AN ALGORITHMIC FRAMEWORK FOR SPECIFYING THE SEMANTICS OF DISCOURSE RELATIONS

ALISTAIR KNOTT
Department of Computer Science, University of Otago, New Zealand [1]

In this paper, a new framework is proposed for defining the semantics of discourse relations, and for defining the semantics of the utterances that relations link together. The proposal is to define relations in terms of the operation of an algorithm simulating the mental state of an agent interacting with the world. The algorithm interleaves perception, theorem-proving and action: the denotation of a complex utterance containing a relation between two simpler utterances is taken to be the description of the operation of the algorithm during the time interval identified by their Reichenbachian reference times. This proposal is presented in detail for two discourse relations. Its potential application in the treatment of mood, tense, aspect and dialogue structure is also discussed in very general terms.

Key words: discourse relations, mood, tense, aspect, grounding.

1. INTRODUCTION

The main goal of this paper is to suggest a new framework for defining the semantics of discourse relations, and consequently, for defining the semantics of the utterances that relations link together. This framework is introduced in the context of two specific relations, both signalled textually by the connective *but*. We begin in Section 2 by describing the two relations and some differences between them. In Section 3, the new framework is introduced, and applied to the two cases. In Sections 4 and 5, we consider some of the wider implications of the framework, and its possible application in the treatment of a range of discourse-semantic phenomena.

2. EXPECTATION-BASED AND PLAN-BASED *BUT*

Analyses of the discourse connective *but* often refer to it as a marker of 'violated expectation'. In the text in Example 1, for instance, the clause introduced by *but* presents information which is unexpected given the information given in the first clause.

(1) Bob had been up all night with the baby, but he was looking fresh as a daisy.

Recently, a number of researchers have formalised the violated-expectation sense of *but* by giving a presuppositional account of the connective (see e.g. Knott, 1996; Lagerwerf, 1998). The idea is that *but* presupposes a general **defeasible rule**, and asserts that the rule is defeated in the situation being described. In the above example, the fact that Bob had been up all night, together with the rule that people who have been up all night are normally tired the next day, generates a defeasible conclusion that he will be tired; the second clause of the text indicates that this conclusion is defeated on this occasion. We will represent the presupposed rule, as Lagerwerf does, using Asher and Morreau's (1991) logic of common-sense entailment, in which the standard logical connectives are augmented with a 'defeasible implication' operator $>$:

[1] The author would like to thank Bonnie Webber and Jerry Hobbs for discussion of an earlier version of this article, and the three PACLING reviewers for their useful comments.

$$\forall p\ stayed\_up\_all\_night(p) > looks(p, exhausted)$$

However, not all occurrences of *but* are easily characterised as violated expectation. Consider the following text:

(2) Mary looked all over the shop, but she didn't find any miso sauce.

The general scenario presented here is one where an agent (Mary) attempts to bring about some desired state of affairs (finding miso), but for some reason it fails to materialise. This scenario must be distinguished in the general case from one of violated expectation. While in some circumstances it might be that we expect a goal being pursued by an agent to be attained, this is by no means always the case: the success or failure of the agent's plan frequently depends on conditions in the world being right as well as on the agent's goal-pursuing actions. For instance, in the above example, the presence or absence of miso in the shop is crucial for determining whether we expect Mary to find miso. But note that even if it is known that there is unlikely to be miso in the shop, the connective *but* is still appropriate.

(3) Mary looked all over the shop, but as expected she didn't find any miso sauce.

In summary, rather than treating the *but* in this case as signalling a violated expectation, it seems preferable to describe it as signalling a frustrated or interrupted plan.

To clarify what we mean by a frustrated plan, we will consider the plan-based sense of *but*, and its relationship to the expectation-based sense, in a little more detail. Note that we can still use a defeasible rule as the basis for a representation of the situation in Example 2. Defeasible rules can function both to generate expectations in cases of incomplete knowledge, and to drive planning algorithms to enable agents to pursue goals. The latter mode of reasoning is termed *abductive planning*: the basic strategy is to look for rules whose right-hand side matches the current goal, and whose left-hand side contains situations which are true, or can be *assumed* to be true if the agent carries out the appropriate actions. A completed proof will then specify the set of actions which need to be carried out. See Shanahan, 1997; Stone, 1998 for some background on this planning paradigm.

Note also that even situations which are not known and are outside the control of the agent can be assumed, if the agent is willing to risk that the plan is not successful. In fact, this is exactly the situation we have in Example 3: Mary is keen enough to get the miso sauce that she is willing to look all over the shop even though she knows she's unlikely to find what she wants. Without going into any detail, we will briefly sketch a method for representing this type of scenario, using a scheme presented in Hobbs *et al.* (1993) called **weighted abduction**. Hobbs *et al* use abduction to generate explanations for observed events: an event is matched to the right-hand side of a rule, and then the likeliness of the circumstances given on its left-hand side is considered. In the weighted abduction scheme, each proposition on the left-hand side of a rule is associated with a numerical 'cost', with higher costs indicating less likely situations; if the truth of a proposition is not known, its cost is calculated, and a cumulative record is kept of the total cost of assumptions made in a proof. The most likely explanation is then the proof with the lowest cost.

In our abductive planning scenario, costs are still used to represent the likelihood of assumptions, but these likelihoods are balanced against the importance of the goal to be achieved. We associate the agent's *goal* with a number too, and stipulate that a plan is worth pursuing if the cost of the goal is lower than the cost of the assumptions which we need to make to prove the goal. In this way, it is possible for an agent to pursue a goal using a plan which she regards as very unlikely to succeed—which is precisely what we want for Example 3. The relevant rule here would be as follows:

$$\forall x, p\ looks\_all\_over(p, shop) \land in(shop, x) > finds(p, x)$$

The agent's goal is $finds(agent, miso)$; the event $looks\_all\_over(agent, shop)$ is one which the agent can bring about herself; and the situation $in(shop, miso)$ is one which will be associated with a cost: even if this cost is very high, provided the priority of the goal is higher, the plan of looking all over the shop will be worth pursuing.

In summary, we can use defeasible rules to represent both varieties of *but* relation: in one case the defeasible rule is being used in a planning paradigm; in the other, it is being used in an inference paradigm. In the next section, we spell out the two required representations in more detail, and discuss what they have in common. In the remainder of the paper we consider some of the wider ramifications of the representations adopted.

3. A FRAMEWORK FOR REPRESENTING PLANS AND EXPECTATIONS

3.1. A definition for plan-based *but*

How should we frame a representation of Example 2 which makes explicit the underlying frustrated plan? The account we propose here derives from the observation that representing a plan which goes wrong involves specifying a temporal succession of events. This fact is just a matter of conceptual analysis: we have to represent a plan, and the initiation of an action which implements the plan, and finally the frustration of the plan—the plan has to exist before it can be pursued by the agent, and it has to be pursued before its success or frustration is known, so a minimum of three time points must be assumed. More concretely, we might state the relevant series of events as follows (assuming a planning agent p).

- T_1: p selects a goal state s to pursue.
- T_2: p initiates planning to determine whether/how s can be achieved, backward-chaining on s to search for a rule whose right-hand side is s and whose left-hand side contains a situation which is known to be achievable through action a (possibly given certain other conditions which p is willing to assume).
- T_3: p finds an action a which meets these requirements (possibly after recursion).
- T_4: p executes a.
- T_5: p checks whether s has been achieved.
- T_6: p discovers that s has not been achieved.

In relation to Example 2, the event associated with *Mary looked all over the shop* is associated with T_4 in the above list, and the event associated with *she didn't find any miso sauce* is associated with T_5 and T_6. A description of the sequence of events T_1–T_6 is certainly sufficient to represent the existence of a frustrated plan.

It is natural to think of the series of events outlined above as stages in the operation of an *algorithm* interleaving planning, action and perception, being executed by Mary, the protagonist of the narrative. Note that these events do not describe the algorithm itself, but its execution in a particular set of circumstances—namely those described in the text. In fact, what we are proposing in this paper is that the denotation of the text in Example 2 is simply a description of a certain sequence of stages during the execution of this algorithm (specifically, stages T_4 to T_6), and that it presupposes a certain sequence of stages prior to this sequence (specifically, stages T_1 to T_3). To be more concrete, we can suggest that the semantics of the text be thought of as a function from the state the algorithm-processor is in before the utterance is made (reflecting the fact that stages T_1–T_3 have just occurred) to the state it is in after the utterance is accepted. We can then frame a general definition of the plan-based *but* relation as the characterisation of a sequence of stages of a certain type—for instance, by replacing the particular goal-state s, action a and agent p in the above description by variables S, A and P of appropriate types.

The general semantic framework just proposed will be discussed in more detail in Section 4. But regardless of the wider implications of the proposed framework, it seems hard to say how we could represent what is going on in Example 2 *without* making reference to a complex process of the kind described. We have argued that the example must be analysed as a frustrated plan, and cannot be reduced to a violated expectation, and the sequence of events described above does not depart far from what must minimally be assumed to represent such a scenario.

3.2. A definition for expectation-based *but*

In this section, we argue that there are good reasons for adopting a similar analysis for the more conventional expectation-based *but* relation found in Example 1. The analysis we are proposing here is to characterise the situation as one where an agent p wants to know whether Bob is tired, frames an expectation about the answer on the basis of the fact that Bob has been up all night, then learns from a more reliable source (let us say perception) that the actual answer is different from the expected one. As a first approximation, the relevant stages of the algorithm look something like this:

T_1: p selects the goal to answer the question 'is Bob tired?'.
T_2: p plans to answer this question by defeasible reasoning.
T_3: p backward-chains on the proposition 'Bob is tired', searching for a rule whose right-hand side is 'Bob is tired' and whose left-hand-side is a fact (or facts) which are known to be true.
T_4: p finds a rule whose left-hand side is 'Bob has been up all night'.
T_5: p ascertains that this is true by consulting a knowledge base.
T_6: p concludes defeasibly that Bob is tired.
T_7: p decides to answer the question 'Is Bob tired?' perceptually as well.
T_8: p discovers that Bob is not tired.
T_9: p deletes the defeasibly-reached conclusion.

This characterisation of expectation-based *but* is considerably more complex than those normally given in the literature. The presuppositional account normally given does not decompose the representation into a sequence of stages in this way. Or at least, only two stages are assumed: the situation before the text is read (in which the presupposed defeasible rule must be true), and the situation after it is read (in which the rule is indicated to be defeated in the scenario being described). Moreover, little is normally said about the identity of the agent doing the theorem-proving: it is typically assumed to be the interpreter of the text. However, a good case can be made for the introduction of the third-party protagonist, and also for the finer granularity of the sequence of stages. To begin with the former point: being specific about the identity of the agent doing the theorem-proving allows us to address a perennial problem with analyses of this type of relation, described below.

A distinction is frequently made between relations which hold in virtue of the propositional content of the linked utterances (termed SEMANTIC) and those which hold in virtue of their illocutionary force or intended effect (termed PRAGMATIC)—see e.g. Sanders *et al.* (1992) for a discussion of these categories. SEMANTIC relations are used in descriptive texts, where a hearer is simply being told about a situation in the world, while PRAGMATIC relations are used in argumentative or persuasive texts, where the hearer's acceptance of the speaker's authority cannot be assumed.

The relation in Example 1 is SEMANTIC: essentially, the hearer is just being told about an unusual circumstance in the world. The difficulty with this analysis is as follows. On the one hand, this 'unusualness' is hard to think of as being in the world itself: it is better thought of as arising from *inferences* made by an observer of the world. On the other hand, we do not want to suggest that the text is in any way argumentative, or that any inferences drawn by the hearer can challenge the speaker's authority as to what actually occurred in the situation described. By introducing the third party p into the picture, we are able to provide a clearer account: what is being narrated is an instance of theorem-proving by p at a certain time in the past; the speaker has complete authority in describing this, but nonetheless the deductive processes are occurring in an agent, rather than being part of the world in some mysterious way. Note that this more complex account of Example 1 is not incompatible with the analysis that a presupposed defeasible rule is being defeated; it just expresses this idea in more detail.

Given the above reasons for identifying the protagonist, what is the reason for describing the protagonist's theorem-proving operations at such a fine level of granularity? One reason is that this points the way towards a general account of plan-based and expectation-based *but*; see Section 3.3 for more on this point. But independently of this, there are indications that a finer-grained analysis would prove useful in a representation of the individual clauses in the text, as well as of the relation between them.

For one thing, the description of the agent's goal to discover how Bob is (in T_1 and also T_7), and the eventual satisfaction of this goal (in T_8) goes a long way towards an explanation of the information structure of the second clause in the text (*[Bob was looking]*$_T$ *[fresh as a daisy]*$_F$). Topic and focus are frequently analysed in terms of a question-answer structure. For instance, in the alternative semantics of Rooth (see e.g. Rooth, 1985; 1995), a sentence with topic-focus structure is to be interpreted relative to a set of propositions which constitute the possible answers to a question: the focussed element of the sentence is the one which specifies which of these alternatives is found to obtain.

As well as providing the framework for an account of topic-focus structure, the fine-grained representation of theorem-proving in Example 1 also provides a framework for an analysis of the tense and aspectual type of clauses. In Reichenbach's (1947) influential account of tense and aspect, the distinction between present and past tense is determined by a value called the 'reference time', which can be defined as 'the time which is being talked about in the text'. The reference time is distinguished from the 'event time', which is the time at which the event described in the text actually occurs, and from the 'speech time', which is the time at which the text is uttered. The distinction between event time and reference time allows us (among other things) to distinguish between the simple past, where event time and reference time coincide, and the past perfect, where the event time predates the reference time. (The distinction between past and present tense is expressed in terms of the relation between reference time and speech time: we will have more to say about this in Section 5.) Many subsequent accounts make use of a concept of the reference time of a sentence, often referred to as its 'temporal index'. But despite the utility of the concept of reference time, it is still not very clear how it is defined. What is meant by 'the time which is being talked about in the text', if this is distinct from the time at which the reported events actually occur? Representing a narrative text in terms of an algorithm being executed by a third-party observer allows a clear definition to be framed. We can simply take the reference time to be the time at which the relevant operation in the agent's algorithm occurs. If the agent is *perceiving* an event, as in the second clause of Example 1, then clearly the event time will be the same as the reference time. If the agent is drawing conclusions from events which have already occurred, as in the first clause of Example 1, then clearly the event time will predate the reference time. These considerations permit an explanation of why the first clause in Example 1 is in the past perfect, and the second clause is in the simple past. Note also that the anaphorically-determined nature of the reference time is something which can be predicted from the algorithmic model. If we assume that successive sentences in a discourse describe successive stages in the operation of the algorithm, then the discourse sensitivity of the reference time follows automatically.

3.3. Generalisations between plan-based and expectation-based *but*

A final piece of evidence in support of the above analyses is that they are expressed in terms of similar frameworks: consequently it is possible to generalise over the frameworks and produce a single definition for the connective *but* which covers both Example 1 and Example 2. Informally, the suggestion is as follows. In both of the proposed algorithm descriptions, three stages are apparent: (i) a stage at which a goal is posted; (ii) a stage at which some reasoning is performed; and (iii) a stage at which the reasoning is tested. For the plan-based example, the goal is *substantive*, concerned with achieving something in the world; the reasoning is a type of abductive planning, and the testing of the reasoning is a monitored substantive action. (There's often no way to test a plan except to carry it out.) For Example 2, we could express these stages like this:

(i) agent's goal: $finds(agent, miso)$
(ii) rule used in reasoning: $looks\_all\_over(p, shop) \wedge in(shop, item) > finds(p, item)$
(iii) method of testing: $execute(looks\_all\_over(agent, shop))$
(iv) result of testing: reasoning does not lead to goal.

For the expectation-based example, the goal is *epistemic*, concerned with finding something out about the world; the reasoning is deductive, and the testing of the reasoning is perceptual. For Example 1, we could express these stages like this:

(i) agent's goal: $knows(agent, looks(bob, ?))$[2]
(ii) rule used in reasoning: $stayed\_up\_all\_night(p) > looks(p, exhausted)$
(iii) method of testing: perceptual action to satisfy $knows(agent, looks(bob, ?))$
(iv) result of testing: reasoning does not lead to goal.

Naturally, there are still plenty of differences between the two cases. For instance, in the case of the epistemic goal ('find out how Bob is'), the perceptual testing of the original reasoning results in this reasoning being found unsatisfactory, *and also* happens to result in a new way of satisfying the goal being found; the agent ends up knowing that Bob was looking fresh as a daisy. Whereas in the case of the substantive goal ('get some miso'), the testing of the reasoning does not result in the goal being satisfied in a different way. (Although it may indicate some other possible plans to pursue.) Such disanalogies are to be expected, as the two senses of *but* are really quite different. Nonetheless, we can envisage a general definition for *but* which picks out what they have in common: the posting of a goal, a piece of reasoning in service of this goal, the testing of this reasoning, and the conclusion that it

[2] The question mark in this notation indicates a free variable: the goal is to bind this variable.

was faulty. There is still a great deal to be done before this definition is formalised; but note that this generalisation is impossible to state without a framework in which the operation of a theorem-prover is explicitly represented, at a fairly fine level of granularity.

4. A DISCUSSION OF THE ALGORITHM-BASED FRAMEWORK

The semantic framework proposed in Section 3.2 needs much clarification. The definitions we have suggested for the relation signalled by *but* all rely on the idea that the semantics of a descriptive text can be represented in terms of an algorithm being run by a planning/acting/perceiving agent. How does this idea relate to existing approaches in semantics? In some ways, it is quite different; however, there are some points of contact.

The clearest point of contact is with theories of dynamic semantics (see Muskens *et al*, 1997 for a review). These theories emphasise the parallels between utterances in natural language and steps in an algorithm. In the same way that operations in a computer program can be represented as functions which map one program state onto another, theories of dynamic semantics see natural language utterances as functions, mapping the information state of an interpreter prior to hearing the utterance onto the information state which results when the utterance has been understood and accepted. The fact that the present model represents the semantics of text units in terms of stages of an algorithm means that the present proposal falls into this class of dynamic theories.

However, there are some significant differences between the present model and other dynamic theories. To illustrate, we will make some comparisons with the dynamic theory presented by Frank Veltman (1996). For Veltman, an information state σ is a representation of the set of situations the agent regards as possible. The **update value** of a sentence S representing an atomic proposition p is defined as a function which takes an information state σ_1, and returns another information state σ_2 in which some of the possible situations in σ_1—namely, those which do not contain p—have been eliminated. The standard logical operations of negation, implication etc can all be defined in terms of set-theoretic operations on the set σ_1.

There are several differences between this model and the one we are proposing. Most importantly, an operation in our algorithm is not—at least in the general case—the action of interpreting a sentence, as it is in the above model. In fact, our operations are not (in the general case) linguistic actions at all; rather, they are 'cognitive' actions. The operations in Veltman's theory can be imagined as specified by an algorithm whose function is to interpret a sequence of sentences. The operations envisaged in our model are much more general-purpose, and include things like the posting of a goal, the execution of a piece of theorem-proving, or the asking of and getting answers to questions. But note that we can still define them as functions from one state of the agent to another. Consequently, we must ensure that the update function which characterises an operation in our model allows the representation of a much wider range of operations than simply the interpretation of sentences. Given that an operation is defined as a function that transforms one state into another, this means that the states in our model must specify not just what the agent knows or expects, but a much wider range of things, including (i) the agent's goals (in particular her epistemic goals), (ii) what goal the agent regards as immediately to be pursued, and (iii) what the outcome of an action in pursuit of some goal is.

Secondly, we allow that the cognitive states between which operations move us are not necessarily those of the hearer, as in Veltman, but can be those of a protagonist being described in the text. This emphasises the fact that we are not modelling the interpreter of a text, but the interpretation itself. Where the agent is the protagonist in a narrative, an interpreter's job will be to infer what the relevant stages of the agent's algorithm are, a task in which plan recognition is likely to feature prominently.

Thirdly, we envisage a more complex relationship between operations in the algorithm and units in the text than in Veltman's model. For Veltman, each update operation 'consumes' a single sentence. In our model, as the analyses of Examples 1 and 2 demonstrate, there are more operations than there are clauses. We can describe the mapping we envisage between operations and textual elements as follows:

- Some operations in the algorithm are represented by the discourse relation in which the clauses feature, rather than by the clauses themselves. In our analysis of Example 1, for instance, stages 6 and 9 do not feature in the representation of either of the atomic clauses.
- A single clause in the text can be represented by more than one operation. For instance, stages 3–5 are all involved in representing the first clause, and stages 7 and 8 are both involved in representing the second.
- A clause can be represented by *non-consecutive* operations in the algorithm. For instance, stage 1 should feature in the representation of the second clause, as well as in that of the first.
- Finally, a single operation in the algorithm can feature in the representation of more than one discourse element. (In fact, this requirement is necessary for the algorithm to describe how one discourse element links to another.)

Note that it is still possible, with the mapping specified above, to think of 'the meaning of a sentence' as a function from one state to another. The mapping just indicates that this function can be decomposed into a number of smaller functions.

5. CONVERSATIONAL PARTICIPANTS AS AGENTS

A final interesting feature of the framework to consider is that the process of participating in a conversation can be described in terms of the same kind of algorithm as the algorithm proposed for representing the semantics of sentences. Speakers and hearers are agents with goals, which they pursue by making or interpreting utterances, and these goals can succeed or fail. If we are able to *identify* the agent p with either the speaker or the hearer, some interesting generalisations emerge.

For one thing, note that our earlier identification of Reichenbach's reference time with the time at which the relevant stages of the agent's algorithm are being executed still makes sense. As already mentioned, Reichenbach distinguished between present and past tenses using the relation between speech time and reference time: for a sentence in the past tense, the reference time predates the speech time, while for a sentence in the present tense, the reference time is equal to the speech time. If we identify the agent with the speaker, then it follows automatically that the reference time (the time at which the agent is executing her algorithm) is the same as the speech time. So modelling the difference between present and past sentences in our account is simply a question of specifying the identity of the agent: third party or speaker.

The identity of the agent could also play a role in the representation of sentences of different moods. For instance, if we again assume the agent is the speaker, then the utterance of an interrogative sentence can be represented as a situation where the agent is initiating an action in pursuit of an epistemic goal (namely, to discover the answer to the question). If this agent indeed receives an answer from her interlocutor (which will be an indicative sentence with topic-focus structure), the situation can be modelled as one where the agent (now the hearer) is monitoring her action, and observes that it has been successful.

Finally, the conception of the agent could play a useful role in a theory of grounding in dialogue (see e.g. Poesio and Traum, 1997, 1998). One of the key concepts in the model we have proposed is that the agent needs to obtain feedback on whether or not an action undertaken in pursuit of a goal is successful. If the agent is the speaker, and her goal is a communicative one, being pursued in a utterance, then her method for obtaining feedback on the success of the plan being pursued is to listen to her interlocutor's response: an indication that the goal behind an utterance has been achieved is what Poesio and Traum call the 'grounding' of the utterance. In the question-answer scenario above, we have one example of a response indicating the satisfaction of a goal; but the phenomenon of grounding is much more general. Consider the following exchange:

(4) A: Go to bed!
B: Okay.

In this case, if we take the agent to be speaker A, then the agent's goal is that B goes to bed. A pursues this goal by uttering an imperative, and receives confirmation that her plan has succeeded in the form of B's assenting statement *okay.* Now consider two exchanges where A's plan fails:

(5) A: Go to bed!
B: But it's only nine o'clock!
(6) A: Bill's not dangerous in the slightest.
B: But he's carrying a gun!

In the former case, A's goal is the same as above: that B goes to bed. In the latter case, A's goal is to persuade B that Bill is not dangerous. In each case, B's response indicates that this goal has not been achieved. These are interesting cases, because in each case the signal of the failed plan is *but.* *But* is frequently noted in the literature on dialogue as signalling a 'blocking' move; however, this sense of *but* has not been tied to either of the narrative senses of *but* which we have been discussing in this paper. The fact that our analysis extends to this 'inter-speaker' variety of *but* is thus a point in its favour.

6. SUMMARY AND FUTURE WORK

This paper provides an informal outline of a new method for specifying the semantics of utterances in discourse, and consequently of framing definitions of discourse relations. There are interesting (though still very speculative) applications for work on tense, aspect, mood and dialogue structure: the framework could potentially provide a useful unified treatment of a number of discourse-semantic phenomena.

Naturally, much remains to be worked out. For one thing, the account needs to be extended to discourse connectives other than plan-based and expectation-based *but.* There are some indications about how this extension could be achieved. For instance, it has been shown (Knott, 1996) that a wide range of causal and contrastive connectives can be analysed in terms of presupposed defeasible rules. The method of reinterpreting presupposed defeasible rules in relation to operations in an agent's algorithm is therefore likely to be applicable to a wide range of connectives.

For another thing, a much more concrete conception of the algorithm underlying relation definitions is needed, which specifies in a clearer way how perception, action and theorem-proving are to be interleaved. Developing this algorithm is of course a research programme in its own right, within the larger AI fields of planning and theorem-proving. We hope that our approach to the task comes from an interesting new angle, in that considerations about what is needed for the semantics of linguistic phenomena and what is needed for a successful algorithm can mutually inform one another. We envisage that simultaneous development of an algorithm and of associated definitions of linguistic concepts will prove a productive research methodology.

References

Asher, N. and Morreau, M. (1991). Commonsense entailment: A modal theory of nonmonotonic reasoning. In *Proceedings of the 12th International Joint Conference on Artificial Intelligence*, Sydney, Australia.

Hobbs, J., Stickel, M., Appelt, D., and Martin, P. (1993). Interpretation as abduction. *Artificial Intelligence*, **63**.

Knott, A. (1996). *A Data-Driven Methodology for Motivating a Set of Coherence Relations*. Ph.D. thesis, Department of Artificial Intelligence, University of Edinburgh.

Lagerwerf, L. (1998). *Causal connectives have presuppositions*. Ph.D. thesis, Katholieke Universiteit Brabant, The Hague.

Muskens, R., van Benthem, J., and Visser, A. (1997). Dynamics. In J. van Benthem and A. ter Meulen, editors, *Handbook of Logic and Language*, pages 587–648. Elsevier, Amsterdam.

Poesio, M. and Traum, D. (1997). Conversational actions and discourse situations. *Computational Intelligence*, **13**(3), 309–347. Also appears as Research paper HCRC/RP-92, HCRC University of Edinburgh.

Poesio, M. and Traum, D. (1998). Towards an axiomatization of dialogue acts. In J. Hulstijn and A. Nijholt, editors, *Proceedings of the Twente Workshop on the Formal Semantics and Pragmatics of Dialogues (13th Twente Workshop on Language Technology)*, pages 207–222.

Reichenbach, H. (1947). *Elements of symbolic logic*. University of California Press, Berkeley, CA.

Rooth, M. (1985). *Association with Focus*. Ph.D. thesis, University of Massachusetts at Amherst.

Rooth, M. (1995). Focus. In S. Lappin, editor, *The handbook of contemporary semantic theory*, pages 271–297. Blackwell, Oxford.

Sanders, T. J. M., Spooren, W. P. M., and Noordman, L. G. M. (1992). Towards a taxonomy of coherence relations. *Discourse Processes*, **15**, 1–35.

Shanahan, M. (1997). Event calculus planning revisited. In *Proceedings of the European Conference on Planning*, pages 390–402. Springer Verlag: Lecture Notes in Artificial Intelligence (1348).

Stone, M. (1998). Abductive planning with sensing. In *Proceedings of the International Conference of the AAAI*, pages 631–636.

Veltman, F. (1996). Defaults in update semantics. *Journal of philosophical logic*, **25**, 221–261.