41 Computational Perspectives on Discourse and Dialog

BONNIE LYNN WEBBER

0 Introduction

Computational work on discourse and dialog reflects the two general aims of natural language processing:

- that of modeling human understanding and generation of natural language in terms of a system of computational processes. Work in this area is usually called **computational linguistics**.
- that of enabling computers to analyze and generate natural language in order to provide a useful service. Work in this area has been called **applied natural language processing**, **natural language engineering**, or more recently **language technology**.

These aims go back as far as the earliest research and development in natural language processing (NLP), which began with work on machine translation in the early 1950s. Early machine translation work pointed out serious problems in trying to deal with unrestricted, extended text in weakly circumscribed domains. This led NLP researchers in the 1960s and early 1970s to focus on question-answering dialogs in restricted domains, such as baseball games in Green et al. (1961), airline schedules in Woods (1968), analyses of lunar rocks in Woods et al. (1972), and a "blocks world" in Winograd (1973). But as the development of meaning representations and reasoning needed for effective language processing became less and less language issues, the attention of NLP researchers shifted from developing natural language systems to solving individual language-related problems – e.g. developing faster, more efficient parsers; developing "weaker" and hence more realistic grammars whose complexity is only slightly more than context-free (cf. Joshi 1999); developing ways of handling referring expressions; modeling communicative goals and plans and their realization in language, etc. But now we have come full circle, and the recent explosion in information available over computer networks, and demands for less frustrating automated telephone-based service facilities made possible by advances in speech technology, have refocused interest on dealing with unrestricted extended text and dialog.

With new attention being paid to discourse and dialog, the aims of computational work in these areas can be seen to be similar to those of NLP in general:

- that of modeling particular phenomena in discourse and dialog in terms of underlying computational processes;
- that of providing useful natural language services, whose success depends in part on handling aspects of discourse and dialog.

By "phenomena" in discourse and dialog, I mean either (1) a word, phrase, and utterance whose interpretation is shaped by the discourse or dialog context, or (2) a sequence of utterances whose interpretation is more than the sum of its component parts. What computation contributes is a coherent framework for modeling these phenomena in terms of resource-limited inferential search through a space of possible candidate interpretations (in language analysis) or candidate realizations (in language generation).

Inference here refers to any form of reasoning. The reasoning may be **nondefeasible**, according to logical principles that guarantee the correctness of its conclusions, as in correctly concluding from "John went to the zoo again" that John had gone to the zoo at least once before. Or the reasoning may be **defeasible**, producing plausible conclusions that are not necessarily correct, as in concluding from "John went to the zoo. He saw an owl" that John had seen the owl at the zoo.

Search refers to how one goes about determining discourse interpretation: there are often several *possible* ways to interpret a word, phrase, utterance, or sequence of utterances in context, and one needs to find the intended, or at least the most likely, one. **Inferential search** refers to the roles that inference can play in this process: it can serve to (1) grow the search space in which the interpretation of an utterance will be found (or alternatively, the search space in which the surface realization of some underlying conceptual form will be found), or (2) provide evidence relevant to evaluating candidate interpretations or surface realizations, or both. For example, in:

(1) a. John arrived at an oasis. He saw the camels around the water hole and . . .
b. John arrived at an oasis. He left the camels around the water hole and . . .

inference can play one or both roles in interpreting the definite noun phrase "the camels." It can be used to link the camels to the oasis or to the means by which John got there. (This use of inference is sometimes called **bridging**.) And it can also be used in choosing which interpretation is more plausible – camels already at the oasis in (1a), since they are something John might observe and whose observation might be mentioned, and camels that John brought with him to the oasis in (1b), since they are something he could then leave.

Resource-limited refers to the fact that the computational processes used in discourse and dialog do not have unlimited time or memory in which to carry out the search. Resource-limited search can manifest itself in terms of restrictions on the context from which search begins and/or as constraints on the way the search space can develop. For example, if there is a cost associated with inference, as in Hobbs et al. (1993) and Thomason and Hobbs (1997), that cost can be used to direct the growth of the search space toward low-cost solutions or to prune more expensive ones from it.

(It can also be used to choose the lowest-cost interpretation among those that can be completed, but that would not be a resource-limited process, as it would require first producing them all.)

This is not to imply that all computational work on discourse and dialog involves resource-limited inferential search. Recent language technology work on discourse (mainly coreference identification) and dialog (mainly call routing and other simple service interactions) exploits probabilistic methods based on frequencies gathered from large tagged corpora. I will say a bit more about this in section 2.

Section 1 of this chapter provides a brief discussion of computational models of discourse and dialog from the perspective of computational linguistics. Section 2 describes language technology in the area of discourse and dialog, while section 3 speculates on future directions and developments.

More extensive discussion of recent computational research and development can be found in the individual papers cited throughout this chapter, in textbooks by Allen (1995) and by Jurafsky and Martin (2000), in a survey by Cohen (1996), and in the websites of the Association for Computational Linguistics' Special Interest Group on Discourse and Dialogue (SIGDial) (http://www.sigdial.org) and the Language Engineering Telematics project, MATE (http://mate.nis.sdu.dk/).

1 Discourse, Dialog, and Computation

1.1 Computational models of cognitive processes in discourse and dialog

Many aspects of language have their use and interpretation shaped by the discourse context:

- forms of reference, such as pronominal anaphora and deixis, and definite and deictic noun phrase (NP) reference;
- certain forms of ellipsis such as VP ellipsis, sluicing – e.g. "I know John goes swimming on New Year's Day but I don't know why" – and background ellipsis – e.g. Q: "Will the shop open in June?" A: "No." Q: "In July?" (Other forms of ellipsis, such as gapping and conjunction reduction, are generally considered purely within the domain of syntax and do not appeal to the resources or processes associated with discourse.)
- the interpretation of clauses in terms of eventualities and their temporal, causal and rhetorical relations to one another.
- aspects of intonation and syntactic choice generally associated with **information structure** (i.e. notions of theme/rheme and background/focus).

What these phenomena share are *constraints* on their use, associated with a continually changing context that they contribute to, and reliance on *inference* to find and/or verify candidate intepretations. These features come from the resource-limited inferential search processes that underlie their generation and interpretation.

Consider, for example, pronominal reference. One of the earliest computational models of pronoun reference appears in LUNAR (cf. Woods 1978), which allowed geologists to pose English-language queries to a large database concerning the A*pollo 11* lunar samples. LUNAR's treatment of pronominal anaphora in follow-up questions such as:

(2) User: Do the breccias contain olivine?
LUNAR: . . .
User: Do they contain magnatite?

(3) User: What is the silicon content of each volcanic sample?
LUNAR: . . .
User: What is its magnesium concentration?

followed Karttunen (1976) in taking pronouns and definite NPs to refer to **entities** in a model of the discourse. In LUNAR, entities could be evoked through indefinite and definite NPs in a user's query, and referenced in the same or a subsequent query. Only the ten most recently evoked or referenced entities were considered possible referents for a subsequent pronoun or definite NP. Entities were tested for semantic fit in order of recency, with the first to fit taken to be the intended referent. This had the side effect of updating the referent's position in the **reference list**, removing it from its current position and inserting it at the start of the list, thereby delaying its dropping off the end. Recent theories of contextual reference based on an approach to contextual modeling called **centering**, developed in the mid-1980s by Grosz et al. (1995), have similar features.

Centering follows work by Sidner (1982) in imposing a finer structure on context than LUNAR, by assigning to each utterance in a discourse both a unique **backward-looking center $\mathbf{C_b}$** and a rank-ordered list of **forward-looking centers $\mathbf{C_f}$**. The C_f-list for one utterance comprises the possible candidate referents for pronouns in the next utterance. One question is how to structure this search, and different ordering metrics have been proposed for different languages (for Italian in Di Eugenio 1997, for Turkish in Turan 1995, and for Japanese in Iida 1997). Another question is how to use the C_b in identifying a preferred solution. For example, Brennan et al. (1987) introduced the idea of **center transition preferences** that prefer interpreting a pronoun in a way that retains the same C_b between utterances, or barring that, only changes it in particular ways. The C_b and C_f-list are then updated at the end of each utterance.

Brennan et al.'s treatment is not incremental. In contrast, Strube (1998) proposed a simpler form of centering that returns to models such as LUNAR in (1) abandoning the backward-looking center C_b and center transitions and (2) using only a finite ordered list of salient candidates. This allows updating to take place as soon as a referring expression is processed, with an entity's insertion into the list determined by how the speaker has chosen to specify it with respect to the "familiarity scale" given in Prince (1981). In this scale, Prince distinguishes between entities presented as **new** to the discourse, entities presented as already **evoked** by the discourse or the outside situation, and entities presented as **inferable** from something already introduced into the discourse. A feature of this scale is that well-known individuals, when first introduced into a discourse, are nevertheless considered new (in Prince's terms, **unused**). In Strube's incremental approach, if an entity is already on the list, its position on the list may change on subsequent reference, reflecting how it has been specified. Besides being simpler, Strube's algorithm better reflects intended interpretations than other centering algorithms, although it still does not provide a complete account of pronominal reference.

It should be noted that centering and earlier focus models have also been used to guide decisions about the use of pronouns in generating text in work by McKeown (1985) and by Dale (1992), though the decision process is not simply the reverse of that used in interpretation. More recently, McCoy and Strube (1999) have considered whether considering changes in temporal focus could explain a speaker's decision to use a name or definite NP where centering allows the use of a pronoun: it is a better model, but still incomplete.

Computational models of other discourse phenomena – including other forms of contextual reference – highlight other features of the resource-limited inferential search that can be seen to underpin their processing.

1.1.1 Definite NPs

The intended referent of a definite NP need not have been explicitly mentioned in the prior discourse, as long as it can be inferred from what has been. For example, in:

(4) Phone "Information". *The operator* should be able to help you.

the definite NP *the operator* refers to the telephone operator you reach when phoning information. The referent of a definite NP can but need not be a member of the set of initial candidates that a reference resolution process begins with. Computational research attempts to specify not just what these additional candidates may be, but the specific search processes by which they will be found and the intended referent correctly identified, as in Bos et al. (1995); Hahn et al. (1996); Hobbs et al. (1993); Markert and Hahn (1997). From the perspective of text generation, choosing whether to use a definite NP (and, if so, choosing one sufficient to refer uniquely to the intended referent) involves both search and inference for other entities in the context that block referential uniqueness, and search for properties that distinguish the intended referent from the remaining others, as in Dale and Haddock (1991); Dale and Reiter (1995); Horacek (1997); Stone and Doran (1997); Stone and Webber (1998).

1.1.2 Demonstrative pronouns

These expressions highlight the need for an augmented candidate set for reference – not only the individuals and/or sets evoked by individual NPs (or sets of NPs) but also properties and eventualities evoked by predicates, clauses and larger units of discourse (**discourse segments**). For example:

(5) Phone "Information". *That* should get you the information you need.

As discussed in Webber (1988, 1991) and later in Asher (1993) and Stone (1994), resolution of demonstrative pronouns appears, in part, parasitic on an update process for discourse segments to provide possible referents. Where demonstratives refer to individuals, Davies and Isard (1972) have pointed out the role of stress in preferring one candidate over another in resolving a demonstrative pronoun versus an anaphoric pronoun:

(6) Think of a number, square it, and then multiply [*it*, *that*] by three.

In NL generation, I am not aware of any more recent attempt to articulate the processes involved in generating demonstrative pronouns than the work of Davey (1974), generating explanations of what happened in a game of tic-tac-toe.

1.1.3 Clausal relations

It has long been noted that a discourse composed of a sequence of clauses requires recognizing intended relations between them (often called **coherence relations**), although Scott and de Souza (1990) and others have pointed out that similar relations hold between phrases and between phrases and clauses as well. Such relations have been taken to contribute to the underlying substructures and their interpretation as explanations, descriptions, proposals, corrections, etc. For example, one must recognize the different relations between the clauses in (7a) and in (7b), in order to understand them correctly:

(7) a. Phone "Information". The operator will have the number you want.
b. Phone "Information". It won't cost you anything.

Rhetorical Structure Theory, as presented in Mann and Thompson (1988), posits a fixed set of relations with constraints on their applicability, but not how they would be used in any kind of process involved in understanding or generation. Identifying clausal relations appears resource-limited in two ways: in establishing what the current clause is related to – the previous clause or some larger segment in which it is embedded – and in establishing what relation(s) hold between them. With respect to the former, while a speaker may be describing more than one event or situation at a time or connecting up many strands into an explanation, the listener, nonetheless, appears limited in terms of how many things she or he can be attending to or keeping in mind simultaneously and on how she or he can use evidence in deciding how a new clause fits in. Computational work here has focused on the updating process, including the role of tense and aspect as evidence for what should be updated and how. Relevant work here includes that of Hitzeman et al. (1995); Kameyama et al. (1993); Kehler (1994); Moens and Steedman (1988); Webber (1988). Different inferential processes that could be used in recognizing the intended relations between clauses within a discourse are described in Hobbs et al. (1993); Lascarides and Asher (1993); Thomason and Hobbs (1997). Discussion of bases for relating clauses in discourse can be found in Grosz and Sidner (1986); Moore and Pollack (1992); Moser and Moore (1996); Webber et al. (1999b, 1999c).

1.1.4 Information structure

Information structure deals with: (1) what a speaker conveys as being the topic under current discussion and, consequently, his or her contribution to that topic (**theme** vs. **rheme**), and (2) what a speaker takes to be in contrast with things a hearer is or can be attending to (**focus** vs. **background**) (cf. Halliday 1967b, 1970; Steedman 2000). Information structure manifests itself in both sentential syntax and intonation.

Just as interpreting a clause as an eventuality requires identifying its temporal, causal, and/or rhetorical relations with others in the discourse, the process involved in recognizing the theme of an utterance also requires recognizing its relation to the theme of the previous utterance or, more generally, to context. This again requires an inferential search process. So too do elements marked intonationally or syntactically as being in contrast require searching through a limited set of elements that could serve as a source of contrast, and inferring the intended alternative set to which both source and contrast item belong. This again is a resource-limited inferential search process. Less work has, to date, been done on characterizing and modeling these processes, but cf. Hajicova et al. (1995); Prevost (1995); Prevost and Steedman (1994); Steedman (1996a). Interest in the area is growing due to its use in improving intonation in spoken language generation.

1.1.5 Repetition and restatement

Speakers have been observed to often restate information already introduced into a dialog. This would contradict the Maxim of Quantity in Grice (1975), unless, as suggested in Walker (1996a, 1996b), there are resource-limitations on the propositions a listener can be attending to and all propositions needed to draw an inference must be attended to simultaneously. In Walker's model, recently introduced or mentioned propositions are held in an unordered cache (rather than an ordered list), and various cache management strategies are explored to see which correspond more closely to observed human behaviors.

There are other discourse phenomena whose interpretation depends on context – from the contextual presuppositions of individual words such as "also" and "other" (cf. Bierner and Webber 1999) to the contextual presuppositions of clauses headed by "when" and "since." Eventually, all such phenomena should be brought within the purview of a computational account framed in terms of resource-limited inferential search.

1.2 Computational models of rational agency

Discourse and dialog pragmatics (including speech acts, relevance, Gricean maxims, etc.), in the procedural view taken here, emerge from considerations of speaker and hearer as rational agents. Rational agency views discourse and dialog as behavior arising from and able to express an agent's beliefs, desires, and intentions (i.e. what the agent is committed to achieving), constrained by its resource limitations, as described in Bratman (1987); Bratman et al. (1988). Both **planning** – the process that maps an agent's intentions into actions, primarily communicative in the case of speakers as agents – and **plan recognition** – the process by which a hearer recognizes what the speaker is trying to accomplish – are resource-limited inferential search processes. However, they are shaped by two factors beyond those discussed in the previous section:

- The context in which they operate *changes continually* in consequence of actions.
- The changing context will only ever be *partially known*.

The former means that any look-ahead or precomputations they do must reflect the fact that beliefs and intentions of speaker and hearer can evolve or even change precipitously during the course of a discourse or dialog. The latter means that these processes must be able to elicit essential information; to provide useful output on the basis of assumptions as well as facts; and to modify or efficiently recompute new output if and when these assumptions are found inconsistent or wrong.

The basic framework for this work comes from the "goal (intention) begets plan" approach to planning developed and used in artificial intelligence since the late 1960s, following ideas in Newell and Simon (1963). The most widely known version is called the STRIPS algorithm, described in Fikes and Nilsson (1971). The data structures used by this algorithm capture such elements of intention and action (including communicative action) as the fact that actions have preconditions that must hold for them to have their intended effects, and that they may therefore be themselves adopted as goals realizable through further communicative actions; and that actions may have several different effects on the world. Later versions incorporated additional features such as a view of actions at different levels of aggregation and abstraction, in work by Di Eugenio (1998), Di Eugenio and Webber (1996), and Moore (1995); actions that can be done to acquire information, which can then affect the further plan or trigger further planning; and the fact that changing an agent's beliefs can cause him or her to adopt particular goals, etc., in the work of Allen (1995); Appelt (1985); Cohen and Perrault (1979); Litman and Allen (1990). This is all well described in Allen (1995).

More recently, researchers have begun to develop more complex computational models of language as rational planned action, reflecting, inter alia:

- that the beliefs of the planner/speaker might differ from those of the hearer and even be incorrect. Pollack (1986) shows how, for a speaker's communicative actions and underlying plan to be understood with respect to her or his beliefs, the hearer must be able to infer or elicit what beliefs support the speaker's inferred plan as well as inferring that plan itself.
- that dialog can be used to explore and negotiate possible courses of action, not just accomplish action, shown in the work of Di Eugenio et al. (1998) and Lambert and Carberry (1992, 1999).
- that dialog involves a collaboration among all its participants. Thus, Grosz and Kraus (1996); Grosz and Sidner (1990); Lochbaum (1998); McRoy and Hirst (1995) all show that the planning process for achieving goals through dialog is more complex than when only a single planning agent is involved.
- that planning agents have preferences shaping the way they choose to realize goals as plans of action. Thus, both Chu-Carroll (1997) and Carberry et al. (1999) show that in an advisory dialog, the participant in the advisory role must be able to infer or elicit those preferences, as well as the advisee's possibly incorrect beliefs.

- that communicative actions – e.g. to justify one course of action over another, to explain how a process works, etc. – may not succeed in their goal, requiring the speaker to use the hearer's feedback to produce a new or augmented plan whose communicative actions will accomplish the goal or support the initial communication in doing so (e.g. through clarification or explanation), as in the work of Moore (1995) and Young et al. (1994).
- that a communicative action conveys information to achieve particular intentions – (cf. Grosz and Sidner 1986; Moore and Paris 1993; Moore and Pollack 1992); that there is a potentially many-to-many relation between information and intention (cf. Di Eugenio and Webber 1996; Moore and Pollack 1992; Pollack 1991; Stone and Webber 1998); and that information and intention must be combined in generating communicative actions and extracted in understanding them. How to do this harkens back to discussions of the **modularity** of syntax and semantics in Fodor (1983). That is, Moore and Pollack (1992) argue that the recognition of informational relations cannot be ordered a priori before the recognition of intentional relations, and vice versa. But whether, in human language processing, the processes operate nondeterministically in parallel on distinct data structures, as in Hobbs (1996), or are integrated into a single process operating on a single integrated database, as in Moore (1995), or something in between, as in Thomason and Hobbs (1997), is not clear. Nor is the optimal form of integration yet known from a purely computational engineering perspective.

The brief discussions in the next two sections will show an ever-increasing number of applications in the areas of discourse and dialog. As in the past, this will also likely act as a spur to increased theoretical understanding of discourse and dialog in terms of cognition and rational agency.

2 Discourse, Dialog, and Language Technology

As noted in section 0, computational work on discourse and dialog has been driven equally by the desire to understand these phenomena as manifestations of intrinsically computational processes and by the desire to satisfy existing or potential consumer needs. In the early days of NLP, those needs were taken to be machine translation (MT) and database question/answering. The latter drove much of the early research on discourse and dialog (cf. section 1.1 and work on "cooperative question answering" such as Cheikes and Webber 1988; Joshi 1982; Joshi et al. 1987; Pollack 1986; Webber 1986). But despite early attempts to provide NL "front ends" to database systems to handle user queries, and NL "back ends" to produce cooperative responses, the consumer base of casual users of database systems, for whom such "wrappers" were designed, never really materialized.

Recently however, there has been renewed interest in cooperative dialogs, made possible by improvements in automated speech recognition and spurred by corporate desires for automated (spoken) telephone and web-based service interactions (cf. Litman et al. 1998; Walker et al. 1998). Similarly, for most of its history, MT ignored discourse and dialog as a relevant factor in translation, but again, speech recognition has made a difference: now the effort to provide "translating telephones" requires making use of whatever sources of knowledge can be brought to bear. Finally, the recent explosion of freely available electronic text and services on the worldwide web (WWW) has become a potent driver of language technology, including work on discourse and dialog.

By and large, language technology methods aim toward broad coverage at low cost. They eschew understanding, tolerating what may, from a theoretical perspective, appear to be a high rate of errors, as long as they individually or together lead to significant improvements in overall task performance. In web-based information retrieval, such improvements may involve either increasing **precision** (i.e. reducing the large number of "false positives" in any search that tries to avoid missing too many "true positives") or increasing **recall** (i.e. increasing the number of "true positives" that might otherwise be missed when anaphora and ellipses replace more lexically "revealing" evidence). The former is being addressed indirectly, by trying to identify what parts of a text might potentially be relevant (subtopic identification) and by trying to identify the sentences in a (short) document that best reveal its content and outputting those sentences as a summary of the text, as in Kupiac et al. (1995) and Mani and Maybury (1998), thereby enabling people to make relevance judgments faster, based on a smaller portion of the text, as in the work of Hearst (1994) and Reynar (1998). Where those sentences themselves contain context-dependent discourse phenomena, efforts are made to include sufficient previous text that people can resolve them. The latter is being tackled by superficial methods of coreference resolution that may guess incorrectly in places or only attempt the easy cases (cf. Baldwin 1997; Kameyama 1997; Kennedy and Boguraev 1996).

Work is also being done on developing and using **dialog models** to support more effective telephone- and web-based computer services, including call-routing (cf. Chu-Carroll and Carpenter, 1999), emergency planning-support systems (cf. Allen et al. 1996; Heeman et al. 1998), and travel information (cf. Bennacef et al. 1996; Carlson and Hunnicut 1996; Flycht-Eriksson and Jonsson 1998; Seneff et al. 1998). A dialog model is an efficient description of standard patterns of action in a dialog, often encoded as a finite-state or probabilistic automaton to reflect the role of the current state in predicting (or constraining) the next one. The development of a dialog model thus requires two things:

- a classification scheme for dialog actions that (1) can be annotated reliably (cf. Carletta et al. 1997) on the basis of superficial evidence, and (2) can support effective predictions. Dialog acts are commonly classified functionally, at some abstract level connected with the type of task being performed (e.g., *greet*, *suggest*, *reject*, etc. as in Samuel et al. 1998 (meeting planning dialogs); *restate plan*, *elaborate plan*, etc. as in Heeman et al. 1998 (complex task-planning dialogs), also (cf. Poesio and Traum 1997; Reithinger and Klesen 1997; Traum and Hinkelman 1992). But they can also be usefully classified by the topic they address, as in Chu-Carroll and Carpenter (1999) and Jokinen et al. (1998).
- a reliable method of correlating evidence from dialog actions and their context with the classification of dialog actions, so that the dialog model can be used in speech recognition, dialog understanding, and/or response generation. The usual problem is that one does not know which combination of which surface features

– including particular vocalizations, particular words and/or phrases, particular surface-syntactic features, the class of the previous utterance(s), etc. – provides reliable evidence, including potentially different features and a different combination for each class in the scheme. So data must first be reliably annotated for features that could serve as evidence. After that, a machine-learning method such as decision-tree induction, transformation-based learning (cf. Samuel et al. 1998), or neural network learning can be used to build the classification scheme. For clarity, one may use sets of probabilistic automata, each trained to a different kind of evidence, combined using the standard calculus of probabilities, as in the dialog managers developed by Stolcke et al. (1998) and Taylor et al. (1998).

A dialog model can also be designed to make use of a **dialog strategy**, embodying decisions for how to respond to dialog actions on the part of the human user that admit a variety of system responses. Here, both machine learning and purely statistical techniques are being used to identify effective strategies and evaluate their effectiveness (cf. Litman et al. 1998; Walker et al. 1998).

3 Speculations on Future Directions and Developments

Before closing this chapter, I would like to add my speculations on where useful future developments are likely to occur in computational work on discourse and dialog.

- The development of a single integrated account of context management (updating, evolution, and retrieval) will provide better understanding of the whole range of resource-bounded, context-linked discourse phenomena, including contextual reference, information structure, and clausal relations.
- The development of an integrated account of both informational and intentional aspects of discourse and dialog will initially support more principled and effective text and speech generation systems and, eventually, understanding systems as well.
- The emergence of new tasks related to discourse and dialog will turn researchers' attention to additional communicative phenomena. For example, broadening communication channels to support "face-to-screen" or even "face-to-face" spoken interaction with computer systems will focus attention on information to be gained from a speaker's gestures and their use in enriching the speaker's message or in disambiguating it, as in the work of André et al. (1998); Cassell et al. (1994); Cassell and Stone (1999); Koons et al. (1993); McGee et al. (1998).
- Improvements in the handling of current phenomena, such as clausal reference and clausal relations, will be needed to support more difficult future tasks involving "mapping" natural language texts to formal specifications (e.g. for software, to support system construction and verification) or to terminologies (e.g. in medicine, to support knowledge discovery and refinement of practice standards).

- Just as at the sentence level, lexical semantics poses more difficult representational and reasoning problems than Montague-style formal semantics, at the discourse level, the semantics of events and actions poses as yet unsolved problems in representation and reasoning. The emergence of solutions to these problems should lead to improved performance on information retrieval and text summarization tasks, and may also support vision systems to use natural language discourse and dialog to talk about what they see as they act in the world.
- Just as sentence-level processing has sought lexically based syntactic/semantic formalisms that can facilitate both understanding and generation (cf. Tree-Adjoining Grammar, described in Joshi 1987; Combinatory Categorial Grammar, described in Steedman 1996b, 2000, etc.), similar efforts by Danlos (1997) and by Webber et al. (1999a, 1999b, 1999c) will contribute to facilitating both discourse understanding and generation.
- As in grammar modeling, where the utterances that people produce are influenced by a wide range of structural and performance factors and where probabilistic models may provide the most reliable predictions, probabilistic models used in discourse and dialog will improve as they move to incorporate more and more sophisticated models of the phenomena they aim to approximate.
- More and more on-line documents are being prepared using mark-up languages like SGML or document-type declarations specified in XML. Mark-up reflecting function (e.g. heading, citation, pie chart, etc.) rather than appearance (e.g. italics, flush right, etc.) will likely facilitate more effective information retrieval and other language technology services such as summarization and multidocument integration.

There seems no doubt that computational approaches are contributing their share to our understanding of discourse and dialog and to our ability to make use of discourse and dialog in building useful, user-oriented systems.

NOTE

I would like to thank Sandra Carberry, Barbara Di Eugenio, Claire Gardent, Aravind Joshi, Mark Steedman, and Michael Strube, who have provided me with useful direction and comments in the orientation, organization, and presentation of this chapter.

REFERENCES

Allen, James, 1995. *Natural Language Understanding*. Redwood City CA: Benjamin/Cummings, second edition.

Allen, James, Miller, Bradford, Ringger, Eric, and Sikorski, Tiresa, 1996. "A robust system for natural spoken dialogue." In *Proceedings of the 34th Annual Meeting, Association for Computational Linguistics*. University of California at Santa Cruz, 62–70.

André, Elisabeth, Rist, Thomas, and Muller, Jochen, 1998. "WebPersona: a life-like presentation agent for the world-wide web." *Knowledge-Based Systems* 11:25–36.

Appelt, Douglas, 1985. *Planning English Sentences*. Cambridge: Cambridge University Press.

Asher, Nicholas, 1993. *Reference to Abstract Objects in Discourse*. Boston, MA: Kluwer.

Baldwin, Breck, 1997. "CogNIAC: high precision coreference with limited knowledge and linguistic resources." In *Proceedings of the ACL/EACL Workshop on Operational Factors in Practical, Robust Anaphora Resolution for Unrestricted Texts*. Madrid, 38–45.

Bennacef, S., Devillers, L., Rosset, S., and Lamel, L., 1996. "Dialog in the RAILTEL telephone-based system." In *Proceedings of the Fourth International Conference on Spoken Language Processing*. Philadelphia, PA, 550–3.

Bierner, Gann and Webber, Bonnie, 1999. "Inference through alternative set semantics." In *Inference in Computational Semantics, Proceedings*. Amsterdam, 39–52. Extended version in *Journal of Language and Computation.* 1(2), Spring 2000, 259–74.

Bos, Johan, Buitelaar, Paul, and Mineur, Anne-Marie, 1995. "Bridging as coercive accommodation." In *Proceedings of the CLNLP Workshop*. Edinburgh.

Bratman, Michael, 1987. *Intentions, Plans and Practical Reason*. Cambridge, MA: Harvard University Press.

Bratman, Michael, Israel, David, and Pollack, Martha, 1988. "Plans and resource-bounded practical reasoning." *Computational Intelligence* 4(4):349–55.

Brennan, Susan E., Friedman, Walker, Marilyn, and Pollard, Carl J., 1987. "A centering approach to pronouns." In *Proceedings of the 25th Annual Meeting, Association for Computational Linguistics*. Stanford University, Stanford, CA, 155–62.

Carberry, Sandra, Chu-Carroll, Jennifer, and Elzer, Stephanie, 1999. "Constructing and utilizing a model of user preferences in collaborative consulation dialogues." *Computational Intelligence* 15(3):185–217.

Carletta, Jean, Isard, Amy, Isard, Stephen, Kowtko, Jacqueline, Doherty-Sneddon, Gwyneth, and Anderson, Anne, 1997. "The reliability of a dialogue structure coding scheme." *Computational Linguistics* 23(1):13–31.

Carlson, Rolf and Hunnicut, Sheri, 1996. "Generic and domain-specific aspects of the Waxholm NLP and dialog modules." In *Proceedings of the Fourth International Conference on Spoken Language Processing*. Philadelphia PA, 677–80.

Cassell, Justine and Stone, Matthew, 1999. "Living hand to mouth: psychological theories about speech and gesture in interactive dialogue systems." In *Proceedings of AAAI Fall Symposium on Psychological Models of Communication in Collaborative Systems*, 34–42.

Cassell, Justine, Badler, Norman, Douville, Brett, Pelachaud, Catherine, Steedman, Mark, and Stone, Matthew, 1994. "Animated conversation: rule-based generation of facial expressions, gesture and spoken intonation for multiple conversational agents." In *Proceedings of SIGRAPH'94*. Anaheim, CA, 413–20.

Cheikes, Brant and Webber, Bonnie, 1988. "The design of a cooperative respondent." In *Proceedings of the Workshop on Architectures for Intelligent Interfaces*. Monterey, CA, 3–17.

Chu-Carroll, Jennifer, 1997. "A plan-based model for response generation in collaborative negotiation dialogues." PhD thesis, Department of Computer and Information Science, University of Delaware.

Chu-Carroll, Jennifer and Carpenter, Bob, 1999. "Vector-based natural language call routing." *Computational Linguistics* 25(3):361–88.

Cohen, Philip, 1996. "Discourse and dialogue: dialogue modelling." In Ron Cole (ed.), *Survey of the State of the Art in Human Language Technology*, Oregon Graduate Institute: NSF/EC/CSLU, ch. 6.3. http://www.cse.ogi.edu/CSLU/HLTsurvey/HLTsurvey.html

Cohen, Philip and Perrault, C. Raymond, 1979. "Elements of a plan-based theory of speech acts." *Cognitive Science* 3(3):177–212.

Dale, Robert, 1992. *Generating Referring Expressions*. Cambridge, MA: MIT Press.

Dale, Robert and Haddock, Nick, 1991. "Content determination in the generation of referring expressions." *Computational Intelligence* 7(4):252–65.

Dale, Robert and Reiter, Ehud, 1995. "Computational interpretations of the Gricean maxims in the generation of referring expressions." *Cognitive Science* 18:233–63.

Danlos, Laurence, 1997. "G-TAG: a formalism for text generation inspired from tree adjoining grammar." In Anne Abeille and Owen Rambow (eds), *Tree-adjoining Grammar*, Stanford, CA: CSLI Press.

Davey, Anthony, 1974. "The formalisation of discourse production." PhD thesis, Department of Artificial Intelligence, University of Edinburgh.

Davies, D. J. M. and Isard, Stephen, 1972. "Utterances as programs." In Donald Michie (ed.), *Machine Intelligence 7*, Edinburgh: Edinburgh University Press. 325–40.

Di Eugenio, Barbara, 1997. "Centering theory and the Italian pronominal system." In Marilyn Walker, Aravind Joshi, and Ellen Prince (eds), *Centering in Discourse*, Oxford: Oxford University Press. 115–38.

Di Eugenio, Barbara, 1998. "An action representation formalism to interpret natural language instructions." *Computational Intelligence* 14(1):89–133.

Di Eugenio, Barbara and Webber, Bonnie, 1996. "Pragmatic overloading in natural language instructions." *International Journal of Expert Systems* 9(2):53–84.

Di Eugenio, Barbara, Jordan, Pamela W., Moore, Johanna D., and Thomason, Richmond H., 1998. "An empirical investigation of proposals in collaborative dialogues." In *Proceedings of COLING/ACL'98*. Montreal, Canada, 325–9.

Fikes, Richard and Nilsson, Nils, 1971. "STRIPS: a new approach to the application of theorem proving to problem solving." *AI Journal* 2:189–208.

Flycht-Eriksson, Annika and Jonsson, Arne, 1998. "A spoken dialogue system using spatial information." In *Proceedings of the International Conference on Spoken Language Processing*. Sydney, 1207–11.

Fodor, Jerry, 1983. *The Modularity of Mind*. Cambridge, MA: MIT Press.

Green, Bert, Wolf, Alice, Chomsky, Carol, and Laughery, Kenneth, 1961. "BASEBALL: an automatic question answerer." In *Proceedings of the Western Joint Computer Conference*. 219–24. Reprinted in B. Grosz et al. (eds) 1986, *Readings in Natural Language Processing*, Los Altos, CA: Morgan Kaufman. 545–50.

Grice, H. P., 1975. "Logic and conversation." In Peter Cole and Jerry Morgan (eds), *Syntax and Semantics*, New York: Academic Press, vol. 3. 41–58.

Grosz, Barbara and Kraus, Sarit, 1996. "Collaborative plans for complex group actions." *Artificial Intelligence* 86(2):269–357.

Grosz, Barbara and Sidner, Candace, 1986. "Attention, intention and the structure of discourse." *Computational Linguistics* 12(3):175–204.

Grosz, Barbara and Sidner, Candace, 1990. "Plans for discourse." In Philip Cohen, Jerry Morgan, and Martha Pollack (eds), *Intentions in Communication*, Cambridge, MA: MIT Press. 417–44.

Grosz, Barbara, Joshi, Aravind, and Weinstein, Scott, 1995. "Centering: a framework for modelling the local coherence of discourse." *Computational Linguistics* 21(2):203–25.

Hahn, Udo, Markert, Katja, and Strube, Michael, 1996. "A conceptual reasoning approach to textual ellipsis." In *Proceedings of the 12th European Conference on Artificial Intelligence*. Budapest, 572–6.

Hajicova, Eva, Skoumalova, Hana, and Sgall, Petr, 1995. "An automatic procedure for topic-focus identification." *Computational Linguistics* 21(1):81–94.

Halliday, Michael, 1967b. "Notes on transitivity and theme in English, part II." *Journal of Linguistics* 3:199–244.

Halliday, Michael, 1970. "Language structure and language function." In John Lyons (ed.), *New Horizons in Linguistics*, Harmondsworth: Penguin. 140–65.

Hearst, Marti, 1994. "Multi-paragraph segmentation of expository text." In *Proceedings of the 32nd Annual Meeting of the Association for Computational Linguistics*. Las Cruces, NM, 9–16.

Heeman, Peter, Byron, Donna, and Allen, James, 1998. "Identifying discourse markers in spoken dialogue." In *Proceedings of the AAAI Spring Symposium on Applying Machine Learning and Discourse Processing*. Stanford University, CA, 44–51.

Hitzeman, Janet, Moens, Marc, and Grover, Claire, 1995. "Algorithms for analysing the temporal structure of discourse." In *Proceedings of the Annual Meeting of the European Chapter of the Association for Computational Linguistics*. Dublin, 253–60.

Hobbs, Jerry, 1996. "On the relation between the informational and intentional perspectives on discourse." In Eduard Hovy and Donia Scott (eds), *Computational and Conversational Discourse: Papers from the NATO Advanced Research Working Group on Burning Issues in Discourse*, Berlin: Springer Verlag. 139–57.

Hobbs, Jerry, Stickel, Mark, Martin, Paul, and Edwards, Douglas, 1993. "Interpretation as abduction." *Artificial Intelligence* 63(1–2): 69–142.

Horacek, Helmut, 1997. "An algorithm for generating referential descriptions with flexible interfaces." In *Proceedings of the 35th Annual Meeting of the Association for Computational Linguistics (ACL97/EACL97)*. Madrid, Spain. Palo Alto, CA: Morgan Kaufmann, 206–13.

Iida, Masayo, 1997. "Discourse coherence and shifting centers in Japanese texts." In Marilyn Walker, Aravind Joshi, and Ellen Prince (eds), *Centering in Discourse*, Oxford: Oxford University Press. 161–80.

Jokinen, Kristina, Tanaka, Hideki, and Yokoo, Akio, 1998. "Context management with topics for dialogue systems." In *Proceedings of COLING/ACL'98*. Montreal, 631–7.

Joshi, Aravind, 1982. "Mutual beliefs in question answering systems." In N. Smith (ed.), *Mutual Belief*, New York: Academic Press. 181–97.

Joshi, Aravind, 1987. "An introduction to tree adjoining grammar." In Alexis Manaster-Ramer (ed.), *Mathematics of Language*, Amsterdam: John Benjamins. 87–114.

Joshi, Aravind, 1999. "Computational linguistics." In R. Wilson and Frank Keil (eds), *MIT Encyclopedia of Cognitive Sciences*, Cambridge, MA: MIT Press. 162–4.

Joshi, Aravind, Webber, Bonnie, and Weischedel, Ralph, 1987. "Some aspects of default reasoning in interactive discourse." In Ronan Reilly (ed.), *Communication Failure in Dialogue and Discourse*, Amsterdam: North-Holland. 213–20.

Jurafsky, Dan and Martin, James, 2000. *Speech and Language Processing*. Englewood Cliffs, NJ: Prentice-Hall.

Kameyama, Megumi, 1997. "Recognizing referential links: an information extraction perspective." In *Proceedings of the ACL Workshop on Operational Factors in Practical, Robust Anaphora Resolution for Unrestricted Texts*. Madrid, 46–53.

Kameyama, Megumi, Passonneau, Rebecca, and Poesio, Massimo, 1993. "Temporal centering." In *Proceedings of the Annual Meeting of the Association for Computational Linguistics*. Columbus, OH, 70–7.

Karttunen, Lauri, 1976. "Discourse referents." In James McCawley (ed.), *Syntax and Semantics*, New York: Academic Press, vol. 7. 363–85. Earlier circulated as a 1966 Indiana University Linguistics Club pre-print.

Kehler, Andrew, 1994. "Temporal relations: reference or discourse coherence." In *Proceedings of the 32nd Annual Meeting of the Association for Computational Linguistics (ACL), Student Session*. Las Cruces, NM, 319–21.

Kennedy, Christopher and Boguraev, Branimir, 1996. "Anaphora in a wider context: tracking discourse referents." In *Proceedings of the 12th European Conference on Artificial Intelligence*. Budapest, 582–6.

Koons, David B., Sparrell, Carlton J., and Thorisson, Kristina R., 1993. "Integrating simultaneous input from speech, gaze and hand gestures." In Mark T. Maybury (ed.), *Intelligent Multimedia Interfaces*, Cambridge, MA: MIT Press. 257–79.

Kupiac, Julian, Pedersen, Jan, and Chen, F., 1995. "A trainable document summarizer." In *Proceedings of the 18th ACM-SIGIR Conference*. Seattle, WA, 68–73.

Lambert, Lynn and Carberry, Sandra, 1992. "Using linguistic, world and contextual knowledge in a plan recognition model of dialogue." In *COLING92, Proceedings of the 14th International Conference on Computational Linguistics*. Nantes, 310–16.

Lambert, Lynn and Carberry, Sandra, 1999. "A process model for recognizing communicative acts and modeling negotiation subdialogues." *Computational Linguistics* 25:1–53.

Lascarides, Alex and Asher, Nicholas, 1993. "Temporal interpretation, discourse relations and commonsense entailment." *Linguistics and Philosophy* 16(5):437–93.

Litman, Diane and Allen, James, 1990. "Discourse processing and commonsense plans." In Philip Cohen, Jerry Morgan, and Martha Pollack (eds), *Intentions in Communication*, Cambridge, MA: MIT Press. 365–88.

Litman, Diane, Pan, Shimei, and Walker, Marilyn, 1998. "Evaluating response strategies in a web-based spoken dialogue agent." In *Proc COLING-ACL'98*. Montreal, 780–6.

Lochbaum, Karen, 1998. "A collaborative planning modal of intentional structure." *Computational Linguistics* 24(4):525–72.

Mani, Inderjeet and Maybury, Mark, 1998. *Advances in Automatic Abstracting*. Cambridge, MA: MIT Press.

Mann, William and Thompson, Sandra, 1988. "Rhetorical structure theory: toward a functional theory of text organization." *Text* 8(3):243–81.
Markert, Katya and Hahn, Udo, 1997. "On the interaction of metonymies and anaphora." In *Proceedings of International Joint Conference on Artificial Intelligence (IJCAI)*. Nagoya, 1010–15.
McCoy, Kathleen and Strube, Michael, 1999. "Taking time to structure discourse: pronoun generation beyond accessibility." In *Proceedings of the 21st Annual Conference of the Cognitive Science Society*. Vancouver, 378–83.
McGee, David, Cohen, Phil, and Oviatt, Sharon, 1998. "Confirmation in multimodal systems." In *Proceedings of COLING/ACL'98*. Montreal, 823–9.
McKeown, Kathleen, 1985. *Text Generation: Using Discourse Strategies and Focus Constraints to Generate Natural Language Texts*. Cambridge: Cambridge University Press.
McRoy, Susan and Hirst, Graeme, 1995. "The repair of speech act misunderstandings by abductive inference." *Computational Linguistics* 21(4):435–78.
Moens, Marc and Steedman, Mark, 1988. "Temporal ontology and temporal reference." *Computational Linguistics* 14:15–28.
Moore, Johanna, 1995. *Participating in Explanatory Dialogues*. Cambridge, MA: MIT Press.
Moore, Johanna and Paris, Cecile, 1993. "Planning text for advisory dialogues: capturing intentional and rhetorical information." *Computational Linguistics* 19(4):651–95.
Moore, Johanna and Pollack, Martha, 1992. "A problem for RST: the need for multi-level discourse analysis." *Computational Linguistics* 18(4):537–44.
Moser, Megan and Moore, Johanna, 1996. "Toward a synthesis of two accounts of discourse structure." *Computational Linguistics* 22(3):409–19.
Newell, Allen and Simon, Herbert, 1963. "GPS, a program that simulates human thought." In Edward Feigenbaum and Julian Feldman (eds), *Computers and Thought*, New York: McGraw-Hill. 279–93. Reprinted by AAAI Press, 1995.
Poesio, Massimo and Traum, David, 1997. "Conversational actions and discourse situations." *Computational Intelligence* 13(3):309–47.
Pollack, Martha, 1986. "Inferring domain plans in question-answering." PhD thesis, Department of Computer and Information Science, University of Pennsylvania.
Pollack, Martha, 1991. "Overloading intentions for efficient practical reasoning." *Noûs* 25:513–36.
Prevost, Scott, 1995. "A semantics of contrast and information structure for specifying intonation in spoken language generation." PhD thesis, Department of Computer and Information Science, University of Pennsylvania. IRCS TR 96-01.
Prevost, Scott and Steedman, Mark, 1994. "Specifying intonation from context for speech synthesis." *Speech Communication* 15:139–53.
Prince, Ellen, 1981. "Toward a taxonomy of given–new information." In Peter Cole (ed.), *Radical Pragmatics*, New York: Academic Press. 223–55.
Reithinger, Norbert and Klesen, Martin, 1997. "Dialogue act classification using language models." In *Proceedings of EuroSpeech'97*. 2235–8.
Reynar, Jeffrey, 1998. "Topic segmentation: algorithms and applications." PhD thesis, Department of Computer and Information Science, University of Pennsylvania.

Samuel, Ken, Carberry, Sandra, and Vijay-Shankar, K., 1998. "Dialogue act tagging with transformation-based learning." In *Proceedings of COLING/ACL'98*. Montreal, 1150–6.

Scott, Donia and de Souza, Clarisse Sieckenius, 1990. "Getting the message across in RST-based text generation." In Robert Dale, Chris Mellish, and Michael Zock (eds), *Current Research in Natural Language Generation*, London: Academic Press. 47–73.

Seneff, Stephanie, Hurley, Ed, Lau, Raymond, Pao, Christine, Schmid, Philipp, and Zue, Victor, 1998. "A reference architecture for conversational system development." In *Proceedings of the International Conference on Spoken Language Processing*. Sydney, 931–4.

Sidner, Candace, 1982. "Focusing in the comprehension of definite anaphora." In Michael Brady and Robert Berwick (eds), *Computational Models of Discourse*, Cambridge, MA: MIT Press. 267–330.

Steedman, Mark, 1996a. "Representing discourse information for spoken dialogue generation." In *Proceedings of the Second International Symposium on Spoken Dialogue*. Philadelphia, PA, 89–92.

Steedman, Mark, 1996b. *Surface Structure and Interpretation*. Cambridge, MA: MIT Press. Linguistic Inquiry Monograph 30.

Steedman, Mark, 2000. *The Syntactic Process*. Cambridge, MA: MIT Press.

Stolcke, Andreas, Shriberg, Elizabeth, Bates, Rebecca, Taylor, Paul, Ries, Klaus, Jurafsky, Dan, Coccaro, Noah, Martin, Rachel, Meteer, Marie, and Ess-Dykema, Carol Van, 1998. "Dialog act modeling for conversational speech." In *Proceedings of the AAAI Spring Symposium on Applying Machine Learning to Discourse Processing*. Stanford, CA. Extended version pub. as: Stolcke, Andreas, Ries, Klaus, Coccaro, Noah, Shriberg, Elizabeth, Bates, Rebecca, Jurafsky, Daniel, Taylor, Paul, Martin, Rachel, Ess-Dykema, Carol Van, and Meteer, Marie, 2000. "Dialog act modeling for automatic tagging and recognition of conversational speech." *Computational Linguistics* 26(3):339–74.

Stone, Matthew, 1994. "Discourse deixis, discourse structure and the semantics of subordination." Manuscript, University of Pennsylvania.

Stone, Matthew and Doran, Christine, 1997. "Sentence planning as description using tree adjoining grammar." In *Proceedings of the 35th Annual Meeting of the Association for Computational Linguistics (ACL97/EACL97)*. Madrid, 198–205.

Stone, Matthew and Webber, Bonnie, 1998. "Textual economy through closely coupled syntax and semantics." In *Proceedings of the Ninth International Workshop on Natural Language Generation*. Niagara-on-the-Lake, 178–87.

Strube, Michael, 1998. "Never look back: an alternative to centering." In *Proceedings, COLING/ACL'98*. Montreal, 1251–7.

Taylor, Paul, King, Simon, Isard, Stephen, and Wright, Helen, 1998. "Intonation and dialogue context as constraints for speech recognition." *Language and Speech* 41:493–512.

Thomason, Richmond and Hobbs, Jerry, 1997. "Interrelating interpretation and generation in an abductive framework." In *Proceedings of the AAAI Fall Symposium on Communicative Action in Humans and Machines*. Cambridge, MA.

Traum, David and Hinkelman, Elizabeth, 1992. "Conversational acts in task-oriented spoken dialogue." *Computational Intelligence* 8(3): 575–99.

Turan, Umit, 1995. "Null vs. overt subjects in Turkish discourse: a centering analysis." PhD thesis, Department of Linguistics, University of Pennsylvania.

Walker, Marilyn, 1996a. "The effect of resource limits and task complexity on collaborative planning in dialogue." *Artificial Intelligence* 85(1–2):181–243.

Walker, Marilyn, 1996b. "Limited attention and discourse structure." *Computational Linguistics* 22:255–64.

Walker, Marilyn, Fromer, Jeanne, and Narayanan, Shrikanth, 1998. "Learning optimal dialogue strategies: a case study of a spoken dialogue agent for email." In *Proceedings of COLING-ACL'98*. Montreal, 1345–51.

Webber, Bonnie, 1986. "Questions, answers and responses: interacting with knowledge base systems." In Michael Brodie and John Mylopoulos (eds), *On Knowledge Base Systems*, New York: Springer-Verlag. 365–401.

Webber, Bonnie, 1988. "Tense as discourse anaphor." *Computational Linguistics* 14(2):61–73.

Webber, Bonnie, 1991. "Structure and ostension in the interpretation of discourse deixis." *Language and Cognitive Processes* 6(2):107–35.

Webber, Bonnie, Knott, Alistair, and Joshi, Aravind, 1999a. "Multiple discourse connectives in a lexicalized grammar for discourse." In *Third International Workshop on Computational Semantics*. Tilburg, 309–25.

Webber, Bonnie, Knott, Alistair, Stone, Matthew, and Joshi, Aravind, 1999b. "Discourse relations: a structural and presuppositional account using lexicalised TAG." In *Proceedings of the 36th Annual Meeting of the Association for Computational Linguistics*. College Park, MD, 41–8.

Webber, Bonnie, Knott, Alistair, Stone, Matthew, and Joshi, Aravind, 1999c. "What are little trees made of: a structural and presuppositional account using lexicalised TAG." In *Proceedings of International Workshop on Levels of Representation in Discourse (LORID'99)*. Edinburgh, 151–6.

Winograd, Terry, 1973. "A procedural model of language understanding." In Roger Schank and Ken Colby (eds), *Computer Models of Thought and Language*, W. H. Freeman. 152–186. Reprinted in B. Grosz et al. (eds), 1986 *Readings in Natural Language Processing*, Los Altos, CA: Morgan Kaufman.

Woods, William, 1968. "Procedural semantics for a question-answering machine." In *Proceedings of the AFIPS National Computer Conference*. Montvale. NJ: AFIPS Press. 457–71.

Woods, William, 1978. "Semantics and quantification in natural language question answering." In *Advances in Computers*, New York: Academic Press, vol. 17. 1–87.

Woods, William, Kaplan, Ron, and Nash-Webber, Bonnie, 1972. "The Lunar Sciences Natural Language Information System: final report." Technical Report 2378, Bolt Beranek and Newman, Cambridge, MA.

Young, R. Michael, Moore, Johanna D., and Pollack, Martha E., 1994. "Towards a principled representation of discourse plans." In *Proceedings of the 16th Annual Conference of the Cognitive Science Society*. Atlanta, GA, 946–51.