

SYNTHETIC DIFFERENTIAL GEOMETRY

MICHAEL SHULMAN

CONTENTS

1. Introduction	1
2. Constructive logic	3
3. Properties of nilsquare infinitesimals	4
3.1. R is a field	4
3.2. Microcancellation	5
3.3. D is not zero	5
3.4. D is not an ideal	5
4. Taylor series	6
5. Multivariable calculus	6
6. Microlinearity	7
7. Differential Geometry	10
7.1. Tangent vectors	10
7.2. Lie Groups	11
7.3. Vector fields	12
8. Sheaf Models	12

1. INTRODUCTION

Remember in my talk first quarter, I told you about Nonstandard Analysis, a way of incorporating infinitesimals into mathematics. We defined a number system ${}^*\mathbb{R}$ called the ‘hyperreal numbers’ which included the real numbers, but also infinitesimal numbers. We could then do calculus this way: for example, we defined the derivative by

$$f'(x) \approx \frac{f(x+d) - f(x)}{d}$$

for all infinitesimals d , where \approx means ‘differs by an infinitesimal’.

Today I want to talk about a totally different way of dealing with infinitesimals called *synthetic differential geometry*. This subject can be done in a very technical (and beautiful) way, but it can also be done in a straightforward axiomatic way, as long as you are willing to accept a little bit of strange behavior. I’ll take the latter path for most of today, and say a little at the end about the technical underpinnings. Experts are free to translate everything I say before then into their own language.

The basic idea is that we want to change the \approx signs to *equalities*. Suppose we want to take the derivative of $f(x) = x^2$ at $x = c$. We argue as follows.

Date: May 31, 2006.

Find $f'(c)$. Let d be so small that $d^2 = 0$. Then

$$\begin{aligned} f(c+d) &= (c+d)^2 \\ &= c^2 + 2cd + d^2 \\ &= c^2 + (2c)d \end{aligned}$$

and so we define $f'(c)$ to be the linear part of this, $f'(c) = 2c$.

This works for any polynomial. For example, if $g(x) = x^3$, we have

$$\begin{aligned} g(c+d) &= (c+d)^3 \\ &= x^3 + 3c^2d + 3cd^2 + d^3 \\ &= x^3 + 3c^2d \end{aligned}$$

so we take $g'(c) = 3c^2$. The general idea is this: **$f'(x)$ is the number such that $f(x+d) = f(x) + f'(x)d$ for all d such that $d^2 = 0$.**

We can also argue heuristically for other functions, like $f(x) = \sin x$. Draw a right triangle whose hypotenuse is 1 and whose angle is d with $d^2 = 0$. Then its height is also d , so its length is $\sqrt{1^2 - d^2} = 1$, and thus

$$\begin{aligned} \sin d &= d \\ \cos d &= 1. \end{aligned}$$

Therefore, we have

$$\sin(c+d) = \sin c \cos d + \cos c \sin d = \sin c + (\cos c)d$$

so $f'(c) = \cos c$, again the linear part of this.

What do we need to make this work? Well, we need there to always exist such an $f'(x)$. Thus, we assume the following axiom.

Axiom Version 1. Let $D = \{d \in R : d^2 = 0\}$. Then any function $g: D \rightarrow R$ is of the form

$$g(d) = a + bd$$

for a unique $a, b \in R$.

Clearly $a = g(0)$. We can then define the *derivative* of any $f: R \rightarrow R$ at $c \in R$ by considering the function

$$g(d) = f(c+d)$$

and write it as

$$g(d) = g(0) + bd = f(c) + bd.$$

We define $f'(c) = b$.

Therefore our Axiom implies that **every function has a derivative**.

Remark 1.1. This is not the same as saying that every function is differentiable in the usual $\varepsilon\delta$ sense. But you actually can show that that's true too.

What I want to do is explore the consequences of this axiom. Unfortunately, the first consequence is a little unnerving.

Theorem 1.2. $R = \{0\}$.

Proof. Consider the function

$$f(x) = \begin{cases} 0 & x = 0 \\ 1 & x \neq 0 \end{cases}$$

We have $f(d) = f(0) + bd = bd$ for some b , by our Axiom.

Consider a $d \in D$. Either $d = 0$ or $d \neq 0$. If $d \neq 0$, then $bd = 1$. But multiplying by d , we get $0 = bd^2 = d$, so $d = 0$. Thus any $d \in D$ is zero, so $D = \{0\}$. But then since $0 = 0 \cdot b$ for any b , the uniqueness in our axiom implies that $R = \{0\}$. \square

Well. That's not so good. (And somebody ought have been throwing things at me already anyway, because we know perfectly well that *not* all functions are differentiable.) Our theory of nilsquare infinitesimals has come to a crashing halt before we even got going. Something must be done; our fundamental axiom is contradictory. But we don't want to throw it away, since we need it for everything we're going to do.

Solution 1: Restrict to considering only differentiable functions. This is the classical thing to do (and also the nonstandard-analysis thing to do). But how do you define differentiable? In some $\varepsilon\delta$ sense. But the whole point of introducing infinitesimals is to *avoid* using $\varepsilon\delta$ arguments. And actually, it would be kind of nice if all functions were differentiable. Maybe some analysts would be out of a job, but for the rest of us, it would make a lot of things simpler.

Solution 2: *weaken logic* so that our axiom is no longer contradictory! The weaker sort of logic we need is called **constructive logic** or **intuitionistic logic** (its unintuitiveness notwithstanding).

2. CONSTRUCTIVE LOGIC

I'm not going to get much into constructive logic; maybe some other time I'll give a pizza seminar about that. But I need to say a little.

Historically, constructivists like Brouwer, Heyting, and Bishop were ideologically motivated and rejected any proof that was not 'constructive'. Of course this means rejecting things like the axiom of choice, but it means much more than that as well. Here's a classic **example of a nonconstructive proof**, which I think all mathematicians should meditate on for a while.

Theorem 2.1. *There exist irrational numbers α, β such that α^β is rational.*

Proof. Consider $\sqrt{2}^{\sqrt{2}}$. If $\sqrt{2}^{\sqrt{2}}$ is rational, then take $\alpha = \beta = \sqrt{2}$ and we're done. If not, then take $\alpha = \sqrt{2}^{\sqrt{2}}$ and $\beta = \sqrt{2}$, so that

$$\alpha^\beta = (\sqrt{2}^{\sqrt{2}})^{\sqrt{2}} = (\sqrt{2})^2 = 2.$$

\square

Classically this is a perfectly good proof, but there's something unsatisfying about it, since we don't actually know which number it is! The constructivists regarded this as a bogus proof because it asserts that something *exists* without actually giving a *construction*. This, in a way, is the fundamental principle of constructivism: **To show that something exists, you must construct it.**

For instance, constructivists reject the axiom of choice. But the above proof doesn't use AC, and it is still nonconstructive. One of the insights of the constructivists

	Classical	Constructive
AC	OK	Not OK
PEM	OK	Not OK
DN	OK	Not OK

TABLE 1. Classical vs. Constructive

was that the culprit which makes this proof nonconstructive is the logical **principle of excluded middle** (PEM):

$$A \vee \neg A.$$

This is, of course, an axiom of classical logic, which we used in an indispensable way in our nonconstructive proof: we said either $\sqrt{2}^{\sqrt{2}}$ is rational or it isn't. But the constructivists refused to assert this without knowing *which*.

This seems a little less strange if you think about some more exotic statements. For example, if R is the Riemann hypothesis, then classically

$$R \vee \neg R$$

is true, even though we don't know whether or not R is true. In fact, I believe people have actually used this to prove theorems: they give one proof assuming R and another proof assuming $\neg R$. So their theorem is true either way. But there's something unsatisfying about this; what if R turned out to be undecidable, like the continuum hypothesis?

None of this may convince you to become a constructivist. It's possible it might at least convince you that the constructivists weren't completely loony. But our motivation is completely different; we *need* to use constructive logic in order to be able to have our Axiom without a contradiction.

Why does using constructive logic solve our problem? We used the PEM all over the place in reaching our contradiction. For instance, we supposed that for any $d \in D$, either $d = 0$ or $d \neq 0$. But even more basically, we used it to *define* the function f ! If we don't know that every number is either zero or nonzero, we can't define a function case-wise—or if we do, it won't be defined on all real numbers.

Basically the only way to understand constructive logic is to use it for a while, so bear with me. I hope that you'll at least let me have it for the moment and let me try to convince you that the power of our Axiom is worth giving up PEM for. I'll leave you to ponder that while you get some more pizza.

3. PROPERTIES OF NILSQUARE INFINITESIMALS

What does the real line look like with these infinitesimals? I'm going to write R instead of \mathbb{R} for this line, to remind you that it behaves a bit differently from the real numbers we're used to.

3.1. R is a field. The first question is what sort of algebraic structure R has. Clearly it has to be a ring. You might expect, since it contains nilpotents, that it's not a field any more, and that might make you a little sad. (Recall that in NSA, the hyperreals ${}^*\mathbb{R}$ were a field, because the infinitesimals were not literally nilsquare, and had multiplicative inverses which were infinitely large. But our Axiom says nothing about whether there are infinitely large numbers.)

However, once again constructive logic comes to the rescue. While we certainly can't divide by anything nilsquare, we can still say that R is a field in the sense that

$$x \neq 0 \implies x \text{ is invertible.}$$

So we *can* divide by anything nonzero. Why does this not contradict our Axiom? We still have the law of contrapositive, but all this tells us is that **anything noninvertible cannot be nonzero**. Since our nilsquare infinitesimals are not invertible, they are *not nonzero*. But this does not necessarily imply they are equal to zero.

Since the **law of double negative** $\neg\neg A \implies A$ is equivalent to PEM, in constructive logic something can be not false without necessarily being true. Thus the infinitesimals are in sort of a 'netherworld' around zero: they aren't all zero, but we can't single out any that are actually *different* from zero either. They 'stick together' too much.

Formally speaking, what I'm doing is introducing the fact that R is a field as another axiom. In fact, I'm going to use all sorts of algebraic properties of R without mentioning where they come from; if I wanted to be completely rigorously formal, I would have to state a bunch more axioms. But they're all more or less intuitive, and the really important one is our Axiom that I've already stated.

3.2. Microcancellation. We can't divide by infinitesimals, but we can cancel them if they appear in an equation universally quantified, which is almost as good.

Proposition 3.1. *If $ad = bd$ for all $d \in D$, then $a = b$.*

Proof. Consider the function $f(d) = ad = bd$. By our Axiom, it is of the form $f(0) + cd = cd$ for a *unique* c . Hence, $a = c = b$. \square

3.3. D is not zero. All the elements of D are not nonzero, so they're 'almost' zero. But we also have

Proposition 3.2. *D is not equal to $\{0\}$.*

Proof. Suppose it were. Then $f(d) = a + bd = a$ and $g(d) = a + b'd = a$ would be the same for any b and b' in R . By the uniqueness clause of our Axiom, $b = b'$, which is absurd. Thus $D \neq \{0\}$. \square

This may look like a contradiction, but it isn't. No *particular* infinitesimal can ever be proved nonzero, but on the other hand we know for sure that zero isn't the only infinitesimal. Infinitesimals are slippery things.

You might object to the seeming use of proof by contradiction in that argument. But it's okay because what we were proving was a *negative* statement. It's okay constructively to prove $\neg A$ by assuming A and deriving a contradiction. It's even okay to prove $\neg\neg A$ by assuming $\neg A$ and deriving a contradiction; you just can't then conclude that A must be true.

3.4. D is not an ideal. Clearly if $d^2 = 0$, then $(ad)^2 = a^2d^2 = 0$, so D is closed under multiplication by elements of R . But it is not closed under addition.

Proposition 3.3. *It is not true that for all $d_1, d_2 \in D$, $d_1 + d_2 \in D$.*

Proof. Addition is the only thing that might fail. Since $d_1^2 = d_2^2 = 0$, we have

$$(d_1 + d_2)^2 = d_1^2 + 2d_1d_2 + d_2^2 = 2d_1d_2.$$

Thus, if $(d_1 + d_2)^2 = 0$ for all $d_1, d_2 \in D$, we would have $d_1d_2 = 0$ for all $d_1, d_2 \in D$ (since $2 \neq 0$, we can divide by it). But by universal microcancellation, this implies $d_2 = 0$ for all $d_2 \in D$, which we know is false. \square

Note we have also proven that although the elements of D square to zero, the product of two *different* nilsquares does not always vanish.

On the other hand, $d_1 + d_2$ does *cube* to zero:

$$(d_1 + d_2)^3 = d_1^3 + 3d_1^2d_2 + 3d_1d_2^2 + d_2^3 = 0$$

So it is a ‘nilcube’ infinitesimal. We can use these nilcube infinitesimals to talk about higher order derivatives.

4. TAYLOR SERIES

I already told you how to define the derivative. What about the second derivative? Well, naively, we take the derivative of the derivative. What does this mean in terms of f ? We have

$$f(x + d) = f(x) + f'(x)d$$

and now we want to ask about $f'(x + d)$. Let’s make the two ds different and write

$$\begin{aligned} f(x + d_1 + d_2) &= f(x + d_1) + f'(x + d_1)d_2 \\ &= f(x) + f'(x)d_1 + f'(x)d_2 + f''(x)d_1d_2 \\ &= f(x) + f'(x)(d_1 + d_2) + \frac{1}{2}f''(x)(d_1 + d_2)^2 \end{aligned}$$

recalling what we just proved. So we have an ‘infinitesimal Taylor formula’ for these nilcube infinitesimals. I’ll let you write down the version for $d_1 + d_2 + d_3$ and check it; clearly the pattern continues.

What we *don’t* know, however, is a Taylor formula for *any* nilcube infinitesimal; only the nilcubes of the particular form $d_1 + d_2$. Obviously, since our axiom only refers to nilsquares, it won’t tell us anything about arbitrary nilcubes. But this is a bit unsatisfying; if all functions on the nilsquares are uniquely affine, all functions on the nilcubes should be uniquely quadratic, and so on. So we can add this as an improvement to our axiom:

Axiom Version 2. Let $D_n = \{d \in R : d^{n+1} = 0\}$ (so $D = D_1$). Then any function $f: D_n \rightarrow R$ has the form

$$f(d) = a_0 + a_1d + a_2d^2 + \cdots + a_nd^n$$

for unique $a_i \in R$.

This shouldn’t be much of a stretch, conceptually. Once we introduce nilsquares that behave this way, it should be clear that nilcubes and so on should behave in a similar way; what’s special about 2?

We can write

$$D_\infty = \{d \in R : (\exists n)d^n = 0\} = \bigcup_{n \geq 1} D_n$$

for the set of all *nilpotent infinitesimals*. Thus, from the above axiom, we conclude that **all functions on D_∞ are given by a unique (formal) power series.**

Note that any formal power series converges on D_∞ , since it reduces to a finite polynomial on anything nilpotent.

There are more general axioms, introducing sets of infinitesimals on which functions are determined by *germs*, but we won't need them today.

5. MULTIVARIABLE CALCULUS

What about multivariable calculus? It's easy to define partial derivatives:

$$f(x + d, y) = f(x, y) + \frac{\partial f}{\partial x}(x, y)d$$

$$f(x, y + d) = f(x, y) + \frac{\partial f}{\partial y}(x, y)d$$

And so on. Now suppose we wanted to talk about the gradient of a function, which includes all the partial derivatives at once. We can consider

$$\begin{aligned} f(x + d_1, y + d_2) &= f(x, y + d_2) + \frac{\partial f}{\partial x}(x, y + d_2)d_1 \\ &= f(x, y) + \frac{\partial f}{\partial y}(x, y)d_2 + \frac{\partial f}{\partial x}(x, y)d_1 + \frac{\partial^2 f}{\partial x \partial y}(x, y)d_1 d_2 \end{aligned}$$

which includes the gradient, but also the mixed second partial. Sometimes we might want that, but sometimes we don't.

But note that if $d_1 d_2 = 0$, then we would just get the gradient. Of course, we know that we don't have $d_1 d_2 = 0$ for *all* d_1, d_2 , but we can just cut down to the set of pairs for which this is true. If we let

$$D(2) = \{(d_1, d_2) \in R^2 : d_1^2 = d_2^2 = d_1 d_2 = 0\}$$

(apologies for the notation! It's standard) then functions on $D(2)$ should be uniquely of the form

$$f(d_1, d_2) = a + b d_1 + c d_2.$$

At this point we start to suspect that there should be a more general version of our axiom. I'm not going to give a formal statement, but here's an informal one.

Definition 5.1. An *infinitesimal object* is a subset $C \subset R^n$ defined by polynomial equations which (among other things) force every component to be nilpotent.

If you like fancy words, you can call an infinitesimal object the 'spectrum of a Weil algebra'.

Axiom Version 3. If C is an infinitesimal object, then every function $f: C \rightarrow R$ is uniquely determined by a suitable polynomial or power series.

This is the final version of our axiom I'm going to introduce today. Infinitesimal objects include D , D_n , $D(2)$, also

$$D(n) = \{(d_1, \dots, d_n) : (\forall i, j) d_i^2 = d_i d_j = 0\}$$

and so on.

6. MICROLINEARITY

So far, we have axioms that let us study the structure of R by mapping into it from infinitesimal objects. To do differential geometry with these things, we'll need to map them into other things too. Functions landing in R^n are easy; we just do it componentwise. But what about more complicated things, manifolds and suchlike? If our manifolds are all embedded in some R^n , we could work in coordinates there, but this is kind of a messy way to do things, and not all our spaces will be embedded in R^n . I'm going to have to push your head through a little bit of abstraction, but it's necessary and not that bad, I think. I'll try to keep it as concrete as possible.

I'm going to be working with things that I'll call 'sets'. From our axiomatic point of view they *are* just sets: remember that once we assumed our Axiom, we didn't have to impose any 'topology' or 'differentiable structure' on R in order to get that all functions are differentiable. R is just a set of numbers, that happens to include some infinitesimals, and the way those infinitesimals behave means that as a set, R behaves 'cohesively' in some way. So in our alternate universe, arbitrary sets behave more like 'spaces' or 'manifolds' do in the world that we're used to. But since they're just sets, we can do all sorts of things with them that we can't necessarily do with manifolds; we'll see an example later.

Anyway, we'd like some property similar to our Axiom for more general sets, but those sets may not have any notion of 'addition' or 'multiplication' that we could use to formulate such a notion. It turns out that the right thing to generalize is not the *absolute* characterization of maps from micro-objects, but the *relative* relationships *between* maps out of different micro-objects.

For example, consider the map $D \times D \xrightarrow{+} D_2$. As we remarked above, not every nilcube need be the sum of two nilsquares, so this map need not be surjective. But our Axiom essentially tells us that they might as well be, from the perspective of the reals, since they are characterized by the same polynomial functions. Even more precisely, we can say this:

Proposition 6.1. *Any function $f: D \times D \rightarrow R$ such that $f(d, 0) = f(0, d)$ is of the form $f(d_1, d_2) = g(d_1 + d_2)$ for a unique $g: D_2 \rightarrow R$.*

Proof. Clearly any g determines an f with the given property. Conversely, we know that any $f: D \times D \rightarrow R$ has the form

$$f(d_1, d_2) = a + b_1 d_1 + b_2 d_2 + c d_1 d_2$$

and the assumption tells us that $a + b_1 d = a + b_2 d$, so by microcancellation, $b_1 = b_2 = b$. Therefore,

$$\begin{aligned} f(d_1, d_2) &= a + b(d_1 + d_2) + c d_1 d_2 \\ &= a + b(d_1 + d_2) + c(d_1 + d_2)^2 \end{aligned}$$

so $g(d) = a + b d + c d^2$ is the unique function that does the trick. \square

I've been trying to keep the category theory to a minimum, but really the best way to think about this is that the diagram

$$D \rightrightarrows D \times D \xrightarrow{+} D_2$$

is not necessarily a coequalizer (in our category of 'sets'), but ' R thinks it is'. If it were a coequalizer, then the above would be true for maps into *any* set X , but we have it only (so far) for maps into R .

Note that this property of maps into R can be stated without reference to addition or multiplication in R (only addition and multiplication in the infinitesimal objects), although its *proof* requires algebra in R .

Here's another similar result:

Proposition 6.2. *For any pair of functions $f_1, f_2: D \rightarrow R$ such that $f_1(0) = f_2(0)$, there is a unique function $g: D(2) \rightarrow R$ such that $f_1(d) = g(d, 0)$ and $f_2(d) = g(0, d)$.*

Proof. Clearly given g , we can define f_1, f_2 as shown and they will have that property. Conversely, given f_1, f_2 , we have

$$\begin{aligned} f_1(d) &= a_1 + b_1d \\ f_2(d) &= a_2 + b_2d \end{aligned}$$

but by assumption $a_1 = a_2 = a$, so define

$$g(d_1, d_2) = a + b_1d_1 + b_2d_2.$$

This works, and clearly is the unique function which does. □

In categorical language, this says that while the diagram

$$\begin{array}{ccc} 0 & \longrightarrow & D \\ \downarrow & & \downarrow \\ D & \longrightarrow & D(2) \end{array}$$

is not necessarily a pushout, ' R thinks it is'. Again, note that this property can be *stated* without reference to the arithmetic on R .

This sort of example can be multiplied indefinitely. In a minute we'll see why these properties are so useful, but for now you can think of them, as I've been saying, as a way of saying " R is well-behaved with respect to maps out of infinitesimal objects" without invoking the arithmetic on R . Thus we can extend this notion to more general sets as follows.

Definition 6.3. A set M is *microlinear* if any diagram of infinitesimal objects which is perceived by R as a colimit is also perceived by M as a colimit.

If you aren't comfortable with the idea of 'colimit', just think of the two examples I gave; this says that they remain true if we replace R by M .

Microlinearity is a property of a set. It essentially measures whether the set is 'cohesive' in 'the same way that R is'.

What are some examples of microlinear objects?

- Clearly R is microlinear.
- Also, R^n is microlinear, since we can argue componentwise.
- More generally, any product of microlinear sets is microlinear, since we can again argue componentwise.
- Even more generally than that, any *limit* of microlinear sets is microlinear.

The last includes, for example, *equalizers*, which allow us to show that **the zero-set of any function is microlinear**. For the non-categorically minded, this goes as follows. Let $\varphi: R^n \rightarrow R$ be a function and define $M = \{x \in R^n : \varphi(x) = 0\}$. I claim that M is microlinear.

Proof by example; skip if time short. Let's just prove it for that pushout. Suppose we have two maps $f_1, f_2: D \rightarrow M$ with $f_1(0) = f_2(0)$. Then composing with the inclusion, we have $f_1, f_2: D \rightarrow R^n$, and hence, since R^n is microlinear, there is a unique $g: D(2) \rightarrow R^n$ with $f_1(d) = g(d, 0)$ and $f_2(d) = g(0, d)$.

$$\begin{array}{ccccc}
 0 & \longrightarrow & D & & \\
 \downarrow & & \downarrow & \searrow & \\
 D & \longrightarrow & D(2) & \dashrightarrow & M \\
 & & & & \downarrow \\
 & & & & R^n \xrightarrow{\varphi} R \\
 & & & & \downarrow 0 \\
 & & & & R
 \end{array}$$

It remains to show that g lands in M , in other words that $\varphi \circ g = 0$. Observe that

$$\begin{aligned}
 \varphi(g(d, 0)) &= \varphi(f_1(d)) = 0 \\
 \varphi(g(0, d)) &= \varphi(f_2(d)) = 0
 \end{aligned}$$

But since R is also microlinear, there is a *unique* map $D(2) \rightarrow R$ which restricts to zero on both copies of D , namely the zero map. Therefore $\varphi \circ g = 0$. \square

So any manifold that we can define as a zero-set in some R^n is microlinear. Moreover, lots of 'infinite-dimensional' things are microlinear.

Proposition 6.4. *If M is microlinear and X is any set, then*

$$M^X = \{\varphi: X \rightarrow M\}$$

is microlinear.

Proof by example. Let's just do the pushout again. Suppose we have two maps $f_1, f_2: D \rightarrow M^X$ with $f_1(0) = f_2(0)$. Then for any $x \in X$, we have two functions

$$f_1(-)(x) \quad \text{and} \quad f_2(-)(x)$$

from $D \rightarrow M$ which agree at 0, so since M is microlinear, there is a unique $g_x: D(2) \rightarrow M$ extending them. Define $g(d)(x) = g_x(d)$; then $g: D \rightarrow M^X$ is our desired function. I'll let you check that it works and is unique. \square

So the class of microlinear sets is closed under 'function spaces'. This is quite an improvement over the classical notion of 'smooth manifold', which becomes much messier when you try to pass to infinite-dimensional function spaces.

7. DIFFERENTIAL GEOMETRY

Here's where the payoff comes.

7.1. Tangent vectors.

Definition 7.1. Let M be microlinear, $m \in M$. A *tangent vector* to M at m is a map $t: D \rightarrow M$ with $t(0) = m$. The set of tangent vectors at m is denoted $T_m M$.

You should think of D as a 'microsegment': a line so short that you can't bend it, but long enough to point in a definite direction (thus its defining property). So this says that a tangent vector is given by drawing a microsegment on M .

Proposition 7.2. $T_m M$ is a vector space.


Proof. Let's construct the operations. For $a \in R$, we let

$$(at)(d) = t(ad)$$

That was easy. In particular, we have the additive inverse defined by

$$(-t)(d) = t(-d).$$

Now let $t_1, t_2 \in T_m M$. Then they are both maps from D into M , which agree on 0. Since M is microlinear, they induce a unique map $D(2) \rightarrow M$.



We define

$$(t_1 + t_2)(d) = (t_1 \oplus t_2)(d, d)$$

Checking that this gives a vector space structure is an exercise for the reader. The most interesting one is the associativity of addition, for which you use $D(3)$ and another diagram that R thinks is a colimit. \square

Definition 7.3. Let $f: M \rightarrow N$ be a map between microlinear sets, and $m \in M$. The *differential* of f at m is $d_m f: T_m M \rightarrow T_{f(m)} N$ defined by

$$(d_m f(t))(d) = f(t(d))$$

This is just *composition*!

Proposition 7.4. $d_m f$ is a linear map.

Proof. Exercise. \square

Definition 7.5. If M is microlinear, its *tangent bundle* is $TM = M^D$.

This is just the set of all tangent vectors to M . It comes with a map $p: M^D \rightarrow M$ which is just evaluation at zero, $t \mapsto t(0)$, and the fiber over any $m \in M$ is just $T_m M$.

Finally, if $f: M \rightarrow N$, we have $df: TM \rightarrow TN$ which is, again, just composition.

We can define the action of a tangent vector on smooth functions by derivations easily too. Let $t \in T_m M$ and $f: M \rightarrow R$; then $f \circ t: D \rightarrow R$; thus we have

$$f \circ t(d) = f(t(0)) + bd = f(m) + bd$$

for a unique $b \in R$, which we define to be $t(f)$.

7.2. Lie Groups.

Definition 7.6. A *Lie group* is a microlinear group.

This includes the classical notion of Lie group, but also many ‘infinite-dimensional’ things, as we’ll see later. But we can do most of the things with it that we would want to.

Definition 7.7. Let G be a Lie group. Its *Lie algebra* is $\mathcal{G} = T_e G$.

For this to make sense, of course, there should be a Lie bracket, which should be the ‘infinitesimal’ version of a commutator in G . Let $X, Y \in \mathcal{G}$, so $X, Y: D \rightarrow G$ with $X(0) = Y(0) = e$, and consider the map

$$X \star Y: D \times D \longrightarrow G$$

given by

$$(d_1, d_2) \mapsto X(d_1)Y(d_2)X(-d_1)Y(-d_2)$$

(using multiplication in the group.

Now it turns out that

$$D \rightrightarrows D \times D \xrightarrow{m} D$$

is thought by R to be a colimit, where $m(d_1, d_2) = d_1 d_2$, and the two left maps send d to $(d, 0)$ and $(0, d)$, respectively. You can check this at home. Since G is microlinear, it also thinks this diagram is a colimit, and we have $(X \star Y)(d, 0) = (X \star Y)(0, d) = e$; thus we get a unique $[X, Y]: D \rightarrow G$ with $[X, Y](0) = e$. I’ll leave it as an exercise for you to check that this is, in fact, a Lie bracket.

7.3. Vector fields.

Definition 7.8. A *vector field* on M is a section of TM , i.e. a map $X: M \rightarrow TM$ such that $p \circ X = \text{id}_M$.

In other words, for each $m \in M$, we are given a tangent vector $X(m) \in T_m M$. We write $\mathcal{X}(M)$ for the set of vector fields on M .

Now, we can rephrase this in some very illuminating ways. A vector field is a map

$$X: M \rightarrow M^D$$

satisfying the property that $X(m)(0) = m$.

But to give a map into a set of functions is the same as to give a map

$$\tilde{X}: M \times D \rightarrow M;$$

a category theorist would call this an *adjointness* property. Thus a vector field can also be thought of as a map \tilde{X} such that $\tilde{X}(m, 0) = m$. This is an *infinitesimal deformation* of M . It is straightforward to check that *addition* of vector fields corresponds to *composition* of deformations:

$$\widetilde{X + Y}(m, d) = \tilde{Y}(\tilde{X}(m, d), d).$$

In particular, every infinitesimal deformation is invertible, since we have

$$\widetilde{-X}(\tilde{X}(m, d), d) = \tilde{X}(\tilde{X}(m, d), -d) = m.$$

Finally, we can use the adjointness property again to see that this is the same as a map

$$\hat{X}: D \rightarrow M^M$$

such that $\widehat{X}(0) = \text{id}_M$. This is a tangent vector to the set of self-maps of M ! (Observe that M^M , although ‘infinite-dimensional’, is microlinear because M is.) Thus we have proven

$$\mathcal{X}(M) = T_{\text{id}}(M^M).$$

In fact, since every infinitesimal deformation is invertible, all these tangent vectors land in the subset

$$\text{Aut}(M) = \{f \in M^M : f \text{ is invertible.}\}$$

which is a group, and in fact a Lie group, since can be defined as a limit from M^M , which is microlinear as M is. Thus we have

Theorem 7.9. $\mathcal{X}(M)$ is the Lie algebra of the Lie group $\text{Aut}(M)$.

We could, if we wanted, define the commutator of two vector fields independently and show that it agrees with that induced by $\text{Aut}(M)$, but I’ll leave that to you as well.

8. SHEAF MODELS

Now we’ve seen all this wonderful stuff that follows from our Axiom, once we agree to use constructive logic. But how do we know this axiom is consistent? All sorts of wonderful stuff can follow from a contradiction.

Of course, we have to take *some* axioms as basic, so we might as well start with this Axiom. But some people are more comfortable starting with the axioms of set theory, so let’s construct a model of our Axiom using set theory. And unfortunately, here I have to stop putting off the category theory.

The main idea is that **if we have a nice enough category \mathcal{C} , we can interpret logic in that category, taking the objects of that category to be ‘sets’**. Constructions such as limits and colimits in \mathcal{C} provide the logical and set-theoretic operations we are used to. For example, for arrows (‘functions’) $f, g: M \rightarrow N$ in our category, the ‘set’ $\{m \in M : f(m) = g(m)\}$ is represented by the equalizer, in \mathcal{C} , of f and g . This is called the *internal logic of \mathcal{C}* . If \mathcal{C} is the category of sets, what we get is the usual logic we are used to.

To interpret all of logic in the best possible way, what we need is for the category to be a *topos*. I’m not going to say any more about that, except to note that **the internal logic of a topos is, in general, constructive** rather than classical. Thus we are *forced* to constructivism, even if we have no philosophical inclination towards it, if we want to do mathematics in a topos.

Now I want to say a few words about how to construct a topos in which our Axiom is true. We want all classical manifolds in it, so let’s start with the category **Mfd** of manifolds. But we also want infinitesimal objects. No algebraic geometer is going to be surprised by what we do, which is to represent spaces by the *rings of functions* on them.

Definition 8.1. A *C^∞ -ring* is a ring A equipped with, for every smooth function $f: \mathbb{R}^n \rightarrow \mathbb{R}$, a map $\widehat{f}: A^n \rightarrow A$, with obvious compatibility conditions.

Examples:

- A manifold M determines a C^∞ -ring $C^\infty(M)$, the set of smooth functions $M \rightarrow \mathbb{R}$, and conversely is uniquely determined by it.

- Moreover, a map of manifolds $f: M \rightarrow N$ is uniquely determined by the map of C^∞ -rings $f^*: C^\infty(N) \rightarrow C^\infty(M)$ given by composition with f .

In particular, the *points* of M , which are precisely the smooth maps $*$ $\rightarrow M$, are the same as C^∞ -ring maps $C^\infty(M) \rightarrow \mathcal{C}^\infty(*) = \mathbb{R}$, which can be identified with certain maximal ideals of $\mathcal{C}^\infty(M)$. In categorical language, we have a full and faithful embedding

$$\mathbf{Mfd} \hookrightarrow C^\infty \mathbf{Rng}^{op}.$$

- In fact, if M is a smooth manifold, any subset $X \subset M$ gives a C^∞ -ring $\mathcal{C}^\infty(X)$.
- The ‘ring of dual numbers’ $\mathbb{R}[\varepsilon] = \mathbb{R}[x]/(x^2)$ is a C^∞ -ring; for a smooth function $f: \mathbb{R} \rightarrow \mathbb{R}$, we define

$$\widehat{f}(a + b\varepsilon) = f(a) + f'(a)b\varepsilon.$$

This ring is what will represent D .

Exercise: (if you’ve never done this) check that C^∞ -ring homomorphisms $C^\infty(M) \rightarrow \mathbb{R}[\varepsilon]$ are the same as tangent vectors to M .

Thus, the category $C^\infty \mathbf{Rng}^{op}$ is a nice one that includes all manifolds, and also infinitesimal objects. (It should be obvious how to define C^∞ -rings to represent D_n , $D(n)$, and so on.) But it isn’t yet a topos: in particular, it doesn’t have ‘exponentials’ like M^M .

Unfortunately, anyone who is likely to understand this last bit is already thinking it and doesn’t need me to say it, but I’ll say it anyway: we equip $C^\infty \mathbf{Rng}^{op}$ (or maybe the finitely-generated ones) with a suitable coverage, or ‘Grothendieck topology’, and take *sheaves* on it to get a topos. And one can then prove that our Axiom (and other axioms) is in fact true, in the internal logic of that topos.

I hope that maybe I’ve whetted your appetite to learn more about this stuff. At least maybe I’ve given you something to be confused about over summer vacation.