TOWARDS AN EVOLUTIONARY SEMIOTICS: THE EMERGENCE OF NEW SIGN-FUNCTIONS IN ORGANISMS AND DEVICES

Peter CARIANI

Eaton Peabody Laboratory, Massachusetts Eye and Ear Infirmary
Boston, USA

Abstract

Signs, symbols, and signals are essential to the survival and evolution of all complex functional organizations that utilize "information." We discuss basic semiotic relations inherent in signalling systems (communication), scientific models (epistemology), adaptive devices (control), and biological organisms (construction). For each of these different functional realms, basic syntactic, semantic, and pragmatic relations are outlined. An evolutionary semiotics seeks to explain how new semiotic relationships can evolve over time. New signalling channels appear in communications systems by the construction of new ways of sending and detecting signals; new observables appear in scientific models through the physical construction of new measuring devices; new "feature primitives" emerge in devices through adaptive construction of sensors. These physical construction and selection processes have analogues in biological evolution. We discuss the crucial role that symbols play in living organizations (the central role of DNA in the self-production of the organism) and in the evolutionary process (inheritability of plans). "Semiotic evolution" and the "semiotics of evolution" thus point us in a common direction, towards a unified, "evolutionary semiotics."

Keywords: Semiotics, biosemiotics, cybernetics, scientific models, adaptive systems, evolutionary robotics, functional emergence, self-reproducing automata, genetic codes

1. Towards a theory of symbols

Signs, symbols, and signals are basic to our existence on many organizational levels, from the biological to the psychological to the social. The ‘semiosphere’, the realm of symbolically-mediated processes, envelopes and incorporates us at every turn (see papers by Hoffmeyer, Umerez, Exteberria, and Joslyn in this volume; Hoffmeyer, 1997). Symbolic nucleotide sequences lie at the root of our biological organizations, neural pulse codes subserve the coherent functional organizations in our brains that permit us to think, while the symbol sequences of our languages afford the complex communications that make human society possible. Semiotic concepts, properly developed, are critical for a deep understanding of the organization of life, the functioning of the brain, and the functional organization of the observer.

This paper presents some of the basic semiotic concepts that are needed in order to understand how symbols function in signalling systems, scientific models, adaptive devices, and biological self-production systems. Using these concepts, we will outline how the symbol-functions in these respective kinds of systems can evolve over time.

The conceptual framework that will be used here incorporates ideas from theoretical biology, cybernetics, general systems theory, information theory, and semiotics. Its foundations come most directly from the work of theoretical biologists Howard Pattee and Robert Rosen on fundamental problems of living organization: what kinds of functional organizations define "life", what roles genetic "codes" play in such organizations, what constitutes biological "information”, and what makes the the emergence of qualitatively new functions possible. Theoretical biology and biological cybernetics have long concerned themselves with the means by which "modelling relations" might be embedded in material systems and how these relations might evolve over time (in other words, how an evolutionary epistemology grounded in the concrete organism or device is possible). Extensive parallels have been drawn between the operational structure of the scientific model, involving measurements, formal computations, and physical constructions, and that of informational processes in biological organisms (Pattee, 1982; Pattee, 1985; Rosen, 1985; Kampis, 1991b; Kampis, 1991a; Rosen, 1991). In my own work in this field, I have strived 1) to analyze the modelling relations in terms of the semiotic framework of Morris (Morris, 1946) 2) to connect these abstract modelling relations with classes of physical adaptive devices that can be built, 3) to outline the inherent epistemic capabilities and limitations of these device classes, 4) to formulate a set of operational definitions sufficient to recognize modelling relations in natural systems, and 5) to apply these operational tests to an observer-based theory of functional emergence (Cariani, 1989; Cariani, 1992a). The result has been to construct a biosemiotic, biocybernetic framework that is grounded in the material organism/device, and centered on the limited observer. The general perspective is not far from the biologically-grounded epistemology of von Uexküll (Uexküll, 1926), and the operationalist, pragmatist philosophies of Mach, Bridgman, Dewey, and Bohr (Murdoch, 1987). It is hoped that this framework can provide additional linkages between theoretical biology and biological cybernetics on one hand and the field of biosemiotics on the other. Biosemiotics has evolved from the study of animal communication to more general considerations of biological codes (see Nöth, 1990; Anderson and Merrell, 1991; Sebeok and Umiker-Sebeok, 1991; Emmeche, 1994; Nöth, 1994; Hoffmeyer, 1997; Uexkull et al, 1993) for entry-points into the field), such that the gaps between these fields appears to be narrowing as time goes on.

2. Functional relations: syntactics, semantics, and pragmatics

It is useful at the outset to define what a sign is and what kinds of relations that it can have with other signs and the world at large. Every sign conveys a distinction (a "difference"), and every sign system utilizes sets of distinctions in order to achieve some purpose. Any object, event, or state-of-affairs can serve as a sign as long as it can be "recognized if it occurs again"[1] and distinguished from other signs. Thus signs are above all functional entities that are bound up in the functional organization of an informational system. In effect, they are the operational states of that system. For the sake of simplicity, we will discussion sign-systems with discrete, well-defined sign-distinctions, leaving iconic, analog processes for other discussions.[2] Such distinctions must incorporate at least two alternative distinguishable sign-states, and in order to be useful, a distinction has to have consequences beyond its mere recognition; in Bateson’s terms, it must be a "difference that makes a difference." While a sign-distinction is physically described in terms of a switch (Pattee, 1973), the description of switching by itself does not capture the sign’s functional role – the manifold effects it has on the world and the purposes it serves.

Signs can engage in several basic kinds of functional, informational transactions, both within the sign-system itself and between the sign system and the external world (Figure 1). A sign can

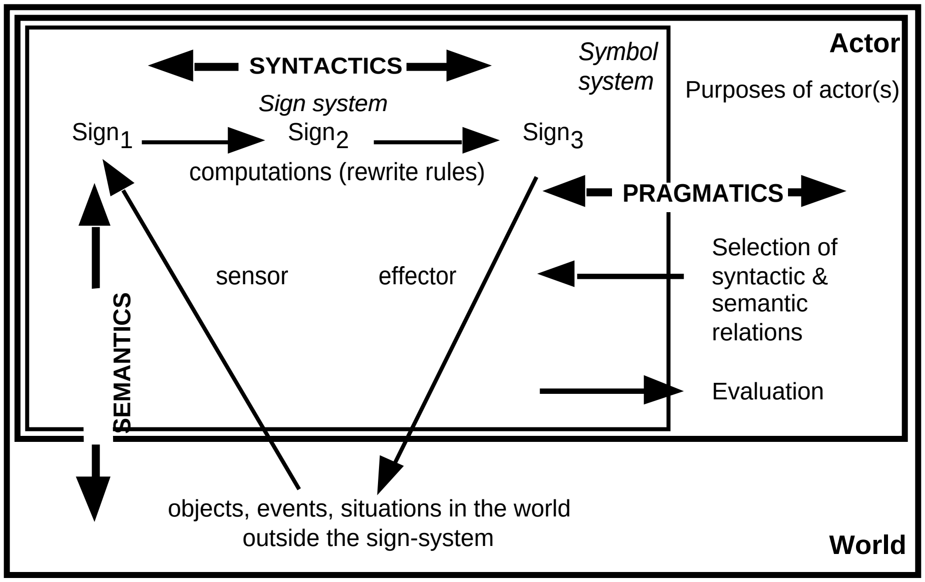

have rule-governed relations to other signs in the system (syntactics), it can have linkages to the world outside the system (semantics), and through its use advance or retard the purposes of those who use it (pragmatics).

Charles Morris first introduced this tripartite set of sign-relations: “...*pragmatics* is that portion of the semiotic which deals with the origin, uses, and effects of signs within the behavior in which they occur; *semantics* deals with the signification of signs in all modes of signifying; *syntactics* deals with combinations of signs without regard to their specific significations or their relation to the behavior in which they occur. When so conceived, pragmatics, semantics, and syntactics, are all interpretable within a behaviorally oriented semiotic, syntactics studying the ways in which signs are combined, semantics studying the signification of signs, and so the interpretant behavior without which there is no signification, pragmatics studying the origin, uses, and effects of signs within the total behavior of the interpretants of signs. The difference does not lie in the presence or absence of behavior but in the sector of behavior under consideration. The full account of signs will involve all three considerations” (Morris, 1946, p. 219; see also Nöth, 1990).

Syntax can be conceived as the set of operationally-deterministic rules which govern the manipulation of signs. A syntactic operation is one in which only the sign's operational state – its*type* – determines the outcome of the procedure (i.e. the manifold material properties of the sign are irrelevant except as they contribute to the recognition of type). Since type is an undifferentiated, unitary property, and the operation's outcome depends only on this property, there can be but one successor state for every predecessor state. All truly syntactic operations thus consist of *rules*, i.e. "operationally determinate" transitions between operational states that depend solely upon the type-identities of other operational states. In effect all physically-realizable syntactic systems are describable in terms of chains of these transitions, i.e. they are formally equivalent to deterministic finite-state automata. All formal procedures and computer programs are purely syntactic operations, wherein the application of the rules governing the manipulation of the signs does not depend upon their "meaning" only upon their symbol type.[3]

Semantics concerns those aspects of signs associated with "meaning," here taken as the linkage between signs on one hand and objects or states of affairs in the world on the other (i.e. "external semantics"). In its broadest sense semantics is the relation of signs to the world beyond the sign-system, encompassing the manifold consequences which flow from their use, both from the world to the sign (measurement, sensing) and from the sign to the world (action). In contrast to the determinate operational structure of syntactic rules (computations), measurement (sensing) has a contingent structure in which an initial ("null", "reset" or "reference") state transits to one or more successor states (Figure 2). Multiple possiblilities are collapsed to one outcome,[4] thereby reducing the uncertainty of the system's subsequent trajectory. The distinction between computation and measurement is therefore that of the (logically) necessary and the (empirically) contingent, between the analytic and the synthetic.

Sometimes the semantics of a sign can be distinguished according to the directionality of the sign-world relation. If interaction with the world determines the sign-state, then the relation is one of "measurement". Similarly, if the directionality flows from the sign to an effect in the world, the relation is one of "action". Syntactic processes in percept-action systems allow particular signs involved in either perception or action to be separated so as to enable their flexible coordination (all possible percept-action mappings). In effect, this decoupling allows for a sign to have a unidirectional semantics, with some signs devoted to measurement (receptor states), others those devoted to action (effector commands), and still others devoted to the coordination of these two semantic sign-types.

Pragmatics is concerned with questions of why particular meanings and syntactic conventions are useful, how they fit into the purposes of symbol-using agents. For human communication this involves what goals or desires an utterance facilitates, why it is useful in a given context. Pragmatics and semantics have often been conflated, but they should not be if purposes are fundamentally distinct from percepts and actions. The linkage of the word "snow" to some particular state of water or weather condition does not tell us why it is useful to have such a linkage (e.g. to communicate whether skiing is possible or flying is advisable), nor does the utility of having a sign for snow determine what that particular sign-world relation will be. While syntactics deals with determinate, rule-governed relations between signs, and semantics deals with contingent, interaction-dependent relations of signs to the world, pragmatics deals with the relations of signs to their intended purposes. Syntactic, semantic, and pragmatic thus are complementary kinds of semiotic relations that a given sign may or may not possess. Uninterpreted formal systems have only syntactics, but they can acquire semantic and pragmatic linkages if their human users provide particular interpretations for their signs, and use them for a particular purpose.

The complementarity of syntactic, semantic, and pragmatic relations means that each provides a different mode of explanation for why things behave as they do. These semiotic relations can be related to Aristotelian modes of explanation, the “four causes” (Graham, 1987; cf. van de Vijver, this volume; Minch, this volume). Behaviors explained in terms of syntactic operations are formal causes, those explained in terms of semantic operations are material causes (because the contingent action of a sensor

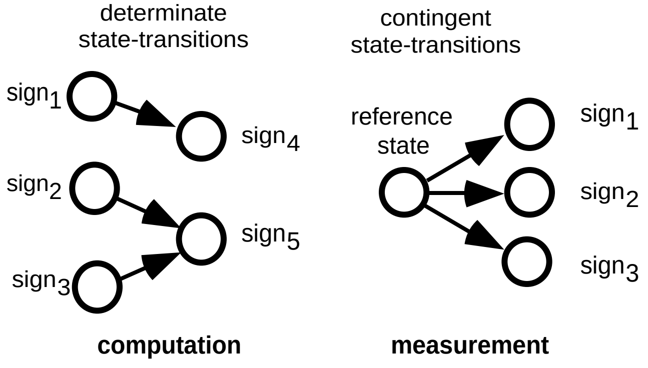

is explainable only in terms of its specific material interactions), and those explained in terms of pragmatic operations are final causes (teleologies). Those explained by the action of a semiotic agent on some other agent or object are efficient causes. Thus constructed, syntactic, semantic, and pragmatic operations involve independent aspects of a sign-system that demand different modes of explanation.

Given the triad of semiotic relations, an "evolutionary semiotics" asks how a "sign" becomes a "symbol", i.e. how new semiotic relations come into being ("semiogeneisis"). This involves both how existing relations are adaptively altered by experience ("learning"), and how new sign-distinctions themselves might arise de novo ("functional emergence").

semantic category	**sign**	**sign states**
	type of intervals in burst	# intervals in burst
city	burst with short intervals (I_1)	1: London 2: Tokyo 3: New York
weather	burst with long intervals (I_2)	1: clear 2: cloudy 3: rain 4: snow

3. The semiotics of signalling

A simple schematic of a signalling system illustrates the semiotic triad (see (Miller, 1951)). A signal-code is shown in Figure 2, while the mappings of the signs and their states to situations in the world, their external semantics, are given in the table. The decoder of such a signal is assumed to be capable of distinguishing inter-pulse intervals, and pulse-bursts are identified as a succession of similar intervals. Intra-burst intervals encode the categories of alternatives while the number of pulse-intervals in a burst encodes the specific alternatives.

The syntax of this system is described by the rules that bursts must alternate between those containing long and short intervals and that each burst of long intervals is preceded by a burst of shorter intervals with which it is associated. (i.e. a "non-grammatical" message would be a string of bursts with only short intervals).

The semantics of the system are that bursts of short intervals (I_1) refer to cities, while bursts of longer intervals (I_2) refer to weather conditions. The pragmatics of this signalling system involve its function to inform travelers of weather conditions in the cities involved. The particular message shown [(3,4)(1,3)] translates as "snow in New York and rain in London". The syntactic structure requires a means of parsing the signal and associating the signs properly. Without syntax, the message would be ambiguous, and it would be necessary to make a distinguishable sign for each of the 12 global message possibilities.

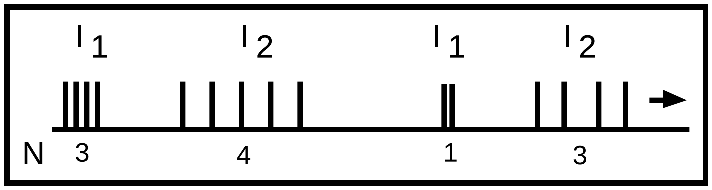

Semantic linkages require reporters (humans or instruments) making observations about the weather in each place. Without semantics, the message would have no empirical connection to what is going on in the world. Here the pragmatics relate to how the use of the sign-system furthers the purposes of those who use it (both senders and receivers of signals). While the syntactics (e.g. signal structure) and semantics (e.g. which cities, which weather conditions) are related to the purposes of the users, by virtue of their design, the syntactic rules and the semantic linkages should not be conflated with the purposes which they subserve (analytic conventions ≠ empirical information ≠ desires).

Traditionally, information theory did not deal with the creation of new signalling channels, but an evolutionary account of this signalling-system could do so. This would involve describing how the particular sign-distinctions (pulses, bursts, intervals) arose (pulse-coded transmitters and receivers were most available), how the particular semantic linkages came to be established (reporters with instruments were stationed at each city), and how the system relates to the purposes of its users (so that travelers could pack appropriate clothing). One could then also explain how new sign-distinctions might arise (perhaps, pulses of different amplitude), how new semantic linkages might be formed (reporters could add thermometers to their repertoire and link their temperature readings to the new distinction of pulse-height), and how the pragmatics might change (the user suddenly wants to predict crop failures).

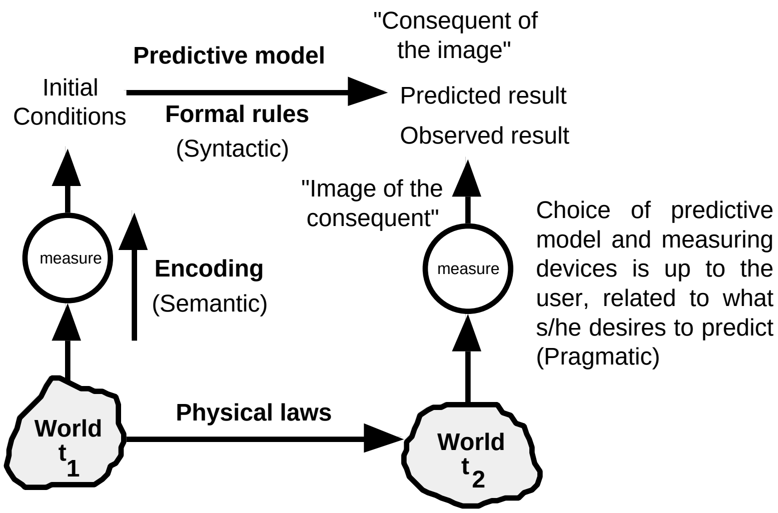

4. The semiotics of scientific models

The semiotic structure of scientific models sheds light on the functional organization of observers and their epistemic relations to the world. Helmholtz, Hertz, and Mach a century ago made explicit the relationships between observations and predictions that constitute the functional correspondences of the symbols in scientific models with events in the world (Cassirer, 1955). In Hertz's scheme, the "commutation diagram" shown in Figure 4, a model must have two parts: 1) a set of measuring devices which, through their interactions with the world, generate a set of signs ("pointer readings") which constitute the "initial conditions" and 2) a mathematical algorithm which takes the initial conditions and generates a prediction. The measuring devices form the bridge between the mathematical part and the undifferentiated world outside the sign-system. Here the model functions not as a mirror of the world, but as a transformation whose detailed structure may not even be completely understood, since the sign-distinctions in the model need not share similarities or even identifiable correspondences with the world (von Helmholtz).

Once the modelling relation is described in operational terms, the semiotic relations inherent in its functional organization become immediately apparent. The formal, mathematical part is described completely in terms of rule-governed syntactic operations on signs, while the measurement part is described in terms of semantic relations between the signs and the physical world. The measuring devices realize the "observables" and the "initial conditions" for the formal model. The model is evaluated by carrying out measurements under two sets of distinguishable situations or experimental arrangements (i.e. "prepared states"), often differing only in time or in space. One of the measurement outcomes is used the initial state of the formal procedure which produces a sign as its output (prediction). The sign generated by the formal part is then compared with the sign generated by the other measurement. If the two always agree, then the model effectively predicts the outcome of the other measurement. To the extent that predictions and measurements diverge, the model becomes less predictive, less reliable in anticipating the outcome of the second measurement, hence less useful to the user. Thus pragmatic relations are described in terms of what predictive goals are achieved with the model.

Scientific models evolve when predictions do not match observations. When the two signs do not correspond, the user has several choices: 1) change the formal part of the model (Δ syntactics),

2) change the observables by altering the measuring devices (Δ semantics), or 3) redefine one's desires/goals (Δ pragmatics). In practice we readily adopt the first solution, less readily the second, and seldom the third (i.e. decide to change what we are trying to predict).

Altering model algorithms and/or observables is an adaptive process, in Spencer's terms "the adjustment of inner to outer relations," i.e. a change in the internal structure that leads to better performance vis-a-vis the external environment. In general, adaptivity requires plasticity of structure coupled with some means of evaluating the performance of the given structure and some means of altering the structure contingent upon that evaluation. When our models don't work properly, we choose to alter algorithms first because it is relatively easy and cheap to perform a different computation, and the set of alternative computations is known to us. Changing the observables, on the other hand, means physically altering the sensors, which can be harder to do and more expensive. Worse still, the set of possible measurements is ill-defined in a way that the set of possible (operationally realizable) computations is not -- we may not understand what physical sensor parameters or measurement conditions need to be altered. In the worst case alteration of model algorithms and observables can be carried out using blind variation, and those configurations that yield better predictions can be selected. This process of structure change, evaluation, and selection is a pragmatic operation, because it is driven by the goals implicit in the evaluative process. Thus the structure of the model evolves through the action of both the observer's mutational-constructional capabilities and the selective pressures brought to bear by the observer's goals.

From this example and through experience, it is obvious that these operations (measurement, computation, evaluation) are independent and complementary. One cannot make measurements and gain empirical information solely by carrying out syntactic operations on tokens (computations); minimally to use computation to predict the state of the world one needs initial conditions, and these must be measured. Similarly one cannot carry reliably out computations using procedures which themselves are contingent upon events in the environment (e.g. if the output of one's computer depended on the amount of car traffic on the streets nearby). By grouping together unreliable elements, one can make more reliable assemblages of elements (a la von Neumann), but then the functional states of the assemblage are no longer contingent on environmental fluctuations. Finally, one cannot create new predictive goals by simply performing computations or by simply making measurements; while an evaluation is a contingent process like a measurement, unlike a measurement, evaluations alter the structure of the modelling relation itself ("feedback to structure") rather than simply changing a particular predictive outcome.

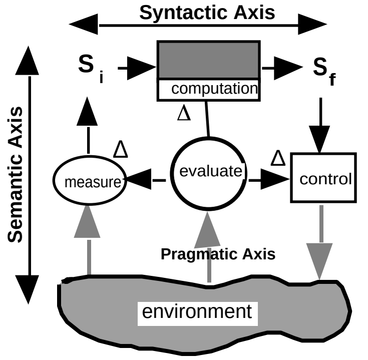

Δ = alter structure → alter function

5. The semiotics of artificial devices

The semiotics of the modelling relation in science can be extended to incorporate organisms and devices that sense the world and act on it contingent upon what they sense. While scientific models are externalizations of the structure of the individual observer, modelling relations are also embedded in the internal structure of individual organisms and devices. Sensory organs and sensors correspond to measuring devices, while coordinative parts of organisms and devices (nervous systems, computational parts) correspond to formal, predictive algorithms. While modelling relations include only measurements and computations, organisms and devices also act directly on the world through effector organs. Effector organs convert signs into action on the material world (labelled "control" in the figures). Thus the basic informational operations of signs (semiotic functionalities) present in organisms and devices can be described in terms of measurement (sensing), computation (coordination), and effecting (action) (Uexküll, 1926; Cariani, 1989; Nöth, 1990). We will touch on several kinds of devices here: computers, robots, trainable machines, and those which construct their own sensors (for fuller discussions see (Cariani, 1989; Cariani, 1992a; Cariani, 1992b; Cariani, 1993)). The semiotic organization of organisms is more difficult to assess and discuss, primarily because the nature of the informational operations performed by nervous systems are presently very poorly understood and may simultaneously involve more than one kind of operation.[5] We are still at the stage where we must make analogies to artifacts whose functioning we understand.

Purely computational device

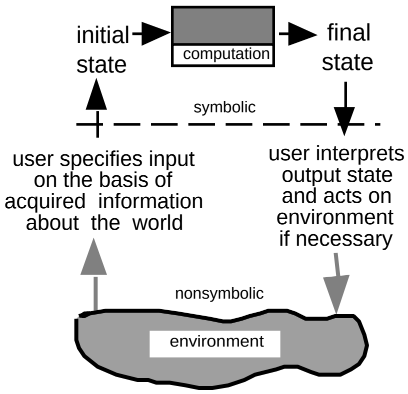

The semiotic of the digital computer, absent its human user and any sensors or actuators, consists solely of syntactic operations (Figure 5). Physically-realized computational devices operate only on signs, and are describable in terms of finite state automata. Thus each total machine state of the device leads to a unique successor state -- what the user does is to load into the machine a program and initial conditions ("data"), and let the machine run from state to state until a terminal state (or sequence) is reached. In order for a computer or formal system to solve a real world problem, the situation must first be encoded into symbolic form by a programmer who has independent access to the real world. Those problems which cannot be effectively encoded into a symbolic notation cannot be solved by computer (Dreyfus, 1979).

Computing devices cannot implement modelling relations because they cannot make measurements. Neither can they act on the world by performing computations; some sorts of effectors are needed. Because they themselves they have no inherent linkages to the world at large, human beings mediate between the world and the computer, choosing the encodings of inputs, measuring their values, interpreting outputs, and taking action. Since they are finite-state devices everything which can occur in the machine is bounded by the finite set of machine states and state-trajectories. Computations within the set of machine states do not create new states, nor can they create new linkages to the world outside the machine. Thus no new syntactic or semantic primitives are created (Carello *et al.*, 1984; Cariani, 1989) and as a result their behavior can be described in a closed notational system ("closed world assumptions").The best that can be done is to combine pre-existing syntactic primitives into logical combinations, and to search the space of combinations for "interesting" syntactic patterns. "Formal times formal is formal." (Kampis, 1991b).

Despite these limitations, pure computation is nevertheless an exceedingly powerful tool. Whatever problems can be effectively encoded, can potentially be solved by purely computational devices, within the physical limits of the device (speed, memory, reliability), the computational demands of the encoded problem, and the foresight of the designer in choosing the appropriate encodings needed.

Robotic devices are computational devices with sensors and/or effectors which connect them directly to the world, giving them inherent external semantics (Figure 6). Robotic devices thus do not need human interpreters to provide meaning for their internal states, or to act directly on the world. By virtue of their sensors and effectors, robotic devices can solve problems and produce

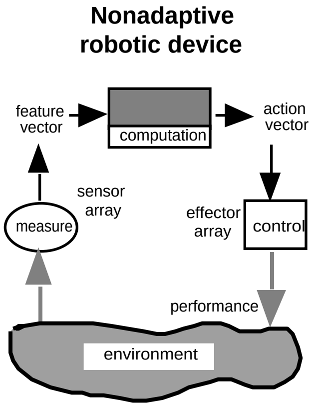

behaviors that are not already encoded in symbolic form, (e.g. recognizing sights and sounds, assembling cars, walking). Whatever, problems can be solved by combinations of sensing, computing, and effecting elements can be solved by the appropriate robotic device. Still, the designer is left the responsibility of finding which sensors, computations, and effectors are suitable for solving a given real world problem. In these nonadaptive devices there is no way of altering the parts, and as a consequence, these devices are limited to the foresight of their designers. If the designer failed to incorporate the sensors needed for a particular task, then the device will not have access to needed sensory information. If the designer failed to foresee the optimal percept-action mappings, then the device will act sub-optimally in some situations. The only solution to this "problem of specification" is to make the device adaptive, to allow it to change its structure contingent upon its performance. The structures and functionalities involved can be either in the syntactic or the semantic realm or both.

Trainable machines are devices that have adaptive computational parts (Figure 7). Altering the computation contingent upon experience is a process of "syntactic-adaptation." Contemporary trainable machines include neural networks, genetic algorithms, adaptive classifiers, Boltzmann machines, and many others. Like the purely computational counterparts, problems must first be symbolically encoded by the designer, but unlike computers, such machines receive (contingent) feedback from their outputs, which then directs the adjustment of their decision function (i.e. the input-output mapping, the computation performed). For example, the designer of a trainable classifier must decide which aspects of the world must be encoded ("feature primitives") such that the machine can find a successful classification rule. The semantics of the states of the adaptive classifier, are therefore supplied by the designer. If the classifier is embedded in a robot, these semantics are determined by the sensors and effectors of the robot. The machine uses feedback from its performance ("supervised learning") in order to steer the search for better classification rules. This feedback is an evaluative process ("test" in the figure) that reflects the goals of the device's designer (i.e. what classification is desired). The consequences of this process are that the internal functional structure of the device changes (i.e. Δ in the input-output mapping, Δ syntactics).

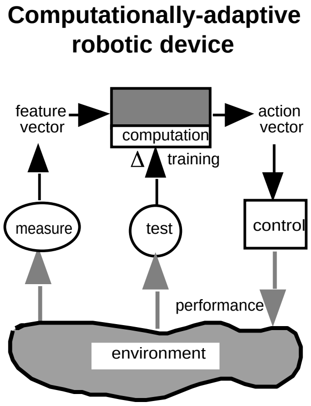

The trainability of the machine means that it can adapt to situations unforeseen by the designer. The designer is freed from having to directly anticipate and specify all appropriate percept-action contingencies. Without the means of altering the sensors and effectors (Δ semantics), trainable machines are ultimately limited by their fixed percepts (feature primitives) and actions. Since these determine the categories in which it operates, the trainable machine can only be as good as its percept and action primitives allow. Since computations alone cannot create new empirical information nor can they create new actions on the world (new semantic linkages), these machines are semantically bounded.

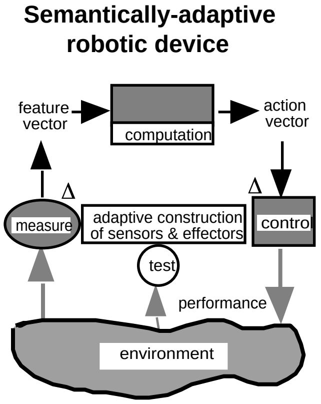

Sensors and effectors are the crucial points at which a real world situation is encoded into a symbolic representation and at which an action-decision is transformed from symbolic representation into physical action. It is only by virtue of actual connection to the world via sensors and effectors that symbolic representations become semantically grounded. Adaptive devices that change their sensors and effectors contingent upon experience are also possible (Figure 8). Semantic categories would then be adaptively selected by altering these parts ("semantic-adaptation"). Such devices would in effect automatically implement what a scientist does when s/he builds a new measuring device or what happens when entirely new sensory and effector organs arise through biological evolution. Within individual organisms, the immune system is a semantically-adaptive process wherein molecular sensors are constructed, evaluated, and selected to recognize foreign antigens.

The various device types outlined above manifest different kinds of emergent behaviors and functions. Semantic adaptation is related to the emergence of new sensory distinctions (e.g. evolving color vision) and new action-alternatives (e.g. evolving capability for flight). Syntactic adaptation is related to emergence of new percept-action mappings ("learning"). In terms of behavior, emergence can be rigorously defined in terms of the deviation of a material system from a model of it (Rosen, 1985); emergent behaviors are those which defy a given model of the system. A given class of devices with emergent functions can then be specified in terms of what an observer has to do to "track" (predict) the device's behavior as it changes its structure and its relations to the world. If the observer must change syntactic state-transitions in his/her model to track the device's behavior, then the device is "syntactically-emergent"; if new observables are needed, then the device is "semantically-emergent". Whenever the device adds a new independent sensory distinction, a new observable is required of the observer, and the dimensionality of the apparent behavior of the device increases (Conrad, 1979; Chen and Conrad, 1994; Conrad, this volume). Trainable machines such as neural nets are thus syntactically emergent, while those that construct their own sensors are semantically-emergent (Cariani, 1989; Cariani, 1992a; Cariani, 1993).

When an organism or device has the capacity to determine its own syntactic and semantic relations we have a situation of semantic closure (Pattee, 1982; Pattee, 1985; Rocha, 1996; Umerez, this volume;), of organizational closure (Maturana, 1981; Pask, 1981; Rosen, 1985; Minch, this volume) wherein a self-modifying system (Kampis, 1991b) attains a degree of epistemic autonomy (Cariani, 1989; Cariani, 1992a) or semiotic freedom (Hoffmeyer, this volume): it chooses its semantic categories as well as its computational states within those categories (semantic and syntactic autonomy). If the device is capable of changing its pragmatic relations by redefining its evaluative criteria (goals, desires), then it attains a degree motivational or pragmatic autonomy (freedom to change its own desires).

6. Evolutionary semiotics and the semiotics of evolution: symbolically-steered construction

Critical to the adaptive process is the modification of the "hardware" that subserves syntactic and semantic functions. In order to alter semantic functions of sensing and effecting, sensors and effectors must be physically altered; in order to augment the mapping capabilities of the computational part, more sign-distinctions must be enabled (e.g. adding RAM to a computer), and this must occur by physically constructing more accessible states (one cannot not create new RAM by running computations, or our computers could expand by continuously "growing themselves"). Optimally the means by which these physical substrates of semantic and syntactic are altered should result in stable alternatives (changes persist, making learning possible) and inheritable structures (new generations can build on old experience). Arguably, symbolic control of the physical construction process is the most efficient means of attaining these ends (Pattee, 1982; Pattee, 1985), and all biological organisms now in existence have by one means or another exploited this mode of functional organization. The cycle (Figure 9) is one of symbol string → construction process → formation of physical parts of the device → action of the device → performance in the environment → differential survival → selection of symbol strings (plans). Within the self-construction process are mutational mechanisms for generating variability in the symbol strings, which permits an expanding the portion of the space of symbol strings to be tested.

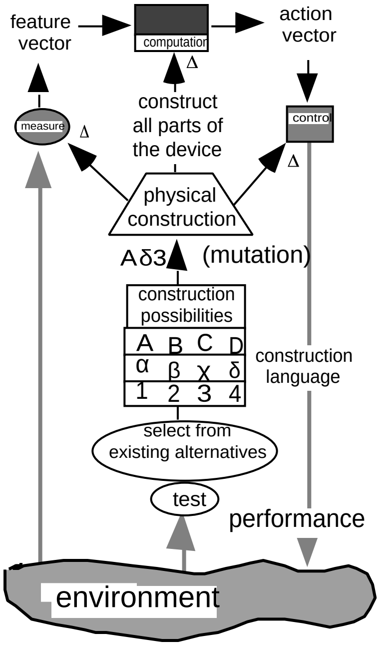

Thus, the most striking aspect of symbols in biological organisms is their central role in self-production, in reproduction and in evolution. Virtually all physical systems that we intuitively recognize as living organisms engage in DNA-directed construction processes. *We cannot understand the organization of living things without understanding the role symbols play in biological organization.*[6] The converse is also true, that virtually all symbols are associated with biological organisms, whether for communication, coordination of action, or construction, and whether at a cellular, organismic or social level. *We cannot understand symbols fully until we understand their role in the organization of life.*

One can outline the semiotic aspects of biological construction languages (Figure 10). The sign-distinctions are DNA sequences, and their most obvious syntactic relations involve transcription-translation rules, the mapping of a DNA nucleotide sequence into an amino acid sequence of a protein.

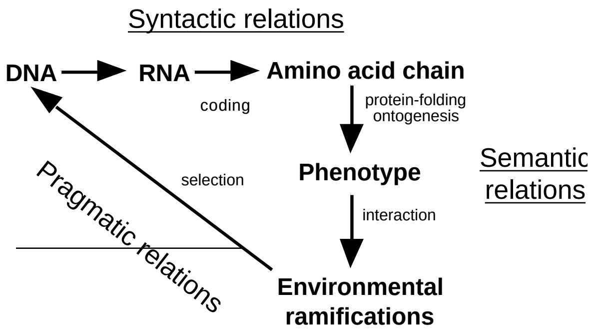

The semantics of a DNA sequence involve the genotype-phenotype relation, i.e. the relation of the nucleotide sequence to the three dimensional protein structure it codes for and all of the manifold consequences of this folded protein on the rest of the cell, the organism, and the world beyond. While the transcription process can be explicitly described in terms of a syntax of discrete nucleotide and amino acid sequences, the semantics of these sequences are bound up the ill-defined analog dynamics of protein folding, enzymatic action, and still more distal consequences. Analog and digital aspects of proteins thus coexist and complement each other (Pattee, 1979, 1982) in what has been termed a “code-duality” (Hoffmeyer, this volume).

The pragmatics of genetic sequences involve the manifold consequences of gene products on the well-being, survival, and reproduction (i.e. the evolution-sculpted "goals" of the organism). These three categories, syntactic, semantic, and pragmatic, are what biologists usually call the genetic, the phenotypic, and the selective aspects of life, and all are essential requirements for evolution. In effect, these categories define the organism at the molecular genetic level as a semiotic entity.

One might ask why living organisms need be semiotic organizations at all. After all, there are self-organizing metabolic systems that are capable of a large range of life-like behaviors (see papers by Minch and Joslyn, this volume). Are genetic codes critical for the stability and evolution of complex biological organizations? While we lack definitive answer to this question, some general functional roles for symbols in cells can be postulated (Cariani, 1989). Some possible advantages of utilizing discrete codes over purely analog dynamics involve the enhancement of internal stability (organizational anchors), greater resistance to perturbation and greater amenability to error repair, provision of memory through gene-switching, more reliable inheritance of acquired information, and the reliable construction of specific, complex amino acid sequences. Discrete genetic coding permits nearly identical copies to be reliably made and inherited over a large number of successive generations, an accomplishment that would be very difficult in an analog system.[7] DNA coding permits an inordinately rich variety of structure by permitting arbitrary (and improbable) amino acid sequences to be reliably manufactured and replicated.

Do symbols enhance evolvability? Perhaps the strongest argument for such a role is that symbols permit universal construction and generalized encoding (Cariani, 1989; Rocha, 1996). This allows for any genetic sequence to code for a product, regardless of whether it is already involved in an autocatalytic loop. Once an apparatus is in place for replicating genetic sequences and the replication apparatus itself, all other sequences can "piggyback" on those which code for the self-reproduction apparatus. This means that a protein need not be connected into its own special, autocatalytic network in order to be replicated in the next generation.

Von Neumann recognized this ability of a self-reproducing apparatus to produce other, non-self-replicating "byproducts" (Figure 11): “...Let X be A+B+C+D, where D is any automaton [and A+B+C is a self-replicating combination of automata]. Then (A+B+C) + F(A+B+C+D) produces (A+B+C+D) + F(A+B+C+D) [the object itself is in parentheses, e.g. (X), and its construction description as a function, e.g. F(X)]. In other words, our constructing automaton is now of such a nature that in its normal operation it produces an object D as well as making a copy of itself. This is a normal function of an auto-reproductive organism: it creates byproducts in addition to

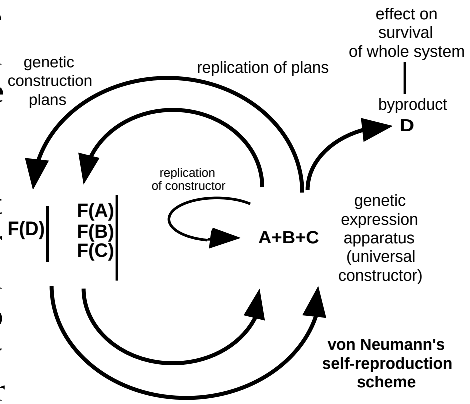

reproducing itself" (Neumann, 1987, p.489). In terms of the genetic code, A+B+C constitute all of the apparatus needed for gene expression, while F(A), F(B), and F(C) are the genetic "plans" which are read by the expression apparatus. Once such a generalized "universal constructor" is in place, new proteins need not be self-reproducing in order to persist long enough to have useful functions for the organism. Their DNA is translated and transcribed along with the generalized machinery for transcription and translation. The ability to enhance stability and survival of the organism and the ability to self-replicate are thus separated for genes and their associated proteins. Coupled with the indefinitely rich possibilities of arbitrary peptide polymers and their combinations, this separation allows DNA sequences to become a common "informational currency" of the organism.

Without generalized encoding, evolution proceeds much more slowly, since every component must both catalyze its own formation and provide some other survival benefit. The alternative to generalized encoding is an autocatalytic ("autopoietic") network in which each component is involved in the production of some other network component, such that all components are produced. Once the requirement that each protein be linked into an autocatalysis network is removed, the number of proteins with useful functions which are propagated to subsequent generations increases enormously. The enormous stabilizing effects and evolutionary advantages conferred by symbols may explain why there are very few natural, complex self-production networks which are not tied in some way to genetic construction. Semiosis, through symbolically-directed construction, is arguably the functional organization that permits complex self-producing biological entities to persist. These complex functional organizations then permit networks of self-producing signals that form the organizational substrates for self-directing epistemic agents such as ourselves.

7. Conclusions

We have discussed basic syntactic, semantic, and pragmatic relations of symbols in signalling, systems (communication), modelling relations (epistemology), adaptive systems (self-adjustment), and constructive systems (self-production). The semiotic relation between symbol and matter forms the basis of the functional organization for both life and information-use. We have outlined basic mechanisms for how new syntactic, semantic, and pragmatic relations can come to be incorporated into organisms and devices, how an "evolutionary semiotics" that spans both the natural and the artificial might be possible.

Index terms

Semiotics, evolution, emergence, syntactics, semantics, pragmatics, adaptive systems, epistemology, genetic code, autopoiesis, self-production systems, self-reproducing systems, scientific models, measurements, computations, biosemiotics, operationalism, autonomy, sensors, effectors, adaptive systems, adaptation, observable, construction, [proper names of references]

Acknowledgments

I would like to acknowledge a profound intellectual debt to Howard Pattee, who brought me into the realm of symbols and their functions.

References

Anderson, Myrdene, and Floyd Merrell. 1991. *On Semiotic Modelling.* Berlin: Mouton de Gruyter.

Ashby, W. Ross. 1956. *An Introduction to Cybernetics.* London: Chapman and Hall.

Carello, Claudia, M. T. Turvey, Peter N. Kugler, and Robert E. Shaw. 1984. Inadequacies of the computational metaphor. In *Handbook of Cognitive Neuroscience*. Edited by M. Gazzaniga. New York: Plenum Press.

Cariani, P., 1995. As if time really mattered: temporal strategies for neural coding of sensory information. *Communication and Cognition - Artificial Intelligence (CC-AI)* 12(1-2): 161-229. Reprinted in: *Origins: Brain and Self-Organization*, (K. Pribram, ed.), Lawrence Erlbaum, Hillsdale, NJ, 1995.

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

Cariani, Peter. 1992b. Some epistemological implications of devices which construct their own sensors and effectors. In *Towards a Practice of Autonomous Systems*. Edited by F. Varela and P. Bourgine. 484-493. Cambridge, MA: MIT Press.

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

Cariani, Peter A. 1989. *On the Design of Devices with Emergent Semantic Functions.* Ph.D., State University of New York at Binghamton.

Cassirer, Ernst. 1955. *The Philosophy of Symbolic Forms, Vols 1-3.* Translated by Manheim, R. New Haven: Yale University Press.

Chen, Jong-Chen, and Michael Conrad. 1994. A multilevel neuromolecular architecture that uses the extradimensional bypass principle to facilitate evolutionary learning. *Physica D* 75 : 417-437.

Conrad, Michael. 1979. Bootstrapping on the adaptive landscape. *Biosystems* 11 : 167-182.

Conrad, Michael. 1997. Towards high evolvability dynamics. This volume.

Dreyfus, Hubert L. 1979. *What Computers Can't Do.* New York: Harper & Row.

Emmeche, Claus. 1994. *The Garden in the Machine.* Princeton: Princeton University.

Exteberria, Arantza. Embodiment of natural and artificial agents. This volume.

Hoffmeyer, Jesper. 1996. *Signs of Meaning in the Universe*. Bloomington: Indiana University Press.

Hoffmeyer, Jesper. The unfolding semiosphere. This volume.

Graham, W.H. 1987. *Aristotle's Two Systems.* Oxford: Oxford University Press.

Kampis, George. 1991a. Emergent computations, life, and cognition. *World Futures* 32 (2-3) : 95-110.

Kampis, George. 1991b. *Self-Modifying Systems in Biology and Cognitive Science.* Oxford: Pergamon Press.

Maturana, Humberto R. 1981. Autopoiesis. In *Autopoiesis: A Theory of the Living*. Edited by M. Zeleny. New York: North Holland.

Miller, George A. 1951. *Language and Communication.* New York: McGraw-Hill.

Minch, Eric. The beginning of the end: on the origin of final cause. This volume.

Morris, Charles. 1946. *Signs, Language, and Behavior.* New York: George Braziller.

Murdoch, Dugald. 1987. *Niels Bohr's Philosophy of Physics.* Cambridge: Cambridge University Press.

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

Nöth, Winfried. 1990. *Handbook of Semiotics.* Indianapolis: Indiana University Press.

Nöth, Winfried. 1994. *Origins of Semiosis: Sign Evolution in Nature and Culture.* Berlin: Mouton de Gruyter.

Pask, Gordon. 1981. Organizational closure of potentially conscious systems. In *Autopoiesis: A Theory of Living Organization*. Edited by M. Zeleny. 265-308. New York: North Holland.

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

Pattee, Howard H. 1973. Discrete and continuous processes in computers and brains. In *The Physics and Mathematics of the Nervous System*. Edited by W. Guttinger and M. Conrad. New York: Springer-Verlag.

Pattee, H. H. 1979. The complementarity principle and the origin of macromolecular information. *Biosystems* 11 : 217-226.

Pattee, Howard 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 *Complexity, Language, and Life: Mathematical Approaches*. Edited by J. L. Casti and A. Karlqvist. 268-281. Berlin: Springer-Verlag.

Rocha, Luis. 1996. Eigenbehavior and symbols. Systems Research 12 (3): 371-384.

Rosen, Robert. 1985. *Anticipatory Systems.* Oxford: Pergamon Press.

Rosen, Robert. 1991. *Life Itself.* New York: Columbia University Press.

Sebeok, Thomas A., and Jean Umiker-Sebeok eds. 1991. *Biosemiotics: The Semiotic Web 1991.* Berlin: Mouton de Gruyter.

von Uexküll, J. von. 1926. *Theoretical Biology.* New York: Harcourt, Brace & Co

von Uexküll, Thure, Werner Geigges, and Jörg M. Hermann. 1993. Endosemiosis. Semiotica 96-1/2: 5-51.

Umerez, Jon. The evolution of the symbolic domain in living systems and artificial life. This volume.

van de Vijver, Gertrudis. Evolutionary systems and the four causes: a real Aristotelian story? This volume.

---

[1] cf. Ross Ashby's operational notion of 'state', (Ashby, 1956).

[2] This is not to say that analog, iconic, or continuous signals are unimportant, only that the definition of what constitutes a distinction in such systems is a much more involved process. Most of the neural codes that have been investigated in the brain have a graded, analog character to them, such that the brain is probably best seen in terms of a mixed digital-analog device (with heavy emphasis on the analog part). To the extent that discharges of primary sensory neurons follow the time structure of their respective stimuli, there exists an iconic representation of the stimulus in the time intervals between spikes (Cariani, 1995).

[3] This is close to Hilbert's original, operational notion of a formal system, wherein communities of observers manipulate sensuously apprehensible tokens to reach convergent conclusions. When unlimited string lengths, tapes, and stacks, are introduced in order to represent the infinite set of "natural numbers," such systems no longer are fully-surveyable by finite, physically-realizable devices. From a strict-finitist perspective, which insists on physical realizability, all of the

---

Gödelian paradoxes are artifacts of the introduction of ill-defined potential-infinities into formal systems. In this view formal procedures must be implementable by observer-participants possessing limited numbers of states and computational capacities. All that can be physically computed is describable in terms of finite-state automata, of some given size. We have access to a given number N as an individuated entity if we have the means to reliably implement N-1 sign-distinctions or N distinguishable tokens. The set of numbers that we can reliably manipulate, the "computational envelope," expands as we acquire the physical means of building larger, faster, more reliable machines (i.e. a “Chuck Yaeger theory of number” is possible).

[4] The "collapse of the wave function" in quantum mechanics, at the point when a measurement is made (the point of observational contingency) is only problematic for realist and reductionist accounts; not for those that take the semiotic operations of the observer as primitive (e.g. Bridgman, Bohr, von Neumann).

[5] See Exteberria’s discussion of the functional decomposition of cognitive systems in this volume.

[6] To quote from Howard Pattee: “The question is whether information is to be treated in biology as just another physical variable, or as the characteristic and exclusive aspect of living systems and their artifacts that distinguish living systems from all other physical systems. ((Pattee, 1979, p.218)” and “Life depends upon records.” (Pattee, 1972).

[7] One difficulty with analog processes is that they tend to accumulate small errors. Digital systems have sets of attractor basins that constantly reduce the effects of small perturbations.