and from this point forward, we will remain in this setting. Following the literature, we often use the letters M and N to denote smooth manifolds.

4 Tangent spaces and the differential

Even though the previous section describes *when* a map $f : M \to N$ between smooth manifolds is differentiable, it gives no process for determining what a derivative should be! The goal of this section is to determine what is meant by a derivative in this setting of smooth manifolds.

Recall that in the finite-dimensional Euclidean setting, the derivative of a map $f : \mathbb{R}^n \to \mathbb{R}^m$ at the point $x \in \mathbb{R}^n$ is a linear approximation $Df_x : \mathbb{R}^n \to \mathbb{R}^m$ to f at the point x. In a similar way, the derivative of a smooth map $f : M \to N$ between smooth manifolds will be an appropriate linear approximation of the map at a point $x \in M$. In this setting, however, it is not apparent what vector spaces should be the domain and codomain of the linear approximation. In the setting of Euclidean spaces, it is easy: we simply take the domain and codomain to be the domain and codomain of the original map f. The setting of smooth manifolds requires us to introduce a new type of vector space, namely, the tangent space T_xM of a manifold M at the point x.

There are many ways to define the tangent space T_xM, and in this section we hope to cover as many of these ways as we can. However, we first discuss an example which provides a geometric picture motivating our forthcoming dicussion.

In the case where the underlying manifold M is a sphere in $\mathbb{R}^3$, each tangent space can be pictured as an affine plane in $\mathbb{R}^3$. Indeed, one can picture the tangent space T_xM as a plane attached to the sphere M at the point x, with the property that the sphere's radius is a normal vector to the plane. This picture is a good one to keep in mind, and motivates the notion of a tangent space for an arbitrary manifold M.

We now begin our discussion of the tangent space. We will discuss three different yet equivalent approaches, in order of increasing abstraction. As is often the case, the most concrete approach is intuitive, but somewhat cumbersome to use in practice. On the other hand, the most abstract approach is elegant, yet somewhat divorced from geometric intuition. Therefore, we find that it is important to understand all three approaches, since they each have their own advantages and disadvantages.

For our discussion, let us fix a smooth n-dimensional manifold M and a point $x \in M$. We will construct the tangent space T_xM to M at x. This will be a vector space of the same finite dimension as M.

Definition using curves. The idea for this approach is to construct a vector space whose elements are equivalence classes of curves passing through the point x.

Let C_x denote the set

$$C_x = \{\gamma : (-a, a) \to M \mid \gamma(0) = x,\ \gamma \text{ is smooth},\ a > 0\}$$

of smooth curves passing through the point x. We define an equivalence relation on the set C_x by the rule $\gamma_1 \sim \gamma_2$ if and only if there is a chart (U, ϕ) containing

x such that

$$\frac{d}{dt}(\phi \circ \gamma_1)(t)|_{t=0} = \frac{d}{dt}(\phi \circ \gamma_2)(t)|_{t=0}.$$

It is easy to see that if the above equality holds for one chart, then it must hold for all of them. For a curve $\gamma \in C_x$, we let $\gamma'(0)$ denote the equivalence class of γ under this equivalence relation, and we let T_xM denote the set of all such equivalence classes.

We now endow T_xM with the structure of a vector space. The idea will be to use a chart to transfer the vector space structure from $\mathbb{R}^n$ to T_xM.

We define addition and we leave the definition of scalar multiplication up to the interested reader. Let (U, ϕ) denote a chart containing x. Up to composing ϕ with a translation, we may assume that $\phi(x) = 0$. For two curves γ_1 and γ_2 in C_x, let γ denote the curve defined by the rule

$$\gamma(t) = \phi^{-1}(\phi(\gamma_1(t)) + \phi(\gamma_2(t)))$$

for whichever t this makes sense. (One should check that this does make sense for at least some t and that γ belongs to C_x!) We then define

$$\gamma_1'(0) + \gamma_2'(0) := \gamma'(0).$$

To show that this addition is well-defined, we need to show that it does not depend on the particular representatives γ_1 and γ_2. We show now that it does not depend on the representative γ_1 and the proof for γ_2 will be similar.

Let $\widetilde{\gamma}_1$ be a curve equivalent to γ_1, meaning that

$$\frac{d}{dt}(\phi \circ \widetilde{\gamma}_1)(t)|_{t=0} = \frac{d}{dt}(\phi \circ \gamma_1)(t)|_{t=0}.$$

Let $\widetilde{\gamma}$ denote the curve defined by

$$\widetilde{\gamma}(t) = \phi^{-1}(\phi(\widetilde{\gamma}_1(t)) + \phi(\gamma_2(t))).$$

Then note that

$$\begin{aligned}
\frac{d}{dt}(\phi \circ \widetilde{\gamma})(t)|_{t=0} &= \frac{d}{dt}(\phi(\widetilde{\gamma}_1(t)) + \phi(\gamma_2(t)))|_{t=0} \\
&= \frac{d}{dt}(\phi(\gamma_1(t)) + \phi(\gamma_2(t)))|_{t=0} \\
&= \frac{d}{dt}(\phi \circ \gamma)(t)|_{t=0},
\end{aligned}$$

showing that $\widetilde{\gamma}$ and γ are equivalent, as desired.

After constructing scalar multiplication in a similar fashion, the reader will be convinced that the set T_xM carries the structure of a vector space. However, note that this vector space structure seems to depend upon our choice of chart (U, ϕ). Indeed, one might be concerned that another choice of chart (V, ψ) would endow the set T_xM with a different vector space structure entirely! This would be bad, for it would mean that we could not really talk about *the* tangent space at the point x; we would need to talk about the tangent space with respect to some chart. However, as is often the case when dealing with smooth manifolds, one can show that the above vector space structure does *not* depend on choice of chart, and hence we are justified in calling T_xM *the* tangent space at x. We outline a proof of this fact now.

Let γ_1 and γ_2 be curves in C_x. For a chart (U, ϕ), construct the path γ representing the sum of $\gamma_1'(0)$ and $\gamma_2'(0)$ as above. Let (V, ψ) be another chart containing x. Again, up to composing with a translation, we may assume that $\psi(x) = 0$. Let ρ be the path in C_x defined by the rule

$$\rho(t) = \psi^{-1}(\psi(\gamma_1(t)) + \psi(\gamma_2(t))).$$

Our goal is to show that ρ is equivalent to γ. We can do this by using the chain rule. In what follows, we use the notation DF_xv to denote the derivative of F at x applied to the vector v. For a single variable map $\tau : (a, b) \to \mathbb{R}^m$, we write $\tau'(c)$ for the derivative of τ at c. Using this notation, we have

$$\begin{aligned}
\frac{d}{dt}(\phi \circ \rho)(t)|_{t=0} &= \frac{d}{dt}(\phi \circ \psi^{-1} \circ \psi \circ \rho)(t)|_{t=0} \\
&= \frac{d}{dt}((\phi \circ \psi^{-1})(\psi(\gamma_1(t)) + \psi(\gamma_2(t)))))|_{t=0} \\
&= D(\phi \circ \psi^{-1})_{\psi(\gamma_1(0))+\psi(\gamma_2(0))}((\psi \circ \gamma_1)'(0) + (\psi \circ \gamma_2)'(0)) \\
&= D(\phi \circ \psi^{-1})_0((\psi \circ \gamma_1)'(0) + (\psi \circ \gamma_2)'(0))
\end{aligned}$$

Recall that the derivative at a point is a linear map. This implies that we have

$$\begin{aligned}
&D(\phi \circ \psi^{-1})_0((\psi \circ \gamma_1)'(0) + (\psi \circ \gamma_2)'(0) \\
&= D(\phi \circ \psi^{-1})_0(\psi \circ \gamma_1)'(0) + D(\phi \circ \psi^{-1})_0(\psi \circ \gamma_2)'(0) \\
&= (\phi \circ \psi^{-1} \circ \psi \circ \gamma_1)'(0) + (\phi \circ \psi^{-1} \circ \psi \circ \gamma_2)'(0) \\
&= (\phi \circ \gamma_1)'(0) + (\phi \circ \gamma_2)'(0) \\
&= \frac{d}{dt}(\phi(\gamma_1(t)) + \phi(\gamma_2(t)))|_{t=0} \\
&= \frac{d}{dt}(\phi \circ \gamma)(t)|_{t=0},
\end{aligned}$$

which shows that ρ and γ are equivalent, as desired.

Now it remains to find the dimension of the vector space T_xM. As expected, it will have the same dimension as the dimension of the manifold M, which, in this case, is n. The idea is to use a chart to construct a linear isomorphism $T_xM \to \mathbb{R}^n$.

Let (U, ϕ) be a chart containing x and assume that $\phi(x) = 0$. We define a map $\widetilde{d\phi}_x : C_x \to \mathbb{R}^n$ by the rule

$$\widetilde{d\phi}_x(\gamma) = \frac{d}{dt}(\phi \circ \gamma)(t)|_{t=0}.$$

This map induces a well-defined map $d\phi_x : T_xM \to \mathbb{R}^n$ defined by

$$d\phi_x(\gamma'(0)) = \frac{d}{dt}(\phi \circ \gamma)(t)|_{t=0},$$

which is a linear map with respect to the vector space structure on T_xM. We argue now that this is a bijective map. Since the map is clearly injective by the definition of T_xM, it remains only to show that the map is surjective.

Let v be a vector in $\mathbb{R}^n$. Let $\alpha_v : (-1,1) \to \mathbb{R}^n$ denote the path defined by $\alpha(t) = tv$. There is an $\epsilon > 0$ so that the image $\alpha(-\epsilon, \epsilon)$ is contained in $\phi(U)$. Let $\gamma_v : (-\epsilon, \epsilon) \to M$ denote the path in M defined by the rule $\gamma_v(t) = (\phi^{-1} \circ \alpha)(t)$. Then it is routine to check that $d\phi_x(\gamma_v'(0)) = v$, showing that $d\phi_x$ is surjective, as desired.

This concludes our discussion of the tangent space T_xM using the approach of curves.

Definition using derivations. The idea for this approach is to construct a vector space whose elements are derivations on the set of smooth real-valued functions on the manifold M.

Let $C^\infty(M)$ denote the vector space of real-valued smooth functions $f : M \to \mathbb{R}$. A *derivation at x* is a linear map $D : C^\infty(M) \to \mathbb{R}$ such that

$$D(f \cdot g) = D(f) \cdot g(x) + f(x) \cdot D(g) \quad \text{for each } f, g \in C^\infty(M).$$

We can define addition and scalar multiplication on the set of derivations at x in the usual way, and we obtain a vector space, which we call the tangent space to M at x. For notation's sake, we denote this vector space by T_xM_{Der}, though it is standard to refer to this vector space as T_xM as well.

We outline now how to show that this approach is equivalent to the previous approach by constructing a bijective correspondence between "tangent vectors" in each approach. The idea is to send a curve to the derivation which differentiates along this curve. More precisely, we define a map $\widetilde{D} : C_x \to T_xM_{\text{Der}}$ by $\gamma \mapsto \widetilde{D}_\gamma$ where $\widetilde{D}_\gamma$ is the derivation defined by the rule

$$\widetilde{D}_\gamma(f) = \frac{d}{dt}(f \circ \gamma)(t)|_{t=0}.$$

One can show that this map only depends on the equivalence class of the path γ and hence gives rise to a well-defined map $D : T_xM \to T_xM_{\text{Der}}$ defined by $D_{\gamma'(0)} = \widetilde{D}_\gamma$. It is somewhat beyond the scope of these notes to show that this map is a linear isomorphism, but this fact can be shown.

Definition using the cotangent space. The idea for this approach is to construct a vector space whose elements are dual vectors to some quotient space constructed from subspaces of $C^\infty(M)$.

The motivation of this construction is based upon the following observation. The space $C^\infty(M)$ is a commutative $\mathbb{R}$-algebra under pointwise multiplication. In this algebra, the set of all functions vanishing at x forms an ideal, which we denote by I. Elements of the square I^2 of this ideal must belong to the kernel of any derivation, by the defining property of derivations. Hence any derivation $D : \mathbb{C}^\infty(M) \to \mathbb{R}$ gives rise to a well-defined linear map $\alpha_D : I/I^2 \to \mathbb{R}$. This means that to every derivation D we can associate an element α_D of the dual space $(I/I^2)^*$.

It turns out that the map $D \mapsto \alpha_D$ is actually bijective, and hence we are justified in defining the tangent space T_xM to be the dual space $(I/I^2)^*$. The vector space structure on this space is canonical and needs no explanation. However, the correspondence between this space and the set of derivations is somewhat subtle, and we let the interested reader supply further details.

This completes our discussion of the different formulations of the tangent space T_xM. Although equivalent to the others, each has its own merits, so it is left to the student to decide which formulation is most applicable in which circumstance.

We now return to our main goal of this section: associating a type of derivative to a smooth map. As in the case of finite-dimensional Euclidean space, the derivative at a point will be a linear map. In this new setting of smooth manifolds, it will be a linear map from one tangent space to another. More precisely, if $f : M \to N$ is a smooth map between manifolds, the differential at $x \in M$ will be a linear map $df_x : T_xM \to T_{f(x)}N$. Because there are different formulations for the tangent space, there are also different formulations for the differential. We cover some of these formulations now.

Differential using curves. We first formulate the differential using curves. For a smooth map $f : M \to N$ and a point $x \in M$, the differential of f at x, denoted df_x, is a linear map $df_x : T_xM \to T_{f(x)}N$ defined by the rule

$$df_x(\gamma'(0)) = (f \circ \gamma)'(0).$$

Recall here that $\gamma'(0)$ denotes the equivalence class of the curve γ, and so $(f \circ \gamma)'(0)$ denotes the equivalence class of the curve $f \circ \gamma$ in N. It is routine to show that this map is well-defined (meaning that it is independent of the choice of representative γ) and that it is linear with respect to the vector space structure on the tangent spaces. We leave these exercises for the interested reader.

Differential using derivations. This is perhaps the most standard approach to defining the differential. In this case, we let $df_x : T_xM \to T_{f(x)}N$ denote the linear map defined by the rule $X \mapsto df_x(X)$ where $df_x(X)$ is the derivation at $f(x)$ defined by

$$df_x(X)(g) = X(f \circ g)$$

for $g \in C^\infty(N)$. Again, we leave it to the reader to check that the map df_x is a linear map.

There is an analogous formulation of the differential when we view the tangent space T_xM as the dual space I/I^2, and we leave it to the reader to construct this formulation.

In summary, we have accomplished our main goal of the section: we have associated to a smooth map a type of derivative. In fact, we have do so in many ways, using various formulations for the tangent space.

As a final remark, we should note that many of the constructions in this chapter carry over to the more general setting of C^k manifolds. To see how such constructions would carry over, it is often sufficient to simply replace the word smooth by the phrase C^k differentiable.