Embodied computation and the autonomy of psychology

Paco Calvo
fjcalvo@um.es

Departamento de Filosofía
Edificio Luis Vives
Campus de Espinardo
Universidad de Murcia
E-30100 Murcia - SPAIN

John Symons
jsymons@utep.edu

Department of Philosophy
University of Texas at El Paso
El Paso, TX 79968 USA

Abstract How do we characterize the computational challenges facing a control system as it controls the behavior of an agent? Part of the answer to this question depends on how we understand the capacities of the agent's body. These capacities, along with the relevant features of the social and physical environment in which it is embedded, serve as the parameters which determine the problems faced by the agent's central control system. In this article, we consider the role of embodiment in computation. We argue that taking embodiment seriously, reorients our thinking with respect to the classical hierarchy of computational, algorithmic and implementational levels of analysis.

1. Introduction

Following David Marr's methodological example, traditional cognitive science has rested on the assumption that the first step in the investigation of mind involves building an information-processing model. This model takes the form of a solution to a well-defined set of information-processing (or computational) problems that the researcher believes the system under consideration must overcome. These problems should, at least initially, be studied in-themselves, apart from any real consideration of embodiment.

Historically, the bias in favor of the functional or computational level of investigation over the evidence of neuroscience did not originate with Marr. It dates back at least to the classic early works in cognitive psychology and is clearly articulated in one of the most influential early works of that tradition, Ulric Neisser's *Cognitive*

Psychology (1967). The dismissal of biological evidence rested, almost from the start on the analogy between minds and computer programs, as Neisser writes:

> The task of the psychologist trying to understand human cognition is analogous to that of a man trying to understand how a computer has been programmed. In particular, if the program seems to store and reuse information, he would like to know by what "routines" or "procedures" this is done. Given this purpose, he will not care much whether his computer stores information in magnetic cores or in thin films; he wants to understand the program not the "hardware." By the same token, it would not help the psychologist to know that memory is carried by RNA as opposed to some other medium. He wants to understand its utilization, not its incarnation. (1967, 6)

This radical denial of biology was not tempered by developments in neuroscience. For example, it is common to read more recent works in cognitive science arguing that:

> Information processes can and must be studied without reference to the physics or biology of the system carrying them out… the algorithm for multiplication can be studied and understood in interesting ways without any knowledge of how it is implemented by neurons in the human brain… In general an important clarity of understanding is attained by separately considering an information process as a formal symbol manipulation process and as a physical or biological process. (Stillings et al. 1987, 4-5)

Even in the 1990s one could still read prominent cognitive scientists arguing along similar lines. In his widely-used textbook, *Mind: An Introduction to Cognitive Science* (1996) Paul Thagard lists 'Neurobiological Plausibility' as the fourth of five criteria for evaluating theories in cognitive science. (1996, 13) In his judgment, only the 'Practical

Applicability' of a theory to educational techniques and design ranks as a less important criterion.

Until relatively recently, workers in cognitive science and the philosophy of mind have consistently portrayed the mind as a symbol manipulating device whose relationship to the body is contingent and ultimately irrelevant. As a result, the information processing approach has encouraged the illusion that cognitive systems constitute a natural kind which can be studied in itself, apart from considerations of embodiment. Information processing models serve to vindicate the claim that psychology constitutes an autonomous science. Even in the case of views which incorporate biological insights, for example in connectionism, computational functionalism continues to support the view that psychology is an autonomous science.

2. Natural kinds, connectionist architectures and the autonomy of psychology

Cognitive systems are thought to constitute a natural kind. The result is the non-reductive pairing of psychological laws with sets of underlying physical mechanisms (Fodor, 1974). While Fodorian arguments are often identified as the principal motivating factors for anti-reductionism, one can retain the autonomy of psychology without endorsing a Fodorian conception of cognition. Even granting a connectionist architecture of the sort challenged by Fodor and Pylyshyn (1988), statistical techniques such as cluster analysis support the view that different networks can be understood to define a *unique class* by dividing their respective hidden spaces into significantly similar sub-spaces. So long as the relevant information is somehow implicit in whatever sensor-input schemes happen to be employed, and so long as the training procedures impose the same requirements on performance, then diverse networks can settle into almost identical abstract organizations. Research by Laakso and Cottrell

(2000) shows how different partitions of hidden space across networks can still reveal a *common* strategy despite a degree of functional divergence inherent to connectionist processing.[1] We may thereby assign genuine kinds to different networks regardless of their architectural idiosyncrasies. Cognitive systems, in short, may well form a natural kind, and the special character of psychology as an autonomous science may thus be vindicated.[2]

Reductionists may nevertheless wish to argue that neural networks still fail to constitute a natural kind. The issue boils down, the argument would run, to the question of whether clusters are causally efficacious in the production of the network's patterns of input/output behavior. What gets activated, at each step of processing, is a component of the cluster, not the cluster itself. Uniting some of these hidden states should not drive us to think that the network actually employs the sort of entities that psychology posits. In short, *post hoc* statistical techniques of the sort deployed in the neuromodeler's methodology are causally inert, thus failing to play any explanatory role as far as the dynamics of connectionist networks is taken as our model of cognition.

1 The upshot of Laakso and Cottrell's neurosimulation is that different dimensionality, architecture or encoding bring no trouble, insofar as correlated distances between points in the respective spaces are preserved. Churchland (1998) has further exploited Laakso and Cottrell's measure in his defence of state space semantics, although see Calvo (2000) for problems with Churchland's conception. Also, the fact that statistical analyses can highlight higher-level commonalities may be called into question, at least in the case of simple recurrent networks (Calvo, 2003), but we shall ignore these matters for present concerns.

2 We are assuming a connectionist architecture for the sake of argumentation, granting that it is an open empirical question what the correct architecture of cognition happens to be. Also, for purposes of illustration, we shall concentrate upon psychology as a special autonomous science. The bearing of the line of argument to be developed in due course may nonetheless be more far-reaching, once the alleged special character of other higher-level disciplines is considered.

While the threat of epiphenomenalism is a serious metaphysical challenge, it is unlikely to resonate with cognitive scientists. Clark (1995) makes a pragmatic appeal to practitioners of the special sciences, arguing that denying the causal efficacy of psychological posits risks falling into microphysics worship: "Unless the sceptic is willing to give up the causal efficacy of chemical properties too, he or she would be unwise to object to the very idea of higher-level constructs figuring in genuine causal claims" (p. 350). It is not clear to us, however, that the choice is so stark. Clearly, Clark's argument is not intended to force us into the dilemma of:

(a) accepting all higher-level causal claims as genuine,

or

(b) rejecting all higher-level causal claims

The challenge for philosophers is to determine whether causal agency can be located at the macro-level in certain cases (say, biology or chemistry), and not in others (say, psychology or economics). There is an ongoing debate in metaphysics concerning the problem of higher-level causation and our criteria for individuation at the macro-level. While we would insist that Clark presents a false dichotomy with respect to the causal agency of psychological posits, the causal agency of the posits is not the critical issue with respect to their explanatory role in revealing the coarse-grained nature of cognition. The notion of causal explanation can be dissociated from the requirement of causal efficacy. Following Jackson and Pettit (1988), extended functionalists including Clark, have relied instead on the distinction between program and process explanations. A program explanation is any high-level explanation that, while gathering a range of cases in terms of certain macrofeatures being shared, abstracts away from the actual

microfeatures which carry the burden of causation. Following that type of explanation, those common macrofeatures are said to 'causally program' a given pattern of behavior, without actually being part of the causal explanation of that behavior. By contrast, an example of a process explanation is an explanation that picks out the microfeatures that are causally efficacious. We may then say that cluster analysis causally programs the network's performance, although it does not play any role as part of the process explanation of the behavior of the network.

Reductionists may respond with a revamped version of the causal/explanatory pre-emption argument, and the Fodorian will surely counter in her turn. Given the terms in which this debate has been formulated, we see the possibility that it simply devolves into a stalemate. In what follows, we shall try to cast a different light upon this set of issues by examining the implications of embodied computation for our characterization of the autonomy of psychological explanation.

3. Generalization in connectionist networks

Instead of relating connectionist clusters and the issue of their causal efficacy to the putative status of cognitive systems as natural kinds, we may consider talk of clusters in a network to the problem of determining the capacity of a connectionist network to generalize. Shea (2007) considers the kinds of entities which can serve as vehicles of content in a neural network. His answer is that unless clusters are the vehicles of content, we will not be able to explain generalizations. He claims, "By activating a cluster the network represents that the new sample has the property common to the training samples in the cluster. This correct intermediate classification is part of the process by which the network makes a correct output classification." (p. 255).

In our view, the typing of clusters as the vehicles of content fails to explain *how* neural networks do actually generalize. In fact, Shea's account amounts to a Fodorian only-game-in-town sort of move: *How would a network be able to generalize after all if it is not by means of exploiting something projectible that is being shared, some common property, between novel and training exemplars?* But notice that there is no positive evidence that some shared property is realistically projected. We merely exploit as an inference to the best explanation that something *must* have been shared, if we are to grasp how a network can extend its competency beyond the training regime. However, if Shea's claim that "the new sample has the property common to the training samples in the cluster" implies that we accept the thesis that the predicates of the special sciences are projectible (cf. Gillett, 2007), an argument is needed in order to show why that is the case. Unfortunately, all we have is the negative only-game-in-town line of argument.

So, what do we really mean when we claim that a neural network is able to generalize? In our view, rather than the activation of a cluster, the critical element in explaining generalization capabilities is the learning algorithm itself. The ability to generalize is a function of the similarity between the function acquired during training and the function that underlies the test set. But for the learning algorithm to furnish the network with the appropriate connection weight adjustments, the global functional profile is irrelevant. Note that all that matters is locally accessible to each single unit during the process of weight adjustment.

It is noteworthy that our reason for denying that clusters can furnish us with the appropriate level of description stems from the lack of shared properties to be typed as

the vehicles of content, and not with their alleged lack of causal efficacy.[3] In the case of connectionist theory, it is the choice of clusters as the vehicles of content, we contend, which continues to support the intuition that psychology is autonomous. However, as we shall see in the section which follows, an analysis of the relationship between the computational and the algorithmic level of description will help us see why positive evidence for the kind of commonalities the functionalist requires is unavailable.

4. Disembodied computation and algorithm

Cognitive scientists usually first provide a computational level description of the cognitive capacity they try to account for. Only later, once the target has been fully characterized in computational terms, they ask how it is carried out algorithmically. Put bluntly, the problem is that, in doing so, the computational profile is carried along the way as a constraint on our discussion at the next lower algorithmic level. Whereas discussion focuses upon the deployment of algebraic or statistical resources (see for example, Marcus, 2001), the computational level of description remains unaltered throughout the whole process (for example, the description of the ability to output the past tense of English verbs). Algorithmic level manipulations on internal representations are said to be appropriate insofar as the cognitive level description of the capacity to be accounted for is matched. For connectionists, cluster analysis is thought to allow us to gauge whether the algorithmic level matches the requirements stated at the computational level.

3 Lack of causal efficacy is precisely what has driven O'Brian and Opie (2006) to reject clusters as types in favour of the local weight matrices that modulate the activity of each particular unit, what they dub *fan-ins*.

This asymmetry between the demands at the computational and algorithmic levels supports the anti-reductionist illusion that something fixed, namely, the abstract profile at the computational level, can be multiply realized, and that something is actually being shared at the algorithmic level of description. It is a short step from here to functionalism and the vindication of the special character of psychology. An example of how this strategy can be successful is the extraction of algebraic rules that represent relationships between placeholders (Marcus, 2001). But let us instead illustrate the point with a paradigmatic example, the approximation of Boolean functions for reasons that will become obvious in section 5, below.

The acquisition of a Boolean function gets specified computationally as a truth table for all possible input-output pairings. This remains fixed. We can then discuss how a system does approximate the function; namely, whether the job gets done symbolically, via inferential processing sub-routines at the subpersonal level, as opposed to subsymbolically, via the local adjustment of connection weights. One way or the other, we know what we are after. The problem, however, with this rendering of the situation is twofold.

(i) It is not clear what it means to say that the function *has been approximated*.

(ii) It is not clear what it means to say that *the function* has been approximated.

Since, as described above, the computational level remains fixed, the reader can see that (ii) will become the focal point of our concern. We will introduce (ii) by first saying something very briefly about problem (i).

As to the issue of when a computational function has been successfully approximated, we may ask somewhat more technically how we decide when a network has mastered a given task. Notice that were we to use for instance the root mean square (rms) error, we could not say that the network has learned the mapping when the error is zero. There will always be a residual error since the sigmoid activation function never reaches 0 or 1. One solution widely employed (Plunkett and Elman, 1997) is to round off the output activation values in such a way that if the target is 1, then activations close to 1 are taken to be correct, and the same goes for 0 as target and activations close to 0.

Philosophically, rounding off *itself* constitutes something of a problem. Rounding off can be seen as involving an arbitrary criterion for determining correctness. In scientific practice we may be able to ignore this worry. The more pressing challenge at the philosophical and scientific level concerns the perspective from which the rounding off takes place. We contend that using a target of 0/1 exemplifies the influence of the disembodied computational perspective. As a result, by rounding off we fail to realize that nothing is actually being shared when we reach the decision that two networks have mastered the task. From an embodied perspective, as we will try to illustrate below, things look very different. Let us then turn to our second problem. Namely, what do we mean when we say that *the function* has been approximated? An example of embodied XOR morphological computation will show us why the autonomy of psychology may be untenable as we learn more about how agents actually tackle the computational challenges that they encounter in their environments.

5. Computing with bodies

In one sense it is obvious that the computational demands facing the control system for a robot body will vary according to the abilities of that body. To begin with, it is well known that the physical configuration of the robot body can reduce the computational burden on its central controller. Rodney Brooks (1991) and Valentino Braitenberg (1984) showed how bodies or machines can overcome complicated behavioral challenges by virtue of their mechanical structure alone. For example, Braitenberg's vehicles are robots whose motors and sensors are connected so as to exhibit behaviors (such as tending to move towards or away from a light source) which are goal directed without the need for any explicit central controller or information processor to determining the body's movement. Brooks and Braitenberg have interpreted these systems as exemplifying ways that the physical structure of a body can reduce the computational burden on the central controller.

In the case of an organism with a relatively limited behavioral repertoire, it is easy to imagine an arrangement whereby all evolutionarily relevant computational challenges are resolved via the physical structure of the body itself. For agents with more complicated sets of behaviors and challenges, it becomes more reasonable to presuppose the involvement of some kind of central control system. However, where one draws the line between the central controller and the body in computational problem solving is not obvious. Without carefully determining the capacities of the body, it remains an open question just how much of the computational burden with respect to some problem is being assumed by the central controller.

Chandana Paul (2004) describes what she calls morphological computations; computations where the mechanical structure of the body itself carries the burden of solving the computational problem. A simple illustration is her morphology-based

treatment of the Boolean function XOR. Paul (2004) designed a robot which, by virtue of its morphology, can exhibit XOR-constrained behavior. The XOR function is not performed by the robot's neural networks (as will be explained below) instead, the agent's body resolves the task of performing the XOR operation.

The significance of the XOR function is connected to the history of Rosenblatt's development of the perceptron. A perceptron is a neural network which takes multiple inputs and gives one output. Rosenblatt describes the perceptron as consisting of three layers: The input layer, the layer of association units, and the output layer. The output of the middle layer is a threshold weighted sum of the inputs. The network has adjustable synaptic weights and fires once it reaches some value determined by its threshold. Rosenblatt proved the perceptron convergence theorem which states that a perceptron can compute any linearly separable function. Historically, the inability of perceptrons to solve functions which are not linearly separable was one of the reasons that popular scientific opinion turned against neural networks for a time.

Put very simply, a linearly separable problem is one which for any output neuron there must be some hyperplane (of dimension n-1) which divides the set of n inputs to that neuron into those which activate the output neuron and those that do not. So, if there are two inputs with values of either 0 or 1, given that the problem is linearly separable, there should be a one-dimensional hyperplane (a line) which partitions the set of inputs into a set consisting exclusively of 0s and another consisting exclusively of 1s.

In the case of Boolean operations like AND or OR, there is linear separability. This is easy to see, given a truth table style representation of these operations.

OR

Input 1	Input 2	Output
1	1	1
1	0	1
0	1	1
0	0	0

Table 1

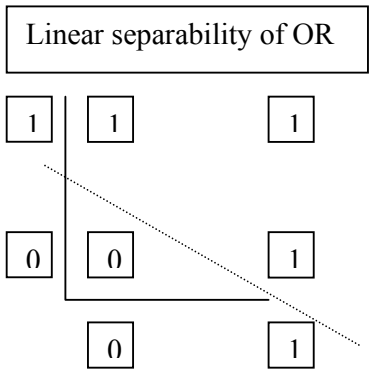

Figure 1

By contrast, functions like XOR or XNOR are not linearly separable as can be seen from the figures below:

XOR

Input 1	Input 2	Output
1	1	0
1	0	1
0	1	1
0	0	0

Table 2

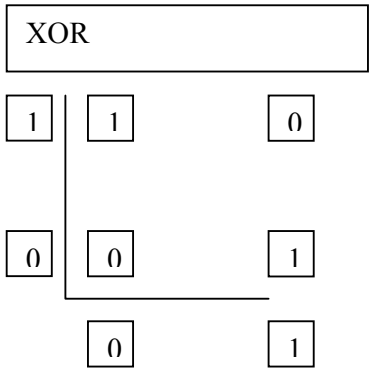

Figure 2

XOR is not linearly separable as can be seen from the figure above. No straight line will partition the space of values into uniform sets of 0s and 1s. Given that functions like XOR and XNOR are not linearly separable, they fall beyond the computational capacity of the perceptron.[4]

4 Although of course it is now possible to construct neural networks which are not subject to this constraint and can solve non-linearly separable problems.

The robot which Paul describes is controlled by two perceptrons, one for a motor M1 that computes OR (figure 3), the other for a second motor M2 that computes AND. Whereas M1 turns an existing single wheel permitting forward motion, M2 serves to lift the wheel off the ground. By means of these two simple perceptron-governed patterns of behavior, the robot is able to perform the XOR function (fig. 3).

A	*B*	Behavior
F	F	stationary
F	T	moving
T	F	moving
T	T	stationary

Figure 3: From Paul, C. (2004), p. 33

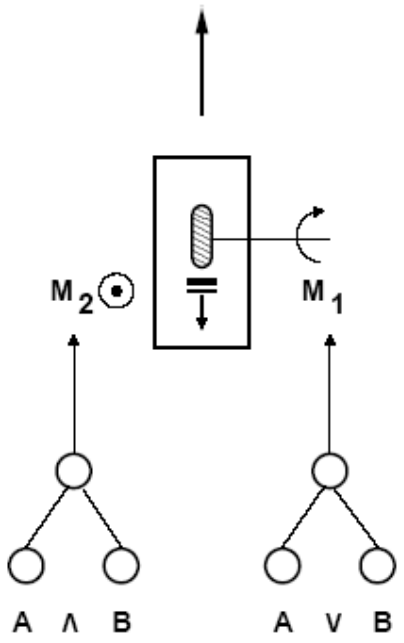

Figure 4: The XOR Robot. This robot has one wheel, with two actuated degrees of freedom. The motor M1 is responsible for turning the wheel so that the robot moves forward. The motor M_2 is responsible for lifting the wheel off the ground. Each motor is controlled by a separate perceptron network, which takes as inputs A and B. M1 is controlled by a network which computes A OR B,

and M_2 by a network which computes A AND B. Using only these controllers, the robot is able to display the XOR function in its behavior. From Paul, C. (2004), p. 34

Since perceptrons by themselves cannot compute the XOR function, her robot relies on the structural features of its body to perform the XOR computation. Furthermore, the structure of the robot provides a way to compute the NOT function which, when followed by the AND and OR functions, becomes computationally equivalent to the standard resolution of the XOR problem with a three-layer feedforward network. The difference here is that it is the body, and not another layer of connectionist processing, which solves the problem. As Paul (2004) points out, the body is not merely an effector but part of the "cognitive" equation itself.

For reasons which we will explain below, let the robot's morphological computation be called the XOR^* function as distinct from the straight XOR Boolean function.

6. Implications for functionalism.

Paul's description of morphological computation provides a precise example of the kind of embodiment which proponents of the embedded-embodied approach in philosophy of mind seek (Calvo and Gomila, in print). Obviously, the mere existence of morphological computation does not, by itself, present a significant challenge to the computational functionalist orthodoxy in the philosophy of mind. Functionalists can argue that morphological computations are just as multiply realizable as central computations albeit within the context of an "extended functionalist" framework. For some advocates of embodied cognition, like Andy Clark, morphological computation is fully compatible with traditional multiple realizability. Different configurations of the body provide different extended ways of implementing *the same* function. In the case of

morphological computations, one might imagine different ways in which the structure of the body could supplement the perceptrons, computationally speaking.

This rendering of the situation however lays the stress upon the different ways in which an algorithm can be implemented. And by focusing upon the implementational and algorithmic levels of description, traditional cognitive science has tended to assume the fixity of the function at the computational level of description.

By reading Paul's example through the lens of radical (non-functionalist) embodiment, things look rather different: According to the perspective we advocate, the unique morphological features that the XOR* robot exemplifies, are constitutive of the computational challenge facing the agent as well as the behavioral repertoire which it exhibits in response.[5] Computing XOR* morphologically is not the same thing as modeling the approximation of the XOR function with a disembodied network. As the robot agent navigates through its environment, the details of its embodiment and the range of feasible physical interactions available to it do matter. In fact, to see how dramatically details external to the cognitive system itself can affect the overall result, consider the environment as well as an extension of the system and place the XOR* robot downhill. Clearly the function, as described at the computational level, is altered and changes *as a whole* (Paul, 2004). Hence XOR* rather than XOR.

The disembodied perspective focuses upon the inner workings of the cognitive architecture, ignoring bodily and environmental constraints. However, once these are acknowledged, our computational accounts of the agent's challenges and responses are subject to revision. Like the body, the environment is part and parcel of cognition, and having changed the computational function itself, failure to account for generalizations is guaranteed unless morphology is considered as well.

5 Noë (2004) and Shapiro (2004) advocate similar positions in the philosophical literature.

But what is the relation between morphological computation and the dismissal of functionalism? Notice that morphological computation is still compatible with the rounding-off procedure discussed above. In most cases, constraints from embodiment are such that, from the higher level point of view, the error does not make a difference in terms of the overall behavioral output (say, "stationary" vs. "moving" in figure 3). Take for instance the pocketing of a ball in a game of pool. Say that you pocket it twice, first by striking the cue ball with a nice solid hit, and second somewhat more loosely. Insofar as the two sequences end up with the pocketing of the ball, the function (understood in terms of an abstract computational problem and solution) can be said to have been approximated. But this abstract characterization of the problem is precisely what encourages the illusion that something is being shared between the two events. The price we pay is a failure to realize that the extended cognitive function itself may in fact have changed. Disembodiment, we contend, obscures principled ways of distinguishing between pairs of functions.

One is happy with the cognitive level description when things work, say, when the two balls get sunk in the pocket. But this is equivalent to the rounding-off method used in evaluating a network's performance. The ball ends up in the pocket, barring some fluke events! But when things go wrong we need to inquire further in order to understand what has happened.

The reader may put it in terms of mechanisms (Bechtel and Richardson, 1993). If we have a cognitive level function to be matched algorithmically, thinking of networks that implement the function better or worse would be equivalent to having a mechanistic explanation where, once the components of the mechanism are identified in relation to the phenomenon in need of explanation, we ignore different tokenings of the

phenomenon, and hold that the very same components and their spatiotemporal organization are preserved, come what may.

When things work we assume something is being shared. It is only when predictions fail that we need to understand the system as a whole. In fact, a disembodied prediction of what the XOR* robot will do drives us to say that the robot is in error if for example it moves forward and downwards when the actuators are off, but once we honor a morphological level of description the former goal misrepresents the situation. It shows that nothing is being shared after all. If embodiment, writ large, matters non-trivially, many environmental and corporal changes (say, placing the robot downhill) will affect the function that describes the embodied cognitive system as a whole.

Once morphological computation is interpreted along the lines we suggest, psychology is not *special* anymore. Details of bodily implementation and environmental coupling not only matter, but are what ultimately explain the behavior of the system.

One line of objection might be that the implementational level does not provide sufficient evidence or criteria by itself for the individuation of psychological kinds. The friend of embodiment *and* reductionism can reply that the relevant pairing is not between implementation and algorithm, but rather between algorithm and computation. From this perspective, the vehicles of content are individuated in a purely instrumental or at best ad hoc manner. This is illustrated in our discussion above of the rounding-off criterion that modelers employ as they maintain the fixity of the computational level function as they try to spell out the underlying algorithm. Advocates for the autonomy of psychology have failed to appreciate this point.

References

Bechtel, W. and Richardson, R. C. (l993) *Discovering Complexity: Decomposition and Localization as Strategies in Scientific Research. Princeton*, NJ: Princeton University Press.

Braitenberg, V. (1984) *Vehicles: Experiments in synthetic psychology*, Cambridge, MA: MIT Press.

Brooks, R. A. (1991) "Intelligence without Representation", *Artificial Intelligence* 47: 139-159.

Calvo, P. (2000) "State space semantics and conceptual similarity: reply to Churchland", *Philosophical Psychology* 13: 77–95.

Calvo, P. (2003) "Connectionist Semantics and the Collateral Information Challenge", *Mind and Language* 18: 77-94.

Calvo, P. and Gomila, T. (eds.) (in print) *Handbook of Cognitive Science: An Embodied Approach*, Elsevier.

Churchland, P (1998) "Conceptual similarity across sensory and neural diversity: the Fodor/Lepore challenge answered", *Journal of Philosophy* 95: 5–32.

Clark, A. (1995) "Connectionist Minds", in G. and C. MacDonald (eds) *Connectionism: Debates on Psychological Explanation*, Blackwell: Oxford, 339-356.

Fodor, J. (1974) "Special sciences and the disunity of science as a working hypothesis", *Synthese* 28: 77-115.

Fodor, J. and Pylyshyn, Z. (1988) "Connectionism and Cognitive Architecture", *Cognition* 28: 3-71.

Gillett, Carl (2007) "Understanding the new reductionism: The metaphysics of science and compositional reduction", *Journal of Philosophy* 104 (4): 193-216.

Jackson, F. and Pettit, P. (1988) 'Functionalism and Broad Content', *Mind* 97: 381-400.

Laakso, A. and Cottrell, G. (2000) "Content and cluster analysis: assessing representational similarity in neural systems", *Philosophical Psychology* 13(1): 47-76.

Marcus, G. (2001) *The Algebraic Mind*, Cambridge: MIT Press.

Marr, D. (1982) *Vision*: A Computational Investigation into the Human Representation and Processing of Visual Information, New York: W.H. Freeman.

Neisser, U. (1967) *Cognitive Psychology*, Englewood Cliffs, NJ: Prentice-Hall.

Noë, A (2004) *Perception in Action*: MIT Press.

O'Brien, and Opie (2006) "How do Connectionist Networks Compute?", *Cognitive Processing* 7: 30-41.

Paul, C. (2004) "Morphology and Computation", Proceedings of the International Conference on the Simulation of Adaptive Behavior Los Angeles, CA, USA: 33-38.

Plunkett, K. and Elman, J. (1997) *Exercises in Rethinking Innateness*, Cambridge, Mass.: MIT Press.

Rosenblatt, F. (1959) *Principles of Neurodynamics*, New York: Spartan Books.

Shapiro, L. (2004) *The Mind Incarnate*, Cambridge, Mass.: MIT Press.

Shea, N. (2007) "Content and Its Vehicles in Connectionist Systems", *Mind & Language* 22: 246-269.

Stillings, N.A. et al. (1987) *Cognitive Science: An Introduction*, Cambridge, Mass.: MIT Press.

Thagard, P. (1996) Mind: An Introduction to Cognitive Science, Cambridge, Mass.: MIT Press.