Logic for Linguists:

Lecture 2

Gregory Wilsenach

University of Cambridge

23rd October 2019

Review: Propositional Logic

We have so far discussed propositional logic, which begins with atomic propositions provides us with a calculus for constructing new propositions by stringing together these atoms using logical connectives ($\wedge$, $\vee$, $\neg$, $\rightarrow$). We recall that a formula looked something like:

$$(a \rightarrow b) \rightarrow (\neg a \wedge b).$$

We also introduced some basic concepts and terminology (e.g. the notion of a tautology).

We showed how one might convert a natural language argument into propositional form. We then showed how one might recognise the validity of a propositional formula and establish the validity of an argument.

Today: Predicate (First-Order) Logic

In first-order logic we suppose that atomic propositions have some structure. Roughly, these propositions are built-up from *atomic formulas* strung together by Boolean connectives and quantifiers.

In other words these propositions are of the form

- "Every Greek is mortal."
- "There exists a man named George W. Bush."
- "Every person had or has a mother."
- "For every number x and there exists a number y such that x + y = 0."

Note at this point that in order to formalise these sentences we will not only need existential and universal quantifiers, but also some notion of assigning a constant (e.g. "George W. Bush"), a relation (e.g. "mother of"), and a function (e.g. "+").

Mathematical Background (1)

We will move onto the formal discussion in a moment, but let's begin with a review of a few basic notions from mathematics.

We are all familiar with the notion of a set of objects. For example

- $S := \{\text{Steve}, \text{Mary}, \text{the flowerpot in my room}\}$,
- $\mathbb{N} := \{0, 1, 2, \ldots\}$, and
- $X := \{x : x \text{ is an English sentence}\}$.

We write $X \subseteq Y$ if for every element $x \in X$ we have $x \in Y$ (in other words X is a subset of Y).

Let X and Y be sets. We write $X \times Y$ to denote the *Cartesian product* of X and Y defined as

$$X \times Y := \{(x, y) : x \in X, \text{ and } y \in Y\}.$$

For example, the Cartesian plane, $\mathbb{R} \times \mathbb{R}$, which consists of pairs of real number.

Mathematical Background (2)

Let $X_1, \ldots, X_r$ be a sequence of non-empty sets. We can similarly define

$$X_1 \times \ldots \times X_r := \{(x_1, \ldots, x_r) : x_1 \in X_1 \text{ and } \ldots \text{ and } x_r \in X_r\}.$$

A *relation* is a subset of some Cartesian product. In other words a relation is a set R such that $R \subseteq X_1 \times \ldots \times X_r$. We say that R has *arity* r.

Let's look at a few examples. Let S be the set of people in this room. Let $R := \{(x, y) : x, y \in S \text{ and } x \text{ and } y \text{ are friends}\}$. Let D be the set of desks in this room. Then $T := \{(x, y) : x \in S \text{ and } y \in D \text{ and } x \text{ is sitting behind } y\}$.

A function $f : X \to Y$ associates every element in x with some element $y \in Y$, which we call $f(x)$. We have that for all $x_1, x_2 \in X$ and $x_1 = x_2$ then $f(x_1) = f(x_2)$. We can more formally understand such a function as a relation $f \subseteq X \times Y$.

Syntax (1)

We now define the syntax of first-order logic formally. We first define the notion of a *vocabulary*. A vocabulary τ is a set containing

- constant symbols, usually denoted by the letters $c, d, e, \ldots$,
- relation symbols, usually denoted by the letters $R, T, S, \ldots$, and
- function symbols, usually denoted by the letters $f, g, h, \ldots$.

We associate every relation and function symbol with some natural number, which we call its arity.

We always have available to us a sequence of variables $x_1, x_2, x_2, \ldots$. We often use $x, y, z, \ldots$ to denote variables.

We define the *terms* of first-order logic in the vocabulary τ such that

- every variable is a term,
- every constant symbol in τ is a term, and
- for every function symbol f in τ of arity r and every sequence of terms $t_1, \ldots, t_r$ we have that $f(t_1, \ldots, t_r)$ is a term.

Syntax (2)

The atomic first-order formulas in the vocabulary τ are defined as follows

- for every relation symbol R in τ of arity r and every sequence of terms $t_1, \ldots, t_r$ we have that $R(t_1, \ldots, t_r)$ is an atomic formula, and
- for every pair of terms t and s we have that $t = s$ is an atomic formula.

The first-order formulas in the formulas in the vocabulary τ are defined as follows

- every atomic formula is a formula,
- if ϕ is a formula then $\neg\phi$ is a formula,
- if ϕ and ψ are formulas then $\psi \wedge \phi$, $\psi \vee \phi$, and $\psi \rightarrow \phi$ are formulas, and
- if ϕ is a formula and x is a variable then $\forall x\phi$ and $\exists x\phi$ are formulas.

We say a variable x appears *free* in ϕ if it appears in a scope not bound by a quantifier, otherwise we say x appears *bound*. We write $\phi(\vec{x})$ to denote that the free variables in ϕ appear amongst $\vec{x}$. We say that ϕ is a *sentence* if no variable appears free in ϕ.

Examples

Let's review our examples from earlier.

1. "Every Greek is mortal."
2. "There exists a man named George W. Bush."
3. "Every person had or has a mother."
4. "For every number x and there exists a number y such that x + y = 0."

Let $\tau := \{G, M\}$ consist of two unary (i.e. arity 1) relation symbols meant to denote the properties of being Greek and mortal, respectively. We can write $\phi_1 := \forall x(G(x) \to M(x))$.

Let $\tau := \{M, c\}$ where M is a unary relation meant to denote the property of being a man and c is a constant symbol mean to denote the person George W. Bush. We can write $\phi_2 := \exists x(M(x) \wedge (x = c))$.

Let $\tau := \{P, M\}$ where P is a unary relation meant to denote the property of being a person and M is a binary (arity 2) relation meant to denote the property that a is the mother of b. We can write $\phi_3 := \forall x(P(x) \wedge (\exists y M(y, x)))$.

What Do We Want From a Model?

You might be asking: So we are saying there "exists this" or "for all that" but from where are we getting these objects? This is particularly important as the existential and universal quantifiers don't specifically quantify over a set.

We shall now, as we did for propositional logic, define the notion of a model and define the semantics of a sentence.

Intuitively, a model is some domain of objects together with some interpretation of each symbol in the appropriate vocabulary. A given formula might be true or false depending on the particular model (i.e. depending on the particular way these symbols are interpreted).

The semantics of a formula will be the set of all models for which the formula is true.

Models

Let $\tau := \{c_1, \ldots, R_1, \ldots, f_1, \ldots\}$ be a vocabulary. We write r_i for the arity of the relation symbol R_i and s_i for the arity of the function symbol f_i.

A τ-model (or τ-structure) is a structure $\mathcal{M}$ consisting of the following

- a non-empty set M (called the *universe of* $\mathcal{M}$),
- for each c_i a designated element $c_i^{\mathcal{M}} \in M$,
- for each R_i a relation $R_i^{\mathcal{M}} \subseteq M^{r_i}$, and
- for each f_i a function $f_i^{\mathcal{M}} : M^{s_i} \to M$.

In other words $\mathcal{M}$ interprets all of these symbols that we've been using in our sentences.

Example

Let τ consist of a single binary relation R meant to denote "favourite composer of".

Let $\mathcal{M}$ be the model where

- $M := \{b, h, s\}$, which are meant to denote Beethoven, Handel, and Scarlattii, respectively, and
- $R^{\mathcal{A}} := \{(h, b), (h, s), (s, h)\}$.

In other words $\mathcal{M}$ has as its universe the three denoted composers and encodes the fact that Handel was Beethoven's favourite composer and that Scarlatti and Handel where each other's favourite composers.

Interpreting a Formula in a Model

We are now ready to formally define how a formula can be interpreted in a model.

Let $\tau := \{c_1, \ldots, R_1, \ldots, f_1, \ldots\}$ be a vocabulary. Let $\mathcal{M}$ be a τ-model. Let $t(\vec{x})$ be a first-order term over τ and $\vec{a}$ be a sequence of elements in M. We define the value of t for $\mathcal{M}$ and $\vec{a}$, which we denote by $t^{\mathcal{A}}[\vec{a}]$, as follows

- if t is equal to some x_i in $\vec{x}$ then $t^{\mathcal{A}}[\vec{a}] = a_i$,
- if t is a constant symbol c then $t^{\mathcal{A}}[\vec{a}] = c_i^{\mathcal{M}}$, and
- if t is of the form $f(t_1, \ldots, t_r)$ for some function symbol f of arity r and sequence of terms $t_1, \ldots, t_r$, then

$$t^{\mathcal{A}}[\vec{a}] = f^{\mathcal{M}}(t_1^{\mathcal{A}}[\vec{a}], \ldots, t_1^{\mathcal{A}}[\vec{a}]).$$

Interpreting a Formula in a Model (cont.)

Now let $\phi(\vec{x})$ be a first-order formula in τ and, again, $\vec{a}$ be a sequence of elements in M. We say that $\mathcal{M}$ is a model of ϕ or satisfies ϕ under the assignment $\vec{a}$, and write $\mathcal{M} \models \phi[\vec{a}]$,

- if ϕ is of the form $t_1 = t_2$ for terms t_1 and t_2 then $\mathcal{M} \models \phi[\vec{a}]$ iff $t_1^{\mathcal{A}}[\vec{a}] = t_2^{\mathcal{A}}[\vec{a}]$,
- if ϕ is of the form $R(t_1, \ldots, t_r)$ where R is a relation symbol and $t_1, \ldots, t_r$ are terms, then $\mathcal{M} \models \phi[\vec{a}]$ iff $R^{\mathcal{M}}(t_1^{\mathcal{A}}[\vec{a}], \ldots, t_1^{\mathcal{A}}[\vec{a}])$,
- if ϕ is of the form $\neg\psi$ then $\mathcal{M} \models \phi[\vec{a}]$ iff it is not the case that $\mathcal{M} \models \psi[\vec{a}]$,
- if ϕ is of the form $\psi \wedge \theta$ then $\mathcal{M} \models \phi[\vec{a}]$ iff $\mathcal{M} \models \psi[\vec{a}]$ and $\mathcal{M} \models \theta[\vec{a}]$,
- if ϕ is of the form $\psi \vee \theta$ then $\mathcal{M} \models \phi[\vec{a}]$ iff $\mathcal{M} \models \psi[\vec{a}]$ or $\mathcal{M} \models \theta[\vec{a}]$ (or both),
- if ϕ is of the form $\exists y \psi$ then $\mathcal{M} \models \phi[\vec{a}]$ iff there is some $b \in M$ such that $\mathcal{M} \models \psi[\vec{a}, b]$ (by which I mean $\vec{x}$ is assigned to $\vec{a}$ and y is assigned to b), and
- if ϕ is of the form $\forall y \psi$ then $\mathcal{M} \models \phi[\vec{a}]$ iff for every $b \in M$ we have $\mathcal{M} \models \psi[\vec{a}, b]$.

Theories and Entailment

Let ϕ be a sentence. We let $\text{mod}(\phi)$ denote the set of all models $\mathcal{M}$ such that $\mathcal{M} \models \phi$. We consider $\text{mod}(\phi)$ to be the semantics of ϕ.

We say that ϕ is a tautology if $\text{mod}(\phi)$ is the set of all model over the vocabulary of ϕ.

We call a set of sentences (all over the same vocabulary) a *theory*. If T is a theory we write $T \models \mathcal{M}$ to denote that $\mathcal{M}$ satisfies all of the formulas in T and write $\text{mod}(T)$ to denote the set of models that satisfy all of the sentences in T. We say that a theory U *axiomatises* T if $\text{mod}(T) = \text{mod}(U)$. We say that T is *consistent* if $\text{mod}(T)$ is non-empty.

We can get from this a purely semantic definition of entailment. We say that T *entails* a sentence ϕ if $\text{mod}(T) \subseteq \text{mod}(\phi)$.

Summary

What you should get from this lecture

- familiarity with basic mathematical concepts and notation that we
- the intuition behind first-order logic,
- the formal syntax and semantics of the logic,
- the notion of a model and how they can be used to define the semantics of first-order logic.

This lecture has been a lot of definitions and examples. Next week we will begin to look at proof theory and how these ideas connect up very neatly with the semantics notions discussed today. We will also look at extensions of first-order logic with generalised quantifiers.

Interesting Reading

Notes on first-order logic:

- https://www.cs.ox.ac.uk/people/james.worrell/lecture9-2015.pdf
- https://www.cs.utexas.edu/ mooney/cs343/slide-handouts/fopc.4.pdf

Discussions on converting natural language to first-order logic:

- https://cs.nyu.edu/faculty/davise/ai/folguide.pdf

Here is a website which automates the conversion from natural language to first-order logic:

- http://attempto.ifi.uzh.ch/race/