

Pauli Salo

## Lexical Atomism and the Principle of (Psycho)semantic Compositionality

### Abstract

What is the constitution of meaning of morphemes (lexical concepts)? Most theories conjecture that they are constituted by molecular or holistic internal structure: prototypes, exemplars, semantic networks, complex schemata, scripts or even classical definitions. Recently, however, a growing opposition has arisen in cognitive science claiming that psychological evidence suggests rather the opposite, namely, that lexical concepts are not internally structured, but that they are primitive unstructured atoms. I will examine these claims in this paper, arguing that such an atomistic hypothesis might indeed be a more plausible alternative.

### 1. Introduction

Lexicon consists of lexical elements, each being a combination of phonological, formal and semantic features. The more or less standard view is that word meanings (meanings of single morphemes) must be internally structured, hence they consist of several interconnected semantic features.<sup>1</sup> Against this received view, some cognitive scientists have argued that word meanings are unstructured, and that there is only one atomistic feature per lexical item

---

<sup>1</sup> Internal structures could be analysed in terms of *prototypes* (Hampton, 2000, Kamp & Partee, 1995, Searle, 1958, Smith & Medin, 1981, Rosch, 1973a); *conceptual roles* (Block, 1986, Cruse, 1986, Field, 1977, Harman, 1982); *complex cognitive schemata* (Barsalou, 1992, Kintsch, 1998); *definitions* (Aristotle, see Charles, 2000, Weitz, 1988; Carnap, 1937); *images or image-prototypes* (Rosch, 1973b, Russell, 1919); *semantic networks* (Findler, 1979, Quillian, 1968); *family-resemblance* (Wittgenstein, 1958); *recognitionist verificationist capacities* (Peacocke, 1992) and even in terms of a whole *theories* (Gopnik, 1988, Gopnik & Meltzoff, 1997, Keil, 1987, Murphy & Medin, 1985).

(morpheme).<sup>2</sup> If these claims are true, then there is no internal structure to constitute lexical meanings because of the so-called “principle of semantic compositionality”: most theories of lexical meanings are not compositional in a way that they must be, as was put in a recent review by L. J. Rips (1995): “if the general criticism is right, it cuts down nearly all models of concepts in psychology, as well as many in linguistic and philosophy” (p. 87). In other words, if these claims are true, then nearly everyone has been wrong about semantic properties in the lexicon: semantic relations claimed to be part of the lexicon must originate from some other source.

## 2. Terminology

“Lexical meaning” is used to refer to meaning as it is assigned to morphemes in a natural language(s), typically in the form of semantic features. Morphemes are taken to be linguistic elements that are atomic from a syntactic and morphological point of view, ignoring whatever is controversial about this notion in linguistics (see, among others, Di Sciullo & Williams, 1987). Hence I will not attempt to discuss the notion of “morpheme,” but assume it throughout. There are some clear instances of this notion (such as *cat*, *dog*, *pet*, *fish*, *run*) which suffice for present purposes. Furthermore, this paper is not concerned with pseudo-productive word formation or derivational relations between words (polysynthesis, compounds, constructions, incorporation). Such processes are not fully productive, systematic or transparent (Chomsky, 1970, Fodor, 1970),<sup>3</sup> while “complete” productivity and systematicity are the properties that matter to the issue at hand. No doubt these idealizations are serious and require justification which, I hope, will emerge as we proceed.

The terms “concepts” or “thoughts” are used to refer to meanings as they are represented in the human mind; more abstract notions “meaning” or “proposition” are used when no mentality is implied. For instance, although the term *water* means H<sub>2</sub>O, and water contains some hydrogen, the concept water does not necessarily contain the concept hydrogen. When referring to concepts, I use underlining; when referring to expressions, I use *italics*. Thus, the term *water* expresses the concept water, which means a chemical substance with its more or less known properties.

---

<sup>2</sup> See Fodor (1975, 1981, 1995, 1998a, 2001), Fodor, Garrett, Walker, & Parkes (1980), Fodor & Lepore (1991), Kintsch (1974), Margolis (1998), among others.

<sup>3</sup> Thus, *musical* - \**artal*, *tidal* - \**waval*, *terrorize* - \**horrorize*, *reality* - \**fakity*, and so on.

I begin by introducing the semantic principle of *compositionality*, first in the mathematicians' sense, then in a more psychological sense. I then continue to show why and how compositionality has led some cognitive scientists to argue that there cannot exist any structure inside lexical items. Finally I discuss some objections to the atomistic theory and conclude that the atomistic theory is still currently the best choice available.

### 3. Semantic compositionality

Frege (1923/1977), Tarski (1935/1956) and Montague (1974) first used the principle of semantic compositionality in the beginning of the 20<sup>th</sup> century. They used it in the following sense:

- (1) The meaning of a complex expression is a function of the meanings of its parts and their syntactic mode of combination, and of these only.<sup>4</sup>

Although (1) is the canonical formulation of the compositionality principle, there is some variation. First, there is currently some debate as to whether or not Frege really assumed compositionality in the sense of (1), or in any sense (see Janssen, 1997). Secondly, Montague originally used a slightly stronger principle than (1), requiring there to exist a homomorphism between syntax and semantics. Finally, we now know that from a mathematical point of view it does not matter whether a linguistic system (expressions plus meanings) is compositional in the sense of (1) or not. Each noncompositional semantic system can be provided a compositional semantics so that the formal properties of that system remain intact (Hodges, 1998, Janssen, 1997, Zadrozny, 1994). Either we must add properties to (1) and attain a more robust principle, or we must conclude that the principle is not important. Indeed, much of the recent debate on compositionality has arisen because of this alleged mathematical vacuousity.

However, there is empirical evidence that a certain more strong notion of compositionality, henceforth "psychosemantic compositionality," must be

---

<sup>4</sup>It is possible to deny the principle of compositionality by negating (1), or by providing a system where the presuppositions of (1) are not satisfied. We can obtain a principle of *contextualism*. On several occasions Hintikka has claimed that we had better reject (1). In particular, his Game-Theoretical Semantics (GTS) and related IF-languages do not obey compositionality (see Hintikka, 2001). In these systems, the meaning of a constituent is determined from its (linguistic) context (Hintikka & Kulas, 1983).

assumed as a *psychological principle of the human mind*. Psychosemantic compositionality must be distinguished from (1), hence, care must be taken to avoid the harmful ambiguity of the term “compositionality” in the literature. The difference in the use of the term “compositionality” is that, whereas the mathematical principle of compositionality (1) requires that there *exists a dependency* between the meanings of the constituents (e.g. morphemes) and their host (sentence), what must be required in addition is a certain *kind of dependency*. For instance, the function must be computable to be of psychological interest; likewise, we shall see that it must be *systematic* as well. Let us first proceed to review the crucial empirical data that has played a key role in the discussion, then ask what kind of a more robust notion of psychosemantic compositionality must be assumed to explain that data.

#### 4. The relevant data: systematicity and productivity

Human behavior is “systematic.” Suppose you have learned to use the expressions *love*, *Mary* and *John*. Now suppose that, in addition, you learn to use one new expression, say *James*. As a consequence, you have come into possession of other new expressions as well, such as *James loves Mary*, *James loves John*, *Mary loves James*, *John loves James*, *James loves James*. In fact, only in finite cases does one learn finite number of new concepts / expressions. Also what has been called “reverse systematicity” holds: if you have learned the expression *James loves Mary*, you must also have learned *James*, *love*, and *Mary*.<sup>5</sup>

It is not, however, just the fact that expressions are learned in infinite ‘clumps’ that has attracted cognitive scientists’ attention, but also the fact that they are “systematically” related to each other. In the general case, learning the expression *James* entails, as a side effect, that one has learned the expression *James loves Mary*, and not, for instance, *Bill was nude in the park*. This means that the truth conditions of (more or less all) complex expressions involving the constituent *John* all involve John and not, say, caterpillar in one context and Bill Paxton in another. Similarly, if *cow* means cows, *black* means black, then *black cow* comes to mean black cows and not the giraffes in the park. Note that it

---

<sup>5</sup>Some scholars think that systematicity is a *constitutive* property of thought, so that anything that fails to be systematic would not be considered as being a thought in the first place (Cocchiarella, 1986, Clark, 1991, Evans, 1982 and Peacocke, 1992). But it is a remarkable cognitive law that our mind/brain is systematic (see Braine, 1963, Bever, Fodor, & Weksel, 1965).

could mean the latter as well; not fully systematic and productive idioms like *red herring* or N-N combinations like *drug campaign* show that this is not irrelevant.

Can we characterize systematicity more exactly? One possibility usually assumed in the literature, implicitly or explicitly, is (2):

- (2) The meaning of the complex expression is constituted by the meanings of its constituents, and by these only.

Thus, the meaning of *black cow* (that is, black cows) is constituted by the meanings of *black* and *cow*, thus by the properties of being black and being a cow, and by nothing else. This seems to agree with how this term is used in cognitive science literature (see Fodor, 1998a: 99). A similar principle can be found from Wittgenstein who wrote that “the reality that corresponds to the sense of the proposition can surely be nothing but its components parts, since we are surely *ignorant* of *everything* else” (Wittgenstein, 1961, 20. Nov. 1914). Principle (2) also implies (1).

Clearly, at all linguistically salient levels (words, phrases, sentences) there are constructions which do not obey (2). One can argue that there might be even an infinite number of exceptions to (2). Yet there are also an infinite number of expressions which *do* obey (2), hence, whatever mechanisms human mind might use to compute the meanings of complex expressions, property (2), or whatever will be assumed to derive it, must be included to the stock.

In addition to being systematic, our behaviour is “creative.” The number of sentences that one can coherently use, and learn, is infinite. A very young child has already mastered well beyond  $10^{20}$  complex or simple expressions and their meanings. We do not want to claim that he has learned these expressions piecemeal, one by one. This calls for a different explanatory mechanism, presumably a ‘generative capacity’ of some sort.

Finally, there is empirical evidence that the generative capacities that are responsible for the production and understanding of expressions somehow harness abstract *constituents* (Chomsky, 1957, 1975, Zwicky, 1978). We cannot describe the properties of language without assuming that linguistic expressions are “made of” abstract constituents like NP, invisible in the corpus or concrete speech acts, but part of the underlying reality of natural language(s),


“I-languages” as they are often called.<sup>6</sup>

Certainly, these properties are rather unique and surprising in the biological world. For this reason alone, we must try to *explain* the data.<sup>7</sup> According to cognitive scientists and psychologists alike, the data must be explained as follows. The learner must first come into possession of a finite stock of *primitive*, unstructured constituents. These constituents are typically referred to as “primitive concepts”. The learner then needs a computational rule that, when several such constituents are put together by some syntactic or associative rule, gives the meaning of the respective complex concept, and does this in a systematic way (2). For example, suppose that the learner has mastered two concepts, the mental representation of *black*, and the mental representation of *cow*. They can be either complex, as usually assumed, or simple. But it is a psychological law that s/he then also possesses the mental representation of *black cow*, which, if this explanation is correct, is then a *construction of the elements that are its parts*. In other words, the concept of black cow is a complex mental representation that inherits its properties systematically from its parts (Fodor, 1998a). The following figure illustrates this process:

---

<sup>6</sup>“I-language” refers to the brain mechanisms that are responsible for the generation and understanding of linguistic expressions; “E-language” refers to a set of expressions that can be said to belong to a language. Thus, a linguistic corpus represents a subset of some E-language.

<sup>7</sup>Fodor thinks that the data (productivity, systematicity, constituency) can be explained by assuming the *principle of compositionality*: “the systematicity and productivity of thought were supposed to trace back to the compositionality of mental representations, which in turn depend on their syntactic constituent structure” (Fodor, 2000: 4). I will not go into this here, but it is easy to see that the principle of semantic compositionality (1) cannot entail any of these properties: a compositional system in the sense of (1) need not be systematic, learnable, or even productive, and it only presupposes the notion of constituency (Hodges, 1998, Janssen, 1997). Assuming (1) would leave the facts as they were. Therefore, Fodor cannot be referring to compositionality in the sense of (1). If so, we still lack a notion of compositionality that could entail those properties, and only those. This is not to claim that Fodor would have claimed that (1) could explain these properties. Rather, Fodor has remain extremely vague about what property he is referring to (see, e.g., Fodor, 1998a: 94 and Fodor, 2001).



**Figure 1.** A black cow and its ‘composition’

In the middle, there is the expression (or the corresponding mental representation) *black cow* with its constituents *black* and *cow* (or black and cow). The meaning of *black cow*, illustrated by the pictures, depends systematically on the meanings of its constituents: whatever *black cow* means is constituted by whatever its constituents mean, together with their ‘logical mode of combination’, and it is constituted by nothing else.<sup>8</sup> This is not the whole theory, but it is certainly true for an infinite number of expressions / concepts. I will now argue that, if these uncontroversial assumptions are right, then the surprising conclusion follows that the lexicon must be unstructured.

### 5. Lexical meaning and compositionality

Consider the following version of the prototype theory. According to this theory, the meaning of a linguistic expression is its prototype: a statistical average of the category, or a collection of typical exemplars of the category (Smith & Medin, 1981). Psychosemantic compositionality requires that we must be able to compute the prototype (meaning) of a complex expression from the

---

<sup>8</sup>These assumptions contribute to what has been called as the “classical cognitive architecture” (Fodor & Pylyshyn, 1988, Pylyshyn, 1984).

prototypes (meanings) of its constituents. Furthermore, there must be a systematic bond between the typical instances of whatever is the meaning of a complex concept and the typical instances of whatever is the meaning of its hosts. Therefore, the prototype of black cow must be such that it can be constructed from the prototypes of black and cow. Otherwise, the learner could not learn expressions in such huge, typically infinite and systematic, sets. Indeed, determining the prototype of black cow from the prototypes of its constituents seems easy: take a typical cow plus a typical black, and what we get is a typical black cow. This was illustrated in the Figure 1.

However, suppose we study all intersective concepts, such as the meaning of *pet fish*. Then, based on the knowledge of the typical pets (perhaps something close to dogs) and typical fish, we must determine what would be the typical pet fish (a prototype of pet fish). But we cannot determine this without presupposing both the meaning of *pet fish* and knowledge about their typical properties: what are typical pets is almost completely arbitrary with respect of what are typical pets and typical fish. What we are looking at here is a random phenomena: for example, pet fish just happen, accidentally, to live typically in a bowl, a fact that is typically true of neither pets nor fish. The property of living in a bowl, plus many other typical properties of pet fish, emerges 'miraculously'. This is illustrated in Figure 2, where pictures represent prototypes.




Figure 2. Pet fish

Figure 2. Pet fish



Why would the combination of a dog-looking creature (a prototype of pet) and a fish-looking creature (a prototype of fish) yield a small goldfish swimming in a bowl (a prototype of pet fish)? Note that it makes no sense to begin to invent *ad hoc* principles that could predict the properties typical pet fish from the properties of typical pet and fish, since there are no general laws concerning such matters: pet fish could have any typical properties, depending on what kind of fish people bother to keep as their pets. So, there cannot be laws or principles that could relate the meanings of complex expressions to the meanings of their constituents, assuming the truth of the prototype theory, and it is nonsensical to try to seek such laws:

What makes something a typical member of the set of Xs needn't be, and generally isn't, what makes something a typical member of some arbitrary sub- (or super-) set of the Xs. And even when it is, it's generally a contingent fact that it is; a fortiori, it isn't a necessary truth that it is; a fortiori, it isn't a linguistic truth that it is, since, I suppose, linguistic truths are necessary whatever else they are. (Fodor, 1998b: 60).

This argument generalises so that no theory of meaning that includes *statistical properties to the meaning of a constituent can satisfy psychosemantic compositionality* (see the references in section §1, footnote 1, for a list of such theories).

Before discussing some of the replies, let me recap the argument presented so far so as to avoid any misunderstanding of the matter at hand. We clearly possess the capacity of systematicity and productivity (2). Although that is only a rudimentary piece in the whole story about the mental capacities related to understanding linguistic expressions in every context of use, it still is *one* part of it. Now it turns out that conceptual roles theories involving statistical attributes – of which most of the current theories of lexical structure are made – are not sufficient to explain that part of our mental capacity. Hence, there is *more* to the meanings of lexical elements than statistical conceptual roles. Conceptual role theories are not sufficient to establish (2); a better alternative is proposed below.

When these problems became evident in the 80's (but, to the best of my knowledge, largely unknown in Finnish linguistic literature), at least four different replies have been offered in print. I will evaluate each in turn.

### 5.1. First proposal: knowledge of the world

First, many cognitive psychologists assumed that, since obviously there is no strict compositional mechanism available, knowledge of the world is involved in the determination of the prototype of complex expressions (concepts) (e.g. Hampton, 2000, Murphy, 1988, Smith & Medin, 1981, Smith & Osherson, 1984). This is a truism: in order to know what typical pet fish look like, one needs knowledge of the world. For instance, finding out what are typical properties of religious rituals of Eskimos, one would need to go out and look, or consult an authoritative book. The point is that these properties could be whatever, and clearly are not subject to any laws. They are just random: there can be no science, no determinate laws, which could tell what kind of religious rituals Eskimos *must* have. This holds for an infinite number of concepts (meanings of expressions mentally understood). But then these properties *cannot be the basis of semantical laws either*. Thus, based on the meaning of its constituents, one can surely understand what *religious rituals of Eskimos* means without knowing what properties religious rituals of Eskimos typically have. Similarly, for *religious rituals of ancient Finns*, *religious rituals of Native Americans*, and so on. This is why psychologists have doubted the prototype theory: most complex concepts lack prototypes, but they do not lack meaning.

Looked at from a slightly different point of view, a learner who uses prototypes cannot learn concepts systematically and productively. But children are both systematic and productive, hence they do not rely on prototypes, although they of course may *have* prototypes.

### 5.2. Second proposal: classical theory

One could assume that the features that make up the meaning of an expression (concept) are all necessary, or defining, for the expression (concept). Whatever is necessary, is not just typical, and whatever is necessary, can't fluctuate randomly. In fact, this theory seems to work well in the case of complex expressions. Thus, *pet fish* has a definition: *x* is a pet fish just in case *x* is pet and *x* is fish. Pet fish is pet, come what may, and it is fish as well. But on the other hand, it need not live in a bowl, or resemble goldfish in any way.

The drawback of this theory is that *lexical concepts lack phrasal definitions*. There is virtually no psychological evidence of such definitions either, obtained by chronological on-line studies or otherwise. Although this is

an empirical hypothesis that could turn out to be differently, it seems that, based on current consensus, the meaning of words cannot be exhaustively defined in terms of other words. But note that, were it the case that lexical concepts would have turned out to be decomposed of classical decompositions, the argument from compositionality could not have ruled *that* theory out. Put in other words, if you can find classical definitions for lexical concepts, then the theory which says that lexical concepts are semantically structured can be saved.

### 5.3. Third proposal: core meaning plus prototypes

According to the third proposal, the lexicon first consists of some kind of “core meaning” that takes care of the requirements of psychosemantic compositionality, but then there are also prototypes, or other conceptual roles, *in addition*. Call this the “dual theory” (see Armstrong et al., 1983, Block, 1986, Cowie, 1999, Horwich, 1997). Since most complex expressions (concepts) lack prototypes, we could claim that primitive concepts are constituted by (i) core meaning plus (ii) prototypes.

It is clear, however, that this theory is not psychosemantically compositional. Suppose you have the concept pet fish, or know the meaning of *pet fish*. Recall that it was a law about our language and thought that we therefore also have the expressions (concepts) *pet* and *fish*. Now suppose that the meaning of *pet fish* would not require knowledge of prototypes of any kind, but that the meaning of *pet* and *fish* would need. Then it also follows that, in general, you could typically know the meaning of *pet fish* without knowing the meaning of *pet* or *fish*: the latter requires something that is “extra” with respect to the former. So prototypes do not constitute lexical meanings, not even in addition to the “core meaning”.

Two claims are easily confused in this connection. One claim says that lexical concepts *have* prototypes, the other says that they are *constituted* by them. It is the latter which is at stake here, not the former. In comparison, it is important to distinguish two claims: viz., that the water is composed of oxygen and hydrogen (it has such an internal structure) from the claim that water is drinkable (it has external relations): the former is constitutive, the latter isn't. Thus, in innumerable circumstances, water has something “extra” beyond its constitution.

#### 5.4. Fourth proposal: lexical atomism

As a fourth option, one could assume that the classical definitional theory is correct, but lexical concepts lack definitions, not because they are constituted by complex prototypes, but because they are *primitive concepts*. Namely, the definitional theory itself does not entail that lexical concepts would need to have definitions: any theory must assume that there are primitives ‘at some level’ (Fodor, 1981). Usually the level of primitive has involved sensory concepts, plus or minus some abstract notions such as Agent, Intention, Object and the like. Yet what if the level of primitives just *is*, more or less, the level of morphemes at the surface language?

If the meanings of lexical words are *not* constituted by definitions, and they *cannot* be constituted by prototypes due to compositionality, then there is no other choice but the atomistic theory: lexical elements are semantically atomistic and they lack internal structure. Call this theory “lexical atomism”. The following table summarizes these findings.

Theory	Main content	Main problems
Classical Theory	Meanings of concepts are determined by definitions.	*Lexical concepts lack definitions.
Prototype Theory	Meanings of concepts are determined by prototypes.	Not compositional, does not work for complex concepts.
Conceptual Role Theories	Meanings of concepts are determined by various conceptual (inferential) roles, thus by various ‘semantic relations’ between concepts.	Not compositional, does not work for complex concepts.
Atomistic Theory	Meanings of concepts are determined by definitions, and lexical concepts are unstructured (see *).	See next chapter.

**Table 1.** Theories of concepts.

## 6. Lexical atomism

If the argument from compositionality is correct, then lexical meaning is in some sense ‘inscrutable’ and cannot be analysed by using other words (/phrases). The meaning of a lexical element cannot be exhausted by resources available in language, whether in terms of definitions, prototypes or conceptual roles.<sup>9</sup> Let us say that lexical meanings involve ‘tacit knowledge’: we ‘know’ about their meaning, but that knowledge is not ‘accessible’ in language, as put by Moore (1903): “If I am asked ‘What is good?’ my answer is that good is good, and that is the end of matter” (p. 6). This tacit knowledge might involve other cognitive faculties, such as kinaesthetic abilities, or, as a behaviorist would have put it, even relations between the mind and the world. To quote Russell: “It is no more necessary to be able to say what a word means than it is for a cricketer to know the mathematical theory of impact and of projectiles. Indeed, in the case of many object-words, it must be strictly impossible to *say* what they mean, except by a tautology, for it is with them that language begins.”<sup>10</sup> If so, this could begin to *explain* why lexical meaning is inscrutable, and why “language begins” from its words, and why philosophical ‘conceptual analysis’ is so hard, if not impossible.<sup>11</sup> These are, roughly, the conclusions reached in Fodor’s classical book *Language of Thought* (Fodor, 1975). Fodor has elaborated this theory in many ways in his subsequent writings, but reviewing this particular theory must be left to another occasion.<sup>12</sup>

---

<sup>9</sup>Innumerable semantic elucidations and lexicographic hints are of course possible, but that is obvious; the problem is what *constitutes* lexical items, not what mutual semantic relations they may have.

<sup>10</sup>Russell, 1940: 26

<sup>11</sup>This is why the present matter, although empirical, is relevant to a philosophy that has greatly occupied itself with so-called “conceptual analysis”. What is truth, love, virtue or a belief? These are all genuine “philosophical questions”, seeking some “ultimate construal”. They are nontrivial problems since, *if* they have an analysis, those analyses are hidden from the surface language. Compare the thoroughly trivial pursuit of the analysis of what it is to be a pet fish – it is simply to be a pet and a fish. This is trivial since the analysis is mirrored in the language. Yet maybe there simply *are no answers* insofar as we stick to the analysis of everyday concepts represented by single morphemes; if there is no answer, then it is a waste of time to try to find one.

<sup>12</sup>If the meaning of lexical elements is unstructured, then what do words such as *cat* mean? According to Fodor, the word *cat* means the property of being a cat (see Fodor, 1998: 107-

There are some objections that reduce this theory's attractiveness, as an anonymous referee puts it: "How can a statement like this be made about the non-existence of lexical structure from a psychological point of view when we are continuously faced with accumulating empirical data about the internal structure of the lexicon?" The referee concludes that "there is definitely some structure within the lexical stock". S/he argues that the lexicon is structured along such semantic axes as animate vs. non-animate, natural kind vs. artefact, and one could surely add more. Many psycholinguistic studies suggest this to be the case as well: numerous priming studies have demonstrated that cognitive processing is sensitive to the semantic structure in the lexicon, and this is just one example among others.

But this is not what anyone holding the atomistic theory would go on to deny. The question is *not* whether lexical elements have mutual relations (syntactic or semantic), the question is what *constitutes* the meaning of lexical elements. That is, what *is* the meaning of a lexical element, what kind of semantic features constitutes it? A human DNA has important relations with various other biological structures, such as RNA, yet RNA does not constitute DNA. They are different chemical compounds, though the relations are important and throughout involved in the function of DNA. The rule is similar for the lexicon and its putative semantic structure: if lexical elements *cat* and *animate* are related, as they obviously are – if only because cats *are* animals for what we know – it does *not thereby follow* that the latter would be a *part* of the former, i.e., part of the concept *cat*. Maybe they are just so related, and nothing more? In other words, perhaps that semantic relation should not be explained by relying on the lexicon and its semantic features? Thus, from the mere existence of some relations, one cannot infer that these relations are constitutive in whatever appears in that relation (i.e., DNA/RNA). Quite the contrary, the important question, at least since Kant, has been precisely which of these many conceptual relations, which are obviously there, are meaning constitutive (analytic) and which are not (synthetic) (See Fodor, 1998a, chapter 3, for a more lengthy argument along these lines).

The referee also points out that such semantically coherent structures emerge in "learning simulations" performed by the Markov-processes, such as self-organisation nets (see Kohonen, 1995). But again, *that* a semantic structure, or a semantic feature, emerges does not establish that the structure or feature is part of the meaning of lexical items. Similarly, an argument that DNA is related

to RNA is not *ipso facto* an argument for the claim that DNA is constituted by RNA, so that RNA would be part of DNA; on the contrary, they are two different chemical compounds. So is it perhaps possible that cat and animate are *two different concepts*, even though they *are related to each other*? This is what compositionality, along with systematicity and productivity, suggests.

The most important objection that I take to be unanswered currently and worthy of closer look is as follows. If lexical meaning is atomistic, how can one *explain* the fact that certain inferences, that seem to be based on lexicon, such as ‘kill entails cause to die’, are true? Why do they seem ‘analytic’ to us? If they do not emerge from the lexicon, then they must emerge from some other source; if so, from where?

Typically, such inferences are explained by relying upon one-way meaning postulates instead of complete decompositions. Thus, the fact that killing entails dying is captured by assuming that it is a necessary truth that killing involves dying:

$$\_ (x, y)(x \text{ kills } y \rightarrow y \text{ dies})$$

So far so good, but meaning postulates are problematic in that, *presupposing meanings*, they come for free as vehicles for analysing meaning-related phenomena: one can assume a meaning postulate whenever needed for whatever purpose. Meaning postulates are therefore *not explanatory*. Consider the fact that languages invoke a productive and systematic system of causatives (*kill* → *cause to die*) with their fairly intricate syntactic, semantic and morphological properties. These properties appear clearly at the “lexical level.” But insofar as this system is subject to laws and principles, there is nothing one could do about it at the lexical level by merely assuming that ‘there could be meaning postulates.’ Nothing of much interest, or any generality, follows from this.<sup>13</sup>

---

<sup>13</sup>Thus, according to Chomsky, “to the extent that anything is understood about lexical items and their nature, it seems that they are based on conceptual structures of a specific and closely integrated type” so that there is “a priori framework of human thought, within which language is acquired”, providing “necessary connections among concepts, reflected in connections of meaning among words” (Chomsky, 2000: 62-63). Prima facie, this is surely so. He then argued, quite correctly in my view, that, assuming that e.g., the causative system is not a result of an internal structure of lexical items “establishes nothing unless it is shown that alternative approach in terms of some [...] theories of belief fixation or semantic importance” can be developed (Ibid., p. 64). He is referring to “theories of belief fixation” since the approach based on meaning postulates must invoke beliefs rather than lexical meanings. In effect, if we rely upon meaning postulates to capture the overwhelming

Then, these inferential relations must originate from some other source than the lexicon. What is that source; how do they emerge, and why? Lexical atomism thus raises new problems, perhaps even more difficult and more numerous than those it solved.

## References

- Armstrong, S. L., Gleitman, L. R., & Gleitman, H. (1983) What some concepts might not be. *Cognition* 13: 263-308.
- Barsalou, L. W. (1992) Frames, Concepts, and Conceptual Fields. In A. Lehrer & E. F. Kittay (eds.), *Frames, Fields, and Contrasts*, pp. 21-74. New Jersey: Lawrence Erlbaum.
- Bever, T. G., Fodor, J. A., & Weksel, W. (1965) On the Acquisition of Syntax: A Critique of "Contextual Generalisation". *Psychological Review* 72: 467-482.
- Block, N. (1986) Advertisement for a semantics for psychology. *Midwest Studies in Philosophy* 10: 615-78
- Braine, M. D. S. (1963) On Learning the Grammatical Order of Words. *Psychological Review* 70: 323-348.
- Carnap, R. (1937) *Logical Syntax of Language*. London: Kegan Paul.
- Charles, D. (2000) *Aristotle on Meaning and Essence*. Oxford: Clarendon Press.
- Chomsky, N. (1957) *Syntactic Structures*. The Hague: Mouton.
- Chomsky, N. (1970) Remarks on Nominalization. In R. A. Jacobs & P. S. Rosenbaum (eds.), *Readings in English Transformational Grammar*, pp. 184-221. Waltham, Mass.: Ginn.
- Chomsky, N. (1975) *The logical structure of linguistic theory*. New York: Plenum Press.
- Chomsky, N. (2000) *New Horizons in the Study of Language and Mind*. Cambridge: Cambridge University Press.
- Clark, A. (1991) Systematicity, Structured Representations and Cognitive Architecture: A Reply to Fodor and Pylyshyn. In T. Horgan & J. Tienson (Eds.), *Connectionism and the Philosophy of Mind*. Boston: Kluwer Academic Press.
- Cocchiarella, N. (1986) *Logical Investigations of Predication theory and the Problem of Universals*. Milano: Bibliopolis.
- Cowie, F. (1999) *What's Within?*. New York: Oxford University Press.
- Cruse, D. A. (1986) *Lexical Semantics*. Cambridge: Cambridge University Press.

---

productivity and systematicity of causative constructions in the world's languages, we need a *theory* which utilizes, in a constrained fashion, such meaning postulates and explains *why* and *how* they originate (in the human mind). Juan Uriagereka put the point as follows: "Part of the job behind discovering the featural makeup of [the logical form] is to determine systematic correspondences in terms of an architecture that ultimately does not need meaning postulates (and relegates them to vaguer, purely world-knowledge relations, like those existing in 'lexical fields' [...]). The task is phenomenally difficult, but recent efforts in very specific domains [...] show that it is not *a priori* impossible, nor is it theoretically different from a similar task on the PF side of the grammar." (Uriagereka, 2000: 583, note 29).



- Di Sciullo, A. & Williams, E. (1987) *On the Definition of Word*. Cambridge, MA.: MIT Press
- Evans, G. (1982) *The Varieties of Reference*. Oxford: Clarendon Press.
- Field, H. (1977) Logic, meaning and conceptual role. *Journal of Philosophy*, 69, 379-408.
- Findler, N. V. (Ed.). (1979) *Associative networks: Representation and use of knowledge by computers*. New York: Academic Press.
- Fodor, J. A. (1970) Three reasons for not deriving "kill" from "cause to die". *Linguistic Inquiry* 1: 429-438.
- Fodor, J. A. (1975) *The language of thought*. New York: Crowell.
- Fodor, J. A. (1981) The Present Status of the Innateness Controversy. In J. A. Fodor (Ed.), *Representations*, pp. 140-177. Cambridge, MA: MIT Press.
- Fodor, J. A. (1995) Concepts: A potboiler. In J. Mehler & S. Franck (Eds.), *Cognition on cognition*, pp. 131-151. Cambridge, MA.: MIT Press.
- Fodor, J. A. (1998a) *Concepts*. Oxford: Clarendon Press.
- Fodor, J. A. (1998b) *In critical condition: Polemical essays on cognitive science and the philosophy of mind*. Cambridge, MA.: The MIT Press.
- Fodor, J. A. (2000) *The mind doesn't work that way. The scope and limits of computational psychology*. Cambridge, MA.: MIT Press.
- Fodor, J. A. (2001) Language, Thought and Compositionality. *Mind & Language*, 16 (1): 1-15.
- Fodor, J. A., Garrett, M. F., Walker, E. C. T., & Parkes, C. H. (1980) Against definitions. *Cognition* 8: 263-367.
- Fodor, J. A., & Pylyshyn, Z. (1988) Connectionism and cognitive architecture. A critical analysis. *Cognition* 28: 3-71.
- Fodor, J. A., & Lepore, E. (1991) Why Meaning (Probably) Isn't Conceptual Role. *Mind & Language* 6: 328-43.
- Frege, G. (1923/1977) *Logische Untersuchungen. Dritter Teil: Gedankenflüge*, Beiträge zur Philosophie des Deutschen Idealismus vol. III, pp. 36-51. Translated as *Compound thoughts, Logical Investigations*, Gottlob Frege, P. T. Geach & R. H. Stoothoff (trans.), Oxford: Basil Blackwell, 1977, pp. 55-78.
- Frege, G. (1948) On Sense and Reference. *The Philosophical Review* 57: 207-230.
- Gopnik, A. (1988) Conceptual and semantic development as theory change. *Mind & Language*, 3 (3): 197-216.
- Gopnik, A., & Meltzoff, A. N. (1997) *Words, thoughts, and theories*. Cambridge, MA.: MIT Press.
- Hampton, J. A. (2000) Concepts and prototypes. *Mind & Language* 15: 299-307.
- Harman, G. (1982). Conceptual role semantics. *Notre Dame Journal of Formal Logic* 23: 242-256.
- Hintikka, J. (2001) Post-Tarskian Truth. *Synthese* 126: 17-36.
- Hintikka, J., & Kulas, J. (1983) *The Game of Language*. Dordrecht: D. Reidel.
- Hodges, W. (1998) Compositionality is not the problem. *Logic and Philosophy* 6: 7-33.
- Horwich, P. (1997) The Compositionality of Meaning. *Philosophical Review* 106: 503-31.
- Janssen, T. M. V. (1997) Compositionality. In J. v. Benthem & A. t. Meulen (Eds.), *Handbook of Logic and Language*, pp. 417-473. Amsterdam: Elsevier.
- Kamp, H., & Partee, B. (1995) Prototype theory and compositionality. *Cognition* 57: 129-

91.

- Keil, F. (1987) Conceptual development and category structure. In U. Neisser (ed.), *Concepts and conceptual development*. Cambridge: Cambridge University Press.
- Kintsch, W. (1974) *The representation of meaning in memory*. Hillsdale, New Jersey: Lawrence Erlbaum.
- Kintsch, W. (1998) *Comprehension*. Cambridge: Cambridge University Press.
- Kohonen, T. (1995) *Self-organizing maps*. Berlin, Springer.
- Margolis, E. (1998) How to acquire a concept. *Mind & Language* 13: 347-369.
- Margolis, E. and S. Laurence (1999) *Concepts: Core readings*. Cambridge, MA.: MIT Press.
- Montague, R. (1974) *Formal Philosophy: Selected Papers of Richard Montague*. New Haven: Yale University Press.
- Moore, G. E. (1903) *Principia Ethica*. Cambridge: Cambridge University Press
- Murphy, G. L. (1988) Comprehending complex concepts. *Cognitive Science*, 12(4): 529-562.
- Murphy, G. L., & Medin, D., L. (1985) The Role of Theories in Conceptual Coherence. *Psychological Review* 92: 289-316.
- Peacocke, C. (1992) *A study of concepts*. Cambridge, MA.: MIT Press.
- Pylyshyn, Z. W. (1984) *Computation and cognition: Toward a foundation for cognitive science*. Cambridge, MA.: MIT Press.
- Quillian, M. R. (1968). Semantic memory. In M. Minsky (ed.), *Semantic information processing*, pp. 216-270. Cambridge, MA.: MIT Press.
- Rips, L. J. (1995). The Current Status of Research on Concept Combination. *Mind & Language* 10: 72-104.
- Rosch, E. (1973a) Natural categories. *Cognitive Psychology* 4(3): 328-350.
- Rosch, E. (1973b) On the internal structure of perceptual and semantic categories. In T. E. Moore (Ed.), *Cognitive development and the acquisition of language*, pp. 17-34. New York: Academic Press.
- Russell, B. (1919) On Propositions: What They are and How they Mean. *Aristotelian Society Supplementary Volume* 2: 1-43.
- Russell, B. (1940) *An Inquiry into Meaning and Truth*. New York: W. W. Norton
- Searle, J. R. (1958) Proper names. *Mind* 67: 166-173.
- Smith, E., & Osherson, D. (1984) Conceptual combination with prototype concepts. *Cognitive Science* 11: 337-361.
- Smith, E. E., & Medin, D. L. (1981) *Categories and concepts*. Cambridge: Harvard University Press.
- Tarski, A. (1935/1956) Der Wahrheitsbegriff in den formalisierten Sprachen, *Logic, Semantics, Metamathematics*. Oxford: Oxford University Press.
- Uriagereka, J. (2000) *Rhyme and Reason: An Introduction to Minimalist Syntax*. Cambridge, MA.: MIT Press.
- Weitz, M. (1988) *Theories of Concepts*. London: Routledge.
- Wittgenstein, L. (1958) *Philosophical Investigations*. Oxford: Basil Blackwell.
- Wittgenstein, L. (1961) *Notebooks, 1914-1916*. (ed.) G. von Wright and G. E. M. Anscombe. Oxford: Basil Blackwell.
- Zadrozny, W. (1994) From compositional to systematic semantics. *Linguistics and*

*Philosophy* 17: 329-342.

Zwicky, A. (1978) Arguing for Constituents. *Papers from the 14th Regional Meeting, Chicago Linguistic Society*: 503-512.