

## Epistemic Autonomy through Adaptive Sensing

Peter Cariani

Eaton Peabody Laboratory for Auditory Physiology,  
Massachusetts Eye and Ear Infirmary, 243 Charles St., Boston, MA 02114  
peter@epl.meel.harvard.edu

### ABSTRACT

Sensors and effectors determine how events in the world at large are related to the internal informational states of organisms and robotic devices. Sensors determine what kinds of distinctions (perceptual categories, features, primitives) can be made on the environment. By "evolving the sensors" perceptual repertoires can be adaptively altered and/or enlarged. To the extent that devices can adaptively choose their own feature primitives for themselves, they gain a greater measure of "epistemic autonomy" vis-à-vis their designers. Such devices are useful in ill-defined situations where the designer does not know a priori what feature primitives are adequate or optimum for solving a particular task.

Several general strategies for adaptively altering or augmenting sensor function are proposed: 1) prosthesis: adaptive fabrication of new front-ends for existing sensors (e.g. telescopes), 2) active sensing: using motor-actions to alter what is sensed through interaction (poking, pushing, bending), 3) sensory evolution: adaptive construction of entirely new sensors (adaptive antibody construction, Gordon Pask's electrochemical device) and 4) internalized sensing: "bringing the world into the device" by creating internal, analog representations of the world out of which internal sensors extract newly-relevant properties (perceptual learning). Since many neural sensory representations appear to be analog and iconic in nature, neural assemblies can be adaptively formed to function as internal sensors that can switch behavior according to new perceptual categories.

**KEYWORDS:** *evolutionary robotics, epistemic autonomy, adaptive sensing, active sensing, semiotics, sensory coding, sensory prosthesis*

### 1. FUNCTIONAL AUTONOMY

Since antiquity, the functional organization of humans and animals has been broadly conceptualized in terms of sensing, thinking, and acting. Sensing is a process of gaining information about the external world; acting, the process of influencing courses of events in that world; and thinking, the coordinative process of choosing appropriate actions given the world as it is sensed. Action leads to changes in the world that subsequently alter perceptions to evoke new actions. When this sequence of percepts, coordinations, actions, and subsequent percepts is iterated, closed loops are formed between an organism and its environment that permit the organism to exert some limited control over its surrounds [19].

The semiotic organization of percept-coordination-action triads with different degrees of internal structural adaptivity are depicted in Figure 1. Strong parallels exist between the functional organization of organisms as informational systems and the operational structure of scientific models, and robotic devices [1, 6, 8, 17, 21]. Sensors and effectors determine the kinds relationships that can exist between internal functional states and the world at large. The sensors determine the perceptual categories that are available, whilst effectors determine the kinds of primitive actions that can be realized. Sensors and effectors thus determine the nature of the external semantics of the internal, informational states of organisms and robotic devices (Fig. 1 A). For sensing, the causal flow is from environment to organism; whereas, for action, causation flows from organism to environment.

Mediating between sensing and acting is a coordinative functionality that Aristotle called "the common sense." This deliberative faculty realizes the percept-action mappings that govern the system's behavior under different perceived circumstances. These coordinative mappings can be simple reflexes that dictate particular responses given particular sensory inputs, or they can be highly elaborated internal models that utilize past experiences as guides for action.

To the extent that a system operates via simple reflexive input-output mappings, its behavior is slave to its (immediate) external inputs. To the extent that more elaborated history-dependent (memory-dependent) mappings are present, the system's behavior becomes less dependent on its immediate past inputs, and more dependent upon its entire history. This history-dependence is a form of memory, and to the extent that percept-action coordinations are mediated by memory, the present behavior of the device depends upon past inputs that are increasingly distant in time and space.

To an external observer, a complex, history-dependent coordinative process makes the behavior of the organism or device more difficult to predict; the system appears to act more autonomously relative to its immediate inputs, to depend more directly on its own internal processes than on its external inputs. Our intuitive sense of what is animate vs. what is inanimate trades heavily on the appearance of autonomous action, e.g. independent, self-initiated movement that is not simply predicted.

Learning is the process by which internal structures and functions are modified through experience in order to improve performance (Fig. 1). For change to improve




Figure 1. Types of adaptivity: linkages of performance-adjustment cycles with percept-coordination-action cycles.

performance over time, structure/function alterations ( $\Delta$  signs) must be based in some way on evaluations of past experience ("test" operations). Adaptive informational functions thus require feedback loops that connect evaluative mechanisms with percept-coordination-action cycles so that more effective percept-action mappings can evolve over time (Fig. 1, B-D). These adjustments are pragmatic operations that bring system structures and functions into congruence with its goals. If sensors and effectors determine the external semantic relations for the internal states of system, these adaptive adjustment mechanisms determine their relations to the goals of the system, i.e. their pragmatic relations [6, 8].


These feedback mechanisms can tune existing analog parameters, as in an autofocusing servomechanism, or switch between alternative discrete percept-action mappings, as in a trainable classifier. Feedback can also be applied to alter the structure of hardware that subserves the functionalities of the percept-action loop: sensors, neural mechanisms, computational substrates, effectors. Such "feedback to structure" can potentially create new parameters by increasing the number of degrees of freedom available to the system [3, 8]. Once such structural self-steering

mechanisms are in place, the performance of the system becomes more directly tied to the goals implicit in the system's evaluative criteria than to the device's particular history (equifinality). The organism or device therefore attains a degree of functional autonomy relative to its surrounds – the device itself determines how and how well it achieves its goals.

## 2. ADAPTIVE SENSING

Classically, theories of learning and approaches to the design of adaptive devices have focused on feedback to coordinative parts (Fig. 1B). It is almost universally assumed that the sensors and effectors of such adaptive devices carry out fixed sensing and effecting functions. Given this assumption, learning then entails experience-dependent alterations in the coordinative part of the system, i.e. finding the right mappings between fixed sets of input features and output actions. A trainable machine learns to improve these mappings with experience. The mappings themselves can be relatively simple, converting current perceptual states into

Figure 2. Potential effect of adding a new feature. More distinctions are better. Adding another independent feature in a classifier system always increases the amount of information available for classification and the dimensionality of the feature space. A hard-to-partition or totally unpartitionable space can potentially be transformed into a easily partitioned one with the addition of a judiciously chosen feature. (Feature 3: plain vs. outlined letter style). Altering semantics of existing features would move positions of objects (A's, B's), around in the feature space, also potentially permitting simplified classification.



actions, or highly elaborate, taking into account past sequences of percepts, actions and internal states. To the extent that percept-action mappings are elaborated, that they are made more dependent on the remote history of the device, the device exhibits increased behavioral autonomy relative to its immediate inputs. For such devices, an external observer needs information concerning the history and current state of the system if its behavior is to be understood or predicted. To an observer confronted with such an elaborated system, its behavior appears more “complex”, appearing to behave more autonomously relative to both observer and surrounds.

There are limits to what can be achieved with fixed sensing and effecting functions. At best, a trainable classifier can only achieve those classifications that its feature primitives can effectively separate. Those structures and functionalities that are fixed from the outset, i.e. not subject to adaptive modification, must be foreseen and specified by a designer. If the designer fails to include particular kinds of features that are necessary for solving the classification problem at hand, such that the requisite information is simply not in the feature set that was provided, then no amount of computation on the provided features can correct the problem. In such situations, the machine's human designer must recognize that the features at hand are not up to the classificatory task, and come up with a new and usually larger feature set.

The process of adjusting sensing functions, finding new feature primitives, or changing observables is another form of learning, the learning of new categories rather than learning within existing categories [2]. In order to realize this kind of learning, a system must be able to adaptively alter its sensing functions, thereby changing the external semantics of its feature primitives (Fig. 1C). This can be accomplished in two basic ways: by redeploying existing internal degrees of freedom, or by creating new ones.

The first process alters sensing function without major changes in hardware, by using sensors in new ways (e.g. active measurement), while the second requires structural modifications that alter existing sensors or add on new ones. When an extra primitive feature is added to the classifier, the dimensionality of the feature space increases by 1; when existing sensors are altered or tuned in new ways, the dimensionality of the space remains constant, but the external semantics of the respective features are changed. Adding a new independent feature or changing the semantics

of an existing one (functionally, one now has a different feature than before) can drastically simplify and render tractable a classification problem that was formerly not soluble (Fig. 2). Thus, changing the features that are available by altering pre-existing sensors or evolving new ones permits improvements in performance that might not be possible under suboptimal sets of fixed features.

Biological evolution provides many examples of the creation and refinement of sensing functions. New sense organs have evolved, and existing ones have been adaptively modified. New sensory distinctions have been added to existing perceptual repertoires over time (e.g. color vision, echolocation, electroreception, magnetoception). Those lineages that survive and persist are able to do so in part because their sensory systems permit them to better detect those aspects of their environments critical for survival and reproduction (e.g. detecting, recognizing, and locating food, mates, and predators). In biological organisms, the expression of selected genes guides the development of sensory organs such that they can be reliably constructed in the individuals of each generation. Selective pressures produce surviving populations with genes that produce sensory systems that tend to be better adapted for particular niches. An evolutionary “feedback to structure” loop thus connects evaluative mechanisms (natural selection) with those that construct sensors (gene expression, developmental processes) that in turn result in adaptive alterations of sensing functions and percept-action mappings [6]. Evolution thus shapes the basic sensory distinctions that are available to an organism, thus in part determining the perceptual categories that constitute its experiential “life-world” (Umwelt) [24]. Sensory evolution is therefore a kind of phylogenetic learning process by which basic categories of perception are selected and refined. Sensory evolution parallels analogous perceptual learning processes that occur over the lives of individuals [12, 22].

Artificial devices can be designed and constructed so that they, too, can evolve the sensing operations that they need to detect relevant features of their surrounds. A traditional adaptive classifier can control which inputs to which it pays attention, but it cannot change the nature of those inputs, their external semantics. When a system has the capacity to adaptively modify or augment its sensors, then it in effect can choose its own perceptual categories. The more degrees of freedom that are available to a device in determining the

structure of its own sensors, the more control that it has concerning the types of empirical information it can access. When the device chooses the nature of its own inputs in this way, it attains a degree of epistemic autonomy relative to its designer. The device no longer must depend solely on its designer for the kinds of information (features, observables) it needs to solve a problem, and the designer is no longer burdened with the problem of foreseeing what perceptual categories are adequate or optimal. A degree of epistemic autonomy means that the system itself can define the terms of a problem in ways that were not anticipated by the designer; the system can come up with more creative solutions that are not simply logical combinations of pre-existing features [7, 22].

Several strategies are available for altering existing sensing functions or adding new ones: external prosthesis, active measurement, new sensor construction, and internal sensing. As means of expanding the repertoire of available sensing operations open to an organism or device, each of these strategies confers upon a system some degree of adaptive control over the nature of its inputs.

### 3. ADAPTIVE PROSTHESIS

Perhaps the easiest means of altering semantics of sensing operations is to interpose external objects (filters, active devices) between existing sensors and the world. Prosthetic devices thus modify the relationship between our sensors and the world beyond the prosthesis. In doing so, they alter sensing function without changing the structure of our sensors. For example, a geiger counter can be used to detect forms of radiation that we cannot apprehend directly with our senses. Internal linkages in the device connect elements sensitive to the radiation (to which we are insensitive) to those that produce physical disturbances to which we are sensitive (visible pointer readings or audible clicks).

While our sensors and the structural boundaries of our bodies are left intact, the prosthesis effectively alters the semantics of some of our perceptual states, such that those states are now linked in a different way to the world beyond the prosthetic sensor. When we are using a geiger counter, the clicks are no longer simply environmental sounds -- they are linked to the presence of radiation nearby. Existing degrees of sensory freedom (e.g. auditory attention, detection of particular sounds) are exchanged for the new ones afforded by the prosthesis (detection of radiation).

The linkage of the outputs of the prosthesis with our own sensors moves the functional boundary of the sensing operation, its point of contingency, outwards, to a point distal to the prosthesis. This is not unlike what happens when we wield a rigid stick: the functional and experiential boundaries of our bodies extend to the end of the stick [10].

Ultimately, all technology is prosthesis, the extension and augmentation of existing biological and informational capabilities. We put eyeglasses in front of our eyes to bring the world into focus, we use telescopes and microscopes to see those realms which our unaided eyes cannot fathom, and

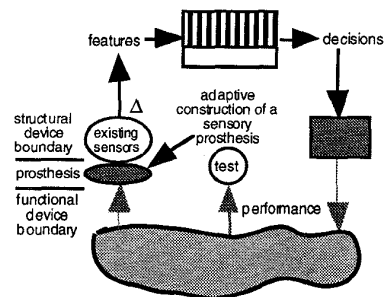


Figure 3. Adaptive construction of a sensory prosthesis.

we employ remote sensors to explore the planets. These are all sensory prostheses. Systems of writing and of calculation function as extensions of our own coordinative, deliberative, memory- and logic-based functionalities. The technological implements that transport us, grow our food, or keep us warm are similarly extensions of our own effectors.

Artificial devices that adaptively construct prosthetic devices for themselves can be envisioned (Figure 3). One general strategy is to use pattern grammars to specify alternative assembly-processes and genetic algorithms to adaptively select alternatives are chosen for realization [11].

### 4. ACTIVE MEASUREMENT

Active measurement is the process of acting on the world and sensing how it behaves as a result of one's actions. For example, an object can be poked to see if it moves (a strategy beloved by small children). The interactive measurement conveys information that may be very different in nature from passive observation with the same sensors. A distinct action or motor-sequence sets up a different active measurement, a different observable. In the physicist's operational terms, each motor-sequence "prepares the system" in a different way. Most experimental science is an active measurement process, as elaborate motor-rituals are played out to set up physical systems in different ways in order to make particular measurements. By changing motor-plans that set up measurements, one changes reference states of sensors, thereby altering what is measured.

While active measurement changes sensing functions without altering sensor structures, it does require additional coordinative and motor resources. For active measurements motor degrees-of-freedom are exchanged for new sensory degrees of freedom, as motor-programs involved with the active measurement pre-empt other uses for their elements. If the system can support adaptive control of motor sequences, then particular combinations of movements can be made contingent on the quality of the information they bring in (relative to some task), and the motor system in effect becomes a part of the perceptual system. Possibilities for active measurement are only limited by the kinds of physical actions that the system can carry out. If external tools are used in motor-sequences, then all manipulations afforded by current technologies become possible (including fabrication of sensory prostheses), and the limits of active




Figure 4. Adaptive construction of an entirely new sensor.

measurement become coextensive with those of science and technology. Needless to say, these boundaries are ill-defined and may be impossible to delineate.

### 5. NEW SENSOR CONSTRUCTION

While the strategies of prosthesis and active measurement leave original sensors unchanged, entirely new sensors can also be adaptively constructed (Figure 4). For example, the immune system constantly produces populations of molecular sensors (antibodies) that are selected (differentially produced) based on their effectiveness in recognizing foreign agents (have higher antigen binding affinities). Molecular sensing functions are continually refined by mutational processes (hypermutation) that produce new molecular sensors by slightly altering genetic plans. Thus as an immune response is mounted, the available repertoire of molecular sensors is adaptively altered and enlarged.

Prostheses and active measurements redeploy available sensory and motor degrees-of-freedom for new sensory degrees-of freedom. Construction of new, independent sensors creates new sensory degrees of freedom (feature spaces increase in dimensionality). New sensing operations not only require new sensor hardware, but also new coordinative capabilities in order to make use of the additional sensory distinctions that have been created.

Devices that adaptively construct new sensors have been built. In the late 1950's British cyberneticist Gordon Pask fabricated electrochemical assemblages that grew their own sensors [3, 15, 16]. These rudimentary devices are apparently the only artefacts constructed thus far that adaptively find their own "relevance criteria," i.e. those observables of feature primitives that are relevant for some task. Albeit in an extremely limited way, the device solved its own "frame problem" by evolving its own relevance criteria.

New sensor construction might be a useful strategy in ill-defined problem domains, where one does not know a priori what kinds of information are needed to solve a problem. In these contexts, self-organizing assemblages such as Pask's would serve as front-ends for trainable classifiers. Whenever desired levels of performance could not be achieved within the current set of feature primitives, a Pask-like assemblage would search for more appropriate ones, and the cycle would begin anew.

### 6. INTERNALIZED SENSING

Prostheses add new sensors distal to an organism or device's structural boundary, while de novo creation of new sensory organs adds them right at the structural boundary. While construction of prostheses exteriorizes the problem of making new linkages with the environment, this problem can also be interiorized by bringing the environment into the device itself. For example, Pask's electrochemical assemblage consisted of a set of platinum electrodes immersed in an aqueous medium that was in extensive mechanical contact with the external world. External vibrations were internalized in the aqueous medium, making them available for interactions with electrodes and attendant filament structures that were under adaptive control.

Devices can thus be built in which internal sensors convert complex analog interactions within the device itself into discrete feature values (Figure 5). The internal milieu is effectively an iconic, analog transformation of the environment that preserves much of its dynamic richness and subtlety. In contact with this internal milieu internal detectors that are sensitive to motions of the medium (e.g. hair cells in the cochlea). Internal detectors then register different aspects of the complex motions of the internal medium that are relevant for a particular classification task.

As in the immune system, this arrangement permits the sensory system to potentially access a wide range of environmental properties, both simple and complex, without explicitly detecting and discretely encoding them all. When a given property becomes relevant to performing some task (e.g. making a particular distinction or detection), then a special-purpose internal sensor can be adaptively constructed to register that property.

One can consider biological sensory systems as encoding the forms of their respective stimuli in patterns of neural discharge. Whether considered in terms of discharge rates, time patterns, or relative latencies, responses of sensory neurons are seldom all-or-none, but show gradations of response to different stimuli [18, 20]. The initial neural representations of sensory information therefore have analog characteristics. In many sensory systems, such as audition, mechanoreception, and vision, the stimulus impresses its own time structure on the timings of neural discharges ("stimulus-locking" or "phase-locking"). In these systems, relative spike arrival times and time intervals between spikes can convey perceptually-relevant information about stimulus qualities [4, 5, 9, 18, 20] in a highly precise and robust manner. Interspike intervals themselves can take on continuous ranges of time duration; to the extent they reflect the stimulus time structure, they constitute iconic, analog representations of the stimulus.

At more central stages of sensory processing, ensembles of neural elements that are sensitive to particular classes of spatio-temporal response patterns can be adaptively formed




Figure 5. Adaptive construction of a new internal sensor.

using the different connectivities, time delays, and intrinsic membrane dynamics that are available [13]. By changing effective interneural connectivities, such "neural assemblies" could be adaptively constructed to respond differentially to particular activity patterns in the neural medium [14, 23]. Differential responses would then switch motor activity (as in the context of a two-alternative matching task). When a newly-formed neural assembly functions as such a switch, then a new sensing function has been created on the internal milieu. A new analog-to-digital mapping is implemented. New feature-primitives [22] and signal-primitives [7] can potentially arise this way.

By virtue of correlation of the internal milieu with the external environment, a new sensing function is thereby also realized on the external world. In Gibsonian terms, the correlational structure of the environment has been brought into the nervous system. Invariant properties of neural activity that correspond to affordances in the environment can then be extracted by emergent assemblies of neurons. Perceptual repertoires can thus be expanded, limited only by the ability of internal milieus to capture the richness of their effective stimuli and to exploit that richness by forming appropriate neural assemblies. Universal correlation-based mechanisms for representing sensory forms and analyzing then permit neural assemblies to be constructed to extract (almost) arbitrary sets of pattern invariances, e.g. [23, 25]. In bringing the environment within, the plasticity of the nervous system can be harnessed to provide greater adaptability of sensing functions. By such mechanisms, adaptive systems can thus extend their range of sensing functions, thereby achieving greater epistemic freedom.

## 7. CONCLUSIONS

Together with other operations that create more powerful and flexible coordinative capabilities and more elaborated internal anticipatory models, adaptive sensing operations permit systems to determine for themselves how they represent the world around them. Simple reflexive reactivity can thus be replaced by a measure of epistemic autonomy.

## 8. REFERENCES

- [1] Cariani, P. "Emergence and Artificial Life." In C. Langton, C. Taylor, J. Farmer, S. Rasmussen, eds., *Artificial Life II*. Addison-Wesley: Redwood City, CA, 1992, pp. 775-798.
- [2] - "Some Epistemological Implications of Devices which Construct their own Sensors and Effectors." In F. Varela, P. Bourguine, eds., *Towards a Practice of Autonomous Systems*. MIT Press: Cambridge, MA, 1992, pp. 484-493.
- [3] - "To Evolve an Ear: Epistemological Implications of Gordon Pask's Electrochemical Devices," *Systems Research*. vol. 10, pp. 19-33, 1993.
- [4] - "As if Time Really Mattered: Temporal Strategies for Neural Coding of Sensory Information." *Communication and Cognition-AI*, vol. 12, pp. 161-229. Reprinted in K. Pribram, ed., *Origins: Brain and Self-Organization*. Lawrence Erlbaum: Hillsdale, NJ, 1995.
- [5] - "Temporal Coding of Sensory Information." In J. M. Bower, ed., *Computational Neuroscience: Trends in Research, 1997*. Plenum: New York, 1997, pp. 591-598.
- [6] - "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*. Kluwer: Dordrecht, Holland, 1998, pp. 359-377.
- [7] - "Emergence of New Signal-Primitives in Adaptive Timing Networks," *Intellectica*. in press.
- [8] - *On the Design of Devices with Emergent Semantic Functions*. Ph.D. thesis, State University of New York at Binghamton, 1989.
- [9] Cariani, P.A. and B. Delgutte. "Neural Correlates of the Pitch of Complex Tones. I. Pitch and pitch salience. II. Pitch shift, pitch ambiguity..." *J. Neurophysiology*. vol. 76, pp. 1698-1734, 1996.
- [10] Churcher, J. "Implications and Applications of Piaget's Sensorimotor Concepts." In O. Selfridge, Rissland, E., Arbib, M., eds., *Adaptive Control of Ill-Defined Systems*. Plenum Press: New York, 1982, pp. 289-304.
- [11] Gause, D. and G. Rogers. "Genetic Pattern Synthesis - An Approach to Man-Machine Symbiotic Design." in Proceedings of Third European Meeting of Cybernetics and General Systems Theory, 1976.
- [12] Gibson, E.J. *Principles of Perceptual Learning and Development*. Appleton-Century-Crofts: Cambridge, 1969.
- [13] Hebb, D.O. *The Organization of Behavior*. Simon & Schuster: New York, 1949.
- [14] Morrell, F. "Electrical Signs of Sensory Coding." In G. C. Quarton, T. Melnechuck, F. Schmitt, eds., *The Neurosciences: A Study Program*. Rockefeller U.: New York, 1967, pp. 452-469.
- [15] Pask, G. "Physical Analogues to the Growth of a Concept." In ed., *Mechanization of Thought Processes, Vol II*. H.M.S.O.: London, 1959, pp. 765-794.
- [16] - *An Approach to Cybernetics*. Harper: New York, 1961.
- [17] Pattee, H.H. "Universal Principles of Measurement and Language Functions in Evolving Systems." In J. Casti, A. Karlqvist, eds., *Complexity, Language, and Life: Mathematical Approaches*. Springer-Verlag: Berlin, 1985, pp. 268-281.
- [18] Perkel, D.H. and T.H. Bullock. "Neural Coding," *Neurosci. Res. Prog. Bull.* vol. 6, pp. 221-348, 1968.
- [19] Powers, W. *Behavior: The Control of Perception*. Aldine: New York, 1973.
- [20] Rieke, F., D. Warland, R. de Ruyter van Steveninck and W. Bialek. *Spikes: Exploring the Neural Code*. MIT Press: Cambridge, 1997.
- [21] Rosen, R. *Anticipatory Systems*. Pergamon: Oxford, 1985.
- [22] Schyns, P.G., R.L. Goldstone and J.-P. Thibaut. "The Development of Features in Object Concepts," *Behavioral and Brain Sciences*. in press.
- [23] Thatcher, R.W. and E.R. John. *Functional Neuroscience, Vol. I. Foundations of Cognitive Processes*. LEA: Hillsdale, NJ, 1977.
- [24] von Uexküll, J. *Theoretical Biology*. Harcourt, Brace & Company: New York, 1926.
- [25] Uttal, W.R. *An Autocorrelation Theory of Form Detection*. Wiley: New York, 1975.