Discourse Representation Theory and the Semantics of Natural Languages:

Contribution to a Panel on "Discourse Theory and Speech Acts"
TINLAP3

F. Guenther
Universität Tübingen/FNS

0. Some Questions

The present panel was asked to elaborate on the following series of questions:

- [1] Is there yet any serious discourse theory with testable computational and empirical consequences?
- [2] What phenomena ought a processing theory of discourse understanding/generation to address itself that are not already being attended to currently?
- [3] What aspects of discourse are language problems and which are general AI/KR problems?
- [4] What makes a theory of discourse a processing theory?
- [5] Does spoken language affect one's theory of discourse?
- [6] Is there any real hope that we will be able to recognize the plans/goals etc. of a speaker?
- [7] How much of conversation is carried on through the linguistic window anyway?
- [8] Do current theories of text and dialogue mesh, and should they?

But one additional question before we try to answer some of the above questions:

- [9] How would the answers to questions 1 - 8 differ if instead of "discourse" or "text" and "dialogue" (or whatever) had we been asked these questions in terms of just the concepts of "sentence" or "language" ?

Let me try and answer this question before coming to (some of) the others.

For anyone who is not convinced that the definition of linguistic meaning must be based on something else than the "sentence", the answer to Question 9 must necessarily and lapidarily be something like "in nothing" - and with good reason. Of course this answer is not right: not because there is a panel that is supposed to give more interesting answers, but because the facts of language elicit another answer.

The basic "unit" of linguistic meaning is the discourse and the basic concept of meaning is one that involves quite different semantic ingredients than those that previous attempts to formulate a semantics for natural language have made us familiar with. We take it that the concept of linguistic meaning should characterize a relation between "information states" (the term - as well as many ideas in this note - originates from two papers by Richard Smaby [1978] and [1981]). In other words we view the (linguistic) meaning of a discourse - and obviously also that of a single sentence - as something like a function which transforms information states into (usually but not always)

new information states. In fact, every linguistic construct below the discourse level, be it a phrase or even a basic expression, should ultimately be viewed in this manner. A systematic characterization of this relation will reveal that for just about every linguistic construct we are led to assume a much more intricate account of the "meaning" they convey than the typical "satisfaction" and "truth clauses" that we once assumed in the context of theories of meaning whose main goal was the systematic characterization of the models which verified isolated sentences. Traditional formal (or model theoretic) semantics is by no means falsified by this view of discourse semantics; on the contrary: what is made clear by the new account is that so-called truth relations (as applied to sentences or stretches of discourse) are no longer the primary relations but rather derived from a more fundamental kind of semantic relation.

The answers to the above questions will as a result be couched in terms of how information states can be changed by the incorporation of new bits of information.

1. Basic Features of a Theory of Discourse Meaning (or a partial answer to Q1)

1.1 What is "Linguistic Meaning"?

Even if it is more than a truism to assert that one of the basic functions of linguistic communication is to convey "information" we should not be misled into assuming too strict a link between the two. Certainly the traditional notion of "information equivalence" in terms of what models (in the logical sense) are characterized by an utterance or sets of utterances is a most useful one. But this notion is too weak as a characterization of linguistic meaning; there are in general many sentences that - given a model - turn out to be equivalent, even though their linguistic properties are quite different. We claim that (logically) equivalent sentences can play very different roles in the setting of discourse. Observe to begin with that for many sentences the traditional approach in terms of truth conditions does not make much sense. Sentences with pronouns or various "elliptical" constructions for example cannot even be evaluated (i.e. semantically characterized) on their own. Nevertheless they have a linguistic meaning; but their meaning must be reconstructed as something else than simple truth conditions. Their meaning is inherently their effect - i.e. the change they can induce - on an antecedently given information state.

1.2 Expressions, Representations, Models

To be more precise we should say that questions of meaning should be treated at three different levels and in a particular order:

- the linguistic meaning of an expression (be it an atomic expression, a phrase, a sentence or a bit of

discourse) is what determines its contribution to information states, or better, the conditions of its contribution to an information state.

- the representational meaning of an expression is the particular "change" the expression effects on a given information state.

- the model-theoretic meaning of an expression are the satisfiability conditions induced by the change given a particular class of models for the information states.

There is a lot to be said about how these "levels" of meaning interact; for a discussion of this issue cf. Guenther [1985a]. What should be pointed out here - albeit briefly - is that the general theory of meaning should characterize the representational and model-theoretic properties of expressions in terms of the linguistic meaning and not the other way around. In other words, it is the (characterization of the) linguistic meaning which gives rise to its model-theoretic (or truth conditional) properties and not vice-versa (as most recent work in formal semantics might have led us to believe).

1.3 Two Types of Semantic Relations

The above distinction between different levels of "meaning" is closely related to a distinction - first drawn in a systematic way by Hans Kamp in the framework of Discourse Representation Theory - between two kinds of semantic relations. Prior to Kamp's work most if not all semantic properties and relations were defined in terms of truth-conditional terms, ultimately therefore in terms of the most basic model-theoretic relation: logical consequence. We shall call semantic relations definable in terms of consequence "T(ruth)-relations" and we shall oppose them to a perhaps more interesting class of semantic relations that we shall call "D(iscourse)-relations". T-relations (most notably consequence, truth and consistency) are relations that characterize the relation of an information state and a model, but T-relations are not the only kind of relations between informations and models. D-relations on the other hand characterize relations between information states. Both types of relations are in general applied to sentences and stretches of discourse, but they can be generalized to expressions below as well as above the sentence. Typical D-relations are presupposition, ambiguity and coherence. But there are many more. The topic of D-relations is pursued in more detail in Guenther [1985b].

1.4 Aspects of the Theory of Communication

All of the above makes a lot more sense once one considers the role of information states in the setting of a theory of communication. Sentences but more generally stretches of spoken and written

discourse are used to "convey" information. In other words they are used to transform information states. But how and with what "intention"? A simple but plausible answer (reminiscent of Kamp [1985]) goes as follows: in communicating a speaker has in general to circumscribe a sub-information state from his own overall representational set-up in such a way that there is a "choice of words" which allows the linguistic encoding of that sub-information state; the decoding of the resulting discourse by the hearer should then result in the installation of an "equivalent" sub-information state in the overall representational set-up of the hearer. The rules which govern both the encoding and the decoding of discourses are called discourse structure construction rules in Kamp [1981]. It is this algorithm which defines the concept of "linguistic meaning" mentioned above. We claim that the regularities of linguistic meaning are exactly captured by this algorithm.

1.5 Understanding "Understanding"

This is a far cry from a theory of understanding, but there is at least something to be said which distinguishes "understanding" from other concepts introduced so far. On our view the result of the successful decoding of a discourse results in an "interpretation" of the discourse with respect to the given antecedent information state. And "understanding" begins where interpretation ends: how one understands depends to a high degree what other available information the result of interpreting a sentence or discourse can interact with. Since no two hearers are exactly in the same information state it can hardly come as a surprise that the ramifications of the incorporated information will not play exactly the same role in their mental states. This is true even for rather banal utterances (like "it is four o'clock") and obviously much more for utterances that involve more intricate conceptual structures be these of the scientific, literary, sociological or whatever sort. Moreover, it seems clear that there can be effects of "understanding" without interpretation. For the inability to provide an interpretation also has side-effects; if the transfer of information fails (e.g. because the hearer simply doesn't understand the language that is used) or because central presuppositions are not in place, the hearer can very well draw conclusions which have little or nothing to do with the linguistic meaning of the discourse. And it would be a great mistake to identify (or even associate) those conclusions with the linguistic meaning of the utterances employed. (For a more extensive discussion of these matters, cf. Chapter 6 in Guenther & Sabatier (to appear).)

2. A Computational Account of DRT

2.1. The Fragment in Kamp [1981]: Some Remarks

An important step towards formulating a theory of discourse analysis was taken in Kamp [1981]

and later work by Kamp and others within the framework of Discourse Representation Theory (DRT). As presented in that paper, the fragment has a straightforward computational interpretation: let d be a discourse consisting of the sentences $s_1, \dots, s_n$, then the Discourse Representation Structure (DRS) associated with d (via the so-called DRS construction algorithm A , the core part of DRT) is $A(d)$, i.e. $A(s) \circ A(s_1) \dots \circ A(s_n)$, where the "result" of applying A to a sentence s depends on the DRS derived from the previous sentences.

An extension of the original fragment of DRT was first implemented in Prolog in Kolb [1985]; this implementation also had the merit of providing a deductive account (restricted to a generalized syllogistic language) for DRSs.

2.2 *Implementing DRT in Prolog*

Prolog lends itself in a natural way as an implementation language for DRT. (A faithful and at the same time optimal implementation of the fragment in Kamp [1981] takes up about 11/2 pages of Prolog code.) Several alternatives to the implementation of DRT (in Prolog) have been investigated (cf. Guenther & Sabatier [to appear] Chapter 6); the most recent implementation takes as its top-level predicate the relation mentioned above between an antecedent context, a resulting context and a discourse. A context is taken to consist of a DRS together with an (ordered) list of possible antecedents for pronouns. (Similar predicates apply to all syntactic categories in the fragment.) The implementation is reversible and thus generates DRSs for bits of discourse as well as bits of discourse from DRSs. Among the new features of this implementation, we should cite among others the treatment of: disjoint reference, possessive and reflexive pronouns, forward anaphora (e.g. the Migs & Pilots variety) as well as a simple but effective treatment of quantifier scope ambiguities. For instance, the scope effects of sentences like **every man loves a woman** are dealt with by treating the "weakest" reading first; if later sentences force a stronger reading, backtracking will induce a DRS where the indefinite noun phrase is given a "topmost" interpretation, i.e. by introducing the discourse reference of that noun phrase in the principal DRS (and nowhere else). There is some linguistic evidence that for this kind of quantifier interaction there are no other plausible possibilities. In future versions of this implementation we shall exploit the fact that the program is reversible for certain applications requiring generation.

Finally, we can extract information from DRSs by translating them into predicate logic; a theorem prover based on Smullyan's tableaux system has successfully been used to derive answers from DRSs in a deductive manner. In addition to semantic evaluation (with respect to a model or database) we thus have another way of using DRSs for interactive information processing.

References

- Barwise, J. & J. Perry **Situations and Attitudes**, MIT Press, 1983.
- Guenthner, F. [1985a] "Linguistic Meaning in Discourse Representation Theory", FNS-Bericht-85-4, to appear in **Synthese**.
- Guenthner, F. [1985b] "Zwei Typen von semantischen Relationen", FNS-Bericht-85-5, to appear in O. Herzog, **Wissensbasierte Systeme und Natürlich Sprache**, Springer.
- Guenthner, F. & P. Sabatier **An Introduction to Natural Language Processing in Prolog**, to appear.
- Kamp, H. [1981] "A Theory of Truth and Semantic Representation", in J. Groenendijk et al. (eds.) **Formal Methods in the Study of Languages**, Foris, 1984.
- Kamp, H. [1983] "SID without Time and Questions", CSLI-Report, to appear.
- Kamp, H. [1985] "Context, Thought and Communication", Aristotelian Society, New Series, 1984/1985.
- Kolb, H.P. [1985] **Aspekte der Implementation der Diskursrepräsentationstheorie**, FNS-Skript-85-1, Universität Tübingen.
- Lewis, D. [1979] "Score-Keeping in a Language Game", **Journal of Philosophical Logic**, 1979
- Smaby, R. [1978] "Ambiguous Co-reference with Quantifiers" in: F. Guenthner & S.J. Schmidt (eds.) **Formal Semantics and Pragmatics for Natural Languages**, Reidel, 1978.
- Smaby, R. [1981] "Pronouns and Ambiguity: A Sample Case", in U. Mönnich (ed.) **Aspects of Philosophical Logic**, Reidel, 1981.