

Literature and Cognition

Jerry R. Hobbs
(SRI International)

Stanford: Center for the Study of
Language and Information (CSLI
Lecture Notes 21), 1990, vii + 180 pp.
Hardbound, ISBN 0-937073-53-9, \$35.00;
Paperbound, ISBN 0-937073-52-0, \$14.95
(distributed by the University of
Chicago Press)

Reviewed by
Janyce Wiebe
University of Toronto

Literature and Cognition applies methods and ideas from computational linguistics and cognitive science to literature. It explores a wide range of topics, and is intended to be read by literary theorists as well as by computational linguists and cognitive scientists. From the perspective of computational linguistics, the largest contributions are to present Hobbs's previously published theory of discourse coherence in the context of an overall framework for discourse interpretation, and then use it to analyze a sonnet in great detail. (Hobbs also analyzes a novella in detail, but much less directly in terms of the theory.) These will be the focus of this review.

The chapter presenting the overall framework (Chapter 3) is entitled "A theory of discourse interpretation," but it outlines what would constitute such a theory, rather than presenting an actual theory itself. A major theme of the presentation is the author's view that a theory of discourse interpretation "must first and foremost be a theory of how knowledge is used in solving the interpretation problems posed by the discourse" (p. 41). Such a theory would include six subtheories. The first four are familiar: the knowledge-representation language; the encoding of background knowledge in this language; the "deductive mechanism"; and syntactic and semantic "translation," which produce logical forms. Notice that the word "discourse" is used with a very broad meaning in the chapter title. This is somewhat confusing, because it is used with a more narrow meaning in the discussions of the fifth and sixth subtheories.

The fifth subtheory would specify the possible interpretations of a sentence, where an interpretation must solve each of the discourse problems of that sentence. Hobbs says:

What counts as a solution [to a discourse problem] can be specified in terms of inferences that can be drawn by the deductive mechanism from the propositional content of the sentence and the knowledge base. A possible interpretation of a sentence is taken to be a consistent combination of individual solutions to all of the sentence's discourse problems. (p. 47)

It appears from the discussion that discourse problems are the interpretation problems that remain after the logical form of a sentence has been derived from syntactic and semantic processing. A list of six categories of discourse problems is given, which is intended to be exhaustive. The descriptions of the categories are brief, and in some

cases not very clear, but they do give the reader an idea of the sorts of problems Hobbs considers to be discourse problems. One is “local coherence,” determining coherence relations between a sentence and the surrounding discourse, and another is “global coherence,” determining relations between the content of a sentence and “some internal representation of the environment.” He mentions plans as being important in this context. It would have been interesting to people working on discourse to see the distinction between local and global coherence spelled out more fully.

The final subtheory would involve choosing the best interpretation out of all those possible.

Hobbs proposes that the inference process should be under the strict control of the processes of finding possible interpretations and choosing the best one, in such a way that the system is able to determine which propositions are relevant to the current discourse context, out of all those entailed by the content of the current sentence and the knowledge base. The relevant inferences, he proposes, are the propositions supporting the interpretation ultimately chosen as the best one. Many questions come to mind, about how solutions to the various discourse problems are to be combined, and about the criteria to be used for choosing the best interpretation. Addressing these questions is outside the scope of the book. As it stands, there is so little detail given that the above proposals are difficult to judge. The real benefits of this chapter are that it lays out many of the author’s assumptions and presents a framework for discourse understanding used in later chapters to place the specific issues addressed into the context of an overall AI system.

Hobbs presents his theory of discourse coherence—what is called “local coherence” above—in Chapter 5. He gives brief definitions of *coherence relations*, such as *contrast*, *elaboration*, and *parallel*, that can hold between segments of a discourse. These relations are defined in terms of propositions that can be inferred from the assertions of discourse segments. The assertion of a clausal discourse segment is, roughly, what is predicated by the main verb (assertions of larger segments are discussed below). For example, the following is the definition of the *parallel* relation (where S_0 and S_1 are adjacent segments of the discourse):

Parallel:

Infer $p(a_1, a_2, \dots)$ from the assertion of S_0 and $p(b_1, b_2, \dots)$ from the assertion of S_1 , where a_i and b_i are similar, for all i . (p. 93)

That is, the *parallel* relation can hold between two segments if propositions predicating the same thing of similar arguments can be inferred from their assertions. The following sentence is an example:

Set stack A empty and set link variable P to T. (p. 93)

In terms of the above schema, the predicate p is *set*, a_1 is *stack A*, a_2 is the value *empty*, b_1 is *link variable P*, and b_2 is the value *T*.

A more complex example (this is a U.S. Congressman’s complaint about communication with the Nixon White House staff, from a *Newsweek* article):

We have nothing to say to Ron Ziegler,
and Al Haig’s never been in politics.
The *parallel* relation here depends on the inference from each clause that Ron Ziegler and Al Haig (similar entities, in that both were advisors to Nixon) are people with whom members of Congress cannot communicate. (p. 95)

Many researchers have found relations such as *contrast* and *parallel* to be valuable in discourse processing. With his relation definitions, Hobbs addresses the question of how domain knowledge is involved in recognizing such relations. In addition, fulfilling the requirements of a relation yields inferences to be drawn in the discourse context, as well as recognition of the relation itself. However, his focus on inference is at the expense of other aspects of discourse processing addressed by other researchers, for example using "surface" information such as cue words to constrain possible interpretations (as in Reichman 1985 and Cohen 1987, for example).

The recognition of coherence relations is just one of the discourse problems included in Hobbs's subtheory 5. Solving the other discourse problems, producing the candidate interpretations, choosing the best interpretation, and the rest of the subtheories are black boxes that are assumed to exist but are not specified here, being outside the scope of the book. The justification for assuming these black boxes is that this is "a way of isolating the problem of interest" (p. 43). People working in NLU, especially those working on discourse, must often make such assumptions. Otherwise, work in this area could not proceed until all problems in syntactic processing, semantic processing, and knowledge representation are solved. Besides the fact that discourse processing is an important problem in itself, work in discourse processing can help to direct work in these other areas, for instance by identifying problems that commonly arise.

So, we can accept in principle Hobbs partitioning off some difficult problems. However, the definitions of the coherence relations are brief, perhaps leaving too much up to the black boxes to truly convince one that the approach is feasible. For example, none of the definitions mentions the knowledge base, and some basic questions are not addressed. What if the assertion of one of the segments is inconsistent with the knowledge base, in which case anything could be inferred from the knowledge base and the assertion? Or, what if propositions appropriate for establishing a relation are entailed by the knowledge base alone and so could be inferred from the assertions of any pair of segments?

Hobbs applies his theory to a number of nontrivial examples, describing ways in which the criteria for particular relations to hold between clauses can be satisfied. However, many of his descriptions are sketchy. For some of the examples, I found it difficult to come up with specific propositions corresponding to the author's brief description, or to identify precisely which discourse units were being related. (Two examples I had particular trouble with are his (8) and (19).) Since this is where the coherence relation theory is applied to "ordinary discourse," laying a foundation for application of the theory to literary works, more detail in discussing the examples in this chapter should have been given.

Hobbs also addresses high-level discourse structure. The coherence relations are intended to apply not only to clauses, but also recursively to larger segments of the discourse. "A clause is a segment of discourse, and when two segments of discourse are discovered to be linked by some coherence relation, the two together thereby constitute a single segment of discourse" (p. 102). A *coherence structure* arises from the recognition of coherence relations between segments, which is a tree in the case of well-organized texts.

Since the relations are specified in terms of assertions of discourse segments, the question arises as to what the assertion is of a segment that is the combination of two segments related by a coherence relation. Hobbs gives a scheme for assigning assertions to such segments, but he says there are exceptions to his scheme, and one of the relations, the *occasion* relation, is not covered by it at all. Since assertions of discourse segments are key components of the relation definitions, how they are determined is an important area for further work.

Hobbs also gives a noncomputational method for analyzing discourse that is intended to be used by researchers working in such fields as ethnography and literary analysis. This method is applied to two literary works: a sonnet, "Lawrence of virtuous father virtuous son," by John Milton (in Chapter 6) and a novella, *Sylvie*, by Gerard de Nerval (in Chapter 7).

One goal motivating the analyses of the literary works is to apply the coherence theory to difficult texts. What "apply" means in this context is for the researcher to give coherence structures for the texts; Hobbs is clear about not trying to specify the processes by which an AI system could produce them. The sonnet is difficult to understand, and is well chosen for this enterprise, because many of the ambiguities that I, at least, had trouble resolving are ambiguities in the discourse structure.

Hobbs gives a detailed discourse analysis of the sonnet that is reasonable and that helped me make basic sense of it. The analysis is largely, but not entirely, derived from applying the coherence theory to the sonnet. The analysis never directly contradicts the coherence theory, but for many of the combined segments of the sonnet, the theory is silent as to what their assertions should be. For example, seven of the fifteen combined segments involve the *occasion* relation (or the *cause* relation, a special case of it), for which rules for determining the assertion are not given in Chapter 5. As another example, some of the combined segments are labeled with more than one relation. While Hobbs states in Chapter 5 that this situation is allowed under his theory, the question of what the assertion of such a segment should be is not addressed. Thus, many of the assertions given do not come from the theory, but are drawn out of the air.

Further, the components of the coherence structure are identified in good, but not complete, detail. A tree is given that spans the sonnet, in which each nonterminal node is labeled with a coherence relation name, and each terminal node is labeled with an unambiguous indication of the clause or clauses of the poem to which it corresponds. For 13 of the 15 combined segments (nonterminal nodes), the propositions supporting their coherence relations and their assertions are discussed, although for two of these segments (nodes 5 and 14), it isn't clear what the propositions supporting the coherence relations really are. Assertions for two of the segments (nodes 9 and 12) are never given. The chapter is a bit difficult for the casual reader to understand, because one must sometimes hunt around for pieces of information. For example, the assertions of four of the segments (nodes 4, 10, 11, and 13) are not given when the segments themselves are discussed, but rather when their parent segments are discussed later.

The analysis reveals some potential difficulties to be faced in further developing the scheme for assigning assertions to segments. For example, as the coherence relations are defined, the only assertions considered are those of the segments being related. Any information from the descendents of those segments that is not incorporated into their own assertions is effectively eliminated from consideration. Yet there is at least one combined segment in the tree (node 3) for which the assertion Hobbs gives seems reasonable only in light of the assertions of the child segments *plus* the assertion of a grandchild segment. One benefit of applying a computational theory to a difficult text could be to identify problems to be investigated further. Unfortunately, Hobbs does not include such a discussion in either of the chapters analyzing literary works.

The second goal of applying the coherence theory to literary works is to show that it can be a useful tool in literary analysis. For the sonnet, the idea is that there is a central meaning that is not explicitly stated, and the coherence theory can be used to help explicate how the reader arrives at it. Specifically, propositions inferred to establish some particular coherence relations are implicatures that are central to

the meaning of the poem. One might argue that it is other implicatures that are most central, or that different implicatures are being drawn, supporting a different literary interpretation; that a theory such as Hobbs's could provide a medium in which to argue such things is the sense in which it could be a tool used by literary analysts. Hobbs digresses in a few places to show how aspects of the poem other than its discourse structure—lexical choice, for example—support the literary interpretation he suggests. This strengthens his analysis and illustrates the integration of discourse and other types of analyses in literary interpretation.

Chapter 7, which analyzes the novella, is co-authored with Patrizia Violi. A number of trees representing pieces of the structure of the novella are given, but many labels on the nonterminal nodes are not from the set of coherence relations given in Chapter 5. The analysis focuses on relationships between the structure of the text and the literary themes of the work, but it is difficult to view it as an application of the coherence theory to the work.

In the other parts of the book, Hobbs speculates on the possible functions of imagining, fiction, and narrative for cognitive agents (in Chapter 2); discusses the relationship between literary and nonliterary discourse (in an afterword); and presents a cognitive science model of the structure of intelligent agents, using it to take positions on issues currently being debated in literary theory about the nature of interpretation (in Chapter 1, reprinted from elsewhere). Hobbs describes his approach to metaphor interpretation (in Chapter 4, a revision of a previously published paper). The basic idea is that metaphor interpretation happens as a matter of course as one seeks solutions to discourse problems—"as a by-product of other interpretation processes" (p. 58). The examples are from "ordinary" discourses, not from literary works; presumably this material is included in the book because metaphor is of great importance in the study of literature.

Readers might not be satisfied with the level of detail and precision in a number of places in *Language and Cognition*. Hobbs, in discussing the fact that the book is written for more than one audience, explicitly says that cognitive scientists might find the book too informal. Because the book is not written solely for computational linguists, it probably is not the appropriate place to further explore computational issues. However, even without addressing such issues, there is room for more precision and detail, which I think would have made the material easier to understand and more convincing for both audiences.

The book is valuable to computational linguistics for a number of reasons. It presents in an integrated framework many of Hobbs's ideas on discourse interpretation that have appeared elsewhere. It gives detailed discourse analyses of two literary works, one of which, that of the sonnet, is largely derived from Hobbs's coherence relation theory. The analysis highlights the central role of implicature in poetry, and includes difficult implicatures sanctioned by the coherence relation definitions. This illustrates the strength of Hobbs's approach to discourse analysis, under which one recognizes discourse structure by considering propositions entailed by segments of the discourse and the knowledge base. The analysis also shows how his theory could form the basis of a tool useful in literary analysis, which will be considered a contribution by computational linguists who are interested in how their theories might be useful to practitioners in other fields. In addition to these particular contributions, the book asks us to step back from the details of projects that are feasible today and consider the computational investigation of the most challenging kinds of texts—sophisticated literary works. I recommend this book not only for these reasons, but also simply because it is very interesting to read.

References

Reichman, Rachel (1985). *Getting Computers to Talk Like You and Me*. Cambridge, MA: The MIT Press.

Cohen, Robin (1987). "Analyzing the structure of argumentative discourse." *Computational Linguistics*, 13(1), 11-24.

Janyce Wiebe is an assistant professor at New Mexico State University. She received her Ph.D. in Computer Science from SUNY at Buffalo and recently completed a postdoctoral fellowship at the University of Toronto. Her current interests are point of view and discourse processing in narrative. Wiebe's address is: Department of Computer Science, Box 30001/Dept. 3CU, New Mexico State University, Las Cruces, NM 88003 USA. e-mail: wiebe@nmsu.edu