PART III

PSYCHOLOGY

https://doi.org/10.1086/psaprocbienmeetp.1986.1.193117 Published online by Cambridge University Press

What Happens to Accounts of Mind-Brain Relations If We Forego An Architecture of Rules and Representations?[1]

William Bechtel

Georgia State University

1. Introduction

While some philosophers have assumed that there are only two options for characterizing the ontological status of mental models in cognitive information processing psychology--treating them as nearly autonomous from theories of brain activity (Putnam 1975 and Fodor 1974) or eliminating them in favor of neuroscience accounts (Churchland 1979)--cognitive scientists have often tacitly assumed a third option. This involves treating the mental models as systems of rules and representations that are instantiated in the nervous system much in the way computer programs are instantiated in computers. While this seems to be the position of those endorsing the autonomy of psychology, it in fact is consistent with a much weaker interpretation if one recognizes that the vehicle in which mental processes or computer programs are instantiated may limit and constrain what kinds of mental processes or computer programs can be instantiated. While the computer or nervous system may not determine the program or mental process, it may yet provide a useful guide to the nature of the system. This way of looking at the relation of cognitive models to neural models or computer instantiations is well characterized by Newell's notion of a physical symbol system and Pylyshyn's notion of a functional architecture. These two concepts point in different directions but together describe a position that allows for a kind of autonomy to psychology while at the same time showing how neuroscience can be relevant to cognitive modelling.

One can explicate the essence of these two concepts by considering the variety of levels at which one can describe operations in a computer. One can describe them in terms of changes in physical components, or in terms of activity in logic circuits, or in terms of manipulation of symbols. Once one moves to the level of symbol manipulation, in fact, a whole variety of levels open up--the level of machine language or the levels of other languages that are interpreted or compiled into machine code. Newell's interest is directed at those levels at which we can characterize the computer in terms of symbol manipulation. His contention is that once one has a system that operates by manipulating symbols one has the necessary and sufficient conditions for general intelligent behavior (Newell and Simon 1976;

PSA 1986, Volume 1, pp. 159-171

https://doi.org/10.1086/psaprocbienmeetp.1986.1.193117 Published online by Cambridge University Press

Newell 1980). One of the features of a physical symbol system, however, is that its symbol processing capacities can be realized in different hardwares. It is this which provides cognitive models a degree of autonomy from models of neuroprocessing. However, this autonomy is not complete in that the nature of the hardware constrains what software can run. It is here that Pylyshyn's concept of a functional architecture is relevant. Whereas Newell's focus is primarily on accounting for the flexibility of human cognition, Pylyshyn's attention is on its fixed features. Pylyshyn (1980 and 1984a) proposes that there are some features of the human cognitive system that are fixed by the way they are instantiated in the nervous system and so provide "the functional architecture of the mind". (Since they are fixed, they cannot be altered by cognition.) Pylyshyn thus speaks of these operations as cognitively impenetrable and uses the criterion that they cannot be altered by cognitive activity as the means for identifying these features.

What Newell's and Pylyshyn's accounts provide is a way for understanding the relationship of cognitive investigations to neuroscience ones. The functional architecture specifies a mapping of the symbols or representations and the procedures and rules by which they are manipulated onto neurological states and neurological processes that modify these states. Although this mapping links the cognitive account that focuses on the representations and the neuroscience account that focuses on neurological states, these two accounts adopt different perspectives. The cognitive story is one that characterizes the processing of symbols provided by the functional architecture, the neurological story explains how such processing is able to occur in the brain. Thus, the functional architecture constitutes a bridge that connects neural processes to cognitive processes. As such, it can be used both to demarcate the two inquiries and to show how they are related to each other.[2]

While this view seems attractive since it avoids the twin perils of eliminative materialism and the neglect of neuroscience, there is growing reason to doubt its empirical adequacy. In this paper I will examine a recent challenge that undercuts the assumption that cognition consists in rule governed processing of stored representations similar to that which occurs in contemporary computers. I shall discuss this challenge in the next section and present some considerations drawn from the history of other sciences that give it plausibility. If this challenge turns out to be correct, the pleasing picture of the relation of cognitive inquiry to neuroscience inquiry will need modification. I will turn to this task in section 3.

2. The Connectionist Challenge to Rule Processing Systems

In invoking the model of the relation of software to hardware to account for the relationship of cognition to the brain one implicitly adopts a model of the mind as a system in which rules govern the manipulation of representations. This model, which historically emerged from Chomsky's generative models in linguistics and successes in programming von Neumann computers, in fact proved quite useful in providing psychologists a means of characterizing internal processing and studying it experimentally and so permitted a break with behaviorist

https://doi.org/10.1086/psaprocbienmeetp.1986.1.193117 Published online by Cambridge University Press

strictures. Until recently it has seemed to be the only conceivable approach that offered much hope of providing a framework in which an explanation of human behavior might be developed. Recently, however, researchers in cognitive psychology and artificial intelligence have become frustrated with certain limitations imposed by assuming that the mind had the character of a rule processing system and have begun exploring alternative types of systems that could explain some characteristics of human cognition that proved difficult to handle within a rule processing format.

The alternative models that have been developed are referred to as connectionist or parallel distributed processing (PDP) models. These models have been constructed by analogy with the structure of neural networks and differ from traditional models in that no rules govern the manipulation of symbols. These systems consist of large numbers of simple processing components ("units"), each of which has a certain degree of activation. These units interact by sending excitatory or inhibitory signals to other units. The strength of these inhibitory and excitatory signals is determined by the degree of activation of the sending unit and the connection strength associated with the pathway. (The connection strength of the pathway is something that can be set up to change depending on the local activity in the system.) Once such a system is set in operation, the units will excite and inhibit each other until the system settles into a roughly stable state, where it will remain until new external inputs are supplied.

To see how these models could account for psychological phenomena one needs to supply an interpretation. One of the simplest interpretations, because it maintains affinity to more classical information processing models, is to let each unit "represent" a hypothesis (or goal). The degree of activation of a unit then represents the probability attached to the hypothesis associated with that unit and excitatory signals from one unit to another represent the support one hypothesis offers to another, while inhibitory signals represent opposition between hypotheses. The activity in the system constitutes the system's attempt to settle on a particular hypothesis. While this approach departs from more traditional information processing approaches in a variety of respects, the most salient one for our purposes is that the activity of the units is not governed by an executive which regulates activity in the system by consulting specifically encoded rules. Rather, the units themselves excite or inhibit each other until one hypothesis "wins" out. This approach has been used to develop quite realistic models of a variety of cognitive activities (see Fahlman 1979; Feldman and Ballard 1982; McClelland and Rumelhart 1981; Rumelhart and Norman 1982; and Cottrell and Small 1983).

Another more radical way of interpreting PDP models does not interpret single units as serving representational functions. Rather, it treats a pattern of activity over a set of units as constituting a representation of an hypothesis or goal. The connection strengths between such units can be set so that the same ensemble of units can have numerous stable patterns of activity into which it will settle depending on initial input. Thus, the ensemble will represent different things on different occasions. While such approaches result in greater difficulties in execution (see Hinton and Sejnowski 1984), they produce some most intriguing properties for modelling cognitive phenomena. For example, these systems are capable of recreating whole patterns from parts of the pattern and of learning certain responses by adjusting the connection strengths between units. The partial pattern that serves as trigger need not even provide a perfect match to part of the whole. As long as it is close to the original pattern, the transmissions through excitatory and inhibitory linkages within the system can serve to restore the original. (See Rumelhart and McClelland (1986) for further discussion of the potential of such systems.) This more radical interpretation marks a further break with the traditional cognitivist approach in that it does not treat the mind as possessing stable representations but as simply having a structure that can recreate patterns when given appropriate input.

In rejecting the view of the mind as possessing rules that govern the manipulation of symbols, connectionist or PDP frameworks are rejecting a position that has a venerable history. Its roots lie in folk psychology, particularly in its philosophical guise, according to which reasoning involves performance of a sequence of transformations on propositions. That approach acquired additional plausibility from work in logic, especially from Church's thesis, which proposed that any decidable process could be performed by a universal Turing machine or other universal machine. While no actual symbolic processing machine is a universal machine since it must lack infinite memory, actual machines generally have sufficient memory to carry out a vast number of decision procedures and thus approximate the behavior of a universal machine. This has made symbol systems seem like the appropriate model for explaining intelligence. (See Newell and Simon 1976; Newell 1980.) Pylyshyn (1984b), for example, objects to connectionist systems on the grounds that without symbols we cannot capture the nearly infinite variability of possible behavior in a finite manner. While these results of logic show that symbol systems provide the capacity for modelling a vast variety of decision procedures, they should not be taken as showing that they provide the only way to do so. There may be other ways to explain this behavior.

Some reasons for questioning symbol processing accounts stem from their origin. Symbol processing accounts of cognition originated not with internal studies of the mind's activities, but as hypothetical models which could account for the data of human behavior. These models are external in the sense that they account for the behavior of the system whether or not they accurately portray the internal procedures responsible for producing the behavior. Other sciences have similarly started with such external accounts of how a system behaved and then tried to figure out the internal operations of that system. A common first move for such sciences has been to assume that the language used to describe the behavior of the system from an external perspective also provided a correct description of the activities within the system that produce the behavior. But subsequent inquiry has often revealed that internally a quite different set of operations are being performed. For example, after the work of Lavoisier, many researchers thought that foodstuffs were quite literally combusted in the animal and much effort was put into finding the site of combustion and to figuring out how oxygen was able to react with foodstuffs at the relatively low temperatures prevailing in the animal body. In actuality, as many decades of research finally showed, a vast array of complex processes

https://doi.org/10.1086/psaprocbienmeetp.1986.1.193117 Published online by Cambridge University Press

are involved in biological oxidation. There is, however, no direct combustion of foodstuffs by oxygen but a host of processes such as the removal or addition of pairs of hydrogen atoms from substrates, the removal of carbon dioxide molecules, phosphorylations and dephosphorylations, and ultimately a process of ion transfer. A very similar story can be told about the history of genetics: whereas Mendelian genetics spoke of genes coding for phenotypic properties, the actual mechanism, as it has been unraveled to date, reveals that coding for traits is accomplished through a vast variety of different kinds of processes.[3]

The point of such examples from other sciences is to show that the language appropriate for describing the overall behavior of a system may not characterize the actual processes responsible for producing that behavior. Applying this to the case of cognitive psychology, it may well be that while we can generally characterize the cognitive behavior of a person in terms of the manipulations of symbols, the mechanism responsible for that cognitive behavior may not employ such symbols. Presumably some operations are occurring within the person producing this behavior, but there is no reason that these operations need correspond to the kinds of operations we might posit in describing the whole system's behavior, just as we now recognize that we do not have to include a combustion process in explaining how animals oxidize their foodstuffs. All these internal operations must do is ensure that the behavior actually produced typically corresponds to those regularities we capture externally and describe in terms of symbol manipulations. This argument, of course, does not rule out the possibility that the behavior of humans is the result of the mind formally manipulating a set of representations in accordance with a set of rules it possesses. It only prepares us for the possibility that the symbol processing account might be wrong. (The point I am making here is quite similar to one Dennett 1978, has made in behalf of his instrumentalistic attitude toward intentional idioms such as statements of belief and desire that ascribe representations to persons. He rejects the claim that people really have beliefs and desires since he sees no reason to think that the processes at what he calls the "design stance" must correspond to the beliefs and desires ascribed to the subject from the "intentional stance". I have argued that this aspect of Dennett's position can be defended without making intentional accounts instrumentalistic by treating intentional ascriptions as real descriptions of how the individual is able to interact with a range of environments. As such they need not characterize the internal processes that make this possible. See Bechtel 1985b.)

There are, moreover, reasons stemming from within cognitive science to take seriously the challenge to rule-based models of cognition. Researchers in domains like perception have faced significant difficulties in developing rule-based systems that could accommodate the flexible capacities of human visual processing. In other domains researchers have claimed a higher degree of stability in cognitive performance, thus making rule-based accounts more plausible. For example, it has seemed like concepts were highly stable structures in terms of which subject store information. Recently, however, Barsalou (in press) has presented evidence to the contrary. He presents data showing variability in concepts used by the same individual over periods of one-month, from which he concludes that concepts are not the basic

https://doi.org/10.1086/psaprocbienmeetp.1986.1.193117 Published online by Cambridge University Press

units of cognition that are stored and retrieved from long-term memory but are constructed by the subject as needed to perform cognitive tasks.[4] One strategy for dealing with such apparent fluidity in cognitive performance is to develop more complex rules working on more basic representations and use these to explain both the deviations and proper instantiations of the incorrect simpler rules. The alternative strategy, adopted by PDP theorists, is to try to produce "a unified account in which the so-called rule-governed and exceptional cases [are] dealt with by a unified underlying [non-rule-based] process." (Rumelhart 1984, p. 60). The fluidity of cognition that seems to defy rule-based accounts then appears as evidence for something like a PDP approach.

3. Recharacterizing the relationship between Cognition and Neuroscience

The discussion of the previous section has at least raised the possibility that an alternative to an architecture of rules and representations might be most useful for explaining cognitive phenomena. That poses a challenge to what seemed to be a quite plausible account of the relation of cognitive inquiry to neuroscience offered in section 1 and raises the question as to whether there is an alternative perspective from which to view the relationship. One way that the PDP models might be viewed is as neuroscience accounts and not as cognitivist accounts at all. This perspective is suggested by the language Feldman uses in introducing a series of papers employing non-von Neumann computational frameworks. He speaks of these papers as "address[ing] the issue of how complex cognitive behavior might be reduced to brain structure." (Feldman 1983, p. 2). Although I do not think Feldman is using the term "reduction" in this sense, this might seem to be an endorsement of eliminative reductionism. The eliminativist interpretation of PDP models is supported by the fact that they were initially motivated by work in neuroscience.[5] Feldman (1983) and Rumelhart (1984) and many others have made it clear that part of the motivation for developing connectionist or PDP models is to account for the clear superiority of the brain over von Neumann computers in carrying out certain kinds of cognitive functions.[6] It is further supported by the way in which PDP theorists sometimes present the relationship of their models to more classical characterizations of cognitive activity. Rumelhart (1984) for example speaks of representations as "emerging" from lower level processing, giving the suggestion that accounts in terms of rules and representations are higher level descriptions of what is causally generated and properly explained at a lower level.

However, there is another prospect, seemingly incompatible with this, that proponents of these models sometimes endorse. Rumelhart and McClelland (1986) present themselves as describing the micro-structure of cognition and indicate that this micro-structure might give rise, at the macro-level, to something more like traditional serial information processing accounts. Thus, they suggest that the relation of their enterprise to traditional cognitive models is like that of sub-atomic physics to atomic chemistry. In the remainder of this paper I will explore what might be gained by recognizing different levels of cognitive activity and consider how PDP accounts might fit into this scheme.

The claim by Rumelhart and McClelland that all they are doing is characterizing the micro-structure of cognition might seem to take the thunder out of the PDP position. PDP accounts seem provocative just because they seem to do away with rules and representations as figuring in the causal nexus involved in producing behavior and treat apparently rule governed behavior as only an emergent regularity. If, instead, all the PDP theorists are up to is showing how the rules and representations architecture might be grounded in the nervous system, then traditional information processing theorists might think they can ignore PDP theories as they have tended to ignore neuroscience. However, that approach misses the point of an exploration of the micro-structure of any phenomenon. What moves to lower levels (and occasionally higher levels) have tended to do in the history of science has been to lead to a reconstrual of the processes at the initial level as well. This is seen clearly in the case of genetics. The discovery of the molecular substrate has resulted in the recognition that the Mendelian account mistakenly conflated different conceptions of the gene (e.g., as the unit coding for an amino acid, as the unit of mutation, as the unit of crossing over, etc.). New accounts at the Mendelian level have to distinguish these different operations (Hull 1974). Examination of a more micro-level of cognition will prove important if it similarly forces a change in the conception adopted at the macro-level. (This, in fact, is one of the reasons PDP theorists offer for examining the micro-level.)

One obvious way in which such theorizing may affect the macro-level is if it shows that representations may not be the fixed entities assumed in classical accounts, but entities that can be modified in a variety of ways. Here we can draw a speculative connection between the proposals of the PDP theorists and Barsalou's results concerning the instability of concepts. Barsalou proposed to explain the difference in an individual's concepts over time as due to an individual having different objectives or recent experiences at different times. If these concepts are the result of activation in a distributed system, one has a ready model of how differences in goals or other cognitive activity might bias the processing: the stable state of an ensemble on any given occasion will be influenced by other activity in the system and need not be precisely the same.

There are additional ways in which research on the microstructure might influence macro-level thinking about cognition. It may not only alter ideas about the basic units of cognition (the representations) but also provide ideas about new kinds of primitive operations that might operate on these representations, thus extending our conception of the rules governing cognition. Finally, it promises to provide an approach to a phenomenon that has posed a challenge to cognitive theorizing (and even been denied by Fodor (1980))--the phenomenon of concept learning. PDP models can demonstrate concept learning, but what is significant is that they locate the mechanism for learning at a lower level than the symbolic level. They explain it in terms of modifications in connection strengths within the system.[7] The shifting of a problem to another level so as to find a solution is a common strategy in science. For example, Darden (1986) shows how Dobzhansky helped solve the problem of speciation by locating it on a new level than had previous inquiries.

https://doi.org/10.1086/psaprocbienmeetp.1986.1.193117 Published online by Cambridge University Press

What the PDP accounts do in exploring processes at a more micro-level than traditional rule-based accounts of cognition is undercut Newell's and Pylyshyn's idea that there is an ultimate architecture for cognition, with cognitive operations all characterizable in symbol processing terms. In its place one may need to recognize a hierarchy of levels which may constrain each other in something like the way Pylyshyn viewed the architecture as constraining cognition.[8] This leaves the question of how neuroscience is to interface with these various levels of cognitive analysis. PDP accounts or other micro-level accounts are not neuroscience accounts, even though they have been motivated by work in neuroscience. Rather, they are abstract processing accounts. This is revealed by the fact that PDP researchers are not particularly interested in the details of the neurophysiological mechanisms that underlie this processing. Yet, insofar as they were motivated by neuroscience work, they show one point of connection between neuroscience and more cognitive investigations.

However, this should not be viewed as the only locus of connection between cognitive inquiries and neuroscience. There is a long tradition, stemming from the phrenologists, of trying to identify the cognitive functions of organs within the brain. These studies are clearly neurological and offer prospects of guiding cognitive theorizing at more macro levels. (For examples of such research with cognitive imports, see O'Keefe and Nadel, 1978, and Grodzhinsky et al. 1985.)

Recognizing that the contributions of neuroscience may come at a variety of levels of cognitive theorizing complicates the view sketched in section 1, for now we have several levels of neuroscience inquiry intermeshing with several levels of cognitive inquiry. However, a final comparison with other sciences indicates that there is nothing seriously problematic about this prospect. One finds a similar kind of interlacing of contributions from various approaches at different levels in the interactions of organic chemistry, biochemistry, and physiology. On the one hand, the functioning of whole organs (which seems to be a physiological problem) is often characterized in chemical terms. On the other hand, one often needs knowledge of physiological modes of organization to explain the chemical activities occurring intracellularly. Physiological and chemical investigations are not distinguished by the level at which they are performed, but by the questions asked and research techniques used. A similar view may be called for in the case of cognitive and neuroscience inquiries. These inquiries differ not in the level of phenomena investigated but in questions investigated and techniques used. What emerges is a picture of a hierarchy of levels with contributions from both neuroscience and cognitive science occurring at various points in this hierarchy. The inquiries from the different disciplines can clearly guide each other as to the nature of the processes occurring at various levels. However, in allowing this interweaving of inquiries, one does not have to worry about the possible eliminative reduction of one inquiry to another. One has a hierarchy of levels and a variety of techniques to study activity at each level in the hierarchy.

Notes

[1]I thank Adele Abrahamsen, James Frame, Robert Richardson, and Robert McCauley for their very helpful comments on earlier drafts of this paper and many discussions of these topics. Work on this paper was supported by the National Endowment for the Humanities and a Georgia State University Research Grant.

[2]As a bridge, the functional architecture may be more useful than Pylyshyn or Newell tend to suggest. For it provides a schema that shows how neuroscience and cognitive inquiries can guide each other. Newell suggests that knowledge that the neurological system must instantiate a symbol system "is a genuine prediction on the structure of the nervous system and should ultimately inform the attempt to understand how the nervous system functions" (Newell 1980, p. 174). Presumably the prediction Newell thinks can be made is that one should find in the nervous system a design that makes symbol processing possible. One should expect that neuroscience may provide similar guidance for cognitive investigations. For example, despite the fact that Pylyshyn carries out his inquiry totally on the cognitive side of the bridge, relying only on the impenetrability criterion for distinguishing the basic operations of the virtual machine in the mind, it is conceivable that neuroscience could point the way to these invariant processes. Thus, neuroscience may point the way to the basic cognitive capacities which the mind can employ. However, in providing such guidance neither inquiry would threaten to subsume the other. The two inquiries would be directed toward different tasks and operate at different levels of organization. Neuroscience would be charged with investigating the physiological processes that provide the basic capacities specified in the functional architecture while cognitive science would explore the ways in which these units interact to produce actual behavior. These inquiries could constrain and guide each other without one subsuming the other. I have developed this view further in Bechtel (1983 1984).

[3]A current problem of great interest is how genes maintain their regulatory system. Some theorists have seen this as a phenomenon that requires explanation in terms of selection forces. Kauffman (1986) has made a proposal quite like that of the PDP theorists--he proposes that there is no specific mechanism maintained by selection but rather there are emergent stable configurations in an interactive genetic ensemble.

[4]See McCauley (in press) for some reservations about the interpretation Barsalou offers of his results.

[5]For an AI proposal within this general tradition that is even more motivated by neuroscience see Gigley (1983).

[6]In trying to account for the brain's advantage over conventional computers in dealing with various cognitive tasks, Rumelhart appeals to the brain's architecture: "What then is the brain's advantage? I suspect that this lies in the kind of computation the brain is able to carry out. Primarily, the brain succeeds because it has an enormous number of processing units all working in parallel and cooperatively <u>settling</u> into a solution--rather than <u>calculating</u> a solution. Processing is done by cooperating coalitions of independent units each working on the information made available to it. It is as if

computation were done by having each little processor carry out its small computation and then vote on the answer to the question. Solutions are reached by majority rule, or by reaching a compromise." (Rumelhart 1984, p.61).

[7]In Bechtel (1985a), I have discussed this capacity of PDP models to model concept learning, which has been a difficulty for traditional information processing models but was a focal point of behaviorist research. I thus proposed PDP models as a basis of reconciliation between cognitive theorists and behaviorists.)

[8]Lycan (1981 and in preparation) defends such a multi-level view of mental activity, defending it as providing answers to some of the traditional objections to functionalism. A PDP type analysis of some of the lower levels in such a hierarchy would seem quite compatible with his view.

https://doi.org/10.1086/psaprocbienmeetp.1986.1.193117 Published online by Cambridge University Press

References

Barsalou, Lawrence W. (in press). "The Instability of Graded Structure: Implications for the Nature of Concepts." In Concepts Reconsidered: The Ecological and Intellectual Bases of Categories. Edited by Ulric Neisser. Cambridge: Cambridge University Press.

Bechtel, William. (1983). "A Bridge Between Cognitive Science and Neuroscience: The Functional Architecture of the Mind." Philosophical Studies 44: 319-330.

----------------. (1984). "Autonomous Psychology: What It Should and Should Not Entail." In PSA 1984, Volume 1. Edited by P.D. Asquith and P. Kitcher. East Lansing, MI: Philosophy of Science Association. Pages 43-55.

----------------. (1985a). "Contemporary Connectionism: Are the New Parallel Distributed Processing Models Cognitive or Associationistic?" Behaviorism 13: 53-61.

----------------. (1985b). "Realism, Instrumentalism, and the Intentional Stance." Cognitive Science 9: 473-497.

Cottrell, G. W. and Small, S. L. (1983). "A Connectionist Scheme for Modelling Word Sense Disambiguation." Cognition and Brain Theory 6: 89-120.

Churchland, P. M. (1979). Scientific Realism and the Plasticity of Mind. New York: Cambridge.

Darden, Lindley. (1986). "Relations Amongst Fields in the Evolutionary Synthesis." In Integrating Scientific Disciplines: Case Studies from the Life Sciences. Edited by William Bechtel. Dordrecht: Martinus Nijhoff.

Dennett, D.C. (1978). Brainstorms. Cambridge: MIT Press (Bradford Books).

Fahlman, S.A. (1979). NETL, A System for Representing and Using Real Knowledge. Cambridge: MIT Press.

Feldman, J.A. and Ballard, D.H. (1982). "Connectionist Models and their Properties." Cognitive Science 6: 205-254.

------------. (1983). "Introduction: Advanced Computational Models in Cognition and Brain Theory." Cognition and Brain Theory 6: 1-3.

Fodor, J. A. (1974). "Special Sciences (or: The Disunity of Science as a Working Hypothesis)." Synthese 28: 97-115.

-----------. (1980). Representations. Cambridge: MIT Press (Bradford Books.)

https://doi.org/10.1086/psaprocbienmeetp.1986.1.193117 Published online by Cambridge University Press

Gigley, H. M. (1983). "HOPE--AI and the Dynamic Process of Language Behavior." Cognition and Brain Theory 6: 39-88.

Grodzhinsky, Y.; Swinney, D.A.; and Zurif, E.B. (1985). "Agramaticism: Structural Deficits and Antecedent Processing Disruptions." In Agrammaticism. Edited by M. L. Kean. New York: Academic Press.

Hinton, G.E. and Sejnowski, T.J. (1984). "Learning Semantic Features." Proceedings of the Sixth Annual Conference of the Cognitive Science Society. Boulder, Colorado. Pages 63-70.

Hull, David. (1974). Philosophy of Biological Science. Englewood Cliffs: Prentice Hall.

Kauffman, S. (1986). "A Framework to Think about Evolving Regulatory Systems." In Integrating Scientific Disciplines: Case Studies from the Life Sciences. Edited by William Bechtel. Dordrecht: Martinus Nijhoff.

Lycan, William G. (1981). "Form, Function, and Feel." The Journal of Philosophy 78. 24-49.

----------------. (in preparation). Consciousness. Cambridge: MIT Press (Bradford Books).

McCauley, Robert M. (in press). "The Role of Theories in a Theory of Concepts." In Concepts Reconsidered: The Ecological and Intellectual Bases of Categories. Edited by Ulric Neisser. Cambridge: Cambridge University Press.

McClelland, J.L. and Rumelhart, D.E. (1981). "An Interactive Activation Model of Context Effects in Letter Perception. Part 1, An Account of Basic Findings." Psychological Review 88: 375-407.

Newell, Allen. and Simon, Herbert A. (1976). "Computer Science as Empirical Inquiry: Symbols and Search." Communications of the ACM 19: 123-126.

-------------. (1980). "Physical Symbol Systems." Cognitive Science 4: 135-183.

O'Keefe, J. and Nadel, L. (1978). The Hippocampus as a Cognitive Map. Oxford: Clarendon.

Putnam, Hilary. (1975). Mind, Language, and Reality. Philosophical Papers, Volume 2. Cambridge: Cambridge University Press.

Pylyshyn, Zenon W. (1980). "Computation and Cognition: Issues in the Foundations of Cognitive Science." Behavioral and Brain Sciences 3: 111-132.

-----------------. (1984a). Computation and Cognition. Toward a Foundation for Cognitive Science. Cambridge: MIT Press (Bradford Books).

----------------. (1984b). "Why 'Computing' Requires Symbols." Proceedings of the Sixth Annual Conference of the Cognitive Science Society, Boulder, Colorado. Pages 71-73.

Rumelhart, D.E. and Norman, D.A. (1982). "Simulating a Skilled Typist: A Study of Skilled Cognitive-Motor Performance." Cognitive Science 6: 1-36.

-------------. (1984). "The Emergence of Cognitive Phenomena from Sub-Symbolic Processes." Proceedings of the Sixth Annual Conference of the Cognitive Science Society. Boulder, Colorado. Pages 59-62.

-------------- and McClelland, J.L. (1986). Parallel Distributed Processing: Explorations in the Microstructures of Cognition, Volume 1: Foundations. Cambridge: MIT Press (Bradford Books).

https://doi.org/10.1086/psaprocbienmeetp.1986.1.193117 Published online by Cambridge University Press