Design Strategies for Open-Ended Evolution

Peter Cariani

Harvard Medical School
cariani@mac.com

Abstract

Open-endedness is an important goal for designing systems that can autonomously find solutions to combinatorically-complex and ill-defined problems. We distinguish two modes of creating novelty: combinatoric (new combinations of existing primitives) and creative (new primitives). Although combinatoric systems may differ in numbers of possible combinations, their set of possibilities is closed. Creative systems, on the other hand, have open-sets of possibilities because of the partial- or ill-defined nature of the space of possible primitives. We discuss classes of adaptive and self-modifying cybernetic robotic devices in terms of these two kinds of processes. We consider material systems constructed from genetically-directed pattern-grammars. Although spaces of accessible structures are closed, function spaces can nevertheless be open. Thus, genome sequence spaces and gene-product structure spaces are regarded as closed, while partially-defined, phenomic function-spaces are potentially open.

Introduction

Intuitively, much of the natural world appears to us to be open-ended in character. When we consider the origins and evolution of life, the appearance and evolutionary elaboration of immune and nervous systems, and add the possible concomitant emergence of consciousness, it is difficult to imagine how the universe of evolved structures, functions, and phenomenal dimensions might be predicted from basic physical laws alone. The world cannot yet be described in closed form: there are too many incommensurable categories, structures, functions, functional organizations, and material/phenomenal distinctions among them to achieve such a grand reduction.

Yet most of us also optimistically believe that a comprehensive theory of life is possible in the future once we fully understand the space of structural and organizational possibilities that physico-chemical systems afford. Perhaps even more optimistically, many of us also believe that once the neural codes and pulse computations that constitute the informational organization of nervous systems are understood, then we will be able to understand and predict the structure and contents of phenomenal experience. But even if the structure of experience is predictable from patterns of neuronal activity, its existence as an aspect of the world is still an emergent if it depends on evolution of particular kinds of complex organizations. If the phenomenal realm emerged over biological evolution, then even if it is closed under physical causation, the material world may nevertheless be irreducibly open-ended in its aspects, i.e. there is more to describing what goes on in the world than in terms of material process alone. Even leaving aside such deep ontological questions, for the foreseeable future, living organisms and nervous systems will remain systems whose structures and functions are only partially-defined for us, and therefore whose behaviors can therefore surprise us in unexpected ways. Until a predictive "theory of everything" is achieved, if one ever is, living systems will continue to appear to us to be capable of open-ended self-modification.

The Importance of Open-Ended Design

Open-endedness is an important goal for designing creative systems. Creative systems are needed when we face ill-defined problems that defy direct solution, when we don't know what observables (sensors, features) and actions (effectors) are needed, and how they should be coupled and controlled (coordinations, computations). In these cases, we want the system itself to come up with a solution that we have not in some sense foreseen (or we would design that solution by fiat). We therefore seek to design and construct devices that act autonomously to go forth into the world to interact with it, to modify themselves in some way in order to find solutions that we cannot already anticipate. Open-ended devices are critical if we are to build robots that autonomously construct their own meanings and artificial immune systems that automate the search for new pharmaceutical agents.

Reliability and Closure

Unlike naturally evolved biological organisms, at present most of our artefacts are designed and constructed to behave in completely reliable and predictable ways that efficiently satisfy our needs. When they are performing within specifications, their structures are well-defined and highly constrained; they are expressly designed *not* to surprise us. The physical hardware of the modern digital electronic computer is the epitome of reliable design – astonishingly complex computations are invariably carried out without error. We manifestly want to avoid surprises (errors) creeping into our computations. When errors do occur, our systems are designed to immediately terminate computations and to indicate that a failure has occurred. In these reliable real-world, finite computational systems, to the extent that we specify all aspects of our devices (structure, operation), we know all possible input-output behaviors. We can circumscribe this closed set of possibilities, and no novel states of input-out behaviors will occur that will lie outside this box.

Open-endedness, Novelty, and Epistemology

Open-endedness requires creation of novel entities. Novelty requires some degree of ignorance – if all the parts and laws of a finite system are perfectly known, then all of the system's possible states and behaviors are known. Novelty (and hence open-endedness) is simply not possible in discourses where one assumes an omniscient, complete, God's-eye view of the world (e.g. realist-materialist and platonic ontologies). By their inherent construction, such discourses in the axiomatic-deductive mode categorically disallow *de novo* creation of new primitives.[1] Effectively, omniscience in a given realm implies closure; partial-knowledge permits the possibility of open-ended surprise. Novelty *is* possible in discourses where a limited observer compares the observed behavior of a system with his/her predictive model of it, since processes unrepresented in the model can cause the system's behavior to deviate from expectations ("emergence-relative-to-a-model"; (Cariani, 1989; Cariani, 1992; Cariani, 1997; Rosen, 1985). We therefore believe that an epistemological stance is necessary when we confront problems involving novelty, creativity, open endedness, and emergence.

In order for open-endedness to be a meaningful and useful criterion for considering natural and artificial systems, it should be principled (not an *ad hoc* construction) and clear; we must be able to construct operational definitions that allow us to unambiguously determine whether a given system is open-ended or not vis-à-vis some criteria. In order for us to ask whether a system has produced novel behavior, we first must ask the question of exactly what are our expectations: "novel relative to what?" In practice, change must be measured relative to some state-of-affairs, some concrete set of expectations we have of the system's structure and organization. Although operational criteria have been developed for restricted kinds of emergent functionalities (see below), open-endedness is a broader and less easily defined attribute than either closure or emergence-relative-to-a-model, mainly because it deals in spaces of possibility rather than the circumscribability of sets of elements.

A simple example (Fig. 1.) is helpful in conveying the differences between closed vs. open-ended realms The set of all 6-digit permutations of digits 0-9 is well-defined and contains 6^{10} elements, which can be enumerated. The set of all permutation sequences of 6 arbitrarily defined objects, however, is ill-defined, because the number of possible objects is indefinite. As a result this latter set is unbounded, ill-defined, and open-ended – one can always augment the set by specifying 6 more objects. In the first case, the primitives are exhaustively described by their token-types; consequently, the set is well-defined and closed. In the second case, the space of possible primitives themselves are not well-defined, and therefore the set of possibilities is ill-defined and open. Like the set of all

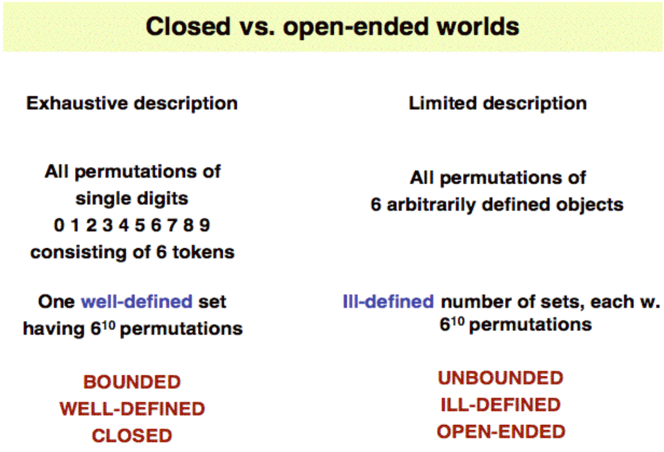

Figure 1. Closed vs. open sets of possibilities.

---

[1] In his 1975 debate with Jean Piaget about the possibility of new ideas (mathematical systems) appearing over (historical) time, Jerry Fodor famously, in platonic-realist fashion, argued for a closed universe in which there are no new, emergent ideas, but instead only selective fixation of previously existing ones (Fodor, 1980)

possible distinguishable objects, the set of possible measurements (observables) and actions that can be carried out respectively by sensors and effectors is ill-defined and open. This means that biological organisms and artefacts that are capable of evolving new sensors and effectors have an open-ended set of possible ways of interacting with the world, and, further, that the space of possible epistemic life-worlds, *umwelts* (Uexküll, 1925), is open-ended.

Combinatoric vs. Creative Novelty

One can envision systems that simply recombine fixed primitives vs. those that somehow create new ones. Emergent novelty can be generated in two ways: *combinatoric emergence* and *creative emergence* (Fig. 2). In a similar vein Lloyd Morgan (Morgan, 1931) distinguished "emergents" from "resultants": emergents being the result of novel creation, resultants, of novel combination. Both kinds of emergent orders are built up from basic sets of possibilities that constitute the most basic building blocks of the order, its “primitives.” Emergence then entails either the appearance of new combinations of previously existing primitives or the formation of entirely new ones. The primitives in question depend upon the discourse; they can be structural, material "atoms"; they can be formal "symbols" or "states"; they can be functionalities or operations; they can be primitive assumptions of a theory; they can be primitive sensations and/or ideas; they can be the basic parts of an observer's model. To say that an entity is "primitive" relative to other objects or functions means it cannot be constructed from combinations of the others, i.e. its properties cannot be logically deduced from those of other entities. Thus, in this way of thinking, simple combinations of “lower-level objects” do not create “higher-level primitives” because the higher-level systems can be decomposed into yet lower-level objects (atoms).

Combinatoric Novelty and Closure

Combinatoric emergence assumes a fixed set of primitives that are combined in new ways to form emergent structures. This is very compatible with the way we often think about structure spaces, where parts can be combined to form larger structures. Thus in biological evolution, new genetic DNA sequences arise from combinations of pre-existing nucleotides, codons, and codon-sequences. Microevolution entails generation of novel combinations of genes; new genes arise through novel combinations of nucleotide sequences. Likewise, new, emergent structures are thought to arise from novel combinations of previously existing molecular, cellular, and organismic structures.

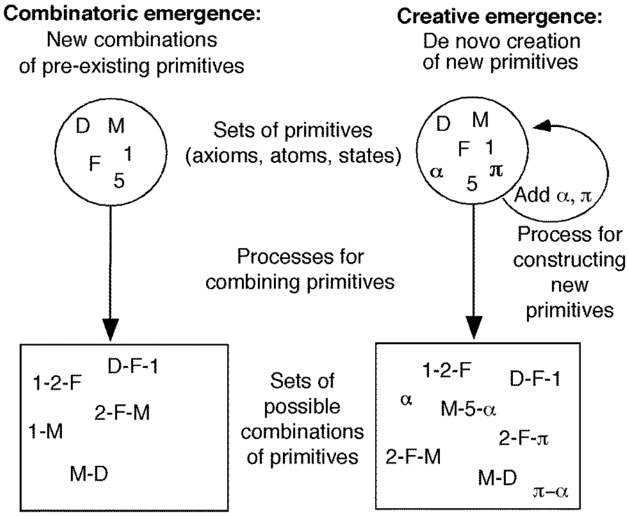

Figure 2. Combinatoric vs. creative emergence.

This strategy for generating variety from combinations of relatively small set of primitive parts is a powerful one that is the basis of the systematicity of human and computer languages. Digital computers are ideally suited for generating combinations of symbol-primitives and logical operations on them that can then be evaluated for useful, interesting, and/or unforeseen formal properties. Correspondingly, in the realm of adaptive, trainable machines, directed searches optimize combinations of pre-specified features and actions (i.e. feature-action mappings, classifications). What formally distinguishes different kinds of trainable machines, such as neural networks or genetic algorithms, are the structures of the respective combination-spaces being traversed, and the rules that direct the search processes through them. In artificial life contexts, genetic algorithms using generative pattern grammars search through complex quasi-organic structure spaces[2] or find more optimal percept-action coordination strategies for simulated robots and organisms. In both types of applications, search spaces are large, but nevertheless closed.

Closure with Ill-defined Elements

We have argued above that well-defined finite sets are closed, while ill-defined, indefinite sets are open-ended. But what about sets of ill-defined elements? The genetic algorithms and pattern grammars mentioned above

---

[2] Dawkins demonstrated his Blind Watchmaker evolutionary graphics program at the first workshop on Artificial Life in 1987 at Los Alamos (Dawkins, 1987).

involve selection of well-defined, discrete entities (in the Blind Watchmaker program, these are discrete graphical elements; in a robotic controller, they are parameter values). However, combinatoric strategies can be used to select combinations of ill-defined parts that can interact in nonlinear, and unpredictable ways. Despite this incorporation of ill-defined elements, the set of possible behaviors is still closed under the set of discrete possibilities of the selection process. For example, if we had considered the set of 6-object permutations of 10 distinguishable, but ill-defined objects in the above example in Figure 1, the set of permutations would have 6^{10} members. Here, even though we don't know all the properties of the objects themselves, we can reliably treat them as individuals (name and distinguish them), and therefore draw a box around the space of possible permutations. We were not able to do this for sets of arbitrarily-defined objects, because there is no clear method by which we can clearly enumerate their elements or circumscribe their content.

Even if the set of possible combinations is closed, it may be useful to consider the respective cardinalities of two different systems as a comparative measure of structural complexity. In biological contexts, many different structural and functional criteria are possible: numbers of cells, cell types, expressed genes, protein conformations, metabolic states, informational states, etc. (Bonner, 1988). Complexity, however, does not by itself beget open-endedness. However staggeringly large the combinatorics become, mere number alone does not transform a closed set into an open one (finite, but large ≠ infinite, indefinite).

Ashby's Homeostat: Combinatoric Adaptivity

An apt historical example of combinatoric novelty using ill-defined elements is the homeostat of Ross Ashby (Ashby, 1960; de Latil, 1956). The homeostat consisted of four subsystems each in dynamic equilibrium with the others (Fig. 3). In each subsystem was a 25-position "uniselector" switch that determined the analog control parameters (capacitance, resistance) of that subsystem's electronic circuit. The circuits were "randomly" constructed and assigned to the uniselector positions, such that their structure and arrangement was not critical to the device's operation and might not even be fully known by the device's designer or user. The homeostat therefore had 25x25x25x25 (390,625) clearly defined uniselector-combination states that determined ill-defined analog control parameters and their associated behaviors. Particular combinations of parameters in interaction with a particular external signal could lead to stability or to chaotic instability.

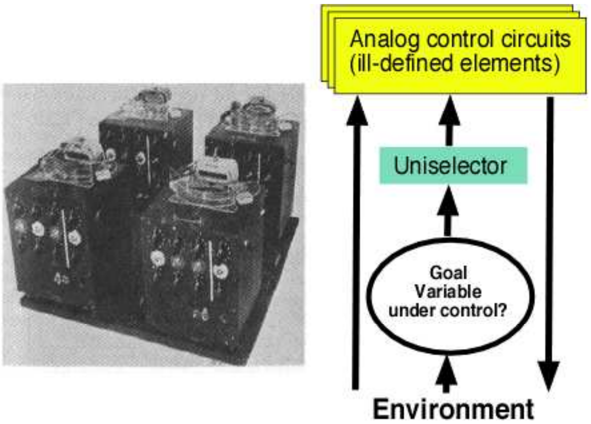

Figure 3. The homeostat and its operational structure.

The goal of the homeostat was to keep the value of a control variable near a given goal state, within specified tolerances. The homeostat thus evaluated whether a particular set of circuit parameters (resistances, capacitances) made a "good controller" vis-à-vis a particular environment. If the controlled variable did not achieve stability within some specified period of time, changing the positions of the uniselector switches would randomly choose another set of parameters to be tested.

The homeostat is a device that has no explicit model either of its environs or its internal workings. As de Latil 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" (p. 308). Ashby also realized that not only could the homeostat be ignorant of the details, so could the designer: a designer need not understand at all how any of the analog controllers worked in order to choose which one worked 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. But 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.

The homeostat may well have been the first artificial adaptive device to incorporate this principle of an "ill-defined" adaptive system, a principle that Gordon Pask was to carry to an extreme a few years later in his electrochemical assemblages (see below).

Limits of Combinatoric Novelty

Combinatoric novelty is a dynamic, creative strategy insofar as it constantly brings into being new combinations of elements. However, its use of fixed sets of primitive elements mean that the set of possible combinations is closed. In the example of Fig. 2, one cannot create new alphabetical letter types by stringing together more and more existing letters – the new notations must be introduced from outside the system by external agents or processes. Similarly, the homeostat could switch between 390k different circuits but it had no way of creating new circuits or of modifying existing ones to carry out new functions. Had the homeostat possessed the means of perturbing the structure of the circuits in an unforeseen way, say contingent on the structure of environmental input, then the device would have had an open-ended structure.

Within a computer simulation, all simulated activity occurs within the state-space and determined by the rules of the simulation program. However, if the observer is ignorant of the program, even partially, or if the computer is connected to unpredictable, external inputs, then novel behaviors vis-à-vis the observer's expectations *can* occur, and new primitives can potentially be created (e.g. a computer suddenly starts displaying Asian ideograms in addition to Roman text.) In such circumstances the computer's behavior would appear open-ended relative to the observer's set of expectations.

Creative Emergence and Open-endedness

Classically, "emergence" has concerned those processes that create new primitives, i.e. properties, behaviors, or functions that are not logical consequences of pre-existing ones. One can always ask how the particular primitives of an existing combinatorial system came into being in the first place. In explaining the origins of new primitives, one must appeal to additional processes that are not the primitives themselves. For example, how were the symbols depicted in Fig. 2 fabricated in the first place? By what process can new symbol types be added? In biological systems, how did nucleotide molecules strung together become the primitives of a genetic code?[3]

Primitive objects in the physical world almost always contain properties not fully known to the observer that can support new functions. These hidden aspects can come into play as primitives interact through the underlying material processes that subserve them. In this latter view, creating a new primitive entails the formation of a new property or behavior that in some strong sense was not predictable (by the limited observer) from what came before.

Open-ended Evolution of New Sensors

It is usually easier to give examples of qualitatively new functions than examples of qualitatively new structures. In our opinion, the most salient examples of the creation of new primitives involve the biological evolution of new sensory capabilities. Where previously there may have been no means of distinguishing colors, odors, or sounds, eventually these sensory capacities evolve in biological lineages. From a set of primitive sensory distinctions, one can list all combinations of distinctions that can be made with those primitives, but there are always yet other possible distinctions that are not on the list. For example, we cannot combine information from our evolution-given senses (sight, hearing, smell, etc.) to directly detect low intensity electrical or magnetic fields in our midst (as is achieved by electroceptive fish and some migratory birds, respectively). Creation of the ability to sense these fields through biological evolution, or artificial construction of measuring instruments (magnetometers, field strength sensors), thus adds new primitives to the set of perceptual distinctions that can be made.

Artificial Sensor Evolution

Artificial devices that create new perceptual primitives have been built. A perspicuous example is a electrochemical device that was constructed by the British cybernetician Gordon Pask in the late 1950's (Cariani, 1993; Pask, 1958, 1959, 1960, 1961). Its purpose was to show how a machine could evolve its own "relevance criteria." The structure of the heart of the analog device itself was hopelessly ill-defined. Current was passed through an array of platinum electrodes immersed in an aqueous ferrous

---

[3] A well-known paper by theoretical biologist Howard Pattee was entitled "How does a molecule become a message?" Dev. Biol. Suppl. 3:1-6, 1969.

sulphate/sulphuric acid equilibrium, such that iron dendritic filaments grew to form bridges between the electrodes. By rewarding iron structures whose conductivity contingently varied with environmental perturbations, the set of structures could be adaptively steered to improve the sensitivity of the whole. Pask's device acquired the ability to sense the presence of sound vibrations and then to distinguish between two different frequencies. In effect, the device had evolved an ear for itself, creating a set of sensory distinctions that it did not previously have. Albeit, in a very rudimentary way, the artificial device automated the creation of new sensory primitives, thereby providing an existence proof that creative emergence is possible in adaptive devices.

Evolvable Cybernetic Systems

Pask's device is a special case of a broader class of devices that are capable of modifying their own internal structure in open-ended ways. One can formulate a taxonomy of possible cybernetic devices and their creative capacities (see Cariani, 1989, 1991, 1998). These robotic devices consist of sensors and effectors coupled together by means of computational coordinative modules with well-defined internal symbolic states (Fig. 4). These devices have an evaluative part that directs the construction and modification of the hardware that subserves faculties of perception, cognition, evaluation & reward, and action. This hardware includes sensors, effectors, and the internal computational mechanisms that mediate sensorimotor coordination by implementing particular percept-action mappings. The evaluative part contains memory, learning, and anticipatory mechanisms for measuring performance, changing percept-action mappings, and adaptively modifying internal structures to improve performance. A methodology has been developed to distinguish between these functionalities and to determine when a new measurement, computation, or action is created. We believe they capture the basic operational structure of the observer-actor.

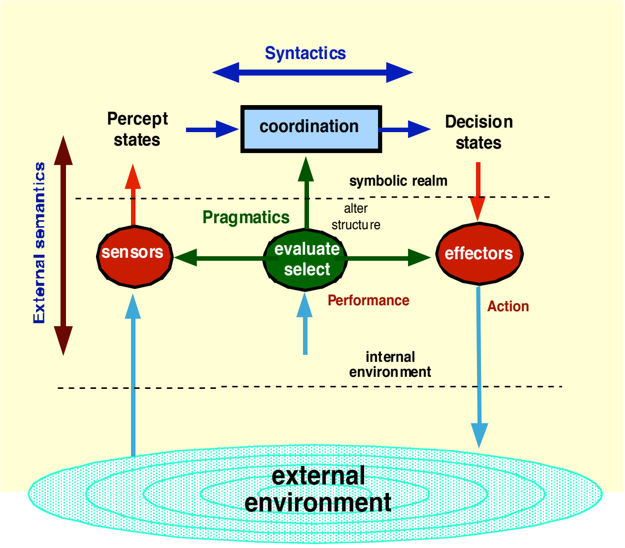

Figure 4. Self-modifying cybernetic devices.

Such cybernetic systems can be described in terms of semiotic categories: syntactic, semantic, and pragmatic dimensions. Syntactics describes rule-governed linkages between signs that are implemented in computational, coordinative portions of devices. Semantics involves the relation of signs to the external world, i.e. causal linkages between internal symbolic states and the world that are mediated by sensors and effectors. Finally, pragmatics involves the purposes for which signs are used: their relation to embedded goal states. Pragmatic relations are implemented by internal evaluation-reward mechanisms that adaptively steer or modify internal device linkages to better achieve embedded goals.

Within such a framework one can envision devices with both mechanisms that switch between existing sets of possible internal states (combinatoric emergence) or mechanisms that adaptively construct new hardware (e.g. new sensors, effectors, internal states) capable of creating new functional primitives (creative emergence). Table I summarizes possible types of adaptivity vis-à-vis combinatoric and creative emergence. In the syntactic realm, creative emergence produces new signs (symbols, internal states). In the semantic realm it produces new observables and actions that make new

Dimension	**Primitives**	**Stable systems** Fixed structure	**Combinatoric systems** Search/optimize existing possibilities	**Creative systems** Add possibilities Evolve
Syntactic	States Computations	Deterministic FSA's (fixed machines)	Change computations (trainable machines)	New states & rules (growing automata)
Semantic	Measurements Actions	Fixed sensors, effectors	Search combinations of existing sensors & effectors	New measurements and/or actions (epistemic autonomy)
Pragmatic	Goals	Fixed goals	Search combinations of existing goals	New goals (creative self-direction)

Table I. Combinatoric and creative emergence in cybernetic devices contingent linkages between internal states and the outer world. In pragmatic realm, it produces new evaluative criteria (new goals).

Each functionality (sensing, effecting, coordinating) can be either be fixed, subject to combinatorial search, or capable of *de novo* creation of new primitives (Table I, above). In this scheme, combinatoric creativity involves new combinations of pre-existing input and output states, sensors, effectors, and goals. Creative emergence requires going outside of the set of existing functionalities to modify material structures ("hardware") in a manner that can create new states, new sensors and effectors, or new goals.

To the degree that a system has control over its own structure and functions, it attains a degree of freedom vis-à-vis both its environment and its own history. When a system can add to its own states and state-transitions, as in a growing automaton, it achieves some degree of computational autonomy. When a system can construct its own sensors, it attains a degree of epistemic autonomy. When it can construct new effectors it attains a greater autonomy of possible actions. Finally, when the system can construct its own set of evaluations and embedded goal states, it becomes self-directing.

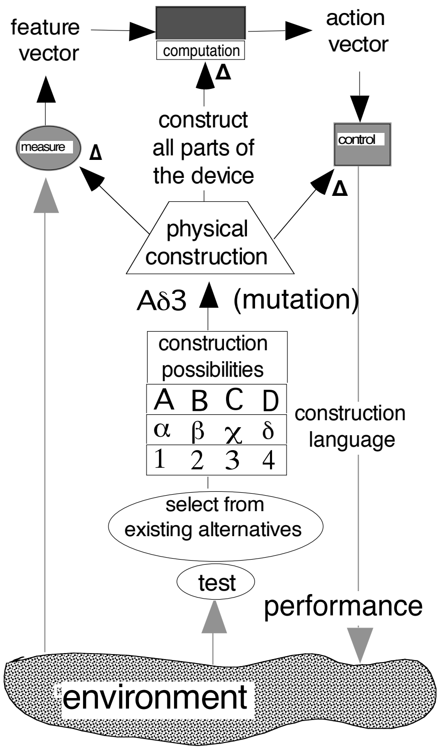

Figure 5. Evolutionary construction of cybernetic devices.

Genetic Construction and Closure

As in biological organisms, adaptive self-construction in these devices can be guided by genetic plans (Fig. 5). In these systems a genetic plan directs the construction of the material hardware of a device. This hardware consists of sensors that implement measurement operations, coordinative parts that implement computational mappings between sensory feature vectors and motor action vectors, effectors that carry out actions on the environment ("control" operations). The construction system also constructs itself and the evaluative sensors that determine which set of construction possibilities is actually realized. Thus, the construction system consists of a set of genetic plans that codes for a pattern grammar of possible material structures that will constitute the hardware that will subserve all the functionalities of the device.

The discrete genetic plans and the analog material hardware of these devices complement each other (Pattee, 1972). This functional organization of symbolic plans that constrain rate-dependent material processes utilizes both the combinatoric possibilities of discrete symbol systems and the open creative possibilities of analog dynamics. The symbolic part is well-defined, steerable, and inheritable but it is bounded by a set of fixed primitives, as was the case with Ashby's homeostat. The analog dynamics of the physical hardware are capable of creating new attractor basins that can subserve new functional states and operations. Pask's analog electrochemical device certainly had a richness of functional possibility, but there were no inheritable plans that could reliably save the process of constructing useful ferrous structures – each assemblage was a one-of-a-kind that had to be grown *de novo*. Genetic plans solve the problem of how to reliably access the rich possibilities inherent in the physical dynamics of matter.

These conceptual examples suggest strategies for open-ended design that involve coupling digital plans with analog dynamics. One needs a physical system that has rich dynamics with a large set of stable accessible states that can subserve useful functions of one sort or another. Means of steering the dynamics such that functional states can sometimes be obtained, are needed. Finally, reliable means of replicating the search for functional states need to be found, and these means themselves need to be controllable through inheritable, symbolic steering mechanisms. Once reliable control/construction structures are in place, then these can in turn be coupled to evaluative mechanisms that can steer the system towards particular goals. Once goals are connected to reliable construction/control processes, then one has an adaptive, self-organizing system. If the system can be made self-replicating, then adaptation can also take place in parallel, amongst populations of systems over many generations. Closing the self-reproduction loop dramatically speeds up the search through genetic and phenotypic spaces. We should note, however, that natural selection itself does not create more variety; it alone does not expand the space of possible genetic sequences or phenotypic structures.

Although spaces of genetic possibilities are well-defined and closed in these systems, spaces of the phenotypic, hardware structures and attendant functions may nevertheless still be open if we have an incomplete description of their environments. In lieu of an exhaustive model of the environment and possible functions within it, phenotypic function spaces are almost always open because of the relational, contextual, environment-dependent nature of functions. If genetically-directed construction occurs independent of external contingencies, then the space of constructed phenotypic structures is closed (1:1 mapping of phenotypes to genotypes). However, if unknown environmental conditions co-modulate the "genetic expression" construction process ("epigenetics"), then the space of possible phenotypes becomes ill-defined and potentially open (>1 phenotype per genotype).

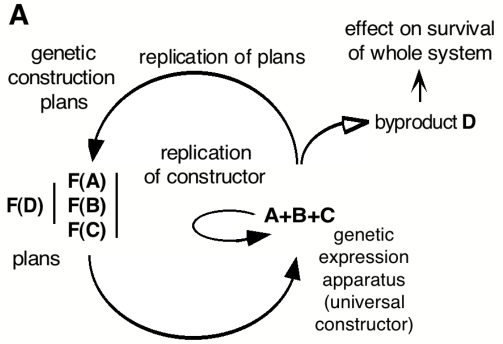

Figure 6. Schematic for von Neumann's kinematic self-reproducing automaton (von Neumann, 1948).

One can ask the analogous question of whether (or in what senses) biological evolution is "open" or "closed." John von Neumann's kinematic self-reproducing scheme (Fig. 6) captures the essence of relations between symbolic, inheritable plans, F(A)…F(D), and material products A…D as well as distinguishing those prodcuts involved in self-construction (A, B, C) from those "byproducts" that are not (D). While the set of possible genetic strings is finite and closed, epigenetic processes can open up somewhat the space of their associated gene-product structures. As with sensors and effectors, the space of intermolecular interactions and possible molecular functions is ill-defined and open-ended, at least until an exhaustive theory of biology is attained. In the meantime, we can reasonably regard biological genomes as closed symbolic realms capable of combinatoric novelty, and biological phenomes as partially-defined, material realms capable of producing both combinatoric and creative novelty in an open-ended fashion.

References

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

Bonner, J. T. (1988). *The Evolution of Complexity*. Princeton: Princeton University Press.

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

Cariani, P. (1992). Emergence and artificial life. In C. G. Langton, C. Taylor, J. D. Farmer & S. Rasmussen (Eds.), *Artificial Life II. Volume X, Santa Fe Institute Studies in the Science of Complexity* (pp. 775-798). Redwood City, CA: Addison-Wesley.

Cariani, P. (1993). To evolve an ear: epistemological implications of Gordon Pask's electrochemical devices. *Systems Research, 10*(3), 19-33.

Cariani, P. (1997). Emergence of new signal-primitives in neural networks. *Intellectica, 1997*(2), 95-143.

Cariani, P. (1998). Towards an evolutionary semiotics: the emergence of new sign-functions in organisms and devices. In G. Van de Vijver, S. Salthe & M. Delpos (Eds.), *Evolutionary Systems* (pp. 359-377). Dordrecht, Holland: Kluwer.

Dawkins, R. (1987). *The Blind Watchmaker*. New York: Norton.

de Latil, P. (1956). *Thinking by Machine*. Boston: Houghton Mifflin.

Fodor, J. (1980). On the impossibility of acquiring "more powerful" structures: fixation of belief and knowledge acqquisition. In M. Piatelli-Palmarini (Ed.), *Language and Learning:The Debate Between Jean Piaget and Noam Chomsky* (pp. 142-162). Cambridge, MA: Harvard.

Morgan, L. (1931). *Emergent Evolution, 3rd Ed.* (Third ed.). New York: Henry Holt.

von Neumann, J. (1948). Re-evaluation of the problems of complicated automata -- problems of hierarchy and evolution. In W. Aspray & A. Burks (Eds.), *Papers of John von Neumann on Computing and Computer Theory (1987)* (pp. 477-490). Cambridge: MIT Press.

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

Pask, G. (1959). Physical analogues to the growth of a concept. In *Mechanization of Thought Processes, Vol II.* (pp. 765-794). London: H.M.S.O.

Pask, G. (1960). The natural history of networks. In M. C. Yovits & S. Cameron (Eds.), *Self-Organzing Systems* (pp. 232-263). New York: Pergamon Press.

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

Pattee, H. H. (1972). Laws and constraints, symbols and languages. In C.H. Waddington (Ed.), *Towards A Theoretical Biology. 4 Essays*. (pp. 248-258). Edinburgh: Edinburgh University Press.

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