

Cite as: Török, B., Janacsek, K., Nagy, D. G., Orbán, G., & Nemeth, D. (2017). Measuring and filtering reactive inhibition is essential for assessing serial decision making and learning. *Journal of Experimental Psychology: General*, 146(4), 529.

**Measuring and filtering reactive inhibition is essential for assessing
serial decision-making and learning**

Balázs Török^{*1,2,3,4}, Karolina Janacsek^{*2,3}, Dávid G. Nagy^{*4,5}, Gergő Orbán^{**4}, Dezső
Németh^{**2,3}

* these authors contributed equally to this work

** these authors contributed equally to this work

1. Department of Cognitive Science, Budapest University of Technology and Economics, *Egry József utca 1. H-1111, Budapest, Hungary*
2. Institute of Psychology, Eötvös Loránd University, *Izabella utca 46., H-1064, Budapest, Hungary*
3. MTA-ELTE NAP B Brain, Memory and Language Research Group, Institute of Cognitive Neuroscience and Psychology, Research Centre for Natural Sciences, Hungarian Academy of Sciences, *Magyar tudósok körútja 2., H-1117, Budapest, Hungary*
4. Computational Systems Neuroscience Lab, Wigner Research Centre for Physics, Hungarian Academy of Sciences, *Konkoly Thege Miklós út 29-33., H-1121, Budapest, Hungary*
5. Institute of Physics, Eötvös Loránd University, *Pázmány Péter sétány 1/A, H-1117, Budapest, Hungary*

Abstract

Learning complex structures from stimuli requires extended exposure and often repeated observation of the same stimuli. Learning induces stimulus-dependent changes in specific performance measures. The same performance measures, however, can also be affected by processes that arise due to extended training (e.g. fatigue) but are otherwise independent from learning. Thus, a thorough assessment of the properties of learning can only be achieved by identifying and accounting for the effects of such processes. Reactive inhibition is a process that modulates behavioral performance measures on a wide range of time scales and often has opposite effects than learning. Here we develop a tool to disentangle the effects of reactive inhibition from learning in the context of an implicit learning task, the alternating serial reaction time task. Our method highlights that the magnitude of the effect of reactive inhibition on measured performance is larger than that of the acquisition of statistical structure from stimuli. We show that the effect of reactive inhibition can be identified not only in population measures but also at the level of performance of individuals, revealing varying degrees of contribution of reactive inhibition. Finally, we demonstrate that a higher proportion of behavioral variance can be explained by learning once the effects of reactive inhibition are eliminated. These results demonstrate that reactive inhibition has a fundamental effect on the behavioral performance that can be identified in individual participants and can be separated from other cognitive processes like learning.

Keywords: reaction time, reactive inhibition, statistical learning, implicit learning, computational modelling

Introduction

Reactive inhibition (RI), accumulative performance deterioration during the execution of a series of tasks, characterizes multiple cognitive processes and spans time scales from tens of seconds to tens of minutes (Hull, 1943; Pan & Rickard, 2015). Reactive inhibition is defined under conditions where the same task is repeated multiple times and is exhibited in the progressive decline of a wide array of performance measures, including motor execution accuracy (Kelso, 2014; Stelmach, 1969), decision accuracy (Tipper et al., 2002; Grison & Strayer, 2001; Houghton et al., 1996), and execution delay (Rickard et al., 2008, Brawn et al., 2010). Situations requiring sustained attention, rapid consecutive inferences and serial decision making are ubiquitous in daily life, for instance during driving a car under busy traffic conditions where inferences are continuously taken about the state of the environment and about other agents in traffic. Therefore the relevance of reactive inhibition is likely to go beyond repetitive task execution. Critically, the same performance measures characterize an opposing process, learning, which induces performance improvements during the course of extended exposure to a set of training trials. This antagonistic relationship between RI and learning makes it difficult to discriminate between the impacts that these processes have on task performance over time, and the effect of RI on performance is often ignored. The aim of the present study is to develop a method for disentangling RI from learning-related processes in reaction time measurements.

Neglecting RI can have a substantial effect on performance measures and therefore on the interpretation of experimental data. A multitude of studies have investigated memory consolidation by exploring the effects of sleep on procedural learning. The sleep-dependent

enhancement hypothesis is based on the observation that within this domain, sleep seems to improve performance relative both to a matched waking period and to the end of the preceding day's training session (Diekelmann & Born, 2007; Robertson, Pascual-Leone, & Miall, 2004; Walker, 2005). However, Rickard et al. (2008) presented evidence that the buildup of RI leads to an underestimation of achieved skill in the task, which could result in illusory gains after sleep. Brawn et al. (2010) demonstrated that such an apparent performance improvement is present after a 5-minute resting period in massed-trial conditions and no further improvement is introduced by sleep. Pan & Rickard (2015) showed in a meta-analysis that indeed, accounting for RI along with a few other non-stimulus dependent factors challenges the evidence for enhanced learning after sleep. Thus, if we want to achieve a detailed understanding of the processes underlying learning, it is essential to be able to dissociate the effects of these concurrent processes.

In this study, we set out to disentangle the contributions of RI and learning to measured performance. Importantly, the relative contributions of these different processes to the changes in performance are expected to be largely varying across participants. As a consequence, in contrast to previous approaches, we performed our analyses for each participant individually, which ensured that opposing effects of RI and learning do not wash out in population averages. We chose a probabilistic sequence learning task as the learning paradigm, which we hypothesized should show a robust RI effect, and would allow us to observe the time course of RI at a fine temporal scale. In this task, stimuli appear sequentially in predefined locations on a screen and participants have to press the button corresponding to the observed stimulus location. Unbeknownst to the participant, the location sequences have statistical structure, which can be used to predict some of the upcoming stimuli. Participants do not become aware of this predictability, however their response times differentially improve for more predictable stimuli indicating that they were able to implicitly capture at least some of the statistical

regularities, a phenomenon termed statistical learning. We suggest that this Alternating Serial Reaction Time (ASRT) paradigm is an appropriate model for tasks requiring sustained attention paired with quick successive inferences and actions where RI is bound to be a significant factor. It also has the advantage that the probabilistic structure of the task is sufficiently hard to learn, and provides a large number of successive measurements over an extended period of time, making fine temporal structure visible even for individual participants.

The rest of the paper is structured in the following way: first, we present the experimental setting in detail. Next, we introduce our parametric model for reaction times and discuss how we estimated the parameters and compared models of different complexity. Finally, we present the analysis of the ASRT task, which is performed on a large sample of participants. Key to this learning paradigm is to be able to reliably assess the capability of individuals to extract statistical patterns. Using the model comparison method, we estimate the relative magnitude of the considered processes and identify key properties of learning measures that make them resistant to the effects of RI. We show that a subject-by-subject fit for RI improves the predictive power of a statistical learning model and that statistical learning can actually explain more of the variance after RI has been filtered out from the raw data.

Methods

Participants

One hundred and eighty healthy young adults participated in the study. All participants had normal or corrected-to-normal vision and none of them reported a history of any neurological and/or psychiatric condition. Participants provided informed consent to the procedure as approved by the institutional review board of the local research ethics committee. The study was

conducted in accordance with the Declaration of Helsinki and participants received course credits for taking part in the experiment.

Alternating Serial Reaction Time Task

The Alternating Serial Reaction Time (ASRT) task was used to measure statistical learning capabilities of individuals (J. H. Howard, Jr. & Howard, 1997). In this task, a stimulus (a cartoon image of a dog's head) appeared in one of four horizontally arranged empty circles on the screen (Nemeth et al., 2010). Participants were instructed to press a corresponding key (Z, C, B, or M on a QWERTY keyboard) as quickly and accurately as they could after the stimulus was presented. The target remained on the screen until the participant pressed the correct button. The response to stimulus interval (RSI) was 120 msec. The ASRT task consisted of 45 presentation blocks in total, with 85 stimulus presentations per block (Fig. 1). After each of these training blocks, participants received feedback about their overall reaction time and accuracy for 5 seconds, and then they were given a 10-second rest before starting a new block. Each of the three sets of 15 training blocks constitutes a training session. Between training sessions, a longer (3-5 min) break was introduced. We used EPRIME 2.0 as a stimulus presentation software (Schneider, 2012).

Unbeknownst to the participants, the presentation of stimuli in a given block was structured. After 5 warm-up random trials, an eight-element sequence was repeated ten times. In the sequence, predetermined (P) and random (r) elements alternated with each other (e.g., 2 – r – 1 – r – 3 – r – 4 – r; where numbers denote the four locations on the screen from left to right, and r's denote randomly chosen locations out of the four possible ones with equal probabilities). The computer program selected a different ASRT sequence for each participant such that in the entire participant pool each of the six unique permutations of the 4 repeating events occurred equal number of times (J. H. Howard, Jr. & Howard, 1997; Nemeth et al., 2010).

	Structure 2 – R – 1 (the last event is pattern)	Structure R – 1 – R (the last event is random)
High-probability triplets	2 – 1 – 1 (50%)	2 – 1 – 1 (12.5%)
Low-probability triplets	Never occurring (pattern events are always high probability)	2 – 1 – 2 (12.5%) 2 – 1 – 3 (12.5%) 2 – 1 – 4 (12.5%)

Table 1. Description of high- and low-probability triplets in the ASRT task. If the current triplet started with 2–1, the following event can either be a pattern with 50% probability (top left cell), in which case the next event is 1; or it can be a random event, where all locations have equal probability (12.5%). Hence, having seen 2–1 as the start of our triplet, there is a high probability (62.5%) of seeing a 1 and a low probability (12.5%) for seeing anything else.

The alternating sequence in the ASRT task gives rise to the occurrence of particular three-element sequences (henceforth referred to as triplets) with higher probability than others. For instance, if the sequence is 2 – r – 1 – r – 3 – r – 4 – r, then 2X1, 1X3, 3X4, and 4X2 (“X” indicates the middle element of the triplet) occurred often, and they did so either by the third element (bold numbers) being derived from the sequence (P - r - **P**) or so that it was a random element (r - P - r). In contrast, infrequent triplets, such as 1X2 and 4X3, could occur only in one way (r - P - r). Following previous studies, we refer to the former as *high-probability triplets* and the latter as *low-probability triplets* (Nemeth, Janacsek, & Fiser, 2013). Note that due to the higher occurrence probability, the third element of a high-probability triplet was more predictable from the first element of the triplet than in the case of low-probability triplets. For instance, in the example shown above (see Table 1, more details in Appendix), Position 2 as the first element of

a triplet is more likely (62.5%) to be followed by Position 1 as the third element, than either Position 2, 3, or 4 (12.5%, each). In accordance with this principle, each item was categorized as either the third element of a high- or a low-probability triplet, and the accuracy and reaction time (RT) of the response to this item were compared between the two categories.

This task overcomes a common challenge of the field in that it allows us to separate pure statistical learning from other modulating factors of RT such as general skill improvements. Statistical learning is defined as faster and more accurate responses to high conditional probability events compared to that to low conditional probability ones (Song et al., 2007). In contrast, general skill improvements refer to the speed-up and changes in accuracy, which are independent of the conditional probabilities of the events. These improvements reflect more efficient visuo-motor and motor-motor coordination due to practice (Hallgató et al., 2013). In addition, in the ASRT task, participants are not generally aware of the alternating structure of the sequences, even after extended practice, or when sensitive recognition tests are used to assess explicit knowledge (D.V. Howard, Howard, Japikse, DiYanni, Thompson & Somberg, 2004; J.H. Howard & Howard, 1997; Song et al., 2007). Thus, in our experiment, an implicit, non-conscious form of learning was tested (Cleeremans & Dienes, 2008; Reber, 1989).

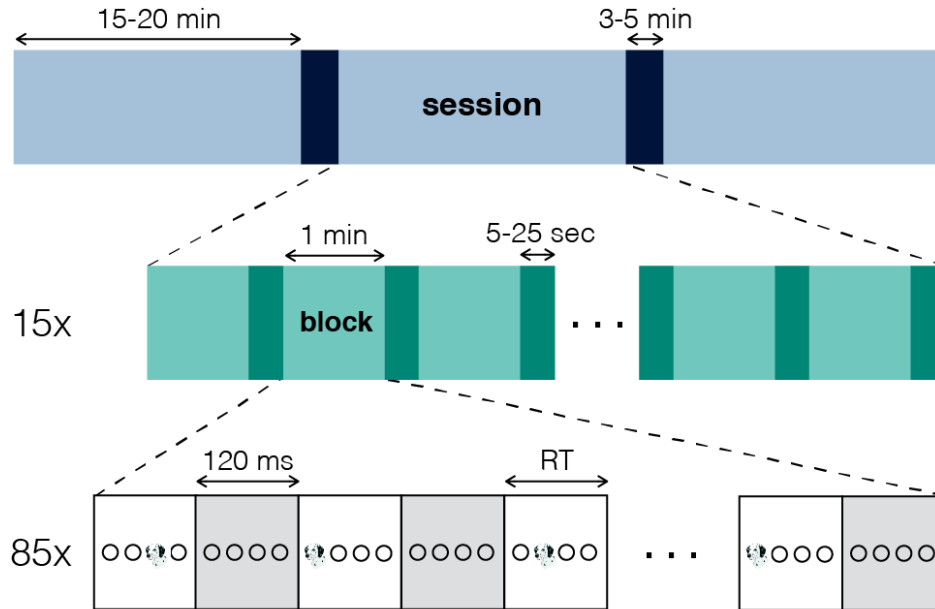


Figure 1. Structure of the ASRT task. Response-to-stimulus interval was fixed 120 msec. Between blocks of 85 trials (about 1 minute long), participants had the opportunity to rest with self-paced break time (about 5-25 seconds). Between training sessions of 15 blocks (15-20 minutes), questionnaires were completed for 3-5 minutes. The predetermined subsequence was held fixed within one participant but varied at random between participants.

Model of reaction times

To evaluate the contributions of different factors, we formulated a parametric model of the measured data. This model includes general skill improvements, statistical learning, reactive inhibition within and across blocks and reaction time changes over the longer breaks between each of the 15 blocks of ASRT. Thus, the complete formulation of the model is:

$$RT = b_1 + b_2 \cdot N^{-b_3} - p_1 \cdot \log P_N + (i_1 + i_2 \cdot B) \cdot n - i_3 \cdot S + \epsilon$$

where plus and minus signs are chosen such that the mean values of parameters are (expected to be) positive. In the following, we provide details and interpretations of the terms in this equation.

General skill improvements: Performance rapidly improves during the early stages of training, but this improvement becomes progressively slower. Building on Rickard et al. (2015), we modeled this with a decreasing power function (later referred to as basic adaptation) as a function of N , that is the number of the current trial in the entire experiment. Parameter b_1 defines an asymptotical performance that the power function converges to after an infinite amount of training, whereas b_2 and b_3 define the rate at which this convergence happens. More specifically, b_2 sets the scale and b_3 the curvature of the power function.

Statistical learning: Statistical learning was characterized by the difference between mean reaction times at high- and low-probability triplets. The mean reaction times are assumed to be a function of the *subjective predictive probability* P_N of the triplets (Carpenter & Williams, 1995). Since there are only two classes of stimuli considered in this study, the actual form of this relationship is not crucial. However, building on Carpenter & Williams (1995), we modelled the relationship by assuming the mean reaction time is proportional to the logarithm of the subjective probabilities. The reaction time associated with one triplet is that of the third element. For high-probability triplets the third element of each triplet has a 62.5% probability conditioned on the previous two elements, whereas for low-probability triplets it has only 12.5%. Parameter p_1 thus describes the magnitude of statistical learning observed in an individual's data. Note, that this formulation of statistical learning does not explicitly account for gradual acquisition of the statistical structure of the stimuli during the experiment.

Reactive inhibition: Subject-averaged reaction time data shows a clear, stimulus-independent residual structure, even though this structure is not visible when looking at individual participants' reaction times. The population average reveals a consistent change in the mean reaction times associated with the position of the stimulus both within one training block and throughout a training session (Figure 3A). We characterized the phenomenological form of this

dependence using three parameters: the rate of increase of reaction times within a training block i_1 , the rate of increase of the previous parameter between blocks i_2 , and the drop in reaction times between training sessions i_3 . Factor i_1 is the *slope* of reactive inhibition, the increase of RT between individual trials, in the first block, the amount with which the mean RT increases between two trials. Parameter i_2 captures the increment of this slope between two blocks, describing how much the reactive inhibition effect accelerates over the course of the experiment. This factor is later referred to as Quadratic Reactive Inhibition since its contribution to reaction times is quadratic in N . Parameter B is the number of the current block within a training session (runs 1-15). Therefore, the slope of reactive inhibition in a block will be $i_1 + i_2 \cdot B$. This slope is multiplied by the number n , the index of the trial within a training block (note n is the remainder of N divided by 85). Finally, i_3 captures a drop in reaction time at session boundaries, which correspond to 3 to 5-minute offline periods where S denotes the number of the current training session (runs 1-3). The role of this factor is to account for the apparent change at session boundaries, which is initially not expected to be explained by the other two reactive inhibition parameters. In this study, the role of the between-session drop factor is not directly interpreted. However, in some experiments, this measure is used for evaluating offline consolidation.

Error term: The error, ϵ , models both measurement error as well as inherent randomness in individual responses. The errors are assumed to be independent and normally distributed. We also tested the model with the assumption that the reciprocal of reaction times come from a normal distribution (Carpenter & Williams, 1995). However, we found that the measurement noise added to RT (lag due to keyboard for example) caused model fitting to be unstable with larger variance imposed on the model parameters.

Trials where the stimulus was identical to the previous (*repeat*) or penultimate (*recurrent*) stimulus were excluded from the analysis, as considered a standard procedure in this paradigm (J. H. Howard, Jr. & Howard, 1997, Song et al., 2007).

Parameter estimation

We estimated the parameters of the generative model using maximum likelihood inference on the reaction time data of each participant separately. Assuming normally distributed errors, this corresponds to least squared error estimation. During parameter estimation as well as during statistical procedures, we only used response times that corresponded to *correct* initial responses. We performed gradient based optimization using the built-in MATLAB[®] function *lsqcurvefit*, however, due to the non-convexity of the likelihood function, convergence to a global optimum cannot be guaranteed. To avoid suboptimal local maxima of the cost function, the optimization was carried out 20 times using randomly sampled uniformly distributed initial values within a bounded set and all but the best fit on the training set were discarded. In addition to the experimental datasets, the algorithm was also tested on synthetic data sampled from the same generative model with known parameters: examining whether the true parameters can be recovered. A substantial proportion (92%) of runs converged to the preset values on synthetic data. On human data, on average 85% of runs converged to the same parameter values with highest log-likelihood. Even with a more conservative estimate of 50% probability of convergence to the global maximum, global maxima are achieved by taking the best estimate out of 20 runs with high probability (probability of error: $0.5^{20} \ll 1\%$).

Selecting models for comparison

There are three sets of parameters whose individual or combined effects are analyzed. These sets are termed *basic adaptation*, *statistical learning* and *reactive inhibition*. In our analyses we

investigate the robustness of the effect of *reactive inhibition*. Therefore we compared predictive powers of models with and without *reactive inhibition* when either, neither or both *statistical learning* and *basic adaptation* are present (see Table 2). Considering these models also allows us to compare relative contributions of each of these components and establish the independence of our statistical learning measure from the other factors: *basic adaptation* and *reactive inhibition*.

	Asymptotic baseline	basic adaptation	Statistical learning	Reactive inhibition		
				Linear (within block)	Quadratic (between blocks)	session-gap
	b_1	b_2, b_3	p_1	i_1	i_2	i_3
Constant	X					
SL Statistical learning	X		X			
LRI Linear reactive inhibition	X			X		
BA basic adaptation	X	X				
BA + SL	X	X	X			
BA + LRI	X	X		X		
BA + LRI + QRI	X	X		X	X	
Complete	X	X	X	X	X	X

Table 2. Variants of the model used in the analyses. Different versions were composed to examine relative contributions of different components of the model as well as relative performance of different

factors. Note that both linear reactive inhibition and statistical learning components were evaluated with and without the presence of basic adaptation. Note that the asymptotic baseline is present in all models.

By assessing the simplest *reactive inhibition* model which only includes its *linear* term and no other model components (model LRI in Table 2), the magnitude of the effect of *reactive inhibition* can be directly compared with that of *statistical learning*.

Measures of performance

To assess the relative contributions of the factors, we compared the predictive performances of variants of the full model obtained by selectively omitting some of the components. Since we compared models of varying complexity, it was important to choose a measure that excludes the possibility of gaining advantage of overfitting noise present in the data by including extra parameters in the model. This was achieved using ten-fold cross-validation (Kohavi, 1995), where the models' predictions are evaluated on data points that were not used during the estimation of model parameters. Ideally, by considering all possible train-test dataset splits, one could have a model-comparison measure independent of train-test splits. Note that cross-validation and Akaike Information Criterion (AIC) have asymptotically equivalent conclusions (Stone, 1977). Model performance was relatively insensitive to the specific splits in train-test datasets, therefore we chose to use ten splits throughout this study (based on Kohavi, 1995). Specifically, reaction time data for the analyzed trials were split into ten disjoint sets for each participant, with each datapoint randomly assigned to one of ten disjoint sets. Then, for each participant, we performed ten least square parameter estimations. In each of these parameter estimations, nine of the disjoint datasets were chosen as training data S_{train} , and the remaining tenth set was chosen as the test data S_{test} . Importantly, a different set was left out as test set each time. To evaluate model performance on a subject-by-subject basis, R^2 values were calculated for the fit of the response time data for every subject and empirical distributions of R^2

values were constructed. The cumulative density function of this distribution provides an accurate characterization of the population in terms of the quality of the model's fit (Fig 6). The formula for the R^2 value for a given model f of variable x (containing all independent variables n, N, B, S, P) for a given test-train split is:

$$R^2 = 1 - \frac{\sum_{j \in S_{test}} (y_j - f(x_j))^2}{\sum_{j \in S_{test}} (y_j - \sum_{k \in S_{train}} \frac{1}{|S_{train}|} y_k)^2}$$

where y_j denotes the actual measured RT for the trial given by x_j .

The denominator corresponds to normalising using a baseline model which predicts the mean reaction time of the training set for all test data points. In statistical tests, we took the mean of the 10 model performance R^2 values for each model for each individual, and used a one-sided paired-sample t-test to assess theoretically implied model improvements.

Residual patterns in human performance that statistical learning can predict was assessed by calculating model improvement using the residual sum of squares in a pair of models: one containing the statistical learning term (extended) and the other deprived of this specific term but keeping all other terms intact (base). Thus, by plotting $1 - \frac{RSS_{extended}}{RSS_{base}}$ for model-pairs BA + SL / BA and Full model / Full model without SL (Table 2), relative improvements induced by inclusion of SL can be directly compared for each participant (Fig 7a).

Results

Statistical learning and reactive inhibition in the ASRT task, classical analysis

The Alternating Serial Reaction Time paradigm has been shown to provide detailed information on implicit learning of complex statistical structure in humans (Howard & Howard, 1997; Nemeth et al. 2013). Importantly, humans are sensitive to subtle details in the statistics, as reflected both by individual performances and population averages (Fig. 2), confirming earlier results of Janacsek et al. (2012).




Figure 2. Learning statistical structure in ASRT task. **A.** Evolution of mean reaction time of correct responses throughout the experiment for a naive participant. Bars show S.E.M. over RT measurements. **B.** Time course of learning for the population of 180 participants of the ASRT task. Since odd numbered stimuli follow a predetermined pattern, some stimuli have higher conditional probability of appearing when conditioned on the previous two stimuli. Third elements of triplets with higher probability are more predictable. There is a widening separation between reaction times to stimuli corresponding to triplet classes with different appearance frequencies, indicating that – even though none of them become consciously aware of it – the participants are sensitive to the statistical regularities present in the stimuli. Means over three blocks along with S.E.M. over subjects are plotted.

Systematic variations in reaction time data is the critical factor in deriving conclusions about learning, therefore understanding the factors contributing to this measure is of paramount importance. Population averaging the reaction time data (Fig 3a) marginalizes over patterns in the stimulus spanning more than eight presentations (i.e., the eight-element probabilistic sequence). As a result, the residual structure on larger time scales can reveal subtle details in human performance that have sources not directly related to stimulus statistics, such as reactive inhibition. The residual variance in our experiment (Fig 3a) exhibited substantial and complex structure spanning the whole period and modulated by the length of both training blocks and training sessions. The reaction times at the start and end of training blocks were significantly different (two-sided, paired-sample t-test on mean RTs over first and second halves of blocks calculated for each participant: $t(179)=-29.25$, $p<0.001$, Cohen's $d=-2.18$, Fig 3b). This increase in RT from first half of the block to the second half is significant despite the first RT in each block having the largest mean in each block (Fig 3A, this pattern of longest first response is also noted in Rickard et al, 2008). A significant difference was also detected between reaction times of end-of-session and start-of-session trials. This latter difference was tested by a two-sided, paired sample t-test on mean RTs in the last and first five blocks of training sessions for each individual (similarly to Song et al., 2007, Hallgató et al., 2013, Csabi et al., 2014). Reaction times significantly decreased both after the first session gap ($t(179)=22.70$, $p<0.001$, Cohen's $d=1.69$) and the second session gap ($t(179)=18.62$, $p<0.001$, Cohen's $d=1.39$, see Fig 3c). Within and across block changes in RT suggest the presence of reactive inhibition (for analogous scalloped RT patterns, see Rickard et al., 2008).

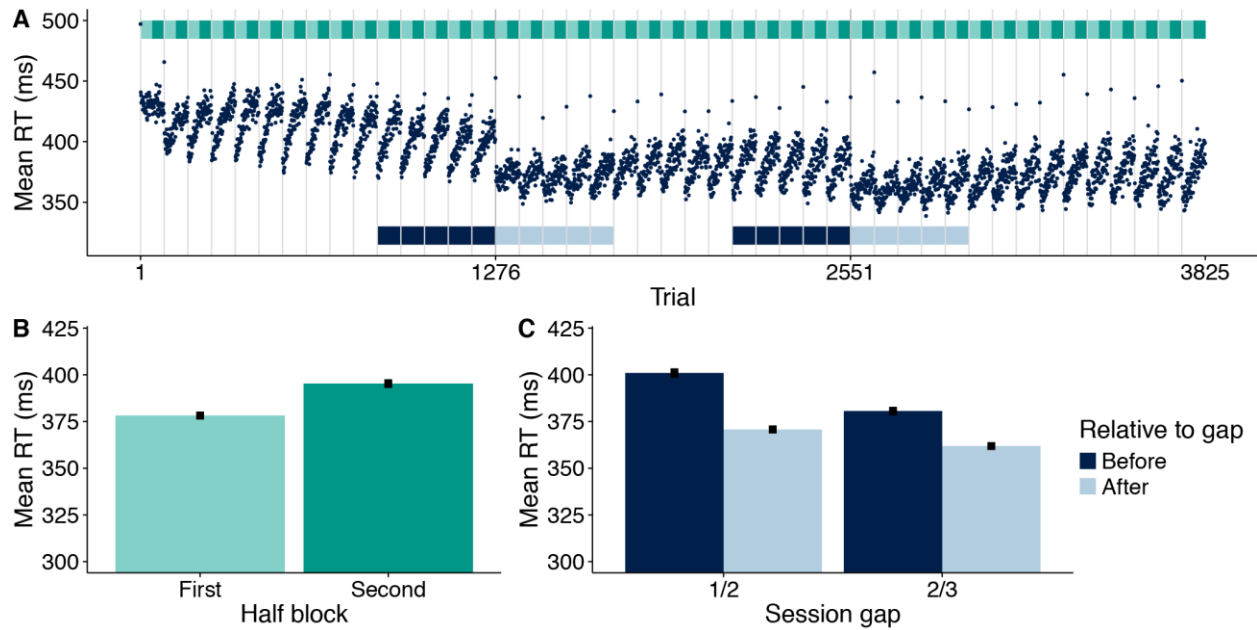


Figure 3. A. Reaction times for each trial averaged over participants show an elaborate structure with periods determined by both block and training session length, indicating the presence of stimulus-independent factors in measured performance. This particular scalloped structure is characteristic of reactive inhibition. **B.** Mean reaction times of first halves of blocks and second halves of blocks (indicated on panel **A** with bars of matching color). Reaction times deteriorate significantly during continuous sequences of task performance in a given block. Means calculated for each participant and then averaged over participants. **C.** Mean reaction times of five blocks (indicated on panel **A** with bars of matching color) before and after a session gap (3-5 minute break). A significant drop in average reaction times after the break is clearly visible. Error bars show S.E.M. over participants (**B, C**).

Model-based analysis of reaction time data

To shed light on the factors contributing to the prominent stimulus-independent variations in the data, we devised a *phenomenological* model, without providing a detailed account of the underlying psychological, physiological or neurological processes. It encompassed processes acting on different time scales to see their individual contributions (Fig 4a). Importantly, while a population-level analysis could reveal the details of overall processes, we devised a subject-by-subject analysis to avoid confounds possibly caused by subject averaging (Fig 4b). The model

can account for the different sources of variability in reaction time data and comprises three main components which are defined by seven parameters. By selectively eliminating some of the components, both their relative contributions and their dependencies can be established. We formulated different variations of the model by eliminating one or some of the components (Table 2). Thus, we obtained a model that included only a baseline reaction time (parameter b_1) and a decrease in reaction times with training (basic adaptation, parameters b_2 and b_3). One model captured the learning of the statistical structure of stimuli by learning the difference between the presentation frequencies of high-probability triplets and low-probability triplets (parameter p_1). Finally, further variants accounted for different forms of reactive inhibition, including an increase in reaction times during a block of trials (parameter i_1 for within-block (linear) changes and i_2 for between-block (quadratic) changes) and a decrease in reaction times between training sessions (parameter i_3).

Fitting 180 participants individually reveals that besides statistical learning - the distinction of high-probability triplets and low-probability triplets (Howard & Howard, 2007; Nemeth et al., 2013) - there are significant contributions of the other factors. One of these factors, reactive inhibition, acting on a short time-scale, is reflected in a systematic increase in reaction times within individual blocks of trials. Its presence is revealed by a significant improvement in the model performance (R^2) when adding a linear reactive inhibition factor (LRI model) to a constant model (paired-sample, one-sided t-test on mean R^2 s of the ten-fold cross-validation of the two models fitted to each participant's data: $t(179)=20.22$, $p<0.001$, Cohen's $d=1.51$). Another factor, namely basic adaptation (modeled using a decreasing power function), is reflected in a systematic decrease in RT over the entire course of the experiment. Its effect is reflected by a significant improvement in model performance (R^2), when adding basic adaptation component (BA model, Table 2) to a single constant model (paired-sample, one-sided t-test on mean R^2 s of

the ten-fold cross-validation of the two models fitted to each participant's data: $t(179)=17.11$, $p<0.001$, Cohen's $d=1.28$).

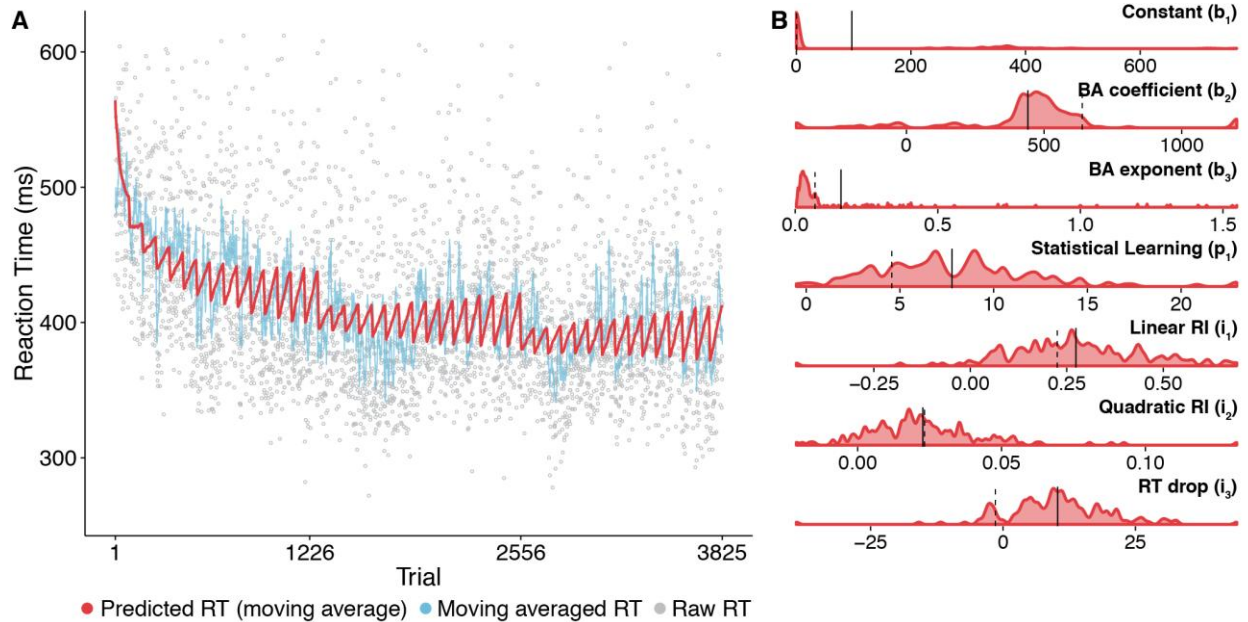


Figure 4. A. Illustration of prediction of the fitted complete model (see Methods) for a typical participant (pink, moving averaged for visual clarity, window size 15). For comparison with measured data, raw (gray points) and moving averaged (light blue, window size 15) reaction times are also plotted. Reaction times shown up to mean plus three standard deviations for visual clarity. **B.** Parameter distributions across all fits across all participants for the complete fits. Vertical dashed lines indicate the fitted parameters on Figure 4A. Vertical solid lines show mean parameter values. It is apparent that each component's parameter comes from a distribution with a positive mean implying that the different components are demonstrably present in the individual reaction time data, not just in the population averages. (BA: basic adaptation, RI: reactive inhibition). The spread of parameters around the mean highlights individual differences across the population of participants.

A reliable measure of learning requires that it is not affected by other factors. We examined the contribution of statistical learning to reaction time data by comparing three models: model SL (parameter p_1), model BA + SL (parameters b_2 , b_3 and p_1) and the full model (see Table 2). Analysis of parameter distributions reveals that the measure of statistical learning (parameter p_1) is invariant to the model choice as demonstrated by the limited change of parameter under

different models (a within-subject ANOVA on the ten estimated values for parameter p_1 per participant per model obtained from the three models showed significant difference, $p < 0.001$, but with minor effect: generalized eta-squared 0.0007, Fig. 5a). As a control, we investigated alternative models and assessed whether independent processes other than statistical learning can be isolated. In particular, parameter of linear reactive inhibition was found to be strongly dependent on the presence of the quadratic reactive inhibition parameter (see Appendix). Because of this degeneracy, reactive inhibition parameters were used jointly to predict reaction time data.

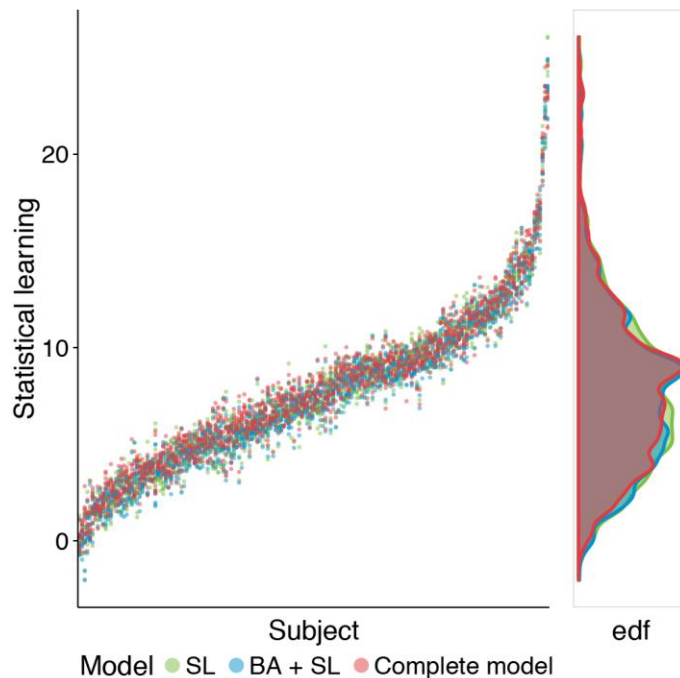


Figure 5. Distribution of fitted statistical learning parameter values (p_1) fits for each participant in different models (green: statistical learning (SL) alone; blue: basic adaptation (BA) and SL; red: complete model). Left panel shows the inferred value of the SL parameter for each participant. Even though there is large variance between participants, the shift between models within subject is low. Right panel is the empirical distribution function (edf) of SL parameters across all participants. A change in edf's would indicate model-dependency of SL parameter. Overlap of edf's indicate the independence of fitted SL parameter of other parameters in the model, suggesting orthogonality.

Predictive power of fitted models

Population-level analysis hinted at processes beyond learning of stimulus statistics. Importantly, a subject-by-subject analysis of the reaction times of human participants could prove that these processes are relevant for predicting individual performances. The predictive power of model variants was tested on cross-validated reaction time data (see Methods and Fig. 6a). We quantified the predictive power of a model by the cumulative distribution of R^2 values over the participant pool. The R^2 value quantifies the residual variance remaining in the test data after removing variance predicted by the model fitted to the train dataset. Later saturation of this measure signifies better predictive power since high R^2 value could be achieved in a larger number of participants when fitting the model. The model that only accounts for experiment-wide changes in reaction times by a power function (blue) serves as a reference for the assessment of model performances. Importantly, inclusion of learning of statistical contingencies (parameter p_1) significantly improved the predictive power of the models, regardless of whether the effects of reactive inhibition were included (solid vs. dashed red lines, paired samples, one-sided t-test for mean R^2 for the models calculated for each participant: $t(179)=-13.72$, $p<0.001$, Cohen's $d=-1.02$) or not (solid vs. dashed blue lines, paired samples, one-sided t-test for mean R^2 values for the models calculated for each participant: $t(179)=-13.27$, $p<0.001$, Cohen's $d=-0.99$).

Regarding the effect of within-block changes in reaction times, the analysis reveals a systematic increase in reaction times (i.e., reactive inhibition) along the span of individual blocks starting over from a baseline level at the beginning of new blocks, thus introducing a periodicity in expected reaction times (Fig 3a). This is confirmed by a significant improvement in predictive power when the reactive inhibition parameter i_1 is included in the model compared to the model

with basic adaptation only (one-sided paired samples t-test on mean R^2 values calculated for each participant and model: $t(179)=37.22$, $p<0.001$, Cohen's $d=2.77$).

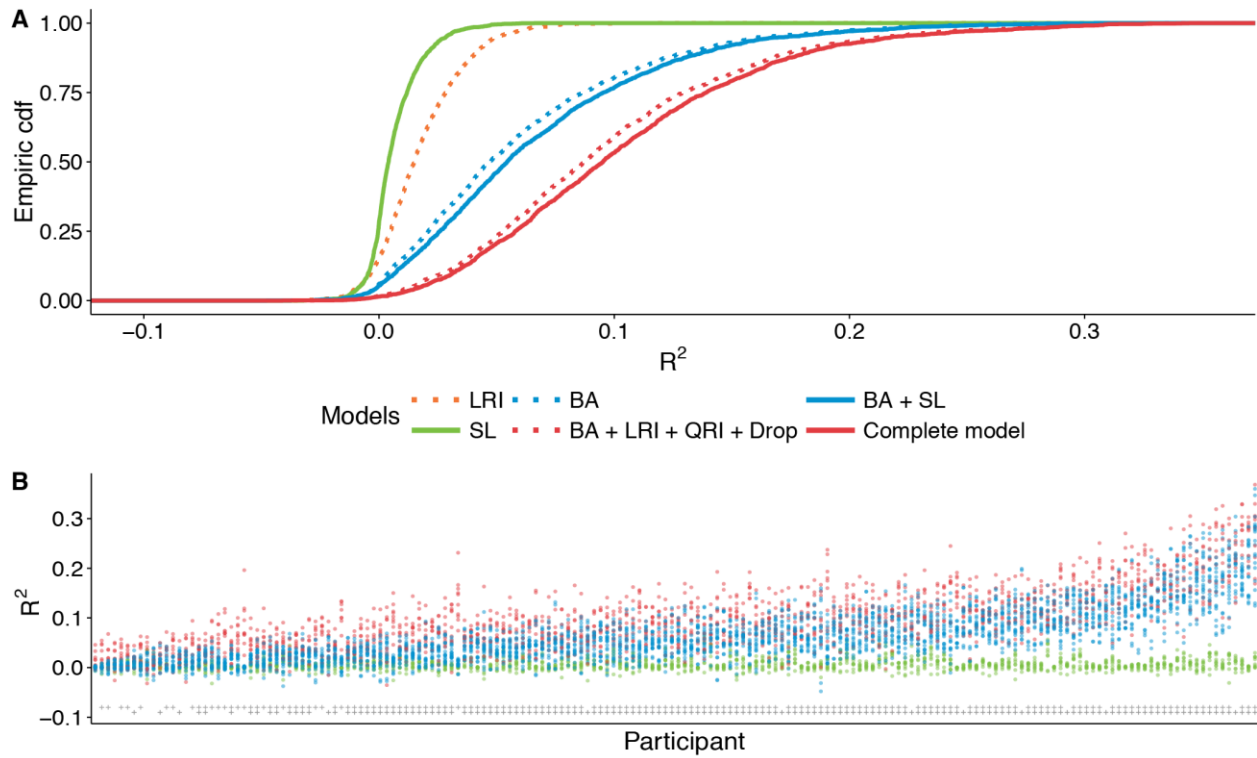


Figure 6. Model comparison. **A.** Predictive power of the fits of different models across the population of participants. Predictive power of a fit on an individual participant is characterized by the R^2 value. Cumulative distribution functions (CDF) of R^2 values quantifies how well the fit performs across the whole population: a CDF with sharp increase at low R^2 values indicates that the model fits the population of participants poorly, while a rightward shift indicates higher R^2 values across the population. Note that even a simplified version of reactive inhibition (LRI) is capable of accounting for more variance in reaction time data than SL (rightward shift of orange line from green line). A complete model containing both SL and reactive inhibition has a larger predictive power than any of the incomplete models. Importantly, the analysis was performed on cross-validated data, therefore this improvement is not due to mere overfitting, but reflecting independent contributions of SL and RI to subject-by-subject performance of the model (see Methods). Three outlier fits (out of 1800 per model) were excluded from the empiric distributions. **B.** Parameter fits for each participant. Colored dots are R^2 values of different models for a particular participant. Gray crosses indicate significant differences of R^2 values (one-sided, paired-sample t-tests,

$p < 0.01$, lower row: difference between dashed blue and green models; upper row: difference between dashed red and solid blue models). Note that predictive powers (R^2) of models are separated for most participants. Models with and without SL parameter are plotted with the same color, participants are ordered according to the mean of predictive power across all fits for visual clarity.

Comparing the predictive power of the models for individuals separately confirms a prominent effect of reactive inhibition (Fig. 6b). Extending the statistical learning model (green on Fig 6b) with basic adaptation (blue) significantly improved model performance for 136 participants (bottom row of crosses on Fig 6b; one-sided, paired samples t-tests on calculated ten R^2 estimates for each model for each participant, significance: $p < 0.01$). Further addition of reactive inhibition significantly improved predictive performance for 163 out of 180 participants (top row of crosses on Fig 6b; one-sided, paired samples t-tests on calculated ten R^2 estimates for each model for each participant, significance: $p < 0.01$; no correction on level of significance is used and hence there is an expected false positive rate of 0.01, resulting in an expected number of ~4 false positives overall for the two tests on the 180 participants).

In summary, reactive inhibition acts independently from statistical learning. Further, subject-by-subject analysis reveals that reactive inhibition is a significant source of variability in the majority of participants and its contribution exceeds that of statistical learning.

Separating the effect of statistical learning from reactive inhibition

A crucial test of the way reactive inhibition affects experimental results is to investigate whether accounting for reactive inhibition on a subject-by-subject basis helps to detect patterns dependent on learning the statistical contingencies. In the ASRT task, statistical learning was assessed by comparing reaction times for high- vs. low-probability triplets (Table 1 and Fig. 2). We aimed at quantifying the amount of variability in the data that can be explained by statistical learning with or without accounting for reactive inhibition (Fig. 7a, see Methods). Improvement

of model performance (R^2) induced by statistical learning component (parameter p_1) relative to a baseline was assessed in two ways. First, a baseline performance measure was established for the raw data (using a constant model); second, a baseline was established by compensating for the effects of reactive inhibition as well as for basic adaptation (full model without statistical learning, see Methods). Analysis of individual participants' data showed a clear advantage of the model using reactive inhibition, thus confirming that statistical learning can account for a larger part of residual variance of reaction times when added to a model with reactive inhibition compared to the case when it is added to a model without reactive inhibition (Fig. 7a, one-sided, paired samples t-test on mean model improvement scores detailed in Methods, calculated for each participant and each model pair: BA to BA + SL, and Full model without SL to Full model: $t(179)=-11.29$, $p<0.001$, Cohen's $d=-0.84$). A direct comparison of the patterns in learning the triplet structure in the ASRT task can reveal the efficiency of statistical learning in accounting for differences in reaction times. Variance of responses to high-probability and low-probability triplets was calculated for raw data as well after taking the remainder of the predictions produced by a model with RI and basic adaptation. The analysis reveals lower variance and thus more prominent learning effect for data adjusted for reactive inhibition and basic adaptation effects (Fig. 7b, for high-probability triplets $t(179)=24.11$, $p<0.001$, Cohen's $d=1.80$, for low-probability triplets: $t(179)=21.57$, $p<0.001$, Cohen's $d=1.61$). Reduction of variance can be as high as 20% and on average 8.8% improvement was achieved.




Figure 7. Error reduction of models. **A**, Comparing error reduction of statistical learning parameter p_1 when being added to a model with only basic adaptation included (x axis) or when being added to a full model including basic adaptation and reactive inhibition parameters (y axis). For 163 participants, statistical learning can account for a larger part of residual variance of reaction times when added to a model with reactive inhibition compared to when it is added to a model without reactive inhibition **B**, Reduction in variance of reaction times when using all factors except SL for prediction, taking the variance of the raw data as a baseline. The reduction is calculated and visualized separately for high- and low-probability triplets. Participants are ordered by mean reduction. Eliminating factors not attributed to learning allows for more accurate measurement of learning.

Between-session shifts in reaction times introduced by reactive inhibition

A significant drop in mean reaction times can be observed after 3-5 minute breaks introduced between ASRT training sessions (Fig. 3a, c). This apparent change in performance has often been interpreted as memory consolidation (for a review, see Walker et al., 2004). In more refined analyses, consolidation is defined as a discontinuity in fitted reaction time curves (Pan & Rickard, 2015). Between-session RT drop is explicitly represented in our parametric model (i_3).

However, our analysis has shown that despite a pronounced drop observed in reaction time data at session breaks (Fig. 4a), i_3 can take a value close to zero (Fig. 4b, bottom panel). Importantly, using a model with basic adaptation and quadratic reactive inhibition in place is capable of producing apparent discontinuities in reaction time curves at session breaks. This discontinuity can emerge without an explicit drop in reaction times, solely as a consequence of the reset of the *slope* of reactive inhibition. We assessed whether a model that includes basic adaptation, within-block (linear) and between-block (quadratic) reactive inhibition (BA + LRI + QRI in Table 2) can account for the session gaps by comparing it to a model with constant RT (Table 2). We calculated five-block averages of residuals and demonstrated a substantial decrease in the difference between end-of-session and start-of-session averages, resulting in a residual session gap close to zero at the second break (Fig 8a). This result is confirmed by directly plotting residuals of subject responses under the two models (Fig. 8c), which demonstrates that a model which accounts for both within-block (linear) and between-block (quadratic) RI but lacks an explicit RT shift at session boundaries effectively eliminates the session gap. Hence, a model that includes basic adaptation, linear and quadratic reactive inhibition may challenge results that are in favor of consolidation even when a discontinuity in reaction times is present after a fitted power function curve.

We also tested whether the effect of consolidation is observable in the *statistical learning* measure on raw data (Figure 8b). There was no significant improvement on a group level either for the first session gap (paired-sample, one-sided t-test on the difference in mean RTs for high- and low-probability triplets calculated for each participant for blocks 13-15 and 16-18: $t(179) = -1.0116$, $p = 0.1565$) nor the second session gap ($t(179) = 1.5725$, $p = 0.9412$).

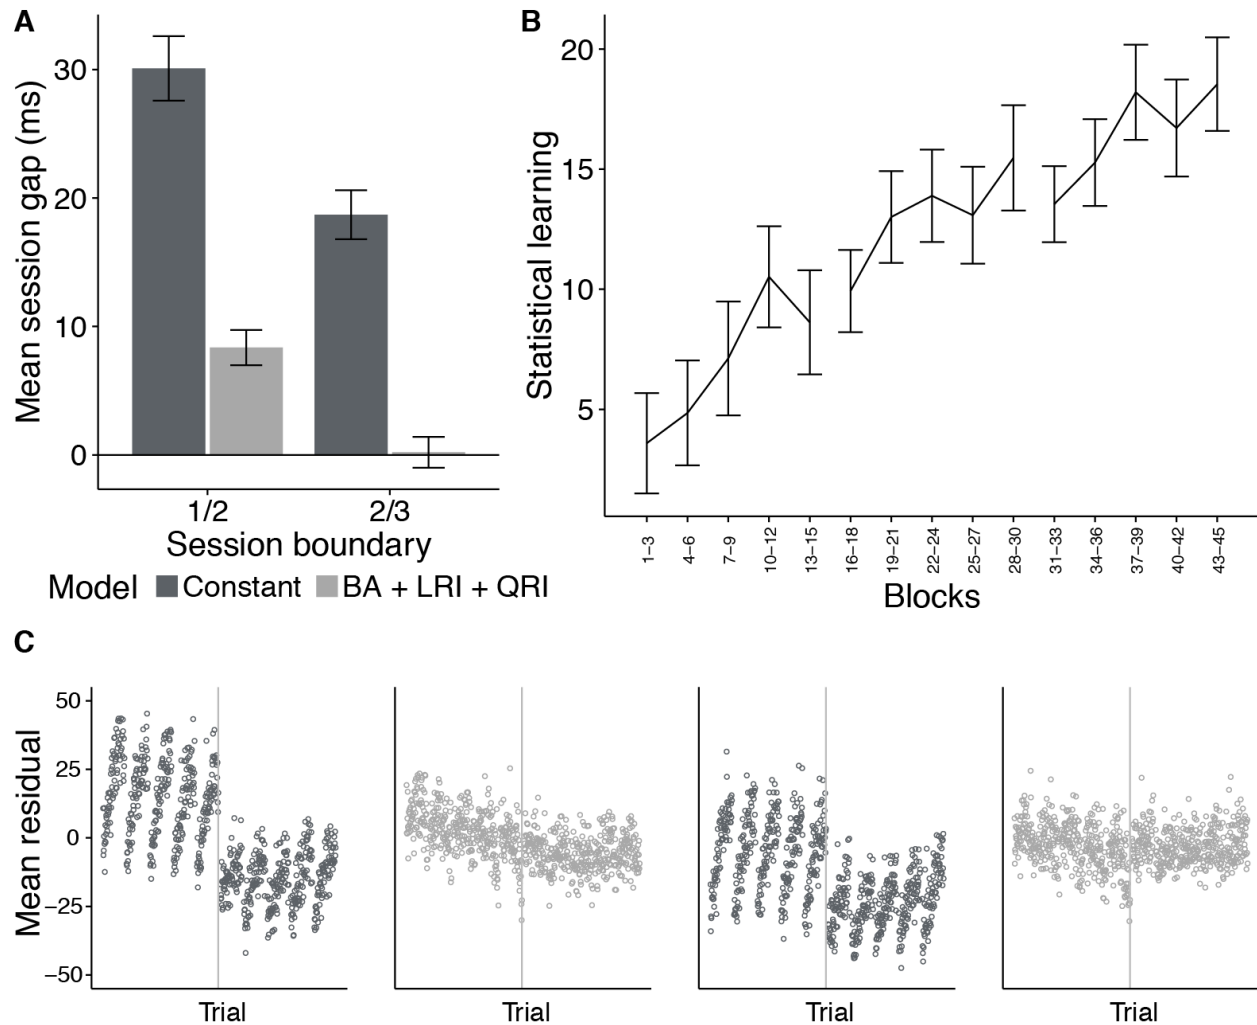


Figure 8. Model residuals. **A**, Drop in mean residual reaction times for last and first five blocks of a training session. Constant model (only parameter b_1) predicts mean of training set. The second model also includes baseline adaptation (parameters b_2 and b_3), linear and quadratic reactive inhibition (parameters i_1 and i_2). **B**, Time course of the statistical learning measure (difference between mean reaction times to high- and low-probability triplets) averaged across all participants. No significant performance improvement on session boundaries (at blocks 13-15 / 16-18 & 28-30 / 31-33). Error bars show S.E.M. **C**, Mean residuals averaged over participants for each trial in the last and first five blocks of a training session. After removing the effects of basic adaptation and reactive inhibition, there is still a significant change between the first two training sessions when five-block averages are calculated (**A**). This difference, however, no longer seems to manifest in a discontinuity.

Discussion

We presented a phenomenological model of reaction times to dissociate the effects of reactive inhibition and learning in an implicit learning paradigm. We have shown that reactive inhibition is a significant effect across a population of participants and can even have a more pronounced effect than learning. Importantly, relying solely on individual participants' data rather than population averages, reactive inhibition was shown to be present for 90% of the participants. We have shown on cross-validated data that prediction of reaction times based on the statistical learning model improved for 90% of participants when reactive inhibition was taken into account. Although the parameters characterizing reactive inhibition jointly determined the within-block, between-block and between-session effects of RI on reaction time data, we found that this modulation was independent of the modulation caused by statistical learning.

It is unknown whether reactive inhibition is due simply to fatigue, a decline in engagement in the task or whether it hints at some more elaborate process. Here we introduced a simple parametric form inspired by the population characteristics of reactive inhibition that implies a linear increase in expected reaction times within a block of trials. This form was sufficient to significantly improve predictive power even on a subject-by-subject basis. However, the linearly modeled performance decline appearing in population-averaged data does not necessarily imply a suitable model form for the dynamics of performance in individuals. For example, the decline in averaged performance could result from arbitrating between purely stimulus-driven reactions and a strategy where an internal model exploiting past experiences is invoked. This could be modeled as a stochastic mixture process where late trials within a single block have a larger probability of resulting from strategy not relying on learning, one that is purely stimulus-driven. Such a mixture process could give rise to a dynamics different from linear RI but result in

linear population-averaged dynamics. Whether a more intricate model such as this mixture process is better able to capture the individual variations requires further investigation.

Rickard et al. (2008) produced evidence that implied that reactive inhibition was a strongly confounding factor when assessing learning progress. However, as we demonstrated here, the previously established measure of statistical learning of Howard & Howard (1997) (also found in Nemeth et al., 2013; Janacsek et al., 2012; Hallgató et al., 2013) which we adopted here is at least partially immune to it. This is a consequence of the fact that while individual reaction times are affected by RI, statistical learning is assessed by comparing reaction times of interleaved trials: learning progress is evaluated by taking the difference of mean reaction times to stimuli of different predictability which on the relevant timescale appear independently of stimulus location within the training session. If during averaging one set of responses is from the first half-block whereas the other is from the second; or if averages of different blocks are compared, an illusory improvement can arise due solely to the buildup of reactive inhibition. Such illusory improvement has been shown in studies focusing on sleep-dependent consolidation (see Rickard, 2008, Brawn et al. 2010, Pan & Rickard, 2015) leading to inaccurate conclusions about the role of sleep in consolidation. In line with this conclusion, there was no statistically significant improvement in our statistical learning measure during the offline periods (session gaps). In light of this, we propose that for analyses where RI is not taken into account explicitly, performance measures should be constructed in such a way that the effect of reactive inhibition is transformed out, e.g. by subtracting mean RTs of interleaved sets of trials. If application of such difference measures is not an option, there are multiple ways to address the issues raised by reactive inhibition. First, our approach provides a method to eliminate the effects of RI. Second, Pan and Rickard (2015) proposed an alternative, by using shorter training blocks with longer rest periods between blocks of trials. The latter has the advantage that it may accommodate a

wide range of possible parametric forms of RI (possibly different from ours), but has the disadvantage that it only attenuates the effects of RI.

In this study, we have not analyzed the acquisition of the statistical structure inherent in the data, relying instead on the asymptotic appearance frequencies of stimulus-triplets which had previously been shown to be useful measures for differentiating between different age-groups and patient populations (Nemeth et al., 2013; Janacsek et al., 2012; Nemeth et al., 2015; Barnes et al., 2010). While using reaction times for probing the characteristics of the internal model of participants is of great interest, such a treatment is predicated on filtering out the non stimulus-dependent part of the variance. Filtering the data - so that it is suitable for making inferences concerning the internal representations or effectively predicting reaction times - requires a model which does not merely describe the population averaged appearance of reactive inhibition but is also capable of capturing individual deviations. The model formulated in this study might provide means to capture the evolution of the internal model .

Problem solving, decision making, and learning are intricately intertwined in complex everyday tasks. In such situations earlier experience and novel evidence are integrated to reach efficient performance. Degraded attention and fatigue are extensively studied processes that act against the various forms of learning but the dynamics of these processes are considered to be relatively slow. In contrast, our results show that reactive inhibition interferes with learning at a considerably shorter time scale: blocks of trials that have a clearly visible effect extend only across a time window of approximately 90 seconds. This short time scale highlights the possible importance of reactive inhibition in complex tasks: within the time frame of the execution of a complex everyday task reactive inhibition can have an effect on decisions. Here we developed a model that was able to dissect reactive inhibition from inference and learning, which is critical for assessing and understanding the dynamics of acquisition of information on a subject-by-subject

basis. This model can be used - with modifications to fit the given task - to other perceptual-motor learning and decision making tasks in which the decisions have to be made consecutively. Moreover, reactive inhibition is especially important in the context of cognitive disorders affecting learning, including ADHD and autism (Nemeth et al., 2010). We expect that a subject-by-subject paradigm for analyzing performance during task execution can help to discover the components of cognitive computations that fail at these disorders.

Acknowledgements

This research was supported by the Research and Technology Innovation Fund, Hungarian Brain Research Program (KTIA NAP 13-2-2015-0002); Hungarian Scientific Research Fund (OTKA NF 105878); by a Lendület Award of the Hungarian Academy of Sciences (G.O., D.G.N., B.T.); and Janos Bolyai Research Fellowship of the Hungarian Academy of Sciences (to K. J.).

References

Barnes, K. A., Howard Jr, J. H., Howard, D. V., Kenealy, L., & Vaidya, C. J. (2010). Two forms of implicit learning in childhood ADHD. *Developmental neuropsychology*, 35(5), 494-505.

Brawn, T., Fenn, K., Nusbaum, H., & Margoliash, D. (2010). Consolidating the Effects of Waking and Sleep on Motor-Sequence Learning. *The Journal of Neuroscience*, 30(42), 13977–13982.

doi:10.1523/JNEUROSCI.3295-10.2010

Carpenter, R. H. S., & Williams, M. L. L. (1995). Neural computation of log likelihood in control of saccadic eye movements. *Nature*, 377(6544), 59-62.

Cleeremans, A., & Dienes, Z. (2008). Computational models of implicit learning. *Cambridge handbook of computational psychology*, 396-421.

Csabi, E., Varszegi-Schulz, M., Janacsek, K., Malecek, N., & Nemeth, D. (2014). The consolidation of implicit sequence memory in obstructive sleep apnea. *PloS one*, 9(10), e109010.

Diekelmann, S., & Born, J. (2007). One memory, two ways to consolidate? *Nature neuroscience*, 10(9), 1085-1086.

Grisson, S., & Strayer, D. L. (2001). Negative priming and perceptual fluency: More than what meets the eye. *Perception & psychophysics*, 63(6), 1063-1071.

Hallgató, E., Győri-Dani, D., Pekár, J., Janacsek, K., & Nemeth, D. (2013). The differential consolidation of perceptual and motor learning in skill acquisition. *Cortex*, 49(4), 1073-1081.

Houghton, G., Tipper, S. P., Weaver, B., & Shore, D. I. (1996). Inhibition and interference in selective attention: Some tests of a neural network model. *Visual cognition*, 3(2), 119-164.

Howard Jr, J. H., & Howard, D. V. (1997). Age differences in implicit learning of higher order dependencies in serial patterns. *Psychology and aging*, 12(4), 634-56.

Hull, C. (1943). *Principles of behavior: an introduction to behavior theory* (p. 278). Oxford, UK: Appleton-Century-Crofts.

Janacsek, K., Fiser, J., & Nemeth, D. (2012). The best time to acquire new skills: age- related differences in implicit sequence learning across the human lifespan. *Developmental science*, 15(4), 496-505.

- Kelso, J. S. (2014). *Human motor behavior: An introduction* (pp. 3-20). NJ, Hillsdale: Psychology Press.
- Kohavi, R. (1995). A study of cross-validation and bootstrap for accuracy estimation and model selection. *IJCAI'95 Proceedings of the 14th international joint conference on Artificial intelligence*, 2, 1137-1143.
- Nemeth, D., Janacsek, K., & Fiser, J. (2013). Age-dependent and coordinated shift in performance between implicit and explicit skill learning. *Front. Comput. Neurosci*, 7(147), 10-3389.
- Nemeth, D., Janacsek, K., Balogh, V., Londe, Z., Mingesz, R., Fazekas, M., ... & Vetro, A. (2010). Learning in autism: implicitly superb. *PloS one*, 5(7), e11731.
- Nemeth, D., Janacsek, K., Király, K., Londe, Z., Németh, K., Fazekas, K., ... Csányi, A. (2013). Probabilistic sequence learning in mild cognitive impairment. *Frontiers in Human Neuroscience*, 7, 318. <http://doi.org/10.3389/fnhum.2013.00318>
- Pan, S. C., & Rickard, T. C. (2015). Sleep and motor learning: is there room for consolidation? *Psychological bulletin*, 141(4), 812-34.
- Reber, A. S. (1989). Implicit learning and tacit knowledge. *Journal of experimental psychology: General*, 118(3), 219-235.
- Rickard, T. C., Cai, D. J., Rieth, C. A., Jones, J., & Ard, M. C. (2008). Sleep does not enhance motor sequence learning. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 34(4), 834-42.
- Robertson, E. M., Pascual-Leone, A., & Miall, R. C. (2004). Current concepts in procedural consolidation. *Nature Reviews Neuroscience*, 5(7), 576-582.

Schneider, W., Eschman, A., & Zuccolotto, A. (2012). E-Prime Reference Guide. Pittsburgh: Psychology Software Tools, Inc.

Song, S., Howard, J. H., & Howard, D. V. (2007). Sleep does not benefit probabilistic motor sequence learning. *The Journal of Neuroscience*, 27(46), 12475-12483.

Stelmach, G. E. (1969). Efficiency of motor learning as a function of intertrial rest. *Research Quarterly. American Association for Health, Physical Education and Recreation*, 40(1), 198-202.

Stone, M. (1977). An Asymptotic Equivalence of Choice of Model by Cross-Validation and Akaike's Criterion. *Journal of the Royal Statistical Society. Series B (Methodological)*, 39(1), 44-47.

Tipper, S. P., Meegan, D., & Howard, L. A. (2002). Action-centred negative priming: Evidence for reactive inhibition. *Visual Cognition*, 9(4-5), 591-614.

Walker, M. P., & Stickgold, R. (2004). Sleep-dependent learning and memory consolidation. *Neuron*, 44(1), 121-133.

Walker, M. P. (2005). A refined model of sleep and the time course of memory formation. *Behavioral and brain sciences*, 28(1), 51-64.

Appendix

Conditional probabilities of third elements of triplets

Assume, for example, that the structure of our sequence is 2 – r – 1 – r – 3 – r – 4 – r, as in the example in Methods. The conditional distribution of the next element (third of the triplet) having seen 2-1 in the first two trials of the triplet can be calculated as follows:

$$\begin{aligned} P(\text{Third} = x \mid \text{First} = 2, \text{Second} = 1) = \\ P(\text{Third} = x \mid \text{First} = 2, \text{Second} = 1, \text{Third is random}) * P(\text{Third is random}) + \\ P(\text{Third} = x \mid \text{First} = 2, \text{Second} = 1, \text{Third is not random}) * P(\text{Third is not random}) \end{aligned}$$

Note that when Third element is random, Third element is any of 1, 2, 3, 4 with equal probability, whereas when Third is not random, it takes value 1 with probability 1. Assuming one does not know at what position they are in the sequence, $P(\text{Third is random}) = 0.5$. Hence, we obtain the probability distribution presented in Table 1.

Model parameter interactions

Dependencies and degeneracies among parameters were analysed using models in which some of the components were removed (for example, see Figure 5). Changes in parameters in the original and truncated model reveal such interactions between model parameters. Besides

tracking changes in individual fits, we also illustrate parameter interactions by constructing the distribution of parameters over participants. This provides a compact representation of parameter interactions. Dependencies were found among the first three parameters (b_1 , b_2 and b_3) describing the basic adaptation component. Predicted reaction times proved to be largely invariant to particular joint changes in this set of parameters. Consequently, in our subsequent analyses we did not investigate further the precise values of this set of parameters.

The three parameters of reactive inhibition (i_1 , i_2 and i_3) were also interdependent. As visible on Figure A1, i_1 is interpreted in a slightly different way in models where i_2 is present compared to when it is not. This is simply due to the fact that i_1 describes the increase of RT block between individual trials in the first training block and i_2 is its increment between blocks. Hence the value of i_1 in a model with only linear reactive inhibition will match the value of reactive inhibition in a model with quadratic reactive inhibition in the middle of a training session ($i_1 + 7.5 \cdot i_2$). This interaction is supported by a significant change in parameter i_1 when comparing its estimated values between-model and within-subject (paired-sample, two-sided t-test: $t(179)=13.149$, $p<0.001$, Cohen's $d=0.98$).




Figure A1. Distribution of fitted linear reactive inhibition parameter values (i_1) as seen for the statistical learning parameter (p_1) in Figure 5. Models investigated are the linear reactive inhibition-only model (LRI, orange) and the complete model (red). The shift in the empirical distribution due to altered model indicates that -- in contrast with the case of the SL parameter -- introducing additional variable changes the interpretation of the previously present variable. For visual clarity, participants are ordered according to the average of the fitted parameter across the investigated models.