

# How to Compose Contents

## A Review of Jerry Fodor's *In Critical Condition: Polemical Essays on Cognitive Science and the Philosophy of Mind*

Markus Werning  
Department of Philosophy  
Chair of Theoretical Philosophy  
Heinrich-Heine-University Düsseldorf  
Universitätsstraße 1, Bldg. 23.21  
40225 Düsseldorf  
GERMANY

werning@phil-fak.uni-duesseldorf.de

Copyright (c) Markus Werning 2002

PSYCHE, 8(07), September 2002  
<http://psyche.cs.monash.edu.au/v8/psyche-8-07-werning.html>

KEYWORDS: concepts, compositionality, systematicity, productivity, connectionism, language-of-thought, functionalism, semantics.

REVIEW OF: Jerry Fodor (2000) *In Critical Condition: Polemical Essays on Cognitive Science and the Philosophy of Mind*. Cambridge, MA: MIT Press, \$18 pbk. 229 pp. ISBN: 026256128X. [Note: the review refers to the hardback edition of this book].

ABSTRACT: The paper critically reviews Jerry Fodor's book *In Critical Condition: Polemic Essays on Cognitive Science and the Philosophy of Mind*. It focuses on Fodor's compositionality arguments and their relevance to the following questions: (i) How should concepts be individuated? (ii) What has semantics to do with epistemology? (iii) Who is right in the debate over classical and connectionist theories of cognition? (iv) How can the semantic properties of a mental state be inherited from the semantic properties of the state's constituents? The paper finally argues that Fodor's opposition to functional role semantics might jeopardize his view that semantic compositionality requires appropriate constituent relations between complex and less complex concepts.

If your research interests, in any respect, are about things with semantic properties, such as language, thought, brains, or computers, many books may concern your particular subject, but only few authors go after the big issues. One of them is Jerry Fodor. Although he is a professional philosopher, his writings bear heavily on disciplines as unlike as linguistics and neurobiology, or literary theory and artificial intelligence. The

reason for this immense scope is his work on what's called compositionality: a necessary property of all systems that show a certain systematic correlation between content bearers. An example for the systematicity of the human mind is the law-like correlation between its capacity to think that a red square is in a green circle and its capacity to think that a red circle is in a green square. A systematic correlation of this sort can only be achieved if the mind comprises primitive content bearers and some means to compose complex content bearers from primitives in a compositional way, i.e., in such a way that the semantic values of the complex state are determined by, and dependent on, the semantic values of the primitive states plus syntax. The semantic value of the complex concept A RED SQUARE IN A GREEN CIRCLE is a function of the semantic values of the primitive concepts RED, GREEN, SQUARE, CIRCLE, and IN as well as the syntax scheme [(adjective noun) preposition (adjective noun)]. The primitive concepts contribute all there is to their semantic value to the complex concept and the complex concept derives all there is to its semantic value from its primitive constituents. Symmetric considerations apply to brain states, natural languages, and computer architectures insofar as they are systematic. Compositionality also is required for representational systems to be productive: A system can impossibly generate infinitely many representations from a finite basis unless complex representations inherit their semantic values exclusively from the semantic values of the primitive constituents and unless the primitive constituents completely contribute their semantic values to the complex whole. The only exception to this principle are finitely many idioms. The expression "red herring" is an idiom precisely because its semantic value - which has to do with its referents to draw attention away from the central issue - is not determined by the semantic values of "red" and "herring".

In Fodor's recent book you learn why compositionality is the critical condition for a hell lot of theories which up to date have been dominating their disciplines. It is amazing how mercilessly the requirement of compositionality is blasting off long-standing traditions in the philosophy of mind and language, in psychology, and computer science. Take only the view that the meaning of words is a function of their inferential role (The following considerations, *mutatis mutandis*, also apply to mental concepts). Many of the most influential philosophers of the 20th century hold that the inclination to draw the right inferences from sentences is constitutive for word meaning: In this regard, not only Wittgenstein, Quine, Davidson, Rorty, and Dennett, but also Heidegger, Derrida, and Saussure play in the same team. Fodor (pp. 27-33) shows that this view ultimately depends on the analytic/synthetic distinction, i.e., the distinction between truth of meaning and truth of empirical fact. The reason is that complex expressions often have inferential roles that cannot be derived from the inferential roles of their primitive constituents. The expression "British cow" may license the inference to "dangerous" if you believe that British cows are dangerous - say because you believe they are likely to be infected by mad-cow disease or foot-and-mouth disease. However, neither "British", nor "cow" licenses any inference to "dangerous". The inferential role of "British cow" is not determined by the inferential roles of "British" and "cow". To reconcile the principle of compositionality with the view, that inferential role constitutes meaning, one cannot but fall back on the distinction between inferences that are constitutive for meaning and those that are not. For, compositionality requires that all that's constitutive for the meaning of a complex expression is inherited from what's constitutive for the meaning of

the primitive constituents (plus syntax). Since the disposition to draw the inference "British cow --> dangerous" is not inherited from the disposition to draw inferences about cows and British things, it cannot be constitutive for the meaning of "British cow". Even though not all inferences are constitutive, there ought to be, at least, some that are. Otherwise, inferential role semantics would lack any justification. Which are the good candidates, then, for meaning-constitutive inferences? Fodor wants to exclude the trivial inferences of the sort "British cows --> cows", of course. What else, then, is left than the analytic truths like "unmarried man --> bachelor"? If you make this move, the view that meaning is constituted by dispositions to draw inferences, at the end of the day, depends on the distinction between analytic and synthetic truth. And this is a distinction most people, following W. V. Quine (1951, 1960), think cannot be made. Fodor's argument goes even further, though. He lays the axe to the root of this misconceived view of semantics. The idea that meaning is inferential role originates in the halfhearted way the linguistic turn was made by most philosophers of the passed century. Modern philosophy was mostly about epistemology whereas recent philosophy is mostly about meaning (or content). The question "How can we know that P?" was substituted by the question "How can we say (think or judge) that P?" The link between both projects were so-called transcendental arguments: They started out epistemologically: "If it weren't the case that P, we couldn't know that Q; but we do know that Q; therefore P"; and ended up semantically: "If it weren't the case that P, we couldn't say (or think or judge) that Q; but we do say (or think or judge) that Q; therefore P." Those semantical question were regarded to be about certain epistemic capacities like the kind of inference subjects are to draw, or about what objects subjects are to recognize under which conditions. This way semantics was contaminated with epistemology. About one quarter of Fodor's book is the attempt to dismiss this historical load. In this context he also argues against the existence of recognitional concepts (pp. 35-62). Not even the concept RED is constituted by the possession of recognitional capacities. The view that there are recognitional concepts is extremely wide-spread and deep-rooted in philosophy and psychology. It's probably the main heritage of empiricism. Again, Fodor employs a compositionality argument to undermine this view. The notion of a capacity to recognize something as a soandso depends on the notion of a good instance of a soandso. The capacity to recognize certain objects as red depends on the fact that these objects are good instances of red. No capacity is so perfect as to recognize all and only red things as red. Goodinstancehood does, however, not distribute, as Fodor calls it: If something is a good instance of red hair, it need not to be a good instance of red and of hair, and usually is not. For, red hair is not typically, but only relatively red. The capacity to recognize something as red hair, thus, is not inherited from the capacities to recognize something as red and as hair, respectively. If recognitional capacities were constitutive for concepts, they, however ought to, because complex concepts do inherit their semantic values from their primitive constituent concepts. If you are pretty good in recognizing hair and red things, you may well be rather bad in recognizing red hair. Since the possession of the concept RED HAIR, according to the principle of compositionality, is determined by the possession of the concepts RED and HAIR (and some syntactic tools), and the capacity to recognize red hair, is not determined by the capacity to recognize red things and hair, the possession of the concepts RED, HAIR, and RED HAIR is independent from the possession of recognitional capacities.

Compositionality also looms large in Fodor's criticism of connectionism (pp. 81-125, pp. 91-111 are co-authored with Brian McLaughlin). Here, his main opponent is Paul Smolensky (1987, 1988) who challenged Fodor and Pylyshyn (1988). Before this challenge connectionism was widely acknowledged to provide explanations of brain functions on a sub-cognitive level. Smolensky, however, put an end to this modesty. He claimed that connectionism was strong enough to contribute to the cognitive level of explanation. In contrast to the chemical or the neurophysiological level, the cognitive level is defined as the level at which representations and processes over representations figure in explanations. If connectionists want to contribute to the cognitive level of explanation, they have to subscribe to the principle of systematicity. Over the years Smolensky has developed various accounts of how connectionist architectures might be considered representational. According to an earlier version (Smolensky 1987, 1988), representations like COFFEE and CUP WITH COFFEE are activity vectors over units which represent micro-features (units like BROWN, LIQUID, MADE OF PORCELAIN, etc.). Units are nodes with a certain value of activation. In a connectionist architecture these units are connected to other units. The connections carry impeding or exciting weights. According to Smolensky, CUP WITH COFFEE representations contain COFFEE representations as non-classical constituents in the sense of a component vector. The method of composition is vector addition. Furthermore, COFFEE and presumably every other representation is supposed to be context-dependent. The activity vector that is the COFFEE representation in CUP WITH COFFEE is supposed to be not identical to the activity vector that contributes COFFEE in GLASS WITH COFFEE. Smolensky calls this a weakly compositional account of representation. Although the exposition of this theory is pretty theoretical, it closely resembles what many neurobiologists have in mind when they talk of neuronal assemblies. Smolensky, thus, expresses a still very popular view. The problem with weak compositionality, as Fodor points out, is that it does not render systematicity. Assume that somebody is capable of the representation CUP WITH COFFEE AND GLASS WITH TEA. Systematicity implies that s/he is also capable of the representation CUP WITH TEA AND GLASS WITH COFFEE. That somebody who is capable of the first representation is also capable of the second is, however, not warranted by weak compositionality. For, somebody who is capable of the vector COFFEE in the context of CUP WITH COFFEE need not be capable of the vector COFFEE in the context of GLASS WITH COFFEE. Recall that COFFEE vectors need not be identical if they occur in different contexts. Weak compositionality, hence, does not suffice for systematicity. Connectionism thus understood does not reside on the cognitive level of explanation.

To match this objection, Smolensky (1995) developed a completely new model, which he calls integrated connectionist/symbolic architecture (ICS) and which he takes to be strongly compositional. Here, complex vector representations are constructed from primitive vector representations in a way that employs vector operations and role vectors. The story is rather complicated and need not be told here. For our purposes only two features of ICS deserve attention: First, ICS is at least syntactically compositional insofar as the form of complex representations is determined by, and dependent on, the form of the primitive representations. By the form of representation I mean the pattern of activity that corresponds to the representation in question and is expressed by a vector. Second, in

ICS primitive representations aren't proper constituents of complex representations. The relation of proper constituency is understood as a necessary co-tokening relation: A state C is a proper constituent of a state W if and only if, necessarily, whenever W is tokened, C is tokened. The dispute between Fodor and Smolensky, then, boils down to two questions: (1) Is syntactic compositionality sufficient for systematicity? (2) Is proper constituency necessary for systematicity? If the answer is yes to the first and no to the second question, Smolensky wins the battle. If the answer is no to (1) and yes to (2), Fodor comes off triumphant. In all other cases, not only Smolensky, but also Fodor seems to be in trouble. If there is a yes to both questions, Fodor's compositionality arguments fall to pieces because he always has in mind semantic compositionality when he talks of compositionality. If syntactic compositionality suffices for systematicity, Fodor lacks the premise that semantic compositionality is necessary. If both questions deserve a negative answer, Fodor is left without a positive argument for his language-of-thought doctrine. Since it is essential to this doctrine that complex concepts have primitive concepts as proper constituents, there will be possible alternatives to the language-of-thought doctrine if proper constituency is not necessary for systematicity.

Unfortunately, Fodor does not really articulate both questions. Though he is a little bit more explicit about the second question, he is altogether beating about the bush. Nevertheless, do I think that Fodor wins. The reason for my optimism is that systematicity requires semantic, rather than syntactic compositionality and semantic compositionality requires primitive representations to be proper constituents of complex representations. Let me first explain the difference between semantic and syntactic compositionality. Consider the following mappings between English expressions:

- (i) ("unmarried", "man") --> "unmarried man",
- (ii) ("unmarried", "man") --> "bachelor",
- (iii) ("red", "herring") --> "red herring".

Number (i) is both syntactically and semantically compositional because the form as well as the semantic value is conveyed through the mapping. The mapping (ii) only is semantically compositional, provided that the sentence "Unmarried men are bachelors" is analytic. The word "bachelor" derives its semantic value, not its form, from the expressions "unmarried" and "man". Finally, (iii) is syntactically, but not semantically compositional because "red herring" is idiomatic.

Consider the two thoughts that the man with a red coat was distracted by an old herring and that the man with an old coat was distracted by a red herring. The thoughts are syntactically composed of the same concepts. However, they aren't nomologically correlated. Someone who is capable of thoughts about red coats and old herrings may well be unable to think of red herrings. This is because the representation of a red herring fails to be semantically composed of the representations of which it is syntactically composed. The concept RED HERRING is idiomatic. Syntactic compositionality does,

thus, not suffice for systematicity. The answer to question (1) is no. What's required for systematicity, apparently, is semantic compositionality.

To win the battle against Smolensky without casualties, Fodor cannot set his heart at rest, already, but has to argue further that proper constituency is required for systematicity. A possible argument, which, however, contains a pill to swallow for Fodor, runs as follows.

(I) The semantic value of a representational state is dependent on nothing but the causal role of the state.

(II) The causal role of a state is dependent on nothing but the causal roles of its proper constituents (and their relation to each other)

Hence:

(III) If the semantic value of a state A is dependent on nothing but the semantic values of the states B and C, then B and C are proper constituents of A.

The argument is valid. Its second premise is backed by deep-rooted metaphysical intuitions. They even apply to microphysics where the causal role of molecules is dependent on nothing but the causal roles of their constituent atoms (and their relation to each other). The first premise, however, is equivocal to functional role semantics (FRS), a doctrine people like Fodor "who hate FRS root and branch" (p. 71) are unlikely to use in arguments of their own. I do, however, not see how the conclusion can be defended without using functional role semantics. Maybe it's this threatening inconsistency that lets Fodor beat around the bush in his argumentation against Smolensky. Notice that Fodor really is in need for a good argument in favor of constituency relations as a requirement for systematicity, and hence compositionality. If there wasn't any argument of this sort, to which Fodor can consistently appeal, the language-of-thought doctrine would come under fire. What's going wrong here, probably, is an equivocation about what functional role semantics means. Fodor uses this term both for inferential role semantics and causal role semantics (cf. p. 70). We have already seen that inferential role semantics does not comply with compositionality. Whereas the causal role of a representational state determines its inferential role, the inferential role need not determine the causal role. There is more to the causal role of a representational state than the fact that this state figures in dispositions to draw inferences. Representations may cause all sorts of things: emotions, behavior, an increase of heart beat etc. They may also be triggered by things other than representations: retinal stimuli, acoustic waves, dopamine, etc. Inferential role and causal role aren't isomorphic to each other. Furthermore, causal role, apparently, composes well, whereas inferential role, as we've seen, doesn't. The causal role of a complex state is entirely determined by, and solely dependent on, the causal role of its proper constituents and their relations to each other.

Fodor himself, from time to time, favors an externalist theory of semantic value, which has causal elements, too, but only external ones. His story goes something like this (cf. p.

63): The fact that a certain internal state - call it DOG - represents dogs as dogs is determined by, and dependent on, the fact that DOG stands in a certain causal or nomic (i.e., law-like) relation to dogs. There are well known difficulties with this position (Why aren't SMOKE and FIRE synonymous, although, by causal law, there is smoke whenever there is fire and vice versa?). Those complications may, however, be accounted for somehow. The problem I want to highlight has to do with compositionality. Does nomic connectedness distribute? Assume that not only DOG is nomologically connected to dogs, but that also ROBOT is nomologically connected to robots. Does this necessitate that ROBOT DOG is nomologically connected to robot dogs? The laws that connect DOG to dogs and ROBOT to robots are at best probabilistic. No biological system is so reliable that it causally connects all and only dogs to the internal state DOG. No landscape is so unruffled that it could not, in principle, hide a dog from your perceptual apparatus. Probability, however, does not distribute properly. Let  $P(\text{SOANDSO})$  be the probability of a certain system to trigger the internal state SOANDSO if and only if there is a soandso in the system's environment (You may well substitute the logical constant "if and only if" by "if" or by "only if" or some other logical operation, e.g., counterfactual dependency). For the connection between an internal state and an external object to count as nomic or causal the corresponding probability should exceed a certain limit value  $L$ . Having said this, it follows from mathematics that probability does not distribute properly. The reason is that the probability of a conjunction is usually lower than the probabilities of each conjunct. Therefore,  $P(\text{ROBOT DOG})$  is smaller than  $P(\text{ROBOT})$  and  $P(\text{DOG})$ . The more conjuncts you have, the less reliable the connection between the concept conjunction and the conjunct object gets. It may well be that  $P(\text{ROBOT})$  and  $P(\text{DOG})$  exceed the limit value  $L$ , but  $P(\text{ROBOT DOG})$  doesn't. In fact, one can always find a conjunction of internal states/concept  $C(1) \dots C(n)$  such that  $P(C(i))$  exceeds  $L$  for each index  $i$ , whereas  $P(C(1) \& \dots \& C(n))$  is lower than  $L$ . The externalist theory of semantic value is in a row with inferential role semantics, the recognitional theory, and other theories not mentioned here (e.g., prototype theory) insofar as it fails to satisfy the compositionality requirement. The best candidate to comply with compositionality seems to be causal role semantics, after all. But Fodor won't agree.

*In Critical Condition* is a collection of reviews, journal articles, and replies most of which have been published elsewhere within the last decade. Thematically it is more heterogeneous than this review reveals. Apart from compositionality, the book deals with metaphysical issues, the innateness of concepts, the modularity of mind, and evolutionary theory. You will also learn why we don't think in English but in Mentalese, and whether science is biologically possible. Most of the time, Fodor either attacks or replies to opponents, giving the book a ping-pong format. Since many chapters originally address an either lay or interdisciplinary audience, you need not be a heavy-duty philosopher to follow the course of argument. Aside from philosophy, Fodor's talent is humor. I haven't read a book written by a professional philosopher that tells more jokes per page. This reflects the sort of pessimism the title conveys: "Nietzsche is righter than Brecht: Sometimes the man who laughs *has* heard the terrible news" (p. x).

## References

Fodor, J. A. & Pylyshyn, Z. W. (1988). Connectionism and cognitive architecture: A critical analysis. *Cognition*, 28, 3-71.

Quine, W. V. (1951). Two Dogmas of Empiricism. *The Philosophical Review*, 60, 20-43.

Quine, W. V. (1960). *Word and Object*. Cambridge, MA: The MIT Press.

Smolensky, P. (1987). The Constituent Structure of Mental States: A Reply to Fodor and Pylyshyn. *Southern Journal of Philosophy*, 26, 137-160.

Smolensky, P. (1988). On the proper treatment of connectionism. *The Behavioral and Brain Sciences*, 11, 1-74.

Smolensky, P. (1995). Connectionism, constituency and the language of thought. In Macdonald, C., & Macdonald, G. (Eds.), *Connectionism*. Cambridge, MA: Blackwell.