

Resolving Dualism: On the Intersection of Connectionist and Computational Networks

Heather Balcom

Traditionally, the object of cognitive science has been the activities of a rule governed non-physical self. This "mind" experiences, believes and chooses; to understand cognition was to learn the rules by which it operates. This view of cognition as computation or information processing has been challenged by a newer account of mind as a dynamic system that combines many computations into something that doesn't act at all like a computer. Then apparent systematicity is incidental; semiotic interpretations of cognitive function should be abandoned to be replaced with an understanding of mind as a subdivision of body. Neither view, I will argue, has accounted, or has the potential to account, for all that should be described as cognitive, yet a view that encompasses both may offer an adequate framework, or at least lead us to some novel questions.

In this paper, I defend what Bechtel calls an "external symbol approach," but propose a set of consequences that differs from the conclusions he offers. Bechtel argues that rules and symbols are, and must be, external to the mind, thus, computational explanations of cognition are approximate models of what is, in fact, a connectionist system. I agree that discreet symbol processing cannot account for the internal workings of the processor, in this case, the mind. However, computational theories offer terrible (not approximately true) explanations of neural networks, and accurate (again, as opposed to approximate) accounts of some aspects of thought. To speak of neural networks as having beliefs, perceptions, etc. is misleading. A neural network is not a person because it can't think or believe or choose; it is not really anything, because it *does nothing*. The actions of a dynamic system are functional. Movement, visuals, sounds, sensations, and all other mediums of cognition are external objects of computation, but functional roles are most evident in human language, and language will be the focus of this paper. External symbols and practices, including the computationally processed components of language, belong in any theory of mind,

right along side, and intertwined with, discussions of neurons and parallel processing.

Let "language" be the system that makes productive and systematic behavior possible. Following Fodor and Pylyshyn, Bechtel defines these capacities as follows: A productive system can always generate a novel state from extant states. For example, embedding clauses allows indefinitely many English sentences to be created. Systematicity is the quality that representations are related. That is, members of each type of symbol are intersubstitutable. If I can say "the cup is on the desk," I can also say "the desk is on the cup," or substitute any known noun for "desk" or "cup" (Bechtel 1993, pp. 127-128). Thus language consists in discreet symbols. The functional roles are not fixed; language changes over time, sometimes quite dramatically, yet, for any given usage, some functional role is fulfilled. Further, truly ambiguous sentences and sentences that are only syntactically ambiguous abound. While the functional role of sentential components is not always determinable by rules, in any use of language, all components have exactly one substitutional role.

Language, as thus defined, is external to the brain. This is not simply the statement that written and spoken words are not in the brain, or even that we get all of our specific words from external sources, that is, from language users around us. That is the uncontroversial part, which, Bechtel demonstrates, should be extended to say that words are not directly linked with anything in the brain that plays the same role in the neural network's activity as does the word in social interaction. *Any* use of language, be it spoken, written, or in thought, is best understood as an object apprehended *by* the brain rather than as a functional part *of* the brain. It is something thought of, and not the means of thought.

It still may be tempting to pin a functional role onto brain states both for intuitive and functional reasons. I mean something by my words, hence, that meaning is in me. This meaning might be described as an image or sound or sensation in some cases. For example, the meaning of the word "green" is, at least in part, an image of something green. But these experiences themselves require a processor to which they are external. In other words, a functionalist account of syntax asserts that input is processed according to rules. But by what? Something must do this processing, and it must recognize instances of the same symbol and apply rules to it. The utility of positing internal symbols in explaining thought has been overestimated; internal symbol processing adds nothing to a discussion of external symbol processing. We should not deny the

importance of brain activity to external symbol processing, but this activity cannot be computational.

Since there are no discreet symbols in the brain, language is processed by a neural network that is acting as such, i.e. not implementing any program. This processor might be some new kind of thing, neither computational nor dynamic, but given the scientific evidence available, the network hypothesis seems more probable. Bechtel explains the processing and production of syntactic structures as a sort of pattern recognition. That is, the network is trained on objects, cups and desks, which is what lets it produce "the cup is on the desk" and "the desk is on the cup," and so on. Just what constitutes an object varies; "the cup is on the desk" may be an indivisible whole with its own functional role, rather than a composite structure. I think that it is fairly clear that language consists in external objects; one couldn't very well use language without some kind of active expression. Likewise all other forms of cognitive activity are just that, *activity*. All cognitive behavior involves objects such as images, sounds, and motions, as well as words.

We have seen that linguistic activity is, in part, discreet symbol processing, but in this aspect, linguistic activity is not brain activity. Thus, beliefs should not be ascribed to the brain. Bechtel argues, and I agree, that words are not manipulated by brain processes. That is, there is nothing in the neural network that stands in for the words in a sentence such that the network may process the sentence as composed of discreet symbols. I want to add that there is nothing in the brain that acts anything *like* a discreet symbol. If language is a formal system with compositional meaning, and neural networks do not implement computations, then it is incorrect to describe a brain as having thoughts, beliefs, perceptions, or any of the other *propositional* attitudes that are, and should be, the focus of cognitive study. Dynamic systems are not independently the objects of cognitive science.

This 'tradition' argument is clearly not sufficient to defend the role of computations in cognition. The bare fact that cognitive science has taken propositional attitudes as its object is not a particularly good reason to believe that they exist to be studied or that they have much to do with cognition. If connectionist pattern recognition explains everything that the supposedly discreet linguistic symbols do, then talk about propositions is a mistake or an approximation. I contend that dynamic systems fail to meet this criterion of autonomy. The externalization of symbols, which makes the connectionist hypothesis viable, necessitates computational modeling of information processing. If dynamic systems, those

masters of pattern recognition, are trained on external structures, then those structures exist. If we deny the functionalist structures, then we deny the dynamic systems the means to exhibit apparently systematic and productive behavior. Connectionist accounts are insufficient to independently explain *any* cognitive functions. Thus, external symbols are not merely a tacked on addition to a theory of mind, designed to appease our intuitions; external symbols are actual and observable things.

It still may be argued that human reasoning, as opposed to how logicians might like humans to reason, often relies far more heavily on content than on form, that is, on things as opposed to rules. We regularly introduce idioms consisting in, but not dependent on, discreet symbols. Proficient language users become puzzled when someone replaces familiar word symbols, which are symbols for something, with new generalized symbols and asks them to do algebra. If the empty set, which doesn't even have the decency to act like a *generic* thing, is introduced, more floundering ensues. Thus the ability to understand and produce language depends on the ability to translate abstract linguistic symbols into connectionist objects. This might tempt us to say that language isn't really a productive, systematic practice after all, or insofar as it is, it isn't really a human practice. For example, Paul and Patricia Churchland would abandon belief talk entirely, while Bechtel keeps it around as a means of representing states of dynamic systems.

It is true that cognitive behavior sometimes breaks every set of computational rules. Rules in any system may be broken without disproving the existence of that system. Rather, the combination of rule breaking and mistake behavior is *evidence* for a computational framework. A connectionist network may be trained to follow syntactic rules; in fact, it must be trained to follow some sort of rules in order to be recognizable as thinking. Part of this training is correction of mistakes. So considering only the specific actions that are performed by the network cannot explain the *goal* of that network's behavior; a thinking dynamic system needs external symbols just as much as those symbols need a dynamic processor. Computations do not exist without something to do the computing, and the brain isn't a neural network if it doesn't do anything. That is, the brain is a connectionist system only in virtue of how it produces behavior, but this behavior, linguistic or not, has its own customs, and they must be understood on their own terms.

REFERENCES

Bechtel, W. (1993). "The Case for Connectionism." *Philosophical Studies* 71:119-154.