6c Lecture 14: May 14, 2014

11 Compactness

We begin with a consequence of the completeness theorem. Suppose T is a theory. Recall that T is *satisfiable* if there is a model $M \models T$ of T. Recall that T is *consistent* if there is no sentence ϕ such that $T \vdash \phi$ and $T \vdash \neg\phi$.

Theorem 11.1. *If T is a theory, then T is satisfiable iff T is consistent.*

Proof. ($\rightarrow$): we prove the contrapositive. If T is inconsistent, there is some ϕ such that $T \vdash \phi$ and $T \vdash \neg\phi$. Then by the soundness theorem, $T \models \phi$ and $T \models \neg\phi$. Hence, if there was a model M of T, then $M \models \phi$ and $M \models \neg\phi$ which is a contradiction.

($\leftarrow$): We showed in our proof of the completeness theorem that if T is consistent, then there is a model M of T. □

Now on your homework, you proved one version of the compactness theorem:

Theorem 11.2 (Compactness theorem, version 1)**.** *Suppose L is a language, T is a theory in, and ϕ is a sentence in L. Then if $T \models \phi$, there is a finite subset $T_0 \subseteq T$ such that $T_0 \models \phi$.*

We'll now give another version of the compactness theorem:

Theorem 11.3 (Compactness theorem, version 2)**.** *If T is a theory, then T is satisfiable iff every finite subset $T_0 \subseteq T$ is satisfiable.*

Proof. ($\rightarrow$): If T has a model M, then clearly this model M is a model of every finite subset of T.

($\leftarrow$): We prove the contrapositive. Suppose that T is unsatisfiable. Then by Theorem 11.1, T is inconsistent, and hence there is some ϕ such that $T \vdash \phi$ and $T \vdash \neg\phi$. But since proofs are finite, there is some finite $T_0 \subseteq T$ such that $T_0 \vdash \phi$ and $T_0 \vdash \neg\phi$. Hence, some finite subset $T_0 \subseteq T$ is inconsistent, and hence unsatisfiable. □

For the rest of the lecture today, we'll discuss consequences of the compactness theorem.

We'll begin by discussing to what extent finiteness can be axiomatized by first-order logic.

Theorem 11.4. *Suppose T is a theory that has finite models of arbitrarily large size. That is, for ever number $n \in \mathbb{N}$, there is a finite model $M_n \models T$ such that the universe of M has at least n elements. Then T has an infinite model.*

Proof. Let ϕ_n be the sentence which states that there are at least n different elements in the universe:

$$\phi_n = \forall x_1 \forall x_2 \ldots \forall x_{n-1} \exists x_n (x_1 \neq x_n \wedge x_2 \neq x_n \wedge \ldots \wedge x_{n-1} \neq x_n).$$

Now consider the theory $T' = T \cup \{\phi_1, \phi_2, \ldots\}$

Given any finite subset $T_0 \subseteq T'$, there is some largest n such that $\phi_n \in T'$. But then we claim that M_n is a model of T_0 since $M_n \vDash T$, and $M_n \vDash \{\phi_1, \ldots, \phi_n\}$. Hence, every finite subset of T' is satisfiable, so T' is satisfiable by the compactness theorem. Hence, there is a model M of T' which will be an infinite model of T. □

A corollary of this theorem is that there is no way in first-order logic to axiomatize finiteness.

Corollary 11.5. *There is no theory T such that every finite structure is a model of T, and no infinite structure is a model of T.*

Proof. Because every finite model would be a model of T, there must be finite models of T of arbitrarily large size. Hence T must have an infinite model by the above theorem. □

Often, concepts involving infinity cannot be expressed in first-order logic. We can often prove this using the compactness theorem. Our next theorem is another result of this type:

Theorem 11.6. *Consider the theory of graphs:*

$$T = \{\forall x(\neg xEx), \forall x \forall y(xEy \to yEx)\}$$

in the language L containing a single binary relation E. Then there is no theory $T' \supseteq T$ such that the models of T' are exactly the connected graphs.

Proof. Lets begin by considering the language $L' \supseteq L$ obtained by adding two constant symbols a and b to the language L. Now for each n consider the sentence:

$$\phi_n = \neg \exists x_1 \exists x_2 \ldots \exists x_n$$
$$(a = x_1 \wedge b = x_n \wedge x_1 E x_2 \wedge \ldots \wedge x_{n-1} E x_n \wedge x_1 \neq x_2 \wedge \ldots x_{n-1} \neq x_n)$$

which says there is no path of length $n-1$ between a and b. Now suppose T' was a theory whose models were precisely the connected graphs. Then we claim that $T' \cup \{\phi_1, \phi_2, \ldots\}$ would be satisfiable.

This is by the compactness theorem. Given any finite subset $T_0 \subseteq T'$, there is some largest n such that $\phi_n \in T'$. But then take a connected graph having two points of distance $> n$ and assign these two points to be a and b. Then the resulting graph is a model of T_0.

Hence, since T' is satisfiable, it has a model M. But then M is not connected, since there is no path of any length between a and b. So regarding M as just a model in the language L by forgetting the constants a and b, $M \vDash T$ which is a contradiction. □

The trick we have used above of adding some constants to the language before we do a compactness argument (and then forgetting them later to derive consequences about the theory we started with) is a very technique.

11.1 Nonstandard analysis and infinitesimals.

When doing calculus, differential equations, mathematical physics, etc. we often pretend that we have infinitesimally small real numbers (we call them things like dx, δ, or Δy, etc).

If you've taken a course in mathematical analysis, you've probably seen how we can often formalize these types of intuitions using limits of increasingly small numbers.

However, there is a different way of working with infinitesimals which realizes them as actual objects (and not just a vague intuitions for limits).

Lets start with the model whose universe is the real numbers $\mathbb{R}$ and where our language L contains:

- The relation $<$
- A constant c_r representing each real number $r \in \mathbb{R}$
- The functions of addition and multiplication $+, \cdot$
- (and optionally) whatever other functions $f_1, f_2, \ldots$ we want.

Lets call this model $\mathcal{R}$, and let T be the theory consisting of all sentences ϕ such that $\mathcal{R} \vDash \phi$. (That is, T is the theory of the real numbers).

Now lets enlarge our language L by adding a constant symbol δ. (We'll think of δ as an infinitesimally small number). For each n, let ϕ_n be the sentence:

$$\phi_n = 0 < \delta < 1/n$$

Now consider the theory $T' = T \cup \{\phi_1, \phi_2, \ldots\}$. We claim that T' is satisfiable.

This is by the compactness theorem. If $T_0 \subseteq T'$ is a finite subset of T', then there is a largest n such that $\phi_n \in T_0$. Now we claim that T_0 is satisfiable, since $\mathcal{R}$ is a model of T_0 if we assign the value $1/n + 1$ to the constant δ.

This means that there is a model ${}^*\mathcal{R}$ of T'. This model looks a whole lot like the real numbers: it still has a constant c_r representing every real number r, and ${}^*\mathcal{R} \vDash T$, the theory of the real numbers. However, this model ${}^*\mathcal{R}$ also has has infinitesimally small real numbers (such as the one represented by δ), which are > 0, but $< 1/n$ for every n. ${}^*\mathcal{R}$ is called a *nonstandard* model of the real numbers, and the field of study which uses such models to prove theorems is called *nonstandard analysis*.

Now given any standard real number $r \in \mathbb{R}$, we can identify r with the number represented by the constant c_r in ${}^*\mathcal{R}$. Literally, this defines an embedding of the model $\mathcal{R}$ into ${}^*\mathcal{R}$. Often we will abuse notation, and not differentiate between standard numbers in $\mathcal{R}$ and their images in ${}^*\mathcal{R}$.

Given any element x in the universe of ${}^*\mathcal{R}$, say that x is *standard* if there is some constant c_r representing a standard real number such that $x = c_r$. Otherwise say x is *nonstandard*. Say that x is infinitesimal if for every $n \in \mathbb{N}$ we have $-1/n < x < 1/n$. Note that 0 is infinitesimal, but there are nonstandard infinitesimal numbers (such as the number represented by δ) in our model ${}^*\mathcal{R}$. We will write $x \approx y$ if $x - y$ is infinitesimal. If there is a standard real number $y \in \mathbb{R}$ such that $x \approx y$, then we say that y is the standard part of x.

Say that x is *finite* if there is an $n \in \mathbb{N}$ such that $-n < x < n$. Say that x is *infinite* otherwise.

We note some examples of facts about nonstandard numbers. To prove these, we heavily use the fact that the theory of ${}^*\mathcal{R}$ is the same as the theory of $\mathcal{R}$:

- There are infinite numbers in ${}^*\mathcal{R}$. To see this, take an infinitesimal $\delta > 0$. Now since $\delta \neq 0$, and $\mathcal{R} \models \forall x(x \neq 0 \rightarrow \exists y(x \cdot y = 1))$, (and hence ${}^*\mathcal{R}$ models this as well) there must be some some number y such that $\delta \cdot y = 1$. Lets use the notation δ^{-1} for this y. Now since $\delta < 1/n$ for every n, we have that $\delta^{-1} > n$ for every n, and hence δ^{-1} is infinitely large.

- If x is a finite number in ${}^*\mathcal{R}$, then there is a standard real number y such that $y \approx x$. To see this[1], note that

  $$\mathcal{R} \models \forall x \forall a \forall b (a < x < b \rightarrow \\ (x = \frac{a+b}{2}) \vee a < x < \frac{a+b}{2} \vee \frac{a+b}{2} < y < b).$$

  Hence, if we start with natural numbers n and $-n$ such that $-n < x < n$, using the above fact, we can get a sequence of standard numbers $a_1, a_2, \ldots$ with $a_i < x$ and $b_1, b_2, \ldots$ with $b_i > x$ such that $\lim_i a_i = \lim_i b_i$. Letting this limit be y, we see that $y - x$ must be infinitesimal.

A common (but imperfect) way to visualize a nonstandard model ${}^*\mathcal{R}$ is to imagine each real number, and then add some some extra infinitesimal "stuff" around each real number, and then some infinite "stuff" at both ends of the number line beyond all the standard numbers. Note though that there's a lot of "layers" involved when doing this. For example, if δ is a nonstandard infinitesimal, and r is a standard real number, then $r\delta$ is a nonstandard infinitesimal (so around each standard number is an infinitesimal copy of the standard reals $\mathbb{R}$). Note though, that there are lost more infinitesimals that just numbers like $r\delta$. For example, δ^2 is also a nonstandard infinitesimal, but it is infinitesimally smaller than δ.

Using nonstandard models, its possible to formally and precisely define lots of concepts from analysis using infinitesimal numbers instead of limits.

For example, a function f is continuous iff for all $x \approx y$, we have $f(x) \approx f(y)$. A fact which is easy to prove from this definition is that if f is continuous on $[0, 1]$ and $0 < x < 1$, then $f(x)$ is finite.

[1] here we are using $y/2$ as an abbreviation for the unique number z such that $2 \cdot z = y$, which must always exist.

Similarly, we can define the derivative of $f(x)$ to be equal to as the standard part of $\frac{f(x+\delta)-f(x)}{\delta}$ for any nonzero infinitesimal δ (provided the standard part is the same for every such nonzero infinitesimal). So for a given infinitesimal δ, if f is differentiable, then the quantity $f(x+\delta)-f(x)$ will be an infinitesimal number, and the fraction $\frac{f(x+\delta)-f(x)}{\delta}$ will actually be a finite number.

For example, if $f(x)=x^2$, then $\frac{(x+\delta)^2-x^2}{\delta}=\frac{2x\delta+\delta^2}{\delta}=2x+\delta$, but then since δ is infinitesimal, the standard part of this is just $2x$ (which is the derivative of x^2).

We have just barely scratched the surface of nonstandard analysis which is a large and important field.