

What details of neural circuits matter?

David W. Tank

Models of biological neural circuits are examined with an emphasis on how the computations they perform are affected by physical details of synaptic connectivity and single unit dynamics. Neural circuits subserving sound localization in the barn owl and range determination of prey in echo-locating bats are used to demonstrate the direct relationship between the computation in a circuit and the combination of the static (time-independent) and dynamic (time-dependent) functions the nervous system uses to perform the computation. The static and dynamic functions are the details that matter for computation in these two circuits. The detailed physical attributes that implement, contribute to or modify these functions are necessary but not necessarily unique: different physical mechanisms can implement the same function. Similar conclusions are demonstrated with several artificial neural network models. Finally, a computational approach to central pattern generator (CPG) circuits is examined where forms of neuromodulation distinguish between CPGs of fundamentally different computational logic that may have identical patterned output in the absence of neuromodulation.

Key words: neural networks / dynamics / computation / neuromodulation

NEUROPHYSIOLOGISTS studying biological neural circuits often ask what details are important for understanding the computational principles of a circuit. Several contemporary theoretical papers discuss the computational power of neural network models where the units (i.e. artificial neurons) are greatly simplified caricatures of biological neurons and, in some cases, communicate through random synaptic connections. If these artificial networks can produce such useful computation, why do real neural circuits have such complex physical mechanisms underlying the simplest cellular properties and why are the patterns of connectivity far from random in the simpler circuits that have been studied? The physiologist quite naturally thinks, 'Either these theorists are crazy or I have been wasting my time

on trivial details that are irrelevant to the functional operation of the nervous system I study'.

Here I consider which physical 'details' are important in several biological neural circuits and artificial neural networks. The analysis is pursued only from the perspective of how information is transformed or utilized by the circuit or network and ignores the equally important and interesting issues of phylogeny and ontogeny. Examples of computational models built on known biological circuits are used to demonstrate the central theme that there is always a direct relationship between the computation performed by a circuit and the combination of static and dynamic functions that the circuit implements. Circuits that have a different organization of these functions compute different things. In some cases, simple mathematical relationships can describe the circuit organization in a concise way. Characterizing the dynamic functions used in a biological circuit is of absolutely central importance to understanding circuit function. *Physical details of the circuit that produce or affect the combination of static and dynamic functions are the essence of computational function. But, a necessary and often overlooked qualification is that several physical mechanisms can implement the same function.* Thus computationally identical neural circuits in different nervous systems could, in principle, appear quite unrelated when examined at the level of the biophysical mechanisms and pattern of synaptic organization. These differences are reduced when the possible homologies between dynamic functions are taken into consideration.

The models examined here demonstrate that it is the relationship between dynamic function and computation that is important. Using one physical mechanism or another one to implement that dynamics is a choice made by the nervous system for other reasons such as evolution, constraints during development, conservation of energy or constituent parts. The building blocks of neural circuit *construction* are the biophysical and biochemical mechanisms found in neurons. The building blocks of neural circuit *computation* are time-independent static functions, time-dependent dynamic functions, and the interrelationships between them. Within this context, neuromodulation can be viewed conceptually as

From the Molecular Biophysics Research Department, AT&T Bell Laboratories, Murray Hill, NJ 07974, USA
©1989 by W.B. Saunders Company
1044-5765/89/0101-0009\$5.00/0

affecting static and dynamic function. The important details of neuromodulation are the classes of changes in function that result and the inherent relationship of these changes to computation.

Computation and dynamic function in specific neural circuits

Sound localization in the barn owl

Although no biological neural circuit is known in sufficient detail for its structure to be related definitively to the computations it performs, several circuits are known in enough detail to begin to see these relationships and what lessons can be learned from them. A good example is provided by the circuitry underlying azimuthal sound localization in the barn owl.¹

The computation is the determination of the position of a sound source in the horizontal plane by comparing the sound wave arriving at the two ears. The basic neural model for how this computation could, in principle, be performed was provided by Jeffress in 1948.² By comparing the arrival times of similar features in the acoustic signal at left and right ears, the lateral position of a sound source can be computed; the difference in arrival times is related to the angular position of the source relative to the head and the velocity of sound in air. The model is based upon a cross-correlation of signals from left and right ears, and Figure 1a shows the anatomy and physiology of its implementation in the nucleus laminaris (NL) of the barn owl, suggested by the work of Sullivan and Konishi¹ and Carr and Konishi.³ The NL is a laminar structure containing isofrequency slabs, each of which receives input fibers tuned to the same characteristic frequency from both contra- and ipsilateral sides of the head. These fibers enter the NL from opposite sides and course along each other, synapsing on intrinsic neurons in the nucleus. Because fibers with a given characteristic frequency from the left and right ears run parallel to each other but propagate action potentials in opposite directions, a temporal conjunction between impulses in these fibers will occur with highest probability when the stimulus is at a unique spatial position. This is where the delay caused by a longer propagation time in the fiber from the ear closer to the source equals the extra time it takes for sound to travel to the more distant ear. Excitatory synapses from the fibers from both right

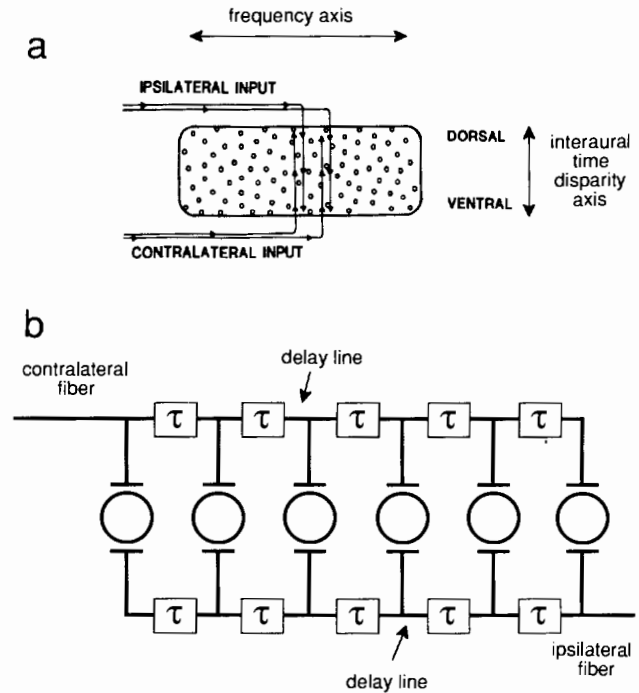


Figure 1. (a) Diagram of the anatomical organization of the nucleus laminaris of the barn owl. Contra- and ipsilateral afferent fibers enter from opposite sides of the nucleus. Interneurons (circles) are situated between fibers with similar frequency tuning but oriented in opposite propagation directions. Along the dorso-ventral axis, different interneurons in a given frequency slab respond best to the interaural time differences that correspond to the difference in axon length between the two sets of fibers. (Adapted from ref 1.) (b) Functional organization of the computation likely to be taking place in the nucleus laminaris. The boxes along the axons represent delays of duration τ . The same organization is used in the sound localization neural network designed and fabricated in analog silicon microelectronics⁴ by Mead's group.

and left ears converging on a neuron at this location should drive it strongly and thus produce a greater response than at other positions along the paths of the parallel axons.

How information in different isofrequency bands in the NL is combined and used by higher parts of the auditory system is not completely understood. Nevertheless it is clear that a decision could be made about the actual time difference for an acoustic stimulus by taking a sum of all of the outputs of NL interneurons in different isofrequency slabs that are tuned to a specific interaural time difference and comparing this sum to that similarly computed for each of the other interaural time differences: the maximum sum determines the azimuthal position of the sound source.

Despite the uncertainty about the operation of this system and the accuracy of this model, I will examine some of the details of the circuit that are essential to its performance and computation.

(1) Does the specific geometrical pattern of synaptic connections matter for the computation being performed? Absolutely: only identical ‘features’ of the acoustic signal can legitimately be compared to determine the interaural time difference, thus only identical information channels should do the temporal correlation computation. Performing a correlation analysis (impulse coincidence detection) on afferents representing different characteristic frequencies would be meaningless. This detail matters; the geometrical pattern of synaptic connection has a direct relationship to the computation being performed. (2) Is it important that voltage-dependent sodium channels are used to produce action potentials in the axons? Here the answer is not simple. From the computational perspective, the important dynamic characteristic of the axons is that the action potential propagates as a non-decrementing pulse at finite velocity: it is the combination of finite velocity with the difference in length between the left and right axons that produces an effective temporal disparity that ultimately relates to the interaural time difference. Any physical mechanism that can implement this dynamics could be used. For example, an action potential based primarily on fast voltage-dependent calcium conductances, as found in some invertebrate nervous systems, could substitute, as long as the connection patterns were appropriately tuned for any resulting alteration in propagation velocity.

The decision about where the sound comes from can be stated as: sum the outputs of neurons in the different isofrequency slabs that give the cross-correlation value for a given temporal disparity; the temporal disparity corresponding to the largest sum is identified with the azimuthal position of the sound source. Mathematically this can be summarized by:

$$\max (\text{over delays } \tau) \left[\sum_{\text{freq}} f(\langle \text{right}_{\text{freq}}(t) \text{left}_{\text{freq}}(t - \tau) \rangle) \right]$$

where, to a first approximation, max, sum and the NL interneuron ‘input-output’ relationship ($f()$) are time-independent static functions and the time delays and $\langle \rangle$ (determine recent coincidences) are time-dependent dynamic functions. The computational organization of this circuit, illustrated for one frequency channel, is shown in Figure 1b. The building blocks for the computation are the static and

dynamic functions and their interrelationships, which consist of an orderly series of delays, a correlation detection mechanism and summing together signals from each of the frequency channels at the same delay times. The computation is identified with, and indeed defined by, these building blocks: the functions and their intercommunication of information.

In the nervous system, these functions are performed by biophysical and biochemical mechanisms, which are the building blocks of neural circuit construction. For example, the resistance and capacitance of axons in combination with voltage-dependent conductances implement a delay line function, although we usually think of this simply as the propagating action potential. Combinations of physical properties, the details in a neural circuit, implement a building block of the neural computation by providing a static or dynamic function. This is the level at which physical details matter. We must look at the physical details and use our growing knowledge of how they could implement static and dynamic functions to see if they represent a functional class.

The example of delay lines resulting from different ionic conductances underlying propagating action potentials illustrates a general principle that different physical mechanisms can produce the same dynamic behavior and thus can substitute for each other to perform the same computation. Any individual circuit must contain a member of this set for its implementation. Dynamically analogous circuits performing the same computation may be found in widely different nervous system as well as in non-biological systems. This is wonderfully illustrated for the sound localization problem by an electronic neural network built in analog microelectronics by Carver Mead’s group.⁴ This integrated circuit performs essentially the same computation as that described above but is built entirely of non-biological components. The computation/dynamical details are conserved; the computational structure is that shown in Figure 1b; the pattern of connections is identical to that described above, with delay lines that propagate non-decrementing pulses at an appropriate velocity and interneurons that sum the activity from different frequency channels at each unit of temporal disparity. But the physical mechanisms used to implement the computational/dynamical details are quite different. For example, transmission lines (axon analogs) are silicon wires on the integrated circuit substrate and the action potential analogs are produced by regenerative transistor circuits at nodes

along the lines. Cross-correlations are performed by silicon neurons that sum the coincidences of current pulses.

This example illustrates the distinction between dynamical/computational class, details that matter for computation and the specific implementation chosen. A given computation ‘maps’ onto an architecture of interactions of static and dynamic functions, in this case the one shown in Figure 1b. In any individual circuit, whether in the owl or on a silicon chip, one member of the class of implementations for each function must be used but within a class the choice may be determined by historical, evolutionary or practical factors, or other factors not readily apparent.

FM/FM sensitive neurons in echo-locating bats

A second model circuit that nicely illustrates the use of different combinations of static and dynamic functions to implement a simple computation comes from the auditory system of bats that echo locate using frequency modulated (FM) signals. Suga⁵ has demonstrated that these signals are used to determine the range of a target such as prey. The bats have neurons that detect complex time-dependent acoustic features. For example, there are neurons that respond only to an FM sweep or chirp of sound. Other neurons in the auditory cortex are sensitive to more complex FM/FM combinations. These respond preferentially to two chirps of sound with

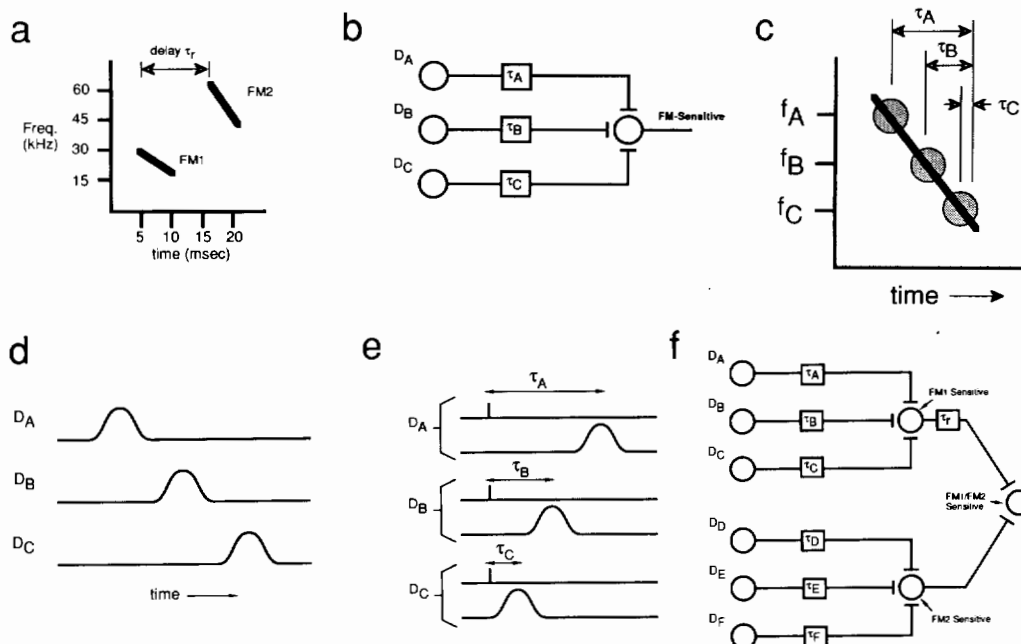


Figure 2. Target range determination using frequency modulated (FM) sweeps. (a) Schematized sound spectrogram of an acoustic signal containing one frequency sweep (FM1) followed with delay (τ_r) by a second sweep (FM2) at higher frequencies. Such a stimulus will elicit a large response from FM/FM combination sensitive neurons in the auditory cortex of some bats. (b) Functional organization of a model neural circuit containing an FM-sensitive output neuron that preferentially responds to an FM sweep. Detectors (D_A , D_B , D_C) with response zones shown as shaded circles in (c), respond to acoustic signals in narrow frequency ranges (f_A , f_B , and f_C , respectively) during an FM sweep. The activation of the detectors as a function of time for this sweep is shown in (d). Detectors project onto the FM-sensitive neuron with different delayed synaptic connections with the delayed responses shown in (e). Lower trace of each pair represents the time dependence of the synaptic current (impulse response function) in the postsynaptic neuron in response to a brief activation of the detector shown by the vertical line in the upper trace. With the set of delays illustrated, the output cell will respond maximally when presented with the downward sweep shown in (c). (f) Functional organization of the hierarchical circuit with an output neuron that is maximally responsive to the acoustic signal shown in (a). Units responsive to FM1 and FM2 frequency sweeps are built as in (b) and the synaptic output of the FM1 sensitive cell is delayed by time τ_r , the echo time for a chirp to travel from bat to prey and return.

different frequency ranges, such that the chirp with the lower range occurs a specific time delay τ_r before the second chirp. An example spectrogram of such an auditory stimulus is sketched in Figure 2a.

The two chirps and the delay τ_r have a computational meaning for the bat when it is detecting prey. It simultaneously emits several chirps (FM1 through FM4) in increasing frequency bands. The chirps stimulate the bat's auditory system immediately when the sounds are emitted; they also travel to the target and are echoed back to the bat. The bat compares the emitted low frequency chirp (FM1 in Figure 2a) with the returned echo (FM2 in Figure 2a) of the emitted chirp in the second frequency band. The time delay τ_r between the detection of the initial chirp and the echo, that is the propagation time of sound from bat to target and back to bat, thus represents the range of the target. Why does the bat use chirps for target ranging instead of brief pulses of constant frequency? It is likely to be for the same reason that chirp-frequency ranging is used in radar: by sending a sweep of frequencies, an emitter can send out a longer signal at less peak power dissipation while still maintaining good temporal discrimination of the difference in time between emitted and received signals.

A model circuit that captures the response characteristics of a neuron sensitive to a single FM1 sweep is shown in Figure 2b. During an FM1 sweep, detectors D_A , D_B , D_C , whose activation zones are illustrated by the shaded circles in Figure 2c, respond to acoustic signals in narrow frequency ranges (f_A , f_B , and f_C , respectively). The activation of these detectors as a function of time for the sweep is shown in Figure 2d. As shown in Figure 2b, if a short delay connects D_C to the FM1-sensitive neuron, with progressively longer delays from D_B and D_A , the sequence of activity f_A , then f_B , then f_C will produce a bigger input pulse in the FM1-sensitive neuron than any other sequence of activation. Likewise it is simple to build a second detector for the higher frequency echoed chirp, FM2, by connecting a set of graded delays from detectors D_D through D_F that respond to the higher set of frequencies. Finally, as illustrated in Figure 2f, the outputs of the FM1 and FM2 sensitive units are connected to an F1/F2 combination sensitive neuron, the FM1-sensitive neuron through a delay of duration τ_r and the FM2 sensitive neuron directly; this results in a unit whose preferential input stimulus is a sweep FM1 followed at delay τ_r with sweep FM2.

Although the hierarchical architecture of the circuit—neurons sensitive to single chirps connected together to form more complex feature-sensitive units—is based upon neurophysiological and anatomical evidence, the detailed form of the circuitry used in the bat is not known. The circuit shown in Figure 2b seems likely to reflect the operation of the real circuit because the structure of combining frequency bands with appropriate delay has a direct correspondence with the computation (detect sequential chirps with specific delay) that is actually performed.

Assuming this model circuit is correct, what details of it are important? First, the structure of the individual frequency sweeps has a direct correlate in the orderly set of time-delayed connections that project to an FM-sensitive unit. The delays must therefore be ordered with this specific form of connectivity to produce the desired response properties. Likewise, the FM/FM sensitive unit is connected to each of the two FM-sensitive units in a specific way: the neuron that measures the first FM sweep has a delayed effect on the FM/FM-sensitive unit, where the time of the delay is related to the distance from the bat to the prey. Similarly, the FM units must sum the activity of the incoming delay lines and produce a signal that is monotonically related to this sum.

But what about the details of physical implementation? The delays appropriate to the detection of the FM sweeps are in the millisecond time scale. Such delays could be generated by a variety of biophysical/biochemical mechanisms, four of which are sketched in Figure 3. Of these, rebound excitation⁷ and inactivation of the potassium A-current are particularly interesting from the computational perspective because the time constants of the delays in both could be tuned to different numerical values in a relatively straightforward manner. For example, in the I_A mechanism the full membrane depolarization expected by the onset of an excitatory synaptic current is delayed because at hyperpolarized membrane potential the I_A channel is activated. The initial partial depolarization produced by the excitatory synaptic current inactivates I_A , but this occurs with a characteristic inactivation time-constant. As inactivation continues, the membrane becomes more depolarized and can eventually reach the threshold for action potential generation, as has been demonstrated experimentally and by computer modeling for the inking circuitry in *Aplysia*.^{8,9} By changing the number of I_A channels




Figure 3. Different physical mechanisms that could contribute to the implementation of a dynamic function, the delayed excitation of the postsynaptic neuron. In general, a static or dynamic function that is used in a neural circuit may have several mechanisms of implementation. I_A is the potassium A current that can implement a delayed excitation (see text.) See ref 6 for a review of slow synaptic connections.

in the membrane, this threshold can occur with different delays from the onset of excitation. The following example demonstrates how such a mechanism could, in principle, be used.

A computational example of neuromodulation

The model circuit for echo location in bats is good for exploring how neuromodulation might be understood in a computational/dynamical framework. The function of the model circuit in Figure 2 is to detect an echoed chirp FM2 that occurs with delay τ_r after an emitted chirp FM1; if activated, prey is detected at a distance equal to (velocity of sound) $\cdot \tau_r/2$. As the bat closes in on its prey, it increases the rate at which it produces range-determining

chirps. It makes sense that it should also shorten the distances to which its auditory range detectors are sensitive by decreasing the delay τ_r detected by a given neuron decreases as the target is approached, an effect that is experimentally observed.⁵

How might this shift be accomplished across the entire set of detectors for different target ranges, each with different characteristic delays? If the time delay representing τ_r , the connection from the output of the FM1-sensitive cell to the FM1/FM2-sensitive cell (Figure 2f), were produced by A-current channels in the postsynaptic cell membrane, reducing the density of these channels would reduce the time of delayed activation. If other parameters stayed equal, the set of FM1/FM2-sensitive neurons representing different τ_r values would then be sensitive to shorter delays between the two sweeps, thus tuned to shorter distances. The mechanisms for regulating the delays in the bat are unknown but a reasonable theoretical model might be that, as the bat increases the rate it produces emitted chirps, changes in active A-current channel density are produced by a concomitant neuromodulatory signal to the circuitry underlying FM1/FM2 combination sensitive neurons. This would shift the timescale of the system and hence the distance it codes for in unison.

By changing the characteristic distances across the population of range-sensitive neurons, the computation of the circuit is modulated. Associated with this is a *change in dynamics*, because the time delays in the circuit are altered. Other mechanisms beside a rapid change in the number of active A-current channels could produce similar changes in dynamics. For example, if the bat's nervous system uses rebound excitation to produce the circuit delays, then changing the kinetics of the underlying ion conductances to shorten the delay in excitation would provide the same kind of change in dynamics, because the postsynaptic cell will recover to spike threshold sooner and thus respond preferentially to shortened delays. From the computational/dynamical class perspective these two forms of neuromodulation are in principle identical, despite the widely different biochemical/biophysical basis. As the mechanisms for the delay in the bat circuit are not understood, a modulatory mechanism like that envisioned here remains hypothetical.

This form of neuromodulation has an interesting correlate in artificial electronic neural networks. In an artificial electronic cochlea,⁴ microelectronic delay lines consist of a series of RC (resistance-

capacitance) filter sections, each producing a short delay. The circuits are built with a global control line that can be used to change the effective values of these time constants in all sections simultaneously. Thus the entire circuit can be modulated to work over longer or shorter time scales in real time.

Computation and circuit details in more complex artificial neural networks

The circuits for locating sound source from interaural time disparity in the owl or for detecting temporal sequences of sensory activation in the bat seem simple in terms of the computations they perform, because we can grasp intuitively the very simple logical combination of algebraic equations that describe the computation. In the study of artificial networks of neuron-like units, dynamic systems performing more complex computational tasks have been constructed, including those described by energy functions and those organized by learning algorithms. As with the systems represented by simple algebraic equations, both map the computation being performed directly onto the organization of the static and dynamic functions that the network implements. Once again, the specific geometry of the synaptic connection pattern matters and the forms of dynamic function (delays versus no delay in a particular connection, for example) can completely change the network's computation: these details are the essence of the way the circuits manipulate information.

Energy minimization

Just as an algebraic equation can express a computational task, an 'energy' function^{10,11} provides an explicit description of the architecture of a network that will solve the problem represented by the minimization of the function. The problems that have been cast in the energy minimization format are quite diverse; those that seem of biological interest include associative memory (i.e., content-addressable memory);¹⁰ problems in visual perception such as determining the shape of an object from its shading or the determination of object boundaries (edges) and surfaces from sparse distance information;¹¹ and detection of a velocity gradient in an image (chapter 14 in ref 4).

The relationship between simple algebraic equations and logical statements and a model circuit that implements the equations and statements was relatively straightforward in the examples described above. It was necessary only to choose an implementation scheme that had the right combination of static and dynamic function building blocks arranged so that the flow of information was in the form specified by the equation. With artificial neural networks designed to minimize energy functions, it is more difficult to see this flow of information, primarily because of its complexity. As in the simpler equations, however, the organization of static and dynamic functions follows from the form of the

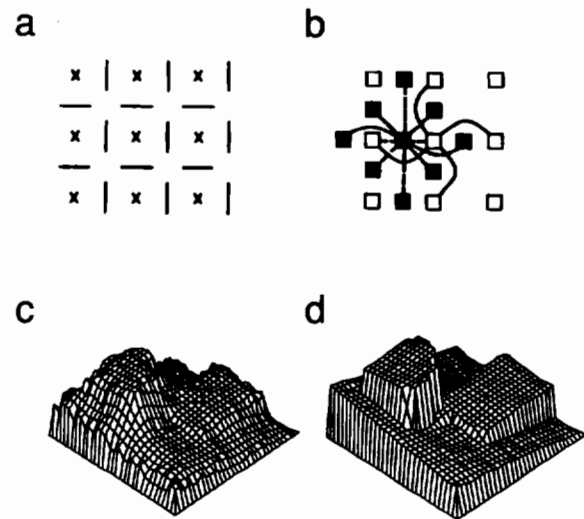


Figure 4. An artificial model neural network for a surface reconstruction problem in vision. The architecture was designed¹² by reference to an energy function.¹¹ As schematically illustrated in (a), the network consists of a lattice of two kinds of units (artificial neurons), one representing the height of a representative section of the surface (depth units, indicated by crosses) and the other representing the presence of a sharp edge or discontinuity in height (indicated by lines). In (b) the lattice of these two kinds of units are represented as filled squares for the line units and unfilled squares for the depth units with the pattern of connections (artificial synapses between units) indicated by the lines joining the boxes. The specific geometrical form and numerical values of these connections are determined by reference to the energy function that defines the circuit, which also predicts how the circuit would behave when presented with noisy information representing the surface of several objects, such as those shown in (c). The network computes that the noisy data represents the smooth surfaces and sharp edges (boundaries between objects) shown in (d). The specific geometrical form of the connections shown in (b) is essential to the computation performed: it is, in a sense, the program of the computation. (Taken from ref 12).

equation, and as the network simply evolves with time, it solves the minimization problem. Here too, the circuit related to the energy minimization can be implemented by several different physical mechanisms, each producing the appropriate circuit dynamics. The map of implementation onto dynamics is not one-to-one and there is a wide latitude in the kind of biological or artificial circuit that could perform the computation.

The structure of a model network constructed from an energy function description of the visual task of surface reconstruction¹¹ is shown in Figure 4a and b. It was designed to determine from the noisy information representing the surfaces of several objects (for example, Figure 4c), the smooth surfaces of the objects and the edges or sharp breaks in smoothness between objects (Figure 4d). The pattern of connections between neurons in this network is

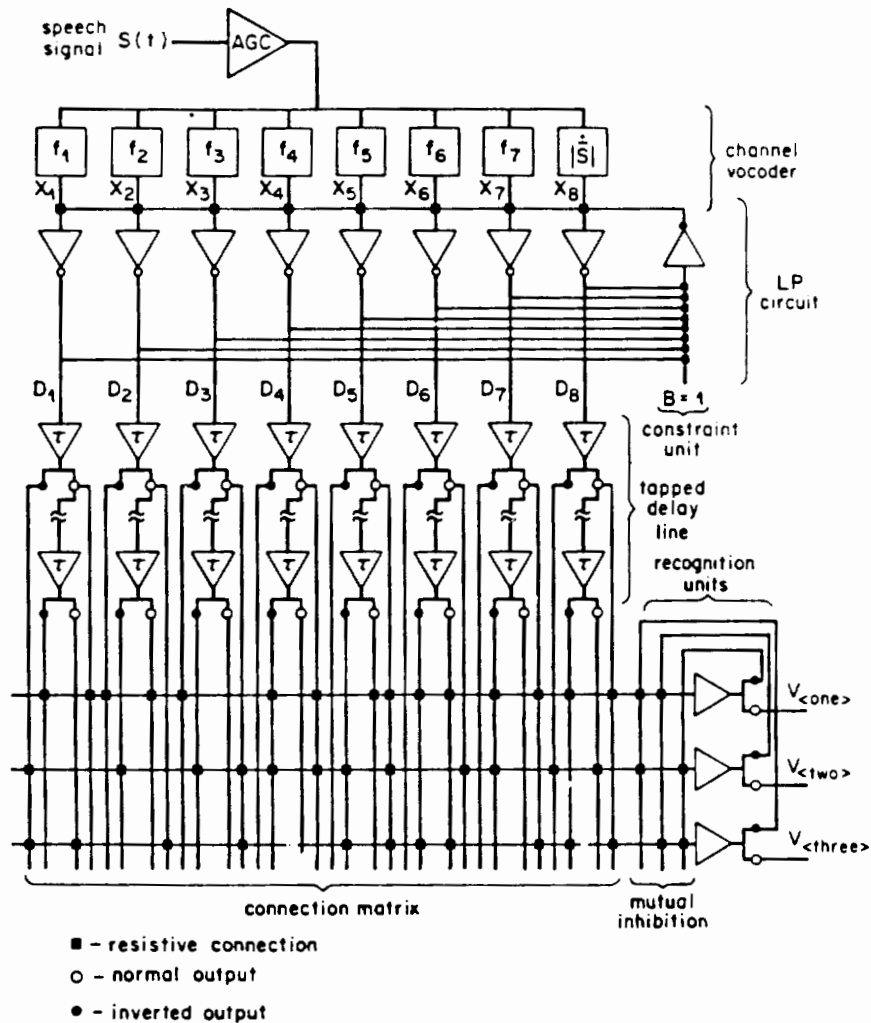


Figure 5. An artificial neural network designed to recognize the spoken digits 0-9, a relatively simple problem in speech recognition. Like the FM-sensitive neurons found in echo-locating bats, detectors for simple features in the acoustic waveform, (resonances or formants in the speech signal, are connected through a set of time-delayed connections to output recognition units whose outputs versus time represent the detection of spoken words in the input sound stream. The connection matrix determines the geometric pattern of time delayed connections. The numerical values of the connections are determined by a learning algorithm. The dynamics of the constituent elements of the circuit (time-delayed and non-time delayed connections) and the specific numerical values of the connections are essential to the performance of the circuit.

far from random: there is a specific description of the connection pattern embodied in the energy function and a circuit with a different pattern of connections would compute something else. Examining the circuit in Figure 4 from the perspective of the details that matter, the conclusions are similar to those for the simpler biological networks. The details of the connection pattern and the details of the specific combinations of static and dynamic functions are critical but, again, there is latitude in the physical mechanisms that could implement the functions and the flow of information through the connection pattern.

The nervous system does not know about the logical combinations of simple equations that we conjured up to describe the computation of sound localization or prey detection; they are mathematical concepts that we use to understand these systems. The energy function is another such mathematical concept, invented because the notion that dynamics can be related to a function being minimized (and thus cast as a variational problem) has a long history as a useful theoretical idea in many areas of physics, including the Hamiltonian formulation of classical dynamics. It seems likely that in the future many other mathematical tools will be brought to bear on the relationship of the dynamics of neural circuits to their function.

The notion of using an energy function or variational approach to describe a neural circuit is useful only because it specifies a relationship between the dynamics of the circuit and the manipulation of information; by reference to the energy function we understand and can *predict* the kind of manipulation of information the actual circuit will perform by virtue of its dynamics. In the same way other theoretical tools that describe dynamics will be useful in understanding the computational aspects of the nervous system only if similar dynamic-description to information-processing relationships can be ascertained. For example, if nervous systems display chaotic dynamics,¹³ the understanding of the nervous system may potentially benefit from the use of the new tools of dynamical systems theory.^{14,15} But this will only be true if a relationship between chaotic dynamics and biologically-relevant computation can be developed. (At present only the vague notion that chaotic dynamics can implement certain search strategies has been proposed but no model demonstrating the detailed method or effectiveness of this idea has been developed.)

Learning algorithms

Another class of tools for circuit design that has been under intensive investigation in artificial neural networks is learning algorithms.^{16,17} Like the energy function, learning algorithms can relate the dynamics of a circuit to the manipulation of information. In the model neural networks in which learning algorithms are employed, a specific form of dynamics of the units and synapses that comprise the network is assumed. The strengths of connections between units are iteratively changed by testing sequentially how the network performs a prespecified task and using an update rule that is a mathematical description of how to modify slightly the present set of connections so that the circuit will perform slightly better. The procedure is repeated until a desired level of performance is reached or failure is apparent.

An example of a circuit with connection strengths determined by a learning algorithm is shown in Figure 5. The circuit is being investigated for artificial speech recognition,¹⁸ and has a basic dynamic structure similar to the FM/FM sensitive neurons used by the bat for prey detection. Instead of recognizing FM sweeps separated by delays, some of its constituent artificial neurons recognize more complex acoustic trajectories of formants in human speech. But the dynamic structure of the circuit is basically the same as that of the bat circuit: the acoustic signal is analysed by a set of units whose levels of activity reflect the amount of acoustic signal in a set of frequency bands as a function of time. The outputs of these units project through a set of time-delayed synaptic connections to higher-order units that recognize individual words in the vocabulary by the circuit: there is one recognition unit for each word. But the pattern of time-delayed connections is not prespecified. A learning algorithm is applied to the basic circuit and the values of the time-delayed connections that produce the highest accuracy of speech recognition are iteratively learned. When the learning algorithm is finished, the resulting network is an artificial speech recognition system based upon neural network computing principles.

What details of this circuit are important? The information that determines which specific words the circuit will recognize has a direct correlate in the pattern of time-delayed synaptic connections. Just like the orderly set of time-delayed connections in the FM-sensitive neurons of the bat, the time-delayed connection set for speech recognition is orderly, although more complicated to describe. The pattern

of time delays is the embodiment of the computation; altering it changes the acoustic signals that are recognized as words. This does not mean that if one synapse is changed or eliminated, the network will completely fail as there is still information in the rest of the connections. But as more and more connections are changed, the performance of the network will eventually decrease. What if the delay connections have the same spatial patterns of connection strengths but the delays are eliminated? Again the circuit would not function. The computation requires the combination of information that is occurring at different times in the acoustic signal; eliminating the delays eliminates the proper combination of information. These are the details that matter: the dynamic functions performed by the delay connections, the summation occurring in the recognition units and the specific pattern of connections. The mechanisms to implement them can be quite varied. A completely artificial network has been constructed out of discrete electronic components to demonstrate the speech recognition principles of the circuit and the learning algorithm.¹⁸

Computation and circuit building blocks in central pattern generators

The examples described so far are all associated with sensory perception. Among the most completely characterized biological neural circuits are central pattern generators (CPGs) that produce the coordinated firing of motor neurons necessary for a specific movement. How can we approach the understanding of these systems in light of the relationships expressed above between computation, dynamic function, and implementation? This is a difficult question that strikes at the heart of what is meant by a computation. In the previous examples the organization of the static and dynamic functions that define the circuit is directly related to the computation it performs, that is, how *information* is transformed or used by the system. To apply these kinds of ideas to CPG circuits, the notion of information must be incorporated, which is not the way that CPG circuits have traditionally been analysed.

There are at least two ways that the notion of information can be brought to bear on the function of a CPG circuit, first through the logic underlying the pattern of activity produced by the circuit in

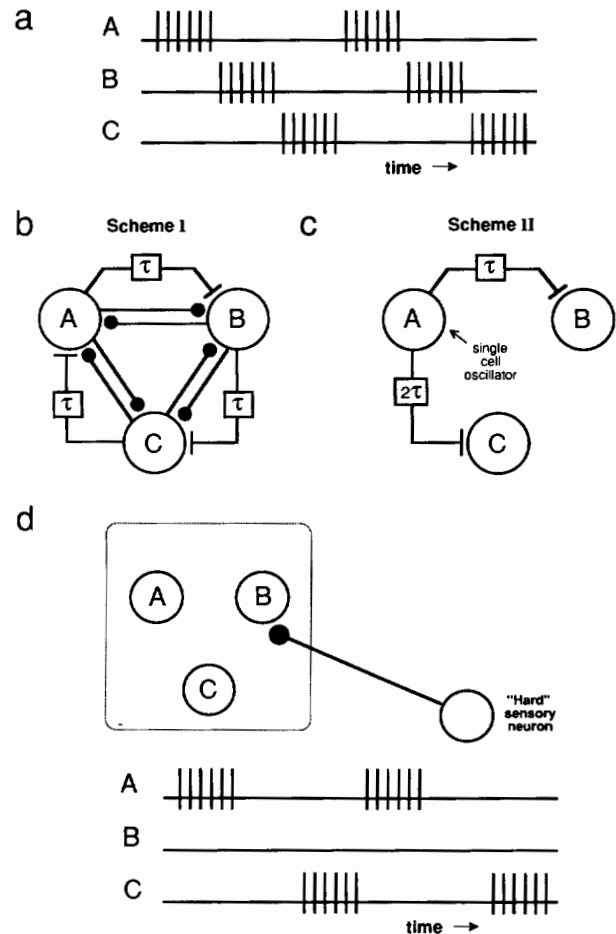


Figure 6. Central pattern generator (CPG) circuits that can produce the same basic patterns with very different patterns of connections. (a) A hypothetical pattern of activity produced by three units (A, B, and C) that comprise a model CPG. (b) A pattern of connections that will produce the activity pattern shown in (a); each cell is an analog nonlinear unit. (c) A different scheme that could also produce the activity pattern in (a). This circuit relies on the intrinsic oscillator activity of unit A. (d) With a sensory neuron that detects when food is hard and that strongly inhibits unit B, scheme II can easily switch between ingestion and regurgitation. The outputs under these conditions are shown below. Scheme I would fail with this feedback arrangement (see text).

isolation, and second by incorporating information from the interaction of the CPG with the rest of the CNS, with the musculature it drives and with associated sensory feedback during different environmental situations. (An example is a circuit controlling feeding where chewing soft food results in ingestion while chewing hard food results in regurgitation.) The logic underlying the pattern of activity of the circuit in isolation cannot be determined from the

pattern alone: many logical computations can produce the same patterned activity (see below). The logic can be ascertained only from the organization of the static and dynamic functions that the circuit embodies and the reason why a particular CPG logic is used in a given circuit may be understood only with reference to the interaction between the circuit and other components of the system, including sensory feedback and neuromodulation.

Some very simple hypothetical CPG model circuits (Figure 6) illustrate these points. There are three neurons (A, B, and C) in the CPG circuit and the circuit goes through a repetitive sequence of three states: (S1) A active, B, C inactive; (S2) B active, A, C inactive; (S3) C active, A, B inactive (Figure 6a). But there are several ways in which the circuit could be organized to produce this pattern, for example, schemes I and II (Figure 6b,c). These have completely different computational logics yet their outputs are identical. They illustrate the concept that there is not a 1 : 1 relationship between basic pattern of activity and the computational logic of the circuit.

The logic of scheme I (Fig. 6b) is a set of conditional (if-then) statements between different states of the network: if the circuit is in state (S1) make state (S2) happen; if in state (S2) make (S3) happen; if in state (S3) make (S1) happen. A neural circuit that will perform this computation and thus produce the desired motor pattern can be constructed from any set of physical mechanisms that embodies dynamics equivalent to this logic. One possible implementation is that suggested by temporal associative network models.¹⁹⁻²¹ The basic idea of these models is that the states of activity (S1), (S2) and (S3) are produced by a specific form of reciprocal connections, without delays, between neurons A, B, and C; the logical transitions, for example, if (S1) then (S2), are produced by the antecedent state, in this case (S1) producing the next state (S2) through a specified pattern of time-delayed connections. The example in Figure 6b shows the synaptic connections between simple model neurons (that is, non-linear threshold elements) that would exhibit this dynamics.

If a real CPG has the pattern of activity states (S1)→(S2)→(S3)→(S1) . . . , but does not have this pattern of synaptic connectivity or single-unit dynamics, there are several possible reasons. First, the specific synaptic form and single-unit dynamics of these models are not unique and other possibilities may produce equivalent circuit dynamics. Second, the circuit, although exhibiting the pattern of activity of Figure 6a, may not be computing the logical

function of scheme I, but something else, for example scheme II (Figure 6c): if neuron A is active, now with delay τ make neuron B active, with delay 2τ make neuron C active, with delay 3τ reactivate neuron A. A circuit implementing scheme II performs a fundamentally different computation because the logic equations that embody it are quite different. The basic pattern of activity of neurons A, B and C will, however, be identical to that produced by scheme I, it will be Figure 6a again. The combination of static and dynamic functions used to implement scheme II is fundamentally different from that of scheme I because the computation is different. For example, neuron A could be a single cell oscillator with period 3τ with a set of time-delayed connections to neurons B and C: A excites B with a time delay τ , and A excites C with longer delay 2τ (Figure 6c).

It is no surprise that the central pattern generators in different animals employ so many different physical circuits²² that seem to produce roughly similar patterns of activity in constituent motor neurons. There are at least two reasons for this: first, the distinction between implementation and computation (different physical mechanisms can implement identical static and dynamic functions); and second, the basic patterns of activity in neurons cannot normally be used to distinguish between fundamentally different computations that must necessarily imply differences in implementation.

The reason that the simple CPG circuit might employ the logic of one scheme instead of the other may become apparent only when operating conditions change or when the interaction of the CPG with its associated musculature and sensory feedback systems is considered. At this point the logic embodied in the two schemes is incorporated into a larger set of computational principles, not all of which may be described by simple if-then statements. Consider the model CPG to control swallowing with a simplified set of musculature and sensory feedback: soft food is ingested with a sequence of contractions in three muscles, activated in orderly sequence by neurons A, B and C. If the food is hard, the sequence of contractions is altered to produce regurgitation, which occurs in the sequence A followed by C, with no intervening contraction of B. To perform this with scheme II is easy: if hard food is sensed after the contraction produced by activity in neuron A, neuron B is directly inhibited by sensory input with sufficient synaptic strength to stop its delayed activation by A. Because the activation of C does

not require the logic of going through the state S2 (B active, A, C inactive), neuron C still becomes active at the appropriate time and the food is regurgitated. This simple feedback scheme will not work with scheme I, because if state (S2) is blocked then state (S3) is never reached. Thus the logic of scheme II becomes apparent when the observed sensory feedback is taken into account.

Now let us imagine a completely different situation: if the food is really good, swallow faster. In scheme I, one global mechanism of neuro-modulation characterized by 'make all delayed synaptic connections occur on a shorter time scale' will make the circuit cycle faster, although still maintaining the necessary pattern of coordination. This is ensured by the logic of the computation. But scheme II requires two modulation signals: speeding up the intrinsic oscillation time of neuron A, the single-cell oscillator, by a modulatory mechanism makes A oscillate faster, but now B and C may be out of synchronization; for example, C might fire after A fires for a second time. Thus it is essential to take into account neuromodulation and sensory feedback under different environmental conditions and their effects on circuit dynamics in order to understand the logical organization of a CPG circuit.

Conclusions

The building blocks of neural circuit *construction* are biophysical and biochemical mechanisms found in neurons. The building blocks of neural circuit *computation* are static and dynamic functions and their interrelationships. The computation performed by a given circuit defines a requisite set of these functions. Various physical mechanisms can be used in principle to implement a particular function and a circuit need use only one of them to provide that function. Although the goal is to understand how these functions are organized to produce useful computation in the nervous system, the computational principles themselves are probably not in any way unique to nervous systems. Man-made artificial devices like the analog microelectronic silicon retina and cochlea⁴ already share computational principles with the nervous system and similar schemes may be found in other forms in the natural world, for example in immunology or computational aspects of insect society.

Selverston²³ has noted that 'much of the fine grain, microscopic data are not necessary for

analyzing the actual output [of the circuit] and must somehow be compacted in lumped categories'. I suggest that static and dynamic functions, the building blocks of computation, are the appropriate categories into which to lump physical details of neural circuit construction. The example in Figure 3 showing various physical mechanisms that could implement delayed excitation illustrates this point. General goals for the neurophysiologist interested in the computational significance and function of the nervous system should be to interpret details of construction in terms of the static and dynamic functions they implement, and to attempt to understand these functions in the context of circuit computation. The tools used to perform the latter will be those that provide a direct map between dynamics and computation. I have given examples of how simple equations can be related to simple computations, how their manifestation by static and dynamic functions can be implemented by neurons, and how variational principles of dynamics (the energy function, for example) can, in certain cases, be related to computation in artificial neural network models. The dynamics of a characterized computation can also be produced by learning algorithms. It seems likely that other methods to relate dynamics to computation will be developed in the future.

What details matter? As I have pointed out here, from a computational perspective details that determine or modify the static and dynamic functions of the circuit are fundamentally important; this is not something unique to the nervous system, but is true of any system that manipulates information. We must, however, be careful not to mistake the determination of dynamic characteristics for uniqueness or necessity of the implementation.

Acknowledgements

I thank John Hopfield, Alan Gelperin, and Alan Kay for many helpful discussions.

References

1. Sullivan W, Konishi M (1986) Neural map of interaural phase difference in the owl's brainstem. *Proc Natl Acad Sci USA* 86:8400-8404
2. Jeffress L (1948) A place theory of sound localization. *J Comp Physiol Psych* 41:35-39
3. Carr CE, Konishi M (1988) Axonal delay lines for time measurement in the owl's brainstem. *Proc Natl Acad Sci USA* 85:8311-8315

4. Mead C (1989) *Analog VLSI and Neural Systems*. Addison-Wesley, New York
5. Suga N (1984) The extent to which biosonar information is represented in the bat auditory cortex, in *Dynamic Aspects of Neocortical Function* (Edelman, GM, Gall WE, Cowan WM, eds), pp 315-374. Wiley, New York
6. Kehoe JS, Marty A (1980) Certain slow synaptic responses —their properties and possible underlying mechanisms. *Ann Rev Biophys Bioeng* 9:437-465
7. Llinas R, Yarom Y (1981) Properties and distribution of ionic conductances generating electroresponsiveness of mammalian inferior olivary neurones in vitro. *J Physiol* 315:569-584
8. Byrne JH (1980) Quantitative aspects of ionic conductance mechanisms contributing to firing pattern of motor cells mediating inking behavior in *Aplysia californica*. *J Neurophysiol* 43:651-668
9. Byrne JH (1982) Cellular and biophysical mechanisms contributing to regulation of reflex excitability of inking behaviour in *Aplysia*. *Fed Proc* 41:2147-2152
10. Hopfield JJ, Tank DW (1986) Computing with neural circuits: a model. *Science* 233:625-633
11. Hopfield JJ, Tank DW (1985) "Neural" computation of decisions in optimization problems. *Biolog Cybernet* 52:141-152
12. Koch C, Marroquin J, Yuille A (1986) Analog "neuronal" networks in early vision. *Proc Natl Acad Sci USA* 83:4263-4267
13. Rapp PE, Zimmerman ID, Albano AM, deGuzman GC, Greenbaun NN, Bashore TR (1987) Experimental studies of chaotic neural behavior: cellular activity and electroencephalographic signals, in *Nonlinear Oscillations in Biology and Chemistry* (Othmer HG, ed.), pp 175-205. Springer, Berlin
14. Cvitanovic P, ed. (1984) *Universality in Chaos*. Adam Hilger, Bristol
15. Thompson JMT, Stewart HB (1986) *Nonlinear Dynamics and Chaos*. John Wiley, New York
16. Rumelhart DE, McClelland JL, eds (1986) *Parallel Distributed Processing: Exploration into the Microstructure of Cognition*, Vols 1,2. MIT Press, Cambridge, MA
17. Touretsky DS, ed (1988) *Advances in Neural Information Processing Systems 1*. Morgan Kaufmann, San Mateo, CA
18. Tank DW, Hopfield JJ (1987) Concentrating information in time: analog neural networks with applications to speech recognition problems. *Proc IEEE Conf Neural Networks*, San Diego, CA
19. Fukushima K (1973) Model of associative memory in brain. *Kybernetik* 12:58-63
20. Hopfield JJ (1982) Neural networks and physical systems with emergent collective computational abilities. *Proc Natl Acad Sci USA* 79:2554-2558
21. Kleinfeld D, Sompolinsky H (1988) Associative neural network model for the generation of temporal patterns: theory and application to central pattern generators. *Biophys J* 54:1039-1051
22. Getting PA (1988) Comparative analysis of invertebrate central pattern generators, in *Neural Control of Rhythmic Movements in Vertebrates* (Cohen AH, Rossignol S, Grillner S, eds), pp 101-127. Wiley, New York
23. Selverston AI (1988) A consideration of invertebrate central pattern generators as computational data bases. *Neural Networks* 1:109-117