Neurobiological modeling: squeezing top down to meet bottom up

P Read Montague[1] **and Peter Dayan**[2]

July 30, 2000

[1]Center for Theoretical Neuroscience
Division of Neuroscience
Baylor College of Medicine, 1 Baylor Plaza
Houston, TX 77030
read@bohr.bcm.tmc.edu

[2]CBCL
Dept of Brain and Cog Sci
E25-210, MIT
Cambridge, MA 02139
dayan@ai.mit.edu

To Appear in: A Companion to Cognitive Science
eds. William Bechtel and George Graham

INTRODUCTION

A cartoon description of the goals of cognitive science and neuroscience might read respectively "how the mind works" and "how the brain works". In this caricature, there would seem to be little overlap in the vocabularies employed by each domain. The cartoon cognitive scientist could speak at length about decision-making and short-term memory in a relatively self-consistent manner without any need to make reference to the language of neuroscience. Likewise, the cartoon neuroscientist could provide an immense body of physical detail about the function of neurons, synapses and their component parts. She could even build models about how collections of neurons work together or even how they might have developed.

In both the cognitive and neural cases, such descriptions are inadequate; some phenomena will appear enduringly complicated, admitting no simple theory at a single level. In the cartoon scenarios, it is possible that many mind-like phenomena are not reducible in any strong way to descriptions enlisting interactions among components in the brain. There are a number of sophisticated arguments suggesting why such a reduction is or is not possible, likely, or fruitful - we do not enter here this debate. Instead, we will illustrate with practical examples the mileage that can be obtained from a kind of 'squeeze' approach to the problem of relating cognitive and neural descriptions. Complexities found by the cognitive scientist find natural explanation in the *neural substrate;* equivalently the natural theoretical context which is necessary for interpreting the neuroscientist's results, comes from hypothesized *purposes.*

This approach is fairly straightforward: take a consistent description of some behavioral or cognitive phenomenon, such as decision-making under simple choice tasks, along with a description of what may be relevant neural constraints and *squeeze.* The squeeze amounts to building a connection from the vocabulary in one to domain to the vocabulary in the other. To the extent that reality and practicality permit, this seat-of-the-pants heuristic for theory construction pushes top-

down and bottom-up constraints toward one another - hence, the title of this chapter. The stated approach clearly exposes our bias: we assume that concepts like attention, reward, memory, decision, *etc.* will find some mapping onto the descriptions of brain function.

The squeeze was first formalised by David Marr using ideas about computational equivalence. He showed how different notions of *computation* can be used in modeling cognitive and neural phenomena. In his view, cognitive tasks have to be specified precisely enough that one can write a computer program that demonstrably solves them. The action of collections of neurons that are believed to be necessary for this task must also be specified precisely enough so that one can write another computer program that captures their behavior. Crudely speaking, the squeeze is successful if the two programs are computationally equivalent.

We illustrate the general approach using two examples. The first concerns how animals learn to predict events in the world that have rewarding consequences and how they can use the predictions to control their actions in order to optimise their rewards. Here, there is an enormous amount of neural, behavioral and computational data that can be used to constrain the model at multiple levels. The second example concerns attention. Attention is less well understood than prediction at almost any level. We therefore explore at some depth just one aspect, namely what it could mean for different sets of neural inputs to control the output of a neuron according to different contents of attention. Before discussing these examples, we describe classes of neurobiological models.

METHODS OF NEUROBIOLOGICAL MODELING

There exist two classes of computational models in neurobiology. One concentrates on capturing closely the substantial information that modern neurobiology has garnered on the processes operating within and between single neurons – *e.g.* the way that current flows through dendritic

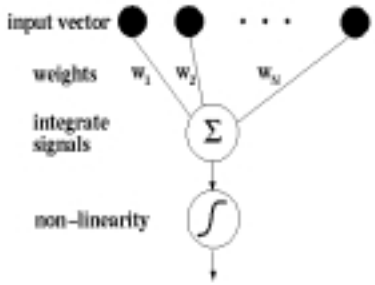

Figure 1: **Simulated neural units.** Artificial neural units take on a variety of forms, however, the general scheme is shown here. A neural unit collects information from other neural units along connections (synapses). Each connection is associated with a weight ($w_1, w_2, \cdots, w_N$) that can be changed according to preset rules. The unit typically has some integration step (here shown as a summation) and some kind of non-linearity before producing an output.

or axonal arbors, the effects of the many different sorts of ion channels, the ways that receptor molecules are influenced by neurotransmitters, etc. These models have been extremely important in understanding certain phenomena including the origin of action potentials, oscillations in membrane potentials, and the integrative function of dendrites. These models have been less illuminating at a systems level because neurobiological models are intrinsically so complicated. Moreover, even these detailed models omit large numbers of phenomena and there is no guarantee that the details on which they focus are the appropriate ones.

Computational models in the other class operate at the level of whole neural systems. In the best examples, the focus is on how collections of neurons cooperate to implement appropriate computations. The neural substrate is represented using artificial neurons that influence one another through modifiable synapses (see figure 1). The neural units are typically extremely simple representations of real neurons and ignore many biophysical details. These representations are then analyzed using mathematical techniques or simulated on digital computers or both.

The models differ in the kinds and number of details incorporated, depending on the problem at hand. For us, the squeeze is key – the models have to represent some known feature(s) of the neural substrate, however, they also have to be simple enough that they admit computational

analysis at another level. We shall see that the resulting models make both biological and behavioral predictions that can be separately tested.

PREDICTION & REWARD

The ability of an animal to anticipate future salient stimuli is a form of *prediction*: representations of sensory events must reliably reflect the likelihood, time, and magnitude of future important events such as food and danger. Experiments have established that both vertebrates and invertebrates are capable of making and using such predictions to modify their internal models of the world and to choose actions appropriately. The concept of prediction is a computational one: a system uses its current state and its past to guess at the likely future state of itself and the world. Hence, prediction can be defined outside the realm of a behaving animal. As we have described, however, behavioral experiments are used to assay this capacity in animals.

Unlike prediction, reward is a concept that is defined, not just assayed, by how an animal behaves. An organism, given multiple behavioral choices, will allocate some portion of its time or visits to each. *Reward* is assumed to be a latent quality of each behavioral choice. The magnitude of the reward content is defined by the relative proportion of time or visits allocated to the choice. This is a very behaviorist notion, however, it permits the easy quantification of many types of behavioral experiments; especially those involving decisions among alternatives. In past behaviorist traditions, these kinds of operational definitions became prohibitively restrictive in the classes of mechanistic explanations that they would permit.

We have been interested in understanding how animals learn to predict and also how they use rewards to (learn to) choose between actions. Substantial constraints are available from the theory of adaptive and optimal control (how systems of any sort can make predictions and choose appropriate actions), animal learning theory (the performance of animals in classical and instru-

mental conditioning tasks) and the neurobiology of reward (studies in electrical self-stimulation and the neuropharmacology of drugs of addiction). We begin by examining a simple neurobiological model which draws on both behavioral data and physiological data to hypothesize that a simple and direct representation of reward exists in the brains of honeybees and humans. These models suggest one way in which predictions are constructed and used by real brains.

Learning & decision-making in honeybees, humans, and networks

We have previously proposed that the representation of reward and expectations about future reward are constructed in part through systems of neurons situated in the midbrain and basal forebrain of mammals and in analogous structures in invertebrates. As discussed below, these groups of neurons are known to be associated with the processing of information about affective values, send extremely long axons to widespread regions of the brain, and deliver neuromodulatory substances including dopamine, serotonin, acetylcholine, and norepinephrine. By appealing to an established body of computational theory called the method of temporal differences (TD), we have constructed a theory in which activity patterns in the cerebral cortex make predictions about future rewarding events through their connections onto subcortical nuclei (figure 2; a **nucleus** is a group of neurons in the mammalian central nervous system; the analogous structures in the peripheral nervous system or in invertebrates are called **ganglia**).

In this TD model, the world for an animal in a conditioning paradigm is cast in simplified form as a Markov decision problem. Here, there are states, which are signalled by stimuli like lights and tones or positions in a maze; transitions between states, like transitions in the maze; and actions like lever presses, which can affect rewards directly and can also affect transitions between states, as in the maze. The task for the animal is to choose appropriate actions that maximise rewards over the long term. The challenge is that actions can have *delayed* affective consequences – *eg* in a maze, a poor early move can make for very long paths, but this fact may only be apparent when

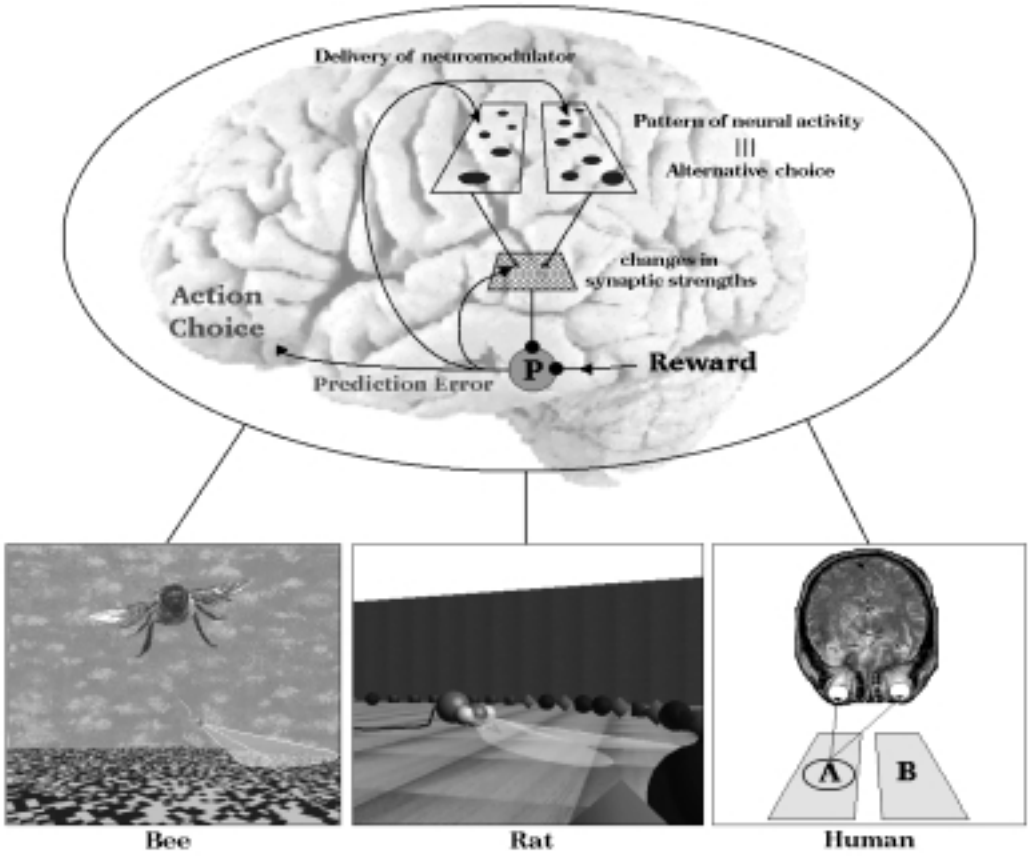

Figure 2: **Neural Representation of reward and expected reward.** Activity patterns in the cortex construct expectations of the time and magnitude of future rewarding events and send this information through highly convergent connections to a group of modulatory neurons (labeled P). Also converging at neuron P is direct information about rewarding events such as the intake of food or pleasurable sensations. Neuron P is a linear unit, and, under the assumption that the cortical activity arrives at P in the form of a temporal derivative, the fluctuating output of P represents ongoing errors in the predicted amount of total future reward and the amount actually received. In the absence of direct reward input, the output of P is used to bias actions. In the presence of direct reward input, the strength of synaptic contacts is modified and this updates the organism's model of the world. This arrangement has been used to model the choices made by flying bees, rats moving in a two-dimensional arena, and humans.

the animal finally reaches the goal. A standard way to solve such decision problems is called dynamic programming and comes from the field of optimal control. In dynamic programming, a system adopts a policy, which is just some consistent way of choosing actions at states, learns to evaluate states under this policy, and then improves the policy on the basis of these evaluations. The value of a state under a policy is the average amount of reward that the system can predict it will receive if it starts in that state and follows its policy. The policy can be improved by choosing actions that lead to high rather than low value states.

TD offers a method of performing dynamic programming in an approximate manner. The values of states are estimated using weights. These estimates can be improved by measuring the degree to which they are inconsistent with each other and with the delivery of rewards and punishments from the world – *eg* if two places in a maze are one step apart, then the estimate of the distance to the goal from one should be one more than the estimate from the other. Any inconsistency is called a *prediction error* and is what is used to drive learning. In the model, the fluctuating outputs of neurons in subcortical nuclei (denoted P in figure 2) represent this prediction error. Hence, the fluctuating delivery of neuromodulator carries information about these errors to widespread target regions. If the inconsistency is in estimation of reward, then increases in neuromodulator delivery literally mean 'things are better than expected' and decreases in neuromodulator delivery mean 'things are worse than expected'. In fact this same error signal can have two roles: (1) training the predictions to be correct - any net bias in the fluctuations indicates an error in the expectations; and (2) choosing and training the choice of good actions - if the expectations are correct on average, then a positive fluctuation indicates that the associated action may be better than average. This second role for the signal implements the approximation to the way that policies are improved in dynamic programming.

As indicated in figure 2, this basic theory has been applied to bee foraging over fields of flowers, rats foraging in a two-dimensional arena, and human decision-making on a simple card choice

task. In the case of the bees and rats, a virtual world was constructed using computers and a virtual rat and virtual bee were permitted to move about in these worlds. This methodology provides a fruitful experimental testground for the behaviors that result from the operation of biological learning rules under the influence of some representation of the environment and sensory apparatus of the animal.

These models make testable predictions about the behavior expected from the bee, rat, and human. The models incorporate biological assumptions, hence, they also offer predictions about the behavior of neurons. In these examples, the unifying factor is a well understood computational theory that permits us to assign computational functions to specific biological constraints. The fact that the behavior of the models on foraging, learning, and decision-making tasks matches the behavior of the appropriate animal provides further support for the approach. As computing technology evolves, it may become possible to use large-scale simulations to make testable predictions about the interaction of multiple organisms in a simulated world.

These experiments primarily addressed the issue of behavioral choice. We have used exactly the same model to study the physiological behavior of neurons that deliver to their targets a neuromodulator called dopamine during the course of experiments that probe the way that animals come to predict events with rewarding consequences. We will describe this in some detail below.

Dopamine and Reward

Dopamine (DA) is a neuromodulator that has long been associated with reward processing. In the mammalian brain, dopamine is produced and distributed by nuclei located in the midbrain. One of the dopamine nuclei is called the ventral tegmental area (VTA). The neurons in the VTA send axons to brain structures known to be involved in reward processing *e.g.* one important site is the nucleus accumbens. As outlined below, three major lines of experimental evidence suggest

that the VTA and dopamine's action in general are involved in reward processing.

First, drugs like amphetamines, which are known to be addictive, are dopamine reuptake inhibitors *i.e.* they prolong the action of dopamine near the sites where it is released. Second, neural pathways connected to the VTA are among the best targets for electrical self-stimulation experiments. In these experiments, rats press bars which delivers an electrical current at the site of the electrode. The rats choose this self-stimulation over food and sex. Third, agents which block the action of dopamine on dopamine receptors lead to extinction behavior in instrumental conditioning: animals that press a bar to get a reward will stop pressing the bar when given the full reward but under measured doses of haloperidol (a dopamine receptor blocker), as if they were no longer being rewarded. In spite of these very concrete results suggesting a role for dopamine in reward processing, the actual relationship between dopamine release and reward delivery is complicated – *e.g.* in many cases, the delivery of reward to an animal is not followed by any increase in the delivery of dopamine.

Dopamine delivery and prediction

Another main source of neurobiological constraints on this issue is a series of experiments performed by Wolfram Schultz and his colleagues (*eg* Schultz, 1992). They characterised the electrophysiological properties of dopamine neurons in the monkey VTA and substantia nigra (another dopamine nucleus involved in motor acts). These workers recorded from dopamine neurons while animals were learning and performing simple behavioral tasks for reward (apple juice). These workers found a subset of dopamine neurons in the VTA whose activity clearly relates to reward processing, but not in a simple fashion.

In one task, monkeys were presented with a light which signaled the delivery of reward (apple juice) provided that the animal performed an action within a pre-specified amount of time. In the

context of this simple task, the light consistently predicts that reward will be delivered consequent on the action. Through training, the animals' reaction times for the action decrease and they clearly use the onset of the light stimulus as a cue that reward will follow if they act correctly. These statements all rely on behavioral assessments; however, Schultz and colleagues found that the dopamine neurons changed their firing rates in ways that consistently related to the learning displayed by the animals. A number of consistent features emerged in these studies:

- Early in training (naive animal), most dopamine neurons increased their firing rate when reward was delivered and showed no change in firing rate upon presentation of the light.
- Later in training (trained animal), most dopamine neurons increased their firing rate when the light came on and showed no change in firing rate upon delivery of reward.
- If two sensory cues consistently precede delivery of reward, then changes in the dopamine neurons' firing rate shift from the reward to the *earliest* consistent predictor of reward.

Remarkably, these neurobiological data mirror the computational requirements of a prediction error signal as specified in a theory based on the method of temporal differences (TD). As we described briefly above, our model for neuromodulatory control of learning and action choice fits into a temporal-difference framework.The striking fact is that the *prediction error* signal in TD has precisely many of the characteristics listed above for the dopamine neurons:

- early on in learning, when the computational agent does not know that a cue predicts the delivery of reward, it is surprised by the delivery of reward, *ie* there is a substantial inconsistency between its predictions and the outcome, and so there is a substantial positive prediction error (the increase in firing upon delivery of reward).
- Once the agent knows that the cue predicts the reward, then the reward itself is expected

and leads to no prediction error (after learning, there is no change in firing upon delivery of reward).

- When the predictive sensory cue appears, it was not predicted, and there is a prediction error consequent on the cue (after learning, there is an increase in firing after the onset of the predictive sensory cue).

- For two sensory cues that both predict reward, a TD model learns that the earliest cue can itself predict the reward that the later cue predicts, and so itself attracts all the net prediction error.

Of course, there are a number of problems with the model and areas in which we made arbitrary choices, have gone beyond the available data, or have brushed aside genuine complexities that might be inconvenient for the model. It is from these problems that behavioral and neurobiological experiments naturally arise. Notable concerns are:

- The basal firing rates of the dopamine cells are low (around 1 Hz) suggesting that increases and decreases in firing rate cannot carry the same amount of detailed information about prediction errors. This fact, along with other theoretical and experimental observations, suggest that there may be an opponent system to dopamine that constructs and delivers information about *punishments* and also *withheld rewards*. There is reason to believe that this might be one of the roles of the serotonin system.

- Through their widespread axons, the dopamine cells distribute information on prediction error to widespread structures. In its simplest form, the model requires dopamine to control synaptic plasticity for synapses that construct the predictions. The location(s) of the memories are also unclear. We have suggested the amygdala as one likely site, based on its pivotal position in the limbic system and evidence that interfering with the amygdala inter-

feres with forms of secondary conditioning. This phenomenon probes the affective values associated with stimuli.

- Some simple learning paradigms are best described in terms of attention: the animal allocates more or less attention to particular stimuli based on its experience with them. This differential allocation results in more or less learning accruing to those stimuli during learning. Possible mechanisms behind selective attention are discussed in the next section but have not been incorporated in the models described above.

These inadequacies nonwithstanding, this model operates at four different levels of description. The temporal-difference model matched animal learning data, however, it also implements known techniques of optimal control: most notably the engineering technique of dynamic programming. This link enables the squeeze – any system implementing an algorithm like temporal-differences can reliably learn to perform appropriate actions that can even require complex sequences of choices. The key signal for temporal-differences is the *prediction error*. Schultz's data strongly support the hypothesis that this error is being carried by the fluctuating output of dopamine neurons. Assessing the appropriateness of behavior lies in the realm of the ethologists, providing the fourth descriptive level.

ATTENTION

One critical element missing from the above discussion is attention: how various sensory cues and prediction errors are marked as being more or less salient. Our examples above have not provided for such effects. The concept of attention originated in the vocabulary of psychology, but it has eluded being made computationally, psychologically or neurobiologically crisp. One category of experimental observation is that, on presenting the same set of stimuli to an animal on different occasions, different stimuli seem to be favored in terms of reaching consciousness,

attracting learning, controlling behavior, and even determining the activities of neurons. Other attentional phenomena such as orienting behavior are not well characterised by this description, but there is no reason to expect that everything we call attention should comprise a natural class.

Unlike the case of prediction, it is hard to specify a precise computational problem that attention is solving. One popular possibility is that attention is important because the way that the neural substrate performs computations is such that it gets confused if multiple cues are processed simultaneously. From the experimental end, attentional effects have been probed in various ways. One is assessing changes in neuronal activity and local blood flow changes in identifiable brain regions. As described and probed by Michael Posner and his colleagues, attentional effects, as measured by brain imaging technology, are not associated with a single brain area nor the whole brain. These workers have sought to determine how attention influences brain activity: either increasing or decreasing it depending on the tasks to which a subject is put.

In addition to brain imaging experiments, detailed electrical recordings from the brains of alert primates indicate that attentional effects, as assayed by changes in behavior or perceptual thresholds, correlate with dramatic changes in the electrical activity of identified neurons in areas of cerebral cortex devoted to vision. The mechanisms and constraints that permit this kind of control of neural activity are in general unknown. These experiments show clearly the net effect that certain synapses become ineffective and/or others become augmented, but they do not distinguish between radically different types of mechanisms.

Generally, theories of attention are significantly underconstrained by the available data at any level of description. Attention is therefore a class of phenomena in which interaction between models and experiment can be particularly crucial in going beyond the phenomenology to the neural mechanisms – they solicit the very evidence that would make squeezing effective. Below, we suggest both the spatial scale and neural loci through which attentional effects could emerge

in a working brain. In all our discussion, we have in mind effects that most likely act at the level of the cerebral cortex and basal ganglia of mammals, particularly primates.

Spatial scale of attentional mechanisms in the brain

In order to investigate how neural tissue in the cerebral cortex could implement constraints that resulted in attentional effects, a decision about scale must be made: do attentional effects emerge at the level of brain regions, neural circuits, single neurons, groups of synapses, single synapses, or perhaps at an even smaller scale ?

We first inquire about the smallest scale at which neural activity could be modulated. It is already clear from brain imaging and electrophysiological studies that changes in the activity of groups of neurons correlate with behaviorally assessed attentional effects. It is therefore reasonable to assume that the activities of single neurons are similarly affected. We suggest here that the physical substrates of attentional effects in the brain could exist at a smaller scale still: the single synapse.

In the mammalian cortex, synapses average about 1 micron in diameter and their density falls somewhere between 0.7 and 1.2 billion synapses per cubic millimeter - this amounts to about 1 synapse in every cubic micron of tissue. That is a billion connections in a region about the size of a match-head. Since the synaptic densities are so high, most notions of 'nearby' include a large number of synapses. It is not a sufficient framing of the problem to assert simply that the function of single synapses is modulated during attentional effects. A number of important points remain: (1) Where does the information originate that modulates the function of a synapse that receives this information? (2) How is the information delivered to the synapse in question? (3) What is the postulated effect on the synapse?

We describe a set of answers to these questions below. We do not delve as deeply into question 1; rather, we assume that some region or regions of cortex become specialized to construct

and distribute the salience of various sensory events. One might expect prefrontal regions to be particularly involved in the *voluntary* control of attention. However, one can also imagine that more local, automatic mechanisms of suppression and enhancement of neural activity are just as important for attentional processing — *e.g.* automatic segmentation of a visual scene that permits recognition of parts of the visual scene. In this latter case, it is unlikely that modular regions of prefrontal cortex would be the final arbiters.

Embedded in question 2 are subtle issues of how the information is coded and the physical substrates used to transmit it to the synapse. There is a wealth of possibilities here – ranging from changes to synaptic function caused by the release and vascular distribution of some humoral factor, to the rapid delivery of a neuromodulator like norephinephrine through activity in axons originating in the locus coruleus (a midbrain nucleus that distributes norepinephrine to the cerebral cortex and other structures). In both cases, synapses can be affected according to a volume effect: hormones and neuromodulators both act at a distance from their sites of initial distribution.

These mechanisms for distributing attentional information share the problem that the delivery mechanism communicates information to very many synapses, therefore, synapses which receive the information need some mechanism to assess whether they should be suppressed or enhanced on its basis. Temporal correlation between the electrical activity of the synapse and delivery of the 'attentional signal' is one way to make this assessment: short-term fluctuations in the 'bottom-up' synaptic activity associated with the object of attention would filter through to the areas controlling attention, and would then be reflected in commensurate short-term fluctuations in the 'top-down' attentional signal they broadcast. The resulting correlations between these two signals are straightforward to measure. The neuromodulators could also be targeted more precisely to particular synapses: the synapses may possess the right combination of receptors making them more sensitive to a particular neuromodulator or the attentional signal could be targeted through a fixed anatomical connection. In any case, comparatively few cells deliver neuromodulators to

enormous numbers of synapses, hence, precision in the delivery of the information is lost and must be recovered by some other mechanistic trick - we have identified two possible general schemes.

Below, we outline in detail how specific synapses and / or cells could be selected by some attentional signal. By making the assumption that the synapses is the smallest scale at which control of neural function can be exercised, we arrive at two different physical schemes for how an attentional signal could be constructed and used in real brains.

The Resource Consumption Principle: attentional selection in volumes of neural tissue

It is well accepted that the synapse is a junction that passes information from one neuron to another neuron or to volumes of neural tissue (white structures in 3 B). In a real nervous system, information travels from one neuron to another in the form of electrical impulses called action potentials. Action potentials travel along thin branched fibers called axons that terminate onto other neurons through enlarged endings called synaptic terminals or boutons. Information about action potential arrival at a synaptic terminal is passed to the recipient neuron through the rapid release and diffusion of chemicals from the synaptic terminal. This transfer of information is rapid because the gap between the end of the axon and the next neuron is very small ($\approx$ 20 billionths of a meter); hence, diffusion of the chemical rapidly influences the next neuron. There is one drawback: the arrival of the impulse causes calcium to flow very rapidly into the synaptic terminal, and, without calcium entry, the terminal will not release its neurotransmitter, *i.e.*, no information is transmitted to the receiving neuron. Normally, the level of calcium inside the terminal is extremely low, but when the action potential (electrical impulse) arrives, the calcium flows in through channels and the levels inside the terminal increase rapidly. This flow of calcium into the terminal is absolutely necessary for the terminal to function - impulses invading a terminal in a region of tissue without calcium will not be transmitted to the next neuron. In this context, the

calcium present outside synapses acts like a limited, shared resource that synapses must obtain in order to operate. One of us (Montague) has suggested that the above facts about calcium and neural transmission amount to an abstract processing principle that permits volumes of brain tissue to select a set of functioning synapses. This idea is called the resource consumption principle (RCP) (see figure 3).

The resource consumption principle appeals to a fluid metaphor when treating the function of synapses. In this theory, there are two classes of fluids. The first class, called the resource (figure 3), must be moved from outside to inside a synapse in order for the synapse to function. The second class of fluid is envisioned as a composite of many separate fluids, each representing different kinds of information delivered throughout a volume of neural tissue through the release of different types of neurotransmitters. The collection of input signals is treated as a vector and called a **key.** As stated, each component of the key is treated as a fluid available homogeneously throughout a volume of neural tissue.

Synapses are pre-equipped with receptors that recognize different combinations of neurotransmitter fluids. The combination of receptors on each synapse is also envisioned as a vector and is called a combination **lock.** Each time a molecule of neurotransmitter binds a receptor, a quantity of the resource is moved from the outside to the inside of the synapse where the binding takes place. The scheme is: (1) key is presented to a volume of neural tissue, (2) the key matches some synapses' locks better than others and causes the matching synapses to consume more of the shared resource, (3) those synapses that have consumed the most resource tend to function (transmit) (4) at a slower timescale, the resource is replenished in the surrounding volume of tissue (see figure 4).

Since the resource is in limited supply, synapses that consume it do so at the expense of neighboring synapses in the local volume of tissue. In this fashion, there is enough resource for only

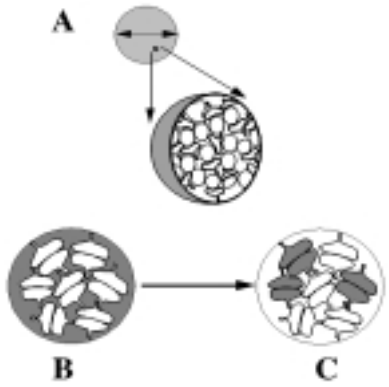

Figure 3: **The Resource Consumption Principle (RCP). A.** Synapses immersed in a homogeneously distributed resource fluid. Subvolume of this region is shown as inset **B.** In order to function, the resource (gray) must be moved into a synapse (white structures). In this particular example, the space enclosed by the synapses represents over 85% of the volume, therefore, the resource is in limited supply and is consumable by a small fraction of synapses in the volume (as shown in **C**).

a subset of synapses to function (transmit) and so a fierce competition for resource is set up because of the way that volumes of tissue are organized and the dependence on resource. In a direct sense, the volume of neural tissue 'attends' to those synapses (locks) that have successfully consumed the resource. This description omits detailed considerations about the dynamics of these processes, but the general idea is communicated.

In these proposed mechanisms, attention to one set of synapses over another is granted as a result of a competition directly through the tissue space. Long distance tissue volumes interact in the same manner through longer-range axonal connections that contribute synaptic terminals into common volumes. The difficult question not answered by this description is how an attentional signal is broadcast widely enough so that this local competition for resource can automatically decide on a set of working synapses.

The capacity for some set of synapses to match a particular key must pre-exist within the tissue *before* the key is experienced. It is intriguing to ask whether basic organizational properties of neural tissue define the primitives out of which various forms of attention are constructed. In the

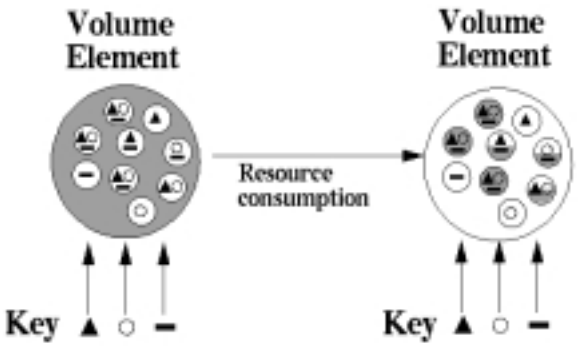

Figure 4: **Matching of locks and keys through a fluid-like connectivity.** The collection of signals that impinge on the volume is represented as a vector $\mathbf{v} = (v_1, v_2, v_3, \cdots)$ where each component of the vector v_i represents a different signal type. We call this collection of signals the **key**. Each component v_i of the key is presented homogeneously throughout the volume. In this fashion, each component of the key acts like a separate fluid homogeneously distributed throughout the volume. Different synapses are sensitive to different keys because they possess various combinations of *receptors* each sensitive to one type of fluid ,*i.e.,* one component of $\mathbf{v}$. The collection of receptors on each synapse can also be envisioned as a vector and is called a combination **lock**. As shown, presentation of a particular key matches the combination locks of some synapses better than others causing the matching synapses to consume resource. At some preset time after presentation of the key, the probability that a synapse is activated is a function of the amount of resource (gray) it has captured. The resource allows the probability that a synapse is activated (ON) to be a function of the distribution of this required resource 'fluid'. The main idea is that those synapses marked as ON define the attention of the engine ,*i.e.,* those synapses that work under the influence of the current key.

case of the resource consumption principle, attentional limitations result first and foremost from a commodity that is literally in limited supply and shared throughout a volume of tissue. There are certainly other strong possibilities for attentional effects that would exist at the level of entire circuits of neurons. We discuss one possibility below.

Attentional selection in recurrent reverberatory loops

The resource consumption principle (RCP) locates attentional suppression of synapses directly as a property of the synapses themselves in conjunction with their three dimensional neighborhood. However, there is a completely different way in which suppression or enhancement could take effect. The key is to consider the net influences of synapses on the postsynaptic cells. Individual synapses that are directly suppressed under the RCP could actually be functioning; however, the postsynaptic cells might not provide a faithful report of their activity *i.e.* the cells are effectively inhibited or just not enhanced. It is unlikely that this inhibition is direct, for instance, there are no long-range GABA-ergic inhibitory axons. The inhibition could however be indirect.

One model for this indirect inhibition or excitation comes from the notion of *self-excitatory loops* through the striatum and the thalamus. These are brain regions involved in motor control that have a critical dependence on dopamine delivery for their proper function. Anatomical data suggest that neurons in somatosensory and motor regions of cerebral cortex project to small clusters of neurons in the striatum, *i.e.* there is *divergence* in the cortex to striatum connections, since neurons from one area of the cortex project to many striatal clusters, each of which also receives information from other cortical areas. Conversely, these clusters send information indirectly back to those same areas of cortex that generated them, *i.e.* there is *convergence* in the striatum to cortex connections. The main neurons in the striatum are inhibitory: they may inhibit each other to a degree partly controlled by dopamine levels (although this is somewhat controversial). The primary source of this dopamine is the substantia nigra (mentioned above).

Based on this suggestive anatomy and the physiology of the inhibitory cells in the striatum, it has been proposed that different motor actions are represented in the cortex and these compete for control of the motor effectors, *e.g.* the limbs, at the level of the striatum. The mutual inhibition in the striatum provides a competitive mechanism by which actions that are incompatible with each other compete, for example, extensor and flexor muscles for the same movement are not simultaneously activated. The effect of winning the competition is that the excitatory loop involving the neurons representing that action in the cortex and the striatum is 'opened' thereby encouraging the relevant action to be performed. Cells in motor cortex representing actions that do not win the competition are not boosted in this manner, and so their actions are not performed.

As pointed out by Ann Graybiel and her colleagues at MIT, the anatomy of this system forces relevant information from the cerebral cortex to compete for control of motor output at the level of the striatum. Moreover, Graybiel notes that the striatum is the target of a large number of neuromodulatory systems, allowing many kinds of information to influence the competition for motor output.

One can conceive of attention directed to one stimulus out of a collection in the same way as attention directed to one action out of a collection – indeed, Neil Swerdlow and George Koob have suggested that the nucleus accumbens or ventral striatum might play the same role for stimuli that the dorsal striatum plays for actions. Cells representing different stimuli in cortex would therefore compete, not at the level of their synapses or even directly with each other, but rather at the level of the striatum, competing to gain access to an excitatory loop. Synapses within the cortex would then appear to be more or less effective according to the victory or defeat of their postsynaptic cells in the striatal competition.

There is a diverse cocktail of neuromodulator receptors in the nucleus accumbens. These molecular signaling pathways would then be available to bias the competition in the light of predictions

of the presence or absence of rewards. There is already direct evidence that dopamine and serotonin exert influence over phenomena in conditioning that have attracted explanations in terms of attention. Also, patients with schizophrenia, a disease that is believed to involve misfeasance in the dopamine system, show symptoms suggesting deficits in their attention.

Similar issues about other means of controlling the focus of attention apply to this model as to the RCP, except that there is now a defined structure on which any manipulations must take their effect. The connections between prefrontal areas and the basal ganglia could mediate explicit attentional control – we imagine the system learning how to pay attention to particular stimuli by learning what outputs to provide to which cell clusters in the striatum.

SUMMARY

We have described both the intents and the processes of neurobiological modeling. The best neurobiological models are computationally explicit enough to perform complex tasks that animals can clearly perform themselves, and close enough to slight abstractions of the neurobiological data that they can be constrained (and therefore falsified) by neural findings. We have described the 'squeeze' that results in the best circumstances. All models of interesting behaviors are radically underconstrained at any given level – by taking on results at multiple different levels, even if they are couched using different vocabularies, the models become much better specified.

We showed two examples of modeling – one rather better worked out than the other. Learning to predict rewards and act appropriately on the basis of those predictions is highly adaptive. Although there are many ways that a vertebrate might go about making these predictions, we pointed to the evidence that they do it in a particular way, by showing the relationship between the firing of dopamine cells in reward related areas of the brain and a key signal in one class of prediction algorithms called temporal difference algorithms. This link not only provided a

rationale behind the otherwise rather perplexing behavior of the dopamine cells in response to rewards, but also suggested why cells in two different dopamine projecting regions might fire the same way, and has led to suggestions for a number of other experiments. The same model applies to the selection and learned selection of good actions in behavioral choice tasks that have been applied from honey-bees to humans. Of course the model is quite abstract and quite incomplete in various important ways – however it shows how one can take advantage of the squeeze.

The other topic for modeling is attention. We were able to progress less far in this direction, mostly because the phenomena are not delimited so well in any of the different vocabularies – it is just not yet clear exactly what features a model should have. To show that modeling is still possible even under these circumstances, we focused on a particular notion, which is common across all models of attention, that synapses change their net efficacies as a result of attentional processing. According to one model, these efficacies change directly. Synapses are constantly competing with their literal neighbors for a shared resource, and the competition can be biased through a number of different message systems in the brain which carry information relevant to the focus of attention. According to the other model, the relevant synapses do not change their efficacies directly, but rather the cells that they influence are part of self-excitatory loops through the striatum. In this latter model, competition happens in the basal ganglia rather than locally in the cortex, and biasing happens there too. The models are not mutually exclusive.

None of these models is complete – they pose more questions than they answer. Nevertheless, neurobiological modeling offers a powerful complement to existing experimental techniques in neurobiology. Our capacity to collect data has far outstripped our capacity to construct appropriate theoretical contexts for them. Experiments are in any case intrinsically theory driven. Neurobiological modeling amounts to specifying those theories precisely enough that they can be programmed or analysed, and specifying them so that they respect the data not only from one level of inquiry but from them all.

Acknowledgements - This work is supported in part by NIMH grants R01MH52797 (to PRM), R29MH5554101 (to PD), and the Center for Theoretical Neuroscience at Baylor College of Medicine.

References

Barto, AG (1994) Reinforcement learning control. **Current Opinion in Neurobiology** 4(6):888-893.

Graybiel AM. (1995) The basal ganglia. **Trends in Neurosciences** 18(2):60-62.

Montague, PR, Dayan, P, Person, C, Sejnowski, TJ (1995) Bee foraging in uncertain environments using predictive hebbian learning [see comments]. **Nature** 377:725-728.

Montague, PR, Dayan, P, Sejnowski, TJ (1996) A Framework for Mesencephalic Dopamine Systems Based on Predictive Hebbian Learning. **Journal of Neuroscience** 16(5):1936-1947.

Pulvirenti, L, Koob, GF (1994) Dopamine receptor agonists, partial agonists and psychostimulant addiction. **Trends in Pharmacological Sciences** 15(10):374-379.

Schultz, W (1992). Activity of dopamine neurons in the behaving primate. *Seminars in the Neurosciences*, **4**, 129-138.

Swerdlow, NR & Koob, GF (1987). Dopamine, schizophrenia, mania, and depression: Toward a unified hypothesis of cortico-striato-pallido-thalamic function. *Behavioral and Brain Sciences*, **10**, 197-245.

Suggested readings

Churchland, PS & Sejnowski (1992). *The Computational Brain*. Cambridge, MA: MIT Press.

Dayan P. (1994) Computational modelling. **Current Opinion in Neurobiology** 4(2):212-217.

Kandel, ER, Hawkins, RD (1992) The biological basis of learning and individuality. **Scientific American** 267(3):78-86.

Montague, PR (1996) The Resource Consumption Principle: attention and memory in volumes of neural tissue. **Proceedings of the National Academy of Science (USA)** 93(8):3619-3623.