

Rhetoric and Intentions in Discourse

Robert Dale

rda@cogsci.edinburgh.ac.uk

This brief note offers some views on the following questions:

What is the precise relationship between rhetorical relations and intentions? Do they perform different (but possibility related) functions, or are rhetorical relations the realizations of intentions, or should rhetorical relations be discarded as a misconstrual of intentions proper?

To make this more concrete, we'll focus on a specific pair of theories. We'll take Grosz and Sidner's Theory of Discourse Structures (henceforth GSDT) to be a good example of a theory whose notion of discourse structure is based on intentions; and we'll take Rhetorical Structure Theory (RST) to be a good example of a theory whose notion of discourse structure is based on rhetorical relations.

On the face of it these theories are rather similar: both provide hierarchical characterisations of texts, and both seem concerned with the purpose of the text – for GSDT, the text realizes the intentional structure that underlies it; and for RST, the text is structured in terms of the rhetorical relations that hold between its parts. But these similarities are only skin deep. The kinds of trees used by the two theories are actually quite different: in GSDT, text can appear at both leaf nodes and internal nodes in the trees; whereas in RST, text can only appear at leaf nodes. So the structures might be isomorphic under some mapping but that mapping remains to be worked out (at the very least, isomorphism could only be achieved by annotating arcs in an intentional structure tree with “rhetorical” labels; there is an interesting parallel here with the difference between dependency grammar and phrase structure grammar). A more significant difference is that, in RST, leaf nodes (the atomic text spans) are typically taken to be single clauses; but in GSDT, if our atomic elements (the bottom level discourse segments) were clauses, then, given the constraints the theory claims on pronominalization and reference, we'd rule out as anomolous a large space of perfectly legitimate pronominalizations (those where the antecedent and anaphor are in clauses at the same level of embedding in the discourse structure: recall that material in a sister discourse segment is inaccessible). Given this, it would be hard to argue that RST's text spans correspond to GSDT's discourse segments.

This point is more important than it might at first seem, since it draws attention to a relatively unnoticed difference between the two approaches: in RST, rhetorical relations are deemed to hold at all levels of discourse, including between individual clauses; whereas in GSDT, intentional relations can hold between individual clauses (precisely in those cases where individual clauses correspond to individual discourse segments), but they need not do so – it's perfectly within the bounds of the

theory to have two clauses between which there is neither a DOM nor an SP relationship. This opens the possibility that there need not be an intention underlying each individual clause or sentence.

RST's desire to postulate rhetorical relations all the way down to the level of clauses can be seen as the root of a problem that gives rise to another more frequently voiced concern about the theory: that many of the so-called rhetorical relations are subject-matter relations, and very few are presentational. If, at the end of the day, we feel the need to postulate some relation between two sentences where the only obvious connection is that the event described by one follows the event described by the other, then it's hardly surprising that we feel the need to include a notion of SEQUENCE in our set of relations. But why stop there? We can often include information which might have been expressed in separate sentences using a noun phrase modifier instead: do we then want to say that there is a rhetorical relation between a head noun and each of its associated adjectives? I would suggest not; and that the first step we should take towards clearing up the confusion here is to banish those rhetorical relations which simply "mirror" underlying knowledge base relations. Doing this will also remove the problem that sometimes appears in the analysis of rhetorical structures, where the analyst is unsure which of several different but equally plausible relations should be postulated between two elements: once we accept that all we are often really doing is identifying knowledge base relations, then it becomes unsurprising that two elements might equally well be related by means of a CAUSE relation or an EVIDENCE relation, since there can be any number of relations between two entities in the knowledge base (precisely because we assume there is no need for structures to be linear or tree-shaped: typically they are thought of as graphs).

Taking this step will have two consequences. First, it will reduce substantially the number of textual relations we are willing to consider as rhetorical relations; and second, it will allow us to contemplate rhetorical structures which do not need to provide an analysis all the way down to individual clauses. This brings us closer to the thinking that lies behind GSDT, and answers some of its objections against RST; but we do not need to adopt as minimalist approach as GSDT when it comes to enumerating the possible relationships in a text. Indeed, there is a sense in which GSDT falls foul of some of the same problems that affect RST: in particular, tying the notion of intentional structure to the structure of the task that underlies a text leads us to postulate communicative intentions that mirror the goals and subgoals in the task itself, but this is little better than having rhetorical relations that mirror the underlying domain relations. A better way forward is to begin with those RST relations which are clearly *not* domain relations, such as MOTIVATION and JUSTIFICATION, and to elaborate their definitions in order to build towards a proper theory of communicative intention and purpose.