

Toward a Synthesis of Two Accounts of Discourse Structure

Megan Moser*
University of Pittsburgh

Johanna D. Moore†
University of Pittsburgh

Among researchers interested in computational models of discourse, there has been a long-standing debate between proponents of approaches based on domain-independent rhetorical relations, and those who subscribe to approaches based on intentionality. In this paper, we argue that the main theories representing these two approaches, RST (Mann and Thompson 1988) and G&S (Grosz and Sidner 1986), make similar claims about how speakers' intentions determine a structure of their discourse. The similarity occurs because the nucleus-satellite relation among text spans in RST corresponds to the dominance relation among intentions in G&S. Building on this similarity, we sketch a partial mapping between the two theories to show that the main points of the two theories are equivalent. Furthermore, the additional claims found in only RST or only G&S are largely consistent. The issue of what structure is determined by semantic (domain) relations in the discourse and how this structure might be related to the intentional structure is discussed. We suggest the synthesis of the two theories would be useful to researchers in both natural language interpretation and generation.

1. Introduction

Within the computational discourse community, there is a long-standing debate between proponents of theories based on domain-independent rhetorical relations (most notably Rhetorical Structure Theory, Mann and Thompson 1988, henceforth RST; see also Hobbs 1985) and those who subscribe to theories based on intentionality (most notably that of Grosz and Sidner 1986, henceforth G&S). While some researchers have tried to integrate the two approaches (Moore and Paris 1993; Asher and Lascarides 1994; Hobbs 1993), the two are usually viewed as competing theories. Here we argue that G&S and RST are essentially similar in what they say about how speakers' intentions determine a structure of their discourse.

Intentional structure describes the roles that discourse actions play in the speaker's communicative plan to achieve desired effects on the hearer's mental state. Intentions encode what the speaker was trying to accomplish with a given portion of discourse. The relations between intentions indicate whether one intention contributes to the satisfaction of another (**dominance**) or whether one intention must be satisfied before another (**satisfaction-precedence**) (Grosz and Sidner 1986). In contrast, **informational structure** is concerned with domain relations among the things being talked about. Moore and Pollack (1992) argue that both intentional and informational analyses are needed simultaneously.

* Learning Research and Development Center, and Department of Linguistics, University of Pittsburgh, Pittsburgh, PA 15260. E-mail: moser@isp.pitt.edu

† Department of Computer Science, and Learning Research and Development Center, University of Pittsburgh, Pittsburgh, PA 15260. E-mail: jmoore@cs.pitt.edu

Prior research has established that recognition of intentional structure, and therefore appropriate generation of cues to such structure, is crucial for many discourse-processing tasks. Grosz and Sidner (1986) argued that intentional structure is crucial for anaphora resolution and plan recognition. Hirschberg et al. (1987) show that intentional structure plays a role in intonation. In addition, Moore (1995, in press) shows that intentional structure is crucial for responding effectively to questions that refer to prior discourse and to communication failures. Further research indicates that intentional structure is an important factor in determining when to generate discourse cues (e.g., “because,” “thus,” “although”), which discourse cues to select, and where to place those cues (Moser and Moore 1995, in preparation).

In this paper, we compare what G&S and RST say about intentional structure. We use the term **Intentional Linguistic Structure**, or ILS, as a theory-neutral way of referring to the structure of a discourse determined by the speaker’s intentions. The definition of ILS comprises one of the major claims in G&S. ILS is not directly addressed in RST, but is implicit in the RST concept of **nuclearity**. We argue that the key to reconciling ILS in the two theories lies in the correspondence between the dominance relation between intentions in G&S and the **nucleus:satellite** relation between text spans in RST. Roughly speaking, an RST nucleus expresses a G&S intention I_n , a satellite expresses another intention I_s and, in G&S terms, I_n dominates I_s . The correspondence between G&S dominance and RST nuclearity helps to clarify the relationship between ILS and informational structure, the structure determined by underlying semantic (domain) relations. In particular, the undesirable possibility that informational structure could be incompatible with intentional structure (Moore and Pollack 1992) is explained. Further, we argue that a synthesis of G&S and RST is possible because the correspondence between dominance and nuclearity forms a great deal of common ground, and because the remaining claims in the two theories are consistent.

2. Intentional Linguistic Structure in G&S

G&S is formulated in terms of the interdependence of three distinct structures. Of the three structures, it is the effect of intentional structure on **linguistic structure** that concerns us in this paper. This effect is an explicit claim about ILS.

In G&S, the intentional structure consists of the set of the speaker’s communicative intentions throughout the discourse, and the relations of dominance and satisfaction-precedence among these intentions. The speaker tries to realize each intention by saying something; i.e., each intention is the purpose behind one or more of the speaker’s utterances. Intentions are thus an extension of the intentions in Grice’s (1957) theory of utterance meaning. Speakers intend for the intentions behind their utterances to be recognized and for that recognition to be part of what makes their utterances effective. A purpose I_m dominates another purpose I_n when satisfying I_n is part of satisfying I_m . A purpose I_n satisfaction-precedes another purpose I_m when I_n must be satisfied first. The dominance and satisfaction-precedence relations impose a structure on the set of the speaker’s intentions, the intentional structure of the discourse, and this in turn determines the linguistic structure.

The linguistic structure of a particular discourse is made up of segments, which are sets of utterances, related by embeddedness and sequential order. A segment DS_n originates with the speaker’s intention: it is exactly those utterances that the speaker produces in order to satisfy a communicative intention I_n in the intentional structure. In other words, I_n is the discourse segment purpose (DSP) of DS_n . DS_n is embedded in another segment DS_m just when the purposes of the two segments are in the




Figure 1
For G&S, dominance in intentional structure determines embedding in linguistic structure.

dominance relation, i.e., I_m dominates I_n . The dominance relation among intentions fully determines the embeddedness relations of the discourse segments that realize them. For example, consider the discourse shown in Figure 1, adapted from Mann and Thompson (1988). The whole discourse is a segment, DS_0 , that attempts to realize I_0 , the speaker's intention for the hearer to adopt the intention of attending the ballet.¹ As part of her plan to achieve I_0 , the speaker generates I_1 , the intention for the hearer to adopt the belief that the ballet will be very entertaining. Then, as part of her plan to achieve I_1 , the speaker generates I_2 , the intention that the hearer believe that the show is made up of all new choreography. As shown on the left in Figure 1, I_0 dominates I_1 , which in turn dominates I_2 . Due to these dominance relations, the discourse segment that realizes I_2 is embedded in the discourse segment for I_1 , which is in turn embedded within the discourse segment for I_0 , as shown on the right in the figure. The dominance of intentions directly determines embedding of segments.

When one DSP I_n satisfaction-precedes another I_m , then DS_n precedes DS_m in the discourse. The satisfaction-precedes relation among intentions constrains the order of segments in the discourse, but it does not fully determine it. In the example in Figure 1, none of the intentions satisfaction-precedes the others.

Here we introduce a concept which is not part of the G&S theory, but which will be important to our discussion below. We coin the term **core** to refer to that part of the segment that expresses the segment purpose. A segment may contain individual utterances as well as embedded segments. Most likely, the core of the segment is found in these unembedded utterances. In the example, (a) is the core of DS_0 , (b) the core of DS_1 and (c) the core of DS_2 . As will be discussed, a core functions to manifest the purpose of the segment, while the embedded segments serve to help achieve that purpose. The defining feature of the core is its function of expressing the purpose of the segment. While the core's position in the G&S linguistic structure is most likely an unembedded utterance, it is also possible that the core could be an embedded segment. This could occur when the expression of the segment purpose is more elaborate than simply stating what the hearer should do or believe. To simplify our discussion, however, we assume the core of a segment is an utterance not embedded in any subsegment.

It should be clear that the theory-independent notion of ILS as it was characterized above is exactly the linguistic structure in G&S. ILS is something G&S makes explicit claims about. By choosing to modify the terminology from simply "linguistic

¹ In subsequent work, Grosz and Kraus have developed a model of collaborative planning in which it is not necessary for one agent to have intentions-to towards an act of another agent (Grosz and Kraus, in press; see also Grosz and Sidner 1990). For the purposes of this paper, however, we will use the notation of G&S.

structure” to “*intentional* linguistic structure,” we mean to suggest that consideration of something other than speaker intentions—for example, semantic relations—could determine another kind of structure to the discourse. Clearly, the semantic (or informational) relations among discourse entities can in principle be the determinant of a separate linguistic structure. Whether or not such an informational structure is useful or is related in an interesting way to ILS is a question requiring further research. We discuss the relationship between ILS and possible approaches to informational structure briefly in Section 5.

3. Intentional Linguistic Structure in RST

In contrast to its explicitness in G&S, ILS is only implicit in RST. To identify the implicit claims about ILS, we must first identify the components of an RST analysis that involve a judgement about the relation between intentions underlying text spans.

The range of possible RST text structures is defined by a set of **schemas**, which describe the structural arrangement of **spans**, or text constituents. Schemas are basic structural units or patterns in the application of RST relations. There are five schema patterns, each consisting of two or more spans, a specification of each span as either **nucleus** or **satellite**, and a specification of the RST relation(s) that exist between these spans. In this paper, we focus on the most commonly occurring RST schema, which consists of two text spans (a nucleus and a satellite) and a single RST relation that holds between them. The nucleus is defined as the element that is “more essential to the speaker’s purpose,” while the satellite is functionally dependent on the nucleus and could be replaced with a different satellite without changing the function of the schema. As we argue below, this functional distinction between nucleus and satellite is an implicit claim about ILS, and is a crucial notion in understanding the correspondence between RST and G&S.

A schema application describes the structure of a larger span of text in terms of multiple constituent spans. Each of the constituent spans may in turn have a structure of subconstituent spans. Thus, the application of RST schemas in the analysis of a text is recursive, i.e., one schema application may be embedded in another. To be an acceptable RST analysis, there must be one schema application under which the entire text is subsumed and which accounts for all minimal units, usually clauses, of the text. In addition, each minimal unit can appear in exactly one schema application, and the spans constituting each schema application must be adjacent in the text. These constraints guarantee that a correct RST analysis will form a tree structure.

An instantiated schema specifies the RST relation(s) between its constituent spans. Each relation is defined in terms of a set of constraints on the nucleus, the satellite, and the nucleus-satellite combination, as well as a specification of the effect that the speaker is attempting to achieve on the hearer’s beliefs or inclinations. An RST analyst must judge which schema consists of RST relation definitions whose constraints and effects best describe the nucleus and satellite spans in the schema application. Mann and Thompson claimed that, for each two consecutive spans in a coherent discourse, a single RST relation will be primary. For reasons discussed in Section 5.1, we consider only the RST presentational relations, or what Moore and Pollack (1992) call intentional relations, in identifying the ILS claims of RST.

To illustrate how a speaker’s intentions determine discourse structure in this theory, consider the RST analysis of the example discourse from Figure 1. As shown in Figure 2, at the top level, the text is broken down into two spans: (a) and (b-c). The span (b-c) forms a satellite that stands in a **motivation** relation to (a). This span can

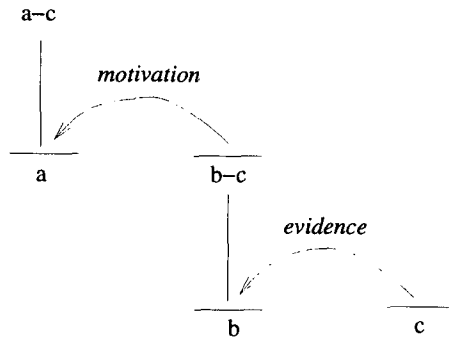


Figure 2
 The RST structure assigned to the example discourse in Figure 1.

be further broken down into the two minimal units (b) and (c), where (c) is a satellite that stands in an **evidence** relation to (b).²

While there is no direct representation of intentions in RST, the asymmetry between a nucleus and its satellite originates with the speaker’s intentions. The nucleus expresses a belief or action that the hearer is intended to adopt. The satellite provides information that is intended to increase the hearer’s belief in or desire to adopt the nucleus. Implicitly, this is a claim that the text is structured by the speaker’s intentions and, more specifically, by the difference between the intention that the hearer adopt a belief or desire expressed in a text span and the intention that a span contribute to this adoption. In the example, the nucleus (a) expresses an action that the speaker intends the hearer to adopt. The satellite (b-c) is intended to facilitate this adoption by providing the hearer with a motivation for doing the suggested action. In the embedded span, the nucleus (b) expresses a belief that the speaker intends the hearer to adopt and the satellite (c) is intended to facilitate this adoption by providing evidence for the belief.

The second implicit RST claim about ILS is a refinement of the first. The intentional relations specify the ways in which a speaker can affect the hearer’s adoption of a nucleus by including a satellite. That is, not only is there a functional distinction between nucleus and satellites, there is also a classification of satellites according to how they help achieve the hearer’s adoption of the nucleus. Translating this into a claim about ILS, text is structured by the ways in which some utterances are intended to help other utterances achieve their purpose.

4. Correspondence between Dominance and Nuclearity

Now we are in a position to compare the explicit claims of G&S about ILS with the implicit ones of RST. Both theories agree that a discourse is structured into a hierarchy of non-overlapping constituents, segments in G&S and spans in RST. Each subconstituent may in turn be structured in exactly the same way as the larger constituent. Superficially, the similarity ends there because the internal structure of segments and spans is different. In G&S, the internal structure of a segment consists of any number of embedded segments plus what we are calling the core, the (usually unembedded)

² As discussed in Mann and Thompson (1988), a motivation relation occurs when a speaker intends the satellite to increase the hearer’s desire to perform the action specified in the nucleus. An evidence relation occurs when a speaker intends the satellite to increase the hearer’s belief in the nucleus.

utterances that express the discourse segment purpose. In RST, the internal structure of a span consists of a nucleus, which we have characterized as expressing a belief or action the hearer is intended to adopt, a satellite, which is intended to facilitate that adoption, and an intentional relation between the nucleus and satellite.

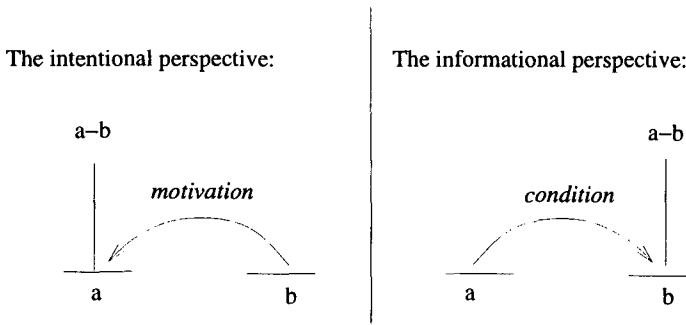
If we look more closely at the correspondence between dominance and nuclearity, we find that the structure of spans and segments is nearly identical. Specifically, an embedded segment corresponds to a satellite, and the core corresponds to the nucleus. Or, because G&S do not have the notion of core in their theory, a more accurate characterization of the correspondence would be that the nucleus manifests a dominating intention, while a satellite manifests a dominated intention. That is, dominance in G&S corresponds closely to nuclearity in RST. There is a relationship, which we can crudely characterize as that of linguistic manifestation, that links the nucleus to a dominating intention and a satellite to a dominated intention. Exactly how to derive a communicative intention from an utterance, and vice versa, is one of the main research issues in computational linguistics. Here we simply assume that an utterance conveys either a belief or an action p and thereby makes manifest the speaker's intention that the hearer adopt belief in or an intention to perform p .

The correspondence suggests a mapping between G&S linguistic structure and RST text structure. An embedded segment in G&S will be analyzed as a satellite in RST, and the segment core will be the nucleus. When there are multiple embedded segments in G&S, each subsegment will be analyzed as an RST satellite. In these cases of multiple subsegments, the RST structure will depend on whether the RST relations are the same or different. The entire segment may be a single RST span with the G&S core as nucleus and each subsegment as a satellite of that nucleus. This occurs when the multiple satellites bear the same RST relation to the nucleus. Alternatively, the G&S core and an adjacent subsegment may be analyzed as an RST nucleus and satellite, forming an RST span. This span is then the nucleus of a higher span in which the satellite is an additional G&S subsegment from the same segment. This occurs when the multiple satellites bear different relations to their nucleus.

Because cores are a central aspect of the mapping between the two theories, and because cores are not part of the G&S proposal, it is natural to ask whether a segment necessarily has a core. Given the nature of segment purposes, a coreless segment seems intuitively unlikely. Recall that segment purposes, like the utterance intentions discussed by Grice, have the property that they are intended to achieve their effect in part from being recognized. The core has an important function: it manifests the purpose of the segment. Without a core, the segment purpose must be inferred from the subsegments alone. In such a case, the speaker intends that the hearer recognize a purpose, but does not supply an utterance that manifests that purpose.

The question of whether or not coreless segments actually occur, however, is best answered by corpus analysis rather than theorizing. For our present purposes, we wish to consider the possibility of a coreless segment only because such a segment would complicate the mapping between the two theories presented above. In G&S, the definition of linguistic structure does not require a segment to contain a core. In the RST schemas considered thus far, a span always consists of a nucleus and satellite. A less common schema pattern, known as the **joint** schema, contains multiple spans with no nucleus-satellite distinction among them joined into a single span. Should a coreless segment occur in a G&S analysis, it can be mapped to a joint schema in RST.

Building on the correspondence between dominance and nuclearity, we raise two issues in the following sections. First, how do informational relations fit into the discourse structure? Second, what synthesis of the two theories emerges when we recognize the correspondence?



(a) Come home by 5.
 (b) Then we can go to the store before it closes.

(a) Come home by 5:00.
 (b) Then we can go to the store before it closes.

Figure 3
 RST intentional and informational relations may determine incompatible structures.

5. Informational Structure

Moore and Pollack (1992) argued that RST defines two types of relations: intentional relations, which arise from the ways in which consecutive discourse elements participate in the speaker’s plan to affect the hearer’s mental state, and informational relations, which obtain between the content conveyed in consecutive elements of a coherent discourse. This is consistent with Mann and Thompson’s (1988, 256) distinction between “presentational” (intentional) and “subject matter” (informational) relations. However, while Mann and Thompson maintain that for any two consecutive elements of a coherent discourse, one rhetorical relation will be primary (i.e., related by an informational *or* an intentional relation), Moore and Pollack showed that discourse interpretation and generation require that intentional and informational analyses exist simultaneously. Thus, in addition to the Intentional Linguistic Structure discussed so far, a discourse may simultaneously have an informational structure, imposed by domain relations among the objects, states, and events being discussed.

5.1 Can Intentional and Informational Structure Differ in RST?

In addition to their claim that intentional and informational analyses must co-exist, Moore and Pollack presented an example in which the intentional and informational relations can impose a different structure on the discourse. It is important to understand, however, that their example shows that the discourse structure determined by informational relations *as defined in RST* can be incompatible with the one determined by intentional relations. Here we argue that the problem is due to the inclusion of nuclearity in the definition of RST subject matter (informational) relations. As shown in Figure 3, the incompatibility arises because the nucleus and satellite of the intentional relation may be inverted in the RST informational relation.³

In Section 4, we argued that nuclearity in an RST analysis is an implicit claim about

³ The particular example given by Moore and Pollack had an additional relation embedded in (b), which had the same incompatibility between the intentional and informational analysis.

Context I

Our current president is far more liberal than the last one.

Context II

Legislation passed this week will not be enacted.

- (a) Clinton supports environmental protection.
 (b) He is sure to veto House Bill 1366.

Figure 4

Either relatum may be the nucleus when an instance of a domain relation is used.

speaker intentions, corresponding to the G&S relation of dominance among intentions. That is, nuclearity rightly belongs in the definitions of intentional relations. In contrast, informational relations, properly construed, should not distinguish between nucleus and satellite in their definitions. As an example, consider the pair of RST relations **volitional-cause** and **volitional-result**. The volitional-cause relation is defined as one in which the nucleus presents a volitional action and the satellite presents a situation that could have caused the agent to perform the action. The effect of this relation is that the reader “recognizes the situation presented in the satellite as a cause of the volitional action presented in the nucleus.” The volitional-result relation is nearly identical except that the cause of the action is the nucleus and the result is the satellite. Why does RST need two relations to capture this? The reason is that the same domain relation, call it **cause-effect**, links a cause and effect regardless of which is the nucleus. In Figure 4, note that, while (a) causes (b), either (a) or (b) can be the nucleus of the relation. For a particular instance of a cause-effect in the domain, it is equally plausible for a speaker to mention the effect to facilitate the hearer’s adoption of belief in the cause, as would be suggested by context I in Figure 4, or to mention the cause to facilitate the hearer’s adoption of belief in the effect, as suggested by context II.

Moreover, this is precisely what the intentional relations capture. By incorporating the nucleus-satellite distinction into the definitions of RST informational relations, these relations include an implicit analysis of intentional structure. As a consequence, strict application of the RST informational relations can result in a different structure than that imposed by the intentional relations, and this is the source of the problem noted by Moore and Pollack. Because nuclearity can only be determined by consideration of intentions, and intentional and informational analyses of a discourse must co-exist, we argue that the solution to the problem is to properly relegate information about nuclearity (intention dominance) to the intentional analysis, and remove it from definitions of informational relations. In this way, these two determinants of discourse structure cannot conflict. In addition, note that this is preferable to adding surplus informational relations to allow either relatum to be the nucleus (as was done in the volitional-cause and volitional-result case) because (1) this obscures the fact that relations such as volitional-cause and volitional-result appeal to the same underlying domain relation and (2) the proliferation of relations weakens the restrictive power of the framework.

5.2 Relationship between ILS and Informational Structure

Once we recognize that an informational analysis is needed simultaneously with ILS and that the informational analysis should be determined by domain relations without reference to how the relations are employed by the speaker, exactly how to determine informational structure becomes an underconstrained question. Should all domain

relations across utterances be analyzed in the informational structure? What patterns of informational relations are employed in realizing various kinds of intentions, and what analysis provides a reliable means for identifying such patterns? Final answers to these questions require further research. Because constraints may be needed in order to make progress on these issues, we point out two approaches to constraining the definition of informational structure. In Section 6.2 we suggest that RST informational relations provide a version of one of these approaches.

The most inclusive definition of informational structure would contain all the domain relations between the things being talked about. Included would be causal relations of various sorts, set relations, relations underlying bridging inferences (Clark and Haviland 1977), and the relation of identity between domain objects underlying coreference of noun phrases across utterances. By this definition, informational structure is a complex network of domain relations that is defined independently of the intentional structure. Keeping track of all domain relations in a discourse is an overwhelming task and is often infeasible. One approach to constraining informational structure is to define it as parasitic on intentional structure. The informational structure would contain an accompanying informational relation for each intentional relation. A second approach to constraining informational structure is to define it as a network of domain relations with type restrictions on the relata. The informational structure would contain only the relations among situations, events, and actions, that is, the types of entities referred to by clauses.

6. A Partial Synthesis

The discussion in Section 4 suggests that RST and G&S share a large amount of common ground. That is, many of the claims in the two theories, although formulated differently, are essentially equivalent. To begin this section, we state the common ground that emerges from relating dominance and nuclearity. Then we briefly review the claims of each theory that are outside this common ground. Each theory has some consistent ground, additional claims that concern issues simply not addressed by the other theory. The actual contentious ground, claims made by one theory that are incompatible with the other, is quite small.

6.1 Common Ground

Building on the correspondence between dominance and nuclearity, a partial synthesis of G&S and RST would be roughly the following: A segment/span arises because its speaker is attempting to achieve a communicative purpose. Such purposes have the feature that they are achieved in part by being recognized by hearers. Thus, the plan for achieving the purpose typically has two distinct parts: (1) one or more utterances that serve to make the purpose manifest by expressing a belief or action for the hearer to adopt (the core/nucleus) and (2) a set of subparts that contribute to achieving the purpose by manifesting subpurposes dominated by that purpose (the embedded segments/satellites).

Note that this synthesis encompasses the ILS claims of both theories regarding the example discourse in Figure 1. DS_0 is a segment/span designed to achieve the purpose I_0 . The plan for achieving I_0 is to first manifest I_0 by expressing the action in (a), the core/nucleus, and then to contribute to the achievement of I_0 by providing the motivation in (b-c), the embedded segment/satellite. In turn, DS_1 is a segment/span designed to achieve the purpose I_1 by first manifesting I_1 in the expression of the core/nucleus (b) and then providing evidence in the embedded segment/satellite (c).

Finally, I_2 is made manifest by (c), though no additional contribution to achieving this intention is provided.

6.2 Consistent Ground

RST and G&S each makes claims about issues not addressed by the other theory. We review these claims briefly in order to establish that they are consistent.

First, the two theories offer different but consistent perspectives on the ordering of segments/spans. In G&S, intentions may be related by satisfaction-precedence in addition to dominance. One intention satisfaction-precedes another when it must be realized before the other. This relation between intentions partially constrains the order of what is said and thus introduces a distinction between necessary order, originating with a satisfaction-precedence relation of the underlying intentions, and artifactual order, additional ordering that must be imposed to produce linearized text. G&S makes no claim about the relative ordering between a core and embedded segments. In RST, because the underlying intentions are not analyzed explicitly, the distinction between necessary and artifactual order is not available. Instead, the relative ordering of core/nucleus and embedded segment/satellite is highlighted. RST's authors claim that many relations have a typical ordering of their nucleus and satellite. The two theories address different aspects of ordering without suggesting any points of contention.

Second, in addition to intentional and linguistic structure, G&S posits an attentional structure. This component determines which discourse entities will be most salient and thereby imposes constraints on available referents for pronouns and reduced definite NPs. This is an important issue, but one that RST simply does not make claims about. As noted earlier, the recognition of intentional structure is crucial for anaphora resolution, among other discourse-processing tasks. By synthesizing RST and G&S, work done using both approaches can be applied to accomplishing these tasks during interpretation and generation.

Finally, while G&S recognize that informational structure is a cue to recognition of intentional structure, the theory does not provide detail. As discussed in Sections 5.1 and 5.2, the analysis of informational relations provided by RST is inadequate and incomplete. In either theory, more research is needed to understand how informational relations are used to achieve discourse intentions.

6.3 Contentious Ground

The claims of G&S and RST discussed so far have been, we argued, either equivalent or compatible. We now turn to a point of contention between the two theories. There are distinctions among the RST intentional relations that, in G&S, would be subtypes of the dominance relation among intentions. However, G&S specifies that the only relations among intentions affecting discourse structure are dominance and satisfaction-precedence. Should the various RST intentional relations be incorporated into a synthesized theory?

The question may be approached from either an empirical or a practical perspective, and the two perspectives may lead to different answers. To answer the question empirically, one could code a corpus for its intentional relations and attempt to identify linguistic cues that correlate with distinctions among the relations. To answer the question practically, one would consider whether distinct intentional relations are useful for computational systems that generate and/or interpret natural language. In fact, the

practical application of these intentional relations may be quite different in generation and interpretation systems. Further research is required to resolve this question.

7. Conclusion

In this paper, we have argued that two of the most important theories of discourse structure in computational linguistics, G&S and RST, are not incompatible, but in fact have considerable common ground. The key to the basic similarity between these two theories is understanding the correspondence between the notions of dominance in G&S and nuclearity in RST. Understanding this correspondence between the theories will enable computational models that effectively synthesize the contributions of the theories, and thereby are useful both for interpretation and generation of discourse.

Acknowledgments

The research described in this paper was supported by the Office of Naval Research, Cognitive and Neural Sciences Division (Grant Number: N00014-91-J-1694), and a grant from the DoD FY93 Augmentation of Awards for Science and Engineering Research Training (ASSERT) Program (Grant Number: N00014-93-I-0812). The authors wish to thank Robert Dale, Barbara Di Eugenio, Donia Scott, Lynn Walker and two anonymous reviews for helpful comments on an earlier draft of this squib.

References

- Asher, Nicholas and Alex Lascarides. 1994. Intentions and information in discourse. In *Proceedings of the 32nd Annual Meeting of the Association for Computational Linguistics*.
- Clark, H. H. and S. E. Haviland. 1977. Comprehension and the given-new contract. In R. O. Freedle, editor, *Discourse Production and Comprehension*. Lawrence Erlbaum Associates, Hillsdale, NJ.
- Grice, H. Paul. 1957. Meaning. *Philosophical Review* 68:377–388.
- Grosz, Barbara and Sarit Kraus. In press. Collaborative plans for complex group action. *Artificial Intelligence*.
- Grosz, Barbara J. and Candace L. Sidner. 1986. Attention, intention, and the structure of discourse. *Computational Linguistics* 12(3):175–204.
- Grosz, Barbara J. and Candace L. Sidner. 1990. Plans for discourse. In Philip R. Cohen, Jerry Morgan, and Martha E. Pollack, editor, *Intentions in Communication*. MIT Press, Cambridge, MA, pages 417–444.
- Hirschberg, Julia B., Diane J. Litman, Janet Pierrehumbert, and Gregory L. Ward. 1987. Intonation and the intentional structure of discourse. In *Proceedings of the Tenth International Joint Conference on Artificial Intelligence*, pages 636–639, Milan, Italy.
- Hobbs, Jerry R. 1985. On the coherence and structure of discourse. Technical Report CSLI-85-37, Center for the Study of Language and Information, Leland Stanford Junior University, Stanford, California.
- Hobbs, Jerry R. 1993. Intention, information, and structure in discourse: A first draft. In *Burning Issues in Discourse, NATO Advanced Research Workshop*, pages 41–66, Maratea, Italy.
- Mann, William C. and Sandra A. Thompson. 1988. Rhetorical Structure Theory: Towards a functional theory of text organization. *Text* 8(3):243–281.
- Moore, Johanna D. 1995. *Participating in Explanatory Dialogues: Interpreting and Responding to Questions in Context*. MIT Press, Cambridge, MA.
- Moore, Johanna D. In press. The role of plans in discourse generation. In Daniel Everett and Sarah G. Thomason, editors, *Discourse: Linguistic, Computational, and Philosophical Perspectives*.
- Moore, Johanna D. and Cécile L. Paris. 1993. Planning text for advisory dialogues: Capturing intentional and rhetorical information. *Computational Linguistics* 19(4): 651–695.
- Moore, Johanna D. and Martha E. Pollack. 1992. A problem for RST: The need for multi-level discourse analysis. *Computational Linguistics* 18(4):537–544.
- Moser, Megan and Johanna D. Moore. 1995. Investigating cue selection and placement in tutorial discourse. In *Proceedings of the 33rd Annual Meeting of the Association for Computational Linguistics*. 130–135.
- Moser, Megan and Johanna D. Moore. In preparation. On the correlation of cue placement with discourse structure: results from a corpus study.

