Abducing Temporal Discourse*

Alex Lascarides and Jon Oberlander

Centre for Cognitive Science *and* Human Communication Research Centre
University of Edinburgh
2, Buccleuch Place, Edinburgh EH8 9LW
Scotland, UK

Abstract

We focus on the following question: given the causal and temporal relations between events in a knowledge base, what are the ways they can be described in extended text? We argue that we want to be able to generate *laconic* text, where certain temporal information remains implicit but pragmatically inferrable. An algorithm for generating laconic text is proposed, interleaving abduction and nonmonotonic deduction over a formal model of pragmatic implicature. We demonstrate that the nonmonotonicity ensures that the generation of laconic text is influenced by the preceding linguistic and extra-linguistic context.

1 Introduction

Given the causal and temporal relations between events in a knowledge base (KB), what are the ways they can be described in extended text?

Elsewhere, we have argued that amongst the various ways, we will prefer those which allow the reader-hearer H to retrieve all and only the causal and temporal relations which the author-speaker S intended to convey. Such utterances we call *temporally adequate*, a notion we define precisely below. Among those temporally adequate utterances will be some which, although not marking the correct relations explicitly, allow H to infer them successfully. Utterances which are both temporally adequate and leave information S intends to convey implicit we call *laconic*.

We'll argue below that in some cases, laconic text is the most natural and so its generation shouldn't be precluded. But generating laconic discourse can be a risky business, since it requires S to assess which implicatures H will draw from it. Having S reason about implicature may be an attractive idea, but without a logical theory to underpin it, it is hard to assess its utility. Lascarides and Asher (1991a) developed a particular formal model of temporal implicature. They extended the formal interpretative framework of discourse representation theory (Kamp 1981) by

* The support of the Science and Engineering Research Council through project number GR/G22077 is gratefully acknowledged. HCRC is supported by the Economic and Social Research Council. We thank Robert Dale, Masato Ishizaki and two anonymous reviewers for their helpful comments.

incorporating discourse (or rhetorical) relations. By representing knowledge of the world and of linguistic pragmatics by defeasible conditionals in a nonmonotonic logic, they showed formally how the knowledge could be used to augment the relatively neutral semantic representations with their richer pragmatic implicatures, including discourse relations and temporal relations.

Here, we consider how to recruit this framework in the service of generation. The current goal is to retain the same defeasible rules, and assess the extent to which the model can inform the generation process. In Oberlander and Lascarides (1991), we investigated one possible strategy, based Joshi *et al*'s (1986) Interactive Defaults. We concluded that it left the calculation of discourse implicatures in a peripheral position, and that a deeper model was required. We have therefore turned to a model which views generation as abduction, as proposed in Hobbs *et al* (1990). We offer a possible abductive generation algorithm. We conclude that the abductive generation of laconic discourse in the defeasible framework requires the interleaving of abduction and deduction; and that the nonmonotonicity of the logic permits context-sensitive generation, which is currently difficult in Hobbs *et al*'s system of weighted abduction.

In the next section, we motivate the need for generating laconic discourse. We then state the temporal constraints that define temporal adequacy, introduce the theory of discourse structure and nonmonotonic inference, and outline the shift from interpretation to generation. We then show in detail how the generation algorithm proposed treats some simple, specific examples. Finally, we contrast our approach with weighted abduction.

2 Laconic Discourse

It is generally agreed that it is not always possible, or even desirable, to syntactically mark all the temporal information that the speaker S wishes to convey to the hearer H (Hobbs 1979, Joshi *et al* 1986, Oberlander and Lascarides 1991). In particular, it is not always appropriate to mark temporal progression with cue phrases such as *after that* or *and then*. Sometimes, simplest is best: (1) is at least as coherent as text (2) if not more so, and one would not want to preclude (1)'s generation.

(1) Max entered the office. John greeted him with a smile. He showed Max to the seat in front of his desk and offered him a cup of coffee.

(2) Max entered the office. Then, John greeted him with a smile. After that, he showed Max to the seat in front of his desk. He then offered Max a cup of coffee.

But leaving temporal progression unmarked has ramifications. Although in some cases it produces the most natural text, in others it can mislead H. Consider the case where Jon switches off the heating, and then Judy comes in and says the room is hot. Unlike text (1), the fact that this temporal progression is left unmarked in text (3) is potentially misleading for H, for the second event described could mistakenly be interpreted by H as the *cause* of the first event: text (4) would be better in this case.

(3) Jon switched off the heating. Judy came in and said the room was too hot.

(4) Jon switched off the heating. Then, Judy came in and said the room was too hot.

If temporal information is to be left implicit, then S must ensure that not only does the semantics of the linguistic structures she uses produce suitable entailments, but that the pragmatic implicatures H infers from the linguistic structures must be appropriate too. A discourse is *laconic* if it keeps some temporal information implicit while retaining the right implicatures. We wish to generate laconic texts like (1) but preclude the generation of misleading texts like (3).

3 Temporal Constraints and Implicature

To build a process that can generate laconic discourse we certainly need a characterisation of what it actually means for a text to be pragmatically appropriate. Equally obviously, we require a formal model of discourse implicature that can inform NL generation. We now address these issues.

3.1 Temporal Constraints

We concentrate on those properties required of the generated utterance that are central to temporal import. As in Lascarides and Oberlander (1992), necessary (but insufficient) properties of an adequate utterance are *temporal coherence* and *temporal reliability*. Following Bach (1986), we take 'eventualities' to cover both events and states. Temporal coherence and temporal reliability are defined in terms of a set $\mathcal{C}$ of relations between eventualities. This set intuitively describes when two eventualities are connected. The relations in $\mathcal{C}$ are: causation, the part/whole relation,[2] temporal overlap, and the immediately precedes relation (where 'e_1 immediately precedes e_2' means that e_1 and e_2 stand in a causal or part/whole relation that is compatible with e_1 temporally preceding e_2).[3] The definitions are:

- `Temporal Coherence`
  A text is *temporally coherent* if H can infer that at least one of the relations $\mathcal{C}$ holds between the eventualities described in the sentences.
- `Temporal Reliability`
  A text is *temporally reliable* if one of the relations in $\mathcal{C}$ which H infers to hold does in fact hold between the eventualities described in the sentences.

A text is temporally *incoherent* if the natural interpretation of the text is such that there are no inferrable relations between the events. A text is temporally *unreliable* if the natural interpretation of the text is such that the inferred relations between the events differ from their actual relations in the world. It follows from these definitions that a text can be coherent but unreliable. On the other hand, there may be

2 We think of 'e_1 is part of e_2' in terms of Moens and Steedman's (1988) event terminology.
3 We assume that an event e_1 precedes an event e_2 if e_1's *culmination* occurs before e_2's. So there are part/whole relations between e_1 and e_2 that are compatible with e_1 temporally preceding e_2.

no question of reliability simply because we cannot establish a temporal or causal relation between the two eventualities. An utterance is temporally *adequate* only if it is both coherent and reliable. Laconic text is adequate, and also leaves some intended temporal information implicit. Ensuring that such a candidate is adequate thus requires a model of implicature. Of course, the set of laconic discourses describing a given temporal-causal structure need not be a singleton: there may be a further choice, or there may be no laconic description. We do not consider how to choose between rival laconic utterances; however, the approach as currently envisaged will proceed from inadequate laconic text to an adequate non-laconic one.

3.2 Discourse Structure and Implicature

The basic model of implicature which we wish to embed in the generation process starts with traditional discourse representation structures (cf. Kamp 1981), but goes on to assume that candidate discourses possess hierarchical structure, with units linked by discourse relations modelled after those proposed by Hobbs (1979, 1985) (cf. also Mann and Thompson 1987, Reichman 1984, Scha and Polanyi 1988). Lascarides and Asher (1991a) use *Narration*, *Explanation*, *Background*, *Result* and *Elaboration*. These are the discourse relations central to establishing the temporal structures, and they are therefore the only ones we consider here.

Lascarides and Asher provide a logical theory for determining the discourse relations between sentences in a text, and the temporal relations between the events they describe. The logic used is the nonmonotonic logic CE proposed by Asher and Morreau (1991). Implicatures are calculated via default rules. For example, they motivate the rules below as manifestations of Gricean-style pragmatic maxims and world knowledge, where the clauses α and β appear in that order in the text, and "α and β are discourse-related" means that one of the five discourse relations holds between α and β. Informally:

- `Narration`
  If clauses α and β are discourse-related, then normally $Narration(\alpha, \beta)$ holds.
- `Axiom on Narration`
  If $Narration(\alpha, \beta)$ holds, and α and β describe events e_1 and e_2 respectively, then e_1 occurs before e_2.
- `Explanation`
  If clauses α and β are discourse-related, and the event described in β caused that described in α, then normally $Explanation(\alpha, \beta)$ holds.
- `Axiom on Explanation`
  If $Explanation(\alpha, \beta)$ holds, then event e_1 described in α does not occur before event e_2 described in β.
- `Push Causal Law`
  If clauses α and β are discourse-related, and α describes the event e_1 of x falling and β the event e_2 of y pushing x, then normally e_2 causes e_1.
- `Causes Precede Effects`
  If event e_2 causes e_1, then e_1 doesn't occur before e_2.

The rules for Narration and Explanation constitute defeasible linguistic knowledge, and the axioms on them indefeasible linguistic knowledge. In particular, Narration and its axiom convey information about the pragmatic effects of the descriptive order of events; unless there is information to the contrary, it is assumed that the descriptive order of events matches their temporal order in interpretation. The Push Causal Law is a mixture linguistic knowledge and world knowledge; given that the clauses are discourse-related somehow, the events they describe must be connected as defined in $\mathcal{C}$; here, given the events in question, they must normally stand in a causal relation. That Causes Precede their Effects is indefeasible world knowledge.

A more formal notation makes clear both the logical structure of these rules, and the problems involved in calculating implicature. Let $\langle\alpha,\beta\rangle$ be a function: the clause β is to be attached to the clause α, where α is part of the discourse structure already. Let $fall(max, e_\alpha)$ mean that the main eventuality described by the clause α is a Max falling event. Let $e_1 \prec e_2$ mean the eventuality e_1 wholly precedes e_2, and $cause(e_1, e_2)$ mean e_1 causes e_2. Finally, we represent the defeasible connective as in Asher and Morreau (1991) as a conditional $>$: so $\phi > \psi$ means 'if ϕ, then normally ψ'. The rules for modelling implicature are then represented as schemas:[4]

- **Narration**
  $\langle\alpha,\beta\rangle > Narration(\alpha,\beta)$
- **Axiom on Narration**
  $Narration(\alpha,\beta) \rightarrow e_\alpha \prec e_\beta$
- **Explanation**
  $\langle\alpha,\beta\rangle \wedge cause(e_\beta, e_\alpha) > Explanation(\alpha,\beta)$
- **Axiom on Explanation**
  $Explanation(\alpha,\beta) \rightarrow \neg e_\alpha \prec e_\beta$
- **Push Causal Law**
  $\langle\alpha,\beta\rangle \wedge fall(max, e_\alpha) \wedge push(john, max, e_\beta) > cause(e_\beta, e_\alpha)$
- **Causes Precede Effects**
  $cause(e_2, e_1) \rightarrow \neg e_1 \prec e_2$

It should be stressed that the formal theory of implicature contains many other rules, including rules concerning information conveyed by aspectual classification, the different tenses featured and temporal adverbials. We highlight the current rules because they are relevant to the simple illustrative examples we consider below. We won't tackle here examples where an eventuality should be described as a state rather than an event, or by a clause modified by the pluperfect tense rather than the simple past. But because the generation process we supply is informed by the *whole* model of implicature, it is powerful enough to tackle such choices.

---

4 Discourse structure is given a model theoretical interpretation in Asher (1991); e_α abbreviates $me(\alpha)$, which is formally defined in Lascarides and Asher (1991b) in an intuitively correct way. For simplicity, we have here ignored the modal nature of the indefeasible knowledge; in fact, an indefeasible rule is embedded within the necessity operator $\Box$.

4 Interpretation by Deduction

Let us now briefly review how CE and the defeasible rules are used deductively, during interpretation, to infer the discourse structures of candidate texts. CE represents monotonic validity as $\models$, and non-monotonic validity as $\approx\!\!\!|$. Two patterns of nonmonotonic inference are particularly relevant:

- **Defeasible Modus Ponens**
  $\phi > \psi, \phi \;|\!\!\approx\; \psi$

e.g. Birds normally fly, Tweety is a bird; so Tweety flies

- **The Penguin Principle**
  $\phi \to \psi, \psi > \chi, \phi > \neg\chi, \phi \;|\!\!\approx\; \neg\chi$

e.g. Penguins are birds, birds normally fly, penguins normally don't fly, Tweety is a penguin; so Tweety doesn't fly.

For example, in interpreting text (5) we assume the text is coherent and so attempt to attach the second clause to the first with a discourse relation (so $\langle\alpha, \beta\rangle$ is true).

(5) Max stood up. John greeted him.

In the absence of further information, the only rule whose antecedent is satisfied in interpreting text (5) is Narration. Other things being equal, we infer via Defeasible Modus Ponens that the *Narration* relation holds between (5)'s clauses. This then yields, assuming logical omniscience, an interpretation where the standing up precedes the greeting. Defeasible Modus Ponens and logical omniscience entail that in (6), the greeting preceded the standing up.

(6) John greeted Max. Max stood up.

Texts (5) and (6) show why it's essential that the temporal information conveyed by textual order should be represented. But it's also essential that this information should be represented as a default, for descriptive order doesn't *always* match temporal order, even if there are no syntactic markers to indicate this:

(7) Max fell. John pushed him.

Text (7) verifies the antecedents to *two* of our defeasible laws: Narration and the Push Causal Law. The consequents of these default laws cannot both hold in a consistent KB. By the Penguin Principle, the law with the more specific antecedent wins: the Causal Law, because its antecedent logically entails that of Narration. Hence (7) is interpreted as a case where the pushing caused the falling. In turn, this entails that the antecedent to Explanation is verified; and whilst conflicting with Narration, it's more specific, and hence its consequent—*Explanation*—follows by the Penguin Principle. Compare this with (5): similar logical forms are assigned to the clauses because of the similar syntax, but different discourse structures, and different temporal structures are derived.[5]

5 The formal details of how the logic CE models the Penguin Principle, and these interpretations, are given in Lascarides and Asher (1991a, 1991b). Note that although the double application of the Penguin Principle, as in (7), is not valid in general, they show that for the particular case considered here, CE validates the double application.

The discourse relations constrain the topic of a discourse segment, as defined in Asher (1991). Since we consider only the temporal domain here, we can replace topics with the simpler notion of *key events*, as used in Maybury (1991). Constraints between discourse structure and key events are used below to explain incoherence.

Intuitively, in describing an event structure, some events are more important in the story than others. Following Obermeier (1985), we assume that keyness is a two place relation on eventualities. The formula $key(e_1, e_2)$ means that e_1 is a key event relative to e_2; intuitively, the text segment containing a description of e_2 is *about* e_1. We assume that deciding a text's key events depends on the text's purpose and S's and H's points of view. Formally stating rules that represent relationships between purpose, point of view and key events are beyond the scope of this paper; this problem is tackled in Maybury (1991).

We provide event- and discourse-constraints on keyness: $key(e_1, e_2)$ entails that e_1 and e_2 stand in one of the relations in $\mathcal{C}$. Furthermore, the law below states that $Narration(\alpha, \beta)$ holds only if the eventualities described by α and β have a distinct common key event (cf. Asher's (1991) topic constraints on *Narration*):

- **Distinct Common Key Event for Narration**
  $Narration(\alpha, \beta) \rightarrow (\exists e)(key(e, e_\alpha) \wedge key(e, e_\beta) \wedge \neg key(e_\alpha, e_\beta) \wedge \neg key(e_\beta, e_\alpha))$

This rule is used to explain the incoherence of (8): as long as WK about cars breaking down and dying hair is represented as intuitions would dictate, one cannot find a common distinct key event, and so *Narration* between the clauses can't be inferred.

(8) ?Max's car broke down. Mary died her hair black.

But no other relation can be inferred given the above defeasible laws. And hence no discourse structure for (8) is constructed.

The relations of *Elaboration* and *Explanation* also constrain the value of *key*: the following rule states that the event that's explained is key relative to the events that explain it.

- **Key Event of Explanation**
  $Explanation(\alpha, \beta) \rightarrow key(e_\alpha, e_\beta)$

Support for this rule comes again from explaining incoherence, this time of a text first cited in Caenepeel (1991):

(9) ?Everyone laughed. Fred told a joke.

In (9), there is likely to be a causal law much like the one for pushes: if laughing and telling a joke are connected, then normally the latter caused the former. Accordingly, H will infer this causal relation, so Explanation's antecedent is verified. But in this minimal context, it is hard to see how laughing (on its own) could be a key for joke telling (on its own). This new constraint therefore blocks *Explanation* from being inferred; in essence laughing in (9) simply doesn't warrant explanation. But *Narration* cannot be deduced either, since this is blocked by the causal relation already inferred. Hence no discourse structure for (9) is constructed.[6] By a parallel argument, *Elaboration* constrains keyness in a similar way.

6 Replacing *Everyone laughed* with *Everyone laughed until their sides split* improves coherence, even though the causal relation is the same. One can envisage a purpose for the

5 Generation by Abduction

Declarative grammars are regarded as appealing partly because they can be used bidirectionally. So how can our theory of implicature for text interpretation inform text generation? One answer is: through abduction.[7] Abduction is inference that is classically logically invalid, but is nonetheless plausible: it is 'inference to the best explanation'. Essentially, abduction permits us, from $p \rightarrow q$ and q, to assume p. p is then an explanation of why q is true. Abduction can be compared to the creation of plans via theorem-proving; generation by abduction thus bears some general similarities to generation by planning.

Making the antecedent of an indefeasible law true yields the consequent, *whatever* other premises are in the KB. But with defeasible laws, we must check that the context, as characterised by H's KB, actually permits the conclusion. If other facts are available, or if other rules are logically stronger, then the conclusion may not follow. Because of this, the algorithm for generation described below interleaves abduction and deduction. A nonmonotonic deductive check (NMDC) guides the abduction process so that a defeasible rule δ is used in abduction only if it is reliable to do so.

Abduction can be viewed as a process of support-location, as in Selman and Levesque (1990). We have a knowledge structure we want to support, and abduction helps us find ways of doing so. So, for example, $Narration(\alpha, \beta)$ supports e_α before e_β (via the Axiom on Narration). Some types of support are more desirable than others; in generation, we will prefer abduction to find supports for event structures that are directly linguistically realisable. That is, we will not rest until we have shown how to prove all the knowledge structure from a set of assumptions A which are actually justified by the text, without any defeasible inferential links between A and the text. Let's call such assumptions *concrete*. Assumptions concerning discourse relations and event-relations are not concrete; those concerning textual order or the presence of temporal connectives are concrete.

The process of abducing a description of a temporal structure will terminate when all the assumptions are directly realisable; that is: abductive generation is a search for concrete support, and it is not complete until all other supports have been shown to be grounded by concrete ones.

The generation process unfolds as follows. The initial knowledge sources for S consists of three sets: Δ, EC and ET. Δ contains the purpose of the text (we here assume it to be to inform) and also the KB H uses to interpret language—in other words, all the defeasible and indefeasible laws introduced earlier, and the facts H already knows about the eventualities to be described. EC is a pair of sets; the first contains the names of the eventualities, and the second contains all causal and part/whole relations between them. These relations induce a hierarchical structure on the eventualities: the part/whole relation is a subordinating one, and the

text that makes the latter event key relative to joke telling, making it worth explaining.

7 Konolige (1991) proves certain equivalence results between defeasible reasoning and abduction. But these results hold only for those systems of defeasible reasoning that do not validate the Penguin Principle, and hence don't directly apply to CE.

cause/effect relation a coordinating one (cf. Nakhimovsky 1987). *ET* is the set of temporal relations between the eventualities.

The first process *S* undertakes in generation is to use the purpose of the text and *S*'s and *H*'s points of view to determine which eventualities in *EC* are the key ones. This produces an additional hierarchical structure *EK*, which describes which eventualities are key relative to others in the database. As we've mentioned, abduction to *EK* from the *S*'s initial knowledge sources is beyond the scope of this paper, but see Maybury (1991).

Given the knowledge sources Δ, *EC*, *ET* and *EK*, the generation algorithm proceeds as described below. But first, some comments about discourse structure will help elucidate the algorithm. A discourse structure *D* is hierarchical on the clauses: *Explanation* and *Elaboration* are subordinating relations, and *Narration*, *Result* and *Background* are coordinating ones. *D* also fixes the textual order of the clauses via the semantics of discourse structure in Lascarides and Asher (1991a): textual order corresponds to the depth-first left-to-right path through *D*. Hence *D* will entail concrete assumptions about textual order, such as $\langle\alpha, \beta\rangle$, which essentially means "utter α before β". The construction of *D* described below will also fix what eventuality each clause describes; so, for example, *D* will entail concrete assumptions such as $fall(max, e_\alpha)$.

1. Start at the topmost leftmost eventuality in *EC*, and proceed in a depth-first, left-to-right manner on *EC*, abducing discourse assumptions pairwise (via Δ, *EC*, *ET* and *EK*) that will prove the relations in *EC*, *ET* and *EK* between the current pair of eventualities. Do this until all events have been covered. This produces a set *D*, which is the discourse structure.
2. We now proceed in a depth-first, left-to-right manner on *D*. We name the set of discourse relations between the current pair of clauses *CDR*; *CER* names the event relations between the eventualities that these clauses describe.
   (a) We add to the set *Conc* the set of concrete assumptions arising from *D* that concern just the current pair of clauses.
   (b) We do an NMDC on Δ and *Conc*. This produces a set of inferences *Inf*.
      i. If *Inf* includes the relations in *CER* and *CDR*, then either (i′) abduce extra concrete assumptions about the current pair of clauses via Δ, and then go on to (i″); or else directly (i″) go on to next pair in *D* and start at (a); if there is no further pair in *D*, send *Conc* to the surface grammar.
      ii. If *Inf* does not include the relations in *CER* and *CDR*, then abduce on any rule in Δ with a relation from *CER* or *CDR* in the consequent, which will add a *further* concrete assumption to the set of existing concrete assumptions in *Conc*. Add these further assumptions to *Conc*, and go back to (b).

Proceeding depth first left to right on *EC* ensures that each abduction is on a pair of events that are in temporal proximity. The output of each abduction is a discourse relation between the clauses. Given the constraints on which clauses in a text can be related by a discourse relation—the so-called openness constraints—the clauses must appear in textual proximity. Hence temporal proximity induces textual proximity.

Sibun (1992) argues in favour of a related 'localist' approach for generating coherent spatial descriptions.[8]

The KB used in the NMDC does not contain all of D, nor any of EC, EK or ET. This is because what the NMDC is trying to prove are the discourse relation for the clause currently being generated, and assertions about its event structure. The fact that these are provable from a KB containing them as premises would be trivial. We would fail to expose the contextual side effects of using default rules.

Notice that the NMDC is nonmonotonic, and the inferences drawn are sensitive to *all* the premises in the relevant KB, including in particular the concrete assumptions for clauses in D that have already been through the abduction process. Abduction for the clause currently being processed is therefore contextually bound, a fact discussed in more detail in the final section.

6 Worked Examples

We now consider two very simple examples in formal detail. These illustrate some general properties of the process we propose, and also expose some of its limitations. We are primarily concerned with the specification of the clausal semantics, textual ordering, and connectives, including cue phrases. So, while basically glossing over the selection of key events, let us consider in detail what is abduced after that.

6.1 The Push and the Fall

Suppose that the event structure to be described consists of two events—John pushing Max and Max falling—and a relation that the push caused the fall (and so preceded it). Thus EC is $\{\{fall(max, e_1), push(john, max, e_2)\}, \{cause(e_2, e_1)\}\}$, and ET is $\{e_2 \prec e_1\}$. The first task is abduce EK, in accordance with the purposes of the discourse.

Suppose that resulting EK is $\{key(e_1, e_2)\}$, so falling is the key. Let α be the clause that is to describe Max falling, and β the clause describing John's pushing him. Then abduction on Δ, EC, ET and EK will yield the discourse structure $Explanation(\alpha, \beta)$, where $fall(max, e_\alpha)$ and $push(john, max, e_\beta)$ must be true of whatever linguistic realisation we eventually choose for α and β. The rules abduced on are Key Event of Explanation and Axiom on Explanation; in this case, no other choice of rules in Δ will do. $Explanation(\alpha, \beta)$ fixes the textual order; α is to be uttered before β.

Now we construct $Conc$: it is $\{\langle\alpha, \beta\rangle, fall(max, e_\alpha), push(john, max, e_\beta)\}$. The NMDC on Δ and $Conc$ yields the inferences Inf, which are $Explanation(\alpha, \beta)$ and $cause(e_\beta, e_\alpha)$. Thus Inf includes CER and CDR. So the NMDC in this case has shown that no additional assumptions need to be made about α and β, although of course they may be made if S wishes: either step (i′) or step (i″) could be now taken.

8 However, unlike Sibun, we concur with Hovy (1988), and assume that calculating discourse structure is an essential part of text generation: there are constraints on which clauses can be related in a coherent discourse; generation should be informed by these constraints; and discourse structure encodes them (cf. also Polanyi 1985, Webber 1991).

Suppose we take the latter. Then *Conc* is sent to the surface grammar for realisation, resulting in the reliable text (7).[9]

(7) Max fell. John pushed him.

Alternatively, suppose we take step (i′).[10] Then abduction on *any* law, defeasible or indefeasible, with elements in *CER* or *CDR* in the consequent is permitted. There may be a number of rules in H's KB that could be used. For example, Non-evidential Because, which is glossed as 'if α and β are discourse related and β features the connective *because*, then $cause(e_\beta, e_\alpha)$ normally holds'.

– `Non-evidential Because`
  $\langle\alpha,\beta\rangle \wedge because(\beta) > cause(e_\beta, e_\alpha)$

Abducing on Non-evidential Because will entail that $because(\beta)$ must be made true; this is added to *Conc*. So the text generated would be (10) instead of (7).

(10) Max fell because John pushed him.

We don't address here the problem of choosing between (7) and (10); but we have fulfilled a task that's necessary to making such choices. We have spelt out in detail exactly what temporal information can remain implicit without misleading H.

6.2 Exception to the Causal Law

The NMDC that preceded abduction on the discourse structure in the previous example may seem superfluous. It didn't constrain the rules in H's KB that could be used for further abduction. In general, however, the NMDC is required, because of the context sensitivity of inference involving defeasible laws. We now consider a simple example that demonstrates the need for the NMDC.

Suppose that the event structure to be described involves the two events of Max falling and John pushing him once more, but this time the temporal relationship between them is different: the falling *immediately precedes* the pushing. This is represented by EC_1 and ET_1, respectively $\{\{fall(max, e_1), push(john, max, e_2)\}, \{\}\}$ and $\{e_1 \prec e_2\}$. Suppose furthermore that the discourse structure arising from abduction on Δ, EC and ET is $Narration(\alpha, \beta)$, where α describes the fall and β the pushing. Now assume for the sake of argument that there is no NMDC on Δ and *Conc*, which is $\{\langle\alpha,\beta\rangle, fall(max, e_\alpha), push(john, max, e_\beta)\}$. Then nothing will block S from using abduction on Narration. Suppose S uses abduction on Narration, *and nothing more*. Then the text generated will be (7).

(7) Max fell. John pushed him.

9 In Oberlander and Lascarides (1991), we discuss how this model of generation indicates why the occurrence of such 'unmarked' temporal reversals is rare in naturally occurring data. It should be stressed that we do not claim (7) is always optimal, although in certain contexts it may be (as urged in Caenepeel 1991).

10 Whether (i′) is used is a question of textual optimality; an issue we don't address here.

But given that the Push Causal Law and Explanation also form part of H's KB, H won't interpret (7) in the way S would hope. In this case, although the antecedent to Narration is verified by uttering (7), the consequent of this defeasible law won't be inferred by H because of the other premises in the KB. Of course, it would be possible for S to generate (7), test how H would interpret it, and then debug it if necessary. But it would be more efficient if the abduction process were guided so that S didn't generate (7) in the first place. The NMDC is there to play this role.

Consider now how the NMDC will block abduction on Narration in this example. The NMDC yields the nonmonotonic inferences from Δ and $Conc$, which as before are $cause(e_\beta, e_\alpha)$ and $Explanation(\alpha, \beta)$. These don't include CER or CDR; they are contrary to what S requires. So according to the above algorithm, abduction on rules that provide *additional* (rather than different) assumptions about the linguistic structure of the clauses must be made, in order to ensure that H will infer $Narration(\alpha, \beta)$. But abduction on Narration would not involve the assumption of additional linguistic structure, since its antecedent is already in $Conc$. So the NMDC precludes abduction on Narration in this case.

S must find another rule whose consequent is $Narration(\alpha, \beta)$, and whose antecedent must be proved by making additional assumptions about linguistic structure. Cue And Then is just such a law:

- `Cue And Then`
  $\langle \alpha, \beta \rangle \wedge andthen(\beta) \rightarrow Narration(\alpha, \beta)$

Abducing on this rule will involve the additional assumption that the clause β features the phrase *and then*. Hence, this path would ultimately yield the following, reliable text, the NMDC having ensured that (7) was never generated, in this context:

(11) Max fell and then John pushed him.

This example may have an air of artificiality: under what circumstances would we want to describe EC_1 and ET_1? The example (13) below provides such a circumstance and also demonstrates the degree of context sensitivity featured in the generation process, arising from the NMDC.

7 Related Work

There is considerable interest in recruiting abduction for natural language interpretation. Advantages and disadvantages of various approaches are investigated in work by Hobbs *et al* (1988, 1990), Charniak and Goldman (1988, 1989), Norvig and Wilensky (1990), Appelt and Pollack (1991). Abduction in generation is considered by Hobbs *et al* (1990:26–28) and by Ishizaki and Ishikawa (1991). We here consider briefly two points of comparison with Norvig and Wilensky and Hobbs *et al*.

7.1 Commensurability and Optimality

Norvig and Wilensky (1990) contrast several abductive interpretation methods. They concede the attraction of measure- (or probability-) based approaches, noting that these enable information from various components of the grammar to be handled in a uniform framework. They go on to point out, however, that a single dimension of measure is not enough to represent cost, goodness of explanation and heuristic search control altogether. In this respect, we should observe that we do not have any measure at all available for representing these factors. We have here been concerned only with the generation of *adequate*, rather than *optimal* texts. To extend our account to deal with optimality, a mechanism—such as weighted abduction—taking into account more complex preferences could be exploited. On the other hand, even without a measure, we do capture part of the appeal of commensurable models: all the information manipulated, from event structures to textual assumptions, is represented in a uniform context-sensitive logical framework.

7.2 Weighted Abduction and Context Sensitivity

Hobbs *et al* (1990) have discussed the use of *weighted abduction*, which allows algorithmic choice between competing sets of assumptions. Weighting guides abduction via assertions about the relative cost of assuming each predicate featured in the antecedent. Costs are assigned to antecedents in a global fashion and are represented by superscripts. For example in (12), the cost of assuming p_1 and p_2 in order to prove q is $\omega_1 c + \omega_2 c$, where c is the cost of assuming q.

(12) $p_1^{\omega_1} \wedge p_2^{\omega_2} \rightarrow q$

We prefer to assume antecedents that prove the result and that have the lowest cost. Such a framework could thus provide a mechanism for preferring one adequate utterance over another.

On the other hand, weighted abduction has problems with the context-sensitive nature of language processing; in this respect, our approach has certain advantages. Consider again the NMDC with the event structure defined by EC_1 and ET_1. It might seem that weighted abduction would obviate the need for a deductive check. Perhaps we could weight the antedecent of the rule for Cue And Then (which is of the form $p_1 \wedge p_2$) so that its assumption cost is lower than that for Narration (of the form p_1). With such a scheme, (11) would certainly be the preferred description of EC_1 and ET_1. However, even if we could arrange this, we wouldn't want to. Under the suggested algorithm, Narration's antecedent could never be assumed at zero cost. The context-free nature of the weighting then means that Narration, as the higher cost rule, would *never* be preferred. And this would mean that the relatively natural text (1) would never be generated.

In fact, an NMDC is ineliminable in a system which permits abduction over a set of rules which incorporate default information. Hobbs *et al* (1990:47–48) observe that incorporating 'et cetera' predicates into their scheme means that certain literals can only be assumed once it has been established that no contradiction results. They further observe that the difficulty of handling the side-effects of assumptions 'may be fundamental, resulting from the fact that the abduction scheme attempts to make global judgments on the basis of strictly local information' [p48].

In this connection, we wish to note two points. First, the use of a nonmonotonic logical framework allows for a relatively graceful treatment of contextual side-effects.

Secondly, we can indicate the means by which we incorporate discourse-contextual factors by turning once more to example (7) with event structure defined by EC_1 and ET_1. Generating descriptions of event pairs *in vacuo* may seem a little strange; but generating them in context seems more natural.

The NMDC exploits the context-sensitivity of nonmonotonic inference, permitting a logical explanation of how linguistic and extra-linguistic context influence what can remain reliably implicit in a clause being generated, and hence when that text can be laconic. This property is, of course, most significant in extended text. By adding rules in H's KB concerning how the preceding discourse structure affects interpretation, we can supply an explanation of why (7) is unreliable in the 'null' context we considered in the previous section, but reliable below:

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

Briefly, the following element of H's KB distinguishes the 'null' context case from (13). We assume H's KB contains a rule of LK that one doesn't describe events in the order: cause, effect, further cause:

- **`Maintain Causal Trajectory`**
  $R(\gamma, \alpha) \land cause(e_\gamma, e_\alpha) \land \langle \alpha, \beta \rangle > \neg cause(e_\beta, e_\alpha)$

For (13), consider γ to be the clause describing Max feeling the sharp blow, and α and β to be the clauses describing the falling and pushing, as before.

The logically unrelated antecedents of Maintain Causal Trajectory and the Push Causal Law are both satisfied in the NMDC; their consequents conflict. Under these circumstances, akin to the standard Nixon Diamond in default logic, CE produces no conclusion. Now, there is arguably a principle of *inertia* in the interpretation of text: if all else fails, assume that the discourse relation for the current sentence is the same as those used already in the discourse. A version of the principle is formally represented in Lascarides, Asher and Oberlander (1992); informally, its says that if you derive a Nixon Diamond of irresolvable conflict when attempting discourse attachment, you should assume the discourse relation is that which was used previously. If the principle is available in processing (13), the preceding discourse relation—in this case *Narration*—will be inferred between α and β. In contrast with the previous example, the NMDC will confirm that the discourse structure is compatible with what S wishes to prove. With the same temporal structure EC_1 and ET_1, the textual pair in (7) is reliable in the context of (13), but not in the 'null' context.

8 Conclusion

A formal model of the temporal implicatures underlying discourse can be incorporated into an NL generation process. By exploiting a nonmonotonic logic, we can extend DRT with a theory of discourse relations, and by employing abduction within this framework, we can show how to generate laconic discourse. A certain amount of nonmonotonic deduction must be interleaved with this abduction. We would argue that this is a price worth paying, if we want to generate texts like (1), where the descriptive order of the events guides the implicatures drawn.[11] Furthermore, although there is a cost, there is also a gain: abduction becomes context sensitive, and so we are able to explain why, in describing the same event structure, generation will yield different linguistic structures in different discourse contexts. The model's most significant current limitation is that it only provides adequate discourses, rather than optimal ones; this is therefore our current area of investigation.

References

Appelt, D. E. and Pollack, M. E.: Weighted Abduction for Plan Ascription. User Modeling and User-Adapted Interaction **1** (4) (1991)

Asher, N. and Morreau, M.: Common Sense Entailment: A Modal Theory of Nonmonotonic Reasoning. Proceedings of the 12th International Joint Conference on Artificial Intelligence. Sydney Australia, August (1991)

Asher, N.: Reference to Abstract Objects in English: A Philosophical Semantics for Natural Language Metaphysics. Book manuscript submitted to Kluwer Academic Publishers (1991)

Bach, E.: The algebra of events. Linguistics and Philosophy **9** (1986) 5-16

Caenepeel, M.: Event Structure versus Discourse Coherence. Proceedings of the Workshop on Discourse Coherence. University of Edinburgh, April (1991)

Charniak, E. and Goldman, R: A Logic for Semantic Interpretation. Proceedings of the 26th Annual Meeting of the Association for Computational Linguistics, State University of New York at Buffalo N.Y., 7-10 June (1988) 87-94

Charniak, E. and Golman, R.: Plan Recognition in Stories and in Life. Proceedings of the Uncertainty Workshop, International Joint Conference in Artificial Intelligence, Detroit USA (1989)

Grice, H. P.: Logic and Conversation. In Cole, P. and Morgan, J. L. (eds.) *Syntax and Semantics*, Volume 3: *Speech Acts* 41-58. New York: Academic Press (1975)

Hobbs, J. R.: Coherence and Coreference. Cognitive Science **3** (1979) 67-90.

Hobbs, J. R.: On the Coherence and Structure of Discourse. Report No. CSLI-85-37, Center for the Study of Language and Information, October (1985)

Hobbs, J. R., Stickel, M., Martin, P. and Edwards, D.: Interpretation as Abduction. Proceedings of the 26th Annual Meeting of the Association for Computational Linguistics, State University of New York at Buffalo, N.Y., 7-10 June (1988) 95-103

Hobbs, J. R., Stickel, M., Appelt, D. and Martin, P.: Interpretation as Abduction. Technical Note No. 499, Artificial Intelligence Center, SRI International, Menlo Park, December (1990)

Hovy, E. H.: Planning coherent multisentential text. Proceedings of the 26th Annual Meeting of the Association for Computational Linguistics. State University of New York at Buffalo, N.Y., 7-10 June (1988) 163-169

Ishizaki, M. and Ishikawa, A.: Taking 'Generation As Abduction' Seriously. Submitted to *COLING92*.

Joshi, A. K., Webber, B. L. and Weischedel, R. M.: Some Aspects of Default Reasoning in Interactive Discourse. LINC LAB 24 No. MS-CIS-86-27, School of Engineering and Applied Science, University of Pennsylvania, Philadelphia, April (1986)

11 Note that changing the descriptive order in text (1) would yield a distinct temporal order of events in interpretation.

Kamp, H.: A theory of truth and semantic representation. In Groenendijk, J. A. G., Janssen, T. M. V. and Stokhof, M. B. J. (eds.) *Formal Methods in the Study of Language* **136** (1981) 277-322. Amsterdam: Mathematical Centre Tracts.

Konolige, K.: Abduction vs. Closure in Causal Theories. Technical Note No. 505, Artificial Intelligence Center, SRI International, Menlo Park (1991)

Lascarides, A. and Asher, N.: Discourse Relations and Defeasible Knowledge. Proceedings of the 29th Annual Meeting of Association for Computational Linguistics, Berkeley, Ca, June (1991a) 55–63

Lascarides, A. and Asher, N.: Discourse Relations and Common Sense Entailment. Research Report HCRC/RP-16, HCRC, University of Edinburgh, January (1991b)

Lascarides, A., Asher, N. and Oberlander, J.: Inferring Discourse Relations in Context. Technical note (1992)

Lascarides, A. and Oberlander, J.: Temporal Coherence and Defeasible Knowledge. Theoretical Linguistics **18** (1992)

Mann, W. and Thompson, S.: Rhetorical Structure Theory: A theory of text organisation. Technical Report ISI/RS-87-190, USC/ISI, June (1987)

Maybury, M.: Planning Multisentential English Text using Communicative Acts. PhD thesis, University of Cambridge (1991)

Moens, M. and Steedman, M. J.: Temporal ontology and temporal reference. Computational Linguistics **14** (1988) 15-28

Nakhimovsky, A.: Temporal Reasoning in Natural Language Understanding: the Temporal Structure of Narrative. Proceedings of the 3rd European Meeting of the Association for Computational Linguistics, Copenhagen (1987) 262-269

Norvig, P. and Wilensky, R.: A Critical Evaluation of Commensurable Abduction Models for Semantic Interpretation. In Karlgren, H. (ed.) *Proceedings of COLING90* (1990) 225-230

Oberlander, J. and Lascarides, A.: Discourse Generation, Temporal Constraints, and Defeasible Reasoning. Paper presented at the AAAI Fall Symposium on Discourse Structure in Natural Language Understanding and Generation, Asilomar, Ca, November (1991)

Obermeier, K.: Temporal Inferences in Medical Texts. Proceedings of the 23rd Annual Meeting of the Association for Computational Linguistics, University of Chicago, Chicago, Illinois, 8-12 July (1985) 9-17

Polanyi, L.: A Theory of Discourse Structure and Discourse Coherence. In Eilfort, W. H., Kroeber, P. D. and Peterson, K. L. (eds.) *Papers from the General Session at the 21st Regional Meeting of the Chicago Linguistics Society*, Chicago, April (1985) 25-27

Reichman-Adar, R.: Extended Person-Machine Interface. Artificial Intelligence **22** (1984) 157-218

Scha, R. and Polanyi, L.: An augmented context free grammar. Proceedings of the 12th International Conference on Computational Linguistics and the 24th Annual Meeting of the Association for Computational Linguistics, Budapest Hungary, 22-27 August (1988) 573-577

Selman, B. and Levesque, H. J.: Abductive and Default Reasoning: A Computational Core. Proceedings of the 8th National Conference on Artificial Intelligence, Boston (1990) 343-348

Sibun, P.: Generating Text without Trees. Computational Intelligence: Special Issue on Natural Language Generation **8** (1992)

Webber, B. L.: Structure and Ostension in the Interpretation of Discourse Deixis. Language and Cognitive Processes **6** (1991) 107-135

This article was processed using the LaTeX macro package with LMAMULT style