The Axiom of Choice

Ryan Maguire

Summer 2022

Throughout our discussion of metric spaces we've repeatedly used the axiom of choice, but neglected mentioning it. Let's review Lebesgue's number lemma. We wish to show that for any open cover $\mathcal{O}$ of a compact metric space $(X,\,d)$ there is a $\delta>0$ such that for all $x\in X$ there is a $\mathcal{U}\in\mathcal{O}$ such that $B_{\delta}^{(X,\,d)}(x)\subseteq\mathcal{U}$. We say *suppose not*. We are negating our claim. Using quantifier notation, our claim says:

$$\forall_{\mathcal{O}\subseteq\tau_d,\bigcup\mathcal{O}=X}\exists_{\delta>0}\forall_{x\in X}\exists_{\mathcal{U}\in\mathcal{O}}\Big(B_{\delta}^{(X,\,d)}(x)\subseteq\mathcal{U}\Big) \tag{1}$$

the negation of this swaps $\forall$ with $\exists$ and vice-versa.

$$\exists_{\mathcal{O}\subseteq\tau_d,\bigcup\mathcal{O}=X}\forall_{\delta>0}\exists_{x\in X}\forall_{\mathcal{U}\in\mathcal{O}}\Big(B_{\delta}^{(X,\,d)}(x)\nsubseteq\mathcal{U}\Big) \tag{2}$$

That is, there exists an open cover $\mathcal{O}$ of X such that for all $\delta>0$ there is an $x\in X$ with the property that for every $\mathcal{U}\in\mathcal{O}$ the δ ball centered at x is not entirely contained in $\mathcal{U}$. In particular, since the negation says this is true *for all* $\delta>0$, it is true for $\delta=1$. It is also true for $\delta=\frac{1}{2}$ and $\delta=\frac{1}{3}$. We define (using the axiom schema of specification) the set A_n to be:

$$A_n=\Big\{\,x\in X\mid\forall_{\mathcal{U}\in\mathcal{O}}\Big(B_{\frac{1}{n+1}}^{(X,\,d)}(x)\nsubseteq\mathcal{U}\Big)\,\Big\} \tag{3}$$

A_n is not empty for all $n\in\mathbb{N}$ since we are assuming the negation of the orignal claim is true. That is, we are assuming for all $\delta>0$ there is an $x\in X$ such that no $\mathcal{U}$ in the open cover $\mathcal{O}$ completely contains the δ ball centered at x. Applying this to $\delta=\frac{1}{n+1}$ we then see that A_n is a non-empty set for all $n\in\mathbb{N}$. We then (using the axiom of the power set and the axiom schema of specification) collect a new set $\mathcal{A}$ defined by:

$$\mathcal{A}=\{\,A_n\in\mathcal{P}(X)\mid n\in\mathbb{N}\,\} \tag{4}$$

To be extremely pedantic and using the axioms precisely as stated, we are defining $\mathcal{A}$ by:

$$\mathcal{A}=\Big\{\,B\in\mathcal{P}(X)\;\Big|\;\exists_{n\in\mathbb{N}}\Big(\big(x\in B\big)\Leftrightarrow\Big(\forall_{\mathcal{U}\in\mathcal{O}}\big(B_{\frac{1}{n+1}}^{(X,\,d)}(x)\nsubseteq\mathcal{U}\big)\Big)\Big)\,\Big\} \tag{5}$$

This reads, cryptically, that $\mathcal{A}$ is the set of all B such that $B = A_n$ for some $n \in \mathbb{N}$. We then we examine the product set. Since the elements of $\mathcal{A}$ are indexed by the natural numbers, we may write:

$$\prod \mathcal{A} = \prod_{n=0}^{\infty} A_n = \big\{\, f : \mathbb{N} \to \bigcup_{n=0}^{\infty} A_n \mid \forall_{n\in\mathbb{N}}\big(f(n) \in A_n\big) \,\big\} \tag{6}$$

Intuitively, an element $f \in \prod \mathcal{A}$ is a *countably infinite tuple*:

$$f = (a_0,\, a_1,\, \dots,\, a_n,\, \dots) \tag{7}$$

with the property that $a_n \in A_n$ for all $n \in \mathbb{N}$. This is for intuition, not rigor. Now we note that A_n is non-empty for all $n \in \mathbb{N}$. By the axiom of countable choice (which is implied by the full axiom of choice), there is an element $a \in \prod \mathcal{A}$. What is this element? It is a sequence $a : \mathbb{N} \to \bigcup_{n=0}^{\infty} A_n$ such that for all $n \in \mathbb{N}$ it is true that $a_n \in A_n$. Realizing that $A_n \subseteq X$ for all $n \in \mathbb{N}$, we see that $\bigcup_{n=0}^{\infty} A_n \subseteq X$, and so a is also a sequence $a : \mathbb{N} \to X$ such that $a_n \in A_n$ for all $n \in \mathbb{N}$. But what does it mean to be in A_n? We now go back to Lebesgue's number lemma, a_n in A_n means for every open set $\mathcal{U}$ in the open cover $\mathcal{O}$ it is not true that the $\frac{1}{n+1}$ ball centered about a_n is contained in $\mathcal{U}$. That is, $a : \mathbb{N} \to X$ is a sequence such that, for all $n \in \mathbb{N}$, and for all $\mathcal{U} \in \mathcal{O}$, we have:

$$B_{\frac{1}{n+1}}^{(X,\,d)}(a_n) \nsubseteq \mathcal{U} \tag{8}$$

In the proof of Lebesgue's number lemma this was entirely swept under the rug. *Obviously* I can pick a point a_n for each n satisfying the negation of our original claim. But just because something is *obvious*, does not make it true. The axiom of countable choice can not be proven using the other axioms of set theory. Defining sequences recursively requires some form of the axiom of choice as well. Since the axiom of choice is confusing and controversial it is all too common to appeal to more intuitive language that *hides* the axiom of choice, but it is important for a mathematician to realize that it is indeed there.

We have appealed to the full axiom of choice as well. For the most part we were satisfied with the axiom of countable choice, which is a far less contraversial axiom, but at times we have used the stronger version. Consider the equivalence of compactness theorem. We said consider a Cauchy sequence $a : \mathbb{N} \to X$ that does not converge in the metric space $(X,\, d)$. To converge means there is an $x \in X$ such that for all $\varepsilon > 0$ there is an $N \in \mathbb{N}$ such that $n \in \mathbb{N}$ and $n > N$ implies $d(x,\, a_n) < \varepsilon$. We denoted this by $a_n \to x$. Using quantifer notation, convergence means:

$$\exists_{x\in X}\forall_{\varepsilon>0}\exists_{N\in\mathbb{N}}\forall_{n\in\mathbb{N},n>N}\big(d(x,\, a_n) < \varepsilon\big) \tag{9}$$

The negation of this means for all $x \in X$ there is an $\varepsilon > 0$ such that for all $N \in \mathbb{N}$ there exists an $n \in \mathbb{N}$ with $n > N$ such that $d(x,\, a_n) \geq \varepsilon$. That is:

$$\forall_{x\in X}\exists_{\varepsilon>0}\forall_{N\in\mathbb{N}}\exists_{n\in\mathbb{N},n>N}\big(d(x,\, a_n) \geq \varepsilon\big) \tag{10}$$

I *conveniently* labelled such an ε as ε_x to denote that ε_x is a positive value that makes the property fail for the point $x \in X$. Why am I allowed to do this? I define the set $A_x \subseteq \mathbb{R}$ as follows:

$$A_x = \big\{\, \varepsilon \in \mathbb{R}^+ \mid \forall_{N\in\mathbb{N}} \exists_{n\in\mathbb{N}, n>N} \big(d(x,\, a_n) \geq \varepsilon\big) \,\big\} \tag{11}$$

Since we are assuming the sequence does *not* converge, for all $x \in X$, the set A_x is non-empty. We then consider the collection of all of these sets $\mathcal{A}$:

$$\mathcal{A} = \{\, A_x \in \mathcal{P}(\mathbb{R}^+) \mid x \in X \,\} \tag{12}$$

Again, being overly formal, we are writing:

$$\mathcal{A} = \Big\{\, B \in \mathcal{P}(\mathbb{R}^+) \;\Big|\; \exists_{x\in X} \Big(\big(\varepsilon \in B\big) \Leftrightarrow \forall_{N\in\mathbb{N}} \exists_{n\in\mathbb{N}, n>N} \big(d(x,\, a_n) \geq \varepsilon\big) \Big) \,\Big\} \tag{13}$$

This is the set of all $B \subseteq \mathbb{R}^+$ such that $B = A_x$ for some $x \in X$. We consider the product set. This time, unlike in the previous example, since we don't know the cardinality of X, we may not be able to index the product set over the natural numbers. We may still write:

$$\prod \mathcal{A} = \prod_{x\in X} A_x = \{\, f : X \to \bigcup_{x\in X} A_x \mid \forall_{x\in X} \big(f(x) \in A_x\big) \,\} \tag{14}$$

Since $A_x \subseteq \mathbb{R}^+$ for all $x \in X$, an element of $\prod \mathcal{A}$ is a function $f : X \to \mathbb{R}^+$. That is, a function from our metric space to the positive real numbers. This function f has the special property that for all $x \in X$, $f(x)$ is an element of A_x. That is, $f(x)$ is a positive number $\varepsilon_x = f(x)$ such that for all $N \in \mathbb{N}$ there is an $n \in \mathbb{N}$ with $n > N$ such that $d(x,\, a_n) \geq \varepsilon_x$. Since none of the elements of $\mathcal{A}$ are empty, by the axiom of choice the product is non-empty. Let $\varepsilon \in \prod \mathcal{A}$ be an element of the product. That is, let $\varepsilon : X \to \bigcup_{x\in X} A_x$ be our *choice* function. Given $x \in X$, instead of writing the image of x as $\varepsilon(x)$, let us write it as ε_x. Then ε_x is a value such that for all $N \in \mathbb{N}$ there is an $n \in \mathbb{N}$ with $n > N$ such that $d(x,\, a_n) \geq \varepsilon_x$, and now we're back to where we started in the proof of the equivalence of compactness. The axiom of choice justifies the notation ε_x where we *choose* a value $\varepsilon_x > 0$ for each x.