Does Classicism explain universality?

Arguments against a pure Classical component of mind

Steven Phillips
Neuroscience Research Institute, AIST

E-mail: Steven.Phillips@aist.go.jp

One of the hallmarks of human cognition is the capacity to generalize over arbitrary constituents. Recently, Marcus (1998b, 1998a) argued that this capacity, called "universal generalization" (universality), is not supported by Connectionist models. Instead, universality is best explained by Classical symbol systems, with Connectionism as its implementation. Here it is argued that universality is also a problem for Classicism in that the syntax-sensitive rules that are supposed to provide causal explanations of mental processes are either too strict, precluding possible generalizations; or too lax, providing no information as to the appropriate alternative. Consequently, universality is not explained by a Classical theory.

Key Words: Classicism, Associativism, Connectionism, systematicity, universality, structural consistency, isomorphism

Author's address: Steven Phillips, Cognitive and Behavioral Sciences Group, Neuroscience Research Institute, AIST, Tsukuba Central 2, 1-1-1 Umezono, Tsukuba, Ibaraki, 305-8568 JAPAN. Fax: +81 298 615857. Email: Steven.Phillips@aist.go.jp. WWW: http://www.etl.go.jp/~stevep.

[0]Phillips, S. (in press). Minds and Machines.

INTRODUCTION

Cognition is complex. And the multitudinous relationship between computation and behaviour makes theory construction all the more difficult. Therefore, constraints relatively loosely defined, but broadly applicable to large classes of theories are just as important as the data-hugging models of narrowly applicable experimental paradigms. For this reason, much is made of the apparent symbol-like properties of human cognition. So much so that entire theoretical frameworks, Associativism and more recently Connectionism, were ruled out, by some, for their failure to explain these properties. Instead, Fodor and Pylyshyn (1988) argued, the only viable alternative is Classicism, where human cognitive architecture *is* a symbol system. Although recent advances in Connectionist models suggest these properties are explainable without postulating symbol systems, new arguments by Marcus (1998b, 1998a) claim this not to be the case.

What follows is a closer examination of Marcus's universal generalization argument and its implication for cognitive architecture. In the first part of this paper (this and the next section), I present the Classical explanation of *systematicity* which motivates the subsequent focus on *universal generalization* (universality). In the second part of this paper, I argue that universality is also a problem for Classicism.

Some constraints on theory

One property of human cognition that has been central to the rejection of Associativism and Connectionism is systematicity.[1] Systematicity is the property whereby the set of behaviours displayable come in "related" groups. A typical example is that if one can infer that *John* is the lover of *Mary* from the statement that *John loves Mary*, then one can also make the related inference that *Mary* is the lover of *John* from the statement the *Mary loves John*. If you have the capacity for one inference, then you have the capacity for others characterized by their common structure. The point of systematicity is that whatever explains one behaviour must also explain the others, because such capacity is *indivisible*.

Classicism explains systematicity because of its insistence on syntactic representations and processes. Assuming processes that recognize *John* and *Mary* as Subject (S) or Object (O), depending on their positions; loves as a verb (V); and the triple $S\ V\ O$ as a well-formed sentence, then by application of those same processes *John loves John*, *Mary loves John*, and so on are also recognized as well-formed sentences. Syntactic compositionality entails systematicity because the *same* sequence of processes is specified for the different combinations. By contrast, Associativism generally identifies processes by the states they transform - one unique process for recognizing *John loves Mary* and another for recognizing *Mary loves John* as well-formed sentences. Since there is no dependency (identification) between these processes, there is no explanation as to why one process is always accompanied by the other; whereas in Classical models they always come together because they are the same process. Associativism is made compatible with systematicity by assuming that the two processes cohabit. Thus, systematicity is not entailed (and so not explained) by Associativism, but may be implemented by it. Similarly, the argument goes for Connectionism (Fodor & Pylyshyn, 1988; Fodor & McLaughlin, 1990)

Many (Connectionists) feel, however, that Fodor and Pylyshyn's characterization of Connectionism as a heavily mathematized version of Associativism was unnecessarily narrow. In a Connectionist model, internal atomic representations are also embedded in a metric space. When these models are specified with suitable activation and learning functions they show generalization over structurally related behaviours. So, in contrast to the atomistic version critiqued by Fodor and Pylyshyn, a (learning) process that changes the mapping of one representation implicitly changes the mapping of many others. In this way, two mapping processes *become* identified as a consequence of some more general learning process. Classical theory explains systematicity by specifying process identification. Connectionism (supposedly) explains systematicity whereby process identification is derived from another (learning) process. It appears that Connectionism provides the framework for a deeper understanding of cognition.

[1]Another property is Productivity, the capacity for unbounded behaviour. But the productivity argument is controversial, because it idealizes to infinite memory. Others (Christiansen & Chater, 1999) have argued for bounded recursion as central to a theory of lingistic competence, not just performance.

The critical question for Connectionism is the degree of generalization: How many examples does it take to acquire capacity over all structurally equivalent behaviours and is this number consistent with humans? When specified as a generalization criterion consonant with humans, general classes of Connectionist models fail to support systematicity (Hadley, 1994; Phillips, 1995, 1998, 1999, 2000). While these results rule out many standard Connectionist models, they leave open the possibility of Connectionist models that incorporate greater internal structure. Some positive results in this direction have been obtained in language processing (Hadley & Hayward, 1997).

Generally speaking, though, Connectionist models without Classical symbol systems are function approximators. Exploiting this feature, Marcus (1998b, 1998a) exposed another problem with Connectionism: its failure to explain a property of cognition he called universal generalization. Universal generalization is systematic capacity over novel constituents. If one can infer that money was stolen from the statement that The thief stole money; one can also infer that kane was stolen from the statement The thief stole kane, even though one had never seen this word before. Like systematicity, universal generalization was motivated by linguistic considerations, but regarded to apply generally. Because Connectionist models of learning are effectively interpolations over some vector space of constituent representations, they cannot be expected to support representations of constituents that lie outside that space. Effectively, the training space provides no information about the correct prediction of such points. Therefore, Marcus argued, such networks cannot support universal generalization.

However, the argument turns on the definition of novelty. At one extreme, if a constituent lies outside the space of all known inputs so that its value is undefined, then any system (including Classical[2]) will fail to support the appropriate generalization. Yet, novelty relative to the task at hand permits generalization based on similarity to internal representations derived for the same constituents on similar tasks (Boden & Niklasson, 2000). (Although, Boden and Niklasson's method of constituent similarity does not support simple analogical transfer, Phillips, 2000, which is a form of universal generalization.) Several Connectionists (e.g., McClelland & Plaut, 1999) have argued for this sort of non-Classical Connectionist explanation of universal generalization in response to claims of symbolic behaviour by seven-month-old infants on a phoneme recognition task (Marcus, Vijayan, Rao, & Vishton, 1999). The retort has been that such Connectionist solutions are, in fact, implementations of Classical systems (Marcus, 1999). Echoing Fodor and Pylyshyn, Marcus acknowledges that Connectionist models can exhibit universal generalization, but only by implementing a Classical system.

CLASSICAL EXPLANATIONS OF SYSTEMATICITY AND UNIVERSALITY

The core element of a Classical architecture is syntactic compositionality. Syntactic compositionality has two fundamental features: (1) sensitivity to constituent role/position information; and (2) capacity to operate on the binding relation between a constituent and its role. The first feature is generally understood as implying an ability to make a representational distinction between, say, John loves Mary and Mary loves John. The second feature is the basis for the Classical account of universal generalization.

Syntactic compositionality accounts for systematicity. The standard way of presenting a Classical system is in terms of a grammar or set of production rules. These are abstract specifications, identifying the before-and-after relations between cognitive states (representations), but not their implementation details. Suppose a set of Classical processes specified by the following grammar:

1. $S\ V\ O \rightarrow G$.
2. John | Mary | Sue $\rightarrow S$.
3. John | Mary | Sue $\rightarrow O$.
4. loves $\rightarrow V$.
5. V er $S\ V\ O$ ³ S.
6. V ed $S\ V\ O$ ³ O.

In this illustration, a sentence is grammatical if the input state can be mapped to the state G. Assuming the sequence 2,4,3,1 of processes for recognizing John loves Mary as grammatical, then the same sequence

[2]For example, English-only wordprocessors given Japanese text.

of processes also recognizes Mary loves John as grammatical. Hence, the capacity to recognize one combination of constituents implies the capacity to recognize the other, because they both involve the same sequence of processes. Similarly, the capacity to infer John from (who is the) lover (in) John loves Mary implies the capacity to infer Mary from (who is the) lover (in) Mary loves John, as both inferences come from the same process sequence 2,4,3,5.

This second set of capacities, involving processes 5 and 6, implicitly specifies a binding relation between representations. In these two processes, the consequent is not S (or O), but the constituent in that role, as specified by 2 (or 3). Suppose the novel constituent Rieko, which lies outside the domains of 2 and 3. Though one may not have encountered this word before, one can still infer it from (who is) loved (in) John loves Rieko, because it is understood that, whatever Rieko stands for, it plays the role of Object and the inference requires retrieving the constituent in that role.[3] Without the capacity for binding relations, a model specified by 5 and 6 simply returns S and O.

It is here that we see how universality differs from systematicity. Fodor and Pylyshyn were careful to account for systematicity in terms of the representability of constituents: "It is, for example, only insofar as 'the', 'girl', 'loves' and 'John' make the same semantic contribution to 'John loves the girl' that they make to 'the girl loves John' that understanding the one sentence implies understanding the other." (p42). But for completely novel constituents, or constituents in completely novel roles, whatever semantic content they had cannot contribute to the semantic content of its compositional host. On the contrary, it is the semantic content of the host that contributes to the semantic content of its novel constituent.

The formulation above does not capture this type of generalization, because the novel constituent lies outside the domains of 2 and 3. Of course, 2 and 3 could be extended to include Rieko, but this is both ad hoc, requiring further extension for all conceivable constituents; and counter to the requirements of universality, that it lie outside the domain of these processes. However, it can be extended by incorporating a definition of cognitive representation by Halford and Wilson (1980), based on structural consistency (i.e., isomorphism). This definition captures other properties of higher cognition (Halford, Phillips, & Wilson, submitted); is the basis for a measure of complexity of higher cognitive processes (Halford, Wilson, & Phillips, 1998); and is one of the principles behind models of analogy (e.g., Gentner, 1983; Hummel & Holyoak, 1997; Wilson, Halford, Gray, & Phillips, 2001).

Suppose two systems $\mathcal{A} = \langle A, F, R\rangle$ and $\mathcal{A}' = \langle A', F', R'\rangle$, where A and A' are sets of states; and F and F' are the state transformations, and R and R' are state relations, respectively. A function ϕ is an isomorphism from A onto A' if $\phi(f_i(a_1, \ldots, a_n)) = f'_i(\phi(a_1), \ldots, \phi(a_n))$ and $r_j(a_1, \ldots, a_n) \Leftrightarrow r'_j(\phi(a_1), \ldots, \phi(a_n))$, for all $a_1, a_2, \ldots \in A$; for all $f_i, f'_i \in F, F'$; and all $r_j, r'_j \in R, R'$. $\mathcal{A}$ has semantic content $\mathcal{A}'$, when $\mathcal{A}'$ is some part of the external environment. $\mathcal{A}'$ may also be another system of representations, in which case $\mathcal{A}$ is an analogy to $\mathcal{A}'$. In this formulation, the binding relation ϕ subsumes both processes over known constituents (e.g., 2 and 3), and unknown constituents. Thus, structural consistency subsumes both systematicity and universality.

This specification implies that unknown constituents become bound to elements in the representing system so as to maintain consistency. In the example (who is) loved (in) John loves Rieko syntactic processes specified by 1, 2 and 4 imply John and loves are aligned to S and V, respectively; and consistency checking processes specified by the structural consistency principle imply Rieko is aligned to O. Rieko is not aligned to V, because it violates consistency:

- $r(\mathsf{John}, \mathsf{loves}, \mathsf{Rieko}) \Leftrightarrow r'(\phi(\mathsf{John}), \phi(\mathsf{loves}), \phi(\mathsf{Rieko}))$
- True $\Leftrightarrow r'(S, V, V)$
- False.

The Classical explanation for universality has the same form as that for systematicity: Given processes specified by 5 and 6, the capacity to infer novel constituent Rieko implies the capacity to infer other novel constituents, because they involve the same structure and therefore the same set of consistency checking processes. However, isomorphism is a symmetric relation. So, the more ways in which a constituent can

[3]In programming language, O takes on the role of a variable.

be interpreted, the more structures that afford consistent interpretation. This flexibility is the source of the problem for Classicism.

WHY CLASSICISM FAILS TO EXPLAIN UNIVERSALITY

The approach taken is to consider two general and exhaustive classes of Classical processes: context free and context sensitive. Context free processes are influenced only by the information contained in the input sentence. Context sensitive processes are also influenced by other (historical) sources of information. I argue that the context free model either does not generate sufficient alternative structural interpretations for a given input, in which case there are some universal generalizations that cannot be supported; or in order to support all universal generalizations all possible structures are generated, in which case the Classical model provides no information about the correct structural interpretation of the input. I then argue that the context sensitive model also does not address the problem, basically because context itself requires a structural interpretation by syntactic rules, and hence the same problem applies.

Universal generalization includes apparently mundane examples like, if one has the capacity to infer Mary as the one loved in John loves Mary, then one also has the capacity to infer xyz from John loves xyz; loves from John loves loves; and so on. The Classical explanation implicitly assumes that the structure derived from the first instance containing only known constituents will be the same structure that gets derived from the other instances where some constituents are novel. But, as we have seen in the previous section, the more universal a generalizer, the more representations that become structurally consistent; and therefore the more unlikely two instances will be interpreted as instances of the same structure. One famous example is Abbot and Costello's comic routine "Who's on first?", where Abbot tries to explain to Costello the field positions of baseball players. Who's on first [base] is given as a statement (of the player's name) by Costello, and taken as the question in want of an answer by Abbot. Universal generalization taken to its extreme means that any constituent can participate in any role. This implies that any n constituent instance is potentially consistent with any n element structure. Returning to the John loves Mary example, if the John loves loves instance is to be included as a possible inference, then Classical processes must admit loves as playing a possible Object role. In this case, John loves loves also becomes structurally consistent with grammaticality rule 1. But, one would not normally consider John loves loves as grammatical. On the one hand, Classicism is supposed to explain systematic capacity over structurally "related" behaviour because it is the same set of processes in operation. But, on the other hand, universal generalization broadens the scope of these processes so that most behaviours are structurally related. The problem is that systematicity and universality are incompatible properties relative to a Classical architecture.

Since Classical models treat all structurally consistent behaviours as instances of the same class, in the most extreme sense of universal generality there can only be one (n constituent) class. But each class was supposed to pick out a cluster of mutually implied (i.e., systematically related) behavioural capacities. That is to say, for example, John gave chocolate to Mary and Mary gave chocolate to John are supposed to be clustered, but not with John gave chocolate to eat. Thus, what explains systematicity is not what explains universality.

Perhaps these two properties are incorporated by different aspects of the same Classical architecture. When syntactic rules over known constituents are not satisfied by the presence of novel constituents, the system could default to the more general consistency checking processes. In this case, the question is one of finding the right balance between two types of processes. There are, however, two problems with this proposal: a balance between the two cannot be specified ahead of time; and the decision on which type of process to apply cannot itself be Classical. These problems are revealed in the following example.

The horse raced past the barn fell fell. Clearly, this sentence is ungrammatical. Or, is it? The horse raced past the barn fell. Garden path sentences, like this one, invite a particular parse with the first few words that turns out to be inconsistent with the rest of the sentence. The first six words suggest the parse as in the complete sentence The horse raced past the barn. But the last word causes this attempt to fail, so an alternative is sort. With some effort the sentence can be read as The horse (that was) raced past the barn fell (down).[4] However, suppose the constraint that fell is a verb is relaxed by informing the reader that its also an object associated with a barn. Without this constraint the original attempt to parse the sentence is no longer blocked by the last word, permitting an interpretation along the line of The horse raced past the barn door. Since universal generalization permits the arbitrary substitution of constituent door by some other unknown constituent, the sentence can also be parsed as The horse raced past the barn fell [part of a barn]. (Coincidently, fell also means animal hide.) Continuing along this line leads to a grammatical and meaningful interpretation of The horse raced past the barn fell fell, as The horse (that was) raced past the barn fell [part] fell (down).

In the previous paragraph, a sentence was presented that most, if not all, English speaking readers would regard as ungrammatical. Yet, by the end of that paragraph, the same readers would regard the same sentence as grammatical, and meaningful. How can a sentence be both grammatical and ungrammatical within a Classical model?

The answer is that it can't, at least not without the model having the capacity to dynamically update its set of processes. A Classical model consists of a set of syntax-sensitive rules that map between compositional representations. In the case of grammaticality it is a classification of sentences into one of two classes, with possibly a structural interpretation. Thus, for any Classical model a sentence is either grammatical, or ungrammatical, but it cannot be both. If the set of rules is such as to yield the first sentence as ungrammatical, then of itself, it cannot explain the grammatical interpretation. If the set of rules accepts the first sentence as grammatical, it cannot explain why English speakers initially regard it as ungrammatical. A fixed set of rules cannot classify sentences as both grammatical and ungrammatical, at the same time.

But even with the capacity to update its own rules, the processes responsible for the change in grammaticality cannot be purely Classical. The critical step from ungrammatical to grammatical came with the relaxation of the constraint that fell must be a verb. But the manipulation of this constraint did not involve syntactic information, and therefore it lies outside the scope of Classical theory. Recall that the core construct of Classical architecture is syntactic compositionality, whereby "... the semantic content of a [complex] representation is a function of the semantic contents of its syntactic [position critical] parts, together with its constituent structure." (Fodor & Pylyshyn, 1988, p12). The re-interpretation of fell involved a contingent fact that it be regarded as a reference to a part of a barn in the current context (that I just said so). Since this fact is arbitrary and independent of the source sentence, it cannot be accommodated by a set of context free rules. Therefore, Classical context free processes cannot support universal generalization.

The alternative is to consider Classical processes that also apply syntactic rules to context. In this example, suppose there is a meta-process that takes contextual information such as treating fell as a noun and modifies its syntactic rules accordingly. But, the arguments that apply to processes also apply to meta-processes. If the meta-process is Classical, then there must be some syntactic compositions that unambiguously index the relevant structures, from which the process of re-interpretation could be bootstrapped. A composition like Treat fell as a noun should be a candidate for the current example. But, this example is not unique either. One can play the same game for this composition as was played in the horse fell fell example. That is, we can introduce another structure and show via analogy that it is alignable to this composition. Examples include, Treat fell (over) like a noun (falls over); and Treat (the person named) 'fell as a noun' (to dinner). As unnatural as these examples may seem, they are still interpretable as such. The point is that this process of constructing universal generalizations requires fixating on a particular structure, despite transformational rules to the contrary. If Treat fell as a noun necessitates only one interpretation, then it prohibits other universal generalizations. If these generalizations are to be accepted, then something other than just a Classical process is needed to keep the corresponding structure in the focus of computation. Consequently, meta-processes cannot be entirely Classical, either.

In general, to support universality is to admit any constituent as potentially interpretable as any structural element. Structural information contained within a representation divorced from its context

[4] See Pinker (1994) for discussion on this and other garden path sentences.

either underdetermines or overdetermines the target structure. The only alternative for Classicism is to suppose that prior context provides the additional structural information to identify the target structure. But, contextual information is not sufficient either, because structural consistency (isomorphism) is both a symmetric and a transitive relation. Symmetric means that if a system (of representations) A is isomorphic to B, then B is isomorphic to A. Transitive means that if A is isomorphic to B, and B is isomorphic to another system C, then A is also isomorphic to C. To see why Classicism cannot rely upon contextual information, suppose that A and B are two alternative, structurally consistent interpretations of some system S, denoted as $I(A, S)$ and $I(B, S)$, respectively; and that C is some relevant context distinguishing these two alternatives by virtue of being isomorphic to one, but not the other. Without loss of generality, assume A is isomorphic to C, denoted as $I(A, C)$. Since I is symmetric, $I(A, S)$ implies $I(S, A)$. Since I is transitive, $I(B, S)$ and $I(S, A)$ implies $I(B, A)$. Now, assuming $I(A, C)$, then $I(B, A)$ and $I(A, C)$ implies $I(B, C)$. Therefore, B is also structurally consistent with context C, and contrary to the original assumption, context C cannot be used to distinguish between A or B by virtue of structural consistency.

DISCUSSION

The problem posed by universality is that the relevant facts can come from anywhere. It is not possible to define a set of syntax sensitive rules to cover all possible generalizations. Either, the set of rules is restricted and fails to capture some generalizations, or is completely unrestricted so that every rule can be applied to every sentence, in which case Classical processes convey no information about the target structures.

Universality is not the problem of simply disambiguating multiple alternative parses. The standard response for preserving Classicism is to say that it is the role of the Classical module to generate all structurally consistent alternatives from which a non-Classical (possibly Connectionist) module filters out those that are contextually irrelevant. But this reproach presumes that the target structure will be in the set of structures generated by the Classical module. A syntax that precludes The a are of I as a grammatical fragment prohibits any possible meaningful interpretation.[5] And, as already argued, the only way to safeguard all possible generalizations is to generate all possible structures regardless of the input, in which case the Classical module is irrelevant. Thus, a pure Classical component cannot support universal generalization. If Classical processes are operating at all, then they must be in conjunction with non-Classical ones.

It may seem counterintuitive that Classicism fails to explain universality. After all, the strength of programming languages is their capacity to realize unbounded behaviour from a single specification through the use of variables for binding arbitrary constituents. But this intuition belies an important assumption that the relevant structure can be made the focus of computation through Classical means. Suppose the example used in Marcus (1998b) "if glork then frum, and glork, infer flum". Since the inference works for any combination of antecedent and consequent elements, one can also do "if if then then, and if, infer then". However, this second example is not even parsable without the assumption that the target structure be taken from the previous instance. That is, without a process that applies to the most recently instantiated structure. But use of recency information is not a Classical operation since it applies to any representation regardless of its structure.

The arguments presented here parallel more general arguments put forth by Fodor (2000) against Classicism. Fodor argued that cognition cannot be Classical because Classical architectures do not support global properties (i.e., properties that are a function of the entire state of a system). The summary elements of Fodor's argument are: 1. Classical processes operate over representations whose syntactic relations are either entirely local (between its constituents only); or additionally global (between other representations). 2a. Locally sensitive processes cannot address behavioural properties (e.g., simplicity, conservatism) that are global. 2b. Globally sensitive processes must be sensitive, in general, to the entire store of representations, a unit of computation which is simply too great to be tractable. 3. Consequently, cognition cannot be entirely Classical. Locally and globally sensitive processes correspond (respectively)

[5] See Abney (1996) for this example, where an are is a parcel of land, a is its modifier, and I is the head noun.

to context free and context sensitive processes discussed here. Context free processes fail to find target structures that rely on contextual information. Context sensitive processes rely on generating all possible structures, which is simply too many to be informative.

If neither Associativism nor Classicism accounts for the various properties of cognition, what then are the alternatives? According to Fodor (2000) there are none available. Associativism is ruled out on the grounds that it doesn't support systematicity; Classicism is ruled out because it can't support globality. The state of affairs was euphemistically summed up as *Fodor's First Law of the Nonexistence of Cognitive Science*, which says that "the more global [informationally unencapsulated, or universally informed] a cognitive process is, the less anybody understands it" (Fodor, 1983, p107). The problem that still remains for Classical theory, in general, is that the context sensitivity of thought cannot be supported by a system of representations that is essentially context independent. Appealing to heuristic, or higher order processes does not help because if a decision process requires global information, then the decision on how to make that decision *also* requires global information (Fodor, 2000). Along similar lines, one might hope for some sort of integration of Classical and non-Classical theory. But, if the solution to universality is to integrate Classical and non-Classical architectures, then the problem becomes one of providing causal explanations for how their respective processes are mediated.

CONCLUSION

Pure Classical models cannot account for universality. Therefore, contrary to Marcus (1998b), Fodor and Pylyshyn (1988), and others, the relationship between non-Classical (e.g., Connectionist) and Classical models is not one of implementation. If Classicism does play a role in universal generalizations, it is not as an isolatable chain in a series of processing steps. Conversely, if there are domains of cognitive behaviour that are governed solely by Classical processes, then universality is not a property of those domains. Universal generalization was supposed to be a *prima facie* example that cognition is a symbol system. Instead, it's the very property that shows that it isn't.

REFERENCES

Abney, S. (1996). Statistical methods and linguistics. In J. L. Klavans & P. Resnik (Eds.), *The balancing act: Combining symbolic and statistical approaches to language* (p. 1-26). Cambridge, MA: MIT Press.

Boden, M., & Niklasson, L. (2000). Semantic systematicity and context in connectionist networks. *Connection Science*, *12*(2), 1-31.

Christiansen, M. H., & Chater, N. (1999). Toward a connectionist model of recursion in human linguistic performance. *Cognitive Science*, *23*, 157-205.

Fodor, J. A. (1983). *The modularity of mind: An essay on faculty psychology.* Cambridge, MA: MIT Press.

Fodor, J. A. (2000). *The mind doesn't work that way: The scope and limits of cognitive psychology.* Cambridge, MA: MIT Press.

Fodor, J. A., & McLaughlin, B. P. (1990). Connectionism and the problem of systematicity: Why Smolensky's solution doesn't work. *Cognition*, *35*, 183-204.

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

Gentner, D. (1983). Structure-mapping: A theoretical framework for analogy. *Cognitive Science*, *7*, 47-59.

Hadley, R. F. (1994). Systematicity in connectionist language learning. *Mind and Language*, *9*(3), 247-272.

Hadley, R. F., & Hayward, M. B. (1997). Strong semantic systematicity from Hebbian connectionist learning. *Minds and Machines*, *7*, 1-37.

Halford, G. S., Phillips, S., & Wilson, W. H. (submitted). Structural complexity in non-symbolic and symbolic cognitive processes: The concept of representational rank. .

Halford, G. S., & Wilson, W. H. (1980). A category theory approach to cognitive development. *Cognitive Psychology*, *12*, 356-411.

Halford, G. S., Wilson, W. H., & Phillips, S. (1998). Processing capacity defined by relational complexity: Implications for comparative, developmental, and cognitive psychology. *Behavioral and Brain Sciences*, *21*(6), 803-831.

Hummel, J. E., & Holyoak, K. J. (1997). Distributed representations of structure: A theory of analogical access and mapping. *Psychological Review*, *104*(3), 427-466.

Marcus, G. F. (1998a). Can connectionism save constructivism? *Cognition*, *66*, 153-182.

Marcus, G. F. (1998b). Rethinking eliminative connectionism. *Cognitive Psychology*, *37*, 243-282.

Marcus, G. F. (1999). Connectionism: with or without rules? response to J.L.McClelland and D.C.Plaut (1999). *Trends in Cognitive Science*, *3*(5), 168-170.

Marcus, G. F., Vijayan, S., Rao, S. B., & Vishton, P. M. (1999). Rule learning by seven-month-old infants. *Science*, *283*, 77-80.

McClelland, J. L., & Plaut, D. C. (1999). Does generalization in infant learning implicate abstract algebra-like rules? *Trends in Cognitive Science*, *3*(5), 166-168.

Phillips, S. (1995). *Connectionism and the problem of systematicity.* (PhD Thesis. Department of Computer Science, University of Queensland, Brisbane, Australia.)

Phillips, S. (1998). Are feedforward and recurrent networks systematic? Analysis and implications for a connectionist cognitive architecture. *Connection Science*, *10*(2), 137-160.

Phillips, S. (1999). Systematic minds, unsystematic models: Learning transfer in humans and networks. *Minds and Machines*, *9*(3), 383-398.

Phillips, S. (2000). Constituent similarity and systematicity: The limits of first-order connectionism. *Connection Science*, *12*(1), 1-19.

Pinker, S. (1994). *The language instinct: How the mind creates language.* New York, NY: HarperPerennial.

Wilson, W. H., Halford, G. S., Gray, B., & Phillips, S. (2001). The STAR-2 model for mapping hierarchically structured analogs. In D. Gentner, K. J. Holyoak, & B. Kokinov (Eds.), *The analogical mind: Perspectives from cognitive science* (p. 125-159). Cambridge, MA: MIT Press.