

## Towards a Semantic Theory of Discourse

C. Raymond Perrault

Artificial Intelligence Center and  
Center for the Study of Language and Information  
SRI International

I don't feel comfortable trying to building reliable, well-understood NLP systems without providing a semantics for their "mental state," including the data structures they encode. One step in that direction is a semantics of sentences, for example that of Montague. However, to handle extended discourses, the methods of model-theoretic semantics need to be extended in (at least) three directions. First, anaphora and deixis require that the interpretation (of phrases, sentences, and entire discourses) should be made sensitive to context, both linguistic and physical. Second, non-declarative sentences must be treated uniformly with declarative ones. Finally, it should be possible to make room in the semantics for interpretation constraints based on the subject-matter of the discourse and on communication itself. I'm thinking particularly of the kind of inferences typically captured computationally by the use of scripts, plans, prototypes, etc. Typically these constraints on interpretations have been kept outside the realm of semantics, and might even be taken to be the distinguishing characteristics of a separate pragmatics component. I'd like to suggest that we already have available many of the necessary elements for a context-sensitive theory of discourse, although substantial work still needs to be done to bring the pieces together and to build implementations faithful to the semantics.

Several proposals have now been made to account for possible relations between anaphora and antecedents within and across sentences, including treatments of the interaction between quantifiers and anaphors [W, K, H, B]. Kamp and Heim both introduce an intermediate level of representation of sentences and give a semantics to that level while Barwise's syntax-directed interpretation relation also takes as arguments input and output contexts which are partial assignments of values to variables. However, none of these treatments takes into consideration the domain constraints, perhaps better described as preferences, or those discussed by Grosz et al. [G] under the name of centering.

I used to think of the interpretation of non-declarative sentences as being outside the domain of semantics; it wasn't clear how semantics was compatible with the attribution of illocutionary force (or speech act type) to utterances. But a uniform semantics for different sentence types is possible if one takes the interpretation of an utterance (of any number of sentences) to be a relation, determined by mood,

intonation, and propositional content, between the joint mental state of the participants before and after the utterance. The speech acts (illocutionary and perlocutionary) applicable to the utterance are then determinable as a function of the initial mental state of the speaker (in the case of illocutionary acts, such as *assert* and *warn*) or of both initial and final mental states of speaker and hearer (in the case of perlocutionary acts, such as *convince* and *scare*). The beginning of such an analysis can be found in [P]. There are other good reasons for this move. It treats utterances as first-class actions, making it possible to use them in complex actions including non-linguistic acts. It also makes it easier to show the relation between illocutionary acts and the related intended perlocutionary effects. It also makes it possible to show what are the consequences of utterances which, intentionally or not, violate the normal conventions, as is the case with ironical, indirect and insincere utterances. The key feature of the analysis is its dependence on *default rules*, formulated within Reiter's non-monotonic *default logic* [R], to express the consequences of utterances, making clear that the conveyance of propositional attitudes constituting a speech act is based on the use of utterances with certain features in mental states appropriate to the attitudes being conveyed, but that, in general, the speaker need not be in the requisite mental state.

Several other discourse constraints have been studied in the framework of various non-monotonic logics (e.g., presupposition [M], anaphora resolution [D], conversational implicature [J].) So have questions with obvious relation to discourse such as plan recognition [KA] and temporal reasoning [S, L]. Script-based reasoning should allow the defeasible inference of the existence of various entities, including events, from the statement of the existence of related entities.

Although non-monotonic reasoning seems to hold the key to a wide range of discourse phenomena, the developments above appeal to a range of different systems: in fact, no two of the papers mentioned above use exactly the same system. One promising area of unification of the various systems is through their use of *preferred* models. It may be useful to start with a familiar case. In Montague grammar (based on intensional logic, a monotonic system), it is possible to specify constraints on lexical items (e.g. the fact that *seek* is equivalent to *try to find*, or that the subject of *seek* is extensional, or human) in a set of meaning postulates. These postulates are used to restrict the class of models considered: a sentence *s* is valid iff it is true in all models satisfying the meaning postulates. Similarly, various non-monotonic theories are given semantics by restricting the models in which they are interpreted. McCarthy's circumscription, for example, makes use of models in which certain predicates are restricted to their smallest extensions; Shoham's logic of chronological ignorance depends on what he calls chronologically maximally ignorant models, and Kautz's logic of plans minimises over several dimensions, including, e.g., the number of steps in a plan.

The whole area of non-monotonic reasoning is in a state of great flux, in part because of the diversity of systems and the technical difficulty of arguments within and between the various systems [HM]. Nevertheless, the simple fact that no other approach comes close to dealing with so many of the relevant problems suggests to me that three questions should be investigated on the way to a semantic theory of discourse:

- whether other "pragmatic" problems can be couched in the same terms (e.g., noun-noun modification, metonymy, word-sense selection),
- whether there is a general enough notion of preference to cover all these cases,
- whether a positive answer to these questions can be translated into processing algorithms.

Maybe we'll know the answers at the next meeting.

### Acknowledgement

This paper was made possible by a gift from the Systems Development Foundation.

### References

- [B] Barwise, K.J., Noun phrases, generalized quantifiers and anaphora, *CSLI report* 86-52, 1986.
- [D] Dunin-Keplicz, B., Default reasoning in anaphora resolution, *ECAI-84*, 1984.
- [G] Grosz, B.J., Joshi, A.K., Weinstein, S., Providing a unified account of definite noun phrases in discourse, *ACL-21*, 1983.
- [HM] Hanks, S., McDermott, D., Default reasoning, nonmonotonic logics, and the frame problem, *AAAI-86*, 1986.
- [H] Heim, I. Definite and Indefinite Noun Phrases, PhD Dissertation, Univ. of Massachusetts, Amhest, 1982.
- [J] Joshi, A.K., Webber, B.L., Weischedel, R.M., Some aspects of default reasoning in interactive discourse, Univ. of Penn. Tech Note MS-CIS-86-27, 1986.
- [K] Kamp, H. A theory of truth and semantic representation, in J. Groenendijk et al (eds.) *Formal Methods in the Study of Language*, Amsterdam, Math. Center, 1981.

[KA] Kautz, H.K., Allen, J.F., Generalized plan recognition, *AAAI-86*, 1986.

[L] Lifschitz, V., Pointwise circumscription, *AAAI-86*, 1986.

[M] Mercer, R. and Reiter, R., The representation of defaults using defaults, UBC Tech. Report, 1982.

[P] Perrault, C.R., An application of default logic to speech act theory, in preparation.

[R] Reiter, R. A logic for default reasoning, *Art. Int.*, 13, 1980.

[S] Shoham, Y., Chronological ignorance, *AAAI-86*, 1986.

[W] Webber, B.L., So what can we talk about now?, in Brady and Berwick (eds.), *Computational Models of Discourse*, MIT Press, 1983.