A Note on Coherence Relations

Jerry R. Hobbs
Artificial Intelligence Center
SRI International
Menlo Park, California

and

Gisela Redeker
Faculteit der Letteren
Vrije Universiteit Amsterdam
Amsterdam, The Netherlands

1 Background

Hobbs et al. (1993; henceforth, IA) elaborated a theory of discourse in which discourse is interpreted by the pragmatic strengthening of its explicit content as provided by weighted abduction. In this approach, the interpretation of a sentence is the minimal abductive proof of its logical form. In Section 6.3 of IA and in Hobbs (1996a) an account was given of the structure of discourse in which relations between successive segments of discourse are seen as pragmatic strengthenings of the relation conveyed by the adjacency of the segments. The most common of these pragmatic strengthenings have been called "coherence relations".

The coherence relations are defined in terms of "what is said" by the segments, the assertions or summaries of the segments, or the principal information conveyed by the segments. For example, a definition of the Explanation coherence relation similar to that given in IA is[1]

[1] In all the logical formulas in this note, axioms of the form

$$(\forall x, y)[p(x, y) \supset [(\exists z)q(x, z)]]$$

are abbreviated

$$p(x, y) \supset q(x, z)$$

(1) $cause(e_2, e_1) \supset CoRel(e_1, e_2, e_1)$

If a segment which asserts the eventuality e_1 is followed by a segment which asserts the eventuality e_2, where e_2 causes e_1, then there is a coherence relation between the two segments. The predicate $CoRel$ expresses the vague relation conveyed by the adjacency of the segments, and the predicate $cause$ is one of its possible pragmatic strengthenings.

For sentences, the assertion is generally the proposition conveyed by the main verb or modal or a top-level adverbial, although this can be altered by high stress or, as we will see below, by other factors.

When two segments are related by a coherence relation, the concatenation of the two itself constitutes a segment, which can in turn enter into coherence relations with other segments. By this recursion, a tree-like structure can be built up for an entire discourse. But in order for this to go through, we must specify the "assertion" of segments of text longer than a single sentence. This is the role of the third argument of $CoRel$. In (1) the assertion of the composed segment is the eventuality e_1. The fact that the assertion of the composed segment is identical to the assertion of one of the component segments is the formal way of specifying in this framework that the relation is "hypotactic". In terms of Rhetorical Structure Theory (Mann and Thompson, 1988), the first segment is the "nucleus". For paratactic relations, like Elaboration, the assertion of the composite arises equally from the assertions of the components.

(2) $imply(e_1, e) \wedge imply(e_2, e) \supset CoRel(e_1, e_2, e)$

If the same eventuality e is derivable from the assertions of the two component segments, there is a coherence relation between the segments and the assertion of the composite is e.

Redeker (in press) raises the question of why coherence relations are needed at all in discourse processing, since the inferences they induce can and must generally be derived by other means. She also puts forward a counterexample to the specific coherence relation rule (1). Suppose Bill calls John and says,

(3) Ann is so happy! She just got a promotion!

In addition, implication outscopes conjunction.

This matches the Explanation pattern (1). The promotion is the cause of Ann's being so happy. But the principal message of the composite segment is the promotion, not the happiness. The assertion of the combined segment, then, is not the explanandum as Axiom (1) specifies, but the explanans.

In this note, we address these two difficulties in the context of this specific example, formalizing to the extent required to elucidate the issues involved. Axiomatizations are simplified where a more elaborate account would only obfuscate. Defeasibility has been simplified out, but could easily be added using the technique of IA, Section 4.2.

Every aspect of discourse structure we discuss is shown to have a close analogue in local pragmatics, in particular, in the problems of coreference resolution, the pragmatic strengthening of vague predicates, and the expansion of metonymies. In this note, we in turn describe the IA solution to a local pragmatics problem and then illustrate an analogous phenomenon in the interpretation of discourse.

Section 2 addresses the question of why coherence relations are needed. We illustrate how the requirements of coreference often force the same inferences that are forced by the requirements of coherence, but argue that the establishment of coherence is necessary nevertheless. Section 3 explores the issue of what is asserted by the composite segment in an explanation relation. Specifically, we show how consideration of the speaker's goals can sometimes alter the main point of a segment of discourse, by a process that is at least analogous to metonymy resolution. In Section 4, we examine some consequences of the approach to coherence sketched here.

2 Why Recognize Coherence?

2.1 Coreference Resolution

As Redeker pointed out, the information obtained from a coherence analysis of texts is generally available by other means. One of the most common means is via coreference resolution. Let us first look at a coreference problem in local pragmatics, the resolution of a definite noun phrase (from IA, Section 5.1).

(4) John walked into the room. The chandelier shone brightly.

What chandelier? Suppose we have axioms in our knowledge base specifying that rooms have lights in them:

(5) $room(x) \supset light(y) \wedge in(y,x)$

and that a light with a branching fixture is a chandelier:

(6) $light(y) \wedge branching\text{-}fixture(y) \supset chandelier(y)$

We know from the first sentence that there is a room. Axiom (5) allows us to infer a light in the room, and by assuming that that light has a branching fixture, we have the chandelier that we needed to find in the second sentence. What we had to assume, that the light in the room had a branching fixture, is the new information conveyed by the definite noun phrase "the chandelier". Figure 1 illustrates this interpretation. Boxes are drawn around propositions that are assumed.

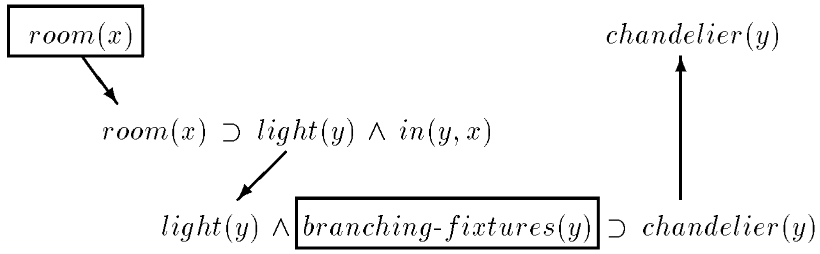

Figure 1: Coreference Resolution in Example (4).

Now let us return to example (3). The relevant explicit content of the text is

(7) $happy'(e_1, A), promotion(e_2, A)$

That is, there is an eventuality e_1 that is Ann's being happy, and there is an eventuality e_2 that is Ann's promotion.

Coreference resolution is not a special process in the IA approach, but rather just an instance of pragmatic strengthening, of finding the best abductive proof of the explicit content of the sentence. It thus works as well for information conveyed by verbs and indefinite noun phrases as it does for definite noun phrases. In this example, part of the content of "promotion" is already implicit in the happiness, just as part of the content of "the chandelier" was already implicit in the room.

The axioms we use to illustrate this example are not necessarily the best way of expressing the facts, but their structure serves to emphasize the similarity of examples (3) and (4). First, we need to encode the fact that if someone is happy, then that is (defeasibly) caused by something that is "better" for that person (for instance, better than it has been or better than it might be):

(8) $happy'(e_1, x) \supset better(e_2, x) \wedge cause(e_2, e_1)$

That is, if e_1 is x's being happy, then there is (defeasibly) an e_2 that is some sort of "betterment" for x and e_2 is the cause of the happiness e_1. In addition, a better job for someone is a promotion:

(9) $better(e_2, x) \wedge job(e_2, x) \supset promotion(e_2, x)$

That is, if e_2 is x's job and also some sort of "betterment" for x, then e_2 is a promotion for x.

In interpreting the second sentence of (3), we try to prove that there is a promotion. Using Axiom (9), we then try to find the existence of a better job. By using Axiom (8), we could establish that there is something better if we can find happiness, and the first sentence provides that. The minimal abductive proof is then obtained by assuming that the something better is a job. That is the new information conveyed by the assertion of a promotion. This interpretation is illustrated in Figure 2.

$$\boxed{happy'(e_1, A)} \qquad\qquad promotion(e_2, A)$$

$$happy'(e_1, A) \supset better(e_2, A) \wedge cause(e_2, e_1)$$

$$better(e_2, A) \wedge \boxed{job(e_2, A)} \supset promotion(e_2, A)$$

Figure 2: Example (3) as Coreference Resolution.

As a side effect of finding the coreferential link between "promotion" and "happy", we have inferred the causal relation that we wanted. No explicit search for or recognition of a coherence relation was necessary. It would seem, therefore, that at least in *some* cases, it is not necessary to look for a coherence relation between two segments. It happens anyway.

2.2 The Interpretation of Vague Predicates and the Recognition of Coherence

The difficulty with this argument, however, is that it does not distinguish between example (3) and the following example.

(10) Ann is so happy! She is tall. She just got a promotion!

Her being tall (except in bizarre contexts) has nothing to do with her happiness or with her promotion. Example (3) is a reasonable thing to say, and (10) is not. Yet the coreference account for "promotion" goes through just the same. Example (10) has cohesion but not coherence, whereas example (3) has both.

For a formal account of the intuitive difference between examples (3) and (10), let us first look at another local pragmatics problem—the interpretation of vague predicates, specifically, the vague predicate conveyed by the adjacency of two nouns in a compound nominal, such as "kitchen light". The explicit content of the phrase "kitchen light" is

(11) $kitchen(x) \wedge nn(x, y) \wedge light(y)$

In addition to the information conveyed by the word "kitchen" and the word "light", there is the information that there is some vaguely specified relation between the kitchen and the light. In IA this relation was called nn, and a pragmatic strengthening of this relation is the solution to the compound nominal resolution problem.

Suppose we know that a kitchen is a room:

(12) $kitchen(x) \supset room(x)$

and that the "in" relation is a possible nn relation:

(13) $in(y, x) \supset nn(x, y)$

These axioms together with Axiom (5) yield the interpretation that the kitchen light is the light that is in the kitchen. This interpretation is shown in Figure 3.

Just as a vague relation is conveyed by the adjacency of the nouns in a compound nominal, so also is a vague relation conveyed by the adjacency

$$kitchen(x) \qquad nn(x,y) \qquad light(y)$$

$$room(x) \supset in(y,x) \wedge light(y)$$

Figure 3: Interpretation of "kitchen light".

of two segments of discourse. We may call this relation $CoRel$. Thus, when we are interpreting example (3) and we must prove its explicit content, it is not enough to abductively prove (7). That is not all that is conveyed. We must prove

(14) $happy'(e_1, A), CoRel(e_1, e_2, e), promotion(e_2, A)$

The $CoRel$ predication encodes the information conveyed by the adjacency of the two sentences. We can prove the $CoRel$ predication with Axiom (1) if we find a causal relation between e_1 and e_2. But that is provided by Axiom (8), used in the establishment of coreference, broadly construed. The fact that axioms get used for more than one problem contributes to the proof's minimality. The proof of (14), that is, the interpretation of example (3) handling both coreference and coherence, is shown in Figure 4.

$$\boxed{happy'(e_1, A)} \qquad CoRel(e_1, e_2, e_1) \qquad promotion(e_2, A)$$

$$happy'(e_1, A) \supset better(e_2, A) \wedge cause(e_2, e_1)$$

$$better(e_1, A) \wedge \boxed{job(e_2, A)} \supset promotion(e_2, A)$$

Figure 4: Interpretation of Example (3) with Coreference and Coherence.

By contrast, in example (10) there is no explanation of the relation conveyed by the adjacency of the successive segments. This is illustrated in Figure 5, where coreference is established but coherence is not. The fact that there is no pragmatic strengthening of the $CoRel$ relations in this example is the formal correlate of our sense of the incoherence of the text. An analogous compound nominal example is "kitchen job light", where the nn relations between "kitchen" and "job" and between "job" and "light" would be unexplained.

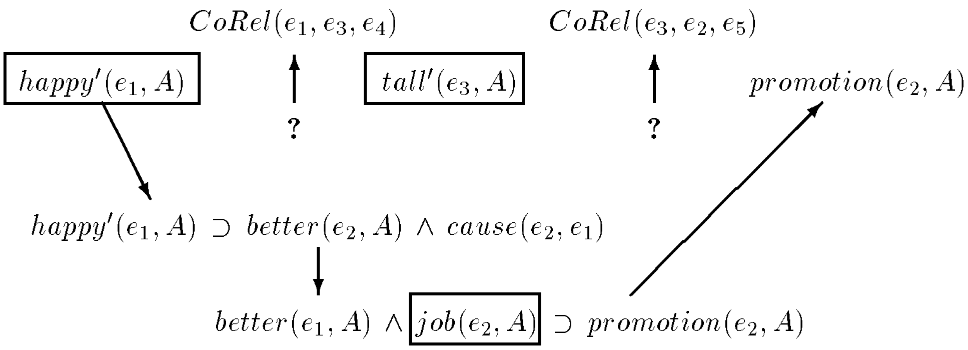

Figure 5: The Incoherence of Example (10).

In summary, considerations of coreference and other aspects of local pragmatics can sometimes provide all the information that would be provided by a coherence analysis, but they leave the coherence (or incoherence) itself unexplained.

3 Assertion in Explanations

3.1 Intentionality in the Abductive Framework

The issue of what the assertion of the composite segment in Axiom (1) should be requires us to build up some more machinery. The assertion of the composite segment is of no significance when that segment is viewed in isolation. It plays a role only when we try to embed the segment in a larger context. We can see the problem without moving beyond the informational aspect of discourse by embedding (3) in a dialogue. Suppose the reply to (3) is

(15) That's wonderful!

We need to be able to figure out whether "that" refers to the happiness or the promotion. Clearly it is the promotion. The Evaluation relation of Hobbs (1985) requires that the second segment (15) be a comment on the assertion of the first (3). This will be satisfied if we resolve "that" to the assertion of composite segment (3). But that means the assertion of composite segment (3) must be e_2, not e_1 as provided by Axiom (1).

However, when we hear or read (3) in isolation, we already know the main point of the composite segment is the promotion; we do not need to wait for the response. The reason is that the promotion is a newsworthy event, whereas happiness is generally not, or at least that promotions are much more newsworthy than happiness.

To accomodate this observation in the formal framework, we have to move beyond the text and consider the environment in which the text is embedded. We have to take into account that discourse is intentional, goal-directed action. An approach that incorporates this basic assumption into the abductive framework is sketched in Hobbs (1996b). What we present here is a very simplified version of that approach.

The key idea is that an intelligent agent's interpretation of the world in general is an abductive process. An agent is presented with a multitude of observables occurring in a history of observables, and the job of the agent is to figure out the best explanation for that set of observables, consistent with the history. As before, this explanation requires that assumptions be made, and optimality involves factors such as minimizing the number of assumptions, minimizing the length of the chains of inference used, maximizing plausibility and salience, maximizing consistency with the rest of what the agent knows, exploiting redundancy in explanation, and so on.

In discourse, the observable to be explained is that someone uttered a string of sounds. The most reasonable explanation for this is that the action was intentional and was intended to convey some meaning to an addressee. It serves the speaker's goals, and one of the speaker's goals is to report newsworthy events. Skipping over several steps, we can represent this consideration by the following axiom:

(16) $Segment(w, e) \wedge newsworthy(e) \supset utter(i, u, w)$

That is, if w is a coherent segment of discourse that describes a situation e where e is newsworthy, then (defeasibly) some speaker i will utter w to some hearer u.

Although our notion of newsworthiness is context-relative and not restricted to eventualities that are *intrinsically* newsworthy, we will simplify the example by assuming that promotions are intrinsically newsworthy:

(17) $promotion(e_2, x) \supset newsworthy(e_2)$

That is, if e_2 is x's promotion, then e_2 is newsworthy.

A coherent segment may consist of a single sentence, or it may consist of more than one sentence. In the latter case, it must have been composed of smaller segments related by coherence relations. The following axiom, from IA, Section 6.3, captures this:

(18) $Segment(w_1, e_1) \wedge Segment(w_2, e_2) \wedge CoRel(e_1, e_2, e)$
$\supset Segment(w_1 w_2, e)$

If the string w_1 is a coherent segment conveying e_1, string w_2 is a coherent segment conveying e_2, and there is a coherence relation relating e_1 and e_2 yielding e, then the concatenation $w_1 w_2$ is a coherent segment conveying e.

A single, grammatical, interpretable sentence is also a coherent segment. In Hobbs (1996c) a substantial fragment of the syntax of English is developed in a manner somewhat similar to the treatment in IA, Section 6.1, integrating syntax with the other processes of interpretation in a thorough fashion. In this treatment, the predication

(19) $Syn(w, e, \ldots)$

says that the string w is a grammatical, interpretable string of words describing the situation or entity e. The arguments of Syn indicated by the dots include information about complements and various agreement features. The interface between syntax and discourse is effected by the axiom

(20) $Syn(w, e, \ldots) \supset Segment(w, e)$

A grammatical, interpretable sentence w describing eventuality e is a coherent segment of discourse describing e. Composition in this framework is effected by axioms of the form

(21) $Syn(w_1, e, \ldots, y, \ldots) \wedge Syn(w_2, y, \ldots) \supset Syn(w_1 w_2, e, \ldots)$

A string w_1 whose head describes the eventuality e and which has an unsaturated argument y can be concatenated with a string w_2 describing y, yielding a string describing e. (This axiom, by the way, is a specialization of Axiom (18), in which the relation conveyed by adjacency is the predicate-argument relation.) The interface between syntax and world knowledge in this approach is effected by "lexical axioms" of a form illustrated by

(22) $read'(e,x,y) \wedge text(y) \supset Syn(\text{“read”}, e, \ldots, x, \ldots, y, \ldots)$

If e is the eventuality of x reading y (the logical form fragment supplied by the word “read”), where y is a text (the selectional constraint imposed by “read”), then e can be described by a phrase headed by the word “read” provided it picks up, as subject and object, phrases of the right sort describing x and y.

In this note, where possible, we avoid the details of syntax by drawing links from the logical form fragment for a sentence or phrase directly to its Syn predication.

With this machinery we can now embed example (3) in its context of utterance. This is illustrated in Figure 6.

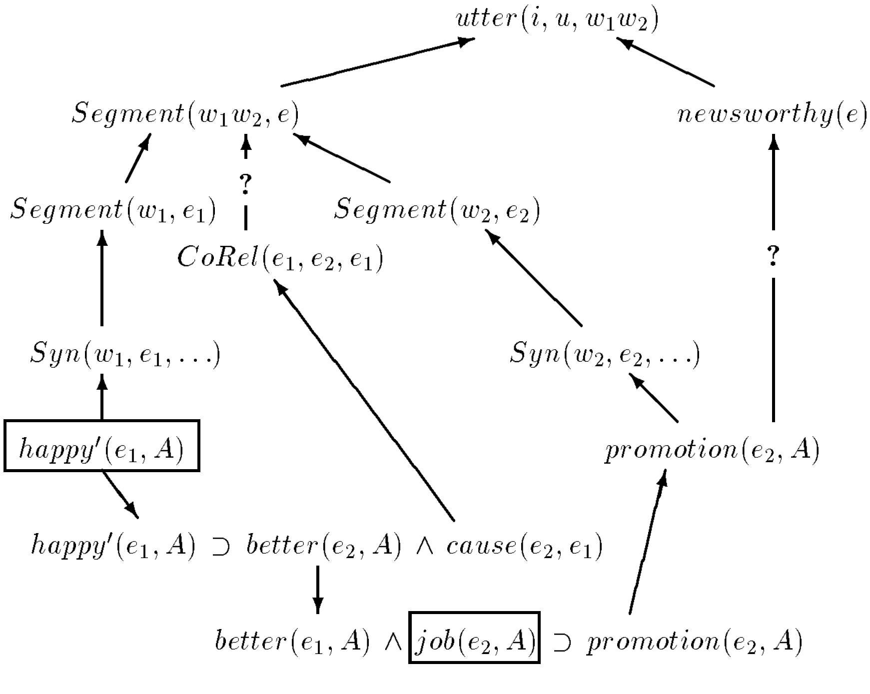

Figure 6: Explanation of the Uttering of (3).

The uttering of example (3) is the observable to be explained. Using Axiom (16) this can be explained by being a coherent segment conveying newsworthy information. By Axiom (18), we can show it is a coherent segment of discourse by showing it is composed of two coherent segments of discourse related by a coherence relation. By Axiom (20) each of the component strings can be shown to be coherent segments by showing them to be grammatical sentences. We skip over the details of syntax by generating the logical forms for the sentences immediately below the Syn predications. The remainder of the proof is as in Figure 4.

There is a problem here however. Axiom (16) insists that the assertion e associated with the segment of discourse $w_1 w_2$ be newsworthy. By Axiom (17) the promotion e_2 is newsworthy, and is thus a good candidate. But the assertion provided by the Explanation pattern, Axiom (1), is the happiness e_1. There is a mismatch. We have now pinpointed in the formal framework exactly where the objection of Redeker (in press) has its force.

3.2 Metonymy and Assertion

To resolve this difficulty, we need to look at a third phenomenon from local pragmatics—metonymy. In

(23) I like to read Proust.

Proust is a person, whereas only texts can be read. We must coerce Proust into some associated text. This can be handled within the syntactic framework sketched above by means of the following kinds of axioms:

(24) $Syn(w, x, \ldots) \wedge rel(x, y) \supset Syn(w, y, \ldots)$
(25) $Syn(w, e, \ldots, x, \ldots) \wedge rel(x, y) \supset Syn(w, e, \ldots, y, \ldots)$

If we wish to express something involving some x and there is a coercion relation rel between x and y, then we can convey the same information by replacing x with y.

We do not deal here with the very hard problem of just what relations can function as a coercion relation. In the right contexts it can be almost anything. However, as Pustejovsky (1992) argues, coercion relations very often involve causation—the Agentive and Telic qualia, in Pustejovsky's terms. Here we encode possible coercion relations explicitly in axioms such as (28) and (33).

For example (23) three axioms are required.

(26) $Proust(y) \supset write(y, x) \wedge novel(x)$
(27) $novel(x) \supset text(x)$
(28) $write(y, x) \supset rel(x, y)$

That is, Proust wrote novels, a novel is a text, and the *write* relation is a possible coercion relation.

The interpretation of example (23) is then shown in Figure 7.

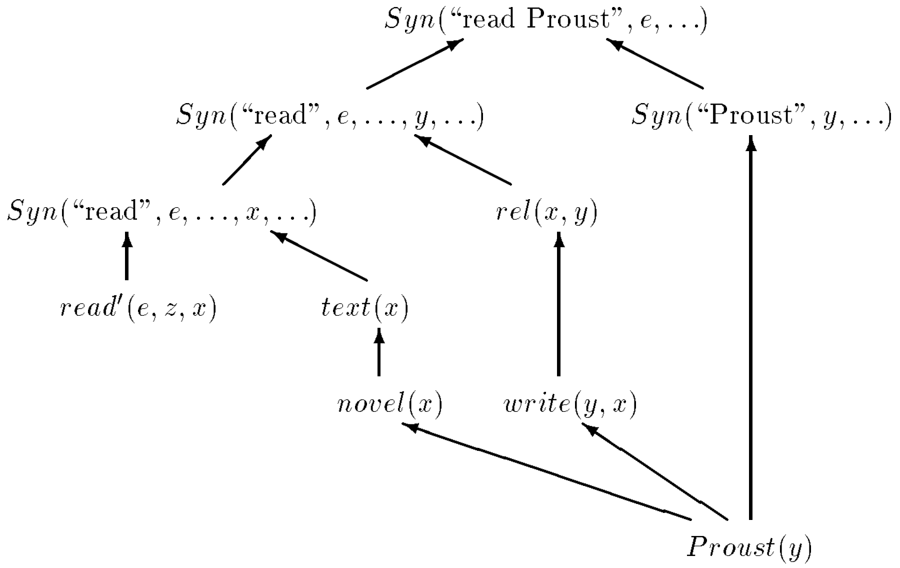

Figure 7: Interpretation of Example (23).

Grammatically subordinated material in sentences, such as adjectives modifying nouns, can sometimes carry the primary information of a sentence, and thus be its assertion. For example, in

(29) I have a sore throat.

it is not the possession of a throat that is being asserted, but the soreness of the throat the hearer already knows the speaker has. This can be viewed as an example of metonymy as well. The explicit assertion of the sentence, the possession, is coerced into the soreness of what is possessed. The possession is related to the throat and the throat is related to the soreness, both by properties that are explicit in the logical form of the sentence and are thus emminently accessible.

This example requires that both possession and soreness be possible coercion relations:

(30) $have'(e_1, z, x) \supset rel(x, e_1)$
(31) $sore'(e_2, x) \supset rel(e_2, x)$

Moreover, coercion relations compose:

(32) $rel(x, y) \wedge rel(y, z) \supset rel(x, z)$

Figure 8 now shows the interpretation of example (29). The having e_1 is taken as the eventuality conveyed by the verb phrase, but that is coerced into the soreness e_2, using as a coercion relation a composite of the having and the soreness. We have not indicated what constraint forces this coercion, but it could be the newsworthiness constraint encoded in Axiom (16), or it could be the weaker constraint that what is said should simply be informative.

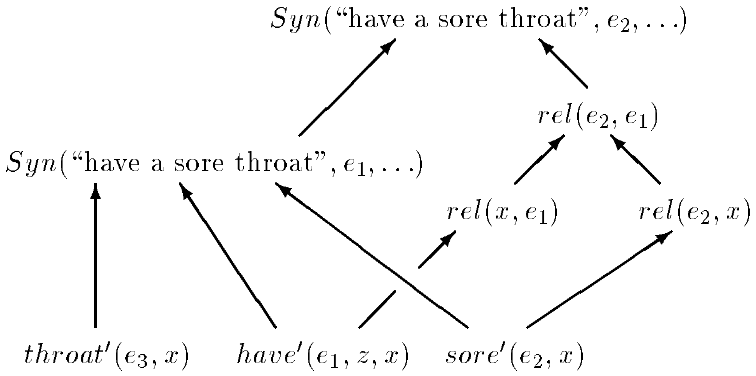

Figure 8: Interpretation of Example (29).

We now have the machinery we need for dealing with example (3). It can be seen as an instance of metonymy. Just as in example (29) we had to coerce from the having to the soreness, in example (3) we have to coerce from the happiness to the promotion. We need a relation linking them, but the causality provides just such a relation.

(33) $cause(e_2, e_1) \supset rel(e_1, e_2)$

In addition, we need to generalize the Metonymy Axiom (24) to work for coherent segments of discourse and not just grammatical fragments of sentences:

(34) $\quad Segment(w, e_1) \wedge rel(e_1, e_2) \supset Segment(w, e_2)$

If w is a coherent segment conveying e_1 and e_1 is related to e_2 by a coercion relation rel, then w can also function as a coherent segment conveying e_2.

The pervasive phenomenon of metonymy provides exactly the solution we need to overcome the difficulty pinpointed in Figure 6. The complete (for this note) interpretation of example (3) is illustrated in Figure 9. The generalized Metonymy Axiom (34) resolves the conflict between the Explanation pattern (1) and the desideratum of Axiom (16) that an utterance convey newsworthy information; it is precisely the requirement of newsworthiness that forces metonymic interpretation to override the default given by Axiom (1). By Axiom (33) the causal relation between the promotion and the happiness provides the required coercion. That the same causal relation figures in the solutions to the coreference, coherence, and metonymy problems all at once is a virtue of the interpretation, because it is an instance of the exploitation of redundancy that leads to minimal abductive proofs.

Metonymy at the discourse level is a common phenomenon. For instance, Moore and Pollack (1992) provide an example that purports to show that the semantic and pragmatic structures of discourse can sometimes be at odds with each other:

(35) Come home by 5:00. Then we can go to the hardware store before it closes. That way we can finish the bookshelves tonight.

Pragmatically the principal message is to come home early tonight, making the first sentence the "nucleus"; semantically, each sentence could be seen as presenting an enabling condition for the state of affairs in the subsequent one, making the final sentence the "nucleus". Note, however, that causal chains are reversed by being embedded in desire predicates. If A enables B, then an agent's wanting B will cause the agent to want A. By coercing the physical eventuality of each of the sentences into the wanting of that eventuality, there is a straightforward backward causal chain, and the pragmatic interpretation of the text is derived by Axiom (1) without difficulty.

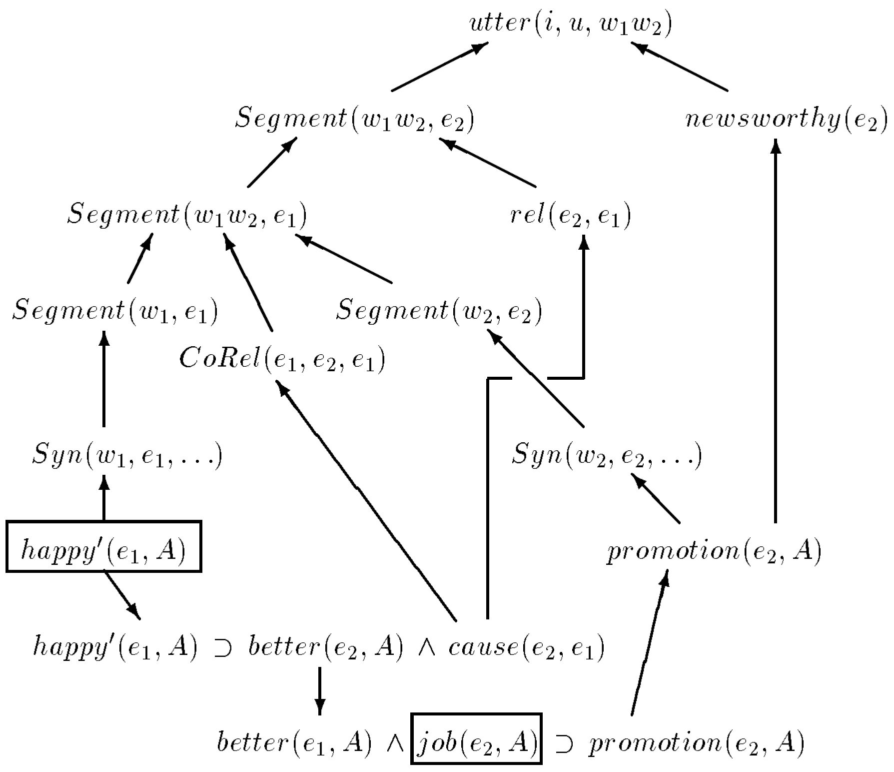

Figure 9: Interpretation of Example (3) as Metonymy.

4 Consequences for the Modelling of Coherence Relations

We have shown in this note how various general mechanisms of discourse interpretation can play a role in the recognition of coherence. In the abduction model, the availability of those general processes allows the specification of the coherence relations themselves to be minimized, thus reducing the need to encode elaborate conditions in the definitions of coherence relations and to distinguish many specific relation types or application cases. A rather limited set such as, for instance, that proposed by Hobbs (1990) would seem quite sufficient. In terms of Redeker's (1990, in press) model, this could be seen as a terse description of the rhetorical and sequential components (the pragmatic structure), with conditions formulated mostly in terms of the ideational component (the semantic structure).

To illustrate how the integration of coherence processing with more general processes may obviate the need to introduce specific subclasses of coherence relations, let us consider again our analysis of example (3) as metonymy (Figure 9). The analysis still contains the Explanation relation, although the final interpretation is equivalent to a coherence relation that could be defined as follows:

(36) $cause(e_2, e_1) \supset CoRel(e_1, e_2, e_2)$

where the assertion of the composite segment is e_2, not e_1 as in Axiom (1). This relation might be interpreted as a special case in the class of Hobbs's (1990) Evaluation relations with the evaluation satellite preceding the evaluandum, as in

Something wonderful happened. Ann got a promotion.

The pragmatic properties of example (3), however, are quite different from those of ordinary evaluations. The implicit presence of an Explanation relation in our analysis of (3) provides one way of describing this difference. It would not seem necessary or even useful, therefore, to introduce the special kind of relation defined in (36).

A final remark on coherence relations is in order. It is sometimes believed that to adopt a framework that includes coherence relations is to commit oneself to a claim that the hearer is consciously aware of the occurrence of the coherence relations. This is not a necessary consequence, as consciousness and interpretation are largely orthogonal issues, just as consciousness and syntax are. But the approach to coherence we have formulated in this note is not even compatible with consciousness of the specific coherence relation. There is no mention of "explanation" in the Explanation pattern (1). To make this explicit we would have to rewrite (1) as

(37) $cause(e_2, e_1) \supset Explanation(e_1, e_2)$
(38) $Explanation(e_1, e_2) \supset CoRel(e_1, e_2, e_1)$

In our view, speakers are sometimes aware of the coherence relation; more often they are not. In any case, from a formal point of view, the difference is slight, whatever its importance on a more intuitive level.

In summary, we have shown that an analysis of the coherence of a text is necessary, even though much of what is discovered thereby will often have been forced from coreference considerations alone. The requirement to establish coherence is similar to analogous interpretation requirements in syntax and local pragmatics. The interpretive moves, such as metonymic interpretation, that are available in local pragmatics are also available in discourse processing, and that fact makes possible a significant economy in the specification of what constitutes a coherent discourse.

Acknowledgements

The authors have profited from discussions with Andrew Kehler about this work. The work described here was supported in part by the National Science Foundation and Advanced Research Projects Agency under Grant IRI-9304961 (Integrated Techniques for Generation and Interpretation).

References

[1] Hobbs, Jerry R., 1985. "On the Coherence and Structure of Discourse", Report No. CSLI-85-37, Center for the Study of Language and Information, Stanford University.

[2] Hobbs, Jerry R., 1990. "Literature and Cognition", CSLI Lecture Notes Number 21. Stanford: Center for the Study of Language and Information.

[3] Hobbs, Jerry R., 1996a. "The Structure of Discourse", to be published in D. Everett, ed., *Discourse: Linguistic, Computational, and Philosophical Perspectives.*

[4] Hobbs, Jerry R., 1996b. "On the Relation Between the Informational and Intentional Perspectives on Discourse", in *Computational and Conversational Discourse: Burning Issues—An Interdisciplinary Account*, NATO ASI Series, Series F: Computer and Systems Sciences, Vol. 151, E. Hovy and D. Scott, editors, Springer, Berlin, Germany, pp. 139-157.

[5] Hobbs, Jerry R., 1996c. "The Syntax of English in an Abductive Framework", manuscript.

[6] Hobbs, Jerry R., Mark Stickel, Douglas Appelt, and Paul Martin, 1993. "Interpretation as Abduction", *Artificial Intelligence*, Vol. 63, Nos. 1-2, pp. 69-142.

[7] Mann, William C., and Sandra A. Thompson, 1988. "Rhetorical Structure Theory: Toward a Functional Theory of Text Organization", *Text*, Vol. 8, No. 3, pp. 243-281.

[8] Moore, Johanna D., and Martha E. Pollack, 1992. "A Problem for RST: The Need for Multi-Level Discourse Analysis", *Computational Linguistics*, Vol. 18, No. 4, pp. 537-544.

[9] Pustejovsky, James, 1992. *The Generative Lexicon: A Theory of Computational Lexical Semantics*, MIT Press, Cambridge, Massachusetts.

[10] Redeker, Gisela, 1990. "Ideational and pragmatic markers of discourse structure", *Journal of Pragmatics*, Vol. 14, No. 3, pp. 305-319.

[11] Redeker, Gisela, in press. "Coherence and Structure in Text and Discourse", to be published in William Black and Harry Bunt, eds., *Computational Pragmatics: Abduction, Belief, and Context.* London: University College Press.