A Large-scale Neurocomputational Model of Task-oriented Behavior Selection and Working Memory in Prefrontal Cortex

George L. Chadderdon and Olaf Sporns

Abstract

■ The prefrontal cortex (PFC) is crucially involved in the executive component of working memory, representation of task state, and behavior selection. This article presents a large-scale computational model of the PFC and associated brain regions designed to investigate the mechanisms by which working memory and task state interact to select adaptive behaviors from a behavioral repertoire. The model consists of multiple brain regions containing neuronal populations with realistic physiological and anatomical properties, including extrastriate visual cortical regions, the inferotemporal cortex, the PFC, the striatum, and midbrain dopamine (DA) neurons. The onset of a delayed match-to-sample or delayed nonmatch-to-sample task triggers tonic DA release in the PFC causing a switch into a persistent, stimulus-insensitive dynamic state that promotes the maintenance of stimulus representations within prefrontal networks. Other modeled prefrontal and striatal units select cognitive acceptance or rejection behaviors according to which task is active and whether prefrontal working memory representations match the current stimulus. Working memory task performance and memory fields of prefrontal delay units are degraded by extreme elevation or depletion of tonic DA levels. Analyses of cellular and synaptic activity suggest that hyponormal DA levels result in increased prefrontal activation, whereas hypernormal DA levels lead to decreased activation. Our simulation results suggest a range of predictions for behavioral, single-cell, and neuroimaging response data under the proposed task set and under manipulations of DA concentration. ■

INTRODUCTION

One of the central problems in cognitive neuroscience is that of executive control, including the selection of behaviors that are consistent with an organism's goals (task-oriented behavior selection [TOBS]). Generally, an associative mapping of immediate stimuli to responses is insufficient for motivating behavior that reflects the reward structure of the environment. Executive control requires *working memory* of sensory stimuli as well as of *tasks* and *task parameters*. A task involves the representation of a goal state or an intended program of action; task parameters are representations of target objects for the actions or parameters that adjust the execution characteristics of the actions. Based on current stimulus representations—including representations of object identity, location in the environment, and associated states of affect—on the basis of current drive state (e.g., hunger, pain, or fatigue), and the currently active task/s and task parameters, a motor program or set of motor programs must be selected to generate coherent, coordinated behavior. An open question concerns the nature of the input cues for the controlling mechanisms of working memory. The model presented in this article suggests that current task representations may serve as triggers for the engaging of working memory.

Executive control and working memory are believed to have their main loci in the prefrontal cortex (PFC) (Funahashi, 2001; Miller & Cohen, 2001). Both primate single-unit recording experiments (Rainer, Asaad, & Miller, 1998; Miller, Erickson, & Desimone, 1996; Funahashi, Bruce, & Goldman-Rakic, 1989; Fuster, 1973; Fuster & Alexander, 1971) and human neuroimaging experiments (Curtis & D'Esposito, 2003; Sakai, Rowe, & Passingham, 2002; Rowe, Toni, Josephs, Frackowiak, & Passingham, 2000; Smith & Jonides, 1999; Courtney, Petit, Maisog, Ungerleider, & Haxby, 1998) involving delayed-response tasks suggest that PFC delay activity serves a critical role in working memory, although the precise mechanisms are still being investigated. Delay activity may be increased or decreased with respect to baseline firing and may be tuned to particular stimuli, trial phases, or tasks (White & Wise, 1999). The PFC, especially the dorsolateral component, is also known to influence motor activity through connections both to the supplementary motor cortex and to the basal ganglia. Working memory performance depends on the level of extracellular prefrontal dopamine (DA), and is impaired if DA levels are elevated or suppressed

Indiana University

(Goldman-Rakic, Muly, & Williams, 2000; Zahrt, Taylor, Mathew, & Arnsten, 1997). One current theory is that PFC working memory representations, which may exert top-down biasing influences on processes as diverse as attention, action-selection, and episodic memory, are actively maintained and/or updated via neuromodulation through midbrain DA release (Braver & Barch, 2002; Cohen, Braver, & Brown, 2002; Dreher & Burnod, 2002; Miller & Cohen, 2001; Durstewitz, Kelc, & Güntürkün, 1999). At the core of our model is a proposed candidate mechanism by which midbrain DA tonic release may modulate PFC activity to allow the maintenance and release of stimulus working memory which, in turn, allows TOBS.

Large-scale neurocomputational models have addressed visual grouping and binding (Tononi, Sporns, & Edelman, 1992), working memory (O'Reilly & Frank, 2004; Frank, Loughry, & O'Reilly, 2001; Tagamets & Horwitz, 1998, 2000), and auditory perception (Husain, Tagamets, Fromm, Braun, & Horwitz, 2004), and they serve as important tools for correlating and interpreting data from multiple neuroscience methodologies (Horwitz, Tagamets, & McIntosh, 1999; Horwitz & Sporns, 1994). This article proposes a large-scale neurocomputational model simulating a simple instance of TOBS: the problem of selecting an appropriate response to a stimulus in a delayed match-to-sample (DMS) or delayed nonmatch-to-sample (DNMS) task. Our model incorporates several key features of cortical anatomy and physiology. Layers of neural population cortical column units model several ventral visual (Tononi et al., 1992) and prefrontal cortical areas (Tagamets & Horwitz, 1998, 2000). The model incorporates a mechanism by which tonic DA levels modulate PFC neuronal response functions (Servan-Schreiber, Printz, & Cohen, 1990), effectively regulating the maintenance of PFC activity patterns (Durstewitz et al., 1999). Tonic DA levels in the PFC are regulated by an intrinsic DA release mechanism whereby ventral tegmental area (VTA) terminal boutons are stimulated by PFC activity, similar to a mechanism proposed in a model of a delayed-alternation task (Dreher & Burnod, 2002). In our model, this mechanism is implemented using a leaky integrator equation with diffusion and reuptake constants governing extracellular DA concentration. In addition to a PFC stimulus representation, the model incorporates PFC task units to represent task state (O'Reilly & Frank, 2004; Frank et al., 2001). Another model component consists of basal ganglia (striatal) circuits that implement a competitive action selection mechanism releasing appropriate motor behaviors (Gurney, Prescott, & Redgrave, 1998, 2001).

In this article, we present an analysis of the model at multiple levels, including behavior, activity profiles of local neuronal populations, aggregate neural activity, integrated synaptic activity (ISA), and synthetic fMRI, in order to investigate the impact of tonic levels of DA on model dynamics and performance. Our analysis reveals that sharply increased or decreased levels of DA lead to specific disruptions in neural dynamics and behavior, with depleted DA resulting in disrupted PFC memory maintenance and elevated DA exerting a strongly suppressive effect on PFC activity. We discuss our results in light of current empirical findings and with an eye towards the role of DA fluctuations during autonomous behavior.

RESULTS

Model Overview and Trial Structure

A functional overview of the model is shown in Figure 1A and the structure of simulation trials in Figure 1B. During a trial, black-and-white "+," "T," or blank stimuli are presented to an object visual system which encodes stimulus identity. *Task working memory*, according to external *task selection cues*, activates a DMS, DNMS, or default (Idle) task. The DMS or DNMS tasks stimulate DA release within networks representing the *stimulus working memory*, thus enabling a representation of a stimulus to be retained until the end of the trial. An *action selection* system uses representations of current stimulus, remembered stimulus, and current task to unambiguously select a motor response.

A single iteration in the model corresponds to ~5 msec of real time (Tagamets & Horwitz, 1998). We simulate separate sets of DMS, DNMS (Figure 1B), and corresponding control trials, each of which consists of a series of trial phases: precue (50 iterations, 250 msec), cue, delay 1, distractor, delay 2, target, and post-trial (all 100 iterations, 500 msec). (For computational efficiency, we use shortened stimulus presentation and delay periods for most simulations shown in the article.) The model's sensory input consists of a 64 × 64 gray-scale image relayed to the model's retina/LGN and consisting of a centered high-contrast visual shape ("+," "T") or a blank stimulus ("_"). Behavior is simulated by the activity of two premotor units, M_{pos} and M_{neg}, whose activation triggers a positive (affirmative) or negative behavioral response, indicating a stimulus match or mismatch. The task schedule is implemented by activation of one of the task units (DMS, DNMS, or Idle) in the PFC.

Control trials involve passive viewing of stimuli (either following DMS or DNMS schedules) in the absence of working memory or DA fluctuations; the Idle task unit remains active during the entire trial. All other trials are *experimental trials* involving the execution of either a DMS or DNMS task. During the precue period, the Idle PFC task unit is active and "_" is presented. During the cue period, the target stimulus ("+") is presented and either the DMS or DNMS PFC task unit is activated (silencing the Idle unit through lateral inhibition) and remains active through the remainder of the trial, until

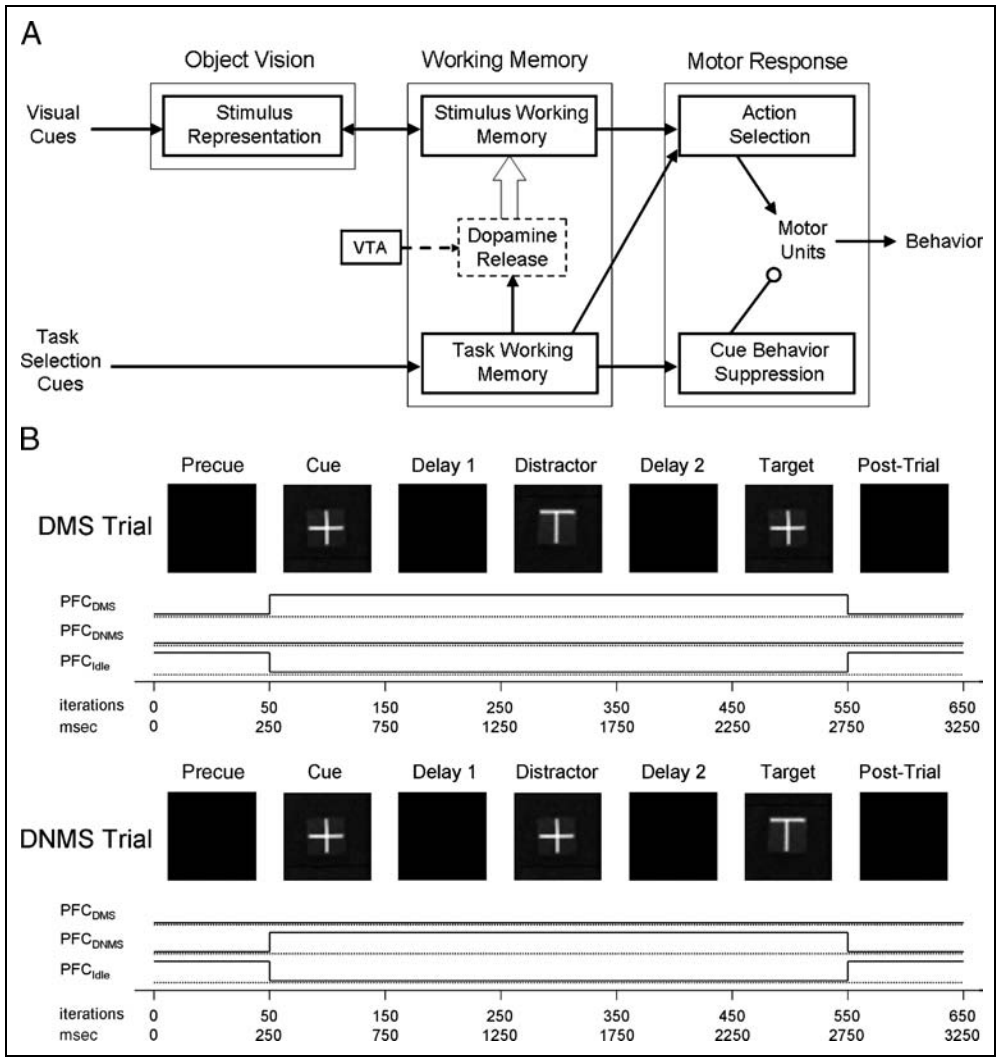

Figure 1. Basic model components and trial structure. (A) Model components include an object vision module, a working memory module, and a motor response module (for details on constituent brain regions, interconnectivity, and dynamic parameters, see Figure 3 and Methods). (B) DMS trial structure (top) and DNMS trial structure (bottom). Plots show task phases, representative visual stimuli, PFC task unit activity profiles, and time scale (1 iteration = 5 msec).

the post-trial period. Blank stimuli are presented for all delay periods. In the distractor period, a stimulus that should evoke a negative behavioral response is presented. After delay 2, during the target period, the stimulus that should evoke a positive behavioral response is presented. Finally, during the post-trial period, "_" is presented and the Idle unit is triggered, effectively resetting the task state of the model. DA effects are manipulated by varying the levels and effectiveness of DA present within the PFC. Subnormal DA levels are achieved by blocking DA effects on PFC neurons, whereas supernormal DA levels are the result of slower reuptake rates (Equations 8 and 9, see Methods).

Summary of Model Architecture

As the neocortex is believed to be organized in a columnar fashion both in posterior sensory cortices and in anterior executive areas such as the PFC (Rao, Williams, & Goldman-Rakic, 1999; Goldman-Rakic, 1996), we model several specialized brain areas as arrays of neural population units. Each unit represents a cortical column containing populations of excitatory and inhibitory neurons and their interconnections (Wilson & Cowan, 1972). Figure 2 shows the architecture of the basic unit which consists of excitatory (E), inhibitory (I), and tonic inhibition (I_0) subunits which interact through local connections. Under default parameter settings, the E subunit output activation—which corresponds to the proportion of currently firing pyramidal cells in the column—tracks the E subunit's net afferent input. Single-unit output results presented in this article (Figures 4 and 5) correspond to E subunit activations of the basic unit. Each columnar unit communicates with other units through specific sets of feedforward excitatory and inhibitory connections, lateral inhibitory connections, and voltage-dependent (modulatory) feedback excitatory connections. The patterns of anatomical connectivity incorporate known characteristics of corresponding primate cortical regions (more details of the cortical unit's architecture and layer connectivity are presented in the Methods section).

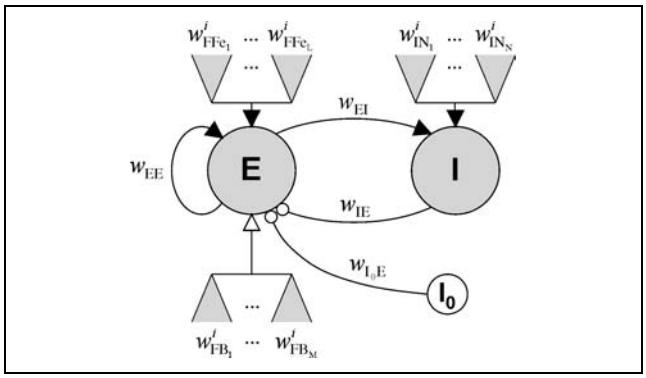

Figure 2. Basic cortical unit model. E = excitatory neuronal population; I = inhibitory neuronal population; I_0 = tonic inhibition; w = synaptic weights; shaded triangles represent connection arbors from feedforward (FF), feedback (FB), and lateral or feedforward inhibitory (IN) connections. Full arrows = excitatory voltage-independent connections; open arrows = excitatory voltage-dependent connections; open circles = inhibitory connections. Dynamic parameters used for all cortical units in the model (for exceptions, see Methods) are $w_{EE} = 0.60$, $w_{EI} = 0.15$, $w_{IE} = -0.15$, $\beta_{VD} = 50$, $\tau_{VD} = 0.3$, $\Delta_E = 1.0$, $\delta_E = 1.0$, $n_E = 0.001$, $\beta_E = 9.0$, $\tau_E = 0.3$, $\Delta_I = 1.0$, $\delta_I = 1.0$, $n_I = 0.001$, $\beta_I = 20.0$, $\tau_I = 0.1$.

During each trial, a ''+,'' ''T,'' or blank stimulus is presented to the object visual system of the model which consists of areas corresponding to the ventral visual stream of the primate cortex (Figure 3). In accordance with visual physiology, receptive field sizes and stimulus selectivity increase from lower (V1, V2/V4) to higher (inferotemporal; IT) levels of the visual hierarchy, whereas visual topography is progressively lost. In area IT, two populations of units, IT_+ and IT_T, selectively respond to either ''+'' or ''T,'' respectively, regardless of the location of these stimuli within the visual field (Figure 4A). These units constitute the object identity representation.

This identity representation serves as the input to the model's working memory architecture (Figure 3). Three PFC layers, PFC_S, PFC_D, and PFC_M, represent, respectively, the *current stimulus*, a *remembered stimulus*, and a *match representation* (a comparison of PFC_S and PFC_D). Task working memory consists of the three PFC task units previously described, which maintain recurrent self-excitation and lateral inhibition. Task selection cues (not explicitly modeled), which correspond to sensory cues that signal specific tasks (e.g., specific auditory ''task command'' stimuli), activate, presumably by a reward-conditioned associative mapping, one of the three tasks (DMS, DNMS, or Idle), which deactivates the others. The DMS and DNMS task units stimulate the release of DA (presumably from VTA terminals) into the extracellular compartment of PFC_D. Under normal DA conditions, a low DA concentration causes PFC_D to mirror the current stimulus representation pattern in PFC_S, and a high DA concentration causes the existing PFC_D pattern to be locked into its current state until the DA level drops. The DA modulation mechanism, described in more detail in the Methods section, operates by shifting the gain of the E subunit sigmoidal activation functions within PFC_D in direct proportion to DA concentration. Because the precue task set is ''Idle,'' the stimulus working memory is initially disengaged. However, during the cue period, PFC_S activation is relayed to PFC_D as DMS or DNMS task units become activated, causing the activity pattern to be maintained throughout the distractor and delay periods. Finally, in the post-trial period, the Idle task unit is reactivated (presumably by cues, not-modeled, that signal either rewards or punishes given in response to behavior), canceling DMS or DNMS activation and signaling the end of the current trial. VTA terminals are no longer stimulated and DA reuptake mechanisms cause the extracellular

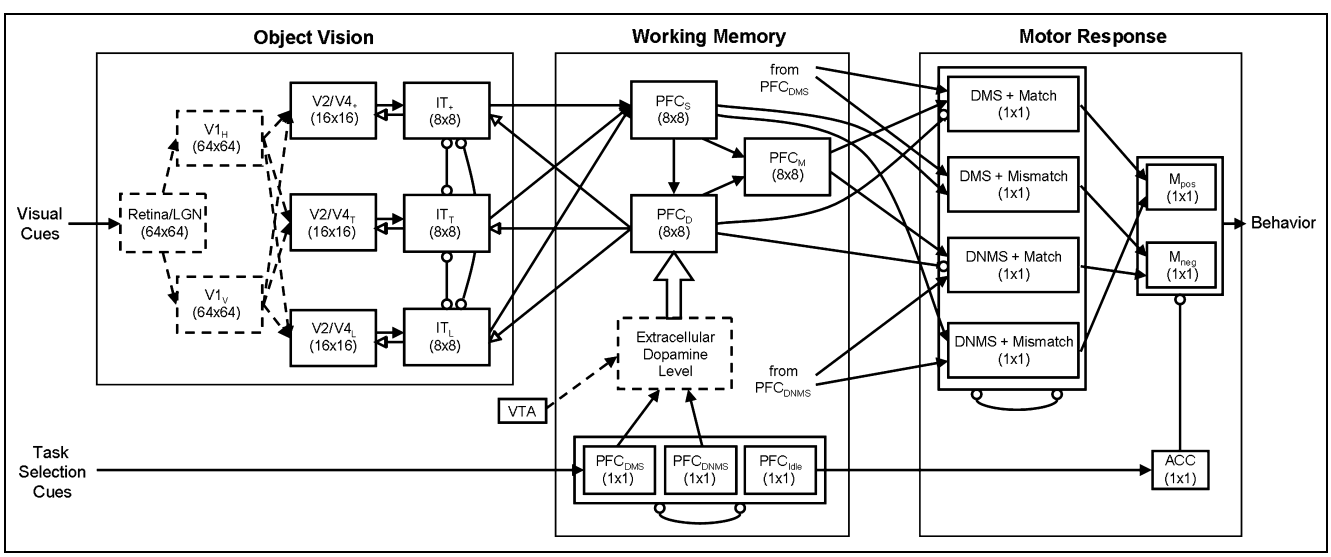

Figure 3. Modeled brain regions and interconnectivity. Numbers refer to sizes of modeled brain regions. Connection types encoded as in Figure 2.

level of PFC_D DA to decrease, which disengages working memory.

Task and stimulus working memory together allow the representation of current stimulus, current task, and current task parameters (i.e., the remembered cue stimulus). Together, these allow the action selection module (Figure 3) to unambiguously select a motor response. The module consists of a set of four striatal Task + Mis/match units which laterally inhibit each other and which uniquely select a positive or negative behavioral response. The Task + Match units are activated by the aggregate amount of matching representation (PFC_M) activity. The Task + Match activation is thresholded by inhibition triggered by aggregate activity of the remembered representation (PFC_D). This allows the match detection to be made on the *relative* activation of PFC_M with respect to the remembered stimulus rather than on the absolute magnitude of the match response which reflects the number of active units in PFC_D. The Task + Mismatch units are triggered by current stimulus activity in PFC_S and inhibited by activity in their equivalent Task + Match units.

During the cue period, the current stimulus is committed to memory, causing the matching activity to activate as well. Because the task is also set during this period, the striatal units will attempt to activate M_{pos} or M_{neg}. It seems plausible that an organism would be tempted to respond during the cue period of the DMS or DNMS and would need to be trained to *suppress* this response. One possible mechanism might be response inhibition by anterior cingulate cortex (ACC) activity. In this model, a single ACC unit activates when the Idle task is active and slowly decays after it is turned off. If the decay period is longer than the cue period, this allows the cue response to be suppressed.

Neural and Behavioral Data from a Single DMS Trial

Figure 4 shows neural and behavioral data obtained from a single representative DMS trial under normal DA conditions. Figure 4A shows an example of object-selective and translation-invariant unit responses in IT, whereas Figure 4B shows a map of the representation of stimulus identity through overlapping distributed patterns of activation within PFC. Both IT populations have random, uniformly distributed connections to the PFC_S layer, which models superficial-layer PFC cortical populations, proposed to be the neural substrate of current stimulus representation within PFC (Durstewitz et al., 1999). Units in PFC_S display a mixture of stimulus selectivity, ranging from highly selective (units shown in red and green in Figure 4B, right), to responsive, but nonselective (units in yellow), to unresponsive (units in black). Each PFC cell location effectively represents a functionally specialized ''macrocolumn'' or hypercolumn, consisting of PFC_S, PFC_D, and PFC_M units arranged in parallel, segregated circuits. Each PFC_S and PFC_D unit innervates solely its corresponding location in PFC_M, causing responses in the latter unit to be the ''logical AND'' of the response of the former two units. Given this structure, PFC_S cells have object-selective *receptive fields* that include none, one, or both of the stimuli used in the simulations. The corresponding PFC_D cells have object-selective *memory fields* (Rainer et al., 1998), whereas corresponding PFC_M cells have object-selective *matching fields* that enable object-related behavioral responses to occur.

Figure 4C shows snapshots of neural activity patterns obtained at the end of each trial phase in a representative DMS trial. During the precue period, there is no stimulus and only the Idle task unit and the ACC unit are active. At the end of the cue period, a ''+'' stimulus is present and the DMS task unit is fully active. Units in IT_+ are active, triggering activity in PFC_S, which propagates to both PFC_D and PFC_M. The DMS task unit's activation causes the extracellular DA level in PFC_D to become elevated. The high match activity in PFC_M is sufficient to overcome the inhibition threshold raised by PFC_D which activates the DMS + Match striatal unit. However, residual activity in the ACC unit prevents a behavioral response. At the end of the delay 1 period, because DA remains elevated, the cue stimulus activation is maintained in PFC_D while all other units (with the exception of the DMS PFC task unit) become inactive. At the end of the distractor period, a ''T'' stimulus has activated IT_T which has, in turn, activated PFC_S. Some PFC_M units (those with object-nonselective matching fields) become active, but the DMS + Match unit is not activated because PFC_D activity greatly exceeds PFC_M activity. Instead, PFC_S activates the DMS + Mismatch striatal unit, which triggers the negative behavioral response signaling a mismatch stimulus in a DMS task. Neural activations at the end of the delay 2 period are very similar to those at the end of delay 1. At the end of the target period, the target ''+'' stimulus has activated

Figure 4. Neural data obtained from a single representative DMS trial. (A) Sample visual stimuli and receptive field (RF) plots for a representative cortical unit in area IT_+. Stimuli are translated across the entire visual field (9 × 9 position grid) and average neural responses are plotted for ''+'' stimuli (top plots) and ''T'' stimuli (bottom plots). RFs exhibit object selectivity and translation invariance. (B) Object identity representation in PFC_S. Plots show neural responses averaged over 9 × 9 stimulus positions (as in A) to ''+'' (red color scale), and ''T'' (green color scale), as well as their superposition (red/green color scale). Representative unit locations, corresponding to PFC hypercolumns, are circled and labeled in blue. Recording traces from units at positions (2,4), (3,3), (4,4), (5,2) are shown in D, and in Figure 5. Scatterplot shows distribution of stimulus selectivity. (C) Representative neural activity patterns for one iteration at the end of seven trial phases. (D) Single-unit activity profiles for units in IT and PFC (q.v. panel B for unit locations). (E) Aggregate activity profiles for excitatory units in areas IT, PFC, action selection units, motor units, and DA (β_E, see Equation 8).

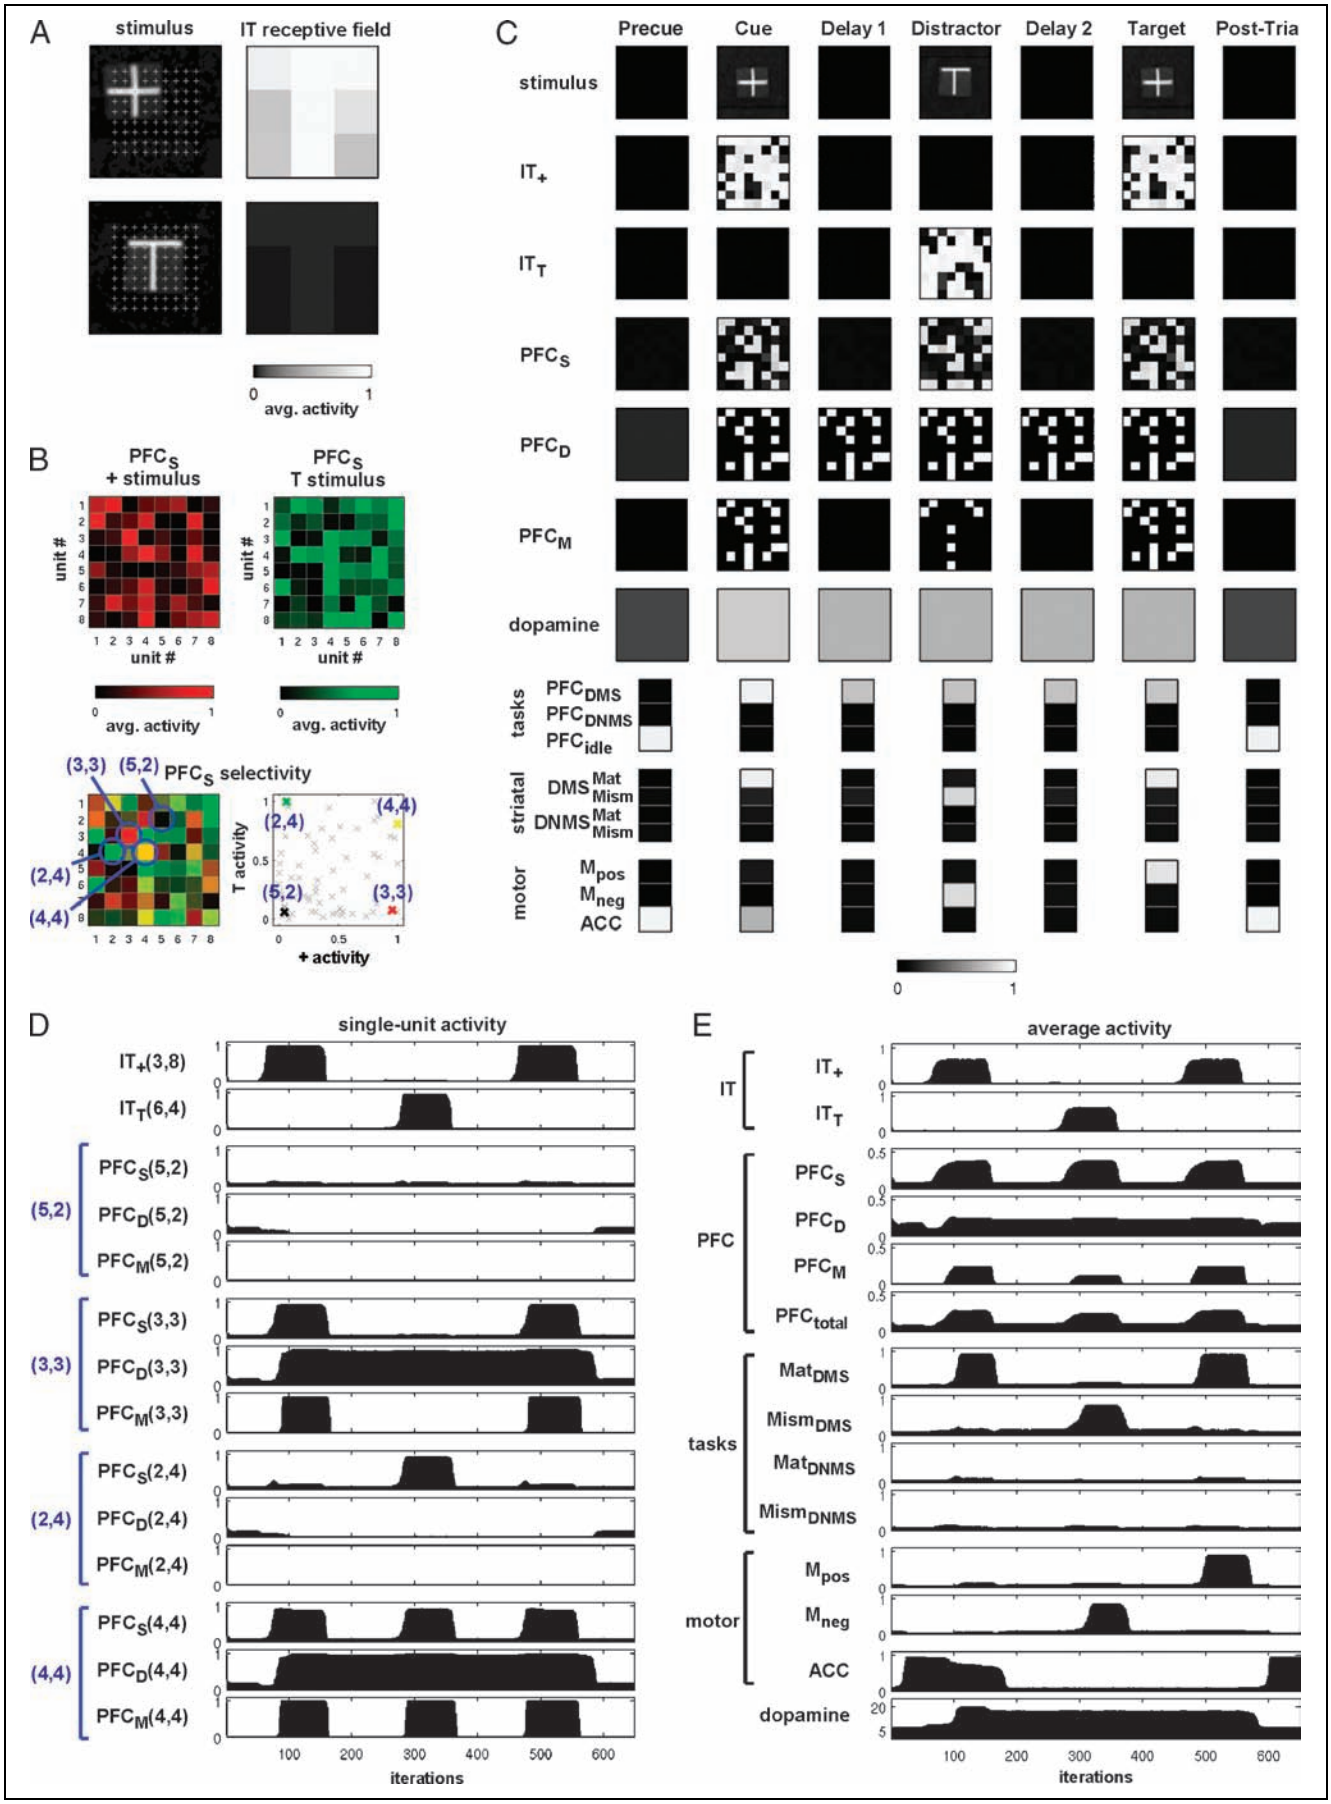
A
stimulus
IT receptive field
avg. activity
B
PFC_S + stimulus
PFC_S T stimulus
unit #
avg. activity
PFC_S selectivity
(3,3)
(5,2)
(2,4)
(4,4)
T activity
+ activity
C
Precue
Cue
Delay 1
Distractor
Delay 2
Target
Post-Tria
stimulus
IT+
ITT
PFCS
PFCD
PFCM
dopamine
tasks
PFCDMS
PFCDNMS
PFCidle
striatal
DMS Mat Mism
DNMS Mat Mism
motor
Mpos
Mneg
ACC
D
single-unit activity
IT+(3,8)
ITT(6,4)
(5,2)
PFCS(5,2)
PFCD(5,2)
PFCM(5,2)
(3,3)
PFCS(3,3)
PFCD(3,3)
PFCM(3,3)
(2,4)
PFCS(2,4)
PFCD(2,4)
PFCM(2,4)
(4,4)
PFCS(4,4)
PFCD(4,4)
PFCM(4,4)
iterations
E
average activity
IT
IT+
ITT
PFC
PFCS
PFCD
PFCM
PFCtotal
tasks
MatDMS
MismDMS
MatDNMS
MismDNMS
motor
Mpos
Mneg
ACC
dopamine
iterations

units in IT_+ and PFC_S. Matching PFC_M units signal the overlap between target and remembered cue stimulus, thus activating the DMS + Match striatal unit as well as the positive behavioral response. Finally, at the end of the post-trial phase, in response to trial termination cues (not explicitly modeled, but presumed to be stimuli that are reward-conditioned to trigger activation of the Idle task unit), the Idle PFC task unit is reactivated which resets the extracellular DA level in PFC_D, and unit activations return to the initial state.

Figure 4D shows ''recordings'' of single neural units obtained during a representative DMS trial from locations within PFC_S marked in Figure 4B. Recording traces show a nonresponsive unit (5,2), a ''+''-selective unit (3,3), a ''T''-selective unit (2,4), and a responsive, but nonselective unit (4,4). The recordings demonstrate that the receptive fields of PFC_S units, the memory fields of PFC_D units, and the matching fields of PFC_M cells show corresponding selectivity, confirming the existence of distributed modules spanning all three layers within the modeled PFC region. PFC_S activity closely follows IT inputs. PFC_D (memory field-driven) activity lasts from the beginning of the cue phase until the end of the target phase, in response to a stimulus in the cue period the memory field is selective for. PFC_M (matching field-driven) activity takes place only when receptive field activity from PFC_S and memory field activity from PFC_D coincide.

Figure 4E shows the aggregate IT, PFC, striatal, and motor area activity in the model, expressed as the sum of excitatory unit activity normalized by the size of the area, during a representative DMS trial. The resulting plots correspond to the time-dependent mean firing rate averaged over the excitatory units within each area of the model. PFC_S and PFC_M units are activated primarily during cue, distractor, and target phases of the trial, whereas PFC_D exhibits elevated neural activity throughout the trial with only relatively small rate modulations. Elevated activity in precue and post-trial phases is due to increased baseline activation resulting from low DA levels. On the other hand, elevated activity in cue, delay, distractor, and target phases reflects stimulus- or memory-specific activity patterns under high DA levels. The activity trace for PFC_{total} is calculated as $PFC_{total} = PFC_S + PFC_D + PFC_M$ representing an approximate measure of average excitatory activity in the PFC. Average PFC activity switches from a low baseline to high levels during the cue period and drops again during delay 1. Activity levels during the distractor period are below those during cue and target periods.

Effects of Dopamine on Delay Unit Activity

Figure 5 examines the effects of DA on memory field activity of selected units in PFC_D (cf. Figure 4D). Figure 5A compares DA profiles and unit activity during Idle and DMS task conditions. The lack of temporal modulation of DA during the Idle condition results in PFC_D responses that effectively mimic PFC_S receptive fields. Elevated DA levels during the DMS task modulate PFC_D unit activity such that cue stimulus representations are maintained during delay and distractor trial phases. Figure 5B shows variations in PFC_D unit activity as DA levels and effectiveness are manipulated. PFC_D units exhibit three sharply different response profiles. Intermediate (''normal'') levels of DA result in stimulus- and task-dependent activity that appropriately reflects memory task demands. Increased DA levels (''hyper DA'') result in strongly diminished PFC_D activity that no longer reflects stimulus or task-dependent inputs. PFC_S activity is effectively blocked from propagating to PFC_D and, in turn, to PFC_M. This tends to cause failures to detect stimulus matches as match striatal units will not be triggered by PFC_M activity. Decreased DA effectiveness (''hypo DA'') results in increased baseline activity and stimulus-dependent responses that reflect associated PFC_S receptive fields, but fail to reflect task demands as PFC_D memory fields lose their sensitivity to task state. The elevated baseline firing rate of PFC_D units results in strong inhibition of striatal match units, leading to a failure to respond appropriately to the target stimulus.

Effects of Dopamine on Behavior and Network Activity

Table 1 shows behavioral frequencies for DMS and DNMS trials under normal DA levels, and for DMS trials under conditions of increased DA levels or decreased DA effectiveness. The model emits positive and negative behaviors consistent with the task set if DA is within ''normal'' range. Some behavioral errors occur, which are largely the result of fluctuations in the distributed pattern of PFC responses and resulting mismatches in PFC_M. As effective DA levels are decreased or increased, behavioral responses become uniformly negative and no longer reflect task demands; although the model continues to exhibit high correct-rejection rates during the distractor phase, the matching target stimulus is generally incorrectly rejected, resulting in a high failure rate for detecting targets.

Figure 6A shows summary plots of aggregate activity of excitatory units (mean, *SD*, $N = 20$ trials), measured as the average of the last 50 iterations (250 msec) of each trial phase. Three different levels of DA are examined, labeled ''normal'' ($r = .4, a = .0$), ''hypo'' ($r = .4, a = 1.0$), and ''hyper'' ($r = .1, a = .0$). For each DA condition, we compare idle to DMS task conditions (DNMS trials show qualitatively similar results; data not shown). Activity in the IT closely follows stimulus presentations, whereas PFC activity reflects both representations of stimulus cues and maintained memory traces. PFC_D and PFC_M aggregate activity is strongly modulated by task condition (Idle vs. DMS) as well as DA levels and effectiveness. Highest levels of PFC activity are seen as

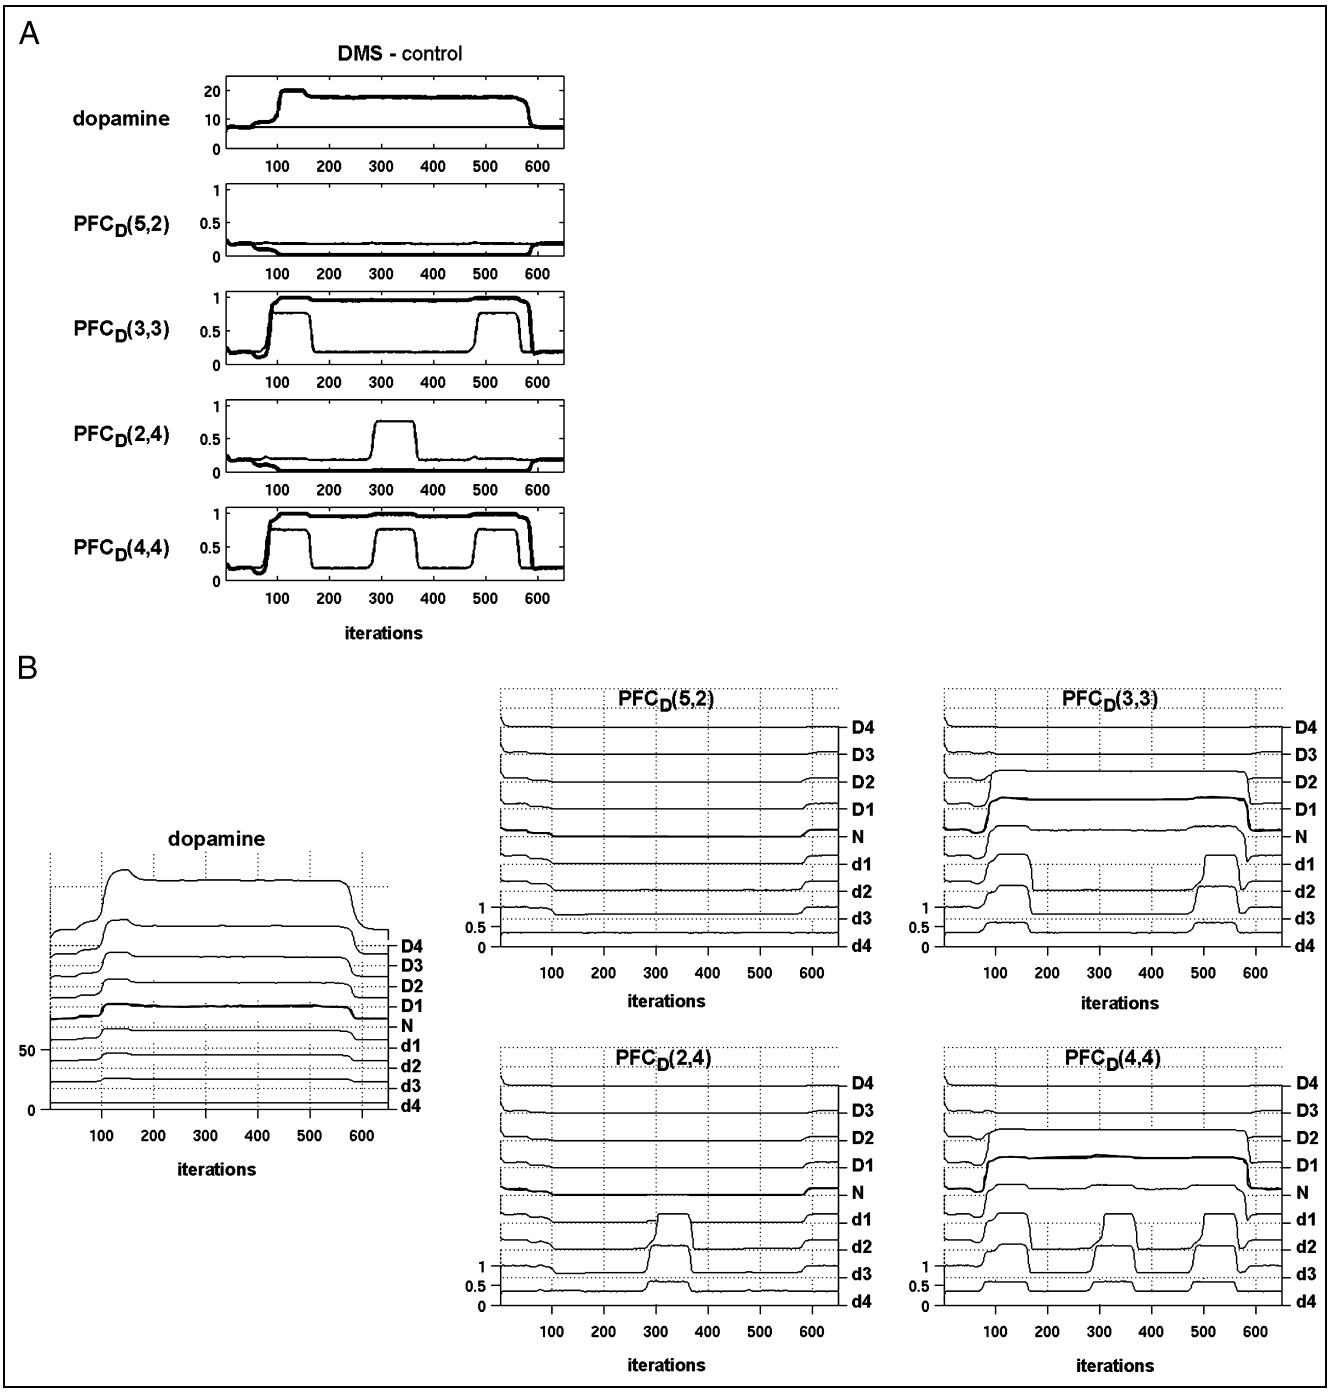

Figure 5. Effects of DA levels on PFC single-unit activity (see Figure 4 for description of single-unit locations). (A) Comparison of DA gain (β_E) and PFC_D single-unit activity profiles during control (thin line) and DMS task conditions (thick line). (B) DA gain (β_E) and PFC_D single-unit activity profiles during a DMS working memory task, with normal (N), elevated (D1–4), and decreased (d1–4) DA levels. DA conditions are labeled N (a = .0, r = .4), d1 (a = .25, r = .4), d2 (a = .5, r = .4), d3 (a = .75, r = .4), d4 (a = 1.0, r = .4), D1 (a = .0, r = .325), D2 (a = .0, r = .25), D3 (a = .0, r = .175), D4 (a = .0, r = .1). a is a DA antagonism parameter that takes values 0 to 1 with the latter setting full DA channel blockage. r is a DA reuptake parameter whose reduction leads to increasing agonist effects (see Equations 8 and 9).

stimuli are presented during Idle trials at ''normal'' DA levels. In the absence of task unit-mediated DA release, all PFC units within a single ''macrocolumn'' will effectively share the same receptive field and become activated. DA modulation during a DMS task results in decreased activity levels during stimulus phases and increased activity during delay phases. ''Hypo DA'' levels lead to robust increases in summed PFC activation, whereas ''hyper DA'' has an inhibitory effect on PFC activity. During delay phases, PFC activity is increased during DMS trials as compared to Idle trials if DA levels are ''normal,'' but it is decreased if DA levels are

Table 1. Behavioral Data for Different DA Conditions

Trial Type and DA Condition (N = 20)	*Distractor (Pos/Neg)*	*Target (Pos/Neg)*	*"Correct Trials"*
DMS			
D4	0/17	0/18	0/20
D3	0/18	0/19	0/20
D2	0/19	0/20	0/20
D1	0/18	10/7	8/20
N	**0/18**	**18/3**	**15/20**
d1	0/18	17/3	15/20
d2	1/18	3/17	3/20
d3	6/13	8/12	2/20
d4	0/20	0/20	0/20
DNMS			
N	**3/15**	**20/0**	**15/20**

a denotes the antagonist parameter and *r* denotes the reuptake rate (see Equations 8 and 9 in the Methods section).

The first column shows the number of M_{pos} and M_{neg} unit responses to distractor stimuli and the second column shows the number of M_{pos} and M_{neg} unit responses to target stimuli. The third column shows the number of "correct trials" out of 20, that is, the number of trials where the distractor triggers only M_{neg} and the target triggers only M_{pos}. DA conditions are labeled N ($a = .0$, $r = .4$), d1 ($a = .25$, $r = .4$), d2 ($a = .5$, $r = .4$), d3 ($a = .75$, $r = .4$), d4 ($a = 1.0$, $r = .4$), D1 ($a = .0$, $r = .325$), D2 ($a = .0$, $r = .25$), D3 ($a = .0$, $r = .175$), D4 ($a = .0$, $r = .1$).

elevated. No change is seen in the "hypo" condition, as DA modulation is absent in either the idle or DMS task.

Significant differences in aggregate neural activity across task or DA levels are seen not only in PFC, but also in IT (Figure 6B). During stimulus presentations, IT activity is slightly suppressed during DMS tasks, as compared to idle tasks. This relationship is reversed as DA effects are blocked, and decreased as DA levels rise. This pattern mirrors corresponding but much larger changes in PFC aggregate activity, communicated to IT via voltage-dependent feedback connections arriving from PFC_D.

Integrated synaptic activity profiles for all units (excitatory and inhibitory) within IT and PFC are shown in Figure 7A, calculated according to Equation 7 and averaged over the last 50 iterations of each trial phase. As in Figure 6, we examine three levels of DA. IT synaptic activity reflects inputs from areas V2/V4, PFC, and intrinsic circuits. PFC synaptic activity is largely composed of IT inputs as well as signals between PFC layers and local circuits. Synaptic activity profiles across trial conditions and phases strongly resemble excitatory aggregate activity (Figure 6).

Integrated synaptic activity was used to derive simple analogs of metabolic BOLD signals, such as might be observed in fMRI experiments. Figure 7B (left) shows the time course of PFC ISA for two sets ($N = 20$) of idle and DMS trials under "normal" DA conditions. Trial phases were lengthened to 10 sec for pre- and post-trial phases, 1 sec for cue, distractor, and target phases, and 3 sec for delay phases. Clearly, DMS synaptic activity is suppressed during stimulus presentations (relative to the control trials) and elevated in the delay phases of the trials. To derive analogs of metabolic signals, we converted the raw time series to 50-msec slices and convolved the resulting idealized instantaneous response with a hemodynamic response function, modeled as a gamma function with a 2-sec (Figure 7B, middle) or 6-sec (Figure 7B, right) peak profile. Resulting time courses are sampled every 1 sec to yield simulated fMRI signals (Horwitz & Tagamets, 1999). Increased levels of DA or decreased DA effectiveness strongly influenced PFC baseline activity in a manner consistent with single-unit (Figure 5) and aggregate activity (Figure 6) data. At "normal" DA levels, simulated fMRI signals are increased for DMS tasks as compared to idle tasks. At "hyper" DA levels, fMRI signals show the opposite pattern, whereas "hypo" DA levels produce no change due to lack of DA modulation in either condition.

DISCUSSION

We presented a large-scale computational model providing a comprehensive framework for TOBS, working memory, and neuromodulation, on the basis of the anatomy and physiology of PFC and associated brain regions. The goal of the model was to investigate the functional role of DA on working memory and behavior selection by implementing a biologically based mechanism for the action of DA on PFC and then examining the effects of varying levels of tonic DA on behavior, as well as local neural populations, aggregate neural activity, and ISA within the PFC and other connected brain regions. Our results indicate that tonic prefrontal DA levels are critical for stimulus working memory and that task-dependent fluctuations in tonic DA are crucial for maintaining and releasing working memory representations.

The model incorporates a number of principles and hypotheses about the neural basis of working memory. We implement two separate domains of prefrontal working memory. Task identity is maintained by a segregated set of recurrently excitatory and mutually inhibitory cell populations (PFC task units). Stimulus feature memory, on the other hand, is maintained by tonic DA level, which modifies PFC dynamics so that activation patterns are maintained even as new stimuli are encountered. Stimulus identity in the PFC is represented through distributed and overlapping activation patterns driven by stimulus-selective and translation-invariant representations of objects arriving from areas of the ventral visual stream, such as the IT. Current

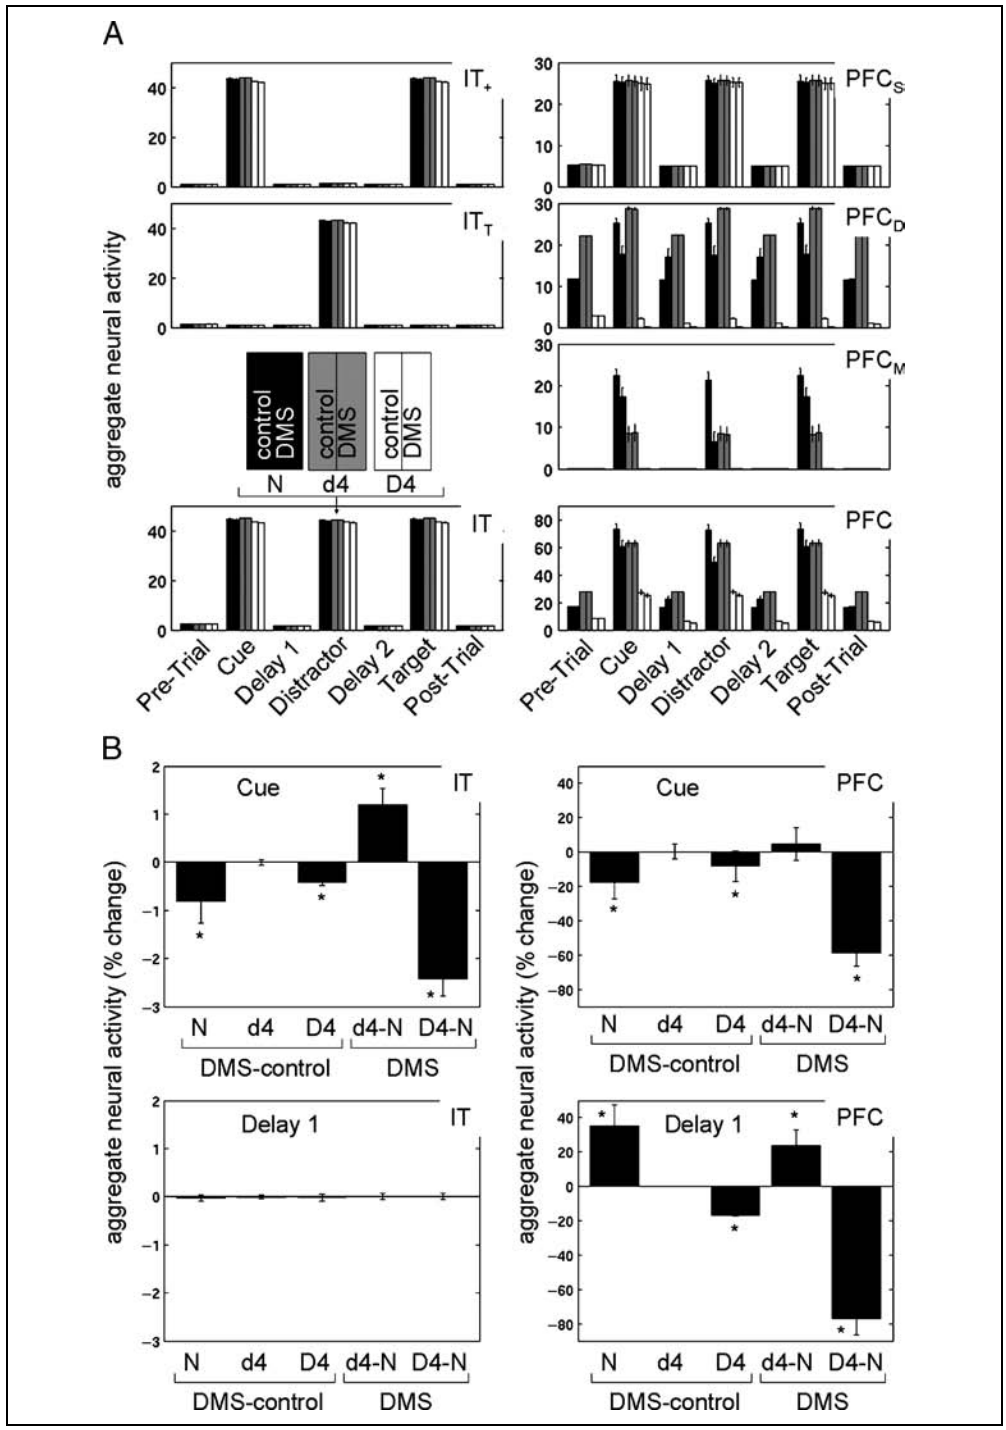

Figure 6. Effect of DA levels on IT and PFC aggregate excitatory activity. (A) Aggregate excitatory activity in the IT and the PFC. DA levels are: N (''normal DA'') (black bars, left: control, right: DMS, $a = .0, r = .4$), d4 (''hypo DA'') (gray, $a = 1.0, r = .5$), D4 (''hyper DA'') (white, $a = .0, r = .1$). (B) Changes in IT and PFC aggregate excitatory activity during the first cue and the first delay phase. Data are presented as percent change over baseline, * indicating $p < .01$ (t test, $n = 20$ trials per condition). The first three bars in each graph subtract control from DMS activity for three DA conditions (N, d4, D4). The last two bars subtract DMS normal DA from DMS ''hypo DA'' (d4) and DMS ''hyper DA'' (D4) activity, respectively.

stimuli may be compared with the maintained stimulus representation within segregated PFC hypercolumns, with more pattern overlap leading to more PFC_M activity. Our PFC hypercolumn architecture predicts that cells in PFC will be found that have stimulus-selective receptive, memory, or matching fields. For behavior selection, we suggest that units in the striatum may combine PFC representations of task, remembered stimulus, and match of current to remembered stimulus to disinhibit appropriate behaviors. Cognitive negation (e.g., of matching) might be accomplished through defaulting to a negative representation (e.g., DMS + Mismatch activation in the model) and allowing positive representations (e.g., DMS + Match) to laterally inhibit it. DA effects are implemented by modulating the excitatory gain of units in PFC_D, which results in changes in their responsiveness to inputs and altered baseline activity levels. Fluctuations in tonic DA result in alterations of the recurrent excitatory dynamics of ensembles of PFC_D neurons, with high levels of DA effectively stabilizing

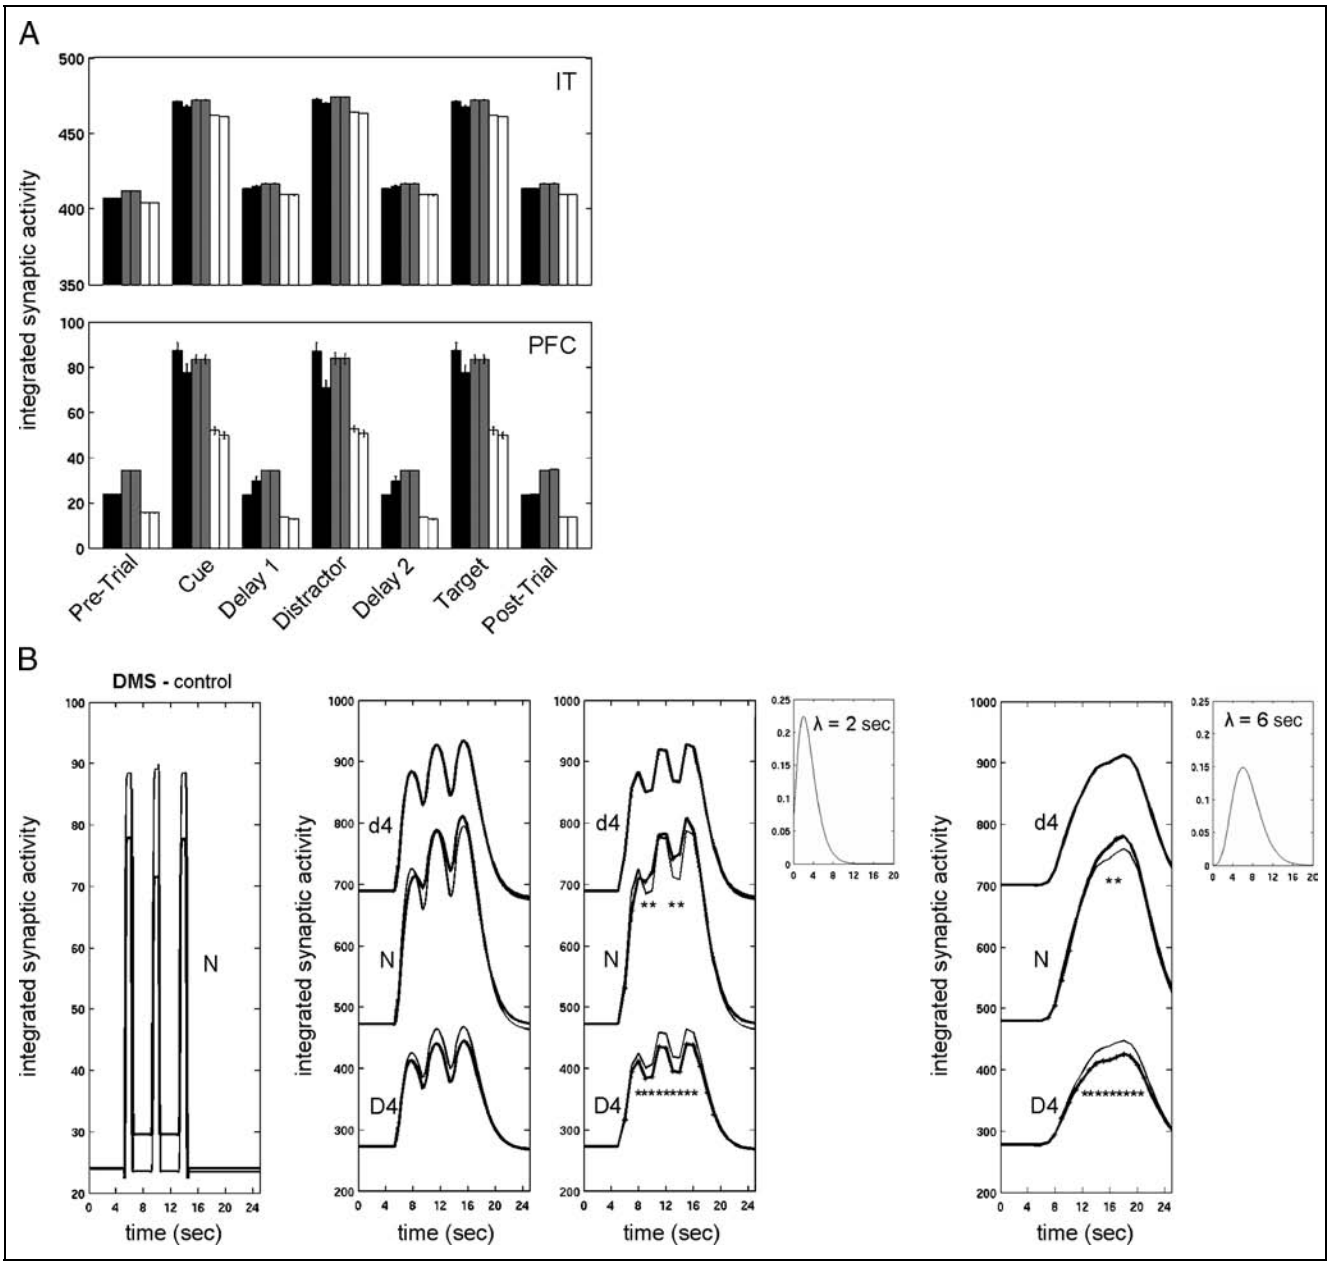

Figure 7. Effect of DA levels on IT and PFC integrated synaptic activity and synthetic BOLD responses. (A) Integrated synaptic activity (in arbitrary units), calculated according to Equation 7, for the IT and the PFC during DMS trial phases. Bar plot arranged as in Figure 6A. (B) (left) Time course of PFC integrated synaptic activity for idle (thin line) and DMS (thick line) task conditions and DA condition N, plotted in 50-msec time slices. Trial phases have been lengthened to 1 sec for stimulus presentations (cue, distractor, target) and 3 sec for delays. (B) (middle) Integrated synaptic activity (DA conditions N, d4, and D4) convolved with a gamma function (inset), derived by fitting to a Poisson function with $\lambda = 2$ sec, representing hemodynamic delay. Plot on the right shows data sampled every 1 sec, with * denoting statistical comparisons with $p < .05$ ($N = 20$ trials per condition). (B) (right) Same as middle panel, except using a gamma function with a longer time course (corresponding to a Poisson function with $\lambda = 6$ sec, see inset).

stimulus pattern attractors that are otherwise vulnerable to disruption by current stimuli (Durstewitz & Seamans, 2002; Durstewitz et al., 1999). Task representations in PFC directly stimulate release of DA. Tonic DA levels are mediated by an *intrinsic* mechanism whereby PFC cells stimulate the dopaminergic terminals of the VTA neurons. Such a mechanism has been shown to regulate tonic DA levels in the ventral striatum (Grace, 1991), and there is evidence that PFC glutamatergic activity in fact regulates PFC DA release (Takahata & Moghaddam, 1998). Due to the intrinsic release mechanism and the distributed hypercolumn architecture, concentrations of tonic DA in PFC may be spatially heterogeneous.

In analyzing model performance, we compared activity profiles from local neuronal populations, aggregate neural activity, and ISA under variations of tonic DA levels. At

the unit (local population) level, most significant differences were found in the response characteristics of PFC_D units. Although receptive fields of PFC_S units were found to be largely independent of DA levels, PFC_D units lost their capacity to maintain or release stimulus working memory representations under both elevated and decreased DA levels. Stimulus-selective receptive and memory fields have been described in PFC single-cell recording experiments using various delayed-response tasks (Rainer et al., 1998; Miller et al., 1996; Funahashi et al., 1989). There is also evidence that these prefrontal memory fields can be degraded by either elevated or suppressed DA levels (Goldman-Rakic et al., 2000; Zahrt et al., 1997). In our results, disruptions of single-unit response profiles lead to sharp declines in behavioral performance under both elevated and decreased tonic DA levels. Pharmacological experiments have provided evidence that performance of working memory tasks is in fact dependent upon ambient levels of prefrontal DA. In these experiments, elevated as well as depleted levels of DA degrade behavioral performance.

The model allowed us to derive synthetic fMRI signals, following a set of computational steps and assumptions previously used in similar studies (Horwitz & Tagamets, 1999). In our model, under normal DA conditions, we observe a significant increase in PFC activity in a DMS working memory task, as compared to an ''idle'' control condition. If a relatively fast hemodynamic delay function is used to derive fMRI signals, the increase is confined to the delay periods of the task, and absent during the cue/distractor/target periods. If tonic DA levels are decreased, we see higher baseline activation of PFC, but the activity differences between idle and working memory conditions are insignificant. If tonic DA levels are elevated, baseline activation of PFC is reduced and activity during a working memory task is decreased with respect to corresponding control trials, thus showing a reversal of the increase found under normal DA levels. This suppressive effect of DA is consistent with observations suggesting an inhibitory role for prefrontal DA (Gao, Krimer, & Goldman-Rakic, 2001; Gulledge & Jaffe, 1998).

A number of open questions remain. One of them concerns the possible role of DA in maintenance and updating of task working memory, which was not considered in the present article. Also, we did not touch upon the issue of learning or long-term plasticity of stimulus representations or behavior selection. Here, neuromodulatory systems may play an important role, for example, through phasic VTA firing resulting from unexpected rewards (Schultz, 2000), which may trigger prefrontal plasticity to strengthen the specific connections along pathways from sensory representations to motor behavior encoding (Miller & Cohen, 2001). Other open issues concern the interplay between PFC and the basal ganglia, the functional role of the ACC (Paus, 2001), and the effects of neuromodulators other than DA, for example, acetylcholine (Hasselmo, 1995) or noradrenaline (Rossetti & Carboni, 2005), in TOBS and working memory.

Future work will involve the implementation of the present model as part of the control architecture of an autonomous robot. Fluctuations in tonic and phasic DA depend on behavior (Phillips, Ahn, & Floresco, 2004), but DA concentration and its modulatory effect within PFC and elsewhere are difficult to assess in behaving animals or humans. A neurorobotic system may allow the recording and correlated analysis of behavioral and neural (including synthetic fMRI) data sets. Furthermore, integration of the present model with a model of phasic DA release in reward conditioning (Sporns & Alexander, 2002) may offer additional insights into the functional roles of neuromodulatory systems in the course of autonomous behavior.

METHODS

Hardware and Software

Stimuli consist of 4×4-in. high-contrast black-and-white patterns presented on a black background. Images were acquired with a stationary CCD camera, cropped and down-sampled to a resolution of 64×64 pixels, covering about 50° visual angle, with pixel values ranging between 0 and 1. Using two different shapes (''crosses'' and ''tees'') resulted in three different stimulus classes, referred to as ''+'' (cross), ''T'' (tee), and ''_'' (blank). Stimuli were translated to different positions without changing their orientation or size.

The simulation model uses MATLAB (Mathworks) run on a PC workstation. Full simulations required about 350 msec/iteration, and all data analysis was carried out off-line.

Cortical Column Unit Architecture

The architecture of our model consists of columnar units as basic computational elements (Wilson & Cowan, 1972) forming segregated cortical areas linked by networks of long-range reciprocal (reentrant) connections (Tononi et al., 1992). The aggregate activity of neural populations within individual columns corresponds to their mean firing rate and is taken as the column's output activity.

Figure 2 shows the architecture of a cortical column unit as used in our model. Each column consists of three units which each stand for different subpopulations of neurons within the column. The first unit (E) represents all of the excitatory neurons in the column. The second (I) represents inhibitory interneurons that respond to either local or afferent excitatory input. Both of these units have output activations ranging from 0 to 1. The third unit (I_0) represents a source of tonic inhibition, with an activity level fixed at 1. The weights connecting

these three units (i.e., w_{EE}, w_{EI}, w_{IE}, and w_{I_0E}) correspond to the relative strengths of the influence of these units on each other, and take on positive values if their ''presynaptic'' units are excitatory (i.e., in w_{EE} and w_{EI}) and negative values otherwise (w_{IE} and w_{I_0E}). Afferent weights into the E and I units are all positive (i.e., excitatory) and represent the relative influence of the various kinds of afferent inputs to the column. Afferent weights are organized into fan-in arbors where all of the ''synapses'' within an arbor come from the same lateral, upstream, or downstream cortical area. There are three arbor types. The excitatory feedforward arbor relays inputs to the E unit, whereas the feedforward (or lateral) inhibitory arbors relay inputs to the I unit. Both of these types of inputs (as well as all of the local connections) are voltage-independent. The third arbor type uses voltage-dependent connections enabling modulatory influences via feedback connections.

Inputs to E and I subunits are a weighted sum of the afferent and local synaptic inputs. A nonlinear sigmoidal ''squashing'' function is applied to yield output activations between 0 and 1. The excitatory input is a sum of voltage-independent and voltage-dependent inputs. Thus,

$$i_E(t) = i_{VI}(t) + i_{VD}(t) \tag{1}$$

where t is the current iteration number. Time constants and dynamics are scaled such that an iteration corresponds to about 5 msec. The voltage-independent input is defined by:

$$i_{VI}(t) = \sum_{i,j} w^j_{FFe_i} o^j_{FFe_i}(t-1) + w_{EE} o_E(t-1) + w_{IE} o_I(t-1) + w_{I_0E} \tag{2}$$

The first term—the summation—is the weighted sum of all of the previous time-step's output activations of every ''presynaptic'' unit from every feedforward arbor. ($w^j_{FFe_i}$ refers to the weight from the jth neuron of the ith feedforward arbor.) The second term is the weighted previous time value of the output activation of the (local) E unit. The third term is the weighted previous time value of the output activation of the (local) I unit. Finally, the last term is a constant tonic inhibition weight.

The voltage-dependent activity is defined by:

$$i_{VD}(t) = \theta_{VD}(i_{VI}(t)) \cdot \sum_{i,j} w^j_{FB_i} o^j_{FB_i}(t-1)$$

$$\theta_{VD}(x) = \frac{1}{1 + \exp(-\beta_{VD}(x - \tau_{VD}))} \tag{3}$$

The first factor in the activation equation is sigmoid-squashed voltage-independent activation, and thus, falls between 0 and 1. This factor weights the second factor which is the actual weighted sum of all of the output activations of every presynaptic neuron from every feedback arbor from the previous time step. The gain β_{VD} and threshold τ_{VD} of the sigmoid determine how much voltage-independent activity the neuron needs to have in order for these voltage-dependent channels to exert an influence. The output activation of the E unit is defined by the difference equation:

$$o_E(t) = o_E(t-1) + \Delta_E \theta_E[i_E(t) + N(0, n_E)] - \delta_E o_E(t-1) \text{ where}$$

$$\theta_E(x) = \frac{1}{1 + \exp(-\beta_E(x - \tau_E))} \tag{4}$$

The first term is the past output activation of the E unit. The second term is the growth term where the input activation (plus noise) is weighted by a growth parameter Δ_E. The third term is the decay term where the previous output activation of the unit is weighted by the decay parameter δ_E. $N(0,n_E)$ denotes noise with mean 0 and standard deviation n_E.

The inhibitory input activation is a sum of input activations of voltage-independent channels. It is defined by:

$$i_I(t) = \sum_{i,j} w^j_{IN_i} o^j_{IN_i}(t-1) + w_{EI} o_E(t-1) \tag{5}$$

The first term here—the summation—is the weighted sum of all of the previous time-step's output activations of every ''presynaptic'' unit from every lateral or upstream arbor. The second term is the weighted previous time value of the output activation of the (local) E unit. The output activation of the I unit is defined by the difference equation:

$$o_I(t) = o_I(t-1) + \Delta_I \theta_I[i_I(t) + N(0, n_I)] - \delta_I o_I(t-1) \text{ where}$$

$$\theta_I(x) = \frac{1}{1 + \exp(-\beta_I(x - \tau_I))} \tag{6}$$

Growth is weighted by a growth parameter Δ_I, whereas decay is weighted by the decay parameter δ_I. $N(0,n_I)$ denotes noise with mean 0 and standard deviation n_I.

Many of the parameters defined by these equations have values that are held common for every cortical column unit of the model (see Figure 2, with exceptions to be noted later). These include the local weights (not including the tonic inhibition weight), the gains and thresholds for all of the sigmoidal functions, the ambient noise parameters, and the growth and decay parameters. Parameter values were chosen for their performance characteristics, and for most of the cortical column units were tuned so that output activation tracks the rise and

fall of the input activation of the units (i.e., with no oscillatory or chaotic dynamics at the single-unit level). When the feedback arbors and the tonic inhibition have their weights set to 0, the result is theoretically close to the architecture proposed by Wilson and Cowan (1972) for modeling populations of neurons. Wilson/Cowan units were also used in the large-scale model of Tagamets and Horwitz (1998, 2000).

The cortical column unit explained in this subsection is used throughout all layers of the network that are modeled with such units, save in three exceptions. First, in the PFC_M layer, the excitatory gain β_E is set to 30.0 rather than 9.0, and the self-excitatory weight w_{EE} is set to 0.30 rather than 0.60. This causes the units to respond in a more Boolean, less graded fashion to the ''logical AND'' of the output activity from the corresponding PFC_S and PFC_D units. Second, the ACC unit has w_{EE} increased to 0.66, which results in longer persistence of ACC unit activity. Third, for the PFC_D layer, the E-subunit gain β_E is allowed to vary over time as a function of the extracellular level of DA (see below for details).

Layer Structure and Interconnectivity

Cortical column units are organized into arrays, each of which models a cortical area. These areas are interconnected through feedforward, lateral, and feedback connections which are organized into arbors of different shapes. Feedforward connections pass information forward from areas of the model that process and represent lower-level features to areas that process and represent higher-level features. Feedback connections are excitatory and are voltage-dependent in their activation, which means that they are only activated when there is some minimal level of activity already in the column unit. Thus, the feedback connections are used by the later-stage layers to sustain rather than initiate a representation at an earlier stage. The lateral connections are inhibitory (through excitatory connections to the I subunits of the column units) and are used to allow competition between rival representations.

There are four afferent arbor shapes that are used for connecting layers together. The first is a full arbor. The entire grid of the source layer feeds afferents into every unit of the target layer. The second type is a uniform sparse connection. Here, the scope of the entire grid of the source layer is covered but only a certain proportion of the neurons in the source layer are actually connected synaptically to each target layer unit. Third, there is a square-topographic connection type. Here, each target neuron is connected to a small rectangular window of the source layer. The center of this window is on the unit in the source layer that corresponds to the location of the target unit within the target layer. Thus, the projection of activation is topographic with larger regions of the source map being projected to smaller regions of the target map. Finally, there are one-to-one connections between layers. Here, each target unit has a single afferent from the topographically corresponding unit in the source layer.

The values of all of the afferent weights are set before the simulation is run and their values remained fixed for the duration of the trials (with the exception of the DA-modulated neurons, as discussed later). In the most general case, these values are chosen randomly with a particular mean and standard deviation.

Aggregate Activity Measures

The aggregate neural activity in Figure 6 is the sum of the graded output activity of all E subunits in a particular brain area (e.g., PFC_D aggregate activity sums all output activations of the E subunits in PFC_D). The ISA (Figure 7) is computed as:

$$\Theta_{\text{area}}(\text{t}) = \sum_{\text{k}\in\text{area}} in_{\text{k}}(t) \text{ where}$$

$$\begin{aligned} in(t) = &\sum_{i,j} w^j_{\text{FFe}_i} o^j_{\text{FFe}_i}(t-1) + w_{\text{EE}} o_{\text{E}}(t-1) \\ &+ \mid w_{\text{IE}} \mid o_{\text{I}}(t-1) + \mid w_{I_0E} \mid \\ &+ \sum_{i,j} w^j_{\text{FB}_i} o^j_{\text{FB}_i}(t-1) \\ &+ \sum_{i,j} w^j_{\text{IN}_i} o^j_{\text{IN}_i}(t-1) + w_{\text{EI}} o_{\text{E}}(t-1) \end{aligned} \tag{7}$$

The ISA for an area corresponds to the sum of the weighted afferent activity for all E, I, and I_0 subunits with all weights treated as positive and voltage-dependent connections treated as voltage-independent. ISA is used to derive synthetic fMRI signals (see Results).

Dopaminergic Modulation of Stimulus Working Memory

The simulated extracellular concentration of DA modifies the dynamic properties of PFC_D units. This mechanism involves modulation of the gain parameter β_E for the E subunits such that low DA concentrations correspond with a low gain and high DA concentrations with a high gain. This mechanism essentially raises the signal-to-noise (SNR) response of the deep-layer PFC E subunits as DA concentrations increase (Servan-Schreiber et al., 1990). The gain used for β_E is a monotonic function of the current simulated extracellular concentration of DA:

$$\beta_{\text{E}}(t) = \phi(t)(1-a)[\beta_{\text{hi}} - \beta_{\text{lo}}] + \beta_{\text{lo}} \tag{8}$$

Here, the desired low-DA gain is defined as $\beta_{lo} = 5.0$, the desired high-DA gain is defined as $\beta_{hi} = 20.0$, a is a DA antagonism parameter ranging from 0 to 1 (see

below), and $\phi(t)$ is the measure of concentration of DA at iteration t. This last function varies according to the difference equation:

$$\phi(t) = \phi(t-1) + d \cdot o_{\text{task}}(t-1) - r\phi(t-1) \quad (9)$$

Here $d = 0.40$ is a diffusion constant that determines how rapidly summed DMS and DNMS task unit output activity $o_{\text{task}}(t)$ gives rise to increased extracellular DA (Dreher & Burnod, 2002). A DA reuptake constant r determines how rapidly DA is depleted from the extracellular space.

Nine different DA conditions, as shown in Table 1, are simulated for the model. For the normal DA condition ($a = .0$, $r = .4$), values for $\phi(t)$ range between 0 and 1 (depending upon the output levels of the PFC task units), which causes β_E to vary between 5.0 and 20.0 (Figures 4E and 5). Antagonist (''hypo DA'') conditions are modeled by keeping the reuptake rate r normal and varying antagonism parameter a between 0 and 1, with $a = 1.0$ modeling maximal DA channel blocking. The effect of a is to determine the ceiling (between β_{hi} and β_{lo}) which β_E will reach during DMS/DNMS tasks. Agonist (''hyper DA'') conditions are modeled by fixing $a = .0$ and lowering the reuptake rate r to a value below 0.4. This causes the upper bound of $\phi(t)$ to rise above 1, which causes β_E to exceed β_{hi} during DMS/DNMS tasks.

Acknowledgments

O. S. was supported by NIH/NIDA grant 1R21DA15647-01. G. C. was supported by NIMH training grant T32MH019879-11.

Reprint requests should be sent to Olaf Sporns, Department of Psychological and Brain Sciences, Indiana University, Bloomington, IN 47405, or via e-mail: osporns@indiana.edu.

REFERENCES

Braver, T. S., & Barch, D. M. (2002). A theory of cognitive control, aging cognition, and neuromodulation. *Neuroscience and Behavioral Reviews, 26,* 809–817.

Cohen, J. D., Braver, T. S., & Brown, J. W. (2002). Computational perspectives on dopamine function in prefrontal cortex. *Current Opinion in Neurobiology, 12,* 223–229.

Courtney, S. M., Petit, L., Maisog, J. M., Ungerleider, L. G., & Haxby, J. V. (1998). An area specialized for spatial working memory in human frontal cortex. *Science, 279,* 1347–1351.

Curtis, C. E., & D'Esposito, M. (2003). Persistent activity in the prefrontal cortex during working memory. *Trends in Cognitive Sciences, 7,* 415–423.

Dreher, J.-C., & Burnod, Y. (2002). An integrative theory of the phasic and tonic modes of dopamine modulation in the prefrontal cortex. *Neural Networks, 15,* 583–602.

Durstewitz, D., Kelc, M., & Güntürkün, O. (1999). A neurocomputational theory of the dopaminergic modulation of working memory functions. *Journal of Neuroscience, 19,* 2807–2822.

Durstewitz, D., & Seamans, J. K. (2002). The computational role of dopamine D1 receptors in working memory. *Neural Networks, 15,* 561–572.

Frank, M. J., Loughry, B., & O'Reilly, R. C. (2001). Interactions between frontal cortex and basal ganglia in working memory: A computational model. *Cognitive, Affective, and Behavioral Neuroscience, 1,* 137–160.

Funahashi, S. (2001). Neuronal mechanisms of executive control by the prefrontal cortex. *Neuroscience Research, 39,* 147–165.

Funahashi, S., Bruce, C. J., & Goldman-Rakic, P. S. (1989). Mnemonic coding of visual space in the monkey's dorsolateral prefrontal cortex. *Journal of Neurophysiology, 61,* 331–349.

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

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

Gao, W.-J., Krimer, L. S., & Goldman-Rakic, P. S. (2001). Presynaptic regulation of recurrent excitation by D1 receptors in prefrontal circuits. *Proceedings of the National Academy of Sciences, U.S.A., 98,* 295–300.

Goldman-Rakic, P. S. (1996). Regional and cellular fractionation of working memory. *Proceedings of the National Academy of Sciences, U.S.A., 93, 13473–13480.*

Goldman-Rakic, P. S., Muly, I. E. C., & Williams, G. V. (2000). D1 receptors in prefrontal cells and circuits. *Brain Research Reviews, 31,* 295–301.

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

Gulledge, A. T., & Jaffe, D. B. (1998). Dopamine decreases the excitability of layer V pyramidal cells in the rat prefrontal cortex. *Journal of Neuroscience, 18,* 9139–9151.

Gurney, K. N., Prescott, T. J., & Redgrave, P. (1998). *The basal ganglia viewed as an action selection device*. Paper presented at the 8th International Conference on Artificial Neural Networks, Skövde, Sweden.

Gurney, K. N., Prescott, T. J., & Redgrave, P. (2001). A computational model of action selection in the basal ganglia: I. A new functional anatomy. *Biological Cybernetics, 84,* 401–410.

Hasselmo, M. E. (1995). Neuromodulation and cortical function: Modeling the physiological basis of behavior. *Behavioural Brain Research, 67,* 1–27.

Horwitz, B., & Sporns, O. (1994). Neural modeling and functional neuroimaging. *Human Brain Mapping, 1,* 269–283.

Horwitz, B., & Tagamets, M. A. (1999). Predicting human functional maps with neural net modeling. *Human Brain Mapping, 8,* 137–142.

Horwitz, B., Tagamets, M. A., & McIntosh, A. R. (1999). Neural modeling, functional brain imaging, and cognition. *Trends in Cognitive Sciences, 3,* 91–97.

Husain, F. T., Tagamets, M. A., Fromm, S. J., Braun, A. R., & Horwitz, B. (2004). Relating neuronal dynamics for auditory object processing to neuroimaging activity: A computational modeling and an fMRI study. *Neuroimage, 21,* 1701–1720.

Miller, E. K., & Cohen, J. D. (2001). An integrative theory of prefrontal cortex function. *Annual Review in Neuroscience, 24,* 167–202.

Miller, E. K., Erickson, C. A., & Desimone, R. (1996). Neural mechanisms of working memory in prefrontal cortex of the macaque. *Journal of Neuroscience, 16,* 5154–5167.

O'Reilly, R. C., & Frank, M. J. (2004). *Making working memory work: A computational model of learning in the prefrontal cortex and basal ganglia (Technical Report No. 03-03)*. Boulder, CO: Institute of Cognitive Science.

Paus, T. (2001). Primate anterior cingulate cortex: Where motor control, drive and cognitive interface. *Nature Reviews Neuroscience, 2,* 417–424.

Phillips, A. G., Ahn, S., & Floresco, S. B. (2004). Magnitude of dopamine release in medial prefrontal cortex predicts accuracy of memory on a delayed response task. *Journal of Neuroscience, 24,* 547–553.

Rainer, G., Asaad, W. F., & Miller, E. K. (1998). Memory fields of neurons in the primate prefrontal cortex. *Proceedings of the National Academy of Sciences, U.S.A., 95,* 15008–15013.

Rao, S. G., Williams, G. V., & Goldman-Rakic, P. S. (1999). Isodirectional tuning of adjacent interneurons and pyramidal cells during working memory: Evidence for microcolumnar organization in PFC. *Journal of Neurophysiology, 81,* 1903–1916.

Rossetti, Z. L., & Carboni, S. (2005). Noradrenaline and dopamine elevations in the rat prefrontal cortex in spatial working memory. *Journal of Neuroscience, 25,* 2322–2329.

Rowe, J. B., Toni, I., Josephs, O., Frackowiak, R. S. J., & Passingham, R. E. (2000). The prefrontal cortex: Response selection or maintenance within working memory? *Science, 288,* 1656–1660.

Sakai, K., Rowe, J. B., & Passingham, R. E. (2002). Active maintenance in prefrontal area 46 creates distractor-resistant memory. *Nature Neuroscience, 5,* 479–484.

Schultz, W. (2000). Multiple reward signals in the brain. *Nature Reviews Neuroscience, 1,* 199–207.

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

Smith, E. E., & Jonides, J. (1999). Storage and executive processes in the frontal lobes. *Science, 283,* 1657–1661.

Sporns, O., & Alexander, W. H. (2002). Neuromodulation and plasticity in an autonomous robot. *Neural Networks, 15,* 761–774.

Tagamets, M. A., & Horwitz, B. (1998). Integrating electrophysiological and anatomical experimental data to create a large-scale model that simulates a delayed match-to-sample human brain imaging study. *Cerebral Cortex, 8,* 310–320.

Tagamets, M. A., & Horwitz, B. (2000). A model of working memory: Bridging the gap between electrophysiology and human brain imaging. *Neural Networks, 13,* 941–952.

Takahata, R., & Moghaddam, B. (1998). Glutamatergic regulation of basal and stimulus-activated dopamine release in the prefrontal cortex. *Journal of Neurochemistry, 71,* 1443–1449.

Tononi, G., Sporns, O., & Edelman, G. M. (1992). Reentry and the problem of integrating multiple cortical areas: Simulation of dynamic integration in the visual system. *Cerebral Cortex, 2,* 310–335.

White, I. M., & Wise, S. P. (1999). Rule-dependent neuronal activity in the prefrontal cortex. *Experimental Brain Research, 126,* 315–335.

Wilson, H. R., & Cowan, J. D. (1972). Excitatory and inhibitory interactions in localized populations of model neurons. *Biophysical Journal, 12,* 1–24.

Zahrt, J., Taylor, J. R., Mathew, R. G., & Arnsten, A. F. T. (1997). Supranormal stimulation of D1 dopamine receptors in the rodent prefrontal cortex impairs spatial working memory performance. *Journal of Neuroscience, 17,* 8528–8535.