J.A. Reggia, E. Ruppin and D. Glanzman (Eds.)
Progress in Brain Research, Vol. 121

CHAPTER 19

Dopamine, cognitive control, and schizophrenia: the gating model

Todd S. Braver[1,*] and Jonathan D. Cohen[2]

[1]*Department of Psychology, Washington University, St. Louis, MO, USA*
[2]*Department of Psychology, Princeton University, Princeton, NJ, USA and Department of Psychiatry, University of Pittsburgh, Pittsburgh, PA, USA*

Introduction

The most prominent behavioral impairments in schizophrenia revolve around the failure to control thoughts and actions. These failures of cognitive control are manifest clinically in symptoms such as distractibility, loosening of associations, and disorganized or socially inappropriate behavior. In the laboratory, these disturbances have been observed as deficits of attention (Zubin, 1975; Kornetsky and Orzack, 1978; Wynne et al., 1978; Nuechterlein, 1991; Cornblatt and Keilp, 1994), working memory (Weinberger et al., 1986; Goldman-Rakic, 1991; Park and Holzman, 1992), and behavioral inhibition (Wapner and Krus, 1960; Chapman et al., 1964; Storms and Broen, 1969; Abramczyk et al., 1983; Wysocki and Sweet, 1985; Manschreck, et al., 1988; Carter et al., 1993). However, it is still not well understood what neurobiological and psychological disturbances contribute to cognitive control impairments in schizophrenia.

In this chapter, we set forth a theory of cognitive control that is formalized as a connectionist computational model. The theory suggests explicit neural and psychological mechanisms that contribute to normal cognitive control, and proposes a specific disturbance to these mechanisms which may capture the particular impairments in cognitive control in schizophrenia. To preview, we propose that cognitive control results from interactions between the dopamine (DA) neurotransmitter system and the prefrontal cortex (PFC). Specifically, we suggest that goal-related information, or context, is actively maintained in PFC, and thus serves as a source of top-down support for controlling behavior. We suggest that the DA projection to PFC serves a 'gating' function, by regulating access of context representations into active memory. As such, DA plays an important control function, by enabling flexible updating of active memory in PFC, while retaining protection against interference. Moreover, we suggest that in schizophrenia, the activity of the DA system is noisier, and that this increased variability leads to disturbances in both the updating and maintenance of context information within working memory. Below, we lay out this theory of cognitive control, motivating it from cognitive, computational, and neurobiological perspectives. Following this review, we present two simulations which establish: (1) the model's computational plausibility; and (2) its success at capturing empirical data regarding the behavioral deficits observed in patients with schizophrenia during performance of a simple cognitive

*Corresponding author. e-mail: tbraver@artsci.wustl.edu

A theory of dopaminergic regulation of active memory

PFC and active memory

Cognitive perspectives. The need for a control mechanism in cognition has long been noted within psychology. Virtually all theorists agree that some mechanism is needed to guide, coordinate, and update behavior in a flexible fashion – particularly in novel or complex tasks (Norman and Shallice, 1986). In particular, control over processing requires that information related to behavioral goals be actively represented and maintained, such that these representations can bias behavior in favor of goal-directed activities over temporally-extended periods. Moreover, goal-related information must be: (1) appropriately selected for maintenance; (2) maintained for arbitrary lengths of time; (3) protected against interference; and (4) updated at appropriate junctures. The recognition that active representation and maintenance of goal-related information is a central component of cognitive control can be seen in many theories. Perhaps the most explicit of these is Baddeley's working memory model (Baddeley, 1986), which includes a specific subcomponent, 'the central executive', responsible for utilizing goal-related information in the service of control. The postulation of a cognitive system involved in executive control closely parallels theorizing regarding the nature of frontal lobe function (Bianchi, 1922; Luria, 1969; Damasio, 1985), based on the clinical observation that patients with frontal lesions often exhibit impairments in tasks requiring control over behavior – the so-called 'dysexecutive syndrome'. However, traditional theories have not specified the mechanisms by which the executive operates.

Theories aimed at providing a more explicit computational account of human behavior have also included goal representations as a central component. In production system models, goal states represented in declarative memory are used to coordinate the sequences of production firings involved in complex behaviors (e.g. Anderson, 1983). One critical feature of goal representations in production systems is that they are actively represented and maintained throughout the course of a sequence of behaviors. Moreover, Shallice (Shallice, 1982, 1988; Norman and Shallice, 1986) has relied upon the production system framework in which to put forth his Supervisory Attentional System (SAS) as a mechanism by which complex cognitive processes are coordinated and non-routine actions are selected.

In our own work, we have suggested that the active maintenance of context information is critical for cognitive control (Cohen and Servan-Schreiber, 1992; Cohen, Braver and O'Reilly, 1996; Braver, 1997). We have defined context to be prior task-relevant information that is internally represented in such a form that it can bias selection of the appropriate behavioral response. In particular, we have suggested that representations of context can include task instructions, a specific prior stimulus, or the result of processing a sequence of prior stimuli (e.g. the interpretation resulting from processing a sequence of words in a sentence). Because context representations are maintained on-line, in an active state, they are continually accessible and available to influence processing. Consequently, context can be thought of as a component of working memory, which is commonly defined as the collection of processes responsible for the on-line maintenance and manipulation of information necessary to perform a cognitive task (Baddeley and Hitch, 1994). Context can be viewed as the subset of representations within working memory which govern how other representations are used. Representations of context are particularly important for situations in which there is strong competition for response selection. These response competition situations may arise when the appropriate response is one that is relatively infrequent, or when the inappropriate response is prepotent (such as in the classic Stroop task). In this respect, context representations are closely related to goal representations within production system architectures. Maintenance of internal goal representations, or goal-related knowledge, is critical for initiating the selection of 'weak' behaviors, and for coordinating their execution over temporally-extended periods, while at the same time suppressing competing, possibly more compelling behaviors. Next, we discuss evidence that context information is actively maintained

Neurobiological perspectives. The PFC has long been an area of particular focus for researchers investigating the neural basis of cognitive control. Over a hundred years of neuropsychological studies have provided strong evidence of the involvement of this brain region in the regulation of behavior. In recent years, a large body of converging evidence from neurophysiology and neuroimaging studies have suggested a more specific role for PFC in the active maintenance of task-relevant information. Single-cell recording studies in non-human primates have typically examined the active maintenance properties of PFC through the use of delayed-response paradigms, in which the animal must maintain a representation of a cue stimulus over some delay, in order to respond appropriately at a later point. It is now well-established that during performance of these tasks, populations of neurons in monkey PFC exhibit sustained, stimulus-specific activity during the delay period (Fuster and Alexander, 1971; Kubota and Niki, 1971). The mnemonic properties of these neurons has been demonstrated by showing that both local and reversible lesions to PFC impair task performance, and that performance errors in intact animals are correlated with reduced delay-period activity (Bauer and Fuster, 1976; Fuster, 1973). Neuroimaging studies have confirmed and extended these findings in humans. In the most recent of these studies, PFC activity has been shown to: (1) increase as delay interval increases (Barch et al., 1997); (2) increase as memory load increases (Braver et al., 1997); (3) be sustained over the entire delay interval (Cohen et al., 1997; Courtney et al., 1997).

In addition to these other properties, PFC also appears to be particularly specialized to maintain information in the face of interference, while still allowing for flexible updating of storage. Recently, Miller and colleagues (Miller et al., 1996) have provided direct evidence for this hypothesis. They trained monkeys to respond to repeats of a prespecified cue (e.g., A) when presented with sequences such as A-B-B-A. This task clearly required the ability to identify the cue on each trial, and maintain it across intervening distractors. They observed cue-specific delay period activity for units in both inferotemporal cortex (IT) and PFC following initial presentation. However, subsequent stimuli obliterated this activity in IT, while it was preserved in PFC until a match occurred. The crucial role of PFC in updating and interference-protection can also clearly be seen in studies of PFC pathology. Increased distractibility and perseveration are hallmarks of PFC damage (Milner, 1963; Damasio, 1985; Stuss and Benson, 1986; Owen et al., 1991; Engle et al., 1999), as well as a classic symptom of schizophrenia, which is thought to involve PFC abnormalities (Malmo, 1974; Nuechterlein and Dawson, 1984). Together, these findings support the idea that there are specialized mechanisms within PFC for active memory, as well as for protecting maintained information from both perseveration and interference, and that these mechanisms are disrupted in schizophrenia.

Computational perspectives. From a computational viewpoint there are a number of different processing mechanisms that could support short-term maintenance of information. The most commonly employed and well-understood of these are fixed-point attractor networks (e.g., Hopfield, 1982; Zipser, 1991). Such networks possess recurrent connections, which 'recirculate' activation among units, and are thus capable of supporting sustained activity. The state of such networks typically settles into 'attractors', defined as stable states in which a particular pattern of activity is maintained. Thus, attractors can be used to actively store information. Indeed, a number of computational models of simple maintenance tasks have demonstrated that both physiological and behavioral data regarding PFC function can be captured using an attractor-based scheme (Dehaene and Changeux, 1989; Zipser et al., 1993; Braver et al., 1995; Moody et al., 1998).

However, simple attractor systems have a number of limitations which create problems in more complex maintenance tasks. These limitations can be traced to the fact that the state of an attractor system is determined by its inputs, so that presentation of a new input will drive the system into a new attractor state, overwriting previously stored information (Bengio et al., 1993; Mozer, 1993). Although attractor networks can be configured to display resistance to disruption from input (i.e. hysteresis), this impairs their ability to be updated

in a precise and flexible manner. One way in which attractor networks can overcome these difficulties is through the addition of a gating mechanism. Such systems only respond to inputs, and change their attractor state, when the 'gate' is opened. Computational analyses suggest that gating mechanisms provide the most effective way to stably maintain information in an active state, while retaining the ability for flexible updating. For example, Hochreiter and Schmidhuber (1997) have compared gated recurrent neural networks with other types of attractor systems, and concluded that networks with a gating mechanism were able to learn and perform complex short-term memory tasks better than simple attractor networks, especially when the tasks involved noisy environments, frequent updating, and relatively long periods of storage. The computational studies have suggested that a gated attractor system is the optimal one for active memory. Moreover, the physiological evidence reviewed above is consistent with the hypothesis that PFC implements such a system. Indeed, in previous work, Zipser and colleagues (Zipser, 1991; Zipser et al., 1993; Moody et al., 1998;) have proposed a gated attractor model and have used it to successfully simulate the pattern of delay period activity observed for PFC neurons. However, the Zipser model has not specified the source of the gating signal. In the following section, we suggest that phasic increases in DA activity serve as a gating signal within PFC.

DA modulation of behavior

Schizophrenia. Disturbances to the DA system have long been regarded as central to schizophrenic pathology. Most of the support for this viewpoint comes from observations regarding the therapeutic efficacy of neuroleptics. The finding that the clinical potency of traditional neuroleptics is highly correlated with its affinity for dopamine receptors (Creese et al., 1976), strongly implicates this neurotransmitter in schizophrenic symptomatology. In addition, long-term usage of drugs which stimulate DA activity in the CNS can lead to schizophreniform psychoses (Snyder, 1972). The DA projection to PFC in particular has been a recent focus of attention in schizophrenia research. Specifically, a viewpoint which is rapidly gaining attention is that many of the cognitive impairments seen in schizophrenia are related to reduced DA activity in PFC (Davis et al., 1991; Goldman-Rakic, 1991).

Motor functions. In addition to its involvement in schizophrenia, the DA system has been implicated in a wide range of effects on behavior. The most prominent of these is the linkage of DA with motor function. It is well-established that disturbances to the subcortical DA system cause severe movement related disorders such as Parkinson's disease. Further, stimulants such as amphetamine and apomorphine (which are thought to act by stimulating DA release – Kelly et al., 1975) have clear effects on motor behavior. In animals, these drugs produce consistent changes in both locomotor activity (Segal, 1975), and the repertoire of behaviors exhibited (Norton, 1973), with high doses inducing species-specific stereotypies (Randrup and Munkvad, 1970). There are also many studies documenting the effect of DA activity on response activity in goal-directed tasks, such as operant conditioning paradigms (Heffner and Seiden, 1980; Louilot et al., 1987). A number of investigators have hypothesized that, together, these findings suggest a function for DA in selecting or initiating new motor response patterns (Iversen, 1984; Oades, 1985).

Reward functions. Another function of DA that has commonly been postulated in the literature, is that of processing rewards. This reward-based account of DA activity is supported by findings which suggest a permissive or facilitory role in a number of primary motivated behaviors, such as feeding, drinking and sexual activity (Willner and Scheel-Kruger, 1991). Conversely, spontaneous engagement in these behaviors has been shown to result in increased DA transmission (Heffner et al., 1980). In addition, innumerable studies have shown that the electrical self-stimulation paradigm is primarily dependent on stimulation of DA pathways (Phillips and Fibiger, 1989; Mora and Cobo, 1990). This finding is consistent with the pharmacological evidence that many drugs of addiction act through the DA system (Koob and Bloom, 1988). Taken together, these findings have led some researchers to postulate a crucial role for DA in

conveying information regarding the rewarding or reinforcing properties of specific behaviors (Wise and Rompre, 1989).

Cognitive functions. The literature on the behavioral effects of DA is not limited to studies of motor and reward-related behaviors. There have also been a number of reports of DA effects on cognitive function. In humans, systemic administration of DA agonists have been associated with improvements on various cognitive tasks (Klorman, et al., 1984; Callaway et al., 1994). In particular, the most consistent effects of DA on cognition have been in tasks relying on 'working' or active memory. These sorts of tasks require subjects to maintain relevant contextual information in an active state (i.e. 'on-line'), such that it can be used to mediate the appropriate response. DA effects in working memory have been seen systemically in humans (Luciana et al., 1992; Luciana et al., 1995), and through local manipulations in non-human primates (Brozoski et al., 1979; Sawaguchi and Goldman-Rakic, 1994). These local effects in primates have focused on DA activity selective to PFC. For example, Goldman-Rakic and colleagues have found that pharmacologically blocking DA receptors in circumscribed areas of PFC produced reversible deficits in task performance (Sawaguchi and Goldman-Rakic, 1991). Moreover, microiontophoresis of DA agonists and antagonists, and even DA itself has been found to directly affect the activity patterns of PFC neurons (Sawaguchi, Matsumara, and Kubota, 1990; Sawaguchi and Goldman-Rakic, 1991). Goldman-Rakic and others have concluded from these findings that DA activity serves to modulate the cognitive functions mediated by PFC (Goldman-Rakic, 1991; Cohen and Servan-Schreiber, 1992).

A unitary function? The literature on DA involvement in motor, reward and cognitive functions reveal the wide-spread influence of this neural system on behavior. Further, the disparate nature of these three domains suggests that DA may perform multiple, unrelated behavioral functions. However, another, more parsimonious explanation is also possible: that DA activity plays a unitary function in the central nervous system which is expressed in different domains as a result of its interaction with the different brain systems to which it projects (i.e. striatal, limbic, and cortical). In this chapter, we put forth the hypothesis that DA does, in fact, play a unitary function in behavior. Specifically, we propose that the function of the DA system is to provide a means for the organism to learn about, predict, and respond appropriately to events that lead to reward. The DA system serves this function through simple neuromodulatory effects in the neural populations which it targets. One effect modulates the responsivity of the target neurons to afferent and local input, and the other effect modulates the synaptic strength between the target and these inputs. The DA effects on synaptic strength serve to drive the learning of temporal predictors of reinforcement, while the effects on responsivity serve to transiently bias on-going processing. Most importantly, we propose that through its projection to PFC, the responsivity effect of DA serves to gate access to active memory, while its coincident learning effect allows the system to discover what information must be actively maintained for performance of a given task. In the remainder of this section, we present the following arguments: (1) DA exerts a modulatory effect on target neurons; (2) This effect is of a type that could be exploited to perform a gating function in PFC; and (3) The role of the DA system in reward-prediction learning provides it with particular activation dynamics and timing that are required of a gating signal.

DA as a modulatory signal. A number of lines of evidence suggest that DA acts in a modulatory fashion in PFC that is consistent with a gating role. The PFC is the most densely innervated cortical target of the DA system. Electron microscopy studies of the local connectivity patterns of DA in PFC have revealed that DA typically makes triadic contacts with prefrontal pyramidal cells and excitatory afferents, and also with inhibitory interneurons (Lewis et al., 1992; Williams and Goldman-Rakic, 1993; Sesack et al., 1995). The triadic synaptic complexes formed in PFC suggest that DA can modulate both afferent input and local inhibition. Electrophysiological data support this view, indicating that DA potentiates both afferent excitatory and local inhibitory signals (Chiodo and Berger, 1986; Penit-Soria et al., 1987). In our previous work, we have simulated this potentiating effect of

DA as a multiplicative change in the slope of the activation function of target processing units (Servan-Schreiber et al., 1990; Cohen and Servan-Schreiber, 1993; Braver et al., 1995; Servan-Schreiber et al., 1998). In this work, we assumed that DA effects were prolonged or tonic. However, for the modulatory action of DA to be useful as a gating signal, it must also be both transient and coincident with the occurrence of task-relevant information.

DA as a gating signal. It has been traditionally assumed that neuromodulatory systems (e.g. dopamine, norepinephrine, serotonin) are slow-acting (tonic), diffuse, and non-specific in informational content (Moore and Bloom, 1978). However, there are a number of recent findings that suggest a revision of this view is needed. Detailed primate studies, involving recordings of norepinephrine-producing neurons in locus coeruleus (LC) during task performance, have demonstrated short bursts of LC activity occurring immediately after target, but not non-target presentation (Aston-Jones et al., 1994). Thus, neuromodulatory nuclei may show responses that are rapid, transient, and specific to behaviorally relevant stimuli. Similar findings have been reported for the DA system. Schultz and colleagues (Schultz, 1992; Schultz et al., 1993), recording from ventral tegmental area DA neurons in behaving primates, have observed transient (~100ms in duration, occurring 80–150 ms after stimulus onset) activity in response to novel or behaviorally relevant stimuli. Furthermore, during learning, task stimuli that failed to activate DA neurons on initial presentation came to elicit transient activity when the animal learned their significance for the task. Specifically, Schultz's group observed transient, stimulus-locked activity to cues predictive of reward during performance of a spatial delayed response task (Schultz et al., 1993). Thus, in neurons projecting to PFC, DA activity was observed at precisely the time that the task required information to be gated into, and maintained, in active memory.

DA as a learning signal. At the same time, these findings regarding DA activity dynamics are consistent with a role for DA in reward learning. Indeed, Montague et al. (1996) recently reported a computational model that treats phasic DA activity as a widely distributed error signal, which drives the learning of temporal predictors of reward. In this model, transient changes in the firing of DA neurons represent mismatches between expected and received rewards. These transient changes in activity serve to modulate the strength of target synapses in proportion to the degree and sign of the activity change. The synaptic changes which result from DA firing, along with the appropriate reciprocal connectivity, allow the neural populations targeted by DA neurons to become configured to respond preferentially to stimuli that predict future rewards. The claim that DA modulates changes in synaptic plasticity is one that has recently received support in the neurophysiological literature, and is hypothesized to occur through changes in intracellular calcium (Law-Tho et al., 1995; Wickens et al., 1996; Calabresi et al., 1997). From a computational perspective, the learning effects proposed for DA by Montague et al. (1996) are similar to that proposed in artificial reinforcement learning systems that use a temporal-difference algorithm to predict delayed rewards (Sutton, 1988; Sutton and Barto, 1990). In these systems, the learning chains backward in time, in order to discover successively earlier environmental predictors of reward, until the earliest possible predictor is found that cannot itself be predicted. Intriguingly, the parameter used in Montague et al. (1996) model to simulate the modulatory effects of DA on learning is very similar to the parameter that seems to best describe its effect on neuronal responsivity, in that both are multiplicative in nature. Thus, it could easily be the case that the effects are mediated by different DA receptor subtypes.

A new theory. Taken together, the properties of DA and PFC reviewed above suggest the outlines of a theory regarding the neural and computational mechanisms of cognitive control. In particular, we refine our previous work on active maintenance in PFC by integrating it with the work of Montague et al. (1996) on reward-based learning. This integration provides a means of accounting for the relevant data regarding DA activity dynamics and reward functions as well as the modulatory role of DA in active memory. Specifically, the following refinements are made to our original theory (Cohen and

- DA gates access to active memory in PFC in order to provide flexible updating while retaining interference protection.
- Phasic changes in DA activity mediate gating *and* learning effects in PFC.
- Both effects rely on similar neuromodulatory mechanisms (possibly through different receptor subtypes).
- The gating effect occurs through the transient potentiation of both excitatory afferent and local inhibitory input.
- The learning effect occurs through Hebbian-type modulation of synaptic weights, and is driven by errors between predicted and received rewards.
- The coincidence of the gating and learning signals produces cortical associations between the information being gated, and a triggering of the gating signal in the future.

The power of this new theory is that it provides a framework which may be able to account for specific patterns of normal behavioral performance across a wide-range of tasks requiring cognitive control. At the same time, by making close contact with the known physiological properties of both the DA system and PFC, it may allow for more detailed and biologically realistic explorations of the neural basis of control. In particular, it provides an explicit framework for testing ideas regarding the particular neurobiological disturbances that may underlie schizophrenia and their consequence for behavior.

Because the theory is conceptualized in terms of explicit computational mechanisms, it can be explored through simulation studies. In recent work we have conducted simulations which tested the computational validity of the theory (Braver and Cohen, in press). In particular, we have provided support for the hypothesis that DA implements both gating and learning effects, and that these can work synergistically to provide a mechanism for how cognitive control might be learned through experience. Our simulation demonstrated that control over active maintenance can be achieved by using a gating signal triggered by reward-based learning dynamics. In this simulation, the timing of the gating signal developed as a function of reward-prediction errors in the temporal-difference algorithm. The algorithm enabled the network to chain backward in time to find the earliest predictor of reward, which was a cue stimulus that also had to be maintained in active memory in order to receive the reward. Because this cue triggered a phasic response in the gating/reward-prediction unit, the information provided by the cue was allowed access to active memory.

In this chapter, we present additional simulations which further test the computational plausibility of the theory and its ability to provide new insights into the pathophysiology of schizophrenia. The first simulation manipulated different parameters in the model associated with gating in order to examine their effects on updating, maintenance, and interference protection. The second simulation tests a particular hypothesis regarding whether disturbances to a DA-mediated gating system can account for the pattern of behavioral impairments observed in schizophrenia patients during performance of a simple cognitive control task.

Simulation 1: gating-based updating of active memory

The goal of this simulation was to examine the effects of relevant parameters associated with gating on active memory regulation within a simple network. Three parameters of gating were explored: connectivity, duration, and strength. With respect to connectivity, we examined whether the dual requirements of memory updating and interference-protection place any constraints on the pattern of connectivity between the gating unit (representing the mesocortical DA system) and the context module (representing the PFC). This issue is an important one because it may provide a potential point of contact between the theory and neuroanatomical studies of DA projections in PFC. With respect to the duration of gating unit activity, we examined whether phasic and tonic activity can be functionally distinguished. This is an explicit assumption of the theory that has implications for our understanding of the physiology of the DA system, and the simulation provides an explicit test of that assumption. With respect to the strength of gating unit activity, we examined the relationship between memory updating and interference-protection. These two functions can be thought of as

opposite ends of a continuum. Specifically, new information associated with a gating signal should produce updating, such that old information is deactivated and replaced with new information, whereas new information that is not associated with a gating signal should be prevented from disrupting the current state of maintenance. We tested how manipulating the strength of the gating signal would affect both functions.

Methods

Architecture and Processing. Gating effects on active memory were explored within the context of a simple memory network (see Fig. 1). The network consisted of a context layer (2 units) which received one-to-one excitatory projections (+3.25 weight) from an input layer (2 inputs), representing two separate input conditions (A or B). Each of the units within the context layer had strong self-excitatory connections (+6 weight),which allowed input activity states to be sustained over time. Further, each context unit received two sources of inhibitory input: (a) lateral inhibition from the other memory unit (−3 weight), which produced competition for representations; and (b) local inhibition from a tonically active bias unit (−2.5 weight), which enforced low-levels of baseline activity.

The context layer also received input from a gating unit. The connections to the context layer from this unit were multiplicative (sigma-pi). The function describing the relationship between gating unit activity and its modulatory effect on connection weights was the following:

$$w'_{ij} = \gamma(t) w_{ij}$$

where

$$\gamma(t) = 1 + ((K-1)/(1+e^{-(a(t)-(C/2))}),\ K>1,\ C>=6$$

and $a(t)$ is the activity of the gating unit at time t, C determines the maximum activity level of the gating unit, and K determines the maximum gain (γ) of the gating unit. Thus, if gating unit activity is less than or equal to 0, the gain on the connection strength is equal to 1. If gating unit activity equals C, gain is approximately equal to K. If gating unit

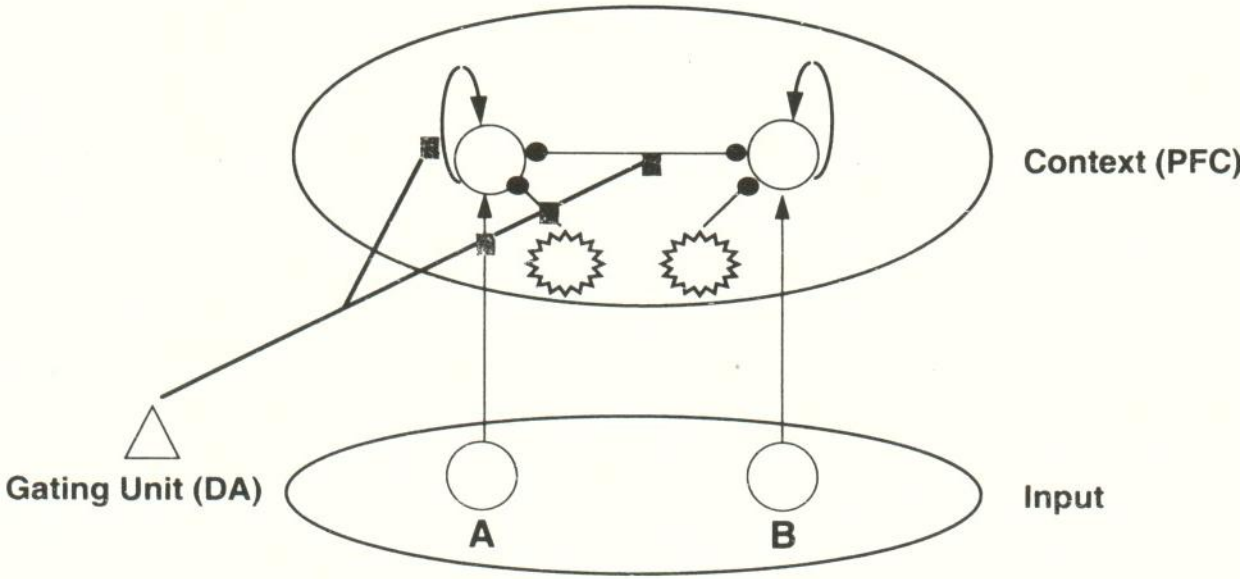

Fig. 1. Simple Attractor Model. This is the architecture of the model used in Simulation 1. Each context unit received two sources of excitatory input (shown with arrowheads) and two sources of inhibitory input (shown with circular heads). Excitatory input was received from a self-connection (recurrent input) and from the input layer (afferent input). Inhibitory input was received from the competing context unit (lateral inhibitory input) and from a tonically active bias unit (local inhibitory input; spiny cell shape). The gating unit (triangular shape) made modulatory connections with context layer inputs (shown with square heads). All four possible gating unit connections were examined in the simulations (shown only for one unit).

activity is greater than 0 and less than C, gain monotonically increases with activity level. In all simulations, C was set to 6, and K was set to 3.

Trials were presented to the network as a sequence of events occurring in the following order: input A, delay, input B, delay. Each input was presented for a duration of 4 time steps, and each delay lasted 50 time steps. Processing occurred continuously over time in the model, with unit activation states governed by temporal difference equations. Specifically, the following equation was used:

$$I_j(t+1) = (\gamma \Sigma w_{ij} y_i + B + I_j(t))\, dt + \sigma Z_i(t)\mathrm{sqrt}(dt)$$

where

$$y_j = 1/(1+e^{-I_j})$$

is the activation of unit j, I_j is the total input to j, dt is the time-step of integration, γ is the gain, B is the bias, $Z_i(t)$ is a standard independent Gaussian random variable, and σ is the variance of the distribution. In the simulations described below, dt was set at 0.5.

Simulations and analysis. The role of three factors were examined in the model: (1) connectivity between the gating unit and memory layer; (2) duration of gating unit activity (tonic vs. phasic); and (3) strength of gating unit activity. The primary hypothesis with respect to connectivity was that gating effects would be optimal when the gating unit made connections only to the afferent and local inhibitory inputs of the context layer. However, the impact of connectivity between the other two input projections (self-excitatory and lateral inhibitory) was also investigated. The first set of simulations examined this question, by looking at effects in the context layer as a result of different patterns of gating unit connectivity. The first simulation also examined differences between tonic and phasic gating unit activity. In particular, three conditions were examined: phasic gating activity during stimulus presentation, no gating during stimulus presentation, and tonic gating activity during delay (i.e. no stimulus presentation). The first condition tested whether a gating signal presented simultaneously with stimulus presentation would produce memory updating. The second condition tested whether the absence of a gating signal during stimulus presentation would prevent memory updating (i.e. interference). The third condition tested the effects of a tonic increase in gating unit activity during the delay period. A second set of simulations manipulated the strength of gating unit activity levels in order to examine the effect of this factor on the active maintenance of context. Moreover, the effects of gating activity strength were examined under both tonic and phasic conditions. In the first simulation, processing was deterministic ($\sigma = 0$). However, in the second simulation, zero-mean Gaussian noise was added to the net input of the context units on every time step ($\sigma = 0.95$), in order to produce variability in processing. One thousand trials were simulated at 10 levels of gating unit activity strength. In the phasic condition, strength varied from 0.0 C to 1.0 C in 0.1 C increments. In the tonic condition, strength varied from 0.0 C to 0.5 C in 0.1 C increments.

In both sets of simulations, phasic activity in the gating unit was simulated by setting the activity level to its maximum value (C) for 2 time steps; otherwise activity was set at 0. Gating occurred during the middle 2 time steps of stimulus presentation. In the first set of simulations, tonic activity in the gating unit was simulated by setting the activity of the gating unit to 50% of its maximum value throughout the last 50 time steps of the delay period. Each simulation run was analyzed by computing the percentage of trials that each context unit was active (greater than 0.5), for every time step of the trial.

Results

The first set of simulations demonstrated that, as expected, gating unit connectivity on both afferent and local inhibitory inputs provided both interference protection as well as updating. As Fig. 2 shows, a new input that is not associated with a gating signal does not disrupt the maintained state, however, if that input is accompanied by gating unit activity, the new information replaces the old state in memory, thus providing successful updating. This dynamics is directly a function of the gating connectivity. The potentiating effect of gating on the inhibitory bias deactivates the current state,

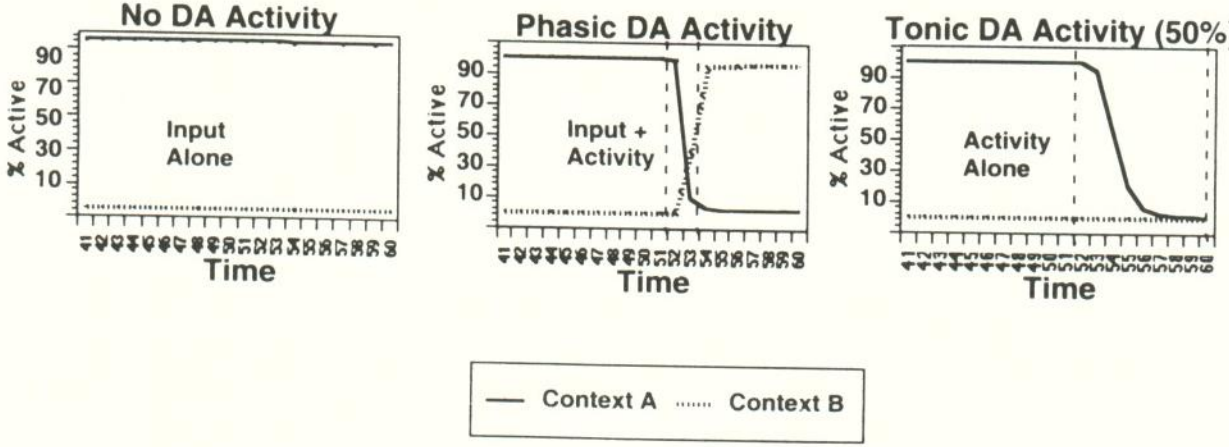

Fig. 2. Gating Effects on Activation Dynamics. These plots illustrate the effects of gating on maintained context information (Context A, solid lines). Afferent input corresponding to a competing context representation (Context B, dashed lines) not accompanied by a gating signal does not disrupt the state of the context layer (left panel), but if the input is accompanied by a gating signal (middle panel), the state of the context layer is updated and a new context representation is maintained. Tonic activity in the gating unit has an inhibitory effect on context activity, causing it to decay (right panel). Vertical dashed lines denote the period during which the gating unit is active. Vertical axis corresponds to the percent of trials in which each context unit was active (greater than 0.5) for the corresponding timestep.

while the potentiation of afferent activity allows the new state to gain access to memory. Interestingly, sustained increases in gating unit activity also had a deactivating effect; when the gating unit was tonically active in the absence of afferent input, it acted in a purely inhibitory manner, causing the maintained state to decay away.

Furthermore, the simulations suggested that this pattern of connectivity from the gating unit – affecting local inhibition and afferent input – was the optimal one. Adding the connection to the lateral inhibitory input did not have any additional affect on activity dynamics. This makes sense in that, during a gating signal, potentiating the lateral inhibitory input would not lead to further deactivation of the old information. Moreover, it would only serve to increase the difficulty by which new information could gain access to memory, by increasing the competition from old information prior to deactivation. Adding a gating connection to the self-excitatory connections in the context layer also had a deleterious effect, in that it prevented updating; this occurred because the potentiation of the self-excitatory input negated the potentiating effect on inhibitory bias. Because the self-excitatory weight was stronger than the inhibitory bias (which is a necessary relationship in order to enable self-sustaining activity), when the gating signal potentiated both connections it served to increase, rather than decrease the net excitatory input to the already active context unit. This prevented the unit from deactivating, and produced a state of increased competition that prevented the new information from replacing the old.

The second simulations demonstrated that gating effects are graded in nature, and are modulated by the strength of the gating signal. As shown in Fig. 3, manipulating the strength of the gating signal affected both the robustness of updating and interference protection. As the gating signal decreased from its maximal value, there was a decreased probability that the new information would replace the previous state. Conversely, as the gating signal increased from its minimum value (i.e. 0), there was an increased probability that information would interfere with the current state. Moreover, as the results make clear, updating and interference lie on a continuum defined by the gating signal; in other words, it is the gating signal that defines which information is relevant or irrelevant. Irrelevant information is information that should not be accompanied by a gating signal.

However, if a partial signal does occur to the information, it will have the opportunity to produce interference, by disrupting the currently maintained state. On the other hand, for information that is task-relevant, a gating signal should occur synchronously with its presentation. If this gating signal is reduced, the new information will not as reliably update active memory, and thus, the previous state will persevere. The simulations also demonstrated that tonic activity in the gating unit has a distinctly different effect than that of phasic activity. Specifically, tonic gating unit activity primarily impacts the robustness of active maintenance (see Fig. 3). When the gating unit has low tonic activity, information can be maintained reliably; however, as tonic activity increases above baseline value it causes memory decay. Within a certain range, the effect interacts with time, such that the longer the delay, the greater the memory decay.

Discussion

The simulation results further establish the computational plausibility of the theory we have put forth regarding the functional role of the gating mechanism in cognitive control, while touching on important issues with regard to the anatomy and physiology of the DA system and their functional consequences. In the simulations, gating occurred through phasic changes in activity that occurred

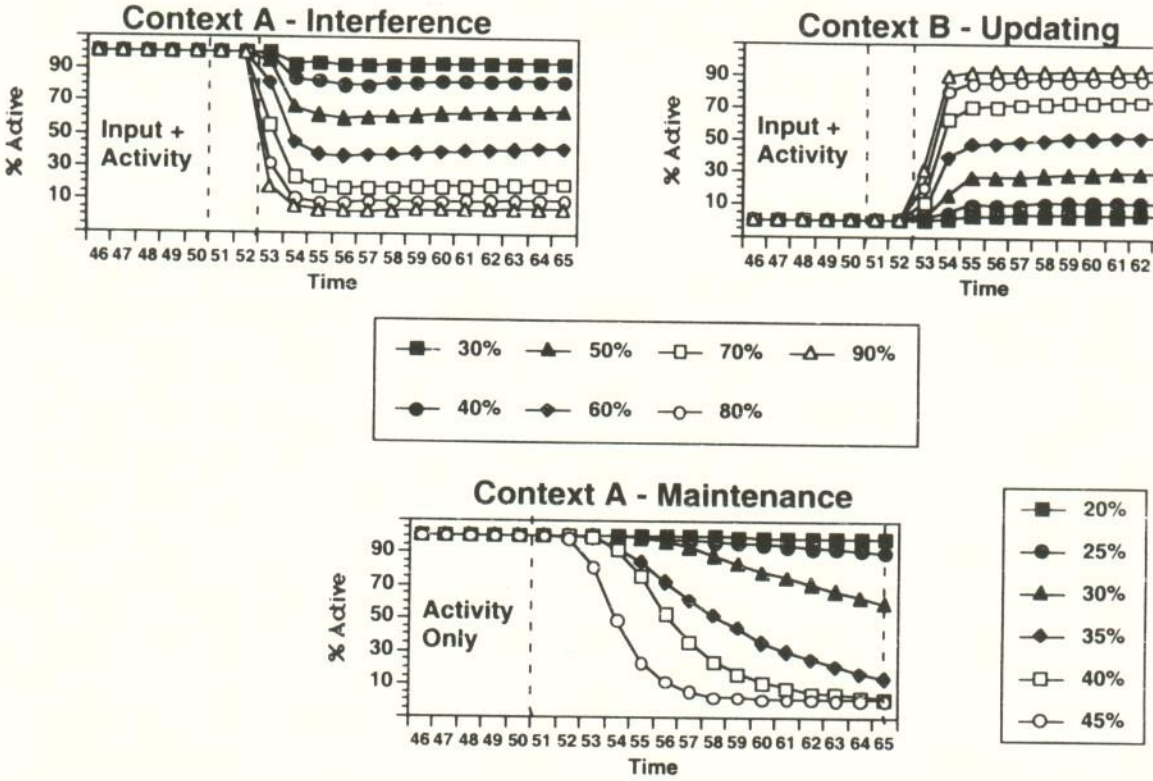

Fig. 3. Parametric Effects on Updating, Interference and Maintenance. The upper two plots show the parametric relationship between updating and interference. In the upper left panel, the maintained context representation (Context A) becomes progressively more perturbed as the strength of the gating signal accompanying afferent input is increased. In the upper right panel, updating to the competing context representation (Context B) becomes progressively more reliable as gating signal strength increases. In the lower panel, the maintained context representation shows progressively greater memory decay as tonic gating unit activity increases. Vertical dashed lines denote the period during which the gating unit is active. The vertical axis corresponds to the percent of trials in which each context unit was active (greater than 0.5) for the corresponding timestep. Each line represents a different value of gating unit strength (expressed as a percentage of baseline strength).

synchronously with task-relevant information, but were absent with information that was not relevant for performance. This pattern of activity dynamics is consistent with the findings from midbrain DA neurons recorded in behaving primates (Schultz, 1992). The effect of gating was one of potentiation, multiplying the strength of both excitatory and inhibitory signals. This effect also matches what has been observed physiologically, through recordings of the post-synaptic effects of DA activity on target cells (Chiodo and Berger, 1986). Moreover, these functions were found to be optimal when the gating unit modulated both excitatory afferent and local inhibitory connections in the context layer. This connectivity pattern is similar to what has been found in anatomical studies, where DA has been observed to affect both pyramidal cells and local inhibitory interneurons in PFC (Goldman-Rakic et al., 1989; Sesack et al., 1995). The simulations may also account for relevant physiological data in terms of the functional distinctions between tonic and phasic activity in the gating unit. Tonic gating unit activity was found to have a primarily inhibitory effect on active memory, causing deactivation of previously maintained information. This effect might provide a resolution of conflicting findings within the neurophysiological literature on DA actions in PFC. In particular, DA has often been observed to have an inhibitory effect on activity, rather than a potentiating effect (e.g. Ferron et al., 1984). Based on the simulation results, DA should have an inhibitory effect unless it occurs synchronously with afferent input.

Although the simulations demonstrate that the dynamics and connectivity of the gating system capture a great deal of the known physiological and anatomical data regarding the DA system in PFC, there are still unresolved issues regarding the nature of the gating mechanism and its relationship to DA. In particular, phasic *decreases* in DA activity have also been observed in the studies of Schultz et al. (1993), and have been interpreted as having functional consequences. Specifically, these decreases in activity occurred in situations where reward was withheld, such as when the animal responded incorrectly. This finding is perfectly consistent with a reinforcement learning account of DA function, in that situations in which a reward was predicted but not received should lead to a temporal-difference error that would be manifest as a decrease in activity. However, it is not clear whether this may also have functional consequences for the gating processes of DA. In the current simulations, a phasic decrease in activity (which would be modelled as negative gating unit activity) did not have any functional consequences, as the multiplicative effect of the gating unit was 1 for activity levels equal to or less than zero. Yet one could imagine how it would be adaptive for phasically decreased DA activity to play a functional role in gating. In particular, given that decreases in activity typically reflect errors, it might be useful to have this signal serve some type of 'reset' function in active memory. To take a real world example, in the Wisconsin Card Sort Test (WCST), receiving error feedback from the experimenter should produce a reset in the currently maintained sorting category, in order to provide new categories with an equal chance to be maintained as the sorting rule. Indeed, a similar type of mechanism has been invoked in previous computational models of the WCST in order to account for the performance data of normal and frontal patients in the task (Levine and Prueitt, 1989; Dehaene and Changeux, 1992).

Interestingly, in exploratory simulations, we observed that using a function which transforms negative gating unit activity into a multiplicative gain less than 1 (i.e. causing a reduction in connection strength) might produce similar types of behavior in the context layer of the model. Specifically, this behavior occurred when the gating unit also affected the lateral inhibitory connections between context units. In this situation, a phasic decrease in activity (which caused the multiplicative gain to go to 0) effectively shut-out the input from all connections except for the self-excitatory ones. Since the self-excitatory connections were strong, both context units rapidly increased their activation level, leading to a transient state where both contexts were simultaneously being maintained. Once the gating unit activity returned to baseline, the lateral inhibitory effects resumed, and the units competed for

equally and highly active when competition resumed, each context had equal opportunity to win the competition. Thus, the phasic decrease in gating activity provided a type of reset function, by allowing all contexts an equal chance to be maintained. Although these effects were not explored in great detail, they warrant further investigation in simulations to determine whether they might capture an additional control function of the DA system in active memory.

Another aspect of the gating mechanism which may not fully capture known properties of the mesocortical DA system is the nature of its inhibitory effects in the context layer. In particular, the anatomical data on DA projections in PFC suggest that the inhibitory actions of DA may occur indirectly, by potentiating the excitatory inputs to local inhibitory neurons (Sesack et al., 1995). In contrast, the current simulations assumed a more direct inhibitory effect of the gating mechanism, by having it modulate the output connection of the bias unit onto the context unit. Moreover, in the simulations the bias units were assumed to be tonically active, so their activity level was not dependent on excitatory input. It is likely that both the direct and indirect connectivity patterns would produce similar inhibitory effects, but this cannot be stated confidently without further exploration through simulations. One possible approach to more fully capturing these anatomical constraints would be to postulate that activity in the bias unit is driven by an external source. For example, connectivity in the model could be modified such that the bias unit and context unit were mutually connected in a negative feedback circuit. In this connectivity pattern, the gating unit would potentiate the output of the context unit and increase the negative feedback provided by the bias unit. Simulations could determine whether this pattern produces the appropriate inhibitory dynamics that allow for updating.

The second set of simulations conducted in this study provided useful insights into the parametric relationship between gating unit activity and active maintenance. In particular, the simulations demonstrated three significant effects: (1) Reduced phasic activity during the presentation of 'task-relevant' stimuli leads to perseveratory behavior, by decreasing the probability that the previous context will be replaced by the current context; (2) Increased phasic activity during the presentation of 'irrelevant' stimuli produces interference effects, by increasing the probability that these stimuli will disrupt the currently maintained context; and (3) Increased tonic activity during delay periods produces a delay-related decay of active memory, by increasing the probability that the current context will deactivate over time. Together, these three effects may provide a potential model of the effects of DA impairment on cognitive control. Indeed, perseveration, poor interference control and maintenance deficits are all three symptoms that have commonly been associated with schizophrenia (Malmo, 1974; Nuechterlein and Dawson, 1984). The theory provides an explicit mechanism which might explain how these symptoms arise. In the next simulation, we test this idea directly, by incorporating the gating mechanism into a model of performance on a simple cognitive control task. We examine whether disturbances to this mechanism can be used to account for the patterns of behavioral impairments observed by schizophrenia patients.

Simulation 2: behavioral impairments in schizophrenic patients

In this simulation, we test whether the disturbances to a DA-mediated gating mechanism can account for the cognitive control impairments seen in patients with schizophrenia. In particular, the model suggests that the pattern of deficits observed in patients are consistent with both decreased phasic and increased tonic DA activity. The gating hypothesis predicts that tonically increased DA activity should produce deficits in active or working memory, while decreased phasic DA activity should produce perseveration and interference-effects. Here, we directly test these predictions by conducting simulations of behavioral performance on a simple cognitive control task which requires both active maintenance and frequent updating of context information. The task is an 'AX' variant of the Continuous Performance Test (CPT, Rosvold et al., 1956). We have collected extensive behavioral data regarding the performance of both healthy subjects and patients with schizophrenia on this

task (Servan-Schreiber et al., 1996; Braver, 1997; Cohen et al., 1999). Simulations of behavioral data were conducted by adding a gating mechanism to an existing computational model of the task (Braver, Cohen, and Servan-Schreiber, 1995; Cohen, Braver, and O'Reilly, 1996; Braver, 1997).

Methods

Task. In the AX–CPT, single letters are visually presented as a sequence of cue-probe pairs (Fig. 4). A target response is required to a specific probe letter (X), but only when it follows a designated cue (A). A manipulation of the delay interval between cue and probe (1 s short delay vs. 5 s long delay) enables an examination of active memory demands. In addition, target trials occur with high frequency (70%), which enables examination of the role of context in biasing response competition and inhibiting response prepotencies. Specifically, control over processing can be examined in the three types of non-target trials, which occur with 10% frequency each (BX, AY, and BY, where 'B' corresponds to any non-A stimulus, and 'Y' to any non-X). Context information must be used on BX trials in order to inhibit the prepotent tendency to make a target response to the X. In contrast, context acts to bias incorrect responding on AY trials, since the presence of the A sets up a strong expectancy to make a target response to the probe. BY trials provide an index of performance in the absence of response competition.

Behavioral data. The data for this simulation were taken from a study first presented in Braver (1997); participants in the study were 16 DSM-IV schizophrenia patients and 16 matched controls. Patients were neuroleptic-naive and experiencing

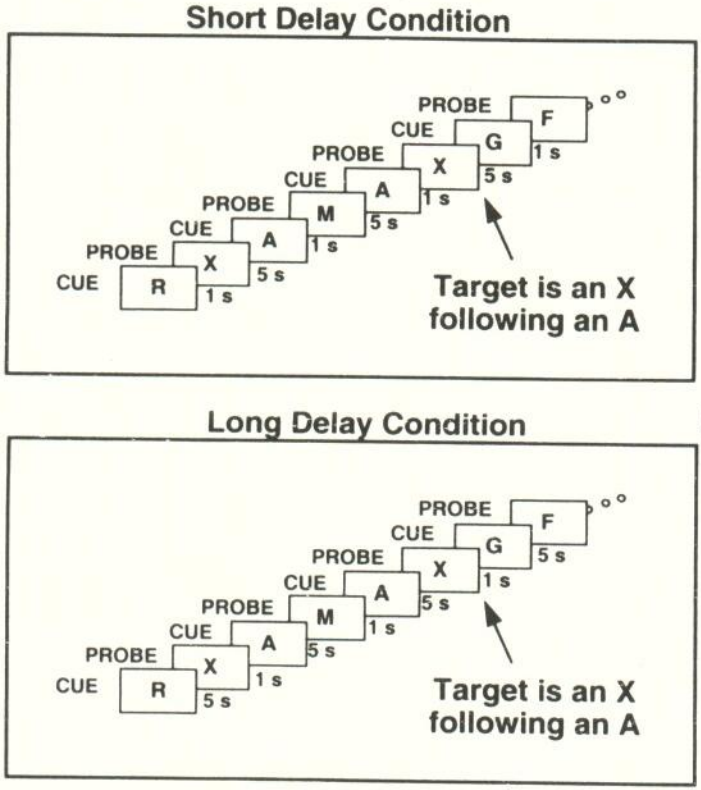

	CUE: VALID	CUE: INVALID
PROBE: TARGET	A-X (70%)	"B-X" (10%)
PROBE: NONTARGET	"A-Y" (10%)	"B-Y" (10%)

Fig. 4. The AX–CPT task. Trials consist of single letters occurring as sequences of cue-probe pairs. In the short delay condition, the delay period is 1s, intertrial interval is 5s. In the long delay conditions, the delay period is 5s, intertrial interval is 1s. A target is defined to be an X immediately following an A. Targets occur with 70% frequency, and the three other trial types (AY,BX, BY) each occur with 10% frequency.

their first hospitalization for psychotic symptoms. Consequently, they formed a select subgroup of participants which are free of many of the confounds and complications associated with studying schizophrenia patients (e.g. medication, chronicity, or institutionalization effects). Both groups performed 200 trials of the AX-CPT evenly divided between short and long delay conditions. Intertrial interval was counterbalanced so that total trial duration was equated across delay conditions. Participants pressed one button of a response box for target probes and a separate button for nontargets. Both accuracy and reaction time data were collected. There were two primary behavioral measures of interest. The first, context sensitivity, indexed the ability to respond correctly to an X probe based on its prior context. Context sensitivity was computed by comparing AX hits to BX false alarms, using the d′ function. The second measure, context cost, indexed the degree of response slowing on non-target trials due to the presence of an A cue. Context cost was computed by calculating the difference in reaction time in AY trials relative to BY trials; these two measures were calculated separately for the long and short conditions in each group.

Computational model. A gating mechanism was incorporated into an existing computational model of the AX-CPT. The original model was found to successfully capture many aspects of both normal and schizophrenic performance in the task (Braver et al., 1995; Cohen et al., 1996; Braver, 1997). The addition of a gating mechanism provided a means to check whether the new model could also account for performance by incorporating a more refined model of DA activity. The architecture of the model is shown in Fig. 5. The model consisted of a direct pathway composed of feed-forward connections between a pool of input units, representing the four stimulus conditions (A, B, X or Y), a pool of four associative units (representing the two possible associations – target or nontarget – activated for each probe stimulus), and a pool of two output units. In addition, the cue inputs also projected to a

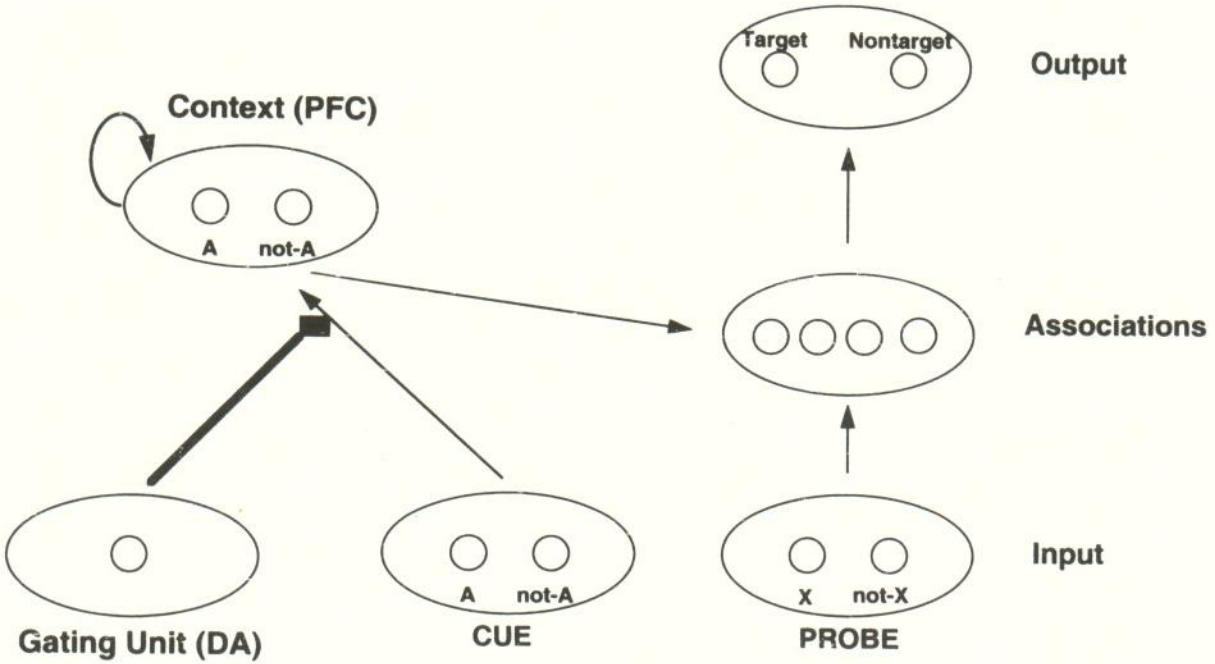

Fig. 5. Diagram of Gating Model. Architecture of model used to simulate the AX–CPT task. Units in the context layer have self-excitatory connections, which provide a mechanism for active maintenance. The gating unit makes a multiplicative connection with both afferent excitatory and local inhibitory (not shown) inputs to the context layer.

layer of context units. The context layer then projected back to the pool of associative units in the direct pathway. Units within the context layer had strong (non-modifiable) self-excitatory connections (+6.0 weight) which provided a mechanism for active maintenance. Additionally, within each pool of units, there were lateral inhibitory connections which produced competition for representations. Finally, each unit was associated with a local inhibitory unit which provided a tonic negative bias (−2.5 weight) on baseline activity states.

Processing evolved continuously over time in the model according to the temporal difference equation described in Simulation 1. The duration of relevant events within the simulation (e.g. cue and probe presentation, delay periods) were scaled to approximate the temporal relationships used in the actual task. Thus, the cue and probe were each presented for 2 time steps, the short delay lasted 7 time steps, and the long delay lasted 33 time steps. The presentation of each stimulus was simulated by adding an external source of activation (i.e. soft-clamping) to units in the input layer for a short duration. Input activation states were then allowed to evolve in response to this external input. All input units were provided this external source of activation during presentation of every stimulus, in order to approximate the effects of distributed representations, and lateral competition at the sensory stage of processing.

Network weights were developed through a backpropagation training procedure consisting of repeated presentations of each of the 8 different trial types of the AX–CPT (AX, AY, BX and BY at both short and long delays), with the presentation frequency of each type matching that of the behavioral task. This learning approach enabled optimization of weight strengths based on both the constraints of task performance and the relative frequencies of task events. Gating was added to the trained model by including an additional unit that had modulatory effects on the local inhibitory and afferent excitatory connections to the context layer which were identical to those described in Simulation 1. As in Simulation 1, the input connections to the gating unit were not trained. Rather, these connections were assumed to already have been learned (a simulation demonstrating how these connections might develop through learning is provided in Braver and Cohen, in press). The only other addition to the model was that the input-to-context connections were adjusted so that the presence of external input alone was strong enough to activate the context module when it was in a resting state (i.e. when no other units in the pool were active), but not strong enough to update it from an active state (i.e. when a competing unit in the pool was already activated).

Simulations. One thousand trials of each of the eight stimulus conditions (4 trial types × 2 delays) were simulated in both the intact and impaired models. Trials were presented to the model as a continuous sequence of events occurring in the following order: cue, delay, probe, ITI. The gating unit became transiently activated during presentation of the cue and probe stimuli. Simulations of performance on each condition were conducted by determining which of the output units was the first to surpass a prespecified threshold value, and then collecting accuracy and RT statistics across each trial. Noise was added to each unit's activation state on each time step in order to simulate variability in processing. Both the noise and threshold parameters were fixed at the levels derived for the original model (noise = 0.95; threshold = 0.65). In order to simulate the disturbances in the mesocortical DA system thought to be present in schizophrenia, we disturbed the activity of the gating unit in the model. The specific disturbance that was implemented was to further increase the noise level in gating unit activity (to a value 5 times that of the rest of the units). This pattern of disturbance causes changes in both tonic and phasic activity levels, as a result of the function which relates gating unit activity to its multiplicative effects on synaptic strength. Specifically, because the function is bounded and monotonic (i.e. a logistic), increases in noise will raise the mean value of gain for baseline (low) levels of gating unit activity (i.e. tonic gating) and decrease the mean value of gain for high levels of gating unit activity (i.e. phasic gating).

Results

Behavioral Data. The behavioral data are shown in Fig. 6. For healthy controls, sensitivity to context

was relatively high ($d' > 3$). Moreover, there were no significant effects of delay on sensitivity. Conversely, the cost of maintaining context was also relatively high in terms of RT slowing (~140 ms), and also did not decrease much with delay. In contrast, in patients with schizophrenia both context sensitivity and context cost were significantly reduced. These effects further interacted with delay, so that the difference between patients and controls was greatest at the long delay. Thus, the performance data suggest patients showed impairments in both the representation and maintenance of context information. Furthermore, the pattern of performance elicited by patients in this task also provides evidence that patients suffer from a specific impairment in cognitive control, rather than a more general deficit pattern (e.g. Chapman and Chapman, 1978). This pattern can be observed by noting that the context disturbance exhibited by patients actually results in a relative benefit in performance, since they show *less* of a context cost, manifest as less response slowing to AY trials relative to BY trials.

Simulation data. The simulations were able to successfully capture the qualitative pattern of the behavioral data (see Fig. 6) and context sensitivity and context cost were both high in the intact model but decreased in the noisy gating model. Moreover, these effects also replicated the interaction with delay observed empirically; in particular, the difference between the two models was greatest at the long delay for both measures. An examination of the dynamics of activity in the context layer during the delay interval revealed the mechanism for these effects and in particular, it was found that in the noisy gating model, there was an increased failure

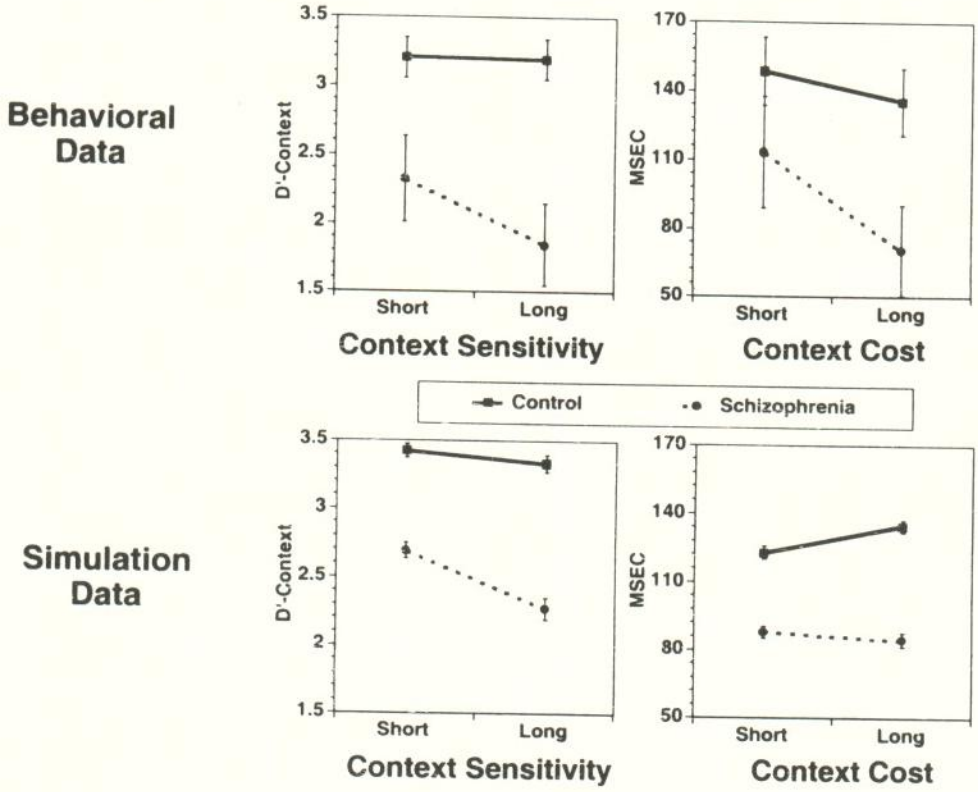

Fig. 6. AX–CPT Data: Behavioral and Simulation. These figures show data for both context sensitivity and context cost performance measures for controls and patients with schizophrenia. The upper plots show the behavioral data and the lower plots show simulation data. The simulation captures both the overall reduction of context sensitivity and context cost in schizophrenia, as well as the interaction with delay.

for the context representation to update following presentation of the cue. Additionally, even in the trials in which the correct representation was properly activated, the maintenance of this representation was less robust. There was an increased tendency for the representation to decay away, and this tendency accumulated over the delay.

Discussion

The results of this study suggest that the gating model of the AX-CPT task was able to successfully capture the specific pattern of behavioral performance observed both in healthy controls and in patients in schizophrenia. Thus, the model compares favorably to the AX–CPT model developed previously, which also accounted for this dataset (Braver, 1997). However, the current model significantly refines and extends the account of the mechanisms hypothesized to underlie schizophrenic deficits in task performance. The earlier model accounted for these performance deficits by suggesting that, in schizophrenia, DA activity is tonically reduced in PFC. In the current model, the mechanism responsible for producing AX–CPT performance deficits is increased noise levels in mesocortical DA. This particular disturbance resulted in both increased tonic activity and decreased phasic activity in the gating system. As suggested by the results of Simulation 1, the increased tonic activity produces deficits in the maintenance of context, while the decreased phasic activity produces in deficits in updating the representation of context.

The functional distinction in the model between disturbances in phasic and tonic DA activity is an important advance in the theoretical account of the pathophysiology of schizophrenia. It is worth noting that the gating account also appears to be more consistent with neurobiological data regarding the etiology of the disease. In particular, Grace (1991) has postulated that schizophrenia is associated with disturbances in both tonic and phasic DA activity, based on an analysis of neuroleptic-effects on DA physiology. Importantly, however, Grace's model predicts that patients with schizophrenia suffer from increased phasic and decreased tonic DA, which is exactly opposite to the account provided by the current model. Thus, further work will be needed to examine these two models in greater detail in order to determine which provides a better account of the data.

In the current simulation, a single disturbance – increased noise levels in gating unit activity – was found to capture the pattern of performance deficits exhibited by patients in the AX–CPT. This occurred because increasing gating unit noise affected both tonic and phasic activity levels. However, the model also holds open the possibility that tonic and phasic DA activity can be independently affected by different mechanisms. Moreover, since tonic DA activity is associated with the active maintenance of context, and phasic DA activity is associated with the updating of context, the model also suggests that deficits in these two processes are also dissociable, at least in principle. This raises the intriguing possibility that different patient subgroups might suffer from independent disturbances in these two components of DA function. If patients from both subgroups were present in the data set, the averaged results would appear as if both deficits were present. This hypothesis could be tested by examining the clinical symptomatology of patients more closely, to examine whether there are relationships between different symptom subtypes and the prevalence of disturbances in updating vs. maintenance of context information. In particular, a specific disturbance in context updating would be revealed as reduced context sensitivity and context cost, but no effect of delay on performance. A specific disturbance in context maintenance would be revealed as normal performance levels at the short delay, but a significant effect of delay, such that both context sensitivity and context cost are reduced at the long delay. Thus, the model provides a means of relating clinical heterogeneity to particular neurobiological mechanisms.

Another advance of the current model over our previous model is that it can potentially account for both normal and schizophrenic behavioral data in a much wider range of cognitive control tasks. In particular, the model suggests how context information in PFC can be actively maintained in the face of interference, and how this function might be disturbed in schizophrenia. In the model, the degree of interference produced by irrelevant items is

directly related to the degree of phasic DA activity that occurs synchronously with the presentation of each item. If phasic DA responses to irrelevant items are increased in patients with schizophrenia, this could account for the increased susceptibility to interference-effects from distractors that are so commonly observed in experimental studies (e.g. Neuchterlein and Dawson, 1984).

In recent preliminary work, we have begun to demonstrate how such a mechanism could also account for interference effects in normal behavior. Specifically, we conducted simulations to try to capture the pattern of behavioral data we observed in a study with healthy subjects performing an interference version of the AX–CPT (Braver, 1997). In terms of performance, we found that irrelevant items presented during the delay period of AX–CPT trials produced significant interference effects when these items were very similar to the cue and probe stimuli (i.e. identical letters presented in different colors). Indeed, the pattern of behavioral impairments observed in healthy subjects under interference conditions was very similar to that observed by patients with schizophrenia (i.e. reduced context sensitivity and context cost). In simulations of this task using the gating model, we found that we could account for these patterns of behavioral impairment by assuming that irrelevant items were associated with a partial gating signal (Braver, Cohen and McClelland, 1997). We made this assumption by hypothesizing that the high degree of featural similarity between relevant and irrelevant stimuli would result in overgeneralization of DA activity. There is support for this assumption from physiological recordings of DA neuronal activity in behaving primates. The primate data suggest that DA neurons do exhibit partial responses to stimuli that are similar to reward-predictive cues (Mirenowicz and Schultz, 1996).

General discussion and conclusions

The studies presented in this paper establish the computational and empirical plausibility of a new theory regarding the role of DA-mediated gating in normal cognitive control and cognitive control impairments in schizophrenia. In particular, the theory postulates that control over the representation and maintenance of goal-related context occurs through phasic gating signals which arise as a consequence of the dynamics of reward-based learning. Because the theory was conceptualized in terms of explicit computational mechanisms, it was possible to investigate these theoretical hypotheses through simulation studies. The simulations demonstrated that: (a) the gating signal can provide a mechanism for both interference protection and flexible updating of stored information; and (b) disturbances to the gating mechanism can account for the behavioral impairments observed in schizophrenia patients during performance of a simple cognitive control task.

As the above discussions has indicated, the simulation results have important implications for: (1) the neurobiology of DA function; (2) pathophysiology and cognitive deficits in schizophrenia; and (3) the nature and mechanisms of cognitive control. With regard to DA function, the theory postulates that dopamine's role in behavior is a unified one, that exploits simple neuromodulatory effects to modulate both learning and on-going processing. Moreover, the simulations make intriguing suggestions regarding constraints on anatomical connectivity (i.e. DA does not project to local recurrent inputs) and differing postsynaptic effects of tonic and phasic activity (i.e. tonic DA effects are inhibitory in nature while phasic DA effects are modulatory in nature). With regards to schizophrenia, the theory suggests that patients suffer from both tonic and phasic DA disturbances and that these may contribute to maintenance and updating deficits, respectively. Additionally, this hypothesis may help to provide a unified account that can explain why patients appear to suffer from three otherwise unrelated impairments: perseveration and switching problems (e.g. Malmo, 1974; Frith and Done, 1983), distractibility and susceptibility to interference (e.g. Nuechterlein and Dawson, 1984), and memory failures (e.g. Servan-Schreiber et al., 1996). The theory may also contribute to an understanding of the pathophysiology of schizophrenia by making close contact with the known physiological properties of both the DA system and PFC. In particular, by specifying the behavioral consequences of detailed physiological disturbances (i.e. tonic vs. phasic DA dysfunction),

the theory can provide a crucial point of contact between behavioral and basic neuroscience research on schizophrenia. This may lead to the development of more refined animal models, and to new ways of examining hypotheses drawn from neurobiologically-based research (e.g. Grace, 1991).

The theory presented here also has the potential to provide a new account of normal cognitive control in terms of neurobiologically plausible mechanisms. In so doing, it provides an illustration of how a system built of simple processing elements can learn to regulate its own behavior in an intelligent and adaptive fashion, without invoking the perennial problem of the homunculus. In particular, this theory extends and refines the account of PFC, DA, and cognitive control developed in our previous modelling efforts (Cohen and Servan-Schreiber, 1992; Braver et al., 1995). Here, we specify the mechanisms by which PFC is able to control processing in a top-down manner, while at the same time remaining responsive to bottom-up input from other parts of the system. We hypothesize that this interplay of bottom-up and top-down processing is mediated by the DA system, through its dual effects on gating and reward-prediction. As a consequence of these dual effects of DA, control emerges in the system through the dynamics of 'regulated interactivity'. DA provides the system with the ability to learn when to trigger the gating signal, thereby controlling the contents of active memory in PFC. In turn, this allows the PFC to use its representations as context for biasing processing in rest of the system.

The claims we have made in this paper are bold and still somewhat speculative. Much work remains to be done in both the development and validation of the theory presented here. The effort to understand the function of DA and its role in schizophrenia promises to be a complex endeavor, and will require the most powerful conceptual tools we have available. We believe that computational modeling represents one such tool. By pursuing our ideas within a computational framework, it is possible to make them explicit within simulation models. This not only provides a check on their conceptual validity, but also provides a means for exploring, in detail, their implications for behavior.

Success in this effort would not only provide insights into the mechanisms underlying some of our highest and uniquely human faculties, but might also allow us to understand better with how these faculties can break down in a disease such as schizophrenia, which has such devastating consequences for behavior.

References

Abramczyk, R.R., Jordan, D.E. and Hegel, M. (1983) 'Reverse' stroop effect in the performance of schizophrenics. *Percept. Motor Skills*, 56, 99–106.

Anderson, J.R. (1983) *The Architecture of Cognition*. Cambridge, MA: Harvard University Press.

Aston-Jones, G., Rajkowski, J., Kubiak, P. and Alexinsky, T. (1994) Locus coeruleus neurons in monkey are selectively activated by attended cues in a vigilance task. *J. Neurosci.*, 14(7), 4467–4480.

Baddeley, A.D. (1986) *Working Memory*. New York: Oxford University Press.

Baddeley, A.D. and Hitch, G.J. (1994) Developments in the concept of working memory. *Neuropsychology*, 8(4), 485–493.

Barch, D.M., Braver, T.S., Nystrom, L., Forman, S.D., Noll, D.C. and Cohen, J.D. (1997) Dissociating working memory from task difficulty in human prefrontal cortex. *Neuropsychologia*, 35, 1373–1380.

Bauer, R.H. and Fuster, J.M. (1976) Delayed-matching and delayed-response deficit from cooling dorsolateral prefrontal cortex in monkeys. *J. Comp. Physiol. Psychol.*, 90(3), 293–302.

Bengio, Y., Frasconi, P. and Simard, P. (1993) *The problem of learning long-term dependencies in recurrent networks*. Paper presented at the Proceedings of the IEEE International Conference on Neural Networks.

Bianchi, L. (1922) *The Mechanism of the Brain and the Function of the Frontal Lobes*. Edinburgh: Livingstone.

Braver, T.S. (1997) *Mechanisms of cognitive control: A neurocomputational model*. Ph.D. Thesis, Carnegie Mellon University.

Braver, T.S. and Cohen, J.D. (in press) On the control of control: The role of dopamine in regulating prefrontal function and working memory. In: S. Monsell and J. Driver (Eds.), *Attention and Performance* (Vol. XVIII.) Cambridge: MA: MIT Press.

Braver, T.S., Cohen, J.D. and McClelland, J.L. (1997) An integrated computational model of dopamine function in reinforcement learning and working memory. *Soc. Neurosci. Abstr.*, 23, 775.

Braver, T.S., Cohen, J.D., Nystrom, L.E., Jonides, J., Smith, E.E. and Noll, D.C. (1997) A parametric study of prefrontal cortex involvement in human working memory. *Neuroimage*.

Braver, T.S., Cohen, J.D. and Servan-Schreiber, D. (1995) A computational model of prefrontal cortex function. In: D.S. Touretzky, G. Tesauro and T.K. Leen (Eds.), *Advances in Neural Information Processing Systems* (Vol. 7, pp. 141–148) Cambridge, MA: MIT Press.

Brozoski, T.J., Brown, R.M., Rosvold, H.E. and Goldman, P.S. (1979) Cognitive deficit caused by regional depletion of dopamine in prefrontal cortex of rhesus monkey. *Science*, 205(31), 929–931.

Calabresi, P., Pisani, A., Centonze, D. and Bernardi, G. (1997) Synaptic plasticity and physiological interactions between dopamine and glutamate in the striatum. *Neurosci. Biobehav. Rev.*, 21, 519–523.

Callaway, E., Halliday, R., Naylor, H. and Yano, L. (1994) Drugs and human information processing. *Neuropsychopharmacology*, 10(1), 9–19.

Carter, C.S., Robertson, L.C., Nordahl, T.E., O'Shora-Celaya, L.J. and Chaderjian, M.C. (1993) Abnormal processing of irrelevant information in schizophrenia: The role of illness subtype. *Psychiatry Res.*, 48, 17–26.

Chapman, L.J. and Chapman, J.P. (1978) The measurement of differential deficit. *J. Psychiat. Res.*, 14, 303–311.

Chapman, L.J., Chapman, J.P. and Miller, G.A. (1964) A theory of verbal behavior in schizophrenia. In: B.A. Maher (Ed.), *Progress in Experimental Personality Research*, Vol. 4. New York: Academic Press, pp. 135–167.

Chiodo, L. and Berger, T. (1986) Interactions between dopamine and amino-acid induced excitation and inhibition in the striatum. *Brain Res.*, 375, 198–203.

Cohen, J.D., Barch, D.M., Carter, C.S. and Servan-Schreiber, D. (1999) Schizophrenic deficits in the processing of context: Converging evidence from three theoretically motivated cognitive tasks. *J. Abnormal Psychol.*, 108, 120–133.

Cohen, J.D., Braver, T.S. and O'Reilly, R. (1996) A computational approach to prefrontal cortex, cognitive control and schizophrenia: recent developments and current challenges. *Philosoph. Trans. Roy. Soc. Lond. Series B*, 351(1346), 1515–1527.

Cohen, J.D., Perlstein, W.M., Braver, T.S., Nystrom, L.E., Noll, D.C., Jonides, J. and Smith, E.E. (1997) Temporal dynamics of brain activation during a working memory task. *Nature*, 386, 604–608.

Cohen, J.D. and Servan-Schreiber, D. (1992) Context, cortex and dopamine: A connectionist approach to behavior and biology in schizophrenia. *Psychol. Rev.*, 99, 45–77.

Cohen, J.D. and Servan-Schreiber, D. (1993) A theory of dopamine function and cognitive deficits in schizophrenia. *Schizophr. Bull.*, 19(1), 85–104.

Cornblatt, B.A. and Keilp, J.G. (1994) Impaired attention, genetics and the pathophysiology of schizophrenia. *Schizophr. Bull.*, 20(1), 31–62.

Courtney, S.M., Ungerleider, L.G., Keil, K. and Haxby, J.V. (1997) Transient and sustained activity in a distributed neural system for human working memory. *Nature*, 386, 608–612.

Creese, I., Burt, D.R. and Snyder, S.H. (1976) Dopamine receptor binding predicts clinical and pharmacological potencies of antischizophrenic drugs. *Science*, 192(April 30), 481–483.

Damasio, A.R. (1985) The frontal lobes. In: K.M. Heilman and E. Valenstein (Eds.), *Clinical Neuropsychology*. New York: Oxford University Press, pp. 339–375.

Davis, K., Kahn, R., Ko, G. and Davidson, M. (1991) Dopamine in schizophrenia: A review and reconceptualization. *Am. J. Psychiat.*, 148(11), 1474–1486.

Dehaene, S. and Changeux, J.P. (1989) A simple model of prefrontal cortex function in delayed-response tasks. *J. Cogn. Neurosci.*, 1, 244–261.

Dehaene, S. and Changeux, J.P. (1992) The Wisconsin card sorting test: Theoretical analysis and modeling in a neuronal network. *Cereb. Cortex*, 1, 62–79.

Engle, R.W., Kane, M.J. and Tuholski, S.W. (1999) Individual differences in working memory capacity and what they tell us about controlled attention, general fluid intelligence and functions of the prefrontal cortex. In: A. Miyake and P. Shah (Eds.), *Models of Working Memory: Mechanisms of Active Maintenance and Executive Control*. New York: Cambridge University Press, pp. 102–134.

Ferron, A., Thierry, A.M., Le Douarin, C. and Glowinski, J. (1984) Inhibitory influence of the mesocortical dopaminergic system on spontaneous activity or excitatory response induced from the thalamic mediodorsal nucleus in the rat medial prefrontal cortex. *Brain Res.*, 302, 257–265.

Frith, C.D. and Done, D.J. (1983) Stereotyped responding by schizophrenic patients on a two-choice guessing task. *Psychological Medicine*, 13, 779–786.

Fuster, J.M. (1973) Unit activity in prefrontal cortex during delayed-response performance: Neuronal correlates of transient memory. *J. Neurophysiol.*, 36, 61–78.

Fuster, J.M. and Alexander, G.E. (1971) Neuron activity related to short-term memory. *Science*, 173, 652–654.

Goldman-Rakic, P.S. (1991) Prefrontal cortical dysfunction in schizophrenia: the relevance of working memory. *Psychopath. Brain*, 1–23.

Goldman-Rakic, P.S., Leranth, C., Williams, S.M., Mons, N. and Geffard, M. (1989) Dopamine synaptic complex with pyramidal neurons in primate cerebral cortex. *Proc. Natl. Acad. Sci. USA*, 86, 9015–9019.

Grace, A.A. (1991) Phasic versus tonic dopamine release and the modulation of dopamine system responsivity: A hypothesis of the etiology of schizophrenia. *Neuroscience*, 41(1), 1–24.

Heffner, T., Hartman, J. and Seiden, L. (1980) Feeding increases dopamine metabolism in the brain. *Science*, 208, 1168–1170.

Heffner, T. and Seiden, L. (1980) Synthesis of catecholamines from h3-tyrosine in the brain during the performance of operant behavior. *Brain Res.*, 183, 403–409.

Hochreiter, S. and Schmidhuber, J. (1997) Long short-term memory. *Neur. Comput.*, 9, 443–453.

Hopfield, J.J. (1982) Neural networks and physical systems with emergent collective computational abilities. *Proceedings of the National Academy of Sciences*, 79, 2554–2558.

Iversen, S. (1984) Cortical monoamines and behavior. In: T. Descarries, R. Reader and H. Jasper (Eds.), *Monoamine Innervation of Cerebral Cortex*. New York: Liss, pp. 321–349.

Kelly, P., Seviour, P. and Iversen, S. (1975) Amphetamine and apomorphine responses following 6-ohda lesions of the nucleus accumbens septi and corpus striatum. *Brain Res.*, 94, 507–522.

Klorman, R., Bauer, L., Coons, H., Lewis, J., Peloquin, L., Perlmutter, R., Ryan, R., Salzman, L. and J., S. (1984) Enhancing effects of methylphenidate on normal young adults cognitive processes. *Psychophamacol. Bull.*, 20, 3–9.

Koob, G.F. and Bloom, F.E. (1988) Celluiar and molecular mechanisms of drug dependence. *Science*, 242, 715–723.

Kornetsky, C. and Orzack, M.H. (1978) Physiological and behavioral correlates of attention dysfunction in schizophrenic patients. *J. Psychiatric Res.*, 14, 69–79.

Kubota, K. and Niki, H. (1971) Prefrontal cortical unit activity and delayed alternation performance in monkeys. *J. Neurophysiol.*, 34, 337–347.

Law-Tho, D., Deuce, J.M. and Crepel, F. (1995) Dopamine favours the emergence of long-term depression versus long-term potentiation in slices of rat prefrontal cortex. *Neurosci. Lett.*, 188, 125–128.

Levine, D.S. and Prueitt, P.S. (1989) Modeling some effects of frontal lobe damage-novelty and perseveration. *Neur. Networks*, 2, 103–116.

Lewis, D.A., Hayes, T.L., Lund, J.S. and Oeth, K.M. (1992) Dopamine and the neural circuitry of primate prefrontal cortex: Implications for schizophrenia research. *Neuropsychopharmacology*, 6(2), 127–134.

Louilot, A., Taghzouti, K., Deminiere, J., Simon, H. and LeMoal, M. (1987) Dopamine and behavior: Functional and theoretical considerations. *Neurotransmitter Interactions in the Basal Ganglia*. New York: Raven Press, pp. 193–204.

Luciana, M., Collins, P.F. and Depue, R.A. (1995) *DA and 5-HT influences on spatial working memory functions of prefrontal cortex*. Paper presented at the Cognitive Neuroscience Society Second Annual Meeting, San Francisco, CA.

Luciana, M., Depue, R.A., Arbisi, P. and Leon, A. (1992) Facilitation of working memory in humans by a D2 dopamine receptor agonist. *J. Cogn. Neurosci.*, 4(1), 58–68.

Luria, A.R. (1969) Frontal lobe syndromes. In P.J. Vinken and G.W. Bruyn (Eds.), *Handbook of Clinical Neurology*, Vol. 2. New York: Elsevier, pp. 725–757.

Malmo, H.P. (1974) On frontal lobe functions: Psychiatric patient controls. *Cortex*, 10, 231–237.

Manschreck, T., Maher, B.A., Milavetz, J.J., Ames, D., Weisstein, C.C. and Schneyer, M.L. (1988) Semantic priming in thought disordered schizophrenic patients. *Schizophr. Res.*, 61–66.

Miller, E.K., Erickson, C.A. and Desimone, R. (1996) Neural mechanisms of visual working memory in prefrontal cortex of the macaque. *J. Neurosci.*, 16(0), 5154–5167.

Milner, B. (1963) Effects of different brain lesions on card sorting. *Arch. Neurol.*, 9, 90–100.

Mirenowicz, J. and Schultz, W. (1996) Preferential activation of midbrain dopamine neurons by appetitive rather than aversive stimuli. *Nature*, 379, 449–451.

Montague, P.R., Dayan, P. and Sejnowski, T.J. (1996) A framework for mesencephalic dopamine systems based on predictive Hebbian learning. *J. Neurosci.*, 16, 1936–1947.

Moody, S.L., Wise, S.P., di Pellegrino, G. and Zipser, D. (1998) A model that accounts for activity in primate frontal cortex during a delayed-match-to-sample task. *J. Neurosci.*, 18, 399–410.

Moore, R.Y. and Bloom, F.E. (1978) Central catecholamine neuron systems: Anatomy and physiology of the dopamine systems. *Ann. Rev. Neurosci.*, 1, 129–169.

Mora, F. and Cobo, M. (1990) The neurobiological bases of prefrontal cortex self-stimulation: A review and integrative hypothesis, *Progress in Brain Research* (Vol. 85, pp. 419–431) Berlin: Elsevier.

Mozer, M.C. (1993) Neural net architectures for temporal sequence processing. In: A. Weigend and N. Gershenfeld (Eds.), *Predicting the future and understanding the past*. (pp. 243–264). Redwood City, CA: Sante Fe Institute Studies in the Sciences of Complexity, Proceedings Volume XVII, Addison-Wesley Publishing.

Norman, D.A. and Shallice, T. (1986) Attention to action: Willed and automatic control of behavior. In: R.J. Davidson, G.E. Schwartz and D. Shapiro (Eds.), *Consciousness and Self-regulation*. (Vol. 4, pp. 1–18). Plenum Press, New York.

Norton, S. (1973) Amphetamine as a model for hyperactivity in the rat. *Physiol. Behav.*, 11.

Nuechterlein, K.H. (1991) Vigilance in schizophrenia and related disorders. In: S.R. Steinhauer, J.H. Gruzelier and J. Zubin (Eds.), *Handbook of Schizophrenia Vol. 5: Neuropsychology, Psychophysiology and Information Processing*. Amsterdam: Elsevier, pp. 397–433.

Nuechterlein, K.H. and Dawson, M.E. (1984) Information processing and attentional functioning in the developmental course of schizophrenia disorders. *Schizophr. Bull.*, 10(2), 160–203.

Oades, R.D. (1985) The role of noradrenaline in tuning and dopamine in switching between signals in the central nervous system. *Neurosci. Biobehav. Rev.*, 9, 261–282.

Owen, A.M., Roberts, A.C., Polkey, C.E., Sahakian, B.J. and Robbins, T.W. (1991) Extra-dimensional versus intra-dimensional set shifting performance following frontal lobe excisions, temporal lobe excisions or amygdalo-hippocampectomy in man. *Neuropsychologia*, 29, 993–1006.

Park, S. and Holzman, P.S. (1992) Schizophrenics show spatial working memory deficits. *Arch. Gen. Psychiatry*, 49, 975–982.

Penit-Soria, J., Audinat, E. and Crepel, F. (1987) Excitation of rat prefrontal cortical neurons by dopamine: An in vitro electrophysiological study. *Brain Res.*, 425, 263–274.

Phillips, A. and Fibiger, H. (1989) Neuroanatomical bases of intracranial self-stimulation: Untangling the gordian knot. *The Neuropharmacological Basis of Reward* (Vol. 66, pp. 1–5) Oxford: Oxford University Press.

Randrup, A. and Munkvad, I. (1970) Biochemical, anatomical and psychological investigations of stereotyped behavior induced by amphetamine. In: E. Costa and S. Garattini (Eds.), *Amphetamine and Related Compounds*. New York: Raven Press, pp. 695–713.

Rosvold, H.E., Mirsky, A.F., Sarason, I., Bransome, E.D. and Beck, L.H. (1956) A continuous performance test of brain damage. *Journal of Consulting Psychology*, 20(5), 343–350.

Sawaguchi, T. and Goldman-Rakic, P. (1994) The role of d1-dopamine receptors in working memory: Local injections of dopamine antagonists into the prefrontal cortex of rhesus monkeys performing an oculomotor delayed-response task. *J. Neurophysiol.*, 63(6), 1401–1412.

Sawaguchi, T. and Goldman-Rakic, P.S. (1991) D1 dopamine receptors in prefrontal cortex: Involvement in working memory. *Science*, 251, 947–950.

Sawaguchi, T., Matsumara, M. and Kubota, K. (1990) Catecholaminergic effects on neuronal activity related to a delayed response task in monkey prefrontal cortex. *J. Neurophysiol.*, 63, 1385–1400.

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

Schultz, W., Apicella, P. and Ljungberg, T. (1993) Responses of monkey dopamine neurons to reward and conditioned stimuli during successive steps of learning a delayed response task. *J. Neurosci.*, 13(3), 900–913.

Segal, D. (1975) Behavioural characterisation of d- and l-amphetamine: Neurochemical implications. *Science*, 190, 475–477.

Servan-Schreiber, D., Bruno, R.M., Carter, C.S. and Cohen, J.D. (1998) Dopamine and the mechanisms of cognition. Part I: A neural network model predicting dopamine effects on selective attention. *Biol. Psychiatry.*, 43, 713–722.

Servan-Schreiber, D., Cohen, J.D. and Steingard, S. (1996) Schizophrenic deficits in the processing of context: A test of a theoretical model. *Arch. Gen. Psychiatry*, 53, 1105–1113.

Servan-Schreiber, D., Printz, H. and Cohen, J.D. (1990) A network model of catecholamine effects: Gain, signal-to-noise ratio and behavior. *Science*, 249, 892–895.

Sesack, S.R., Snyder, C.L. and Lewis, D.A. (1995) Axon terminals immunolabeled for dopamine or tyrosine hydroxylase synapse on GABA-immunoreactive dendrites in rat and monkey cortex. *J. Comp. Neurol.*, 363, 1735–1780.

Shallice, T. (1982) Specific impairments of planning. *Phil. Trans.R. Soc. Lond.*, 298, 199–209.

Shallice, T. (1988) *From Neuropsychology to Mental Structure*. Cambridge: Cambridge University Press.

Snyder, S. (1972) Catecholamines in the brain as mediators of amphetamine psychosis. *Arch. Gen. Psychiatry*, 27, 169–179.

Storms, L.H. and Broen, W.E. (1969) A theory of schizophrenic behavioral disorganization. *Arch. Gen. Psychiatry*, 20(Feb), 129–144.

Stuss, D.T. and Benson, D.F. (1986) *The Frontal Lobes*. New York: Raven Press.

Sutton, R. (1988) Learning to predict by the method of temporal difference. *Machine Learning*, 3, 9–44.

Sutton, R.S. and Barto, A.G. (1990) Time-derivative models of Pavlovian reinforcement. In: M. Gabriel and J. Moore (Eds.), *Learning and Computational Neuroscience: Foundations of Adaptive Networks*. Cambridge, MA: MIT Press, pp. 497–537.

Wapner, S. and Krus, D.M. (1960) Effects of lysergic acid diethylamide and differences between normals and schizophrenics, on the Stroop color-word test. *J. Neuropsychiatry*, 2(Nov/Dec), 76–81.

Weinberger, D.R., Berman, K.F. and Zec, R.F. (1986) Physiological dysfunction of dorsolateral prefrontal cortex in schizophrenia: I. Regional cerebral blood flow evidence. *Arch. Gen. Psychiatry*, 43, 114–125.

Wickens, J.R., Begg, A.G. and Arbuthnott, G.W. (1996) Dopamine reverses the depression of rat corticostriatal synapses which normally follows high frequency stimulation of cortex in vitro. *Neuroscience*, 70, 1–5.

Williams, M.S. and Goldman-Rakic, P.S. (1993) Characterization of the dopaminergic innervation of the primate frontal cortex using a dopamine-specific antibody. *Cereb. Cortex*, 3(May-June), 199–222.

Willner, P. and Scheel-Kruger, J. (1991) *The mesolimbic dopamine system: From Motivation to Action*. New York: Wiley.

Wise, R.A. and Rompre, P.-P. (1989) Brain dopamine and reward. *Ann. Rev. Psychol.*, 40, 191–225.

Wynne, L.C., Cromwell, R.L. and Matthysse, S. (1978) *The Nature of Schizophrenia: New Approaches to Research and Treatment*. New York: John Wiley and Sons, Inc.

Wysocki, J.J. and Sweet, J.J. (1985) Identification of brain damaged, schizophrenic and normal medical patients using a brief neuropsychological screening battery. *An International Journal of Clinical Neuropscyhology*, 7(1), 40–44.

Zipser, D. (1991) Recurrent network model of the neural mechanism of short-term active memory. *Neur. Comput.*, 3, 179–193.

Zipser, D., Kehue, B., Littlewort, G. and Fuster, J. (1993) A spiking network model of short-term active memory. *J. Neurosci.*, 13, 3406–3420.

Zubin, J. (1975) Problem of attention in schizophrenia. In: M.I. Kietzman, S. Sutton and J. Zubin (Eds.), *Experimental Approaches to Psychopathology*. New York: Academic Press, pp. 139–166.