

Compositionality and systematicity

Abstract

A feasible model of natural language semantics should also account for the systematicity displayed by language understanding. The idea that compositionality alone is sufficient to guarantee systematic reasoning within a semantic system is rejected. An analysis of systematic inferences involving adjective–noun combinations with intersective adjectives illustrates that more than a compositional semantics is required for their understanding, and that even simple concepts involve hidden structures. It is argued that the same processes used in determining context-dependent meaning are at work in understanding complex concepts where world-knowledge plays a key role. In the latter cases however accumulated experience with the given concept is argued to serve as a contrast set against which complex concepts are interpreted.

1 What is systematicity?

The term ‘systematicity’ was first introduced in Fodor and Pylyshyn (1988) as a cognitive capacity on par with productivity which a compositional linguistic system is guaranteed to support. Fodor and Pylyshyn refer to systematicity as ‘a feature of cognition, inferential coherence’. In the following quote they describe what they mean (Fodor and Pylyshyn, 1988, p. 37):

The ability to produce/understand some sentences is intrinsically connected to the ability to produce/understand certain others.

However, Fodor and Pylyshyn give no formal definition (see van Gelder, 1990, for an attempt to work out the concept formally), instead choosing to illustrate the concept with examples. Here are examples of what types of systematic inferences we can make with language.

When an agent understands the sentence ‘John loves the girl’,
she understands the sentence ‘The girl loves John’ as well (Fodor and Pylyshyn, 1988, pp. 42, 48–49). (1)

When an agent understands the expressions ‘brown triangle’ and ‘black square’, she understands the expressions ‘brown square’ and ‘black triangle’ as well. (Based on examples given in Fodor and Pylyshyn, 1988.) (2)

Initially, these two ‘systematicity clauses’ seem intuitively to be true. Surely understanding ‘brown triangle’ and ‘black square’ means that you will be able to infer what ‘brown square’ means. Systematicity is argued to follow from compositionality because understanding a complex concept like ‘brown triangle’ is achieved through an understanding of ‘brown’ and ‘triangle’, and therefore an agent must be able to use these lexical items in novel combinations. It is true that we understand all the above expressions, but how do we do it and what type of linguistic model will support these inferences? The ability to predict and confirm systematic inferences like these would seem to be a prerequisite for any language model that claims to model natural language understanding.

Does understanding the results of the modification of one concept really insure the understanding of the same modification when applied to another concept? We will argue that the answer is not as simple as it first looks, not only for complex cases, but even for so-called simple compounds. Additionally, we will argue that the compositionality of a semantic theory is not sufficient to insure that the semantics predicts systematicity in language understanding.

In Section 2 we examine several systematicity inferences that range from those requiring some reasoning to those that lead to compounds that seem only interpretable as idiomatic expressions, or even seem uninterpretable. Section 3 looks at simple adjective–noun compounds involving intersective adjectives. The classic compositional approach is argued to be insufficient for correctly producing even the simplest systematicity clauses. Complex expressions with intersective adjectives are examined in more detail in Section 4. It is first shown that even the simplest concepts have hidden structures that are relevant to their interpretation. Second, it is shown that context affects both how adjective–noun phrases are produced and how they are interpreted. We then argue that the mechanisms at work in cases where context-sets play a role in the meaning of adjective–noun compounds, are the same mechanisms that were at work in determining in what way modification applies to a given object relative to other instances of that object in the world. Finally, Section 5 summarizes what implications this work has for other work in lexical semantics and natural language understanding.

2 Concepts do not relate systematically

Lexical items evoke or refer to concepts, but these concepts don’t necessarily relate in predictable ways. Consider the following examples of systematicity clauses where determining the actual meaning of each complex expression is not straightforwardly computable.

An agent who understands ‘good writer’ and ‘bad teacher’
also understands ‘bad writer’ and ‘good teacher’. (3)

An agent who understands ‘within an hour’ and ‘without a watch’
also understands ‘without an hour’ and ‘within a watch’. (4)

An agent who understands ‘Mom drives me to kindergarten’
also understands ‘I drive mom to kindergarten’. (5)

Consider example (3). Do all who understand ‘good writer’ and ‘bad teacher’ also understand ‘bad writer’ and ‘good teacher’? The answer seems to be *yes*. However, determining the meaning of the new complex is a more involved process. For most of us, a writer is good or bad based on characteristics that are different from those under which a teacher is considered good or bad. Both characteristics can be identified from understanding the first two compound expressions. In order to understand the second two we seem to switch from the one end of the scale to the other.

Example (4) seems unlikely. Do all who understand ‘within an hour’ and ‘without a watch’ also understand ‘within a watch’ and ‘without an hour’? (Szabó, 2000.) An interpretation requires some creativity, given that ‘within a watch’ does not seem to have any type of conventional meaning, and even if it did, it would not be part of combination of the meaning of the words ‘within’ and ‘watch’. On hearing it, listeners are more likely to assume that it has some sort of idiosyncratic, noncompositional meaning that they are unaware of.

Finally, example (5) illustrates a case where the problem is conceptual. Would a child who understands ‘Mom drives me to kindergarten’ also be understand ‘I drive mom to kindergarten’? Here we are dealing with another type of example where there is a problem of conceptualization and likelihood. A child of kindergarten age may not even initially be able to conceive that he or she could drive the car.

All three of these examples illustrate that systematic inferences seem to be heavily dependent on world knowledge, and not just linguistic knowledge. This raises the question what is actually involved in systematicity. If the systematicity that we sometimes find is a feature of language, entailed by virtue of language being a compositional system, then it should be in some ways independent of world knowledge. If however, systematic inferences are in part possible because of the way the world is, and are not really a reflection of linguistic structures, then whether or not your semantic system is compositional would seem to be beside the point. Consider what Clark (1993, p. 148) says:

Instead of treating [systematicity] as a property to be directly induced by a canny choice of basic architecture, it may be fruitful to try to treat it as intrinsic to the knowledge we want a system to acquire.

Clark seems to suggest then that our understanding of some complex concepts does not carry over to other complex concepts because the complex concept is modulated by what we know about the world. We find this a fruitful perspective to consider, but it necessarily leads to the obvious question of how this modulation takes place and how it relates to language.

The complex concepts examined here are all exceptional and rather difficult. For the rest of the paper we limit the scope of the analysis to adjective–noun combinations with intersective adjectives, which are generally considered to be very tractable. Adjectival modification is a well-studied phenomenon that has often been argued to be compositionally determinable and in fact the set of intersective adjectives by definition, derives the meaning of a concatenation of an intersective adjective and the noun it modifies by simply taking the intersection of the sets that the adjective and the noun are an extension of. Formal treatments of adjectives such as Kamp (1975) and Kamp and

Partee (1995) developed from these early ideas.¹ What we will show is that correctly interpreting complex expressions made with these ‘easy’ adjectives are as challenging as the examples just given above, and are strong evidence that modeling systematicity in a semantic theory requires much more than just compositionality.

3 Compositionality and ‘simple’ adjective–noun combinations

Fodor and Pylyshyn (1988) claim that systematicity of understanding follows from compositionality. Thus by virtue of having a semantics that is compositional, valid systematicity inferences should be naturally available. Fodor and Pylyshyn subscribe thus to a definition of compositionality like the following (taken from Szabo, 2000)²:

The meaning of a complex expression is determined by
the meanings of its constituents and by its structure. (C)

This is because compositionality should insure that all there is to the meaning of a complex construction can be derived solely based on the meaning contained in the individual words, stored somewhere in a mental lexicon. If we know what the meaning of individual constituents are, here words, then we can construct the meaning of complex concepts. The systematicity of inferences like (1) and (2) are then plausible for two reasons. First, the meaning of the individual constituents of each phrase (word) is treated as an unstructured, context-independent meaning, which can be modeled by a set. Second, the syntactic operation of modification is handled by the operation of set intersection in the semantics. Of course, it is well known that this only works for intersective adjectives, but does it even work for these?

According to this classical view, understanding complex concepts in natural language proceeds in the following way. To understand a phrase such as ‘brown triangle’ we roughly follow the following steps:

- (i) we retrieve the conceptual representation corresponding to each lexical item, i.e., ‘brown’ and ‘triangle’;
- (ii) these conceptual components have known truth-functional content which is represented in terms of sets of objects that can be truthfully referred to by the lexical item;
- (iii) with the help of the simple logical operation of conjunction it is possible to determine the truth-conditions of the whole phrase under discussion ($\text{brown} \cap \text{triangle}$).

Under this view similar steps can be followed to derive systematicity clauses. Understanding the expressions ‘brown triangle’ and ‘black square’ is done by first deriving the meanings of $\text{brown} \cap \text{triangle}$ and $\text{black} \cap \text{square}$, as done above then:

¹Note however that many researchers in cognitive semantics no longer consider sets to be an appropriate representation for lexical meaning, e.g., Rosch (1977), Fillmore (1982, 1985), among others.

²This is only one of many ways in which compositionality can be defined but seems to be most consistent with the view Fodor and Pylyshyn take. Another possible way is to take a functional definition such as the definition of functional compositionality given by Szabo: ‘(F) The meaning of a complex expression is a function of the meaning of its constituents and of its structure’ (p. 484).

1. Extract the lexicon entries ‘brown’ $\rightarrow$ *brown*, ‘black’ $\rightarrow$ *black*, ‘triangle’ $\rightarrow$ *triangle*, ‘square’ $\rightarrow$ *square*. Assume that the extracted meaning components are grounded (i.e., they have a fixed truth-conditional content).
2. Calculate the corresponding conceptual representations for ‘brown square’ and ‘black triangle’: *brown* $\cap$ *square* and *black* $\cap$ *triangle*.
3. Determine the truth-conditional impact of these forms by using the truth-conditional impact of the grounded constituents involved and the operation that connects them.

There are two major problems with this classical view of meaning. First, the above proposal does not recognize or take into account the fact that concepts (as well as features that modify them), can have structure. Instead, the approach tries to treat all modifiers and objects modified as if they were basically equivalent. But the modification of one object can be very different from the modification of another, and often the value of a modifier is only relevant to certain aspects of a given object, and is inappropriate for others. Without knowing these details about how the modification applies, understanding one adjective–noun combination often does not entail that one can understand how the same adjective would be interpreted when applied to another noun.

Second, the global context of use also influences when a complex compound is appropriate and what its interpretation will be. The simplistic classical view does not take seriously the possibility that in some cases these contextual effects can also influence truth-conditional meaning. These two major problems with the classical view are addressed in the next section, and a simple outline of how to approach our actual understanding of these complex concepts is given.

4 Failures of the classical view

4.1 Even simple concepts have hidden structure

It is well accepted that very abstract concepts such as events have internal structure, and there is even a rich literature on dealing with unrealized event arguments. Expressions with agentive nouns and change of state verbs cannot be treated with simple functional application and require much more apparatus, sometimes postulating the existence of hidden event variables (see, e.g., Larson, 1998, on agentive nouns and Dowty, 1979, on change of state verbs).

The fact that even simple objects have structures or features is often ignored, considered to be a problem having to do with world knowledge rather than a linguistic matter. But this structure can have an effect on semantic interpretation and plays an important role when the object is modified. Sometimes the modification applies to only part of the structure (or is only relevant to part of the structure). Consider how the modifier in the following noun–adjective combination applies to only one part of the concept evoked by the noun in the following examples:

- (a) a red apple [has a red peel]
- (b) a sweet apple [is sweet inside]

- (c) a pink grapefruit [the color of the pulp is pink]
- (d) a white room/a white house [the inside is white/the outside is white]

We see in each case that the modifier applies to some relevant part of the object, in some cases determined by the object, but in some cases determined by the modifier as well. This phenomenon is well known. Quine (1960) was the first to note the contrast between *red apple* (red on the outside) and *pink grapefruit* (pink on the inside). In a similar vein, Lahav (1993, p. 76) argues that an adjective such as *brown* does not make a simple and fixed contribution to any composite expression in which it appears:

In order for a cow to be brown most of its body's surface should be brown, though not its udders, eyes, or internal organs. A brown crystal, on the other hand, needs to be brown both inside and outside. A brown book is brown if its cover, but not necessarily its inner pages, are mostly brown, while a newspaper is brown only if all its pages are brown. For a potato to be brown it needs to be brown only outside, . . .

In cognitive semantics the part of a concept for which a given modification is relevant is often referred to as an 'active zone', first discussed as such in Langacker (1991). In the case of an apple, the color is only relevant for the skin of the apple, which is its active zone.

What we have now seen is that even simple adjective–noun compounds as 'red apple' require more than merely knowing what things are red and what things are apples, i.e., a semantic treatment using sets and set intersection will not work. Without ascribing more descriptive structure to the meaning of nouns than the set of objects that can be truthfully described by that noun, we cannot seriously model the understanding of complex concepts described by adjective–noun compounds. We need to know additionally that a color term is relevant generally to the skin of an apple to be able to say that we have correctly captured the meaning of 'red apple' with our semantics.

If the classical view fails for simple adjective–noun combinations then it cannot be a possible semantics for predicting true systematicity clauses. Consider the following:

When an agent understands the expressions 'red apple' and 'sweet grapefruit', then it's likely that she understands 'red grapefruit' and 'sweet apple' as well. (6)

We do not understand 'red grapefruit' as referring to the peel of the grapefruit. If asked, most people will interpret a 'red grapefruit' as referring to a grapefruit where the pulp is red, and a sweet apple refers to the sweetness of the skin and meat of the apple. This interpretation is achieved by knowing that colors refer to the skins of an apple, but to the pulp of a grapefruit (i.e., we talk of white or pink grapefruits), and that sweetness applies to the edible parts of a fruit. These inferences cannot be captured by treating the individual words as simple sets and taking their intersection for the complex concept.

For some adjective–noun combinations the classical view seems to work, e.g., for 'black square' and 'brown triangle'. However, this is only a coincidence. Systematicity clauses derived with modificational adjective–noun combinations will fail in all cases where the adjectives and nouns cannot be interpreted as two intersective sets. Adjectives and nouns can only be interpreted as two intersective sets when the concept that

the adjective modifies has no internal structure. Since this is seldom the case, systematicity clauses that appear to be correct are actually only valid in cases where by sheer coincidence the world actually mirrors what intersective adjectival modification would look like.

The actual structure that objects have is not readily described because it differs greatly by object type. Information about relevant modification seems to be deeply grounded in our experience with the world. If the inside of apples were different colors, then 'red apple' could have a very different meaning. In the same way different cultures may have divided up the referential possibilities of the same object in different ways. Thus in English we speak of 'brown sugar' for unprocessed sugar, whereas in Japanese this is called 'kurozatto', literally 'black sugar'. These compounds could be dismissed as noncompositional idioms, and are perhaps argued to be outside the realm of what our theory has to explain. But regarding such examples as exceptions puts us on a slippery slope and we become forced to establish an arbitrary border between 'brown sugar', on the one hand, and 'red apples', on the other, classifying one as compositional and the other as idiomatic. Clark's (1993) view that our systematic understanding has more to do with the features or characteristics of the world that language describes than languages itself, is only partially right. There is a degree of arbitrariness in the applicability of individual linguistic terms that has led to different lexical choices in different languages, even if both choices have a logic to them. Even though we strive to identify general principles, there will always be a certain degree of arbitrariness that cannot be ignored.

Of course, even such pervasive factors as the influence of context and the underdetermination of most lexical meaning have been ignored by those who try to rely on a classical compositional theory that accounts for some systematicity conditions under additional (questionable) conditions. Breheny (2004, p. 8) makes the following remark in regard to this type of analysis:

... they propose that a compositional semantics for English can be given – it is just that English speakers choose to ignore the semantics of their language on almost all occasions of use. This view gives rise to a theory which is of dubious explanatory value. If meaning is cut off from understanding in this way, it is unclear how the proposed semantic theory is meant to meet any of the basic desiderata for semantic theories which invoke speakers' understanding of natural language (that it is productive and systematic) or learnability. What is needed is a semantic theory which can explain how the meanings of constituent expressions contribute to a compositional determination of the semantic content of an utterance in context.

In summary, merely having a compositional semantics isn't informative enough to model the systematicity of language understanding that is such an essential part of our language competence. This is not to say that there is no system involved in understanding these compounds. We do have strong intuitions about what, where or in which way a given modifier should apply to some object. These intuitions come from our world experience and could easily be otherwise. However, the way we have acquired the information was with general principles, which explains why human understanding seems often to be systematic. A hint as to what these general processes are like can be seen by looking at the context dependence of meaning.

4.2 Context dependence of even the simplest concepts

In the previous section we showed that a treatment of adjective–noun combinations based on intersection alone does not give the systematic inferences expected or desired when interpreting simple compounds in isolation. It is further the case that intersection also fails to give a plausible account for simple adjective–noun combinations in a context. The reason is because meaning is context-dependent. In his “Grundlagen der Mathematik” Frege (1884) noticed the context-dependence of words. ‘One should ask for the meaning of a word only in the context of a sentence, and not in isolation’ (Frege, 1884). Frege considered this context-dependency as evidence of noncompositionality in natural language.

Fodor and Pylyshyn do not use a compositional model that takes context-dependence into account, though such models exist. However, even with such a model the phenomenon of context-dependence provides an argument against assuming that compositionality will guarantee the systematicity. Consider the following two diagrams (taken from Bosch, 2002) in Figure 1. Note that in both diagrams the dark gray areas of the middle triangle should be interpreted as red. Consider Figure 1A. The leftmost triangle in Figure 1A can best be referred to as ‘the white triangle’. However, in Figure 1B, a better referring expression would be ‘the black triangle’. The contrast set with which the triangle is displayed influences what referring expression is the most appropriate. Here again we see that the modifier that would be the most distinguishing is the most appropriate modifier for the objects.

It could be argued that the contrast set is only relevant for choosing an appropriate referring expression. Classical semantics will generate a set of possible expressions and bi-directional reasoning of some kind will help a speaker to choose, forcing a speaker to choose as the most appropriate referring expression the expression that will help the hearer identify the intended triangle. This means that we would refer to the leftmost triangle as both the ‘white triangle’ and the ‘black triangle’ but with our classical theory we would have no way to express that this refers to different perspectives. The same object would be referred to as ‘black’ and ‘white’, seemingly the result of the same process of interpretation, and that would seem strange. It seems as if we are rel-




Figure 1.

egating a key component in the semantics to processes we then consider to be outside it.

Alternatively, we could accept that utterance meaning is context-dependent. There are many proposals in the literature for handling context-dependent meaning which give a formal treatment of indexical expressions such as *I, you, he, here, now, that, that book*, etc. within model-theoretic semantics (e.g., Montague, 1970; Kaplan, 1979). A Kaplan context can be straightforwardly augmented with further components by including a list of contextual elements. This was done for instance by Nunberg (1995) and Sag (1981) in order to describe predicate transfer. This could be easily extended to deal with contrast sets of the type discussed above. Another possible approach is to make use of neo-Gricean and post-Gricean ideas to include such effects. In this connection relevance theory, e.g., Sperber and Wilson (1986, 1998), Carston (1999, 2002) and blending theory in cognitive linguistics, Fauconnier and Turner (2002), should be mentioned.

In either approach agents are argued to evaluate expressions relative to the contrast set in real time, i.e., the correct referring expression relative to the contrast set is not stored someplace in the mental lexicon. As for systematicity inferences, context will have to be incorporated into the inferences predicted because an individual who knows what a white triangle and a black square are in one context, might not necessarily be able to predict what a black triangle and a white square will be, without knowing the context.

4.3 Same processes, different level of conventionalization

In both of the above cases we saw that the classical view of the meaning of adjective–noun compounds as intersective gets into trouble. On the one hand, we saw that detailed, structured information from world experience plays a role in even the simplest compounds. On the other hand, we saw that the context can strongly affect the interpretation of compounds that without a context first appeared to be straightforwardly interpretable.

But if we look closely, we can see that the processes involved in both of these cases are actually the same. The process used in determining the best expression or interpretation from a contextual set involves analyzing the set to choose those features that are useful in distinguishing the intended triangle from the others, and then using those features in generating a referential expression. The same type of reasoning is involved in the way in which active zone type information for objects and their modifiers are determined. The difference is that the contrast set is made up of our accumulated experience with the objects involved. How do we know that a color term applies to the skin of an apple? If presented with a set of apples we would quickly see that color is only distinguishing for the skin. If we collected all our experiences with apples into an ‘experience’ contrast set, we would be able to come to the same conclusion. We can determine that color is only relevant for the skin, and thus the red in ‘red apple’ must refer to the skin.

The difference between the two cases has to do with when and how the process of interpretation proceeds. That color in apples refers to skins is something that for

most of us has been integrated into the sum of knowledge we have about apples as a concept, and does not seem to be something we need to calculate repeatedly on demand. It is possible that for common concepts this information is stored. On the other hand, context-dependent aspects of references seem to be dealt with on the fly. But in both cases the same general processes seem to be used.

5 Conclusions and implications

Fodor and Pylyshyn (1988) brought up the important point that there is a systematicity in our natural language understanding. They argue that our ability to make systematic inferences is part of our communicative competence, and any model of human language that attempts to account for human language production and interpretation must also be able to handle these systematic inferences.

However, we have shown here that compositionality is not a sufficient requirement for a semantic system that claims to model systematicity, at least not when we adopt a conception of ‘understanding’ that relates to *natural* language understanding. We use the term natural language understanding in opposition to an *artificial* notion of ‘understanding’ such as the one that is limited to the construction of an abstract representation that deals with the formal semantic side of language, say with handling analyticity. Such a theory can always be right, and is thus not very interesting, nor very successful as a model of natural language.

We have illustrated our point by examining seemingly simple examples of adjective modification with intersective adjectives, and showing that even these examples demand a deeper analysis than what is usually given. This is because even these simple items have unarticulated structure.

In this work we therefore go one step further than work in cognitive semantics, which only identifies the problem of structured concepts with structure derived from world experience. Cognitive semantics also identifies the solution to this problem, the necessary identification of the active-zone relevant for a given type of modification for each concept. However this work stops short of suggesting any general process by which the solution could be worked out on any large scale.

Instead, we have argued that the same general process used in reference with respect to a context set can be used to determine the relevant hidden structure based on world knowledge that is necessary to understand concept modifications. Our solution also looks at the problem of adjective–noun interpretation as the result of a dynamic process, rather than striving to characterize their interpretation as the result of the concatenation of a finite list of static features. Our perspective of lexical meaning as procedural and dynamic, rather than static, is novel. What remains is to formalize this process. Proposals like van Rooy’s (2000) work with a probabilistic theory of entropy seems a promising approach, as does work using connectionist models along the lights of the *selective modification model* proposed by Smith *et al.* (1988), though this is an area where much research remains to be done (see also Blutner *et al.*, 2004).

We also anticipate that some of the ideas found here can be fruitfully applied to work on hidden structures done on agent predicates and change of state verbs

(e.g., Dowty, 1979; Larson, 1998). Most of this research has concentrated on syntax–semantic interface issues, focusing on deriving representations for underspecified syntactic structures where there is a mismatch with semantics. But many of the examples show the same characteristics as our examples and we would in future work like to more fully explore the similarities and differences between the two research streams which have examined these two related problems so differently.

References

- Blutner, R., P. Hendriks, H. de Hoop and O. Schwartz. 2004. When compositionality fails to predict systematicity. In S.D. Levy and R. Gaylor (eds), *Compositional Connectionism in Cognitive Science*. Papers from the AAAI Fall Symposium, pp. 6–11.
- Bosch, P. 2002. Explaining semantic productivity. The Symposium on Logic and Creativity: Integrating Categorical Rules and Experiences, Osnabrück.
- Breheny, R. 2004. On the dynamic turn in the study of meaning and interpretation. Unpublished manuscript, University of Cambridge.
- Carston, R. 1999. The semantics/pragmatics distinction: a view from relevance theory. In K. Turner (ed.), *The Semantics/Pragmatics Interface from Different Points of View*. Elsevier, Oxford.
- Carston, R. 2002. *Thoughts and Utterances: The Pragmatics of Explicit Communication*. Blackwell, Oxford.
- Clark, A. 1993. *Associative Engines: Connectionism, Concepts and Representational Changes*. MIT Press, Cambridge, MA.
- Dowty, D. 1979. *Word Meaning and Montague Grammar: The Semantics of Verbs and Times in Generative Semantics and in Montague's PTQ*. Reidel, Dordrecht.
- Fauconnier, G. and M. Turner. 2002. *The Way We Think. Conceptual Blending and the Mind's Hidden Complexities*. Basic Books, New York.
- Fillmore, C. 1982. Frame semantics. In *Linguistics in the Morning Calm*, pp. 111–137. Hanshin Publishing Co., Seoul, South Korea.
- Fillmore, C. 1985. Frames and the semantics of understanding. *Quanderni di Semantica* 6 (2): 222–254.
- Fodor, J.A. and Z.W. Pylyshyn. 1988. Connectionism and cognitive architecture: A critical analysis. *Cognition* 28: 3–71.
- Frege, G. 1884. *Grundlagen der Arithmetik. Eine logisch-mathematische Untersuchung über den Begriff der Zahl*. W. Koebner, Breslau.
- Kamp, H. 1975. Two theories about adjectives. In E. Keenan (ed.), *Formal Semantics of Natural Language*, pp. 123–155. Cambridge University Press, Cambridge.
- Kamp, H. and B. Partee. 1995. Prototype theory and compositionality. *Cognition* 57: 129–191.
- Kaplan, D. 1979. On the logic of demonstratives. *J. Philosophical Logic* 8: 81–89.
- Lahav, R. 1993. The combinatorial–connectionist debate and the pragmatics of adjectives. *Pragmatics and Cognition* 1: 71–88.
- Langacker, R. 1991. *Foundations of Cognitive Grammar. Volume II: Descriptive Application*. Stanford Univ. Press, Stanford.
- Larson, R. 1998. Events and modification in nominals. In D. Strolovitch and A. Lawson (eds), *Proceedings from SALT VIII*. Cornell University, Ithaca.
- Montague, R. 1970. Universal grammar. *Theoria* 36: 373–398.
- Nunberg, G. 1995. Transfers of meaning. *J. Semantics* 12: 109–132.
- Quine, W. 1960. *Word and Object*. MIT Press, Cambridge, MA.
- Rosch, T. 1977. Principles of categorization. In E. Rosch and B.B. Lloyd (eds), *Cognition and Categorization*. Lawrence Erlbaum Associates, Hillsdale, NJ.
- Sag, I. 1981. Formal semantics and extralinguistic context. In P. Cole (ed.), *Radical Pragmatics*, pp. 273–294. Academic Press, New York.
- Smith, R., Osherson, D., Rips, L., and Keane, M. 1988. Combining prototypes: a selective modification model. *Cognitive Science* 12: 485–527.

- Sperber, D. and D. Wilson. 1986. *Relevance*. Blackwell Sci., Oxford.
- Sperber, D. and D. Wilson. 1998. The mapping between the mental and the public lexicon. In P. Carruthers and J. Boucher (eds), *Language and Thought*, pp. 184–200. Cambridge Univ. Press, Cambridge.
- Szabó, Z.G. 2000. Compositionality as supervariance. *Linguistics and Philosophy* **23**: 475–505.
- van Gelder, T. 1990. Compositionality: A connectionist variation on a classical theme. *Cognitive Science* **14**: 355–384.
- van Rooy, R. 2000. Comparing questions and answers: a bit of logic, a bit of language, and some bits of information. Unpublished manuscript, University of Amsterdam.