Proceedings, Dagstuhl Seminar on Computational Creativity, Schloss Dagstuhl – Lebnitz Center for Informatics, Wadern, Germany, July 12-17, 2009, Dagstuhl Workshop number 09291.

Strategies for creating new informational primitives in minds and machines

Peter Cariani[1]

Abstract

Open-endedness is an important goal for designing systems that can autonomously find new and expected solutions to combinatorically-complex and ill-defined problems. Classically, issues of open-ended generation of novelty in the universe have come under the rubric of the problem of emergence. We distinguish two modes of creating novelty: combinatoric (new combinations of existing primitives[2], “resultants”) and creative (new primitives, “emergents”).[3] Although combinatoric systems may differ in numbers of possible combinations, their set of possibilities is closed. Creative systems, on the other hand, have open-sets of possibilities because of the partial- or ill-defined nature of the space of possible primitives. The two conceptions provide two modes for describing and understanding change and creativity: as the unfolding consequences of fixed combinatorial rules on bounded sets of pre-defined primitives or as new processes and interactions that come into play over time to define new primitives.

We face several related problems. We want to know how to recognize creative novelty when it occurs (the *methodological problem*), we want to understand the creative process in humans and other systems (*the scientific problem*) such that creativity in human-machine collaborations can be facilitated and that creativity in autonomous devices can be usefully enabled (the *design problem*).

The methodological problem can be solved by the “emergence-relative-to-a-model” approach in which an observer forms a model of the behavior of a system. Novelty and creativity are inherently in the eye of the observer, i.e. relative to some model that specifies possible alternatives. If the behavior changes, but it can still be predicted/tracked in terms of the basic categories/state set of the model, one has combinatorial creativity. If it changes, but requires new categories/observables/states for the observer to regain predictability, then one has emergent creativity (creation of new primitives).

We argue that pure computation by itself can generate new combinations of symbol primitives, but, absent states or processes that are hidden from the observer, it cannot autonomously generate new primitives. Breakout strategies are therefore required. In order for a computational system to transcend the limitations of its own primitive symbol set, it must be coupled to some other non-symbolic, material system. In order to increase its effective dimensionality, it can couple to the world outside its internal symbol-states by three means:

1) *via human-machine interactions* (facilitate novel insights in humans, use humans to create new primitives that expand systems, develop tools for creativity),
2) *via sensors and effectors on an external world* (epistemically-autonomous evolutionary robots), and
3) *via internal analog dynamics* (adaptive self-organization in mixed analog-digital devices or biological brains).

When a computational system is augmented and opened up in these ways, it is transformed from a formal system that is informationally isolated from its surrounds to one that is self-organizing, self-complexifying, and in informational interaction with its surrounds.

We discuss classes of adaptive and self-modifying cybernetic evolutionary robotic devices in terms of combinatoric and creative novelty and in terms of new functionalities that are created (new syntactic states, new semantic observables & actions, new pragmatic goals). If adaptive sensors and effectors are internalized in the form of signal generators and receivers, it is possible to think of neural networks in these terms. What this view of biological brains might look like is sketched out. Adaptively-tuned neuronal assemblies would conceivably function as internal sensors and signal generators, such that new signal types could be produced (i.e. new concepts). Emergence of new signal types would increase the effective dimensionality of internal signal spaces over time in an apparently open-ended manner.

---

[1] Academic affiliation: Department of Otology & Laryngology, Harvard Medical School. Mailing address: 629 Watertown St., Newton, MA 02460; cariani@mac.com, www.cariani.com. For deeper discussion see, Emergence of new signal-primitives in neural networks, Intellectica 2:95-143: 1997 (available on my website).

[2] By primitive, we mean an indivisible, unitary entity, atom, or element in a system that has no internal parts of structure from the perspective of that system. Individual symbols are the primitives of string rewrite systems, binary distinctions are primitives in flip-flop-based digital computers, total machine states are primitives in finite state automata.

[3] Lloyd Morgan (1931) distinguished "emergents" from "resultants." Emergents are the result of novel creation, resultants, of novel combination.

Dagstuhl Seminar Proceedings 09291
Computational Creativity : An Interdisciplinary Approach
http://drops.dagstuhl.de/opus/volltexte/2009/2192

Emergence and creativity

Emergence concerns the means by which novelty arises in the world. Intuitively, emergence is the process by which new, more complex order arises from that which is, in some sense, simpler or more predictable. As such, images of birth, development, and evolution infuse our notions of emergence. These images provide explanations for how novelty, spontaneity, and creativity are possible and how complex organizations arise and become further elaborated.

All around us we see the complex organizations that are the emergent products of biological, psychological and social processes. Our current discourses on emergence consequently encompass a wide range of phenomena: the appearance of new material structures (thermodynamic emergence), formal structures (computational emergence), biological structures and functions (emergent evolution), scientific theories (emergence vs. reduction), modeling relations in observers, percepts, ideas, notational systems, and economic and social relations.

Two fundamental conceptions of emergence can be distinguished: *combinatoric emergence* and *creative emergence*.[4] These two accounts of the origin of novelty parallel notions of the origin of order: "order-from-order" vs. "order-from-noise."[5] Where order comes from order, novelty is but a preformationist unfolding of latent possibility; where order arises from noise, chaos, or formlessness, novelty entails de novo formation of new realms of possibility. Both kinds of emergent orders are built up from basic sets of possibilities that constitute the most basic building blocks of the order, its “primitives.” Emergence then entails either the appearance of new combinations of previously existing primitives or the formation of entirely new ones. The primitives in question depend upon the discourse; they can be structural, material "atoms"; they can be formal "symbols" or "states"; they can be functionalities or operations; they can be primitive assumptions of a theory; they can be primitive sensations and/or ideas; they can be the basic parts of an observer's model. To say that an entity is "primitive" relative to other objects or functions means it cannot be constructed from combinations of the other entities, i.e. its properties cannot be logically deduced from those of other entities.

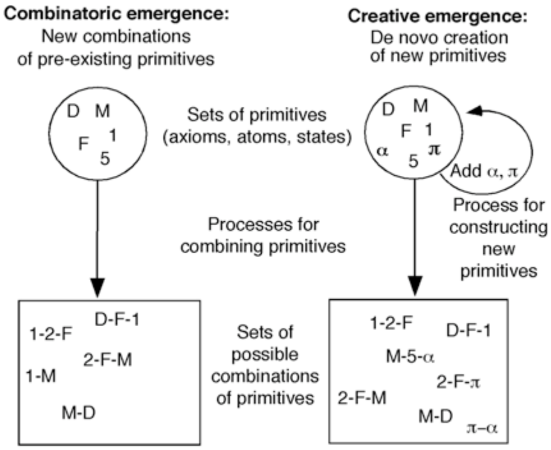

Figure 1. Combinatoric and creative emergence.

Novel combinations of closed sets of primitives

Combinatoric emergence assumes a fixed set of primitives that are combined in new ways to form emergent structures. Thus in biological evolution, new genetic DNA sequences arise from combinations of pre-existing nucleotides, codons, and codon-sequences. Microevolution proceeds through novel combinations of genes; new genes through novel combinations of nucleotide sequences. Likewise, new,

---

[4] This fundamental distinction is paralleled in Margaret Boden’s division of explorative vs. transformational creativity (M. Boden, 1990a, 1990b; M. A. Boden, 1991, 1994, 2006).

[5] As Piatelli-Palmarini so elegantly pointed out (Piatelli-Palmarini, 1980), the debate that occurred between Piaget, Chomsky, and Fodor ostensibly over the origins of new ideas was really a debate over the existence and nature of emergent novelty in the world. The two poles of the debate were held by Fodor, who defended an extreme preformationist view (all learning is belief-fixation, from a fixed repertoire of possible beliefs), and Piaget, who defended an emergentist view (qualitatively new concepts are created anew). The combinatorial-creative distinction parallels that between ontological and epistemological perspectives on emergence. For those who assume an ontological, realist, and/or reductionist stances that provisionally take a "God's-eye view" of objects in a universe, all novelty is necessarily of the combinatoric sort – questions of the origins of ontological entities are arguably incompatible with an atemporal ontological framework itself.

emergent structures and functions are thought to arise from novel combinations of previously existing molecular, cellular, and organismic structures and functions. In psychology, associationist theories hold that emergent perceptual states arise from novel combinations of pre-existing primitive sensations. Whether cast in terms of platonic forms, material atoms, or mental states, combinatoric emergence is compatible with reductionist programs for explaining macroscopic structure through microscopic interactions.

This strategy for generating structural and functional variety from a relatively small set of primitive parts is a powerful one that is firmly embedded in many of our most advanced informational systems. In the analytic-deductive mode of exploration and understanding, one first adopts some set of axiomatic, primitive assumptions, and then explores the manifold logically-necessary consequences of those assumptions. In the realm of logic and mathematics, the primitives are axioms and their consequences are deduced by means of logical operations on the axioms. Digital computers are ideally suited for this task: to generate combinations of symbol-primitives and logical operations on them that can then be evaluated for useful, interesting, and/or unforeseen formal properties. In the field of symbolic artificial intelligence (AI) these kinds of symbolic search strategies have been refined to a high degree. Correspondingly, in the realm of adaptive, trainable machines, directed searches optimize prespecified combinations of features and actions (feature-action mappings). What formally distinguishes different kinds of trainable machines, e.g. neural networks or genetic algorithms, are the structures of the respective combination-spaces being traversed, and the rules that direct the search processes.

Limitations of closed sets of primitives

Combinatoric novelty is a dynamic, creative strategy insofar as it constantly brings into being new combinations of elements. However, such combinatoric realms are inherently limited by their fixed, closed sets of primitive elements.[6] Arguably, all that can happen within such universes are recombinations of existing, prespecified symbols – there is no means by which new primitive symbols can be created by simply recombining existing ones. One does not create new alphabetical letter types by stringing together more and more existing letters – the new notations must be introduced from outside the system by external agents. Likewise, in our computer simulations, we set up a space of variables and their possible states, but the simulation cannot add new variables and states simply by traversing the simulation-states that we have given it. This "closed-world" character of computational systems poses fundamental problems for purely symbolic approaches to artificial intelligence and artificial life. Godel's Undecidability theorems do not evade this closure.[7] Far more limiting to pure computation is the inability to make measurements and act on the world directly -- one cannot ascertain whether it is raining outside or physically throw a ball simply by performing a computation on symbols.

---

[6] Consider the set of all the 6-digit sequences of digits 0-9, (one set containing 10^6 elements) vs. the set of all sequences of 6 arbitrarily defined objects (an indefinite number of sets containing 10^6 elements). The first set is well-defined and closed, while the latter is ill-defined and open.

[7] Godel's impotency principles and Turing's Halting Problem only apply to formal systems with potentially-infinite symbol strings or machine states (e.g. arithmetic operations on the potentially-infinite set of natural numbers). However, all physically-realizable formal-computational systems by necessity have finite state sets, and therefore are functionally equivalent to deterministic finite state automata (FSA). For finite systems, the set of computations is finite and bounded and therefore surveyable, such that consistency can always be tested within a finite number of steps. An FSA either reaches a terminal state (halts) or repeats a total machine state (in which case we know then that it will never halt). Practically speaking we are always limited by computational complexity and our own limited temporal existence rather than by imaginary computability constraints. Rather than tolerating or celebrating undecidability, we should recognize that logical consistency can be proved and paradoxes avoided if systems are kept small (finite) and computationally tractable.

Thus, computations play functional informational roles that are completely disjoint from and complementary to sensing (measurements) and physical action.

As entities in and of themselves, digital computers and formal systems are therefore bounded and closed, but in collaboration with human beings, they can greatly facilitate formation of entirely novel ideas in their human collaborators. In turn their human collaborators can add new primitives to expand their state-sets. Thus human-machine combinations can be open-ended systems that generate new primitives. Enhancing human creativity using flexible human-machine interfaces and other "tools for creativity" is a much more efficient route at present for generating open-ended novelty than attempting to build autonomous self-organizing systems that are creative in their own right.

Creation of new primitives

Classically, “emergence” has concerned those processes that create new primitives, i.e. properties, behaviors, or functions that are not logical consequences of pre-existing ones. Primitive-creation is the central issue for creative emergence, but it also can be asked how the particular primitives of an existing combinatorial system came into being in the first place.[8] Creative emergence, on the other hand, comes naturally to those who adopt the epistemic perspective of a limited, but expandable observer. Primitive objects in such a world almost always contain properties that are not be fully known to the observer. These hidden aspects can come into play as primitives interact through the underlying material processes that subserve them.

In this latter view, creating a new primitive entails the formation of a new property or behavior that in some strong sense was not predictable (by the limited observer) from what came before. The most salient examples of this kind of emergence involve the biological evolution of new sensory capabilities. Where previously there may have been no means of distinguishing colors, odors, or sounds, eventually these sensory capacities evolve in biological lineages. From a set of primitive sensory distinctions, one can list all combinations of distinctions that can be made with those primitives, but there are always yet other possible distinctions that are not on the list. For example, we cannot combine information from our evolution-given senses (sight, hearing, smell, etc.) to detect gamma radiation. Creation of the ability to sense gamma rays, through biological evolution or artificial construction of measuring instruments, thus adds a new primitive to the set of perceptual distinctions that can be made.

Observables are the perceptual primitives of scientific models. If a given model fails to accurately predict the observed behavior of some material system, we may very well require additional observables to fully predict or explain its behavior. In this case we cannot arrive at new observables simply by making computations on the states of existing ones; we must go out and construct a new kind of measuring instrument that will give us yet another independent window on the world. Each independent observable yields a different perspective that is not completely translatable into the states of other ones. Each independent observable represents a different category (e.g. mass,voltage, current, temperature, velocity, barometric pressure, humidity, tensile strength) and therefore must be given a separate unit-dimension in a model. Models with disjoint sets of observables thus may not be reducible to each other because of different, incommensurable categories.

Evolvable cybernetic devices

Artificial devices that create their own perceptual primitives can be built. The best example– and perhaps the only one – is a electrochemical device that was constructed by the British cybernetician Gordon Pask in the late 1950’s (Bird & Di Paolo, 2008; Peter Cariani, 1993; Pask, 1958, 1960, 1961).

Its purpose was to show how a machine could evolve its own “relevance criteria.” Current was passed through an array of platinum electrodes immersed in an aqueous ferrous sulphate/sulphuric acid medium, such that iron filaments grew to form bridges between the electrodes. By rewarding structures whose conductivity covaried in some way with an environmental perturbation, structures could be adaptively steered to improve their sensitivity. Pask’s device acquired the ability to sense the presence of sound vibrations and then to distinguish between two different frequencies. In effect, the device had evolved an ear for itself, creating a set of sensory distinctions that it did not previously have. Albeit, in a very rudimentary way, the artificial device automated the creation of new sensory primitives, thereby providing an existence proof that creative emergence is possible in adaptive devices.

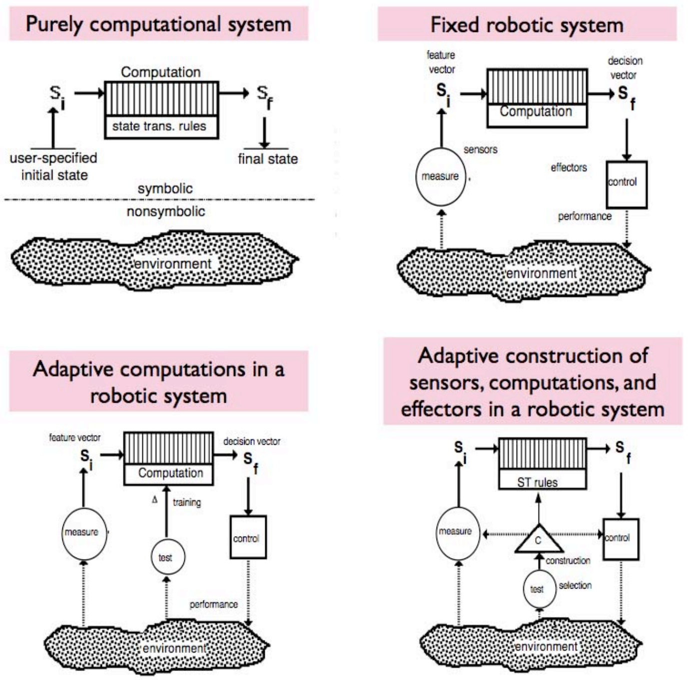

Figure 2. A taxonomy of adaptive devices based on primitive functionalities of observer-actors.

One can formulate taxonomies of possible mixed analog-digital cybernetic devices and their creative capacities (P. Cariani, 1989; Peter Cariani, 1992, 1998; de Latil, 1956; Pask, 1961). The devices in my own taxonomy have internal states and operations that link states with each other and to the external world (Figs. 2 & 3). The basic functionalities are “computation” (coordination, including memory mechanisms), “measurement” (sensors), “action” (effectors), and goal-based “evaluations” (measurements of performance that adaptively switch behavior or self-constructions). We think these functionalities account for the basic operational structure of the observer-actor. There is the cycling of signals from sensors to coordinative elements to effectors (outer loop in the diagram) and "feedback to structure" (inner loops) in which evaluative mechanisms steer the modification and/or construction of hardware (sensors, computational, coordinative structures, effectors).

We initially considered the state-transition structure in simple computational devices, then added sensors and effectors to produce fixed robotic devices that couple the internal, symbol-states to the world in nonarbitary ways.[9] One can then add evaluative sensors and steering mechanisms that switch the behavior of the computational part to produce adaptive computational machines. This is the basic

---

[9] See (P. Cariani, 1989) for the extended methodological discussion of state-transition structure of digital computers, computers with sensors, and the rest of the taxonomy. We strongly believe that biological brains are mixed analog-digital systems (and not deterministic finite state automata or Turing machines), and hence they are not the digital computational corrdinative systems depicted in the taxonomy.Nevertheless, we use digital state-determined discrete-state machines to simplify discussion of the operational, methodological problems of distinguishing between types of adaptivity. If the internal functional states of a system are analog and consequently, in some respect ill-defined (signal vs. noise), then it becomes more difficult and tedious to recognize clearly when a new sensor or effector has been created. This is not a problem if one has discrete states whose recognition can be rendered unambiguous.

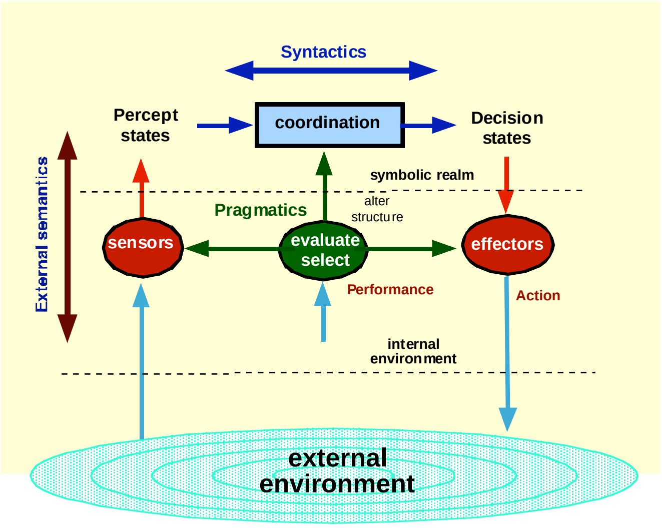

Figure 3. Basic functionalities of adaptive systems in the taxonomy and their semiotic dimensions.

high-level operational structure of virtually all contemporary trainable machines that use supervised learning feedback mechanisms (adaptive classifiers and controllers, genetic algorithms, neural networks, etc.). Here the internal states and their external semantics are fixed, such that the evaluative-steering mechanism merely switches input-output (percept-action, feature-decision) mappings using the same set of possible states. This is a form of combinatorial creativity, since the machine searches through percept-action combinations to find more optimal ones.

Consider the case, however, where the evaluation mechanism steers the construction of the hardware of the device rather than simply switching input-output mappings. If sensors are adaptively constructed contingent on how well they perform a particular function, then the external semantics of the internal states of the device are now under the device's adaptive control. When a device has the ability to construct and therefore to choose its sensors – which aspects of the world it can detect – it attains a partial degree of "epistemic autonomy." In effect the device has changed its own observables, added a new perceptual primitive, while retaining its existing set of internal states. When the construction process adds a new independent sensor, then the dimensionality of the internal percept feature space increases by one. Internal states associated with other sensors or other functionalities would need to be reallocated to register the inputs of the new sensor, so while the state set remains the same size, its effective dimensionality has increased.

Construction of the hardware of the device also allows for the expansion of internal states as well as the numbers of sensors and effectors that couple them to the world at large. Adaptive self-constructing devices can add new internal states (much as we add more memory to digital computers) that can be coupled to new sensors and effectors to form new semantic linkages. They can be coupled to newly constructed evaluative and steering mechanisms to form new pragmatic, goal-seeking relations.

Table I. Modes of creativity with respect to semiotic dimensions

Dimension	Primitives	Stable systems *Maintain structure*	Combinatoric systems *Search existing possibilities*	Creative systems *Add possibilities*
Syntactic	states & operations	deterministic-finite-state-automata formal systems (fixed machines)	Adaptive changes in state-transition rules (trainable machines)	Evolve new states & rules (growing automata)
Semantic	measurements & actions	Fixed sensors & effectors (fixed robots)	Adaptive search for optimal combinations of existing sensors & effectors	Evolve new observables, actions (epistemic autonomy)
Pragmatic	goals	Fixed goals (fixed self-direction)	Search combinations of existing goals (Adaptively prioritize goals)	Evolve new goals (creative self-direction)

The functional organization of these various devices has syntactic, semantic, and pragmatic aspects.[10] Syntactics describes rule-governed linkages between signs, semantics, the relation of signs to the external world, and pragmatics, the relation of signs to purposes (goal states). Creative emergence in the syntactic realm involves creation of new signs (symbols, internal states) that link internal states with other internal states (e.g. perceptual states with decision states, percept-action mappings). Creative emergence in the semantic realm involves creating new observables and actions (e.g. sensors, effectors) that contingently link the outer world with internal states. Creative emergence in the pragmatic realm involves creating new goals and evaluative criteria. These various functionalities can be either stable (fixed), subject to combinatorial search, or capable of de novo creation of new primitives (Table I).

Combinatoric creativity requires the ability to switch internal state-transition rules ("software") and/or sets of existing sensors and effectors and therefore does not require new states per se. Arguably creative emergence requires the ability to modify material structures or organizations (e.g. "hardware"), be it to create new states, sensors, effectors, or goals. It is obvious enough that one cannot enlarge the set of internal states by switching among them; one cannot bring new internal machine states into being by computing on the old state set. All growing automata require some material growth process that lies outside the symbolic description of digital states. In this context, we could speak of analog and digital aspects of these devices and their functionalities, the analog aspect being related to rate-dependent processes, and the digital aspect being related to those behaviors that can be described by stable rate

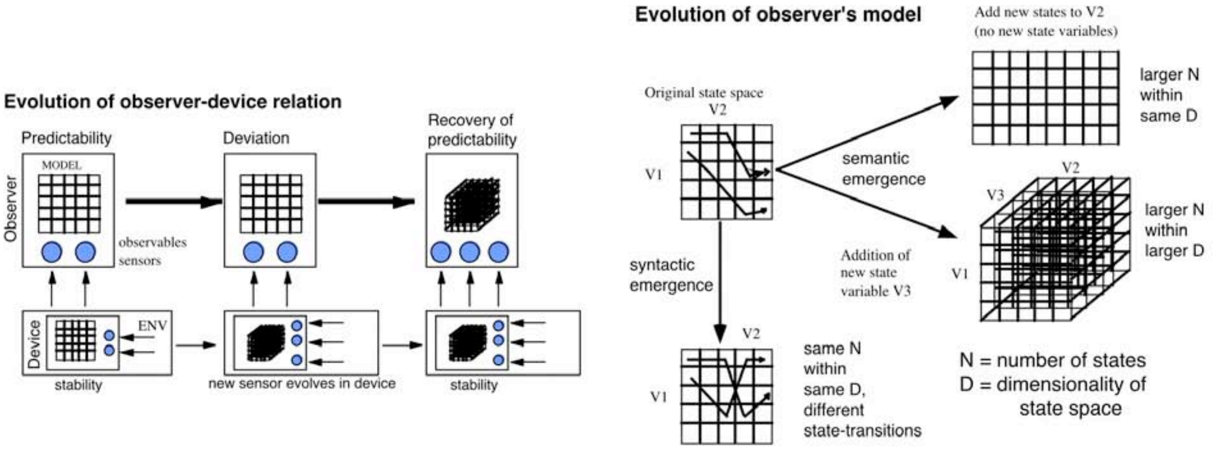

Figure 3. Emergence-relative-to-a-model. Top: Coevolution of a device embedded in an external environment and an observer attempting to predict its behavior. Initially the observer can predict the behavior of the device, but as device changes its internal structure to evolve a new sensor, the device behaves unpredictably relative to the observer. Observer eventually recovers predictability only by adding an observable that covaries with the readings of the device's new sensor. Bottom: Possible routes by which the observer can modify a predictive model when observed behavior deviates from prior expectations. There are three possibilities: 1) Alterations of transitions between existing states, adding more states to existing observables, and adding new observables.

Figure 4 (above) A solution to the methodological problem of recognizing emergence when it occurs. Concepts are only well defined when we have a reliable way of making the distinctions, i.e. in the context of a specificd observational framework.

---

[10] These distinctions originally come from Charles Morris, *Signs, Language, and Behavior* (1946)(Morris, 1946; Nöth, 1990).

independent macro-states and macro-state-transition rules.[11] In the language of dynamical systems, analog processes are described in terms of rate-dependent processes, while digital, discrete macro-states are the macro-reflections of underlying attractor basins. Here new states, sensors, effectors, and goal structures must arise out of rate-dependent processes that create new attractor basins that were not part of the original high-level digital description.

Recognizing combinatoric and emergent creativity (the methodological problem)

How does one distinguish these two forms of creativity operationally? Theoretical biologist Robert Rosen proposed a definition of emergence as the deviation of the behavior of a material system from the behavior predicted by a model of that system. I have operationalized this concept (Figure 4) -- one can distinguish fixed, combinatoric, and creative systems by what one needs to do as an observer in order to successfully "track" the behavior of the system over time, i.e. whether one must add states, new state-transition rules, observables, actions, goals to maintain predictability.

Towards neural networks that create new signals

For roughly two decades I have been contemplating what creation of new signal primitives would mean for neural networks and our theories of brain function. If individual elements can evolve new sensors and effectors, then new signal channels (or modes of signaling) can be created within a network, and the effective dimensionality of the network increases. In the brain we can think of Hebb's neural assemblies (Hebb, 1949; Orbach, 1998) as internal sensors on an analog internal milieu. The creation of a new neural assembly through a activity-dependent modification of neuronal structures (e.g. synapses) is the equivalent to adding a new internal observable on the system.

What follows is an outline of how a system based on these principles might work. I believe that brains require the ability to multiplex neuronal signals and to create new ones. I think it is also possible that the signals themselves are temporal patterns of spikes (a temporal pattern code) rather than the conventional assumption that these are all firing rate codes. The core issue is whether the information resides in temporal patterns of spikes or in "spatial" patterns of activated neural elements (which particular neurons fire how much).

I also want to caution the reader of the highly provisional nature of this working hypothesis. It is a daunting task to propose an alternative brain theory. Although there is some evidence for temporal coding in every sensory system and in many diverse parts of the brain, in contrast to our study of pitch in the auditory nerve, the evidence at the cortical level is scant and generally underwhelming. On the other hand, when carefully scrutinized most rate codes and connectionist accounts explain much less than the headlines would have us believe. The nature of the central neural code is still one of science's biggest unsolved mysteries, a situation akin to biology before DNA-based mechanisms of inheritance were understood.

---

[11] Howard Pattee has written extensively on the complementarity between "linguistic" (symbolic, digital, rate-independent) vs. "dynamical" (analog, rate-dependent) modes of describing physical devices and organisms. See (P. Cariani, 2001b) for discussions of symbols and dynamics in the brain.

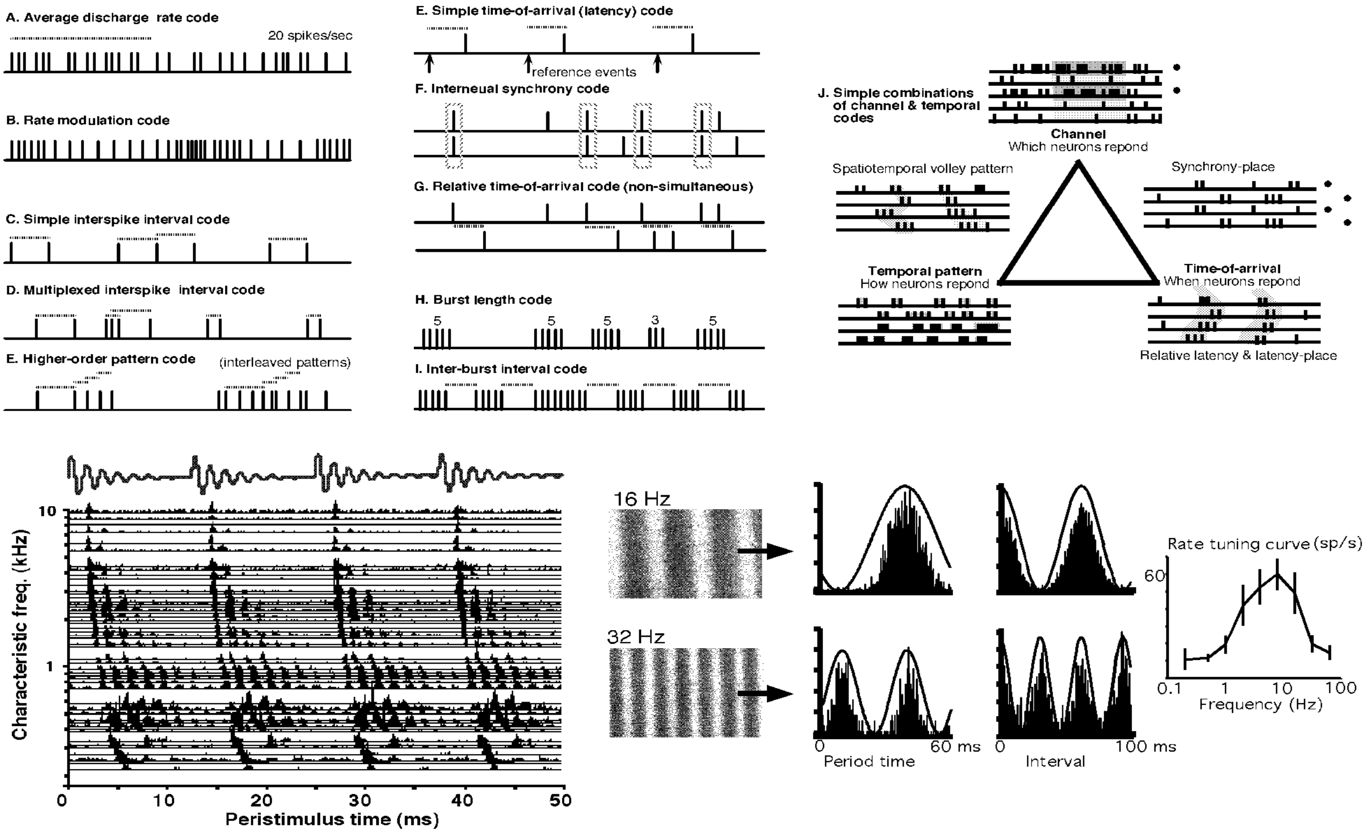

Figure 5. The neural coding problem (from Cariani, 2004).

A pulse coding scheme that supports multidimensional signal types permits each Hebbian neuronal assembly to have its own dedicated signal type. The most obvious vehicle for multidimensional signaling and signal creation lies in temporal coding, where information is conveyed through temporal patterns of spikes rather than through across-neuron patterns of activation (which neurons fire how much). In parallel with these theoretical concerns I have also spent most of the last two decades working on the concrete neurophysiological problems of neural coding of pitch, timbre, and consonance in the auditory system (P. Cariani, 1999; Peter A. Cariani & Delgutte, 1996). Here the early representation of these perceptual attributes is clearly based on temporal codes. Figure 5 above shows examples of different types of possible pulse codes that have been found (upper left), a taxonomy of basic coding types (upper right), and temporal firing patterns from both the auditory nerve in response to a harmonic complex (lower left) and from a thalamic visual unit in response to moving gratings (lower right). Importantly, in the interspike interval representation in the auditory nerve, information is conveyed in population-wide interspike interval statistics that support multiplexing of signals related to different perceptual attributes and correlation-based mechanisms for separating multiple auditory objects (e.g. two notes a third apart played by two instruments at the same time).

The existence of population-based interspike interval representations begs the question of whether information in other parts of the brain could be represented using population-based temporal codes, albeit of different kinds. It also provokes the question of what kinds of neuronal operations and processing architectures might conceivably make use of information in this form. This is still very much an open question in the central auditory system. Time-delay neural nets (e.g. the Jeffress model for binaural localization, Licklider's temporal autocorrelation network for pitch, Braitenberg's cerebellar timing model) are possible, but these ultimately require connectionist networks after time patterns are converted to neuronal activation profiles.

Several years ago, as an alternative to time-delay and connectionist approaches, I proposed neural timing nets that operate entirely in the time domain to process information encoded in temporal patterns of pulses (spikes) (P. Cariani, 2001a; Peter Cariani, 2002; P Cariani, 2005; P. A. Cariani, 2004). In formulating timing nets, I had also been searching for general coding schemes and computational architectures that would permit multiplexing, broadcast,[12] demultiplexing, and elaboration of signals. If one can get beyond scalar signals (e.g. spike counts or firing rates), then what kind of information a given spike train signal contains can be conveyed in its internal structure. On which particular input line the signal arrives is no longer critical to its interpretation, i.e. one now has an information processing system in which signals can be liberated from particular wires. Although one still has definite neuronal pathways and regions where particular kinds of information converge, these schemes obviate the need for the ultra-precise and stable point-to-point synaptic connectivities and transmission paths that purely connectionistic systems require.

In several papers (Cariani, 2001, 2004) I demonstrated in simple simulated coincidence networks how temporally-coded signals can be detected, recognized, and extracted from signal mixtures in a manner that can support multiplexing and demultiplexing of signals. Recurrent timing nets readily build up spike correlation patterns and separate objects on the basis of temporal pattern invariances. These kinds of correlation-based separation strategies obviate the need for explicit feature detections that are then segregated and bound into separate objects. If such alternatives to connectionist networks are at all viable, then we may need to rethink our basic theory of neural networks. Here we will outline how these primitive operations might form the basis for larger networks that could subserve major brain functions.

Design for a brain as a network of adaptive pattern-resonances

Brains are simultaneously communications networks, anticipatory correlation machines, and purposive, semantic engines that analyze their sensory inputs in light of previous experience to organize, direct, and coordinate effective action. Multiplexing permits signals to be combined nondestructively, broadcast, and then demultiplexed by local assemblies that are tuned to receive them. The brain can thus be reconceptualized, from the connectionist notion of a massive switchboard or telegraph network to something more like a radio broadcast network or even an internet (John, 1972). We believe that these basic functionalities can be implemented by global architectures that consist of many reciprocally connected neuronal populations. The loops support pattern amplification cycles of neuronal signals that allow signals to be dynamically regenerated. Ability of a system to regenerate alternative sets of neural signals means that the system can support alternative persistent informational states. This is critical for working short term memory, informational integration in global workspaces (Baars, 1988) and possibly also conscious awareness (P. Cariani, 2000). Stable, persistent mental states are formed from self-stabilized regenerative signal productions (complex signal attractor states). These states constitute neural "pattern resonances" or an autopoiesis of neural signal productions. Some micropatterns of neural activity are self-reinforcing, while others are self-extinguishing. Neural resonances form in re-entrant pathways: in thalamocortical, cortico-cortical, hippocampal, and subcortical loops.

---

[12] In his seminal paper on adaptive timing networks , D. M. MacKay (1961) pointed out that temporal patterns could be broadcast, a coordinative strategy he called "the advertising principle."

Sensory information comes into the system through modality-specific sensory pathways. Neural sensory representations are built up through basic informational operations that integrate information in time by establishing circulating patterns which are continuously cross-correlated with incoming ones (i.e. bottom-up/top-down interactions). When subsequent sensory patterns are similar to previous ones, these patterns are amplified and stabilized; when they diverge, new dynamically-created “templates” are formed from the difference between expectation and input ("mismatch negativities"). Tuned neural assemblies can thus provide top-down facilitation or inhibition of particular temporal patterns by adding them to circulating signals.

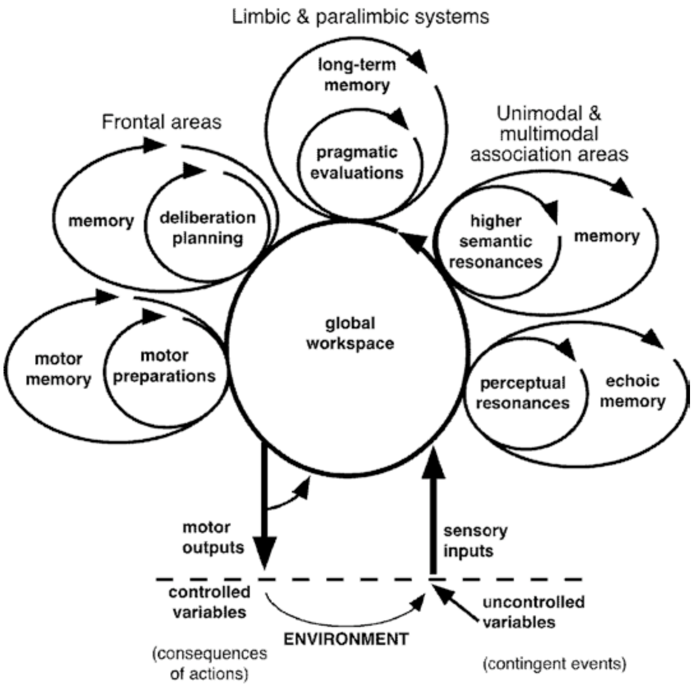

Figure 6. Functional schematic of the brain as a set of pattern-resonances coupled to an external environment. Each loop represents a pattern-resonance amplification-selection process. From (P. Cariani, 1997).

In addition to stimulus-driven temporal patterns, stimulus-triggered endogenous patterns can be evoked by conditioned neural assemblies (John, 1967; Morrell, 1967) (see also the "cognitive timing nodes" of (MacKay, 1987)). Coherent temporal, spatially-distributed and statistical orders (“hyperneurons”) consisting of stimulus-driven and stimulus-triggered patterns have been proposed as neural substrates for global mental states ((John, 1988)). Stimulus-driven patterns encode the stimulus, while stimulus triggered patterns reflect its meaning, semantic and pragmatic, for the organism. The stimulus-triggered endogenous patterns can function as higher-level annotative "tags" that are added to a signal to indicate that it has a particular cognitive attribute. Neural assemblies could be adaptively tuned to emit new tag patterns that would mark novel combinations of perceptual, cognitive, conative, and mnemonic activation. New tags would constitute new symbolic primitives (P. Cariani, 1997) and new concepts. This may be a means by which new dedicated "perceptual symbols" can be formed from semantically and pragmatically meaningful iconic sensory representations (Barsalou, 1999).

Signal multiplexing, nondestructive superposition, and nonlocal storage of information permit broadcast strategies of neural integration. The global interconnectedness of cortical and subcortical structures permits widespread sharing of information that has built-up to some minimal threshold of global relevance, in effect creating a “global workspace” (Baars, 1988; Dehaene & Naccache, 2001). The contents of such a global workspace would become successively elaborated, with successive sets of neurons contributing their own annotations to the circulating pattern in the form of characteristic pattern-triggered signal-tags. Such tags could then be added on to the evolving global pattern as unique indicators of conjunctions of meaningful events (higher-order associations).

Traditionally, the brain has been conceived in terms of sequential hierarchies of decision processes, where signals represent successively more abstract aspects of a situation. As one moves to higher and higher centers, information about low-level properties is presumed to be discarded. A tag system, on the other hand, elaborates rather than reduces, continually adding additional annotative dimensions.

Depending upon attentional and motivational factors, such a system would distribute relevant information over wider and wider neural populations. Rather than a feed-forward hierarchy of feature-detections and narrowing decision-trees, a system based on signal-tags would more resemble a heterarchy of correlational pattern-amplifiers (a la Lashley, (Orbach, 1998)) in which neural signals are competitively facilitated, stabilized, and broadcast to produce one dominant, elaborated pattern that ultimately steers the behavior of the whole. There would then be bi-directional influence between emergent global population-statistical patterns and those of local neural populations.

Neural signal tags characteristic of a given neural assembly would signify that a particular assembly had been activated, be it characteristic of a sensory, motor, executive, or limbic generator. Tags produced by sensory association cortical areas would connote sensory attributes and conjunctions; those produced by limbic circuits would indicate hedonic, motivational, and emotive valences, such that these neural signal patterns would bear pragmatic content. Linkages between particular sensory patterns and motivational evaluations could be formed that add tags related to previous reward or punishment history, thereby adding to a sensory pattern a hedonic marker. In this way, these complex, elaborated neural signal productions could be imbued with pragmatic meanings ("intentionality") which could be conferred on sensory representations that have causal linkages with the external world ("intensionality").

Neural signal tags with different characteristics could thus differentiate patterns that encode the syntactic, semantic, and pragmatic aspects of an elaborated neural activity pattern. In the wake of an action that had particular hedonic salience, associations between all such co-occurring tags would be continuously rebroadcast by hippocampal circuits and consolidated in long term memory. The system would thus build up learned expectations of the manifold hedonic consequences of percepts and actions. When similar circumstances presented themselves, temporal memory traces containing all of these consequences would be read out faster-than-real-time to facilitate or inhibit particular action alternatives, depending upon whether percept-action sequences in past experience had resulted in pleasure or pain. Such a reward prediction system, which computes conditional probabilities weighted by hedonic relevance, is capable of rapid and flexible learning. Dopamine-mediated circuits deep in our brains play this role of anticipatory reward prediction and the basal ganglia appear to steer both attention and behavior contingent on their hedonic predictions.

Although this framework is bound to be incorrect in many respects, it does give us a way of thinking about the generation of new primitives in neural networks and hence also a way of thinking more concretely about how new concepts might arise in our brains. At this point, the theory is highly qualitative, somewhere between metaphor and model. It can serve both to generate hypotheses about the neural basis of new concepts and to provide heuristics for designing artificial, mixed analog-digital self-organizing systems that create new internal signal primitives in an apparently open-ended fashion.

References

Baars, B. J. (1988). *A Cognitive Theory of Consciousness*. Cambridge: Cambridge University Press.

Barsalou, L. W. (1999). Perceptual symbol systems. *Behavioral and Brain Sciences, 22*, 577-660.

Bird, J., & Di Paolo, E. (2008). Gordon Pask and his Maverick Machines. In P. Husbands, O. Holland & M. Wheeler (Eds.), *The Mechanical Mind in History* (pp. 185-211). Cambridge, Mass.: MIT Press.

Boden, M. (1990a). *The Creative Mind*. New York: Basic Books.

Boden, M. (1990b). What is creativity? *Dimensions of Creativity* (pp. 75-117). New York: MIT Press.

Boden, M. A. (1991). *The creative mind : myths & mechanisms*. New York, N.Y.: Basic Books.

Boden, M. A. (1994). *Dimensions of creativity*. Cambridge, Mass.: MIT Press.

Boden, M. A. (2006). *Mind as machine : a history of cognitive science*. New York: Oxford University Press.

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

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

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

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

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

Cariani, P. (1999). Temporal coding of periodicity pitch in the auditory system: an overview. *Neural Plasticity, 6*(4), 147-172.

Cariani, P. (2000). Regenerative process in life and mind. In J. L. R. Chandler & G. Van de Vijver (Eds.), *Closure: Emergent Organizations and their Dynamics* (Vol. 901, pp. 26-34). New York: Annals of the New York Academy of Sciences.

Cariani, P. (2001a). Neural timing nets. *Neural Networks, 14*(6-7), 737-753.

Cariani, P. (2001b). Symbols and dynamics in the brain. *Biosystems, 60*(1-3), 59-83.

Cariani, P. (2002). Temporal codes, timing nets, and music perception. *J. New Music Res., 30*(2), 107-136.

Cariani, P. (2005). Recurrent timing nets for F0-based speaker separation. In P. Divenyi (Ed.), *Perspectives on Speech Separation* (pp. 31-53). New York: Kluwer Academic Publishers.

Cariani, P. A. (2004). Temporal codes and computations for sensory representation and scene analysis. *IEEE Trans Neural Netw, 15*(5), 1100-1111.

Cariani, P. A., & Delgutte, B. (1996). Neural correlates of the pitch of complex tones. I. Pitch and pitch salience. II. Pitch shift, pitch ambiguity, phase-invariance, pitch circularity, and the dominance region for pitch. *J. Neurophysiology, 76*(3), 1698-1734.

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

Dehaene, S., & Naccache, L. (2001). Towards a cognitive neuroscience of consciousness: basic evidence and a workspace framework. *Cognition, 79*(1-2), 1-37.

Hebb, D. O. (1949). *The Organization of Behavior*. New York: Simon & Schuster.

John, E. R. (1967). Electrophysiological studies of conditioning. In G. C. Quarton, T. Melnechuk & F. O. Schmitt (Eds.), *The Neurosciences: A Study Program* (pp. 690-704). New York: Rockefeller University Press.

John, E. R. (1972). Switchboard vs. statistical theories of learning and memory. *Science, 177*, 850-864.

John, E. R. (1988). Resonating fields in the brain and the hyperneuron. In E. Basar (Ed.), *Dynamics of Sensory and Cognitive Processing by the Brain* (pp. 368-377). Berlin: Springer-Verlag.

MacKay, D. G. (1987). *The Organization of Perception and Action*. New York: Springer-Verlag.

Morrell, F. (1967). Electrical signs of sensory coding. In G. C. Quarton, T. Melnechuck & F. O. Schmitt (Eds.), *The Neurosciences: A Study Program* (pp. 452-469). New York: Rockefeller University Press.

Morris, C. (1946). *Signs, Language, and Behavior*. New York: George Braziller.

Nöth, W. (1990). *Handbook of Semiotics*. Indianapolis: Indiana University Press.

Orbach, J. (1998). *The Neuropsychological Theories of Lashley and Hebb*. Lanham: University Press of America.

Pask, G. (1958). *Physical analogues to the growth of a concept.* Paper presented at the Mechanization of Thought Processes, Symposium 10, National Physical Laboratory, November 24-27, 1958.

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

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

Piatelli-Palmarini, M. (1980). How hard is the hard core of a scientific paradigm? In M. Piatelli-Palmarini (Ed.), *Language and Learning. The Debate between Jean Piaget and Noam Chomsky*. Cambridge, MA: Harvard University Press.