1

Laconic discourses and total eclipses: abduction in DICE

JON OBERLANDER
AND ALEX LASCARIDES

1.1 Introduction

The goal of this chapter is demonstrate one particular use of abduction in the processing of natural language discourse. DICE (Discourse In Commonsense Entailment) can be used to model both interpretation and generation. For interpretation, it uses defeasible deduction to determine the discourse structures and event structures underlying multisentential texts. For generation, it uses abduction to build up specifications of texts from underlying event structures. Thus, abduction is the primary mechanism of generation, and has a smaller role to play in interpretation.

To demonstrate how the approach works in generation and interpretation, we introduce two exemplary problems, which involve brevity and accommodation respectively. We then outline a formal model of implicature, and indicate how it is recruited for interpretation by deduction and generation by abduction. We show how brief but accurate discourses can be generated under this model, and then turn to a possible role for abduction in changing the context of interpretation. Finally, we mention one of the differences between this approach, and weighted abduction in Hobbs, Stickel, Martin, and Edwards (1988) and Hobbs, Stickel, Appelt, and Martin (1990).

1.2 Two exemplary problems

1.2.1 The problem of laconic discourse

It is a commonplace that brevity is the soul of wit. Natural language generation systems have, on occasion, striven for it. In developing the approach to discourse structuring in DICE, our primary focus has been on domains where temporal and causal relations between eventualities have to be communicated briefly but accurately. That is to say, we have been investigating methods whereby temporal information can be left unstated but inferrable. We call the texts which possess this putatively desirable property *laconic* discourses.

Let's consider in turn three questions. First, *are* laconic discourses desirable? If they are, do they introduce any risks? And finally, what basic mechanism do we need if we are to deal with the risk effectively?

Are laconic discourses a good thing? Contrast (1) and (2), which are supposed to describe the same course of events.

(1) 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.

(2) 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) seems somewhat more direct than the relatively verbose (1), and is perhaps preferable as a result. Part of the difference is that (1) makes all the temporal relations explicit; there is, for example, no doubt as to the temporal relation between Max's entry and John's greeting. (2), by contrast, leaves this temporal relation implicit. The entry is mentioned first, and the greeting mentioned second, but no explicit ordering is mentioned. In the current genre, the order of mention generally correlates with the order of events, and we would take the ordering of these events to be *implicitly* conveyed by text (2), and explicitly by (1).

It is tempting to generalise from such cases, and suggest that brevity is invariably preferable to verbosity. However, brevity has its risks, as we can see when we consider (3).

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

What does this mean? On the one hand, knowledge of the correlation between order of mention and order of events would lead a hearer/reader H to believe that the switching off preceded Judy's statement. On the other hand, general world knowledge about heating mechanisms and cooperative behaviour might well lead H to believe that the switching off was a consequence of Judy's statement. Thus, the knowledge H exploits in interpretation will have a crucial effect on whether (3) is taken to be a simple narrative, or a causal explanation. The problem is clear: if a speaker/writer S wishes to exploit brevity, she runs the risk of misleading H, unless she is sure what knowledge H will bring to bear. We may conclude, then, that leaving temporal progression unmarked is natural but risky.

What mechanism does S require if she is to avoid risky utterances like (3) while still generating desirable realizations such as (2)? We assume that one way for S to test for misleading utterances is for her to reason about the implicatures the hearer H will infer, given the discourse context and likely background knowledge. We conclude that we therefore require a formal model of implicature to inform a speaker or writer's generation process.

1.2.2 The problem of total eclipses

It has been observed in the literature (Harkness 1985; Hamann 1989) that some discourses set up an expectation of a contingent link between two eventualities, and that there are differing ways of responding to utterances which frustrate that expectation.

In particular, Hamann has observed that temporal connectives, such as *before*, *after* and *when* appear to have this effect:

> The [temporal clause] event not only informs of time location but has a narrative relevance in its own right . . . only some of [the many possible co-locating events] make sense in context . . . total unconnectedness . . . leads to a purely temporal reading, or to some profound statement about the way of the world – brought about by the interpreter's willingness to see a narrative connection (Hamann 1989, 43)

On her account, (4a) and (4b) can be taken to describe the same pair of events, but (4b), by yoking the events together with an *after*-clause, carries an additional implicature that the leaving was contingent upon the homecoming – perhaps by being caused by it. The fact that *after* leads to such an expectation is involved in the explanation of the oddness of (5).[1]

(4) a. John came home. Mary left.

b. Mary left after John came home.

(5) ? The moon eclipsed the sun after John came home.

The explanation goes along the following lines. Sentences of the form *A after B* usually link eventualities which can cause one another, where B is a clause, but not when it is an adverb. However, we also have a strong prior intuition that eclipses are not contingent on human activities. When a conflict like this arises, we either (i) derive a purely temporal reading; or (ii) derive a 'profound' reading. Following strategy (i), we would read (5) as simply locating the solar eclipse at a time after the time of the homecoming; here, the clause *after John came home* is functioning just as *after 7pm* would function. Following strategy (ii), we would read (5) as pointing out a new and

[1]Hamann's original example involved the event of Jupiter's completing its orbit, and the version of this chapter delivered at the Alghero meeting used this same example. It was pointed out there, however, that the Jupiter event is somewhat spurious, there being no begin- and end-points to an elliptical orbit. We have therefore varied the example to involve an eclipse, a real (and dateable) event.

surprising fact about the world; there *is* a human whose movements influence planetary motion.

Is there an assumption of connection in *after*-sentences? Is Hamann's account the correct one? Even if it is not, an account which seeks to explain implicit assumptions about temporal structure ought to be able to specify what is going on in a case such as (5). We therefore conclude that we require a formal model of implicature to inform a hearer or reader's interpretation process.

1.3 The formal model of implicature

In recent work, we have developed DICE as a formal model of some of the implicatures relating to discourse and temporal structure (Lascarides and Asher 1991, (1992), (1993); Lascarides and Oberlander 1992a, (1992b); Lascarides, Asher, and Oberlander 1992, Oberlander and Lascarides 1992).

In this section, we first define more precisely the concept of laconic discourse, before introducing our representation of discourse structure. We then exemplify the varieties of knowledge which, encoded as defeasible rules, influence discourse processing in this model; these include Gricean pragmatic maxims (cf. Grice 1975), causal knowledge, and discourse contextual effects. We then indicate how the processes of interpretation and generation are logically modelled; in the former, an underlying logic uses the rules to construct discourse structures and temporal structures from NL text; in the latter, the same rules are used under a different inference regime, to construct discourse structures and NL text from temporal structures.

1.3.1 Temporal constraints and laconic discourse

First, then, let us be more precise about what makes a text laconic. We consider a set C of relations between eventualities (events and states), and say that it includes the following: a causal relation, a part–whole relation, an overlap relation, and an immediate precedence relation. Since all of these except overlap are asymmetric, the set actually specifies seven possible relationships between a pair of eventualities e_1 and e_2. Three points about this set should be noted. First, causal and part–whole relations are – in some sense – more ontologically committed than the overlap and precedence relations; the former *require* eventualities in a way that the latter need not. Secondly, the relation of immediate precedence is taken to be a temporal precedence compatible with a causal or part–whole or whole–part relation. Finally, the bias towards relations between events distinguishes the current approach from those, such as Allen's theory (Allen 1984), which are based on relations between time intervals; one effect of taking eventualities to be basic is that we need a relatively small set of possible relations between them.

Now, let us specify the notions of *temporal coherence* and *reliability*, which we first discussed in Lascarides and Oberlander (1992b):[2]

[2] There we also introduced the notion of *temporal precision*, which we do not touch upon here. Note that these clauses should be read as full definitions.

- Temporal Coherence
  A text is *temporally coherent* if its reader can infer that at least one of the relations in C holds between the eventualities described in the sentences.

- Temporal Reliability
  A text is *temporally reliable* if one of the relations in C which its reader infers to hold does in fact hold between the eventualities described in the sentences.

Although we have characterised reliability in terms of what *actually* holds, what matters is, of course, what the text's speaker/writer S believes to hold – or what S intends H to come to believe holds. We don't here distinguish these possibilities, because we are naively assuming that a writer sets out to transfer information. This is, of course, a simplification, but it is one which is tenable in the simple cases we examine.

We are now in a position to say that discourse is *laconic* when it is: (i) temporally reliable; (ii) temporally coherent; and (iii) not fully explicit with respect to the set C. Full explicitness is only comparative, of course. For current purposes, we will say that a text is not fully explicit with respect to C when an expression specifying one of the relations could be added to the text without affecting the text's reliability or coherence. Hence, we can say that (1) could result from adding expressions in this way to (2). For this reason, (2), but not (1), is laconic.

(1) 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.

(2) 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.

1.3.2 Discourse structure

Here we introduce three key elements of our approach to discourse structure: coherence relations combined with DRT; structural openness, and the concept of key events.

Our basic approach is to follow (Asher 1992), and augment Kamp's discourse representation theory with a number of Hobbsian discourse coherence relations (Kamp 1981; Hobbs 1979; Hobbs 1985; Polanyi 1985; Mann and Thompson 1987). The coherence relations we have examined in most detail are *Narration*, *Explanation*, *Elaboration*, *Result* and *Background*; these are the discourse relations most relevant to temporal structure, and the only ones we discuss below. On this representation, discourse representation structures DRSs are *segmented*, so that they include both traditional DRSs, and a set of discourse relations between DRSs. (6–9) provide intuitive examples of four of the relations, here (as in all the examples we discuss) holding between DRSs corresponding to simple clauses.

(6) Max stood up. John greeted him.
Narration(α, β)

(7) Max fell. John pushed him.
Explanation(α, β)

(8) John pushed Max. He fell
Result(α, β)

(9) Max opened the door. The room was dark.
Background(α, β)

The different discourse relations are related to differing temporal-causal relations. (6) describes two successive events. (7), by contrast, introduces one event, and then describes another, which probably caused it (and thus helps explain it). (8) is rather like (6), except that the connection between the two events is a little stronger: it's not merely temporal succession, but more like a causal relation. (9) is also a little like (6), except that this time, we're dealing with an event-state pair, rather than an event-event pair. Intuitively at least, the room was already dark at the time of door opening, and the second clause functions as background for the first, rather than describing something which follows on from it.

Structural openness

The discourse relations we have introduced determine a hierarchical structure upon text, (cf. Grosz and Sidner 1986, Scha and Polanyi 1988). The subordinating relations are *Elaboration* and *Explanation*; the others are coordinating. This hierarchical structure constrains where the DRSs of successive sentences can be attached. Essentially, the only attachment sites are those provided by the previous clause or the clauses to which it is subordinate. Thus, we maintain a 'right frontier' analysis (cf. Webber 1991).

Topical keyness

We wish to acknowledge that some of the events mentioned in a story or description are of greater importance than others. A natural way of talking about this relative importance is in terms such as 'what the discourse is about', and so the concept of a discourse topic might well seem relevant. However, topic is a big topic, so we will steer round it here, at the price of introducing a new item to our working vocabulary: keyness (but cf. Obermeier 1985). We will use $key(e_1, e_2)$ to say that e_1 is a key event relative to e_2. Keyness constrains both event and discourse relations. In the first case, $key(e_1, e_2) \rightarrow C(e_1, e_2)$; and in the second, we have the following rules:

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

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

1.3.3 The rules

So, we have a hierarchical discourse structure, containing DRSs connected by coherence relations. How are the coherence relations computed? We view this computation as a piece of reasoning, so that in interpretation, we reason about updating the discourse structure built so far, exploiting H's knowledge base (KB), which contains specific facts and defeasible and indefeasible rules. Translated into a notation, we say that we must compute the value of the function $\langle \tau, \alpha, \beta \rangle$, where β is to be attached to α with a discourse relation, and α is part of the discourse structure τ already. Much of our interest in interpretation and generation lies in investigating the interactions between the differing types of knowledge involves in the reasoning about the function $\langle \tau, \alpha, \beta \rangle$. It seems that a considerable fraction of this knowledge can be described as defeasible, or contextual, or nonmonotonic, in the sense that it can be overridden by additional information. We notate defeasible conditional rules in the following way: $\phi > \psi$ means 'If ϕ then normally ψ', or 'In the absence of information to the contrary, ϕ implies ψ'.

First of all, a KB contains rules which encode various preferences that language users seem to hold. Take the notion that, other things being equal, textual order gives some information about the order of events:

- **Narration**
  $\langle \tau, \alpha, \beta \rangle > \mathit{Narration}(\alpha, \beta)$
- **Axiom on Narration**
  $\mathit{Narration}(\alpha, \beta) \rightarrow e_\alpha \prec e_\beta$

There's both a defeasible preference for *Narration* (in the absence of any other information), and an indefeasible rule which tells us that *if Narration* holds between two clauses, *then* the first-mentioned event immediately precedes the second.

On the other hand, there is also the notion that world knowledge – particularly knowledge about its causal structure – can influence our interpretations:

- **Explanation**
  $\langle \tau, \alpha, \beta \rangle \wedge cause(e_\beta, e_\alpha) > \mathit{Explanation}(\alpha, \beta)$
- **Axiom on Explanation**
  $\mathit{Explanation}(\alpha, \beta) \rightarrow \neg e_\alpha \prec e_\beta$

The defeasible rule here states that, if we believe that two clauses are discourse related, and that the events they describe are such that the second-mentioned caused the first, then (in the absence of information to the contrary), we will also believe that there's an *Explanation* relation between the two clauses. The indefeasible rule tells us that if the latter relation holds, then it can't be the case that the first mentioned event preceded the second.

Now, such Gricean-style maxims interact with other preferences in a knowledge base; in particular, information about the causal structure of the world can ultimately make the difference between a text's functioning as a narrative, and as a causal explanation:

- **Push Causal Law**
  $\langle \tau, \alpha, \beta \rangle \wedge fall(m, e_\alpha) \wedge push(j, m, e_\beta) > cause(e_\beta, e_\alpha)$
- **Causes Precede Effects**
  $cause(e_1, e_2) \rightarrow \neg e_2 \prec e_1$

The defeasible rule here says that if we believe that two clauses are discourse related, and we also know that the first-mentioned event is a Max falling event, and that the second is a pushing by John of Max, then we will have a preference for believing that the pushing caused the falling. The indefeasible rule encodes the metaphysical principle that effects never precede their causes. The defeasible rule is a slightly unusual creature, in the sense that it folds together information about possible discourses and the world. This is partly a matter of notational convenience: the knowledge could kept more rigorously separate, but it would be at the price of requiring a regular piece of inference to bring them back together. The example causal rule is also arguably over-specific; we use it here because it's relevant to the simple examples we are discussing; but it's clear that in a working system, it would be subsumed under a more general causal preference, making no reference to the individuals we are currently discussing.

In a case where H can draw the conclusion of the Push Causal Law, they will also be able to draw the conclusion of the Explanation rule; the conclusion of this rule conflicts with that of the Narration rule. What happens under such circumstances will depends upon the inference regime used to manipulate the representations we have sketched.

1.3.4 The logic

In order to deal with the interactions between the various pieces of defeasible and indefeasible knowledge, we need an inference regime. And for the interactions to be dealt with in a principled way, we require a formal system which effectively models the nonmonotonic nature of the interactions. We have used Commonsense Entailment (CE), a logic due to Asher and Morreau (1991). We here mention three of the inference patterns captured in CE, and illustrate each with a familiar example from commonsense reasoning.

- `Defeasible Modus Ponens`: $\phi > \psi, \phi \mathrel{|\!\approx} \psi$ e.g. Birds normally fly, Tweety is a bird; so Tweety flies
- `The Penguin Principle`:
  $\phi \rightarrow \psi, \psi > \chi, \phi > \neg\chi, \phi \mathrel{|\!\approx} \neg\chi$
  e.g. Penguins are birds, birds normally fly, penguins normally don't fly, Tweety is a penguin; so Tweety doesn't fly.

If the premises of one rule logically imply the premises of another rule, we can say that the former rule is 'more specific' than the latter. Now suppose the conclusions of the rules conflict; then what the Penguin Principle tells us is that the preferred conclusion is that of the more specific rule.

- `Nixon Diamond`:
  Not: $\phi > \chi, \psi > \neg\chi, \phi, \psi \not\approx \chi$ (or $\neg\chi$)
  e.g. Not: Quakers are pacifists, Republicans are not, Nixon is both a quaker and republican
  $\not\approx$ Nixon is a pacifist/Nixon is a non-pacifist.

The Nixon Diamond is quite similar to the Penguin Principle, except that neither of the rules is more specific; that is, the premises of one rule do not logically imply the premises of the other. Again, suppose the conclusions of the rules conflict; then what the Nixon Diamond tells us is that there is no preferred conclusion.

CE provides an inference regime; encoding various rules related to language interpretation within CE then allows us to probe how the candidate rules interact with one another, and with specific information. We now turn to the deployment of CE in discourse interpretation.

1.4 Interpretation by deduction

The basic idea is that discourse coherence structure is computed via defeasible deduction. Recall the simplest kinds of case, where two events are described, and there is no special background knowledge:

(6) Max stood up. John greeted him.

Of the rules we have introduced, the only one verified in interpreting this text is Narration; by Defeasible Modus Ponens, we therefore conclude *Narration*(α, β), and $e_\alpha \prec e_\beta$. The first sentence is discourse-related to the second by the relation of *Narration*; and the event denoted by the first sentence immediately precedes the event denoted by the second.

Now consider the slightly more complex case where there is background knowledge which seems to lead to conflict with the straightforwardly linguistic knowledge:

(7) Max fell. John pushed him.

This time, two rules are verified: the Push Causal Law, and Narration. By the Penguin Principle, we conclude that $cause(e_\beta, e_\alpha)$, and so the rules now verified are Explanation and Narration. By a second application of the Penguin Principle, we conclude *Explanation*(α, β). So the pushing explains the falling (and must temporally precede it, by the indefeasible rule, Causes Precede Effects).

Finally, consider the slightly odd discourse in (10). Interpreting it will involve some linguistic knowledge about the typical relation between events and states mentioned together, and a piece of world knowledge, which we encode as a causal preference about where races can't be won:

(10) ? Max won the race in record time. He was home with the cup.

- **States Overlap**:
  $\langle \tau, \alpha, \beta \rangle \wedge state(e_\beta) > overlap(e_\alpha, e_\beta)$

- **Win Law**:
  $win(max, race, e_1) \wedge athome(max, e_2) > \neg overlap(e_1, e_2)$

In this case, the premises of three rules are satisfied: States Overlap, the Win Law and Narration. States Overlap is more specific than Narration, and so its conclusion will be preferred by the Penguin Principle; however, the conclusions of States Overlap and the Win Law conflict, and neither is more specific than the other; thus a Nixon Diamond pattern is formed, and there is no conclusion. When there is no conclusion about discourse relations, we must say that the two clauses are *not* discourse related; we can therefore say that we have a case of *local discourse incoherence*. But such local incoherence is not fatal; we can attempt to discourse pop. If there were open attachment sites higher up in the discourse tree, then we could try to attach *John was home with the cup* elsewhere. If we succeeded in doing so, then the discourse would still be coherent as a whole. However, if there is no discourse context prior to (10), then the local incoherence is fatal.

1.5 Generation by abduction

The previous section discussed the use of defeasible deduction for interpretation. Suppose that, when we turn to generation, we wish to use a framework that shares many of the same characteristics with this defeasible deductive framework. Several options are feasible.

First, we could continue to use defeasible deduction, but combine it with a new set of rules. Consider, for example, the rule of Narration; it took us from information concerning textual order to information concerning discourse structure. We could perhaps construct a new rule, taking us from information concerning discourse structure to information concerning textual order. We could then generate by deduction. The major task to accomplish in pursuing this strategy would be the compilation of a new set of rules. A corollary of this method would probably be a strict separation between the rule sets – and hence processing – of interpretation and generation. Generation rules could not be 'visible' to the interpretation process; otherwise, we would infer discourse structure from textual ordering, and then re-infer textual structure from discourse structure.

Secondly, we could continue to use defeasible deduction over the original set of rules. This would happen, for example, if we were to use the interpretation model to check the product of some other generation process; but, of course, the inferences and the rules would not constitute a generation model as such, and the connection between the generator and the tester could be very loose indeed. This is essentially the course we pursued in Oberlander and Lascarides (1992), where the intention was to check for unintended inferences flowing from differences between the speaker's KB, and their model of their audience's KB. Indeed, we continue with a version of this strategy below; but as should become apparent, this time the deductive testing is tightly integrated into the generation model.

Finally, then, we can maintain a common set of rules for interpretation and generation, at the price of switching to a different inference regime for generation. This is the course we pursue here: generation by abduction. We thus exploit defeasible and indefeasible rules, of the forms $p > q$, and $p \rightarrow q$. In interpretation, from p, we concluded q; in generation, from q, we conclude p. The relationship between defeasible deductive systems, and abductive systems is quite complex. There are results indicating the equivalence between certain defeasible deductive systems, and certain abductive systems with essentially indefeasible rules (cf. Konolige 1991). But it's not clear that CE in general constitutes one of the kind of defeasible systems all of whose deductions can be shown to be equivalent to abductions in another system. And anyway, we are here proposing something else again: abduction within a mixed set of defeasible and indefeasible rules. Whether this is equivalent to defeasible deduction in another type of system is an open question, to which we return briefly in section 1.7.

Assuming the position sketched above, the basic process is this: we abduce from event structure to realisable semantic structure via key event structure and discourse structure. It turns out that in a *defeasible* framework, abduction must be checked by deduction, to prevent unwanted side-effects. Since whether q follows from p and $p > q$ depends on context, we must ensure the context we have abduced is of the appropriate kind. We have argued in Lascarides and Oberlander (1992a) that this is a cost, but that the correlative benefit is context-sensitive generation.

1.5.1 Overview of the process

We search for a proof of the event structure via abduction. We continue until the set of assumptions constituting the proof falls within the set $Conc$, the concrete assumptions, which are those that can be made true through linguistic realisation – by saying them. For example, the following propositions all represent concrete assumptions. $\langle \tau, \alpha, \beta \rangle$ means that we say α before β; $fall(m, e_\alpha)$ means that α must describe Max falling; $SP(\alpha)$ means that α must feature a simple past tense; $because(\alpha)$ means that α must feature the word *because*. On the other hand, neither of the following are concrete: $e_\alpha \prec e_\beta$ and *Narration*(α, β).

The relevant inputs to the generation process are: the set EC (actually, a pair of sets) representing the eventualities and their causal structure; the set ET representing the temporal structure; and the set Δ representing H's KB plus S's purpose. For simplicity, we assume that S's purpose is to inform H of the events and states and the constraints on their causal and temporal relations, encoded in EC and ET. Note that the causal relations induce a hierarchical structure on the eventualities: the part/whole relation subordinates, and the cause/effect relation coordinates (cf. Nakhimovsky 1987).

There are three abductive steps: first, from EC, ET and Δ, we abduce the key event structure EK commensurate with S's purpose; secondly, from EC, ET, EK and Δ, we abduce the discourse structure D that will prove the event structures; and thirdly, from D and Δ, we abduce linguistic realisations that will prove D. During this last stage, we must constrain abduction to avoid contextual side effects. A more detailed description is given in Figure 1.1; note that, given a pair of clauses, the set

1. From the purpose of text, construct EK.
2. Depth-first left-to-right on EC, abduce D via Δ, EC, ET and EK.
3. Depth-first left-to-right on D, we build up $Conc$
   (a) Add to $Conc$ the set of concrete assumptions arising from D concerning the current pair of clauses.
   (b) We do a nonmonotonic deductive check (NMDC) on Δ and $Conc$. This produces a set of inferences *Inf*.
      i. If *Inf* includes the relations in *CER* and *CDR*, then go to next pair in D and (3a), or abduce extra concrete assumptions about current pair of clauses and add them to $Conc$.
      ii. If not, then abduce on any rule in Δ with a relation from *CER* or *CDR* in the consequent that will *add* further concrete assumptions about current clause-pair. Then go back to (b).

Figure 1.1: A recipe for abductive discourse generation

of discourse relations between them is *CDR*, and *CER* is the set of event relations between the eventualities that the given clauses describe.

1.5.2 Worked examples: push and fall

To illustrate aspects of this algorithm, let's consider two simple scenarios. In both, we ignore the first stage of the algorithm, in which we derive the set EK; the relation between intentions and textual topics is beyond the scope of the current discussion, although we do have some related comments on the matter in section 1.6. In the first scenario, we have e_1, an eventuality in which Max falls, and e_2, in which John pushes Max, and the latter is known to cause (and precede) the former. We stipulate that the falling is the key event.

EC $\{\{fall(m, e_1), push(j, m, e_2)\}, \{cause(e_2, e_1)\}\}$

ET $\{e_2 \prec e_1\}$

EK $key(e_1, e_2)$

D *Explanation*(α, β), where $fall(m, e_\alpha)$ and $push(j, m, e_\beta)$. D is obtained by abduction on the rules: Explanation and Key Event of Explanation

$Conc$ $\{\langle \tau, \alpha, \beta \rangle, fall(m, e_\alpha), push(j, m, e_\beta)\}$

NMDC $cause(e_\beta, e_\alpha)$ and *Explanation*(α, β)

– Either send $Conc$ to surface grammar:
  Max fell. John pushed him.

– Or abduce some further concrete assumptions:
  $\langle \alpha, \beta \rangle \wedge because(\beta) > cause(e_\beta, e_\alpha)$
  Max fell because John pushed him.

The second scenario involves the same two eventualities, but this time, there is no causal relation between them, and the falling is known to precede the pushing; neither event is a key.

EC $\{\{fall(m, e_1), push(j, m, e_2)\}, \{\}\}$

ET $\{e_1 \prec e_2\}$

EK Neither e_1 nor e_2 key

D *Narration*(α, β), $fall(m, e_\alpha)$ and $push(j, m, e_\beta)$
D is obtained by abduction on the rules: Narration and Key event of Narration

$Conc$ $\{\langle \tau, \alpha, \beta \rangle, fall(m, e_\alpha), push(j, m, e_\beta)\}$

NMDC $cause(e_\beta, e_\alpha)$ and *Explanation*(α, β). The check indicates that the existing set of concrete assumptions will lead to unreliable text. We must therefore add to $Conc$ further assumptions about α and β:

– Narration will not suffice, since it doesn't add any further assumptions to $Conc$; but if we use this rule about *and then*, we will add a further concrete assumption.

– $\langle \tau, \alpha, \beta \rangle \wedge andthen(\beta) \rightarrow$ *Narration*(α, β)
  Max fell and then John pushed him.

This case indicates the crucial need for the nonmonotonic deductive check; if this element of deduction were not interleaved within the general abductive process, there would be no control over the discourse contexts being built up by abduction, and as in this example, a speaker S would find that their texts might be brief, but they would be unreliable. Hence, the NMDC is a necessary part of our abductive generation process. This is particularly important if we further assume that a generator will start out by attempting *not* to use verbose text. In this case, the rules would effectively be ordered by their syntactic complexity, and the process would try to abduce on rules with fewer premises before trying those with extra premises (this is, in fact, what we assumed in the second scenario). Of course, this 'simple rules first' method is open to question. It may make for the briefest texts, but it may also be computationally expensive, since it is likely to lead to a larger number of NMDC-based iterations than other methods (cf. Reiter and Dale 1992 for a discussion of the computational costs involved in generating brief texts).

1.6 Solar eclipses: abduction in interpretation

The algorithm outlined above requires deduction to play a role in a primarily abductive generation process. It's very natural to ask whether abduction has a role to play in the primarily deductive interpretation process. Let us therefore return to the solar eclipse to frame some speculative remarks about abduction in interpretation.

Hamann (1989, 70) observed that in sentences of the form *A after B* there is an implicature of causal connection when B is a clause, but not when it is an adverb. Elsewhere, we have recently suggested that, viewed from a discourse perspective, we might recharacterise the position in the following way (cf. Lascarides and Oberlander [in prep]). There is always an assumption that two adjacent sentences in a discourse are discourse-related; and this means in turn that the eventualities they denote will somehow by temporally or causally related. An *after*-sentence is a presuppositional environment, and hence, when we try to discourse attach a complex sentence of this kind, we will first of all attempt to accommodate the DRS corresponding to the *after*-sentence within the discourse structure; this process of discourse accommodation is just discourse attachment itself. If the attachment succeeds, we will then have to attach the matrix DRS to the new discourse structure. If the attempted attachment of the presupposition fails, then the situation is reversed, and the presupposition must be attached to the matrix, and the matrix then attached to previous discourse structure.

Let us take the cases where a sentence of the form *A after B* is discourse-initial. The practical consequences of the view just outlined are then as follows. We do not predict that there is any assumption of causal connection which is not already present for simple sentences in a discourse. However, we do predict that B's DRS must be attached to a structure in which A's DRS is the only possible attachment site.

(4) a. John came home. Mary left.

b. Mary left after John came home.

(5) ? The moon eclipsed the sun after John came home.

On this account, the reason why a case of 'total unconnectedness', like (5) is odd is because we have to attach the DRS about John's homecoming to a DRS about an eclipse. These two eventualities have no common key event, and hence violate the relevant axiom on the discourse relation of *Narration*. If, as seems plausible, one event is a key for the other, then we could support a discourse relation between them of *Explanation*, where the eclipse is key, and is explained by the homecoming. But this itself would only be possible if we had a piece of background world knowledge which, by permitting us to infer a causal relationship between the eventualities, licenses the inference to the discourse relation. In fact, of course, we lack world knowledge linking eclipses to human activity, so we just don't have the world knowledge to permit a discourse relation between *John came home* and *the moon eclipsed the sun*.[3]

[3] It is, of course, conceivable that we have knowledge about how people who believe in their horoscopes act; but this information would be irrelevant here, since the direction of causal influence has to be *from* humans *to* planets.

Without a discourse relation to link them, the discourse is locally incoherent. We can now see how the two responses to this strange discourse mentioned by (Hamann 1989, 43) can emerge. Firstly, even though the reasoning component has failed to deliver a discourse relation, there is still a temporal relation, given by semantics of *after*. We know that the eclipse started to occur after the point in time at which John came home. Thus, on this interpretation, *after John came home* functions in approximately the same way as *after 7pm*: as a temporal adverbial, locating, but not discourse relating to, the other eventuality.

Secondly, if an interpreter H's desire to find a discourse relation is stronger than their desire to maintain a minimal set of beliefs, then they can, following Lewis (1979) change the score in the language game, by altering the context of interpretation. This is a different, more far-reaching type kind of accommodation than the presupposition accommodation we described above. This is, in a sense, 'incoherence accommodation'; metaphorically, we must pull a new rabbit out of the hat, rather than find a home for an old one. There are essentially two types of trick available: H can change the input data (the string of natural language text), or change the context of interpretation, by *abducing* specific missing knowledge. Assuming that the textual string is given, H can thus make two possible changes to the context of interpretation. H can (i) retrieve a *Narration* relation by assuming that there is a common key event for John's homecoming and the solar eclipse; or else maintain as before that the solar eclipse is a key event, and that an explanation has a key event, and so (ii) retrieve a *Explanation* relation by assuming that John is the kind of person for whom the heavens wait. A more general assumption permitting *Explanation* might be that humans in general affect the stars; but a specific piece of knowledge is probably more conservative.

So, let us conclude that the lack of a discourse relation leads in the first instance to a purely temporal reading; and that if an interpreter 'clamps' a key relation, they can abduce a possible discourse relation; and to permit this discourse relation to follow from the context, they must add a new rule to their knowledge base, which, given two eventualities, states a preference for a determinate causal relationship between them. Arguably, the addition of this new causal hypothesis to a knowledge base is a case of induction, rather than abduction: we have added a rule, not a fact. Nonetheless, its addition is triggered by an abductive line of reasoning, from key event structure, to discourse structure, and thence to a specification of missing world knowledge. Thus, the profundity reading discussed by Hamann is to be explained in terms of the acquisition of a new (and surprising) piece of knowledge about the causal structure of the world.

1.7 Conclusions

In Lascarides and Oberlander (1992a, 178–80), we contrasted the current abductive approach with that employed by Hobbs et al. (Hobbs, Stickel, Martin, and Edwards 1988; Hobbs, Stickel, Appelt, and Martin 1990) for both generation and interpretation. Their *weighted abduction* system offers a method whereby we can constrain which rules are abduced on, when there's a choice. We argued that weighted rules would augment, rather than replace, the NMDC. In Lascarides, Asher, and Oberlander (1992), we showed how, given a sentence-pair embedded in a discourse, our model permits the prior context to influence which coherence relations are retrieved for the pair. Obviously, we would want generation to respect this effect, but the global nature of weight assignment would not permit discourse generation to exhibit the required context sensitivity. The NMDC allows us to capture it.

In section 1.5, we canvassed a number of ways of possible ways of transforming a defeasible deductive mechanism for interpretation into a mechanism for generation. Our strategy has been to maintain the same mixed set of defeasible and indefeasible rules, but switch to abduction. We observed that this placed the system in a potentially interesting position, in the sense that systems of this type have not yet been shown equivalent to pure defeasible deductive systems. Where, then, does this system stand?

From discussions at Alghero, it emerged that, although there are real differences from Hobbs et al.'s weighted abduction, there are also some striking similarities. First, the differences in the syntactic form of the rules relating textual features to discourse relations should not mask the degree of agreement about the relationships at issue. Secondly, and more particularly, their use of *et cetera* predicates to weaken various rules essentially means that they too are exploiting abduction in a defeasible framework. This has lead them to something like the NMDC, since on their scheme too, certain literals can only be assumed once it has been established that no contradiction results.

To summarise: we focussed on two phenomena which demonstrate aspects of abduction in DICE. The treatment of laconic discourse exhibited DICE's abductive approach to discourse generation; the more speculative discussion of solar eclipses exhibited a possible abductive approach to the accommodation of missing knowledge in interpretation. Both phenomena required an articulated model of implicature: generation primarily uses abduction over that model, but interleaves deduction; interpretation primarily uses deduction over that model, but could also interleave abduction.

The emphasis has been on the cases where the simple abductive scheme can make a difference. Of course, there are many problems which require a more detailed algorithm. For example: we have shown when it's safe to generate laconic discourse; but we have shown nothing about how to choose between non-laconic realizations, and we have said next to nothing about methods for minimizing the number of iterations within the generation process. This issues raise both formal and computational questions, which we hope to address in further research.

Acknowledgements

This chapter reports a combination of new and existing work: the discussion of temporal connectives has not been reported elsewhere, but the work on abductive generation has been discussed before, in Lascarides and Oberlander (1992a). 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.

Bibliography

Allen, J. (1984). Towards a general theory of action and time. *Artificial Intelligence 23*, 123–54.

Asher, N. (1992). *Reference to abstract objects in English: a philosophical semantics for natural language metaphysics*. Dordrecht: Kluwer Academic Publishers.

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

Grice, H. (1975). Logic and conversation. In P. Cole and J. Morgan (Eds.), *Syntax and semantics 3: speech acts*, pp. 41–58. New York: Academic Press.

Grosz, B. and C. Sidner (1986). Attention, intentions and the structure of discourse. *Computational Linguistics 12*(3), 175–204.

Hamann, C. (1989). English temporal clauses in a reference frame model. In A. Schopf (Ed.), *Essays on tensing in English II: time, text and modality*, pp. 31–153.

Harkness, J. (1985). unknown. In A. Schopf (Ed.), *Essays on tensing in English II: time, text and modality*.

Hobbs, J. (1979). Coherence and coreference. *Cognitive Science 3*, 67–90.

Hobbs, J. (1985). On the coherence and structure of discourse. Technical Report CSLI-85-37, CSLI Center for the Study of Language and Information, Stanford.

Hobbs, J., M. Stickel, D. Appelt, and P. Martin (1990). Interpretation as abduction. Technical Report Technical Note 499, SRI International, Menlo Park, CA.

Hobbs, J., M. Stickel, P. Martin, and D. Edwards (1988). Interpretation as abduction. In *Proceedings of the 26th Annual Meeting of the Association for Computational Linguistics*, pp. 95–103. Stanford.

Kamp, H. (1981). A theory of truth and semantic representation. In J. Groenendijk, T. Janssen, and M. Stokhof (Eds.), *Formal methods in the study of language*, pp. 277–322. Amsterdam: Mathematical Centre.

Konolige, K. (1991). Abduction vs. closure in causal theories. Technical Report Technical Note No. 505, Artificial Intelligence Center, SRI International, Menlo Park, CA.

Lascarides, A. and N. Asher (1991). Discourse relations and defeasible knowledge. In *Proceedings of the 29th Annual Meeting of the Association for Computational Linguistics*, pp. 55–63.

Lascarides, A. and N. Asher (1992). A semantics and pragmatics for the past perfect. ms.

Lascarides, A. and N. Asher (1993). Discourse relations and common sense entailment. *Linguistics and Philosophy*.

Lascarides, A., N. Asher, and J. Oberlander (1992, June). Inferring discourse relations in context. In *Proceedings of the 30th Annual Meeting of the Association for Computation Linguistics*, Newark, Delaware, pp. 1–8.

Lascarides, A. and J. Oberlander (1992a). Abducing temporal discourse. In R. Dale, E. Hovy, D. Rösner, and O. Stock (Eds.), *Aspects of automated natural language generation*, pp. 167–182. Berlin: Springer-Verlag.

Lascarides, A. and J. Oberlander (1992b). Temporal coherence and defeasible knowledge. *Theoretical Linguistics 18*.

Lewis, D. (1979). Scorekeeping in the language game. *Journal of Philosophical Logic 8*, 339–359.

Mann, W. and S. Thompson (1987). Rhetorical structure theory: a theory of text organisation. Technical Report Technical Report ISI/RS-87-190, Information Sciences Institute, University of South California.

Nakhimovsky, A. (1987). Temporal reasoning in natural language understanding: the temporal structure of narrative. In *Proceedings of the 3rd European Meeting of the Association for Computational Linguistics*, pp. 262–269.

Oberlander, J. and A. Lascarides (1992, July). Preventing false temporal implicature: interactive defaults for text generation. In *Proceedings of the 14th International Conference on Computational Linguistics*, pp. 721–7. Nantes.

Obermeier, K. (1985, July). Temporal inferences in medical texts. In *Proceedings of the 23rd Annual Meeting of the Association for Computational Linguistics*, pp. 9–17. Chicago, Illinois.

Polanyi, L. (1985, April). A theory of discourse structure and discourse coherence. In W. Eilfort, P. Kroeber, and K. Peterson (Eds.), *Papers from the General Session at the 21st Regional Meeting of the Chicago Linguistics Society*, Chicago, pp. 25–27.

Reiter, E. and R. Dale (1992, July). A fast algorithm for the generation of referring expressions. In *Proceedings of the 14th International Conference on Computational Linguistics*, Nantes, France, pp. 232–238.

Scha, R. and L. Polanyi (1988, August). An augmented context-free grammar. In *Proceedings of the 12th International Conference on Computational Linguistics and the 24th Annual Meeting of the Association for Computational Linguistics, COLING 88*, pp. 573–7. Budapest, Hungary.

Webber, B. (1991). Structure and ostension in the interpretation of discourse deixis. *Language and Cognitive Processes 6*, 107–135.