A New Hypothesis on Compositionality

Reinhard Blutner (blutner@hum.uva.nl)
Institute for Language, Logic and Computation (ILLC)
University of Amsterdam, Nieuwe Doelenstraat 15, 1012 CP Amsterdam, The Netherlands

Petra Hendriks (P.Hendriks@let.rug.nl)
Center for Language and Cognition Groningen (CLCG)
University of Groningen, P.O. Box 716, 9700 AS Groningen, The Netherlands

Helen de Hoop (H.deHoop@let.kun.nl)
Department of Linguistics, CLS
University of Nijmegen, P.O. Box 9103, 6500 HD Nijmegen, The Netherlands

Abstract

In this paper we put forward a new hypothesis on compositionality of meaning, namely 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. In contrast, the bidirectional optimization approach we advocate proceeds from 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. Hence, this approach avoids a number of well-known empirical problems associated with concatenative compositionality.

Styles of Compositionality

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)*. Obviously, the thesis of context independency is based on this type of compositionality (cf. Hintikka, 1983, cited in Janssen, 1997): "The meaning of an expression should not depend on the context in which it occurs."

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 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. Smolensky (1996) 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. In that sense, Smolensky's weak compositionality is reminiscent to Frege's (1884) contextuality principle, cited in Janssen (1997): "A word has a meaning only in the context of a sentence, not in separation."

Compositionality and the Role of Context

Notoriously, there are a number of empirical problems with the semantic principle of compositionality. For example, compare the sentence *Most people sleep at night* to the sentence *Most people drink at night*. Although the syntactic structures of these two sentences are absolutely identical, the preferred (default) interpretations differ truth-conditionally. The preferred reading of the first sentence can be paraphrased as "At night, most people sleep" whereas the preferred reading of the second sentence is not "At night, most people drink" but rather, "Most people who drink (alcohol), drink at night".

Because these sentences have the same subject, namely *most people*, at first sight they appear to express conflicting assertions about what most people do. However, they can be simultaneously true in a certain situation for a certain set of people, even if noone of this set of people drinks and sleeps at the same time. In a situation when 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. In the same situation, the sentence *Most people sleep at night* can be true as well, namely if the people who don't drink at night (that is, the non-drinkers as well as the day-drinkers), sleep. These two sentences can be true in the same situation because they have different preferred readings.

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 intonational structure. In addition, their syntactic structures are completely identical. Nevertheless, different preferred readings arise for these sentences. We conclude that in the absence of further context, intonational clues and syntactic differences, these different preferred readings arise as the result of the differences in lexical content in relation to our world knowledge. Although these different readings actually involve different truth conditions, there is no structurally based, mechanical way in which the correct interpretations can be derived. Therefore, sentences like the two above are considered a problem for the principle of compositionality.

The above two sentences are quantificational in nature: they contain a quantificational determiner, in this case *most*. Quantificational determiners are known to establish a relation between two sets. The first set, the domain of quantification, is generally expressed by the noun (together with possible modifiers of the noun) with which the determiner forms a syntactic constituent. However, it can be argued that context always restricts this domain of quantification. Hence, in order to calculate the truth conditions of a quantificational expression, one always has to take into account the context.

The truth conditions of a sentence such as *Most people drink* depend already on what subset of people counts as the domain of quantification of *most*. Usually, a sentence like this will not be about people in general but rather about some relevant subset of people determined by the context. Obviously, a sentence such as *Most of them drink* will have varying truth conditions depending on what *them* refers to. And if the sentence *Most drink* is uttered out of the blue, the hearer even totally depends on the context in order to determine the relevant set of individuals for the quantifier to live on.

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. The content of this empty structure denotes the whole domain of individuals and gets intersected with a context set variable (cf. Westerståhl, 1985). But in fact, we need two context set variables then. One would be equated with the generalized union over the set of alternatives for the syntactic argument that contains the sentential accent, or focus (cf. de Hoop & Solà, 1996), so that the quantificational domain would be the set of individuals who drink at certain times. The other one 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 a *major* 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 Rooy (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 compositionality 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, 1997). Optimality Theory applied to the domains of syntax and semantics involves two closely related issues: 1. given a semantic input, what is its optimal expression?, and 2. given a syntactic input, 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 the ranked constraints in a parallel fashion. One of the advantages of such an approach is that constraints of various nature (syntactic, pragmatic, 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. 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).

Blutner (2000) extensively argues in favour of bidirectional optimization, where speaker's and hearer's optimization are carried out simultaneously. In Blutner's version of bidirectional OT, a form-meaning pair <f,m> is called *super-optimal* if and only if there is no other super-optimal pair <f',m> such that <f',m> » <f,m> (» is an ordering relation which can be read as *'being more harmonic'* or *'being more economical'*) and there is no other super-optimal pair <f,m'> such that <f,m'> » <f,m>. Under the assumption that the relation » is transitive and well-founded, Jäger (2000) 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, to appear). 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, we also have to check whether the underlying content is the optimal interpretation of the reduced string.

Compositionality and Underspecification

In Blutner's (1998, 2000) theory, a view of radical underspecification augmented with contextual enrichment is assumed to account for features of compositionality. In this view, every lexical unit determines an underspecified representation. The combinatorial system determines how lexical units are combined into larger units. There is a system of type and sortal restrictions that determines whether structures of a certain degree of (under-)specification are well-formed. And finally, there is a mechanism of contextual enrichment (pragmatic strengthening based on contextual and encyclopedic knowledge) which is controlled by factors of economy.

So, in an underspecification approach we take the meaning of *most* to be a relation between two sets of individuals, such that the cardinality of the intersection of these two sets exceeds the cardinality of the difference between the first and the second set. Syntactic structure, lexical material, and context can all contribute to the determination of the two sets of individuals in the real world that are related by *most*, thereby all influencing the eventual interpretation of a sentence containing *most*. However, it does not seem to be possible to come up with an algorithm to compute the interpretation of a sentence such as *Most people drink at night* in a certain context, as we cannot predict in what precise way and to what degree context is involved in the interpretation of that sentence from one situation to the other.

In other words, Blutner's theory does not really clarify how the border line between the underspecified representation and the contextual enrichment is ever to be determined. That is, it does not really help us in determining the optimal interpretations of sentences such as *Most people drink at night* or *Most people sleep at night* in the presence or absence of a certain context. This is in general the problem with underspecification approaches, as already pointed out above.

Compositionality and Bidirectionality

In Optimality Theory the procedure that provides us with an optimal interpretation of a given form within a certain context can also be viewed in a radically different way. 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), who uses an OT approach to interpretation that incorporates the Strongest Meaning Hypothesis of Dalrymple *et al.* (1994). Under such an approach, again 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 in fact, 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 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 intonational 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 a 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 among different factors.

Recoverability

As we pointed out earlier, a principle relating 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, possibly conflicting, information to arrive at the intended meaning of this utterance. Importantly from the perspective of bidirectional OT, the hearer must also take into account all options and information available to a speaker. Similarly, recoverability is crucial to a speaker who wishes to express a certain meaning. He or she must use all information to arrive at a certain form for this meaning. Again, the speaker must also take into account the way a hearer would interpret the sentence.

So compositionality and recoverability appear to be two sides of the same coin. While compositionality relates to the task of the hearer, recoverability relates to the task of the speaker. Both principles require that the perspective of the other conversational partner is also taken into account.

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."

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 should be 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 ranked set of 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 a form-meaning pair within a context goes hand in hand with the decomposition of that form-meaning pair 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).

Acknowledgments

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.

Buchwald, A., Schwartz, O., Seidl, A., & Smolensky, P. (2002). Recoverability Optimality Theory: Discourse anaphora in a bidirectional framework. *Proceedings of the sixth workshop on the semantics and pragmatics of dialogue (EDILOG 2002)*.

Dalrymple, M., Kanazawa, M., Mchombo, S., & Peters, S. (1994). What do reciprocals mean? *Proceedings of SALT 4*.

De Hoop, H., & Solà, J. (1996). Determiners, context sets, and focus. *Proceedings of the Fourteenth West Coast Conference on Formal Linguistics*. Stanford: CSLI.

Dekker, P., & van Rooy, 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: Willhelm Koebner.

Hendriks, P., & de Hoop, H. (1997). On the interpretation of semantic relations in the absence of syntactic structure. *Proceedings of the 11th Amsterdam Colloquium*. Amsterdam: ILLC.

Hendriks, P., & de Hoop, H. (2001). Optimality Theoretic Semantics. *Linguistics and Philosophy,* 24, 1-32.

Hintikka, J. (1983). *The game of language*. Dordrecht: Reidel.

Jäger, G. (2000). Some notes on the formal properties of Bidirectional Optimality Theory. In R. Blutner & G. Jäger (Eds.), *Studies in Optimality Theory*. Potsdam: Linguistics in Potsdam.

Janssen, T. (1997). Compositionality. In J. van Benthem & A. ter Meulen (Eds.), *Handbook of logic and language*. Amsterdam: Elsevier Science.

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?* Cambridge, MA: MIT Press.

Prince, A., & Smolensky, P. (1993). *Optimality Theory: Constraint interaction in generative grammar*. Technical Report RuCCSTR-2, Rutgers Center for Cognitive Science.

Prince, A., & Smolensky, P. (1997). Optimality: From neural networks to Universal Grammar. *Science,* 275, 1604-1610.

Smolensky, P. (1996). Constituent structure and explanation in an integrated connectionist/symbolic architecture. In C. Macdonald & G. Macdonald (Eds.), *Connectionism: Debates on psychological explanation*. Oxford, UK: Blackwell.

Van Gelder, T. (1990). Compositionality: A connectionist variation on a classical theme. *Cognitive Science,* 14, 355-384.

Vogel, R. (to appear). Remarks on the architecture of OT syntax. In R. Blutner & H. Zeevat (Eds.), *Optimality Theory and pragmatics*. Palgrave/Macmillan.

Westerståhl, D. (1985). Determiners and context sets. In J. van Benthem & A. ter Meulen (Eds.), *Generalized quantifiers in natural language*. Dordrecht: Foris.

Zeevat, H. (2000). The asymmetry of Optimality Theoretic Syntax and Semantics. *Journal of Semantics*, 17, 243-262.

Zwarts, J. (2003). *Paths round a prototype*. Unpublished manuscript, Utrecht University.