To evolve an ear: epistemological implications of Gordon Pask's electrochemical devices

Peter Cariani

Systems Research 1993; 10 (3):19-33

"With this ability to make or select proper filters on its inputs, such a device explains the central problem of epistemology. The riddles of stimulus equivalence or of local circuit action in the brain remain only as parochial problems." --Warren McCulloch, preface, [29]

Abstract

In the late 1950's Gordon Pask constructed several electrochemical devices having emergent sensory capabilities. These control systems possessed the ability to adaptively construct their own sensors, thereby choosing the relationship between their internal states and the world at large. Devices were built that evolved *de novo* sensitivity to sound or magnetic fields. Pask's devices have far-reaching implications for artificial intelligence, self-constructing devices, theories of observers and epistemically-autonomous agents, theories of functional emergence, machine creativity, and the limits of contemporary machine learning paradigms.

Keywords: epistemology, cybernetics, self-organizing systems, emergence, sensory evolution, machine creativity, evolutionary robotics, artificial intelligence, artificial life, autonomous agents

Creativity and structural autonomy

All attempts at artificial intelligence inevitably confront what W. Ross Ashby called Descartes' Dictum: how can a designer build a device that outperforms the designer him/herself [2]? Ultimately it is a problem of specification: if the designer specifies all of the parts of the device and what it will do under all circumstances, it will not do any better than its maker. We know, however, that such devices can be made, and indeed they have been made. We now have chess-playing devices which can outplay all but the very best human grandmasters, playing far better than their creators. What allows them to outperform their designers? As Ashby noted, in order to achieve better performance over its initial specification, a device must be informationally open, capable of interacting with the world independently of its designer. The device must have some degree of epistemic autonomy in order to improve itself, but epistemic autonomy is not achievable without some degree of structural autonomy.

By now the latest wave of connectionist, neural net devices has made us all aware of the multitude of possibilities inherent in trainable machines. Such machines improve on their (initial) designs by altering their decision functions contingent upon evaluation of past performance. But even with these machines, the designer must foresee the basic categories of percepts (i.e. primitive features) and actions which will be adequate to solve the problem at hand. Once these are given, the device attempts to optimize its performance by finding better mappings between the perceptual states it has been given and its available action alternatives.

For completely symbolic realms, such as chess or problems of mathematical logic, a set of basic categories is given by the formal description of the problem. For these problems, finding appropriate mappings within predefined alternatives is all that can be done. For real world tasks, however, there is no such set of basic categories that is defined beforehand, so that in addition to finding appropriate mappings there is also the problem of deciding what the basic categories will be.

Essentially, contemporary trainable machines have the freedom to adapt within a set of percept and action categories, but they do not have the freedom to modify those categories. Aspects of the device that are not left plastic and subject to adaptive modification must be pre-specified. Hence the designer is left with the ill-defined task of coming up with appropriate sensors and effectors (or in other terms, "relevance criteria", "observables", "controls", "primitive features") for a given task. As Ross Ashby put it:

> "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 the correct subset?" [6]

Could one go further? Could one construct devices that have the capacity to adaptively construct their own perceptual categories and their own means of influencing the world? Such devices would find their own "relevance criteria", by adaptively constructing sensors to gather the information that they needed to solve a given real world problem. Out of "an infinity of variables" such a device would come up with a set of variables adequate for a specific task. Such a device would be the analogue of both the scientist searching for the right observables for his/her model and the biological evolution of a new sensory modality.

It seems that there is but one person thus far who has recognized this fundamental question and has taken on the task of automating a process for finding its answer. In the 1950's Gordon Pask conceived and built a series of electrochemical devices deliberately designed to find their own "relevance criteria".

Organic analogues to the growth of a concept

Through the early and mid 1950's Pask experimented with electrochemical assemblages, passing current through various aqueous solutions of metallic salts (e.g. ferrous sulphate) in order to construct an analog control system. The system would be different from others in existence in that its design would not be completely well-defined: no explicit specification would be given for its parts. Pask was specifically looking for a machine that would create its own "relevance criteria", one which would find the observables that it needed to perform a given task. The device would go one step beyond Pask's earlier Musicolor system (see papers by Pangaro and McKinnon-Woods, this volume), by evolving sensors to choose, independent of the designer, those aspects of its external environment to which it would react. Not only would particular input-ouput combinations be chosen but the *categories* of input and of output would be selected by the device itself.

To carry out this research program, Pask needed a medium rich in structural possibilities, one which could be adaptively steered. What kind of medium could support this kind of self-organization?

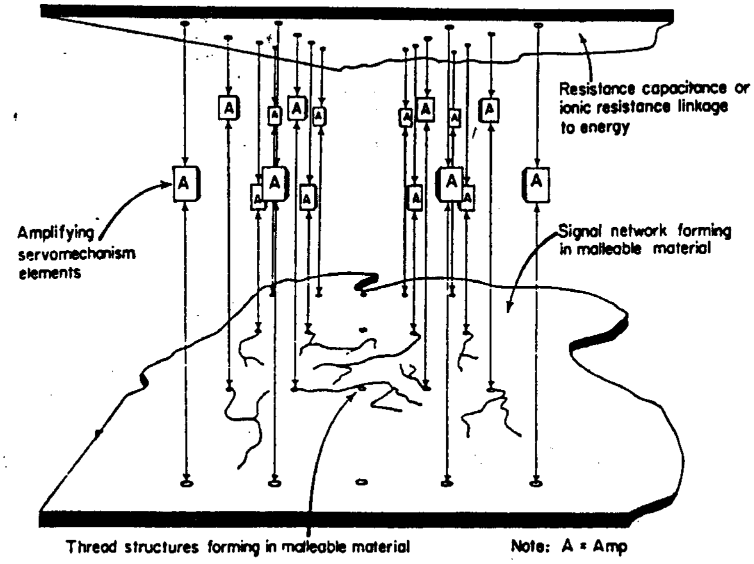

Figure 1. Pask's schematic indicating the relationship between the electrode array and the ferrous sulphate medium. From [28].

Some close friends of Pask's, like Stafford Beer, were attempting to use populations of biological organisms (such as the water flea Daphnia) to compute complex functions. The advantages of biologically-based elements revolve around their ability to self-regulate and self-proliferate; their disadvantages involve the difficulties of steering such elements in directions contrary to their natural homeostatic tendencies. Today's nanotechnologists face similar dilemmas as to which strategy to pursue: biological-evolutionary elaboration vs. mechanical, direct specification [15]. Whether biological or inorganic, it was important that the elements could be grown in great numbers so that large scale adaptive networks (analog and/or digital) could potentially be built. This strategy would start with a plastic medium with a rich set of possible structures and let the medium self-organize guided by appropriately structured reward system. The elements could proliferate themselves and the reward constraints could then mold their connections to form a functioning device.

At the time there were also people who were contemplating the prospects of having to wire up extremely large computing machines and were looking for cheap, "self-wiring" analog elements which could be grown to do the job (D.M. MacKay, pp. 924-925, in [27]; see also [1]). Remarks from the mid-1960's give the flavor of this strategy:

> "We believe that if the 'complexity barrier' is to be broken, a major revolution in production and programming techniques is required, the major heresies of which would mean weakening of machine structural specificity in every possible way. We may as well start with the notion that with 10 000 000 000 parts per cubic foot (approximately

> equal to the number and denity of neurons in the human brain), there will be no circuit diagram possible, no parts list (except possibly for the container and the peripheral equipment), not even an exact parts count, and certainly no free and complete access with tools or electrical probes to the "innards" of our machine or for possible later repair.....We would manufacture 'logic by the pound', using techniques more like those of a bakery than of an electronics factory." [39]

Many of these early forays into self-organizing devices passed current through metallic structures (iron, tin, silver) immersed in an acidic milieu (sulphuric, nitric acid), often in capillary tubes or dishes. The potential complexity of the behavior of these electrochemical assemblages was well appreciated by those familiar with the "iron-wire" neural models that had been around since the turn of the century. These physical models were capable of astonishingly nerve-like properties [20]. From 1909 into the mid-1930's R.S. Lillie investigated these properties as a potential model for nervous conduction. His iron wires in nitric acid propagated electrical disturbances down their lengths, causing refractoriness and recovery in their wake, they had thresholds for initiating these travelling pulses, they could be excited or inhibited by electric currents, they exhibited threshold accomodation and oscillatory, rhythmic behavior. Like myelinated nerve fibers, when these wires were intermittently shielded to expose only discrete nodes to their nitric acid milieu, the wires exhibited a rapid saltatory conduction, their pulses jumping from node to node. The interplay between the iron-wire physical model and the developing theories of the neuron continued well into the 1950's.

Perhaps because of these and other considerations Pask embarked on a long series of electrochemical experiments plating out metals in solution. Most of these used an array of platimum electrodes immersed in a dish containing an acidic aqueous metal-salt solution (e.g. ferrous sulphate). By passing current through the electrode array (either transiently, through capacitative discharge or through a more slowly changing source), dendritic metallic threads could be grown (Figure 1). By choosing which electrodes to pass the current between, one could control the growth of the dendritic structures. For several years Pask tried various kinds of aqueous environments, temperatures and catalysts (e.g. vanadium). Unfortunately, details concerning the many particular conditions that were tried remain obscure, although a few anecdotes concerning (potential and actual) explosions and the first emergence of sound amplification survive. By 1958, Pask had a rudimentary demonstration device working, one which could serve as an existence proof that a control system could be built which evolved its own relevance criteria [26].

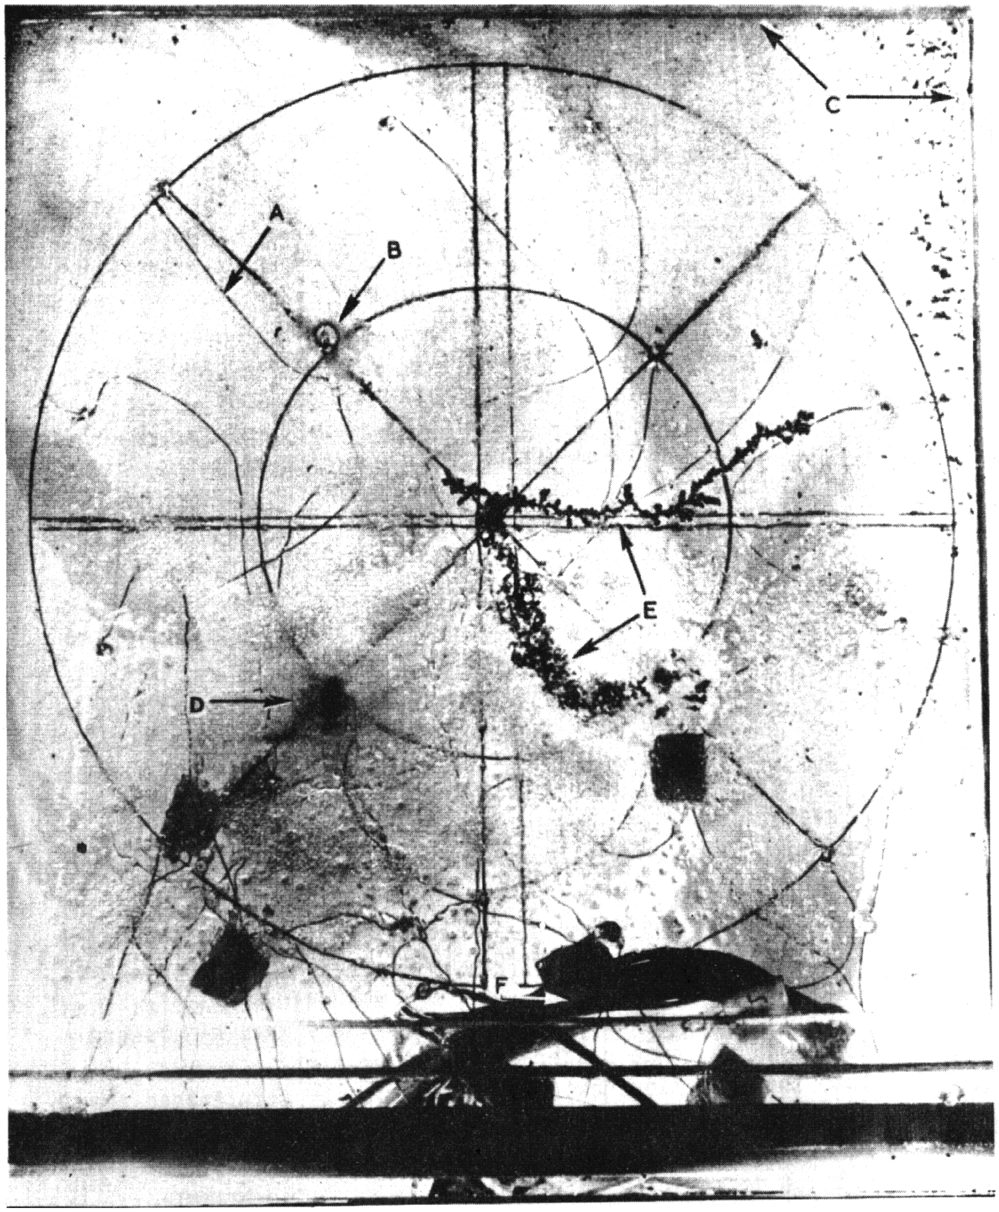

Figure 2. Photograph of Pask's electrochemical assemblage. The electrodes run perpendicular to the page. The circular wires are a support frame. Two dendritic iron thead structures can be seen in the righthand quadrants. The large dark area in the lower left quadrant is undissolved ferrous sulphate. From [27].

Pask's device premiered at the seminal "Mechanization of Thought Processes Conference" sponsored by the National Physical Laboratory in November, 1958, possibly the last large meeting to encompass representatives from all of the various approaches to the general problem of artificial intelligence, from direct programming (McCarthy, Minsky, Backus, Hopper, Bar-Hillel) to neural nets (Rosenblatt, Selfridge, Uttley) to cybernetics and self-organizing systems (Ashby, Pask) to neurophysiology (Barlow, McCulloch, Whitfield). Fittingly, Pask called his presentation "Organic analogues to the growth of a concept" [27].

As Pask pointed out, one could physically implement an analog perceptron with such an assemblage: the conductances between electrodes in the array would correspond to their connection weights. From his point of view, however, this would have been beside the point. Instead, the thread structures could be steered and selected to become sensitive to other kinds of perturbations, such that they could be tuned with the appropriate rewards. By rewarding conductance changes associated with a particular kind of environmental disturbance, the assemblage could evolve its own sensitivities. In roughly half a day ferrous threads could be adaptively grown to become sensitive either to sound or to magnetic fields:

> "We have made an ear and we have made a magnetic receptor. The ear can discriminate two frequencies, one of the order of fifty cycles per second and the other on the order of one hundred cycles per second. The 'training' procedure takes approximately half a day and once having got the ability to recognize sound at all, the ability to recognize and discriminate two sounds comes more rapidly....The ear, incidentally, looks rather like an ear. It is a gap in the thread structure in which you have fibrils which resonate at the excitation frequency." --Gordon Pask, [28], p. 261.

Design principles for self-organizing devices

> "the U-Machine must be enabled to construct its own components, and this fluid and evolutionary, self-designing process should not be irreversible...A high-variety material is required which can be topologically constrained, and reversibly, by low energy inputs. Structuring of the fabric thus obtained must supply requisite variety for absorbing input variety. The structuring and its associated measures of information must be "readable": not indeed in the (by now) trivial sense of offering a digitized output, but as mapping itself onto an external situation from which feedback can be supplied to the inputs. None of these activities needs to be a linear function, nor even a definable function, of input. The whole assembly is a black box, and it needs no designing. In it solutions to problems simply grow, as Pask's metallic threads grow." [7]

> "What I am trying to bring out is a basic distinction which exists between reward, as used in computer programme to mean that a rewarded event becomes more probable, and reward as it used in a system able to create fresh components and parts of itself. In this latter case, reward means ability to develop, ability to expand, and ability to become stable by becoming a larger system."[27], p. 928

> "But there is no possible way in which a control mechanism built of elements with well specified functions to perform, can acquire special sensitivity to an input not originally specified as relevant. On the other hand, it would be surprising if an extensive control mechanism built of elements with an initially unspecified function did no behave in this manner. Thus, for example, vibration may not be included in the list of relevant inputs, but vibration may so modify the state of an extensive control mechanism that it elicits a particular decision. Suppose this decision is favored, so that when it is made, more current is allowed to pass. In this case the current or currency will be used to construct a region in which the elements have acquired the function of vibration receptors -- in other words -- the extensive control system seeks current for building itself by forming a region sensitive to vibration.

> In general, the extensive control mechanism is able to develop relevance criteria, and to examine initially unspecified attributes of its surroundings. It does so because it is building up initially functionless elements which acquire a function as components of the system. This is characteristic of organic assemblages and decision makers who are able to laugh, when asked to make decision about bags of black and white golf balls." [26]

Some of the design principles embodied in Pask's device are: 1) construction, reconstruction and repair of its own parts (structural closure) 2) proliferation of alternative connected structures through branching, dendritic structural forms (increasing structural variety), 3) reward to useful structures in the form of material (i.e. current & iron) to build more structure (economic allocation of resources), 4) dynamic stabilization and de-stabilization of functional structures (performance-contingent survival) 5) finite amount of building resources (zero-sum competition and recycling of materials) 6) ill-defined structural elements (structural autonomy vis-a-vis the designer) 7) openness of structures to perturbations in their external environments (informationally open).

In some ways the assemblage resembled a coherer, the evacuated tube filled with iron filings that served as the tuning component for early radios [14, 40]. In both Pask's device and the coherer, the evolution of the conducting pathway (the iron filament) is shaped by those perturbations it is designed to detect. As Oliver Selfridge noted at the time, like the coherer, Pask's assemblage was also one of the first devices ever to construct itself without macroscopic motion:

> "It would be very nice to have a machine build another machine electronically without any physical motion involved, and this is the second such mechanism which has actually worked. The first one probably most of you are, like myself, too young to have ever heard of. It was, I think the way the first radios worked, with coherers." Oliver Selfridge, in [27], p. 926

In the dynamics of their growth, however, Pask's ferrous threads were constructed quite differently from the filaments of a coherer. Pask apparently had explicitly considered coherer-like devices, but had rejected them because they were not dynamically stabilized, hence not under the control of the reward system: "... in the case of the coherer, there is no sense in which the existence of that structure depended upon an obliterating tendency" [27], p. 928. While the filaments of a coherer aligned themselves according to the electric field they were to detect (perceptual input), Pask's threads were steered through a reward system with many more degrees of freedom.

One might imagine yet additional design principles for enhancing structural adaptivity. As Howard Pattee has often emphasized, "life depends upon records" [35]. Symbolic constraint of nonsymbolic analog processes (like the steering of dendrites) is essential if this structural search process is to have a memory. Without memory or inheritability of specification, structural information which has been obtained the hard way, through physical search and selection, must be garnered anew with each individual device. One therefore wants a system that can impart its structure to other systems so that each subsequent generation does not have to undergo the same long, adaptive steering. Thus, one might want to add also that the structure be steerable through some sort of (inheritable or communicable) symbolic control, enabling structural knowledge acquired by one generation to be passed on to others [9].

Direct technological implications: self-constructing sensors

What might be the contemporary technological implications of Pask's device? On the current scene, the closest relatives to his electrochemical assemblages would be those devices which utilize tunable sensors. At the low end of the scale would be autofocussing mechanisms for cameras; at its high end would be the adaptive silicon retinas and cochleas of Carver Mead and associates [24]. While these latter analog-VLSI devices move in the direction of making sensors more flexible and adaptive, they still optimize their tuning parameters within a well-defined set of possibilities (which have been foreseen by the designers). While such devices can find optimal tuning parameters, they cannot find relevant observables that were not, in some sense, designed into them. Purposely designed as well-specified devices constructed to behave predictably, they fall short of the kind of open-ended evolutionary possibilities afforded by Pask's ill-defined assemblages.

There is undoubtedly much still to be discovered concerning the malleable electrochemical media that Pask and others used. Today one might immerse an analog-VLSI chip in a medium like a ferrrous sulphate solution and adaptively build analog iron structures which would interact with the chip. This is not so far from experiments that have been conducted in recent years where real neurons are grown in issue culture over chips with many electrodes on their surface. If the loop is closed and the neurons are also stimulated by the electrode array, and some reward mechanism is implemented, then Pask's structurally adaptive configuration is achieved. Once the tissue and organ culture techniques are worked out, there is no reason that powerful adaptive devices could be grown via large scale bio-silicon adaptive assemblages. When -- if -- the time comes when large networks can be grown artificially, we will then need to come to grips with the profound moral responsibilities posed by bringing such autonomous entities into existence [44]. On the other hand, they are the responsibilities encountered by bringing any child, human or otherwise, into this world.

Closer to the present, artefacts which adaptively evolve their own sensors might potentially be useful as front-ends for trainable digital machines such as neural networks. Devices of this type could be combined with a computed neural network: the self-organizing assemblage would evolve the feature primitives that form the feature space within which the neural network operates [9, 10, 12]. At each step the neural network would attempt to partition the feature space. If the desired levels of performance could not be achieved with a given set of feature primitives, then a Pask assemblage would be put into operation to find more appropriate features and the cycle would begin anew.

The fate of his experiments

For almost 30 years Pask's device has languished in obscurity. This is partly because few people at the time understood what the device was supposed to do. Many mathematicians, scientists, and engineers understand very clearly what it means to perform a computation in the sense of a formal operation (mainly because of the ubiquity of arithmetic, calculators and computers), while few understand with comparable clarity what it means to make a measurement. And while we can readily see and agree on what happens when we run a program or train a perceptron, it is much more difficult to understand what is going on in an (purposely) ill-defined system like an electrochemical assemblage.

Structural factors in the intellectual history of information processing also contributed to the device's neglect. The device was meant more as an existence proof than a new technology that would compete with existing off-the-shelf sensors and effectors. As was then and is still now the case, conceptual demonstrations lacking obvious market potential are not highly valued, except as esoteric curiosities. And devices such as the digital computer which evolve into large industries create entire worldviews and mold the thinking of the armies of engineers that design, build, manage, and maintain them. Once the digital electronic computer had gained hegemony in information processing, it became difficult if not impossible for large segments of the engineering community to conceive of devices based on radically different design principles. Today anyone attempting to develop such alternatives must contend with the predominance of the digital worldview.

Along with the capture of imaginations, economically nascent technologies tend to gradually dominate the governmental structures which fund research, thereby further consolidating their hegemony. By the mid-1960's, much of the funding for alternative, bottom-up approaches to artificial intelligence (e.g. neural nets, evolutionary programming, cybernetics, biological computation, bionics) had dried up. This came about, in part, after a campaign against such alternatives was waged by advocates of symbolic, logic-driven artificial intelligence [16].

When funding for alternatives disappeared, many researchers found that they either had to adapt (go digital) or perish (get out of the field). Like many other researchers in cybernetics and adaptive machines, by the mid-1960's Pask had moved on to other realms that could be implemented on a digital computer: computer-aided learning and conversation theory.

While he rarely made explicit references to his earlier wet work in his subsequent papers, the lessons that Pask and his contemporaries learned from his electrochemical experiments did seem to influence many of the basic concepts that were to be used later on. These revolve around 1) how an external observer determines when a device or agent has acquired a new distinction/concept/sensor (the problem of recognizing functional emergence) 2) the functional structure of the observer, 3) the notion of the self-constructing, epistemically-autonomous ("organizationally closed") observer and 4) networks of such observer-participants.

Descriptions of observers and their interactions

If one is to build a device which purportedly has the capacity to evolve new functions, then one must grapple with problem of recognizing when this has occurred. Only a dadaist would think of building a device without some recognizable purpose, or without some perceptible effect whose utility could be evaluated. How *does* one recognize an emergent property? One needs appropriate operational criteria for functional emergence, such that a community of observers can independently verify for themselves that an emergent event has ocurred. Pask was very careful in this respect; many of his discussions about the electrochemical experiments delve deeply into systems theoretic considerations (a la Ashby [3, 4, 5]) of exactly how the situation and its evolution could be described and observed [27, 28, 29, 30, 34]. The ill-defined nature of the electrochemical milieu necessitated a sophisticated account of role of the external observer attempting to "track" what is going on in the assemblage and its environment by altering his/her model [27, 28, 29, 34]. An emergent event can be defined as the point where the observer's model breaks down, or in Rosen's terms, the deviation of the observed behavior from the behavior predicted by a model [8, 9, 10, 37, 38]. When devices alter their structure to evolve new informational linkages with their environments (new observables), their behavior becomes contingent upon new factors, and consequently changes relative to a fixed model. Thus "self-organization" (in the sense of structural rearrangement) and "emergence" are related.

Both Pask and Ashby took great pains (and apparent delight) in rigorously defining these terms. Pask's assemblage is "self-organizing", "emergent", and implements a set of distinctions on the world, thereby realizing a "system". However, it is not a "self-organizing system", because "self-organizing system" itself is a contradiction in terms, as they both argued. Since a fixed set of observable distinctions defines a "system" [3, 29], the evolutionary addition of new observables does not modify the existing "system", it creates an entirely new one [30]. Further, new observables come about by factors from outside (rather than from within) the set of existing distinctions, so it is not a *self*-organizing system [5]. The clash is an incompatibility between the notion of a model or system as a fixed point-of-view versus a materially-embedded set of modelling relations (observables) that are changing over time.

The structure of the observer-participant

All of the basic properties of the observer can be discerned in the relationship between the electrochemical assemblage and its environment. Above all the assemblage has the capacity to interact with the world, to classify the apparent state of the world based on this interaction, in short, to make measurements, to draw distinctions, to form concepts. As such the device is physically evolving its own observables, not unlike the building of the measuring devices that carry out observations for use in a scientific model. A reward system motivates this search for appropriate "relevance criteria". This is not so far removed from the psychological construction of concepts and the systems of internalized rewards that drive that process.

A set of these concepts or distinctions forms a "reference frame" through which an observer ("observer-participant") apprehends the world. This idea is thus intimately related to Ashby's theory of systems (of observable distinctions)[3, 4, 6], Uexkull's *Umwelt* ("life-world" [43]), "frames of reference" in quantum mechanics [25], and various accounts of "modelling relations" embedded in biological systems [9, 10, 12, 17, 18, 19, 36, 37, 38].

Subserving the epistemic functionalities ("distinctions", "observables") of the observer is a material substrate interacting with the rest of the material world. The material substrate makes possible the distinctions the observer makes on the external world and the influences that the "observer-*participant*" can have on it. Thus the limits of the observer-participant are the physical limits of the underlying material structures. This point is made vividly clear by Pask's assemblage: while it is "ill-defined" and thus "open-ended", it is nevertheless bounded by its own structural possibilities. At any given time the assemblage can only make those distinctions and carry out those actions that can be implemented with the system of ferrous threads that is in place.

The self-construction of the observer

When a device gains the ability to construct its own sensors, or in McCulloch's terms "this ability to make or select proper filters on its inputs" [29], it becomes *organizationallyclosed*. The device then controls the distinctions it makes on its external environment, the perceptual categories which it will use. On the action side, the device acquires the ability to construct its own effectors, and with them gains control over the kinds of actions it has available to influence the world. The self-construction of sensors and effectors thus leads to an *epistemic autonomy*, where the organism or device itself is the major determinant of the nature of its relations with the world at large [9, 10, 37]. This basic concept of structural closure and its consequent, functional autonomy,underlies many of the closely related notions of *semantic closure* [36], autopoiesis [21, 22, 41], self-modifying systems [13, 19], self-reproducing automata [42], anticipatory systems [38] and the recurrent, "nets with circles" of McCulloch and Pitts [23].

Pask has proposed this organizational closure as one of the constituitive conditions for consciousness: "A process is potentially conscious if it is organizationally-closed, informationally open, and if information is tranferred across distinctions that are computed *as required to permit the execution* of the process." ([32], p. 214; see also [33]).

Networks of observer-participant elements

Another idea implicit in Pask's assemblages is the notion of a network of elements (a society of actors) adaptively constructing their own observables, their own ways of seeing the world, as they interact and communicate with each other. In the early 1960's Pask proposed a series of evolutionary models (predating contemporary artificial life models by 25 years) in which mortal, reproductive actor-automata searched for food and formed communicative coalitions for finding it [29, 30, 31]. Here the currency is food rather than electrical current and the connectivities are in terms of inter-actor communicative capabilities rather than conductances, but the basic framework of the interactions is still like that of the electrochemical devices. While simulations such as these are very useful in undertanding in a generic way how these processes could work, they are not ill-defined systems, since all simulation states and rules are known. While Pask never intended for these simulations to be substitutes for physical realizations, it is thus easy for the unwary (e.g. many contemporary artificial life researchers) to believe that the discrete simulation has all of the actual and potential properties of an analog construction such as Pask's assemblage [9, 11, 19, 37, 38]. Rather than realizing new observables in an open-ended way, such simulations are bound by their closed set of state variables (i.e.their "state space"). They fail to achieve epistemic autonomy relative to their designers because all possibilities have been prespecified by their designers; they have no structural autonomy. A digitally simulated agent (as opposed to a material device or a human-machine system) thus loses its “ill-defined” character relative to its programmer-designer, and with it the ability to defy the designer’s categories for describing its behavior. We are back to Descartes’ Dictum. If this line of reasoning is valid, then we need to return to Pask’s earlier strategy of*building*, rather than simulating, actual physical devices.

In contrast to his evolutionary simulations and his later highly formalized, abstract discussions of entailment meshes and consciousness, Pask's electrochemical assemblage has the distinct pedagogical advantage of concreteness: it grounds us in the realm of the sensuously apprehendable material world. It is too easy to lose the broader research program, the *realization* of self-constructing, conscious, autonomous agents capable of open-ended learning, when one descends into the infinite labyrinth of formal *descriptions* of how they might work.

Were we to go back to building physical devices, replication of his electrochemical assemblages would be a good first step. Eventually we would want to make networks using collections of electrochemical assemblages ("Paskian elements"). A set of materially realized networks of Paskian elements would have properties radically different from contemporary neural networks. Such networks cannot fail to have implications for how we think about the brain [10]. Pairs of elements could evolve their own modes of signalling by evolving compatible effector-receptor combinations. By virtue of the its particular sensitivities and capabilities for producing disturbances in the common medium, an element could be tuned to preferentially sense the actions of other particular elements. Similarly tuned elements could thus act together. To borrow a metaphor from radio, elements operating "on the same wavelength" could selectively (inter)act. In addition, new tunings orthogonal to those already in the network could be formed. The network would thus be self-organizing in a way that the current neural networks are not: the dimensionality of the signal space can increase over time as new informational channels evolve. Hill-climbing is thus accomplished not only by following gradients upwards, but also by changing the dimensionality of the problem landscape when one can go no further using those dimensions already available.

This is not unlike what goes on as we engage in conversations with each other: our concepts (especially our own word-meanings and our models for the word-meanings of others) are continually being constructed and reconstructed as our real world interactions progressively add degrees of constraint. We are led to a conception of evolving actors and their interactions ("conversations"). When our concepts of the other actors break down, we are forced to come up with new concepts that allow us to make sense of (and perhaps to better predict) their behavior. Sometimes these new concepts enable new modes of communication and interaction. Like the electrochemical demonstration, this evolution takes place not in a circumscribed space of well-defined possibilities, but in an ill-defined, and therefore "open-ended" space of possible distinctions and actions. Open-ended possibility confers upon us the option "to laugh, when asked to make decision about bags of black and white golf balls." [26]

Such is the nature of biological creativity in its most fundamental form.

References

1. Ainsworth, W.A. Electrolytic growth of silver dendrites. *Science.* 146 (1964), 1294-1295.

2. Ashby, W.R. Can a mechanical chess-player outplay its designer? *British J. for the Philosophy of Science.* 3 (1952), 44-57. Reprinted in *Mechanisms of Intelligence: Ross Ashby's Writings on Cybernetics*, R. Conant (ed.), Intersystems Publications, Salinas, California (1981)

3. Ashby, W.R. *An Introduction to Cybernetics.* Chapman and Hall, London (1956)

4. Ashby, W.R. *Design for a Brain.* Chapman and Hall, London (1960)

5. Ashby, W.R. Principles of the self-organizing system. In H. von Foerster and G. W. Zopf (ed.), *Principles of Self-Organization.* Pergamon Press, New York (1962)

6. Ashby, W.R. Analysis of the system to be modeled. In (ed.), *The Process of Model Building in the Behavioral Sciences.* Ohio State University Press, (1970), 94-114. Reprinted in *Mechanisms of Intelligence: Ross Ashby's Writings on Cybernetics,* R. Conant (ed.), Intersystems Publications, Salinas, California (1981), 335-355

7. Beer, S. Towards the cybernetic factory. In H. von Foerster and G. W. Zopf (ed.), *Principles of Self-Organization.* Pergamon Press, New York (1962)

8. Cariani, P. Adaptivity, emergence, and machine-environment dependencies. 33rd Annual Metting of the International Society for Systems Sciences (ISSS), Edinburgh, Scotland, July 2-7, 1989. P. W. J. Ledington (ed.), (1989), 31-37.

9. Cariani, P. *On the Design of Devices with Emergent Semantic Functions.* Ph.D., State University of New York at Binghamton (1989)

10. Cariani, P. Adaptivity and emergence in organisms and devices. *World Futures.* 32 (1991), 111-132.

11. Cariani, P. Emergence and artificial life. In C. G. Langton Taylor, C., Farmer, J.D. and Rasmussen, S. (ed.), *Artificial Life II. Volume X, Santa Fe Institute Studies in the Science of Complexity.* Addison-Wesley, Redwood City, CA (1992)

12. Cariani, P. Some epistemological implications of devices which construct their own sensors and effectors. In F. Varela and P. Bourgine (ed.), *Towards a Practice of Autonomous Systems.* MIT Press, Cambridge, Massachusetts (1992)

13. Csanyi, V. *Evolution in Biological and Social Systems: A General Theory*. Duke University Press, Durham, North Carolina (1989)

14. Del Mar, W.A. The coherer: how to make a modern coherer and a brief review of its history and theory. *Scientific American Supplement.* 1381 (1902), 22131-22132.

15. Drexler, K.E. Biological and nanomechanical systems: contrasts in evolutionary capacity. In C. G. Langton, C. Taylor, J. D. Farmer and S. Rasmussen (ed.), *Artificial Life II. Volume X, Santa Fe Institute Studies in the Science of Complexity*. Addison-Wesley, Redwood City, CA (1992)

16. Dreyfus, H.L., and Dreyfus, Stuart E. Making a mind versus modeling the brain: artificial intelligence at a branchpoint. *Daedalus.* (1988), 15-43.

17. Heylighen, F. Formulating the problem of problem-formulation. In R. Trappl (ed.), *Cybernetics and Systems '88.* Kluwer Academic Publishers, Dordrecht, Holland (1988)

18. Heylighen, F. *Representation and Change.* Communication & Cognition, Ghent, Belgium (1989)

19. Kampis, G. *Self-Modifying Systems in Biology and Cognitive Science.* Pergamon Press, Oxford (1991)

20. Lewis, E.R. The iron wire model of the neuron: a review. In H. L. Oestreicher and D. R. Moore (ed.), *Cybernetic Problems in Bionics (Bionics Symposium 1966).* Gordon and Breach, New York (1968)

21. Maturana, H. The biology of cognition. In H. Maturana and F. Varela (ed.), *Autopoiesis and Cognition.* D. Reidel, Dordrecht, Holland (1970)

22. Maturana, H. and F. Varela. Autopoiesis: the organization of the living. In H. Maturana and F. Varela (ed.), *Autopoiesis and Cognition (1980).* D. Reidel, Dordrecht, Holland (1973)

23. McCulloch, W.S. and W. Pitts. A logical calculus of the ideas immanent in nervous activity. *Bulletin of Mathematical Biophysics.* 5 (1943), 115-133. Reprinted in *Collected Works of Warren S. McCulloch*, R. McCulloch (ed.), Intersystems Publications, Salinas, California (1989)

24. Mead, C. *Analog VLSI and Neural Systems.* Addison-Wesley, Reading, MA (1989)

25. Murdoch, D. *Niels Bohr's Philosophy of Physics.* Cambridge University Press, Cambridge (1987)

26. Pask, G. The growth process inside the cybernetic machine. Second International Conference on Cybernetics, Namur, Belgium. (1958), 765-794.

27. Pask, G. Physical analogues to the growth of a concept. Mechanization of Thought Processes, Symposium 10, National Physical Laboratory, November 24-27, 1958. H.M.S.O. (London) (1958), 765-794.

28. Pask, G. The Natural History of Networks. In M. C. Yovits and S. Cameron (ed.), *Self-Organzing Systems*. Pergamon Press, New York (1960)

29. Pask, G. *An Approach to Cybernetics.* Harper & Brothers, New York (1961)

30. Pask, G. The cybernetics of evolutionary processes and of self-organizing systems. Third International Conference on Cybernetics, Namur, Belgium. (1961), 27-74.

31. Pask, G. A proposed evolutionary model. In H. von Foerster and G. W. Zopf (ed.), *Principles of Self-Organization.* Pergamon Press, New York (1962)

32. Pask, G. Consciousness. *Journal of Cybernetics.* 9 (1979), 211-258.

33. Pask, G. Organizational closure of potentially conscious systems. In M. Zeleny (ed.), *Autopoiesis: A Theory of Living Organization.* North Holland, New York (1981)

34. Pask, G. and H. von Foerster. A predictive model for self-organizing systems. *Cybernetica.* 4 (1960), 1(1961) (1960),

35. Pattee, H.H. Physical problems of decision-making constraints. *International Journal of Neuroscience.* 3 (1972), 99-106.

36. Pattee, H.H. Cell psychology: an evolutionary view of the symbol-matter problem. *Cognition and Brain Theory.* 5 (1982), 325-341.

37. Pattee, H.H. Simulations, realizations, and theories of life. In C. G. Langton (ed.), *Artificial Life. Volume VI, Santa Fe Institute Studies in the Science of Complexity.* Addison-Wesley, Redwood City, CA (1989)

38. Rosen, R. *Anticipatory Systems.* Pergamon Press, Oxford (1985)

39. Stewart, R.M. Electrochemically active field-trainable pattern recognition systems. *IEEE Transactions on Systems Science and Cybernetics.* SSC-5 (1969), 230-237.

40. Unknown. Theory of coherer action. (Reprinted from English Mechanic and World of Science). *Scientific American Supplement.* 1558 (1905), 24958-24960.

41. Varela, F. *Principles of Biological Autonomy*. North Holland, New york (1979)

42. von Neumann, J. Re-evaluation of the problems of complicated automata -- problems of hierarchy and evolution. (1948), 477-490. Reprinted in *Papers of John von Neumann on Computing and Computer Theory*, W. Aspray and A. Burks (ed.), MIT Press, Cambridge, Massachusetts (1987)

43. von Uexküll, J. *Theoretical Biology*. Harcourt, Brace & Co, New York (1926)

44. Winner, L. *Autonomous Technology: Technics-out-of-Control as a Theme in Political Thought.* MIT Press, Cambridge, MA (1977)