Switch costs in inhibitory control and goal-directed behavior: A computational study of the antisaccade task

Eduardo A. Aponte[1,], Klaas E. Stephan[1,2,3], Jakob Heinzle[1]*

[1] Translational Neuromodeling Unit (TNU), Institute for Biomedical Engineering, University of Zurich and ETH Zurich. Wilfriedstrasse 6, 8004, Zurich, Switzerland.

[2] Wellcome Trust Centre for Neuroimaging, University College London. 12 Queen Square London WC1N 3BG,.

[3] Max Planck Institute for Metabolism Research. Gleueler Strasse 50, 50931, Cologne, Germany.

[*]Corresponding author:

Eduardo A. Aponte; aponte@biomed.ee.ethz.ch

Running title: Switch costs in inhibition and voluntary behavior

EAA: Designed the experiment, analysed the data, developed analytical tools and wrote the manuscript.

KES: Provided the funding, designed the experiment, and edited the manuscript.

JH: Design the experiment and wrote the manuscript.

Number of pages: 35

Number of figures: 8

Number of tables: 2

Number of words Abstract: 232

Number of words Introduction: 641

Number of words Discussion: 1486

Conflict of interest: None.

Acknowledgement: This work was supported by the René und Susanne Braginsky Foundation and the University of Zurich.

Abstract

When instructed to alternate between different tasks, humans require more time and are less accurate than when they repeat the same task. This decrease in performance is referred to as the 'task switch cost'. Several behavioral studies have investigated this phenomenon in the mixed antisaccade task, in which participants are cued to saccade either in the same or in the opposite direction of a peripheral stimulus. These reports have provided conflicting answers as to whether a cue to saccade away from a target leads to positive, negative, or no switch costs. These contradictory findings are paralleled by two opposing theoretical hypotheses which focus on *task-set inertia* and *oculomotor inhibition*, respectively. Here, we applied a computational, generative model to data from a mixed antisaccade task in which the peripheral visual stimulus also served as a trial type cue. Behaviorally, we found reaction time and error rate switch costs on pro- and antisaccade trials. Modeling revealed that these costs were due to two different effects. First, antisaccade trials enhanced inhibitory control on the following trial, resulting in fewer but faster inhibition failures. Second, goal-directed actions displayed a task-inertia effect, in that rule these were slower on switch trials compared to repeat trials. Our results shed light on the general phenomena of task switching, by demonstrating two types of switch cost that affect differently the inhibition of habitual responses and the initiation of goal-directed actions.

Significance statement

A pervasive phenomenon in human cognition is that task switching decreases performance, compared to repeating the same task. This phenomenon can be investigated in the mixed antisaccade task, in which subjects are required to alternate between saccades in the same and opposite direction of a visual target. Here, we used a novel computational model to disentangle how task-switching affects the inhibition of habitual actions -a saccade towards a cue- and the initiation of goal-directed behavior -a saccade in the opposite direction. Our model demonstrates that opposite goal-directed actions interfere with each other. However, stopping a habitual action leads to increased inhibitory control. This finding could shed light on the cost of switching in other domains.

Introduction

One hallmark of higher-order cognition is the ability to alternate between habitual and non-habitual, goal-directed actions (Isoda and Hikosaka, 2008). However, alternating between different tasks engenders costs in terms of reaction time (RT) and error rate (ER), and switching from a non-habitual to a habitual response sometimes leads to larger costs than the opposite transition (Allport et al., 1994). An attractive paradigm to study these phenomena in the oculomotor domain is the antisaccade task (Hallett, 1978; Munoz and Everling, 2004), in which a habitual response – a prosaccade towards a salient peripheral stimulus – needs to be overwritten by a non-habitual action, i.e., an antisaccade in the opposite direction. Importantly, this paradigm has received much attention in psychiatric research, because changes in ER and RT constitute a stable finding in schizophrenia (Greenwood et al., 2007; Reilly et al., 2014; Radant et al., 2015; Myles et al., 2017).

Behaviorally, switch costs in the antisaccade task have been investigated in great detail (Barton et al., 2002; Cherkasova et al., 2002; Manoach et al., 2002; Bojko et al., 2004; Fecteau et al., 2004; Barton et al., 2006a; 2006b; Rivaud-Pechoux et al., 2007; Ansari et al., 2008; Mueller et al., 2009; Lee et al., 2011; Weiler and Heath, 2012a; 2012b; DeSimone et al., 2014; Weiler and Heath, 2014; Weiler et al., 2014; Heath et al., 2015; Pierce et al., 2015; Weiler et al., 2015; Heath et al., 2016; Chan et al., 2017). Despite the large number of studies, no unified picture of the cost of switching in this paradigm has emerged. Concretely, all studies we are aware of have reported that switch prosaccades, i.e. correct prosaccades that follow an antisaccade trial, have higher latencies than repeat trials, whereas the costs associated with switch antisaccades are less clear. For example, some studies have indicated that switch antisaccades display lower RTs (e.g. Cherkasova et al., 2002), while others have reported both lower and higher RTs (e.g., Barton et al., 2006a), and yet others indicate no switch costs (e.g., Weiler and Heath, 2012b).

One path to clarify the relationship between seemingly contradictory experimental evidence is the application of generative models to empirical data (Monsell, 2003; Karayanidis et al., 2010; Heinzle et al., 2016) which might help

disentangle the underlying mechanisms that explain switch costs. In this direction, we recently developed the Stochastic Early Reaction, Inhibition and late Action (SERIA) model (Aponte et al., 2017) of the antisaccade task. In essence, SERIA combines the 'horse-race' model usually applied in the context of the countermanding saccade task (Logan et al., 1984; Camalier et al., 2007) to explain the inhibition of habitual, fast prosaccades, with a second race between two voluntary, or goal-directed actions that can generate both pro- and antisaccades. In contrast to previous models (Noorani and Carpenter, 2013), we acknowledge the possibility that prosaccades can be generated not only as the result of reactive (habitual) saccades, but also by a rule-guided decision process.

In two recent studies (Aponte et al., 2017; 2018), we showed that SERIA can be used to account for RT and ER distributions with great accuracy, and that it is able to predict corrective antisaccades. Here, we employ the model to investigate switch costs and whether they could be attributed either to the inhibition of habitual responses or to the generation of voluntary saccades. In other words, we compare the two dominant theories of antisaccade switch costs – the task-set inertia hypothesis (Allport et al., 1994) and the oculomotor inhibition hypothesis (Barton et al., 2006a; Weiler and Heath, 2014)– by a modeling approach that inherently separates these two processes. Our results demonstrate that switch costs are largely explained by the interference between voluntary pro- and antisaccades, irrespective of trial type. In addition, antisaccade trials lead to increased inhibitory control, which reduces the number of reactive saccades following an antisaccade trial, leading to an apparent switch cost on prosaccade trials.

Methods

In this study, we analyzed a subset (one of two tasks) of the data presented in detail in Aponte et al. (2018). In the following, we briefly summarize the experimental procedures relevant for the analysis presented here. This study was approved by the ethics board of the Canton of Zurich, Switzerland (KEK-ZH-Nr.2014-0246) and was conducted according to the Declaration of Helsinki.

Participants

We used the data of all twenty-four, healthy male subjects that were analyzed in Aponte et al., 2018. All subjects had normal or corrected to normal vision and provided a written inform consent to participate in the study.

Apparatus

The experiment was conducted in a dimly illuminated room. Subjects sat 60cm in front of a computer screen (41.4x30cm; Philips 20B40; refresh rate 85Hz). Eye position was recorded at a sampling rate of 1000Hz with a remote, infrared eye tracker (Eyelink 1000; SR Research, Ottawa, Canada). Head position was stabilized using a chin rest. The experiment was controlled by in-house software written in the Python programming language (2.7) using the PsychoPy package (1.82.02) (Peirce, 2007; 2008).

Experimental design

Here, we only considered the data from the synchronous task condition in Aponte et al. (2018). Subjects participated in three blocks of mixed pro- and antisaccade trials. Each block consisted of 200 trials of which either 20, 50, or 80% were prosaccade trials. Before the main experiment, subjects underwent a training block of 50 prosaccade trials followed by 50 antisaccade trials. In the training phase (but not during the main experiment) subjects received feedback about their performance.

As shown in Fig. 1A, two red circles (radius 0.25°) were presented throughout the experiment at an eccentricity of ±12°. Each trial started with a central fixation cross (0.6x0.6°). Subjects were required to fixate for at least 500ms, after which a random interval of 500 to 1000ms started. Completed this period, the fixation cross disappeared, and a green bar (3.48x0.8°) centered on one of

the red circles was presented in either horizontal or vertical orientation for 500ms. Subjects were instructed to saccade to the red circle ipsilateral to a horizontal green bar (prosaccade trials), and to saccade to the contralateral circle in the opposite case (antisaccade trials). After 1000ms the next trial started. Horizontal and vertical bars were randomly interleaved, but the same sequence was presented across all subjects. The location (left of right) of the peripheral cue was also randomly permuted, such that the number of pro- and antisaccade trials in each direction was the same.

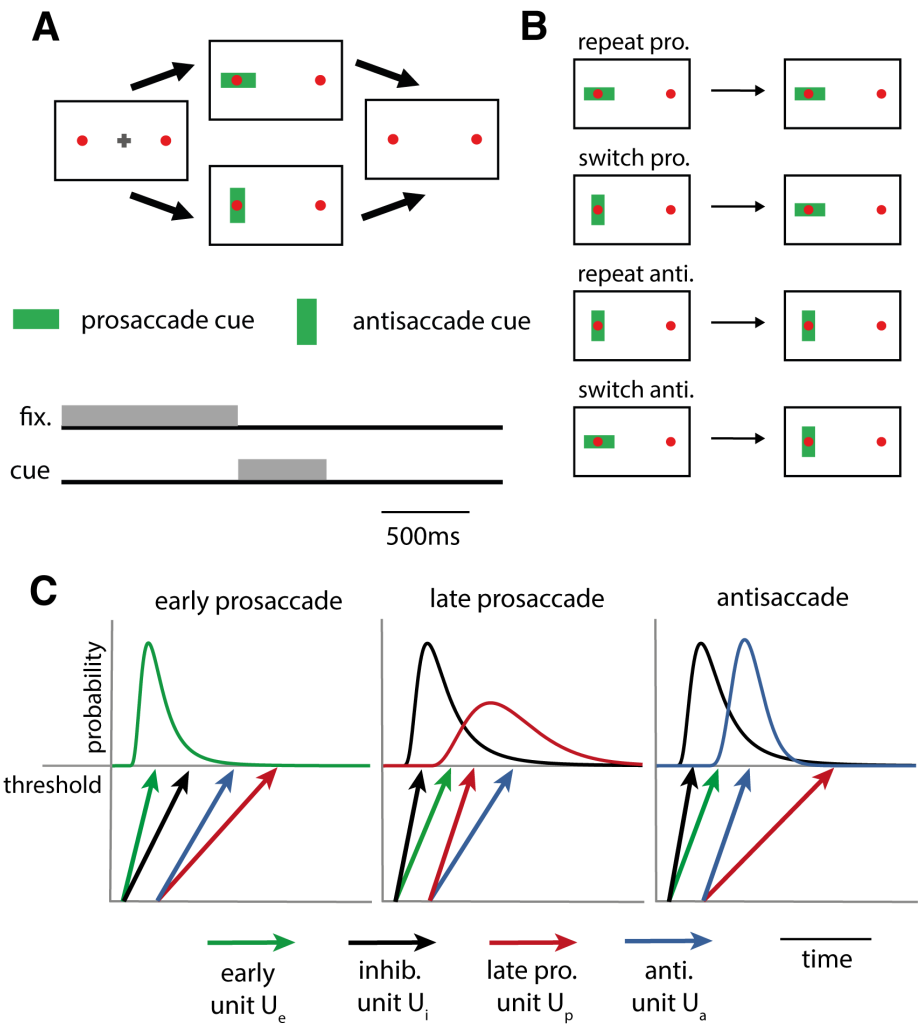

Figure 1: Experimental design and model. Adapted with permission from Aponte et al. 2018. **A. Experimental design**. Subjects were instructed to first fixate to a central cross. After a variable interval (500-1000ms), a cue indicating the trial type was presented for 500 ms. **B. Trial types.** Depending on the cue presented in the previous trial, four types of trials were considered. From top to bottom: repeat prosaccade, switch prosaccade, repeat antisaccade, and switch antisaccade. The cue was presented to the right or left pseudorandomly. **C. SERIA model.** The SERIA model is a race model that incorporates four different units: an early prosaccade unit, an inhibitory unit, a late prosaccade and an antisaccade unit (see legend). The RT distributions are a function of the hit time distribution of the four units. Early reactions occur when the early unit hits threshold before all other units. Late reactions occur mainly when early reactions are stopped by the inhibitory unit. Note that late pro- and antisaccade compete with each other in a similar manner. In principle, different trial types could affect the hit time distribution of each unit.

167 *Data processing*

Saccades were detected with the software provided by the eye tracker manufacturer (Stampe, 1993), which uses a 22°/s and 3800°/s2 threshold to define the start of a saccade. Only saccades with a magnitude larger than 2° were included in the analysis. Trials were rejected in case of an eye blink, if subjects failed to maintain fixation before the peripheral green bar was presented, and if their latency was above 900ms or below 50ms, and, in the case of antisaccades, below 95ms.

Statistical Analysis

As variables of interest, we investigated mean RT of correct responses and mean ER. These were analyzed with a generalized, linear mixed effect (GLME) model implemented in MATLAB 9.3 (*fitglme.m*). Independent variables were prosaccade probability (PP) with levels 20, 50 and 80%; trial type (TT); last trial (LT) with levels *switch* and *repeat* (Fig. 1B); and SUBJECT entered as a random effect. In the case of ER, the probit function acted as link function in the GLME. The first trial of a block was not included in this analysis.

The SERIA model

Briefly, SERIA (Aponte et al., 2017) models the race of four independent accumulators or units: an early (u_e), an inhibitory (u_i), a late prosaccade (u_p), and an antisaccade (u_a) unit. The pair of an action $A \in \{pro., anti.\}$ and its latency $T \in [0, \infty[$ is treated as a random variable, whose distribution is a function of the hit times of each of the units, U_e, U_i, U_p, and U_a respectively. Conceptually, SERIA can be decomposed into two different competitions: First, the early unit, which models reactive, habitual responses, generates a prosaccade at time t if it hits threshold at time t (i.e., $U_e = t$) and all the other units hit threshold afterwards. An early response can be stopped by the inhibitory unit if the latter hits threshold at some earlier point. In that case, either a late prosaccade or an antisaccade is generated. This second decision process is modeled as the race between the corresponding units, such that, for example, a late prosaccade at time t is generated only if the late prosaccade unit hits threshold at $U_p = t$ before the antisaccade unit (i.e., $U_a > t$).

More concretely, SERIA provides an explicit formula for the probability of an action A and its RT. First, a prosaccade at time t is generated when either the

early unit u_e hits threshold at time t (i.e., $U_e = t$) before all other units. The probability of this event is given by

$$p(U_e = t)p\big(U_p > t\big)p(U_a > t)p(U_i > t). \tag{1}$$

Furthermore, a prosaccade at time t can be triggered when the late prosaccade unit hits threshold at t and before all other units

$$p\big(U_p = t\big)p(U_e > t)p(U_a > t)p(U_i > t) \tag{2}$$

or an early response is stopped by the inhibitory unit (i.e., $U_i < t$ and $U_i < U_e$), and the late prosaccade unit hits threshold before the antisaccade unit

$$p\big(U_p = t\big)p(U_a > t)\int_0^t p(U_i = \tau)p(U_e > \tau)d\tau\,. \tag{3}$$

Similarly, an antisaccade at time t is generated when the antisaccade unit hits threshold at t ($U_a = t$), before all other units

$$p(U_a = t)p(U_e > t)p\big(U_p > t\big)p(U_i > t) \tag{4}$$

or after an early prosaccade has been stopped

$$p(U_a = t)p\big(U_p > t\big)\int_0^t p(U_i = \tau)p(U_e > \tau)d\tau\,. \tag{5}$$

To fit the model, we assumed a parametric form for the hit times of each of units: the hit times of the early (U_e) and inhibitory unit (U_i) were modeled with an inverse Gamma distribution, while the hit times of the late units (U_p and U_a) were modeled using a Gamma distribution (Aponte et al., 2017). Thus, each unit could be fully characterized by two parameters controlling the mean and variance of the hit times. Accordingly, to fully specify the distribution of actions and RTs in a condition, 8 (unit) parameters were required.

Model space

We aimed to answer three different questions through Bayesian model comparison (Kass and Raftery, 1995; Stephan et al., 2009): First, are models that included information about the previous trial superior in explaining experimental data compared to models which did not include this factor? Second, are inter-trial effects driven by either the trial type, the action, or the correctness of the action performed in the previous trial? Third, can inter-trial

effects be accounted for by changes in either the generation of goal-directed actions, inhibitory control, or a combination of both?

To answer these questions, we fitted models that explained the totality of the data not only as a function of the current trial type, but also as a function of the previous trial. For this, all trials were divided into four different conditions, according to the cue displayed (pro- or antisaccade) and the previous trial, as explained below. Although a different set of parameters could operate in each condition, this seems biologically implausible and our goal was to identify which parameters could be fixed across conditions, without compromising the ability of the models to parsimoniously explain the behavioral data. As illustrated in Table 1, we first evaluated models in which the parameters of the units were constant irrespective of the previous trial but could vary depending on the cue presented on the current trial. Second, we evaluated models that accounted for effects of the last trial in three manners: either the parameters of the model depended on (i) the previous trial type (Fig. 1B), (ii) the previous action, or on (iii) whether the previous action was an error or not. In principle, one could also consider interactions among these three factors, but this would require 2x2x2x2x8=128 parameters. Hence, we limited the analysis to the three options mentioned above and ignored their interaction, acknowledging that such effects have been reported in the literature (DeSimone et al., 2014).

These families of models were divided into three nested versions that represented different types of inter-trial effects: either the parameters of the inhibitory and early unit were fixed across conditions, but the late units were allowed to vary across conditions, or only inhibitory control changed across conditions, or both the late and the inhibitory units could differ in all conditions, but the early units were equal across them.

Table 1

	Model space					
	Late units		**Inhib. units**		**Late+inhib. units**	
	Model	# Pars.	Model	# Pars.	Model	# Pars.
No switch	m_1	15	m_2	13	m_3	17
Switch trial	m_4	23	m_5	21	m_6	29
Last action	m_7	23	m_8	21	m_9	29
Last error	m_{10}	23	m_{11}	21	m_{12}	29

Model space. In addition to the 8 parameters controlling the units, all models included three parameters that accounted for the 'no response' time, the late response delay, i.e., a constant delay associated with all goal-directed saccades and, finally, the probability of an outlier, i.e., the probability of a saccade faster than the no-response-time. In model m_1, we assumed that the early and inhibitory units were equal in pro- and antisaccade trials (4 parameters), but the late units were allowed to vary across pro- and antisaccade trials (4x2 parameters). Thus, the total number of parameters was 3+4+4x2=15. In the case of models m_4, m_7 and m_{10}, the parameters of the late units varied in all four possible conditions, resulting in 2x4 (two switch trials times four parameters) additional degrees of freedom (15+2x4=23). In model m_2, the inhibitory unit was allowed to vary between pro- and antisaccade trials, but the late units were fixed across the two conditions. Thus, m_2 had 13 parameters, i.e., two parameters less than m_1. In models m_5, m_8 and m_{11}, we allowed the inhibitory unit to vary across all four conditions. Note that m_2 had the identical late decision process in pro- and antisaccade trials. We relaxed this severe restriction in models m_5, m_8 and m_{11} and allowed the late units to vary between pro- and antisaccade trials, but kept them fix across switching. Thus, theses three models had 3x2+4+13=21 parameters. Lastly, in the *late+inhib.* family we allowed both the inhibitory and late units to vary, keeping the parameters of the early unit fixed. In model m_3 this implied two more parameters when compared to m_1, while in models m_6, m_9 and m_{12} this led to 3x2 more parameters than e.g. m_4.

Model fitting

The models were estimated using the techniques presented in our previous studies (Aponte et al., 2017; 2018). Data from all subjects were entered simultaneously into a hierarchical model presented in Aponte et al. (2017; 2018). Samples from the posterior distribution were drawn using the Metropolis-Hasting algorithm. The evidence or marginal likelihood of models

was computed using thermodynamic integration with 32 parallel chains ordered according to the temperature schedule suggested by Aponte et al. (2018). The algorithm was run for 60000 iterations, from which only the last half was used to compute summary statistics. The software used here to implement the models and inference is available in the TAPAS toolbox (http://translationalneuromodeling.org/tapas/).

We were interested in several model-based statistics derived from the fits. First, we evaluated the probability of an inhibition failure, defined as the probability that the early unit hits threshold before all the other units:

$$p(inhib.fail.) = \int_0^\infty p(U_e = t)p(U_i > t)p\big(U_p > t\big)p(U_a > t)dt\,. \qquad (6)$$

Inhibition failures are fast, reflexive prosaccades, which are correct on prosaccade trials and errors on antisaccade trials. The expected RT of an inhibition failure is

$$E[inhib.fail.RT] = \frac{1}{p(inhib.fail.)}\int_0^\infty t\,p(U_e = t)p(U_i > t)p\big(U_p > t\big)p(U_a > t)dt\,. \qquad (7)$$

We also report the conditional probability of a late prosaccade, defined as the probability that the late prosaccade unit hits threshold before the antisaccade unit:

$$p(late\,pro.) = \int_0^\infty p\big(U_p = t\big)p(U_a > t)dt\,. \qquad (8)$$

Note that the conditional probability (given no inhibition failure) of a late antisaccade is defined as

$$p(anti.) = 1 - p(late\,pro.). \qquad (9)$$

We were also interested in the expected hit times of the late units, defined as

$$E[late\,pro.hit\,time] = \frac{1}{p(late\,pro.)}\int_0^\infty t\,p\big(U_p = t\big)p(U_a > t)dt \qquad (10)$$

and analogously so for antisaccades. This quantity is the expected hit time of the late prosaccade unit, conditioned on the antisaccade unit arriving at a later point. We report this statistic, as it conveys an interpretable quantity that can be readily compared to experimental data. The derivation for these terms can

be found in Aponte et al. (2017). Finally, we computed the probability of a saccade at time t to be an early response:

$$p(early|T = t, A = pro) = \frac{p(early\ pro, T=t)}{p(early\ pro, T=t) + p(late\ pro, T=t)}. \quad (11)$$

Results

The results are based on all trials of the experiment, after roughly 2.2% of the trials were discarded due to blinks or other artefacts, as explained in detail in Aponte et al. (2018). We first report a classical statistical analysis of ER and RT. Only correct trials entered the analysis of RT. Then, we proceed to the model-based findings.

Switch costs: Error rate

Mean ERs (Fig. 2A-C) were submitted to a binomial regression model. On prosaccade trials, there was a significant effect of both PP ($F(2,138) = 36.4, p < 10^{-3}$) and LT ($F(1,138) = 178.7, p < 10^{-3}$). On average, switch prosaccade trials showed 15% more errors than repeat trials. In the case of antisaccades, the main effects of PP ($F(2,138) = 17.1, p < 10^{-3}$) and LT ($F(1,138) = 147.7, p < 10^{-3}$) as well as the interaction between PP and LT ($F(2,138) = 4.3, p = 0.01$) were significant. Switch trials led to 18% more errors on antisaccades trials. When we included all trials in a single model, there was a significant effect of TT (anti. ER – pro. ER=9%; $F(1,276) = 112.4, p < 10^{-3}$) and a significant interaction between TT and LT ($F(1,276) = 4.8, p = 0.02$), indicating a significant difference in switch cost between anti- and prosaccade trials ($\text{anti. switch cost} - \text{pro. switch cost} = 2.8\%$).

Reaction time

To evaluate the effect of previous trials, we submitted the mean RT of correct trials (Fig. 2D-F), to a GLME that included PP, LT, and SUBJECT as independent variables. There was a significant effect of LT on both pro, and antisaccades trials (pro. switch cost $= 14ms, F(1,138) = 37.1, p < 10^{-3}$; anti. switch cost $= 12ms, F(1,138) = 32.7, p < 10^{-3}$). The effect of PP was significant on antisaccade ($F(1,138) = 6.6, p < 10^{-3}$) but not on prosaccade trials ($F(1,138) = 0.9, p = 0.39$). In both cases, the interaction between PP and LT was not significant. We then considered pro- and antisaccades in a single model with factors PP, LT, TT and all possible interactions. Our main interest was the interaction between the factors LT and TT, which would indicate a significant difference in the switch cost between pro- and antisaccades. This interaction was not significant ($F(2,276) = 0.2, p = 0.64$).

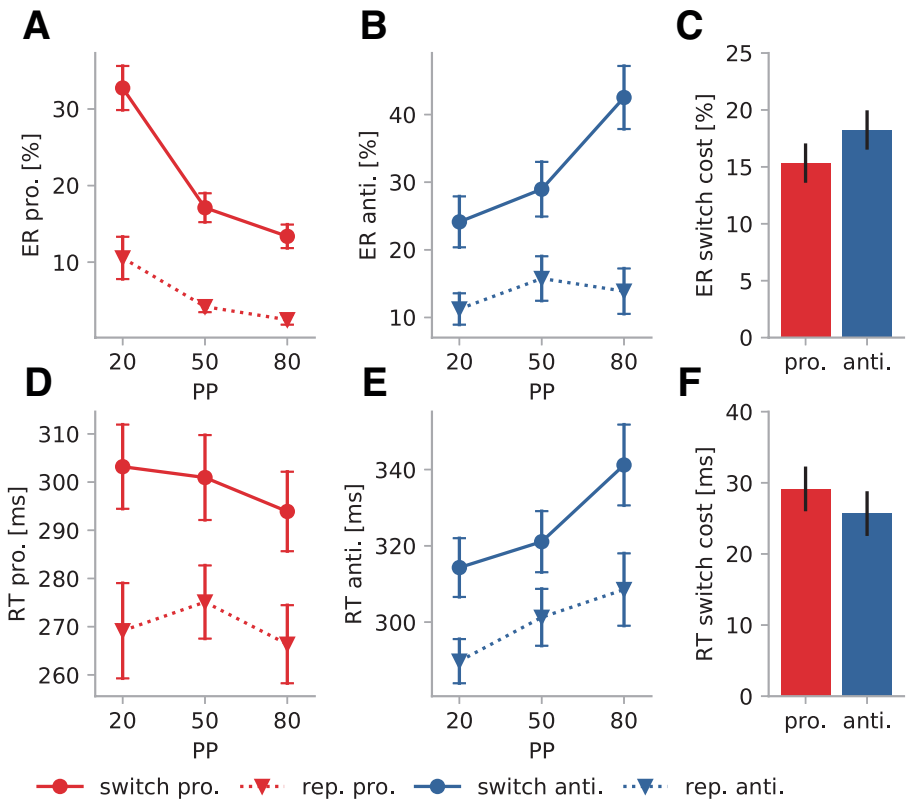

Figure 2: A. Mean ER on prosaccade trials. **B.** Mean ER on antisaccade trials. **C.** ER switch costs. **D.** Mean RT on prosaccade trials. **E.** Mean RT on antisaccade trials. **E.** RT switch cost. Error bars display the s.e.m.. PP: prosaccade probability

SERIA results – model comparison

All models were initially evaluated according to their log evidence or marginal likelihood, which corresponds to the accuracy or expected log likelihood adjusted by the complexity of the model (Stephan et al., 2009). Table 2 (top) reports the evidence of all models in log units. The model with the highest evidence (m_6, LME = -17467.6, ΔLME>8.0 log units) allowed for differences in the late and inhibitory units across all conditions. In addition, the accuracy or expected log likelihood is reported in Table 2 (bottom). In general, models in which switch costs depended on the last trial type, as opposed to the last action or error, obtained higher evidence.

Table 2 illustrates the penalization for complexity (or number of parameters) in terms of the difference between log model evidence (top) and expected log likelihood (bottom) in each model. For example, in the case of the *no-switch:late+inhib.* model (m_3) the penalty was 817.8 log units. Similarly, the

penalty for *switch:inhib* (m_5) was 857.7 units, whereas for the best model (*switch:late+inhib*; m_6) the penalty was 985.5 log units. This demonstrates that the *switch:late+inhib* model provided the most parsimonious explanation of the data as the increase in expected log likelihood outweighed the increased complexity.

Table 2

	Log model evidence		
	Late	Inhib.	Late+inhib.
No switch	-17596.0	-19571.0	-17591.4
Switch trial	-17481.1	-17475.8	**-17467.6**
Last action	-17599.1	-17577.1	-17634.8
Last error	-17657.5	-17623.0	-17706.1
	Expected log-likelihood		
No switch	-16842.8	-18806.0	-16773.6
Last trial type	-16532.2	-16618.1	**-16482.1**
Last action	-16661.9	-16797.4	-16598.0
Last error	-16751.1	-16804.4	-16715.5

Model comparison. The model with the highest evidence is high lightened in bold face. A difference of 3 log units is typically regarded as strong evidence (Kass and Raftery, 1995). The expected log-likelihood is also displayed for comparison. Note that within all model families (rows) the difference in expected log-likelihood between the most complex and the other models is strongly reduced (for last trial type) or even inverted (last action or last error) in log model evidence due to the penalization for complexity.

To illustrate qualitative differences between models, the fits of m_4 (*switch:late*), m_5 (*switch:inhib.*) and m_6 (*switch:late+inhib.*) are displayed in Fig. 3 and 4. Fits represent the expected posterior predictive probability estimated from posterior samples. Visual inspection suggests that the *switch:inhib.* model failed to capture the distribution of late responses, and particularly so on prosaccade switch trials. The *switch:late* model made a better job regarding late saccades, but it did not seem to capture early responses on, for example, prosaccade repeat trials (Fig. 3 bottom row). Finally, the *switch:late+inhib.* model was able

to accommodate most of the features of the behavioral data, with perhaps the exception of the early peak on prosaccade switch trials.

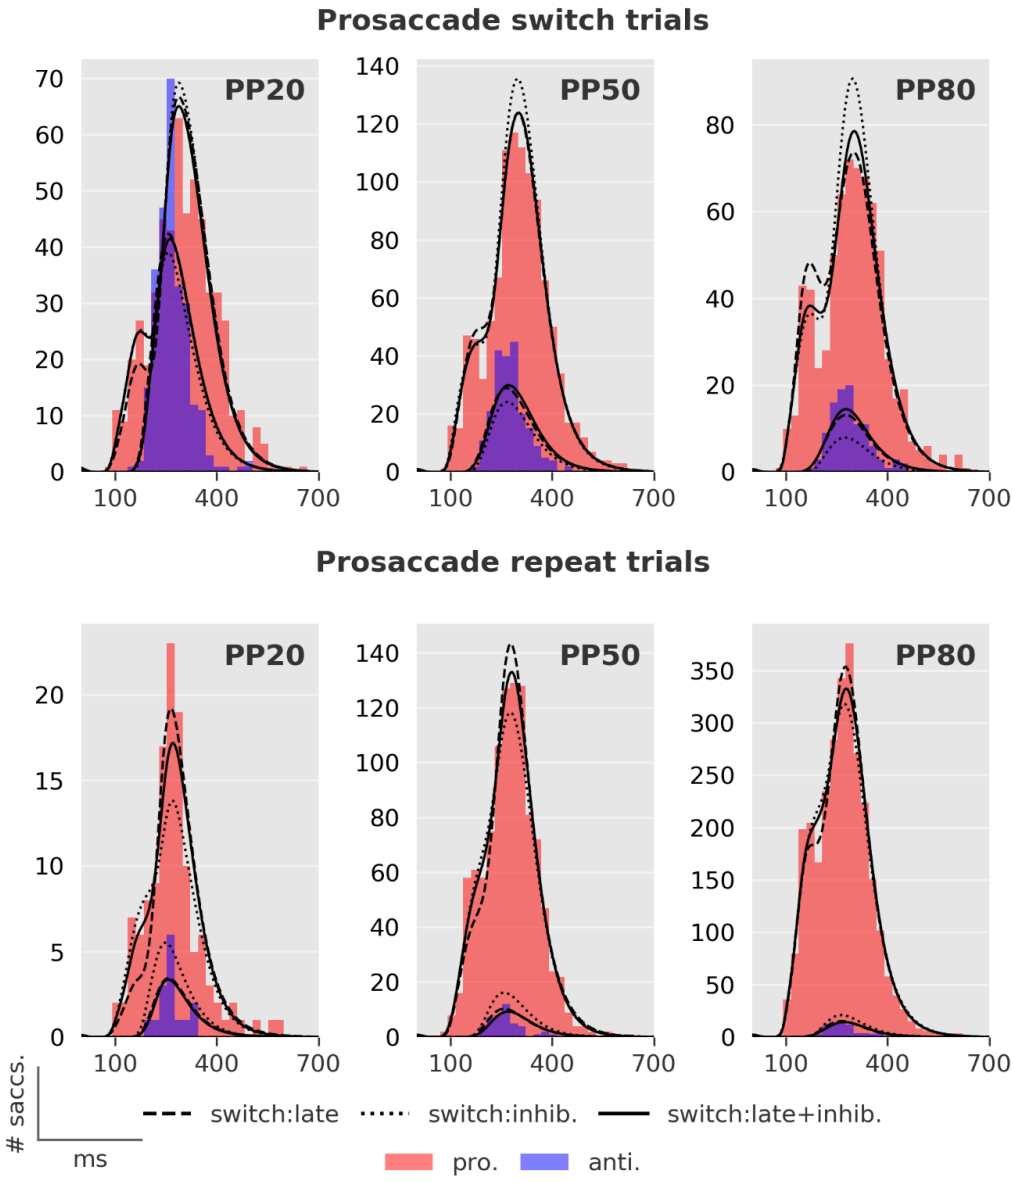

Figure 3: Histogram of prosaccade trials and model fits. Prosaccades are displayed in red, and antisaccades in blue. For comparison, the weighted posterior predictive distribution of models m_4 (*switch:late*), m_5 (*switch:inhib.*) and m_6 (*switch:late+inhib.*) computed from posterior samples are displayed. The *switch:late+inhib.* model fitted both the early and late peaks in the prosaccade RT distribution better.

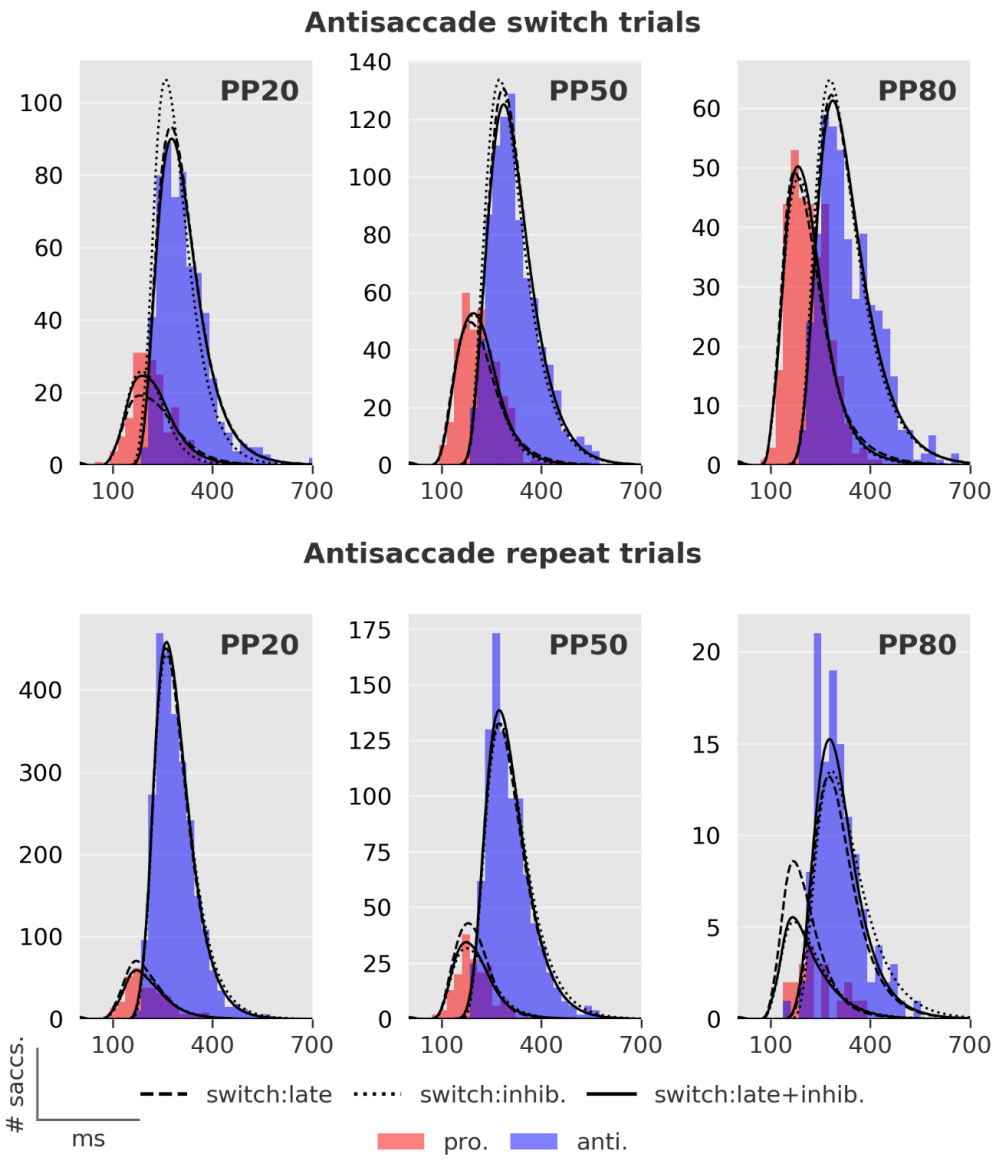

Figure 4: Histogram of antisaccade trials and model fits. For a detailed legend see Fig. 3. Note that all three models showed a relatively poor model fit for errors in the PP80 condition in repeat trials. This can be explained by the very low number of errors (N=19 across all subjects) in this condition.

SERIA results – parameter estimates

In the second part of the analysis, we investigated how model-based parameter estimates differed across conditions. For this, we used the parameter estimates from the model with the highest evidence, namely the *switch:late+inhib.* model. The focus of this analysis was the hit times of the late units and the early responses, as well as the expected proportion of late errors, and the number of inhibition failures in switch and repeat trials.

Threshold hit times

The expected hit time of the late pro- (327ms, std. 35ms) and antisaccade (314ms, std. 33ms) units (Fig. 5A-B) were first submitted to two separate GLMs. The main effect of LT was significant for the two late units (pro. switch cost = 16ms, std. 19ms, $F(1,138) = 28.6, p < 10^{-3}$; anti. switch cost = 23ms, std. 21ms, $F(1,138) = 46.6, p < 10^{-3}$). When considered together, we found no significant interaction between LT and TT ($F(1,276) = 2.4, p = 0.11$). Finally, the interaction between PP and TT was significant ($F(1,276) = 5.0, p = 0.007$).

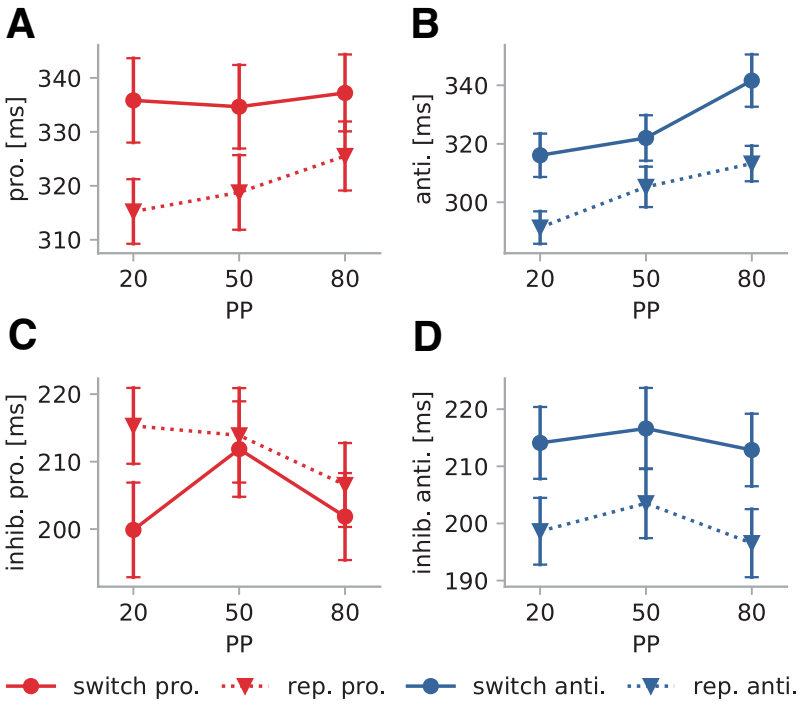

Figure 5: **A.** Hit times of the late prosaccade unit (see Eq. 10). **B.** Hit times of the antisaccade unit. **C.** RT of inhibition failures on prosaccade trials (see Eq. 7). **D.** RT of inhibition failures in antisaccade trials. Error bars depict the s.e.m..

In addition to the late units, we investigated the latency of inhibition failures (i.e., non-stopped early reactions; Fig. 5C-D). In SERIA, these are always prosaccades and occur on both pro- and antisaccade trials when the early unit hits threshold before all other units. Note that inhibition failures are correct in prosaccade trials, but errors in antisaccade trials. When pro- and antisaccade trials were analyzed together, we found a significant effect of PP ($F(2,276) = 16.2, p < 10^{-3}$), LT ($F(1,276) = 4.0, p = 0.045$), and the interaction between TT and LT was significant, as well ($F(1,276) = 46.9, p < 10^{-3}$). As shown in Fig.

5D, inhibition failures had a higher mean latency on antisaccade switch trials (214ms, std. 32ms) than on repeat trials (199ms, std. 29ms), whereas on prosaccade trials (Fig. 5D), the opposite pattern occurred (switch 204ms, std. 33ms; rep. 211ms, std. 31ms). Thus, according to our model fits, inhibition failures following an antisaccade trial were faster than inhibition failures following a prosaccade trial.

To illustrate the above effect, we compared the mean RT of errors on switch (227ms, std. 36ms) and repeat antisaccade trials (221ms, std. 48). We have shown previously that most of the errors in this paradigm are due to inhibition failures (Aponte et al., 2017), and we therefore hypothesized that errors on antisaccade trials that follow prosaccade trials should have a higher latency than errors that follow antisaccade trials. In agreement with our prediction, we found a positive switch cost (4ms, std42) with respect to errors on antisaccade trials. Note that the scant number of errors on antisaccade repeat trials in the PP80 condition (N=19, see Fig. 4, bottom right panel) prevented us from conducting a statistical analysis. However, if the PP80 block was excluded, the switch cost was still positive (11ms; std. 30ms) and significant ($F(1,86) = 5.1, p = 0.025$).

Error types

In order to characterize different sources of errors, we investigated early and late errors separately (Fig. 6). First, we submitted the probability of late errors (Eq. 8 and 9) on pro- (mean 19%, std. 14) and antisaccade (mean 4%, std.4) trials to a single GLME. This revealed a positive switch cost as demonstrated by a significant effect of LT ($F(1,276) = 55.9, p < 10^{-3}$). Moreover, there was a significant interaction between LT and TT ($F(1,276) = 14.7, p < 10^{-3}$). The mean switch cost for late prosaccades was 11% (std. 11), whereas for antisaccades, it was 1% (std. 4.8). When late antisaccades were analyzed separately, the effect of the previous trials was still significant ($F(1,138) = 9.87, p = 0.002$).

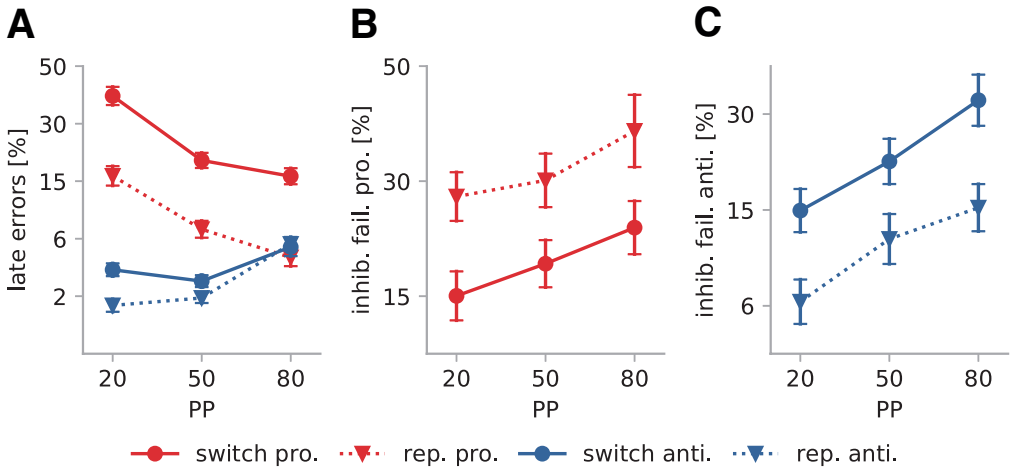

Figure 6: A. Percentage of late errors (see. Eq. 8 and 9)**. B.** Percentage of inhibition failures in prosaccade trials (cf. Eq. 6). **C.** Percentage of inhibition failures in antisaccade trials. Percent errors are presented in the probit scale and summary statistics (mean and s.e.m.) were computed from probit transformed probabilities. Error bars represent the s.e.m..

We then investigated the probability of an inhibition failure (see Eq. 6), defined as the probability that the early unit hits threshold before all other units. According to our model, on prosaccade trials 28% (std. 19) of saccades were inhibition failures, whereas this number was lower on antisaccade trials (mean 21%, std. 18). The effect of LT on pro- ($F(1,138) = 105.5, p < 10^{-3}$) and antisaccades trials ($F(1,276) = 149.2, p < 10^{-3}$) was significant. When considered together, we found a significant interaction between the factors TT and LT ($F(1,276) = 219.8, p < 10^{-3}$). Concretely, prosaccade trials induced more inhibition failures on the next trial regardless of trial type (pro. switch cost=-12%, std. 10; anti. switch cost=11%, std. 9). Finally, while we found a weak but significant interaction between PP and TT ($F(1,276) = 3.2, p = 0.04$), the three-way interaction between PP, TT and LT was not significant ($F(2,276) = 1.3, p = 0.25$).

In Fig. 7, the effects of switch and repeat trials on early responses are exemplified. We depict the histogram of all responses in the PP50 condition sorted into trial types, and switch and repeat trials. We restrict this illustration to the PP50 case because the number of repeat and switch trials was similar (pro. switch: 2479, pro. repeat: 2150, anti. switch: 2435, anti. repeat 2150). The predicted histograms of early responses are overlaid on the empirical

distributions. The former were computed by weighting each saccade's RT by its probability of being an early response (Eq. 11). Clearly, on prosaccade trials, the early peak of repeat trials (orange) is larger than on switch trials (green), as emphasized by the predicted histograms. The opposite pattern is observable on antisaccade trials, in which switch trials have a larger early peak than repeat trials.

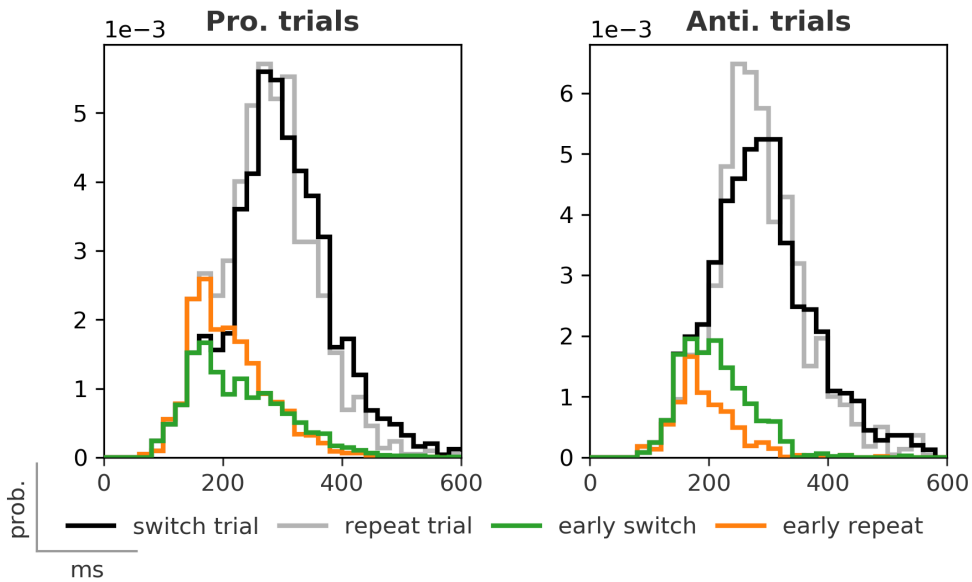

Figure 7: Normalized histograms of all saccades in the PP50 condition and predicted early responses. The histograms of all responses on pro- and antisaccade trials are sorted into switch and repeat trials. Histograms were normalized to have unit area. Overlaid are the predicted histogram of early responses, computed by weighting each saccade by its probability of being an early response (Eq. 11). In prosaccade trials, there were more early saccades in repeated trials than on switch trials, while the opposite pattern could be observed on antisaccade trials. Note that the black and gray histograms represent the empirical distributions, while orange and green histograms show modeling results.

Fig. 8 summarizes the effect of switching on the RT of early and late reactions, late ER, and the percentage of inhibition failures. Our findings indicate a strong positive switch cost in RT and ER for late responses, with the exception of late errors on antisaccade trials. Regarding inhibition failures, antisaccade trials decreased the frequency and latency of early reactions on the subsequent trial compared to prosaccade trials. This observation is compatible with the claim that an antisaccade trial leads to enhanced inhibitory control on the following

trial. The reason is that in a race model, when the distribution of the hit times of the stop unit is shifted to the right, more early reactions with *higher* RT are possible. In other words, faster inhibitory control leads to *fewer but faster* errors.

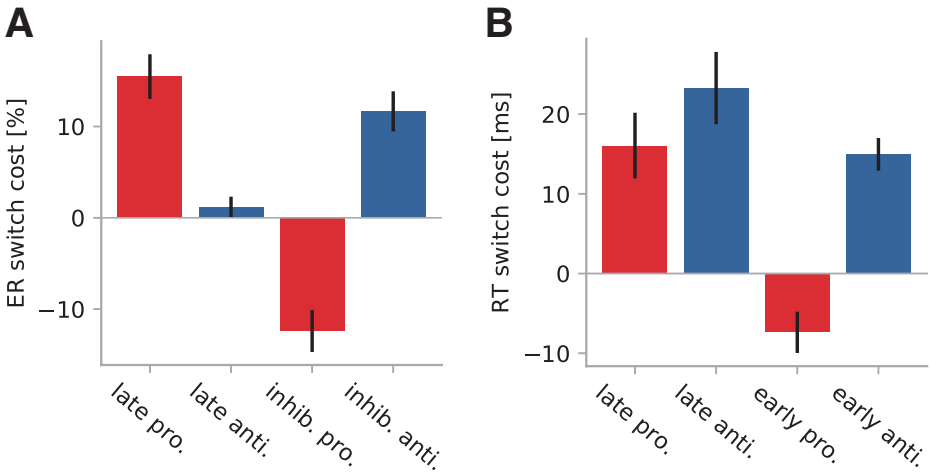

Figure 8: Summary of switch costs. A. Error rates: Task switching increase the probability of late errors in pro- and antisaccade trials. This cost was higher on prosaccade trials when compared to antisaccade trials. In contrast, we found that there was a negative switch cost in the number of inhibition failures on prosaccade trials, while we observed the opposite effect on antisaccade trials. Thus, prosaccade trials lead to more inhibition failures on the subsequent trial irrespective of trial type. **B. Hit times.** There was a positive and significant switch cost for late pro- and antisaccades. Early reactions that followed a prosaccade trial showed a higher latency than those following an antisaccade trial. Error bars display the s.e.m..

Discussion

In the present study, we investigated switch costs in a mixed pro- and antisaccade task with the SERIA model. This allowed us to quantify to what extent switch costs affected the inhibition of habitual reactions (early prosaccades) and goal-directed (rule-guided) behavior (late pro- and antisaccades), respectively. Modeling revealed two distinguishable effects: on the one hand, switch trials engendered RT costs in goal-directed behavior. On the other hand, antisaccade trials enhanced inhibitory control compared to prosaccade trials, as demonstrated by fewer inhibition failures with lower latency following antisaccade trials. In the following, we discuss these findings.

Behavioral findings

Two types of switch costs have been commonly reported for the antisaccade task. Early studies (e.g., Barton et al., 2002; Cherkasova et al., 2002; Manoach et al., 2002; Fecteau et al., 2004) reported positive prosaccade RT switch costs, negative antisaccade RT switch costs, as well as an increase in ER following switch trials in both type of trials. More recently, Heath and Weiler (e.g. Weiler and Heath, 2012b; Weiler et al., 2015) have reported a positive switch cost on prosaccade trials, and no switch cost on antisaccade trials.

Our behavioral analysis based on standard statistics supports positive switch costs in pro– and antisaccades, both in terms of ER and RT. Although our study seems at odds with previous reports, we believe that our findings complement rather than contradict them, because the design used here differs in key aspect from the tasks used by other researchers. Notably, in our paradigm (inspired by Sato and Schall, 2003), the peripheral spatial cue also signaled the action to be performed, whereas in most studies the trial type cue precedes the spatial cue by hundreds of milliseconds. We refer to these two paradigms (Aponte et al., 2018) as the synchronous cues (SC) and asynchronous cues (AC) designs, respectively. Importantly, it has been shown (e.g., Weiler and Heath, 2014; Aponte et al., 2018) that displaying the task cue before the peripheral saccadic target has a dramatic effect on ER and RT when compared to the SC task. Supporting the idea that different designs can trigger different switch costs, Barton et al. (2006a) showed that when the task cue was presented only 200ms

before the peripheral cue, there was a positive switch cost in both pro- and antisaccades. However, if the peripheral cue was delayed by 1800ms, a negative antisaccade cost could be observed. Here, we replicated the observation that a very short cue-to-peripheral-stimulus delay (0ms in our case) causes switch costs regardless of trial type.

Weiler and Heath (2014) have also noted methodological differences across studies. For example, in studies reporting a negative antisaccade switch cost (e.g., Cherkasova et al., 2002), the saccadic targets (in our case, the two peripheral red points) were displayed during the entire experiment. This group has attributed the difference in findings to this manipulation (Barton et al., 2006a).

One possible explanation for the contrast between AC and SC designs is that if the trial type cue is presented early enough relative to the peripheral cue, the decision processes that triggers voluntary actions can be consolidated to the point that contextual factors, such as the previous trial or the probability of a task cue, become ineffective (Weiler and Heath, 2014). Thus, the AC might abolish task switching effects in voluntary action generation by allowing participants to overwrite contextual effects during the cue-to-stimulus interval.

We speculate that seemingly contradictory reports could be explained by a single computational mechanism. In other words, a unified explanation of behavioral findings might be attained through the lens of quantitative modeling. Thus, we proceed to discuss our computational results.

Computational Modelling

The main goal of our study was to account for switch costs in the antisaccade task using the SERIA model. The conclusions presented here are based on quantitative Bayesian model comparison, as well as on qualitative posterior predictive fits (Fig. 3 and 4). Our results indicate that the bulk of effects of the last trial can be accounted for by alternations of trial type, and not by the previous actions. Note, however, that we cannot rule out interactions between trial type and actions that have been previously reported (Tatler and Hutton, 2007; DeSimone et al., 2014).

Qualitative inspection of the predictive fits (Fig. 3 and 4) clearly indicates that accounting for switch costs in both habitual and goal-directed actions improved the fit of the RT distributions. The second best model ($\Delta LME = 8.2$) accounted for switch costs only through changes in inhibitory control (*switch:inhib.*). Examination of the predictive fits clearly shows why this model was inferior: The RT distribution of late responses could not be properly fitted.

The analysis of parameter estimates led to two conclusions. First, there was a positive switch cost in late pro- and antisaccades in terms of RT and ER. However, the switch cost associated with late errors on antisaccade trials was small (1%). Second, antisaccade trials led to fewer and faster habitual responses on the following trial. Therefore, we conclude that goal-directed behavior is facilitated by the application of the same rule over trials. Moreover, the enhanced inhibitory control induced by antisaccade trials led to more efficient inhibition on the following trial.

It is not straightforward to compare our model-based findings to predictions from conceptual models, because these do not always translate into precise empirical predictions. In particular, our results support two theories that are considered opposites (Weiler and Heath, 2014). On the one hand, the switch cost observed in goal-directed behavior supports the task-set inertia hypothesis (Weiler et al., 2014), according to which the activation of a cue-action mapping facilitates the activation of the same rule on the next trial, while interfering with other mappings.

On the other hand, our results also support the oculomotor inhibition hypothesis (Allport et al., 1994), according to which the type of inhibitory control required to execute an antisaccade leads to enhanced inhibition on the following trial, as reflected by a fewer inhibition failures following an antisaccade trial, as well as faster early prosaccades. The latter is a natural consequence of more efficient inhibitory control, which allows only inhibition failures with very short latencies. Note that the existence of inter-trial effects on inhibitory control is supported by evidence in the countermanding saccade task (e.g., Barton et al., 2006a; Weiler and Heath, 2014). In summary, our model

supports both theories, depending on whether the interpretation focusses on the inhibition of reactive responses or on the generation of voluntary actions.

Finally, we note that our modeling approach cannot distinguish whether antisaccade trials enhance inhibitory control, or whether prosaccade trials disengage it. Here, we have opted to assert that antisaccade trials enhance inhibition; however, in the present context, both options are undistinguishable.

Other models of inter-trial effects

To our knowledge, this is the first computational, trial-by-trial model used to investigate switch costs in the antisaccade task. Although SERIA seems to accommodate the most salient features of our data, one conceptual limitation of our approach is its phenomenological or descriptive nature. For instance, SERIA does not postulate any specific neurophysiological mechanism to explain the different types of switch costs. This role could be better filled by biologically inspired models (Emeric et al., 2007; Pouget et al., 2011), that can be compared to behavioral data as well as to physiological recordings.

A further limitation is that learning effects were not modeled explicitly. Because no information about the trial type probability was provided to participants, the PP effects reported here and by others (Cutsuridis et al., 2007; Heinzle et al., 2007; Wiecki and Frank, 2013; Lo and Wang, 2016) can only be explained by adaptation to statistical regularities in the stimuli. Here, we made the assumption that subjects quickly reach a steady state such that dynamic effects can be ignored. However, since the pioneering work of Carpenter and Williams (1995), it has been hypothesized that changes in saccadic RT can be explained by the principles of optimal statistical inference (assuming, for example, that humans act like ideal Bayesian observers; Brodersen et al., 2008). More recently, Vossel et al. (2014) (but see also Anderson and Carpenter, 2006; Brodersen et al., 2008) used a Bayesian model (Mathys et al., 2011) to demonstrate trial-by-trial variations in saccadic RT in a modified Posner paradigm, in which the average validity of the cue changed over time. In future work we plan to account for learning effects, in addition to the last trial effects modeled here.

Summary

Our quantitative modeling suggests that conceptual theories of switch costs in the antisaccade task can profit from a more precise formulation in computational terms, as seemingly contradictory statements can be correct at the same time. This is possible because of the non-trivial interactions between a habitual response mechanism that can be subject to fast inhibitory control and the ability to generate goal-directed, context sensitive behavior. Concretely, our analysis indicates that alternating between goal-directed behaviors engenders sizeable switch costs, whereas increased inhibitory demands on one trial enhance inhibitory control on the following trial.

References

Allport A, Styles EA, Hsieh S (1994) Shifting Intentional Set: Exploring the Dynamic Control of Tasks. In: Attention and performance series. Attention and performance 15: Conscious and nonconscious information processing (Umilta C, Moscovitch M, eds), pp 266–290. Cambridge, MA, US: The MIT press.

Anderson AJ, Carpenter RH (2006) Changes in expectation consequent on experience, modeled by a simple, forgetful neural circuit. J Vis 6:822–835.

Ansari TL, Derakshan N, Richards A (2008) Effects of anxiety on task switching: evidence from the mixed antisaccade task. Cogn Affect Behav Neurosci 8:229–238.

Aponte EA, Schobi D, Stephan KE, Heinzle J (2017) The Stochastic Early Reaction, Inhibition, and late Action (SERIA) model for antisaccades. PLoS Comput Biol 13:e1005692.

Aponte EA, Tschan DG, Heinzle J, Stephan KE (2018) Inhibition and late errors in the antisaccade task: Influence of task design. bioRxiv.

Barton JJ, Cherkasova MV, Lindgren K, Goff DC, Intriligator JM, Manoach DS (2002) Antisaccades and task switching: studies of control processes in saccadic function in normal subjects and schizophrenic patients. Ann N Y Acad Sci 956:250–263.

Barton JJ, Greenzang C, Hefter R, Edelman J, Manoach DS (2006a) Switching, plasticity, and prediction in a saccadic task-switch paradigm. Exp Brain Res 168:76–87.

Barton JJ, Raoof M, Jameel O, Manoach DS (2006b) Task-switching with antisaccades versus no-go trials: a comparison of inter-trial effects. Exp Brain Res 172:114–119.

Bojko A, Kramer AF, Peterson MS (2004) Age equivalence in switch costs for prosaccade and antisaccade tasks. Psychol Aging 19:226–234.

Brodersen KH, Penny WD, Harrison LM, Daunizeau J, Ruff CC, Duzel E, Friston KJ, Stephan KE (2008) Integrated Bayesian models of learning and decision making for saccadic eye movements. Neural Netw 21:1247–1260.

Camalier CR, Gotler A, Murthy A, Thompson KG, Logan GD, Palmeri TJ, Schall JD (2007) Dynamics of saccade target selection: race model analysis of double step and search step saccade production in human and macaque. Vision Res 47:2187–2211.

Carpenter RH, Williams ML (1995) Neural computation of log likelihood in control of saccadic eye movements. Nature 377:59–62.

613 Chan JL, Koval MJ, Johnston K, Everling S (2017) Neural correlates for task
614 switching in the macaque superior colliculus. J Neurophysiol 118:2156–
615 2170.

616 Cherkasova MV, Manoach DS, Intriligator JM, Barton JJ (2002) Antisaccades
617 and task-switching: interactions in controlled processing. Exp Brain Res
618 144:528–537.

619 Cutsuridis V, Smyrnis N, Evdokimidis I, Perantonis S (2007) A neural network
620 model of decision making in an antisaccade task by the superior colliculus.
621 Neural Networks 20:690–704.

622 DeSimone JC, Weiler J, Aber GS, Heath M (2014) The unidirectional prosaccade
623 switch-cost: correct and error antisaccades differentially influence the
624 planning times for subsequent prosaccades. Vision Res 96:17–24.

625 Emeric EE, Brown JW, Boucher L, Carpenter RH, Hanes DP, Harris R, Logan GD,
626 Mashru RN, Pare M, Pouget P, Stuphorn V, Taylor TL, Schall JD (2007)
627 Influence of history on saccade countermanding performance in humans
628 and macaque monkeys. Vision Res 47:35–49.

629 Fecteau JH, Au C, Armstrong IT, Munoz DP (2004) Sensory biases produce
630 alternation advantage found in sequential saccadic eye movement tasks.
631 Exp Brain Res 159:84–91.

632 Greenwood TA et al. (2007) Initial heritability analyses of endophenotypic
633 measures for schizophrenia: the consortium on the genetics of
634 schizophrenia. Arch Gen Psychiatry 64:1242–1250.

635 Hallett PE (1978) Primary and secondary saccades to goals defined by
636 instructions. Vision Res 18:1279–1296.

637 Heath M, Gillen C, Samani A (2016) Alternating between pro- and
638 antisaccades: switch-costs manifest via decoupling the spatial relations
639 between stimulus and response. Exp Brain Res 234:853–865.

640 Heath M, Starrs F, Macpherson E, Weiler J (2015) Task-switching effects for
641 visual and auditory pro- and antisaccades: evidence for a task-set inertia. J
642 Mot Behav 47:319–327.

643 Heinzle J, Aponte EA, Stephan KE (2016) Computational models of eye
644 movements and their application to schizophrenia. Current Opinion in
645 Behavioral Sciences 11:21–29.

646 Heinzle J, Hepp K, Martin KA (2007) A microcircuit model of the frontal eye
647 fields. J Neurosci 27:9341–9353.

648 Isoda M, Hikosaka O (2008) Role for subthalamic nucleus neurons in
649 switching from automatic to controlled eye movement. J Neurosci
650 28:7209–7218.

Karayanidis F, Jamadar S, Ruge H, Phillips N, Heathcote A, Forstmann BU (2010) Advance preparation in task-switching: converging evidence from behavioral, brain activation, and model-based approaches. Front Psychol 1:25.

Kass RE, Raftery AE (1995) Bayes factors. Journal of the american statistical association 90:773–795.

Lee AK, Hamalainen MS, Dyckman KA, Barton JJ, Manoach DS (2011) Saccadic preparation in the frontal eye field is modulated by distinct trial history effects as revealed by magnetoencephalography. Cereb Cortex 21:245–253.

Lo CC, Wang XJ (2016) Conflict Resolution as Near-Threshold Decision-Making: A Spiking Neural Circuit Model with Two-Stage Competition for Antisaccadic Task. PLoS Comput Biol 12:e1005081.

Logan GD, Cowan WB, Davis KA (1984) On the ability to inhibit simple and choice reaction time responses: a model and a method. J Exp Psychol Hum Percept Perform 10:276–291.

Manoach DS, Lindgren KA, Cherkasova MV, Goff DC, Halpern EF, Intriligator J, Barton JJ (2002) Schizophrenic subjects show deficient inhibition but intact task switching on saccadic tasks. Biol Psychiatry 51:816–826.

Mathys C, Daunizeau J, Friston KJ, Stephan KE (2011) A bayesian foundation for individual learning under uncertainty. Front Hum Neurosci 5:39.

Monsell S (2003) Task switching. Trends Cogn Sci (Regul Ed) 7:134–140.

Mueller SC, Swainson R, Jackson GM (2009) ERP indices of persisting and current inhibitory control: a study of saccadic task switching. Neuroimage 45:191–197.

Munoz DP, Everling S (2004) Look away: the anti-saccade task and the voluntary control of eye movement. Nat Rev Neurosci 5:218–228.

Myles JB, Rossell SL, Phillipou A, Thomas E, Gurvich C (2017) Insights to the schizophrenia continuum: A systematic review of saccadic eye movements in schizotypy and biological relatives of schizophrenia patients. Neurosci Biobehav Rev 72:278–300.

Noorani I, Carpenter RH (2013) Antisaccades as decisions: LATER model predicts latency distributions and error responses. Eur J Neurosci 37:330–338.

Peirce JW (2007) PsychoPy--Psychophysics software in Python. J Neurosci Methods 162:8–13.

Peirce JW (2008) Generating Stimuli for Neuroscience Using PsychoPy. Front Neuroinform 2:10.

Pierce JE, McCardel JB, McDowell JE (2015) Trial-type probability and task-switching effects on behavioral response characteristics in a mixed saccade task. Exp Brain Res 233:959–969.

Pouget P, Logan GD, Palmeri TJ, Boucher L, Pare M, Schall JD (2011) Neural basis of adaptive response time adjustment during saccade countermanding. J Neurosci 31:12604–12612.

Radant AD et al. (2015) Robust differences in antisaccade performance exist between COGS schizophrenia cases and controls regardless of recruitment strategies. Schizophr Res 163:47–52.

Reilly JL, Frankovich K, Hill S, Gershon ES, Keefe RS, Keshavan MS, Pearlson GD, Tamminga CA, Sweeney JA (2014) Elevated antisaccade error rate as an intermediate phenotype for psychosis across diagnostic categories. Schizophr Bull 40:1011–1021.

Sato TR, Schall JD (2003) Effects of stimulus-response compatibility on neural selection in frontal eye field. Neuron 38:637–648.

Stampe D (1993) Heuristic filtering and reliable calibration methods for video-based pupil-tracking systems. Behavior Research Methods, Instruments, \& Computers 25:137–142.

Stephan KE, Penny WD, Daunizeau J, Moran RJ, Friston KJ (2009) Bayesian model selection for group studies. Neuroimage 46:1004–1017.

Tatler BW, Hutton SB (2007) Trial by trial effects in the antisaccade task. Exp Brain Res 179:387–396.

Vossel S, Mathys C, Daunizeau J, Bauer M, Driver J, Friston KJ, Stephan KE (2014) Spatial attention, precision, and Bayesian inference: a study of saccadic response speed. Cereb Cortex 24:1436–1450.

Weiler J, Hassall CD, Krigolson OE, Heath M (2015) The unidirectional prosaccade switch-cost: electroencephalographic evidence of task-set inertia in oculomotor control. Behav Brain Res 278:323–329.

Weiler J, Heath M (2012a) The prior-antisaccade effect influences the planning and online control of prosaccades. Exp Brain Res 216:545–552.

Weiler J, Heath M (2012b) Task-switching in oculomotor control: unidirectional switch-cost when alternating between pro- and antisaccades. Neurosci Lett 530:150–154.

Weiler J, Heath M (2014) Oculomotor task switching: alternating from a nonstandard to a standard response yields the unidirectional prosaccade switch-cost. J Neurophysiol 112:2176–2184.

Weiler J, Mitchell T, Heath M (2014) Response suppression delays the planning of subsequent stimulus-driven saccades. PLoS ONE 9:e86408.

Wiecki TV, Frank MJ (2013) A computational model of inhibitory control in frontal cortex and basal ganglia. Psychol Rev 120:329–355.