Ling 596D Compositionality Jan. 14, 2004
Heidi Harley

1.0 Form and meaning

Almost any linguist can agree that 'language' consists of a mapping of some kind from form to meaning.

In the unmarked declarative situation:

Proposition A represented in Speaker A's head is encoded into natural language utterance *a* and transmitted.

Hearer B, a speaker of the same natural language[1], upon decoding utterance *a*, comes to have a representation of Proposition A' in his head. If all goes well, A=A'.

1.1 Underspecification

Now, it seems obvious that certain parts of the meaning of Prop A may not be encoded, or only partially encoded, in the natural language representation.

E.g. if Speaker A estimates that Hearer B is equipped to evaluate the reference of a pronoun correctly, and assuming that the syntactic conditions on pronoun use are satisfied, Speaker A may choose to use a pronoun, not specifying exactly the referent s/he intends.

(1) She left.

E.g. Cases of deixis: the tense system (present, past, future are evaluated with respect to speech time, whenever that is); deictic locations and times (today, here, yesterday, then, there); possibly even the cases like the interpretation of 'killer' vs. 'singer' that Massimo mentioned.

(2) She will leave tomorrow.

So the form may not be a full representation of the meaning. However usually, there are very rigid rules about how to use context to fill in the underspecified parts. They in general do not allow for ambiguity. (These rules can be and are studied in their own right, e.g. Discourse Representation Theory)

Confusion may arise if Hearer B is assuming a different context for the utterance than Speaker A, but confusion never arises because Hearer B and Speaker A disagree about the procedure for filling in the underspecified information.

---

[1] Speaker A and Hearer B, as speakers of the 'same' natural language, are idealized as having identical grammars and lexicons.

E.g. although to fix an absolute reference for any given use of the word 'tomorrow' requires additional information, namely the time the word was uttered (in the simplest case), once that information is known, there is only one possible absolute referent for 'tomorrow'; e.g. "January 15, 2004". Hearer B could be confused about what day it is today, and thereby be confused about what absolute day Speaker A is referring to by 'tomorrow', but Hearer B is NOT confused about what he has to do to interpret 'tomorrow'.

Hypothesis A: The decoding process is *always* determinate in this way. If Hearer B's and Speaker A's understanding of the context is identical, utterance *a* can only be interpreted in one way, as Prop A.

Misunderstanding arises when A and B's representations of the (grammar-independent) context are different.

1.2 Structural Ambiguity

Maybe hypothesis A is too strong. What about:

(3) Mary saw the boy with the telescope.
LOT A: MARY SAW THE BOY WHO HAD THE TELESCOPE
LOT B: MARY SAW THE BOY BY LOOKING THROUGH THE TELESCOPE

Linguists' answer: problem here is not with Hypothesis A, but rather with the idea that (3) are adequate representations of utterance *a*. The encoding of Proposition A into utterance *a* involves more than a simple string of phonological words.

(4) a. Mary saw John.
b. Mary saw John with the telescope.

Because (4a) could be exactly equivalent to LOT A but not LOT B , and (4b) could be equivalent to LOT B but not LOT A, utterance *a*, when equivalent to LOT A, must involve additional distinguishing information than when equivalent to LOT B. The additional distinguishing information is *constituent structure*, not usually represented in the written string, and often only incompletely represented in the spoken string (via, e.g. intonation patterns).

(5) a. Mary saw [the boy [with the telescope]]
b. Mary saw [the boy] [with the telescope]

Utterance *a*, then, involves a hierarchical structure; part of Hearer B's job is to make an accurate representation of the structure of utterance *a*, guided by the grammar of the natural language. In many cases, the string will only allow for a single structure, but sometimes there will be ambiguity. Again, however, there are grammar-independent rules which govern the hearer's construction of structural representations when there is

potential ambiguity; these can be and are studied in their own right. (E.g. Janet Fodor's 'Attach Right' rule.)

So we can write a description of what's involved in utterance *a* into Hypothesis A:

Hypothesis A: The decoding process is *always* determinate in this way. If Hearer B's and Speaker A's understanding of the context is identical, an utterance *a* of a hierarchically structured string of words can only be interpreted in one way, as Prop A.

Misunderstanding arises when A and B's representations of the context are different. Misunderstanding arises if B has mis-analyzed the structure of utterance *a*.

1.3 Lexical ambiguity:

One last obvious case where Hypothesis A seems too strong:

(6) [Mary [saw [the bank]]].
LOT A: MARY SAW THE [financial institution]
LOT B: MARY SAW THE [side of the river]

Linguists' answer: Again, the problem here is not with Hypothesis A, but with the fact that (6) is an inadequate representation of utterance *a*, even though (6) has constituent structure.

What (6) fails to represent is the homophony of the word *bank*: there are (at least) two distinct lexical items represented by the phonetic string /bæŋk/.

Part of Hearer B's job, again, is to construct a representation of utterance *a* that does not just involve a structured string of phonological words, but a structured string of *semantic* words — words that themselves have meanings independent of the utterance they occur in.

If B has accidentally represented utterance *a* as containing a different semantic word than A put in, misunderstandings will occur.

In the ideal case, the phonological word will uniquely identify one or more semantic words, but when there is homophony, as is (6), there will be ambiguity.

Again, however, there are grammar-independent rules that govern how Hearer B will decide which semantic word to represent in cases of lexical ambiguity, having to do with context, priming, and so on. These rules can be, and are, studied in their own right. (E.g. semantic priming studies).

So, rewriting Hypothesis A to indicate that utterance *a* must be a representation of semantic words

Hypothesis A: The decoding process is *always* determinate in this way. If Hearer B's and Speaker A's understanding of the context is identical, an utterance *a* of a hierarchically structured string of semantic words can only be interpreted in one way, as Prop A.

Misunderstanding arises when A and B's representations of the context are different.
Misunderstanding arises if B has mistakenly represented the structure of utterance *a* differently than A intended.
Misunderstanding arises if B has mistakenly represented utterance *a* as containing different semantic words than A intended.

2 Compositionality

Let's look at the first possible source of misunderstanding for a minute: different representations of discourse context.

→ By hypothesis, if A&B's representations of the context are identical, rigid rules concerning the specification of deictic material will ensure that B arrives at the same semantics for utterance *a*'s deictic material as A intended.[2]
→ Contextual representations are by hypothesis grammar independent — they're part of one's internal representation of the world.
→ In fact, relative to a given context, hypothesis A states that the interpretation of a structured string of semantic words, utterance *a*, is completely determined.

Let's re-write hypothesis A to eliminate the reference to speakers and hearers:

Hypothesis A: Relative to a particular discourse context, an utterance *a* of a hierarchically structured string of semantic words of a natural language can only be interpreted in one way, as Prop A.

If the last two ingredients above that can cause misunderstanding of the content of a natural language utterance — mistakes in the structure and mistakes in the meaning of the semantic words — are the only sources of misunderstanding that can arise between two speakers with identical contexts, grammars and lexicons, then the following is a corollary of Hypothesis A:

Compositionality Corollary: The context-independent meaning of a hierarchically structured string of semantic words of a natural language is a function of the meanings of the semantic words (lexemes) and the way in which they are structured.

[2] Actually, what's really at issue here is if B's internal representation of the context is the same as what A *thinks* B's internal representation of the context is, not whether A's and B's are the same. A will try to correct for differences between A's representation of the context and what A thinks B's representation of the context is.