Finite Model Theory
Unit 1

Dan Suciu

Spring 2018

Welcome to 599c: Finite Model Theory

- Logic is the foundation of Mathematics (see Logicomix).
- Logic is the foundation of computing (see Turing Machines).
- Finite Model Theory is Logic restricted to finite models.
- Applications of FMT: Verification, Databases, Complexity
- This course is about:
  - Classic results in Mathematical Logic
  - Classic results in Finite Model Theory
  - New results in Finite Model Theory
  - Most results are negative, but some positive results too.
- This course is not about: systems, implementation, writing programs.

Course Organization

Lectures:

- Regular time: MW 10 - 11:20, CSE 303
- Canceled: April 9, 11; May 14, 16.
- Makeup (all in CSE 303):
  4/6 (10-11:20), 4/20 (10-11:20), 5/17 (9:30-10:50), 5/18 (10-11:20)

Homework assignment:

- 6 Homework assignments
- Short problems, but some require thinking.
- Email them to me by the due date.
- Ignore points: I will grade *all 6 together* as Credit/No-credit.
- Discussion on the bboard encouraged!
- Goal: *no stress*, encourage to participate and think.

Resources

- Required (fun) reading: Logicomix.
- Libkin *Finite Model Theory*.
- Enderton *A Mathematical Introduction to Logic*.
- Barnes and Mack *An Algebraic Introduction to Logic*.
- Abiteboul, Hull, Vianu, *Database Theory*
- Several papers, talks, etc.
- Course on Friendly Logics from UPenn (by Val Tannen and Scott Weinstein) (older version: `http://www.cis.upenn.edu/~val/CIS682/`)

Course Outline

Unit 1 Classical Model Theory and Applications to FMT.

Unit 2 Games and expressibility.

Unit 3 Descriptive Complexity.

Unit 4 Query Containment.

Unit 5 Algorithmic FMT.

Unit 6 Tree Decomposition. Guest lecturer: Hung Ngo.

Unit 7 Provenance semirings. Guest lecturer: Val Tannen.

Unit 8 Semantics of datalog programs.

Structures

A vocabulary σ is a set of relation symbols $R_1, \dots, R_k$ and function symbols $f_1, \dots, f_m$, each with a fixed arity.

A structure is $\boldsymbol{D} = (D, R_1^D, \dots, R_k^D, f_1^D, \dots, f_m^D)$,
where $R_i^D \subseteq (D)^{\text{arity}(R_i)}$ and $f_j^D : (D)^{\text{arity}(f_j)} \to D$.

D = the *domain* or the *universe*.
$v \in D$ is called a *value* or a *point*.
$\boldsymbol{D}$ called a *structure* or a *model* or *database*.

Examples

A graph is $G = (V, E)$, $E \subseteq V \times V$.

A field is $\mathbb{F} = (F, 0, 1, +, \cdot)$ where

- F is a set.
- 0 and 1 are constants (i.e. functions $F^0 \to F$).
- $+$ and $\cdot$ are functions $F^2 \to F$.

An ordered set is $\boldsymbol{S} = (S, \leq)$ where $\leq \subseteq S \times S$.

A database is $\boldsymbol{D} = (\text{Domain}, \text{Customer}, \text{Order}, \text{Product})$.

Discussion

- We don't really need functions, since $f : D^k \to D$ is represented by its graph $\subseteq D^{k+1}$, but we keep them when convenient.
- If f is a 0-ary function $D^0 \to D$, then it is a constant D, and we denote it c rather than f.
- $\boldsymbol{D}$ can be a finite or an infinite structure.

First Order Logic

Fix a vocabulary σ and a set of variables $x_1, x_2, \ldots$

Terms:

- Every constant c and every variable x is a term.
- If $t_1, \ldots, t_k$ are terms then $f(t_1, \ldots, t_k)$ is a term.

Formulas:

- $\boldsymbol{F}$ is a formula (means *false*).
- If $t_1, \ldots, t_k$ are terms, then $t_1 = t_2$ and $R(t_1, \ldots, t_k)$ are formulas.
- If φ, ψ are formulas, then so are $\varphi \to \psi$ and $\forall x(\varphi)$.

Discussion

$\boldsymbol{F}$ often denoted: `false` or $\perp$ or 0.

$=$ is not always part of the language

Derived operations:

- $\neg\varphi$ is a shorthand for $\varphi \to \boldsymbol{F}$.
- $\varphi \vee \psi$ is a shorthand for $(\neg\varphi) \to \psi$.
- $\varphi \wedge \psi$ is a shorthand for $\neg(\varphi \vee \psi)$.
- $\exists x(\varphi)$ is a shorthand for $\neg(\forall x(\neg\varphi))$.

Formulas and Sentences

We say that $\forall x(\varphi)$ *binds* x in φ. Every occurrence of x in φ is *bound*. Otherwise it is *free*.

A sentence is a formula φ without free variables.

E.g. formula $\exists y(E(x,y) \wedge E(y,z))$.

E.g. sentence $\exists x \forall z \exists y(E(x,y) \wedge E(y,z))$.

Truth

Let φ be a formula with free variables $\boldsymbol{x} = (x_1, \ldots, x_k)$.
Let $\boldsymbol{D}$ be a structure, and $\boldsymbol{a} = (a_1, \ldots, a_k) \in D^k$.
We say that φ is true in $\boldsymbol{D}$, written:

$$D \vDash \varphi[\boldsymbol{a}/\boldsymbol{x}]$$

if:

- φ is $x_i = x_j$ and a_i, a_j are the same value.
- φ is $R(x_{i_1}, \ldots, x_{i_n})$ and $(a_{i_1}, \ldots, a_{i_n}) \in R^D$.
- φ is $\psi_1 \rightarrow \psi_2$ and $\boldsymbol{D} \nvDash \psi_1[\boldsymbol{a}/\boldsymbol{x}]$, or $\boldsymbol{D} \vDash \psi_1[\boldsymbol{a}/\boldsymbol{x}]$ and $\boldsymbol{D} \vDash \psi_2[\boldsymbol{a}/\boldsymbol{x}]$.
- φ is $\forall y(\psi)$, and, forall $b \in D$, $\boldsymbol{D} \vDash \psi[(a_1, \ldots, a_k, b)/(x_1, \ldots, x_k, y)]$.

Problems

- Classical model theory:
  - *Satisfiability* Is φ true in *some* structure $\boldsymbol{D}$?
  - *Validity* Is φ true in *all* structures $\boldsymbol{D}$?
- Finite model theory, databases, verification:
  - *Finite satisfiability/validity* Is φ true in some/every *finite* structure $\boldsymbol{D}$?
  - *Model checking* Given φ, $\boldsymbol{D}$, determine whether $\boldsymbol{D} \models \varphi$.
  - *Query evaluation* Given $\varphi(\boldsymbol{x})$, $\boldsymbol{D}$, compute $\{\boldsymbol{a} \mid \boldsymbol{D} \models \varphi[\boldsymbol{a}/\boldsymbol{x}]\}$.

What do these sentences say about *D*?

$$\exists x \exists y \exists z (x \neq y) \wedge (x \neq z) \wedge (y \neq z)$$

"There are at least three elements", i.e. $|D| \geq 3$

$$\exists x \exists y \forall z (z = x) \vee (z = y)$$

"There are at most two elements", i.e. $|D| \leq 2$

What do these sentences say about $\boldsymbol{D}$?

$$\forall x \exists y E(x,y) \vee E(y,x)$$

"There are no isolated nodes"

$$\forall x \forall y \exists z E(x,z) \wedge E(z,y)$$

"Every two nodes are connected by a path of length 2"

$$\begin{gathered}\exists x \exists y \exists z (\forall u (u = x) \vee (u = y) \vee (u = z)) \\ \wedge \neg E(x,x) \wedge E(x,y) \wedge \neg E(x,z) \\ \wedge \neg E(y,z) \wedge \neg E(y,y) \wedge E(y,z) \\ \wedge E(z,x) \wedge \neg E(z,y) \wedge \neg E(z,z)\end{gathered}$$

It completely determines the graph: $D = \{a, b, c\}$ and $a \to b \to c \to a$.

Logical Implication

Fix a set of sentences Σ (may be infinite).

Σ implies φ, $\Sigma \models \varphi$, if every model of Σ is also a model of φ:
$\boldsymbol{D} \models \Sigma$ implies $\boldsymbol{D} \models \varphi$.

$\text{Con}(\Sigma) \stackrel{\text{def}}{=} \{\varphi \mid \Sigma \models \varphi\}$. Somtimes called the *theory* of Σ, $\text{Th}(\Sigma)$.

Σ finitely implies φ, $\Sigma \models_{\text{fin}} \varphi$ if every *finite* model of Σ is also a model of φ.

Discussion

- $\boldsymbol{F} \models \varphi$ for any sentence φ why?.
- $\Sigma \models \boldsymbol{F}$ iff Σ is unsatisfiable why?.
- If $\Sigma \models \varphi$ and $\Sigma, \varphi \models \psi$ then $\Sigma \models \psi$ why?.
- If $\Sigma \models \varphi$ then $\Sigma \models_{\text{fin}} \varphi$, but the converse fails in general why?.
  Let λ_n say "there are at least n elements, and $\Sigma = \{\lambda_n \mid n \geq 1\}$.
  Then $\Sigma \models_{\text{fin}} \boldsymbol{F}$ but $\Sigma \not\models \boldsymbol{F}$ why?.
- If $\models \varphi$ then we call φ a *tautology*.

Theory

A theory is a set of sentences Σ closed under implication, i.e. $\Sigma = \text{Con}(\Sigma)$.

A theory Σ is complete if, for every sentence φ, either $\varphi \in \Sigma$ or $\neg\varphi \in \Sigma$.

The theory of a set of structures $\mathcal{D}$ is

$$\text{Th}(\mathcal{D}) \stackrel{\text{def}}{=} \{\varphi \mid \varphi \text{ is true in every } \boldsymbol{D} \in \mathcal{D}\}$$ closed under implication?

For a single structure $\boldsymbol{D}$, $\text{Th}(\boldsymbol{D})$ is complete why?

Discussion

Which of the following theories are complete?

- The theory of fields $\mathbb{F} = (F, 0, 1, +, \cdot)$. No: $\exists x(x^2 + 1 = 0)$
- The theory $\mathrm{Th}(\mathbb{R})$ (vocabulary $0, 1, +, \cdot$). yes
- The theory of total orders:

$$\forall x \forall y \neg((x < y) \wedge (y < x))$$
$$\forall x \forall y((x < y) \vee (x = y) \vee (y < x))$$
$$\forall x \forall y \forall z((x < y) \wedge (y < z) \to (x < z))$$

  No: $\forall x \exists y(x < y)$.

- The theory of dense total orders without endpoints: axioms above plus

$$\text{Dense:} \quad \forall x \forall y(x < y \to \exists v(x < v < y))$$
$$\text{W/o Endpoints:} \quad \forall x \exists u \exists w(u < x < w)$$

  Yes! Will prove later

The Sentence Map

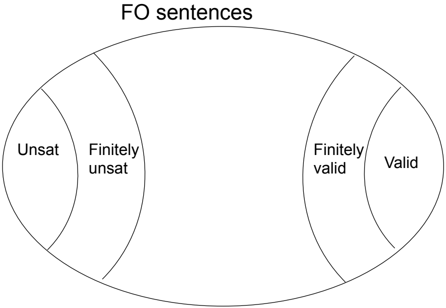

Give examples for each of the five classes

$\exists x(\neg(x = x))$ "< is a dense total order" $\exists x \exists y(E(x, y))$

"if < is a total order, then it has a maximal element" $\forall x(x = x)$

The Zero-One Law for FO

- Some sentences are neither true (in all structures) nor false.
- The Zero-One Law says this: over *finite* structures, every sentence is true or false *with high probability*.
- Proven by Fagin in 1976 (part of his PhD thesis).
- Although the statement is about *finite* structures, the proof uses theorems on *finite and infinite* structures.

The Zero-One Law for FO

Consider a relational vocabulary (i.e. no functions, no constants).
Let φ be a sentence. Forall $n \in \mathbb{N}$ denote:

$$\#_n\varphi \stackrel{\text{def}}{=} |\{\boldsymbol{D} \mid D = [n], \boldsymbol{D} \vDash \varphi\}$$

$$\#_n\boldsymbol{T} \stackrel{\text{def}}{=} \text{number of models with universe } [n]$$

$$\mu_n(\varphi) \stackrel{\text{def}}{=} \frac{\#_n\varphi}{\#_n\boldsymbol{T}}$$

Theorem (Fagin'1976)

For every sentence φ, either $\lim_{n\to\infty}\mu_n(\varphi) = 0$ or $\lim_{n\to\infty}\mu_n(\varphi) = 1$.

Informally: for every φ, its probability goes to either 0 or 1, when $n \to \infty$; it is either almost certainly true, or almost certainly false.

Examples

Vocabulary of graphs: $\sigma = \{E\}$. Compute these probabilities:

$\varphi = \forall x \forall y E(x, y)$ $\quad \#_n(\varphi) = 1$ $\quad \mu_n = \frac{1}{2^{n^2}} \to 0$

$\varphi = \exists x \exists y E(x, y)$ $\quad \#_n(\varphi) = 2^{n^2} - 1$ $\quad \mu_n = \frac{2^{n^2} - 1}{2^{n^2}} \to 1$

$\varphi = \forall x \exists y E(x, y)$ $\quad \mu_n = \frac{(2^n - 1)^n}{2^{n^2}} \to 1$

The Sentence Map Revised

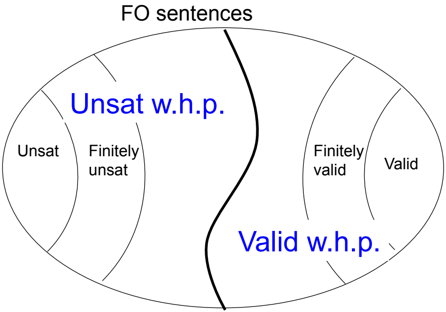

Discussion

Attempted proof: Derive the general formula $\#_n\varphi$, then compute $\lim \#_n\varphi/2^{n^2}$ and observe it is 0 or 1.

Problem: we don't know how to compute $\#_n\varphi$ in general: there is evidence this is "hard"

Instead, we will prove the 0/1 law using three results from classical model theory.

Three Classical Results in Model Theory

We will discuss and prove:

- Compactness Theorem.
- Lövenheim-Skolem Theorem.
- Los-Vaught Test.

Then will use them to prove Fagin's 0/1 Law for First Order Logic.

Later we will discuss:

- Gödel's completeness theorem.
- Decidability of theories.
- Gödel's incompleteness theorem.

Compactness Theorem

Recall: Σ is satisfiable if it has a model, i.e. there exists $\boldsymbol{D}$ s.t. $\boldsymbol{D} \vDash \varphi$, forall $\varphi \in \Sigma$.

Theorem (Compactness Theorem)

If every finite subset of Σ is satisfiable, then Σ is satisfiable.

Short: if Σ is finitely satisfiable[1], then it is satisfiable.

Considered to be the most important theorem in Mathematical Logic.

[1] Don't confuse with saying "Σ has a finite model"!

Compactness Theorem - Alternative Formulation

The following is equivalent to the Compactness Theorem:

Theorem

If $\Sigma \models \varphi$ then there exists a finite subset $\Sigma_{fin} \subseteq \Sigma$ s.t. $\Sigma_{fin} \models \varphi$.

Proof: assume Compactness holds, and assume $\Sigma \models \varphi$. If $\Sigma_{\text{fin}} \not\models \varphi$ for any finite subset, then the set $\Sigma \cup \{\neg\varphi\}$ is finitely satisfiable, hence it is satisfiable, contradiction.

In the other direction, let Σ be finitely satisfiable. If Σ is not satisfiable, then $\Sigma \models \boldsymbol{F}$, hence there is a finite subset s.t. $\Sigma_{\text{fin}} \models \boldsymbol{F}$, contradicting the fact that Σ_{fin} has a model.

Warmup: The Propositional Case

Let Σ be a set of Boolean formulas, a.k.a. Propositional formulas.

Theorem (Compactness for Propositional Logic)

If every finite subset of Σ is satisfiable, then Σ is satisfiable.

Application: $G = (V, E)$ is an infinite graph s.t. every finite subgraph is 3-colorable. Prove: G is 3-colorable.
Boolean Variables: $\{R_i, G_i, B_i \mid i \in V\}$ ("i is colored Red/Green/Blue").

$$
\begin{aligned}
\Sigma = & \{R_i \vee G_i \vee B_i \mid i \in V\} && \text{every node gets some color} \\
& \cup \{\neg R_i \vee \neg R_j \mid (i,j) \in E\} && \text{adjacent nodes get different colors} \\
& \cup \{\neg G_i \vee \neg G_j \mid (i,j) \in E\} \\
& \cup \{\neg B_i \vee \neg B_j \mid (i,j) \in E\}
\end{aligned}
$$

Every finite subset of Σ is satisfiable, hence so is Σ.

Warmup: The Propositional Case

Two steps:

- Extend Σ to $\bar{\Sigma}$ that is both complete and finitely satisfiable.
- Use the Inductive Structure of a complete and finite satisfiable set.

Step 1: Extend Σ to a complete $\bar{\Sigma}$

Enumerate all formulas $\varphi_1, \varphi_2, \ldots$, and define:

$$\Sigma_0 = \Sigma \qquad \Sigma_{i+1} = \begin{cases} \Sigma_i \cup \{\varphi_i\} & \text{if } \Sigma_i \cup \{\varphi_i\} \text{ is finitely satisfiable} \\ \Sigma_i \cup \{\neg\varphi_i\} & \text{if } \Sigma_i \cup \{\neg\varphi_i\} \text{ is finitely satisfiable} \end{cases}$$

One of the two cases above must hold, because, otherwise both $\Sigma_i \cup \{\varphi_i\}$ and $\Sigma_i \cup \{\neg\varphi_i\}$ are finitely UNSAT, then $\Sigma_{\text{fin}} \cup \{\varphi_i\}$ and $\Sigma'_{\text{fin}} \cup \{\neg\varphi_i\}$ are UNSAT for $\Sigma_{\text{fin}}, \Sigma'_{\text{fin}} \subseteq \Sigma_i$, hence $\Sigma_{\text{fin}} \cup \Sigma'_{\text{fin}}$ is UNSAT, contradiction.

Then $\bar{\Sigma} \stackrel{\text{def}}{=} \bigcup_i \Sigma_i$ is complete and finitely satisfiable

Step 2: Inductive Structure of a Complete Set

Lemma

If $\bar{\Sigma}$ is a complete, and finitely satisfiable set, then:

- $\varphi \wedge \psi \in \bar{\Sigma}$ *iff* $\varphi, \psi \in \bar{\Sigma}$.
- $\varphi \vee \psi \in \bar{\Sigma}$ *iff* $\varphi \in \bar{\Sigma}$ *or* $\psi \in \bar{\Sigma}$.
- $\neg\varphi \in \bar{\Sigma}$. *iff* $\varphi \notin \bar{\Sigma}$

Proof in class

To prove Compactness Theorem for Propositional Logic, define this model:

$$\theta(X) \stackrel{\text{def}}{=} 1 \text{ if } X \in \bar{\Sigma}$$

$$\theta(X) \stackrel{\text{def}}{=} 0 \text{ if } X \notin \bar{\Sigma}$$

Then $\theta(\varphi) = 1$ iff $\varphi \in \bar{\Sigma}$ (proof by induction on φ).
Hence θ is a model for $\bar{\Sigma}$, and thus for Σ.

Proof of the Compactness Theorem for FO

In addition to the propositional case, we need to handle $\exists$

Σ is *witness-complete* if, forall $\exists x(\varphi) \in \Sigma$, there is some c s.t. $\varphi[c/x] \in \Sigma$.

Extend Σ to a complete *and* witness-complete set $\bar{\Sigma}$, by adding countably many new constants $c_1, c_2, \ldots$ proof in class

Define a model $\boldsymbol{D}$ for $\bar{\Sigma}$ as follows:

- Its domain D consists of all terms[2].
- For each relation R, $R^D \stackrel{\text{def}}{=} \{(t_1, \ldots, t_k) \mid R(t_1, \ldots, t_k) \in \bar{\Sigma}\}$.
- Similarly for a function f.

Check this is a model of $\bar{\Sigma}$ (by showing $\boldsymbol{D} \vDash \varphi$ iff $\varphi \in \bar{\Sigma}$), hence of Σ.

[2] Up to the equivalence defined by $t_1 = t_2 \in \bar{\Sigma}$.

Discussion

- Compactness Theorem is considered the most important theorem in Mathematical Logic.
- Our discussion was restricted to a finite vocabulary σ, but compactness holds for any vocabulary; e.g. think of having infinitely many constants c
- Gödel proved compactness as a simple consequence of his completeness theorem.
- We will later prove Gödel's completeness following a similar proof as for compactness.

Application of the Compactness Theorem

Can we say in FO "the world is inifite"? Or "the world is finite"?

- Find a set of sentences Λ whose models are precisely the infinite structures.
  $\Lambda = \{\lambda_1, \lambda_2, \ldots\}$ where λ_n says "there are $\geq n$ elements":

$$\lambda_n = \exists x_1 \cdots \exists x_n \bigwedge_{i<j} (x_i \neq x_j)$$

- Find a set of sentences Σ whose models are precisely the finite structures.
  Imposible! If we could, then $\Sigma \cup \Lambda$ were finitely satisfiable, hence satisfiable, constradiction.

Löwenheim-Skolem Theorem

Suppose the vocabulary σ is finite.

Theorem (Löwenheim-Skolem)

If Σ admits an infinite model, then it admits a countable model.

In other words, we can say "the world is infinite", but we can't say how big it is.

Background: Cardinal Numbers

If there is a bijection $f : A \to B$ then we say that A, B are *equipotent*, or *equipollent*, or *equinumerous*, and write $A \cong B$.

We write $|A|$ for the equivalence class of A under $\cong$.

Definition

A *cardinal number* is an equivalence class $|A|$.
We write $|A| \leq |B|$ if there exists an injective function $A \to B$;
equivalently, if there exists a surjective function $B \to A$.

Background: Cardinal Numbers

- 4 is a cardinal number, why? The equivalence class of $\{a, b, c, d\}$.
- $4 < 7$, why? $\{a, b, c, d\} \to \{x, y, z, u, v, w, m\}$: $a \mapsto x$, $b \mapsto y$ etc.
- $\aleph_0$ is the *infinite countable cardinal*; equivalence class of $\mathbb{N}$.
- $\mathfrak{c}$ is the *cardinality of the continuum*; equivalence class of $\mathbb{R}$.
- What is the cardinality of the even numbers $\{0, 2, 4, 6, \ldots\}$? $\aleph_0$.
- What is the cardinality of $[0, 1]$? $\mathfrak{c}$.
- What is the cardinality of $\mathbb{Q}$? $\aleph_0$
- Is there a cardinal number between $\aleph_0$ and $\mathfrak{c}$? Either yes or no! (Recall Logicomix!)
- What is the cardinality of the set of sentences over a finite vocabulary? $\aleph_0$

Löwenheim-Skolem Theorem: Proof

Suppose the vocabulary σ is finite or countable.

Theorem

If Σ admits an infinite model, then it admits a countable model.

Proof in four steps:

- Write each sentence $\varphi \in \Sigma$ in prenex-normal form: $(\forall|\exists)^*\psi$.
- "Skolemize" Σ: replace each $\exists$ with a fresh "Skolem" function f, e.g.

$$\forall x \exists y \forall z \exists u(\varphi) \mapsto \forall x \forall z(\varphi[f_1(x)/y, f_2(x,z)/u])$$

  Let Σ' be the set of Skolemized sentences.
- Property of Skolemization: Σ satisfiable iff Σ' satisfiable. In class
- Proof of Löwenheim-Skolem. Let $\boldsymbol{D} \models \Sigma$; then $\boldsymbol{D} \models \Sigma'$ (by interpreting the Skolem functions appropriately).
- Let: D_0 be any countable subset of D,
  $D_{i+1} = \{f^D(d_1, \ldots, d_k) \mid d_1, \ldots, d_k \in D_i, f \in \sigma\}$.
  Then $\bigcup_i D_i$ is countable and a model of Σ' why?.

Discussion

- We have assumed that σ is finite, or countable.
- If σ has cardinality κ, then the Löwenheim-Skolem Theorem says that there exists a model of cardinality κ.
- The *upwards* version of the Löwenheim-Skolem Theorem[3] if Σ has a model of infinite cardinality κ and $\kappa < \kappa'$ then it also has a model of cardinality κ'.

  Proof: add to σ constants $c_k, k \in \kappa'$, add axioms $c_i \neq c_j$ for $i \neq j$. By compactness there is a model; then we repeat the previous proof of Löwenheim-Skolem.

[3] Called: Löwenheim-Skolem-Tarski theorem.

The Los-Vaught Test

Simple observation: if $\boldsymbol{D}_1, \boldsymbol{D}_2$ are *isomorphic* then $\mathrm{Th}(\boldsymbol{D}_1) = \mathrm{Th}(\boldsymbol{D}_2)$.

Call Σ $\aleph_0$-categorical if any two countable models of Σ are isomorphic.

Theorem (Los-Vaught Test)

If Σ has no finite models and is $\aleph_0$ categorical then it is complete.

Proof. Suppose otherwise: there exists φ s.t. $\Sigma \not\models \neg\varphi$ and $\Sigma \not\models \varphi$. Then:

- $\Sigma \cup \{\varphi\}$ has a model $\boldsymbol{D}_1$; assume it is countable why can we?
- $\Sigma \cup \{\neg\varphi\}$ has a model $\boldsymbol{D}_2$; assume it is countable.
- Then $\boldsymbol{D}_1, \boldsymbol{D}_2$ are isomorphic.
- Contradiction because $\boldsymbol{D}_1 \models \varphi$ and $\boldsymbol{D}_2 \models \neg\varphi$.

Application of the Los-Vaught Test

The *theory of dense linear orders without endpoints* is complete.

$$\forall x \forall y \neg((x < y) \wedge (y < x))$$
$$\forall x \forall y((x < y) \vee (x = y) \vee (y < x))$$
$$\forall x \forall y \forall z((x < y) \wedge (y < z) \rightarrow (x < z))$$

Dense: $\forall x \forall y(x < y \rightarrow \exists v(x < v < y))$

W/o Endpoints: $\forall x \exists u \exists w(u < x < w)$

Note: just "total order" is not complete!

Proof: we apply the Los-Vaught test.

Let $\boldsymbol{A}, \boldsymbol{B}$ be countable models. Construct inductively $A_i \subseteq A$, $B_i \subseteq B$, and isomorphism $f_i : (A_i, <) \rightarrow (B_i, <)$, using the Back and Forth argument.

The Back-and-Forth argument

$\boldsymbol{A} = (\{a_1, a_2, \ldots\}, <)$, $\boldsymbol{B} = (\{b_1, b_2, \ldots\}, <)$ are total orders w/o endpoints. Prove they are isomorphic.

$A_0 \stackrel{\text{def}}{=} \emptyset$, $B_0 \stackrel{\text{def}}{=} \emptyset$.
Assuming $(A_{i-1}, <) \cong (B_{i-1}, <)$, extend to $(A_i, <) \cong (B_i, <)$ as follows:

- Add a_i and any $b \in B$ s.t. $(A_{i-1} \cup \{a_i\}, <) \cong (B_{i-1} \cup \{b\})$.

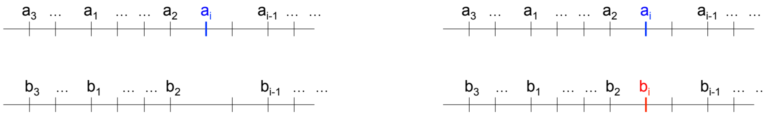

- Add b_i and any matching $a \in A$.

Then $A = \bigcup A_i$, $B = \bigcup B_i$ and $(A, <) \cong (B, <)$.

Discussion

The Los-Vaught test applies to any cardinal number, as follows:

- If Σ has no finite models and is categorical in some infinite cardinal κ (meaning: any two models of cardinality κ are isomorphic) then Σ is complete.

Useful for your homework.

Recap: Three Classical Results in Model Theory

We proved:

- Compactness Theorem.
- Lövenheim-Skolem Theorem.
- Los-Vaught Test.

Next, we use them to prove Fagin's 0/1 Law for First Order Logic.

Proof of the Zero-One Law: Plan

Zero-one Law: $\lim_{n\to\infty} \mu_n(\varphi)$ is 0 or 1, for every φ

For simplicity, assume vocabulary of graphs, i.e. only binary E.

- Define a set Σ of *extension axioms*, $EA_{k,\Delta}$
- We prove that $\lim_n \mu_n(EA_{k,\Delta}) = 1$
- Hence Σ is finitely satisfiable.
- By compactness: Σ has a model.
- By Löwenheim-Skolem: has a countable model (called **the Rado graph *R***, when undirected).
- We prove that all countable models of Σ are isomorphic.
- By Los-Vaught: Σ is complete.
- Then $\Sigma \models \varphi$ implies $\lim \mu_n(\varphi) = 1$ and $\Sigma \not\models \varphi$ implies $\lim \mu_n(\varphi) = 0$.

The Extension Formulas and the Extension Axioms

For $k > 0$ denote $S_k = ([k] \times \{k\}) \cup (\{k\} \times [k])$ and $\Delta \subseteq S_k$.

$$EF_{k,\Delta}(x_1, \ldots, x_{k-1}, x_k) = \bigwedge_{(i,j) \in \Delta} E(x_i, x_j) \wedge \bigwedge_{(i,j) \in S_k - \Delta} \neg E(x_i, x_j)$$

$$EA_{k,\Delta} = \forall x_1 \ldots \forall x_{k-1} (\bigwedge_{i<j<k} (x_i \neq x_j)) \rightarrow \exists x_k (\bigwedge_{i<k} (x_k \neq x_i) \wedge EF_{k,\Delta})$$

Intuition: we can extend the graph as prescribed by Δ.

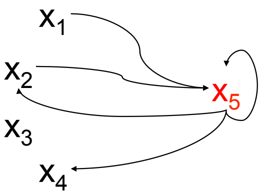

$E(x_1, x_5) \wedge \neg E(x_5, x_1) \wedge$
$E(x_2, x_5) \wedge E(x_5, x_2) \wedge$
$\neg E(x_3, x_5) \wedge \neg E(x_5, x_3) \wedge$
$\neg E(x_4, x_5) \wedge E(x_5, x_4) \wedge$
$E(x_5, x_5)$

How many extension axioms are there for $k = 5$?

Proof of $\lim_n \mu_n(EA_{k,\Delta}) = 1$

$$EF_{k,\Delta}(x_1, \ldots, x_{k-1}, x_k) = \bigwedge_{(i,j)\in\Delta} E(x_i, x_j) \wedge \bigwedge_{(i,j)\in S_k-\Delta} \neg E(x_i, x_j)$$

$$EA_{k,\Delta} = \forall x_1 \ldots \forall x_{k-1}\Big(\bigwedge_{i<j<k}(x_i \neq x_j)\Big) \to \exists x_k\Big(\bigwedge_{i<k}(x_k \neq x_i) \wedge EF_{k,\Delta}\Big)$$

$$\mu_n(\neg EA_{k,\Delta}) = \mu_n\Big(\exists x_1 \ldots \exists x_{k-1}\Big(\bigwedge(x_i \neq x_j) \wedge \forall x_k\Big(\bigwedge(x_k \neq x_i) \to \neg EF_{k,\Delta}\Big)\Big)\Big)$$

$$\leq \sum_{a_1,\ldots,a_{k-1}\in[n], a_i\neq a_j} \mu_n\left(\bigwedge_{a_k\in[n]-\{a_1,\ldots,a_{k-1}\}} \neg EF_{k,\Delta}(a_1, \ldots, a_{k-1}, a_k)\right)$$

$$= \sum_{a_1,\ldots,a_{k-1}\in[n], a_i\neq a_j} \prod_{a_k\in[n]-\{a_1,\ldots,a_{k-1}\}} \mu_n(\neg EF_{k,\Delta}(a_1, \ldots, a_k)) \qquad \text{why?}$$

$$= \sum_{a_1,\ldots,a_{k-1}\in[n], a_i\neq a_j} \prod_{a_k\in[n]-\{a_1,\ldots,a_{k-1}\}} c \qquad \text{where } c = 1 - \frac{1}{2^{2k-1}} < 1$$

$$\leq n^{k-1} c^{n-k+1} \to 0$$

Extension Axioms Have a Countable Model

Let $\Sigma = \{EA_{k,\Delta} \mid k > 0, \Delta \subseteq S_k\}$ be the set of extension axioms.

Σ is finitely satisfiable why?

Because forall $\varphi_1, \ldots, \varphi_m \in \Sigma$, $\mu_n(\varphi_1 \wedge \cdots \wedge \varphi_m) \to 1$

Hence, when n is large, there are *many* finite models for $\varphi_1, \ldots, \varphi_m$!

By compactness, Σ has a model.

By Löwenheim-Skolem, Σ has a countable model.

Extension Axioms have a Unique Countable Model

Need to prove: any two countable models $\boldsymbol{A}, \boldsymbol{B}$ of Σ are isomorphic.

Will use the Back-and-Forth construction!

Let $\boldsymbol{A} = \{a_1, a_2, \dots\}$, $\boldsymbol{B} = \{b_1, b_2, \dots\}$.

By induction on k, construct $(A_k, E_k) \cong (B_k, E'_k)$, using the back-and-forth construction and the fact that both $\boldsymbol{A}$, $\boldsymbol{B}$ satisfy Σ.

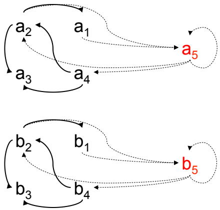

Hence, there is a unique (up to isomorphism) countable model. Called *The Random Graph* or *Rado Graph*, $\boldsymbol{R}$ for undirected graphs. See Libkin.

Proof of the Zero-One Law

Let φ be any sentence: we'll prove $\mu_n(\varphi)$ tends to either 0 or 1.

Σ is complete, hence either $\Sigma \vDash \varphi$ or $\Sigma \vDash \neg\varphi$.

Assume $\Sigma \vDash \varphi$.

By compactness, then there exists a finite set $\{\psi_1, \ldots, \psi_m\} \vDash \varphi$

Thus, $\mu_n(\varphi) \geq \mu_n(\psi_1 \wedge \cdots \wedge \psi_m) \to 1$ why?

Assume $\Sigma \vDash \neg\varphi$: then $\mu_n(\neg\varphi) \to 1$, hence $\mu_n(\varphi) \to 0$.

Discussion

- The 0/1 law does *not* hold if there constants:
  e.g. $\lim \mu_n R(a, b) = 1/2$ (neither 0 nor 1).
  Where in the proof did we use this fact? (Homework!)
- The Random Graph $\boldsymbol{R}$ satisfies precisely those sentences for which $\lim \mu_n(\varphi) = 1$.
- We proved the 0/1 law when every edge $E(i, j)$ has probability $p = 1/2$.
  The same proof also holds when every edge has probability $p \in (0, 1)$ (independent of n).
- The Erdös-Rényi random graph $G(n, p)$ allows p to depend on n. 0/1 law for FO may or may not hold. discuss more in class

A Cool Application: Non-standard Analysis

"Infinitezimals" have been used in calculus since Leibniz and Newton.

But they are not rigorous! Recall Logicomix.

Example: compute the derivative of x^2:

$$\frac{dx^2}{dx} = \frac{(x + dx)^2 - x^2}{dx} = \frac{2 \cdot x \cdot dx + (dx)^2}{dx} = 2x + dx \simeq 2x$$

because dx is "infinitely small", hence $dx \simeq 0$.

Robinson in 1961 showed that how to define infinitezimals rigorously (and easily) using the compactness theorem!

The Nonstandard Reals

$\mathbb{R}$ = the true real numbers.

- Let σ be the vocabulary of all numbers, functions, relations:
  - Every number in $\mathbb{R}$ has a symbol: $0, -5, \pi, \ldots$
  - Every function $\mathbb{R}^k \to \mathbb{R}$ has a symbol: $+, *, -, \sin, \ldots$
  - Every relation $\subseteq R^k$ has a symbol: $<, \geq, \ldots$
- Let $\text{Th}(\mathbb{R})$ all true sentences, e.g.:

$$\forall x(x^2 \geq 0)$$
$$\forall x \forall y(|x+y| \leq |x| + |y|)$$
$$\forall x(\sin(x+\pi) = -\sin(x))$$

- Let Ω be a new constant, and $\Sigma \stackrel{\text{def}}{=} \text{Th}(\mathbb{R}) \cup \{n < \Omega \mid n \in \mathbb{N}\}$. "$\Omega$ is bigger than everything".
- Σ has a model $^*\mathbb{R}$. WHY?

What exactly is $^*\mathbb{R}$???

The Nonstandard Reals

- Every number in $\mathbb{R}$ also exists in ${}^*\mathbb{R}$: $0, -5, \pi, \ldots$
- Every function $\mathbb{R}^k \to \mathbb{R}$ has an extension $({}^*\mathbb{R})^k \to {}^*\mathbb{R}$.
- Every relation $\subseteq \mathbb{R}^k$ has a corresponding $\subseteq ({}^*\mathbb{R})^k$.
- $\omega \stackrel{\text{def}}{=} 1/\Omega$; the, $0 < \omega < c$ forall real $c > 0$. Infinitezimal! others?
- The infinitezimals are $\mathcal{I} \stackrel{\text{def}}{=} \{v \in {}^*\mathbb{R} \mid \forall c \in \mathbb{R}, c > 0 : |v| < c\}$
  The finite elements are $\mathcal{F} \stackrel{\text{def}}{=} \{v \in {}^*\mathbb{R} \mid \exists c \in \mathbb{R}, |v| < c\}$
- $2\omega, \omega^3, \sin(\omega)$ are infinitezimals; 0.001 is not.
- $\pi, 0.001, 10^{10^{10}}$ are finite; $\Omega, \Omega/1000, \Omega^{\Omega^{\Omega}}$ are not.

The Nonstandard Reals

Infinitezimals closed under $+, -, *$; $x, y \in \mathcal{I}$ implies $x + y, x - y, x * y \in \mathcal{I}$

Finite elements closed under $+, -, *$; $x, y \in \mathcal{F}$ implies $x + y, x - y, x * y \in \mathcal{F}$

Call $x, y \in {}^*\mathbb{R}$ infinitely close if $x - y \in \mathcal{I}$; write $x \simeq y$.

Fact: $\simeq$ is an equivalence relation. Exercise!

Now we can work with infinitezimals rigorously:

$$\frac{dx^2}{dx} = \frac{(x + dx)^2 - x^2}{dx} = \frac{2 \cdot x \cdot dx + (dx)^2}{dx} = 2x + dx \simeq 2x$$

Two Other Classical Theorem (which everyone should know!)

- Gödel's completeness theorem.
- Gödel's incompleteness theorem.

We discuss them next

Gödel's Completeness Theorem

- Part of Gödel's PhD Thesis. (We need to raise the bar at UW too.)
- It says that, using some reasonable axioms:
  $\Sigma \models \varphi$ iff there exists a syntactic proof of φ from Σ.
- Completeness $\Leftrightarrow$ Compactness ($\Rightarrow$ is immediate; $\Leftarrow$ no easy proof).
- Instead, proof of Completeness Theorem is similar to Compactness.
- The Completeness Theorem depends on the rather ad-hoc choice of axioms, hence mathematicians consider it less deep than compactness.

Axioms

There are dozens of choices[4] for the axioms[5]. Recall $\neg\varphi$ is $\varphi \rightarrow \boldsymbol{F}$.

$$
\begin{aligned}
A_1 :& \varphi \rightarrow (\psi \rightarrow \varphi) \\
A_2 :& (\varphi \rightarrow (\psi \rightarrow \gamma)) \rightarrow ((\varphi \rightarrow \psi) \rightarrow (\varphi \rightarrow \gamma)) \\
A_3 :& \neg\neg\varphi \rightarrow \varphi \\
A_4 :& \forall x \varphi \rightarrow \varphi[t/x] && \text{for any term } t \\
A_5 :& (\forall x(\varphi \rightarrow \psi)) \rightarrow (\forall x(\varphi) \rightarrow \forall x(\psi))) \\
A_6 :& \varphi \rightarrow \forall x(\varphi) && x \notin \text{FreeVars}(\varphi) \\
A_7 :& x = x \\
A_8 :& (x = y) \rightarrow (\varphi \rightarrow \varphi[y/x])
\end{aligned}
$$

These are axiom *schemas*: each A_i defines an infinite set of formulas.

[4] $A_1 - A_8$ are a combination of axioms from Barnes&Mack and Enderton.

[5] Fans of the Curry-Howard isomorphisms will recognize typed λ-calculus in A_1, A_2.

Proofs

Let Σ be a set of formulas.

Definition

A *proof* or a *deduction* is a sequence $\varphi_1, \varphi_2, \ldots, \varphi_n$ such that[a], for every i:

- φ_i is an Axiom, or $\varphi_i \in \Sigma$ or,
- φ_i is obtained by modus ponens from earlier φ_j, φ_k ($\varphi_k \equiv (\varphi_j \to \varphi_i)$.)

[a] There is no Generalization Rule since it follows from A_6 (Enderton).

Definition

We say that φ is *provable*, or *deducible* from Σ, and write $\Sigma \vdash \varphi$, if there exists a proof sequence ending in φ.
If $\vdash \varphi$ then we call φ a *theorem*.

$\mathrm{Ded}(\Sigma)$ is the set of formulas φ provable from Σ.

Discussion

- $\Sigma \models \varphi$ is semantics: it says something about truth.
- $\Sigma \vdash \varphi$ is syntactic: an application of ad-hoc rules.
- Example: prove that $\varphi \to \varphi$:

$$\begin{aligned}
A_1 :& \varphi \to ((\varphi \to \varphi) \to \varphi) \\
A_2 :& (\varphi \to ((\varphi \to \varphi) \to \varphi)) \to ((\varphi \to (\varphi \to \varphi)) \to (\varphi \to \varphi)) \\
\text{MP} :& (\varphi \to (\varphi \to \varphi)) \to (\varphi \to \varphi) \\
A_1 :& (\varphi \to (\varphi \to \varphi)) \\
\text{MP} :& (\varphi \to \varphi)
\end{aligned}$$

- Prove at home $\boldsymbol{F} \to \varphi$ and $\varphi \to \psi, \psi \to \omega \vdash \varphi \to \omega$.

Consistency

Definition

Σ is called inconsistent if $\Sigma \vdash \boldsymbol{F}$. Otherwise we say Σ is consistent.

Σ is inconsistent iff for every φ, $\Sigma \vdash \varphi$
Proof: $\vdash \boldsymbol{F} \rightarrow \varphi$.

Σ is inconsistent iff there exists φ s.t. both $\Sigma \vdash \varphi$ and $\Sigma \vdash \neg\varphi$
Proof: $\varphi, \neg\varphi \vdash \boldsymbol{F}$.

Soundness and Completeness

Theorem (Soundness)

If Σ is satisfiable (i.e. $\Sigma \not\models \boldsymbol{F}$), then it is consistent (i.e. $\Sigma \nvdash \boldsymbol{F}$).

Equivalent formulation: if $\Sigma \vdash \varphi$ then $\Sigma \models \varphi$.

Prove and discuss in class

Theorem (Gödel's Completeness Theorem)

If Σ is consistent ($\Sigma \nvdash \boldsymbol{F}$), then it has a model ($\Sigma \not\models \boldsymbol{F}$).

Equivalent formulation: if $\Sigma \models \varphi$ then $\Sigma \vdash \varphi$.

The Completeness Theorem immediately implies the Compactness Theorem why?.

Proof of Gödel's Completeness Theorem

Follow exactly the steps of the compactness theorem.

- Extend a consistent Σ to a consistent $\bar{\Sigma}$ that is complete and witness-complete
- Use the Inductive Structure of a complete and witness-complete set.

Two Lemmas

Lemma (The Deduction Lemma)

If $\Sigma, \varphi \vdash \psi$ then $\Sigma \vdash \varphi \rightarrow \psi$.

Proof: induction on the length of $\Sigma, \varphi \vdash \psi$. Note: converse is trivial.

Lemma (Excluded Middle)

If $\Sigma, \varphi \vdash \psi$ and $\Sigma, (\varphi \rightarrow \boldsymbol{F}) \vdash \psi$ then $\Sigma \vdash \psi$.

$\Sigma \vdash \varphi \rightarrow \psi$	Deduction Lemma
$\Sigma, \psi \rightarrow \boldsymbol{F} \vdash \varphi \rightarrow \boldsymbol{F}$	by $\varphi \rightarrow \psi, \psi \rightarrow \boldsymbol{F} \vdash \varphi \rightarrow \boldsymbol{F}$
$\Sigma \vdash (\varphi \rightarrow \boldsymbol{F}) \rightarrow \psi$	Deduction Lemma
$\Sigma, \psi \rightarrow \boldsymbol{F} \vdash (\varphi \rightarrow \boldsymbol{F}) \rightarrow \boldsymbol{F}$	As above
$\Sigma, \psi \rightarrow \boldsymbol{F} \vdash \boldsymbol{F}$	MP: $\varphi \rightarrow \boldsymbol{F}, (\varphi \rightarrow \boldsymbol{F}) \rightarrow \boldsymbol{F} \vdash \boldsymbol{F}$
$\Sigma \vdash (\psi \rightarrow \boldsymbol{F}) \rightarrow \boldsymbol{F}$	Deduction Lemma
$\Sigma \vdash \psi$	by A_3

Step 1: Extend Σ to a (witness-) complete $\bar{\Sigma}$

Enumerate all formulas $\varphi_1, \varphi_2, \ldots$, and define:

$$\Sigma_0 = \Sigma \qquad \Sigma_{i+1} = \begin{cases} \Sigma_i \cup \{\varphi_i\} & \text{if } \Sigma_i \cup \{\varphi_i\} \text{ is consistent} \\ \Sigma_i \cup \{\neg\varphi_i\} & \text{if } \Sigma_i \cup \{\neg\varphi_i\} \text{ is consistent} \end{cases}$$

At least one set is consistent, otherwise:
$\Sigma_i, \varphi_i \vdash \boldsymbol{F}$ and $\Sigma_i, \neg\varphi_i \vdash \boldsymbol{F}$, thus $\Sigma_i \vdash \boldsymbol{F}$ by Excluded Middle.

To make it witness-complete, add countably many fresh constants $c_1, c_2, \ldots$, and repeatedly add $\neg\varphi[c_i/x]$ to Σ whenever $\neg\forall x(\varphi) \in \Sigma$; must show that we still have a consistent set (omitted).

Step 2: Inductive Structure of a (Witness-)Complete Set

Lemma

If $\bar{\Sigma}$ is complete, witness-complete, and consistent, then:

- $\varphi \to \psi \in \bar{\Sigma}$ *iff* $\varphi \notin \bar{\Sigma}$ *or both* $\varphi, \psi \in \bar{\Sigma}$.
- $\neg\varphi \in \bar{\Sigma}$ *iff* $\varphi \notin \bar{\Sigma}$.
- $\neg\forall x(\varphi) \in \bar{\Sigma}$ *iff there exists a constant s.t.* $\neg\varphi[c/x] \in \bar{\Sigma}$.

Sketch of the Proof in class

Now we can prove Gödel's completeness theorem:

- If Σ is consistent ($\Sigma \nvdash \boldsymbol{F}$), then it has a model.

Simply construct a model of $\bar{\Sigma}$ exactly the same way as in the compactness theorem.

Discussion

- Gödel's completeness theorem is very strong: everything that is true has a syntactic proof.
- In particular, $\text{Con}(\Sigma)$ is r.e.
- If, furthermore, Σ is complete, then $\text{Con}(\Sigma)$ is decidable!
- Gödel's completeness theorem is also very weak: it does not tell us how to prove sentences that hold in a particular structure $\boldsymbol{D}$.
- Gödel's incompleteness proves that this is unavoidable, if the structure is rich enough.

Application to Decidability

Corollary

If Σ is r.e. and complete (meaning: $\Sigma \vDash \varphi$ or $\Sigma \vDash \neg\varphi$ forall φ), then $Con(\Sigma)$ is decidable.

why?

Proof: given φ, simply enumerate all theorems from Σ:

$$\Sigma \vdash \varphi_1, \varphi_2, \varphi_3, \ldots$$

Eventually, either φ or $\neg\varphi$ will appear in the list.

Example 1: total, dense linear order without fixpoint is decidable

Example 2: $\mathrm{Th}(\mathbb{N}, 0, \mathtt{succ})$ is decidable (on your homework).

Gödel's Incompleteness Theorem

- Proven by Gödel in 1931 (after his thesis).
- It says that no r.e. Σ exists that is both consistent and can prove all true sentences in $(\mathbb{N}, +, *)$.
- The proof is actually not very hard for someone who knows programming (anyone in the audience?).
- What is absolutely remarkable is that Gödel proved it before programming, and even computation, had been invented.
- Turing published his *Turing Machine* only in 1937, to explain what goes on in Gödel's proof.
- ... and 81 years later, we have Deep Learning!

Gödel's Incompleteness Theorem

Theorem

Let Σ be any r.e. set of axioms for $(\mathbb{N}, +, *)$. If Σ is consistent, then it is not complete.

What if Σ is not consistent?

In particular, there exists a sentenced A s.t. $(\mathbb{N}, +, *) \models A$ but $\Sigma \not\vdash A$.

We will prove it, by simplifying the (already simple!) proof by Arindama Singh `https://mat.iitm.ac.in/home/samy/public_html/mnl-v22-Dec2012-i3.pdf`

Computing in $(\mathbb{N}, +, *)$

Lemma

Fact: for every Turing computable function $f : \mathbb{N} \to \mathbb{N}$ there exists a sentence $\varphi(x, y)$ s.t. forall $m, n \in \mathbb{N}$, $\mathbb{N} \vDash \varphi(m, n)$ iff $f(m) = n$.

In other words, φ represents f.

The proof requires a lot of sweat, but it's not that hard.

Sketch on the next slide.

Computing in $(\mathbb{N}, +, *)$

- Express exponentiation: $\mathbb{N} \vDash \varphi(m, n, p)$ iff $p = m^n$. This is hard, lots of math. Some books give up and assume exp is given: $(\mathbb{N}, +, *, E)$.
- Encode a sequence $[n_1, n_2, \ldots, n_k]$ as powers of primes: $2^{n_1} 3^{n_2} 5^{n_3} \cdots$ You might prefer: a sequence is just bits, hence just a number.
- Encode the API: concatenate, get i'th position, check membership.
- For any Turing Machine T, write a sentence $\varphi_T(x, y, z)$ that says[6]: "the sequence of tape contents z is a correct computation of output y from input x."
- The function computed by T is $\exists z(\varphi_T(x, y, z))$.

[6] We will do this in detail in Unit 3.

The Checker and the Prover

Fix an r.e. set of axioms[7], $(\mathbb{N}, +, *) \models \Sigma$. Construct two sentences s.t.:

- $(\mathbb{N}, +, *) \models \text{Checker}(x, y, z)$ iff
  - x encodes a formula φ,
  - y encodes a sequence $[\varphi_1, \varphi_2, \ldots, \varphi_k]$,
  - z encodes a finite set Σ_{fin}, and
  - $[\varphi_1, \varphi_2, \ldots, \varphi_k]$ is proof of $\Sigma_{\text{fin}} \vdash \varphi$.
- $\text{Prover}_\Sigma(x) \equiv \exists y \exists z(\text{"}z \text{ encodes } \Sigma_{\text{fin}} \subseteq \Sigma\text{"} \wedge \text{Checker}(x, y, z))$.
  Here we assume Σ is r.e.

By Soundness, $(\mathbb{N}, +, *) \models \text{Prover}_\Sigma(\varphi)$ implies $\Sigma \vdash \varphi$.

[7] E.g. Endetron pp. 203 describes 11 axioms

Gödel's Sentence

- Let $\varphi_1(x), \varphi_2(x), \ldots$ be an enumeration[8] of all formulas with one free variable.
- Consider the formula $\neg\text{Prover}_\Sigma(\varphi_x(x))$ this requires some thinking!
- It has a single variable x, hence it is in the list, say on position k: $\varphi_k(x) \equiv \neg\text{Prover}_\Sigma(\varphi_x(x))$.
- Denote $\alpha \equiv \varphi_k(k)$.
- In other words: $\alpha \equiv \neg\text{Prover}_\Sigma(\alpha)$ (syntactic identity)
- α says "I am not provable"!
- Next: prove two lemmas which imply Gödel's theorem.

[8] Computable!

Lemma 1

$\alpha \equiv \neg\text{Prover}_\Sigma(\alpha)$ (syntactic identity)

Lemma (1)

$\Sigma \vdash \mathit{Prover}_\Sigma(\alpha) \rightarrow \mathit{Prover}_\Sigma(\neg\alpha)$

Proof. Assume Σ is rich enough to prove:

$$P_1 : \Sigma \vdash \varphi \text{ implies } \Sigma \vdash \text{Prover}_\Sigma(\varphi)$$
$$P_2 : \Sigma \vdash (\text{Prover}_\Sigma(\varphi \rightarrow \psi)) \rightarrow (\text{Prover}_\Sigma(\varphi) \rightarrow \text{Prover}_\Sigma(\psi))$$
$$P_3 : \Sigma \vdash \text{Prover}_\Sigma(\varphi) \rightarrow \text{Prover}_\Sigma(\text{Prover}_\Sigma(\varphi))$$

The lemma follows from the last two lines:

$\vdash \neg\neg\text{Prover}_\Sigma(\alpha) \rightarrow \neg\alpha$	by $\varphi \rightarrow \varphi$
$\vdash \text{Prover}_\Sigma(\alpha) \rightarrow \neg\alpha$	$\psi \rightarrow \neg\neg\psi$
$\Sigma \vdash \text{Prover}_\Sigma(\text{Prover}_\Sigma(\alpha) \rightarrow \neg\alpha)$	P_1
$\Sigma \vdash \text{Prover}_\Sigma(\text{Prover}_\Sigma(\alpha)) \rightarrow \text{Prover}_\Sigma(\neg\alpha)$	P_2
$\Sigma \vdash \text{Prover}_\Sigma(\alpha) \rightarrow \text{Prover}_\Sigma(\text{Prover}_\Sigma(\alpha))$	P_3

Lemma 2

$\alpha \equiv \neg\text{Prover}_\Sigma(\alpha)$ (syntax) $\Sigma \vdash \text{Prover}_\Sigma(\alpha) \to \text{Prover}_\Sigma(\neg\alpha)$ (Lemma 1)

Lemma (2)

$\Sigma \vdash \textit{Prover}_\Sigma(\alpha) \to \textit{Prover}_\Sigma(\boldsymbol{F})$

Assume Σ is rich enough to also prove:

$$P_4 : \Sigma \vdash \text{Prover}_\Sigma(\varphi) \wedge \text{Prover}_\Sigma(\psi) \to \text{Prover}_\Sigma(\varphi \wedge \psi)$$

Lemma 2 follows from the last line:

$\Sigma, \text{Prover}_\Sigma(\alpha) \vdash \text{Prover}_\Sigma(\neg\alpha)$	Lemma 1 and Deduction Lemma
$\Sigma, \text{Prover}_\Sigma(\alpha) \vdash \text{Prover}_\Sigma(\neg\alpha \wedge \alpha)$	P_4
$\Sigma, \text{Prover}_\Sigma(\alpha) \vdash \text{Prover}_\Sigma(\boldsymbol{F})$	$\neg\alpha \wedge \alpha \to \boldsymbol{F}$

Proof of Gödel's First Incompleteness Theorems

$\alpha \equiv \neg\text{Prover}_\Sigma(\alpha)$ (syntax) $\Sigma \vdash \text{Prover}_\Sigma(\alpha) \rightarrow \text{Prover}_\Sigma(\boldsymbol{F})$ (Lemma 2)

Theorem (Σ Is Not Complete)

If Σ is consistent ($\Sigma \nvdash \boldsymbol{F}$), then $\Sigma \nvdash \alpha$ and $\Sigma \nvdash \neg\alpha$.

Proof:

Suppose $\Sigma \vdash \alpha$:

$\Sigma \vdash \text{Prover}_\Sigma(\alpha)$	P_1
$\Sigma \vdash \neg\text{Prover}_\Sigma(\alpha)$	syntax
$\Sigma \vdash \boldsymbol{F}$	$\varphi, \neg\varphi \vdash \boldsymbol{F}$

Suppose $\Sigma \vdash \neg\alpha$:

$\Sigma \vdash \neg\neg\text{Prover}_\Sigma(\alpha)$	syntax
$\Sigma \vdash \text{Prover}_\Sigma(\alpha)$	A_3
$\Sigma \vdash \text{Prover}_\Sigma(\boldsymbol{F})$	Lemma 2
$\Sigma \vdash \boldsymbol{F}$	soundness

Proof of Gödel's Second Incompleteness Theorems

$\alpha \equiv \neg \text{Prover}_\Sigma(\alpha)$ (syntax) $\quad\quad \Sigma \vdash \text{Prover}_\Sigma(\alpha) \rightarrow \text{Prover}_\Sigma(\boldsymbol{F})$ (Lemma 2)

Theorem (Σ Cannot Prove its Own Consistency)

$\Sigma \not\vdash \neg \mathit{Prover}_\Sigma(\boldsymbol{F})$

Proof: suppose $\Sigma \vdash \neg \text{Prover}_\Sigma(\boldsymbol{F})$

$\Sigma \vdash \neg \text{Prover}_\Sigma(\boldsymbol{F}) \rightarrow \neg \text{Prover}_\Sigma(\alpha)$	Lemma 2
$\Sigma \vdash \neg \text{Prover}_\Sigma(\alpha)$	Modus Ponens
$\Sigma \vdash \alpha$	Syntax
$\Sigma \vdash \boldsymbol{F}$	First Incompleteness Theorem

Discussion

- We only proved that neither α nor $\neg\alpha$ is provable. Can we get a complete theory by adding α or $\neg\alpha$ to Σ (whichever is true)? In class
- Not all theories of $\mathbb{N}$ are undecidable. Examples[9]:
  - $(\mathbb{N}, 0, \mathtt{succ})$ is decidable (homework!).
  - $(\mathbb{N}, 0, \mathtt{succ}, <)$ is decidable; can define finite and co-finite sets.
  - $(\mathbb{N}, 0, \mathtt{succ}, +, <)$ is decidable and called Presburger Arithmetic; can define eventually periodic sets.
  - $(\mathbb{N}, 0, \mathtt{succ}, +, *, <)$ is undecidable (Gödel).
  - $(\mathbb{C}, 0, 1, +, *)$ is decidable.

[9] Enderton pp. 187, 197, 158