

Incremental and predictive interpretation: Experimental evidence and possible accounts

Adrian Brasoveanu & Jakub Dotlačil

1. The main question: The nature of real-time semantic interpretation. The main question we will address in this talk is whether meaning representations of the kind that are pervasive in formal semantics are built up incrementally and predictively when language is used in real time, in much the same way that the real-time construction of syntactic representations has been argued to be (Steedman 2001, Lewis and Vasishth 2005, Lau 2009, Hale 2011 among many others).

There is a significant amount of work in psycholinguistics on incremental interpretation (Haagoort et al. 2004, Pickering et al. 2006 among many others), but this research focuses on the processing of lexical semantic and syntactic representations, as well as the incremental integration of world knowledge into the language interpretation process. The processing of logical representations of the kind formal semanticists are interested in is much less studied. Similarly, the significant amount of work in natural language processing / understanding on incremental interpretation (Poesio 1994, Hough et al. 2015 among many others) usually discusses it from a formal and implementation perspective, and focuses much less on the cognitive aspects of processing semantic representations. (Steedman 2001 and related work is a notable exception.)

2. The phenomena: Cataphoric presupposition resolution in conjunctions vs. conditionals. Investigating the incremental processing of formal semantics representations will first require us to identify phenomena that can tease apart the syntactic and semantic components of the interpretation process. In particular, the pervasive aspects of meaning composition that are syntax based / driven cannot provide an unambiguous window into the nature of semantic representation building: the incremental and predictive nature of real-time compositional interpretation could be primarily or exclusively due to our processing strategies for building syntactic representations.

The interaction of *presupposition resolution* with *conjunctions vs. conditionals* (where conditionals have a sentence-final antecedent) promises to provide us with the right kind of evidence. Consider the contrast between the ‘cataphoric’ examples in (1) and (2) below:

- (1) Tina will have coffee with Alex again AND she had coffee with him at the local café.
- (2) Tina will have coffee with Alex again IF she had coffee with him at the local café.

Now assume that the construction of semantic representations is *incremental*, i.e., the human language interpreter processes IF as soon as it is encountered, and *predictive*, i.e., the interpreter builds a semantic evaluation structure wherein the upcoming *if*-clause provides (some of) the interpretation context for the previously processed matrix clause. We would then expect to see a facilitation / speed-up when the second clause (*she had coffee with him ...*) appears after IF (2) vs. AND (1) – which is what the experimental results actually show. We expect (1) to be less acceptable / more difficult than (2) right after the presupposition trigger again is interpreted. The conjunction AND in (1) signals that an antecedent that could resolve the again-presupposition is unlikely to come after this point, since the second conjunct is interpreted relative to the context provided by the first conjunct. In contrast, the conditional IF in (2) leaves open the possibility that a suitable resolution for the again-presupposition is forthcoming, since the first clause is interpreted relative to the context provided by the second clause.

Crucially, the different expectations triggered by the interaction of the presupposition trigger again and the operators AND vs. IF are semantically driven: there is nothing about the syntax of conjunction vs. *if*-adjunction that could make the possibility of a successful presupposition resolution more or less likely.

The talk will present the full experimental setup for the self-paced reading experiment we ran to investigate these predictions. The experimental setup is more complicated than initially expected; we will discuss why that is and in the process, identify methodological issues specifically related to using self-paced reading to study this kind of semantic issues. We will also outline a second self-paced reading experiment that uses pronominal anaphora/cataphora (rather than cataphoric presupposition resolution) to probe the nature of real-time semantic representation construction.

3. Two accounts. Finally, we discuss two very different strategies that could capture the incremental and predictive nature of real-time construction of meaning representations.

3.1. Enriched semantics: Incremental Dynamic Predicate Logic (IDPL). On one hand, we can enrich semantic evaluation contexts – basically, we replace single variable assignments with trees of variable assignments. This refinement enables us to provide a recursive definition of truth and satisfaction that derives the correct truth conditions for both (1) and (2) above while interpreting both of them in a strictly left-to-right, incremental, and predictive way. We dub the resulting system Incremental Dynamic Predicate Logic (IDPL). IDPL builds extensively on the propositional dynamic logic introduced in Vermeulen (1994).

3.2. Processing enriched logical forms: An ACT-R based left-corner parser for DRs. On the other hand, we show that we can maintain a simpler notion of evaluation contexts and a simpler definition of truth and satisfaction as long as we assume that:

- i.* the human language processor makes crucial use of an intermediate, Discourse Representation Theory (DRT) style level of representation (Kamp 1981, Kamp and Reyle 1993), and
- ii.* it incrementally and predictively builds Discourse Representation Structures (DRs) in much the same way that a left-corner parser incrementally and predictively build syntactic structures (Lewis and Vasishth 2005, Hale 2011 among others).

We briefly outline a computational implementation of such a DRT parser in Python ACT-R (see Anderson and Lebiere 1998 for ACT-R, and Stewart and West 2007 for the Python implementation), and compare the results of our experiment and the quantitative predictions made by this ACT-R based DRT parser.

3.3 Semantics and processing. We might be able to empirically distinguish between these two (types of) accounts, but irrespectively of that, a methodological point can be made. When we analyze phenomena in purely syntactic and / or semantic terms, we implicitly classify them as essentially belonging to the grammar of the language under investigation. But we should probably try to consider alternative, processing based / laced explanations more systematically. It is hard to know *a priori* what the best explanation for a phenomenon is – consider the variety of accounts of negative polarity item (NPI) licensing proposed over the last 50 years or so.¹

Enriching semantic evaluation contexts and providing a finer-grained recursive definition of truth and satisfaction for natural language interpretation (as in IDPL, for example) might be the right thing to do. But an independently needed theory of the real-time processing of semantic (and syntactic) representations might also be able to account for crucial aspects of the phenomena under investigation. Or it might turn out that a ‘hybrid’ semantic and processing approach might be the right way to go. But the only way to begin exploring this space of alternative explanations is to explicitly consider and try to formulate mathematically explicit theories of how formal semantic representations are processed in real time.

¹For example, the mainly syntactic approach in Klima (1964), the semantic approach in Ladusaw (1979), the recent discussion of NPI processing effects in Vasishth et al. (2008) – among many others.