Segmented Discourse Representation Theory: Dynamic Semantics with Discourse Structure

Pierre Lison
pierrel@coli.uni-sb.de

Seminar on Computational Models of Discourse, WS 2007-2008
Department of Computational Linguistics & Phonetics
Universität des Saarlandes

Introduction - What are discourse structures?

- As we'll already seen, the ability to extract and handle **discourse structure** is crucial for many NLP applications
- Discourse structures can be analysed and represented in different - albeit complementary - ways [Sporleder 07]:
  1. **Linguistic Structure**: linguistic manifestation of discourse structure, e.g., lexical cohesion, discourse connectives/cue words, intonation, gesture, referring expressions etc.
  2. **Intentional Structure**: each discourse segment fulfils a purpose (why does a speaker/write make a given utterance in a given form?)
  3. **Informational Structure**: how do the different segments of a discourse relate to each other (which discourse relations hold)?
  4. **Focus/Attentional Structure**: which entities are salient at a given point in discourse?

Introduction - What are discourse structures?

- As we'll already seen, the ability to extract and handle **discourse structure** is crucial for many NLP applications
- Discourse structures can be analysed and represented in different - albeit complementary - ways [Sporleder 07]:
  1. **Linguistic Structure**: linguistic manifestation of discourse structure, e.g., lexical cohesion, discourse connectives/cue words, intonation, gesture, referring expressions etc.
  2. **Intentional Structure**: each discourse segment fulfils a purpose (why does a speaker/write make a given utterance in a given form?)
  3. **Informational Structure**: how do the different segments of a discourse relate to each other (which discourse relations hold)?
  4. **Focus/Attentional Structure**: which entities are salient at a given point in discourse?

Introduction - Goal of this talk

In this talk, we'll introduce the core ideas of **Segmented Discourse Representation Theory** [Asher 03, Lascarides 07]:

- A formal approach to **discourse interpretation**,
- ... grounded in **dynamic semantics** - notably DRT [Kamp 93];
- ... and extended with **rhetorical relations**.

- In other words, SDRT is an attempt to model the **semantics-pragmatics interface**.

Outline of the talk

2 Background
- Dynamic Semantics
- Motivating rhetorical relations
- The SDRT approach

3 Segmented Discourse Representation Structures
- Syntax
- Detailed example
- Availability: the right frontier
- Semantics

4 Constructing Logical Forms
- The glue logic
- Defeasible reasoning
- Discourse Update

5 Summary

Traditional formal semantics

- In traditional formal semantics [Montague 88], the content of a discourse is defined as the **set of models** (''possible worlds'') that it satisfies.
- They are typically unable to model how the interpretation of the current sentence is **dependent on the interpretations of those that precede it** [Lascarides 07].
- Trouble handling most intersentential phenomena, like temporal and pronominal anaphora:
  (1) The man walked in.
  (2) He ordered a beer.
  $\Rightarrow$ *How to express the fact that the man who ordered a beer is the same as the one who walked in?*

Dynamic semantics

- Dynamic semantics views the meaning of a given discourse as a **relation** (or more precisely, a *function*) between contexts. This function is called *Context Change Potential*.
- Contrarily to Montagovian semantics, dynamic semantics is generally **non-compositional** (ie. you can't define the meaning of a discourse as a simple, static composition of its parts)
- In addition to contributing to the "'static"' content of a discourse, expressions like indefinite NPs also contribute *dynamically* to it by introducing new referents.
- Most well-known theory based on dynamic semantics: *Discourse Representation Theory* [Kamp 93]

Dynamic semantics - Simple example of a DRS

(1) The man walked in.
(2) He ordered a beer.

Box-style notation of the final **Discourse Representation Structure** of (1)-(2):

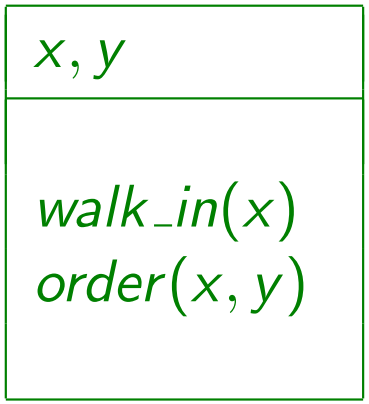

Shortcomings of dynamic semantics

- Dynamic semantics theories typically explore a relatively restricted set of pragmatic phenomena, mainly focusing on the effects of logical structure on anaphora.
- They typically fail to take into account the **discourse structure** (ie. rhetorical relations between discourse segments).
- And, as we've already seen last week when we examined *Rhetorical Structure Theory* [RST], understanding discourse structure is important for discourse interpretation.
- In order to motivate the need for rhetorical relations, let's analyse in detail 2 types of discourse phenomena: *pronoun resolution* and *temporal structure*.

Motivating rhetorical relations: Pronouns (1)

- Consider this simple discourse:

  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.

- In DRT, nothing would prevent the pronoun it inπ_6 to pick the salmon as the referent. The theory clearly overgenerates the possible interpretations.
- If we had some notion of a rhetorical structure for this discourse, we would be able to specify more precisely the possible antecedents for the pronoun.

Motivating rhetorical relations: Pronouns (2)

- Rhetorical structure for the given example:

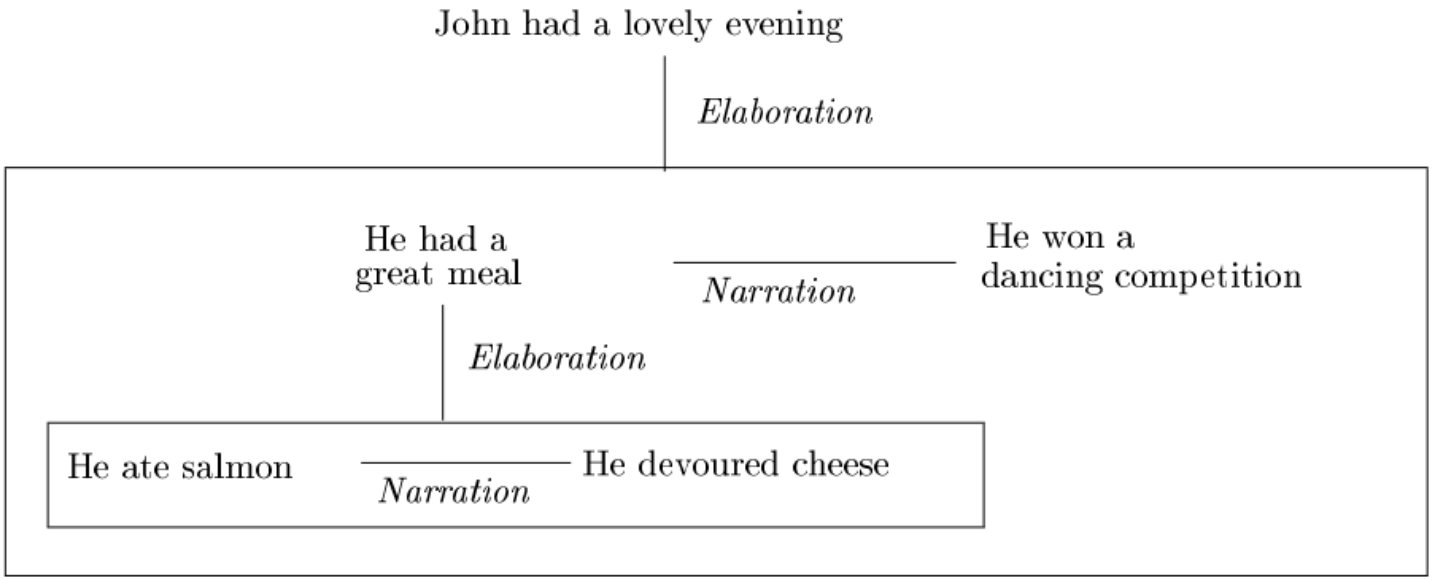

- Using the so-called *right frontier constraint* (more detail on this later), we can then easily rule out the salmon as an antecedent for it.

Motivating rhetorical relations: Temporal structure

- Consider these two examples:

  (1) John fell. Mary helped him up.
  (2) John fell. Mary pushed him.

- In (1), the textual order reflects the temporal one, whereas (2) doesn't.
- The compositional semantic forms of (1) and (2) are insufficient for distinguishing their interpretations: they have the same tense and the same aspectual classes.
- The additional bit of information we need resides in rhetorical relations: *Narration* for (1) and *Explanation* for (2).

The SDRT approach

- SDRT seeks to combine two paradigms in discourse interpretation: dynamic semantics and discourse analysis.
- To put it shortly: *SDRS = DRT + discourse structure*
- The theory attempts to explicit the interactions between the semantic content of the segments and the global, pragmatic structure of the discouse.
- It can thus be seen as a model of the **semantics-pragmatics interface**.
- The basic units are segmented and analysed according to their propositional content, and not eg. on their attentional or intentional content, like in [Grosz 86].

Segmented Discourse Representation Structures (1)

- The formal representations derived for a given discourse according to SDRT are called *Segmented Discourse Representation Structures*.
- Formally, a SDRS is a structure $\langle A, F, LAST \rangle$, where:
  - A is a set of labels (speech acts discourse referents)
  - F maps labels to SDRS-formulae (i.e., labels tag content)
  - $LAST$ is a label (of the last utterance)

Segmented Discourse Representation Structures (2)

- A SDRS-Formula can be either:
  - A DRS,
  - $R(\pi, \pi')$, where R is a rhetorical relation and π and π' are labels.
  - Boolean combinations of these.
- In addition, the following constraint is imposed on A: Let $Succ(\pi, \pi')$ means that $R(\pi'', \pi')$ or $R(\pi', \pi'')$ is a literal in $F(\pi)$. Then A must form a partial order under $Succ$ with a unique root.

SDRSs allow plurality

- **Of relations**: Contrast(π_1 , π_2), Narration(π_1 , π_2)
  π_1: Did you buy the apartment?
  π_2: Yes, but we rented it.

- **Of attachment site**s: Correction(π_2 , π_3), Elaboration(π_1 , π_3)
  π_1: Max owns several classic cars.
  π_2: No he doesn't.
  π_3: He owns two 1967 Alfa spiders.

- A single utterance can make more than one **illocutionary contribution** to the discourse. [Lascarides 06]

Example of SDRS (1)

- Let's consider this example again:

  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 associated SDRS is defined as $\langle A, F, LAST \rangle$, where:
  $A = \{ \pi_0 , \pi_1 , \pi_2 , \pi_3 , \pi_4 , \pi_5 , \pi_6 , \pi_7 \}$
  $F(\pi_1) = K_{\pi_1}$, $F(\pi_2) = K_{\pi_2}$, $F(\pi_3) = K_{\pi_3}$,
  $F(\pi_4) = K_{\pi_4}$, $F(\pi_5) = K_{\pi_5}$,
  $F(\pi_0) = Elaboration(\pi_1 , \pi_6)$
  $F(\pi_6) = Narration(\pi_2 , \pi_5) \wedge Elaboration(\pi_2 , \pi_7)$
  $F(\pi_7) = Narration(\pi_3 , \pi_4)$
  $LAST = \pi_5$

Example of SDRS (2)

- Or graphically:

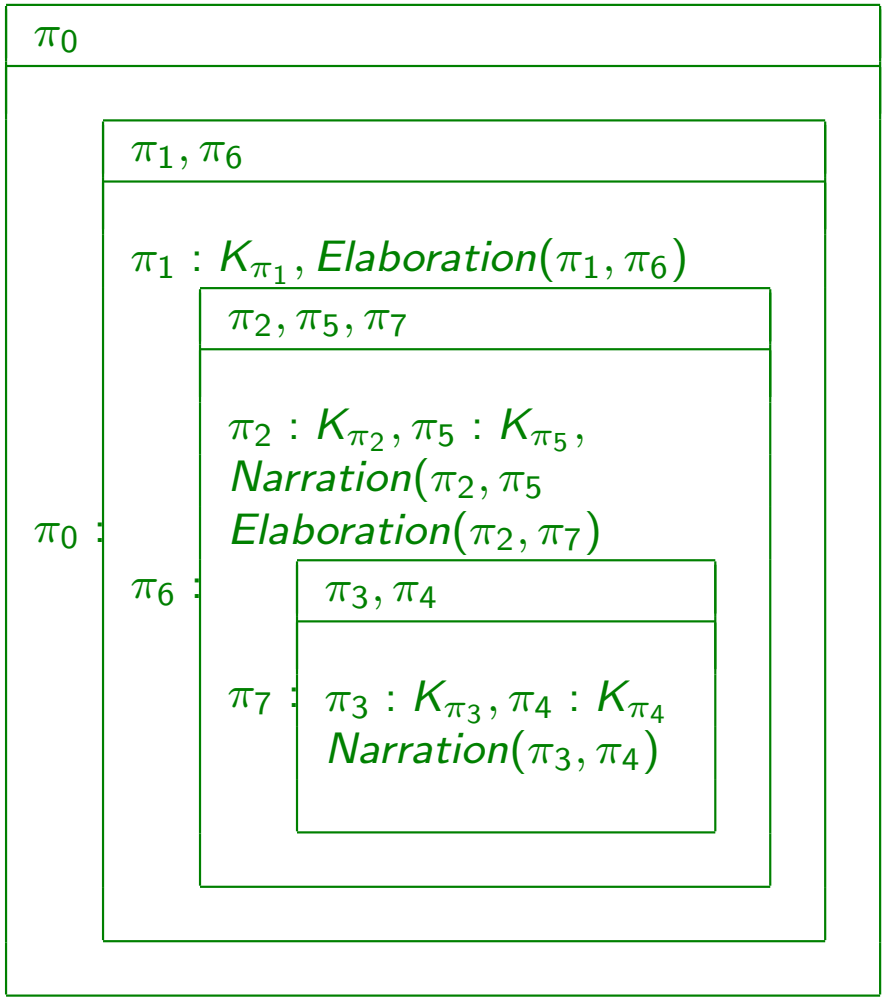

Example of SDRS (3)

- Or even [Lascarides 06]:

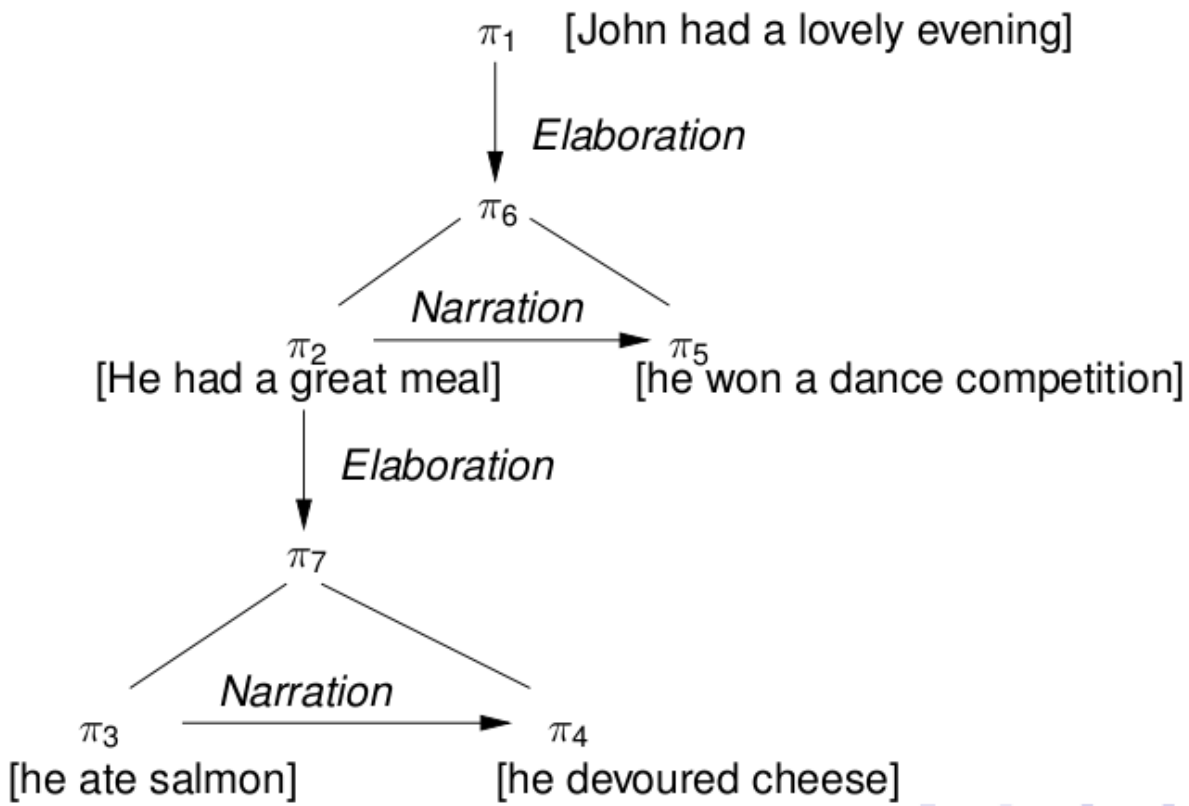

Availability: the right frontier

- Given a SDRS and a new SDRS-formula to insert, we need to know "'where"' this formula can be attached
- The right frontier constraint enables us to restrict the potential places where a formula can be attached
- Formaly, the right frontier constraint is expressed as follows: New information β can attach to:
  1. The label α = LAST ;
  2. Any label λ such that:
     1. $Succ(\lambda, \alpha)$; or
     2. $F(l) = R(\lambda, \alpha)$ for some label l, where R is a subordinating discourse relation (Elaboration, Explanation, or $\Downarrow$). We gloss this as $\alpha < \lambda$
  3. Transitive Closure: Any label λ that dominates α through a sequence of labels λ_1 , . . . , λ_n such that $\alpha < \lambda_1$, $\lambda_1 < \lambda_2$, . . . , $\lambda_n < \lambda$.

Semantics of SDRS (1)

- We now need to assign a semantics to rhetorical relations
- For the sake of simplicity, we will restrict attention here to rhetorical relations that can be assigned an extensional semantics
- *Satisfaction Schema for Veridical Relations*:
  $f[R(\pi_1, \pi_2)]_M g$ iff $f[K_{\pi_1}]_M \circ [K_{\pi_2}]_M \circ [\phi_{R(\pi_1,\pi_2)}]_M g$
- **Veridical**: *Explanation, Elaboration, Background, Contrast, Parallel, Narration, Result, Evidence*
- **Non-veridical**: *Alternation, Consequence*
- **Divergent**: *Correction, Counterevidence*

Semantics of SDRS (2)

- $\phi_{R(\pi_1,\pi_2)}$ expresses the semantic constraints pertinent to the particular rhetorical connection $R(\pi_1, \pi_2)$.
- How to define it? We have to specify **meaning postulates** (or axioms).
- Axiom for Narration:

  $\phi_{Narration(\pi_1,\pi_2)} \Rightarrow e_{\pi_1} < e_{\pi_2}$
  Example: John fell. Mary helped him up.
- Axiom for Explanation:

  $\phi_{Explanation(\pi_1,\pi_2)} \Rightarrow e_{\pi_1} \not< e_{\pi_2}$
  Example: John fell. Mary pushed him.

Constructing logical forms (1)

- We have now introduced the language of SDRSs and their dynamic semantic interpretation.
- The question now arises as to how one **constructs** these logical forms for discourse.
- SDRT distinguishes between the SDRSs themselves (expressed in a so-called "'*logic of information content*"') and a language in which we describe them (the "'*glue logic*"').
- We have already seen the logic of information content in the previous section, we now turn to the glue logic, used to incrementally **build** logical forms.

Constructing logical forms (2)

- The grammar produces only partial (or **underspecified**) descriptions of logical forms.
- Why? They must confront many ambiguities: semantic scope ambiguities, anaphora of various kinds such as pronouns and presuppositions, lexical ambiguities, etc.
- The glue logic performs the following co-dependent inferences:

  1. Infer (preferred) values of underspecified conditions generated by the grammar;
  2. Infer what's rhetorically connected to what;
  3. Infer the values of the rhetorical relations

- The glue logic of SDRS is based on **nonmonotonic** reasoning (more on this in the next section).

Infer rhetorical relations

- Rhetorical Relations aren't always linguistically marked.
- They depend on:
  1. Compositional and lexical Semantics
  2. World Knowledge
  3. Cognitive states
- We need to:
  1. Encode knowledge used to infer rhetorical relations.
  2. Use a logic that supports the inferences we need.

Defeasible reasoning (1)

- The presence of clue words sometimes suffices to compute the appropriate discourse relation, but not always.
- Often, we must also exploit information about the semantic content of the constituents, pragmatic principles and domain
- But, even with all these information sources, we are still for the most part making defeasible inferences as to what discourse relation the author intended.
- Thus the underlying logic for this computation must be a nonmonotonic logic.
- Defeasible reasoning system used for SDRT: *Commonsense reasoning* [Daver 95].

Defeasible reasoning (2)

- $A > B$ is used to denote "'If A then normally B"'. The nonmonotonic validity $\mid\!\sim$, supports intuitive patterns of commonsense reasoning.
- **Defeasible Modus Ponens**: $A > B, A \mid\!\sim B$

If Tweety is a bird,
then normally
Tweety flies
Tweety is a bird

Tweety flies

Knowledge conflict

- **Penguin Principle**:
  If $C \vdash A$ then
  $A > B, B, C > \neg B, C \mid\!\sim \neg B$

If Tweety is a penguin,
then Tweety is a bird
If Tweety is a bird,
then normally Tweety flies
If Tweety is a penguin,
then normally
Tweety doesn't fly
Tweety is a Penguin

Tweety doesn't fly

Discourse update with glue logic axioms

- The glue logic axioms that we just introduced will be used to infer which rhetorical relations to use when inserting a new constituent
- To this end, we specify a set of axioms. We note these axioms as follows:
  $(?(\alpha, \beta, \lambda) \wedge \text{Info}(\alpha, \beta, \lambda)) > \text{R}(\alpha, \beta, \lambda)$
- *Human translation*: if β is to be attached to α with a rhetorical relation and the result is labelled λ, and information $\text{Info}(\alpha, \beta, \lambda)$ about α, β and λ holds, then normally, the rhetorical connection is R.
- $\text{Info}(\alpha, \beta, \lambda)$ expresses information retrieved from rich knowledge sources (world knowledge, cognitive states, linguistic ressources, etc.)

Example of glue logic axioms

- **Narration**:
  $(?(\alpha, \beta, \lambda) \wedge \text{occasion}(\alpha, \beta)) > \text{Narration}(\alpha, \beta)$
- **Scripts for Occasion** :
  $(?(\alpha, \beta, \lambda) \wedge \phi(\alpha) \wedge \psi(\beta)) > \text{occasion}(\alpha, \beta)$.
- **Explanation** :
  $(?(\alpha, \beta, \lambda) \wedge \text{causeD } (\beta, \alpha)) > \text{Explanation}(\alpha, \beta)$
- **Causation and Change** :
  $(\text{change}(e_\alpha \text{ , y }) \wedge \text{cause-change-force}(e_\beta \text{ , x, y }))$
  $\rightarrow \text{causeD } (\beta, \alpha)$

  [Lascarides 06]

Summary

- Many discourse phenomena cannot be analysed without taking **discourse structure** into account
- SDRT does precisely that: it combines **dynamic semantics** with a discourse structure defined via **rhetorical relations** between segments
- It has a well-defined **syntax** and model-theoretic **semantics**
- In order to construct logical forms, SDRT employ **defeasible axioms** specified via a glue logic.

Thanks for your attention!
Questions, comments?

Bibliography I


Nicholas Asher & Alex Lascarides.
Logics of conversation.
Cambridge University Press, 2003.


Joan Busquets, Laure Vieu & Nicholas Asher.
La SDRT : une approche de la cohérence du discours dans la tradition de la sémantique dynamique.
Verbum, vol. XXIII, no. 1, pages 73–101, 2001.


Nathalie Daver.
An Implementation of Segmented Discourse Representation Theory.
In CITE-95, 1995.

Bibliography II


Barbara J. Grosz & Candace L. Sidner.
Attention, Intentions, and the Structure of Discourse.
Computational Linguistics, vol. 12, no. 3, pages 175–204, 1986.


H. Kamp & U. Reyle.
From discourse to logic: Introduction to model-theoretic semantics of natural language, formal logic and discourse representation theory (studies in linguistics and philosophy).
Springer, 1993.


Alex Lascarides.
Semantics and Pragmatics of NLP course, 2006.

Bibliography III


Alex Lascarides & Nicholas Asher.
Segmented Discourse Representation Theory: Dynamic Semantics with Discourse Structure.
In H. Bunt & R. Muskens, editeurs, Computing Meaning: Volume 3. Kluwer Academic Publishers, 2007.


R. Montague.
The Proper Treatment of Quantification in Ordinary English.
In J. Kulas, J. H. Fetzer & T. L. Rankin, editeurs, Philosophy, Language, and Artificial Intelligence: Resources for Processing Natural Language, pages 141–162. Kluwer, Boston, 1988.


Caroline Sporleder.
Computational Models of Discourse Seminar, 2007.