

Empirical Evidence for Embodied Semantics

Petra Hendriks

Center for Language and Cognition Groningen, University of Groningen, P.O. Box 716,
9700 AS Groningen, The Netherlands

P.Hendriks@rug.nl

Abstract. This paper addresses the question whether and under which conditions hearers take into account the perspective of the speaker, and vice versa. Empirical evidence from computational modeling, psycholinguistic experimentation and corpus research suggests that a distinction should be made between speaker meanings and hearer meanings. Literal sentence meanings result from the hearer's failure to calculate the speaker meaning in situations where the hearer's selected meaning and the speaker meaning differ. Similarly, non-recoverable forms result from the speaker's failure to calculate the hearer meaning in situations where the speaker's intended meaning and the hearer meaning differ.

Keywords: Bidirectional Optimality Theory, Embodied Semantics, Perspective Taking, Processing Efficiency, Pronouns, Word Order.

1 Introduction

If we were to interpret all sentences literally, we would frequently misunderstand others. We wouldn't understand metaphors such as *The car died on me*, we would have trouble responding appropriately to indirect speech acts like *Can you tell me the time?*, and we would fail to understand the implicated meaning 'Not all people like soccer' for the utterance *Some people like soccer*. Fortunately, many hearers are quite capable of going beyond the literal meaning of these utterances to grasp the meaning that was intended by the speaker. However, despite hearers' remarkable ability to avoid misunderstanding, how hearers arrive at the intended meaning is still the subject of a lively debate.

Traditionally, a sharp distinction is made between sentence meaning (i.e., the literal meaning of the sentence) and speaker meaning (i.e., what the speaker intended to communicate) (see, e.g., [1] for discussion). Sentence meaning is assumed to be explained by a theory of grammar, whereas speaker meaning is assumed to be explained by a theory of pragmatics. It is thus believed that semantics and pragmatics are distinct domains, with the only uncertainty being where exactly the distinction should be drawn. Contrasting with this traditional view on meaning, this paper argues in favor of *embodied semantics*, the view that meaning does not exist independently of speakers and hearers. Consequently, the relevant distinction is argued to be between speaker meanings and hearer meanings. In this paper, empirical evidence of

various sorts will be provided to support this alternative view. The central claim is that hearers always aim at calculating the speaker meaning. However, if they fail to do so, perhaps because they have insufficient processing resources or cognitive abilities, they may assign an intermediate, for example literal, meaning instead. Similarly, speakers are argued to always aim at calculating the hearer meaning. This guarantees that the produced sentence conveys the intended meaning. If speakers fail to do so, they may produce a non-recoverable form instead.

A distinction between speaker meanings and hearer meanings presupposes a linguistic theory that distinguishes the speaker's perspective from the hearer's perspective. The next section introduces different approaches to perspective taking in semantics and pragmatics. Section 3 considers the question whether and under which conditions hearers calculate the speaker meaning. This question is addressed on the basis of experimental investigations of the pronoun interpretation problem in language acquisition. Section 4 considers the inverse question and asks whether speakers calculate the hearer meaning. This possibility is investigated by looking at semantic factors determining word order in Dutch.

2 Perspective Taking in Semantics and Pragmatics

In his influential William James lectures at Harvard in 1967, Grice [2] proposed that speakers are guided by a Cooperative Principle, backed by a set of Maxims of Conversation that specify speakers' proper conduct. For example, the Maxim of Relation tells speakers to be relevant, and the Maxim of Quantity tells speakers to make their contribution as informative as is required for the purposes of the exchange, but not more informative than that. By choosing a particular form to express their intentions, speakers assume that hearers will be able to infer the intended meaning on the basis of this form. Grice formulates this as follows: "[Speaker] meant something by x' is (roughly) equivalent to '[Speaker] intended the utterance of x to produce some effect in an audience by means of the recognition of this intention.'" (p. 220).

Several later studies have sought to reduce Grice's maxims, while maintaining the division of labor between speakers and hearers in the sense that speakers choose the sentence to be uttered, while hearers must do a certain amount of inferencing to determine the speaker's intended meaning. However, given Grice's formulation of the Maxim of Quantity, speakers also have to do some inferencing, as they have to determine how much information is required for the purposes of the exchange. Are the inferences that speakers draw of the same sort as the inferences that hearers draw, or are they fundamentally different?

A fully symmetric account of conversational inference, according to which hearers and speakers make similar inferences about the effects of their choices, has been proposed within the framework of optimality theory (OT) [3]. According to Blutner's definition of bidirectional optimality theory (biOT) [4], speakers select the best form for a given meaning, thereby taking into account the hearer's perspective, and hearers select the best meaning for a given form, thereby taking into account the speaker's perspective. Contrasting with Blutner's symmetric conception of bidirectional optimization, various asymmetric models have been proposed within OT. For

example, Zeevat proposes an asymmetric model according to which hearers take into account the speaker's perspective, while speakers do not take into account the hearer's perspective to the same degree [5]. A similar position is adopted by Franke in his game theoretic model of conversational inference [6]. Jäger, on the other hand, develops a bidirectional learning algorithm in which speakers take into account hearers when evaluating form-meaning pairs, but not vice versa [7]. These different positions are mainly based on theoretical arguments and have not been tested by looking at the actual processes of speaking and understanding. Therefore, a relevant question is whether it is possible to find empirical evidence for the symmetry or asymmetry of conversational inference by considering how actual hearers and speakers comprehend and generate sentences.

A second question is whether the proposed conversational inferences are automatic word-by-word interpretational processes (as is believed to be the case for grammatical processes) or additional end-of-sentence processes (as is assumed by some to be true for pragmatic processes). This question is independent of the symmetry or asymmetry of perspective taking and conversational inference, but is relevant in relation to the distinction between semantics and pragmatics. According to Blutner and Zeevat, (weak) bidirectional optimization, and in fact the whole domain of pragmatics, should be seen as reflecting offline interpretation mechanisms [8]. In their view, perspective taking through bidirectional optimization only occurs at the end of the utterance. An alternative view on bidirectional optimization is that it is a grammatical mechanism that is applied in an online fashion during incremental sentence interpretation.

The remainder of this paper aims to shed new light on these two issues by discussing empirical evidence from computational modeling, psycholinguistic experimentation, and corpus research. Section 3 considers a phenomenon that has been argued to require hearers to take into account the speaker's perspective, and addresses the question whether this conversational inference is a local and online interpretational process, or a global and offline process. Whether speakers also take into account hearers is the topic of Section 4.

3 Speaker Effects on the Hearer

A well-studied phenomenon in language acquisition is the interpretation of pronouns and reflexives. Many studies have found that children make errors interpreting pronouns in sentence sequences such as *This is Mama Bear and this is Goldilocks. Mama Bear is washing her* until the age of five or six (see, e.g., [9]). This contrasts with children's interpretation of reflexives, which is adult-like from the age of four onward. Many explanations of children's pronoun interpretation delay appeal to non-syntactic factors, such as children's inability to compare the various interpretational possibilities for pronouns (see [10] for an influential approach).

3.1 A Bidirectional Account of Pronoun Interpretation

In [11], an explanation is proposed of children's pronoun interpretation delay in terms of biOT. Whereas the distribution of reflexives is subject to Principle A, which

requires reflexives to corefer with the local subject, it is argued that pronouns are *not* subject to a complementary Principle B which forbids pronouns to corefer with the local subject. Rather, pronouns are essentially free in their interpretation. As a consequence, children allow both a coreferential and a disjoint interpretation for pronouns. This would explain children's guessing behavior with pronouns in experimental tasks. In contrast to children, adults are argued to optimize bidirectionally ([11], cf. [12]) and hence block the coreferential meaning for the pronoun. Adults reason that a speaker, due to a weaker constraint preferring reflexives to pronouns, would have used a reflexive to express a coreferential meaning. As a consequence, for a hearer the coreferential meaning is blocked for the pronoun. This leaves only the disjoint meaning as the meaning of the pronoun.

The biOT explanation of children's errors in pronoun interpretation proposed in [11] predicts children's production of pronouns, in contrast to their comprehension, to be adult-like. If Principle A is stronger than the preference for reflexives over pronouns, a disjoint meaning is expressed best by a pronoun. Choosing a reflexive to express a disjoint meaning would violate the stronger constraint Principle A. On the other hand, if the meaning to be expressed is a coreferential meaning, the optimal form is a reflexive. In this situation, choosing a reflexive satisfies both constraints, whereas choosing a pronoun would violate the weaker constraint. This remarkable prediction of a guessing pattern in comprehension but correct performance in production was confirmed in a psycholinguistic experiment. Testing comprehension and production of the same type of sentences with pronouns and reflexives in the same children, Spenser et al. [13] found that children who made errors interpreting pronouns performed correctly on pronoun production.

3.2 A Cognitive Model of Pronoun Interpretation

Although the biOT explanation proposed in [11] accounts for children's delay in pronoun interpretation, it is compatible with a local as well as a global view on bidirectional optimization. Children may compare the pronoun to the alternative reflexive form as soon as the pronoun is encountered, or they may wait until the end of the sentence to compare the sentence containing the pronoun with the alternative sentence containing a reflexive.

To test the biOT explanation and to compare it to non-OT accounts of children's delay in pronoun interpretation, the biOT explanation was implemented in the cognitive architecture ACT-R [14][15]. The cognitive architecture ACT-R is both a theory of cognition and a computational modeling environment. The cognitive architecture imposes cognitive constraints on the computational models, based on a wide range of experimental data on information processing, storage and retrieval. By constructing a cognitive model, concrete and testable predictions can be generated regarding children's development and online comprehension of pronouns.

Two aspects of ACT-R are of crucial importance to constructing a cognitive model of pronoun interpretation. First, every operation in ACT-R takes a certain amount of time. Because operations can be executed in parallel if they belong to different modules of the architecture, the total time that is necessary to perform a cognitive process is not simply the sum of the durations of all constituting operations. Rather,

the total time critically depends on the timing of the serial operations within a module, and how the various modules interact. To generate predictions about the timing of cognitive processes, computational simulation models can be constructed and run. A second aspect of ACT-R that is essential to constructing a cognitive model of pronoun interpretation is that higher processing efficiency can be obtained through the mechanism of production compilation. If two cognitive operations are repeatedly executed in sequence, production compilation integrates these two operations into one new operation. This new operation will be faster than the two old operations together. This process of production compilation can continue until the cognitive process has been integrated into a single operation. As a consequence of production compilation, cognitive processes become faster with experience.

Bidirectional optimization combines the speaker's direction of optimization with the hearer's direction of optimization. In the cognitive model, bidirectional optimization is therefore implemented as two serial processes of unidirectional optimization:

$$(1) f \rightarrow m \rightarrow f'$$

Interpreting a pronoun thus consists of a first step of interpretation ($f \rightarrow m$), followed by a second step of production ($m \rightarrow f'$), in which the output of the first step (the unidirectionally optimal meaning) is taken as the input. If the output of production f' is identical to the initial input in interpretation f , a bidirectionally optimal pair results. If the output of the production step is different, the unidirectionally optimal meaning m must be discarded and another meaning m' must be selected in the first optimization step. Because pronouns are ambiguous according to the biOT explanation discussed in Section 3.1, discarding the coreferential meaning results in selection of the disjoint meaning. The interpretation process in (1) formalizes the assumption that hearers take into account the choices of the speaker.

If unidirectional optimization needs a given amount of time, the serial version of bidirectional optimization in (1) will initially need about twice this amount of time. When time for interpretation is limited, the model will initially fail to complete the process of bidirectional optimization. So at first, the output of the model will be a unidirectionally optimal meaning rather than a bidirectionally optimal meaning. However, over time the model's performance will become more and more efficient as a result of the mechanism of production compilation. As soon as processing efficiency is high enough to perform bidirectional optimization within the given amount of time, the model will do so, resulting in a bidirectionally optimal meaning as the output. As production compilation results from the repeated sequential execution of particular operations, such as retrieval of particular lexical items from declarative memory, it is dependent on the frequency of these lexical items in the language spoken to the child. As a consequence, the speed of development of bidirectional optimization is different for different lexical items.

Simulations of the cognitive model show a pattern of interpretation that is similar to the pattern displayed by English- and Dutch-speaking children [14]. Already from the beginning of the simulated learning period, when the constraints are already in place but bidirectional optimization is not mastered yet, the interpretation of reflexives is correct. In contrast, the proportion of correct interpretations for pronouns

hovers around 50% during the first half of the simulated learning period, and then gradually increases to correct performance. The model's correct performance on reflexives is not surprising because unidirectional and bidirectional optimization both yield the correct meaning. The model's performance on pronouns follows from the gradual increase in processing efficiency, as a result of which bidirectional optimization can be performed more frequently.

3.3 Testing the Cognitive Model Experimentally

In incremental interpretation, time limitations arise from the speed at which the next word of the sentence arrives. In the previous section it was argued that children need less time for interpretation if their processing has become more efficient. If bidirectional optimization is a local process which takes place as soon as the pronoun is encountered, there is a second way to facilitate bidirectional optimization: by slowing down the speech rate, so that it takes longer for the next word to arrive. This prediction was tested in a study with 4- to 6-year-old Dutch children, who were presented with sentence sequences such as *Look, a penguin and a sheep are on the sidewalk. The penguin is hitting him with a pan* (translated from Dutch) at a normal speech rate as well as at a speech rate that was artificially slowed down [15]. Crucially, the pronoun does not appear at the end of the sentence but is always followed by a prepositional phrase. It was found that slower speech improved children's performance with pronouns but not with reflexives, and only improved children's performance with pronouns if they made errors with pronouns at normal speech rate. In all other situations, slowing down the speech rate had a negative effect. Because children who make errors in pronoun interpretation succeed in arriving at the correct interpretation when they are given more time, the experimental results suggest that insufficient processing speed is the limiting factor in children's comprehension of pronouns. If this is true, children's interpretation is already aimed at computing the speaker's meaning before they have acquired sufficient processing speed to actually do so.

Apparently, taking into account the speaker as a hearer requires sufficient processing efficiency. If, initially, children's processing is too slow, they may fail to optimize bidirectionally, and select a unidirectionally optimal meaning instead. With experience in pronoun interpretation, children's processing of pronouns becomes more efficient until correct performance is reached. These results suggest that bidirectional interpretation of pronouns must be viewed as a local rather than a global process, since slowing down the speech rate gave the child participants in the experiment more time within the sentence (immediately after they heard the pronoun), while they still had the same amount of time at the end of the sentence. Thus, these results seem incompatible with approaches to bidirectional optimization and perspective taking advocating a purely global view (e.g., [8]). Also, the results provide a challenge to alternative accounts of children's pronoun interpretation delay that attribute children's errors with pronouns to their lack of pragmatic knowledge, limitations in perspective taking, or task effects (see [15] for discussion). The results even seem to undermine the processing explanation proposed by Reinhart [10], who attributes children's errors to their insufficient working memory capacity to perform a

global comparison operation, as child hearers compare the pronoun to the alternative reflexive form as soon as the pronoun is encountered.

4 Hearer Effects on the Speaker

In the previous section, empirical evidence was presented for the view that hearers take into account the speaker's perspective to arrive at the intended meaning for object pronouns. If conversational inference is fully symmetric, we expect speakers to also take into account the hearer's perspective, perhaps in the following way:

$$(2) m \rightarrow f \rightarrow m'$$

According to (2), producing a form f consists of a first step of production ($m \rightarrow f$), followed by a second step of interpretation ($f \rightarrow m'$), in which it is checked whether the initial meaning m is recoverable on the basis of form f . In the next section, the issue is addressed whether sentence generation actually proceeds in this manner.

4.1 Constituent Fronting in Dutch

Word order in Dutch is characterized by the fact that in declarative main clauses the finite verb must occur in second position. In addition, however, Dutch allows for a moderate amount of word order variation with respect to what can appear in front of this finite verb. Although the first position of the sentence is most frequently (in roughly 70% of cases, according to an estimation [16]) occupied by the subject, this position can also be occupied by direct objects, indirect objects and other constituents.

In a large scale corpus study, Bouma [16] investigated the factors determining what constituent comes first in a Dutch main clause. To this end, Bouma conducted a logistic regression analysis of data from the spoken Dutch corpus *Corpus Gesproken Nederlands* (CGN). The factors grammatical function, definiteness and grammatical complexity were found to independently influence the choice of constituent in first position. Regarding grammatical function, subjects have the strongest tendency to occur in first position, followed by indirect objects and direct objects. Regarding definiteness, definite full NPs are more likely to appear in first position than indefinite full NPs. Although pronouns as a group show a strong tendency to appear in first position, this is only visible in the fronting behavior of demonstrative pronouns, which front more often than definite full NPs. Reduced personal pronouns are strongly discouraged from appearing in first position, perhaps because they express highly predictable material. Finally, more complex material is preferably placed at the right periphery of the clause, thus resulting in an avoidance of the first position.

4.2 Partial Word Order Freezing

Although speakers of Dutch may place non-subjects in first position under the influence of factors such as the ones mentioned above, in certain situations placing a

non-subject in first position makes it difficult for the hearer to infer the intended meaning. If a hearer encounters a sentence such as *Fitz zag Ella* ('Fitz saw Ella'), he can in principle assign an SVO interpretation or an OVS interpretation to this sentence, as both word orders are possible in Dutch. Under the first interpretation, Fitz is the subject. Under the second interpretation, Fitz is the object. However, presented out of context and in the absence of any intonational clues, most hearers will interpret this sentence as conveying an SVO interpretation. Their preferred interpretation thus reflects the observation that the first constituent most likely is the subject. This observation about hearers' preference may have consequences for speakers' freedom of word order variation. If the speaker wishes to convey the meaning that Ella did the seeing, the sentence *Fitz zag Ella* is a poor choice because hearers will have a preference for Fitz as the subject.

This type of conversational inference is implicit in the biOT model of word order variation proposed by Bouma [16] (cf. [17]). In this model, the speaker's choice for a particular word order is influenced by the hearer's ability to recover the subject and object. If speakers take into account the perspective of the hearer, they are expected to limit the freedom of word order variation in situations such as the one sketched above, where subject and object can only be distinguished on the basis of word order. On the other hand, if other clues are present that allow the hearer to distinguish the subject from the object, speakers are expected to have more freedom of word order variation. Such clues may include definiteness. Subjects tend to be highly definite, whereas direct objects tend to be indefinite. Indeed, Bouma's analysis of the transitive sentences in the CGN confirmed the prediction that a non-canonical word order occurs more frequently in sentences with a definite subject and an indefinite object [16]. A preliminary analysis of a manually annotated subset of the CGN suggests that animacy may have a similar effect, as a non-canonical word order occurs more frequently in sentences with an animate subject and an inanimate object [16]. These hearer effects on the speaker's choice of word order were found on top of the factors discussed in Section 4.1. So the possibility of word order variation increases if subject and object can be distinguished on the basis of other clues than word order. Speakers limit word order variation in situations where a non-canonical word order would make it more difficult for the hearer to recover the intended meaning.

Bouma's corpus study thus provides evidence for a tendency toward partial freezing of word order variation in spoken Dutch discourse, parallel to the observation of partial blocking in the domain of interpretation. In the previous section, we saw that hearers restrict the interpretational possibilities of pronouns in situations where a better form is available to the speaker for expressing one of the meanings. The corpus study provides evidence for the assumption that speakers take into account the hearer's perspective, and limit word order variation in situations that would result in unrecoverability of the expressed meaning.

5 Embodied Semantics

Section 3 addressed the question whether hearers take into account the speaker's perspective in interpretation. A biOT account of conversational inference in pronoun

interpretation, according to which hearers also consider alternative forms the speaker could have used but did not use, was shown to be supported by results from cognitive modeling and psycholinguistic experimentation. Section 4 addressed the inverse question whether speakers take into account the hearer's perspective when producing a sentence. Bouma's corpus study of word order in Dutch seems to provide evidence that speakers consider how hearers will interpret potential forms. Empirical evidence of various sorts thus suggests that hearers take into account speakers, and vice versa. Hence, the empirical evidence discussed in Section 3 and Section 4 supports a symmetric conception of biOT. Moreover, hearers appear to consider the speaker's choices locally, as soon as they encounter the relevant linguistic form, and do not wait until the end of the utterance.

This paper argued for a distinction between speaker meaning and hearer meaning. Hearers select an initial meaning for the utterance they hear on the basis of the constraints of the grammar. This hearer meaning usually is the literal meaning of the utterance (although other interpretations are possible if the hearer still has a non-adult constraint ranking, or under strong contextual pressure). Hearers then go on to compute the speaker meaning (i.e., the intended meaning) on the basis of the hearer meaning and the constraints of the grammar. These two steps can be formally modeled by bidirectional optimization. In some situations the speaker meaning differs from the hearer meaning. If the hearer is not (yet) able to optimize bidirectionally, he will assign a non-intended meaning to the utterance in this situation. Note that the speaker meaning not necessarily is the meaning that is actually intended by the speaker. Rather, it is the meaning that the hearer assumes is intended by the speaker by considering the speaker's perspective. The same two meanings, speaker meaning and hearer meaning, also play a role in production, with speakers aiming to compute the hearer meaning but sometimes failing to do so. Given this distinction between speaker meaning and hearer meaning, there is no need to distinguish a separate sentence meaning. In fact, under the proposed view sentences do not have meanings by themselves. Sentences have meanings only in so far as these meanings are assigned to them by speakers and hearers. This view of semantics as embodied in speakers and hearers and their tasks of speaking and understanding is a departure from traditional thinking about meaning.

If no distinction is made between sentence meaning and speaker meaning, it becomes difficult to distinguish semantics and pragmatics. The difference between assigning a literal meaning to a sentence and assigning a speaker meaning to this sentence is argued to lie in the hearer's processing efficiency. As acquiring higher processing efficiency is a gradual process, the distinction between semantics and pragmatics (if there is any) must also be gradual. The traditional distinction between semantics and pragmatics is blurred even more by the fact that this paper addressed two phenomena that are not immediately associated with conversational inference, namely pronoun interpretation and constituent fronting. Nevertheless, evidence was presented that supports analyses of these phenomena in terms of conversational inference. Cognitive modeling of the development of pronoun interpretation illustrates that it is possible for these processes of conversational inference to become automatic in such a way that their output cannot be distinguished from the output of regular grammatical processes. This suggests that at least some conversational inferences that start out as slow and effortful processes can become automatic and

indefeasible over time, not only in the course of diachronic language change (cf. [8]) but also in synchronic language development.

Acknowledgments. This investigation is supported by a grant from the Netherlands Organization for Scientific Research (NWO grant no. 277-70-005). I thank the audience of the Amsterdam Colloquium for their useful comments.

References

1. Levinson, S.: *Presumptive Meanings: The Theory of Generalized Conversational Implicatures*. MIT Press, Cambridge, MA (2000)
2. Grice, P.H.: *Studies in the Ways of Words*. Harvard University Press, Cambridge, MA (1989)
3. Prince, A., Smolensky, P.: *Optimality Theory: Constraint Interaction in Generative Grammar*. Blackwell, Malden, MA (2004)
4. Blutner, R.: Some Aspects of Optimality in Natural Language Interpretation. *Journal of Semantics* 17, 189--216 (2000)
5. Zeevat, H.: The Asymmetry of Optimality Theoretic Syntax and Semantics. *Journal of Semantics* 17, 243--262 (2000)
6. Franke, M.: *Signal to Act: Game Theory in Pragmatics*. PhD Thesis, University of Amsterdam (2009)
7. Jäger, G.: Learning Constraint Sub-hierarchies: The Bidirectional Gradual Learning Algorithm. In Blutner, R., Zeevat, H. (eds.) *Optimality Theory and Pragmatics*, pp. 251--287. Palgrave Macmillan, Hampshire (2004)
8. Blutner, R., Zeevat, H.: Optimality-Theoretic Pragmatics. In Maienborn, C., von Stechow, K., Portner, P. (eds.) *Semantics: An International Handbook of Natural Language Meaning*. Mouton de Gruyter, Berlin (to appear)
9. Chien, Y.C., Wexler, K.: Children's Knowledge of Locality Conditions on Binding as Evidence for the Modularity of Syntax and Pragmatics. *Language Acquisition* 1, 225--295 (1990)
10. Reinhart, T.: *Interface Strategies: Optimal and Costly Computations*. MIT Press, Cambridge, MA (2006)
11. Hendriks, P., Spender, J.: When Production Precedes Comprehension: An Optimization Approach to the Acquisition of Pronouns. *Language Acquisition* 13, 319--348 (2005/6)
12. De Hoop, H., Krämer, I.: Children's optimal interpretations of indefinite subjects and objects. *Language Acquisition* 13, 103--123 (2005/6)
13. Spender, J., Smits, E.J., Hendriks, P.: Coherent Discourse Solves the Pronoun Interpretation Problem. *Journal of Child Language* 36, 23--52 (2009)
14. Hendriks, P., van Rijn, H., Valkenier, B.: Learning to Reason about Speakers' Alternatives in Sentence Comprehension: A Computational Account. *Lingua* 117, 1879--1896 (2007)
15. Van Rij, J., van Rijn, H., Hendriks, P.: Cognitive Architectures and Language Acquisition: A Case Study in Pronoun Comprehension. *Journal of Child Language* (in press)
16. Bouma, G.: *Starting a Sentence in Dutch: A Corpus Study of Subject- and Object-Fronting*. PhD Thesis, University of Groningen (2008)
17. Lee, H.: *Optimization in Argument Expression and Interpretation: A Unified Approach*. PhD Thesis, Stanford University (2001)