

UC Merced

Proceedings of the Annual Meeting of the Cognitive Science Society

Title

Connectionism, Systematicity, and Nomic Necessity

Permalink

<https://escholarship.org/uc/item/5kv2p52d>

Journal

Proceedings of the Annual Meeting of the Cognitive Science Society, 18(0)

Author

Hadley, Robert F.

Publication Date

1996

Peer reviewed

Connectionism, Systematicity and Nomic Necessity

Robert F. Hadley

School of Computing Science
and Cognitive Science Program
Simon Fraser University
Burnaby, B.C., V5A 1S6
hadley@cs.sfu.ca

Abstract

In their provocative 1988 paper, Fodor and Pylyshyn issued a formidable challenge to connectionists, viz., to provide a *non-classical* explanation of the empirical phenomenon of *systematicity* in cognitive agents. Since the appearance of F&P's challenge, a number of connectionist systems have emerged which *prima facie* meet this challenge. However, Fodor and McLaughlin (1990) advance an argument, based upon a *general* principle of *nomological necessity*, to show that one of these systems (Smolensky's) could not satisfy the Fodor-Pylyshyn challenge. Yet, if Fodor and McLaughlin's analysis is correct, it is doubtful whether any existing connectionist system would fare better than Smolensky's. In the view of Fodor and McLaughlin, humans and classical architectures display systematicity as a matter of nomological necessity (necessity by virtue of natural law), but connectionist architectures do not. However, I argue that the Fodor-Pylyshyn-McLaughlin appeal to nomological necessity is untenable. There is a sense in which neither classical nor connectionist architectures possess nomological (or 'nomic') necessity. However, the sense in which classical architectures *do* possess nomic necessity applies equally well to at least *some* connectionist architectures. Representational constituents can have causal efficacy within both classical and connectionist architectures.

1. Introduction

In their provocative 1988 paper, Fodor and Pylyshyn issued a formidable challenge to connectionists, viz., to provide a non-classical explanation of the empirical phenomenon of *systematicity* in cognitive agents. Fodor and Pylyshyn (F&P) acknowledge that connectionism might provide an implementational foundation for a *classical* explanation of systematicity, but *that*, they observe, would not provide an alternative to the classical account.

Although the precise definition of systematicity is a matter of some dispute (see below), we may, for the moment, ignore subtleties and assume that 'systematicity' refers to the fact that cognitive capacities are systematically related, and come in 'clumps'. As F&P insist, 'you don't find people who can *think the thought* that John loves the girl but can't think the thought that the girl loves John'. F&P maintain that systematicity occurs not only in thought, but in language understanding and production.

Since the appearance of F&P's challenge, a number of connectionist systems have emerged which *prima facie* meet this challenge (cf. Chalmers, 1990; Elman, 1990;

Smolensky, 1990; St. John and McClelland, 1990; Niklasson and van Gelder, 1994). However, Fodor and McLaughlin (1990) advance an argument, based upon a *general* principle of *nomological necessity*, to show that Smolensky's methods in particular could not satisfy F&P's challenge. If Fodor and McLaughlin's (F&Mc) analysis is correct, it is doubtful whether any of the connectionist systems just cited would fare better than Smolensky's. In fact, McLaughlin later argues (1993a) that neither the results of Chalmers nor Elman constitute counterexamples to the F&P position (contrary to the explicit claims of Chalmers, 1990).

Now, the crux of F&Mc's argument, and indeed of F&P's original thesis, lies in their appeal to nomological necessity. In their view, humans and classical architectures display systematicity as a matter of nomological necessity (necessity by virtue of natural law), but connectionist architectures do not. In what follows, I argue that the Fodor-Pylyshyn-McLaughlin appeal to nomological necessity is untenable. There is a sense in which neither classical nor connectionist architectures possess nomological (or 'nomic') necessity. However, the sense in which classical architectures *do* possess nomic necessity applies equally well to at least *some* connectionist architectures. Moreover, Fodor, Pylyshyn, and McLaughlin all stress the *causal* efficacy of constituents (atomic elements) in complex *classical* representations. They maintain that complex connectionist representations (such as Smolensky's tensor-product representations) lack atomic constituents, and so, they believe that connectionists cannot appeal to constituent structure to *explain* why systematicity should be nomically necessary. By contrast, I argue (in section 4) that F&Mc ignore the manner in which atomic constituents can causally determine mathematical properties of complex connectionist representations. Yet, those very properties can engender systematicity within the context of some *particular* connectionist architecture. The F&Mc stance seems to arise from the fact that they regard particular connectionist systems as nomically arbitrary, while they regard classical architectures as nomically necessary.

2. Systematicity.

Before proceeding, it would be well to have in mind some more definite notion of systematicity than described above. However, as van Gelder and Niklasson (1994) have observed, F&P do not offer a precise definition of systematicity. Instead, they provide examples to support their

contention that certain important cognitive capacities are systematically related. Most of their examples follow the pattern described above; the capacity to think (or understand) aRb is systematically related to the capacity to think (or understand) bRa , where a and b are referential terms, and R is some relation, e.g., 'loves'. However, F&P also include *systematicity of inference* in their discussion of systematicity. They contend, for example, that 'it's a psychological law that thoughts that P & Q tend to cause thoughts that P and thoughts that Q , all else being equal'. F&P's examples of systematicity in inference all involve rather immediate inferences, and it is unclear whether they would agree, for example, that thoughts of the form $P \rightarrow Q$ tend to cause thoughts of the form $\neg Q \rightarrow \neg P$. In any case, van Gelder and Niklasson have taken F&P to task on the question (a) whether humans do in fact exhibit systematicity in inference, and (b) whether F&P have produced a workably clear conception of systematicity. In light of problems raised by van Gelder and Niklasson, and because the issues which here concern us do not obviously involve systematicity in inference, and I shall not address this aspect of F&P's discussion. It is noteworthy that both Fodor and McLaughlin (1990) and McLaughlin (1993a) focus almost entirely upon the aRb , bRa class of mental capacities when presenting their case regarding nomic necessity.

Once we set aside concerns about systematicity in inference, other questions arise. For example, it is not clear whether F&P intend their observations on systematicity to apply to children in the early stages of language learning. At one point, they remark that *grammatically competent* (my emphasis) humans display systematicity in language use, but in a footnote citing Pinker (1984) they seem to suggest that children never lack the combinatorial architecture which engenders systematicity. Given Fodor's repeated stress on the nomic necessity of systematicity, it would seem incumbent upon him, at least, to present a more careful examination of the empirical evidence relevant to these claims. For there is substantial evidence that children do *not* exhibit strong forms of systematicity in very early stages of language learning (see Ingram, 1985).

Now, in Hadley (1992, 1994a) a learning-based hierarchy of *degrees* of systematicity is defined, where levels are distinguished according to the degree of *novelty* of sentences which the learning agent can interpret. Ingram has informed me (personal communication) that children do indeed pass through stages of systematicity which correspond fairly closely to the levels distinguished in those publications. Having noted this, I propose for the present to ignore these complications, and to focus upon F&P's approach to systematicity, since *that* is the basis of F&Mc's claims regarding nomic necessity.

The question we should now consider is whether F&P's conception of systematicity, *vis-a-vis* the ability to represent and process objects of the form aRb and bRa , can be clearly defined. McLaughlin (1993a) provides an analysis of systematicity, which, I believe, provides an affirmative answer. For brevity, I offer the following gloss on McLaughlin's analysis. In doing so, I omit certain details of McLaughlin's presentation which seem to me to go beyond F&P's original conception.

- I. Systematicity: A cognitive agent, C , exhibits systematicity just in case its cognitive architecture *causally ensures* that C has the capacity to be in a propositional attitude (A) towards proposition (or sentence) aRb if and only if C has the capacity to be in attitude (A) towards proposition (sentence) bRa .

McLaughlin's examples of propositional attitudes (A) are 'thinks', 'believes', 'prefers that', and 'understands'. By contrast, I propose to restrict the scope of (A) to 'thinks' and 'understands', because these are the attitudes that F&P employ in their examples. Also, to extend (A) to include 'believes', and 'prefers that' would seem to raise unwanted complications. For example, it's far from clear that agents who have the capacity to believe, say, that *Kim sees a tree* also have the capacity to believe that *a tree sees Kim*. Of course, even when we replace 'believes' with 'understands', difficulties may arise. However, it is arguable that anyone who understands *Kim sees the tree* can make *sufficient sense* of *A tree sees Kim* to realize that it describes a factually impossible situation. Also, one can still understand the individual words in the latter sentence and imagine a cartoon that would illustrate the factually problematic situation.

3. Nomic Necessity.

A crucial aspect of F&P's view of systematicity, which is expressed in (I), is the idea that an integrated underlying mechanism *explains* the fact that certain cognitive capacities 'come in pairs'. I have followed McLaughlin (1993a) in building into the *definition* of systematicity a presupposition about the causal genesis of systematically related capacities. In Fodor and McLaughlin (F&Mc, 1990), this presupposition is not a matter of definition, but is described as a separate assumption. However, taken either way, the assumption is crucial to the arguments presented in F&Mc (1990) and McLaughlin (1993a). In F&Mc, the assumption is phrased in terms of nomic necessity (natural law).

The Fodor-McLaughlin stance, which re-emerges in McLaughlin (1993a), is that within *humans* and classical architectures it is a matter of nomic necessity that cognitive capacities are systematically related. By contrast, they hold that in connectionist architectures it is *at best* an accident of nature (i.e., nomically arbitrary) if cognitive capacities are so related. I shall argue, however, that neither F&Mc nor McLaughlin (1993a) succeed in establishing a relevant asymmetry between classical and connectionist architectures. Both architectures can, in principle, exhibit the relevant causal powers, and both architectures depend upon processing mechanisms which (since they might have arisen through evolution) may be viewed as nomically arbitrary to the same degree.

To begin, let us consider *why* F&Mc regard systematicity as a nomically necessary consequence of classical architecture. In (1990) they say:

Whereas, as we keep on saying, in the Classical architecture, if you meet the conditions for being able to represent aRb , YOU CANNOT BUT MEET THE CONDITIONS FOR BEING ABLE TO REPRESENT bRa ; the architecture won't let you do so because (i) the representation of a , R and b are constituents of the repre-

sentation of aRb , and (ii) you have to token the constituents of the representations that you token, so Classical constituents can't be just imaginary. So then, it is *built into* the Classical picture that you can't think aRb unless you are able to think bRa , but the Connectionist picture is *neutral* on whether you can think aRb even if you can't think bRa . But it is a law of nature that you can't think aRb if you can't think bRa . So the Classical picture explains systematicity and the Connectionist picture doesn't. So the Classical picture wins.

Now, this passage contains at least two tendentious claims. These are: (a) 'But it is a law of nature that you can't think aRb if you can't think bRa ', and (b) 'it is *built into* the Classical picture that you can't think aRb unless you are able to think bRa .'

Regarding (a), it is not clear whether F&Mc take this 'law of nature' to be a *fundamental* psychological law about cognitive agents, or to be a derived law of some kind. In either case, (a) is likely to be contested by anyone who believes (i) that language acquisition is a precondition for the capacity to think, or (ii) that children exhibit only partial systematicity in the early stages of language learning. Both (i) and (ii) raise complex and interesting issues, but I propose to focus instead upon what seems to me the more central issue, namely (b). Is it clearly true that, on the Classical picture, 'you can't think aRb unless you are able to think bRa '? Granted, in a classical symbol system, you cannot token aRb without tokening the constituents, 'a', 'R', and 'b'. But, how is this fact supposed to *causally ensure* that you can think bRa . F&Mc do not explain in any detail, but they occasionally allude to structure-sensitive processes, which are a topic of emphasis in F&P, 1988.

Now, Smolensky (1994) maintains that classical theorists have no explanation of how the capacity to think aRb is supposed to *entail* the capacity to think bRa . He says,

The point here is that *systematicity* is a basic part of the definition of a symbol *system*; the Classical theory gets systematicity by *assuming* it, not by deriving it from more fundamental principles. The necessity of systematicity is not *explained* by the classical theory in any sense, it is simply *described* by it. (original emphasis)

Later, Smolensky says that systematicity 'is achieved by stipulation in the Classical theory, which, on principle, refuses to posit lower-level principles from which systematicity might be derived.' Given the incompleteness of Fodor's, Pylyshyn's and McLaughlin's remarks about how, on the Classical theory, systematicity is supposed to arise, it is not surprising that Smolensky should conclude that systematicity 'is achieved by stipulation in the Classical theory'. However, I believe that Smolensky's analysis is not accurate. For, given F&P's repeated reference to structure-sensitive processes, and tokening of constituents, the outline of a theory can be discerned. In McLaughlin (1993a, p. 170), the theory is succinctly stated, but only in the most general terms. McLaughlin asserts that, 'classical architectures have complex symbols

and algorithmic operations on the constituent structures of such symbols. Even though the symbols do not participate in the algorithmic operations in virtue of their semantic properties, the algorithms will be such that symbol transitions make sense given the meanings of the symbols'.

Now, although details are *not* provided by F&P or McLaughlin about the nature of the relevant algorithmic operations (i.e., the structure-sensitive processes), it is not difficult to imagine the *kind* of algorithms that *might* serve the purpose. The fields of natural language 'comprehension' and generation (in AI) are rife with examples of algorithms that possess the kind of combinatorial properties that could, in principle, explain the *syntactic* requirements of systematicity. To be sure, these algorithms have only been applied to comparatively small languages, and their adequacy for the totality of natural language and mental representation remains uncertain. Nevertheless, it is reasonable to suppose that Fodor, Pylyshyn, and McLaughlin may have had algorithms of this *general type* in mind when they refer to 'algorithmic operations' and 'structure-sensitive' processes. The fact that these authors wish to remain neutral on the details of such algorithms, and upon how such algorithms might be realized in the brain, merely reflects good judgment on their part. It does not, I think, reflect a refusal 'to posit lower-level principles from which systematicity might be derived', as Smolensky suggests. The general nature of the lower-level principles is not hard to surmise. Only the details of such principles remains unspecified.

However, the crucial point I wish to underscore here is that, on the classical account, the systematicity of representations arises *only* in the presence of *assumed* algorithmic processes. Whatever 'causal powers' classical constituents may possess, and whatever nomological necessity may reside in some classical architecture, these exist only against a *background* of algorithms. It follows, then, that when the nomological characteristics of a connectionist architecture are considered, we must permit the connectionist to assume that *correspondingly general* processing mechanisms are in place. (Within connectionism, sets of weighted links between nodes, activation functions, and firing thresholds can all enter into these processing mechanisms.) Yet, F&Mc, and later McLaughlin (1993a) seem unwilling to allow Smolensky the connectionist mechanisms which would permit a network to process his tensor-product representations (activation vectors, for present purposes) in a manner that would engender systematic relations between those representations. For example, F&Mc say,

No doubt it is possible for Smolensky to wire a network so that it supports a vector that represents aRb if and only if it supports a vector that represents bRa ; and perhaps it is possible for him to do that without making the imaginary units explicit (though there is, so far, no proposal about how to ensure this for *arbitrary* a , R and b). The trouble is that, although the architecture permits this, it equally permits Smolensky to wire a network so that it supports a vector that represents aRb if and only if it supports a vector that represents zSq ; or, for that matter, if and only if it supports a vector that repre-

sents *The Last of The Mohicans*. The architecture would appear to be absolutely indifferent as among these options.

F&Mc seem willing, for the sake of argument, to grant that Smolensky *could* 'wire a network so that it supports a vector that represents aRb if and only if it supports a vector that represents bRa ', for arbitrary a , R , and b . Moreover, results reported in Chalmers (1990), Phillips (1994), and Miyata, Smolensky, Legendre (1993) establish that such *general purpose*, holistic wiring is not merely a theoretical possibility, but is now a reality.¹ Given this, it seems odd that F&Mc should insist upon the *arbitrariness* of the wiring (weighted links) that Smolensky would need to employ. For, it is difficult to see how the algorithmic operations presupposed by classical theorists are any *less* arbitrary than the general purpose processing mechanisms (or *unified* weighted link structures) that Smolensky, Chalmers, and other connectionists have employed. Presumably, both the classical and the connectionist processing mechanisms could (for all Fodor or McLaughlin say to the contrary) be the *products of natural selection*. Moreover, once either of these kinds of mechanisms is integrated within its respective architecture, the causal effects produced by representations via those mechanisms would be nomically necessary in precisely the same sense. The representations (classical or connectionist), together with the processing mechanisms, cause the entire system to behave in a systematic fashion. So, unless F&Mc can provide some reason, beyond natural selection (and beyond verbal stipulation) for saying that classical architectures *must* have the algorithmic mechanisms that they do in fact have, it is difficult to see that they have uncovered any *relevant* asymmetry between classical and connectionist architectures.

Now, F&Mc might rejoin that two relevant asymmetries can yet be found. These are: (1) that the actual physical presence of classical constituents in complex representations makes it possible to *explain* how systematicity arises within classical architectures, whereas no analogous explanation is possible within connectionist architectures; (2) that distributed activation patterns that represent complex meanings in a connectionist system are only *arbitrarily* related to the representations of their constituents, whereas just the opposite is true in classical architectures. Let us consider these two possible rejoinders in turn.

Concerning (1), it may well be true that *classical* constituent structure leads to transparent, obvious explanations of systematicity. The fact that classical algorithmic operations can (among other things) simply recombine atomic constituents in different orderings yields an elegant explanation of systematicity. However, some connectionists (e.g., van Gelder and Niklasson, 1994) have argued that such classical explanations are simplistic and fail to account for the subtle departures from straightforward systematicity that humans exhibit. In any case, the most elegant and obvious explanation of a natural phenomenon may not be the most accurate, and connectionists have argued that it is indeed possible to explain systematicity within connectionist architectures (cf. Chalmers, 1990;

Smolensky, 1994; Niklasson and van Gelder, 1994). At present, the connectionist account may be very abstract and schematic, but, as we have seen, the same is presently true of the classical explanation. As McLaughlin (1993a) observes, 'There are a number of promissory notes in classicism's proposal for explaining systematicity'.

Within connectionism, the explanation appeals to two fundamental truths: (a) it is possible to *generate* distributed representations in such a way that striking *mathematical regularities* exist between those distributed patterns that we would want to describe as systematically related (e.g., the patterns for 'dogs chase cats' and 'cats chase dogs'); (b) it is possible to create holistic (or unified) weight vectors (sets of weighted links) which exploit those mathematical regularities in a way that *causally ensures* that aRb is representable if and only if bRa is representable. Clear verification of both (a) and (b) can be found in Chalmers (1990), Phillips (1994), and Miyata, Smolensky, Legendre (1993). In addition, cluster analyses of distributed representations reported by Elman (1991) and others dramatically underscore the truth of (a).²

I will not attempt to recapitulate arguments already provided by Chalmers, Niklasson and van Gelder, Phillips, and Smolensky. However, the truth of (a) and (b) is implicit in the fact that a single unified weight vector is capable of *transforming* distributed representations of *novel* sentences (both active and passive voice) in the same fashion that it transforms representations of syntactically isomorphic, previously encountered sentences. This point emerges clearly from Chalmers' discussion of his active-passive transformation network. McLaughlin (1993a) disputes Chalmers' claim to systematicity, but he does so, in my view, by raising the hurdle. McLaughlin does *not* dispute the fact that Chalmers' network is capable, in effect, of re-ordering constituents. Rather, his complaint centers upon the fact that Chalmers has not dealt with certain subtle aspects of the active-passive transformation (such as the fact that verbs acquire past tense endings during passivization). However, F&P's examples of systematicity do not involve such subtleties; they focus upon the reordering and decomposition of representations. It is the reordering problem that Chalmers addressed. Whether

²In a 1991 paper, Elman describes cluster analyses of distributed activation patterns which appear on the hidden layer of the recursive backpropagation networks he employs. These patterns are quite complex, and contain information not only about the most recent input word, but about the syntactic context in which that word occurs. Elman's analyses clearly reveal that distributed patterns, which represent a word in one syntactic context, cluster very close, in vector space, to patterns representing the same word in *other* syntactic contexts. Moreover, Elman discovered that *not only* do nouns form separate clusters from verbs, but subclasses can be discovered within these clusters, corresponding to animate nouns, transitive verbs, etc. This clustering establishes beyond doubt that significant mathematical regularities can exist in distributed connectionist representations. In the case of Smolensky's tensor-product representations, which are *defined by* mathematical formulae, there can be no doubt whatsoever that important mathematical regularities exist. These regularities are exploited both in the active-passive transformations reported in Miyata, Smolensky, Legendre (1993) and in the systematicity results of Phillips (1994).

¹I am assuming that 'a', 'R', and 'b', though arbitrary, are not entirely novel to the agent.

connectionism can explain *all* linguistic phenomena, including language acquisition, remains unclear, but that is not the point at issue. It is true that McLaughlin also objects to Chalmers' use of *training* to produce an association between active and passive representations. However, even ignoring the empirical likelihood that systematicity is partially a result of learning, Chalmers could very well reply that *certain* weight vectors, which are produced by training in artificial networks, might be innately wired in creatures which have evolved naturally.

Let us now return to the second possible rejoinder (2), introduced above. The rejoinder was: 'that distributed activation patterns that represent complex meanings in a connectionist system are only *arbitrarily* related to the representations of their constituents, whereas just the opposite is true in classical architectures'

Now, a connectionist might respond to this claim in a number of ways. One approach would be to challenge its relevance. For, as long as distributed connectionist representations can be transformed, via weighted link structures, into *systematically related* new representations, what does it matter whether activation patterns for complex representations are arbitrarily related to representations of constituents? Systematicity and decomposability must be present, but the shapes of the representations involved are not important as long as transformations can be achieved in a holistic (non-piecemeal) fashion.

Perhaps a deeper reply to (2) would be to note that *within the context of the processing mechanisms that could generate* distributed representations of the type employed by Chalmers (1990) and Pollack (1990), complex representations are *not* arbitrarily related to representations of their constituents. For, by means of Recursive Auto-Associative Memories (RAAM), one can generate complex distributed representations that are *functionally related* in a systematic fashion to representations of their constituents (cf. van Gelder, 1990). Indeed, it is for this very reason that complex distributed representations exhibit those mathematical regularities which enable systematic transformations of the activation patterns to occur; witness the transformations described by Chalmers (1990).

The classicist might now object that, although distributed representations can be generated in a systematic fashion, the connectionist networks which *generate* those representations are themselves *nominally arbitrary*. However, the same can be said of classical algorithms which generate complex phrases from atomic constituents. The elegance and simplicity of the rules of formation for well-formed formulae in predicate logic, say, can scarcely be taken as evidence that *natural law* forces our brains to employ these rules. The only argument so far advanced by classicists, to support the claim that our brains (nominally) *must* employ classical rules of formation, is that no other explanation of systematicity is possible. But that, we have seen, is an argument unsupported by the facts.

4. The Causal Power of Constituents.

Both F&P and F&Mc make much of the causal power of constituents in classical representations. In their view, it is because classical constituents are physically present in complex constituents that they can have causal efficacy. Since constituents of distributed representations (including those of Smolensky, Elman, and Chalmers) are *not*

physically present in those representations, they cannot have causal powers (or so it is argued). For this reason, F&Mc have concluded that any connectionist system which does *not* implement a classical architecture cannot create systematicity as a matter of nomic necessity.

In the preceding section, I argued that F&Mc fail to establish any *relevant* asymmetry between classical and connectionist architectures *vis-a-vis* nomic necessity. Both kinds of architectures can achieve such necessity only in the presence of background algorithms or processing mechanisms. In the present section I pursue the Fodor-Pylyshyn-McLaughlin claim that constituents can have causal powers only if they are physically present in complex representations. The issue seems worthy of pursuit for a couple of reasons. First, as suggested above, it appears that confusion on this issue has increased confusion about the alleged lack of nomic necessity in connectionist systems. Second, I believe that Smolensky concedes too much when he says (1988), 'It may be that a good way to characterize the difference [between classical and connectionist architectures] is in terms of whether the constituents in a mental structure are causally efficacious in mental processing'. Although Smolensky's representations can be defined in terms of equations involving the identity of atomic constituents in essential ways, and though he has shown that his composite (tensor-product) representations are *decomposable* into their constituents under certain assumptions, Smolensky views the presence of atomic constituents in the composite representations as *merely imaginary*. F&Mc, of course, welcome this description, for what is merely imaginary cannot have causal efficacy (or so it would seem).

However, consider for a moment the phenomenon of *implicit information*. I have argued elsewhere (Hadley, 1995) that information which is recoverable from complex structures, even when background mechanisms must be assumed, may be regarded as *implicit* in a special sense. One could argue further that information which is implicit in this sense is not *merely* imaginary, since the complex structures in question must possess specific properties that reflect the derivable information. Analogously, Smolensky's tensor-product representations possess special properties that reveal the identity of their (purportedly) 'imaginary' atomic constituents. Thus, some *trace* of the atomic constituents is present even in the complex representations. In this sense those constituents might be regarded as implicitly present. However, I shall not insist upon this view, since I think a stronger, related argument exists. For the crucial point is that certain complex (distributed) connectionist representations can *not only be decomposed* into their atomic constituents, via assumed background processes, but can be *generated* from their atomic constituents in a systematic fashion. That they can be so generated is evident from results reported in Chalmers (1990), Elman (1990), Niklasson and van Gelder (1994), Pollack (1990), and Smolensky (1990) (see previous footnote).

Now, suppose that a given connectionist system employs both atomic representations (for 'cat', 'sees', etc.) and complex distributed representations that are generated from those atomic constituents. We have seen that such complex representations can, if they are *suitably gen-*

erated, exhibit striking mathematical regularities. Also, as argued in the preceding section, these mathematical regularities can engender patterns exhibiting systematicity, provided background processes are in place (and we know that classicists must assume comparable background processes). So, atomic constituents can (in the presence of background structures) cause complex representations to exhibit mathematical regularities. These regularities, in turn, cause systematic relationships to arise in the presence of suitable processing mechanisms. It would seem, therefore, that within *some* connectionist networks, atomic constituents can cause mathematical properties which entail systematicity. In such systems, atomic constituents *do have causal efficacy* because causality is a transitive relation.³ Moreover, whatever causal efficacy *classical* constituents possess in a classical system, they possess only in the presence of assumed background processes. So, again, there appears to be no relevant asymmetry between classical and connectionist architectures *vis-a-vis* the causal efficacy of constituents. In both cases, much depends upon the nature of the background processing mechanisms that are in place.

It might, perhaps, be objected that connectionist constituents have a less *direct* causal efficacy than classical constituents. However, even if this were true, it is far from obvious that this is a *relevant* difference. Moreover, it's not clear that any appreciable difference in the directness of the causal relationships actually exists. For, a classical constituent in a complex structure may have efficacy only after some algorithm analyses the entire structure and determines both the identity of its constituents and their relationship to one another. It may often happen that neither the connectionist nor the classical constituent will have an immediate causal effect in producing some systematically related effect.

6. Summary.

In the foregoing, I have argued that, contrary to the claims of Fodor, Pylyshyn, and McLaughlin, there is no relevant asymmetry between classical and connectionist architectures with regard to (a) the nomic necessity of the systematicity of cognitive capacities, or (b) the potential for those architectures to *explain* systematicity. We have seen that the mere presence of classical representations within a system does not, by itself, entail systematicity; appropriate processing mechanisms must be in place. Yet, the situation appears to be precisely analogous within connectionism. Provided the *right kinds* of representations are employed, and that appropriate weight vectors (processing mechanisms) are in place, connectionist systems exhibit systematicity as a matter of nomic necessity. Furthermore, an *explanation* of systematicity appears feasible within both architectures. In both cases, the explanation appeals to how the processing mechanisms (whether algorithmic or weight vectors) exploit the structure present within the representations.

References.

> Chalmers, D. (1990). Why Fodor and Pylyshyn were wrong: the simplest refutation. *Proceedings of the Twelfth Annual Conference of the Cognitive Science Society*, Cambridge, Mass.

³I have recently learned that the preceding argument, in its essentials, is given by van Gelder in (van Gelder, 1991).

- > Elman, J.L. (1990). Finding structure in time. *Cognitive Science*, 14, 179-212.
- > Elman, J.L. (1991). Representation and structure in connectionist models. In G. Altmann (Ed.) *Computational and psycholinguistic approaches to speech processing*. New York: Academic Press.
- > Fodor, J.A. and McLaughlin, B.P. (1990). Connectionism and the problem of systematicity: Why Smolensky's solution doesn't work. *Cognition*, 35, 183-204.
- > Fodor, J.A. and Pylyshyn, Z.W. (1988). Connectionism and cognitive architecture: a critical analysis. *Cognition*, 28, 3-71.
- > Hadley, R.F. (1992). Compositionality and systematicity in connectionist language learning. *Proceedings of the Fourteenth Annual Conference of the Cognitive Science Society*, Bloomington, Indiana, 659-664.
- > Hadley, R.F. (1994a). Systematicity in connectionist language learning. *Mind and Language*, 9, 247-272.
- > Hadley, R.F. (1995). The explicit-implicit distinction. *Mind and Machines*, 5, 219-242.
- > Ingram, D. (1985). The psychological reality of children's grammars and its relation to grammatical theory. *Lingua*, 66, 79-103.
- > McLaughlin, B.P. (1993a). The connectionism/classicism battle to win souls. *Philosophical Studies*, 71, 163-190.
- > McLaughlin, B.P. (1993b). Systematicity, Conceptual Truth, and Evolution. In C. Hookway and D. Peterson (Eds.) *Philosophy and Cognitive Science*. Royal Institute of Philosophy, Supplement No. 34.
- > Miyata, Y., Smolensky, P. and Legendre, G. (1993). Distributed Representation and Parallel Processing of Recursive Structures. *Proceedings of the Fifteenth Annual Conference of the Cognitive Science Society*, Boulder, Colorado, 759-764.
- > Niklasson, L.F. and van Gelder, T. (1994). On being systematically connectionist. *Mind and Language*, 9, 288-302.
- > Phillips, S. (1994). Strong systematicity within connectionism: the tensor-recurrent network. *Proceedings of the Sixteenth Annual Conference of the Cognitive Science Society*, Atlanta, GA, 723-727.
- > Pinker, S. (1984). *Language learnability and language development*. Cambridge, MA: Harvard University Press.
- > Pollack, J.B. (1990). Recursive distributed representations. *Artificial Intelligence*, 46, 77-105.
- > Smolensky, P. (1988). Connectionism, constituency, and the language of thought, in Loewer, B. and Rey, G. (Eds.), *Meaning in Mind: Fodor and his critics*. Oxford: Blackwell.
- > Smolensky, P. (1990). Tensor product variable binding and the representation of symbolic structures in connectionist systems. *Artificial Intelligence*, 46, 159-216.
- > Smolensky, P. (1994). Constituent structure and explanation in an integrated connectionist/symbolic cognitive architecture. *The Philosophy of Psychology: Debates on Psychological Explanation*, C. Macdonald and G. Macdonald (eds.), Oxford: Basil Blackwell.
- > St. John, M.F. and McClelland, J.L. (1990). Learning and applying contextual constraints in sentence comprehension. *Artificial Intelligence*, 46, pp. 217-257.
- > van Gelder, T. (1990). Compositionality: A Connectionist Variation on a Classical Theme. *Cognitive Science*, 14, 355-384.
- > van Gelder, T. (1991) Classical questions, radical answers: Connectionism and the structure of mental representations. In T. Horgan and J. Tienson (Eds.), *Connectionism and the Philosophy of Mind*. Dordrecht: Kluwer.
- > van Gelder, T. and Niklasson, L.F. (1994). Classicism and cognitive architecture. *Proceedings of the Sixteenth Annual Conference of the Cognitive Science Society*, Atlanta, Georgia, 905-909.