CDM

FOL Theories

Klaus Sutner

Carnegie Mellon University

42-fol-theories 2017/12/15 23:21

cc BY NC SA

1. Theories and Models
    - Decidability and Completeness
    - Derivations and Proofs
    - Compactness
    - Incompleteness and Undecidability

Structures and Models

Recall from last time that we have a notion of a structure being a model of a sentence, or a whole set of sentences.

$$\mathcal{A} \models \Gamma$$

To really make sense out of this, we need some background theory that allows us to talk about the corresponding first-order structures. Typically this is handled via Zermelo-Fraenkel set theory.

Definition

Given a set Γ of sentences and a sentence φ, we say that φ is valid in Γ or a semantic consequence of Γ if any structure $\mathcal{A}$ that satisfies all formulae in Γ also satisfies φ: $\mathcal{A} \models \Gamma$ implies $\mathcal{A} \models \varphi$.

Notation:

$$\Gamma \models \varphi$$

For example, for a ground term t

$$\varphi(t) \models \exists\, x\, \varphi(x)$$

$$\varphi, \varphi \to \psi \models \psi$$

Less Obvious

The last two examples are purely logical, but semantic entailment is not always so easy to see.

For example, the standard group axioms

$$
\begin{aligned}
&x * (y * z) = (x * y) * z \\
&x * e = e * x = x \\
&x * x^{-1} = x^{-1} * x = e
\end{aligned}
$$

have the semantic consequence

$$(x * y)^{-1} = y^{-1} * x^{-1}$$

The insistence on sentences (formulae without free variables) should be taken with a grain of salt, one can deal with the general situation at some modest cost.

Theories

Fix some first-order language L once and for all. As always, we are only interested in languages with a finite or at most countable supply of non-logical symbols, and decidable syntax.

Definition

A theory (over L) is an arbitrary collection of sentences (over L).

A theory T is satisfiable if there is some structure $\mathfrak{A}$ such that $\mathfrak{A} \models T$.

Warning: This is the skinny definition of theory, some authors insist that a theory be closed under semantic consequence: $\Gamma \models \varphi$ implies $\varphi \in \Gamma$.

Satisfiability

Obviously, theories that have no models at all are of little interest, they are self-contradictory. Alas, it is surprisingly difficult to avoid contradictions when one constructs some logical system (not necessarily first-order logic).

For Hilbert, satisfiability was the key to existence:

> If the arbitrarily given axioms do not contradict each other through their consequences, then they are true, then the objects defined through the axioms exist. That, for me, is the criterion of truth and existence.

Constructivists recoil at this idea, but is hard to see how one could convince mainstream mathematicians to let go of Hilbert's dream.

Example: Monoid Axioms

Here is a closer look at a standard example from algebra: axioms for semigroups, monoids and groups. Initially, let's use the language

$$\mathcal{L}(*) \quad \text{of type} \quad (2)$$

So there is only one binary function symbol $*$, and no constants. We will write $*$ in infix notation to keep things readable and drop the universal quantifiers up front.

Definition

A semigroup is any structure satisfying the following associativity axiom:

$$x * (y * z) \approx (x * y) * z \tag{1}$$

If, in addition, there is a neutral element

$$\exists\, y \,\forall\, x\, (x * y \approx x \;\wedge\; y * x \approx x) \tag{2}$$

then we have a monoid.

A Better Language

The quantifiers in the second monoid axiom are unnecessarily complicated, we can change our language to

$$\mathcal{L}(*, e) \text{ of type } (2, 0)$$

and replace the second axiom by

$$x * e \approx x \;\wedge\; e * x \approx x \tag{3}$$

Much easier to read.

Example Monoids

In order to specify a (small) monoid, we only need to write down a Cayley table (multiplication table) for the operation, and the neutral element.

Here are the Cayley tables of two simple examples that we are already familiar with:

$*$	0	1
0	0	1
1	1	1

$*$	0	1
0	0	0
1	0	1

This is just logical disjunction and conjunction and usually written as $\langle \mathbb{B}; \vee, 0 \rangle$ and $\langle \mathbb{B}; \wedge, 1 \rangle$.

Note that it is easy to determine the whole atomic diagram from the Cayley tables. E.g., $(1 * 1) * (1 * 1) = (1 * 1)$ is in the diagram.

Six Elements

A slightly larger Cayley table with elements $\{1, 2, 3, 4, 5, 6\}$.

$*$	1	2	3	4	5	6
1	1	2	3	4	5	6
2	2	4	6	1	3	5
3	3	6	2	5	1	4
4	4	1	5	2	6	3
5	5	3	1	6	4	2
6	6	5	4	3	2	1

Clearly we can test in time cubic in the size of the structure whether is satisfies the associativity axiom. To check for the neutral element only requires linear time.

Alas, a Cayley table obfuscates the structure of the monoid, there are other descriptions that are much more useful.

Some Infinite Monoids

For infinite carrier sets things become more interesting. Again, we have to specify the carrier set, the operation $*$, and the constant e.

Here are some monoids:

$$\langle\, \mathbb{N}; +, 0 \,\rangle$$
$$\langle\, \mathbb{N}^+; *, 1 \,\rangle$$
$$\langle\, \mathbb{R}^+; *, 1 \,\rangle$$
$$\langle\ \text{strings};\ \text{concat},\ \text{empty word}\ \rangle$$
$$\langle\ \text{lists};\ \text{join},\ \text{empty list}\ \rangle$$

Also note that it is not at all clear how one would go about verifying that these are indeed models of the monoid axioms: a brute force test is not possible here. We have to argue within a more powerful system.

Two Strange Monoids

Here is a less obvious example. Let $\mathcal{F}$ be the collection of propositional formulae modulo equivalence. Then the following two structures are monoids.

$$\langle\, \mathcal{F};\ \wedge,\ \top\,\rangle$$
$$\langle\, \mathcal{F};\ \vee,\ \bot\,\rangle$$

Note that for this to work we must identify equivalent formulae: $p \wedge \top \neq p$ but certainly $p \wedge \top \equiv p$.

Exercise

What is the size of these monoids if we consider only formulae with n propositional variables?

Non-Models

Not just any old set with a binary operation is a monoid.

$$\langle\, \mathbb{N}^+ ; + \,\rangle$$
$$\langle\, \mathbb{R}^+ ; \exp \,\rangle$$
$$\langle\, \text{full binary trees} \;;\; \text{join} \,\rangle$$

In the first example there is no neutral element. The second operation, exponentiation, is not associative. Likewise, joining together full binary trees is not associative (figure out what exactly this means).

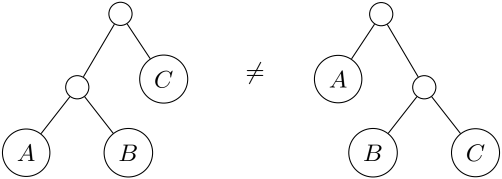

Example: Groups 1

To hammer this home, here is another standard example from algebra: groups. First, we use the language

$$\mathcal{L}(*) \quad \text{of signature} \quad (2)$$

of one binary operation. Groups can then be specified as monoids that have inverse elements.

$$x * (y * z) \approx (x * y) * z$$

$$\exists\, z\, \big(\forall\, x\, (x * z \approx x \ \wedge\ z * x \approx x) \wedge$$
$$\forall\, x\, \exists\, y\, (x * y \approx z \wedge y * x \approx z)\big)$$

Note that our minimalist language makes the formula for the the existence of a neutral element and corresponding inverses a bit complicated.

Example: Groups 2

It is better to have a language with a constant for the neutral element:

$$\mathcal{L}(*, e) \ \ \text{of signature} \ \ (2, 0)$$

The theory of groups then looks like

$$x * (y * z) \approx (x * y) * z$$

$$x * e \approx x \ \wedge \ e * x \approx x$$

$$\exists\, y \,(x * y \approx e \wedge y * x \approx e)$$

We are down to one quantifier for the existence of an inverse element.

Example: Groups 3

If we extend our language one more time by a unary function symbol for inverses we get an even clearer description over

$$\mathcal{L}(*, {}^{-1}, e) \text{ of signature } (2,1,0)$$

Group theory now is

$$x * (y * z) \approx (x * y) * z$$

$$x * e \approx x \ \wedge \ e * x \approx x$$

$$x * x^{-1} \approx e \wedge x^{-1} * x \approx e$$

This is arguable a better representation: all the formulae are now universal; in fact, we have a purely equational characterization.

Example: Groups 3.5

The last step is to give up on the pedantic $\approx$ and write $=$ instead, and to allow for chains of identities.

$$x * (y * z) = (x * y) * z$$

$$x * e = e * x = x$$

$$x * x^{-1} = x^{-1} * x = e$$

This is the form that is familiar from any (modern) algebra textbook. But note that the models are essentially the same as for the first form.

Equivalence

Definition

Two theories T_1 and T_2 are equivalent if they have the same models.

Note that T_1 and T_2 have to be over the same language (essentially, extensions are fine).

All the examples for groups can be considered as theories over $\mathcal{L}(*, {}^{-1}, e)$.

One can check that they are all equivalent.

Definition

Given a class $\mathfrak{C}$ of structures we define the theory of $\mathfrak{C}$ to be the collection of sentences valid in all the structures in $\mathfrak{C}$:

$$\mathrm{Th}(\mathfrak{C}) = \{\, \varphi \mid \text{ for all } \mathcal{A} \in \mathfrak{C} : \mathcal{A} \models \varphi \,\}$$

If there is only one structure $\mathcal{A}$ in the class we write $\mathrm{Th}(\mathcal{A})$.

For example, $\mathrm{Th}(\mathcal{N})$ is the theory of arithmetic.

$\mathrm{Th}(\mathrm{Grp})$ is the theory of groups, $\mathrm{Th}(\mathrm{BoolAlg})$ is the theory of Boolean algebras and so on.

This is quite different from trying to understand the theory of a specific group, or a specific Boolean algebra.

Definition

Let T be a theory. The semantic closure of T is defined by

$$\mathrm{Th}(T) = \{\, \varphi \mid T \models \varphi \,\}$$

In other words, if $\mathfrak{M}$ is the class of all models of T, then $\mathrm{Th}(T) = \mathrm{Th}(\mathfrak{M})$.

It is entirely possible that $\mathrm{Th}(T)$ is much more complicated than T itself: in general, we know little about the class of all possible models.

Models

Definition

For any theory T, the class of models of T is defined to be

$$\mathsf{Mod}(T) = \{\, \mathcal{A} \mid \mathcal{A} \models T \,\}.$$

where all the structures have the appropriate signature.

It follows that $\mathrm{Th}(T) = \mathrm{Th}(\mathsf{Mod}(T))$.

For example, the collection of all groups is described concisely by the theories from above.

Similarly we can describe Abelian groups or infinite groups (warning: this does not work for the collection of all finite groups).

It has become one of the standard methods in math and TCS to describe (classes of) structures $\mathfrak{C}$ by selecting a theory Γ such that $\mathfrak{C} = \mathsf{Mod}(\Gamma)$.

Definition

Let T be any theory. A theory Γ is a set of axioms for T if T and Γ are equivalent.

A theory Γ is a set of axioms for $\mathfrak{C}$ if T is equivalent to $\mathrm{Th}(\mathfrak{C})$.

This is a white lie: axioms are supposed to be few in number and self-evident, they should encapsulate the fundamental truths of the domain of discussion, and there should be no point in trying to analyze them any further (fat chance).

Elegance matters greatly in the selection of axioms.

Axioms versus the RealWorld

Axiom systems can have two entirely different purposes:

- One may try to describe a large class of models; e.g., all groups or all fields of characteristic 2.
- One may try to pin down one particular structure; e.g., the natural numbers, the rationals or the reals.

As we will see, FOL is not particularly good at the second task, one often has to contend with unintended models. These bad models are not due to a poor choice of axioms, they just cannot be avoided in FOL.

Independence

There is a tacit requirement for axioms: they should all be logically independent, no one axiom should follow from the others.

This is a clean technical condition, though not always easy to check: one has to show that dropping any one axiom produces a larger class of models.

Still, famous and well-studied axiom systems like Zermelo-Fraenkel set theory or Peano arithmetic are independent in this sense.

The question arises whether axioms can always be kept simple in a technical sense. Let T be any consistent theory.

Definition

T is finitely axiomatizable if it admits a finite set of axioms.

T is recursively axiomatizable if it admits a decidable set of axioms.

We use the same terminology for classes of structures. In addition, $\mathfrak{C}$ is axiomatizable if there is some set of axioms for it.

Theories that fail to be recursively axiomatizable are essentially too complicated for FOL (see below for proofs in FOL).

Also note that a theory is always axiomatizable, but a class of structures may not be.

Finitely Axiomatizable

- standard structures from algebra
- ordered fields
- lattices
- Boolean algebras
- Neumann-Bernays-Gödel set theory

Recursively Axiomatizable

- torsion-free groups
- fields of characteristic 0
- algebraically closed fields
- finite fields
- Peano arithmetic
- Zermelo-Fraenkel set theory

Alas, not all reasonable classes $\mathfrak{C}$ are axiomatizable, no matter how complicated the axioms. One needs to resort to other logics to capture these classes.

Example

The collection of finite groups FinGrp is not axiomatizable: statements that hold for arbitrarily large groups also hold for some infinite groups.

Example

The collection of all connected graphs ConGrf is not axiomatizable: there is no way to express the existence of a path of unbounded length in FOL.

Example

The class of all well-orderings is not axiomatizable: we cannot quantify over sets or sequences of individuals.

Example: Fields

Rings and fields are finitely axiomatizable.

Here are axioms for fields in the language $\mathcal{L}(+, *, -, ^{-1}, 0, 1)$ of type $(2, 2, 1, 1, 0, 0)$.

$$\begin{array}{cc}
x + (y + z) \approx (x + y) + z & x * (y * z) \approx (x * y) * z \\
x + y \approx y + x & x * y \approx y * x \\
x + 0 \approx x & x * 1 \approx x \\
x + (-x) \approx 0 & x \not\approx 0 \rightarrow x * x^{-1} \approx 1 \\
x * (y + z) \approx x * y + x * z & 0 \not\approx 1
\end{array}$$

Exercise

What would these axioms look like using the language $\mathcal{L}(+, *)$ *of type* $(2, 2)$*?*

Characteristic

Write $\underline{n}$ for the ground term $\underbrace{1 + 1 + \ldots + 1}_{n}$, $n \geq 1$.

Consider the sentences $\chi_n = \underline{n} \approx 0$, $n \geq 2$.

Adding a single axiom χ_n, n prime, to the field axioms produces axioms for fields of characteristic n.

Adding all axioms $\neg\chi$ produces axioms for fields of characteristic 0.

Exercise

What happens if we adjoin an axiom χ_n for n composite?

Exercise

How can one axiomatize the theory of algebraically closed fields?

Example: Boolean Algebras

Boolean algebras are also finitely axiomatizable.

Here the language is $\mathcal{L}(\sqcup, \sqcap, ', 0, 1)$ of signature $(2, 2, 1, 0, 0)$.

$$\begin{aligned} x \sqcup (y \sqcup z) &= (x \sqcup y) \sqcup z & \qquad x \sqcap (y \sqcap z) &= (x \sqcap y) \sqcap z \\ x \sqcup y &= y \sqcup x & x \sqcap y &= y \sqcap x \\ x \sqcup 0 &= x & x \sqcap 1 &= x \\ x \sqcup (y \sqcap z) &= (x \sqcup y) \sqcap (x \sqcup z) & x \sqcap (y \sqcup z) &= (x \sqcap y) \sqcup (x \sqcap z) \\ x \sqcup x' &= 1 & x \sqcap x' &= 0 \end{aligned}$$

Exercise

What would these axioms look like we adopted as primitive notion a partial order $x \leq y \iff x \sqcap y = x$*?*

Example: Linear Orders

Yet another finitely axiomatizable class.

Here the language is $\mathcal{L}(<)$ of signature (2). For a strict total order we need transitivity, trichotomy and anti-symmetry.

$$x < y \land y < z \to x < z$$
$$x < y \lor x \approx y \lor y < x$$
$$x < y \to \neg y < x$$

To get a dense order we add

$$x < y \to \exists\, z\, (x < z < y)$$

Axioms versus Schemata

When one constructs a logic, one has to strike a balance with a view towards derivations in the system:

- Use many axioms, keep inference rules simple.
- Have few axioms, but strong inference rules.

In particular when the number of axioms is infinite, one typically introduces them via an axiom schema. For example, we may admit all formulae

$$\varphi \to \psi \to \varphi$$

as axioms. Here φ and ψ are meta-variables that range over formulae. Alternatively, we could have just one axiom

$$p \to q \to p$$

and allow substitutions to obtain others of the same form.

Digression: Elementary Classes

Unfortunately, there is a parallel terminology:

Definition

A class of structures $\mathfrak{C}$ is elementary if there is a theory T such that $\mathfrak{C} = \mathsf{Mod}(T)$.

A class of structures $\mathfrak{C}$ is basic elementary if there is a finite theory T such that $\mathfrak{C} = \mathsf{Mod}(T)$.

On top, some authors use “elementary” in place of “basic elementary.”

We will avoid this terminology and stick with axiomatizable, finitely axiomatizable and recursively axiomatizable.

- Theories and Models

2. Decidability and Completeness

- Derivations and Proofs

- Compactness

- Incompleteness and Undecidability

As before, we fix some (at most countable, decidable syntax) first-order language L once and for all.

Definition

A theory T is decidable if $\mathrm{Th}(T)$ is decidable.

Terminology Warning: a set T of axioms may be decidable as a set of sentences, but $\mathrm{Th}(T)$ may still be undecidable.

Example: Arithmetic

As an example of a highly undecidable theory, consider

$$T = \text{ all true statements of arithmetic}$$

in the language $\mathcal{L}(+, *, 0, 1; <)$ of signature $(2, 2, 0, 0; 2)$. In fact, not even the universal sentences in T (such as the Riemann hypothesis) admit a decision procedure.

Peano arithmetic (PA) is a small, sub-theory of T that has a recursive set of axioms. Alas, (PA) is still undecidable and thoroughly fails to axiomatize T (through no fault of Peano's, FOL is too weak).

Decidable Theories

- $\mathrm{Th}(f, \approx)$
- Presburger arithmetic
- Skolem arithmetic
- Abelian groups
- Boolean algebras
- Algebraically closed fields
- Real closed fields

Undecidable Theories

- $\mathrm{Th}(f, g, \approx)$
- $\mathrm{Th}(R^2, \approx)$
- Arithmetic (plus and times)
- Robinson arithmetic
- Algebraic structures (semigroups, groups, finite groups, rings, fields)

Since axiom systems are typically small the problem arises whether they contain enough information to really pin down the objects under discussion. Ideally we would like the axioms to settle all possible questions (at least questions that can be phrased in the particular language we have chosen).

Definition

A set of sentences Γ is complete if for every sentence φ of the language we have either $\Gamma \models \varphi$ or $\Gamma \models \neg\varphi$.

A cheap example for a complete theory is $\mathrm{Th}(\mathfrak{C})$.

Completeness is a rather strong property; for example, Γ has to pin down statements about cardinality. Most standard axiom systems (such as the group axioms) are indeed incomplete.

Example

The theory of algebraically closed fields of characteristic 0 is complete, as is the theory of real closed fields.

Another way of thinking about completeness is to consider the models, they are all very similar in the following sense:

Definition

Two structures $\mathcal{A}$ and $\mathcal{B}$ are elementarily equivalent if $\mathrm{Th}(\mathcal{A}) = \mathrm{Th}(\mathcal{B})$.

As always, we assume that the structures have the same signature.

Thus, as long as we restrict ourselves to the limitations of FOL we cannot distinguish between the two structures. But note that this notion is much weaker than the existence of an isomorphism between the structures.

Proposition

Γ is complete if, and only if, any two models of Γ are elementarily equivalent.

Complete Theories

- Presburger arithmetic.
- Dense linear orders.
- Algebraically closed fields of a given characteristic.
- Real closed fields.
- Tarski's axioms for Euclidean geometry
- Any $\aleph_0$-categorical theory.

- Theories and Models

- Decidability and Completeness

3. Derivations and Proofs

- Compactness

- Incompleteness and Undecidability

Derivations

To really get mileage out of FOL, though, we need a better grip on $\mathrm{Th}(\Gamma)$: we need to be able to find the consequences φ of Γ directly, without complete knowledge of all the models.

We would like to argue solely from the axioms themselves, without any reference to the (possibly complicated) assortment of models.

For example, we know informally how to derive the assertion

$$\forall\, x, y\, ((x * y)^{-1} = y^{-1} * x^{-1})$$

directly from the group axioms, using some basic equational reasoning and "obvious" rules for quantification.

How do we formalize these rules?

How do we construct proofs in FOL?

We need to come up with a notion of proof or derivation in FOL that allows us to draw inferences from a set of given formulae. The goal is to develop a notion of provability

$$\Gamma \vdash \varphi$$

analogous to provability in propositional logic. There are many different ways of doing this, for the sake of brevity let us only indicate how the natural deduction system from the lecture on propositional logic can be expanded to FOL.

Needless to say, we can retain the propositional rules, they are still sound in the new context.

To extend natural deduction from propositional logic to FOL, we need to add rules for quantifiers. Intuitively, that's not hard.

$$\frac{\phi(t)}{\exists\, x\, \phi(x)}\ (\exists e) \qquad\qquad \frac{\exists\, x\, \phi(x)}{\phi(c)}\ (\exists i)$$

$$\frac{\phi(c)}{\forall\, x\, \phi(x)}\ (\forall i) \qquad\qquad \frac{\forall\, x\, \phi(x)}{\phi(t)}\ (\forall e)$$

where x is a variable, c a constant, and t a term.

Alas, while these rules are correct in spirit, as stated they are not sound.

Counterexample

Suppose we adopt the unconstrained quantifier rules from above. Then we can perform a derivation along the following lines (this is Hilbert style).

$\forall x \, \exists y \, (x < y)$	premiss
$\exists y \, (c < y)$	$(\forall e)$
$(c < d)$	$(\exists e)$
$\forall x \, (x < d)$	$(\forall i)$
$\exists y \, \forall x \, (x < y)$	$(\exists i)$

Disaster! This is exactly the wrong direction of the valid implication $\exists x \, \forall y \, \varphi(x, y) \to \forall y \, \exists x \, \varphi(x, y)$. The problem is that the "constant" d really depends on c.

Counterexample, contd.

To address this and similar problems, one has to add some technical conditions to the quantifier rules that rule out the derivation from above.

- For $(\forall i)$ we insist that c is “fresh”: it must no occur in any undischarged assumptions in the derivation of $\varphi(c)$ nor in the conclusion $\forall\, x\, \phi(x)$ itself.

- For $(\exists e)$ again we insist on a fresh constant c:

$$\dfrac{\exists\, x\, \phi(x) \qquad \begin{matrix}[\phi(c)]\\ \vdots \\ \chi\end{matrix}}{\chi}\ (\exists e)$$

- For $(\forall e)$ and $(\exists i)$ we insist that term t is substitutable for x in φ: no variable in t becomes bound as a result of the replacement process.

Syntactic Consequence

We won't belabor the details, though they are extremely important when one tries to actually implement a deduction system in FOL. What we need here is the ability to derive consequences from a given collection of formulae, in a purely syntactic fashion.

Definition

The set of all provable theorems of a collection of sentences in FOL is called its (syntactic) theory or its set of consequences.

Notation:

$$\mathsf{Cn}(\Gamma) = \{\, \varphi \mid \Gamma \vdash \varphi \,\}$$

Note that on the face of it, $\mathsf{Cn}(\Gamma)$ is a much less complicated notion than $\mathrm{Th}(\Gamma)$.

Proofs versus Truth

As with propositional logic the central problem is that we need

- Soundness: $\Gamma \vdash \varphi \Rightarrow \Gamma \models \varphi$, and
- Completeness: $\Gamma \models \varphi \Rightarrow \Gamma \vdash \varphi$.

at the same time.

Note that two worlds interact here: proofs are merely syntactic objects (data structures) whereas truth and semantic consequence depend on actual structures. There can be very many such structures and they are possibly infinite, so it is difficult to get a clear understanding of what they look like.

As a consequence, it is far from clear that there even exists a set of non-logical axioms, which, together with the logical axioms and the rules of inference, would achieve this goal.

Proofs in FOL

The first proof of completeness is due to Gödel in 1930.

Theorem (Completeness Theorem)

There is a formal system for FOL that is sound and complete.

Deduction theorem

As far as implications are concerned, FOL behaves just like propositional logic.

Theorem (Deduction Theorem)

Let Γ, φ, ψ *be sentences. Then* $\Gamma, \varphi \vdash \psi$ *if, and only if,* $\Gamma \vdash \varphi \to \psi$*.*

But note the condition on sentences here, free variables cause problems. For example, $P(x) \vdash P(y)$ but $\vdash P(x) \to P(y)$ is false.

So if there is a proof $\Gamma \vdash \varphi$ between sentences, then

$$\vdash \psi_1 \wedge \psi_2 \ldots \wedge \psi_n \to \varphi$$

for some ψ_i in Γ.

Recall that a theory Γ is consistent if it has at least one model. Could it ever happen that some set of axioms Γ has no models at all? Sure, $\Gamma = \{a \approx b, b \approx c, a \not\approx c\}$ or, more radically, $\Gamma = \{\bot\}$ will do.

Of course, no one would use these directly as axioms. But, it might happen that $\bot$ is provable from an ill-chosen set of axioms: there will be no models, but that may not at all be obvious from the axioms.

Definition

Γ is inconsistent if $\Gamma \vdash \bot$, and consistent otherwise.

Note that in an inconsistent theory all sentences are provable.

Clearly, any inconsistent set of axioms has no models; the properties we were trying to pin down contradict each other.

Surprisingly, this is the only thing that can go wrong.

Theorem (Completeness Theorem I)

A sentence of a theory T is a theorem of T if, and only if, it is valid in T.

Theorem (Completeness Theorem II)

A theory has a model if, and only if, it is consistent.

In symbols:

$$T \models \varphi \ \text{ iff } \ T \vdash \varphi$$

$$\mathrm{Th}(T) = \mathsf{Cn}(T)$$

Proof Sketch

It is not hard to see that the second version implies the first:

$$\begin{aligned} T \models \varphi &\Leftrightarrow T + \neg\varphi \text{ inconsistent} \\ &\Leftrightarrow T + \neg\varphi \text{ not satisfiable} \qquad \text{CT II} \\ &\Leftrightarrow \text{for all } \mathcal{A} : \mathcal{A} \models T \text{ implies } \mathcal{A} \models \varphi \\ &\Leftrightarrow T \models \varphi \end{aligned}$$

Exercise

Show the CT I also implies CT II.

Henkin's Argument

It remains to establish CT II.

Without going into details, the idea of a proof by L. Henkin is to

- add many constants c to the language, and
- extend T to a certain complete set Γ of sentences.

The constants are used to make sure that $\exists\, x\, \varphi(x) \in \Gamma$ implies $\varphi(c) \in \Gamma$ for some c.

Then we can construct a term model for Γ: the collection of all terms, modulo equality provable in Γ, turns out to be a model of Γ and thus of T.

Enforcing Completeness

The fact that consistency is the only requirement for the existence of a complete extension of an axiom system is really a separate theorem by Tarski.

Theorem (Tarski)

Every consistent set of sentences T can be extended to a complete set of sentences $T' \supseteq T$.

Of course, T' might be very complicated. For example, we might lose decidability.

Exercise

What might a complete extension of the theory of fields look like? Note that it has to settle on a value for 0^{-1}.

- Theories and Models

- Decidability and Completeness

- Derivations and Proofs

4. Compactness

- Incompleteness and Undecidability

Another crucial property of first-order theories is based on the following simple observation: any single derivation in FOL uses only finitely many formulae. One can certainly imagine stronger systems where a single proof involves infinitely many formulae, but in FOL this is not the case.

Hence, Γ inconsistent means that there is a finite subset Γ_0 of Γ that is already inconsistent. By completeness we get the following surprising consequence:

Theorem (Compactness theorem)

Γ *has a model if, and only if, every finite subset* Γ_0 *of* Γ *has a model.*

This is positively wild, because infinitely many axioms can be used to construct weird conditions, when every finite subset is perfectly harmless.

Lemma

The class of all finite structures (of some signature) is not axiomatizable.

For assume that Γ is some axiom system that characterizes finite structures. Hence Γ has arbitrarily large finite models. But then Γ must have an infinite model.

To see this, add new constants c_i, $i \in \mathbb{N}$, to the language and add new axioms

$$c_i \not\approx c_j \qquad \text{for } i < j$$

to Γ to obtain an extension Γ'. Every finite subset of Γ' has a model; by compactness Γ' has a model which must be infinite by choice of the additional axioms.

Exercise

Use a similar argument to show that there is an infinite field of characteristic 2. By contrast, give an algebraic construction of such a field.

Lemma

The class of all infinite structures (of some signature) is not finitely axiomatizable.

For otherwise the class of finite structures would also be elementary.

But note that the class of infinite structures is FO definable: we need infinitely many sentences of the form “there are at least n elements”.

$$\mathsf{EX}_{\geq n} = \exists\, x_1, \ldots, x_n \,(x_1 \neq x_2 \land x_1 \neq x_3 \land \ldots \land x_{n-1} \neq x_n)$$

Exercise

Fill in the details in the proofs above.

Exercise

Write down axioms for fields.

Exercise

Write down axioms for finite fields, for infinite fields and for fields of characteristic p (both for $p = 0$ and p prime).

Exercise

How would you axiomatize vector spaces where one has to deal with a field and the actual vector space at the same time?

Exercise

Axiomatize the real numbers as best you can, starting from the field axioms. Repeat for complex numbers.

Exercise

Try to write down axioms for a stack of elements of some ground type. Repeat for queues, lists and trees.

Peano Arithmetic (PA)

Here is an example of a hugely important axiom system due to G. Peano in 1889 that describes the salient properties of the structure of natural numbers $\mathcal{N} = \langle \mathbb{N}; +, \cdot, S, 0, 1, < \rangle$, signature $(2, 2, 1, 0, 0; 2)$, with the usual arithmetic operations and order.

Actually, R. Dedekind, Gauss's last student, developed the same system a year before Peano but never published.

Peano's Axioms

successor

$$S(x) \neq 0 \qquad S(x) \approx S(y) \rightarrow x \approx y$$

addition

$$x + 0 \approx x \qquad x + S(y) \approx S(x + y)$$

multiplication

$$x \cdot 0 \approx 0 \qquad x \cdot S(y) \approx (x \cdot y) + x$$

order

$$\neg(x < 0) \qquad x < S(y) \leftrightarrow x \approx y \vee x < y$$

Computational Aspects

The Peano axioms are over a century old. Surprisingly, they are almost like programs.

More precisely, the axioms provide primitive recursive definitions of plus, times and less-than in terms of the atomic successor function succ.

```
int add( int x, int y ) {
        if( y == 0 )  return x;
        return   succ( x, pred(y) );
}

int mult( int x, int y ) {
        if( y == 0 ) return 0;
        return  add( mult( x, pred(y) ), x );
}
```

Here pred(y) stands for the predecessor $y - 1$.

Induction

Arithmetic operations alone are not enough, we are missing one essential feature of the natural numbers: induction. To capture induction we add the Induction Axiom:

$$\varphi(0) \land \forall\, x \,\big(\varphi(x) \to \varphi(S(x))\big) \to \forall\, x\, \varphi(x)$$

Strictly speaking, this is not an axiom but an axiom schema: we get one axiom for each choice of φ.

At any rate, (PA) is a very succinct representation of the essential features of arithmetic.

When someone says “prove that every number is divisible by a prime”, they really mean: “prove in (PA) that every number is divisible by a prime”.

OK, this is a white lie: a great many mathematicians and even more computer scientists are blissfully unaware of (PA).

But the point is that even if one is not interested in formal derivations, Peano arithmetic still provides a natural framework for reasoning about the natural numbers.

There are more powerful systems such as set theory or type theory which allow us to prove more complicated facts, but when dealing with the natural numbers (PA) is almost always all you need.

Lemma

(PA) *proves that* $\forall\, x\, (0 + x \approx x)$

Proof.

Consider the formula $\varphi(x) \equiv (0 + x \approx x)$.

Then $\varphi(0)$ is the first addition axiom (more precisely, replace x by 0 there).

Now assume $\varphi(x)$. Then by the second addition axiom

$$0 + S(x) \approx S(0 + x) \approx S(x)$$

Hence we have shown $\varphi(0) \wedge \forall\, x\, (\varphi(x) \to \varphi(S(x)))$.

By the Induction Axiom and modus ponens we get $\forall\, x\, \varphi(x)$.

□

The important point in all of these exercises is to argue from the axioms, not any intuitive understanding of the structure.

Exercise

Prove the following assertions in (PA).

$$\forall\, x, y, z\ (x + (y + z) \approx (x + y) + z)$$
$$\forall\, x, y\ (x + y \approx y + x)$$

Conclude that $\langle\, \mathbb{N}; +, 0\,\rangle$ *is a commutative monoid.*

Exercise

Define the GCD and describe the Euclidean algorithm in (PA).

Exercise

Prove that there are infinitely many primes in (PA).

One could argue that some $50+$ years ago, no one cared except some mathematicians and logicians.

Some were even rather proud of their work being totally impractical (G. H. Hardy, "A Mathematician's Apology", CUP 1948).

As recent history shows, Hardy totally underestimated the impact of apparently obscure branches of mathematics (in his defense: the role of computers was not clear when he wrote the book).

Modern cryptography is unthinkable without number theory, group theory, finite field theory. Graphics without projective geometry is a sad affair. The notion of efficient algorithm is derived from recursion theory (classical theory of computation). Verification of IEEE protocols uses formal logic.

And so on.

One would hope that the Peano axioms allow one to derive all true statements of arithmetic.

But note that there is a little problem here: the true statements of arithmetic are the theory of a single structure:

$$\mathrm{Th}(\mathcal{N}) = \{\, \varphi \mid \mathcal{N} \models \varphi \,\}$$

But all we get from our notion of provability is the statements that hold in all models of (PA).

That would be fine if $\mathcal{N}$ were the only such model (up to isomorphism or even elementary equivalence). Unfortunately, there are others as we shall see shortly.

A Weird Model for (PA)

If we think of (PA) as pinning down axiomatically the properties of the natural numbers, here is some very bad news: $\mathcal{N}$ is not the only model of (PA).

To see this, introduce a new constant c and add the following new axioms to (PA):

$$0 < c, \underline{1} < c, \underline{2} < c, \ldots, \underline{n} < c, \ldots$$

Call the new set of axioms (PA^∞).

Claim

(PA^∞) *has a model.*

Proof. To see this, exploit compactness. Any finite subset Γ of (PA^∞) contains only finitely many of the new axioms. So there is a largest n such that $S^n(0) < c \in \Gamma$.

But then we can simply interpret c as $n+1$ in the standard model $\mathcal{N}$, done. □

So let $\mathcal{A}$ be a model of (PA^∞). Since $\mathcal{A}$ is in particular a model of (PA) we have a full copy of $\mathcal{N}$ inside $\mathcal{A}$: there are distinct elements for 0, $S(0)$, $S(S(0))$ and so forth.

But, this structure also contains an "infinitely large" element $c^{\mathcal{A}}$: since c is a constant in the language it is interpreted by some element of $\mathcal{A}$, and it follows from the extra axioms that

$$\mathcal{A} \models \underline{n} < c$$

for all $n \geq 0$.

Of course, $c^{\mathcal{A}}$ is not infinitely large from the perspective of $\mathcal{A}$, it's just a "natural number" there. As are $c^{\mathcal{A}} + 1$, $c^{\mathcal{A}} + c^{\mathcal{A}}$, $(c^{\mathcal{A}})^2$ and so on.

Note that the atomic diagram diag$(\mathfrak{N})$ of the natural numbers is very simple (atomic, we are not making any claims about the complete diagram).

A typical formula in the atomic diagram is $\underline{2} + \underline{3} \approx \underline{5}$ or $\neg \underline{3} < \underline{2}$.

Certainly diag$(\mathfrak{N})$, given some reasonable encoding, is decidable, even primitive recursive (at a very low level of the hierarchy).

Theorem (Tennenbaum, 1960)

No non-standard model $\mathcal{A}$ of (PA) *is computable: the atomic diagram of $\mathcal{A}$ is always undecidable.*

Note that we are dealing with the atomic diagram here, not quantified sentences. The problem comes from prime divisors of integers in the model:

$$\{\, n \in \mathbb{N} \mid \mathcal{A} \models p_n \text{ divides } a \,\}$$

If a here is non-standard, these sets can be very complicated.

Non-standard models of Peano arithmetic may seem like a mere curiosity, vaguely interesting but essentially useless, in particular in view of Tennenbaum's theorem.

A. Robinson realized in 1960, though, that this type of unintended model is the perfect framework for analysis: one can construct strange models of the reals that contain infinitesimal elements.

As a consequence, there is no need for limits in such a model.

In essence, differentiation is just a quotient operation $\frac{f(x+h)-f(x)}{h}$ and integrals are just sums.

The drawback is, of course, that someone trying to learn or apply calculus this way must already have a solid background in logic – perhaps unsurprisingly, non-standard analysis never really took off.

True arithmetic (TA) is defined as the theory of $\mathcal{N}$.

Of course, $\mathsf{Cn}(\mathrm{PA}) \subseteq (\mathrm{TA})$ but because of non-standard models we are nowhere near equality: (PA) only proves theorems that are valid in all models.

$$\forall\, \mathcal{A} \in \mathsf{Mod}(\mathrm{PA})\, \mathcal{A} \models \varphi \quad \text{implies} \quad (\mathrm{PA}) \vdash \varphi$$

This is most emphatically NOT the same as

$$\varphi \in (\mathrm{TA}) \quad \text{implies} \quad (\mathrm{PA}) \vdash \varphi$$

There are statements of arithmetic that hold in $\mathcal{N}$ but not in non-standard models.

Exercise

Show that if Γ has arbitrarily large finite models then Γ must have an infinite model.

Conclude that there are infinite fields of characteristic $p > 0$.

Exercise

Show that every partial order can be embedded into a total order.

Exercise

Show that there is no FO set of axioms Γ that characterizes finiteness in the sense that $\mathcal{A} \models \Gamma$ if, and only if, $\mathcal{A}$ is infinite.

- Theories and Models

- Decidability and Completeness

- Derivations and Proofs

- Compactness

5. Incompleteness and Undecidability

So how about (PA)? Is (PA) complete, and if not can we augment it a bit to obtain a complete theory?

Sadly, Kurt Gödel showed in 1931 that this is impossible: there are lots of facts about the natural numbers that cannot be proven in (PA), nor in any other reasonable axiom system for arithmetic (e.g., $\mathrm{Th}(\mathcal{N})$ is not a good choice).

Theorem

Suppose (PA) *is consistent. There is a sentence of arithmetic such that* (PA) *neither proves nor refutes this sentence.*

But of course $\mathcal{N}$ must be a model of this sentence φ or its negation $\neg\varphi$, it's just that (PA) is not strong enough to determine which.

One might still hope that such sentences are horribly complicated, but in fact a single universal quantifier suffices.

Gödel's second Incompleteness theorem pins down one such sentence explicitly.

Theorem

If (PA) *is consistent then its consistency cannot be proven in* (PA)*.*

This does not mean that we cannot prove (PA) to be consistent, we just cannot do it inside of the system (transfinite induction to ε_0 suffices, though, or you can just believe that $\mathcal{N}$ is a model).

Still, the missing theorems tend to be a bit weird, they don't seem to matter much in "real life", that is, in applications of number theory.

Note the hedge here, in 1975 Paris and Harrington showed how to construct statements of finite combinatorics that are true but not provable in (PA).

Programs and (PA)

How does this relate to computer science?

Suppose P is a program in some standard programming language that computes a numerical function $\widehat{P} : \mathbb{N} \to \mathbb{N}$.

As always write $\underline{n}$ for the numeral representing $n \in \mathbb{N}$ in (PA). Then there is a formula $\phi(x, y)$ of arithmetic such that

- $\widehat{P}(m) = n$ implies $(\mathrm{PA}) \vdash \phi(\underline{m}, \underline{n})$.
- $(\mathrm{PA}) \vdash \phi(x, y) \land \phi(x, z) \to y \approx z$.

But even though $\widehat{P}$ is total, (PA) may well not be powerful enough to prove it: in general

$$(\mathrm{PA}) \not\vdash \forall\, x\, \exists\, y\, \phi(x, y).$$

If we cannot prove all true statements of arithmetic, can we at least give a decision algorithm for them?

It is a direct consequence of the way proof systems are built that all formulae provable from a decidable set of axioms are semi-decidable: we can systematically enumerate all axioms and all proofs in the system based on these axioms – in principle, efficiency is not a consideration here.

Theorem

Let Γ be decidable. Then $\mathsf{Cn}(\Gamma)$ *is semi-decidable.*

So for $\mathsf{Cn}(\Gamma)$ to be decidable we only need a semi-algorithm for the non-theorems, the sentences that do not follow from Γ. Alas, there is no obvious general way to obtain such an algorithm. But sometimes it happens that the set of axioms is very rich and "knows" facts about non-theorems in the following sense:

$$\Gamma \not\vdash \varphi \qquad \text{implies} \qquad \Gamma \vdash \neg\varphi$$

Lemma

Suppose T is a complete and axiomatizable theory. Then T is decidable.

Proof. It suffices to show that the complement of T is semi-decidable. But $\varphi \notin T \iff \neg\varphi \in T$, done. □

Unfortunately, completeness of a theory is not all that easy to check in general. Here is one valuable test.

Let's assume throughout that the language we are using is countable. A set of sentences Γ is κ-categorical for some cardinal $\kappa \geq \aleph_0$ if all models of Γ of size κ are isomorphic.

Theorem (Los-Vaught Test)

Let T be a theory with no finite models. If T is κ-categorical for some cardinal $\kappa \geq \aleph_0$ then T is complete.

Algebraically Closed Fields

Let (ACF_0) be the theory of algebraically closed fields of characteristic 0.

Clearly, all models of (ACF_0) are infinite.

Suppose $\mathcal{A}$ is a model of (ACF_0) of cardinality $\kappa = |\mathbb{R}|$. Then $\mathcal{A}$ must be an extension field of $\mathbb{Q}$. The transcendence degree of $\mathcal{A}$ is κ. But then it is easy to construct an isomorphism between any two such models.

Hence (ACF_0) is complete.

Exercise

Fill in the details of the last argument.

Exercise

Show that the theory of atomless Boolean algebras is complete.

Exercise

Show that the theory of $\langle \mathbb{N}, S \rangle$ *is complete.*

Peano arithmetic is not complete and indeed undecidable.

Theorem (K. Gödel 1931, J.B. Rosser, 1936)

Peano arithmetic is undecidable.

Two questions arise naturally:

- Is there are large subsystem of Peano arithmetic that is still decidable?
- Is there are small subsystem of Peano arithmetic that is still undecidable?

To avoid undecidability we have to remove multiplication. Presburger arithmetic uses the language $\mathcal{L}(+, -, 0, 1; <)$ of signature $(2, 2, 0, 0; 2)$ and axiomatizes these operations over the integers $\mathbb{Z}$.

Theorem (M. Presburger, 1929)

Presburger arithmetic is decidable.

But note that the algorithm is not practical (triple exponential). In fact, even for quantifier-free formulae the problem is $\mathbb{NP}$-hard.

For the second question it turns out that indeed the full power of Peano arithmetic is not even needed for undecidability, the following rather weak fragment, called Q suffices:

successor

$$S(x) \not\approx 0 \qquad S(x) \approx S(y) \to x \approx y$$

$$x \not\approx 0 \to \exists\, y\, (x \approx S(y))$$

addition

$$x + 0 \approx x \qquad x + S(y) \approx S(x + y)$$

multiplication

$$x \cdot 0 \approx 0 \qquad x \cdot S(y) \approx (x \cdot y) + x$$

Theorem

The theory of Abelian groups is decidable. But the theory of general groups is undecidable.

This result is rather remarkable since the axioms are nearly identical: only one commutativity axiom is needed to enforce decidability.

This decidability result reflects nicely the fact that Abelian groups are much less complicated than general groups.

In general, many familiar structures in algebra are undecidable: rings, commutative rings, integral domains, fields, fields of characteristic 0.

Boolean algebra, on the other hand, is decidable.

For important structures $\mathcal{A}$ such as the rationals and reals, the question arises whether $\mathrm{Th}(\mathcal{A})$ is decidable (regardless of attempts at axiomatization).

Theorem (J. Robinson, 1948)

The theory of the rationals with addition and multiplication is undecidable.

This is in sharp contrast to a famous theorem by Tarski concerning the real numbers.

Theorem (A. Tarski, 1948)

The theory of the reals with addition and multiplication is decidable.

As a consequence, basic geometry is decidable.

The theorem is proved by a very interesting technique that provides a direct decision algorithm: quantifier elimination.

Tarski's original method was highly inefficient, though (not bounded by a stack of exponentials).

Definition

A theory T admits quantifier elimination if for every formula φ there is a quantifier-free formula φ_0 such that $T \vdash \varphi \leftrightarrow \varphi_0$.

It actually suffices to show that one can eliminate existential quantifiers in a formula

$$\exists\, x\, (A_1 \wedge A_2 \wedge \ldots \wedge A_n)$$

where all the A_i are literals (atomic formulae or negations thereof).

Note that on the face of it this may seem utterly hopeless: the existence of x may depend on the values of the free variables. It is not clear how to capture this dependence without quantifiers.

Exercise

Prove that full quantifier elimination follows from the weaker form.

Consider the theory of the simple structure $\mathcal{N}_S = \langle \mathbb{N}, S, 0 \rangle$.

We can give an axiom system Γ for this theory as follows:

$$\begin{aligned} S(x) &\not\approx 0 \\ S(x) \approx S(y) &\to x \approx y \\ y \not\approx 0 &\to \exists\, x\, S(x) \approx y \\ S^n(x) &\not\approx x \qquad \text{schema } n \geq 1 \end{aligned}$$

It is quite obvious that $\mathrm{Cn}(\Gamma) \subseteq \mathrm{Th}(\mathcal{N}_S)$ but because of non-standard models the obvious direction is far from obvious.

What does an arbitrary, non-standard model $\mathcal{A}$ of Γ look like?

Intuitively, we must have

$$A = \mathbb{N} \cup \mathbb{Z} \cup \mathbb{Z} \cup \mathbb{Z} \ldots$$

where the unions are meant to be disjoint: There is one copy of $\mathbb{N}$, the standard model, plus a number of copies of $\mathbb{Z}$.

We can make this precise by showing that $\exists\, n \geq 0\, (S^n(x) = y \vee S^n(y) = x)$ is an equivalence relation on A with the equivalence classes as indicated.

At any rate, the cardinality of A is $\aleph_0 + \aleph_0 \kappa$ where κ is the number of copies of $\mathbb{Z}$.

But then any two models of Γ with the same κ must be isomorphic. Hence by Los-Vaught Cn$(\Gamma) = \mathrm{Th}(\mathcal{N}_S)$ is complete. So we have decidability.

To get a real algorithm we show that quantifier elimination holds. Consider

$$\exists\, x\, (A_1 \wedge A_2 \wedge \ldots \wedge A_n)$$

We may safely assume that x occurs in each A_i. Now the positive atomic formulae must be of the form

$$S^n(x) \approx t.$$

If $t \approx S^m(x)$ the literal can easily be replaced by $\bot$ or $\top$. So we only have to deal with $S^n(x) \approx S^m(y)$ and $S^n(x) \approx 0$.

If all the literals are negative then we can replace the whole formula by $\top$.

So suppose $A_1 = S^n(x) \approx t$ is positive. But then we can replace A_1 by

$$t \not\approx 0 \wedge t \not\approx S(0) \wedge \ldots \wedge t \not\approx S^{n-1}(0).$$

Any other literal $S^r(x) \approx z$ is replaced by $S^r(t) \approx S^m(z)$ (and likewise for negative ones).

But then none of the literals contains x any more, so we can drop the quantifier.

Note that if we apply our elimination process to a sentence we wind up with a quantifier-free sentence, which must be a Boolean combination of pieces $S^m(0) \approx S^n(0)$.

Exercise

Fill in all the details in the argument.

In fact, even if we disregard attempts to axiomatically describe some reasonable class of structures and only consider the logic itself we run into undecidability.

Theorem

First-order logic is undecidable: it is undecidable whether a sentence in FOL is derivable from the empty set of axioms.

Theorem

Satisfiability in first-order logic is undecidable: it is undecidable whether a sentence in FOL has a model.

One has to be a bit careful about having enough function and relation symbols around for these results to hold. E.g., one binary relation symbol suffices, as does one unary and one ternary function symbol.

The point is: for any underlying language that is strong enough to be useful provability and satisfiability are undecidable.

From the point of view of verification this is a disaster: one cannot hope in general to automatically check specifications that are written in FOL.

For some limited areas of discourse there are decision algorithms. For example, the theory of real numbers is decidable (quantifier elimination, a result by Tarski). As a result, certain aspects of geometry are decidable (though the algorithms are not particularly efficient).

To get around this problem one has to regroup and develop languages that are expressive enough to state correctness properties, but still weak enough to allow for decision algorithms.