Semantics for Predicate Logic: Part I

Spring 2004

1 Interpretations

A sentence of a formal language (e.g., the propositional calculus, or the predicate calculus) is neither true nor false. By definition, an *interpretation* of a sentence of a formal language is a specification of enough information to determine whether that sentence is true or false. It follows from this and our understanding of how predicate logic works than a predicate logic interpretation must state:

1. The **domain** of quantification; i.e., what the quantifiers (or variables) range over. The domain must be nonempty. (The domain is sometimes also called the **universe of discourse**.)

2. For each **name** (i.e., $m, n, o, \ldots$), which object in the domain it denotes. If m denotes α, we write $m \mapsto \alpha$. (A name denotes only one object; but different names can denote the same object.)

3. For each **predicate** (i.e., $F, G, H, \ldots$), those objects in the domain to which it truly applies.

We need to say a bit more about how one can specify the interpretation of a predicate (such as $F, G, \ldots$). We give three of these ways below.

1.1 Interpreting predicates via descriptions

Since we already know which things natural language predicates apply to, we can interpret a predicate calculus (PC) predicate by identifying it with some natural language predicate. For example:

1. If the universe of discourse consists of natural numbers $1, 2, 3, \ldots$, we could specify:
$$Fx \equiv x \text{ is even.}$$
2. If the universe of discourse consists of currently living people, we could specify:
$$Fx \equiv x \text{ is male.}$$
In this case, $(\exists x)Fx$ is true, and $(x)Fx$ is false.

1.2 Interpreting predicates via tables

The "extension" of a predicate (relative to an interpretation) is simply the collection of all things to which it truly applies. If the universe of discourse happens to be finite, and indeed rather small, the extension of predicates can be given by means of a table. Here we put "+" in the table if the predicate applies to the object and "−" if it doesn't.

Example: Give an interpretation for a problem involving the predicate letters F, G, H, and the names m, n, o.

$$\text{Domain} = \{\alpha, \beta, \gamma\}$$

$$m \mapsto \alpha, \quad n \mapsto \alpha, \quad o \mapsto \beta$$

	F	G	H
α	+	−	−
β	+	+	−
γ	+	−	−

1.3 Interpreting predicates via set-theoretic equalities

In the domain $\{2, 3, 4, 5\}$, the extension of the predicate "is even" is $\{2, 4\}$. If the domain is finite, it is sometimes easiest just to write down equalities that give the extension of the predicate letters. For this purpose, we use "Ext(F)" to denote the extension of the predicate F. Thus, here is another way to write down the interpretation we gave in the previous section.

$$\text{Domain} = \{\alpha, \beta, \gamma\}$$

$$m \mapsto \alpha, \quad n \mapsto \alpha, \quad o \mapsto \beta$$
$$\text{Ext}(F) = \{\alpha, \beta, \gamma\}$$
$$\text{Ext}(G) = \{\beta\}$$
$$\text{Ext}(H) = \emptyset$$

(This last equality means that the extension of H is empty; i.e., H does not truly apply to anything in the domain.)

2 Truth is relative (to an interpretation)

We now give a completely precise definition of what it means for a sentence ϕ to be true relative to an interpretation $\mathfrak{I}$. We have to look at three cases:

Simple sentences: Fm is true on the interpretation $\mathfrak{I}$ just in case $\mathfrak{I}$ assigns m to some object in the extension of F. It is false otherwise.

Universally quantified sentences: $(x)\phi(x)$ is true on the interpretation $\mathfrak{I}$ just in case $\phi(a)$ is true no matter which element of the domain of $\mathfrak{I}$ is assigned to a.

Existentially quantified sentences: $(\exists x)\phi(x)$ is true on the interpretation $\mathfrak{I}$ just in case we can assign the name a to some element of the domain of $\mathfrak{I}$ such that $\phi(a)$ comes out true.

Truth functional compounds: The truth-value of a compound sentence ϕ is determined by the truth-values of its component parts. (We already know how it's determined.)

Example: Let $\mathfrak{I}$ be the interpretation with domain $\{\alpha, \beta, \gamma, \delta\}$ and:

$$m \mapsto \alpha, \quad n \mapsto \alpha, \quad o \mapsto \beta$$
$$\text{Ext}(F) = \{\alpha, \beta, \gamma\}$$
$$\text{Ext}(G) = \{\gamma, \delta\}$$
$$\text{Ext}(H) = \emptyset$$

Then, we have the following:

1. $(\exists x)(Fx \,\&\, Gx)$ is true on this interpretation, since γ lies in the extension of both F and G.
2. $(x)(Fx \rightarrow Gx)$ is false on this interpretation, since α is an F that is not a G.
3. $(x)(Fx \vee Gx)$ is true on this interpretation, since every object in the domain lies in the extension of either F or G.
4. $(x)(Hx \rightarrow Gx)$ is true on this interpretation. Since no object in the domain is an H, it is trivially true that every H in the domain is also a G.

3 Semantic properties of sentences, relations between sentences

Here we list a bunch of definitions that relate everyday notions to precise counterparts in the semantics of our formal language.

Definition. If a sentence ϕ is true relative to an interpretation $\mathfrak{I}$, we write $\models_{\mathfrak{I}} \phi$.

Definition. If ϕ is true on every interpretation, then we say that ϕ is a **logical truth**. In this case we write $\models \phi$, with no subscript. If ϕ is false on some interpretation, then we say that ϕ is **falsifiable**, and we sometimes write $\not\models \phi$ (which simply means that ϕ is not a logical truth).

Definition. A sentence false on every interpretation is **inconsistent**. A sentence true on some interpretation is **consistent**. Also, a collection of sentences such that at least one is false on every interpretation is said to be **inconsistent**, and such a collection is said to be **consistent** if there is an interpretation on which all sentences in the collection are true.

Definition. Two sentences ϕ, ψ are **logically equivalent** if they have the same truth value relative to every interpretation. In this case we write $\phi \Leftrightarrow \psi$, with no subscript.

Definition. A collection Γ of sentences **logically implies** a sentence ϕ if there is no interpretation that makes all of Γ true while making ϕ false (there is no "counterexample"). In that case, the argument with premises Γ and conclusion ϕ is **valid**, and we write $\Gamma \models \phi$.

Example. $(x)(Fx \vee -Fx)$ is a logical truth, because in any interpretation, each object in the domain is either in the extension of F or it isn't.

Example. $(\exists x)(Fx \,\&\, -Fx) \rightarrow Gm$ is a logical truth, because $(\exists x)(Fx \,\&\, -Fx)$ is inconsistent, and a conditional is true whenever its antecedent is false.

4 How to answer questions about semantic properties and relations

An answer to a problem requiring the presentation of an interpretation is best seen as having three parts, as follows. (1) State the domain of your proposed interpretation. (2) Present (using one of the above methods) the interpretation of the names and predicate symbols. (3) State the truth values of the various sentences, defend your claim that they have those truth values, and be explicit as to how this information solves the problem you began with. We now consider a couple of examples of solved problems.

Problem: Show that (a) $(\exists x)Fx \rightarrow (x)Fx$ is falsifiable (not logically true).

1. Domain = $\{\alpha, \beta\}$.
2. Ext(F) $= \{\alpha\}$.
3. (a) is false on this interpretation: Since α is in the extension of F, $(\exists x)Fx$ is true on this interpretation. However, since β is not in the extension of F, $(x)Fx$ is false on this interpretation. Therefore (by truth tables) $(\exists x)Fx \rightarrow (x)Fx$ is false on this interpretation. Since (a) is false on some interpretation, (a) is falsifiable. □

Problem: Show that the same sentence (a) is consistent (true on some interpretation).

1. Domain = $\{\alpha\}$.
2. Ext(F) $= \{\alpha\}$.
3. (a) is true on this interpretation: Since every object in the domain is an F, $(x)Fx$ is true on this interpretation. Thus, truth tables tell us that $(\exists x)Fx \rightarrow (x)Fx$ is true on this interpretation. Since (a) is true on some interpretation, (a) is consistent. □

Solving a semantic problem (e.g., "Does ϕ logically imply ψ?") always requires you to answer the following question:

> *Is there an interpretation of a certain sort?*

(And, if there is such an interpretation, you may be asked to display it.) In general, it is a *very hard* task to answer these sorts of questions.[1] However, when quantifiers are not nested within each other, it turns out to be *easy* to answer these sorts of questions.[2] By "easy", we mean that there are algorithms that you can employ that will always provide you with the right answer to the question, in a finite amount of time. We will learn two such algorithms.

[1] The claim that these sorts of questions are "hard" to answer is not a subjective claim: it has been proved (in fact, by a logician at Princeton) that no one can ever write an algorithm that reliably answers these sorts of questions.

[2] As shown by Lowenheim.