from *SIAM News*, Volume 31, Number 5

Cooperativity and Parallelism
In Mathematical Models of Brain Function

By Martin Beckerman

Cooperative processes, or algorithms, are procedures that generate large-scale, or global, effects through sequences of small-scale, or local, operations. Such procedures underlie life as we know it. They are encountered in a variety of physical and chemical systems, and are found in biological systems at all levels of organization. In a cooperative system, global computations are performed by many units operating in parallel. Cooperativity emerges in these systems through interactions among component units. Each unit generates a new state from its previous one by combining information available internally with signals received from neighboring units and from outside.

This article considers cooperativity and parallelism in two classes of mathematical models of brain function developed by researchers in the field over the past 10 to 15 years. The human brain, containing approximately 10^{11} neurons and perhaps as many as 10^{15} distinct synaptic connections, is the archetype of a highly cooperative, massively parallel system. The first class of models and accompanying solution methods were constructed to carry out tasks in perceptual inferencing, such as the segmentation of a scene into its component surfaces, figure–ground segregation, and object recognition. In these computer-based approaches we construct a Markov random field (MRF) to capture correlations in a visual image and then use simulated annealing, or one of its siblings, to generate sequences of small changes. The solution methods are stochastic in character; the computations are intrinsically parallel and are inspired by our thinking about how the early stages of visual processing might occur in the brain.

The second class of mathematical models and methods has as its central theme the notion of assembly coding. These models and methods were developed to carry out the same tasks as those in the first class, but we now construct families of coupled first-order ordinary differential equations that describe circuits of interconnected neurons. We then evolve these nonlinear dynamical systems through sequences of states that converge to fixed points, limit cycles, or chaotic attractors. The solution methods can be either deterministic or stochastic; the computations are parallel and highly distributed, and they model circuits found in the mammalian brain.

Perceptual Inferencing with MRFs

In the mammalian brain, most visual signals are relayed from the retina through the lateral geniculate nucleus to the primary visual cortices. From there information is parceled out to several distinct regions, each specializing in a particular modality, such as color, motion, or depth. As is the case elsewhere in the brain, the supporting architecture is complex—cells both within and across regions are linked to one another through networks of forward and reciprocal connections. These observations support the view that the early processing of visual information takes place in parallel. There are many open questions on how this is done, and especially on just how the information is reintegrated to form a percept.

In the MRF-based computational models, we design a process, or algorithm, that takes as its input a two-dimensional projection of a three-dimensional visual scene and produces as its output an image description in terms of surfaces, their attendant discontinuities, and their stable, physical properties. To reconstruct an image we remove noise, blur, and any other distortions and artifacts from the input image while preserving surface boundaries. In segmenting an image we both reconstruct and identify the individual surfaces in the scene. In integration, we have the additional task of combining two-dimensional image data from a number of perceptual domains into an internally self-consistent representation. In many instances we attach labels to each surface that describe their physical characteristics. We execute the integration tasks concurrently with the reconstruction and segmentation tasks, from the beginning, using coupled Markov random fields.

A Markov random field model is a random field with the special property that all joint probability distributions have associated conditional probabilities that are local in character. For an image modeled as an MRF, this means that the probability of a pixel's assumimg a particular gray-value is dependent on the gray-values of neighboring pixels, but not on those of pixels at more remote sites. The utility of Markov random fields for describing the short-range correlations in visual images is derived from the Hammersley–Clifford theorem, which tells us that the global character of a Gibbs random field possessing potentials of a particular form is equivalent to the purely local character of a Markov random field. We can therefore use easy-to-implement Gibbs distributions and their accompanying potentials in place of difficult-to-implement conditional probabilities.

The mathematical details are straightforward. A Gibbs distribution, or Gibbs random field, is a random field if and only if its joint probability distribution is of the form

$$p(x)=\frac{1}{Z}\exp\left(-\frac{1}{T}U(x)\right).$$

Here Z is a normalization constant, T is a control parameter called the temperature, the vector x denotes an array of pixel values, and $U(x)$ is the global potential, sometimes called a Hamiltonian, energy, or a cost function. The potentials of interest to us are those that can be decomposed into a sum of potentials, one for each lattice site. That is, $U(x)$ assumes the form

$$U(x) = \sum_{i \in \Lambda} V_i(x_i)$$

with site potentials, V_i, that can be written as sums of clique potentials, V_c:

$$V_i(x_i) = \sum_{c \in C_i} V_c(x_c).$$

If the potentials can be so expressed, then the Hammersley–Clifford theorem is applicable.

The clique potentials, which describe interactions among pixel elements in the neighborhood of the *i*th pixel element, contain terms of two classes: (1) single-element terms involving the potential at the site in question and (2) terms describing cooperative interactions between the site element and its nearest neighbors of various orders. The decomposition of the total potential into a sum of site terms allows us to express the joint probability as a product of quasi-independent site potentials, and thus to devise any number of sequential or parallel algorithms for updating pixel values.

In the framework of image reconstruction, the role of the single-site potentials is to guarantee that changes in pixel values during the updating are small at each step. The primary function of the cooperative interactions is to express constraints that, from the many possible solutions, allow only those that are physically meaningful. The problem being addressed can be stated as follows:

The light convergent onto the retina, or, equivalently, the intensity encoded as a pixel value at each location in a digital image, is influenced by many environmental factors. These factors include the character and geometry of the illumination source(s), the orientation and reflectance properties of the surfaces, and the optics of the imaging system. One solution to this inverse, ill-posed problem of recovering properties of three-dimensional surfaces from two-dimensional images is to use prior knowledge about the physical properties of surfaces, i.e., how they are spatially organized. For instance, away from boundaries, neighboring portions of a surface tend to have the same physical characteristics. We model our prior knowledge as constraints that loosely enforce such characteristics as piecewise smoothness, linearity, and spatiotemporal continuity.

In the MRF approach we introduce two coupled spatial lattices for each modality—one describing the intrinsic image and the other representing the associated discontinuities. This dual modeling allows us not only to treat surface and boundary aspects simultaneously, but also to suspend the smoothness requirements at the boundaries in a natural way. We couple image modules to each other so that information is shared. Data present in one modality aid in the integration of the information in the others. Our emerging coupled random field picture, in summary, is characterized by a three-dimensional parallelism, a cooperative low-level integration of multiple sources of information through simple coupling interactions, and an encoding of general information about the physical world through constraints.

Once we have constructed an appropriate Hamiltonian for the problem under consideration, we use simulated annealing (SA) to carry out a pixel-by-pixel iterative reconstruction of the image. The SA algorithm endows the system with an equilibrium dynamics that avoids trapping in local minima while evolving the system to a near-optimal low-energy state. More specifically, the sequence of states generated by a slow lowering of the temperature control parameter forms a nonhomogeneous Markov chain that converges in distribution to a near-optimal minimum in the cost function. The algorithm will converge if we lower the temperature sufficiently slowly, which can be accomplished by following a logarithmic cooling schedule. The solutions so found represent maximum a posteriori estimates of the true image given the corrupted data.

A number of approximations and alternatives to the SA algorithm have more rapid convergence properties. Two of them—iterated conditional modes (ICM) of the posterior distribution and maximizer of the posterior marginals (MPM)—have been widely used for image reconstruction. ICM corresponds to taking the $T = 0$, or quenched, limit of the annealing algorithm. MPM corresponds to taking a fixed temperature limit. Another class of methods includes the mean-field approaches, where we approximate the posterior Hamiltonian with a mean-field Hamiltonian. The mean-field approaches employ convex approximation and a deterministic dynamics, such as gradient descent, to find useful minima.

Still another approach is to employ a stochastic differential equation of the general form

$$dy(t) = -\nabla U(y) + \sqrt{2T(t)}\,dW,$$

where y is used in place of x to denote our image, U is again the potential, and W is a Wiener, or Brownian motion, process. The gradient of the potential appearing in the first term controls the drift of the system toward equilibrium. This gradient term replaces the friction term in the conventional Langevin equation. In the second term we introduce a time-dependent temperature, $T(t)$, that controls the magnitude of the random fluctuations, or stochastic noise, in the system. The result is a Langevin or Brownian dynamics that can be used for a global optimization in which large and random fluctuations dominate the short-term behavior of the system and permit escape from local minima or metastable states, while a slow drift controls the long-term behavior and pushes the system into a near-optimal low-energy state.

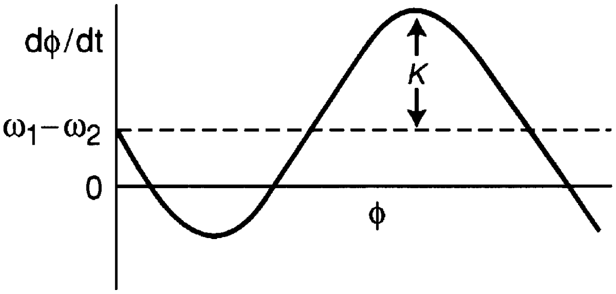

Figure 1. *In rotator models, sychronization occurs whenever the absolute value of the phase difference is less than or equal to the coupling constant (K).*

Assembly Coding

The processing of early visual information, such as contours, depth, motion, and color, is distributed, in parallel, among large populations of neurons in different cortical areas such that each population encodes a particular aspect of the visual scene. As mentioned earlier, the way in which these attributes are integrated to produce an unambiguous representation of the component features and objects in the scene is an open question. The issue of how attributes are integrated to produce a segmentation of the scene into its component surfaces, and a segregation of objects from their backgrounds, is known as the “binding problem.”

One possible approach to this problem, which is strikingly similar in philosophy to the MRF approach, is through assembly coding. According to the assembly-coding hypothesis, neurons that encode attributes that belong together are bound together by their synchronous firing. When occurring in cells located in different columns within a given cortical area, synchronized firing would encode local information; when occurring in cells distributed across different cortical areas, it would bind features. A neural assembly is a collection of cells that cooperate with one another. In assembly coding, temporal clusters form and reform in response to input signals. If there are multiple objects in a scene, then there are several co-active clusters. Individual neurons belong at various times to different assemblies and can rapidly associate into a functional group while at the same time disassociating from a different functional group. In distributed coding of this type, it is temporal coherence rather than spatial proximity that signifies membership in a cluster.

During the past decade a large number of experiments have provided us with evidence that clusters of synchronously discharging cells form within one or more columns located in a particular cortical area, and within columns located in different cortical areas and across hemispheres. These results tell us that at least some of the conditions for assembly coding are satisfied. The experimental studies have been accompanied by mathematical modeling studies demonstrating that cluster formation is an intrinsic property of networks of model neurons.

A number of cellular, network, and combined cellular–network mechanisms can generate collective behavior in neurons. Some neurons are endowed with intrinsic electrophysiological properties that support rhythmic firing. These cells serve as pacemakers. Others do not, but they still have intrinsic properties that promote rhythmicity when driven by pacemakers. Cells that are not intrinsically oscillatory can become so as a consequence of network properties. Synchronized oscillations can arise in populations of excitatory and inhibitory neurons reciprocally coupled to one another.

A system as small as two cells can exhibit cooperative behavior. Mutual synchronization, or mutual entrainment, is an example of cooperativity in which temporal order emerges in populations of oscillatory units coupled to one another. When isolated, each oscillator moves according to its native frequency. When coupled, these oscillatory units are able to synchronize their motions, provided that the native frequencies are not too dissimilar. Systems exhibiting mutual entrainment include lasers and large populations of chirping crickets and flashing fireflies. Among other examples of mutual entrainment are the synchronized rhythmic firing of neurons in central pattern generators, the synchronized firing of pacemaker neurons in the heart, oscillations and waves in the mammalian small intestine, and circadian rhythms.

One broad class of mathematical models of mutual entrainment consists of the rotator models. In rotator models of mutual entrainment, we are able to probe what will happen in a large population of similar, limit-cycle oscillators, each with its own native frequency, when these elements are weakly coupled through interactions that communicate phase information from one member of the population to the others. The simplest, two-element system is described by the rotator equations

$$\frac{d\theta_1}{dt} = \omega_1 + (K/2)\sin(\theta_2 - \theta_1)$$
$$\frac{d\theta_2}{dt} = \omega_2 + (K/2)\sin(\theta_1 - \theta_2).$$

In this system, K is the coupling constant, and ω_1, and ω_2 are the native frequencies of the two rotators. To see how entrainment occurs even in this simple system, suppose that the phase difference $\theta_1 - \theta_2$ is positive; that is, the phase of the first oscillator is ahead of that of the second. The coupling term in the first equation will then be negative, so that $d\theta_1/dt < \omega_1$, and the first oscillator will slow down. Similarly, if the phase of the first oscillator is behind that of the second, the first oscillator will speed up. The same considerations hold for the second oscillator.

This mechanism for achieving synchrony, called “phase pulling,” has the property that the force exerted by one oscillator on the other is the same at all points along the limit-cycle trajectory, and the amount of attraction depends on the coupling strength and phase difference. Entrainment will occur whenever the absolute value of the phase difference is less than or equal to the coupling constant (see Figure 1). Neurons in the central nervous system exhibit a number of prominent forms of rhythmic and synchronous activity. Some of these forms are of low frequency and large amplitude, implying that large populations of cells are participating. Others, of high frequency and low amplitude, involve smaller populations of neurons. Multiple steady states, bifurcations, hysteresis, global oscillations, local clustering, and propagating waves of synchrony are characteristic properties of cooperative behavior in neural systems.

Electroencephalographic (EEG) recordings of oscillatory wave patterns in the brain are a well-established tool for identifying abnormal conditions, as well as various stages of sleep and arousal. In EEG recordings we observe delta waves in the range of 0.1–4 Hz. These oscillations are associated with deep sleep and abnormal function. We find theta-wave oscillations (6–7 Hz) in components of the limbic system, such as the hippocampus; we encounter alpha waves, rhythmic neural activity, in the 10-Hz range, during drowsiness and relaxed activity; and spindle activity, associated with the early stages of quiescent sleep, is observed in the range of 7–14 Hz. Repetitive firing at high frequencies appears in EEG recordings as irregular fluctuations in the beta band, from 15 to 30 Hz, and in the gamma band, from 30 to 60 Hz. These correlated activity patterns correspond to states of arousal and attentiveness.

Using simultaneous multielectrode recording devices, researchers have recently developed techniques that provide us with a far more detailed picture of cortical activity patterns at high frequencies. Repetitive firing patterns in the 40-Hz range were first observed in the 1950s in studies of the olfactory bulb. Synchronized, repetitive firing activity in the visual cortex of the cat was discovered more recently, in 1989, by researchers using arrays of microelectrodes. Still another form of high-frequency synchrony, hippocampal sharp waves spanning a broad range of frequencies in the vicinity of 200 Hz, have been observed in the hippocampus during the past five years. These activity patterns are associated with states of consciousness and immobility.

Rhythmic activity in the gamma band is associated with binding phenomena and assembly coding. One possible mechanism for this form of rhythmicity is that of reciprocal reentrant connections between groups of excitatory and inhibitory neurons distributed across feature-encoding areas. In reentrant signaling, we find that cells belonging to a neuronal group exchange signals in parallel, repeatedly, and in register with cells in other groups. The reentrant signals drive the integration process, promoting the elimination of conflicting and ambiguous responses within a single area, permitting outputs in one area to be used in another area, and making possible the reuse of outputs back into the originating feature-encoding region. Assembly coding and MRF-based integration-by-labeling are self-organizing processes that reinforce and improve the integration of features from one iteration to the next and that are robust against noise. Both use multiple cues in an iterative manner to resolve conflicting or ambiguous situations arising from single cues, and both maintain a distributed representation of the stimulus.

As mentioned earlier, cells located in different cortical regions and hemispheres are able to synchronize their discharges. Several physical mechanisms can promote coherence over long distances. In one set of recent experimental and theoretical studies, we have found that networks of inhibitory interneurons are able to generate 40-Hz oscillations through mutual inhibition. When linked to a network of excitatory pyramidal cells, the cells in the inhibitory network fire spike doublets, with a time interval that serves to entrain the excitatory pyramidal cells. In this picture, the doublets provide a local clock for the system. More generally, many local clocks, operating on a number of frequencies and serving a variety of functions, have been observed throughout the brain.

Phase transition phenomena are associated with the emergence of synchronous activity in the central nervous system. Dynamic processes occurring far from equilibrium, such as mutual synchronization in oscillator communities, can be regarded as temporal analogs of order–disorder phase transition phenomena encountered in equilibrium and nonequilibrium lattice systems in materials science and other disciplines.

A laser can be viewed as an assembly of coupled (quantum mechanical) oscillators. Laser light in the form of coherent radiation is produced by the cooperative interactions between the oscillating atomic dipoles. In devices of this type, an inversion in the population distribution of atomic energy levels is produced by energy from an external source. At low pumping rates, the laser emission is incoherent, i.e., noisy. When the pumping rate exceeds the laser threshold, however, there is transition to a temporally ordered state in which coherent emission of radiation occurs.

In a ferromagnet, the spontaneous magnetization provides us with a measure of the global order emerging through the sequence of local spin–spin interactions. This order parameter is zero at and above the critical temperature and increases to unity as the temperature approaches zero. Order parameters, symmetry breaking (for a ferromagnet, the selection of a particular direction in space by the aligned spins), and critical fluctuations (large fluctuations in energy and magnetization in the vicinity of the phase-transition temperature) accompany the transition to the synchronous or coherent activity state in the laser, and in neural assemblies.

Learning and Memory

Neural signals are transmitted from cell to cell at sites called synapses. Presynaptic and postsynaptic junctions are electrically isolated from one another. At chemical synapses a change in potential at the presynaptic site triggers the release of neurotransmitter molecules, which diffuse across the synaptic cleft to trigger a change in potential at the postsynaptic site. The fact that the two cells are not physically joined to one another is crucial. It means that the efficiency of synaptic transmission can be modified through use, thereby making learning possible.

“Synaptic plasticity” is the name given to the adaptive changes in the efficiency of synaptic transmission triggered by use. The modifications may emerge in the form of anatomical and morphological changes, such as sprouting and pruning of axons and dendrites, naturally occurring cell death, and/or alterations in synaptic membrane properties. The term "synaptic plasticity" encompasses activity-dependent processes, such as long-term potentiation, long-term depression, and use-dependent modifications produced either surgically or by manipulations of afferent input. The mechanisms that promote these forms of synaptic plasticity are thought to underlie learning and memory.

The changes in postsynaptic potential produced by chemical transmission at any single synaptic site are small. The cooperative actions of many cells providing convergent input to a target cell within a short time span are required to evoke an action potential. It appears that changes in the frequency of presynaptic and postsynaptic activity determine not only the amplitude but also the sign of the changes in synaptic efficiency. Synapses whose activation is strongly correlated with the firing of postsynaptic cells are strengthened, while synapses that are silent during postsynaptic firing suffer reductions in efficiency. Synapses whose activation is not accompanied by postsynaptic firing exhibit reductions in efficiency as well. Synapses whose states can be modified by use-dependent, spatiotemporally correlated afferent input in a manner consistent with these rules are called Hebbian synapses.

Two forms of synaptic plasticity mentioned earlier, long-term potentiation (LTP) and long-term depression (LTD), have been studied intensively during the past 25 years, mostly in the region of the mammalian brain called the hippocampus. This region is associated with short-term memory. Injuries to the hippocampus result in an inability to form new memories, while long-term memories remain intact. These forms of synaptic plasticity seem to follow Hebbian rules. The threshold and magnitude of LTP and its inverse, LTD, are dependent on frequency and prior use. Rapid use produces LTP, slow extended use generates LTD, and the crossover from one to the other depends on prior activity.

The simplest mathematical formulation of a dynamical theory of synaptic modification is one in which we simply take the product of presynaptic and postsynaptic activities. If we construct a model of this type, we find that there is no way to prevent unlimited growth in the synaptic strengths. In the theory of Bienenstock, Cooper, and Munro (BCM), the modification of the strength m_i of the ith synaptic connection at time t is proportional to the product of the input activity d_i to the ith synaptic junction and a nonlinear modification function ϕ. The synaptic modification or learning rule assumes the form

$$\frac{dm_i(t)}{dt} = \eta\phi\big(c(t),\Theta_M(t)\big)d_i(t),$$

where η is a positive constant that controls the maximum change in strength during a single iteration, c is the summed, cooperative activity in the postsynaptic cell, and Θ_M is a dynamic (adaptive) modification threshold. The modification threshold Θ_M separates the modification function into two regimes, one positive and the other negative, depending on whether the postsynaptic activity is higher or lower than the modification threshold (see Figure 2). The dynamic modification threshold provides the stabilization of the strengths, allowing the strengths to converge to a fixed point, or equilibrium state, in phase space. It does so by sliding as a function of the history of the cell's activity and therefore adapts that cell's response properties.

In humans and other mammals, the neural circuitry in the visual cortex and elsewhere is profoundly influenced by experience during early postnatal life. The number of synaptic connections in children increases twenty-fold, from 50 trillion to 1000 trillion, during the first few months following birth. A cell's receptive field is the region of the sensory periphery (the area on the retina for the visual system,) whose illumination, or stimulation, influences the cell's activity. Researchers can study the early postnatal development of synaptic connections experimentally by manipulating the visual environment in such subjects as newborn kittens and monkeys, and observing the resulting changes in receptive fields.

Cells in the primary visual cortex develop the ability to discriminate between certain features in the visual environment. In particular, they become position-sensitive, develop well-defined ocularity properties, and have receptive fields that are timed to certain correlates in the visual scene. In BCM theory we are able to describe how cells develop their selectivity to features in the visual environment during early postnatal life. We can describe the emergence of orientation specificity, ocularity, and the interactions between orientation and binocular properties in newborn kittens.

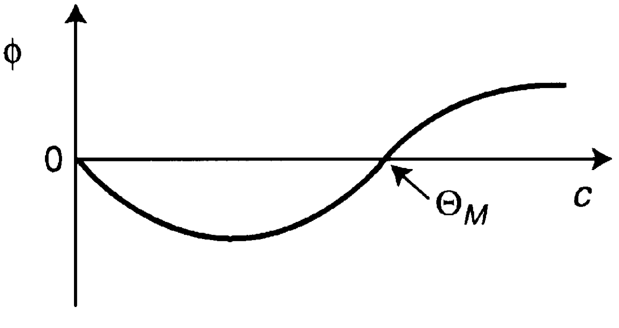

Figure 2. *Dynamical theory of synaptic modification: A nonlinear modification function ϕ is separated into two regimes, one positive and the other negative, depending on whether the postsynaptic activity is higher or lower than the adaptive modification threshold Θ_M.*

One of the findings that emerges from BCM dynamics is that the neurons in the primary visual cortex carry out a sophisticated statistical procedure involving moments higher than second. We can obtain this result by using an objective function approach to construct a stochastic differential equation whose solution converges to a local minimum in the risk in L^2. The local minima in the risk are the interesting features in the data, and our procedure yields a highly parallel feature extractor.

Another set of findings allows us to link synaptic modification in the visual system with hippocampal LTP and LTD, which are associated with learning and memory. When applied to these forms of synaptic plasticity, the mathematical model predicts that the threshold and magnitude of LTP and its inverse, LTD, should depend on frequency and prior use. The nonlinear modification function ϕ and the dynamic threshold Θ_M, both determined by mathematical arguments, imply that rapid use produces LTP, while slow, extended use generates LTD, with a crossover from one to the other that depends on prior activity, as observed experimentally.

Computation

The human brain contains multiple, interacting levels of organization, from genes and the proteins they express to hosts of molecular agents, synapses, and cells, to small circuits and large systems. We encounter many forms of cooperativity and parallelism both within and across these levels of organization. It is noteworthy that many of the results described in this article were obtained through simulations on parallel computers.

In our efforts to understand how the brain and its many subsystems work, mathematical models, and accompanying simulations on advanced architecture computers, are essential. High-performance computing will enable us to increase the amount of anatomical detail that can be included in models of brain circuitry. It will permit us to incorporate information pertaining to the biophysical substrate and will enable us to exploit the full integrative and explanatory capabilities of the mathematical models.

Acknowledgment

The author acknowledges many useful discussion with his fellow group member Judson P. Jones.

Martin Beckerman is a member of the Complex Dynamical Systems Group at Oak Ridge National Laboratory. He is the author of Adaptive Cooperative Systems *(John Wiley and Sons, 1997), in which recent findings on Markov random fields, annealing algorithms, BCM theory, and rhythms and synchrony in neurobiological systems are discussed in detail.*