Math 1c	TA: Padraic Bartlett
Recitation 1: Open and Closed Sets; Limits and Continuity	
Week 1	*Caltech 2013*

1 Administrivia

Here are most of the random administrative details for the course:

- Recitation time: 2-3pm, in 102 Steele.
- My email: padraic@caltech.edu
- My office: 156 Sloan.
- My office hours: 7-8pm on Sunday night, in either my office or an ambiently open classroom on the first floor of Sloan (depending on people.) Also by appointment.
- My website: www.its.caltech.edu/∼padraic. Course notes for every recitation will be posted here, usually within a few days of the recitation. The course notes for this quarter will likely be similar to last year's: I'll try to update/change things around to match the problem sets and lectures, but if you're curious for what I'm likely to talk about, 2012's Ma1c rec webpage will likely be of some use.
- HW policy: The course-wide policy is that you're allowed one free extension, as long as you let us know ahead of time. Subsequent extensions can/will only be granted with a note from the deans or the health center. It bears noting that both entities are reasonably kind; if your reason for needing more time isn't something like "up all night playing LoL," they'll typically grant an extension.

2 Level Curves

The idea behind level curves is pretty simple: suppose that we have a function $f(x, y) : \mathbb{R}^2 \to \mathbb{R}$. If we consider $f(x, y)$ to be a "height" function, we can interpret the "graph" of $f(x, y) = z$ as the set of all values $\{(x, y, f(x, y)) : (x, y) \in \mathbb{R}^2\}$:

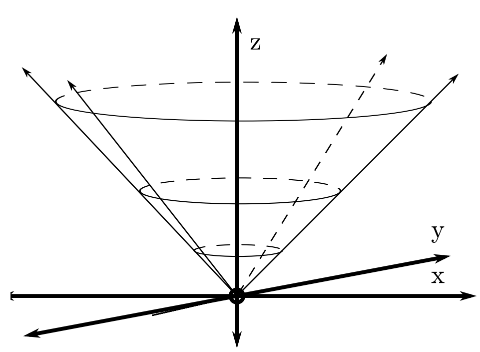

(The graph of $f(x, y) = \sqrt{x^2 + y^2}$, a cone centered around the positive z-axis.)

So: given a function $f(x, y)$, how can you quickly and accurately get an idea of what the graph of $f(x, y)$ looks like?

- Snarky answer: Mathematica.
- Serious answer: level curves!

Definition. Take a function $f(x, y)$, and some height value h. The **level curve** of $f(x, y)$ at height h is just the set of points (x, y) in $\mathbb{R}^2$ that satisfy the equation

$$f(x, y) = h.$$

To quickly get an idea of what the graph of a function $f(x, y)$ is, it usually suffices to find the level curves of $f(x, y)$ for a handful of height values, and then to draw each level curve in $\mathbb{R}^3$ on its corresponding 2-dimensional plane $\{(x, y, h) : (x, y) \in \mathbb{R}^2\}$. To illustrate this process, we calculate a quick example below:

Question 1 *Draw some of the level curves of*

$$f(x, y) = \sqrt{25 - x^2 - y^2}.$$

What does this shape look like?

Solution. Pick a few promising-looking values of h: say, $h = 0, 1, 2, 3, 4, 5$. For each value of h, then, let's find the solutions to the equation

$$h = \sqrt{25 - x^2 - y^2}.$$

If we square both sides, subtract h^2 from both sides, and add $x^2 + y^2$ to both sides we have

$$x^2 + y^2 = 25 - h^2.$$

The solutions of this equation is precisely a circle of radius $\sqrt{25 - h^2}$; therefore, for our given values of h, we have that these level curves are circles of radius $5, \sqrt{24}, \sqrt{21}, 4, 3, 0$. Graphing these values shows us that $f(x, y)$ is a hemisphere of radius 5, as drawn below:

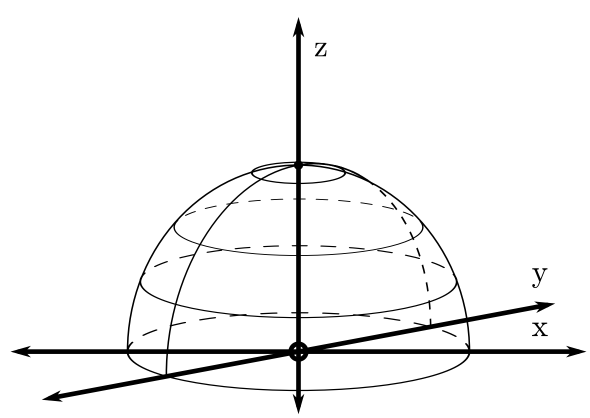

Level curves are useful for much more than just drawing graphs. One quick example comes from physics, and can be described as follows:

- Suppose that you have a particle p with mass m whose position is constrained to the parabola $y = x^2$:

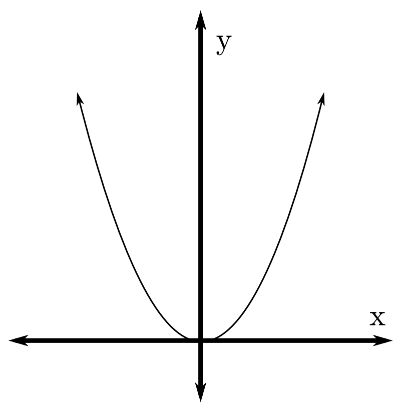

- At any time, roughly speaking, we can break the **total energy** possessed by this particle into two quantities: its **kinetic energy**, which is given by $\frac{1}{2}mv^2$, and its **potential energy**, which (if we assume that our parabola is in a lab on the surface of the earth, and that the only force acting on it is gravity) is roughly $9.8m \cdot x^2$.
- Fun Fact of Physics: in any closed system, energy is **conserved.** In other words, if we assume that our particle's motion is frictionless and that our system is otherwise closed, the total energy posessed by our particle is always constant.
- So, if you want to view this in terms of level curves, consider the following way to describe our particle in three dimensions:

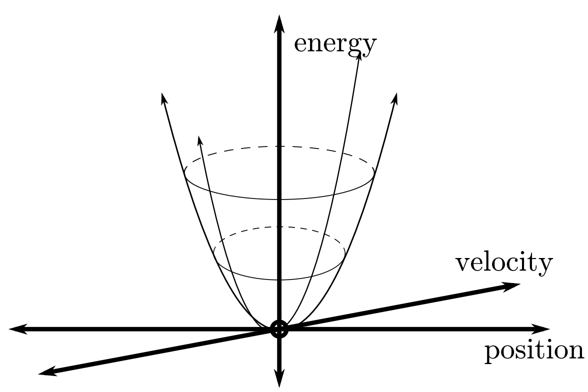

Suppose that our particle starts off with energy e: then the set of possible position/velocity pairs it can have is precisely the level curve of

$$\frac{1}{2}mv^2 + 9.8mx^2.$$

at height e. In particular, if our particle starts with energy 0, we can see that it must be at rest at the bottom of our parabola!

This is a relatively simple toy version of this problem; in general, however, you can use these methods to transform positional knowledge of particles into knowledge of its velocity (and vice-versa.)

3 Limits and Continuity in $\mathbb{R}^n$

3.1 Basic definitions.

We now turn to a discussion of limits and continuity as they exist in $\mathbb{R}^n$:

Definition. For a set $D \subseteq \mathbb{R}^n$, values $\mathbf{a} \in \mathbb{R}^n, \mathbf{L} \in \mathbb{R}^m$, and a function $f : D \to \mathbb{R}^m$, we say that

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

if

- (informally:) As $\mathbf{x}$ goes to $\mathbf{a}$, $f(\mathbf{x})$ goes to $\mathbf{L}$.
- (symbolically:) $\forall \epsilon > 0, \exists \delta > 0$ s.t. $\forall \mathbf{x} \in D, \mathbf{x} \neq \mathbf{a}$, we have that $((||\mathbf{x} - \mathbf{a}|| < \delta)$ implies $(||f(\mathbf{x}) - \mathbf{L}|| < \epsilon))$.

Notice that this definition is almost completely identical to the one we came up with in Ma1a, for $\mathbb{R}$! The only difference is that we've replaced $|x-a|$ and $|f(x)-L|$'s with $||\mathbf{x}-\mathbf{a}||$ and $||f(\mathbf{x}) - L||$; this is because in $\mathbb{R}^n$, we measure distance using the Euclidean norm $|| \cdot ||$, which happened to be equal to taking the absolute value $| \cdot |$ when $n = 1$.

Definition. A function $f : D \to \mathbb{R}^n$ is **continuous** at $\mathbf{a} \in D$ iff

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

Note that this definition is *exactly* the same as it was for $\mathbb{R}$.

Just like in Math 1a, it's easy to feel like there's a big gap between seeing the definitions for limits and continuity and being able to actually **use** them. Here, we describe a few common strategies for how you can **use** these definitions:

- **Combining limits:** Often, the easiest way to compute limits is just to use your past knowledge from Ma1a, alongside the fact that we know how limits behave under operations like composition, product, sum, and quotient. Specifically, we know that if $f, g : \mathbb{R}^n \to \mathbb{R}^m$ are a pair of functions whose limits exist as $\mathbf{x} \to \mathbf{a}$, we have
  - $\lim_{\mathbf{x}\to\mathbf{a}}(f(\mathbf{x}) + g(\mathbf{x})) = (\lim_{\mathbf{x}\to\mathbf{a}} f(\mathbf{x})) + (\lim_{\mathbf{x}\to\mathbf{a}} g(\mathbf{x}))$.
  - $\lim_{\mathbf{x}\to\mathbf{a}}(f(\mathbf{x}) \cdot g(\mathbf{x})) = (\lim_{\mathbf{x}\to\mathbf{a}} f(\mathbf{x})) \cdot (\lim_{\mathbf{x}\to\mathbf{a}} g(\mathbf{x}))$.
  - $\lim_{\mathbf{x}\to\mathbf{a}} \left(\frac{f(\mathbf{x})}{g(\mathbf{x})}\right) = (\lim_{\mathbf{x}\to\mathbf{a}} f(\mathbf{x})) \,/\, (\lim_{\mathbf{x}\to\mathbf{a}} g(\mathbf{x}))$,, if $g(x) \neq 0$ near $\mathbf{a}$, and $\lim_{\mathbf{x}\to\mathbf{a}} g(\mathbf{x}) \neq 0$.

Similarly, we can compose limits as well: if $f : \mathbb{R}^n \to \mathbb{R}^m$ is a function such that $\lim_{\mathbf{y}\to\mathbf{a}} f(\mathbf{y}) = \mathbf{L}$, and $g : \mathbb{R}^k \to \mathbb{R}^n$ is a function such that $\lim_{\mathbf{x}\to\mathbf{b}} g(\mathbf{x}) = \mathbf{a}$, then

$$\lim_{\mathbf{x}\to\mathbf{b}} f(g(\mathbf{x})) = \mathbf{L}.$$

In particular, this tells you that the composition of continuous functions is continuous, as is their product, sum, and quotient (provided the denominator is nonzero.)

Using these tools, you can break lots of functions down into individual components which you know are continuous / whose limits you know from Ma1a, and use this to deduce the original limit.

- **Using the definition:** If the above strategy doesn't work, you can always turn to the definition of the limit. One "blueprint" for a proof that $\lim_{\mathbf{x}\to\mathbf{a}} f(\mathbf{x}) = \mathbf{L}$ via the definition is the following:

  1. First, examine the quantity

     $$||f(\mathbf{x}) - \mathbf{L}||.$$

     Specifically, try to find a simple upper bound for this quantity that depends only on $|\mathbf{x} - \mathbf{a}|$, and goes to 0 as $\mathbf{x}$ goes to $\mathbf{a}$ – something like $|\mathbf{x} - \mathbf{a}| \cdot$ (constants), or $|\mathbf{x} - \mathbf{a}|^3 \cdot$ (some bounded functions).
  2. Using this simple upper bound, for any $\epsilon > 0$, choose a value of δ such that whenever $|\mathbf{x} - \mathbf{a}| < \delta$, your simple upper bound $|\mathbf{x} - \mathbf{a}| \cdot$ (bounded things) is $< \epsilon$. Often, you'll define δ to be $\epsilon/$(upper bound on bounded things), or something like that.
  3. Plug in the definition of the limit: for any $\epsilon > 0$, we've found a δ such that whenever $|\mathbf{x} - \mathbf{a}| < \delta$, we have

     $$|f(\mathbf{x}) - \mathbf{L}| < (\text{simple upper bound depending on } |\mathbf{x} - \mathbf{a}|) < \epsilon.$$

     Thus, we've proven that $\lim_{\mathbf{x}\to\mathbf{a}} f(\mathbf{x}) = \mathbf{L}$, as claimed.

- **Proving discontinuity/that a limit DNE**: Often, you will want to prove that some function is in fact discontinuous at a point, or that it has no limit at a point. A simple blueprint for showing that a function f has no limit at a given point $\mathbf{a}$ is the following:

  1. Find a pair of paths (i.e. continuous functions $\gamma_1, \gamma_2 : \mathbb{R} \to \mathbb{R}^n$) such that

     $$\gamma_1(0) = \gamma_2(0), \text{ and}$$

     the single-variable limits

     $$\lim_{t\to 0} f(\gamma_1(t)) \neq \lim_{t\to 0} f(\gamma_2(t))$$

     Common choices of γ are $\gamma(t) = (t, 0, 0, \ldots 0) + \mathbf{a}$, the path that approaches $\mathbf{a}$ along the first coördinate and holds all the others constant; other popular choices are moving along a different coördinate than the first, as well as the path $\gamma(t) = (t, t + \ldots t) + \mathbf{a}$ in which you move along all coördinates simultaneously.

2. If you have found two such paths, then you know that f cannot have a limit at $\mathbf{a}$; along these two paths, our function seems to be going to two different values, which is impossible if a limit exists.

Similarly, to show that a function is not continuous, it suffices to simply find one path $\gamma(t), \gamma(0) = \mathbf{a}$ along which

$$\lim_{t\to 0} f(\gamma(t)) \neq f(0).$$

To illustrate these techniques, we work a few examples:

3.2 Some worked examples.

Question 2 *Consider the function $f(x, y) : \mathbb{R}^2 \to \mathbb{R}$, defined as follows:*

$$f(x, y) = \frac{y^2}{x^2 + y^2}, \qquad (x, y) \neq (0, 0)$$

Does $f(x, y)$ have a limit at $(0, 0)$?

Solution. As it turns out (and as you might be able to guess after attempting to graph or sketch this function), no! This function is not continuous at $(0, 0)$.

To prove this, we turn to the idea of **paths** that we discussed above: can we find two paths $\gamma_1(t), \gamma_2(t) : \mathbb{R} \to \mathbb{R}^2$ such that the function $f(x, y) = \frac{y^2}{x^2+y^2}$ goes to one value on γ_1, and a different value on γ_2?

After trying out a few paths, it's not hard to find a few which go to different values. A few examples: if $\gamma_1(t) = (t, 0)$, we have that $\gamma_1(0) = (0, 0)$, and that

$$\lim_{t\to 0} f(\gamma_1(t)) = \lim_{t\to 0} \frac{0^2}{t^2 + 0^2} = 0$$

conversely, if $\gamma_2(t) = (t, t)$, we have that $\gamma_1(0) = (0, 0)$, and that

$$\lim_{t\to 0} f(\gamma_1(t)) = \lim_{t\to 0} \frac{t^2}{t^2 + t^2} = \frac{1}{2}.$$

So these two paths go to different values: therefore, our function does not have a limit as $(x, y) \to (0, 0)$.

Question 3 *Consider the function $f(x, y, z) : \mathbb{R}^3 \to \mathbb{R}$, defined as follows:*

$$f(x, y) = \frac{xyz}{x^2 + y^2 + z^2}, \qquad (x, y) \neq (0, 0)$$
$$0, \qquad (x, y, z) = (0, 0).$$

Is $f(x, y, z)$ continuous everywhere?

Solution. Given the earlier problem, you might expect this function to be discontinuous at $(0,0,0)$; however, after about ten minutes of trying to find sequences that converge to any other nonzero value, you might begin to doubt this intuition.

Which, as it turns out, is the correct move – because this function is continuous! To prove that this is continuous at every point (a,b,c), we can consider two cases: either $(a,b,c) \neq (0,0,0)$, or $(a,b,c) = (0,0,0)$. In the first case, we have that our function is just something created by taking the three continuous functions $\pi_1(x,y,z) = x, \pi_2(x,y,z) = y, \pi_3(x,y,z) = z$ and their products/sums/quotients. Because continuity is preserved under products/sums/quotients where the denominator is nonzero, we're looking at nonzero points, and the denominator of $\frac{xyz}{x^2+y^2+z^2}$ is nonzero whenever we're looking at a nonzero point, we know $\frac{xyz}{x^2+y^2+z^2}$ is continuous at every nonzero point.

So we only have to look at the point $(0,0,0)$. Here, we cannot use our earlier methods, because we're looking at a quotient where the denominator is going to 0; instead, we must use the definition to prove our function is continuous here. Specifically, we can use the blueprint we discussed above for proving things are continuous:

- We start by taking the quantity $||f(\mathbf{x} - L)||$, and try to come up with a simple upper bound on it. In this case, we have for all $(x,y,z) \neq (0,0,0)$,

$$\begin{aligned} ||f(x,y,z) - (0,0,0)|| &= \left| \frac{xyz}{x^2+y^2+z^2} \right| \\ &\leq \left| \frac{\max\{|x|^3,|y|^3,|z|^3\}}{x^2+y^2+z^2} \right| \\ &\leq \left| \frac{\max\{|x|^3,|y|^3,|z|^3\}}{\max\{x^2,y^2,z^2\}} \right| \\ &= |\max\{|x|,|y|,|z|\}|. \end{aligned}$$

  (In order to observe the inequalities above, we used two very useful tricks. The first is where we bounded a polynomial expression xyz from above by assuming all of your variables $x,y,z\ldots$ were just the largest one $\max\{|x|,|y|,|z|\}$. The second is where we bounded another polynomial $x^2+y^2+z^2$ from below by only taking the largest monomial $\max\{|x|^2,|y|^2,|z|^2\}$ in that polynomial. These are both very useful and frequently-occuring tricks!)

- Now, we want to bound the $||\mathbf{x} - \mathbf{a}||$ portion of our proof from below, so that it is related to the simple upper bound we just got. In this case, we can use the observation that

$$\begin{aligned} ||(x,y,z) - (0,0,0)|| &= \sqrt{x^2+y^2+z^2} \\ &\geq \sqrt{\max\{x^2,y^2,z^2\}} \\ &= |\max\{|x|,|y|,|z|\}|. \end{aligned}$$

- Now, given any $\epsilon > 0$, use this knowledge to pick a value of $\delta > 0$ such that whenever

$||\mathbf{x}-\mathbf{a}|| < \delta$, $||f(\mathbf{x}-L)|| < \epsilon$! In particular, for our example, we've shown the following:

$$||f(x, y, z) - (0, 0, 0)|| \leq |\max\{|x|, |y|, |z|\}|, \text{ and}$$
$$|\max\{|x|, |y|, |z|\}| \leq ||(x, y, z) - (0, 0, 0)||.$$

Therefore, if we pick $\delta = \epsilon$, we will have

$$||f(x, y, z) - (0, 0, 0)|| \leq |\max\{|x|, |y|, |z|\}| \leq ||(x, y, z) - (0, 0, 0)|| < \delta = \epsilon,$$

which concludes our proof. So we're done!