

See discussions, stats, and author profiles for this publication at: <https://www.researchgate.net/publication/226373693>

Segmented Discourse Representation Theory: Dynamic Semantics With Discourse Structure

Chapter · October 2007

DOI: 10.1007/978-1-4020-5958-2_5

CITATIONS

63

READS

605

2 authors:



[Alex Lascarides](#)

The University of Edinburgh

118 PUBLICATIONS 4,427 CITATIONS

[SEE PROFILE](#)



[Nicholas Asher](#)

Institut de Recherche en Informatique de Toulouse

237 PUBLICATIONS 7,427 CITATIONS

[SEE PROFILE](#)

Some of the authors of this publication are also working on these related projects:



philosophy of science [View project](#)

Segmented Discourse Representation Theory: Dynamic Semantics with Discourse Structure*

Alex Lascarides
Division of Informatics,
University of Edinburgh
alex@inf.ed.ac.uk

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

Abstract

This paper motivates and describes a dynamic semantic theory of discourse interpretation called SDRT, which uses rhetorical relations to model the semantics/pragmatics interface. We describe the syntax and dynamic semantics of the language in which logical forms are represented, a separate but related language in which semantic underspecification is expressed as partial descriptions of logical forms, and a glue logic which uses commonsense reasoning to construct logical forms, relating the semantically underspecified forms that are generated by the grammar to their pragmatically preferred interpretations. We apply the framework to some examples involving anaphora and other kinds of semantic ambiguities.

1 Introduction

At least two important ideas emerged from research on discourse interpretation in the 1980s. First, dynamic semantics changed the way linguists think about meaning: instead of viewing the content of a discourse as the set of models that it satisfies (e.g., Montague (1974), Davidson (1980)), dynamic semantics views it as a *relation* between contexts known as the *context change potential* or CCP (e.g., Kamp (1981), Groenendijk and Stokhof (1991)). Secondly, AI-based research demonstrated that *discourse structure* is a necessary component of discourse interpretation (e.g., Hobbs (1985), Mann and Thompson (1987), Grosz and Sidner (1986)). Both these insights address the need to model how the interpretation of the current sentence is dependent on the interpretation of the sentences that precede it, but they differ in their aims and execution. Dynamic semantics typically explores a relatively restricted set of pragmatic phenomena, focusing on the effects of logical structure on anaphora of various kinds. For example, it predicts the difference in acceptability of the pronouns in (1a) vs. (1b):

- (1) a. A man walked in. He ordered a beer.
 b. Every man walked in. ??He ordered a beer.

*Much of the work described here is taken from Asher and Lascarides (2003).

Discourse structure in dynamic semantics is thus determined entirely by the presence of certain linguistic expressions such as *if*, *not*, *every* and *might*. The process of constructing logical form is equally simple, either using only syntax and the *form* of the logical forms of clauses but not their interpretations (e.g., Kamp and Reyle (1993)), or using in addition notions such as *consistency* and *informativeness* (e.g., van der Sandt (1992)).

In contrast, many AI approaches to discourse interpretation aim to model implicatures generally, including the interpretation of pronouns (e.g., Hobbs et al. (1993)). These theories emphasise the role of commonsense reasoning with non-linguistic information such as domain knowledge and cognitive states. For example, Hobbs (1985) argues that such reasoning is necessary for inferring that *he* in (2b) binds to Bill rather than John:

- (2) a. John can open Bill's safe.
- b. He's going to have to get the combination changed soon.

He argues persuasively that this interpretation occurs as a byproduct of working out how (and why) the discourse is *coherent*, where a discourse is defined to be coherent only if the contents of its utterances are rhetorically connected in a discourse structure. In this case, the rhetorical relation is *Result*, and this is inferred via commonsense reasoning with domain knowledge since no cue phrases such as *therefore* are present. Dynamic semantics predicts that John and Bill are possible antecedents to *he* in (2) but doesn't rank these alternatives. Discourse (2) is also a counterexample to theories of anaphora which utilise only grammatical information; e.g., Centering Theory (Grosz et al., 1995) predicts that *he* binds to John since antecedents in subject position are preferred to those in object position.

By eschewing insights from dynamic semantics on how logical form is constructed and interpreted, Hobbs *inter alia* (Hobbs, 1979, 1985, Hobbs et al., 1993) tend to exploit commonsense reasoning in cases where the simpler mechanisms from dynamic semantics would do. Indeed, it's not even clear that one can explain (1a) vs. (1b) by relying on commonsense reasoning and ignoring dynamic semantics, especially since *every man* in (1b) implicates that at least one man exists (who walked in). Further, Hobbs et al. (1993) assume a highly unmodular architecture: any piece of information from any knowledge source can be accessed at any time. But we believe this has drawbacks. First, an unmodular approach misses certain generalisations. For example, one cannot express within weighted abduction that the preferences for interpreting pronouns which are predicted by Centering Theory are overridden if the semantics of the rhetorical relation that's predicted by other information sources conflicts with them (see Stone and Thomason (2002) for motivation of such a rule). Indeed, information from grammar and from domain knowledge aren't distinguished at all, and weighted abduction is unable to express laws about how weights are assigned anyway.

Secondly, allowing the process for constructing logical form to have full access to their interpretations, as Hobbs *et al.* do, confuses *constructing* what is said with *evaluating* what is said. To see the difference consider (3), where (3b) plainly elaborates (3a):

- (3) a. There are some unsolvable problems in number theory.
- b. Every even number greater than two is expressible as the sum of two primes is undecidable, for instance.

Suppose that we were to infer this in a system of defeasible reasoning that has full access to the interpretations of the two clauses. This default reasoning demands a consistency test, and given that the semantics of *Elaboration* is such that it's true only if the propositions it connects are also true, testing the *Elaboration* connection for consistency will entail a test as to whether (3a) and (3b) are satisfiable. That is, we would need to test whether Goldbach's Conjecture is in fact undecidable or not, something which we have no idea how to do! But even the most mathematically inept interpreter can easily understand the discourse structure of (3) and construct its logical form; one has a clear picture of *what* is being said without necessarily being able to *evaluate* what is said. Unlike Hobbs *et al*'s framework, this distinction between constructing logical form and interpreting it is clearly marked in dynamic semantics.

We will describe here a theory of discourse interpretation that *integrates* dynamic semantics and AI-approaches, in an attempt to ameliorate the disadvantages of one framework with the advantages of the other. The theory is called Segmented Discourse Representation Theory or SDRT, and it is something that we have working on for over a decade.

SDRT provides both a logic for representing (and interpreting) the logical forms of discourse, and a logic for constructing logical forms. The former logic is known as the *logic of information content*; the latter is the *glue logic*. SDRT is wedded to dynamic semantics in that the logic of information content is assigned a dynamic semantic interpretation, using CCP in the familiar way. It extends prior work on discourse structure by assigning rhetorical relations a precise dynamic semantics, which explains how the content of the discourse augments the compositional semantics of its clauses. We'll see several examples of this in Sections 4 and 5.

The glue logic extends dynamic semantics' mechanisms for constructing logical form by encoding an interaction between semantics and pragmatics: it involves commonsense reasoning with both linguistic and non-linguistic information, which extends the partial information about content that's generated by the grammar to a more complete semantic representation of the discourse. For example, it will compute the value of a rhetorical relation (and its arguments) that was absent from compositional semantics; and/or it identifies the antecedent to a pronoun, resolves the semantic scope of a presupposition, disambiguates a word sense, yields a bridging inference etc. The glue logic therefore contributes an important additional element to current research in semantic underspecification: instead of relating an underspecified semantic representation to all its possible interpretations, as current formalisms do (e.g., Reyle (1993), Koller et al. (2000)), SDRT relates an underspecified semantic representation to its *pragmatically preferred* interpretations.

The glue logic is also distinct from commonsense reasoning as it's used in other work. For example, unlike Hobbs *et al*'s abductive approach, it works over *partial descriptions* of logical forms, so that constructing logical form proceeds in a constraint-based fashion. Secondly, for the reasons given earlier, SDRT has a highly modular architecture, which leads to a more constrained approach. Each knowledge source that contributes to discourse interpretation—compositional and lexical semantics, domain knowledge, cognitive states etc.—is represented in a distinct language with its own distinct logic. The glue logic has only *restricted* access to the information within these logics; for example, it has access only to *descriptions* of formulae in the logic of information content, but not to what those formulae entail (in the dynamic logic where logical forms for discourse are interpreted). This separation ensures we don't confuse constructing what is said with evaluating what is said. It also ensures that we can

represent the interaction between the information from different knowledge resources within the glue logic’s consequence relation.

While SDRT can be used to extend any dynamic semantic theory, for the sake concreteness we will use Discourse Representation Theory (DRT, Kamp and Reyle (1993)) as the starting point in this paper. We will briefly describe DRT and use simple texts to motivate the need for rhetorical relations (e.g., see (2)). Accordingly, we’ll extend the logic of information content in DRT to include rhetorical relations, to which we assign a compositional and dynamic semantic interpretation. We’ll then introduce SDRT’s mechanisms for constructing logical form. As we mentioned, this takes place at the *description* level, following current practice in composing logical forms for clauses within the grammar; e.g., Copestake et al. (2001), Asudeh and Crouch (2001). We’ll show how the resulting theory provides a unifying explanatory mechanism for a number of different discourse phenomena, overcoming some problematic predictions in both dynamic semantics and AI-based accounts of discourse.

2 Dynamic Semantics

Montague semantics (Montague, 1974) wasn’t designed to construct logical forms for multi-sentence discourse, but extending it in obvious ways falls short of handling the phenomena we want to analyse. Consider (1a) again. Appending the logical forms of the clauses to yield (1a’) incorrectly predicts that the man that walked can be different from the one that ordered the beer:

- (1) a’. $\exists x(man(x) \wedge walk-in(x)) \wedge \exists y(beer(y) \wedge order(z, y))$
 a''. $\exists x(man(x) \wedge walk-in(x) \wedge \exists y(beer(y) \wedge order(x, y)))$

The formula (1a'') is an improvement in this respect, but assigning (1a) this logical form would make constructing logical form overly complex, since the scope of quantifiers would extend beyond sentence boundaries. Moreover, since anaphoric binding must be blocked in (1b), one would need to block constructing a logical form for (1b) that’s similar to (1a''). In any event, this misses the point: (1a'') fails to represent the fact that uttering the first sentence *changes* the context in which the second sentence is interpreted.

Dynamic semantics redefines meaning to address these problems: a sentence S is interpreted as a *relation* between an input context and an output one. Assuming that the model M is fixed, these contexts consist of variable assignment functions. Roughly put, the input context is the set of variable assignment functions which make the content of the discourse prior to S true in M ; the output context is the subset of variable assignment functions from the input context which (also) make S true. Thus the set of functions always gets smaller, capturing the (simplifying) idea of monotonically accumulating information as discourse is interpreted.

Viewing meaning this way provides an elegant account of the anaphoric dependency in (1a). The input context for the first sentence is the set of all variable assignment functions; the output one consists of just those variable assignment functions which are defined for the individual or *discourse referent* x that’s introduced in the grammar by *a man* and that make $man(x)$ and $walk(x)$ true. Like all NPs, pronouns introduce a discourse referent but they

also introduce a condition that this discourse referent be identified with an *accessible* prior discourse referent of appropriate number and gender. As we'll see, *accessibility* is defined in terms of the *form* of logical form, but semantically, the accessible discourse referents amount to those for which each variable assignment function in the output context is defined. This captures the anaphoric binding in (1a), and the output contexts consist of variable assignment functions which are defined for x , and which satisfy the conditions that x is a man, walked in and ordered a beer. In (1b), there is no accessible discourse referent which can be identified with that introduced by the pronoun, making the discourse uninterpretable.

Now to formal details, focusing on Discourse Representation Theory (or DRT, Kamp and Reyle (1993)) although our evaluation of it in Section 3 applies to dynamic semantics generally. The logical forms of discourse in DRT are discourse representation structures or DRSS. A DRS is a pair: a set of discourse referents, and a set of DRS-conditions. Their syntax (for a very simple fragment) is defined as follows:

Definition 1 The Syntax of DRSS


The set of DRSS is defined by:

$$K := \langle U, \emptyset \rangle \mid K \oplus \langle \emptyset, \gamma \rangle$$

Where:

1. U is a set of discourse referents;
2. γ is a DRS-condition; i.e., if x_1, \dots, x_n are discourse referents and R is an n -place predicate, then $\gamma := R(x_1, \dots, x_n) \mid \neg K \mid K_1 \Rightarrow K_2$; and
3. \oplus is an 'append' operation on DRSS: that is, if K_1 is the DRS $\langle U_1, C_1 \rangle$ and K_2 is the DRS $\langle U_2, C_2 \rangle$, then $K_1 \oplus K_2 = \langle U_1 \cup U_2, C_1 \cup C_2 \rangle$.

DRSS are sometimes written in a 'box-style' notation, as shown in (4a) (for *a man walks*) and (4b) (for *every man walks*):



Thus one DRS can *subordinate* another, and this is used to define the accessibility constraint on anaphora mentioned earlier:

Definition 2 Subordination

A DRS K_1 is immediately subordinate to K_2 iff:

1. K_2 contains the DRS-condition $\neg K_1$; or
2. K_2 contains the DRS condition $K_1 \Rightarrow K_3$ or $K_3 \Rightarrow K_1$ for some DRS K_3 .

Transitive Closure: A DRS K_1 is subordinate to K_2 iff there is a DRS K_3 such that K_1 is immediately subordinate to K_3 and K_3 is subordinate to K_2 .

Definition 3 Accessibility

A discourse referent x is *accessible* to an anaphoric DRS condition in K_1 (e.g., the condition introduced by a pronoun) iff x is introduced in U_{K_2} where:

1. K_1 is subordinate to K_2 ; or
2. $K_2 \Rightarrow K_3$ is a DRS condition in a DRS K_4 , such that K_1 is subordinate to K_3 .

Thus the discourse referent y that's introduced by a pronoun in a DRS K_1 can be bound to any discourse referent (of appropriate number and gender) that is introduced in a DRS on the following path: starting at K_1 , if there is a DRS K_2 immediately to your left (i.e., $K_2 \Rightarrow K_1$ is DRS-condition) then move to that; if not, but there is a DRS K_2 that you're immediately subordinate to, move to that; otherwise stop.

Let's focus first on how logical form is *constructed*, rather than how it's *interpreted*. Logical-form construction is encapsulated in the process of *discourse update*. In its simplest form, it consists of the following steps:

1. Construct the logical form of the current sentence, leaving the anaphoric conditions unresolved. We'll examine how unresolved conditions are represented shortly; for now we gloss the condition introduced by a pronoun as $x = ?$ (the “?” showing that the antecedent to x is unknown). Constructing such DRSS can be done compositionally within the grammar (Muskens, 1996, Asher, 1993).
2. Use \oplus to append this logical form to the DRS of the discourse context; and
3. Resolve any conditions of the form $x = ?$ to conditions of the form $x = y$, where y is accessible to $x = ?$.

Observe that this construction procedure uses only the *form* of the DRSS—i.e., subordination—and not their interpretation. Nevertheless, it makes the right predictions about (1a) and (1b), as shown in Figures 1 and 2 (ignoring tense for now).

The dynamic interpretation of DRSS makes semantic sense of the accessibility constraint on anaphora. The *introduction* of new discourse referents into a DRS K causes a *transition* from an input context to an output one, while DRS-conditions impose *tests* on the input context (observe the conditions $f = g$ in clauses 3–6):

Definition 4 The Interpretation of DRSS

Assuming a first order model M consisting of a set of individuals D_M and an interpretation function I_M :

1. $f[\langle U, \emptyset \rangle]_M g$ iff $dom(g) = dom(f) \cup U$.
2. $f[K \oplus \langle \emptyset, \gamma \rangle]_M g$ iff $f[K] \circ [\gamma]_M g$
3. $f[R(x_1, \dots, x_n)]_M g$ iff $f = g$ and $\langle f(x_1), \dots, f(x_n) \rangle \in I_M(R)$.
4. $f[\neg K]_M g$ iff $f = g$ and there is no function h such that $f[K]_M h$.

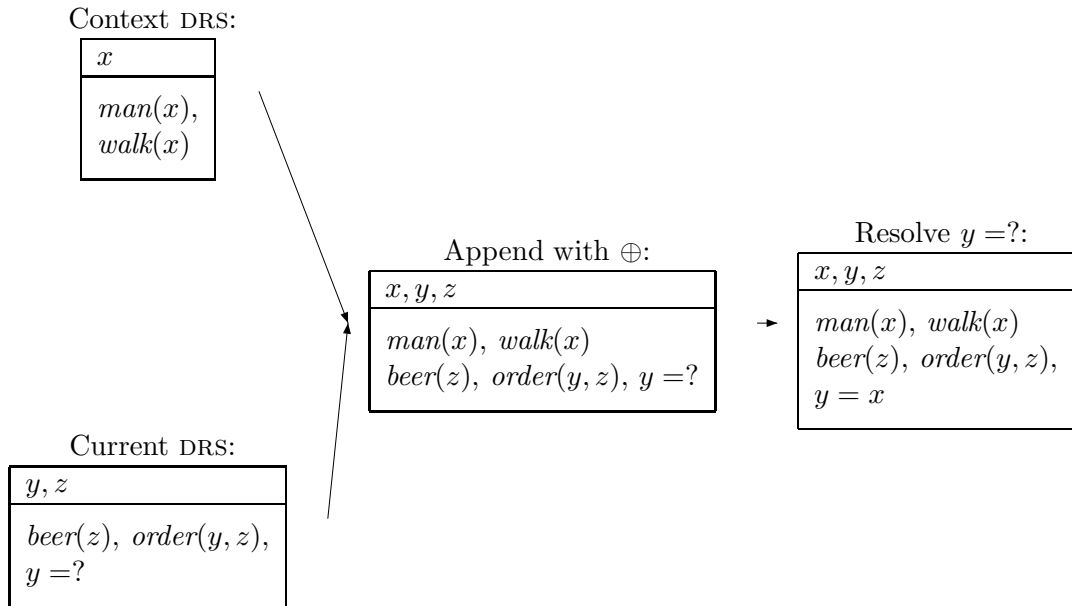


Figure 1: Constructing the DRS for (1a)

5. $f \llbracket K_1 \Rightarrow K_2 \rrbracket g$ iff $f = g$ and for every function h such that $f \llbracket K_1 \rrbracket_M h$ there is a function k such that $h \llbracket K_2 \rrbracket_M k$.

This looks like a small change to the Tarskian notion of satisfaction: instead of defining semantics in terms of one variable assignment function, we use two functions. And indeed, there is a close correspondence between first-order logic and basic fragments of DRT (Fernando, 1994). However, the change is more dramatic than it seems: $\neg\neg K$ is not dynamically equivalent to K —indeed, the former simply imposes a test on the input context while the latter transforms it. This semantic difference surfaces in the treatment of anaphora: a discourse referent that’s introduced inside a double negation isn’t an accessible antecedent to subsequent anaphora while a discourse referent that’s not inside a double negation is.

3 Why Dynamic Semantics needs Rhetorical Structure

Dynamic semantics is used to account for various anaphora: pronouns (e.g., Kamp and Reyle (1993)), tense (e.g, Kamp and Rohrer (1983)) and presupposition (e.g., van der Sandt (1992)), among others. We review this work here to motivate the introduction of rhetorical relations.

3.1 Pronouns

Consider text (5) from Lascarides and Asher (1993):




Figure 2: Constructing the DRS for (1b)

- (5)
- π_1 . John had a great evening last night.
 - π_2 . He had a great meal.
 - π_3 . He ate salmon.
 - π_4 . He devoured lots of cheese.
 - π_5 . He won a dancing competition.
 - π_6 . ??It was a beautiful pink.

This discourse contains no expressions such as *every* or *not* that block discourse referents from being antecedents. DRT therefore over-generates the possible interpretations of *it* in π_6 , allowing it to bind to the salmon in π_3 . Rhetorical relations can help overcome this problem, however. They allow one to reflect the capacity of a discourse to describe things at different levels of detail: for example, one can introduce a relation *Elaboration*(π_1, π_2) whose semantics entails that the events described in π_2 describe in more detail those described in π_1 ; in contrast, *Narration*(π_3, π_4) reflects temporal progression between the events, rather than a change in the granularity of description. These relations therefore provide a way of thinking about the content of (5) that's shown in Figure 3.

This figure follows Hobbs (1985) and Asher (1993) in assuming that *Elaboration* induces subordination (to reflect its semantic function of changing granularity of description) whereas *Narration* induces coordination. The resulting structure affects anaphora. Most research on discourse structure assumes what's known as a *right-frontier* constraint (e.g., Grosz and Sidner (1986), Polanyi (1985), Webber (1991) and others): anaphora in the current clause




Figure 3: The discourse structure of (5)


must be bound to an antecedent which is on the right frontier of the structure. This blocks *it* in π_6 from binding to the salmon in π_3 , since π_3 isn't on the right frontier.

DRT doesn't introduce discourse referents which denote abstract objects such as propositions, and it therefore under-generates the possible interpretations of *this* in (6):

- (6)
- π_1 . One plaintiff was passed over for promotion three times.
 - π_2 . Another didn't get a raise for five years.
 - π_3 . A third plaintiff was given a lower wage compared to males who were doing the same work.
 - π_4 . But the jury didn't believe this.

However, simply extending DRT to include such referents would replace the under-generation problem with an over-generating one. Since there are no linguistic expressions such as *every*, *not* and *if* that block discourse referents from being antecedents to anaphora, DRT's accessibility constraint would incorrectly predict that *this* can refer to the second claim alone. But in fact, *this* can only refer to the last claim or to the sum of the claims (differences in intonation would facilitate these differences in interpretation).

Rhetorical relations and the right-frontier constraint help here too: π_2 forms a *Continuation* with π_1 , the continuation segment elaborating some linguistically implicit topic (such as *three plaintiffs made three claims that they are ill-treated*), and π_3 continues this continuation as shown in (6').



Thus according to the right-frontier constraint π_4 can either be rhetorically connected to the topic (in which case *this* resolves to the three claims) or to π_3 (in which case *this* resolves to the third claim). This right-frontier constraint also explains why inserting a sentence expressing the topic between π_3 and π_4 changes the interpretation of *this*: now the only proposition on the right frontier is the topic, and so *this* must bind to the three claims.

3.2 Temporal Anaphora

Intuitively, (7) is true only if the events all occur in the past and in a definite sequence: i.e., the sequence in which they are mentioned.

(7) Max fell. John helped him up.

Kamp and Reyle (1993) and others use DRT’s mechanisms for anaphoric binding to account for this. Thus only syntax and the form of the DRSS affect how temporal anaphora are interpreted. As Kamp and Reyle (1993) themselves observe, however, factors other than syntax and form influence temporal interpretation. Lascarides and Asher (1993) use (8) as a counterexample to DRT’s analysis:

(8) Max fell. John pushed him.

The natural interpretation of (8) is one where the temporal order of the events *mismatches* their textual order, and the rules for constructing logical form in Kamp and Reyle (1993) yield a DRS with the wrong truth conditions. Default world knowledge—about pushing causing fallings, for example—might seem a plausible basis for distinguishing (7) from (8). But in general it’s not sufficient. On its own, default world knowledge would predict the wrong truth conditions in (9) (from Lascarides et al. (1992)):

(9) π_1 . John gave Max a blow to the back of his neck.
 π_2 . Max fell.
 π_3 . John pushed him.
 π_4 . Max rolled over the edge of the cliff.

Reasoning about discourse structure—and in particular, the right frontier constraint mentioned earlier—can explain these textual ordering effects: the proposition π_3 cannot be interpreted as a cause of π_2 , because that would require it to be rhetorically connected to the other cause π_1 , and the right-frontier constraint blocks this.

3.3 Presuppositions

Many recent accounts of presupposition have exploited the dynamics in dynamic semantics (Beaver, 1996, Geurts, 1996, Heim, 1982, van der Sandt, 1992). Presuppositions impose tests on the input contexts: either the context must satisfy the presuppositions of the clause (e.g., Beaver (1996), Heim (1982)), or the presuppositions are anaphoric (e.g., van der Sandt (1992))

and so must be bound to elements in the context. When presuppositions fail this test, as they do in (10) (where *Jack's son* generates the presupposition that Jack has a son), the presupposition is accommodated or *added* to the context, provided various constraints are met (e.g., the result must be satisfiable).

(10) If baldness is hereditary, then Jack's son is bald.

(11) If Jack has a son, then Jack's son is bald.

One of the success stories of dynamic semantics has been to use the *structure* of the contexts to place constraints on accommodation. Van der Sandt (1992) stipulates that presuppositions are accommodated into the part of the context that gives them the widest scope possible, subject to the result being consistent and informative (which means that no part of the asserted content becomes logically redundant once the presupposition is added). In contrast, Beaver (1996) stipulates that they are accommodated in the accessible site that produces the most plausible pragmatic interpretation, though he doesn't formalise this constraint. Nevertheless, both these theories predict that the presupposition projects from the embedding inside the conditional in (10). Since the context in (11) satisfies the test on the input context (both the satisfaction test and the anaphoric binding one), the presupposition isn't added to the context and so it doesn't project out from the embedding inside the conditional.

This doesn't explain some of the data, however. Pragmatics and rhetorical structure also affect presuppositions. To illustrate this, consider texts (12a) and (12b):

(12) a. If David scuba dives, he'll bring his regulator.

b. If David scuba dives, he'll bring his dog.

In both cases, the context fails the tests imposed by the presupposed content that's generated by the possessive NP in the consequent (that David has a regulator and that David has a dog respectively). So this content has to be accommodated. According to van der Sandt's constraints on accommodation, the presupposed content projects out from the conditional in both cases. But although this is correct for (12b), it's clearly wrong for (12a), which is interpreted as: *If David scuba dives then he has a regulator and he'll bring it*. In other words, the constraints on accommodation are too weak, yielding wide scope readings in cases where it should be narrow scope.

Beaver's (1996) plausibility constraint seems to do a better job. According to domain knowledge, there is no logical dependency between scuba diving and owning a dog. But one implies there is if one assigns the presupposed content narrow scope relative to the conditional, for then (12b) would mean: *If David scuba dives, then he has a dog and he'll bring it*. In contrast, domain knowledge suggests there *is* a dependency between being a scuba diver and owning a regulator (i.e., you're much more likely to own a regulator if you scuba dive than if you don't). And the narrow scope reading of the presupposition for (12a) reflects this dependency.

However, Beaver doesn't formalise this story, and further inspection reveals that it's not so simple to do so. In particular, measuring the plausibility of the content *as a whole* that results from some particular interpretation of a presupposition can't be right: this would always make the narrow scope reading win over the wide scope one, because the former are entailed by

the latter and are therefore necessarily more plausible/probable to be true. In particular, measuring plausibility this way (rather than in the way we described above) would predict the wrong reading of (12b): one where the presupposition has narrow scope because this reading doesn't require *John owns a dog* to be true while the wide-scope reading does require this (and is therefore necessarily less plausible). But if we don't measure the plausibility of the whole content, then what are we measuring the plausibility of?

One might try to fix this by only blocking interpretations which entail something highly implausible. Since owning a dog is not highly implausible, the wide scope reading of (12b) would not (necessarily) be ruled out. This contrasts with the plausibility of owning a regulator, providing the basis for preferring the narrow scope reading of (12a). But avoiding interpretations that are highly unlikely to be true isn't right either, because this strategy predicts the wrong interpretation of (13):

- (13) I doubt that the knowledge that this seminal logic paper was written by a computer program running on a PC will confound the editors.

The factive noun *knowledge* generates the presupposition that this seminal logic paper was written by a computer program running on a PC. Given world knowledge about the current state of NLP technology, this is very unlikely to be true! But interpreting the presupposition so that the result reflects default world knowledge results in the wrong prediction—one where the presupposition takes narrow scope and is embedded within the referentially opaque context generated by *doubt*. Thus unless one suitably constrains what one measures the plausibility of, Beaver's constraints on accommodation are in danger of predicting a narrow scope reading where intuitively we get a wide scope one.

Rhetorical relations can offer the means to constrain accommodation appropriately. In essence, they determine what we should be measuring the plausibility of, for they allow us to cache out plausibility in terms of the 'naturalness' or overall quality of the different rhetorical links in different candidate interpretations. Let's assume that, just as asserted content is coherent only if it's rhetorically connected to something in the context, presupposed content is also coherent only if it's rhetorically connected to the context. So the semantic scope of a presupposition depends on which part of the context the presupposition binds to with a rhetorical relation. Now, we can assume that on the one hand, presuppositions have a default tendency to project from embeddings and rhetorically connect to propositions which have widest scope. But on the other hand, to capture the effects of plausibility on interpretation, let's assume a *monotonic* principle that we prefer discourse interpretations that maximise the naturalness of the rhetorical links between its constituents. This in fact follows from a more general principle of interpretation that we motivate in Asher and Lascarides (2003): the principle of *Maximising Discourse Coherence* (or MDC). This principle rests on the observation that coherence is not a yes/no matter, but it can vary in quality. And MDC states that one (monotonically) prefers discourse interpretations that are consistent with the compositional semantics of the clauses and maximise coherence.

But how does one measure the degree of coherence of a discourse interpretation? SDRT takes a very conservative view on how interpretations are ranked, which is as roughly follows (for a full definition of MDC, see Asher and Lascarides (2003)):

Definition 5 Maximise Discourse Coherence (or MDC)

Discourse is interpreted so as to maximise discourse coherence, where the ranking among interpretations are encapsulated in the following principles:

1. All else being equal, the more rhetorical connections there are between two items in a discourse, the more coherent the interpretation.
2. All else being equal, the more anaphoric expressions whose antecedents are resolved, the higher the quality of coherence of the interpretation.
3. Some rhetorical relations are inherently scalar. For example, the quality of a *Narration* is dependent on the specificity of the common topic that summarises what went on in the story; the quality of a *Contrast* is dependent on the extent to which the semantics of the connected propositions are dissimilar (to see this, consider *John loves to collect classic cars. But his favourite car is a 1999 Ford Mondeo*, which is a ‘better’ contrast than *John loves to collect classic cars. hates football*). All else being equal, an interpretation which maximises the quality of its rhetorical relations is more coherent than one that doesn’t.

Now consider (12a) again. It seems intuitively plausible that the quality or ‘naturalness’ of the rhetorical relation *Consequence* in (12a)—as triggered by the word *if*—is improved if it connects a proposition that John scuba dives to one that includes the content that he has a regulator compared to the alternative where no such logical dependency between scuba diving and owning a regulator is recorded. So by clause 3. of MDC we predict that the presupposition has narrow scope relative to the conditional, the default that presuppositions take wide scope being overridden because MDC is monotonic. In (12b), on the other hand, world knowledge doesn’t support a logical dependency between scuba diving and owning a dog, and so the naturalness of the *Consequence* link wouldn’t be enhanced by assigning the presupposition narrow scope, and so the default for it to take wide-scope wins.

Overall, the strategy we advocate here, contra Hobbs et al. (1993), is to separate the task of constructing the interpretation of an utterance from the likelihood that this interpretation is true. Instead, we aim to construct the interpretation of a discourse by reasoning about the demands that are imposed on it by discourse coherence. This strategy allows us to apply preferences based on likelihood more selectively.

3.4 Some other Phenomena

Interpreting definite descriptions typically involves computing a *bridging relation* to some antecedent, and rhetorical relations affect this. Observe that contrary to world knowledge, the waitress in (14c) is the waitress in the hotel rather than the waitress in the Thai restaurant:

- (14)
- a. We had dinner at a Thai Restaurant
 - b. and then drinks at a fancy hotel on 5th Avenue.
 - c. The waitress was from Bangkok.

The right-frontier constraint mentioned earlier would predict this, since (14a) is connected to (14b) with the coordinating relation *Narration* and so only (14b) is on the right frontier when (14c) is interpreted. Rhetorical relations also explain the minimal pair (15ab) vs. (15ab’):

- (15) a. John moved engine E1 from Avon to Dansville.
- b. He picked up the boxcar and took it to Broxburn.
- b’. He also took the boxcar.

The boxcar in (15ab) is interpreted as a boxcar in Dansville, but changing the rhetorical relation (from *Narration* to *Parallel*) through the introduction of *also* in (15ab’) changes that interpretation: now the boxcar is in Avon when it’s picked up.

Word sense disambiguation is similarly effected by discourse structure:

- (16) a. John bought an apartment.
- b. But he rented it.
- (17) a. The judge asked where the defendant was.
- b. The barrister said he was in the pub drinking.
- c. The bailiff found him slumped beneath the bar.
- c’. But the bailiff found him slumped beneath the bar.

The intuitive interpretation of (16ab) is one where *rent* is interpreted as *rent out* (i.e., John is the landlord). This is predicted by MDC, because this interpretation supports not only a *Contrast* relation between the constituents (as required by *but*) but also *Narration* (which entails that the buying happened before the renting). The rent-from sense of *rent* doesn’t support *Narration* and moreover the quality of the *Contrast* is worse (though coherent, because one isn’t usually a tenant in an apartment before one buys it).

In (17a–c), the noun *bar* is ambiguous; it has (at least) a pub/place-to-drink sense and a courtroom sense. But it is not ambiguous in this discourse, where the preferred interpretation of *bar* is its pub sense. Arguably, the most detailed models of word sense disambiguation are stochastic models which are trained on corpora and other on-line resources (e.g., Guthrie et al. (1991), Dagan et al. (1997)). But these models cannot fully explain the data in (17): if it predicts the right sense disambiguation of *bar* in (17c), then it will get the wrong results when (17c) is replaced with (17c’) (and where *bar* now means the courtroom bar), for the word *but* is statistically speaking independent of the word *bar*.

Using rhetorical relations as a clue to word sense disambiguation yields a different story. Roughly, (17abc) is a narrative: the proposition expressed by each sentence is related by the discourse relation *Narration* to the proposition expressed by the previous sentence. *Narration* imposes strong spatio-temporal constraints on the actors and events involved (see Section 4.1): the narrative links are better if the locations of objects in the interpretation of the current clause comply with expectations from the context vs. when they don’t so comply (e.g., if *bar* in (17c) is replaced with *courtroom bar*). So MDC predicts that we interpret (17c) so that the defendant is found in the pub and not in court, making *bar* disambiguate to its pub sense. But

(17abc') has a different interpretation, because the rhetorical role of (17c') is different. This is related with *Contrast* (plus perhaps other discourse relations too) to (17b) thanks to the cue phrase *but*,¹ and this has different spatio-temporal effects to the cases where the clauses are related only with *Narration* (e.g., (17abc)). Roughly, because of MDC, the spatial-trajectory of the objects must be such as to maximise the 'semantic differences' of the propositions connected (e.g., the expectations that arise from the content of one proposition should be asserted as false in the other). In this case, this means that interpreting *bar* as the courtroom bar is preferred. For then the expectation arising from (17b)—that the defendant is not in the courtroom—is asserted as false in (17c'). Thus interpreting *bar* as the courtroom bar in (17abc') is predicted by the principle MDC.

Using rhetorical relations to analyse (17) is complementary to using world knowledge. The presence of *but* is crucial to the distinct interpretations of (17abc) and (17abc'), and the fact that *but* favours interpretations where expectations in the context get violated is a matter of linguistic convention rather than world knowledge.

4 The Logic of Information Content

Having motivated the need for rhetorical relations, we will now extend the language of DRSS accordingly. We introduce two new expressions: *speech act discourse referents* label content (either of a clause or of text segments), and keep track of the token utterances in the discourse; and *rhetorical relations* relate speech act discourse referents. The resulting structures are known as segmented DRSS or SDRS. We start with a definition of SDRS-formulae, which express the content tagged by speech act discourse referents (or "labels") in SDRS:

Definition 6 SDRS-Formulae

The well-formed SDRS-formulae are constructed from the following vocabulary:

1. The vocabulary of DRSS.
2. Labels or *speech act discourse referents* π, π_1, π_2, \dots
3. A set of relation symbols for discourse relations: e.g., *Explanation*, *Contrast*, *Narration* etc.

The set Φ of well-formed SDRS-formulae consists of:

1. The set of DRSS
2. If R is a (2-place) rhetorical relation symbol and π_1 and π_2 are labels, then $R(\pi_1, \pi_2)$ is an SDRS-formula.
3. If ϕ and ϕ' are SDRS-formulae, then so are: $(\phi \wedge \phi')$ and $\neg\phi$ (where \wedge and \neg are interpreted dynamically, as in Definition 4).

¹Asher and Lascarides (2003) argue that explicit cues, such as the presence of a cue phrase *but*, must be present when the *Contrast* relation conveys a denial of expectation. This is why the clauses (17bc) cannot be interpreted so that they are connected with *Contrast*.

Definition 7 SDRS or Discourse Structure

An SDRS or a discourse structure is a triple $\langle A, \mathcal{F}, LAST \rangle$, where:

- A is a set of speech act discourse referents;
- $LAST$ is a member of A (intuitively, this is the label of the content of the last clause that was added to the logical form); and
- \mathcal{F} is a function which assigns each member of A an SDRS-formula.

In addition, the following **constraint** is imposed on A : let $Succ(\pi, \pi')$ hold just in case $\mathcal{F}(\pi)$ contains the literal $R(\pi', \pi'')$ or $R(\pi'', \pi')$. Then the transitive closure of $Succ$ (which we also call *outscofes*) from a partial order over A and there is a unique $Succ$ -supremum $\pi_0 \in A$.

When there is no confusion, we may write $\langle A, \mathcal{F} \rangle$ instead of $\langle A, \mathcal{F}, LAST \rangle$.

Note how \mathcal{F} assigns to labels SDRS-formulae that contain labels (indeed, this yields the partial order *outscofes* on A). This captures the intuition that the contents of clauses ‘group together’ to form coherent text segments. Having the unique supremum π_0 corresponds to assuming that the content of the discourse *overall* receives a single label.

Let’s illustrate the definition with a couple of examples. (8’) is the SDRS for the two-sentence discourse (8):

(8) Max fell. John pushed him.

(8’) $\langle A, \mathcal{F}, LAST \rangle$, where:

- $A = \{\pi_0, \pi_1, \pi_2\}$

$$\mathcal{F}(\pi_1) = \begin{array}{|l} x, e_{\pi_1} \\ \hline max(x), \\ fall(e_{\pi_1}, x), e_{\pi_1} \prec n \end{array}$$

$$\mathcal{F}(\pi_0) = Explanation(\pi_1, \pi_2)$$

- $LAST = \pi_5$

$$\mathcal{F}(\pi_2) = \begin{array}{|l} y, e_{\pi_2} \\ \hline john(x), push(e_{\pi_2}, y, x) \\ e_{\pi_2} \prec n \end{array}$$

The temporal relation between e_{π_1} and e_{π_2} is not explicitly encoded, but instead follows from the semantics of $Explanation(\pi_1, \pi_2)$, as we’ll see in Section 4.1. The more complex SDRS (5’) is the logical form for the first five sentences of (5), so long as $K_{\pi_1} - K_{\pi_5}$ are the DRSS representing the contents of those sentences respectively:

(5’) $\langle A, \mathcal{F}, LAST \rangle$, where:

- $A = \{\pi_0, \pi_1, \pi_2, \pi_3, \pi_4, \pi_5, \pi_6, \pi_7\}$
- $\mathcal{F}(\pi_1) = K_{\pi_1}, \mathcal{F}(\pi_2) = K_{\pi_2}, \mathcal{F}(\pi_3) = K_{\pi_3}, \mathcal{F}(\pi_4) = K_{\pi_4}, \mathcal{F}(\pi_5) = K_{\pi_5}$
 $\mathcal{F}(\pi_0) = Elaboration(\pi_1, \pi_6)$
 $\mathcal{F}(\pi_6) = Narration(\pi_2, \pi_5) \wedge Elaboration(\pi_2, \pi_7)$
 $\mathcal{F}(\pi_7) = Narration(\pi_3, \pi_4)$




Figure 4: The SDRS (5') in DRT-style notation

- $LAST = \pi_5$

In words, the overall content of the text, which is labelled π_0 , consists of π_1 (Max having a lovely evening) being elaborated by the narrative (and hence complex proposition) π_6 , which consists of the content of π_2 (having a lovely meal) and π_5 (winning a dancing competition), where π_2 is elaborated by the content of the narrative of π_3 (eating salmon) and π_4 (devouring cheese). We'll show how to construct this SDRS for (5) in Section 5.

The SDRS (5') makes use of a **convention** that for any label π , we call the formula $\mathcal{F}(\pi)$ that it labels K_π . We will make use of this convention from now on. We also adopt other conventions for writing SDRSs. For example, one can use the ‘box-style’ notation familiar from DRT and our earlier work (e.g., Asher (1993), Asher and Lascarides (1998)). First, we convey $\mathcal{F}(\pi) = \phi$ by writing $\pi : \phi$. We then convey the “immediate outscores” relation *Succ* as follows: If $Succ(\pi, \pi')$, then π' appears in the top strip of K_π and $\pi' : K_{\pi'}$ in the main part of K_π . Thus Figure 4 is just another way of coding up the SDRS (5'). Another way of coding up SDRSs explicitly shows which rhetorical relations are subordinating and which are coordinating, as shown in Figure 5. SDRT's constraints on anaphora (see Definition 8) ensures that antecedents must be DRS-accessible on the right frontier of this structure (unless *Contrast* or *Parallel* are present, where this principle breaks down).

The definition of SDRSs allows two labels π_1 and π_2 to be related by more than one rhetorical relation. This plurality allows an utterance to make more than one illocutionary contribution to the discourse. For example, π_1 and π_2 in (16) are related by both *Contrast* (as indicated by *but*) and *Narration* (because as we will see shortly, this ensures the right temporal effects, that the buying precedes the renting):

- (16) π_1 . John bought an apartment
 π_2 . but he rented it.

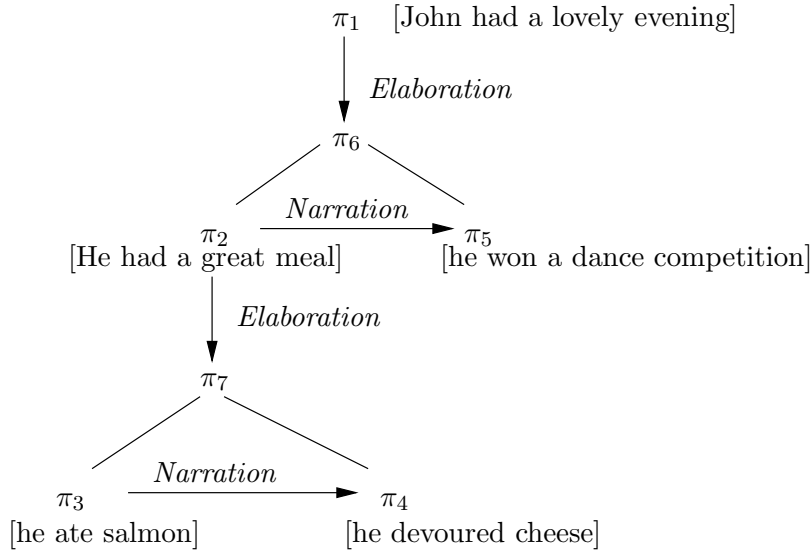


Figure 5: A Graphical Representation of the SDRS (5')

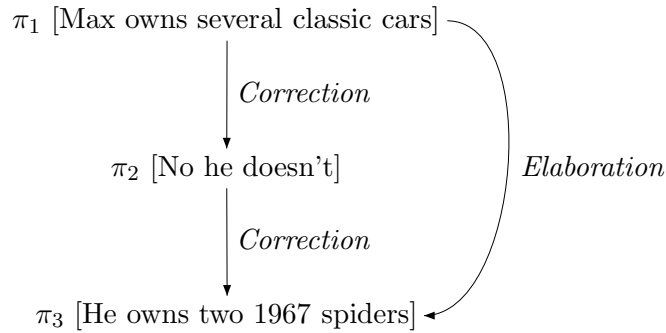


Figure 6: The SDRS for (18)

It also allows a given utterance to be rhetorically connected to more than one proposition in the context. This also allows an utterance to make more than one illocutionary contribution to the discourse, this time to more than one part of the context. For example, an adequate interpretation of dialogue (18) requires *both* the relations $Correction(\pi_2, \pi_3)$ and $Elaboration(\pi_1, \pi_3)$ to be identified, yielding the SDRS depicted in Figure 6, for logically inferring one of these relations is co-dependent on inferring the other.

- (18) π_1 . A: Max owns several classic cars.
 π_2 . B: No he doesn't.
 π_3 . A: He owns two 1967 Alfa spiders.

Before examining the semantics of SDRSs, let's define the constraints on anaphora, for like the accessibility constraints of DRT, these are defined in terms of the *form* of an SDRS but not its interpretation. We first define which labels in the SDRS new information can attach to with a rhetorical relation, and then define constraints on antecedents in terms of that. The following

definition doesn't apply when the rhetorical relations *Contrast* or *Parallel* are present; for details of the special effects of these relations, see (Asher, 1993, Asher and Lascarides, 2003).

Definition 8 Availability

Let $\langle A, \mathcal{F}, LAST \rangle$ be an SDRS, and let K_β (which we label β) be new information. Then β can attach with a rhetorical relation to:

1. The label $\alpha = LAST$;
2. Any label γ such that:
 - (a) $Succ(\gamma, \alpha)$; or
 - (b) $\mathcal{F}(\lambda) = R(\gamma, \alpha)$, for some label λ , where R is a subordinating discourse relation (*Elaboration*, *Explanation* etc.).

We gloss this as $\alpha < \gamma$.

3. *Transitive Closure*: Any label γ that dominates α through a sequence of labels $\gamma_1, \dots, \gamma_n$ such that $\alpha < \gamma_1 < \dots < \gamma_n < \gamma$.

Let K_β contain an anaphoric condition φ . Then the available antecedents are:

1. in K_β and DRS-accessible to φ ; or
2. in K_α , DRS-accessible to any condition in K_α , and there is a condition $R(\alpha, \gamma)$ in the SDRS such that $\gamma = \beta$ or γ outscopes β (i.e., γ is related to β by a sequence of *Succ* relations).

This definition ensures that antecedents to anaphora must be DRS-accessible on the right frontier of the discourse structure (unless *Contrast* and *Parallel* are present, where in certain cases DRS-inaccessible discourse referents become available; see Asher and Lascarides (2003) for details). This means that *It was a beautiful pink* in (5) isn't acceptable: the discourse referent introduced by *salmon* is in K_{π_3} , and π_3 is not on the right frontier and so *It was a beautiful pink* cannot attach to it. Thus SDRT's constraints on anaphora refines those of DRT. It also refines those of AI-based theories of discourse structure which adopt the right-frontier constraint (e.g., Grosz and Sidner (1986), Webber (1991)). SDRT correctly predicts that (1b) is odd, because although π_2 can attach to π_1 , the discourse referent introduced by *every man* in π_1 isn't DRS-accessible in K_{π_1} , and hence it's not available to *he* in π_2 :

- (1) b. π_1 . Every man walked in.
 π_2 . ??He ordered a beer.

4.1 The Dynamic Semantics of SDRSs

By supplying truth definitions for all SDRS-formulae, one can assign an interpretation to the formula $\mathcal{F}(\pi_0)$, where the label π_0 outscopes all other labels in the SDRS (a unique such label must exist by the definition of SDRSs). Typically, $\mathcal{F}(\pi_0)$ is some boolean combination of formulae of the form $R(\pi_i, \pi_j)$ (e.g., see (5')). We also define the interpretation of $R(\pi_i, \pi_j)$

in terms of K_{π_i} and K_{π_j} (i.e., $\mathcal{F}(\pi_i)$ and $\mathcal{F}(\pi_j)$). Thus, we can recursively unpack the semantics of an SDRS, culminating in interpreting the contents of the clauses. These contents are typically DRSS, with the semantics given in Definition 4. SDRT’s semantics in fact extends DRT’s semantics.

We need to assign a semantics to rhetorical relations. For the sake of simplicity, we will restrict attention here to an extensional approximation of the intensional semantics for rhetorical relations (to handle intensional relations one needs to make contexts into pairs of the form (w, f) , where w is a possible world and f is an assignment function, and this induces more complications than we want to go into here). Unlike predicate symbols such as *love* within the DRT vocabulary, rhetorical relations are not interpreted as imposing *tests* on the input information state. Rather, they define a real transition across information states. For example, *veridical* rhetorical relations satisfy the schema given below:

- **Satisfaction Schema for Veridical Rhetorical Relations**

$$f\llbracket R(\pi_1, \pi_2) \rrbracket_M g \text{ iff } f\llbracket K_{\pi_1} \rrbracket_M \circ \llbracket K_{\pi_2} \rrbracket_M \circ \llbracket \phi_{R(\pi_1, \pi_2)} \rrbracket_M g$$

where $\phi_{R(\pi_1, \pi_2)}$ expresses the semantic constraints pertinent to the particular rhetorical connection $R(\pi_1, \pi_2)$.

Veridical relations include *Narration*, *Explanation*, *Elaboration*, *Background*, *Contrast* and *Parallel*. This schema ensures that these relations entail the two propositions they connect. It contrasts with non-veridical relations such as *Alternation* (which is the SDRT way of expressing *or*), and relations such as *Correction*, where $f\llbracket Correction(\pi, \pi') \rrbracket g$ entails $f\llbracket \neg K_\pi \rrbracket f$ (for a full definition of *Correction* see Asher and Lascarides (2003)). More generally, rhetorical relations act semantically like complex update operators, and their interpretation reflects the special semantic influence that they have on the propositions they connect. It also reflects their status as *speech acts*: like other kinds of actions, they change the context. We’ll focus attention on veridical relations from now on.

Interpreting SDRSS involves defining the values of $\phi_{R(\pi_1, \pi_2)}$ for various relations R . For most relations, $\phi_{R(\pi_1, \pi_2)}$ is defined in terms of K_{π_1} and K_{π_2} or the discourse referents introduced in them. For example Asher and Lascarides (2003) define $\phi_{Narration(\pi_1, \pi_2)}$ to mean the end of the main eventuality (or semantic index in HPSG terms) e_{π_1} in K_{π_1} overlaps, both in space and time, with the beginning of the main eventuality e_{π_2} in K_{π_2} . This ensures that so long as the logical form of (7) contains *Narration*(π_1, π_2), the interpretation of this logical form entails the temporal progression of the events (i.e., Max falling precedes John helping him up). It also places the boxcar in the narrative (15ab) in Dansville (for the event of picking up the boxcar starts in Dansville, since this is where the event in (15a) ended). The location of the boxcar is different in (15ab’) because *Parallel*(π_1, π_2) imposes different spatio-temporal constraints from *Narration*.

In contrast, $\phi_{Explanation(\pi_1, \pi_2)}$ entails *cause*(e_{π_2}, e_{π_1}). Thus the logical form (8’) of (8) entails that the pushing caused the falling, even though the compositional semantics of the clauses don’t entail this (note that we have assumed in (8’) that *him* binds to *Max*; by availability this is the only choice).

(8) Max fell. John pushed him.

We'll examine how one constructs this logical form for (8) in Section 5. Here we show how it's interpreted. This SDRS relates the input variable assignment function f to g iff the content K_{π_0} that \mathcal{F} assigns to the highest label π_0 does this. I.e., $f \llbracket K_{\pi_0} \rrbracket_M g$. K_{π_0} is $Explanation(\pi_1, \pi_2)$. So $f \llbracket K_{\pi_0} \rrbracket_M g$ iff $f \llbracket Explanation(\pi_1, \pi_2) \rrbracket_M g$. According to the semantics of $Explanation$ this holds iff:

1. There is an h such that $f \llbracket K_{\pi_1} \rrbracket_M h$, and
2. There is an i such that $h \llbracket K_{\pi_2} \rrbracket_M i$ and
3. $i \llbracket cause(e_{\pi_2}, e_{\pi_1}) \rrbracket_M g$.

Clause 1. holds iff h is defined for x and e_{π_1} and $\langle h(e_{\pi_1}), h(x) \rangle \in I_M(fall)$ etc. Clause 2. holds iff i extends h such that and $\langle i(e_{\pi_2}), i(y), i(x) \rangle \in I_M(push)$ etc. Finally, clause 3. holds iff $i = g$ and $\langle i(e_{\pi_2}), i(e_{\pi_1}) \rangle \in I_M(cause)$. In other words, Max fell, John pushed him, and the latter event caused the former.

We have now introduced the language of SDRSs and their dynamic semantic interpretation, which in turn makes sense of the availability constraint on anaphora which we defined in terms of rhetorical structure. The question now arises as to how one constructs these logical forms for discourse.

5 Constructing Logical Form

To construct a conceptually clean account of how to reason about discourse structure and construct logical forms, SDRT distinguishes between the SDRSs themselves and a language in which we describe them. As interpreters attempt to reconstruct the intended logical form of a discourse, they must confront many ambiguities: the grammar and lexical semantics typically underdetermines the intended logical form thanks to semantic scope ambiguities, anaphora of various kinds such as pronouns and presuppositions, and lexical ambiguities. SDRT contains a description-language \mathcal{L}_{ulf} which allows us to analyse and to reason about such underdetermination (ULF stands for *underspecified logical form*). SDRT's glue logic then defines the pragmatically preferred ways of resolving semantic underspecification.

5.1 The Language \mathcal{L}_{ulf} of Semantic Underspecification

The language \mathcal{L}_{ulf} partially describes the form of SDRSs. It allows us to express how a given knowledge source, such as the grammar, yields only partial information about content. Let's clarify the idea with an example. Ignoring rhetorical relations for now, sentence (19) contains a two-way semantic scope ambiguity between the indefinite NP and *might* and an anaphoric ambiguity, as given by *him*.

(19) A man might push him.

Let's assume that the discourse context is such that there are two available antecedents for *him*: z_1 and z_2 . Then there are four fully determinate logical forms for (19). Two of these




Figure 7: Two logical forms for (19), shown as trees

are shown in Figure 7 (the first one corresponds to \exists outscoping *might* and *him* resolving to z_1 , and the second corresponds to *might* outscoping \exists and *him* resolving to z_2); note that these trees show the *form* of the determinate logical forms. We want this form to be all that the description-language \mathcal{L}_{ulf} ‘knows’ about. In fact, the two trees in Figure 7 will each correspond to a *model* of \mathcal{L}_{ulf} , so that $M \models_{\mathcal{L}_{ulf}} \phi$ corresponds to: the ULF ϕ (partially) describes the unique determinate logical form that corresponds to M .

But how do we express partial descriptions of such trees? In this example, what’s the formula or ULF ϕ in \mathcal{L}_{ulf} that describes just the four trees or determinate logical forms for (19) and no others? Well, following the usual strategy (Bos, 1995, Asher and Fernando, 1997, Copestake et al., 1999), \mathcal{L}_{ulf} ’s vocabulary consists of labels which pick out nodes in the trees of Figure 7. These labels allow one to talk independently about on the one hand the logical connectives, predicate symbols and variables that are present in the determinate logical form and on the other hand the way they are combined. Thus labels tag bits of content (as expressed in the SDRS-language); in fact, all constructors (\wedge , $=$, *man*, *x* etc.) in the SDRS-vocabulary become predicate symbols over labels in \mathcal{L}_{ulf} .

We can then express partial information about semantic scope by underdetermining the outscopes constraints on labels (in this case, the ULF will underdetermine the relative semantics scopes of the label that tags \exists and the label that tags *might*). Information about anaphoric conditions amounts to not knowing the value of an SDRS-discourse referent (at least, for pronouns referring to individuals). Discourse referents become *one-place predicates* in \mathcal{L}_{ulf} , the argument of the predicate being reserved for the label that tags its position in the ‘trees’ of the kind shown in Figure 7. So the compositional semantics of a pronoun involves not knowing the value of a one-place predicate in \mathcal{L}_{ulf} , and is thus represented with a high-order variable. For simplicity, we gloss the anaphoric condition as $x = ?$ as given earlier, although in fact one should think of x and $?$ as one-place predicate symbols in \mathcal{L}_{ulf} , and one should also bear in mind that this gloss ignores the labels indicating their position in the trees of Figure 7. So the ULF for (19) (in simplified notation, where labels that don’t contribute to the semantic ambiguities are ignored) is (19’); we’ve shown this graphically in Figure 8 (where curved arrows convey the outscopes conditions).

$$\begin{aligned}
 (19') \quad & l_1 : \exists(x, \text{man}(x), l_2) \wedge \\
 & l_3 : \text{might}(l_4) \wedge \\
 & l_5 : \wedge(l_6, l_7) \wedge l_6 : \text{push}(x, y) \wedge l_7 : x = ? \wedge \\
 & \text{outscopes}(l_4, l_5) \wedge \text{outscopes}(l_2, l_5)
 \end{aligned}$$




Figure 8: A graphical representation of the ULF (19')

\mathcal{L}_{ulf} can also express underspecified information about rhetorical connections. For example, $?(π_1, π_2, π_0)$ expresses the information that $π_1$ and $π_2$ are rhetorically connected and the resulting connection is labelled $π_0$, but the value of the rhetorical relation is unknown (as shown by the higher-order variable ‘?’). The compositional semantics of sentence-initial *but* includes the underspecified condition $Contrast(?_1, π_2, ?_2)$ (again, this is a notational gloss for the formula involving higher-order variables), where $π_2$ is the top label of the clause that’s syntactically outscoped by *But* in the grammar. This indicates that sentence-initial *but* generates a *Contrast* relation between the label of a proposition that’s not determined by the grammar (although it may be by some other knowledge source) and the label of the proposition denoted by the clause that’s syntactically outscoped by *but*, and the label that’s assigned to this *Contrast* connection in the SDRS is also unknown.

The satisfaction relation $\models_{\mathcal{L}_{ulf}}$ is defined relative to finite first-order models (i.e., trees like those in Figure 7) and the higher-order variables which are used to express unknown values of predicate symbols etc. in the SDRS are interpreted *substitutionally*. In fact, $\models_{\mathcal{L}_{ulf}}$ is monotonic, extensional, static and decidable. This contrasts with the logical of SDRSs themselves, which is dynamic and undecidable. The difference comes from the fact that $\models_{\mathcal{L}_{ulf}}$ consists of reasoning only about the *form* of SDRSs but not their (dynamic) interpretation. In essence, $\models_{\mathcal{L}_{ulf}}$ relates a ULF to all possible ways of resolving the underspecification, making it unnecessary to define separately a notion of *supervaluation*, contra Reyle (1993). However, the SDRT framework not only defines the possible ways of completing an underspecified logical form, it defines the pragmatically preferred ways of doing it. This is part of the definition of discourse update, which is defined in terms of the glue logic.

5.2 The Glue Logic

The *glue logic* of SDRT defines a nonmonotonic consequence relation \vdash_g over \mathcal{L}_{ulf} . Together with the principle MDC described earlier, this defines the *pragmatically preferred* interpretations of ULFs. In general, pragmatically preferred interpretations are more informative: they are a subset of the possible interpretations. Or to put it another way, \vdash_g generates more consequences than $\models_{\mathcal{L}_{ulf}}$.

The glue logic has only limited access to the logic of SDRSS: since it accesses only ULFs it knows about the forms of SDRSS but not their (dynamic) interpretations. It also has only restricted access to information in domain knowledge, the lexicon and cognitive states. The relationship between these richer knowledge sources and their shallow form in the glue language is very similar to the relationship between SDRSS proper and their corresponding ULFs. Building these ‘porous’ fences between the information sources that contribute to discourse interpretation and the logic in which logical form is constructed is the only way of ensuring that constructing logical form—or equivalently, computing what is said—is computable.

The glue logic in combination with MDC then determines the following logically dependent information:

1. the (pragmatically preferred) values of certain underspecified conditions that are generated by the grammar;
2. which labels are rhetorically connected to which other labels (this is equivalent to the task of text segmentation);
3. the values of the rhetorical relations.

This information is computed on the basis of inferences over default axioms within the glue logic, written $A > B$ (which is read as *If A then normally B*). These express information about pragmatically preferred values of underspecified conditions in a given ULF. SDRT thus enriches dynamic semantics with contributions from pragmatics in a constrained way. It’s a contribution from pragmatics in that the default axioms are justified on the basis of pragmatic information such as domain knowledge and cognitive states; it’s constrained because of its limited access to these information sources.

Many glue-logic axioms are schemata of the form (20) (where α , β and λ are metavariables over SDRS-labels π_1 , π_2 etc.):

$$(20) \quad (?(\alpha, \beta, \lambda) \wedge Info(\alpha, \beta, \lambda)) > R(\alpha, \beta, \lambda)$$

In words, if β is to be attached to α with a rhetorical relation and the result is labelled λ , and information $Info(\alpha, \beta, \lambda)$ about α , β and λ , that is transferred into the glue logic from more expressive languages such as that of SDRSS, the lexicon, domain knowledge and cognitive states holds, then normally, the rhetorical connection is R . Observe that $Info(\alpha, \beta, \lambda)$ expresses information from rich knowledge sources that contribute to discourse interpretation in a *shallow* form: for example the discourse content present in SDRSS is transferred into the glue logic in a shallow form, as expressed in \mathcal{L}_{ulf} .

For example, *Narration* stipulates that if β is to be connected to α and α occasions β , then normally the relation is *Narration*:

- **Narration:** $(?(\alpha, \beta, \lambda) \wedge occasion(\alpha, \beta)) > Narration(\alpha, \beta, \lambda)$.

Scriptal knowledge can be used to infer *occasion* predicates (by default), and such knowledge takes the following form: if two event types of a certain kind (ϕ and ψ) are to be related, then *occasion* can normally be inferred:

- **Scripts for Occasion:** $(?(\alpha, \beta, \lambda) \wedge \phi(\alpha) \wedge \psi(\beta)) > occasion(\alpha, \beta)$

Of course this isn't a general schema in the sense that *any* ϕ and ψ will do. Instances of **Scripts for Occasion** will depend on particular semantic contents of the clauses involved. For example, we assume there's scriptal information or 'domain knowledge' that if x 's falling and y 's helping x up are connected somehow, then the former occasioned the latter (we forego giving the formal axiom here, but see Asher and Lascarides (2003)).

By contrast, $Explanation(\alpha, \beta)$ can be inferred when there's evidence in the discourse that β causes α . *Evidence* of a causal relation is distinct from a causal relation actually holding; the glue logic expresses evidence in the discourse of a causal relation with $cause_D(\beta, \alpha)$ and the *actual* causal relation between events with $cause(e_\beta, e_\alpha)$; note that the former does *not* entail the latter. However, given the default rule **Explanation** and the semantics of $Explanation$ given earlier, evidence in the discourse of a causal relation non-monotonically yields an SDRS which does entail an actual causal relation (e.g., the SDRS in (8')):

- **Explanation:** $(?(\alpha, \beta, \lambda) \wedge cause_D(\beta, \alpha)) > Explanation(\alpha, \beta, \lambda)$.

Glue-logic axioms for inferring $cause_D(\beta, \alpha)$ are *monotonic*, for either the discourse contains evidence of a causal connection or it doesn't. For example, **Causation and Change** stipulates that if e_α describes a change (of some kind, e.g., a change in location) in y , and e_β describes a force that would cause a change (of that same kind, in this case a change in location), then $cause_D(\beta, \alpha)$ holds:

- **Causation and Change:**
 $(change(e_\alpha, y)) \wedge cause-change-force(e_\beta, x, y) \rightarrow cause_D(\beta, \alpha)$

This rule applies in the analysis of (8):

- (8) π_1 . Max fell.
 π_2 . John pushed him.

Lexical semantics stipulates that *fall* is a verb describing a change in location and *push* is a verb describing a force that causes a change in location (see Asher and Lascarides (2003) for detailed motivation for this position). Moreover, if the discourse is coherent, then the π_1 must be connected to π_2 with a rhetorical relation. Hence (π_1, π_2, π_0) holds. By Definition 8, this means that the only available antecedent for the pronoun is *Max*. Thus the information about content that's transferred into the glue logic from \mathcal{L}_{ulf} (and the lexicon) verifies the antecedent of **Causation and Change**, and so $cause_D(\pi_2, \pi_1)$ is inferred. Thus the antecedent to **Explanation** is verified, and so $Explanation(\pi_1, \pi_2, \pi_0)$ is non-monotonically inferred. The definition of SDRT update given in the next section uses this output of the glue logic to ensure that the (pragmatically preferred) logical form for (8) is the SDRS in (8'). Given the semantics of this SDRS—and in particular the dynamic interpretation of $Explanation(\pi_1, \pi_2)$ —discourse (8) is correctly predicted to mean: Max fell, John pushed Max, and the latter caused the former.

The relation *Elaboration* is inferred via axioms that are analogous to the ones for *Explanation*, save now the discourse evidence is for subtype information (written $subtype_D$):

- **Elaboration:** $(?(\alpha, \beta, \lambda) \wedge subtype_D(\beta, \alpha)) > Elaboration(\alpha, \beta, \lambda)$

5.3 Discourse Update

The consequences of the glue logic are used to build the SDRS for discourse (8). Discourse update in SDRT is entirely declarative, and in the absence of divergent relations such as *Correction* (which we won't consider here), it's also a monotone decreasing function thereby reflecting the idea that one monotonically accumulates more information as the discourse is interpreted.

Discourse update is defined as a sequence of simple (and monotone decreasing) update operations $+$, where $+$ is defined in terms of the glue-logic consequence relation \sim_g . This simple update operation works over the set σ of SDRSs which represents the content of the discourse context. The ULF which describes this set of SDRSs is the theory of σ in \mathcal{L}_{ulf} , written $Th(\sigma)$. $+$ also takes as input some new information: this is either a ULF \mathcal{K}_β (e.g., this could be the ULF of a clause as generated by the grammar), or it is an assumption $?(\alpha, \beta, \lambda)$ about attachment where $Th(\sigma) \models_{\mathcal{L}_{ulf}} \mathcal{K}_\beta$ (in other words, the ULF \mathcal{K}_β is part of the description of logical form already). The result of $+$ over these arguments is a set σ' of discourse structures which (a) is a subset of the old information σ in that it satisfies this old information and also the new information; and (b) it also ensures that any \sim_g -consequences of the old information and the new are satisfied too. Note that all monotonic consequences are nonmonotonic consequences, and so ensuring that the updated context satisfies (b) will make it satisfy (a) as well (because the old and the new information follows monotonically from itself). More formally:

Definition 9 The Simple Update $+$

Let σ be a set of (fully-specified) discourse structures. And let ψ be either (a) a ULF \mathcal{K}_β , or (b) a formula $?(\alpha, \beta, \lambda)$ about attachment, where $Th(\sigma) \models_{\mathcal{L}_{ulf}} \mathcal{K}_\beta$. Then $\sigma + \psi$ is a set of SDRSs defined as follows:

1. $\sigma + \psi = \{\tau : \text{if } Th(\sigma), \psi \sim_g \phi \text{ then } \tau \models_{\mathcal{L}_{ulf}} \phi\}$, provided the result is not \emptyset ;
2. $\sigma + \psi = \sigma$ otherwise.

Recall that SDRSs are in effect models of the ULF-logic \mathcal{L}_{ulf} . $\tau \models_{\mathcal{L}_{ulf}} \phi$ means that ϕ (partially) describes τ . Simple update is thus defining the set of SDRSs which comply with a *partial description* of the logical form of the discourse, this partial description being the conjunct of the \sim_g -consequences as shown (making it a formula of \mathcal{L}_{ulf}). In essence, $+$ defines a constraint-based approach to constructing logical form: it uses the old information and the new to accumulate constraints on the form of the SDRS which ultimately represents the interpretation of the updated discourse.

Discourse update $update_{SDRT}$ itself is very conservative: it remains neutral about what's attached to what. In other words, suppose that A is the set of available attachment sites in the old information σ for the new information β . Then the power set $\mathcal{P}(A)$ represents all possible choices for what labels α_i in σ the new label β is actually attached to. $update_{SDRT}$ is neutral about which member of $\mathcal{P}(A)$ is the 'right' choice, for $update_{SDRT}(\sigma, \mathcal{K}_\beta)$ is the *union* of SDRSs that result from a sequence of $+$ -operations for each member of $\mathcal{P}(A)$ (we forego giving

the formal definition here, but see Asher and Lascarides (2003) for details). Since $update_{\text{SDRT}}$ is defined in terms of $+$, it is also a monotone decreasing function, reflecting the idea that interpreting discourse amounts to a (monotonic) accumulation of information.

Any satisfiable set of statements in \mathcal{L}_{ulf} describe a countably infinite set of equivalence classes of SDRSS (where equivalence is alphabetic variance). To see why, simply observe that a discourse can continue in an indefinite number of ways. So the output of $+$ can be a countably infinite set. This has no adverse computational effects on SDRT update however. Performing updates is simply a matter of accumulating more and more constraints in the description language \mathcal{L}_{ulf} as \sim_g -consequences, as shown above. If at any point during discourse processing one wants to actually *interpret* the discourse (so far), then one needs to construct all *pragmatically preferred* SDRSS which satisfy the description (that’s accumulated so far). Note that while the glue logic uses pragmatic information to compute rhetorical relations, thereby ensuring that $+$ eliminates some pragmatically inadmissible logical forms, *ranking* the models in the update is done via the principle MDC given in Definition 5 together with the following additional factor that determine ranking: we prefer models (or SDRSS) with a minimum number of labels. The content of a discourse at any given point will be those things that follow from the highest ranked SDRSS in the update. In essence, only a subset of SDRSS in the update are the ones that ‘matter’, and because of the minimality constraint there are a finite number of these (up to alphabetic variance).

This ranking of models by MDC (plus minimality) is in fact what influences our inferences about what’s rhetorically connected to what in more complex discourses such as (5).

- (5)
- π_1 . John had a great evening last night.
 - π_2 . He had a great meal.
 - π_3 . He ate salmon.
 - π_4 . He devoured lots of cheese.
 - π_5 . He won a dancing competition.

The output of $update_{\text{SDRT}}$ for (5) contains many different SDRSS, each with different assumptions about which labels are rhetorically connected. However, the highest ranked SDRSS in the update according to MDC are those with the minimum number of labels, the maximum number of rhetorical connections, the fewest unresolved semantic ambiguities (including anaphoric conditions) and no inconsistencies. These principles determine that the SDRS in Figure 4 is the preferred model in the ranking. A full analysis is given in Asher and Lascarides (2003), but to illustrate the point we focus on one particular decision: what π_5 attaches to.

Given the interpretation of the prior context π_1 – π_4 , Definition 8 means that there are five available labels for π_5 to attach to: π_4 (because this is the *LAST* label), π_7 (because it immediately outscopes π_4), π_2 (because it’s attached to π_7 with the subordinating relation *Elaboration*), π_1 (because it’s attached to π_2 with *Elaboration*) and π_0 (because it immediately outscopes π_1 and π_2). Note that π_3 isn’t available. Thus $update_{\text{SDRT}}$ will output the SDRSS that follow from any combination of these attachment assumptions, and MDC must then rank these choices. There are no glue-logic axioms which allow us to infer *occasion*, *subtype_D* or *cause_D* for linking π_5 and π_4 , and so if π_4 is one of the actual attachment sites, then the update would include the underspecified condition $?(\pi_4, \pi_5, \pi')$. The same holds for the link

between π_6 and π_5 . However, attempting to attach π_5 to just π_1 and π_2 yields something more coherent according to our assumptions, in that the update won't include these (rhetorical) underspecifications. There is subtype information we can exploit in attaching π_5 to π_1 , yielding *Elaboration*(π_1, π_5). So eating the meal and winning the dance competition are both part of the evening. This additional information verifies an *occasion*-axiom for π_5 and π_2 , yielding *Narration*(π_2, π_5). So MDC determines that the SDRSS with highest ranking are those where π_5 attaches to π_1 and to π_2 , but *not* to π_0 (such an attachment would not allow us to resolve the pronoun *he* in π_5), π_4 or π_6 . This is exactly the SDRS shown in Figure 4.

As we explained earlier, similar reasoning involving MDC predicts the correct interpretations of the presuppositions in (12a) vs. (12b), the lexical sense disambiguations in (16) and (17), and the bridging inferences in (14) and (15) though we forego spelling out the formal details here.

6 Conclusion

We have presented brief highlights of Segmented Discourse Representation Theory or SDRT. SDRT is distinct from other dynamic semantic theories in that it enriches logical forms to include rhetorical relations, to which SDRT assigns a semantics, making them complex update operators. Indeed, all logical forms are interpreted compositionally and dynamically.

SDRT refines the accessibility constraint on anaphora, replacing it with the notion of availability which takes *both* logical structure *and* rhetorical structure into account. Because logical structure is a factor in blocking off antecedents, it also refines the right-frontier constraint from AI-based work on discourse structure which ignores this information source.

SDRT includes a language \mathcal{L}_{ulf} in which logical forms of discourse are *described*. This language essentially knows about the form of an SDRS, but not its (dynamic) interpretation. It's a language in which one can express information about semantic underspecification, and its consequence relation captures the relation between an underspecified logical form and all its possible interpretations. Discourse update in SDRT takes this further, defining a relation between an underspecified logical form and its *pragmatically preferred* interpretations. This is achieved via (a) the glue logic, which consists of axioms describing default values for certain underspecified semantic conditions; and (b) the principle MDC, which imposes a ranking on the set of SDRSS that are output by $update_{SDRT}$, as determined by the glue logic consequence relation \sim_g . The process of computing these pragmatically preferred logical forms is decidable, unlike the interpretation mechanisms described in much of the AI-research (e.g., Hobbs et al. (1993), Lochbaum (1998), Grosz and Sidner (1990), Traum and Allen (1994)). We believe that this is crucial for an adequate model of semantic competence, since it's essential to explaining why competent language users by and large agree on what was said, if not its consequences.

References

N. Asher. *Reference to Abstract Objects in Discourse*. Kluwer Academic Publishers, 1993.

- N. Asher and T. Fernando. Labelling representation for effective disambiguation. In *Proceedings of the 2nd International Workshop on Computational Semantics (IWCS-2)*, pages 1–4, Tilburg, 1997.
- N. Asher and A. Lascarides. The semantics and pragmatics of presupposition. *Journal of Semantics*, 15(2):239–299, 1998.
- N. Asher and A. Lascarides. *Logics of Conversation*. Cambridge University Press, 2003.
- A. Asudeh and R. Crouch. Glue semantics for HPSG. In *Proceedings of HPSG-2001*, Trondheim, Norway, 2001.
- D. Beaver. Local satisfaction preferred. In P. Dekker and M. Stokhof, editors, *Proceedings of the 10th Amsterdam Colloquium*, ILLC, Amsterdam, 1996.
- J. Bos. Predicate logic unplugged. In *Proceedings of the 10th Amsterdam Colloquium*, pages 133–143, Amsterdam, 1995.
- A. Copestake, D. Flickinger, I. A. Sag, and C. Pollard. Minimal recursion semantics: An introduction. available from <http://www-csli.stanford.edu/~aac>, 1999.
- A. Copestake, A. Lascarides, and D. Flickinger. An algebra for semantic construction in constraint-based grammars. In *Proceedings of the 39th Annual Meeting of the Association for Computational Linguistics (ACL/EACL 2001)*, pages 132–139, Toulouse, 2001.
- I. Dagan, L. Lee, and F. C. N. Pereira. Similarity-based methods for word sense disambiguation. In *Proceedings of the 35th Annual Meeting of the Association for Computational Linguistics and the 8th Meeting of the European Chapter for the Association for Computational Linguistics (ACL97/EACL97)*, pages 56–63, Madrid, 1997.
- D. Davidson. *Essays on Actions and Events*. Clarendon Press, 1980.
- T. Fernando. Bisimulations and predicate logic. *Journal of Symbolic Logic*, 59(3):924–944, 1994.
- B. Geurts. Local satisfaction guaranteed. *Linguistics and Philosophy*, 19:259–294, 1996.
- J. Groenendijk and M. Stokhof. Dynamic predicate logic. *Linguistics and Philosophy*, 14: 39–100, 1991.
- B. Grosz, A. Joshi, and S. Weinstein. Centering: A framework for modelling the local coherence of discourse. *Computational Linguistics*, 21(2):203–226, 1995.
- B. Grosz and C. Sidner. Attention, intentions and the structure of discourse. *Computational Linguistics*, 12:175–204, 1986.
- B. Grosz and C. Sidner. Plans for discourse. In J. Morgan P. R. Cohen and M. Pollack, editors, *Intentions in Communication*, pages 365–388. MIT Press, 1990.
- J. Guthrie, L. Guthrie, Y. Wilks, and H. Aldinejad. Subject-dependent co-occurrence and word sense disambiguation. In *Proceedings of the 29th Association for Computational Linguistics*, pages 146–152, Berkeley, 1991.

- I. Heim. *The Semantics of Definite and Indefinite Noun Phrases*. PhD thesis, University of Massachusetts, 1982.
- J. R. Hobbs. On the coherence and structure of discourse. Technical Report CSLI-85-37, Center for the Study of Language and Information, Stanford University, 1985.
- J. R. Hobbs, M. Stickel, D. Appelt, and P. Martin. Interpretation as abduction. *Artificial Intelligence*, 63(1–2):69–142, 1993.
- Jerry R. Hobbs. Coherence and coreference. *Cognitive Science*, 3(1):67–90, 1979.
- H. Kamp. A theory of truth and semantic representation. In M. B. J. Stokhof J. A. G. Groenendijk, T. M. V. Janssen, editor, *Formal Methods in the Study of Language*, pages 277–322. Mathematisch Centrum, Amsterdam, 1981.
- H. Kamp and U. Reyle. *From Discourse to the Lexicon: Introduction to Modeltheoretic Semantics of Natural Language, Formal Logic and Discourse Representation Theory*. Kluwer Academic Publishers, 1993.
- H. Kamp and C. Rohrer. Tense in texts. In C. Schwarze R. Bauerle and A. von Stechow, editors, *Meaning, Use and Interpretation of Language*, pages 250–269. Berlin: de Gruyter, 1983.
- A. Koller, K. Mehlhorn, and J. Niehren. A polynomial-time fragment of dominance constraints. In *Proceedings of the 28th Annual Meeting of the Association for Computational Linguistics (ACL2000)*, Hong Kong, 2000.
- A. Lascarides and N. Asher. Temporal interpretation, discourse relations and commonsense entailment. *Linguistics and Philosophy*, 16(5):437–493, 1993.
- A. Lascarides, N. Asher, and J. Oberlander. Inferring discourse relations in context. In *Proceedings of the 30th Annual Meeting of the Association for Computational Linguistics (ACL92)*, pages 1–8, Delaware, 1992.
- K. E. Lochbaum. A collaborative planning model of intentional structure. *Computational Linguistics*, 24(4):525–572, 1998.
- W. C. Mann and S. A. Thompson. Rhetorical structure theory: A framework for the analysis of texts. *International Pragmatics Association Papers in Pragmatics*, 1:79–105, 1987.
- R. M. Montague. The proper treatment of quantification in ordinary english. In *Formal Philosophy: Selected Papers of Richard Montague*, pages 247–270. Yale University Press, 1974.
- R. Muskens. Combining montague semantics and discourse representation. *Linguistics and Philosophy*, 19:143–186, 1996.
- L. Polanyi. A theory of discourse structure and discourse coherence. In P. D. Kroeber W. H. Eilfort and K. L. Peterson, editors, *Papers from the General Session at the 21st Regional Meeting of the Chicago Linguistics Society*. 1985.
- U. Reyle. Dealing with ambiguities by underspecification: Construction, interpretation and deduction. *Journal of Semantics*, 10:123–179, 1993.

- M. Stone and R. Thomason. Context in abductive interpretation. In *Proceedings of the 6th International Workshop on the Semantics and Pragmatics of Dialogue (Edilog)*, Edinburgh, 2002.
- D. Traum and J. Allen. Discourse obligations in dialogue processing. In *Proceedings of the 32nd Annual Meeting of the Association for Computational Linguistics (ACL94)*, pages 1–8, Las Cruces, New Mexico, 1994.
- R. van der Sandt. Presupposition projection as anaphora resolution. *Journal of Semantics*, 9(4):333–377, 1992.
- B. L. Webber. Structure and ostension in the interpretation of discourse deixis. *Natural Language and Cognitive Processes*, 6(2):107–135, 1991.