

COMPUTATIONAL APPROACHES TO NEURAL REWARD AND DEVELOPMENT

P. Read Montague^{1*} and Steven R. Quartz²

¹Center for Theoretical Neuroscience, Division of Neuroscience, Baylor College of Medicine, Houston, Texas

²Computational Neurobiology Lab, The Salk Institute for Biological Studies, La Jolla, California

Despite much progress in brain and cognitive sciences, attempts to connect brain function to cognition are hampered by the large explanatory gap between psychology and neurobiology. In recent years, a neurocomputational perspective has emerged as the most promising approach to integrating brain and mind. According to this perspective, the brain is a special sort of computer, a system of many parallel neural networks whose operation underlies cognition. In this paper, we present this neurocomputational perspective and examine the ways in which this new approach to explaining our mental skills differs from earlier ones. In particular, we examine its emerging insights into two domains. First, we explore the neurocomputational approach to decision-making, the adaptive guidance of behavior in the satisfaction of life maintenance goals. Decision-making is central to all mobile creatures in an uncertain environment, and this approach reveals a surprising conservation of decision-making strategies across many species. We then examine the neurocomputational approach's new insights into characterizing cognitive development. In particular, this approach offers the new framework of self-organization to characterize the complex interaction between neural developmental programs and the environment, a framework that has important implications for understanding early intervention.

MRDD Research Reviews 1999; 5:86-99.

© 1999 Wiley-Liss, Inc.

Key Words: prediction; temporal differences; learning; development; dynamic programming

There is something decidedly strange and immaterial about our thoughts. They seem to appear from nowhere, are vivid while present, and we have no access to where they were before they appeared. Despite these immaterial qualities, all empirical evidence available today suggests that physical processes in our nervous systems generate our thoughts. This same nervous system is the product of a billion years of evolution and is likely to be extraordinarily complex. It is responsible for a vast array of tasks required to keep our bodies running properly, and an equally vast array of information processing tasks related to our adaptive movements in the world. The exciting prospect suggested by these observations presents itself quite naturally as a question: How does the brain give rise to the mind? This question represents a deep problem that no single approach or theoretical construct will solve alone. It seems apparent that each small piece of the answer will itself engender more insight, and more questions.

The current ambition to connect brain function to cognition is fueled by a renaissance in the experimental techniques available to study brains and behavior. At the lowest

levels, cognitive neuroscience roots its theories in a molecular and cellular understanding, while at the highest level it roots its theories in new scanning technologies for peering into working human brains. It would seem that we are on the verge of connecting brain function to mental function, however, there is a hold-up—we lack ideas. As strange as it may seem, the empirical data from molecular neuroscience and behavioral cognitive science far outstrips our ability to provide satisfying explanations that connect these two domains. For example, even though we can identify brain regions where activity is associated with something interesting like an imagined word, we do not know how the word is represented in the brain region, which specific neurons are involved, or how or why they are involved. What we lack are tools and methods that bridge the gap between the molecular and the behavioral. In this paper, the framework is presented for modern computational approaches to brain function and how its use in two domains has successfully led to insights about how neural function relates to aspects of cognition and learning.

The Computational Hypothesis

The most fruitful unifying idea that promises to provide the bridging concepts between brain and mind is that of computation. The term as used here is more general than what we commonly associate with desktop computers. The main idea, introduced in a general way by Alan Turing [Turing, 1937] in the 1930s, is that a physical system undergoing transitions from one state to another can be viewed as executing computations under an appropriate set of assumptions. Rather than give a general definition of computation and its use in various fields, computational hypothesis will be presented here as it applies to neural function, and will show how this can yield heuristics for discovering how neural tissue provides a substrate for cognition.

For nervous systems, the computational hypothesis has two parts:

- (1) The physical states of the brain represent the external world, the body, and other parts of the nervous system.

Grant Sponsor: Center for Theoretical Neuroscience at Baylor College of Medicine; Grant Sponsor: NIH; Grant numbers: MH52797 and DA11723.

*Correspondence to: Dr. P. Read Montague, Center for Theoretical Neuroscience, Division of Neuroscience, Baylor College of Medicine, 1 Baylor Plaza, Houston, TX 77030. E-mail: read@bcm.tmc.edu

- (2) Transitions between these physical states constitute computations or transformations on these representations.

How can the computational hypothesis serve to connect neural function to cognitive function? The answer is that it gives a kind of heuristic for discovering explanations that connect the two domains. The main idea is to establish the computational equivalence of two systems. It is best illustrated with an example: Suppose that there was some cognitive task, like a simple decision-making task, and one suspected that activity in a particular group of neurons was involved in the decisions made by a subject on this task. The basic idea is to decompose the cognitive task into constituent parts in a manner that allows one to write a computer program to solve the task. Now do the same for the neurons that you suspect are involved in the task, i.e., write a program (or specify a set of equations) to account for their activity during the execution of the task. Crudely, if the two programs are equivalent, then the connection between the neural and cognitive substrate has been made. This might be called the *principle of computational equivalence* [Churchland and Sejnowski, 1992]. This equivalence does not of course ensure that the connection is real. Experiments must be carried out in order to establish whether the action of the neural circuit actually causes the decisions that are made.

The Incompleteness of “Brain-Neutral Cognition”

The computational heuristic outlined above implicitly suggests a distinction between neural implementation (the hardware) and cognition (the software). The strong separation of function and implementation is a misleading metaphor for the brain; however, it is an idea common to modern computer science—computer programs and their functions are independent of the way they are implemented in particular machines. This same point of view has been taken for the brain, and generally falls under the term *functionalism* [Fodor, 1981; Pylyshyn, 1984]. The functionalist position has been very effective in guiding ideas about computational functions; however, the clean hardware/software distinction that it conjures is misleading in seductive ways.

First, it is not clear if one can simply discover how cognition works based on some guiding computational notions. This was the point of view taken in the early days of artificial intelligence, i.e.,

early investigators thought that the really important problem was to find the functions or computations being implemented by the brain independent of the specifics of their implementation using biological components. This view is now seen as impoverished because as structures constructed by evolution, most creatures are tightly woven into particular environmental and social niches, and are the ‘answers’ to manifold questions posed by their environs. It seems reasonable therefore that we must cheat a little. That is, evolution has already discovered how to self-organize minds; therefore, why not investigate the tricks played by the neural substrates to try to guess at its computational functions?

Secondly, there is a misleading and implicit assumption about the nature of hardware and software in the brain. By asking why cognition must run on neural hardware or by using phrases like *neural implementation*, one implicitly buys into the seductive computer metaphor where cognition is software and neural components are hardware. The division of hardware and software in the brain is not that clear.

Lastly, even if it is possible to generate a *brain-neutral* or strictly functionalist account of cognition, which seems unlikely, such an explanation would be grossly incomplete. Events in the brain and their effects on behavior and cognition are a central concern to humans. The effects of disease, injury, birth trauma, drugs of abuse, and impoverished environments are all known to perturb brain structure and function. These changes in brain function are associated with dramatic and deleterious changes in mental function that bring with them enormous fiscal and social costs: Parkinson’s disease, schizophrenia, depression, addiction, learning disabilities, behavioral problems, etc. Moreover, pharmacological and behavioral therapies for these disorders engage specific neural systems; therefore, our main means of intervention operate in large part at the level of the neural substrate. Although these therapies seek to achieve a behavioral endpoint, the changes take place in the neural substrate. These facts alone make imperative the search for explanations that connect neural function to mental function even if the connection is remote. It must also be emphasized that the connection is not necessarily remote.

With these issues in mind let us review where we are. We have now summarized how and why the computational framework may provide a link from the operation of neural substrates to

the operations involved in human cognition. This framework for investigating both cognitive and neural function has been developed by a number of philosophers, cognitive scientists, neuroscientists, and physicists. It cuts across traditional knowledge domains, and as an explanatory framework it is still under active investigation. In fact, the extent to which all physical systems can be viewed as computational systems remains a hot topic all over the world [e.g., Gershenfeld and Chuang, 1998]. It is clear that computation is more than some series of manipulations carried out on a set of data, rather, it is one way to explain cognition as resulting from the operations carried out by neural substrates.

We have now set up the computational framework as a means of explaining various connections between mind and brain. We explore how this computational framework can be applied to the brain for a specific class of neuromodulatory neuron present in the midbrain. We then explore computational studies of development and consider some lessons it may hold for understanding development and early intervention.

Biological Decision Making

We make decisions every day. Should we choose the blue flowers or the yellow flowers? Should we go to this restaurant or that one? Should we move away from an aggressive person or do they present no real threat? Some decisions require immediate actions, which deliver immediate consequences, while the future impact of other decisions may not be known for many years, e.g., should I marry this person or not? Casting decision making at a conscious, egocentric level is misleading, and masks the fact that our nervous system is constantly making important, life-preserving decisions moment by moment. In general, decision making is a much more automatic process than our intuitions would first reveal, and necessarily so since there are simply too many crucial decisions to make moment to moment.

Different parts of the nervous system must decide when to increase heart rate, change the resistance of the blood vessels in limbs, direct an eye movement toward an interesting or potentially threatening region of visual space, and so on. These and many other automatic decisions must occur quickly and accurately with the consequences ranging from the successful acquisition of food and sex to some dismal event like an injury that results in death. Even a task as complex as choosing a mate is subject to




Fig. 1. Neuromodulatory neurons in bee brain and dopamine projections in human brain play homologous roles. Systems like the dopamine system in humans and the octopamine system in bees are called diffuse neuromodulatory systems, diffuse because the axons of the neurons are diffusely projecting, making synaptic connections throughout widespread brain regions, neuromodulatory because the neurotransmitters released from these axons are thought to modulate global brain states. Computational models show that neural activity in some of these neurons distributes information about expected reward based on previous sensory experience. In both species, the diffuse neurons receive pre-categorized information about rewarding events and combine this with sensory information to construct a scalar signal that represents errors between the expected amount of reward and the amount actually received. Using this signal to control long-term changes in synaptic weights allows the system to learn and store predictions rather than correlations.

certain automatic mechanisms that assay a host of measurable features in the potential mate. Although humans appear to have the capacity to interject conscious control and evaluation into many decisions, automatic neural processes also play an important, yet unconscious role. The algorithms employed by such decision-making mechanisms have been engineered by the many survival constraints placed on any mobile creature. There is

one evolutionary constraint placed on every adaptive system on this planet, the uncertainty of the future.

For any physical system capable of recording its experiences, uncertainty exists primarily in one direction—the future, not the past. In principle, this system can know the events it has experienced and recorded, while its future lies before it uncertain. To be reproductively successful, any organism on this planet must possess mechanisms to deal with the uncertain future. One way an organism can deal with uncertainty is to reduce it by using information gathered from the past to predict future events important for survival. It is easy to see that uncertainty about future events has a real cost: If events from my past do not help me acquire resources in the future, then I am not likely to be reproductively successful. It is therefore not surprising that all adaptive organisms have evolved numerous predictive mechanisms that guide both learning and decision making.

Neural Mechanisms That Anticipate Future Reward

One important prediction made by a host of creatures is the time and magnitude of future rewarding events. As outlined above, prediction is a computational idea; a system uses its past experience to make guesses at the likely future state of itself and its surroundings. Reward is a behavioral idea. It is used to characterize the value that a creature attributes to objects, internal physical states, actions, and so forth. For example, for bees, nectar is rewarding, and for humans, food and sex are rewarding. The reward associated with an object in the world can be measured behaviorally. Given two objects like an apple and a block of wood, the reward content of each is defined by the relative amount of time an animal spends with each object. In this case, reward is conventionally considered a latent property of the object. Reward is not, however, a static quality of objects or behavioral acts, but can depend on the internal state of an animal as well as its history of experience.

To recap briefly, uncertainty about the future has acted as a powerful selective pressure on all adaptive systems. Prediction is a computational idea that provides a way to reduce uncertainty about the future by using the past to prognosticate the future. Good predictions about the future permit an adaptive system to acquire more resources based on its past experience with the world. Such prediction may take place ‘on-line’ in a single mobile organism, or may

express itself as something more indirect like the simple strategy of overproducing progeny. Reward is associated with resources, like food and sex, known to be directly related to the reproductive success of individual organisms. Prediction about the likely time and magnitude of future rewards is therefore a fundamental task faced by all adaptive systems on this planet [Montague & Sejnowski, 1994].

The Existence and Function of Predictor Neurons

It has been discovered that humans and bees share a common computational mechanism for making predictions about the likely time and magnitude of future reward [Schultz and Romo, 1990; Quartz et al., 1992; Schultz et al., 1993; Hammer, 1993]. Experiments also show that this common computational mechanism appears to be implemented by homologous neural structures present in the brains of both humans and bees. These findings show a fascinating conservation of both computational strategy and neural substrate.

There are neurons situated at the base of the brain (Fig. 1) that use information from the world and other parts of the nervous system to make sophisticated guesses about future reward. These guesses are apparently used by bees and humans in at least two roles:

- (1) to guide learning, which improves future guesses, and
- (2) to guide decision making about which actions lead to future reward [Montague et al., 1993, 1995, 1996; Schultz et al., 1997].

These neurons appear to be responsible for sending out an ongoing ‘economic evaluation’ of whether the current state of the organism predicts that the near future is likely to be rewarding (Fig. 2).

More specifically, changes in the firing rate of these neurons represent a prediction error about future reward, i.e., the ongoing difference between the expected amount of reward and the actual amount received. These same changes in firing rate are associated with changes in neurotransmitter release from the synaptic connections made by these neurons (Fig. 3). Hence, the prediction error is probably translated into changes in neurotransmitter delivery. In humans, this neurotransmitter is dopamine, and in bees, the neurotransmitter is octopamine (a close chemical relative of dopamine). These neurons make synaptic connections throughout widespread regions of the nervous system (Figs. 1, 2); therefore,

Architecture of predictor neurons

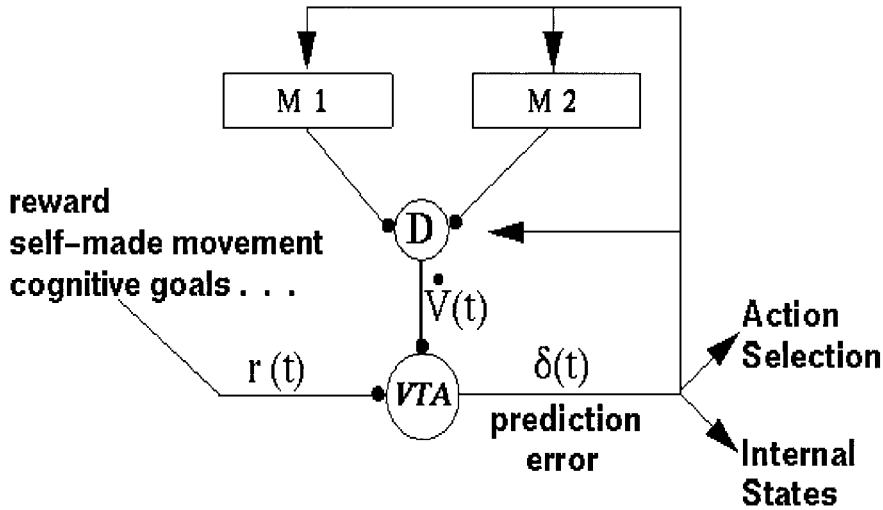


Fig. 2. Temporal difference model of midbrain dopamine neurons. Neurons in the ventral tegmental area (VTA) receive two classes of input: (1) input from sensory representations, which arrives in the form of a surprise signal, and (2) information about reward, self-made movements, etc. M1 and M2 represent different modules or modalities. The surprise signal means that activity along these inputs encodes the degree to which the current state is different from the last state. This kind of signal could even arise from modules that analyzed complex features of a situation. The signals along $r(t)$ are construed as *reward* signals (see text), but are meant to represent information already categorized as important—it is this signal which the predictor neurons learn to anticipate. By combining (adding) these signals, the output of the VTA neuron, $\delta(t) = r(t) + \dot{V}(t)$ can be interpreted as a prediction error [Sutton and Barto, 1981; Sutton, 1988; Montague et al., 1996; Schultz et al., 1997]. The synaptic weights along the path to the VTA represent stored predictions.

the prediction error is simultaneously broadcast to many brain structures at once.

How is the prediction error signal constructed and improved in a real brain? In both bees and humans, it appears that input from sensory representations may construct the predictions through a set of adaptable synaptic strengths (weights) (Fig. 2). If the prediction error is carried by dopamine fluctuations (as in primates), then synaptic plasticity that depended on such fluctuations would in principle allow a target neural structure to improve predictions about future reward. In regions of the nervous system related to action selection, this signal can be used directly to evaluate the likely future “goodness” of a planned action. As described below, this computational framework is derived from well-understood computational methods of adaptive control.

Detailed Requirements for Constructing a Prediction Error Signal in Real Brains

One way for an animal to make predictions about future reward is to possess a system that reports its current best guess, and have learning be contingent on errors in this guess. The construction, delivery, and use of an error signal related to predictions about future stimuli would require the following [Montague et al., 1996]:

- i) access to a representation of the phenomenon to be predicted such as the amount of reward or food;
- ii) access to the current predictions so that they can be compared to the phenomenon to be predicted;
- iii) capacity to influence plasticity in structures responsible for constructing the predictions;
- iv) sufficiently wide broadcast of the error signal so that stimuli in different modalities can be used to make and respond to the predictions.

These general requirements are met by a number of neuromodulatory systems including dopamine systems in humans and primates, and octopamine system in bees (Fig. 1). Neuromodulatory neurons receive synaptic input that carries information about reward as well as information from sensory representations. This input originates from widespread regions of the nervous system and converges on a relatively small number of neuromodulatory neurons. This high

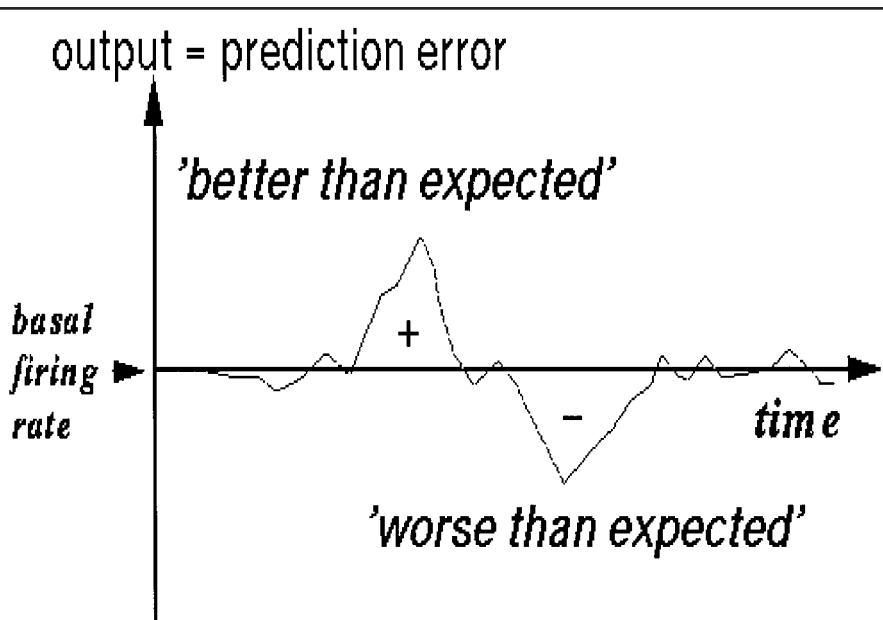
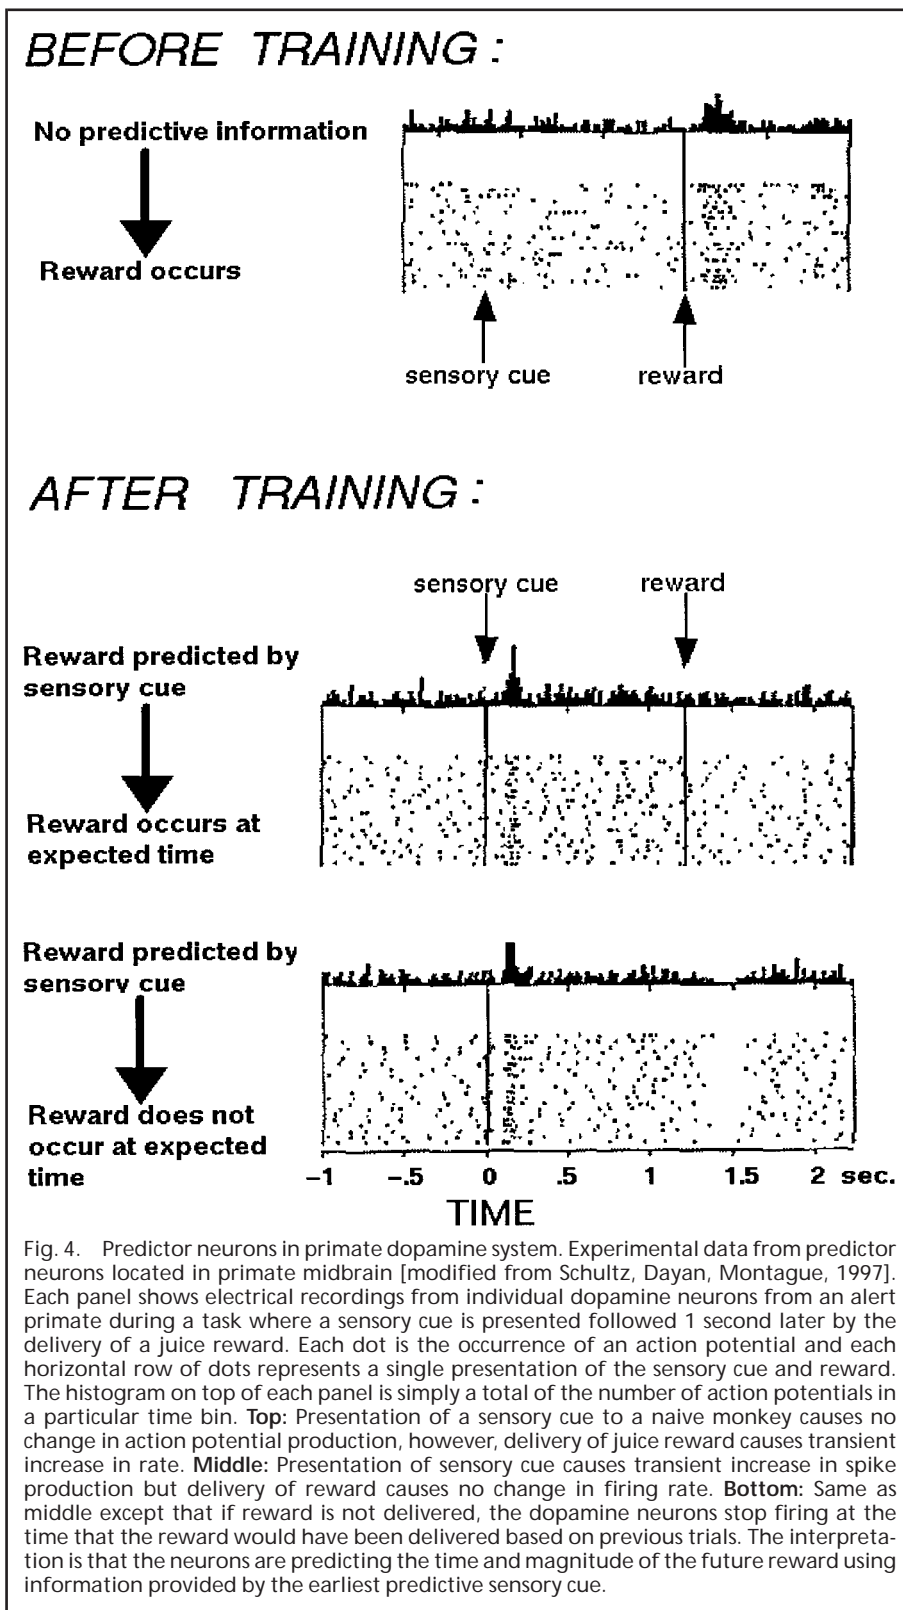


Fig. 3. Biological interpretation at the synaptic terminals of predictor neurons. Positive prediction errors are encoded as an increase in spike production and dopamine delivery. They indicate that the current state is “better than expected.” Negative prediction errors are encoded as a decrease in spike production and dopamine delivery. They indicate that the current state is “worse than expected.” This kind of signal can be used in dual roles: (1) to improve predictions, and (2) to control on-line selection of actions



change in its output $\delta(t) = r(t) + \dot{V}(t)$. In doing so, the change in spike output $\delta(t)$ of the predictor neuron reflects an ongoing *prediction error* in future expected reward (see legend for Fig. 2).

The most detailed data about the nature of neural processing in the dopamine system comes from extensive electrophysiological studies over more than fifteen years by Wolfram Schultz and his collaborators [e.g., Schultz, 1992; Schultz and Romo, 1990; Schultz et al., 1993]. In these experiments, the activity of single dopamine neurons is recorded in alert monkeys while they learn to perform behavioral tasks and receive rewards. For untrained monkeys, dopamine neurons respond with transient increases in their firing rate when the monkeys are presented with appetite stimuli like a morsel of apple (Fig. 3). The neurons act surprised by the unanticipated arrival of reward. The monkeys also act surprised (assayed behaviorally). However, if a sensory cue like a light or tone consistently precedes (predicts) the delivery of reward, a remarkable change occurs in both the electrical behavior of the neurons and the behavior of the monkeys.

After repeated pairings of a sensory cue followed by reward, dopamine neurons change the time of their phasic activation from just after the time of reward delivery to the time of sensory cue onset (Fig. 4). Moreover, the neurons no longer respond to the delivery of reward if it occurs at the expected time. During this pairing process, if the reward is not delivered at the expected time, the dopamine neurons dramatically decrease their firing rate at the time the reward would have been delivered. The transfer of the dopamine neuron activity is paralleled by a behavioral transfer in the monkey: After training, the onset of the predictive sensory cue causes the monkey to act like it expects reward at a particular time.

The neurons are actually more sophisticated than this previous description implies. This claim is illustrated by the behavior of the neurons in the presence of multiple predictors of reward. For example, repeated presentation of the sequence

light 1 $\xrightarrow{1 \text{ sec}}$ light 2 $\xrightarrow{2 \text{ sec}}$ reward

is associated with the following changes. During early trials, the delivery of the reward produces increased firing rate as before; however, after training, only the earliest cue is associated with increased firing rate. The activity transfers to the earliest predictive sensory cue, and ap-

degree of convergence effectively discards the identity of the source of information so that it is reasonable to assume that neuromodulatory neurons produce only a *scalar output signal*.

The computational model of the dopamine system receives sensory information separately in the form of a surprise

signal (denoted $\dot{V}(t)$ in Fig. 2) and a reward signal ($r(t)$ in Fig. 2). The idea of the surprise signal is straightforward: It encodes the degree to which the current sensory state is different from the immediately preceding sensory state. The predictor neuron (labeled VTA) simply adds these two signals together to produce a

pears to ignore intermediate cues that predict the same future event. As described below, this neural behavior can be understood using computational theories of adaptive optimizing control.

A Temporal Difference Model Explains Activity of Predictor Neurons

The computational model that we have developed predicts this exact pattern of dopamine neuron activity. In addition, the model predicts activity patterns in the dopamine neurons during behavioral tasks that have not yet been tested in experiments. The model is based on a computational theory called temporal difference learning developed by Sutton & Barto in the early 1980s [Sutton and Barto, 1981,1990; Sutton, 1988].

In temporal difference learning, a system uses adaptable weights to make predictions of the time and magnitude of the reward it expects to receive in the future and provides a continual output of the error in the prediction. Instead of learning just about the reward expected in the near future, this method learns to predict the total amount of reward expected for long times into the future. These prediction errors drive learning (synaptic weight changes in our neural case) that improve predictions made by the system. The activity recorded in the dopamine neurons during learning tasks (Fig. 4) suggests strongly that these cells are reporting *prediction errors* in the future delivery of reward, which is exactly the signal that temporal difference learning requires to operate.

The architecture of the model mimics the architecture of neuromodulatory systems in bees and humans. As detailed in Figure 5, the output of the model yields exactly the experimental results taken from experiments by Wolfram Schultz (Fig. 4). The model also accounts for many more complicated experiments, but the example included here makes all the basic points.

The model learns to improve its predictions by using the broadcast prediction error signal encoded in the changes in dopamine delivery. This learning takes place in a simple way: Changes in synaptic strength are proportional to the product of the prediction error signal (dopamine fluctuation) and activity in target neurons. Simple learning rules analogous to this one have been used for many years by psychologists, engineers, and neurobiologists. In a biological context, these previous rules have generally focused on learning associations or correlations between experienced events. Us-




Fig. 5. Response of temporal difference model. This surface shows the changes in the prediction error signal (output of dopamine neuron) during training. The time base is arbitrary, but the sensory cue and reward were always separated by the same number of time steps. Early trials: model gives positive response to delivery of reward and no response to sensory cue. Late trials: model gives positive response to sensory cue and no response to reward. If the reward is not delivered, the model gives negative response at the expected time of reward. Compare these responses to those in Figure 4.

The same responses will develop in the presence of multiple predictive sensory cues. Other more complex profiles are also known, and they are accounted for by this model.

ing the prediction error signal as one factor in a simple multiplicative rule allows the model to learn and store predictions rather than correlations (associations).

Using the Prediction Error Signal to Select Actions

The computational framework presented above would be interesting if it only accounted for a common neural processing strategy in bees and humans. There is another important use to which the prediction error signal can be put, the evaluation of possible future actions. We have shown that the same signal is likely to be used by bees and primates to select planned actions, i.e., to make decisions. The first hint that this is true comes from the monkey work—the transfer of firing observed for dopamine neurons is paralleled by changes in behavior of the monkeys.

Perhaps the monkeys are actually using the fluctuation in dopamine delivery to help select their actions. The anatomy of the dopamine axons supports this hypothesis; the dopamine neurons make connections throughout neural structures well known to be involved in motor planning and the sequencing of motor actions. This possibility in monkeys is mirrored by similar data from invertebrates where neuromodulator delivery directly influences motor behavior. These observations show that, in humans and bees, the prediction error signal is

likely to be available to neural structures involved in action selection. Not only is the signal available in the relevant neural structures, it is precisely the correct kind of signal for directing the selection of actions. This latter assertion is explained below.

There is a principled way to use just such an on-line prediction error to choose sequences of actions or to learn how to choose optimal sequences of actions. This method is called *dynamic programming*, and was discovered in the late 1950s by Richard Bellman [Bellman, 1957]. The problem taken on by dynamic programming is called the *temporal credit assignment problem*, i.e., actions at one point in time can affect important future events in complicated ways.

For example, an organism may have taken many actions prior to some important future event like a bump on the head or the acquisition of food. How does the organism assign credit to the prior actions most responsible for its current state? Consider the cartoon example in Figure 6. The bug walking along a tabletop falls off and bumps its head; however, its last few steps are not primarily responsible for the bump. Once it passes through the narrow opening in the barricade, the number of paths that lead to a bumped head far outnumber the number of paths that lead to freedom. One wrong move plus statistics gave the bug its bump. Although this is an overstatement, one very important action




Fig. 6. Temporal credit assignment. Actions taken at one point in time can be related in complicated ways to futures consequences both good and bad. A bug walks along a table, falls, and bumps its head. Which of its previous actions was most responsible for the bump? The step through the barricade was certainly important.

was the one where the bug stepped through the narrow opening in the barricade; this action should get a lot of credit for the bump.

Dynamic programming solves the temporal credit assignment problem using the idea of a *policy* [Howard, 1960]. A policy indicates how specific actions are assigned to specific states. For our purposes, the state of the bug is its collection of sensory cues at each position on the tabletop. The goal of the bug is to improve its policy through interaction with its environment. To perform dynamic programming, the bug requires an on-line evaluation of its current state, along with the capacity to use the evaluation to choose an appropriate action. The evaluation used by dynamic programming is the total future rewards expected from each position on the table. This is precisely what the temporal difference method learns. Hence, the output of the predictor neurons is ideally positioned to select appropriate sequences of actions in a computationally reasonable way.

Roughly, the predictor neurons emit a signal that means 'things are better than expected' (prediction error positive), 'things are the same' (prediction error 0), or 'things are worse than expected' (prediction error negative). A simple mechanism for choosing actions is readily apparent: Take actions correlated with increased dopamine neuron activity (positive prediction error), and avoid actions correlated with decreased dopamine neuron activity (negative prediction error).

As simple as this mechanism sounds, its action reproduces the strategies used

by humans and bees on simple decision-making tasks. We see that the notion of *computational equivalence* described earlier is helpful. Using the prediction error in this neural algorithm is computationally equivalent to executing a form of dynamic programming. This connection gives us a handle on what to expect from systems that employ the prediction error signal to select actions.

The Cost of Future Uncertainty Encoded in Dopamine Fluctuations?

The behavioral tasks that we have addressed all involve uncertainty about future rewards, and the cost of this uncertainty to the organism. The game is clear: We have a computational description of a signal that evaluates the state of the organism and provides estimates of the likelihood that one action will yield more total future reward than another. This signal is constructed and distributed by homologous neural structures in the brains of insects and primates. If this signal is actually used by the organisms that possess it, then we should be able to predict their behavior on a variety of tasks that involve estimating likely future reward.

Leslie Real [Harder and Real, 1987] has examined the role of uncertainty using bee foraging behavior. He studied how the mean and variability of nectar returns from flowers affects the way bees decide to forage on flowers. An artificial field of flowers with two flower colors (blue and yellow) was used. Two units of nectar were placed in all the blue flowers, six units were placed in one third of the yellow flowers, and zero units were placed in the remaining two thirds of the yellow flowers. Bees were released over the enclosed field, and the fraction of visits to each flower color was recorded. Flower color was the only physical feature that predicted the nectar reward. Bees behave very much like humans on similar tasks; they make about 85% of their visits to blue flowers despite the same mean return (in nectar) from the more variable yellow flowers. This behavior is called *risk aversion* [Harder and Real, 1987]. Real also tested the value of uncertainty for the bees. He increased the mean return on the yellow flowers while holding the variance constant. When the bees sampled indifferently from each flower type (sampled 50% from each), he recorded the amount of increase in the mean return required to elicit this indifference. This measures the value of the variance (uncertainty) in units of nectar.

We equipped the temporal difference model with a simple visual system, and used computational techniques to construct a simulated field of blue and yellow flowers on which the model bee could forage. This computational experiment was not a model of flight, but tested the way in which the bee could use its prediction error signal to choose its next action. At each point in time the model bee had two choices: (1) continue in the current direction, or (2) change directions. The sign of the prediction error selected which of these choices was more likely: prediction error positive, continue in same direction; prediction error negative, change to new random direction.

When presented with the reward distributions described above, the virtual bees behave just like the real bees and sample the blue flowers (constant return) on 73% to 85% of their landings. We were also able to test the cost of uncertainty using the virtual bees. The model again captured the real bee behavior. So why should bee foraging behavior have anything to do with the human decision making? The hint originates in the fact that the dopamine neurons in primates and the octopamine neurons in the bee appear to be implementing a similar computational strategy with homologous neural components. In addition, learning displayed by the monkeys parallels the learning displayed by the dopamine neurons. Therefore, we set out to test whether the *same action selection mechanism* described above would account for human behavior on decision tasks analogous to those faced by the bees.

We gave humans a related task, but instead used two buttons on a computer screen (Fig. 7). The rewards associated with the selection of each button varied as a function of a short history of selections. This is illustrated in the three tasks shown in Figure 8—the reward from button A and button B depends on the fraction of previous selections made from button A. The fraction is computed over the last 30 selections. So the rewards received vary depending on how one's selections are allocated to each alternative.

In designing the reward functions, we introduced features in the rewards that would cause the model to get stuck in behaviors that were suboptimal, i.e., the model's sensitivities to changing reward can "blind" it to more profitable strategies. Our intuition was that humans would almost certainly behave differently or at least with much more variability than various incarnations of the model. Our expectations were wrong in a

number of informative ways. For example, the humans act just like the bee would have if the bee were capable of executing the choice tasks given to the humans. We suspect that although the humans have more memory capacity and more sophisticated ways to categorize tasks, they appear to use the prediction error signal in the same way as bees.

Task 1: Flat returns for all decision strategies. Figure 8 (top) shows the average results for the networks and human subjects on a task where all strategies give the same long-term returns. Every allocation of choices yields the same return (dotted line) over the 250 total selections allotted to this task. This task is modified from one originally used by Richard Herrnstein in experiments on rational choice behavior in humans [Herrnstein, 1991]. Humans and model converge quickly to a stable strategy, making choices that tend to equalize returns from each button. For both humans and network, the mean allocation to button A settled near the crossing point in the reward functions with a slight central tendency (tendency to play 50% from each alternative). This central tendency has been tested in experiments, is real, and is predicted from the behavior of the model. Our human results confirm Herrnstein's results on a very similar task; however, we should note that this is exactly the behavior expected based on the model. The model provides a plausible mechanistic explanation for why humans tend to stick at the crossing point in the reward functions.

Task 2 and 3: Optimal strategy lies on and off crossing point. In task 2 (Figure 8 middle & bottom), the attraction to the crossing point allowed 18 of 24 subjects to maximize their long-term return; however, in task 3, the same attraction to the crossing point blinds humans to higher long-term returns. In this latter task, over half (14 of 25) the subjects converged to the crossing point even though other, easy to discover allocations yielded higher long-term returns. The remaining 11 of 25 displayed a mean of around 70% allocated to button A. Choosing roughly 80% from button A is close to optimal behavior. All the networks converge to the crossing point.

Task 4 (not shown): No systematic reward structure. In the presence of a pseudo-random sequence of rewards, both the networks and the humans play almost exactly 50% from each button. In this case, the numbers are striking: network 0.498 ± 0.007 ($n=19$), humans 0.501 ± 0.002 ($n=19$).

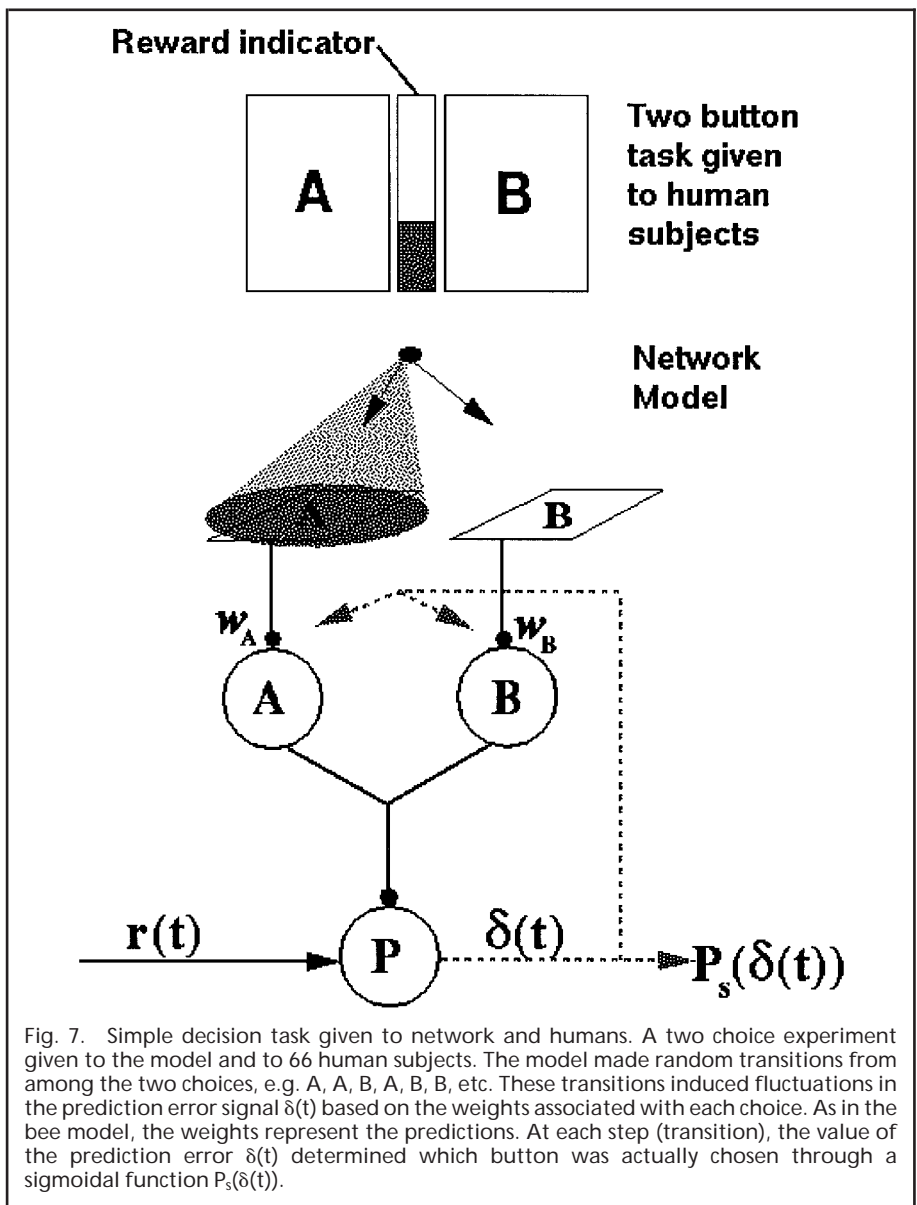


Fig. 7. Simple decision task given to network and humans. A two choice experiment given to the model and to 66 human subjects. The model made random transitions from among the two choices, e.g. A, A, B, A, B, B, etc. These transitions induced fluctuations in the prediction error signal $\delta(t)$ based on the weights associated with each choice. As in the bee model, the weights represent the predictions. At each step (transition), the value of the prediction error $\delta(t)$ determined which button was actually chosen through a sigmoidal function $P_s(\delta(t))$.

Computational Bridges Dopamine Neuron Activity to Simple Decision-Making

We have described at length the reasons that we think dopamine systems are in part responsible for distributing a sophisticated evaluatory signal throughout the nervous system. We have suggested how and why we think that this signal is a scalar *prediction error* in the expected amount of future reward. Moreover, we have also argued that reward can be a relatively complex quantity that depends on many variables; however, it is clear that successfully and consistently obtaining reward confers a selective advantage.

The findings described above highlight a number of interesting issues. The evolutionary conservation of both the computational and neural mechanisms is striking. Perhaps once the solution was

found (for the problems we considered), there were no dramatically different solutions that were better. Another striking result is the degree to which the model predicts behavioral performance in humans. It did not have to be this straightforward. The computational model explains the output of dopamine neurons and was not engineered to explain choice behavior on the decision tasks that we examined.

This example illustrates how the computational approach provides a framework to link cellular and behavioral work in an explanation of mental operation. Another area of fertile computational research is the study of development; computational studies of development are producing new lessons for thinking about the connection between mind and brain and offer an important perspective on early experience and intervention.




Fig. 8. Network and human performance on Task 1, Task 2, Task 3. Averages for humans and network are shown. Each graph plots the reward received against the fraction of choices allocated to button A. For humans, the reward received was encoded as the height of the central bar between the buttons. The horizontal position of the symbols (asterisk, diamond) encodes the average allocation to button A, and the vertical position encodes the average reward actually obtained over the task (250 button presses). Subjects were instructed to try to maximize their returns over the course of the task.

Development From a Computational Perspective

We have already mentioned Turing's monumental work on the theory of computability. In 1938, Claude Shannon showed how to implement Turing's logic with electrical switching elements. Shannon went further and invented information theory, of which the basic unit is the binary digit or bit, the amount of information obtained in knowing which of two states a switch is in (on or off, corresponding to true or false). In 1943,

Warren McCulloch and Walter Pitts [McCulloch and Pitts, 1943] applied these two ideas—information and Turing's work—to the brain.

They suggested that since neurons appear to be in one of two states, quiet or firing, they were equivalent to logical switches. They then showed how neurons could be arranged in the brain to make the brain a Universal Turing Machine.

In 1958, Frank Rosenblatt at Cornell [Rosenblatt, 1958] took the idea a step further when he developed a brain-style computer he called the perceptron. Based loosely on the retina, Rosenblatt trained his system to recognize letters by changing the connections between units during learning. Although Rosenblatt's model contained the seeds of a brain-style approach to computing known today as neural networks, his model was severely limited in what it could learn, as Marvin Minsky and Seymour Papert pointed out in their 1969 book, *Perceptrons* [Minsky and Papert, 1969].

For these and other reasons, neural network research was largely dormant until the mid-1980s. By 1985, a group of researchers known as the Parallel Distributed Processing (PDP) research group had achieved learning in several types of multilayer networks, showing that direct descendants of the perceptron could overcome the limitations Minsky and Papert had shown [Rumelhart and McClelland, 1986]. Rosenblatt's model had been very simple, and Minsky and Papert's critique applied only to this simple version of neural networks. The publication of the PDP Research group's two-volume work in 1986 sparked an explosion of interest in neural networks. After years of hibernation, almost overnight a new scientific discipline was founded. Associations, conferences, journals, and academic departments emerged, making the rise of brain-style computing one of the most significant scientific advances of that decade.

The resurgence of neural network models has been especially revealing for development. As we explore, computational studies of development are shifting how we view the relationship between developing cortex and learning, a shift in perspective that is very much aligned with recent experimental work in developmental neurobiology. In particular, models of "self-organization" reveal how complex structures can emerge through the interaction between developing cortex and temporal patterns of neural activity reflecting environmental structure [Bienenstock et al., 1982; Miller et


al., 1989; Montague et al., 1991; Goodhill et al., 1997]. These models suggest that complex cortical structures do not require a strong pre-specification, but rather may emerge reliably through the rich interaction between developmental mechanisms and environmentally-derived information (Fig. 9). As recent work in developmental neurobiology likewise suggests, this computational work suggests that the traditional distinction between learning and cortical maturation no longer holds. The computational perspective provides a means of characterizing the significance of this interaction. In particular, it suggests that activity-dependent cortical development solves a fundamental problem of learning (Figs. 9–11). We suggest that understanding this learning problem from a computational perspective highlights the importance of early intervention for at-risk populations. More generally, it suggests a needed realignment between cognitive and neural development and how maturation and learning are related. Historically, the study of cognitive development has remained largely isolated from the study of neural development. Despite their relative independence, however, there are important historical parallels, as we explore.

Cognitive and Neural Development

Roger Sperry's [1963] landmark developmental work produced an influential view of brain development. Culminating in his "chemoaffinity hypothesis," Sperry regarded neural development as a process whereby neural connectivity was determined in a molecular lock-and-key fashion; axons were viewed as carrying molecular addresses for particular target sites. The wiring that resulted from such a scheme was thought to delimit the behavioral repertoire mediated by the developed neural circuits. Sperry's interpretation of his work was even more expansive; he believed it demonstrated that developmental psychology had become a part of neurobiology, thus eliminating the notion that the acquisition of cognitive skills was a process of learning. Although the work was exciting in its potential scope and demystified a number of preconceptions of the day, this interpretation is now seen to be unnecessarily restrictive.

Surprisingly, a similar view also dominated cognitive science at around the same time. Developed independently by the linguist Noam Chomsky [1965] with very different kinds of evidence, Chomsky too surmised that the neural structures underlying complex skills

Fig. 9. Activity-dependent mechanisms and their influence on neural mappings. Activity-dependent computational models use some form of long-term synaptic change to allow patterns of neural activity to influence mappings from the sensory world to more central structures like the cerebral cortex. These models are also thought to capture the kinds of ongoing synaptic modification that takes place at connections between central neural structures. The main idea in all these models is that neural activity in the source region drives synaptic changes in the target region through the action of synaptic learning rules, i.e., mechanisms that influence the strength or lifetime of a synapse. One favorite mechanism is a correlational learning rule—this rule strengthens synapses that experience correlated presynaptic and postsynaptic activity (spikes), and otherwise weakens them. A sufficiently weak synapse is eventually removed. Such a mechanism extracts statistical correlations present in the neural activity and represents them in the pattern of synaptic strengths. The model shown here illustrates the refinement of a topographic mapping from a source layer into a three-dimensional target layer. Initially, the connections originating from each point in the source layer were made throughout the recipient target region with crude topography. This is illustrated in the upper panel (“Before training”); a horizontal pattern of activity activates a broad band of synaptic connections. In this particular model, a correlational learning rule either stabilized or removed synaptic connections based on the correlation of presynaptic activity and the ambient levels of the gas nitric oxide. The pattern of activation in the input layer followed a simple rule—nearby neurons tended to fire together. This pattern of activation together with the synaptic learning rule allowed the mapping to sharpen so that the same horizontal stimulus pattern now activates a much more refined portion of the target region (**bottom**). Active synapses are shown as light circles; the two-dimensional plot of the target is taken from the slice as indicated. [Adapted from Montague et al., 1991.]



emerged through an intrinsic process of maturation. Chomsky reached his conclusions through the burgeoning field of computer science, which provided him with computer languages and a study of their complexity. Out of this grew computational learning theory, a mathematical treatment of learning centered on the study of assessing the difficulty of learning various classes of formal languages. In 1967, Mark Gold [Gold, 1967] published a seminal paper in mathematical learning theory, which seemed to demonstrate very pessimistic results against learning. Essentially, Gold argued that learning a language was possible only if the learner brought very specific prior knowledge to the problem. Otherwise, learning is prohibitive—the problem is simply too difficult without a lot of prior knowledge.

In one of the intellectual tour de forces of this century, Chomsky applied these emerging ideas to human languages. Transforming the field of linguistics from a cataloging endeavor to one concerned with the deep structure common to all human languages, Chomsky argued that languages were too complex to be learned in any meaningful sense. Chomsky's arguments were not based simply on intuition. They were bolstered by mathematical learning theory and psychological studies revealing that children received only limited feedback from parents regarding the grammaticality of their utterances [Brown, 1973]. Children, argued Chomsky, must come equipped with innate knowledge of language in the form of a universal grammar, a rule

system capable of generating all the grammatical instances of human language. Exposure to an environment simply set parameters in this template to the appropriate values for the particular language confronting the child. Combined, this led to the image of the language system of the brain as a mental organ or module, which unfolded analogously to the development of any other organ. In 1967, Eric Lenneberg [Lenneberg, 1967] brought the view of intrinsic cortical maturation and language acquisition together in his now classic work, *Biological Foundations of Language*. Subsequently, this developmental view has spread to many other domains in cognitive science.

Under the influence of this view, the interaction between a child and his or her environment was devalued. Since the unfolding of cognitive skills was largely a matter of the intrinsic biological maturation of prespecified structures, as long as certain minimal conditions were met, such as nutritional ones, there would be no effect on outcomes. As many of the papers in this volume illustrate, however, developmental research over the last decade is giving rise to a markedly different view. A new appreciation for the progressive nature of development [Purves et al., 1996] and its dependence on environmentally-derived activity are its hallmarks. In light of this new research, the innate modules of Chomsky and their intrinsic maturation seem highly problematic and are coming under increasing revision [Elman et al., 1996; Quartz and Sejnowski, 1997]. Recall, however, that

the motivation for modules stemmed from the results of mathematical learning theory. It is important to emphasize this key insight from mathematical learning theory, which has become more forceful in the decades since Gold's work [For review, see Quartz, 1993]. We can illustrate this with an example. Suppose we want to design a system to automatically sort letters by recognizing handwritten addresses. One way to do this would be by “automated programming.” This technique would equip the computer system with a powerful learning algorithm and a large set of samples to discover on its own how to categorize handwriting into discrete letters and digits. Unfortunately, such a system would take prohibitively long to learn because it does not have any prior knowledge that the patterns are supposed to be letters and digits. There are arbitrarily many consistent ways of sorting this data and the computer could never know on its own which was the right way. A more powerful system would come equipped with built-in knowledge about letters and digits, perhaps as dynamic templates that it uses to fit the data. Now the computer knows what it is looking for and so imposes its knowledge on the data to classify it. This built-in knowledge is known as inductive bias. In recent years, the prospect of building general-purpose systems has waned, as researchers realize that systems must be tailored to specific domains for engineering applications.

What computational studies of learning reveal, then, is that the most

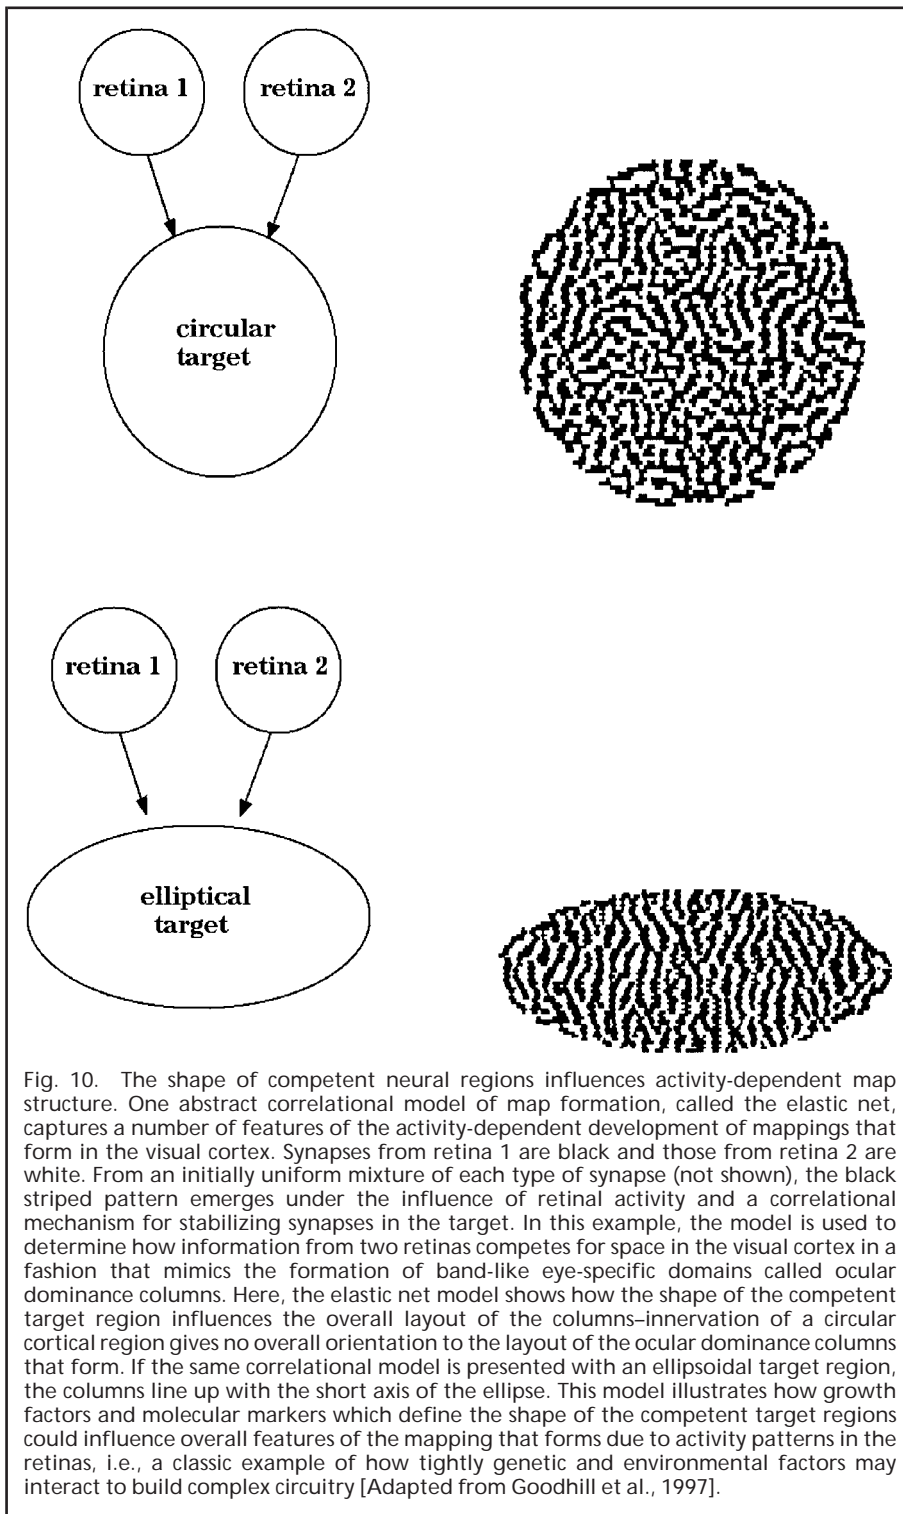


Fig. 10. The shape of competent neural regions influences activity-dependent map structure. One abstract correlational model of map formation, called the elastic net, captures a number of features of the activity-dependent development of mappings that form in the visual cortex. Synapses from retina 1 are black and those from retina 2 are white. From an initially uniform mixture of each type of synapse (not shown), the black striped pattern emerges under the influence of retinal activity and a correlational mechanism for stabilizing synapses in the target. In this example, the model is used to determine how information from two retinas competes for space in the visual cortex in a fashion that mimics the formation of band-like eye-specific domains called ocular dominance columns. Here, the elastic net model shows how the shape of the competent target region influences the overall layout of the columns—innervation of a circular cortical region gives no overall orientation to the layout of the ocular dominance columns that form. If the same correlational model is presented with an ellipsoidal target region, the columns line up with the short axis of the ellipse. This model illustrates how growth factors and molecular markers which define the shape of the competent target regions could influence overall features of the mapping that forms due to activity patterns in the retinas, i.e., a classic example of how tightly genetic and environmental factors may interact to build complex circuitry [Adapted from Goodhill et al., 1997].

natural systems get this background knowledge, the repertoire of representations that facilitate learning? If it is not built-in according to Chomsky, where is it?

Neural network models of development and their analysis suggests that these representations are constructed during development through an interaction with a information-rich environment. Traditional learning theory did not allow a system to construct representations in this manner. Instead, it limited learning to a search through a pre-defined class of representations. Allowing a system to construct its representations as it learns appears to be a powerful form of learning, one with unique learning properties [Quartz, 1993].

In other words, viewed from a computational perspective, early development is a period of representation construction that makes the subsequent acquisition of cognitive skills, such as language, possible. One of the key insights of computational studies [e.g., Miller et al., 1989; Berns et al., 1993; Montague et al., 1991] is that complex neural structures can emerge out of the interaction among richly structured neural activity, developmental mechanisms, the constraints imposed by subcortical organization, the hierarchical organization of cortex, and a basic common circuitry that O'Leary [1990] refers to as "protocortex" (Figs. 10,11). A major concern of computational studies is to understand this interaction and to characterize the algorithms that build the neural structures underlying complex representations. Complementary to the study of networks that build their representations as they experience the world is the increased interest in characterizing algorithms capable of extracting more information latent in the environment than previous accounts believed existed.

Seeing early development as a period of representation construction making efficient learning possible makes sense of another apparent paradox involving human development. The duration of human development has often been seen as little more than a period of increased vulnerability. However, as a period of environmentally-guided representation construction, its duration suggests that it is a powerful strategy for constructing the complex representations underlying human cognitive skills. Extending the duration of representation construction maximizes the capacity of the world to participate in this process of construction, allowing increasingly powerful representations to emerge [Quartz, 1999].

basic problem of learning is not the traditional one known as statistical inference [Geman et al., 1992]. Rather, it is one of finding the efficient representations that make statistical inference possible at all. Further, how difficult it is to learn something is really a question that only makes sense relative to this background knowledge. Our previous example shows that learning strategies controlled by neurons in humans and bees are sensitive to particular contingencies

through time; hence, they are very much pre-equipped to deal with specific kinds of learning. That is, they provide background knowledge for the tasks they influence.

Combining these results with the emerging view of neural development appears to lead to a paradox: How can we reconcile the emerging view of development with the requirement that learning is possible only against a background of appropriate representations? Where do

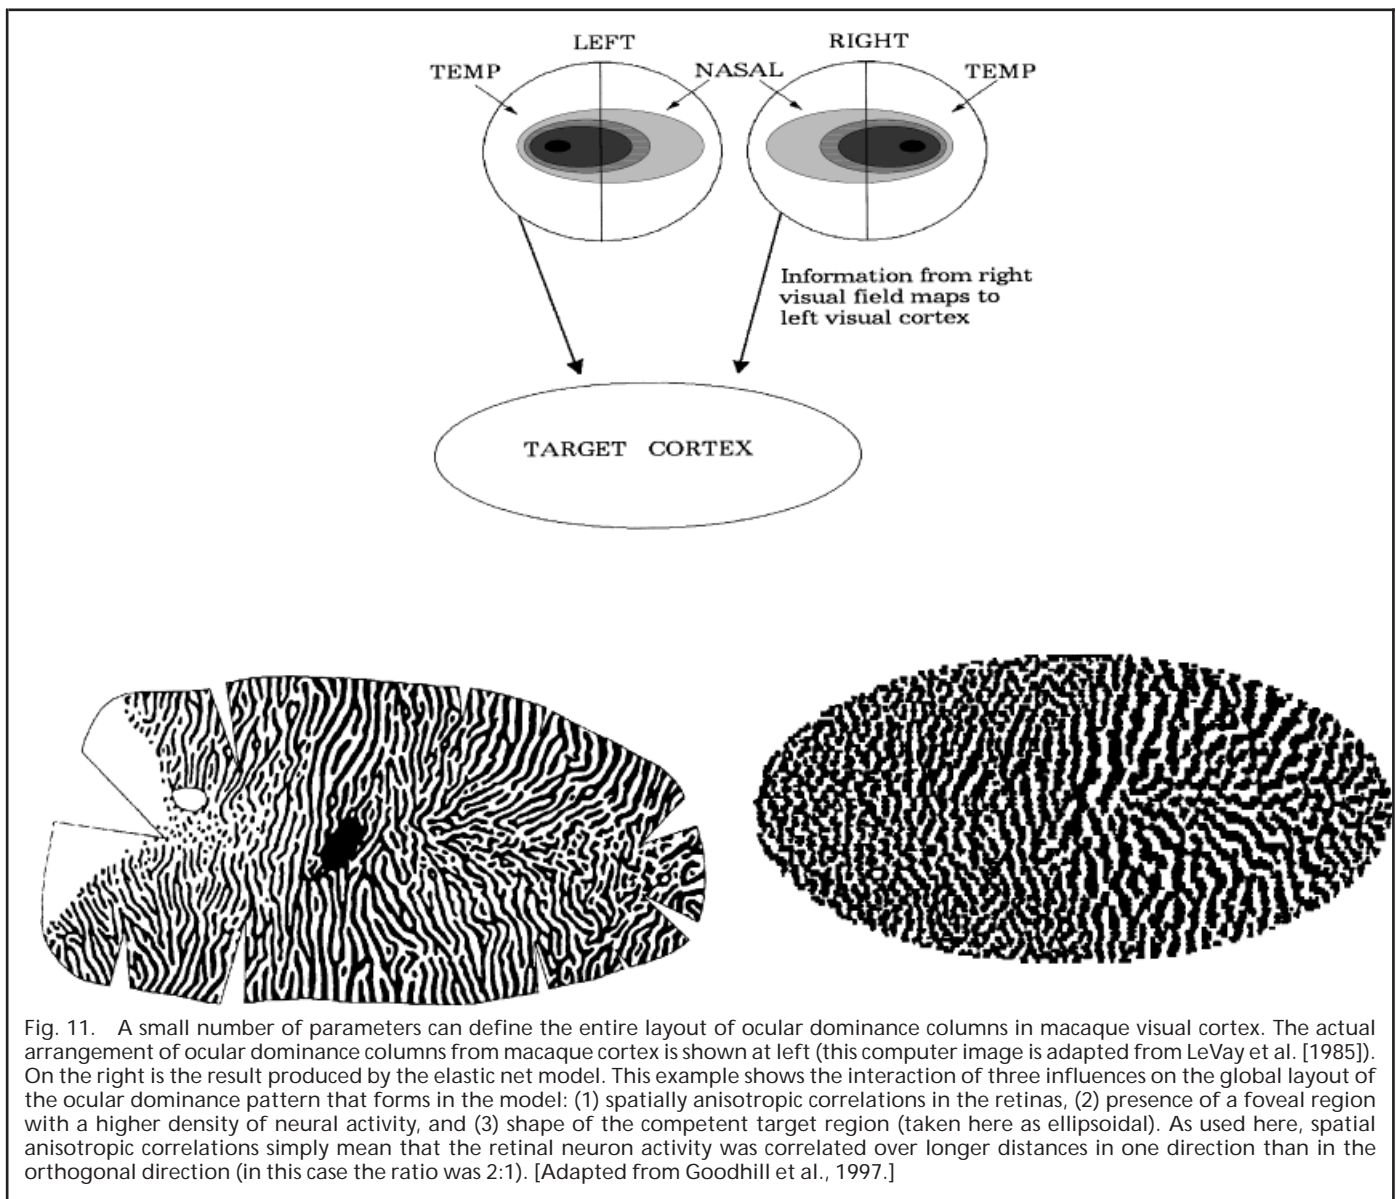


Fig. 11. A small number of parameters can define the entire layout of ocular dominance columns in macaque visual cortex. The actual arrangement of ocular dominance columns from macaque cortex is shown at left (this computer image is adapted from LeVay et al. [1985]). On the right is the result produced by the elastic net model. This example shows the interaction of three influences on the global layout of the ocular dominance pattern that forms in the model: (1) spatially anisotropic correlations in the retinas, (2) presence of a foveal region with a higher density of neural activity, and (3) shape of the competent target region (taken here as ellipsoidal). As used here, spatial anisotropic correlations simply mean that the retinal neuron activity was correlated over longer distances in one direction than in the orthogonal direction (in this case the ratio was 2:1). [Adapted from Goodhill et al., 1997.]

There is another reason why human cortical development is so protracted. Earlier we mentioned that artificial systems are typically made with a great deal of tailored information regarding the tasks they must perform. This makes a system efficient, but it also makes it inflexible. That is, a system designed for a particular task is ill suited to a dynamic environment in which demands are constantly altering. Suppose the post office adds letters to zip codes, for example. In that case, our letter sorter will fail and will have to be reprogrammed by hand. Thus, protracted human cortical development, while coming at the expense of increased vulnerability and parental investment, appears to be both a powerful means of constructing the representations underlying cognition and a flexible strategy for dealing with a dynamic environment. It appears there

are no other feasible means of obtaining flexibility other than maximizing the period of environmentally-guided representation construction.

Computational studies of development thus illustrate the consequences of the breakdown between learning and maturation, which the emerging view of neural development suggests. This breakdown leads to a uniquely powerful and flexible strategy for creating the representations underlying both cognitive skills and efficient learning. Development, then, is solving the fundamental problem facing a cognitive system—how to obtain the appropriate representations underlying complex skills. Viewed from this perspective, development is more aptly described as a special kind of learning, not statistical inference per se but a prior period of representation construction. Quartz and Sejnowski [1997] call this representation

construction “constructive learning.” Uncovering its nature and mechanisms is a major research direction in computational neurobiology.

The Argument for Early Intervention

Constructive learning puts a premium on the interaction between developing neural systems and a richly structured environment. Contrary to classical arguments regarding the poverty of environmental information, current computational research highlights the richness of environmental structure. As brain research has demonstrated, a reduction in the complexity of the early environment brings about a reduction in the complexity of neural structures. The most likely explanation now is that this deprivation causes a failure to grow rather than the failure to stabilize exuberant structures.

Because early experience is literally a form of environment-dependent learning that builds the representations underlying cognition and subsequent learning, deprivation places great risks on the future prospect of cognitive skill acquisition and learning in general. Without an appropriate set of representations, subsequent learning is impaired, although it remains an outstanding research question to assess the question of critical period phenomena for this capacity. The recent results of Eric Knudsen's [1998] research on the owl highlight the importance of early experience as a constructor of representations facilitating learning later in life. The optic tectum of barn owls contains a multimodal map of space. In particular, auditory-visual neurons in the optic tectum associate particular values of auditory spatial cues with locations in the visual field. This association is accomplished through matching the tuning of tectal neurons for interaural time differences with their visual receptive fields. During development but not adulthood, there is considerable plasticity in this system, allowing for a wide range of associations to be learned. When juvenile animals are fitted with goggles that shift the visual field, the resulting abnormal associations are learned. Knudsen demonstrated that the range of associations adult owls could learn is greatly expanded in those animals who had learned abnormal associations during development.

This is a striking demonstration of the role of early experience in constructing the representations that facilitate later learning. From a computational perspective it highlights the importance of separating learning into distinct problems: learning as statistical inference and the construction of an efficient set of representations that make statistical inference possible at all. Further, Knudsen's results demonstrated that the capacity to learn abnormal associations as an adult remains even though the associations had not been used for an extended period of time. While generalizing from such systems requires a great deal of circumspection, the lesson that early experience is a time of representation construction facilitating later learning and the acquisition of cognitive skills has important implications for intervention. Foremost, it suggests that the period of cortical maturation be viewed as a special kind of learning, one of "constructive learning" in which cognitive representations are constructed under the influence of environmental information. It also suggests that later remediation strategies that attempt to replicate the environmental

input of early experience may be more effective than ones based on traditional methods of learning as statistical inference. For example, a promising response to language learning impairment is based on acoustically modifying speech to resemble that of the early language environment. Modifying the acoustic properties of speech appears to facilitate the re-construction of cortical representations underlying the discrimination of phonetic elements [Merzenich et al., 1996]. Such an approach takes advantage of the plastic properties of cortex to address the fundamental representational issues underlying language learning impairment. Understanding this central stage of cognitive skill acquisition as a period of representation construction thus has far-reaching implications for both how we view early experience and how we approach the challenges facing at risk populations. ■

ACKNOWLEDGEMENTS

We thank Dr. Michael Wiest for helpful comments/criticisms. We also thank David Egelman and Richard King for comments on an earlier version of the manuscript.

REFERENCES

- Bellman RE. 1957. *Dynamic Programming*. Princeton: Princeton University Press.
- Berns GS, Dayan P, Sejnowski TJ. 1993. A correlative model for the development of disparity selectivity in visual cortex that depends on prenatal and postnatal phases. *Proceedings of the National Academy of Sciences U.S.A.* 90:8277-8281.
- Bienenstock EL, Cooper LN, Munro PW. 1982. Theory for the development of neuron selectivity: Orientation specificity and biocular interaction in visual cortex. *J Neuroscience* 2:32-48.
- Brown R. 1973. *A first language: The early stages*. Cambridge: Harvard University Press.
- Chomsky N. 1965. *Aspects of the theory of syntax*. Cambridge: MIT Press.
- Churchland PS, Sejnowski TJ. 1992. *The computational brain*. MIT Press.
- Egelman DM, Montague PR. 1998. A computational role for dopamine delivery in human decision-making. *J Cognitive Neuroscience*. In press.
- Elman JL, Bates EA, Johnson MH, Karmiloff-Smith A, Parisi D, Plunkett K. 1996. *Rethinking innateness: A connectionist perspective on development*. Cambridge: MIT Press.
- Fodor JA. 1981. *Representations*. Cambridge: MIT Press.
- Geman S, Bienenstock E, Doursat R. 1992. Neural networks and the bias/variance dilemma. *Neural Computation* 4:1-58.
- Gershenfeld N, Chuang IL. 1998. Quantum computing with molecules. *Sci Am* 278: 66-71.
- Gold EM. 1967. Language identification in the limit. *Information and Control* 278 10:447-474.

- Goodhill GJ, Bates KR, Montague PR. 1997. Influences on the global structure of cortical maps. *Proceedings of the Royal Society London B* 264:649-655.
- Hammer M. 1993. An identified neuron mediates the unconditioned stimulus in associative olfactory learning in honeybees. *Nature* 366: 59-63.
- Harder LD, Real L. 1987. Why are bumble bees risk averse? *Ecology* 68:1104-1108.
- Herrnstein RJ. 1991. Experiments on stable suboptimality in individual behavior. *Am Econ Rev Papers Proc* 83:360-364.
- Howard RA. 1960. *Dynamic programming and Markov processes*. Cambridge: MIT Press.
- Knudsen EI. 1998. Capacity for plasticity in the adult owl auditory system expanded by juvenile experience. *Science* 279:1531-1533.
- Lenneberg EH. 1967. *Biological foundations of language*. New York: Wiley.
- LeVay S, Connolly M, Houde J, Van Essen DC. 1985. The complete pattern of ocular dominance stripes in the striate cortex and visual field of the macaque monkey. *J Neurosci* 5:486-501.
- McCulloch WS, Pitts WH. 1943. A logical calculus of ideas immanent in nervous activity. *Bulletin of mathematical biophysics* 5:115-133.
- Merzenich MM, Jenkins WM, Johnston P, Schreiner C, Miller SL, Tallal P. 1996. Temporal processing deficits of language-learning impaired children ameliorated by training. *Science* 271:77-81.
- Miller KD, Keller JB, Stryker MP. 1989. Ocular dominance column development: Analysis and simulation. *Science* 245:605-615.
- Minsky M, Papert S. 1969. *Perceptrons*. Cambridge: MIT Press.
- Montague PR, Gally JA, Edelman GM. 1991. Spatial signaling in the development and function of neural connections. *Cerebral Cortex* 1:199-220.
- Montague PR, Dayan P, Nowlan SJ, Sejnowski TJ. 1993. Using aperiodic reinforcement for directed self-organization. In: *Advances in Neural Information Processing Systems*. Giles CL, Hanson SJ, Cowan JD, editors. San Mateo: Morgan Kaufman.
- Montague PR, Sejnowski TJ. 1994. The predictive brain: Temporal coincidence and temporal order in synaptic learning mechanisms. *Learning and Memory* 1:1-33.
- Montague PR, Dayan P, Person C, Sejnowski TJ. 1995. Bee foraging in uncertain environments using predictive Hebbian learning. *Nature* 376:725-728.
- Montague PR, Dayan P, Sejnowski TJ. 1996. A framework for mesencephalic dopamine systems based on predictive Hebbian learning. *J Neurosci* 16(5):1936-1947.
- O'Leary DDM. 1990. Do cortical areas emerge from a protocortex? *Trends Neurosci* 12:400-406.
- Purves D, White LE, Riddle DR. 1996. Is neural development Darwinian? *Trends Neurosci* 19:460-464.
- Pylshyn Z. 1984. *Computation and cognition*. Cambridge: MIT Press.
- Quartz SR, Dayan P, Montague PR, Sejnowski TJ. 1992. Expectation learning in the brain using diffuse ascending projections. *Soc Neuroscience. Abstract* 18:210.
- Quartz SR. 1993. Nativism, neural networks, and the plausibility of constructivism. *Cognition* 48:123-144.
- Quartz SR, Sejnowski TJ. 1997. The neural basis of cognitive devils: A constructivist mani-

- festo. Behavioral and Brain Sciences, 20:537–596.
- Quartz SR. 1999. The constructivist brain. Trends in Cognitive Sciences 3:48–57.
- Real L. 1991. Animal choice behavior and the evolution of cognitive architecture. Science 253:980–986.
- Rosenblatt, F. 1958. The perceptron: a probabilistic model for information storage and organization in the brain. Psychol Rev 65:386–408.
- Rumelhart DE, McClelland J, eds. 1986. Parallel Distributed Processing: Explorations in the Microstructure of Cognition. Cambridge: MIT Press.
- Schultz W. 1992. Activity of dopamine neurons in the behaving primate. Semin Neurosci 4:129–138.
- Schultz W, Romo R. 1990. Dopamine neurons of the monkey midbrain: contingencies of responses to stimuli eliciting immediate behavioral reactions. J Neurophysiol 63:607–624.
- Schultz W, Apicella P, Ljungberg T. 1993. Responses of dopamine neurons to reward and conditioned stimuli during successive steps of learning a delayed response task. J Neurosci 13:900–913.
- Schultz W, Dayan P, Montague PR. 1997. A neural substrate of prediction and reward. Science 275:1593–1599.
- Shannon C. 1949. The mathematical theory of communication. Urbana: University of Illinois Press.
- Sperry R. 1963. Chemoaffinity in the orderly growth of nerve fiber patterns and connections. Proc Nat Acad Sci USA 50:703–710.
- Sutton RS. 1988. Learning to predict by the method of temporal differences. Machine Learning 3:9–44.
- Sutton RS, Barto A. 1981. Toward a modern theory of adaptive networks: Expectation and prediction. Psychol Rev 88:135–170.
- Sutton, RS, Barto A. 1990. Time-derivative models of Pavlovian reinforcement. In: Learning and computational neuroscience. Gabriel M, Moore J, editors. Cambridge: MIT Press.
- Turing A. 1937. On computable numbers, with an application to the Entscheidungsproblem. Proceedings of the London Mathematical Society 42:230–265.