Logics for Computation

Lecture #8: We like it Complete and Compact

(and have a Soft Spot for Löwenheim-Skolem)

Carlos Areces and Patrick Blackburn

{carlos.areces,patrick.blackburn}@loria.fr

INRIA Nancy Grand Est
Nancy, France

ESSLLI 2008 - Hamburg - Germany

The story up till now

- We put together various bits-and pieces we have played with (names, :, diamonds) added the $\langle x \rangle$ and $[x]$ operators, and reached the most expressive language we have seen so far.
- The approach might seem somewhat shake-and-bake (or as the French would probably put it: “Bricolage!”), but we arrived at somewhere very important.
- For we reached “first-order logic”, the logic often considered to “classical logic”, and one of the most distinctive and important spots on the logical landscape. For our $\langle x \rangle$ and $[x]$ are really just the familiar $\exists$ and $\forall$ quantifiers.

Not called classical logic for nothing

- Some logicians (notably Quine) view first-order logic as the be-all and end-all of logic.
- We disagree with this viewpoint, but there is no getting away from a very stubborn fact — first-order logic certainly is special.
- Being "first-order logic" is not a matter of notation or symbolism. It's about something much deeper. It's about finally being able to get to grips with each and every element in the domains of our models.
- And as we shall learn today and tomorrow, this has some deep consequences. Today we're going to take a (rather abstract) look at inference and expressivity in first-order logic.

In today's lecture

- First we'll look at inference. We'll briefly discuss what we have to do to get a tableaux system for the quantifiers, and then discuss the concepts of soundness and completeness. We'll avoid computational issues —- though we will observe that model checking, the simplest of our inference tasks, is starting to get computationally difficult.
- But then we turn to the heart of the lecture: expressivity. We have clearly gained a lot of expressivity, but interesting gaps remain. We will discuss two results, the Compactness Theorem and the Löwenhein-Skolem theorem(s) that pin down the crucial expressive limitations of first-order logic. Indeed, as we shall learn, these results actually characterize first-order logic.

To infinity and beyond. . . ?!

The distinction between finiteness and the different degrees of infinity will play an important role in the later part of this lecture. Will we get to infinity? Will we get beyond? Or is first-order logic really just a toy?

Example: Theory of linear order

We now have a lot of expressive power — enough to state some interesting theories. Here, for example, is a simple theory of linear order:

- Axiom Irr: $[x](x : \neg\langle R\rangle x)$
- Axiom Tran: $[x][y](x : \langle R\rangle\langle R\rangle y \rightarrow x : \langle R\rangle y)$
- Axiom Lin: $[x][y](x : y \vee x : \langle R\rangle y \vee y : \langle R\rangle x)$

So if we had a proof system for our first-order language, we could proved some non-trivial theorems:

$$\models \mathit{Irre} \wedge \mathit{Tran} \wedge \mathit{Lin} \rightarrow [x][y](x : \langle R\rangle y \rightarrow y : \neg\langle R\rangle x).$$

That is (recall Lecture 1) we would be able to handle Euclid-style reasoning.

Two Tableaux Rules

Here are the two key tableau rules we need:

$$\frac{s : \langle x \rangle \varphi}{s : \varphi[x \leftarrow n]} \; (\langle x \rangle)$$

(*where n is an new name*)

$$\frac{s : \neg \langle x \rangle \varphi}{s : \varphi[x \leftarrow o]} \; (\neg \langle x \rangle)$$

(*where n is an old name*)

Actually, we need other rules too — as yet we have given no rules for : or for names — but we won't bother with these here.

Instead we ask again the question raised yesterday: how to we know that what we're doing is right? And what does "right" mean anyway?

Proof theory

- Proof theory is the syntactic approach to logic.
- It attempts to define collections of rules and/or axioms that enable us to generate new formulas from old. That is, it attempts to pin down the notion of inference syntactically.
- Given some proof system P, we write $\vdash_P \phi$ to indicate that a formula ϕ is provable in the the proof system. (Incidentally, $\nvdash_P \phi$ means that ϕ is *not* provable in proof system P.).

Many types of proof system

- Natural deduction
- Hilbert-style system (often called axiomatic systems)
- Sequent calculus
- Tableaux systems
- Resolution

Why so many different proof systems?

- Well, one of the most important may simply be that logicians love to play with such systems — and every logician has his or her own favourite pet system!
- A more serious reason is: different proof systems are typically good for different purposes.
- In particular, some systems (notably tableau and resolution) are particularly suitable for computational purposes.

But what does all this have to do with semantics and inference?

- Note: nothing we have said so far bout proof systems makes any connection with the model-theoretic ideas previously introduced.
- All we have done is talk about provability and vaguely said that we want to "generate" formulas syntactically. What does this have to do with relational structures and semantics?

But what does all this have to do with semantics and inference?

- Note: nothing we have said so far bout proof systems makes any connection with the model-theoretic ideas previously introduced.
- All we have done is talk about provability and vaguely said that we want to "generate" formulas syntactically. What does this have to do with relational structures and semantics?
- Answer: we insist on working with proof systems with two special properties, namely soundness and completeness.

Soundness

- Recall that we write $\models \phi$ to indicate that the formula ϕ is valid (that is, satisfied in all models under all assignments).
- Recall that we write $\vdash_P \phi$ to indicate that ϕ is provable in proof system P.
- We say that a proof system P is sound if and only if

$$\vdash_P \phi \text{ implies } \models \phi$$

.

Explanation

- That is, soundness means that syntactic provability implies semantic validity.
- To put it another way: P does not produce garbage.
- And another: P is "safe".
- Needless to say, all the standard proof systems are sound.

Remark

- Soundness is typically an easy property to prove.
- Proofs typically have some kind of inductive structure. One shows that if the first part of proof is true in a model, then the rules only let us generate formulas that are also true in a model.

Completeness

- Recall that we write $\models \phi$ to indicate that the formula ϕ is valid (that is, satisfied in all models under all assignments).
- Recall that we write $\vdash_P \phi$ to indicate that ϕ is provable in proof system P.
- We say that a proof system p is complete if and only if

$$\models \phi \text{ implies } \vdash_P \phi$$

.

Explanation

- That is, completeness means that our proof system is strong enough to prove everything that is provable.
- To put it another way: if some formula really is true in all models, then our proof system P really is powerful enough to generate it.
- And another: no valid formula is out of reach of our proof system.
- The standard proof systems are complete.

Remark

- Completeness is a much deeper property that soundness, and is a lot more difficult to proof.
- It is typically proved by contraposition. We show that if some formula is not provable ($\not\vdash \phi$) then ϕ is not valid ($\not\models \phi$). This is done by building a model for $\neg\phi$.
- And our first-order language is very expressive — so it can describe some pretty intricate models. And they certainly won't all be trees!

Soundness and completeness together

- Recall: proof system P is sound if and only if

$$\vdash_P \phi \text{ implies } \models \phi$$

.

- Proof system P is complete if and only if

$$\models \phi \text{ implies } \vdash_P \phi$$

.

- So if a proof system is both sound and complete (which is what we want) we have that:

$$\models \phi \text{ if and only if } \vdash_P \phi$$

.

- That is, syntactic provability and semantic validity coincide.

A remark on first-order model checking

- I've carefully avoided saying anything about computational issues — because there is a lot to say and that is the topic of tomorrows talk! We've brought a lot of expressivity — and we're going to have to pay for it!
- But things are getting computationally complex. To see this, we remark that even the (usually easy) model checking task is getting tough.
- In fact the model checking task for first-order logic is PSPACE-complete! That is, it is as tough as validity/satisfiability checking for the diamond language!.
- To see why, reflect on what we have to do to check a formula of the form $[x]\langle y\rangle[z]\langle w\rangle\cdots[v]\langle x\rangle\phi$.

Expressivity of first-order logic

- We turn now to the theme of expressivity of first-order logic.
- Our discussion revolves around two famous results: the Compactness Theorem, and the Löwenheim Skolem Theorem(s). I'm going to be state (but won't prove) these results and discuss some of their consequences.
- As we said at the start of the lecture, our discussion will have a lot to do with finiteness, and the different grades of infinity. So before going any further let's remind ourselves what in infinite set is . . .

Finite and Infinite

Finite and Infinite

- The most fundamental infinite set in set theory is the natural numbers **N**. This is the set

$$\{0, 1, 2, 3, 4, 5, \ldots\}$$

.

Finite and Infinite

- The most fundamental infinite set in set theory is the natural numbers **N**. This is the set

$$\{0, 1, 2, 3, 4, 5, \ldots\}$$

.

- A set S is infinite if there is one-to-one (injective) function from **N** to S.

Finite and Infinite

- The most fundamental infinite set in set theory is the natural numbers **N**. This is the set

$$\{0, 1, 2, 3, 4, 5, \ldots\}$$

.

- A set S is infinite if there is one-to-one (injective) function from **N** to S.
- A set is finite if it is not infinite. That is, a set is finite it is not possible to define such a function. This amounts to saying that that it is the same size as some finite set

$$\{1, 2, 3, 4, 5, \ldots, n\}$$

Such a set, incidentally, is called a an initial segment of the natural numbers.

Infinite Axiom Sets!

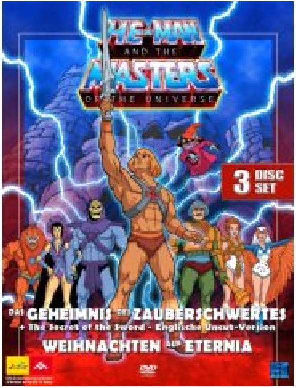

- We already seen a simple set of axioms, namely the three axioms that defined the theory of linear order.
- But note: any finite set of axioms can be replaced by one single axiom — for we simply need to form their conjunction!
- So, very early in the history of logic, logicians started to consider the consequences of working with infinite sets of axioms After all, why not? Remember: Logicians are the Masters of the Universe!
- Now, using infinite sets of axioms does give us more power. . .

An infinite set of axioms defining infinity!

An infinite set of axioms defining infinity!

- Try to write down a single sentence of first-order logic that is true in every infinite model (and false on every finite one). But don't try too hard, for I promise that you won't find one...

An infinite set of axioms defining infinity!

- Try to write down a single sentence of first-order logic that is true in every infinite model (and false on every finite one). But don't try too hard, for I promise that you won't find one...
- But it's easy to write down an infinite set of sentences with this property.

An infinite set of axioms defining infinity!

- Try to write down a single sentence of first-order logic that is true in every infinite model (and false on every finite one). But don't try too hard, for I promise that you won't find one. . .
- But it's easy to write down an infinite set of sentences with this property.
- Let $AtLeast_2$ be $\langle x\rangle\langle y\rangle(x : \neg y)$.

An infinite set of axioms defining infinity!

- Try to write down a single sentence of first-order logic that is true in every infinite model (and false on every finite one). But don't try too hard, for I promise that you won't find one...
- But it's easy to write down an infinite set of sentences with this property.
- Let $AtLeast_2$ be $\langle x \rangle \langle y \rangle (x : \neg y)$.
- Let $AtLeast_3$ be $\langle x \rangle \langle y \rangle \langle z \rangle (x : \neg y \wedge x : \neg z \wedge y : \neg z)$.

An infinite set of axioms defining infinity!

- Try to write down a single sentence of first-order logic that is true in every infinite model (and false on every finite one). But don't try too hard, for I promise that you won't find one...
- But it's easy to write down an infinite set of sentences with this property.
- Let $AtLeast_2$ be $\langle x \rangle \langle y \rangle (x : \neg y)$.
- Let $AtLeast_3$ be $\langle x \rangle \langle y \rangle \langle z \rangle (x : \neg y \wedge x : \neg z \wedge y : \neg z)$.
- Following this pattern, we can write down sentences $AtLeast_4$, $AtLeast_5$, $AtLeast_6$, $AtLeast_7$,...,$AtLeast_{100}$,... ....

An infinite set of axioms defining infinity!

- Try to write down a single sentence of first-order logic that is true in every infinite model (and false on every finite one). But don't try too hard, for I promise that you won't find one...
- But it's easy to write down an infinite set of sentences with this property.
- Let $AtLeast_2$ be $\langle x\rangle\langle y\rangle(x : \neg y)$.
- Let $AtLeast_3$ be $\langle x\rangle\langle y\rangle\langle z\rangle(x : \neg y \wedge x : \neg z \wedge y : \neg z)$.
- Following this pattern, we can write down sentences $AtLeast_4$, $AtLeast_5$, $AtLeast_6$, $AtLeast_7$,...,$AtLeast_{100}$,... ....
- Let **INF** be the set of all these sentences. Then we have that

$$\mathcal{M} \models \textbf{INF} \text{ iff } \mathcal{M} \text{ is an infinite model.}$$

Compactness Theorem

Compactness Theorem: Let Σ be an infinite set of first-order sentences. If every finite subset of Σ can be made true (in some model or other) then there is at least one model that simultaneously makes every sentence in Σ true.

We can paraphrase this as follows. To check whether an infinite set of first-order sentences Σ, has a model, we don't need to try and find a model that make all the sentences in Σ true at once. If it enough to show that any finite subset of Σ can be made true in some model. For if we can show this, the Compactness Theorem guarantees that there is some model that makes all the (infinitely many) sentences in Σ true all at once.

How do you prove the Compactness Theorem?

Many proofs are known, but two are worth mentioning...

- Actually, one can prove the Compactness Theorem for first-order logic more-or-less simultaneously with the Completeness Theorem. More precisely, completeness for first-order logic is a reasonably simple extension of compactness.
- And there is another (in a sense more revealing) proof. The ultraproduct construction lets us "multiply together" the finite models into one big model.

A powerful theorem

The Compactness Theorem is central to mathematical model theory:

- A powerful theorem — and a two-sided one.
- On the positive side, it allows us to build many interesting and unexpected model (such as non-standard models of arithmetic).
- And it has negative uses too — it can also show that we cannot define certain things. And that's the kind use we will put it now.

Finiteness is not first-order definable

The question we will ask is the following:

> *Is there a single sentence of first-order logic that defines finiteness?*

That is, is there a single first order sentence (let's call it **fin** such that:

$$\mathcal{M} \models \mathbf{fin} \textit{ iff } \mathcal{M} \textit{ is finite}$$

As we shall see, the answer is no.

The argument

The argument

- Suppose there is such a sentence **fin**. We are going to show (with the help of the Compactness Theorem) that this assumption leads to a contradiction.

The argument

- Suppose there is such a sentence **fin**. We are going to show (with the help of the Compactness Theorem) that this assumption leads to a contradiction.
- Consider the set of sentences $\mathbf{INF} \cup \mathbf{fin}$

The argument

- Suppose there is such a sentence **fin**. We are going to show (with the help of the Compactness Theorem) that this assumption leads to a contradiction.
- Consider the set of sentences $\mathbf{INF} \cup \mathbf{fin}$
- That is: $\{AtLeast_1, AtLeast_2, AtLeast_3, AtLeast_4, \ldots\} \cup \mathbf{fin}$

The argument

- Suppose there is such a sentence **fin**. We are going to show (with the help of the Compactness Theorem) that this assumption leads to a contradiction.
- Consider the set of sentences $\mathbf{INF} \cup \mathbf{fin}$
- That is: $\{AtLeast_1, AtLeast_2, AtLeast_3, AtLeast_4, \ldots\} \cup \mathbf{fin}$
- Claim: every finite subset of this has a model. Why is this?

And now we have a problem

And now we have a problem

- Since every finite subset of **INF** $\cup$ **fin** is true, by the Compactness Theorem, there is some model $\mathcal{M}$ that simultaneously makes every sentence in **INF** $\cup$ **fin**

And now we have a problem

- Since every finite subset of $\mathbf{INF} \cup \mathbf{fin}$ is true, by the Compactness Theorem, there is some model $\mathcal{M}$ that simultaneously makes every sentence in $\mathbf{INF} \cup \mathbf{fin}$
- Hence, as $\mathcal{M}$ makes $\mathbf{INF}$ true, $\mathcal{M}$ is infinite.

And now we have a problem

- Since every finite subset of **INF** $\cup$ **fin** is true, by the Compactness Theorem, there is some model $\mathcal{M}$ that simultaneously makes every sentence in **INF** $\cup$ **fin**
- Hence, as $\mathcal{M}$ makes **INF** true, $\mathcal{M}$ is infinite.
- But as $\mathcal{M}$ also makes **fin** true, so $\mathcal{M}$ is finite!

And now we have a problem

- Since every finite subset of **INF** $\cup$ **fin** is true, by the Compactness Theorem, there is some model $\mathcal{M}$ that simultaneously makes every sentence in **INF** $\cup$ **fin**
- Hence, as $\mathcal{M}$ makes **INF** true, $\mathcal{M}$ is infinite.
- But as $\mathcal{M}$ also makes **fin** true, so $\mathcal{M}$ is finite!
- Ooooooooooooops!!!!!!!!!!!

And now we have a problem

- Since every finite subset of **INF** $\cup$ **fin** is true, by the Compactness Theorem, there is some model $\mathcal{M}$ that simultaneously makes every sentence in **INF** $\cup$ **fin**
- Hence, as $\mathcal{M}$ makes **INF** true, $\mathcal{M}$ is infinite.
- But as $\mathcal{M}$ also makes **fin** true, so $\mathcal{M}$ is finite!
- Oooooooooooops!!!!!!!!!!!
- From this contradiction we deduce that **fin** does not exist. That is, there is no sentence of first-order logic that expresses the concept of finiteness.

Countably infinite and bigger. . .

Countably infinite and bigger...

- A set S is countably infinite if there is one-to-one (injective) and onto (surjective) function from $\mathbf{N}$ to S.

Countably infinite and bigger...

- A set S is countably infinite if there is one-to-one (injective) and onto (surjective) function from $\mathbf{N}$ to S.
- Countable infinite sets are the smallest infinite sets. Examples of countable infinite sets include: $\mathbf{N}$, $\mathbf{E}$ (the even numbers), $\mathbf{Z}$ (the integers), and $\mathbf{Q}$ (the fractions).

Countably infinite and bigger. . .

- A set S is countably infinite if there is one-to-one (injective) and onto (surjective) function from $\mathbf{N}$ to S.
- Countable infinite sets are the smallest infinite sets. Examples of countable infinite sets include: $\mathbf{N}$, $\mathbf{E}$ (the even numbers), $\mathbf{Z}$ (the integers), and $\mathbf{Q}$ (the fractions).
- The set of real numbers $\mathbf{R}$ is not countably infinite — it is bigger.

Countably infinite and bigger...

- A set S is countably infinite if there is one-to-one (injective) and onto (surjective) function from $\mathbf{N}$ to S.
- Countable infinite sets are the smallest infinite sets. Examples of countable infinite sets include: $\mathbf{N}$, $\mathbf{E}$ (the even numbers), $\mathbf{Z}$ (the integers), and $\mathbf{Q}$ (the fractions).
- The set of real numbers $\mathbf{R}$ is not countably infinite — it is bigger.
- In general, the power set $\mathcal{P}(S)$ of S is bigger than S.

Countably infinite and bigger...

- A set S is countably infinite if there is one-to-one (injective) and onto (surjective) function from $\mathbf{N}$ to S.
- Countable infinite sets are the smallest infinite sets. Examples of countable infinite sets include: $\mathbf{N}$, $\mathbf{E}$ (the even numbers), $\mathbf{Z}$ (the integers), and $\mathbf{Q}$ (the fractions).
- The set of real numbers $\mathbf{R}$ is not countably infinite — it is bigger.
- In general, the power set $\mathcal{P}(S)$ of S is bigger than S.
- The universe of sets is huge. It contains infinitely many ever-bigger grades of infinity.

Löwenhein-Skolem Theorems

The Löwenheim Skolem Theorems tell us that first-order logic is completely blind to all these distinctions:

> **Upward Löwenhein-Skolem Theorem:** *If a set of sentences Σ has an infinite model, it has infinite models of all larger infinite cardinalities.*
>
> **Downward Löwenhein-Skolem Theorem:** *If a set of sentences Σ has an infinite model, it has infinite models of all lower infinite cardinalities. In particular, it always has a countable model.*

It follows that is is not possible to fully describe the real numbers $\mathbf{R}$ (that is, describe them up to isomorphism) using first-order logic. Any description will have many models (indeed models of every infinite cardinality) and a model that (being countable) is too small!

Lindström's Theorem

- There is a field of logic called abstract model theory which works with very general and abstract definitions of what logics are.
- There is a celebrated result from abstract model theorem called Lindtröm's Theorem which tells us that first-order logic is the only logic for which both the Compactness and the Downward Löwenheim-Skolem Theorems hold.
- That is, if you invent a very strange logical formalism, but can prove that it has these two properties, then you have invented a logic with exactly the same expressive power as first-order logic. You "new" logic, when you get right down to it, is in the business of quantifying over individuals.
- And is there really anything else out there to quantify over....?

What we Covered Today

- Brief remarks on inference for the $\langle x \rangle$ language: how to build a tableaux system, and why even model checking is starting to get hard (PSPACE-complete).
- Soundness and completeness and why they are important.
- Expressivity for the $\langle x \rangle$ language: Compactness and the Löwenheim Skolem Theorem.
- We also learned that "first-order logic" occupies a fundamental place in the expressivity hierarchy. Being "first-order" is not about this notation or that notation; it's something more fundamental.
- But there remains a big gap in our discussion: computability.

Relevant Bibliography

Georg Cantor (1845-1918) created modern set theory, which is the setting for virtually all of modern mathematics. As David Hilbert, the famous mathematician and logician once remarked, "No one shall expel us from the Paradise that Cantor has created."

Relevant Bibliography

The Completeness Theorem for first-order logic was proved by Kurt Gödel (1906–1978) in his 1928 doctoral thesis. As it turned out, however, this was merely the first of many great results that he was to prove.

Relevant Bibliography

The modern form of the Compactness Theorem seems to trace back to the work of the Russian mathematician Anatoly Maltsev (1909–1967).

Relevant Bibliography

The downward Löwenheim-Skolem Theorem was first proved by Leopold Löwenheim (1878–1957; top picture) in 1915. In 1920 his proof was greatly simplified and generalised by Thoralf Skolem (1887 -1963).

The Next Lecture

Computing the Uncomputable