



Hebb and Darwin

PAUL ADAMS

Department of Neurobiology, SUNY, Stony Brook, NY 11794, U.S.A.

(Received on 31 March 1997, Accepted in revised form on 19 December 1997)

It is proposed that vertebrate brains, especially those of mammals, operate according to an algorithm subsumable as “synaptic Darwinism”. The key postulate is that genes and synapses follow the same rules, because they act as autocatalytic, hypercyclic, units of selection. Synapses replicate by quantally strengthening, and mutate by connecting new cells. Because synapses relate pre- and post-synaptic firing, they perform a translation operation. Furthermore the product of this operation, conjoint firing, favors replication (by Hebb’s Rule). The result is that variants are selected and patterns of connection automatically adopt optimal configurations. These configurations are determined by scalar neuromodulatory “reward” signals applied globally to layers of neurons, which reduce spike frequency adaptation and enhance Hebbian replication. Global or local control of mutation rates provides further improvements in the Darwinian algorithm. All the processes and circuits postulated have plausible, and often obvious, implementations. The result is that brains evolve and adapt like large ecosystems.

© 1998 Academic Press

Introduction

Two of the most influential books in the history of biology are Darwin’s *On the Origin of Species* (1964) and Hebb’s *The Organisation of Behavior* (1949). The books are famous because of the simple but powerful ideas that they contain. In this paper it will be argued that the key idea of each book is essentially the same—hypercyclic autocatalysis with selection of favorable variants. It is because this Hebb–Darwin algorithm is so powerful that brains and species can efficiently adapt to everchanging environments.

The unity of the Hebb–Darwin algorithm can be illustrated most clearly in the modern language of genes and synapses. Genes are the fundamental unit of biology, and synapses of neurobiology.

In each science there are important higher and lower level entities (nucleotides, chromosomes,

genomes, populations, etc.; channels, neurons, circuits, etc.) but genes and synapses are fundamental because they are the units of replication, mutation and selection. Carl Sagan (1977) has pointed out that around 200 million years ago the complexity of brains (roughly, the information content of synapses) first exceeded the complexity of genomes (roughly, the information content of genes). At this point synaptic adaptation became in a sense more powerful than genetic adaptation. However neither form of adaptation is universal: when complexity exceeds some level, self-organisation can no longer be maintained, and there can be no further advance until some new, more efficient implementation of the Darwinian algorithm is implemented (such as language). Genes and synapses are also fundamental in several other more trivial senses. DNA is the most important

constituent of the most prominent part of a cell, the nucleus, and synapses are the most important part of the most prominent constituent of the brain, the neuropil. Genes and synapses are currently the focus of the most rapid developments in biology and neurobiology, and recent evidence shows that synapses are the basic functional units of neuronal integration (Yuste & Denk, 1995; Denk *et al.*, 1995).

The framework proposed here centers on the notion that genes and synapses are autocatalytic units each subject to the same basic operations of replication, mutation, translation, variation and selection. As a result, brains and ecosystems efficiently and continuously adapt to their environment. This framework integrates a number of disparate processes or structures in neurobiology, such as details of synaptic transmission and neuromodulation, thalamocortical circuitry, and awareness and sleep. The framework differs quite radically from that proposed by Edelman (1987), where the units of selection are neuronal groups, and repertoires of variants are not continuously generated.

Selfish Genes and Selfish Synapses

The two essential operations that a gene performs are (1) coding for another identical gene (replication) and (2) coding for a protein (translation; see Fig. 1). The essence of Darwinism (“Darwin’s Rule”) is that the encoded protein acts, directly or indirectly, as a replicase for the gene that encodes it. In a primitive hypothetical single-gene organism (to which the $Q\beta$ virus is an approximation; Dawkins, 1986; Eigen, 1992) the gene encodes a replicase (or other protein that is necessary for replication) which, because of compartmentation within a plasmamembrane, acts exclusively to catalyse the replication of that particular gene itself, and not other genes. Both the gene and the protein replicate, and because each assists each other’s replication, the arrangement is “hypercyclic” (Eigen, 1992), and very efficient. Errors in replication lead to replicase variants which may be more or less efficient at catalysing the replication of the encoding gene. As a result, a population of organisms modifies its gene pool composition to optimise overall replication

efficiency. “Adaptation” (or “survival of the fittest”) is this apparently purposeful optimisation of the gene pool composition. However, as Dawkins has emphasised, adaptation reflects the selfish, or individual, behavior of genes following a blind local rule that knows nothing about the goal or direction of adaptation. Darwin’s supreme insight was that this local algorithm could generate structures of incredible appropriateness and sophistication.

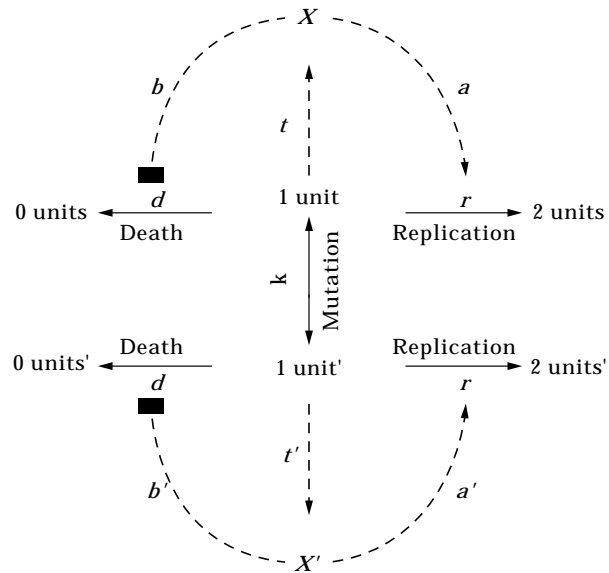


FIG. 1. Replication, Death and Mutation of Selection Units. This figure has two interpretations. In one, the unit is a gene: it can die, or it can duplicate itself. It also specifies (i.e. translation: t) an autocatalytic protein X , which acts as a replicase (a , arrow) or reduces (■) the likelihood of the gene disappearing (b ; since this stabilises the gene, it also enhances gene replication, acting indirectly as a replicase). In the other interpretation, the unit is a synapse. Its replication is enhanced by conjoint pre- and post-synaptic activity (X ; Hebb’s Rule), which is promoted (t) by the presence of the synapse, since the presence of the synapse enhances post-synaptic firing when the pre-synaptic cell fires (see Fig. 2 and text). Conjoint activity diminishes the probability that the synapse will disappear (b); however, again this process is formally similar to the enhancement of replication. The number of units grows or falls with a rate constant $w = (ar - bd)$, where r and d are the intrinsic replication and death rates (in the absence of autocatalysis) and a and b are the factors by which X multiplies the intrinsic rates. In addition, units occasionally undergo mutation, with probability k . Mutated units’ encode slightly altered products (X') which change replication and death rates by small amounts a' and b' . Note that mutation is bidirectional and can proceed to generate further variant units', etc. (not shown). Also note that the top and bottom of the figure constitute separate, compartmented hypercycles.

How does Hebb's Rule (synaptic strengthening by conjoint pre- and post-synaptic activity) lead to the same algorithm? First, it is necessary to define "synapse" a little more closely. Synapses perform logical or storage operations. Because the energies available in the rapid operations of the brain (such as action potentials) are only of the order of a few kT (the thermal energy), single molecules cannot be used to perform reliable logic or storage operations, and clusters of molecules—synapses—are required. The exact nature and arrangement of these molecules are not too important. For the moment, let us consider only excitatory synapses. An excitatory synapse has two key features. The first is that it physically connects two specific neurons, out of the very large number of neurons that it could connect. The second is that it relates the firing of the pre-synaptic neuron to the probability of firing of the post-synaptic neuron. In fact, we can define the synapse ji as a connection which increases the probability of firing of the i -th neuron in a post-synaptic array of equivalent neurons by an amount δq for a duration δt when the j -th neuron in a pre-synaptic array of equivalent neurons fires an action potential. If we assume that all synapses have the same δq and δt , we are essentially assuming a rather strong form of quantal transmission (von Kitzing *et al.*, 1994). Each synapse would be equivalent to a single release site, or to a unitary patch of post-synaptic receptors, or whatever turns out to be the basis of the "quantum" (Bekkers, 1994; Kuno, 1995). Although there is much evidence that synaptic transmission is quantal, the ideal synapse assumption made here is only an approximation: quantal size does show some variation, as a result of cable filtering, vesicle size scatter, etc., and quanta are released stochastically. However, neither of these deviations is likely to be practically important for Darwinian behavior, because the charge delivered to the soma by single quanta occurring at different dendritic locations is roughly constant, and the Hebb–Darwinian algorithm is itself probabilistic. The definition given implies *linearity*, since the amounts δq and δt are assumed to be independent of the post-synaptic firing level. This would be true if (uniquantal)

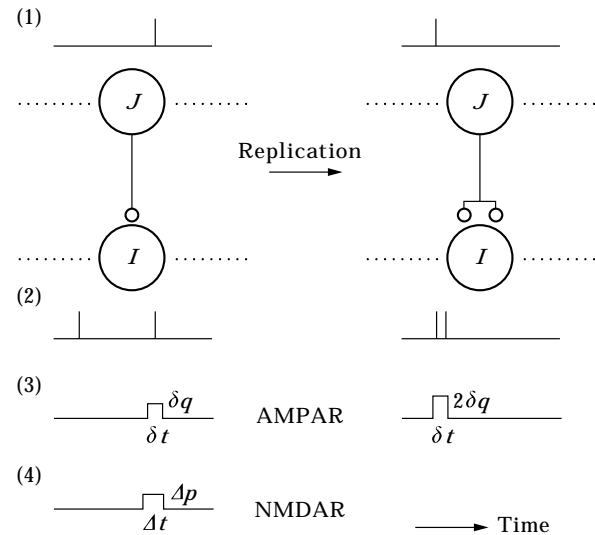


FIG. 2. Replication of Synapses. On the left, a synapse couples the j -th cell (J) in the pre-synaptic array with the i -th cell (I) in the post-synaptic array. A spike in J (row 1) increases the probability of firing of I by an amount δq for a time δt (row 3). This might induce a spike in I (row 2), which can also fire at other times (spontaneously or as a result of firing of other synapses). If spikes occur in both J and I within the Hebb interval Δt , then replication occurs with a probability Δp (right), i.e. the probability that a spike in J induces a spike in I is enhanced by an amount $\delta q \delta t$ (row 3), perhaps causing additional spikes in I leading to further replication. Note that trace 3 corresponds roughly to the duration of the AMPA component of the synaptic current, and trace 4 to the duration of the NMDA component.

synapses are all very weak and/or the post-synaptic cell is a leaky random walk integrator. The first assumption is very likely true, because unquantal currents are in the pA range, far short of typical nA rheobases. The second assumption has been recently discussed (Shadlen & Newsome, 1994), and is considered further below.

Given this definition of an ideal excitatory synapse, the synaptic equivalent of replication is straightforward (Fig. 2). It corresponds to strengthening. If a synapse becomes biquantal, it has replicated. The new synapse is identical to the old. It also connects the same j -th and i -th neurons, and increases the post-synaptic firing rate by the same amount $\delta q \delta t$ in response to the same pre-synaptic activity as does the "parent" synapse. The replication can be considered semiconservative, in the sense that there is no operational distinction between the "parent" and the two "offspring". The two operations encoded by a gene (self-replication

and production of a self-replicase) are also encoded by a synapse, provided that synaptic strengthening follows Hebb's Rule. The "replicase" that allows replication of the synapse is conjoint pre- and post-synaptic firing (of the j -th and i -th neurons), and it is precisely this "replicase" that the synapse promotes. The synapse "translates" the pre-synaptic activity into the post-synaptic activity (using the stochastic $\delta q \delta t$ formula), and the conjoint activity promotes replication, according to the classic Hebb rule. To be more specific, the Hebb rule is formulated as "whenever the pre- and post-synaptic cell fire together (within an interval Δt), then the synapse replicates (i.e. strengthens) by an amount $\delta q \delta t$ with a probability Δp ". The two basic intervals, δt and Δt , correspond roughly to the durations of the AMPA and NMDA components of the synaptic currents respectively (e.g. Kuno, 1995).

It may be objected that Hebb's Rule, as thus formulated, is probabilistic, whereas Darwin's Rule is not. However, of course gene-based Darwinian evolution is highly stochastic. A particular organism has only a certain *probability* of replicating. Its genes may encode highly efficient replicases, but nevertheless the vicissitudes of life may extinguish it prematurely.

Autocatalytic replication of synapses is constrained by several factors, such as the presence of adequate levels of pre- and post-synaptic activity, and competition with other synapses within the same "niche" (or post-synaptic surface), for space, growth factors, etc. (Purves & Lichtman, 1985). This engenders a struggle for survival. In a case like a cerebellar Purkinje cell, where two different types of synapse (contributed by two different pre-synaptic arrays, the olivary climbing fibers and the granule cell parallel fibers) coexist, one must assume that they require different pre- and post-synaptic resources. Each synapse can be viewed as selfishly striving to monopolise the activity of its post-synaptic cell. Each synapse on a given post-synaptic cell which originates from a given pre-synaptic cell is a copy of itself. Thus each climbing fiber synapse is a member of a large uniform clone ($> 10^2$), while each parallel fiber synapse is a member of a large ($> 10^5$) diverse population (Eccles *et al.*, 1967).

Mutation, Variation and Selection

Hypercyclic replication must be imperfect for adaptation to occur. Replication inaccuracies (i.e. mutations) ensure a supply of variant genes which replicate at different rates (assuming that the replicase is a true catalyst, and influences a rate limiting replication step). Let us assume that sequence differences between variants can be represented by a Hamming-like distance x , which determine a replication rate w (defined in Fig. 3) which is a function of x . The rate of change of the number of units at a particular value of x and t , (z), will be given by the sum of a selection/growth term and a diffusion term: $\partial z / \partial t = (w - \bar{w}) \cdot z + k \partial^2 z / \partial x^2$, where k is the mutation rate and \bar{w} is the population average net replication rate or "fitness" (Fig. 1; Fisher-Eigen equation; Volkenstein, 1994). The result is that the population, though remaining constant in number, migrates to a local fitness optimum (Fig. 3; see Ginzburg, 1983; Futuyma, 1986), roughly as a Gaussian cloud (Ebeling *et al.*, 1984).

Since the mutation rate is normally fixed, the population at the local optimum (a "quasispecies": Eigen, 1992) has an average fitness lower than the fittest member of the population. If the mutation rate is itself a genetically determined character, then a population can further increase its average fitness by selection of lower mutation rate variants. However, if the position of the fitness optima (with respect to the gene's location in sequence space) varies with time, selection will *increase* mutation rates. Most variation in sexually reproducing species is procured through heterozygosity and recombination. Essentially, non-lethal mutations are stockpiled in shielded forms which, in a perfectly uniform, constant environment, slightly lower average fitness. However, the principle is essentially the same as for asexual reproduction. To what extent levels of mutation, dominance, recombination, etc. are set by "group selection" mechanisms is a matter of controversy, hingeing much on precise definitions of "fitness" (see e.g. Eldridge, 1995). Perhaps the simplest view is that all selection is "group" selection, with a group size of 1 being a special, important case.

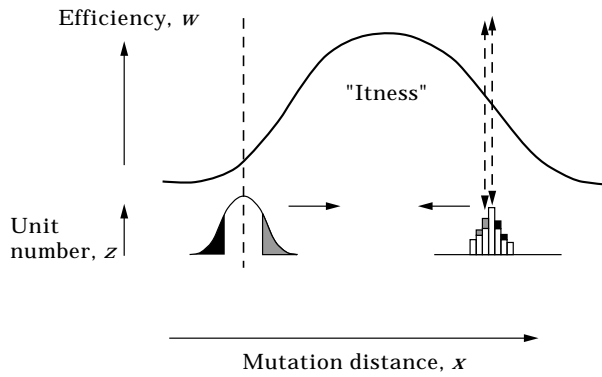


FIG. 3. Individual and Ensemble Selection. The top curve illustrates how the replication rate constant efficiency $w(= (ar - bd)$; "itness": fitness for genes or britness, for synapses) of X s encoded by different versions (alleles) of a unit (see Fig. 1: X is replicase activity or conjoint firing; the units are different versions, related by mutation, of a gene or synapse; mutation distance is defined in the legend to Fig. 4) varies depending how much the unit differs from some arbitrary starting value (---). The frequency distribution, shown under the itness curve, shows the actual numbers of the different versions of a unit after some arbitrary time, relative to the starting distribution (all genes or synapses identical), again as a function of mutation distance. Individual Selection (left, smooth distribution): in the absence of selection, mutation results in a gradual diffusive broadening of the frequency distribution. However, since the replication rate of a unit at a given distance from the starting value is determined by the itness curve (i.e. Darwinian selection occurs), the frequency distribution will shift towards the peak of the itness curve, maintaining a spread that depends on the mutation rate (i.e. the diffusion constant) and the selection strength (i.e. the width and height of the itness curve). This occurs because units with higher itness (such as those shown shaded) multiply faster than those with lower efficiency (solid). Note that if the total population does not grow as it migrates (normalisation), this corresponds to clamping the mean fitness (e.g. by supplying some fixed rate-limiting resource) or the conjoint rate (by inhibiting the entire post-synaptic array at a level proportional to the average level of conjoint activity). Ensemble Selection (right histogram): here the population is smaller, so that temporal fluctuations in the position of the population mean (arrows) are significant. If the population fluctuates to the left (increase, shaded, in fitter versions; decrease, solid, in less fit versions) a global reward is delivered to the whole population to increase its replication rate. The reward is proportional to the slope of the itness curve between the old and new population means (arrows). Note that a very simple itness curve is shown. In reality, gene sequence space forms a hypercube (Eigen, 1992). The two-dimensional curve shown would arise if replicase efficiency depended only on the *fraction* of bases of one type in a binary gene, not their *sequence*. Also, real fitness (or britness) landscapes will show many peaks, and trapping at local maxima can occur. It is assumed that real landscapes can be approximated as sums and products of Gaussians, so the process shown corresponds to migration to local optima.




FIG. 4. Synaptic Mutation. Two alternative (but equivalent) definitions of synaptic mutation. On the left, a mutation of the ji synapse is formed by j onto a *not- i* cell [in this case the neighboring $(i + 1)$ -th cell]. If the j -th pre-synaptic cell synapses onto the $(i + x)$ -th cell, this is considered to be a "larger" mutation. A mutation of the ji synapse to form a $j(i + x)$ synapse is a Hamming-like "distance" x away from the original synapse. For example, the synapse $j(i + 1)$ is unit distance away from the ji synapse. For a gene, mutation distance can be defined analogously, as the number of bases by which two genes differ (Eigen, 1992; Volkenstein, 1995). In both cases mutation as a function of x is modelled as a diffusion process. On the right, an alternative, equivalent, definition of mutation, not used in this paper, is shown.

but it is far too simple. Usually there are two or more peaks (and many dimensions, see below). However this does not affect the qualitative behavior (see Ebeling *et al.*, 1984 and below).

What are the synaptic equivalents of "mutation" and "selection"? Given the above definitions of synapse, replication and translation, it follows that there are two possible definitions of mutation (Fig. 4). Consider the synapse ji . A mutation must be either a connection from pre-synaptic cell j to a different post-synaptic cell *not- i* , or a connection from a different pre-synaptic cell *not- j* to the same post-synaptic cell i . The behavior of the network does not depend on the definition (because of symmetry), but to be consistent we will choose the first. As long as the projective fields (Churchland & Sejnowski, 1989) of the i and not i cells differ, synaptic mutation changes the nature of the coupling between the pre-synaptic neuron and the post-synaptic array, albeit (given the weakness assumption), infinitesimally. If it is assumed that the chances of a synaptic mutation (an erroneous replication) are greater the closer the "wrong" post-synaptic cell $(i + x)$ is to the "correct" post-synaptic cell (i) (Fig. 3), then the Fisher-Eigen equation will describe synaptic evolution. In this case $w(x)$ is set by the conjoint

firing rate of the j -th pre-synaptic cell and the $(i + x)$ -th post-synaptic cell.

Consider a particularly simple pair of pre- and post-synaptic arrays, containing one pre-synaptic neuron (A), and two post-synaptic neurons (0 and 1). Let us imagine the *de novo* birth of “mental life”—the formation of the first synapse (Fig. 5). Equivalently, we can imagine the formation of the first single gene organism, a gene which has only two alleles (0 and 1). Let us suppose that these two alleles encode replicases of different efficiency (in a particular, fixed, environment). 1 is more efficient than 0. Whether the first organism is a 0 or a 1, eventually 1s will be far more numerous. The population of organisms has adapted to the environment. It is important to note that the direction of evolution has been biased by the assumption that allele 1 encodes a better replicase than allele 0 (within the particular environment).




FIG. 5. Individual Selection of Synapses. This shows the evolution of a very simple array pair. The single pre-synaptic neuron A can connect either to post-synaptic cell 0 or to cell 1. The bars below show the firing rate of the post-synaptic cells in the two cases (open: actual; filled: ideal). If A connects to 0 then firing of 0 is more likely than firing of 1. Since the observed firing of 0 is slightly less than the target, the selection signal (or britness, i.e. the value read off the efficiency curve in Fig. 3) should be small and negative (■). Conversely, since the firing of 1 is much less than the target, a strong positive selection signal is applied (arrow). The selection signals, or rewards, measure how far the array is from the target synaptotype. If A connects to cell 1, then the firing rate of 0 will be lessened, and the firing of 1 increased, resulting in appropriately modified selection signals. Note that if the observed output matches the ideal, selection signals are zero and the synaptotype is stable (corresponding to the peak of the britness curve in part A: note that in this case since there are only two versions of the synapse—A0 and A1—the britness curve in this case has only two values, and is not continuous.)

So far the evolution of our simple ideal synapses is not biased to one outcome or another. They will simply grow at a rate dependent on the average firing rates of the pre- and post-synaptic neurons. (The post-synaptic neurons are assumed to have random input from other synapses outside the array pair, or alternatively to show spontaneous random firing). To bias the array to one outcome or another, the outputs have to be compared to the desired, or target outputs. The target outputs essentially define a brain equivalent of fitness, which can be called “britness”. If a post-synaptic neuron’s output falls short of the appropriate target output, then synapses on to that neuron from the pre-synaptic array (in the present case, just the one pre-synaptic neuron) should strengthen or replicate. This will boost the firing of the post-synaptic cell, bringing it closer to the target. Conversely, if the neuron’s output exceeds the target, synapses on that neuron should be weakened. We have not yet considered synaptic weakening (long term depression). It corresponds to death of individual organisms, extinguishing that particular copy of the gene (Fig. 1). Death can simply be random—organisms or synapses have a certain probability of disappearing, which does not depend on their information content. However, evolution is more efficient if death *does* reflect inadequate adaptation, i.e. it is coupled probabilistically to a particular gene sequence or connection pattern. The signal that represents the difference between the target britness output and the actual output can be considered as a “reward”. The reward can be positive or negative. If the actual output exactly matches the target output, the reward signal is zero, since no change in the synaptic strengths to that cell is required. Clearly the effect of a positive reward signal must be to increase the probability of synaptic replication. It may be objected that since pre-synaptic firing rates can vary over a wide range, there will be many possible synaptic strengths that correctly match the various input levels to the desired output. This is related to the notorious dilemma in neurophysiology, the relationship of spatial and temporal codes (e.g. Groh & Sparks, 1992). Obviously, the nervous system has solved this dilemma. The dilemma, and the objection,

largely evaporate if the firing rate of a cell represents the *probability* that the feature it represents is present (Gnadt & Breznen, 1996), and if features are coarsely coded, with approximately Gaussian receptive or projective fields. This is closely related to Poggio's (1990; Poggio & Girosi, 1993) argument that the brain uses radial basis functions to exploit the "smoothness", or redundancy, of the world, and the proposed mechanism, retaining random changes that reduce the error, is similar to his proposed learning method. As noted above, the Fisher–Eigen equation is usually written so as to clamp the total population number. For similar reasons, strong global inhibition should be added to a post-synaptic array, to clamp the average coincidence rate (and thus the total number of synapses). These inhibitory neurons should measure the average amount of pre-synaptic–post-synaptic coincidence, and subtract this signal from the post-synaptic array. This can be accomplished if each post-synaptic axon makes a feedforward connection on the distal dendrites of the coincidence detecting interneuron, and each post-synaptic cell makes a recurrent connection on its proximal dendrites. Under these conditions synaptic evolution will result in *rearrangement* rather than growth of connections.

So far our description of Darwinian synaptic evolution matches the classic "perceptron learning rule", and our array is a simple linear "perceptron". This learning rule is well known to be Hebbian in spirit (Hertz *et al.*, 1991) and it essentially underlies more sophisticated procedures such as backpropagation (Minsky & Papert, 1988). However the analogy to Darwinian evolution has not often been stressed. A single-gened asexually reproducing organism can adapt to any arbitrary one-dimensional fitness landscape, where each local optimum represents a different niche occupied by a different race (within the physical constraints set by the gene-replicase system), provided that the gene can adopt a continuum of states that map onto fitness. In the absence of sex the distinction between race and species is blurred. Normally, of course, gene sequence space is high-dimensional, but the number of possible values along each dimension is quite small, and it is still reasonable

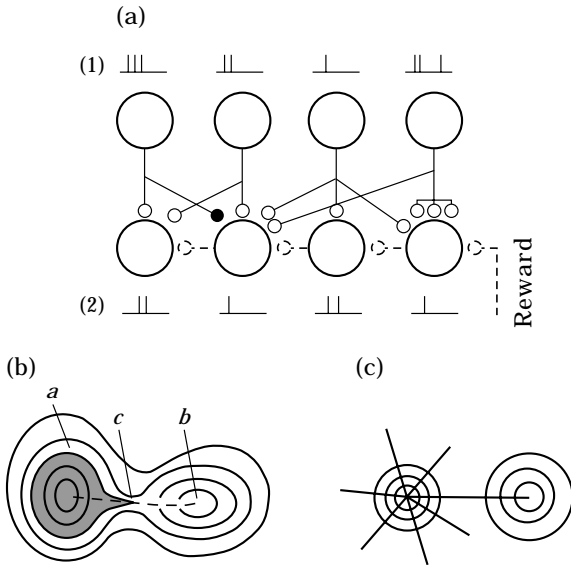


FIG. 6. Complex Landscapes. (a) shows one particular possible synaptic matrix or synaptotype (out of an astronomical possible number; see text). This synaptotype maps the vector representing the pre-synaptic activities of the array pair (row 1) onto the vector representing the post-synaptic activities (row 2). A tiny change (e.g. the mutated synapse shown as solid) slightly varies this vector mapping. If this improves the array output, the global reward system (dotted line) fires. The conjunction of a Hebb Rule and a global reward signal ensures that, in the long run, the mutated synapse will preferentially replicate. (b) shows a contour plot of fitness (or britness) where diffusion along two separate x - y distance axes confers varying degrees of fitness (the vertical axis). In the example shown, there is a simple ridge (---) connecting a local optimum *a* to a global optimum *b*. Darwinian selection (individual or group) promotes migration along this ridge, since the downward fitness gradient is least here. However, in order to explore the ridge the population (shown in gray) has to have a high enough mutation/selection ratio to reach the col (at *c*). This inevitably lowers the average fitness; (c) shows schematically a more realistic situation (with many genes or synapses, and strong epistatic or nonlinear interactions). Many steep ridges (---) lead away from a local fitness peak; however only one leads to a better optimum. Other things being equal, this ridge descends less steeply than the others, and it can thus be found using a local rule (such as topographic layer VI feedback control of mutation rate).

to map fitness in low-dimensional space, provided that it is remembered that there are far more "ridges" connecting fitness peaks than can be shown in a two-dimensional map (Fig. 6). As neurons are added to the post-synaptic array, the set of possible "alleles" increases, as does the set of possible britnesses that the array pair can learn (Fig. 6). Likewise a simple perceptron can solve the large class of problems known as "linearly separable", although it may take a long

time to do so. The similarity to perceptrons will become clearer when we consider multigene evolution.

Array Evolution

Single gene organisms do not exist, nor do single celled arrays (except perhaps in invertebrates, where connections are more likely to be specified genetically than by learning). Early in evolution single-gene organisms fused because genes can cooperate to enhance their own replication. Of course an inefficient gene that finds itself on the same chromosome or in the same genome as a highly efficient gene can hitch a ride (“linkage disequilibrium”), although not for ever. This dilutes the efficiency of the Darwinian algorithm, but this is largely offset by the widened possibilities of cooperation. Let us discuss “cooperation”. If two one-gene organisms fuse, it could be that each encoded replicase continues to act only upon its own encoding gene. Such an organism would have no advantage over its progenitors. However, each replicase might also be able to copy the other gene. The redundancy of such an arrangement would be useful, because each replicase could now vary to match not just the gene but also the availability of other essential ingredients. Furthermore, two different replicases, each perhaps separately suboptimal compared to the replicases of single-gene organisms, might come together as subunits of a superreplicase.

Consider now the equivalent situation in “Synaptic Darwinism”. A neuron is added to the pre-synaptic array, which can also form synapses on either post-synaptic cell. Given our previous definitions, the synapses formed by this cell cannot be considered as either replicas or mutants of the original synapses. These synapses are not alleles or replicas of each other, and we thus have an informationally new synapse. We can call all the synapses formed by the first cell A synapses and all the synapses formed by the second cell B synapses, and synapses formed by the j -th cell, J synapses. If all the synapses formed by the A cell are on post-synaptic neurons $0-k$, and those formed by the B cell are all on post-synaptic neurons $i-z$, then we have the equivalent of two one-gene species. If the

two pre-synaptic neurons can synapse on *any* post-synaptic cell we have the equivalent of one two-gene species. Since this is often the case, the entire array pair can be considered to be the equivalent of a single species (or an isolated population thereof). An individual post-synaptic cell is similar to an individual organism *and its progeny* (though synapses can die individually, whereas genes in an individual genome die collectively).

The 2×2 array can learn the simplest possible discrimination—to associate activity in cell A with activity in 1 or 2 or both, and to associate activity in B with 1 or 2 or both. To learn this the target britnesses for each mapping have to be provided. As is well known, it cannot (and no simple perceptron can) learn to be inactive (1 and/or 2) when both A and B are active if it is active when either A or B are active—the exclusive OR. $m \times n$ arrays can however learn quite complex tasks. They can learn to recognise unique combinations of their inputs, just as multigene organisms can adapt to utilise substrates that single gene organisms cannot use. This is a particularly strong form of cooperation, and requires that the strengthening produced by conjoint firing of cell j and i should be much greater than the sum of the strengthening produced by their separate firing. This is achieved by replacing the simple linear leaky integrator firing mechanism with a nonlinear, usually sigmoidal relation (Hertz *et al.*, 1991). It corresponds to the phenomenon of epistasis in genetics (Futuyma, 1986), where the contribution of genes to fitness is not additive, but depends on each other. If each gene contributes additively to fitness, then in principle the fitness contributed by each gene can be optimised separately, and if the fitness landscape for each gene is separately defined, so is the landscape for the ensemble (see also Poggio & Girosi, 1993). However if n genes interact, then the overall fitness landscape is of higher dimension.

What happens when an arbitrary starting synaptotype is placed in a complex britness landscape? The population initially diffuses into a Gaussian cloud that migrates to the local britness optimum. It then migrates coherently to nearby britness peaks along connecting ridges (Eigen, 1992; Volkenstein, 1994). However,

rather remarkably, this migration is invisible if the population is large—it essentially pseudo-quantum-mechanically “tunnels” or “catwalks” instantaneously to the new peak (Lande, 1985; Ebeling *et al.*, 1988)! The evolutionary equivalent is “punctuated equilibrium” (Eldridge, 1995).

Suppose that each post-synaptic cell in an $m \times n$ array pair receives p synapses. The number of possible different combinations of connection is $[(m + p)!/m!p!]^n$, which even for small values of m, n and p becomes superastronomical (and even larger if there is a *maximum* of p synapses). Of course the weakness assumption means that many similar configurations have similar average britness. However the britness landscape is sufficiently rich that a perceptron, especially a nonlinear one, can solve many problems. Because the britness landscape is less correlated for nonlinear perceptrons (corresponding to a high degree of epistasis), the danger of getting stranded on false optima is increased, something called by geneticists the “complexity-catastrophe” (Kaufmann, 1993). Possible strategies that the brain may use to overcome stranding are considered below.

Individual and Ensemble Selection

So far it has been assumed that the desired or target activity of each output neuron is known for each task. If this were the case, there would not be much point in training the array pair—one could just use the known targets as templates. This ideal situation is rather like artificial selection, where a breeder can examine each individual for conformity with some target, and correspondingly adjust its probability of reproduction (Perfect Artificial Individual Selection; PAIS). However the brain does not have this luxury. Its quandary resembles that of a breeder who wants to grow a strain of red bacteria, but cannot isolate individuals, or even colonies. The best he can do is to take a small rapidly mutating culture and measure its redness from time to time. It will fluctuate around some mean value (Fig. 3). Whenever its redness slightly exceeds this mean value, he briefly rewards it with extra substrate, leading to a small growth spurt. This procedure is repeated

whenever the colony turns slightly redder than the previous best recorded value. In order to keep the colony at constant size, it is punished whenever the redness drops by briefly withholding substrate. Inevitably the colony will gradually become redder (provided of course that redness is accessible within its protein sequence space). The point is that the reward does not have to be restricted to those individual mutants that are redder on average. It can be shared equally between all members of the colony. Spontaneous fluctuations in the average position of the colony along the artificial redness coordinate can be selected just as individual outliers at the leading edge of the population are selected under normal Darwinian evolution (Fig. 3). This can be seen more clearly when it is realised that this is just the temporal equivalent of the selection of different ensembles. If a large number of bacterial colonies were grown up, then one could just choose the reddest for further growth. It does not matter whether the fluctuations are temporal or spatial, as long as the memory of the fluctuation is preserved and amplified. We can refer to this procedure as Differential Artificial Ensemble Selection (DAES). Note that the breeder does not know, or care, whether the colony fluctuations toward redness were caused by an excess of redder mutations, or a chance increase in replication by redder organisms. DAES is closely related to Wright’s (1982; see also Lenski & Travisano, 1994) “shifting balance” model, where ensemble selection acts on finite populations whose means are fluctuating. DAES reduces to PAIS when the population size is one.

This argument is at first sight somewhat surprising because the redness of bacteria in the model does not influence the fitness of individual bacteria. However, it *does* influence the fitness of the entire colony. Consider an ensemble of small colonies that all start out equally pink. Some of these colonies will, by chance, die. The colonies that, by chance, become redder are less likely to die than ones that become, by chance, less red. Thus over time the entire collection of surviving colonies will become redder. Of course, if the colonies are allowed to grow too much, fluctuations will become negligible, and DAES will fail.

So to train an array pair without using specific target values for each neuron, one needs to provide a global reward signal that increases synapse replication across the entire post-synaptic array whenever the array's output moves closer to the desired goal, essentially the local slope of the brittness curve (Fig. 3). Normally, of course, an animal's brain measures its entire error, which is how far the animal is currently from its goal. This error measurement is itself quite a sophisticated operation, perhaps performed by the limbic system (Damasio, 1995). However, it is in principle much simpler than estimating how far each neuron in an array is from its best output—a scalar rather than a vector. In fact, a binary reward signal will often be adequate. The implementation of replication, mutation and reward is considered in the next section.

In summary, an array pair maps the pre-synaptic vector (the pattern of pre-synaptic firing, corresponding to the input to the array pair) onto the post-synaptic vector (the pattern of post-synaptic firing). The mapping is influenced by the pattern of synaptic connections within the pair (the synaptotype) and the nonlinearity of the post-synaptic neuron. The latter is assumed to be the same for all neurons, and influenced by a multiplicative scaling factor proportional to a “reward” signal. Mutation within the array pair changes the vector mapping; mutations that improve the array performance are rewarded. In addition, “punishment” signals can reduce the post-synaptic gain factor, which coupled with random synapse death eliminates bad synapses. DAES is closely related to the Associative Reward Penalty scheme (ARP) advocated by Barto (1990).

Implementation

Mutation ultimately represents the formation of new physical connections, presumably by initially random growth of pre-synaptic neurites within the array (Antonini & Stryker, 1993; O'Rourke *et al.*, 1994), as well as dendritic expansion (Ziv & Smith, 1996; Dailey & Smith, 1996). However neurite growth is a rather slow process, typically 1 mm per day, and is unlikely to generate the constant rapid mutation that is

necessary for most learning. Therefore I suggest that growth to establish new physical connections mostly occurs “off-line”, during sleep (Horne, 1988; Karni *et al.*, 1994). This phase could be called “pre-mutation”. Pre-mutation itself is presumably under genetic, hormonal and neuromodulatory control. The new physical connections are almost complete synapses, but lack an essential component. This component is very likely spine head NMDA receptors. If an array is performing satisfactorily, then no mutation or replication is immediately needed (though, because of mutation when other tasks are being solved, some array updating will be required). The pre-mutant synapses can be held in reserve for future use (reducing the need for subsequent nights' sleep). However if errors consistently occur, a neuromodulatory system is activated that increases the rate of insertion of post-synaptic NMDA receptors (Rs). A related situation for genes is the evolution of high mutation rates for antibodies or bacteria (Sniegowski *et al.*, 1996). Synapses where NMDARs have been recently inserted but which so far lack AMPARs are known as “silent synapses” (Liao *et al.*, 1995; Isaac *et al.*, 1996). An attractive candidate for this role is the activation of metabotropic glutamate receptors located on the distal dendrites of cortical and thalamic neurons (Godwin *et al.*, 1995, 1996). mGLUR activation is known to enhance NMDAR responses (Bleakman *et al.*, 1992; O'Connor *et al.*, 1993; Kinney & Slater, 1993; Bortoletto *et al.*, 1994). If, after a brief spurt of mutation, error is lessened, the newly recruited synapses replicate. In many systems, especially those closest to output and input, mutation at high rates only occurs during critical periods, when the overall patterns of synaptic circuitry are being laid down. Minor ongoing mutation keeps refining circuitry to account for ongoing cell death. Such ongoing rearrangement is well documented in sensory cortex (Wong *et al.*, 1995; Darian-Smith & Gilbert, 1994). However, arrays that compute more abstract stimulus features, or that formulate plans, continue to show mutation in adult life, until dendrites eventually become clogged up with synapses. Such deep arrays are probably much more “isomutational” (that is, individual neurons

have access to every member of the post-synaptic array) than superficial arrays, both because of higher mutation rates, and because they are more compact.

Conjoint pre- and post-synaptic activity causes insertion of AMPARs into the spine head of recently mutated synapses (Isaac *et al.*, 1996). The change from $n = 1$ to $n = 2$ constitutes synaptic replication. The initial change from $n = 0$ (silent) to $n = 1$ can be considered as the zeroth replication. It is subject to the same Hebb rule as subsequent replications. The insertion of a dollop of AMPA receptors increases the effective strength of the synapse from $n = 0$ to $n = 1$. Further conjoint firing may add more dollops of AMPA receptor. Further replication can increase n to 3 or 4, but AMPAR insertion could only occur for a few cycles before structural pre- and post-synaptic elaboration becomes necessary. This may correspond to the process of perforation and splitting described by Geneisman *et al.* (1993), and may also occur off-line during sleep. The newly formed synapses (with $n = 0$ or 1) would have the same connectivity as their parent, except for rare mutations. Silent synapses that fail to achieve this step might eventually be subject to NMDAR removal during sleep; some disassembly of pre-mutated synapses may also occur. However, under normal conditions conjoint activity is rather rare, because (a) both feedforward and feedback inhibition is very strong (Traub & Miles, 1991); (b) somadendritic potassium currents ensure rapid and powerful spike adaptation (e.g. Madison & Nicoll, 1986; Adams, 1987); and (c) action potentials generated in the axon hillock region do not efficiently back propagate into dendrites because of a combination of (a) and (b); Stuart & Sakmann, 1994; Markram *et al.*, 1995; Jaffe *et al.*, 1992; Spruston *et al.*, 1995). I refer to this condition as “Hebb-locked”. Most neurons are in the Hebb-locked state, otherwise the brain would rapidly explode with synapses, just as an exponentially growing bacterial culture will soon overflow an incubator. However, if a mutation spurt favorably, though slightly, changes the array pair synaptotype, and hence its vector mapping, a brief reward signal temporarily disengages the Hebb lock, and synapses replicate

(Fig. 7). Alternation of brief episodes of mutation and replication will produce convergence on a solution (not necessarily a globally optimal one).

There is in fact much evidence that neuromodulatory reward systems, releasing acetylcholine, noradrenaline, serotonin and dopamine, have the required properties. They globally innervate the telencephalon (Robbins & Everitt, 1995). They appear to be activated during the reward phase of learning, though not simply in response to the reward (Schultz *et al.*, 1993). Indeed they appear to have exactly the desired feature—they require the combination of an actual reward (which defines the desired direction of change) and some measure of the task performance error. If either the reward is absent or the error is zero, the dopamine neurons do not fire (Schultz *et al.*, 1993). Finally, and most persuasively, these reward systems do seem to specifically disable the Hebb lock (Madison & Nicoll, 1986, 1988; Malenka & Nicoll, 1986; Andrade & Nicoll, 1987). These neuromodulators reduce transmitter release from inhibitory neurons and reduce the size of potassium currents, especially those




FIG. 7. The Triangle of Awareness. The sketches show a cell's response to a short train of synaptic input. A cortical neuron can be in three different states: (a) depolarised, tonic but adapting (“aware” but Hebb-locked); (c) depolarised, tonic but non-adapting (“aware” but Hebb unlocked); (b) hyperpolarised and bursting (unaware). The ground state is (b); maintained depolarisation (e.g. via the metabotropic GluR) inactivates T-current and brings the neuron to (a); release of cholinergic and/or aminergic neuromodulators brings the cell to state (c). Note that in both states (c) and (b) the input–output gain is high; however, close spacing of action potentials in (b) is unfavorable for Hebb strengthening, because the burst precedes the main phase of NMDA receptor activation. In state (a), input–output gain is low, also unfavorable for Hebb strengthening. Reward pushes the aware cell to (c), punishment to (a).

involved in spike frequency adaptation and dendritic backpropagation.

A synaptotype which produces minimal error (maximum average britness) in a stable environment can be maintained by two alternate strategies: (1) ongoing intense selection, replication, death and mutation or (2) slowing or cessation of replication, death and mutation. The former is the bacterial strategy, the latter the metazoan. Obviously the former cannot maintain the synaptotype when the environment changes or disappears. A simple way to switch off replication, death and mutation is to convert synapses to a change-resistant state. A short-term way to do this would be to switch the pre-synaptic cells from tonic firing to burst firing (see below). In the longer term synapses might be consolidated to a less plastic physical form. This process could be called post-mutation. Post-mutation itself may involve various increasingly stable states (DeZazzo & Tully, 1995). Indeed, any memory mechanism is likely to involve a continuum of stability levels, since the more stable the memory the more energetic (and slower) the “write” process must be.

Catastrophes

The most difficult problem faced by any Darwinian system is the double catastrophe posed by complexity. As the number of units expands, the replication signal provided to the individual genotype (PAIS) or the evolving group (DAES) is diluted at the level of the individual units of selection, especially when fitnesses are merely additive. Selection is unable to hold complex genotypes at a local fitness optimum, and the population undergoes a phase-transition to much lower average fitness: the Eigen “error catastrophe” (Kaufmann, 1993; Volkenstein, 1995). The only simple solution to this problem, (other than restricting the number of units), is to increase epistasis, or in the synaptic case, to increase the nonlinearity of the post-synaptic transfer function. However, as already noted, this in turn increases the dimensionality of the fitness landscape, greatly increasing the number of local, but false, optima. Such an array pair can wander endlessly in an astronomically rich synaptic labyrinth without

ever encountering the global optimum. However, there are two possible solutions to this problem.

Mental Life

The first solution, advocated by Kaufmann (1993), is to maintain the evolving population at the error threshold. Actually, this is a sort of special biological case of “simulated annealing”, which in turn is connected to a fact about the world that provides a second solution—it exhibits regularity. Thus the fitness landscape (though high-dimensional) is somewhat correlated, and in particular fitness peaks tend to cluster in ranges or massifs. World correlations or regularities imply redundancy. The best search strategy therefore is to locate ranges at low resolution, and then refine within the range at finer grain. The Darwinian interpretation is straightforward—“melt” the population over the entire landscape by greatly increasing the mutation rate, until it finds the coarsest orographical features, and then gradually “cool” the population by lowering the mutation rate to find progressively finer, and higher, features within the landscape. Of course in both biology and neurobiology one cannot completely melt the population because some minimal level of fitness (basically the replacement level) must be maintained. However during active array evolution one should maintain the mutation rate as close to the error catastrophe threshold as possible. Fortunately, in “synaptic Darwinism”, continuous control of the global mutation rate (via a metabotropic receptor like the mGluR) is easy. Viruses also seem to operate at the error threshold (Eigen, 1992; Volkenstein, 1994).

The synaptotypic space of a large array pair is so vast that even at the error threshold convergence to the current britness optimum is unlikely. However, because of world redundancy, it may be necessary to mutate only *part* of the array pair. For example, following a small retinal lesion it is necessary to only rewire a small thalamic or cortical area centered on the lesion location. One of the roles of the massive topographic corticocortical and corticothalamic feedback projections is probably to control the effective local mutation rate, by releasing glutamate onto mGLURs located on distal

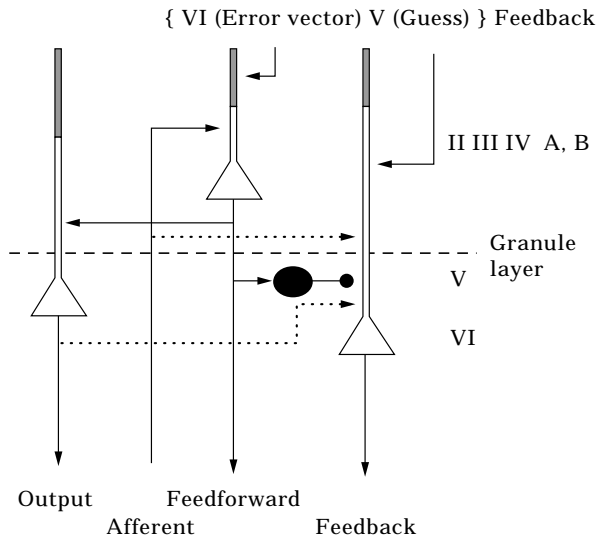


FIG. 8. The Basic Cortical Operation (---) shows the level of the spiny stellate cells (layer IV) which relay afferent information (from either first order or higher order thalamic relay nuclei; Sherman & Guillery, 1996) to the dendrites of supragranular (SG) cells (mainly Layers II/III). This pre-conditioning step is omitted, so that afferents are shown as synapsing directly on SG neurons. SG cells act as a hidden layer, extracting features of the input vector, and sending them to higher cortical areas (as afferents). Feedback to this cortical field from higher fields is of two types. Layer VI feedback controls the mutation rate in SG cells by activating mGluRs on apical dendrites (hatched). Layer V feedback (and possibly SG feedback) is compared to the output of SG cells (after sign inversion) in layer VI cells, which compute an error vector representing the difference between the features extracted in this field and that proposed by higher fields (the “guess”, “interpretation” or “expected feature”). This error vector is sent back to the distal dendrites of lower order SG or thalamic relay cells, where it in turn controls the local mutation rate, and hence allows adjustment of feature extraction. Recent evidence indicates that distal (apical) dendrites form a special physiological zone (Yuste *et al.*, 1994). Note, extensive lateral connections omitted! Layer V provides output to the rest of the brain, including second order thalamic nuclei (Sherman & Guillery, 1996). Some layer VI cells may compute an error vector based on the mismatch between afferent inputs and the layer V output.

dendrites (Godwin *et al.*, 1995, 1996; Zhou *et al.*, 1994). This would depolarise the target cells, and, at least in the case of the thalamus, the resulting tonic type of activity would result in a higher effective mutation rate (to a maximum set by developmental and hormonal regulation). It is likely that layer VI cells, the main origin of the backprojection, compute a local error signal which is used to control mutation rates topographically in lower order arrays (Fig. 8). If the britness landscape is high-dimensional, many

ridges lead downwards from local false optima. Use of an error vector to control melting effectively steers the evolving population along the least steeply descending ridges, the most efficient local search strategy in an extremely rugged terrain. This must have been one of the key innovations that allowed mammals to become intelligent (see also Shimizu & Karten, 1993). Note that local control of mutation rate and of “awareness”, procured by an mGluR-like device, operate efficiently in tandem. Feedback control of the local mutation rate is unlikely to be 1:1 because at least in the thalamus feedback axons greatly exceed feedforward axons (Sherman & Guillery, 1996). An efficient arrangement would be to have topographic control of mutation rate over variable sized patches of thalamus or cortex, with the largest layer VI cells innervating the largest patches, and the smallest cells innervating single neurons or even individual dendrites. Recruitment would follow a size principle (Henneman *et al.*, 1965). It is also possible that layer VI feedback provides a local reward signal, enhancing replication rather than mutation. It is known that mGluR activation inhibits adaptation-promoting potassium currents (Pedarzani & Storm, 1996). Such a local reward would promote individual, rather than group, selection. Furthermore, local feedback control of both mutation and selection could act in tandem. The pattern of these feedback connections would itself be established by Hebbian mechanisms.

The second type of solution to the error-complexity catastrophe also explores ranges before individual peaks. The solution is to actively cluster into ranges (i.e. to generalise or abstract) by using hierarchies of arrays, again exploiting the regularity of the world. A simple example is the clustering of units that represent purely topographic information to represent orientation (Das, 1996). Of course, this only works if the world indeed contains bars and edges. The ultimate aim is to cluster the whole world into two broad supermassifs, to measure the relative average heights, and to take a binary decision based on that measure. The entire space of vision is almost infinitely rich (despite the world’s considerable redundancy) and because of the

error catastrophe each array can only do a limited amount of clustering.

Array hierarchies as described are just multilayer neural networks. It has been pointed out that synaptic Darwinism, in its DAES form, is akin to ARP, whereas in its PAIS form it is like the perceptron learning rule. Therefore, since the perceptron rule generalised to multilayer networks is backpropagation, it is not surprising that ARP performs equivalently (Mazzoni *et al.*, 1991). Note that there is nothing in the definitions proposed which prevents a post-synaptic array from being its own pre-synaptic array—an autoassociative, rather than feedforward, network. Furthermore, a post-synaptic array can be the target of several separate pre-synaptic arrays, though there should be some mechanism for applying rewards selectively within such overlapping array pairs.

Neuromodulatory reward systems fire proportionately to how much closer a recent mutation episode has brought an array hierarchy to errorless task performance, so disabling the Hebb lock and allowing replication. Any real brain has thousands of arrays, and no convergence will occur if mutation occurs simultaneously in all, or many, of them—it would be like changing a thousand variables at one in a complex experiment. How is mutation and selection to be confined to a few arrays so that the problem space can be intelligently explored? There must be something like an *attentional* mechanism, which puts a few arrays into a susceptible, “aware”, state. Attention is an inevitable attribute of Darwinian evolution in a multiarray system.

Recent work in thalamus has revealed the probable mechanism for the attentional spotlight (Sherman & Koch, 1996). What should such a mechanism do? In the unaware state, neurons should be relatively quiescent or hyperpolarised. However they should be able to respond vigorously if a particularly clear signal arrives (high reliability). Furthermore, when so responding the response must (a) contain a signature so that downstream neurons can appropriately classify the information as coming from an “unaware” neuron; (b) need not encode every (appropriately transformed) detail of their input, since that information is not being directly used

to evaluate the all-important error function; finally (c) the unaware neuron should show strong spike frequency adaptation, so that Hebbian replication is not engaged. All these features are characteristic of the “burst” or low threshold spike mode of thalamic relay cells (Jahnsen & Llinas, 1984; Scharfman *et al.*, 1990; Sherman & Koch, 1996). In “tonic” mode, in contrast, neurons are depolarised, and their spike activity accurately encodes stimulus parameters (Guido *et al.*, 1995). Furthermore, those distal cortical feedback synapses which were proposed above to enhance mutation also shift cells from burst to tonic mode (McCormick & von Krosigk, 1992; Godwin *et al.*, 1996). This purely electrical change may be the main short-term control of the effective mutation rate. Thalamic cells do not show strong spike frequency adaptation in tonic mode, probably because in the adult there is little thalamic plasticity, and thus little need for a Hebb lock. However, cortical pyramidal cells show both spike frequency adaptation and T-current. These cells seem to exist in three basic modes: burst, tonic adapting (Hebb-locked) and tonic un-locked (Fig. 7). Recent *in vivo* recording shows that these three basic modes occur, though to what extent they are freely interconvertible is as yet unclear (Gray & McCormick, 1996). It is also possible that cortical cells are placed in “aware” or “unaware” mode not by intrinsic mechanisms, but by the pattern of their thalamic inputs (tonic or burst).

The suggested basic cortical operation is shown in Fig. 8 (cf. the “canonical cortical circuit” of Douglas & Martin, 1990). Layer IV is considered to be a modified version of Layers II/III, which pre-conditions lower order input and transfers it to Layers II/III. These function as a “hidden layer” and feed forward to both higher cortical areas and to layer V, the output layer. It is important that all cortex has an output layer, because selection guided by output-driven errors becomes progressively weaker the deeper the selected layer lies to output (a similar problem exists in “bucket-brigade” classifier systems in computer science (Holland, 1992). Weights in the hidden and output layers are optimised by the Hebb–Darwin reward algorithm, on the basis of rewards

calculated by the limbic system from the current and previous outputs, and the goal, and applied by the aminergic diffuse modulatory systems. Outputs from II/III and V are fed forward to higher level arrays for extraction of higher order features. They are also compared (probably after sign reversal) in layer VI neurons with layer V feedback from higher levels, which constitutes a sort of “guess” about what is “actually out there”. If the guess is locally inconsistent with what layer II/III actually reports, local error signals are sent back to lower order arrays, especially the thalamus, to initiate further mutation. The system thus automatically compares input signals, and a cascade of progressively more abstract extracted features, with previous interpretations of the world. If no consistent interpretation is possible, mutation is automatically initiated and tested to optimise feature extraction. Learning and development would be essentially the same process, though occurring in different epochs at different levels. Also, feature extraction and interpretation would be essentially the same process, at different levels of abstractness. Often, especially in non-humans, the system stops at a level of abstraction that generates behavior consistent with some biologically Darwinian goal.

Fetal Development and Sleep

How can useful connections be made or maintained while feedback from the environment is absent, during fetal life or sleep? If the world is redundant, then topographic mapping, which can be done in the absence of detailed feedback, can be a useful general principle. Such mapping means that neighboring neurons in a pre-synaptic array (whose activity, because of Gaussian receptive and projective fields, tends to be correlated) should project most strongly to neighboring cells in the post-synaptic array, provided that other correlations (provided by reward signals) are absent. It is well known that a simple Hebb principle promotes such correlations and neighborhood mappings (Cline, 1991; Peretto, 1992). Obviously, a sine wave pattern of activity sweeping across the pre-synaptic array will provide the optimal combination of high local correlation and modulation

depth. If the topographic mapping is to be accurate in two dimensions, these travelling waves should occur along random directions. These features seem to account for retinotectal and retinothalamic mappings (Meister *et al.*, 1991; Wong *et al.*, 1993; Wong, 1993; Wong & Oakley, 1996). Similar travelling waves of thalamic origin occur in slow wave sleep (Kim *et al.*, 1995). They might be involved in maintaining topographic mappings which would otherwise degrade during sleep.

A topographic mapping is most obviously useful as a “map” of the external world, which contains coherent objects. However, it can also provide a way of organising any set of extracted features (such as tone). Furthermore, the simple Hebbian generation of *locally* topographic mappings, in the absence of selection, may also underlie the generation of arrays that extract more abstract features. Consider the model suggested by von der Malsburg (1972) for orientation selectivity. This essentially uses real input “bars” to generate, via a Hebbian mechanism, cells that are tuned to bars. However, cortical cells show orientation selectivity before the animal experiences any coherent visual input (Chapman *et al.*, 1996). There are two possible, bizarre, solutions to this enigma. The first is that the appropriate wiring is somehow genetically specified. The second is that the brain creates virtual oriented travelling bars or stripes. Consider what would happen if regular travelling waves originate (presumably thalamically) at several points simultaneously. Their regular collision will create interference patterns, which will resemble slowly moving or stationary slightly curved fringes. These fringes, though virtual, will create local domains of orientation selectivity.

Thalamic travelling waves, from single or multiple origins, will not just topographically program sensory and motor maps (movement maps are really just the output equivalent of sensory bars), but the entire “association” cortex. However, to be useful, the basic structure imposed by simple non-selective Hebbian mapping has to be sculpted by detailed experience. It seems possible that mental “universals” reflect the combination of spatially correlated spontaneous activity, and the genetically-determined

TABLE 1.
Rules for genetic and synaptic Darwinism

Genes	Rules	Synapses
Unique nucleotide sequence	1. Information	Unique connection
Crick–Watson base pairing	2. Replication	Synaptic strengthening (LTP)
Disappearance of a copy of a gene	3. Death	Synaptic weakening (LTD)
Gene sequences specifies protein	4. Translation	A synapse specifies an input/output relation
The fate of certain genes are interdependent (e.g. organisms)	5. Bundling	All neurons within an array receive same neuromodulatory signals
Genotype specifies phenotype	6. Execution	Input vector specifies output vector
Erroneous replication leads to altered sequence	7. Mutation	Neurite growth leads to new connections
Slightly altered genotype leads to slightly altered phenotype	8. Variation	Slightly altered synaptotype leads to slightly altered output vector
Darwin’s Rule: “Survival of the Fittest”	9. Selection	Modified Hebb Rule: “Survival of the Brittest”

developmental sequence of successive temporal waves of heightened mutation. Of course array pairs exposed to a real, redundant world, will, because Hebb synapses are correlation detectors, also generate useful mappings, such as coordinate transformations (Mazzoni *et al.*, 1991; Salinas & Abbott, 1995), in the absence of spontaneous waves or biasing rewards.

Conclusion

Synapses and genes are both information-rich autocatalytic hypercyclic units that program complex behavior of the “bundles” of which they form part. The “bundle” is an association of units on which selection acts in tandem. The most important bundles are, for genes, the individual, and for synapses, the array pair, although selection can act at the ensemble (population or race) level, and at the individual (neuron) level. Selection acts to maximise the fitness (or britness) of the individual and/or group (and is weighted by the inverse of the group size; Williams, 1996). Synapses undergo operations like replication, translation, mutation, variation and selection (Table 1). These operations have straightforward neural implementations. Particularly important is the selection operation, which occurs by the

conjunction of mechanisms operating at the single synapse (Hebb’s Rule) and whole array (global neuromodulatory aminergic reward) levels.

Synaptic mutation is likely to occur by growth of new connections during sleep (premutation) and rapid insertion of NMDA receptors to create silent synapses (mutation proper). Replication (including the zeroth replication) is likely to occur by rapid (1 s), AMPAR insertion or disclosure (though early pre-synaptic changes are possible), followed by a type of consolidation involving pre- and post-synaptic elaboration and splitting. While the brain is disconnected from the real world during fetal development or sleep, yet synapse pre-mutation, consolidation and elimination is still occurring, accidental mutation and replication can degrade mappings set up in the awake state. Mappings can be maintained, or created, by travelling waves, which may interfere to create moving bar-like patterns.

Multiple hierarchical arrays, which extract progressively more abstract features, coevolve and compete. Most arrays are in a non-evolving “sketch” mode corresponding to “burst” firing; arrays which are actively evolving are in a detailed, “tonic”, “aware”, firing mode, upon which selective, global neuromodulatory reward systems can act whenever the animal moves closer to some “goal”. Burst mode arrays are still

capable of automatically performing already learned tasks, but cannot immediately modify themselves to meet unexpected contingencies.

The basic cortical module corresponds to a stochastic two layer backpropagation network (Mazzoni *et al.*, 1991). However, the computations of the hidden layer (II/III) are compared to feedback signals from higher order arrays, and an error vector computed in layer VI. This is returned to the distal dendrites of preceding cortical layer II/III cells, or thalamic relay cells, where it locally increases the mutation rate by switching cells to tonic mode. The system automatically extracts and tunes parallel hierarchies of abstract features, and updates feature analysis to incorporate novelty (cf. Logothetis *et al.*, 1995).

Considerable skepticism about “neural Darwinism” (Edelman, 1987) exists (Crick, 1989; Purves *et al.*, 1996). However, the present formulation, which may be dubbed “synaptic Darwinism”, seems to be free of many otherwise objectionable features. The unit of selection has been defined, and operations with units, such as replication, translation, mutation and selection, have been shown to correspond to broadly-understood brain processes. Because synapses are so abundant, and because they operate and rewire quite rapidly, synaptic evolution is likely to be much more efficient than evolution based on neuronal groups and selection from repertoires laid down during early development, and *a fortiori* than gene-based evolution. A great many brain processes and structures, such as neuromodulation, adaptation, bursting, electrical waves, sleep and thalamocortical circuitry, receive plausible interpretations. Furthermore, theoretical tools like Hebb’s Rule, backpropagation and ARP blend harmoniously with a Darwinian viewpoint. Perhaps the most attractive aspect of synaptic Darwinism is its close linkage between emergent phenomena in biology and neurobiology, without invoking unrealistic mechanisms. The key feature is amplification of microscopic fluctuations by autocatalysis, to produce dissipative macroscopic structures which are stabilised by selection to exploit the redundancy of the world.

Is the approach outlined above truly Darwinian? Here it is useful to distinguish

between the *units* of selection and the *level* of selection (Williams, 1992). Clearly the unit of selection is that entity which undergoes imperfectly accurate autocatalytic replication—the gene or the synapse. Selection is then applied in tandem to collections of units, or bundles. Single-gened organisms or linear array pairs with only one pre-synaptic neuron constitute minimal bundles. The largest possible gene bundle is the entire world biota, and the largest synaptic bundle, the brain itself. Darwinian selection is based on fluctuations in the composition of an ensemble of bundles. In the usual, biological, case, these fluctuations are spatial—at any one time part of the ensemble is fitter than other parts, and replicates faster. In the case of the brain, it is suggested that these fluctuations are temporal—the ensemble may be fitter at one time rather than another, again replicating faster. However, because replication and mutation are asynchronous, this difference is more apparent than real. Even in conventional natural selection, fluctuations are essentially compared to the mean of the ensemble. If they are fitter, they are amplified. Although the comparison is synchronic, the mean reflects the previously attained population mean, since replication “memorises” the composition of the gene pool, and is thus also diachronic. And in both cases, selection is based on the regularity of the world. Darwinian gene-based evolution creates the objects studied by biologists, and synapse-based evolution those studied by neurobiologists. It is because synapses are much more nimble and numerous than genes that mental adaptation is so much faster than physical adaptation.

Although the framework proposed here starts from biological considerations, it is quite similar in essence to that implicit in neural networks theories. The present framework places emphasis on Darwinian evolution as a fitness maximisation algorithm, while network theory uses analogies to statistical mechanics to stress that Hebbian learning is an energy minimisation procedure (Hertz *et al.*, 1991; Peretto, 1992). It is hoped that because more biologists are familiar with evolution theory than with statistical mechanics, the approach sketched here will be illuminating.

Note Added in Proof

Experimental evidence for synaptic replication has been recently reported by Petersen *et al.* [Peterse, C. C. H., Malenka, R. C., Nicoll, R. A. & Hopfield, J. J. (1998). All-or-none potentiation at CA3-CA1 synapses. *Proc. Nat. Acad. Sci. U.S.A.* **95**, 4732–4737].

I thank Don O'Malley, Lev Ginzburg, George Williams, Paul Bingham and, especially, Kingsley Cox for help, discussion and advice. I thank SUNY, Stony Brook for granting me sabbatical leave during which the idea for this paper was conceived.

REFERENCES

- ADAMS, P. R. (1987). Cholinergic hypothesis of Alzheimer's disease: biophysical aspects. In: *Molecular Neurobiology in Neurology & Psychiatry* (Kandel, E., ed.), N.Y.: Raven Press.
- ALGAR & NICOLL (1982).
- ANDRADE, R. & NICOLL, R. A. (1987). Pharmacologically distinct actions of serotonin on single pyramidal neurones of the rat hippocampus recorded *in vitro*. *J. Physiol.* **394**, 99–124.
- ANTONINI, A & STRYKER, M. P. (1993). Rapid remodeling of axonal arbors in the visual cortex. *Science* **260**, 1819–1821.
- BARTO, A. (1990). From chemotaxis to cooperativity: abstract exercises in neuronal learning strategies. In: *The computing neuron* (Durbin, R. M., Miall, R. C. & Mitchison, G. J., eds), pp. 73–98. New York: Addison-Wesley.
- BEKKERS, J. M. (1994). Quantal analysis of synaptic transmission in the central nervous system. *Curr Opin. Biol.* **4**, 360–365.
- BLEAKMAN, D., RUSIN, K. I., CHARD, P. S., GLAUM, S. R. & MILLER, R. J. (1992). Metabotropic glutamate receptors potentiate ionotropic glutamate responses in the rat dorsal horn. *Mol. Pharmacol.* **42**, 192–196.
- BORTOLETTO, Z. A., BASHIR, Z. I., DAVIES, C. H. & COLLINGRIDGE, G. L. (1994). A molecular switch activated by metabotropic glutamate receptors regulates induction of long-term potentiation. *Nature* **368**, 740–743.
- CHAPMAN, B., STRYKER, M. P. & BONHOEFFER, T. (1996). Development of orientation preference maps in ferret primary visual cortex. *J. Neurosci.* **16**, 6443–6453.
- CHURCHLAND, P. S. & SEJNOWSKI, T. J. (1989). Neural Representation and neural computation. In: *Neural Connections, Neural Computation*. (Nadel, L., Cooper, L. A., Cullicover, P. & Harnish, R. M., eds). Cambridge, MA: MIT Press.
- CLINE, H. T. (1991). Activity-dependent plasticity in the visual systems of frogs and fish. *Trends. Neurosci.* **14**, 104–111.
- CRICK, F. H. C. (1989). Neural Edelmanism. *Trends Neurosci.* **12**, 240–248.
- DAILEY, M. E. & SMITH, S. J. (1996). The dynamics of dendritic structure in developing hippocampal slices. *J. Neurosci.* **16**, 2983–2994.
- DAMASIO, A. R. (1995). *Descartes' Error*. New York: Putnam.
- DARIAN-SMITH, C. & GILBERT, C. D. (1994). Axonal sprouting accompanies functional reorganisation in adult cat striate cortex. *Nature* **368**, 737–740.
- DARWIN, C. (1964). *On the Origin of Species*. (Facsimile Edition). Cambridge, MA: Harvard University Press.
- DAS, A. (1996). Orientation in visual cortex: a simple mechanism emerges. *Neuron* **16**, 477–480.
- DAWKINS, R. (1986). *The Blind Watchmaker*. New York: OUP.
- DENK, W., SUGIMORI, M. & LLINAS, R. (1995). Two types of calcium response limited to single spines in cerebellar Purkinje cells. *Proc. Nat. Acad. Sci.* **92**, 8279–8282.
- DEZAZZO, J. & TULLY, T. (1995). Dissection of memory formation: from behavioral pharmacology to molecular genetics. *Trends in Neurosciences* **18**, 212–218.
- DOUGLAS, R. J. & MARTIN, K. A. C. (1990). Neocortex. In: *The Synaptic Organisation of the Brain* (Shepherd, G. M., ed.), New York: OUP.
- EBELING, W., ENGE, A., ESSER, B. & FEISTEL, R. (1984). Diffusion and reaction in random media and models of evolution processes. *J. Stat. Phys.* **37**, 369–384.
- ECCLES, J. C., ITO, M. & SZENTAGOTHAJ, J. (1967). *The Cerebellum as a Neuronal Machine*. New York: Springer-Verlag.
- ELDRIDGE, N. (1995). *Reinventing Darwin*. New York: Wiley.
- EDELMAN, G. M. (1987). *Neural Darwinism*. London: Basic Books.
- EIGEN, M. (1992). *Steps Toward Life*. London: OUP.
- FUTUYMA, D. (1986). *Evolutionary Biology*. Sunderland, MA: Sinauer.
- GENEISMAN, Y., DETOLEDO-MORELL, L., MORELL, F., HELLER, R. E., ROSSI, M. & PARSHALL, R. F. (1993). Structural synaptic correlate of long-term potentiation: formation of axospinous synapses with multiple, completely partitioned transmission zones. *Hippocampus* **3**, 435–446.
- GINZBURG, L. R. (1983). *Theory of Natural Selection and Population Growth*. Menlo Park: Benjamin Cummings.
- GNADT, J. W. & BREZNEK, B. (1996). Statistical analysis of the information content in the activity of cortical neurons. *Vision Res.* **36**, 3525–3537.
- GODWIN, D. W., VAN HORN, S. C., GUNLUK, A. E., SESMA, M. A., ROMANO, C. & SHERMAN, S. M. (1995). Localisation of two metabotropic glutamate receptors (mGluR1a and mGluR5) in cat LGN. *Soc. Neurosci. Abstr.* **21**, 658.
- GODWIN, D. W., VAUGHAN, J. W. & SHERMAN, S. M. (1996). Metabotropic glutamate receptors switch visual response mode of lateral geniculate nucleus cells from burst to tonic. *J. Neurophysiol.* **76**, 1800–1816.
- GRAY, C. M. & MCCORMICK, D. A. (1996). Chattering cells: superficial pyramidal neurons contributing to the generation of synchronous oscillations in the visual cortex. *Science* **274**, 10–19.
- GROH, J. M. & SPARKS, D. L. (1992). Two models for transforming auditory signals from head-centered to eye-centered coordinates. *Biological Cybernetics* **67**, 291–302.
- GUIDO, W., LU, S. M., VAUGHAN, J. W., GODWIN, D. W. & SHERMAN, S. M. (1995). Receiver operating characteristic (ROC) analysis of neurons in the cat's lateral geniculate nucleus during tonic and burst response mode. *Vis. Neurosci.* **12**, 723–741.

- HEBB, D. O. (1949). *The Organization of Behavior*. New York: Wiley.
- HENNEMANN, E., SOMJEN, C. & CARPENTER, D. O. (1965). Functional significance of cell size in spinal motoneurons. *J. Physiol.* **28**, 560–580.
- HERTZ, J., KROGH, A. & PALMER, R. G. (1991). *Introduction to the Theory of Neural Computation*. Redwood City: Addison-Wesley.
- HOLLAND, J. H. (1992). *Adaptation in Natural and Artificial Systems*. Cambridge, MA: MIT Press.
- HORNE, J. A. (1988). *Why We Sleep*. Oxford: Oxford University Press.
- ISAAC, J. T., HJELMSTAD, G. O., NICOLL, R. A. & MALENKA, R. C. (1996). Long-term potentiation at single fiber inputs to hippocampal CA1 pyramidal cells. *Proc. Nat. Acad. Sci.* **93**, 8710–8715.
- JAFFE, D. B., JOHNSTON, D., LASSER-ROSS, N., LISMAN, J. E., MIYAKAWA, H. & ROSS, W. N. (1992). The spread of Na⁺ spikes determines the pattern of dendritic Ca²⁺ entry into hippocampal neurons. *Nature* **357**, 244–246.
- JAHNSEN, H. & LLINAS, R. (1984). Electrophysiological properties of guinea-pig thalamic neurones: an *in vitro* study. *J. Physiol.* **349**, 205–226.
- KARNI, A., TANNE, D., RUBENSTEIN, B. S., ASKENASY, J. J. M. & SAGI, D. (1994). Dependence on REM sleep of overnight improvement of a perceptual skill. *Science* **265**, 679–682.
- KAUFMANN, S. (1993). *The Origins of Order*. New York: Oxford University Press.
- KIM, U., BAL, T. & MCCORMICK, D. A. (1995). Spindle waves are propagating synchronised oscillations in the ferret LGNd *in vitro*. *J. Neurophysiol.* **74**, 1301–1323.
- KINNEY, G. A. & SLATER, N. T. (1993). Potentiation of NMDA receptor-mediated transmission in turtle cerebellar granule cells by activation of metabotropic glutamate receptors. *J. Neurophysiol.* **69**, 585–594.
- KITZING, E. VON, JONAS, P. & SAKMANN, B. (1994). Quantal analysis of excitatory postsynaptic currents at the hippocampal mossy-fiber-CA3 pyramidal cell synapse. *Adv. Second Messenger & Phosphoprotein Res.* **29**, 235–260.
- KUNO, M. (1995). *The Synapse: Function, Plasticity and Neurotrophism*. Oxford: OUP.
- LANDE, R. (1985). Expected time for random genetic drift of a population between stable phenotypic states. *Proc. Nat. Acad. Sci.* **91**, 6608–6814.
- LENSKI, R. E. & TRAVISANO, M. (1994). Dynamics of adaptation and diversification: a 10 000-generation experiment with bacterial populations. *Proc. Nat. Acad. Sci.* **91**, 6808–6814.
- LIAO, D., HESSLER, N. A. & MALINOW, R. (1995). Activation of postsynaptically silent synapses during pairing-induced LTP in CA1 region of hippocampal slice. *Nature* **375**, 400–404.
- LOGOTHETIS, N. K., PAULS, J. & POGGIO, T. (1995). Shape representation in the inferior temporal cortex of monkeys. *Current Biol.* **5**, 552–563.
- MADISON, D. V. & NICOLL, R. A. (1986). Actions of noradrenaline recorded intracellularly in rat hippocampal CA1 pyramidal neurones *in vitro*. *J. Physiol.* **372**, 221–244.
- MADISON, D. V. & NICOLL, R. A. (1988). Norepinephrine decreases synaptic inhibition in the rat hippocampus. *Brain Res.* **442**, 131–138.
- MALENKA, R. C. & NICOLL, R. A. (1986). Dopamine decreases the calcium-activated afterhyperpolarisation in hippocampal CA1 pyramidal cells. *Brain Res.* **379**, 210–215.
- MALSBERG, C. VON DER (1972). Self-organisation of orientation sensitive cells in the striate cortex. *Kybernetik* **14**, 85–100.
- MARKRAM, H., HELM, P. J. & SAKMANN, B. (1995). Dendritic calcium transients evoked by single back-propagating action potentials in rat neocortical pyramidal neurons. *J. Physiol.* **485**, 1–20.
- MAZZONI, P., ANDERSON, R. A. & JORDAN, M. I. (1991). A more biologically plausible learning rule for neural networks. *Proc. Nat. Acad. Sci.* **88**, 4433–4437.
- MCCORMICK, D. A. & KROSIGK, M. VON. (1992). Corticothalamic activation modulates thalamic firing through glutamate “metabotropic” receptors. *Proc. Nat. Acad. Sci.* **89**, 2774–2778.
- MEISTER, M., WONG, R. O. L., BAYLOR, D. A. & SHATZ, C. J. (1991). Synchronous bursts of action potentials in ganglion cells of the developing mammalian retina. *Science* **252**, 939–943.
- MINSKY, M. L. & PAPERT, S. A. (1988). *Perceptrons*. Cambridge, MA: MIT Press.
- O’CONNOR, J. J., ROWAN, M. J. & ANWYL, R. (1993). Long-lasting enhancement of NMDA receptor mediated synaptic transmission by metabotropic glutamate receptor activation. *Nature* **367**, 557–559.
- O’ROURKE, N. A., CLINE, H. T. & FRASER, S. E. (1994). Rapid remodeling of retinal arbors in the tectum with and without blockade of synaptic transmission. *Neuron* **12**, 921–934.
- PEDARZANI, P. & STORM, J. F. (1996). Evidence that Ca/calmodulin-dependent protein kinase mediates the modulation of the Ca²⁺-dependent K⁺ current, I_{AHP}, by acetylcholine, but not by glutamate, in hippocampal neurons. *Pflugers Archiv.* **431**, 723–728.
- PERETTO, P. (1992). *An Introduction to the Modeling of Neural Networks*. Cambridge: Cambridge University Press.
- POGGIO, T. (1990). A theory of how the brain might work. *Cold Spring Harbor Symposia on Quantitative Biology*, LV, pp. 899–910. New York: CSHL Press.
- POGGIO, T. & GIROSI, F. (1993). Learning algorithms and network architectures. In: *Exploring Brain Functions: Models in Neuroscience*. (Poggio, T. A. & Glaser, D. A., eds), pp. 77–96. Chichester: Wiley.
- PURVES, D. & LICHTMAN, J. W. (1985). *Principles of Neural Development*. Sunderland, MA: Sinauer.
- PURVES, D., WHITE, L. E. & RIDDLE, D. R. (1996). Is neural development darwinian?. *TINS* **19**, 460–464.
- ROBBINS, T. W. & EVERITT, B. J. (1995). Arousal systems and attention. In: *The Cognitive Neurosciences*. (Gazzaniga, M. S., ed.). Cambridge, MA: MIT Press.
- SAGAN, C. (1977). *The Dragons of Eden*. NY: Random House.
- SALINAS, E. & ABBOTT, L. F. (1995). Transfer of coded information from sensory to motor networks. *J. Neurosci.* **15**, 6461–6474.
- SCHARFMAN, H. E., LU, S. M., GUIDO, W., ADAMS, P. R. & SHERMAN, S. M. (1990). N-methyl-D-aspartate (NMDA) receptors contribute to excitatory postsynaptic potentials of cat lateral geniculate neurons recorded in thalamic slices. *Proc. Natl Acad. Sci. USA* **87**, 4548–4552.

- SCHULTZ, W., APICELLA, P. & LJUNGBERG, T. (1993). Responses of monkey dopamine neurons to reward and conditioned stimuli during successive steps of learning a delayed response task. *J. Neurosci.* **13**, 900–913.
- SHADLEN, M. N. & NEWSOME, W. T. (1994). Noise, neural codes and cortical organisation. *Curr. Opin. Neurobiol.* **4**, 569–579.
- SHERMAN, S. M. & GUILLERY, S. M. (1996). Functional organisation of thalamocortical relays. *J. Neurophysiol.* **76**, 1367–1395.
- SHERMAN, S. M. & KOCH, C. (1996). Thalamus. In: *The Synaptic Organisation of the Brain* (Shepherd, G.M., ed.). Oxford.
- SHIMIZU, T. & KARTEN, H. J. (1993). The avian visual system and the evolution of the neocortex. In: *Vision, Brain and Behavior in Birds*. (Zeigler, H. P. & Bischof, H.-J., eds). Cambridge, MA: MIT Press.
- SNIEGOWSKI, P. D., GERRISH, P. J. & LENSKI, R. E. (1997). Evolution of high mutation rates in experimental populations of *E. coli*. *Nature* **387**, 703–705.
- SPRUSTON, N., SCHILLER, Y., STUART, G. & SAKMANN, B. (1995). Activity-dependent action potential invasion and calcium influx into hippocampal CA1 dendrites. *Science* **268**, 297–300.
- STUART, G. J. & SAKMANN, B. (1994). Active propagation of somatic action potentials into neocortical pyramidal cell dendrites. *Nature* **367**, 69–72.
- TRAUB, R. D. & MILES, R. (1991). *Neuronal Networks of the Hippocampus*. Cambridge: Cambridge University Press.
- VOLKENSTEIN, M. V. (1994). *Physical Approaches to Biological Evolution*. Berlin: Springer-Verlag.
- WANG, X., MERZENICH, M. M., SAMESHIMA, K. & JENKINS, W. M. (1995). Remodelling of hand representation in adult cortex determined by timing of tactile stimulation. *Nature* **378**, 71–75.
- WILLIAMS, G. C. (1992). *Natural Selection: Domains, Levels and Challenges*. New York: Oxford University Press.
- WILLIAMS, G. C. (1996). *Adaptation and Natural Selection*. Princeton: Princeton University Press.
- WONG, R. O. L. (1993). The role of spatio-temporal firing patterns in neuronal development of sensory systems. *Curr. Opin. Neurobiol.* **3**, 595–601.
- WONG, R. O. L. & OAKLEY, D. M. (1996). Changing patterns of spontaneous bursting activity of On and Off retinal ganglion cells during development. *Neuron* **16**, 1087–1095.
- WONG, R. O. L., MEISTER, M. & SHATZ, C. (1993). Transient period of correlated bursting activity during development of the mammalian retina. *J. Neurosci.* **15**, 2696–2706.
- WRIGHT, S. (1982). Character change, speciation and the higher taxa. *Evolution* **36**, 427–443.
- YUSTE, R. & DENK, W. (1995). Dendritic spines as basic functional units of neuronal integration. *Nature* **375**, 682–684.
- YUSTE, R., GUTNICK, M. J., SAAR, D., DELANEY, K. R. & TANK, D. W. (1994). Ca²⁺ accumulations in dendrites of neocortical pyramidal neurons: an apical band and evidence for two functional compartments. *Neuron* **13**, 23–43.
- ZHOU, Q., GODWIN, D. W., BICKFORD, M. E., SHERMAN, S. M. & ADAMS, P. R. (1994). Relay cells and local GABAergic cells contribute to responses mediated by metabotropic glutamate receptors in cat LGN. *Soc. Neurosci. Abstr.* **20**, 133.
- ZIV, N. E. & SMITH, S. J. (1996). Evidence for a role of dendritic filopodia in synaptogenesis and spine formation. *Neuron* **17**, 91–102.