A Topological Viewpoint of Logic and the Compactness Theorem
University of Victoria, British Columbia
T. Schmid

Consistent Theories, Topologically

With consequences such as the Upward Löwenheim-Skolem Theorem, the Compactness Theorem is one of the most foundational results in the study of first-order logic. It is common to see only a Henkin-style syntactic proof or a lay-it-on-the-anvil ultraproduct proof of first-order compactness. However illuminating these proofs are, neither of them illustrate a topological picture of what the theorem represents. As one might guess, this picture is where the theorem gained its name.

A few clarifications before we begin: We define an $\mathcal{L}$-*theory* to be a set of first-order logic sentences written in the language $\mathcal{L}$ which is closed under logical consequence. For example, the uniqueness of the group-theoretic inverse

$$\forall x \exists y \forall z (z \cdot x = x \cdot z = 1 \rightarrow z = y)$$

is a sentence in the $\{\cdot, 1\}$-theory of groups, because it follows from the group axioms. An $\mathcal{L}$-theory T is called *consistent* if $\neg\varphi$ does not belong to T when φ belongs to T. Note that I have chosen to draw no distinction between two sentences which are logically equivalent, and that "iff" should read "if and only if."

Compactness Theorem. Fix a language of symbols $\mathcal{L}$, and let T be an $\mathcal{L}$-theory. Then T is consistent iff every finite subset T_0 of T is consistent.

Considering the compactness theorem is a statement about consistent $\mathcal{L}$-theories, it should seem sensible to topologize this collection. There turns out to be a natural way of doing this, but the fact of the matter is this: The set of all consistent $\mathcal{L}$-theories is unneccesarily complicated for our purposes! All we need to encode in our open sets are which finite subsets belong to an $\mathcal{L}$-theory. It is sufficient to focus our attention on a much better behaved collection of $\mathcal{L}$-theories.

Definition. An $\mathcal{L}$-theory T is *complete* if for any $\mathcal{L}$-sentence φ, either $\varphi \in T$ or $\neg\varphi \in T$.

In other words, complete theories decide on the truth-value of every sentence. To motivate the attention we give this special class of theories, one should verify for themself the following properties of complete theories, omitting the parenthetical content either altogether or not at all: (1) Each (consistent) $\mathcal{L}$-theory is contained in some complete (and consistent) $\mathcal{L}$-theory, (2) any intersection of (consistent) $\mathcal{L}$-theories is also a (consistent) $\mathcal{L}$-theory, and (3) an $\mathcal{L}$-theory is determined precisely by the set of complete $\mathcal{L}$-theories which contain it. In particular, (3) holds for single sentences.

I will denote the collection of complete and consistent $\mathcal{L}$-theories with S. We define a topology on S as follows: For each $\mathcal{L}$-sentence φ, let $[\varphi]$ denote the collection of $\mathcal{L}$-theories containing φ. Then

$$\mathcal{B} = \{[\varphi] \mid \varphi \text{ is an } \mathcal{L}\text{-sentence}\}$$

is a basis for a topology called the *Stone topology on* S (named after the mathematician M. H. Stone). To see that this is true, notice that $[\varphi] \cap [\psi] = [\varphi \wedge \psi]$, that $\emptyset = [\varphi \wedge \neg\varphi]$, and that $[\varphi \vee \neg\varphi] = S$. The theorem which follows is exactly the connection between topological and logical compactness.

Lemma. A subset $\mathcal{O} \subseteq \mathcal{B}$ covers S iff $\{\neg\varphi : [\varphi] \in \mathcal{O}\}$ is inconsistent.

Proof.

	$\mathcal{O} \subseteq \mathcal{B}$ covers S
iff	for every $T \in S$ there is a $[\varphi] \in \mathcal{O}$ such that $\varphi \in T$
iff	for no $T \in S$ is $\varphi \notin T$ for every $[\varphi] \in \mathcal{O}$
iff	for no $T \in S$ is $\neg\varphi \in T$ for every $[\varphi] \in \mathcal{O}$
iff	$\{\neg\varphi \mid [\phi] \in \mathcal{O}\}$ is inconsistent (see (1) above).

□

Theorem. S is compact iff the Compactness Theorem holds.

Proof. ($\Leftarrow$) Suppose that the Compactness Theorem holds, and let $\mathcal{O}$ be a collection from $\mathcal{B}$ which covers S. Every element of $\mathcal{O}$ is of the form $[\varphi]$, so we can form a set B out of the $\mathcal{L}$-sentences $\{\neg\varphi \mid [\varphi] \in \mathcal{O}\}$. By the lemma, B is an inconsistent set of sentences, so the Compactness Theorem tells us there is a finite subset B_0 of B which is inconsistent. We are left with a finite subset $\{[\neg\varphi] \mid \varphi \in B_0\}$ of $\mathcal{O}$ which covers S. $\mathcal{O}$ was chosen arbitrarily, so S must be compact.
($\Rightarrow$) Suppose that S is compact, and let T be an $\mathcal{L}$-theory. If T is inconsistent, then the lemma tells us $\mathcal{O} = \{[\neg\varphi] \mid \varphi \in T\}$ is an open cover of S. S is compact, so a finite subset $\mathcal{O}_0 \subseteq \mathcal{O}$ covers S. The collection $B_0 = \{\neg\varphi \mid [\varphi] \in \mathcal{O}_0\}$ is a finite subset of T which is inconsistent. On the other hand, if T has a finite inconsistent subset T_0, then T is inconsistent and there is nothing to show.

□

What amazes me about the illustration above is that it approaches consistency in first-order logic from a purely topological perspective.

Filters and Topology

The topology we introduced in the previous chapter occurred rather naturally for it to have given such a clear picture of the relevant geometric interpretation. Our next step toward proving the Compactness Theorem requires a bit more than just a topology, so we're also going to introduce a "measure"!

Definition. Let X be a set. A subset $\mathcal{F} \subseteq P(X)$ is called a *filter* if $X \in \mathcal{F}$, $\mathcal{F}$ is closed under finite intersections, and if $F_2 \in \mathcal{F}$ whenever $F_1 \subseteq F_2$ and $F_1 \in \mathcal{F}$. Moreover, $\mathcal{F}$ is called an *ultrafilter* if for any $F \in P(X)$ either $F \in \mathcal{F}$ or $F^c \in P(X)$.

Ultrafilters provide a way of talking about "large" subsets of a set (in fact, this is probably the motivation for their name). An example might help: Let $X = [0, 1]$, and μ Lebesgue measure. Then the family $\mathcal{F} \subseteq P(X)$ defined by

$$A \in \mathcal{F} \text{ iff } A \text{ is measurable, and } \mu(A) = 1$$

is an ultrafilter on X. Notice that the first condition is neccessary!
When seeing filters from this perspective, it becomes much easier to see their connection with the geometry of a set. Our first step toward drawing these connections out explicitly is the follwing hybrid definition incorporating the topology of a given set and the ultrafilters found on the set.

Definition. Let X be a set, $\mathcal{U}$ be an ultrafilter on X, and B be a basis for a topology on X. We say that $\mathcal{U}$ *converges to a point* $x \in X$, written $\mathcal{U} \to x$, if for any basic open subset $\beta \in B$ containing x, $\beta \in \mathcal{U}$.

Recalling the definition of sequential compactness, one might think of the following as analogous to the Hein-Borel theorem: For any subset $X \subseteq \mathbb{R}^n$, if every sequence in X contains a convergent subsequence converging to a point in X, then X is compact.

Lemma 2.3. A topological space $(X, \mathcal{T})$ is compact if every ultrafilter on X converges to at least one point.

Proof. Suppose that every ultrafilter on a topological space (X, B) converges to at least one point, and let $\mathcal{S}$ be a collection of closed sets with the finite intersection property. We let $\mathcal{U}$ be the ultrafilter on X,

$$\mathcal{U} := \{U : \text{ for some finite } \mathcal{S}_0 \subseteq S, \ \bigcap S_0 \subseteq U\}.$$

We may confirm that $\mathcal{U}$ is an ultrafilter by noticing that $\emptyset$ is (by definition) not in $\mathcal{U}$, that it is closed under finite intersections, contains its supersets, and that it decides on its containment of every subset of X.
Assume that for some $x \in X, \mathcal{U} \to x$. Every basic open set $O \in B$ is in $\mathcal{U}$, and thus so is every open set containing x. This means that every open set containing x also contains $\bigcap \mathcal{S}$, by definition of $\mathcal{U}$. Since $\bigcap \mathcal{S}$ is a closed set, if its compliment contained x, then there would be an open set containing x but not $\bigcap \mathcal{S}$. Hence $x \in \bigcap \mathcal{S}$ and $\bigcap \mathcal{S} \neq \emptyset$. $\mathcal{S}$ therefore has the finite intersection property, and X is compact. □

Filters and Logic

We have seen the connection between a more geometrical viewpoint of first-order logic and discussed a notion of largeness for sets. The next step is a construction which makes use of ultrafilters to define the generic model amongst a collection of models.
Let X be a set, $\mathcal{L}$ a language, and $\mathcal{M}_x$ (the x should just be thought of as an index) a structure interpreting the symbols from $\mathcal{L}$. We will use an ultrafilter $\mathcal{U}$ on the set X to define a new model, which I will denote by $\mathcal{M}^X/\mathcal{U}$.

Definition. Let $M^* = \Pi_{x\in X} M_x$. We define the *ultraproduct* $\mathcal{M}^X/\mathcal{U}$ to be the unique structure interpreting $\mathcal{L}$ with a universe of M^* such that if $\phi(v_1, ..., v_k)$ is an atomic formula, and $f_1, ..., f_k \in M^*$,

$$\mathcal{M}^X/\mathcal{U} \models \phi(f_1, ..., f_k) \text{ iff } \{x \in X : \mathcal{M}_x \models \phi(f_1(x), ..., f_k(x))\} \in \mathcal{U}.$$

That it is well-defined is a technical point, and the reader should confirm this for themself. More details on this can be found in [1].

Compactness of Stone Space

We continue our discussion of logic, but now with all of the previously mentioned ideas in mind. Again, we let $\mathcal{L}$ be a language, and X be the set of complete and consistent first-order $\mathcal{L}$-theories.

Lemma 4.2. If $\mathcal{U}$ is an ultrafilter on X, and we let $\mathcal{M}_t \models t$ for each $t \in X$ and $\mathcal{M} := \mathcal{M}^X/\mathcal{U}$, then $\mathcal{U} \to \mathrm{Th}(\mathcal{M})$.

Proof. In order to make sense of this, recall that $\mathrm{Th}(\mathcal{M})$ is a complete and consistent theory, and hence lies somewhere in X. Let $[\phi]$ be a basic open set containing $\mathrm{Th}(\mathcal{M})$. Then $\{t \in X : \mathcal{M}_t \models \phi\}$ belongs to $\mathcal{U}$ by definition of ultraproduct. Also, given that every $t \in X$ is complete, $\mathrm{Th}(\mathcal{M}_t) = t$. Now, since $\mathcal{M}_t \models \phi$ if and only if $\phi \in t = \mathrm{Th}(\mathcal{M}_t)$, we have

$$\begin{aligned}[\phi] &= \{t \in X : \phi \in t\} \\ &= \{t \in X : \mathcal{M}_t \models \phi\} \\ &\in \mathcal{U}.\end{aligned} \tag{1}$$

We chose $[\phi]$ arbitrarily, so $\mathcal{U} \to \mathrm{Th}(\mathcal{M})$. ☐

Theorem 4.3]. X is compact.

Proof. Every ultrafilter on X converges to at least one point, by the last lemma. Thus, the result follows from lemma 2.3. ☐
The Compactness of first-order logic follows!

Concluding Remarks and References

What I became interested in before I began this project was the sort of shape given to the syntax of first-order logic by the semantical properties of deduction. It truly is fascinating to consider the homeomorphism class of complete and consistent sets of first order sentences under the Stone topology – an approach to wrangling such a large and structured set I now consider quite natural. If you also feel this way and are interested in more detail than my overview, I encourage you to explore the following documents, or at least any others you happen to find. In fact, if you do find one that is not on this list, email me at *twsch@uvic.ca*.

[1] *Ultrafilters, ultraproducts, and the Compactness Theorem*, A. Caicedo, 2009. Link:
`https://caicedoteaching.files.wordpress.com/2009/10/502-compactness.pdf`
[2] *Filters in Analysis and Topology*, D. MacIver, 2004. Link:
`http://www.efnet-math.org/~david/mathematics/filters.pdf`.
[3] *Model Theory: An Introduction*, D. Marker, 2002.