

INFERRING DISCOURSE RELATIONS IN CONTEXT*

Alex Lascarides

Human Communication
Research Centre,
University of Edinburgh,
2 Buccleuch Place, Edinburgh
alex@cogsci.ed.ac.uk

Nicholas Asher

Center for Cognitive Science,
University of Texas,
Austin, Texas 78712
asher@cgs.utexas.edu

Jon Oberlander

Human Communication
Research Centre,
University of Edinburgh,
2 Buccleuch Place, Edinburgh
jon@cogsci.ed.ac.uk

Abstract

We investigate various contextual effects on text interpretation, and account for them by providing contextual constraints in a logical theory of text interpretation. On the basis of the way these constraints interact with the other knowledge sources, we draw some general conclusions about the role of domain-specific information, top-down and bottom-up discourse information flow, and the usefulness of formalisation in discourse theory.

Introduction: Time Switching and Amelioration

Two essential parts of discourse interpretation involve (i) determining the rhetorical role each sentence plays in the text; and (ii) determining the temporal relations between the events described. Preceding discourse context has significant effects on both of these aspects of interpretation. For example, text (1) *in vacuo* may be a non-iconic explanation; the pushing caused the falling and so explains why Max fell. But the same pair of sentences may receive an iconic, narrative interpretation in the discourse context provided by (2): John takes advantage of Max's vulnerability while he is lying the ground, to push him over the edge of the cliff.

- (1) Max fell. John pushed him.
- (2) John and Max came to the cliff's edge. John applied a sharp blow to the back of Max's neck. *Max fell. John pushed him.* Max rolled over the edge of the cliff.

¹The 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 two anonymous reviewers for their helpful comments.

Moreover, the text in (3) *in vacuo* is incoherent, but becomes coherent in (4)'s context.

- (3) ?Max won the race in record time. He was home with the cup.
- (4) Max got up early yesterday. He had a little bite to eat. He had a light workout. He started the tournament in good form. *He won the race in record time. He was home with the cup.* He celebrated until late into the evening.

So we can see that discourse context can *time switch* our interpretation of sentence pairs, (cf. (1) and (2)); and it can *ameliorate* it, (cf. (4)'s improvement of (3)). The purpose of this paper is two-fold: we attempt to capture formally these aspects of discourse context's impact on clausal attachment; and in the process, we assess whether the structure of the domain being described might be sufficient alone to account for the phenomena.

Of course, the idea that discourse context constrains the discourse role assigned to the current clause is by no means new. Reference resolution is influenced by discourse structure (cf. Grosz and Sidner 1986:188 for a very clear case); and it in turn influences discourse structure. Now, on the one hand, Polanyi and Scha (1984), Hobbs (1985), and Thompson and Mann (1987) have argued that 'genre' or 'rhetorical schemata' can influence the relations used in discourse attachment. On the other hand, Sibun (1992) has recently argued that domain-specific information, as opposed to domain-independent rhetorical information, plays the central role. Both ideas are intriguing, but so far only the latter has been specified in sufficient detail to assess how it works in general, and neither has been applied to time switching or amelioration in particular.

We limit our discussion to temporal aspects of discourse interpretation; our strategy here is to explore two possible *contextual constraints*; these state how the discourse context filters the set of discourse relations and temporal relations which may be used to attach the current clause to the representation of the text so far. We then frame contextual constraints in a logical theory of text interpretation, where their effects and interactions can be precisely calculated. We therefore first introduce a domain-specific contextual constraint, following Sibun, and then place it in a formal theory of discourse attachment called DICE, developed in Lascarides and Asher (1991a). We then show how the proposed *domain*-constraint is insufficient, and demonstrate how it can be augmented by adding a rhetorical, or *presentational* constraint to the theory.

Constraints from the Domain Context

In the field of NL generation, Sibun (1992) has recently argued that coherent text must have a structure closely related to the domain structure of its subject matter; naturally, her remarks are also relevant to NL interpretation. She pursues a view that task structure, or more generally, domain structure, is sufficient to account for many discourse phenomena (but cf. Grosz and Sidner 1986:182). She examines in detail the generation of paragraph-length texts describing the layout of a house. Houses have structure, following from a basic relation of spatial proximity, and there are also hierarchical levels to the structure (rooms can be listed without describing what's in them, or the objects within each room can be detailed).

Either way, one constraint on text structure is defined in terms of the description's *trajectory*: the spatial *direction* the description moved in the domain, to get from the objects already described to the current one. The constraint is: don't change trajectory. Sibun argues that in the temporal domain, the basic relation is temporal proximity. But Lascarides and Oberlander (1992a) urge that the temporal coherence of text is characterised in terms of, among other things, the stronger basic relation of *causal* proximity. So in the latter domain, Sibun's domain constraint precludes textual descriptions which proceed from a cause to an effect to a further cause of that effect, or from effect to cause to effect.

This *Maintain Causal Trajectory* (MCT) constraint has two important attributes: first, it is domain-specific; secondly, it introduces into discourse interpretation an element of top-down processing. To investigate these properties, and see how far they go towards explaining discourse time switch, and discourse amelioration, we now incorporate MCT into DICE's formal model of discourse structure, where its interaction with other causal information and strategies for interpretation can be precisely calculated.

Discourse Interpretation and Commonsense Entailment

DICE (Discourse and Commonsense Entailment) starts with traditional discourse representation structures (cf. Kamp 1981), but goes on to assume with Grosz and Sidner (1986) that candidate discourses possess hierarchical structure, with units linked by discourse relations modelled after those proposed by Hobbs (1979, 1985) (cf. also Thompson and Mann 1987, Scha and Polanyi 1988).¹ Lascarides and Asher (1991a) use *Narration*, *Explanation*, *Background*, *Result* and *Elaboration*. These are the discourse relations central to temporal import and they are the only ones we consider here. Full coverage of text would require a larger set of relations, akin to that in Thompson and Mann (1987).

DICE is a dynamic, logical theory for determining the discourse relations between sentences in a text, and the temporal relations between the eventualities they describe. The logic used is the nonmonotonic logic Commonsense Entailment (CE) proposed by Asher and Morreau (1991). Implicatures are calculated via default rules. The rules introduced below are shown in Lascarides and Asher (1991a) to be manifestations of Gricean-style pragmatic maxims and world knowledge.

Discourse Structure and Implicature

A formal notation makes clear both the logical structure of these rules, and the problems involved in calculating implicature. Let $\langle \tau, \alpha, \beta \rangle$ be the update function, which means "the representa-

¹Lascarides and Asher (1991a) introduces the general framework and applies it to interpretation; Oberlander and Lascarides (1992) and Lascarides and Oberlander (1992b) use the framework for generation.

tion τ of the text so far (of which α is already a part) is to be updated with the representation β of the current clause via a discourse relation with α . Let $\alpha \Downarrow \beta$ mean that α is a topic for β ; let e_α be a term referring to the main eventuality described by the clause α ; and let $fall(m, e_\alpha)$ mean that this event is a Max falling. Let $e_1 \prec e_2$ mean the eventuality e_1 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 ’) and \rightarrow is the material conditional. The maxims for modelling implicature are then represented as schemas:²

- **Narration:** $\langle \tau, \alpha, \beta \rangle > Narration(\alpha, \beta)$
- **Axiom on Narration:**
 $Narration(\alpha, \beta) \rightarrow e_\alpha \prec e_\beta$
- **Explanation:** $\langle \tau, \alpha, \beta \rangle \wedge cause(e_\beta, e_\alpha) > Explanation(\alpha, \beta)$
- **Axiom on Explanation:**
 $Explanation(\alpha, \beta) \rightarrow \neg e_\alpha \prec e_\beta$
- **Push Causal Law:**
 $\langle \tau, \alpha, \beta \rangle \wedge fall(m, e_\alpha) \wedge push(j, m, e_\beta) > cause(e_\beta, e_\alpha)$
- **Causes Precede Effects:**
 $cause(e_2, e_1) \rightarrow \neg e_1 \prec e_2$
- **States Overlap:**
 $\langle \tau, \alpha, \beta \rangle \wedge state(e_\beta) > overlap(e_\alpha, e_\beta)$
- **Background:** $\langle \tau, \alpha, \beta \rangle \wedge overlap(e_\alpha, e_\beta) > Background(\alpha, \beta)$
- **Axiom on Background:**
 $Background(\alpha, \beta) \rightarrow overlap(e_\alpha, e_\beta)$

The rules for Narration, Explanation and Background 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

²Discourse structure and $\alpha \Downarrow \beta$ are given model theoretical interpretations in Asher (in press); 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 .

temporal order in interpretation. The Push Causal Law is a mixture of linguistic knowledge and world knowledge; given that the clauses are discourse-related somehow, the events they describe must normally be connected in a causal, part/whole or overlap relation; here, given the events in question, they must normally stand in a causal relation. That Causes Precede their Effects is indefeasible world knowledge.

We also have laws relating the discourse structure to the topic structure (Asher, in press): for example, A Common Topic for Narrative states that any clauses related by *Narration* must have a distinct, common (and perhaps implicit) topic:

- **A Common Topic for Narrative**
 $Narration(\alpha, \beta) \rightarrow (\exists \gamma)(\gamma \Downarrow \alpha \wedge \gamma \Downarrow \beta) \wedge \neg(\alpha \Downarrow \beta) \wedge \neg(\beta \Downarrow \alpha)$

The hierarchical discourse structure is similar to that in Scha and Polanyi (1988): *Elaboration* and *Explanation* are subordinating relations and the others are coordinating ones. Equally, this structure defines similar constraints on attachment: the current clause must attach to the previous clause or else to the clauses it elaborates or explains. In other words, the open clauses are those on the right frontier. We do not directly encode the nucleus/satellite distinction used in RST (Thompson and Mann, 1987).

Interpretation by Deduction

CE and the defeasible rules are used to infer the discourse—and temporal—structures of candidate texts. CE represents nonmonotonic validity as \approx . Three 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:**
 $\phi \rightarrow \psi, \psi > \chi, \phi > \neg\chi, \phi \approx \neg\chi$
e.g. Penguins are birds, birds normally fly, penguins normally don’t fly, Tweety is a penguin; so Tweety doesn’t fly.
- **Nixon Diamond:**
Not: $\phi > \chi, \psi > \neg\chi, \phi, \psi \approx \chi$ (or $\neg\chi$)
e.g. Not: Quakers are pacifists, Republicans are not, Nixon is both a quaker and republican
 \approx Nixon is a pacifist/Nixon is a non-pacifist.

Iconic and Non-iconic text: In interpreting text (5) we attempt to attach the second clause to the first (so $\langle \alpha, \alpha, \beta \rangle$ holds, where α and β are respectively the logical forms of the first and second clauses).

- (5) Max stood up. John greeted him.
 (1) Max fell. John pushed him.

In the absence of further information, the only rule whose antecedent is satisfied is Narration. So we infer via Defeasible Modus Ponens that the *Narration* relation holds between its clauses. This then yields, assuming logical omniscience, an iconic interpretation; the standing up precedes the greeting. In contrast, text (1) 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 Narration's. Hence (1) is interpreted as: 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.³ Notice that deductions about event structure and discourse structure are interleaved.

Incoherence and popping: Consider the incoherent text (3).

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

The Win Law captures the intuition that if Max wins the race and he is at home, then these events normally don't temporally overlap—regardless of whether they're connected or not.

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

The appropriate knowledge base in the analysis of (3) satisfies States Overlap, the Win Law and Narration. The first two of these conflict, but their antecedents aren't logically related. They

³The formal details of how the logic CE models these interpretations are given in Lascarides and Asher (1991b). Although the double application of the Penguin Principle, as in (1), is not valid in general, they show that for the particular case considered here, CE validates it.

therefore form a pattern out of which a Nixon Diamond crystallises: no temporal or discourse relation can be inferred. We stipulate that it is incoherent to assume that $\langle \tau, \alpha, \beta \rangle$ if one can't infer which discourse relation holds between α and β . So the assumption that the clauses are connected must be dropped, and hence no representation of (3) is constructed.

DICE exploits this account of incoherence in its approach to *discourse popping*. When a Nixon Diamond occurs in attempting to attach the current clause to the previous one, they don't form a coherent text segment. So the current clause must attach to one of the other open clauses, resulting in discourse popping (Lascarides and Asher, 1991a).

Trajectory in DICE

It should be clear DICE's devices, while formal, are also quite powerful. However, the maxims introduced so far cannot actually explain either discourse time switching (cf. (1) *vs* (2)) or amelioration (cf. (3) *vs* (4)). Incorporating some form of contextual constraint may be one way to deal with such cases. Because DICE makes essential use of nonmonotonic inference, adding contextual constraints will alter the inferences without requiring modification of the existing knowledge representation. We now investigate the consequences of adding MCT.

Maintain Causal Trajectory

Suppose $R(\alpha, \beta)$ holds for some discourse relation R ; then α appears in the text before β , and we use this fact to define MCT. The default law below states that if the existing discourse context is one where a cause/effect relation was described in that order, then the current clause should not describe a further cause of the effect:

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

In using this rule, an interpreter brings to bear 'top-down' information, in the following sense. Up to now, discourse and temporal relations have been determined by using the input discourse as data, and predicting the relations using general linguistic and world knowledge. Now, the interpreter is permitted to 'remember' which predic-

tion they made last time, and use this to constrain the kind of relation that can be inferred for attaching the current clause; this new prediction needs no data to drive it. Of course, incoming data can prevent the prediction from being made; MCT is just a default, and (6) is an exception.


- (6) Max switched off the light. The room went pitch dark, since he had drawn the blinds too.

Time Switching

MCT says how the event structures predicted for preceding context can affect the temporal relations predicted for the current clause. But how does it interact with other causal knowledge in DICE? Does it account for time switching? Since MCT is a contextual constraint, it will only interact with causal knowledge in a discourse context. So consider how it affects the attachment of (2c) and (2d).

- (2) a. John and Max came to the cliff's edge. α
 b. John applied a sharp blow to the back of Max's neck. β
 c. Max fell. γ
 d. John pushed him. δ
 e. Max rolled over the edge of the cliff. ϵ

Suppose that the logical forms of the clauses (2a–e) are respectively α to ϵ , and suppose that the discourse structure up to and including γ has been constructed in agreement with intuitions:



Furthermore, assume, in line with intuitions, that the interpreter has inferred that e_β caused e_γ . Consider how δ is to be attached to the above discourse structure. γ is the only open clause; so $\langle \tau, \gamma, \delta \rangle$ must hold. The antecedents to three defeasible laws are verified: the Push Causal Law and Narration just as before, and also MCT. The consequents of the Push Causal Law and MCT conflict; moreover, their antecedents aren't logically related. So by the Nixon Diamond, we can't infer which event—or discourse—relation holds. Accordingly, the discourse is actually incoherent. Yet intuitively, a relation can be inferred: the push happened after the fall, and the clauses γ and δ must be related by *Narration*.

On its own, MCT cannot account for time switching (or, indeed, amelioration). In one sense this isn't surprising. Causal knowledge and MCT were in conflict in (2), and since both laws relate to the domain, but in incommensurable ways, neither logic nor intuition can say which default is preferred. This suggests that using domain structure alone to constrain interpretation will be insufficient. It seems likely that presentational issues will be significant in cases such as these; where domain-specific knowledge sources are in irresolvable conflict, aspects of the existing discourse structure may help determine current clause attachment. Since MCT has some motivation, it would be preferable to let presentational information interact with it, rather than replace it.

Constraints from the Presentational Context

To what degree does existing rhetorical structure determine clause attachment? It's plausible to suggest that a speaker-writer should not switch genre without syntactically marking the switch. Thus, if the preceding context is narrative, then a hearer-reader will continue to interpret the discourse as narrative unless linguistic markers indicate otherwise; similarly for non-narrative contexts (cf. Caenepeel 1991, Polanyi and Scha 1984).

This constraint relies on the continuation of a characteristic pattern of *discourse* relations, rather than on maintaining trajectory on some *domain* relation. Let's call this a presentational constraint; it may be able to get the right analyses of (2) and (4). In (2), for example, the context to which *John pushed him* is attached is narrative, so according to the constraint this clause would be attached with *Narration* in agreement with intuitions. But clearly, this constraint must be a soft one, since discourse pops can occur without syntactic markers, as can interruptions (Polanyi 1985:306). Both of these cause a change in the discourse 'pattern' established in the preceding context.

Patterns in DICE

Can we use presentational constraints without accidentally blocking discourse popping and interruptions? The problem is to represent in formal

terms exactly *when* an interpreter should try to preserve the pattern of rhetorical structure established in the context. Because DICE provides a formal account of how discourse popping occurs—the Nixon Diamond is the key—we are in a good position to attempt this.

Discourse Pattern and Inertia

First, we define the discourse pattern established by the context in terms of a function DP . This takes as input the discourse structure for the preceding context, filters out those discourse relations which would break the pattern, and outputs the remaining set of relations. This is similar to Hobbs' (1985:25–26) notion of genre, where, for example (in his terms) a story genre requires that the type of *occasion* relation can be only *problem-solution* or *event-outcome*.

How much of the preceding discourse context does DP take as input? At one extreme, it could be just the discourse relations used to attach the previous clause; the output would be those same discourse relations. At the other extreme, the *whole* discourse structure may be input; DP would have to establish the regularity in the configuration of discourse relations, and evaluate which discourse relation would preserve it when the new clause is added. We leave this question open; for the examples of time switching and amelioration we consider here, DP would produce the same result whatever it takes as input—*Narration*.

Using DP , we can represent the discourse pattern constraint. The intuition it captures is the following. If the sentence currently being processed can't attach to any of the open nodes because there's a Nixon Diamond of irresolvable conflict, then assume that the discourse relation to be used is defined by DP . In other words, discourse pattern preservation applies only when all other information prevents attachment at all available open nodes. To express this formally, we need a representation of a state in which a Nixon Diamond has formed. In CE, we use the formula \perp (meaning contradiction) and the connective $\&$, whose semantics is defined only in the context of default laws (cf. Asher and Morreau 1991b). Intuitively, $(A\&B) > \perp$ means 'A and B are antecedents of default rules that lead to a conflict that can't be resolved'.

We use this to represent cases where the information provided by the clauses α and β (which

are candidates for attachment) form a Nixon Diamond. Let $Info(\alpha)$ be glossed 'the information $Info$ is true of the clause α '. It is an abbreviation for statements such as $fall(max, e_\alpha)$, $cause(e_\alpha, e_\beta)$, and so on. If a Nixon Diamond occurs when attempting to attach α to β on the basis of information other than DP , the following holds:

- $Info(\alpha) \wedge Info(\beta) \wedge \langle \tau, \alpha, \beta \rangle \wedge$
 $((Info(\alpha) \wedge Info(\beta)) \& \langle \tau, \alpha, \beta \rangle) > \perp$

We will use $ND(\alpha, \beta)$ as a gloss for the above schema, and $open(\tau, \alpha)$ means α is an open clause in the discourse structure τ ; assume that $DP(\tau)$ returns some discourse relation R . So the presentational constraint for preserving discourse pattern is defined as follows:⁴

- **Inertia:** $(\forall \alpha)(open(\tau, \alpha) \wedge ND(\alpha, \beta)) >$
 $(\exists \alpha')(open(\tau, \alpha') \wedge DP(\tau)(\alpha', \beta))$

The antecedent to Inertia is verified only when all the information available—except for the preceding discourse pattern—yields a Nixon Diamond in attempting the attachment of β at *all* open nodes. Inertia thus won't prevent discourse popping, because there a Nixon Diamond is averted at a higher-level open node. The model of text processing proposed here restricts the kind of information that's relevant during text processing: the discourse pattern is relevant only when all other information is insufficient. Like MCT, Inertia is top-down, in the sense that it relies on earlier predictions about other discourse relations, rather than on incoming data; but unlike MCT, the 'theory-laden' predictions are only resorted to if the data seems recalcitrant.

Time Switching

We now look at text (2) in detail. Suppose as before that the discourse structure τ for the first three clauses in (2) is (2'), and the task now is to attach δ (i.e. *John pushed him*). The only open clause is γ , because the previous discourse relations are all *Narration*. Moreover, $DP(\tau)$ is *Narration*. As before, a Nixon Diamond forms between MCT and the Push Causal Law in attempting to attach δ to γ . Where A_{mct} is the antecedent to MCT, and A_{pcl} the antecedent to the Push Causal Law substituted with γ and δ :

⁴Inertia features an embedded default connective. Only two nonmonotonic logics can express this: Circumscription and CE.

- $A_{mct} \wedge A_{pcl} \wedge ((A_{pcl} \& A_{mct}) > \perp)$

So $ND(\gamma, \delta)$ is verified, and with it, the antecedent to Inertia; substituting in the Inertia schema the value of $DP(\tau)$, the Nixon Diamond, and the open clauses yields the following:

- **Inertia for (2):**
 $(A_{mct} \wedge A_{pcl} \wedge ((A_{pcl} \& A_{mct}) > \perp)) >$
 $Narration(\gamma, \delta)$

The antecedent to Inertia entails that of Maintain Trajectory (A_{mct}) and that of Push Causal Law (A_{pcl}). In CE the most specific law wins. So the discourse context in this case determines the relation between the fall and the push: it is *Narration*. Hence even though WK yields a causal preference for the pushing causing the falling, given the discourse context in which the pushing and falling are described in (2), *Narration* is inferred after all, and so the falling precedes the push. In this way, we can represent the presentational, and domain-specific, information that must be brought to bear to create a time switch.⁵

Amelioration

Now consider texts (3) and (4). A Nixon Diamond formed between *Narration*, *States Overlap* and the *Win Law* in the analysis of (3) above, leading to incoherence. Now consider attaching the same clauses (4e) and (4f).

- (4) a. Max got up early yesterday.
 b. He had a little bite to eat.
 c. He had a light workout.
 d. He started the tournament in good form.
 e. He won the race in record time.
 f. He was home with the cup.
 g. He celebrated until late into the evening.

Given the discourse (4a–e), (4e) is the only open clause to which (4f) can attach. Moreover, as in (3), attempting to attach (4f) to (4e) results in a Nixon Diamond. So the antecedent to Inertia is verified. DP delivers *Narration*, since the discourse context is narrative. So (4e–f) is interpreted as a narrative. Compare this with (3),

⁵If a speaker-writer wanted to avoid this contextual inference pattern, and sustain the non-iconic reading, then they could switch to the pluperfect, for example.

where no discourse relation was inferred, leading to incoherence.

Inertia enables discourse context to establish coherence between sentence pairs that, in isolation, are incoherent. It would be worrying if Inertia were so powerful that it could ameliorate *any* text. But incoherence is still possible: consider replacing (4f) with (4f’):

- f. ?Mary’s hair was black.

If world knowledge is coded as intuitions would suggest, then no common topic can be constructed for (4e) and (4f’); and this is necessary if they are to be attached with *Narration* or *Background*—the only discourse relations available given the defensible laws that are verified. Moreover, Inertia won’t improve the coherence in this case because it predicts *Narration*, which because of Common Topic for *Narration* cannot be used to attach (4f’) to (4e). So the text is incoherent.

Hobbs et al (1990) also explore the effects of linguistic and causal knowledge on interpretation, using abduction rather than deduction. Now, Konolige (1991) has shown that abduction and nonmonotonic deduction are closely related; but since Hobbs et al don’t attempt to treat time-switching and amelioration, direct comparison here is difficult. However, the following points are relevant. First, weighted abduction, as a system of inference, isn’t embeddable in CE, and vice versa. Secondly, the weights which guide abduction are assigned to predicates in a context-free fashion. Hobbs et al observe that this may make the effects of context hard to handle, since ‘the abduction scheme attempts to make global judgements on the basis of strictly local information’ [p48].

Conclusion

We examined instances of two types of contextual constraint on current clause attachment. These were *Maintain Causal Trajectory*, a domain constraint; and *Inertia*, a presentational constraint. We argued that domain constraints seemed insufficient, but that presentational constraints could constructively interact with them. This interaction then explains the two discourse interpretation phenomena we started out with. Context can switch round the order of events; and it can ameliorate an otherwise incoherent interpretation.

Both of the constraints allow predictions about

new discourse relations to be driven from previous predictions. But MCT simply adds its prediction to the data-driven set from which the logic chooses, whereas discourse pattern and Inertia are only relevant to interpretation when the logic can otherwise find no discourse relation.

This formalisation has also raised a number of questions for future investigation. For example, the discourse pattern (or Hobbsian ‘genre’) function is important; but how much of the preceding discourse structure should the *DP* function take as input? How do we establish—and improve—the linguistic coverage? What is the relation between communicative intentions and contextual constraints? How do we actually implement contextual constraints in a working system?

The idea of contextual constraints is a familiar and comfortable one. In this respect, we have merely provided one way of formally pinning it down. Naturally, this requires a background logical theory of discourse structure, and we have used DICE, which has its own particular set of discourse relations and implicature patterns. However, the process of logically specifying the constraints has two important and general benefits, independent of the particular formalisation we have offered. First, it demands precision and uniformity in the statement both of the new constraints, and of the other knowledge sources used in interpretation. Secondly, it permits a program-independent assessment of the consequences of the general idea of contextual constraints.

References

- Asher, Nicholas [in press] Reference to Abstract Objects in English: A Philosophical Semantics for Natural Language Metaphysics. Dordrecht: Kluwer Academic Publishers.
- Asher, Nicholas and Morreau, Michael [1991] Common Sense Entailment: A Modal Theory of Non-monotonic Reasoning. In *Proceedings of the 12th International Joint Conference on Artificial Intelligence*.
- Caenepeel, Mimo [1991] Event Structure versus Discourse Coherence. In *Proceedings of the Workshop on Discourse Coherence*, Edinburgh, 4–6 April, 1991.
- Grosz, Barbara and Sidner, Candy [1986] Attention, Intentions, and the Structure of Discourse. *Computational Linguistics*, 12, 175–204.
- Hobbs, Jerry [1979] Coherence and Coreference. *Cognitive Science*, 3, 67–90.
- Hobbs, Jerry [1985] On the Coherence and Structure of Discourse. Report No. CSLI-85-37, Center for the Study of Language and Information.
- Hobbs, Jerry, Stickel, Martin, Appelt, Doug and Martin, Paul [1990] Interpretation as Abduction. Technical Note No. 499, Artificial Intelligence Center, SRI International, Menlo Park.
- Kamp, Hans [1981] 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*, Volume 136, pp277–322. Amsterdam: Mathematical Centre Tracts.
- Konolige, Kurt [1991] Abduction vs. Closure in Causal Theories. Technical Note No. 505, Artificial Intelligence Center, SRI International, Menlo Park.
- Lascarides, Alex and Asher, Nicholas [1991a] Discourse Relations and Defeasible Knowledge. In *Proceedings of the 29th Annual Meeting of Association for Computational Linguistics*, pp55–63.
- Lascarides, Alex and Asher, Nicholas [1991b] Discourse Relations and Common Sense Entailment. DYANA deliverable 2.5B, available from Centre for Cognitive Science, University of Edinburgh.
- Lascarides, Alex and Oberlander, Jon [1992a] Temporal Coherence and Defeasible Knowledge. *Theoretical Linguistics*, 18.
- Lascarides, Alex and Oberlander, Jon [1992b] Abducting Temporal Discourse. In Dale, R. et al (eds.) *Aspects of Automated Natural Language Generation*, pp167–182. Berlin: Springer-Verlag.
- Polanyi, Livia and Scha, Remko [1984] A Syntactic Approach to Discourse Semantics. In *Proceedings of the 22nd Annual Meeting of the Association for Computational Linguistics*, pp413–419.
- Polanyi, Livia [1985] A Theory of Discourse Structure and Discourse Coherence. In *Papers from the General Session at the Twenty-First Regional Meeting of the Chicago Linguistics Society*, pp 25–27.
- Oberlander, Jon and Lascarides, Alex [1992] Preventing False Temporal Implicatures: Interactive Defaults for Text Generation. In *Proceedings of COLING92*.
- Scha, Remko and Polanyi, Livia [1988] An augmented context free grammar. In *Proceedings of the 24th Annual Meeting of the Association for Computational Linguistics*, pp573–577.
- Sibun, Penelope [1992] Generating Text without Trees. To appear in *Computational Intelligence: Special Issue on Natural Language Generation*, 8.
- Thompson, Sandra and Mann, William [1987] Rhetorical Structure Theory: A Framework for the Analysis of Texts. In *IPRA Papers in Pragmatics*, 1, pp79–105.