

UC Merced

Proceedings of the Annual Meeting of the Cognitive Science Society

Title

Constraints associated with cognitive control and the stability-flexibility dilemma

Permalink

<https://escholarship.org/uc/item/7nm3c7xj>

Journal

Proceedings of the Annual Meeting of the Cognitive Science Society, 40(0)

Authors

Musslick, Sebastian

Jang, Seong J

Shvartsman, Michael

et al.

Publication Date

2018

Constraints associated with cognitive control and the stability-flexibility dilemma

Sebastian Musslick^{1,*}, Seong J. Jang¹, Michael Shvartsman¹, Amitai Shenhav², and Jonathan D. Cohen^{1,3}

¹Princeton Neuroscience Institute, Princeton University, Princeton, NJ 08544, USA.

²Department of Cognitive, Linguistic, and Psychological Sciences,
Brown Institute for Brain Science, Brown University, Providence, RI 02912, USA.

³Department of Psychology, Princeton University, Princeton, NJ 08544, USA.

*Corresponding Author: musslick@princeton.edu

Abstract

One of the most compelling characteristics of controlled processing is our limitation to exercise it. Theories of control allocation account for such limitations by assuming a cost of control that constrains how much cognitive control is allocated to a task. However, this leaves open the question of why such a cost would exist in the first place. Here, we use neural network simulations to test the hypothesis that constraints on cognitive control may reflect an optimal solution to the stability-flexibility dilemma: allocating more control to a task results in greater activation of its neural representation but also in greater persistence of this activity upon switching to a new task, yielding switch costs. We demonstrate that constraints on control impair performance of any given task but reduce performance costs associated with task switches. Critically, we show that optimal control constraints are higher in environments with a higher probability of task switches.

Keywords: cost of cognitive control; capacity constraint; neural networks; task switching

Introduction

Everyday we are confronted with tasks that require us to bias processing towards task-relevant information and actions, while avoiding processing interference from distracting tasks (e.g. writing a paper while ignoring incoming email notifications). This ability to 'focus' on a task is referred to as cognitive control and is engaged across various domains of cognition (Cohen, 2017)

Despite its tremendous utility in daily life, cognitive control is subject to fundamental processing limitations that manifest themselves in two qualitatively different ways: we appear to be constrained in both the *number* of control-demanding tasks that we can execute at the same time (Posner & Snyder, 1975; Shiffrin & Schneider, 1977), as well as in the *intensity* of control we are willing to allocate to any given task (Padmala & Pessoa, 2011; Botvinick & Braver, 2015; Shenhav, Botvinick, & Cohen, 2013; Shenhav et al., 2017). Constraints on the number of control-dependent tasks that can be executed have been attributed to the sharing of local resources (process-specific representations) required to execute the different tasks (Navon & Gopher, 1979; Meyer & Kieras, 1997; Allport, 1980; Salvucci & Taatgen, 2008; Feng, Schwemmer, Gershman, & Cohen, 2014; Musslick et al., 2016). From this perspective, constraints on multitasking can be viewed as a purpose of control – to limit processing to a single task among

ones that share representations and therefore are subject to interference (Cohen, Dunbar, & McClelland, 1990; Botvinick, Braver, Barch, Carter, & Cohen, 2001) – rather than a limitation of the control system itself. However, constraints on the *intensity* of control allocated to a single task remain less well understood. These constraints seem puzzling from a normative perspective: Why would a system refrain from allocating maximal control to a task to which it is already committed, assuming that performance scales with the intensity of control allocated? One hypothesis is that the allocation of control is associated with a cost, that subjects factor into their decisions about control allocation (Botvinick & Braver, 2015; Shenhav et al., 2013, 2017). For instance, participants respond faster and more accurately on a cognitive control task (e.g. name the ink of a color word instead of reading the word) when offered a greater reward for their performance (Krebs, Boehler, & Woldorff, 2010; Padmala & Pessoa, 2011), suggesting that they can increase the intensity of control allocated to a task if it is worth the incentive, but otherwise hold back from doing so. Recent computational modeling work has demonstrated that including such a cost can help integrate a wide range of empirical findings concerning the allocation of control (Musslick, Shenhav, Botvinick, & Cohen, 2015; Manohar et al., 2015; Lieder, Shenhav, Musslick, & Griffiths, 2018). However, the *reason* for this cost remains a mystery. If it is assumed that there is an overall "budget" of control available, then it is possible that allocating control to one task is associated with an opportunity cost with respect to others (Kurzban, Duckworth, Kable, & Myers, 2013). However, this does not explain why there is a budget in the first place. Put another way, once a commitment has been made to perform a given task (i.e., allocate cognitive control to it), and that precludes the performance of others, then the opportunity cost has already been paid, so why not allocate control maximally to the selected task? Here, we explore a potential answer to this question.

Specifically, we explore the hypothesis that constraints on control intensity (i.e., encoded as cost) reflect, at least in part, an optimal solution to the stability-flexibility dilemma. This dilemma arises from a tension between allocating control maximally to a task currently being performed (to minimize distraction and optimize performance), and the ability to quickly and flexibly reconfigure the system to perform a

different task when the environment changes. This is evident empirically in the form of costs to performance when switching from one task to another, which are magnified with increases in the allocation of control to the initial task (Goschke, 2000). The dilemma is consistent with the task-set inertia hypothesis, according to which the task-set of a previously performed task persists and interferes with initial performance of a subsequent task following a switch (Allport, Styles, & Hsieh, 1994).

We use a recurrent neural network model of task performance and control to explore how different choices of a global control parameter (gain modulation) that determines its maximal intensity, influence the stability and flexibility of performance in a task switching environment. Critically, we determine the optimal value of this parameter as a function of the demand for flexibility in the task environment, and show how these results can explain differences in human task switch costs as a function of task switch probability. Finally, we conclude with a discussion about how computational dilemmas such as the stability-flexibility tradeoff may help provide a normative account of previously unexplained – and what may otherwise appear to be irrational – constraints on cognitive control. The code for all simulations used in this work is available at github.com/musslick/CogSci-2018b.

Recurrent Neural Network Model

To explore the effect of constraints on the intensity of control, we simulate control configurations as activity states of processing units in a neural network (control module) that unfold over the course of trials. Within each trial, the processing units of the network engage an evidence accumulation process that integrates information about the stimulus, and is used to generate a response (decision module). In this section we describe the processing dynamics for both the control and decision modules, as well as the environments in which the model is tasked to perform.

Control Module

We simulate the intensities of two different control signals as activities of two processing units indexed by $i, i \in \{1, 2\}$ in a recurrent neural network. The activity of each processing unit i at a given trial T represents the intensity of the control signal for one of two tasks and is determined by its net input

$$net_i^T = w_{i,i}act_i^{T-1} + w_{i,j}act_j^{T-1} + I_i \quad (1)$$

that is a linear combination of the unit's own activity at the previous trial act_i^{T-1} multiplied by the self-recurrent weight $w_{i,i}$, the activity of the other unit $j \in \{1, 2\}, j \neq i$ at the previous trial act_j^{T-1} multiplied by an inhibitory weight $w_{i,j}$, as well as an external input I_i provided to the unit (see Figure 1). The self-recurrent and mutually-inhibitory weights induce attractors within the control module, such that it can maintain its activity over time in one state or the other, but not both. The latter implements a capacity constraint on control with

regard to the number of control-dependent tasks the network can support. These weights also determine the activity of the units in each attractor state, that is also regulated by a gain parameter that we use to implement the intensity constraint, as discussed further below. The external input acts as a gating signal to the corresponding control unit (Braver & Cohen, 1999) and is set to 1 if the task represented by the control unit needs to be performed and set to 0 otherwise. The net input of each unit is averaged over time

$$\overline{net}_i^T = \tau \cdot (net_i^T) + (1 - \tau)\overline{net}_i^{T-1} \quad (2)$$

where $\overline{net}_i^{T-1}$ corresponds to the time averaged net input at the previous trial and $\tau \in \mathbb{R} : 0 \leq \tau \leq 1$ is the rate constant. A higher τ leads to a faster change in activation for both units. Finally, the activity of each unit is a sigmoid function of its time averaged net input

$$act_i^T = \frac{1}{1 + e^{-g \cdot \overline{net}_i^T}} \quad (3)$$

where g is a gain parameter that regulates the slope of the activation function. The sigmoid activation function constrains the activity of both units to lie between 0 and 1. The gain of the activation function effectively regulates the distance between the two control states, with lower gain leading to a lower activation of the currently relevant control unit. In this model, we use g to implement a constraint on the intensity of control, and explore its effect on the network's balance between stability and flexibility.




Figure 1: Recurrent neural network model used in simulations. Each of the two processing units in the control module (blue) receive an external input signal I_1, I_2 that indicates the currently relevant task. The dynamics of the network unfold over the course of trials and are determined by recurrent connectivity $w_{1,1}, w_{2,2}$ for each unit, as well as mutual inhibition $w_{1,2}, w_{2,1}$ between units. The activity of each control unit biases the processing of a corresponding stimulus dimension on a given trial. On each trial, the decision module accumulates evidence for both stimulus dimensions towards one of two responses until a threshold is reached.

Decision Module

On each trial the decision module integrates information along two stimulus dimensions S_1 and S_2 of a single stimulus to determine a response. Each dimension (e.g., color or

shape) can take one of two values (e.g., red or green; round or square), each of which is associated with one of two responses (e.g. pressing left or right button). Each of the two tasks requires mapping the current value of one of the two stimulus dimensions to its corresponding response, while ignoring the other dimension. Since both tasks involve the same pair of responses, stimuli can be congruent (stimulus values in both dimensions associated with the same response) or incongruent (associated with different responses). We simulate the response integration process using a drift diffusion model (DDM, Ratcliff, 1978), in which the drift is determined by the combined stimulus information from each dimension, weighted by input received from the control module (as described below), and evidence is accumulated over time until one of two response thresholds is reached. The drift rate is decomposed into an automatic and controlled component

$$drift = \underbrace{(S_1 + S_2)}_{\text{automatic}} + \underbrace{act_1^T S_1 + act_2^T S_2}_{\text{controlled}} \quad (4)$$

where the automatic component reflects automatic processing of each stimulus dimension that is unaffected by control. The absolute magnitude of S_1, S_2 depends on the strength of the association of each stimulus with a given response and its sign depends on the response (e.g. $S_1 < 0$ if the associated response is to press the left button, $S_1 > 0$ if the associated response is to press the right button). Thus, for congruent trials S_1 and S_2 have the same sign, and the opposite sign for incongruent (conflict) trials. For the simulations described below, the strength of the associations was equal along the two stimulus dimensions. The controlled component of the drift rate is the sum of the two stimulus values, each weighted by the activation of the corresponding control unit. Thus, each unit in the control module biases processing towards one of the stimulus dimensions, similar to other computational models of cognitive control (e.g. Cohen et al., 1990; Mante, Sussillo, Shenoy, & Newsome, 2013; Musslick et al., 2015). As a result, progressively greater activation of a control unit improves performance – speeds responses and improves accuracy – for the corresponding task. Mean reaction times (RTs) and error rates for a given parameterization of drift rate at trial T are derived from an analytical solution to the DDM (Bogacz, Brown, Moehlis, Holmes, & Cohen, 2006).

Task Environment and Processing Dynamics

We used the model to simulate performance while switching between 100 mini-blocks of the two tasks (Meiran, 1996). Each mini-block consisted of six trials of the same task. A task cue presented before and throughout each mini-block instructed the model to adjust control signals to the currently relevant task i , by setting I_i to 1 for the task-relevant control unit and $I_{j \neq i}$ to 0 for the other. On each trial, a stimulus was presented comprised of a value along each dimension S_1 and S_2 , the decision module integrated the input, and generated a response that was deemed correct if it corresponded to the

one associated with the stimulus value along the dimension indicated by the task cue.

Effects of Network Gain on Stability and Flexibility

The intensity of the control signals is functionally constrained by the gain parameter of the activation function in Equation 3: lowering gain lowers the activity of a control unit for a given positive net input, and thus constrains the maximum signal intensity of the task-relevant control unit. Here, we examine how manipulations of gain influence the model's performance, and in particular measures of stability and flexibility in the task switching design described above.

To do so we varied g from 0.1 to 3 in steps of 0.1. The control module of each model was parameterized with balanced recurrent and inhibitory weights, $w_{i,i} = 1, w_{i,j} = -1$ and a rate constant of $\tau = 0.9$. The decision module (DDM) was parameterized¹ with a threshold of $z = 0.0475$, a non-decision time of $T_0 = 0.2$ and a noise of $c = 0.04$. We simulated performance of each model on 10 different randomly permuted task switching sequences of the type described above. Each sequence was generated with a 50% task switch rate (i.e. one half of the mini-blocks required to switch to another task with respect to the previous mini-block whereas the other half of the mini-blocks required to repeat the previous task). Task transitions were counterbalanced with each task and response congruency conditions.

For each simulation, we assessed the mean difference in reaction times (RTs) and error rates between incongruent and congruent trials as a measure of cognitive stability. Incongruent trials typically lead to slower reaction times and higher error rates than congruent trials due to response conflict, referred to as the incongruency effect. The stability-flexibility dilemma suggests that increased control intensity (implemented here by higher values of g) should augment sensitivity to the task-relevant stimulus dimensions and diminish it to the task irrelevant dimension, thereby reducing the incongruency effect (Goschke, 2000). We also measured mean task performance as an index of task stability. Finally, we assessed the flexibility in terms of performance costs associated with a task switch relative to a task repetition from one mini-block to another. Specifically, we computed the performance difference between the first trial of switch vs. repetition mini-blocks. We predicted that increasing g would increase task switch costs, as this would increase the distance between the attractors in the control module and thus make it harder to switch between them. We evaluated the effect of gain on each of these measures by regressing each against the gain of the network across all task sequences and networks.

The results in Figure 2a indicate that all tested models

¹The parameter values were chosen to yield reasonable model performance, i.e. an average task accuracy higher than chance. While each parameter has a quantitative effect on the results described below, the qualitative (direction) of the effects remains robust across a wide range of parameter values (for details, see github.com/musslick/CogSci-2018b).




Figure 2: Effects of network gain. (a) Overall mean performance, mean incongruity effects and mean switch costs for both reaction time and error rate are shown as a function of network gain. Error bars indicate the standard error of the mean across different task sequences for each network. (b) Activation trajectory for models with different gain is shown as a series of connected black dots from the control attractor for task 1 (red) to the control attractor for task 2 (blue). Contour lines and arrows indicate the energy and shape of the attractor landscape after a task switch from task 1 to task 2.

showed an incongruity effect. Moreover all models exhibited switch costs with respect to both RTs and error rates. More interestingly, higher values of gain lead to increases in cognitive stability as reflected in a lower incongruity effect for both RTs, $b = -0.3508$, $t(289) = -49.93$, $p < 10^{-143}$ and error rates, $b = -0.1194$, $t(289) = -24.60$, $p < 10^{-72}$. Models with higher gain also exhibit overall faster reaction times, $b = -0.1742$, $t(289) = -47.68$, $p < 10^{-138}$, and lower error rates, $b = -0.0597$, $t(289) = -24.60$, $p < 10^{-72}$, for a given task. Conversely, increases in gain lead to higher costs of switching between tasks for both RTs, $b = 0.1866$, $t(289) = 67.19$, $p < 10^{-177}$, and error rates, $b = 0.1014$, $t(289) = 251.21$, $p < 10^{-140}$. To investigate these effects in more detail, we plotted the change in activity after a task switch for two models with different gains ($g_1 = 1, g_2 = 3$). Figure 2b illustrates that the control states for both tasks are closer together for the model with lower gain, reflecting overall lower control signal intensities for both control units. However, this also shortened the trajectory of activity from one control state to the other, thus requiring less time steps to traverse. The opposite effects were observed for higher gain. Together these effects illustrate that lower values of g functionally constrain the amount of control, reducing stability and overall performance, but affording greater flexibility.

Optimal Network Gain as a Function of Flexibility Demand

The findings above suggest that there may be an optimal value of gain for a given task environment, depending on the degree of flexibility it requires. We examined this directly, by

varying the probability of task switches across experiment sequences from 10% to 90% in steps of 20%, while counterbalancing tasks, task transitions and congruency conditions within each sequence. We used the same values for all model parameters as in the previous section. The network was tested on 10 different task sequences at a given level of switch frequency. For each sequence, we optimized the gain parameter of the network so to maximize its accuracy across all trials in the sequence. We assessed the mean activity of the control units associated with optimal gain at each level of task switch frequency, as a measure of constraints on control intensity. We measured average performance, incongruity effects and switch costs associated with the optimal values of gain at each task switch frequency, to test whether this could account for effects observed in human performance: Sequences with a higher frequency of task switching typically produce impaired overall performance but smaller task switch costs (Mayr, 2006; Monsell & Mizon, 2006). All measurements were linearly regressed against the probability of a task switch in the sequence.




Figure 3: Gain optimization for experiments with different task switch probabilities. (a) The optimal gain is plotted as a function of task switch probability in the experiment sequence. (b) The average intensity of the control unit dedicated to the currently relevant task (orange), as well as the intensity of the control unit associated with the irrelevant task (blue) are shown as a function of task switch probability in a given experiment sequence.

Simulation results are shown in Figure 3. The optimal gain of the network decreases with the frequency of task switches, $b = -0.4520$, $t(49) = -29.59$, $p < 10^{-31}$. That is, the model achieves an overall higher accuracy on sequences with high switch rate if it imposes a higher constraint on the amount of control it can allocate to a single task in exchange for the benefit of improved performance of task switches. The decrease in gain is reflected in the model's control allocation policy: The higher the probability of a task switch, the lower the average control unit activity allocated to the relevant task, $b = -0.0376$, $t(49) = -32.94$, $p < 10^{-33}$, and the more control unit activity is allocated to the irrelevant task, $b = -0.0890$, $t(49) = -29.58$, $p < 10^{-31}$. That is, it is better to bring the control attractors closer together when the tasks switch frequently. Note that this occurs at the cost of reduced stability: Models optimized for higher switch rates exhibit higher incongruity effects (RT, $b = 0.0974$, $t(49) = 31.87$, $p < 10^{-34}$; and error rate, $b = 0.0084$, $t(49) = 30.92$, $p < 10^{-33}$). Moreover, overall RTs, $b = 0.1409$, $t(49) = 128.87$, $p < 10^{-61}$, and error rates, $b = 0.0322$, $t(49) = 138.81$,

$p < 10^{-64}$, increase with the frequency of switching. However, task switch costs decrease as a function of switch probability for RT, $b = -0.0865$, $t(49) = -31.22$, $p < 10^{-33}$ and error rates, $b = -0.0595$, $t(49) = -36.96$, $p < 10^{-36}$, as a consequence of lower gain. The latter results are qualitative replications of empirical observations made by Mayr (2006), as well as Monsell and Mizon (2006).

General Discussion and Conclusion

Cognitive control allows us to flexibly reconfigure processing in accord with current task goals. However, controlled processing is subject to fundamental limitations (Posner & Snyder, 1975; Shiffrin & Schneider, 1977). Understanding these limitations is critical for understanding human processing and its failure. One limitation is the intensity of control people are able (or willing) to allocate to a given task, recently described in terms of the cost of control (Shenhav et al., 2013, 2017). Here, we described neural network simulations that suggest that this constraint may reflect a normative solution to a fundamental tradeoff between the intensity of control given to a single task (cognitive stability) and the ability to switch quickly between tasks (cognitive flexibility). We demonstrated that a meta-control parameter, the gain of the network's activation function, can regulate this tradeoff in a task switching paradigm. Specifically, while lower gain (higher constraints on control) degraded performance on each task (lower stability), it reduced performance costs associated with task switches (higher flexibility). We showed that the optimal level of gain (constraint on control) varied as a function of the demand for flexibility in the task environment, and that this pattern qualitatively matches observations of human performance made under similar conditions.

These findings provide the first normative account for constraints on cognitive control from the perspective of the stability-flexibility dilemma. They suggest why people may take account of costs when deciding how much control to allocate to a given task (Botvinick & Braver, 2015; Shenhav et al., 2013, 2017). The simulation results described in this paper may also provide an explanation for behavioral effects of manipulations in task switch probability observed by Mayr (2006); Monsell and Mizon (2006), namely that higher switch rates lead to an overall decrease in performance but also lower switch costs. The present work suggests that this may reflect adaptation to higher frequency of task switches by decreasing the amount of control allocated to a task. The model also predicts that higher task switch frequency should produce an increase in incongruency effects, a prediction that can be tested in future empirical work.

The conclusions of this work are limited to the model assumptions and parameter ranges considered. Also, we limited our analyses to the gain of the network's activation function and its effect on the stability-flexibility tradeoff. However, previous computational work has identified alternative mechanisms that mediate the balance between stability and flexibility, including dopaminergic (Braver & Cohen, 1999; Cools,

2015) and acetylcholinergic (Liljenström, 2003) neuromodulation, as well as GABA channel conductance (Ueltzhöffer, Armbruster-Genç, & Fiebach, 2015), all of which are assumed to regulate the excitability of processing units in recurrent neural networks. An important step for future work is to test the generality of the stability-flexibility tradeoff across different model architectures, as well as different mechanisms within a given architecture. The latter may involve inhibitory bias units that selectively suppress a corresponding control unit, rather than a global modulation parameter. Such selective constraints may help to explain limitations of cognitive control that are task specific. Finally, future work will have to test whether the constraints on control allocation are also optimal with respect to alternative objective functions, such as reward rate maximization (Simen et al., 2009).

The presence of a stability-flexibility tradeoff often relies on the implicit assumption that tasks cannot be executed in parallel. Interestingly, computational work suggests that such parallel processing limitations can result from another computational dilemma that neural networks face: the tradeoff between learning efficiency that is promoted by the use of shared task representations, and interference-free parallel processing that requires the separation of task representations (Musslick et al., 2016, 2017). From this perspective, parallel processing limitations reflect a preference for learning efficiency that is associated with shared representations. The stability-flexibility tradeoff arises from the enforcement of such parallel processing limitations – that is, the serial execution of tasks – and thereby the value of switching quickly between them. The study of such computational dilemmas in neural systems holds promise to uncover normative explanations for the seemingly irrational constraints on cognitive control, as well as human cognition in general.

Acknowledgement

This project was made possible through the support of a grant from the John Templeton Foundation. The opinions expressed in this publication are those of the authors and do not necessarily reflect the views of the John Templeton Foundation. We would also like to thank Nathaniel Daw and the PDP meeting group at Princeton University for helpful discussions.

References

- Allport, D. A. (1980). Attention and performance. *Cognitive psychology: New directions*, 1, 12–153.
- Allport, D. A., Styles, E. A., & Hsieh, S. (1994). Shifting intentional set: Exploring the dynamic control of tasks.
- Bogacz, R., Brown, E., Moehlis, J., Holmes, P., & Cohen, J. D. (2006). The physics of optimal decision making: a formal analysis of models of performance in two-alternative forced-choice tasks. *Psychological review*, 113(4), 700.
- Botvinick, & Braver, T. (2015). Motivation and cognitive control: from behavior to neural mechanism. *Annual Review of Psychology*, 66.

- Botvinick, Braver, T. S., Barch, D. M., Carter, C. S., & Cohen, J. D. (2001). Conflict monitoring and cognitive control. *Psychological Review*, 108(3), 624.
- Braver, T. S., & Cohen, J. D. (1999). Dopamine, cognitive control, and schizophrenia: the gating model. In *Progress in brain research* (Vol. 121, pp. 327–349). Elsevier.
- Cohen, J. D. (2017). Cognitive control. *The Wiley Handbook of Cognitive Control*, 1–28.
- Cohen, J. D., Dunbar, K., & McClelland, J. L. (1990). On the control of automatic processes: a parallel distributed processing account of the stroop effect. *Psychological Review*, 97(3), 332–361.
- Cools, R. (2015). The cost of dopamine for dynamic cognitive control. *Current Opinion in Behavioral Sciences*, 4, 152–159.
- Feng, S. F., Schwemmer, M., Gershman, S. J., & Cohen, J. D. (2014). Multitasking vs. multiplexing: toward a normative account of limitations in the simultaneous execution of control-demanding behaviors. *Cognitive, Affective, & Behavioral Neuroscience*, 14(1), 129–146.
- Goschke, T. (2000). Intentional reconfiguration and involuntary persistence in task set switching. *Control of cognitive processes: Attention and performance XVIII*, 18, 331.
- Krebs, R. M., Boehler, C. N., & Woldorff, M. G. (2010). The influence of reward associations on conflict processing in the stroop task. *Cognition*, 117(3), 341–347.
- Kurzban, R., Duckworth, A., Kable, J. W., & Myers, J. (2013). An opportunity cost model of subjective effort and task performance. *The Behavioral and Brain Sciences*, 36(6), 661–679.
- Lieder, F., Shenhav, A., Musslick, S., & Griffiths, T. L. (2018). Rational metareasoning and the plasticity of cognitive control. *PLoS computational biology*, 14(4), e1006043.
- Liljenström, H. (2003). Neural stability and flexibility: a computational approach. *Neuropsychopharmacology*, 28(S1), S64.
- Manohar, S. G., Chong, T. T.-J., Apps, M. A., Batla, A., Stamelou, M., Jarman, P. R., ... Husain, M. (2015). Reward pays the cost of noise reduction in motor and cognitive control. *Current Biology*, 25(13), 1707–1716.
- Mante, V., Sussillo, D., Shenoy, K. V., & Newsome, W. T. (2013). Context-dependent computation by recurrent dynamics in prefrontal cortex. *nature*, 503(7474), 78.
- Mayr, U. (2006). What matters in the cued task-switching paradigm: Tasks or cues? *Psychonomic Bulletin & Review*, 13(5), 794–799.
- Meiran, N. (1996). Reconfiguration of processing mode prior to task performance. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 22(6), 1423.
- Meyer, D. E., & Kieras, D. E. (1997). A computational theory of executive cognitive processes and multiple-task performance: Part I. Basic mechanisms. *Psychological review*, 104(1), 3.
- Monsell, S., & Mizon, G. A. (2006). Can the task-cuing paradigm measure an endogenous task-set reconfiguration process? *Journal of Experimental Psychology: Human Perception and Performance*, 32(3), 493.
- Musslick, S., Dey, B., Özcimder, K., Patwary, M. M. A., Willke, T. L., & Cohen, J. D. (2016). Controlled vs. automatic processing: A graph-theoretic approach to the analysis of serial vs. parallel processing in neural network architectures. In *Proceedings of the 38th annual conference of the Cognitive Science Society* (pp. 1547–1552). Philadelphia, PA.
- Musslick, S., Saxe, A., Özcimder, K., Dey, B., Henselman, G., & Cohen, J. D. (2017). Multitasking capability versus learning efficiency in neural network architectures. In *Proceedings of the 39th Annual Meeting of the Cognitive Science Society* (pp. 829–834). London, UK.
- Musslick, S., Shenhav, A., Botvinick, M. M., & Cohen, J. D. (2015). A computational model of control allocation based on the expected value of control. In *The 2nd multidisciplinary conference on reinforcement learning and decision making*.
- Navon, D., & Gopher, D. (1979). On the economy of the human-processing system. *Psychological Review*, 86(3), 214.
- Padmala, S., & Pessoa, L. (2011). Reward reduces conflict by enhancing attentional control and biasing visual cortical processing. *Journal of cognitive neuroscience*, 23(11), 3419–3432.
- Posner, M., & Snyder, C. (1975). attention and cognitive control. In *Information processing and cognition: The Loyola symposium* (pp. 55–85).
- Ratcliff, R. (1978). A theory of memory retrieval. *Psychological review*, 85(2), 59.
- Salvucci, D. D., & Taatgen, N. A. (2008). Threaded cognition: an integrated theory of concurrent multitasking. *Psychological Review*, 115(1), 101.
- Shenhav, A., Botvinick, M. M., & Cohen, J. D. (2013). The expected value of control: an integrative theory of anterior cingulate cortex function. *Neuron*, 79(2), 217–240.
- Shenhav, A., Musslick, S., Lieder, F., Kool, W., Griffiths, T. L., Cohen, J. D., & Botvinick, M. M. (2017). Toward a rational and mechanistic account of mental effort. *Annual review of neuroscience*, 40, 99–124.
- Shiffrin, R. M., & Schneider, W. (1977). Controlled and automatic human information processing: II. Perceptual learning, automatic attending and a general theory. *Psychological Review*, 84(2), 127.
- Simen, P., Contreras, D., Buck, C., Hu, P., Holmes, P., & Cohen, J. D. (2009). Reward rate optimization in two-alternative decision making: empirical tests of theoretical predictions. *Journal of Experimental Psychology: Human Perception and Performance*, 35(6), 1865.
- Ueltzhöffer, K., Armbruster-Genç, D. J., & Fiebach, C. J. (2015). Stochastic dynamics underlying cognitive stability and flexibility. *PLoS computational biology*, 11(6), e1004331.