Bridging

Nicholas Asher
Department of Philosophy,
University of Texas at Austin,
Austin, Texas 78712,
USA
nasher@bertie.la.utexas.edu

Alex Lascarides
Centre for Cognitive Science and HCRC,
University of Edinburgh,
2, Buccleuch Place,
Edinburgh, EH8 9LW,
Scotland, UK
alex@cogsci.ed.ac.uk

Abstract

In this paper, we offer a novel analysis of bridging, paying particular attention to definite descriptions. We argue that extant theories don't do justice to the way different knowledge resources interact. In line with Hobbs (1979), we claim that the rhetorical connections between the propositions introduced in the text plays an important part. But our work is distinct from his in that we model how this source of information interacts with compositional and lexical semantics. We formalise bridging in a framework known as SDRT (Asher, 1993). We demonstrate that this provides a richer, more accurate interpretation of definite descriptions than has been offered so far.

1 Introduction

We aim to offer a formal model of bridging. We take bridging to be an inference that two objects or events that are introduced in a text are related in a particular way that isn't explicitly stated, and yet the relation is an essential part of the content of the text in the sense that without this information, the lack of connection between the sentences would make the text incoherent. Examples of bridging are illustrated in texts (1–4):

(1) I met two interesting people last night at a party.
The woman was a member of Clinton's Cabinet.

(2) In the group there was one person missing. It was Mary who left.

(3) John partied all night yesterday. He's going to get drunk again today.

(4) Jack was going to commit suicide. He got a rope.

In (1), *the woman* generates the presupposition that there's a unique salient woman in the context. The context doesn't supply one explicitly. However, the hearer draws the implicature that the woman is one of the two people the speaker met last night, and therefore, to guarantee the uniqueness of this antecedent, the other person must have been a man. In fact, without this inference the text would be incoherent, because there would be no connection between the objects or events described in the two sentences. So this implicature is a bridging inference.

While work on bridging inferences has typically concentrated on definite descriptions (e.g., Poesio 1994, Poesio, Vieira and Teufel 1997), other presupposition triggers generate bridging inferences too (Clark, 1977). For example, the *it*-cleft in (2) conveys the presupposition *someone left*. The hearer draws an implicature that the person missing from the group left, and indeed, Mary is that person. This inference is a bridging inference, since (2) would be incoherent without it: there would be no connection between the events or the objects. In (3), the presupposition triggered by *again* is that John got drunk before today. A bridging inference occurs here too: one infers that this previous occurrence of getting drunk is *concurrent* with the event of partying mentioned in the first sentence. Without this inference, one cannot compute how the events are connected, resulting in incoherence.

Karttunen (1974), Heim (1983, 1992) and van der Sandt (1992) have developed accounts of how presuppositions are satisfied in context. But these theories don't handle bridging, and so they don't explain the relevant inferences for (1–3). Indeed, it won't be possible to model all cases of bridging by refining presupposition satisfaction, because bridging occurs in the absence of presupposition triggers (Clark, 1977). Consider the example (4) taken from Charniak (1983). Here, there is an inference connected with the indefinite description *a rope*: one infers that it is to be used in the suicide. Without this link, there is no connection between the contents of the sentences, leading to text incoherence. As such, it's a bridging inference. And yet since it occurs in the absence of presupposition triggers, it can't be explained in terms of presupposition satisfaction.

In this paper, we will provide a formal theory of bridging based on the conjecture that it is a byproduct of discourse interpretation. In particular, bridging is part of the task of computing rhetorical connections between propositions introduced in a discourse. For example in (4), information conveyed by the second sentence is computed to be an *elaboration* of the information given by the first sentence. Part of this computation involves the inference that getting a rope is part of the plan to commit suicide: the rope is the intended instrument. A similar inference is involved with (2): the information in the second sentence serves to elaborate the first, and computing this involves inferring that Mary is the member of the group that's missing.

Our theory will be specified in a formal representation of discourse semantics known as SDRT (Asher, 1993), which incorporates rhetorical relations. An accompanying formal theory of pragmatics known as DICE (Lascarides and Asher, 1993) models how the construction of this discourse semantics is influenced by a wide variety of information. By mixing these ingredients, we hope to furnish a richer theory of bridging than has been attempted so far, where domain knowledge, compositional semantics, lexical semantics and rhetorical relations all play a central role.

This conjecture that bridging is a byproduct of discourse interpretation isn't new. Hobbs (1979), Hobbs *et al.* (1993) and Sperber and Wilson (1986) also propose this. But we approach discourse interpretation differently. Bridging for Hobbs *et al.* and Sperber and Wilson is part and parcel of figuring out the intended message or full understanding of the message. They equate the semantics of discourse with the task of integrating the clause that's currently being processed with the interpreter's beliefs. For Hobbs *et al.* (1993), this integration is a matter of abduction, whereas for Sperber and Wilson (1986) it is a matter of relevance.

We approach discourse interpretation differently. For us, bridging is a byproduct of computing the discourse structure of a discourse, which we view as a necessary *precondition* for discourse interpretation, as the interpretation of a discourse is for us compositional: a function of interpretation of the discourse's parts and how they are put together (viz. the discourse structure).[1]. We have argued elsewhere and will largely presuppose here that we need a logic different from the simple lambda calculus of standard semantics in order to construct discourse structure. But our notion of interpretation is still essentially tied to the goals of truth conditional accounts of meaning. For us there is a big distinction between getting the semantic form of the message and full understanding of it. A theory of discourse interpretation as we see it has two tasks: first, to specify a structure that has a coherent interpretation, and second to offer a model-theoretic interpretation of that structure. Full understanding takes that full structure and integrates it with the beliefs of the interpreter, and as such comes *after* discourse interpretation. In our view, we're after the linguistic content of the message (pragmatically and semantically determined). In contrast, Hobbs *et al.* and Wilson are after an integration of the content with beliefs—a theory of how beliefs are updated as a result of information present in the discourse. They are more ambitious than we are, but in turn we think that what they're after can't be analyzed illuminatingly in detail with the general ideas about inference that they have. From a computational perspective, there are also differences between our approach and theirs: full interpretation as pursued by Hobbs *et al.* and Sperber and Wilson involves inferences which aren't recursively enumerable (and perhaps shouldn't be). But the task of building a coherent discourse structure for interpretation—which encompasses bridging inferences—must be feasible for computational agents, if understanding is possible. As we will indicate below in section 4, the problem of computing bridging inferences is a decidable one our theory.

Bridging also occurs in the absence of definite descriptions, but in line with most research, we will focus our attention on cases involving definite descriptions. We will assume an existing compositional analysis of definite descriptions (Chierchia, 1995) and build a formal theory of bridging which is compatible with it. Although we think that from our discourse perspective Chierchia's analysis isn't quite right, we won't argue for that here. And our underlying theory of bridging in SDRT won't depend on the details of Chierchia's semantics.

2 Preliminaries and Some Simple Examples

We aim to provide a theory of how objects denoted by definite descriptions are related to previously described objects. For example:

(5) a. Lizzie met a dog yesterday.

b. The dog was very friendly.

(*The dog* in (5b) is identical to the dog mentioned in (5a)).

[1] In fact, we view the resolution of anaphora and the interpretation of presuppositions this way too (Asher and Lascarides, 1998)

(6) a. I took my car for a test drive.

b. The engine made a weird noise.

(*The engine* in (6b) is part of the car mentioned in (6a)).

(7) a. I've just arrived.

b. The camel is outside and needs water.

(*The camel* in (7b) is used as transport in the arrival mentioned in (7a)).

As we've stated, we will use Chierchia's (1995) compositional semantics of definite descriptions as input to the bridging which occurs at the discourse level.

Chierchia treats definite descriptions as anaphoric: *The N* denotes an N that's related in some anaphorically determined way B to an antecedent u. Chierchia (1995) and von Fintel (1994) have suggested that the Russellian uniqueness condition holds for definite descriptions so long as one includes this relation B, because it serves to restrict the domain. So Chierchia's analysis of *the N* is given in (8a). We will exploit the anaphoric resolution processes that already exist in DRT (Kamp and Reyle, 1993) to model bridging. So we will assume the (roughly) equivalent representation of definites in (8b):

(8) a. $\lambda Q.Q(\iota x(B(x,u) \wedge N(x)))$

b. $$\lambda e \lambda Q \begin{array}{|l|} \hline x, u, B \\ \hline Q(x,e) \\ N(x) \\ B(x,u) \\ B = ? \\ u = ? \\ \begin{array}{|l|} \hline z \\ \hline N(z) \\ B(z,u) \\ \hline \end{array} \Rightarrow \begin{array}{|l|} \hline \\ \hline z = x \\ \hline \end{array} \\ \hline \end{array}$$

B is an underspecified relation (as marked by the condition $B =?$), which must be further specified through connecting to the discourse context. Chierchia doesn't spell out this process. We intend to do this.[2]

Taking van der Sandt's (1992) view that presuppositions are anaphora (and so presupposed content can be viewed as those DR-conditions containing '?'), this analysis assumes that the presupposed part of definites is minimal: there is some antecedent (u) which is related in some way (B) to the individual referred to by the definite.

[2]We are aware that the proposed Russellian uniqueness condition is controversial, even when it comes in tandem with the restriction provided by $B(x,u)$. We believe that one can uphold Russellian uniqueness in these circumstances, but it isn't essential to our account of bridging itself. We have also assumed here that the uniqueness condition is part of the asserted content, rather than being presupposed; the latter case would be represented by making the uniqueness condition anaphoric in some respect. We are in fact agnostic about what the correct status is for the uniqueness conditions of definites, but see Asher and Lascarides (1998) for more detailed discussion of this issue.

How does one compute the value of B? Van der Sandt's (1992) theory of presupposition satisfaction in DRT gives us one clue. He suggests that presupposed content binds to an antecedent of the same content which is in an accessible part of the DRS representing the prior discourse context, if it can. This amounts to a preference for resolving B to identity. We will formally encode this preference. It provides a nice account of (5), for example. It predicts that B and u get resolved respectively to identity and the discourse referent introduced by the indefinite *a dog*, thereby capturing the intuition that the dog mentioned in (5b) is the same one that's mentioned in (5a).

But there are alternatives to B being identity. Clark (1977) provides a taxonomy of relations that include, among others: set membership (as in (1)); necessary parts; probable parts (as in (6)); inducible parts (as in (7)); reasons (as in (9)); causes (as in (10)); consequences (as in (11)); and concurrences (as in (3)).

(9) John had a suit on. It was Jane that he hoped to impress.

(10) John had a suit on. It was Jane who told him to wear it.

(11) John fell. What he broke was his arm.

We will build on Chierchia's analysis by spelling out a detailed formal theory via SDRT (Asher, 1993) and DICE (Lascarides and Asher, 1993) of exactly how B gets resolved to such connections. In contrast to von Fintel (1994), we will use rhetorical relations to do this. We explain why in the next section.

3 The Need for Rhetorical Relations

Bos *et al.* (1995) develop a theory of bridging by extending van der Sandt's work with lexical knowledge. The strategy is to include more information about word meaning in the discourse context, so that definite descriptions can link to objects that are introduced as part of this additional information. They assume a generative lexicon (Pustejovksy 1991, 1995), where lexical semantic information and real world knowledge are not seen as necessarily distinct. Instead, linguistic processes have limited access to world knowledge, which could therefore interact with knowledge of language and become conventionalised in various ways. In particular, lexical entries for artifacts have a *qualia structure*, which represents a limited amount of information about that artifact: what it's made up of, what one does with it, and so on.

Bos *et al* use the qualia structure to perform bridging inferences. They amend van der Sandt's model of presuppositions as follows: if it cannot be bound by identity to an accessible antecedent, then one tries to link it to elements of the qualia structure of entries in the accessible parts of the DRS. So in (6), the engine links successfully to the QUALIA:CONSTITUENCY value of the lexical entry for *car*, which in turn is in the accessible DRS representing the discourse context (6a), because this value in the lexical entry contains an engine (to reflect the fact that cars have engines as parts).

However, this extension to van der Sandt's theory has shortcomings. First, it fails to model bridging inferences in the absence of presupposition triggers (e.g., (4)). Secondly, although lexical semantics is a useful source of information for modeling bridging, it isn't sufficient. To illustrate the problem, consider (7). It's implausible to assume that the inference that I arrived by camel is achieved solely through lexical semantic information. For then the lexicon would essentially contain arbitrary domain knowledge, and consequently productive lexical phenomena would in general overgenerate word senses (cf. Verspoor, 1996).

There is a wide variety of knowledge that's used to support the bridging inference in (7). First, one uses the meanings of the words: for example, *arrive* is a motion verb, and so it is plausible to assume that there was a mode of transport. Second one uses world knowledge: for example, camels can be used as a mode of transport. But crucially, one uses the above lexical knowledge and world knowledge, as opposed to other knowledge, because this knowledge must be utilised to meet the coherence constraints imposed by the way (7b) connects to (7a). (7a) is stative, and according to Lascarides and Asher (1993), states normally provide background information. If this were the case here, however, then the camel being outside would temporally overlap the arrival, thereby blocking the camel from being part of the arrival. But another coherence constraint on *Background* is that the constituents must have a common topic (Lascarides and Asher, 1993). And if one is forced to assume that the camel has nothing to do with the arrival, then a suitable topic can't be constructed, leading ultimately to discourse incoherence. Intuitively, one tries to interpret constituents to obtain the best possible discourse coherence. Here, assuming the camel isn't the mode of transport leads to discourse incoherence. On the other hand, assuming the camel is the mode of transport allows us to interpret the discourse coherently—my arrival caused the camel to be outside, and so the propositions are connected by *Result*. Thus, if we formalise the coherence constraints of different rhetorical relations, together with the principle that you aim for discourse coherence, one can compute the link between the camel in (7b) and its discourse context.

Verifying coherence constraints imposed by the rhetorical relation that connects the sentences together has two important effects. First, it brings certain lexical knowledge and world knowledge into play. Second, it adds semantic content to the constituents that are connected (cf. Asher, 1993). We now know that the object described in (7b) isn't just a camel; it's a camel that I used as a mode of transport in the arrival event mentioned in (7a). Thus the added semantic content is a bridging inference in this case.

Grosz and Sidner (1986) offer an account of how connections between sentences in discourse serve to constrain the world knowledge that is brought into play in discourse interpretation; a feature we have just claimed is essential to bridging. They define a close relationship between the discourse segmentation of task oriented dialogues and the intentional structure of the plan that underlies the task described. Poesio (1993, 1994) merges Grosz and Sidner's framework with a situation theoretic semantics to account for how focus affects the denotation of definite descriptions. Tracking focus and allowing this to influence the available antecedents is a compelling idea. It enables one to capture the intuition that the uniqueness constraint on definite descriptions is closely related to the notion of saliency. For example, Poesio (1994) tracks the motion in (12) below, to infer that the focus of attention at the time when (12b) is processed is Dansville:[3]

[3] In fact, this is a slightly modified version of the example in Poesio (1994), in that we have put it in the past tense, rather than having a sequence of instructions. We modify the example here because we want to ignore speech acts in this paper.

(12) a. John took engine E1 from Avon to Dansville.

b. He picked up the boxcar and took it to Broxburn.

By doing this, he is able to infer that the boxcar is in Dansville—that is, he infers additional semantic content for (12b) as a result of tracking focus through the discourse structure.

Such an account is fine as far as it goes. However, it lacks a detailed formal, *general* theory of how the semantic content of constituents can be modified in the light of the way they connect together in the discourse structure.[4] But this flow from discourse structure to the addition of further semantic content is an essential feature of bridging. Moreover, Poesio's account of how motion determines focus produces the wrong results for other examples that feature other rhetorical relations. This is because Grosz and Sidner's model of discourse structure includes only two discourse relations—*dominance* and *satisfaction precedence*. This is too coarse grained to handle the different semantic effects that different rhetorical relations can have on bridging. So, for example, the rhetorical relation in (12a,b′) is *Parallel* rather than *Narration*:

(12) a. John took the engine E1 from Avon to Dansville.

b′. He also took the boxcar.

In contrast to (12a,b), the natural reading of (12a,b′) is one where the boxcar is in Avon. Presumably this is because of the different way that the sentences connect together, which in turn results in different spatio-temporal effects in the semantic content. But these spatial differences between *Narration* and *Parallel* aren't represented in the theory of discourse structure that Poesio adopts. Just as before, tracking the motion in (12a) leads to the focus of attention being Dansville at the point when (12b′) is processed. And so as in (12a,b), this predicts that the boxcar mentioned in (12b′) is in Dansville, contrary to intuitions. Computing that the boxcar was in Avon by recognising John's commonsense plan won't help either, since to recognise this plan involves computing the rhetorical connection that we've described between the sentences, and yet in Grosz and Sidner's theory, recognising commonsense plans is primary to constructing discourse structure.

One can view changes to semantic content caused by rhetorical connections as closely related to the concept of focus. The added content affects what's being talked about, and hence what's salient. So a general theory of how discourse structure affects semantic content can be viewed as contributing towards a general theory of focus. We will use this feature to model bridging inferences, by formalising the process in SDRT (Asher, 1993). Note that these inferences about the content of the description remain when *the boxcar* is replaced by *a boxcar*. So once again, bridging occurs in the absence of presupposition triggers.

We've given texts where different rhetorical relations have different effects on bridging. Text (13) provides evidence that rhetorical coherence can even override default world knowledge during bridging.

[4] Perhaps more seriously, these accounts also lack a general inference procedure for computing intentional structures from commonsense plans, and hence the ultimate discourse segmentation, which is assumed to be isomorphic to this intentional structure, is inferred by theory bound intuitions. For a detailed critique of this, see Asher and Lascarides (in press).

(13) a. John moved from Brixton to St. John's Wood.

b. The rent was less expensive.

Matsui (1995) tested subjects' judgements on where the rent was less expensive in (13). All the subjects knew the world knowledge that rents tend to be less expensive in Brixton than in St. John's Wood. But in spite of this, the majority of informants judged that in (13), the rent being talked about was in St. John's Wood, thereby drawing conclusions which conflicted with their world knowledge. Arguably, information about how the sentences connect together conflicts with the world knowledge, and ultimately wins over it. So if computing bridging ignores discourse structure, then the world knowledge would trigger the wrong results in (13).

We will explain (13) in terms of the rhetorical relation that's used to connect the constituents. (13b) is stative, and so supports a *Background* relation. However, intuitively, one prefers explanations of intentional changes (in this case, moving house), to simple background information that sets the scene for the change. Assuming that we *always* want to maximise discourse coherence, then even if default world knowledge conflicts with this, we infer both *Background* and *Explanation* for these texts. But the *Explanation* that John moved because the rent was less expensive is plausible only if the rent was less expensive in the place he went to: St. John's Wood.

The above texts where rhetorical information affects bridging pose challenges for extant theories. We need to analyse definite descriptions in a theory where information flow from rhetorical relations to the semantic content of constituents is taken into account. So we propose to use SDRT (Asher, 1993), where this information flow is a distinguishing feature. SDRT is a theory of discourse semantics designed to explore systematically the interface between semantics, pragmatics and discourse structure. To date it has been used to model several phenomena on the semantics/pragmatics interface (e.g., Asher 1993; Asher and Lascarides 1994, 1995, in press; Lascarides and Asher 1993, Lascarides and Copestake 1997, Lascarides and Oberlander 1993). Here, we will use it to interpret definite descriptions and to offer a new picture of bridging in general.

SDRT has three main advantages for our purposes. First, the way discourse structure affects and is affected by semantic content has already been studied extensively in this framework, and an adequate account of definite descriptions must make use of these effects. Second, the basic semantic framework which underlies SDRT (DRT), has already proved useful in specifying constraints on the interpretation of definite descriptions (van der Sandt 1992, Bos *et al* 1995). We will build on this work here. Finally, one of the main features of SDRT is the underlying axiomatic theory DICE (Discourse in Commonsense Entailment) which allows us to infer rhetorical relations, using semantic content and world knowledge as clues (Lascarides and Asher, 1993). DICE is distinctive in that it deals in a principled way with cases where different knowledge sources give conflicting clues about how to interpret a text. We will use this axiomatisation to provide a novel analysis of bridging that records the influence of background knowledge on the process, and we will use DICE's tools for conflict resolution to model why the default world knowledge is 'ignored' in (13).

4 A Crash Course in SDRT

Broadly speaking, there are two components to SDRT. First, there is a formal language with a compositional semantics, in which the content of discourse is represented (Asher, 1993). This is an extension of Discourse Representation Theory (DRT): discourse is represented as a segmented DRS (SDRS), which is a recursive structure of labelled DRSs that represent the clauses, and these labels are linked together with rhetorical relations, such as *Narration* and *Parallel* (cf. Hobbs (1985), Polanyi (1985), Thompson and Mann (1987) and others). The second component to SDRT is a formal theory of pragmatics known as DICE (Discourse in Commonsense Entailment) (Lascarides and Asher, 1993), which is used to build the SDRS of the text or dialogue. It uses a variety of knowledge sources to do this: for example, lexical and compositional semantics, domain knowledge and cognitive states.

DICE is a type of 'glue' logic, because it specifies how SDRSs connect together with rhetorical relations. The glue logic differs from the logic of 'information content' (i.e., the logic of the SDRSs themselves), whose validity problem is at least recursively enumerable (Asher, 1996). DICE exploits a much weaker language (Lascarides and Asher, 1993): it's a quantifier free fragment of a first order language augmented by a weak conditional operator $>$ ($P > Q$ means *If P, then normally Q*). The logic is decideable.

All axioms in DICE for computing rhetorical relations are of the form given in (14), where τ, α and β label SDRSs, $\langle\tau,\alpha,\beta\rangle$ means β is to be attached with a rhetorical relation to α, where α is available in the SDRS labelled τ that's built so far; *some stuff* is a gloss for relevant information, and R is a rhetorical relation:

(14) $(\langle\tau,\alpha,\beta\rangle \wedge \text{ some stuff}) > R(\alpha,\beta)$

While the glue logic and language are distinct from their counterparts at the level of information content, the glue language nevertheless exploits some aspects of information content in axioms of the form just given. To this end, we have devised an information transfer function μ from SDRSs into the DICE language, which allows DICE to use information about content to compute the rhetorical relation. Roughly, for each labelled SDRS $\pi : K_\pi$, μ takes conditions inside the SDRS K_π and turns them into predicates of its label π. So $\mu(K_\pi)(\pi)$ is a set of formulae of the form $\phi(\pi)$, where ϕ is a predicate. *some stuff* in (14) will be formulae of this kind.

For example, the schema `Narration` states: If β is to be attached to α and α and β describe events, then normally the rhetorical relation is *Narration*.[5] The `Temporal Consequence of Narration` is a coherence constraint on *Narration* in that it constrains the contents of the connected constituents: if *Narration*(α,β) holds, then α's event precedes β's.

- `Narration:` $(\langle\tau,\alpha,\beta\rangle \wedge event(e_\alpha) \wedge event(e_\beta)) > Narration(\alpha,\beta)$
- `Temporal Consequence of Narration:` $Narration(\alpha,\beta) \rightarrow e_\alpha \prec e_\beta$

Narration also constrains spatio-temporal trajectories of objects. Asher *et al.* (1996) derive the following constraint from `Narration` and commonplace assumptions about eventualities:

[5]Formulae like e_α and $event(e_\alpha)$ are a notational 'gloss' for propositional formulae of the form $\phi(\alpha)$ (Lascarides and Asher, 1993).

- `Spatial Consequence of Narration:`
  $(Narration(\alpha,\beta) \wedge actor(x,\alpha) \wedge actor(x,\beta)) \rightarrow loc(x, source(e_\beta)) = loc(x, goal(e_\alpha))$

In words, if $Narration(\alpha,\beta)$ holds and α and β share an actor x then the location of x is the same at the end of e_α and the onset of e_β.[6] There's also an axiom which states that narratives have a distinct common topic. We will introduce further axioms in later sections of this paper.

A distinctive feature of SDRT is that if the DICE axioms yield a nonmonotonic conclusion that $R(\alpha,\beta)$ holds, and information that's necessary for this to hold isn't already in the constituents K_α or K_β (e.g., $Narration(\alpha,\beta)$ is nonmonotically inferred, but the formula $e_\alpha \prec e_\beta$ and information about the spatial location of actors are not in K_α or in K_β), then this content is added to K_β in a constrained manner through the SDRS *Update* process. Asher and Lascarides (1998) give the detailed formal definition of discourse update for hierarchically structured contexts. An informal, simpler definition does for our purposes, however. Informally, $Update(K_\tau, K_\alpha, K_\beta)$ is an SDRS in which three things are added to the SDRS K_τ: (a) β is added to K_τ's list of discourse referents; (b) $R(\alpha,\beta)$ is added to K_τ's conditions, where $R(\alpha,\beta)$ follows nonmonotonically from DICE; and (c) $\beta : K_\beta^\varphi$ is also added to K_τ's conditions, where K_β^φ is just like the SDRS K_β, save that information φ that's necessary for $R(\alpha,\beta)$ to hold and that wasn't already in K_α or K_β has been added. In what follows, we will specify constraints on *Update*. And in certain cases, we will replace one update task with another. So $Update(K_\tau, K_\alpha, K_\beta) := Update(K_{\tau'}, K_{\alpha'}, K_{\beta'})$ means: replace the task of updating K_τ with K_β via attachment to K_α with the task of updating $K_{\tau'}$ with $K_{\beta'}$ via attachment to $K_{\alpha'}$.

As an illustrative example, consider (12a,b″):

(12) a. John took engine E1 from Avon to Dansville.
b″. He picked up a boxcar
c. and took it to Broxburn.

First, we use the grammar to build DRSs K_α and K_β for the (12a) and (12b″), and these receive the labels α and β respectively. The pronoun in K_β is resolved to *John* because in SDRT, the only available antecedents to pronouns are those that are DRS-accessible in the current constituent (in this case, K_β), or those that are DRS-accessible in the constituent K_α to which K_β is going to be attached. So *John* is the only choice. Defeasible Modus Ponens on `Narration` yields $Narration(\alpha,\beta)$. Modus Ponens on `Axiom on Narration` yields $e_\alpha \prec e_\beta$ (i.e., John's taking engine E1 from Avon to Dansville precedes his picking up a boxcar), and Modus Ponens on the `Spatial Consequence on Narration` yields that the shared actor John is in Dansville when he begins to pick up a box car, because this is the location of the goal of e_α. By the lexical semantics of *picking up* (see Asher and Sablayrolles, 1995), the location of the source of this event is the same as the location of its goal, and the object that's picked up is at this location. So the boxcar is in Dansville when it's picked up. The definition of SDRT *Update* guarantees that the content that's inferred as a result of the

[6] Some narratives imply that the actor x is in motion and so his location at the end of e_α is different from his location at the time of the onset of e_β. Our hypothesis is that these transitions are due to the presence of frame adverbials. Asher *et al.* (1996) are currently verifying this hypothesis for French with an extensive corpus-based search for counterexamples.

DICEinference that the text is narrative is added to K_β in the SDRS for (12a,b$''$). In particular, the information that the boxcar is in Dansville is added to K_β, and this can be viewed as a bridging inference, because it amounts to a relation between an object mentioned in the current clause and one mentioned previously, which arose out of coherence constraints on the discourse. Thus in contrast to Bos *et al.* (1995), SDRT can model bridging inferences in the absence of presupposition triggers.

5 Bridging with SDRT

We will use SDRT to resolve the underspecified conditions in Chierchia's analysis of definite descriptions. In effect, computing the bridging inference will occur as a byproduct of SDRT update.

5.1 Building the Bridges in SDRT

We now define how the anaphoric binding relation B and antecedent u, which are introduced by the compositional semantics of definites, are resolved in terms of the function *Update* introduced in section 4. There are four rules that define this. They are not part of the DICE language. Rather, they are meta-rules about how the semantic content of underspecified constituents and the function *Update* interact. The first rule captures van der Sandt's intuition that one uses identity to resolve bridging if one can. The second captures the intuition that bridging inferences must be plausible. The third captures the intuition that if updating the discourse with (underspecified) information adds semantic content which can act as a bridging implicature, then this added information is indeed a bridging implicature. And the last rule captures the intuition that we favour bridging implicatures that maximise discourse coherence.

First some notation: $\downarrow K$ means that the SDRS K is well-defined; that is, it contains no unresolved conditions of the form $x =?$ and every DRS in K is attached to another with a rhetorical relation. Furthermore, $K[\phi]$ is a formula, which is true if the SDRS K contains the condition ϕ, and $K[\phi'/\phi]$ is a term which denotes the SDRS which results from replacing ϕ in K with ϕ'. The first rule is given below. It states that if SDRS update with the binding relation B specified to identity is well-defined, then SDRS update must set B to identity.

- `If Possible Use Identity:`
  $(K_\beta[B =?] \wedge \downarrow \mathit{Update}(K_\tau, K_\alpha, K_\beta[\lambda x \lambda y x = y/B])) \rightarrow$
  $(\mathit{Update}(K_\tau, K_\alpha, K_\beta) := \mathit{Update}(K_\tau, K_\alpha, K_\beta[\lambda x \lambda y x = y/B]))$

This axiom reflects the preference noted by van der Sandt, for standard anaphoric binding over the alternatives. However, the condition this axiom imposes on standard anaphoric binding is stronger than van der Sandt's. In van der Sandt's theory, a presupposition will bind in any context where there's an accessible discourse referent satisfying the same content, and the result is satisfiable and informative. In contrast, `If Possible Use Identity` permits this binding only if van der Sandt's conditions hold, *and* one can compute a rhetorical relation with the result. Van der Sandt's weaker condition on binding is problematic in an example such as (15):[7]

(15) a. Boggs stood calmly by as Ryan struck out the hitter with a 95-mph pitch,
b. then he stepped up to the plate and
c. he hit the pitch out of the park.

In van der Sandt's analysis *the pitch* in (15c) will bind to the 95mph pitch mentioned in (15b), because his theory fails to account for the effects of temporal constraints. Moreover, we have shown elsewhere (Lascarides and Asher 1991, 1993) that an adequate account of the temporal constraints on discourse requires reasoning about discourse structure. In contrast, our theory will detect that the binding relation B in the representation of *the pitch* in (15b) cannot be identity, because the result will violate the temporal coherence constraints on *Narration* which, by Defeasible Modus Ponens on `Narration`, binds the propositions together in this discourse. Instead of B resolving to identity, the three axioms below for computing B will ensure that B resolves to 'thrown-by' and u to Ryan.

Note that `If Possible Use Identity` is monotonic rather than default. Giles Fauconnier (pc.) has offered (16) as a potential counterexample to its monotonicity: Resolving the binding relation to identity in (16) doesn't produce the intended reading.

(16) A foreign president visited the White House, but the President was busy.

But we believe resolving B to identity in (16) doesn't produce a well-defined SDRS, and so `If Possible Use Identity` doesn't apply in this case: If we do identify *the President* with the president mentioned in the first sentence, then the coherence constraints required by the relation *Contrast*, which is monotonically inferred from the cue word *but*, are violated, much in the same way as they're violated in (17), if one assumes that *he* refers to the foreign president.

(17) ?A foreign president$_i$ visited the White House, but he$_i$ was busy.

As we've seen, specifying B as identity doesn't always yield a well-defined SDRS. In this case, we allow the discourse context to guide us to a suitable specification for B. All the following rules suppose that $\neg\downarrow(Update(K_\tau, K_\alpha, K_\beta[\lambda x\lambda y x = y/B]))$ holds.

In general, there are many ways the underspecified parameter B could be made precise; some of these may be more plausible than others. We see here an important role for world knowledge. It specifies certain plausible ways of filling in the underspecified parameters in the presupposed material (cf. Beaver, 1994). To represent this we introduce a conditional operator: $P >_\diamond Q$ should be read as "If P, then it's plausible to assume Q". This specifies a weaker connection than $>$; it stipulates what is plausibly the case, rather than what is normally the case. In essence `Bridges are Plausible` below will restrict bridging as follows: the bridge must be built from $>_\diamond$ consequences of the semantic content of the constituents. That is, a bridge must be plausible:

[7] Thanks to Geoff Nunberg for this example.

- `Bridges are Plausible:`
  $(\beta[B = \phi; u = x/B =?; u =?] \wedge \langle \tau, \alpha, \beta \rangle \wedge R(\alpha, \beta)) \rightarrow$
  $((\mu(K_\tau)(\tau) \wedge \mu(K_\beta)(\beta) \wedge R(\alpha, \beta)) >_\diamond (B = \phi \wedge u = x))$

In words, if B and u are resolved to ϕ and x respectively, and β is attached to the constituent α in τ with a rhetorical relation R, then the semantic content of this (updated) discourse must make these bindings plausible. We'll see in section 7.2 that this rule will prove important when distinguishing (7a,b) from (7a,b′) (it's not plausible to assume fleas were the mode of transport):

(7) a. I just arrived.
b. The camel is outside and needs water.
b′. ?The fleas are outside and need water.

An axiomatisation of $>_\diamond$ would involve extensive discussion of commonsense reasoning with world knowledge, and so we gloss over it here.[8] However, if one believes that all bridging relations are constrained to fall within Clark's (1977) taxonomy, then one could capture this within this axiom `Bridges are Plausible`: one could assume that $>_\diamond$ is constrained so that the formula on the RHS of $\rightarrow$ in `Bridges are Plausible` holds only if the bridging relation ϕ is one of those that falls within Clark's taxonomy; i.e., ϕ must be a part-whole relation, or a set membership relation, or a causal relation, etc. This would amount to the assumption that only those relations within Clark's taxonomy form plausible candidates for bridging. There would be computational advantages to restricting ϕ this way, because this would provide a monotonic restriction on the search space of candidates for bridging. However, we remain agnostic as to whether Clark's taxonomy of bridging relations provides an *exhaustive* list of plausible bridging relations. There may be rich discourse contexts in which world knowledge permits a plausible bridging relation that lies outside this taxonomy.

Our third rule governing bridging inferences is `Discourse Structure (DS) Determines Bridging`. This rule captures the intuition that when the rhetorical relation used to connect the constituents gives us a particular way of resolving B, we do it that way. More formally, let $\mu(K_\beta)(\beta) \rightsquigarrow_* \mu(K_\phi)(\phi)$ mean: K_ϕ is a DRS which represents one way of resolving the underspecification in K_β. Then `DS Determines Bridging` is given below:

- `DS Determines Bridging:`
  Suppose: (a) $\mu(K_\tau)(\tau) \wedge \mu(K_\beta)(\beta) \wedge \langle \tau, \alpha, \beta \rangle \mathrel{|\!\approx} R(\alpha, \beta)$
  (b) $\mathrel{|\!\approx} \mu(K_\beta)(\beta) \rightsquigarrow_* \mu(K_\phi)(\phi)$; and
  (c) $\mathrel{|\!\approx} (R(\alpha, \beta) \wedge \mu(K_\tau)(\tau)) > \mu(K_\phi)(\phi)$
  Then $Update(K_\tau, K_\alpha, K_\beta) := Update(K_\tau, K_\alpha, K_\phi)$

In words, if we can infer the rhetorical connection R between the discourse context τ and the *underspecified* constituent β, and this relation R allows us to infer a particular resolution K_ϕ of the underspecified elements in β, then these specifications are incorporated into the

[8] Note that this constraint involving $>_\diamond$ is monotonic, and that $>_\diamond$ can be axiomatised within a decidable system (e.g., conditional probability theory). If $>_\diamond$ is axiomatised using conditional probabilities, then the decidability of $>$ remains unaffected.

SDRS update. This rule is called `DS Determines Bridging`, because computing the discourse structure serves to resolve B and u in β.

To see how `DS Determines Bridging` models the information flow from discourse structure to the content of definite descriptions, consider (12).

(12) a. John took engine E1 from Avon to Dansville.

b. He picked up the boxcar and took it to Broxburn.

We can use DICE to infer (12a,b) is narrative even before determining the underspecified elements B and u in (12b); we then use *Narration*'s coherence constraints to infer the boxcar is in Dansville, and this added content suffices to produce a plausible way of resolving $B =?$ and $u =?$ (B resolves to *in* and u to Dansville). `DS Determines Bridging` ensures we resolve them this way. The details of this analysis are given in the next section.

`DS Determines Bridging` deals with the case when the coherence constraints imposed by the rhetorical relation that's inferrable from the *underspecified* constituent β produces a plausible bridging inference. But the underspecified constituent β doesn't always contain sufficient information to determine the rhetorical relation; hence it may not be enough to determine the bridging inference. To deal with such cases, we state a rule which captures the intuition that people interpret text so as to maximise discourse coherence. It is a more restricted version of the Interpretation Constraint in DICE that was introduced in Lascarides *et al* (1996) for modeling word sense disambiguation, and this more restricted rule suffices for our purposes.

As background to this rule, we assume that rhetorical relations between constituents may be partially ordered with respect to the semantic content of the context. This reflects the fact that given the semantic content of the clauses, some rhetorical relations will produce a 'closer connection' or 'better coherence' than other rhetorical relations. We encapsulate this by introducing the following partial order: *Explanation* $>_{\tau,\alpha}$ *Background* means that it would be preferable to interpret β as an explanation for α, rather than background information—although both alternatives may be coherent, one is better than the other—and this is partly because of the content of τ and α.[9] The following rule then captures the following: resolve the underspecified element B so as to maximise discourse coherence:

- `Maximise Discourse Coherence:`
  If (a) $\mu(K_\beta)(\beta) \rightsquigarrow_* \mu(K_{\beta_1})(\beta_1)$; and
  (b) $\langle\tau, \alpha, \beta_1\rangle \wedge \mu(K_\tau)(\tau) \wedge \mu(K_{\beta_1})(\beta_1) \mathrel{|\!\!\approx} R_1(\alpha, \beta_1)$; and
  (c) R_1 is the $>_{\tau,\alpha}$ maximal rhetorical relation of attachment

  Then $Update(K_\tau, K_\alpha, K_\beta) \;:=\; Update(K_\tau, K_\alpha, K_{\beta_1})$.

It does this because in words, the rule ensures that if β_1 resolves B and produces the best coherence, then one must replace β with β_1 in the update. `Maximise Discourse Coherence` will be used in the analysis of (1) and (7) in section 7.2.

[9] Note that this won't affect the worst case complexity of DICE, and indeed from a practical perspective it may on occasion improve it because it will guide choices about which rhetorical relation to aim for first when computing the discourse update.

Note that these rules for computing bridging by reasoning about SDRT update are fully declarative and monotonic. They therefore don't make any assumptions about whether rhetorical relations are inferred first, or whether bridging relations are inferred first. However, such orders could be imposed in an implementation of this theory: for example, one could guide the implementation so that one attempts to compute rhetorical relations on the underspecified constituent *before* one computes a bridging relation; and failing that, one reasons about bridging relations, and then tries to compute rhetorical relations on the resolved constituents.

6 Modularity of Discourse Processing

Both our theory and Hobbs *et al.*'s theory use rhetorical relations to help compute briding inferences, and they are quite similar in spirit. However, there are several important differences. First, Hobbs ignored compositional semantic information and lexical semantics in computing the antecedents to definite descriptions, and he doesn't specify how to translate NL definite descriptions into logical form. We do.

The main difference, however, concerns modularity. For both linguistic and computational reasons, DICE exploits a logic that is distinct from the logic of information content (that is, the logic of SDRT). Indeed, the former logic is not only separate, but weaker than the latter logic. In contrast, in Hobbs *et al*'s abductive framework, the logic of the information content and the logic for computing rhetorical relations are one and the same. Hobbs *et al.* (1993) use weighted abduction to interpret discourse: one makes assumptions that explain the data at least cost, from a knowledge base that includes *all* information, both linguistic and non-linguistic.

Using abduction on semantic content and background knowledge to guide pragmatic inference is intuitively compelling. But there are two technical reasons for splitting the logics of information content and information cohesion in the way we do. First, all the nonmonotonic frameworks, including Hobbs *et al*'s abductive one, require some appeal to consistency tests to draw conclusions. But if one's base logic of information content is already that of first order logic, then adding consistency tests goes beyond the boundary of what is recursively enumerable. Our framework for computing rhetorical relations is also nonmonotonic. But the base logic is propositional rather than first order logic, because it is kept separate from the logic of information content of discourse (which is first order logic). So the logic for information cohesion we use here is decidable.

Second, by modeling compositional semantics, background knowledge and discourse coherence principles within a single logic as Hobbs *et al* do, one cannot separate the process of anaphora binding from the semantic content of the discourse as one would wish. Abduction requires some additive measure of cost on the various assumptions made to compute a proof of the discourse, and so inconsistent interpretations will always have the highest overall cost, and will be avoided if possible. Consequently, it's unclear how one should handle discourses where definite descriptions receive an *unambiguous* interpretation, which results in an inconsistency in the semantic content of the discourse (thereby making the discourse sound odd). For example, *the woman* and *the election* in (18b) unambiguously denote one of the people I met last night and the vote denoted in (18a) respectively, even though this results in an inconsistency that makes the discourse sound strange:

(18) a. I met two interesting people last night who voted for Clinton.

b. The woman abstained from voting in the election.

It's not clear that Hobbs *at al*'s abductive framework can account for examples like these, because the account will prefer accommodating the definite descriptions to binding it, in order to preserve consistency. In our account, binding definite descriptions to the discourse context is essential, because the compositional semantics of the definite article will demand it. In the above example, one would infer *Elaboration* between the constituents because of the relationship between the woman and the two people. The coherence constraints on this relation *won't* be violated by the fact that one can't abstain and vote at the same time. However, the discourse is still predicted to be odd in SDRT, because its representation is unsatisfiable.

Finally, Hobbs *et al.* assign different weights to different predicates, in order to deal with cases like (13), where there are choices about what bridging inferences to draw, because of the conflicting clues from different knowledge sources. A notion of cost for inferring information is very intuitive. But the the meaning of the weights in the abductive logic is unclear, and so there are no general principles that explain when and how (default) information about rhetorical relations overrides default world knowledge. In contrast, the logic we use is designed to resolve conflicting clues about semantic content from different knowledge resources logically, rather than through the use of weights (see Lascarides and Asher, 1993 for details). Reasoning among the knowledge resources will be handled 'automatically' by the logic (though we must take care in representing the axioms, so that the logic does this appropriately). So our approach is computationally more tractable while being more fine tuned to the linguistic phenomena.

Sperber and Wilson's approach to bridging also deserves some comment, though the comparision between the two approaches is more difficult here than in Hobbs *et al.*'s case. Relevance theorists could, though they have not done so, adopt our linguistic assumptions and most of our framework. Their view is compatible with our modular view of discourse interpretation, in a way that Hobbs's approach is not. Their claim would then be that it is the principle of relevance that guides the resolution of the underspecified elements in our treatment of definite descriptions. But then detailed comparison at this point would be highly speculative, given that we are not sure how to use the relevance principle in reasoning about underspecification.

7 Applications to Examples

We now examine some examples in detail. In sections 7.1 and 7.2, we will concentrate on bridging inferences involving definite descriptions. In section 7.3, we will briefly discuss cases that involve other expressions.

7.1 Bridging Through Discourse Attachment

First, consider a case where discourse structure determines bridging:

(12) a. John took engine E1 from Avon to Dansville.

b. He picked up the boxcar

c. and took it to Broxburn.

The DRSs representing (12a) and (12b) are α and β respectively:

$$(\alpha)\quad \begin{array}{|l|}\hline j, E1, a, d, e_1, t_1, n \\ \hline \mathit{John}(j) \\ \mathit{engine\text{-}E1}(E1) \\ \mathit{Avon}(a) \\ \mathit{Dansville}(d) \\ \mathit{take}(e_1, j, E1) \\ \mathit{from}(e_1, a) \\ \mathit{to}(e_1, d) \\ \mathit{hold}(e_1, t_1) \\ t_1 \prec n \\ \hline \end{array}$$

$$(\beta)\quad \begin{array}{|l|}\hline n, B, u, y, e_2, t_2, n \\ \hline \mathit{pick\text{-}up}(e_2, j, y) \\ \mathit{hold}(e_2, t_2) \\ t_2 \prec n \\ B =? \\ u =? \\ B(y, u) \\ \mathit{boxcar}(y) \\ \begin{array}{|l|}\hline z \\ \hline \mathit{boxcar}(z) \\ B(z, u) \\ \hline \end{array} \Rightarrow \begin{array}{|c|}\hline \\ \hline z = y \\ \hline \end{array} \\ \hline \end{array}$$

Note that *he* in β resolves to John. This is because anaphoric constraints in SDRT make John the only choice, regardless of the rhetorical relation which connects α and β.

In this example, resolving B to identity makes the update undefined, because there is no boxcar in α, and so no resolution of $u =?$. So according to `DS Determines Bridging`, we should check to see if we can attach β (as it stands) to α with a rhetorical relation, and if the results of this give us other values for u and B. The antecedent to `Narration` is verified, since both e_α and e_β are events. So by Defeasible Modus Ponens on `Narration`, $\mathit{Narration}(\alpha, \beta)$ is inferred.

Further inferences follow from this. First, by Modus Ponens and the `Temporal Consequent of Narration`, e_α occurs before e_β; that is, the taking of the engine from Avon to Dansville occurs before a boxcar is picked up. Furthermore, as we showed in section 4, by the semantics of the phrases *take to* and *pick up* and the `Spatial Consequence of Narration`, one infers that the source of the picking up event is in Dansville and the object that is picked up is therefore also in Dansville. Hence, the boxcar is in Dansville. Thus, the coherence constraints on *Narration* allows us to infer a particular way of resolving B and u—viz. B is *in* and u is d or Dansville (for simplicity, we have ignored conditions on when these relations hold, but they could be added to the formal representation of content). So `DS Determines Bridging` leads to the following revision of β, and this gets attached to α with *Narration*:

$$(\beta_1) \quad \begin{array}{|l|}\hline d, e_2, t_2, y, B, u, n \\ \hline pick\text{-}up(e_2, j, y) \\ hold(e_2, t_2) \\ t_2 \prec n \\ in(y, d) \\ boxcar(y) \\ Dansville(d) \\ source(e_2, d) \\ location(t_2, y, d) \\ \begin{array}{|l|}\hline z \\ \hline boxcar(z) \\ in(z, d) \\ \hline \end{array} \Rightarrow \begin{array}{|c|}\hline \\ \hline z = x \\ \hline \end{array} \\ \hline \end{array}$$

Note that our final result β_1 includes added content. We have resolved anaphoric conditions that were conventionally triggered by the definite. This added content was inferred in order to meet constraints on discourse coherence. It amounts to: the boxcar is located in Dansville and moreover, it's the only one in Dansville.

Poesio accounts for (12a,b), but fails to model cases involving different rhetorical relations:

(12) a. John took the engine E1 from Avon to Dansville.

b′. He also took the boxcar.

His theory doesn't predict the boxcar in (12a,b′) is in Avon. In contrast, our analysis captures the intuitive interpretation of (12a,b′). Briefly, as in the previous example, the attempt to specify the binding relation B to identity fails. The similarity in syntactic structure and the cue word *also* are clues in DICE that the discourse relation between (12a) and (12b) is *Parallel*. This doesn't have a spatial constraint like that represented in `Spatial Consequence of Narration`. Rather, the spatial constraints are computed on the basis of the way the different parts of the DRSs related in the parallel relation are mapped onto each other. This mapping is an essential feature of the coherence constraints on *Parallel* (Asher, 1993). For the sake of brevity, we omit the details of constructing this mapping here, but informally, the taking event in (12b′) is matched with that in (12a). The consequence is that, by the spatial constraints on *Parallel*, their sources and goals are taken to be the same, unless there's information to the contrary. This adds semantic content to the DRS representing (12b′); the source of the taking event in (12b′) is Avon. So by lexical semantics, the boxcar is in Avon at this source. One adds this to the representation of the given information via `DS Determines Bridging` as before. And so one obtains an interpretation where the boxcar is in Avon rather than Dansville, and it's the only boxcar in Avon.

7.2 Bridging Before Discourse Attachment

We have looked at cases where inferring a rhetorical relation helps specify bridging inferences. The rule `Maximise Discourse Coherence` specified in section 5.1 enables us to specify bridging inferences so as to gain discourse coherence that wouldn't be there otherwise.

In example (1), we fail to get a well-defined update if we specify the binding relation to identity. Furthermore, in contrast to texts like (12a,b), there isn't enough information in the *underspecified* constituent β representing (1b) to infer a particular rhetorical relation between it and α representing (1a).

(1) a. I met two interesting people last night at a party.

b. The woman was a member of Clinton's Cabinet.

This is because only `Background` in DICE applies, and so the only candidate relation is *Background*. But constituents related by *Background* must have a common topic. We can compute this using the technique discussed in Grover *et al.* (1994). That is, we generalise over the predicates and arguments in the propositions. Since we haven't resolved B and u, the woman is unconnected with the two people. And so computing a common topic in this way isn't possible, because the result is too general; something like *things that were true yesterday.*[10] Hence *Background* can't be inferred between α and the underspecified β. Neither can any other relation. Hence `DS Determines Bridging` won't apply.

Instead, we must use `Maximise Discourse Coherence`. That is, we must investigate which resolution of β produces the best discourse, and resolve β to that. Suppose that β_2 is a resolution of β where B and u are defined so that *the woman* y is *separate* from the two people mentioned in the first sentence. Then this produces just as bad a discourse as that between α and β itself, for the same reasons. On the other hand, suppose that β_1 is the resolution of β where the woman y in the DRS β is one of the two people I met last night. In other words, the binding relation B in β_1 resolves to *member-of*, and u resolves to the discourse referent denoting the two people I met in α. Then the rules in DICE given in Asher and Lascarides (1995) allow us to compute *Elaboration* between these constituents α and β_1. This comes with different coherence constraints from *Background*: the topic is α. The discourse coherence is therefore much improved. So, the antecedent to `Maximise Discourse Coherence` applies with respect to β_1, and so the discourse context α is updated via *Elaboration* with β_1. As before, we have gained further information: we now know that the woman is one of the two people I met last night, and only one of the people I met last night was a woman by the uniqueness condition that forms part of the compositional semantics of the definite. So the other one must have been a man.

Our analysis of (7a,b) also uses the principle `Maximise Discourse Coherence`.

(7) a. I just arrived.

b. The camel is outside and needs water

b′. The fleas are outside and need water.

[10] We don't formalise here the conditions under which a topic is poor. For such a formalisation, see Lascarides *et al.* (1996).

Again, B can't be identity. The antecedent to `Background` is verified, but notice the difference with the following variants (7a′,b″) and (7a′,b‴):

(7) a′. John arrived at 3pm.
b″. A camel was outside and needed water.
b‴. ?A camel is outside and needs water.

Background requires a distinct common topic, and one is readily able to construct this in (7a′,b″): a camel's being outside and needing water can be understood to be a property of the place John arrives at, a description perhaps of the scene that he sees. The operation of generalization then would yield a topic like: properties of the place that John arrives at. But this seems to be blocked in the case of (7a,b) and (7a′,b‴). We need an analysis of the effects of tense shift (from past to present) and words like *just* on discourse topics to model this. But exploring these effects would take us too far afield, and so we'll simply assume that *Background* is blocked in (7a,b) because a common topic can't be constructed. So we have to find another connection.

Just as in (1), we must entertain various resolutions of the underspecified parameters in β and see which option maximises discourse coherence. Suppose B and u are resolved so that the camel had some role in the arrival. By the constraint `Bridges are Plausible` given in section 5.1, this must be a plausible role. The only one is that the camel is the mode of transport by which I arrived. This content enables us to infer a new rhetorical relation, with improved discourse coherence. We can infer that the camel being outside was caused by my arrival thanks to the spatial information in the compositional semantics of the change of location phrase *arrive here*, and so the rhetorical relation is *Result*. So `Maximise Discourse Coherence` is used to infer this new content to the definite description *the camel*, together with the *Result* relation between the constituents.

(7a,b′) is odd because one cannot infer that the fleas are the mode of transport. This is implausible, and so it's ruled out by `Bridges are Plausible`. Indeed, there is no plausible resolution of B and u that produces a coherent discourse, and so the SDRS can't be updated. (7a′, b‴) is odd because the antecedent to `Maximise Discourse Coherence` isn't verified—the semantic representation of (7b‴) contains no underspecified elements. Therefore, even though (7b‴) as it stands cannot attach to (7a′), we lack the means to change its content. This demonstrates that although we capture bridging inferences for certain indefinites (e.g., (12a,b″)), we don't overgenerate bridging inferences for them, resulting in discourse coherence where there shouldn't be any.

Now consider the text (13):

(13) a. John moved from Brixton to St. John's Wood.
b. The rent was less expensive.

Let the sentences (13a,b) be represented by the DRSs α and β respectively. Once again attempting to resolve B to identity fails. But *rent* is a functional noun, and so in and of itself it suggests a value for B: it should be *of*, and the other term of the binding relation should be some object that can have rents. But there are no places that are mentioned in (13a) that have rents. So we must construct one through attempting to attach β to α.[11]

As in the previous examples, one cannot compute a rhetorical relation between α and the (underspecified) β. We need to know more about the connection between the rent mentioned in β and the content of α. There are at least two possible resolutions of u in β. The first, β_1, is such that the constituent means: *the rent of the place that John moved to, which is in St. John's Wood, is less expensive than the rent he paid in Brixton.* The second, β_2, is such that the constituent means: *the rent he paid in Brixton is less expensive than the rent of the place he moved to, which is St. John's Wood.* β_1 together with the content of α yield *Explanation*(α, β_1) in DICE. They also yield *Background*(α, β_1), because *Background* is compatible with *Explanation*, and β_1 describes a state (i.e., the rent in St. John's would being less expensive). Moreover, in contrast to α and β, we can compute a good topic for α and β_1, since we now know the rent is connected to St. John's Wood. In contrast, β_2 and the content of α yields only *Background*(α, β_2), but it cannot support *Explanation* (since moving to a more expensive house doesn't explain why one moved; at least, not on its own). Intuitively, one prefers an interpretation of a discourse that offers explanations of intentional behaviour that's described in the text—such as moving house—to an interpretation of the discourse where such behaviour is left unexplained. In essence, interpreters don't like miracles, or unexplained changes. We can model this via the partial order of rhetorical relations: *Explanation* $>_{\tau,\alpha}$ *Background* in this case. Therefore, the antecedent to the monotonic rule `Maximise Discourse Coherence` is verified and one updates β to β_1. In other words, one infers the rent referred to in (13b) is the rent that John pays in the place he moved to, which is in St. John's Wood.

This consequent of `Maximise Discourse Coherence` is incompatible with the default world knowledge that rents in Brixton are typically less expensive than those in St. John's Wood. However, since `Maximise Discourse Coherence` is a monotonic rule, it overrides this default world knowledge. This is as required, given the evidence in Matsui's experiments. In essence, `Maximise Discourse Coherence` guarantees that maintaining discourse coherence takes priority over default world knowledge; a principle of discourse interpretation for which we have argued elsewhere in modeling word sense disambiguation (Lascarides and Copestake 1997, Lascarides *et al* 1996).

7.3 Beyond Definite Descriptions

Bridging can occur in the absence of definites. We have already discussed how SDRT captures the bridging relation in (12a,b″):

(12) a. John took engine E1 from Avon to Dansville.
b″. He picked up a boxcar
c. and took it to Broxburn.

The bridging in (4), which we discussed in §1, in modelled in a similar manner:

[11] For the sake of simplicity, we ignore the comparative nature of *less*, and gloss over the way one computes from the discourse context the set over which the comparison (or rental cost) is measured.

(4) a. Jack was going to commit suicide.

b. He got a rope.

The proposition representing (4b) must be attached to the one representing (4a) with a rhetorical relation. Let's assume that the content of (4a) allows us to infer by default that Jack has a plan to commit suicide. Let us further suppose that if Jack has such a plan, and he gets a rope, and we know these events are connected somehow (as they must be for a rhetorical relation to hold), then normally, getting a rope is *part of* the plan, and the rope is the suicide instrument. These defaults will lead to an inference in DICE that the rhetorical relation is *Elaboration*. And the definition of SDRT Update will add the information that the rope is an instrument in the suicide to the representation of (4b), since this content is essential for the coherence of the *Elaboration*. So just as in (12a,b$''$), the coherence constraints on rhetorical relations trigger additions to the the semantic representation of (4), which amount to bridging inferences between the objects described in the text.

Bridging inferences also occur with presupposition triggers other than the definite, e.g., the *it*-cleft in (2):

(2) In the group there was one person missing. It was Mary who left.

Let us suppose that in line with Chierchia's analysis of definite descriptions, the compositional semantic analysis of *it*-clefts reflects the fact that they're anaphoric, demanding a relationship B between the event e corresponding to the content of the presupposed information (here, that someone left) and an antecedent event e' in the discourse context. Let us further suppose that by default someone leaving a group causes him to be missing from that group. Then this can be exploited to connect the two sentences in (2) with a rhetorical relation, and it also provides a way of resolving B via `DS Determines Bridging`. By the DICE axioms in Lascarides and Asher (1993), $cause(e, e')$ is inferred, where e' is the eventuality that someone's missing from the group, described in (2a). Moreover, this resolution of B to *cause* yields discourse coherence: the second sentence specifies who left, and so DICE supports the inference that this elaborates content of the first sentence.

Now consider the discourse (3):

(3) John partied all night yesterday. He's going to get drunk again today.

As with *it*-clefts, we assume *again* is anaphoric, in that its content includes the conditions $B(e, e')$, $B =?$ and $e' =?$, where e is the event that forms part of the presupposed content triggered by *again*; in this case, e is the event that John got drunk (before today). B and e' are resolved through discourse update. By generalising over the two properties of times given in the two DRSs that represent the two sentences, we can construct a common theme that supports a *Parallel* relation between them (for more details see Asher (1993)). To maximise the common theme, we infer that John got drunk at the party yesterday. And so computing the rhetorical structure of the discourse produces values for B and e' via `DS Determines Bridging`: B is *concurrent* and e' is the event described in the first sentence.

We have only hinted here at how our theory of bridging contributes to the analysis of cases involving other expressions. For the formal details of how our axioms introduced in section 5.1

are involved in the analysis of presupposition triggers in general, see Asher and Lascarides (1998).

8 Conclusion

Bridging inferences involve a complex interaction between lexical and compositional semantics, world knowledge and discourse structure. We have shown that the coherence constraints imposed by different rhetorical relations have an effect on bridging, which cannot be accounted for purely in terms of focus or domain knowledge.

We have modeled this effect in SDRT, a theory of discourse structure with the distinguishing feature that rhetorical connections can trigger a change to the semantic content of the propositions introduced in the text. Bridging inferences are a byproduct of computing how the current sentence connects to the previous ones in the discourse. Our account fully integrates compositional and lexical semantics and discourse structure. We use a well-defined logic which combines various knowledge sources to compute how new information integrates with the discourse context, paying particular attention to when these knowledge resources conflict. We demonstrated that by integrating compositional semantics and pragmatic reasoning in this way, we provide a more refined account of bridging inferences than either compositional semantic accounts or AI accounts that exploit background knowledge in discourse interpretation can achieve on their own.

Acknowledgements

Various versions of this paper have been presented at the International Workshop on Underspecification which was held at Berlin in 1996, the International Workshop on lexical semantics and acquisition which was held in Courmayeur, Italy in 1996, CUNY 1996, and seminars at the University of Texas at Austin and the University of Edinburgh. We would like to thank the people that attended these talks for their feedback. We would also like to thank David Beaver, Janet Hitzeman, Ali Knott, Rob van der Sandt, Frank Veltman and two anonymous reviewers for their helpful comments and suggestions on previous drafts of this paper.

9 References

Asher, N. [1993] *Reference To Abstract Objects in Discourse*, Kluwer Academic Publishers.

Asher, N. [1996] Mathematical Treatments of Discourse Contexts, in P. Dekker and M. Stokhof (eds.) *Proceedings of the Tenth Amsterdam Colloquium on Formal Semantics*, pp. 21–40, ILLC publications, University of Amsterdam.

Asher, N. [forthcoming] The Logical Foundations of Discourse Interpretation, in J. M. Larrazabal (ed.) *Logic Colloquium 1996*, Springer Verlag.

Asher, N., M. Aurnague, M. Bras, L. Vieu [1996] De l'espace-temps dans l'analyse du discours", *Semiotique : Numero Special "Theories sémantiques et modalisation*, **9**.

Asher, N. and A. Lascarides [1994] Intentions and Information in Discourse, in *Proceedings of the 32nd Annual Meeting of the Association of Computational Linguistics*, Las Cruces, USA, June 1994, pp34–41.

Asher, N. and A. Lascarides [1995] Lexical Disambiguation in a Discourse Context, *Journal of Semantics*, **12** number 1, pp69–108, Oxford University Press.

Asher, N. and A. Lascarides [in press] Questions in Dialogue, to appear in *Linguistics and Philosophy*.

Asher, N. and A. Lascarides [1998] The Semantics and Pragmatics of Presupposition, manuscript available from `http://www.cogsci.ed.ac.uk/ alex`.

Asher, N. and P. Sablayrolles [1995] A Typology and Discourse Semantics for Motion Verbs and Spatial PPs in French, *Journal of Semantics*, **12**.

Beaver, D. [1994], An Infinite Number of Monkeys, Technical Report, ILLC, Universiteit van Amsterdam.

Bos, J., A-M. Mineur and P. Buitelaar [1995] Bridging as coercive accommodation, technical report number 52, Department of Computational Linguistics, Universität Saarbruücken.

Briscoe, E. J., A. Copestake and B. Boguraev [1990] Enjoy the Paper: Lexical Semantics via Lexicology, *13th International Conference on Computational Linguistics (*COLING*-90)*, Helsinki, pp. 42–47.

Charniak, E. [1983] Passing Markers: A Theory of Contextual Influence in Language Comprehension, *Cognitive Science*, **7**, pp. 171–190.

Chierchia, G. [1995] *Dynamics of Meaning : anaphora, presupposition, and the theory of grammar*, University of Chicago Press.

Clark, H. [1977] Bridging, in P. N. Johnson-Laird, P. C. Wason (eds.) *Thinking: Readings in Cognitive Science*, pp. 411–420, Cambridge University Press.

Fernando, T. [1994] What is a DRS?, in R. Cooper and J. Groenendijk (eds.) *Integrating Semantic Theories II, Dyana Deliverable R2.1.B*, 1994, also appeared in the Proceedings of the First International Workshop on Computational Semantics, Tilburg 1994.

von Fintel, K. [1994] *Restrictions on Quantifier Domains*, PhD Thesis, University of Massachussetts, Amherst Mass.

Grosz, B. and C. Sidner [1986] Attention, Intentions and the Structure of Discourse. *Computational Linguistics*, **12**, 175–204.

Grover, C., C. Brew, M. Moens, S. Manandhar [1994] Priority union and generalisation in discourse grammar, *Proceedings of the 32nd Annual Meeting of the Association for Computational Linguistics (ACL94)*, Las Cruces, USA, pp 17–24.

Heim, I. [1992] Presupposition Projection and the Semantics of Attitude Verbs, *Journal of Semantics*, **9** (3).

Hobbs, J. [1979] Coherence and Coreference, *Cognitive Science*, **3**, pp.67–90.

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

Hobbs, J. R., M. E. Stickel, D. E. Appelt, P. Martin [1993] Interpretation as Abduction, *Artificial Intelligence*, **63**, pp69–142.

Kamp, H. and U. Reyle [1993] *From Discourse to Logic: Introduction to Modeltheoretic Semantics of Natural Language, Formal Logic and Discourse Representation Theory*, Kluwer Academic Publishers.

Karttunen, L. [1974] Presupposition and Linguistic Context, *Theoretical Linguistics* **1**, pp. 181–194.

Lascarides, A. and N. Asher [1993] Temporal Interpretation, Discourse Relations and Commonsense Entailment, in *Linguistics and Philosophy*, **16**, pp.437–493.

Lascarides, A. and A. Copestake [1997] The Pragmatics of Word Meaning, to appear in *Journal of Linguistics*.

Lascarides, A., A. Copestake and E. J. Briscoe [1996] Ambiguity and Coherence, *Journal of Semantics*, **13** number 1, pp41–65, Oxford University Press.

Lascarides, A. and J. Oberlander [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.

Lewis, D. [1979] Scorekeeping in a Language Game, *Journal of Philosophical Logic*, **8**, pp. 339–59.

Matsui, T. [1995] *Bridging and Relevance* University of London Ph.D. Thesis

Poesio, M. [1993] A Situation-Theoretic Formalization of Definite Description Interpretation in Plan Elaboration Dialogues, In P. Aczel, D. Israel, Y. Katagiri, and S. Peters, (eds.), *Situation Theory and its Applications*, vol. 3, 1993, p.339–374.

Poesio, M. [1994] Definite Descriptions, Focus Shift and a Theory of Discourse Interpretation, in *Proceedings of the Conference in Focus on Natural Language*, June 1994.

Poesio, M., R. Vieira and S. Teufel [1997] Resolving Bridging Descriptions in Unrestricted Text, in *Proceedings of the ACL97 Workshop on Operational Factors in Robust Anaphora Resolution*, also HCRC technical report number 87, HCRC, University of Edinburgh.

Polanyi, L. [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.

Pustejovsky, J. [1991] The Generative Lexicon, *Computational Linguistics*, **17**(4), pp. 409–441.

Pustejovsky, J. [1995] *The Generative Lexicon*, MIT Press, Cambridge, Mass.

Russell, B. [1905] On Denoting, *Mind*, **14**, pp. 479–493.

van der Sandt, R. [1992] Presupposition Projection as Anaphora Resolution, *Journal of Semantics*, **19**(4).

Sanfilippo, A. [1992] Grammatical Relations in Unification Categorial Grammar, *Lingua & Stile*, Fall issue.

Sperber, D. and Wilson, D. [1986] *Relevance*, Blackwells.

Stalnaker, R. [1978] Assertion, in Cole, Peter (ed.) *Syntax and Semantics volume 9: Pragmatics*, Academic Press.

Thompson, S. and W. Mann [1987] Rhetorical Structure Theory: A Framework for the Analysis of Texts. In *IPRA Papers in Pragmatics*, **1**, 79–105.

Verspoor, C. [1996] Lexical Limits on the Influence of Context, in G. W. Cottrell (ed.) *Proceedings of the 18th Annual Conference of the Cognitive Science Society* pp 116–120, Lawrence Erlbaum Associates.