Dissertation proposal: Overview for nonlinguists

Lucas Champollion
University of Pennsylvania

February 9, 2009

Abstract

This document describes my dissertation proposal, "Aspect, plurality and quantification", for a nonlinguist audience for the purposes of Penn's SAS Dissertation Completion Fellowship. The dissertation proposal is available online on my website, `http://www.ling.upenn.edu/~champoll/`, as well as a handout with a description targeted at a linguist audience.

The supervisor of this dissertation is Dr. Cleo Condoravdi from the Palo Alto Research Center (PARC) and Stanford University; the chair of the dissertation committee is Dr. Aravind K. Joshi from the University of Pennsylvania.

1 Introduction

My dissertation proposal contributes to the field of *linguistics*, the study of human language, more specifically *theoretical linguistics*, which concerns itself with modeling what people must know in order to speak and understand language – as opposed to, say, the study of particular languages. The proposal is situated in the subfield of *semantics*, the scientific study of meaning in language. The proposal uses the methods and tools of *formal semantics*, in which analytical techniques from logic are adapted to human language in order to describe what individual words mean, and how the meaning of a sentence is derived from the meaning of the words and larger parts of which it consists. Some aspects of the proposal relate to the neighboring subfield of *syntax*, the study of the structure of sentences and of the principles to which speakers of a language unconsciously conform when they say or write sentences. This section introduces a few key concepts necessary to understand what my proposal is about. My proposal itself is described in the next section.

Semantic theories make predictions on whether certain sentences are true, false or nonsensical. Syntactic theories make predictions on whether certain sentences are grammatical or ungrammatical. These theories can be empirically tested by comparing these predictions with the intuitions, or "judgments", of native speakers of a language. A common convention in linguistic texts is to mark sentences which native speakers judge to be unacceptable (that is, either ungrammatical or nonsensical) with a star (*). For example, sentence (1) below is grammatical, while (2) is unacceptable:

(1) I made up the example you're reading right now.

(2) *I made up it.

The contrast between (1) and (2) is a typical example of a fact for which linguists account using syntactic theories.

Many sentences are *ambiguous*, or can be understood in more than one way. Linguists call these different ways of understanding a sentence *readings*. The following sentence, for example, has at least two different readings:

(3) Twenty families own fifty of the nation's major companies.

This sentence could be understood to mean that there was a total of twenty families who between them own a total of fifty of the nation's major companies (so that the sentence mentions fifty companies in total). This is called a *cumulative reading*. It could also mean that twenty families *each* own a potentially different set of fifty companies (so that there might be a total of up to 1000 companies involved). This is called a *distributive reading*. In that reading, the phrase "fifty companies" does not necessarily have a constant meaning: it picks out a potentially different set of fifty companies per family. Linguists describe this phenomenon by saying that the phrase "fifty companies" *depends on* the phrase "twenty families". The sentence can be made unambiguously distributive by adding "each" to it: Sentence (4) below only has a distributive reading. Whenever a reading can be paraphrased using "each", this is an indicator that the reading is distributive.

(4) Twenty families each own fifty of the nation's major companies.

Semanticists have identified other parts of language in which distributivity is also at work, for example, sentences containing phrases like *for a year*. These phrases are called *for*-adverbials. Certain sentences that would be otherwise grammatical are judged unacceptable by many speakers of English when they are followed by a *for*-adverbial. Linguists have tried since the 1950s to understand exactly which sentences these are and why they become unacceptable. The examples in (5) illustrate the problem:

(5) a. *John built a house for a year.
b. *John built two houses for a year.
c. John built houses for a year.
d. John owned a house for a year.
e. John owned two houses for a year.
f. John owned houses for a year.

To see why this problem is hard, note that the reason why (5a) and (5b) are unacceptable cannot be traced to the words that occur in these sentences, because each of these words also occurs in at least one of sentences (5c-f), and these sentences are grammatical. So whatever property is responsible for the contrast between (5a-b) and (5c-f), it must be a property of larger parts that consist of more than one word – possibly the whole sentence. At the same time, the pattern in (5) is not random, but it correlates with an important property of the descriptions of events in these sentences:

Each of the states or events described in (5c-f) – building houses, owning a house, etc. – can be thought as being made up of smaller parts to which the same description applies. For

example, suppose that John owned a house from 2000 to 2005. Then it follows that John also owned a house from 2000 to 2001. We might think of that as a part of John's owning the house from 2000 to 2005. By contrast, the events described in (5a-b) don't have any smaller parts that match their description. For example, suppose that John built a house from 2000 to 2005. Then it does *not* follow that he built a house from 2000 to 2001. (We might at best say that he *was building* a house during that period, but note that in the previous case, we felt no need to reformulate "John owned a house" to "John was owning a house".) So there are no smaller parts of building two houses that match the description of building two houses.

This distinction has been the object of linguistic and philosophical study starting with Aristotle (Metaphysics 1048b). Nowadays, it is usually referred to as *telicity*: Whenever a telic description matches an event, there are no parts of that event which match the same description. (5a-b) contain telic descriptions, while (5c-f) do not. Now, (5a-b) are also incompatible with *for*-adverbials, while (5c-f) are not. This has prompted semanticists to claim that *for*-adverbials are incompatible with telic descriptions. One way of formulating this claim in detail is that *for*-adverbials require event descriptions to be true not just at the time period that they name, but also at each shorter time period. For example, the sentence "John built a house for a year" could only be true if it John built a house at each shorter time period during that year, but this is impossible. By contrast, the sentence "John owned a house for a year" is true if John owned a house at each shorter time period during that year.

Since that paraphrase contains "each", the sentence is distributive. Specifically, "owned" *depends* on the phrase "for a year", just as in "twenty families each own fifty of the nation's major companies", the phrase "fifty of the nation's major companies" depends on the phrase "twenty families".

2 This proposal

The main objective of this proposal is to find out how many kinds of distributivity exist in human language.

My hypothesis is that there is only one kind of distributivity in language. In other words, "distributivity" is not just a descriptive label by which linguists refer to different phenomena that look superficially similar, but a common factor responsible for all these phenomena. My hypothesis is based on the fact that I found that distributive sentences across different parts of language have several properties in common besides being distributive, and on more specific findings that result from applying the tools of formal semantics to various instances of distributivity.

What does it mean to say that something is not just a "descriptive label" but a "common factor"? To make this more understandable, here is a parallel with scientific classification in biology, a method by which biologists categorize species of organisms and group them together to larger categories or taxonomic ranks, such as classes and orders. Now, certain categories of species are commonly grouped together by tradition or by the conventions of human language, but biologists do not consider them taxonomic ranks. For example, "shrubs" are small woody plants with multiple stems, "bugs" includes most insects but also spiders and centipedes, but biologists do not recognize these categories. "Insects" and "spiders", by contrast, are recognized as taxonomic ranks. For a biologist, to speak of

"shrubs" as a taxonomic rank would be making a false scientific claim (namely that all shrubs are part of the same rank, have common DNA, etc.). But a biologist who wants to refer to small woody plants with multiple stems might still use the word "shrub", as long as it is clear from the context that he or she is not making a scientific claim of any sort. In this case, the biologist is using the word "shrub" as a descriptive label.

Suppose now that there is only one kind of distributivity. Then distributivity is not just a descriptive label, and theories of semantics should model it as just one phenomenon. The specific framework in which I work is formal semantics. Moreover, I consider implications of my work on linguistic theories which have also been formulated in the framework of formal semantics. In formal semantics, the meaning of each word and phrase is specified as a mathematical expression. The meaning of larger phrases and sentences is defined as the result of combining these expressions using a small number of rules. The expressions that represent sentences describe what must be the case for these sentences to be true or false. This way, formal semantics makes predictions regarding the conditions under which people understand sentences as being literally true or false.

In formal semantics, every word, as well as every phenomenon like distributivity or plurals, is represented as a mathematical expression, no matter in which sentence it occurs. So the claim that there is only one kind of distributivity translates to the claim that the mathematical expression representing distributivity (the "distributivity operator", in semantic jargon) is the same in any sentence involving distributivity.

In the previous section, I have described distributivity in the case of sentences like "twenty families own fifty of the nation's major companies" (on the distributive reading) and sentences like "John owned a house for a year". In my dissertation proposal, for each of these sentences, I identify the distributivity operator, and I show that it must be the same in both types of sentences, or else formal semantics will make empirically wrong predictions about the conditions under which people understand these sentences as true or false. I show that the difference between the cumulative and the distributive reading of "twenty families own fifty of the nation's major companies" can be predicted by manipulating a single variable: the presence or absence of the distributivity operator. With respect to *for*-adverbials, I show that the specific properties of the distributivity operator provide a novel argument in a long-standing dispute on the precise meaning of these adverbials, and the precise reason why they are incompatible with telic descriptions. This dispute concerns the question whether (i) *for*-adverbials require event descriptions to be true both of a given event and of the parts of which it is made up (in which case "John built a house" is incompatible with "for a year" because no part of an event described as building a house can itself be described as building a house), or whether (ii) *for*-adverbials require event descriptions to be true at every shorter time period throughout the time mentioned by the *for*-adverbial (in which case "John built a house" is incompatible with "for a year" because it will not generally be true that John built a house in every time period shorter than the year in question). My findings are that only (ii) is compatible with the properties of the distributivity operator; moreover, only (ii) is compatible with the properties of cumulative readings.

In addition, my proposal has a number of methodological goals. I plan to integrate existing formal semantic theories for cumulative and distributive readings and for phenomena similar to *for*-adverbials and to check their mutual compatibility. I generalize existing theories that account for the cumulative and distributive readings of sentences with two noun phrases, like "twenty families own fifty companies", so that they also apply to sentences with

more than two noun phrases, like "twenty families gave fifty companies two orders". The point is that since any two noun phrases can stand either in a cumulative or in a distributive relation to each other, sentences with more than two noun phrases can have mixed readings, where a noun phrase stands in a cumulative relation with another one and in a distributive relation with a third. These mixed readings cannot occur in sentences with only two noun phrases (there are not enough noun phrases to go around), so their existence is not always correctly predicted by theories that restrict their attention to the case of two noun phrases.

Finally, my dissertation proposal considers the question whether there is any difference between the meaning of nouns and verbs – the noun "love" and the verb "love", to take a concrete example. This kind of consideration has important consequences theoretical linguistics, because if nouns and verbs do not differ in their meaning, then it is theories of syntax and not of semantics that need to account for the fact that they do differ in other ways (for example in the way they are used in sentences).

In formal semantics, each word is represented as a mathematical expression. In the case of the noun "love", the expression on which there is consensus among semanticists represents the set of all states of love – where a state of love is what we talk of when we say, for example, "My love is like a red red rose". In the case of the verb "love", semanticists have proposed various expressions, standing for the set of all states of love; or the set of all pairs of a lover and a loved one; or the set of all triples of a state of love, a lover, and a loved one; and so on. Each of these possibilities is, in effect, a semantic theory, and while there are good linguistic reasons supporting each of these theories, they are mutually incompatible, and linguists have looked for arguments to choose among them. Since the meaning of the noun and the verb are clearly related, a semanticist's first impulse will be to represent them with just one expression, i.e. the meaning of "love" should be represented as the set of all states of love, no matter if the word is used as a noun or as a verb. However, linguists have looked for other arguments that would not just allow but in fact force a choice between the different possibilities. These other arguments have turned out to be quite difficult to find. The dissertation proposal presents one such argument based on the mathematical properties of the different expressions that semanticists have proposed for the meaning of the verb "love". In a nutshell, only the expression standing for the set of all states of love is compatible with the distributivity operator, and this expression happens to be the same as the one that represents the noun "love". The argument does not depend on the choice of the verb, that is, a similar argument can be made for any other verbs. This means that as far as linguistic theory is concerned, a word has the same meaning whether its function in the sentence is that of a noun or that of a verb. To put it differently, any differences between the way nouns and verbs are used in language cannot be due to their meaning.

3 Conclusion

In sum, this dissertation proposal aims to contribute to the field of theoretical linguistics, using the methods of formal semantics. It sets out to test the hypothesis that there is only one kind of distributivity in human language, to integrate existing formal semantic theories for cumulative and distributive readings and for phenomena similar to *for*-adverbials, and to argue that any differences between the way nouns and verbs are used in language are not due to their meaning.