

Section 14.1 Functions of Two or More Variables

In Calculus I and II, we spend (almost) all of our time discussing functions with one input and one output. In chapter 13, we discussed vector-valued functions which we can think of as functions with one input and several outputs. In this section, we will explore functions of two or more variables, which have multiple inputs and a single (scalar) output. These functions are difficult to draw when we have $z = f(x, y)$, and they become essentially impossible to draw with three or more input variables. In order to visualize many of these functions, we consider level curves (and with more than two inputs level surfaces). If you have access to Matlab, Maple, or Mathematica, you can use them to draw many of these functions. Not on a quiz or exam, of course, but while you are trying to visualize what is going on with examples or to check your work on textbook problems.

Yes, we will certainly discuss functions which have both multiple inputs and multiple outputs. These are vector fields which are the focus of chapters 16 and 17 at the end of this course.

Find the following definitions/concepts:

- function of n variables
- What does the domain of a function of n variables look like? (Think mostly about $n = 2$ and $n = 3$ or your head will start to hurt. Trust me on this one.)
- What does the range of a function of n variables look like?
- Not in the book: do you know the difference between codomain and range? If not, you should ask me.
- Graphing a function of two variables
- vertical trace
- horizontal trace
- level curve
- contour map
- contour interval
- average rate of change (as it pertains to a function of two variables)
- path of steepest ascent (teaser: this relates to the gradient in section 14.5)
- level surface

The motivating examples above example 1 are intended to give you a feel for when you will need to use multivariable functions to model real-world phenomena.

Example 1 describes how to find the domain of a function of several variables. This is a must-read (and understand).

Examples 2 through 7 are all about being able to visualize graphs of functions of two variables. Any time you have $z = f(x, y)$, you are going to get some graph in \mathbb{R}^3 which includes exactly one point above, at, or below each point in the xy -plane which is in your domain. If f is continuous, the graph is a surface. You should spend some serious time working through these examples. Practicing the techniques in these examples will help develop your intuition about what these functions are doing at various points in space, which will certainly make the rest of the chapter more manageable.

Example 8 is the first hint about what derivatives and slopes are going to look like for functions of several variables. When you were in middle school, your math teacher talked about average rate of change. At some later time, your math teacher started calling it slope and wrote equations of lines on the board. Guess where we're headed...

The motivation for example 9 is in the paragraph above it. There are many real-world situations where we are interested in what some quantity is doing at every point in some 3-dimensional space. We could be thinking about the temperature of a room, a planet, or a universe. We could be interested in how much light will reach every point in a hotel ballroom for each of several different possible ways to install lighting fixtures. (We might also be interested in things like velocities and electric fields, but for those you have to wait until chapter 16). The figures on pages 796-797 are level surfaces. When we saw level curves, those were two-dimensional representations of what was happening in the third dimension (each curve representing a function value above/below that point in the xy -plane). Here, each surface is a collection of points at which the function has the same value. The actual graph has four dimensions, since the graph consists of tuples $(x, y, z, f(x, y, z))$. But these three dimensional surfaces are supposed to help us understand the function the same way a contour map is supposed to help us visualize terrain.

Section 14.2 Limits and Continuity in Several Variables

Back in Calculus I (probably in Precalculus as well), you learned about limits and continuity. When you defined $\lim_{x \rightarrow c} f(x)$, you were told that in order for this limit to exist that the two one-sided limits had to exist and be equal. And then you were told that in order for a function to be continuous at some point $x = c$ that $\lim_{x \rightarrow c} f(x)$ and $f(c)$ both had to exist and they had to be equal. The continuity part of that works pretty much the same way for functions of several variables.

But limits are very different in one particular way. When you only have one real input variable, your domain is just part of the number line. If you look at a point on the number line, there are only two ways to approach it: from below and from above. But if we are talking about $f(x, y)$, the domain is some part of the plane \mathbb{R}^2 . For an arbitrary point in the middle of your domain, there are infinitely many paths that you can follow to get there. The limit only exists if it is the same on any of these infinitely many paths.

Find the following definitions/concepts/formulas/theorems:

- limit of a function of two variables ($\epsilon - \delta$ definition)
- Theorem: limit laws (functions of two variables)
- continuity (functions of two variables)
- Theorem: composition of continuous functions (functions of two variables)
- techniques for showing that the limit of a function of two variables does not exist

Examples 1-4 are all short and direct applications of rules to (reasonably) well-behaved functions. Make sure that you look at the justifications carefully, though. One of the most common errors on this type of question is applying some limit law or theorem incorrectly and getting a value for a limit when the limit does not actually exist. You have been warned.

Example 6 shows how careful you have to be when you are trying to show that a limit exist as you approach a hole in the domain. You may need to read this one a few times.

Examples 5 and 7 show various ways to prove that a limit does not exist. It's very nice when you find two different limits as you approach the point on two different lines or curves. The frustrating part is that sometimes you consider lots of different paths and get the same limit every time, but you can't find one that's different. And you also can't find a theorem that tells you that the function actually has a limit. This is what a lot of higher level math (and much of mathematical research) is like. You try to prove that something is true for a while and can't. Then you spend some time trying to find a counterexample that shows the something is false and you can't. But some of the counterexamples you tried gave you a new technique that you use to try and prove the theorem... and you still can't prove the theorem. But it is very satisfying when you end up with either a valid proof or a valid counterexample after a long struggle.