


This article was downloaded by: [Cariani, Peter A.]

On: 8 April 2009

Access details: Access Details: [subscription number 908384570]

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>

The homeostat as embodiment of adaptive control

Peter A. Cariani^a

^a Department of Otology and Laryngology, Harvard Medical School, Newton, MA, USA

Online Publication Date: 01 February 2009

To cite this Article Cariani, Peter A.(2009)'The homeostat as embodiment of adaptive control',International Journal of General Systems,38:2,139 – 154

To link to this Article: DOI: 10.1080/03081070802633593

URL: <http://dx.doi.org/10.1080/03081070802633593>

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.

The homeostat as embodiment of adaptive control

Peter A. Cariani*

Department of Otolaryngology, Harvard Medical School, Newton, MA, USA

(Received 13 February 2008; final version received 30 September 2008)

W. Ross Ashby was a founder of both cybernetics and general systems theory. His systems theory outlined the operational structure of models and observers, while his cybernetics outlined the functional architecture of adaptive systems. His homeostat demonstrated how an adaptive control system, equipped with a sufficiently complex repertoire of possible alternative structures, could maintain stability in the face of highly varied and challenging environmental perturbations. The device illustrates his ‘law of requisite variety’, i.e. that a controller needs at least as many internal states as those in the system being controlled. The homeostat provided an early example of how an adaptive control system might be ill-defined *vis-à-vis* its designer, nevertheless solve complex problems. Ashby ran into insurmountable difficulties when he attempted to scale up the homeostat, and consequently never achieved the general purpose, brainlike devices that he had initially sought. Nonetheless, the homeostat continues to offer useful insights as to how the large analogue, adaptive networks in biological brains might achieve stability.

Keywords: W. Ross Ashby; cybernetics; general systems theory; homeostat; requisite variety

1. Ashby’s intellectual legacy: cybernetics and systems theory

W. Ross Ashby is rightly regarded as a founder of both cybernetics and general systems theory, those two highly intertwined intellectual movements that flowered in the 1950s and 1960s. From their inception to the present day, both cybernetics and systems theory have been deeply imbued with epistemological concerns – how we know what we know. Cybernetics focused on adaptive organisations and informational mechanisms, the materially embedded mechanisms that make observers and actors possible. As its complement, systems theory concerned the structure of knowledge and ‘observer mechanics’, how an observer–actor can make use of information gleaned through a set of observables.

In his 1948 book Wiener had called cybernetics ‘the science of communication and control’, which subsumed the two strands of adaptive control and information theory (Wiener 1948, Heims 1991, Conway and Siegelman 2005). Ashby’s two books *Design for a Brain* (Ashby 1952b, 1960) and *An Introduction to Cybernetics* (Ashby 1956) also mixed the two, taking under consideration both cybernetic mechanisms and systems methodology. Some sections of these books dealt with regulation and control, and their purpose was to outline a cybernetic theory for designing and constructing adaptive systems. Other sections operationally defined models (‘systems’) and their parts, and their aim was to lay out the groundwork for a general theory of models, which subsequently

*Email: cariani@mac.com

became 'general systems theory' (for its historical development, see Klir 1969, Klir and Elias 2003) and the archives of this journal). The two concerns, mechanism and methodology, were intertwined in Ashby's earlier works; conceiving and building the homeostat coevolved with Ashby's ideas about modelling, which rigorously described what it did. In later writings, these concerns became more separated in his writings as he transitioned from building physical adaptive devices to formalising his analytical systems framework. The shift paralleled similar divergences in the broader movements of cybernetics and systems theory that occurred over the decade of the 1960s, through which separate, but overlapping research communities and professional societies for systems theory and cybernetics emerged. Despite their extensive intellectual similarities and shared membership, one can nevertheless draw a few useful distinctions between the two research programmes.

Cybernetics concerns questions of how material systems, appropriately organised, can support informational mechanisms that make coherent, effective, and adaptive behaviour possible. Cybernetics delineates the functional organisations that permit ensembles of inert material components to form purposive, goal-seeking entities. If mathematics and logic concern formal causation, and physics the theory of material and efficient cause, then cybernetics is the theory of final causes (teleology), how organisations of processes realise functions and purposes. The most fundamental cybernetic strategy is adaptive 'retroaction': to incorporate 'ends' (goals) into 'means' (mechanisms) in a manner that makes their attainment as inevitable as possible (de Latil 1956). Cybernetics as the science of adaptive mechanisms and organisations is a general, highly self-conscious, 'rational' approach to effective design that spans organisms, machines, and societies.

Systems theory, on the other hand, focuses on the functional structure of the observer without regard to the specific material implementations (Buckley 1968, Klir 1969, Weinberg 1975, Klir and Elias 2003, for systems overviews). If cybernetics answers the question, 'how is it possible that organisms and devices can learn what they need to know to act effectively', systems theory asks 'what is the structure of that knowledge such that useful information for effective action can be derived from it?' Systems theory is a theory of observer-mechanics, of the operational structure of scientific models. Ashby's underlying philosophical attitude, as expressed in his writings, is a very clearly operationalist, empiricist, and pragmatist philosophy of scientific models. His philosophy properly lies in the tradition of Hertz (1894), Mach, Bridgman (Bridgman 1931), and Bohr (Murdoch 1987) and, more recently, van Fraassen (1980). Rather than conceptualising state variables in realist terms, as aspects of the material world itself, or in platonist terms, as ideal, abstract, purely mathematical entities, Ashby was always careful to define 'system' operationally, and his operational definitions are cast in terms of observables connected to concrete measurements. Despite the formality of the presentation of his systems theory, his grounding in observation and verifiability sets him apart philosophically from those who conflate the behaviour of the material world with ideal mathematical and physical descriptions of it.

It is only through the lens of Ashby's operationalism, his obsession with definitional clarity, that his critiques of 'self-organisation' (Ashby 1962), 'complexity', and the inability of deterministic automata to self-complexify make any sense. Ashby believed in the existence of material systems that evolve more complex structure over time, but he rigorously defined 'machine' in such a way (correctly we think, Cariani (1989, 1992) that precluded any possibility of finite state automata (digital computers) that could amplify their capacities without the benefit of external inputs. Thus, despite deep differences with Dreyfus' Heideggerian critique (Ashby 1967a), Ashby can also be seen as an early strong critic

of symbolic artificial intelligence (cf. Dupuy 2000, pp. 148–155), who nevertheless believed in the feasibility and even inevitability of constructing artificial minds.

In Ashby's view, observables and system variables are not intrinsic properties of a material realm under study, but windows on that realm that are chosen by the modeller according to what aspects of the observed world s/he seeks to predict:

The would-be model maker is now in the extremely common situation of facing some incompletely defined "system," that he proposes to study through a study of "its variables." Then comes the problem: of the infinity of variables available in this universe, which subset shall he take? What *methods* can he use for selecting them? (Ashby 1972)

Although the selection of variables depends upon the purposes of the modeller – what aspects of the world are to be predicted and/or controlled – there is no general effective method for 'finding the right observables' other than selection from an existing repertoire of available measurements or through the construction of additional measuring devices. How one decides which prospective observables one will test (whether from existing measuring devices or newly constructed ones) are decisions guided by theoretical assumptions. In cases where a material system is already well-characterised and predicted through existing theory, one can use the theory to rapidly find which measurements are needed. In contrast, if the material system is not well-characterised, then existing theoretical knowledge is of less direct use, and one must proceed more through trial-and-error.

In Ashby's pragmatist systems theory, there are no correct or incorrect variables, only those appropriate or inappropriate vis-à-vis some purpose. This generality and pragmatism illustrates the fundamental strengths and weaknesses of general systems theory – on one hand the methodological framework can be applied to any problem in any domain; on the other, being general, the framework is not designed and optimised for any particular set of problems. Modelling success often depends on domain-specific knowledge (which observables are most informative and robust, which goals are practically attainable). Although modelling and simulation are now important computer-based tools in many engineering and policy-oriented disciplines, each discipline uses its own set of techniques particularly adapted to its problem domain. Thus a consequence of system theory's generality is that it cannot easily lay claim to exclusive expertise in any particular concrete problem domain, which in turn makes it difficult to establish academic, professional, institutional, and economic niches for its practice and propagation. Perhaps for this reason, thus far neither cybernetics nor systems theory has managed to achieve more than a few footholds in academia. In this respect the social status of general systems is similar to those of other general, cross-cutting transdisciplinary movements, such as cybernetics, semiotics, general theory of design, hermeneutics, and rhetoric. Many academic departments draw on the ideas and techniques of these movements, but few are explicitly organised around them. For Ashby's work, this has had the consequence that his ideas have not been as widely recognised and propagated there as they might have been had he worked within a more traditional academic discipline.

2. Ashby's homeostat

Ashby conceived and developed his famous homeostat in the early 1940s and published his first papers on it later that decade (Ashby 1948, Pickering 2002, 2008, Asaro 2006, 2008a, 2008b, forthcoming). One of the clearest accounts of its purpose, design, and behaviour can be found in de Latil's (1956) book *Thinking by Machine*. The device embodied and explicated many general principles of cybernetic design and function, such as the Law of Requisite Variety, that Ashby would later formalise and articulate more fully

in his books. The biography in the Ashby Archives (<http://www.rossashby.info/>), based on Ashby's personal notes, lists his explication and proof of the law as occurring during 1952–1953, well after the homeostat's design and construction. It appears likely to us, however, that Ashby had more intuitive notions of variety and selection much earlier than this, and that he designed the homeostat as a concrete illustration of them. At the same time, the experience of building the device and contemplating its workings must have sharpened and refined his theoretical ideas.

The homeostat consisted of four module subsystems each in dynamic equilibrium with the others (Figure 1, photograph and top right schematic). A 25-position 'uniselector' switch determined the analogue control parameters (capacitance and resistance) of each of the four homeostat modules. Given a particular uniselector switch position, the analogue control parameters of the circuits were set arbitrarily (Figure 1). Ashby insisted that these should be random, and adopted parameter values that were based on the Fisher and Yates Table of Random Numbers (Ashby 1960, p. 103). Importantly, the designer or user of the device could be entirely ignorant of the nature of the circuits involved.

Although, as electrical circuits go, each of the modules was a fairly simple affair, the dynamics of the coupled electromagnetic systems, however, could be quite complex. Some of the details of the particular interacting circuits are shown in the bottom panel of Figure 1. These are described, though not as clearly as one might hope, in Chapter 8 of *Design for a Brain*. Each box-shaped homeostat module had a pivoting magnet at its heart that moved a plate. The plate swung within a semicircular chamber filled with distilled water that was situated on the top of the box. Electrodes were situated at the ends of the chamber. Depending on the orientation of the magnet, the plate's position *vis-à-vis* the electrodes would change, and hence its voltage V with respect to them. The four plate voltages constituted the four control variables to be stabilised. In each module, the orientation of the magnet itself depended on the summed effect of the magnetic fields created by four coils (A–D), which were driven by currents from their respective plates in the different modules (amplified via triodes). If the control voltage V of the plate in a given module deviated from its intermediate null state, which was about halfway between the electrodes, a 'uniselector' stepping switch was activated and a new, randomly chosen circuit was put in action in place of the previous one. The uniselector for a given module switched several circuit parameters at once: four resistances that effectively determined the relative strengths of magnetic fields in the four coils, and four capacitances that changed the temporal characteristics of the circuits in which the coils were embedded. Thus at any given time the structure of the homeostat could be characterised in terms of 32 parameters, i.e. 4 modules \times (4 resistances + 4 capacitances) per module.

Given the four uniselector switches, each with 25 distinct positions, the homeostat therefore had $25 \times 25 \times 25 \times 25$ (390,625) clearly defined combinations of uniselector states that determined the analogue control parameters and the associated behaviour of the device. Particular combinations of parameters in interaction with a particular external signal could lead to stability or to chaotic instability. While the 'digital' part of the device, the uniselector switch, had clear, discrete functional states, the analogue part was less exhaustively specified. Although Ashby notes that the parameter space of the homeostat at any given time could be characterised in terms of 32 parameters (16 resistances + 16 capacitances), in fact other factors, such as the friction of the magnet pivot or external magnetic fields or subtle nonlinearities in the operations of the triode, could also come into play to influence the dynamics. Since the homeostat proved to be an ultrastable device, designed to be impervious to almost all conceivable perturbations, including signal inversions, these additional factors likely made no difference in the device's adaptive

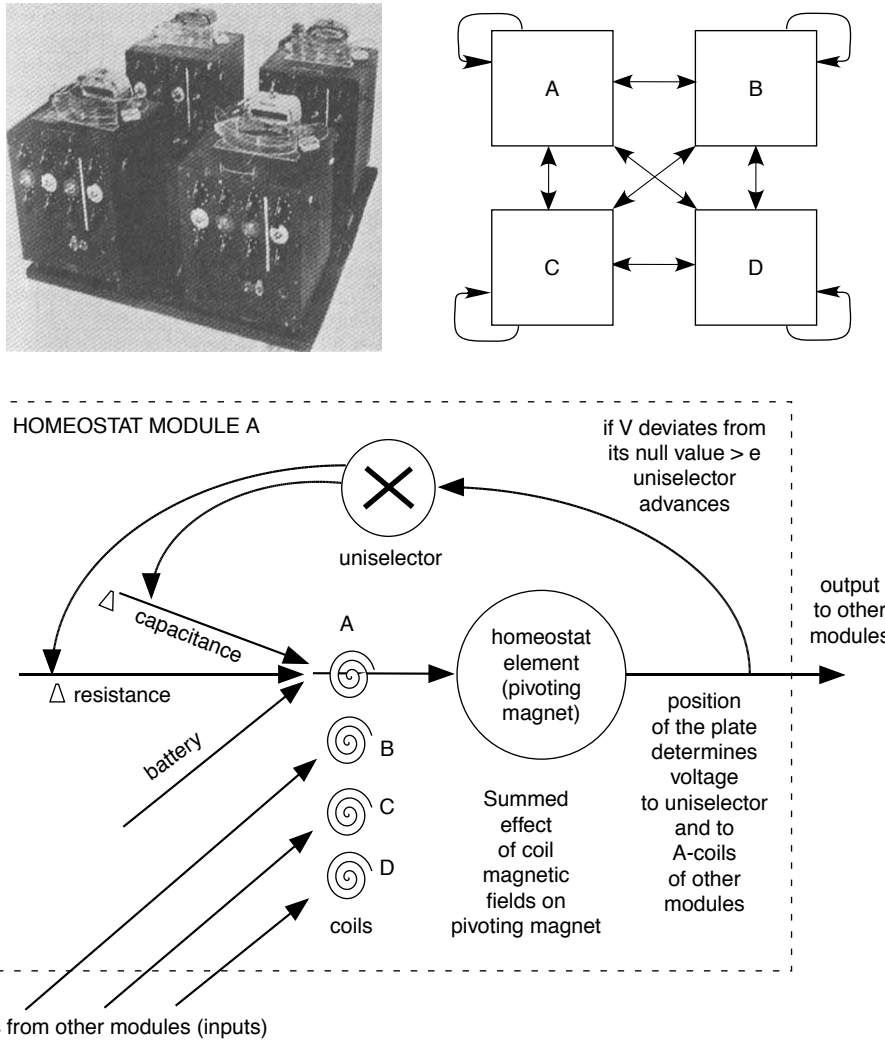


Figure 1. The homeostat. Top left: photograph of the entire device (from Ashby, 1960, used with permission). Top right: interconnections between the four homeostat modules (A, B, C, D). Bottom: functional schema of one homeostat module (after de Latil, 1956). The summed magnetic fields of the four coils A–D influence the position of the pivoting magnet and plate that determines the voltage V to the uniselector, to module A’s own coil (A) and to corresponding A-coils in the other modules (not shown). The four circuits interact such that when they are at equilibrium, voltage V is near zero. When the voltage V exceeds a critical deviation ϵ from its zero null point, the uniselector advances, randomly changing the capacitance and resistance values of module A’s control circuit.

control functioning, since the device always eventually successfully settled on a stable configuration. In terms of specification, in contrast with the digital uniselector states, strictly speaking, the analogue portions of the device were only partially defined.

The purpose of the homeostat was to keep the value of a control variable near a given goal value, within some specified acceptable range. The homeostat thus evaluated whether a particular set of parameters made a ‘good controller’ *vis-à-vis* a particular environment. If the controlled variable did not achieve stability within some specified period of ‘settling’ time, which could be set by the operator, the uniselector switch (in effect, randomly) chose

another set of parameters to be tested. The homeostat would keep testing circuits specified by successive uniselector combinations, its rate of searching through the 390,625 possible circuits limited only by the time needed for mechanical switching and for the electro-mechanical circuit to equilibrate. Apparently, this search process could be agonisingly slow:

It is the slowness of the reactions of the homeostat that makes it difficult to for us to regard the exploration of eventual possibilities as a mechanism for thought. We see the homeostat hesitate, we see it explore a number of solutions one after the other and we find this takes too long. The explanation is simple; the homeostat is an electromechanical device; it is a machine that is retarded by the inertia of its mechanical elements. (de Latil 1956, p. 309)

In order to overcome this speed limitation and expand the scope of the homeostat's goal seeking mechanisms, Ashby subsequently attempted in the 1950s to develop 'multistats', multiple goal-seeking machines that would be based on large numbers of faster, homeostat-like modules (de Latil 1956, pp. 314–316). Some of these efforts are discussed in the last section of this paper.

3. Principles of adaptive design in a mixed digital–analog device

The homeostat embodied several general principles of adaptive design. As with all adaptive systems, the device modifies its internal organisation to better perform some external function. One could couple it to an external environment by means of sensor and effectors, and if the environment was a simple enough dynamical system, the device could eventually find a way to stabilise its interactions. In contrast with the existing special-purpose servomechanisms and adaptive controllers of the day, the homeostat had a much larger, more general repertoire of alternative circuit states and could therefore stabilise more possible situations. Like other adaptive systems, there is a division between parts of the device that continuously interact with the environment (or in the case of the homeostat, with the other modules), and those parts that change the internal organisation of the device itself (in the homeostat, within a given module). Further, the parts responsible for reorganisation require an evaluative part (the uniselector threshold) and a prescriptive part (the implementation of new circuit connections by the movement of the switch). The evaluative part carries within it an implicit goal-state to be satisfied or optimised and has the operational structure of a measurement (Cariani 1989). The prescriptive part, on the other hand, has the operational structure of an internal effector that alters the structure of the device itself.

Ashby's device was a mixture of well-defined digital alternatives (positions of uniselectors), and partially defined analogue dynamics (the circuitry of the controllers). The mixed digital–analogue character of the device partly reflected the electronics that was available in the decade following the Second World War. Digital electronic computers were still in their infancy, while analogue servomechanisms and fire-control systems had reached a high state of technological development. In the late 1940s, there was a keen appreciation of the differences between and respective strengths and weaknesses of analogue and digital systems (von Neumann 1948) that was largely lost amidst the ascendancy of the digital electronic computer a decade later. We believe that Ashby was also driven to this particular mixed analogue–digital design, because he wanted to clearly quantify the number of alternatives available to the device (variety) and the amount of information gained (uncertainty reduction) by the random search process (learning, information pump). Thus, he sought to combine information theory with adaptive control using the dominant, analog electronics technology of the day.

Information theory is easily applied when alternatives are discrete and their transition probabilities are known, but it can become problematic when applied to real world situations in which the universe of alternatives is not fully known or defined. For example, how does one go about quantifying in bits the meaning of a message (Mackay 1969)? This clearly depends on its effect on constraining the functional states of the receiver. For well-defined digital machines, defining the set of alternative functional states is almost trivial, but for animal or human receivers with complex, ill-defined nervous systems, it can be highly problematic.

A key problem Ashby faced is how to quantify variety if one has a continuous set of alternatives. This is the typical case when specifying an analogue control system with continuous-variable parameters. How does one operationalise distinctions between different classes of adaptive devices? One is forced to incorporate discrete functional states if one wants to make rigorous, operational definitions. Ashby was clearly of this spirit, on one hand attempting to explain how creativity is possible (e.g. 'How can a mechanical chess player outplay its designer?') (Ashby 1952a), and on the other striving for clear distinctions and measures.

One answer to the apparent dilemma of creativity vs. rigour is to couple a creative, expansive realm of variety-amplification in which possible solutions are multiplied, to a very hard-nosed evaluative realm in which there is selective contraction of possibilities. In science, the expansive phase engages imagination in formulating questions and hypothetical answers, and enjoys some freedom in choosing which measurements (observables) will be employed to test the hypotheses posed. In this realm, 'anything goes' – random or arbitrary inputs can expand the variety of possibilities. The homeostat was apparently criticised as a 'learning' device because it relied on chance and arbitrary parameters (Dupuy 2000). The contractive phase that follows, however, subjects such hypotheses to severe critical evaluation by comparing concrete predictions with the results of empirical observations.

Deterministic digital devices, when performing to specification, have exhaustive, 'symbolic' descriptions, i.e. the state-transition behaviour of their discrete functional states can be described and reliably predicted in terms of symbols and finite rules (see discussions in Pattee 1985, Cariani 1989). While the states of the uniselector are discrete in number, clearly distinguishable from simple observations, and exhaustively enumerable, this is manifestly not the case for states in the analogue portions of the device. In a real analogue device, one always has the practical problem of distinguishing what are the internal signals and states of the device and what is noise, error, or external perturbation. One can assess with near-certainty the functional state of a discrete switch, but it is *a priori* unclear in an analogue device what levels of signal fluctuations can eventually lead to observable differences in behaviour.

Ashby solved the specification–description problem by introducing the uniselector switch, which selects between different, arbitrary, even 'ill-defined' alternatives. Once a uniselector is embedded in the device, then one can quantify the number of internal structural states of the device, and it is then possible to quantify the effects of selection on this set of clear alternatives. Information gain can then be defined in terms of 'uncertainty reduction,' the process of finding uniselector states that yield successful control systems in a given environmental context. The amount of information gained is the ratio of the number of successful states found to the total number of states searched.

Ashby's Principle of Requisite Variety is therefore well-illustrated by the homeostat, which may even have been the guiding principle of its design. In order to control an external system, a controller needs at least as many internal states as the system being

controlled. In the words of *Principia Cybernetica* (<http://pespmc1.vub.ac.be/>), ‘The larger the variety of actions available to a control system, the larger the variety of perturbations it is able to compensate’. The 390,625 internal circuit configurations of Ashby’s device can compensate for at most 390,625 different environmental contexts. This conclusion seems obvious enough, but it depends critically on how one defines and quantifies the variety of behaviour in a real world situation. Ashby, of course, took great methodological pains to provide the concrete operational definitions and quantifications that could support his ideas (Ashby 1956, Chapter 7).

The homeostat is a device that has no explicit model either of its environs or its internal workings. As de Latil (1956, p. 308) says, ‘The homeostat works through the exploration of possibilities and the sifting of eventualities. The machine itself cannot “know” the best solution of its problems, so it tries either systematically or at random, all possible solutions’. Ashby also realised that not only could the homeostat be ignorant of the details, so could the designer: a designer need not understand at all the workings of the analogue controllers in order to choose those that perform better. This use of constrained random search of ill-defined substrates is a departure from the dominant engineering philosophy of conscious, ‘rational’ design, where designers are guided by some model of the processes they seek to control. The epistemic context of the homeostat is obviously the normal case in biological organisms and brains in homeostasis, learning and evolution – the parts of the system that do the selecting need not (and almost as a rule never do) have any understanding or model of the detailed processes they control. Biological evolution is blind in this sense, genetic mechanisms possess no anticipatory models of themselves or their environs that would guide which mutations would enhance survival and reproduction and which would not. As long as one has a rich source of alternatives (high in variety), and an evaluative process that steers a selective mechanism, one can find solutions to real world problems without understanding how they work or why they succeed. As long as a system is steerable, by selection or feedback, performance can be improved even if the agent steering the system has no model of the underlying processes that are being chosen or modified. With high variety and redundancy, such systems are also highly robust in their operation. William Powers remarked that ‘The homeostat could survive something that no computer program, however adaptive, could survive – an attack with a pair of wire cutters. If one operational connection was destroyed, the uniselector could substitute another one’ (Powers 1973).

Ashby’s homeostat may well have been the first artificial adaptive device to incorporate the principle of an ‘ill-defined’ adaptive system (the coherer of early radio being another candidate). This principle was to be carried to an extreme in Gordon Pask’s electrochemical systems of the late 1950s in which the growth of conductive ferrous thread structures were adaptively steered in order to create new sensors and ‘relevance criteria’ appropriate to their purposes (Pask 1958, 1960, 1961, Cariani 1993, Bird and Di Paolo 2008).

4. The homeostat and biological organisms

There is a basic mapping of the homeostat to the mixed analogue–digital nature of biological organisms. At the centre of every organism are, of course, discrete genetic sequences whose expression determines the parameters that determine the dynamics of cell processes. Although the existence of genetic inheritance mechanisms and discrete units of inheritance had been known to biologists long before the 1940s, the natures of the specific informational substrate of inheritance (DNA nucleotide sequences) and its units (genes) were not yet understood. In that decade, Schrödinger (1945) had postulated an aperiodic crystal storage system for inheritance, and von Neumann (1951) had laid out

a schematic for a self-reproducing kinematic automaton based on informational plans and material parts.

Long after Ashby's homeostat, theoretical biologist Pattee (1982) clarified and extended the analogue–digital distinction and its functional role in living organisation. All material dynamics can in principle be described in terms of classical mechanics, i.e. continuous differential equations that embody the laws of motion for the system. However, those systems that have stable, discrete functional states also have an alternative description in terms of time- or rate-independent states and state-transitions (nonholonomic constraints). This higher level description is couched in terms of discrete states and state transitions, i.e. symbol–strings and transition rules. The genetic code is a 'code' because it can be described in these discrete-state, symbolic terms, rate-independent configurations (DNA nucleotide sequences), that can be mapped to other rate-independent configurations (protein amino acid sequences). Arguably, all memory requires the construction of such special constraints that 'take states out of time' and so provide the basis for their persistence through time. In contrast, the phenotype of an organism includes all those rate-dependent material processes that do not have an alternative, 'symbolic' description. Genes provide constraints or boundary conditions for the chemical dynamics that constitute the phenotype. Typically, gene expression modulates cell processes by switching on production of proteins that function as enzymes that in turn increase the rates of specific reactions.

In the homeostat, one can regard the uniselector states as the device's symbolic genotype that chooses the analogue circuit phenotypes of the homeostat modules (Figure 2). The uniselector's switching of circuit parameters in the homeostat is thus analogous to patterns of gene expression that switch control parameters of cells. The uniselector positions are not determined by rate-dependent processes; they only change step-wise when stability conditions are violated; as with DNA sequences and states of gene expression, patterns of expression change relatively more slowly than, and are effectively buffered from, the much faster analogue dynamics they control. Like biological phenotypes, the analogue circuits that are specified by the uniselector interact with the external world (stippled lines) and with each other in a manner that can be highly unpredictable and nonlinear.

There are some similarities and differences between adaptation or learning in the homeostat and biological evolution (Figure 2). In *Design for a Brain*, Ashby (1960, Chapter 9) had spent considerable effort discussing possible analogues of ultrastability and step-functions in biological organisms and nervous systems. Both homeostat learning and biological evolution involve an evaluative, selection operation followed by variation and testing of new phenotypes. The end result is that the internal structure of the device or organism is altered in a manner that improves its performance in its environment. Both involve random search through a large space of possibilities. The differences between the two stem mainly from differences in variation and self-construction. In the homeostat, the search process is one of iterative search through a sequence of possibilities, whereas in biological evolution the search is highly parallel, simultaneously involving many organisms. The homostat evaluates for stability of one set of variables, whereas natural selection effectively evaluates for complex, general ability to survive and successfully reproduce. The homostat randomly selects from the variety that is built into it from the start; in contrast, biological variety is constantly expanded by mutational and recombinatorial processes that provide for more constrained search. While the homeostat's uniselector switches into a given, pre-established circuit, biological systems undergo a genetically guided developmental process that constructs new phenotypes.

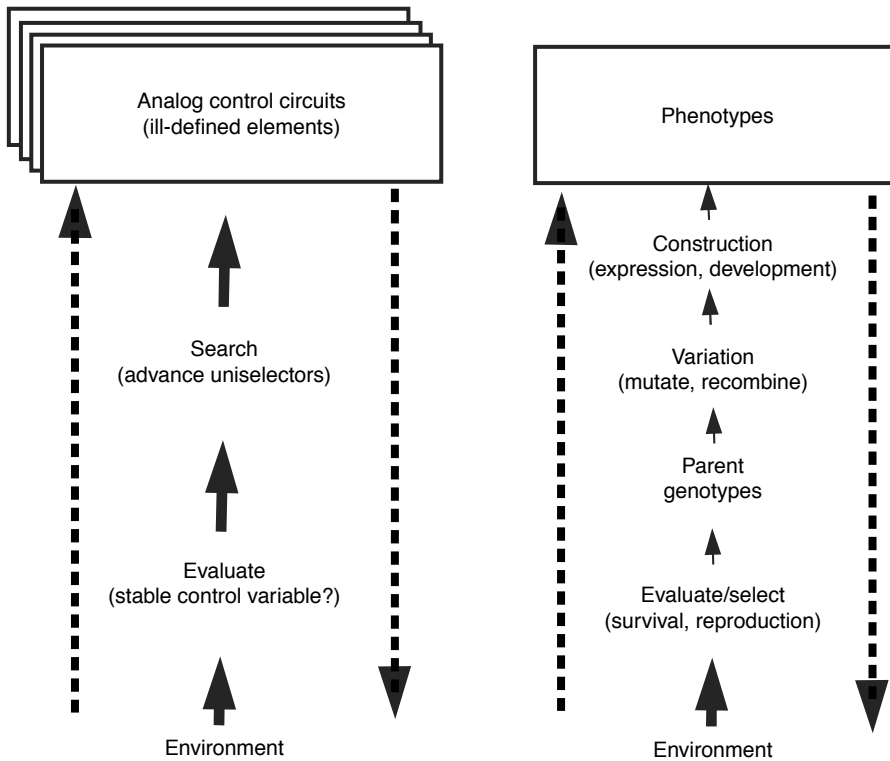


Figure 2. Operational structure of adaptation in the homeostat and evolution of biological organisms. Left: homeostat. Bidirectional interactions with the environment (stippled lines) result in fluctuations of a control variable in each homeostat module. When the variable's value exceeds prescribed bounds over a threshold time duration, the uniselector switch associated with that module is advanced and another analogue control circuit is chosen. The search process stops when stability is achieved. Right: biological evolution. Bidirectional interactions of a population of organisms with their environment result in differential survival and reproduction (natural selection as implicit evaluation *vis-à-vis* survivability and reproductive success). Those parent genotypes thus selected are mutated and recombined to generate ensembles of new offspring genotypes, which guide the physical construction of organisms (phenotypes) in the succeeding generation.

5. The homeostat and the brain

Ashby began his career as a neuroscientist and psychiatrist, studying the anatomy of the cerebral cortex in the search for possible neuroanatomical clues to mental disorders. His M.A. thesis at Cambridge was entitled 'The Thickness of the Cerebral Cortex and its Layers in the Mental Defective' (1935). For most of his formative years (1936–1960), he worked within mental hospitals, but his theoretical ideas appear to have evolved largely independently of psychiatric practice. Ashby's ideas concerning brain function focused on questions concerning the dynamic stability and adaptability of large ensembles of elements rather than psychodynamics, neuronal biophysics, or neural information processing. As a consequence, the points of correspondence between his theories and constructed adaptive devices on one hand and biological brains on the other are almost completely at the level of overt behaviour. One might be tempted to attribute this disconnection from neural mechanisms to the relatively undeveloped state of the neuroscience of the period, but many other neuroscientists of the time, such as McCulloch, Pitts, Rashevsky, Lashley, Gerard,

Licklider, Jeffress, and Hebb, did develop neural net theories that explained particular aspects of brain function. Instead, the homeostat was intended (much like Pask's electrochemical device) as an illustration of principles of stability and adaptation. In his book, *An Approach to Cybernetics*, Gordon Pask defends the relevance of the homeostat for brain theory:

It is easy to cite brain models which are merely imitations; most well-behaved robots, most of the tidy automata that imitate a naughts and crosses [tic-tac-toe] player, nearly all of the maze solving machines (though there are some like Deutsch's Rat which are used explicitly to illustrate an organizational principle rather than to imitate a response). There are not so many cybernetic models to choose from, but one of them, made by Ashby and called the Homeostat admirably illustrates the distinction. It is made up of four interacting regulators and an independent switching mechanism which changes the interconnections between these elements until a stable arrangement is reached. It can (from the viewpoint of psychology and engineering respectively) be dubbed a "brain-like analogue" and "a device for solving differential equations", for it does rather imperfectly, display a brain-like behaviour and it will, rather eccentrically, solve differential equations. Its imperfections as an equation solver (which it is not meant to be) are obvious from its construction and have met with a good deal of heavy-handed criticism. Its imperfections as a brain-like analogue (which, once again, it is not meant to be) occur because at the level of functional analogy the organization of the homeostat is not particularly brain-like. It is only when we come to the level intended in the cybernetic abstraction that the self-regulation in a homeostat is identical with the self-regulation in the brain, and with reference to this feature the homeostat is a cybernetic model of all brains. (Pask 1961, p. 17).

In the early 1950s, Ashby attempted to speed up and scale up the homeostat by building larger networks of interacting, purely electronic, inertia-less elements. Although the simple homeostat is not immediately recognisable as being particularly brain-like in its internal structure, one can conceive of brains as networks of mutually interacting homeostatic neuronal assemblies that settle into complex global resonance states following some external perturbation. Ashby called his electronic implementation, the dispersive-and-multistable system (DAMS). The effort seems to have failed to produce usable results, and little mention is given about this phase of Ashby's research. Apparently references to DAMS in the first 1952 edition of *Design for a Brain* were omitted from the second edition eight years later (Ashby 1960); see Asaro (2008b) and contributions by Asaro and Pickering in this issue. Writing in the period between the first and second editions, de Latil reported,

Such a system is designed to have a hundred units instead of the four of the original homeostat it is a question of how many different "organizations" will enter into each element . . . The new universe that would be created by such machines consisting of thousands of units is hardly conceivable. Ashby already considers that the present DAMS machine is too simple and is planning another with even more complex action. Unfortunately, its construction would be an extremely costly undertaking and is not to be envisaged for the present. It may be that some co-operative alliance between binary calculating machines and machines of the DAMS type is possible. It would seem to be a natural sequence to integrate fifth-degree [goal-seeking] mechanisms with the complexes of machines with receptive units and servo-mechanisms with retroactive regulation. (de Latil 1956, p. 310)

If the problem of how to construct ultrastable neural networks from interacting homeostat-like elements could be solved by DAMS machines, then one could embed these networks in robotic devices with sensor and effector arrays that would provide rich coupling with the external world and goal mechanisms that could direct behaviour. A decade after Ashby, Soviet engineers also attempted unsuccessfully to build large-scale general-purpose analogue computers (Carello *et al.* 1984).

Aside from similarity at high levels of cybernetic abstraction, one can ask whether there is merit in considering brains as some species of homeostat or large scale network

of homeostatic elements. Habituation, accommodative responses (Morrell 1967), and enhanced responses to unexpected or novel stimuli (e.g. mismatched negativities) that are commonly seen at the cortical level are suggestive of analogue compensatory mechanisms. Brains have been modelled in terms of nested networks of feedback control processes (Reichardt 1961, Arbib 1972, Powers 1973, Sommerhof 1974). In terms of its network properties, the brain is now seen as an interconnected, 'small world' network of recurrent pathways with dense local connections and sparse long-range ones (Basset and Bullmore 2006). Perhaps Ashby, from his experience with DAMS, could have told us why such small networks are essential for global stability, and why more completely connected networks would tend towards instability. Thus the basic premise of the homeostat as a goal-directed, self-regulating complex feedback control network of relatively few interconnected, resonating modules seems consistent with some essential aspects of the functional organisation of the brain.

The discrete switching of analogue dynamics is harder to envision within nervous systems. What brain structures and processes might loosely correspond to the uniselector step-function switch of the homeostat? Is there any structure that implements a discrete search for appropriate sets of neuronal assemblies?

Some possible high-level candidates would be neural mechanisms that switch behavioural modes. There are brain circuits and processes that do perform switching functions for qualitative behavioural states that faintly resemble Ashby's uniselector switch. Kilmer and McCulloch (1969) proposed a model for the 'reticular formation command and control system' in which the brainstem nucleus switches the rest of the nervous system into different operating modes appropriate for a given context. They listed 17 'mutually incompatible modes of vertebrate behaviour', including sleep, eat, drink, fight, flee, hunt, search, urinate, defecate, groom, engage in sex, lay eggs/give birth, suckle/hatch, build nests, and possibly also migrate, hibernate, and engage in other instinctual behaviour. Like Ashby's uniselector switch, 'The reticular core cannot invent new modes and it cannot perform what it commands. Its logic is abductive, the Apagoge of Aristotle. Given the rules, it must when confronted with the fact, guess which rule applies'. (Kilmer and McCulloch 1969, p. 1331) Also like the uniselector, if a goal state is not reached, the system eventually switches into another behavioural mode (e.g. if 'flight' does not succeed in evading a predator, then the system switches into 'fight' or 'play dead' mode). Unlike the homeostat, however, the mode transition is not random, and there is no one, global goal state. Instead, each behavioural mode has its own goals, and behaviour may be switched to another mode once the goals of the present mode are achieved.

Other possibilities that Ashby considered for neuronal step mechanisms are pulsatile action potentials themselves, switchings between alternative sets of self-exciting reverberating circuits, and adaptive formation of neuronal interconnections (Ashby 1960, pp.124–126). Were he alive today, he might also consider 'neural Darwinist' theories of learning (Edelman 1987) in which intrinsic variation in parallel neural circuits is selectively stabilised.

An alternative path to ultrastability that does not rely on step functions for internal reorganisation is that of adaptive self-modification and self-construction. The homeostat can be seen as a complex dynamic, resonant system that has hard constraints (uniselector position, hard-wired circuits) that control its boundary conditions. In living organisms, self-producing, autopoietic systems achieve organisational closure and stability through active regeneration of parts and relations (Maturana and Varela 1973, Zeleny 1981, Mingers 1995, Geoghehan and Pangaro 2008). Here resetting of boundary conditions, the role of the uniselector, is mediated by hard genetic constraints. It is possible to think of brains in analogous, albeit

informational, terms, as organisations that actively regenerate their own signals, i.e. in terms of an ‘autopoiesis of neural signals’ (Cariani 2000, 2001). Here the resonances are between complex patterns of spikes, and stability is achieved by the regenerative amplification of mutually compatible sets of neuronal signals. Stable global brain states are thus similar to the global stabilities achieved by the interplay of the homeostat control circuits, except that the numbers of elements and circuits of course are many orders of magnitude larger, and the signals are complex, multidimensional spike volley patterns rather than scalar voltages and magnetic field intensities. In the brain, long term memory sets the boundary conditions that can regenerate the characteristic neuronal signal dynamics of an individual, even after drastic alterations of state that accompany sleep, seizure, general anesthesia, and coma. Thus, from this perspective, in the brain, long term memory and its modification through learning plays the role of the unisector in selecting and changing the constraints that regulate moment-to-moment dynamics.

Even if homeostat-like, mixed analogue–digital systems do successfully solve general problems of stability and adaptive control, one is still left with unknown processes of form perception and integration that subserve complex information processing (de Latil 1956, pp. 334–336). It is difficult to see how homeostat-like control networks can implement perceptual and cognitive equivalence categories and operations. In addition to adaptive control, one needs generalised mechanisms that can effect transformations (Pitts and McCulloch 1947) and extract invariances (Cariani 2004, Hawkins and Blakeslee 2004). This apparent dichotomy notwithstanding, there are possible unifying strategies for merging adaptive control and pattern-recognition operations in Gibsonian concepts of affordances, direct perception, analogue ‘smart machines’, and coordinative transformations (Carello *et al.* 1984) and Somerhoff’s idea of ‘directive correlations’ (Ashby 1967b, Sommerhoff 1968, 1974).

Had Ashby been successful in developing DAMS, then general self-organising analogue computers might have been developed alongside the emerging digital programming approaches of symbolic artificial intelligence, and a better balance between adaptive robotics and pure computation might have been achieved in the decades that followed. Similarly, theoretical neuroscience might have developed a better balance between brain models based on large-scale distributed feedback control and those based on sequential hierarchies of feature detectors.

Notes on contributor



Peter A. Cariani (B.S. 1978, biology MIT, 1978, M.S. 1983, Ph.D. 1989, systems science Binghamton University; www.cariani.com) has had the good fortune to be engaged in a wide variety of empirical and theoretical scientific and philosophical investigations: aging in nematodes, computer modelling of protein folding, theoretical biology/biological cybernetics, biosemiotics, neural coding in the auditory system, auditory scene analysis, neural networks, and spinal cord regeneration. His doctoral work under the supervision of Howard Pattee at Binghamton involved the epistemology of emergent measurement functions in adaptive self-modifying systems. He subsequently investigated temporal coding of pitch and consonance in the

auditory system with Bertrand Delgutte and Mark Tramo. Most recently, he has proposed new kinds of neural architectures for information processing, which he calls ‘neural timing nets’. Dr Cariani currently serves as an external scientific consultant for the John Templeton Foundation on funding initiatives related to emergence in biological systems and the material/neuronal basis of

consciousness. He is a Clinical Instructor in the Otolaryngology Department of Harvard Medical School and teaches a graduate course in music perception and cognition (HST.725) in the Speech, Hearing, Biosciences and Technology Program (SHBT) in the Harvard-MIT Division of Health Sciences and Technology (HST) at MIT (<http://ocw.mit.edu/OcwWeb>).

References

- Arbib, M.A., 1972. *The metaphorical brain: an introduction to cybernetics as artificial intelligence and brain theory*. New York: Wiley-Interscience.
- Asaro, P., 2006. Working models and the synthetic method: electronic brains as mediators between neurons and behavior, Special Issue on Models and Simulations. In: T. Knuuttila, E. Mattila, and M. Merz, eds. *Science studies*, 19 (1), 12–34.
- Asaro, P.M., 2008a. Information and regulation in robots, perception and consciousness: Ashby's embodied minds. *International journal of general systems*, 38, 111–128.
- Asaro, P.M., 2008b. From mechanisms of adaptation to intelligence amplifiers: the philosophy of W. Ross Ashby. In: P. Husbands, O. Holland, and M. Wheeler, eds. *The mechanical mind in history*. Cambridge, MA: MIT Press, 149–184.
- Asaro, P.M., forthcoming. Computers as models of the mind: on simulations, brains and the design of early computers. In: S. Franchi and F. Bianchini, eds. *Towards an archaeology of artificial intelligence*. New York: Springer.
- Ashby, W.R., 1948. The homeostat. *Electronic engineering*, 20, 380.
- Ashby, W.R., 1952a. Can a mechanical chess-player outplay its designer? *British journal for the philosophy of science*, 3, 44–57, reprinted in Conant, R., ed., 1981.
- Ashby, W.R., 1952b. *Design for a brain*. 1st ed. London: Chapman & Hall.
- Ashby, W.R., 1956. *An introduction to cybernetics*. London: Chapman & Hall.
- Ashby, W.R., 1960. *Design for a brain*. 2nd ed. London: Chapman & Hall. Unless otherwise noted, all general references to this book are to this edition.
- Ashby, W.R., 1962. Principles of the self-organizing system. In: H. von Foerster and G.W. Zopf, eds. *Principles of self-organization*. New York: Pergamon Press, 255–278, reprinted in Buckley, W., ed., 1968.
- Ashby, W.R., 1967a. The brain of yesterday and today. *1967 IEEE convention record*, 15 (9), 30–33, reprinted in Conant, ed., 1981.
- Ashby, W.R., 1967b. The set theory of mechanism and homeostasis. In: D.J. Stewart, ed. *Automaton theory and learning systems*. London: Academic Press, 27–51, reprinted in Conant, ed., 1981.
- 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, 78–97, reprinted in Conant, R., ed., 1981.
- Bassett, D.S. and Bullmore, E., 2006. Small-world brain networks. *Neuroscientist*, 12 (6), 512–523.
- Bird, J. and Di Paolo, E., 2008. Gordon Pask and his maverick machines. In: P. Husbands, O. Holland, and M. Wheeler, eds. *The mechanical mind in history*. Cambridge, MA: MIT Press, 185–211.
- Bridgman, P.W., 1931. *Dimensional analysis*. New Haven: Yale University Press.
- Buckley, W., ed., 1968. *Modern systems research for the behavioral scientist: a sourcebook*. Chicago: Aldine.
- Carello, C., Turvey, M.T., Kugler, P.N. and Shaw, R.E., 1984. Inadequacies of the computational metaphor. In: M. Gazzaniga, ed. *Handbook of cognitive neuroscience*. New York: Plenum Press, 229–248.
- Cariani, P., 1989. *On the design of devices with emergent semantic functions*. Thesis (PhD). State University of New York, Binghamton.
- Cariani, P., 1992. Emergence and artificial life. In: C.G. Langton, C. Taylor, J.D. Farmer, and S. Rasmussen, eds. *Artificial life II. Volume X, Santa Fe institute studies in the science of complexity*. Redwood City, CA: Addison-Wesley, 775–798.
- Cariani, P., 1993. To evolve an ear: epistemological implications of Gordon Pask's electrochemical devices. *Systems research*, 10 (3), 19–33.
- Cariani, P., 2000. Regenerative process in life and mind. *Annals of the New York Academy of Sciences*, 91, 26–34.
- Cariani, P., 2001. Symbols and dynamics in the brain. *Biosystems*, 60 (1–3), 59–83.

- Cariani, P., 2004. Temporal codes and computations for sensory representation and scene analysis. Special Issue on Temporal coding for neural information processing. *IEEE transactions on neural networks*, 15 (5), 1100–1111.
- Conant, R., ed., 1981. *Mechanisms of intelligence: Ross Ashby's writings on cybernetics*. Seaside, CA: Intersystems Publications.
- Conway, F. and Siegelman, J., 2005. *Dark hero of the information age: in search of Norbert Wiener, the father of cybernetics*. New York: Basic Books.
- de Latil, P., 1956. *Thinking by machine*. Boston, MA: Houghton Mifflin.
- Dupuy, J.P., 2000. *The mechanization of the mind: on the origins of cognitive science*. Princeton, NJ: Princeton University Press.
- Edelman, G.M., 1987. *Neural Darwinism: the theory of neuronal group selection*. New York: Basic Books.
- Geoghehan, M.C. and Pangaro, P., 2008. Design for a self-regenerating organization. *International journal of general systems*, 38, 155–173.
- Hawkins, J. and Blakeslee, S., 2004. *On intelligence*. 1st ed. New York: Times Books.
- Heims, S.J., 1991. *The cybernetics group*. Cambridge, MA: MIT Press.
- Hertz, H., 1894. *Principles of mechanics*. New York: Dover, 1956 reprint edition.
- Kilmer, W. and McCulloch, W.S., 1969. The reticular formation command and control system. In: K.N. Leibovic, ed. *Information processing in the nervous system*. New York: Springer-Verlag, 297–307.
- Klir, G., 1969. *An approach to general systems theory*. New York: Van Nostrand.
- Klir, G.J. and Elias, D., 2003. *Architecture of systems problem solving*. 2nd ed. New York: Kluwer Academic/Plenum Publishers.
- Mackay, D.M., 1969. *Information, mechanism and meaning*. Cambridge, MA: MIT Press.
- Maturana, H. and Varela, F., 1973. Autopoiesis: the organization of the living. In: H. Maturana and F. Varela, eds. *Autopoiesis and cognition*. Dordrecht, Holland: D. Reidel.
- Mingers, J., 1995. *Self-producing systems*. New York: Plenum Press.
- Morrell, F., 1967. Electrical signs of sensory coding. In: G.C. Quarton, T. Melnechuck, and F.O. Schmitt, eds. *The neurosciences: a study program*. New York: Rockefeller University Press, 452–469.
- Murdoch, D., 1987. *Niels Bohr's philosophy of physics*. Cambridge, MA: Cambridge University Press.
- Pask, G., 1958. Physical analogues to the growth of a concept. *Mechanization of thought processes*. Symposium 10, National Physical Laboratory, 24–27 November. London: HMSO, 765–794.
- Pask, G., 1960. The natural history of networks. In: M.C. Yovits and S. Cameron, eds. *Self-organizing systems*. New York: Pergamon Press, 232–263.
- Pask, G., 1961. *An approach to cybernetics*. New York: Harper & Brothers.
- Pattee, H.H., 1982. Cell psychology: an evolutionary view of the symbol-matter problem. *Cognition and brain theory*, 5, 325–341.
- Pattee, H.H., 1985. Universal principles of measurement and language functions in evolving systems. In: J.L. Casti and A. Karlqvist, eds. *Complexity, language, and life: mathematical approaches*. Berlin: Springer-Verlag, 268–281.
- Pickering, A., 2002. Cybernetics and the mangle: Ashby, Beer and Pask. *Social studies of science*, 32 (3), 413–437.
- Pickering, A., 2008. Psychiatry, synthetic brains and cybernetics in the work of W. Ross Ashby. *International journal of general systems*, 38, 213–230.
- Pitts, W. and McCulloch, W.S., 1947. How we know universals: the perception of auditory and visual forms. *The bulletin of mathematical biophysics*, 9, 127–147.
- Powers, W., 1973. *Behavior: the control of perception*. New York: Aldine.
- Reichardt, W., 1961. Autocorrelation, a principle for the evaluation of sensory information by the central nervous system. In: W.A. Rosenblith, ed. *Sensory communication*. New York: MIT Press/John Wiley, 303–317.
- Schrödinger, E., 1945. *What is life? The physical aspect of the living cell*. Cambridge, MA: Cambridge University Press.
- Sommerhoff, G., 1974. *Logic of the living brain*. London: John Wiley.
- Sommerhoff, G., 1968. Purpose, adaptation, and “directive correlation” (excerpt from *Analytical Biology*, 1950, Oxford University Press). In: W. Buckley, ed. *Modern systems research for the behavioral scientist: a sourcebook*. Chicago: Aldine, 281–295.

- van Fraassen, B.C., 1980. *The scientific image*. Oxford: Oxford University Press.
- von Neumann, J., 1948. Re-evaluation of the problems of complicated automata—problems of hierarchy and evolution. In: W. Aspray and A. Burks, eds. *Papers of John von Neumann on computing and computer theory*. Cambridge, MA: MIT Press, 477–490.
- von Neumann, J., 1951. The general and logical theory of automata. In: L.A. Jeffress, ed. *Cerebral mechanisms of behavior (The Hixon Symposium)*. New York: Wiley, 1–41 (Also reprinted in Buckley, W., 1968).
- Weinberg, G., 1975. *An introduction to general systems thinking*. New York: Wiley-Interscience.
- Wiener, N., 1948. *Cybernetics*. New York: Wiley.
- Zeleny, M., ed., 1981. *Autopoiesis, a theory of living organizations*. New York: North Holland.