

Using Cue Phrases to Determine a Set of Rhetorical Relations

Alistair Knott

Department of Artificial Intelligence, University of Edinburgh

80, South Bridge, Edinburgh EH1 1HN, Scotland

email: A.Knott@aisb.ed.ac.uk

‘Relation based’ approaches to discourse analysis and text generation suffer from a common problem: there is considerable disagreement between researchers over the set of relations which is proposed. Few researchers use identical sets of relations, and many use relations not found in any other sets. This proliferation of relations has been pointed out before (eg Hovy [1]), and several methods for justifying a standard set of relations have been proposed: this paper reviews some of these, and presents a new method of justification which overcomes some awkward problems.

1 Current Approaches to Justifying a Set of Relations

Descriptive Adequacy

Clearly, a set of relations must have ‘good coverage’—it must be possible to analyse all the texts of the kind targeted using the specified relations. At the same time, this cannot be the *only* requirement on a set of relations, because many different sets of relations can be used to describe the same set of texts. For instance, the level of detail of the description is not constrained: how do we decide whether or not to subdivide RESULT into VOLITIONAL RESULT and NON-VOLITIONAL RESULT? Again, different cuts through the space of relations are possible: why distinguish between VOLITIONAL and NON-VOLITIONAL result, and not between, say, IMMEDIATE and DELAYED result? In fact the notion of ‘descriptive adequacy’ seems to make little sense in isolation: in addition, it is necessary to specify a *purpose* for which the proposed description is adequate.

Psychological Reality

One way that a purpose can be specified is by stipulating that relations model ‘cognitive’ constructs—that is, constructs which people *actually use* when they create and interpret text. On this conception, a description of the relations in a text becomes part of a *theory* of how the text originated, and why it is the way it is. This stipulation gives relations real explanatory power in a theory of discourse coherence: we could argue that it is *because* a

given set of relations is involved in text processing that we are able to use them to describe text.

Claiming psychological reality for relations makes them amenable to empirical investigation. Sanders, Spooren and Noordman [3] make the claim explicitly, and seek evidence for their proposed set of relations in psychological experiments on readers and writers. But these experiments are not without their own problems—it is questionable whether empirical experiment is a sharp enough tool to reveal fine-grained distinctions between relations. (Automatic text generation is one area where such fine-grained distinctions might well be necessary.)

Cue Phrases

Cue phrases (sentence connectives such as *because* and *nevertheless*) have been another influence on the choice of a set of relations. Even in RST, where relations are defined without reference to surface linguistic phenomena, many correlations exist between relations and particular cue phrases. Hovy [1] makes more explicit use of cue phrases, taking them as ‘nonconclusive’ evidence for a taxonomy of relations.

However, while it is clear that a fine-grained classification of relations could indeed be constructed using cue phrases, the question of *why* relations should be linked to cue phrases has not been addressed in detail. Hovy’s rationale concerns the practicalities of designing text planning systems; such a pragmatic approach has its advantages, but it would nonetheless be useful to think of a *theoretical* reason for linking relations to cue phrases. Without one, relations lack the kind of explanatory power they receive when thought of as psychological constructs.

2 Cue Phrases as Evidence for Cognitive Constructs

I have suggested (Knott and Dale [2]) that it is possible to think of cue phrases as evidence for relations *precisely if* they are conceived as psychological constructs. The argument is basically that we can *expect* language to contain ways of making explicit any relations which are actually used by people when they process texts. If identifying relations is a component of text understanding, then it makes sense for there to be ways of signalling relations directly in text: it facilitates the tasks of both the reader and the writer.

Of course it is often possible for a reader to identify a relation without the need for textual signals. But it seems unlikely that any relation exists that *never* needs to be textually marked. Unmarked relations can only be recognised if the reader has a certain amount of background knowledge about a text; and this knowledge can hardly be guaranteed in advance for all texts in which the relation can appear. If relations do play a part in human text comprehension, it is reasonable to suggest that there exists a particular linguistic formula or expression which can be used to distinguish each relation from all others.

This claim, if accepted, would allow a taxonomy of relations to be built which has both

fine-grained detail and explanatory power. The explanatory power comes from the conception of relations as psychological entities: this conception in turn legitimises the use of cue phrases as evidence for relations, allowing a detailed taxonomy to be constructed.

3 A Methodology for Deciding on a Set of Relations

On the basis of the above argument, a new methodology for determining a set of relations can be proposed: in essence, a relation is included in the set if a cue phrase can be found which picks it out. The starting point for this methodology is the assembly of a corpus of cue phrases.

A study beginning from this point is reported in Knott and Dale [2]. We began by devising an informal test for identifying cue phrases in naturally-occurring discourse. Using the test, 120 pages of academic articles and books were analysed, yielding a corpus of over 100 cue phrases. Following an idea by Hovy [1], this corpus was arranged into a taxonomy of synonyms and hyponyms: a portion of this taxonomy is shown in Figure 1.

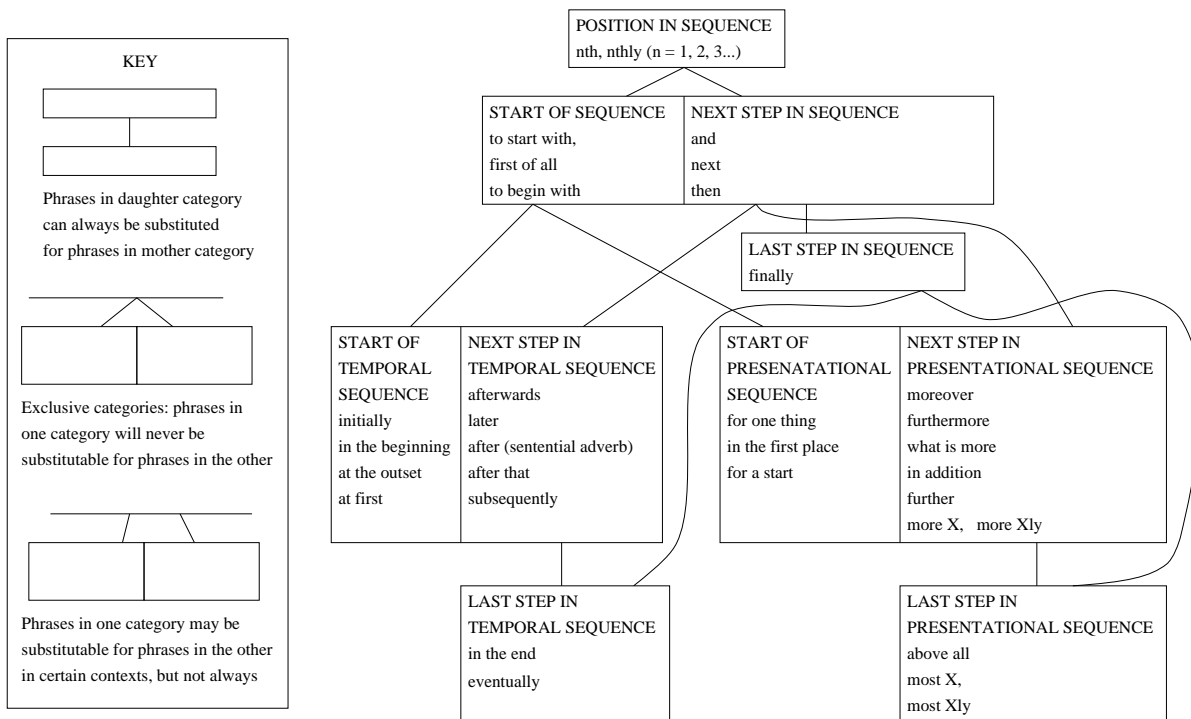


Figure 1: A Portion of the Taxonomy of Cue Phrases

To illustrate the working of this taxonomy, consider the following patterns of substitutability. *Initially* and *in the first place* can both be substituted by *first of all*; but they cannot be substituted for each other: *in the first place* is specific to ‘presentational’ sequences as *initially* is to ‘temporal’ sequences.

The labels in the boxes of the taxonomy are still just informal specifications: we are currently using the taxonomy to work out an isomorphic classification of relations, complete with more formal definitions. It is productive to think of relations as feature-based constructs, and of the taxonomy as an **inheritance hierarchy** for features, such that daughter nodes inherit the features of their mothers, and are in addition specified for extra features. The substitutability data could then be seen as informing a decision about which features to use in relation definitions.

The new taxonomy of relations is likely to have much in common with other taxonomies in the literature—although there may be some revealing additions and omissions. Indeed, the taxonomy of cue phrases can be seen as an interesting testbed for other theories of relations: can they explain the patterns of cue phrases which it describes?

4 Conclusion

This paper has presented a new way of justifying a set of relations, by viewing them as modelling psychological constructs and using cue phrases as evidence for these constructs. A corpus of cue phrases has been gathered and worked into a taxonomy, from which an isomorphic taxonomy of relations can be derived.

The benefits of this methodology are considerable. To begin with, it sets out a *systematic* way to decide on a set of relations: any disagreement can be traced to a particular stage in the process, such as the decision that a word is a cue phrase, or the decision that two phrases are synonymous. Furthermore, the assumption of psychological reality gives the relations in the taxonomy a clear explanatory role in a theory of discourse coherence: they are not just ‘purely descriptive’ constructs. Finally, the methodology being proposed is *incremental*. Many important decisions about relation definitions—whether parameters should be used, whether intentions need to be specified separately—can be deferred until after the taxonomy of cue phrases has been constructed: at this point, the taxonomy serves as a useful source of evidence for such decisions.

References

- [1] E H Hovy. Parsimonious and profligate approaches to the question of discourse structure relations. In *Proceedings of the 5th International Workshop on Natural Language Generation*, Pittsburgh, 1990.
- [2] A Knott and R Dale. Using linguistic phenomena to motivate a set of coherence relations. To appear in *Discourse Processes*. Also available as Technical Report RP-39, Human Communication Research Centre, University of Edinburgh, 1992.
- [3] T J M Sanders, W P M Spooren, and L G M Noordman. Towards a taxonomy of coherence relations. *Discourse Processes*, 15:1–35, 1992.