

Kent Academic Repository

Full text document (pdf)

Citation for published version

Johnson, Colin G. (2005) Does a functioning mind need a functioning body? In: Davis, Darryl N., ed. Visions of Mind. Idea Group Publishing, pp. 307-321. ISBN 1-59140-483-5.

DOI

Link to record in KAR

<https://kar.kent.ac.uk/14379/>

Document Version

UNSPECIFIED

Copyright & reuse

Content in the Kent Academic Repository is made available for research purposes. Unless otherwise stated all content is protected by copyright and in the absence of an open licence (eg Creative Commons), permissions for further reuse of content should be sought from the publisher, author or other copyright holder.

Versions of research

The version in the Kent Academic Repository may differ from the final published version.

Users are advised to check <http://kar.kent.ac.uk> for the status of the paper. **Users should always cite the published version of record.**

Enquiries

For any further enquiries regarding the licence status of this document, please contact:

researchsupport@kent.ac.uk

If you believe this document infringes copyright then please contact the KAR admin team with the take-down information provided at <http://kar.kent.ac.uk/contact.html>

Chapter 14.

Does a Functioning Mind Need a Functioning Body? Some Perspectives from Postclassical Computation

Colin G. Johnson
University of Kent, UK

Abstract

In recent years, the idea that somatic processes are intimately involved in actions traditionally considered to be purely mental has come to the fore. In particular, these arguments have revolved around the concept of somatic markers, i.e., bodily states that are generated by mind and then re-perceived and acted upon. This chapter considers the somatic marker hypothesis and related ideas from the point of view of postclassical computation, i.e., the view that computing can be seen as a property of things-in-the-world rather than of an abstract class of mathematical machines. From this perspective, a number of ideas are discussed: the idea of somatic markers extending into the environment, an analogy with hardware interlocks in complex computer-driven systems, and connections with the idea of “just-do-it” computation.

Introduction

One of the main topics of this book is the computational requirements for the existence of a functioning mind. This could either be a “purely rational” mind, or it could be a mind with affective capacity. In this chapter, I would like to consider to what extent it is possible for such a mind to exist in isolation from some form of “body.” In particular, this question will be considered from the point of view of postclassical computation, which attempts to ground computation not in mathematical theories of abstract machines but by an analysis of the computational capabilities of the real world.

We will take a nondualist perspective as an axiom. Therefore, there is a requirement for the mind to be realized in some fashion in the physical world. The aim here is to consider the relationship between those parts of the body that act as a substrate for mind (in the sense that they could be replaced in a functionalist fashion by another substrate with no difference) and those parts of the body that influence mind yet that cannot/are not part of a substitutable substrate. The “cannot/are not” in the previous sentence can be interpreted usefully at a number of levels. A strong notion may be that there are no physically possible ways of realizing the same phenomenon. Some phenomena may admit a weaker notion in that it is “easier” in some sense (for example, faster, more energy efficient) for the mind to process this phenomenon using an alternative process rather than processing it on the neural substrate.

This chapter is structured as follows. The first section consists of a review that gives a context for the current work. This consists of an outline of notions of embodiment from the cognitive science and robotics literature and the core ideas of postclassical computation. The main sections of the chapter are concerned with in-the-world extensions to Damasio’s theory of somatic markers (1994), connections between notions of the embodied mind and ideas of hardware interlocks in computing systems, and connections between the mind–body relationship and “just do it” computation. A final conclusion summarizes the arguments of the chapter.

Background

This section reviews the two main ideas on which this chapter is based. In the first part, notions of embodiment from the robotics literature are discussed. The second part discusses the idea of postclassical computation, i.e., computation that is based on the properties of physical objects in the world rather than on

specific abstract models of computing machines. The final part looks at ways in which these two areas can be linked.

Notions of Embodiment

In recent years, much has been written about the importance of embodiment in the study of robots and autonomous machines. A definition of embodiment is given by Quick et al. (1999): basically, the embodied object is able to perturb some states of the environment and vice versa. There are a number of ways in which such notions are important for cognitive science research (Quick et al., 1999; Wilson, 2002; Ziemke, 2001, 2003).

In order for cognition to occur in the world, it has to be realized in some worldly “stuff.” This alone is sufficient reason for considering notions of embodiment to be important to the arguments in this chapter. However, there is a second reason that is yet more relevant: the stuff in which cognition is realized will influence the cognitive capacity of the system. The brain-substrate, on which neural-network models of mind reside, opens certain pathways for, and places certain constraints on, the kinds of computations that can be carried out. By delegating parts of cognition to the body, a different set of computational affordances (Gibson, 1977; Norman, 1988) is provided.

A third perspective is that some procedures make use of their being-in-the-world as part of their functioning. Brooks (1991) said that “the world is its own best model,” and many artificial cognitive systems consist of a reactive model, taking information from the world in an implicit fashion and reacting to that information, rather than building up an internal model of the world.

Another aspect of embodiment is that the actions of an embodied cognitive system are potentially unlimited. An action carried out on the mental substrate is limited by the computational capacity of that substrate. However, once a computation is sent-off-substrate into the world, this restriction is removed.

Postclassical Computation

The aim of this chapter is largely to explore how the perspective offered by postclassical computation (a.k.a., nonclassical, nonstandard, physically grounded computation) can inform issues concerned with the embodiment of mind. Broadly speaking, postclassical computation is concerned with ideas of computation that come from the world, for example, studying the information-processing capabilities or the information storage capacity of some physical, chemical, or biological system. This contrasts with the (equally valid) “classical” view of

computing as a property of certain classes of abstract machines. Readers should see Stepney et al. (2003) and Stonier (1990) for overviews of some of the ideas in this area.

The standard view of computing is that it is not concerned with things in the natural world. Grundy (1998) compared the existence of a science of computers to “a science called motorcarology.” The gist of this argument is that computing is a synthetic subject. It is concerned with discovering (empirically grounded, theoretically supported) guidelines for the creation of computing machines. This contrasts with the analytic perspective on the world taken by a subject such as physics or psychology. In such a subject, the aim is to analyze things in the world using the tools and perspectives offered by the subject in question.

Most of the mature sciences have both an analytic side and a synthetic side to them. For example, the subjects of chemistry (analytic) and chemical engineering (synthetic) support each other in understanding and enabling the construction of objects at the molecular level. However, within computer science, there is traditionally no analytic perspective. One of the key approaches in postclassical computation is to consider a computational stance toward objects in the world. We can ask computational questions, such as the following: What is the information-processing capacity of that object? How much information can an object store? How quickly can a particular object transform information, and how does that contrast with how quickly a traditional computing device might do it? How is information compressed within that object? How does it achieve a certain kind of parallel processing of information? What constraints are placed on the object’s function by the requirement to calculate certain characteristics, retrieve certain pieces of information, etc.?

It might be the case that computation has only accidentally been discovered through the creation of computing machines. Feasibly, computational notions could have been discovered by the analysis of certain objects in the world, before being applied in a synthetic fashion.

There are a number of areas in which we might see this kind of thinking being applied in the near future, or where we are beginning to see evidence of this approach already. One area is in immunology—the immune system can be viewed as a learning system that takes information from various invasions into the body and stores information about these in a network that helps the body distinguish self from nonself (deBoer et al., 1993).

This perspective on immunology allows us to ask computational questions about the immune system. What is the memory capacity of the system? How does the system retrieve information about previous infections quick enough to respond to reinfections by the same or similar antigen? How can the system recognize a wide variety of infections, and yet still have the capacity to swamp a particular infection with large amounts of specific lymphocytes when needed? Is informa-

tion passed between lymphocytes to enable them to deal with infection, or is the apparent change in specificity a property of the population changing in composition? How quickly does a rapidly mutating antigen such as the HIV virus have to change to be able to change quicker than the rate of change in the immunological memory of the system?

Other areas of physiology can also benefit from being viewed in a computational fashion. Paton (1996) discussed the benefits to be gained from viewing certain processes in tissues as being a parallel and distributed system that processes molecular information. As an example, he discussed the processes found in the lining tissues of glands in terms of information processing. Using this, he was able to formulate testable hypotheses about why certain structures in glands have the forms that they do.

An interesting question is the extent to which computational capabilities of systems are able to influence the development of those systems. It has already been demonstrated that the dynamics of a system can influence the overall behavior of the system. An example comes from the study of certain blood diseases (Haurie et al., 1998). The distinction between two forms of a disease comes not from radically different concentrations of substances in the blood or from a large error in the system, but instead from small changes that upset periodic feedback patterns in the concentration of platelets in the blood, which then become very hard to return to equilibrium. Could there be similar problems that are best understood in terms of computational constraints?

Reasoning such as this can also be applied to physics. An example from cosmology is the work of Schmidhuber (1997, 2000) who has argued that it may make sense to choose between different “theories of everything” on grounds of computational complexity.

We demonstrated a number of different ways in which a computational attitude toward the material being studied allows us to ask and answer questions that are distinct from traditional scientific questions. This has been shown to be a valuable perspective in zoology, molecular biology, medicine, psychology, and physics. Answers to questions about the existence of structures in worldly phenomena that store and process information, and their information capacity and computational speed of action, provide valuable scientific information. Without looking at the world with computational eyes, these questions are unlikely to arise.

It is important to note that what we are doing here is different than simply simulating these systems on the computer. In traditional modelling, the questions being asked are identical to those that would be asked if the methodology being used was a conventional experimental one. In contrast, in these arguments, the nature of the model influences the kinds of questions being asked. Another difference is that in this approach, we can learn a lot from when the simulation “goes wrong” and does not reproduce the behavior seen in the real world.

Another important consequence is that this opens the possibility of exploiting the computational capabilities of the world to create new kinds of computation. For example, Feynman (1982) argued that we could respond in two ways to the “problem” that quantum systems are difficult to simulate on computers: we can regard this as a limitation on what computers can do, or we can regard this as an opportunity to build new kinds of computers that are grounded in quantum mechanical concepts.

A generalization of this argument is as follows. If we are faced with a problem in the world that appears to be inherently complex to compute, then we can take one of two stances. First, we can say that the process is “not computable,” “difficult to compute,” “computationally complex,” etc. However, a second possibility arises—we can use that process as the basis of a new form of computation. One of the main points of this chapter is that such an exploitation may already have happened during the course of evolution of the mind, i.e., that the mind has learned to exploit certain somatic systems to do computational processes that might be too slow to carry out if done using on-substrate processes.

Combining these Perspectives

The main aim of this chapter is to consider some of the consequences of postclassical computation for the problems of embodiment. In particular, we can ask the question of when the embodied mind might make use of modes of computation that are not part of the substrate on which the (classically) “computational” parts of the mind are being implemented. These processes could be termed “off-substrate” computations. There are a number of reasons why we might expect such phenomena to be observed. The most radical is that there is some aspect of mental functioning that cannot in any way be carried out using conventional computing machines. Arguments of this kind have been made by Penrose (1989, 1995), who proposed that certain aspects of mental processing cannot be carried out in a (classical) computational fashion. Instead, he proposed that these functions might be carried out by quantum processes working alongside traditional notions of mental functioning.

However, there are a number of other levels of arguments that are weaker than those that show there are mental phenomena that are not computable in the classical sense. One important reason may be that the result is computable, but the use of some alternative process may facilitate faster computation. For example, it was shown by Adamatzky (2001) that a diffusive chemical can solve the problem of finding the shortest route through a maze. The chemical is released through the maze and takes all possible routes. When it reaches a

branch-point, some of the chemical will diffuse in one direction, some in a different direction. This provides a kind of “parallelism on demand” that vastly improves the efficiency of the search contrasted to traditional search. It is feasible that off-substrate processing within the body is carried out for this reason.

Importantly, the potential presence of such off-substrate processes is supplementary to traditional on-substrate computation. One of the ways in which such computation may work is that the traditional on-substrate processing prepares information for input into an off-substrate process, and the output from that process is then further processed using conventional neural processes.

An additional reason for using processes other than conventional computation is the additional input/output capacities that are created by such processes. One feature of the neural-network brain that is typically advantageous is its capacity for processing many aspects of the world in parallel. However, this occasionally is to its disadvantage, for example, when the mind needs to “pull together” to dedicate most of its resources to attending to a danger signal. In such cases, there is no on-substrate mechanism to force the attention of various mind mechanisms toward the danger. By triggering a body state that will require attention by many different mental mechanisms (for example, a sudden feeling of nausea or a rapid heartbeat), there is a “massive synchronization” of mental resources, which provides a counterweight to the usual “massive parallelism” of mental functioning. This will be discussed below in the context of somatic markers.

Extended Somatic Markers

Damasio (1994) introduced the notion of the somatic marker. Somatic markers are bodily states that play a role in cognition, in particular, in the direction of attention. Specifically, a somatic marker is some bodily state that is generated as the consequence of some mental process. This state is then re-perceived by the mind, and as a consequence, the mental state is changed. An example of such a marker is the rapid onset of nausea upon witnessing an act of violence. This bodily state does not have any immediate relevance to the mental state that has generated it, in contrast, say, to a feeling of nausea generated by viewing a plate of rotting food. Some such states might be explained away as side effects, for example, a rapid change of hormone levels upon witnessing violence in preparation for running from the danger might also trigger nausea.

The somatic marker hypothesis suggests that such reactions are not mere side effects. Instead, they are ways of generating rapid shifts of attention, using the body state in an arbitrary fashion to draw mental attention to the current situation.

The presence of the marker in the body draws the mind's attention toward it, and as a consequence, the mind is focused on the meaning of that marker. It is plausible that such phenomena are exaptations [i.e., co-options of previously evolved functions to new ends (Gould & Lewontin, 1979; Gould, 1991)] from unwanted physical reactions to changes in body state, as discussed above.

We can see this as an example of off-substrate computation. The somatic response is being used as a way of carrying out a process (bringing the attention of many mental processes together to focus on a single danger point) that cannot be carried out within the computational model implemented on the substrate.

An interesting question is whether it is important that such markers be somatic, i.e., need they be internal body states? There are two questions to be answered here. First, we can consider why it is important for the marker to exist in the body and not just in the mind. Reasons for this are detailed in Damasio's book, and a particular viewpoint on this is given in the next section of the chapter. This section will focus on the contrasting question: why does the marker need to be constrained to reside within the body? One approach to this draws on ideas from Dawkins' (1992) book, *The Extended Phenotype*.

In biology, the phenotype is the expression of a gene or set of genes in the world, for example, through physical structure or through influences on behavior. For example, we can talk about the "blue-eyed" phenotype versus the "brown-eyed phenotype" of some animal. This is distinguished from the "genotype," i.e., the set of genes of interest. Sometimes, more than one genotype can give rise to the same phenotype (for example, where there are regressive traits).

The difficulty starts when we want to say where the boundary of the phenotype lies. Clearly, certain things are in the phenotype, for example, the sequence of proteins associated with a particular expression of a particular gene. A standard definition would extend this to the whole body—genes influence the growth, development, and activity of the body (alongside other influences).

Dawkins' argument is that it is naive to simply say "everything inside the body, phenotype; everything outside, not." As an example, consider an imaginary species of bird in which the male has a gene that predisposes itself to mate with females that have blue feathers; it could be said that this gene is also a gene for blue feathers in the female, and as a result of the presence of the gene, blue feathers will spread through the female population. To abstract this, the genotype in the male bird is having a phenotypic effect in the female bird. Why should we regard the gene's effect on the feathers of the female bird in any different way than we regard another gene that causes the male bird to have red eyes?

A similar kind of argument can be made about the somatic marker hypothesis. Damasio argued for a body-minded brain in which we create emotions via "somatic markers." These work when parts of the brain recognize an emotionally

charged stimulus, and rather than create a direct link to an action on that stimulus, the “marker,” consisting of a bodily reaction, is created. This is then reperceived by the brain as the basis for action or for rapid alteration of emotional state. Why do these markers have to be physically internal to the body? It would seem that the same reasoning could be applied to markers that I leave in the external world when I have an emotion. For example, if I am anxious, then I might scribble on the pad of paper in front of me, without attending to this scribbling. This could then become a marker, in this case, perceived via the eyes rather than through proprioception. Why should it matter whether I use a bodily state or an external state as the substrate for the marker?

It may be that there are reasons why somatic markers need to be somatic. One could be that the speed of reaction required is just too quick to be capable of being carried out by the external perceptive system. Another more convincing explanation is that the reason we use somatic markers is to communicate with multiple brain regions in a simultaneous and coordinated way, and therefore, we need something that can be perceived in a direct way by different parts of the brain.

This might be a continuum effect. An example of a phenomenon that might be seen as either an external or somatic marker is biting nails when anxious. This is, in many ways, an external physical process, but, nonetheless, we can perceive the nail state internally via soreness of fingers. There must be other similar examples. Perhaps nail-chewing is “causing” the anxiety (in the sense of being part of the causal chain between subconscious perception of an anxiety-producing stimulus and the affective response) rather than being an epiphenomenon of the emotional state.

Hardware Interlocks

In the previous section, we asked why the somatic marker needs to be constrained to the body, and whether it is important to make a body-nonbody distinction. In this section, we address the opposite question: why is it not sufficient for the marker to be a mental marker? Why not just make a “mental note”? While there are circumstances in which a truly somatic marker can get transformed into a mental process in the limbic system, it is interesting to consider whether there might be reasons why the evolution of the mind might have led to the markers being body-centered rather than mind-centered.

One reason may be for safety. In the design of complex systems involving computer-controlled mechanical and electrical devices, it is common for there to be conservative safety devices included in the system, known as hardware

interlocks (Leveson & Turner, 1993; Leveson, 1995). A hardware interlock is a device that is independent of the main control system and that is designed to monitor one small aspect of the system, typically by using its own sensor system. For example, in a radiotherapy device, an interlock might exist that monitors the output of radiation, and if more than a certain amount is let out in 1 minute, the interlock shuts down the device completely.

Hardware interlocks are designed to be parts of the overall system that do not depend on the abstraction offered by the overall control system. For example, they do not take information from the main system sensors, they do not use the main control system (for example, for timing), and they do not sit upon the operating system abstraction used by the controlling structure. To do this would compromise their role as safety-critical components. They provide a reassurance of safety because they are separate; they are independent from the main abstraction. If the main sensors go wrong, or the builder of the controller has misunderstood the relationship between the abstraction offered by the operating system and the real hardware and software, it does not matter.

One important role of the body-mind system is to react quickly and reliably to dangerous phenomena. There would seem to be a *prima facie* case for thinking that if engineers consider the use of such hardware interlocks as an important way of responding to danger in computer-controlled systems, evolution may have created such interlock systems for dangers to animals. It may be that our body-grounded response to danger is a response of this kind. Instead of making a mind-centered judgement about the danger of a situation, we make a rapid decision based on a few simple cues. One characteristic of hardware interlocks is that they typically work on a small number of basic sensors that facilitate a conservative approximation to safety. The same may be true of interlocks in the mind-body system: our sensory system perceives a small number of simple “danger signals” (such as a rapid movement) and triggers an action within the body immediately. This “massive synchronization” acts as a counterpart to the more commonly discussed “massive parallelism” of the neural-network-based mind.

Typically, the fact that the brain is a unified system with all aspects connected and mutually accessible is seen to be to its advantage. Similarly, the unity found in a complex software system is often seen as being to its advantage; instead of having to connect individual components as needed (as might be the case in an electronic system), all information is passed to a central repository and accessed as needed. In some situations, it is necessary, for computers and for minds, for the complete attention of the system to be directed toward one thing. Hardware interlocks provide a way for such responses to “leap out” of the complexity of the control software for certain emergency situations. This nondecomposability, and the consequent need for a powerful way of leaping out of the complex

interactions, would seem to be particularly strong for neural-network-based systems, where the system is highly nondecomposable.

“Just Do It” Computation

Traditional computing is concerned with the construction of machines based on various mathematical models of computing machines. One of the ideas in postclassical computing is that many kinds of transformations in the world can be regarded as computations by consistently ascribing informational values to the objects involved in those transformations. Such “computers” carry out their computations without regard to traditional notions of computational complexity.

For example, proteins fold consistently into complex three-dimensional shapes, despite the complexity of this process. It was shown that a simplified model of the protein-folding problem is NP complete (Berger & Leighton, 1995; Crescenzi, 1998; Fraenkel, 1993), and that an exhaustive search of all configurations would take on the order of 10^{45} years (Levinthal, 1969). Nonetheless, real proteins fold reliably within seconds. Similarly, adding new stars to a galaxy does not slow it for computational reasons.

In such processes, the information (the positions of things in the world) is transformed without any explicit computational effort; the object “just does it.” We can imagine a new kind of computer-based problem solving based on this. In traditional computation, the role of the computer is to compute the solution to a problem directly, by applying algorithms that transform some representation of the input into some representation of the output. However, if a sufficiently flexible set of “just do it” (JDI) devices can be brought together, then there remains the possibility of a new kind of computation. Instead of building a single computing device, computational problem solving is seen as being about the preparation of input for JDI devices, which are then allowed to complete their calculations, output read offs, and interpret them. More complex processes may require a number of JDI processes to be carried out.

A toy example of this is given by Dewdney (1988). This is an $O(n)$ sorting algorithm. A number of lengths of (uncooked) spaghetti are cut to the lengths of the input to the algorithm. These lengths are gathered together and stood upright on the table. The lengths can then be read off one-by-one from the longest downwards. Note that the complexity of the process increases linearly with the number of items being sorted. By contrast, a traditional computer increases in complexity by $O(n \log n)$ while carrying out a sorting process. The process of standing the spaghetti on the table is a massively parallel JDI process.

It is possible that some of the off-substrate processes in the somatically extended mind are of this type, i.e., the on-substrate computation is preparing the input for certain somatic processes to compute. In particular, decision making may be of this kind. The “rational” outcomes of various options are being computed on-substrate, and the final decision is made by a “gut feeling.” A related argument has been made by Evans (2002). His argument is that one of the roles of emotion is to solve the “search problem,” i.e., the problem of knowing where to stop when calculating the consequences of a decision, that emotions “prevent us from getting lost in endless explorations of potentially infinite search spaces.”

These “just do it” processes are similar to what Kauffman (1993, 1995) described as “order for free.” For example, in this viewpoint, the occurrence of the Fibonacci sequence in the phyllotaxis of pinecones is not some miracle of nature; it is simply the energetically cheapest way of generating the desired structure.

A main part of his argument is that evolution will exploit such “order for free” as an energetically cheap way of constructing complex organisms. Instead of constructing new devices to achieve action, evolution pieces together various devices available in the world to construct complex biological machines. In this section, we made a similar argument with regard to the evolving mind exploiting available computational devices in the world.

Conclusion

Does a functioning mind need a functioning body? Perhaps yes, if some parts of mental functioning are delegated to off-substrate processes, either for the purposes of computational efficiency or because certain processes (such as the “massive synchronization” required to respond efficiently to danger) are not available in the on-substrate model.

One perspective from which to view this is that of postclassical computation. From this point of view, computation is seen as a property of things in the world, rather than only as a property of specially constructed computing machines. This allows us to ask questions about the computational capabilities of many different objects. In particular, we can ask questions about the computational capabilities of on-substrate processes and compare these to off-substrate processes. In some cases, the off-substrate processes may have properties that are not available on-substrate or that operate in a more efficient fashion than their on-substrate equivalents. In such cases, we have a *prima facie* case for considering that the off-substrate way of realizing that process might have been favored during the evolution of mind.

A particular example of an off-substrate process is the somatic marker, where a body state is used as shorthand for some mental state that needs to be rapidly appreciated by a number of parallel mental processes. This postclassical computation stance allows us to consider the relationship between these on- and off-substrate processes, and it provides the beginning of a way of determining whether particular processes are likely to be carried out on-substrate using connectionist networks, or whether they are likely to be delegated to other parts of the body.

References

- Adamatzky, A. (2001). *Computing in nonlinear media and automata collectives*. Bristol: Institute of Physics Publishing.
- Berger, B., & Leighton, T. (1998). Protein folding in the hydrophilic–hydrophobic (HP) model is NP-complete. *Journal of Computational Biology*, 5(1), 27–40.
- Brooks, R. A. (1991). Intelligence without representation. *Artificial Intelligence*, 47, 139–159.
- Crescenzi, P., Goldman, D., Papadimitriou, C., Piccolboni, A., & Yannakakis, M. (1998). On the complexity of protein folding. *Journal of Computational Biology*, 5, 423–465.
- Damasio, A. (1994). *Descartes' error: Emotion, reason and the human brain*. New York: Gosset/Putnam Press.
- Dawkins, R. (1982). *The extended phenotype*. Oxford: Oxford University Press.
- de Boer, R. J., van der Laan, J. D., & Hogeweg, P. (1993). Randomness and pattern scale in the immune network: A cellular automata approach. In W. Stein & F. J. Varela (Eds.), *Thinking about biology*, Santa Fe Institute Studies in the Sciences of Complexity: Lecture Notes (Vol. III, pp. 231–252). Reading, MA: Addison-Wesley.
- Dewdney, A. K. (1988). *The armchair universe: an exploration of computer worlds*. New York: Freeman.
- Evans, D. (2002). The search hypothesis of emotion. *British Journal for the Philosophy of Science*, 53, 497–509.
- Feynman, R. P. (1982). Simulating physics with computers. *International Journal of Theoretical Physics*, 21, 467–488.

- Fraenkel, A. (1993). Complexity of protein folding. *Bulletin of Mathematical Biology*, 55(6), 1199–1210.
- Gibson, J. J. (1977). The theory of affordances. In R. E. Shaw & J. Bransford (Eds.), *Perceiving, acting and knowing*. Mahwah, NJ: Lawrence Erlbaum.
- Gould, S. J. (1991). Exaptation: A crucial tool for evolutionary psychology. *Journal of Social Issues*, 47, 43–65.
- Gould, S., & Lewontin, R. (1979). The spandrels of San Marco and the Panglossian paradigm: A critique of the adaptationist program. *Proceedings of the Royal Society of London, Series B*, 205(1161), 581–598.
- Grundy, F. (1998). Computer engineering: Engineering what? *AISB Quarterly*, 100, 24–31.
- Haurie, C., Dale, D. C., & Mackey, M. C. (1998). Cyclical neutropenia and other periodic haematological disorders: A review of mechanisms and mathematical models. *Blood*, 92(8), 2629–2640.
- Johnson, C. G. (2003). What kinds of processes can be regarded as computations? In R. Paton, H. Bolouri, M. Holcombe, J. H. Parish, & R. Tateson (Eds.), *Computing in cells and tissues: Perspectives and tools of thought*. Heidelberg: Springer.
- Kauffman, S. (1993). *The origins of order*. Oxford: Oxford University Press.
- Kauffman, S. (1995). *At home in the universe*. New York: Penguin.
- Leveson, N. (1995). *Safeware: System safety and computers*. Reading, MA: Addison-Wesley.
- Leveson, N., & Turner, C. (1993). An investigation of the Therac-25 accidents. *IEEE Computer*, 26(7), 18–41.
- Levinthal, C. (1969). How to fold graciously. In J. T. P. DeBrunner & E. Munck (Eds.), *Mossbauer spectroscopy in biological systems* (pp. 22–24). Illinois: University of Illinois Press.
- Norman, D. A. (1998). *The psychology of everyday things*. New York: Basic Books.
- Paton, R. (1996). Metaphors, models and bioinformation. *BioSystems*, 38, 155–162.
- Penrose, R. (1989). *The emperor's new mind: Concerning computers, minds, and the laws of physics*. Oxford: Oxford University Press.
- Penrose, R. (1995). *Shadows of the mind*. New York: Vintage.
- Quick, T., Dautenhahn, K., Nehaniv, C. L., & Roberts, G. (1999). On bots and bacteria: Ontology independent embodiment. In *Proceedings of the Fifth European Conference on Artificial Life*.

- Schmidhuber, J. (1997). A computer scientist's view of life, the universe, and everything. In C. Freksa, M. Jantzen, & R. Valk (Eds.), *Foundations of computer science: Potential — theory — cognition* [Lecture Notes in Computer Science] (pp. 201–208). Heidelberg: Springer.
- Schmidhuber, J. (2000). *Algorithmic theories of everything*. [Technical Report 20-00.] IDSIA.
- Stepney, S., Clark, J., Tyrell, A., Johnson, C., Timmis, J., Partridge, D., Adamatsky, A., & Smith, R. (2003). *Journeys in non-classical computation*. National E-Science Centre Grand Challenge report.
- Stonier, T. (1990). *Information and the internal structure of the universe*. Heidelberg: Springer.
- Wilson, M. (2002). Six views of embodied cognition. *Psychological Bulletin and Review*, 9(4), 625–636.
- Ziemke, T. (2001). Are robots embodied? In C. Balkenius, J. Zlatev, C. Breazeal, K. Dautenhahn, & H. Kozima (Eds.), *Proceedings of the First International Workshop on Epigenetic Robotics: Modeling Cognitive Development in Robotic Systems*. Lund University Cognitive Studies (Vol. 85). Lund, Sweden.
- Ziemke, T. (2003). What's that thing called embodiment? In *Proceedings of the 25th Annual Meeting of the Cognitive Science Society*. Mahwah, NJ: Lawrence Erlbaum.