9.12.1991/NOB

TWO GAMES IN TOWN

Systematicity in distributed connectionist systems

Niels Ole Bernsen and Ib Ulbæk
Centre of Cognitive Science. Roskilde University, Denmark

Summary: The paper firstly analyses Fodor and Pylyshyn's (1988) central claims that (a) cognition generally requires the properties of systematicity and compositionality which (b) are unobtainable in connectionist systems. Secondly, a distributed connectionist simulation is described which exhibits the properties of systematicity and compositionality. Thirdly, an analysis is provided for the mechanism responsible for these properties of the described system.

1. Introduction

Connectionism now appears as if it has come to stay as a second computational paradigm for cognitive science in addition to the paradigm of classical AI (for the latter, see, e.g., Pylyshyn 1984). Connectionism with distributed representations throughout has been presented as a complete, self-sufficient, and non-hybrid alternative to the classical paradigm in accounting for the representational states of an organism or its cognitive architecture (Smolensky 1988). There seems, however, to remain a fundamental unsolved problem with distributed connectionism, namely, that of accounting for the constituent structure of thought. In this paper, we want to take another look at this problem and demonstrate that it can be solved both theoretically and in a working connectionist simulation. Distributed connectionism, therefore, can be considered a second general computational paradigm for cognitive science.

Fodor and Pylyshyn (F&P, 1988) have argued that thoughts or representational mental states, like language, have *combinatorial syntactic and semantic structure* and that the utilisation of such structural properties is crucial to inference and reasoning. It is because mental representations have combinatorial structure that it is possible for mental operations or processes to apply to them by reference to their form. Thus, mental processes have *structure sensitivity*. Structures of expressions can have causal roles because structural relations are encoded (or implemented) by physical properties of brain states in appropriate ways. Localist connectionist networks, F&P claim, do not have combinatorial syntactic and semantic structure and the processes operating on them do not have structure sensitivity. Such networks, therefore, are at most vehicles for the implementation of cognition rather than accounts of cognition at the proper theoretical level which to F&P is the level of complex symbol structures. Distributed representation networks are no better off in this respect than are localist networks. Briefly, the argument runs as follows with respect to localist networks: such networks need one set of elements to represent, e.g., the thought that John loves Mary and a different set of elements to represent the thought that Mary loves John. No one set of elements is able to represent the combinatorial syntax of the thought that John loves Mary for the simple reason that such a set of elements does not have combinatorial syntactic and semantic structure. In particular, the node which fires whenever the system entertains the thought that John loves Mary does not have syntactic and semantic constituent structure. It is a simple, atomic, and therefore unstructured node which fires whenever the system has the thought that John loves Mary. And of course, processes involving this node cannot be sensitive to a structure that the node doesn't have.

However, it now appears possible to provide localist networks with constituent structure and variable binding (Ajjanagadde and Shastri 1991). Unfortunately, localist networks are highly implausible from a cognitivist point of view so we will not go into them here. This brings us back to the unsolved problem of how to represent constituent structure in distributed networks (Smolensky 1987, Fodor and McLaughlin 1990). In what follows, and for the reasons already indicated, we shall only discuss distributed connectionist systems.

2. The problem

The problem raised by F&P is a fundamental one for the following reasons. A large and central part of cognitive representation and processing is the part which makes us able to think, reason, and understand and produce natural language. Whatever its exact nature, this part of cognition involves *thoughts* as a form of mental representation. If thoughts have a set of properties, *P*, which cannot be accounted for by distributed connectionism then distributed connectionism cannot account for the central part of cognition to which thoughts are basic. And if distributed connectionism cannot do this, it cannot be a candidate for a general computational paradigm for cognititive science. So the question becomes, what is the set of properties *P* ? As we saw, a first characterisation of *P* is "combinatorial syntactic and semantic structure" and "structure sensitivity". However, this characterisation is too close to the classical paradigm for cognitive science that F&P want to defend. In their actual development of the argument, F&P are rather careful not to presuppose the classical paradigm and arrive at a characterisation of the set *P* which clearly does not presuppose their own explanatory theory. It is this non-circular characterisation in terms of "systematicity" and "compositionality" which is crucial to their argument and which will be discussed in what follows.

The core of F&P's (1988) argument is an example of more or less the following kind. The sentence "John loves Mary" is built compositionally according to syntactic and semantic rules from the constituents "John", "loves" and "Mary". According to the Language of Thought hypothesis inherent to F&P's version of the classical paradigm for cognitive science, the same is true of thoughts. The syntactic constituents of thoughts are represented in the system through some suitable physical implementation. Given this, a system which is able to represent the thought that John loves Mary is necessarily (or "inherently") also able to represent the thought that Mary loves John. This is obvious on the stated premises: both thoughts involve exactly the same lexical entries and can be correctly formed using exactly the same syntactical rules. Just as you don't find native speakers who know how to say in English that John loves Mary but don't know how to say in English that Mary loves John, you don't find systems who are able to think the thought that John loves Mary but are unable to think the thought that Mary loves John. Thus, such systems are characterised by syntactic and semantic *systematicity*, that is, their ability to understand or produce (and therefore think) some sentences is intrinsically connected to their ability to understand or produce (and therefore think) certain others. In fact, systematicity follows from the postulation of constituent structure.

The systems' mental representations as just described are also characterised, to some considerable extent, at least, by syntactic and semantic *compositionality* . Systematicity and compositionality may be two aspects of a single phenomenon, F&P speculate. "Compositionality", in this context, is simply the fact that the systematically related thoughts that the described systems are able to have are not only related syntactically but also semantically. Which sentences are systematically related is not arbitrary from a semantic point of view. Normally, a lexical item makes approximately the same semantic contribution to each expression in which it occurs. Thus, a system which is able to represent the thought that John loves Mary is necessarily also able to represent the thought that Mary loves John, but it does not have to be able to represent other thoughts having the same syntactical structure but a different semantics. Compositionality, in this sense, is so closely related to systematicity that we need not spend much time on considering it separately in what follows. In conclusion, the set *P* fundamentally comprises the property of systematicity. The property of (some measure of) compositionality follows from systematicity both on F&P's account and on our alternative account. However, we reject F&P's claim that compositionality presupposes syntactic/semantic structure in sentences.

As said earlier, F&P are careful (most of the time) not to make the case for the systematicity of thought rest on the Language of Thought hypothesis. Nor does the case rest on systems' possession of linguistic capacity. We are dealing with a much more fundamental property of thought and mental representation. Non-linguistic animals and infraverbal cognition also demonstrate systematicity of thought and F&P claim that the inadequacy of connectionist models as cognitive theories follows quite straightforwardly from this empirical fact. So, basically, the thesis of the systematicity of thought is a claim about the systematicity of representations underlying a great deal of the observable behaviour in humans and animals. If some human or animal is able to think certain thoughts (or have certain mental representations), they are necessarily also able to think certain other thoughts which can be seen to be systematically related to the former. We have no reason to dispute this claim. What is at stake is rather the conclusion which is being derived from it, namely, that thought has combinatorial syntactic and semantic constituent structure in the sense of the physical symbol systems hypothesis. We will *not* go deeply into the question of the structure sensitivity of thought in this paper but consider this a matter for future work. We simply hypothesise that once we have an alternative account of the systematicity of thought, i.e., an account which does not explain systematicity as a consequence of combinatorial syntactic and semantic constituent structure in the sense of the physical symbol systems hypothesis, but provides a different explanation, then it will also be possible to develop an alternative account of what F&P call the structure sensitivity of thought. Once we have non-syntactic combinatorial structure, the next step will be to define structure sensitive operations over the structural items. It should be remarked here that the structure sensitivity of thought does not form part of F&P's central argument.

F&P claim that since connectionist systems, or at least (we add) connectionist systems with distributed representations evidently don't have combinatorial constituent structure in the sense of the physical symbol systems hypothesis, the necessity inherent in the property of systematicity does not obtain in them. If distributed connectionism is the right paradigm for the description of biological cognition, it follows that there might be biological systems which are able to think that John loves Mary but unable to think that Mary loves John. And this is extremely unplausible, or even preposterous. On distributed connectionist principles, the systematicity of thought is a *mystery*. Therefore, cognition requires combinatorial constituent structure and the systematicity that goes with it, and, so far at least, the classical paradigm remains the only game in town which provides an intelligible account of this fundamental aspect of cognition.

3. The core issues

The core issue is not one of syntax. Connectionist systems with distributed representations do not have a physically implemented syntax the way classical systems have. The standard notion of syntax is difficult or impossible to apply to distributed systems. If the issue were one of having classical syntax there would be no need for argument. As we have seen, the issue is more fundamental than that and the claim that we have to consider is the claim that *only* classical syntax or the existence of an internal syntactic structure of mental representations can explain the systematicity of mental representation that is evident from much of human and animal behaviour. F&P claim that "systematicity arguments infer the internal structure of mental representations from the patent fact that nobody has a *punctuate* intellectual competence." So the core issues are the following two: (1) We need to demonstrate that distributed systems *are able to* produce the required systematicity of mental representation. This will be a behavioural demonstration following standard criteria for connectionist demonstrations of system behaviour. Such demonstrations of the behaviour of connectionist systems on given tasks are notorious for being subject to many different types of objection. We have tried to anticipate some of those in the description of the simulation below. (2) We need to explain, at the appropriate level(s) for the description of cognitive representation and processing, including the semantic level, the *mechanism* which allows connectionist systems to achieve systematicity (and compositionality) as a matter of necessity. This mechanism, of course, cannot and should not be a classical syntactic and semantic one. Requirement (2) turns out to be an interesting and difficult one to completely meet at this point. We are confident to have met it in a way which is sufficient for our purpose. Meeting it completely is a different matter, since that might lead one into the (so far) obscure terrain which some call that of "providing a semantic account of distributed connectionist representation" and which is known from, i.a., Smolensky's papers.

Clearly, the core issue is not one of representing the examples concerning John's love for Mary and Mary's love for John. At the end of the day, if successful, distributed connectionist systems would have to be able to also represent such examples, but the issue over systematicity and its explanation clearly does not hinge on this particular case. As F&P point out "... linguistic capacity is a paradigm of systematic cognition, but it's wildly unlikely that it's the only example". In other words, the systematically related thoughts of a cat or a chimpanzee would suffice. We are thus free to choose a different example for experimental demonstration and subsequent interpretation. On the other hand, the issue *does* seem to be one of representing central cognition or linguistically expressible *thoughts* rather than possibly unconscious and linguistically inexpressible peripheral cognitive states and processes.

So we need an example of a complex thought. It could be a two-place relational thought just like the ones about John and Mary. It should have systematicity in the sense that a system should not be able to have the thought that *aRb* (or *R(a,b))* without, necessarily, being able to have the different, but systematically related thought that *bRa* (or *R(b,a)*), *a* and *b* being individuals. If a connectionist system with distributed representations is able to entertain such systematically related thoughts then such systems demonstrate systematicity. They also demonstrate compositionality in the sense introduced above. And, *ex hypothesis,* they do so without having syntactical representations in the classical sense of the term. Finally, we would like to be able to explain why the mentioned necessity obtains. If this is possible, then the classical syntactical paradigm is no longer the only game in town which is capable of providing an intelligible account of central cognition possibly including language, inference and reasoning. And if that is true, then we might have identified at least part of a lead as to how to go about building biologically plausible, distributed connectionist models of central cognition.

4. A different example

The thought that John is to the right of Mary would seem to fill the bill as well as any. What is required for a system to be able to have such a thought ? Clearly, the system should be able to correctly apply the concept "to the right of" to many different kinds of object, not just to John and Mary when these two persons happen to stand in the proper spatial relationship. Applying the concept correctly requires an ability to correctly distinguish between situations in which two arbitrary objects stand in the proper spatial relationship and situations in which they do not stand in the proper spatial relationship. And if the system is really able to correctly apply the concept to *arbitrary* sets of spatial objects, then *it does not matter whether the situation is one of John's standing to the right of Mary or one of Mary's standing to the right of John*. In other words, once the system is able to apply the concept to Mary's standing to the right of John, it is *necessarily* also able to apply the concept to John's standing to the right of Mary. This is what systematicity is all about. Finally, we would like to have a system which is able to learn the concept "to the right of" from experience since this is undoubtedly the way biological systems come to master such spatial concepts. In other words, we need a system which is able from experience to *abstract* the relation "*right-of(x,y)* " and subsequently to *instantiate* "*x* " and "*y* " to arbitrary individual objects in space in order to determine whether the "right-of" relationship obtains between them.

A system of the kind just described is able to handle abstract concepts and variables on the basis of experience and training. It is able to somehow *bind* those variables to concrete particulars. And it is able to determine, once the variable-binding has taken place, whether the concept is true of those particulars or not. In virtue of these capabilities the system will necessarily be able to have the thought that *R(b,a)* once it is able to have the thought that *R(a,b)* , since the only difference between these two thoughts is the different bindings of the variables *x* and *y*.

Note that nothing has been said about syntax here. We are not saying that the described system does acquire, through experience and training, a syntactic representation of the form "*right-of(x,y)* ". We are not saying that on the basis of such a representation the system performs a formal syntactic operation of binding, through the operation of substitution which we call instantiation, the variables *x* and *y* to, say, John and Mary. Such representations and operations form part of *one particular* (syntactic) algorithmic way of describing what the system might be doing. A distributed connectionist system does not do things this way. It does learn the abstract concept that *something* is spatially related to *something else* through the relation "right-of". But instead of the variables *x* and *y* it has a pattern of weighted connections between its units of activation which perform *as if* they were variables like *x* and *y*, or, rather, which perform the same task as that performed by a syntactic system with variables *x* and *y*, but differently. And the distributed system does not formally bind the variables *x* and *y* (which it doesn't have) to concrete individual objects through the formal syntactic operation of substitution. Rather, the "something" and "something else" parts of the network (which are represented by its weighted connections and units of activation) become activated by input representing individual objects in space. This activation allows the network to determine whether or not those objects stand in the right spatial relationship for the relationship "right-of" to obtain between them.

The system described does not, strictly speaking, know of formal logic and does not represent the world in terms of formal logic. But it does represent abstract concepts and knows how to apply them to individuals that it perceives in its world. It represents abstract concepts of two kinds. First, it represents the concept of a spatial object in general, more or less. Second, it represents the concept "right-of" in general, more or less, since it is able to correctly describe objects in different positions as being or not being to the right of other objects. The powerful mechanism of abstraction from experience and subsequent instantiation to experienced objects is what is responsible for these capabilities and thus for the system's mastery of systematicity (and compositionality). We would not hesitate in claiming that the system does master (non-syntactic) combinatorial semantic constituent structure. This point merits a little more reflection.

The abstract representation which our hypothesised system has "that some spatial object is to the right of some other spatial object" clearly is a semantically *complex* representation. Being abstract, this representation is able to *generate* infinitely many different instantiations. In virtue of its abstractness, it is also to a large extent *context-independent* (contrast Smolensky 1988). We obtain all these classical properties without having to assume a syntactic level of representation consisting of atomic symbols and complex symbols having atomic symbols as their parts. F&P argue that the representations *aRb* and *bRa* literally have the same parts and call this "real constituency". Does our system have real constituency ? If the criterion for real constituency is that the system does represent, in both cases, the external objects *a* and *b* and the relation *R* between them, then the answer is clearly affirmative. If the criterion is the syntactic symbol system hypothesis, then the answer is just as clearly negative. The same argument applies to the question whether our system has *combinatorial* constituent structure. In both cases, it's an empirical question that is totally independent of the basic issues dealt with in this paper, which patterns of activation we would find in a network when the network has the complex representations *aRb* and *bRa*, respectively. Finally, we are justified in claiming that the *structure* of the complex representation which the system has (e.g., *aRb* or *bRa*) has a causal role in the production of the system's behaviour: the system responds differently depending on which one of these complex representations it has. In F&P's terminology, when the system's representations of *a, b,* and *R* are simultaneously active and the system has the complex representation *aRb,* then we also have to admit that the system's representations of *a*, *b*, and *R* enter into a specific kind of "construction" with each other. This construction differs from the construction among *a*, *b*, and *R* when the system has the complex representation *bRa*. So, the constituency relations are themselves semantically significant as F&P claim they should be. The distributed connectionist representation "a is to the right of b" is a *non*-atomic mental representation having *non*-syntactic structure. It is therefore just false to maintain that "... we cannot have both a combinatorial representational system and a connectionist architecture *at the cognitive level* ".

This is what we have to say about the system's mastery of systematicity and the mechanism by which this mastery is achieved. We believe this to be sufficient for our purpose without having to go into issues of "sub-symbolic semantic representation", "semantic microfeatures", and the like. It turns out not to be too difficult to build a distributed connectionist system with these capabilities. The system is a kind of micro-world animal, but in contrast to the animal described by F&P (1988) this animal masters systematicity:

> "Such animals would be, as it were, *aRb* sighted but *bRa* blind, since presumably, the representational capacities of its mind affect not just what an organism can think, but also what it can perceive. In consequence, such animals would be able to learn to respond selectively to *aRb* situations but quite *un*able to learn to respond selectively to *bRa* situations. (So that, though you could teach the creature to choose the picture with the square larger than the triangle, you couldn't for the life of you teach it to choose the picture with the triangle larger than the square)."

5. A distributed network with constituent structure and systematicity

We have built a network with distributed representations having the properties just described. The simulated network learned how to apply the concept "to the right of" through being trained on pictures of discriminably different 2-D objects. There was a semantic difficulty which had to be overcome. The concept "to the right of" is closer to the description of perception than is "loves" and it has a simpler and less exciting semantics, but the semantics of "to the right of" is not that simple either, since it has an asymmetrical "trajectory-landmark structure" (Langacker 1987). When an object, *a*, is said to be to the right of another object, *b*, then object *b* acts as landmark for the trajector, *a*. To capture this property, we placed one object at the centre of the 2-D array whenever the presentation contained a landmark object. This placement gave landmark status to the object without the need for separate labels for any of the objects used in the simulation in addition to their different visual appearances and positions. (Presumably, the concept "to the right of" is normally learned only by creatures which have an independent vocabulary for the description of objects. The setup described circumvents this difficulty without giving way on the crucial issue of systematicity). It might be objected that the system does not learn the completely general concept "to the right of" but only learns the concept "to the right of a fixed landmark". This is true, but we don't consider the objection serious with respect to the principles we want to demonstrate. Our actual "to the right of" concept is even more complicated since it also allows us to change coordinate systems from a viewer-dependent coordinate system to an object-centered coordinate system. Again, this does not affect the central point of the demonstration.

The system also has to learn that "to the right of" is a two-place predicate. When there is only one object (or marked position) in the scene, or when there are more than two objects (or marked positions), the question whether "this object is to the right of that object (the landmark)" either does not make sense or is ambiguous. In such cases, the system answers "no" to the question posed to it. On all other presentations one object is placed at the landmark site. A second object is then placed in one of four different positions around the landmark object (right, left, above, or below). We did not teach the network to discriminate among all those positions but simply to respond with a "yes" if and only if the trajector was positioned to the right of the landmark, and to respond with a "no" otherwise. In this way, the network is answering the question: "Is the trajector to the right of the landmark ?" If there is no landmark, it responds with a "no" and if there are three or more objects, it also responds with a "no".

The network is a standard one-layer backpropagation network with graphics facilities for the display of presented objects and running on a PC. The training tolerance for output is 0.1, which means that on a scale from 0 to 1 the network will count 0.9 as correct and stop training when all exemplars in the training set perform above 0.9. The testing tolerance is 0.4 which is sufficient for mechanically distinguishing success and failure. The 2-D picture array measures 8 x 20 (160 input units). The hidden layer has 30 units and the output layer has 2 units for "yes" and "no". The training set consists of 6 different objects which are placed in different numbers and combinations and sometimes as landmark, sometimes as trajector. The landmark site and each trajector site consists of a field of 4 units. The different objects occupy different numbers and combinations of units at a site. The test object set included 3 objects different from the 6 in the training set (see fig. 1).

[Insert fig. 1 around here]

To demonstrate that the network can handle the systematicity of *aRb* and *bRa*, we only trained the network on one of these relations for a given pair of objects while saving the second relation between the pair for the test. Thus, (1) if the network had been trained on "*a* is to the right of *b* " it was not trained on "*b* is to the right of *a* ". In the test phase, the network was shown already familiar objects in combinations it had not encountered before. In addition (2), the network was shown objects it had not encountered before in order to verify that the network was able to abstract a "general" concept of "spatial object". Taken together, (1) and (2) offer sufficient evidence that the network is able to master systematicity from *aRb* to *bRa*; abstraction to the "right of" concept which we are used to representing as *xRy*; and abstraction over all possible objects in its world, thus successfully taking the set [*a*,*b*,.....*i*] as instantiations of *x* and *y* or as legitimate arguments of the relation *R*.

The training file consisted of 84 training exemplars. The network converged on the desired output in 24 epochs with the mentioned training tolerance of 0.1. The test file consisted of 65 testing exemplars. The network was able to generalise successfully within the testing tolerance of 0.4. In other words, systematicity is so simple that a mouse could probably achieve it if its cognitive architecture consisted of distributed connectionist networks. It is, therefore, *no* mystery why nature contrives to produce only systematic minds.

6. Concluding remarks

Perhaps the most interesting point about systematicity is that it is a function of the abilities to abstract general concepts from the perception of particular instances and to instantiate those concepts to new (or old) instances. Whether or not the network had actually met with a training case of *bRa* having already been shown a case of *aRb*, is of much less importance than its ability to perform concept abstraction since it is the latter which achieves systematicity and compositionality. As for compositionality, the necessary capability to handle *bRa* once *aRb* can be handled follows directly from the mastery of abstraction and instantiation since it is one and the same abstract concept which is being applied in both cases. And the mechanisms in distributed connectionist networks for handling abstraction and instantiation provide alternatives to the syntactic mechanisms of variables and variable binding.

One or two further questions present themselves. Answering these questions is not essential to the "only game in town" issue as raised by F&P. The questions might, however, come in focus in a subsequent discussion of whether there are one or two games in town. In such a discussion, it goes without saying, the "only game in town" point of view would start out from a dramatically weakened position. The position is dramatically weakened for the following reason: the position rests on the assumption that thoughts in general have a set of properties, *P*, which cannot be accounted for by distributed connectionism, notably the properties of (representational) systematicity and compositionality. We have already accounted for these on distributed connectionist principles *and there aren't that many other candidate properties which might belong to P*.

The first question is how distributed networks can demonstrate *structure sensitivity* of processing. It seems clear that networks can only handle systematicity because they exhibit a distributed version of combinatorial semantic constituent structure. They are able to represent individuals, generalised individuals (the "spatial objects in general" of the simulation above), abstract concepts, and applications of these concepts to individuals. What more is required to demonstrate structure sensitivity of processing ? If, for instance, as already demonstrated, (1) a network is able to verify that *aRb* ; if (2) we can make the landmark move so that the network could also verify that, *cRa* ; and if (3) the network could then conclude that *cRb* ; then it would seem that we had demonstrated structure sensitivity of thought processes in distributed networks. However, F&P are not very explicit on the requirements for a demonstration of structure sensitivity of processing.

The second question concerns the *systematicity of inferential processes* which, in contrast to our first question, is actually discussed by F&P. This kind of systematicity is the one we find in systems with the following capability. If the system is able to deduce *a* from the conjunction *a&b*, and if the system is also able to deduce *a* from the conjunction *a&b&c&d*, then it would be extremely unlikely, if not preposterous, to find that the system was not able to deduce *a* from the conjunction *a&b&c*. We have not demonstrated inferential systematicity above. However, our solution to the representational systematicity of thought might already have provided the mechanism needed for inferential systematicity. What a system with inferential systematicity has to learn is an abstract schema for propositional inference from conjunctions. The abstraction is the following. It doesn't matter how many conjuncts you have in a set or which they are: you are always allowed to infer a subset of the conjuncts from the set. We have already seen the general character of a distributed connectionist mechanism needed for this purpose. We speculate that this mechanism will also work for inferential systematicity.

References

Ajjanagadde, V. and Shastri, L.: Rules and variables in neural nets. *Neural Computation* 3, 1991, 121-34.

Fodor, J.A. og Z.W. Pylyshyn: "Connectionism and Cognitive Architecture: A Critical Analysis", *Cognition* 28, 1988, 3-71.

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

Langacker, R.W.: *Foundations of Cognitive Grammar*. Stanford CA: Stanford University Press 1987.

Pylyshyn, Z.W.: *Computation and Cognition: Toward a Foundation for Cognitive Science.* Cambridge MA: MIT Press 1984.

Smolensky, P.: The constituent structure of connectionist mental states: A reply to Fodor and Pylyshyn. *Southern Journal of Philosophy* 26, 1987, supplement, 137-161.

Smolensky, P.: On the proper treatment of connectionism. *Behavioral and Brain Sciences* 11, 1988, 1-74.

Niels Ole Bernsen and Ib Ulbæk, November 1991