REU'09·Transfinite combinatorics·Lecture 7

Instructor: László Babai
Scribe: Thomas Church

July 27, 2009

7.1 Model theory - continued

Definition 7.1. A set of sentences Σ is *satisfiable* if it has a model:

$$\exists \mathcal{A} \colon \mathcal{A} \models \Sigma$$

Definition 7.2. A set of sentences Σ is *consistent* if no contradiction can be deduced from Σ by a sequence of applications of "rules of formal deduction". ("Rules of formal deduction" are things like: if we have $\varphi \to \psi$ and φ, then we can deduce ψ.) "Deducing a contradiction" means deducing both φ and $\neg\varphi$ for some sentence φ.

Syntactic inference: $\Sigma \vdash \varphi$ if φ follows syntatically (by the "rules of formal deduction") from Σ.
Semantic inference: $\Sigma \models \varphi$ if any model for Σ models φ:

$$\Sigma \models \varphi \qquad \iff \qquad \forall \mathcal{A}\colon \quad \mathcal{A} \models \Sigma \implies \mathcal{A} \models \varphi$$

Theorem 7.3 (Gödel's Completeness Theorem)**.**

$$\big(\Sigma \models \varphi\big) \iff \big(\Sigma \vdash \varphi\big)$$

"Whatever is true is provable"; true for first-order logic, propositional calculus.

Exercise 7.4. The completeness theorem can be equivalently stated:

$$\Sigma \text{ is satisfiable} \iff \Sigma \text{ is consistent.}$$

The compactness theorem follows from completeness:

Theorem 7.5 (Compactness theorem, reworded)**.** If every finite subset of Σ is satisfiable, then Σ is satisfiable.

Proof, using completeness theorem. By the completeness theorem, it suffices to prove: "If every finite subset of Σ is consistent, then Σ is consistent." Assume that Σ is not consistent. Then a contradiction can be deduced from Σ. This deduction has finite length, and thus involves only finitely many sentences from Σ. Thus that finite subset of Σ is itself inconsistent. $\square$

Definition 7.6. Recall that $\mathcal{A}$ and $\mathcal{B}$ are *elementary equivalent* if

$$\forall \text{ sentences } \varphi: \qquad \mathcal{A} \models \varphi \iff \mathcal{B} \models \varphi$$

Theorem 7.7 (Upward Löwenheim–Skolem)**.** If Σ has an infinite model, then it has models of arbitrarily large cardinality. (Formally, if $|\mathcal{A}| \geq \aleph_0$, then for any cardinal m, $\exists \mathcal{B}$ so that $\mathcal{B}$ is elementary equivalent to $\mathcal{A}$ and $|\mathcal{B}| \geq m$.)

Proof. New constants! (Constants are nullary operations.) Choose I with $|I| \geq m$. Take $\mathcal{A}$, and add constants f_α for each $\alpha \in I$. Let Σ be the theory of $\mathcal{A}$ (all φ so that $\mathcal{A} \models \varphi$ in the original language), and let Σ' be obtained by adding the (essentially trivial) sentences $f_\alpha \neq f_\beta$ for each $\alpha \neq \beta \in I$. These add no expressive power to the language, they just force any model to have cardinality $\geq m$, since it must have $|I|$ different constants.

We use the compactness theorem to show that Σ' is satisfiable; as above, any model will have cardinality $\geq |I| \geq m$. We need to show that any finite subset of Σ' is satisfiable. Any finite subset consists of some subset of Σ (which we know is satisfiable, because $\mathcal{A}$ models it) plus some finite number of the sentences $f_\alpha \neq f_\beta$. Since $\mathcal{A}$ is infinite, we may find an interpretation by choosing distinct elements for each f_α that appears in our *finite* number of inequalities, and interpreting the remaining constants by any element of $\mathcal{A}$. Then each of these finitely many inequalities will be satisfied. ☐

Exercise 7.8. If Σ has aribtrarily large *finite* models, then it has arbitrarily large infinite models. (Hint 1: use ultraproducts to reduce to the previous problem. Hint 2: Don't use ultraproducts, just Compactness as in the previous solution.)

Theorem 7.9. If Σ has an infinite model and the language is at most countable, then Σ has a countable model; moreover, Σ has models of all infinite cardinalities.

The proof follows from the previous result (which gave us models of arbitrarily large cardinalities) and the following result (which allows us to go down with the cardinality). First a definition:

Definition 7.10. $\mathcal{B}$ is an *elementary submodel* of $\mathcal{A}$ if for every formula $\varphi(\mathbf{x}, \mathbf{y})$ (where $\mathbf{x} = (x_1, \ldots, x_k)$ and $\mathbf{y} = (y_1, \ldots, y_\ell)$ are the only free variables in φ) and for every $\mathbf{b} \in \mathcal{B}$, if

$$\mathcal{A} \models (\exists \mathbf{y})\big(\varphi(\mathbf{b}, \mathbf{y})\big) \tag{1}$$

then

$$\mathcal{B} \models (\exists \mathbf{y})\big(\varphi(\mathbf{b}, \mathbf{y})\big)$$

Note that if $\mathcal{B}$ is an elementary submodel of $\mathcal{A}$ then $\mathcal{A}$ and $\mathcal{B}$ are elementary equivalent, but not conversely, as the next exercise shows.

Exercise 7.11. Find a model $\mathcal{A}$ and a submodel $\mathcal{B}$ such that $\mathcal{B}$ and $\mathcal{A}$ are elementary equivalent but $\mathcal{B}$ is not an elementary submodel of $\mathcal{A}$. (Hint: ordered sets.)

Theorem 7.12 (Downward Löwenheim–Skolem)**.** If the language is at most countable, and $\mathcal{A}$ is infinite, then $\mathcal{A}$ has a countable elementary submodel. Moreover, for any infinite cardinality $\mathfrak{m} \leq |\mathcal{A}|$, $\mathcal{A}$ has an elementary submodel of cardinality $\mathfrak{m}$.

Proof. The construction of the elementary substructure which we use is something like the "substructure generated by a subset", analogous to the subgroup generated by a subset, or the subfield generated by a subset. The operations under which we need to close the subset are the "Skolem functions" which for each existentially quantified formula (1) select a solution. If we start with a countable subset, the elementary substructure it "generates" will still be countable. More generally, if we start from a subset of any infinite cardinality $\mathfrak{m}$, the closure will also have cardinality $\mathfrak{m}$. □

In particular, set theory has a countable model. Such a model is a countable set, with a single relation "$a \in b$". We can talk about subsets by saying that a' is a subset of a if for all b such that $b \in a'$, $b \in a$. This leads to apparent contradiction, since the power set of $\mathbb{N}$ is uncountable, while our whole model itself is countable. This is resolved when we remember that "countable" means "has a bijection with $\mathbb{N}$"; in this model, we will have sets which externally can be seen to be countable, but internally are not countable, because there is no *internal* bijection with the element representing $\mathbb{N}$.

7.2 The Random Graph

Exercise 7.13 (Exercise from before)**.** Two countable random graphs are isomorphic with probability 1.

Proof. Build an isomorphism incrementally. We can send $1 \mapsto A$. Let's say that 2 is adjacent to 1 in the first graph; then we need to find some vertex in the second graph to which we may send 2. Any vertex which is adjacent to A in the second graph will suffice; the probability that no vertex is adjacent to A is 0, so with probability 1 we may find such a vertex. Perhaps $2 \mapsto E$. Now, we need to find some vertex in the first graph which we may send to B in the second graph. Say that B is adjacent to E, but not to E. Then we need to find some vertex in the first graph which is adjacent to 2 but not to 1. The probability that no such vertex exists is 0, so with probability 1 we may find such a vertex. Now alternating back and forth, we build a bijection elementwise; since a probability measure is countably additive, the probability of trouble at any step is 0, so with probability 1 our algorithm constructs a bijection.

(The reason that we must alternate between $\{1, 2, \ldots\}$ and $\{A, B, \ldots\}$ is that we want a *bijection*. If we just worked with $\{1, 2, \ldots\}$ we would construct an injection, but could not guarantee that it would be surjective.) □

Let's capture the first-order property of a random graph that makes this proof work:

$$\varphi_{k,\ell}\colon \quad (\forall x_1,\dots,x_k,y_1,\dots,y_\ell)(\exists v)\Big(\bigwedge_{i,j} x_i \neq y_j \to \\ v\sim x_1 \wedge \dots \wedge v \sim x_k \wedge v \not\sim y_1 \wedge \dots \wedge v \not\sim y_\ell\Big)$$

Claim. The theory $R = \{\varphi_{k,\ell}\colon k,\ell \in \omega\}$ is categorical in $\aleph_0$.

(Recall that a theory is categorical in a cardinality $\mathfrak{m}$ if it has only one model of that cardinality.) We have essentially just proved this claim: if two countable graphs both satisfy R then with the back-and-forth method described, we find an isomorphism between them. We also proved that the probability $\Pr\big(\text{countable graph } \models R\big) = 1$.

Definition 7.14. A set of sentences Σ is *complete* if for every sentence φ, either

$$\Sigma \models \varphi \qquad\qquad \text{or} \qquad\qquad \Sigma \models \neg\varphi$$

Exercise 7.15 (Tarski–Vaught principle)**.** If Σ is categorical in some cardinality $\mathfrak{m} \geq \aleph_0$ and Σ has an infinite model then Σ is complete.

Proof. Let $\mathcal{A}$ be the unique model of Σ of cardinality m. Suppose $\mathcal{A} \models \varphi$.

Claim. $\Sigma \models \varphi$. Assume not. Then $(\exists \mathcal{B})(\mathcal{B} \models \Sigma \cup \{\neg\varphi\})$. If $\mathcal{B}$ is inifinite, there exists $\mathcal{B}'$ which is elementary equivalent to $\mathcal{B}$ and has cardinality $\mathfrak{m}$, a contradiction because $\mathcal{B}' \models \neg\varphi$ so $\mathcal{B}'$ cannot be isomorphic to $\mathcal{A}$. – What if $\mathcal{B}$ is finite? The T-V principle as stated is not exactly correct; Exercise: fix it. □

Theorem 7.16 (Fagin)**.** For any sentence φ, let

$$p_n := \Pr\big(\mathcal{A}_n \models \varphi\big) = \frac{\#\{\dots\}}{2^{\binom{n}{2}}}$$

where $\mathcal{A}_n$ is a random graph on n vertices. Then $\lim_{n\to\infty} p_n$ is either 0 or 1.

In fact $\lim p_n = 1 \iff R \models \varphi$; otherwise $R \models \neg\varphi$ because R is complete.

Proof. Assume $R \models \varphi$. Then a finite subset R' of R already implies φ: $R' \models \varphi$. So it suffices to consider the case when $\varphi \in R$ (why?). Let $\varphi = \varphi_{k,\ell}$ with $k+\ell = t$.

For a fixed sequence $x_1,\dots,x_k,y_1,\dots,y_\ell$ of $t = k+\ell$ distinct vertices in a graph on N vertices, the chance that a given other vertex v is "bad," i. e., it does not satisfy $\varphi_{k,\ell}$, is $\big(1-\frac{1}{2^t}\big)$. Thus the chance that every vertex is bad is $\big(1-\frac{1}{2^t}\big)^{N-t}$. The number of sequences of t distinct vertices is $N(N-1)\dots(N-t+1) < N^t$. Thus the chance there is some sequence of length t for which all vertices v are bad, is $\leq N^t\big(1-\frac{1}{2^t}\big)^{N-t}$. For fixed t, the first term grows polynomially, while the second decays exponentially; thus as N goes to ∞, the chance $\varphi_{k,\ell}$ is false goes to 0. □