Notes for Section 2.2

This is a section with some rather sophisticated concepts. I hope your single variable calculus instructor gave you a good foundation in limits.

The first concept to tackle is open sets. You are familiar with what an open interval on the real line is:
$$(a, b) = \{x \in \mathbb{R} \mid a < x < b\}.$$
What makes this different from a closed interval $[a, b]$ or a half-open (half-closed) interval $(a, b]$ or $[a, b)$ is that it does not contain any of its boundary points. Another way to say this is that all of its points are interior points. Boundary point and interior point both have formal definitions, but I think you can make good sense of these based on the everyday meanings of those words. If you need a little clarification, a boundary point is one that has both points that belong to the set and points that do not belong to the set "next to" it. Of course, there is no such thing on the real line as two points that are next to each other; between any two distinct real numbers, there are plenty more real numbers. But I think this will still make intuitive sense if you think about the points a and b in relation to the open interval (a, b). An interior point on the other hand is one that is some distance away from the boundary. That distance could be very small, but is more than 0. We can make this a little more precise. We will now make this idea little more precise.

Definition 1. Let $x \in \mathbb{R}$ and let ϵ be a positive real number. The *ϵ-neighborhood of x* is $(x-\epsilon, x+\epsilon)$.

You can also call the ϵ-neighborhood of x, the neighborhood of radius ϵ centered at x, or simply the ϵ-ball centered at x. If you think about this a bit, the $(x-\epsilon, x+\epsilon)$ contains exactly those points of $\mathbb{R}$ whose distance from x is less than ϵ:

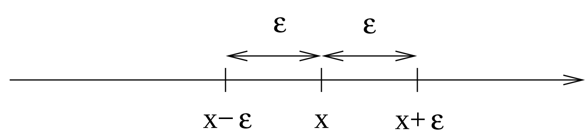

We are saying that the number y is in $(x-\epsilon, x+\epsilon)$ exactly if the distance $|y - x|$ is less than ϵ. In set notation,
$$(x-\epsilon, x+\epsilon) = \{y \in \mathbb{R} \mid |y - x| < \epsilon\}.$$
Hence the radius and ball terminology: a sphere (ball) of radius ϵ contains exactly those points that are within a distance of ϵ from its center. A note on the use of ϵ: we referred to ϵ as a radius, so why not use r? That would be alright, and in fact, our textbook does use r. But it is also customary in calculus and analysis to use ϵ to mean a small positive number and you will see that we will often think about the radius of that neighborhood as small.

Equipped with this notion, we can now say that an interior point of a subset $S \subseteq \mathbb{R}$ is point that has some ϵ-neighborhood that is also contained in S. Notice that it is indeed true that every point $x \in (a, b)$ is an interior point: if we choose ϵ to be small enough that it is no bigger than the distance of x from either endpoint of the interval, that is $\epsilon \leq b - x$ and $\epsilon < x - a$, then $(x-\epsilon, x+\epsilon) \subseteq (a, b)$:

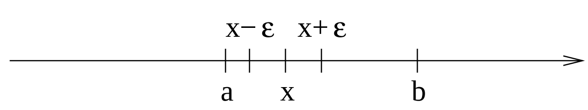

Now, we can generalize the idea of an open interval to describe what it means for a set to be open.

Definition 2. A subset S of $\mathbb{R}$ is open if every point $x \in S$ is an interior point of S, that is every point $x \in S$ has an ϵ-neighborhood that is contained in S.

Open intervals are certainly open sets, but they are not the only open sets. As an example of how a set can fail to be open, consider the closed interval $[a, b]$. No matter how small an ϵ-neighborhood of a you take, it will always contain some numbers, just a little less than a that are not in S. So a is not an interior point.

These ideas are easy to generalize to $\mathbb{R}^n$. Remember that we know how to calculate the distance of two points $(x_1, \ldots, x_n)$ and $(y_1, \ldots, y_n)$:

$$d\big((x_1, \ldots, x_n), (y_1, \ldots, y_n)\big) = ||\vec{x} - \vec{y}|| = \sqrt{(x_1 - y_1)^2 + \cdots + (x_n - y_n)^2}.$$

We can use this to define an ϵ-neighborhood of $\vec{x}$ much the same way as in $\mathbb{R}$.

Definition 3. Let $\vec{x} \in \mathbb{R}^n$ and let ϵ be a positive real number. The *ϵ-neighborhood of $\vec{x}$* is the set

$$D_\epsilon(\vec{x}) = \{\vec{y} \in \mathbb{R}^n \mid ||y - x|| < \epsilon\}.$$

And now we can define interior points and open sets in $\mathbb{R}^n$ exactly the same way as in $\mathbb{R}$.

Definition 4. A subset S of $\mathbb{R}^n$ is open if every point $\vec{x} \in S$ is an interior point of S, that is every point $\vec{x} \in S$ has an ϵ-neighborhood that is contained in S.

It is easy enough to see that if $\vec{x} \in \mathbb{R}^n$ and $r > 0$, then the r-neighborhood of $\vec{x}$, $D_r(\vec{x})$ is itself an open set. See Theorem 1 in the textbook.

We can now make the notion of neighborhood more general.

Definition 5. Let $\vec{x} \in \mathbb{R}^n$. A *neighborhood of x* is any open set $U \subseteq \mathbb{R}^n$ which contains $\vec{x}$.

So the difference between an ϵ-neighborhood of $\vec{x}$ and just any neighborhood of $\vec{x}$ is that the former has a nice regular spherical shape with $\vec{x}$ sitting in the middle, while the latter can have any irregular shape. Finally, we can give precise meaning to a boundary point.

Definition 6. Let $S \subseteq \mathbb{R}^n$. A point $\vec{x}$ is a boundary point of S if every neighborhood of $\vec{x}$ contains some point $\vec{y} \in S$ and some point $\vec{z} \notin S$.

Notice that this definition says very much the same thing as the informal definition of a boundary point I gave at the beginning of these notes. Study Examples 1 and 2 in the text to better understand these concepts.

Let us talk about limits. Let f be a function of real numbers (that is the domain and the codomain are both subsets of $\mathbb{R}$). You will (hopefully) recall from single variable calculus that

$$\lim_{x \to a} f(x) = b$$

means informally that we can make the values of $f(x)$ arbitrarily close to b if we make x sufficiently close (but not equal) to a. The formal definition is

Definition 7. Let $f : A \to B$ where $A, B \subseteq \mathbb{R}$ and let a be an interior point of A. We say

$$\lim_{x \to a} f(x) = b$$

if for every $\epsilon > 0$ there exists a corresponding $\delta > 0$ such that

$$|f(x) - b| < \epsilon$$

whenever

$$0 < |x - a| < \delta.$$

If a is a boundary point of the domain, things are a little more subtle since every ϵ-neighborhood of a, no matter how small, contains some points that are not in S. We sometimes salvage things by using one-sided (left and right) limits instead. We could easily generalize this definition to multivariable functions by using ϵ and δ-neighborhoods. But that is not what our textbook chooses to do. The approach there is even more general and uses general neighborhoods, as in

Definition 8. Let $f : A \to B$ where $A \subseteq \mathbb{R}^n$ and $B \subseteq \mathbb{R}^m$. Let $\vec{a}$ be an interior point of A. We say

$$\lim_{\vec{x}\to\vec{a}} f(\vec{x}) = \vec{b}$$

if for every neighborhood N of $\vec{b}$ there exists a corresponding neighborhood U of $\vec{a}$ such that

$$|f(x) - \vec{b}| \in N$$

whenever

$$\vec{x} \in U$$

and $\vec{x} \neq \vec{a}$.

In fact, it does not make much of a difference whether we use general neighborhoods or ϵ and δ-neighborhoods. This is because any neighborhood of $\vec{b}$ also contains a (smaller) ϵ-neighborhood of $\vec{b}$ and if we can make the values of $f(\vec{x})$ land in that smaller ϵ-neighborhood, then we have also made sure that they land in N. And for U, you can always just choose a small enough δ-neighborhood of $\vec{a}$. Theorem 6 near the end of the section says the same thing, but there is no proof given to explain why.

The book also uses some funky (and nonstandard) language about $f(\vec{x})$ being eventually in N. If you think about it carefully (and this may take some deep reflection), that means the same thing as what I wrote. I prefer what I wrote because it is more parallel to what you may have seen in single variable calculus. In fact, I would suggest that you skip the definition of limit in the book and Examples 3-5 instead to help you understand how this definition works in practice. These are sophisticated concepts and probably quite challenging to understand, especially if your prior calculus courses treated limits lightly. But limits are at the very heart of calculus, so if you really want to understand calculus, you really need to understand limits.

At the end of Section 2.1, there is a part about ϵ's and δ's. Look at the examples in that part for how you can use ϵ and δ-neighborhoods to work with limits. I am not sure why these are example are left to the end of the section. They are not any more technical than the discussion of limits on pp. 91-94. In fact, doing actual calculations with δ and ϵ may make it easier to understand how one can show the existence of a neighborhood U of $\vec{x}$, which corresponds to a given neighborhood N of $\vec{b}$ in the definition of the limit.

Now, the reason the definition of the limit is so long and so complex in the textbook is that they want to consider not only limits at points that lie well inside the domain of a function but also limits at boundary points. There is good reason for this: often the motivation for looking at a limit is exactly that $f(\vec{a})$ does not exist (because $\vec{a}$ is not in the domain) but you can look at the values of f at points really close to $\vec{a}$ and see where they are headed. This is the reason why the definition says $\vec{x} \neq \vec{a}$ (and the single variable version says $0 < |x - a|$). In fact, if you have a sharp eye for detail, you may already have noticed that in Examples 5, $x \to 1$ and 1 is not in the domain of g. But it is a boundary point of the domain. The points that $\vec{x}$ approaches in Examples 13-15 are not in the domain either but are boundary points of the domain. If $\vec{a}$ is a boundary point of A, we can still define the limit at $\vec{a}$ but we need to be careful to exclude points near $\vec{a}$ that are not in A from being used as inputs to f. Here is how we could do that:

Definition 9. Let $f : A \to B$ where $A \subseteq \mathbb{R}^n$ and $B \subseteq \mathbb{R}^m$. Let $\vec{a}$ be a point in A or a boundary point of A. We say

$$\lim_{\vec{x}\to\vec{a}} f(\vec{x}) = \vec{b}$$

if for every neighborhood N of $\vec{b}$ there exists a corresponding neighborhood U of $\vec{a}$ such that

$$|f(\vec{x}) - \vec{b}| \in N$$

whenever

$$\vec{x} \in U \cap A$$

and $\vec{x} \neq \vec{a}$.

I don't think you should stress a lot about these technicalities for now. If you can make sense of the previous definition, where $\vec{a}$ is an interior point of A, that's a really good start. We can worry about the technical subtleties when we need them.

Look at the part about properties of limits. These should look familiar and should also make good intuitive sense. You have seen the corresponding properties in your single variable calculus courses and probably used them to actually calculate limits in most cases. In fact, other than the more daunting notation, which results from having multiple input and sometimes multiple output variables, these properties are the same as in single variable calculus. Property (v) in Theorem 3 talks about component functions. These have so far not been mentioned. Here is what they mean through an example. Let $f : \mathbb{R}^2 \to \mathbb{R}^2$ be the function $f(x, y) = (x^2 - 2y, 5xy + sin(x))$. Then we can refer to $f_1(x, y) = x^2 - 2y$ and $f_2(x, y) = 5xy + sin(x)$ as the component functions of f. Now, it is certainly not immediately clear that every vector-valued function can be expressed in a form that each coordinate of the output is a separate real-valued function. In fact, it is not hard to show that this can always be done, but I would say it is not worth worrying about it at this point. I just want you to understand what those component functions are that are mentioned in property (v). No proof of Theorem 3 or 2 is given and that is just fine. Given the scope and one-semester duration of our course, we can focus on how to use limits than on the details of their properties. If you are itching to play with ϵ and δ neighborhoods, and neighborhoods in general, you can take MCS 220 (unless you already have).

Example 6 may restore your faith that limits are (sometimes) not that hard to work with.

Now the we have some idea of how limits of multivariable functions work, we can tackle the question of continuity. In single-variable calculus, you learned that

Definition 10. Let f be a function of real numbers and x_0 an interior point of its domain. We say f is *continuous at* x_0 if

$$\lim_{x \to x_0} f(x) = f(x_0).$$

First, the reason the definition is for an interior point of the domain is that if x_0 is not an interior point, then any neighborhood of x_0 will have some points in it that are not in the domain, and you will not be able to make $|f(x) - f(x_0)| < \epsilon$ at such a point x because $f(x)$ does not even exist. This is a bit of a technicality and you do not need to spend a lot of your time thinking about it for now. It may be worth pointing out that you can think about continuity in terms of three necessary ingredients: f must have a value at x_0, the limit of f as x approaches x_0 must exist, and these two things must be equal. This is the formal way to capture the intuitive idea that a continuous function is one whose graph can be drawn without lifting up the pencil. This intuitive idea does not easily extend to multivariable functions. After all how can you draw any surface in some higher dimensional space with one stroke of your pencil? But the formal definition above extends just fine and is in fact the definition of continuity for multivariable functions too.

Definition 11. Let $f : A \to B$ where $A \subseteq \mathbb{R}^n$ and $B \subseteq \mathbb{R}^m$. Let $\vec{x}_0$ be a point in A. Then f is continuous at $\vec{x}_0$ if

$$\lim_{\vec{x} \to \vec{x}_0} f(\vec{x}) = f(\vec{x}_0).$$

Here we did not have to require that $\vec{x}_0$ must be an interior point because the way we defined limits for multivariable functions works even if $\vec{x}_0$ is a boundary point. But again, these are technicalities. We can say that f is continuous on a subset S of its domain if it is continous at every point of S. If we simply say f is continuous, we mean that it is continuous at every point in its domain. While the notion of drawing the graph with one stroke of a pencil did not generalize from single-variable to multivariable calculus, many of the typical ways that a function can be discontinuous do, such as holes and sudden jumps in the graph, and places where the values of f go off toward $\pm\infty$.

The properties of continuous functions listed in Theorem 4 are also mostly generalizations of the corresponding properties you learned in single-variable calculus, except for the last one about component functions. It is not surprising that these properties hold in the multivariable setting too as the definition of continuity is basically the same as in single-variable calculus. In fact, they could be proved exactly the same way as their single-variable versions. The composition of multivariable functions is also defined the same way as in single-variable calculus, and Theorem 5 about the continuity of a composite function is a generalization of the corresponding property you learned in single-variable calculus and it holds for the same reason.

Here is a more complex example of a function that is discontinuous at $(0,0)$:

$$f(x,y) = \begin{cases} \frac{2x^2-y^2}{xy} & \text{if } x \neq 0 \text{ and } y \neq 0, \\ 0 & \text{if } x = 0 \text{ or } y = 0. \end{cases}.$$

To see that f is discountinuous at $(0,0)$, we will show that

$$\lim_{(x,y)\to(0,0)} f(x,y) \neq f(0,0) = 0.$$

Let $\epsilon = 1/2$ and let $N = (-1/2, 1/2)$. We will show that there is no neighborhood U of $(0,0)$ such that $f(x,y) \in N$ for every $(x,y) \in U \setminus \{(0,0)\}$. First, note that for any $x \neq 0$,

$$f(x,x) = \frac{2x^2 - x^2}{x^2} = 1.$$

Now, let U be any neighborhood of $(0,0)$. Since $(0,0)$ is an interior point of U, there must be some $\delta > 0$ such that

$$D_\delta(0,0) = \{(x,y) \mid ||(x,y) - (0,0)|| < \delta\}$$

is contained in U. Now, take $x = \delta/2$. Then

$$||(x,x) - (0,0)|| = ||(x,x)|| = \sqrt{x^2 + x^2} = \sqrt{2\left(\frac{\delta}{2}\right)^2} = \sqrt{\frac{\delta^2}{2}} = \frac{\delta}{\sqrt{2}} < \delta,$$

and so $(x,x) \in D_\delta(0,0)$. But as we noted above, $f(x,x) = 1$. So every neighborhood of $(0,0)$ contains a point $(x,x) \neq (0,0)$ where the value of f is 1. Therefore there is no neighborhood U of $(0,0)$ such that every point $(x,y) \in U \setminus \{(0,0)\}$ is in $N = (-1/2, 1/2)$. Try graphing f on the computer to get an idea what surface defined by f looks like near the origin.