Notes on Ultrafilters

Alex Kruckman

Toolbox Seminar 11/7/12

1 Basic theory

Let X be a set. An (ultra)filter on X is a consistent choice of which subsets of X are "large".

Definition. A *filter* on X is $\mathcal{F} \subseteq \mathcal{P}(X)$ such that

1. $X \in \mathcal{F}$ (the whole set is large).
2. If $A \in \mathcal{F}$ and $A \subseteq B$, then $B \in \mathcal{F}$ (any set containing a large set is large).
3. If $A, B \in \mathcal{F}$, then $A \cap B \in \mathcal{F}$ (large sets have large intersection).

A filter $\mathcal{F}$ is *proper* if

4. $\emptyset \notin \mathcal{F}$ (the empty set is not large).

Examples.

- The improper filter: $\mathcal{F} = \mathcal{P}(X)$
- The trivial filter: $\mathcal{F} = \{X\}$
- The principal filter generated by $x \in X$: $\mathcal{F} = \{A \subseteq X \mid x \in A\}$
- The cofinite filter (X infinite): $\mathcal{F} = \{A \subseteq X \mid X \setminus A \text{ is finite}\}$

Note that proper filters have the *finite intersection property* (FIP):

$$\text{If } A_1, \ldots, A_n \in \mathcal{F}, \text{ then } \bigcap_i A_i \neq \emptyset.$$

Lemma (FIP)**.** *Any subset $S \subseteq \mathcal{P}(X)$ with the FIP has a minimal proper filter containing it, the filter generated by S.*

Proof. Close S downward under finite intersections, then upward under supersets. □

Definition. A proper filter $\mathcal{F}$ on X is an *ultrafilter* if

5. For any $A \subseteq X$, $A \in \mathcal{F}$ or $X \setminus A \in \mathcal{F}$ (every set is either large or co-large).

Exercise 1. Given a proper filter $\mathcal{F}$, show that $\mathcal{F}$ is an ultrafilter iff $\bigcup_{i=0}^{n} A_i \in \mathcal{F}$ implies that $A_i \in \mathcal{F}$ for some i (a large set cannot be a finite union of small sets).

Note that principal filters are ultrafilters, but the trivial filter and the cofinite filter are not. Are there any nonprincipal ultrafilters? The answer is yes, but not constructively! The following lemma is a weak form of the Axiom of Choice, i.e. it is not provable in ZF.

Lemma (Ultrafilter lemma)**.** Every proper filter $\mathcal{F}$ is contained in an ultrafilter.

Proof. Zorn's Lemma on the poset of filters on X containing $\mathcal{F}$.

We need to check the following:

- The union of a chain of filters is a filter. (Easy.)
- A maximal filter is an ultrafilter. (Let $\overline{\mathcal{F}}$ be maximal. If $\mathcal{F}$ is not an ultrafilter, take A with $A \notin \overline{\mathcal{F}}$ and $X \setminus A \notin \overline{\mathcal{F}}$. Then $\overline{\mathcal{F}} \cup \{A\}$ has the FIP, and $\overline{\mathcal{F}}$ is properly contained in the filter generated by $\overline{\mathcal{F}} \cup \{A\}$, contradicting maximality.) □

Corollary. Any subset $S \subseteq \mathcal{P}(X)$ with the FIP is contained in an ultrafilter.

Exercise 2. Every nonprincipal ultrafilter contains the cofinite filter.

2 Generalized limits

Consider a function $f : \mathbb{N} \to [0,1]$, determining a sequence $a_0, a_1, a_2, \ldots$ We say that $\lim_{n\to\infty} a_n = L$ if for every open set U containing L, all but finitely many natural numbers are mapped into U by f, i.e. $f^{-1}(U)$ is in the cofinite filter on $\mathbb{N}$.

Now given distinct points a and b in $[0,1]$, the sequence $a, b, a, b, \ldots$ does not converge. If we take disjoint neighborhoods $a \in U$, $b \in V$, $f^{-1}(U)$ and $f^{-1}(V)$ are both infinite, coinfinite subsets of $\mathbb{N}$. An ultrafilter on $\mathbb{N}$ would give preference to one of these two sets and decide whether a or b should be the limit of the sequence.

Definitions

We'll take a more general perspective by transporting (ultra)filters from the index set to the target space (from $\mathbb{N}$ to $[0,1]$ in the example above).

Exercise 3. Let $\mathcal{F}$ be an (ultra)filter on the set X, and let $f : X \to Y$ be a map of sets. Then $f_*\mathcal{F} = \{A \subseteq Y \mid f^{-1}(A) \in \mathcal{F}\}$ is an (ultra)filter on Y.

Definition. A filter $\mathcal{F}$ on a topological space Y converges to a point $y \in Y$ if for all open sets U containing y, $U \in \mathcal{F}$.

Now we can define generalized limit points of maps. With this definition, the usual notion of convergence of a sequence is $\mathcal{F}$-convergence, where $\mathcal{F}$ is the cofinite filter on $\mathbb{N}$.

Definition. Let X be a set, $\mathcal{F}$ an filter on X, Y a topological space, and $f : X \to Y$ a map of sets. Then $y \in Y$ is an $\mathcal{F}$-limit point of f if $f_*\mathcal{F}$ converges to Y.

Filter convergence works best with ultrafilters on compact Hausdorff spaces, as illustrated by the following theorem.

Theorem (Ultrafilter convergence theorem)**.** Let Y be a topological space.

1. Y is compact if and only if every ultrafilter $\mathcal{F}$ on Y converges to at least one point.

2. Y is Hausdorff if and only if every ultrafilter $\mathcal{F}$ on Y converges to at most one point.

Proof. 1. Suppose for contradiction that Y is compact, but $\mathcal{F}$ has no limit points. Then for all $y \in Y$, there is an open set U_y containing y such that $U_y \notin \mathcal{F}$. So $Y = \bigcup_y U_y$, and by compactness, $Y = \bigcup_{i=1}^n U_{y_i}$. But $Y \in \mathcal{F}$, so some $U_{y_i} \in \mathcal{F}$, contradiction.

Conversely, suppose that Y is not compact. Then there is an open cover $Y = \bigcup_i U_i$ with no finite subcover. So $\bigcap_i (Y \setminus U_i) = \emptyset$, but no finite subintersection is empty. Then $\{(Y \setminus U_i)\}_i$ has the FIP, so we can take an ultrafilter $\mathcal{F}$ containing it. Now for any point $y \in Y$, y is contained in some U_i, and $U_i \notin \mathcal{F}$, since $(Y \setminus U_i) \in \mathcal{F}$. So y is not a limit point of $\mathcal{F}$.

2. Suppose for contradiction that Y is Hausdorff, but $y \neq y'$ are both limit points of $\mathcal{F}$. Take disjoint open sets $y \in U$, $y' \in U'$. Now $U, U' \in \mathcal{F}$, but $U \cap U' = \emptyset$, contradition.

Conversely, suppose that Y is not Hausdorff. Then there are points $y \neq y'$ such that every open neighborhood of y intersects every open neighborhood of y'. This means that $\{U \,|\, y \in U \text{ open}\} \cup \{U' \,|\, y' \in U' \text{ open}\}$ has the FIP. Let $\mathcal{F}$ be an ultrafilter containing it. Then y and y' are both limit points of $\mathcal{F}$.

$\square$

So if $f : X \to Y$ is a function from a set to a compact Hausdorff space, f has a unique $\mathcal{F}$-limit for every ultrafilter $\mathcal{F}$ on X. If $\mathcal{F}$ is the principal filter generated by $x \in X$, then the $\mathcal{F}$-limit of f is $f(x)$.

Application: Stone-Cech compactification

The Stone-Cech compactification of a set X is a compact Hausdorff space βX along with a map of sets $\eta : X \to \beta X$ satisfying the following universal property: Given a compact Hausdorff space Y and a map of sets $f : X \to Y$, there is a unique continuous map $\phi : \beta X \to Y$ such that $\phi \circ \eta = f$.

$$\begin{tikzcd} X \arrow[r, "\eta"] \arrow[dr, "f"'] & \beta X \arrow[d, dotted, "\exists!\,\phi"] \\ & Y \end{tikzcd}$$

For those initiated into category theory, the Stone-Cech compactification is a functor $\beta : \text{Set} \to \text{CHaus}$ which is left-adjoint to the forgetful functor $\text{CHaus} \to \text{Set}$, and βX can be seen as the free compact Hausdorff space on X. As usual, βX is unique up to unique isomorphism.

The construction is as follows.

Underlying set: $\beta X = \{\mathcal{F} \,|\, \mathcal{F} \text{ is an ultrafilter on } X\}$.

Topology: Given $A \subseteq X$, let $U_A = \{\mathcal{F} \,|\, A \in \mathcal{F}\}$. The sets U_A are a basis (of clopen sets).

The universal map η: $x \mapsto$ the principal ultrafilter generated by $\{x\}$.

Given $f : X \to Y$, define $\phi : \beta X \to Y$ by $\mathcal{F} \mapsto$ the $\mathcal{F}$-limit of f.

Exercise 4. Check the details.

- βX is a compact Hausdorff space for all X. In fact, it is a *Stone space*: compact, Hausdorff, and totally disconnected.
- β is a functor (given a map $g : X \to X'$, there's only one reasonable choice for $\beta g : \beta X \to \beta X'$ taking ultrafilters on X to ultrafilters on X')
- The diagram commutes.
- ϕ is continuous.
- ϕ is unique.

Note that since every ultrafilter on a finite set is principal, if X is finite, $\beta X = X$ with the discrete topology.

Application: Tychenoff's theorem

The convergence theorem allows us to give a snappy proof of Tychenoff's theorem.

Theorem (Tychenoff's theorem)**.** A product of compact spaces is compact.

Proof. Let $\langle X_i \rangle_i$ be a collection of compact spaces. To show that $X = \prod_i X_i$ is compact, it suffices to show that every ultrafilter $\mathcal{F}$ on X has a limit point.

Let $\pi_i : X \to X_i$ be the i^{th} projection map. Then $(\pi_i)_*\mathcal{F}$ is an ultrafilter on X_i, and since X_i is compact, it has a limit point $x_i \in X_i$. I claim that $x = (x_i) \in X$ is a limit point of $\mathcal{F}$.

The topology on X is generated by sets of the form $V = \pi_i^{-1}(U)$, where U is an open set in X_i. So any open set containing x contains a finite intersection of sets of this form which contain x. Since $\mathcal{F}$ is closed under supersets and finite intersections, it suffices to show that if $x \in V = \pi_i^{-1}(U)$, then $V \in \mathcal{F}$.

Suppose $x \in V$. Then $x_i \in U$, so $U \in (\pi_i)_*\mathcal{F}$, since x_i is a limit point of $(\pi_i)_*\mathcal{F}$. But then by definition of $(\pi_i)_*$, $V \in \mathcal{F}$. This completes the proof. □

3 Ultraproducts

Application: Prime ideals in products of rings

Exercise 5. In a finite product of rings $\prod_{i=0}^n A_i$, all prime ideals have the following form: $\{(a_0, \ldots, a_i, \ldots, a_n) \mid a_i \in p\}$, where p is a prime ideal in A_i for some i. Hence a prime ideal in the product is determined by a choice of a prime idea in one of the factors, and $\mathrm{Spec}(\prod_{i=0}^n A_i) \cong \coprod_{i=0}^n \mathrm{Spec}(A_i)$ (you can think of this as an isomorphism of sets, topological spaces, or schemes).

The situation is not so simple for an infinite product of rings. For example, if $\{A_i \mid i \in I\}$ is an infinite collection of rings, the set $\{(a_i) \mid a_i = 0 \text{ for all but finitely many } i\}$ is an ideal in $\prod_{i \in I} A_i$. It can be extended to a maximal ideal, which cannot be contained in any prime ideal of the form described in the exercise.

This situation is analagous to one we've seen before: on a finite set, all ultrafilters are principal. But on an infinite set, there are other ultrafilters, obtained by extending the cofinite filter. One can give a complete description of the primes in an infinite product of rings using ultrafilters, but to keep things simple, we'll focus on products of fields here.

Theorem. Let $\{F_i \,|\, i \in I\}$ be a collection of fields. The prime ideals in the ring $\prod_{i\in I} F_i$ are in bijection with the ultrafilters on I. The ultrafilter $\mathcal{F}$ corresponds to the prime ideal $\{(a_i) \,|\, \text{the set of indices } i \text{ such that } a_i = 0 \text{ is in } \mathcal{F}\}$. In the same way, the proper ideals in the ring are in bijection with the proper filters on I.

Proof. Given an element $a = (a_i) \in \prod_i F_i$, let $Z_a = \{i \in I \,|\, a_i = 0\}$. Notice that $Z_{ab} = Z_a \cup Z_b$, and $Z_{a+b} \supseteq Z_a \cap Z_b$.

Let $\mathcal{F}$ be a proper filter on I. We claim that $p = \{a \,|\, Z_a \in \mathcal{F}\}$ is a proper ideal. Indeed, $0 \in p$ since $I \in \mathcal{F}$, and $1 \notin p$ since $\emptyset \notin \mathcal{F}$.

Suppose $a, b \in p$. Then $Z_a, Z_b \in \mathcal{F}$. Now $Z_{a+b} \supseteq Z_a \cap Z_b$, so $Z_{a+b} \in \mathcal{F}$, and $a + b \in p$.

Suppose $a \in p$ and c is another ring element. Then $Z_{ca} = Z_c \cup Z_a \supseteq Z_a$, so $Z_{ca} \in \mathcal{F}$, and $ca \in p$.

We have established that p is an ideal. Suppose now that $\mathcal{F}$ is an ultrafilter, and let a and b be ring elements such that $ab \in p$. Now $Z_{ab} = Z_a \cup Z_b$ is in $\mathcal{F}$. Then $Z_a \in \mathcal{F}$ or $Z_b \in \mathcal{F}$, so $a \in p$ or $b \in p$, and hence p is prime.

Conversely, let p be a proper ideal. If $Z_b = Z_a$, then a and b differ by a unit: define $c = (c_i)$ by $c_i = 1$ if $i \in Z_a$ and $c_i = b_i/a_i$ otherwise. Then $ac = b$. So $a \in p$ if and only if $b \in p$, and p is completely determined by the set $\mathcal{F}_p = \{Z_a \subseteq I \,|\, a \in p\}$, as $p = \{a \,|\, Z_a \in \mathcal{F}_p\}$. We'll check that $\mathcal{F}_p$ is a proper filter on I.

1. $I \in \mathcal{F}_p$, since $0 \in p$.

2. $\emptyset \in \mathcal{F}_p$, since $1 \notin p$.

3. Suppose $A \in \mathcal{F}_p$ and $A \subseteq B$. Then there is an element $a \in p$ such that $Z_a = A$. Define the element c by $c_i = 0$ if $i \in B$ and $c_i = 1$ otherwise. Then $Z_{ca} = A \cup B = B$, and $ca \in p$, so $B \in \mathcal{F}_p$.

4. Suppose $A, B \in \mathcal{F}_p$. Take $a, b \in p$ with $Z_a = A$ and $Z_b = B$. Then $Z_{a+b} \supseteq A \cap B$. This may be a proper containment because of unexpected cancellation in some indices i such that $b_i = -a_i \neq 0$. To remedy this problem, multiply b by c with $c_i = 0$ for the problem indices i and $c_j = 1$ for the other indices j. Then $Z_{a+cb} = A \cap B$, and $a + cb \in p$, so $A \cap B \in \mathcal{F}_p$.

5. We have established that $\mathcal{F}_p$ is a proper filter. Suppose now that p is a prime ideal, and let A be any subset of I. Let a and b be ring elements such that $Z_a = A$ and $Z_b = I \setminus A$. Then $ab = 0 \in p$, so one of a or b is in p, and hence one of A or $I \setminus A$ is in $\mathcal{F}_p$, so $\mathcal{F}_p$ is an ultrafilter.

□

Syntax and semantics

The ultraproduct construction is a very general method for putting together algebraic structures to obtain structures with specified properties. To speak in the correct level of generality, we need a quick introduction to the language of universal algebra and model theory.

Definition. A language $\mathcal{L}$ is a set of symbols. Each symbol is specified to be a constant symbol, a function symbol (of a specified finite arity), or a relation symbol (of a specified finite arity). An $\mathcal{L}$-structure is a set M, together with interpretations of the symbols in $\mathcal{L}$:

- For each constant symbol c, an element $c^M \in M$.
- For each function symbol f of arity n, a function $f^M : M^n \to M$.
- For each relation symbol R of arity m, a subset $R^M \subseteq M^m$.

Examples.

- Posets and linear orders are $\mathcal{L}_\leq$-structures, where $\mathcal{L}_\leq = \{\leq\}$, and $\leq$ is a binary relation.
- Rings and fields are $\mathcal{L}_r$-structures, where $\mathcal{L}_r = \{+, -, \cdot, 0, 1\}$, $+$ and $\cdot$ are binary functions, $-$ is a unary function, and 0 and 1 are constants.

There is an obvious notion of homomorphism of $\mathcal{L}$-structures (a map of underlying sets which preserves the interpretations of the symbols). There is also a natural way to define a product of $\mathcal{L}$-structures, which, as one might hope, gives the categorical product in the category of $\mathcal{L}$-structures.

But this situation is unsatisfying: the class of $\mathcal{L}$-structures is too broad. At this point, we have no way of specifying that an $\mathcal{L}_r$-structure is a ring, or that $\leq$ is interpreted as an order in an $\mathcal{L}_\leq$-structure.

In order to impose more interesting conditions on our structures, we need more expressive syntax. We'll use the syntax of first-order logic.

Definition. Fix a language $\mathcal{L}$. A *term* is built up from the constant symbols of $\mathcal{L}$ and a supply of variables $x_0, x_1, x_2, \ldots$ by applications of the function symbols of $\mathcal{L}$.

An *atomic formula* is of the form $t_1 = t_2$ or $R(t_1, \ldots, t_n)$, where the t_i are terms and R is a relation symbol of arity n.

A *formula* is built up from atomic formulas by Boolean connectives $\wedge$, $\vee$, and $\neg$, and quantifiers $\forall x_i$ and $\exists x_i$.

A *sentence* is a formula with no free variables. That is, every variable appearing in the formula is quantified.

Sentences in a language $\mathcal{L}$ have a familiar semantics in $\mathcal{L}$-structures, in which they express properties of these structures. In particular, given a sentence ϕ and an $\mathcal{L}$-structure M, ϕ is either true or false in M. In the same way, formulas with free variables express properties of elements. Given a formula $\phi(x)$ with free variable x and an element $a \in M$, $\phi(a)$ is either true or false in M.

For example, the following sentence asserts that the interpretation of $\leq$ is a transitive relation:

$$\forall x_1 \, \forall x_2 \, \forall x_3 \, (x_1 \leq x_2 \wedge x_2 \leq x_3) \to x_1 \leq x_3.$$

Here $\phi \to \psi$ is being used as a convenient abbreviation for $\neg\phi \vee \psi$.

The following formula in the free variable x_1 expresses that x_1 has a multiplicative inverse:

$$\exists x_2\, x_1 \cdot x_2 = 1.$$

Definition. If the sentence ϕ is true in the $\mathcal{L}$-structure M, we say that M satisfies ϕ and write $M \models \phi$.

A *theory* is a set of sentences in a language $\mathcal{L}$.

If T is a theory and the $\mathcal{L}$-structure $M \models \phi$ for all $\phi \in T$, then M is a *model* for T, written $M \models T$.

Exercise 6. Write down the field axioms as a first-order theory. Now find a theory whose models are exactly the algebraically closed fields.

Note that the definition of formulas has an inductive structure: each formula is built from strictly simpler formulas. This property allows us to conduct arguments by induction on the complexity of formulas.

Since $\phi \vee \psi$ is equivalent to $\neg(\neg\phi \wedge \neg\psi)$ and $\forall x\, \phi(x)$ is equivalent to $\neg\exists x\, \neg\phi(x)$, it suffices to consider formulas built up by $\neg$, $\wedge$, and $\exists x$.

Ultraproducts and Łoś's theorem

Now that we have the syntax of first-order logic at our disposal, we can talk about the class of models for any theory. Unfortunately, given an arbitrary theory, its class of models may lack nice categorical properties. For example, the class of fields does not have products, coproducts, initial or final objects, etc.

As a compromise, there is a nice class of theories, the algebraic theories, which give rise to well-behaved categories. The study of algebraic theories and their models is called universal algebra, and it's really a delight!

But what we're interested in here is the ultraproduct construction, which will allow us to build models for arbitrary (consistent) first-order theories.

Definition. Let $\{M_i \mid i \in I\}$ be a collection of $\mathcal{L}$-structures, and let $\mathcal{F}$ be an ultrafilter on I. The ultraproduct $M = \prod_I M_i/\mathcal{F}$ is an $\mathcal{L}$-structure defined as follows:

- The underlying set is $\prod_I M_i/\sim$, where $(a_i) \sim (b_i)$ if they agree on a large set, i.e. $\{i \in I \mid a_i = b_i\} \in \mathcal{F}$. We will denote the equivalence class of an element (a_i) by $[(a_i)]$.
- If c is a constant symbol, $c^M = [(c^{M_i})]$.
- If f is a function symbol of arity n, $f^M([(a_{i1})], \ldots, [(a_{in})]) = [(f^{M_i}(a_{i1}, \ldots, a_{in}))]$.
- If R is a relation symbol of arity n, $([(a_{i1})], \ldots, [(a_{in})]) \in R^M$ if and only if $\{i \in I \mid (a_{i1}, \ldots, a_{in}) \in R^{M_i}\} \in \mathcal{F}$.

Exercise 7. Check that this is well-defined: that $\sim$ is an equivalence relation and that the interpretations of f and R are independent of the choice of representative for each equivalence class.

Theorem (Łoś's theorem)**.** Let $\{M_i \mid i \in I\}$ be a collection of $\mathcal{L}$-structures, and let $\mathcal{F}$ be an ultrafilter on I. Let $\phi(\overline{x})$ be a first-order formula in the free variables $\overline{x}$, and let $\overline{[(a_i)]}$ be a tuple of elements from the ultraproduct $\prod_I M_i/\mathcal{F} \models \phi$. Then $\prod_I M_i/\mathcal{F} \models \phi(\overline{[(a_i)]})$ if and only if $\{i \in I \mid M_i \models \phi(\overline{a_i})\} \in \mathcal{F}$.

Exercise 8. Prove Łoś's theorem. The proof is by induction on the complexity of ϕ.

As a consequence of Łoś's theorem, if $\{M_i \mid i \in I\}$ is a collection of models of some first-order theory T, then their ultraproduct will also be a model of T. So an ultraproduct of fields is a field, and ultraproduct of linear orders is a linear order, etc.

Exercise 9. Let $\mathcal{F}$ be the principal ultrafilter on I generated by $j \in I$. Show that $\prod_I M_i/\mathcal{F} \cong M_j$.

Exercise 10. Let $\{F_i \mid i \in I\}$ be a collection of fields. Explain the connection between the ultraproduct of these fields with respect to the ultrafilter $\mathcal{F}$ and the prime ideal in the product of these fields corresponding to $\mathcal{F}$ identified in the previous section.

Application: Building funny structures

Given an $\mathcal{L}$-structure M, the theory of M, written $\mathrm{Th}(M)$, is the set of all sentences true in M. One of the most famous applications of ultraproducts is the construction of nonstandard models for $\mathrm{Th}(\mathbb{R})$.

Consider $\mathbb{R}$ as a structure in the language $\mathcal{L}_r$. Let $\{\mathbb{R} \mid i \in \mathbb{N}\}$ be a countable collection of copies of $\mathbb{R}$, and let $\mathcal{F}$ be a nonprincipal ultrafilter on $\mathbb{N}$. The ultraproduct $\mathcal{R} = \prod_i \mathbb{R}/\mathcal{F}$ is called an *ultrapower* of $\mathbb{R}$, since all the factors are the same. Note that since each factor satisfies $\mathrm{Th}(\mathbb{R})$, $\mathcal{R} \models \mathrm{Th}(\mathbb{R})$.

But $\mathcal{R}$ has "infinitesimal" elements. Consider the element $\epsilon = [(1, \frac{1}{2}, \frac{1}{3}, \frac{1}{4}, \ldots)] \in \mathcal{R}$. For any integer n, $\mathcal{R} \models \epsilon < \frac{1}{n}$, since the number of factors in which this is true is cofinite, and hence in $\mathcal{F}$.

Abraham Robinson was able develop the theory of calculus in fields like $\mathcal{R}$ by dispensing with the limit definition of derivative and interpreting the concept of infinitesimal change quite literally. His theory is called "Nonstandard Analysis". It is important to note that while an ultrapower like $\mathcal{R}$ satisfies the same first-order sentences as $\mathbb{R}$, many of the important properties of $\mathbb{R}$ are second-order in nature, that is, they quantify over *subsets* of the domain, not just elements. In order to get his hands on enough of the second order properties to develop analysis, Robinson actually worked in an expansion of $\mathbb{R}$ with extra elements representing subsets of the line, sets of subsets of the line, etc.

Exercise 11. Use the ultrapower construction to exhibit a nonstandard model for $\mathrm{Th}(\mathbb{N})$, i.e. a semiring which satisfies all the same first-order statements as $\mathbb{N}$ in the language of rings, but which has infinite elements.

Exercise 12. Use the ultraproduct construction to construct a field of characteristic 0 which has exactly one algebraic extension of each degree.

Application: The compactness theorem of first-order logic

The compactness theorem is a powerful tool for showing that first-order theories are satisfiable, i.e. that structures with certain properties exist. There are several proofs of this theorem, of which the ultraproduct proof given here is the slickest.

Theorem (Compactness theorem)**.** A theory T is satisfiable if and only if every finite subset of T is satisfiable.

Proof. If T is satisfiable, then there is a model $M \models T$. Now $M \models \Delta$ for any finite $\Delta \subseteq T$, so every finite subset of T is satisfiable.

Conversely, suppose every finite subset of T is satisfiable. Then we have a collection of structures $\{M_\Delta \mid \Delta \in I\}$ indexed by the collection of all finite subtheories of T, with $M_\Delta \models \Delta$ for all $\Delta \in I$.

The idea is to take the ultraproduct of the M_Δ with respect to some ultrafilter on I, in such a way that the ultraproduct $M = \prod_I M_\Delta / \mathcal{F}$ is a model for T. By Łoś's theorem, it suffices to find an ultrafilter $\mathcal{F}$ such that for all $\phi \in T$, $\{\Delta \mid M_\Delta \models \phi\} \in \mathcal{F}$.

Now certainly $\{\Delta \mid M_\Delta \models \phi\} \supseteq \{\Delta \mid \phi \in \Delta\}$, since $M_\Delta \models \Delta$. So we just need to pick $\mathcal{F}$ so that $A_\phi = \{\Delta \mid \phi \in \Delta\} \in \mathcal{F}$ for all $\phi \in T$. We can do this if $\{A_\phi \mid \phi \in T\}$ has the FIP. And it does: $\bigcap_{i=0}^{n} A_{\phi_i} \neq \emptyset$, since it contains the finite theory $\{\phi_0, \ldots, \phi_n\}$. $\square$

To apply the compactness theorem, just write down any theory you like. If you can show that there are no inconsistencies arising from finite pieces of the theory, the entire theory has a model. Write down what you want, and then get what you want!

Exercise 13. For each of the examples in the "Building funny structures" section, show how to use the compactness theorem to prove that a structure with the desired properties exists.

4 Two more perspectives on ultrafilters

Ultrafilters on Boolean algebras

A Boolean algebra is an algebraic structure with two binary operations $(\wedge, \vee)$, one unary operation $(\neg)$, and two distinguished elements $(1, 0)$, satisfying the rules of classical propositional logic.

Given any set X, the power set algebra $(\mathcal{P}(X), \cap, \cup, ^c, X, \emptyset)$ is a Boolean algebra. This fact leads to an observation and a question.

Observation: The definition of an ultrafilter on X does not refer to the elements of X, only to the Boolean algebra operations on the powerset algebra ($A \subseteq B$ is equivalent to $A \cup B = B$). So we can extend the definition in a natural way to consider ultrafilters on arbitrary Boolean algebras, not just on powerset algebras.

In fact, if we view the elements of a Boolean algebra B as logical propositions, an ultrafilter on B is just a consistent assignment of truth values to the propositions (the elements in the ultrafilter are true, the others are false). Equivalently, it is the preimage of 1 (true) under a Boolean homomorphism to the 2 element Boolean algebra, $B \to \{0, 1\}$.

Question: Is every Boolean algebra isomorphic to the powerset algebra of a set X?

The answer to the question is no, but the same idea that led to the construction of the Stone-Cech compactification can lead us to a positive result: every Boolean algebra can be embedded as a subalgebra of a powerset algebra.

Theorem (Stone representation theorem)**.** There is an equivalence of categories $\beta : \text{Bool} \to \text{Stone}$ between the category of Boolean algebras and the category of Stone spaces (compact, Hausdorff, totally disconnected topological spaces with continuous maps).

The functor β takes a boolean algebra B to the space of ultrafilters on B, with a basis for the topology given by (the clopen sets) $U_x = \{\mathcal{F} \mid x \in \mathcal{F}\}$ for all $x \in B$. B can be recovered from βB as the Boolean algebra of clopen sets in βB, so B is a subalgebra of $\mathcal{P}(\beta B)$.

Note that if B is the powerset algebra of the set X, then βB is the Stone-Cech compactification of X.

Measurable cardinals

Another way to view ultrafilters is as finitely additive 0-1-valued measures on the σ-algebra $\mathcal{P}(X)$ (assigning measure 1 to "large" subsets and measure 0 to "small" subsets).

Since the usual convention in measure theory is to consider countably additive measures, it is natural to ask whether we can find countably additive 0-1-valued measures. These would be ultrafilters satisfying the *countable intersection property* (CIP), instead of the FIP.

It turns out that any ultrafilter with the CIP in fact must be an ultrafilter on a set X of cardinality κ and satisfy the $(<\kappa)$IP (any intersection of fewer than κ many subsets is nonempty), where κ is a *measurable cardinal*.

Measurable cardinals are one of the mysterious large cardinals studied in set theory. They must necessarily be much much larger than any of the cardinalities encountered in everyday mathematics, so large that their existence cannot be proven in ZFC. Think of the comparison of cardinals like $\aleph_0$ and $2^{\aleph_0}$ with κ as being analagous to the comparison of cardinals like 2 and 47 with $\aleph_0$.

In fact, the situation is not just that we cannot prove measurable cardinals exist: we cannot prove that the consistency of ZFC implies the consistency of ZFC + "there exists a measurable cardinal" (so adding this axiom could make the system inconsistent), and moreover, we can prove that such a thing cannot be proven. It's entirely possible that ZFC proves that there are no measurable cardinals.

But we have no such proof, and most set theorists believe that this and other "large cardinal axioms" are not only consistent, but also philosophically motivated, and provide an extension of the ZFC set theory which is too elegant and well-structured to be ignored. But that's another story...