

Meaning at a Crossroads
Louise McNally
Universitat Pompeu Fabra
ESSLLI 2016 Evening Lecture
Draft: Please do not cite text without permission

ESSLLI evening lectures are a tough gig. Basically, they compete with dinner. In addition, according to the Program Committee guidelines, the evening lecturer is expected to (and I quote) “discuss a topic which is central to their work, putting it in an interdisciplinary perspective, and presenting it as much as possible in an entertaining way.”

I chose the title of this lecture in that spirit - I’m at a point in my career where I’m feeling a bit like the guy in the picture there.

When I chose the title, I did what any responsible academic does and checked on its originality through a Google search. And, as might be expected, I did not coin the phrase “meaning at a crossroads.” Julia Kristeva beat me to it by 43 years in her 1973 book *The system and the speaking subject*. She wrote:

The theory of meaning now stands at a crossroad:

either it will remain an attempt at formalizing meaning-systems by increasing sophistication of the logico-mathematical tools which enable it to formulate models on the basis of a conception (already rather dated) of meaning [...]; or else [...]

For those of you for whom Julia Kristeva does not ring a bell, she’s an extremely influential literary theorist, heavily influenced by the psychoanalytic tradition. So, I’m probably the first and last ESSLLI evening speaker who will ever quote her, and that explains why I’ve deleted much of the quote, and indeed, taken it out of context. I’m not here to endorse anything she’s said - in fact, somewhat ironically, I got into formal linguistics after a frustrating early exposure to literary theory.

But I’ve chosen to begin with this anecdote because it underscores one of the more general take-home points of this talk. When one specializes and opts for a particular methodology – as is inevitable if you want to be on the cutting edge in research – practical limitations, and perhaps also both implicit and explicit biases, lead you to set aside data that you consider not to be central, and methods that you consider too distant from your own. Julia Kristeva is clearly concerned about meaning and language, but those of us who work on natural language meaning using tools logico-mathematical tools set her aside. We tell ourselves that she is concerned with other issues, that she is part of another academic discourse. This strategy works – it has been extremely helpful in the study of natural language meaning.

But it has its limits, and the rest of this talk is about how differently I began to see natural language meaning once I hit those limits and was forced to look at a problem

using a different set of tools.

So what was that problem? I've worked a lot on how lexical meaning interacts with morphosyntax. Most people in formal linguistics don't want to bother with the lexicon. It's messy. It's arbitrary. You get problems like this one: How do we characterize the denotation of even a simple color term in an interesting way? We say it's the set of red things, or a function that takes us from hair to red hair and wine to red wine, but obviously that leaves a lot of work for someone else to do. [red slide]

Like a good, well-trained formal semanticist, this problem only slightly nagged at me, but it did not keep me up at night. I started to worry a bit more when I began to notice things like the following. One day my daughter asked me if she could borrow my jacket, and I said "which one?", and she said, "the blue one", and she meant this one. [blue jacket] You might ask why she chose that description. I think this is why: [green zipper]. Now, we all know that modifiers serve to distinguish objects that fit a particular description, but if you're a nuts and bolts linguist and ask yourself: what is the semantic rule that combines a color adjective with a noun, you see that "the set of blue things" or even "a function from jackets to blue jackets" is not looking obvious. Now, you might say, "aw this problem isn't so hard - you could do something like Kennedy & McNally did for the red traffic light in a 2010 paper." The color term can be interpreted as an indexical property correlated with the color blue. But what is that constant "blue" picking out? We said it denoted an abstract entity - a kind - following Greg Carlson's analysis of generics. But what IS a kind? That question also slightly nagged at me, but I pressed on with my research.

Meanwhile, another problem had been bothering me. It had to do with the syntax/semantics interface of modification. I'll give you just two examples, though there are many, in many languages, involving not just noun phrases but also verb phrases.

The first has to do with the ordering of modifiers. Various factors play a role in the ordering of adjectives and other modifiers, including especially discourse pragmatic ones. I will not get into them all here. But one generalization emerges over and over. I've put Denis Bouchard's formulation here. [slide] This begs the question: what is a concept and what does it mean to "form a concept" - a complex concept? But that there is something to it can be seen in the second example, which comes from a very interesting phenomenon in Catalan that has been called "pseudo-incorporation". [slide] In Catalan, singular count noun phrases normally take an article. However, with verbs related to having, you can drop the article. This is productive, as long as the expression describes a contextually relevant characterizing property. These bare nouns are not often modified, but they can be, and when they are, they have to be those kinds of modifiers that we could describe as "concept forming". [slide] Otherwise, an article is needed. Why should that be?

I want to emphasize two things here, based on these last two examples. First, syntax has come to care about "concept formation" enough to signal it structurally. Second, the concepts are expressed by phrases that need some kind of interpretation rule, just

like any other phrases. You can't really say that the phrases that appear in such constructions are fixed phrases - the phenomenon is too productive, and even if you did, how would a meaning have been associated with the fixed phrase in the first place?

Again, I had ideas about this, but they were not particularly satisfying. They looked like this: [slide]. Following work I did some years ago with Gemma Boleda, M. Teresa Espinal and I proposed that the noun describes a set of kinds (a basic one and all its subkinds) - it doesn't **denote** one - so, different from what I said a moment ago about "blue". Modifiers can be distinguished according to whether they apply to kinds or not. Those that do can restrict the interpretation of the noun, those that don't, cannot. What does the article do? Here we follow basically the same insight that Roberto Zamparelli developed in his thesis. Very simply, the article indicates that the noun is being interpreted differently, as picking out a set of token objects, rather than kinds. The Realize relation here, due to Greg Carlson, introduces those tokens. And the token descriptive adjective can modify the token noun description. The details are not important because I am going to recast this analysis.

I'm not going to get into the pros and cons of this proposal. Let me just say that I didn't feel like I was getting any closer to understanding this notion of complex concept formation or how it might differ from forming complex descriptions of token objects. My toolkit needed some different tools.

At the time, I thought it would be enough to have better tools for analyzing lexical meaning and combining lexical contents. I didn't really worry that much about what a kind was or whether it was the same as or different from a concept. I tried various things, mainly Generative Lexicon theory and Asher's Type Construction Logic. I didn't try any formalizable cognitive linguistic approach such as Frame Semantics because I thought these were too unwieldy to use, not because I had any problem with thinking about meaning in cognitive terms. In fact, I was quite frustrated with the referentialist view of semantics.

And then, around 2009, I learned about Distributional Semantics.

There have been courses and workshops on distributional semantics for awhile now at ESSLLI, so many of you may be at least somewhat familiar with it. You might know it under the terms vector space semantics, or latent semantic analysis, or word embedding. The idea is simple and goes back at least to Zelig Harris and JR Firth: It involves representing the lexical content of a word in terms of other words that it co-occurs with in some contextual window, which could be a sentence, or something larger or smaller. With this [slide] sort of example you can see why this might be appealing: Looking at these words you get a sense of what things are red, and that red can but need not be a gradable property, and so on.

Let's say you want to make a representation for "red". The simplest approach is to build a vector by counting co-occurrences of different words with "red" within some window. You can then compare the vectors that resulted for different words and show,

for example, that “red” is more like “green” than like “hair”. [slide] You can also give distributional representations to phrases. The simplest way to do that is just to follow same process as one uses to build a representation for a word [slide].

But very quickly techniques were developed to compose vectors for complex expressions from simpler ones - I won't go into these here, but it's enough to know that they do a fairly good job. You can show this by comparing the output of composing these representations with what one actually finds for the corresponding phrases, for example. And these models keep getting better. This got me really excited because I saw hope for more interesting models of semantic composition for the sorts of expressions I was interested in.

However, the distributional approach and tools in their simplest form differ in some deep ways from the formal approach and tools. I've schematized these here [slide].

One crucial difference is what goes in a lexical representation. In formal semantics, it is at most the lexical entailments of a word. For instance, for *red* this would include being a color, but beyond that it starts to become hard to say exactly what you could put into the lexical entry. The distributional representation includes all kinds of things that are not entailed when the word is used.

Symbolic cognitive approaches like Frame Semantics might be more like the distributional approach on the issue of lexical content, but I think the important thing to point out has to do with the what the tools lead you to do. If you're using discrete representation, as most cognitivists do, you are not conditioned as to what you put into the lexical content. You can make it as rich or lean as you consider appropriate. The distributionalist doesn't (or didn't, when I started getting interested in this approach) have much of a choice.

A second and related difference has to do with how indeterminacy in word meaning is handled. We've seen that *red* has many uses, corresponding to different shades and distributions of color in combination with different words. Let's call this indeterminacy. In formal semantics, there are two ways to handle indeterminacy: Either you enumerate all the senses - an approach that is computationally problematic and also not at all explanatory - or you put as little as possible into the lexical representation, so that it is compatible with all of the possible uses of the word, and have general inferential processes fill in the rest - these processes not being, so the story goes, the linguist's concern.

The distributional strategy is to put everything in the lexical representation, and have irrelevant information filtered or suppressed in composition. Since everything is in the representation at the outset, it's not surprising that the distributional strategy will do better with composition of content words - the formalist's strategy is to remain silent about precisely that information that will help.

Now I want to turn from content words to two other differences between formal and

distributional approaches. These are things that can make distributional semantics a non-starter for most formal semanticists. One is that it is far from obvious how to treat function words using distributional methods. Think about *and* and *or*. The articles are not much easier. The problem is that these words are sensitive to syntactic category, but arguably not to specific lexical items. Relatedly, if you ask yourself what concept does *and* or *the* correspond to, it's not that easy to answer. These are the kinds of words that formal semantic methods offer sophisticated analyses for.

Finally, there is the problem of token reference. You can see the problem very easily by considering a short text like the Three Little Pigs. [slide] Let's focus on not even on a word, but a phrase, *the house*. The phrase appears 4 times in this version of the story, but only twice does it refer to the same house. How could a distributional vector distinguish these different uses of *the house*? Again, though resolving anaphora in discourse is far from an easy problem, formal methods can easily represent the differences between these different uses of *the house*.

So, distributional semantics does well with the lexicon and “complex concept” formation; it doesn't (or didn't) do well with function words and reference. People are working on this. I spent some time thinking about how to deal with especially function words myself. But gradually I opted to take another route. What moved me to do this was in part the same language data that I mentioned earlier, and that got me interested in distributional semantics in the first place:

1. The restrictions on adjective order, and the fact that they arguably involve a “concept-forming layer” combined with a “descriptive layer”. Grammar distinguishes these two.
2. Data like the Catalan pseudo-incorporation I showed earlier, or this example from Hindi, and other so-called incorporation constructions. It is often noted that incorporation constructions carry some kind of condition that they describe a “prototypical” situation or a “well-established” type of event. It's striking that precisely the construction in which the nominal is non-referential carries conditions related to general knowledge that the referential condition does not. This suggests that composition that is NOT mediated by reference is different from composition that IS mediated by reference.
3. Work in what Hagit Borer calls “exoskeletal” approaches to morphosyntax. Here's a description of such approaches (which include, broadly speaking) so-called Distributed Morphology as well as arguably at least some versions of Construction Grammar. [slide] It is not surprising that Borer would defend something like this as a speaker of a Semitic language, where the morphology consists of roots and templates, even for the most basic words. It turns out that there is even more syntactic work that goes strikingly in the direction of composing content words independently of function words, despite traditional assumptions about constituent structure.

Here is how I'm proceeding. [slides]