

A PROOF OF THE BOUNDEDNESS THEOREM BY INDUCTION

Theorem (Boundedness Theorem). *A continuous real-valued function on a closed interval is bounded.*

Proof. Suppose f is such a function and $[a, b]$ is its domain. First, observe that for every $c \in [a, b]$, since f is continuous at c , there is a positive $\delta \in \mathbb{R}$ such that for every $x \in [a, b] \cap (c - \delta, c + \delta)$, we have $f(x) \in (f(c) - 1, f(c) + 1)$. So there is a neighbourhood of c , namely $(c - \delta, c + \delta)$, on which f is bounded. Thus f is, in a sense, *locally* bounded at every point in its domain; the problem is to prove that this local boundedness implies *global* boundedness.

In textbook proofs of the boundedness theorem, this is generally done using what I would regard as a trick, such as supposing f isn't bounded and using the Bolzano-Weierstrass theorem to obtain a contradiction. More advanced texts may appeal to the compactness of $[a, b]$, but the proof that $[a, b]$ is compact (the Heine-Borel theorem) amounts to basically the same logic, and is usually no less tricky¹.

However, if we think about how to construct an *algorithm* to find a global bound (often a helpful move to make—the un-situated-in-time-ness of ordinary mathematical language can be quite thought-limiting) then there is a procedure which, in my opinion, is quite obvious and immediately suggests itself. Just start on the left, with the singleton interval $[a, a] = \{a\}$, on which f is certainly bounded, and repeatedly apply local boundedness to the right endpoint, gradually expanding the subinterval of $[a, b]$ on which we know f to be bounded. At each step the existing subinterval has a bound, and the neighbourhood of the right endpoint has a bound; taking the maximum of these two bounds gives us a bound of the whole expanded subinterval formed by unioning² the neighbourhood with the existing subinterval. Eventually the right endpoint will reach b and we will no longer be able to expand it further, at which point we stop and observe that we have bounded the whole of $[a, b]$. (Sanity check: why would this fail for (a, b) ? Because then we wouldn't be able to take the right endpoint to b ; there would be nowhere to stop the procedure, without leaving a part of the right of (a, b) unbounded.)

Of course, this algorithm will not in general terminate in a finite number of steps, but this is no issue; don't let your thinking be limited by the arbitrary limitations of physical machines. Obviously, it is still necessary to prove that this algorithm *will* terminate, even if it takes infinitely many steps. Since we're now dealing with the distinctions among the infinite quantities, we're going to need to use some set-theoretic machinery (which is why you won't see this proof in introductory real analysis textbooks). I'll assume that you are familiar with the ordinals, the techniques of transfinite induction and recursion, and Burali-Forti's paradox (which says that there is no set of all ordinals).

¹This word doesn't appear to exist currently, but it needs inventing.

²Another word that needs inventing. OK, I could use "uniting", but I suspect people wouldn't immediately make the connection with the union operation on sets if I used that word.

Here's the plan. Using transfinite recursion, we shall construct an ordinal-indexed sequence $\langle x_\alpha \rangle$ of members of $[a, b]$ such that every ordinal α has the following properties:

- (1) The function f is bounded on $[a, x_\alpha]$.
- (2) We have $x_\alpha \leq x_{\alpha+1}$, and if $x_\alpha = x_{\alpha+1}$, then $x_{\alpha+1} = b$.

Then, since there are a proper class's worth of ordinals and $[a, b]$ is just a set, the sequence $\langle x_\alpha \rangle$ will have to repeat³, so there will be ordinals α and β such that $\alpha < \beta$ and $x_\alpha = x_\beta$. By (2), the sequence $\langle x_\alpha \rangle$ will be nondecreasing, so we will have $x_\alpha \leq x_{\alpha+1} \leq x_\beta$ and hence $x_\alpha = x_{\alpha+1} = x_\beta = b$. Then by (1), it will follow that f is bounded on $[a, b]$.

All we need to do to complete the proof is describe the recursion and make it clear that it is a valid recursion (i.e. all of the assumed properties used to construct the next term of the sequence are preserved by the constructed sequence with the next term added).

Base case: Let $x_0 = a$ and observe that f is bounded on $[a, a] = \{a\}$ by $|f(a)|$.

Successor case: For every ordinal α , if we assume that $x_\alpha \in [a, b]$, then since f is continuous, there is a positive $\delta \in \mathbb{R}$ such that f is bounded on $[x_\alpha, x_\alpha + \delta]$ by some M . Let $x_{\alpha+1} = \min(b, x_\alpha + \delta)$, and observe that:

- If $b \leq x_\alpha + \delta$, then $x_{\alpha+1} = b$ and hence $a \leq x_\alpha \leq x_{\alpha+1} \leq b$. Otherwise, we have $x_{\alpha+1} = x_\alpha + \delta$ and hence $a \leq x_\alpha < x_{\alpha+1} \leq b$, with the middle inequality strict, so that we can only have $x_\alpha = x_{\alpha+1}$ when $x_{\alpha+1} = b$.
- We have $x_{\alpha+1} \leq x_\alpha + \delta$, so f is bounded by M on $[x_\alpha, x_{\alpha+1}]$ as well as $[x_\alpha, x_\alpha + \delta]$. Assuming f is bounded on $[a, x_\alpha]$ by some L , it follows that f is bounded on $[a, x_{\alpha+1}]$ by $\max(L, M)$.

Limit case: For every limit ordinal λ , let $A = \{x_\alpha : \alpha < \lambda\}$. Then A contains x_0 and hence is nonempty, and is bounded above by b , so it has a supremum. Let $x_\lambda = \sup A$, and observe that:

- If we assume that $A \subseteq [a, b]$, then since A is nonempty and x_λ is ge every member of A , we have $x_\lambda \geq a$; and since b bounds A above and x_λ is the smallest upper bound of A , we also have $x_\lambda \leq b$.
- Since $x_\lambda \in [a, b]$ and f is continuous, there is a positive $\delta \in \mathbb{R}$ such that f is bounded on $[x_\lambda - \delta, x_\lambda]$ by some M . And since $x_\lambda = \sup A$, there is an ordinal $\alpha < \lambda$ such that $x_\lambda - \delta < x_\alpha$ and hence f is bounded on $[x_\alpha, x_\lambda]$ by M . Assuming f is bounded on $[a, x_\alpha]$ by some L , it follows that f is bounded on $[a, x_\lambda]$ by $\max(L, M)$.

And that's it.

Of course, this proof is a bit tedious. There might be ways to make it shorter (maybe we can use Zorn's lemma instead of transfinite induction). But the advantage it has over other proofs of the boundedness theorem I've seen is that it falls out more or less automatically, without requiring any flashes of insight. $\square$

³This is a sort of infinitary pigeonhole principle.