

BIROn - Birkbeck Institutional Research Online

Sexton, N.J. and Cooper, Richard P. (2015) Task-set inhibition, conflict, and the n-2 repetition cost: a computational model of task switching. In: Dale, R. and Jennings, C. and Maglio, P. and Matlock, T. and Noelle, D. and Warlaumont, A. and Yoshimi, J. (eds.) Proceedings of the 37th Annual Conference of the Cognitive Science Society. Pasadena, U.S.: Cognitive Science Society.

Downloaded from: <https://eprints.bbk.ac.uk/id/eprint/12077/>

Usage Guidelines:

Please refer to usage guidelines at <https://eprints.bbk.ac.uk/policies.html> or alternatively contact lib-eprints@bbk.ac.uk.

The role of conflict in the n-2 repetition cost in task switching: a computational model

Nicholas J. Sexton

n.sexton@mail.bbk.ac.uk

Centre for Cognition, Computation and Modelling
Department of Psychological Sciences
Birkbeck, University of London,
Malet Street, London, UK, WC1E 7HX

Richard P. Cooper

r.cooper@bbk.ac.uk

Centre for Cognition, Computation and Modelling
Department of Psychological Sciences
Birkbeck, University of London,
Malet Street, London, UK, WC1E 7HX

Abstract

In task switching, the n-2 repetition cost (informally, the elevation in RT associated with performing a recently abandoned task) is an indicator of residual task-set inhibition. One suggestion is that such inhibition is triggered by conflict between task-set elements. We present a novel computational model instantiating this proposal, by adding task-conflict monitoring units to an existing, interactive activation model of task switching. The model produces the empirical pattern, n-1 switch costs and n-2 repetition costs, as an intrinsic property of its architecture, but dependent on the inhibition of task demand units by the conflict detection mechanism. In a further simulation, we make predictions about n-2 repetition costs for asymmetric tasks, and show that one functional benefit of such a conflict-based, task inhibition mechanism is to facilitate top-down control of tasks by automatically reducing cross-task interference.

Keywords: backward inhibition; conflict monitoring; interactive activation model; task inhibition; task switching

Introduction

When switching task, switch costs, in terms of increased reaction time (RT) and errors, are robustly observed when compared with successive performance of the same tasks (Kiesel et al., 2010). Much debate has focused on whether switch costs reflect the operation of executive processes specific to switch trials, such as reconfiguring the cognitive system appropriate to the new task, or facilitation and interference due to residual activation and/or inhibition from the preceding trial. Recent consensus is that switch costs reflect an interplay of both control and interference (Vandierendonck, Liefooghe, & Verbruggen, 2010).

In order to more conclusively ascertain whether task inhibition occurs, one approach is to see if a cost is associated with re-activating a recently abandoned task. Mayr and Keele (2000) devised an experiment involving three tasks (A,B,C) in which repeating a recently switched-away-from task (e.g., the final trial in the sequence ABA, henceforth *n-2 repeats*) are contrasted against tasks abandoned less recently (e.g., CBA *n-2 switches*). They hypothesized that if task-set inhibition occurs when abandoning a task, assuming that inhibition dissipates slowly, there should be a cost associated with n-2 repeats compared to n-2 switches. In fact, this is typically observed in human participants, and is taken as evidence for a cognitive task inhibition mechanism. In contrast to the (n-1) switch cost, these n-2 repetition costs have, to date, been resistant to non-inhibitory explanations.

As yet, however, there is no agreed-upon mechanistic explanation of task inhibition. In one proposal, Grange, Juvina, and Houghton (2013) presented a computational model based in the ACT-R architecture. Task activation and inhibition were simulated using a modified form of the equation used to model the activation of items in declarative memory, in which the activation initially increases (simulating the decay of inhibition), peaks, and decays. This form of task inhibition is sufficient to produce n-2 repetition costs, with its absence predicting n-2 facilitation. The authors argue that lateral inhibition between task-sets, alone, is not a sufficient mechanism to produce persistent effects lasting more than one trial. Instead, task sets self-inhibit following their execution. A limitation of this model in its current form, therefore, is that it predicts n-1 switch facilitation rather than costs. Overall, whether self-inhibition represents a viable theoretical proposal remains under debate (Koch, Gade, Schuch, & Philipp, 2010, offer a critical review).

Moreover, a second line of behavioural research suggests that task inhibition is variably recruited by conflict generated during task processing. For example, increasing conflict during various stages of task processing, including response generation (e.g., by manipulating overlap of response sets) has been found to affect n-2 repetition costs, suggesting that task inhibition may occur in response to conflicting elements of multiple task-sets (Koch et al., 2010). Any complete model of task switching should parsimoniously explain both behavioural effects (i.e. n-1 switch costs and n-2 repetition costs) and their modulation by conflict between task-sets.

This paper presents a cognitive computational model of switching between three tasks, by adding a novel task inhibition mechanism, triggered by task conflict, to an earlier model of two-task switching (Gilbert & Shallice, 2002). We present two simulations in which the model reproduces the main empirical effects, namely costs for both n-1 switches and n-2 repeats. Importantly, the model demonstrates that lateral inhibition alone is not sufficient to produce these effects. Moreover, the model makes specific predictions regarding the asymmetry of n-2 repetition costs given tasks of different difficulty. In addition, the simulations suggest that a conflict/task-set inhibition mechanism provides benefits in a multitask environment, by smoothing performance during task switches, and by shielding task processing from residual activation that can occur following highly controlled tasks.

Behavioural and computational studies of mechanisms in task switching

The basis for our model is the phenomenon of asymmetric switch costs, which has been studied and modelled in the two-task switching literature. When two tasks are of different difficulties, such as word reading and colour naming of Stroop-type stimuli, *asymmetric* switch costs are frequently reported. Although the colour naming task is slower than word reading, counterintuitively the switch cost is smaller when switching to it than to word reading (e.g., Allport, Styles, & Hsieh, 1994; Yeung & Monsell, 2003). This asymmetry has been attributed to between-task interference. For example, Allport et al. (1994) proposed that performance of the weaker task requires inhibiting the stronger task to prevent it being performed inadvertently. If residual inhibition affected processing on the next trial, on switch trials it would interfere more with switching to the stronger task than the weaker task. While an inhibition-based account is appealing, similar activation-only accounts are possible, such as the mathematical model of Yeung and Monsell (2003), in which separate task processing pathways compete in terms of activation, and are influenced by residual activation from previous trials, top-down control, and intrinsic task strength. Similar levels of task activation create interference, and thus longer response times.

Gilbert and Shallice (2002) present an interactive activation model of task switching in which switch costs have a similar explanation. Processing in two task pathways is affected by the current activation state of *task demand (TD) units*, or task representations. These units receive top-down (control) input, and unlike other units in the model, a proportion of their activation is carried over from trial to trial. Common to this class of model, units have lateral inhibitory connections to other units at the same level. Switch costs occur due to residual TD unit activation favouring the repeat task. The switch cost asymmetry occurs because activating the weaker pathway (i.e., colour naming) in the face of strong irrelevant-task (word reading) interference, takes longer to produce a response, by which time the relevant TD unit tends to be more highly active at the end of the trial, than vice versa. A portion of this activation is carried forward to the next, switch trial, where TD unit activation of the previous (now irrelevant) task causes interference in the early stages of task processing, which is therefore greater for word reading trials than colour naming trials.

From the models of Yeung and Monsell (2003) and Gilbert and Shallice (2002), we note: a) switch cost asymmetries are an intrinsic result of an interplay between task strength and cross-task interference from a residually active, alternative task. b) a dedicated, explicit task inhibition mechanism is not required. In extending these findings into the current line of research, we assume that cross-task interference is a form of task conflict. Similar to these previous explanations of switch costs, our model aims to explain a complex pattern of behavioural effects in terms of a relatively simple task inhibition




Figure 1: Architecture of the model. Excitatory connections in black (sharp arrows), inhibitory in red (circular arrows). Arrowheads show the direction of the connection. Not shown are within-module connections (e.g., lateral inhibition).

mechanism, triggered by between-task conflict/interference.

Little systematic attention has been paid to $n-2$ repetition costs for asymmetric tasks. One exception is the study of Arbutnott (2008), which examined switching between three digit judgement tasks: participants judged whether a given digit was odd or even (easy), greater or less than 5 (easy), or prime or non-prime (hard). In two experiments, involving either separate or overlapping response sets respectively, asymmetric $n-2$ repetition costs were observed, with greater costs for easy-hard-easy triplets than hard-easy-hard triplets. That is, the $n-2$ task received greater backward inhibition when it was easy than when it was hard. However, the effect on RTs was not robust, and only reached statistical significance for one pairing of tasks, and then for non-overlapping response sets only. Additionally, unexpected effects occurred, such as the reversal of direction of the switch cost asymmetry for one pairing of tasks, in both experiments (i.e., greater cost when switching to the more difficult task) — a result which it is difficult to fully predict using only a verbal model. Therefore, one useful role of modelling is the integration and explanation of these disparate effects.

A conflict-based model of task-set inhibition

Our model architecture is illustrated in figure 1. The lower portion of the figure is equivalent to the model of Gilbert and Shallice (2002) applied to three tasks. The upper level corresponds to conflict monitoring units, similar to that of Botvinick, Braver, Barch, Carter, and Cohen (2001). The input to these units is somewhat different to elsewhere in the model. Each monitors the conflict (i.e., simultaneous activation) between two Task Demand (TD) units, by taking the product of the TD activations as an input¹, multiplied by a *gain* parameter. Thus, if two TD units have activation greater

¹The range of TD activations are linearly rescaled from (-1,1) to (0,1) for this calculation only.

than zero, one conflict unit will receive a positive input. Conflict units also receive a constant, negative *bias* input, hence activation decreases in the absence of positive input. Unlike the model of Botvinick et al. (2001), conflict units bias model processing interactively, via inhibitory connections to both respective task demand units, multiplied by a *weight* parameter.² Unless otherwise specified, connection weights are fixed and take the default values used by Gilbert and Shallice (2002).

Simulations were run on blocks of three trials. On the first trial of each block, all units are initialised with zero activation. On subsequent trials, TD units, which carry over 20% of their final, previous-trial activation (as in the original model) and conflict units, which carry over 50% of their activation, modelling the effects of residual task inhibition. All other units are initialised as for the first trial. In a simulated trial, one input unit in each task pathway (representing a trivalent stimulus), and a top-down control unit (representing the currently cued task) are set to 1. Activation then iteratively propagates throughout the model. A response is made when the most active output unit exceeds that of the next most active, non-congruent output unit by a response threshold of 0.15, and the number of cycles taken for this to occur is the simulated response time (RT). In sum, the model instantiates a theoretical position similar to the proposal by Koch et al. (2010), i.e., that task inhibition is recruited by conflict generated during task processing.

General simulation methods

The model was tested using an analog of the paradigm of Arbuthnott (2008). Blocks of three tasks are classified according to the number and type of task switches, with the dependent variable being the RT of the final trial. The n-1 switch cost is the difference between 1-switch (1SW) and 0-switch blocks (0SW), in which the final trial is a task switch (e.g., AAB) or a repeat (ABB), respectively. The n-2 repetition cost is the difference between final trial RT on alternating-switch (ALT) blocks (ABA), and 2-switch (2SW) blocks (CBA). If no response is made within 500 cycles, the trial is classified as an error. RTs are only analysed from blocks with no errors.

Running the model requires a number of parameter values to be specified. In addition to those shared with the model of Gilbert and Shallice (2002), which took default values, an additional parameter controls the amount of residual conflict activation (50% for all simulations). As described above, three further parameters are required for the conflict monitoring layer: *gain*, *bias*, and *weight*. One approach to parameter setting would simply be to fit the model to the empirical data pattern. However, it might be that with an alternative set of parameters, the model could fit any arbitrary pattern of

²Given that unit activation varies between -1 and 1, only above-zero conflict unit activations are allowed to inhibit task demand units to prevent negative activation from exciting task demand units (due to the negatively weighted connection).

behaviour (Roberts & Pashler, 2000). It is important, therefore, to show what behaviour is predicted across a wide range of possible parameter values, and examine whether a specific behaviour is intrinsic to the model's theoretical content, or dependent on specific parameter values. Accordingly, we pursue a methodology similar to parameter space partitioning (Pitt, Kim, Navarro, & Myung, 2006). By varying three parameters across a wide range, dependent variables were generated and compared for each region (voxel) of a 3D grid.

Simulation 1

Simulation 1 tests the hypothesis that n-2 repetition costs are dependent on a conflict-driven task-inhibition mechanism.

Method

This simulation varied three parameters of the conflict system: *gain*, (0 to 100); *bias* (-40 to 0); and *weight* (-30 to 0). *Gain* and *bias* both affect the rate at which conflict unit activation builds up, and decays, respectively. *Weight* affects the amount of biasing that conflict units exert on TD units. A weight of zero is functionally equivalent to a model with no conflict mechanism (thus, only lateral inhibition of TD units). The effects that the task inhibition/conflict mechanism has on behaviour is assessed by comparing stronger levels of *weight* with this baseline.

Mean switch costs and n-2 repetition costs, in model cycles, were calculated for 3000 blocks for each voxel of parameter space, for each condition (0SW, 1SW, 2SW, ALT). DV's were compared for each voxel using a Welch two-samples t-test, and the resulting effect sizes (*r*) for were plotted in figure 2. The intersection of both empirical effects (figure 2 lower panel) was taking the geometric mean of both effect sizes, for voxels with both effects in the correct direction.

Results and discussion

Figure 2 (upper panel) shows RT switch costs are robustly predicted over a wide region of the model's parameter space, except for a small region in the upper right of the plot for stronger *weight* values. Here, a high gain and weak bias means that activation of conflict units increases irrespective of the degree of actual conflict. Conversely, behaviour in the bottom left of each plot (i.e., strong bias and low gain) is relatively uniform, because these settings mean conflict unit activation decreases irrespective of input, thus no biasing of model processing occurs.

N-2 repetition costs (figure 2, centre panel), in contrast, are less robust than switch costs, partly because the difference in sequences (occurring on the n-2 trial) must affect processing even after one intervening trial. Nevertheless, systematic effects did occur. The model did not produce N-2 repetition costs for near zero *weight* values, demonstrating that lateral inhibition between TD units, alone, is insufficient to produce n- 2 repetition costs. However, for stronger *weight* values, inhibition of TD units by the conflict units was sufficient to produce the effect for a contiguous region of parameter space.




Figure 2: Simulation 1, visualisation of parameter space. Horizontal panels show levels of the *weight* parameter, with zero (baseline) at far right. Upper: switch cost effect size (r) of switch cost. Central: n2-repetition cost effect size (r). Lower: geometric mean of switch cost and n-2 repetition cost effect sizes for voxels with positive costs only.

The main region of parameter space (below a top left - bottom right diagonal) predicts n-2 facilitation, rather than costs. Here, the combination of strong *bias* and low *gain* means that conflict decays too quickly, with the units insufficiently sensitive to their inputs to produce residual conflict effects. In the region above this diagonal, where *bias* is weaker and/or *gain* stronger, n-2 repetition costs are consistently predicted for non-zero *weight* values.

The two empirical phenomena, costs for n-1 switches and n-2 repetitions (figure 2 lower panel) co-occur in a well-defined region for non-zero *weight* values. Informally, this behavioural pattern is obtained with the constraints that the activation of conflict units must increase or decrease, given conflict or lack of conflict, respectively. Outside these regions, other behaviour (e.g., switch costs but n-2 repetition facilitation) may be understood either in terms of implementational failure of the model (the parameters are inappropriate for the functioning of conflict units) or in terms consistent with theory (TD unit processing must be biased by inhibitory connections from conflict units). In conclusion, the empirical pattern is a feature of the model architecture, and not of a specific set of parameter values.

Simulation 2

While simulation 1 tests whether or not the model can reproduce empirically observed phenomena, simulation 2 examines two questions. Firstly, why does a conflict/task-inhibition mechanism affects performance in this way? More specifically, do n-2 repetition costs reflect a performance improvement for n-2 switches, or an impairment of n-2 repeats, compared to a system lacking such a mechanism? Secondly,

while conflict detection might beneficially be used to regulate performance by trading off speed and accuracy (Botvinick et al., 2001), what functional advantages are provided by task inhibition triggered on this basis? Simulation 2 manipulates the between-trial conflict by using two tasks of identical, fixed, intermediate difficulty, while varying the difficulty of the third task.

Method

In the model of Gilbert and Shallice (2002), task difficulty is specified by two parameters — *stimulus input strength* (*SIS*), representing the automatic, bottom-up activation of a response by a stimulus (greater for stronger tasks), and *top-down control strength* (*TDCS*), specifying the control needed to ensure the task is performed (greater for weaker tasks). As top-down control provides a constant positive input to the cued task demand unit, a variable *TDCS* is a confounding factor, in that the same degree of task inhibition has a stronger influence on the processing of units with a low *TDCS* (i.e., easier tasks). Thus, rather than using a single weight value for all inhibitory conflict-TD connections (as in simulation 1), the *weight* parameter in simulation 2 was multiplied by *TDCS* for each task demand unit. The *bias* and *gain* parameters were fixed, at -10.0 and 75.0 respectively.

This simulation varies the *weight*, *SIS* and *TDCS* parameters to create a three-dimensional space. The task parameters (*TDCS*, *SIS*) of task A were manipulated, while B and C were left at default. For asymmetric tasks, each task sequence (e.g., 0SW) has various permutations — (e.g., ABB, BAA and BCC). Here, we considered only switches from task A (variable *SIS* and *TDCS*) to task B (fixed). Hence, 0SW sequences are all ABB, 1SW are AAB, 2SW are CAB and ALT are BAB, with only the n-1 task being of variable difficulty (except for the 0SW condition). By varying the parameters of task A, we test the effect on behaviour for both hard-easy-hard (HEH) and easy-hard-easy (EHE) switches.

By varying *SIS* and *TDCS* of task A factorially, such that either may be greater or less than that for task B, the resulting two-dimensional parameter space is divided into four quadrants. The upper-left represents the region in which task A is stronger, but less controlled, than task B, as in a stronger task (e.g., word reading). In the lower-right, A is weaker, but more controlled than task B, indicating a weaker task (e.g., colour naming). In the upper-right both the input and control strength are greater for task A, hence the task has more control than is needed to perform the task. Finally, in the lower-left a weak task is coupled with insufficient control.

Results and discussion

To determine the effect of the conflict/task-inhibition mechanism on performance, figure 3 plots switch costs (panel 1) and n-2 repetition costs (panel 2) and RTs (panels 3 to 6) relative to a baseline of *weight* = zero.

The plot of relative switch costs (figure 3, panel 1) suggests that stronger *weight* values produce smaller switch costs, especially for HEH switches (upper quadrants). The effect is

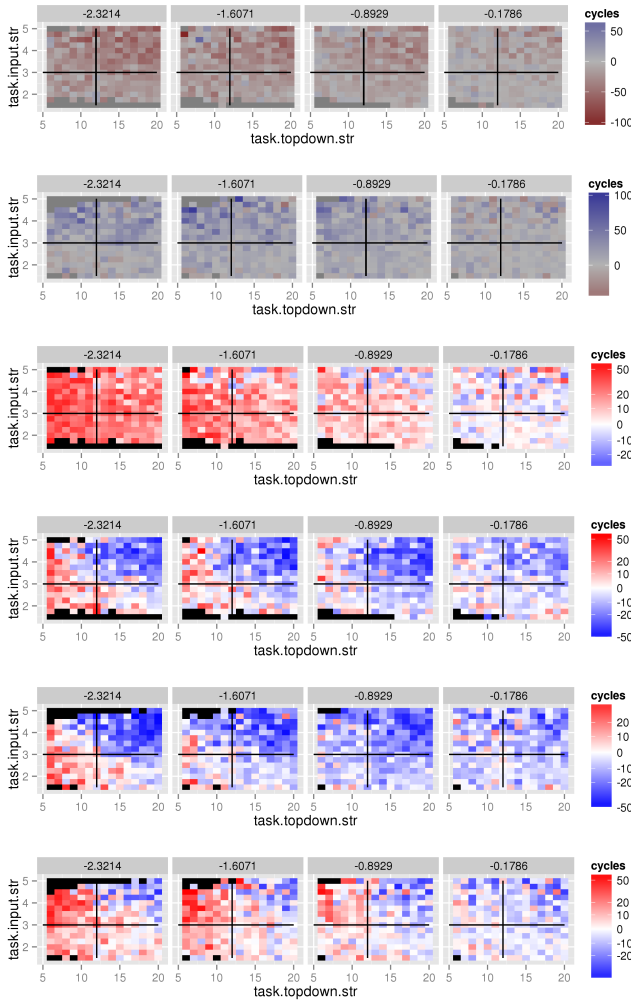


Figure 3: Simulation 2. Horizontal panels show parameter space for successive values of *weight*, from strong (left) to weak (right) biasing. All values are relative to a baseline where weight equals zero. (1) switch costs (i.e., 1SW - 0SW), (2) n-2 repetition costs (i.e., ALT - 2SW), (3 - 6) RTs for individual 0SW, 1SW, 2SW and ALT conditions respectively.

qualitatively modulated by the strength of inhibitory biasing: for the weakest *weight*, the effect is minimal. The reduction in switch costs is most pronounced when the A task is over-controlled (top right quadrant). In this case, higher levels of control on the trial preceding the switch produce greater task demand activation, leading to more residual conflict on the subsequent (switch) trial. Interestingly, this selective reduction in switch costs exaggerates the switch cost asymmetry (i.e., it reduces costs more for EH than HE switches) suggesting that in a task-switching system with such a mechanism, a component of the switch cost asymmetry may be attributable to task inhibition. In contrast, for n-2 repetition costs (figure 3, panel 2), stronger *weight* values produce larger costs. However, increased costs are also modulated by input con-

trol strength, with a greater increase in n-2 repetition cost for HEH switches (i.e., upper quadrants). To understand why, we next consider each sequence individually.

In the 0SW condition (figure 3, panel 3), intermediate or stronger values of *weight* predominantly produce longer RTs. The figure suggests topdown control strength modulates this increase — the greatest increase occurs during switches from a less controlled task of a similar difficulty (centre left). In the 1SW condition (panel 4), lower *weight* levels produce RT facilitation, particularly for HE switches (lower right). This is due to residual task inhibition helping to overcome the residual task activation which contributes to the asymmetric switch cost. At higher *weight* values, slowing occurs for switches from undercontrolled tasks (left centre), with some effect on easy-hard switches (upper left). This occurs due to conflict on the undercontrolled (n-1) trial, and thus inhibition of the non-relevant task demand unit, which becomes the relevant task demand unit on the switch trial. Taken together, the reduction in switch cost, greater for EH switches, occurs for two reasons: firstly, 1SW trial facilitation, particularly for switches from more controlled tasks (including HE switches); secondly, 0SW trial interference, particularly for switches from less controlled tasks (including EH switches). Overall, the switch cost is reduced for both HE and EH switches, but the effect is greater for EH switches, exaggerating the switch cost asymmetry. In general, weak *weight* values produce more generalised effects, with effects becoming more specific to task asymmetries for higher *weight* values.

In the 2SW condition (figure 3 panel 5), mild *weight* values produce generalised facilitation. Interestingly, this extends to switches from under-controlled tasks (bottom left). For strongest *weights*, the effect on RTs is highly modulated by top-down control, with interference and facilitation caused by low-control and high-control n-1 trials, respectively, with the effects most exaggerated for under- and over-controlled trials. That facilitation dominates irrespective of *SIS* or *TDCS* for all but the highest *weight* values, suggests one benefit of this mechanism is to reduce the amount of control required to achieve good performance when switching. In the ALT condition (panel 6), the effect is modulated by n-1 task difficulty. For low *weight* values, conflict units facilitate performance. For intermediate and stronger values, they cause interference — greatest for EH switches, but also for HE switches. In general, the interference effect is modulated by both *SIS* and *TDCS* of the n-1 task. Together, these results explain the larger n-2 repetition costs observed in HEH alternations — it is a composite of stronger facilitation for switches from easier/more controlled tasks in the 2SW condition, and greater interference when switching from those same tasks in the ALT condition.

General discussion

N-2 repetition costs are typically attributed to residual task inhibition. Here, task demand units receive both inhibitory and excitatory inputs, hence ‘task inhibition’ may be too simplis-

tic. However, consistent with the original hypothesis, conflict units effectively smooth performance in switch trials, at the cost of interference when resuming a recently abandoned task. The reduction in switch cost had a side-effect of contributing to the switch cost asymmetry, suggesting that one component of the switch cost may be due to task inhibition, modulated by inter-task conflict.

A beneficial effect of the proposed mechanism is facilitation of performance following under-controlled tasks, seen in the 2SW condition of simulation 2. This suggests that conflict units insulate switching performance against deterioration when top-down control is lower than ideal, such as in the case of distraction or divided attention. However, the trade-off is weaker performance when repeating the same task. Thus, conflict units might serve the function of an intermediate control layer — ‘dumb’ units that are unselective/uncontrolled as to the target of inhibition, but effectively facilitating performance in contexts requiring control, such as task switching. Such units might provide an automatic, low-level control layer, reserving top-down attentional biasing for the ‘heavy lifting’. Additionally, the effect of conflict units is heavily modulated by top-down control. Specifically, in switching conditions (1SW, 2SW) it protects performance following a highly controlled task (such as a simple task with a high cost of failure — imagine carrying an antique vase across a polished floor), effectively protecting subsequent tasks against distracted attention.

Three issues remain. Firstly, while the simulations explore the effect of various *weight* values, it remains an empirical question whether this parameter models something fixed or variable in a human cognitive system. Does the conflict system exert more or less biasing on task representations in different contexts? The sensitivity of the n-2 repetition cost to task parameters suggests that it may.

A second issue concerns that fact that the model predicts that in response to asymmetric task difficulties, n-2 repetition costs should be greater for hard-easy-hard switches than easy-hard-easy switches. This is the opposite direction to that observed in the only empirical data available, that of Arbuthnott (2008), although as previously noted, the switch cost asymmetry found in that study was not robustly observed. One difference between the model and that study concerns the overlap of response sets: in the model, response sets are mutually connected, that is, compatible responses are mutually excitatory while incompatible responses are inhibitory. Arbuthnott (2008) found a statistically significant n-2 repetition cost asymmetry only when response sets did not overlap, perhaps suggesting that any effect is modulated by response conflict. The status of the present model and simulations is considered a tentative hypothesis, therefore, to be empirically tested as a priority.

Finally, in theoretical terms, the model only considers conflict between task representations, as a trigger for task inhibition. However, some evidence suggests response processes have a critical role. Accommodation of these findings within

the current model would seem to require an elaboration of the model’s response processes, at least. Alternatively, these phenomena may be better explained by a model in which task inhibition is triggered by response, not task, conflict. Development of such a model, and detailed behavioural comparisons on a range of simulated experimental paradigms, elucidating the role of response or task conflict in task inhibition, is a goal for future research.

Acknowledgements

This work was supported by the Economic and Social Research Council [grant number ES/J500021/1].

References

- Allport, A., Styles, E., & Hsieh, S. (1994). Shifting intentional set - exploring the dynamic control of tasks. In C. Umiltà & M. Moscovitch (Eds.), *Attention and performance xv: Conscious and nonconscious information processing* (Vol. 15, pp. 421–452). Cambridge: MIT Press.
- Arbuthnott, K. D. (2008). Asymmetric switch cost and backward inhibition: Carryover activation and inhibition in switching between tasks of unequal difficulty. *Canadian Journal of Experimental Psychology/Revue canadienne de psychologie expérimentale*, 62(2), 91–100.
- Botvinick, M., Braver, T., Barch, D., Carter, C., & Cohen, J. (2001). Conflict monitoring and cognitive control. *Psychological Review*, 108(3), 624–652.
- Gilbert, S. J., & Shallice, T. (2002). Task switching: A PDP model. *Cognitive Psychology*, 44(3), 297–337.
- Grange, J. A., Juvina, I., & Houghton, G. (2013). On costs and benefits of n-2 repetitions in task switching: towards a behavioural marker of cognitive inhibition. *Psychological Research-Psychologische Forschung*, 77(2), 211–222.
- Kiesel, A., Steinhauser, M., Wendt, M., Falkenstein, M., Jost, K., Philipp, A. M., & Koch, I. (2010). Control and interference in task switching: A review. *Psychological Bulletin*, 136(5), 849–874.
- Koch, I., Gade, M., Schuch, S., & Philipp, A. M. (2010). The role of inhibition in task switching: A review. *Psychonomic Bulletin & Review*, 17(1), 1–14.
- Mayr, U., & Keele, S. W. (2000). Changing internal constraints on action: The role of backward inhibition. *Journal of Experimental Psychology: General*, 129(1), 4–26.
- Pitt, M. A., Kim, W., Navarro, D. J., & Myung, J. I. (2006). Global model analysis by parameter space partitioning. *Psychological Review*, 113(1), 57. (00064)
- Roberts, S., & Pashler, H. (2000). How persuasive is a good fit? a comment on theory testing. *Psychological Review*, 107(2), 358–367.
- Vandierendonck, A., Liefvooghe, B., & Verbruggen, F. (2010). Task switching: Interplay of reconfiguration and interference control. *Psychological Bulletin*, 136(4), 601–626.
- Yeung, N., & Monsell, S. (2003). Switching between tasks of unequal familiarity: The role of stimulus-attribute and response-set selection. *Journal of Experimental Psychology: Human Perception and Performance*, 29(2), 455–469.