

A rule-to-rule semantics for a simple language

PHIL 43916
August 29, 2012

1. Sentences, names, and V_i 's.....	1
2. Transitive verbs.....	2
3. Sentence operators and connectives.....	4
4. Some examples.....	5
5. Relativizing to circumstances.....	6
6. Entailment and contradiction.....	7

1. SENTENCES, NAMES, AND V_i 'S

We know that $\llbracket S \rrbracket$ is a truth-value. So what are $\llbracket N \rrbracket$ and $\llbracket V_i \rrbracket$?

The semantic value of a name is the object for which the name stands. So, for example, $\llbracket \text{Pavarotti} \rrbracket = \text{Pavarotti}$, $\llbracket \text{Sophia Loren} \rrbracket = \text{Sophia Loren}$, and so on.

The semantic value of an intransitive verb will be a set of individuals. So $\llbracket \text{is boring} \rrbracket$ will be the set of individuals that are boring, $\llbracket \text{is cute} \rrbracket$ will be the set of cute individuals, and so on. We will refer to the set of boring individuals using the notation

$$\{x: x \text{ is boring}\}$$

Intuitively: the set of all x 's which are such that x is boring.


Note that so far we have said what the semantic values of sentences, names, and intransitive verbs are, but have *not* provided a rule for determining the semantic value of a sentence consisting of a N and a V_i on the basis of the semantic values of the latter. That we can do using (to follow the numbering in the text) rules 31 (a) and (e):¹

$$(a) \llbracket S N VP \rrbracket = 1 \text{ iff } \llbracket N \rrbracket \in \llbracket VP \rrbracket \text{ and } 0 \text{ otherwise}$$

Let's pause on this rule for a second. What does it say? Consider some examples of sentences consisting of names and V_i 's. What does it indicate about the truth conditions of sentences of this sort?

¹ Here I'm simplifying by ignoring the relativization to circumstances. We un-simplify in §5 below.

But note that rule (a) by itself, plus the above remarks about the semantic values of names, sentences, and V_i 's, does *not* tell us how to derive the truth or falsity of a sentence. Consider our tree diagram for 'Pavarotti is boring':



We know what the semantic values of the N and the V_i are, and we know how to figure out the semantic value of the S once we have the semantic value of the N and the VP; but so far we have no way of determining the semantic value of the VP.

You might think that this is pretty obvious: surely, in this case, $\llbracket \text{VP} \rrbracket = \llbracket \text{Vi} \rrbracket$. And this is correct. But we need some rule to state clearly when we're allowed to 'pass up' the semantic value of an expression from a child node to its parent. That is the point of rule 31 (e):


(e) If A is a category and a is a lexical entry or category and $\Delta = [A \ a]$, then $\llbracket \Delta \rrbracket = \llbracket a \rrbracket$.

Here we use ' $[A \ a]$ ' to mean 'the tree dominated by A, whose only child is a.' More generally, ' $[A \ b \ c]$ ' means 'the tree dominated by A, whose only children are b and c.'

We can think of the process of determining the semantic value — i.e. truth value — of a sentence as working in steps. First, enter the semantic values of the leaves. Then, we consult the rules of our semantics to determine the semantic values of the parents of the leaf nodes, continuing to work from child to parent until we have assigned a semantic value — 1 or 0 — to the S.

2. TRANSITIVE VERBS

Now consider a sentence like our example from last time,



This is still a sentence of the form $[_S N VP]$, so rule (a) above should apply. This means that, as above, the semantic value of our VP must be a set — in this case, it will be

$$\{x: x \text{ likes Sophia Loren}\}$$

One idea would be to simply add this fact about $\llbracket \text{likes Sophia Loren} \rrbracket$ to our semantic theory. Why might this be a bad idea?

Better would be to give a rule for determining $\llbracket \text{likes Sophia Loren} \rrbracket$ on the basis of $\llbracket \text{likes} \rrbracket$ and $\llbracket \text{Sophia Loren} \rrbracket$. We already know that $\llbracket \text{Sophia Loren} \rrbracket = \text{Sophia Loren}$. So our question is: what is $\llbracket \text{likes} \rrbracket$?

Remember that the semantic value of ‘boring’ was the set of boring things. So one might think, by extension, that the semantic value of ‘likes’ is a set of sets: the set of sets of things which are such that one likes the other:

$$\llbracket \text{likes} \rrbracket = \{\{x,y\}: x \text{ likes } y\}$$

What would be wrong with this? (Keep in mind that if sets S_1, S_2 have the same members, then $S_1 = S_2$; so, in particular, $\{a,b\} = \{b,a\}$.)

Better to take the semantic value of an intransitive verb to be a set of *ordered pairs*, namely

$$\llbracket \text{likes} \rrbracket = \{\langle x,y \rangle: x \text{ likes } y\}$$

This leaves open the possibility that $\langle \text{Pavarotti}, \text{Sophia Loren} \rangle$ will be an element of $\llbracket \text{likes} \rrbracket$, whereas $\langle \text{Sophia Loren}, \text{Pavarotti} \rangle$ will not.

But now we are in a situation like the one above: we have an assignment of semantic values to ‘likes’ and ‘Sophia Loren,’ but we need an extra rule to tell us how to get from these semantic values to the semantic value of the complex VP ‘likes Sophia Loren.’

That is the point of rule 31 (d):

$$(d) \llbracket [_{VP} V_t N] \rrbracket = \{x: \langle x, \llbracket N \rrbracket \rangle \in \llbracket V_t \rrbracket\}$$

What does this say? What does it imply about the case of ‘likes Sophia Loren’?

3. SENTENCE OPERATORS AND CONNECTIVES

We’re almost done with the semantics for our simple language: all that’s left is to explain the semantic values of sentence operators and connectives.

Consider first our lone sentence operator, ‘it is not the case that.’ We know that our language permits sentences of the form $[_S \text{neg } S]$, so $\llbracket \text{neg} \rrbracket$ must be something which combines with a truth-value — which is $\llbracket S \rrbracket$ — to give us a truth-value.

A natural choice for $\llbracket \text{neg} \rrbracket$ is a *function*. A function is a relation between a set of inputs — the function’s *arguments* — and a set of outputs — its *values* — which has the property that any argument is related to exactly one value.

A familiar example of a function is addition. Its arguments are pairs of numbers, and its values are individual numbers — the sum of the arguments. Addition is a function, rather than some other sort of relation, because it is never the case that, for any a, b , $a+b=c$ and $a+b=d$ for $c \neq d$.

What sorts of things should the arguments and values of $\llbracket \text{it is not the case that} \rrbracket$ be? Which arguments should get mapped to which values?

We write this as:

$$\llbracket \text{it is not the case that} \rrbracket = \begin{bmatrix} 1 \rightarrow 0 \\ 0 \rightarrow 1 \end{bmatrix}$$

Now, as before, in addition to specifying the semantic value of $\llbracket \text{neg} \rrbracket$, we need an extra rule telling us how to compute the semantic value of $[_S \text{neg } S]$ on the basis of $\llbracket \text{neg} \rrbracket$ and $\llbracket S \rrbracket$. That is rule 31 (c):

$$(c) \llbracket [_S \text{neg } S] \rrbracket = \llbracket \text{neg} \rrbracket (\llbracket S \rrbracket)$$

This follows the standard notation for functions, where we express the claim that ‘function f applied to argument a has value v ’ as ‘ $f(a)=v$ ’ — as in ‘ $+(2,3)=5$.’

How would you extend this treatment of ‘it is not the case that’ to ‘and’ and ‘or’?

Since these two members of the category *conj* combine with two sentences to form a sentence, it is natural to treat these as functions from pairs of truth-values to truth-values. In particular:

$$\begin{aligned} \llbracket \text{and} \rrbracket = [& \langle 1,1 \rangle \rightarrow 1 \\ & \langle 1,0 \rangle \rightarrow 0 \\ & \langle 0,1 \rangle \rightarrow 0 \\ & \langle 0,0 \rangle \rightarrow 0] \end{aligned}$$

$$\begin{aligned} \llbracket \text{or} \rrbracket = [& \langle 1,1 \rangle \rightarrow 1 \\ & \langle 1,0 \rangle \rightarrow 1 \\ & \langle 0,1 \rangle \rightarrow 1 \\ & \langle 0,0 \rangle \rightarrow 0] \end{aligned}$$

and we derive the semantic values of sentences involving a conj using rule 31 (b):

$$(b) \llbracket [s \text{ S1 conj S2}] \rrbracket = \llbracket \text{conj} \rrbracket(\langle \llbracket \text{S1} \rrbracket, \llbracket \text{S2} \rrbracket \rangle)$$

It is worth pausing for a moment over the case of ‘or.’ It might seem that, whatever is true of our simple language, the semantic value given to ‘or’ can’t possibly be the semantic value of the English word ‘or.’ For consider a sentence like ‘Jim will go to bed early or Jim will fail the exam.’ Surely this means that *exactly* one — not *at least one* — of the two sentences connected by ‘or’ is true.

This is a good case to bring up the distinction between what *sentences* mean and what *speakers* mean by using those sentences. What we’re trying to capture is, in the first instance, facts about sentence meaning.

Some evidence that ‘or’ in English has $\llbracket \text{or} \rrbracket$ as its semantic value is given by the way that ‘or’ sentences behave as a part of more complex discourses. There is, for example, no contradiction in saying

Jim will go to bed early or Jim will fail the exam — indeed, he’s not very bright, so he might well do both.

And the sentence

It is not the case that Jim will go to bed early or Jim will fail the exam.

the sentence seems to be false, not true, if Jim does both.

4. SOME EXAMPLES

Let’s work through some examples, and try to derive the truth-values of some sentences using the rules of our semantic theory. To do this we will have to be clear about exactly what the semantic values of our V_i ’s and V_t ’s are; we know that $\llbracket \text{is boring} \rrbracket =$ the set of boring things, but we don’t know what things are in that set. So let’s suppose that:

$\llbracket \text{is boring} \rrbracket = \{\text{James Bond, Pavarotti}\}$
 $\llbracket \text{is cute} \rrbracket = \{\text{Pavarotti}\}$
 $\llbracket \text{likes} \rrbracket = \{\langle \text{Sophia Loren, Pavarotti} \rangle, \langle \text{James Bond, Pavarotti} \rangle\}$

And consider the following sentences:

It is not the case that Sophia Loren is cute.
 Sophia Loren is boring or Pavarotti is cute.
 It is not the case that James Bond is boring and Pavarotti is cute.

In each case, our theory allows us to derive semantic values — truth-values — for the sentences on the basis of the semantic values of the simple terms, plus facts about how they are combined, plus our semantic rules for combining expressions to form complex expressions.

5. RELATIVIZING TO CIRCUMSTANCES

So far, in introducing our semantic theory, I've suppressed the need to relativize semantic values to different circumstances of evaluation; I've been talking, e.g., about $\llbracket \text{Pavarotti is cute} \rrbracket$ but not $\llbracket \text{Pavarotti is cute} \rrbracket^v$.

It's now time to re-introduce this. For some expressions we've discussed, this makes no difference. Ignoring some complications to which we will return later, the semantic values of names and of connectives will be the same with respect to every circumstance of evaluation; for any v , $\llbracket \text{Pavarotti} \rrbracket^v = \text{Pavarotti}$, and $\llbracket \text{neg} \rrbracket^v = [1 \rightarrow 0, 0 \rightarrow 1]$.

But this is not true of our V_i 's and V_t 's — can you see why?

However, the modification of their semantic values which this requires is, in one sense, not so great. Rather than the simple

$\llbracket \text{boring} \rrbracket = \{\text{the } x: x \text{ is boring}\}$

we will now have

$\llbracket \text{boring} \rrbracket^v = \{\text{the } x: x \text{ is boring in } v\}$

The important thing about this change, for our purposes, is that it now allows us to derive not just the truth-values of sentences of our language, but their truth-conditions — i.e., their truth-value with respect to different circumstances of evaluation. And this is important because, plausibly, this is what competent language users know about sentences they understand — not *whether* they are true or false, but the conditions under which they would be true or false.

This requires a modification of our semantic rules 31 (a)-(e) — in each case, simply replace every reference to a semantic value $\llbracket x \rrbracket$ with a relativized $\llbracket x \rrbracket^v$ and everything else remains the same. The relativized rules — which are the versions in the text — are:

$$(a) \llbracket [S N VP] \rrbracket^v = 1 \text{ iff } \llbracket N \rrbracket^v \in \llbracket VP \rrbracket^v \text{ and } 0 \text{ otherwise}$$

$$(b) \llbracket [S S1 conj S2] \rrbracket^v = \llbracket conj \rrbracket^v(\langle \llbracket S1 \rrbracket^v, \llbracket S2 \rrbracket^v \rangle)$$

$$(c) \llbracket [S neg S] \rrbracket^v = \llbracket neg \rrbracket^v(\llbracket S \rrbracket^v)$$

$$(d) \llbracket [VP V_t N] \rrbracket^v = \{x: \langle x, \llbracket N \rrbracket^v \rangle \in \llbracket V_t \rrbracket^v\}$$

$$(e) \text{ If } A \text{ is a category and } a \text{ is a lexical entry or category and } \Delta = [A a], \text{ then } \llbracket \Delta \rrbracket^v = \llbracket a \rrbracket^v.$$

Using these rules, let's derive $\llbracket \text{Pavarotti is boring} \rrbracket^v$.

Looking at 31 (a)-(e), you can see why in the text this theory is referred to as an example of 'rule-to-rule' semantics. 31 (a)-(e) mirror the syntactic rules 21 (a)-(e) of our language. Each of those syntactic rules gives one type of case when it is possible in our language to grammatically combine expressions of two types. For any such case, we then need to add to our semantics a rule which tells us how, in cases of that type, the relevant semantic values combine to give us the semantic value of the complex expression. For each syntactic rule, we have a corresponding semantic rule.

6. ENTAILMENT AND CONTRADICTION

Another benefit of our relativization of semantic values to circumstances is that it enables us to define *entailment*.

To a first approximation, one sentence S1 entails another sentence S2 if and only if, necessarily, if S1 is true, then S2 is true — or, to put the same point another way, S1 entails S2 if and only if the truth of S1 guarantees the truth of S2.

Often, just on the basis of understanding sentences, and without knowing whether either is true, we can see that one sentence entails another. For example, many have claimed that any competent speaker can see that if

Pavarotti is boring and James Bond is cute.

is true, so must be

Pavarotti is boring.

If we relativize semantic values to circumstances, then we can define entailment as a relation between individual sentences as follows:

S1 entails S2 iff for any v , if $\llbracket S1 \rrbracket^v=1$, then $\llbracket S2 \rrbracket^v=1$.

In a related way, we can define the relation of *contradiction* between sentences:

S1 entails S2 iff for any v , if $\llbracket S1 \rrbracket^v=1$, then $\llbracket S2 \rrbracket^v=0$.

And in many cases we can use our semantic theory to prove that one sentence entails (or contradicts) another. Consider the example above, about Pavarotti and James Bond. How, using our semantic rules, could you prove that the first of these sentences entails the other?

Next, try to prove that

It is not the case that Pavarotti is boring or James Bond is cute

(on one interpretation) contradicts

Pavarotti is boring.

These facts about entailment (and contradiction) are connected to the question of how we can tell that whether a semantic theory for a language like English is correct. What, in semantics, is supposed to play the role of experimental results in physics? Many have thought that the answer is, at least in part, given by the following two tests:

- Competent speakers of a language know the truth conditions of sentences of their own language. The correct semantic theory should therefore assign truth conditions to those sentences which fit the beliefs of competent speakers.
- Competent speakers of a language know when one sentence of their language entails (or contradicts) another. A semantic theory should explain this ability by providing an explanation, in something like the above way, of these entailment relations.

In the end, we'll see that, plausibly, no theory can quite meet these tests — every theory makes some surprising claims about truth conditions, and no theory can explain every entailment. But these at least provide reasonable starting points for evaluating semantic theories.