Intellectica, 1997/2, 25, pp. 95-143

Peter CARIANI[*]

Emergence of new signal-primitives in neural systems

Emergence is the process by which new structures and functions come into being. There are two fundamental, but complementary, conceptions of emergence: *combinatoric emergence,* wherein novelty arises by new combinations of pre-existing elements, and *creative emergence,* wherein novelty arises by de novo creation of new kinds of elements. Combinatoric emergence is exemplified by new strings constructed from existing alphabetic letters, whereas creative emergence is exemplified by the addition of new kinds of letters to an alphabet. The two conceptions are complementary, providing two modes for describing and understanding change: as the unfolding consequences of a fixed set of rules or as new processes and interactions that come into play over time.

Within an observer-centered, operational framework, the two kinds of emergent novelty can be distinguished by what an external observer must do in order to successfully predict the behavior of an evolving system. Combinatoric and creative emergence can be operationally distinguished by changes in apparent effective dimensionality. Whenever a new independent observable is added to a model, its dimensionality increases by one. A system that only recombines requires no new observables, and does not expand in effective dimension. In contrast, a system that creates new primitives requires new observables for its description, such that its apparent dimensionality increases over time.

Dimensional analysis can be applied to signaling systems. Signals have two basic functional properties: signal-type (category, variable, type) and signal-value (state, value, token). These properties can be conveyed by a variety of means: by the signal's physical channel, by the internal form of the signal (waveform, Fourier spectrum), by its time of arrival, and by its magnitude (average power). Neural coding schemes can similarly be based on which neurons fire, which temporal patterns of spikes are produced, when volleys of spikes arrive, or how many spikes are produced. Traditional connectionist networks are discussed in terms of their assumptions about signal-roles and neural codes. For the most part, connectionist networks

---

[*] *Eaton Peabody Laboratory for Auditory Physiology - Massachusetts Eye & Ear Infirmary, 243 Charles St, Boston MA 02114 USA*
tel (617) 573-4243; FAX (617) 720-4408; E-mail: peter@epl.meei.harvard.edu

are conceptualized in terms of new linkage combinations rather than in terms of new types of signals being created.

Neural networks that increase their effective dimensionalities can be envisioned. Some kinds of neural codes, such as temporal pattern and time-of-arrival codes, permit encoding and transmission of multidimensional information by the same elements (multiplexing). We outline how synchronous time-division and asynchronous code-division multiplexing might be realized in neural pulse codes. Multidimensional temporal codes permit different kinds of information to be encoded in different time patterns. Broadcast-based coordination strategies that obviate the need for precise, specified point-to-point connections are then made possible.

In such systems new signal types arise from temporal interactions between time-coded signals, without necessarily forming new connections. Pitches of complex tones are given as examples of temporally-coded, emergent Gestalts that can be seen either as the sums of constituent micro-patterns (combinatoric emergence) or as the creation of new ones. Within these temporally-coded systems, interacting sets of neural assemblies might ramify existing, circulating signals to construct new kinds of signal primitives in an apparently open-ended manner.

Key words: combinatoric and creative emergence; dimensionality; observables; signaling systems; temporal coding; neural assemblies.

L'émergence de nouveaux signaux primitifs dans les systèmes neuronaux. L'émergence est le processus par lequel de nouvelles structures et fonctions se mettent à exister. Il y a deux conceptions fondamentales de l'émergence, qui sont en fait complémentaires : l'émergence combinatoire, où la nouveauté provient de nouvelles combinaisons d'éléments préexistants ; et l'émergence créatrice, où la nouveauté naît de la création de nouveaux genres d'éléments. Ces deux conceptions sont complémentaires en ce qu'elles procurent deux manières de décrire et de comprendre le changement : comme le déploiement des conséquences d'un ensemble préétabli de règles d'un côté ; et comme de nouveaux procédés et interactions qui entrent en jeu au cours du temps, de l'autre.

Emergence combinatoire et émergence créatrice peuvent être distinguées de manière opérationnelle par des changements dans la dimensionalité effective telle qu'elle apparaît à un observateur. Un système qui se contente de recombiner ne requiert aucun élément observable nouveau, et sa dimension effective n'augmente pas. Par contre, un système qui crée de nouveaux éléments primitifs requiert de nouveaux observables pour sa description, de telle manière que sa dimensionalité s'accroît au cours du temps.

L'analyse dimensionnelle peut être appliquée aux systèmes de signaux. Les réseaux connexionistes sont, le plus souvent, conceptualisés en termes de combinaisons de nouveaux cablages plutôt qu'en termes de création de nouveaux types de signaux. Il est cependant possible de concevoir des réseaux neuronaux qui augmentent leur dimensionalité effective. Certains codes neuronaux, comme les motifs temporels et les codes de temps d'arrivée, permettent l'encodage et la transmission d'une information multidimensionnelle par les mêmes éléments. Dans le cadre de ces systèmes temporellement codés, des séries d'assemblages neuronaux en interaction peuvent s'appuyer sur des signaux existants de manière à construire des ramifications de nouveaux genres de signaux primitifs, et ce indéfiniment semble-t-il.

Mots-clés : émergence combinatoire et émergence créatrice ; dimensionalité ; éléments observables ; systèmes de signaux ; codage temporel ; assemblées de neurones

TWO BASIC CONCEPTIONS OF EMERGENCE

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.[1] These images provide explanations for how novelty, spontaneity, and creativity are possible. They provide us with accounts of how complex organizations arise and become further elaborated, be they physical, chemical, biological, psychological, or social formations.

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)[2], new formal structures (computational

---

[1] *The domain of emergence is qualitative, unpredicted change. Obviously, all is not change, and all change does not involve novelty, so that emergentist accounts are but one mode of explanation among many. Emergence does not explain why things remain the same or recur in familiar patterns or why things have the basic properties that they do, e.g. why there should be gravity or consciousness.*

[2] *(Denbigh, 1975; Prigogine, 1980)*

emergence), new biological structures and functions (emergent evolution)[3], new scientific theories (emergence vs. reduction)[4], new modeling relations in observers,[5] new percepts,[6] new ideas[7], new notational systems, and new economic and social relations.[8] Despite very different subject matters, these discussions share many formative concepts and internal debates.

---

[3] *(Bergson, 1911; Morgan, 1931)*

[4] *(Nagel, 1961)*

[5] *(Rosen, 1978; Rosen, 1985; Kampis, 1988; Cariani, 1989b; Kampis, 1991b; Rosen, 1991; Cariani, 1992a)*

[6] *See essays on perceptual learning in (Gibson, 1969; Gibson, 1991).*

[7] *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). (van de Vijver, 1991) discusses emergence in the context of connectionist representations. Issues related to the formation of new cognitive primitives are currently under discussion in (Schyns, Goldstone and Thibaut, in press).*

[8] *The emergence of new social structures has been discussed in terms of structural-evolutionary processes (J. Habermas), andin terms of self-organizing systems and spontaneous orders (F. Hayek).*

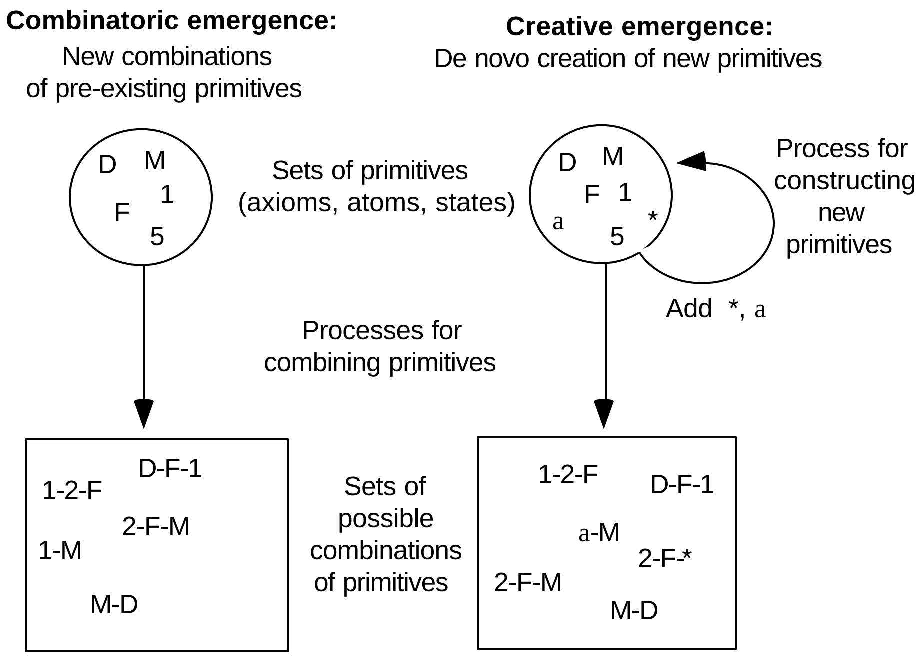

Figure 1. Combinatoric emergence and creative emergence.

Two fundamental conceptions of emergence can be distinguished: *combinatoric emergence* and *creative emergence*. These two accounts of the origin of novelty parallel notions of the origin of order: "order-from-order" vs. "order-from-noise."[9] 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

---

[9] *(Maruyama, 1977; Piatelli-Palmarini, 1980)*

from combinations of the other entities, i.e. its properties cannot be logically deduced from those of other entities.

NOVEL COMBINATIONS WITHIN SETS OF EXISTING 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 is seen in terms of novel combinations of pre-existing genes. Likewise, new, 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.[10]

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

---

[10] *See (Klee, 1984; van de Vijver, 1991) for discussions of micro-macro relations and emergence. The alternative to microscopic rules unilaterally producing macroscopic orders is to envision mechanisms by which macroscopic orders can alter microscopic interactions (e.g. a lizard crawling onto a sunny rock to alter its body temperature, thereby changing the boundary conditions for its own microscopic biomolecular processes). This bidirectional causation is a means by which qualitatively new kinds of structures and functions, such as new control structures, can be created (Pattee, 1973a; Pattee, 1973b; Rosen, 1973). Combinatoric emergence assumes exhaustive descriptions of micro-processes, whereas creative emergence invokes previously undescribed interactions that come into play once aggregates form.*

(feature-action mappings). What formally distinguishes different kinds of trainable machines, e.g. neural networks, genetic algorithms, Boltzmann machines, are the structures of the respective combination-spaces being traversed, and the rules that direct the search processes. In the realm of scientific models, the observables of the experimental scientist correspond to the feature-primitives of trainable classifiers, while the formal computations that generate predictions from initial conditions correspond to decision rules.[11] Once a set of measuring devices and their associated observables is specified, then different formal predictive algorithms can be generated and tested.[12]

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 sets of primitive elements. 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. Various related arguments concerning problems of predictability, complexity, emergence, and closure in computer simulations have been summarized, discussed and debated in depth elsewhere.[13]

---

[11] *For discussions of the operational structure of scientific models (Hertzian commutation diagrams), see (Cassirer, 1955; Rosen, 1978; Rosen, 1985; Kampis, 1988; Cariani, 1989b; Kampis, 1991b; Cariani, 1992a).*

[12] *The realm of combinatorial search is the formal, computational part of the process. Given a finite set of observables with discrete, distinguishable states, general systems theory outlines what formal models (what parameters and parameter-combinations) are possible. Of course, if adequate predictions cannot be obtained with a given the set of observables, then the scientist must go outside of the model to change its assumptions, by altering observables and/or formal relations.*

[13] *(Cariani, 1989a; Pattee, 1989; Kampis, 1991a; Cariani, 1992a; Emmeche, 1994; Bonabeau, Desalles and Grumbach, 1995; Bonabeau and Theraulaz, 1995; Steels, 1995; Clark, 1996) The more platonically-minded counter-arguments assert that computer*

While computationalists in artificial intelligence and artificial life have correctly recognized the generative power of combinatoric novelty, they have generally failed to recognize the need for processes that create new primitives.[14] The informational and economic power of the digital computer is constantly increasing by leaps and bounds, so it is entirely natural to look to more powerful machines and complex programs as generators of emergent functionality. Powerful techniques rapidly become totalizing world views that redefine the world in their image, and the digital computer is no exception. The dramatic appearance of formal theories of computation in the 1930's swiftly followed by the development of digital electronic computers in subsequent decades effectively revived platonic ideas, eventually giving birth to universalist computationalist ideologies with broad annexationist aspirations.[15]

---

simulations can be made potentially-infinite in one way or another. Since little can be proven concerning what is possible or impossible in infinite systems (e.g. Godel's Undecidability Theorems), one can make the argument that such systems might not be closed systems. However, indefinitely expanding systems are not physically realizable – once we invoke potential-infinities we are no longer in the realm of logics or computer simulations that can be actually implemented, but in the universe of platonic objects. There is a qualitative gap between the very-large-but -finite, and the potentially-infinite. Other break-out strategies are more incrementalist: adding logical depth, increasing string lengths, growing automata, creating meta-rules. I would argue that these moves only delay the problem, that while new combinatorial states are created this way, still no new primitive categories are added, hence the dimensionality of possibility-spaces does not increase. They are akin to adding new distinctions within existing observables (Figure 2, bottom panel). These maneuvers by themselves do not obviate the need for primitive-creation and dimensional increase.

[14] *e.g. (Boden, 1990a; Boden, 1990b; Kauffman, 1993; Holland, 1998). It is invariably assumed that hill-climbing take places on hills that don't change their shapes as populations climb them. In evolutionary landscapes, dimensional increase means that new factors come into play over time. The modes of interaction between organisms can depend on their perceptual capabilities, and these in turn shape the adaptive landscape. Percept and action repertoires shape "effective connectivities" between organisms. An adaptive landscape where pigments and color vision have evolved has a higher effective dimensionality than an otherwise similar world without color. This is related to Conrad's notions of "bootstrapping the adaptive landscape"(Conrad, 1983) and "extradimensional bypass"(Chen and Conrad, 1994).*

[15] *Many of the tenets of the computationalist worldview coalesced at the 1956 Dartmouth workshop on computation at which the disciplines of cognitive science and artificial intelligence were born (McCorduck, 1972). From the mid-1950's onward, computationalists argued that discrete digital computations could* in principle *simulate -- and presumably supplant -- any material process. This led to the burying of the analog-digital and hardware-software distinctions, and occurred in parallel with the discarding of distinctions between analytic and empirical truths in philosophy (logic vs. observation, computation vs. measurement). In AI and later, in artificial life, neo-Laplacian computational microdeterminisms could always be postulated as "physical"*

It is imperative to keep the role of the formal – mathematics, logic, and formal computation – in perspective. While formal systems are effective, almost indispensable tools of thought and practice, this does not necessarily make them ideal universal prototypes for the structure of mind and matter. In focusing on the generative power of computational combinatorics, it is easy to overlook the critical, formative role of the human designer in setting up formal notations in the first place, and in expanding them when needed. A wise practitioner understands the capabilities and limitations of the tools that s/he uses; a practitioner with vision understands the dangers of limiting the imagination to one kind of tool. In the focus on recombination, the temptation is great to truncate conceptions of novelty, creativity, and emergence to pure recombination and evaluation. These notions then are defined in terms that exactly fit the capabilities of digital computers.[16] This Procrustean approach to creativity chops off those capacities that distinguish human beings from digital computers: the ability to autonomously interact with the external, material world through sensing and acting, and in doing so, to formulate new conceptual primitives in an open-ended manner.[17]

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 the formation of entirely novel

---

descriptions for any problem (see critiques by (Rosen, 1991; Rosen, 1994; Pattee, 1995)), despite the difficulties of ever testing these hypotheses experimentally. The early conceptual developments paved the way for hegemony of platonic, model-theoretic logic systems in the cognitive sciences that had little interest in how meaning might be embedded in real brains such that it could be formed through material interaction with the external world via perception and action. It goes to show that, for better or worse, every sufficiently powerful intellectual revolution inevitably produces its own peculiar "excesses." The computer revolution is no exception.

[16] *Reductions of creativity to discrete generative systems, e.g. (Boden 1990a; Boden 1990b; Holland 1998), carry out such truncations, and in doing so, unecessarily narrow our vision.*

[17] *Questions of semantic closure of computational systems have been discussed elsewhere (Dreyfus, 1979; Carello* et al.*, 1984; Cariani, 1989b; van de Vijver, 1991; Cariani, 1992a). Rather than redefine creativity to accomodate what computers can do, as an alternative, we could rethink what kinds of artefacts are possible that might embody alternative notions of fundamental creativity. Rather than designing devices that operate on our prespecified categories, devices need to have their own means of adaptively altering their internal structure, independent of their designers. They must be capable of constructing their own categorical primitives, their own relevance criteria, in order to attain full epistemic autonomy (Cariani, 1992b; Cariani, 1992a; Cariani, 1993). The answer to the 'frame problem' (Dreyfus, 1979; Dreyfus and Dreyfus, 1987) is to give the devices the means to adaptively determine what is contextually-relevant, e.g. by constructing their own sensory linkages with the world.*

ideas. Although computers are but externalized extensions of our minds, human-machine collaborations taken as a whole system can manifest creative emergent functionalities that cannot be seen in computers alone.

EMERGENCE AS 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.[18] Primitive-creation is also an issue for combinatoric, analytic-deductive conceptions of emergence – one can always ask how the particular primitives of an existing combinatorial system came into being in the first place. Combinatoric-emergence assumes a "God's-eye view" of the objects that inhabit its universe, but questions of the origins of this universe lie outside its realm.[19] Creative emergence, on the other hand, adopts 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.

---

[18] *In a very similar vein (Morgan, 1931) distinguished "emergents" from "resultants". Emergents are the result of novel creation, resultants, of novel combination.*

[19] *Questions concerning the ultimate origins of gods and numbers are usually avoided entirely by theologians and mathematicians.*

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 obervables 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.[20] 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.

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.[21] 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 a 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. The artificial device automated, in a very rudimentary way, the creation of new sensory primitives, thereby providing an existence proof that creative emergence is possible in adaptive devices.

OPERATIONALIZING DEFINITIONS OF EMERGENCE

Practically, how does one distinguish combinatoric emergence from creative emergence? How do we know when a new primitive has been created? An operational definition is needed to make the concept

---

[20] *i.e. the observable is not commensurable with the others.*

[21] *(Pask, 1958; Pask, 1959; Pask, 1960; Pask, 1961; Cariani, 1993)*

explicit.[22] The theoretical biologist Robert Rosen has 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. (Figure 2). Rosen’s formal definition can be operationalized. An observer watches the behavior of a material system over time, noting the sequences of observed-states that the material system traverses (Figure 2, top). The observer begins to form expectations about what states the material system will subsequently pass through given that it is presently in some particular state. This set of expectations forms the observer's model of the material system.

---

[22] *An operational definition specifies, in a step-by-step manner, how one carries out a procedure. An operational definition clarifies conceptual distinctions by giving a method that others can use to unambiguously classify objects, events, or states-of-affairs. We understand the definition of something well only when we can provide an explicit method by which others can make the same observations and classifications that we do. Operational definitions avoid endless confusion over what is "emergent." and what is not. There is an unfortunate tendency in too many contemporary discussions to rely on intuitive understandings of "emergent computation", “life”, and “complexity’ without ever attempting to provide rigorous definitions.*

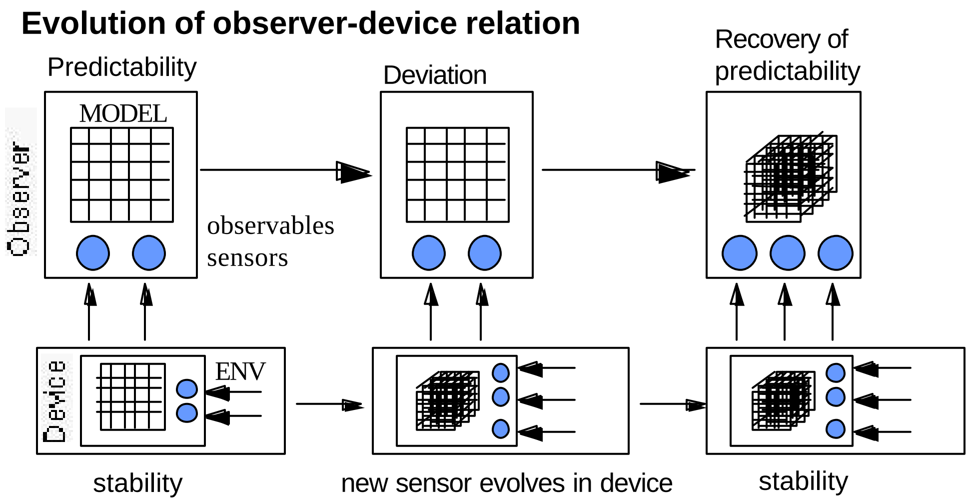

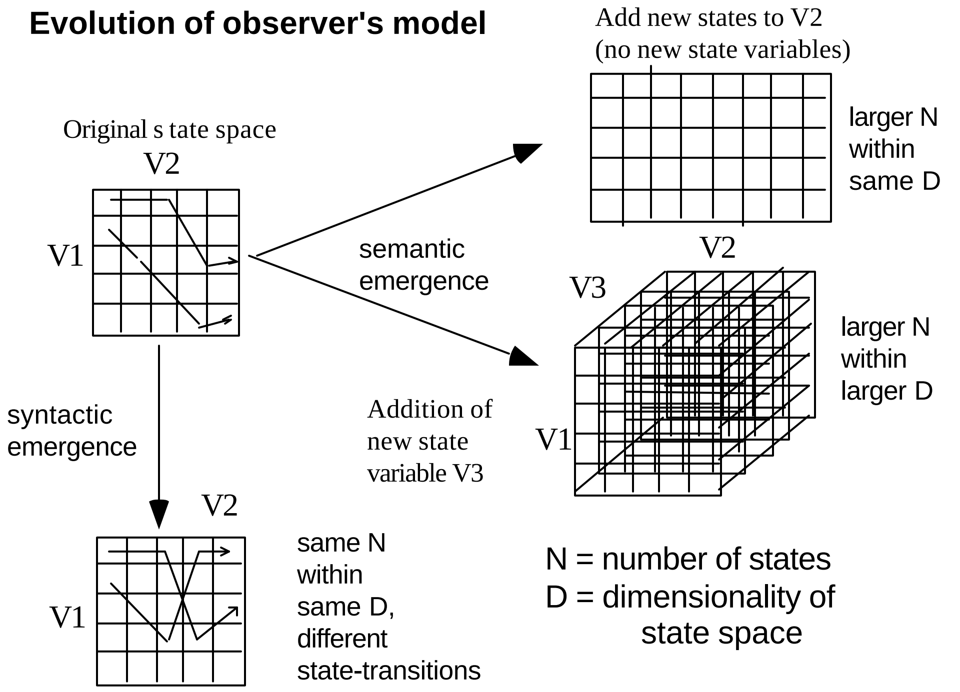

Figure 2. Emergence-relative-to-a-model. Top: Coevolution of a device embedded in an external environment and an observer attempting to predict its behavior. Initially the observer can predict the behavior of the device, but as device changes its internal structure to evolve a new sensor, the device 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.

A full description of these expectations includes the measuring devices and experimental conditions that define the observable-states of the material system. Thus, the model tells us that if we look at the material system in a particular way, we should expect to see the specified state-transition behavior. In effect the model of the system makes explicit the expectations of the observer, such that an emergent event takes place when those expectations are violated.[23]

This system-theoretic definition of emergence ("emergence-relative-to-a-model") can be operationalized and applied to various kinds of natural systems and artificial devices to determine whether a given material system has emergent behavior. An emergent event involves a structural change in the relationship between the observer and the physical system under observation. Like any other measurement,the detection of an emergent event is a joint property of both observer and system. If we are observing an adaptive device, we assume that we initially have an adequate model of the behavior of the device. Over time as the device adapts (learns from experience), its internal structure changes and consequently its behavior changes. As a consequence, if we are to continue to predict its behavior, we then must change our model in some way to "track" the internal changes in the device. Observer and device then evolve in parallel. The observer infers what is going on in the device from what changes the observer finds he or she must make in order to continue to account for the device's behavior.

When the device's behavior deviates from the observer’s model, the observer has two complementary strategies for changing the model (Fig. 2, bottom). The observer can either modify the model’s predictive algorithm or s/he can change the observables used to measure the states of the device and its environment. If the device has merely changed its internal computational structure, but is still using the same set of operational, internal states as before, then modifications to the predictive algorithm will suffice to track the device. No observables need be changed or added. On the other hand, if the device suddenly grows a sensor that detects some new property of its environment (as in Pask’s device), then the observer must add a sensor (i.e. another observable whose readings are roughly congruent with those of the device’s emergent sensor) in order to track the device's behavior.

---

[23] *Outside of purely formal realms, operational definition requires a specification of the observational frame, since replicability of result across observers depends on the holding the observational context constant. Experimental scientists who deal with messy, ill-defined real-world systems are familiar and comfortable with the necessity of interpreting observations taking into account the method by which they were made.*

Intermediate between these scenarios is the case when the device refines an existing sensor to make finer discriminations of the same observable (e.g. refining a temperature sensor so that it can make 1 degree disriminations rather than 5 degree ones). Here the observer similarly must refine the corresponding observable.

These different "tracking" scenarios produce models with different dimensionalities. In the formal, computational part of the model we have one state variable for each independent observable, and the dimensionality of the "model space" ("phase space") represented is equal to the number of state variables. When we keep the same number of observables (new computation), the dimensionality of the model space stays constant. When we need to add another observable (new sensor), the dimensionality of the model space increases by one. In this case it can be said that the "effective dimensionality" of the device (relative to us) has increased as the device's internal structure evolves. We will assume that our initial observables in some sense capture the behavior of the primitives of the system being observed, and that this rough congruence is the reason that our model initially made successful predictions in the first place.

Combinatoric emergence occurs when the behavior of the device changes but no new observables are needed to regain predictability. From this we infer that only recombination of primitives is occurring, since our own model needs no new observables or system variables to explain the behavior of the system. The old framework will do, albeit with some adjustments. Effective dimensionality is unchanged.

Creative emergence occurs when new observables are required — and from this we infer that new primitives have been created. Effective dimensionality increases. Thus under a given set of observables (or observational frame or informational frame) we can unambiguously determine whether we are witnessing emergence-as-new-combination or emergence-as-de novo-creation. It should always be kept in mind, however, that this decision is relative to the observer's frame, so that it is in no sense an absolute distinction. In this framework, emergence lies in the changing relationship of the observer to the material system that is observed.

DIMENSIONAL ANALYSIS OF SIGNALING NETWORKS

The dimensionality of the descriptive notation needed to adequately capture the behavior of a system thus informs us of the structure of the space of possibilities. It tells us how many categorical primitives, how many independent state-variables are in play at any given time. It tells us how many different properties or observables are needed to describe our system. In our notations, each kind of categorical primitive has its own dimensional label so that independent categories can be kept distinct. The formal mechanics of categorical primitives constitute “dimensional analysis.”[24] Take any physics or chemistry problem and examine how the values of the different state-variables are manipulated. The different unit-dimensions function to enforce commensurability relations between the variables. One has to keep different "units of measurement" separate, and these unit-dimensions must accompany all (scalar) magnitudes as one proceeds through the computations. In physics, variables with units of mass, velocity, and time must all be combined only in particular ways. Adding magnitudes for mass and time is like adding apples and elephants – the resulting dimensions may not be meaningful. In the stochiometric equations that describe chemical equilibria, each chemical species has its own unit-category that accompanies its numeric concentration value (e.g. 100 nM-[Ca^{++}]) . The unit-dimensions preserve the semantic coherence of the computations and in doing so improve both the syntactic and semantic reliability of the symbol manipulations.[25] The same advantages hold for typed logics over untyped or "flat" ones: logical types guarantee coherent manipulation of sets of entities whose categories may or may not be commensurable.[26]

What does dimensionality mean for a system as complex as the brain? For almost a century neuroscientists have analyzed electrical signals from the brain in hopes of understanding the underlying neural processes. Electroencephalographic (EEG) recordings, electrical potentials from the scalp of awake humans, reflect in a rough way the

---

[24] *(Bridgman, 1931)*

[25] *Combining magnitudes of different units willy-nilly can easily create syntactically-correct, but semantically-incoherent computational results. At the end of a computation, if all unit-dimension relations have been preserved, one knows what the numerical result means by its accompanying unit-dimension. One can also detect a computational error if at the end one is left with unexpected or meaningless dimensions. Similar kinds of relations arise in computer software design.*

[26] *Any typed logic can be expressed in an untyped one, in the same way that, given enough states, any higher dimensional state-space can be mapped onto one-dimension. The cost of fewer dimensions lies in the complexity of book-keeping operations needed to keep mutually-exclusive states from occuring. All information, whatever its dimensionality, can be recoded into ever-longer strings of 1's and 0's. However, the greater the dimensional reductions, the more brittle encodings become, syntactically and semantically, in the face of minor bit-errors.*

electrical activities of large populations of neurons.[27] In early EEG studies, it was found that when the brain was at rest, in a relaxed state, there were relatively fewer periodicities present, with more large-scale synchronized behavior.[28] When sensory stimuli were presented or when subjects were asked to carry out some mental task, more complex, less synchronized patterns were observed. Intuitively, one might conclude that the complexity of the EEG patterns reflects in some crude way the complexity of the underlying neural processing. In the last decade, some theorists have attempted to estimate this complexity in terms of "effective dimensionality", i.e. the number of independent dimensions needed for a dynamical system model to account for the observed electrical behavior.[29] This measure of the degrees of freedom available to the system is also called its "correlation dimension." For example, the lowest correlation dimensions in EEGs are found for sleep states (roughly 4-6), with higher ones for awake, resting subjects with eyes closed (roughly 10) and still higher values when eyes are opened.[30]

In many ways analyzing EEG signals is akin to trying to deduce the workings of a computer from sets of voltage readings taken on the outside case. The residual voltages that one sees have something to do with the processing that is taking place inside the machine, but the relationship may be only tangentially related to the signals that the system itself is using.[31] Ideally, we want, whenever possible, to discuss

---

[27] *The EEG is produced by the summed electrical activity of large numbers of neurons and glial cells. As a result one sees only those temporal patterns in the response that are the result of widespread synchronization of discharges across a population or local ensemble, i.e. asynchronous temporal microstructure is averaged out. It is often far from clear what EEG patterns mean in terms of neural codes and representations (Cariani, 1997a).*

[28] *(Walter, 1959b)*

[29] *(Gershenfeld, 1988; Basar, 1990; Rapp, 1994)*

[30] *(Graf and Elbert, 1989). While these methods are quite powerful in theory and concept, in practice the estimation process is frought with many methodological pitfalls and uncertainties (Rapp, 1993b; Rapp, 1993a), some involving the difficulties of distinguishing deterministic processes from stochastic ones.*

[31] *How are the correlation dimensions from EEG's related to the effective dimensionalities of neural coding spaces? It may be that only some of the processes that are modelled in the dynamical system have functional significance - some may be byproducts of particular structural parameters of the system, such as head size or pH, that play only an incidental role in perception and action. It would be as if we had a detailed physical model of DNA: would we recognize the "code" amidst the detailed interactions of its constituent atoms? At some point those aspects of the system that are functionally essential must be identified: one examines how well a human or animal performs a specified task, say discriminating pitches, and one looks for those aspects of neural activity that covary with the performance of the task. One then has a working hypothesis*

the operation of the brain in terms of the neural codes that it uses rather than in terms of the more convoluted external signs of these processes. We want to describe the brain in terms of a signaling system and to discuss how new signals might be created in such a system. We are interested in processes that increase the effective dimensionality of our signal-space.[32]

What determines the dimensionality of a signaling system? As in the physics or chemistry problem, a signaling network must also maintain coherent relations between the different kinds of information being processed in each operation. How the network accomplishes this is a matter of how the information (units, magnitudes) is encoded in the various signals that are sent within the network. Each distinct kind of information, each informational dimension, constitutes a distinct signal type. Within a given signal type, the particular form of the signal conveys the signal's value. Thus every signal has two different aspects, its type and its value.

As an illustrative example, consider a clothes dryer that contains both a temperature sensor that registers two different states (high, low) and a timer that registers two states (not done, done). The outputs of the temperature sensor and the timer are connected to the control module of the dryer by means of separate wires. Here the signal types are defined by their semantics, i.e. one conveys a temperature distinction and one conveys a time distinction, each by virtue of the kind of element that produced it. The form of each signal here is a voltage value, relative to ground, in each respective wire. The controller "knows" which signal represents temperature and which one represents time by the wire on which the signal is conveyed. This is known as a "labelled-line" coding scheme. But knowing the signal type without the signal value, or the signal value without the signal type is useless – one must know both. This means that there must exist two independent aspects of the signal that convey signal type and signal value. In this example, these two aspects are which-wire (type) and which-voltage (value), but obviously the same information could have been encoded and sent to the controller in a different way. For example, instead of separate wires, the temperature sensor and the timer could have been

---

concerning which aspects of neural activity subserve a particular function. In the case of the functional organization of the brain, we believe in some sense that the brain can be described in terms of informational, signaling processes, i.e. in terms of neural representations and their transformations.

[32] *See (Pask, 1959; Pask, 1960; Pask, 1961; Carello et al, 1984; Chen and Conrad, 1994) for other related approaches to this problem.*

given radio transmitters with their own special frequencies with the amplitudes of their respective output waveforms encoding signal values. For each incoming signal the controller would need to somehow distinguish between the two frequencies in order to determine signal type and to measure the amount of power associated with that frequency to determine signal value.

General, very basic aspects of signals can be used to encode signal type and signal value. Four such properties of a time-series signal are: 1) the channel (or medium) through which it is conveyed (e.g. vibrations through air or electrical current through a particular wire), 2) the intrinsic form or pattern that constitutes the signal, independent of magnitude (e.g. a sine wave or a square wave or a particular spectral shape), 3) the time-of-arrival of a signal-event, such as its onset, and 4) properties related to its absolute magnitude, such as its mean or variance.

Any of these aspects of signals can convey signal type and/or signal value. Often signal type is conveyed by a physical channel. One dedicates particular channels to particular kinds of information ("labeled lines") so that the receiver can infer the kind of information arriving in a signal by where it came from (which input channel). An example of such a system is a telegraph network or a telephone switchboard, where each separate signal is conveyed over a separate wire.

Alternately, one can encode signal-type in the intrinsic form of the signal itself. Examples of this are AM and FM radio, where the identity of each signal (the station that sent it) is conveyed by its carrier frequency. Here the average magnitude of the signal is largely irrelevant as long as the signal is detectable; what matters most is its frequency. Internet headers that route messages to their destinations similarly encode their own identity in their form – the path through which they arrive at the receiver is irrelevant to their interpretation. Encoding signal-type in the signal’s form means that many signals can share the same transmission medium (multiplexing), and this in turn means that signals are no longer bound to their own separate physical channels. It makes "broadcast" strategies of information transmission and processing possible.

A third strategy is to use time-of-arrival to indicate signal identity. At one second after a reference time marker, the arriving signal may be a temperature reading, at two seconds, a humidity reading, and so on. This strategy is also used in contemporary telephone systems to concurrently transmit the waveforms of multiple phone calls over the same wire ("time-domain multiplexing").

Lastly, properties related to the magnitude of a signal can be used to encode its type. A waveform with a particular mean voltage could encode one kind of signal, while the same waveform with a different mean could encode another. Mean voltage could encode signal type while variance could encode its value, or vice-versa. By far the most common signal-coding strategy pairs channel-coding of signal-type with magnitude-coding of signal value.

The neural coding problem is the task of understanding which aspects of neural activity encode and convey specific kinds of information.[33] From the perspective of neural coding, the brain is regarded as an adaptive signaling network in which the form and organization of neural signals determine its informational capabilities. Most of the signals in the brain that subserve our perceptual, cognitive, and motor functions appear to be conveyed from neuron to neuron by trains of pulses (action potentials, discharges, spikes). These neural, pulse-coded signals share the same general properties with other kinds of time-series signals (Figure 3): 1) a channel through which they are conveyed, i.e. which neuron produced the spike train, 2) a time pattern or waveform that constitutes the internal form of the signal, i.e. time patterns of spikes in a spike train, 3) the time-of-arrival of signal-events, i.e. when particular spikes arrive in relation to some external event, their latency, and 4) their magnitude, i.e. the number of spikes, the discharge rate and/or the variability of that firing rate.

---

[33] *(Troland, 1929b; Boring, 1933; Boring, 1942; Mountcastle, 1967; Perkell and Bullock, 1968; Uttal, 1973; Cariani, 1995; Rieke et al., 1997)*

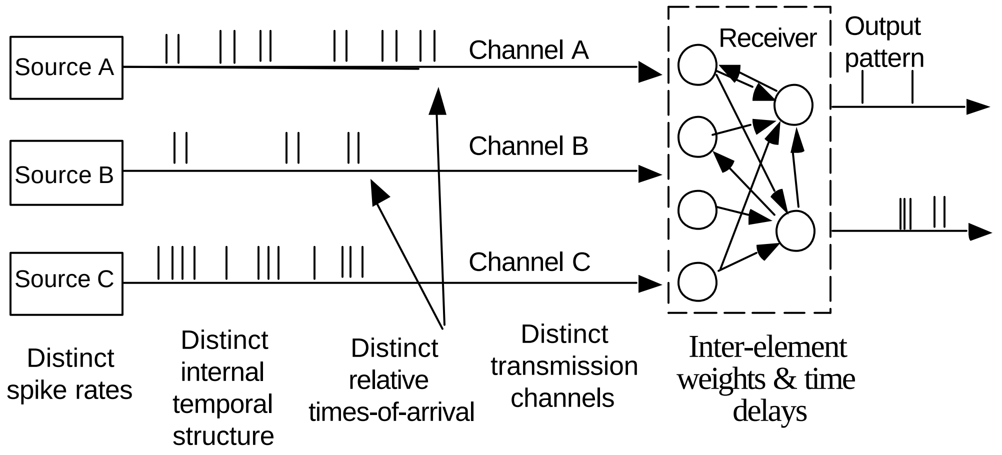

Figure 3. Basic properties of signals in the context of spike trains. A time-series signal has fourcomplementary aspects: the channel through which it arrives (A, B, C), its internal form (interspikeinterval structure), when it arrives (latency), and its magnitude (spike rate). Spike trains A and B arrive through different channels at different times and have different spike rates, but share similar interspike intervals). Compared with trains A and B, train C arrives in a different channel and has a different temporal form, but has the same latency as B and the same spike rate as A. Sources A-C could be different receptors or neurons. Incoming channel, time pattern, arrival time and magnitude can all be used to convey to the receiver what kind of signal is being sent and what its value is. Which aspects of the signal convey signal type and signal state depend upon the interpretation implemented by the receiver (a neural ensemble or assembly). The interpretation depends upon the internal structure of the receiver (properties of elements and their effective connectivities inmagnitude and time). The receiver then produces output signals that themselves can be potentially organized by specific combinations of channels, time patterns, latencies, and magnitudes.

SIGNAL PRIMITIVES IN CONNECTIONIST SYSTEMS

By far, most contemporary models of information-processing in the brain are connectionist models (Figure 4, top).[34] These networks embody basic assumptions about neural coding in their functional organization. Each node in a connectionist network represents a single neuron. Contingent upon the neuron’s inputs, the neuron sends signals to other neurons. Different kinds of information are represented by particular nodes having specific connectivities to the rest of the network and to the outside world. In this interpretation, an auditory neuron conveys information about the acoustic environment by virtue of its specific connections to organs of hearing. It may respond best to only a particular range of frequencies or some particular property or "feature" of sounds, so that the average number of spikes that it produces within a given time window provides an indication of the presence or absence of those factors in the acoustic environment. Thus, in this representational system, particular channels convey signal type (what sound feature is represented) and a scalar quantity (average rate) conveys the signal state (e.g. relative presence or absence of the feature). Each node receives many signals from other nodes, weights them according to their particular signal type, sums them and thresholds the result to produce an output signal. The output signal is then sent, "fanned out", to other nodes in the network. Adaptation in the network takes place by modifying, via a learning rule, the weights associated with each input channel (signal type) according to past experience.

These fundamental assumptions about neural representations mean that connectionist networks depend upon the specificity of connections between the elements to maintain the coherence of function. One can use a "telegraph" or "switchboard" metaphor to crudely describe the normal ways that we visualize the operation of these networks.[35] The particular connections of the wires maintain the identities of the senders (signal types, which convey the semantics of the signals). If one scrambles the wirings, the coherence of the network is lost. The system must then be adaptively rewired through training. In their encoding of signals, connectionist networks are like the clothes dryer with its particular wirings; where they differ from most contemporary devices is in their ability to adaptively rearrange their internal wiring.

---

[34] *See (Churchland and Sejnowski, 1992), (Arbib, 1989), and (Feldman, 1990) for overviews.*

[35] *(John 1972)*

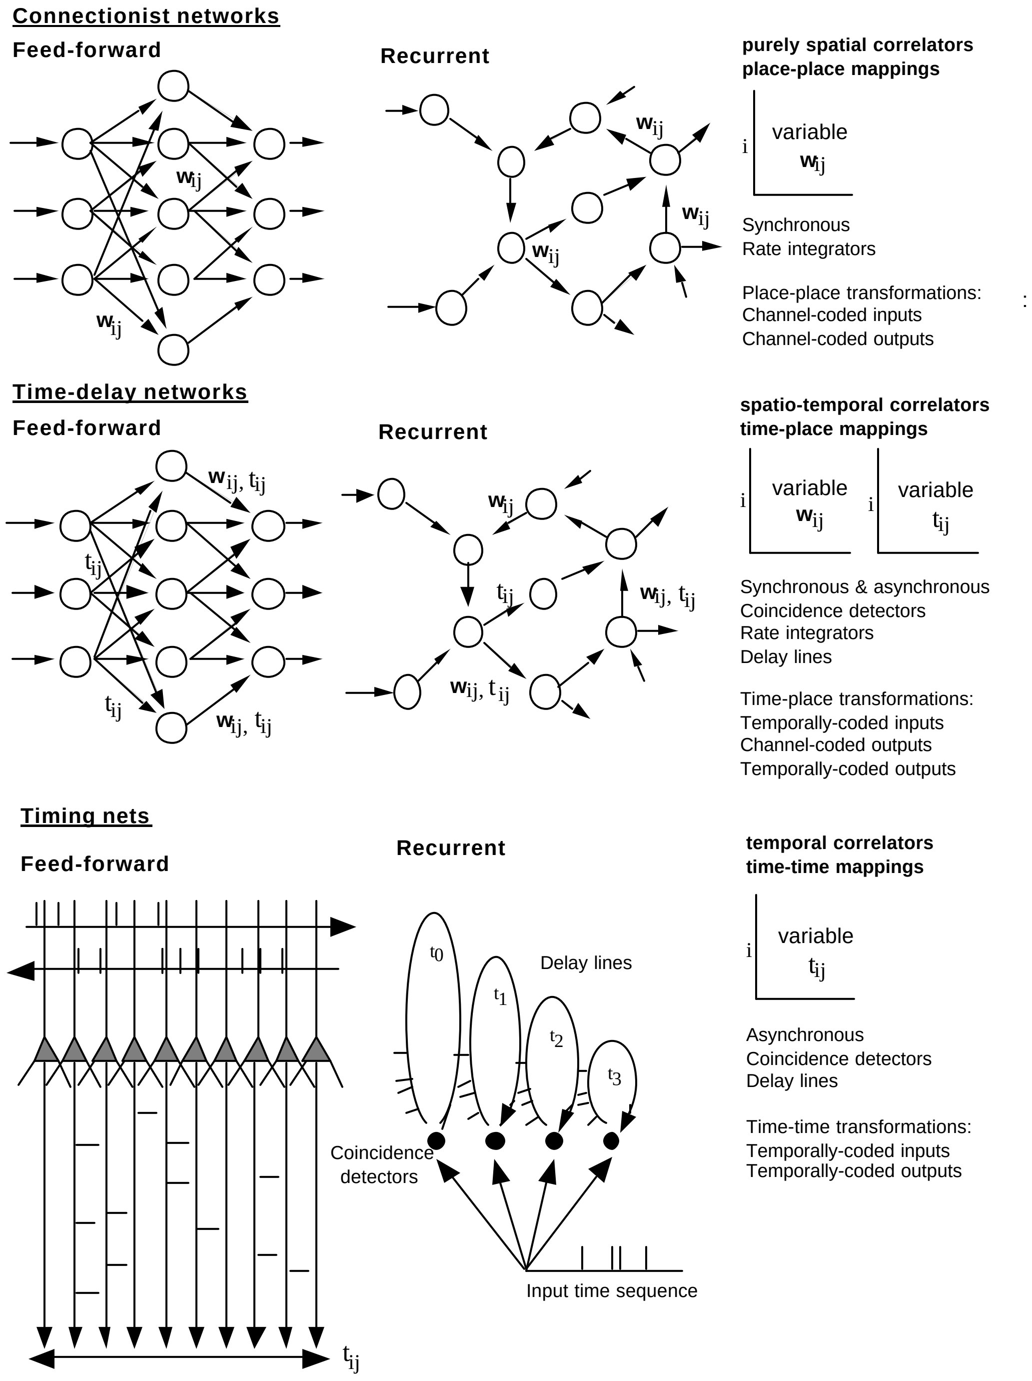

Figure 4. Three kinds of neural networks: connectionist, time-delay, and timing networks.

Connectionist networks thus effectively reflect many of the standard operating assumptions of contemporary neuroscience. The framework is associationist: complex ideas are built up from associations between primitive ideas or sensations. In its most basic form, rate-channel coding is assumed. Different sensations are encoded in different neurons that are connected to different sensory receptors, i.e. signal-type is channel-coded. Different values of these sensations, such as intensity, are encoded in the discharge rates. Systematic changes in stimulus qualities, such as frequency or color are thought to be encoded in across-neuron profiles of neural discharge (which neurons are firing how much). In this system, the channels within which a given neural signal originates determine its signal type. For example, in this view a particular spike train encodes information about a particular frequency by virtue of its connectivity, however distant, to that “place” on the cochlea most sensitive to that frequency.[36] Discharge rates of neurons with particular frequency tunings are then thought to encode the amount of stimulus energy in the corresponding best-frequency regions. Thus, in these rate-channel or rate-place schemes, at each stage of processing all signals are converted into a common neural currency, discharge rate. But it is the specific pattern of interconnections, that makes the network a coherent functional entity. Roughly put, if the brain is a connectionist system, the discharge rates of all its neurons together with their interconnectivities (connection weights) is sufficient to characterize its functional state.

In connectionist networks, mechanisms for adaptive adjustment of connection weights ensure functional stability over time. These are also mechanisms by which new combinations of connections can be realized. Given liberal constraints on how many nodes and connections are available and what interconnections are possible, an astronomical number of possible configurations can be accessed. Emergent-connectivity is thus one way that the generativity of brains can potentially be explained. The mechanisms that subserve this combinatorial richness are generally consistent with a body of neurophysiological and neurological evidence that shows the growth and modification of connections during development, learning, and recovery from injury.[37] In this perspective, conceptual novelty arises purely

---

[36] *For discussions of “place” vs. temporal theories of hearing, see (Troland, 1929a; Boring, 1933; Boring, 1942; Wever, 1949; Evans, 1978). Very similar issues apply for spatial vision and somatoception.*

[37] *There are other ways that novelty could potentially be generated in a connectionist network, some being more physiologically plausible than others. New elements could be added to the system, in effect, creating a growing automaton whose state space expands; this would be a clear example of the emergence of new primitives. Alternately, if we were not restricted to axonal transmission of information via discrete spiking events, then it*

through new combinations of connectivities between pre-existing neural elements, almost excluding out of hand the possibility of creating new kinds of neural signals.

Connectionism, however, is a broad and flexible umbrella that encompasses a very wide range of neural network approaches. Dynamical systems theory is a particularly fertile approach to conceptualizing neural networks, one that lends itself to many alternative interpretations.[38] If the discreteness of the states and the state-determined (rule-governed) nature of the system is emphasized, then such systems resemble discrete, logic-based mechanisms, albeit more distributed and microscopic ones. If, on the other hand, the similarity between alternative neighboring states and trajectories is emphasized, then the dynamical system is seen as a close approximation to a continuous, analog, law-like physical process. In these descriptions one can either see the reductionist world, where macroscopic patterns emerge from deterministic microscopic computations (macro-order from micro-rules), or the formation of discrete attractor basins and rule-like transitions between them from "continuous" dynamics. The difference lies in how the micro-states are interpreted: as surrogates for micro-computations or as approximations to continuous process. Dynamical systems thus can serve either as adjuncts to computationalist world views or as images of material processes.[39] The ways in which such systems are seen to exhibit emergent behaviors, e.g. combinations of microcomputations vs. creation of new discrete attractor-based primitives from continuous dynamics, depends upon the relative success or failure of reductionist explanations: how far the global dynamics can be decomposed into interactions of independent parts.[40]

---

would be possible to conceive of new channels that utilize entirely different media (e.g. neuromodulators, hormonal diffusion, ionic and/or chemically-mediated action through glial cells, volume conduction (Köhler, 1951; Pellionisz, 1991), spike shape and/or amplitude, even changes in cell ultrastructure), to convey messages, and such events would also constitute the emergence of alternative pathways that would create new modes of (analog) signaling, thereby increasing the effective dimensionality of the system. This is not dissimilar to molecular strategies for "extradimensional bypass" that have been proposed (Chen and Conrad, 1994). Widespread use these alternate signaling modes to convey specific information (rather than to modulate some other processing) appears unlikely, but since relatively little attention has been paid to them, it is hard to completely rule out any functional, informational role for them.

[38] *(van de Vijver, 1991; van Gelder and Port, 1995; Horgan and Tienson, 1996)*

[39] *See also Rosen's discussions of computation in terms of dynamical systems theory (Rosen, 1986; Rosen, 1987).*

[40] *See cited works by Rosen, Kampis, and van de Vijver.*

Functional emergence in connectionism thus depends on the observer's perspective.[41] We can first adopt the perspective that the space of possible connectivities is given by the number of pre-existing elements, such that the system is operating within an exceedingly large, but closed space of possibilities. For example, all classifiers are bound by their feature spaces and the set of parameter combinations that implement their decision rule. Seen this way, connectionist systems exhibit emergence by recombination. On the other hand, instead of looking at the parameter space, we can instead observe the functional space of a classifier, i.e. the number of equivalence classes that are distinguished by a classifier. If there is parallel, redundant structure in the underlying parameter space of the classifier, then training can increase the number of classes that are distinguished, thereby increasing the effective dimensionality of the system. A classifer network with many elements might start out in a configuration that distinguishes 20 classes, a number far lower than its capacity. During training, parameter alterations permit the classifier to distinguish 50 classes by reducing the degree of redundancy. Such a system in effect increases the number of global functional states that it has available, so that on a macroscopic scale such a system appears to be adding new states and new observables.[42] In the first interpretation, the exhaustive micro-description of the system in terms of all feature — and parameter-states produces combinatoric emergence. In the second interpretation, from the functionally-oriented perspective of the limited observer, new primitives appear to arise. The two perspectives are complementary; they do not contradict each other outright because they are based on different sets of observables, but being based on different sets of observables, both perspectives cannot be seen at the same time.

SIGNAL PRIMITIVES IN TEMPORALLY-CODED SYSTEMS

The prevailing neural coding assumption in neuroscience is that average rates of neural discharges, not their temporal patterns, convey information.[43] Some of the consequences of these assumptions for connectionist networks have been outlined above. However, if precise timings of spikes are used in the nervous system to convey information, then there are a variety of alternative ways that neural networks can be built. Like our clothes dryer, it is possible that point-to-point wiring can

---

[41] *(van de Vijver, 1991; Kolen and Pollack, 1995; van Gelder and Port, 1995).*

[42] *e.g. as in a self-organizing Kohonen map (Kohonen, 1988).*

[43] *e.g. Barlow's Neuron Doctrine (Barlow, 1972; Feldman, 1990; Barlow, 1995).*

be replaced by a more flexible coding scheme. There do exist many neurophysiological and psychophysical examples that point to time-codes in the brain.[44] Many of these examples are particularly striking.[45] At present the body of evidence, while considerable and very suggestive, is not yet extensive enough to be anything more than a tentative general hypothesis for temporal coding of sensory information. Consequently, the broader functional implications of temporal codes remain largely unexplored. We will discuss some of these implications for neural processing architectures – time-delay networks and timing nets – and what signal-creation would mean in such information processing systems.

The simplest temporal pattern code is an interspike interval code, where times between spikes convey sensory information or motor commands.[46] To the extent that voltages produced by sensory receptors follow stimulus time patterns, then spike trains of subsequent sensory neurons in the pathway will reflect the time pattern of the stimulus waveform. The resulting neural temporal discharge patterns then form an analog, iconic image of the stimulus, albeit in the medium of neural discharges rather than in the physical medium of the stimulus.[47] Rather than converting all information into channel-coded features, the brain might instead make use of these time patterns to represent and analyze the stimulus. In the early evolution of theories of sensory coding, Rutherford's temporal "telephone theory" of hearing embodied this strategy, in opposition to Helmholtz’s channel-based "place principle."[48] In a temporal theory, the stimulus impresses its own iconic, spatiotemporal form on the neural representation, while in a channel-coded “place” theory, neural representations take their form from sets of spatially-distributed channels that are provided by sensory organs,

---

[44] *(Bullock, 1967; Mountcastle, 1967; Perkell and Bullock, 1968; Uttal, 1973; Wasserman, 1992; Cariani, 1995; Cariani, 1997c).*

[45] *(Reichardt, 1961; Young, 1977; Covey, 1980; Emmers, 1981; Bialek et al., 1991; Langner, 1992; Carr, 1993; Di Lorenzo and Hecht, 1993; Mountcastle, 1993; Heiligenberg, 1994; Cariani and Delgutte, 1996; Lestienne, 1996; Simmons, 1996)*

[46] *The interspike interval is the joint property of two arrival times, and as such it is inherently a relational entity rather than an absolute event-in-itself. Correlation-based representations are constructed from these relational atoms rather from registrations of absolute, relation-less pixel-like perceptual atoms.*

[47] *It should be noted that while the spikes that make up an interspike interval are discrete events, the interval itself can vary in a continuous manner. Thus interspike interval codes are analog codes. This makes them suitable for conveying a continuous range of possible values.*

[48] *(Boring, 1942).*

e.g. spatial positions within retinal, somatosensory or cochlear surfaces.

Temporal pattern and sensory channel can potentially interact in various ways to codetermine sensory quality. For example, if most incoming spike trains have many 10 millisecond (ms) intervals between spikes, neural central processors may infer a source in the environment that is producing a 100 Hz periodicity. This might be interpreted as a 100 Hz mechanical vibration of the skin if the pattern arrives through somatosensory channels or a 100 Hz pitch if it arrives through auditory ones. Whether the stimulus is a sound or a mechanical vibration, 100 Hz can be distinguished from 120 Hz by differences in the interspike intervals that are produced, i.e. by comparing the relative numbers of 10 ms vs. 8.25 ms intervals.[49] Here stimulus frequency is encoded by temporal-pattern and stimulus modality is encoded by channel. Alternatively, if each sensory pathway adds its own characteristic delays and/or additional time patterns[50], then a central processor could differentiate different types of information solely on the basis of incoming time patterns, so that both frequency and modality could be temporally-coded in a manner that depended on incoming transmission channels. One of the differences between channel-coding of signal-types and temporal-coding of signal-types is that in the former, the set of types is given by the channels themselves (which presumably are already created) while in the latter the stimulus impresses its time pattern on the neural responses and effectively creates signal-types in its image. If the underlying temporal dynamics are rich enough, then new time patterns and hence signal-types can be created through the interaction of time patterns.

In the context of the neural coding problem, pitch is a useful phenomenon to examine because a great deal is known about both its psychophysics and at least some of the neural representations that may

---

[49] *Periodic stimuli impress their time structure upon neurons at many stations in the ascending sensory pathways of the auditory and somatosensory systems. The result is that many interspike intervals that correspond to stimulus periodicities are found in both systems (Wever, 1949; Kiang et al., 1965; Rose et al., 1967; Mountcastle et al., 1969; Evans, 1978; Langner, 1983; Langner, 1985; Morley et al., 1990; Langner, 1992; Mountcastle, 1993; Cariani and Delgutte, 1996; Cariani, 1997b; Cariani, 1997c).*

[50] *(Emmers, 1981) found modality-specific latency and interval patterns for the coding of somatosensory qualities in single units at the level of the thalamus. Stimulus qualities such as these are almost certainly encoded in the discharge patterns of large ensembles of sensory neurons. The population-discharge pattern evoked by a 100 Hz mechanical vibration on the skin will share some characteristics with that evoked by a 100 Hz tone (many 10 msec intervals), but differ in other respects (how many neurons are activated).*

subserve it. We have investigated the neurophysiology of the pitches produced by complex tones (i.e. musical pitch) at the earliest stages of auditory processing, at the level of the auditory nerve and the cochlear nucleus. At these auditory stations there appear to be precise, robust, and pervasive correspondences between interspike interval distributions and the complex patterns of pitches that are heard by human listeners for particular stimuli.[51] "Population interval distributions" are the interspike interval distributions that exist over a whole population of auditory neurons (e.g. over the 30,000 nerve fibers that make up the auditory nerve). Acoustic stimuli are presented to an anesthetized animal and the trains of spikes that are evoked by the stimulus in many single auditory neurons are recorded one by one. From the responses of 50-100 neurons to 100 presentations of the stimulus, very good estimates can be made of the population interval distribution. With very few exceptions, at the level of the auditory nerve, the pitch that is heard corresponds to the interspike intervals that are most prevalent across the population. The population interval distributions for two different stimuli that produce the same 160 Hz pitch are shown in Figure 5, where the major interval peaks lie at the pitch period (6.25 ms) and its multiples. On a deeper level, the shapes of population interval distributions bear close resemblances to the autocorrelation functions of the stimuli that produce them.[52] Because the autocorrelation function contains the same information about the stimulus as its power spectrum, population interval statistics potentially can encode auditory qualities related to spectral shape (e.g. timbre, vowel quality, other phonetic contrasts, musical intervals) as well as pitch. Unlike power spectra, however, autocorrelation preserves underlying temporal relationships between stimulus components.[53]

---

[51] *(Evans, 1978; Langner, 1985; Meddis and Hewitt, 1991; Cariani and Delgutte, 1996; Cariani, 1997b)*

[52] *(Cariani, 1997b)*

[53] *For example, the pitches of complex sounds are more readily seen in autocorrelation functions than power spectra.*

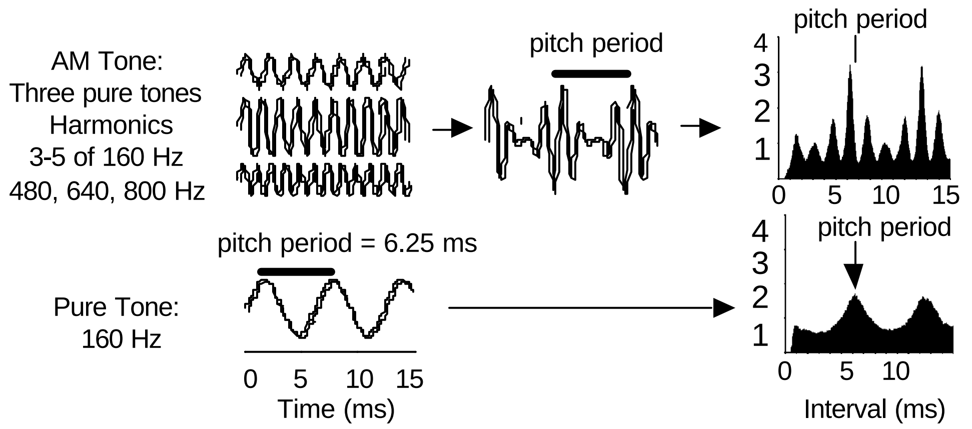

Figure 5. Emergence of a new pitch percept. Top: Three pure tones, 480, 640 and 800 Hz (left) sounded individually produce pitches at their respective frequencies. When they are are sounded together (middle), they produce a new, lower pitch at their common fundamental frequency, 160Hz. This pitch is the same as would be produced by a 160 Hz pure tone (lower waveform). At rightare interspike interval distributions compiled from the recorded neural responses of 50-100 catauditory nerve fibers to these stimuli. The major peaks in both all-order interval distributions correspond to the common pitch, at 160 Hz, that would be heard. See (Cariani & Delgutte, 1996) for detailed description of methods.

The pitch of complex tones is a prime example of an emergent percept.[54] For example, when single pure tones are played we hear pitches associated with each pure tone. If many harmonics of a 100 Hz pure tone (e.g. sinusoids with frequencies of 500, 600, 700, 800, 900, and 1000 Hz) are sounded together, we hear a “low” pitch at their common fundamental frequency (F0 = 100 Hz). The fundamental frequency is the largest common denominator of the frequencies. We clearly hear this low pitch at the fundamental even if none of the pure tones we played were themselves at the fundamental (i.e. we hear the "missing fundamental"). Further, all of the individual tones may be fused together to such a degree that it may even be difficult for us to hear them as separate entities. Whether one thinks of this perceptual phenomenon in terms of a top-down pattern-completion process, or in terms of bottom-up interactions between underlying neural

---

[54] *(Morgan, 1931) cites a related example of musical chords (p.4). He says "Such emergence of the new is now widely accepted where life and mind are concerned. It is untiringly advocated by Professor Bergson. Wundt pressed its acceptance under his 'principle of creative resultants' (i.e. what we distinguish as emergents) which, he says, 'attempts to state the fact that in all psychical combinations the product is not a mere sum of the separate elements...but that it represents a new creation."*

representations of the pure tones, the outcome is an emergent Gestalt that is the result of the whole stimulus.

What neurophysiological processes might underlie the creation of this new, low pitch? When each pure tone of frequency f is played alone, it produces many interspike intervals whose duration matches its period (1/f) and its integer multiples (n/f). When all of the intervals associated with all of the individual tones are combined, as they are when the tones are sounded together, the result is that the longer intervals common to all of the individual tones are most numerous. These will be intervals at the fundamental period (1/F0), the period that is an integer multiple of each of the periods of the individual pure tones (the harmonics). The resultant global interval pattern, with its major peak at the fundamental, then most resembles the interval pattern that is produced by a pure tone at the fundamental frequency. Thus, if the nervous system used such a population-interval representation for pitch, then one would expect the pitch of the pure tone at F0 and the low pitch of the harmonic complex to be similar. Such representations would account for the appearance of such low pitches when multiple harmonically-related pure tones were sounded together.

We can translate this into signal-types and signal-values. On the level of signal-types corresponding to each pure tone component, it appears (perceptually and neurally) that a new signal-type has been created — there is a global time pattern that is related to, but different from its constituents. This emergence can occur because the representation of each pure tone component has a microstructure (subharmonics) that is not included when we describe the situation in terms of signal A + signal B + signal C. A property that is not included in the original description thus comes into play to create an new emergent property. On the other hand, we would not see this as emergent if we took the whole underlying sinusoidal pattern of each pure tone as primitive — the summation of these patterns would be seen as predictable combinations of these primitives.

In effect, because the discharges of auditory nerve fibers follow the temporal pattern of the stimulus waveform, the neural interspike interval representation of the stimulus mirrors its time structure. If there were a central auditory processor that analyzed which intervals are most prominant in the discharges of populations of auditory neurons, then its actions could account for the pitches that we hear. The existence of such timing information and of a central processor capable of handling it would provide a straightforward explanation for many complex pitch phenomena. Since interspike intervals themselves preserve temporal relationships and retain the inherent harmonic structure that is present in the stimulus, they might also explain the strong role that harmonic relationships play in hearing. If this is how the auditory system represents sounds, then the perception of harmonic relationships is an innate perceptual facility that is given by the basic neural codes that the system uses, rather than an acquired perceptual capability that is built up solely through experienced associations.

A comprehensive account of pitch and the creation of new pitches that is based on temporal pattern codes thus exists at the level of the auditory nerve. It appears likely that such population-based interspike interval representations parallel pitch percepts at least up to the level of the midbrain. Many puzzles remain, however. As one ascends further up the auditory pathway, each neuron discharges fewer and fewer spikes for each stimulus presentation, so that fine timing information (periodicities above 200 Hz) becomes more difficult to detect. While a good deal of timing information is still present at the level of the auditory thalamus,[55] exactly how much is available for cortical processing remains a critical question.

TEMPORAL PROCESSING NETWORKS

Purely connectionist networks operate on channel-coded inputs. Although the properties of connectionist networks have been widely explored, comparatively little work has been done on temporal processing networks, which operate on temporally-coded inputs. Temporal processing networks can be divided into time-delay networks and timing networks (Figure 4). Time-delay networks convert incoming time patterns to spatialized, channel-coded output patterns, while timing nets convert incoming time patterns to temporal output patterns.

Most temporal processing schemes in neural nets have concerned time-delay neural networks, where not only neural connection weights but also time delays between elements are represented. The neural analogs of connection weights are synaptic efficacies, while those of time delays are dendritic and axonal conduction times and various membrane and cytoplasmic recovery times for neurons to rebound from past events.

Some of the earliest proposals for how neural networks could account for human perceptual abilities, such as the Jeffress model for

---

[55] *(de Ribaupierre, 1997)*

sound localization and the Licklider model for pitch perception[56], were auditory time-delay networks. The Jeffress model was a simple time-delay network that used time-structured inputs, tapped delay lines, and sets of coincidence detectors to carry out binaural cross-correlations to localize sources on the basis of interaural time-of-arrival differences. Over the last half-century, the corresponding neuroanatomical structures and neural response properties in the auditory brainstem that subserve the cross-correlation analysis have been identified and characterized.[57] For the analysis of auditory forms, Licklider's duplex model proposed time-delay networks that carried out temporal autocorrelation operations in order to compute the pitch of pure and complex tones.[58] In a time-delay network, one can either change connection weights to alter patterns of time delay or change time delays (conduction velocities) to modify effective connectivity. Both mechanisms have been suggested as means for making such networks adaptive.[59] More recently, following the resurgence of interest in neural networks in the last decade, more abstract, general purpose time-delay network architectures have been proposed.[60]

Time delay networks are usually utilized to effect time-to-place transformations, wherein time-structured inputs are converted to spatial patterns of activation. The outputs of such networks are then analyzed by conventional connectionist networks. Corresponding physiological models for pitch processing posit time-to-place transformations at the level of the auditory midbrain, in populations of neurons tuned to particular periodicities.[61] It remains to be seen whether these conversion strategies really work, whether they can account for the robustness and precision of the pitch percept.[62] Once the time-to-place

---

[56] *(Jeffress, 1948; Licklider, 1951; Licklider, 1956; Licklider, 1959)*

[57] *(Casseday and Covey, 1995)*

[58] *(Licklider, 1951; Licklider, 1956). Subsequent temporal processing models for the binaural fusion of sounds ("auditory scene analysis") combined auto-correlation and cross-correlation mechanisms (Licklider, 1959; Cherry, 1961). Mechanisms along these lines that incorporate inhibitory inputs seen in brainstem physiology have been proposed for processing of sound duration cues in bats (Casseday and Covey, 1995).*

[59] *Licklider's 1959 triplex model proposed a self-organizing neural network in which particular time and frequency patterns could be recognized by adaptively altering interconnectivities. (MacKay, 1962) proposed timing nets in which conduction velocities could be adaptively controlled.*

[60] *(Tank and Hopfield, 1987; Chappelier and Grumbach, 1994).*

[61] *(Langner, 1992).*

[62] *Major problems arise when one looks closely at how the putative periodicity-detectors actually behave. Their tuning is coarse, while pitch discrimination is much more precise;*

conversions are effected, then standard, non-temporal neural networks perform an analysis of which periodicity channels were activated. By converting all time-structure into the common currency of spatial patterns of activation, such a theory preserves the fundamental assumption that the cerebral cortex is purely a spatial pattern analyzer.

On the other hand, there is no inherent reason why cortical structures with their many layered surfaces and slow horizontal-fiber systems cannot be viewed as time-delay neural networks that are capable of both analyzing and generating arbitrary spatiotemporal patterns.[63] This conception of the cortex is consistent with a functional role for cortical pyramidal cells as coincidence detectors rather than rate integrators.[64] When neurons have short integration windows (e.g. a few milliseconds or less), they operate as coincidence detectors in which relative timings of incoming spikes become more important in determining the neuron's discharge pattern than average rates of incoming spikes or synaptic efficacies. As integration windows span longer durations, coordinated timing of inputs becomes much less critical, while sheer numbers of incoming spikes and their respective synaptic efficacies become more important.

Ranges of temporal integration thus form a continuum of possible neural networks: from purely connectionist networks to time-delay networks to timing networks (Figure 4). For purely connectionist networks, spatial patterns of inputs are critical and fine timing is irrelevant. For time-delay networks, both spatial pattern and timing codetermine behavior. For timing nets, coherence of timing is relatively

---

their selectivity degrades at higher levels, whereas pitch frequency and pitch strength is remarkably stable at such levels. Further, the range of best periodicity-tunings becomes lower and lower in frequency as one proceeds to higher auditory stations. By the time the auditory cortex is reached, most units have best-periodicities below 20 Hz. There is something wrong with this basic picture. There is currently no remotely adequate physiologically-grounded account, temporal or place-based, of how the pitches of complex tones are represented at the level of the auditory cortex.

[63] *(Braitenberg, 1961; Braitenberg, 1967; Abeles, 1982a; Abeles, 1990; Braitenberg, Heck and Sultan, in press). Often it is argued that the conductions are not slow enough to achieve longer delays on the order of tens or hundreds of milliseconds. But these cortical structures are all connected to each other and to other subcortical structures by means of recurrent pathways that create multisynaptic loops limited only by time-jitter. The number of multisynaptic return paths of a given duration increases combinatorically with the number of nodes. It is an open theoretical question whether reverberating time patterns can be maintained in randomly-connected timing networks given some rules governing fan-in, fan-out, temporal integration windows, and thresholds or whether connection-rules can be specified so as to arrive at timing networks that propagate temporal structure with minimal degradation rates.*

[64] *(Abeles, 1982b; Softky and Koch, 1993)*

more important than spatial pattern. Where we place specific neural structures along this continuum depends both on how much relevant information is temporally encoded in their inputs and to what extent, given the temporal integration windows of their elements, these neural architectures can make effective use of the timing information that is present.

Thus far discussions of timing issues within the neurosciences have mainly been confined to problems of perceptual organization: how the brain coherently organizes the many elements of a perceptual scene into multiple discrete objects. Almost a century ago, the Gestaltists brought these problems of perceptual organization to the forefront of perceptual psychology. A well known metaphor, the “cocktail party effect,” begs the question of how we manage to follow a single conversation amidst the background din of many other conversations.[65] Nowadays these concerns fall under the rubric of “scene analysis.”[66]

Perceptual scenes are thought to be segmented, through bottom-up similarities between local perceptual features (e.g. common motion, texture, color, rhythm, voice pitch) and/or via top-down perceptual and cognitive expectations. Channel-coded local features within segments are thought to be bound together to form distinguishable stable objects. A fundamental difficulty for these representational strategies is that information from a given (cochleotopic or retinotopic) channel may participate in multiple objects. How does a central processor determine which channels go with which object when? In their critiques of the local features of associationist psychology, the Gestaltists provided many demonstrations of perceptual effects that depended on global relations rather than local features.[67]

In the last decade many of these questions have been revived within the context of feature-based theories of visual processing.[68] In accord with current coding assumptions in vision, local features are place-coded, with common timing providing a mechanism by which related local features can be grouped together. In this view, the representations of multiple objects are spread out over the same sets of neurons, so that the information linked with these objects is effectively multiplexed on the same processing and transmission lines.

---

[65] *(Cherry, 1961; von der Malsburg and Schneider, 1986; von der Malsberg, 1995).*

[66] *e.g.(Bregman, 1990)*

[67] *(Wertheimer, 1923; Uttal, 1975; Uttal, 1988). In audition, pitch, melody, and rhythm were Gestaltist examples for the primacy of global relations over local features.*

[68] *(Eckhorn and Reitboeck, 1990; Singer, 1990; Eckhorn, 1991; Singer, 1995)*

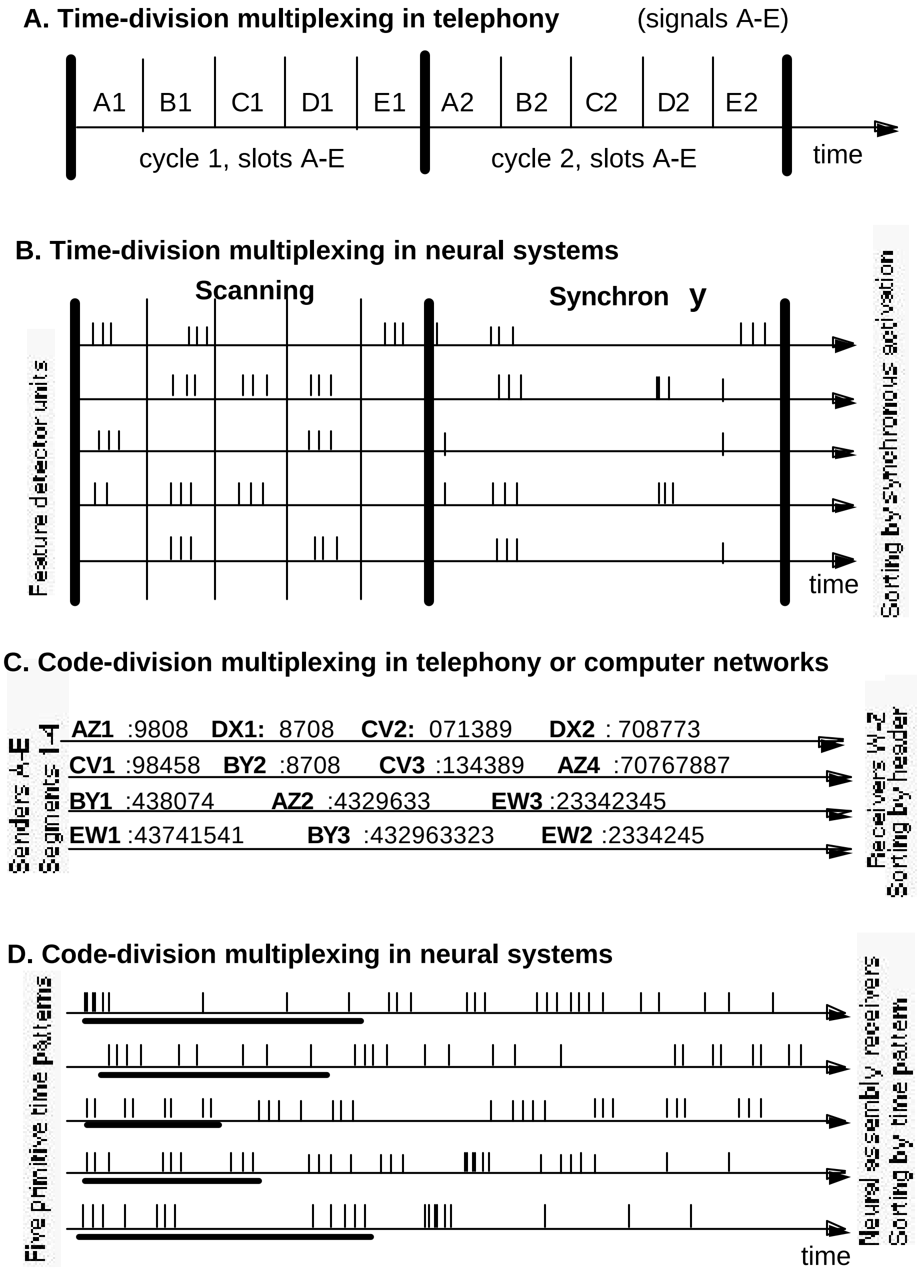

Figure 6. Signal multiplexing: concurrent transmission of multiple signals over the same channels. A.Time-division multiplexing allocates particular time-slots to each signal (A-E). B. Time-division multiplexing models in the brain assume that different objects are represented by channel-activation patterns that are grouped according to common phase within an oscillation (scanning) or by temporal proximity (synchrony). C. Code-division multiplexing encodes signal-type in the form of the signal,using special header sequences that direct asynchronous and nonsequential transmission and reconstruction of the signal by each receiver. D. Code-division multiplexing can be implemented using asynchronous, temporal pattern codes if

different kinds of information (e.g. encoding different sensory qualities) have characteristically different forms, and if neural assemblies respond to particular complex temporal patterns (underbars). Broadcast is possible, because patterns are no longer channel-based.

This is a strategy, called "time-division multiplexing," is used in some telecommunications systems (Figure 6). In the neuroscientific discussion, there are two basic types of time-division multiplexing models: those based on interneural synchronies (perceptual properties associated with those elements that fire together bind together) and those based in intrinsic neural periodicities such as the "alpha rhythm" (properties associated with elements firing within a particular phase of a scanning period bind together).[69] Considerable neurophysiological evidence exists for both interneural synchronies and neural oscillations, although there are many unanswered questions and a great deal of debate over the functional roles that these neural processes play.[70] One potential difficulty with time-division multiplexing is that many different kinds of stimuli, such as trains of light flashes, acoustic clicks, or electrical shocks can drive large parts of the brain in arbitrary rhythmic patterns, yet in general our perceptions are relatively undisturbed by such interventions. Similarly, the timing of particular parts of the stimulus can be manipulated, presumably, to desynchronize parts of the response, generally without disrupting perception.[71]

All of these considerations argue in favor of asynchronous neural representational mechanisms that cannot be disrupted by small shifts in time-of-arrival of signals. This has always been taken as one of the major strengths of average rate codes, that a great deal of temporal imprecision can be accommodated without disrupting the coherence of neural representations. Unfortunately, it has also been commonly assumed that all temporal codes must be synchronous ones, either rigidly locked to the time-course of the stimulus or to the discharges of

---

[69] *Scanning models were first proposed by [Pitts, 1965 #1023] (McCulloch, 1951) and (Walter, 1953; Walter, 1959b; Walter, 1959a). More recent versions have been proposed by (Llinas and Pare, 1996).*

[70] *See discussions in (John, 1967b; John et al., 1973; Thatcher and John, 1977; John and Schwartz, 1978; Cariani, 1997a)*

[71] *Many of these general counterarguments to scanning mechanisms were raised at the Hixon Symposium in response to (McCulloch, 1951). There is currently a great deal of healthy skepticism about all of the mechanisms for segmentation and binding that have been proposed. On the other hand, there are a number of striking and odd perceptual and cognitive effects that were reported by (Walter, 1959b) when specific time patterns of stimulation were presented and/or triggered by endogenous brain rhythms. A cottage industry that produces brain stimulation devices exploits these effects (Hutchinson, 1991).*

other neurons in a population. Conflation of temporal coding with synchrony sets up a false dichotomy between asynchronous rate- and channel-based codes on one hand, and synchronous time codes on the other. But a third possibility exists, that of asynchronous time codes. In such codes, multidimensional signals can be represented and conveyed in the statistics of time patterns that are propagated asynchronously through neural populations.[72] The population-interval representations for pitch that were described above would be examples of such asynchronous codes. A particular message is encoded via an interspike interval distribution or through the distributions of more complex time patterns. Multiple time patterns can be sent on the same lines, interspersed or interleaved with other patterns. If the time patterns are sparse enough in time, representations based on them are effectively "transparent" with respect to each other.[73] Here precise synchrony across transmission lines is not necessarily required; receivers are assumed to have rich sets of relative delays that allow them to integrate similar time patterns arriving at slightly different times.

There is some neurophysiological evidence that supports the possibility of such asynchronous temporal pattern codes, where precise spike time patterns have been found mixed in with other spike

---

[72] *Other examples of interval-based multiplexing and possible processing mechanismsfor decoding multiplexed time patterns are (Chung, Raymond and Lettvin, 1970; Raymond and Lettvin, 1978; Waxman, 1978; Pratt, 1990; Wasserman, 1992). Optican, Gawne,Richmond and coworkers have proposed multidimensional temporal codes for the representation of visual forms (Richmond et al., 1987a; Richmond et al., 1987b; Richmond, Optican and Gawne, 1989; Richmond and Gawne, 1998). The group has a deep appreciation of the importance of code-dimensionality and present many strong information-theoretic and functional arguments for the necessity of multiplexing visual information. However, their temporal codes are variations in discharge rates over tens or hundreds of milliseconds rather than the more precise stimulus-driven spike arrival times that are suggested by the visual system of the fly (Reichardt, 1961; Reichardt and Poggio, 1976; Bialek et al, 1991) or cortical encoding of acoustic transients (Phillips, 1989), where sub-millisecond precisions are found.*

[73] *Because the time patterns associated with multiple objects can be interleaved, they do not automatically interfere with each other in the way that representations based on local spatially — or temporally — contiguous receptive fields do. This means that grouping can occur by common asynchronous time structure rather than through synchrony, freeing the system of the temporal coordination all of the information related to a particular object. If the basic sensory representation itself has this "transparent" nature, then the segmentation and binding problem is no longer so formidible. Our impressive ability to separate auditory and visual objects would then be part and parcel of the kinds of temporal correlation-based representations on which they were based (Cariani, 1995; Cariani and Delgutte, 1996).*

patterns.[74] Since these patterns are also not rigidly synchronized to the stimulus (e.g. the patterns of Figure 8), virtually all traditional neurophysiological methods for analyzing spike trains and evoked potentials miss them. As a consequence, the extent of such temporal microstructure in the brain remains largely unexplored.

If temporal microstructure is more ubiquitous than commonly thought, what implications would this hold for the brain as a signaling system? We can outline how a putative asynchronous signaling system might be organized using complex temporal pattern primitives.[75] Rather than converting all time structure into spatial activation patterns, time structure would be the organizing currency of the system. Time structure would be preserved in sparse form, distributed over neurons in local ensembles. Information processing would be statistical-mechanical in nature, implementing temporal correlation operations on the all-order interval statistics of larger ensembles of neurons.[76]

In sensory maps time structure would exist embedded within the classical orderings of cortical neural maps, e.g. by retinal position (retinotopy), by body position (somatotopy), by cochlear position or frequency (tonotopy). Within local patches and across whole cortical fields the all-order interval statistics of neural populations would encode information concerning sensory qualities such as pitch, timbre, rhythm, spatial frequencies, and tactile texture. These would statistics provide running, parallel, local, multidimensional, autocorrelation-like representations. The outputs of local patches would be cross-correlated with each other, such that all-to-all cross-correlations of autocorrelations are computed.[77] The canonical neural computations are carried out in feed-forward and recurrent temporal coincidence arrays

---

[74] *(Emmers, 1981; Strehler and Lestienne, 1986; Lestienne and Strehler, 1987; Lestienne et al., 1990; Abeles et al., 1993; Mountcastle, 1993; Vaadia et al., 1995; Lestienne, 1996)*

[75] *In terms of temporal mechanisms and timing nets, we are still working in an unformed realm that is not unlike the world of neural networks before Rashevsky, McCulloch, and Pitts. We are currently outlining what basic operations are possible, and what logics might be built up from their action.*

[76] *(John, 1972; John and Schwartz, 1978; John, 1988; Cariani, 1995; Cariani, 1997a) REM*

[77] *Analogous running spatial auto — and cross — correlation operations have been proposed for vision, e.g. (Kabrisky, 1966; Uttal, 1975; Uttal, 1988). Related, temporal correlation models for visual motion detection and texture are (Reichardt, 1961; Reitboeck, Pabst and Eckhorn, 1988; Pabst, Reitboeck and Eckhorn, 1989; Reitboeck, 1989).*

(Figure 4, bottom).[78] These arrays in effect function as temporal sieves, propagating in their outputs those interspike intervals that are common to many inputs. In doing so, they extract those time patterns that are common to their inputs. If patches of cortical maps are interconnected by horizontal delay lines, then coincidence arrays extract and potentially reinforce commonalities of interval statistics present in corresponding frequency lamina or retinotopic patches. For example, different frequency regions driven by harmonically-related component frequencies would share common intervals at the fundamental, such that they would mutually reinforce each other. Iterated, reciprocal correlation operations then provide a basis for grouping by commonality of interval statistics rather than by strict synchrony. Perception thus would rely on the statistical mechanics of temporal coincidences between spikes rather than on specific ensembles of firing rates.[79] The present temporal correlation account is related to the "synfire chains" proposed by Abeles[80], except that functionally, the outputs of coincidence operations are temporal patterns (collective interval statistics) rather than ensembles of particular neurons firing in synchrony.

Temporal coding permits different modes of signal multiplexing. Particular thalamic and cortical regions might also generate characteristic temporal signatures[81] that signal the type of information being conveyed in the spike train, much in the same way that internet headers carry information about the nature of the message and its sender. If signal types can be encoded in characteristic time patterns

---

[78] *In the last year, we have begun to investigate the basic properties of these coincidence arrays (Figure 4, bottom) (Cariani, 1998a; Cariani, 1998b). (Longuet-Higgins, 1987; Longuet-Higgins, 1989) has shown how simple coincidence arrays can carry out convolutions in the time domain. These simple coincidence networks can be used to multiply autocorrelation functions, such that multidimensional comparisons can be made, and common periodicities extracted.*

[79] *Spatially-distributed temporal structures have been proposed in the past as functional organizations for the brain athat do not require specific point-to-point wiring, i.e. "switchboards" (Lashley, 1951; John, 1967a; John, 1967b; John, 1972; Thatcher and John, 1977; John and Schwartz, 1978; John, 1988). These alternatives have been ignored, in some part because of the current technological pre-eminence of digital and symbolic information processing strategies in computers, such as spatially-localized memory, sequential-hierarchical processing, and discrete local feature detectors.*

[80] *(Abeles, 1990)*

[81] *(Emmers, 1981) found thalamic spike patterns that consisted of onset bursts followed by later interspike intervals. He found that the latencies of the later intervals as well as the duration of the intervals themselves encoded the kind of somatosensory information that was being transmitted.*

that serve as temporal "stamps" or "tags", then different kinds of time patterns can be sent over the same transmission lines, at different times or even interleaved together. This scheme is somewhat like "code-division" multiplexing, a mode of signal multiplexing that is used in some cellular telephone and computer networks (Figure 6, bottom). In such networks receivers that can potentially receive all transmissions, but only respond to those messages that are relevant. In the context of neural assemblies, this might plausibly involve the selection of particular relative delay configurations, such that the assembly reacts differentially to particular time patterns.[82] One then has potential response-specificity based on the presence of specific temporal patterns in the inputs, in the manner of a lock-and-key or a matched filter. Neurophysiological support for such adaptive timing mechanisms comes from conditioning experiments where associations are formed for time structured stimuli, e.g. different rhythms. It has been found in single neurons and population responses that there is an "assimilation of the rhythm" of the stimulus, such that temporal response properties change over the course of conditioning.[83] Behaviorally, there is evidence from conditioning studies suggesting that the time courses of events, both rewarded and unrewarded, are stored in memory,[84] and that cross-modal stimulus generalizations are readily made for stimuli having common time structure (e.g. 3 vs 10 Hz clicks, flashes, or shocks).[85] These observations suggest very general facilities for storing and comparing time patterns.

There are many potential advantages of coordinating large numbers of interacting neural processes through such asynchronous, complex time patterns. Because different kinds of information can be kept separate by virtue of time pattern alone, the function of preserving signal identity can be dissociated from particular labeled transmission lines, such that highly specific, point-to-point wiring is no longer strictly required. This in turn permits "broadcast" modes of coordination of neural assemblies,[86] whereby time patterns are propagated widely

---

[82] *Or different neural assemblies can have distinct time courses of activation and de-activation that can support processing of temporal and sequential patterns, as in D. G. MacKay's theory of cognitive timing nodes (MacKay, 1987).*

[83] *(John, 1967a; John, 1967b; Morrell, 1967; Thatcher and John, 1977; John and Schwartz, 1978)*

[84] *(Miller and Barnet, 1993)*

[85] *(John, 1967b)*

[86] *In his paper on adaptive timing networks (MacKay, 1962), D. M. MacKay (1961) pointed out that temporal patterns could be broadcast, a coordinative strategy he called*

through the system. Memory traces are then themselves reverberating time patterns, albeit not necessarily synchronized at the level of whole neural populations. Circulating memory traces would then be temporally cross-correlated with incoming time patterns in primary and secondary sensory areas to facilitate the detection of particularly relevant information.[87] The basic organization is a build-up, reverberatory process mediated through recurrent loops of reciprocal connections between cortical areas. In turn different sets of neural assemblies are activated according to their complex tunings that were acquired through previous experience. Thus the tunings represent the "relevance criteria" through which a given assembly can be weakly or strongly activated. When activated, such assemblies would be capable of producing characteristic stimulus-triggered time patterns. Unlike the pitch example above, where the emergent neural time pattern was a stimulus-driven time pattern, these internally-generated patterns need not be direct reflections of the time structure of the stimulus. Such patterns are seen in late evoked potentials in which characteristic temporal responses are released by particular stimuli.[88] These kinds of neural tags or signatures would then be added to the time structure of the spike patterns that caused their production, either through synchronous concatenation or asynchronous combination (Figure 8). In effect the tags would become independent markers for combinations of stimulus properties[89], and their addition to the stimulus representation

---

"the advertising principle." Other contemporary broadcast-based, coordinative frameworks assume connectionist substrates (Baars, 1988) to account for global integration of information and the unity of conscious awareness.

[87] *For example, rhythms set up temporal patterns of expectations (Jones, 1976) that might then be cross-correlated with incoming sensory signals to either build up or extinguish particular expectations. These mechanisms are similar in spirit to the adaptive-resonance models of Grossberg, e.g. (Grossberg, 1988; Grossberg, 1995), except that they are implemented using temporal codes and temporal correlation mechanisms rather than through channel-codes and spatial cross-correlation mechanisms.*

[88] *(John et al, 1973; Thatcher and John, 1977; John and Schwartz, 1978).*

[89] *This is not unlike the means by which major and minor peaks associated with the fundamental period (and the pitch that is heard) appeared in population-interval distributions when several harmonics were sounded together (Figure 6). The pattern of major and minor peaks arises from common periodicities present in the individual harmonics. Thus this "emergent" pattern is a property of the whole harmonic complex, with the major peaks corresponding to pitch and the minor peak patterns corresponding to timbre (Cariani, 1997b). The global pattern serves as a marker for the perceptual properties (pitch, timbre) that emerge when the individual harmonics are sounded together. The autocorrelation-interval pattern associated with the whole coexists with those associated with the individual harmonics, such that an elaboration of representation is carried out rather than information compression into simplified decisions (symbols).*

would allow other neural assemblies to operate on them. Instead of discarding earlier, "lower-level" representations of the stimulus, as occurs in sequential, hierarchical processing, such time patterns would continue to be propagated throughout the system (Figure 7).

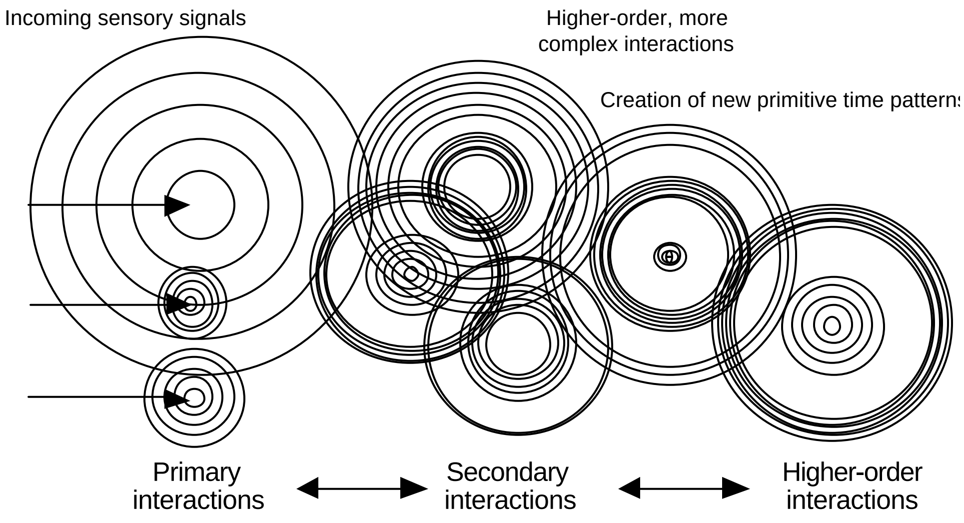

Figure 7. Time-coded broadcast schema for asynchronous, heterarchical global integration.

As a result, unique aspects of associations (combinations of stimulus properties) can be represented by themselves without rendering the lower-level primitives inaccessible (Figure 8).

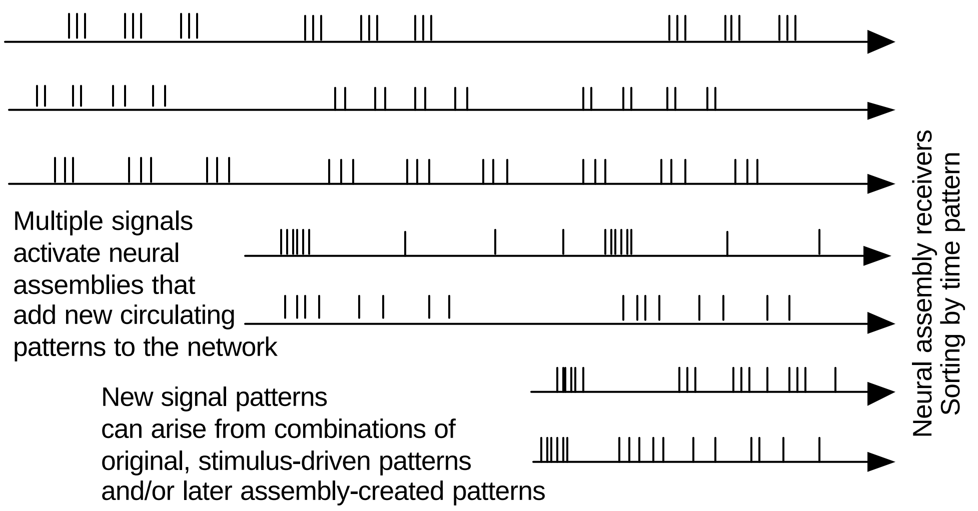

Figure 8. Elaboration of signals through successive and recurrent interactions.

In such a system new kinds of temporal tags would be constantly evolving, concomitant with the formation of new concepts. Such signal-primitives would have many of the advantages of both traditional symbolic logics and distributed connectionist ones.[90] Their production would be distributed across many neural elements, as in a connectionist network, yet the specific time patterns themselves would be unitary, as are the symbols of classical logics. Complex time codes thus could conceivably implement many of the same kinds of multidimensional, multivalent logics that are beginning to be actively explored in cognitive science.

Acknowledgments

I would like to thank Mrs. Anne Wiseman, who taught me the value of dimensional analysis through her chemistry problems. This work was supported by NIH Grant DC03054 from the National Institute for Deafness and Communications Disorders, a part of the National Institutes of Health.

References

---

[90] *(Horgan and Tienson, 1996). This was an issue at the heart of a debate almost a decade ago between Fodor and Pylyshyn vs. Smolensky over the respective capabilities and limitations of connectionist vs. symbolic approaches to artificial intelligence.*

Abeles, M. 1982a. *Local Cortical Circuits. An Electrophysiological Study*. Berlin: Springer-Verlag.

Abeles, Moshe. 1982b. Role of the cortical neuron: integrator or coincidence detector. *Israel Journal of Medical Sciences* 18: 83-92.

Abeles, Moshe. 1990. *Corticonics*. Cambridge: Cambridge University Press.

Abeles, M., H. Bergman, E. Margalit, and E. Vaadia. 1993. Spatiotemporal firing patterns in the frontal cortex of behaving monkeys. *J. Neurophysiol*. 70 (4): 1629-1638.

Arbib, Michael A. 1989. *The Metaphorical Brain* 2: Neural Nets and Beyond. New York: John Wiley.

Baars, Bernard J. 1988. *A Cognitive Theory of Consciousness*. Cambridge: Cambridge University Press.

Barlow, H.B. 1972. Single units and sensation: a neuron doctrine for perceptual psychology. *Perception* 1: 371-392.

Barlow, H.B. 1995. The neuron doctrine in perception. In *The Cognitive Neurosciences*. Edited by M. S. Gazzaniga. 415-435. Cambridge: MIT Press.

Basar, E. 1990. Chaotic dynamics and resonance phenomena in brain function: progress,, perspectives, and thoughts. In *Chaos in Brain Function*. Edited by E. Basar. 1-30. Berlin: Springer-Verlag.

Bergson, Henri. 1911. *Creative Evolution*. Modern Library Edition, 1944 ed. Translated by Mitchell, Arthur. New York: Henry Holt, and Company.

Bialek, W., F. Rieke, R. R. van Stevenink, and Warland de Ruyter, D. 1991. Reading a neural code. *Science* 252 (28 June): 1854-1856.

Boden, Margaret. 1990a. *The Creative Mind*. New York: Basic Books.

Boden, Margaret. 1990b. What is creativity? In *Dimensions of Creativity.* 75-117. New York: MIT Press.

Bonabeau, Eric, Jean-Louis Desalles, and Alain Grumbach. 1995. Characterizing emergent phenomena (1): A critical review. *Revue Internationale de Systemique* 9 (3): 327-371.

Bonabeau, Eric W., and Guy Theraulaz. 1995. Why do we need artifical life? In *Artifical Life: An Overview*. Edited by C. G. Langton. 75-110. Cambridge, MA: MIT Press.

Boring, Edwin G. 1933. *The Physical Dimensions of Consciousness*. New York: Dover.

Boring, Edwin G. 1942. *Sensation and Perception in the History of Experimental Psychology*. New York: Appleton-Century-Crofts.

Braitenberg, V. 1961. Functional interpretation of cerebellar histology. *Nature* 190: 539-540.

Braitenberg, Valentino. 1967. Is the cerebellar cortex a biological clock in the millisecond range? *Prog. Brain Res.* 25: 334-346.

Braitenberg, V., Detlef Heck, and Fahad Sultan. in press. The detection and generation of sequences as a key to cerebellar function. Experiments and theory. *Brain and Behavioral Sciences* 20(2):229-277.

Bregman, Albert S. 1990. *Auditory Scene Analysis: The Perceptual Organization of Sound*. Cambridge, MA: MIT Press.

Bridgman, P. W. 1931. *Dimensional Analysis*. New Haven: Yale University Press.

Bullock, T.H. 1967. Signals and neural coding. In *The Neurosciences: A Study Program*. Edited by G. C. Quarton, T. Melnechuck and F. O. Schmitt. 347-352. New York: Rockefeller University Press.

Carello, Claudia, Michael T. Turvey, Peter N. Kugler, and Robert E. Shaw. 1984. Inadequacies of the computer metaphor. In *Handbook of Cognitive Neuroscience*. Edited by M. S. Gazzaniga. 229-248. New York: Plenum Press.

Cariani, Peter. 1989a. Adaptivity, emergence, and machine-environment dependencies. 33rd Annual Metting of the International Society for Systems Sciences (ISSS), Edinburgh, Scotland, July 2-7, 1989. Edited by P. W. J. Ledington. III: 31-37.

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

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

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

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

Cariani, P. 1995. As if time really mattered: temporal strategies for neural coding of sensory information. *Communication and Cognition - Artificial Intelligence* (CC-AI) 12 (1-2): 161-229. Reprinted in: K Pribram, ed. Origins: Brain and Self-Organization, Hillsdale, NJ: Lawrence Erlbaum, 1994; 208-252.

Cariani, P. 1997a. Consciousness and the organization of neural processes: comments on John et al. *Consciousness and Cognition* 6 (1): 56-64.

Cariani, P. 1997b. Neural representation of pitch through autocorrelation. Proceedings, Audio Engineering Society Meeting (AES), New York, September, 1997, Preprint #4583 (L-3).

Cariani, P. 1997c. Temporal coding of sensory information. In *Computational Neuroscience*: Trends in Research, 1997. Edited by J. M. Bower. 591-598. New York: Plenum.

Cariani, P. 1998a. Neural computations in the time domain. *Assoc. Res. Otolaryn. Abstr.*

Cariani, P. 1998b. Neural timing nets for auditory computation. Internal proceedings, NATO Adv. Study Inst. on Computational Hearing, Lucca, Italy, July, 1998.: 6.

Cariani, Peter A., and Bertrand Delgutte. 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.

Carr, Catherine E. 1993. Processing of temporal information in the brain. *Annu. Rev. Neurosci*. 16: 223-243.

Casseday, John H., and Ellen Covey. 1995. Mechanisms for analysis of auditory temporal patterns in the brainstem of echolocating bats. In *Neural Representation of Temporal Patterns*. Edited by E. Covey, H. L. Hawkins and R. F. Port. 25-51. New York: Plenum.

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

Chappelier, J.C., and A. Grumbach. 1994. Time in neural networks. *SIGART Bulletin* 5 (3): 3-11.

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

Cherry, Colin. 1961. Two ears – but one world. In *Sensory Communication*. Edited by W. A. Rosenblith. 99-117. New York: MIT Press/John Wiley.

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

Churchland, Patricia S., and Terrence J. Sejnowski. 1992. *The Computational Brain*. Cambridge: MIT Press.

Clark, Andy. 1996. Happy couplings: emergence and explanatory interlock. In *The Philosophy of Artificial Life.* Edited by M. Boden. 262-281. Oxford: Oxford University Press.

Conrad, Michael. 1983. *Adaptability.* New York: Plenum Press.

Covey, Ellen. 1980. *Temporal Neural Coding in Gustation*. Ph.D., Duke University.

de Ribaupierre, Francois. 1997. Acoustical information processing in the auditory thalamus and cerebral cortex. In *The Central Auditory System*. Edited by G. Ehret and R. Romand. 317-397. New York: Oxford University Press.

Denbigh, K. G. 1975. *An Inventive Universe*. New York: George Braziller.

Di Lorenzo, Patricia M., and Gerald S. Hecht. 1993. Perceptual consequences of electrical stimulation in the gustatory system. *Behavioral Neuroscience* 107: 130-138.

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

Dreyfus, Hubert L., and Stuart E. Dreyfus. 1987. How to stop worrying about the frame problem even though it's computationally insoluable. In *The Robot's Dilemma: The Frame Problem in Artificial Intelligence*. Edited by Z. Pylyshyn. 95-112. Norwood, NJ: Ablex.

Eckhorn, R. 1991. Stimulus-specific synchronization in the visual cortex: linking of local features into global figures? In *Neuronal cooperativity*. Edited by J. Kruger. 184-224. Berlin: Springer-Verlag.

Eckhorn, R., and H. Reitboeck. 1990. Stimulus-specific synchronization in cat visual cortex and its possible roles in visual pattern recognition. In *Synergetics of Cognition*. Edited by H. Haken and M. Stadler. 99-111. Berlin: Springer-Verlag.

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

Emmers, Raimond. 1981. *Pain: A Spike-Interval Coded Message in the Brain*. New York: Raven Press.

Evans, E.F. 1978. Place and time coding of frequency in the peripheral auditory system: some physiological pros and cons. *Audiology* 17: 369-420.

Feldman, Jerome A. 1990. Computational constraints on higher neural representations. In *Computational Neuroscience*. Edited by E. L. Schwartz. 163-178. Cambridge: MIT Press.

Gershenfeld, Neil. 1988. An experimentalist's introduction to the observation of dynamical systems. In *Directions in Chaos*, Vol. II. Edited by H. Bai-lin. Singapore: World Scientific.

Gibson, Eleanor J. 1969. *Principles of Perceptual Learning and Development*. Cambridge: Appleton-Century-Crofts.

Gibson, Eleanor J. 1991. *An Odyssey in Learning and Perception*. Cambridge: MIT Press.

Graf, K. E., and T. Elbert. 1989. Dimensional analysis of the waking EEG. In *Brain Dynamics*. Edited by E. Basar and T. E. Bulloch. 176-191. Berlin: Springer-Verlag.

Grossberg, Stephen. 1988. *The Adaptive Brain*, Vols I. and II. New York: Elsevier.

Grossberg, Stephen. 1995. Neural dynamics of motion perception, recognition learning, and spatial attention. In *Mind as Motion: Explorations in the Dynamics of Cognition*. Edited by R. F. Port and T. van Gelder. 449-490. Cambridge: MIT Press.

Heiligenberg, Walter. 1994. The coding and processing of temporal information in the electrosensory system of the fish. In *Temporal coding in the brain*. Edited by G. Buzaki, R. Llinas, W. Singer, A. Berthoz and Y. Christen. 1-12. Berlin: Springer-Verlag.

Holland, John. 1998. *Emergence. Reading,* Massachusetts: Addison-Wesley.

Horgan, Terence, and John Tienson. 1996. *Connectionism and the Philosophy of Psychology*. Cambridge, MA: MIT Press.

Hutchinson, Michael. 1991. *Megabrain*. New York: Ballantine.

Jeffress, L. A. 1948. A place theory of sound localization. *J. Comp. Physiol. Psychol.* 41: 35-39.

John, E. R. 1967a. Electrophysiological studies of conditioning. In *The Neurosciences: A Study Program*. Edited by G. C. Quarton, T. Melnechuk and F. O. Schmitt. 690-704. New York: Rockefeller University Press.

John, E. R. 1967b. *Mechanisms of Memory*. New York: Wiley.

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 *Dynamics of Sensory and Cognitive Processing by the Brain*. Edited by E. Basar. 368-377. Berlin: Springer-Verlag.

John, E. R., F. Bartlett, M. Shimokochi, and D. Kleinman. 1973. Neural readout from memory. *J. Neurophysiol*. 36 (5): 893-924.

John, E. Roy, and Eric L. Schwartz. 1978. The neurophysiology of information processing and cognition. *Ann. Rev. Psychol*. 29: 1-29.

Jones, Mari Riess. 1976. Time, our lost dimension: toward a new theory of perception, attention, and memory. *Psychological Review* 83 (5): 323-255.

Kabrisky, Matthew. 1966. A Proposed Model for Visual Information Processing in the Human Brain. *Urbana*, Ill.: University of Illinois.

Kampis, György. 1988. On the modelling relation. *Systems Research* 5 (2): 131-144.

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

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

Kauffman, Stuart. 1993. *The Origins of Order*. New York: Oxford University Press.

Kiang, N.Y.S., T. Watanabe, E.C. Thomas, and L.F. Clark. 1965. *Discharge Patterns of Single Fibers in the Cat's Auditory Nerve.* Cambridge: MIT Press.

Klee, Robert L. 1984. Micro-determinism and concepts of emergence. *Philosophy of Science* 51: 44-63.

Köhler, Wolfgang. 1951. Relational determination in perception. In *Cerebral Mechanisms in Behavior:* The Hixon Symposium. Edited by L. A. Jeffress. 200-243. New York: Wiley.

Kohonen, T. 1988. *Self-organization and associative memory*. New York: Springer-Verlag.

Kolen, J. F., and J. B. Pollack. 1995. The observer's paradox: apparent computational complexity in physical systems. *J Experimental and Theoretical* AI 7: 253-277.

Langner, G. 1983. Evidence for neuronal periodicity detection in the auditory system of the Guinea Fowl: implications for pitch analysis in the time domain. *Exp. Brain Res*. 52: 33-355.

Langner, G. 1985. Time coding and periodicity pitch. In *Time Resolution in Auditory Systems.* Edited by A. Michelsen. 108-121. Berlin: Springer-Verlag.

Langner, Gerald. 1992. Periodicity coding in the auditory system. *Hearing Research* 60: 115-142.

Lashley, K. S. 1951. The Problem of Serial Order in Behavior. In *Cerebral Mechanisms of Behavior*. Edited by L. A. Jeffress. 112-146. New York: Wiley.

Lestienne, R. 1996. Determination of the precision of spike timing in the visual cortex of anesthetized cats. *Biol. Cybern*. 74: 55-61.

Lestienne, R., E. Gary-Bobo, J. Przybyslawski, P. Saillour, and M. Imbert. 1990. Temporal correlations in modulated evoked responses in the visual cortical cells of the cat. *Biol. Cybern*. 62: 425-440.

Lestienne, R., and B.L. Strehler. 1987. Time structure and stimulus dependence of precise replicating patterns present in monkey cortical neuronal spike trains. *Brain Research* 43: 214-238.

Licklider, J.C.R. 1951. A duplex theory of pitch perception. *Experientia* VII (4): 128-134. Experientia.

Licklider, J.C.R. 1956. Auditory frequency analysis. In *Information Theory*. Edited by C. Cherry. 253-268. London: Butterworth.

Licklider, J.C.R. 1959. Three auditory theories. In Psychology: A Study of a Science. *Study I. Conceptual and Systematic*. Edited by S. Koch. 41-144. New York: McGraw-Hill.

Llinas, R., and D. Pare. 1996. The brain as a closed system modulated by the senses. In *The Mind-Brain Continuum*. Edited by R. Llinas and P. S. Churchland. 1-18. Cambridge: MIT Press.

Longuet-Higgins, H.C. 1987. *Mental Processes: Studies in Cognitive Science*. Cambridge, Mass.: The MIT Press.

Longuet-Higgins, H.C. 1989. A mechanism for the storage of temporal correlations. In *The Computing Neuron*. Edited by R. Durbin, C. Miall and G. Mitchison. 99-104. Wokingham, England: Addison-Wesley.

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

MacKay, D.M. 1962. Self-organization in the time domain. In *Self-Organizing Systems* 1962. Edited by M. C. Yovitts, G. T. Jacobi and G. D. Goldstein. 37-48. Washington, D.C.: Spartan Books.

Maruyama, Margaroh. 1977. Heterogenistics: an epistemological restructuring of biological and social sciences. *Cybernetica* 20: 69-86.

McCorduck, Pamela. 1972. Machines Who Think. San Francisco: W.H. Freeman.

McCulloch, W.S. 1951. Why the Mind is in the Head. In *Cerebral Mechanisms of Behavior* (the Hixon Symposium). Edited by L. A. Jeffress. 42-111. New York: Wiley.

Meddis, R., and M. J. Hewitt. 1991. Virtual pitch and phase sensitivity of a computer model of the auditory periphery. I. Pitch identification. *J. Acoust. Soc. Am.* 89 (6): 2866-2882.

Miller, Ralph R., and Robert C. Barnet. 1993. The role of time in elementary associations. *Current Directions in Psychological Science* 2 (4): 106-111.

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

Morley, J.W., J.S. Archer, D.G. Ferrington, M.J. Rowe, and A.B. Turman. 1990. Neural coding of complex tactile vibration. In *Information Processing in Mammalian Auditory and Tactile Systems*. 127-140. Alan R. Liss, Inc.

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

Mountcastle, Vernon. 1967. The problem of sensing and the neural coding of sensory events. In *The Neurosciences: A Study Program*. Edited by G. C. Quarton, Melnechuk, T., and Schmitt, F.O. New York: Rockefeller University Press.

Mountcastle, Vernon. 1993. Temporal order determinants in a somatosthetic frequency discrimination: sequential order coding. *Annals New York Acad. Sci.* 682: 151-170.

Mountcastle, V.B., W.H. Talbot, H. Sakata, and J. Hyvärinen. 1969. Cortical neuronal mechanisms in flutter-vibration studied in unanesthetized monkeys. Neuronal periodicity and frequency discrimination. *J. Neurophysiol.* 32: 452-485. Journal of Neurophysiology.

Nagel, Ernest. 1961. *The Structure of Science*. New York: Harcourt, Brace & World.

Pabst, M., H.J. Reitboeck, and R. Eckhorn. 1989. A model of preattentive texture region definition based on texture analysis. In *Models of Brain Function*. Edited by R. M. J. Cotterill. 137-150. Cambridge: Cambridge University Press.

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

Pask, Gordon. 1960. The natural history of networks. In *Self-Organzing Systems*. Edited by M. C. Yovits and S. Cameron. 232-263. New York: Pergamon Press.

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

Pattee, Howard H. 1973a. The physical basis of the origin of hierarchical control. In *Hierarchy Theory: The Challenge of Complex Systems*. Edited by H. Pattee. New York: George Braziller.

Pattee, Howard H. 1973b. Physical problems in the origin of natural controls. In *Biogenesis, Homeostasis, Evolution*. Edited by A. Locker. New York: Pergamon Press.

Pattee, Howard H. 1989. Simulations, realizations, and theories of life. In *Artificial Life*. Volume VI, Santa Fe Institute Studies in the Science of Complexity. Edited by C. G. Langton. 63-78. Redwood City, CA: Addison-Wesley.

Pattee, Howard H. 1995. Artificial life needs a real epistemology. In *Advances in Artificial Life*. Edited by F. Morán, A. Moreno, J. J. Merelo and P. Chacón. 23-38. Berlin: Springer-Verlag.

Pellionisz, A. J. 1991. Geometry of massively parallel neural interconnectedness in wired networks and wireless volume transmission. In *Volume Transmission in the Brain: Novel Mechanisms for Neural Transmission*. Edited by K. Fuxe and L. F. Agnati. 557-568. New York: Raven Press.

Perkell, D.H, and T.H. Bullock. 1968. Neural Coding. *Neurosciences Research Program Bulletin* 6 (3): 221-348.

Phillips, D.P. 1989. Timing of spike discharges in cat auditory cortex neurons: implications for encoding of stimulus periodicity. *Hearing Research* 40: 137-146.

Piatelli-Palmarini, Massimo. 1980. *How hard is the hard core of a scientific paradigm?* Cambridge, MA: Harvard University Press.

Pratt, Gill. 1990. *Pulse Computation*. Ph.D., M.I.T.

Prigogine, Illya. 1980. *From Being to Becoming*. San Francisco: W. H. Freeman.

Rapp, Paul E. 1993a. Chaos in the neurosciences: cautionary tales from the frontier. *Biologist* 40 (2): 89-94.

Rapp, Paul E. 1993b. Filtered noise can mimic low-dimensional attractors. *Physical Review* E 47 (4): 2289-2297.

Rapp, Paul E. 1994. A guide to dynamical analysis. *Integrative Physiological and Behavioral Science* 29 (3): 311-327.

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. G. Waxman. New York: Raven Press.

Reichardt, Werner. 1961. Autocorrelation, a principle for the evaluation of sensory information by the central nervous system. In *Sensory Communication*. Edited by W. A. Rosenblith. 303-317. New York: MIT Press/John Wiley.

Reichardt, Werner, and Tomaso Poggio. 1976. Visual control of orientation behaviour in the fly. Part I. A quantitative analysis. *Quart. Rev. Biophys.* 9 (3): 311-375.

Reitboeck, Herbert J. 1989. Neural mechanisms of pattern recognition. In *Sensory Processing in the Mammalian Brain*. Edited by J. S. Lund. 307-330. Oxford: Oxford University Press.

Reitboeck, H.J., M. Pabst, and R. Eckhorn. 1988. Texture description in the time domain. In *Computer Simulation in Brain Science*. Edited by R. M. J. Cotterill. Cambridge, England: Cambridge University Press.

Richmond, B.J., and T.J. Gawne. 1998. The relationship between neuronal codes and cortical organization. In *Neuronal Ensembles: Strategies for Recording and Decoding*. Edited by H. B. Eichenbaum and J. L. Davis. 57-80. New York: Wiley-Liss.

Richmond, B.J., L.M. Optican, and T.J. Gawne. 1989. Neurons use multiple messages encoded in temporally modulated spike trains to represent pictures. In *Seeing Contour and Colour*. Edited by J. J. Kulikowski and C. M. Dickenson. 705-713. New York: Pergamon Press.

Richmond, B.J. , L. M. Optican, M. Podell, and H. Spitzer. 1987a. Temporal encoding of two-dimensional patterns by single units in primate inferior temporal cortex. I. Response characteristics. *J. Neurophysiol.* 57 (1): 132-146. Journal of Neurophysiology.

Richmond, B.J. , L. M. Optican, M. Podell, and H. Spitzer. 1987b. Temporal encoding of two-dimensional patterns by single units in primate inferior temporal cortex. I. Response characteristics. II. Quantification of response waveform III. Information-theoretic analysis. *J. Neurophysiology* 57 (1): 132-178. Journal of Neurophysiology.

Rieke, Fred, David Warland, Rob de Ruyter van Steveninck, and William Bialek. 1997. *Spikes: Exploring the Neural Code*. Cambridge: MIT Press.

Rose, J.E., J.R. Brugge, D.J. Anderson, and J.E. Hind. 1967. Phase-locked response to low-frequency tones in single auditory nerve fibers of the squirrel monkey. *J. Neurophysiol.* 30: 769-793. Journal of Neurophysiology.

Rosen, Robert. 1973. On the generation of metabolic novelties in evolution. In *Biogenesis, Homeostasis, Evolution*. Edited by A. Locker. New York: Pergamon Press.

Rosen, Robert. 1978. *Fundamentals of Measurement and Representation of Natural Systems*. New York: North-Holland.

Rosen, Robert. 1985. Organisms as causal systems which are not mechanisms: an essay into the nature of complexity. In *Theoretical Biology and Complexity. Three Essays on the Natural Philosophy of Complex Systems*. Edited by R. Rosen. Orlando, FL: Academic Press.

Rosen, Robert. 1986. Causal structures in brains and machines. *International Journal of General Systems* 12: 107-126.

Rosen, Robert. 1987. On the scope of syntactics in mathematics and science: the machine metaphor. In *Real Brains Artificial Minds*. Edited by J. Casti and A. Karlqvist. New York: North Holland.

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

Rosen, Robert. 1994. Effective processes and natural law. In *The Universal Turing Machine. A Half Century Survey*. Edited by R. Herken. 485-498. Vienna: Springer-Verlag.

Schyns, Phillippe G., Robert L. Goldstone, and Jean-Pierre Thibaut. in press. The development of features in object concepts. *Behavioral and Brain Sciences* 21(1):1-54.

Simmons, J. A. 1996. Formation of perceptual objects from the timing of neural responses: target-range images in bat sonar. In *The Mind-Brain Continuum*. Edited by R. Llinas and P. S. Churchland. 219-250. Cambridge: MIT Press.

Singer, W. 1990. Search for coherence: a basic principle of cortical self-organization. *Concepts in Neuroscience* 1 (1): 1-26.

Singer, W. 1995. Putative functions of temporal correlations in neocortical processing. In *Large-Scale Neuronal Theories of the Brain*. Edited by C. Koch and J. L. Davis. 201-237. Cambridge: MIT Press.

Softky, William R., and Christof Koch. 1993. The highly irregular firing of cortical cells is inconsistent with temporal integration of random EPSPs. *J. Neuroscience* 13 (1): 334-350.

Steels, Luc. 1995. The artificial life roots of artificial intelligence. In *Artifical Life: An Overview*. Edited by C. G. Langton. 75-110. Cambridge, MA: MIT Press.

Strehler, Bernard L., and Remy Lestienne. 1986. Evidence of precise time-coded symbols and memory of patterns in monkey cortical neuronal spike trains. *Proc. Nat. Acad. Sci.* 83: 9812-9816.

Tank, D. W., and J.J. Hopfield. 1987. Neural computation by concentrating information in time. *Proc. Natl. Acad. Sci.* USA 84: 1896-1900.

Thatcher, Robert W., and E. Roy John. 1977. *Functional Neuroscience,* Vol. I. *Foundations of Cognitive Processes*. Hillsdale, NJ: Lawrence Erlbaum.

Troland, Leonard T. 1929a. *The Principles of Psychophysiology,* Vols I-III. New York: D. Van Nostrand.

Troland, L.T. 1929b. The psychophysiology of auditory qualities and attributes. 2: 28-58. *The Journal of General Psychology*.

Uttal, W.R. 1973. *The Psychobiology of Sensory Coding*. New York: Harper and Row.

Uttal, W.R. 1975. *An Autocorrelation Theory of Form Detection*. New York: Wiley.

Uttal, W.R. 1988. *On Seeing Forms*. Hillsdale, NJ: Lawrence Erlbaum.

Vaadia, E., I. Haalman, M. Abeles, H. Bergman, Y. Prut, H. Slovin, and A. Aertson. 1995. Dynamics of neuronal interactions in monkey cortex in relation to behavioral events. *Nature* 373: 515-518.

van de Vijver, Gertrudis. 1991. The emergence of meaning and the antinomy of naturalism. *Uroboros* I (2): 153-175.

van Gelder, Timothy, and Robert F. Port. 1995. It's about time: an overview of the dynamical approach. In *Mind as Motion: Explorations in the Dynamics of Cognition*. Edited by R. F. Port and T. van Gelder. 1-44. Cambridge: MIT Press.

von der Malsberg, C. 1995. The binding problem of neural networks. In *The Mind-Brain Continuum*. Edited by R. Llinas and P. S. Churchland. 131-146. Cambridge: MIT Press.

von der Malsburg, C., and W. Schneider. 1986. A neural cocktail party processor. *Biological Cybernetics* 54: 29-40.

Walter, Grey. 1959a. Intrinsic rhythms of the brain. In *Handbook of Physiology: Neurophysiology*. Volume II. Edited by J. Field, H. W. Magoun and V. E. Hall. 279-298. Washington, D.C.: American Physiological Society.

Walter, Grey. 1959b. *The Living Brain*. New York: Norton.

Walter, Grey W. 1953. Theoretical properties of diffuse projection systems in relation to behavior and consciousness. In *Brain Mechanisms and Consciousness*. Edited by E. D. Adrian, F. Bremer and H. H. Jasper. 345-375. Springfield, IL: Charles C. Thomas.

Wasserman, Gerald S. 1992. Isomorphism, task dependence, and the multiple meaning theory of neural coding. *Biol. Signals* 1: 117-142.

Waxman, S.G. 1978. Regional differentiation of the axon: A review with special reference to the concept of the multiplex neuron. In *Physiology and Pathobiology of Axons*. Edited by S. G. Waxman. New York: Raven Press.

Wertheimer, Max. 1923. Untersuchungen zur Lebre von der Gestalt. Psychol. Forsuch. 4: 301-350. Abridged translation reprinted in: Beardslee, David C. and Wertheimer, Michael, *Readings in Perception*. Princeton, NJ: vanNostrand, pp. 115-135.

Wever, Ernest Glen. 1949. *Theory of Hearing*. New York: Wiley.

Young, R.A. 1977. Some observations on temporal coding of color vision: psychophysical results. *Vision Research* 17: 957-965.