CONVERGENCE OF SEQUENCES

MATH 153, SECTION 55 (VIPUL NAIK)

Corresponding material in the book: Section 11.3.

What students should already know: The $\epsilon - \delta$ definition of limit. The definition and basic terminology related to a sequence.

What students should definitely get: All the ϵ-type definitions of limits over the reals and for all sequences. The statements of the uniqueness theorem, pointwise combinations, composition theorem, and pinching theorem. How changing finitely many terms, left shift, infinitary permutations, and taking subsequences do not change the limit. The relation between $\mathbb{N}$-limits and $\mathbb{R}$-limits. The relationship between boundedness, monotonicity, and convergence.

Executive summary

Words ...

(1) We say that $\lim_{n \to \infty} a_n = L$ if, for every $\epsilon > 0$, there exists $n_0 \in \mathbb{N}$ such that for all $n > n_0$ with $n \in \mathbb{N}$, we have $|a_n - L| < \epsilon$. We say that the sequence (a_n) *converges* to L.
(2) Similarly, we say that $\lim_{n \to \infty} a_n = \infty$ if, for every $a \in \mathbb{R}$, there exists $n_0 \in \mathbb{N}$ such that for all $n > n_0$, we have $a_n > a$.
(3) Similarly, we say that $\lim_{n \to \infty} a_n = -\infty$ if, for every $a \in \mathbb{R}$, there exists $n_0 \in \mathbb{N}$ such that for all $n > n_0$ we have $a_n < a$.
(4) If, for a continuous function f, $\lim_{x \to \infty} f(x)$ is finite, then the limit of the corresponding sequence $f(n), n \in \mathbb{N}$ is the same finite number. Similarly, if $\lim_{x \to \infty} f(x) = \infty$, then the limit of the sequence $f(n), n \in \mathbb{N}$ is ∞, and if $\lim_{x \to \infty} f(x) = -\infty$, then the limit of the sequence $f(n), n \in \mathbb{N}$ is also $-\infty$.
(5) However, it is possible that the sequence $f(n), n \in \mathbb{N}$ has a finite limit or that its limit is $+\infty$ or $-\infty$, but that the limit $\lim_{x \to \infty} f(x)$ is neither finite nor $+\infty$ or $-\infty$.
(6) *Pointwise combination theorems for limits*: The limit of the sum is the sum of the limits, the limit of the difference is the difference of the limits, the limit of the product is the product of the limits, and the limit of the quotient is the quotient of the limits.
(7) There is a pinching theorem for limits of sequences, just as there is a pinching theorem for limits of functions.
(8) There is a composition theorem for limits of sequences: if f is a continuous function, and $a_n \to L$, then $f(a_n) \to f(L)$.
(9) A non-increasing sequence bounded from below converges to its greatest lower bound. Similarly, a non-decreasing sequence bounded from above converges to its least upper bound.
(10) A sequence is bounded if and only if it is eventually bounded.
(11) Any convergent sequence is bounded.
(12) If a sequence is bounded and eventually monotonic, then it is convergent.

Actions ...

(1) For a sequence that is obtained by iterating a continuous function f, i.e., for a sequence given by the recursion $f(a_n) = a_{n+1}$, if the limit exists and is finite, then $f(L) = L$.

1. Limits and convergence: a review

These review notes assume that you are familiar with the $\epsilon - \delta$ definition of limits. If you are a little shaky about these definitions, please review the notes titled "informal introduction to limits" and "formal definition of limit" from the Math 152 course. These explain the $\epsilon - \delta$ definition from a number of angles in a lot more details.

1.1. **What does limit mean?** As A goes to B, C goes to D. In the language of limits, we say that $\lim_{A\to B} C = D$. This is very similar to an analogy, and we here try to understand what such a limit statement means.

Let's first restrict ourselves to functions, i.e., we are interested in:

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

In other words, the image variable is a function of the domain variable that we are sending to a particular point.

What does this mean? One incorrect interpretation is that we can get as close to L as we want. That interpretation would be that:

> For any neighborhood of L, however small, there exist points x close to c such that $f(x)$ is in that neighborhood of L.

The reason this definition is inadequate is because, to capture a notion of limit or convergence, we need more than just occasional or even frequent proximity. We need *definitive* proximity. We need to make sure that we can *trap* the function value close to L.

Thus, a better description is:

> For any neighborhood of L, however small, there exists a neighborhood of c such that for all $x \neq c$ in that neighborhood of c, $f(x)$ is in the original neighborhood of L.

This *is* the $\epsilon - \delta$ definition, albeit without an explicit use of the letters ϵ and δ. Rather, I have used the term *neighborhood* which has a precise mathematical meaning. Making things more formal in the language we are familiar with, we can say:

> For any open ball centered at L, however small, there exists an open ball centered at c such that for all $x \neq c$ in that open ball, $f(x)$ lies in the original open ball centered at L.

An open ball is described by means of its radius, so if we use the letter ϵ for the radius of the first open ball and the letter δ for the radius of the second open ball, we obtain:

> For any $\epsilon > 0$, there exists $\delta > 0$ such that for all $x \in (c - \delta, c + \delta) \setminus \{c\}$, we have $f(x) \in (L - \epsilon, L + \epsilon)$.

Or, equivalently:

> For any $\epsilon > 0$, there exists $\delta > 0$ such that for all x satisfying $0 < |x - c| < \delta$, we have $|f(x) - L| < \epsilon$.

Although it is the final formulation that we use, the first two formulations are conceptually better because they avoid unnecessary symbols and are also easier to generalize to other contexts.

1.2. **Neighborhoods of infinity.** A literal interpretation of the $\epsilon - \delta$ definition at ∞ is problematic. However, going back to the neighborhood definition, we see that if we can somehow define a notion of *neighborhood* of ∞, we can make sense of cases where c and/or L is ∞. Similarly, a notion of *neighborhood* of $-\infty$ allows us to plug in the value $-\infty$ for c and/or L.

Here is the idea:

(1) The neighborhoods of ∞ of interest to us are sets of the form (a, ∞) for $a \in \mathbb{R}$, i.e., the sets $\{x : x > a\}$.
(2) The neighborhoods of $-\infty$ of interest to us are sets of the form $(-\infty, a)$ for $a \in \mathbb{R}$, i.e., the sets $\{x : x < a\}$.

We can now come up with all the precise definitions:

(1) $\lim_{x\to c} f(x) = \infty$ means that for every $a \in \mathbb{R}$, there exists $\delta > 0$ such that if $0 < |x - c| < \delta$, then $f(x) > a$.
(2) $\lim_{x\to c} f(x) = \infty$ means that for every $a \in \mathbb{R}$, there exists $\delta > 0$ such that if $0 < |x - c| < \delta$, then $f(x) < a$.
(3) $\lim_{x\to\infty} f(x) = L$ means that for every $\epsilon > 0$, there exists $a \in \mathbb{R}$ such that for $x > a$, $|f(x) - L| < \epsilon$.
(4) $\lim_{x\to-\infty} f(x) = L$ means that for every $\epsilon > 0$, there exists $a \in \mathbb{R}$ such that for $x < a$, $|f(x) - L| < \epsilon$.
(5) $\lim_{x\to\infty} f(x) = \infty$ means that for every $a \in \mathbb{R}$, there exists $b \in \mathbb{R}$ such that if $x > b$, then $f(x) > a$.
(6) $\lim_{x\to\infty} f(x) = -\infty$ means that for every $a \in \mathbb{R}$, there exists $b \in \mathbb{R}$ such that if $x > b$, then $f(x) < a$.

(7) $\lim_{x\to-\infty} f(x) = \infty$ means that for every $a \in \mathbb{R}$, there exists $b \in \mathbb{R}$ such that if $x < b$, then $f(x) > a$.
(8) $\lim_{x\to-\infty} f(x) = -\infty$ means that for every $a \in \mathbb{R}$, there exists $b \in \mathbb{R}$ such that if $x < b$, then $f(x) < a$.

Intuitively, the neighborhoods now are sets that go off the deep end, rather than nice nests that cocoon the element.

1.3. **One-sided limits.** For one-sided limits, the notion of neighborhood gets replaced by a corresponding one-sided notion. Specifically, instead of an interval *centered* at the point, we look for an interval that *ends* at the point. Specifically:

(1) $\lim_{x\to c^-} f(x) = L$ means that for every $\epsilon > 0$, there exists $\delta > 0$ such that for $0 < c - x < \delta$, we have $|f(x) - L| < \epsilon$. This is called the *left-hand limit*.
(2) $\lim_{x\to c^+} f(x) = L$ means that for every $\epsilon > 0$, there exists $\delta > 0$ such that for $0 < x - c < \delta$, we have $|f(x) - L| < \epsilon$. This is called the *right-hand limit*.

1.4. **Corresponding notions of continuity.** We say that:

(1) f is *continuous* at c if $\lim_{x\to c} f(x) = f(c)$.
(2) f is *left-continuous* at c if $\lim_{x\to c^-} f(x) = f(c)$.
(3) f is *right-continuous* at c if $\lim_{x\to c^+} f(x) = f(c)$.
(4) f has a *removable discontinuity* at c if $\lim_{x\to c} f(x)$ exists but is not equal to $f(c)$.
(5) f has a *jump discontinuity* at c if $\lim_{x\to c^-} f(x)$ and $\lim_{x\to c^+} f(x)$ exist but are not equal.
(6) f has an *infinite discontinuity* at c if either or both one-sided limits is $+\infty$ or $-\infty$.

1.5. **A question begging to be asked.** We looked at a notion of one-sidedness that led us to define left-hand limits and right-hand limits. However, we applied the one-sidedness modification only to domain approach. What happens if we apply it to range approach? What we get is the usual notion of limit with the additional constraint that the value of the function approaches its limiting value from one side.

There is also a related notion of *semicontinuity* that is very important in the real world but that we do not have the time to explore here.

2. Convergence of sequences

We now turn to the question of when a sequence $a_n, n \in \mathbb{N}$, can be said to have a limit as $n \to \infty$. To do this, we need to get a somewhat better understanding of $\mathbb{N}$ as a set.

2.1. **Natural numbers: discrete, but clustering at infinity.** The set of natural numbers is discrete, i.e., each one is far from every other one. Thus, for a given natural number m, it does not really make sense to take the limit as $n \to m$. Any such attempts will just end up looking at the function values at $m - 1$ and $m + 1$. The problem is that *we do not have arbitrarily small neighborhoods of natural numbers containing other natural numbers*.

However, it is still feasible to develop an interesting theory of what happens as $n \to \infty$. Although the natural numbers form a discrete set, we can nonetheless think of them as clustering at ∞. The formal definitions we use are precisely the same as the definitions that we use for real variables. Specifically:

(1) $\lim_{n\to\infty} f(n) = L$ for some finite L means that for all $\epsilon > 0$, there exists $n_0 \in \mathbb{N}$ such that $|f(n)-L| < \epsilon$ for all $n > n_0$. (Note: we can replace either of the $>$ signs with $\geq$ signs without changing the spirit of the definition.)

    The notation $\lim_{n\to\infty} a_n = L$ is sometimes abbreviated as $a_n \to L$. We also say that the sequence (a_n) *converges* to L. Eliding L from the sentence, we say that (a_n) is a *convergent sequence*.

    Note that when we just say $\lim_{n\to\infty} f(n)$, it is potentially ambiguous if f makes sense as a function both on the real numbers and on the natural numbers. The best thing is to clearly indicate that $n \in \mathbb{N}$. Sometimes, this is clear from the context.
(2) $\lim_{n\in\mathbb{N}, n\to\infty} f(n) = \infty$ means that for all $a \in \mathbb{R}$, there exists $n_0 \in \mathbb{N}$ such that $f(n) > a$ for all $n > n_0$.

    We can write $\lim_{n\to\infty} a_n = \infty$ as $a_n \to \infty$. In words, we say that a_n *tends to infinity*.
(3) $\lim_{n\to\infty} f(n) = -\infty$ if for all $a \in \mathbb{R}$, there exists $n_0 \in \mathbb{N}$ such that $f(n) < a$ for all $n > n_0$.

We sometimes say that a sequence is *divergent* if it is not convergent. However, some people use that term only for sequences that go to $+\infty$ or $-\infty$. Thus, avoid using the term unless you're sure how your listeners will interpret you.

2.2. How do limits in $\mathbb{N}$ and $\mathbb{R}$ compare?

Suppose f is a function on $\mathbb{R}$. We can consider the limit:

$$\lim_{x \in \mathbb{R}, x \to \infty} f(x)$$

and the limit:

$$\lim_{n \in \mathbb{N}, n \to \infty} f(n)$$

What can we say about these limits? The following:

(1) If the $\mathbb{R}$-limit exists, the $\mathbb{N}$-limit exists and the two limits are equal.
(2) If the $\mathbb{R}$-limit is $+\infty$, so is the $\mathbb{N}$-limit. If the $\mathbb{R}$-limit is $-\infty$, so is the $\mathbb{N}$-limit.
(3) It is possible for the $\mathbb{N}$-limit to exist but the $\mathbb{R}$-limit to not exist. This may happen, for instance, for a $\mathbb{R}$-periodic function that is constant on $\mathbb{N}$, such as $f(x) := \sin(\pi x)$. More generally, if the $\mathbb{N}$-limit exists and is finite, the $\mathbb{R}$-limit must either equal that value or be undefined – and it cannot head to $+\infty$ or $-\infty$.
(4) It is possible for the $\mathbb{N}$-limit to be $+\infty$ and the $\mathbb{R}$-limit to not exist. For instance, consider $f(x) := x\cos(2\pi x)$. Restricted to $\mathbb{N}$, this coincides with the identity function, and goes to ∞. However, on $\mathbb{R}$, it is a wildly oscillatory function and the oscillation ends approach $\pm\infty$ as $x \to \infty$.

2.3. Basic ideas for computing limits in $\mathbb{N}$.

Here are some of the key ideas:

(1) For a sequence that is constant or eventually constant, the limit is that constant value. For a nonconstant periodic sequence, there is no limit.
(2) If we know how to take the $\mathbb{R}$-limit for the function that describes the sequence, then that gives the $\mathbb{N}$-limit, if the $\mathbb{R}$-limit is finite, $+\infty$, or $-\infty$. Note that if the $\mathbb{R}$-limit is not defined because the function is oscillatory, the $\mathbb{N}$-limit may still be defined.
(3) Note that we have flexibility about how to extend the function from $\mathbb{N}$ to $\mathbb{R}$, and we should exercise this with discretion (in the *discreet* sense, not the *discrete* sense). For instance, the functions $x\cos(2\pi x)$ and x both restrict to the same function on $\mathbb{N}$. However, when thinking of limits at infinity, it is the latter form in which we should consider the function.

3. Theorems on limits

3.1. The uniqueness and pointwise combination theorems.

These theorems are unsurprising analogues of the theorems we already saw for limits in the finite real context in the first quarter. Specifically:

(1) The *uniqueness theorem* for limits asserts that if $\lim_{n\to\infty} a_n$ exists and is finite, then it is unique.
(2) The limit of the sum is the sum of the limits, the limit of the difference is the difference of the limits, the limit of the product is the product of the limits, and the limit of the quotient is the quotient of the limits. However, there is a unidirectionality to these claims. For instance, when we say that $\lim_{n\to\infty}(a_n + b_n) = \lim_{n\to\infty} a_n + \lim_{n\to\infty} b_n$, this statement holds true *if both sides are defined and finite*. However, something stronger is true: if the right side is defined and finite, so is the left side and they are equal. On the other hand, it is possible for the left side to be defined and finite but the right side to not make sense. Analogous observations hold for differences and products. For quotients, an additional complication arises because of the $0/0$ form, which we deal with explicitly a little later.
(3) The composition theorem, which says that *for f a continuous function*, if $\lim_{n\to\infty} a_n = L$, then $\lim_{n\to\infty} f(a_n) = f(L)$. In other words, the function value at the limit is the limit of the function values.

3.2. **The pinching theorem.** Recall the *pinching theorem* for limits, also called the *sandwich theorem* or the *squeeze theorem*. For the real numbers, the *left hand* version says that if $f(x) \le g(x) \le h(x)$ for $x \in (c - \delta, c)$, and if $\lim_{x \to c^-} f(x) = \lim_{x \to c^-} h(x) = L$, then $\lim_{x \to c^-} g(x) = L$. A similar version holds for right hand limits, two-sided limits, and limits at infinity or taking values at infinity. And, the same idea holds for $\mathbb{N}$-limits at infinity.

3.3. **Sufficient condition for the existence of a limit.** These statements are useful to show that the limit exists and is finite, and to obtain an approximate value, even if the precise value is elusive:

(1) A sequence that is bounded and non-decreasing has a limit. Moreover, this limit equals the least upper bound of the underlying set of the sequence.
(2) A sequence that is bounded and non-increasing has a limit. Moreover, this limit equals the greatest lower bound of the underlying set of the sequence.
(3) A sequence that is bounded and monotonic has a limit. This just combines the last two statements. Remember that *monotonic* means non-increasing or non-decreasing.

4. Perturbations and deflections

4.1. **Changing finitely many terms.** Here are some observations:

(1) If (a_n) and (b_n) are two sequences such that, for all but finitely many n, $a_n = b_n$, i.e., the two sequences are *eventually equal*, then their limiting behavior is the same
(2) If one sequence is obtained by applying a left shift to another sequence, their limiting behavior is the same.

In other words, *changes in the finite parts do not affect limits.*

4.2. **Taking subsequences.** Suppose S is an infinite subset of $\mathbb{N}$ and (a_n) is a sequence of real numbers. The *subsequence* of (a_n) corresponding to S is the sequence obtained by retaining only those a_n for which $n \in S$, and re-indexing. So, the first term of the subsequence is a_k where k is the smallest element of S, and more generally, the m^{th} element of the subsequence is a_n where n is the m^{th} smallest element of S.

For instance, the subsequence corresponding to the subset of even numbers is the subsequence whose m^{th} term is a_{2m}. The subsequence corresponding to the subset of odd numbers is the subsequence whose m^{th} term is a_{2m-1}.

The following are true:

(1) If a sequence has a limit, then every subsequence of the sequence has the same limit. In other words, every subsequence of a convergent sequence is convergent and to the same limit.
(2) It is possible for a subsequence of a sequence to have a limit but for the sequence itself to have no limit.
(3) if two different subsequences of a sequence have different limits, the sequence has no limit. An example of this is nonconstant periodic sequences, which can be split up into subsequences each of which is constant, but where the constants differ.

4.3. **Infinitary permutations.** Another fact that is surprising at first, but obvious after reflection, is this:

> A permutation (i.e., rearrangement), even an infinitary one, of the terms of a sequence does not change the limiting behavior of the sequence.

This means that even if we shuffle infinitely many terms of the sequence, the limit is unaffected.

This is at first surprising because, in the definition of limit, we have a condition of the form: "there exists n_0 such that for all $n > n_0$, some condition holds." In other words, the "neighborhoods" of ∞ that we are using are sets of natural numbers greater than a particular natural number. It seems that an infinitary permutation, which could make large elements of $\mathbb{N}$ small and small elements of $\mathbb{N}$ large, could alter the nature of neighborhoods of ∞.

In fact, the alteration is not significant and does not affect our definition. The reason? We can reconceptualize neighborhoods of ∞ as subsets of $\mathbb{N}$ whose complement is finite, i.e., subsets that miss only finitely many elements. This definition is clearly unaffected by permutations.

5. Convergence, boundedness and eventual behavior

For sequences, we have often given one property of sequences and then added another property by tacking on the adverbial modifier *eventually*. In the previous section, we observed, essentially, that tacking on the *eventually* modifier does not alter the limiting behavior. For instance, we had the notion of eventually constant and eventually periodic.

We can similarly define *eventually increasing*, *eventually decreasing*, *eventually non-decreasing*, *eventually non-decreasing*, and *eventually monotonic*. We make the following observations:

(1) A sequence is bounded if and only if it is *eventually* bounded: Here's the explanation. If a sequence is eventually bounded, then that means that, throwing out finitely many of the initial terms, the rest of the sequence is bounded. But now, since we have thrown out only finitely many terms, we can consider the all those terms and the bounds on the remaining infinite sequence, and we get a global bound on the sequence.

For instance, suppose we have a sequence $a_1, a_2, \dots$, and we know that if we throw off the first four terms, the rest of the sequence is bounded with lub U and glb L. Then, the lub for the whole sequence is $\max\{a_1, a_2, a_3, a_4, U\}$ and the glb for the whole sequence is $\min\{a_1, a_2, a_3, a_4, L\}$.

(2) Any convergent sequence is eventually bounded, and hence, is bounded: Here's the explanation. Suppose (a_n) is a convergent sequence. So, there exists some limit L. Now, pick $\epsilon = 1$. We know that there exists n_0 such that $|a_n - L| < 1$ for all $n > n_0$, so $a_n \in (L - 1, L + 1)$ for all $n > n_0$. Thus, the sequence (a_n) is eventually bounded. From the previous observation, the sequence must be bounded.

(3) If a sequence is bounded and eventually monotonic, it is convergent: Throwing off the first few terms, we get a sequence that is monotonic. Since the original sequence was bounded, so is this truncated subsequence. Now, if it is non-decreasing, we see that it converges to its least upper bound, and if it is non-increasing, we see that it converges to its greatest lower bound. This follows from a little tinkering with definitions.