

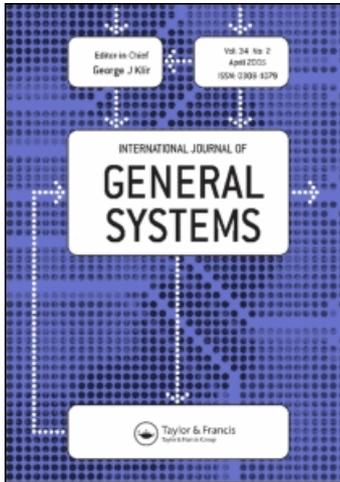
This article was downloaded by: [State University of New York]

On: 11 May 2009

Access details: Access Details: [subscription number 908120893]

Publisher Taylor & Francis

Informa Ltd Registered in England and Wales Registered Number: 1072954 Registered office: Mortimer House, 37-41 Mortimer Street, London W1T 3JH, UK



International Journal of General Systems

Publication details, including instructions for authors and subscription information:

<http://www.informaworld.com/smpp/title-content=t713642931>

Information and regulation in robots, perception and consciousness: Ashby's embodied minds

Peter Mario Asaro ^{ab}

^a Center for Cultural Analysis, Rutgers University, New Brunswick, USA ^b Department of Media Studies and Film, New School University, New York, USA

Online Publication Date: 01 February 2009

To cite this Article Asaro, Peter Mario(2009)'Information and regulation in robots, perception and consciousness: Ashby's embodied minds',International Journal of General Systems,38:2,111 — 128

To link to this Article: DOI: 10.1080/03081070802633643

URL: <http://dx.doi.org/10.1080/03081070802633643>

PLEASE SCROLL DOWN FOR ARTICLE

Full terms and conditions of use: <http://www.informaworld.com/terms-and-conditions-of-access.pdf>

This article may be used for research, teaching and private study purposes. Any substantial or systematic reproduction, re-distribution, re-selling, loan or sub-licensing, systematic supply or distribution in any form to anyone is expressly forbidden.

The publisher does not give any warranty express or implied or make any representation that the contents will be complete or accurate or up to date. The accuracy of any instructions, formulae and drug doses should be independently verified with primary sources. The publisher shall not be liable for any loss, actions, claims, proceedings, demand or costs or damages whatsoever or howsoever caused arising directly or indirectly in connection with or arising out of the use of this material.

Information and regulation in robots, perception and consciousness: Ashby's embodied minds

Peter Mario Asaro^{ab*}

^aCenter for Cultural Analysis, Rutgers University, New Brunswick, USA; ^bDepartment of Media Studies and Film, New School University, New York, USA

(Received 13 February 2008; final version received 6 October 2008)

This article considers W. Ross Ashby's ideas on the nature of embodied minds, as articulated in the last five years of his career. In particular, it attempts to connect his ideas to later work by others in robotics, perception and consciousness. While it is difficult to measure his direct influence on this work, the conceptual links are deep. Moreover, Ashby provides a comprehensive view of the embodied mind, which connects these areas. It concludes that the contemporary fields of situated robotics, ecological perception, and the neural mechanisms of consciousness might all benefit from a reconsideration of Ashby's later writings.

Keywords: representation; embodiment; cognition; information theory; requisite variety

1. Introduction

My interest in the work of W. Ross Ashby stems from his historical legacy and lasting impact on the future of research into artificial intelligence (AI) and cognitive neuroscience. There is, however, an ambiguity in this legacy. Even though some of Ashby's early ideas on learning mechanisms and the use of phase space descriptions of complex systems were enormously influential, his particular vision of the mind and the brain seems to have been largely ignored by both his contemporaries and our own. I want to try to give Ashby his due, and consider how certain aspects of his work have been rediscovered, somewhat painfully, in fields like AI and cognitive neuroscience, and some of his other highly promising ideas have yet to be properly explored.

In particular, I wish to examine the conception of embodied representation that Ashby developed towards the end of his career, some of it in collaboration with Roger Conant. This conception was elaborated in several papers published in the period from 1967 to 1972 (Ashby 1967, 1968a, 1968b, 1972, Conant and Ashby 1970). In those papers, Ashby articulated a conception of embodied representation which was influential in control theory and systems engineering, but largely overlooked by AI and cognitive neuroscience. Perhaps because it appears on the surface to be merely an extension of control theory, it was not seen as a candidate for being a cognitive theory of mental representation. Yet it was these insights into embodied representation that were partially rediscovered in what might be called the reformation of orthodox AI in the late 1980s, by researchers like Agre and Chapman (1987) and Brooks (1991a, 1991b) in their work on activity theory and the

*Email: peterasaro@sbcglobal.net

subsumption architecture respectively. Further, I believe that Ashby's conception of embodied representation could be fruitfully brought into dialogue with the powerful ecological theory of perception developed by a contemporary of Ashby's, Gibson (1950, 1975, 1979). Such a dialogue could inform and advance the implementation of new computational and robotic systems. Finally, the way in which Ashby envisions representations as embodied and related to sensory-motor activities is compatible with the architecture of a recent mechanistic neurobiological theory of consciousness put forward by Cotterill (1997, 1998), which respects the role of motor activity in organising perception and planning. At the very least, I want to argue that these more recent theories of situated robotics, ecological perception and the mechanisms of consciousness are compatible with Ashby's approach to the embodied mind, and are in some sense part of Ashby's legacy in the study of the mind and brain but I also want to argue that future work in these areas of research could benefit from a rereading of Ashby's later papers, and a consideration of their implications and suggestions for further work, a few of which I will outline.

These papers from the last half-decade of Ashby's life interest me because they come at the most mature stage of his thought, and because they are among the most reflective and speculative pieces in his long list of publications. Thus, they contain Ashby's hopes for the extension of his earlier analyses of learning and adaptation towards a more robust theory of perception and mind, both in terms of theoretical understanding and practical engineering, as these were always closely related in his mind. These papers also exemplify his approach in that they clearly and explicitly consider the activity of the mind within the everyday demands of human life in the natural world, a set of problems that AI continues to have little success in dealing with.

Much of the power of Ashby's (1952) approach in *Design for a Brain* is the way in which the brain learns to deal with the world through the control of its actions and interactions with an environment, rather than through constructing elaborate symbolic representations (Asaro 2008). The Homeostat, an adaptive system built by Ashby in 1948 and described in this book, seeks an ultra-stable state in which it is able to act reliably to maintain the stability of its essential variables against disturbances. It does not represent the world in any formal sense – it contains no logical expressions or syntactic symbols – but it can learn to adjust its parameters for interacting with the world so as to adapt to various disturbances and maintain certain desirable internal states. I will investigate what kinds of problems are met in trying to extend this low-level physiological insight into a theory of perception. I will also consider how such a theory might provide a basis for more complex symbolic representations such as language or images. I argue that recent interest and research in embodied intelligence and situated robotics have rediscovered these ideas some 20 years after Ashby first published them, although contemporary AI could still benefit from developing other aspects of Ashby's approach that still languish.

Before continuing, I also want to remark that the notion of the embodied mind, and the motor-centric view of the brain, were not unique to Ashby. McCulloch's (1965) anthology was entitled *Embodiments of Mind*, while *Embodied Cognition* by Varela *et al.* (1991) and numerous other works from the cybernetic perspective have articulated views of the mind as an embodied information processor. There are also many researchers from neuroscience and neurophysiology – notable examples include Sir Charles Sherrington, Lord Adrian, and Roger Sperry – who take a similar motor-centric view of the mind, or at least consider perception, memory and planning largely in the context of potential motor actions. What I want to emphasise in Ashby's work is not the novelty of the idea of embodied representation itself, but the specific character of embodiment in terms of information

flows, feedback coordination and its relation to essential, vital states, and the structures in which these are organised in adaptive and active systems. In the next few sections I want to address specific approaches to situated autonomous robotics, perception and consciousness, first as works continuous with the critical aspects of embodied representation. Then to consider the specific insights from Ashby's work that suggest ways of extending these theories in directions which have not, to my knowledge, yet been fully explored. In that respect, it is an effort to show how his work continues to be relevant despite having been overlooked as a theory of embodied mental representation.

2. Embodied representations

I want to begin by sketching out the conception of an embodied representation which Ashby developed in the late 1960s. I should say that it was an extension of his earlier work in *Design for a Brain*, but that his earlier work was more concerned with promoting certain techniques of analysis. The underlying view of the mind was as a complex system seeking various equilibriums, and emphasised the separation of stable 'essential variables' from the dynamic variation of active variables. While these were important contributions at the time, *Design for a Brain* does not actually go very far in helping to design artificial brains for robots, or helping to explain the details of perception and consciousness in biological brains. I believe, however, that his later work goes much further in these directions, though the ideas were published in isolated articles and were never brought together by Ashby himself into a unified view. The appropriate place to start describing this work would be with the most influential of his later papers, co-authored with his student Roger Conant.

In their classic paper, Conant and Ashby (1970) present a compelling case that every good regulator of a system must be a model of that system. I believe this paper contains within it the foundations for a practical embodied theory of mind with significant philosophical implications. What this paper demonstrates is how we should think about the relationships between the world, a regulator (or any system which manages its interactions with the world) and the informational transactions between them.

Let us begin by considering the two kinds of systems described in that article. In the first of these (see Figure 1), we see a disturbance connected directly to the regulator prior to reaching the system being controlled. In the second (see Figure 2), the disturbance acts on the system before that information reaches the regulator. There are several interesting things to note in the difference of these flows of information.

A regulator of the second sort will never be perfect. This is because it must wait for the system to depart from a stable or desired state before it can be returned to that state by the regulator. It only recognises actual changes in its immediate field of sensitivity. The thermostat and governor are examples of this type – the furnace will not kick on until the temperature of the room drops, and the steam pressure will not begin to build until the steam engine slows and the valves are closed. We can call such systems 'reactive'. The other type of system detects potential disturbances before they actually disturb the system. That is, they obtain information about the *causes* of the system's movement from the desired state, and can act on this information rather than the information about the

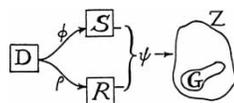


Figure 1. Cause-controlled regulation.

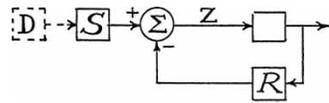


Figure 2. Error-controlled regulation.

effects on the system. With this relation to information, the regulator can act *before* the effects are realised. We can call such systems ‘proactive’. Examples of this type include predictive regulators which employ inverse dynamics. It follows that such proactive and predictive regulators are in some sense models of the relationships that exist between potential disturbances and regulated systems, at least if they are *good* regulators.

Notice that I say they *are* models, rather than that they *use* models. This is intended to intimate that the model is implicit in the organisation of the regulator and not an explicit symbolic model. The governor represents engine speed in the sense that it embodies a regulator which effectively controls engine speed. Nor should we be confused by the numerical representations on the face of a thermostat, which are really just an indication of the state of the thermostat at the moment (not the temperature of the room) so as to help us in regulating the regulator. Nor should we be confused by the fact that, as observers of the system, we could use the regulator as a model of the system for purposes other than regulation. The sense of model employed here is one I have discussed elsewhere (Asaro 2006) as a *working model* – it models a system in virtue of its dynamic material actions over time, not in virtue of a syntactic structure. There is a sense of analogy and isomorphism implicit in any model, however, and in many models it may be possible to articulate that isomorphism in a syntactic formalism – but this is not the essential character of a working model.

Now let us consider where perception fits into Figures 1 and 2. What appear here as simple arrows ‘p’ indicate a flow of information. Aspects of D and/or S transmit some amount of information about environmental states to the regulator. Just how this happens is, in fact, the very essence of perception. There are, however, several ways we might construe this process even here, so let us consider some. First, let us consider the perfect and optimally efficient regulator. Such a regulator would always act to maintain the system in its ideal state, without disturbances. As I just explained, the regulator must be proactive and predictive in order to achieve perfect regulation. By saying a regulator is optimally efficient, we mean that it requires only the minimum amount of information from D and S necessary to select the proper action. Saying just what that is for any particular case, however, is not so easy. In Ashby’s (1968a, 1968b) paper ‘Information Processing in Everyday Human Activity’, he addresses these difficulties directly. Following Shannon (1948), the quantity of information in a signal depends upon the set of messages from which the signal is distinguished as being the one intended message, as well as the probabilities of the receipt of each particular signal.

For example, a thermometer might receive only four bits of information, such as the temperature of the room in degrees Fahrenheit, and from that determine whether to turn on the air conditioning, or the furnace, or neither. To be optimal, we could give it just two bits, one for turning the AC on or off, and one for turning the furnace on or off (such as simply allowing a column of mercury to open and close the appropriate circuits). This is the case of the trivially reactive controller, which reacts directly to its inputs with an action. While this kind of thermostat might be sufficient, but cannot be a perfect regulator because the temperature will fluctuate around the ideal, its variance depending on the force of the

disturbances and the effectiveness and latency of the furnace and AC in altering the room temperature.

How could we make it proactive? Perhaps by measuring the outside temperature, the humidity, the wind speed, the airflow through opening doors and windows, the current and near-future amounts of sunlight coming through windows or falling on the exterior of the building, and performing a function which indicates what to do with the furnace and AC as outputs. That is to say that, while our thermostatic regulator still only has three possible actions (cool, heat, nothing) we have introduced a much more complicated set of potential inputs. These may not be all the relevant inputs, indeed there may be a limitless number of possible disturbances. Still, the point is that there could be a very large amount of information which a proactive regulator might need to obtain in order to choose the appropriate action, and the choice itself might be quite difficult given the relevant information. The 'processing' of this information amounts to making the appropriate selection of actions from these preceding states. There is a sense in which the proactive system is a model of the potential future states of a system – it is anticipatory or predictive. As such, it must model not only the states of the system it regulates, but the transitions between states – the dynamics of the system. Thus, the regulator must model a *dynamic* system. In many cases, such as the governor, this need not entail a complex symbolic representation, especially if the regulator is itself a dynamic system.

There are several things I want to emphasise from this brief sketch of embodied representations. First and foremost, the notion of the regulator is fundamentally a matter of control of action. In the absence of action, or at least potential action, it does not make sense to talk about mind, as it is fundamentally bound up with behaviour. Compared to more traditional theories of mind built upon mental images, representation and perception, subjective experience, or structured knowledge of the world – all open to passive observers and not requiring actors – this conception redistributes the relative weights placed upon perception and motor control. Some philosophers might even think we have succeeded in explaining mind once we have explained how thoughts get into the mind, without worrying at all about what consequences they have in action apart from 'decisions to believe or act'. Only such a philosopher could think that motor control is a trivial achievement while perception presents a real philosophical 'problem', even though the two are deeply interconnected. Even if they would not assert such things, the casually held assumption that inner thought resembles perception more than it does action is a reflection of this bias.

It should be noted that the position I believe Ashby to be endorsing aims to balance the two – that perception and action have equal shares in the structure of embodied representation, and where the role of perception is largely to modulate and regulate action. That is also to say that the sensory and the motor are both equally implicated in the feedback loops which regulate behaviour. This, of course, has relevance for how we design robots. Which brings me to my second point – that the recognition of the significance of motor control ought to influence how we think about perception, as well as memory and consciousness. Even though Ashby did not have time to fully develop these implications himself, I believe he was aware of them.

While the individual regulator is a good model for understanding an individual reflex, behaviour is often much more complicated than this. It is not surprising that, in Ashby's view, a system will need a robust variety of behaviours in order to respond to a complex environment, lest we forget his Law of Requisite Variety (Ashby 1956, p. 207) and it is precisely in the extension and scaling-up of these simple mechanisms to account for the

behaviour of the large and complex system of the brain that the real difficulties lay. I think Ashby recognised this quite early, and this recognition motivated his studies of information flow and organisation in large systems. I think he also began to see some of the directions in which the solutions would be found. One aspect of the solution was that even as systems get very large, they tend to naturally organise themselves into smaller, stable subsystems. Thus, we need not bother trying to explain unnecessarily complex systems, as long as we can explain their organisation into systems of interrelated subsystems (Ashby 1972).

The basis of all such systems, or the atomic mechanism, is the feedback loop of sensation, reflex action and adaptation and he saw that these atomic reflex mechanisms can be built up artificially, or allowed to self-organise naturally, into layers. These layers can form much more complex organism behaviours such as foraging, long-distance navigation and sophisticated multi-step tasks. The use of atomic reflex mechanisms, and their layering, is the basis of recent work in autonomous robotics, called the subsumption architecture. I will consider this work in autonomous robotics in the next section, but first want to mention how the subsequent sections will return to the theme of the mind as a regulator of motor activity.

Layers of feedback motor control are also essential in the production of complex perceptions such as imagined states of the world, and objects and events distant in time and space. The basic insight of Gibson's ecological theory of perception is that even abstract mental representation of distant times and places is built upon the same embodied interactions as simple sensory-motor mechanisms. This is to say that ecological perception shares Ashby's fundamental perspective on the importance of motor control, and thus the structure of visual perception is more likely to be shaped by the features of a creature's body and its interactions with an environment, than it is to be shaped by abstract geometric properties, such as Euclidian geometry. I will return to this point in the fourth section of the paper.

Another key element of Ashby's view was appropriate selection and so at some point in the elaborately braided networks of layered reflexes – probably somewhere at the higher-levels though not independently of the lower – the human brain must have an opportunity to select among alternative actions. Recent work on the neurological basis of consciousness has elaborated a notion of a triple-feedback loop of perception, potential actions and memory in the regulation of performed actions, which appears to provide a mechanistic explanation of how the highest level of perception and motor control in sophisticated neural systems actually requires consciousness in order to achieve appropriate selection and again, the guiding principle of this theory is that the brain developed as an elaborate regulator of muscular control. I will examine this in the fifth section of the paper.

3. AI and autonomous robotics

It might seem that a discussion of autonomous robotics would be more apt for a paper on the legacy of W. Grey Walter, a colleague of Ashby's at the Burden Neurological Institute, than for a paper on the legacy of W. Ross Ashby (Asaro 2006). However, while Walter actually built some autonomous robots at about the same time that Ashby built his Homeostat, it was Ashby who came closest to formulating a theory of embodied representation strikingly close to the one that has been taking over the field of autonomous robotics for the past 15 years. The irony is that neither Walter nor Ashby had a direct

influence on this recent work, and so the basic theory of autonomous robotics had to await independent rediscovery some 20 years after Ashby published his formulation of it.

The approach of AI from the late 1960s, throughout the 70s, and well into the 80s, was to model the world explicitly in centralised representations, commonly a 3D model of the world, or a propositional model of its states and potential transitions. These were mostly static models, not dynamic apart from the actions taken by the computer program or robot. These were usually partial worlds, or micro-worlds limited and greatly simplified to make them easier to deal with. This is not what Ashby believed the mind was doing. While even Marvin Minsky (1963) recognised that Ashby's early work on learning mechanisms had led to the widespread use of search procedures, I do not think Ashby believed that studying this aspect of intelligence in isolation, as AI was doing in the 1960s, was a productive route to understanding or designing brains. To design brains, one ought to consider the 'place of the brain in the natural world' and in 'everyday human activity'.

Various critiques of the traditional symbolic approach to AI have been levelled over the years. In the 1980s connectionism, with its neural network models, revived some aspects of the cybernetic work that had predated symbolic AI. But connectionism was also overly limited in its attempts to develop systems with the essential properties of brains, and settled for somewhat superficial properties of brain-like network architectures that exhibited some interesting and useful properties. However, these approaches still failed to build systems which could effectively engage with the complex world of human activity.

Then in 1987, two important critiques were presented which focused on the importance of atomic reflex mechanisms for intelligent behaviour. Philip Agre was the first to level severe criticisms of AI along these lines, and to offer a workable alternative. In a 1987 paper with Chapman, he outlined a new approach called activity theory, and exemplified it with a computer program called Pengi, in which the behaviour of a computer-game penguin is directly mapped to states of its world (Agre and Chapman 1987). Pengi is essentially a programmed mechanism which displays reflexes, reacting directly to the state of its environment, unconditioned by learning or even previous states of the environment (i.e. it has no memory). This is precisely what Ashby had in mind when he describes reflex mechanisms in his later work: a mapping from disturbances to actions. Agre (1997) has written an excellent and comprehensive critique of AI, and the alternative he proposes, but today I will focus on another approach, which extends this basic concept by layering reflexes.

Rodney Brooks also presented his scathing critique of AI in 1987, though the most often cited version was published later (Brooks 1991b). In his now classic paper 'Intelligence without Representation' he points out that many simple animals, such as flies, could not possibly be constructing centralised 3D models of the world in their head and yet, these flies can outperform the most sophisticated robots when it comes to dealing with a complex dynamic world in real time. Based on this observation, he suggested that rather than trying to model the most sophisticated aspects of human intelligence on computers, we should instead try to build very simple animals with insect-like intelligence and work up. He likened the current situation in AI to 19th century aviation engineers trying to build the first plane by copying the things they saw during a time-machine visit to see a modern 747 jet airliner.

To achieve this goal, Brooks (1991b) proposed a new approach which he called the subsumption architecture. The basic idea was that robotic creatures could be built which were quite adept at getting along in the world if we changed our ideas about how they do it. Instead of centralised representation, planning and decision-making, robots should be viewed as bundles of reflexes. Instead of a single look-up table that maps a vast range

of sensory inputs to complex sets of motor actions, robots should have a layered architecture with many simple finite-state machines linking senses to motor controls:

In fact, we hypothesize that all human behaviour is simply the external expression of a seething mass of rather independent behavior without any central control or representations of the world . . . The two key aspects of the subsumption architecture are that (a) it imposes a layering methodology in building intelligent control programs, and that (b) within each network the finite state machines give the layer some structure and also provide a repository for state. (Brooks 1991b, 140)

Compare the proposed subsumption architecture of Brooks, to the proposals made by Ashby, 20 years earlier in an effort to explain how ‘instincts’ appear more complex than ‘reflexes’:

It is now known, however, that this property of reacting to combinations and *relations* between stimuli, is readily obtained from the mechanism, if the mechanism works in stages or levels so that the first level “computes” various functions of the primary stimuli, then the later levels compute functions of these functions, and the final stage acts only if these “computational” processes have resulted in some actual physical event at the penultimate stage. In this way, any defined function over the primary stimuli, however complex or subtle it may be, can be transformed, in a purely mechanistic way, to a physical event suitable to act as physical cause for the instinctive action. The apparent distinction between reflex and instinct were based, mostly unconsciously, on a one-level model: stimulus-to-response, without intermediate processing. (Ashby 1967, p. 101)

We can see in this passage that he recognised the fundamental importance of layering simple mechanisms in order to achieve complex behaviours. His framework is also that of state-determined mechanisms, and distributed, rather than centralised, representation and control.

Brooks’ approach was successfully implemented in a series of mobile robots. It was so successful, in fact, that Brooks later became head of the AI Laboratory at MIT. Yet, neither Agre nor Brooks deal with the more fundamental processes of learning and adaptation which concerned Ashby. Neither activity theory nor the subsumption architecture can explain or offer a method for how a robot could learn, or automatically construct, the finite-state machines which transform sense data into behaviours. This requires the careful work of system designers. Brooks indicates that he was ‘working on it,’ but he provides no details and says only that it appears to require some form of centralised representations because the necessary techniques are not available.

The problem of learning is not difficult for simple sets of relations and single-layered architectures, but rapidly becomes very difficult with multi-layer architectures. Researchers working on adaptive multi-layered architectures would be well advised to reread Ashby’s papers on adaptation and information flows in large-scale adaptive architectures. In those papers Ashby presents many interesting ideas and useful equations for how large networks of such mechanisms naturally organise themselves into distinct subsystems. As we will see in a moment, learning to regulate motor activities at a high enough level complexity, and in a sufficient time to keep the system engaged with its environment, will ultimately require a mechanism of coordination and deliberation – consciousness. But first, let us consider perception more carefully.

4. Perception

Now that we have seen that embodied representations are essentially layered reflexes, tying together sensation with action, what does this mean for perception? As was the case for AI in general in the 1970s and 1980s, the computational approaches to vision and

perception followed an approach aiming to construct elaborate 3D models of the world by extracting information from 2D digital images. These approaches were largely unconcerned with action, seeking only to build the internal model, which could be used, presumably, for all sorts of things, but in practice were usually not used at all. The most highly articulated version of this approach came in the form of David Marr's (1982) computational theory of vision. Marr's approach was to use sophisticated mathematical techniques to extract depth information, edges, and plane orientations from 2D images. From these, schematic 3D structures could be assembled and serve as a computational model of the world.

Marr's models give no consideration to the nature or structure of motor activity, or its regulation. There is a sense in which it is sensitive to the structure of the world, but this is only to the extent that the world is made up of regular geometric surfaces, and he thus argues that perception must have evolved ways of extracting this 'intrinsic structure of the world'. In philosophy, this is called the subject/object distinction, according to which perception is the process whereby a representation of the object enters into the subject's mind. The computational approach starts and ends with the information contained in a single, or perhaps two, images. It does not consider the interactions of the subject with its environment over time, or what this might contribute to the perceptual information. While more contemporary approaches seek to extract additional information from image sequences, they are only beginning to explore the value of correlating that information with motor activity (Wörgötter *et al.* 2004).

The theory which offered the strongest challenge to the computational theory was James J. Gibson's ecological theory of perception. Gibson first began to publish his ideas in 1950, and his final formulation came in 1979. This theory, I argue, is an elaboration of the concept I have been calling embodied representation. That is, it takes the relationship between the perceiving system and its environment as primary. And in so doing, it places the proper significance on motor activity – both for its ability to inform the system apart from a single sensation or image, and for the specific ways in which it structures perception of the world. Gibson's theory is best remembered for introducing the concept of 'affordances,' a term he used to describe how the mind perceives the world in terms of the various activities it affords to the body. This is exactly what we should expect once we understand that perception evolved primarily to serve action, i.e. that it is an aspect of embodied representation. I want to take a moment to give a very brief sketch of some critical aspects of Gibson's theory, before remarking on how it might be extended based upon some additional insights from Ashby's work.

Gibson begins his account of perception by considering Berkeley's puzzle of depth perception from the *New Theory of Vision*¹ and dissolves the problem by denying that we perceive the absolute time and space of Newton. He states in response to this puzzle that,

Distance therefore is *not* a line endwise to the eye as Bishop Berkeley thought. To think so is to confuse abstract geometrical space with the living space of the environment. It is to confuse the Z-axis of a Cartesian coordinate system with the number of paces along the ground to a fixed object. (Gibson 1979, p. 117)

The basic notion is that the mental representation of space is not absolute in the Newtonian sense, as Kant believed, but is instead a consequence of our interaction with the world. It is the same for time as it is for space:

There is no such thing as depth perception, or the perception of distance, or the third dimension, or in fact the perception of space. There is only the perception of textured surfaces and what I call the "layout" of these surfaces . . . The true question is how we perceive all these surfaces with their inclination to one another, and their curvatures and their edges . . . This problem is the

crucial one, not the problem of the third dimension. How do we perceive that portion of the layout of the world that is temporarily *out of sight*? This question, please note, has reference to time as well as space . . . I now want to argue that the perception of time is a puzzle of the same sort that the perception of space has been – an insoluble one. There is no such thing as the perception of time, but only the perception of events and locomotions. These events and locomotions, moreover, do not occur in space but in the medium of the environment that is rigid and permanent. Abstract space is a sort of ghost of the surfaces of the world, and abstract time is a ghost of the events of the world. (Gibson 1975, p. 295)

From this we can see how interactions with the environment, such as movement through it coupled with memory traces, can lead to representations of things distant in time and space. That is to say, it is bodily activity itself which structures the perception of time and space – the essence of the most sophisticated and complex representations of the world are fundamentally embodied representations.

Gibson's theory considers visual sensation as the impending *ambient optical array* – the totality of the information-bearing light that falls upon both eyes. Seeing an object, for Gibson, amounts to discerning a particular *closed-contour form* from the ambient optical array that is completely filled with various forms and textures. The individual discernible components of the array consist of *solid visual angles* which are packed together. This is very different conceptually than the computational view that holds that the image falling on the eye is a bit-map array of the discrete rays of light taken up by the rods and cones in a fashion similar to a video camera or digital image:

There are several advantages in conceiving the optic array in this way, as a nested hierarchy of solid angles all having a common apex instead of as a set of rays intersecting at a point. Every solid angle, no matter how small, has form in the sense that its cross-section has a form, and a solid angle is quite unlike a ray in this respect. Each solid angle is unique, whereas a ray is not unique and can only be identified arbitrarily, by a pair of coordinates. Solid angles can fill up a sphere in the way that sectors fill up a circle, but it must be remembered that there are angles within angles, so that their sum does not *add* up to a sphere . . . The structure of an optic array, so conceived, is without gaps. It does not consist of points or spots that are discrete. It is completely filled. Every component is found to consist of smaller components. Within the boundaries of any form, however small, there are always other forms. This means the array is more like a hierarchy than like a matrix and that it should not be analyzed into a set of spots of light, each with a locus and each with a determinate intensity and frequency. (Gibson 1979, p. 68)

Gibson is arguing directly against digital imagery as a model of vision in saying that there are no atomic sensations like bits taken up by the eye, but rather solid angles which are continuous wholes and which can be further divided into smaller wholes but never atomic or discrete parts. It should also be noted that optical invariants arise from consistencies in these solid angles. This gives us a much better way of thinking about visual processing in the brain, which we know from the very first layers of neurons in the retina to be using lateral inhibition, i.e. looking at discontinuities and contrasts in the visual field both in terms of light and motion. Lateral inhibition also occurs and the next level of visual processing in the optic nerve and the layers of the lateral geniculate nucleus, and acts not only upon light properties such as intensity contrasts, but appears to do this across the various colour channels, and at multiple scales of spatial and temporal resolution. That is to say, lateral inhibition is not just a trick for finding edges in a visual scene, but rather extracts the multi-dimensional contrasts that *are* the visual scene in the form of nested solid visual angles. Thus, the information the visual system extracts from the ambient visual array is really just the disturbances and discontinuities in quality, space and time. When these disturbances and discontinuities occur persistently and reliably, they come to be treated as the invariant structures of the perceived world.

But perception does not end with the information provided by vision. This information must be assimilated with other information, memory, the current state of the body and the potential actions of the body. While Brooks is right that much of what creatures do can be done by independent reflexes interacting in only limited ways, the kind of perception humans are capable of indicates that they are able to integrate information from multiple modalities, and use this information to coordinate very complicated motor activity sequences. A tennis player combines information from the sound of their opponent's racket, visual information about speeds and angles, elastic collisions, the positions of their own feet and arms, balance and bodily momentums, etc., in order to execute a sequence of muscular contractions which places their body in a suitable position to swing a racket in a precise collision course with the moving ball. How does this kind of integration of multiple perceptions in the generation of motor activities happen?

Gibson's account of unified perception rests on two interdependent notions. The first notion is self-awareness, which consists of *proprioception* (usually meant to apply to the states of the skeletal muscles) and what he calls *egoreception*, or an awareness of the internal states of the body. The second notion is that of orientation – the relation between the internal states of the body and the external states of the environment. If perception is to be unitary, the perception of self or internal states cannot be fundamentally different than the perception of the external environment. Gibson therefore dispels the subjective/objective distinction of self to the environment:

The supposedly separate realms of the subjective and objective are actually only poles of attention. The dualism of observer and environment is unnecessary. The information for the perception of "here" is of the same kind as the information for the perception of "there," and a continuous layout of surfaces extends from one to the other . . . the gradients of increasing density of texture, of increasing binocular disparity, and of decreasing motility that specify increasing distance all the way from the observer's nose out to the horizon, are actually variables between two limits, implying just this complementarity of proprioception and exteroception in perception. Self-perception and environment perception go together. (Gibson 1979, p. 116)

Gibson's notion of self-awareness is a combination of the awareness of one's environment with the awareness of oneself as existing in the midst of it, as the here and now. This is how perception brings together our representation of time and space in a unified consciousness:

The puzzle of past, present, and future is not relevant to the problems of event perception. The feeling of *now* is nevertheless often strongly experienced, and we often speak of the *present moment*. Whence comes this compelling experience? I suggest it comes from proprioception, that is, from the perception of the body of the observer himself as distinguished from his environment. It comes particularly from locomotion, and very largely from the visual perception of the locomotion of the observer through the environment. One who makes a journey sees himself moving relative to a stable and rigid world. The flowing perspective in the ambient array of light is ordinarily not noticed. The visual sensations of motion are paid no attention and the underlying invariants that specify the layout of surfaces are what gets perceived. But the traveler also sees his body in the environment and its momentary position. He perceives *here* and, in fact, the very perception of the environment entails the perceiving of *here*. The traveler perceives the path to be traveled if he looks ahead, the path that has been traveled if he looks behind, and the position in between is called *here*. The traveler is tempted to think of the linear path as the dimension of time and to see the path traveled as the past, that to be traveled as the future, and the division point as the present. The point *here* and the moment *now* coincide. (Gibson 1975, p. 300)

Gibson clearly thinks that representing something as being 'over there' implies that you are representing yourself as being 'here', otherwise you would be stuck on the

question of ‘over there *from where?*’ without any reply and the question of what grounds these representations, of internal and external states, are the motor activities which are possible – the repositionings of the body and the distances in time and space imagined as potentials to muscular actions, movements in and through the world.

Work in robotics in the past decade has begun to exploit some aspects of Gibson’s theory. This work has been advanced by the development of algorithms that are able to extract motion information from video sequences to construct vector flow fields, and coupling these to motor activities. Srinivasan and Venkatesh (1997) have done elaborate and careful studies of how honeybees use visual flow in navigation. They have shown that bees flying through narrow tunnels control their distance from each side of the tunnel using a neural circuit which compares the speed of the optic flows in each eye and automatically adjusts the wing beats to keep the bee in the middle of the tunnel. They have also shown that bees use optic flow to measure distance to pollen targets, a strategy which is invariant to wind speed. Using these models, Srinivasan and others have built robotic bees which navigate tunnels by coordinating the visual flow rates between two cameras. Lewis (1998) has used a similar embodied technique to create a one-eyed robot which can navigate a 3D world by recognising differential rates of flow and treating these as indications of occlusion, and hence obstacles.

So the ecological theory of perception is consistent with Ashby’s conception of the mind as embodied representation but one might still ask, what can be gained by pointing out this shared perspective? To that, we could answer that there is much that might be gained by employing Ashby’s formal techniques to the framework and problems presented in the Ecological Theory of Perception. Primarily, what I have in mind is to begin thinking about how the brain reduces the vast amounts of sensory information into quantities of information which are both manageable and useful. If we wish to extend the application of the ecological theory to robotics in more sophisticated ways than just visual flow fields, it will be necessary to find techniques by which ambient sensory information can be partitioned into meaningful parts. This is the specific problem which Ashby addresses in his 1968 paper, ‘Information Processing in Everyday Human Activity’. I want to briefly sketch out the problem and Ashby’s solution before moving on to consciousness.

Information, and quantities of information, is too often treated in a casual way which fails to respect the deep insights provided by Shannon’s (1948) formal theory. By this I mean that we need to understand that as a regulator improves, it actually deals with less information, rather than more. A naïve, untrained regulator will not yet have learned which aspects of the environment are relevant to it. Thus, it must pay attention to as many aspects as it can. One significant part of learning is to discover which set of messages are the relevant ones that the regulator needs to select among. As Ashby points out, it does not even make sense to talk about information in the absence of the set of possible messages and while there are currently many techniques for machine learning algorithms to learn the structure of the messages within a set, or to associate them to sets of actions, there seem to be no methods for learning the boundaries of a set of messages. The set must always be given by the system designer. This is one of the major differences between natural and artificial learning systems, and one of the major impediments to achieving a semblance of life-long learning, and general purpose learning in machines.

In this rather short paper from 1968, Ashby considers a typical action in everyday life: a man is reading when he encounters an unfamiliar French word, gets up, walks across the room avoiding a chair in his path, finds his French dictionary among 100 other books, finds the word, reads the English translation, and writes down the corresponding English word.

Ashby asks us to consider: How much information is required to complete this task? Ashby first notes that without defining the sample space, information cannot be quantified:

Now “information” as understood today, has meaning only when defined over some sample space (Shannon), or over a set of frequencies (McGill): the *multiplicity* of possibilities is essential. If therefore we think of this Action being performed by a particular person in a particular room on a particular day, then *this* event is unique in the universe, has no multiplicity, and makes any question about its informational properties merely improper. To bring this event into some relation to a measure of information, we must extend it to a *set* of Actions. It is this extension, in our opinion, which is the critical and essential step in the development of a logically defensible method. (Ashby 1968b, pp. 190–191)

Given that a brain is unlikely to have access to the statistical structure of similar events experienced by other brains – the objective probabilities – a reasonable way to measure information is from the perspective of the brain, and thus it ought to be measured against the set of possible actions the brain can take. In fact, he makes the much stronger argument that information can *only* be defined in terms of possible actions and the method of measuring this information ought to apply to robots as well as brains:

Once the sample space or set (over which the transmission is to be computed) has been defined, the computation proceeds in just the same way, and must arrive at the same number, whether the subject is an intelligent *Homo sapiens* or is a Robot designed to perform just that set of actions and nothing more. The approach through this axiom may reduce greatly one’s initial intuitive estimate of what is necessary. In particular, it removes from our consideration all the activities within the nervous system, for these activities are neither described nor varied in the defined set of actions. (If the reader prefers to introduce neuronically variations into those listed above, his numerical answer would be different from ours; the method, however, would be the same. Essentially, he would be asking a different question.) (Ashby 1968b, p. 191)

It should now begin to become clear why I think Ashby’s later work meshes so nicely with work on ecological perception a decade later, and work on autonomous robotics two decades later. Ashby recognised that we should not think of robots in different terms than we do human brains. When one considers each from the perspective of performing a task involving controlled action, the processing of information will be the same in each. This is the very essence of information processing in the brain, though he also acknowledges that things become much more complicated when the behavioural repertoire increases:

The question asks, in effect: If a robot is built to carry out the defined Action successfully, with the coordinations and corrective actions necessary, how much transmission *must* be provided? The answer cannot be far from our estimate, for either the machine will not be able to give a tolerable imitation of this Action (by being excessively clumsy), or it will demonstrably be using transmission wastefully. Yet even if it (or the human counterpart) were only 1% efficient, and used 300 bits per second, one would still want to know (say) why man’s optic nerve, with about 500,000 fibers, offers at least that latter number of bits per second. We may well ask: Why do man’s sense organs accept all the extra information? A possible answer is suggested as soon as we realize that the two systems we are comparing are a Robot (or a man) performing the defined Action *and nothing more*, and the man of real life, who can perform not merely this Action (call it A_1) but who can also perform a great number of other Actions $A_2, A_3, A_4 \dots$ Even while engaged in A_1 , the normal man is able to respond to the intrusion of other variations – the ringing of the telephone, the discovery that the Dictionary is missing, the collapse of the bookshelves, and a host of variants not given in our list of “everyday variants” above. These choices *between* A_1, A_2, A_3 , etc., will require a “higher level” activity with information-processing extra to that used *within* any particular A . Our estimate suggests that this “higher level” activity, not detectable while the Action is in progress, is, in fact, requiring much larger quantities of transmission than that used in the more obvious Action itself. One is reminded here of the modern computer, which differs from the older computer largely in the amount of organizational activity it undertakes, activity concerned not with direct

computation but with *which* computation shall occur, and how and where. (Ashby 1968b, pp. 191–192)

Thus, we can now see why the human brain gathers so much information, not to construct a complete 3D model of the world, but to sense all of the affordances to action available in its immediate environment, and even in the not-so-immediate environment. As we seek to increasingly apply the ecological theory of perception to the construction of robots, we would thus be well advised to make use of Ashby's methods of quantifying information in the development of algorithms for processing video images, and coordinating multiple sensory modalities. It is upon this last point that I want to linger. It is the 'higher level' activities, primarily choosing actions, that demand the integration of different senses, of states both internal and external, and of possible motor activities. As I will argue, this turns out to require a capacity of the brain to focus attention, imagine and choose, and this is precisely the function of the mechanism of consciousness.

5. Consciousness

Consciousness is generally considered to be the quintessential mental property, yet raises many unanswered questions. The philosophical question is: How should we define consciousness? The scientific question is: What mechanisms make it possible? And the engineering question is: How do we design a machine to be conscious? For his own part, Ashby had some rather pessimistic things to say about consciousness. Primarily, he was responding to the trend he observed of identifying consciousness with subjective experience, self-awareness in the absence of choice and action, whether something really feels pain, and the qualities of inner mental life. As such, he saw it as something inaccessible to public observation and thus to science. It is with this in mind that he says:

The work of the last twenty years seems to me only to have repeatedly emphasized the profound difference between those aspects of a system that an observer can discover from its outside, by interacting with it (giving it stimuli and receiving stimuli in return from it) and those aspects accessible to the system itself. The difficulty seems to be that science deals only with what is communicable (to other scientists and thus to the body of collective knowledge). A system can thus yield to science only such aspects of itself as are communicable. Some aspects, e.g. its weight, are readily communicable, but what Eddington described as: "my taste of mutton" is not so: he can transmit to another only his reaction to mutton. (Ashby 1968a)

So while the definition of consciousness as something inaccessible to scientific observation was unacceptable, Ashby did not rule out that other definitions might be possible. In particular, it seems he would have been quite pleased with a definition of consciousness given in terms which were observable, measurable, and even more so by one framed in terms of the communication of information, mechanisms of regulation, and embodied representations.

In the past decade, a picture of how consciousness emerges from the mechanisms of the brain has begun to form. The basic idea of these accounts is to identify conscious with the inner communications of the brain which provide it with information about the state of its body and the environment, and thus inform the coordination of action. By far, the most comprehensive account of how this happens has been presented by Rodney Cotterill (1997, 1998). Like Gibson's theory of perception, Cotterill's mechanistic theory of consciousness is compatible with Ashby's theory of embodied representation. I also want to argue that Ashby would have embraced the kind of approach to the scientific study

of consciousness taken by Cotterill. I want to briefly sketch out some key elements of this theory of consciousness, highlight its connections to Ashby's ideas, and then consider how we might proceed in exploring these connections.

Cotterill's work draws on a vast array of otherwise isolated studies of neural circuits, anatomical and behavioural studies, and synthesises them into a unified mechanistic view of consciousness which sidesteps many of the more treacherous philosophical problems associated with this concept. The centrepiece of his theory is what he calls the 'Vital Triangle', a network of three interlocking feedback loops which links together perception, memory, and motor activity. The first aspect of these feedback loops is that they interconnect the sensory cortex at the highest levels, thus bringing together information from all the internal and external senses. This includes proprioceptive information about the state of the body, which is the essence of self-awareness in this view. These sensory elements are also interconnected with the hippocampus and its short term memory, which allows the integration of these senses over time. While it is too complicated to explain in detail here, much of Cotterill's empirical evidence rests on essential aspects of the timing of messages between the areas of the vital triangle which allow for the coordination of complex activities, and for dealing with sudden contingencies – just as the man looking for his French dictionary had to do. Also tied into the vital triangle is the pre-motor and motor cortex. A crucial part of Cotterill's account rests upon a distinction between covert and overt motor activity. That is, he specifies the neural circuits in the brain which express *potential* muscular contraction sequences, imagined by the pre-motor cortex though not actually carried out.

Thus, the vital triangle consists of the motor cortex, sensory cortex, and short term memory. Through the structure of the multiple-feedback loops and their timing, it can be shown that the brain constantly imagines possible motor activities, but the sensory and memory areas can *veto* these plans before the signals are sent to the nerve spindles which coordinate the actual muscular contractions. When we read, and when we think with words, we are vocalising internally, or sub-vocalising, without actually moving our lips, larynx and breath or producing sounds. The same is true, on this view, for all thought – it is motor activity not fully realised. When we imagine running, our motor areas are planning out running movements, but our senses tell us we are sitting, and our memory reminds us we are in a lecture and running is inappropriate behaviour, and so the behaviour is suppressed. Everything which is conscious is a potential action, potential perception, or recalled memory of some kind. Even the most fanciful and impossible acts of the imagination are only conceivable insofar as they are potential actions and perceptions of the body.

To get technical for a moment, the network of neural feedback loops interlaces the cerebral cortex, basal ganglia and the limbic system. The central feedback network involves the coordination of the sensory and motor cortex, mediated by the hippocampus, anterior cingulate nucleus, prefrontal cortex and pre-motor cortex. Using such a mechanistic explanation of consciousness, it could be shown which animals do, and which do not, possess the appropriate neural circuitry and activity to have consciousness. Consciousness is made susceptible to scientific investigation without becoming subjective, yet also without discrediting the role of conscious experience or denying the richness of its quality – it simply does not address the internal perspective apart from the internal flow of information between key structures in the brain.

There is plenty of elaboration to be done in spelling out what this means, but in short it says that we are continually aware of our sensory inputs, as well as aware

of a schema of bodily motor activity which we are currently attempting to perform and moreover, we are constantly projecting the possible future implications of those two structures, evaluating them, and comparing them to our memories. This enables us to recognise when the action we are about to perform is going to have undesirable consequences (hopefully), and thus allows us to abort or veto that action in favour of another action schema. Consciousness is an evaluative or reactive choice tied directly into a constant stream of real-time perception of the world, expressed as embodied motor activities.

What is compelling about Cotterill's thesis is that he provides a simple yet powerful definition (rather close to that presented by William James, actually), where the elements of this simple definition can be identified with concrete mechanisms known to be working in brain structures which are interacting with each other according to the appropriate dynamics, and this is able to account for an incredibly broad range of mental phenomena. It also acknowledges that what the brain does is intimately tied to the environment – mostly it is using sensation to modulate the activity of the body in its constant investigation of the world. Thus, consciousness is not disembodied. Moreover, consciousness isn't just some epiphenomenal accident of evolution. In fact it is causally significant, and a crucial cognitive ability of those animals which have the necessary brain structures to support consciousness. Whereas non-conscious creatures are stuck taking just those actions which their brains reactively select as appropriate for given sensory states and biological goals, the conscious creatures imagine various sorts of consequences of their current actions with respect to their intentions, and make choices about what to do and this happens at time scales ranging from microseconds in the case of the motor control adjustments of a tennis racquet angle or the steering wheel of a car in traffic, to minutes, hours or days in the more reasoned deliberation of such life choices as which car to buy or career to pursue, and that can actually modify the intentional structure of future actions.

All of this is highly time-dependent. Perceptions must be temporally organised along with memories, sequences of muscular contraction must be carefully timed, and actions must be nominated and selected rapidly enough to allow the system to act in time:

It would seem that conscious awareness is possible if temporal changes in sensory input (or in internal patterns of signals) can be detected while the signals resulting from that input are still reverberating in the system. And we could go on to note that such detection will permit the system to capture correlations between cause and effect so rapidly that their significance for the organism can immediately be appreciated, and also immediately used to modify the organism's repertoire of responses. (Cotterill 1998, pp. 333–334)

In fact, Cotterill goes further to argue that consciousness actually arose as a mechanism to coordinate the selection of actions in a timely manner. The basic thrust of the argument is that having a mechanism like consciousness has an evolutionary advantage for two reasons. First, it allows much more rapid forms of learning. A conscious circuit can learn from a single experience, if it is important enough, by focusing its attention on the event, rehearsing it in the imagination, and considering alternative responses and their potential consequences. Second, it allows the system to actively explore its environment, rather than merely react to it. The conscious being can probe for stimuli, and seek out information from the environment. This active form of perception is still poorly understood from a computational perspective, and so I believe it is this area which could most gain from reconsidering Ashby's ideas. It is with a brief discussion of how this might go that I want to conclude.

6. Conclusions

So what is to be gained from a re-examination of Ashby's later work? First and foremost I believe that Ashby gives us a way of looking at information theory and its application to biological cognition which is simultaneously more rigorous and more general than other approaches. That is, we ought to make use of the quantitative and analytic tools that information theory provides but at the same time, we need not limit ourselves to the kinds of symbolic and syntactic representations that have been developed for computer science or communications protocols. In his own words:

Yet the fact remains that information theory is essentially the science of complex dynamic systems, with complex weavings of causes and effects in great numbers. Those who would study such systems need information theory (in some form) just as surveyors need some form of geometry. What has happened, I think, is that too much effort has gone into its development for the telephone and computer that it has developed along lines little suited to the real needs of workers in the biological sciences. Take for instance the fact that almost all information theory developed so far deals with the very tidy case in which the system is going to use an exactly defined set of symbols – the 26 letters of the alphabet, the 10 digits, the distinct voltages between -3 and $+3$, for instance – a very useful case in much engineering. In the biological cases, however, the “alphabet” is not sharply limited, but tails off almost indefinitely. Again, Shannon's basic method is to consider one set (e.g. all possible ten-word phrases) with the messages as the population sampled. But in “content analysis” one has the essentially opposite situation: a unique message has been received and one wants to discuss the various sets from which it might have come. (Ashby 1968a, p. 1497)

The problem to which Ashby points in this passage is one which still challenges us: How are we to understand the flow of information in a rigorous, quantitative way, when the scope and set of possible messages is unknown? This problem is intimately related to the problem of determining which elements of the sensorium are significant, and determining which alternatives are open for action. It is a problem facing explanations of attention, creativity and life-long learning in artificial systems. I believe the solution to these problems might be found through a more careful consideration of consciousness, behavioural repertoires and embodied representations, utilising an information theory retooled for open-systems. In particular, the mechanisms of exploration – active perception and the construction of alternatives – present themselves as key avenues for future research.

Note

1. Gibson was apparently unaware that the original source of Bishop Berkeley's puzzle was William Molyneux.

Notes on contributor



Peter Asaro earned his Doctorate in the History, Philosophy and Sociology of Science, Master of Computer Science, and Master of Arts in Philosophy at the University of Illinois at Urbana-Champaign. He has worked at the National Center for Supercomputer Applications (NCSA), the Beckman Institute for Advanced Science and Technology, and Iguana Robotics, Inc. in the areas of virtual reality, artificial intelligence, machine learning, robot vision, and neuromorphic robotics. His dissertation examined the relationships between brain modelling, the development of early computers, and cybernetic and cognitive theories of mind in the 1940s and 1950s. He has held post-doctoral fellowships at the Austrian Academy of Sciences in Vienna, Umeå University in Sweden, and at the Center for Cultural Analysis

at Rutgers University. He is currently teaching media theory at The New School, and investigating the ethical dimensions of autonomous and tele-operated military robotics, as well as the concepts of agency and responsibility in distributed socio-technical systems. (Homepage: www.cybersophe.org)

References

- Agre, P., 1997. *Computation and human experience*. Cambridge: Cambridge University Press.
- Agre, P. and Chapman, D., 1987. Pengi: an implementation of a theory of activity. *Proceedings of AAAI-87*, 268–272.
- Asaro, P.M., 2006. Working models and the synthetic method: electronic brains as mediators between neurons and behavior. *Science studies*, 19 (1), 12–34.
- Asaro, P.M., 2008. From mechanisms of adaptation to intelligence amplifiers: the philosophy of W. Ross Ashby. In: M. Wheeler, P. Husbands, and O. Holland, eds. *The mechanical mind in history*. Cambridge, MA: MIT Press.
- Ashby, W.R., 1952. *Design for a brain*. London: Chapman & Hall.
- Ashby, W.R., 1956. *Introduction to cybernetics*. London: Chapman & Hall.
- Ashby, W.R., 1967. The place of the brain in the natural world. *Currents in modern biology*, 1 (2), 95–104.
- Ashby, W.R., 1968a. The contribution of information theory to pathological mechanisms in psychiatry. *British journal of psychiatry*, 114, 1485–1498.
- Ashby, W.R., 1968b. Information processing in everyday human activity. *BioScience*, 18 (3), 190–192.
- Ashby, W.R., 1972. Analysis of the system to be modeled. In: R.M. Stogdill, ed. *The process of model-building in the behavioral sciences*. New York: W.W. Norton, 94–114.
- Brooks, R.A., 1991a. How to build complete creatures rather than isolated cognitive simulators. In: K. Van Lehn, ed. *Architectures for intelligence*. London: LEA Publications, 225–239.
- Brooks, R.A., 1991b. Intelligence without representation. *Artificial intelligence*, 47, 139–159.
- Conant, R. and Ashby, W.R., 1970. Every good regulator of a system must be a model of that system. *International journal of systems science*, 1 (2), 89–97.
- Cotterill, R., 1997. On the mechanism of consciousness. *Journal of consciousness studies*, 4 (3), 231–247.
- Cotterill, R., 1998. *Enchanted looms: conscious networks in brains and computers*. Cambridge: Cambridge University Press.
- Gibson, J.J., 1950. *The perception of the visual world*. Boston, MA: Houghton Mifflin.
- Gibson, J.J., 1975. Events are perceivable but time is not. In: J.T. Fraser and N. Lawrence, eds. *The study of time II*. New York: Springer-Verlag, 295–301.
- Gibson, J.J., 1979. *The ecological approach to visual perception*. Boston, MA: Houghton Mifflin.
- Lewis, M.A., 1998. Visual navigation in a robot using zig-zag behavior. *Neural information processing systems*, 10.
- Marr, D., 1982. *Vision: a computational investigation into the human representation and processing of visual information*. New York: W.H. Freeman and Company.
- McCulloch, W.S., 1965. *Embodiments of mind: a collection of papers*. Cambridge, MA: MIT Press.
- Minsky, M., 1963. Steps toward artificial intelligence. In: E.A. Feigenbaum and J. Feldman, eds. *Computers and thought*. New York: McGraw Hill, 405–456.
- Shannon, C.E., 1948. A mathematical theory of communication. *The Bell system technical journal*, 27, (July, October), 379–423, 623–656.
- Srinivasan, M.V. and Venkatesh, S., eds, 1997. *From living eyes to seeing machines*. Oxford: Oxford University Press.
- Varela, F., Thompson, E., and Rosch, E., 1991. *The embodied mind: cognitive science and human experience*. Cambridge, MA: MIT Press.
- Wörgötter, F., et al., 2004. Early cognitive vision: using Gestalt-Laws for task-dependent, active image-processing. *Natural computing*, 3 (3), 293–321.