

Introduction to Logic

First-Order Logic

Michael Genesereth
Computer Science Department
Stanford University

Objects with Multiple Names

People (Nicknames):

$michael \leftrightarrow$ 
 $mike \leftrightarrow$ 

Arithmetic:

$plus(s(0), s(0)) \leftrightarrow$
 $times(s(s(0)), s(0)) \leftrightarrow 2$
 $s(s(0)) \leftrightarrow$


In real life, we often find it convenient to use more than one term refer to the same real world object. For example, we sometimes have nicknames for the same person - Michael, Mike. And, in Elementary Arithmetic, we frequently use different terms to refer to the same number - $2+2$, $2*2$, and $s(s(s(s(0))))$. This contradicts our assumption about isomorphic representation. However, we can deal with this in the context of Herbrand Logic.

Objects with No Names

People

$michael \leftrightarrow$ 
 $maureen \leftrightarrow$ 

Unnamed People



The reverse is also true, We sometimes find that there are objects in an application area for which we have no names. For example, we have names for many real numbers - integers, decimal numbers, and some transcendental numbers, like pi and e. However, we do not have names for all of them because we have only countably many names and uncountably many objects.

Mathematical Objects Without Names

Named real numbers (countably many):

123
34.12
 π
 e

All the others (uncountably many):

3.141544878723489184093893477809489084...
6.878989783975975738975379875837593358...

How many floating point numbers are there?

First Order Logic

Syntax

Identical to Herbrand Logic

Semantics

Herbrand Logic - defined in terms of language

First Order Logic - defined in terms of external worlds

Consequences

Differences in ease of understanding

Differences in some results (ground terms)

Differences in what can be defined

Differences in completeness and decidability

Language of First Order Logic

$p(a)$ $q(b,c)$

$\neg p(b)$ $\forall x.(p(x) \Rightarrow q(x,f(x)))$

$p(c) \vee p(d)$ $\exists x.q(x,d)$

Identical to Syntax of Functional Logic

Interpretations


A vocabulary is a set of symbols.

$$\{a, b, f, q\}$$

A *universe of discourse* is an *arbitrary* set of objects.



An *interpretation* is an assignment to symbols in language.



Tarskian semantics is the traditional semantics for First-Order Logic. In Tarskian semantics, a model consists of an arbitrary set of objects (called the universe of discourse) and an interpretation function (1) that maps object constants into elements of this set, (2) that maps function constants into functions on this set, and (3) that maps relation constants into relations on this set. A model of this sort completely determines the truth or falsity of all sentences in the language. And it gives us a definition of logical entailment. Note, however, that there are unboundedly many interpretations for any language, and entailment is defined over all conceivable universes - finite, countably infinite, uncountable, and beyond. It also requires an understanding of relations as set of tuples of objects.

Interpretations

A vocabulary is a set of symbols.

$$\{a, b, f, q\}$$

A *universe of discourse* is an *arbitrary* set of objects.

$$\{1, 2, 3, 4\}$$

An *interpretation* is an assignment to symbols in language.

$$\begin{aligned} a &= 1, b = 2 \\ f &= \{1 \rightarrow 2, 2 \rightarrow 3, \dots\} \\ q &= \{\langle 1, 2 \rangle, \langle 2, 3 \rangle, \langle 2, 4 \rangle, \dots\} \end{aligned}$$

Entailment defined in terms of *all* interpretations over
all possible universes of *all* possible sizes

Variable Assignments

A *variable assignment* for a universe of discourse U is a function assigning variables to objects in U .

$$v: \text{Variable} \rightarrow U$$

Universe of Discourse:

$$U = \{1, 2, 3\}$$

Example:

$$\begin{aligned} v(x) &= 1 \\ v(y) &= 2 \\ v(z) &= 3 \end{aligned}$$

Example:

$$\begin{aligned} v(x) &= 2 \\ v(y) &= 2 \\ v(z) &= 2 \end{aligned}$$

A variable assignment is for variables what an interpretation is for object constants. Given a universe of discourse U , a *variable assignment* is an arbitrary mapping from the variables of the language into U .

Value Assignments

A *value assignment* s_{iv} based on interpretation i and variable assignment v is a mapping from the terms of the language into the universe of discourse.

$$\begin{aligned}s_{iv}(\sigma) &= i(\sigma) \\ s_{iv}(v) &= v(v) \\ s_{iv}(\pi(\tau_1, \dots, \tau_n)) &= i(\pi)(s_{iv}(\tau_1), \dots, s_{iv}(\tau_n))\end{aligned}$$

A value assignment based on interpretation i and variable assignment v is a mapping from the terms of the language into the universe of discourse. The mapping must agree with i on constants, and it must agree with v on variables. The value of a functional term is the result of applying the function corresponding to the function constant to the objects designated by the terms.

Relational Sentences

A truth assignment satisfies a relational sentence if and only if the tuple of objects denoted by the arguments is a member of the relation denoted by the relation constant.

$$\begin{aligned}t_{iv}(\rho(\tau_1, \dots, \tau_n)) &= \text{true} \text{ if } \langle s_{iv}(\tau_1), \dots, s_{iv}(\tau_n) \rangle \in i(\rho) \\ &= \text{false} \text{ otherwise}\end{aligned}$$

Given an interpretation i and a variable assignment v , a relational sentence is true if and only if the tuple of objects denoted by the arguments is contained in the relation denoted by the relation constant.

Logical Sentences

$$\begin{aligned}t_{iv}(\neg\varphi) &= \text{true} \text{ iff } t_{iv}(\varphi) = \text{false} \\ t_{iv}(\varphi \wedge \psi) &= \text{true} \text{ iff } t_{iv}(\varphi) = \text{true} \text{ and } t_{iv}(\psi) = \text{true} \\ t_{iv}(\varphi \vee \psi) &= \text{true} \text{ iff } t_{iv}(\varphi) = \text{true} \text{ or } t_{iv}(\psi) = \text{true} \\ t_{iv}(\varphi \Rightarrow \psi) &= \text{true} \text{ iff } t_{iv}(\varphi) = \text{false} \text{ or } t_{iv}(\psi) = \text{true} \\ t_{iv}(\varphi \Leftrightarrow \psi) &= \text{true} \text{ iff } t_{iv}(\varphi) = t_{iv}(\psi)\end{aligned}$$

The conditions for truth of logical sentences in First Order Logic are analogous to those for the truth of logical sentences in Propositional Logic and Herbrand Logic. A negation $\neg\phi$ is true if and only if ϕ is false. A conjunction $(\phi \wedge \psi)$ is true if and only if both ϕ and ψ are true. A disjunction $(\phi \vee \psi)$ is true if and only if ϕ is true or ψ is true (or both). An implication is true unless the antecedent is true and the consequent is false. A reduction is true unless the antecedent is true and the consequent is false. An equivalence is true if and only if the truth values of the two embedded sentences are the same.

Quantified Sentences

Intuitively, a universally quantified sentence is true if and only if it is true no matter what *value* we assign to the universally quantified variable.

Intuitively, an existentially quantified sentence is true if and only if it is true for some *value* of the existentially quantified variable.

Stating these definitions precisely is a little tricky due to the possibility of nested quantifiers.

$$\forall x.(\exists y.r(x,y) \Rightarrow \forall x.r(x,x))$$

Intuitively, a universally quantified sentence is satisfied if and only if the embedded sentence is satisfied for all values of the quantified variable. Intuitively, an existentially quantified sentence is satisfied if and only if the embedded sentence is satisfied for some values of the quantified variable.

Sadly, expressing this intuition rigorously is a little tricky. First, we need the notion of a version of a variable assignment. Then, we can give a rigorous definition to our intuition by talking about some versions or all versions of a given variable assignment.

Versions

A *version* $v[\omega \leftarrow x]$ of a variable assignment v is the variable assignment that agrees with v on all variables except ω , which is assigned the value x .

$$\begin{aligned} v[\omega \leftarrow x](\theta) &= x && \text{if } \theta = \omega \\ v[\omega \leftarrow x](\theta) &= v(\theta) && \text{if } \theta \neq \omega \end{aligned}$$

A version $v[\omega \leftarrow x]$ (read v with ω gets x) of a variable assignment v is a variable assignment that assigns object x as value of the variable v and agrees with v on all other variables. v with ω gets x applied to θ is x if $\theta = \omega$; otherwise, the value is the same as $v(\theta)$.

Quantified Sentences

A universally quantified sentence is true in interpretation i and variable assignment v if and only if the scope is true for i and *every* version of v .

$$t_{iv}(\forall \omega. \varphi) = \text{true} \text{ iff } t_{iv[\omega \leftarrow x]}(\varphi) = \text{true} \text{ for all } x \in |i|.$$

An existentially quantified sentence is true in interpretation i and variable assignment v if and only if the scope is true for i and *some* version of v .

$$t_{iv}(\exists \omega. \varphi) = \text{true} \text{ iff } t_{iv[\omega \leftarrow x]}(\varphi) = \text{true} \text{ for some } x \in |i|.$$

With this notion, we can formalize the semantics of quantified sentences as follows. An interpretation i and a variable assignment v satisfy a universally quantified sentence $\forall v. \varphi$ if and only if i and v satisfy φ for all versions of v that can be defined over the universe of discourse of i . An interpretation i and a variable assignment v satisfy an existentially quantified sentence $\exists v. \varphi$ if and only if i and v satisfy φ for at least one version of v that can be defined over the universe of discourse of i .

Comparison

First Order Logic: A universally quantified sentence is true in interpretation i and variable assignment v if and only if the scope is true for i and *every* version of v compatible with interpretation i .

Herbrand Logic: A universally quantified sentence is true in a truth assignment if and only if every instance is true.

Does this look complicated to you? If so, you are not alone. Many students, seeing the semantics of first order logic for the first time, are bewildered. The intuition behind the semantics is simple enough, but the details are quite complex. Much more so than for Herbrand Logic. To see this, compare the semantics of universally quantified sentences in the two logics. In first order logic, we say that a universally quantified sentence is true in interpretation i and variable assignment v if and only if the scope is true for i and every version of v compatible with i . In Herbrand Logic, we say that a universally quantified sentence is true in a truth assignment if and only if every instance of the scope is true. That is all. The Herbrand Logic definition is simpler, and for many people it is more intuitive as well. We talk more about the two logics in a later segment.

Herbrand Logic vs First-Order Logic

In Herbrand Logic, if $\Delta \models p(\tau)$ for *every* ground term τ , does $\Delta \models \forall x.p(x)$?

Yes.

Let's look at a few definitions and observations before we close out this segment. An interpretation i and variable assignment v satisfy a sentence if and only if it is true according to the rules we have just seen. An interpretation of a sentence is a *model* if and only if it satisfies the sentence for every compatible variable assignment.

Herbrand Logic vs First-Order Logic

In Herbrand Logic, if $\Delta \models p(\tau)$ for *every* ground term τ , does $\Delta \models \forall x.p(x)$?

Yes.

In First-Order Logic, if $\Delta \models p(\tau)$ for *every* ground term τ , does $\Delta \models \forall x.p(x)$?

No. There can be objects without names.

Let's look at a few definitions and observations before we close out this segment. An interpretation i and variable assignment v satisfy a sentence if and only if it is true according to the rules we have just seen. An interpretation of a sentence is a *model* if and only if it satisfies the sentence for every compatible variable assignment.

Herbrand Logic and Uncountable Worlds

Can we describe uncountable world in Herbrand Logic?

No. There only countably many terms and countably many ground sentences in our language.

Upshot: It is not possible to axiomatize uncountable worlds in Herbrand Logic.

Let's look at a few definitions and observations before we close out this segment. An interpretation i and variable assignment v satisfy a sentence if and only if it is true according to the rules we have just seen. An interpretation of a sentence is a *model* if and only if it satisfies the sentence for every compatible variable assignment.

FOL and Uncountable Worlds

Can we describe uncountable worlds in First-Order Logic?

Yes. There can be objects without names.

Lowenheim-Skolem Theorem: If a set of sentences in First Order Logic has a model of one infinite cardinality, then it has a model of every infinite cardinality. (This striking result is true but not obvious.)

Upshot: It is not possible to write a sentence in First Order Logic that is true in an uncountable world and not true in any countably infinite world or vice-versa.

Let's look at a few definitions and observations before we close out this segment. An interpretation i and variable assignment v satisfy a sentence if and only if it is true according to the rules we have just seen. An interpretation of a sentence is a *model* if and only if it satisfies the sentence for every compatible variable assignment.

Completeness of Herbrand Logic

Peano Arithmetic

Transitive Closure

Both of these are finite axiomatizations and are complete (i.e. they precisely define which sentences are true and which are false). There are no *non-standard* models.

More disturbing is that there is no complete proof procedure for Relational Logic with Herbrand semantics. Godel's incompleteness theorem tells us that the set of all true sentences of Peano arithmetic is not computably enumerable. Our axiomatization is complete using Herbrand Semantics. If Herbrand entailment were semi-decidable, the set of all true sentences would be enumerable. Consequently, there is no complete (semi-decidable) proof procedure for Relational Logic with Herbrand semantics.

Incompleteness of First-Order Logic

First-Order Logic (FOL) theories with infinite universes have *nonstandard models* (unintended models that *cannot be excluded*).

Upshot: FOL is *weaker* than Herbrand Logic. Some notions that can be defined exactly in Herbrand Logic cannot be defined in FOL without allowing nonstandard models, e.g. Peano Arithmetic, transitive closure.

More disturbing is that there is no complete proof procedure for Relational Logic with Herbrand semantics. Godel's incompleteness theorem tells us that the set of all true sentences of Peano arithmetic is not computably enumerable. Our axiomatization is complete using Herbrand Semantics. If Herbrand entailment were semi-decidable, the set of all true sentences would be enumerable. Consequently, there is no complete (semi-decidable) proof procedure for Relational Logic with Herbrand semantics.

Inferential Completeness of First Order Logic

There **is** a proof procedure for First Order Logic that is both sound and complete. (Spoiler Alert: Fitch without Domain Closure and Induction does the trick.)

Moreover, by systematically applying the procedure, possible to compute all logical consequences of any enumerable set of premises. (Apply Fitch to produce all finite proofs in systematic way.)

Upshot: provability and logical entailment are semi-decidable (though not decidable).

More disturbing is that there is no complete proof procedure for Relational Logic with Herbrand semantics. Godel's incompleteness theorem tells us that the set of all true sentences of Peano arithmetic is not computably enumerable. Our axiomatization is complete using Herbrand Semantics. If Herbrand entailment were semi-decidable, the set of all true sentences would be enumerable. Consequently, there is no complete (semi-decidable) proof procedure for Relational Logic with Herbrand semantics.

Inferential Incompleteness of Herbrand Logic

The axiomatization of Peano Arithmetic in Herbrand Logic completely defines Peano Arithmetic.

If Herbrand entailment were semi-decidable, the set of all true sentences would be enumerable.

Godel's incompleteness theorem tells us that the set of all true sentences of Peano arithmetic is not computably enumerable.

Consequently, there is no complete (semi-decidable) proof procedure for Herbrand Logic.

More disturbing is that there is no complete proof procedure for Relational Logic with Herbrand semantics. Godel's incompleteness theorem tells us that the set of all true sentences of Peano arithmetic is not computably enumerable. Our axiomatization is complete using Herbrand Semantics. If Herbrand entailment were semi-decidable, the set of all true sentences would be enumerable. Consequently, there is no complete (semi-decidable) proof procedure for Relational Logic with Herbrand semantics.

Comparison

Theorem: Any sound proof procedure for First Order Logic is sound for Herbrand Logic.

Even though there is no complete proof procedure, Herbrand logic is *not* weaker. In fact, Herbrand logic is **stronger** than FOL. There are simply *more* things that are true.

We cannot prove them all, but we can prove everything we could prove in First Order Logic; and, by building in induction, we can prove **more** things.

However, this is not as bad as it seems. It turns out that everything that is true under Tarskian semantics is also true under Herbrand semantics, so we can use the same rules of inference. The upshot here is that we lost nothing by switching to Herbrand semantics. In fact, we can add some additional rules of inference. It is not that Herbrand logic is weaker. In fact, it is stronger. There are more things that are true. We cannot prove them all, but we can prove everything we could prove before.

Some may be disturbed by the fact that Herbrand entailment is not semi-decidable. But a similar argument could be leveled against Tarskian semantics. Semi-decidability is not that great either, as a proof procedure might still run forever if the proposed conclusion does not follow from the premises.

Summary

First Order Logic:

- Compact
- Complete proof procedure
- Semi-decidable

Herbrand Logic:

- Not compact
- No complete proof procedure
- Not even semi-decidable

Comparison of Herbrand Logic over First Order Logic:

- Herbrand Semantics simpler and more intuitive
- FOL can be used to describe uncountable worlds
- More things definable in Herbrand Logic
- Greater inferential power in HL but not complete

In order to understand the trade-off between these two approaches to semantics. Eric Kao and I decided to turn things upside down, focussing Herbrand semantics in its own right instead of as a special case of Tarskian semantics. The results were interesting. On the one hand, we realized quite quickly that we no longer have many of the nice features of Tarskian semantics - compactness, inferential completeness, and semidecidability. On the other hand, we found some real benefits as well. Most importantly, Herbrand semantics is conceptually a lot simpler than Tarskian semantics; as a result, it is easier to teach and learn. It has equivalent or greater inferential power. And more things are definable with Herbrand semantics than with Tarskian semantics.

My goal today is to go over some of these results with you, as a way of explaining my newfound enthusiasm for Herbrand semantics and my view that it ought to supplant Tarskian semantics in much CS research and most logic education.

Herbrand Manifesto



<http://logic.stanford.edu/herbrand/manifesto.html>

Manifesto being nailed to the door of the Stanford philosophy department.

