To Abduce or not to Abduce – that is the Question.

On the Role of Abductive Inference within Discourse Construction

Matthias Hartung
Computational Linguistics Department, Heidelberg University
hartung@cl.uni-heidelberg.de

Philipp Cimiano
Institute AIFB, University of Karlsruhe
cimiano@aifb.uni-karlsruhe.de

December 2007

1 Introduction

In this paper, we discuss the role that abductive reasoning can play in the task of computing discourse relations for the construction of discourse representation. In order to introduce our line of argumentation, let's consider the following examples[1]:

(1) Fred jumped off a plane. He died.

(2) Fred forgot his parachute. He died.

(3) Peter sneezed. The napkin fell off the table.

In all these examples, we tend to infer a resultative connection between the first and second sentences (this assumption will be discussed and experimentally corroborated in more detail below). In particular, in the first and second examples, the *dying* of Fred is a consequence, and thus a *result* of his forgetting of the parachute or the jumping off the plane, respectively[2]. In the third example, the fact that the napkin *fell* of the table is clearly a result of the *sneezing*. However, due to the lack of surface discourse clues (cf. Marcu and Echihabi (2002)) which allow to infer a resultative connection, we should ask ourselves: what exactly are the properties of these discourses

[1] Examples (1) and (2) are taken from Danlos (2001), slightly paraphrased. Example (3) originally stems from Goldberg (1995), where the connection between both events is realized in a construction on sentence level: *Peter sneezed the napkin off the table.*

[2] Arguably, it is not an immediate result, but the consequence of a chain of events eventually leading to his death.

which trigger the inference that both sentences are connected via the discourse connection *Result* ? In the absence of surface lexical clues, there should definitely be other features which allow to infer a resultative relation. But which are these? As Asher and Lascarides (2003) have correctly argued, in many cases such inferences are due to lexical knowledge, and more precisely due to knowledge which has been linguistically conventionalized.

Such knowledge includes, for example, the fact that *pushing* typically results in an underspecified *movement*, that *sinking* presupposes some event causing the sinking, that the cause for somebody being annoyed is some (underspecified) experience which the person being annoyed is affected by. Such knowledge allows the resolution of examples as the following (from Asher and Lascarides (2003)):

(4) Max fell. John pushed him.

(5) The boat sank. The enemies torpedoed it.

(6) Mary was annoyed. Fred didn't call her.

However, in the above examples (1)–(3), there seems not to be such linguistically conventionalized knowledge available. We could assume that there is knowledge that if somebody *dies* then something happened which affected the person in question in such a way that she or he died. This might explain example (1) but surely not (2). For example (3), no linguistic triggers in the sense of linguistically conventionalized knowledge about sneezing and falling are involved.

In these cases, Asher and Lascarides argue, arbitrary knowledge needs to be incorporated into the process of computing discourse structure. However, it is doubtful that there is world knowledge available which will allow us to infer that jumping from a plane or forgetting a parachute causes death. And it is even less likely that there will be world knowledge available saying that sneezing causes napkins to fall off tables. What we can expect instead is world knowledge of the form: *If somebody falls from a high altitude (without using a parachute), he or she will be likely to die.* On the other hand, we can expect an axiom saying that *light forces will cause light objects to move.* This knowledge will however only lead to the appropriate inferences if we assume that the plane was flying and that the napkin was made of paper or some other light material. Such complex inferences can for sure not be assumed to be lexicalized. Thus, we argue in what follows that the computation of the resultative relation in the above examples is only possible because we make certain assumptions, i.e. that the plane was flying and that the napkin was made of a very light material (e.g. paper). We will argue that these assumptions can be understood as some sort of abduction to the best explanation.

The paper is structured as follows: we start off in the following Section 2 with a review of two important discourse theories which we actually aim

to reconcile, i.e. the *Interpretation as Abduction* approach of Hobbs et al. (1993) as well as *Segmented Discourse Representation Theory* (SDRT) of Asher and Lascarides (2003). As we will build on SDRT in our analyses in this paper, we introduce the necessary ingredients of SDRT in Section 3. In Section 4, we then present some motivating data of some psycholinguistic experiments we carried out which serve to corroborate our claims. In Section 5, we suggest how abduction can be integrated into SDRT. In Section 6, we discuss how our approach can be applied to some of the examples mentioned above. Section 7 concludes the paper with a discussion of open questions we will have raised and which remain for future work.

2 Modular vs. Monolithic Discourse Processing

In the field of discourse processing, we find those researchers who adhere to what we will call the *modular discourse processing paradigm* and those who adhere to the *monolithic discourse processing paradigm.*

Segmented Discourse Representation Theory (SDRT), for example, adheres to the *modular paradigm* as it clearly separates the processes of computing discourse structure from the process of interpreting it. In particular, Asher and Lascarides introduce two different logics: the *logic of information content* and the *logic of information packaging.* The logic of information content is a dynamic logic in the DRT tradition (Kamp and Reyle, 1993) which allows to model-theoretically interpret Segmented Discourse Representation Structures (SDRSs). The so-called *glue logic*, which is part of the logic of information packaging, is a non-monotonic logic with restricted quantification which is used to infer discourse relations. The glue logic is thus simpler than the logic of information content and has only restricted access to world knowledge and cognitive states. The rationale for adhering to the modular paradigm is the argumentation that discourse computation should be decideable, while discourse interpretation might not be so (we refer the interested reader to Asher and Lascarides (2003) for a more detailed elaboration of this idea).

The *monolithic* paradigm typically models the construction of discourse, in particular the inference of discourse relations, in the same logic which is also used to interpret the discourse, i.e. the one which describes truth conditions of the discourse. Approaches following the monolithic paradigm are the ones of Hobbs et al. (1993) as well as the one presented in Cimiano (2006). The basic idea of Hobbs et al. is summarized as follows:

To interpret a sentence: prove the logical form of the sentence, together with the constraints that predicates impose on their arguments, allowing for coercions, merging redundancies where possible, and making assumptions where necessary. (Hobbs et al. 1993, p70).

The last point – augmenting the discourse by additional assumptions

which are not explicitly mentioned in the context – is essentially what abductive reasoning within discourse interpretation amounts to.
What distinguishes abduction from other logical principles, e.g. deduction, is the validity of the following conclusion:

(7) $\forall x : p(x) \supset q(x)$
$\underline{q(a)}$
$p(a)$

Further, in Hobbs et al. (1993), abduction is weighted in the sense that there is a cost associated with assuming information. So ultimately we infer those ground atoms (such as $p(a)$ in the above implication) which have the least costs associated with them, thus providing the "cheapest" explanation for why $q(a)$ is true.

The common ground between Asher and Lascarides (2003) and Hobbs et al. (1993) is their approach to discourse interpretation via a coherent discourse structure which is modelled by means of discourse relations connecting single discourse segments. One major cause of debate between them, however, is the question (i) as to how the computation of such a coherent discourse structure should be carried out, and (ii) as to whether abductive reasoning takes part in this task or not.

The main criticism of Asher and Lascarides against the abductive approach of Hobbs et al. concerns its unmodular nature. Moreover, they argue that there is no principled way to set the weights in order to select the preferred interpretations in all contexts.

In any case, it should be clear from the above discussion that the approach by Hobbs et al. is more amenable to account for the assumptions that are needed in some cases to infer a certain discourse relation and to which we have already referred to in the introduction.

In this article, we investigate the question whether SDRT is reconcilable with an abductive approach in the flavour of Hobbs et al.

In fact, we adhere to the general idea purported in Asher and Lascarides (2003) that the construction of discourse structure and discourse interpretation are separate processes with different underlying logics.[3] However, we argue that in some cases, in order to license the appropriate inferences, discourse update needs to incorporate additional information which is not explicitly given in the respective discourse in an abductive manner. While we agree with Asher and Lascarides that linguistic knowledge is essential for

[3]In particular, Asher and Lascarides argue that discourse construction needs to be decidable on the basis of the observation that humans are able to construct a coherent discourse even in cases where truth conditions cannot be evaluated. From this they conclude that the logic modeling discourse construction needs to be decideable. We would like to emphasize that such an argumentation assumes that human discourse processing follows the laws of logic, which is not granted. Nevertheless, we do not question this assumption and will follow SDRT's architectural choices.

discourse update, we show that in some cases complex world knowledge is also necessary. Though this is a claim which is also put forward by Asher and Lascarides, they do not make explicit the process by which complex world knowledge can be used to infer the appropriate discourse relations. In general, our strategy will be to remain as close as possible to SDRT and introduce our notions and changes on top of the theory as described in Asher and Lascarides (2003). We thus describe in more detail how complex world knowledge and discourse update interact. In particular, we rely on first-order reasoning for discourse update. However, we still adhere to the modular architecture in SDRT as the results of reasoning within a first-order logic (using the DRT calculus described in Kamp and Reyle (1996)) need not be part of the glue language but can be incorporated into it (as an additional knowledge source).

3 Segmented Discourse Representation Theory

As mentioned above, SDRT features a modular architecture which clearly separates the processes of discourse construction and discourse interpretation. The former is modeled by the logic of information packaging, while the latter is modeled through the logic of information content. Essentially, the logic of information packaging consists of the *glue logic*, which is used to infer discourse relations, the *discourse update* operation as well as the principle of *Maximizing Discourse Coherence* (MDC). We will briefly explain these ingredients of the logic of information packaging in the following.

The glue logic mainly consists of *glue axioms* of the form:

$$(?(\alpha, \beta, \lambda) \wedge \textit{some stuff}) > R(\alpha, \beta, \lambda)$$

In words, the above axiom can be circumscribed as follows: If β is to be attached via some (underspecified) relation to α in the constituent λ and *"some stuff"* holds, then we nonmonotonically infer that R is the relation in question. In doing so, $>$ is to be interpreted as a non-monotonic implication. In what follows, we give some examples of instantiations of such axioms[4]:

(8) $(?(\alpha, \beta, \lambda) \wedge \textit{occasion}(\alpha, \beta)) > \textsc{Narration}(\alpha, \beta, \lambda)$

(9) $(?(\alpha, \beta, \lambda) \wedge \textit{cause}_D(\sigma, \alpha, \beta)) > \textsc{Result}(\alpha, \beta, \lambda)$

(10) $(?(\alpha, \beta, \lambda) \wedge \textit{cause}_D(\sigma, \beta, \alpha)) > \textsc{Explanation}(\alpha, \beta, \lambda)$

(11) $(?(\alpha, \beta, \lambda) \wedge \textit{subtype}_D(\sigma, \beta, \alpha)) > \textsc{Elaboration}(\alpha, \beta, \lambda)$

[4]The axioms are slightly simplified for the purposes of this paper, in particular omitting the *Top* and *Aspect* predicates which are not crucial for our purposes (see Asher and Lascarides (2003), p. 206 for the complete axioms).

Here, $occasion(\alpha, \beta)$ stands for the fact that there is a natural sequence between events of the sort described by α and events of the sort described by β in the sense that the former typically precede the latter.

$cause_D(\sigma, \alpha, \beta)$ stands for *discourse permissible cause* meaning that the discourse σ provides evidence that α caused β. However, $cause_D$ merely represents evidence and says nothing about whether α actually caused β in the situation described by the discourse in question. The axioms above state that if it is the case that we have evidence that α could have caused β, then we will non-monotonically infer *Result* as rhetorical relation. Finally, $subtype_D(\sigma, \beta, \alpha)$ states that β is a part of the event α, such that *elaboration* is inferred. It is important to note that all predicates are subscripted with a D indicating that there is evidence in the (D)iscourse which permits to infer a corresponding relation. The inference that the relation in question holds is thus a non-monotonic one, performed on the basis of the available evidence in the discourse.

While some axioms lead to non-monotonic conclusions, other axioms are monotonic. This is, for example, the case for axioms which infer some rhetorical connection on the basis of surface clues, i.e.

$$(?(\alpha, \beta, \lambda) \wedge \textit{and-then}(\alpha, \beta)) \rightarrow Narration(\alpha, \beta, \lambda)$$

The above axiom should be read as follows: if β is to be attached via some rhetorical relation to α in the constituent λ and α and β are connected via the surface clues *'and then'*, then (montonically) infer that the relation in question is *Narration*.

The core of SDRT is the *discourse update* function; it is defined as follows (see Asher and Lascarides (2003), p. 218)[5]:

Definition 1 (SDRT Update) *Let avail-pairs(σ) be the set of pairs of labels:*

$$\{\langle \alpha, \lambda \rangle | \alpha \in \textit{avail-sites}(\sigma) \textit{ and } Succ_D(\lambda, \alpha)\}$$

Moreover, let S_σ be the set of all possible sequences of all possible subsets of avail-pairs(σ). And let $X \in S_\sigma$. Then:

1. $\Sigma_X(\sigma, K_\beta)$ *is the sequence of updates:*

$$\sigma + K_\beta + ?(\alpha_1, \beta, \lambda_1) + ... + ?(\alpha_i, \beta, \lambda_i)$$

   where $\langle \alpha_i, \lambda_i \rangle \in X$ is the i^{th} element of X; and

2. $update_{SDRT}(\sigma, K_\beta) = (\bigcup_{X \in S_\sigma} \Sigma_X(\sigma, K_\beta)) + LAST = \beta$

[5]Here we simplify the definition and assume that K_β is not presupposed. See Asher and Lascarides (2003) for the full definition taking into account that K_β can also be presupposed.

In the above definition, *avail-pairs* is a subset of all the sites to which β could attach (see Asher and Lascarides (2003), 212ff.), e.g. those accessible DRSs subject to other conditions such as the right frontier constraint (see Asher and Lascarides (2003), pp. 8-18). $Succ_D(\alpha, \lambda)$ denotes that α is dominated by λ in the discourse structure; the '+'-operator is the simple update operator which updates the current DRS with new information (see Asher and Lascarides (2003), pp. 216).

In essence, $update_{SDRT}(\sigma, K_\beta)$, i.e. the result of updating the DRS σ with the DRS K_β, is the union of all updates which result from all possible attachment sites for β together with labels for the resulting rhetorical connections. In some sense, SDRT update thus only imposes constraints on the possible attachments, leaving all the possibilities open which respect these constraints. Thus, in contrast to earlier versions of SDRT, this formulation of the update operation is completely declarative. We will keep this declarative nature in our extensions proposed in this paper.

The principle of *Maximize Discourse Coherence* (MDC) is then in charge of selecting the most *coherent* update from the set of all possible updates. MDC can be described as follows. Given two possible updates $\mathcal{K}$ and $\mathcal{K}' \in update_{SDRT}(\sigma, K_\beta)$, we will regard $\mathcal{K}$ as more coherent than $\mathcal{K}'$ (i.e. $\mathcal{K} \leq_{\sigma,\beta} \mathcal{K}'$) in the following cases:

- if $\mathcal{K}'$ is satisfiable, then so is $\mathcal{K}$ (in particular, this means that consistent SDRSs are preferred over inconsistent ones. Inconsistent SDRSs can arise from the fact that some updates violate (monotonic) consequences of discourse relations)
- every rhetorical connection verified by $\mathcal{K}$ is at least as maximal in its context as the rethorical connections in $\mathcal{K}'$ and there are at least as many discourse relations in $\mathcal{K}$ as in $\mathcal{K}'$.
- $\mathcal{K}$ resolves as many underspecifications as $\mathcal{K}'$ does.

With respect to the partial order introduced above, Asher and Lascarides (2003) formalize the principle of MDC as follows:

Definition 2 (Maximize Discourse Coherence) $\textit{Best-update}_{SDRT}(\sigma, \mathcal{K}_\beta) = \{\tau \in update_{SDRT}(\sigma, \mathcal{K}_\beta) : \tau \textit{ is } \leq_{\sigma,\beta}\textit{-maximal}\}$

The rules for inferring rhetorical relations in the glue logic have (restricted) access to the following information sources:

- lexically specified knowledge
- defeasible inferences from lexical information
- subcategorization frames
- subtype information

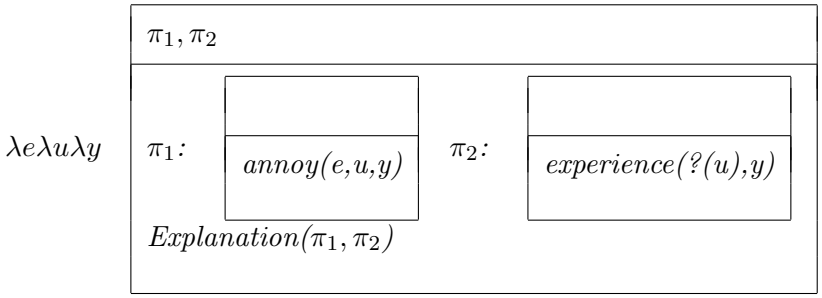

Figure 1: Lexical entry for *annoy*

- world knowledge

The following examples from Asher and Lascarides (2003) illustrate cases in which discourse relations can be inferred due to purely linguistic knowledge:

(12) a. Mary was annoyed. Peter didn't call.
b. Mary was annoyed. Peter called her at midnight.

According to the lexical entry given in Fig. 1, *annoy* presupposes an underspecified proposition u which can be resolved to any event which bears in its semantics an experiencer role that is filled by the same syntactic argument y that also participates in *annoy*. From the perspective of discourse processing, *annoy* opens a slot which can be filled by an underspecified event which licenses the corresponding semantic roles. If this slot-filling process is realized across sentence boundaries, these purely linguistic criteria give rise to a specific discourse relation.
The following example illustrates another case in which a discourse relation can be inferred due to specific linguistic knowledge, namely defeasible inferences which can be derived from lexical information (cf. Fig. 2):

(13) Max fell. John pushed him.

With regard to examples like these, Asher and Lascarides (2003) argue that the transitive use of a verb like *push* has to be distinguished from its intransitive use. This linguistic distinction affects the defeasible entailments which can be derived from these events: Being used transitively, *push* and similar verbs indicate movement of the object being pushed. As opposed to that, such an entailment is not licensed by intransitive patterns of those events, e.g. *Peter pushed against the door.*
What the examples (12) and (13) have in common is that in both cases a discourse relation can be inferred by way of a representation of linguistic knowledge which is rich enough to comprise semantic role information

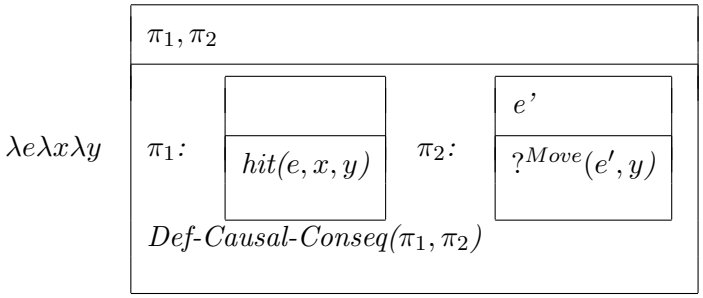

Figure 2: Example of defeasible inferences from lexical information

and syntactic argument structures. The problem is, however, that this very principled linguistic approach is not sufficient for resolving cases like (1)–(3) as mentioned in the introduction, which require deeper inferences based on common-sense knowledge.

4 Motivating Data

In order to get an intuition as to how people interpret the short discourses mentioned in the introduction as well as about what kind of inferences they draw, we carried out a small experiment involving the following examples (repeated here for presentation purposes):

(14) Fred jumped off a plane. He died.

(15) Peter sneezed. The napkin fell off the table.[6]

Within our experimental setting, we used (14) and (15) as sample cases for a set of questions we presented to a small number of subjects ($n = 20$). These questions are listed below, where the items in (16) relate to (14) and the one in (17) is related to (15).

(16) a. Was the plane flying when Fred jumped off ?
b. Did Fred use a parachute ?
c. What was the cause of his death ?

(17) What material was the napkin made of ?[7]

[6]This example originally stems from Goldberg (1995), where the connection between both events is realized in a construction on sentence level: *Peter sneezed the napkin off the table.*

[7]This question related to the example in (3) was brought up by Robert Porzel (personal communication).

Our interview yielded rather consistent results: A clear majority of subjects preferred a *resultative* interpretation for (14) and (15) as well. Regarding (14), we used (16c) as a control variable: 87.5% of our subjects considered reasons which are immediately related to Fred's jump off the plane (e.g. jump injuries, skull and brain damages, neck fractures) as causing his death. Among this subset of subjects, 92.9% were of the opinion that the plane was actually flying when Fred jumped off – as opposed to being on the ground for service, for instance.

Question (16b) yielded a slightly broader distribution of answers. Nevertheless, a majority of 64.2% of the subjects showing a preference for the resultative interpretation considered the option of Fred having used a parachute as rather unlikely.

On the other hand, a *narrative* interpretation of (14) seems not to trigger any inferences at all. Under these circumstances, the information explicitly given in our example discourses was regarded as insufficient in order to find reasonable answers to the questions.

One of our subjects described the reasons for his decision concerning (14) as follows: *"If both occurences are not directly related to each other, both possibilities for (16b) are equally likely. If he died as a result of jumping off the plane, it is clearly more likely that he did not use a parachute or that he did in fact use one but was not able to handle it."*

Concerning question (17), 83.3% of our subjects held the opinion that the napkin must have been made of paper or a light material. Even though we did not explicitly ask for justification of answers to (17), numerous subjects among those who considered the napkin to consist of paper argued that otherwise the sneezing might not have caused the napkin to fall off the table.

From our experiments, we can first conclude that people indeed tend to interpret our example sentences in a resultative way. Further, they draw interesting conclusions beyond what is explicitly stated in the text, i.e. that the plane was flying and that the napkin was made of paper. A crucial question is certainly whether they draw these inferences *because* they infer a resultative relation or they infer the resultative relation *because* they make certain assumptions. In the first case, we still need to explain how a resultative explanation is inferred in the absence of triggers for such a relation. In this article we explore the second possibility (without having clear empirical evidence for it), adopting the view that people make certain assumptions on the basis of common sense knowledge of such a nature that they are able to infer a discourse relation in order to establish coherence.

On a more general note, our experiments corroborate that there is strong evidence that human interpreters are in fact capable of interpreting discourse in such a way that a coherence structure can even be established in cases where the presence of a discourse relation is *not* predictable by the inventory of knowledge sources proposed by SDRT. In case it is true that in the above

discourses people make assumptions and draw inferences based on common sense knowledge to establish a discourse relation, we have to conclude that the architecture and mechanisms provided by SDRT to infer discourse relations are not enough. In particular, due to the modular architecture of SDRT, it is unclear as to how to incorporate additional assumptions or common sense reasoning.

5 Building the Bridge between SDRT and Abduction

With our extension to SDRT, we thus need to achieve two things: (i) make sure that the right discourse relation is inferred even in absence of surface clues in the case that some information needs to be assumed and (ii) update the discourse structure with the assumed (accommodated) information. Certainly, we could argue if the information should be really accommodated. In the tradition of dynamic semantics, we will assume that this is the right way as it is information implicitly implied by the discourse and thus an essential part of the meaning of the discourse. The right place to accomplish (i) is certainly the glue logic, while the right place to accomplish (ii) seems to be the discourse update operation as this is the operation which actually adds new information. We discuss our extension to the glue logic in subsection 5.1 as well as the extension to the SDRT update operation in subsection 5.2.

5.1 Extension of the Glue Logic

In this section, we discuss how the glue logic axioms need to be extended to take additional assumptions into account. In this article, we particularly focus on the $cause_D$ relation as all the examples we discuss involve resultative relations.

In general, within our approach we infer discourse relations from a specific configuration between DRSs K_α and K_β when K_α implies an alphabetic variant of K_β.

(18) $[(K_\alpha \models K'_\beta) \wedge K_\beta$ *is homomorphically embeddable in* $K'_\beta] > cause_D(\alpha, \beta)$

We rely on the DRS calculus described in Kamp and Reyle (1996) for inferences. Such inferences lead to (non-monotonically) inferring a discourse permissible cause in the sense that events of type α imply (as a consequence) events of type β. However, this says nothing about whether this actually holds for the events reported in the discourse in question; instead, the causal relation between those events is simply assumed in the form of a discourse permissible cause. The definition of *homomorphical embedding* is given in the following:

Definition 3 (Homomorphical Embedding) *A DRS K is homomorphically embeddable in K' if there is a 1-1 function* $m : U_K \rightarrow U'$ *where* $U' \subseteq U_{K'}$ *such that m(K) is an alphabetic variant of K and* $m(K) \subseteq K'$*, where the inclusion relation* $\subseteq$ *between DRSs is defined according to Kamp and Reyle (1996) as inclusion of the conditions.*

Putting these pieces together, the core of our system consists of the interplay between formally proving DRS-implications and the integration of axiomatic world knowledge. In order to infer causal relations between discourse segments, our system augments the linguistic information contained in the DRSs by implicit information coming from additional assumptions until a discourse relation can be derived in the glue logic. In the next section, we give an example of how formal operations on DRSs and abductive techniques based on axiomatic world knowledge are interrelated in our approach.

5.2 Extending SDRT update

We now have to extend the SDRT update function in such a way that it is able to accommodate assumptions:

Definition 4 (SDRT Update) *Let avail-pairs*(σ) *be the set of pairs of labels:*

$$\{\langle \alpha, \lambda \rangle | \alpha \in \textit{avail-sites}(\sigma) \textit{ and } Succ_D(\lambda, \alpha)\}$$

Moreover, let S_σ *be the set of all possible sequences of all possible subsets of avail-pairs*(σ)*. And let* $X \in S_\sigma$*. Then:*

1. $\Sigma_X(\sigma, K_\beta)$ *is the (set) of sequences of updates:*

$$\bigcup_{K'_\beta, K'_{\alpha_1}, ..., K'_{\alpha_{|X|}}} \{\sigma[K_{\alpha_i}/K_{\alpha_i} \oplus K'_{\alpha_i}] + [K_\beta \oplus K'_\beta] + ?(\alpha_1, \beta, \lambda_1) + ... + ?(\alpha_i, \beta, \lambda_i)\}$$

   where $\langle \alpha_i, \lambda_i \rangle \in X$ *is the* i^{th} *element of X; and*

2. $update_{SDRT}(\sigma, K_\beta) = (\bigcup_{X \in S_\sigma} \Sigma_X(\sigma, K_\beta)) + LAST = \beta$

In the definition above, the DRSs $K'_{\alpha_1}, ..., K'_{\alpha_{|X|}}$ contain the information which is assumed (and accommodated technically) by merging it with the DRSs corresponding to the attachment points in X.

The best-update then is defined as follows:

Definition 5 $Best{-}update_{SDRT}(\sigma, \mathcal{K}_\beta) = \{\tau \in update_{SDRT}(\sigma, \mathcal{K}_\beta) : \tau$ *is* $\leq_{\sigma,\beta}$ $-$*maximal and* τ *is minimal in the number of assumptions made*$\}$*.*

6 Application to Examples

Consider example (1) again. One obvious update according to the standard SDRT update operation is the following:

(19) π_1 : Fred jumped off a plane. π_2 : He died.

(20)

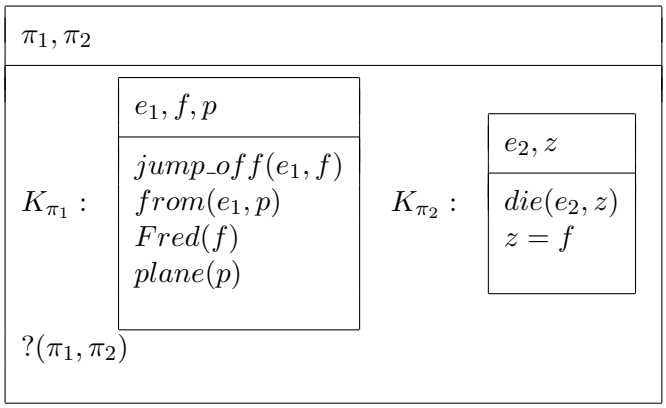

In our formulation of the SDRT update operation, this corresponds to an empty set of assumptions, i.e. $K'_{\pi_1} = \emptyset$. However, this does not allow to infer an appropriate discourse relation in the glue logic as there are neither discourse cues nor linguistic knowledge nor prototypical knowledge which can be used to infer such a relation. In particular, axiom (18) does not apply here as jumping from a plane does not necessarily imply dying. Instead of rejecting the discourse as incoherent, though, we try to establish coherence by updating the discourse with recourse to additional assumptions:

(21) K'_{π_1}:

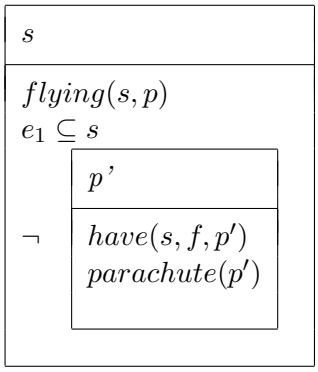

This means that in one possible update (which incorporates the additional assumptions as stated in K'_{π_1}) we assume that the plane p that Fred jumped off was flying during state s in which the jumping occurred and in which Fred did not have a parachute.

Suppose the following inference rules are given in the background knowledge:

(22)

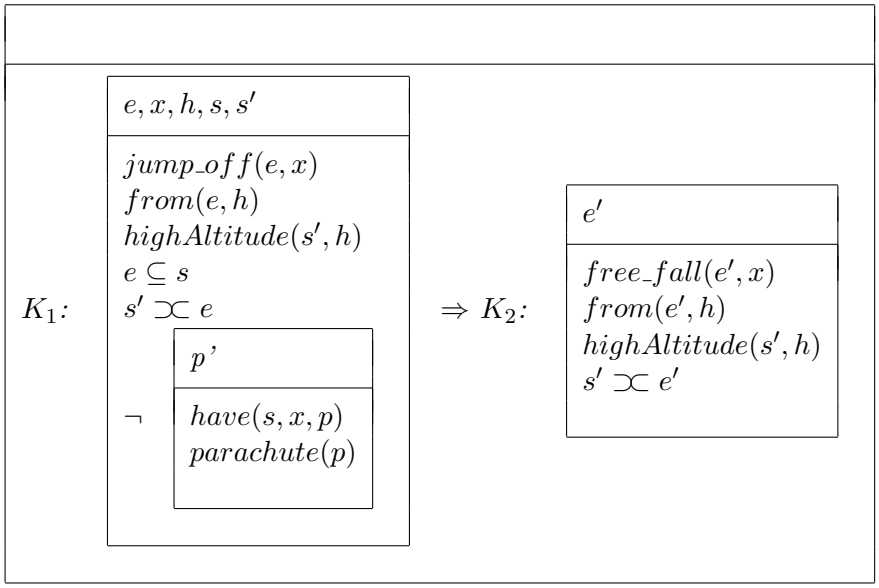

(23)

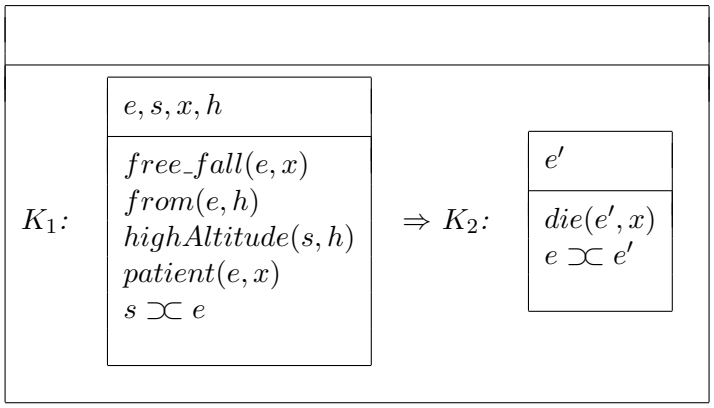

(24)

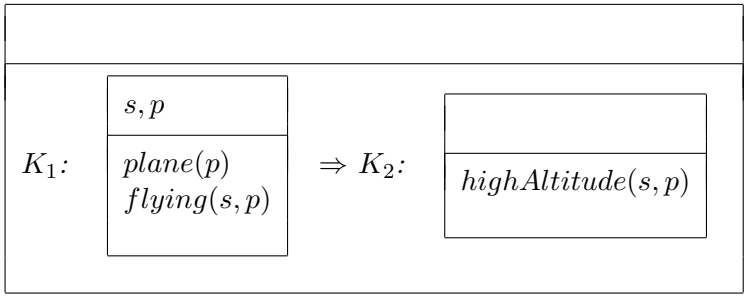

Applying these inference rules to $K_{\pi_1} \oplus K'_{\pi_1}$ leads to the following result. Note that this satisfies the first condition our axiom (18) imposes on inferring a discourse relation:

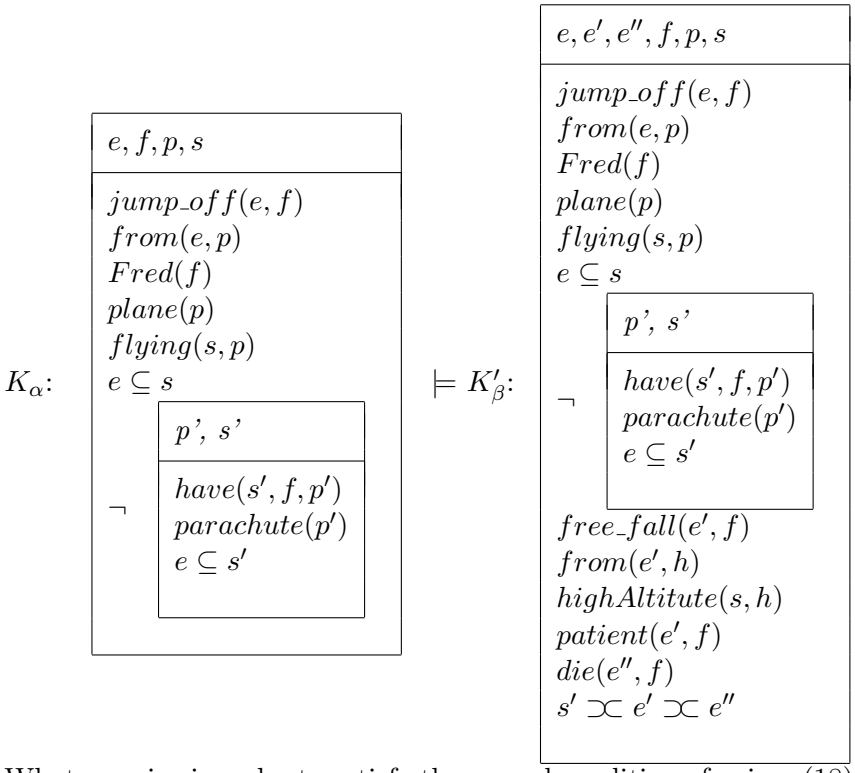

What remains in order to satisfy the second condition of axiom (18) is to prove that the discourse contains a DRS which is homomorphically embeddable into K'_β above. Obviously, this is the case for the DRS stating that Fred died (see K_{π_2} in (26); repeated as K_β below for the sake of convience and accordance with the notation introduced in (18)).[8]

(25) K_β:

e_2, z
$die(e_2, z)$ $z = f$

In more formal terms, we can state a function m mapping the discourse referents in K_β to their counterparts in K'_β with $m(e_2) = e''$ and $m(z) = f$. This mapping yields an alphabetic variant of K_β which is homomorphically embeddable into K'_β. Hence, both conditions that are necessary to infer a discourse permissible cause between e and e'' – and thus a resultative relation between π_1 and π_2 – are satisfied:

[8] Abstracting from the discourse referent z and the condition $z = f$ which are present in K_{π_2} in (26).

(26)

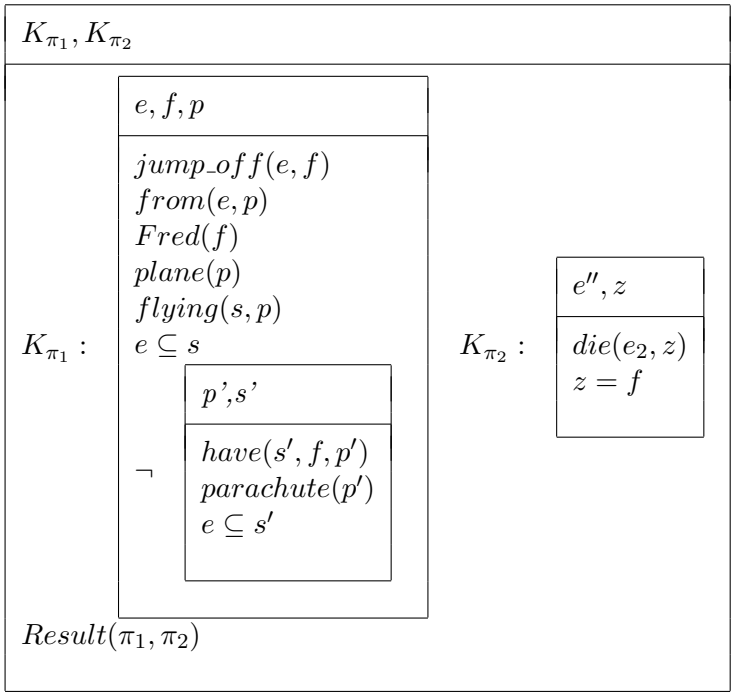

In the resolution above, the correct discourse relation has been inferred due to our new glue axiom (18) and the additional assumptions have been made explicit (accommodated). These assumptions correspond to those inferences that our experiments have shown that people actually draw when confronted with such a discourse.

The following example is more interesting in that it involves a presupposition, i.e. the verb 'forget' triggers the presupposition that there was some event in which Fred would have needed the parachute, but did not have it:

(27) π_1 : Fred forgot the parachute. π_2 : He died.

The semantic representation for π_1 would be:

(28)

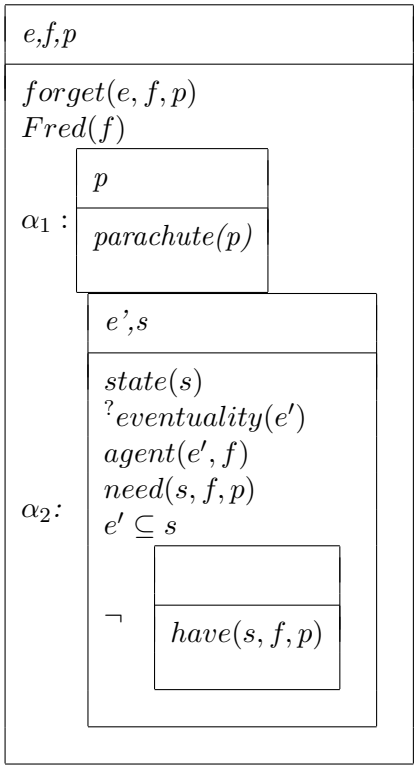

It is our intuition here that the construction of a coherent discourse depends on the interrelation between resolving the presuppositions α_1 and α_2, on the one hand, and inferring some additional assumptions, on the other. In the course of the resolution of α_1, the parachute p is accomodated to the global universe. Analogously, an underspecified event e' is accomodated, but remains still unresolved. We suggest that, linguistic criteria being absent, the resolution of e' is performed with regard to some background knowledge from which additional assumptions can be derived by means of abductive reasoning.

In this case, actually two abductive steps are required in order to arrive at a maximally coherent discourse: One possible explanation our knowledge base provides for someone's death is a free fall from high altitude – see (23). Note, however, that $free\_fall(e', f)$ is not a valid resolution for e' because of an incompatibility of the semantic roles of e' and $free\_fall$: the underspecified event e' requires f to be an *agent* whereas the corresponding syntactic argument in $free\_fall$ plays the *patient* role (compare (28) and (23)). Further reasoning being required, we proceed by consulting the knowledge base for a possible explanation for a free fall which satisfies the semantic role constraint as imposed by e'. Indeed, such an explanation can be found by assuming that Fred intentionally *jumped off* from high altitude.

(29)

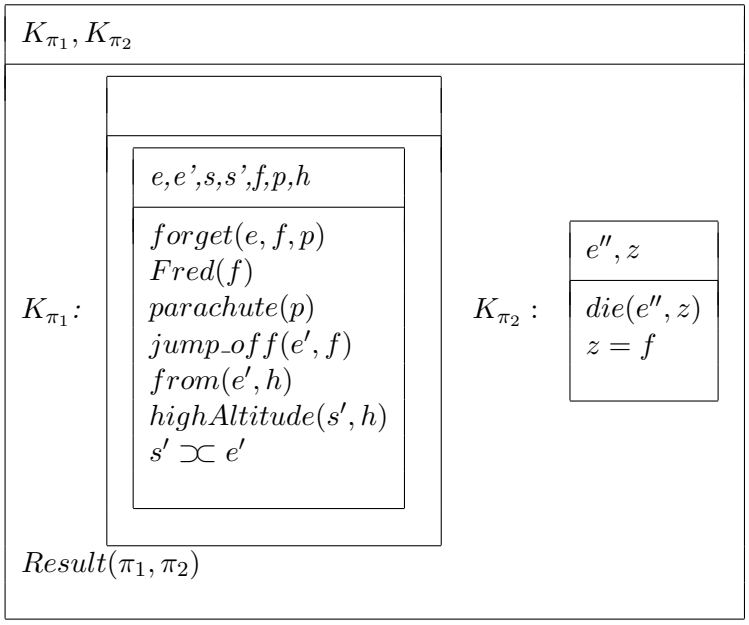

Note that (29) is only one among several possible solutions. Of course, other sequences of events are imaginable: For instance, Fred only might have had the plan to jump off a plane, but noticed he had forgotten his parachute and was involved in a car accident when he had been on his way back home in order to fetch it. In line with SDRT, however, we argue that updates can be ranked with regard to their degree of coherence. Indeed, our extension to SDRT Update disprefers the somewhat far-fetched scenario involving the car accident due to the greater number of assumptions it requires to establish coherene.

Certainly, the showcase demonstration this paper is intended to be suffers from the sparsity of the background knowledge base we considered. For example, we provided only one possible cause of death, namely free falls from high altitude. Practical reasoning, however, requires to deal with a huge number of concurrent explanations for one and the same event. This in turn underlines the importance of the minimality constraint we adopt from SDRT in order to prune solutions which are too far-fetched. The question as to how to balance the conflicting priorities between incorporating additional background knowledge necessary in order to establish maximal coherence, on the one hand, and restricting the approach to a minimal number of assumptions, on the other, is left for future work.

7 Conclusion and Outlook

It seems to us that there are many discourses for which there are no clear (linguistic) triggers to infer a specific rhetorical relation. Nevertheless, in the absence of such clear triggers, people are able to establish discourse co-

herence by inferring appropriate discourse relations. We have offered one possible explanation: People are able to make assumptions and draw inferences on the basis of common sense knowledge in order to infer a certain discourse relation. We presented an extension of SDRT (in particular of SDRT Update and the MDC principle) which allows to account for this observation and select that update including assumptions which maximizes discourse coherence and minimizes the number of assumptions made. In this line we reconciled SDRT with an abductive approach in the tradition of Hobbs et al. We maintained the declarative flavour of the SDRT update operation. Thus, our model does not say anything about procedural aspects related to the type of information to be assumed. Proposing an algorithmic framework to accomplish this seems an avenue for future work.

Overall, we feel that we have raised more questions than we have actually answered. First of all, our investigation suffers from the fact that we considered quite artificial examples from the literature. In particular, it would be an interesting question whether the examples discussed in this paper seem natural to speakers, i.e., whether discourses similar to the following ones actually occur in "real-world" text or conversation:

(30) Peter jumped off a plane. He died.

(31) ?Peter talked to his mother. He died.

According to our intuitions, a resultative interpretation is more likely for the first example, whereas in the second example the resultative interpretation is only possible if the fact that Peter talked to his mother initiated a chain of events which might or might not be causally related to the talking but eventually lead to his death. Assuming that both sentences are acceptable (which we doubt), the crucial question is as to which properties allow a resultative interpretation in the first case but prevent or at least make this interpretation less obvious in the second. Is it the fact that more assumptions need to be made in the second case? This would go in line with our proposal. However, such arguments ultimatively depend on the naturalness of examples such as the above. Our experimental investigations show that people indeed infer that the plane was flying in the first case. What we have not been able to clarify (and in fact it seems quite hard to create experimental conditions to verify this) is whether this inference is a consequence of the resultative interpretation or a precondition of it. We intend to further investigate this question in future experiments.

References

Asher, N. and Lascarides, A. (2003). *Logics of Conversation.* Cambridge University Press.

Cimiano, P. (2006). Ingredients of a First-Order account of Bridging. In *Proceedings of the International Workshop on Inference in Computational Semantics (ICoS)*.

Danlos, L. (2001). Event Coreference in Causal Discourses. In Bouillon, P. and Busa, F., editors, *The Language of Word Meaning*, pages 216–245. Cambridge University Press.

Goldberg, A. (1995). *Constructions. A Construction Grammar Approach to Argument Structure*. The University of Chicago Press.

Hobbs, J. R., Stickel, M., Appelt, D., and Martin, P. (1993). Interpretation as Abduction. *Artificial Intelligence*, 63:69–142.

Kamp, H. and Reyle, U. (1993). *From Discourse to Logics. Introduction to Modeltheoretic Semantics of Natural Language, Formal Logic and Discourse Representation Theory*. Kluwer, Dordrecht.

Kamp, H. and Reyle, U. (1996). A Calculus for First Order Representation Structures. *Journal of Logic, Language and Information*, 5:297–348.

Marcu, D. and Echihabi, A. (2002). An Unsupervised Approach to Recognizing Discourse Relations. In *Proceedings of the 40th Annual Meeting of the Association for Computational Linguistics (ACL-02)*, pages 365–372.