

9 Chemical, Neuronal, and Linguistic Replicators

Chrisantha Fernando and Eörs Szathmáry

The Modern Synthesis can be extended laterally and vertically. Lateral extensions transfer the thought patterns and the methodology of evolutionary theory to different, previously nonevolutionary disciplines. Examples include replicator and systems chemistry, linguistics, cultural selection (Boyd and Richerson 2005), memetic replicators (Aunger 2002; Dawkins 1976), somatic selection (Edelman 1994), and microeconomics. Vertical extensions deepen our knowledge in traditional areas of evolution research such as EvoDevo, niche construction, epigenetic inheritance (Jablonka and Lamb 2005), and multiple levels of selection (Okasha 2006). In this chapter we first review the relevance of evolutionary thinking in certain areas of chemistry. Then we present, as the major novel contribution, a lateral extension to neuroscience, in that we outline a truly evolutionary approach to brain function in the higher vertebrates. Finally, some relevant aspects of language will be considered.

Evolution by natural selection is perhaps the most important process acting in populations of living systems. This is one of the reasons why it is so tempting to equate units of evolution (i.e., an abstract generalization that makes no reference whatsoever to any particular level of biological organization) to units of life. Another reason is that units of evolution can be much more readily defined. There are a few known alternative formulations of the concept of units of evolution; here we stick to the version outlined by Maynard Smith (1986): such units must multiply, show heredity across generations (like begets like), and heredity should not be exact. If some of the hereditary traits affect the chance of reproduction and/or survival of the units, evolution by natural selection can take place in a population of such units. The combination of survival and reproduction (translating into the expected number of descendants) is called fitness. The above characterization of Darwinian dynamics is deliberately general: note that it is not restricted to cover

living systems only. (As a matter of fact, some living systems do not—sometimes cannot—multiply: mules and neurons normally do not reproduce.) Hence it is potentially applicable to molecules and cultural traits as far as the criteria really apply.

A general point about definitions is that they cannot be falsified. They have to be internally consistent, of course, but there can be an arbitrary number of such definitions for life, for example. It is the use of the alternative definitions that makes the difference: some definitions are found helpful because they categorize natural phenomena in a way that is conducive to further insights. There is always an ingredient of arbitrariness in definitions: we have to live with this fact.

Chemical Origin of Evolvability and Systems Chemistry *in Statu Nascendi*

Biology has its roots in chemistry (Von Kiedrowski 2001). Gánti (e.g., 2003) emphasized that contemporary living systems always have (1) some metabolic subsystem, (2) some systems for heritable control, and (3) some boundary system to keep the component together. We consider it unlikely that a chemical system satisfying all the constraints from this abstraction could have appeared just out of chemical chaos. This observation led to the formulation of the concept of infrabiological systems (Szathmáry 2005; Fernando et al. 2005). Infrabiological systems always lack one of the key components just listed. For example, in the original formulation of Gánti (1971), a model of minimal life did not include a boundary system. The combination of a metabolic cycle and a membrane was also conceived by Gánti (1978), and called a self-reproducing microsphere. In contrast, Szostak et al. (2001) conceived a protocell-like entity with a boundary and template replication but no metabolic subsystem. Such systems show a crucial subset of interesting biological phenomena. The three subsystems can be combined to yield three different doublet systems (figure 9.1).

The emerging field of systems chemistry deals with the analysis and synthesis of coupled autocatalytic systems (e.g., Kindermann et al. 2005; Ludlow and Otto 2008). Chromosomes made of DNA come in different lengths. They can harbor a small or a large number of genes. During replication of the bacterial chromosome, it makes perfect sense to say that replication is half complete when one half is already present in two copies. This sharply contrasts with the following example. Imagine a molecule A, which reacts with a number of compounds to yield two molecules of A after one turn of the cycle. Molecular systems of this




Figure 9.1

Elementary combinatorics of infrabiological systems (Fernando et al. 2005). The chemoton is a minimal biological system comprising three qualitatively different subsystems (metabolism, membrane, and template). The analysis and synthesis of such systems is the aim of the emerging field of systems chemistry (Kindermann et al. 2005; Ludlow and Otto 2008). (From Fernando et al. 2005)

kind do exist; examples include the formose reaction (figure 9.2), the reductive citric acid cycle (which is almost the exact reverse of the citric acid cycle and is used for carbon fixation by some bacteria), and the Calvin cycle (fixing carbon dioxide in plants). One, or a few, autocatalysts are sufficient to seed the system, and parts (the chemical moieties) of the autocatalytic molecules are held together by covalent bonds (and are thus sterically constrained). They are also stoichiometric in the sense that the elementary steps are simple chemical reactions (transformations). Two questions must be asked about such systems: (1) Are they feasible as autonomous replicators (self-replicators)? (2) Is there hereditary information stored in them? We discuss these questions in turn.

It is important to emphasize that all sufficiently well described metabolic networks contain at least one (sometimes several) autocatalytic metabolic seeds without which the cell cannot start running, despite the presence of all genes and enzymes (Kun et al. 2008). However, the reductive citric acid cycle and the Calvin cycle (which is in fact a complex network) are not autonomous, in the sense that they require the operation of enzymes that are not produced by them. This is in contrast to the formose reaction, which does not require enzymes.

Heredity requires alternative types of cycles. Currently, there are only hypothetical suggestions, put forward by Wächtershäuser (1988, 1992): they are various extensions of the (equally hypothetical) “archaic” reductive citric acid cycle. Even if alternative forms of such systems can exist, most changes will be mere fluctuations and will not lead to hereditary alterations (“mutations” in the general sense). It is expected that the system will flip from the basin of one attractor into that of another attrac-

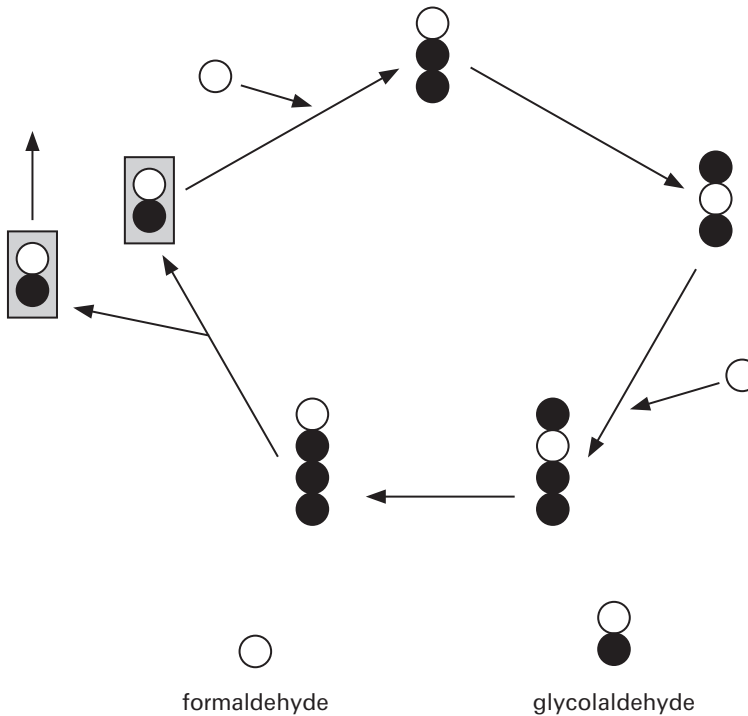


Figure 9.2

The autocatalytic core or seed of the formose reaction (Fernando et al. 2005). Each circle represents a chemical group including one carbon atom. Black and white circles denote different internal chemical structures.

tor very rarely; hence there will be infrequent “macromutations” only (Wächtershäuser 1988). Nevertheless, such macromutations may have been of paramount importance in chemical evolution. Models suggest that chemical evolution can occur by natural selection if a generative chemical system capable of producing satellite autocatalytic cycles is enclosed in a compartment. Variation is by “chemical avalanches” (macromutations), as suggested by Wächtershäuser, and selection is at the compartment level (Fernando and Rowe 2007, 2008). This idea is open to experimental test. Such systems are holistic replicators (Maynard Smith and Szathmáry 1999). If one looks at the core of the formose reaction (figure. 9.2), one sees that there is no real sense in which one could say that replication is “halfway through,” in sharp contrast to a piece of RNA or DNA. This is because replication here is not template replication (copying) that rests on a modular polymerization of monomers.




Figure 9.3

Scheme of simple modular self-replication (from Bag and Von Kiedrowski 1996). A and B, building blocks; T, template; C, catalytic complex; D, duplex. Note the reversible and irreversible reactions.

The first modular type of self-replicator (figure 9.3) was synthesized by Von Kiedrowski (1986). The palindromic arrangement of the template ensures that the copy will be identical to the template, despite complementary base pairing. There is now a large number of such experimentally produced replicators (for a review, see Von Kiedrowski 1999). A common criterion for the replication process is that the two strands (template and copy) must spontaneously separate. Since they are held together by hydrogen bonds (also necessary for replication), the strands cannot be too long; otherwise they will stick together for too long a time. Long pieces of nucleic acids can be replicated in the cell because enzymes of the replicase complex also ensure the unwinding of the strand—this cannot be assumed in nonenzymatic systems. These artificial replicators hence must generally be rather short. Although replication is modular, heredity is still limited because of size limitation. These replicators have only a didactic relevance to evolution since they are not feasible in prebiotic environments. Chemical evolutionists nevertheless do believe in prebiotically feasible counterparts.

The difficulty of long-strand replication has been analyzed in depth by Fernando et al. (2007). A novel stochastic model of nucleic acid chemistry was developed to allow rapid prototyping of chemical experiments

designed to discover sufficient conditions for template replication. Experiments using the model brought to attention a robust property of nucleic acid template populations, the tendency for elongation to outcompete replication. Externally imposed denaturation-renaturation cycles did not reverse this tendency. For example, it has been proposed that fast tidal cycling could establish a TCR (tidal chain reaction) analogous to a PCR (polymerase chain reaction) acting on nucleic acid polymers, allowing their self-replication (Lathe 2003). However, elongating side reactions that would have been prevented by the polymerase in the PCR still occurred in the simulation of the TCR. The same finding was found with temperature and monomer cycles. It is only the short nucleic acids (like the ones similar to the structure in figure 9.3) that replicate readily, until they are mopped up by ever-elongating long polymers that show no sign of replication under the investigated conditions. Thus the chemical origin of evolvability is still obscure. A possible way out could be the self-construction of active replicase enzymes from smaller parts (Hayden et al. 2008). Alternatively, an early replicase ribozyme may have been a restriction ribozyme capable of cutting itself out of elongating strands.

Unlimited heredity has arisen from limited heredity on multiple occasions: in the origin of life, in the adaptive immune system, and in language. The next section proposes that unlimited heredity also arose in the form of symbolic neuronal replicators inhabiting the brains of humans, and that this was a prerequisite for language.

Neuronal Replicators in Cognition

We propose that the scope of natural selection should be further widened to cognitive phenomena, in order to address the following open questions. How can the remarkable creativity of human thought and problem solving be explained, namely, how does cognitive search work? How does the brain solve the delayed reinforcement problem, that is, on what basis do we decide to behave, given that sometimes external rewards are far removed in time or even absent, and environments are non-Markovian? How does long-term memory formation and retrieval interact with working memory? How do we form causal models of the environment? How can perceptual and conceptual operations be applied to sensory data in a position-independent manner? How is language transmitted?

We hypothesize that natural selection occurs in the brain at rapid timescales (e.g., overnight), and that this contributes to ~~aspects of cogni-~~

~~tion involving thought~~, a process we equate with a kind of structured heuristic search. Furthermore, we propose that the human brain contains mechanisms for generating fitness functions based not only on the presence of extrinsic reward objects but also on internal value functions, such as the first derivative of predictability, causal coherence, and information-based criteria.

Just as fitness landscape terminology can be used to describe an evolutionary search, so it can also be used to describe a cognitive search through a problem space. We show how the *efficacy* of natural selection is increased by neuronal mechanisms that are capable of transforming a random search into a structured search, thus highlighting that there exists a continuum between blind variation and thought, the underlying process that structures thought being the evolution of evolvability in neuronal replicators (Toussaint 2003).

Many aspects of cognition appear to involve structured heuristic search rather than random search or pure hill-climbing; for example, insight and transformation problems (Chronicle et al. 2004) are in many respects more complex versions of problems such as the Stroop Task (Dehaene et al. 1998), the Wisconsin Card Sorting Task (Dehaene and Changeux 1991), and the Tower of London Task (Dehaene and Changeux 1997). These latter tasks involve choosing from a set of behaviors on the basis of positive and negative reward feedback. The search space of these tasks is typically small. More complex cognitive problems, in contrast (e.g., puzzles such as the nine-dot problem), have a larger search space that limits the effectiveness of exhaustive and random search. Also, they do not provide explicit feedback, making it necessary for the subject to depend on internally generated success criteria. We propose that the process of solving such problems involves natural selection of neuronal replicators in the brain. We suggest plausible mechanisms for neuronal replicators, neurophysiological evidence for these mechanisms, and which candidate behavioral tasks may be explained by natural selection. The characterization given here is a “law of qualitative structure” in that it describes a framework in which more detailed knowledge may potentially be obtained (Newell and Simon 1976).

The neuronal replicator hypothesis states that several types of replicator exist in the human brain: synaptic replicators, topological neuronal replicators (i.e. groups of synapses) (Fernando et al., 2008), and dynamical neuronal replicators. The hypothesis permits multiple realizability (Marr 1982), that is, it does not commit itself to a particular implementation of a neuronal replicator, although several are suggested here.

A crucial aspect of the extended evolutionary synthesis is a focus on the evolution of evolvability (G. B. Müller 2007; Pigliucci 2008). In neuronal replicators, simple yet powerful mechanisms exist to structure the exploration distribution of a neuronal replicator, that is, the phenotypic distribution of variants produced upon replication (Toussaint 2003). This permits a neuronal replicator system to modify a search based on the outcomes of previous searches, in a way that is not easily possible at the genetic replicator level. In addition, Lamarckian evolution takes place in neuronal replicators because neuronal changes due, for example, to reinforcement learning can be copied directly. Finally, neuronal topology copying has other functions that may be important in causal inference algorithms (Gopnik and Schulz 2004) for forming internal models (Craik 1943) and emulators (Grush 2004).

Precursors of the Neuronal Replicator Hypothesis

There is a long history of application of selectionist and Darwinian dynamics to intrabrain processes, from early thoughts on the evolution of ideas (James 1890 [1950]), to memetics (Dawkins 1976, 1982), neural selectionism (Changeux et al. 1973; Edelman 1987; Marr 1969; Young 1979), synaptic replicators (Adams 1998), synfire chains (Abeles 1991), and hexagonal replicators (Calvin 1996). These theories attempt to explain neural information processing, conscious thought, and human problem solving. All have been influential for the neuronal replicator hypothesis (see Fernando et al., 2008 for a full review). To this we may add the philosophical tradition of evolutionary epistemology, with some of its roots in Mach (1897, 1910), and the full exposition by Campbell (1974), Perkins (1995), Popper (1972), and Simonton (1995). An outline of these influences is given below, except for the philosophical part, which will be dealt with elsewhere (cf. Dennett 1981).

Psychologists such as William James were already writing about the selection of ideas (James 1890 [1950]). A century later, Monod discussed mutation and recombination as applied to ideas (Monod 1971). Dawkins revitalized the concept by introducing the “meme” (Dawkins 1976, 1982). The memetic paradigm has arguably remained in a poised state, at least for two reasons: first, it could not demonstrate a phenomenon that conclusively required memes; and second, it could not demonstrate a physical basis for memes (Aunger 2002). Typically, memes were not considered to replicate *within* a single brain; however, Aunger (2002) is a notable exception because he believes that historically intra-brain neu-

romeme replication would have preceded interbrain neuromeme transmission, a view to which we adhere.

Influenced by Donald Hebb's notion of neuronal assemblies (Hebb 1949), intrabrain selectionist mechanisms were proposed by Edelman and Changeux for neuronal groups (Changeux 1985; Changeux et al. 1973; Edelman 1987).¹ These mechanisms were intended to explain, among other things, executive function (Dehaene and Changeux 1997) and perceptual categorization (Edelman 1990). In Changeux's mechanism, the inherent dynamics of neural networks produces transient pre-representations, some of which get stabilized by resonance with perceptual inputs. Resonance is thought to arise in loops between cortex and thalamus. Dehaene and Changeux (1997) write that "in the absence of specific inputs, prefrontal clusters activate with a fringe of variability, implementing a 'generator of diversity.'" However, this idea appears to be poorly extendable due to one version of the curse of dimensionality (Belman 1957): If there is a large space to search, how can adaptive pre-representations be produced sufficiently rapidly? Changeux addresses this by allowing heuristics to act on the search through pre-representations; notably, he allows recombination between neuronal assemblies, writing that "this recombining activity would represent a 'generator of hypotheses,' a mechanism of diversification essential for the geneses of pre-representations." However, it is not clear how these heuristics are optimized. Changeux never mentions multiplication of pre-representations, that is, he does not consider pre-representations to be units of evolution. Recent models of neuronal networks exhibit structures and dynamics that correspond to Changeux's pre-representations, that is, polychronous groups (Izhikevich 2007); these are groups of neurons defined by the fact that they fire in temporal correspondence with each other. It is not clear how polychronous groups can be used to produce adaptive behavior.

Edelman (1987) proposed a theory similar to Changeux's in order to explain how an organism "decides" how to behave, given a set of sensory inputs and reward feedbacks as it interacts in a situated and embodied manner with an environment. To test his theory, Edelman has implemented what he calls a Darwinian system within the computer controller of a robot. In Edelman's theory, elements of a primary repertoire of neuronal groups within the brain are thought to compete with each other for stimulus and reward resources. This results in selection of a secondary repertoire of behaviorally proficient groups (Izhikevich et al. 2004). The theory has been the subject of a large number of models and vari-

ants that attempt to explain a wide range of behavioral and cognitive phenomena at various levels of abstraction (Gisiger et al. 2000), such as category formation (Edelman 1987), reinforcement learning using spike-time-dependent plasticity modulated by dopamine reward (Izhikevich 2007), and visual-motor control in a robotic brain-based device (Krichmer and Edelman 2005).

Again, Edelman's neuronal groups are not units of evolution, for there is no multiplication (Crick 1989, 1990). The most modern version of Edelman's neuronal group selection is presented by Izhikevich et al. (2004). In Izhikevich's model, neurons are arranged in a recurrent network with axonal conduction delays and weights being modified by spike-time-dependent plasticity (STDP). Polychronous groups (PCGs) with stereotypical temporal firing patterns self-organize when a particular firing set of neurons is activated in a spatiotemporal pattern, resulting in the convergence of spikes in downstream neurons. STDP subsequently reinforces these causal relationships. Because the same neuron can occur in many groups, and because delays produce an effectively infinite-sized dynamical system (subject to temporal resolution constraints), the number of PCGs far exceeds the number of neurons in the network, allowing a very high memory capacity for stored PCGs. In a group of 1,000 neurons, approximately 5,000 PCGs are found, with a distribution of sizes. If inputs are random, the PCGs are formed and lost transiently, over the course of minutes and hours. Structured input can stabilize certain PCGs. Izhikevich attempts to understand the functioning of polychronous groups within Edelman's framework of neuronal group selection (Izhikevich 2006). Michod explains that in neuronal group selection, synaptic change rules replace replication as a mechanism of variability of the "unit of selection": there is correlation between the parental and offspring states of the same neuronal group even without multiplication (Michod 1988). Here "parental" is used to mean the group at time t , and "offspring" refers to the same group at time $t + x$. However, no mechanism is described showing how a beneficial trait of one PCG could be transmitted to another PCG. We claim that this absence of multiplication is a fundamental limitation of Edelman's Neural Darwinism.

To summarize, because neither Edelman's nor Changeux's model includes multiplication, we argue that they lack the algorithmic capabilities of natural selection. The problem of transmission of a favorable trait from one group to nongroup material (or another group) is a ubiquitous bind in selectionist theories. Replication is the most natural way to envisage this transfer of function operation.

What advantage has Darwinism over selectionism? First, adaptations can be transmitted between structures. This allows search to be undertaken by modifying one's offspring and not oneself, an important factor if most variants are harmful. Second, replication allows good solutions to obtain more search resources. Third, replication allows selection to act on variability properties, that is, the evolution of evolvability (Pigliucci 2008), by hitchhiking of neutral variability-structuring neuronal changes (mutations) that are selected because they shape the exploration distribution of variants (Kirschner and Gerhart 1998; Toussaint 2003). This is because, if there is variation in variability, then selection can act on that variation (Wagner and Altenberg 1996).

Consider how the evolution of evolvability is possible. A minimal natural selection algorithm lacking the capacity for the evolution of evolvability (except by neutral drift) is a 1 + 1 evolutionary strategy (ES) in which a parent makes one mutated offspring (Beyer 2001). A 1+1 ES is so called because it involves a "population" of only one parent producing only one offspring. If the offspring is immediately superior to the parent, the offspring replaces the parent. If it is not, the offspring is destroyed and the parent makes another offspring. One iteration is a generation. Imagine two offspring are produced that have equal fitness, yet differ in the following feature only: the first offspring is capable of producing fit offspring itself, whereas the second offspring can produce only low-fitness offspring. The 1+1 ES cannot "see" the grandchildren's average fitness, and so both offspring will have equal likelihood of being selected. The minimal selective scenario required for selection for the more evolvable first offspring would have been a 1+1+1 ES in which the grandparent is potentially replaced after assessing the grandchildren, rather than replacing the parent after assessing the children. If a population of genotypes has nontrivial neutrality, that is, if there exist two offspring with the same fitness but different exploration distributions, selection on variability properties is possible.

James Baldwin already in 1898 was applying concepts of structured variability to how "valuable thought-variations" could be generated by a process of trial and error, as a basis of intelligence (Baldwin 1898, 1909). The Baldwin Effect is an example of how developmental exploration can structure intergenerational genetic variability, but clearly Baldwin was concerned with structuring variability in other domains. Nontrivial neutrality may be a necessary condition for open-ended evolution (Bedau 1998). This leads us back to the meme. No one disputes that social communication involves duplication of information, but

Robert Aunger asks, “are people the exclusive agents behind this process,” or “is an information-bearing replicator” acting as an agent?” (Aunger 2002). The concept of an agent is notoriously difficult to formalize (Barandiaran and Ruiz-Mirazo 2008), and we do not define it here, but one empirical finding is clear: entities that are known to possess agenthood are produced by the subset of units of evolution capable of nontrivial neutrality and whatever other features are necessary for these units to be capable of open-ended evolution. Self-referential units of evolution are distinct from simple units of evolution in being “reproducers” (Szathmáry and Maynard Smith 1997), that is, in specifying the means of their own production. Self-referential replication is one mechanism that produces the phenomenon of nontrivial neutrality. The memetic question can be reformulated: *Are there self-referential units of evolution that copy themselves between brains?* The best candidates may be linguistic replicators (Maynard Smith and Szathmáry 1999).

We suggest that self-referential replicators allow structured heuristic search, which is crucial in many aspects of cognition, and that the brain contains many such agents that compete and cooperate for extrinsic and intrinsic reward resources.

The idea of neuronal replicators has already arisen twice before (Adams 1998; Calvin 1996). The most promising suggestion is by Paul Adams, who describes synapses as replicating. This would make them the minimal neuronal units of evolution. Synapses replicate by increasing the amount of quantal release from the presynaptic to the postsynaptic neuron, according to a Hebbian rule. Mutations are noisy quantal Hebbian learning events where a synapse is made to contact a nearby postsynaptic neuron rather than to enhance the connection to the current postsynaptic neuron (Adams 1998). Synapses compete with each other for correlation resources and other reward resources if, for example, dopamine acts to modulate Hebbian learning. Adams demonstrates how arrays of synaptic weights can be selected, and how error-correction mechanisms in cortical layer VI can adjust the synaptic mutation rates. As described in Fernando et al. (2008), there is a mathematical isomorphism between Hebbian learning and Eigen’s replicator equations, a standard model of natural selection dynamics in chemical and ecological systems. Hebbian learning can be said to *select* between weights on the basis of correlations in activity. Synaptic weights can be said to multiply in proportion to the product of presynaptic and postsynaptic activity. The synaptic equivalent of mutation is a shifting of synaptic weight

resources between synapses afferent on the same postsynaptic neuron. Specifically, the Oja version of Hebbian learning is (Oja 1982):

$$\tau_w \frac{d\mathbf{w}}{dt} = v\mathbf{u} - \alpha v^2 \mathbf{w}, \quad (9.1)$$

where \mathbf{w} is the synaptic weight vector, v is the output rate, \mathbf{u} is the input rate vector, and the rest are constants. This is isomorphic to Eigen's replicator equation (Eigen 1971):

$$\frac{dx_i}{dt} = A_i Q_i x_i + \sum_{j \neq i} m_{ij} x_j - \frac{x_i}{c} \sum_{j=1}^N \sum_{k=1}^N m_{ij} x_j, \quad (9.2)$$

where x_i is the concentration of sequence i , m_{ij} is the mutation rate from sequence j to i , A_i is the gross replication rate of sequence i and Q_i is its copying fidelity, N is the total number of different sequences, and formally $m_{ii} = A_i Q_i$ (Eigen 1971).

The second proposal for neuronal replicators appeared in *The Cerebral Code* (Calvin 1996). Calvin outlined an algorithm for the copying of neural dynamics (spatiotemporal patterns) between hexagonally shaped regions of the cortex. Pyramidal cells in the superficial layers of the cortex have a circular field of excitatory efferents projecting to a standard length of 0.5mm from the central cell. Due to geometric constraints on a 2-D surface, this means that if six such cells are arranged in a hexagon and have the same receptive field, they form self-reexciting loops. If the cells behave as nonlinear relaxation oscillators (like fireflies), they can entrain each other and fire spikes in synchrony. The system may start with two cells firing in synchrony, which recruits a third cell to form an equilateral triangle, and so on. Due to long-term potentiation (LTP), the underlying synaptic weights would change to reinforce this synchrony in the future. A fundamental limitation of Calvin's proposal is that he does not explain how the circuitry between cells is to be replicated between hexagons.

Copying of neural receptive fields is already known to occur in the brain. The formation of topographic maps (from an initial random connection-weight matrix) by using spike-time-dependent plasticity (STDP) and lateral inhibition has been demonstrated (Song and Abbott 2001). This is effectively copying of the receptive fields in one layer onto a second, parallel layer. J. M. Young et al. (2007) have extended this work to explain the formation of new receptive fields in de-afferented regions of the cat's visual cortex. The extent of one-to-one topographic copying of receptive fields can be tuned by altering a neuronal gain parameter

in the lesioned area. At high gain, instead of obtaining a one-to-one copying of adjacent receptive fields (copying with no selective amplification and low information loss), the neurons in the entire de-afferented area receive the receptive field of just one of the adjacent neurons (copying with selective amplification and high information loss). Nonselective copying can be obtained not just with STDP but with Hebbian learning as well, with gain being adjusted according to activity-dependent scaling, for example (Van Rossum et al. 2000). Calvin's proposal does only half the job of copying, just as hydrogen bond formation in DNA replication does only half the job of semi-conservative replication; phosphodiester bond formation is required to re-create the original topology of the parent strand; similarly, in neuronal topology copying, an extra mechanism would be required, one that Calvin has not described.

Finally, concepts from synfire chains can be modified to include replication of neuronal dynamics. A synfire chain is a feed-forward network of neurons with several layers (or pools). Each neuron in one layer feeds many excitatory connections to neurons in the next pool, and each neuron in the receiving layer is excited by many neurons in the previous one. When activity in such a cascade of layers is arranged like a packet of spikes propagating synchronously from layer to layer, it is called a synfire chain (Abeles 1982, 1991). There have been reports in the literature about observations of precisely repeating firing patterns (Abeles and Gat 2001). An excellent summary is provided by Abeles et al.² One can see that almost trivially, replication of the spike packet is possible simply if the synfire chain branches into two or more. If we have branches like this, then one can use them for the spread of spike packets gated by reward. If we imagine a lattice where every arrow between the neuronal groups can work both ways, but in a reward-gated fashion, then fitter and fitter packets can fill up the lattice (Fernando et al., 2008). The snag, for the time being, is the limited heredity potential due to the limited information a spike packet can propagate. Recombination is easy to imagine when two roughly equally fit packets are transmitted to the same neuron group. The topology of the synfire network influences the outcome of selection. Neuronal evolutionary dynamics could turn out to be the best application field of evolutionary graph theory (Lieberman et al. 2005). It has been shown that some topologies speed up, whereas others retard, adaptive evolution. The brain could well influence the replacement topologies by gating, thereby realizing the most rewarding selection topologies (Fernando et al., 2008).

Mechanisms of Neuronal Replication

Neuronal Topology Replicators

A mechanism for copying of neuronal topology from one neuronal topographic layer to another is outlined. The mechanism utilizes topographic map formation (Song and Abbott 2001; Willshaw and Von der Malsburg 1976), coupled with spike-time-dependent plasticity (STDP; Markram et al. 1997), neuronal resetting (Crick and Mitchison 1995), topological error correction, and activity reverberation limitation. Neuronal topology copying works by using a neuronal implementation of a causal inference algorithm, in the sense that the offspring layer “observes” the activities of the parental layer in order to infer the connectivity of the parental layer.

Figure 9.4 shows a minimal example of neuronal topology copying (Fernando et al., 2008). It works by first establishing a topographic map between the parental and the offspring layers (Song and Abbott 2001; Willshaw and Von der Malsburg 1976). Spike-time-dependent plasticity (Markram et al. 1997) in the offspring layer is then used to infer the underlying topology of the parental layer, on the basis of activity received from the parental layer neurons as they are randomly sparsely activated.

To improve copying fidelity of neuronal topology, error-correction neurons are hypothesized that measure the difference in activity between corresponding neurons in parent and offspring layers (Adams and Cox 2002). On the basis of this discrepancy of activity, they modify the afferents to the offspring neuron accordingly. Two types of error-correction neurons are hypothesized, false-positive and false-negative error correctors. Activity reverberation limitation is required to allow the correct inference of Markov-equivalent causal graphs, and thus the copying of larger networks. This is implemented using inhibitory neurons that allow spikes to pass only if its associated excitatory neuron was principally depolarized from outside its layer.

The simplest way to demonstrate natural selection using the above copy operation is by using the 1+1 Evolutionary Strategy (1+1 ES; Beyer 2001) shown in figure 9.5. A 1+1 ES is a simple evolutionary algorithm that works as follows. If the offspring does not have fitness higher than the parent, then the offspring is erased and another attempt at copying the parent can be made (not shown). If the offspring has fitness higher than the parent, then the parent is erased and the offspring becomes the new parent and makes a new offspring in what was previously the parental layer (see parts 5, 6, 7, 8).




Figure 9.4

An outline of the neuronal topology replication mechanism with error correction. The parental layer is on the bottom. The offspring layer is on the top. In this example, each layer has three neurons. Topographic collaterals (vertical arrows) connect parent layer to offspring layer. Copying is the reproduction of the intralayer topology of the parent to the offspring. Error correction mechanisms are shown. STDP operates in the offspring layer. There are two error correction mechanisms; EC1 (right) is a false-positive error correction mechanism implemented using an “observer” neuron (EC1) that negatively neuromodulates neuron A' in the copy layer on the basis of differences in firing between the parental (A) and copy (A') layer neuron. We assume C is undergoing stimulation (1) when EC1 acts. EC2 (left) is a false-negative error correction mechanism implemented using an “observer” neuron that positively neuromodulate inputs that pass to a poorly firing neuron (C') in the copy layer from the neuron that is undergoing interventional stimulation (in this case we assume B is undergoing stimulation (2)) when EC2 acts. EC1- and EC2-type neurons are required for each neuron pair—A, A', B, B' and C, C'—and their neuromodulatory outputs must pass *widely* to all synapses in the child layer.

Note that neuromodulation is critical in the function of the 1+1 ES circuit in three ways. First, the direction of copy making depends on modulation to open and close vertical up and down gates at different times. Second, neuromodulation is necessary to switch on and off STDP-based plasticity in L0 and L1 alternately. Third, a mechanism is necessary to *reset* the layer (i.e., reduce weights in the layer) that is to be overwritten.

Figure 9.6 shows an example of an evolutionary run in which a 10-node neural network is evolved using the above algorithm. The replicators in the 1+1 ES are units of evolution as defined previously. This corresponds to a




Figure 9.5

A neurally implemented 1+1 ES using the STDP-based copying mechanism. (1) The circuit to be copied exists in the lower layer L0. The black connections in L0 show the original circuit. (2) Horizontal UP connections are activated (e.g., by opening neuromodulatory gating). These are the equivalent of the h-bonds in DNA copying. (3) A copy of the topology of L0 is made in L1, using STDP and error correction. (4) The layers are functionally separated by closing neuromodulatory gating of the UP connections. The fitness of each layer is assessed independently. (5) The layer with the lowest fitness is erased or reset (i.e., strong synaptic connections are reduced). In the above diagram we see that L1 fitness $>$ L0 fitness, so L0 experiences weight unlearning. (6) DOWN vertical connection gates are opened. (7) STDP in layer 0 copies the connections in L1. (8) After DOWN connections are closed, fitness is assessed and the cycle continues.

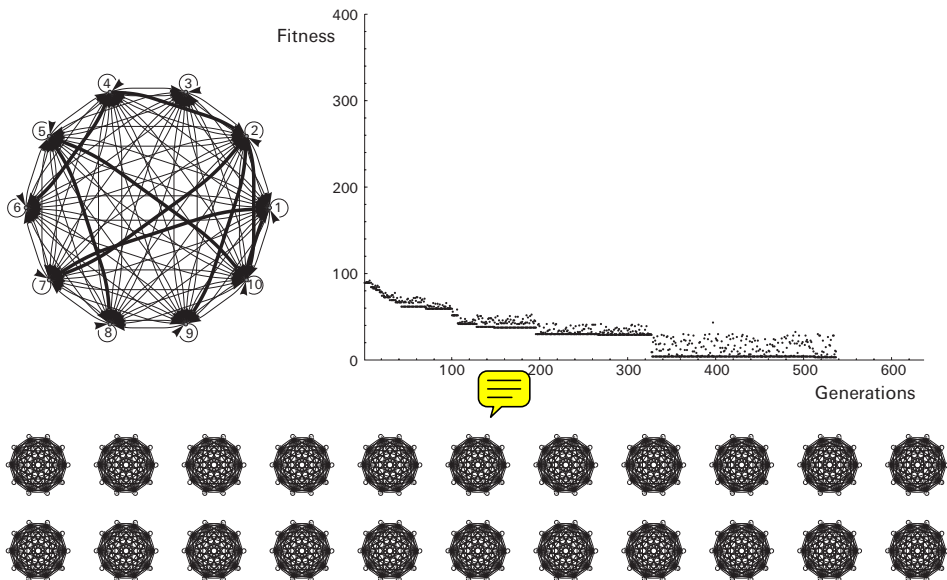


Figure 9.6

An example evolutionary run in which a particular desired topology is selected for. The neuronal copying algorithm is capable of sustaining evolution by natural selection to optimize the topology of a 10-node motif. Desired topology is on the top left (randomly initialized with 10% connectivity). Fitness (Euclidean distance between desired and actual topology) of parent (line) and offspring (dots) over 600 generations is shown on the top right. Bottom graphs show 11 parent–offspring pairs taken at intervals of 50 generations.

“genome size” of 90 synapses (self-synapses are not permitted). The network is selected for a particular topology of strong weights, the fitness being the Euclidean distance between the actual and the desired topology.

The network was initialized with all weak weights (fully connected; see bottom left). Explicit mutation operators were used after each copying event. A mutation involved the modification of a weight to be either very strong or very weak. In this implementation, the copying operation was often so effective that without an explicit mutation operator, there was insufficient variability for rapid evolution of the desired topology. The above model is capable of copying a sparsely connected neuronal topology of any size from one layer to another, if the two layers are connected topographically, and STDP and anisotropic reverberation limitation operates in the offspring layer. However, the rate of topology copying would be limited by the rate of synapse formation and resetting.³ In the next section, dynamical neuronal replicators are described that are capable of orders of magnitude faster evolution.

Dynamic Neuronal Replicators

The minimal unit of dynamical neuronal replication consists of a bidirectional coupled and gated pair of bi-stable neurons (Izhikevich 2007a), as in figure 9.7. We propose that dynamical replicators operate in working memory (Baddeley and Hitch 1974), for which bi-stability (Wang 1999) and recurrence (Zipser et al. 1993) have previously been proposed as mechanisms.

Grouping these pairs together, one can make two layers coupled initially by a topographic map. The parental layer has neurons initialized




Figure 9.7

An outline of how dynamical neuronal replicators can implement a 1+1 evolutionary strategy. The two bi-stable neurons are shown as large circles (light = not firing, dark = firing). They are coupled bidirectionally by gated axons. Each neuron can be reset and activated by external control. An array of such pairs is shown on the right. Initially the vector is all off (not spiking). Layer 0 is randomly initialized; here the neuron on the left becomes active. The up gates are opened, allowing that neuron to activate its corresponding neuron in layer 1. The gates are closed, and the fitness of each layer is assessed. The least fit layer is reset, and is overwritten by the fitter layer.

randomly as follows. If a bi-stable neuron is given some depolarizing current, it begins to fire repeatedly, whereas if a bi-stable neuron is given some hyperpolarizing current (in the correct phase), it stops firing. The state of neurons in the offspring layer is reset (i.e., all neurons are hyperpolarized to switch off spiking). Activity gates are opened for a brief period from the parental layer to the offspring layer, allowing the spiking neurons in the parental layer to switch on the corresponding neurons in the offspring layer. Activity gates between layers are then closed. Thus the vector of activities in the parental layer is copied to the offspring layer. As in figure 9.5, a 1+1 ES can be implemented if the two layers have their fitness assessed, the higher fitness layer is defined as the parent, and the offspring layer is reset. Figure 9.8a shows the result of selecting for a particular desired vector of activity using the above protocol.

How might dynamical and topological neuronal replicators interact? The consolidation of memories into long-term stores is known to involve the changing of synaptic weights in the hippocampus and cortex (Abraham and Robins 2005; Nadel et al. 2007; Nadel and Moscovitch

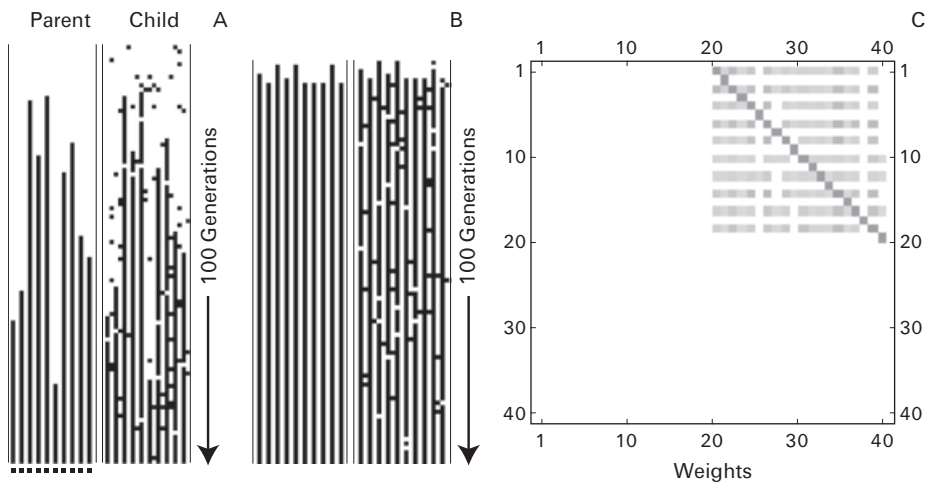


Figure 9.8

An example of dynamical replicators evolved to produce a desired activity vector with and without Hebbian learning. (A) shows that within 100 generations (each generation lasting 7 seconds), the desired activity vector can be selected for (bottom left sequence). (B) shows the much faster discovery of the desired activity vector when the copying operation is biased by previously found solutions using Hebbian learning. (C) shows the weight vector connecting parent to child that arises due to Hebbian learning. Note that the diagonal weights are the initially strong topographic weights. These are not adjusted by Hebbian learning.

1997; Nadel et al. 2000). The multiple-trace theory of memory consolidation suggests the conversion of solutions evolved in working memory, using dynamical replicators, into more permanent topological replicators that compete for synaptic resources.

Structuring Search Using Hebbian Learning

A remarkable capacity for structuring exploration distributions emerges if instead of limiting between-layer connections to a topographic map, one starts with a strong one-to-one topographic map and allows all-to-all Hebbian connections to develop once a local (or global) optimum has been reached. Imagine that evolution has been run to the end point of figure 9.8a, so that the optimum activity vector has been obtained and is present in both the parent layer and the child layer. Hebbian learning is then permitted between these two vectors (for all synapses except the original one-to-one topographic connections). If the activity vectors are then reset and a new evolutionary run is started, then copying will be biased by the Hebbian learning that took place in previous evolutionary runs. An active neuron in the parental layer will tend to activate not only the corresponding one-to-one topographic neuron, but also other neurons in the offspring layer that were previously active when the optimal solution had been found. Oja's rule is used to control the Hebbian between-layer synapses. Figure 9.8b shows that if Hebbian learning is permitted, then later evolutionary searches can converge faster, because they learn from previous evolutionary searches. The Hebbian weight vector that evolved is shown in figure 9.8c.

Richard Watson and colleagues in Southampton have described a set of search problems that are particularly well suited to neuronal copying biased by Hebbian learning (Watson 2006), and have proposed that "symbiotic evolution" can effectively solve these problems. ~~These are problems in which~~ there is interdependency between problem variables, that is, where the fitness contribution of one variable is contingent upon the state of other variables. An archetypical example is the Hierarchical IF-and-only-IF problem (HIFF; Watson et al. 1998), illustrated in figure 9.9.

The lowest level of fitness contributions comes from looking at adjacent pairs in the vector and applying the transfer function and the fitness function. The transfer function is $\{0,0\} \rightarrow 0$, $\{1,1\} \rightarrow 1$, and all other pair types produce a NULL (N). The fitness function for each level just sums the 0 and 1 entries. The second level is produced by applying the same




Figure 9.10

(Left) A 16-bit HIFF problem. Parental activity vectors are on the left and child activity vectors are on the right of each column of results. The number next to the column shows how many Hebbian learning and resetting loops have elapsed. Hebbian learning is permitted only in the last 10 generations of each loop (i.e., when it is most likely that a local optimum will have been found). After 50 generations, all initial conditions lead to one of the global optima (i.e., the all 1s are optimum). The final weight matrix is shown on the bottom left. (right) Control case, as before, except with no Hebbian learning. There is no improvement each time the genotype is reinitialized.


a stable state, given a perturbation to a large and highly connected network (Ashby 1960). Many tasks for testing prefrontal cortex function do not require a structured search. In fact, often they can be solved by random search. Consider Rougier et al.'s (2005) model of a classification task. The system receives reward if it classifies a set of feature vectors according to a particular element of that vector. The network is given a cue instructing which element it should base the classification upon. A slightly more complex version involves feature matching, that is, doing XOR (the non-linear exclusive OR logic function which returns 1 if given inputs [0,1] or [1,0] but returns 0 if given inputs [0,0] or [1,1]) on the relevant element pair in two feature vectors. Reward has the effect of gating a stochastic search, stabilizing the activities in prefrontal cortex if an unexpected reward is obtained, and destabilizing them if the system does not get an expected reward. Similarly, in temporal difference learning (Sutton and Barto 1998) the network undertakes a random search, biased

by instantaneous differences in predicted reward. A cue defines distinct environments in which a different behavior is rewarded. Rougier et al.'s task is fundamentally the same as the Stroop Task, in which the subject must either name the word "red" or name the color in which the word "red" is painted, or the Winonsin Card Sorting Task (WCST), in which the subject must sort the cards on the basis of a particular feature dimension. Although switching is rapid, it can be achieved by random (unstructured) search in the space of classification dimensions. The authors admit that the system does not deal with generativity: "Generativity also highlights the centrality of search processes to find and activate the appropriate combination of representations for a given task." They describe the mechanism as "random sampling with delayed replacement."

Similar approaches have been taken by Changeux's group (Dehaene et al. 1998) for the Stroop Task, but the language used to describe the process is different. They talk of stabilization (selection) of internal pre-representations by external stimuli. Configurations in a global workspace are stabilized by internal reward and attention signals (Dehaene et al. 1998). Chialvo and Bak's paper "Learning from Mistakes" (Chialvo and Bak 1990) and variants (Bosman et al. 2004) also describe a similar stochastic hill-climbing type of algorithm. Seung has explicitly described the process of reward-biased search in spiking neurons as stochastic hill climbing (Seung 2003).

Selective attention in complex tasks (Desimone and Duncan 1995) and solving insight problems (Chronicle et al. 2004) may require a mechanism for structuring search. This is because for many problems it would take too long to exhaustively search all possible solutions, and hill climbing may be inefficient due to local optima. Rougier et al. write that "an important feature of future research will be to identify neural mechanisms that implement more sophisticated forms of search" (Rougier et al. 2005).

We do not yet know what neural mechanisms underlie human creativity in problems such as the Nine-Dot Problem: "Draw four continuous straight lines, connecting all the dots without lifting your pencil from the paper" (MacGregor et al. 2001).



The difficulty of finding the right solution depends upon the representation of solutions, the variability operators, and the selection algorithm.

It is likely that for most representations of the Nine-Dot Problem the fitness landscape is rugged (Perkins 1995). That is, move operators will result in getting stuck on local minima. Some problem representations may be so poor as not to contain the correct solution. Several authors describe that insight problems require “restructuring of the initial problem representation” (Chronicle et al. 2004) or sculpting the response space (Frith 2000) to encompass the goal state (Öllinger et al. 2006) because of strong constraints that “prevent one from considering and evaluating the correct solutions” (Reverberi et al. 2005; Knoblich et al. 2005). For a problem to be an insight problem, performance cannot be explained by innate, random, or exhaustive search; for example, it has not been conclusively demonstrated that New Caledonian crows use insight when bending wires to lift peanuts in buckets out of bottles. Simpler search methods may be sufficient for this task (Weir et al. 2002).

Most people’s phenomenal concomitants to solving this problem are something like undertaking trial-and-error search with constraint relaxation, that is, increasing the space of possible solutions, as well as hill climbing to improve some intermediate goal function. However, one must be careful not to assume that the brain works in the same way as the mind thinks. “The ghost has been chased further back into the machine, but it has not been exorcised” (Fodor 1983: 127). The moment of insight is instantaneous, and the correct solution does not seem to be arrived at consciously (Sternberg and Davidson 1995; Metcalfe and Wiebe 1987). What neuronal process may underlie performance in these tasks (Kaplan and Simon 1990; Simon and Reed 1976)?

It is proposed that a role for neuronal natural selection is to structure exploration distributions in order to bias cognitive search. An exploration distribution is defined by both the representation of solution space *and* the move set. The example of the HIFF problem shows how the exploration distribution of a dynamical neuronal replicator can be structured by Hebbian learning. Hebbian learning alters only the move set, not the underlying representation of solution space, not the selection method, and not the fitness function. The neuronal replicator hypothesis proposes that the brain configures all these processes to implement an effective natural selection algorithm for a given search problem. The frame problem (Pylyshyn 1987) then becomes the problem of how to choose an initial population of solutions. An example is provided below of an idealized method by which a neuronal natural selection algorithm could solve an “insight” problem.

Consider how to apply the neuronal replicator approach to the 10-coin problem; 10 coins are arranged in a triangle, and the aim is to make the triangle face the opposite direction by moving only three coins.

Figure 9.11 shows a possible representation of this problem. Assume that when the problem is defined to the subject verbally, the subject's brain is capable of representing the positions of the coins as activations of an array of neurons. The figure shows a simple representation consisting of a square neuronal array with coin positions shown. This representation is capable of emulating various properties of the real coins, such as their relative spatial location, and constraints about how they can be arranged. Assume that there is a population of at least two solutions. A solution is a probabilistic neuronal description of three coin moves, implemented, for example, as an array of neurons with different firing rates. When a solution is applied to the coin emulation, the emulation is transformed into the final state. Another neuronal system is then capable of assigning a subjective utility to this final state. The assignment of subjective utility is not trivial, and depends on previous experience. Two subjective utility functions are used here. In the first, a simplifying assumption is made that subjective utility corresponds to the Hamming distance from the correct solution. In the second, subjective utility is the maximum number of coins corresponding to an inverted triangle convolved over the final state of coins after three moves. Performance using both functions was very similar. A potential solution consists of six probabilistic transformation matrices representing which coin is to be moved and where it is to be moved. Other potential solution representations may arise that are more or less likely to be fit; however, we do not consider the mechanism by which an initial population of neuronal replicators is first formed. A mechanism applies a solution to the emulator's representation of the coin positions by choosing a valid coin to move, and making sure it moves to a valid place. Because the solutions are initialized randomly according to a uniform distribution, many different coin moves can be produced by the same solution. Therefore, to accurately determine the fitness of a solution, the subjective utility of many final states arising from many (in this case 1,000) applications of the rule must be summed (Cooper 2001). Over time, the solution representation becomes more deterministic and fewer runs through the emulator are required to obtain a fitness assessment. Figure 9.11 shows the above algorithm implemented with a population size of 2. At each round, the fitness of both solutions is calculated, and the fitter solution replaces the less fit solution, with mutation. Mutation is defined as choosing each site in the

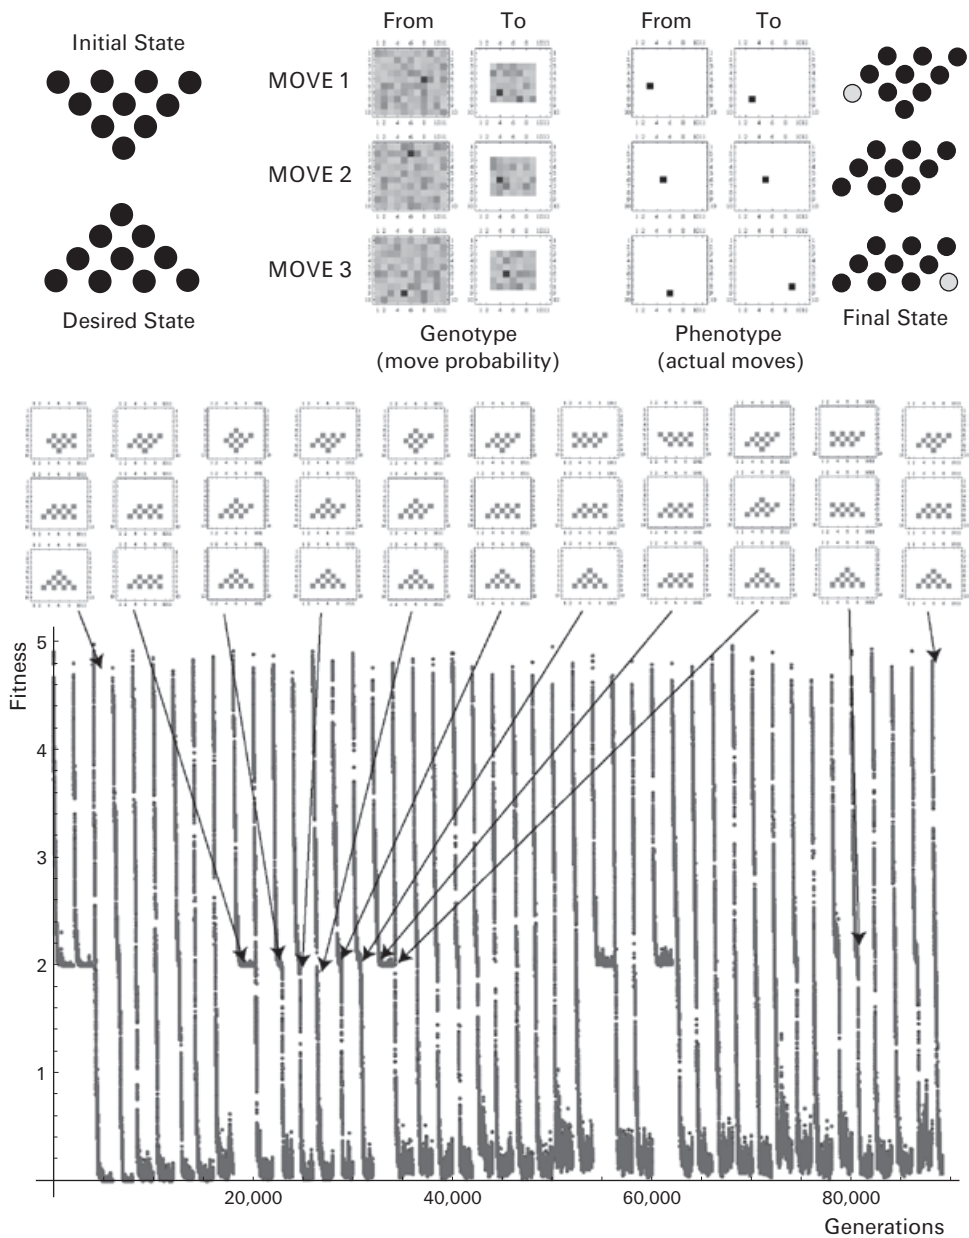


Figure 9.11

(Top left) The initial and the desired states of the 10 coins. (Top right) An example of the genotype, the chosen move, and the emulated coin positions for 3 moves that result in a partial solution. (Middle and bottom) An evolutionary run consisting of reinitialization of the 2 genotypes every 3000 generations. Hebbian learning is not used between generations. Most runs result in the global optimum being reached. Some runs get stuck with an error of 2 (i.e., one coin in the wrong place). The intermediate and final coin positions are shown in the boxes above the fitness plot.

transformation matrices with a probability of 1/10 and either doubling it or halving it, and then renormalizing that matrix. Sometimes the system can find the correct solution by a process of hill climbing, but sometimes it gets stuck in local optima. The local optima occur due to null neutral moves (i.e., where the initial and final positions of a coin are the same in one of the moves) and due to circular neutral moves (i.e., where a coin contributing to fitness in the correct inverted triangle is moved to another place where it provides equal fitness to the inverted triangle).

The ability to solve the problem by hill climbing alone calls into question whether this problem (as formulated above) is really an insight problem. The model presented says nothing about how such goal functions are chosen and modified. A more realistic model of human performance would include an algorithm for generating predicted goal states. The Hebbian learning mechanism is not required to find the solution to the 10-coins problem here; however, it can *save* the solution and allow this saved solution to bias mutation on another round of natural selection from random initial conditions.⁴ This process of Hebbian learning of local optima with subsequent biasing of exploration distributions may be one of the means by which associations can be learned and applied (Simonton 1995).

An objection to the neuronal plausibility of the above stochastic hill-climbing type of algorithms may be that many runs through an emulator were required in order to obtain an accurate fitness assessment of a rule (e.g., 1,000 emulations per rule). Clearly, such a large number of applications of a probabilistic rule are not made consciously. In fact, consciously we are only aware of testing deterministic rules a relatively small number of times (e.g., we move a coin physically, then see what we can do next). Failure to find the correct solution may arise due to various limitations imposed upon the above algorithms; for example, emulation of a rule may be possible only consciously or by use of real coin moves in the external environment, in which case far fewer assessments of a probabilistic rule can be carried out. In this circumstance, fitness assessment of a probabilistic rule may be so noisy that no increase in population fitness is possible from a random initial probabilistic rule initialization. Where fewer emulations are allowed, a rule must either be defined deterministically, or the fitness function must be redesigned to decrease the variance in the outcomes of emulations.

This method of solving the coin problem differs from those discussed in Chronicle et al. (2004). There it is assumed that the only configurations tested are those that are physically moved in the real world. Here it is

assumed that internal emulation of the effect of coin movements takes place unconsciously in the brain when one is told the problem. Evidence for unconscious mechanisms in insight tasks comes from an experiment showing that sleep doubles the rate at which explicit knowledge is gained about a hidden correlation in a stimulus-response task (Wagner et al. 2004). In reality, a complex interaction between what are described as conscious and unconscious processes is likely. In addition, people may vary greatly in their capacity for emulation.

The heuristic search hypothesis of Newell and Simon (1976) assumes that “solutions to problems are represented as symbol structures. A physical symbol system exercises its intelligence in problem solving by search—that is, by generating and progressively modifying symbol structures until it produces a solution structure.” The physical symbol system hypothesis has been heavily criticized (Brooks 1990). But we note that there appears to be a parallel in the requirement for “symbols” both in the brain and in organismal natural selection, as became evident in the problem of how to maintain information and diversity by blending inheritance (Gould 2002: 622). The only solution was to allow symbolic (i.e., particulate Mendelian) inheritance (Fisher 1930). Symbols seem to be a crucial requirement for natural selection with unlimited heredity, irrespective of its implementation.

How Is Fitness Defined in the Brain?

How does the brain determine the criteria by which to select neuronal replicators? Midbrain dopamine systems signal the error in predicted reward (Izhikevich 2007b; Oudeyer et al. 2007); however, dopamine also signals more complex value functions, such as prediction error (Horvitz 2000), and some have claimed that the brain tries to minimize prediction error (Friston and Stephan 2007). But then why seek novelty, why play? Oudeyer et al. (2007) propose that the brain attempts to maximize learning progress (i.e., to produce continual increases in predictivity). Other formulations of intrinsic fitness functions exist (Lungarella and Sprons 2005), such as maximization of information flow in the sensorimotor loop (Klyubin et al. 2007), and maximization of mutual information between the future and the past (Bialek et al. 2001), and these have been tested in robotic control (Der et al. 2008). Alternatively, co-evolutionary approaches suggest that separate populations of neuronal replicators in the brain may have different functions, some units operating as fitness functions for other units that act as solutions or as predictors or as perceptual agents (De Jong and

Pollack 2003). In an almost evolutionary system, the function being maximized in Copycat (a program for solving analogy-based insight problems) is the activation of a conceptual network (Hofstadter and Mitchell 1995), which is effectively a neuronal network with neuromodulation (Sporns and Alexander 2002) and gating (Steriade and Pare 2007).

We return now to Aunger's claim that neuronal replicators (neuro-memes) escaped the confines of the single brain and became capable of being copied between brains (Aunger 2002).

Language and Neuronal Replicators

Prions are an example of molecular *phenotypic* replicators. Prions can have alternative conformations; molecules with bad conformation (phenotypes) transform peptides with the right conformation into ones with bad conformation (Mestel 1996). There is a direct phenotype-to-phenotype transmission, without modular copying of constituents, which is in sharp contrast to the case of RNA, for which the phenotypes are correlated because the parental sequence is replicated. Interestingly, further evolution of this initially purely selfish system has been co-opted by yeast where it transmits a certain phenotypic trait to the read-through of all three nonsense codons (Patino et al. 1996). The sequences of prions are coded for by genes. It is intuitively clear that such molecular replicators can exist in a few alternative states only; hence they belong to the class of limited hereditary replicators.

Can there be phenotypic replicators with unlimited heredity? Memes are proposed within this category (Dawkins 1976), although the fact that they are typically phenotypic replicators was recognized only recently (Maynard Smith and Szathmáry 1999). Consider, for example, the neuronal correlates of an understanding of Newton's Second Law. When teachers teach it to their students, there is no copying involved whatsoever. Copying would require the transmission of the synaptic configuration of the neural network storing the piece of information in question. There are reasons to believe that such a copying would produce no meaningful result. Instead, the emerging hypotheses in the students are tested according to performance (phenotype), until performance in students and teacher is sufficiently similar (figure 9.12).

Why are we able to sustain an indefinitely large number of memes? We think the answer is human language. Language is a cultural inheritance system with indefinitely large semantic coverage (Maynard Smith and Szathmáry 1995). It is also digital, since an indefinitely large number

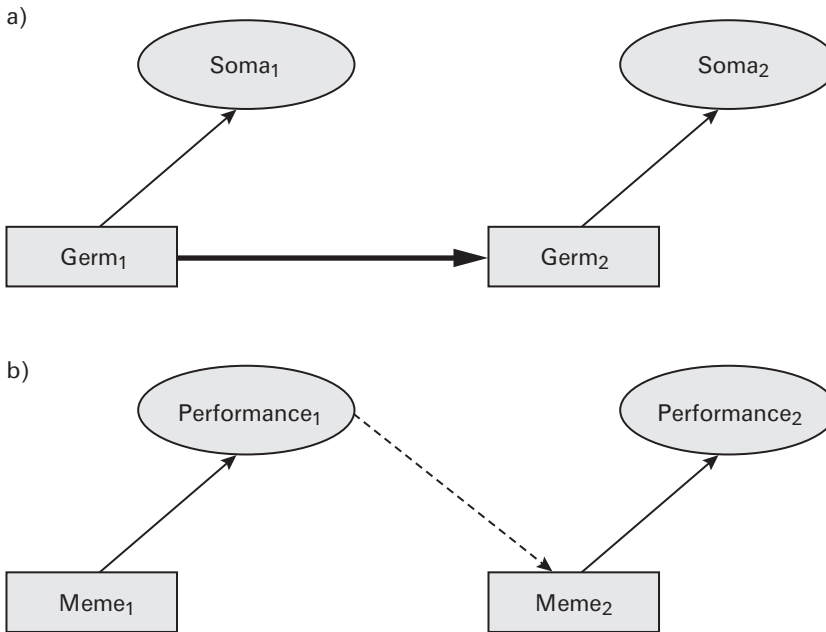


Figure 9.12

Memos and Lamarckian inheritance. (a) The Weissmanist segregation of soma and germ line. (b) Transfer of memes passes through the performance level, which is mostly absent in the molecular world (Szathmáry 2002).

of sentences can be generated with a limited alphabet. Although it is true that alphabets were superimposed on languages long after their invention, the number of basic phonemes we use in any language is a finite set. We construct words using this set. Even the number of words does not exceed 10^5 .

It is relevant to ask whether the neuronal structures underlying language are true replicators or not. Given what we wrote above about neuronal replicators, we believe the answer is affirmative. Memos are genotypic replicators inside brains and phenotypic replicators between brains.

Accepting that memos within brains are also digital replicators, whose replication from brain to brain is phenotypic, could one think of a didactic molecular analogy? Here is one (Szathmáry 2000). There are two individuals, A and B. Take protein X from individual A. Suppose you want to enable individual B to develop a molecule with the same phenotypic effect (enzymatic function, for example). If gene transfer from A to B is not allowed, one then must have (1) some generative mecha-

nism for proteins in B, and (2) some method for the assessment of phenotype. This comes very close to an immune system in B. The crucial difference is that the task now is to produce “antibody” Y, in individual B, that shares crucial phenotypic properties with “antigen” X, from individual A. Although both molecules would have sequences, it is most unlikely that they would be close to one another in protein space (cf. Maynard Smith 1970). In all probability the pleiotropic effects of the two proteins would differ. This is why cultural heredity is bound to be inexact and why cultural evolution is faster than biological evolution. By analogy, the transmission of language may involve the phenotypic transmission of linguistic constructions, inferred by a process of natural selection within the brain of the receiver.

Indeed, there is evidence that compositionality in language may be an adaptation for phenotypic transmissibility. Henry Brighton in his Ph.D. thesis (2003) describes an evolutionary model that was intended to demonstrate that languages were constrained by the fact that they had to be copyable. The basic task involves a conceptual space of a “speaker” modeled as a 2-D surface, onto which 100 randomly positioned points are assigned. Each point is classified into one of five classes. This constitutes the initial “concept” space. Then, in series, a further 100 unclassified points are presented to the speaker. The speaker classifies each point based on the class of the already existing point that is closest to the new point. The “listener” has access *only* to the 100 newly classified points, which are painted onto the listener’s 2-D space. The listener then becomes the speaker, and the process iterates. The question is, how well can the pattern of classes over the 2-D surface be transmitted?

Because the sampling of the class space is stochastic, because the class distribution of learners is biased, and because new classes cannot be reintroduced once lost, eventually, after many generations, only one class of points exists, having taken over the whole surface. Heredity of more than one class type is not possible with this algorithm. The system has reached one of five global asymptotic stable points. How can the system be modified to allow many classes to be stably transmitted?

An extension of the above experiment was devised to test how a *mapping* from one 2-D space (M) to another 2-D space (S) could be inherited. The first speaker starts by being given a random one-to-one map from 100 points in M to 100 points in S. As before, 100 new points in M are presented to the speaker. The speaker decides what to say by finding the three closest points in M to the new point, which leads him to the corresponding three points in S. The speaker then uses some

interpolation of the positions of S_1 , S_2 , and S_3 to produce a new point, s . It is these 100 new s points that are heard by the listener, who uses them to construct her set of 100 points in M , and the process iterates.

Brighton defines compositionality as the degree of correlation between distance between points in M and the distances between the corresponding points in S . A compositional map is a topographic map with $c = 1$, and a random $M \rightarrow S$ map has $c = 0$. For cases with the system started with a random map and with a fully compositional map, the above algorithm could sustain only weak compositionality, $c = 0.3$, with agents often having different $M \rightarrow S$ maps between generations. Thus a mapping could not be stably inherited using the above algorithm. This was shown to be because there was no error-correcting mechanism acting to counter small rounding errors in the interpolation calculation.

To solve this problem, Brighton introduced the *obverter procedure*, originally also devised to study the evolution of communication. The speaker uses “introspection,” that is, when trying to produce a signal to represent a point in M , the speaker works out which signal, if received by itself, would maximize the probability of its inferring the correct point in M . The speaker is thus using himself as a model of the listener. An alternative method could be to use other techniques to infer the $M \rightarrow S$ mapping of the listener. By doing this (and by using new points to improve the classification of subsequent new points, “production memory”), a highly compositional ($c = 0.92$) and highly stable mapping was achieved.

What is the relevance of these dynamics to phenotypic copying? A simple “autistic” phenotypic transmission algorithm may be insufficient; a predictive mechanism may be necessary to ensure that stable and unlimited $M \rightarrow S$ maps are copied accurately between brains.

Conclusion

In this chapter we have ventured into the domain of bona fide neuronal replication, in the hope of extending the synthesis even further in breadth. We are happy to point out that while doing so, we are standing on the shoulders of Maynard Smith, Changeux, and Edelman. Without the influence of these scholars we would have been unlikely to arrive even at an initial research program on neuronal evolution proper. We also hope, at least in the long run, to be able to contribute ~~also~~ to the depth of the theory with our research program. From a general point of view it is likely that this program will shed new light on evolvability,

exploration distributions in evolution, and the interplay of Lamarckian and Darwinian mechanisms, but strictly within the constraints of genetically based evolution of the component brain mechanisms. We do not know what all this will deliver. Maybe the particular neuroevolutionary mechanisms presented in this chapter will not survive, but the general research program as such will, and in the future some neurobiologists will also be evolutionary biologists at the same time.

The replication mechanisms are proposed to exist at the neuronal level and not at the level of thought. The existence of noise is a major constraint to neuronal copying of all kinds. Just as in the study of the origin of life, it is a nontrivial question to understand how unlimited heredity could have arisen in the brain when noise is taken into account (Szathmáry and Maynard Smith 1997). However, before the question of origins is addressed, it remains to be seen whether empirical evidence can be found for topological or dynamical neuronal replicators.

Acknowledgments

Thanks to a Career Development Fellowship at the MRC National Institute for Medical Research, Mill Hill, London, and to a Marie Curie Inter-European Grant to work at Collegium Budapest, Hungary. Partial support of this work has generously been provided by the Hungarian National Office for Research and Technology (NAP 2005/KCKHA005) and by the Regional Knowledge Centre of Eötvös University (National Office for Research and Technology grant no. RET2.4/2005). The work has also been supported by the COST Action CM0703 on systems chemistry. Thanks to Luc Steels, Richard Watson, Richard Goldstein, James McDonald, Anil Seth, and Zoltán Szatmáry for useful discussions.

Notes

1. Dawkins (1971) also proposed neuronal elimination as a basis of memory.
2. (http://www.scholarpedia.org/article/Synfire_chain).
3. An interesting parallel development is von Kiedrowski's (Von Kiedrowski et al. 2003) "connectivity copying" mechanism whereby selectable nano-robots can be constructed in the chemical domain.
4. Introducing Hebbian learning to the copy operator in a manner analogous to the previous approach with the HIFF problem allows evolution to structure the exploration distribution of variants based on the structure of solutions found in previously explored optima. A 660×660 Hebbian copying matrix is updated to represent the within genotype correlations after an optimum has been reached (a genotype is 660 units in length) and is used to bias the copy operation in subsequent evolution. However, since the system can find a global optimum without Hebbian learning, this system is largely redundant (see figure 9.12).

References

- Abeles M (1982) *Local Cortical Circuits: An Electrophysiological Study*. Berlin: Springer.
- Abeles M (1991) *Corticonics: Neural Circuits of the Cerebral Cortex*. New York: Cambridge University Press.
- Abeles M, Gat I (2001) Detecting precise firing sequences in experimental data. *Journal of Neuroscience Methods* 107: 141–154.
- Abraham WC, Robins A (2005) Memory retention—the synaptic stability versus plasticity dilemma. *Trends in Neurosciences* 28: 73–78.
- Adams P (1998) Hebb and Darwin. *Journal of Theoretical Biology* 195: 419–438.
- Adams PR, Cox KJA (2002) A new interpretation of thalamocortical circuitry. *Philosophical Transactions of the Royal Society of London B* 357: 1767–1779.
- Ashby WR (1960) *Design for a Brain: The Origins of Adaptive Behavior*. New York: Wiley.
- Aunger R (2002) *The Electric Meme: A New Theory of How We Think*. New York: The Free Press.
- Baddeley AD, Hitch GJ (1974) Working memory. In *The Psychology of Learning and Motivation: Advances in Research and Theory*. GH Bower, ed.: 8: 47–89. New York: Academic Press.
- Bag BG, Von Kiedrowski G (1996) Templates, autocatalysis and molecular replication. *Pure and Applied Chemistry* 68: 2145–2152.
- Baldwin JM (1898) On selective thinking. *The Psychological Review* 5: 4.
- Baldwin JM (1909) The influence of Darwin on theory of knowledge and philosophy. *Psychological Review* 16: 207–218.
- Barandiaran X, Ruiz-Mirazo K (2008) Introduction to the special issue: Modelling autonomy. *Biosystems* 91: 295–304.
- Bedau MA (1998) Four puzzles about life. *Artificial Life* 4: 125–140.
- Belman RE (1957) *Dynamic Programming*. Princeton, NJ: Princeton University Press.
- Beyer H-G (2001) *The Theory of Evolution Strategies*. Berlin: Springer.
- Bialek W, Nemenman I, Tishby N (2001) Predictability, complexity and learning. *Neural Computation* 13: 2409–2463.
- Bosman RJC, Van Leeuwen WA, Wemmerhove B (2004) Combining Hebbian and reinforcement learning in a minibrain model. *Neural Networks* 17: 29–36.
- Boyd R, Richerson PJ (2005) *The Origin and Evolution of Cultures*. Oxford: Oxford University Press.
- Brighton H (2003) *Simplicity as a driving force in linguistic evolution*. Ph.D. thesis, University of Edinburgh.
- Brooks RA (1990) Elephants don't play chess. *Robotics and Autonomous Systems* 6: 3–15.
- Calvin WH (1996) *The Cerebral Code: Thinking a Thought in the Mosaics of the Mind*. Cambridge, MA: MIT Press.
- Campbell DT (1974) Evolutionary epistemology. In *The Philosophy of Karl Popper*. PA Schillpp, ed.: 412–463. LaSalle, IL: Open Court.
- Changeux J-P (1985) *Neuronal Man: The Biology of Mind*. Princeton, NJ: Princeton University Press.
- Changeux J-P, Courrègne P, Danchin A (1973) A theory of the epigenesis of neuronal networks by selective stabilization of synapses. *Proceedings of the National Academy of Sciences of the USA* 70: 2974–2978.

- Chialvo DR, Bak P (1990) Learning from mistakes. *Neuroscience* 90: 1137–1148.
- Chronicle EP, MacGregor JN, Ormerod TC (2004) What makes an insight problem? The roles of heuristics, goal conception, and solution recording in knowledge-lean problems. *Journal of Experimental Psychology: Learning, Memory, and Cognition* 30: 14–27.
- Cooper WS (2001) *The Evolution of Reason: Logic as a Branch of Biology*. Cambridge: Cambridge University Press.
- Craik K (1943) *The Nature of Explanation*. Cambridge: Cambridge University Press.
- Crick F, Mitchison G (1995) REM sleep and neural nets. *Behavioral and Brain Research* 69: 147–155.
- Crick FHC (1989) Neural Edelmanism. *Trends in Neurosciences* 12: 240–248.
- Crick FHC (1990) Reply. *Trends in Neurosciences* 13: 13–14.
- Dawkins R (1971) Selective neurone death as a possible memory mechanism. *Nature* 229: 118–119.
- Dawkins R (1976) *The Selfish Gene*. Oxford: Oxford University Press.
- Dawkins R (1982) *The Extended Phenotype: The Gene as the Unit of Selection*. Oxford: W.H. Freeman.
- De Jong ED, Pollack JB (2003) Learning the ideal evaluation function. *Genetic and Evolutionary Competition* 2003: 203. No. 2723 of LNCS.
- Dehaene S, Changeux J-P (1991) The Wisconsin Card Sorting Test: Theoretical analysis and modeling in a neuronal network. *Cerebral Cortex* 1: 62–79.
- Dehaene S, Changeux J-P (1997) A hierarchical neuronal network for planning behavior. *Proceedings of the National Academy of Sciences of the USA* 94: 13293–13298.
- Dehaene S, Kerszberg M, Changeux J-P (1998) A neuronal model of a global workspace in effortful cognitive tasks. *Proceedings of the National Academy of Sciences of the USA* 95: 14529–14534.
- Dennett DC (1981) *Brainstorms*. Cambridge, MA: MIT Press.
- Dennett DC (1995) *Darwin's Dangerous Idea*. New York: Simon & Schuster.
- Der R, Güttler F, Ay N (2008) Predictive information and emergent cooperativity in a chain of mobile robots. In *Proceedings of Alife XI*: 166–172. Cambridge, MA: MIT Press.
- Desimone R, Duncan J (1995) Neural mechanisms of selective visual attention. *Annual Review of Neurosciences* 18: 193–222.
- Edelman GM (1987) *Neural Darwinism: The Theory of Neuronal Group Selection*. New York: Basic Books.
- Edelman GM (1990) *The Remembered Present: A Biological Theory of Consciousness*. New York: Basic Books.
- Edelman GM (1994) The evolution of somatic selection: The antibody tale. *Genetics* 138: 975–981.
- Eigen M (1971) Selforganization of matter and the evolution of biological macromolecules. *Naturwissenschaften* 58: 465–523.
- Fernando C, Rowe J (2007) Natural selection in chemical evolution. *Journal of Theoretical Biology* 247: 152–167.
- Fernando C, Rowe J (2008) The origin of autonomous agents by natural selection. *Biosystems* 91: 355–373.
- Fernando C, Santos M, Szathmáry E (2005). Evolutionary potential and requirements for minimal protocells. *Topics in Current Chemistry* 259: 167–211.
- Fernando C, Karishma, KK, Szathmáry E (2008) Copying of neuronal topology by spike-time dependent plasticity and error-correction. *PLoS ONE* 3: e3775 [↓](#)



- Fernando C, Von Kiedrowski G, Szathmáry E (2007) A stochastic model of nonenzymatic nucleic acid replication: "Elongators" sequester replicators. *Journal of Molecular Evolution* 64: 572–585.
- Fisher RA (1930) *The Genetical Theory of Natural Selection*. Oxford: Clarendon Press.
- Fodor JA (1983) *The Modularity of Mind*. Cambridge, MA: MIT Press.
- Friston KJ, Stephan KE (2007) Free-energy and the brain. *Synthese* 159: 417–458.
- Frith CD (2000) The role of dorsolateral prefrontal cortex in the selection of action, as revealed by functional imaging. In *Control of Cognitive Processes: Attention and Performance*. S Monsell, J Driver, eds. Cambridge, MA: MIT Press.
- Ganti T (1971) *The Principles of Life* (in Hungarian). Budapest: Gondolat. (English translation 1986. Budapest. OMIKK)
- Ganti T (2003) *The Principles of Life*. Oxford: Oxford University Press.
- Gisiger T, Dehaene S, Changeux J-P (2000) Computational models of association cortex. *Current Opinion in Neurobiology* 10: 250–259.
- Gopnik A, Schulz L (2004). Mechanisms of theory formation in young children. *Trends in Cognitive Sciences* 8: 371–377.
- Gould SJ (2002) *The Structure of Evolutionary Theory*. Cambridge, MA: Belknap Press of Harvard University Press.
- Grush R (2004) The emulation theory of representation: Motor control, imagery, and perception. *Behavioral and Brain Sciences* 27: 377–442.
- Hayden EJ, Von Kiedrowski G, Lehman N (2008) Systems chemistry on ribozyme self-reconstruction: Evidence for anabolic autocatalysis in a recombination network. *Angewandte Chemie (international ed.)* 47: 8424–8428.
- Hebb DO (1949) *The Organization of Behavior: A Neuropsychological Theory*. New York: John Wiley.
- Hofstadter DR, Mitchell M (1995) The Copycat Project: A model of mental fluidity and analogy-making. In *Fluid Concepts and Creative Analogies: Computer Models of the Fundamental Mechanisms of Thought*. DR Hofstadter and the Fluid Analogies Research Group, eds.: 205–267. New York: Basic Books.
- Horvitz J-C (2000) Mesolimbocortical and nigrostriatal dopamine responses to salient non-reward events. *Neuroscience* 96: 651–656.
- Izhikevich EM (2006) Polychronization: Computation with spikes. *Neural Computation* 18: 245–282.
- Izhikevich EM (2007a) *Dynamical Systems in Neuroscience: The Geometry of Excitability and Bursting*. Cambridge, MA: MIT Press.
- Izhikevich EM (2007b) Solving the Distal Reward Problem through linkage of STDP and dopamine signaling. *Cerebral Cortex* 17: 2443–2452.
- Izhikevich EM, Gally JA, Edelman GM (2004) Spike-timing dynamics of neuronal groups. *Cerebral Cortex* 14: 933–944.
- Jablonka E, Lamb MJ (2005) *Evolution in Four Dimensions: Genetic, Epigenetic, Behavioral, and Symbolic Variation in the History of Life*. Cambridge, MA: MIT Press.
- James W (1890 [1950]). *The Principles of Psychology*. New York: Dover.
- Kaplan C, Simon HA (1990). In search of insight. *Cognitive Psychology* 22: 374–419.
- Kindermann M, Stahl I, Reimold M, Pankau WM, Von Kiedrowski G (2005) Systems chemistry: Kinetic and computational analysis of a nearly exponential organic replicator. *Angewandte Chemie (international ed.)* 44: 6750–6755.
- Kirschner M, Gerhart J (1998) Evolvability. *Proceedings of the National Academy of Sciences of the USA* 95: 8420–8427.

- Knoblich G, Öllinger M, Spivey MJ (2005) Tracking the eyes to obtain insight into insight problem solving. In *Cognitive Processes in Eye Guidance*. G Underwood, ed. Oxford: Oxford University Press.
- Klyubin AS, Polani D, Nehaniv CL (2007) Representations of space and time in the maximization of information flow in the perception-action loop. *Neural Computation* 19: 2387–2432.
- Könnnyű B, Czárán T, Szathmáry E (2008) Prebiotic replicase evolution in a surface-bound metabolic system: Parasites as a source of adaptive evolution. *BMC Evolutionary Biology* 8: 267.
- Krichmer JL, Edelman GM (2005) Brainbased devices for the study of nervous systems and the development of intelligent machines. *Artificial Life* 11: 63–77.
- Kun Á, Papp B, Szathmáry E (2008) Computational identification of obligatorily autocatalytic replicators embedded in metabolic networks. *Genome Biology* 9: R51.
- Lathe R (2003) Fast tidal cycling and the origin of life. *Icarus* 168: 18–22.
- Lieberman E, Hauert C, Nowak W (2005) Evolutionary dynamics on graphs. *Nature* 433: 312–316.
- Ludlow RF, Otto S (2008) Systems chemistry. *Chemical Society Reviews* 37: 101–108.
- Lungarella M, Sprons O (2005) Information self-structuring: Key principle for learning and development. In *Proceedings of the 14th IEEE International Conference on Development and Learning*: 25–30.
- MacGregor JM, Ormerod TC, Chronicle EP (2001) Information-processing and insight: A process model of performance on the nine-dot and related problems. *Journal of Experimental Psychology: Learning, Memory, and Cognition* 27: 176–201.
- Mach E (1897) *Contributions to the Analysis of Sensations*. Chicago: Open Court.
- Mach E (1910) *Popular Scientific Lectures*. 4th edition. Chicago: Open Court.
- Markram H, Lubke J, Frotscher M, Sakmann B (1997) Regulation of synaptic efficacy by coincidence of postsynaptic APs and EPSPs. *Science* 275: 213–215.
- Marr D (1969) A theory of cerebellar cortex. *Journal of Physiology* 202: 437–470.
- Marr D (1982) *Vision: A Computational Investigation into the Human Representation and Processing of Visual Information*. San Francisco: W.H. Freeman.
- Maynard Smith J (1970) Natural selection and the concept of a protein space. *Nature* 225: 563–564.
- Maynard Smith J (1986) *The Problems of Biology*. Oxford: Oxford University Press.
- Maynard Smith J, Szathmáry E (1995) *The Major Transitions in Evolution*. Oxford: Oxford University Press.
- Maynard Smith J, Szathmáry E (1999) *The Origins of Life*. Oxford: Oxford University Press.
- Mestel R (1996) Putting prions to the test. *Science* 273: 184–189.
- Metcalfe J, Wiebe D (1987) Intuition and insight and noninsight problem solving. *Memory and Cognition* 15: 238–246.
- Michod RE (1988) Darwinian selection in the brain. *Evolution* 43: 694–696.
- Monod J (1971) *Chance and Necessity: An Essay on the Natural Philosophy of Modern Biology*. New York: Knopf.
- Müller GB (2007) Evo-devo: Extending the evolutionary synthesis. *Nature Reviews Genetics* 8: 943–949.
- Muller HJ (1966) The gene material as the initiator and the organizing basis of life. *American Naturalist* 100: 493–517.
- Nadel L, Campbell J, Ryan L (2007) Autobiographical memory retrieval and hippocampal activation as a function of repetition and the passage of time. *Neural Plasticity*: 90472.

- Nadel L, Moscovitch M (1997) Memory consolidation, retrograde amnesia and the hippocampal complex. *Current Opinion in Neurobiology* 7: 217–227.
- Nadel L, Samsonovich A, Ryan L, Moscovitch M (2000) Multiple trace theory of human memory: Computational, neuroimaging, and neuropsychological results. *Hippocampus* 10: 352–368.
- Newell A, Simon HA (1976) Computer science as empirical inquiry: Symbols and search. *Communications of the Association for Computing Machinery* 19: 113–126.
- Oja E (1982) Simplified neuron model as a principal component analyzer. *Journal of Mathematical Biology* 15: 267–273.
- Okasha S (2006) *Evolution and the Levels of Selection*. Oxford: Oxford University Press.
- Öllinger M, Jones G, Knoblich G (2006) Heuristics and representational change in two-move matchstick arithmetic tasks. *Advances in Cognitive Psychology* 2: 239–253.
- Oudeyer P-Y, Kaplan F, Hafner VV (2007) Intrinsic motivation systems for autonomous mental development. *IEEE Transactions on Evolutionary Computation* 11: 265–286.
- Patino MM, Liu J-J, Glover JR, Lindquist S (1996) Support for the prion hypothesis for the inheritance of a phenotypic trait in yeast. *Science* 273: 622–626.
- Perkins DN (1995) Insight in minds and genes. In *The Nature of Insight*. RJ Sternberg, JE Davidson, eds.: 495–533. Cambridge, MA: MIT Press.
- Pigliucci M (2008) Is evolvability evolvable? *Nature Reviews Genetics* 9: 75–82.
- Popper KR (1972) *Objective Knowledge: An Evolutionary Approach*. Oxford: Oxford University Press.
- Polyshteyn ZW, ed. (1987) *The Robot's Dilemma: The Frame Problem in Artificial Intelligence*. Norwood, NJ: Ablex.
- Reverberi C, Toraldo A, D'Agostino S, Skrap M (2005) Better without (lateral) frontal cortex? Insight problems solved by frontal patients. *Brain* 128: 2882–2890.
- Richards RJ (1977) The natural selection model of conceptual evolution. *Philosophy of Science* 44: 494–501.
- Rougier NP, Noelle DC, Braver TS, Cohen JD, O'Reilly RC (2005) Prefrontal cortex and flexible cognitive control: Rules without symbols. *Proceedings of the National Academy of Sciences of the USA* 102: 7338–7343.
- Seung SH (2003) Learning in spiking neural networks by reinforcement of stochastic synaptic transmission. *Neuron* 40: 1063–1073.
- Sievers D, Von Kiedrowski G (1994) Self-replication of complementary nucleotide-based oligomers. *Nature* 369: 221–224.
- Simon HA, Reed SK (1976) Modeling strategy shifts in a problem-solving task. *Cognitive Psychology* 8: 86–97.
- Simonton DK (1995) Foresight in insight? A Darwinian answer. In *The Nature of Insight*. RJ Sternberg, JE Davidson, eds. Cambridge, MA: MIT Press.
- Song S, Abbott L (2001) Cortical development and remapping through spike timing-dependent plasticity. *Neuron* 32: 339–350.
- Sporns O, Alexander, WH (2002) Neuromodulation and plasticity in an autonomous robot. *Neural Networks* 15: 761–774.
- Steriade M, Pare D (2007) *Gating in Neuronal Networks*. Cambridge: Cambridge University Press.
- Sternberg RJ, Davidson JE, eds. (1995) *The Nature of Insight*. Cambridge, MA: MIT Press.
- Sutton SR, Barto AG (1998) *Reinforcement Learning: An Introduction*. Cambridge, MA: MIT Press.

- Szathmáry, E. (2000) The evolution of replicators. *Philosophical Transactions of the Royal Society of London B355*: 1669–1676.
- Szathmáry E (2002) Units of evolution and units of life. In *Fundamentals of Life*. G Pályi, L Zucchi, L Caglioti, eds.: 181–195. Paris: Elsevier.
- Szathmáry E (2005) Life: in search of the simplest cell. *Nature* 433: 469–470.
- Szathmáry E (2006) The origin of replicators and reproducers. *Philosophical Transactions of the Royal Society of London B361*: 1761–1776.
- Szathmáry E (2007) Coevolution of metabolic networks and membranes: The scenario of progressive sequestration. *Philosophical Transactions of the Royal Society of London B362*: 1781–1787.
- Szathmáry E, Maynard Smith J (1997) From replicators to reproducers: The first major transitions leading to life. *Journal of Theoretical Biology* 187: 555–571.
- Szostak JW, Bartel DP, Luisi PL (2001) Synthesizing life. *Nature* 409: 387–390.
- Toussaint M (2003) The evolution of genetic representations and modular adaptation. Ph.D. thesis, Institut für Neuroinformatik, Ruhr-Universität, Bochum, Germany.
- Van Rossum MCW, Bi GQ, Turrigiano GG (2000) Stable Hebbian learning from spike-timing-dependent plasticity. *Journal of Neuroscience* 20: 8812–8821.
- Von Kiedrowski G (1986) A self-replicating hexadeoxy nucleotide. *Angewandte Chemie (international. ed.)* 25: 932–935.
- Von Kiedrowski G (1999) Molekulare Prinzipien der artifiziiellen Selbstreplikation. In *Gene, Neurone, Qubits & Co.: Unsere Welten der Information*. D Ganten, ed.: Stuttgart: S. HirzelVerlag.
- Von Kiedrowski G (2001) Chemistry: A way to the roots of biology. *ChemBioChem*. 2: 597–598.
- Von Kiedrowski G, Eckhardt L, Naumann K, Pankau WM, Reinhold M, Rein M (2003) Toward replicatable, multifunctional, nanoscaffolded machines: A chemical manifesto. *Pure and Applied Chemistry* 75: 609–619.
- Wächtershäuser G (1988) Before enzymes and templates: Theory of surface metabolism. *Microbiological Reviews* 52: 452–484.
- Wächtershäuser G (1992) Groundworks for an evolutionary biochemistry: The iron^Δsulfur world. *Progress in Biophysics and Molecular Biology* 58: 85–201.
- Wagner GP, Altenberg L (1996) Complex adaptations and evolution of evolvability. *Evolution* 50: 967–976
- Wagner U, Gais S, Haider H, Verleger R, Born J (2004) Sleep inspires insight. *Nature* 427: 352–355.
- Wang X-J (1999) Synaptic basis of cortical persistent activity: The importance of NMDA receptors to working memory. *Journal of Neuroscience* 19: 9587–9603.
- Watson RA (2006) *Compositional Evolution: The Impact of Sex, Symbiosis, and Modularity on the Gradualist Framework of Evolution*. Cambridge, MA: MIT Press.
- Watson RA, Buckley CL, Mills R (2009) The Effect of Hebbian Learning on Optimisation in Hopfield Networks. Technical Report, ECS, University of Southampton.
- Watson RA, Hornby GS, Pollack JB (1998) Modelling building-block interdependency. In *Parallel Problem Solving from Nature*, no. 1498 of LNCS. A Eiben, T Back, M Schoenauer, HP Schwefel, eds.: 97–106. Berlin: Springer.
- Weir AAS, Chappell J, Kacelnik A (2002) Shaping of hooks in New Caledonian crows. *Science* 297: 981.
- Willshaw D, Von der Malsburg C (1976) How patterned neural connections can be set up by self-organisation. *Proceedings of the Royal Society of London B194*: 431–445.

Young JM, Waleszcyk WJ, Wang C, Calford MB, Dreher B, Obermayer K (2007) Cortical reorganization consistent with spike timing—but not correlation-dependent plasticity. *Nature Neuroscience* 10: 887–895.

Young JZ (1979) Learning as a process of selection and amplification. *Journal of the Royal Society of Medicine* 72: 801–814.

Zipser D, Kehoe B, Littlewort G, Fuster GM (1993) A spiking network model of short-term active memory. *Journal of Neuroscience* 13: 3406–3420.

