

Vector Analysis—Math 471/571

Spring 2000

Tennessee Technological University

Jeffrey Norden & Matthias Eller

§ 1 REVIEW OF VECTORS (and a gentle introduction to linear algebra)

It is certainly correct to say that a vector has both magnitude and direction. It is also correct to say that we can picture vectors as arrows drawn between a pair of points in space. In the greatest generality, the correct way to define a vector is to say that it is a member of a “vector space,” and to then define exactly what a vector space is. For this course, however, we can adopt the following fairly simple definition for what a vector is.

1.1 Definition. Let n be a positive integer. An n -dimensional vector is an ordered list of n real numbers.

We usually use angle bracket notation to denote a specific vector. For example $\langle 2, 7 \rangle$ is a two-dimensional vector (but more precisely, the vector itself is simply “2 followed by 7”), $\langle 3, -1, 8.2 \rangle$ is a three-dimensional vector, and $\langle 1, -2, 3.9, 42, \sqrt{7} \rangle$ is a five-dimensional vector. There are two common notation for indicating that a variable represents a vector. One can use bold face as our text does: $\mathbf{v} = \langle 1, 2, 3 \rangle$, or one can use an arrow as we will do in class and in these notes: $\vec{v} = \langle 1, 2, 3 \rangle$.

We *visualize* a vector as an arrow between two points in the appropriate space (this is a bit of a stretch if the dimension is four or more). So we can picture $\langle 2, 7 \rangle$ as an arrow from the origin which goes right two units and up seven. But we can also picture it starting at another point—the arrow between the points $(1, 4)$ and $(3, 11)$ would also be labeled by $\langle 2, 7 \rangle$. The direction of the vector is just the direction of the corresponding arrow. The magnitude is computed via the Pythagorean Theorem by taking the square root of the sum of the squares, i.e., $\|\langle 2, 7 \rangle\| = \sqrt{2^2 + 7^2} = \sqrt{53}$, while $\|\langle 1, -2, 3.9, 42, \sqrt{7} \rangle\| = \sqrt{1791.21}$. Many authors, including our textbook, just use simple absolute value signs to denote vector length.

You should already be quite familiar with several operations on vectors. If \vec{v} and \vec{w} are vectors of the same dimension, then $\vec{v} + \vec{w}$ is another vector of this same dimension formed by adding coordinate-wise. If t is a number and \vec{v} is a vector, then $t\vec{v}$ is a vector formed by multiplying each coordinate of \vec{v} by t . In this context we usually refer to the number t as a *scalar*, which is just a fancy way of saying that it is a quantity which is not a vector.

So $5\langle 2, 3 \rangle = \langle 10, 15 \rangle$ and $\langle 1, 2, 3 \rangle + \langle 4, 5, 6 \rangle = \langle 5, 7, 9 \rangle$. But note that $\langle 1, 2, 3 \rangle + \langle 4, 5 \rangle$ is undefined since the dimensions are different. These two operations (*vector addition* and *scalar multiplication*) are the most fundamental. You should also be familiar with two vector products. If \vec{v} and \vec{w} are vectors of the same dimension, then we form the *dot product* $\vec{v} \cdot \vec{w}$ by summing the result of coordinate-wise multiplication. This is sometimes also called the scalar product of \vec{v} and \vec{w} , since the result is a scalar. If \vec{v} and \vec{w} are three dimensional vectors, then you should also be familiar with the vector product $\vec{v} \times \vec{w}$. It is important to note that this cross product is defined only in dimension three.

There is another notational device for dealing with vectors. If we let \vec{i} , \vec{j} , and \vec{k} denote the unit coordinate vectors, then we can uniquely decompose any vector into a sum of scalar multiples of these vectors. For example, $\langle 1, 2, 3 \rangle = 1\langle 1, 0, 0 \rangle + 2\langle 0, 1, 0 \rangle + 3\langle 0, 0, 1 \rangle = \vec{i} + 2\vec{j} + 3\vec{k}$. However, one has to be a bit careful with this notation, since it can sometimes make the dimension ambiguous. We write $\langle 2, 3 \rangle = 2\vec{i} + 3\vec{j}$, but $\langle 1, 0 \rangle$ and $\langle 1, 0, 0 \rangle$ are *not* the same vectors. You must tell from the context which of the unit vectors \vec{i} and \vec{j} actually represent. So, when you see $2\vec{i} + 3\vec{j}$ you must be able to decide whether this means the vector $\langle 2, 3 \rangle$ or the vector $\langle 2, 3, 0 \rangle$. (Fortunately, this sort of notation is never used in dimensions greater than three.)

You may have noticed that something about our definition of vector seems very familiar or even a bit suspicious. Recall that \mathbb{R}^2 denotes the set of all points in the plane, and \mathbb{R}^3 denotes the set of all points in three dimensional space. Well, the points in \mathbb{R}^2 were *also* defined to be ordered pairs of real numbers, and points in \mathbb{R}^3 were defined as ordered triples of real numbers. How can we use the same definition for for points and vectors? It *seems* like vectors and points are very different. In fact the opposite is actually true:

1.2 Note. *Vectors and points are exactly the same mathematical objects!*

In other words, there is only one “ordered pair consisting of the number 2 followed by the number 3.” This isn’t changed by the fact that we now have two different notations for it: $(2, 3)$ and $\langle 2, 3 \rangle$. How can an ordered pair simultaneously represent both a point and arrow? This sort of thing actually happens all the time. The number “2” might represent 2 apples or 2 oranges, we don’t use a different number system for each type of thing we want to count. For a better example, consider a simple function like $f(x) = x^2$. When we write $f(\sqrt{2}) = 2$, the 2 on the right hand side represents the idea that “ f has a magnitude of 2” at a certain x -value. When we write $f'(1) = 2$, this exact same

“number 2” represents the rate at which f is changing at a certain x -value. In the same way, the ordered pair “2 followed by 3” can represent either a point or the idea of moving 2 units to the right and then 3 units up.

Actually we could make do with just one notation, and many mathematicians use just $\langle 2, 3 \rangle$, leaving it to the context in which it is used to determine whether it should be a point or a vector. More important is that we use a different notation for variables: when you see either $\mathbf{x} = \mathbf{y} + \mathbf{z}$ (in the textbook) or $\vec{x} = \vec{y} + \vec{z}$ (in these notes) it *means* that in addition to whatever else is said, \vec{x} , \vec{y} , and \vec{z} are vectors. This saves a lot of tedious writing of phrases like “let x , y , and z be vectors.”

Why is this important to understand? The fact that we can interchange vectors and points is the key to using vectors effectively. Consider the derivative example again. Even though each $f'(x)$ represents a rate, we still draw the graph of f' in order to understand its properties. We can do this because, regardless of how we interpret them, the values of $f'(x)$ are numbers. Suppose we are working in \mathbb{R}^2 , and consider solving the vector equation $\|\vec{x}\| = 1$. More precisely, consider the set defined by

$$S = \{ \vec{x} : \vec{x} \in \mathbb{R}^2 \text{ and } \|\vec{x}\| = 1 \}$$

Then S is a set of vectors, but how should we interpret and picture this set? We will do so by considering S to be a set of points, so $\|\vec{x}\| = 1$ becomes an equation for the unit circle in \mathbb{R}^2 . This is just the boundary of the circle—to include the inside, you would use $D = \{ \vec{x} : \vec{x} \in \mathbb{R}^2 \text{ and } \|\vec{x}\| \leq 1 \}$. Note that if we change \mathbb{R}^2 to \mathbb{R}^3 , the same equation defines the unit sphere in three dimensions. In even higher dimensions we take this equation to be the definition of the unit sphere, even though it is virtually impossible to picture it (such things are sometimes called hyper-spheres). Even more important are vector definitions for lines and planes such as $L = \{ t \langle 2, 4 \rangle : t \in \mathbb{R} \}$ or $P = \{ \vec{x} : \vec{x} \cdot \langle 1, 5, 7 \rangle = 0 \}$.

There is another way to view the distinction between points and vectors. The field of mathematics called linear algebra studies the general notion of a vector space. A vector space is a set V together with two operations defined on it: scalar multiplication and vector addition. So for any number t and any \vec{v} in V there must be a well defined $t\vec{v}$ in V , and for any \vec{v} and \vec{w} in V , there must be a well defined $\vec{v} + \vec{w}$ in V . The set V and the multiplication and addition defined on it can be quite general and abstract. But, in order to earn the title of “vector space” they are required to satisfy certain algebraic

properties, similar to the commutative and distributive properties that the real numbers have. (If you have your copy of Stewart's calculus textbook handy, these properties are exactly the ones which are listed in the box at the bottom of page 692).

For an example of an abstract vector space, consider letting V denote the set of all continuous functions from \mathbb{R} to \mathbb{R} , and define tf and $f + g$ by $(tf)(x) \stackrel{\text{def}}{=} tf(x)$ and $(f + g)(x) \stackrel{\text{def}}{=} f(x) + g(x)$. This does satisfy the required algebraic properties (check this), but it is not a 2-dimensional, or even a 200-dimensional space. In an important sense, this space is *infinite* dimensional.

We won't be able to go very far into the general theory of linear algebra this semester, but you are likely to encounter it in later classes. For now, the important idea is very simple. When we think of \mathbb{R}^2 as the plane with its geometric properties (such as distance) we call it "Euclidean space" and refer to its elements as "points." But if we think of \mathbb{R}^2 as a set on which scalar multiplication and vector addition are defined, we call it a "vector space" and refer to its elements as "vectors."

1.3 Review Problems.

Decide which of the following are well defined, and evaluate those that are. Give a geometric interpretation for each of the defined answers.

- 1) $2\langle 1, -2, \pi \rangle - \sqrt{2}\langle 8, 4, -7 \rangle$
- 2) $\langle 1, 3 \rangle + \langle 5, 2, 1 \rangle$
- 3) $\langle 1, 4 \rangle \langle 6, 2 \rangle$
- 4) $\langle 1, 4 \rangle \cdot \langle 6, 2 \rangle$
- 5) $\langle 1, 4 \rangle \cdot \langle 6, 2, 3 \rangle$
- 6) $\langle 1, 4 \rangle \times \langle 6, 2 \rangle$
- 7) $\langle 1, 4, 1 \rangle \times \langle 6, 2, 3 \rangle$
- 8) $\langle 1, 4, 1 \rangle \cdot \vec{i}$
- 9) $(2\vec{i} + 3\vec{j}) \cdot (7\vec{i} - 8\vec{j})$
- 10) $(2\vec{i} + 3\vec{j}) \times (\vec{i} - \vec{j} + 2\vec{k})$
- 11) $\langle 1, 2, 3, 4, 5, 6, 7 \rangle \cdot \langle 7, 6, 5, 4, 3, 2, 1 \rangle$
- 12) $\| \langle 1, 2, 3, 4, 5, 6, 7 \rangle \|$
- 13) $\| \langle \sqrt{2}, 10, 1/\pi \rangle \|$
- 14) $\vec{r} \cdot \vec{r} - \|\vec{r}\|^2$

§ 2 VECTORS FUNCTIONS

We start with the *general* mathematical concept of a function. There are several ways to formalize the idea of a function. The following is not the most precise possible definition, but it will suffice for this class.

2.1 Definition. Let X and Y be arbitrary sets. A *function* from X to Y is an assignment of a specific member of Y to each member of X .

Like everything else in mathematics, we give functions names in order to talk about them. We often use the letters f , g , and h for functions, but actually any name can be used. More important is the standard use of “functional notation.” If f is a function from X to Y and if $x \in X$ then we use the notation $f(x)$ to denote the member of Y which is assigned to x by f . (The element symbol \in denotes membership— $x \in X$ is read “ x is in X ” and just means that x belongs to the set X .)

2.2 Example. To make this concrete, we review a trivial example. Let f be the function from \mathbb{R} to \mathbb{R} which multiplies each number in \mathbb{R} by itself. Then $f(0) = 0$, $f(2) = 4$, $f(3) = 9$, etc. In general, for any number x , $f(x) = x^2$, and of course this equation is a more convenient way to define this f than the sentence we used about multiplication.

Often, books define a function as a “rule” which makes an assignment of a member of Y to each member of X . This is a useful idea, but it can be misleading. Consider the function g from \mathbb{R} to \mathbb{R} defined by $g(x) = (x + 1)^2 - 2x - 1$. The rule for g seems to be very different from the one for f in the previous paragraph, since we have to add one, square, and then subtract two other quantities. But notice that $g(0) = 0$, $g(2) = 4$, and $g(3) = 9$. In fact, for any number x , $g(x) = (x + 1)^2 - 2x - 1 = (x^2 + 2x + 1) - 2x - 1 = x^2$. This shows that f and g are identical functions, which we express simply by writing $f = g$. So, the function is something a bit more abstract than the “rule,” rather it is the thing that is the same when considering the rules x^2 and $(x + 1)^2 - 2x - 1$.

2.3 Notation. We will often use the notation $f : X \rightarrow Y$ as a shorthand for the statement “ f is a function from X to Y .” When $f : X \rightarrow Y$, we call the set X the *domain* of f and Y is called the *co-domain* of f . The *range* of f is the subset of Y consisting of the members of Y which the function actually uses. More precisely $\text{ran}(f) \stackrel{\text{def}}{=} \{f(x) : x \in \text{dom}(f)\}$. When we define a function by a simple equation it is

understood that its domain is taken to be the largest possible X which makes sense. So for $f(x) = x^2$, $\text{dom}(f) = \mathbb{R}$, while for $f(x) = 1/x$, $\text{dom}(f) = \{x \in \mathbb{R} : x \neq 0\}$.

You should already be very familiar with the ideas of continuity and differentiability for functions from \mathbb{R} to \mathbb{R} . You should also be familiar with the idea of composing two functions, the notion of a one-to-one function, and the concept of the inverse of a one-to-one function.

Now we finally bring vectors into the picture. This semester, we will be dealing with functions whose domains and/or ranges are sets of vectors instead of sets of numbers. In fact, this is what Calculus-III was really all about, but you might not have realized it.

2.4 Vector functions of a scalar variable. We first consider the simple case where the domain is still in \mathbb{R} , but the range is some \mathbb{R}^n . How can we define such a function? All that is needed are n functions from \mathbb{R} to \mathbb{R} which specify the coordinates of the range. For example, $f(x) = \langle x, x^2, x^3 + 1, \sin(x) \rangle$ defines a function $f : \mathbb{R} \rightarrow \mathbb{R}^4$. For a function like this, we often use $\vec{f}(x)$ instead of $f(x)$ to indicate that the values of \vec{f} are vectors. However, since it is usually clear from the context where the range of a function will be, the use of the arrow on the function name is sometimes considered optional.

How can we interpret these functions? One way is to think of the domain as measuring time and the value of \vec{f} as giving the position of an object in space (recall that a vector itself can be interpreted as an arrow or as a point). When we do this, we usually change the domain variable to t , and in fact $\vec{r}(t)$ is a common notation. Of course the dimension of the range must be either 2 or 3 if the object we are modeling lives in the same universe that we do. On the hand, if you are building a multi-jointed robot arm whose configuration is determined by 15 measurements, then the state of the arm at time t gives a function $\vec{s} : \mathbb{R} \rightarrow \mathbb{R}^{15}$. An alternate interpretation for these functions is to “throw away” the domain of the function and only consider its range as a subset of \mathbb{R}^n . Interpreted in this way, we really just have another way of viewing the “parametric” curves you first saw in Calculus-II.

2.5 Scalar functions of several variables. Next we consider the case of functions $f : \mathbb{R}^n \rightarrow \mathbb{R}$. Since the value of f is a scalar, we don't put an arrow over the function name. We can consider an element of the domain to be a single vector, in which case we might use the notation $f(\vec{r})$. Equivalently, we can explicitly show the dimension of the domain by listing the coordinates of a point in the domain with separate variables. We might use

$f(x, y)$, $f(x, y, z)$, and $f(x, y, z, w)$ for the cases where $n = 2$, $n = 3$, and $n = 4$. It is important to realize that these are just different notations. For example, we could define $f : \mathbb{R}^2 \rightarrow \mathbb{R}$ by $f(x, y) = \sqrt{x^2 + y^2}$, but the equation $f(\vec{r}) = \|\vec{r}\|$ defines exactly the same function. For $n > 4$, we usually use subscripts on the variables instead of separate letters: $f(x_1, x_2, x_3, x_4, x_5)$. In fact, subscripts can be used even in lower dimensions.

In the case when $n = 2$ the simplest (and most important) interpretation is to look at the graph of $z = f(x, y)$, which will be a surface in \mathbb{R}^3 . When $n = 3$, the corresponding “graph” would have to be placed in \mathbb{R}^4 , which we can’t visualize. So, instead we can just think of f as “assigning” a number to each point in \mathbb{R}^3 . For example, if we measure the temperature at each point in space, we obtain a function $T(x, y, z)$. Another important example is to measure density (in, say, gm/cm³) at each point. When $n > 3$, we can no longer visualize f at all. But, for example, if we want to express how the height of the end of the robot arm mentioned above varies with the state of the arm, we would use a function $h : \mathbb{R}^{15} \rightarrow \mathbb{R}$.

You should already be very familiar with the concept of partial derivatives and of iterated integrals, which apply to functions of two or more variables.

2.6 Vector functions of several variables. We now consider the general case of functions $f : \mathbb{R}^n \rightarrow \mathbb{R}^m$, which are a main focus for this course. It is easy enough to define such functions by combining the ideas from 2.4 and 2.5. We just need to give m scalar functions each of which has n variables. For example, $f(x, y) = \langle x^2, y^3, x + y \rangle$ defines $f : \mathbb{R}^2 \rightarrow \mathbb{R}^3$, and $f(x, y, z) = \langle x^2z, x + y - z \rangle$ defines $f : \mathbb{R}^3 \rightarrow \mathbb{R}^2$. Note that these definitions can also be written as $x^2\vec{i} + y^3\vec{j} + (x + y)\vec{k}$ and $x^2z\vec{i} + (x + y - z)\vec{j}$ (although the second equation could be accidentally interpreted as having a 3-dimensional range).

Visualizing these functions can be a challenge. First we will consider the cases where n and m are equal and no bigger than 3, i.e., $f : \mathbb{R}^2 \rightarrow \mathbb{R}^2$ and $f : \mathbb{R}^3 \rightarrow \mathbb{R}^3$. One can think of these functions in a simple “pointwise” fashion— f “moves” each point of the plane (or \mathbb{R}^3) into a new position. The term “move” is a little misleading, since the motion isn’t smooth or continuous with respect to time. Instead, we think each point \vec{r} jumping into a new position $f(\vec{r})$. Actually, you have been using this interpretation since high school. To make it more obvious, let’s use the letters r and θ (instead of x and y or \vec{p}) for the coordinates of the domain, and consider the function defined by $f(r, \theta) = \langle r \cos(\theta), r \sin(\theta) \rangle$. This is just the “coordinate transformation” which corresponds to polar coordinates. As a simple exercise, you should also write down the functions which correspond to cylindrical and

spherical coordinates. Since complex numbers are really just another interpretation of two-dimensional vectors, complex-valued functions are usually also viewed in this way. However, there is another visualization for $f : \mathbb{R}^2 \rightarrow \mathbb{R}^2$ and $f : \mathbb{R}^3 \rightarrow \mathbb{R}^3$ which is much more important for this class. We take the domain of f to represent points, but interpret the range of f as vectors visualized as *arrows*. We imagine the base of each arrow $f(\vec{r})$ as being placed at the point \vec{r} . The result is a picture of the plane (or \mathbb{R}^3) which is continuously covered up with arrows of varying length pointing in various directions. When we interpret f in this way we call f a *in vector field*. We also usually use $\vec{f}(\vec{r})$ or $\vec{f}(x, y, z)$ to emphasize that we are using this interpretation. It is important to realize that the arrows don't really overlap or intersect—we think of each one as being independent of the other. It would require 4 or 6 dimensions to really “draw” the arrows in an independent way. Since it is basically impossible to picture what these higher dimensions look like, we have to make due with just pretending that the arrows don't cross each other.

A vector field is exactly what you need in order to model the situation of a varying force at each point in space. For this reason, our book concentrates on the case of the electric field, which is a simple situation where this kind of force occurs. Another important use for vector fields is to measure varying velocity at each point in space (i.e., velocity as a vector, not just speed). This occurs when you want to model the flow of a fluid. The power (and perhaps even beauty) of mathematics is that the same mathematical object can be used to describe very different sorts of problems. So the concepts of vector calculus will apply equally well to electromagnetic theory and to fluid mechanics. For a somewhat challenging exercise, take functions for the polar, cylindrical, and spherical coordinate transformations, and interpret them as vector fields. What do these fields look like?

As a final case, consider $f : \mathbb{R}^2 \rightarrow \mathbb{R}^3$. To visualize such a function, we use the idea of 2.4 and “throw away” the domain and just concentrate on the range. Since the domain is just two-dimensional, the range is likely to be a “thin” two-dimensional surface in three-dimensional space. This sort of a description is called a *parametric surface*. For example, $f(r, s) = \langle 2r - s, s + 1, 3r + s + 1 \rangle$ would describe just a flat plane (since all the coordinates change linearly). On the other hand, see if you can show that $f(r, s) = \langle \sin(r) \cos(s), \sin(r) \sin(s), \cos(r) \rangle$ is a parametric representation for the unit sphere.

2.7 Even more general functions.

To finish this up, we briefly introduce a more abstract example. Let V be the vector space example from the end of the first handout: V denotes the set of all continuous functions from \mathbb{R} to \mathbb{R} . Consider the equation

$$f(t)(x) \stackrel{\text{def}}{=} t^3x - \frac{1}{t^2 + 1}$$

If we had written $f(t, x)$, then we could interpret the equation on the right side as defining $f : \mathbb{R}^2 \rightarrow \mathbb{R}$. But with the above notation, we consider the equation as instead defining $f : \mathbb{R} \rightarrow V$. In other words, for each number $t \in \mathbb{R}$, we fix t and consider the equation $t^3x - \frac{1}{t^2+1}$ as defining a function in V by using x as the variable. Note that $f(t)$ is always a function whose graph is a line, even though t^2 and t^3 appear in the definition!

Here is a more natural example. Fix the simple parabola $y = x^2$. Now, for each number t , let $f(t)$ be the *function* whose graph is the line tangent to the parabola at the point on the graph whose x -coordinate is t . As an exercise, you should write out a definition like the one above for this function f .

What has been hidden throughout Calculus I, II, and III is that derivatives are really functions of this type. However functions whose graphs are lines are determined by just two numbers—the slope and y -intercept. From many points of view, the slope of a tangent line is far more important than its intercept, so we can ignore the latter and just consider the simpler “linear” function of the form $y = mx$ whose graph is parallel to the corresponding tangent line. Since these functions are entirely determined by one number (the slope) we can get away with saying that the derivative f' is just a function from \mathbb{R} to \mathbb{R} . We would much rather deal with $f'(x) = 2x$ than $f'(x)(t) = (2x)(t)$. But the vector situation is more complex, and really the *best* way to consider derivatives in higher dimensions is in this abstract setting. For now, though, we’ll concentrate on studying vector fields. We’ll come back to this topic in more detail later in the semester.

§ 3 MOTIVATION FOR DIFFERENTIABILITY

We have studied partial derivatives, directional derivatives, and the gradient. In the next few sections of these notes we will discuss the more basic notion of a *differentiable* function $f : \mathbb{R}^n \rightarrow \mathbb{R}^m$, and define a single mathematical object which will be the *derivative* of such a function. To indicate that we are working in a general setting, we will use the notation $\vec{f}'(\vec{x})$, but our treatment will apply to cases where n and/or m are equal to one, in which case you can just mentally erase the corresponding arrow(s).

In the simple case where $n = m = 1$, the notion of derivative we define must reduce to the one from first semester calculus. It is worth noting what goes wrong if we try to simply use the old limit definition for derivative:

$$(*) \quad f'(x) \stackrel{?}{=} \lim_{\vec{h} \rightarrow \vec{0}} \frac{\vec{f}'(\vec{x} + \vec{h}) - \vec{f}'(\vec{x})}{\vec{h}}$$

Since \vec{x} and \vec{h} are both n -dimensional, $\vec{f}'(\vec{x} + \vec{h})$ is ok. Since $\vec{f}'(\vec{x} + \vec{h})$ and $\vec{f}'(\vec{x})$ are both m -dimensional, the numerator of the fraction is ok. The limit as $\vec{h} \rightarrow \vec{0}$ would be a well-defined concept, if we were taking the limit of some well-defined expression. But there is no way to reasonably define the division of an m -dimensional by an n -dimensional one. The equation in (*) makes sense only when $n = 1$.

3.1 Exercise. Let $f : \mathbb{R} \rightarrow \mathbb{R}^3$. Show that the definition in (*) reduces to the usual “derivative vector” for vector-valued functions of a single variable. (Hints: use $\vec{f}(t)$ to denote the function in this case, and recall that we can interpret $\vec{f}'(t)$ as the velocity vector for an object whose position at time t is given by \vec{f} . The velocity is computed by just differentiating each of the coordinate functions for f with respect to t .)

One possible way to fix (*) when $n > 1$ is to replace the denominator with $\|h\|$. This would give a coherent definition—but not a very useful one. In the one-dimensional case, the denominator would become the absolute value of h , and the resulting definition would not be the same as the usual derivative. Much worse, however, is the fact that the resulting limit is undefined in even very simple cases.

3.2 Exercise. Let $f : \mathbb{R} \rightarrow \mathbb{R}$ be defined by $f(x) = 2x$. Show that $\lim_{h \rightarrow 0} \frac{f(x+h) - f(x)}{|h|}$ is undefined for every x .

In fact, there is no simple way to patch up this limit. The “definition” in (*) is just no good. The key to remedying the problem is differentials. Let’s recall the idea behind differentials

for functions of one variable. To make the idea clearer, we'll avoid using the dx and dy notations. The concept is simply that, if h is “small,” then $f'(x) \approx \frac{f(x+h)-f(x)}{h}$. If we solve this “approximate equation” for $f(x+h)$, we get $f(x+h) \approx f(x) + f'(x)h$. This should also look familiar from Calculus II, since it is just the beginning of the Taylor series for f —the Taylor series really just improves on the differential approximation by adding the terms $\frac{f''(x)}{2}h^2$, $\frac{f'''(x)}{6}h^3$, etc.

It turns out that the differential approximation idea can be used to *characterize* when a function of one variable is differentiable. We can then generalize the resulting idea in order to *define* the derivative in higher dimensions. We need two important concepts in order to accomplish this. The first is the general notion of a *linear function* from \mathbb{R}^n to \mathbb{R}^m . Linearity will allow us to generalize the “ $f'(x)h$ ” term in the differential approximation. The second concept is that of the “order” of an error. This will let us improve the fuzzy idea of approximately equal, so that we will arrive at an actual definition.

§ 4 LINEARITY

In high-school (or perhaps before), you learned to call equations of the form $y = mx + b$ linear (a name which makes sense because their graphs are lines). In our more general setting, we will say a function is linear if it satisfies two simple, but abstract conditions:

4.1 Definition. Let $f : \mathbb{R}^n \rightarrow \mathbb{R}^m$ be a function. Then f is *linear* means that f satisfies the following two conditions:

- (1) $\vec{f}(\vec{x} + \vec{y}) = \vec{f}(\vec{x}) + \vec{f}(\vec{y})$ for all $\vec{x}, \vec{y} \in \mathbb{R}^n$
- (2) $\vec{f}(\alpha\vec{x}) = \alpha\vec{f}(\vec{x})$ for all $\alpha \in \mathbb{R}$ and $\vec{x} \in \mathbb{R}^n$

We have given this definition only for functions from \mathbb{R}^n to \mathbb{R}^m . However, the definition of linearity extends to the even more abstract setting of general vector spaces, since it refers only to vector addition and scalar multiplication. In fact, this is the primary definition underlying the field of mathematical research called linear algebra. Note also that the above definition does make sense for the simplest case of a function $f : \mathbb{R} \rightarrow \mathbb{R}$ (just remove all the arrows). Unfortunately, this abstract definition doesn't quite correspond to what you learned in high-school:

Theorem 1. Let $f : \mathbb{R} \rightarrow \mathbb{R}$. Then f is linear (in the sense of definition 4.1) if, and only if, there is some real number m such that $f(x) = mx$ for all x .

Proof. First, suppose that f is linear. Let $m = f(1)$. By condition (2) of the definition we have that, for all $x \in \mathbb{R}$, $f(x) = f(x \cdot 1) = x f(1) = mx$. Now suppose that $m \in \mathbb{R}$ and that $f(x) = mx$ for all x . Then $f(x + y) = m(x + y) = mx + my = f(x) + f(y)$ holds for all $x, y \in \mathbb{R}$, and $f(\alpha x) = m(\alpha x) = \alpha(mx) = \alpha f(x)$ holds for all $\alpha, x \in \mathbb{R}$. Thus, f is linear. \square

So not every “ $mx + b$ ” function is really a linear function—only those with $b = 0$ satisfy definition 4.1. In linear algebra, a function is called *affine* if it is equal to the sum of a linear function and a constant function. A function is called *constant* when its value never changes, i.e., when $f(x) = f(y)$ for all x and y in its domain. Thus, the general “ $mx + b$ ” functions should really be called affine, since only some of them are linear, but this distinction has been unimportant in most of your previous classes. Note, however, that the affine functions do include all the linear ones, since the function which is always zero is certainly a constant function.

Now recall the differential approximation equation: $f(x + h) \approx f(x) + f'(x)h$. Think of x as being fixed and consider h to be the variable in this equation. The right hand side of the equation is an affine function of h . More importantly, the term $f'(x)h$ is a linear function of h (with $m = f'(x)$). When $\vec{f}: \mathbb{R}^n \rightarrow \mathbb{R}^m$, we just replace this equation with $\vec{f}(\vec{x} + \vec{h}) \approx \vec{f}(\vec{x}) + \vec{L}_{\vec{x}}(\vec{h})$, where $\vec{L}_{\vec{x}}$ is a linear function from \mathbb{R}^n to \mathbb{R}^m . The subscript \vec{x} indicates that the linear function will change from point to point, just as the tangent slope changes from point to point on a curve. For each \vec{x} , there will be a possibly different linear function $\vec{L}_{\vec{x}}$ corresponding to \vec{x} .

While writing the equation $\vec{f}(\vec{x} + \vec{h}) \approx \vec{f}(\vec{x}) + \vec{L}_{\vec{x}}(\vec{h})$ is simple enough to do, it leaves many things unclear. What does the $\vec{L}_{\vec{x}}$ look like? How can we compute it? When will the approximation be a good one? We start to answer these by first showing that, in general, linear functions from \mathbb{R}^n to \mathbb{R}^m are fairly simple objects. We first consider the $m = 1$ case (i.e., a scalar linear function of several variables).

Theorem 2. *Let $f: \mathbb{R}^n \rightarrow \mathbb{R}$. Then f is linear if, and only if, there is some n -dimensional vector \vec{m} such that $f(\vec{x}) = \vec{m} \cdot \vec{x}$ for all \vec{x} .*

Theorem 2 looks exactly like Theorem 1 except for the arrows and the dot product. It is worth noting that $f(\vec{x}) = \vec{m} \cdot \vec{x}$ can also be written with variable notation as $f(x_1, x_2, \dots, x_n) = m_1 x_1 + m_2 x_2 + \dots + m_n x_n$, where $\vec{x} = \langle x_1, x_2, \dots, x_n \rangle$ and $\vec{m} = \langle m_1, m_2, \dots, m_n \rangle$.

4.3 Exercise. Prove Theorem 2. Hints: find a way to mimic the proof of Theorem 1 on each coordinate. Try the case of $f : \mathbb{R}^3 \rightarrow \mathbb{R}$ first.

In order to handle the most general case of a linear function $f : \mathbb{R}^n \rightarrow \mathbb{R}^m$ we need to first review multiplication of matrices. Recall that if A and B are matrices, then the product AB can be formed only if the number of columns in A is equal to the number of rows in B . More precisely, if A is an $m \times n$ matrix (i.e., A has m rows and n columns) and B is a $n \times k$ matrix, then AB is the $m \times k$ matrix defined by computing the entry at row i and column j by looking at the i^{th} row of A and the j^{th} column of B , multiplying the corresponding entries (n pairs of entries) and then summing the results. Note that if we consider the rows of A and the columns of B to be n -dimensional vectors, then each entry in AB is just the dot product of the corresponding pair of vectors.

In order to use matrices for vector calculus, it is important to realize that a vector can be thought of as a matrix. In fact, there are two ways to do this. Given an n -dimensional vector \vec{v} , we can form a $1 \times n$ matrix which has just one row and whose n columns are taken from the vector. Notationally, this just amounts to replacing the vector brackets with square brackets, i.e., $\langle 1, 5, 7 \rangle$ becomes $[1, 5, 7]$. When we create this matrix, it is often referred to as a *row vector*. Alternatively, we can create an $n \times 1$ matrix from \vec{v} , i.e., $\langle 1, 5, 7 \rangle$ becomes $\begin{bmatrix} 1 \\ 5 \\ 7 \end{bmatrix}$. We now adopt the convention that whenever a vector and matrix are to be multiplied, we will simply interpret the vector as either a column vector or a row vector as necessary. The context will determine how to do this so that the multiplication will make sense. We will then take the matrix that results from the multiplication and interpret it as a vector again. For an example, let $\vec{v} = \langle 1, 5, 7 \rangle$ and let $A = \begin{bmatrix} 1 & 2 & 3 \\ 4 & 5 & 6 \end{bmatrix}$. Then

$$A\vec{v} = \begin{bmatrix} 1 & 2 & 3 \\ 4 & 5 & 6 \end{bmatrix} \begin{bmatrix} 1 \\ 5 \\ 7 \end{bmatrix} = \begin{bmatrix} 32 \\ 71 \end{bmatrix} = \langle 32, 71 \rangle.$$

With this convention in place, we can now see that linear functions from \mathbb{R}^n to \mathbb{R}^m are really just $n \times m$ matrices.

Theorem 3. Let $\vec{f} : \mathbb{R}^n \rightarrow \mathbb{R}^m$. Then \vec{f} is linear if, and only if, there is some $n \times m$ matrix M such that $\vec{f}(\vec{x}) = M\vec{x}$ for all \vec{x} .

4.4 Exercise. Prove Theorem 3. Hint: Just apply Theorem 2 to each of the “coordinate functions” for \vec{f} .

§ 5 ORDER OF ERROR TERMS

The final tool needed to define the general derivative is a way to say that the differential approximation is “very good.” We certainly want the approximation error to go to zero as $\vec{h} \rightarrow \vec{0}$, but we want it to do so “quickly.” There is a general way of comparing the rates at which functions go to zero.

5.1 Definition. Suppose that \vec{f} and \vec{g} are functions from \mathbb{R}^n to \mathbb{R}^m such that $\lim_{\vec{x} \rightarrow \vec{0}} \vec{f}(\vec{x}) = \vec{0}$, $\lim_{\vec{x} \rightarrow \vec{0}} \vec{g}(\vec{x}) = \vec{0}$, and $\vec{g}(\vec{x}) \neq \vec{0}$ for all $\vec{x} \neq \vec{0}$. The notation

$$\vec{f} = o(\vec{g})$$

means that

$$\lim_{\vec{x} \rightarrow \vec{0}} \frac{\|\vec{f}(\vec{x})\|}{\|\vec{g}(\vec{x})\|} = 0.$$

The notation

$$\vec{f} = O(\vec{g})$$

means that there is some positive number M and some $\epsilon > 0$ such that

$$\frac{\|\vec{f}(\vec{x})\|}{\|\vec{g}(\vec{x})\|} \leq M \text{ whenever } \|\vec{x}\| < \epsilon.$$

The ideas behind definitions in 5.1 are simpler than the definitions look. When $\vec{f} = o(\vec{g})$, \vec{f} is getting small “much faster” than \vec{g} is. When $\vec{f} = O(\vec{g})$, \vec{f} is getting small at a rate which is either “faster than” or at least “not much slower than” the rate at which \vec{g} is getting small. We will actually only need the $o(g)$ notation for differentiability (which is fortunate because the definition is simpler), but contrasting the two notions can help to understand them.

You should realize that these notations are an abuse of the equal sign, since they don’t actually assert that any two things are equal. We will carry this abuse even further and write, for example

$$(1) \quad f_1(x) = f_2(x) + o(x^2)$$

to mean that $f_1(x) - f_2(x) = o(x^2)$, or more precisely that

$$(2) \quad \lim_{x \rightarrow 0} \frac{|f_1(x) - f_2(x)|}{x^2} = 0$$

(here we assume that $f_1, f_2 : \mathbb{R} \rightarrow \mathbb{R}$). While equation (2) gives the precise meaning, equation (1) is easier to interpret—it says that f_1 is equal to the function f_2 plus “stuff that goes to zero faster than x^2 does.”

The best way to understand “ o ” and “ O ” is to work a few problems. Be sure you understand (c) and (d) below, since they are the most important ones for what will follow.

5.2 Exercises. Show the following (assume $f, g : \mathbb{R} \rightarrow \mathbb{R}$):

a) $\sin(x^2) = o(x)$

b) $\sin(x) = O(x)$

c) $\sin(x) = x + o(x)$

d) $e^x = 1 + x + \frac{x^2}{2} + o(x^2)$

e) $e^{-\frac{1}{x}} = o(x^k)$ for every positive integer k .

f) If P is an n^{th} degree polynomial, then $P(x) = O(x^n)$ and $P(x) = o(x^k)$ for every positive integer $k < n$.

g) If $f = o(g)$, then $f = O(g)$.

h) If $M > 0$ and $\lim_{x \rightarrow 0} \frac{|f(x)|}{|g(x)|} = M$, then $f = O(g)$ and $g = O(f)$.

A final note on “ o ” and “ O ” is that these notations are also used to define the rate at which functions grow as $x \rightarrow \infty$. If f, g satisfy $\lim_{x \rightarrow \infty} f(x) = \infty$ and $\lim_{x \rightarrow \infty} g(x) = \infty$ then $f = o(g)$ (i.e., $\lim_{x \rightarrow \infty} \frac{f(x)}{g(x)} = 0$) means that f grows more “slowly” than g does, and $f = O(g)$ means that f grows “not much faster than” g does. An especially important use of this notation is to study the efficiency of computer algorithms. Suppose $f(n)$ measures the time that a subroutine will take to sort a list of n numbers. Then $f = O(n)$ would be much better than $f = O(n^2)$, but $f = O(n^2)$ would be much better than $f = O(2^n)$. In reality, one can show that it is impossible to sort in $O(n)$ time, but the best known algorithms come very close by working in $O(n \log(\log(n)))$ time. On the other hand, the security of virtually every data encryption system is based on the belief (still unproven at this point) that the solution to certain problems cannot be computed in less than $O(2^n)$ time.

§ 6 DIFFERENTIABILITY

We have now assembled the necessary tools to define what it means for a function to be differentiable.

6.1 Definition. Let $\vec{f} : \mathbb{R}^n \rightarrow \mathbb{R}^m$ and fix some $\vec{x} \in \mathbb{R}^n$. Then \vec{f} is *differentiable at* \vec{x} means that there is some linear function $\vec{L}_{\vec{x}} : \mathbb{R}^n \rightarrow \mathbb{R}^m$ such that

$$\vec{f}(\vec{x} + \vec{h}) = \vec{f}(\vec{x}) + \vec{L}_{\vec{x}}(\vec{h}) + o(\vec{h}).$$

The linear function $\vec{L}_{\vec{x}}$ is called the *derivative of f at \vec{x}* .

The equation in Definition 6.1 says exactly that the error in the differential approximation $\vec{f}(\vec{x} + \vec{h}) \approx \vec{f}(\vec{x}) + \vec{L}_{\vec{x}}(\vec{h})$ must go to zero *faster* than \vec{h} does. The first thing to mention is that this does, in fact, generalize the Calculus-I definition of derivative.

Theorem 4. *Let $f : \mathbb{R} \rightarrow \mathbb{R}$ and fix some $x \in \mathbb{R}$. Then f is differentiable in the sense of Definition 6.1 if, and only if,*

$$f'(x) \stackrel{\text{def}}{=} \lim_{h \rightarrow 0} \frac{f(x+h) - f(x)}{h}$$

exists. Furthermore, when this occurs, the function $L_x(h)$ from Definition 6.1 must be $L_x(h) = f'(x)h$.

So, for the $\mathbb{R} \rightarrow \mathbb{R}$ case, the only difference between the “new” derivative and the one from elementary calculus is that the derivative at x used to be the number $f'(x)$, while it is now the linear function which is determined by that number.

While Definition 6.1 provides a good way of thinking about differentiability, it doesn't provide much help in figuring out what the linear function $\vec{L}_{\vec{x}}$ should be. We can generalize one direction Theorem 4 in order to help the situation out.

Theorem 5. Let $\vec{f} : \mathbb{R}^n \rightarrow \mathbb{R}^m$, fix some $\vec{x} \in \mathbb{R}^n$, and suppose that \vec{f} is differentiable at \vec{x} and that the derivative of \vec{f} at \vec{x} is $\vec{L}_{\vec{x}}$. Give names for all the coordinate functions and variables involved by letting $\vec{f} = \langle f_1, f_2, \dots, f_m \rangle$, where each $f_i = f_i(x_1, x_2, \dots, x_n)$. Then $\vec{L}_{\vec{x}}(\vec{h}) = J\vec{h}$, where J is the $m \times n$ matrix defined by $J_{i,j} = \frac{\partial f_i}{\partial x_j}(\vec{x})$.

The matrix in Theorem 5 is called the Jacobian matrix for \vec{f} at \vec{x} . It determines the derivative of a function $\mathbb{R}^n \rightarrow \mathbb{R}^m$ in the same way that $f'(x)$ does for functions from $\mathbb{R} \rightarrow \mathbb{R}$. It is important to note that Theorem 5 is lacking an important feature which Theorem 4 has. Theorem 5 does *not* use the phrase “if, and only if.” It would be nice if the existence of all the entries in the Jacobian matrix would imply the differentiability of \vec{f} , but this turns out to not be true! To understand the correct relationship, we need to first review the connections between differentiability and continuity. From Calculus-I you should recall that every differentiable function is continuous. This fact holds in the general case as well.

Theorem 6. If $\vec{f} : \mathbb{R}^n \rightarrow \mathbb{R}^m$ is differentiable at \vec{x} , then \vec{f} is also continuous at \vec{x} .

While the existence of $f'(x)$ implies the continuity of f at x , it does not imply the continuity of $f'(x)$ at x (a fact which is often missed in Calculus-I). In other words, there is a distinction between functions which are differentiable and those which are *continuously differentiable*. The next exercise gives an example of this behavior.

6.2 Exercise. Let $f : \mathbb{R} \rightarrow \mathbb{R}$ be defined by $f(x) = x^2 \sin(1/x)$ for $x \neq 0$ and $f(0) = 0$. Show that f is continuous and differentiable at $x = 0$, but that the function $f'(x)$ is not continuous at $x = 0$.

With more than one dimension, it is even easier to get examples where things go awry. This is shown in the next example.

6.3 Exercise. Let $f : \mathbb{R}^2 \rightarrow \mathbb{R}$ be defined by $f(x, y) = \frac{xy}{\sqrt{x^2 + y^2}}$ for $\langle x, y \rangle \neq \vec{0}$ and $f(\vec{0}) = 0$. Show that f is continuous at $\vec{0}$ and that all the entries of the Jacobian matrix J for f exist at $\vec{0}$, but that f is *not* differentiable at $\vec{0}$. Note also that the partial derivative functions in J are not continuous at $\vec{0}$. (Hint: the Jacobian at $\vec{0}$ only looks along the x and y axes. But the “bad behavior” of f can be seen along the $y = x$ line.)

Exercise 6.3 shows that the existence of partial derivatives is not enough to imply differentiability. The fact that the partial derivatives in 6.3 are discontinuous is important, as Theorem 7 will show. On the other hand, the discontinuity of the partial derivatives does not imply non-differentiability either.

6.4 Exercise. Define a function $f : \mathbb{R}^2 \rightarrow \mathbb{R}$ which has the property that neither of its partial derivatives is continuous at $\vec{0}$, but also has the property that f is differentiable at $\vec{0}$. (Hint: find a “two dimensional” version of the example in exercise 6.2.)

We now come to the most important point: if the partial derivatives are continuous, then \vec{f} must be a differentiable function.

Theorem 7. Let $\vec{f} : \mathbb{R}^n \rightarrow \mathbb{R}^m$, where $\vec{f} = \langle f_1, \dots, f_m \rangle$, and each $f_i = f_i(x_1, \dots, x_n)$. Fix some $\vec{x} \in \mathbb{R}^n$. Suppose that for each i and j , the function $\frac{\partial f_i}{\partial x_j}$ is defined at \vec{x} and is continuous at \vec{x} . Then \vec{f} is differentiable at \vec{x} .

To summarize, we can say the following. If \vec{f} is differentiable at \vec{x} , then \vec{f} must be continuous at \vec{x} and every partial derivative of \vec{f} must exist at \vec{x} . (From this it follows that if \vec{f} is either discontinuous at \vec{x} or any of its partial derivatives are undefined at \vec{x} , then \vec{f} will be non-differentiable at \vec{x} .) On the other hand, if every partial derivative of \vec{f} exists at \vec{x} and is continuous at \vec{x} , then \vec{f} must be differentiable at \vec{x} .

But, there are (rare) difficult cases of continuous functions whose the partial derivatives exist but are not continuous. For such functions, there is no “easy” way to test for differentiability, since examples 6.3 and 6.4 show that such a function may fail to be differentiable, or it may turn out to be differentiable despite the discontinuity of the partial derivatives. For such examples, the only way to test for differentiability is to use the definition directly.

6.5 Exercise. Let $f : \mathbb{R}^3 \rightarrow \mathbb{R}^2$ be defined by

$$\vec{f}(x, y, z) = \langle x^2y + z^3, \cos(x - 1) + 5y^3z \rangle$$

- Use Theorem 7 to argue that \vec{f} is differentiable at every point in \mathbb{R}^3 .
- Find the Jacobian matrix for the derivative of \vec{f} at the point $(1, 1, 1)$.
- Use the matrix from (b) to approximate $\vec{f}(1.01, 1.02, 0.97)$.