5

Computing with Bodies: Morphology, Function, and Computational Theory

John Symons and Paco Calvo

Contemporary philosophers inherit an anti-psychologistic tradition. The central figures in the early history of both the continental and analytic movements opposed what they saw as the encroachment of psychologists and their fellow travelers on the territory of philosophers (see Kusch 1995 and Dummett 1993).† Most prominently, both Frege and Husserl argued that we should avoid corrupting the study of thought with psychologism. As they understood it, psychologism is the view that the best way to understand thought is to look to the empirical study of what we (or our brains) happen to do when we're thinking. Thought itself, on their view should be understood apart from the empirical investigation of mind, let alone the study of the gory details of the brain and nervous system.

The emergence of the computational model of mind in the 20th century, and specifically the conceptual distinction between structural and functional properties of computational systems seemed to provide a non-reductive account of mind, and computationalism proved quite compatible with the anti-psychologistic tendencies of mainstream analytic philosophers. Treating pain, belief, desire, etc. as functionally individuated concepts allowed philosophers to resolve the tension between anti-psychologism and a commitment to the progress of empirical science. Computationalism maintained the autonomy of mental properties as functional states of a physical system, thereby protecting our commonsense understanding of mental life from revision at the hands of the empirical sciences. Since computational functionalism is usually presented as a non-reductive *physicalist* theory, adherents get autonomy for their account of mental life, without being forced to commit to what they regarded as a metaphysically problematic form of ontological dualism. In this way, computationalism supports autonomy without abandoning broader scientific principles. This, at least, is the way the story is usually told.

In this paper we examine some of the implications of our understanding of the body for computational theory. We shall argue that one of the most important constraints on computational theorizing in the study of mind is the initial determination of the challenges that an embodied agent faces. One of the lessons of an area of robotics known as morphological computation is that these challenges are inextricably linked to an understanding of the agent's body and environment. The simple conceptual point is that computational theory is part of the science of mind insofar as it serves to answer questions about the kinds of minds that we and other creatures have. David Marr, one of the most important early figures in Cognitive Science and a pioneering figure in the computational theory of vision, provided an apt summary of the questions that computational theory works to answer:

> Computational Theory: What is the goal of the computation, why is it appropriate, and what is the logic of the strategy by which it can be carried out? (Marr 1982, 25)

Contrary to the way Marr's view of computational theory is sometimes read, and contrary to many of his other methodological claims, we suggest that these questions point to the study of the body and the environment as indispensable parts of computational theorizing in the science of mind.[1] Answering these questions, we argue, requires a detailed understanding of the agent's body and environment. We differ with traditional computationalist philosophers of mind with respect to the starting point of inquiry. On our view, computational theorizing must begin with an understanding of the challenges faced by the agent. It also involves understanding how the morphological characteristics of the body can serve as solutions to computational problems. Knowing how the body solves problems will be part of what a computational theory tells us about an agent.

Taking this approach to computational theory allows a more precise treatment of the kinds of insights that have been presented by advocates of embodied theories of cognition. For the most part, when such theories figure in philosophical debates, they have had relatively little scientific content. This paper offers a way of understanding how philosophical arguments for the importance of the body and environment can be translated into scientific questions for a computational theory of mind. As we argue here, no good computational theory of mind ignores the body and environment. Pace Hilary Putnam, if we were to discover that we are made of Swiss cheese, it actually would matter for the science of mind.

5.1 Anomalous monism, ordinary language philosophy, and conceivability arguments

In this section we briefly review three lines of argument that run counter to our central claim before turning to morphological computation in detail. We regard each of these arguments as flawed and have addressed some of them directly elsewhere. We do not claim to decisively refute these arguments here, but present them simply in order to explain to non-philosophical audiences why it is even necessary to defend the importance of embodiment in computational theory.[2]

In contemporary philosophy of mind, resistance to the empirical study of the brain and body comes in varying degrees and is motivated by three basic kinds of argument. All of these assume that some protected domain of propositions concerning the mind is not subject to revision. The scope and limits of this insulated domain differs from thinker to thinker, and there may always be ways to construct gerrymandered or increasingly anemic conceptions of mental life that can be protected from revision. Our goal is not to block the possibility that there might be something that is not subject to revision via empirical investigation. Instead, we will begin by examining some of most prominent anti-empirical strategies before comparing them with a revised model of computational theory that takes biological, and more specifically morphological, considerations seriously.

Donald Davidson's well-known anomalous monism arguments (see Davidson 1970 in particular) have been understood to license the claim that psychology has no intrinsic relationship to the so-called 'lower-level' sciences. These arguments have led some philosophers to criticize, in principle, any attempt to connect neuroscientific inquiry with the investigation of mental life. The trouble with anomalous monism is that it seems to leave psychological phenomena disconnected from the causal economy of the physical world. Davidson attempted to answer the charge that contrary to his stated views, anomalous monism leaves mental life epiphenomenal (See for example Davidson 1993), but his later defense of anomalous monism has not satisfied critics. (See for example McIntyre 1999.) Nevertheless, in spite of their apparently unpalatable consequences, Davidson's arguments for the autonomy of psychology have been highly influential among philosophers.

A less influential line of argument derives from ordinary language philosophy. For example, in their *Philosophical Foundations of Neuroscience* (2003), Bennett and Hacker argued that neuroscience is motivated by faulty reductionist presuppositions and simple logical fallacies. Following Wittgenstein, they argue that scientists are simply mistaken when they claim that brains believe, interpret, decide, etc. On their view, only human beings are capable of doing such things. Brains are not. Their sweeping critique of contemporary neuroscience rests on the idea that ascribing information processing capacity to brains involves a confusion of a part (the brain) for a whole (the person). They call this the mereological fallacy.[3] Bennett and Hacker assert that the brain cannot be treated as a possible subject of belief and desire and there is nothing that the brain does the can be predicted on the basis of its beliefs and desires. This contention represents the most extreme form of those views that seek to insulate our understanding of the mind from possible revision. The heart of this argument seems to be something like the following: Since we don't ordinarily talk about brains as having beliefs, desires, and the like, any such talk has violated the norms governing the use of these terms. Since the source of the norms governing the use of the terms is ordinary usage, deviant usage, or attempts to revise that usage, are simply misuses or misunderstandings of those terms. Such a position is implausibly conservative insofar as it seems to assume that ordinary usage fixes the meanings of terms in ways that do not permit revision in light of future evidence. The history of science is replete with examples that would undermine this faith in the authority of ordinary usage.

Conceivability arguments have served as a third strategy for attempting to set the mind apart from the rest of the natural world by reference to one or another of its allegedly essential characteristics. According to proponents, these essential characteristics are *conceivably* separable and therefore *possibly* separable. If they are possibly separable, then they are not necessarily identical to the brain and body. Since ontological identity must be necessary identity, these properties of the mind cannot be identical to the brain. So, for example, many philosophers are convinced of the ontological peculiarity of phenomenal experience because of what they take to be the impossibility of necessary a posteriori identity statements linking minds and bodies. Saul Kripke provided the crucial components of this argument in *Naming and Necessity* (Kripke 1980, 148–155).

While these three lines of argument differ, they all support the view that biological or other empirical evidence can have no direct bearing on those aspects of the mind that they regard as autonomous. Of course, as John Bickle points out, even the most hardcore "autonomist" will concede that some aspects of mental life admit of a biological explanation (Bickle 1998, 1).[4] When philosophers assert the autonomy of the mental, they are usually excepting only a subset of the phenomena we ordinarily associate with mental life from reduction or revision. What anti-neuroscientific philosophers take this genuinely cognitive or phenomenal subset to be varies from thinker to thinker. For Putnam, for example, reason's normative powers fall beyond the purview of neuroscience, while for Fodor and others, our ordinary folk psychological talk of belief, desire, and action stands as an inviolable and unrevisable benchmark for all claims about the mind. For Chalmers, the content of phenomenal judgments will never be reducible to biology, etc. (1996), for these philosophers, biological mechanisms can safely be ignored insofar as their particular sub-domain of the mental is, for one reason or another, autonomous.

While this paper cannot respond directly to all the anti-biological arguments in the literature, our goal is to critically examine the idea that the concepts that we use to talk about mental life – at least when those concepts are characterized in a computational theory – can legitimately float free of our knowledge of the body and the environment. Insofar as computational theory treats the mind in functional terms, it seems like functions should float free of the body. For the most part, the functions that we associate with mental life can (at least in principle) be implemented by a variety of structures – this is the nature of functions. Putnam, the most important philosopher in the early days of computational functionalism, famously emphasized this point when he claimed that as far as our study of mental life is concerned "we could be made out of Swiss cheese and it wouldn't matter" (Putnam 1975, 291). From the perspective of a computational functionalist like Putnam, too much emphasis on neuroscience and embodiment run the risk of confusing the mind's 'software' with its 'hardware.' As we noted above, the hardware-software and structure-function distinctions have held great appeal for philosophers. In his classic papers from the early 1960s, Putnam was one of the first to identify mental life with a set of functions that are contingently implemented in human brains. Mental states, he claimed, are functional states that just happen to be implemented by biological systems but that could just as easily be implemented in a suitably organized machine, cloud of gas, or anthill.

On our view, the most charitable interpretation of this kind of functionalism is that it is premature. If the Swiss cheese approach made sense, then we would have a finished computational theory of mind. However, given that we don't have a finished theory of mind, then we cannot accept the Swiss cheese view, and we will need to attend to the constraints of the body and the environment in order to take the first step towards a functional account.[5] On our view computational theory needs to begin with an understanding of the problems that the mind must solve before we can claim to have a functional characterization of the mind.

Of course, we are free to stipulate that the concepts we use to talk about mental life are descriptively adequate; that they correctly capture the relevant features of their subject matter. This latter assumption would certainly help, a priori, to make them invulnerable to revision. While many philosophers have held the view that our commonsense folk psychology is perfect just the way it is, we take this to be an implausibly strong assumption for healthy and ongoing inquiry in the science of mind. Rather than assuming that we have the story of mind perfectly well under control, we suggest that, at least for the time being, one ought to be committed to the possibility of inquiry into the nature of mind. If we are correct, then given the questions that computational theory should answer, the idea that our psychological concepts are unrevisable is highly doubtful.

5.2 Morphological computation

So, how should computational theory take account of the properties of the body? Thankfully, we do not have to rely on thought experiments to guide our reasoning here. For the remainder of this paper, we discuss an emerging subfield of robotics known as morphological computation. Work in this area has implications for cognitive science and philosophy of mind. In particular, debates concerning the extended mind and embodied cognition can be brought into sharper focus by attending to some relatively simple cases in the morphological computation literature.

One obvious methodological consequence of this work is that it challenges efforts to establish a non-arbitrary boundary between the properties of the central controller and the properties of the rest of the agent's body. In one sense it is obvious that the computational demands facing the control system for a body will vary according to the abilities of and the constraints upon that body. Furthermore, it is well known that the physical configuration of, for example, a robot body can reduce the computational burden on its central controller. Rodney Brooks (1991) and Valentino Braitenberg (1986) showed how bodies or machines can overcome apparently complicated behavioral challenges by virtue of their mechanical structure alone without recourse to representations. Most famously, for example, Braitenberg's vehicles are simple 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 determine the body's movement.

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. The morphology of a bacterium is suited to the set of computational challenges that a bacterium might encounter in its evolutionary niche.

For agents with more complicated sets of behaviors and challenges, it is reasonable to assume the involvement of some kind of central control system. However, drawing the line between the central controller and the body in computational problem solving is difficult. 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. Likewise, the capacities and constraints associated with the body itself can be understood as part of the problem landscape to be navigated by the central controller. In robotics, Rolf Pfeifer drew attention to this set of questions in his discussion of the morphology and control trade-off problem (Pfeifer and Scheier 1999).

The morphology and control problem involves determining an agent's bodily capacities and constraints, but it can also be understood as the problem of determining the boundaries between the computations that are performed by the central controller and those that are performed by the rest of the agent's body.

This problem raises very traditional philosophical questions concerning embodiment. Some advocates of a view known as morphological computation, most notably Chandana Paul, have argued that bodies should be understood as part of the agent's cognitive process

(Paul 2004). Paul defined morphological computations as computations where the mechanical structure of the body itself carries some of the burden of solving the agent's computational problem. The goal of her research is to provide "a common currency between the realm of the physical body and the controller" in robotics (Paul 2006), and she writes:

Controllers lie in the realm of abstract computation, and as such are usually implemented in computational hardware. However, if physical interactions can also perform computation, it becomes possible for the dynamics of the morphology to play a computational role in the system, and in effect subsume part of the role of the controller. (ibid., 620)

So, for example, Rolf Pfeifer describes a robot hand that is designed with flexible and soft gripping surfaces, artificial tendons such that a single instruction from the hand's control system can initiate a range of kinds of gripping actions with a range of different kinds of objects. The key here is that the physical structure of the robot hand carries the burden of coping with a wide range of behavioral challenges. The manner by which the robot hand grips a wide variety of objects could be understood as a highly complex information theoretic challenge. Much of the complexity of that challenge is overcome by means of the morphological structure of the hand.

5.3 Perceptrons, bodies, and computation

Chananda Paul's morphology-based treatment of the Boolean XOR function provides a clear way of understanding the central conceptual problems in morphological computation. Paul (2004) describes a robot which, by virtue of its morphology, can exhibit XOR-constrained behavior. The robot has a central control system composed of perceptrons, but that system, for reasons to be explained below, cannot implement the XOR function. Instead, the rest of the agent's body implements a function that looks a great deal like XOR. As we shall see, it is not quite correct to straightforwardly identify the morphologically computed solution with XOR.

The significance of the XOR function is connected to the history of cognitive science, and specifically to Rosenblatt's development of the perceptron and the subsequent criticisms by Minsky and Papert (1969). A perceptron is a 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 learn to solve functions which are not linearly separable was one of the reasons that popular scientific opinion turned against neural networks for a time.

A linearly separable problem is one in which for any output neuron there must be some hyperplane (of dimension n-1) which divides the set of n inputs to that neuron between 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 outputs 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

Linear separability of OR

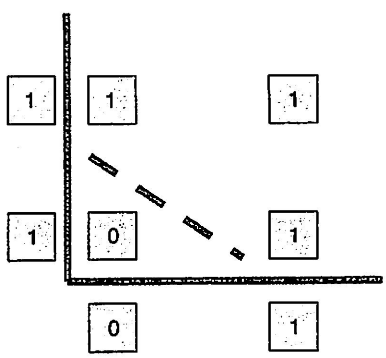

By contrast, functions like XOR or XNOR are not linearly separable as can be seen in the following table and figure:

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

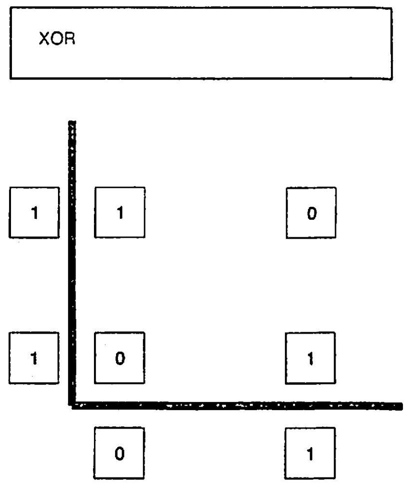

XOR is not linearly separable as can be seen from the figure here. No straight line will partition the space of values into homogenous sets of 0s and 1s. Given that functions like XOR and XNOR are not linearly separable, they fall beyond the computational capacity of a single perceptron.[6]

The robot which Paul describes is controlled by two perceptrons, one for a motor M1 that computes OR, the other for a second motor M2 that computes AND. Whereas M1 turns a 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.

A	B	**Behavior**
T	T	Stationary
T	F	Moving
F	T	Moving
F	F	Stationary

The XOR 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. (See Paul 2004, 34.)

Since an AND perceptron and an OR perceptron by themselves cannot compute the XOR function, her robot relies on the structural features of its body to perform the XOR computation. Specifically, we can understand the structure of the robot as providing a way of generating a function which is 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.

How should we understand the role played by the body in the computational equation? Well, we can expand on her description of the truth table for the various parts of the robot as shown in Table 5.1.

However, we might be concerned with providing an account of the functional role of the morphology of the body itself (call it morph). For, example, we might want to specify the function as shown in Table 5.2.

Table 5.1 Truth values for functions and associated robot actions

A or B	M1	A and B	M2	A xor B	Robot
T	GO!	T	Lift	F	Don't move
T	GO!	F	Don't lift	T	Move
T	GO!	F	Don't lift	T	Move
F	STOP	F	Don't lift	F	Don't move

Table 5.2 Truth table for the function "morph"

(a	or	b)	morph	(a	and	b)
T	T	T	F	T	T	T
T	T	F	T	T	F	F
F	T	T	T	F	F	T
F	F	F	F	F	F	F

Notice that though this specification of morph captures the behavior of the robot under all possible configurations of inputs for the combination of the AND and OR perception, it does not exhaustively characterize morph as a logical operator. Specifically, the table above lists only three of the 2^n combinations of truth values for morph. Filling in the final line in the truth table for the operator morph is pretty straightforward, as can be seen in Figure 5.1.

Paul's description of morphological computation provides a precise example of the kind of embodiment that proponents of the embedded-embodied approach in philosophy of mind envision.

So, how should we interpret Paul's characterization of morphological computations? To begin with, we should be cautious about drawing metaphysical implication from examples like this. The mere existence of morphological computation does not, by itself, present a significant challenge to the view that morphological computations are just as multiply realizable as central computations albeit within the context of an "extended functionalist" framework. Functions realized via morphological computation are fully compatible with traditional multiple realizability. In principle, different configurations of the body could provide different extended ways of implementing *the same* function. In the case of morphological computations, one might imagine a variety of different ways in which the structure of the body could supplement the perceptrons, computationally speaking.

So, the point here is not to present a metaphysical challenge to multiple realizability taken as an in principle argument. Rather, examples like this support a proposal to modify classical constraints on computational theory in cognitive science. To begin with, by fixating on multiple realizability, traditional cognitive science has tended to assume the fixity

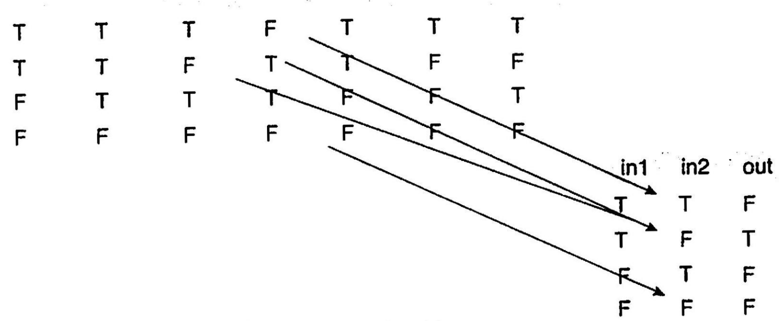

Figure 5.1 Completing "morph's" truth table

of the function at the computational level of description. According to the perspective we'd advocate here, the unique morphological features that the XOR* robot (as we'll call it) exemplifies, are constitutive of the computational challenge facing the agent as well as the behavioral repertoire which it exhibits in response. 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 its environment, the details of its embodiment and the range of feasible physical interactions available to it are constitutive of the manner in which it learns to implement XOR*.

Furthermore, details external to the cognitive system itself can affect the computational problem facing the system; consider how the environment can play a role in the problem, when, for example, we place the XOR* robot on a steep slope. In this case, the function, as described at the computational level, is altered, and as Paul notes such changes happen *as a whole* (Paul 2004). This is why we would call the function XOR* rather than XOR. The XOR function is impervious to hills.

For cognitive science, ignoring bodily and environmental constraints is a methodologically risky strategy insofar as it can miss the problems that serve as the starting point for computational theory. Once these are acknowledged, our computational accounts of the agent's challenges and responses are subject to revision and can be updated in light of increased understanding of the body and the environment.

In this sense, the computational theory of how the agent responds to the environment can be understood as an ongoing research topic, rather than a matter of settled functions at the computational level.

On the view we are advocating here, computational theory should reveal precisely what Marr said it does: It should illuminate *the goal of the computation,* explain *why it's appropriate, and* discover *the logic of the strategy by which it can be carried out.* The computation itself *is* the result of the agent finding or *falling into* a strategy for achieving some goal.

This strategy is constrained by the body, the environment, and the goals of the agent. However, one might object that the morphological computation movement risks confusing the kinds of problems that minds must overcome with the kinds of problems that bodies or organisms as a whole must overcome. The mind (for the most part) solves its own problems in the service of solving the agent's problems.

One problem is the claim that we could think of this as pushing the problem back a bit. If we were, for example to identifying the mind with some aspect of the sub-agential level, we might, for example, push back to identifying the activity of the mind with the functions performed at the level of the perceptrons. If so, then how does the central controller (read the perceptions) tackle the problem of dealing with this bodily morphology? But isn't there something arbitrary about identifying the central controller with the perceptrons rather than any other aspects of the machine? Moreover, there are two distinct perceptrons connected by a body; which of the two is to be identified with the central controller? The AND perception or the OR perceptron? Notice too that they are jointly participating in the effort of producing the XOR operation..

Marr's criticism of the study of feature detecting neurons and his criticism of Gibson certainly encouraged the view that Marr understood the computational level as autonomous with respect to the biology of perceptual systems. However, the *problems and strategies* view of computation which we should take from Marr is more fruitful and, strikingly, not alien from the embodied-embedded. Reflecting on morphological computation encourages the view that computational theory, as Marr understood it, is already the study of embodied computation.

Historically, the bias in favor of the functional or computational level of investigation 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). Clearly, the doctrine of multiple realizability makes plenty of room for the view of the mind as a symbol manipulating system whose relationship to the body is at least contingent and most likely irrelevant. However, our quarrel in this paper is not with multiple realizability per se. Rather, it's with the idea that we have a true and complete story with respect to the way the mind works at the computational level. If we had such a story, then it's quite possible that the body would be irrelevant. However, we will only arrive at a good computational account by understanding the kinds of problems faced by cognitive agents, and we will only understand those problems by understanding the capacities of the body.

The most obvious point we can take from examples like the XOR* robot is that computational theory can generate distinct stories about the inner workings of the agent even when the resulting function is understood. This is not very surprising. Where methodological considerations are more dramatically altered is with respect to the conditions defining a problem for the organism or agent. Understanding what can be morphologically computed by an agent requires more attention to the way the body works than philosophers of mind have been used to paying.

Notes

*We are very grateful to audiences at the University of Pittsburgh and at the New Mexico-Texas Philosophical Society for helpful feedback. Thanks especially to Sarah Robins, whose commentary on an earlier version of this paper helped us to improve it considerably. We gratefully acknowledge the support of Fundación Séneca (Agencia de Ciencia y Tecnología de la Región de Murcia), through project 11944/PHCS/09, and The Ministry of Economic Affairs and Competitiveness of Spain, through Project FFI2009-13416-C02-01.

1. For more on Marr's methodology see Symons 2007.
2. We are grateful to an anonymous reviewer for encouraging us to add this section.
3. They focus on what they see as the mereological fallacy involved in claiming that the brain rather than the whole person "gathers information, anticipates things, interprets the information it receives, arrives at conclusions, etc."
4. The example that Bickle cites is Jerry Fodor's antireductionism. He writes:

   > Citing associative processes and the effects of emotion on perception and belief, even the Jerry Fodor of 1975 insisted that explaining them is left to lower-level (probably biological) investigation (1975, 203). Yet Fodor has long denied that reduction is viable for theories about genuinely cognitive phenomena. (Bickle 1998, 3)
5. Some philosophers insist that we have a finished theory of mind, namely our folk psychological theory of mind. We assume that our readers share our hope for more from the science of mind than the belief-desire-action model of mind has provided so far. If you do, then you do not think that the game is over yet and accept that we do not have a finished theory of mind.
6. 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.

References

Bennett, M. R. & Hacker, P. M. S. (2003). *Philosophical Foundations of Neuroscience*. Oxford: Blackwell.

Bickle, J. (1998). *Psychoneural Reduction: The New Wave*. Cambridge, Mass.: The MIT Press.

Braitenberg, V. (1986). *Vehicles: Experiments in Synthetic Psychology*. Cambridge, Mass.: The MIT press.

Brooks, R. A. (1991). Intelligence without Representation. *Artificial Intelligence*, 47(1), 139–159.

Chalmers, D. J. (1996). *The Conscious Mind: In Search of a Fundamental Theory*. Oxford: Oxford University Press.

Davidson, D. (1970). Mental Events. In *Actions and Events*. Oxford: Clarendon Press.

Davidson, D. (1993). Thinking Causes. In John Heil and Alfred Mele (eds), *Mental Causation*, Oxford: Clarendon Press.

Dummett, M. A. (1993). *The Logical Basis of Metaphysics*. Cambridge, Mass.: Harvard University Press.

Fodor, J. A. (1975). *The Language of Thought*. Cambridge, Mass.: Harvard University Press.
Kripke, S. A. (1980) *Naming and Necessity*. Cambridge, Mass.: Harvard University Press.
Kusch, M. (1995). *Psychologism: The Sociology of Philosophical Knowledge*. London: Routledge.
Marr, D. (1982) *Vision: a Computational Investigation into the Human Representation and Processing of Visual Information*. New York: Henry Holt and Co.
McIntyre, L. (1999). Davidson and Social Scientific Laws. *Synthese*, 120(3), 375–394.
Minsky, M. and Papert, S. (1969). *Perceptrons*. Cambridge, Mass.: The MIT Press.
Neisser, U. (1967). *Cognitive Psychology*. New York: Appleton-Century-Crofts.
Paul, C. (2004). Morphology and Computation. *Proceedings of the International Conference on the Simulation of Adaptive Behavior* (pp. 33–38), Los Angeles.
Paul, C. (2006). Morphological computation: A Basis for the Analysis of Morphology and Control Requirements. *Robotics and Autonomous Systems* 54(8), 619–630.
Pfeifer, R. and Scheier, C. (1999) *Understanding Intelligence*. Cambridge, Mass.: The MIT Press.
Putnam, H. (1975). *Mind, Language and Reality. Philosophical Papers* (Vol. 2). Cambridge: Cambridge University Press.
Symons, J. (2007). The Complexity of Information Processing Tasks in Vision. In Carlos Gershenson, Diederik Aerts, and Bruce Edmonds (eds) *Philosophy and Complexity* (pp. 300–314). Singapore: World Scientific.