NATURAL RELATIONSHIPS BETWEEN NOUNS, VERBS, AND ADJECTIVES

JOHN THICKSTUN

I mentioned in slack that I thought I saw a natural equivalence between nouns and adjectives. I'm going to attempt to clarify what I meant by that and to highlight the close connection between nouns, verbs, and adjectives. Along the way, I'm also going to try to hint at some deeper categorical structure. I'm not a very proficient category theorist, so I'd welcome some feedback. Also, I haven't done much of a literature survey; if I continue to be interested and this generates interest, the next step might be to find out what's already been said.

We're going to need to start by developing some of the machinery of linguistics. Let's start with sentences. Declarative sentences can usually be thought of as propositions. By declarative sentence I mean stuff like "the dog howls at the moon" or "Barack Obama is the president," but I am excluding sentences like "kill the boy" (imperative) or "when does the tide come in?" (question). I'll stick to declarative sentences throughout; they are rich enough to discuss about nouns, verbs, and adjectives but avoid the need to develop a broader linguistics theory.

Because we will be using the English language to study itself, we will need to draw a careful distinction between our use of English as a meta-language of discussion, and as the object-language of study. To do so, I will enclose any of of English object-language in quotations (as I did in the preceeding paragraph). I'm also going to introduce two operators, a syntax operator $|\cdot|$ and a semantics operator $[[\cdot]]$. The latter is a standard is standard in the literature; I made up the syntax operator since the literature doesn't seem to write it out explicitly. I wanted to highlight it because I think these operators may actually be functors[1] (which we would obviously be interested in studying). The English language then consists of triples $\langle``\alpha", |\alpha|, [[\alpha]]\rangle$ where α denotes a word or fragment or sentence in the object language.

Let's take a closer look at syntax. A sentence like "the dog howls at the moon" has a rich syntactic structure. I can, for example, replace "dog" with "cat" or "ambulance" and the sentence will still make sense (although it may no longer be true). But I cannot replace "dog" with "red"; "the red howled at the moon" is neither false nor true; it is not a well formed sentence. This leads us to a theory of syntactic types, also called parts of speech (or sometimes syntactic categories, a term that I will avoid because I do not believe it agrees with mathematical use of the word).

[1] There are a few hints on the internet that Lawvere has a theory of syntax and semantics as adjoint functors, which would be pretty cool. Regardless, if $|\cdot|$ and $[[\cdot]]$ are in fact functors, they may provide the proper context in which to formulate the naturality claims I've been hand-waving about.

The codomain of $|\cdot|$ consists of parts of speech, for example sentences. We will denote sentences by S, so $|\text{the dog howls}| = S$. We will denote nouns by N and write $|\text{dog}| = N$. We can also write $|\text{cat}| = N$ and $|\text{ambulance}| = N$, but not $|\text{red}| = N$ (we will later identify "red" as an adjective). We begin to develop a superficial understanding of why we could substitute "cat" or "ambulance" for "dog", but not "red": the former terms have a common part of speech, not shared by the latter. But this is somewhat circular: the whole point of lumping nouns together in N is that they are syntactically substitutable; "noun" is just a name for a particular set of words with a common syntactic role.

Before developing our theory of syntax further, let's take a look at some basic semantics. A declarative sentence like "Obama is the president" reads like a proposition: it is true or false. We might be inclined to say that $[[\text{Obama is the president}]] = \textbf{True}$ but this is problematic. For one, every true sentence in the English language would have the same meaning! This is clearly not right. More subtly, someday this sentence will no longer be true. This gives us a clue for modifying our semantics: $[[\text{Obama is the president}]]$ is actually a function from states of the world to truth values. It is true of the present world, but was false in the world of 2004 and will be false again in the world of 2020. If we denote worlds by W and $T = \{\textbf{True}, \textbf{False}\}$ then $[[\text{Obama is the president}]] : W \to T$. While this distinction is important, we will often find it convenient to omit the worlds parameter when discussing semantics.

A noun like "dog" can be modeled as a set of objects or alternatively, as a function from objects to truth values. If we denote objects by E, then $[[\text{dog}]] : E \to T$. It is an empirical fact that all nouns can be modeled this way. This should surprise you! Recall that "noun" is a term we created to denote a syntactic type. But now we see that nouns also share a semantic type signature. This is our first encounter with Frege's principle: there is a homomorphism between the syntactic and semantic structure of language. This principle is defensible from both normative and positive perspectives. We have seen it here in its positive form; in what follows we will exploit it as a normative tool for modeling.

In addition to nouns, we have a related syntactic type of noun phrases, denoted NP. This type includes proper nouns like "Obama" or "New York." These fragments do not have the same syntactic properties as N (try writing down a sentence and substituting an NP for an N). They also have different semantics. A proper noun isn't a $\text{Hom}(E,T)$; it's just an E. Obama isn't a set, and neither is New York (at best they are single-element sets). This syntactic type also includes phrases like "the dog" (again, exercise, try substituting "the dog" for a proper noun in a sentence). We won't attempt to give a syntax and semantics to "the" yet, but feel free to think about it.

What are the syntax and semantics of a verb? The syntax is easy. We would like to call words like "howls" verbs. Clearly "howls" is not an N or an NP (try substituting it for a noun in a sentence) nor is it an S, so it merits its own syntactic type V. Furthermore, we can characterize how it behaves syntactically: a V combines with an NP to form an S. For example, "Lassie howls." We can express this graphically as a syntax diagram:

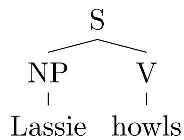

Semantically, we know that [[Lassie howls]] is a truth value (well, really a $\mathrm{Hom}(W, T)$ but remember we drop the worlds when it's convenient) and [[Lassie]] is an E. By Frege's principle (now the normative version!) [[Lassie]] and [[howls]] should compose to form [[Lassie howls]]. With this motivation, it should be obvious that we need [[howls]] : $E \to T$. In particular, [[howls]] is the function that maps objects that howl to **True** and everything else to **False**.

If you've been paying attention, you might have just noticed that nouns and verbs have the same semantic type. They are both $\mathrm{Hom}(E, T)$! The distinction between nouns and verbs is (mostly) syntactic. It is an artifact of the structure of the English language. From a mathematical perspective, we make the observation that Frege's principle is no more than a homomorphism; it cannot be an isomorphism because multiple syntactic types map to the same semantic type. From a practical perspective, semantic equivalence between nouns and verbs enables "verbing," the impromptu conversion between nouns into verbs. For example, "talk" and "dream" can be used as both nouns and a verbs: "the talker talked at the talk," or "I dreamed a dream in my dreams." I can't resist linking my favorite comic strip.

Before moving on to adjectives, let's revisit the word "the." Clearly "the" doesn't belong to any of the syntactic types introduced so far. We will introduce a new type *Det* (short for "determiner") for this purpose. This type includes words like "this" and "that." A *Det* combines with an N to form an NP. So for example

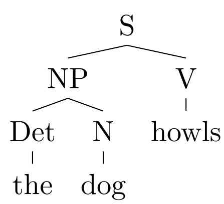

We know that [[the dog]] $\in E$ so and [[dog]] $\in \mathrm{Hom}(E, T)$ so we must have [[the]] : $\mathrm{Hom}(E, T) \to E$. This function doesn't have a neat mathematical formulation (that's English for you) but if you think of $\mathrm{Hom}(E, T)$ as a set of E then [[the]] picks out the contextually salient E in that set. If there's no contextually salient dog, or many, then talking about "the dog" doesn't make much sense.

Let's look at a (prenominal) adjective now. Repeating the exercise we did with verbs, clearly a word like "red" isn't any part of speech that we've already described (try experimenting by substituting it for words with parts of speech we've already identified). So we'll give it a new syntactic type *Adj*. Its syntactic role is to combine with an N to form an N. For example, "the red balloon flies" has the following syntax diagram:

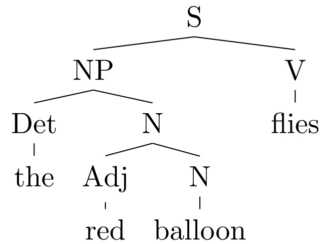

By Frege's principle, [[red]] and [[ballon]] must compose to form [[red balloon]]. It follows that [[red]] : $\mathrm{Hom}(E,T) \to \mathrm{Hom}(E,T)$. Given an input $\alpha : E \to T$, [[red]](α) is a function $\beta : E \to T$ that satisfies $\beta(e) = \textbf{True}$ iff $\alpha(e)$ and e is red.

We are now in a position to state the claim I made in slack about a natural equivalence between adjectives and nouns (if that is indeed what it is–obviously I haven't clearly specified the categories or functors involved). Notice that [[red]] is a conjuction of two criteria. The first criterion is determined by the input to [[red]]. The second is the criterion "is red." This is sometimes called a predicate. Notice that it must be a $\mathrm{Hom}(E,T)$ (by Frege's principle). The transformation I have in mind maps a $\mathrm{Hom}(E,T) \to \mathrm{Hom}(E,T)$ of this form to a $\mathrm{Hom}(E,T)$ by extracting this predicate from the conjunction. In mathematical detail, if γ is an adjective then $\gamma : \mathrm{Hom}(E,T) \to \mathrm{Hom}(E,T)$ maps $\alpha : E \to T$ to a function $\beta : E \to T$ by the rule $\beta(e) = \textbf{True}$ iff $\alpha(e) \wedge \Gamma(e)$ for some predicate $\Gamma : E \to T$. The transformation I'm talking about is $\gamma \mapsto \Gamma$.

I should disclose that the above discussion of prenominal adjectives doesn't appear to be standard. Standard practice seems to model adjectives as $\mathrm{Hom}(E,T)$ and put the conjunctive logic into the combinatory rule for composing the semantics of adjectives and nouns. I suspect that the math works out better if we push this work into the semantics of the word (as I did above) and uniformly apply a single combinatory rule: function composition. Of course, in the standard model, the connection between nouns, verbs, and adjectives is even more apparent.

I'll leave this for now with a comment I discovered in my notes from a linguistics class with Professor Pauline Jacobson at Brown. Polly made the empirical observation that $\mathrm{Hom}(E,T)$ in English tend to sort themselves into the types N, V, and *Adj* based on how stable they are over time. Verbs are very unstable. I am typing right now, but I won't be typing later today. Verbs denote activities and activities are transient. Adjectives, on the other hand, are pretty stable. If I have long hair today, then I will probably have long hair tomorrow. If a car is red today, it will probably still be red tomorrow. Nouns, finally, are the most stable class. If I am a dog today, I would be astonished if I woke up as a cat tomorrow.