

On compositionality and bidirectional optimization*

Helen de Hoop¹, Petra Hendriks², Reinhard Blutner³

¹ *Centre for Language Studies (CLS), Radboud University Nijmegen,
H.deHoop@let.ru.nl*

² *Center for Language and Cognition Groningen (CLCG), University of Groningen,
P.Hendriks@rug.nl*

³ *Institute for Language, Logic and Computation (ILLC), University of Amsterdam,
K.R.Blutner@uva.nl*

Abstract

In this paper we revisit the semantic principle of compositionality and argue that compositionality is bidirectional optimization. Underspecification approaches to natural language interpretation generally start with an underspecified or weak meaning, which is strengthened by contextual information. By contrast, the bidirectional optimization approach we advocate proceeds from the strongest possible meaning. This meaning can be changed or weakened by contextual information. Under this approach, the meaning of an utterance is composed in a functional rather than a concatenative way, while contextual sources of information play a major role. Yet, because the context of any utterance is in principle the same for the speaker and the hearer, composition and decomposition proceed hand in hand. Hence, bidirectional optimization ultimately guarantees (functional) compositionality.

The Principle of Compositionality

One of the key principles in formal approaches to natural language interpretation is the principle of compositionality, which expresses the idea that the meaning of a complex expression can be derived from the meanings of its parts and the way these parts are syntactically linked. That is, a syntactic structure can be mapped onto a semantic structure in which the meanings of the elements of that structure are inserted. This principle of compositionality guarantees that syntax and semantics go hand in hand as each syntactic combinatorial rule is associated with a semantic one. For example, the meaning of a determiner is defined as a two-place relation between two sets of individuals. In the sentence *Most cats are black*, the set of cats (C) is thus related to the set of black individuals (B) by the determiner *most*. The sentence *Most cats are black* is true if and only if the cardinality of the intersection of these two sets exceeds the cardinality of their difference. That is, the sentence is true if there are more cats which are black than cats which are not black. In a formula: *most CB* is true iff $|C \cap B| > |C - B|$. This semantic rule can be generalized to all sentences of the form *Most AB* where A and B represent the two sets of individuals related by the determiner *most*. Where do these sets of individuals come from? In accordance with the semantic principle of compositionality, they are assumed to come with the syntactic structure of these sentences. In sentences of the form *Most NP VP* (where NP stands for ‘noun phrase’ and VP for ‘verb phrase’, the predicate), the NP refers to set A (also called the domain of quantification) and the VP to set

* We thank three anonymous reviewers and Chungmin Lee for their stimulating comments and questions. Helen de Hoop and Petra Hendriks gratefully acknowledge the Netherlands Organisation for Scientific Research (NWO) for financial support of their research.

B. The sentence *Most black cats have some sort of white spot on them* is syntactically decomposed as *Most [black cats]_{NP} [have some sort of white spot on them]_{VP}* and thus set A is the set of black cats, while set B is the set of individuals that have some sort of white spot on them, and semantically the sentence is true if and only if there are more black cats which have some sort of white spot on them (set $A \cap B$) than black cats which do not have some sort of white spot on them (set $A - B$).

Other Views on Compositionality

Smolensky (1991) notes that “It would constitute significant progress to be able to reduce the (symbolic) principle of semantic compositionality to more basic connectionist principles (...). Developing such a connectionist semantics might well involve formalizing the weak notion of compositionality.” With *weak compositionality* Smolensky refers to compositionality in an “approximate” sense: a non-concatenative way of combining contextually dependent (representations of) elements of a compound expression. Smolensky’s (1988) discussion of the (in)famous “cup with coffee”-example may serve to illustrate the point. Note that if the meaning of “cup with coffee” would have been built up compositionally, then we may expect in a connectionist model that the activity vector representing “cup with coffee” would be composed of the vectors representing “coffee” and “cup”. However, as Smolensky points out, the pattern representing “coffee” in the context of “cup” is quite different from the pattern representing “coffee” in the context of “can”. Roughly, the distributed semantic representation of “cup with coffee” should involve the activation of features such as ‘brown liquid with curved sides’ or ‘brown liquid contacting porcelain’ which are not ‘part of’ the representation of “cup”. Hence, if we extract the representation of “cup” from the representation of “coffee” then these features would still be activated, and therefore the remaining activity vector does not represent a context-independent concept of “coffee” but crucially “coffee in the context of cup”. Thus, the representation of “coffee” is context-dependent. Its internal structure is influenced by the structure of which it is a part. Smolensky (1991) calls this *weak compositionality*: The activation vector consists of ‘parts’ which influence each other. In that sense, Smolensky’s weak compositionality is reminiscent of Frege’s (1884) contextuality principle, cited in Janssen (1997): “A word has a meaning only in the context of a sentence, not in separation.” Obviously, this latter idea of compositionality conflicts with the thesis of context independency (cf. Hintikka, 1983, cited in Janssen, 1997): “The meaning of an expression should not depend on the context in which it occurs.” It should be noted that Smolensky’s conception closely resembles what many neurobiologists have in mind when they talk of neuronal assemblies. Smolensky, thus, expresses a still very popular view. A potential problem with weak compositionality is that it does not automatically render systematicity (Blutner *et al.*, 2004).

Van Gelder (1990) distinguishes between concatenative and functional compositionality. He describes the essence of a concatenative mode of combination informally as “a way of *linking* or *ordering* successive constituents without altering them in any way as it forms the compound expression.” For example, tokens of the symbol “P” are the same whether appearing standing alone, *P*, or in the context of an expression such as (*P&Q*). Yet, although formal languages of mathematics, logic, and computer science are all compositional in this concatenative sense, concatenation is not the only way of implementing the combination of elements in getting a compound expression. Van Gelder (1990) points out that functional compositionality can be obtained whenever there are general, effective, and reliable processes for (a) producing an expression given its constituents, and (b) decomposing the expression back into those constituents. Whereas concatenative schemes are always functionally compositional as well, it is possible to have merely functionally compositional schemes that are *not* concatenative. One useful example for a coding scheme that is merely

functionally compositional is the Gödel numbering schema for a formal language. A crucial feature of this scheme is that it is completely reversible. In addition, using the prime decomposition scheme it is possible to calculate the Gödel numbers of the (primitive) constituents of the expression under discussion (cf. van Gelder, 1990, p. 362).

There is an empirical problem with the semantic principle of compositionality, as it was introduced above. For example, compare the sentence *Most people sleep at night* to the sentence *Most people drink at night*. The syntactic structures of these two sentences are identical. Following the syntactic-semantic rules given in the introductory section, we define the meanings of these sentences as follows. In both sentences the set of people gives the set A (the domain of quantification) while the rest of the predicate (*sleep at night* and *drink at night*) provides us with set B. This means that *Most people sleep at night* becomes true if and only if there are more people who sleep at night than people who do not. This seems to be the right result. However, this procedure does not give the right result for the syntactically identical sentence *Most people drink at night*. That is, in a situation when only half of the people drink (they are drinkers) and 80% of these drinkers only drinks at night, the sentence *Most people drink at night* will be judged true. Yet, in such a situation it is strictly speaking not true that at night, more than half of the people drink, since only half of the people actually drink. That is, the preferred reading of the latter sentence is rather “Most people who drink, drink at night”. The sentence is true if there are more drinking people who drink at night than drinking people who do not drink at night. But that means that in the sentence *Most people drink at night* the domain of quantification (set A) is not just provided by the NP *people* but is further restricted by the verb which is part of the VP, and hence of set B.

The above two sentences were presented out of context. Because the sentences were presented in written form, no clues were provided with respect to their prosodic structure. In addition, their syntactic structures appear to be completely identical. Although these different readings actually involve different truth conditions, there is no structurally based, mechanical way in which the correct interpretations can be derived. As these sentences seem to suggest, structural information can be overruled by other information sources, presumably our world knowledge, as we will argue below. Therefore, sentences like the two above are considered a problem for the semantic principle of compositionality.

In order to calculate the truth conditions of a quantificational expression, one always has to take into account the context. At this point, consider the sentence *Most drink AT NIGHT*. Here, the capitals indicate the accented part of the sentence. If we want to derive the interpretation of this example sentence compositionally, we must assume the presence of an empty noun. Thus the sentence provides us with an underspecified or weak meaning, which may be strengthened by adding contextual information. This can be modelled by assuming that the content of the empty noun, which denotes the whole domain of individuals, gets intersected with a context set variable (cf. Westerståhl, 1985). But in fact, we need two context set variables then because of the effects of sentential accent, or focus. One context set variable would be equated with the generalized union over the set of alternatives for the syntactic argument that contains the focus (cf. de Hoop & Solà, 1996), so that the quantificational domain would be the set of individuals who drink at certain times. The other context set variable would be equated with some additional context set, for example, the set of linguists in Sydney. Hence, in this case we get as the domain of quantification the set of drinking linguists in Sydney. But how many contextual restrictions can or should we add before we may calculate the truth conditions of a quantificational sentence?

The question arises when, how and to what extent hearers use different guiding principles to arrive at the proper interpretation of a quantified expression in a given context. As we have seen, different readings do in fact involve different truth conditions. Therefore,

we may say that the problem for compositionality just pointed out is also a problem for linguistic theory.

Compositionality and Optimization

The problem of compositionality, as we pointed it out in the above discussion, is stated by Dekker and van Rooij (2000) as follows: “(...) we cannot systematically determine the semantic content of a sentence in a *compositional* way based on its syntactic structure, without making reference to the attitudes of speakers and hearers, if we equate the semantic content of a sentence with its truth-conditions (...). So what should we do? Give up compositionality, or give up the assumption that what should be determined compositionally are the truth-conditions of a sentence?” In fact, we will argue that neither of these assumptions has to be given up *if* we take a broader view on compositionality (van Gelder’s 1990 functional compositionality) in addition to Blutner’s bidirectional optimization view on the relation between form and meaning. But before we explore this hypothesis, let us examine the question what is the current view on compositionality within Optimality Theory.

Optimality Theory was developed in the 1990s by Alan Prince and Paul Smolensky as a general theory of language and grammar (cf. Prince & Smolensky, 1993/2004, 1997). Optimality Theory applied to the domains of syntax and semantics involves two closely related issues: 1. Given a semantic input (a meaning), what is its optimal expression?, 2. Given a syntactic input (a form), what is its optimal interpretation? In OT semantics, developed by Hendriks and de Hoop (1997, 2001), each grammatical expression is associated with an, in principle, infinite number of interpretations. These candidate interpretations are tested against a set of ranked constraints in a parallel fashion. The candidate interpretation that satisfies the constraints best is the optimal interpretation and hence the preferred interpretation for the given expression. Crucially, the constraints of the grammar differ in strength. When two constraints are in conflict, it is better to satisfy the stronger constraint than it is to satisfy the weaker constraint. So the optimal candidate need not satisfy all constraints, but merely must satisfy the constraints of the grammar better than its competitors do. For each input, this optimization procedure provides the optimal output. Note that the input need not be complete or fully grammatical to be provided with an output. This allows for interpretation to proceed incrementally as the sentence is built up.

One of the advantages of such an approach is that constraints of various nature (syntactic, contextual, etc.) interact with each other in a truly cross-modular way. This view crucially differs from the classical compositional approach, where interpretation is computed on the basis of the syntactic input, making use of context only when necessary. In OT, on the other hand, constraints of various nature apply in parallel, and it is conceivable that syntactic constraints are violated in order to satisfy stronger contextual constraints.

Whereas OT syntax optimizes syntactic structure with respect to a semantic input (the so-called speaker’s perspective), OT semantics optimizes interpretation with respect to a syntactic input (the hearer’s perspective). Although it is possible to treat OT syntax and OT semantics as two separate optimization procedures, the one concerned with generation and the other with interpretation, Blutner (2000) extensively argues in favour of an integration of the two perspectives into a simultaneous optimization procedure. In Blutner’s version of bidirectional OT, speaker’s and hearer’s optimization are carried out simultaneously over pairs consisting of a form and a meaning. A form-meaning pair $\langle f, m \rangle$ is called *super-optimal* if and only if there is no other super-optimal pair $\langle f', m' \rangle$ such that $\langle f', m' \rangle \gg \langle f, m \rangle$ (\gg is an ordering relation which can be read as ‘*being more harmonic, being more economical*’) and there is no other super-optimal pair $\langle f, m' \rangle$ such that $\langle f, m' \rangle \gg \langle f, m \rangle$. Under the assumption

that the relation » is transitive and well-founded, Jäger (2002) proved the above to be a sound recursive definition.

With respect to the relation between compositionality and optimization, Zeevat (2000) discusses two relevant constraints, one that prohibits adding material to the content or context of utterance and one that requires us to interpret all that the speaker has said. Satisfaction of these two constraints means interpreting all and only the material available in the utterance and so, their combination “restores important aspects of compositional semantics (not the full principle, but essential aspects)” (Zeevat, 2000). Additionally, in OT syntax, a principle called recoverability relates to compositionality, the idea being that the semantic content of elements that are not pronounced must be recoverable from local context (Pesetsky, 1998; Buchwald *et al.*, 2002; Kennedy, 2002; Vogel, 2004). Kuhn (2001) shows that recoverability effects automatically follow in a (weak) bidirectional optimization model, as in such a system we not only have to check whether a reduced string is the optimal way of expressing the underlying content, but we also have to check whether the underlying content is the optimal interpretation of the reduced string.

Compositionality and Bidirectionality

In Optimality Theory, the procedure that provides us with an optimal interpretation of a given form within a certain context can be viewed in two different ways. The first approach combines the view of radical underspecification with a mechanism of contextual enrichment. This approach is taken, for example, in Blutner (1998, 2000). The second approach takes the opposite position in crucial respects. Rather than strengthening a weak (underspecified) meaning with contextual knowledge, we may take as our point of departure the strongest possible meaning and have it weakened by contextual information. This is the approach advocated in Zwarts (2003) and Hogeweg (2005), who use an OT approach to interpretation that incorporates the Strongest Meaning Hypothesis of Dalrymple *et al.* (1994). In the following we adopt the second position without excluding that also the first position may be appropriate under certain conditions.

Under the second approach, syntactic structure, lexical material, context and world knowledge may all help in arriving at the correct interpretation of sentences such as *Most people drink at night*. But crucially, these different factors do not just function one after the other as reducers of the presumably infinite set of interpretations given by a highly underspecified representation. In fact, the different factors can be *in conflict*. For example, the lexical meaning of *most* gives us the relation between two sets such that the intersection contains more elements than the difference between the two does. The lexical material within the sentence and the syntactic structure of the sentence give us these two sets, in this case, the set of people and the set of individuals that drink at night. This would give us the optimal (hence, preferred) interpretation in the absence of further context, and as we have argued above, that would give us the optimal interpretation in the case of *Most people sleep at night*.

For the sentence *Most people drink at night*, however, this interpretation gives rise to a conflict with our world knowledge. As it is probably not even true that most people drink (where *drink* is generally understood as *drink alcohol*), it is hard to believe that it holds for most of the people that they drink at night. So, there is a conflict between the information provided by the syntactic structure of the sentence and the information provided by our world knowledge. This conflict is resolved by considering the next optimal interpretation (that is, “next optimal” purely on the basis of the syntactic structure of the sentence). This is the interpretation such that the set of people gets intersected with the generalized union over the set of alternatives for a certain constituent in the sentence.

If prosodic information is available, then the constituent that gives rise to this set of alternatives is the syntactic argument containing the focus (where focus is marked by sentential accent). In the absence of intonation, we may consider what would be the unmarked constituent to bear the focus. In the case of *Most people drink at night*, the default position of the sentential accent seems to be on *at night*. This gives us as the domain of quantification of *most* the set of people who drink (at certain times). Hence, the interpretation for the entire sentence is that for most of the people who drink (alcohol) it holds that they drink at night. This interpretation is not in conflict with our world knowledge, and it is in fact the optimal (that is, preferred) interpretation against an empty context.

Of course, in the presence of an actual context, another interpretation might become optimal. So, the sentence *Most people DRINK at night* might be used as an answer to the question why there are so many empty beds in the middle of the night, with a concomitant interpretation. Again, this interpretation would deviate from the interpretation dictated by the syntactic structure of the sentence alone.

In these cases, in the competition between a syntactically optimal but pragmatically unlikely interpretation and a pragmatically optimal but syntactically suboptimal interpretation, the latter wins. The advantage of an optimization approach to interpretation is clearly that it can deal with actual conflicts between different factors. In addition, the advantage of a bidirectional optimization approach is that it can deal with the influence of speaker choices on interpretation. Suppose a hearer encounters the form *Most drink* in a context where the topic of the discourse is linguists. In this case, the hearer will interpret the missing noun as the set of linguists. Taking the selected meaning as the input of optimization and determining the optimal form for this meaning, the hearer will be able to check whether the selected meaning is a meaning that could have been intended by the speaker. In a context where the topic of the discourse is linguists, the meaning that most linguists drink can indeed be expressed using the form *Most drink*. Thus the hearer may conclude that the selected meaning could have been intended by the speaker.

In the above example, the meaning selected by the hearer could have been intended by the speaker, since it gives rise to the encountered form in production. In other cases, the selected meaning is *not* identical to the meaning that gives rise to the encountered form. De Hoop and Krämer (2005/6) and Hendriks and Spenader (2005/6) have shown that mismatches between the selected meaning and the meaning that gives rise to the encountered form can be observed in child language. For example, children may interpret the object pronoun *him* in the sentence *Bert washed him* as referring back to the subject *Bert*, thereby violating Principle B of Binding Theory (Chomsky, 1981). However, when they select the optimal form for expressing reference to the subject, they correctly use the reflexive form *himself*. And for expressing reference to some other element, these children correctly use a pronoun. As a result, they display a mismatch between the selected (coreferential) meaning for the pronoun *him* in interpretation and the (non-coreferential) meaning that gives rise to the encountered pronoun *him* in production. Thus bidirectional optimization may be a guiding mechanism in the avoidance of non-compositional meanings, a mechanism which still has to develop during the course of language acquisition.

Because interpretation of the given form and subsequent generation of the selected meaning occur in the same context, contextual influences on interpretation are unproblematic. Since it is the same contextual information that influences interpretation and generation, contextual influences do not render the process of decomposition irreversible.

Recoverability

As we pointed out earlier, a principle related to compositionality is recoverability: Only elements whose semantic content can be recovered from the local context may be left unpronounced. If a speaker wishes to express the meaning that most linguists drink, and if the topic of the discourse is linguists, then the speaker may utter the sentence *Most drink*. A hearer will then be able to infer that the missing noun must be interpreted as the set of linguists. On the other hand, if the topic of the discourse is some other entity, for example people present at the conference, and if the speaker again wishes to express the meaning that most linguists drink, then he or she cannot leave the noun unpronounced. If the speaker would utter the sentence *Most drink* in this context, the hearer would mistakenly interpret the missing noun as the set of people present at the conference.

Recoverability is usually assumed as a meta-restriction on syntactic analyses. However, its status appears to be similar to the status of the principle of compositionality in semantics. Compositionality is crucial to a hearer who wishes to interpret a certain utterance. He or she must use all information that is necessary to arrive at the intended meaning of this utterance. According to the semantic principle of compositionality, this information is initially limited to the meanings of the words in the sentence and the syntactic structure of the sentence, which may be enriched with contextual and prosodic information if the resulting meaning is still underspecified. According to the OT approach, on the other hand, the information used by a hearer includes all, possibly conflicting, information sources (syntactic structure, lexical information, context, intonation) already from the start. Importantly from the perspective of bidirectional OT, the hearer must also take into account the options and information available to a speaker. In particular, the hearer must take into account the form that would express the selected interpretation best. If this best form is different from the form that was encountered, the hearer may conclude that the speaker could not have intended the selected interpretation. As a consequence, the hearer must select a different interpretation. The effect of bidirectional optimization in interpretation is that only compositional interpretations are assigned, that is, only interpretations that the speaker may have intended.

So compositionality is crucial to a hearer taking into account the speaker's perspective. Similarly, recoverability is crucial to a speaker who wishes to express a certain meaning. She must use all information that is necessary to arrive at a certain form for this meaning. Again, from the perspective of bidirectional OT, the speaker must also take into account the way a hearer would interpret the sentence. If the interpretation the hearer would assign to the sentence is different from the meaning the speaker intended to express, this meaning is not recoverable for the hearer. In that case the speaker must select another form. So the effect of bidirectional optimization in production is that no unrecoverable forms are produced. Evidence for bidirectional optimization as a guiding mechanism in the avoidance of unrecoverable forms in a language comes from cross-linguistic analysis (de Swart, 2007), psycholinguistic experiments (Hendriks *et al.*, 2007) and corpus investigation (Bouma, in prep.), and pertains to linguistic phenomena such as object case marking, pronoun selection and word order freezing.

Compositionality and recoverability appear to be two sides of the same coin. While compositionality relates to the (bidirectional) task of the hearer, recoverability relates to the (bidirectional) task of the speaker. Both principles require that the perspective of the other conversational partner is also taken into account. As a result, generation and interpretation are completely reversible without being fully determined by structural properties (cf. van Gelder's example of the Gödel numbering schema for a formal language discussed earlier).

The effect of bidirectional optimization is that whatever speakers can produce, they are able to understand, and vice versa for hearers.

Compositionality is Bidirectional Optimization

Connectionist models are often criticized for their lack of compositionality, since interpretation is assigned to activity patterns but not to individual units. But as van Gelder (1990) points out, “The absence of strictly syntactic structure, however, does not imply the absence of significant structure of any kind.” A similar point is made in Blutner *et al.* (2004). Connectionist approaches to language, such as Optimality Theory, provide the necessary tools to combine different pieces of information (from context, world knowledge, lexicon, syntax) in a precisely defined way. Information provided by the meaning of the lexical items or the syntactic structure can interact or even compete with information given by the context. But in each case the optimal solution is predictable as the different constraints are ranked with respect to each other. Thus, within OT the interpretation of a complex expression is brought out by an optimization procedure that takes into account syntactic and contextual information simultaneously on the basis of a set of ranked constraints of various nature.

Bidirectional OT adds to this general procedure that the hearer takes into account the speaker’s perspective (and, the other way around, that the speaker takes into account the hearer’s perspective). That is, if a form is associated with a certain interpretation within a certain context by a hearer, then within that same context, the same meaning would have been expressed by the same form if the hearer would have been the speaker. To put it differently, the composition of an output form on the basis of an input meaning within a context goes hand in hand with the decomposition of that same form into an output meaning within that same context. By evaluating form-meaning pairs against a set of ranked (cross-modular) constraints, bidirectional OT guarantees a general procedure of optimization from form to meaning and from meaning to form such that a speaker’s optimal expression of a meaning and a hearer’s optimal interpretation of a form depend on each other in each context in a well-defined way.

To conclude, bidirectional OT provides a general, effective, and reliable process for producing and comprehending complex expressions, therefore it *is* compositionality in van Gelder’s sense (i.e., non-concatenative composing and decomposing of complex expressions in a nontrivial and independent way).

References

- Blutner, R. (1998). Lexical pragmatics. *Journal of Semantics*, 15, 115-162.
- Blutner, R. (2000). Some aspects of optimality in natural language interpretation. *Journal of Semantics*, 17, 189-216.
- Blutner, R., Hendriks, P., de Hoop, H., & Schwartz, O. (2004). When Compositionality Fails to Predict Systematicity. In S. D. Levy & R. Gayler (Eds.), *Compositional Connectionism in Cognitive Science. Papers from the AAI Fall Symposium*. Arlington: The AAI Press, 6-11.
- Bouma, G.J. (in prep.). Starting a sentence in Dutch: A corpus study of subject- and object-fronting. Doctoral dissertation, University of Groningen.
- Buchwald, A., Schwartz, O., Seidl, A., & Smolensky, P. (2002). Recoverability Optimality Theory: Discourse anaphora in a bidirectional framework. In J. Bos, M.E. Foster, & C. Matheson (Eds.), *Proceedings of the sixth workshop on the semantics and pragmatics of dialogue (EDILOG 2002)*, Edinburgh, UK, 37-44.
- Chomsky, N. (1981). *Lectures on government and binding*. Dordrecht: Foris.

- Dalrymple, M., Kanazawa, M., Mchombo, S., & Peters, S. (1994). What do reciprocals mean? In M. Harvey & L. Santelmann (Eds.), *Proceedings of SALT 4*. Ithaca: Cornell University, 61-78.
- Dekker, P., & van Rooij, R. (2000). Bi-directional Optimality Theory: An application of game theory. *Journal of Semantics*, 17, 217-242.
- Frege, G. (1884). *Die Grundlagen der Arithmetik. Eine logisch-mathematische Untersuchung über den Begriff der Zahl*. Breslau: W. Koebner. Reprint published by Georg Olms, Hildesheim, 1961.
- van Gelder, T. (1990). Compositionality: A connectionist variation on a classical theme. *Cognitive Science*, 14, 355-384.
- Hendriks, P., Englert, C., Wubs, E., & Hoeks, J. (2007). Age differences in adults' use of referring expressions. Manuscript, University of Groningen,
- Hendriks, P., & de Hoop, H. (1997). On the interpretation of semantic relations in the absence of syntactic structure. In P. Dekker, M. Stokhof, & Y. Venema (Eds.), *Proceedings of the 11th Amsterdam Colloquium*. Amsterdam: ILLC, 157-162.
- Hendriks, P., & de Hoop, H. (2001). Optimality Theoretic Semantics. *Linguistics and Philosophy*, 24, 1-32.
- Hendriks, P., & Spenader, J. (2005/6). When production precedes comprehension: An optimization approach to the acquisition of pronouns. *Language Acquisition*, 13, 319-348.
- Hintikka, J. (1983). *The game of language. Studies in game-theoretical semantics and its applications*. Dordrecht: Reidel. In collaboration with J. Kulas.
- Hogeweg, L. (2005). *Well, about wel. On the diversity and unity of the Dutch particle wel*. MA thesis, Utrecht University.
- de Hoop, H., & Solà, J. (1996). Determiners, context sets, and focus. In J. Camacho, L. Choueiri, & M. Watanabe (Eds.), *Proceedings of the Fourteenth West Coast Conference on Formal Linguistics*. Stanford: CSLI, 155-167.
- de Hoop, H. & Krämer, I. (2005/6). Children's optimal interpretations of indefinite subjects and objects. *Language acquisition*, 13, 103-123.
- Jäger, G. (2002). Some notes on the formal properties of Bidirectional Optimality Theory. *Journal of Logic, Language and Information*, 11, 427-451.
- Janssen, T. (1997). Compositionality. In J. van Benthem & A. ter Meulen (Eds.), *Handbook of logic and language*. Amsterdam: Elsevier, 417-473.
- Kennedy, C. (2002). Comparative deletion and optimality in syntax. *Natural Language & Linguistic Theory*, 20, 553-621.
- Kuhn, J. (2001). *Formal and computational aspects of Optimality-Theoretic Syntax*. Doctoral dissertation, University of Stuttgart.
- Pesetsky, D. (1998). Some optimality principles of sentence pronunciation. In P. Barbosa, D. Fox, P. Hagstrom, M. McGinnis, & D. Pesetsky (Eds.), *Is the best good enough? Optimality and competition in syntax*. Cambridge, MA: MIT Press, 337-383.
- Prince, A., & Smolensky, P. (1993/2004). *Optimality Theory: Constraint interaction in generative grammar*. Malden, MA: Blackwell Publishing, 2004. Also appeared as Technical Report, Rutgers University, New Brunswick, NJ, and University of Colorado at Boulder, 1993.
- Prince, A., & Smolensky, P. (1997). Optimality: From neural networks to Universal Grammar. *Science*, 275, 1604-1610.
- Smolensky, P. (1988). On the proper treatment of connectionism. *Behavioral and Brain Sciences*, 11, 1-23.
- Smolensky, P. (1991). Constituent structure and explanation in an integrated connectionist/symbolic architecture. In C. Macdonald & G. Macdonald (Eds.),

- Connectionism. Debates on psychological explanation.* Oxford, UK: Blackwell, 221-290.
- de Swart, P. (2007). *Cross-linguistic Variation in Object Marking*. Doctoral dissertation, Radboud University Nijmegen.
- Vogel, R. (2004). Remarks on the architecture of OT syntax grammars. In R. Blutner & H. Zeevat (Eds.), *Optimality Theory and Pragmatics*. Houndmills, Basingstoke: Palgrave Macmillan, 211-227.
- Westerståhl, D. (1985). Determiners and context sets. In J. van Benthem & A. ter Meulen (Eds.), *Generalized quantifiers in natural language*. Dordrecht: Foris, 45–71.
- Zeevat, H. (2000). The asymmetry of Optimality Theoretic Syntax and Semantics. *Journal of Semantics*, 17, 243-262.
- Zwarts, J. (2003). Paths round a prototype. In P. Saint-Dizier (Ed.), *ACL-SIGSEM workshop: The Linguistic Dimensions of Prepositions and their Use in Computational Formalisms and Applications*. Toulouse: IRIT, 228-238.