Advanced Review

Rethinking the language of thought

Susan Schneider[1]* and Matthew Katz[2]

In this piece, we overview the language of thought (LOT) program, a currently influential theory of the computational nature of thought. We focus on LOT's stance on concepts, computation in the central system, and mental symbols. We emphasize certain longstanding problems arising for the LOT approach, suggesting resolutions to these problems. Many of the solutions involve departures from the standard LOT program, i.e., the LOT program as developed by Jerry Fodor. We close by identifying avenues for future work.

How to cite this article:
WIREs Cogn Sci 2011. doi: 10.1002/wcs.1155

INTRODUCTION

The computational paradigm in cognitive science aims to provide a complete account of our mental lives, from the mechanisms underlying our memory and attention to the computations of the singular neuron. However, what does the computational paradigm amount to, philosophically speaking? In this piece, we discuss one influential approach to answering that question: the *language of thought* (LOT) approach. LOT is one of two leading positions on the computational nature of thought, the other being connectionism. According to LOT, humans and even non-human animals think in a *lingua mentis*, an inner mental language that is not equivalent to any natural language. This mental language is computational in the sense that thinking is regarded as the algorithmic manipulation of mental symbols. In this piece, we outline the main features of the LOT position. We then consider several problems that have plagued the LOT approach for years: the concern that LOT is outmoded by connectionism; an objection by LOT's own philosophical architect, Jerry Fodor, that urges that it is likely that cognition is non-computational; and the failure of LOT to specify the nature of mental representations that it invokes in its theory (such representations are called 'mental symbols'). We respond to these problems, and in so doing, we sculpt a version of LOT that departs in certain ways from Fodor's influential approach. *Inter alia*, we outline an account of symbols that gives rise to a new theory of the nature of concepts (pragmatic atomism). We close with a discussion of an issue that we believe requires further investigation, that of mathematical cognition.

*Correspondence to: sls@sas.upenn.edu

[1]Department of Philosophy, University of Pennsylvania, Philadelphia, PA, USA

[2]Department of Philosophy and Religion, Central Michigan University, Mt. Pleasant, MI, USA

THE BASIC ELEMENTS OF THE LOT HYPOTHESIS

The idea that there is a LOT was mainly developed by Jerry Fodor, who defended the hypothesis in his influential book, *The Language of Thought* (1975).[1] The view that cognition is a species of symbol processing was in the air around the time that Fodor wrote the book. For instance, Allen Newell and Herbert Simon suggested that psychological states could be understood in terms of an internal architecture that was like that of a digital computer.[2] Human psychological processes were said to consist in a system of discrete inner states (symbols) that are manipulated by a central processing unit (see also Ref 3). The LOT or 'symbol processing' position became the paradigm view in information processing psychology and computer science until the 1980s, when the competing connectionist view gained currency. At that time the symbol processing view came to be known as 'Classical Computationalism' or just 'Classicism'.[4,5]

LOT consists in several core claims. (1) A first claim is that cognitive processes are the causal sequencing of tokenings of symbols in the brain.

(2a) LOT also claims that mental representations have a *combinatorial syntax*. A representational system has a combinatorial syntax just in case it employs a finite store of atomic representations that may be combined to form compound representations, which may in turn be combined to form further compound representations. (2b) Relatedly, LOT holds that symbols have a *compositional semantics*—the meaning of compound representations is a function of the meaning of the atomic symbols, together with the grammar.[a] (2c) LOT further claims that thinking, as a species of symbol manipulation, often preserves semantic properties of the thoughts involved.[1,6] An important example of one such property is *being true*. Consider the mental processing of an instance of *modus tollens*. The internal processing is purely syntactic, but it is nevertheless *truth preserving*. Given true premises, the application of the rule will result in further truths.

(3) Finally, LOT asserts that mental operations on internal representations are causally sensitive to the syntactic structure of the symbols. That is, computational operations work on any symbol or string of symbols satisfying a certain structural description, transforming the symbol/string into another symbol/string that satisfies another structural description. For instance, the system may employ an operation in which it transforms any representation that it recognizes as having the form (Q&R) into a representation with the form (Q). In such an operation, the meaning of the symbols plays no role at all. In addition, the physical structures onto which the symbol structures are mapped are the very properties that cause the system to behave in the way it does.[6]

(1)–(3) combine to form a position called the 'Computational Theory of Mind' (CTM). CTM holds that thinking is a computational process involving the manipulation of semantically interpretable strings of symbols that are processed according to algorithms.[2,4,7–11] Together, LOT and CTM aim to answer the age old question, 'how can rational thought be grounded in the brain'? Their answer is that rational thought is a matter of the causal sequencing of symbol tokens that are realized in the brain. And further, these symbols, which are ultimately just patterns of matter and energy, have both representational and causal properties. And as noted, the semantics mirrors the syntax. So thinking is a process of symbol manipulation in which the symbols have an appropriate syntax and semantics (roughly, natural interpretations in which the symbols systematically map to states in the world[b]).[4,10]

LOT is controversial, to be sure. For one thing, LOT is by no means the only approach to the format of thought. As we shall now see, some regard connectionism as rendering LOT obsolete.

RESPONDING TO THE CONNECTIONIST CHALLENGE

Cognitive science is only in its infancy, but various cognitive scientists, such as Jeffrey Hawkins (computational neuroscience) and Paul and Patricia Churchland (philosophy) claim that they already see the outlines of a final theory, glimmers of a singular format of thought exhibited throughout the entire brain, from the simplest sensory circuit to the most complex computation of the prefrontal cortex. In essence, the mind is a species of connectionist network. For instance, Paul Churchland writes:

> ... we are now in a position to explain how our vivid sensory experiences arise in the sensory cortex of our brains: how the smell of baking bread, the sound of an oboe, the taste of a peach, and the color of a sunrise are all embodied in a vast chorus of neural activity... And we can see how the matured brain deploys that framework almost instantaneously: to recognize similarities, to grasp analogies, and to anticipate both the immediate and the distant future (Ref 12, p. 3).

Connectionists claim that thinking is determined by patterns of activation in a neural network. Consider the following extremely simple connectionist network (Figure 1).

Each circle (or 'unit') can represent either a single neuron or a group of neurons. There are three layers of

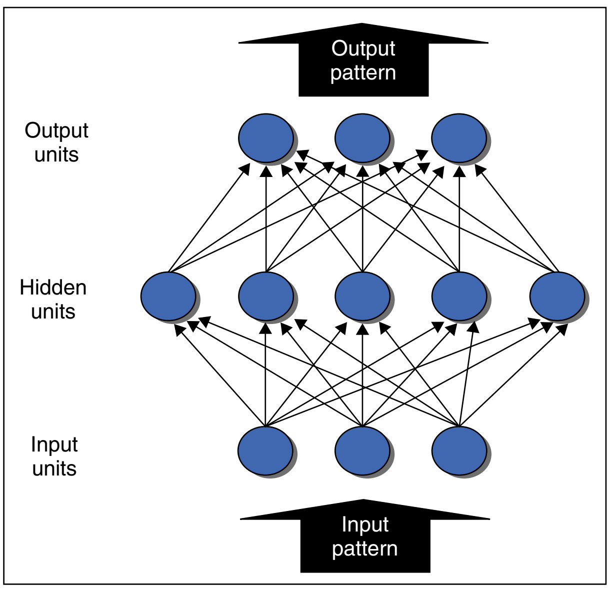

FIGURE 1 | A simple network.

units: an input layer, a middle or 'hidden' layer, and an output layer. Computation flows upward, with the smaller arrows specifying connections between units. Each hidden or output unit carries a numerical activation value that is computed given the values of the neighboring units in the network, according to a function. The input units' signals thereby propagate throughout the network, determining the activation values of all the output units. Of course, actual models of perceptual and cognitive functions are far more complex than the simple illustration given here, exhibiting multiple hidden layers and feedback loops. But this is the bare bones idea.

It is important to note some of the differences between such networks and classical computers, from which LOT was originally inspired. First, while classical computers possess distinct memory and processing units, connectionist networks do not, being comprised entirely of simple nodes and the connections between them. Second, while classical computers are serial machines, processing representations in step-wise fashion, network processing occurs in parallel. Activation values of the nodes in a network are continuously updated until the network 'settles down' into a steady state in which activation values are no longer being updated. Third, while processing in classical computers is concentrated in the central processing unit, it is widely distributed in a connectionist network. Indeed, semantic interpretation is distributed across nodes, and compound representations are not formed by concatenation, as in classical machines. This makes it difficult, if not impossible, to explain features of thought such as productivity and systematicity within a connectionist framework (more on this below).

The connectionist framework has had many success stories, especially in domains involving the recognition of perceptual patterns, and indeed, we find connectionist work on pattern recognition promising. However, when it comes to considering the implications of connectionism for the format of higher thought, even connectionists are wont to admit that existing models are only highly simplified depictions of how a given perceptual or cognitive capacity works. As new information emerges about the workings of the brain, connectionists will of course add more sophistication to their models. And thus connectionists suspect that higher thought will be describable given the resources of a computational neuroscience that is based on, and represents a more sophisticated version of, connectionist theorizing. Put bluntly, it will be 'more of the same' but with added bells and whistles.

But will it really be more of the same? Many advocates of LOT deny this possibility. Their main objection arises from the very reason they believe the brain computes in LOT: thought has the crucial and pervasive feature of being combinatorial. First, consider the thought *the beer in Munich is better than on Mars*. You probably have never had this thought before, but you were able to understand it. The key is that the thoughts are built out of familiar constituents, and combined according to rules. It is the *combinatorial* nature of thought that allows us to understand and produce these sentences on the basis of our antecedent knowledge of the grammar and atomic constituents (e.g., *beer, Munich*). More specifically, thought is *productive*: in principle, one can entertain and produce an infinite number of distinct representations because the mind has a combinatorial syntax. Relatedly, thought is *systematic*. A representational system is systematic when the ability of the system to entertain and produce certain representations is intrinsically related to its ability to entertain and produce other representations.[6] For example, one does not find normal adult speakers who entertain and produce 'David likes Mitsuko' without also being able to entertain and produce 'Mitsuko likes David'. As with productivity, systematicity arises from the fact that thoughts display combinatorial structure.[1,4,6,13]

It is commonly agreed that LOT excels in explaining these language-like features of thought. From the connectionist vantage point, however, thought is not *primarily* language-like; instead, cognition is a species of pattern recognition consisting in associative relationships between units. The language-like and combinatorial character of cognition has therefore been a thorn in the side of the connectionist. Whether any connectionist models, and especially, any models that are *genuinely non-symbolic* (more on this shortly), can explain these important features of thought are currently a source of intense controversy.[6,13–18] Relatedly, it has been argued that the symbol processing view best explains how minds distinguish between representations of individuals and of kinds, encode rules or abstract relationships between variables, and have a system of recursively structured representations.[16] As with systematicity and productivity, which are closely related to the language-like structure of higher thought, these features of thought are more naturally accommodated by the symbolic paradigm.

Not only is it controversial whether purely connectionist models can fully explain higher thought, but it is also unclear how very simple models of isolated neural circuits are supposed to 'come together' to give rise to a larger picture of how the mind works.[19] Existing 'big picture' accounts are fascinating, yet patchy and speculative. Furthermore, higher cognition is the domain in which we would see validation of the symbol processing approach, if validation is to come at all.[4] The discrete, combinatorial representations of the prefrontal cortex are distinct from the more distributed modality-specific representation of the posterior cortex; *prima facie*, this latter representation seems more straightforwardly amenable to traditional connectionist explanation.[c]

Furthermore, the relationship between LOT and connectionism is subtle. The advocate of LOT has an important rejoinder to the connectionist attempt to do without mental symbols: to the extent that the connectionist can explain the language-like, combinatorial nature of thought, the connectionist model would, at best, merely be a model in which symbols are implemented in the mind. So connectionism is not really a genuine alternative to the LOT picture, for the networks would just be the lower-level implementations of symbolic processes. This position is called 'implementational connectionism'.[6,16] As Gary Marcus explains:

> The term *connectionism* turns out to be ambiguous. Most people associate the term with the researchers who have most directly challenged the symbol-manipulation hypothesis, but the field of connectionism also encompasses models that have sought to explain how symbol-manipulation can be implemented in a neural substrate... The problem is that discussions of the relation between connectionism and symbol manipulation often assume that evidence for connectionism automatically counts as evidence against symbol-manipulation (Ref 16, p. 2).

If connectionism and symbolicism represent genuine alternatives, a view called 'Radical Connectionism' must be correct. But existing connectionist models of higher cognitive function are few, and there are persuasive arguments that putative radical connectionist models in fact make *covert* use of symbolic representations.[16]

Furthermore, the mind may be a sort of 'hybrid' system, consisting in both neural circuits that satisfy symbolic operations and other neural circuits that compute according to connectionist principles, and do not satisfy symbolic operations at all. (These other neural circuits may lack representations that combine in language-like ways, and in this case, the proponent of LOT would not even claim that they are implementations of symbolic systems.) Indeed, certain connectionists currently adopt hybrid positions in which symbol processing plays an important role.[20]Many of these models employ connectionist networks to model sensory processes and then rely on symbol processing models for the case of cognition, but perhaps even the cognitive mind will be a mix of different formats, a point we consider in the final section of this piece.

The upshot is that connectionist success stories in the domain of pattern recognition do not suggest that LOT is false. LOT is compatible with both the success and failure of connectionism. Instead, the proponent of LOT should look to connectionism with great interest, as it seeks to uncover the neurocomputational basis of thought. Pursuing a better understanding of LOT is thus crucial to contemporary cognitive science.

INTEGRATING THE PHILOSOPHICAL LOT WITH COGNITIVE SCIENCE: BEYOND FODORIAN PESSIMISM

Yet there is an additional reason why connectionists and other critics of LOT reject LOT: the mind is not analogous to a digital computer. It does not have a CPU through which every mental operation is shuttled sequentially, and more generally, the brain has a complex and unique functional organization that is not like the functional organization of a standard computer. We agree with both of these observations, but we believe that they do not speak against LOT. While the idea that the mind has a CPU was held by early Classicists, most contemporary advocates of the symbol processing approach regard the mind as computational because they believe cognition is the algorithmic computation of discrete symbols, where the algorithm that the brain runs is to be discovered by a completed cognitive science.[4,7,10] These algorithms compute on symbols in a manner that is sensitive to the constituent structure of symbolic strings, enabling the compositionality, productivity, and systematicity of thought. This position, rather than the outdated claim that the brain has the functional organization of a standard computer, best captures the sense in which the brain is said to compute in LOT.

But LOT's chief philosophical architect, Jerry Fodor, has argued the cognitive mind is likely non-computational, for the brain's 'central system' will likely defy computational explanation.[8,9,21] Fodor calls the system responsible for our ability to integrate material across sensory divides and generate complex, creative thoughts 'the central system'. The central system is 'informationally unencapsulated'—its operations can draw from information from outside of the system, in addition to its inputs. And it is *domain general*, with inputs ranging over diverse subjects. Fodor urges that the central system's unencapsulation gives rise to insuperable obstacles. One longstanding worry is that the computations in the central system are not feasibly computed within real time. For if the mind truly is computational in a classical sense, when one makes a decision one would never be able to determine what is relevant to what. For the central system would need to walk through every belief in its database, asking if each item was relevant. Fodor concludes from this that the central system is likely non-computational. Shockingly, he recommends that cognitive science stop working on cognition.

Fodor's pessimism has been influential within philosophical circles, situating LOT in a tenuous dialectical position: since LOT itself was originally conceived of as a computational theory of deliberative, language-like thought (herein, 'conceptual thought'),[1,3] and the central system is supposed to be the domain in which conceptual thought occurs, it is unclear how, assuming that Fodor is correct, LOT could even be true. In any case, advocates of LOT in cognitive science proper reject Fodor's pessimism about computation in the central systems.[10,16] And some recent philosophical discussions of LOT have responded to Fodor's anti-computationalism. For instance, Peter Carruthers responds to Fodor's relevance challenge by granting that if the central system is amodular, as Fodor contends, the problem would be insurmountable. According to Carruthers, the central system must be *massively modular*, that is, the central system features special purpose, innate information processing modules.[7,22,23] Carruthers summons research in areas such as evolutionary psychology, animal cognition, and dual systems theory to develop such a view.[7]

Schneider urges that an amodular central system can compute what is relevant. She draws from the work of Murray Shanahan and Bernard Baars, who sketch the beginnings of a solution to the Relevance Problem that is based upon the Global Workspace (GW) theory.[24–27] According to the GW theory, a pancortical system (a 'GW') facilitates information exchange among multiple parallel, specialized unconscious processes in the brain. When information is conscious there is a state of global activation in which information in the workspace is 'broadcast' back to the rest of the system. At any given moment, there are multiple parallel processes going on in the brain that receive the broadcast. Access to the GW is granted by an attentional mechanism and the material in the workspace is then under the 'spotlight' of attention. When in the GW the material is processed in a serial manner, but this is the result of the contributions of parallel processes that compete for access to the workspace. (This view is intuitively appealing, as our conscious, deliberative, thoughts introspectively appear to be serial.) Schneider summons the GW theory as the basis for a computational account of the central system.

It is also noteworthy that Schneider's development of the central system requires that LOT turn its attention to cognitive and computational neuroscience, despite Fodor's well-known repudiation of these fields. First, she urges that integration with neuroscience will enable a richer explanation of the structure of the central system. In her book, *The Language of Thought: A New Philosophical Direction*, she illustrates how neuroscience can provide a deeper understanding of the central system by appealing to certain recent work on the central system in neuroscience and psychology, such as, especially, the aforementioned GW theory.[4] Second, on her view, cognitive and computational neuroscience are key to determining what kinds of mental symbols a given individual has (e.g., whether one has a symbol for bulldogs, jazz, or Chianti), for these fields detail the algorithms that the brain computes, and on her view, the type that a given mental symbol falls under is determined by the role it plays in the algorithms that the brain computes. Third, she stresses that given that LOT is a *naturalistic* theory, that is, one that seeks to explain mental phenomena within the domain of science, it depends on integration with neuroscience for its own success.

Now let us turn to yet a further problem that LOT faces: LOT and CTM purport to be theories of the symbolic nature of thought, but as we shall see, their notion of a mental symbol is poorly understood.

TOWARD AN UNDERSTANDING OF SYMBOLIC MENTAL STATES

LOT's commitment to the idea that thinking involves the manipulation of mental symbols is often misunderstood. What are mental symbols? Clearly, there is no orthography or 'brain writing' in which strings of symbols are encoded. Rather, LOT claims that at a high level of abstraction, the brain can be accurately described as employing a representational system that has a language-like structure. A second source of misunderstanding is that LOT is frequently wedded to a form of radical nativism in which all symbols are said to be innate. While certain symbols may be innate, it is hard to believe that cavemen had innate symbols such as [photon] and [internet]. We suspect that this implausible view has strengthened the sense of mystery surrounding LOT's notion of a mental state, and it has led many to reject LOT as implausible. However, LOT does not require radical nativism: indeed, even Fodor no longer endorses radical nativism,[28] and other proponents of LOT have for years taken a more moderate approach to concept nativism. A third source of misunderstanding is more serious, however. Proponents of LOT have not specified whether LOT's symbols are typed with respect to their semantic properties, their neural properties, their computational roles, or with respect to some other feature entirely.[4,29–32]

Without an understanding of how symbols are typed it is difficult to see how neural activity can, at a high level of abstraction, even be described as the manipulation of mental symbols. It is also difficult to determine whether a given connectionist model truly implements symbolic manipulations. Furthermore, the proponent of LOT summons symbols to do important philosophical work, and this work cannot be done without a plausible conception of a symbol. Symbols are supposed to be neo-Fregean 'modes of presentation': roughly, one's way of conceiving things. As such, LOT looks to symbols to explain how an individual can have different ways of conceiving of the same object, and to explain the causation of thought[d] and behavior.[4,33,34]

Without a clear notion of a LOT symbol, none of this work can be accomplished. However, Schneider has recently developed a theory of symbols that is designed to play the aforementioned roles for LOT. She begins by ruling out competing theories of the nature of symbols. For instance, she rules out approaches that identify symbols by what they refer to because they are not finely grained enough to differentiate between different ways one can conceive of the same object (e.g., consider thinking of Clark Kent as being Superman, versus merely being a reporter for the *Daily Planet*). She then argues that LOT requires a theory of symbols that assigns a symbol token to its type in virtue of the role the symbol plays in the algorithms employed by the central system (she calls this 'the algorithmic conception'). This position is not new: in the past both Jerry Fodor and Stephen Stich have appealed to it.[34,35] However, neither philosopher developed the idea in any detail, and Fodor came to repudiate it. Schneider provides three arguments for the algorithmic conception. The first is that the classical view of computation, to which LOT is wedded, requires that primitive symbols be typed in this manner. The second is that without this manner of individuating symbols, cognitive processing could not be symbolic. The third is that cognitive science needs a natural kind that is defined by its total computational role. Otherwise, either explanation in cognitive science will be incomplete, or its generalizations will have counterexamples. Finally, she explores how symbols, thus understood, figure in explanations of thought and behavior in cognitive science.[4,31,32] Her theory of symbols gives rise to a theory of the nature of concepts, to which we shall now turn.

A 'PRAGMATIST' THEORY OF CONCEPTS

One outgrowth of the philosophical LOT program is a theory of the nature of concepts, *conceptual atomism,* a view of the nature of concepts devised by Jerry Fodor, Eric Margolis, Stephen Laurence, and others.[4,28,36] Conceptual atomism holds that lexical concepts lack semantic structure, being in this sense 'atomic'. It further holds that a concept has a twofold nature, consisting in both its broad content (roughly, what it refers to in the world) and its symbol type.[e] Much of Fodor's work on concepts situates conceptual atomism, and indeed, LOT itself, in opposition to positions that he labels 'pragmatist'—positions on the nature of concepts that hold one's psychological abilities (e.g., classificatory or inferential capacities) determine the nature of concepts (Ref 37, p. 34). Fodor urges that pragmatism is a '…catastrophe of analytic philosophy of language and philosophy of mind in the last half of the twentieth century'.[38]

Schneider claims that both LOT and conceptual atomism must embrace pragmatism. For as noted, she argues that the nature of symbols is determined by the role they play in computations in the central system. So a symbol type is determined by role a symbol plays in an individual's mental life, including the role it plays in inference and classification. Furthermore, even on Fodor's view, symbols are elements of a concept's nature.[4,39] Clearly, Fodor did not anticipate that conceptual atomism is pragmatist; this is likely because symbol natures were left underspecified. To distinguish her position from Fodor's she calls this brand of conceptual atomism *Pragmatic Atomism*.

Schneider contends that pragmatic atomism is a superior version of conceptual atomism. Critics have long charged that conceptual atomism is too skeletal, merely being a semantic theory in which concepts are a matter of the information they convey, and saying next to nothing about the role concepts play in our mental lives. Indeed, the standard conceptual atomist cannot even distinguish between corefering concepts that have the same grammatical form (e.g., *Cicero/Tully*). This defect stands in stark contrast to psychological theories of concepts, such as the prototype theory and the theory theory, which focus on how we reason, learn, categorize, and so on. However, pragmatic atomism can say that the features of concepts that psychologists are traditionally interested in are built into concepts' very natures. Symbols, being individuative of concepts, capture the role the concept plays in thought. For example, symbols have computational roles that distinguish corefering concepts, characterize the role the concept plays in categorization, determining whether, and how rapidly, a person can verbally identify a visually presented object, confirm a categorization judgment, identify features an object possesses if it is indeed a member of a given category, and so on. Here, pragmatic atomism can draw from research from any of the leading psychological theories of concepts. Consider the prototype view, for instance. In the eyes of pragmatic atomism, the experimental results in the literature on prototypes are indications of features of certain symbols' underlying computational roles, and these roles determine the relevant concept's natures.[4,39]

Our discussion has isolated certain longstanding problems arising LOT; and we have urged that in dealing with these problems one must rethink key elements of the standard LOT program, namely, LOT's stance on concepts, symbols, and computation. Now, we shall conclude by asking: where should the LOT program go from here? We believe that the following topics deserve future attention: (1), understanding the relationship between symbols, typed in the manner outlined herein, and accounts of higher cognitive function in cognitive and computational neuroscience. This will facilitate a better understanding of how symbols and concepts are grounded in the brain, and determine whether the brain indeed computes symbolically. (2) Understanding the relation between linguistically formatted and non-linguistically formatted representations in light of an emerging challenge presented by the literature on numerical representation. We close by commenting on this latter issue.

CONCLUDING REMARKS: FUTURE INVESTIGATION ON NON-LINGUISTIC MATHEMATICAL COGNITION

An interesting development has emerged from research on the representation and processing of numerical information. A wealth of experimentation[f] suggests that humans as well as certain non-human animals possess an innate cognitive system that represents cardinality in terms of *mental magnitudes*. This system represents the cardinality of groups of objects as quantities that are proportional in size to the cardinalities they represent. For example, if a subject counts a set of objects, the system tokens one increment for each object counted. The more objects in the group, the more increments are tokened, and the larger the resulting representation. However, the increments themselves are variable in size, and as the increments are combined to form larger and larger representations, the variability compounds. At a certain point, the system is unable to distinguish a given number from its nearby neighbors, where what counts as 'nearby' is proportional to the number represented. This indicates that the system cannot 'recall' the number of increments that were used to form a compound representation. Rather, it only has access to the final product. In other words, the system forms compound representations by combining increments and not by concatenating them. So the components of compound representations are not discrete. This is a fundamental difference between mental magnitudes and linguistic representations.[g]

The existence of mental magnitudes may indeed provide evidence that LOT is not true of all of cognition. But if so, the issue is limited to those domains that employ magnitude representations, and even within those domains the research does not necessarily challenge symbolic representation. For instance, the presence of magnitude representations of cardinality does not call into question the view that mathematical operations are also symbolic, for most authors hold that higher thought about numbers is mediated by linguistically formatted representations. For instance, Dehaene has argued for a 'triple code' theory, according to which numerical information is processed by three separate cognitive systems.[40] One of these systems employs mental magnitudes, one employs a visual number form (such as the Arabic numerals), and one employs auditory verbal representations. These latter two kinds of representations are discrete, and indeed, can be individuated by their computational roles. Thus, research on mental magnitudes (and other non-linguistic forms of mental representation) can be viewed as being complimentary to the LOT program, such that these various avenues of investigation together seek to answer questions regarding the scope of and relationship between linguistic and non-linguistic representational systems.

Indeed, because mental magnitudes are present in human infants and non-human animals, most researchers have been concerned with whether and how mental magnitudes play a role in the formation of higher, conceptual, thought about numbers. For example, Spelke has argued that mental magnitudes must combine with another innate representational system—the system of object files—in order to achieve the precision that is characteristic of mature conceptual thought about numbers.[41] Carey has argued that natural number concepts first arise only as small number concepts, provided by object-file representations that are 'enriched' by quantificational markers in natural language. She argues that mental magnitudes only later play a role, as they become linked with the enriched object-file representations to thus provide large number concepts.[42,43] In contrast, Leslie, Gelman, and Gallistel and Laurence and Margolis have argued that even combining mental magnitudes with other systems is insufficient to supply mature numerical concepts, and that humans must in fact possess innate representations of the first few natural numbers.[44,45] Whichever, if any, of these accounts turns out to be accurate, the question being addressed can be thought of in terms of whether and how symbolic language-like mental representations develop from non-symbolic non-language-like mental representations.[h]

NOTES

[a]A proponent of LOT need not claim that LOT has a semantics, for one could reject meanings. But in practice, an appeal to meaning has generally been a key facet of LOT. For an important exception, see Stephen Stich (1983).[35] Given that this section follows the standard elaborations of LOT, which appeal to a semantics, LOT is basically indistinguishable from a view called 'CTM' or the 'Classical Computational Theory of Mind', a position we discuss later in this section.

[b]For another overview of LOT, see Ref 46.

[c]For a discussion of the distinct processing in the prefrontal cortex and posterior cortex, see O'Reilly and Munakata (Ref 47, p. 214–219).

[d]Fregeans consider mode of presentations (MOPs) to be semantic. Many advocates of LOT appeal to a referential semantics and do not consider MOPs to be semantic in nature.

[e]Because symbol natures have been neglected, only the semantic dimension of the theory has been developed. Indeed, the conceptual atomists' concepts are often mistakenly taken as being *equivalent* to broad contents, despite the fact that they are individuated by their symbol types as well.[4,39]

[f]For useful reviews, see for example, Dehaene (1997), Feigenson, Dehaene, and Spelke (2004), Gallistel, Gelman, and Cordes (2006), and Cantlon, Pratt, and Brannon (2009).[48–51]

[g]In the above description we have omitted several areas of disagreement about mental magnitudes. First, there is debate about whether the system of mental magnitudes functions serially or in parallel (see Refs 52 and 53). Second, there is debate about whether variability is present in the magnitudes themselves or in memory (see Ref 50). Third, there is debate about whether the system exhibits scalar variability or logarithmic compression (see Refs 54 and 55). Despite these disagreements, there is wide consensus that positing magnitude representations explains the fact that human and non-human discrimination of sets based on cardinality obeys Weber's Law, which states that $\Delta I/I = k$, where I is the intensity of a stimulus, ΔI is the minimal change in intensity required for discrimination of the stimulus from another, and k is a constant. In other words, whether a subject is able to discriminate between two stimuli depends on the difference in intensity of the two, not their absolute values. Positing linguistically formatted representations would not explain why discrimination of sets based on cardinality obeys Weber's Law, and hence the argument that there is a fundamental difference between magnitude representations and representations in LOT (see also Ref 56).

[h]For an overview of recent work documenting magnitude representations and more on their relationship to LOT (see Ref 57).

ACKNOWLEDGMENTS

We are grateful to Carlos Montemayor, two anonymous reviewers, and the editors for their helpful comments on this piece.

REFERENCES

1. Fodor JA. *The Language of Thought*. Cambridge: Harvard University Press; 1975.
2. Newell A, Simon HA. *Human Problem Solving*. Englewood Cliffs, NJ: Prentice-Hall; 1972.
3. Harman G. *Thought*. Princeton, NJ: Princeton University Press; 1973.
4. Schneider S. *The Language of Thought: A New Philosophical Direction*. Boston, MA: MIT Press; 2011.
5. Boden M. *Mind as Machine: A History of Cognitive Science*, vol 2. New York: Oxford University Press; 2008.

6. Fodor JA, Pylyshyn ZW. Connectionism and cognitive architecture: a critical analysis. *Cognition* 1988, 28:3–71.
7. Carruthers P. *The Architecture of the Mind: Massive Modularity and the Flexibility of Thought*. Oxford: Oxford University Press; 2006.
8. Fodor JA. *The Mind Doesn't Work That Way: The Scope and Limits of Computational Psychology*. London: MIT Press; 2000.
9. Fodor JA. *LOT 2: The Language of Thought Revisited*. Oxford: Oxford University Press; 2008.
10. Pinker S. *How the Mind Works*. New York: W.W. Norton; 1999.
11. Rey G. *Contemporary Philosophy of Mind*. London: Blackwell; 1997.
12. Churchland P. *Engine of Reason, Seat of the Soul*. Boston, MA: MIT Press; 1996.
13. Fodor JA, McLaughlin B. Connectionism and the problem of systematicity: why Smolensky's solution doesn't work. *Cognition* 1990, 35:183–204.
14. Elman J. Generalization, simple recurrent networks, and the emergence of structure. In: Gernsbacher MA, Derry S, eds. *Proceedings of the 20th Annual Conference of the Cognitive Science Society*. Mahwah, NJ: Lawrence Erlbaum; 1998.
15. van Gelder T. Why distributed representation is inherently non-symbolic. In: Dorffer G, ed. *Konnektionismus in Artificial Intelligence und Kognitionsforschung*. Berlin: Springer-Verlag; 1990, 58–66.
16. Marcus G. *The Algebraic Mind*. Boston, MA: MIT Press; 2001.
17. Smolensky P. On the proper treatment of connectionism. *Behav Brain Sci* 1988, 11:1–23.
18. Smolensky P. Reply: constituent structure and explanation in an integrated connectionist/symbolic cognitive architecture. In: Macdonald C, Macdonald G, eds. *Connectionism: Debates on Psychological Explanation*, vol 2. Oxford: Basil Blackwell; 1995, 221–290.
19. Anderson J. *How Can the Human Mind Occur in the Physical Universe?* Oxford: Oxford University Press; 2007.
20. Wermter S, Sun R, eds. *Hybrid Neural Systems*. Heidelberg: Springer; 2000.
21. Ludwig K, Schneider S. Fodor's critique of the classical computational theory of mind. *Mind Lang* 2008, 23:123–143.
22. Samuels R. Evolutionary psychology and the massive modularity hypothesis. *Br J Philos Sci* 1998, 49:575–602.
23. Sperber D. In defense of massive modularity. In: Dupoux I, ed. *Language, Brain and Cognitive Development*. Cambridge: MIT Press; 2002, 47–57.
24. Baars BJ. *In the Theater of Consciousness*. New York: Oxford University Press; 1997.
25. Dehaene S, Changeux JP. Neural mechanisms for access to consciousness. In: Gazzaniga MS, ed. *Cognitive Neurosciences*. 3rd ed. Cambridge: MIT Press; 2004, 1145–1158.
26. Dehaene S, Naccache L. Towards a cognitive neuroscience of consciousness: basic evidence and a workspace framework. *Cognition* 2001, 79:1–37
27. Shanahan M, Baars B. Applying global workspace theory to the frame problem. *Cognition* 2005, 98:157–176.
28. Fodor JA. *Concepts: Where Cognitive Science Went Wrong*. Oxford: Oxford University Press; 1998.
29. Aydede M. On the type/token relation of mental representations. *Facta Philos Int J Contemp Philos* 2000, 2:23–49.
30. Prinz J. *Furnishing the Mind: Concepts and Their Perceptual Basis*. Cambridge: MIT Press; 2004.
31. Schneider S. LOT, CTM and the elephant in the room. *Synthese* 2009, 170:235–250.
32. Schneider S. The nature of symbols in the language of thought. *Mind Lang* 2009, 24:523–553.
33. Frege G. Über Sinn und Bedeutung. *Zeitschrift Philosophie Philosophische Kritik* 1892, 100:25–50. Trans. Black M, as Sense and reference. *Philos Rev* 1948, 57:209–230.
34. Fodor JA. *The Elm and the Expert: Mentalese and Its Semantics*. Cambridge: MIT Press; 1994.
35. Stich S. *From Folk Psychology to Cognitive Science: The Case Against Belief*. Boston, MA: MIT Press; 1983.
36. Laurence S, Margolis E. Radical concept nativism. *Cognition* 2002, 86:25–55.
37. Fodor JA. Having concepts: a brief refutation of the twentieth century. *Mind Lang* 2004, 19:29–47.
38. Fodor JA. *Hume Variations*. Oxford: Oxford University Press; 2003.
39. Schneider S. Conceptual atomism rethought. *Behav Brain Sci* 2010, 33:224–225.
40. Dehaene S. Varieties of numerical cognition. *Cognition* 1992, 44:1–42.
41. Spelke ES. What makes us smart? Core knowledge and natural language. In: Gentner D, Goldin-Meadow S, eds. *Language in Mind: Advances in the Investigation of Language and Thought*. Cambridge: MIT Press; 2003, 277–311.
42. Carey S. Bootstrapping and the origin of concepts. *Daedalus* 2004, 133:59–64.
43. Carey S. Where our number concepts come from. *J Philos* 2009, 106:220–254.
44. Leslie AM, Gelman R, Gallistel CR. The generative basis of natural number concepts. *Trends Cogn Sci* 2008, 12:213–218.
45. Laurence S, Margolis E. Linguistic determinism and the innate basis of number. In: Carruthers P, Laurence S, Stich S, eds. *The Innate Mind: Foundations*

and the Future. Oxford: Oxford University Press; 2007, 139–169.

46. Katz M. Language of thought hypothesis. In: Feiser J, Dowden B, eds. *Internet Encyclopedia of Philosophy*. 2009.
47. O'Reilly R, Munakata Y. *Computational Explorations in Computational Neuroscience*. Boston, MA: MIT Press; 2000.
48. Dehaene S. *The Number Sense*. Oxford: Oxford University Press; 1997.
49. Feigenson L, Dehaene S, Spelke ES. Core systems of number. *Trends Cogn Sci* 2004, 8:307–314.
50. Gallistel CR, Gelman R, Cordes S. The cultural and evolutionary history of the real numbers. In: Levinson S, Jaison P, eds. *Evolution and Culture*. Cambridge: MIT Press; 2006, 247–274.
51. Cantlon JF, Platt ML, Brannon EM. Beyond the number domain. *Trends Cogn Sci* 2009, 13:83–91.
52. Meck WH, Church RM. A mode control mechanism of counting and timing processes. *J Exp Psychol Anim Behav Process* 1983, 9:320–334.
53. Dehaene S, Changeux J. Development of elementary numerical abilities: a neuronal model. *J Cogn Neurosci* 1993, 5:390–407.
54. Dehaene S, Izard V, Spelke ES, Pica P. Log or linear? Distinct intuitions of the number scale in western and Amazonian indigene cultures. *Science* 2008, 320:1217–1220.
55. Cantlon JF, Cordes S, Libertus ME, Brannon EM. Comment on 'log or linear? Distinct intuitions of the number scale in western and Amazonian indigene cultures'. *Science* 2009, 323:38b.
56. Montemayor C, Balci F. Compositionality in language and arithmetic. *J Theor Philos Psychol* 2007, 27:53–72.
57. Katz M. Numerical Competence and the Format of Mental Representation. PhD Dissertation, The University of Pennsylvania, 2007.