THE UNIVERSITY *of* EDINBURGH

Edinburgh Research Explorer

Intentions and information in discourse

INTENTIONS AND INFORMATION IN DISCOURSE

Nicholas Asher
IRIT, Université Paul Sabatier,
118 Route de Narbonne,
31062 Toulouse, CEDEX,
France
asher@irit.fr

Alex Lascarides
Department of Linguistics,
Stanford University,
Stanford,
Ca 94305-2150,
USA,
alex@csli.stanford.edu

Abstract

This paper is about the flow of inference between communicative intentions, discourse structure and the domain during discourse processing. We augment a theory of discourse interpretation with a theory of distinct mental attitudes and reasoning about them, in order to provide an account of how the attitudes interact with reasoning about discourse structure.

INTRODUCTION

The flow of inference between communicative intentions and domain information is often essential to discourse processing. It is well reflected in this discourse from Moore and Pollack (1992):

(1)a. George Bush supports big business.

b. He's sure to veto House Bill 1711.

There are at least three different interpretations. Consider Context 1: in this context the interpreter I believes that the author A wants to convince him that (1b) is true. For example, the context is one in which I has already uttered *Bush won't veto any more bills*. I reasons that A's linguistic behavior was intentional, and therefore that A believes that by saying (1a) he will convince I that Bush will veto the bill. Even if I believed nothing about the bill, he now infers it's bad for big business. So we have witnessed an inference from premises that involve the desires and beliefs of A (Moore and Pollack's "intentional structure"), as well as his linguistic behavior, to a conclusion about domain information (Moore and Pollack's "informational structure").

Now consider Context 2: in this context I knows that A wants to convince him of (1a). As in Context 1, I may infer that the bill is bad for big business. But now, (1b) is used to support (1a).

Finally, consider Context 3: in this context I knows that House Bill 1711 is bad for big business, but doesn't know A's communicative desires prior to witnessing his linguistic behaviour. From his beliefs about the domain, he infers that supporting big business would cause Bush to veto this bill. So, A must have uttered (1a) to support (1b). Hence I realises that A wanted

him to believe (1b). So in contrast to Contexts 1 and 2, we have a flow of inference from informational structure to intentional structure.

This story makes two main points. First, we agree with Moore and Pollack that we must represent both the intentional import and the informational import of a discourse. As they show, this is a problem for current formulations of Rhetorical Structure Theory (RST) (Thompson and Mann, 1987). Second, we go further than Moore and Pollack, and argue that reasoning about beliefs and desires exploits different rules and axioms from those used to infer rhetorical relations. Thus, we should represent intentional structure and discourse structure separately. But we postulate rhetorical relations that express the discourse function of the constituents in the communicative plan of the author, and we permit interaction between reasoning about rhetorical relations and reasoning about beliefs and desires.

This paper provides the first steps towards a formal analysis of the interaction between intentional structure and informational structure. Our framework for discourse structure analysis is SDRT (Asher 1993). The basic representational structures of that theory may be used to characterise cognitive states. We will extend the logical engine used to infer rhetorical relations—DICE (Lascarides and Asher 1991, 1993a, 1993b, Lascarides and Oberlander 1993)—to model inferences about intentional structure and its interaction with informational structure.

BUSH'S REQUIREMENTS

We must represent both the intentional import and the informational import of a discourse simultaneously. So we need a theory of discourse structure where discourse relations central to intentional import and to informational import can hold simultaneously between the same constituents. A logical framework in which all those plausible relations between constituents that are consistent with each other are inferred, such as a nonmonotonic logic like that in DICE (Lascarides and Asher, 1993a), would achieve this. So conceivably, a similar nonmonotonic logic for RST might solve the problem of keeping track of the intentional and informational

structure simultaneously.

But this would work only if the various discourse relations about intentions and information could simultaneously hold in a *consistent* knowledge base (KB). Moore and Pollack (1992) show via discourse (2) that the current commitment to the nucleus-satellite distinction in RST precludes this.

(2)a. Let's go home by 5.

b. Then we can get to the hardware store
before it closes.

c. That way we can finish the bookshelves tonight.

From an intentional perspective, (2b) is a satellite to (2a) via *Motivation*. From an informational perspective, (2a) is a satellite to (2b) via *Condition*. These two structures are incompatible. So augmenting RST with a nonmonotonic logic for inferring rhetorical relations would not yield a representation of (2) on multiple levels in which both intentional and informational relations are represented. In SDRT, on the other hand, not all discourse relations induce subordination, and so there is more scope for different discourse relations holding simultaneously in a consistent KB.

Grosz and Sidner's (1986) model of discourse interpretation is one where the same discourse elements are related simultaneously on the informational and intentional levels. But using their framework to model (1) is not straightforward. As Grosz and Sidner (1990) point out: "any model (or theory) of the communication situation must distinguish among beliefs and intentions of different agents," but theirs does not. They represent intentional structure as a stack of propositions, and different attitudes aren't distinguished. The informal analysis of (1) above demands such distinctions, however. For example, analysing (1) under Context 3 requires a representation of the following statement: since A has provided a reason why (1b) is true, he must *want* I to *believe* that (1b) is true. It's unclear how Grosz and Sidner would represent this. SDRT (Asher, 1993) is in a good position to be integrated with a theory of cognitive states, because it uses the same basic structures (discourse representation structures or DRSs) that have been used in Discourse Representation Theory (DRT) to represent different attitudes like beliefs and desires (Kamp 1981, Asher 1986, 1987, Kamp 1991, Asher and Singh, 1993).

A BRIEF INTRODUCTION TO SDRT AND DICE

In SDRT (Asher, 1993), an NL text is represented by a segmented DRS (SDRS), which is a pair of sets containing: the DRSs or SDRSs representing respectively sentences and text segments, and discourse relations between them. Discourse relations, modelled after those proposed by Hobbs (1985), Polanyi (1985) and Thompson and Mann (1987), link together the constituents of an SDRS. We will mention three: *Narration, Result* and *Evidence*.

SDRSs have a hierarchical configuration, and SDRT predicts points of attachment in a discourse structure for new information. Using DICE we infer from the reader's knowledge resources *which* discourse relation should be used to do attachment.

Lascarides and Asher (1991) introduce default rules representing the role of Gricean pragmatic maxims and domain knowledge in calculating the value of the update function $\langle \tau, \alpha, \beta \rangle$, which means "the representation β of the current sentence is to be attached to α with a discourse relation, where α is an open node in the representation τ of the text so far". Defaults are represented by a conditional—$\phi > \psi$ means 'if ϕ, then normally ψ. For example, Narration says that by default *Narration* relates elements in a text.

• **Narration**: $\langle \tau, \alpha, \beta \rangle > Narration(\alpha, \beta)$

Associated axioms show how *Narration* affects the temporal order of the events described: Narration and the corresponding temporal axioms on *Narration* predict that normally the textual order of events matches their temporal order.

The logic on which DICE rests is Asher and Morreau's (1991) Commonsense Entailment (CE). Two patterns of nonmonotonic inference are particularly relevant here. The first is Defeasible Modus Ponens: if one default rule has its antecedent verified, then the consequent is nonmonotonically inferred. The second is the Penguin Principle: if there are conflicting default rules that apply, and their antecedents are in logical entailment relations, then the consequent of the rule with the most specific antecedent is inferred. Lascarides and Asher (1991) use DICE to yield the discourse structures and temporal structures for simple discourses. But the theory has so far ignored how A's intentional structure—or more accurately, I's model of A's intentional structure—influences I's inferences about the domain and the discourse structure.

ADDING INTENTIONS

To discuss intentional structure, we develop a language which can express beliefs, intentions and desires. Following Bratman (forthcoming) and Asher and Singh (1993), we think of the objects of attitudes either as plans or as propositions. For example, the colloquial intention to do something—like wash the dishes—will be expressed as an intention toward a plan, whereas the intention that Sue be happy is an intention toward a proposition. Plans will just consist of sequences of basic actions $a_1; a_2; \ldots; a_n$. Two operators—$\mathcal{R}$ for *about to do* or *doing*, and $\mathcal{D}$ for *having done*—will convert actions into propositions. The attitudes we assume in our model are *believes* ($\mathcal{B}_A \phi$ means 'A believes ϕ'), *wants* ($\mathcal{W}_A \phi$ means 'A wants ϕ'), and *intends* ($\mathcal{I}_A \phi$ means 'A intends ϕ'). All of this takes place in a modal, dynamic logic, where the propositional attitudes are supplied with a modal semantics. To this we add the modal conditional operator $>$, upon which the logic of DICE is

based.

Let's take a closer look at (1) in Context 1. Let the logical forms of the sentences (1a) and (1b) be respectively α and β. In Context 1, I believes that A wants to convince him of β and thinks that he doesn't believe β already. Following the DRT analysis of attitudes, we assume I's cognitive state has embedded in it a model of A's cognitive state, which in turn has a representation of I's cognitive state. So $\mathcal{W}_A \mathcal{B}_I \beta$ and $\mathcal{B}_A \neg \mathcal{B}_I \beta$ hold in I's KB. Furthermore, $\langle \tau, \alpha, \beta \rangle \wedge Info(\alpha, \beta)$ holds in I's KB, where $Info(\alpha, \beta)$ is a gloss for the semantic content of α and β that I knows about.[1] I must now reason about what A intended by his particular discourse action. I is thus presented with a classical reasoning problem about attitudes: how to derive what a person believes, from a knowledge of what he wants and an observation of his behaviour. The classic means of constructing such a derivation uses the practical syllogism, a form of reasoning about action familiar since Aristotle. It expresses the following maxim: Act so as to realize your goals *ceteris paribus*.

The practical syllogism is a rule of defeasible reasoning, expressible in CE by means of the nonmonotonic consequence relation $\approx$. The consequence relation $\phi \approx \psi$ can be stated directly in the object language of CE by a formula which we abbreviate as $\beth(\phi, \psi)$ (Asher 1993). We use $\beth(\phi, \psi)$ to state the practical syllogism. First, we define the notion that the KB and ϕ, but not the KB alone, nonmonotonically yield ψ:

- **Definition:**
  $\beth_{kb}(\phi, \psi) \leftrightarrow \beth(KB \wedge \phi, \psi) \wedge \neg \beth(KB, \psi)$

The Practical Syllogism says that if (a) A wants ϕ but believes it's not true, and (b) he knows that if ψ were added to his KB it would by default make ϕ true eventually, then by default A intends ψ.

- **The Practical Syllogism:**
  - (a) $(\mathcal{W}_A(\phi) \wedge \mathcal{B}_A(\neg\phi) \wedge$
  - (b) $\mathcal{B}_A(\beth_{kb}(\psi, eventually(\phi)))) >$
  - (c) $\mathcal{I}_A(\psi)$

The Practical Syllogism enables I to reason about A's cognitive state. In Context 1, when substituting in the Practical Syllogism $\mathcal{B}_I \beta$ for ϕ, and $\langle \tau, \alpha, \beta \rangle \wedge Info(\alpha, \beta)$ for ψ, we find that clause (a) of the antecedent to the Practical Syllogism is verified. The conclusion (c) is also verified, because I assumes that A's discourse act was intentional. This assumption could be expressed explicitly as a $>$-rule, but we will not do so here.

Now, abduction (i.e., explanatory reasoning) as well as nonmonotonic deduction is permitted on the Practical Syllogism. So from knowing (a) and (c), I can conclude the premise (b). We can state in CE an 'abductive' rule based on the Practical Syllogism:

- **The Abductive Practical Syllogism 1 (APS1)**
  $(\mathcal{W}_A(\phi) \wedge \mathcal{B}_A(\neg\phi) \wedge \mathcal{I}_A(\psi)) >$
  $\mathcal{B}_A(\beth_{kb}(\psi, eventually(\phi)))$

---

[1] This doesn't necessarily include that House Bill 1711 is bad for big business.

APS1 allows us to conclude (b) when (a) and (c) of the Practical Syllogism hold. So, the intended action ψ must be one that A believes will eventually make ϕ true.

When we make the same substitutions for ϕ and ψ in APS1 as before, I will infer the conclusion of APS1 via Defeasible Modus Ponens: $\mathcal{B}_A(\beth_{kb}(\langle \tau, \alpha, \beta \rangle \wedge Info(\alpha, \beta), eventually(\mathcal{B}_I \beta)))$. That is, I infers that A believes that, by uttering what he did, I will come to believe β.

In general, there may be a variety of alternatives that we could use to substitute for ϕ and ψ in APS1, in a given situation. For usually, there are choices on what can be abduced. The problem of choice is one that Hobbs *et al.* (1990) address by a complex weighting mechanism. We could adopt this approach here.

The Practical Syllogism and APS1 differ in two important ways from the DICE axioms concerning discourse relations. First, APS1 is motivated by an *abductive* line of reasoning on a pattern of defeasible reasoning involving cognitive states. The DICE axioms are not. Secondly, both the Practical Syllogism and APS1 don't include the discourse update function $\langle \tau, \alpha, \beta \rangle$ together with some information about the semantic content of α and β in the antecedent, while this is a standard feature of the DICE axioms for inferring discourse structure.

These two differences distinguish reasoning about intentional structures and discourse structures. But discourse structure is linked to intentional structure in the following way. The above reasoning with A's cognitive state has led I to conclusions about the *discourse function* of α. Intuitively, α was uttered to support β, or α 'intentionally supports' β. This idea of *intentional support* is defined in DICE as follows:

- **Intends to Support:**
  $Isupport(\alpha, \beta) \leftrightarrow (\mathcal{W}_A(\mathcal{B}_I \beta) \wedge \mathcal{B}_A(\neg\mathcal{B}_I \beta) \wedge$
  $\mathcal{B}_A(\beth_{kbh}(\langle \tau, \alpha, \beta \rangle \wedge Info(\alpha, \beta), eventually(\mathcal{B}_I \beta))))$

In words, α intentionally supports β if and only if A wants I to believe β and doesn't think he does so already, and he also believes that by uttering α and β together, so that I is forced to reason about how they should be attached with a rhetorical relation, I will come to believe β.

$Isupport(\alpha, \beta)$ defines a relationship between α and β at the discourse structural level, in terms of I's and A's cognitive states. With it we infer further information about the particular discourse relation that I should use to attach β to α. $Isupport(\alpha, \beta)$ provides the link between reasoning about cognitive states and reasoning about discourse structure.

Let us now return to the interpretation of (1) under Context 1. I concludes $Isupport(\alpha, \beta)$, because the right hand side of the $\leftrightarrow$-condition in Intends to Support is satisfied. So I passes from a problem of reasoning about A's intentional structure to one of reasoning about discourse structure. Now, I should check to see whether α actually does lead him to believe β. This is a check on the coherence of discourse; in order for an SDRS τ to

be coherent, the discourse relations predicated of the constituents must be satisfiable.[2] Here, this amounts to justifying A's belief that given the discourse context and I's background beliefs of which A is aware, I will arrive at the desired conclusion—that he believes β. So, I must be able to infer a particular discourse relation R between α and β that has what we will call the Belief Property: $(\mathcal{B}_I\alpha \wedge R(\alpha, \beta)) > \mathcal{B}_I\beta$. That is, R must be a relation that would indeed license I's concluding β from α.

We concentrate here for illustrative purposes on two discourse relations with the Belief Property: $Result(\alpha, \beta)$ and $Evidence(\alpha, \beta)$; or in other words, α results in β, or α is evidence for β.

- **Relations with the Belief Property:**
  $(\mathcal{B}_I\alpha \wedge Evidence(\alpha, \beta)) > \mathcal{B}_I\beta$
  $(\mathcal{B}_I\alpha \wedge Result(\alpha, \beta)) > \mathcal{B}_I\beta$

The following axiom of Cooperation captures the above reasoning on I's part: if α $Isupports$ β, then it must be possible to infer from the semantic content, that either $Result(\alpha, \beta)$ or $Evidence(\alpha, \beta)$ hold:

- **Cooperation:**
  $(Isupport(\alpha, \beta) \wedge \langle \tau, \alpha, \beta \rangle) \rightarrow$
  $\quad (\exists_{kb}(\langle \tau, \alpha, \beta \rangle \wedge Info(\alpha, \beta), Result(\alpha, \beta)) \vee$
  $\quad \exists_{kb}(\langle \tau, \alpha, \beta \rangle \wedge Info(\alpha, \beta), Evidence(\alpha, \beta)))$

The intentional structure of A that I has inferred has restricted the candidate set of discourse relations that I can use to attach β to α: he must use $Result$ or $Evidence$, or both. If I can't accommodate A's intentions by doing this, then the discourse will be incoherent. We'll shortly show how Cooperation contributes to the explanation of why (3) is incoherent.

(3)a. George Bush is a weak-willed president.
  b. ?He's sure to veto House Bill 1711.

FROM INTENTIONS TO INFORMATION: CONTEXTS 1 AND 2

The axioms above allow I to use his knowledge of A's cognitive state, and the behaviour of A that he observes, to (a) infer information about A's communicative intentions, and (b) consequently to restrict the set of candidate discourse relations that are permitted between the constituents. According to Cooperation, I must infer that one of the permitted discourse relations does indeed hold. When clue words are lacking, the semantic content of the constituents must be exploited. In certain cases, it's also necessary to infer further information that wasn't explicitly mentioned in the discourse,

in order to sanction the discourse relation. For example, in (1) in Contexts 1 and 2, I infers the bill is bad for big business.

Consider again discourse (1) in Context 1. Intuitively, the reason we can infer $Result(\alpha, \beta)$ in the analysis of (1) is because (i) α entails a generic (Bush vetoes bills that are bad for big business), and (ii) this generic makes β true, as long as we assume that House Bill 1711 is bad for big business.

To define the Result Rule below that captures this reasoning for discourse attachment, we first define this generic-instance relationship: $instance(\phi, \psi)$ holds just in case ϕ is $(\forall x)(A(x) > B(x))$ and ψ is $A[x/d] \wedge B[x/d]$. For example, $bird(tweety) \wedge fly(tweety)$ (Tweety is a bird and Tweety flies) is an instance of $\forall x(bird(x) > fly(x))$ (Birds fly).

The Result Rule says that if (a) β is to be attached to α, and α was intended to support β, and (b) α entails a generic, of which β and δ form an instance, and (c) δ is consistent with what A and I believe,[3] then normally, δ and $Result(\alpha, \beta)$ are inferred.

- **The Result Rule:**
  (a) $(\langle \tau, \alpha, \beta \rangle \wedge Isupport(\alpha, \beta) \wedge$
  (b) $\exists_{kb \wedge \tau}(\alpha, \phi) \wedge \exists_{kb \wedge \tau \wedge \delta}(\beta, \psi) \wedge instance(\phi, \psi) \wedge$
  (c) $consistent(\text{KB}_I \cup \text{KB}_A \cup \delta))$
  $\quad\quad > (Result(\alpha, \beta) \wedge \delta)$

The Result Rule does two things. First, it allows us to infer one discourse relation ($Result$) from those permitted by Cooperation. Second, it allows us to infer a new piece of information δ, in virtue of which $Result(\alpha, \beta)$ is true.

We might want further constraints on δ than that in (c); we might add that δ shouldn't violate expectations generated by the text. But note that the Result Rule doesn't choose between different δs that verify clauses (b) and (c). As we've mentioned, the theory needs to be extended to deal with the problem of choice, and it may be necessary to adopt strategies for choosing among alternatives, which take factors other than logical structure into account.

We have a similar rule for inferring $Evidence(\beta, \alpha)$ ("β is evidence for α"). The Evidence rule resembles the Result Rule, except that the textual order of the discourse constituents, and the direction of intentional support changes:

- **The Evidence Rule:**
  (a) $(\langle \tau, \alpha, \beta \rangle \wedge Isupport(\beta, \alpha) \wedge$
  (b) $\exists_{kb \wedge \tau}(\alpha, \phi) \wedge \exists_{kb \wedge \tau \wedge \delta}(\beta, \psi) \wedge instance(\phi, \psi) \wedge$
  (c) $consistent(\text{KB}_I \cup \text{KB}_A \cup \delta))$
  $\quad\quad > (Evidence(\beta, \alpha) \wedge \delta)$

We have seen that clause (a) of the Result Rule is satisfied in the analysis of (1) in Context 1. Now, let δ be the proposition that the House Bill 1711 is bad for big

---

[2]Asher (1993) discusses this point in relation to *Contrast*: the discourse marker *but* is used coherently only if the semantic content of the constituents it connects do indeed form a contrast: compare *Mary's hair is black but her eyes are blue*, with *?Mary's hair is black but John's hair is black.*

[3]Or, more accurately, δ must be consistent with what I himself believes, and what he believes that A believes. In other words, KB$_A$ is I's model of A's KB.

business (written as $bad(1711)$). This is consistent with $\text{KB}_I \cup \text{KB}_A$, and so clause (c) is satisfied. Clause (b) is also satisfied, because (i) α entails Bush vetoes bills that are bad for big business—i.e., $\perp_{KB_{A\tau}}(\alpha, \phi)$ holds, where ϕ is $\forall x((bill(x) \wedge bad(x)) > veto(bush, x))$; (ii) $\beta \wedge \delta$ is $bill(1711) \wedge veto(bush, 1711) \wedge bad(1711)$; and so (iii) $instance(\phi, \beta \wedge \delta)$ and $\perp_{KB_{A\tau\wedge\delta}}(\beta, \beta \wedge \delta)$ both hold.

So, when interpreting (1) in Context 1, two rules apply: Narration and the Result Rule. But the consequent of Narration already conflicts with what is known; that the discourse relation between α and β must satisfy the Belief Property. So the consequent of the Result Rule is inferred: δ (i.e., House Bill 1711 is bad for big business) and $Result(\alpha, \beta)$.[4]

These rules show how (1) can make the knowledge that the house bill is bad for big business moot; one does not need to know that the house bill is bad for big business prior to attempting discourse attachment. One can infer it at the time when discourse attachment is attempted.

Now suppose that we start from different premises, as provided by Context 2: $\mathcal{B}_A\mathcal{B}_I\beta$, $\mathcal{B}_A\neg\mathcal{B}_I\alpha$ and $\mathcal{W}_A\mathcal{B}_I\alpha$. That is, I thinks A believes that I believes Bush will veto the bill, and I also thinks that A wants to convince him that Bush supports big business. Then the 'intentional' line of reasoning yields different results from the same observed behaviour—A's utterance of (1). Using APS1 again, but substituting $\mathcal{B}_I\alpha$ for ϕ instead of $\mathcal{B}_I\beta$, I concludes $\mathcal{B}_A(\perp_{kb}(\langle\tau, \alpha, \beta\rangle \wedge Info(\alpha, \beta), eventually(\mathcal{B}_I\alpha))$.[5] So $Isupports(\beta, \alpha)$ holds. Now the antecedent to Cooperation is verified, and so in the monotonic component of CE, we infer that α and β must be connected by a discourse relation R' such that $(\mathcal{B}_I\beta \wedge R'(\alpha, \beta)) > \mathcal{B}_I\alpha$. As before, this restricts the set of permitted discourse relations for attaching β to α. But unlike before, the textual order of α and β, and their direction of intentional support mismatch. The rule that applies this time is the Evidence Rule. Consequently, a different discourse relation is inferred, although the same information δ—that House Bill 1711 is bad for big business—supports the discourse relation, and is also be inferred.

In contrast, the antecedents of the Result and Evidence Rules aren't verified in (3). Assuming I knows about the legislative process, he knows that if George Bush is a weak willed president, then normally, he won't veto bills. Consequently, there is no δ that is consistent with his KB, and sanctions the *Evidence* or *Result* relation. Since I cannot infer which of the permitted discourse relations holds, and so by contraposing the axiom Cooperation, α doesn't *Isupport* β. And so I has failed to conclude what A intended by his discourse action. It can no longer be a belief that it will eventually

lead to I believing β, because otherwise $Isupport(\alpha, \beta)$ would be true via the rule Intends To Support. Consequently, I cannot infer what discourse relation to use in attachment, yielding incoherence.

FROM INFORMATION TO INTENTIONS: CONTEXT 3

Consider the interpretation of (1) in Context 3: I has no knowledge of A's communicative intentions prior to witnessing his linguistic behaviour, but he does know that the House Bill 1711 is bad for big business. I has sufficient information about the semantic content of α and β to infer $Result(\alpha, \beta)$, via a rule given in Lascarides and Asher (1991):

- **Result**
  $(\langle\tau, \alpha, \beta\rangle \wedge cause(\alpha, \beta)) > Result(\alpha, \beta)$

$Result(\alpha, \beta)$ has the Belief Property, and I reasons that from believing α, he will now come to believe β. Having used the information structure to infer discourse structure, I must now come to some conclusions about A's cognitive state.

Now suppose that $\mathcal{B}_A\mathcal{B}_I\alpha$ is in I's KB. Then the following principle of Charity allows I to assume that A was aware that I would come to believe β too, through doing the discourse attachment he did:

- **Charity:** $\mathcal{B}_I\phi > \mathcal{B}_A\mathcal{B}_I\phi$

This is because I has inferred $Result(\alpha, \beta)$, and since *Result* has the belief property, I will come to believe β through believing α; so substituting β for ϕ in Charity, $\mathcal{B}_A\mathcal{B}_I\beta$ will become part of I's KB via Defeasible Modus Ponens. So, the following is now part of I's KB: $\mathcal{B}_A(\perp_{kb}(\langle\tau, \alpha, \beta\rangle \wedge Info(\alpha, \beta)), eventually(\mathcal{B}_I\beta))$. Furthermore, the assumption that A's discourse behaviour was intentional again yields the following as part of I's KB: $\mathcal{I}_A(\langle\tau, \alpha, \beta\rangle \wedge Info(\alpha, \beta))$. So, substituting $\mathcal{B}_I\beta$ and $\langle\tau, \alpha, \beta\rangle \wedge Info(\alpha, \beta)$ respectively for ϕ and ψ into the Practical Syllogism, we find that clause (b) of the premises, and the conclusion are verified. Explanatory reasoning on the Practical Syllogism this time permits us to infer clause (a): A's communicative goals were to convince I of β, as required.

The inferential mechanisms going from discourse structure to intentional structure are much less well understood. One needs to be able to make some suppositions about the beliefs of A before one can infer anything about his desires to communicate, and this requires a general theory of commonsense belief attribution on the basis of beliefs that one has.

IMPERATIVES AND PLAN UPDATES

The revision of intentional structures exploits modes of speech other than the assertoric. For instance, consider another discourse from Moore and Pollack (1992):

---

[4]We could have a similar rule to the Result Rule for inferring $Evidence(\alpha, \beta)$ in this discourse context too.

[5]Given the new KB, the antecedent of APS1 would no longer be verified if we substituted ϕ with $\mathcal{B}_I\beta$.

(2)a. Let's go home by 5.

   b. Then we can get to the hardware store
      before it closes.

   c. That way we can finish the bookshelves tonight.

Here, one exploits how the imperative mode affects reasoning about intentions. Sincere Ordering captures the intuition that if A orders α, then normally he wants α to be true; and Wanting and Doing captures the intuition that if A wants α to be true, and doesn't think that it's impossible to bring α about, then by default he intends to ensure that α is brought about, either by doing it himself, or getting someone else to do it (cf. Cohen and Levesque, 1990a).

- **Sincere Ordering:**
  $Imp(\alpha) > \mathcal{W}_A(\alpha).$

- **Wanting and Doing:**
  $(\mathcal{W}_A\alpha \land \neg\mathcal{B}_A\neg eventually(\mathcal{R}\alpha)) > \mathcal{I}_A(\mathcal{R}\alpha)$

These rules about A's intentional structure help us analyse (2). Let the logical forms of (2a-c) be respectively α, β and γ. Suppose that we have inferred by the linguistic clues that $Result(\alpha,\beta)$ holds. That is, the action α (i.e., going home by 5pm), results in β (i.e., the ability to go to the hardware store before it closes). Since α is an imperative, Defeasible Modus Ponens on Sincere Ordering yields the inference that $\mathcal{W}_A\alpha$ is true. Now let us assume that the interpreter I believes that the author A doesn't believe that α's being brought about is impossible. Then we may use Defeasible Modus Ponens again on Wanting and Doing, to infer $\mathcal{I}_A(\mathcal{R}\alpha)$. Just how the interpreter comes to the belief, that the author believes α is possible, is a complex matter. More than likely, we would have to encode within the extension of DICE we have made, principles that are familiar from autoepistemic reasoning. We will postpone this exercise, however, for another time.

Now, to connect intentions and plans with discourse structure, we propose a rule that takes an author's use of a particular discourse structure to be *prima facie* evidence that the author has a particular intention. The rule Plan Apprehension below, states that if α is a plan that A intends to do, or get someone else to do, and he states that δ is possible as a *Result* of this action α, then the interpreter may normally take the author A to imply that he intends δ as well.

- **Plan Apprehension:**
  $(Result(\alpha,\beta) \land \mathcal{I}_A(\mathcal{R}\alpha) \land \beta = can(\delta)) > \mathcal{I}_A(\mathcal{R}(\alpha;\delta))$

We call this rule Plan Apprehension, to make clear that it furnishes one way for the interpreter of a verbal message, to form an idea of the author's intentions, on the basis of that message's discourse structure.

Plan Apprehension uses discourse structure to attribute complex plans to A. And when attaching β to α, having inferred $Result(\alpha,\beta)$, this rule's antecedent is verified, and so we infer that δ—which in this case is to go to the hardware store before it closes—as part of A's

plan, which he intends to bring about, either himself, or by getting another agent to do it.

Now, we process γ. *That way* in γ invokes an anaphoric reference to a complex plan. By the accessibility constraints in SDRT, its antecedent must $[\alpha;\delta]$, because this is the only plan in the accessible discourse context. So γ must be the DRS below: as a result of doing this plan, finishing the bookshelves (which we have labelled ϵ) is possible:

$(\gamma) Result([\alpha;\delta], can(\epsilon))$

Now, substituting $[\alpha;\delta]$ and ϵ for α and β into the Plan Apprehension Rule, we find that the antecedent to this rule is verified again, and so its consequent is non-monotonically inferred: $\mathcal{I}_A(\mathcal{R}(\alpha;\delta;\epsilon))$. Again, I has used discourse structure to attribute plans to A.

Moore and Pollack (1992) also discuss one of I's possible responses to (2):

(4) We don't need to go to the hardware store.
    I borrowed a saw from Jane.

Why does I respond with (4)? I has inferred the existence of the plan $[\alpha;\delta;\epsilon]$ via Plan Apprehension; so he takes the overall goal of A to be ϵ (to finish the bookshelves this evening). Intuitively, he fills in A's plan with the reason why going to the hardware store is a subgoal: I needs a saw. So A's plan is augmented with another subgoal ζ, where ζ is to buy a saw, as follows: $\mathcal{I}_A(\mathcal{R}[\alpha;\delta;\zeta;\epsilon])$. But since ζ holds, he says this and assumes that this means that A does not have to do α and δ to achieve ζ. To think about this formally, we need to not only reason about intentions but also how agents update their intentions or revise them when presented with new information. Asher and Koons (1993) argue that the following schema captures part of the logic which underlies updating intentions:

- $Update\mathcal{I}_A(\mathcal{R}[\alpha_1;\ldots;\alpha_n]), \mathcal{D}(\alpha_1;\ldots;\alpha_j) \models$
  $\mathcal{I}_A(\mathcal{R}[\alpha_{j+1};\ldots;\alpha_n]) \land \neg\mathcal{I}_A(\mathcal{R}[\alpha_1;\ldots;\alpha_j])$

In other words, if you're updating your intentions to do actions α_1 to α_n, and α_1 to α_j are already done, then the new intentions are to do α_{j+1} to α_n, and you no longer intend to do α_1 to α_j.

The question is now: how does this interact with discourse structure? I is attempting to be helpful to A; he is trying to help realize A's goal. We need axioms to model this. Some key tools for doing this have been developed in the past couple of decades—belief revision, intention and plan revision—and the long term aim would be to enable formal theories of discourse structure to interact with these formal theories of attitudes and attitude revision. But since a clear understanding of how intentions are revised is yet to emerge, any speculation on the revision of intentions in a particular discourse context seems premature.

CONCLUSIONS AND FURTHER WORK

We have argued that it is important to separate reasoning about mental states from reasoning about discourse structure, and we have suggested how to integrate a formal theory of discourse attachment with commonsense reasoning about the discourse participants' cognitive states and actions.

We exploited a classic principle of commonsense reasoning about action, the Practical Syllogism, to model I's inferences about A's cognitive state during discourse processing. We also showed how axioms could be defined, so as to enable information to mediate between the domain, discourse structure and communicative intentions.

Reasoning about intentional structure took a different form from reasoning about discourse attachment, in that explanatory reasoning or abduction was permitted for the former but not the latter (but cf. Hobbs *et al*, 1990). This, we argued, was a principled reason for maintaining separate representations of intentional structure and discourse structure, but preserving close links between them via axioms like Cooperation. Cooperation enabled I to use A's communicative intentions to reason about discourse relations.

This paper provides an analysis of only very simple discourses, and we realise that although we have introduced distinctions among the attitudes, which we have exploited during discourse processing, this is only a small part of the story.

Though DICE has used domain specific information to infer discourse relations, the rules relate domain structure to discourse structure in at best an indirect way. Implicitly, the use of the discourse update function $\langle \tau, \alpha, \beta \rangle$ in the DICE rules reflects the intuitively obvious fact that domain information is filtered through the cognitive state of A. To make this explicit, the discourse community should integrate work on speech acts and attitudes (Perrault 1990, Cohen and Levesque 1990a, 1990b) with theories of discourse structure. In future work, we will investigate discourses where other axioms linking the different attitudes and discourse structure are important.

REFERENCES

Asher, Nicholas (1986) Belief in Discourse Representation Theory, *Journal of Philosophical Logic*, 15, 127–189.

Asher, Nicholas (1987) A Typology for Attitude Verbs, *Linguistics and Philosophy*, 10, pp125–197.

Asher, Nicholas (1993) *Reference to Abstract Objects in Discourse*, Kluwer Academic Publishers, Dordrecht, Holland.

Asher, Nicholas and Koons, Robert (1993) The Revision of Beliefs and Intentions in a Changing World, in *Precedings of the* AAI *Spring Symposium Series: Rea-soning about Mental States: Formal Theories and Applications*.

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

Asher, Nicholas and Singh, Munindar (1993) A Logic of Intentions and Beliefs, *Journal of Philosophical Logic*, 22 5, pp513–544.

Bratman, Michael (forthcoming) *Intentions, Plans and Practical Reason*, Harvard University Press, Cambridge, Mass.

Cohen, Phillip R. and Levesque, Hector J. (1990a) Persistence, Intention, and Commitment, In Philip R. Cohen, Jerry Morgan and Martha E. Pollack (editors) *Intentions in Communication*, pp33–69. Cambridge, Massachusetts: Bradford/MIT Press.

Cohen, Phillip R. and Levesque, Hector J. (1990b) Rational Interaction and the Basis for Communication, In Philip R. Cohen, Jerry Morgan and Martha E. Pollack (editors) *Intentions in Communication*, pp221–256. Cambridge, Massachusetts: Bradford/MIT Press.

Grosz, Barbara J. and Sidner, Candice L. (1986) Attention, Intentions and the Structure of Discourse. *Computational Linguistics*, 12, 175–204.

Grosz, Barbara J. and Sidner, Candice L. (1990) Plans for Discourse. In Philip R. Cohen, Jerry Morgan and Martha E. Pollack (editors) *Intentions in Communication*, pp417–444. Cambridge, Massachusetts: Bradford/MIT Press.

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

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*, 277–332.

Kamp, Hans (1991) Procedural and Cognitive Aspects of Propositional Attitude Contexts, Lecture Notes from the Third European Summer School in Language, Logic and Information, Saarbrücken, Germany.

Lascarides, Alex and Asher, Nicholas (1991) Discourse Relations and Defeasible Knowledge, in *Proceedings of the 29th Annual Meeting of Computational Linguistics*, 55–63, Berkeley California, USA, June 1991.

Lascarides, Alex and Asher, Nicholas (1993a) Temporal Interpretation, Discourse Relations and Commonsense Entailment, in *Linguistics and Philosophy*, 16, pp437–493.

Lascarides, Alex and Asher, Nicholas (1993b) A Semantics and Pragmatics for the Pluperfect, in *Proceedings of the European Chapter of the Association for Computational Linguistics (EACL93)*, pp250–259, Utrecht, The Netherlands.

Lascarides, Alex, Asher, Nicholas and Oberlander, Jon (1992) Inferring Discourse Relations in Context, in *Proceedings of the 30th Annual Meeting of the Asso-*

ciation of Computational Linguistics, pp1–8, Delaware USA, June 1992.

Lascarides, Alex and Oberlander, Jon (1993) Temporal Connectives in a Discourse Context, in *Proceedings of the European Chapter of the Association for Computational Linguistics (EACL93)*, pp260–268, Utrecht, The Netherlands.

Moore, Johanna and Pollack, Martha (1992) A Problem for RST: The Need for Multi-Level Discourse Analysis *Computational Linguistics*, **18** 4, pp537–544.

Perrault, C. Ray (1990) An Application of Default Logic to Speech Act Theory, in Philip R. Cohen, Jerry Morgan and Martha E. Pollack (editors) *Intentions in Communication*, pp161–185. Cambridge, Massachusetts: Bradford/MIT Press.

Polanyi, Livia (1985) A Theory of Discourse Structure and Discourse Coherence, in Eilfor, W. H., Kroeber, P. D., and Peterson, K. L., (eds), *Papers from the General Session a the Twenty-First Regional Meeting of the Chicago Linguistics Society*, Chicago, April 25–27, 1985.

Thompson, Sandra and Mann, William (1987) Rhetorical Structure Theory: A Framework for the Analysis of Texts. In *IPRA Papers in Pragmatics*, **1**, 79–105.