

3. Ontological Semantics and the Study of Meaning in Linguistics, Philosophy and Computational Linguistics

This chapter contains a very brief survey of the study of meaning in linguistics and philosophy (see Raskin 1983 and 1986 for a more detailed discussion). Its purpose is limited to placing ontological semantics in the realm of linguistic and philosophical semantics.

3.1 Prehistory of semantics

Before the study of meaning emerged as a separate linguistic discipline in the late 19th century, a number of disjoint ideas about meaning had accumulated over the millennia. For instance, Plato's "Kratylos" is devoted essentially to a discussion about whether words are natural and necessary expressions of notions underlying them or merely arbitrary and conventional signs for these notions, that might be equally well expressed by any other collection of sounds. The closely related problem of sound symbolism has recurred ever since. In modern times, de Saussure's (1916), Jakobson's (1965) and Benveniste's (1939) debate on the arbitrariness of the linguistic sign develops the same issue. The currently active area of word sense disambiguation can be traced back at least to Democritus who commented on the existence of polysemy and synonymy (1717; cf. Lurfler 1970). Modern work on diachronic changes in word meaning was anticipated by Proclus (1987, 1989). Aristotle (1968) contributed to the definition of what we would now call the distinction between open- and closed-class lexical items, a taxonomy of parts of speech and another one for metaphors (or tropes).

An ancient Indian (see, for instance, Zvegintzev 1964) school of linguistic thought was preoccupied with the question of whether the word possesses a meaning in isolation or acquires it only in a sentence. This argument was taken up by Gardiner (1951) and Grice (1957). Practical work with meaning can be traced back to the Middle Ages and the trailblazing lexicographic and thesaurus-building work by Arab scholars (see, for instance, Zvegintsev 1958).

3.2 Diachrony of word meaning

In 1883, a French classical philologist Michel Bréal (1832-1915) published an article (see Bréal 1997) which contained the following passage: "The study where we invite the reader to follow us is of such a new kind that it has not even yet been given a name. Indeed, it is on the body and the form of words that most linguists have exercised their acumen: the laws governing changes in meaning, the choice of new expressions, the birth and death of idioms, have been left in the dark or have only been casually indicated. Since this study, no less than phonetics and morphology, deserves to have a name, we shall call it semantics (from the Greek verb σημαίνειν 'to signify'), i.e., 'the science of meaning.'"³⁸

Semantics was thus originally established as a historical discipline. This was not surprising in the

38. The word 'semantics' had, in fact, existed before. In the seventeenth century it was used by philosophers to denote 'the science of prediction of Fate on the basis of weather signs.' Larousse's French dictionary defined 'sémantique' only a century ago as a science of directing troops with the help of signals. See Read (1948) for more information on the term.

post-Darwin era, when the historical approach was dominant in science. What Bréal, Hermann Paul (1886), and Arsène Darmesteter (1887) initiated, and what was later continued by Wundt (1921), Meillet (1922), Wellander (1973) and Sperber (1958), was: studying changes of meaning, exploring their causes, classifying them according to logical, psychological and/or other criteria, and, if possible, formulating the general ‘laws’ and tendencies underlying such changes. The examples below illustrate the types of phenomena discussed by Bréal and his colleagues.

Table 3: Examples of Meaning Change

Type of Change	Language	Word	Old meaning	New meaning
Restriction	Latin	felix	female of any animal	pussycat
	Latin	fenum	produce	hay
	Greek		possessions	cattle
	German	Mut(h)	soul, intelligence	courage
	English	meat	food	meat
Expansion	French	gain	harvest	produce, result
	French	temps	temperature	weather
	French	briller	beryl	shine
	English	dog	dachshund	dog
Metaphor	Latin	putare	count	think
	Latin	aestimare	weigh the money	evaluate
	English	bead	prayer	bead
Concretion	Latin	vestis	the action of dressing	vest, jacket
	Latin	fructus	enjoyment	fruit
	Latin	mansio	stopping	mansion
	English	make love	court	have sex
Abstraction	Latin	Caesar	Caesar	caesar, emperor
	English	Bismarck	Bismarck	great statesman

Bréal (1897) was also the first to introduce what we would now call lexical rules (“laws,” in his terminology). Thus, he talks about the diachronic law of specialization (lexicalization or degrammaticalization, in current terminology), according to which words undergo change from synthetic to analytical expression of grammatical meaning, e.g., Latin: *fortior* > French: *plus fort*. Bréal’s

law of differentiation says that synonyms tend to differentiate their meaning diachronically: thus, the Swiss French *païle* changed its neutral meaning of 'room' for that of 'garret,' after the French *chambre* had ousted it. The law of irradiation (analogy, in modern terms) deals with cases when an element of a word assumes some component of the word's meaning and then brings this meaning over to other words in which it occurs: e.g., the Latin suffix *-sco* acquired its inchoative ('beginning') meaning in such words as *adolesco*, 'to grow up to maturity,' and later irradiated that meaning into *maturesco*, 'to ripen,' or *marcesco*, 'to begin to droop' (in a contemporary American English example, *-gate* acquired the meaning of 'scandal' in *Watergate* and contributed this meaning to many other names of scandals, e.g., *Koreagate* or *Monicagate*).

3.3 Meaning and reference.

The next major question that interested semanticists was the relation between word meaning and the real world (that is, the entities to which words referred). The distinction between meaning and reference was introduced in logic by Frege (1892).³⁹ To illustrate the difference between meaning and reference, Frege used the following example: the expressions *Morning Star* and *Evening Star* have a different meaning (stars appearing in the morning and evening, respectively) but refer to the same entity in the world, the planet Venus.

The distinction was introduced into linguistic semantics by Ogden and Richards (1923) who presented it as the triangle:




Figure 17. Ogden and Richards' original word meaning triangle. A language symbol (a word) does not directly connect with its referent in the world. This connection is indirect, through a mental representation of the element of the world.

According to Ogden and Richards, the thought of reference symbolizes the symbol and refers to the referent. The relationship between the symbol and the referent is, thus, indirect ("imputed").

By postulating the disconnect between the word (symbol) and the thing it refers to (referent), a revolutionary idea at the time, Ogden and Richards attempted to explain the misuse and abuse of language. For instance, language is often used to refer to things that do not, in fact, exist. As prescriptive linguists, they believed that, if only people used words right, many real world problems would disappear. In this, they anticipated the concerns of the general semanticists, such as

39. Frege (1952b) actually used the term *Sinn* 'sense' for meaning and *Bedeutung* 'meaning' for reference.

Korzybski, whose most well-known book was, in fact, entitled “Language and Sanity” (1933), and Hayakawa (1975). Ogden and Richards, thus, proceeded from the assumption that speakers can avoid “abusing” language, that is, that language can and should be made, in some sense, logical. Carnap (1937) was, independently, sympathetic to this concern and tried to develop principles for constructing fully logical artificial languages for human consumption. Wittgenstein (1953: 19^e) would make a famous observation that “philosophical problems arise when language *goes on holiday*” that resonates with the original thinking of Ogden and Richards.⁴⁰

3.4 The Quest for Meaning Representation I: From Ogden and Richards to Bar-Hillel

While Ogden and Richards identified the symbols with words and the referents with things in the world, they made no claim about the nature of the thought of reference (that is, meaning). Stern (1931) placed the latter in the domain of ‘mental content’ situated in the mind of the speaker. In this, he anticipated work on mental models (e.g., Miller and Johnson-Laird 1976), mental spaces (Fauconnier 1985) and artificial believers (e.g., Ballim and Wilks 1991). Over the years, there have been several types of reaction to the task of meaning representation, and various researchers have opted for quite different solutions.

3.4.1 Option 1: Refusing to Study Meaning

Stern postulated the nature of meaning but said nothing about how to explore it. Of course, it is not at all clear how to go about this task of describing something which is not as directly observable as words or real-world objects. In the behaviorist tradition, ascendant in the USA roughly between 1920 and 1960, the study of unobservable objects became unacceptable. That is why Bloomfield (1933) declared that meaning is but a linguistic substitute for the basic stimulus-response analysis of human behavior. In his classical example, he described the behavior of a human being, Jill. When she is hungry (stimulus) and sees an apple (another stimulus), she picks it up and eats it (response). Stimuli and responses need not be real-life states of affairs and actions. They can be substituted for by language expressions. Thus, in the situation above, Jill may substitute a linguistic response for her action by informing Jack that she is hungry or that she wants the apple. This message becomes Jack’s linguistic stimulus, and he responds with a real-life action. Thus, Bloomfield does not reject the concept of meaning altogether. However, it is defined in such a way that the only methodology for discovering and describing, for instance, the meaning of a particular word, is by observing any common features of the situations in which this word is uttered (cf. Dillon 1977).

Without any definition of the features or any methods or tools for recording these features, this program is patently vacuous. Bloomfield considered the task of providing such definitions and methods infeasible. As a result, he did the only logical thing: he declared that semantics should not be a part of the linguistic enterprise. This decision influenced the progress of the study of meaning in linguistics for decades to come. Indeed, until Katz and Fodor (1963), meaning was

40. Another foresight of Ogden and Richards that took wing in later years was the idea of expressing meanings using a limited set of primitives (“Basic English”). This idea anticipates componential analysis of meaning (e.g., Bendix, 1966). A similar direction of thought can be traced in the works of Hjelmslev (e.g., 1958) and some early workers in artificial intelligence (Wilks 1972, Schank 1975).

marginalized in linguistics proper, though studied in applied fields, such as anthropology—which contributed to the genesis of componential analysis of word meaning—or machine translation—which has maintained a steady interest in (lexical) semantics. Thus, a pioneer of machine translation stated: “...MT is concerned primarily with *meaning*, an aspect of language that has often been treated as a poor relation by linguists and referred to psychologists and philosophers. *The first concern of MT must always be the highest possible degree of source-target semantic agreement and intelligibility.* The MT linguist, therefore, must study the languages that are to be mechanically correlated in the light of source-target semantics.” (Reifler 1955: 138).

3.4.2 Option 2: Semantic Fields, or Avoiding Metalanguage

Before componential analysis emerged as a first concrete approach to describing word meaning, Trier (1931), Weisgerber (1951) and others distinguished and analyzed ‘semantic fields,’ that is, groups of words whose meanings are closely interrelated. A simple topological metaphor allowed the authors to position the words with ‘contiguous’ meanings next to each other, like pieces of a puzzle. The original semantic fields defined contiguity on a mixture of intuitive factors including, among others, both the paradigmatic (synonymy, hyperonymy, antonymy, etc.) and the syntagmatic (what we today would call thematic or case-role) relations among word meanings. Characteristically, none of these relations were either formally defined or represented in the semantic fields: in other words, the semantic field approach explored semantics without an overt metalanguage. In this sense, semantic fields anticipated a direction of work in corpus linguistics in the 1990s, where paradigmatic relations among word meanings are established (but once again, with neither word meanings nor semantic relations overtly defined or represented) by automatically matching the contexts in which they are attested in text corpora. It is not surprising that the same corpus linguists have widely used thesauri (originating in modern times with Roget 1852), practical lexicographic encodings of the intuitive notion of semantic fields that, in fact, predated the work on semantic fields by almost a century.

Hjelmslev (1958) compared semantic fields across different languages. This gave him the idea about determining the minimal differentiating elements (‘semes,’ in Hjelmslev’s terminology) of meaning which would allow to describe word meaning in any language. Not only do the semes provide a bridge to componential analysis, they also anticipate modern work in ontology. The notion of semantic fields was given an empirical corroboration when Luria (e.g., Vinogradova and Luria 1961) showed through a series of experiments that human conditional reflexes dealing with associations among words are based on the speaker’s subconscious awareness of structured semantic fields.

3.4.3 Option 3: Componential Analysis, or the Dawn of Metalanguage

The anthropologists Kroeber (1952), Goodenough (1956) and Lounsbury (1956) suggested a set of semantic features (components) to describe terms of kinship in a variety of cultures. Using an appropriate combination of these features, one can compose the meaning of any kinship term. Thus, the meaning of ‘father’ is the combination of three feature-value pairs: {GENERATION: -1; SEX: male; CLOSENESS-OF-RELATIONSHIP: direct}. If the approach could be extended beyond closed nomenclatures to cover the general lexicon, this would effectively amount to the introduction of a parsimonious metalanguage for describing word meaning, as relatively few features could be used in combinations to describe the hundreds of thousands of word meanings, presu-

ably, in any language. Leaving aside for the time being the unsolved (and even unstated) issue of the nature of the names for the component features (are they words of English or elements of a different, artificial, language?), the componential analysis hypothesis promised exciting applications in practical lexicography, language training and computer processing of language.

It was shown later by Katz and Fodor (1963) that the general lexicon could be represented using a limited number of semantic features only if one agreed to an incomplete analysis of word meaning. They called the ‘residue’ of the word meaning after componential analysis ‘the semantic distinguisher’ and did not analyze that concept any further. Thus, one of the senses of the English word *bachelor* was represented by the set of componential features (‘semantic markers’ to Katz and Fodor) of (Human) (Adult) (Male) and the semantic distinguisher [Who has never married]. This meaning is, for Katz and Fodor, a combination of the meaning of *man*, derived fully componentially, and an unanalyzed residue. Katz and Fodor realized, of course, that each such residue could be declared another marker. However, this would have led to unconstrained proliferation of the markers, which would defeat the basic idea of componential analysis: describing many in terms of few.

3.4.4 Option 4: Logic, or Importing a Metalanguage

Greenberg (1949) introduced first-order predicate calculus as the metalanguage for componential analysis. As a result, various features (components) were assigned different logical status. Some were predicates, others, arguments; still others, functors. Thus, if xPy is defined as ‘ x is a parent of y ,’ f is defined as ‘female,’ $u \neq v$ and $x \neq y$, then $(\exists u)(\exists v)[uPx \ \& \ uPy \ \& \ vPx \ \& \ vPy \ \& \ x \neq f]$ means ‘ x is a sister of y .’ Greenberg demonstrated that his system was, indeed, capable of expressing any kind of kinship relationship. It was not important for him that his formulae could be expressed in a number of ways in natural language, not always using strictly synonymous phrases; e.g., the formula above can be expressed as ‘ y has a sister,’ ‘ y is a brother or sister of x ’ or even ‘ u and v have at least two children, and one of them is a girl.’ If a relationship—for instance, equivalence—is posited for two formulae, the result is a true or false statement. Also, formulae usually have entailments, e.g., that u and v in the formula above are not of the same sex. The categories of truth and entailment, while peripheral for an empiricist like Greenberg, are central to any approach to semantics based on logic.

While Greenberg used mechanisms of logic to analyze word meaning, the main thrust of the logical tradition in the study of language had been to apply its central notion, the proposition, to the study of the sentence. Extending the Ogden and Richards’ triangle to sentence level from word level, we obtain the following relationships:




Figure 18. Ogden and Richards' triangle extended to sentence level from word level

The logicians renamed the labels of the nodes in this triangle with terms defined inside their system:




Figure 19. The meaning triangle at the sentence level, using logicians' terms.

The main difference between the logical triangle in Figure 19 and that the one in Figure 18 is that, in the former, none of the elements relates directly to natural language. A proposition is the result of a translation of a sentence into the metalanguage of logic. Its extension (also referred to as 'denotation') is formally defined as the truth value of the proposition, realized as either 'true' or 'false.' The intension of a proposition is defined as a function from the set of propositional indices, such as the speaker, the hearer, the time and location of the utterance and a 'possible world' in which it is uttered, to the proposition's extension (see, e.g., Lewis 1972). While these definitions are very natural from the point of view of logic, we will argue later that, outside of it, they are not necessarily so.

Bar-Hillel (1970: 202-203) characterized the overall program of exploring language using the tool of formal logic as follows: "It seems that... the almost general attitude of all formal logicians was to regard [semantic analysis of natural language] as a two-stage affair. In the first stage, the original language formulation had to be rephrased, without loss, in a normalized idiom, while in the second stage, these normalized formulations would be put through the grindstone of the formal logic evaluator.... Without substantial progress in the first stage even the incredible progress made by mathematical logic in our time will not help us much in solving our total problem." The first

stage may have been motivated by the desire—shared by such very different scholars as Ogden and Richards, on the one hand, and Carnap, on the other—to make natural language more logical and thus to avoid obfuscation through polysemy, use of metaphor and other phenomena that make semantic analysis difficult. Another related goal was to cleanse language of references to non-existent entities that make analysis through logic impossible. Indeed, had this goal been achieved, Russell (1905; see also Frege 1952a) would not have had to devote so much thought to the issue of the truth value of the proposition contained in the utterance *The present king of France is bald*.

The implementation of the first of Bar-Hillel's two stages of the logic program for semantics would have enabled the second stage to express a complete analysis of the meaning of natural language utterances in logical terms. The development of the second stage proved much more attainable (provided one assumed the success of the first stage). Given this assumption, the second stage was able to concentrate on such purely technical issues in logic as the calculation of truth values of complex propositions, given the truth values of their components; truth preservation in entailments; or the assignment of appropriate extensions to entities other than objects and propositions (for instance, events or attributes).

Bar-Hillel's charge concerning the first stage of the program of logic vis-a-vis language could, in fact, be mitigated if one took into account the attempts by logicians to account at least for the syntactic properties of natural language sentences. Ajdukiewicz's (1935) work that eventually led to the development of categorial grammar (Bar-Hillel 1953), was the first attempt to describe phrase and sentence structure formally. The grammar introduces two basic notions—the sentence (S) and the noun (N)—and presents the syntactic value of the sentence as the product of its constituents. Thus, a one-place predicate, such as *sleep* in *George sleeps* obtains the value of S/N, which means that it is the element which, when a noun is added to it, produces a sentence ($N \times S/N = S$). Similar formulae were built for other types of predicates, for modifiers, determiners and other lexical categories. This work was the first example of the logical method applied to a purely linguistic concern, falling outside the program of logic proper. Indeed, it deals, though admittedly not very well, with the syntax of natural language, which is much more complex than the formal syntax of a logical system.

Ajdukiewicz's work seems also to have first introduced into linguistics and logic the idea of a process through which one can compose a characterization of a complex entity out of the characterizations of its constituents. After Ajdukiewicz, Bar-Hillel and Chomsky, among others, applied this method to syntax of natural language without necessarily preserving the original formalism. Later, Katz and Fodor in linguistic semantics and Montague within the logic camp extended this method to deriving the meaning of a sentence from the meanings of its constituents. Work on compositional syntax led to ideas about the compositional derivation of sentence meaning from meanings of phrases and the latter, from meanings of words.

3.5 The Quest for Meaning Representation II: Contemporary Approaches

3.5.1 Formal Semantics

Semantic compositionality (see, for instance, Partee 1984a) deals with the contribution of sentence constituents to the truth value of a proposition expressed by a sentence. The basic process of calculating truth values resembles syntactic analysis in categorial grammar, with sentence consti-

uents being assigned labels in which the syntactic category *S* is replaced by the truth value *t*. Thus, the extension of a simple proposition like **George snores**, denoted (cf. Heim and Kratzer 1998) as $[[\text{George snores}]]$, is defined as a function of $[[\text{snores}]]$ called with the argument $[[\text{George}]]$, or $[[\text{George snores}]] = [[\text{snores}]] ([[\text{George}]])$. If the proposition **George snores** is true (which it is if George, in fact, snores), the formula becomes $t = [[\text{snores}]] ([[\text{George}]])$. More generally, for one-place predicates like *snores*, $t = [[\text{predicate}]] ([[\text{argument}]])$. Conflating logical terms with lexical categories, as is customary in formal semantics, we can write $t = [[V]]([[PrN]])$, where *V* stands for verb and *PrN*, for proper noun.

It is precisely this operation of assigning appropriate extensions to the components of a proposition that is described as “...a principle of compositionality, which states that the meaning of a complex expression is determined by the meaning of its constituents and the manner in which they are combined” (Ladusaw 1988:91). Let us see how this program for formal semantics handles the following four central issues on its agenda (*op.cit.*: 92): “1. What is the formal characterization of the objects which serve as semantic representations? 2. How do these objects support the equivalence and consequence relations which are its descriptive goal? 3. How are expressions associated with their semantic representations? 4. What are semantic representations? Are they considered to be basically mental objects or real-world objects?”

The formal characterization of semantic representations refers to the metalanguage of double brackets for representing extensions. By contributing correctly to the calculation of the truth value of the propositions, these representations clearly support such truth value-based relations as equivalence, consequence (entailment) and all the others. The expressions are associated with their semantic representations by the act of assignment. Whether semantic representations are mental or real-world objects does not directly influence the compositional process, though this issue is the object of active research and debate (with, e.g., Fodor and Lepore 1998 and Fauconnier 1985 arguing for the mentalist position; and, e.g., Barwise and Perry 1983 contributing to the opposing view).

Thus, on their own terms, formal semanticists can declare that their program indeed responds to the four questions they consider central to the semantic enterprise. As a result, the bulk of the research focuses on the refinement of the logical formalism and extension assignments, and on extending the range of linguistic examples that can illustrate the appropriateness of the logical formalism. Over the years, formal semantics has concentrated on studying the meaning of the syntactic classes of nouns and verbs, thematic roles, space (including deixis), aspect, tense, time, modality, negation and selected types of modification, with the greatest amount of effort devoted to the issue of quantification. Practically any book or article on formal semantics has been devoted to a subset of this inventory (see Montague 1974; Dowty 1979; Dowty *et al.* 1981; Partee 1973, 1976; Hornstein 1984; Bach 1989; Chierchia and McConnell-Ginet 1990; Frawley 1992; Cann 1993; Chierchia 1995; Heim and Kratzer 1998).

As was already mentioned, the truth value of a proposition establishes a direct relation between the sentence containing the proposition and the state of affairs in the world, that is, between language and the extralinguistic reality that language “is about” (Ladusaw 1988: 91; Chierchia and McConnell-Ginet 1990:11). This tenet is so basic and essential to the formal semantics program that the truth values assume the dominant role in it: only issues that lend themselves to truth-con-

ditional treatment are added to the inventory of formal semantics tasks. As a result, many issues escape the attention of formal semanticists, in other words, are declared to be outside the purview of this approach. Among the important issues that cannot be treated using truth values are conversion of natural language sentences into logical propositions⁴¹ (cf. Bar-Hillel's comment on the subject discussed in 3.1.4.4 above); representation of lexical meanings for most open-class lexical items,⁴² which would enable a substantive representation for the meaning of a sentence; as well as the resolution of most kinds of semantic ambiguity, notably, every ambiguity not stemming from a syntactic distinction.

The insistence on using truth values as extensions for propositions leads to assigning the same extension to all true propositions, and thus effectively equating, counterintuitively, all sentences expressing such propositions. The formal semanticists perceived both this difficulty and the need for overcoming it: "... if sentences denote their truth values, then there must be something more to sentence meaning than denotation, for we don't want to say that any two sentences with the same truth value have the same meaning" (Chierchia and McConnell-Ginet 1990:57). So, the category of intension was introduced to capture the differences in meaning among propositions with the same extension. If one uses the standard definition of intension (see 3.1.4.4 above), such differences can only be represented through different values of the intensional indices. As the set of values of the speaker, the hearer, the time and place of the utterance is insufficient to capture realistic semantic differences, the set of all objects mentioned in the propositions is added as another index (see, e.g., Lewis 1972). This addition preempts the necessity to explain the semantic difference between two sentences pronounced in rapid succession by the same speaker in the same place and intended for the same hearer simply by the minuscule difference in the value of the time index. For example, if Jim says to Rémi in Las Cruces, NM, on September 15, 1999 at 14:23:17, *The new computer is still in the box* and, at 14:23:19, *Evelyne is still in Singapore*, the index values {computer, box} and {Evelyne, Singapore}, respectively, distinguish the propositions underlying these utterances much more substantively than the two-second difference in the value of the time index.

The sentence *The new computer is still in the box* shares all the index values with such other sentences as *The computer is in the new box*, *The old computer is in the box*, *The box is behind the new computer*, *The new computer resembles a box*, and many others. These sentences obviously differ in meaning, but the intensional analysis with the help of the indices, as defined above, fails to account for these differences. The only method to rectify this state of affairs within intensional analysis is to introduce new indices, for instance, a predicate index, an index for each attribute of each predicate and object, etc. In other words, for an adequate account of all semantic differences among sentences, the framework will need an index for every possible meaning-carrying linguistic entity that might occur in the sentence. When this is achieved, it will appear that the original indices of speaker, hearer, time and place prove to contribute little, if anything, to the representa-

41. On the one hand, the same proposition can be expressed in a language using any sentence from an often large set of paraphrases. On the other hand, the same sentence expresses a proposition and all of its logical equivalents.

42. Marconi (1997: 1) seems to make a similar argument: "...I concentrated on the understanding of *words*: not words such as 'all,' 'and,' and 'necessarily' but rather words such as 'yellow,' 'book,' and 'kick' [because] the research program generated within the traditional philosophical semantics stemming from Frege... did not appear to adequately account for word meaning."

tion and disambiguation of sentence meaning.⁴³

While this method of extending the intensional analysis of meaning is plausible, it has not been pursued by formal semantics.⁴⁴ This is not because formal semanticists did not recognize the problem. Kamp (1984:1) formulated it as follows:

“Two conceptions of meaning have dominated formal semantics of natural language. The first of these sees meaning principally as that which determines conditions of truth. This notion, whose advocates are found mostly among philosophers and logicians, has inspired the disciplines of truth-theoretic and model-theoretic semantics. According to the second conception, meaning is, first and foremost, that which a language user grasps when he understands the words [sic!] he hears or reads. The second conception is implicit in many studies by computer scientists (especially those involved with artificial intelligence), psychologists and linguists—studies which have been concerned to articulate the structure of the representations which speakers construct in response to verbal input.”

Kamp adhered to both of these conceptions of meaning. His Discourse Representation Theory (DRT) proposed to combine the two approaches, specifically, by adding to the agenda of formal semantics a treatment of co-reference and anaphora. He suggested that, in the mind of the speaker, there exists a representation that keeps tabs on all the arguments of all predicates that helps to recognize deictic antecedents and referents of all definite descriptions. This proposal amounts to adding another index to intensional semantics, which is definitely useful. However, the same discourse representation structure will still represent sentences with different meanings. In other words, even after Kamp’s enhancements, formal semantics will still assign the same sets of index values to sentences with different meanings.

Barwise and Perry (1983) took a completely different road to obviating the difficulties stemming, in the source, from the foundational tenet of reliance on truth values. They declared that the extension of a proposition is not a truth value but rather a complex entity they called the ‘situation.’ This extension was rich enough to allow for semantically different sentences to have different extensions, which made the account much more intuitive and closer to what “a language user grasps” about meaning, thus bridging the gap mentioned by Kamp. Their approach ran into two kinds of difficulties. First, there are no tools to describe actual situations within the arsenal of formal semantics, including neither a methodology nor a tradition of large-scale descriptive work, and Barwise and Perry did not attempt to borrow that expertise from elsewhere, e.g., field linguistics. Second, they came under attack from fellow logicians and philosophers of language for using a category, situation, which was dangerously close to the category of fact, which, in turn, had long been known to philosophers as practically impossible to define and manipulate properly (Austin 1962, cf. 1961a,b).⁴⁵

43. It is possible, however, that these indices may prove very important, for example, in applications, such as systems devoted to question answering based on inferences about facts in the Fact DB.

44. Instead, when intension is discussed at all in formal semantics (e.g., Ladusaw 1988, Chierchia and McConnell-Ginet 1990), it is typically limited to the issue of truth values in the so-called ‘opaque’ contexts, such as the belief sentences

45. The problem with the category of fact in philosophy has been essentially that any candidate fact could be easily shown to be an aggregate of other facts. This search for the elementary (or primitive) fact stemmed, of course, from the axiomatic theory paradigm which requires a postulated finite set of primitives.

3.5.2 Semantic vs. Syntactic Compositionality

Sentences are syntactically compositional because they consist of clauses, which, in turn, consist of phrases, which, in turn, consist of other phrases and words. In other words, saying that sentences are syntactically compositional is tantamount to saying that they have syntactic structure. Sentence meaning is compositional because, to a large extent, it depends on a combination of the meanings of sentence constituents, which implies the concept of semantic structure. That both syntactic structure and semantic structure are compositional does not imply that the two structures are in any sense isomorphic or congruent: in other words, it does not follow that the syntactic and semantic constituents are the same.

Formal semanticists are aware of the possible distinctions between the shape of the syntactic and semantic structures. “In theory, the semantically relevant structure of a complex expression like a sentence may bear little or no relation to the syntactic structure assigned to it on other linguistic grounds (on the basis, for example, of grammaticality judgments and intuitions about syntactic constituency)” (Chierchia and McConnell-Ginet 1990: 91).

Having observed a parallelism between the (morphological) lexicon and phrase structure rules in syntax, on the one hand, and the (semantic) lexicon and compositional rules in semantics, on the other, Ladusaw observes that “[t]he distinction between lexical and compositional in semantics is not necessarily the same as between lexical and phrasal in syntax. Polymorphemic words may have completely compositional meanings and apparently phrasal constituents may have idiomatic meanings. See Dowty (1978) and Hoeksma (1984) for a discussion of the relationship between compositionality and the lexical/syntactic distinction.”

We basically agree with this observation, though we believe that it does not go far enough in stating the inherent discrepancies between syntactic and semantic compositionality. First, experience in multilingual descriptive work clearly shows that word boundaries and, therefore, the demarcation lines between morphology and syntax, are blurred and unimportant for grammatical description (see, e.g., Kornfilt 1997 on Turkish agglutination or Dura 1998 on Swedish compounding). Second, even a non-polymorphemic word may have a compositional meaning, as Postal (1971) showed on the example of the English *remind*, which he analyzed as STRIKE + SIMILAR. Raskin and Nirenburg (1995) identifies many cases of syntactic modification (such as adjective-noun constructions), in which no semantic modification occurs: thus, *occasional pizza* actually means that somebody eats pizza occasionally, and *good film* means that somebody watches the film and likes it.

Unfortunately, as formal semanticists readily admit, the reality of research in the field with regard to the relationship between syntactic and semantic compositionality is different: “In practice, many linguists assume that semantics is fed fairly directly by syntax and that surface syntactic constituents will generally be units for purposes of semantic composition. And even more linguists would expect the units of semantic composition to be units at some level of syntactic structure, though perhaps at a more abstract level than the surface” (Chierchia and McConnell-Ginet 1990: 91). We could not have said this better ourselves (see, however, Nirenburg and Raskin 1996; see also Chapter 4).

3.5.3 Compositionality in Linguistic Semantics

Similarly to formal semanticists, Katz and Fodor (1963) believed that semantic compositionality is determined by syntactic compositionality. Their semantic theory, the first linguistic theory of sentence meaning, was conceived as a component of a comprehensive theory of language competence which had at its center a syntactic component, specifically, the transformational generative grammar. The comprehensive theory implied an order of application of the constituent theories, with the output of the syntactic component serving as the input for the semantic component.

Having realized that Chomsky's syntax was a model of the speakers' grammatical competence, more specifically, their ability to judge word strings as well-formed or not well-formed sentences of a language, Katz and Fodor extended the same approach to semantics. Only instead of well-formedness (or grammaticality), they were interested in the speakers' judgments of meaningfulness. They defined semantic competence as a set of four abilities:

- determining the number of meanings for each sentence;
- determining the content of each meaning;
- detecting semantic anomalies in sentences; and
- perceiving paraphrase relations among sentences.

Their semantic theory consists of two components: the dictionary and the compositional projection (or amalgamation) rules. In the dictionary, each entry contains a combination of lexical category information, such as common noun, with a small number of general semantic features (see 3.1.4.3 above). Starting at the terminal level of the phrase structure represented as a binary tree, the projection rules take pairs of lexical entries that were the children of the same node and amalgamate their semantic markers. A special rule is devised for each type of syntactic phrase. The procedure continues until the semantics of the root node of the tree, *S*, is established. For example, the head-modifier projection rule essentially concatenates the semantic features in the entries for the head and the modifier. A more complex verb-object rule inserts the entry for the object *NP* into the slot for object in the verb's entry. A special slot in the entries for nominal modifiers and verbs lists selectional restrictions (represented as Boolean combinations of semantic features) that constrain the modifier's capacity to combine with particular heads and the verb's capacity to combine with certain verbal subjects and objects, respectively. Projection rules fire only if selectional restrictions are satisfied. Otherwise, the sentence is pronounced anomalous.

Katz and Fodor's was the first theory that combined lexical and compositional semantics. They were also the first to address explicitly the purview of their enterprise and deliberately to constrain it. While semantic competence, as the authors defined it, obviously includes the speaker's capacity to understand each sentence in context, Katz and Fodor saw no way of accommodating this capability within a formal theory. Instead, they declared the sentence meaning "in isolation" to be the only viable goal of their, and any other, theory. Without the disambiguating role of the context, this results in a counterintuitive treatment of virtually any sentence as ambiguous. In other words, they did not have a procedure for determining which of the potential meanings of a sentence was appropriate in a text. They could claim, however, that this latter task was not one of the four aspects of semantic competence that their theory was set up to model. While this claim was correct, it led to a serious discrepancy between the goal of their theory and the actual semantic competence of the speakers. This amounted to trading a real and necessary but seemingly unat-

tainable goal for a well-defined and specially designed objective that seemed attainable. In this respect, there is no theoretical difference between Katz and Fodor's substitution and the decision to study truth values in lieu of meaning on the part of formal semanticists, except that Katz and Fodor were aware of the substitution and open about it. It matters also, of course, that their theory produced a list of possible meanings out of which the desired one could be selected.

The appearance of Katz and Fodor's article, followed by Katz and Postal (1964), had the effect of energizing research on compositional semantics within linguistics. Many leading linguists commented on this theory, often criticizing quite severely its various tenets, with the curious exception of the above meaning-in-isolation flaw. Thus, Weinreich (1966) perceptively accused Katz and his co-authors of having no criteria for limiting the polysemy in their dictionary entries. Lakoff (1971) convincingly showed that in order for the proposed semantic theory to work, the overall "architecture" of the linguistic theory needed to be changed. Staal (1967) and Bar-Hillel (1967) observed that the proposed theory could not accommodate such important semantic relation as the conversives, e.g., *buy / sell*. Nonetheless, no critic of Katz and his co-authors (see, however, Raskin 1986) attacked their four-part agenda (even though the issue of paraphrases was manifestly ignored in the theory⁴⁶), and it has proved useful to gauge any subsequent semantic proposals against the background of Katz and Fodor's theory.

Remarkably, Katz and Fodor achieved their compositional semantic goals without feeling any need for truth values, which is, of course, directly opposite to the formal semantics approach. Another related difference is Katz and Fodor's emphasis, often exaggerated by their critics, on disambiguation while formal semantics has no interest and no tools for dealing with the problem. The response to Katz and Fodor's theory from formal semanticists was seminally formulated by Lewis (1972), who pointed out the failure of their semantic features, markers and distinguishers (which, for him, were just words in "Markerese"), as failing to relate language to the extralinguistic reality. It was as an alternative to Katz and Fodor's theory that Lewis formulated the first cohesive proposal of intensional semantics.

As we discuss in 2.6.2.2 above and 3.3.3.2 below, the position of ontological semantics is different from both Katz and Fodor's and Lewis'. We only partially agree with Jackendoff (1983: x) that "the standard notions of truth and reference play no significant role in natural language semantics."⁴⁷ First, we maintain that reference is relevant for the study of co-reference and anaphora (both of which, in ontological semantics, are subsumed by the phenomenon of reference) relations in text. Second, while we agree that truth plays no role in the speaker's processing of

46. Contrary to the initial implication by Katz and Fodor, paraphrases would not get identical semantic interpretations in the theory, and an additional apparatus would be necessary to establish the appropriate equivalences. Formal semanticists are right in claiming an advantage in this respect because their "semantic representations are logical formulas from an independently defined logic [, which] allows the theory to incorporate all of the familiar logic equivalences" (Ladusaw 1988: 92).

47. The linguistic tradition of rejecting truth-conditional semantics dates back at least to Wilson (1975) who accused it of impoverishing the treatment of meaning in language, of using entailment and truth conditions in ways that are too wide for linguistic semantic purposes and of being unable to treat non-declaratives. Even more devastatingly, we think, is the fact that using truth values creates pseudo-problems in linguistic semantics: thus, the sentence *The present king of France is bald* is seen as highly problematic by formal semantics because it has no truth value; it is, however, perfectly meaningful and problem-free from the point of view of linguistic semantics.

meaning, we are also aware of the need to “anchor” language in extralinguistic reality. Formal semanticists use truth values for this purpose. We believe that this task requires a tool with much more content, and that an ontology can and should serve as such a tool. On the other hand, we find the “Markerese” accusation spurious: there is no legitimate way to confuse semantic markers with words of English. We deflect a similar criticism concerning the use of English labels for ontological concepts by explicitly setting up these labels as language-independent entities with their own content and by training the personnel working with these labels to distinguish between elements of the ontology and elements of language.

3.6 A Trio of Free-Standing Semantic Ideas from Outside Major Schools

Ontological semantics contains elements that reverberate against a few interesting semantic ideas that have been proposed outside of the major semantic approaches and that have never been fully incorporated by those approaches.

The intuition that each utterance carries a reference to information already known to the hearer as well as information that is new to the hearer was first formulated as the basis of the so-called functional perspective on the sentence by the founders of the Prague Linguistic Circle (Mathesius 1947). It has been a recurring issue in semantics and pragmatics ever since, under different terminological systems (see, for instance, Kuno 1972; Chafe 1976; Clark and Haviland 1977; Prince 1979, 1981). The distinction, while definitely useful, cannot provide a comprehensive representation of sentential meaning—it can only contribute as an add-on to a full-fledged semantic system. Before generative grammar, however, this phenomenon was studied essentially in isolation. In generative grammar, the distinction, introduced as presupposition and focus (Chomsky 1971), was supposed to be added to the semantic component, but the idea was never implemented. More recently, work has been done on incorporating the topic/focus dichotomy in formal syntax and semantics (e.g., Krifka 1991, Rooth 1992, Birner and Ward 1998, Hajičová *et al.* 1998) and in the study of prosody and intonation (e.g., Féry 1992, Hajičová 1998). In computational linguistics, information about focus and presupposition was used primarily, though not exclusively, in natural language generation, and was implemented through a set of special clues (e.g., McKeown 1985 but also Grosz 1977). Ontological semantics accommodates the distinction between old and new information using the mechanism of the saliency modality parameter. The microtheory of saliency includes several clues for establishing the appropriate values (XREF).

Humboldt (1971) and Whorf (1953) introduced the intriguing idea that different languages impose different world views on their speakers. Humboldt spoke of the magic circle drawn by the language around the speaker, a metaphor characteristic of Romanticism in science, art and culture that was the dominant contemporary world view, at least in Germany. Whorf, on the other hand, amassed empirical data on such crucial, for him, differences among languages as the circular notion of time in Hopi as opposed to the linear notion of time in “Standard Average European.” Whorf’s claims of this nature depended primarily on the availability of single-word expressions for certain ideas: the unavailability of such an expression for a certain idea was interpreted by him as the absence of this idea in the world of the speaker of that language. Taking this claim absurdly far, one arrives at the conclusion that an Uzbek, whose language reportedly has only three words for color, can distinguish fewer colors than the speakers of languages with a larger color taxonomy. Whorf’s own and subsequent research failed to produce any justification for the prime

nature of the single-word claim (XREF to Footnote in 4). As most other approaches, ontological semantics subscribes to the principle of effability (XREF) which directly contradicts the Whorf hypothesis. Moreover, ontological semantics is based on ontology that is language independent and thus assumes the conceptual coherence of all natural languages. The lexicon for every language inside ontological semantics uses the same ontology to specify meanings, and, as it must cover all the meanings in the ontology, some of the entry heads in the lexicon will, for a particular language, end up phrasal.

Among Alfred Korzybski's (1933) many bizarre ideas about semantics, completely marginalized by the field, there was a persistent theme of instantiating a mention of every object. He claimed that no mention of, say, a table, could be made without its unique numbered label, no mention of a person, without an exact date in the life of this person about which the statement is made. This idea is a precursor for instantiation in ontological semantics, a basic mechanism for meaning analysis.

3.7 Compositionality in Computational Semantics.

When Katz and Fodor described semantic processes, they had in mind mathematical processes of derivation. With the advent of computational processing of language, a natural consequence was algorithmic theories of language processing, often with the idea of using their results as the bases of some computational applications, such as machine translation or text understanding. The goals of computational semantics have been, by and large, compatible with those of linguistic semantics, that is, representing the meaning of the sentence in a manner which is equivalent to human understanding (as aspired to by linguistic semanticists) or as close to human understanding as possible or, at least, complete, coherent and consistent enough to support computational applications of language processing (as computational semanticists would have it).

The reason the computational goals are much more modest is that, unlike linguistic semantics, computational semantics develops algorithms which produce meaning representations for texts (analysis) or texts realizing meaning representations (generation). It is not surprising, in view of the above, that Wilks and Fass (1992b: 1182; see also the longer version in Wilks and Fass 1992a; cf. the earlier work in Wilks 1971, 1972, 1975) states that “[t]o have a meaning is to have one from among a set of possible meanings” and posits as the central goal of a computational semantic theory “the process of choosing or preferring among those,” which is why Wilks’ theory is called ‘preference semantics.’ While the second goal is missing from Katz and Fodor’s theory—and from linguistic theory in general—entirely, there is also a significant difference between treating meaning as a set of possible meanings, as they do, and realizing that actually meaning is always only one element from this set. This was acceptable in a theory that explicitly and deliberately concerned itself mostly with potential meaning rather than with calculating the meaning of a particular sentence in a particular text. The latter goal is, of course, the overall goal of computational semantics.

Wilks (1992b: 1183) sees preference semantics as “a theory of language in which the meaning of a text is represented by a complex semantic structure that is built up out of components; this compositionality is a typical feature of semantic theories. The principal difference between [preference semantics] and other semantic theories is in the explicit and computational treatment of

ambiguous, metaphorical and nonstandard language use.” The components of the theory include up to 100 semantic primitives including case roles, types of action, types of entities and types of qualifiers; word senses expressed in terms of the primitives; a hierarchy of templates corresponding to phrases, clauses and sentences; inference rules used for resolving anaphora; and some text-level structures. Preferences are essentially procedures for applying heuristics to selection restrictions and other constraint satisfaction statements, as well as for selecting the outcome (that is, a semantic representation) with the greatest semantic ‘density’ and ‘specificity’ (*op.cit.*: 1188). There is no expectation in the approach that all preferences will somehow “work,” and provisions are made for such eventualities, so that some meaning representation is always guaranteed to obtain. In other words, this approach is based on a realistic premise that the computer program will have to deal with an incomplete and imprecise set of resources such as lexicons and grammars.

Preference semantics is a comprehensive approach to meaning in natural language not only because it combines lexical semantics with compositional semantics but also because it aspires to a full meaning representation of each sentence. Other approaches in computational semantics were—deliberately or otherwise—less general and concentrated on particular issues. Thus, Schank’s (e.g., 1975, 1981; Lehnert 1978; Wilensky 1983) school of computational semantics, conceptual dependency, used a different and more constrained set of semantic primitives to represent the meaning of both words and sentences but eventually concentrated on story understanding based on the idea of a progressively more abstract hierarchy of text-level level knowledge structures—scripts, plans, goals, memory organization packets, etc. Hirst (1987), following Charniak (e.g., 1983a), further developed the mechanism to calculate preferences, and each computational-semantic project (e.g., Hobbs and Rosenschein 1977, Sowa 1984, among many) propounded a different representation formalism for both text meaning and lexical semantics.

Over the years of work in linguistic and then computational semantics, the early aspirations for parsimony of primitive elements for describing lexical meaning have gradually given way to a more realistic position, first stated by Hayes (1979), that in computational semantics (and, for that matter, in all of artificial intelligence) a much more realistic hope is to keep the ratio of description primitives, a' , to entities under description, a , as small as possible: $a'/a \ll 1$. Experience shows that if the number of primitives is kept small, descriptions tend to become complex combinations of the primitives that are hard to interpret and use. Given the additional fact that such primitives are rarely explicitly described, let alone formally defined, there is a strong pressure to expand the range of each primitive, resulting in vagueness of primitive meaning. This issue strikes us as being of primary importance. While many approaches use primitives (whether overtly or implicitly), very few expend sufficient energy on their explicit characterization, which is essential for reliability of knowledge acquisition and meaning representation. We see ontologies as the loci for precisely such characterizations.

Much valuable experience, both positive and negative, has been accumulated in formal, linguistic and computational semantics. Ontological semantics aspires to take advantage of the results available in the field. We see the principal differences between ontological semantics and other semantic theories as follows. First, besides introducing ontology as a locus for establishing a rich set of primitives, we see it also as the best means of supporting multilingual NLP applications because ontological information is—by definition and by practice of acquisition—language-independent.

Second, ontological semantics is a comprehensive theory integrating lexical semantics with compositional semantics and moving into pragmatics. Third, ontological semantics is designed to adjust semantic description depth to the needs of an application (see 2.5.4). Fourth, ontological semantics has an emphasis on full-coverage description of text at a predetermined level of granularity because a computational procedure has no tolerance for what has become a staple in the mainstream linguistic literature—assumed similarities of descriptions of many phenomena with those few that were actually illustrated, extrapolations to adjacent phenomena, and tempting have-no-more-patience-for-this etceteras in vitally important lists.