

Banishing the Homunculus: Making Working Memory Work

Thomas E. Hazy, Michael J. Frank, and Randall C. O'Reilly*

Department of Psychology

University of Colorado Boulder

345 UCB

Boulder, CO 80309

*oreilly@psych.colorado.edu

Abstract:

The prefrontal cortex (PFC) has long been thought to subserve both working memory and “executive” function, but the mechanistic basis of their integrated function has remained poorly understood, often amounting to a homunculus. This paper reviews the progress in our lab and others pursuing a long-term research agenda to deconstruct this homunculus by elucidating the precise computational and neural mechanisms underlying these phenomena. We outline six key functional demands underlying working memory, and then describe the current state of our computational model of the PFC and associated systems in the basal ganglia (BG). The model, called PBWM (prefrontal-cortex, basal-ganglia working memory model), relies on actively maintained representations in the PFC, which are dynamically updated/gated by the BG. It is capable of developing human-like performance largely on its own by taking advantage of powerful reinforcement learning mechanisms, based on the midbrain dopaminergic system and its activation via the BG and amygdala. These learning mechanisms enable the model to learn to control both itself and other brain areas in a strategic, task-appropriate manner. The model can learn challenging working memory tasks, and has been corroborated by several important empirical studies.

Introduction

This article reviews an ongoing research agenda in our lab and others that is attempting to elucidate the precise computational and neural mechanisms underlying working memory and “executive” function. Our approach represents an attempt to understand these phenomena in terms of a set of biologically based, computational mechanisms. This approach has resulted in the identification of a core set of six functional demands that collectively help define the fundamental nature of working memory from a neuro-mechanistic

Draft Manuscript: Do not cite or quote without permission.
Supported by ONR grant N00014-03-1-0428, and NIH grants MH069597 and MH64445.

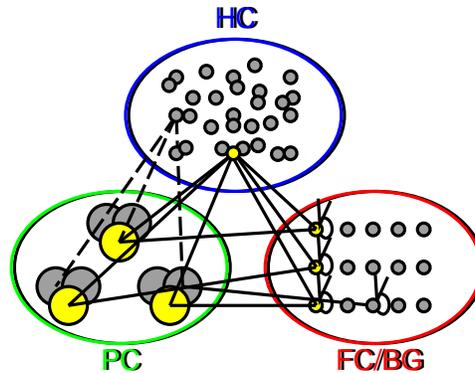


Figure 1: Tripartite cognitive architecture defined in terms of different computational tradeoffs associated with Posterior Cortex (PC), Hippocampus (HC) and Frontal Cortex/Basal Ganglia (FC/BG) (with motor frontal cortex constituting a blend between FC and PC specializations). Large overlapping circles in PC represent overlapping distributed representations used to encode semantic and perceptual information. Small separated circles in HC represent sparse, pattern-separated representations used to rapidly encode (“bind”) entire patterns of information across cortex while minimizing interference. Isolated, self-connected representations in FC represent isolated stripes (columns) of neurons capable of sustained firing (i.e., active maintenance or working memory). The basal ganglia also play a critical role in the FC system by modulating (“gating”) activations there based on learned reinforcement history.

perspective. Our proposed mechanisms for dealing with these functional demands collectively explain many of the same phenomena as traditional working memory constructs, but in a manner that contrasts with them in important ways and does so in a comprehensive, integrated way.

The overall format for the article is as follows. After a brief introduction of our approach to working memory in terms of developing a biologically-based architecture for understanding human cognition, we then describe the current version of our computational model of the prefrontal cortex (PFC) and basal ganglia (BG) in working memory (PBWM), with special emphasis on six key *functional demands* underlying working memory. To try to make things as transparent as possible, we describe both the functional demands and the model itself in the context of a concrete working memory task. We then review some of the empirical data that has tested predictions of our model, and then outline our research trajectory that is attempting to simulate many of the most important task paradigms of working memory and executive function in a single instantiation of a comprehensive model built around the core PBWM mechanisms. Finally, we discuss some overall implications for future work.

Biologically-Based Cognitive Architecture

Our PBWM working memory model is motivated by a number of considerations derived from an overarching biologically-based cognitive architecture for understanding human cognition (Atallah, Frank, & O'Reilly, 2004; O'Reilly & Munakata, 2000). This *tripartite* architecture is comprised of three functionally complementary brain systems that can be understood in terms of a set of computational tradeoffs, which provide a more precise and often subtle set of functional properties for these areas (Figure 1). These systems are as follows:

1. The posterior cortex (PC) system that performs the vast majority of the “automatic” sensory and motor processing in the brain. This system exhibits slow, integrative learning that extracts the long-term statistical structure of the environment, thereby efficiently representing accumulated knowledge and skills. In a sense, the PC system provides the “substrate” upon which the other two (“higher level”) systems operate to produce working memory phenomena.
2. The hippocampal system (HC) system that is specialized for rapid (e.g., one trial) learning that binds together arbitrary information, which can be subsequently recalled in the service of controlled processing. The neural specializations required for this rapid learning without interference are incompatible with the integrative statistical learning of the posterior cortical system, thus motivating the need for two separate neural systems.
3. The PFC/BG system that is specialized for active maintenance of internal contextual information (PFC), which can be dynamically updated by the basal ganglia (BG). This system can bias (control) ongoing processing throughout the cortex according to actively maintained information in the PFC (e.g., goals, instructions, partial products). It includes a dopamine (DA)-based learning system that uses a version of reinforcement learning mechanisms that are widely thought to be supported by the basal ganglia and related mid-brain structures.

Because it represents the canonical form of cortical processing, our model of the posterior cortex lies at the foundation of much of our work. It is based upon a computational framework called *Leabra* (O'Reilly, 1998; O'Reilly & Munakata, 2000; O'Reilly, 2001) that contains all of the basic mechanisms and properties

required by our models. These mechanisms have mostly been developed separately by other researchers over many years. Our specialized model of the hippocampal system has been developed over numerous publications (O'Reilly & McClelland, 1994; McClelland, McNaughton, & O'Reilly, 1995; O'Reilly & Rudy, 2001; Norman & O'Reilly, 2003; O'Reilly & Norman, 2002), and also includes many themes going back to earlier work (e.g., Marr, 1971; McNaughton & Morris, 1987; Rolls, 1989) that are widely adopted in the literature. Our hippocampal model (implemented in the Leabra framework) has been used to simulate a wide range of learning and memory phenomena, including human recognition memory (Norman & O'Reilly, 2003), and animal learning paradigms (e.g., contextual fear conditioning, nonlinear discrimination learning, transitive inference) (O'Reilly & Rudy, 2001; Frank, Rudy, & O'Reilly, 2003). Our models of the PFC/BG system are elaborated in the remainder of this paper.

In the context of this overall architecture, we can define working memory as an emergent property of the interactions between these three specialized brain areas, involving both active maintenance of task-relevant information (PFC/BG), and rapid learning of arbitrary associations (HC). These mechanisms support basic memory functions associated with working memory (i.e., memory for partial products of ongoing processing, task goals, etc), but also more complex *controlled processing* functions that are typically ascribed to a “central executive” in other working memory frameworks (e.g., Baddeley, 1986). Controlled processing emerges from the biasing influence of actively maintained and updated PFC representations on other parts of the system (e.g., Cohen, Dunbar, & McClelland, 1990; Cohen, Braver, & O'Reilly, 1996; O'Reilly, Braver, & Cohen, 1999; Miller & Cohen, 2001). This “top-down” biasing (which can also be supported by the HC to some extent) supports the performance of task-relevant processing in the face of competition from habitual or more well-practiced forms of processing that may not be task relevant.

This contrast between theoretical models that explicitly distinguish between working memory and executive function, typically designating them as separate “modules” (e.g., Baddeley, 1986), and our more emergent, interactive approach, is critical. Our view that working memory and executive function are really two sides of the same coin provides a more parsimonious model, that we believe is more consistent with extant biological data on the nature of neural specializations in the PFC/BG system. These contrasting approaches may be due in part to traditional models relying on a computer-like mental architecture, where processing is centralized and long-term memory is essentially a passive store. For those models, it

makes sense to have a separate set of cache-like working memory buffers dedicated to the temporary storage of items that are needed during processing by a completely separate central executive module (Baddeley, 1986). However, these architectures do not correspond to the known micro-architecture of the brain, in which processing and memory functions are typically distributed within and performed by the very same neural substrates (Rumelhart, McClelland, & PDP Research Group, 1986).

Thus, instead of thinking about the moving of information from long-term memory into and out of working memory buffers, we think that information is distributed in a relatively stable configuration throughout the cortex, and that working memory amounts to the controlled activation of these representations. This view shares some similarities with the view of working memory offered by production-system accounts (e.g., ACT, Anderson, 1983; Lovett, Reder, & Lebiere, 1999). However, it does not include the structural distinction between declarative and procedural knowledge assumed by such accounts.

To summarize then, our contention is that activation-based working memory is best thought of as the primary mechanism behind the PFC/BG system, and that controlled processing is a resulting emergent function that depends critically on this mechanism. Furthermore, we believe it is likely that working memory may represent a kind of phylogenetic extension of the same kinds of mechanisms that underly all forms of complex motor coordination and planning.

This conceptualization has important implications for forming theories and building models. While it certainly remains perfectly valid (as well as heuristically productive) to formulate theories to explain a narrow subset of working memory phenomena, in doing so one runs the distinct risk that those theories will be inherently incapable of dealing with other, equally important phenomena in the domain. Thus, any given theoretic model might best be judged not only by its fit to the data it was formulated for, but also by its ability to account for other working memory phenomena — or at least its potential to do so. In other words, we believe that a truly robust theory of working memory is best formulated with the big picture in mind.

The PBWM Model of Working Memory

Based on our cumulative work on a wide variety of working memory tasks, we have identified a core set of six functional demands, enumerated below, that are required by tasks involving working memory and executive function. Taken together, these functional demands provide a basic set of constraints for

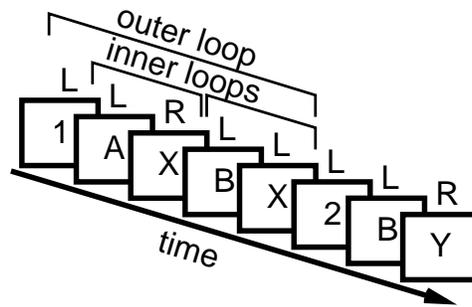


Figure 2: The 1-2-AX task. Stimuli are presented one at a time in a sequence. The participant responds by pressing the right key (R) to the target sequence, otherwise a left key (L) is pressed. If the subject last saw a 1, then the target sequence is an A followed by an X. If a 2 was last seen, then the target is a B followed by a Y. Distractor stimuli (e.g., 3, C, Z) may be presented at any point and are to be ignored. The maintenance of the task stimuli (1 or 2) constitutes a temporal outer-loop around multiple inner-loop memory updates required to detect the target sequence.

our biologically-based PBWM model. The 1-2-AX task, which is an extension of the widely-studied AX version of the continuous performance task (AX-CPT), provides a nice demonstration for these information-processing demands on the working memory system.

The AX-CPT is a standard working memory task that has been extensively studied in humans (Cohen, Perlstein, Braver, Nystrom, Noll, Jonides, & Smith, 1997; Braver & Cohen, 2000; Braver, Barch, & Cohen, submitted; Braver, Barch, & Cohen, 1999; Frank & O'Reilly, submitted). The subject is presented with sequential letter stimuli (A, X, B, Y), and is asked to detect the specific sequence of an A followed on the very next event by an X, by pushing the target (right) button. All other combinations (A-Y, B-X, B-Y) should be responded to with a non-target (left) button push. This task requires a relatively simple form of working memory, where the prior stimulus must be maintained over a delay until the next stimulus appears, so that the subject can discriminate the target from non-target sequences. This is the kind of activation-based working memory that has often been observed for example in electrophysiological studies of working memory in monkeys (e.g., Fuster & Alexander, 1971; Kubota & Niki, 1971; Miyashita & Chang, 1988; Funahashi, Bruce, & Goldman-Rakic, 1989; Miller, Erickson, & Desimone, 1996).

In the 1-2 extension of the AX-CPT task (1-2-AX; Figure 2; Frank, Loughry, & O'Reilly, 2001; Kroger, Nystrom, O'Reilly, Noelle, Braver, & Cohen, in preparation), the target sequence varies depending on prior *task demand* stimuli (a 1 or 2). Specifically, if the subject last saw a 1, then the target sequence is A-X. However, if the subject last saw a 2, then the target sequence is B-Y. Thus, the task demand stimuli define an *outer loop* of active maintenance (maintenance of task demands) within which there can be any number

of inner loops of active maintenance for the A-X level sequences.

Six Key Functional Demands Underlying Working Memory

With the 1-2-AX task as a concrete example, the six functional demands upon the working memory system are:

1. *Rapid updating*: The working memory system should be able to rapidly encode and maintain new information as it occurs. In the 1-2-AX task, as each relevant stimulus is presented, it must be rapidly encoded in working memory.
2. *Robust maintenance*: Information that remains relevant should be maintained in the face of the interference from ongoing processing or other stimulus inputs. In the 1-2-AX task, the task demand stimuli (1 or 2) in the outer loop must be maintained in the face of the interference from ongoing processing of inner loop stimuli and irrelevant distractors. Also, a specific A or B must also be maintained for the duration of each inner loop.
3. *Multiple, separate working memory representations*: To maintain the outer loop stimuli (1 or 2) while updating the inner loop stimuli (A or B), these two sets of representations must be distinct within the PFC (i.e., they must not be in direct mutual competition with one another, such that only one such representation could be active at a time).
4. *Selective updating*: Only some elements of working memory should be updated at any given time, while others are maintained. For example, in the inner loop, A or B should be updated while the task demand stimulus (1 or 2) is maintained.
5. *Top-down biasing of processing*: For working memory representations to achieve controlled processing, they must be able to bias (control) processing elsewhere in the brain. For example, whichever outer loop stimulus (1 or 2) is active at a given time must bias processing in the PFC/BG system itself, to condition responses and working memory updates as a function of the current target sequence.
6. *Learning what and when to gate*: Underlying all successful working memory task performance is the need to learn when to gate appropriately. This is a challenging problem because the benefits of having

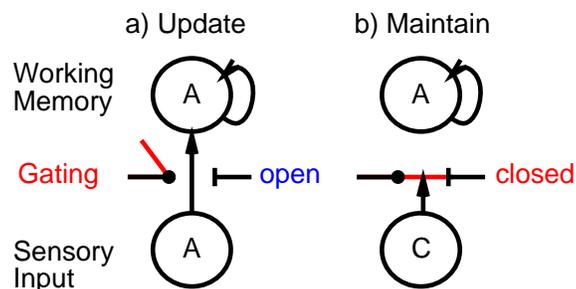


Figure 3: Illustration of active gating. When the gate is open, sensory input can rapidly update working memory (e.g., encoding the cue item A in the 1-2-AX task), but when it is closed, it cannot, thereby preventing other distracting information (e.g., distractor C) from interfering with the maintenance of previously stored information.

gated something in are only available later in time (e.g., encoding the 1 task demand stimulus only affects overt behavior and error-feedback later when confronted with an A-X sequence).

Earlier computational work by our group of collaborators has instantiated and validated a number of aspects of our overall theory, including the graded nature of controlled processing (Cohen et al., 1990); the ability of PFC representations to bias subsequent processing (Cohen & Servan-Schreiber, 1992); the role of PFC in active maintenance (Braver, Cohen, Jonides, Smith, Awh, Schumacher, Lauber, & Noll, 1995); the ability of the BG to update PFC working memory representations (Frank et al., 2001); and the role of the HC in rapid learning (O'Reilly & Rudy, 2001; Norman & O'Reilly, 2003). Most recently, we have been focused on the role of the PFC/BG system, and most specifically, how, mechanistically, it can learn to do what it has to do to support working memory.

Dynamic Updating via Basal Ganglia Gating

It is important to note that the first two demands (rapid updating versus robust maintenance) are in direct conflict with each other when viewed in terms of standard neural processing mechanisms. This motivates the obvious need for a dynamic gating mechanism to switch between these two modes of operation (Cohen et al., 1996; Braver & Cohen, 2000; O'Reilly et al., 1999; O'Reilly & Munakata, 2000). When the gate is open, working memory can get updated by incoming stimulus information; when it is closed, currently active working memory representations are robustly maintained even in the face of potential interference as from intervening distractor stimuli (Figure 3).

A central feature of our PBWM model is that the basal ganglia provide this requisite dynamic gating

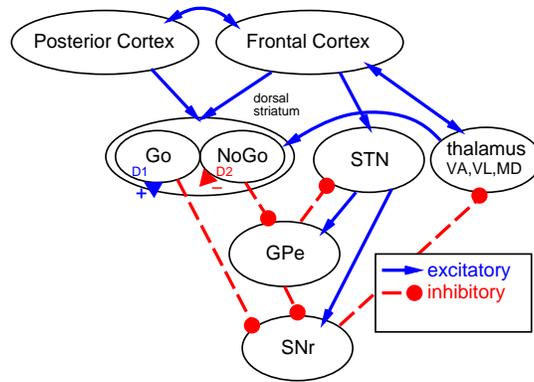


Figure 4: The basal ganglia are interconnected with frontal cortex through a series of parallel loops, each of the form shown. Working backward from the thalamus, which is bidirectionally excitatory with frontal cortex, the SNr (substantia nigra pars reticulata) is tonically active and inhibiting this excitatory circuit. When direct pathway “Go” neurons in dorsal striatum fire, they inhibit the SNr, and thus disinhibit frontal cortex, producing a gating-like modulation that we argue triggers the update of working memory representations in prefrontal cortex. The indirect pathway “NoGo” neurons of dorsal striatum counteract this effect by inhibiting the inhibitory GPe (globus pallidus, external segment). The STN (subthalamic nucleus) provides an additional dynamic background of inhibition (NoGo) by exciting the SNr.

mechanism for information maintained via sustained activation in the PFC. The basal ganglia are interconnected with frontal cortex through a series of parallel loops (Figure 4). When direct pathway “Go” neurons in dorsal striatum fire, they inhibit the SNr, and thus disinhibit frontal cortex, producing a gating-like modulation that we argue triggers the update of working memory representations in prefrontal cortex. The indirect pathway “NoGo” neurons of dorsal striatum counteract this effect by inhibiting the inhibitory GPe (globus pallidus, external segment). The STN (subthalamic nucleus) provides an additional dynamic background of inhibition (NoGo) by exciting the SNr. As reviewed in Frank et al., 2001, this idea is consistent with a wide range of empirical data and other computational models that have been developed largely in the domain of motor control, but also for working memory as well (e.g., Wickens, 1993; Houk & Wise, 1995; Wickens, Kotter, & Alexander, 1995; Dominey, Arbib, & Joseph, 1995; Berns & Sejnowski, 1995; Jackson & Houghton, 1995; Beiser & Houk, 1998; Berns & Sejnowski, 1998; Nakahara & Doya, 1998; Kropotov & Etlinger, 1999; Amos, 2000; Nakahara, Doya, & Hikosaka, 2001). Our ideas regarding just how the PFC and BG might work together to accomplish this complex coordination are outlined below, along with a brief description of the specific biologically-realistic computational mechanisms that our PBWM model uses to instantiate them.

1. *Rapid updating* occurs when direct pathway spiny neurons in the dorsal striatum fire (*Go* units). *Go* firing directly inhibits the substantia nigra pars reticulata (SNr), and releases its tonic inhibition of

the thalamus. This thalamic disinhibition enables, but does not directly cause (i.e., gates), a loop of excitation into the corresponding PFC “stripe” (see *Multiple, separate working memory representations* below). The effect of this net excitation is to toggle the state of bistable currents in the PFC neurons. Striatal Go neurons in the direct pathway are in competition (downstream in the SNr, if not actually in the striatum; Mink, 1996; Wickens, 1993) with a corresponding *NoGo* (indirect) pathway that promotes *greater* inhibition of thalamic neurons, thereby working to block gating.

2. *Robust maintenance* occurs via two intrinsic PFC mechanisms: 1) recurrent excitatory connectivity (O’Reilly et al., 1999; O’Reilly & Munakata, 2000), and; 2) bistability (Fellous, Wang, & Lisman, 1998; Wang, 1999; Durstewitz, Kelc, & Gunturkun, 1999; Durstewitz, Seamans, & Sejnowski, 2000), the latter of which is toggled between a maintenance state and a non-maintenance state by the Go gating signal from the BG.
3. *Multiple, separate working representations* are possible because of the “striped” micro-anatomy of the PFC, which is characterized by small, relatively isolated groups of interconnected neurons, thereby preventing undo interference between representations in different (even nearby) stripes (Levitt, Lewis, Yoshioka, & Lund, 1993; Pucak, Levitt, Lund, & Lewis, 1996). We estimate there may be as many as 20,000 such stripes in human prefrontal cortex (Frank et al., 2001).
4. *Selective updating* occurs because of the existence of independently updatable parallel loops of connectivity through different areas of the basal ganglia and frontal cortex (Alexander, DeLong, & Strick, 1986; Graybiel & Kimura, 1995; Middleton & Strick, 2000). We hypothesize that these loops are selective to the relatively fine-grained level of the anatomical stripes in PFC. This stripe-based gating architecture has an important advantage over the *global* nature of a purely dopamine-based gating signal (e.g., Braver & Cohen, 2000; Rougier & O’Reilly, 2002; Tanaka, 2002), which appears computationally inadequate for supporting a selective updating function by itself.
5. *Top-down biasing of processing* occurs via projections from actively-maintained representations in PFC to relevant areas throughout the brain, most typically the posterior cortex, but also the HC and the PFC/BG itself (Cohen & Servan-Schreiber, 1992; Fuster, 1989).
6. *Learning what and when to gate* is accomplished by a dopamine-based reinforcement-learning mech-

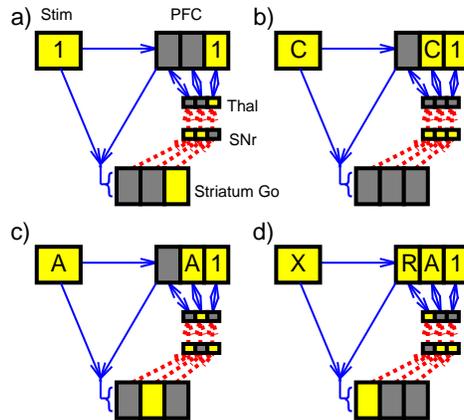


Figure 5: Illustration of how the basal ganglia gating of different PFC stripes can solve the 1-2-AX task (light color = active; dark = not-active). **a)** The 1 task is gated into an anterior PFC stripe because a corresponding striatal stripe fired Go. **b)** The distractor C fails to fire striatal Go neurons, so it will not be maintained; however, it does elicit transient PFC activity. Note that the 1 persists because of gating-induced robust maintenance. **c)** The A is gated in. **d)** A right keypress motor action is activated (using same BG-mediated disinhibition mechanism) based on X input plus maintained PFC context.

anism that is capable of providing temporally-appropriate learning signals to train gating update activity in the striatal Go and NoGo synapses. Thus, each spiny neuron develops its own unique pattern of connection weights enabling separate Go vs. NoGo decisions in each stripe.

Figure 5 shows how the BG-mediated selective gating mechanism can enable performance of the 1-2-AX task (see Frank et al., 2001 for a working simulation). When a task demand stimulus is presented (e.g., 1), a BG gating signal (i.e., a Go signal) must be activated to enable a particular PFC stripe to gate in and retain this information (panel a), and no stripe (or NoGo firing) should be activated for a distractor such as C (Panel b). A *different* stripe must be gated for the subsequent cue stimulus A (panel c). When the X stimulus is presented, the combination of this stimulus representation plus the maintained PFC working memory representations is sufficient to trigger a target response R (panel d).

Learning When to Gate in the BG

Of all the aspects of our model that purport to deconstruct the homunculus, learning when to gate is clearly the most central. For any working memory model, either the knowledge of when to update working memory must be built in by the model's designer, or, somehow, a model must learn it on its own, relying only on its builtin constraints interacting with its training experience. That is, without such a learning mechanism,

our model would have to resort to some kind of intelligent homunculus to control gating.

Our approach for simulating how the BG learns to update task-relevant versus irrelevant working memory information builds on our prior work for how the same system modulates the selection of motor responses. Specifically, the BG are thought to facilitate the selection of the most appropriate response, while suppressing all competing responses (Mink, 1996). In our models, the BG learn the distinction between good and bad responses via changes in dopamine firing during positive and negative reinforcement (for details, see Frank, 2005). In brief, our model leverages the observation that dopamine D1 and D2 receptors are relatively segregated in Go and NoGo cells, respectively (Gerfen, 1992; Wise, Murray, & Gerfen, 1996; Aubert, Ghorayeb, Normand, & Bloch, 2000). This is of interest because dopamine effects on neuronal excitability and synaptic plasticity are dependent on whether it acts via D1 or D2 receptors (Hernandez-Lopez, Bargas, Surmeier, Reyes, & Galarraga, 1997; Hernandez-Lopez, Tkatch, Perez-Garci, Galarraga, Bargas, Hamm, & Surmeier, 2000; Nishi, Snyder, & Greengard, 1997; Centonze, Picconi, Gubellini, Bernardi, & Calabresi, 2001). The net effect is that increases in dopamine during positive reinforcement enhance BG Go firing and learning via simulated D1 receptors, whereas decreases in dopamine during negative reinforcement have the opposite effect, enhancing NoGo firing and learning via simulated D2 receptors. This functionality enables the BG system to learn to discriminate between subtly different reinforcement values of alternative responses, and is consistent with several lines of biological and behavioral evidence (Frank, 2005).

For the PBWM model, we have extended these ideas to include BG modulation of PFC working memory representations. Thus, increases in dopamine reinforce BG Go firing to update information that is adaptive to store in working memory, while decreases in dopamine allow the model to learn that its current working memory state is maladaptive. In this manner, the BG eventually come to update information that is task-relevant, because maintenance of this information over time leads to adaptive behavior and reinforced responses. Conversely, the system learns to ignore distracting information, because its maintenance will interfere with that of task-relevant information and therefore lead to poor performance.

As Figure 5 illustrates, the learning problem in the basal ganglia boils down to learning when to fire a Go vs. NoGo signal in a given stripe based on the combination of current sensory (PC) input and maintained PFC activations. From a computational perspective, there are two fundamental problems that must be solved by the learning mechanism:

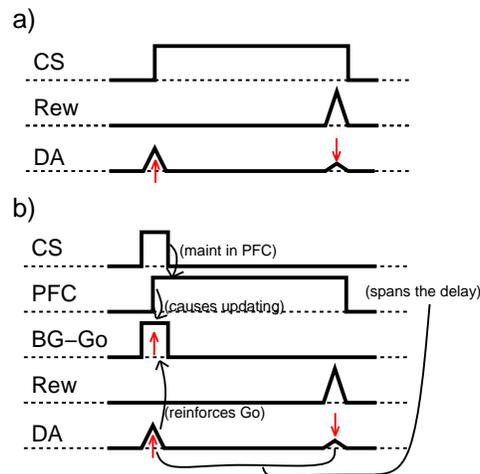


Figure 6: **a)** Schematic of dopamine (DA) neural firing for an input stimulus (Input, e.g., a tone) that reliably predicts a subsequent reward (Rew). Initially, DA fires at the point of reward, but then over repeated trials learns to fire at the onset of the stimulus. **b)** This DA firing pattern can solve the temporal credit assignment problem for PFC active maintenance. Here, the PFC maintains the transient input stimulus (initially by chance), leading to reward. As the DA system learns, it can predict subsequent reward at stimulus onset, by virtue of PFC “bridging the gap” (in place of a sustained input). DA firing at stimulus onset reinforces the firing of basal ganglia Go neurons, which drive updating in PFC.

Temporal credit assignment: The benefits of having encoded a given piece of information into prefrontal working memory are typically only available later in time (e.g., encoding the 1 task demand stimulus can only really help later (in terms of getting an actual reward) when confronted with an A-X sequence). Thus, the problem is to know which prior events were critical for subsequent good (or bad) performance.

Structural credit assignment: The network must decide which stripes should encode which different pieces of information at a given time, and when successful performance occurs, it must reinforce those stripes that actually contributed to this success. This form of credit assignment is what neural network models are typically very good at doing (e.g., the backpropagation algorithm), but clearly this form of structural credit assignment interacts with the temporal credit assignment problem and with the unique micro-anatomical structure of the stripe-loop architecture of the PFC and BG, making the technical problem considerably more complex.

The firing patterns of midbrain dopamine (DA) neurons (ventral tegmental area, VTA; substantia nigra pars compacta, SNc; both strongly innervated by the basal ganglia) exhibit the properties necessary to solve the temporal credit assignment problem, because they learn to fire for stimuli that predict subsequent

rewards (e.g., Schultz, Apicella, & Ljungberg, 1993; Schultz, 1998). This property is illustrated in schematic form in Figure 6a for a simple Pavlovian conditioning paradigm, where a stimulus (e.g., a tone) predicts a subsequent reward. Figure 6b shows how this predictive DA firing can reinforce BG Go firing to gate in and subsequently maintain a stimulus, when such maintenance leads to subsequent reward. Specifically, the DA firing can move discretely from the time of a reward to the onset of a stimulus that, if maintained in the PFC, leads to the subsequent delivery of this reward. Because this DA firing occurs at the time when the stimulus comes on, it is well timed to facilitate the storage of this stimulus in PFC. In our model, this occurs by reinforcing the connections between the stimulus and the Go gating neurons in the striatum, which then cause updating of PFC to maintain the stimulus.

The apparently predictive nature of the DA firing has most often been explained in terms of the temporal differences (TD) reinforcement learning mechanism (Sutton, 1988; Sutton & Barto, 1998; Schultz, Romo, Ljungberg, Mirenowicz, Hollerman, & Dickinson, 1995; Houk, Adams, & Barto, 1995; Montague, Dayan, & Sejnowski, 1996; Suri, Bargas, & Arbib, 2001; Contreras-Vidal & Schultz, 1999; Joel, Niv, & Ruppin, 2002). However, extensive exploration and analysis of these models has led us to develop a non-TD based account, which moves away from the prediction framework upon which it is based (O'Reilly & Frank, submitted-a, submitted-b). In brief, TD depends on being able to predict exactly when in time subsequent rewards will be received, based on a consistent sequence of stimulus states. However, in many of the tasks faced by our models (e.g., the 1-2-AX task), the precise sequence of stimulus states is highly unpredictable, and this dramatically undermines the efficacy of the TD learning mechanism. Our alternative learning mechanism, called PVLV (perceived value and learned value) involves two separable but inter-dependent learning mechanisms, each of which is essentially a simple delta-rule or Rescorla-Wagner mechanism (Rescorla & Wagner, 1972; Widrow & Hoff, 1960). This PVLV mechanism shares several features in common with the model of Brown, Bullock, and Grossberg (1999).

The further details of these PBWM and PVLV learning mechanisms are beyond the scope of this paper, but the basic results are that the resulting model can learn complex working memory tasks such as the 1-2-AX task based purely on trial-and-error experience with the task. We provide more examples of the learning and performance of this model below.

Empirical Tests of the Model

In addition to incorporating a large amount of existing empirical data, the overall PBWM model of the role of the PFC and BG in working memory makes a number of further predictions, several of which have been tested empirically. We review some of these below, and then discuss ongoing applications of the PBWM model to understand working memory behavior.

PFC Organization

The PBWM model strongly predicts that information that must be updated at different points in time during a working memory task (e.g., the outer loop vs. the inner loop of the 1-2-AX task) should be represented in different parts of the PFC. This prediction has been tested directly in the 1-2-AX task (Kroger et al., in preparation), and in a number of other tasks that also share this same kind of inner/outer loop structure (e.g., Braver & Bongiolatti, 2002; Koechlin, Corrado, & Grafman, 2000; Koechlin, Ody, & Kouneiher, 2003). In all of these cases, more anterior areas of PFC, specifically the frontal pole (Broadman's area 10) was selectively activated for outer-loop information, while more posterior areas (dorso-lateral PFC, Broadman's areas 46/9) were active for inner-loop information.

Basal Ganglia / Dopamine Mechanisms

Recently, we have tested various aspects of the hypothesized roles of the basal ganglia / dopamine system across both reinforcement learning and working memory processes. First, we demonstrated striking support for a central prediction of our model regarding dopamine involvement in "Go" and "NoGo" cognitive reinforcement learning (Frank, Seeberger, & O'Reilly, 2004; Frank, 2005). We tested Parkinson patients on and off medication, along with healthy senior control participants matched for age, education and a measure of verbal IQ. We predicted that decreased levels of dopamine in Parkinson's disease would lead to spared NoGo learning, but impaired Go learning (which depends on DA bursts). We further predicted that dopaminergic medication should alleviate the Go learning deficit, but would block the effects of dopamine dips needed to support NoGo learning, as was simulated to account for other medication-induced cognitive deficits in Parkinson's disease (Frank, 2005). Results were consistent with these predictions. In a probabilistic learning task, all patients and aged-matched controls learned to make choices that were more likely

to result in positive rather than negative reinforcement. The difference was in their strategy: patients taking their regular dose of dopaminergic medication implicitly learned more about the positive outcomes of their decisions (i.e., they were better at Go learning), whereas those who had abstained from taking medication implicitly learned to avoid negative outcomes (better NoGo learning). Age-matched controls did not differ in their tendency to learn more from the positive/negative outcomes of their decisions.

We have also tested predictions for a more a general role for BG/dopamine in cognitive function by administering low doses of dopamine agonists/antagonists to young, healthy participants (Frank & O'Reilly, submitted). The drugs used (cabergoline and haloperidol) were selective for D2 receptors, which are by far most prevalent in the BG. By acting on presynaptic D2 receptors, cabergoline reduces, while haloperidol enhances, the amount of phasic dopamine that is released during dopaminergic cell bursting (e.g., Wu, Reith, Walker, Kuhn, Carroll, & Garris, 2002). Again, results were consistent with our model: Increases in dopamine during learning caused participants to learn more about the positive outcomes of their decisions (as in medicated Parkinson's patients), whereas decreases in dopamine caused the same participants to learn more about negative outcomes (as in non-medicated patients). Notably, these same effects were borne out in the context of a modified version of the AX-CPT working memory task. In our version, a variable number of task-irrelevant distractor stimuli were presented during the delay period, and participants were told to ignore these distractors for the purpose of target (A-X) detection. We also included an attentional set-shifting condition, in which the previously task-relevant letters (A, X, B, Y) became distractors, while previous distractors were now task-relevant.

Interestingly, increases in dopamine by haloperidol enhanced selective working memory updating of task-relevant (i.e., "positively-valenced"), but not distracting ("negatively-valenced") information. By our model's account, dopamine release evoked during the presentation of task-relevant information reinforces BG Go firing to update this information. Consistent with this analysis, increased dopamine release also caused difficulty *not* updating (i.e., ignoring) this information when it subsequently became distracting in the set-shift. Conversely, under cabergoline (decreased dopamine release) set-shifting deficits were observed that were consistent with impaired, rather than enhanced, Go learning. In particular, whereas set-shifting deficits under haloperidol were only observed in trials for which participants had to ignore previously task-relevant distractors, deficits were observed under cabergoline when they only had to update the new task-

relevant set (i.e, in trials without distractors). Finally, and perhaps most suggestive for a role of BG dopamine in working memory, participants with low baseline working memory span were most subject to the effects of increases in dopamine by haloperidol, while those with high span were most subject to decreases in dopamine by cabergoline (Frank & O'Reilly, submitted). These latter results are consistent with the notion that individual differences in working memory span are partially characterized by underlying differences in dopamine levels (Kimberg, D'Esposito, & Farah, 1997), but extend this hypothesis in a more mechanistic fashion consistent with our modeling.

Taken together, these results provide strong support that BG signals, under modulation by dopamine, are critical for the updating of PFC working memory representations. Further, the model's success in capturing subtle cognitive effects in both Parkinson's disease and controlled dopamine manipulation suggests that it can also be applied to mechanistically understand cognitive deficits in those with more complex disorders involving BG/dopamine dysfunction, such as ADHD (attention deficit hyperactivity disorder) and schizophrenia. However, it is also the case that, although the PBWM model has been specifically designed to include many biological aspects, it obviously also goes beyond what is currently known. For example, the model ascribes specific roles to subsets of neurons in the nucleus accumbens, which provides testable hypotheses about the biology and function of these systems in the brain. It will be interesting to see how some of these ideas implemented in the PBWM model stand up to further biological investigations.

Simulating Multiple Working Memory Tasks in a Single Integrated Model

The PBWM model is quite complex. Although it is motivated by a wide range of empirical data, and some of its predictions have been tested and confirmed as described above, it is nevertheless important to constrain the complexity of the model further by subjecting it to increasingly stringent tests. One strategy that we have employed successfully in the past with both our hippocampal and posterior cortical models is to apply these models to as wide a range of cognitive neuroscience phenomena as possible. To the extent that the same basic model can account for a wide range of data, it provides confidence that the model is capturing some critical core elements of cognitive function. The virtues of this general approach have been forcefully argued by Newell (1990).

For these reasons, one important current research goal is to attempt to simulate a wide range of work-

Task	Brief Description
Stroop	Color words printed in same or different ink colors, response is either color or word. Differential slowing is observed for conflicting word/color in color naming condition, esp with PFC damage.
AX-CPT	Letters appear one at a time on the screen. Target sequence of A followed by X must be detected, requiring maintenance of A (impaired with PFC damage).
1-2-AX	Like AX-CPT, except target depends on prior 1 or 2 stimulus; 1 = AX, 2 = BY target sequence.
WCST	Wisconsin Card Sorting Task: multidimensional stimuli (color, shape, number) are sorted according to one dimension; Sorting “rule” learned by trial-and-error, and changes unexpectedly. Increased perseveration on initial rule observed with PFC damage.
ID/ED	Intradimensional/Extradimensional task, which is like WCST except multiple stimuli are presented simultaneously, allowing rule changes within a given dimension, or across different dimensions. PFC-related perseveration also observed.
Eriksen	A central, to-be-named stimulus is flanked by consistent or inconsistent stimuli. PFC damage causes more intrusion from flankers.
ABCA/ABBA	A sequence of stimuli are presented; a response is required when the 1st stimulus repeats. PFC damage causes false alarms in ABBA condition.
Serial recall	A sequence of verbal items must be recalled in order.
Sternberg	Digits are presented in an array, then removed. One location is probed & measured capacity is often greater than serial recall.
N-Back	Repetitions separated by N in a continuous stimulus sequence must be detected. PFC damage impairs performance.

Table 1: Brief descriptions of the tasks simulated in the multitask (MT) model, and the evidence implicating the PFC in them.

ing memory tasks with one instantiation of the PBWM model (Table 1). This research builds upon earlier work simulating many of the paradigmatic tasks thought to be characteristic of working memory and executive function, including: the Stroop effect (Cohen et al., 1990; O’Reilly & Munakata, 2000); the AX-CPT (Braver et al., 1995); the 1-2-AX (O’Reilly & Frank, submitted-a); the WCST (Rougier & O’Reilly, 2002); the ID/ED dynamic categorization task (O’Reilly, Noelle, Braver, & Cohen, 2002); and the Eriksen flanker task (Eriksen & Eriksen, 1974; Cohen, Servan-Schreiber, & McClelland, 1992; Yeung, Botvinick, & Cohen, 2004; Bogacz & Cohen, 2004). We also plan to simulate the ABCA/ABBA task (Miller et al., 1996), serial recall (phonological loop) (Burgess & Hitch, 1999), Sternberg task (Sternberg, 1966), and the N-Back task (Braver, Cohen, Nystrom, Jonides, Smith, & Noll, 1997).

In addition to the basic goal of simulating all of these tasks within a single model, this model will allow us to explore fundamental questions about the nature and origins of cognitive control, and intelligence more generally, by varying the training experiences that we give to the model prior to testing on the full set of standard experimental tasks. Specifically, we plan to build up a repertoire of complex cognitive skills by

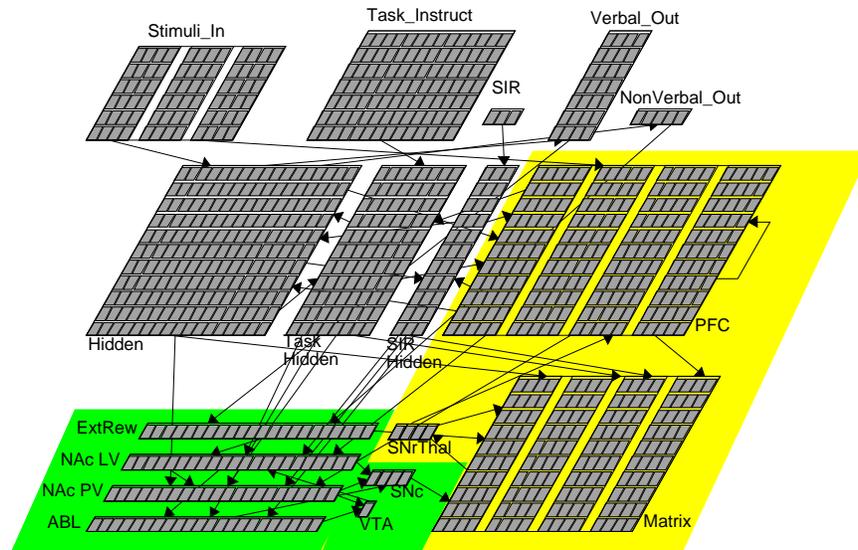


Figure 7: Overall structure of the MT model. Similar to our other PBWM models, the input/output layers are at the top-left of the diagram. The PFC and basal ganglia layers are at the bottom-right. The PVLV learning algorithm layers are highlighted in green while the full PBWM component is in yellow + green. Depending on the particular task, single or multiple featured stimuli are presented in one or more 'slots' in the Stimuli_In layer, along with task instructions in the Task_Instruct and SIR (store, ignore, recall) layers. Based on these inputs, plus context provided by PFC input, the Hidden layer determines the correct output in verbal or nonverbal form, or both. The shown model has four 'stripes' reflected in the four subgroups of the PFC and Matrix (striatal matrixosomes) layers, and the four units of the SNc and SNrThal layers. SNc = substantia nigra, pars compacta; SNrThal = abstracted layer reflecting direct and indirect pathways via substantia nigra, pars reticulata and thalamus; VTA = ventral tegmental area; ExtRew = external reward; NAcLV = nucleus accumbens, learned value; NAcPV = nucleus accumbens, perceived value; ABL = basolateral nucleus of the amygdala.

initial training on much simpler tasks, and to explore the extent to which performance across a range of tasks can build efficiently upon a common set of shared task elements. Before discussing these issues, we first describe our comprehensive MT (multi-task) model (designed around the basic PBWM mechanisms) that we are developing for this project, and our initial efforts demonstrating basic competency for the paradigmatic working memory tasks listed in Table 1.

The Full MT Model

Figure 7 shows the MT (multitask) model, with input/output layers appearing at the top of the network, posterior cortical "Hidden" layers and PFC layer in the middle, and basal ganglia/midbrain areas for learning and gating of PFC at the bottom. The input/output representations were designed to accommodate the vagaries of each individual task in a way that achieves a high level of surface validity.

Size:	Small	Medium	Large	X-Large
Color:	Black	White	Red	Green
Object:	Circle	Square	Triangle	Diamond
	“Red”	“Green”	1	2
	A	B	X	Y

Figure 8: Perceptual input features, organized along 3 separate dimensions. Three separate locations of these features are provided as input to the network.

The perceptual input representations in the MT model (Figure 8) assume a high level of perceptual preprocessing, such that different stimulus items (“objects”) are represented with consistent and unique activity patterns. We encode three separate (orthogonal) stimulus dimensions: object identity, color, and size, and we also provide three spatial locations in which a given object may appear. The task instruction layer tell the network what to do with the input stimuli, including the overall task and any more specific pieces of information that might be required (e.g., whether to do word reading or color naming in the Stroop task). We are exploring different ways of presenting these task instructions, including presenting them only at the beginning of a block of trials on a given task, at the beginning of a trial but not during it, and constantly throughout each trial. Furthermore, we will explore the presentation of a given task in a more surface valid manner, where after training on a set of basic cognitive operations, a given task is then “described” to the model in terms of the combination of these operations, instead of having it simply memorize the meaning of a distinctive task label. We have also included a sub-category of instruction inputs in the form of the store/ignore/recall (SIR) layer, which can be used to provide explicit working memory update signals that are encoded in a variety of different ways in different tasks, and may also be present via implicit timing signals via the cerebellum (e.g., Mauk & Buonomano, 2004; Ivry, 1996). The outputs include both verbal and non-verbal responses, the latter including button presses and pointing to locations.

The task instruction and SIR input each have a separate hidden layer associated with them, which enables the network to develop a more systematic internal representation of the task demands. These TaskHidden and SIRHidden layers, along with the perceptual inputs, all feed into one large common hidden layer (representing posterior association cortex), which in turn projects to the two output layers. The PFC is bidirectionally connected to all relevant high-level processing layers (sensory input, task hidden, central hidden, and output), and its associated BG layers receive from all of these layers as well to provide control over the

learning and execution of the dynamic gating signals. Note that the shown PFC/BG system has four stripes, with each stripe representing a selectively updatable component of working memory. More stripes facilitate faster learning, but result in a larger, more computationally costly model, so the exact number of stripes is a matter of pragmatic optimization in the model (in the brain, we estimate that many thousands of stripes are present).

When sensory inputs are presented, activation flows throughout the network in a bidirectional manner, so that internal posterior cortical “hidden” layers are affected by both these bottom-input and maintained top-down activations in the PFC. In the Leabra algorithm that we use, individual units are modeled as point neurons, with simulated ion channels contributing to a membrane potential, which is in turn passed through a thresholded nonlinear activation function to obtain a continuous instantaneous spike rate output that is communicated to other units. The inhibitory conductances are efficiently computed according to a k-winners-take-all algorithm (kwta), which ensures that no more than some percentage (typically between 15-25%) of units within a layer are active at a time.

Outside of the BG system, learning occurs as a result of both Hebbian and error-driven mechanisms, with the error-driven learning computed in a biologically-plausible fashion based on the GeneRec learning algorithm (O'Reilly, 1996). The learning mechanisms for the BG components (PVLV algorithm) were described earlier.

The model can be systematically lesioned in a wide variety of ways and its performance tested on each of the behavioral paradigms. In some cases, there are clear patient data that the model will be expected to simulate. However, much of the time, these lesions will not clearly map onto any existing piece of empirical data, and will therefore stand as important testable predictions of the model. In addition to the standard techniques of damaging units and connections, we can also manipulate the dopaminergic and other pathways in the BG and their projections to the PFC to simulate conditions such as Parkinson's, ADHD, and schizophrenia.

Current Progress and Future Directions

The model is currently able to perform a set of core tasks, including the Stroop, AX-CPT, 1-2-AX, and WCST, in addition to a set of more primitive component tasks (e.g., naming, matching, and comparing stim-

ulus features, dimensions, and locations) that also involve basic working memory capacities. The network has replicated basic features of these tasks as highlighted in Table 1 (e.g., the differential slowing of the color-naming conflict condition in the Stroop task), but more extensive detailed testing has yet to be performed, awaiting further training of the model on more tasks, and all of the tasks integrated together. Thus, although we are optimistic that the network will succeed in simulating all of these tasks, considerable work remains to be done.

As mentioned earlier, a single instantiation of our comprehensive MT model will provide important opportunities for exploring a wide range of issues in cognitive development and human intelligence generally. Perhaps the greatest mystery in cognitive processing is where all the “smarts” come from to control the system in a task-appropriate manner. Specifically, a fundamental question we face is, How is it that people quickly adapt to performing novel cognitive tasks, when it takes monkeys months of highly-focused training to learn a single new task? We hypothesize that people develop an extensive repertoire of basic cognitive operations throughout the long developmental period into adulthood, and then, are able to rapidly and flexibly combine these basic “building block” operations to solve novel tasks. This may sound like common sense, but demonstrating even rudimentary capabilities in this regard in a mechanistically explicit computational model remains an important challenge.

To explore these issues, we are planning to train the MT model on a large corpus of primitive task-oriented experiences (e.g., object naming, color naming, etc.) so as to establish a basic set of capabilities underlying the more complex task paradigms targeted above. A critical issue will be how much specific training on each task will still be required, in comparison to prior training on more basic cognitive operations that may be combinatorially applied to each task. One extreme of this dimension is represented by studies conducted on monkeys, where many months of highly specific training are required to achieve task performance. In contrast, adult human participants typically require a simple set of verbal instructions, followed by a small number of practice trials, to achieve high levels of performance. In the past, neural network models have been much closer to the monkey end of this spectrum (or worse), requiring extensive task-specific training. Other cognitive models (e.g., production systems) typically rely on the intelligence of the modeler to program in all the necessary task components by hand, while non-computational theories often invoke constructs that amount to homunculi, for example the “central executive” of Baddeley (1986).

A major goal of our overall research program is to move our models closer to the human end of the spectrum. A key hypothesis to be tested is that our model will learn complex tasks significantly faster after being pre-trained on simpler, relevant ones. If we are successful in achieving this goal, it would represent a critical qualitative step forward in the modeling of human-like intelligence.

Some preliminary work using an earlier version of our basic model provides reason to be optimistic regarding this overall approach. In simulations of a task paradigm we called cross-task generalization (XT; Rougier, Noelle, Braver, Cohen, & O'Reilly, in press), we explored the ability of training on one set of tasks to generalize (transfer) to other, related tasks. In general, the key to generalization in a neural network is the formation of *abstract* (e.g., categorical) representations (O'Reilly & Munakata, 2000; Munakata & O'Reilly, 2003). For example, the abstract representation of the category “dog” (which must be abstracted over experiences with specific dogs) can enable one to surmise that Poodles eat dog food based on seeing a German Shepherd eating dog food, even though this might not be obvious given the raw perceptual differences among members of the dog category (e.g., you might otherwise be tempted to feed a Poodle cat food).

In the Rougier et al. (in press) model, we trained a PFC model using a simple form of dynamic gating mechanism on a varying number of related tasks operating on simple visual stimuli (e.g., *name* a “feature” of the stimulus along a given “dimension” such as its color, shape, or size; *match* two stimuli along one of these dimensions; *compare* the relative size of two stimuli). To test for generalization, we only trained a given task on a small percentage (e.g., 30%) of all the stimuli, and then tested that task on stimuli that were trained in other tasks. We found that only the model with an intact PFC and dynamic gating mechanism was capable of significant levels of generalization. Furthermore, this model developed discrete rule-like representations in the PFC that clearly and uniquely encoded the task-relevant stimulus dimensions, and the generalization performance and formation of these rule-like representations was strongly correlated throughout all the models.

We think this pattern of results reflect a general principle for why the prefrontal cortex should develop more abstract representations than posterior cortex, and thus facilitate flexible generalization to novel environments: abstraction derives from the maintenance of stable representations over time, interacting with learning mechanisms that extract commonalities over varying inputs. Supporting this view are data showing

that damage to prefrontal cortex impairs abstraction abilities (e.g., Dominey & Georgieff, 1997), and that prefrontal cortex in monkeys develops more abstract category representations than those in posterior cortex (Wallis, Anderson, & Miller, 2001; Freedman, Riesenhuber, Poggio, & Miller, 2002; Nieder, Freedman, & Miller, 2002).

As a preliminary test of our overall approach, we have successfully implemented the core findings of the cross task generalization model (Rougier et al., in press) in a version of the MT model described here to validate that the latter can replicate the core results generated by that model. Specifically, we have been able to easily demonstrate that the latest PBWM mechanisms, using the input framework and training environment of the Rougier et al. (in press) model, could quickly learn the core tasks (*name* feature, *match* feature, *compare* feature) and could also successfully learn to generalize across tasks just as in the original.

Future Developmental Directions for the Model

The question of how the PFC is functionally organized is also prominent in the literature, and remains largely unresolved. We think this path of research can shed considerable light on this issue as well. Previously, we have proposed that the anterior-posterior (and perhaps dorsal-ventral) axis of the PFC might be organized along a gradient from abstract to concrete, respectively (O'Reilly et al., 2002; O'Reilly & Munakata, 2000). However, the PBWM model currently has no mechanism for encouraging such a gradient (or any other kind of gradient) to develop in the stripes of the PFC (all stripes are equipotent in their access to information). Therefore, an additional avenue of research we plan to explore is to look at various ways of biasing the model to develop a gradient of organization along its PFC stripes, and see what types of gradients actually develop in response to the battery of training provided, while assessing any behavioral implications this organization might have (e.g., does the organizational principle actually facilitate processing, and if so, how?). One particular organizational bias suggested by the biology is to have only the more posterior areas of PFC connected (bidirectionally) with posterior cortical areas, while more anterior PFC areas connect only with these posterior PFC areas. Thus, anterior PFC areas might be able to serve as more abstract biasing inputs to more posterior PFC areas, which in turn bias more specific processing in posterior cortex.

Understanding the human capacity for generativity may be one of the greatest challenges facing the field of “higher-level” cognitive function. We think that the mechanisms of the PBWM model, and in particular its ability to exhibit limited variable-binding functionality, may be critical steps along the way to

such an understanding. It may be that, over the 20 or so years it takes to fully develop a functional PFC, people have developed a systematic and flexible set of representation that support dynamic reconfiguration of input/output mappings according to maintained PFC representations. Thus, these PFC “variables” can be activated by task instructions, and support novel task performance without extensive training. This and many other important problems remain to be addressed in future research.

Finally, while our model addresses the computational role of basal ganglia dopamine in working memory, we are only beginning to explore dopamine effects in PFC. In brief, we think that phasic dopamine effects are critical for learning in the BG, whereas longer lasting dopamine effects in PFC support robust maintenance of working memory representations (Durstewitz et al., 2000; Seamans & Yang, 2004; Tanaka, 2002). In the model described here, these dopaminergic effects in PFC were abstracted and subsumed by a simple intracellular maintenance current — but these currents are known to depend on a healthy level of tonic dopamine.

Conclusion

Although many theoretical models have been developed purporting to explain aspects of working memory and executive function, the mechanistic basis underlying them has remained inadequately described, often amounting to a homunculus. In this paper, we have reviewed some of the progress made by our colleagues and others in attempting to deconstruct this implicit homunculus by elucidating the precise computational and neural mechanisms underlying them, particularly the role of the PFC. These ideas can be specified at multiple levels. At a more abstract level, we outlined six key functional demands that we see underlying working memory, which need to be satisfied by the neural system. We also described a detailed implementation of these functional demands in the PBWM (prefrontal cortex basal ganglia working memory) computational model. This model attempts to incorporate detailed biological constraints in addition to the more abstract functional demands, and is capable of learning complex working memory tasks strictly as a function of experience (without task-specific knowledge having to be built in by the modeler).

We are currently applying this computational model to a range of different working memory tasks, which should strongly test the cognitive neuroscience validity of the model. For example, the model can be used to explore roles of the individual neural systems involved by perturbing parameters to simulate development,

aging, pharmacological manipulations, and neurological dysfunction, and it promises to be extensible to a broad array of other relevant manifestations of working memory and executive function. Furthermore, we hope to use this platform to explore fundamental questions regarding the neural basis and origins of our uniquely flexible human intelligence and capacity for cognitive control.

References

- Alexander, G. E., DeLong, M. R., & Strick, P. L. (1986). Parallel organization of functionally segregated circuits linking basal ganglia and cortex. *Annual Review of Neuroscience*, *9*, 357–381.
- Amos, A. (2000). A computational model of information processing in the frontal cortex and basal ganglia. *Journal of Cognitive Neuroscience*, *12*, 505–519.
- Anderson, J. R. (1983). *The architecture of cognition*. Cambridge, MA: Harvard University Press.
- Atallah, H. E., Frank, M. J., & O'Reilly, R. C. (2004). Hippocampus cortex and basal ganglia: insights from computational models of complementary learning systems. *Neurobiology of Learning and Memory*, *82*(3), 253–267.
- Aubert, I., Ghorayeb, I., Normand, E., & Bloch, B. (2000). Phenotypical characterization of the neurons expressing the D1 and D2 dopamine receptors in the monkey striatum. *Journal of Comparative Neurology*, *418*, 22–32.
- Baddeley, A. D. (1986). *Working memory*. New York: Oxford University Press.
- Beiser, D. G., & Houk, J. C. (1998). Model of cortical-basal ganglionic processing: encoding the serial order of sensory events. *Journal of Neurophysiology*, *79*, 3168–3188.
- Berns, G. S., & Sejnowski, T. J. (1995). How the basal ganglia make decisions. In A. Damasio, H. Damasio, & Y. Christen (Eds.), *Neurobiology of decision-making* (pp. 101–113). Berlin: Springer-Verlag.
- Berns, G. S., & Sejnowski, T. J. (1998). A computational model of how the basal ganglia produces sequences. *Journal of Cognitive Neuroscience*, *10*, 108–121.
- Bogacz, R., & Cohen, J. D. (2004). Parameterization of connectionist models. *Behavioral Research Methods, Instruments & Computers*, *36*(4), 732–741.

- Braver, T. S., Barch, D. M., & Cohen, J. D. (1999). Cognition and control in schizophrenia: A computational model of dopamine and prefrontal function. *Biological Psychiatry*, *46*, 312–328.
- Braver, T. S., Barch, D. M., & Cohen, J. D. (submitted). Mechanisms of cognitive control: Active memory, inhibition, and the prefrontal cortex.
- Braver, T. S., & Bongiolatti, S. R. (2002). The role of frontopolar cortex in subgoal processing during working memory. *Neuroimage*, *15*, 523–536.
- Braver, T. S., & Cohen, J. D. (2000). On the control of control: The role of dopamine in regulating prefrontal function and working memory. In S. Monsell, & J. Driver (Eds.), *Control of cognitive processes: Attention and performance XVIII* (pp. 713–737). Cambridge, MA: MIT Press.
- Braver, T. S., Cohen, J. D., Jonides, J., Smith, E. E., Awh, E., Schumacher, E., Lauber, E., & Noll, D. C. (1995). A parametric study of frontal cortex involvement in human working memory using functional MRI. *Society for Neuroscience Abstracts*. San Diego, CA: Society for Neuroscience.
- Braver, T. S., Cohen, J. D., Nystrom, L. E., Jonides, J., Smith, E. E., & Noll, D. C. (1997). A parametric study of frontal cortex involvement in human working memory. *NeuroImage*, *5*, 49–62.
- Brown, J., Bullock, D., & Grossberg, S. (1999). How the basal ganglia use parallel excitatory and inhibitory learning pathways to selectively respond to unexpected rewarding cues. *Journal of Neuroscience*, *19*, 10502–10511.
- Burgess, N., & Hitch, G. J. (1999). Memory for serial order: A network model of the phonological loop and its timing. *Psychological Review*, *106*, 551–581.
- Centonze, D., Picconi, B., Gubellini, P., Bernardi, G., & Calabresi, P. (2001). Dopaminergic control of synaptic plasticity in the dorsal striatum. *European Journal of Neuroscience*, *13*, 1071–1077.
- Cohen, J. D., Braver, T. S., & O'Reilly, R. C. (1996). A computational approach to prefrontal cortex, cognitive control, and schizophrenia: Recent developments and current challenges. *Philosophical Transactions of the Royal Society (London) B*, *351*, 1515–1527.
- Cohen, J. D., Dunbar, K., & McClelland, J. L. (1990). On the control of automatic processes: A parallel distributed processing model of the Stroop effect. *Psychological Review*, *97*(3), 332–361.

- Cohen, J. D., Perlstein, W. M., Braver, T. S., Nystrom, L. E., Noll, D. C., Jonides, J., & Smith, E. E. (1997). Temporal dynamics of brain activation during a working memory task. *Nature*, *386*, 604–608.
- Cohen, J. D., & Servan-Schreiber, D. (1992). Context, cortex, and dopamine: A connectionist approach to behavior and biology in schizophrenia. *Psychological Review*, *99*, 45–77.
- Cohen, J. D., Servan-Schreiber, D., & McClelland, J. L. (1992). A parallel distributed processing approach to automaticity. *American Journal of Psychology*, *105*, 239–269.
- Contreras-Vidal, J. L., & Schultz, W. (1999). A predictive reinforcement model of dopamine neurons for learning approach behavior. *Journal of Comparative Neuroscience*, *6*, 191–214.
- Dominey, P., Arbib, M., & Joseph, J.-P. (1995). A model of corticostriatal plasticity for learning oculomotor associations and sequences. *Journal of Cognitive Neuroscience*, *7*, 311–336.
- Dominey, P. F., & Georgieff, N. (1997). Schizophrenics learn surface but not abstract structure in a serial reaction time task. *Neuroreport*, *8*, 2877.
- Durstewitz, D., Kelc, M., & Gunturkun, O. (1999). A neurocomputational theory of the dopaminergic modulation of working memory functions. *Journal of Neuroscience*, *19*, 2807.
- Durstewitz, D., Seamans, J. K., & Sejnowski, T. J. (2000). Dopamine-mediated stabilization of delay-period activity in a network model of prefrontal cortex. *Journal of Neurophysiology*, *83*, 1733–1750.
- Eriksen, B. A., & Eriksen, C. W. (1974). Effects of noise letters upon the identification of a target letter in a nonsearch task. *Perception and Psychophysics*, *16*, 143–149.
- Fellous, J. M., Wang, X. J., & Lisman, J. E. (1998). A role for NMDA-receptor channels in working memory. *Nature Neuroscience*, *1*, 273–275.
- Frank, M. (2005). Dynamic dopamine modulation in the basal ganglia: A neurocomputational account of cognitive deficits in medicated and non-medicated Parkinsonism. *Journal of Cognitive Neuroscience*, *17*, 51–72.
- Frank, M., & O'Reilly, R. (submitted). A mechanistic account of striatal dopamine function in cognition: Psychopharmacological studies with cabergoline and haloperidol.
- Frank, M. J., Loughry, B., & O'Reilly, R. C. (2001). Interactions between the frontal cortex and basal

- ganglia in working memory: A computational model. *Cognitive, Affective, and Behavioral Neuroscience*, *1*, 137–160.
- Frank, M. J., Rudy, J. W., & O'Reilly, R. C. (2003). Transitivity, flexibility, conjunctive representations and the hippocampus: II. a computational analysis. *Hippocampus*, *13*, 341–354.
- Frank, M. J., Seeberger, L., & O'Reilly, R. C. (2004). By carrot or by stick: Cognitive reinforcement learning in Parkinsonism. *Science*, *306*, 1940–1943.
- Freedman, D. J., Riesenhuber, M., Poggio, T., & Miller, E. K. (2002). Visual categorization and the primate prefrontal cortex: Neurophysiology and behavior. *Journal of Neurophysiology*, *88*, 929–941.
- Funahashi, S., Bruce, C. J., & Goldman-Rakic, P. S. (1989). Mnemonic coding of visual space in the monkey's dorsolateral prefrontal cortex. *Journal of Neurophysiology*, *61*, 331–349.
- Fuster, J. M. (1989). *The prefrontal cortex: Anatomy, physiology and neuropsychology of the frontal lobe*. New York: Raven Press.
- Fuster, J. M., & Alexander, G. E. (1971). Neuron activity related to short-term memory. *Science*, *173*, 652–654.
- Gerfen, C. (1992). The neostriatal mosaic: multiple levels of compartmental organization in the basal ganglia. *Annual Review of Neuroscience*, *15*, 285–320.
- Graybiel, A. M., & Kimura, M. (1995). Adaptive neural networks in the basal ganglia. In J. C. Houk, J. L. Davis, & D. G. Beiser (Eds.), *Models of information processing in the basal ganglia* (pp. 103–116). Cambridge, MA: MIT Press.
- Hernandez-Lopez, S., Bargas, J., Surmeier, D. J., Reyes, A., & Galarraga, E. (1997). D1 receptor activation enhances evoked discharge in neostriatal medium spiny neurons by modulating an L-type Ca^{2+} conductance. *Journal of Neuroscience*, *17*, 3334–42.
- Hernandez-Lopez, S., Tkatch, T., Perez-Garci, E., Galarraga, E., Bargas, J., Hamm, H., & Surmeier, D. (2000). D2 dopamine receptors in striatal medium spiny neurons reduce L-type Ca^{2+} currents and excitability via a novel PLC β 1-IP $_3$ -calcineurin-signaling cascade. *Journal of Neuroscience*, *20*, 8987–95.
- Houk, J. C., Adams, J. L., & Barto, A. G. (1995). A model of how the basal ganglia generate and use

- neural signals that predict reinforcement. In J. C. Houk, J. L. Davis, & D. G. Beiser (Eds.), *Models of information processing in the basal ganglia* (pp. 233–248). Cambridge, MA: MIT Press.
- Houk, J. C., & Wise, S. P. (1995). Distributed modular architectures linking basal ganglia, cerebellum, and cerebral cortex: their role in planning and controlling action. *Cerebral Cortex*, *5*, 95–110.
- Ivry, R. (1996). The representation of temporal information in perception and motor control. *Current Opinion in Neurobiology*, *6*, 851–857.
- Jackson, S., & Houghton, G. (1995). Sensorimotor selection and the basal ganglia: A neural network model. In J. C. Houk, J. L. Davis, & D. G. Beiser (Eds.), *Models of information processing in the basal ganglia* (pp. 337–368). Cambridge, MA: MIT Press.
- Joel, D., Niv, Y., & Ruppin, E. (2002). Actor-critic models of the basal ganglia: new anatomical and computational perspectives. *Neural Networks*, *15*, 535–547.
- Kimberg, D. Y., D’Esposito, M., & Farah, M. J. (1997). Effects of bromocriptine on human subjects depend on working memory capacity. *Neuroreport*, *8*, 3581–3585.
- Koechlin, E., Corrado, G., & Grafman, J. (2000). Dissociating the role of the medial and lateral anterior prefrontal cortex in human planning. *Proceedings of the National Academy of Sciences*, *97*, 7651.
- Koechlin, E., Ody, C., & Kouneiher, F. (2003). Neuroscience: The architecture of cognitive control in the human prefrontal cortex. *Science*, *424*, 1181–1184.
- Kroger, J., Nystrom, L., O’Reilly, R. C., Noelle, D., Braver, T. S., & Cohen, J. D. (in preparation). Multiple levels of temporal abstraction in the prefrontal cortex: converging results from a computational model and fmri.
- Kropotov, J. D., & Etlinger, S. C. (1999). Selection of actions in the basal ganglia-thalamocortical circuits: Review and model. *International Journal of Psychophysiology*, *31*, 197–217.
- Kubota, K., & Niki, H. (1971). Prefrontal cortical unit activity and delayed alternation performance in monkeys. *Journal of Neurophysiology*, *34*, 337–347.
- Levitt, J. B., Lewis, D. A., Yoshioka, T., & Lund, J. S. (1993). Topography of pyramidal neuron intrinsic

- connections in macaque monkey prefrontal cortex (areas 9 & 46). *Journal of Comparative Neurology*, 338, 360–376.
- Lovett, M. C., Reder, L. M., & Lebiere, C. (1999). Modeling working memory in a unified architecture. In A. Miyake, & P. Shah (Eds.), *Models of working memory: Mechanisms of active maintenance and executive control*. New York: Cambridge University Press.
- Marr, D. (1971). Simple memory: A theory for archicortex. *Philosophical Transactions of the Royal Society (London) B*, 262, 23–81.
- Mauk, M. D., & Buonomano, D. V. (2004). The neural basis of temporal processing. *Annual Review of Neuroscience*, 27, 307–340.
- McClelland, J. L., McNaughton, B. L., & O'Reilly, R. C. (1995). Why there are complementary learning systems in the hippocampus and neocortex: Insights from the successes and failures of connectionist models of learning and memory. *Psychological Review*, 102, 419–457.
- McNaughton, B. L., & Morris, R. G. M. (1987). Hippocampal synaptic enhancement and information storage within a distributed memory system. *Trends in Neurosciences*, 10(10), 408–415.
- Middleton, F. A., & Strick, P. L. (2000). Basal ganglia output and cognition: Evidence from anatomical, behavioral, and clinical studies. *Brain and Cognition*, 42, 183–200.
- Miller, E. K., & Cohen, J. D. (2001). An integrative theory of prefrontal cortex function. *Annual Review of Neuroscience*, 24, 167–202.
- Miller, E. K., Erickson, C. A., & Desimone, R. (1996). Neural mechanisms of visual working memory in prefrontal cortex of the macaque. *Journal of Neuroscience*, 16, 5154.
- Mink, J. W. (1996). The basal ganglia: Focused selection and inhibition of competing motor programs. *Progress in Neurobiology*, 50, 381–425.
- Miyashita, Y., & Chang, H. S. (1988). Neuronal correlate of pictorial short-term memory in the primate temporal cortex. *Nature*, 331, 68–70.
- Montague, P. R., Dayan, P., & Sejnowski, T. J. (1996). A framework for mesencephalic dopamine systems based on predictive Hebbian learning. *Journal of Neuroscience*, 16, 1936–1947.

- Munakata, Y., & O'Reilly, R. C. (2003). Developmental and computational neuroscience approaches to cognition: The case of generalization. *Cognitive Studies*, *10*, 76–92.
- Nakahara, H., & Doya, K. (1998). Near saddle-node bifurcation behavior as dynamics in working memory for goal-directed behavior. *Neural Computation*, *10*, 113–132.
- Nakahara, H., Doya, K., & Hikosaka, O. (2001). Parallel cortico-basal ganglia mechanisms for acquisition and execution of visuomotor sequences—a computational approach. *Journal of Cognitive Neuroscience*, *13*, 626–647.
- Newell, A. (1990). *Unified theories of cognition*. Cambridge, MA: Harvard University Press.
- Nieder, A., Freedman, D. J., & Miller, E. K. (2002). Representation of the quantity of visual items in the primate prefrontal cortex. *Science*, *298*, 1708–1711.
- Nishi, A., Snyder, G., & Greengard, P. (1997). Bidirectional regulation of DARPP-32 phosphorylation by dopamine. *Journal of Neuroscience*, *17*, 8147–8155.
- Norman, K. A., & O'Reilly, R. C. (2003). Modeling hippocampal and neocortical contributions to recognition memory: A complementary learning systems approach. *Psychological Review*, *110*, 611–646.
- O'Reilly, R. C. (1996). Biologically plausible error-driven learning using local activation differences: The generalized recirculation algorithm. *Neural Computation*, *8*(5), 895–938.
- O'Reilly, R. C. (1998). Six principles for biologically-based computational models of cortical cognition. *Trends in Cognitive Sciences*, *2*(11), 455–462.
- O'Reilly, R. C. (2001). Generalization in interactive networks: The benefits of inhibitory competition and Hebbian learning. *Neural Computation*, *13*, 1199–1242.
- O'Reilly, R. C., Braver, T. S., & Cohen, J. D. (1999). A biologically based computational model of working memory. In A. Miyake, & P. Shah (Eds.), *Models of working memory: Mechanisms of active maintenance and executive control*. (pp. 375–411). New York: Cambridge University Press.
- O'Reilly, R. C., & Frank, M. J. (submitted-a). Making working memory work: A computational model of learning in the frontal cortex and basal ganglia.

- O'Reilly, R. C., & Frank, M. J. (submitted-b). Rewards are timeless: The perceived value and learned value (pvlv) pavlovian learning algorithm.
- O'Reilly, R. C., & McClelland, J. L. (1994). Hippocampal conjunctive encoding, storage, and recall: Avoiding a tradeoff. *Hippocampus*, *4*(6), 661–682.
- O'Reilly, R. C., & Munakata, Y. (2000). *Computational explorations in cognitive neuroscience: Understanding the mind by simulating the brain*. Cambridge, MA: MIT Press.
- O'Reilly, R. C., Noelle, D., Braver, T. S., & Cohen, J. D. (2002). Prefrontal cortex and dynamic categorization tasks: Representational organization and neuromodulatory control. *Cerebral Cortex*, *12*, 246–257.
- O'Reilly, R. C., & Norman, K. A. (2002). Hippocampal and neocortical contributions to memory: Advances in the complementary learning systems framework. *Trends in Cognitive Sciences*, *6*, 505–510.
- O'Reilly, R. C., & Rudy, J. W. (2001). Conjunctive representations in learning and memory: Principles of cortical and hippocampal function. *Psychological Review*, *108*, 311–345.
- Pucak, M. L., Levitt, J. B., Lund, J. S., & Lewis, D. A. (1996). Patterns of intrinsic and associational circuitry in monkey prefrontal cortex. *Journal of Comparative Neurology*, *376*, 614–630.
- Rescorla, R. A., & Wagner, A. R. (1972). A theory of Pavlovian conditioning: Variation in the effectiveness of reinforcement and non-reinforcement. In A. H. Black, & W. F. Prokasy (Eds.), *Classical conditioning II: Theory and research* (pp. 64–99). New York: Appleton-Century-Crofts.
- Rolls, E. T. (1989). Functions of neuronal networks in the hippocampus and neocortex in memory. In J. H. Byrne, & W. O. Berry (Eds.), *Neural models of plasticity: Experimental and theoretical approaches* (pp. 240–265). San Diego, CA: Academic Press.
- Rougier, N. P., Noelle, D., Braver, T. S., Cohen, J. D., & O'Reilly, R. C. (in press). Prefrontal cortex and the flexibility of cognitive control: Rules without symbols. *Proceedings of the National Academy of Sciences*.
- Rougier, N. P., & O'Reilly, R. C. (2002). Learning representations in a gated prefrontal cortex model of dynamic task switching. *Cognitive Science*, *26*, 503–520.

- Rumelhart, D. E., McClelland, J. L., & PDP Research Group (Eds.). (1986). *Parallel distributed processing. Volume 1: Foundations*. Cambridge, MA: MIT Press.
- Schultz, W. (1998). Predictive reward signal of dopamine neurons. *Journal of Neurophysiology*, *80*, 1.
- Schultz, W., Apicella, P., & Ljungberg, T. (1993). Responses of monkey dopamine neurons to reward and conditioned stimuli during successive steps of learning a delayed response task. *Journal of Neuroscience*, *13*, 900–913.
- Schultz, W., Romo, R., Ljungberg, T., Mirenovicz, J., Hollerman, J. R., & Dickinson, A. (1995). Reward-related signals carried by dopamine neurons. In J. C. Houk, J. L. Davis, & D. G. Beiser (Eds.), *Models of information processing in the basal ganglia* (pp. 233–248). Cambridge, MA: MIT Press.
- Seamans, J., & Yang, C. (2004). The principal features and mechanisms of dopamine modulation in the prefrontal cortex. *Progress in Neurobiology*, *74*, 1–57.
- Sternberg, S. (1966). High speed scanning in human memory. *Science*, *153*, 652–654.
- Suri, R. E., Bargas, J., & Arbib, M. A. (2001). Modeling functions of striatal dopamine modulation in learning and planning. *Neuroscience*, *103*, 65–85.
- Sutton, R. S. (1988). Learning to predict by the method of temporal differences. *Machine Learning*, *3*, 9–44.
- Sutton, R. S., & Barto, A. G. (1998). *Reinforcement learning: An introduction*. Cambridge, MA: MIT Press.
- Tanaka, S. (2002). Dopamine controls fundamental cognitive operations of multi-target spatial working memory. *Neural Networks*, *15*, 573–582.
- Wallis, J. D., Anderson, K. C., & Miller, E. K. (2001). Single neurons in prefrontal cortex encode abstract rules. *Nature*, *411*, 953–956.
- Wang, X.-J. (1999). Synaptic basis of cortical persistent activity: The importance of NMDA receptors to working memory. *Journal of Neuroscience*, *19*, 9587.
- Wickens, J. (1993). *A theory of the striatum*. Oxford: Pergamon Press.
- Wickens, J. R., Kotter, R., & Alexander, M. E. (1995). Effects of local connectivity on striatal function: simulation and analysis of a model. *Synapse*, *20*, 281–298.

- Widrow, B., & Hoff, M. E. (1960). Adaptive switching circuits. *Institute of Radio Engineers, Western Electronic Show and Convention, Convention Record, Part 4* (pp. 96–104).
- Wise, S. P., Murray, E. A., & Gerfen, C. R. (1996). The frontal cortex-basal ganglia system in primates. *Critical Reviews in Neurobiology, 10*, 317–356.
- Wu, Q., Reith, M., Walker, Q., Kuhn, C., Carroll, F., & Garris, P. (2002). Concurrent autoreceptor-mediated control of dopamine release and uptake during neurotransmission: an in vivo voltammetric study. *Journal of Neuroscience, 22*, 6272–81.
- Yeung, N., Botvinick, M. M., & Cohen, J. D. (2004). The neural basis of error detection: Conflict monitoring and the error-related negativity. *Psychological Review, 111*(4), 931–959.