Some epistemological implications of devices which construct their own sensors and effectors

Peter Cariani

37 Paul Gore St, Boston, MA 02130.
eplunix!peter@eddie.mit.edu

Abstract

Various classes of physical devices having adaptive sensors, coordinative parts, and/or effectors are considered with respect to the kinds of informational relations they permit the device to have with its environment. Devices which can evolve their own physical hardware can expand their repertoires of measurements, computations, and controls in a manner analogous to the structural evolution of sensory, coordinative, and effector organs over phylogeny. In particular, those devices which have the capacity to adaptively construct new sensors and effectors gain the ability to modify the relationship between their internal states and the world at large. Such devices in effect adaptively create their own (semantic) categories rather than having them explicitly specified by an external designer. An electrochemical device built in the 1950's which evolved the capacity to sense sound is discussed as a rudimentary exemplar of a class of adaptive, sensor-evolving devices. Such devices could potentially serve as semantically-adaptive front-ends for computationally-adaptive classifiers, by altering the feature primitives (primitive categories) that the classifier operates with. Networks composed of elements capable of evolving new sensors and effectors could evolve new channels for inter-element signalling by creating new effector-sensor combinations between elements. The new channels could be formed orthogonal to pre-existing ones, in effect increasing the dimensionality of the signal space. Such variable-dimension signalling networks might potentially lead to more flexible modes of information processing and storage. Devices having the means to both choose their sensors (primitive perceptual categories) and effectors (primitive action categories) as well as the coordinative mappings between the two sets would acquire a degree of epistemic autonomy not yet found in contemporary devices.

Modelling relations in biological organisms and artificial devices

A very ancient and biological notion is that the sensory organs of an organism determine the basic categories through which it perceives the world and that its effector organs determine the basic categories through which an organism can act to alter its environment. Between sensors and effectors lies some mechanism for their coordination, thereby permitting the organism to respond appropriately to some perceived state of the world. Such a framework can also be applied to analyzing functional relations embedded in artificial devices. One can talk about the informational relationships between organisms and devices and their environments in terms of the basic functionalities of *sensing*, *coordinating*, and *effecting* (or, more abstractly in terms of *measurement, computation*, and *control* (Cariani 1989)). Taken together these functionalities allow an organism or device to sense some perceived state of the world, to make a prediction for what is the best action to take given that perception, and to act accordingly. Given a set of sensors, a set of effectors, and a mechanism for coordinating between the two, such *modelling relations* can be embedded in an organism or artificial device (Uexküll 1926; Pattee 1982; Rosen 1985; Cariani 1989, 1991ab; Kampis 1991; Emmeche 1991).

Within this functional, informational framework (Figure 1), the sensory organs determine the relation between the world at large and those internal states which are affected by sensory inputs. A sensory organ makes primitive distinctions on the world by producing two or more possible outcomes, depending upon its interaction with the external world. For example, a receptor may signal the presence or absence of some chemical substance (relative to some threshold concentration). There must be at least two

Cariani P. Some epistemological implications of devices which construct their own sensors and effectors. In: Varela F, Bourgine P eds. Towards a Practice of Autonomous Systems. Cambridge, MA: MIT Press, 1991;484-493.

states possible for the receptor to function as a discriminatory element -- if the receptor were always in one state, no information regarding the external world would be conveyed, the organism could not act contingent upon some perceived change in the world. The distinction or discrimination made forms a perceptual *category* having two or more possible, mutually-exclusive *states*. To the extent that different receptors interact with the world differently, so as to produce different outputs in the same situation, those receptors implement different categories, they function as different windows on the world outside the organism.

Effector organs also determine the relation of internal states and the world at large, but the directionality of immediate causation is different. While the output state of a sensory element is contingent upon some state of affairs in the external world, the external world in some way is contingent upon the input state of the effector organ. For example, the electric field strength in the vicinity of an electric eel is highly contingent upon the internal states of the eel's nervous system; whenever the eel decides to fire, a large change in the field will follow. Similarly, when a set of muscles is activated for moving, the relative position of the organism relative to the rest of the world is altered -- the positional configuration of objects in the world becomes contingent upon inputs to the organism's motor system.

Sensors and effectors thus mediate between the internal states of the organism and the world outside of it. Together they determine the external semantics of the internal states of the organism, the immediate consequences of world-states on organism-states and vice-versa. Here, the word "immediate" is stressed because the causality between sensing, coordinating, and acting is a circular one -- an action has an impact upon subsequent perceptions which influence subsequent actions, and so on (Uexküll 1926 and the "circular causal relations" in the cybernetics of Wiener 1961 and McCulloch 1989). An animal moves its head in order to better see something moving in its visual periphery which causes yet other objects to come into view, prompting another movement and more perceptions. An entire theory of adaptive control has been developed with this circular organization as its organizing principle, that effective regulators adaptively adjust their behavior so as to bring about a desired perception (Powers 1973, 1979).

Coordinative organs mediate between the output states of sensory organs and the input states of effector organs, in effect providing a mapping between particular perceptual outcomes and particular action alternatives. The more reliable the coordination, the more this mapping resembles a function (i.e. for a given distinct perceptual state, one particular action will always be taken). When this coordination is completely reliable, it is possible to describe the state-transitions in terms of a "state-determined system" (Ashby 1954). Such (finite) state-determined systems can be described completely in terms of syntactic, formal operations, and it is in these circumstances that one is justified in characterizing what is going on in the natural system in terms of "computations". In the ideal case of a completely reliable coordinating organ (as might be found in a contemporary robot with a digital electronic computer taking inputs from sensors and activating effectors), coordination implements a purely syntactic relation within the set of internal states whilst sensing and effecting organs implement purely semantic relations with the world at large. Here the coordinative rules must operate within the set of perceptual categories given by sensors and those behavioral categories determined by effectors.

The evolution of sensors, coordinators, and effectors

In biological evolution all three functionalities evolve over time. Receptor organs coevolve with coordinative structures and effector organs to produce organisms which can make the sensory distinctions, coordinations, and actions needed for survival and reproduction. Sensors evolve to make those distinctions which enhance the survival of their bearers; those sensors which make discriminations important for survival and reproduction will over generations prevail within a population. Effectors similarly evolve to enable those actions which enhance survival and hence reproduction. Those body configurations of bones, muscles, and secretory organs which are particularly effective in getting food, evading predation, eluding disease, and reproducing will tend to prevail over time in the population. Coordinative organs (in higher organisms this is usually the nervous system) evolve to more effectively map the perceived state of the world to an appropriate action. Those individuals which can coordinate percepts with possible actions more effectively can make better use of the information provided by the sensory array, and will tend to survive longer to have more offspring.

Phylogenetic vs. ontogenetic learning

By constructing their own sensors, effectors, and coordinative organs biological species alter their relationships to the world around them. Over phylogenetic time periods the categories through which organisms apprehend the world and act on it evolve with the changing needs of the population. If learning is taken in its broadest sense as the alteration of internal structures so as to better perform in an external environment, then this biological evolution of sensing, coordinating, and effecting organs represents a kind of phylogenetic learning in which the ability to model the environment, and consequently survival and reproduction, is improved.

> Two of the phenomena which we consider to be characteristic of living systems are the power to learn and the power to reproduce themselves. These properties, different as they appear, are intimately related to one another. An animal that learns is one which is capable of being transformed by its past environment into a different being and is therefore adjustable to its environment within its individual lifetime. An animal which multiplies is able to create other animals in its own likeness, or at least approximately, although not so completely in its own likeness that they cannot vary in the course of time. If this variation itself is inheritable, we have the raw material on which natural selection can work. If hereditary invariability concerns manners of behavior, then among the varied patterns of behavior which are propagated some will be found advantageous to the continuing existence of the race and will establish themselves, while others which are detrimental to this continuing existence will be eliminated. The result is a certain sort of racial or phylogenetic learning, as contrasted with the ontogenetic learning of the individual. Both ontogenetic and phylogenetic learning are modes by which the animal can adjust itself to its environment. (Wiener, 1961)

Usually phylogenetic learning is thought to involve the optimization of the kinds of structures which constitute the organism, while ontogenetic learning is an optimization within those (relatively invariant) sets of structures and plasticities given the individual organism as the product of phylogeny. Here the basic semantic categories of the organism (percept, action repertoires) are usually thought to be fixed by phylogenetic learning while the particular mappings within those categories (percept-action coordinations) are thought to lie in the domain of ontogenetic learning. While this may be valid as a first cut, there may be circumstances (particularly in higher organisms) where new sensory and behavioral categories are formed within the lifetime of the individual organism. The ontogenetic, performance-contingent selection of antibodies in the immune system is an obvious example of this: an array of molecular sensors is evolved through differential reproduction of those cells producing antibodies that recognize foreign molecules. The set of distinctions (the equivalence classes of antigens which can be recognized) that the immune system can make on the world around it evolves over time. Here the functional interrelationships mediating between evaluation of particular sensors (how strongly a given antibody binds to an antigen, thereby "recognizing" it) and structural adjustment are similar to those found in morphological evolution. Less clear-cut is the question of how well this description might apply to the formation of new conceptual structures in the brain. Here one would look toward adaptive alterations in the ways in which continuous, analog neural discharge patterns (e.g. interspike interval distributions) are transformed into discrete global modes of representation and action. While the particular mechanisms of evaluation, inheritance and construction of structural alternatives may turn out to be very different for individual nervous systems and evolving populations of organisms, the abstract functional organization of adaptation can still, nevertheless, be qualitatively similar. In all cases this organization consists of a cycle of performance evaluations, modification of a constructive plan (through blind variation or directed alteration), and the subsequent construction of new structures to enable new performances.

Optimization within categories and optimization of categories

From the standpoint of analyzing what general kinds of functional adaptations are possible, it is useful to distinguish between the optimization of the percept-action categories themselves and the optimization of particular mappings between states formed by the categories. The corresponding processes in the optimization of a scientific model would involve the selection and optimization of appropriate measuring devices and observables vs. the optimization of predictions made given a particular set of observables. One is a semantic search process, finding the appropriate relations between the world and the state variables of the model (the right measurements to make), while the other is a syntactic search process, finding the algorithm which generates the best predictions given those measured initial states. These two kinds of optimization, of the categories themselves vs. between states within fixed categories, represent two kinds of acquired knowledge.

Each kind of knowledge entails a different kind of learning; each corresponds to a possible kind of adaptive device that we could potentially build (Figure 2). If the sensors of our devices correspond to sensory organs of organisms, computational parts correspond to biological coordinative organs (nervous systems), and effectors correspond to the various biological effector organs, then we can envision all of the types of device adaptivity which would correspond to the biological evolution of these various organ types.

The optimization of categories would correspond to the structural evolution of sensors, computational networks, and effectors. The artificial analogy of the biological evolution of sense organs would be devices which evolved (or adaptively constructed) their own sensors to make the primitive distinctions necessary to solve a problem. The artificial analogy of the biological evolution of effector organs would be devices which evolved (or adaptively constructed) their own effectors appropriate to some given task. The artificial analogy of the evolution of coordinating systems would be a device which evolved the material structures subserving computations: new memory states and/or new processing element types.

Optimization within categories would correspond to flexible switching between the set of states given by structural evolution. One could think of structural evolution (optimization of the categories) in terms of "hardware" design, construction, and augmentation while this latter kind of optimization (within pre-established categories) is performed by "software" algorithms operating within the hardware constraints. While we now have adaptive devices which can optimize the particular computations they perform (trainable machines, neural nets, genetic algorithms), we do not as yet have devices which evolve their own hardware – their sensors, their computational hardware, their effectors. And while much attention has been paid to adaptation in the computational realm, exceedingly little has been paid to adaptation in the sensing and effecting realms.

The evolution of new perceptual distinctions and action alternatives

How can a new sensory distinction or a new possible action evolve? In general it will entail the adaptive construction of a new material structure which interacts with the device's environment in a way that is different from existing sensors and effectors. If the structure enhances the performance of the device, then it is stabilized and incorporated into the device's constructive specifications.

This new structure could be a part of the device proper or it could take the form of a prosthesis, modifying the relationship(s) between existing sensors and effectors and the world beyond the prosthesis. The immune system was mentioned above as a natural example of internal structural evolution. Most of the examples of prosthesis are artificial. We ourselves build tools (chain saws, steam shovels, chemical plants, bicycles) which, when directed by our biological effectors, radically alter our relationship to the world around us. We build sensory prosthetics (cat-scanners, telescopes, microscopes, molecular assays, hearing aids) which allow us to amplify the number of distinctions we can make on the world. We build coordinative prosthetics (computers, pencil-and-paper, written symbol strings) that allow us to amplify the complexity of our coordinations by extending our memory and our ability to effectively manipulate large numbers of symbols.

In all of the artificial examples cited above, however, human beings are part of the loop-- human beings construct the new types of sensors, the larger computers, the more powerful effectors. On the whole we have very few devices which carry out these adaptive constructions on their own. If we look around us at the kinds of robotic devices available, the vast majority of them have fixed sensors, effectors, and coordinative programs. A small minority of them are adaptive robots having trainable computational parts mediating between fixed sensor and effector arrays (via neural nets, genetic algorithms, Bayesian methods, or other training algorithms). The devices which come closest to adaptively evolving their own sensors and effectors are those which tune the parameters controlling their sensors (as in Carver Mead's analog-VLSI devices (1989) or as in autofocusing mechanisms for cameras), but these parameters are designed explicitly into these devices -- the device itself does not have the structural autonomy necessary to evolve its own categories.

Artificial evolution of a primitive ear

In the late 1950's, the cybernetician Gordon Pask (Pask 1958, 1959, 1960) constructed an electrochemical assemblage to demonstrate that a machine could in fact evolve its own sensors. The device consisted of an array of electrodes immersed in a ferrous sulfate/sulphuric acid solution. By passing current through the array, iron dendritic thread structures could be grown between selected pairs of electrodes. By adaptively allocating current to the electrodes (by changing their relative share of current), the growth of these threads between the electrodes could be steered; structures could be constructed and rewarded on the basis of how they interacted with the world at large. All of this could be accomplished without a specific physical theory of how the threads form and extend themselves, much in the same way the "blind" search for effective sensory substrates operates in biological evolution. One could implement a contemporary neural network this way -- the inter-element weights of the neural net would be proportional to the inter-electrode conductances. Instead Pask adaptively steered the assemblage to construct structures sensitive to sound. If particular configurations arose which changed their resistive properties in response to sound, then they were rewarded. In about half a day they were able to train the device to discriminate between presence and absence of sound, and then between two frequencies of sound. The device represents the artificial evolution of a new sense modality, the emergence of sensory distinctions which were not present in the device at the start. The structural adaptivity of the device enables the emergence of a new function. This relation between adaptivity and emergence is discussed in more depth elsewhere (Rosen 1985, Cariani 1989, 1991ab, and Kampis 1991). In the process of training, a new perceptual category (or primitive distinction) was created by the device *de novo*. This process is fundamentally distinct from the building up of complex distinctions from logical combinations of pre-existing ones, such as one finds in virtually all contemporary artificial intelligence strategies.

Pask's device is an example of a general class of evolutionary robotic devices which evolve their own hardware. The potential existence of evolutionary robots which construct their own sensors and effectors has a number of far-reaching epistemological consequences. Warren McCulloch was one of the few people at the time to see the broader implications of Pask's device:

> With this ability to make or select proper filters on its inputs, such a device explains the central problem of experimental epistemology. The riddles of stimulus equivalence or of local circuit action in the brain remain only as parochial problems. (McCulloch, in the preface of Pask 1960).

Finding the right feature primitives for a neural net

Pask's assemblage was the first artificial device when given a task, to automatically find the observable(s) or feature primitives necessary to solve the problem with which it was confronted. While Pask's device was quite rudimentary, it would be possible to build other devices along the same general organizational principles. We could then use these devices as front-ends for our computationally adaptive neural nets and trainable machines (Cariani 1990).

In contrast to the ability to evolve new sensory distinctions, all existing classifiers operate within fixed sets of feature primitives which the designer has determined is appropriate to the classification problem at hand. The device then adaptively searches for an appropriate algorithm to partition the space of feature combinations. If the set of feature primitives is obvious or the problem can be specified in purely formal terms, then feature selection is unproblematic. In the case of an ill-defined real world problem (Selfridge 1982), however, the appropriate feature primitives may not be obvious, and the designer probably will not have at his/her disposal a set of features which can correctly classify all cases. Typically in these cases the primitive feature set is tuned by the designer -- the designer thinks of all the possible features which might plausibly enhance performance and some subset of those features is chosen. If the features chosen are informationally inadequate to perform the task (e.g. to classify perfectly), then the performance of the adaptive device will plateau at this limit. In this case a device capable of adaptively constructing new sensors and hence new feature primitives would be able to surpass this plateau by further improving the basic features. A feature space which was difficult or impossible to correctly partition might well become more tractable under a change of observables. If an observable is added, the feature space increases in dimension by one, and what was a local minimum in the old feature space might easily be a saddle point in the augmented space. Hill-climbing, besides involving following local gradients upwards, can also thus involve increases in the dimensionality of the hill itself. Because it is usually tacitly assumed that the sensors (and hence the feature primitives) are fixed and that computational adaptation is the only form of learning possible, these alternatives are usually not explicitly discussed in the literature. However, in real life tasks, when a neural net fails to improve beyond a given level of performance, despite the best efforts of its designers to optimize the algorithm, often those designers go looking for more (or more suitable) observables, thereby altering the dimensionality and external semantics of the space they are searching. Sometimes this proves to be a successful strategy.

Increasing the number of independent signalling modes in a network

Networks of devices capable of evolving new sensors and effectors also have the means of raising the dimensionality of their signal spaces. Whenever a new effector-sensor combination arises between two elements there is the possibility that the sensor of one element can detect the actions of the new effector. If the new sensor and effector are different from pre-existing ones, then there will be some informational independence between signals sent through that channel and pre-existing signalling channels. For example, if we had one element evolving a primitive sensitivity to sound (or hormones or electrical pulses) while another evolved a primitive means of producing sound (or hormones or electrical pulses), then the network will have evolved a signalling channel which is (at least partially) orthogonal to pre-existing ones. This is a general strategy for increasing bandwidth – proliferate new orthogonal signalling modes. One major advantage of increasing dimensionality (categories) over proliferating states within existing dimensions is that different signals can be kept separate from each other if desired; different patterns need not interfere with each other (as in many kinds of connectionist networks), and new signals need not be logical combinations of pre-existing ones. This strategy also enables simultaneous transmission of different signal types (multiplexing). These two features allow for much more flexible adaptive signalling possibilities and are potentially relevant to our conceptions of the biological neuron as an adaptive information processing device (Cariani 1991b). Rather than discrete logic elements adaptively summing various combinations of synaptic inputs to produce one scalar output signal, we might also consider neurons in terms of mixed digital-analog elements capable of transmitting multidimensional information through interspike interval distributions (Chung, Raymond and Lettvin 1970, Raymond and Lettvin 1978). By tuning membrane recovery processes and electrical resonances, a neuron could adaptively proliferate new temporal discharge patterns, thereby increasing the dimensionality of the information encoded in its spike trains. It is possible that such tunings could be implemented by adaptive allocation of ion channels and pumps to local patches of excitable membrane, contingent upon their local past histories of excitation. In addition to adaptively altering effective connectivities through synaptic weights (conceived as scalars), in effect we would also have a learning mechanism which operated in the frequency domain, with all the advantages conferred by (partial) orthogonalities of different interspike intervals. Recurrent networks of these elements, each one adaptively tuning itself, might then be capable of proliferating new resonant network modes, thereby increasing the number of stable global states available to the network as a whole. Whether or not such functional organization proves to be operant in the nervous system, useful artificial mixed digital-analog neural nets can be constructed along these lines by utilizing trees of adaptively tuned pulse oscillators connected together in various ways (Pratt 1990).

The construction of an epistemically autonomous observer

Finally, once we have devices which can both construct their categories for sensing and effecting and also optimize the mapping between those categories, we have devices which are *epistemically-autonomous* -- within the limits imposed by their material structure, they determine the nature of what they see and how they act on the world. This conception of epistemic autonomy lies close to the concepts of *semantic closure* (Pattee 1982) *organizational closure* (Maturana 1981; Varela 1979), the *self-modifying device* (Kampis 1991; Csanyi 1989) and the *anticipatory system* (Rosen 1985, 1991). Once we have the concept of epistemic autonomy and understand the material structures needed to subserve this kind of functional organization, we then have the means for constructing an autonomous subject, a point of view independent of our own, one capable of telling us about that which we do not already know.

References

Ashby, W.R. 1956. *An Introduction to Cybernetics*. London: Chapman & Hall.

Cariani, Peter. 1989. On the design of devices with emergent semantic functions. Ph.D. dissertation, State University of New York at Binghamton.

Cariani, Peter. 1990. Implications from structural evolution: semantic adaptation. *Proceedings, International Joint Conference on Neural Networks, January, 1990, Washington, D.C.* I: 47-51.

Cariani, Peter. 1991a. Emergence and artificial life. In *Artificial Life II. Proceedings of the Second Workshop on Artificial Life,* edited by C.G. Langton, C.E.Taylor, J.D.Farmer, and S.Rasmussen. Reading, MA: Addison-Wesley.

Cariani, Peter. 1991b. Adaptivity and emergence in organisms and devices. *World Futures*, 31: 49-70.

Chung, S.H., S.A. Raymond, and J.Y. Lettvin. 1970. Multiple meaning in single visual units. *Brain Behavior and Evolution* 3:72-101.

Csanyi, Vilmos. 1989. *Evolutionary Systems and Society*. Durham, NC: Duke.

Emmeche, Claus. 1991. A semiotical reflection on biology, living signs, and artificial life. *Biology and Philosophy* 6: 325-340.

Kampis, George. 1991. *Self-Modifying Systems in Biology and Cognitive Science: A New Framework for Dynamics, Information, and Complexity*. New York: Pergamon.

Maturana, Humberto. 1981. Autopoiesis. In *Autopoiesis: A Theory of Living Organization,* edited by Milan Zeleny. New York: North-Holland.

McCulloch, Warren S. 1989. *The Collected Works of Warren S. McCulloch*. Salinas, CA: Intersystems Publications.

Mead, Carver. 1989. *Analog VLSI and Neural Systems*. Reading, MA: Addison-

Wesley.

Pask, Gordon 1958 Physical analogues to the growth of a concept. In *Mechanization of Thought Processes: Proceedings of a Symposium, National Physical Labs, November, 1958*. London: H.M.S.O.

Pask, Gordon. 1959. The natural history of networks. In *Self-Organizing Systems,* proceedings of a conference, Chicago, May 5-6, 1959, edited by M.C. Yovits and S. Cameron. New York: Pergamon Press, 1960.

Pask, Gordon . 1961. *An Approach to Cybernetics*. New York: Harper & Brothers.

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

Powers, William. 1973. *Behavior: The Control of Perception*. New York: Aldine.

Powers, William. 1989. *Living Control Systems. Selected Papers of William T. Powers*. Gravel Switch, KY: Control Systems Group.

Pratt, Gill. 1990. Pulse Computation. Ph.D. dissertation. Massachusetts Institute of Technology.

Raymond, S. A. and J. Y. Lettvin. 1978. Aftereffects of activity in peripheral axons as a clue to nervous coding. In *Physiology and Pathobiology of Axons,* edited by S. Waxman. New York: Raven Press.

Rosen, Robert. 1985. *Anticipatory Systems*. New York: Wiley.

Rosen, Robert. 1991. *Life Itself: A Comprehensive Inquiry into the Nature, Origin and Fabrication of Life* . New York:: Columbia University Press.

Selfridge, Oliver. 1982. Some themes and primitives in ill-defined systems. In *Adaptive Control in Ill-Defined Systems,*. edited by O. Selfridge, E.L. Rissland and M.A. Arbib. New York: Plenum Press.

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

Varela, Francesco. 1979. *Principles of Biological Autonomy*. New York: North Holland.

Wiener, Norbert. 1961. On Learning and Self-Reproducing Machines.*Cybernetics or Control and Communication in the Animal and the Machine*. Cambridge: MIT Press.

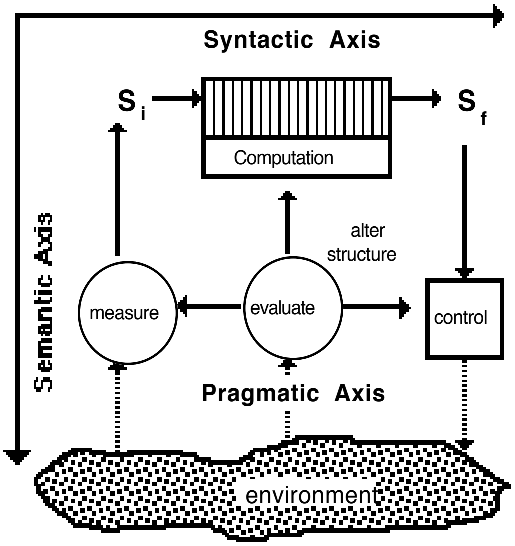

Figure 1. Basic informational relationships between an organism or device and its environment. The sensors, effectors and computational parts of a device correspond to the sensory organs, effector appendages (muscles, secretory organs), and coordinative organs (hormonal, nervous systems) of organisms. These material structures implement the informational functionalities of measurement (sensing), control (acting), and computation (coordinating). When embedded in an organism or device in the appropriate way, these three functionalities implement a modelling relation. The semiotics of this relation involve three independent axes: a semantic axis, a syntactic axis, and a pragmatic axis. The external semantics of the internal states of the organism/device are determined by its sensors and effectors, while the internal relations between those states can (often) be described in terms of rule-governed, syntactic relations. The pragmatics of devices are those evaluative structures which cause a change in internal structure so as to improve performance. In the case of biological organisms, the evaluative test is survival and reproduction, while in artificial devices the evaluative criteria are set by the designer-user to solve some task at hand.

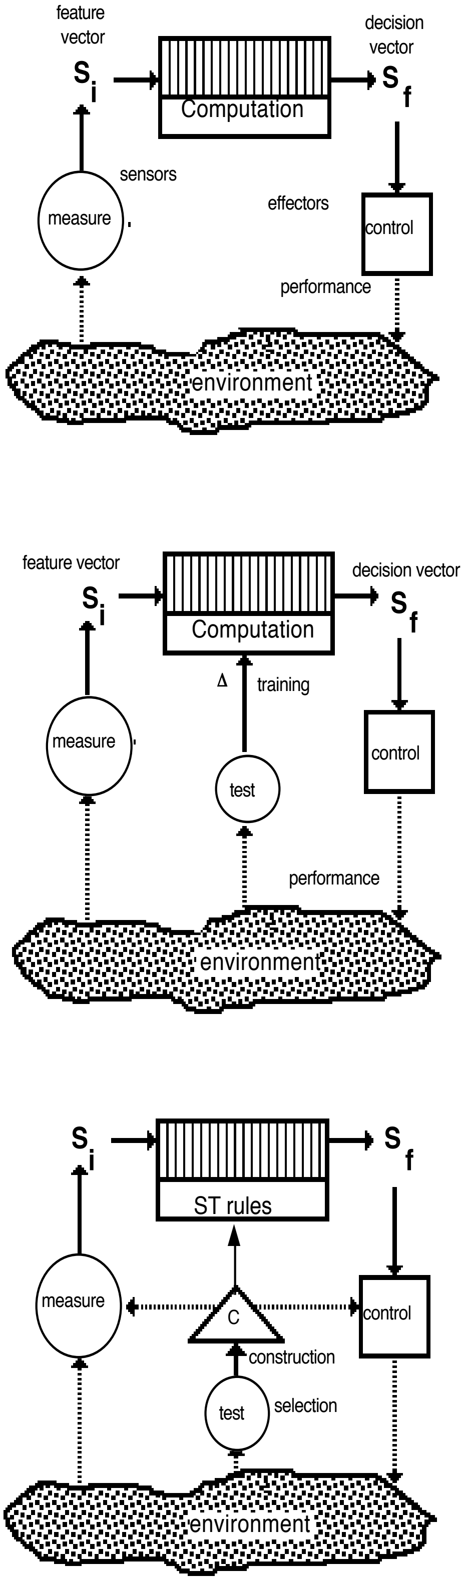

Figure 2. General types of adaptivity. The solid arrows represent digital or symbolic processes, while the dashed arrows represent analog or undifferentiated processes. (Top) *Nonadaptive robotic devices*. These devices do not modify their internal structure contingent upon their experience (past performance). There is no feedback from performance to enhanced structure. (Middle) *Adaptive computational devices*. These devices alter the input-output algorithm of their computational part contingent upon their performance. They operate within fixed sets of functional states, switching from alternative input-output functions (fixed hardware, variable software). These devices improve their input-output mappings over time but are constrained by the fixed, nonadaptive nature of their sensors and effectors. Because sensors and effectors determine the linkage of device states and the world at large and there is no means here of adaptively altering them, these devices are semantically-bounded. Such devices operate within the set of semantic categories defined by their sensors and effectors and within the set of syntactic states defined by their computational hardware. (Bottom) *Structurally adaptive devices*. These devices construct new material structures (new hardware), enabling new functional states and operations to arise. Such devices are capable of evolving new semantic categories through the adaptive construction of sensors and effectors. Those which evolve new sensors and effectors change their external semantics in order to optimize their performance. This is analogous to the evolution of new sensory and effector organs in biological evolution. Structually adaptive devices are also capable of augmenting the state spaces of their computational parts by building new material structures (more memory, processors) which increase the number of syntactic distinctions which can be manipulated. This is analogous to the evolution of more complex coordinative structures in biological organisms. The two adaptive device types are not mutually-exclusive: there could be devices which were both computationally-adaptive and structurally-adaptive.