Minds & Machines (2015) 25:37–56
DOI 10.1007/s11023-015-9362-1

Behavioural Explanation in the Realm of Non-mental Computing Agents

Bernardo Aguilera

Received: 27 July 2014 / Accepted: 10 February 2015 / Published online: 17 February 2015

Abstract Recently, many philosophers have been inclined to ascribe mentality to animals (including some insects) on the main grounds that they possess certain complex computational abilities. In this paper I contend that this view is misleading, since it wrongly assumes that those computational abilities demand a psychological explanation. On the contrary, they can be just characterised from a computational level of explanation, which picks up a domain of computation and information processing that is common to many computing systems but is autonomous from the domain of psychology. Thus, I propose that it is possible to conceive insects and other animals as mere computing agents, without having any commitment to ascribe mentality to them. I conclude by sketching a proposal about how to draw the line between mere computing and genuine mentality.

Keywords Mentality · Psychological explanation · Computation · Computing agents · Animal cognition

Introduction

If science is supposed to carve nature at its joints, then psychology should be the discipline that carves the domain of minded entities at its joints. Among the entities that have mentality we naturally include human beings and some animals. But which animals? Could, for example, some arthropods, fish or reptiles become the subject of any serious psychological inquiry? If we were to draw a line between minded and non-minded creatures, where should we do so in the evolutionary tree of life?

B. Aguilera (✉)
University of Chile, Santiago, Chile
e-mail: bedobardo@gmail.com

A common way to deal with these questions has been to identify animals whose behaviours do not demand a psychological explanation and then exclude them from the realm of minded creatures. A paradigmatic example is the food-dragging behaviour observed in the digger wasp *Sphex* (Wooldridge 1971). When it comes to lay its eggs, the insect brings food to its burrow nest, leaves the food near its opening, and proceeds to check inside the burrow for the presence of intruders. If nothing is found disturbed, then the wasp emerges from its burrow and drags in the food. Interestingly, if an experimenter moves the food while the wasp is still inside the burrow, it will invariably repeat the same routine of leaving the food near the burrow and going into it for inspection. This procedure can be repeated again and again, without the wasp altering its behaviour.

The example of the digger wasp is often put forward as a case of a rigid, stereotyped behaviour, that does not deserve to explained in psychological terms (e.g. Dennett 1984; Sterelny 1990). However, recently some philosophers have argued that insects such as bees and ants do have mentality, on the basis that they exhibit behaviours that are much more complex and flexible than the kind observed in the digger wasp (Carruthers 2004, 2006; Fitzpatrick 2008). More precisely, they argue that some behaviours of insects have to be explained by appeal to capacities of computation and information processing that far outstrip fixed innate patterns or associative conditioning. Therefore, the authors conclude, we are justified to shift to psychological explanations and so ascribe mentality to those insects.

In this paper I will argue that this argument is misleading since it rests on a false dilemma. It assumes that apart from fixed action patterns or associative conditioning, the only alternative to explain complex animal behaviour is psychology. But I claim this is wrong insofar as it is possible to distinguish an alternative explanatory level that is capable of doing all the work of explaining complex animal behaviour without any appeal to mentality. This explanatory level describes what I will call the *computational domain*, which can be couched independently from the domain of psychological explanation. My main aim in this paper will then be to argue that the computational level of explanation is autonomous in this sense, and claim that it must be considered as a non-psychological alternative to account for animal behaviour. If my arguments are sound, then it will prove possible to conceive of insects as marvellous biological computers without having any commitment to ascribe mentality to them. Having considered a number of objections and related views, by the end of the paper I shall outline a proposal of what is the difference between mere computing agents and those that should be viewed in psychological terms.

Levels of Explanation and the Scientific Study of the Mind

I will start by setting forth some background I take for granted in this paper. Succinctly, I shall be assuming that the mind is a physical and a computing system, and that—as is customary with scientific explanations of other complex natural phenomena—the mind can be described from the viewpoint of multiple explanatory levels.

According to a standard conception (Glymour 1999), scientific theories consist in a group of predicates, formulated in the vocabulary of the respective science, that describe generalisations or law-like regularities concerning certain natural phenomenon. This is normally framed in terms of the deductive-nomological model of explanation, where explanations consist in subsumption of events under natural laws, and supplemented by a causal account of explanation, which specifies the causal mechanisms that contribute in bringing about the phenomena under study (Salmon 1989). Additionally, it has also become customary in science to study complex systems by appeal to multiple levels of analysis, each specifying a natural domain. As Fodor and Pylyshyn (1988) observe:

> It seems certain that the world has causal structure at very many different levels of analysis, with the individuals recognized at the lowest levels being, in general, very small and the individuals recognized at the highest levels being, in general, very large. Thus there is a scientific story to be told about quarks; and a scientific story to be told about atoms; and a scientific story to be told about molecules… ditto rocks and stones and rivers… ditto galaxies. And the story that scientists tell about the causal structure that the world has at any one of these levels may be quite different from the story that they tell about its causal structure at the next level up or down. (p. 5)

Each "story" corresponds to a scientific level of analysis, with its own explanatory vocabulary and laws. In principle, explanations at each level are supposed to be autonomous because they capture a genuine level of organisation in nature which would be missed if described from lower levels and thus cannot be reduced to them. So according to the standard model of multiple levels of analysis, complex systems such as the mind are members of many natural domains, each specified in terms of some particular explanatory level. Furthermore, levels are hierarchically structured, in the sense that the processes of each ascending level are being implemented or realised by the processes of the next level down (McClamrock 1991).

Since the advent of the cognitive revolution the mind has also begun to be analysed in terms of levels of organisation (Marr 1982; Pylyshyn 1984; Sterelny 1990). Though using different terminologies, theorists generally distinguish three levels of explanation for the mind. At the top there is a *psychological level* at which we describe the mental symbols and reasoning processes that cause behaviour.[1] The next level down is the *computational level*, at which we specify the informational structures and computational processes that underlie mental capacities, and at the base we have a *physical level* which describes the mind directly in terms its physical structure. Behind this analysis of levels is the assumption that the mind can be studied by postulating internal functional states, that mediate causally between sensory inputs and behavioural outputs. Both the psychological and the computational levels adopt a functional characterisation of the mind, and thus formulate their explanations at a higher level of abstraction that is independent of the physical mechanisms that implement them.

[1] In addition to mental states that play a causal role in behaviour, the mind is normally understood as involving consciousness. But for the purposes of this paper—and following many philosophers persuaded by the computational theory of mind, I shall assume that consciousness is not essential to psychological explanation, and that important progress can be made on the nature of the mind without addressing the phenomenal character of mental states.

But greater clarification of these levels is in order. The psychological level roughly corresponds to what is known as folk psychology (Von Eckardt 1998). It involves the ascription of beliefs, desires and other symbolic structures that take part in inferences which rationalise behaviour. I shall assume that both the mental and the computational levels pose symbolic structures that carry information from the environment. I call those structures *mental symbols* and *computational symbols*, respectively.[2] In the case of mental creatures, their *mental symbols* are realised by computational symbols (from the computational level). However, I prefer to treat them separately because, as I shall explain later, *computational symbols* can exist without being the implementational base of mental symbols.

A note of caution regarding the extent to which these explanatory levels are independent. Each level is relatively autonomous in terms of their taxonomies and generalisations, but as previously noted each describes a natural domain that depends for its existence on more fundamental domains. In the case of the mind, the domain described by the psychological level is realised by the domain of the computational level, and the computational domain is then realised by the physical domain. When studying mental creatures, it might then be more accurate to speak of a "computational *implementation*-level" and a "physical *implementation*-level". But I prefer to omit "implementation" because I will argue that computational- and physical-level explanations can constitute behavioural explanations in their own right, without being descriptions of the implementational level of a higher natural domain.

It is also important to note that this analysis of levels is compatible with a monist metaphysics as well as with the generality of physics as an account of the natural world. The fact that computational and mentalistic explanations treat their own explanatory domains by using concepts and generalisations that cannot be, or do not need to be, expressed by the vocabulary of physics is compatible with the generality of physics insofar as they denote physically constituted entities whose processes are not in conflict with the principles of physics (Crane 2001).

Whether the computational level is an independent level of explanation is a controversial issue, however. It could be argued, for instance, that computational explanations are always part of psychology, and that it therefore makes no sense to deploy a computational level of analysis to describe the behaviour of agents without a mind. In the following sections I will argue that this is not the case, however, and that the computational level represents a natural domain that is not necessarily restricted to mental creatures.

The False Dilemma

Before the emergence of cognitive science, animal behaviour was normally explained without appeal to internal functional states. Standard forms of explanation

[2] I believe it innocuous to describe symbolic structures at the computational level for two reasons. First, computational theory is inherently symbolic. Second, I adopt an informational approach to symbols according to which computational structures can in some way refer to environmental properties by virtue of carrying and using information about them, without necessarily having fully-fledged mental content. I explain these points in Sects. 4 and 5.1.

were fixed action patterns and associative conditioning, which can roughly be considered types of physical-level explanation given that their neuronal mechanisms are well understood (see e.g. Hawkins and Kandel 1984).[3] But since cognitive science emerged philosophers and cognitive ethologists have become optimistic about the prospects of ascribing internal functional states to animals (e.g. Fodor 1975; Gallistel 1990). This has also been encouraged by ethological studies that show that many animals appear to possess complex computational capacities.

In this respect, insects are a good example. As the aforementioned case of the digger wasp shows, insect behaviour was generally considered to be the result of fixed action patterns or at best basic forms of associative conditioning. But after careful study of certain striking insect behaviours, theorists have become convinced that their explanations should generalise over many possible physical-level descriptions and be couched in a higher-level vocabulary. As a consequence, cognitive ethologists have opted to use the psychological level to account for those complex behaviours, often behind the—sometimes implicit—assumption that without the availability of mental vocabulary they would be left without alternative ways of explaining them (e.g. Jamieson and Beckoff 1996; Carruthers 2004, 2006; Fitzpatrick 2008). Staying with the case of insects, the argument put forward to ascribe mentality to them can be formulated in the following way:

P1 We are justified in placing animals within the psychological domain iff their behaviours cannot be best explained in terms of other explanatory domains
P2 Certain behaviours of insects cannot be best explained in terms of the physical domain
C Insects that exhibit those behaviours can be placed within the psychological domain

Let me briefly comment on the premises. P1 is an assumption about the epistemic grounds that justify the description of animal behaviour in terms of certain explanatory domain. It states the common idea that the same animal behaviour can normally be explained at different levels, but that we should only adopt a positive epistemic attitude towards those explanations that are better than the others. How to precisely spell out what makes an explanation better than another is not an easy task, but most cognitive ethologists agree in that best explanations are those that have more explanatory coverage and predictive power (Allen and Bekoff 1997).

P2 is a consequence of the well settled fact that some insect behaviour is far too complex to be explained at only the physical level. The most common examples are their navigational abilities, which I will explain in some detail in Sect. 6. Others are communication and non-associative forms of learning (Gallistel 1990; Carruthers 2006). But even if we take both P1 and P2 as true, the argument is not sound insofar as it has a hidden premise, which I shall argue is false. It is the following:

[3] It has been proposed that associative conditioning can be accommodated within a computational framework. For example, Gallistel and Gibbon (2001) suggest that associative learning can be better explained in terms of computational operations such as extracting temporal regularities between information-bearing structures. If this is the case, associationist explanations would be part of the computational level. But insofar as they can be formulated without appeal to psychological notions, this approach is compatible with the argument put forward in this paper.

P3 The physical domain is the only alternative to the psychological domain to explain animal behaviour

By omitting this premise theorists have implicitly assumed that the behaviour of animals has to be described by using either of two explanatory levels: the physical or the psychological. But this is a false dilemma, since it ignores one alternative explanatory level, *viz*. the one that describes the computational domain. To support this claim, though, it has to be argued that computational-level explanations are autonomous, which will be the task for the following sections.

The Computational Level Specifies An Autonomous Natural Domain

Now I shall start defending a key claim of this paper: that computational level explanations characterise a natural domain that supervenes on the physical domain, but is not dependent on the psychological domain. The basic idea is that it is at least possible to conceive entities whose behaviour can be satisfactorily explained in terms of computational processing over information-bearing structures, while there is no reason to add psychological notions to explain them. Therefore, the computational domain can be regarded as autonomous from the domain of psychology. I will start by exploring what makes for being a computer before discussing this idea in the context of behavioural explanation. Then, in the following sections I shall address some objections.

On first consideration, a computation is just a functional mapping between two domains (e.g. inputs and outputs). According to this fundamental notion, however, computations are being realised in every entity whose behaviour can be described by a mathematical function. For example, a planet would be computing its orbit, or an enzyme the chemical reaction it is catalysing. It soon becomes clear that in this broad sense a computation is nothing beyond physical-level descriptions that use mathematical models to explain their phenomena. As Crane (1995) points out, in cases like this we could describe the planet and the enzyme as instantiating mathematical functions, but not as computing them. To be a computer is, then, a more demanding notion.

One way to develop the notion of computer is by noting that computational explanations are not just an account of the functional mapping between domains, but also a description of the mechanisms involved in performing that mapping. Computational explanations normally describe those mechanisms in terms of algorithms, that is, specifications of the successive states required for getting from the input to the output of the respective function. Importantly, according to most cognitivist approaches computational explanations are realistically construed, in the sense that the algorithms performed by a computer have to be mirrored at the physical level, by a causal series of state transitions capable of implementing the complexities of the algorithms performed (Haugeland 2003). In this way, to include the specification of algorithms into the definition of a computer constrains the possible physical devices that could possibly realise a computational architecture, insofar as not any entity that dynamically changes over time is capable of mirroring the causal structure of a computer (Chalmers 1996).

But the definition of a computer is even more subtle than this, since not any system capable of mechanising an algorithmic procedure is a computer. A mechanical clock with its cogs and gears can mirror an algorithm for striking the correct number of times for the hour of the day, but it is not a computer. Even simple pocket calculators are not, strictly speaking, computers. This is because they are special-purpose devices which, though able to compute certain algorithms, cannot compute anything else. Computers, on the other hand, are much more powerful devices capable in principle of computing any algorithm. It was Babbage and Turing who discovered that it is possible to build a machine which can perform some fundamental symbol-manipulating operations and in this way be instructed to execute a boundless amount of algorithms (Turing 1937; Haugeland 1987). The conclusion of their research was that there is a certain class of machine—a *universal* or *general-purpose machine*—with the capacity to be programmed to compute any function computable by algorithm. This corresponds to an idealised version of what we normally regard as a computer.

Even though concrete computing machines cannot be universal (given obvious limitations of time and memory) the idea of being programmable remains a fundamental constraint on the definition of a computer (Chalmers 2011). Thus, a computer is a mechanical device that can be programmed to manipulate symbolic structures by means of performing a basic set of fundamental operations (such as encoding, storing, deleting and transforming data-structures) required for implementing any algorithm. It comprises symbols because computations act upon information-bearing structures that can be interpreted as having significant relations with the outside world. It is important to emphasise that this definition does not assume that any actual computer has to be a multi-purpose system. The idea is just that it must have an internal structure that is flexible enough to be in principle programmed to run different algorithmic procedures (Cf. Haugeland 1981; Copeland 1993; Pylyshyn 1984). I will return to this point in the next section.

Being programmable, computing machines can run different algorithms and branch into alternative computational processes depending on current input. This makes them flexible and adaptable to environmental circumstances, features that were identified during the cognitive revolution as mirroring key aspects of human intelligence. Before the development of computers, the only known computer was the brain. But from then on, machines became capable of performing complex tasks that were before only known to be achievable by minded creatures. The idea that computers were the key to understanding human intelligence led proponents of the cognitive revolution to propose that computers qua programmable symbol-manipulators had at least the sufficient means for intelligent action (Newell and Simon 1981). The working hypothesis from this computational approach to the mind, then, is that the mind is a kind of computer, and that it is at least possible to build an artificial computing machine complex enough to be capable of thinking.

Now we are in conditions to make an important distinction for the purposes of this paper: namely, between those entities that are computers and those that are not. Computers correspond to a particular class of machines whose behaviour can be properly described from the computational level, *viz*. a level of explanation that maps onto the natural domain of computation and information processing that is common to computing machines. But can we say that this domain is also independent from the domain of mental entities? I believe the response is yes, and that the most straightforward way of arguing for this is by showing that some entities can be considered genuine computers and that at the same time there is no justification for ascribing minds to them. Let me start elaborating this idea with an example.

As mentioned above, computers can mirror the algorithmic structure of mental processes and generate flexible and adaptive behaviour. Further, computers can be equipped with transducers to pick up information from their environments and the algorithms required for processing that information in a way that leads to intelligent action. In those cases, I contend, it is plausible to ascribe to a computer at least basic forms of symbolic algorithmic processing. A straightforward example of this are robotic computing machines that can behave in natural environments without any help. Take the case of the two robotic rovers that are now exploring the surface of Mars (i.e. Opportunity and Curiosity). Among their multiple capacities are autonomous navigation, thermoregulation, and the capacity to collect detailed geological information from the planet and transmit it to the earth.[4] It appears highly intuitive to say that these vehicles are fairly intelligent computing machines, even though they lack mentality. But in any case, and this is the main point, they exemplify how computational level explanations can be autonomous, since in order to explain how the robot behaves in such a desolated place we have to appeal to the syntactic and symbolic components of its computational architecture, realistically construed. I will return to the case of these robotic rovers in the following sections when addressing objections to the idea that the computational domain is autonomous.

Two Objections to the Autonomy of the Computational Domain

The Computational Level is Just Syntax

This first objection directly challenges the autonomy of the computational domain. It claims that when cognitive scientists describe the mind from the computational level they are not picking out any domain distinct from the psychological domain. Instead, computational explanations are regarded as nothing beyond a formal characterisation, or a syntactic description, of the computational operations carried out by mental creatures (Fodor 1980; Egan 1995). According to this objection the computational level describes the manipulation of symbolic structures, but from a perspective that focuses on their formal or syntactic properties, not their representational or informational properties. Therefore, the objection goes, the computational level does not map onto any particular ontological level of organisation but just “provides a formal, environment-independent, characterisation of a process”. (Egan 1995, p. 199).

[4] Of course, the Mars rovers receive instructions from earth. But since those instructions take some time to reach the rover, they are designed to carry out many tasks in a rather autonomous way, such as self-monitoring, navigating and making some decisions without human intervention.

One motivation[5] behind this position is the *formality condition*, the view that computational explanations of the behaviour of mental agents can only advert to formal (nonrepresentational) properties of symbolic structures (Fodor 1980). Given that psychological explanations do deploy mental symbols with representational properties, a consequence of the formality condition is that the computational domain is an incomplete implementation of the psychological domain, in the sense that "syntactic processes at the computational level implement causal laws [situated] at the intentional [psychological] level" (Fodor 1991, p. 280), whereas symbolic structures at the computational level simply cannot implement the representational properties of the psychological level. This is not to say that computational-level explanations do not involve the manipulation of symbolic structures, but that this manipulation is effected only by means of the syntactic structure of symbols, while their contents are borrowed, so to speak, from the psychological level. If this is the case, then the computational level cannot map on to a separate—autonomous—domain in a hierarchy of supervenient levels of organisation.

I believe the formality condition is too restrictive, however, and that computational-level explanations should not respect it. First, this condition rests on an incomplete characterisation of computing systems, focusing on their syntactic architecture but abstracting from their behaviour as embodied agents in the real world. This point can be made clearer by reflecting on the nature of algorithms. In principle, an algorithm is a sequence of formal operations that could be run by a computer to perform certain tasks (in the same sense as a single algebraic operation can be used for different purposes). However, when an algorithm is implemented in embodied computers behaving in the real world, computational-level explanations do not describe these algorithms as mere formal operations, but as inferential procedures engaged in genuine causal commerce with the environment.[6]

Secondly, it should be noted that computational-level explanations do not lack the resources to account for the representational dimension of computational symbols. Computational theory is often coupled with informational approaches to representation (Dretske 1981; Adams 2003), according to which computational symbols are information-bearing structures, while computational-level explanations normally deal with the coding and transformation of those structures. These computational processes are typically subpersonal; that is, they underlie psychological (belief-desire) explanations. So if information is taken as an objective commodity that can be picked up from the environment, coded and transmitted through subpersonal—computational-level—explanations, then there is nothing mysterious in assuming that computational symbols do possess contents (although

[5] A second motivation can be *internalism*, the (more general) view that representational notions cannot play any genuine role in a scientific psychology (Kim 1982; Stich 1983). But for the purposes of this paper I assume *externalism*, according to which it is plausible to formulate psychological explanations that advert to representational contents.

[6] To satisfy the demands of explaining how concrete computing agents behave in real environments, the computational level might have to be supplemented by computational frameworks distinct from Turing models, such as interactive computation or hypercomputation (see Dodic-Crnkovic 2011 for a review). Nothing in this paper depends crucially on which computational framework we adopt.

distinct from the contents of mental symbols couched at the next level "up", viz. the psychological level).

A last motivation for conceiving the computational level as a distinct, autonomous level of description, is that without it we would be left with no theoretical resources to explain the behaviour of non-mental computing agents. It is impossible to deny the existence of autonomous computing robots capable of engaging in informational transactions with their environment and behaving effectively within it. But if, as the view presented above contends, the computational level is just a formal description of the syntax of psychological processes, then how could we account for these robots' successful negotiation with their environment? The formality condition might be adequate for the purposes of explaining the behaviour of mental creatures, given that the content of computational symbols is provided by the psychological domain. But when it comes to computational explanations beyond the domain of psychological agents, this approach is little able to account for those agents' behaviours.[7]

By way of example, imagine that we want to study one of the robotic Mars rovers mentioned in the previous section, in particular its behaviour related with searching and examining a Martian stone. If we adopt a purely syntactic approach we would be able to describe the algorithms implemented by the robot throughout the process. But of course, this computational account would have to recognise that these algorithms are not just formal abstractions but effective procedures that actually manipulate symbolic structures that carry information from Mars. What could we do to account for these symbolic structures then? One way could be to shift to the psychological level and call those structures mental symbols. However, this would imply ascribing mentality to the robot, something that seems implausible. The only alternative appears to be, then, to simply ascribe the robot with symbolic structures from an autonomous computational domain. This would suffice to explain how information picked up from Mars is relevant in accounting for its behaviour, and also in accounting for counterfactual situations such as implementation of alternative algorithms if the stone examined had a different composition. Therefore there seems to be no reason for treating computational level explanations as purely formal, or as dependent on there being mentality in those agents that implement a computer. They might just be non-mental computing agents in their own right.

Objection 2: The Computational Level of Artefacts has Derived Representations

The second objection to the autonomy of the computational domain can be formulated as a rejoinder to my reply to the first objection. According to this second objection, the case of the Mars rovers does not count as an example of autonomous computation because they are human-designed machines. The general idea behind this argument is clearly stated in the following quote from Haugeland (1981):

[7] When discussing a Fodorian approach (called the "semantic account of computation") Piccinini (2012) arrives at a similar conclusion: the dependency of this account on a psychological notion of content, makes it unfit for providing computational explanations beyond philosophy of mind, e.g. computer science.

> [symbolic structures] only have meaning because we give it to them; their intentionality, like that of smoke signals and writing, is essentially borrowed, hence *derivative*. To put it bluntly: computers themselves don't mean anything by their tokens (any more than books do)—they only mean what we say they do. Genuine understanding, on the other hand, is intentional "in its own right" and not derivatively from something else. (pp. 32–33).

The author uses the term *intentionality* to refer to the semantic or representational properties of symbolic structures. For present purposes the point is that any computing artefact, insofar as it is the product of the purposeful design of a human being, inherits its intelligence from the purposes and intentions of its creator. So, the objection goes, robots such as the Mars rovers have intelligence and other capacities only in a derived, non-original sense. This objection can be linked to theorists who defend teleological approaches to the mind and adopt a historical approach to cognitive functions. They also make the positive claim that the only way an entity could be endowed with intelligence or any kind of function is by have been designed by non-purposeful mechanisms such as natural selection (Millikan 1984; Papineau 1987).

One problem with this teleological objection is that it rests on controversial assumptions about how history determines the nature of functions and cognitive capacities (Crane, 1995; Fodor 2000). For instance, a creature that happens to lack an evolutionary history of selection would be incapable of instantiating any function, regardless of how complex its current behaviour is. But even more problematic is the fact that a teleological approach would have to rule out, as a matter of conceptual analysis, the possibility that an artefact of the right sort could ever be capable of developing intelligent behaviour. On the contrary, it seems likely that even if we have not created genuinely intelligent cognitive machines so far, things might change as computer technology develops (Copeland 1993).

This is not the place to settle this issue, however. The idea that artefacts have intelligence or process information in a way that does not depend on human beings might still be resisted by some authors. In any case, I believe that the objection that artefacts cannot count as instantiations of an autonomous computational domain can be overcome in a different way: by arguing that there are biological systems that instantiate computational capacities yet do not deserve to be ascribed with mentality. Note that since biological systems are the product of evolutionary processes, they are invulnerable to the second objection. I elaborate this point with examples of real biological computers in the next section.

The Case of Biological Computers

In Sect. 4 I sketched the minimal requirements for instantiating a computing system. They basically consist in possessing symbolic structures that engage in informational relations with the environment, and being potentially programmable systems by virtue of performing certain fundamental symbol-manipulating operations. Now I shall present two examples of biological entities that appear to meet these requirements without there being a clear justification for ascribing minds to them. With this, I wish ultimately to make the point that the computational domain is autonomous and independent of the psychological domain, and that theorists who ascribe mentality to animals often make the mistake of ignoring this point as an alternative explanation.

My first example is concerned again with the digger wasp. As mentioned in the introduction, the food-dragging behaviour of this insect has become a commonplace for showing how rigid and stereotypical insect behaviour could be. But interestingly, the wasp also has some more clever facets, in particular, its navigational abilities. Wasps are able to forage over long distances and then find their way home. In order to orientate themselves, wasps can memorise visual properties of landmarks present near their nests and during their displacements, and rely on them to determine the direction of their flight back. In addition, wasps can also navigate by keeping record of the distance and direction traveled, through a process known as *path integration*. They appear to shift to this second mode of navigation when landmarks are not available, as for example when flying in unfamiliar terrain (Healy 1998).

The navigational capacities of the wasp exhibit many of the hallmarks of a computer. They can encode and store information about landmark cues and carry out computational operations on that information. And those operations can be quite complex, notably when doing path integration. In those cases the insect has to monitor its angular and linear displacements and continually update the current distance and direction from their present location to their starting point, and sometimes recalculate their vector flight, for instance when they find themselves lost. That involves basic computational operations such as storing, deleting and transforming data-structures, as well as manipulating them through algorithmic steps that branch into alternative courses of action.

It is important to note is that the navigational capacities of the wasp are probably modular and specific for that task (Carruthers 2006). That explains why the remarkable intelligence exhibited when flying cannot be used for a different task, such as altering their food-dragging behaviour when fooled by an experimenter. So the wasp does exhibit intelligent behaviour, however it is restricted to its navigational capacities. Should we then ascribe psychological states to the wasp? It is at this point that some authors have fallen into the false dilemma of regarding the computational capacities of the wasp as suitable for psychological explanation. On the contrary, I contend, a computational-information processing framework provides us with enough explanatory and predictive power to account for such insect behaviour, without there being any justification, on explanatory grounds, for adding psychological notions. Thus, wasps might just be, if you like, marvellous biological robots that instantiate computation but not mentality.

One could object, however, that my skepticism about the ascription of mentality to the wasp is question begging, since I am just assuming that the computational processes that control its navigational behaviour lack any form of mentality. Someone might claim, for instance, that any creature that instantiates computational processes in its brain deserves to be explained by the psychological level. I believe this objection fails, though, since it draws the line for instantiating mentality too low. I will present a final example to illustrate how implausible it is to ascribe mentality solely on the grounds of biological computation.

My final example relates to the *enteric nervous system*. It consists in a large network of a hundred million neurones (two thousands times more than the wasp's brain!) located in the wall of the human gastrointestinal tract. It accomplishes a variety of functions such as regulating processes of secretion and absorption, blood flow, and controlling of motility of the intestine (Wood 2011). This motility control involves coordinated patterns of contraction and relaxation at different parts of the intestine, related with segmentation, peristalsis and motility cycles. The enteric nervous system has shown to be quite complex and to operate rather autonomously from the brain, to the extent that it has become known as "the gut brain". A remarkable aspect of this system is that it can receive and integrate information coming from different sources; in particular inputs from the brain and information about the mechanical and chemical conditions of the intestine. Then, the enteric nervous system can produce organised motor patterns that are generated, coordinated and modulated by the system itself (Thomas et al. 2004).

Arguably, capacities of the enteric nervous system such as encoding and integrating information, or modulating complex patterns of intestinal motility, make it a good candidate for instantiating computational processes. Indeed, according to a review by Hansen (2003), "The ENS [enteric nervous system] acts as a microcomputer with its own independent software and is organised for program operations independent of the input from the central nervous system (CNS)". This example serves to illustrate the plausibility of biological systems capable of instantiating the computational domain despite their lack of mentality. If there is no impediment for the implementation of computational capacities by non-mental creatures, then the claim that insects have minds cannot simply be grounded in the fact that they are computing systems. Something else must be said about what makes them part of the selected group of computing systems that possess a mind.

Some Related Views: Allen, Dretske and Burge

Various other philosophers have developed the idea that there might be alternative ways of explaining animal behaviour that do not necessarily commit to the taxonomy of folk psychology. Here I critically review proposals of Allen, Dretske and Burge, before pointing out some problems in them.

As I do in this paper, Allen (1997) worries about the tendency of some authors writing on animal cognition to equate complex forms of information processing with mentality. For instance, he draws our attention to the notion of "cognitive map" often employed to describe honeybee navigation (e.g. Gould and Gould 1994), which despite of being used interchangeably with "mental map", does not seem to entail that honeybees possess mentality. Consequently, Allen proposes to distinguish between (mere) cognitive and mentalistic explanations, in a manner analogous to my distinction between computational- and psychological-level explanations. However, contrary to my view, he does not ground this distinction in the use of different explanatory taxonomies (nor in the contrast between subpersonal and personal levels, as I suggest in the next section). Allen's proposal is that what distinguishes mentality is that it "involves consciousness or awareness of the (semantic) *contents* of internal states" (Allen 1997, p. 231).

One problem with this proposal (as noted in footnote 1) is that cognitive scientists often regard consciousness as inessential to psychological explanation, at least when concerned with the capacity of thinking and its role in causing behaviour (see, for example, Cummins 1989, pp. 19–20). Allen is aware of this tradition and thus attempts to operationalise the notion of consciousness in a way compatible with a functionalist (computational) framework. He then goes on to suggest that genuine mentality is conferred by metarepresentational capacities such as having alternative ways of representing the occurrence of the same event, or being able to detect when something has been misrepresented. As Allen puts it, "creatures with minds are sensitive not just to changes in the world, but are sensitive to how well they manage to represent those changes" (Allen 1997, p. 241).

Allen may be on the right track in suggesting that metarepresentational capacities should be among the essential properties of mentality. However, his appeal to consciousness is misleading, since he ends up formulating a functionalist framework where the conscious character of mental states (i.e. "what it is like" to be in them) plays no explanatory role. In any case, as a proposal for setting the bar for mentality, I believe Allen's appeal to metarepresentational capacities is insufficient. We could perfectly conceive of computational architectures capable of sensing and monitoring their own informational states without mentality. Thus a meta-symbolic architecture might be necessary for having a mind, but is not obviously sufficient. In any case, I believe the personal-level framework I put forward in the next section does a better job in providing an intuitive view of the architecture of the psychological domain.

Dretske (1981, 1999) is another philosopher who has noted that not all behaviour caused by complex information processing demands a psychological explanation. In his view, mental symbols originate from information-bearing states that acquire a functional role in behaviour, thanks to a process of codification and integration with other informational structures. But not all systems possessing information-bearing states that play an explanatory role in behaviour count as having a mind. Dretske (1999) gives the example of machines as a paradigmatic case of mindless information-processing systems. This would be so because their internal functional states were designed, by their creators, to perform the way they do. They would then lack thought and purposeful behaviour since to explain their capacities, we have to appeal to the minds and purposes of their creators. Dretske argues that genuine mental symbols, in contrast, must have to be acquired by the agent's own means through processes of learning such as operant conditioning.

I believe Dretske's proposal is problematic since it ends up setting the bar for mindedness too low. If the threshold for purposeful action and thus for mental symbols is situated at the capacity to learn by operant conditioning, then we would have to ascribe mental symbols to some artefacts and animals to which the attribution of mentality is implausible. Starting with artefacts, it is certainly possible to program some robots with learning algorithms that make them able to develop behavioural effects similar to classical and operant conditioning (e.g. Touretzky and Saksida 1997; Weng 2004). With respect to animals, Dretske's proposal would lead us to ascribe mentality to sea slugs such as *Aplysia*, which is capable of associative learning via classical conditioning (Hawkins and Kandel 1984). More generally—as noted in Sect. 3—the problem with this view is that learning by associative conditioning can be adequately explained within a physical- or a computational-level framework, without the use of psychological notions.

Finally, I shall discuss a proposal recently put forward by Burge (2010). As here, he adopts a form of scientific realism and suggests that we should restrict mentalistic taxonomy to account for phenomena that are better explained by the psychological level. Burge argues that the most basic forms of representation[8] appear in perceptual explanations, to the extent that "in studying perception, representational psychology begins. With perception, one might even say, mind begins" (p. 367). Burge goes on to criticise some philosophers (e.g. Dretske) on the grounds that they have been too liberal in using the term representation to describe the information-bearing states of certain creatures and artefacts, since those philosophers have assimilated representation "to notions that have no *distinctive* theoretical relation to psychology as it is ordinarily understood" (p. 293). Accordingly, Burge proposes to draw the line for mentality at creatures where perceptual capacities can be properly accounted for in psychological terms.

Let us briefly examine Burge's view of perception. Following the cognitive tradition he claims that the primary form of perceptual representation originates in the detection of perceptual invariants; that is, through the computational capacity to yield a constant perception of distal environmental properties. According to Burge, perceptual capacities are also teleological, and so have the function of generating perceptual representations. This gives rise to a normative dimension in the sense that when perception functions properly, perceptual representations are accurate or veridical. A further aspect of Burge's account is that perceptual functions are part of a larger system of distinct functions that work towards the fulfilment of the functions of an organism taken as a whole. Briefly, what integrates the different functions of an organism is what Burge calls "agency", viz. the capacity to generate "functioning, coordinated behaviour by the whole organism, issuing from the individual's central behavioural capacities, not purely from subsystems" (p. 331). The notion of agency is important because it is a precondition for the emergence of perception and representation. More precisely, thanks to the possession of agency perceptual representations can play a genuine role in guiding the actions of the whole organism, and thus be the object of psychological explanations.

Contrary to my view, Burge does not attempt to formulate a level of behavioural explanation between (what I call) the physical and the psychological levels. When perceptual explanations do not apply, Burge tends to shift towards a biological (i.e. physical-level) model of explanation. For example, he suggests that when the information-gathering mechanisms of certain creatures do not fit within a psychological explanatory framework, the notion of perceptual representation "should be dropped in favor of other notions, notions of sensitivity or discrimination, or co-variation, or causal co-variation, or structurally isomorphic causal co-

[8] Burge understands the term "representation" as strictly psychological, meaning that it can be equated with what I call "mental symbols" in this paper.

variation, or information-carrying—together with the notion of biological function" (p. 294).

Burge presents empirical evidence that suggests that information-gathering mechanisms which satisfy his criteria for having perceptual functions appear to be widespread in the animal kingdom, even in phylogenetically primitive animals such as arthropods. I believe, however, that his view ends up being too liberal, and that his reasons for using psychology to describe the behaviour of those animals is not compelling. First, it seems plausible to ascribe the capacity of detecting environmental invariants and processing that information through computational operations to mindless robots, such as the Mars rovers. And as far as their information-gathering mechanisms are fallible and thus allow for malfunctioning, they might also be characterised in teleological terms. Second, the notion of agency does not help much to set the bar for mentality since, according to Burge, agency can be found in very primitive organisms, even some which lack a central nervous system. For example, he claims the paramecia's eating and swimming behaviour counts as agency.

In sum, I concur with Burge on his claim that psychological explanations should set the criteria for mapping out the domain of mental agents. However, I believe his proposal regarding minimal forms of psychological explanation (i.e. perceptual psychology) is unconvincing and does not appear to rule out an alternative explanatory framework, viz. the computational-level. Another point that deserves consideration is Burge's notion of agency, in particular the idea that typical psychological explanations presume that representational states are guiding the behaviour of whole-agents, instead of their (subpersonal or computational) parts. I find the general idea appealing, yet I believe it is better accommodated within the personal-level framework I put forward in the next section.

Final Remarks

In this paper I have argued that there is some conceptual space to be mapped out between the kind of complex information processing required to explain the behaviour of insects (and other computing agents) and genuine mentality. This conceptual space is specified by what I have called computational level, which describes a natural domain distinct from—yet overlapping with—the domain of mental agents. But even if the reader grants all of this, she might still be wondering what makes the difference between mere computing agents and those that should be viewed in psychological terms. This concern takes us back to our initial question: where should we draw the line between minded and non-minded creatures?

One common way to proceed is by following what Lurz (2009) calls a *bottom-up approach*, which begins by taking an intuitively plausible ascription of mentality at face value and then proceeds towards a theory of behavioural explanation for non-human agents that includes psychological terms. This applies to philosophers who have considered it plausible to ascribe "simple minds" to animals on the basis of their possession of rather complex computational mechanisms linking their information-gathering and action systems together. But as I have contended here, this approach is problematic since it ignores computational-level explanations as an alternative to psychology. Unless they offer some principled way of distinguishing mere computational explanations from genuine psychological explanations, their grounds for ascribing mentality remain unjustified.

In concluding I shall outline how we might better draw the line between minded and non-minded computing agents based on a *top-down approach*. Instead of furnishing computational explanations of animal behaviour with psychological notions, I propose to take psychological explanations of human behaviour as the paradigm for judging whether other computational agents have minds. I believe the most straightforward way to proceed in this respect is by identifying the contrast between psychological- and computational-level explanations with the distinction—well entrenched in the philosophical literature—between *personal* and *subpersonal* levels of explanation of the mind.[9]

What characterises personal-level explanations is that they describe behaviour in terms of activities and states that belong to the person as a whole. In contrast, subpersonal explanations target a distinct subject matter: they go deep into the underlying cognitive or neural mechanisms that enable a person to have the mental or behavioural properties under scrutiny. In Dennett's words, "subpersonal theories proceed by analysing a person into an organization of subsystems… and attempting to explain the behavior of the whole person as the outcome of the interaction of these subsystems" (1979, p. 153). Personal- and subpersonal-level explanations also differ in their theoretical vocabulary. While personal-level phenomena are couched in terms of psychological (belief-desire) explanation, subpersonal processes are expressed in terms of computation and informational theories.

A final contrast between both explanatory approaches is that personal-level explanations are governed by normative principles of (instrumental) rationality which constrain how the person's beliefs and desires come together to bring about actions and thereby satisfy desires and goals. This sort of normativity is supposed to constrain the behaviour of the whole rational agent and thus not to be operative at the subpersonal level. Even though subpersonal-level explanations can have a normative dimension—in particular when computational capacities are understood teleologically and therefore aimed at certain ends or goals—this normativity is concerned with the functioning of particular subsystems and not with explanations that rationalise how a (whole) person attains desires and goals.

Subpersonal explanations can be regarded as a subclass of computational-level explanations; while the former maps onto mental computing agents, the latter maps onto non-mental computing agents. But even though subpersonal explanations always apply to psychological creatures, it is important to keep in mind that the subpersonal level is distinct from the personal (i.e. psychological) level of description. As a kind of computational-level explanation, the subpersonal level has its own taxonomy and captures principles of operation that differ from those of the personal level. Thus, while subpersonal processes are always part of mental agents,

[9] The personal-subpersonal distinction was first formulated by Dennett (1969) and since then it has had considerable influence within the philosophy of pscyhology literature (see e.g. McDowell 1994; Hornsby 2000; Bermúdez 2005; Frankish and Evans 2009).

this does not undermine my claim that computational-level explanations (in general) are autonomous from psychology, given that the computational domain can sometimes be instantiated in computing agents where psychological explanations do not apply.

It is not the purpose of this section to offer a fully-fledged set of criteria for mindedness, but just to show how we might advance in drawing distinctions between computing agents with and without mentality. Of course, the personal-subpersonal distinction has to be worked out in order to be suitable for present purposes. First, to deal with anthropomorphic concerns related to defining the mind in terms of persons, the personal-level approach should abstract from human-specific features by focusing just on paradigmatic aspects of psychological explanation so as to adapt it to explaining the behaviour of animals and even machines. This might leave out, for example, language and self-consciousness, but should maintain fundamental aspects of the personal level such as taking whole agents as its subject matter, using a distinctive theoretical vocabulary, and being constrained by norms of rationality.

Second, in order to map the personal-subpersonal distinction onto the contrast between psychological and computational levels of explanation, the former distinction should be reflected in some constraints on cognitive architecture. For example, computing agents endowed with mentality should possess symbolic structures that function in a coordinated way towards achieving goals that concern the agent as a whole. Regarding the theoretical vocabulary of personal-level descriptions, this involves the attribution of mental states with conceptual structure, which can be expected to satisfy requirements for concept possession such as the generality constraint. And finally, the application of instrumental rationality to the whole agent typically demands some degree of consistency and coherence between its mental states. This, in turn, might constrain the way the computational-inferential structure of the system is organised, for example, with respect to the integration between different domain-specific processing units.[10]

As with any account committed to scientific realism, my proposal is a blend of epistemology and metaphysics. It assumes that it is through our best psychological theories that we can gain (observer-independent) knowledge about the nature of the mind. I have put forward the personal-level approach as a convenient way to formulate psychological explanations and thus reveal the main features and constraints of any computing agent to which we ascribe mentality. Admittedly, this section provides a preliminary pass through those constraints. My purpose here has been to show how we can make headway towards setting the bar for mentality, provided that there is a real distinction to be made between computing agents with and without a mind.

[10] To determine which animals might actually satisfy these computational constraints exceeds the scope of this paper. However, I believe that these conditions are rather demanding and cast doubt about whether computing agents such as insects could satisfy them. For example, honeybees might fail to satisfy the generality constraint due to their massively modular computational architecture (Aguilera 2011; for a dissenting view, see Carruthers 2009).

Acknowledgments The author would like to thank Stephen Laurence, Dominic Gregory, Asa Cusack and anonymous referees for helpful comments on earlier drafts of this paper.

References

Adams, F. (2003). The informational turn in philosophy. *Minds and Machines, 13*, 471–501.
Allen, C. (1997). Animal cognition and animal minds. *Mindscapes: Philosophy, science and the mind* (pp. 227–244). Pittsburg: University of Pttsburg press.
Allen, C., & Bekoff, M. (1997). *Species of mind: The philosophy and biology of cognitive ethology*. Cambridge, MA: MIT Press.
Aguilera, B. (2011). Do honeybees have concepts? *Disputatio*, *4*(30), 1–19.
Bermúdez, J. L. (2005). *Philosophy of psychology: A contemporary introduction*. New York: Routledge.
Burge, T. (2010). *Origins of objectivity*. New York: Oxford University Press.
Carruthers, P. (2004). On being simple minded. *American Philosophical Quaterly, 41*(3), 205–220.
Carruthers, P. (2006). *The architecture of the mind*. New York: Oxford University Press.
Carruthers, P. (2009). Invertebrate concepts confront the generality constraint (and win). In R. Lurz (Ed.), *The philosophy of animal minds* (pp. 89–107). Cambridge: Cambridge University Press.
Chalmers, D. J. (1996). Does a rock implement every finite-state automaton? *Synthese, 108*(3), 309–333.
Chalmers, D. (2011). A computational foundation for the study of cognition. *Journal of Cognitive Science, 12*, 323–357.
Copeland, J. (1993). *Artificial intelligence: A philosophical introduction*. Oxford: Blackwell.
Crane, T. (1995). *The mechanical mind: A philosophical introduction to minds, machines and mental representation*. Harmondsworth: Penguin Books.
Crane, T. (2001). *Elements of mind*. New York: Oxford University Press.
Cummins, R. (1989). *Meaning and mental representation*. London: MIT Press.
Dennett, D. (1969). *Content and consciousness*. London: Routledge.
Dennett, D. (1979). *Brainstorms*. London: Penguin Books.
Dennett, D. (1984). *Elbow room: The varieties of free will worth wanting*. Cambridge, MA: MIT Press.
Dodig-Crnkovic, G. (2011). Significance of models of computation, from Turing model to natural computation. *Minds and Machines, 21*(2), 301–322.
Dretske, F. (1981). *Knowledge and the flow of information*. Cambridge, MA: MIT Press.
Dretske, F. (1999). Machines, plants and animals: The origins of agency. *Erkenntnis, 51*, 19–31.
Egan, F. (1995). Computation and content. *The Philosophical Review, 104*(2), 181–203.
Fitzpatrick, S. (2008). Doing away with Morgan's canon. *Mind and Language, 23*(2), 224–246.
Fodor, J. (1975). *The language of thought*. Cambridge, MA: Harvard University Press.
Fodor, J. (1980). Methodological solipsism considered as a research strategy in cognitive psychology. *Behavioral and Brain Sciences, 3*, 63–109.
Fodor, J. (1991). Replies. In B. Loewer & G. Rey (Eds.), *Meaning in mind: Fodor and its critics* (pp. 255–319). Cambridge: Blackwell.
Fodor, J. (2000). *The mind doesn't work that way*. Cambridge, MA: MIT Press.
Fodor, J. A., & Pylyshyn, Z. W. (1988). Connectionism and cognitive architecture: A critical analysis. *Cognition, 28*(1–2), 3–71.
Frankish, K., & Evans, J. S. B. T. (2009). The duality of mind: An historical perspective. In J. S. B. T. Evans & K. Frankish (Eds.), *In two minds: Dual processes and beyond* (pp. 1–32). Oxford: Oxford University Press.
Gallistel, C. R. (1990). *The organization of learning*. Cambridge, MA: Bradford Books/MIT Press.
Gallistel, C. R., & Gibbon, J. (2001). Computational versus associative models of simple conditioning. *Current Directions in Psychological Science, 10*(4), 146–150.
Glymour, C. (1999). Realism and the nature of theories. In M. H. Salmon, C. Earlman, C. Glymour, J. G. Lennox, J. Machamer, J. McGuire, J. Norton, et al. (Eds.), *Philosophy of science* (pp. 104–131). Indianapolis: Hackett Pub Co Inc.
Gould, J. L., & Gould, C. G. (1994). *Animal mind*. New York: NY: Scientific American Library.
Hansen, M. B. (2003). The Enteric nervous system I: Organisation and classification. *Pharmacology and Toxicology, 92*, 105–113.
Haugeland, J. (1981). Semantic engines: An introduction to mind design. In J. Haugeland (Ed.), *Mind design: Philosophy, psychology, and artificial intelligence* (pp. 1–34). Cambridge, MA: MIT Press.

Haugeland, J. (1987). *Artificial intelligence: The very idea*. Cambridge, MA: MIT Press.
Haugeland, J. (2003). Syntax, semantics, physics. In J. M. Preston & M. A. Bishop (Eds.), *Views into the Chinese room: New essays on Searle and artificial intelligence* (pp. 379–392). New York: Oxford University Press.
Hawkins, R. D., & Kandel, E. R. (1984). Is there a cell-biological alphabet for simple forms of learning? *Psychological Review, 91*(3), 375–391.
Healy, S. D. (1998). *Spatial representation in animals*. Oxford: Oxford University Press.
Hornsby, J. (2000). Personal and sub-personal: A defence of Dennett's early distinction. *Philosophical Explorations, 1*, 6–24.
Jamieson, D., & Beckoff, M. (1996). On aims and methods of cognitive ethology. In *Readings in animal cognition* (pp. 65–78). Cambridge, MA: MIT Press.
Kim, J. (1982). Psychophysical supervenience. *Philosophical Studies, 41*, 51–70.
Lurz, R. (2009). *The philosophy of animal minds*. Cambridge: Cambridge University Press.
Marr, D. (1982). *Vision*. New York: W. H. Freeman.
McClamrock, R. (1991). Marr's three levels: A re-evaluation. *Minds and Machines, 1*, 185–196.
McDowell, J. (1994). The content of perceptual experience. *The Philosophical Quaterly, 44*(175), 190–205.
Millikan, R. (1984). *Language, thought, and other biological categories*. Cambridge: MIT Press.
Newell, A., & Simon, H. A. (1981). Computer science as empirical enquiry: Symbols and search. In *Mind design: Philosophy, psychology, and artificial intelligence* (pp. 35–66). Cambridge, MA: MIT Press.
Papineau, D. (1987). *Reality and representation*. Oxford: Blackwell.
Piccinini, G. (2012). Computation in physical systems. In E. N. Zalta (Ed.), *The Stanford encyclopedia of philosophy*. Retrieved from http://plato.stanford.edu/archives/fall2012/entries/computation-physicalsystems/.
Pylyshyn, Z. W. (1984). *Computation and cognition*. Cambridge MA: MIT Press.
Salmon, W. C. (1989). Four decades of scientific explanation. In P. Kitcher & C. Salmon, Wesley (Eds.), *Scientific explanation* (pp. 3–219). Minneapolis: University of Minnesota Press.
Sterelny, K. (1990). *The representational theory of mind: An introduction*. Oxford: Basil Blackwell.
Stich, S. (1983). *From folk psychology to cognitive science*. Cambridge MA: MIT Press.
Thomas, E. A., Sjövall, H., & Bornstein, J. C. (2004). Computational model of the migrating motor complex of the small intestine. *American Journal of Physiology. Gastrointestinal and Liver Physiology*, *286*(4), G564–G572.
Touretzky, D. S., & Saksida, L. M. (1997). Operant conditioning in Skinnerbots. *Adaptive Behavior, 5*(3), 219–247.
Turing, A. (1937). On computable numbers, with an application to the Entscheidungsproblem. *Proceedings of the London Mathematical Society*, *s2-43*, 544–546.
Von Eckardt, B. (1998). Folk psychology. In S. Guttenplan (Ed.), *A companion to the philosophy of mind* (pp. 300–307). Oxford: Blackwell.
Weng, J. (2004). Developmental robotics: Theory and experiments. *International Journal of Humanoid Robotics, 1*(2), 199–236.
Wood, J. D. (2011). *Enteric nervous system (the brain-in-the-gut)*. Morgan & Claypool Life Sciences Series: Princeton.
Wooldridge, D. (1971). *The machinery of the brain*. New York: McGraw-Hill.