Calculus 251:C3 Reading Guide - 6/11/2020

Section 14.8 Lagrange Multipliers: Optimizing with a Constraint

The basic idea of Lagrange multipliers is this: if you want to maximize a function, you *want* to move in the direction of the gradient from wherever you are. If you are subject to a constraint, you may not be able to move in the direction of the gradient while staying "in bounds" for the constraint. So you do the best you can and move along your constraint curve (or surface) in the direction closest to the gradient direction. When you get to a place where the gradient is orthogonal to the constraint where you are, moving in either direction won't increase the function value (recall: orthogonal to gradient=along level curve/surface). So where you are is at least a local maximum subject to the constraint.

Find the following definitions/concepts/formulas/theorems:

- constraint (you have seen this before)
- Lagrange multiplier - general concept
- Theorem: Lagrange Multipliers
- Lagrange condition
- critical point/critical value (specifically for optimization with a constraint)

The Graphical Insight on page 866 is another way of saying what I wrote in the opening paragraph.

The proof of the Lagrange Multipliers theorem is not too hard to follow. Again, if you plan to be a math major, you should spend some time considering it. I would suggest doing so after going through a few examples, because that will help with both the motivation and with the intuition for why the proof works.

Examples 1 and 2 are both quite long, but they are about as simple as Lagrange Multiplier examples can be. The setup isn't too bad in either example, and they perform the same steps in the same order. You should certainly try to work your way through them. The Assumptions Matter comment between the examples talks about the types of constraints where this method will always work. What are they?

Example 3 is very quick even though it involves a function of 3 variables. The method doesn't really change much when you increase the number of input variables. But...

Example 4 gets tricky because we have more than one constraint. When we have multiple constraints, we have to set up one Lagrange multiplier for each constraint. In this case, the Lagrange condition is that the gradient of our objective function is a linear combination of the gradients of the constraints. You may not want to look at this one until after the lecture. Hopefully it will make some sense then.