27

The Role of Prefrontal Cortex in Normal and Disordered Cognitive Control: A Cognitive Neuroscience Perspective

TODD S. BRAVER, JONATHAN D. COHEN, AND DEANNA M. BARCH

The advent of cognitive neuroscience as a discipline has accelerated research on the functions of the prefrontal cortex (PFC). The cognitive neuroscience perspective has led to greater consideration and integration of multiple different methodologies and research approaches, including basic neuroscience (neurophysiology and neuroanatomy), clinical neuroscience (neuropsychology and psychiatry), and cognitive science (experimental psychology and computer science). Moreover, researchers now have the opportunity to study the human brain "in action" through the use of functional neuroimaging methods. However, the greater focus on convergence between these different approaches and traditions has also revealed some of the conflicts between them, in terms of their traditional views of PFC function.

The first conflict concerns whether the PFC should be considered a storage buffer or an executive controller. In the neurophysiology literature, a commonly held view of PFC function is that of a short-term storage mechanism, actively holding information on-line through neural activity (Goldman-Rakic, 1995). These findings have been supported by human neuroimaging data, suggesting sustained PFC activation specifically associated with maintenance periods in tasks requiring short-term or working memory (Fiez et al., 1996; Cohen et al., 1997; Courtney, et al., 1997). In contrast, the neuropsychological literature has tended to focus on the role of the PFC in behavioral regulation and cognitive control and on the impairments in these functions following damage to the PFC (Hecaen & Albert, 1978; Damasio, 1985; Stuss & Benson, 1986). Although these differing views of PFC function are not by themselves incompatible, a prominent cognitive psychological model put forward by Baddeley and colleagues has suggested that storage and control processes should be considered architecturally distinct and strictly segregated components of a working memory system (Baddeley & Hitch, 1974; Baddeley, 1986).

A second conflict among the various literatures concerns whether the PFC functions more as a mnemonic, inhibitory, or attentional structure. Neurophysiological and neuroimaging studies have tended to focus on the role of PFC in short-term or working memory (e.g., Fuster, 1973; Funahashi et al., 1993; D'Esposito et al., 1998; Smith & Jonides, 1999). However, there is clear evidence that the PFC is critically involved in selective attention as well (Cabeza & Nyberg, 2000). The PFC is seen as playing a central role in the "anterior attentional system," in which the se-

lection, direction, and allocation of attentional resources is endogenously controlled (Posner & Petersen, 1990). In contrast, the developmental and clinical literatures have viewed the PFC as primarily geared towards inhibitory functions, such as reducing sensory interference and overriding dominant but inappropriate response tendencies. In this literature, much has been made of findings that inhibitory control throughout the lifespan seems to parallel the time course of PFC maturation and decline (Diamond, 1990; Dempster, 1992), and that damage to the PFC often leads to a behavioral "disinhibition syndrome" in which the normal control over social and sexual behavior is released (Hecaen & Albert, 1978). In previous reviews attempting to synthesize these different literatures, some theorists have suggested that memory and inhibition should also be thought of as distinct and anatomically segregated functions of PFC (Fuster, 1989). This hypothesis has been based on neuroanatomical and neurophysiological data suggesting functional segregation of dorsolateral and ventromedial regions of PFC, with dorsolateral PFC being associated with mnemonic functions and ventromedial PFC with inhibitory function. Less focus has been placed on the anatomical locus of attention within PFC, yet the implicit view seems to be that attentional functions are one component of an anatomically based modular organizational structure within the PFC.

These functional distinctions between storage, control, and memory, inhibition, and attention present a confusing and often inconsistent picture for the theorist attempting to develop a coherent theory of PFC function. Thus, an important open question is whether and how these distinctions can be reconciled. In this chapter, we shall provide such an attempt at reconciliation, by focusing on the potential computational mechanisms associated with storage, control, inhibition, and attention. This computational approach suggests a resolution of the apparent conflicts between the different perspectives on PFC function. We suggest that a common set of computational mechanisms allow for PFC mediation of mnemonic, inhibitory, and attentional functions, where each is preferentially observable under different task situations, and each occurs in the service of cognitive control. Specifically, we argue that the control functions of the PFC emerge as a direct consequence of two specific mechanisms: active maintenance of task-relevant context and top-down biasing of local competitive interactions that occur during processing.

We have developed our theory of PFC function using the connectionist computational modeling framework. This modeling approach involves three components: *(1)* computational analysis of the critical processing mechanisms required for cognitive control; *(2)* use of neurobiologically plausible principles of information processing; and *(3)* implementation and simulation of cognitive tasks and behavioral performance. The modeling work is complemented by a series of convergent empirical studies relying on multiple experimental methodologies. First we will describe behavioral and neuroimaging data on healthy young adults that provide validation for critical components of the model. We will then summarize our work applying the model to the clinical domain; we have tested its predictions in different population groups (older adults, schizophrenia patients) thought to suffer from PFC dysfunction. These studies highlight the power of the cognitive neuroscience approach by demonstrating how a single, integrated account of PFC function can capture a wide range of data from different methodologies and multiple populations.

A THEORY OF PREFRONTAL CORTEX FUNCTION IN COGNITIVE CONTROL

CONTEXT AND COGNITIVE CONTROL

A basic and fundamental function of cognitive control is to flexibly adapt behavior to the demands of particular tasks by facilitating processing of task-relevant information over other sources of competing information and by inhibiting habitual or otherwise prepotent responses that are inappropriate to the task. Because this control function is such a fundamental one, it occurs in even very simple task situations. As a specific example, take a

situation in which a speeded response is required to a particular stimulus, but only in a particular context (e.g., respond to the letter *X* only if immediately following the letter *A*). If the context–stimulus pairing occurs frequently, the cognitive system should begin to exploit the context to prime or facilitate processing of the stimulus. In contrast, in the rare situations in which the stimulus occurs in a different context (e.g., *X* following the letter *B*), the system must rely on the information provided by the context to inhibit the tendency to respond. This example raises the question of what types of processing mechanisms could perform such a cognitive control function.

We would argue that there are at least three minimal components required of this type of cognitive control mechanism: *(1)* a representational code that conveys implications of the behavioral goal or prior context for future behavior; *(2)* a capability for actively maintaining this representation in an accessible form; and *(3)* a means of conveying an activation signal that can directly influence ongoing processing by directing attention, biasing action selection, or resolving perceptual ambiguities. In reference to the example task above, the contextual cue should be translated into a code that appropriately distinguishes between different upcoming inputs (*X* versus *not-X*) and/or their behavioral consequences (respond versus don't respond). This representation should be accessible even under conditions in which the contextual cue is temporally isolated from the relevant stimulus or response, and thus no longer externally available. Finally, the contextual cue should directly modulate processing in perceptual and/or motor pathways, such as by priming the expected response or by focusing attention towards the expected stimulus features.

Over the past 10 years we have been developing a theory which suggests that these minimal elements of cognitive control are subserved by specific neural mechanisms that are dependent on the function of the PFC (Cohen & Servan-Schreiber, 1992; Braver et al., 1995; Cohen et al., 1996; O'Reilly et al., 1999; Miller & Cohen, 2001). Moreover, we have described these neural mechanisms in terms of explicit computational principles and have implemented these principles within connectionist neural network models. Our theory is composed of three central hypotheses: *(1)* the PFC is specialized for the representation and maintenance of context information; *(2)* context information is maintained in the PFC as a stable and self-sustaining pattern of neural activity; and *(3)* context representations in the PFC mediate control through interactions that modulate the flow of information in other brain systems that more directly support task performance.

A critical aspect of our hypothesis regarding the role of the PFC in cognitive control relates to the notion of context representation. We define *context* as any task-relevant information that is internally represented in such a form that it can bias processing in the pathways responsible for task performance. Goal representations are one form of such information, which have their influence on planning and overt behavior. However, we use the more general term context to include representations that may have their effect earlier in the processing stream, on interpretive or attentional processes. For example, in the Stroop Task, the context provided by the task instructions must be actively represented and maintained to bias attentional allocation and response selection towards the ink color dimension of a visually presented word. Thus, context representations may include a specific prior stimulus or the result of processing a sequence of stimuli, as well as task instructions or a particular intended action. Representations of context are particularly important for situations in which there is strong competition for response selection. These situations may arise when the appropriate response is one that is relatively infrequent or when the inappropriate response is dominant and must be inhibited (such as the word *name* in the Stroop Task). Importantly, context representations can be maintained on-line, in an active state, such that they are continually accessible and available to influence processing. Thus, context can be thought of as a component of working memory. Specifically, context can be viewed as the subset of representations within working memory that govern how other rep-

resentations are used. In this manner, context representations simultaneously subserve both storage and control functions. As described above, this aspect of the model differentiates it from Baddeley's model of working memory (Baddeley, 1986; 1993), which postulates a strict separation of representations for storage versus control.

A COMPUTATIONAL MODEL OF PREFRONTAL CORTEX FUNCTION AND CONTEXT PROCESSING

The theory of PFC involvement in context processing and cognitive control described above was developed within the connectionist, or parallel distributed processing, framework (Rumelhart & McClelland, 1986; McClelland, 1993). The connectionist framework is a natural one for concomitantly studying the neural and psychological mechanisms of cognitive control, since it provides a computational architecture that is specified in neurobiological terms and can be used to quantitatively simulate performance in cognitive tasks. In this framework, information is represented as graded patterns of activity over populations of simple units, processing takes place as the flow of activity goes from one set of units to another, and learning occurs through the modification of the connection strengths between these. From one perspective, connectionist models are highly simplified, capturing brain-style computation, without necessarily committing to the details of any particular neural system or subsystem. However, with appropriate refinement, such models offer the opportunity to build bridges between our understanding of the low-level properties of neural systems and their participation in higher-level (system) behavior.

The theory of cognitive control put forward here can be schematized in the form of a simple canonical model in which a context module serves as an indirect pathway that modulates processing in a direct stimulus–response pathway (see Fig. 27–1). This context-processing module represents the functions of the PFC. There are two critical features of this module that provide it with the capacity for control over processing. The first is that there is strong, recurrent connectivity within the context layer, which allows for the active maintenance of information. Thus, input to the context layer can be sustained through activity recirculation along mutually excitatory connections, even when the external source of input is no longer present. The second critical feature of the context pathway is its feedback connection to the direct pathway. This provides a means for activity within the context module to provide an additional source of input, which can modulate the flow of processing within the direct pathway. In particular, feedback from the context layer serves to bias the local competition for representation that exists within each module, favoring one activation pathway or set of representations over their competitors. This biasing action of the context module can produce inhibitory effects

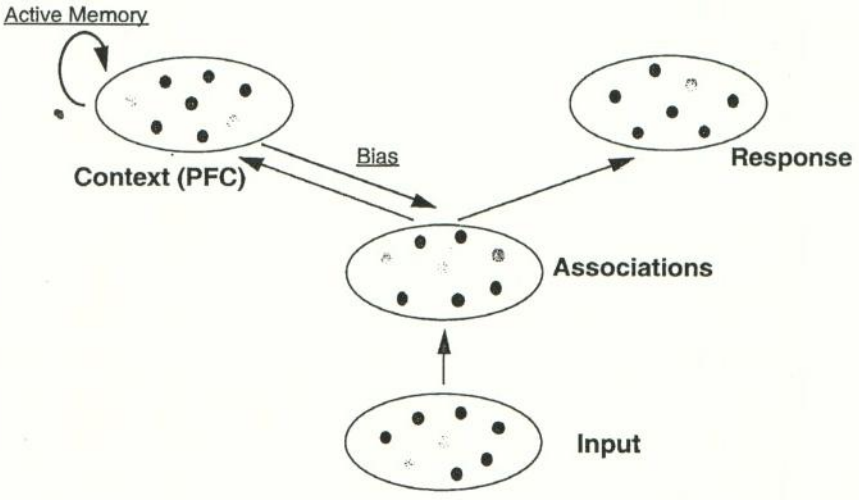

Figure 27–1. Diagram of canonical model. Key computational principles of context processing mechanism are shown: active memory through recurrent connections, and top-down bias through feedback connections. PFC, prefrontal cortex.

on processing by allowing a weak pathway to inhibit the more dominant one.

An important insight that has emerged from our work with this model is that it demonstrates how context-processing mechanisms might jointly support three cognitive functions—working memory, inhibition, and attention—that have all been suggested to be subserved by PFC, but that are often treated as independent. When a task involves a delay between a cue and a later contingent response, it is usually assumed that a working memory function is involved. However, there is no dedicated mechanism for working memory in the model. Rather, the mechanism that is used to initially represent context information can also maintain this information against the interfering and cumulative effects of noise over time. In contrast, when a task involves competing, task-irrelevant processes (as in the Stroop Task), it is often assumed that a dedicated inhibitory function is responsible for suppressing or overriding these irrelevant processes. Once again, in the model, there is no dedicated mechanism for inhibition. Rather, context representations indirectly provide an inhibitory effect by providing top-down support for task-relevant processes, allowing these to compete effectively against irrelevant ones. It is important to note that this competition is expected to occur locally, in posterior regions, rather than in the PFC itself. Finally, attention is thought to be required in tasks during which an internal or external cue signals that a specific stimulus feature or dimension has increased salience relative to others. In the model, contextual information is translated into a representational code and fed back into the system, facilitating information processing in specific stimulus–response pathways via selective support of those pathways. Yet the context mechanism is not a dedicated "attentional module," but rather just an extra source of top-down input that can be sustained over time.

Thus, the same PFC mechanism can be involved in tasks that alternatively tap working memory, inhibition, or attention; it is simply a matter of the behavioral conditions under which it operates (i.e., the source of interference) and the information it selects (task-relevant versus task-irrelevant) that lead us to label it as having a working memory, inhibitory, or attentional function. Consequently, the model suggests a clear resolution of the supposedly disparate findings of PFC involvement in working memory, attentional, and inhibitory functions by suggesting how a single mechanism might commonly subserve all three domains. Below, we discuss the computational and empirical studies we have conducted to test the model.

THE AX-CPT PARADIGM

Our investigations have focused on testing whether this model of cognitive control provides a useful account of both normal and disordered PFC function. To examine this question, we have conducted a series of studies using multiple methodologies and populations, but employing a single experimental paradigm for probing cognitive control function. Our research strategy has been to systematically examine and understand the properties and consequences of PFC activity within a single paradigm, before testing the model further in additional paradigms. The paradigm we have studied, known as the AX-CPT, was selected on the basis of a number of favorable properties. First, the AX-CPT, is derived from the well-known Continuous Performance Test (CPT; Rosvold et al., 1956), which in the clinical and neuropsychological literatures is widely used as a test of attentional control and vigilance (e.g., Nuechterlein, 1991). The more demanding versions of the CPT have been shown to strongly rely on PFC function, as evidenced by performance deficits observed in patients with frontal lesions (Glosser & Goodglass, 1990) and other syndromes thought to involve prefrontal dysfunction, such as schizophrenia (Cornblatt & Keilp, 1994) and attention-deficit disorder (ADHD) (Losier et al., 1996). Many neuroimaging studies have also used these difficult versions of the CPT to elicit PFC activity (Cohen et al., 1987; Rezai et al., 1993; Siegel et al., 1995; Seidman et al., 1998). Second, the AX-CPT appears to be sensitive to individual differences in PFC function and/or cognitive control. In previous studies, we have shown that

AX-CPT performance is correlated with performance on other widely used probes of PFC and cognitive control function, such as the N-back, Stroop, and reading span tasks (Cohen et al., 1999; Keys et al., submitted). Third, the AX-CPT probes key aspects of cognitive control, while distilling them into a task paradigm that is as simple and interpretable as possible. Because of the simplicity of the task, it can be used with many different subject populations and under a wide variety of task environments. Indeed, the task is very similar in structure to delayed-response tasks used in the neurophysiological literature on working memory (e.g., Fuster, 1973), and thus allows easy comparison with this literature. Moreover, because the task is relatively simple, it can be simulated in computational studies. Lastly, although the task is simple, it nevertheless produces multiple performance measures, which generate a rich set of data on which to base and constrain theoretical interpretations.

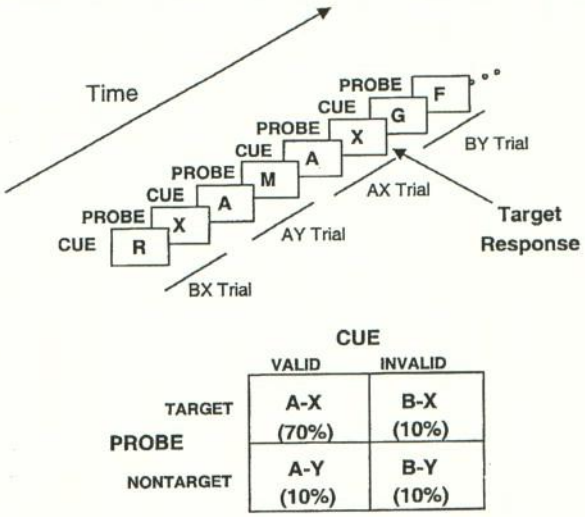

		CUE: VALID	CUE: INVALID
PROBE	TARGET	A-X (70%)	B-X (10%)
	NONTARGET	A-Y (10%)	B-Y (10%)

Figure 27–2. Schematic of AX-CPT paradigm. Single letters are visually displayed as a series of cue–probe pairs. A target is defined as the occurrence of an *X* probe immediately following an *A* cue. Three types of nontarget trials occur, each with equal frequency (10%): BX, AY, and BY (where *B* refers to any non-*A* cue, and *Y* refers to any non-*X* probe).

In the AX-CPT, subjects view sequences of letters presented one at a time as a series of cue–probe pairs in the center of a visual display. The object of the task is to make a target response to a specific probe letter (*X*), but only when it follows a valid cue (*A*), and a nontarget response in all other cases (see Fig. 27–2). Performance is dependent on the representation and maintenance of context information, insofar as the correct response to the probe depends on knowledge of the previous cue (*A* or not-*A*). In the model, the context provided by the cue is represented and maintained within the PFC. We further designed the task to selectively measure different aspects of the context processing functions subserved by PFC (working memory, attention, and inhibition).

In the task, target (*AX*) trials occur with high frequency (70%). This induces two types of biases in subjects. The first is a bias to make a target response to the occurrence of an *X* probe. On those trials in which a target response should *not* be made to the *X* probe (i.e., *BX* trials, where *B* refers to any non-*A* cue), context information must be used in an *inhibitory* fashion to override the tendency to false alarm. The second bias that occurs in the AX-CPT is an expectancy to make a target response following the occurrence of an *A* cue. In this case, the context provided by the cue serves a predictive function that directs *attention* to a particular response (i.e., attention-to-action; Norman & Shallice, 1986; Allport, 1989). On those trials in which the cue is an invalid predictor of the response (i.e., *AY* trials, where *Y* refers to any non-*X* probe), the attentional function of context creates the tendency to false alarm. This type of cue validity effect is similar to others that have been extensively studied in the attentional literature (e.g., Posner, 1980). Thus, the integrity of context processing can be examined not only through performance on AX target trials but also through an examination of performance on nontarget trials.

A key element of our theory is that both attentional and inhibitory functions in the AX-CPT should be subserved by a single underlying mechanism—the internal representation of context information within the PFC. This assumption can be tested by examining the relationship of AY to BX performance. Note that on BX trials, failure to internally represent context should impair performance, by not

suppressing the inappropriate response bias. Consequently, BX false alarms can be considered *context-failure errors*. However, on AY trials, representation of context should create an inappropriate expectancy bias that leads to an increased tendency to false alarm. AY false alarms can thus be considered *context-induced errors*. Thus, if context representations are intact, more AY (context-induced) than BX (context-failure) errors should be made (with a similar pattern observable in reaction time on non-error trials). Conversely, if context representations are impaired, the opposite pattern should occur (more BX than AY errors). Performance on AX target trials should also be poorer if context processing is impaired, since determination of targets is dependent on the context provided by the cue. However, target performance may not be as impaired as BX performance, since on target trials, the response bias works in subjects' favor, by increasing the tendency to make the correct target response. Finally, a third type of nontarget trial, *BY,* provides a useful internal control, since in this condition the influence of context on performance should be relatively small (given that both the cue and the probe always map to a nontarget response).

The AX-CPT paradigm also provides a means for examining the *mnemonic* role of context information through manipulations of the cue–probe delay duration. Specifically, under conditions in which there is a long cue–probe delay (e.g., 5–10 seconds), context information must be actively maintained within working memory. Our theory suggests that context information is both represented and actively maintained within the PFC. Thus, the same context-processing mechanism that subserves inhibitory and attentional functions should also subserve active maintenance in the AX-CPT. Consequently, a strong prediction of the theory is that the effect of delay will interact with performance on AY and BX trials. If context maintenance is intact, then the strength of context representations should either hold constant or increase with delay (i.e., if it takes some period of time for context representations to reach full activation strength). Consequently, BX performance should remain constant or improve at long delays, while AY performance should remain constant or worsen with delay. Conversely, if context maintenance is impaired, then context representations should lose strength over time. This should lead to a worsening of BX performance with a delay, but an improvement in AY performance.

As the above description of the task makes clear, our model of PFC function in cognitive control leads to a number of specific predictions for the AX-CPT that relate to normal behavioral performance, brain activation, and clinical/neuropsychological populations. A central focus of our recent work has been to systematically test these predictions through a series of studies (Braver et al., 1995, in press; Barch et al., 1997, 2001; Servan-Schreiber et al., Braver, 1997; Braver & Cohen, 1999, 2000, in press; Cohen et al., 1999). The first phase of the research was to provide initial validation of the model by examining how it can capture aspects of normal AX-CPT performance, in terms of both brain activity and behavior. The second phase of the research was to apply the model as a predictive tool for studying the consequences of impairments in PFC function and/or context processing. In this work, we have attempted to show how the model can be used to make rather specific and sometimes highly counterintuitive predictions regarding AX-CPT performance in different populations (healthy individuals performing under load, schizophrenia patients, and older adults). The remainder of the chapter discusses these two phases of our research.

MODEL VALIDATION: NORMAL COGNITIVE CONTROL

Over the past 10 years we have tested numerous healthy young adults in the AX-CPT task. From this testing it has emerged that there is a set of fairly standard characteristics of behavioral performance on the task (Braver et al., 1990b). Figure 27–3 shows error data on nontarget trials averaged from over 200 subjects performing the AX-CPT under both short (1–2 seconds) and long delay (5–10 seconds) conditions. A number of notable features can be observed. First, there are very

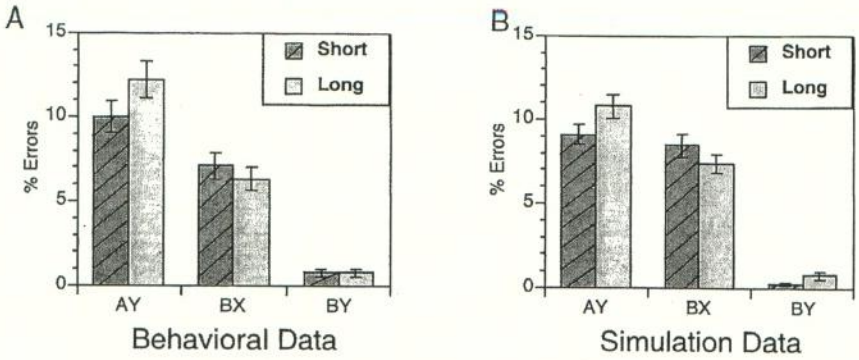

Figure 27–3. Simulation and behavioral performance data from the AX-CPT. *A:* Behavioral data obtained from a normative sample of over 200 healthy young adults performing the AX-CPT under standard conditions. Performance on each of the three nontarget trial types is shown for short- (dark hatched bars) and long-delay (light bars) conditions. *B:* Simulation data from computational model showing performance on each the three nontarget trial types (AY, BX, and BY) under short- (dark hatched bars) and long-delay (light bars) conditions. (*Source:* Data from Braver et al, 1999b)

few errors ever made on BY trials, which serve as an internal control for baseline performance. This suggests that, overall, the task is not difficult for healthy adults to perform. Relative to the BY error rate, there are significantly more errors on BX and AY trials. This is consistent with the conflict that is present on BX and AY trials between the contextual information and the current input (Carter et al., 1998). Thus, in these conditions greater cognitive control is required. Moreover, healthy adults produce more AY (context-induced) errors than BX (context-failure) errors. As described above, this pattern indicates that context exerts a strong influence over responding—subjects are more likely to overrely on contextual information than they are to fail to utilize it. Finally, it is clear that the error pattern also interacts with delay. Specifically, at long delays AY errors are increased, BX errors are decreased, and BY errors are unchanged. This pattern indicates that context information is accurately maintained in the system over the delay period, and even appears to exert a stronger influence over performance at the longer delay interval (i.e., at the long delay there is a greater likelihood of context inducing an error and a reduced likelihood of context failing to prevent an error).

In computer simulations with our model, we found that we were able to closely capture this pattern of behavioral performance, both in qualitative and quantitative terms (see Fig. 27–3). In particular, the model exhibits all of the primary effects observed in the empirical data, including effects related to reaction time (RT) and AX target performance (not shown; see Braver et al., 1999b). However, it is important to note that a number of free parameters were adjusted in the model to optimize its fit to the data. Thus, the simulations do not provide a strong test of the theory's sufficiency or explanatory power. Nevertheless, they do serve as an initial validation of the model, suggesting that it is capable of capturing the major behavioral phenomena associated with normal AX-CPT performance. Moreover, the results increase our confidence in using the model as a starting point for generating hypotheses and predictions regarding the effects of other variables.

In a second set of studies, we directly validated the role of PFC in context processing and cognitive control postulated in our model. Specifically, brain activity was measured during AX-CPT performance through the use of functional magnetic resonance imaging (fMRI) methods. Our model suggested that in the AX-CPT task the contextual information provided by the cue should be represented in PFC and actively maintained there over an intervening delay. To test this hypothesis, we compared brain activity in the AX-CPT under short- and long-delay conditions, holding all other aspects of the task constant (such as the total trial duration). Consequently, the only

thing that should have differed across conditions was the proportion of the trial period devoted to active maintenance of context information.

Our results indicated that activity within left PFC, in a dorsolateral (DL-PFC) region corresponding to Brodmann's area (BA) 46/9, was significantly increased in the long-delay condition relative to the short one (see Fig. 27–4; Barch et al., 1997). This finding has now been replicated in three subsequent studies using independent cohorts of subjects (Barch et al., 2001b; Braver & Cohen, 2001; Braver & Bongiolatti, in press). Recently, we extended our initial findings by conducting a study using newly developed event-related fMRI methods (Braver & Cohen, 2001). Event-related techniques enable measurement of the temporal dynamics and evolution of brain activity over the course of a trial (Buckner et al., 1996; Buckner & Braver, 1999). By plotting the activation time course in this left DL-PFC region across the trial, we observed a response pattern that was highly consistent with that predicted by the model. Specifically, we found that left DL-PFC activity increased immediately following cue presentation and remained high over the duration of the delay period (10 seconds), then declined following probe presentation (Fig. 27–4). Critically, we found no other brain regions outside of DL-PFC showing this characteristic activity pattern.

The neuroimaging studies of the AX-CPT provide important validation for our claims regarding the role of PFC in context processing and cognitive control. Moreover, they point to the specific region within the PFC that is activated by context processing demands during AX-CPT performance. This is critical, because the model in its current scope is rather anatomically nonspecific, in that it does not differentiate between specific PFC subregions. In one aspect, this lack of anatomical specificity is intentional, because we believe that a unifying principle of PFC function is its involvement in aspects of context processing. However, human PFC takes up nearly a third of the cerebral cortex, and there are clear anatomical subdivisions within it (ventrolateral, dorsolateral, frontopolar, medial, orbital) (Barbas & Pandya, 1989b; Fuster, 1997a). Given the structure–function relationships that exist in the brain, it is likely that these anatomic subdivisions within the PFC reflect relevant functional specialization. An important question that we shall return to at the end of the chapter is whether the model can be extended

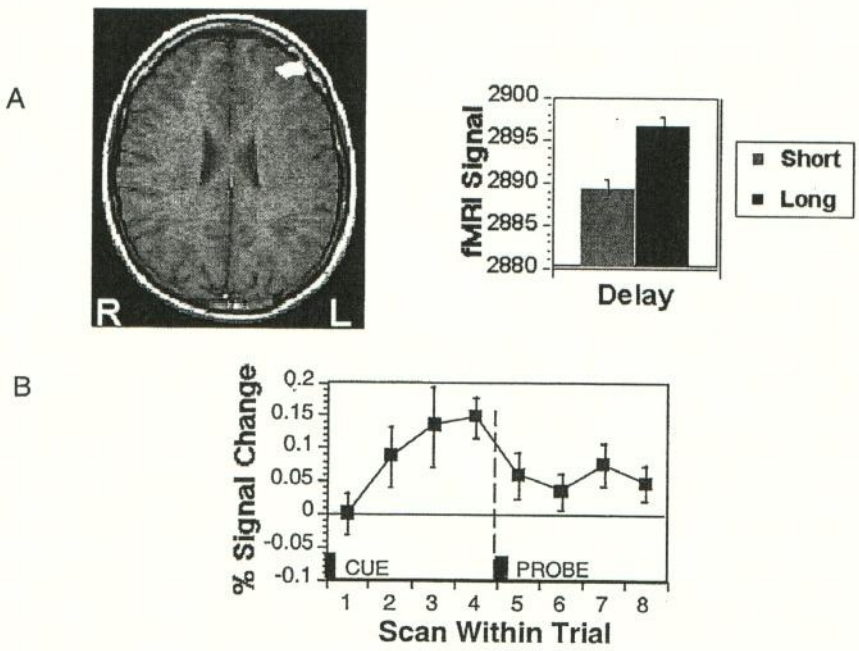

Figure 27–4. *A:* Activity in left dorsolateral (DL prefrontal cortex PFC) for short- (light bars) and long-delay (dark bars) conditions. Brain activation data are shown for a representative axial slice 24 mm superior to the AC–PC plane (centroid: -37, 42, 29). *B:* Time course of event-related activity in left DL-PFC during a long-delay trial (10 seconds). Note that there is a sustained increase in activity during the cue–probe delay period which the decays during the intertrial interval. (*Source:* From Barch et al., 1997, and Braver & Cohen, 2001)

in scope to help address and explain these functional specializations within PFC. For the current purposes, the neuroimaging studies suggest that our model of PFC context-processing functions in the AX-CPT should be considered to pertain to DL-PFC regions most directly. Moreover, the consistent observation of left DL-PFC activity during the AX-CPT in healthy young adults suggests that studies examining AX-CPT context-processing impairments in different populations and/or task conditions should focus on this region as the neural locus of dysfunction. In the next section, we will focus on a series of studies designed to test this hypothesis.

APPLYING THE MODEL: DISORDERED COGNITIVE CONTROL

Our primary research strategy has been to apply our model of cognitive control as a predictive tool for generating hypotheses regarding the relationship between PFC dysfunction and cognitive control impairments related to context processing. In particular, we adopted a cognitive neuroscience approach that sought to provide convergent evidence for the model using multiple methods. Our first step was to determine what biologically relevant parameters in the model affect context processing functions. By manipulating these parameters in computational simulations, we determined their effects on performance of the AX-CPT task, both in terms of behavior and PFC activity. The simulation results were then used to generate empirically testable predictions for populations or experimental conditions that were hypothesized to correspond to the same change in the biologically relevant parameter. Our results are from three different populations: healthy subjects performing under interference, schizophrenia patients, and older adults.

SIMULATION STUDIES

To examine the relationship between PFC dysfunction, context processing impairment, and AX-CPT performance in the model, we simulated changing the function of the dopamine (DA) neurotransmitter projection into the PFC. A detailed description of the motivation and theory behind these simulations is beyond the scope of this chapter, and is described elsewhere (Braver & Cohen, 1999; 2000). Briefly, we hypothesized that DA provides a modulatory input into the PFC, which serves to regulate the access of incoming information. In particular, we have suggested that the connection serves as a gating mechanism. When the gate is opened, as is thought to occur following a phasic burst of DA activity, incoming information can gain access to the PFC, and thus update the current state of context representation. Conversely, when the gate is closed, access to the PFC is restricted, protecting context representations from the interfering effects of noise or other irrelevant inputs. We hypothesized that the timing of gating signals is learned through a reward prediction learning mechanism associated with the midbrain DA system (Schultz et al., 1997), which enables task-relevant information to be selected as context, because of its association with the potential for future reinforcement.

In our simulations, we found that when DA effects are reduced in the context module, the representation of context becomes less reliable (since access is partially blocked). Under conditions when DA input is noisier (i.e., more variable over time) both the representation and maintenance of context is disrupted (see Fig. 27–5; Braver, et al., 1999). The maintenance effects occur because context information is more susceptible to the interfering effects of noise and task-irrelevant inputs, and thus is more likely to decay over time. These conditions of dysfunctional context processing were also associated with clear changes in model performance on the AX-CPT (Fig. 27–5). First, we found that BX performance worsened. Recall that BX trials require inhibition of a response bias, based on context information. Thus, if the context representation is weak, then so is the ability to inhibit the response bias. In contrast, we found that AY performance was actually improved in the disturbed model. Remember that on AY trials,

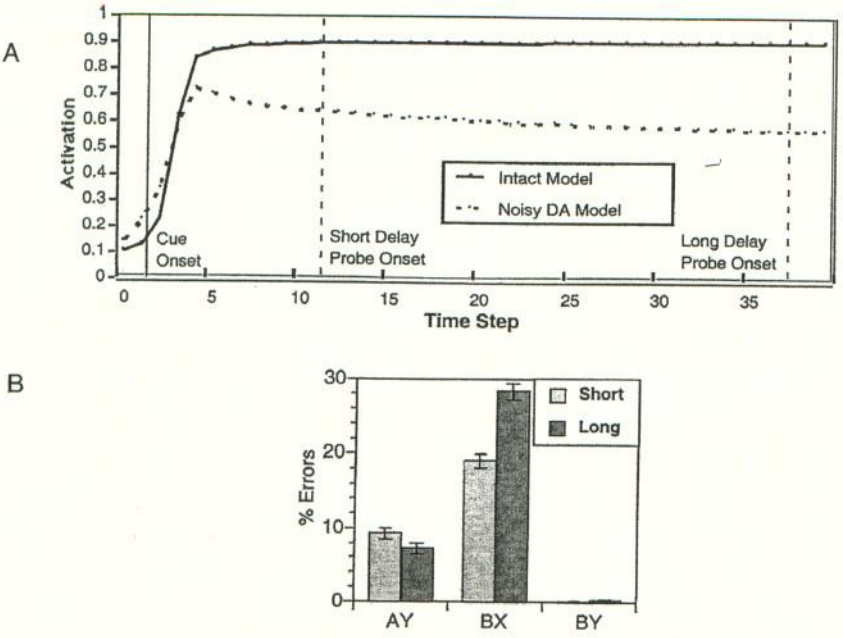

Figure 27–5. *A:* Plot of average activity level in context units during the course of an AX-CPT trial in the intact and noisy DA models. Note that context activity is attenuated in the noisy dopamine (DA) model and further declines with delay. *B:* Simulation data from computational model with noisy DA. (*Source:* Data from Braver et al., 1999b)

context sets up an expectancy that is violated by the nontarget probe letter. When context representations are weaker, so is the expectancy. The lower the expectancy, the less likely it is that a false alarm will be made.

It is important to note that this improvement of AY performance under impaired context-processing conditions represents a highly counterintuitive prediction of the model. It is typically much harder to obtain improvements in performance due to cognitive dysfunction than it is to obtain impairments. Moreover, the predicted effects on AY and BX performance suggest a crossover interaction when comparing normal and dysfunctional context-processing conditions. Predicted crossover interactions are optimal from an experimental point of view because *(1)* they are statistically powerful, and thus easier to detect (Wahlsten, 1991), and *(2)* because they obviate nonspecific interpretations of their cause, such as differential difficulty across conditions (e.g., Chapman & Chapman, 1978).

INTERFERENCE CONDITIONS

Our first test of these model predictions regarding context-processing disruption was in healthy young subjects performing the AX-CPT under interference conditions (Braver et al., 1999b). Specifically, we hypothesized that presentation of irrelevant distractor letters during the cue–probe delay interval would disrupt context processing by increasing the probability of an inappropriate DA-mediated gating signal (causing an improper update of context). Indeed, we observed that healthy subjects under interference displayed a performance pattern similar to that predicted by the model under context disruptions. At the short delay, interference effects were minimal, with the typical effect of more AY than BX errors. However, at the long delay, the effects of interference appeared to accumulate. Specifically, BX errors increased but AY errors decreased, leaving significantly more BX than AY errors (see Fig. 27–6). We further examined the effect of interference in a neuroimaging study to determine whether performance in the interference condition was associated with changes in the activation of left DL-PFC (Braver & Cohen, 2001). We again used event-related fMRI to track DL-PFC activity dynamics over the course of a trial. Our results indicated a significant effect of interference in the left DL-PFC region that we had previously found to show delay-related activity during the task (Fig. 27–6). Under interference there is an activation response to the cue that initially increases, but then shows a rapid decay over the delay period. This finding provides fairly strong support for the idea that the

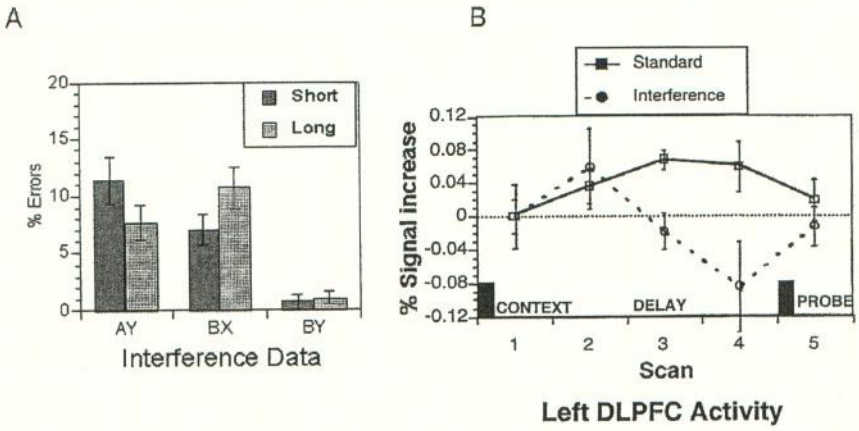

Figure 27–6. *A:* Behavioral data from young adults performing the AX-CPT under interference conditions. *B:* Time course of event-related activity in the left DLPFC under both standard and interference AX-CPT conditions. Activity declines more rapidly under the interference condition. (*Source:* Data from Braver et al., 1999, and Braver & Cohen, 2001).

behavioral impairments related to context processing observed in AX-CPT performance under interference are linked to a change in the activity dynamics of DL-PFC.

SCHIZOPHRENIA PATIENTS

Our second set of studies examined AX-CPT performance in schizophrenia patients. Schizophrenia is a disease associated with marked cognitive disturbance. A common view is that schizophrenia patients suffer from a failure of cognitive control (Callaway & Naghdi, 1982; Nuechterlein & Dawson, 1984). There is a wealth of neuropsychological, neuroanatomical, and neuroimaging evidence suggesting that PFC dysfunction is a central component of the disease, and may be the underlying cause of the cognitive deficits (Goldman-Rakic, 1991). Additionally, the role of DA in schizophrenia is well known, and most pharmacological treatments for clinical symptoms involve agents that affect the DA system (Creese et al., 1976). More recently, investigators have suggested that some of the cognitive deficits present in schizophrenia may result from altered DA activity in PFC (Davis et al., 1991).

We have examined the performance of schizophrenia patients on the AX-CPT across a number of studies (Servan-Schreiber et al., 1996; Cohen et al., 1999; Barch et al., 2001b; Braver et al., 1999b). A consistent pattern across all of these studies is that, relative to matched controls, patients show a performance pattern that is indicative of a selective deficit in context processing and that appears to be magnified in the long-delay condition. More recently, we have studied patients suffering their first psychotic episode and thus were not yet medicated and had never been hospitalized. This produced a very clean sample for examining schizophrenia cognitive impairments, because many of the confounding clinical variables that typically affect cognitive performance were not present (e.g., length of illness, chronic medication status, and institutionalization). We again found that this group made significantly more BX errors relative to a tightly matched control sample, but in fact made fewer AY errors, and that this effect was most pronounced at the long delay (see Fig. 27–7; Barch et al., 2001b). These findings provide support for the hypothesis that a selective context-processing deficit underlies some of the cognitive impairments observed in schizophrenia. We simulated this data in our model and found that when the DA gating mechanism was disrupted, the model produced performance patterns that were qualitatively very similar to that observed empirically (Braver et al., 1999).

We also acquired measurements of fMRI activity in this cohort of first-episode patients and matched controls during AX-CPT performance (Barch et al., 2001b). The healthy controls showed normal delay-related activation of left DL-PFC. However, the patients failed to show a delay-related increase, yielding a significant group x delay interaction in left DL-PFC (Fig. 27–7). It is noteworthy that this was the only brain region found to show such a pattern. Moreover, it was found that many other regions showed normal task-related ac-

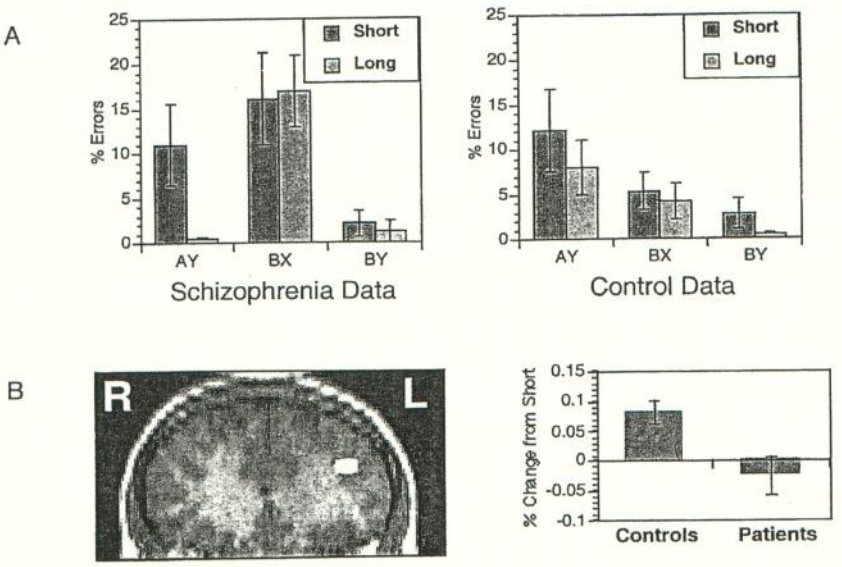

Figure 27–7. *A:* Behavioral data from first-episode schizophrenia patients and matched controls. Patients show more BX errors, but fewer AY errors; this effect is more pronounced at long delay. *B:* Left dorsolateral prefrontal cortex (centroid: -34, 25, 26) shows delay-related activation increase in controls but not in patients. (*Source:* Data from Barch et al., 2001b)

tivation in patients, including more posterior and inferior regions of PFC (e.g., Broca's area).

OLDER ADULTS

A final population that we have studied is healthy older adults. Like schizophrenia, healthy aging is associated with impairments in both PFC and DA functions, albeit at a much less severe level. Older adults perform similarly to patients with frontal lobe damage on neuropsychological batteries (West, 1996), and in some cases neuropsychological test performance has been found to correlate to the degree of age-related reductions in PFC gray matter (Raz et al., 1998). The functioning of the DA system also declines in older adults, with DA receptor concentration regularly decreasing with increasing age (de Keyser et al., 1990). Moreover, recent work has linked age-related DA receptor decreases with cognitive decline in tests sensitive to PFC function (Volkow et al., 1998). We tested whether these aging effects on DA and PFC function would be reflected in terms of selectively decreased context-processing function.

In a large sample of older adults performing the AX-CPT (in the long-delay condition only), under both baseline and interference conditions, we found the same pattern predicted by the model to be indicative of selective context-processing impairment—opposite patterns of performance change on BX and AY trials (see Fig. 27–8; Braver et al., 2001). On BX trials, older adults made significantly more errors and showed substantial slowing of RT. However, on AY trials, older adults showed significantly fewer errors and equivalent RT. The AY RT pattern is especially striking, given that an almost ubiquitous finding in the cognitive aging literature is that older adults are slower than younger adults on any cognitive task (Cerella, 1985; Myerson & Hale, 1993). Moreover, the longer the response latency for younger adults, the greater the amount of slowing typically observed. In contrast, our results indicate that there was no age-related slowing on AY trials, even though these were the conditions that produced the longest RTs in younger adults. Without the benefit of a model, this would seem to be an anomalous and potentially difficult result to interpret. Yet it is fully consistent with our hypotheses that context-processing deficits should produce a paradoxical improvement in AY performance.

In our most recently completed work, we have followed up on the behavioral results with a fMRI study conducted in an independent sample of younger and older adults (Barch et al., 2001a). Subjects performed the long- and short-delay AX-CPT conditions in our standard design. Once again, in confirmation of our predictions, we observed a significant delay-related reduction of PFC activ-

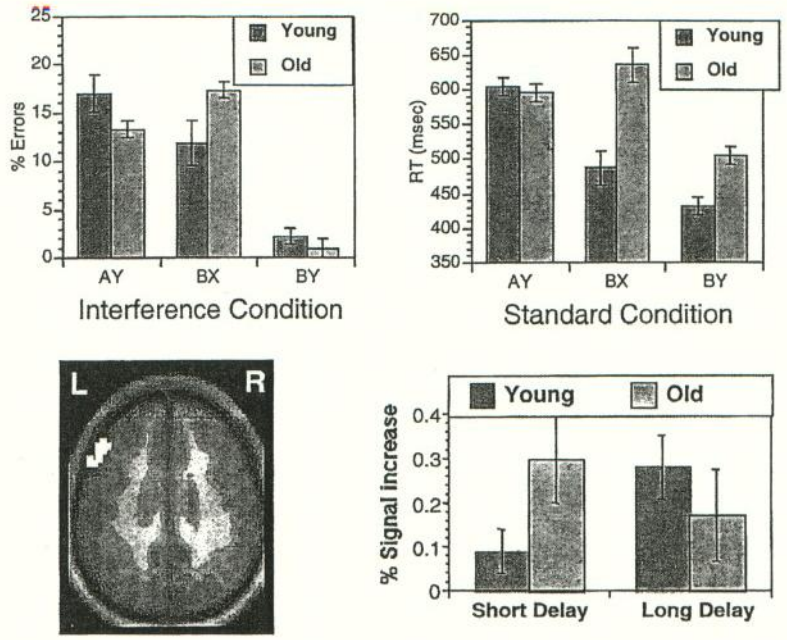

Figure 27–8. *A:* Behavioral data from healthy older (light bars) and younger (dark bars) adults, including both errors (left) and reaction time (RT). Older adults show more BX errors, but less AY errors. The RT data show significant age-related slowing on BX trials, but no slowing on AY trials. *B:* Left dorsolateral prefrontal cortex (centroid: -43, 22, 28) shows delay-related activation increase in young but not older adults. (*Source:* Data from Barch et al., 2001a, and Braver et al., 2001)

ity in older adults, located within the same left dorsolateral region identified in the previous studies (Fig. 27–8). Interestingly, this reduced DL-PFC activity contrasted with the general pattern observed across other brain regions, in which older adults actually showed greater task-related activity than younger adults. This finding is consistent with other neuroimaging studies of cognition in aging (Cabeza, 2001), and could indicate that older adults recruit additional brain regions during tasks requiring cognitive control as a means of compensating for impaired context processing functions in the PFC.

SUMMARY AND DISCUSSION

Across a series of studies, we have attempted to draw a tight link between context-processing function, AX-CPT performance, and PFC activity. In four independent samples of healthy young adults performing the AX-CPT under standard conditions, the same region of left DL-PFC (BA 46/9) was found to show increased activity when context information needed to be actively maintained over a delay period (Barch et al., 1997, 2001b; Braver & Cohen, 2001; Braver & Bongiolatti, in press). This highly replicable pattern of results suggests that the AX-CPT provides a selective and focal probe of left DL-PFC function. The finding that this left DL-PFC region showed reduced activity in two different population (older adults and schizophrenia patients) thought to suffer from PFC dysfunction further supports this suggestion. Moreover, in both populations, the PFC activity pattern co-occurred with a distinctive and similar pattern of behavioral performance. This behavioral pattern included both performance impairments in some task conditions (BX trials), but also relative improvements in others (AY trials). The relationship between AX-CPT performance and left DL-PFC activity suggests that the AX-CPT paradigm has potential clinical utility as a selective neuropsychological marker. An obvious prediction of the model that is worth testing in future studies is that patients with focal left DL-PFC lesions would show a performance profile similar to that of schizophrenia patients and older adults. Another possible use of the AX-CPT worth exploring in future studies is as a diagnostic tool for suspected left DL-PFC dysfunction (i.e., when this is not immediately clear from other neuroanatomical markers). However, for this to occur, a clear amount of additional work is needed to determine whether the AX-CPT shows appropriate psychometric properties (i.e., sensitivity, specificity, and reliability) in both its behavioral and

neural activity measurements to warrant such usage.

It is important to recall that not only were we able to detect differences between-groups in PFC activity associated with context-processing dysfunction, but we also observed within-subjects PFC activity differences related to experimentally induced disruptions of context processing. In particular, the findings from the interference AX-CPT studies in healthy subjects indicate that reductions in left DL-PFC activity during task performance do not necessarily indicate the presence of a neurobiological impairment. Conversely, the results also strengthen the hypothesis that the reduced left DL-PFC activity observed among older adults and schizophrenia patients was not a causally irrelevant population difference, such as being a byproduct of a more global pathophysiology. Instead, we suggest that in the AX-CPT, delay-related decreases in left DL-PFC signal the failure to appropriately represent and maintain the contextual information provided by task cues.

Although the clinical implications of our work are certainly important, our primary goal in this chapter has been to highlight the benefits of a multimethod cognitive neuroscience approach and, most importantly, the role that computational modeling can play in advancing theories of PFC function. The critical point of our studies is not that the AX-CPT activates the left DL-PFC, or that impaired activity can be observed in different populations and/or conditions, but to what these findings imply about the underlying cognitive mechanisms supporting task performance. We began with two related hypotheses: *(1)* a central function of PFC is to represent and actively maintain context or goal-relevant information; and *(2)* these context representations serve to control processing by biasing the flow of information through posterior brain pathways. These hypotheses are similar in many ways to those expressed by other investigators (Goldman-Rakic, 1987; Fuster, 1997b). In particular, the recent work of Miller and colleagues is highly consistent with our ideas. Miller suggests that representations in the PFC are not directly tied to sensory features but instead code more abstract properties such as rules or a task set (Miller, 2000; Miller & Cohen, 2001). However, a distinguishing feature of our approach is that we have tried to make the hypotheses explicit, by implementing them within computer simulation models (c.f., Kimberg & Farah, 1993; Dehaene & Changeux, 1995). These simulation models are admittedly highly simplified and somewhat abstract in relation to many of the important neurobiological details that characterize real neural systems. Nevertheless, by developing and implementing a running simulation model, it is possible to examine how a postulated set of mechanisms interact with each other in terms of their impact on information-processing dynamics and their implications for behavioral performance within specific task contexts.

The critical insight we developed from this simulation work is that actively maintained context representations, when fed back into the rest of the system, can support multiple functions: short-term information storage (working memory), suppression of inappropriate responses (inhibition), and selection and enhancement of task-relevant input (attention). Moreover, although these functions are all subserved by the same mechanism, they are most apparent in different circumstances—storage under delay conditions, suppression under conditions involving dominant response biases, and attentional facilitation under conditions requiring robust, efficient processing (e.g., speeded conditions) or conditions with perceptually weak (or ambiguous) inputs. From this insight we designed the AX-CPT paradigm to probe each of these aspects of PFC function through different measures of performance (i.e., delay effect = working memory, BX = inhibition, AY = attention). We hypothesized that if all three functions were subserved by a common set of PFC mechanisms, then performance on the different measures should interact. Because the AX-CPT paradigm was conceptually so simple, we were able to simulate performance of the task in a quantitative manner and under multiple different conditions. The simulation results supported the hypothesis of interaction, showing that when context-processing mechanisms were disturbed (by simulating a change in DA modulation of PFC inputs), performance

changed on both BX and AY trials, and the effects were modulated by delay. In addition, the behavioral effects were associated with a pattern of altered activity dynamics, such that activity was reduced in the simulated PFC units, and tended to decay over time. Thus, the primary benefit of the model was that it provided us with a means of generating explicit and highly specific predictions about the consequences of context-processing disturbances for both brain activity and behavioral performance.

The confirmation and replication of the model's predictions in different populations and with different methods suggest that it is capturing important principles related to PFC involvement in cognitive control during AX-CPT performance. It is our goal, however, to extend the model beyond the AX-CPT, to determine whether it can help explain PFC involvement in cognitive control more generally. Specifically, we have tried to show how a very simple processing mechanism of context representation and maintenance can subserve both storage and control functions, both memory and inhibition, as well as certain aspects of attentional selection. Thus, the context-processing model may provide a resolution of the various theories of PFC function that have either argued for a preferential role of PFC in one of the functions, or have suggested that the different functions are anatomically segregated and independent. Our model argues that neither of these alternatives needs to be true. A single region within PFC may be involved with multiple cognitive control functions and be critical for multiple cognitive domains, because these functions are not necessarily independent or computationally incompatible.

LIMITATIONS AND FUTURE DIRECTIONS

An important direction for future research is to determine the appropriate scope of the context-processing model. Obviously, the work presented to date is only a small first step. A clear direction for future research will be to determine if the model can be extended to examine context effects beyond the AX-CPT paradigm, in different tasks, and across different cognitive domains. Early stages of this effort have already begun (Cohen & Servan-Schreiber, 1992; Cohen et al., 1992; O'Reilly et al., in press). Another direction of future research will be to enrich our model of PFC functions. Our current model is highly simplified and, in particular, does not differentiate among PFC subregions. However, regional specializations clearly appear to be present in the PFC, in terms of both anatomy and function (Goldman-Rakic, 1987; Barbas & Pandya, 1989a; Fuster, 1997a). Thus a critical question is whether the model should only apply to DL-PFC regions (and even more specifically to left hemisphere DL-PFC), or whether it can be applied more generally.

The suggestion that the context-processing model applies only to DL-PFC regions is a plausible one and, in fact, is the only interpretation strictly supported by the results. Moreover, the dorsolateral region of PFC is most reliably implicated in cognitive control functions, in both human (Cabeza & Nyberg, 2000; Duncan & Owen, 2000; Smith & Jonides, 1999) and primate studies (Fuster, 1989; Goldman-Rakic, 1987; Smith & Jonides, 1999); We would like to speculate, however, that participation in context representation and maintenance functions may be a unifying dimension that cuts across PFC subregions. Even if this hypothesis is true, the dimensions and organizational structure that define the regional specialization of PFC remain to be discovered. The question of PFC organizational structure is being actively pursued by a number of investigators (Goldman-Rakic, 1996; Petrides, 1996; D'Esposito et al., 1998; Owen et al., 1998; Christoff & Gabrieli, 2000). One view commonly adopted is that the PFC is organized according to the information content of representations, with distinctions being drawn according to psychologically relevant categories, such as identity versus location, verbal versus nonverbal, and sensory modality (Wilson et al., 1993; Smith & Jonides, 1999; Levy & Goldman-Rakic, 2000). A second view suggests that the organizational structure is primarily based on processing rather than content specializations, such as cognitive versus emotional/motivational functions (Dias et al., 1996; Bechara et al., 1998; Rolls, 2000), and

maintenance versus manipulation of information in working memory (D'Esposito et al., 1999; Postle & D'Esposito, 2000). Each of these views has met with mixed success at accounting for the full range of empirical data.

Our own view is that representational content and processing functions are likely to be tightly intertwined within the brain. Moreover, given the complexity of the computational functions performed in PFC, it is unlikely that the representational distinctions between PFC subregions will map onto easily labeled content categories. For example, the relevant dimensions of specialization within the PFC may include the temporal duration over which representations are typically activated and maintained or the number of stimulus dimensions that are simultaneously integrated. We have begun exploring models that make use of such computational specializations (O'Reilly et al., in press). Yet, regardless of which dimensions of specialization eventually prove to be the most relevant for PFC function, we feel confident that computational modeling analyses have the potential to strongly contribute to our understanding of this brain region.

REFERENCES

Allport, A. (1989). Visual attention. In: M.I. Posner (Ed.), *Foundations of Cognitive Science* (pp. 631–682). Cambridge, MA: MIT Press.

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

Baddeley, A.D. (1993). Working memory or working attention? In: A.D. Baddeley & L. Weiskrantz (Eds.), *Attention: Selection, Awareness, and Control: A Tribute to Donald Broadbent* (pp. 152–170). Oxford: Clarendon Press.

Baddeley, A.D. & Hitch, G.J. (1974). Working memory. In: G. Bower (Ed.), *The Psychology of Learning and Motivation* Vol. VIII (pp. 47–89). New York: Academic Press.

Barbas, H. & Pandya, D.N. (1989a). Architecture and intrinsic connections of the prefrontal cortex in the rhesus monkey. *Journal of Comparative Neurology, 286*, 353–375.

Barbas, H. & Pandya, D.N. (1989b). Architecture of intrinsic connections of the prefrontal cortex in the Rhesus monkey. *Journal of Comparative Neurology, 286*, 353–375.

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

Barch, D.M., Braver, T.S., Racine, C.A., & Satpute, A.B. (2001a). Cognitive control deficits in healthy aging: neuroimaging investigations. *NeuroImage, 13*, S1025.

Barch, D.M., Carter, C.S., Braver, T.S., McDonald, A., Sabb, F.W., Noll, D.C., & Cohen, J.D. (2001b). Selective deficits in prefrontal cortex regions in medication naive schizophrenia patients. *Archives of General Psychiatry, 50*, 280–288.

Bechara, A., Damasio, H., Tranel, D., & Anderson, S.W. (1998). Dissociation of working memory from decision making within the human prefrontal cortex. *Journal of Neuroscience, 18*, 428–437.

Braver, T.S. (1997). Mechanisms of Cognitive Control: A Neurocomputational Model. Ph.D. Thesis, Carnegie Mellon University, Pittsburgh, PA.

Braver, T.S., & Bongiolatti, S.R. (in press). The role of frontopolor prefrontal cortex in subgoal processing during working memory. *NeuroImage*.

Braver, T.S. & Cohen, J.D. (1999). Dopamine, cognitive control, and schizophrenia: the gating model. Progress in Brain Research, 121, 327–349.

Braver, T.S. & Cohen, J.D. (2000). On the control of control: The role of dopamine in regulating prefrontal function and working memory. In: S. Monsell & J. Driver (Eds.), *Attention and Performance XVIII* (pp. 713–738). Cambridge, MA: MIT Press.

Braver, T.S. & Cohen, J.D. (2001). Working memory, cognitive control, and the prefrontal cortex: Computational and empirical studies. *Cognitive Processing 2*, 25–55.

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

Braver, T.S., Barch, D.M., & Cohen, J.D. (1999). Cognition and control in schizophrenia: a computational model of dopamine and prefrontal function. *Biological Psychiatry, 46*, 312–328.

Braver, T.S., Barch, D.M., Keys, B.A., Carter, C.S., Kaye, J.A., Janowksy, J.S., Taylor, S.F., Yesavage, J.A., Mumenthaler, M.S., Jagust, W.J., & Reed, B.R. (2001). Context processing in older adults: evidence for a theory relating cognitive control to neurobiology in healthy aging. *Journal of Experimental Psychology: General, 130*, 746–763.

Braver, T.S., Barch, D.M., & Cohen, J.D. (1999b). Mechanisms of cognitive control: active memory, inhibition, and the prefrontal cortex. (Technical Report PDP.CNS.99.1). Pittsburgh: Carnegie Mellon University.

Buckner, R.L., & Braver, T.S. (1999). Event-related functional MRI. In: P. Bandettini & C. Moonen (Eds.), *Functional MRI* (pp. 441–452). Berlin: Springer-Verlag.

Buckner, R.L., Bandettini, P.A., O'Craven, K.M., Savoy, R.L., Petersen, S.E., Raichles, M.E., & Rosen, B.R. (1996). Detection of cortical activation during averaged single trials of a cognitive task using functional magnetic

resonance imaging. *Proceedings of the National Academy of Science USA, 93,* 14878–14883.

Cabeza, R. (2001). Functional neuroimaging of cognitive aging. In: R. Cabeza & A. Kingstone (Eds.), *Handbook of Functional Neuroimaging of Cognition* (pp. 331–337). Cambridge, MA: MIT Press.

Cabeza, R. & Nyberg, L. (2000). Imaging cognition II: an empirical review of 275 PET and fMRI studies. *Journal of Cognitive Neuroscience, 12,* 1–47.

Callaway, E. & Naghdi, S. (1982). An information processing model for schizophrenia. *Archives of General Psychiatry, 39,* 339–347.

Carter, C.S., Braver, T.S., Barch, D.M., Botvinick, M.M., Noll, D.C., & Cohen, J.D. (1998). Anterior cingulate cortex, error detection, and the online monitoring of performance. *Science, 280,* 747–749.

Cerella, J. (1985). Information processing rates in the elderly. *Psychological Bulletin, 98,* 67–83.

Chapman, L.J. & Chapman, J.P. (1978). The measurement of differential deficit. *Journal of Psychiatry, 14,* 303–311.

Christoff, K. & Gabrieli, J.D.E. (2000). The frontopolar cortex and human cognition: evidence for a rostrocaudal hierarchical organization within the human prefrontal cortex. *Psychobiology, 28,* 168–186.

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

Cohen, J.D., Servan-Schreiber, D., & McClelland, J.L. (1992). A parallel distributed processing approach to automaticity. *American Journal of Psychology, 105,* 239–269.

Cohen, J.D., Braver, T.S., & O'Reilly, R. (1996). A computational approach to prefrontal cortex, cognitive control, and schizophrenia: recent developments and current challenges. *Philosophical Transactions of the Royal Society of London Series B, 351*(1346), 1515–1527.

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

Cohen, J.D., Barch, D.M., Carter, C., & Servan-Schreiber, D. (1999). Context-processing deficits in schizophrenia: converging evidence from three theoretically motivated cognitive tasks. *Journal of Abnormal Psychology, 108,* 120–133.

Cohen, R.M., Semple, W.E., Gross, M., Nordahl, T.E., Delisi, L.E., Holcomb, H.H., King, A.C., Morihisa, J.M., & Pickar, D. (1987). Dysfunction in a prefrontal substrate of sustained attention in schizophrenia. *Life Sciences, 40,* 2031–2039.

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

Courtney, S.M., Ungerleider, L.G., Keil, K., & 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., & Snyder, S.H. (1976). Dopamine receptor binding predicts clinical and pharmacological potencies of antischizophrenic drugs. *Science, 192,* 481–483.

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

Davis, K.L., Kahn, R.S., Ko, G., & Davidson, M. (1991). Dopamine in schizophrenia: a review and reconceptualization. *American Journal of Psychiatry, 148,* 1474–1486.

Dehaene, S. & Changeux, J.-P. (1995). Neuronal models of prefrontal cortical function. In: J. Grafman, K. Holyoak, & F. Boller (Eds.), *Structure and Functions of the Human Prefrontal Cortex*, Vol. 769 (pp. 305–321). New York: New York Academy of Sciences.

de Keyser, J., De Backer, J.-P., Vauquelin, G., & Ebinger, G. (1990). The effect of aging on the D1 dopamine receptors in human frontal cortex. *Brain Research, 528,* 308–310.

Dempster, F.N. (1992). The rise and fall of the inhibitory mechanism: towards a unified theory of cognitive development and aging. In: G.J. Whitehurst (Ed.), *Developmental Review: Perspectives in Behavior and Cognition*, Vol. 12 (pp. 45–75). San Diego: Academic Press.

D'Esposito, M., Aguirre, G.K., Zarahn, E., Ballard, D., Shin, R.K., & Lease, J. (1998). Functional MRI studies of spatial and nonspatial working memory. *Cognitive Brain Research, 7,* 1–13.

D'Esposito, M., Postle, B.R., Ballard, D., & Lease, J. (1999). Maintenance versus manipulation of information held in working memory: an event-related fMRI study. *Brain and Cognition, 41,* 66–86.

Diamond, A. (1990). The development and neural bases of memory functions as indexed by the A-not-B and delayed response tasks in human infants and infant monkeys. In: A. Diamond (Ed.), *The Development and Neural Bases of Higher Cognitive Functions* (pp. 267–317). New York: New York Academy of Science Press.

Dias, R., Robbins, T.W., & Roberts, A.C. (1996). Dissociation in prefrontal cortex of affective and attentional shifts. *Nature, 380,* 69–72.

Duncan, J. & Owen, A.M. (2000). A similar frontal lobe network recruited by diverse cognitive demands. *Trends in Neurosciences, 23,* 475–483.

Fiez, J.A., Raife, E.A., Balota, D.A., Schwarz, J.P., Raichle, M.E., & Petersen, S.E. (1996). A positron emission tomography study of the short-term maintenance of verbal information. *Journal of Neuroscience, 16,* 808–822.

Funahashi, S. Bruce, C.J., & Goldman-Rakic, P.S. (1993). Dorsolateral prefrontal lesions and oculomotor delayed-response performance: evidence for mnemonic "scotomas". *Journal of Neuroscience, 13,* 1479–1497.

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

Fuster, J.M. (1989). *The Prefrontal Cortex: Anatomy, Physiology and Neuropsychology of the Frontal Lobe.* New York: Raven Press.

Fuster, J.M. (1997a). Anatomy of the prefrontal cortex. In: M. Placito, M. Bailer, & K. Bubbeo (Eds.), *The Prefrontal Cortex: Anatomy, Physiology, and Neuropsychology of the Frontal Lobe* (3rd ed. (pp. 6–42). Philadelphia: Lippincott-Raven.

Fuster, J.M. (1997b). Overview of prefrontal functions: The temporal organization of behavior. In: M. Placito, M. Bailer, & K. Bubbeo (Eds.), *The Prefrontal Cortex: Anatomy, Physiology, and Neuropsychology of the Frontal Lobe*, (3rd ed. (pp. 209–252). Philadelphia: Lippincott-Raven.

Glosser, G. & Goodglass, H. (1990). Disorders in executive control functions among aphasic and other brain-damaged patients. *Journal of Clinical and Experimental Neuropsychology, 12,* 485–501.

Goldman-Rakic, P.S. (1987). Circuitry of primate prefrontal cortex and regulation of behavior by representational memory. In: F. Plum & V. Mountcastle (Eds.), *Handbook of Physiology, The Nervous System V,* Vol. 5 (pp. 373–417). Bethesda, MD: American Physiological Society.

Goldman-Rakic, P.S. (1991). Prefrontal cortical dysfunction in schizophrenia: the relevance of working memory. In: B.J. Carroll & J.E. Barrett (Eds.), *Psychopathology and the Brain* (pp. 1–23). New York: Raven Press.

Goldman-Rakic, P.S. (1995). Cellular basis of working memory. *Neuron, 14,* 477–485.

Goldman-Rakic, P.S. (1996). The prefrontal landscape: implications of functional architecture for understanding human mentation and the central executive. In: A.C. Roberts, T.W. Robbins, & L. Weiskrantz (Eds.), *The Prefrontal Cortex: Executive and Cognitive Functions* (pp. 87–103). Oxford: Oxford University Press.

Hecaen, H. & Albert, M.L. (1978). *Human Neuropsychology*. New York: John Wiley & Sons.

Keys, B.A., Barch, D.M., Braver, T.S., & Janowsky, J.S. (submitted). Task sensitivity to age differences in working memory: relative superiority of the *N*-back paradigm.

Kimberg, D.Y. & Farah, M.J. (1993). A unified account of cognitive impairments following frontal lobe damage: the role of working memory in complex, organized behavior. *Journal of Experimental Pscyhology: General, 122,* 411–428.

Levy, R. & Goldman-Rakic, P.S. (2000). Segregation of working memory functions within the dorsolateral prefrontal cortex. *Experimental Brain Research, 233,* 23–32.

Losier, B.J., McGrath, P.J., & Klein, R.M. (1996). Error patterns on the continuous performance test in non-medicated and medicated samples of children with and without ADHD: a meta-analytic review. *Journal of Child Psychology & Psychiatry, 37,* 971–987.

McClelland, J.L. (1993). Toward a theory of information processing in graded, random, and interactive networks. In: D.E. Meyer & S. Kornblum (Eds.), *Attention and Performance XIV: Synergies in Experimental Psychology, Artificial Intelligence, and Cognitive Neuroscience* (pp. 655–688). Cambridge, MA: MIT Press.

Miller, E.K. (2000). The prefrontal cortex and cognitive control. *Nature Reviews Neuroscience, 1,* 59–65.

Miller, E.K. & Cohen, J.D. (2001). An integrative theory of prefrontal cortex function. *Annual Review of Neuroscience, 21,* 167–202.

Myerson, J. & Hale, S. (1993). General slowing and age invariance in cognitive processing: the other side of the coin. In: J. Cerella & J.M. Rybash (Eds.), *Adult Information Processing: Limits on Loss* (pp. 115–141). San Diego: Academic Press.

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

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

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

O'Reilly, R.C., Braver, T.S., & Cohen, J.D. (1999). A biologically based computational model of working memory. In: A. Miyake & P. Shah (Eds.), *Models of Working Memory: Mechanisms of Active Maintenance and Executive Control* (pp. 102–134). Cambridge, UK: Cambridge University Press.

O'Reilly, R.C., Noelle, D.C., Braver, T.S., & Cohen, J.D. (in press). Prefrontal cortex and dynamic categorization tasks: representational organization and neuromodulatory control. *Cerebral Cortex.*

Owen, A.M., Lee, A.C.H., Williams, E.J., Kendall, I.V., Downey, S.P.M.J., Turkheimer, F.E., Menon, D.K., & Pickard, J.D. (1998). Redefining the functional organisation of working memory processes within human lateral frontal cortex. *NeuroImage, 7,* S12.

Petrides, M. (1996). Specialized systems for the processing of mnemonic information within the primate frontal cortex. *Philosophical Transactions of the Royal Society of London Series B, 351,* 1455–1462.

Posner, M.I. (1980). Orienting of attention. *Quarterly Journal of Experimental Psychology, 32,* 3–25.

Posner, M.I. & Petersen, S.E. (1990). The attention system of the human brain. *Annual Review of Neuroscience, 13,* 25–42.

Postle, B.R. & D'Esposito, M. (2000). Evaluating models of the topographical organization of working memory function in frontal cortex with event-related MRI. *Psychology, 28,* 132–145.

Raz, N., Gunning-Dixon, F.M., Head, D., Dupuis, J.H., & Acker, J.D. (1998). Neuroanatomical correlates of cognitive aging: Evidence from structural magnetic resonance imaging. *Neuropsychology, 12,* 95–114.

Rezai, K., Andreasen, N.C., Alliger, R., Cohen, G., Swayze, V., & O'Leary, D.S. (1993). The neuropsychology of the prefrontal cortex. *Archives of Neurology, 50,* 636–642.