Uniform Continuity

Ryan Acosta Babb
r.acosta-babb@warwick.ac.uk

MA260 Support Class – Week 3

Summary

A brief motivation and review of uniform continuity in $\mathbb{R}$. We discuss the importance of quantifier order and variable dependence through an example that helps motivate the need for uniform continuity. This review is intended as a "warm up" before tackling compactness.

Let $f : (0,1) \to \mathbb{R}$ be a continuous function and $0 < a_n < \frac{1}{n}$ for all $n \in \mathbb{N}$.

Question Does $(f(a_n))_{n\in\mathbb{N}}$ converge?

A rough attempt follows:

"Proof 1". Note that $a_n \to 0$ and f is continuous, so $f(a_n) \to f(0)$. □

This clearly will not do, as f is not necessarily defined at $x = 0$! Indeed,

Example 1. Take $f(x) := \frac{1}{x}$. Then $f(a_n) > n$ for all $n \in \mathbb{N}$, so it clearly does not converge.

What if we strengthen our assumptions to prevent the blow-up at the end-points?

Question 2 Suppose $f : (0,1) \to \mathbb{R}$ is continuous *and bounded*, with a_n as above. Does $(f(a_n))_{n\in\mathbb{N}}$ converge?

Our first "proof" failed because we didn't know the behaviour of f at the "limit point" $f(0)$. If we want to show that $f(a_n)$ converges without knowing *a priori* what the limit might be, we can check whether it is a Cauchy sequence!

Proof 2. Let $\varepsilon > 0$. Then,

(i) By continuity of f, there exists a $\delta > 0$ such that

$$|x - y| < \delta \implies |f(x) - f(y)| < \varepsilon.$$

(ii) By convergence of $(a_n)_{n\in N}$, it is a Cauchy sequence, so there is an $N \in \mathbb{N}$ such that, for all $n, m \geqslant N$,

$$|a_n - a_m| < \delta.$$

(iii) Hence, by (i) and (ii),

$$|f(a_n) - f(a_m)| < \varepsilon \quad \text{for all} \quad n, m \geqslant N.$$

This proves that $(f(a_n))_{n\in\mathbb{N}}$ is Cauchy in $\mathbb{R}$ and, therefore, convergent. □

Does it? Well, no. Our proof only used the assumption that f is *continuous* (we never invoked boundedness!), and we saw in Example 1 that $(f(a_n))_{n\in\mathbb{N}}$ does not converge in general under this assumption. So let us take a closer look at our "proof" to see what is going on.

In (i) we applied continuity to get a $\delta > 0$, which we then fed into the definition of Cauchy sequence to produce the index $N \in \mathbb{N}$. But we run into trouble when we try to apply the continuity of f to the condition $|a_n - a_m| < \delta$, since this is expected to hold *for all* a_n, a_m with $n, m \geqslant N$. However, the δ only tells us about points y within δ of a *specific* x!

Let us run through the argument again, noting carefully the dependence of the different constants:

(i) For any $x \in (0, 1)$ there is a $\delta_x > 0$ *depending on* x (and ε) such that

$$|x - y| < \delta_x \implies |f(x) - f(y)| < \varepsilon.$$

(ii) For each δ_x there is an $N_{\delta_x} \equiv N_x \in \mathbb{N}$ *depending on* x (actually δ_x, which depends on x and ε, but we care about the x dependence) such that

$$|a_n - a_m| < \delta_x \quad \text{for all} \quad n, m \geqslant N_x.$$

We therefore have possibly *uncountably many* N_x to choose from! How are we to pick one N to rule them all?

The issue is that regular continuity is a *local property*: it tells us how f changes near a single point, say x. So while we may know how to bound $f(y)$ to within ε of $f(x)$ for y close to x (say within δ of x), if we now move to a different point x' we have to start all over again, since getting δ-close to x' may not be enough to bound $|f(x') - f(y)| < \varepsilon$ (same ε and δ as before).

A better behaved function might not care about the specific choice of x, but only the closeness δ.

Example 2. Let $f(x) := |x|$. Then, by the reverse triangle inequality, we have

$$|f(x) - f(y)| \leqslant |x - y| \quad \text{for all} \quad x, y \in \mathbb{R}.$$

Thus, choosing $\delta = \varepsilon$ for given $\varepsilon > 0$, we have

$$|f(x) - f(y)| < \varepsilon \tag{1}$$

at *any two points* with $|x - y| < \varepsilon$. It doesn't matter whether we pick $x = 0$ and $y = \varepsilon/2$ or if we take $x = -400 + \varepsilon/67$ and $y = -400$: as long as they span an interval no greater than ε, the bound (1) holds. (See Figure 1.)

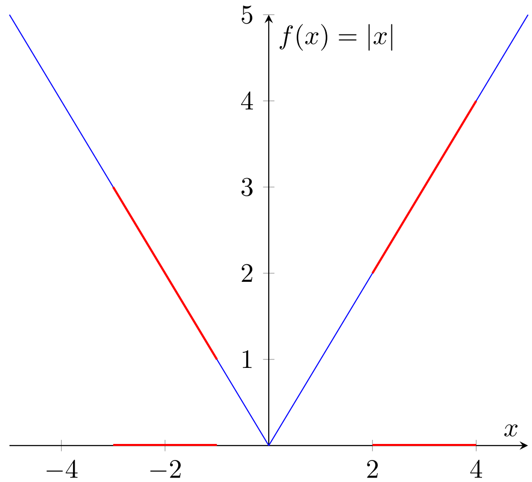

Figure 1: The red marks identify intervals of width ε. No matter where we take them on the x-axis, the range of the function f stays within a vertical interval of height ε.

Example 3. Consider now $g(x) := \frac{1}{x}$ defined on $(0, 1)$. Take $\varepsilon = 1/2$. Then, for any $\delta \in (0, 1)$ we can choose $x = \delta/4$ and $y = 3\delta/4$ so clearly

$$|x - y| < \delta \quad \text{and} \quad \left|\frac{1}{x} = \frac{1}{y}\right| = \frac{8}{3\delta} \geqslant \varepsilon$$

(as $0 < \delta < 1$). Provided we pick x and y close to 0, where the function g blows up, we can make their difference as small as we like, and still $|g(x) - g(y)| \geqslant \varepsilon$.

If, however, we fix $x \in (0, 1)$, for any $\varepsilon > 0$ we can take $\delta = \frac{x^2\varepsilon}{2}$ (depending on x!). This "pinning of x in place" gives us more control over $g(y)$ for y close to x. Compare this to the case above, were we needed to squish *both x and y* close to 0 to get the blow up. (See Figure 2.)

This you should recognise as *uniform continuity* from Analysis III. (If you don't, go back and revise your notes!) Compare:

(i) A function $f : (a, b) \to \mathbb{R}$ is *continuous* if

$$\forall \varepsilon > 0 \forall x \in (a, b) \exists \delta_{x,\varepsilon} > 0 \forall y \in (a, b) \left(|x - y| < \delta_{x,\varepsilon} \implies |f(x) - f(y)| < \varepsilon\right)$$

(ii) A function $f : (a, b) \to \mathbb{R}$ is *uniformly continuous* if

$$\forall \varepsilon > 0 \exists \delta_\varepsilon > 0 \forall x \in (a, b) \forall y \in (a, b) \left(|x - y| < \delta_\varepsilon \implies |f(x) - f(y)| < \varepsilon\right)$$

Note the quantifier flip! Thus, the δ no longer depends on x.

We can now state and answer our question correctly and affirmatively:

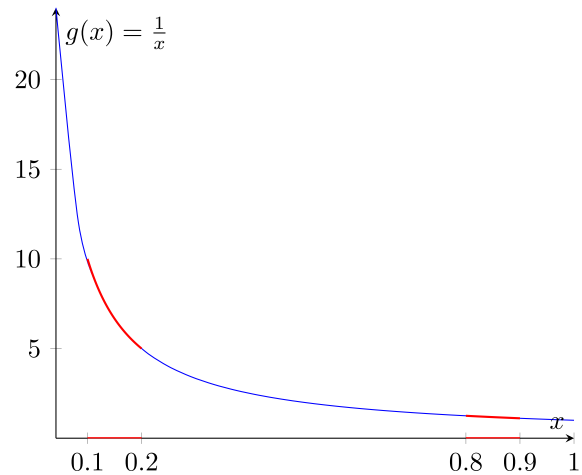

Figure 2: The interval of width 0.1 from 0.1 to 0.2 yields a vertical jump of 5 units, while an interval of the same length measured at 0.8 yields a jump of about 0.14.

Proposition 4. *Suppose* $f : (0, 1) \to \mathbb{R}$ *is* uniformly continuous*, with* $0 < a_n < \frac{1}{n}$*. Then* $(f(a_n))_{n\in\mathbb{N}}$ *converges in* $\mathbb{R}$*.*

Proof. We use the uniform continuity of f to free our δ from the restriction of x.

(i) Let $\varepsilon > 0$. By uniform continuity of f, find $\delta > 0$ such that

$$|x - y| < \delta \implies |f(x) - f(y)| < \varepsilon \quad \text{for all} \quad x, y \in (a, b).$$

(ii) With this *single* $\delta > 0$, apply the Cauchy criterion to $(a_n)_{n\in\mathbb{N}}$ to get an $N \in \mathbb{N}$ (depending on δ, but that's fine!) such that

$$|a_n - a_m| < \delta \quad \text{for all} \quad n, m \geqslant N.$$

(iii) Taking $x = a_n$ and $y = a_m$ in (i), (ii) implies that

$$|f(a_n) - f(a_m)| < \varepsilon \quad \text{for all} \quad n, m \geqslant N.$$

This proves that $(f(a_n))_{n\in\mathbb{N}}$ is Cauchy in $\mathbb{R}$, and therefore convergent. □

The key step in the proof is finding a δ that works *uniformly* on the a_n, a_m, without being tied down to a particular choice of x, since the a_n are changing (what if we pick $n = N + 1$ or $n = N + 2$ or $n = N + 3$ or ...?) and we need to control all of them in one go. This sort of idea is of crucial importance in analysis, as you will have ample opportunity to discover in later modules. For now, keep it in mind, as it will be immensely helpful in understanding the notion of *compactness*.