1 Systematicity: An Overview

John Symons and Paco Calvo

Introduction

In 1988, Jerry Fodor and Zenon Pylyshyn published "Connectionism and Cognitive Architecture: A Critical Analysis." Their article presented a forceful and highly influential criticism of the explanatory relevance of neural network models of cognition. At the time, connectionism was reemerging as a popular and exciting new field of research, but according to Fodor and Pylyshyn, the approach rested on a flawed model of the human mind. Connectionism is the view that the mind can be understood in terms of an interconnected network of simple mechanisms. Its proponents contend that cognitive and behavioral properties can be modeled and explained in terms of their emergence from the collective behavior of simple interacting and adaptive mechanisms. According to Fodor and Pylyshyn, connectionist approaches neglect an essential feature of thought—its systematic nature. On their view, the basic psychological fact that thoughts are intrinsically related to other thoughts in systematic ways becomes inexplicable if one denies that representations are structured in a syntactically and semantically classical combinatorial manner.[1] Connectionism, they argued, inevitably fails to provide a meaningful explanation of cognition insofar as it confuses the intrinsically systematic nature of thought with a system of associations.[2] Connectionism might shed some light on the way that cognitive architectures happen to be implemented in brains, but the explanation of cognition does not take place at the level of biology or hardware. A cognitive architecture must be systematic to the core in order to shed light on the intrinsically systematic character of cognition.

One prominent message of their article, that a cognitive architecture must explain systematicity in order to explain human cognition, came to be called *the systematicity challenge*. The meaning and implications of the challenge can be interpreted in more than one way, and it quickly generated a vigorous response in philosophy and cognitive science. The ensuing debate resulted in an enormous literature from a variety of perspectives.[3] Philosophers and scientists on all sides of the issue generally agree that the paper helped to sharpen central questions concerning the nature of explanation in cognitive science and that Fodor and Pylyshyn have encouraged the scientific community to think carefully about what the goals of cognitive science should be. Nevertheless, in cognitive science and philosophy, opinion is sharply divided concerning the role of systematicity in a mature science of mind.

The criticism of connectionism in the 1988 paper harks back to Fodor's earlier arguments for the idea that cognition should be understood by analogy with classical computational architectures as a system of rules and representations. In the 1970s, Fodor had argued that the language of thought hypothesis explained the systematic features of thought. On his view, all thoughts that have propositional content are representational in nature, and these representations have syntactical and semantic features that are organized in a way that is similar to the transformational rules of natural languages. Insofar as tokens of mental representation figure in thoughts and insofar as we can judge those thoughts to be true or false, they must be organized in a language-like way (Fodor 1975). Fodor presents the language of thought hypothesis as the best way to account for a range of features of our psychology. The three basic explananda that Fodor highlights in his work are:

(a) The productivity of thought: we have an ability to think and understand new thoughts and previously unheard sentences.
(b) The systematicity of thought: to genuinely understand a thought is to understand other related thoughts.
(c) The principle of compositionality: the meaning of sentences results from the meanings of their lexical parts.

Fodor and Pylyshyn's criticism of connectionism is shaped by Fodor's early articulation of the language of thought hypothesis and by their view that competing explanatory strategies miss what is distinctively cognitive about cognition. Their 1988 article applied a challenging philosophical argument to a lively and ongoing scientific controversy. By any measure, their paper has served as a focal point for one of the most active debates in the philosophy of cognitive science over the past twenty-five years.

In our view, the scientific landscape has changed in ways that call for a fresh look at this influential set of arguments. Most obviously, the quarter-century since Fodor and Pylyshyn's paper has seen the development of new approaches to connectionism that depart in a number of important respects from the modeling strategies that were the direct target of their criticism in the 1980s. More generally, the broader scientific context has changed in ways that are relevant to the kinds of explanations that are available to cognitive science. Fodor and Pylyshyn presented an argument about what should count as the right kind of explanation of cognition. In the intervening years, a range of scientific and philosophical developments have supported alternative approaches to explanation in the study of cognition. Dynamical, embodied, situated, ecological, and other methodologies are no longer exotic or marginal options for cognitive scientists. At the other end of the spectrum, the majority of researchers in neuroscience adopt highly reductionist approaches to the brain, focusing increasingly, and very fruitfully, on the cellular and subcellular details.[4]

Systematicity tends to be conflated with classical computational kinds of cognitive architecture and traditional research in artificial intelligence. However, the contemporary landscape with respect to artificial intelligence has also shifted in important ways. Researchers can choose from a range of classical, hybrid, and nonclassical architectures, along with a growing set of noncognitive architectures inspired by developments in robotics. Much has changed since Fodor and Pylyshyn published their article, and this volume is intended as a way of taking stock of one of the most important debates in the history of cognitive science from a contemporary perspective. The question of what counts as a good explanation of cognition clearly has not been settled decisively.

Setting the Stage for the Systematicity Challenge

It is helpful to locate the systematicity challenge in the context of the development of cognitive science in the second half of the twentieth century. To begin with, it is important to note that Fodor and Pylyshyn were not the first to challenge the explanatory status of nonclassical approaches, including network models, for the study of cognition. They were well aware of the historical context and pointed back to what they saw as the decisive defeat suffered by advocates of perceptrons, behaviorism, and the like. They claimed that the discussion of what the correct architecture of cognition looks like "is a matter that was substantially put to rest about thirty years ago; and the arguments that then appeared to militate decisively in favor of the Classical view appear to us to do so still"

(Fodor and Pylyshyn 1988, 3). Indeed, many of the central conceptual questions were already coming into focus in the early 1960s.

Regarding the big-picture philosophical concerns, the relative status of associationism and rationalism, for example, were in play in throughout the early history of cognitive science. These concerns came into focus, for example, in debates over the explanatory status of perceptron models. The inability of single-layered perceptrons to learn how to approximate non-linearly separable functions was one of the reasons that popular scientific opinion had turned against neural networks two decades earlier. The combination of Chomsky's (1959) review of Skinner's *Verbal Behavior* and Minsky and Papert's (1969) critique of Rosenblatt's (1958) perceptrons encouraged Fodor and Pylyshyn to see the fight against network models and behaviorism as having been settled decisively in favor of classicism and rationalism.

The neural network approach was delayed but not derailed. In addition to the development of well-known parallel distributed processing strategies that culminated in the two volumes of Rumelhart and McClelland's (1986) *Parallel Distributed Processing: Explorations in the Microstructure of Cognition*, developments in the 1970s and early 1980s include Stephen Grossberg's (1982) adaptive resonance theory and self-organizing maps and Kunihiko Fukushima's (1980) neocognitron, among others. Although neural network approaches may not have been as prominent during this period, Grossberg and others were developing increasingly sophisticated formal techniques that allowed researchers to sidestep many of the objections from the 1960s.

Several well-known milestones marked the revival of popular interest in connectionism during the late 1970s and early '80s.[5] The most prominently cited is the development of the backpropagation learning algorithm for multilayered perceptrons (Rumelhart, Hinton, and Williams 1986) that permitted researchers to address problems that had previously been regarded as intractably difficult challenges for network modelers. The debate over the acquisition of the past tense in English played a historically important role in this regard. Heated discussion as to how much of the developmental psycholinguistic data could be accounted for by statistical means alone continued throughout the 1980s and '90s. The phenomenon to be explained takes the following form. As they learn the past tense of English, children develop their abilities in a familiar "U-shaped developmental profile" (Berko 1958; Ervin, 1964). Initially, they correctly produce the past tense forms of regular and irregular verbs. However, they soon go through a period where they make predictable errors in the inflection of verbs, for example, "goed" and "falled" instead of "went" and "fell." Errors of this kind, at this stage in the developmental process, are widely understood as being the result of the overgeneralization of rules. After this decline in their performance, children begin to correct these overgeneralization errors and their performance improves.

According to connectionists, explanations of the characteristic U-shaped pattern of acquisition can be provided by network models that mimic the pattern of acquisition with a reasonable level of accuracy. (See, e.g., Plunkett and Juola 1999.) Against the connectionists, Pinker and Prince (1988) maintained that researchers would need to resort to more abstract symbol-and-rule mechanisms if they were ever to model the phenomenon adequately. At this point, the trade-offs were becoming clear. The kind of approach championed by Pinker, Prince, and others offered a simple and precise solution to the narrow challenge of explaining the U-shaped developmental profile. Connectionist models performed in less precise ways, but had the virtue of being applicable to a broad range of cases rather than being a tailor-made system of symbols and rules that applied narrowly to the specific problem of learning the English past tense. In addition to being more general, connectionists claimed that their models provided the required level of generalization and productivity, without the unparsimonius proliferation of modules that came with symbol-based approaches. It was in this atmosphere of excitement about the prospects of connectionist theory that Fodor and Pylyshyn published their article. As for the debate over the English past tense, we continue to see a lack of consensus as to whether statistical mechanisms embedded in artificial neural networks are genuinely explanatory (Ramscar 2002; Pinker and Ullman 2002).

The terms of this disagreement echo Fodor and Pylyshyn's discussion of what it means to give an explanation of a cognitive phenomenon. Research on language acquisition in infants (Marcus, Vijayan, Bandi Rao, and Vishton 1999; Seidenberg and Elman 1999) and speech processing in adults (Peña et al. 2002; Endress and Bonatti 2007; Laakso and Calvo 2011) are areas where many of the central themes of the original debate continue to be explored.[6] The set of considerations that were raised in debates over the past tense have spread throughout the field of psycholinguistics and the rest of the cognitive sciences, with contemporary versions of the dispute being held between "probabilistic" (Griffiths, Chater, Kemp, Perfors, and Tenenbaum 2010) and "emergentist" (McClelland, Botvinick, Noelle, Plaut, Rogers, Seidenberg, and Smith 2010) models of thought.

While the so-called great past tense debate was affected directly by Fodor and Pylyshyn's criticisms, their argument was aimed at a very general set of questions rather than at the specifics of a particular debate in science. They were unmoved by the fact that this or that area of cognition could actually be modeled statistically. On their view, a simulation that mimics some feature of human cognition is not an explanation of that feature. Their criticism applied equally, for example, to single-layered and multi-layered perceptrons; they were relatively unconcerned with the intricacies of particular battle lines in particular cases. Their criticisms were fundamentally philosophical, targeting the underlying associative character of neural network processing per se. In their view, the "parallel distributed processing" of the 1980s was basically equivalent to the kind of associationism that philosophers would associate with Locke or Hume. Insofar as it claims to be a theory of cognition, connectionism is simply associationism dressed up in the jargon of vectorial patterns of activation, matrices of weighted connections, and gradient descent learning. Any claim to being neurobiologically plausible was also irrelevant to their criticisms of connectionism. Fodor and Pylyshyn's chief concern was the very nature of human cognition, rather than the details of how cognition happened to be implemented in the nervous system. Insofar as connectionist theory echoed British empiricist philosophy of mind, the inferential treatment of classical cognitive science took the side of the rationalists. Thus, while they sided with Pinker and Prince (1988) against Rumelhart and McClelland (1986) in the dispute over the English past tense, Fodor and Pylyshyn had their eye on traditional philosophical questions. They famously conclude their essay by acknowledging that these debates have a venerable heritage in the history of philosophy as well as the more recent history of psychology: "We seem to remember having been through this argument before. We find ourselves with a gnawing sense of déjà vu" (Fodor and Pylyshyn 1988, 70).

By acknowledging the historical precedents for the debate, Fodor and Pylyshyn certainly did not mean to imply that the situation is a stalemate or that this debate exemplifies perennial problems in philosophy that will never be resolved. In their view, there is a clear winner: when it comes to cognition, rationalism is, as Fodor (1975) had earlier claimed of the language of thought hypothesis, *the only game in town.*

The Systematicity of Thought

At the heart of the systematicity debate is a basic disagreement concerning what, precisely, needs to be explained in the science of mind. For example, in the case of language, is it the actual linguistic parsing and production performed by cognitive agents or their competence-level characterization in generative linguistic terms (Chomsky 1965)? What is the correct level of description that captures the distinctively cognitive core of the phenomenon? Another basic difference between the classicists and the connectionists was their attitude toward the autonomy of psychology relative to the details of implementation (Fodor 1974). Fodor and Pylyshyn's strategy was to focus on what they took to be a clear-cut and pervasive set of phenomena that they regarded as uncontrovertibly cognitive and to downplay questions of implementation and performance.

Fodor and Pylyshyn focused on the productive and systematic features of thought along with its inferential coherence. Notwithstanding constraints on human hardware capacities, the productivity of thought refers to our capacity to entertain or grasp an indefinitely large number of thoughts. Thought processes, on the other hand, are systematic to the extent that our capacity to entertain or grasp a thought appears to be intrinsically connected with our capacity to entertain or grasp a number of other semantically related thoughts. Likewise, thought exhibits inferential coherence. Our capacity to follow a pattern of inference appears intrinsically connected to our capacity to draw certain other inferences. The productivity, systematicity, and inferential coherence of thought strongly suggest that mental representations possess a constituent structure without which it is difficult to come to terms with the interconnections among thought-related capacities.[7]

According to Fodor and Pylyshyn, an inference to the best explanation should lead us to regard these explananda as involving operations performed on a stock of representations that can be combined and recombined in accordance with a set of rules. On their view, unless we postulate syntactic and semantic combinatorial relations and unless thought is compositional, the way we are endowed with these abilities remains a mystery. Unstructured connectionist networks would, at first sight, lack the resources to explain the productivity, systematicity, or inferential coherence of thought.

Fodor and Pylyshyn developed a number of parallel lines of argument in support of the superiority of a symbol-and-rule-based approach. In the case of productivity, as Fodor and Pylyshyn observe, the existence of structured representational schemata is inferred from the fact that our competence does not seem to be finite; we appear to be able to entertain an indefinitely large number of thoughts. Arguments from productivity invite a number of responses. Calling the very existence of a competence–performance divide into question was one common line of response. A less direct approach that was favored by some scientists was to simply consider performance and competence separately, approaching one side of the divide and not the other on methodological grounds (see Elman, this volume; Frank, this volume). Although productivity has been a central topic that figured in responses to the 1988 paper, it strikes us that the core of Fodor and Pylyshyn's argument is the notion that our capacity to think is not *punctate*; once we have entertained a thought, we have the resources to deal with others that are semantically related. This is the critical datum that compels us, on their view, to see thought as requiring structured representational schemata.

The simplicity of the argument's starting point is powerful and compelling. It is intuitively obvious that a speaker's capacity to understand native sentences of her language is related to her understanding of a number of other semantically related sentences. Pathologies aside, it is difficult to imagine that someone could understand the sentence "John loves Mary" without having *ipso facto* the resources to produce or understand "Mary loves John."[8]

Imagine attempting to learn a language using only a phrase book. Punctate understanding is what one gets when a phrase book is one's only access to a foreign language. What you can say and understand depends on the particular phrases you happen to look up. This contrasts sharply with the way competent speakers of a language understand sentences. The question is whether this intuitive starting point is sufficient to license Fodor and Pylyshyn's claim that plausible models of human thought must have a classical combinatorial structure.

According to Fodor and Pylyshyn, the systematicity of thought argument against connectionism as an architectural hypothesis of cognition runs as follows:

(i) It is a fact of psychology that thought is systematic insofar as our thoughts are intrinsically related to one another in such a way that having one thought means having the capacity to access an indefinitely large set of other thoughts. So, to take the canonical illustration, someone who can think JOHN LOVES MARY must have the capacity to think MARY LOVES JOHN.

(ii) An explanation of systematicity requires syntactic and semantic constituency relations among mental representations and a set of processes that are sensitive to such internal structure such as those provided by the language of thought (LOT) hypothesis.

(iii) Connectionist theory posits neither syntactic and semantic constituency relations among mental representations, nor any set of processes that are sensitive to the internal structure of mental representations.
(iv) Therefore, connectionism is unable to account for the systematicity of thought.

It should be noted that the fact that LOT provides the explanatory framework required is meant to imply that it *guarantees* the phenomenon and not merely that it is compatible with its occurrence. The explanatory framework is meant to satisfy the sense that the cognitive architecture and cognition itself share a common core. On Fodor and Pylyshyn's view, a genuine explanation involves a robust constraint on acceptable architectures: the demand is that the model accords with systematicity in a way that is not merely the product of an exercise in data-fitting (see Aizawa 2003). Instead, Fodor and Pylyshyn are eager to emphasize that systematicity follows from classical architecture as a matter of nomological necessity and is not simply a contingent fit between the architecture and the explanandum. Of course, this does not mean that it is conceptually impossible that other architectures might play the same explanatory role or even that some other architecture might also have a necessary connection with the explanandum. Indeed, there may be psychological theories that play an equivalent role, but in order to satisfy the Fodorian demand it is not enough to show that the framework is compatible with the explanandum. Insofar as the systematicity of thought is understood to be a basic psychological fact, or even a psychological law (McLaughlin 2009), it would not suffice for a neural modeler to hit upon a configuration of weights that *happens to* allow a connectionist network to mimic the cognitive explanandum.[9] It would need to be shown how it follows of necessity from that architecture, as is supposedly the case for LOT (but see Phillips, this volume). Fodor and Pylyshyn's empirical bet is that we will not be able to find successful explanations of the systematicity of thought that do not involve full-fledged compositional semantics.

Cognitive Architecture

For Fodor and Pylyshyn, the issue is whether connectionism can be understood to serve as an explanatory *cognitive* theory, rather than as a high-level description of the underlying neural substrate of cognition. But this raises some core issues with regard to the very notion of "cognitive architecture."

A cognitive architecture provides the theoretical framework that constrains and aims at explaining the putative mental capacities of a physical system. However, fixing precisely what is meant by a "cognitive architecture" is a basic conceptual problem in itself. Fodor and Pylyshyn adopted Newell's (1980, 1982) distinction between a knowledge level and a physical-symbol-system level in their treatment of neural networks. Newell noted that "given a symbol level, the architecture is the description of the system in whatever system-description scheme exists immediately next below the symbol level" (1990, 81). Appealing in a physical symbol system to the level *immediately below* the symbol level implies a clear division of labor in a cognitive architecture. Clearly, this division of labor is congenial with the autonomy of psychology as championed by Fodor (1974). According to Fodor and Pylyshyn:

> The architecture of the cognitive system consists of the set of basic operations, resources, functions, principles, etc. (generally the sorts of properties that would be described in a "user's manual" for that architecture if it were available on a computer), whose domain and range are the *representational states* of the organism. (1988, 5)

Although they make no reference to David Marr in their essay, the latter promoted a particular account of explanation and a distinction of levels that was common currency in the 1980s, and which has helped to frame the discussion of the architecture of cognition ever since. What Fodor and Pylyshyn (and Newell) are after is exemplified by Marr's (1982) well-known tripartite approach to the description of cognitive systems in terms of the computational, the algorithmic, and the implementational. We may thus read them as endorsing a Marrian view of the appropriate analysis of levels. Their aim is to explain the systematicity of thought *algorithmically*, once the phenomenon has been defined at the computational level, and with details of substrate implementation being entirely left aside.

Fodor and Pylyshyn's critical analysis targets neural network modeling only insofar as it presumes to address the cognitive level of explanation. It is in this *top-down* spirit that their challenge is to be read. Fodor and Pylyshyn argue that accounting for the systematicity of thought is only achieved to the extent that connectionist models import the structural features of classical combinatorial processes. However, insofar as this is the case, no alternative algorithm or explanatory framework is actually being provided. On this view, *structured* neural networks are relegated to the status of "implementational connectionism" (Pinker and Prince 1988). Thus, classicists will happily concede that connectionism may well be able to unearth details of the neural substrate that allow for the physical implementation of structured sets of mental representations. Clearly, having the right story about the neural hardware would be a welcome supplement to psychological theorizing. The target for the classicist is possession of an algebraic, symbolic, and rule-based explanation of systematicity, on the one hand, and a neurological understanding of how the capacity to remain structurally sensitive to compositional processes is implemented in the brain and nervous system, on the other.

Let's consider the trade-off between neural structure and cognitive function. Traditionally, the discussion has taken the form of determining what the correct level of description or explanation is (Broadbent 1985). The tendency of both the "eliminative connectionists" and the classicists (Pinker and Prince 1988) has been to focus exclusively on either neural structure or cognitive function. On the cognitive side, the mission has been to explain systematicity in terms of the set of "operations, resources, functions and principles" that Fodor and Pylyshyn regard as governing the representational states of a physical system. By contrast, the eliminative connectionist tendency has been to focus on the structure and processes of the brain and nervous system rather than on cognition.

The idea that there are cognitive mechanisms that explain the systematicity of the human mind, and that any empirically adequate theory should incorporate them, is not at issue. Instead, disagreements center on the question of the relationship between the cognitive and the biological levels of analysis. Although Marr's understanding of the methodological relationship between computational, algorithmic, and implementational levels was highly influential for cognitive science, many philosophers and scientists have argued for a reevaluation of his tripartite methodological framework (Symons 2007). Nowadays, most cognitive scientists find it difficult to accept, without significant qualification, the top-down recommendation that understanding the goal of computations should take priority over the investigation of the implementational level. Marr understood investigations of the computational level to involve the determination of the problem that a system was forced to solve. However, the classical hierarchy of (autonomous) computational, algorithmic, and implementational levels of analysis is not uncontroversially assumed by contemporary cognitive scientists. Neurobiological constraints have become centrally important to the characterization of the computational level. In general, there is a growing tendency to merge top-down and bottom-up considerations in the determination of the architecture of cognition.[10] In this way, contemporary treatments of the correct architecture of cognition concern the relationship between implementation structure and cognitive function.

Another important feature of these debates is the central place given to representation in classicist arguments against connectionism. This topic is far too subtle and complex for an introductory overview, but it is important to note that Fodor and Pylyshyn presented their challenge to connectionism in representationalist terms:

If you want to have an argument about *cognitive* architecture, you have to specify the level of analysis that's supposed to be at issue. If you're *not* a Representationalist, this is quite tricky since it is then not obvious what makes a phenomenon cognitive. (1988, 5)

It is not at all clear to us how one ought to read this appeal to representationalism in the contemporary context. The connectionist of the 1980s would not see this as particularly problematic insofar as neural network modeling was generally speaking committed to some sort of representational realism, typically in the form of context-dependent subsymbols. But it is no longer the case that contemporary connectionists would accept the kind of representationalist view of the mind that Fodor and Pylyshyn assumed. At the very least, the classicist would need to provide a more developed argument that the denial of standard representationalism is equivalent to some form of noncognitivist behaviorism.[11] While the challenge is not directly aimed at nonrepresentational connectionists, this does not prevent nonrepresentational connectionism from having something to say about systematicity.

Of course, Fodor and Pylyshyn assume that the etiology of intentional behavior must be mediated by representational states. As far as explanation is concerned, this assumption should not be understood as placing an extra, asymmetrical burden of proof on someone wishing to provide nonrepresentationalist accounts of the phenomenon in question. Calvo, Martín, and Symons (this volume) and Travieso, Gomila, and Lobo (this volume) propose neo-Gibsonian approaches to systematicity. While Fodor and Pylyshyn were committed representationalists, neo-Gibsonians can take the systematicity of thought or the systematicity of behavior as an explanandum for cognitive and noncognitive architectural hypotheses alike, irrespective of whether those hypotheses include representations.

Twenty-Five Years Later: Taking Stock of the Architecture of Cognition

Systematicity arguments have figured prominently in discussions of cognitive architecture from the heyday of connectionism in the 1980s and '90s to the advent of a "post-connectionist" era in the last decade. Many of the aforementioned concerns are either explored or cast in a new light in the chapters of this volume.

The present volume represents a collective effort to rethink the question of systematicity twenty-five years after Fodor and Pylyshyn's seminal article in light of the wide variety of approaches in addition to connectionist theory that are currently available. As more implementational details are being honored, artificial neural networks have started to pay closer attention to the neurobiology (Borensztajn et al., this volume; O'Reilly et al., this volume). Overall, what we find is a greater sensitivity in the scientific literature to the aforementioned trade-off between structure and function. Memory, for instance, is not understood as the general configuration of a weight matrix, but instead is modeled by means of specific, neurobiologically plausible attractor networks in which components and their activities are organized with great precision and with considerable effort to maintain biological plausibility.

We may then consider whether networks of these kinds either meet the systematicity challenge or change the terms of the debate in significant ways. In what sense does Fodor and Pylyshyn's critical analysis extend to nonclassical forms of connectionism? When they objected to the possibility of having "punctate minds," they were thinking of models like those one finds in the work of Hebb, Osgood, or Hull; connectionism, according to Fodor and Pylyshyn, was *nothing but old wine in new bottles*. Whether or not they are correct in their judgment that nothing of substance had changed in the period between the precursors of connectionism and the 1980s,[12] superficially at least, there seem to be many reasons to reassess the debate today in light of the widespread use of "nonclassical connectionist" approaches. Fodor and Pylyshyn would acknowledge that connectionist networks have grown in sophistication, but that the basic principles remain the same. Had they written their paper in 2013, the argument would have probably been that at the cognitive level, the architecture of the mind is not a "nonclassical connectionist parser" (Calvo Garzón 2004), and that nonclassical connectionism may at best provide an implementational account of thought. The basic lines of the argument still apply, in spite of impressive scientific developments. This, in itself, is an interesting feature of the debate. Fodor and Pylyshyn are presenting an argument for a particular conception of explanation. Connectionists, no matter how sophisticated, have simply missed what is most important and interesting about cognition. The fact that connectionists are likely to have grown very weary of responses of this kind does not mean that Fodor and Pylyshyn are simply wrong.

In addition to connectionism, a constellation of methodologies and architectures have entered the field, many of which explicitly tackle basic questions concerning the nature of explanation in cognitive science. Modern post-connectionist viewpoints include dynamical, embodied, and situated cognitive science, the enactive approach, and neo-Gibsonian approaches. By focusing their criticism on connectionism, Fodor and Pylyshyn limited the range of hypotheses under consideration in their original article. Throughout his career, Fodor has engaged directly with alternative conceptions of psychological explanation.[13] However, it is fair to say that the Fodorian side of the systematicity debate has maintained (rightly or wrongly) a very fixed picture of what counts as a genuinely cognitive explanation.

The present volume represents the state of play in 2013. Aizawa (chapter 3, this volume) argues that, if one is to pay close attention to what ecological psychology, enactivism, adaptive behavior, or extended cognition actually say, it is unclear what the dividing line between cognition and behavior is—either because these methodologies are behaviorally inclined in themselves, downplaying their relevance to questions concerning cognition, or (even worse) because they identify or conflate cognition itself with behavior. These constitute new challenges to the systematicity arguments in the post-connectionist era. Of course, if that is the case, the rules of the game may no longer be clear. One way or another, Aizawa concludes, the post-connectionist era brings about "tough times to be talking systematicity."

In the immediate aftermath of Fodor and Pylyshyn's article, classicists and connectionists focused primarily on the dichotomy between context-free versus context-dependent constituency relations. Context-dependent constituents that appear in different thoughts as syntactically idiosyncratic tokens were first discussed by Smolensky (1987) and Chalmers (1990). This idea is revisited by Brian McLaughlin (chapter 2, this volume), who takes issue with Smolensky and Legendre's most recent views as presented in *The Harmonic Mind*. According to McLaughlin, Smolensky and Legendre's integrated connectionist symbolic architecture is only able to explain systematicity and productivity, despite its hybridity, by collapsing into a full-fledged LOT model.

Gary Marcus (chapter 4, this volume) defends the idea that the mind has a neurally realized way of representing symbols, variables, and operations over variables. He defends the view that the mind is a symbol system against eliminativist varieties of connectionism. On Marcus's view, minds have the resources to distinguish types from tokens and to represent ordered pairs and structured units, and have a variety of other capacities that make the conclusion that minds are symbol systems unavoidable. In his view, connectionist architectures have proved unable to exhibit this set of capacities. Connectionism therefore continues to be subject to the same sorts of considerations that Fodor and Pylyshyn raised twenty-five years ago. Marcus, however, identifies one particular capacity where classicism has not been able to succeed, namely, in the representation of arbitrary trees, such as those found in the treatment of syntax in linguistics. Unlike computers, humans do not seem to be able to use tree structures very well in mental representation. While we can manipulate tree structures in the abstract, our actual performance on tasks requiring manipulation of tree-like structures is consistently weak.

Marcus appeals for an integrative approach to problems of this kind, arguing that the symbolic and statistical features of mind should be modeled together. Nevertheless, on Marcus's view, the human mind is an information-processing system that is essentially symbolic, and none of the developments in the years since Fodor and Pylyshyn's paper should shake that conviction.

Fodor and Pylyshyn focused on rules as a source of systematicity in language. In chapter 5, Jeff Elman points out that the intervening years have seen an increased interest in the contribution of lexical representations to the productivity of language. The lexicon was initially thought to be a relatively stable set of entities with relatively little consequence for cognition. Elman notes that the lexicon is now seen as a source of linguistic productivity. He considers ways in which systematicity might be a feature of the lexicon itself, and not of a system of rules. He proceeds to provide a model of lexical knowledge in terms of performance and distributed processing without positing a mental lexicon that is independently stored in memory.

Frank (chapter 6, this volume) considers that neural network success in accounting for systematicity cannot rely on the design of toy linguistic environments, and he scales up simple recurrent networks (SRNs; Elman 1990) by using more realistic data from computational linguistics. In his chapter, he sets out to compare empirically a connectionist recurrent neural network with a probabilistic phrase-structure grammar model of sentence processing in their systematic performance under more or less realistic conditions. As Frank reports, the performance of both models is strikingly similar, although the symbolic model displays slightly stronger systematicity. In his view, nevertheless, real-world constraints are such that in practice performance does not differ susbtantially across models. As a result, the very issue of systematicity becomes much less relevant, with the litmus test residing in the learning and processing efficiency of the models in dispute, in their flexibility to perform under adverse conditions, or in the explanation of neuropsychological disorders. This range of phenomena, among others, is what in Frank's view "getting real about systematicity" boils down to. Interestingly enough, systematicity, as presented by Fodor and Pylyshyn, may not be that relevant after all.

Overall, what this type of neural modeling hints at, regardless of the level of realism involved, is that it is the nonlinear dynamics that result from a word's effect on processing that counts. Constituency may be understood in terms of dynamical basins of attraction, where convergence toward stability is compatible with dynamic states involved in combinatorial behavior being transient. It is then a step forward that falls between SRNs and other connectionist networks, and dynamical systems theory.[14] In this way, if SRNs exploit grammatical variations as dynamical deviations through state space, the explanatory framework and formal tools of dynamic systems theory (Port and van Gelder 1995) provides yet a more solid avenue of research. The working hypothesis is that there is no need to invoke information-processing concepts and operations, with combinatorial behavior grounded in sensorimotor activity and the parameter of time. However, although there is a trend to replace the symbols and rules of classical models with quantities, different types of attractors and their basins may furnish different dynamical means of implementing combinatorial structure. Thus, monostable attactors (globally stable controllers inspired in cortical circuitry and which hold single basins of attraction; Buckley et al. 2008), for instance, may hint toward different sets of solutions than those in terms of the trajectories that get "induced by sequences of bifurcations ('attractor chaining')" (van Gelder 1998). Either way, the dynamical setting of monostable attractors, attractor basins, or attractor chaining points toward alternative ways to understand cognition and its temporal basis. Other connectionist proposals that have exploited some of the toolkit of dynamic systems theory use articulated attractors (Noelle and Cottrell 1996), including the deployment of wide enough basins of attraction to capture noisy patterns, and stable enough attractors to remember input indefinitely. But humans do not implement arbitrarily long sequential behavior. If compositionality is to be modeled, it seems it have to depend on other sort of resources, than memory resources per se.[15]

Of course, connectionist "structure-in-time" models are incomplete in a number of respects, most notably in being disembodied and in the fact that the vectorial representations they make use of cannot be taken as theoretical primitives. If we pay attention to ontogenesis in the developmental psychology experimental literature, we find that models point toward the decentralization of cognition (Thelen et al. 2001). It is then a step forward to move from SRNs and dynamical systems theory to an embodied and situated cognitive science (Calvo and Gomila 2008; Robbins and Aydede 2008). Hotton and Yoshimi (2011), for instance, exploit "open" (agent-*cum*-environment) dynamical systems to model embodied cognition with dynamic-based explanations of perceptual ambiguity and other phenomena. According to embodied cognitive science, we should be phrasing *cognitively* the following question: what is it that adults *represent* from the world that allows them to behave systematically and productively? This rendering of the situation presupposes an answer where a connectionist or dynamical phase space obtains stable representational states (regardless of whether they are context-dependent or collapse into context-free states). But embodied constituents are not hidden manipulable states, but rather states that change continuously in their coupling with the environment. A system converges to stable states from nearby points in phase space as a result of external conditions of embodiment of the system and endogenous neurally generated and feedback-driven activity. The appropriate question may then be what sensory-to-neural continuous transformations permit adults to exhibit a combinatorial behavior.

In contrast with connectionist, more or less traditional lines of response, Coram (chapter 11, this volume) focuses on the extended theory of cognition, a framework that has proved valuable in informing dynamic systems models of the mind. Putative explanations of systematicity reside in the wider cognitive system, something that includes language and other structures of public representational schemes. Coram compares this extended shift to a strategy that Pylyshyn has tried out in his research on imagistic phenomenon, and argues that embodied and embedded cognitive science need not redefine the phenomenon of systematicity itself, but rather can account for it in its classical clothes with some revisions to the concept of representation at play. Her proposed explanation combines extended explanatory structures with internal mechanisms.

The systematic features of visual perception appear to be an example of a non-linguistic forms of systematicity (Cummins 1996). Following this thread, in chapter 16, Calvo, Martín, and Symons show how systematicity may also emerge in the context of simple agents, taking a neo-Gibsonian perspective to the explanation of this form of systematicity. The objective of this chapter is to provide an explanation of the emergence of systematic intelligence per se rather than providing a defense of a particular cognitive architecture. To this end, Calvo, Martín, and Symons examine marginal cases of behavioral systematicity in the behavior of minimally cognitive agents like plants and insects, rather than beginning with the linguistically mediated cognition of adult human beings, with the intention to provide a basis for understanding systematicity in more sophisticated kinds of cognition.

Other authors such as Travieso, Gomila, and Lobo (chapter 15, this volume) favor a dynamical, interactive perspective and discuss the alleged systematicity of perception as illustrated by the phenomenon of amodal completion (see Aizawa, this volume). According to Travieso et al., amodal completion emerges globally out of context-dependent interactions and cannot be explained compositionally. They further discuss systematicity in the domain of spatial perception, and argue that although systematic dependencies are not found in perception in general, a Gibsonian ecological approach to perception that recurs to higher-order informational invariants in sensorimotor loops has the potential to explain a series of regularities that are central to perception, despite remaining unsystematic. Research on sensory substitution and direct learning serves to make their case. One way or another, it seems that the fact that extracranial features (bodily or environmental) play a constitutive role for the sake of cognitive processing is compatible both with cognition being extended and with cognition reducing to behavior while the latter is accounted for directly in neo-Gibsonian terms.

As we've already mentioned, probably every single aspect of Fodor and Pylyshyn's argument has been questioned. If from the connectionist corner, rejoinders included buying into constituent structure with an eye to unearthing implementational details of an otherwise LOT cognitive architecture or developing an alternative form of context-dependent constituent structure, more recent post-connectionist responses include variations such as the development of spiking neural network and plasticity implementational models (Fernando 2011) or realistic linguistic settings to feed SRNs without resorting to constituency internalization. It might be possible to bypass classical compositionality by individuating neural network internal clusters in a hierarchical manner (Shea 2007). Work with a SINBAD (set of interacting backpropagating dendrites; Ryder 2004) neural model that exploits cortical hierarchies for the purpose of allowing for increasing generalization capabilities by scaffolding variables as we move cortically away from the periphery of the system may be read in this light. The Hierarchical Prediction Network (HPN) of Borensztajn, Zuidema, and Bechtel (chapter 7, this volume) points in this direction, contributing a notable combination of functional abstraction with substrate-level precision.

In particular, Borensztajn et al., inspired by Hawkins's Memory Prediction Framework, and departing from the emergentist view of unstructured connectionist modeling, explore dynamic binding (Hummel and Biederman 1992) among processing units, and develop a neurobiologically plausible version that is structured and can thus account for systematicity. Encapsulated representations that result from the hierarchical organization of the cortex enable categories to play causal roles. In their chapter, they elaborate on how encapsulated representations can be manipulated and bound into complex representations, producing rulelike, systematic behavior. This, Borensztajn et al. argue, does not make their proposal implementational, because despite the fact that encapsulated units can act as placeholders, the value of encapsulated representations only gets set in the system's interaction with an external environment. As to how to combine encapsulated representations into complex representations, Borensztajn et al. show how temporary linkages between representations of the sort allowed by dynamic binding (Hummel and Biederman 1992) may deliver the goods.

If a hierarchical category structure can play causal roles and account for systematicity by treating constituents as "substitution classes," with an eye to exploiting encapsulated representations for the purpose of respecting compositionality but without retaining classical constituents, Phillips and Wilson (chapter 9, this volume) rely on "universal constructions" for the same purpose. According to them, neither classicism, nor connectionism, nor dynamicism, among other methodologies, for that matter, has managed thus far to fully explain the systematicity of human thought in a way that is not ad hoc. They propose instead a *categorial* cognitive architecture: a category theoretic (Mac Lane 2000) explanation based on the concept of a "universal construction" that may constitute the right level of description to inform empirical sciences. Whereas substitution classes do the trick in hierarchical prediction networks, their model, relying on a formal theory of structure, relates systematically maps of cognitive processes that are structurally preserved, allegedly meeting Fodor and Pylyshyn's challenge.

On the other hand, the SAL framework of O'Reilly, Petrov, Cohen, Lebiere, Herd, and Kriete (chapter 8, this volume) provides a synthesis of ACT-R (Adaptive Control of Thought—Rational; Anderson and Lebiere 1998) and the Leabra model of cortical learning (O'Reilly and Munakata, 2000) and is a plea for pluralism in the cognitive sciences in the form of a biologically based hybrid architecture, where context-sensitive processing takes place first, on the ground of evolutionary and online-processing considerations. Their systems neuroscience approach is then aimed at illuminating how partially symbolic processing, to the extent that human performance happens to approximate a degree if systematicity, is the result of complex interactions, mainly in the prefrontal cortex/basal ganglia (PFC/BG) system.

Although Fodor and Pylyshyn presuppose architectural monism to be the default stance, a commitment to some form of pluralism is shared by a number of authors in this volume. In fact, strategies inspired by "dual-process" theories have gained increasing support in recent years (Evans and Frankish 2009). The dual-process working hypothesis is that the architecture of cognition is split into two processing subsystems, one older than the other, evolutionary speaking. Whereas the former puts us in close relation to our fellow nonhuman animals (e.g., pattern-recognition), the latter system is in charge of abstract reasoning, decision making, and other competencies of their ilk. Gomila et al. (2012), for instance, adopt a dual-process framework to argue that systematicity only emerges in the restricted arena of the newer subsystem. In their view, thought's systematicity is due to the fact that human animals are *verbal* (Gomila 2012). Language, in the external medium, underlies our ability to think systematically (see also Coram, this volume; Travieso et al., this volume). Martínez-Manrique (chapter 12, this volume), in turn, argues that connectionist rejoinders to Fodor and Pylyshyn's challenge have never been satisfactory. Systematicity, nevertheless, need not be a general property of cognition, and in that sense there is some room to maneuver. Martínez-Manrique motivates a variety of conceptual pluralism according to which there are two kinds of concepts that differ in their compositional properties. Relying in part on the dual-process approach, he suggests a scenario of two processing systems that work on different kinds of concepts. His proposal boils down to an architecture that supports at least two distinct subkinds of concepts with different kinds of systematicity, neither of which is assimilable to each other.[16]

Architectural pluralism retains, in this way, at least partially, a commitment to representations, but there are other options. In his contribution, Ramsey (chapter 10, this volume) contends that the fact that systematicity is a real aspect of cognition should not be seen as bad news for connectionism. Failure to explain it does not undermine its credibility since the mind need not have only one cognitive architecture. Ramsey is again calling for some form of architectural pluralism via dual-process theories that would allow connectionist theory to shed its distinctive light on those aspects of cognition that remain unsystematic. However, his proposal is radical enough to allow for the vindication of connectionism despite not just their inability to account for systematicity, but also their not constituting a representational proposal.[17]

In chapter 13, Edouard Machery applies Fodor and Pylyshyn's early criticism of connectionist models to neo-empiricist theories in philosophy and psychology. After reviewing the central tenets of neo-empiricism, especially as presented in the work of Jesse Prinz and Lawrence Barsalou, Machery focuses on its characterization of occurrent and non-occurrent thoughts, he argues that amodal symbols are necessary conditions for non-occurrent thoughts. On Machery's view, if occurrent thoughts are individuated by their origins then some feature of the architecture of cognition other than the contingent history of the learning process is needed to account for thought's inferential coherence. Machery's chapter reflects the continuing influence that systematicity arguments have in contemporary philosophy of psychology.

Still more radical departures come from systems neuroscience, a field that, by integrating the scales of specific neural subsystems with constraints of embodiment for cognition and action (Sporns 2011), provides a further twist in the tale. In fact, from the higher point of view of vastly interconnected subnetworks, the brain as a complex system appears to some authors to defy a part-whole componential reading. With the focus placed in the brain-body-environment as a complex system, Chemero (chapter 14, this volume) calls our attention to a *radical embodied cognitive science* (2009) that suggests that cognitive systems are interaction dominant to some extent, and that this requires that we fully revisit Fodor and Pylyshyn's notion of systematicity. He describes a number of examples and argues that interaction dominance is inconsistent with the compositionality of the vehicles of cognition. Since compositionality underlies the phenomenon of systematicity, cognition happens not to be systematic at least to the extent that cognitive systems are interaction dominant. In addition, Silberstein (chapter 17, this volume) combines systems neuroscience and psychopathology to shed light on theories of standard cognitive functioning. In particular, he proposes to make some empirical progress by studying the effects on cognition and behavior when inferential coherence fails to obtain in patients with schizophrenia. According to Silberstein, the fact that the absence of dynamical subsymbolic properties of biological neural networks correlates with the breakdown in systematic inferential performance tells against a symbol-and-rule approach.

One can imagine a variety of approaches to the evaluation of theories of mind in addition to the systematicity criterion. So, for example, behavioral flexibility, faithfulness to developmental considerations, or performance in real time could all constitute plausible functional constraints on the architecture of cognition.[18] It goes without saying that there are many more approaches that could potentially shed light on these questions than those presented in the pages of this book. Our purpose here has been to consider a sample of nonclassical connectionist empirical and theoretical contestants in light of the conceptual challenge that Fodor and Pylyshyn articulated. The central question for readers is whether the symbol-and-rule stance is required for genuine explanations in cognitive science. In the late 1980s, it was clear to Fodor and Pylyshyn what the answer should be. It may be the case twenty-five years later that their verdict would remain the same, but assessing whether their original arguments continue to have the same force for the range of approaches included in this volume is something we leave to the reader to judge.

Acknowledgments

Many thanks to Alicia Coram for very helpful comments. This research was supported by DGICYT Project FFI2009-13416-C02-01 (Spanish Ministry of Economy and Competitiveness) to PC, and by Fundación Séneca-Agencia de Ciencia y Tecnología de la Región de Murcia, through project 11944/PHCS/09 to PC and JS.

Notes

1. The most thorough examination of the arguments associated with the systematicity challenge is Kenneth Aizawa's 2003 book, *The Systematicity Arguments.*

2. As David Chalmers wrote, their trenchant critique "threw a scare into the field of connectionism, at least for a moment. Two distinguished figures, from the right side of the tracks, were bringing the full force of their experience with the computational approach to cognition to bear on this young, innocent field" (1990, 340).

3. At the time of this writing, "Connectionism and Cognitive Architecture: A Critical Analysis" has been cited over 2,600 times, according to Google Scholar.

4. See, e.g., John Bickle's (2003) defense of the philosophical significance of developments in cellular and subcellular neuroscience.

5. For an account of the history of connectionism, see Boden 2006, ch. 12.

6. See Marcus 2001 for a critical appraisal from the classicist perspective.

7. According to Fodor (1987), constituents appear in different thoughts as syntactically identical tokens. "The constituent 'P' in the formula 'P' is a token of the same representational type as the 'P' in the formula 'P&Q', if 'P' is to be a consequence of 'P&Q'" (Calvo Garzón 2000, 472).

8. On pathological cases, see Silberstein, this volume.

9. After all, neural networks are universal function approximators (Hornik, Stinchcombe, and White 1989). Thus, since they are Turing equivalent (Schwarz 1992), the worry is not whether they can compute, but rather whether they can compute "systematicity functions" without implementing a classical model in doing so.

10. Clark (2013) and Eliasmith (2013) are recent illustrations.

11. Ramsey (2007), for example, has recently argued that only classical cognitive science is able to show that a certain structure or process serves a representational role at the algorithmic level. Connectionist models, Ramsey argues, are not genuinely representational insofar as they exploit notions of representation that fail to meet these standards (but see Calvo Garzón and García Rodríguez 2009).

12. One could argue that already in the 1980s they were working with an incomplete picture of the state of network theory. They make no reference to Stephen Grossberg's adaptive resonance approach to networks, for example.

13. See, e.g., Fodor and Pylyshyn's (1981) response to Gibson.

14. As a matter of fact, the line between connectionist and dynamicist models of cognition is anything but easy to draw (see Spencer and Thelen 2003).

15. In the case of some "structure-in-time" models, such as "long short-term memory" models (LSTM; Schmidhuber Gers and Eck 2002), the implementational outcome is more clearly visible. Long short-term memories are clusters of nonlinear units arranged so that an additional linear recurrent unit is places in the middle of the cluster, summing up incoming signals from the rest. The linear unit allows the system to maintain a memory of any arbitrary number of time steps, which apparently would make the model collapse into our original context-free versus context-dependent dichotomy. In addition, the linear units that LSTM models employ are unbiological.

16. Interestingly, as Martínez-Manrique discusses, if cognitive processes happen not to be systematic in Fodor and Pylyshyn's sense, a nonclassical systematicity argument may be run by analogy to their systematicity argument.

17. This is something Ramsey has argued for elsewhere (Ramsey 2007): that cognitive science has taken a U-turn in recent years away from representationalism and back to a form of neobehaviorism (see Calvo Garzón and García Rodríguez 2009 for a critical analysis).

18. For other constraints and further discussion, see Newell 1980 and Anderson and Lebiere 2003.

References

Aizawa, K. 2003. *The Systematicity Arguments*. Dordrecht: Kluwer Academic.

Anderson, J. R., and C. Lebiere. 1998. *The Atomic Components of Thought*. Mahwah, NJ: Erlbaum.

Anderson, J. R., and C. Lebiere. 2003. The Newell test for a theory of cognition. *Behavioral and Brain Science* 26:587–637.

Berko, J. 1958. The child's learning of English morphology. *Word* 14:150–177.

Bickle, J. 2003. *Philosophy and Neuroscience: A Ruthlessly Reductive Account*. Dordrecht: Springer.

Boden, M. A. 2006. *Mind as Machine: A History of Cognitive Science*. Oxford: Oxford University Press.

Broadbent, D. 1985. A question of levels: Comment on McClelland and Rumelhart. *Journal of Experimental Psychology: General* 114:189–192.

Buckley, C., P. Fine, S. Bullock, and E. A. Di Paolo. 2008. Monostable controllers for adaptive behaviour. In *From Animals to Animats 10: The Tenth International Conference on the Simulation of Adaptive Behavior*. Berlin: Springer.

Calvo Garzón, F. 2000. A connectionist defence of the inscrutability thesis. *Mind and Language* 15:465–480.

Calvo Garzón, F. 2004. Context-free versus context-dependent constituency relations: A false dichotomy. In *Compositional Connectionism in Cognitive Science: Proceedings of the American Association for Artificial Intelligence Tech. Report FS-04-03*, ed. S. Levy and R. Gayler, 12–16. Menlo Park, CA: AAAI Press.

Calvo Garzón, F., and A. García Rodríguez. 2009. Where is cognitive science heading? *Minds and Machines* 19:301–318.

Calvo, P., and A. Gomila, eds. 2008. *Handbook of Cognitive Science: An Embodied Approach*. Amsterdam: Elsevier.

Chalmers, D. 1990. Why Fodor and Pylyshyn were wrong: The simplest refutation. In *Proceedings of the 12th Annual Conference of the Cognitive Science Society*, 340–347. Hillsdale, NJ: Erlbaum.

Chemero, A. 2009. *Radical Embodied Cognitive Science*. Cambridge, MA: MIT Press.

Chomsky, N. 1959. Review of B. F. Skinner's *Verbal Behavior*. *Language* 35:26–58.

Chomsky, N. 1965. *Aspects of the Theory of Syntax*. Cambridge, MA: MIT Press.

Clark, A. 2013. Whatever next? Predictive brains, situated agents, and the future of cognitive science. *Behavioral and Brain Sciences* 36 (3):1–73.

Cummins, R. 1996. Systematicity. *Journal of Philosophy* 93:591–614.

Eliasmith, C. 2013. *How to Build a Brain: A Neural Architecture for Biological Cognition*. New York: Oxford University Press.

Elman, J. L. 1990. Finding structure in time. *Cognitive Science* 14:179–211.

Endress, A. D., and L. L. Bonatti. 2007. Rapid learning of syllable classes from a perceptually continuous speech stream. *Cognition* 105 (2):247–299.

Ervin, S. M. 1964. Imitation and structural change in children's language. In *New Directions in the Study of Language*, ed. E. H. Lenneberg, 163–189. Cambridge, MA: MIT Press.

Evans, J. S. B. T., and K. Frankish, eds. 2009. *In Two Minds: Dual Processes and Beyond*. Oxford: Oxford University Press.

Fernando, C. 2011. Symbol manipulation and rule learning in spiking neuronal networks. *Journal of Theoretical Biology* 275:29–41.

Fodor, J. 1974. Special science, or the disunity of science as a working hypothesis. *Synthese* 28:97–115.

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

Fodor, J. 1987. *Psychosemantics*. Cambridge, MA: MIT Press.

Fodor, J., and Z. W. Pylyshyn. 1981. How direct is visual perception? Some reflections on Gibson's "Ecological Approach." *Cognition* 9:139–196.

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

Fukushima, K. 1980. Neocognitron: A self-organizing neural network model for a mechanism of pattern recognition unaffected by shift in position. *Biological Cybernetics* 36 (4):193–202.

Gomila, A. 2012. *Verbal Minds: Language and the Architecture of Cognition*. Amsterdam: Elsevier.

Gomila, A., D. Travieso, and L. Lobo. 2012. Wherein is human cognition systematic? *Minds and Machines* 22 (2):101–115.

Griffiths, T. L., N. Chater, C. Kemp, A. Perfors, and J. B. Tenenbaum. 2010. Probabilistic models of cognition: Exploring the laws of thought. *Trends in Cognitive Sciences* 14:357–364.

Grossberg, S. 1982. *Studies of Mind and Brain: Neural Principles of Learning, Perception, Development, Cognition, and Motor Control*. Dordrecht: Kluwer Academic.

Hornik, K., M. Stinchcombe, and H. White. 1989. Multilayer feedforward networks are universal approximators. *Neural Networks* 2:359–366.

Hotton, S., and J. Yoshimi. 2011. Extending dynamical systems theory to model embodied cognition. *Cognitive Science* 35:444–479.

Hummel, J. E., and I. Biederman. 1992. Dynamic binding in a neural network for shape recognition. *Psychological Review* 99:480–517.

Laakso, A., and P. Calvo. 2011. How many mechanisms are needed to analyze speech? A connectionist simulation of structural rule learning in artificial language acquisition. *Cognitive Science* 35:1243–1281.

Mac Lane, S. 2000. *Categories for the Working Mathematician*, 2nd ed. New York: Springer.

Marcus, G. F. 2001. *The Algebraic Mind: Integrating Connectionism and Cognitive Science*. Cambridge, MA: MIT Press.

Marcus, G. F., S. Vijayan, S. Bandi Rao, and P. M. Vishton. 1999. Rule learning by seven-month-old infants. *Science* 283:77–80.

Marr, D. 1982. *Vision: A Computational Investigation into the Human Representation and Processing of Visual Information*. New York: Freeman.

McClelland, J. L., M. M. Botvinick, D. C. Noelle, D. C. Plaut, T. T. Rogers, M. S. Seidenberg, and L. B. Smith. 2010. Letting structure emerge: Connectionist and dynamical systems approaches to cognition. *Trends in Cognitive Sciences* 14: 348–356.

McLaughlin, Brian P. 2009. Systematicity redux. *Synthese* 170:251–274.

Minsky, M. L., and S. A. Papert. 1969. *Perceptrons*. Cambridge, MA: MIT Press.

Newell, A. 1980. Physical symbol systems. *Cognitive Science* 4:135–183.

Newell, A. 1982. The knowledge level. *Artificial Intelligence* 18:87–127.

Newell, A. 1990. *Unified Theories of Cognition*. Cambridge, MA: Harvard University Press.

Noelle, D., and G. W. Cottrell. 1996. In search of articulated attractors. In *Proceedings of the 18th Annual Conference of the Cognitive Science Society*. Mahwah, NJ: Erlbaum.

O'Reilly, R. C., and Y. Munakata. 2000. *Computational Explorations in Cognitive Neuroscience: Understanding the Mind by Simulating the Brain*. Cambridge, MA: MIT Press.

Peña, M., L. Bonatti, M. Nespor, and J. Mehler. 2002. Signal-driven computations in speech processing. *Science* 298:604–607.

Pinker, S., and M. T. Ullman. 2002. The past and future of the past tense. *Trends in Cognitive Sciences* 6 (11):456–463.

Pinker, S., and A. Prince. 1988. On language and connectionism: Analysis of a parallel distributed processing model of language acquisition. *Cognition* 28:73–193.

Plunkett, K., and P. Juola. 1999. A connectionist model of English past tense and plural morphology. *Cognitive Science* 23 (4):463–490.

Port, R., and T. van Gelder. 1995. *Mind as Motion*. Cambridge, Mass.: MIT Press.

Ramscar, M. 2002. The role of meaning in inflection: Why the past tense does not require a rule. *Cognitive Psychology* 45 (1):45–94.

Ramsey, W. 2007. *Representation Reconsidered*. New York: Cambridge University Press.

Robbins, P., and M. Aydede. 2008. *The Cambridge Handbook of Situated Cognition*. Cambridge: Cambridge University Press.

Rosenblatt, F. 1958. The perceptron: A probabilistic model for information storage and organization in the brain. *Psychological Review* 65 (6):386–408.

Rumelhart, D. E., G. E. Hinton, and R. J. Williams. 1986. Learning internal representations by error propagation. In *Parallel Distributed Processing: Explorations in the Microstructure of Cognition*, vol. 1: *Foundations*, D. E. Rumelhart, J. L. McClelland, and the PDP Research Group, 318–362. Cambridge, MA: MIT Press.

Rumelhart, D. E., and J. L. McClelland. 1986. On learning the past tenses of English verbs. In *Parallel Distributed Processing: Explorations in the Microstructure of Cognition*, vol. 1: *Foundations*, D. E. Rumelhart, J. L. McClelland, and the PDP Research Group. Cambridge, MA: MIT Press.

Rumelhart, D. E., and J. L. McClelland, and the PDP Research Group. 1986. *Parallel Distributed Processing: Explorations in the Microstructure of Cognition*, vol. 1: *Foundations*. Cambridge, MA: MIT Press.

Ryder, D. 2004. SINBAD neurosemantics: A theory of mental representation. *Mind & Language* 19 (2):211–240.

Schmidhuber, J., F. Gers, and D. Eck. 2002. Learning nonregular languages: A comparison of simple recurrent networks and LSTM. *Neural Computation* 14 (9): 2039–2041.

Schwarz, G. 1992. Connectionism, processing, memory. *Connection Science* 4: 207–225.

Seidenberg, M. S., and J. L. Elman. 1999. Do infants learn grammar with algebra or statistics? *Science* 284:435–436.

Shea, N. 2007. Content and its vehicles in connectionist systems. *Mind and Language* 22:246–269.

Smolensky, P. 1987. The constituent structure of connectionist mental states: A reply to Fodor and Pylyshyn. *Southern Journal of Philosophy* 26:137–163.

Spencer, J. P., and E. Thelen eds. 2003. Connectionist and dynamic systems approaches to development. *Developmental Science* (Special Issue) 6:375–447.

Sporns, O. 2011. *Networks of the Brain*. Cambridge, MA: MIT Press.

Symons, J. 2007. Understanding the complexity of information processing tasks in vision. In *Philosophy and Complexity: Essays on Epistemology, Evolution, and Emergence*, ed. C. Gershenson, D. Aerts, and B. Edmonds, 300–314. Singapore: World Scientific.

Thelen, E., G. Schöner, C. Scheier, and L. Smith. 2001. The dynamics of embodiment: A field theory of infant perseverative reaching. *Behavioral and Brain Sciences* 24:1–86.

van Gelder, T. 1998. The dynamical hypothesis in cognitive science. *Behavioral and Brain Sciences* 21:615–665.