

Computational and neural mechanisms of task switching

Jeremy R. Reynolds^{a,*}, Todd S. Braver^a, Joshua W. Brown^a, Stefan Van der Stigchel^b

^aDepartment of Psychology, Washington University, Saint Louis, MO 63130-4899, USA

^bDepartment of Psychology, Vrije Universiteit, Amsterdam, The Netherlands

Available online 3 February 2006

Abstract

Switching between tasks that overlap in perceptual and response characteristics is assumed to rely upon the maintenance of task representations in prefrontal cortex (PFC). However, task-switching studies demonstrate “switch costs,” even when there is sufficient time to prepare for a new task. These costs suggest that task-switching performance reflects a complex interplay between priming and the updating and maintenance of task representations. We describe a computational model in which this interaction is made explicit and linked to the dynamics of PFC. Simulation results account for a variety of empirical phenomena and predict a double dissociation in lateral PFC that was subsequently identified.

© 2006 Elsevier B.V. All rights reserved.

Keywords: Behavior; Prefrontal cortex; Cognitive control; Functional imaging; Task switching

1. Introduction

PFC is hypothesized to subserve the active maintenance of task representations in situations that require rapidly switching between multiple demands [8]. However, one potential problem for this account is that these active maintenance processes appear to be nonoptimal: trials in which the task switches demonstrate increased response times and error rates relative to trials in which the task remains the same (i.e., “switch costs”) [1]. These switch costs persist even when sufficient preparation time is given for a new task [1,5,7]. One potential explanation for this finding is that the active maintenance process subserved by PFC is optimal, but is only engaged on a subset of trials in which the task switches. Distributional analyses of response times support this hypothesis. When there is sufficient time to prepare for a task switch, response times appear to come from two stochastic distributions: a prepared distribution of fast trials in which there are no switch costs, and an unprepared distribution of slow response times in which there are severe switch costs [5]. These data suggest that performance on fast trials is governed by a process that effectively suppresses proces-

sing of the irrelevant stimulus dimension (i.e., the one associated with the other task), and that the absence of this process on slow trials reveals the lasting effects of previous stimuli (e.g., priming), which are likely subserved by associative learning. To investigate this claim, we developed a connectionist model of task switching that linked the interactions between priming effects and controlled processing to the dynamics of the PFC.

2. Methods

The model used the connectionist architecture seen in Fig. 1. The simulations were conducted from within the LEABRA framework of the `pdp++` software. Model units had a nearly sigmoidal point neuron activation function that maintains key aspects of the electrophysiological process of firing neurons, such as different ion channel types (for details, see [9]).

Activations were calculated for each unit by clamping the inputs at the beginning of each event and allowing activity to propagate through the network until equilibrium was reached (see Fig. 1). Lateral inhibition was used in each layer, and zero-mean gaussian noise ($\sigma = 0.0004$) was added to the membrane potential of each of the units to provide variability in the behavioral measures. All weights started with an initial value of 0.5, but were subject

*Corresponding author. Tel.: 314 935 8547; fax: 314 835 8790.

E-mail address: jeremy.reynolds@colorado.edu (J.R. Reynolds).




Fig. 1. Model architecture, task, and activity dynamics. Each schematic represents the activity state and pattern of connectivity of the model at the beginning and end of the three different events composing each trial. Darker circles represent higher levels of activity. During the cue period, a cue stimulus is presented and propagates activity through the PFC to the hidden layer and the other layers of the network, activating both attributes of the relevant dimension. If a phasic DA signal occurs during the cue period (panel A), then the task information is actively maintained in the PFC over the delay period, but in the absence of such a signal (panel B), this information decays rapidly over the empty delay. During the target period, a target stimulus with one attribute from each dimension is presented, and activity propagates through the network (potentially re-activating the PFC, see panel B) until the network's activity state settles into a stable state (maximum Δ activity < 0.0003) and a response is achieved (maximum activity in the response layer > 0.65). At the end of each event, learning strengthens the connections between those units that were co-active (indicated by heavier solid lines).

to associative learning processes simulated via post-synaptically gated Hebbian learning ($\Delta w_{ij} = 0.1 y_j (x_i - w_{ij})$) on the connections projecting from the hidden layer to the PFC and the response layer. This learning rule was applied after equilibrium was reached on each event. After each trial, the weights decayed exponentially towards their original value of 0.5 ($w_{t+1} = 0.5 + 0.3(w_t - 0.5)$, where t is the current trial).

In order to simulate a random distribution of prepared and unprepared trials, a gating signal occurred on a random 50% of task cue events. This gating signal was hypothesized to reflect a phasic dopamine (DA) response [2]. Whereas previous research has considered the phasic DA response to occur on all trials after training [2], the current research capitalizes upon the well documented stochastic properties of DA activity [6]: The noise in the DA response is reflected in the probabilistic behavior of the gating signal, and it is thought to be the mechanistic source underlying the probabilistic nature of task-switch costs. The gating signal modulated intracellular currents to enable active maintenance: if a gating signal occurred, a hysteresis current was turned on such that activity could be maintained throughout the trial [4]. This actively maintained task cue was able to bias subsequent processing to produce an appropriate response [8]. If the gating signal

did not occur, there was no internal mechanism for actively maintaining PFC activity states in the absence of external input. Under this circumstance, the empty delay period indirectly flushed the activity out of the PFC units, leaving the network unprepared to process the upcoming target stimulus (see Fig. 1). Sixteen sequences of 1024 trials were performed, each initialized with a different random seed.

3. Results and discussion

In order to evaluate the model's performance, data produced by the network (performance measures as well as PFC activity dynamics) were compared to data produced by 13 participants whose brain activity was measured with functional magnetic resonance imaging (fMRI) while performing a random-cueing task-switching paradigm [3].

The network captured several general phenomena in the task-switching literature. Task switch trials were slower and less accurate than task repeats (13.0 cycles; effect size = 0.38sd; 3.6% errors; effect size = 0.22sd), with effect sizes comparable to those seen in the empirical data (74 msec; effect size = 0.29sd; 3.2% errors; effect size = 0.13). Further analyses indicated that the model was also capturing other common but subtle effects. For example, the global switch cost was mediated by a task




Fig. 2. Behavioral measures and associative learning mechanism underlying switch costs. Panel A presents the empirical (bars) and model (open circles) data demonstrating greater switch costs on response-repeat trials relative to response-alternation trials. Error bars represent 99% confidence interval (CI) of the behavioral data. Model data were transformed to be on the same scale as the empirical data by the regression equation: $BehavioralRT = 6.8ModelRT - 184$. Panel B presents one mechanism underlying switch costs: the changes in connection strength between the hidden layer and PFC as a function of previous experience. When the target stimulus appears on a task-repeat trial, there is a large difference between the weights from the hidden layer to the appropriate and inappropriate task units in PFC, due to the accumulation of consistent weight changes on both the current and previous trial. However, on task-switch trials, this difference is reduced, and this reduction results in more competition between the two PFC representations (leading to increased response times).

switch $\times$ response repetition interaction, such that switch costs were greater on trials in which the correct response repeated relative to trials in which the correct response changed [7](see Fig. 2).

Separating trials by whether they involved active maintenance in PFC illuminated several other findings. First, the performance of the model was above chance even when the model did not actively maintain the current task in PFC (5.3% errors). In particular, the model performed well above chance on trials in which the two dimensions of the target stimulus led to different responses, even though the network did not have an actively maintained signal to select the task-relevant dimension (10.6% errors). The ability to make an appropriate response in the absence of an actively maintained task was subserved by the associative learning mechanism. The weights from the hidden layer to the PFC and response layers were

strengthened on the most recent task-cue event. These weight changes biased the network to perform the appropriate task, even when an actively maintained task was not present. In particular, the change in weight space enabled the re-activation of the appropriate task representation in PFC during target presentation (see Fig. 1, panel B and Fig. 3), such that activity from the presented target stimulus would retrace the pathways that had just been strengthened during the cue and delay periods. These strengthened pathways ultimately led to increased activation values in the appropriate task representation, which could then serve to bias on-going processing.

Further analysis of trials without active maintenance revealed that there were clear differences in the model's performance between these trials and those that involved active maintenance. Trials with active maintenance had smaller switch costs than those without active maintenance (switch cost with active maintenance = 6.1 cycles, effect size = 0.18sd; switch cost without active maintenance = 20.4 cycles, effect size = 0.60sd), demonstrating that




Fig. 3. Behavior and PFC responses as a function of response speed. Panel A demonstrates increased switch costs for slower trials relative to faster trials in both human data (bars) and the model (overlaid open circles). Error bars represent 99% CI. Model data were transformed to be on the same scale as the empirical data by the regression equation: $BehavioralRT = 6.8ModelRT - 127$. Panel B demonstrates a double dissociation in a region of lateral PFC (Brodmann area 44/9; center of mass: $-46, 15, 21$; volume = 324 mm^3). The fastest trials have greater delay related responses than the slowest trials, but the slowest trials have greater target-related activity than the fastest trials. Error bars represent the 99% CI of the imaging data. Model data were transformed to be on the same scale as the empirical data by the regression equation: $ImagingSignal = 0.05ModelSignal - 0.1$.

although the associative learning mechanism enabled successful performance without active maintenance, a price was paid for this flexibility: performance was strongly influenced by previous history. Active maintenance in PFC diminished the impact of previous history by providing another source of input to the appropriate target dimensions. Without active maintenance, both target dimensions activated their respective hidden layer representations and then competed for activity. Due to prior learning, one target dimension would gain an advantage over the other by re-activating the appropriate task representation. However, because this re-activation process depended on the learning mechanism, it was sensitive to the history of prior trials, such that switching from one task to another resulted in less differentiation in the weights from the hidden layer to the PFC. This smaller difference led to the re-activation process taking longer than if the task had repeated (see Fig. 2).

The finding that the presence or absence of active maintenance in model PFC units had a large impact on switch costs suggests that the two distributions underlying residual switch costs may be characterized by their use of lateral PFC. This hypothesis led to a novel empirical prediction regarding PFC activity. Specifically, the model results suggest that there should be greater delay-related responses in PFC on trials in which there are minimal switch costs relative to trials with larger switch costs (reflecting active maintenance). Conversely, there should be greater target-related responses in PFC on trials in which there are larger switch costs relative to trials with smaller switch costs (reflecting task-set reactivation). Because one cannot independently determine which trials in an empirical data set correspond to active maintenance and which trials do not, the modeling data were re-analyzed in order to generate concrete predictions that could be tested in the imaging data. “Prepared” trials were taken to be the fastest 10% of task switch and task-repeat trials, and “unprepared” trials were taken to be the slowest 10% [5]. Re-analyzing the data in this way revealed that the model produced small switch costs on the fastest trials but large switch costs on the slowest trials (3.2 vs. 40.4 cycles, effect sizes = 0.09 vs. 1.20sd; see Fig. 3). This effect of response speed on switch costs corresponded well to that observed in the empirical data (15 vs. 140 ms, effect sizes = 0.05 vs. 0.52sd). To investigate our hypotheses regarding PFC activity, the average area under the timecourse of activity in the PFC units was computed during the delay and target periods for both “prepared” and “unprepared” trials. These values revealed a double dissociation: there was an increased response in the PFC units over the delay interval in the “prepared” trials relative to the “unprepared” trials (reflecting active maintenance of the appropriate task representation; effect size = 0.55), and there was an increased response in the PFC units after the target presentation on the “unprepared” trials relative to the “prepared” trials (reflecting re-activation of the appropriate task representation; effect size = 4.33; see Fig. 3).

The imaging study confirmed the predicted pattern of responses in lateral PFC [3]. The fastest trials displayed a greater hemodynamic response to the task cue relative to the slowest trials ($t(12) = 2.4, p < 0.05$; effect size = 0.27sd), whereas the slowest trials displayed a greater hemodynamic response to the target stimulus relative to the faster trials ($t(12) = 8.1, p < 0.001$; effect size = 1.27sd), resulting in a significant response speed $\times$ period interaction ($F(1, 12) = 44.1, p < 0.001$).

4. Conclusions

This model demonstrates that a combination of active maintenance processes and associative learning mechanisms accounts for several relevant behavioral phenomena in the task-switching literature. The model provides a mechanistic account of how PFC (and DA) may contribute to task-switching performance through the active maintenance of task sets, while also capturing a behavioral phenomenon (i.e., “residual switch costs”) that, at first glance, appears to be at odds with this hypothesis. Further, it generated a concrete prediction concerning responses in PFC that were subsequently confirmed. This model enhances our understanding of how learning mechanisms and active maintenance interact to produce complex, goal-directed behavior.

Acknowledgments

We are grateful to Andy Jones and Nicole Speer for helpful comments and suggestions. Supported by the Office of Naval Research (grant N00014-00-1-0715) awarded to TSB and a National Defense Science and Engineering Graduate Fellowship awarded to JRR.

References

- [1] D.A. Allport, E.A. Styles, S. Hsieh, Shifting intentional set: exploring the dynamic control of tasks, in: C. Umiltà, M. Moscovitch (Eds.), *Attention and Performance*, vol. XV, MIT Press, Cambridge, MA, 1994, pp. 421–452.
- [2] T.S. Braver, J.D. Cohen, 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*, vol. XVIII, MIT Press, Cambridge, 2000, pp. 713–737.
- [3] T.S. Braver, J.R. Reynolds, D.I. Donaldson, Neural mechanisms of transient and sustained cognitive control during task-switching, *Neuron* 39 (2003) 713–726.
- [4] M.J. Frank, B. Loughry, R.C. O'Reilly, Interactions between frontal cortex and basal ganglia in working memory: a computational model, *Cogn. Affect. Behav. Neurosci.* 1 (2) (2001) 137–160.
- [5] R. de Jong, E. Berendsen, R. Cools, Goal neglect and inhibitory limitations: dissociable causes of interference effects in conflict situations, *Acta Psycho.* 101 (1999) 379–394.
- [6] R. King, J.D. Barchas, B.A. Huberman, Chaotic behavior in dopamine neurodynamics, *Proc. Nat. Acad. Sci.* 81 (4) (1984) 1244–1247.
- [7] N. Meiran, Reconfiguration of processing mode prior to task performance, *J. Exp. Psycho.: Learn. Mem. Cogn.* 22 (6) (1996) 1423–1442.

- [8] E.K. Miller, J.D. Cohen, An integrative theory of prefrontal cortex function, *Annu. Rev. Neurosci.* 21 (2001) 167–202.
- [9] R.C. O'Reilly, Y. Munakata, *Computational Explorations in Cognitive Neuroscience: Understanding the Mind by Simulating the Brain*, MIT Press, Cambridge, 2000.



Jeremy R. Reynolds is a research associate in the Computational Cognitive Neuroscience Lab at the University of Colorado at Boulder. He received his Ph.D. in Psychology from Washington University in Saint Louis in 2005. His research uses converging methods to understand how the prefrontal cortex subserves controlled behavior.



Todd S. Braver is an associate professor and director of the Cognitive Control and Psychopathology laboratory at the Washington University in Saint Louis. He received his Ph.D. in Cognitive Neuroscience from Carnegie Mellon University in 1997. His research uses multiple methods to understand how humans exert control over their thoughts and behavior.



Josh W. Brown is a research scientist in the Cognitive Control and Psychopathology laboratory at the Washington University in Saint Louis. He received his Ph.D. in Cognitive and Neural Systems from Boston University in 2000. His research focuses on the neural mechanisms of cognitive control in healthy and clinical populations, using integrated computational neural modeling, cognitive psychology, functional neuroimaging, and systems neurophysiology.



Stefan Van der Stigchel received his master's degree in Cognitive Artificial Intelligence in 2003 from Utrecht University, The Netherlands. He is currently pursuing a Ph.D. in Cognitive Psychology at the Vrije Universiteit, Amsterdam. His research interests include the relation between attention and eye movement behavior and computational models of cognitive control.