Notes on Compactness

Phil 140A | Spring 2016

§1 Proof of Compactness

Recall that we've proven the following theorem earlier:

Theorem 1 *(Soundness and Completeness).* $\Gamma \models \sigma$ iff $\Gamma \vdash \sigma$.

This is the key theorem needed for proving compactness. First, some definitions.

Definition 2 *(Satisfiability).* Let Γ be a set of sentences.

- Γ is ***satisfiable*** if there's a model $\mathfrak{A}$ such that $\mathfrak{A} \models \Gamma$. In other words, Γ is satisfiable if $\Gamma \not\models \bot$. If Γ is not satisfiable, we'll say it's ***unsatisfiable***.
- Γ is ***finitely satisfiable*** if for all finite subsets $\Delta \subseteq \Gamma$, Δ is satisfiable.

Note the following equivalences between satisfiability and implication:

$$\begin{array}{lcl} \Gamma \models \bot & \Leftrightarrow & \Gamma \text{ is unsatisfiable} \\ \Gamma \not\models \bot & \Leftrightarrow & \Gamma \text{ is satisfiable} \\ \Gamma \models \varphi & \Leftrightarrow & \Gamma \cup \{\neg\varphi\} \text{ is unsatisfiable} \\ \Gamma \not\models \varphi & \Leftrightarrow & \Gamma \cup \{\neg\varphi\} \text{ is satisfiable.} \end{array}$$

Lemma 3 *(Easy Finite Satisfiability).* If Γ is satisfiable, then it is finitely satisfiable.

Proof: If $\mathfrak{A} \models \Gamma$, and if $\Delta \subseteq \Gamma$ where Δ is finite, then $\mathfrak{A} \models \Delta$. ■

Theorem 4 *(Compactness v1).* If Γ is finitely satisfiable, then it is satisfiable.

Proof: By contraposition. Suppose Γ is unsatisfiable. That means $\Gamma \models \bot$. By completeness, this implies that $\Gamma \vdash \bot$. But since derivations are all finite, there can only be finitely many members of Γ that are used in any given derivation. So that means there must be a finite $\Delta \subseteq \Gamma$ such that $\Delta \vdash \bot$. But then by soundness, $\Delta \models \bot$, i.e., Δ is unsatisfiable. So there's a finite $\Delta \subseteq \Gamma$ that's unsatisfiable, and hence Γ is not finitely satisfiable. ■

§2 Consequences of Compactness

The following is an *equivalent* statement of compactness.

Corollary 5 *(Compactness v2)*. If $\Gamma \models \varphi$, then for some finite $\Delta \subseteq \Gamma$, $\Delta \models \varphi$.

You're asked in the exercise set to prove that Compactness v1 implies Compactness v2. Note, then, that compactness is just encoding the fact that all semantic implications are "finite" in the same way that all derivations are finite (i.e., only involve a finite number of premises).

The first major consequence of compactness is that any set of sentences with arbitrarily large models has an infinite model.

Definition 6 *(Arbitrarily Large Finite)*. Let Γ be a set of sentences. We'll say Γ has ***arbitrarily large finite models*** if for all $n \geqslant 1$, there is a model $\mathfrak{A}_n$ where $|\mathfrak{A}_n| \geqslant n$ such that $\mathfrak{A}_n \models \Gamma$. In other words, Γ has arbitrarily large finite models if there's no upper bound on how big the finite models of Γ can get.

Corollary 7 *(Arbitrarily Large Finite Implies Infinite)*. If Γ has arbitrarily large finite models, then it has an infinite model.

Proof: Suppose Γ has arbitrarily large finite models. Consider the following set:

$$\Gamma^* := \Gamma \cup \{\lambda_n \mid n > 1\}.$$

Claim: Γ^* is finitely satisfiable.

Subproof: Let $\Delta^* \subseteq \Gamma^*$ be finite. That means for some $\gamma_1, \ldots, \gamma_n \in \Gamma$ and some $k_1, \ldots, k_m$, $\Delta^* = \{\gamma_1, \ldots, \gamma_n, \lambda_{k_1}, \ldots, \lambda_{k_m}\}$. Let $k = \max(k_1, \ldots, k_m)$. Since Γ has arbitrarily large finite models, there's a model $\mathfrak{A}_k \models \Gamma$. But then $\mathfrak{A}_k \models \gamma_1 \wedge \cdots \wedge \gamma_n$ and $\mathfrak{A}_k \models \lambda_{k_1} \wedge \cdots \wedge \lambda_{k_m}$. So $\mathfrak{A}_k \models \Delta^*$, and thus Δ^* is satisfiable. Hence, Γ^* is finitely satisfiable. ∎

So by compactness, it follows that Γ^* is finitely satisfiable. But as you showed in your exercise set, any model of $\{\lambda_n \mid n > 1\}$ must be infinite. So there's an infinite model of Γ^*, and hence of Γ. ∎

The second major consequence of compactness is that there's no sentence in first-order logic that is true in all and only finite models. That is, having a finite number of elements if not "first-order definable".

Theorem 8 *(Undefinability of Finitude)*. There's no sentence φ such that for all models $\mathfrak{A}$, $\mathfrak{A} \models \varphi$ iff $|\mathfrak{A}|$ is finite.

Proof: Here's a quick proof. If φ is true on all and only finite models, then $\{\varphi\}$ has arbitrarily large finite models (indeed, every finite model). So that means $\{\varphi\}$ has an infinite model, contradiction. So $\{\varphi\}$ cannot be true on all and only finite models.

Here's a more direct proof. Suppose for *reductio* there were such a sentence φ. Consider the following set of sentences:

$$\Gamma := \{\varphi\} \cup \{\lambda_n \mid n > 1\}.$$

Claim: Γ is finitely satisfiable.

Subproof: Let $\Delta \subseteq \Gamma$ be finite. Then for some $\lambda_{k_1}, \ldots, \lambda_{k_n}$, $\Delta \subseteq \{\lambda_{k_1}, \ldots, \lambda_{k_n}\} \cup \{\varphi\}$. Let $k = \max(k_1, \ldots, k_n)$. Consider any model $\mathfrak{A}_k$ where $|\mathfrak{A}_k| = k$. Then $\mathfrak{A}_k \models \lambda_{k_1} \wedge \cdots \wedge \lambda_{k_n}$ since there are at least k-many elements in $\mathfrak{A}_k$ (in fact, there are exactly k-many). And since $k \in \mathbb{N}$, $\mathfrak{A}_k$ is finite. So by assumption, $\mathfrak{A}_k \models \varphi$. Hence, $\mathfrak{A}_k \models \Delta$, i.e., Δ is satisfiable. So Γ is finitely satisfiable. ■

By compactness, Γ is satisfiable. So there's a model $\mathfrak{A} \models \Gamma$. But then $\mathfrak{A}$ is infinite, even though $\mathfrak{A} \models \varphi$. This contradicts our initial supposition. ■

§3 Löwenheim-Skolem

The Löwenheim-Skolem theorems state that if you have an infinite model of Γ, then you can get a model of "any infinite size", as long as it's at least as big as the language.

Theorem 9 *(Downward Löwenheim-Skolem, Simple)*. Let $L(\Gamma)$ be the language of Γ, and $|L(\Gamma)|$ be the cardinality of the set of formulas in $L(\Gamma)$. Suppose there is an infinite model $\mathfrak{A} \models \Gamma$. Then there is a model $\mathfrak{B} \models \Gamma$ such that $|\mathfrak{B}| = |L(\Gamma)|$.

Proof: This follows from Lemma 4.1.11. Consider the set:

$$\Gamma^* := \Gamma \cup \{c_i \neq c_j \mid i \neq j \text{ and } i, j < |L(\Gamma)|\}$$

where c_i for $i < |L(\Gamma)|$ is a new constant not in $L(\Gamma)$. It's easy to check that Γ is satisfiable (just take the appropriate expansion of $\mathfrak{A}$) and hence consistent. Thus, we can use Lemma 4.1.11 to construct a model of Γ using the terms of the language. And just note that this canonical model will be of size $|L(\Gamma)|$. ■

Theorem 10 *(Upward Löwenheim-Skolem, Simple)*. Suppose $\mathfrak{A} \models \Gamma$, where $|\mathfrak{A}| = |L(\Gamma)|$. Then for any $\kappa > |L(\Gamma)|$, there is a model $\mathfrak{B} \models \Gamma$ such that $|\mathfrak{B}| = \kappa$.

Proof: This follows from compactness and the version of Downward Löwenheim-Skolem given above. Suppose $\mathfrak{A} \models \Gamma$ where $|\mathfrak{A}| = |L(\Gamma)|$. For each $i < \kappa$, let c_i be a new constant not in $L(\Gamma)$. Consider the set:

$$\Gamma^* := \Gamma \cup \{c_i \neq c_j \mid i \neq j \text{ and } i, j < \kappa\}$$

Γ^* is finitely satisfiable (note that since $\mathfrak{A}$ is infinite, $\mathfrak{A}$ can be your model for each finite subset of Γ^*; but you should check this yourself). Hence, by compactness, Γ^* is satisfiable. So there's a model $\mathfrak{B} \models \Gamma^*$.

Now, $\mathfrak{B}$ must be of at least size κ, since there are κ-many constants that denote different objects in $\mathfrak{B}$. But it might not be exactly of size κ (it might be bigger). But now notice that $|L(\Gamma^*)| = \kappa$, since there are now κ-many constants in $L(\Gamma^*)$. So by Downward Löwenheim-Skolem, there's a model $\mathfrak{B}' \models \Gamma^*$ such that $|\mathfrak{B}'| = |L(\Gamma^*)| = \kappa$. Hence, $\mathfrak{B}'$ is a model of Γ that's of exactly size κ. ■

Theorem 11 *(Downward Löwenheim-Skolem, General)*. Suppose $\mathfrak{A} \models \Gamma$, where $|\mathfrak{A}| > \kappa \geqslant |L(\Gamma)|$. Then there's a model $\mathfrak{B} \models \Gamma$ such that $|\mathfrak{B}| = \kappa$.

Proof: If $\mathfrak{A} \models \Gamma$ where $|\mathfrak{A}| > \kappa \geqslant |L(\Gamma)|$, then by the simple version of Downward Löwenheim-Skolem (since $|\mathfrak{A}| > |L(\Gamma)|$), there's a model $\mathfrak{B} \models \Gamma$ such that $|\mathfrak{B}| = |L(\Gamma)|$. But then by the simple version of Upward Löwenheim-Skolem, there's a model $\mathfrak{B}' \models \Gamma$ such that $|\mathfrak{B}'| = \kappa$. ■

Theorem 12 *(Upward Löwenheim-Skolem, General)*. Suppose $\mathfrak{A} \models \Gamma$ where $|\mathfrak{A}| \geqslant |L(\Gamma)|$. Then for any $\kappa > |\mathfrak{A}|$, there is a model $\mathfrak{B} \models \Gamma$ such that $|\mathfrak{B}| = \kappa$.

Proof: Again, use the simple Downward Löwenheim-Skolem to get a model $\mathfrak{B}' \models \Gamma$ where $|\mathfrak{B}'| = |L(\Gamma)|$. Then use the simple Upward Löwenheim-Skolem to get $\mathfrak{B} \models \Gamma$ so that $|\mathfrak{B}| = \kappa$. ■

In the simple case where $L(\Gamma)$ is countable (which is usually the case), we get the following result: if Γ has an infinite model, then it has a model of any infinite size.

An interesting theorem that we won't prove is the following (stated informally):

Theorem 13 *(Lindström's Theorem)*. First-order logic is the strongest logic satisfying both (countable) compactness and Downward Löwenheim-Skolem.