

Natural Language and NLP

If having formal axioms for the everyday world is too difficult, then how about having artificial agents use a natural language like English? After all, that language already exists and works great for humans. So we'd avoid the need to do all that hard work of picking an adequate ontology. Not only that, humans could then talk to the agents without special training in an artificial language.

This of course is an attractive idea and one that has been getting huge attention for a long time. But it has plenty of hard aspects itself. Most glaringly, we do not have precise rules for how English (or any natural language) is used by humans. This turns out to be enormously complicated. And since we do not understand this very well, then we cannot (at present) write it all down to go into an agent KRR system.

Linguistics. Linguists have been hard at work for a long time, trying to come up with better and better rules about how humans actually use language. Traditional subareas include these:

Phonology – the study of the sound-pieces (phonemes) that make up speech

Morphology – the study of words and the word-pieces they are composed of

Syntax – the study of how words fit together into sentences

Semantics – the study of how word-meaning contributes to sentence-meaning

Pragmatics – the study of everything else that contributes to meaning

Syntax is perhaps the best-known part of linguistics to most of us. It includes the key notion of *grammar*. However, this is not the same as the grammar that we were taught in middle-school. The grammar that linguists study is the actual patterns of words that people in fact use, not what someone says they *should* use. And a big claim (by Noam Chomsky and many others) is that how people actually form and understand sentences in natural language (i) actually does follow definite patterns that can be discovered, and (ii) those patterns lend insights into how the mind works at an innate level – that sentence-structure is not simply a matter of imitating whatever we hear as children.

We shall not enter deeply into this topic here. But a few more comments will be helpful. Here is a very famous sentence, due to Chomsky:

(a) *Colorless green ideas sleep furiously.*

The point he makes with this sentence is that English speakers immediately recognize that this is in fact a sentence. Compare it to, say,

(b) *Green sleep colorless furiously ideas.**

which no takes to be a sentence at all. The * at the end of (b) is a common way to mark an *infelicitous* combination of words, a sentence wannabe that does not make it.

Yet (a) is utter nonsense! I don't mean nonsense as in *The Moon is made of green cheese*. This latter sentence is false, but entirely intelligible. We know what it means, what specifics about the world it conveys, what the Moon would be like if it were true. But (a) above is not like that; it is total gibberish, as also is (b); and yet—quite unlike (b)—(a) stands out as a full-fledged sentence of English. Indeed, (a) is almost more of a template—*subject, verb, object*—for what a meaningful sentence ought to be, a meaning-conveying *format*. And so the Chomskyian innateness hypothesis in effect says that to convey meaning, a natural language has certain special formats that the brain can recognize (when we are spoken to) and produce (when we speak); but those formats themselves are not meaning.

The above discussion suggests we try to form a hierarchy of sorts:

Gibberish: Green sleep colorless furiously ideas.

Grammatical: Colorless green ideas sleep furiously.

Meaningful: The Moon is made of green cheese.

True: The Moon is a satellite of the Earth.

Syntax, then, is the study of (or search for) rules that distinguish Gibberish word-sequences from Grammatical word-sequences. And semantics is the study of (or search for) rules that distinguish merely grammatical from meaningful word-sequences, and (perhaps more importantly) what the meaning of the latter sort of sentence actually is. However, the above hierarchy is misleading in some ways.

More examples will help. First consider the garden-path sentence

(c) *The cat curled up in the basket was friendly.*

This is both grammatical and meaningful; but it takes a little work to determine that it is grammatical – there is a temporary puzzle as to whether the words fit together into a proper sentence, especially concerning the role played by the word “curled”. It is tempting to think this is simply a matter of whether the word sequence makes sense, has an interpretation that we can picture (like a green-cheese Moon). But Chomsky's sentence (a) above makes clear that being meaningful is *not* the criterion for sentence-hood.

Here is a sentence wannabe – but less severe than (b) – that may clarify this more:

(d) *The house beside the lake with the dragons eat people**

(d) almost makes sense. We can picture something: there is a house and a lake and people-eating dragons. But (d) does not make it as a sentence, despite a more-or-

less unambiguous meaning that can be extracted from it, with some effort. So (d) is not a sentence, and yet (largely) meaningful.

Now consider:

(e) *The double-agent squashed the bug.*

Here there is not much worry about its being grammatical; but whether the bug is an insect or a miniature microphone is uncertain (although the context of double-agent perhaps tilts us toward the microphone interpretation. This illustrates how meaning is not just a matter of individual words that contribute their individual meanings to an overall sentence meaning. The word meanings interact contextually, and can even depend on world-knowledge (double-agents may not want to be overheard). How world-knowledge (beyond that of what the individual words can mean) contributes to the meaning of the sentence, is the area of pragmatics.

Here are two more examples:

(f) *The iron ball crashed through the table because it was thin plywood.*

What does “it” mean in (f)? The table, for more than one reason: (i) iron is not wood, so the iron ball cannot be thin plywood; (ii) crashing suggests violent motion and iron is hard and dense and thin plywood is not and so might not be able to support the moving iron ball; and (iii) the word “because” offers thin plywood as an explanation of *why* the table did not stop the ball – and so from (ii) this again points to the table – not the ball – as being plywood. Quite a bit of knowledge about iron, plywood, tables as supports, falling motion, and so on, comes into play in the determination of what *it* refers to in (f).

(g) *Mary asked Ann if she could help with her homework.*

Here we cannot tell whom *she* and *her* refer to; but we can be fairly sure that they do not refer to the *same* person; people do not *help* with their *own* homework. Again this is world knowledge about people and homework.

Examples (e), (f), and (g) illustrate the issue of *reference*: what a given word (usually a noun or pronoun) refers to. In some cases this is settled simply by unambiguous shared understanding of how a particular word is used in a given language (information that could be looked up in the agent’s lexicon); in some cases it requires substantial general knowledge as in the three cases just seen; and in some cases it seems to be built into the syntactic rules that we all share. Here is an example of the latter:

(h) John told him never to do that again.

In (h) the pronoun *him* refers to someone *other* than the subject *John*. This is not something we can figure out by world knowledge. After all, John certainly could very reasonably tell *himself* not to do something again. It is apparently a *rule* about how English “works” (or how minds process English) that does not allow a pronoun in direct object role to refer to the subject unless it is a *reflexive* pronoun. So this is not a matter of pragmatics, nor even very much of semantics *per se*, but more of syntactic rules dictating semantic roles.

And one final example:

(i) *It’s awfully windy in here, with that open window behind you, and papers are blowing all over the room.*

Here we see a case of pragmatics at work *par excellence*. When people make utterances to each other, they have reasons. Mary utters “Call Sue” to John because she wants him to do something – and in fact her utterance says just that: it is a command (imperative). Sometimes the reason is simply to usefully inform, such as when one says “Clocks were set back one hour last night.” In these two cases, the meaning is pretty much right there in the sentence. But in (i), while on the one hand there is a “conventional” meaning in the words themselves (temperature related to location of open window and to blowing papers), there is an implicit meaning: please shut that window. And that (the implicit request to do something) is nowhere in the sentence at all; nor is it part of knowledge about the world in general, not windows and wind and papers in general. It is quite specifically about people and what they do and don’t tend to like, and why they say things to one another, and what we are expected to do in certain circumstances.

We have barely scratched the surface of linguistics. Among the many topics not even mentioned above is that of *developmental linguistics*—how people come to acquire the *competence* to use a language as they grow up (or as they learn a new language as adults). This in particular has strong ties to neuroscience and psychology.

Philosophy of language. This is a huge subject as well, and again one we will barely touch on here, just enough to mention two aspects: reference and pragmatics.

Reference has to do with the relation between an expression and what it stands for or refers to, as in

(j) “*The morning star*” refers to the planet *Venus*.

We saw earlier examples, for instance concerning pronominal reference—concerning what *it*, *she*, and *him* pick out in (f), (g), (h) above. But in (j) the question is not about how one word relates to another item in a sentence, but to an item in the world; this is so-called *problem of external reference*, and has been the subject of

a vast amount of study, at least from John Stuart Mill on (including philosophers such as Brentano, Meinong, Frege, Russell, Kripke, Putnam and many others).

Among the complications are these: two words can co-refer (have the same referent) and yet not be totally interchangeable; a word can refer to two or more different things; reference appears somehow to be a kind of social agreement within a community sharing a common language; and it often is the case that *no one* in such a community knows what a given expression in fact picks out even though they all may use that expression effectively. In addition, external reference has been linked to the philosophy of mind (something we will take a quick look at near the end of the semester).

Paul Grice (1913—1988) wrote extensively about pragmatics and meaning, including what he called “speaker meaning” – what the speaker is trying to *convey* in making an utterance. One of his famous conclusions is that “A meant M by X” is roughly equivalent to “A uttered X with the intention of inducing a belief that M by means of the recognition of this intention”. The idea in part is that the listener gains fuller understanding of the speaker’s intent by considering the speaker’s purpose in making the utterance. If I tell you that $2+2=4$, you may well wonder why I bothered: perhaps I noticed an arithmetic error in something you wrote, or perhaps I am saying something obvious (and not in need of saying) on purpose to suggest that something said a moment ago was (also) obvious and not in need of being said.

A related phenomenon Grice studied was that of *implicature*. Consider this dialog fragment:

A: Are the roses fresh?

B: I put them in the fridge, but they are not fresh.

B’s initial response (I put them in the fridge) is then clarified (but they are not fresh). That is a strange clarification, on the face of it, to an equally strange initial response (what has the fridge got to do with the question?). But our world knowledge allows us to infer that the clarification is aimed at *preventing a conclusion by A* that the roses are fresh (since they are in the fridge, which would help them stay fresh longer). That is, B realizes A might quite reasonably infer something (that is the so-called implicature) from the initial response, and so B is being helpful in trying to prevent A from that mistake.

Grice formulated several *maxims* of conversation, that are intended to capture aspects not how people *should* converse but how we (by and large) *do* converse. Thus B above sets the record straight (“but they are not fresh”) because B knows A is likely to think otherwise based on the earlier portion of the response. And B knows that because people (A in this case) assume their questions are being answered in ways intended to be relevantly informative. So B’s clarifying behavior in part depends on B’s beliefs about A’s beliefs about B’s intentions regarding A’s question and B’s initial response. Hard stuff!

Computational aspects. OK, the above was a whirlwind tour of issues in linguistics and philosophy of language. But what about NLP—the *processing* of language by computers? Here I shall be even briefer, just saying two things:

1. Almost all the above topics (and more, such as machine translation between languages) have been and are being studied intensively in terms of how the behaviors in question might be carried out computationally—and this is being done not only by computer scientists (AIers) but also by psychologists and linguists, since the brain apparently must somehow manage to “compute” whatever behaviors it carries out (in a sense we will consider a little later).
2. One very hot topic is that of *statistical methods* in which reference and so on are computed in terms of what is the most likely one in a given context. And as it turns out, statistical methods in AI (probability and decision-making) are the topic for next week!