THE UNIVERSITY *of* EDINBURGH

Edinburgh Research Explorer

Knowledge, Causality and Temporal Representation

Citation for published version:

Lascarides, A 2009, 'Knowledge, Causality and Temporal Representation', *Linguistics: An Interdisciplinary Journal of the Language Sciences*, vol. 30, no. 5, pp. 941-973. <https://doi.org/10.1515/ling.1992.30.5.941>

Digital Object Identifier (DOI):

[10.1515/ling.1992.30.5.941](https://doi.org/10.1515/ling.1992.30.5.941)

Link:

[Link to publication record in Edinburgh Research Explorer](#)

Document Version:

Publisher's PDF, also known as Version of record

Published In:

Linguistics: An Interdisciplinary Journal of the Language Sciences

General rights

Copyright for the publications made accessible via the Edinburgh Research Explorer is retained by the author(s) and / or other copyright owners and it is a condition of accessing these publications that users recognise and abide by the legal requirements associated with these rights.

Take down policy

The University of Edinburgh has made every reasonable effort to ensure that Edinburgh Research Explorer content complies with UK legislation. If you believe that the public display of this file breaches copyright please contact openaccess@ed.ac.uk providing details, and we will remove access to the work immediately and investigate your claim.

Knowledge, causality, and temporal representation*

ALEX LASCARIDES

Abstract

In this paper, a formal semantic account of the simple past tense in text is offered. The contributions to the interpretation of text made by the text's syntactic structure, semantic content, aspectual classification, world knowledge of the causal relations between events, and Gricean pragmatic maxims are all represented within a single logical framework. This feature of the theory gives rise to solutions to several puzzles concerning the relation between the descriptive order of events in text and their temporal relations in interpretation.

1. The problem

If John hits Max, causing Max to turn round (to face John), then text (1) reflects this while (2) distorts it:

- (1) John hit Max. Max turned round.
- (2) Max turned round. John hit him.

At least, (2) distorts it in the "null" context in which I've represented it.

So the order in which such clauses appear is crucial. It gives rise to an *ordering question*: given a particular order in which events are described, what are the constraints in interpretation on their temporal and causal relations in the world?

The above seems to show that the ordering constraints are satisfied if the textual order of simple past tensed clauses reflects the temporal order of the events they describe. But this constraint isn't necessary. The second event described in (3) doesn't follow the first event, and in (4) the second event PRECEDES the first.¹

- (3) The council built the bridge. The architect drew up the plans.
- (4) Max fell. John pushed him.

Furthermore, the event and state described in text (5) temporally OVERLAP, whereas in (6) they do not.

- (5) Max opened the door. The room was pitch dark.
- (6) Max switched off the light. The room was pitch dark.

Our aim is to explain this: what forms the basis of the different interpretations of (5) and (6), and why can't the events of hitting and turning round be described in the reverse to their causal order using just the simple past tense, while falling and pushing can?

Existing formal treatments of tense, such as Kamp and Rohrer (1983), Partee (1984), Hinrichs (1986), and Dowty (1986) don't account for the backward movement of time in (3) and (4), since their theories are concerned only with narrative texts such as (1), (2), (5), and (6). Webber (1988) can account for the backward movement of time in (3), but her theory is unable to predict that mismatching the descriptive order of events and their temporal order is allowed in some cases (such as [3] and [4]), but not in others (such as [2]). Our aim is to extend the coverage of these existing formal treatments, so that an account of the distinct interpretations of (2) and (4) is provided in a declarative framework. In particular, I will extend the discourse representation theory (DRT) account of tense by deploying a strategy proposed in Partee (1984).

The basis for distinguishing (2) and (4) will involve encoding causal knowledge about the relation that typically holds between the events described. In this respect, my approach refines that outlined by Hobbs (1979, 1985) and Dahlgren (1988). Hobbs encodes causal knowledge in a declarative framework and this guides discourse interpretation. Dahlgren characterizes the causal knowledge in terms of an episodic knowledge representation scheme. One difference from Hobbs is that in addition to causal knowledge, we place emphasis on temporal information conveyed by the TEXTUAL order of events. And one difference from Dahlgren is that we place more emphasis on the need to declaratively specify the relations between the representations. The main difference from both approaches, in fact, lies in the utilization of the underlying relation of logical consequence, which yields the interactions required among the various knowledge resources.

In Hobbs's and Dahlgren's theories as they stand, it is not completely clear that the requisite notion of logical consequence could be defined, since there are no obvious relations between defeasible laws that ought to interact in certain specific ways. Consequently, conflicts that arise among the knowledge sources their theories recruit are not resolved in a systematic way. For example, Hobbs (1979) uses causal knowledge to

choose the antecedent to the pronoun in text (7). That knowledge is still relevant for (8), but further conflicting knowledge is assumed to override it. There's no explanation of HOW it overrides it.²

(7) John can open Bill's safe. He knows the combination.

(8) John can open Bill's safe. He must change the combination.

Dahlgren (1988) represents linguistic knowledge concerning the relation between textual order and temporal order. This knowledge explains the difference between (1) and (2); but in the analysis of (4), it is overridden by conflicting causal knowledge. Again, the resolution of conflict lacks principled justification. In both theories, the representation of knowledge is unconstrained by the need to comply with inferences supported by a notion of logical consequences. As a result, particular resolutions of conflict, like those that arise in the above examples, appear arbitrary. But I assume that resolution of conflict among knowledge sources must have logical justification. So causal knowledge and linguistic knowledge will be placed in a logical context where the implications can be precisely calculated. The representation of such knowledge is therefore constrained in that the underlying logic must enable the representations of the various knowledge sources to interact in the appropriate way.

2. Tense in discourse representation theory

The current theories on tense in text expressed in discourse representation theory (DRT) (Kamp and Rohrer 1983; Partee 1984; Hinrichs 1986) address the ordering question in the domain of narrative texts such as (1), (2), (5), and (6). It's instructive to examine their analyses.

In DRT the representation of text is called a *discourse representation structure* (DRS). A DRS is a pair of sets, where the first set constitutes the DISCOURSE ENTITIES that are introduced by the discourse, and the second set constitutes conditions on those discourse entities.³ The DRS representing the entire discourse is built by DRS construction rules, and only after this is achieved does any semantic interpretation occur. The DRS construction rules process the sentences of text one at a time, taking into account the syntactic structure of the sentence currently being processed to extend the representation of text built up so far.

The temporal relations between the eventualities⁴ are encoded in the DRS by the DRS construction rules. This is achieved with essentially two innovations. First, tense is characterized anaphorically as a tripartite relation between speech time, eventuality time, and reference time. The

reference time is the anaphor (see Reichenbach 1947). It is evaluated with respect to context, in that the rules for constructing DRSs relate the reference time of the clause being processed to the reference times for previous clauses. The second innovation is to make DRS construction sensitive to the aspectual classification of the clauses concerned. These properties enable one to capture the apparent forward movement of time in texts (1) and (2), and the difference between (1) and (5).

More specifically, the rules in Partee (1984) adhere to the following. Simply put, event clauses move the reference time forward; stative clauses leave the reference time where it is:

- a. If the clause α_n is in the simple past, then the current reference time r_n precedes the speech time *now*.
- b. If α_n is an event clause, then its eventuality time e_n is contained in the current reference time r_n , and the clause updates the reference time to r_{n+1} , where $r_n < r_{n+1}$, in preparation for analyzing the next clause of the text, should there be one.
- c. If α_n is a stative clause, then the eventuality time e_n contains the current reference time r_n . The reference time is not updated.

The DRS construction rules in Kamp and Rohrer (1983) and Hinrichs (1986) are similar. The comments made here about Partee's analysis also apply to Kamp and Rohrer's and Hinrichs's theories of tense.

In simplified form, the DRSs for (2), (5), and (6) are given below; I have ignored the discourse entities referred to by NPs and nominal anaphora resolution, since these are not our concerns here.

- (2) Max turned round. John hit him.
- (2') $[e_1, e_2, r_1, r_2, \text{now}][\text{turnaround}(\text{max}, e_1), e_1 \subseteq r_1, r_1 < \text{now}, r_1 < r_2, \text{hit}(\text{john}, \text{max}, e_2), e_2 \subseteq r_2, r_2 < \text{now}]$
- (5) Max opened the door. The room was pitch dark.
- (5') $[e_1, e_2, r_1, \text{now}][\text{open}(\text{max}, \text{door}, e_1), e_1 \subseteq r_1, r_1 < \text{now}, r_1 < r_2, \text{dark}(\text{room}, e_2), r_2 \subseteq e_2, r_2 < \text{now}]$
- (6) Max switched off the light. The room was pitch dark.
- (6') $[e_1, e_2, r_1, \text{now}][\text{switchoff}(\text{max}, \text{light}, e_1), e_1 \subseteq r_1, r_1 < \text{now}, r_1 < r_2, \text{dark}(\text{room}, e_2), r_2 \subseteq e_2, r_2 < \text{now}]$

(2') entails that the turning round precedes the hitting as required. But although the preferred readings of (5) and (6) are compatible with the truth conditions of (5') and (6'), no explanation is offered of why (5)'s preferred reading is distinct from (6)'s. Nor do the above DRS construction rules apply to nonnarrative texts like (4). The aim of this paper is to extend the DRT account to a more general theory of temporal structure, thereby solving the above puzzles concerning texts (1) to (6).

2.1. *Extending the account*

Since the DRS is built by construction rules that take into account only syntactic structure, all significant temporal relations must be read off the syntax by these construction rules. It is for this reason that DRT adopts a syntactic-based approach to aspectual classification. That is, whether a clause describes an event or state is determined solely by syntactic considerations. Dowty (1986) argues against this view. Partee (1984: 281) acknowledges that Dowty's criticisms pose "serious problems" for the analysis. We consider here how the ordering question poses problems for the syntactic approach to DRS construction. Can extending the theory to deal with texts like (4) preserve the relation between logical form and syntax that currently holds in DRT?

We have already seen that (5) and (6) have distinct preferred interpretations in spite of the similar syntax of the sentences. Partee preserves a close relation between logical form and syntax in this case by keeping the temporal relations in the logical forms of (5) and (6) "vague"; the temporal relations in the DRS don't preclude overlap or precedence. Presumably, in order to distinguish (5) and (6) one would have to calculate which of the several temporal relations that are compatible with those described in logical form is the preferred one given the context, ideally without sacrificing formality at this stage. Intuitively, the meanings of entering a room, switching off the light, and darkness would be used to calculate the preferred readings. Very roughly, world knowledge dictates that, in the absence of information to the contrary, that the switching off the light caused the darkness is the most plausible relation between the eventualities mentioned in (6). This knowledge is lacking for the events in (5).

A similar strategy could conceivably be used for extending the theory to cover texts like (4). Namely, revise the logical form of text (2) so that the temporal relation is vague, and compatible with the preferred readings of (2) and (4). Then a syntactic-based approach to DRS construction as currently deployed in DRT can be used to build the logical forms of (2) and (4). A formal account of how world knowledge and the context affect which temporal relation holds in the preferred reading would then have to be provided, presumably along similar lines to an account of the distinct readings of (5) and (6).

It's worth examining this strategy for representing texts (2) and (4) since this approach is already used in DRT to avoid problems with (5) and (6). Keeping as close as possible to the current logical form Partee provides for (2), its revised logical form would be (2'), or some logically equivalent formula.⁵

- (2') $[e_1, e_2, r_1, r_2, now][turnround(max, e_1), e_1 \subseteq r_1, r_1 < now, hit(john, max, e_2), e_2 \subseteq r_2, r_2 < now, r_1 < r_2 \vee r_2 < r_1]$

This is because the first $r_1 < r_2$ corresponds to the preferred reading of (2), and the second disjunct $r_2 < r_1$ corresponds to the preferred reading of (4), which has similar sentential syntax to (2).

Achieving (2') as the logical form of (2) would mean revising the way reference times get updated in the DRS construction rules: (b) must be replaced by (b').

- (b') If α_n is an event clause, then its eventuality time e_n is contained in the current reference time r_n , and the clause updates the reference time to r_{n+1} , where $r_n < r_{n+1} \vee r_{n+1} < r_n$, in preparation for analyzing the next clause of the text, should there be one.

But by using (b'), the logical form of (1) is logically equivalent to (2').

- (1) John hit Max. Max turned round.
 (2) Max turned round. John hit him.

One would now have to provide an account of why (1) and (2) have distinct natural interpretations in spite of having the same logical forms. Presumably, textual order is a crucial piece of information that must play a role in distinguishing (1) and (2), for textual order is the only difference between them. But the textual order is lost in the new representation of (1) and (2), since the order of construction is lost in (2'). Note that textual order is not lost in Partee's original analysis of (2), because the temporal progression of reference times reflected textual order. Because we are attempting to provide a representation of (3) and (4), however, while preserving the relationship between logical form and syntax, the temporal progression among reference times no longer reflects textual order. And it is unclear how one could represent rules about how a text's preferred reading is affected by textual order, if the textual order is not retrievable from the text's representation.

The above discussion indicates that simple changes to Partee's DRS construction rules will not yield the whole story concerning the interpretation of the above texts. Assigning "vague" temporal relations in the analysis of (5) and (6) may avoid problems, but this strategy cannot be applied unproblematically to provide an analysis of both (2) and (4). This in a sense isn't surprising, since there is intuitively a fundamental difference between (2) and (4): unless context indicates otherwise (2) is narrative and (4) is not.

Partee (1984: 260) observes the problems imposed by text like (4) for her analysis, and although she doesn't provide a solution, she suggests a

strategy for solving it. She proposes that the way forward is to introduce a HIGHER-ORDER PARAMETER of discourse structure; a parameter that can have different values for different subparts of the discourse. She suggests that the value of this parameter determines the ordering of reference times and is in turn influenced by an elaborate theory of the interaction among semantic content, context, inference, implicature, and causal knowledge. The DRS construction rules will in the initial stage impose a FREE relation between the reference times because no value will be assigned to the parameter at this stage. The setting of specific values of the parameters will yield specific temporal relations and will be the result of subsequent stages in processing, which will involve integrating the information in the DRS so far with background knowledge and contextual information (Partee 1984: 283).

This extension that Partee suggests involves a radical change to the way DRSs are constructed. No longer does semantic interpretation occur only when the DRS for the whole text is constructed. Rather, assessing the values of the parameters would have to be influenced by the semantic interpretation of the DRS built so far, and this information would have to interact in certain specific ways with background knowledge. The change to the way DRSs are constructed forms a large part of the theory outlined here, and so in a sense we aim to explore in detail Partee's suggestion that the way DRSs are constructed should be affected by semantic content and background knowledge, as well as syntax.⁶

The general approach I adopt will be the following: the logical form of the text is built up through the processing of the successive clauses. The problem of constructing the logical form of text thus becomes a problem of updating the logical form of the text processed so far with the clause currently being processed. Suppose Δ is the logical form of the first n clauses of the text, and β is the DRS representing the semantic content of the $n + 1^{\text{th}}$ clause. Then we must provide rules that, given Δ and β , produce a new DRS representing the $n + 1$ -clause text. The updated DRS will be the value of a partial function $\langle \Delta, \alpha, \beta \rangle$, which informally should be thought of as updating Δ with β via a relation between the eventualities described by β and α , where α is a DRS that's already part of Δ .

To understand the import of defining the update function, I now present the syntax and semantics of the DRT language I use.

The language. The language consists of first-order expressions, together with discourse entities sorted into terms denoting individuals, times, and eventualities. The language is defined inductively in the usual way:

- All first-order well-formed formulae (WFF) are WFF.
- If $x_1, \dots, x_n$ are discourse entities and ϕ is an n -plane predicate, then

$\phi(x_1, \dots, x_n)$ is an (atomic) condition on $x_1, \dots, x_n$. If x_1 and x_2 are eventuality or time discourse entities, then $x_1 < x_2$ is an atomic condition on x_1 and x_2 .

- If $x_1, \dots, x_n$ are discourse entities and $Con(x_1, \dots, x_n)$ are conditions on $x_1, \dots, x_n$, then $[x_1, \dots, x_n][Con(x_1, \dots, x_n)]$ is a WFF. These WFFs are called DRSs.

- If K_1 and K_2 are DRSs, then $K_1 \rightarrow K_2$, $K_1 \wedge K_2$, and $\neg K_1$ are DRSs.
- If α is a DRS then $me(\alpha)$ is a term (referring to an eventuality).
- If Δ , α , and β are DRSs, then $\langle \Delta, \alpha, \beta \rangle$ is a partial function whose value is a DRS.

The truth definitions. The interpretations of the well-formed expressions of the language are defined with respect to a first-order valuation v and an embedding function f that provides suitable assignments to the discourse entities. We use $\llbracket \phi \rrbracket_{\langle v, f \rangle}$ to be the interpretation of the well-formed expression ϕ with respect to the valuation v and embedding function f . And we use $\models_f \phi$ to mean that the WFF ϕ is true with respect to the valuation v and embedding function f .

- Where β is a first-order expression, $\llbracket \beta \rrbracket_{\langle v, f \rangle}$ is defined in the usual first-order way.
- Where β is an atomic condition $\phi(x_1, \dots, x_n)$ on discourse entities $x_1, \dots, x_n$, $v \models_f \phi(x_1, \dots, x_n)$ iff $\langle f(x_1), \dots, f(x_n) \rangle \in \llbracket \phi \rrbracket_{\langle v, f \rangle}$. (Similarly for $e_1 < e_2$.)
- If β is a WFF of the form $[x_1, \dots, x_n][Con(x_1, \dots, x_n)]$, $v \models_f \beta$ iff there is a proper embedding g of f such that $v \models_g Con(x_1, \dots, x_n)$. (Proper embeddings are defined in the traditional DRT way.)
- If β is a WFF of the form $K_1 \vee K_2$ where K_1 and K_2 are DRSs, $v \models_f \beta$ iff there is a proper embedding g of f such that $v \models_g K_1$ or $v \models_g K_2$.
- If β is a WFF of the form $K_1 \rightarrow K_2$ where K_1 and K_2 are DRSs, $v \models_f \beta$ iff there is a proper embedding g of f such that if $v \models_g K_1$ then there is a proper embedding h of g such that $v \models_h K_2$.
- If β is a WFF of the form $K_1 \wedge K_2$ where K_1 and K_2 are DRSs, $v \models_f \beta$ iff there is a proper embedding g of f such that $v \models_g K_1$ and $v \models_g K_2$.
- If β is a WFF of the form $\neg K_1$ where K_1 is a DRS, $v \models_f \beta$ iff there is no proper embedding g of f such that $v \models_g K_1$.

2.2. Defining the updating function

Let $\Delta = [U_\Delta][Con_\Delta]$, $\alpha = [U_\alpha][Con_\alpha]$ and $\beta = [U_\beta][Con_\beta]$. Then $\langle \Delta, \alpha, \beta \rangle$ is defined only if $U_\alpha \subseteq U_\Delta$, $Con_\alpha \subseteq Con_\Delta$, and $Con_\beta \not\subseteq Con_\Delta$. That is, α must

have already been added to the DRS for the text so far, and β must contain new information. So given that clauses are processed successively, the function is defined only if the clause represented by α appeared in the text before that represented by β . In this way, the update function keeps track of the textual order of the clauses. And if we reason about the value of this function directly in the logical language, then we are able to represent the effect of textual order on the temporal structure of text. This property of $\langle \Delta, \alpha, \beta \rangle$ will play an important role in distinguishing (1) from (2).

We also assume that if $\langle \Delta, \alpha, \beta \rangle$ is defined, then $\langle \Delta, \alpha, \beta \rangle$ is $[U_\Delta, U_\beta][Con_\Delta, Con_\beta, r(me(\alpha), me(\beta))]$, where $me(\alpha)$ and $me(\beta)$ are respectively the eventualities described by α and β , and r is some relation on these eventualities. In words, extending Δ with β involves adding β 's discourse entities and conditions to those in Δ and adding a relation between β 's eventuality and α 's where α is part of Δ already.

The DRS construction rules must determine the value of the event relation r for various Δ , α , and β . To this end, we provide rules in the logical language that establish how the function $\langle \Delta, \alpha, \beta \rangle$ is calculated. These rules are influenced by the semantic content of Δ , α , and β , unlike the traditional DRT approach for building logical form. Moreover, the rules for updating are also influenced by context and by what the reader knows about the typical relations between fall and push, say. The context sensitivity of DRS construction will be captured by the nonmonotonic nature of the underlying inference regime. I will show that through making obvious assumptions about world knowledge and language use, the representations of (2) and (4) are distinct. Essentially, the function $\langle \Delta, \alpha, \beta \rangle$ will play the role of Partee's discourse parameter, its value determines the temporal relations, and calculating its value depends on background knowledge.

Partee (1984) shows for expository purposes how the logical form of (2) is constructed. But this construction is not part of the language or the logic. It is carried out at a metalinguistic level. Essentially, Partee uses parenthesized numbers in the DRS to show the order of construction steps, but these numbers aren't part of the logical form; the information they convey is "lost" in the final representation. The extension of DRT in this paper reasons about the way the text's representation is constructed directly in the language; we provide rules in the language about the value of $\langle \Delta, \alpha, \beta \rangle$ for various Δ , α , and β . So in a general sense the approach adopted here spells out in the language what Partee was showing for expository purposes alone. In doing this, it is possible to allow textual order to interact with world knowledge of causal relations in a single logical framework.

2.3. *Caveats*

Before I proceed any further, I should note some provisos concerning the intended coverage of this paper. First, I am concerned only with temporal structure and not discourse structure. Clearly, calculating temporal structure isn't sufficient for calculating discourse structure. There will be many cases where a pair of sentences is provided with an analysis describing the temporal relations between the events they describe, but the text is intuitively incoherent and the current account cannot explain what is wrong with them (see Caenepeel 1991 for a discussion of such cases).

Nevertheless Lascarides and Asher (1991a) show how the account presented here of why (2) and (4) have distinct temporal structures can contribute to an account of why they have distinct discourse structures. So the theory presented here can be viewed as forming the foundations for an explanation of why (2) and (4) have distinct discourse structures.

Furthermore, the constraints on the relation between eventualities described in NONconsecutive clauses requires some notion of hierarchical discourse structure, as discussed in Polanyi (1985), Thompson and Mann (1987), and Grosz and Sidner (1986). Therefore, I also restrict my concerns to the temporal relations derived from two consecutive sentences in a text; that is, LOCAL temporal structure.

Partee (1984) also restricts coverage to local temporal structure, for her rules encode temporal relations only between the reference times of the current clause and the next clause. This is sufficient for characterizing global temporal structures of narrative text; time progressively "moves forward" as the text unfolds, and so the global temporal structure follows from the local temporal structures:

- (9) Max arrived at the house. Mary ran into the driveway. She greeted him with a huge smile. They walked into the house arm in arm.

The correct global temporal structure also follows from the local temporal structures in the following nonnarrative text, since time "moves backward" from one clause to the next:

- (10) Max died. John poisoned him. He put arsenic in Max's wine.

But an account of local temporal structure will not tell one how or when discourse popping occurs, and so it isn't sufficient for calculating global temporal structure. For example in text (11), the event in (11e) must be related to that in (11b) and not (11d), and the theory presented here isn't sufficient to explain why.

- (11) a. Guy experienced a lovely evening last night.
- b. He had a fantastic meal.
- c. He ate salmon.
- d. He devoured lots of cheese.
- e. He won a dancing competition.

Nevertheless, Lascarides and Asher (1991a) argue that calculating local temporal structure is necessary if one is to build discourse structure, which in turn is required to calculate global temporal structure. They extend the treatment of local temporal relations explored here to construct discourse structure and thereby present an explanation of why (11e) is related to (11b) rather than (11d).

A further caveat is that there is a feeling of artificiality upon reading the above texts in isolation of any real linguistic context. In my favor, however, I would argue that one of the major advantages of the account presented here is that it provides a rather precise explanation of how interpretation can be changed when context is enriched. More specifically, the framework is rich enough to explain that the preferred reading of (12) is one where the pushing caused the falling, but in (13) the preferred reading is one where the falling PRECEDED the pushing:

- (12) Max had a horrible accident yesterday. *He fell. John pushed him.*
- (13) John and Max were at the edge of the cliff. Max felt a sharp blow on the back of his neck. *He fell. John pushed him.* Max rolled over the edge of the cliff.

Changes in interpretation due to changes in context are captured through nonmonotonic inference. The discourse context effects encapsulated in (12) and (13) are examined in detail in Lascarides et al. (1992).

Finally, we consider only simple past tensed texts in this paper, and so it may seem that the rich source of natural language expressions available for encoding temporal relations has been marginalized. For instance, the temporal structure underlying (14) is the same as that of (4), but the interpretation of (15) is different from (2):

- (4) Max fell. John pushed him.
- (14) Max fell. John had pushed him.
- (2) Max turned round. John hit him.
- (15) Max turned round. John had hit him.

Similarly, the temporal structures underlying (16) and (6) are the same, but those underlying (17) and (5) are different.

- (6) Max switched off the light. The room was pitch dark.
- (16) Max switched off the light. The room was pitch dark thereafter.

- (5) Max opened the door. The room was pitch dark.
 (17) Max opened the door. The room was pitch dark thereafter.

We fail in this paper to offer a detailed account of why the pluperfect and adverbials like *after* have these discourse effects. In my favor, however, I would argue that one of the virtues of the theory is that semantic entailments, which are characterized by indefeasible inference, and pragmatic implicatures, which are characterized by defeasible inference, are both represented in this theory in a single framework. This provides a solid foundation on which to build representations of BOTH the semantics and the pragmatics of not just the simple past, but a variety of temporal expressions. Details of the analysis of the pluperfect in a framework similar to the one presented here are provided in Lascarides and Asher (1991b). The semantics and pragmatics of the connective *after* form the focus of research currently being undertaken jointly with Jon Oberlander and Nicholas Asher.

3. The semantics of sentences

In line with Partee's suggested strategy for extending the DRT account of tense, there are two stages to constructing the logical form of text. The first stage assigns "free" temporal relations between the eventualities described in a text, and the second stage uses background knowledge to refine the logical form constructed so far. In our theory, constructing the DRSs that represent the semantics of the sentences in a text essentially corresponds to the first stage. No temporal relations between eventualities are imposed at this level, and the logical form of sentences is calculated using syntactic information alone, in line with traditional DRT. The second stage will involve calculating the value of the function $\langle \Delta, \alpha, \beta \rangle$, for various values of Δ , α , and β . As we've mentioned, this function will play the role of Partee's suggested discourse parameter, and the value of this function will be influenced by background knowledge. This is probably not the only way of pursuing Partee's suggestions; it's just one way.

We examine now the first stage of DRS construction. The sentences in (2) are represented by the DRSs α and β respectively:⁷

- (2) Max turned round. John hit him.
 (α) $[e_1][turnround(max, e_1), e_1 < now]$
 (β) $[e_2][hit(john, max, e_2), e_2 < now]$

In words, α introduces a discourse entity e_1 , where e_1 is the event of Max turning round that occurs earlier than the time of speech *now* (which for

simplicity's sake is treated as a constant here). β is similar, save that e_2 is John hitting Max.

The logical forms of the sentences in (4) will be exactly the same as α and β above, save that the predicates *turnround* and *hit* are replaced with *fall* and *push*. This reflects the similar syntax of the sentences in (2) and those in (4).

Contrary to Partee's (1984) analysis of tense, the relation between e_1 and e_2 is not encoded into the representation of the tenses. Indeed, there are no reference times at all. Tense is still deictic, because the event is a discourse entity and hence determined deictically. So it's still possible to explain why sentence (18) refers to a particular event occurring at a particular time.

(18) I failed to turn off the stove.

Partee argues that reference times are essential for modeling the effect of tense on the temporal structure in narrative discourse. I'll show that reference times aren't needed to model this in the case of the simple past, at least. Nor are they needed to model the role of the simple past in nonnarrative text. In this paper, I leave open the issue of whether reference times are needed to model any NL temporal expressions apart from the simple past. But see Lascarides and Asher (1991b) for arguments that one can abandon reference times even when one broadens the coverage of data to include the pluperfect.

In traditional DRT, a reference time is introduced into the DRS for each temporal expression featured in the NL sentence, including the tense. In contrast, a corollary of our approach is that reference times are introduced in logical form only if the natural language sentence features a relational temporal adverbial, such as *afterwards*, *two minutes later*, *the week before*. Because of this, it should be possible for reference times to capture generalizations between temporal and nontemporal adverbials; in particular the common properties of relations invoked by temporal and nontemporal adverbs will be reflected in the behavior of reference times.

Partee's representation of sentences differs from ours in at least two further respects. First, she interprets the event term e as naming both the culmination and the preparatory phase that leads to the culmination. Lascarides (1988) shows that problems arise from this, which are solved if e refers to the CULMINATION ALONE, and another term, say $PR(e)$ where PR is a function from terms to terms, refers to the preparatory phase (this representation of preparatory phases is explored in Lascarides 1991). So according to our interpretation, e is a punctual entity.

Second, and more significantly, the interpretations of α and β must be

modal. For world knowledge and linguistic knowledge, which will interact with the semantic content of α and β via the update function, will be represented in modal terms. All constants and discourse entities will be treated as RIGID DESIGNATORS, so as to avoid a lot of the traditional puzzles of identifying objects in different modal contexts (see Kripke 1972). Details of how the semantics is extended to a modal interpretation are given in section 5.

We will use *hold*(e, t) to mean that e holds at the point of time t . The way e behaves with respect to *hold* determines whether e is an event or a state. In the case of an event e , *hold*(e, t) entails that e 's culmination holds at the point of time t , and e holds nowhere else in that modal context. That is, events hold on at most a unique point of time in any given modal context (although an event may hold at distinct times in different modal contexts). Because of this, events are best thought of as event tokens rather than event types. Events are tokens, for if they occur in a particular modal context at all, then they do so only once.

The formalism also reflects the intuition that, in contrast to the punctual structure of culminations, STATES occur over an extended period of time with no definite endpoints (see Moens and Steedman 1988). If s is a state, then all the points t for which *hold*(s, t) is true relative to a particular modal context form either the empty set (in which case s doesn't hold in that modal context), or an open interval. Whether or not a clause describes an event or a state is thus determined by the semantic interaction between e and *hold*, where e is the discourse entity of the clause. Thus contrary to traditional DRT, aspectual classification is SEMANTICALLY determined. I'll use the formula *state*(e) and *event*(e) to mean respective " e is an event" and " e is a state".

4. Temporal relations and defeasible reasoning

We have so far given the representations of isolated sentences. These must be used to update the representation of text so far. So we now turn to the second stage in building the DRSSs; calculating the update function $\langle \Delta, \alpha, \beta \rangle$.

We first concentrate on (2) and (4). What forms the basis for their distinct temporal structures? As I've mentioned, one possibility is that the difference in interpretation is derived from the difference in the relations that typically hold between the events being described. More specifically, Lascarides and Oberlander (1992) argue that the temporal coherence of text should be characterized in terms of the following definition of when events are connected:

- *Event connection*

Two events e_1 and e_2 are connected somehow if: e_1 causes e_2 or vice versa; or e_1 is part of the preparatory phase of e_2 or vice versa; or e_1 is part of the consequent phase of e_2 or vice versa; or e_1 and e_2 temporally overlap.

Using *event connection* above, one can state a piece of world knowledge (WK), gained from perception and experience, that relates falling and pushing.⁸

- *Causal law 4*

If the events e_1 and e_2 are connected somehow where e_1 is x falling and e_2 is y pushing x , then unless there's information to the contrary, e_2 caused e_1 .

There is no similar law for hitting and turning round; if there is a connection (but we don't know which connection) between hitting and turning round, then we are unable to conclude, even tentatively, exactly how the events are related.

The claim that our perception and experience of the world yield a causal law for falling and pushing but none for turning round and hitting is an empirical claim requiring justification, presumably from psychological experiments testing commonsense reasoning about causation. Such experiments are beyond the scope of this paper, however. For now, suffice to say that it is intuitively plausible that our knowledge about falling, pushing, turning round, and hitting is compatible with what I have stated.

The force of the phrase *unless there's information to the contrary* in causal law 4 means that this is a defeasible law. If it forms the basis for distinguishing between (2) and (4), then defeasible reasoning must underly discourse interpretation. In particular, the rules for DRS construction must be guided by defeasible reasoning.

We will shortly explore how causal law 4 could be used in building the representation of (4). But we first turn our attention to the interpretation of (2). Since there is no WK like causal law 4 that's available to guide (2)'s interpretation, how do we infer that it is iconic?

In addition to rules like causal law 4, the reader has defeasible LINGUISTIC knowledge (LK). If there is no temporal information at all gained from context, genre, WK, or syntactic markers (apart from the simple past tense, which is the only temporal "expression" we consider here), then the descriptive order of events provides the only vital clue as to their temporal order. In such cases, the descriptive order of events is typically assumed to match their temporal order. That is, if e_1 is described before e_2 , then unless there's information to the contrary, e_1 holds at a point

of time t_1 , e_2 holds at t_2 , and $t_1 < t_2$. There are two motivations for this defeasible LK: the first concerns the order of perception of events, and the second concerns one of Grice's pragmatic maxims.

As regards the first, a plausible reason for proposing this LK is that it corresponds — for some text genres at least — to the Dowtian protagonist's "order of discovery" (see Dowty 1986). In such genres, a narrator typically describes events in the order in which the protagonist or narrator views them: this is predicted by the above LK. For states, which occur over extended periods, the protagonist may well not be aware of the state until after it has started to occur. In these situations, the narrator typically indicates the protagonist's inability to view the state at the point when it started by introducing it in the text only at the point where the protagonist perceives it. Text (5) is an example of this:

(5) Max opened the door. The room was pitch dark.

Note that the piece of defeasible LK we're proposing does not exclude the possibility that a state described after an event actually started to occur before the event did; it would merely describe the state as still holding sometime after the event occurred. But it must be emphasized that there are exceptions to this LK; (4) is one of them.

The second motivation is more general. As I've mentioned, it will sometimes be the case that the ONLY information available to an interpreter regarding temporal structure will be textual order. In such cases, the above LK will be the only thing available. In practice such cases may indeed be rare, since the interpreter will have information available about discourse structure (see Lascarides and Asher 1991a).⁹ In any case, this LK observes Grice's (1975) maxim of manner, in which it is suggested that text should be orderly. One way of interpreting this is that events should be described in the right order. In essence, the story about the above LK suggests that the theory should represent Grice's pragmatic maxims as defeasible rules. Such an approach to pragmatics has been suggested in, for instance, Joshi et al. (1984).

The above LK is about the temporal information conveyed by textual order. Textual order is reflected in $\langle \Delta, \alpha, \beta \rangle$ because this is defined only if the clause represented by α appears in the text before that represented by β . So we represent the above LK in terms of $\langle \Delta, \alpha, \beta \rangle$, as a default law:

- *Narration*

If $\langle \Delta, \alpha, \beta \rangle$ is true, then unless there is information to the contrary, there are times t_1 and t_2 such that $me(\alpha)$ holds at t_1 , $me(\beta)$ holds at t_2 , and $t_1 < t_2$.

In words, *narration* states that if Δ is to be updated with β via a relation with α 's eventuality, then α 's eventuality holds at a time before β 's holds, unless there's information to the contrary.

Narration is crucially a rule that encodes reasoning about DRS construction. More generally, the theory proposed here reasons about the way the logical form of text is constructed directly in the language via default rules about the value of the function $\langle \Delta, \alpha, \beta \rangle$.

Narration is similar to Dowty's (1986) temporal discourse interpretation principle (TDIP), but there are two important differences. First, *narration* is expressed purely in terms of event times rather than reference times. And second, the applicability of *narration* is dependent on the reader's WK. Assuming that the reader has INDEFEASIBLE knowledge that poisoning precedes death, the temporal structure given in *narration* won't be inferred for (19) because there's information to the contrary.

(19) Max died. John poisoned him.

Dowty's TDIP is expressed as a principle that is independent of WK, because he's not concerned with the difference in the natural interpretations of (2) and (19). Such differences should be captured in a general theory of text, however, and so *narration* will be formally defined so that its consequent isn't inferred for texts like (19).

Finally, as mentioned earlier, sentences featuring temporal connectives such as *afterward* would yield logical forms that contain reference times in this account. The relationships between event times encoded by *narration* and the relations between reference times encoded in logical form will interact in nontrivial ways. If the relation between reference times is incompatible with the temporal structure predicted by *narration*, then *narration*'s consequent won't hold since the logical form provides the appropriate *information to the contrary*. This would explain why the temporal structure for (20) is different from (2)'s: the relation between reference times encoded in the logical representation of *before* would conflict with the default conclusion provided in *narration* and so the default conclusion would not be inferred (details are given in Lascarides et al. n.d.):

(2) Max turned round. John hit him.

(20) Before Max turned round, John hit him.

The default rules representing LK and WK that will play a role in constructing the DRSs must be represented in a logic that can model suitable patterns of defeasible inference. We will represent the above rules in *hierarchic autoepistemic logic* (or HAEL), as developed in Konolige (1988) and used in Appelt and Konolige (1988) to model speech acts.

I use HAEL because it is one of the default logics that support the intuitively compelling patterns of defeasible reasoning I use to build the DRSs.

5. A brief introduction to HAEL

HAEL (Konolige 1988) is a logic for defeasible reasoning that can model, among other things, the following two patterns of inference.

- *Defeasible modus ponens*
for example, Birds typically fly, Tweety is a bird; so Tweety flies.
But not: Birds typically fly, Tweety is a bird, Tweety doesn't fly; so Tweety flies.
- *The penguin principle*
Penguins are birds, Birds typically fly, Penguins typically don't fly, Tweety is a penguin; so Tweety doesn't fly.

This feature of HAEL will be exploited for constructing DRSs for NL text.

HAEL is derived from autoepistemic logic (Moore 1985). An autoepistemic logic is a first-order language augmented by a model operator L , where iteration of L is precluded. L is interpreted intuitively as self-belief. Its interpretation is defined with respect to a belief set Γ , where Γ is the modal index:

- If v is a first-order valuation, then
 $\langle v, \Gamma \rangle \models L\phi$ iff $\phi \in \Gamma$

An autoepistemic base set A is a set of formulae and is intuitively a statement of the agent's partial knowledge. A STABLE EXPANSION of an autoepistemic base set A is a set of formulas T satisfying the following conditions:

1. T contains all the sentences of the base theory A .
2. T is closed under first-order consequence.
3. If $\phi \in T$, then $L\phi \in T$
4. If $\phi \notin T$, then $\neg L\phi \in T$

Default laws are represented as $L\phi \wedge \neg L\neg\psi \rightarrow \psi$. The base set A defeasibly entails ψ if ψ holds in all the stable expansions of A . The above definition of stable expansions is sufficient to support defeasible modus ponens. More formally, if A is the following (think of $B(x)$ as “ x is a bird” and $F(x)$ is “ x flies”):

$$(A) \quad \{LB(x) \wedge \neg L\neg F(x) \rightarrow F(x), B(t)\}$$

Then all stable expressions of A contain $F(t)$. Konolige (1987) demonstrates certain equivalence results between autoepistemic logic and Reiter's (1980) default logic.

Hierarchic autoepistemic logic (HAEL) is an extension of autoepistemic logic motivated in part by the need to validate the "penguin principle." Autoepistemic logic lacks the facility to state preferences among various conflicting knowledge sources; HAEI introduces this facility. Certain expansions are preferred on the basis that the defaults used in them have a higher priority in some well-defined sense than the ones used in the alternative expansions.

In HAEI, the primary structure is not a single uniform theory, but a collection of subtheories linked in a hierarchy (the ordering in the hierarchy is represented by $<$). Subtheories represent different sources of information available to an agent, while the hierarchy expresses the way in which this information is combined. HAEI includes indexed modal operators L_i ; intuitively $L_i\phi$ means that ϕ is in the subtheory τ_i . And default laws are represented as $L_i\phi \wedge \neg L_j\neg\psi \rightarrow \psi$, where $\tau_j < \tau_i$.

Each subtheory τ_i is associated with a base set A_i . Within the base set A_i , the occurrence of L_j is restricted by the following condition:

- If L occurs positively (negatively) in A_i , then $\tau_j \preceq \tau_i$ ($\tau_j < \tau_i$)

So A_i cannot refer to subtheories that succeed τ_i , and τ_i is forbidden from representing what it does NOT contain.

A COMPLEX STABLE EXPANSION of a HAEI structure τ is a set of sets of sentences T_i corresponding to the subtheories τ_i of τ . It obeys the following conditions (ϕ is a sentence without a modal operator):

1. Each T_i contains A_i
2. Each T_i is closed under first-order consequence
3. If $\phi \in T_j$, and $\tau_j \preceq \tau_i$, then $L_j\phi \in T_i$
4. If $\phi \in T_j$, and $\tau_j < \tau_i$, then $\neg L_j\phi \in T_i$
5. If $\phi \in T_j$, and $\tau_j < \tau_i$, then $\phi \in T_i$.

These conditions ensure that information present in the lowest subtheories of the hierarchy percolates to its top. More specific evidence, or preferred defaults, should therefore be placed lower in the hierarchy, so that their effects will block the action of more general defaults. We'll shortly give a simple example illustrating this point: the line of reasoning behind the "penguin principle."

But first note that unlike autoepistemic base sets, each HAEI base set will have a UNIQUE minimal, complex stable expansion (see Konolige 1988). So one can talk of "the" theory of a HAEI structure. And one

can identify a subtheory τ_i with the sentences in its complex stable expansion T_i .

Konolige (1988) shows how HAEL validates the penguin principle. One starts with a HAEL structure representing the appropriate premises and proves that Tweety doesn't fly in the stable complex expansion. The following HAEL structure represents the premises, where $F(x)$ means x flies, $B(x)$ means x is a bird, and $P(x)$ means x is a penguin:

$$(21) \quad \begin{aligned} \tau_0 &< \tau_1 < \tau_2 \\ A_0 &= \{P(a), B(a)\} \\ A_1 &= \{L_1 P(a) \wedge \neg L_0 F(a) \rightarrow \neg F(a)\} \\ A_2 &= \{L_2 B(a) \wedge \neg L_1 \neg F(a) \rightarrow F(a)\} \end{aligned}$$

Intuitively, the penguin law in A_1 , the bird law in A_2 , and $\tau_1 < \tau_2$ represent the fact that information about taxonomic subkinds takes precedence over that about kinds. Moreover, the facts are in A_0 , where τ_0 precedes τ_1 and τ_2 , and this reflects the idea that facts have precedence over defaults. One must prove that $\neg F(a)$ is in the complex stable expansion.

We now build the complex stable expansion of (21) using the conditions listed above. τ_0 contains all logical consequences of $P(a)$, $B(a)$, $L_0 B(a)$, and $L_0 P(a)$. $\neg L_0 F(a)$ is NOT in τ_0 but is in τ_1 , as is $L_0 P(a)$. So $\neg F(a)$ is in τ_1 . Hence $L_1 \neg F(a)$ is in τ_2 , so $F(a)$ cannot be derived in τ_2 . Instead, $\neg F(a)$ is in τ_2 because it's in τ_1 . So the complex stable expansion of (21) contains $\neg F(a)$, as required. In words, from knowing that a is a penguin and a bird, penguins don't fly and birds do, that a doesn't fly has been inferred.

The penguin principle reflects a constraint on the order defined by the heirarchy, that default information about subkinds must take precedence over kinds. A further constraint on the hierarchy is the following: a default law $L_i \phi_1 \wedge \neg L_{i'} \neg \psi_1 \rightarrow \psi_1$ would take precedence over the default law $L_j \phi_2 \wedge \neg L_{j'} \neg \psi_2 \rightarrow \psi_2$ if ψ_1 and ψ_2 cannot both hold in a consistent expansion, and ϕ_1 logically entails ϕ_2 . In words, default laws with logically more specific antecedents take precedence over conflicting defaults with less specific antecedents. So in the above example the subtheories τ_i and τ_j must be ordered accordingly. This constraint on the hierarchy ensures that conflict between defeasible laws whose antecedents are logically related will be resolved. Conflict between default laws whose antecedents aren't logically related will fail to be resolved, however. Thus the logic will supply a principled means for resolving conflict among knowledge sources.

HAEL supplies a constructive semantics for the modal operators L_i . The details of this are explored in Konolige (1988). It is relatively straight-

forward to augment this semantics with the interpretation of well-formed expressions in DRT. This involves augmenting the first-order valuations v used in the semantic interpretations of ordinary HAEL with the embedding function f used in DRT. $v| =_f \phi$ for formulae ϕ containing no modal operators has already been defined. It corresponds to the standard DRT truth definition.

The augmented HAEL semantics is then characterized as follows. First, we add the modal operators to the language:

- If ϕ is a WFF without a modal operator, then $L_i\phi$ is a WFF.

Now we define the truth conditions for these WFF. Let $\Gamma_1, \dots, \Gamma_n, \dots$ be sets of sentences, to be thought of as belief subsets. The interpretation of indexed operators L_i is defined with respect to the sequence of belief subsets $\Gamma_1, \dots, \Gamma_n, \dots$, which Konolige calls a COMPLEX BELIEF SET. The complex belief set is essentially the modal index. The interpretation rules for HAEL valuations are then defined as follows, where ϕ is a sentence lacking a modal operator:

$$\begin{aligned} \langle v, \Gamma_1, \dots, \Gamma_n, \dots \rangle &|=_f \phi \quad \text{iff } v| =_f \phi \\ \langle v, \Gamma_1, \dots, \Gamma_n, \dots \rangle &|=_f L_i \phi \quad \text{iff } \phi \in \Gamma_i \end{aligned}$$

Konolige shows that the validity associated with the semantics for L_i bears the appropriate soundness and completeness relations with the construction of complex stable expansions.

6. Building the DRSs for extended texts

As we've mentioned, the rules for constructing DRSs of extended text are default rules about the value of the updating function $\langle \Delta, \alpha, \beta \rangle$. These rules reflect WK and LK, and they will be represented in HAEL.

For convenience, we assume as a notational convention in what follows that $\tau_i \leq \tau_j$ iff $i \leq j$. The representation of *narration* in HAEL is given below together with the definition of $pre(e_1, e_2)$. And as before, $me(\alpha)$ refers to the eventuality described in α .¹⁰

- *Definition of $pre(e_1, e_2)$*
 $pre(e_1, e_2) \leftrightarrow (\exists t_1, t_2)(hold(e_1, t_1) \wedge hold(e_2, t_2) \wedge t_1 < t_2)$
- *Narration*
 $L_i \langle \Delta, \alpha, \beta \rangle \wedge L_{i-1} \neg pre(me(\alpha), me(\beta) \rightarrow pre(me(\alpha), me(\beta)))$

This rule will be applied to build the DRS representing text (2). Recall that the logical forms of the sentences in (2) are α and β :

(α) $[e_1][turnround(max, e_1), e_1 < now]$

(β) $[e_2][hit(john, max, e_2), e_2 < now]$

To see how *narration* applies, we make some assumptions about the reader's knowledge base (KB) when interpreting text. The KB is defined as a HAEL structure, and in words, the following holds:

- The logical forms α and β of the sentences of the text form part of A_0 .

- The reader assumes text is coherent: in other words one can update the text so far with the sentence currently being processed. So in this case $\langle \alpha, \alpha, \beta \rangle$ is in A_0 .

- WK and LK, such as causal law 4 and *narration*, are in the base sets representing the reader's KB, subject to the constraints on the hierarchy described earlier.

So the reader's KB when interpreting (2) in the "null" discourse context is KB_2 , where *narration* has the appropriate indices on the modal operators relative to A_i (that is, i in *narration* is 1):

(KB_2) $\tau_0 < \tau_1$
 $A_0 = \{\alpha, \beta, \langle \alpha, \alpha, \beta \rangle\}$
 $A_1 = \{narration\}$

The natural interpretation of (2) is derived by calculating the complex stable expansion of KB_2 , using the logic HAEL.

Given the above conditions on complex stable expansions, the antecedent to *narration* is verified in τ_1 , so by logical closure its consequent, $pre(me(\alpha), me(\beta))$, is in τ_1 . This is simply an instance of defeasible modus ponens. Furthermore, by the punctual nature of events $pre(me(\alpha), me(\beta))$ entails $me(\alpha) < me(\beta)$. So using HAEL, the required temporal precedence relation between the turning round and the hitting is inferred, producing the following logical form for (2):

(2') $[e_1, e_2][turnround(max, e_1), e_1 < now,$
 $hit(john, max, e_2), e_2 < now,$
 $e_1 < e_2]$

That is, (2') is the value of $\langle \alpha, \alpha, \beta \rangle$ with respect to the HAEL structure KB_2 .¹¹ So the interpretation of (2) given the above reader's KB is one where the turning round preceded the hitting, in agreement with (2)'s natural reading. It should be stressed, however, that if the reader's KB had contained the knowledge that the hitting preceded the turning round, then the representation of (2) would be different: the antecedent to *narration* would not be verified in the subtheory τ_1 because the $\neg L_0 \neg$ conjunct would be false. This blocks the inference to *narration*'s conse-

quent in τ_1 and hence in the complex stable expression. With this new KB, the value for $\langle \alpha, \alpha, \beta \rangle$ is different, and the logical form of (2) is (2'') rather than (2').

(2'') $[e_1, e_2][\text{turnround}(\text{max}, e_1), e_1 < \text{now},$
 $\text{hit}(\text{john}, \text{max}, e_2), e_2 < \text{now},$
 $e_2 < e_1]$

So the representation of text is sensitive to the reader's knowledge of the context in which the text is uttered.

Now consider text (4).

(4) Max fell. John pushed him.

The logical forms of the two sentences are respectively α and β :

(α) $[e_1][\text{fall}(\text{max}, e_1), e_1 < \text{now}]$
 (β) $[e_2][\text{push}(\text{john}, \text{max}, e_2), e_2 < \text{now}]$

I've suggested that causal law 4 will play a central part in calculating the DRS representing (4).

• *Causal law 4*

If the events e_1 and e_2 are connected somehow where e_1 is x falling and e_2 is y pushing x , then unless there's information to the contrary, e_2 caused e_1 .

For intuitively, it is the WK reflected in (4) that distinguishes it from (2). The consequent of causal law 4 cannot hold simultaneously with the consequent of *narration* given that *causes precede effects*:

• *Causes precede effects*

$(\forall e_1, e_2)(\text{cause}(e_2, e_1) \rightarrow \neg e_1 < e_2)$

We make the same assumptions about the reader's KB for interpreting (4) as we made before but with the new α and β . As before, given these assumptions, the antecedent to *narration*, except for perhaps the $\neg L \neg$ conjunct, is verified when interpreting (4). We wish to infer the consequent of causal law 4, and so this law must be represented so that the logic ensures it "overrides" *narration*. For otherwise, the inference pattern will be one where the consequent of *narration* is inferred, contrary to intuitions. Given that causal law 4 is defeasible, there is only one way in which it can be made to override *narration* in HAEL, and that is to make its antecedent more specific than that of *narration* in the appropriate sense, so that causal law 4 takes precedence in the hierarchy.

The antecedent of *narration* includes the conjunct $\langle \Delta, \alpha, \beta \rangle$. So the antecedent of causal law 4 must entail $\langle \Delta, \alpha, \beta \rangle$ in order to make it more

specific in the appropriate sense. Having the two laws logically related like this will ensure that resolving conflict between them is done in a systematic logical way. Moreover, causal law 4 must be represented in terms of $\langle \Delta, \alpha, \beta \rangle$, since in line with Partee's suggestion, we aim to use the causal knowledge it encapsulates to reason about how the logical form of text is updated. As it stands, however, causal law 4 is not stated in terms of updating DRSs. But by the definition of the temporal coherence of text given in Lascarides and Oberlander (1992), $\langle \Delta, \alpha, \beta \rangle$ can be defined only if α 's event is connected somehow to β 's. In other words, if $r(e_1, e_2)$ means that the events e_1 and e_2 are connected somehow, we must have the following infeasible law:

$$(22) \quad \langle \Delta, \alpha, \beta \rangle \rightarrow r(me(\alpha), me(\beta))$$

(22) helps explain the awkwardness of texts like (23) (first cited in Moens 1987) in the context where no connection between the car breaking down and the sun setting is retrievable:

(23) ?Max's car broke down. The sun set.

(22) motivates restating causal law 4 in HAEL as follows:¹²

• *Causal law 4*

$$L_i(\langle \Delta, \alpha, \beta \rangle \wedge fall(x, me(\alpha)) \wedge push(y, x, me(\beta))) \wedge \neg L_{i-1} \neg cause(me(\beta), me(\alpha)) \rightarrow cause(me(\beta), me(\alpha))$$

We have represented the causal law as a mixture of WK and LK; it states that given that sentence β is to be added to the text Δ by relating β 's eventuality to α 's, and given what α and β describe, the reader believes that the event in β caused that in α , unless there is information to the contrary.

The above representation of causal law 4 and *causes precede effects* ensures that causal law 4's consequent conflicts with that of *narration*. Moreover, by the semantics of conjunction, causal law 4's antecedent except for the $\neg L \neg$ conjunct entails that of *narration*, and thus the causal law has priority over *narration* in the hierarchy. So any statement of the reader's KB must place causal law 4 in A_i and *narration* in A_j where $i < j$. This is exactly the relation between default laws in the penguin principle. Given our assumptions, the reader's KB when interpreting text (4) is KB_4 , where *narration* and the causal law have appropriate indices on the modal operators:

$$\begin{aligned} (KB_4) \quad & \tau_0 < t_1 < t_2 \\ & A_0 = \{\alpha, \beta, \langle \alpha, \alpha, \beta \rangle, \text{causes precede effects}\} \\ & A_1 = \{\text{causal law 4}\} \\ & A_2 = \{\text{narration}\} \end{aligned}$$

Note that *causes precede effects* is in A_0 because it is an INDEFEASIBLE law and so has precedence over ALL defeasible laws. The line of reasoning to produce the stable expansion from KB_4 is just like that used in the penguin principle example above. The antecedent to causal law 4 is verified in A_1 because α and β respectively entail *fall(max, me(α))* and *push(john, max, me(β))*. So the consequent of the most specific defeasible law is inferred: *cause(me(β), me(α))*. Thus the above reasoning about the value of $\langle \alpha, \alpha, \beta \rangle$ produces the following representation of (4):

- (4') $[e_1, e_2][fall(max, e_1), e_1 < now$
push(john, max, e_2), e_2 < now
cause(e_2, e_1)]

The interpretation of (3) is worked out in a very similar way to that of (4).

- (3) The council built the bridge. The architect drew up the plans.

The interpretation exploits the following knowledge: if building the bridge and drawing the plans are connected, then unless there's information to the contrary, the latter is part of the preparatory phase of the former. This knowledge is represented in a similar fashion to causal law 4. Just like causal law 4, it is more specific than *narration* and so by the penguin principle it overrides it in the interpretation of (3).

The construction of DRSs representing extended text is reader-specific and context-specific. Working out event relations crucially uses nonmonotonic inference from premises that represent the reader's KB. The nonmonotonic nature of the inference entails that as this KB is enriched with further LK and WK, or further knowledge about discourse context, some inferences about the event relations may be retracted and others added. Lascarides et al. (1992) show how the nonmonotonic construction of logical form yields a formal account of the difference between the non-iconic (4) IN VACUO, and the iconic (4) in the discourse context of (13).

- (13) John and Max were at the edge of the cliff. Max felt a sharp blow on the back of his neck. *He fell. John pushed him.* Max rolled over the edge of the cliff.

In Partee (1984), semantic structure is built from syntax alone and is thus independent of the reader's knowledge. I have extended her account in line with her suggestions, to enable the construction of temporal structure to be influenced by semantic content and the reader's background knowledge, so as to distinguish the interpretation of texts (2) and (4), which have similar syntax. The crucial difference between (2) and (4) is not that (2) is always interpreted as a narrative whereas (4) is always interpreted

as a nonnarrative. Rather, IN THE ABSENCE OF FURTHER INFORMATION, (2) is interpreted as a narrative and (4) is interpreted as a nonnarrative.

7. How aspectual classification affects DRS construction

How can the distinct semantics of events and states yield distinct temporal structures in text, as demonstrated in (24) and (25)?

(24) Max arrived at the house. Mary stepped outside.

(25) Max arrived at the house. Mary was outside.

And how is the distinction between the temporal structures of (5) and (6) to be explained?

(5) Max opened the door. The room was pitch dark.

(6) Max switched off the light. The room was pitch dark.

In text (6), one assumes that the event causes the state: this is derived from a defeasible causal law as before. But in text (5), one assumes a different causal structure; entering a room does not typically cause darkness. And one also assumes the event and state overlap. So text (5) indicates that perhaps the following holds:

- *States overlap*

Contrary to other information, one assumes that the state overlaps previously described eventualities.

Informally, according to the above *states overlap*, the event and state hold at a common time in (25) and (5) because there is no information to the contrary. In (6) on the other hand, a causal law conflicts with *states overlap*, given the indefeasible law *causes precede effects*.

There is intuitive motivation for *states overlap*: it is a realization of Grice's maxim of relevance. The argument goes as follows: in order to construct a full picture of the situation described by NL text, one should know the relative occurrences of all the eventualities, including where they start and stop relative to each other. In the case of events, which are punctual, one can fully determine the start and stop through inferences that use rules like *narration*. But because states extend in time, knowing a POINT of time where the state holds, as provided in *narration*, is not sufficient for determining where it starts relative to other eventualities described in the text.

There are several linguistic mechanisms one can use to indicate where a state starts relative to the other eventualities. One could explicitly refer in the text to what caused the state, and so from the law that causes

precede effects one knows the relative place where the state starts. This is what's going on in (6). Alternatively, one can use temporal adverbials to indicate where a state starts. But these two mechanisms are not enough: what about determining where states start in texts that do NOT feature adverbials or causes? *States overlap* is another vital mechanism for determining where a state starts relative to other eventualities described in text. For it essentially asserts that if there is no "explicit" indication of where the state starts, in the form of mentioning causes or temporal adverbials for example, then the start of the state is assumed to be irrelevant; that is, the state started to hold BEFORE the situation that the text is concerned with occurred. Given the principle that the descriptive order of events typically reflects the order of their perception (compare the motivation for *narration*), this results in temporal overlap between a state and previously described eventualities. Thus *states overlap* can be viewed as a manifestation of Grice's maxim of relevance, for it asserts that unless there is indication in the text to the contrary, the point where a state starts is assumed to be irrelevant.

We represented Grice's maxim of orderliness as defeasible LK. Similarly, *states overlap*, which we claim is closely related to Grice's maxim of relevance, is defeasible LK. It is represented as follows:

- *States overlap*

$$L_i(\langle \Delta, \alpha, \beta \rangle \wedge \text{state}(\text{me}(\beta))) \wedge \neg L_{i-1} \neg \text{overlap}(\text{me}(\alpha), \text{me}(\beta)) \rightarrow \text{overlap}(\text{me}(\alpha), \text{me}(\beta))$$

States overlap is added to the reader's KB. The predicates *overlap* and *hold* are related by the following nonlogical axiom, so that e_1 overlaps e_2 means that e_1 and e_2 hold at a common point of time:

- *Overlap and hold*

$$(\forall e_1 e_2)(\text{overlap}(e_1, e_2) \leftrightarrow (\exists t)(\text{hold}(e_1, t) \wedge \text{hold}(e_2, t)))$$

The antecedent of *states overlap* is more specific than that of *narration*, but the consequents don't conflict. Indeed, given the above definitions of *overlap* and *pre*, and the way events and states behave with respect to the predicate *hold* that we described earlier, the following law holds:

- *Overlap and pre*

$$(\forall e_1, e_2)(\text{overlap}(e_1, e_2) \wedge \text{event}(e_1) \wedge \text{state}(e_2) \rightarrow \text{pre}(e_1, e_2))$$

So because the consequents don't conflict, *states overlap* can be added to the KB at the same level in the hierarchy as *narration*.

If $\text{me}(\beta)$ is an event, then updating the text with β will not verify the antecedent of *states overlap*. So *states overlap* will play no part in determining where events start. Thus this new defeasible LK does not affect

the reader's beliefs about texts (1), (2), (3), or (4), as they describe events alone.

However, *states overlap* plays a central role in the reader's assessments of the temporal structures described by texts (25), (5), and (6).

(24) Max arrived at the house. Mary stepped outside.

(25) Max arrived at the house. Mary was outside.

(5) Max opened the door. The room was pitch dark.

(6) Max switched off the light. The room was pitch dark.

Let's look at the difference between (24) and (25) first. The logical forms of the sentences in (24) are respectively α and β_1 , and the logical forms of (25)'s sentences are α and β_2 .

(α) $[e_1, t_1][arrive(max, e_1), e_1 < now]$

(β_1) $[e_2, t_2][stepout(mary, e_2), e_2 < now]$

(β_2) $[e_2, t_2][beout(mary, e_2), e_2 < now]$

The truth conditions of β_1 and β_2 are the same, save that e_2 in β_1 is the EVENT of Mary stepping outside, and e_2 in β_2 is the STATE of Mary being outside. I assume *state(me(β_2))* holds because of the temporal behavior of the predicate *beout*. That is, if *beout*(x, e) is true, then e holds on a set of times that form an open interval, thus ensuring e is a state. Similarly, that *event(me(β))* holds is determined by the temporal behavior of the predicate *stepout*.

Despite the fact that the semantics of events and states are distinct in that events are punctual and states are not, β_1 and β_2 on their own do not account for the distinct natural interpretations of (24) and (25). This is because no conditions are imposed on the temporal relation between e_1 and e_2 in the above.

But the reader's KB contains defeasible LK, namely *narration* and *states overlap*. Using HAEL, one infers from KB_{24} that the logical form of (24) is (24'), and one infers from KB_{25} that the logical form of (25) is (25').

(KB_{24}) $\tau_0 < \tau_1 < \tau_2$

$A_0 = \{\langle \alpha, \alpha, \beta_1 \rangle, \alpha, \beta_1, \text{causes precede effects}\}$

$A_1 = \{\text{states overlap, narration}\}$

$A_2 = \{\text{causal law 4}\}$

(24') $[e_1, e_2][arrive(max, house, e_1), e_1 < now$

$stepout(mary, e_2), e_2 < now$

$e_1 < e_2]$

(KB_{25}) $\tau_0 < \tau_1 < \tau_2$

$A_0 = \{\langle \alpha, \alpha, \beta_2 \rangle, \alpha, \beta_2, \text{causes precede effects}\}$

$A_1 = \{\text{states overlap, narration}\}$

$A_2 = \{\text{causal law 4}\}$

- (25') $[e_1, e_2][\text{arrive}(\text{max}, \text{house}, e_1), e_1 < \text{now}$
 $\text{beout}(\text{mary}, e_2), e_2 < \text{now}$
 $\text{overlap}(e_1 < e_2)]$

The inference pattern in both cases is defeasible modus ponens.

Now for texts (5) and (6): the sentences in (5) are represented by α_1 and β and the sentences in (6) by α_2 and β .

- (α_1) $[e_1][\text{open}(\text{max}, \text{door}, e_1), e_1 < \text{now}]$
 (α_2) $[e_1][\text{switchoff}(\text{max}, \text{light}, e_1), e_1 < \text{now}]$
 (β) $[e_2][\text{dark}(\text{room}, e_2), e_2 < \text{now}]$

Causal law 6 represents the knowledge that if switching off the light and the room being dark are connected, then the switching off the light caused the darkness unless there's already information to the contrary:¹³

- Causal law 6
 $L_i(\langle \Delta, \alpha, \beta \rangle) \wedge \text{switchoff}(\text{light}, \text{me}(\alpha)) \wedge \text{dark}(\text{room}, \text{me}(\beta)) \wedge \neg L_{i-1}$
 $\neg \text{cause}(\text{me}(\alpha), \text{me}(\beta)) \rightarrow \text{cause}(\text{me}(\alpha), \text{me}(\beta))$

This defeasible knowledge must be added to the KB. It is more specific than *states overlap*: by the stative classification of the predicate *dark*, $\text{dark}(x, e)$ entails $\text{state}(e)$. Moreover, the consequents of *states overlap* and causal law 6 conflict in the light of *causes precede effects*. So causal law 6 must be added to the hierarchy of defaults so that it has precedence over *states overlap*.

The relevant KB for text (5) verifies the antecedents of *narration* and *states overlap*. So similarly to (25), the following logical form of (5) is constructed as a result of the inferences in HAEL:

- (5') $[e_1, e_2][\text{open}(\text{max}, \text{door}, e_1), e_1 < \text{now}$
 $\text{dark}(\text{room}, e_2), e_2 < \text{now}$
 $\text{overlap}(e_1, e_2)]$

On the other hand, the appropriate KB for text (6) verifies the antecedents to causal law 6 as well as *states overlap* and *narration*, save perhaps the $\neg L \neg$ conjuncts. So if the defeasible rules are ordered as described above, the penguin principle yields the following logical form for (6):

- (6') $[e_1, e_2][\text{switchoff}(\text{max}, \text{light}, e_1), e_1 < \text{now}$
 $\text{dark}(\text{room}, e_2), e_2 < \text{now}$
 $\text{cause}(e_1, e_2)]$

It must be stressed that these representations are context-sensitive. Suppose that the reader has knowledge that Max opening the door CAUSED the room to be dark, so that the relevant KB contains

$cause(me(\alpha_1), me(\beta))$ in the base set A_0 . Then $cause(me(\alpha_1), me(\beta))$ is in the stable expansion rather than $overlap(me(\alpha_1), me(\beta))$.¹⁴ So the representation of (5) is different; we replace *overlap* in (5') with *cause*. Thus in essence, the reader's interpretation of (5) is represented by (5') unless the context is one where the reader already knows something to the contrary.

8. Conclusion

In this paper the ordering question has been addressed: given a particular order in which eventualities are described, what are the constraints in interpretation on their temporal and causal relations in the world?

The logical framework developed was one where the sentences' truth conditions, defeasible WK, and defeasible LK all contribute to the construction of the representation of text. As a result, the theory solved many of the problems associated with the ordering question. This extends the approach deployed in current DRT theories on tense, which offer an account where temporal relations between eventualities are determined from syntactic structure alone. It also refines theories of discourse interpretation that exploit causal knowledge, such as those of Hobbs and Dahlgren, because the causal knowledge was put to work by a well-defined notion of logical entailment. Consequently, conflicts among the knowledge sources recruited during interpretation were resolved in a systematic way.

The constraints on the relation between the descriptive order of eventualities and their temporal order described in this paper are incomplete. The question of how temporal adverbials, quantifiers, and the pluperfect behave is not addressed here. Neither have the factors that contribute to discourse "popping" been explored. Lascarides and Asher (1991b) show that the strategy pursued here can be extended to provide an account of the pluperfect and can be used to give an account of when and how discourse popping occurs. Lascarides et al. (n.d.) explore the discourse roles of temporal adverbials and quantifiers.

Received 17 December 1991

University of Edinburgh

Revised version received

15 June 1992

Notes

- * This paper is a rewritten version of an HCRC research report published in June 1990. I would like to thank Nicholas Asher, Mimo Caenepeel, Lex Holt, Marc Moens, Jon

Oberlander, and an anonymous reviewer for useful comments on earlier drafts. The research was funded by the Science and Engineering Research Council (UK) under grant number GR/G22077. HCRC is supported by the Economic and Social Research Council (UK). Correspondence address: Human Communications Centre, University of Edinburgh, 2 Buccleuch Place, GB-Edinburgh EH8 9LW, England.

1. These texts were first cited in Moens (1987).
2. More recently, Hobbs et al. (1990) use weighted abduction in order to choose among conflicting knowledge sources. But the semantics of the weights remain undefined, and their values are calculated in an arbitrary way.
3. For the simple texts we consider here, conditions on the discourse entities are atomic conditions rather than (embedded) DRSs. Partee analyzes examples where embedded DRSs are needed to encode quantification over times, as in *Whenever Mary rings, Max is asleep*. But we won't be concerned with such examples here. The account could be straightforwardly extended to deal with these examples.
4. An eventuality is the general term for an event or state, due to Bach (1986).
5. For simplicity, we're treating $r_1 < r_2 \vee r_2 < r_1$ as an atomic condition on discourse entities. It holds for an embedding f of discourse entities into a model M if $f(r_1)$ is earlier than $f(r_2)$ in M , or $f(r_2)$ is earlier than $f(r_1)$ in M .
6. A similar view is proposed in Asher (i.p.).
7. For the sake of simplicity, nominal anaphora resolution has been ignored.
8. The number on the causal law corresponds to the number of the text it is relevant to. Note that causal law 4 does NOT say that pushings cause fallings; this would be far-fetched.
9. At any rate, it is important not to confuse default with frequency.
10. For the simple fragment considered in this paper, $me(\alpha)$ is defined to be the (unique) eventuality discourse entity in the first set of α .
11. Recall that the value of $\langle \Delta, \alpha, \beta \rangle$ involves adding β 's discourse entities and conditions to Δ 's and adding the relation between $me(\alpha)$ and $me(\beta)$ inferred in HAEL.
12. For simplicity, I have ignored quantifying over x and y since it's not the focus of interest here.
13. For the sake of simplicity, the problem of inferring that the light is in the room is ignored.
14. The antecedent to *states overlap* isn't verified in the KB because $L_0 \sqcap \text{overlap}(me(\alpha), me(\beta))$ will hold.

References

- Appelt, Doug; and Konolige, Kurt (1988). A practical nonmonotonic theory for reasoning about speech acts. In *Proceedings of the 25th Annual Meeting of the Association for Computational Linguistics*, 170–178. Buffalo: State University of New York at Buffalo.
- Asher, Nicholas (i.p.). *Reference to Abstract Objects in English: A Philosophical Semantics for Natural Language Metaphysics*. Dordrecht: Kluwer.
- ; and Morreau, Michael (1991). Common sense entailment: a modal theory of nonmonotonic reasoning. In *Proceedings to the 12th International Joint Conference on Artificial Intelligence*. Sydney.
- Bach, Emmon (1986). The algebra of events. *Linguistics and Philosophy* 9, 5–16.
- Caenepeel, Mimo (1991). Event structure versus discourse coherence. In *Proceedings of the Workshop on Discourse Coherence*. Edinburgh: University of Edinburgh.

- Dahlgren, Kathleen (1988). *Naive Semantics for Natural Language Understanding*. Boston: Kluwer.
- Davidson, Donald (1967). Causal relations. *Journal of Philosophy* 64.
- Dowty, David (1986). The effects of aspectual class on the temporal structure of discourse: semantics or pragmatics? *Linguistics and Philosophy* 9, 37–61.
- Grice, H. Paul (1975). Logic and conversation. In *Syntax and Semantics*, vol. 3: *Speech Acts*, Peter Cole and Jerry L. Morgan (eds.), 41–58. New York: Academic Press.
- Grosz, Barbara J.; and Sidner, Candice L. (1986). Attention, intentions, and the structure of discourse. *Computational Linguistics* 12, 175–204.
- Hinrichs, Erhard (1986). Temporal anaphora in discourses of English. *Linguistics and Philosophy* 9, 63–82.
- Hobbs, Jerry R. (1979). Coherence and coreference. *Cognitive Science* 3, 67–90.
- (1985). *On the Coherence and Structure of Discourse*. Report No. CSLI-85-37. Stanford: Center for the Study of Language and Information.
- ; Stickel, Mark; Appelt, Doug; and Martin, Paul (1990). *Interpretation as Abduction*. Technical Note No. 499. Menlo Park, CA: Artificial Intelligence Center, SRI International.
- Joshi, Aravind; Webber, Bonnie; and Weischedel, Ralph (1984). Default reasoning in interaction. In *Proceedings of the Non-Monotonic Reasoning Workshop*, 144–150. New York: AAAI.
- Kamp, Hans; and Rohrer, Christian (1983). Tense in texts. In *Meaning, Use and Interpretation of Language*, Rainer Bauerle, Christian Schwarze, and Arnim von Stechow (eds.), 250–269. Berlin: de Gruyter.
- Konolige, Kurt (1987). *On the Relation Between Default and Autoepistemic Logic*. Report No. CSLI-87-105. Stanford: Center for the Study of Language and Information.
- (1988). *Hierarchic Autoepistemic Theories for Nonmonotonic Reasoning: Preliminary Report*. SRI International Technical Note No. 446. Menlo Park, CA: SRI International.
- Kripke, Saul (1972). Naming and necessity. In *Semantics of Natural Language*, Donald Davidson and Gilbert Harman (eds.), 253–355. Dordrecht: Reidel.
- Lascarides, Alex (1988). A formal semantic analysis of the progressive. Unpublished Ph.D. thesis, Centre for Cognitive Science, University of Edinburgh.
- (1991). The progressive and the imperfective paradox. *Synthese* 87 (6), 401–447.
- ; and Asher, Nicholas (1991a). Discourse relations and defeasible knowledge. In *Proceedings to the 29th Annual Meeting of the Association of Computational Linguistics*. Berkeley.
- ; and Asher, Nicholas (1991b). Discourse relations and common sense entailment. DYANA deliverable 2.5b, available from Centre for Cognitive Science, University of Edinburgh.
- ; Asher, Nicholas; and Oberlander, Jon (1992). Inferring discourse relations in context. In *Proceedings to the 30th Annual Meeting of the Association of Computational Linguistics*. Delaware.
- ; and Oberlander, Jon (1992). Temporal coherence and defeasible knowledge. *Theoretical Linguistics* 19.
- ; Oberlander, Jon; and Asher, Nicholas (n.d.). A semantics and pragmatics for temporal connectives. Unpublished manuscript.
- Moens, Marc (1987). Tense, aspect and temporal reference. Unpublished Ph.D. thesis, Centre for Cognitive Science, University of Edinburgh.
- ; and Steedman, Mark (1988). Temporal ontology and temporal reference. *Computational Linguistics* 14, 15–28.
- Moore, Robert C. (1985). *A Formal Theory of Knowledge and Action*. Report No. CSLI-85-31. Stanford: Center for the Study of Language and Information.

- Partee, Barbara H. (1984). Nominal and temporal anaphora. *Linguistics and Philosophy* 7, 243–286.
- Polanyi, Livia (1985). A theory of discourse structure and discourse coherence. In *Papers from the General Session at the Twenty-First Regional Meeting of the Chicago Linguistic Society*, William H. Eilfort, Paul D. Kroeber, and Karen L. Peterson (eds.). Chicago: CLS.
- Reichenbach, Hans (1947). *Elements of Symbolic Logic*. London: Macmillan.
- Reiter, Ray (1980). A logic for default reasoning. *Artificial Intelligence* 13, 81–132.
- Thompson, Sandra A; and Mann, William C. (1987). Rhetorical structure theory: a framework for the analysis of texts. *International Pragmatics Association Papers in Pragmatics* 1, 79–105.
- Webber, Bonnie L. (1988). Tense as discourse anaphor. *Computational Linguistics* 14, 61–73.