
The Principle of Semantic Compositionality

Francis Jeffry Pelletier

ABSTRACT. The Principle of Semantic Compositionality (sometimes called 'Frege's Principle') is the principle that the meaning of a (syntactically complex) whole is a function only of the meanings of its (syntactic) parts together with the manner in which these parts were combined. This principle has been extremely influential throughout the history of formal semantics; it has had a tremendous impact upon modern linguistics ever since Montague Grammars became known; and it has more recently shown up as a guiding principle for a certain direction in cognitive science.

Despite the fact that The Principle is vague or underspecified at a number of points – such as what meaning is, what counts as a part, what counts as a syntactic complex, what counts as combination – this has not stopped some people from viewing The Principle as obviously true, true almost by definition. And it has not stopped other people from viewing The Principle as false, almost pernicious in its effect. And some of these latter theorists think that it is an empirically false principle while others think of it as a methodologically wrong-headed way to proceed.

In fact, there are approximately 318 arguments against The Principle which can be found in the literature, whereas there are only three (or maybe four) arguments proposed in favor of The Principle. This paper will adjudicate among these arguments. And at the end it will suggest some other way to look at what proponents of compositionality really want.

1. Introduction: What is semantic compositionality?

The Principle of Semantic Compositionality is the principle that the meaning of an expression is a function of, and only of, the meanings of its parts together with the method by which those parts are combined.¹ As stated, The Principle is vague or underspecified at a number of points such as 'what counts as a part', 'what is a meaning', 'what kind of function is allowed' and the like. But this hasn't stopped some people from treating it as an obviously true principle, true almost by definition, nor has it stopped some others from attacking it both on 'empirical grounds' and on theoretico-methodological grounds. It seems to me that

many of these discussions fail because of a lack of precision on the abovementioned points and that other discussions are best described as 'how compositionality can/cannot be accommodated within theory *X*' rather than whether The Principle is or is not true. In its most general form, for instance as stated above, The Principle makes no assumptions about what meaning is, nor does it say how one can tell whether two expressions have the same or different meanings. It makes no assumptions about what the parts of a complex expression are, nor does it put any restrictions on what is the function on the parts and the mode of combination.²

In general, The Principle has received a good press. It is rather difficult to find anyone who has other than warm feelings towards The Principle, at least in the philosophical literature. However, many writers in the linguistic literature have pointed to certain types of sentences for which they think it is difficult to give a compositional account. Some of these sentences will be considered below. My own feeling is that although many of these sentences *can* be handled compositionally (indeed perhaps all of the ones I'll mention), still, and nonetheless, the overall conclusion was right: compositionality is false. I shall not attempt to prove this here. Instead I will merely sketch what I take to be an alternative conception of semantics, one that relies on 'groundedness' rather than on 'functionality'. To this end I will survey some of the arguments both for and against The Principle of Compositionality and will argue that none of them achieve what they have set out to achieve; that is, none of them show that compositionality either is or isn't false, and thus all we are left with is either the 'warm fuzzy feeling' that arises when a theory claims to be compositional, or the 'challenge for the establishment' feeling that arises when we claim our theory is non-compositional.

2. Some senses of ‘compositionality’

There have always been other notions of ‘compositionality’ that are appealed to in scholarly work concerning whether The Principle is or isn’t true. I mention here a few of them and try to give a feel for what is understood within the community by the notion of compositionality.

1. The only way to combine meanings is by function application.

The sentiment in (1) is attributed to Richard Montague (see Thomason, 1974) by Brian Smith (1988), and has been called ‘the intuitive version of The Principle’ by Theo Janssen (1983).

2. Meanings of complex symbols are systematically determined by their composition.

(2) is the understanding of ‘compositionality’ given in Haugeland (1985).

3. By ‘compositionality’ we mean that the meaning of the whole is a systematic function of the meaning of its parts.

(3) is taken from Graeme Hirst (1987). Together, (2) and (3) is what many people would understand by the notion of compositionality.

4. ‘Compositionality’ is taken to mean that with any piece you can associate something such that, given a whole made of parts, there is a way of systematically deriving the ‘meaning’ of the whole from the ‘meanings’ of the parts.

This is taken from Brian Smith (1988), who is giving a very general account of what the notion of compositionality means in terms of algebras for syntax and semantics. Together, these last three concepts of compositionality probably capture the popular notion that the functions which are used to combine meanings of parts into meanings of wholes have to in some sense always ‘work the same way’. If they combine to form a given semantic category in one case where a noun phrase is combining with a verb phrase in a certain way, for example, then in all cases where a noun phrase is combining in that way with a verb phrase, it has to form an element of the same category. An extremely strong version of that very notion of compositionality is prevalent in the linguistics folklore. I have heard it on numerous occasions, and instead of citing anyone in

particular, let me just attribute it to an anonymous participant in a recent Linguistic Society of America meeting.

5. ‘Compositionality’ means that the grammar obeys the ‘rule-to-rule hypothesis’.

We see here that we’re getting more and more specific understandings of what kinds of functions are going to be permitted in order for it to be considered a compositional function. The general thrust is that a compositional semantics has to be ‘systematic’, in some or other sense of the term.

My own view is that one person’s anarchy is another person’s favorite form of systematic organization. What the above authors have just cited as being systematic (and by implication, the kinds of things they think are not systematic and hence not to be understood as falling under The Principle) might very well be opposite to sorts of things that I think are systematic. The Principle itself makes no claims *at all* about what sorts of functions there are. It would seem to me that it might be quite easy to argue against the notion of compositionality, if you insist that the functions have to be ‘systematic’ in the sense that our authors have just given us.³ However, it would be much more difficult to argue against the notion of semantic compositionality if you allowed *any* conceivable function whatsoever. Yet this is precisely what I intend to do later in the paper.

But before embarking on that task, I would like to turn to the topic of the *extent* to which The Principle applies. Our first quotation is taken from Allan (1986, pp. 61–62). We’ll see that Allan believes compositionality to be an extremely broad notion indeed.

The meaning of *S*’s Utterance delivered to *H* in context *C* is composed from:

- i. The meaning of sentence Σ that *S* uses.
- ii. The meaning contributed by the prosody ϕ with which Σ is spoken.
- iii. The meaningful input to the interpretation of Σ spoken with ϕ in *C* and from background information. . . . Each of these components of *U*’s meaning is itself compositional. . . . Sentence meaning is compositional. It manifests a compositional hierarchy such that the meaning of a sentence is composed from the meanings of its constituent clauses (and their connectives), the meaning of a clause is composed from the meanings of its constituent phrases, the meaning of a phrase is composed from the meanings of its constituent words, and those in turn from the meanings of their constituent lexemes and morphemes (semantic primitives).

As we see from this quotation, Allan, a champion of

The Principle, believes in compositionality “all the way up and all the way down”. But it is not just champions of The Principle who have this view of the scope of compositionality. Theorists who don’t believe in the notion of semantic compositionality also have the view that the extent of The Principle is “compositionality all the way up and down”. They just deny that it occurs. Here is a quotation taken from Kamp (1990).

The question here depends to a large extent at what we take the data a semantic theory should account for. Thus, some people would question whether anaphora resolution . . . is something that a semantic theory should address . . . I think that a theory of language should account for *all* linguistic data . . . What are the facts that a comprehensive theory of language should explain? . . . My own answer is that linguistic meaning is to be understood in terms of the potential to modify states of information . . . The theory should include an analysis of the inference mechanisms that create, modify and exploit the information states in terms of which linguistic meaning is defined. . . . To sum up, someone who takes seriously the idea that meaning is to be extricated in terms of information change potential, will find the change for a compositional theory as extraordinarily slim.

We see here that Kamp and Allan agree about what a compositional theory would look like. It’s a wide-ranging theory, starting with the individual lexemes and going all the way to the utterance in context. Allan believes that all of that is compositional. Kamp thinks that it is extraordinarily unlikely to be compositional.

3. A psychological argument against The Principle

Certain arguments that have been raised against The Principle of Semantic Compositionality seem to be less successful than others. A certain group of these arguments point out that human minds bring to bear a lot of in-built interpretations to things that they perceive. For example, Gestalt psychologists have long ago pointed out that people will tend to see many things which ‘aren’t there’ because their mind ‘fills in the gaps’, or alternatively put, the mind makes certain inferences about the whole that is expected when only certain parts of the whole are seen. Since this can happen in ordinary situations, such as vision for example, some theorists conclude that the same thing might be happening in linguistic contexts. The mind might bring to bear many things that are not present in the actual linguistic input. So, having brought these things to bear, the meaning that is associated with the combination of these linguistic inputs might not be

present in any of the parts, and so compositionality is false.

It seems to me that this sort of argumentation is not very good. One tactic to argue against it might be to say that, when the mind infers these kinds of things that ‘aren’t there in the input’, then in fact it is the action of ‘putting together’ of these smaller parts which introduces ‘the things that the mind brings to bear’. (Recall that The Principle did not only appeal to the meanings of the parts but also it appealed to the method by which the parts are composed; and one might say that it is here that these other aspects of the meaning of the whole are to be found.)

Another tactic, one that I am here in favour of, is to argue that many things which might be thought to fall under the rubric of ‘being brought to bear by the mind even though they aren’t there in the parts’ are not really involved in the meanings of the composed whole. Let me approach this by means of an analogy. Gestalt psychologists have used examples such as the ‘square’ shown in Figure 1.




Fig. 1. A Gestalt ‘square’.

According to Gestalt psychologists, we see a square in Figure 1, despite the fact that a careful examination reveals that the only things in the picture are emphasized corners. There is no square at all in Figure 1, despite what we ‘see’. From this some theorists claim that the content of the entire picture in Figure 1 is something over and above the contents of the parts of the picture, since the contents of the parts are only the four corners.

One style of objection to such a claim was raised in the last paragraph: that the method of combining these four parts into a whole includes putting them in a certain relationship to one another, namely the specific relationship that would lead us to see a square. This is

certainly a possible position for a compositionalist to take on the Gestalt figures, and so far as I can see there has never been an adequate rejoinder to this response-to-Gestalt-objections-to-compositionality. But personally, I would prefer to use the second tactic to argue against the Gestaltist argument. I would like to say that the square is *not* really there *at all*. In this regard I'm reminded of a cartoon posted on the Psychology Department bulletin board which had the corners of a Gestalt square such as Figure 1, each being held by a fireman outside of a burning building in which a woman was trapped on the second floor. The firemen were shouting 'Jump, Jump!' The title of the cartoon was 'Gestalt firemen play a joke'. The point is that even though the trapped person might *perceive* a square of safety netting being held by the Gestalt firemen, the fact of the matter is that there was *no* square there. Similarly one might say in Figure 1 that there simply is *no* square there, regardless of what we seem to think. Figure 1 then represents no challenge to The Principle because there is nothing in it other than the four corners. And similarly, no Gestalt challenge to compositionality can work because we deny that the 'inferred meanings' are there in the whole.

4. Some linguistic arguments against The Principle

In the linguistic literature there have been numerous attempts to show that compositionality is wrong, by actually presenting cases where the parts of a sentence have meanings that 'just can't' be combined to form the correct meaning of the whole.⁴

One type of challenge might be said to come about because 'there are things that we know are in the meaning of the whole that just are not in the parts'. Consider the sentence:

6. A man is in this room

and its representation:

7. $(\exists x) (\text{Man}(x) \ \& \ \text{In}(x, \text{this-room}))$

Note that in this representation there are variables which are not in the original sentence. Furthermore note that there is no 'and' in the original sentence yet (7) asserts that there is an ampersand in the representation. Thus there are things in the representation that are not in the sentence. It is not clear to me whether this kind of an

argument should be taken seriously, although it has been mounted by certain theorists who should remain nameless. I will only point out that there is no reason in the world that the composition function could not insert variables or &s.

Here is another type of example.

8. Jane isn't liked by many men.

(8) is ambiguous. It could mean either of

8a. Most men do not like Jane.

8b. There is a specific group (of many men) who do not like Jane.

So the original sentence is ambiguous and can be understood in either one of those two ways. However, if you were to embed (8) into a larger context the ambiguity disappears. Consider

9. Jane isn't liked by many men and Sally isn't either.

10. Jane isn't liked by many men and Sally isn't liked by them either.

Each of these last two sentences is unambiguous; in each of those sentences the meaning of this first clause is unique. But this first clause is just the ambiguous (8)! In (9) the phrase 'Jane isn't liked by many men' only has the meaning of (8a) whereas that very same phrase in (10) has only the (8b) meaning. At least that's the claim that some anti-compositional theorists make.

It seems to me that this is not a very successful argument against compositionality. For one thing, the meaning of the entire sentence in either of these latter two cases is conditioned in part by the second conjunct, and The Principle only predicts that the *entire* sentence should mean whatever it is that (9) and (10) mean. It does not say anything about the meaning of the component parts of those sentences. It neither has anything to say about *what* the meaning of the first clause will be as parts of (9) and (10), nor does it say what the meaning of these clauses will be in isolation, nor does it have anything to say about the relationship between the clause's meaning in isolation and its meaning in some context. So, clearly the component parts *could* be ambiguous in isolation, but in context they are not ambiguous; or rather, the entire sentence does not have an ambiguous first conjunct. So examples like (8)–(10) in no way defeat The Principle.

A similar sort of example that might be seen as more challenging is:

11. Dogs get fleas.

This sentence is a ‘generic’ statement about the propensities or the dispositions (etc.) of dogs – namely that they are able to get fleas, or that they typically get fleas, or that they get fleas under usual or normal circumstances. However, if you embed (11) into a larger context, that portion of the sentence gets interpreted differently. Consider:

12. When dogs get fleas it is best to keep them away from children.

In this case the embedded sentence, ‘dogs get fleas’, is interpreted as actual episodes or instances of dogs and getting fleas, rather than about the propensities or the dispositions of dogs. From this, some have concluded (just as they have concluded from (8)–(10)) that there has been a change of meaning of (11), brought about by the embedding into a larger context; and thus it is not solely the meaning of the component parts that goes to making the meaning of the whole, as The Principle predicts, but rather something in addition. Once again however, this is a misunderstanding of what The Principle is committed to. The Principle says only that there is some meaning which can be attributed to the entire sentence based solely on the meanings of its parts. One relevant question, then, is: what are the parts of this sentence? Well for sure, one of the parts is ‘dogs get fleas’. But equally for sure another part is that this sentence is embedded inside of a ‘when’ clause. So it is open for one to say that if a generic statement gets embedded inside of such a clause, then the entire statement has as its meaning that actual episodes of the generic disposition are being manifested. This is all that The Principle is committed to; and certainly this is within the realm of compositional theories.

Another group of examples comes when the subject term is a superlative construction (or other construction that picks out an endpoint on some scale). For example, consider the following sentences:

- 13a. *The first person* landed on the moon in 1969.
 b. *The person with the biggest grant* usually supports the rest of the department.
 c. *The tallest person* usually plays center for the team.
 d. *The first case of AIDS* was reported in 1975.

In these examples, whatever the meaning is of the italicized subject terms – whether you take it to be referential and hence it picks out the actual first person or the actual person with the biggest grant or the actual

tallest person or the actual first case of AIDS, or whether you’re more indirect and think that a phrase such as ‘the first person’ designates a function on possible worlds which in each possible world picks out the first person of that world – nonetheless, in *none* of these sentences is it this that is being talked about. For, in each of these sentences there is something about the predicate which, in one way or another, gives an idiosyncratic understanding as to how the subject should be understood. It’s *not* that the first person landed on the moon in 1969 but rather that the first person who landed on the moon did so in 1969. It’s *not* that the tallest person, whoever that might be, has the property of usually playing center; but rather that, given a team, the tallest person in that team usually plays center for the team. Challenges such as these have formed the basis for a number of attacks on The Principle.

Another attack has come from such sentences as:

- 14a. Every philosopher in the world can fit into this room.
 b. Every philosopher in New Zealand can fit into this room.

It seems quite clear, so the argument against compositionality goes, that in (14a) the meaning is ‘distributive’; that is, it says of each philosopher individually, that he or she could fit into this room. However, (14b) could have a collective meaning, in which case it would mean that all of the philosophers in New Zealand taken together could collectively fit into this room at the same time. (14a) seems to talk about the fatness of individual philosophers, (14b) seems to talk about the size or the number of philosophers in New Zealand. Once again this difference in the understanding of the sentence (whether distributive or collective) seems to be not traceable to either the verb phrase or the subject phrase, but rather is somehow dependent upon our knowledge of how many philosophers there are in the world and how many philosophers there are in New Zealand – that is to say, it relies on facts other than the meaning of the component parts.

Non-restrictive relative clauses also have been seen as forming challenges to The Principle of Semantic Compositionality. Consider, for example:

15. Kim, whose paper was rejected from *Linguistics and Philosophy*, began to rant and rave.

It seems clear that the meaning of the non-restrictive relative clause construction is more than merely an *and*.

In this case, perhaps it means *because*. However, in other cases the non-restrictive relative clause construction might mean *despite* or sometimes *although*. We see that there is no one meaning for a non-restrictive relative clause construction; and this has been seen to be a challenge for compositional semantics.

Another challenge for compositionality has been seen to reside in the adverbials which change the meaning of a given sentence in idiosyncratic ways. Consider the following four sentences:

- 16a. Laszlo has been reading Kim's diary *again*.
- b. Laszlo *even* suggested that Kim spend the weekend at his apartment.
- c. Laszlo hasn't *actually* made it with Kim *yet*.
- d. He hasn't touched her, *let alone* kissed her.

The meaning of any of those italicized adverbs is dependent on the discourse in which they are found and dependent upon world information that we each have about the kind of scenario that this discourse is likely to be describing; and it has been claimed that this has nothing to do with the lexical meaning of these individual words, and so compositionality is incorrect.

Another example along the same lines is provided by idioms and other tropes. Consider

- 17. The symbolists are really *on the ropes* now!
- 18. While I was holding Kim's baby, *it/the little sweetie* wet itself.

We see in these sentences that the literal meaning of the italicized phrases 'on the ropes' and 'the little sweetie' just are *not* what is relevant to the understanding of these sentences, and it seems that there is no compositional functional way to compute the meaning of the whole – or so the argument goes.

5. Gricean responses to challenges concerning compositionality: Some difficulties

I would like now to consider one kind of response that many compositionlists would give to such challenges as those of the last section. I only outline this response, I do not state it in any great detail. It is not, in any case, a response that I would care to give, although I think it will work in certain cases. After giving the response I will mention the reason why I think that the response cannot be used in full generality; why there seem to be cases not covered by this line of defense.

Compositionlists are tempted to use a 'Grice-y' strategy when responding to such challenges. Consider the examples (14a, b). The phrases 'Every philosopher in the world' and 'Every philosopher in New Zealand' in fact are each able to carry both the distributive and the collective meanings, the Grice response would say, and therefore each of the sentences (14a, b) are literally or really or underneath it all, ambiguous. But it is our world knowledge (our knowledge of empirical facts, not facts about language and meaning) which tells us that no *actual* room can hold all of the world's philosophers. Therefore, by Gricean mechanisms, we cannot be trying to communicate that meaning; so we must be trying to communicate the other meaning, the meaning that says of each philosopher, he or she can fit into the room – the distributive meaning. We get the *illusion* of non-ambiguity because we are not able to use one of the meanings for any reasonable communicative purpose. But the Grice position is that the sentence *does* have the meaning – it is just that it can't use it.

This is the kind of response that many compositionlists attempt to give to the previously mentioned sorts of challenges, especially the challenges concerning idioms, ambiguities and the like. But there are some problems for Grice-y explanations. Consider a sentence like:

- 19. There are green ideas.

Such a sentence *never* gets used as a simple sentence with this literal meaning. Therefore, according to a Grice-y explanation, analyzing the meaning of a sentence like

- 20. Laszlo would never believe that there are green ideas

will always fail. If you were faced with (2), then in accordance with the Grice-y explanation you would say that the sub-sentence 'there are green ideas' can't possibly be used with its literal meaning. So it must be used with some other meaning; and now it is our job to figure out this other meaning (using other Gricean mechanisms). And it is that other, discovered meaning which will turn out to be what Laszlo never believes – contrary to what our intuitions say about sentence (20).

Another reason to believe that there is something wrong with Gricean explanations in general, is that we *do* understand outrageous and absurd tales. Consider the following (cited in Hirst, 1987):

21. Go on, have your fun, it's always the children that suffer later: Los Angeles secretary Jannine Swift married a fifty pound rock in a formal ceremony in Lafayette Park yesterday.

We understand this sentence entirely in its literal meaning despite the fact that it is never the case that the verb phrase 'marries a fifty pound rock' can have its literal meaning. (Certainly in none of our lives have we encountered a sentence which truthfully and literally asserts that a person has married a rock; yet we, all of us, understood that sentence.) If the Gricean mechanisms were right in their full generality we would never be able to understand a *National Enquirer* headline. But we do, and that's why the newspapers sell so well.

6. The argument from synonymy

The source of challenges to compositionality that we have so far been considering, all have turned on an expression's meaning something X_1 (which might be ambiguous, or even null) in context C_1 and meaning some other X_2 (perhaps a resolution of ambiguity or a creation of new meaning or a change to some related meaning, etc.) in context C_2 . I myself am not very taken with these arguments; I think all of them can be overcome one way or another and some of them are even downright silly. But I shall not myself attempt to refute the arguments; I'll leave it to the reader to determine what he or she thinks about the correctness of the Gricean response or some other way of understanding compositional functions. I think there are other arguments that show that compositionality is not a very attractive prospect. So let us now turn to a different strategy for arguing against compositionality. A fact about semantic compositionality that has not often been noticed is the following:

If there are two different expressions which have the same meaning (i.e. if there is any synonymy in the language), then if there is any case where embedding one of these expressions into a context yields a different meaning than embedding the other expression into that same context, then semantic compositionality is false.

(This would be a case where expressions X and Y mean the same but $C(X)$ means something different from $C(Y)$, where C is some linguistic context.) The reason

that this shows that semantic compositionality would be false is that, *ex hypothesi*, the two expressions X and Y have the same meaning and therefore they contribute the same thing in whichever context they occur. If they occur in the same context then the entire constructions must mean the same thing, according to semantic compositionality. So if there is ever a case where two different expressions mean the same thing, but when they are embedded the two wholes mean something different, then compositionality is false. Is there any such case? Consider the following argument.⁵ Given these three assumption:

- (A) If ϕ and ψ are sentences that have the same meaning, then they have the same truth value.
- (B) For a given syntactic theory, there is only one rule or sequence of rules which creates or analyzes sentences of the form
Kim + believes + that + Sentence
- (C) If ϕ and ψ are syntactically distinct sentences then it is possible that exactly one of (i) and (ii) is true.
 - (i) Kim believes that ϕ .
 - (ii) Kim believes that ψ .

Then, if there are any synonymous ϕ and ψ , The Principle has to be false. The argument for this conclusion is quite simple. Suppose that there are synonymous sentences S_1 and S_2 . Then by assumption (B), it is the same rule or sequence of rules that analyzes both of

- a. Kim believes that S_1
- b. Kim believes that S_2

and therefore these two sentences, (a) and (b), mean the same thing – according to The Principle. But by assumption (C), it is always possible that sentence (a) is true and sentence (b) is false; and hence by assumption (A), sentences (a) and (b) do *not* mean the same thing. Therefore, given assumptions (A)–(C), if there is any synonymy of sentences, The Principle of Semantic Compositionality is false.⁶ If we believe in the existence of sentential synonymy and we believe these three assumptions, then we cannot believe The Principle of Semantic Compositionality. This strikes me as a rather powerful argument against The Principle, for surely these assumptions and the existence of synonymy are more plausible than the highly theoretical and methodologically-motivated Principle.

7. Arguments in favor of semantic compositionality

So far as I have been able to tell, there are only four considerations that have been brought to bear in favor of compositionality in the literature. This is not to say of course that there are only a few people who believe in semantic compositionality. In fact, as I indicated earlier, almost everyone has a ‘warm and fuzzy feeling’ towards compositionality. Almost everyone assumes it to be true. Almost everyone assumes that it is a desideratum of any adequate theory. Yet when one looks for arguments in favor of it, one finds very few. Here are the four that I have found.

Argument 1: Compositionality is the only way to get an account of such semantic notions as truth, validity and inference, etc.

Argument 2: Semantics is a mirror of our cognitive states. Our cognitive states are compositional, and therefore, semantics must be compositional.

Argument 3: If a language lacked compositionality it would be *unlearnable*.

Argument 4: Compositionality is the only explanation of how a finite mechanism (such as the human brain/mind) can *understand* an infinite set of sentences.⁷ (Without compositionality, novel utterances would be non-*understandable*.)

These arguments seem to me to differ in their strength. The first argument merely asserts a falsehood: it is just not true that compositionality is the only way to get an account of these semantic notions. Indeed, the traditional Tarski notion defining validity and truth in terms of satisfiability is not compositional in *any* straightforward sense.⁸ The second argument seems to me to beg the question. We have even less knowledge of the status of our cognitive states – whether they are compositional or not – than we do about semantics, and so it seems rather premature to assert that we can use facts about cognitive states to prove facts about semantics.

The third and fourth of these arguments seem to me to be strongest. If the meaning of the whole were not a function of the meaning of its parts, so these arguments say, then we would not be able to *learn* the language. We would not be able to *understand* all of the language. How else, so these arguments ask us, are we to be able

to figure out the meaning of an arbitrary, new, novel sentence if it isn't by the fact that we've learned some finite numbers of parts and finite number of ways of putting them together. How is it that we can understand a novel sentence, except by predicting its meaning by our understanding of the meaning of its parts, so these arguments say. Towards the end of this paper I will try to sketch a different conception of semantics that allows languages to be learnable and understandable, even though they are not, strictly speaking, compositional.

8. Language and representation

Figure 2 gives us a picture of the organization of the semantic component of a linguistic theory. It is a picture that is very familiar in philosophy of language and also in many different linguistic theories.

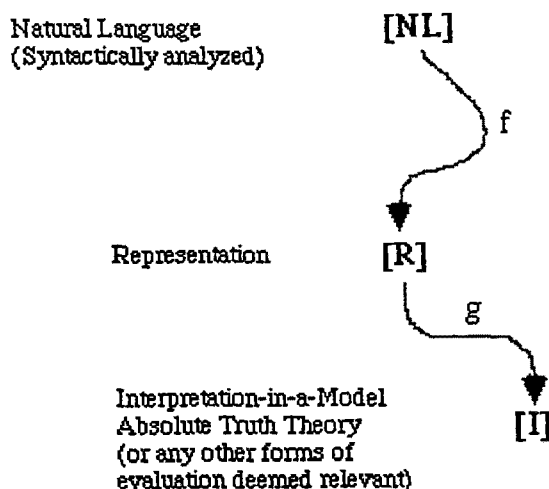


Fig. 2. The structure of the semantic component of a theory of language, according to certain views.

We see here a rather standard view, where the semantics is defined in two stages. One starts with a syntactically analyzed sentence (or other fragment of natural language) and one determines a representation of this piece of natural language, and from that one generates an interpretation of the representation. For example, one might take an individual sentence of English and represent this by means of a sentence of first order logic and then interpret first order logic in terms of Tarski-style truth conditions.

Some theorists prefer to take a short-cut in this method. These are the Eliminativists depicted in Figure 3.



Fig. 3. Eliminativism.

Some Eliminativists start with the theory given in Figure 2, but then take each of the f and g to be a function. Hence the representation [R] in the middle of Figure 2 is theoretically dispensable: one could just use function composition on the two functions and have a generated function of $f \circ g$. So the level of representation [R] is dispensed with and we are left with Figure 3. Montague, for example, took this view. In his (1973), the level of representation was Intensional Logic but he held that this level was theoretically dispensable and that we could eliminate the representation level of Intensional Logic. (This was laid out in his 1970.)

Another way of eliminating one of the levels is illustrated in Figure 4. Here the theory might be called Straight Representationalism. The idea is that there is *no* interpretation of natural language over and above the level of representation. The kind of theorist that have this picture are cognitive grammarians, generally speaking, for instance Jackendoff (1985) and Langacker (1980), to name a few. These theorists deny the necessity of having *any* 'external' interpretation of [R]; indeed they view the translation (as we might call it) of natural language into some other representational language as itself sufficient for us to understand the meaning of natural language. In this belief these theorists reject the arguments in Lewis (1972) to the effect that translation into another language (the 'Language of Thought' or into 'Markerese') just isn't semantics in any real sense.



Fig. 4. Straight representationalism.

It should be noted that both Eliminativism and Straight Representationalism have exactly the same structure in that they are two-level theories wherein the levels are related by functions. Therefore they make the same sort of predictions for certain semantic data. In particular, for either one of these views to maintain compositionality (that is, for either one of these views to hold that the relationship between the natural language and whatever it gets mapped into – whether

it's an interpretation as in Eliminativism or just a representation as in Straight Representationalism – is a function), these theories must *deny* the existence of any 'serious ambiguity'. (A 'serious ambiguity' occurs when a sentence of natural language has exactly one syntactic analysis but has two meanings, and this ambiguity is not traceable to a lexical ambiguity.) For, if the relationship between these two levels were a function, as compositionality demands, then there is no opportunity for 'serious ambiguity' to arise. For compositionality to be true in these theories, all ambiguity must be 'minor': either traceable to lexical ambiguity or else to differing syntactic analyses of the same English string.

Another view of the relationship between natural language and the interpretation is illustrated in Figure 5. In Crooked Representationalism, a syntactically analyzed sentence of natural language is associated with a level of representation, R_1 , and then, in some way or other, we are allowed to manipulate this representation and 'do various things to it' until we reach a final representation, R_n , which is then interpreted. This 'doing something' might be to bring in context, to resolve anaphora, to eliminate ambiguities, and to draw inferences (possibly from other sentences, possibly from the background theory, possibly from prototype theory and possibly from other areas). In general these 'doings' will not be functional, for they will not be unique. If they were unique then all of these intermediate levels between R_1 and R_n could be theoretically dispensed with; but most theorists believe that they are not theoretically dispensable, and that they are essential to the interpretation of the sentence. As examples of theorists who have held such views I might point to, for instance, Schubert and Pelletier (1982). In this theory the relationship between the natural language and R_1 was functional, that is R_1 was compositionally determined on the basis of the syntactic features of a natural




Fig. 5. Crooked representationalism.

language sentence. However, then various ‘doings’ were alleged to happen, such as scope ambiguity resolution and anaphora resolution, and various inferences were drawn. None of this was envisaged as being functional, that is to say, it was possible that one and the same level of representation might have different things happen to it at different times, and hence it was rather unpredictable as to what the outcome of a level R_i becoming R_{i+1} would be.

Having mentioned Crooked Representationalism we are in a position to give yet another version of Eliminativism, which, for definiteness, we call the GB Strategy. Essentially what the GB Strategy does is take the Crooked Representationist point of view about everything, except for where the boundary is between analyzing a natural language utterance and doing the semantics on that utterance. In Crooked Representationalism, as we can see from Figure 5, the boundary is at the point where we have a syntactically analyzed surface structure of natural language. The rest of the diagram is viewed as part of the semantic theory for understanding the sentence. However the GB Strategy usurps all of the diagram up to R_n as part of the syntactic theory. After all of these ‘doing somethings’ (in GB Theory this is applying the Quantifier Raising rules, for example) one arrives at a new syntactic level which is called LF. In some of the advanced versions of this theory, for example Higginbotham (1987), it is possible to interpret these LF structures in a straightforwardly compositional way. So we can see that the GB strategy, at least as augmented by a real semantics *à la* Higginbotham, is essentially the same as Eliminativism. This reveals a hitherto unnoticed close affinity between Montague and Chomsky.




Fig. 6. The GB strategy.

9. The argument from ambiguity

Once again let me emphasize a point that was made earlier: that any theory of either the Eliminativist type, the Straight Representationist type, or the GB Strategy type is committed to denying one of The Principle of Semantic Compositionality or the existence of ‘serious ambiguity’.⁹ But can it really be true that there is *no* serious ambiguity? How about the following sentences?

22. Every linguist knows two languages.
23. John wondered when Alice said she would leave.
24. When Alice rode a bicycle, she went to school.
25. The philosophers lifted the piano.
26. The Canadian family used less water last year than the preceding year.

These sentences all are ambiguous and all ambiguous in the way that I mentioned before, namely they have only one syntactic analysis, but have at least two meanings. In (22) there is the obvious ambiguity between every linguist knowing the same two languages or whether merely that every linguist knows two possibly different languages. Now, it is traditional to analyze such sentences as containing scope ambiguities, that is to say, as having two different semantic analyses, one where *every* has wide scope and one where *two* has wide scope. But doesn’t it seem really implausible to say that (22) has two different *syntactic* analyses?

Sentence (23) clearly only has one syntactic analysis, but there are two different meanings. One meaning is that John wonders when Alice said something and the other meaning is that John wonders when Alice is going to leave. (24) is ambiguous between stating (i) that on those occasions when Alice rode a bicycle, she rode it to school, and (ii) that back in the times when Alice had the disposition to be a bicycle rider, she also had the disposition or propensity to go to school (i.e., she was also a student). (25) manifests a familiar distributive/collective ambiguity about whether it was all of the philosophers together who lifted the piano or whether they each did it separately. And (26) is ambiguous between whether it is the average Canadian family that used less water last year than the preceding year, or whether it’s the total of Canadian family water usage that is less this year than the preceding year. (Contrast ‘The Canadian family bought 9500 BMWs last year’ and ‘The Canadian family owns 2.3 TV sets’). These two can have different truth conditions depending on how many Canadian families there are.

Besides the abovementioned strategy of trying to attribute different syntactic analyses to account for the scope ambiguities, there are other syntactic strategies that have been used in trying to account for some of these ambiguities. For instance, Montague had ‘qualifying in’ rules which would analyze the first of the sentences differently in the two meanings depending on which order the ‘quantifying in’ rules were applied. And Cooper (1983) had a mechanism for storing quantified noun phrases so as to allow different analyses of their scopes. Other strategies for some of these sentences might involve ‘traces’ and ‘gaps’. It seems to me that all of these are rather desperate measures in that they try to invent a syntactic ambiguity when we know perfectly well that in reality there is no syntactic ambiguity. Admittedly there is a *semantic* ambiguity, but the only reason that this fact should call for a syntactic ambiguity is if you antecedently believe The Principle. I mentioned in the last section that compositionists must find all ambiguities to be ‘minor’. The syntactic strategies suggested in this paragraph take the route of saying that the type of minor ambiguity to be found is the one where the same string has differing syntactic analyses. But isn’t that really silly for these examples? My recommendation is: don’t force yourself to accept spurious syntactic ambiguities.

Another strategy that is often applied to some of the other sentences is to invent semantic features (as opposed to the syntactic features mentioned in the last paragraph) that are involved with some of the noun phrases. For example, one might have semantic Distributivity vs. Collectivity features and claim that the verb *lifted* is itself an ambiguous verb. One might similarly try to invent semantic Genericity features to account for the two different meanings of (24) and claim that *when* is ambiguous. The general methodology of this semantic strategy to defend compositionality is to describe the two different meanings by means of some semantic features, associating them with some phrase, and to claim that some lexical item in the sentence is ambiguous between a meaning using one vs. using the other of these features. I think this semantic strategy leaves a lot to be desired. True enough, these sentences are semantically ambiguous and if you wish to use such features to describe what that ambiguity is, well and good. But there is no independent reason to say that these features are somehow inherent in the literal meanings of the lexical items themselves. It is unjustified and only motivated by a desire to maintain semantic

compositionality. In fact there is a hidden difficulty with the whole semantic enterprise. Consider the collective vs. distributive meanings of (25). It is not enough simply that *lifted* be viewed as ambiguous. In addition, the subject term, *the philosophers*, has to also be ambiguous. But where did that ambiguity come from? Are we to posit an underlying distributive/collective ambiguity in every lexical noun? (It has to be in the lexicon, for otherwise there would be a violation of compositionality in forming the meaning of the noun phrases.) Better to give up the semantic strategy.

10. Synonymy again

Recall from before that The Principle of Semantic Compositionality is incompatible with any form of synonymy. Can it really be true there is no synonymy? Can it really be true that *attorney* and *lawyer* are not synonymous, or if you think that those are not quite synonymous then that there are *no* lexical items that are synonymous? What about phrases? Can you really deny that *a circle* and *a locus of all points on a plane equidistant from a given point* are synonymous? What about sentences? Can you really deny that *Dentists usually need to hire an attorney* and *Tooth doctors commonly require the services of a lawyer* are synonymous? Isn’t it better to give up The Principle of Semantic Compositionality? Isn’t it better to just face the fact that, ‘warm and fuzzy’ though it may be, semantic compositionality is just inconsistent with any reasonable interpretation of linguistic data?

But, I hear you say, if there really is synonymy, if there really is serious ambiguity, and therefore semantic compositionality is false, then how is it possible to learn language? How is it possible to understand an infinite (or hugely finite) number of these linguistic structures? This is the big question for non-compositional semantic theories. To answer this question will take a bit of a detour. I don’t intend to answer it in great detail but I do hope to give a feeling for why I think that these difficulties can be overcome.

11. The Principle of semantic groundedness

The Principle of Semantic Compositionality is closely related to inductive definitions, a general form of which is:

Definition of $\phi(n)$:

$$\begin{aligned} \phi(n) &=_{df} a, b, c \dots, \\ &\quad \text{if } n = 0, 1, 2 \dots \quad \} \text{ basis clauses} \\ &= \psi(\phi(n-1)), \\ &\quad \text{otherwise} \quad \} \text{ inductive clause} \end{aligned}$$

Some examples of inductive definitions are:

$$\begin{aligned} n! &=_{df} 1, \quad \text{if } n = 0 \\ &= (n-1)! * n, \quad \text{otherwise.} \\ n^m &=_{df} 1, \quad \text{if either } n = 1 \text{ or } m = 0 \\ &= n^{(m-1)} * n, \quad \text{otherwise} \end{aligned}$$

The general feature of these inductive definitions in defining some operation ϕ is that after one has done some basic cases (the so called ‘basis clauses’) one defines the operation ϕ on the arbitrary number n in terms of that very same operation on some lesser number(s). (It is the very same operation ϕ on the lesser number, but then we ‘do something’ to this result. In the abstract definition of $\phi(n)$ this ‘doing something’ was represented as ψ . In the examples of defining the factorial and exponentiation, the ‘doing something’ was to multiply by n .) The reason that these are all called inductive definitions is that they apply the concept being defined to a number less than the number that is being defined. But in fact, there is no logical necessity for this clause to be an operation on the *next* smaller number ($n-1$) or indeed on *any* number which is less than n . Think for example of programming a recursive procedure in some suitable programming language. There is no syntactic requirement that the recursive clause be a function of some number less than n . All that’s required is that the function doing the evaluation ultimately be grounded for any value we wish to evaluate. For example,

$$\begin{aligned} \phi(n) &= 2, \quad \text{if } n = 0 \\ &= 3, \quad \text{if } n = 1 \\ \text{else} \\ &= n * \phi(n-2), \quad \text{if } n \text{ is odd} \\ &= n * \phi(2n+1), \quad \text{if } n \text{ is even.} \end{aligned}$$

Note in this definition that the last clause defines $\phi(n)$ in terms of $\phi(2n+1)$, and $(2n+1)$ is of course larger than n . Therefore, this definition is not an inductive definition of the sort described above. We might instead call this type of definition by a different name; let’s use ‘recursive definition’.

A problem with these non-inductive, recursive defi-

nitions (which isn’t a problem with just the pure inductive definitions) is that they are not always ‘grounded’. For example

$$\begin{aligned} \phi(n) &=_{df} 1, \quad \text{if } n = 1 \\ \text{else} \\ &= 1 + \phi\left(\frac{n}{2}\right), \quad \text{if } n \text{ is even} \\ &= 1 + \phi(3n-1), \quad \text{if } n \text{ is odd.} \end{aligned}$$

(The last clause is a non-inductive form of definition, where ϕ is being applied to $3n-1$.) I show here how this series evaluates n for values from one to five.

$$\begin{aligned} \phi(1) &= 1 \\ \phi(2) &= 1 + \phi(1) = 1 + 1 = 2 \\ \phi(3) &= 1 + \phi(8) = 2 + \phi(4) = 3 + \phi(2) = 5 \\ \phi(4) &= 1 + \phi(2) = 1 + 2 = 3 \\ \phi(5) &= 1 + \phi(14) = 2 + \phi(7) = 3 + \phi(20) \\ &= 4 + \phi(10) = 5 + \phi(5) \end{aligned}$$

Note, in the evaluation of $\phi(3)$, that we needed to discover the value of $\phi(8)$; yet there was no problem here, this was evaluated without any difficulties. Yet, when we did the same thing with $\phi(5)$, it called $\phi(14)$ which called $\phi(7)$ which called $\phi(20)$ which called $\phi(10)$, and this finally called $\phi(5)$ again and thus generated a circle, a ‘non-grounded’ evaluation. Inductive definitions are always grounded because of the structure of the integers; however, these arbitrary recursive definitions are not always grounded. Sometimes, as here, this is a difficulty. But at other times there is no difficulty, as the following example shows.

$$\begin{aligned} \phi(n) &=_{df} 0, \quad \text{if } n = 0 \\ \text{else} \\ &= 0, \quad \text{if } m \text{ is even} \\ &= 1 + \phi(n+1), \quad \text{otherwise} \end{aligned}$$

$$\begin{aligned} [\text{note: } \phi(0) &= \phi(2) = \phi(4) = \dots = 0 \\ \phi(1) &= \phi(3) = \phi(5) = \dots = 1] \end{aligned}$$

This function is recursive but non-inductive, as can be seen from the last clause of the definition. However, it can be seen from the evaluation that $\phi(n)$ is just a perfectly good alternative way of defining $(n \bmod 2)$. There is no ungroundedness in this definition at all.

It is this notion of non-inductive but recursive definitions that I think lends a key to the understanding of how there might be a non-compositional semantics. Let’s take a look at an example first put forward by

Kaplan (1968) (under considerably different circumstances to make a considerably different point). The example is from elementary semantics of the sentence logic in which each of the connectives $\{\sim, \rightarrow, \leftrightarrow, \&, \vee\}$ are viewed as primitives and not defined by the others. We are going to construct an assignment of truth values function. The standard way of doing this, the way that you can always find in any textbook, is the following:

Let f be an assignment of truth values to atomic sentences.

$$\begin{aligned}
 f^*(\phi) &=_{df} : \\
 f^*(p) &= f(p), \quad \text{if } p \text{ is atomic} \\
 f^*(\neg\phi) &= 1, \quad \text{if } f^*(\phi) = 0 \\
 f^*(\phi \rightarrow \psi) &= 1, \quad \text{if either } f^*(\phi) = 0 \text{ or } f^*(\psi) = 1 \\
 &= 0, \quad \text{if both } f^*(\phi) = 1 \text{ and } f^*(\psi) = 0 \\
 f^*(\phi \& \psi) &= 1, \quad \text{if both } f^*(\phi) = 1 \text{ and } f^*(\psi) = 1 \\
 &= 0, \quad \text{if either } f^*(\phi) = 0 \text{ or } f^*(\psi) = 0 \\
 f^*(\phi \vee \psi) &= 1, \quad \text{if either } f^*(\phi) = 1 \text{ or } f^*(\psi) = 1 \\
 &= 0, \quad \text{if both } f^*(\phi) = 0 \text{ and } f^*(\psi) = 0 \\
 f^*(\phi \leftrightarrow \psi) &= 1, \quad \text{if } f^*(\phi) = f^*(\psi) \\
 &= 0, \quad \text{if } f^*(\phi) \neq f^*(\psi)
 \end{aligned}$$

This definition of f^* is completely compositional. Furthermore, it is completely inductive. One defines f^* applied to any formula in terms of how it is applied to the parts of that formula. Now, it seems completely clear that we could replace the very last clause of that definition by the following clause:

$$f^*(\phi \leftrightarrow \psi) = f^*((\phi \rightarrow \psi) \& (\psi \rightarrow \phi))$$

Here we have defined the f^* when applied to a certain category of formulas (the \leftrightarrow formulas), in terms of something entirely different – in terms of an entirely different formula that has only a very tenuous relationship, syntactically speaking, to the \leftrightarrow formula. (Recall that they were all primitive connectives.) Still, we all know that this is a correct definition, and that it will work. This new definition, as can be seen, is non-compositional. It defines the \leftrightarrow in terms of things that are not part of the double arrow. Similarly, it is not inductive in the sense of defining things on the basis of parts.

It is this picture of semantics that I wish to urge upon the reader. A semantic evaluation, in general, can bring into play all kinds of facts, all kinds of information; it could bring in context, it could bring in inferences, it could bring in world knowledge to evaluate an expression, where these facts, etc. are not part of the meanings

of the parts of the expression (and they are furthermore not dictated by the ‘method of combination’ used to construct the expression). So long as this evaluation is always grounded – or perhaps more weakly, that it’s grounded in the cases in which we actually employ it – then this will serve just as well as any compositional analysis. And as I have argued before, it is more in accord with the facts of language. It allows for ambiguity and it allows for synonymy, unlike semantic compositionality.

12. Why do people cling to compositionality?

In closing, I think that the apparent distaste that’s manifested in the world of semantics for any denial of The Principle of Semantic Compositionality is due to a combination of father worship and disgust at certain theorists. The fact of the matter is that a certain group of us all admired the rigor, the clarity and tough-mindedness that surrounded the initial appearance of Montague Grammar. And didn’t we all sneer at those who had logic anxiety? And, when we look at the loud voices of non-compositionality in the wider profession, we find a number of researchers who are opposed to doing anything with rigor, doing anything tough-mindedly. And we say to ourselves that we would never want to be identified with them. And so we continue to hold on to The Principle of Semantic Compositionality. But these are not worthy reasons to adopt The Principle. Once it is recognized that what we really want is to avoid these other people’s commitment to anti-formalism, and we recognize that this is separate from their opposition to compositionality, then I think that any fear of being mistaken for one of these anti-theoretical, flaky researchers will subside. That, together with giving up not our love but rather our blind adulation of Father Montague, will help us cheerfully embrace non-compositionality for the obvious truth that it is.

Notes

¹ The Principle is often stated incorrectly. Although their intent is to express what I’ve given in the text, many writers use the phrase ‘is a function of the meaning of its parts’. They intend to say that the meanings of the parts actually take a role, indeed the *only* role, in determining of the whole. However, it is infelicitous to use the phrase ‘is a function of’ to express this. Rather, what we should say is ‘is *only* a function of’. Otherwise we could allow the function to invoke other items than the meanings of the parts, and we might allow

the meaning of the parts to appear vacuously as mere 'dummy arguments' of the function. The intent of The Principle is that the meanings of the parts are *real arguments* of the function that determines the meaning of the whole. (Another question arises as to whether the function is to be total; that is, whether *every* meaning of parts can give rise to a meaning of wholes. It seems pretty clear that it is implausible, and that it is intended that this function be partial. The issue of partiality will not concern us in this paper.)

² The Principle is often said to trace back to Frege, and indeed many textbooks call The Principle of Semantic Compositionality 'Frege's Principle'. However, it is extraordinarily difficult to find The Principle in Frege. As Cresswell (1973: 75fn) says, "it is more of a tribute to the general tenor of his views on the analysis of language" that we attribute The Principle to him. It seems to me, though, that even this much is false; for what we find in Frege, rather than The Principle of Semantic Compositionality, is instead *The Principle of Contextuality*. This latter principle can be found in Frege's earliest writings (namely Frege, 1884) and it says "it is only in the context of a sentence that a word has meaning". It is rather difficult to see how this sympathy can be combined with The Principle of Semantic Compositionality, although various writers (for example Dummett, 1974) have attempted to do so. On a separate occasion I hope to trace the real history of The Principle of Semantic Compositionality.

³ Indeed this is argued for by Zadrozny (1992).

⁴ I shall not attribute any of these arguments to anyone in particular. If I were to do so I would have to be much more circumspect in my comments than I care to be. Careful research would find all of them in the literature, however. (But I *am* pleased to acknowledge taking examples from Hirst (1987), and I agree with most of what he has to say about them.)

⁵ This argument is elaborated at more detail in Pelletier (1994).

⁶ In Pelletier (1994) I used similar arguments to show that The Principle of Semantic Compositionality is also inconsistent with phrasal synonymy and with lexical synonymy.

⁷ Actually, the argument need not be stated in terms of an infinite set of sentences; as Grandy (1991) has pointed out, the argument seems to work even if we just presume a hugely large but finite set of sentences.

⁸ In fact, Tarski's favorite way of defining these notions is not so well known – it depends on the notion of cylindrical algebras, and in these the notions *are* defined compositionally. (Thanks to David Israel for pointing this out to me.) But, as I say, the traditional account attributed to Tarski, where 'satisfaction' is taken as the primitive notion, is not compositional.

⁹ Once again, let me remark on what I mean by a serious ambiguity here. It is where one and the same syntactic structure, using identical basic parts, can have two different meanings. This is different from sentences like *Visiting professors can be fun* which arguably has two different syntactic analyses. With two different syntactic analyses, of course a compositional theory can assign two different meanings. But what I mean here is a sentence which is *not* syntactically ambiguous, but is nonetheless semantically ambiguous without the ambiguity being traceable to any lexical item.

References

Allan, K.: 1986, *Linguistic Meaning*, Vol. I, London: Routledge & Kegan Paul.

- Avramides, A.: 1989, *Meaning and Mind*, Cambridge: MIT Press.
- Cooper, R.: 1983, *Quantification and Syntactic Theory*, Dordrecht: D. Reidel.
- Cresswell, M.: 1973, *Logics and Languages*, London: Methuen.
- Dummett, M.: 1981, *Frege: Philosophy of Language*, 2nd ed.: Cambridge: Harvard UP.
- Field, H.: 1978, 'Mental Representation', *Erkenntnis* 13, 9–61.
- Frege, G.: 1884, *Die Grundlagen der Arithmetik* translated as *Foundations of Arithmetic* by J. L. Austin, NY: Bantam Books.
- Grandy, R.: 1990, 'Understanding and Compositionality', *Philosophical Perspectives* 4, 557–572.
- Grice, P.: 1975, 'Logic and Conversation' in D. Davidson and G. Harman (eds.), *The Logic of Grammar*, Encino: Dickenson Publishers.
- Haugeland, J.: 1987, *Artificial Intelligence: The Very Idea*, Cambridge: MIT Press.
- Higginbotham, J.: 1987, 'On Semantics' in LePore (Ed.), *New Directions in Semantics*, London: Academic Press, pp. 1–54.
- Hirst, G.: 1987, *Semantic Interpretation and the Resolution of Ambiguity*, Cambridge: CUP.
- Jackendoff, R.: 1987, *Consciousness and the Computational Mind*, Cambridge: MIT Press.
- Janssen, T.: 1983, *Foundation and Applications of Montague Grammar*, Ph.D. Dissertation, University of Amsterdam.
- Kamp, H.: 1990, 'Response to Groenendijk and Stockhof', unpublished manuscript, University of Stuttgart.
- Kaplan, D.: 1968, 'What is Russell's Theory of Descriptions?' reprinted in D. Pears ed.: *The Philosophy of Bertrand Russell*, Oxford: Oxford UP.
- Lakoff, G.: 1987, *Women, Fire and Dangerous Things*, Chicago: University of Chicago Press.
- Langacker, R.: 1987, *Foundations of Cognitive Grammar*, Vol. 1, Stanford: Stanford UP.
- Lewis, D.: 1972, 'General Semantics' in D. Davidson and G. Harman (Eds.), *Semantics of Natural Language*, Dordrecht: D. Reidel.
- Montague, R.: 1970, 'Universal Grammar', in Thomason (1974), pp. 222–246.
- Montague, R.: 1973, 'The Proper Treatment of Quantification in English', in Thomason (1974), pp. 247–270.
- Pelletier, F. J.: 1994, 'On an Argument against Semantic Compositionality' in D. Prawitz and D. Westerstaal (Eds.), *Logic, Methodology and Philosophy of Science*, Dordrecht: Kluwer.
- Schubert, L. K. and Pelletier, F. J.: 1982, 'From English to Logic: Context-free Computation of "Conventional" Logic Translation', *Journal of Computational Linguistics* 8, 27–44.
- Smith, B.: 1987, 'Philosophy 266: Representations, Formality, and the Foundations of Computation', unpublished manuscript, Stanford Univ.
- Thomason, R.: 1974, *Formal Philosophy*, New Haven: Yale Univ. Press.
- Zadrozny, W.: 1992, 'On Compositional Semantics', *Proceedings of COLING-92*, pp. 260–266.

*Departments of Philosophy,
Computing Science
University of Alberta
Edmonton, Alberta T6G 2E1
Canada*