

Rhetorical Relations: Necessary But Not Sufficient

Sandra Carberry Jennifer Chu Nancy Green
carberry@cis.udel.edu jchu@cis.udel.edu green@cis.udel.edu

Department of Computer Science
University of Delaware
Newark, Delaware 19716

and

Lynn Lambert
lambert@pcs.cnu.edu
Department of Physics and Computer Science
Christopher Newport University
Newport News, Virginia 23606

It is generally agreed that coherent discourse consists of segments that are related to one another. A number of researchers have argued for the use of rhetorical[Gri75] or coherence relations [Hob79], and the rhetorical relations specified by RST [MT87] have been used in structuring text [Hov88, MP90]. In this paper, we examine rhetorical relations in the context of dialogue, rather than single-speaker text. We argue that reasoning about relational propositions is necessary but not sufficient for structuring dialogue, point out several problems of applying RST to dialogue, and argue for the necessity of recognizing the intentions underlying utterances and the rich relationships among these intentions.

Our research on recognizing expressions of doubt and interpreting indirect replies provides evidence that what Moore and Pollack call informational level relations [MP92] play an important role in identifying intentions in dialogue. Consider some continuations of the following dialogue sequence:

- (1) S1: "Who is teaching CS360 (a systems course)?"
- (2) S2: "Dr. Smith is teaching CS360."
- (3) S1: "Isn't Dr. Smith on sabbatical?"
- (4) S2: "No, he's not on sabbatical."

If S1 responds with "Wasn't Dr. Smith awarded a Fulbright?", then since it is plausible that Dr. Smith having a Fulbright would result in his going on sabbatical, S1's utterance should be interpreted as expressing doubt at the proposition that Dr. Smith is not on sabbatical. Similarly, if being a systems person is necessary for teaching CS360 and S1 responds with "Isn't Dr. Smith a theory person?", then since Dr. Smith being a theory person is an alternative to Dr. Smith being a systems person, S1's utterance should be interpreted as expressing doubt at the proposition that Dr. Smith is teaching CS360. But if S1 had instead asked "Isn't CS240 a prerequisite for CS360?", S1's utterance should be interpreted as seeking new information since there is no plausible informational relation suggesting that the new utterance contributes to determining whether Dr. Smith is on sabbatical or to identifying the instructor of CS360. Thus, identifying the informational relationship between utterances is essential for responding appropriately. [LC92] provides an algorithm that utilizes such relations in recognizing expressions of doubt.

Our research on interpreting indirect replies [GC92] provides additional evidence that informational relations play an important role in discourse interpretation. Consider the following dialogue sequence:

- (5) S1: *“Did Harry see his girlfriend last weekend?”*
(6) S2: *“He left for New York on Friday.”*

We have shown [GC92] that S1 and S2’s shared assumption that S2’s indirect answer in (5) is coherently related to some possible direct answer is necessary to enable S1 to identify S2’s intended (but implicit) answer of “*Yes*” or “*No*”. For example, suppose S1 believes that it is mutually believed that Harry and his girlfriend both live in Denver. Then, S1 would interpret (6) as describing an alternative to Harry seeing his girlfriend and thus would infer a negative answer to (5). On the other hand, suppose S1 believes that it is mutually believed that Harry lives in Denver but his girlfriend lives in New York. Then, S1 would interpret (6) as describing an action enabling Harry to see his girlfriend and thus would infer a positive answer. In other words, S1’s recognition of the relational proposition which S2 intended to convey is necessary for interpreting S2’s intended answer. If S1 failed to recognize such a relation, then S1 would be unable to answer a question such as *“Why didn’t Harry see his girlfriend last weekend?”* Thus in our model [GC92], the applicability conditions of discourse plan operators for answering a Yes/No question include conditions requiring that components of the response (whether implicit or explicit) be related by coherence relations similar to informational-level rhetorical relations. Since the same set of operators is used in generating answers, informational relations also play a role in ensuring that an indirect answer is appropriate. In other words, the informational relations constrain what *extra* information (information not specifically requested by the questioner) may be included in an appropriate response. An indirect answer is generated by use of the *extra* information alone.

In our view, this extra information may serve a rhetorical function such as to increase belief. For example, S2 might have decided to include the extra information in (6) (that Harry left for New York on Friday) in his negative answer because S2 anticipated that S1 would have doubted a simple “*No*”. In our current research, we are defining stimulus conditions for use of the extra information in an answer. A stimulus condition describes a situation in which it may be beneficial to include information that was neither requested nor what the speaker had a prior intention to convey. Thus, stimulus conditions are used to select an appropriate informational relation.

However, we contend that rhetorical relations, especially as defined by RST, are insufficient for characterizing dialogue. Although one might consider developing a larger set of rhetorical relations [Hob79], the following dialogue can not be completely characterized by rhetorical relations between segments.

- (7) S1: *“The AI class must be cancelled today.”*
(8) *“John is taking the course and he just went home.”*
(9) S2: *“John often cuts classes.”*

Although one might argue that there is a Motivate-Volitional-Action relation between John cutting class in (9) and John going home in (8), the function of utterance (9) is to attack the evidence relationship between utterances (8) and (7), which RST cannot account for since RST captures relations between spans of text whereas (9) is refuting the implicit proposition between (7)-(8), not the text itself.

In addition, rhetorical relations as defined in RST relate satellites to a nucleus, which together form an uninterrupted text span. However, in dialogue, an utterance may relate to an earlier

utterance U yet be separated from U by other utterances that do not relate to U. In such cases, expectations about speaker intentions play a major role in understanding, as illustrated by the following example.

- (10) S1: *"AI is the best course to take next semester."*
- (11) S2: *"You should take Computer Graphics instead."*
- (12) S1: *"Dr. Smith is teaching AI and he's wonderful."*
- (13) S2: *"Dr. Brown is teaching Graphics and he's a great teacher."*
- (14) S1: *"The AI projects sound like a lot of fun."*
- (15) S2: *"Computer Graphics has you drawing all kinds of objects."*

S1 and S2 each declare a claim in utterances (10) and (11). Although one can argue that an alternative relation holds between utterances (12) and (13) and between utterances (14) and (15), the structure of the dialogue is not completely captured by relating these consecutive pairs of utterances. Utterances (12) and (14) support S1's claim in (10), and utterances (13) and (15) support S2's claim in (11). Although one might suggest dropping the requirement that rhetorical relations relate a span of utterances, it is unclear how rhetorical relations alone could handle complex dialogues. We believe that in identifying the structure of dialogues such as the above, expectations about the discourse goals of the participants must be taken into account (in this case, a speaker *supporting* his/her own claim and *indirectly attacking* that of the other) and a more elaborate intentional structure allowed.

Furthermore, utterances can simultaneously serve more than one function, which would require two distinct RST analyses. Consider the following dialogue sequence:

- (16) S1: *"Can you come to a party at my house Saturday night?"*
- (17) S2: *"No, I can't."*
- (18) *"I have to work."*
- (19) *"So I am vry sorry, but I must decline your invitation."*

Utterance (16) conveys both a literal question and an invitation. Two RST analyses are required for S2's response, one relating (17) and (18) to the literal question and the other relating (19) to the invitation. However, RST requires a single analysis and a hierarchical structure. In addition, RST cannot account for the fact that if (19) is omitted from S2's response, it may still be implicated.

Recognizing intentions is essential for dialogue understanding, since these intentions provide expectations used in interpreting subsequent utterances and identifying the structure of the dialogue. For example, we show in [GC92] that expectations about discourse goals play a role in the interpretation of indirect replies to Yes/No questions. After S1's request for information in the exchange (5) - (6), S1 and S2 share the expectation that S2's response will convey the requested information. This expectation is used to focus on a certain set of discourse plan operators representing mutually accessible knowledge of standard forms for giving a positive or negative answer. Furthermore, our dialogue model [LC91] captures not only communicative intentions but also domain and problem-solving intentions, and these intentions result in a set of expectations that facilitate understanding subsequent utterances and generating appropriate responses. In [Chu93], we explore response generation in collaborative dialogue.

Although current research has been concerned with recognizing the intentions that a speaker is trying to convey, we believe that an effective and intelligent system must do more. If a system is to handle naturally occurring dialogue, which can range from completely cooperative to non-cooperative in a single interaction, then the system must be able to recognize hidden intentions,

such as the intention to lie or deceive. Although recognition of such hidden intentions may not be essential for identifying the structure of the discourse, it is necessary for the system's responses to be intelligent, natural, and effective.

So what makes a dialogue coherent? We believe that dialogue coherence depends upon both informational and intentional level properties of the dialogue. We contend that a natural language dialogue system must be able to recognize a speaker's intentions, that this recognition of intention is often aided by identification of informational level relational propositions, and that in many cases a speaker intends for these informational level relations to be recognized (as in the exchange given in (5) - (6)). Our research has led us to conclude that rhetorical relations as specified by RST are necessary but not sufficient for handling dialogues — the rich relationships among discourse intentions must also be captured.

References

- [Chu93] Jennifer Chu. Responding to user queries in a collaborative environment. In *Proceedings of the 31st Annual Meeting of the Association for Computational Linguistics*, Student Session, 1993. To appear.
- [GC92] Nancy Green and Sandra Carberry. Conversational Implicatures in Indirect Replies. In *Proceedings of the 30th Annual Meeting of the Association for Computational Linguistics*, pages 64–71, 1992.
- [Gri75] Joseph E. Grimes. *The Thread of Discourse*. Mouton, The Hague, Paris, 1975.
- [Hob79] Jerry R. Hobbs. Coherence and Coreference. *Cognitive Science*, 3:67–90, 1979.
- [Hov88] Eduard H. Hovy. Planning Coherent Multisentential Text. *Proceedings of the 26th Annual Meeting of the Association for Computational Linguistics*, pages 163–169, 1988.
- [LC91] Lynn Lambert and Sandra Carberry. A Tripartite Plan-based Model of Dialogue. In *Proceedings of the 29th Annual Meeting of the ACL*, pages 47–54, Berkeley, CA, June 1991.
- [LC92] Lynn Lambert and Sandra Carberry. Using Linguistic, World, and Contextual Knowledge in a Plan Recognition Model of Dialogue. In *Proceedings of the 14th International Conference on Computational Linguistics (COLING-92)*, pages 310–316, Nantes, France, July 1992.
- [MP90] Johanna Moore and Cecile Paris. Planning Text for Advisory Dialogues. In *Proceedings of the 27th Annual Meeting of the Association for Computational Linguistics*, pages 203–211, Vancouver, Canada, 1990.
- [MP92] Johanna Moore and Martha Pollack. A Problem for RST: The Need for Multi-Level Discourse Analysis. *Computational Linguistics*, 18(4):537–544, 1992.
- [MT87] William C. Mann and Sandra A. Thompson. Rhetorical Structure Theory: A Theory of Text Organization. Technical Report ISI/RS-87-190, ISI/USC, June 1987.