

The Language of Thought and the Embodied Nature of Language Use

Norman Y. Teng

Philosophical Studies, 94(3), 237-251.

Abstract

Fodor argues that the systematicity of language use provides a solid ground for his language of thought (LOT) hypothesis. This paper re-examines the notion of systematicity and argues that Fodor underestimates the systematicity. The systematicity cannot be explained merely by what is in the brain, but is inherently connected to an environment, in the sense that (1) language use depends on the perceptual and motor skills for detecting and confronting certain recurring patterns in an environment; and (2) the compositionality of language use is grounded in our bodily interactions with the environment, which provide a set of highly correlated principles that account for our basic sense of the systematicity and function as a general guideline for what a particular language will look like.

The Language of Thought and the Embodied Nature of Language Use*

This paper attempts to clarify and critically examine Fodor's language of thought (LOT) hypothesis, focusing on his contention that the systematicity of language use provides a solid ground for the LOT hypothesis. The usual response to Fodor's systematicity arguments is to argue that systematicity is much less than Fodor assumes. And the sort of systematicity that Fodor draws our attention to is a language-based artifact constructed from a top-level task analysis of what a cognitive system does. The

second tact, which I am willing to support, is to argue that it is by mastering an external symbol system that one acquires the systematicity Fodor identifies (more on this later). But this paper offers something more radical -- there is much more systematicity in language, and this cannot be accounted for by Fodor's LOT hypothesis. (1) The systematicity, I will argue, cannot be explained merely by what is in the brain, but is inherently connected to an environment, in the sense that (1) language use depends on the perceptual and motor skills for detecting and confronting certain recurring patterns in an environment; and (2) the compositionality of language use is grounded in our bodily interactions with the environment, which provide a set of highly correlated principles that account for our basic sense of the systematicity and function as a general guideline for what a particular language will look like. (cf. Johnson, 1987, 1992, 1993).

Arguments for the Language of Thought Hypothesis

The most straightforward and convincing argument for the linguistic nature of thoughts concerns the following three basic features of language use (see Fodor, 1975, 1980, 1987; Fodor & Pylyshyn, 1988; Fodor & McLaughlin, 1989, among others). First, language use is productive in that the range of well-formed linguistic expressions we can produce is indefinite. We don't memorize them one by one. There have to be some sort of principles that guide our linguistic production and comprehension. Second, language use is systematic in that the ability to understand some sentences is intrinsically connected with the ability to understand certain other sentences with semantically related contents. Here, too, some sort of principles that guide our linguistic production and comprehension must be in place. Third, certain inferential patterns are embedded in the ways we use a language. We don't, for example, find people who are prepared to infer 'John went to the store' from 'John and Mary and Susan and Sally went to the store' and from 'John and Mary went to the store' but not from 'John and Mary and Susan went to the store'. Some sort of principles that sanction such inferences must figure in everyday language.

What then are the principles that support the productivity, systematicity, and inferential coherence of language use? It is clear that language use is compositional. Without going into details, we can assume that principles of compositionality in language use account for these three features: productivity, systematicity, and inferential coherence. To illustrate this point, consider the following example. The sentence 'John loves the girl' is systematically related to the sentence 'the girl loves John' in that you don't find a native speaker who has the ability to understand the first sentence but does not have the ability to understand the second. The systematicity can be explained if both sentences are composed from the same constituents in accord with the same set of compositional rules. Presumably, the constituents in this example are 'John', 'loves', 'the', and 'girl', which function as invariants in changing contexts and make the same semantic contribution to both sentences in accord with the same compositional rules. Moreover, one can produce an indefinite number of linguistic expressions by combining invariant constituents in accord with the compositional rules. Also one can infer 'John went to the store' from 'John and Mary went to the store', because, when appropriately analyzed, the first expression is a constituent of the second and retains its semantic value under simplification of conjunction.

These three basic features of language use, according to Fodor, must in some way exhibit certain features of the thoughts that one can entertain, and can be best accounted for if we postulate that, prior to language, one has a mental representation system that is compositional. And the principles of compositionality in everyday language are derived from, or at least supported by, a mental representation system that is inherently compositional.(2)

So far this argument does not show that the supposed language of thought is a different language from the one that we ordinarily use. In fact it highlights a similarity between the LOT and everyday language in that both are compositional (Kaye, 1995). But the similarity stops here. The mental process that

operates on mental representations is supposed to be a symbol-computing process. What matters in this computational process is the formal structure of mental representations, which has to be amenable to algorithmic characterization so that the computational brain can operate on it. So, in addition to compositionality, the LOT has to have a recursive representational format. For the brain to process linguistic information, the LOT has to be realized in the physical structures or has to emerge as a distinctive level of cognitive organization from the interaction of neurons in the brain. It follows that productivity, systematicity, and inferential coherence are further constrained by a recursive representational format, which by its nature has predictable syntactic consequences and thus constitutes an autonomous representational scheme, whose structure is impervious to any influence from peripheral systems and invariant with changing contexts.

The Embodied Nature of Language Use

Let me clarify what is in dispute. I agree that the systematicity of everyday language must in some way exhibit (and, in my view, also structure) certain features of the thoughts we entertain. I also agree that systematicity depends on compositionality, and language is compositional. I contend, however, that Fodor underestimates systematicity of language use, in that the systematicity is to be located not just within the language system itself, but also in the recurring patterns that arise from the process of human-environment interactions. It follows that the LOT hypothesis is not as well-established as Fodor intended to demonstrate in his analysis of systematicity.

Let us start with the following example (cf. Jackendoff & Landau, 1992; Talmy, 1983):

- (1) a. The cup is under the table.
b. ??The table is under the cup.

If we restrict our investigation to the constituent structures of (1a) and (1b) as suggested by Fodor, it seems unavoidable that we would be forced to say that these two sentences are systematically related to each other in that we wouldn't find a native speaker who has the ability to understand (1a) but does not have the ability to understand (1b). However, (1b) hardly makes good sense to us.

Let us take a closer look at the above linguistic phenomenon before examining a number of possible ways that Fodor might respond. When we say that the cup is under the table, we locate the cup as directly below the table but some distance away from it. The table here serves as a reference point that anchors a region within which the cup is to be located. Given the perceptual and motor capacities we have in detecting and confronting environmental features, the table, in serving as a reference point, has to be more salient than the target we are searching for, other things being equal. But if it is the table that is the target to be located, we can hardly use the cup as a reference point. In short, language use exhibits some sort of systematicity sensitive to our perceptual and motor interactions with the environment (see Langacker, 1991a, pp. 167-180).

Before we proceed, it should be noted that believers in the LOT hypothesis might not consider the above evidence to be important. (This might stem from the nature of the dialectic between the generative grammarians and the cognitive grammarians -- that evidence that seems basic to one side tends to be dismissed as irrelevant or unconvincing by the other.)⁽³⁾ The LOT hypothesis, as they propose, can be considered as part of a set of unifying, explanatory principles, in which the mind is viewed as operating upon symbolic representations. If the alignment of target and reference point described above is indeed adequate, the alignment is to be thought of as a reflection of a more fundamental structure of our cognitive ability, which will eventually explain away the alignment as an incidental product of our cognitive processes. Some proponents of the LOT hypothesis might therefore suggest, "Don't be distracted by incidental products, concentrate on important issues." In my view, if the LOT hypothesis were not yet sufficiently clear, that might be a cogent reason to stop the debate. But since the LOT hypothesis is already sufficiently clear and well developed, it is now profitable to

examine whether or not the alignment is merely an incidental product.

Now, proponents of the LOT hypothesis can respond to the above analysis of the alignment of target and reference point in one of four possible ways.

First Possibility

The first way is to argue that the word `under' is in fact ambiguous. There are two mental lexicons, UNDER1 and UNDER2, that account for the different uses of `under' in (1a) and (1b), and thus dissociate (1a) from (1b). This response looks implausible, for (1b) does not make good sense to us and thus is unmotivated in actual usage events. It is used here only to demonstrate the point I want to make by reversing the order of the constituent structure of (1a). We don't need to postulate a distinct mental lexicon to legitimize this unmotivated usage.

Second Possibility

The second way the proponents of the LOT hypothesis can respond here is to claim that there is indeed only one mental lexicon UNDER, but the alignment of target and reference point is part of its semantic content. The advantage of this second response is that it preserves the intuition that (1a), but not (1b), makes good sense to us. And we can generalize this case over an array of the supposed mental lexicons related to locative expressions, like ABOVE, BELOW, BEHIND, BESIDE, IN FRONT OF, and so on.

Note, however, that the alignment of target and reference point not only shows up in locative expressions, but also manifests itself in other linguistic constructions. For example, a house as a whole is more salient than the roof which is part of it, and thus can serve as a reference point relative to the roof. This accounts for the fact that the possessive expression `the roof of the house' is well motivated,

but 'the house of the roof' is unmotivated. For the same reason, the alignment of target and reference point also accounts for the fact that 'the girl's neck' is well motivated, but 'the neck's girl' is unmotivated. And so on and so forth. Moreover, we can sometimes reverse the alignment and use 'the fleas' cat' to say something from the viewpoint of the fleas' welfare, although it is the expression 'the cat's fleas' that conforms to the requirement of the alignment (Langacker, 1991a, pp. 167-180). Those linguistic phenomena show that the alignment of target and reference point cuts across a diverse array of linguistic constructions and is sensitive to systematic variations of contexts. It is thus implausible that the alignment is part of the semantic content of any mental lexicons conforming to Fodor's LOT hypothesis.

Third Possibility

The third way that the proponents of the LOT hypothesis can respond is to argue that, given the pervasiveness of the alignment of target and reference point, the alignment is part of the compositional rules that apply to a set mental lexicons. If this indeed is a case, a LOT cannot constitute an autonomous representational scheme. Notice that the structure of a LOT that accounts for our linguistic capability is impervious to the recurring patterns emerging from our perceptual and motor interactions with the environment. But the requirement that a reference point has to be more salient than the target is in large measure dependent on the perceptual and motor skills we develop in coping with the environment. That is to say, the compositionality of language use that respects the alignment is in large measure constrained by how we negotiate the environment with our perceptual and motor skills.

The last point is worth pursuing further. The alignment of target and reference point is only part of the recent discovery that many structural patterns of everyday language are based on a small number of schemas that govern our perceptual and motor interactions with the environment, plus certain metaphorical cross-domain conceptualization. For example, consider the following parallel structural

pattern of spatial and temporal semantic fields (Jackendoff, 1983, p. 189):

- (5) a. at 6:00
from Tuesday to Thursday
in 1976
on my birthday
- b. at the corner
from Denver to Indianapolis
in Cincinnati
on the table

Spatial schemas that govern our experience of places and paths are mapped onto our temporal experience, and thus motivate a parallel structural pattern of our linguistic expressions about space and time. Indeed, a vast number of language usages that reflect such patterns are highly regimented and do not vary in changing contexts.

Pinker (1989, pp. 370-373) observes that those highly regimented structural patterns can be used to push speculations about the nature and evolution of a LOT. Granted that such spatial schemas and their metaphorical extensions are grounded in how we negotiate the environment with our perceptual and motor skills, they can be recruited to new uses in the direction of greater specificity to support more abstract cognition, and thus become part of an autonomous representational system, that is, a LOT.

Pinker's speculation fits the standard explanation of the emergence of novel capabilities in evolution: old parts can be recruited to new uses or reconfigured to suit novel tasks. Thus, this speculation can in principle demystify a fundamental puzzle in the study of the mind: how evolution could have produced

a brain capable of intricate and abstract cognition given the absence of selection pressure for such ability in the course of evolutionary history. In short, although the recurring patterns of our spatial experience are deeply grounded in our perceptual and motor interactions with the environment, they form a base on which more specific structural regimentation can be shaped and copied to the representational format of a LOT.

Notice, however, that under this speculation the syntactic structure of a LOT is highly sensitive to the conceptual contents of spatial schemas, such as PLACE and PATH. The regimentation of those linguistic patterns may simply reflect the fact that those spatial schemas deemed to be parts of a LOT are highly entrenched in our cognitive organizing activity. To render Pinker's speculation plausible, there has to be a principled way to distinguish those linguistic structures that have predictable syntactic consequences from those that are only motivated in virtue of the schemas and their metaphorical extensions.

It is unlikely, I think, that there can be such a principled distinction. There is in fact a diverse array of spatial schemas and metaphorical extensions that are highly regimented and do not vary in changing contexts, but they don't have predictable syntactic consequences required by the LOT hypothesis. The systematicity of such language usages are deeply conceptually oriented.

Take the spatial orientation up and down for illustration. One important environmental feature concerning this spatial orientation is this: gravity pulls downwards; but things can be piled up, and a container can be filled up; and if you add more of a substance or of physical objects to a container or pile, the level goes up. The recurring pattern that emerges from those sorts of human-environment interactions makes up a spatial schema of up and down, and motivates a metaphorical organization MORE IS UP; LESS IS DOWN, which systematically links together a diverse array of language usages by way of cross-domain mapping (Lakoff & Johnson, 1980, pp. 15-16):

The number of books printed each year keeps going up. His draft number is high. My income rose last year. The amount of artistic activity in this state has gone down in the past year. The number of errors he made is incredibly low. His income fell last year. He is underage. If you're too hot, turn the heat down.

Note that the range of those metaphorical usages is indefinite. We don't memorize them one by one. And we can always produce novel metaphorical expressions, say, 'the economy was launched by the government using low interest rates and a rising dollar', that are still within the range. Notice that this expression also exemplifies a very sophisticated reasoning supported by the spatial schema of up and down and the metaphorical organization MORE IS UP; LESS IS DOWN. On the other hand, certain expressions, like 'the price of computers plummets, so you have to pay much more money to buy one', are unmotivated and hardly interpretable, because the implicit inferential pattern embedded in such expressions violates the metaphorical organization MORE IS UP; LESS IS DOWN.

The actual usages of those metaphorical expressions are indeed sensitive to contexts, but the underlying schemas and the basic metaphorical organizations are highly entrenched and do not vary in changing contexts; they are parts of grammar in the sense that they are parts of the knowledge that underlies our linguistic competence and supports the systematicity of language use (Lakoff, 1993). Given the diverse structural patterns of the sentences permitted by such spatial metaphors, it should be clear that they do not have the syntactic consequences required by Fodor's LOT hypothesis.

It is worth mentioning that Langacker (1982) first labeled his theory of cognitive grammar "space grammar". Viewed from the important role of spatial schemas in grammatical constructions, this label might not sound so bizarre as it appeared to be. This space, it should be noted, is not a Euclidean space, but an ecological space in which we live, with all sorts of invariants that emerge from our bodily

interactions with the environment (see Gibson, 1966, 1970, 1979, Ch. 2). Many linguistically relevant spatial invariants, such as place, path, trajectory, landmark, ground, figure, region, alignment, etc., coupled with their metaphorical elaborations, profoundly shape our grammatical constructions (see Lakoff, 1987; Langacker, 1987, 1991a, 1991b; Talmy, 1983, among others). Those environmental features, as they emerge from the process of human-environment interactions, provide a set of highly correlated principles that account for our basic sense of the systematicity of language use and function as a general guideline for what a particular language will look like. Given the spatial schemas we have, it is unlikely that we will have, say, 'a table under a cup'.

Fourth Possibility

The above analysis shows that the compositionality of language use is sensitive to the recurring patterns emerging from our bodily interactions with the environment. A moderate defender of the LOT hypothesis can still reply, "Given that our knowledge of spatial and other real world features affect how we use our language, the systematicity of language use might not be explained in the way Fodor has proposed. But, in order to process these structural features of the world and to make proper inferences about them, we still need a language of thought. How else could we construct the necessary hypothesis and test them? And how could we make sure the highly regimented, syntactic kinds of systematicity are preserved in our language and thought?"(4)

A preliminary comment on this reply is in order. The discussion of the previous three possibilities shows that the systematicity of language use does not really support Fodor's LOT hypothesis. It further suggests that general cognitive activity can occur without a LOT. The defender's reply puts this suggestion into question. Unlike the previous three replies, which are made from a defensive position, this one counterattacks; it demands more working out of the position on which my argument is based. To avoid disappointment, I should confess right away that I do not have a fully adequate answer.

Instead, I shall simply try to show that the position on which my argument is based is a coherent, defensible position. (My defence, as one might expect, is in line with cognitive grammar and connectionism; see Clark, 1993, for a detailed account of how a connectionist network can account for structured thought; see Langacker, 1991a, Ch. 12, on relations between cognitive grammar and connectionism.)

First, the spatial schemas discussed above need not be encoded in a preexisting symbolic representational system. They can be considered as emerging patterns from our interactions with the environment. The idea can be illustrated by the following example. Suppose you want to find a cup. You do not know exactly where it is. But you remember that it is on a table in the kitchen room. You go to the room, find the table, and look for the cup. Now, a pattern emerges from this process of navigating your way around the environment. The kitchen room serves as a reference point for you to find the table, which in turn serves as a reference point for you to find the cup. Your behavior in the environmental setting shows forth a structural pattern of an alignment of target and reference point. The point here is that the alignment is not encoded as a preexisting rule within the brain, but is constitutively tied to a way of orienting oneself in an environment (Clark, 1989, pp. 63-66; Rowlands, 1995).

Second, before you go to the room, find the table, and look for the cup, you might first think about how to find the cup. Again, the room serves as a reference point for you to find the table, which in turn serves as a reference point anchoring a region where you might find the cup. An alignment of target and reference point constrains the way you think about the situation. That does not mean you must have the alignment encoded as a preexisting rule in order to follow it in your thinking. Instead, it suggests that one can process information as if one were acting on or manipulating physical structures in an environment, which in turn suggests that the pattern can be stored in the brain, and thus, when activated, constrains the way one thinks. Presumably, the pattern is stored in superpositional,

distributed neural networks (see Clark, 1993, pp. 17-23, on superpositional storage).

Third, in learning a language, we do not use preexisting, context-free representational elements to construct some hypothesis about the relations between the emerging patterns and the language we learn to use. The relations are not problems to be solved. Instead, in constraining the way we think, those patterns set up problems, creating problem spaces that bias children to entertain certain assumptions about the nature of linguistic structures and about what kinds of evidence can be processed as effective linguistic evidence. (See Clark, 1993, Ch. 9, on how a connectionist network can create new problem spaces as part of its ongoing activity in an environment.) The phrase 'a table under a cup' is a case in point. My analysis predicts that a child in early stages of learning a language would dismiss such a phrase as irrelevant, for it violates an alignment of target and reference point that is so basic to human cognition. (This is an empirical prediction, to be sure. I do not claim that it is already well corroborated.) Note also that the assumptions that a child may entertain need not be encoded in a LOT. And one does not need to follow Fodor's presumption that language learning consists solely in hypothesis formation and confirmation. The assumptions can be considered to be activation patterns across a population of neurons in the brain, which is situated in an environment. And making assumptions in learning a language can be considered to be part of a problem-solving activity, in which the assumptions that can be "targeted" on effective linguistic evidence are used and accepted by default. (See Clark, 1997, for a detailed account of the embodied and environmentally embedded problem-solving activity.)

Fourth, the problem about the preservation of the highly regimented, syntactic kinds of systematicity, as I take it, comes down to this: Without a set of internal context-free representational elements, how can we have combinatorial language systems and structured thoughts? Language, I agree, is combinatorial. Note, however, that the combinations that we may anticipate in using words are in large measure dependent on the construals they symbolize. The spatial metaphor is a case in point, as shown

in the above discussion. We can further purposefully suspend the symbolic relations in our thinking and doing, and use an external linguistic medium as if we were simply using its physical shape so as to highlight certain patterns that would be otherwise difficult to detect. For example, we can train ourselves to see and utilize physical patterns of the following sort to do logical calculation:

A -> B

A

B

This act of suspending requires subtle manoeuvring of perceptual and motor skills and a reflective act of disengagement from our surroundings, often with a purpose to form an "objective" view of what there is or what is to be done. The supposed syntactic regimentation of our language is to be explained in terms of our capability to detect, reproduce, and manipulate external linguistic entities. The external linguistic entities make up a discrete combinatorial system in the way we use them. But the internal resources that are deployed when we are using the external linguistic entities need not be comprised of a system of discrete representational elements. The same account can be applied to structured thoughts if we further agree that we have the capacity to model mentally our environment and the ways we manipulate those external entities. The structuring of our thoughts might be achieved when we manipulate those linguistic entities in the mind. In this view, structured thoughts are not something built upon the supposed representational elements of a LOT, but are emergent features of the cognitive processes that can exploit discrete and structural features of the physical environment, including the external linguistic entities that our ancestors bequeathed to us (cf. Bechtel & Abrahamsen, 1991, Ch. 7; Rumelhart, 1992; Rumelhart et al., 1986; Smolensky, 1987, 1988, 1991; Sweetser, 1990, Chs. 4-5; Vygotsky, 1934/1962). Much more has to be done to make this proposal complete, but I hope it is sufficiently clear for the present purposes.

Concluding Remarks

I examine four possible ways that proponents of the LOT hypothesis might respond to the systematicity that reflects the richly structured, emerging patterns from our interactions with the environment. The first three are non-satisfactory. The fourth response demands a coherent explanation of the position on which my argument is based. I show that the position is defensible. Therefore, I conclude, (1) the systematicity of language use does not really support Fodor's LOT hypothesis, and (2) it is a defensible position that general cognitive activity can occur without a LOT.

Notes

* I am very grateful to William Bechtel and the reviewers for their very helpful suggestions and comments on earlier drafts. Part of the present paper in slightly different form, titled "The language of thought hypothesis: A critique", was presented at the APA Pacific Division meetings in March 1997. My commentator on that occasion was Harry Deutsch, whose defence of the LOT hypothesis helped me to further clarify a number of conceptual issues related to the LOT hypothesis.

1. I am indebted here to an anonymous referee for pointing out this significant contrast.
2. I am indebted to Harry Deutsch and an anonymous referee for helping to clarify the point that the issue here is whether, prior to language, one has a mental representation system that is compositional so as to allow for systematicity, productivity, and inferential coherence in thought.
3. I am indebted here to an anonymous referee.
4. This question was raised by an anonymous reviewer of an earlier draft.

References

Bechtel, W., & Abrahamsen, A. (1991) *Connectionism and the Mind: An Introduction to Parallel Processing in Networks*. Cambridge, MA: Basil Blackwell.

Clark, A. (1989) *Microcognition: Philosophy, Cognitive Science, and Parallel Distributed Processing*. Cambridge, MA: MIT Press.

Clark, A. (1993) *Associative Engines: Connectionism, Concepts, and Representational Change*. Cambridge, MA: MIT Press.

Clark, A. (1997) *Being There: Putting Brain, Body, and World Together Again*. Cambridge, MA: MIT Press.

Fodor, J. A. (1975) *The Language of Thought*. Cambridge, MA: Harvard University Press.

Fodor, J. A. (1980) "Methodological Solipsism Considered As a Research Strategy in Cognitive Science", *Behavioral and Brain Science* 3, 63-109.

Fodor, J. A. (1987) "Why There Still Has to Be A Language of Thought", in *Psychosemantics: The Problem of Meaning in the Philosophy of Mind*. Cambridge, MA: MIT Press.

Fodor, J. A., & McLaughlin, B. P. (1989) "Connectionism and the Problem of Systematicity: Why Smolensky's Solution Doesn't Work", *Cognition* 35, 183-204.

Fodor, J. A., & Pylyshyn, Z. (1988) "Connectionism and Cognitive Architecture: A Critical Analysis",
Cognition 28, 3-71.

Gibson, J. J. (1966) *The Senses Considered as Perceptual System*. Boston: Houghton Mifflin.

Gibson, J. J. (1970) "On Theories for Visual Space Perception", *Scandinavian Journal of Psychology* 11,
67-74.

Gibson, J. J. (1979) *The Ecological Approach to Visual Perception*. Boston: Houghton Mifflin.

Jackendoff, R. (1983) *Semantics and Cognition*. Cambridge, MA: MIT Press.

Jackendoff, R., & Landau, B. (1992) "Spatial Language and Spatial Cognition", in R. Jackendoff,
Languages of the Mind: Essays on Mental Representation. Cambridge, MA: MIT Press.

Johnson, M. (1987) *The Body in the Mind: The Bodily Basis of Meaning, Imagination, and Reason*.
Chicago, IL: University of Chicago Press.

Johnson, M. (1992) "Philosophical Implications of Cognitive Semantics", *Cognitive Linguistics* 3,
345-366.

Johnson, M. (1993) "Why Cognitive Semantics Matters to Philosophy", *Cognitive Linguistics* 4, 62-74.

Kaye, L. J. (1995) "The Language of Thought", *Philosophy of Science* 62, 92-110.

Lakoff, G. (1987) *Women, Fire, and Dangerous Things: What Categories Reveal about the Mind*.

Chicago, IL: University of Chicago Press.

Lakoff, G. (1993) "The Contemporary Theory of Metaphor", in A. Ortony, ed., *Metaphor and Thought*, 2nd ed.. New York: Cambridge University Press.

Lakoff, G. & Johnson, M. (1980) *Metaphors We Live By*. Chicago, IL: University of Chicago Press.

Langacker, R. W. (1982) "Space Grammar, Analysability, and the English Passive", *Language* 58, 22-80.

Langacker, R. W. (1987) *Foundations of Cognitive Grammar, Vol. I: Theoretical Prerequisites*. Stanford, CA: Stanford University Press.

Langacker, R. W. (1991a) *Foundations of Cognitive Grammar, Volume II: Descriptive Application*. Stanford, CA: Stanford University Press.

Langacker, R. W. (1991b) *Concept, Image, and Symbol: The Cognitive Basis of Grammar*. New York: Mouton de Gruyter.

Pinker, S. (1989) *Learnability and Cognition: The Acquisition of Argument Structure*. Cambridge, MA: MIT Press.

Rowlands, M. (1995) "Against Methodological Solipsism: The Ecological Approach", *Philosophical Psychology*, 8, 5-24.

Rumelhart, D. E. (1992) "Toward a Microstructural Account of Human Reasoning", in S. Davis, ed.,

Connectionism: Theory and Practice. Oxford: Oxford University Press.

Rumelhart, D. E., Smolensky, P., McClelland, J. L., & Hinton, G. E. (1986) "Schemata and Sequential Thought Processes in PDP Models", in J. L. McClelland, D. E. Rumelhart and the PDP Research Group, eds., *Parallel Distributed Processing, Vol. 2: Explorations in the Microstructure of Cognition*. Cambridge, MA: MIT Press.

Smolensky, P. (1987) "The Constituent Structure of Connectionist Mental States: A Reply to Fodor and Pylyshyn". *Southern Journal of Philosophy*, 26, Supplement, 137-161.

Smolensky, P. (1988) "On the Proper Treatment of Connectionism", *Behavioral and Brain Sciences*, 11, 1-74,

Smolensky, P. (1991) "Connectionism, Constituency and the Language of Thought", in B. Loewer & G. Rey, eds., *Meaning in Mind: Fodor and His Critics*. Oxford: Blackwell.

Sweetser, E. E. (1990) *From Etymology to Pragmatics: Metaphorical and Cultural Aspects of Semantic Structure*. Cambridge: Cambridge University Press.

Talmy, L (1983) "How Language Structures Space", in H. Pick & L. Acredolo, eds., *Spatial Orientation: Theory, Research, and Application*. New York: Plenum.

Vygotsky, L. S. (1962) *Thought and Language* (E. Hanfmann & C. Vakar, Trans.). Cambridge, MA: MIT Press. (Original work published 1934)