



HAL
open science

Visual Predictions Operate on Different Timescales

Álvaro Darriba, Sandrien van Ommen, Yi-Fang Hsu, Florian Waszak

► **To cite this version:**

Álvaro Darriba, Sandrien van Ommen, Yi-Fang Hsu, Florian Waszak. Visual Predictions Operate on Different Timescales. *Journal of Cognitive Neuroscience*, 2021, 33 (6), pp.984-1002. 10.1162/jocn_a_01711 . hal-03438918

HAL Id: hal-03438918

<https://hal.science/hal-03438918v1>

Submitted on 22 Nov 2021

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Title: Visual predictions operate on different timescales

Running title: Predictions on different timescales

Authors names and affiliations:

Álvaro Darriba (alvarodarriba@gmail.com; alvaro.darriba@parisdescartes.fr)^a

Sandrien Van Ommen (sandrienvanommen@gmail.com; sandrien.vanommen@parisdescartes.fr)^a

Yi-Fang Hsu (yi-fang.hsu@cantab.net; yifanghsu@ntnu.edu.tw)^{b,c}

Florian Waszak (f.waszak@gmx.net; florian.waszak@parisdescartes.fr)^{a,d}

^a Université de Paris, CNRS, Integrative Neuroscience and Cognition Center, F-75006 Paris, France

^b Department of Educational Psychology and Counselling, National Taiwan Normal University, 10610, Taipei, Taiwan

^c Institute for Research Excellence in Learning Sciences, National Taiwan Normal University, 10610, Taipei, Taiwan.

^d Fondation Ophtalmologique Rothschild, 25 rue Manin, F-75019 Paris, France

Corresponding author:

Álvaro Darriba (alvarodarriba@gmail.com)

Integrative Neuroscience and Cognition Center (INCC)

CNRS (UMR 8002) – Université de Paris

45 rue des Saints-Pères, 75006 Paris, France

Abstract

Humans live in a volatile environment, subject to changes occurring at different timescales. The ability to adjust internal predictions accordingly is critical for perception and action. We studied this ability with two EEG experiments in which participants were presented with sequences of four Gabor patches, simulating a rotation, and instructed to respond to the last stimulus (target) to indicate whether or not it continued the direction of the first three stimuli. Each experiment included a short-term learning phase in which the probabilities of these two options were very different ($p = .2$ vs $p = .8$, rules A and B, respectively), followed by a neutral test phase in which both probabilities were equal. Additionally, in one of the experiments, prior to the short-term phase participants performed a much longer long-term learning phase where the relative probabilities of the rules predicting targets were opposite to those of the short-term phase. Analyses of the reaction times and P3 amplitudes showed that in the neutral test phase, participants initially predicted targets according to the probabilities learned in the short-term phase. However, while participants not pre-exposed to the long-term learning phase gradually adjusted their predictions to the neutral probabilities, for those who performed the long-term phase, the short-term associations were spontaneously replaced by those learned in that phase. This indicates that the long-term associations remained intact while the short-term associations were learned, transiently used, and abandoned when the context changed. The spontaneous recovery suggests independent storage and control of long-term and short-term associations.

Keywords

Electroencephalography (EEG), Multiple timescales, P300, Prediction.

Introduction

The ability to make predictions about future states of the environment allows humans to adapt their perception and optimize their behaviour. According to predictive coding models (Friston, 2005; Rao & Ballard, 1999), predictions are represented in the brain as probability distributions that are continuously compared with actual evidence and adjusted correspondingly. Predictions are primarily based on general knowledge and experience, that is, on global probabilities about how events succeed in a given context. These probabilities can be learned over time, improving our adaptation to the environment as we are exposed to it. Such an adaptation may be easily achieved under stable contextual conditions. However, we live in a mutable environment in which the relationships between predictive and predicted events are subject to changes. Moreover, these changes occur at different timescales, from very transient to very long-lasting, and the ability to adapt our predictions consequently is critical for optimizing perception and action. The aim of the present study was to shed light on the adaptability of predictions by investigating into the dynamics of the acquisition of relationships between predictive and predicted events learned at different timescales.

The general idea that the nervous system is adapted to the statistical properties of the environment is a longstanding principle in Neuroscience (Fiser et al., 2010). This principle is related to an equally long tradition emphasizing the role of prediction in perception and cognition, a view that has roots in Helmholtz's work in late XIX century (Dayan et al., 1995; Helmholtz, 2013; Swanson, 2016). Over the last two decades, this "predictive brain" view has regained considerable strength, leading to the emergence of theoretical proposals such as the hierarchical predictive coding models (Friston, 2005; Rao & Ballard, 1999). Predictive coding models characterize the brain as an organ fundamentally dedicated to actively infer the causes of the inputs it receives and to predict future inputs accordingly. Furthermore, these models claim that all brain-based behaviours, including those of high-level cognition, can be explained in terms of suitably organized hierarchical prediction processes (Adams et al., 2013; Thornton, 2017). Predictive processes are generative, that is, they are based on prior beliefs about how causes interact and lead to a particular input. This information is estimated from sensory data, and thus needs to be learned. This learning process seeks inferring the causes of the sensory inputs by minimising the difference between the actual sensory data and the sensory data predicted on the basis of preceding events, i.e., prediction error (PE). In predictive coding, backward projections from one hierarchical level to its subordinate level furnish predictions of the lower level's representation, while reciprocal forward projections convey PE that report the difference between the representation and the prediction (Bastos et al.,

2012). Error signals received by the higher level are then used to correct its representation so that its predictions progressively improve. This recurrent exchange of signals continues until PE is minimised, at which point the hierarchy contains a representation as accurate as possible of the causes of sensory input.

Gathering evidence to uncover the probabilistic structure of the environment and reduce PE necessitates time. A consequence of the continuous PE-reduction process is that the more stable the context is the faster an optimal predictive state is achieved, i.e., the faster the associations between predictive and predicted inputs are learned and hence the faster PE is reduced. In experimental sets, time needed for learning can indeed be very short, taking only a few trials under particularly stable conditions (Bekinschtein et al., 2009; Garrido et al., 2009; Hsu et al., 2015; Schadow et al., 2009; Todorovic & Lange, 2012), but increases significantly when apprehending the probabilistic structure of the environment requires more complex computations (Bidet-Caulet et al., 2012; Domenech & Dreher, 2010). Once learned, the associations between predictive and predicted events are strengthened in the long term when participants are exposed to a given context beyond what is necessary for unveiling the probabilistic relationships between those events. Moreover, our experience depends not only on immediate information from the environment but also on our prior knowledge and expectations. Human behaviour is shaped by past experience on multiple timescales, and successful performance in our dynamic environment critically depends on the brain's capacity to adapt their predictions based not only on current contextual information, but also on past experience (Sohoglu & Davis, 2016). In other words, the brain needs to be able to learn from transient conditions and to adapt its predictions consequently without weakening the associations between predicted and predictive inputs that hold valid on a longer-term basis.

The current study addresses this issue by assessing PE in a dynamic learning environment as described in detail below. To do so we harness the fact that a number of ERP components have been linked to PE (Garrido et al., 2009; Stefanics & Czigler, 2012; Todorovic & Lange, 2012). This is the case of the P3, a positive deflection in the ERP waveform peaking between 300 and 800 ms and with a broad but varying topographical distribution depending on the task employed. P3 has been related to a wide variety of processes, including context updating (Donchin & Coles, 1988), decision confidence (Sawaki & Katayama, 2006), evidence accumulation and the updating of perceptual evidence (O'Connell et al., 2012), all arguably related to prediction processing. In the context of prediction research, modulations of P3 amplitude have been extensively related to mismatch, surprise, and novelty processing, and consequently to PE (Ehinger et al., 2015; Feldman & Friston, 2010;

Kolossa et al., 2015; Mars et al., 2008; Waszak & Herwig, 2007) and to the updating of an internal prediction model (Fischer & Ullsperger, 2013). More specifically, P3 has been linked to high-level error associated not to specific sensory features but to the processing of global contextual deviations (Bekinschtein et al., 2009). Mars and colleagues, for instance (Mars et al., 2008), instructed participants to learn the associations between four arbitrary stimuli and four response buttons and later presented them with a series of experimental blocks in which participants were required to respond to each stimulus with the previously associated button as quickly and accurately as possible. Unbeknownst to participants, the probability of the occurrence of each event was manipulated between blocks such that the relative probabilities of events were either low, medium or high. Their results demonstrated that trial-by-trial fluctuations in P3 amplitudes could be explained in terms of participants keeping track of the global probabilities of visual events, so that its amplitude was reduced or enhanced as a function of surprise or, in other words, PE. In another experiment, Wacongne and colleagues (Wacongne et al., 2011) employed an auditory paradigm to dissociate two types of predictions, based on local probabilities versus global rules. In a given block, a frequent sequence of five tones was presented (in 75% of trials), interspersed with rare violations (in 15%) in which the frequency of the fifth tone deviates from the expected, and with rare omissions (10%) in which the fifth tone is simply omitted. Authors found that in blocks in which the frequent sequences consisted of four identical tones followed by a deviant one, this last tone elicited a mismatch negativity (MMN), that is, an early PE response related to local regularity violations, but not a P3. However, in sequences in which the last tone was identical to the four preceding ones, this last tone did not elicit an MMN, but an enhanced P3 response. Authors concluded that two different levels of predictions operate in that context: a first low-level expectation, based on local transition probabilities and reflected by the MMN, and a second, higher-level prediction, based on the knowledge about the global overall rule or pattern followed by the stimuli which, when violated, elicits a PE response reflected in the P3 component (Wacongne et al., 2011). Hence, due to its sensitiveness to PE, P3 amplitude offers insight into the predictions that participants are making at any given moment and, consequently, into their learning about the statistical structure of the environment and the use they make of that information to adapt to changes in that environment over time.

Here we used event-related potentials (ERP) and reaction times (RT) to investigate the adaptability of predictions learned at different timescales. Specifically, we focused on the P3 ERP component as a measure of PE and thus as a means to track the learning and adaptation of participants to different contexts. The study comprised two experiments in which participants were presented with sequences of four succeeding Gabor

patches, simulating either a clockwise or an anticlockwise rotation, and were instructed to respond to the last stimulus (target) by pressing one of two buttons to indicate whether it followed or not the direction indicated by the first three stimuli. Each experiment included a short-term learning phase in which the probabilities of the targets following or not following the direction indicated by the first three stimuli were very different ($p = .2$ vs. $p = .8$, designated as predictive rules A and B, respectively), followed by a neutral test phase in which both rules predicted the targets with equal probability. Additionally, in one of the experiments participants performed a much longer long-term learning phase, prior to the other short-term learning and test phases, in which the relative probabilities of the two possible rules predicting the targets were opposite to those of the short-term learning phase (p rule A = $.8$, p rule B = $.2$). In every learning phase, targets presented according to the least predictive rule should generate a larger PE response, measured as the amplitude of the P3 component, and yield slower RTs. Moreover, if long-term predictions remain intact while participants learn the short-term associations, in the neutral test phase a recovery of the PE response observed in the long-term learning phase should occur for participants who were pre-exposed to that phase. For participants not pre-exposed, however, differences in PE responses should simply disappear in the neutral test phase. This pattern of results would parallel those demonstrated in perceptual (Bao & Engel, 2012b; Vul et al., 2008), motor (Kording et al., 2007; Smith et al., 2006) and Pavlovian (Bouton, 1993) learning, and would indicate that predictions valid at different timescales are independently stored and controlled.

Materials and Methods

Participants

40 subjects (27 females, 26.2 ± 3.8 years of age) took part in the study and received monetary compensation for their participation. All participants were right-handed students with normal or corrected-to-normal vision and reported no history of neurological or psychiatric disorders. Informed consent was obtained from them. Experimental procedures were undertaken in accordance with the Declaration of Helsinki and with the approval by the Comité de Protection des Personnes Ile de France II. Participants were randomly assigned to either Experiment 1 (20 participants, 13 females, 25.7 ± 3.9 years of age) or Experiment 2 (20 participants, 14 females, 26.4 ± 3.6 years of age).

Stimuli and trials

Visual stimuli were presented on a 27-in. 60 Hz LCD monitor. Stimuli were presented against a grey background at the centre of the screen and consisted of displays containing a Gabor patch (.4 Michelson contrast, 7 cycles, 4° visual angle), presented within the limits of a central fixation box (5° visual angle) and presented with one of eight possible orientations with respect to the vertical meridian (0°, 45°, 90°, 135°, 180°, 225°, 270°, 315°). A small dark grey dot was placed on the internal edge of each Gabor patch to increase the possible number of the rotation's starting and ending points (otherwise there would be only four possible orientations, given that 0° and 180°, 45° and 135°, 90° and 270°, and 225° and 315° would be identical) to help participants differentiate between supplementary orientation angles and to favour their engagement in the task (Figure 1). Trials consisted of sequences of four succeeding Gabor patches. The first three succeeded rapidly (duration 100 ms, 166 ms SOA) and were presented in a manner in which they simulated either a clockwise or anticlockwise rotation in 45° steps. Both rotation directions were equally probable. The initial point of this rotation (i.e., the orientation of the first stimulus in the sequence) was randomized. These first three Gabor patches were followed after a 1000 ms delay by a fourth one that could either be presented 45° following the same rotation direction indicated by the first three or be presented 45° against the rotation direction. Participants were instructed to respond to that fourth stimulus by pressing one of two buttons to indicate whether this target stimulus followed or not the direction indicated by the first three. Participants were encouraged to respond as fast as possible, and given a maximum of 1000 ms to do so. The association between response button (left, right) and response meaning (same, opposite direction) was counterbalanced between participants. Intertrial interval (ITI), defined as the interval between the presentation of a target on a given trial and the presentation of the first stimulus in the sequence of the next trial, was set to 2500 ms. The stimuli were created, and the experiments written, using the Psychophysics Toolbox (Psychtoolbox-3) extensions (Brainard, 1997; Kleiner et al., 2007)

Procedures

Experiment 1

Experiment 1 consisted of two phases, containing 600 trials each. In the first phase (short-term learning), targets could be presented either following the direction indicated by the predictive sequence (Rule A) or the opposite direction (Rule B). The predictive value of rules A and B was different (.2 and .8, 120 and 480 trials, respectively). In the second phase (test), rules A and B predicted targets with equal probability (.5, 300 trials each). EEG activity was recorded while participants performed in both phases.

Experiment 2

Experiment 2 was similar to Experiment 1 regarding the short-term learning and test phases explained above, but in this case those phases were preceded by a long-term learning phase. This phase was much longer than the short-term learning and test phases, and was divided into two sessions to reduce participants' fatigue. All through the long-term learning phase (sessions I and II) the predictive values of rules A and B were the opposite of those in the short-term learning phase. In the first session (long-term learning Session I) participants were presented with 1000 trials (800 presented according to one rule, 200 according to the other), and the EEG was not recorded. The second session (long-term learning Session II) started 90 minutes after the end of the first one. In this second session subjects were presented with 600 additional trials (480 according to one rule, 120 according to the other), while their EEG activity was recorded. After this two-session long-term learning phase, participants performed a short-term learning phase and a test phase similar to those in Experiment 1.

EEG recording and pre-processing

Continuous EEG data (0.1–250 Hz band-pass) were collected from 60 actiCAP EEG electrodes (BrainProducts GmbH) mounted on an elastic cap. EEG electrodes were placed following the extended 10–10 position system (Acharya et al., 2016), and were referenced to right mastoid. Four additional electrodes were placed above and below the left eye and on the outer canthi of both eyes to monitor blinks and eye movements. EEG and EOG data were collected using the PyCorder system and actiCHamp amplifiers (BrainProducts GmbH, Gilching, Germany) in DC recording mode with a sampling rate of 2000 Hz.

EEG data were processed using EEGLAB (Swartz Center for Computational Neurosciences, La Jolla, CA: <http://www.sccn.ucsd.edu/eeglab>) running under MATLAB R2012b (Mathworks, Navick, MA).

Pre-processing was performed as follows. EEG data were re-referenced offline to linked mastoids. Bad channels were then identified by visual inspection and excluded from processing. Epochs for each stimulus type were extracted from –2000 to +2000 ms with respect to the target stimulus in each trial, and were inspected for non-stereotyped artefacts and removed if present (4.87 % ± 4.22 of trials removed). Stereotyped artefacts, including blinks, eye movements and muscle artefacts were deleted via independent component analysis (ICA) using the extended infomax algorithm (Bell & Sejnowski, 1995). The average number of independent components removed was 2.71 (±1.27 SD). The remaining components were then projected back into electrode

space. After ICA, channels that were deemed bad were reintroduced by interpolating data between neighbouring electrodes using spherical spline interpolation (Perrin et al., 1987).

ERP analyses were performed on ICA-corrected epochs time-locked to the onset of each target (−200 to +1000 ms). To minimize the influence of individual differences in topographies as well as the effects of performing multiple statistical comparisons, the analyses of the P3 was performed on a central cluster including C1, Cz, C2, CP1, CPz, and CP2 in a 20 ms window with regards to the most positive point in the latency range of 315–375 ms (356 ms in the Long- and Short-term learning phases and 340 ms in the Test phase), selected on the basis of both the grand average visual detection of the electrodes showing maximal P3 amplitude and the topographical distribution of the activity on the scalp (see figures 8 and 10). Baseline was designated from −200 to 0 ms relative to stimulus onset.

Statistical analyses

Results were analysed with a Bayesian linear mixed-model (LMM) analysis using the package *brms* (Bürkner, 2017), a high-level interface on Stan (Carpenter et al., 2017) in R (RCore, 2016). Plots were made using *brms* and *ggplot2* (Wickham, 2016). An advantage of LMMs over traditional approaches such as repeated measures ANOVA and paired t-tests is that a single model can take all sources of variance into account simultaneously, and comparisons between conditions can be implemented in a single model. LMMs (of which t-tests and ANOVA are specific examples) allow for modelling complex data structures, interacting continuous and categorical variables, and taking inherent correlations in data structures into account. Bayesian LMMs give insight in the range of possible effect sizes, and as such allow for direct comparison of parameter estimates. Bayesian statistics is theoretically distinct from frequentist statistics in its inferences. The coefficient estimates are expressed in credible intervals, which reflect the intuitive notion of the value of a parameter falling within that interval with a given probability, 95% in this case.

Main analyses focused on comparing the two experiments, and therefore only behavioural and EEG data obtained in the Short-term learning and Test phases were included. Behavioural and EEG data from the Long-Term learning phase in Experiment 2 were analysed independently (see the end of this section), both as data sanity check and for descriptive purposes. For the main analyses, we used a predefined model reflecting our experimental design (Barr et al., 2013), and we kept this model structure the same for the behaviour and

amplitude models. In the behavioural model, participant's RTs in milliseconds were transformed to their logarithmic function in order to normalize the residuals (R. Harald Baayen & Milin, 2010). In the P3 amplitude model, participant's P3 amplitudes were normally distributed and did not need transformation. Both dependent variables were scaled to z-scores (making parameter estimates interpretable as effect sizes) for ease of interpretation and comparison. In the model (Figures 3 and 6), we implemented our hypothesis with a full interaction of Experiment (1, 2), Rule (A, B) and Trial order (1:599, scaled) in each Phase (short-term learning, test). The model additionally included individual participant intercepts and slopes of Rule (A, B) in order to account for individual variation. Contrasts of all categorical factors were centred (R. H. Baayen, 2008), so the intercept of the model represents the grand mean. We tested specific hypotheses with the 'hypothesis' function in brms. We used a generic weakly informative normally distributed prior with mean 0 and 1SD for each fixed parameter (Lemoine, 2019) and kept all other priors at default (see Appendix I for the full prior specification of all models). This way, our models are explicit about expected small effect sizes, conservative and robust to unrealistically large effects caused by noise, i.e., Type I errors. We furthermore used 4 chains of 3000 iterations each per model, of which 1000 per chain were used for warm-up only, a maximum tree depth of 15 and a target acceptance rate (adapt delta) of .95. Convergence was verified through visual inspection of trace plots, and the Rhat of 1.00 for each parameter.

The model was specified as follows,

*brms formula: [log reaction time, amplitude] ~ phase / experiment * rule * trial order + (1 + rule | participant)*

Our main hypothesis was that the associations acquired during the Long-term learning phase are preserved during the transient Short-term learning phase, in which they are replaced by newly learned associations, and influence predictions (measured as RTs and P3 amplitude) in the subsequent neutral Test phase. This hypothesis was tested with the three-way interaction of Experiment, Rule and Trial order in the Test phase. More specifically, we expected behavioural and EEG data at the beginning and the end of the Test phases to show different patterns in Experiments 1 and 2, due to the influence of the Long-term learning phase in Experiment 2. In order to further determine how these patterns differed, we subsequently performed simple-slope comparisons with 'Trial' to investigate the relevance of the effect of 'Rule' at the beginning and end of each Phase in each experiment (Figures 4 and 7).

Participants were excluded if they had less than 50% accuracy in the test phase. Two participants were additionally excluded for presenting periods of no-response. Our final sample consisted of 18 subjects in Experiment 1 and 17 subjects in Experiment 2.

For the analysis of data from the Long-term learning phase of Experiment 2 we used the same model as specified above, with the following differences. There was no variable 'Experiment', since the data were only collected for Experiment 2, and there was no variable 'Phase' in the model of the EEG data, since EEG data was only collected in Long-term Learning phase Session II. The variable Session was included in the model of the behavioural data, where it had two levels: Long-term Learning Session I, and Long-term Learning Session II.

Results

In order to show the differences observed between Experiments 1 and 2, this section focuses on describing behavioural and ERP results obtained in the Short-term learning and Test phases, present in both experiments. Behavioural and ERP results corresponding to the Long-term learning phase, present only in Experiment 1, are described together at the end of the Results section.

Short-term learning and Test phases: behavioural results

Results obtained in the Short-term Learning and Test phases are illustrated in Figure 2 (left and right panel, respectively). Medians and credible intervals of values of each parameter can be found in Figures 3 and 4¹. The model showed that RTs were shorter (Figure 3) in the Test (Mean: 504 ms) than in the Short-term Learning phase (Mean: 510 ms), presumably reflecting a general effect of practice. Below, we describe the different phases separately.

Short-term learning phase

Results obtained in the Short-term Learning phase are illustrated in Figure 2 (left panel) and described in Figures 3 (general model) and 4 (grouped by Experiment). A main effect of Experiment revealed slower RTs in Experiment 1 (Mean RT: 534 ms) than in Experiment 2 (Mean RT: 480 ms). This probably reflects the effect of

¹ Credible intervals depict the probability that a population effects would fall somewhere on that interval. As such, when an interval does not cross the 0 line, the probability is high that the effect found with the sample is a real effect.

practice, since participants in Experiment 2 performed a similar task in the Long-term learning phase prior to the Short-term learning and Test phases. A main effect of Trial showed that RTs progressively decreased over this phase. However, the significant interaction of Experiment and Trial revealed that RTs indeed decreased over trial in Experiment 1, but actually increased, albeit slightly, in Experiment 2 (Figure 4, left panel). Taken together, these results suggest that, while there was apparently a general effect of practice in Experiment 1, with RTs getting faster over the course of the phase, such an effect was not present in Experiment 2, presumably due to some interference caused by the Long-term Learning phase that participants in this experiment, but not those in Experiment 1, performed prior to the Short-term Learning phase. This interference could partly consist of a ceiling effect resulting from previously performing the extensive Long-term learning phase in Experiment 2. As suggested in the Experiment 2 Long-term learning phase result description (see the end of the Results section and Figure 10), RT improvement over the Long-term learning phase seems to reach a ceiling throughout the second part of the Long-term learning phase. This ceiling effect seems to have consolidated and can also be seen in Figure 2, where it is evident that RTs in the Short-term learning phase are from the very beginning much faster in Experiment 2 than in Experiment 1.

Finally, a main effect of Rule revealed slower RTs to Rule A (Mean RT: 525 ms) than B (Mean RT: 498 ms). This effect was expected, since it indicates that participants learned about the relative probabilities of the two rules in both experiments and consequently responded faster to targets presented according to the most predictive one (Figure 3). Interestingly, an interaction of Rule and Trial further characterized this effect by showing that RTs decreased over this phase for targets presented according to Rule B while RTs to targets presented according to Rule A did not change. This would indicate that targets presented according to the less probable rule remained relatively surprising all through this phase, while participants progressively learned to anticipate targets following the most probable rule.

Test phase

Results obtained in the Test phase are illustrated in Figure 2 (right panel) and described in Figures 3 and 4. No main effects were found in the Test phase. There was however an interaction of Experiment and Trial showing decreasing RTs over the course of this phase in Experiment 1 but not in Experiment 2 (Figure 4, left panel). Furthermore, an interaction of Rule and Trial was also observed, with decreasing RTs over trial for Rule A, and increasing RTs over trial for Rule B, as shown in Figure 2. Lastly and crucially, results showed a three-way

interaction of Experiment, Rule and Trial (Figure 3). This difference between experiments consisted of a stronger interaction of Rule and Trial in Experiment 2 than Experiment 1, as well as a main effect of Trial indicating that RTs decreased over trial in Experiment 1 but not in Experiment 2 (Figure 4, left panel). Visual inspection of the RT plots (see Figure 2) suggests that this difference lies in RTs to Rules A and B converging over trial in Experiment 1, but showing a X-shaped development in Experiment 2. To investigate the significance of this difference we probed the effect of ‘Rule’ at different stages of each experiment. To this end, we made simple-slope comparisons referencing the trial-predictor to the start and end (Figure 4, right panel) of each phase. Results revealed that at the start of the Test phase there was a main effect of Rule in both Experiment 1 and 2, with longer RTs to Rule A than B, while at the end of this phase an effect of Rule was observed in Experiment 2, but not in Experiment 1. This effect showed an opposite pattern to that observed at the beginning of the phase, with longer RTs to Rule B than Rule A. Altogether, these results suggest that in the neutral context of the Test phase participants started performing according to the relative probabilities they learned in the immediately preceding Short-term Learning phase. Over the course of the phase, however, RT differences between rules disappeared in Experiment 1, but reverted in Experiment 2 to show faster RTs to targets presented according to Rule A than to Rule B, a pattern similar to that observed in the Long-term Learning phase, as we describe at the end of the Results section.

Short-term learning and Test phases: ERP results

Short-term learning phase

Results obtained in the Short-term Learning phase are illustrated in Figure 5 (left panel) and medians and credible intervals of parameter values can be found in Figures 6 (general model) and 7 (grouped by Experiment). A main effect of Trial showed that the amplitude of P3 decreased over the course of the phase. However, an interaction of Experiment and Trial further clarified this effect by showing that amplitude decreased in Experiment 2, but not in Experiment 1 (see Figure 7, left panel), hence suggesting that participants in Experiment 2 were somehow influenced by their pre-exposure to the Long-term Learning phase. This effect seems to be essentially driven by the pronounced reduction of P3 amplitude to targets presented according to the most predictive rule (B) in Experiment 2, which presumably reflects a practice effect produced by participants’ previous exposure to the extensive Long-term Learning phase. That previous experience may have allowed them to learn about the general structure of the Experiment (one rule is much more predictive than the other). As a consequence, the predictability of targets presented according to the most predictive rule would increase at a

higher rate over the phase for participants pre-exposed to the Long-term learning phase than for participants not pre-exposed. This difference would be reflected in P3 amplitude reduction.

Results also revealed a main effect of Rule (Figure 6), with higher amplitudes for Rule A (Mean amplitude: 3.75 μV) than B (Mean amplitude: 1.14 μV), that is, higher for targets presented according to the less predictive rule and therefore more surprising or unexpected. This anticipated effect was further explained by an interaction of Rule and Trial, with decreasing amplitudes for Rule B but not A suggesting that participants became progressively better at anticipating targets presented according to the most predictive rule, while targets presented according to the alternative rule remained surprising throughout the phase. These results parallel those observed in the analysis of RTs. The interaction of Experiment and Rule, showing a bigger difference in amplitudes between rules in Experiment 2 (Mean amplitude Rule A: 4.69 μV ; Mean amplitude Rule B: 0.98 μV) than in Experiment 1 (Mean amplitude Rule A: 2.86 μV ; Mean amplitude Rule B: 1.29 μV), suggests again the influence of the Long-term Learning phase on participants' performance in the Short-term Learning phase in Experiment 2. ERP waveforms obtained in response to Rules A and B in Experiments 1 and 2 in the Short-term learning phase are shown for illustration purposes in Figure 8 (A).

Test phase

Results obtained in the Test phase are illustrated in in Figures 5 (right panel) and 8, means and credible intervals of parameter values can be found in Figure 6 and 7. A main effect of Trial was observed, showing that amplitudes decreased over the course of the phase. Crucially, results revealed a three-way interaction of Experiment, Rule and Trial (Figure 6). The difference between experiments (see Figure 7, left panel) consisted of an interaction of Rule and Trial in Experiment 2 and no interaction in Experiment 1. To investigate the significance of this difference we proceeded the same way as we did for the RTs, with simple-slope comparisons. No effect of Rule was observed in Experiment 1 nor 2 at the beginning of this phase. Over the course of the phase, however, differences in P3 amplitude between rules emerged in Experiment 2 but not in Experiment 1. Interestingly, as observed in RT analyses these differences showed a pattern similar to that observed in the Long-term Learning phase, as we describe in the following section. No such effect was observed in Experiment 1. ERP waveforms obtained in response to Rules A and B in Experiments 1 and 2 in the Test phase are shown in Figure 8 (B).

Long-term learning phase

Behavioural results

Behavioural results obtained in the Long-term Learning sessions of Experiment 2 are described in Figure 9 (left panel) and illustrated in Figure 10 (A). There was a main effect of Session, with longer RT in Session I (Mean RT: 466 ms) than II (Mean RT 427 ms), indicating a general effect of practice. There was also an effect of Rule in both Session I and II, with longer RTs to Rule B than A, showing that participants learned about the relative predictive value of the two rules and anticipated the most probable targets. Additionally, in Session I there was an effect of Trial, with decreasing RTs over the course of the session. No such effect was observed in Session II. This would be indicative of a ceiling effect on learning. Lastly, there was an interaction of Rule and Trial in Session I, with a sharper decrease of RTs in response to Rule A than to Rule B. Again, there was no such interaction in Session II. On the one hand, this interaction would also be indicative of the ceiling effect suggested above. On the other hand, it would show that participants became better over the course of the session at anticipating targets presented according to the most predictive rule, while targets presented according to the alternative rule remained relatively surprising.

ERP results

ERP results obtained in the Long-term Learning Session II of Experiment 2 are described in Figure 9 (right panel) and illustrated in Figure 10 (B). There was a main effect of Rule, showing higher amplitudes in response to the least predictive Rule B (Mean amplitude: 6.76 μV) than to Rule A (Mean amplitude: 0.75 μV). There was also a main effect of Trial, revealing that amplitudes decreased over the course of the phase. This effect was further clarified by an interaction of Rule and Trial, showing decreasing amplitudes over the course of the phase in response to targets presented according to Rule B, with no amplitude modulations for targets presented according to Rule A.

Discussion

Here we investigated the adaptability of predictions to different timescales. We measured RT and P3 amplitude as behavioural and neural correlates of PE. In the two experiments reported above we found effects on both measures, showing probability-related differences between targets following the most and the least predictive rules in every learning phase of both experiments. Over and above these effects, we found that for participants who performed the long-term learning phase prior to the short-term learning and test phases, both RTs and P3

amplitude showed a spontaneous transition from the short-term to the long-term learned associations in the test phase: differences observed in the short-term learning phase faded throughout the test phase and were replaced by those learned in the long-term learning phase. No such transition was observed in participants that were not pre-exposed to the long-term learning phase. This pattern of results indicates that long-term associations remained intact while short-term associations were learned and transiently used while valid, and then no longer used when the context changed. The spontaneous recovery suggests independent storage and control of long-term and short-term associations. Before we further consider the implications of our main finding, we will briefly discuss the pattern observed in the long- and short-term learning phases of the two experiments.

P3 amplitude was larger, and RTs slower, in response to targets presented according to the less probable rule in every learning phase of both experiments (short-term in both experiments, long-term in Experiment 2). Furthermore, these differences increased throughout every learning phase, indicating that participants learned about the probabilities and consequently predicted the most probable outcome, while targets presented according to the alternative rule remained surprising throughout each of those phases. At a behavioural level, this interpretation is supported by numerous studies showing that RTs are faster for targets that can be anticipated and slower for unpredicted, surprising stimuli (Huettel et al., 2002; Hyman, 1953; Meyniel et al., 2016). Regarding ERP modulations, previous studies on P3 have shown that its amplitude is sensitive to subjective probability, i.e., to estimation of the environment on the basis of previous observations and learning (Donchin & Coles, 1988; Mars et al., 2008). In other words, P3 is sensitive to the discrepancy between the predicted state of the environment and its actual state, that is, to PE. However, in contrast to other error-related components, such as N1 or the mismatch negativity (Stefanics & Czigler, 2012; Todorovic et al., 2011), which are related to sensory processing, P3 is considered a high-level error correlate, associated not to specific sensory features but to the processing of global contextual deviations (Bekinschtein et al., 2009) and to the update of perceptual evidence (Ehinger et al., 2015; Kelly & O'Connell, 2013; O'Connell et al., 2012; Sutton et al., 1965). The modulations of RTs and P3 amplitude observed in the long-term and short-term learning phases of the present experiments are therefore coherent with these views. Such modulations would reflect the violation of an abstract rule (same or opposite direction), rather than a mismatch between the physical attributes of the predicted and actual targets, which are different on a trial basis and thus would not create a strong sensory model on the basis of which physical characteristics of the target could be predicted. This differs from most other studies, in which sensory predictions are based on strong contextual regularities that favour the creation of

such a model (Bekinschtein et al., 2009; Garrido et al., 2009; Hsu et al., 2015; Schadow et al., 2009; Todorovic & Lange, 2012). This explains why no modulations in components related to sensory prediction, such as N1, were found.

Our results are also congruent with previous works that have studied the relationship between P3 amplitude and stimulus improbability expressed in terms of statistical surprise (Kolossa et al., 2015; Mars et al., 2008; Strange et al., 2005). In this regard, Mars and colleagues (Mars et al., 2008) demonstrated that trial-by-trial fluctuations in P3 amplitudes could be explained in terms of participants keeping track of the global probabilities of visual events, so that its amplitude was reduced or enhanced as a function of surprise. Although the design used in the present study is much simpler in terms of probabilities, our results fit this interpretation, with P3 amplitudes in response to targets presented according to the infrequent rule being systematically larger on a trial basis and P3 amplitudes in response to targets presented according to the dominant rule reducing throughout every single learning phase as participants learned about their high probability. Results observed in the test phase of Experiment 1, where differences in P3 amplitude rapidly disappeared as participants learned that both rules were equally predictive, are also coherent with this interpretation.

Despite the similarities, there were clear differences between experiments when comparing the short-term learning phases. While in Experiment 1 RTs became faster throughout the phase in response to targets presented according to any rule, the difference increase being due to RT to targets that followed the dominant rule getting faster at a higher rate, in Experiment 2 RT difference between rules had a different origin. Indeed, rather than speeding up, RTs overall slowed down, and the difference increase between rules relied on RTs to surprising targets becoming progressively slower while RTs to predictable targets became slightly but significantly faster. This pattern suggests that in Experiment 2 the previous exposition to the Long-term learning phase influenced performance in the new context participants encountered in the Short-term learning phase. This influence may be twofold. First, it may reflect a general effect of practice, so that participants reached a performance level close to ceiling throughout the Long-term learning phase and consequently there was very little room for further RT improvement in the Short-term learning phase. This would in part explain the fact that RT do not further decrease during this phase. Second, and most importantly, the RTs increase in response to targets presented according to the infrequent rule in Experiment 2 suggests a more direct interference of the associations learned in the Long-term learning phase on performance in the Short-term learning phase. Specifically, the rule that was

the most predictive in the Long-term learning phase is in the Short-term learning phase the most unlikely to predict the target, and participants need to learn about that. This learning process relies on practice and on accumulating enough contextual data (Bogacz, 2007; Gold & Shadlen, 2007), which takes time, and participants could have relied on their recent past experience in the Long-term learning phase to make their predictions while there is not enough information available yet, needing to make an effort to inhibit their responses to the formerly most, now least predictive rule, which would be reflected in the RT increase.

The pattern observed in the long- and short-term learning phases allows us to interpret our main findings, namely the differences between experiments in the test phase and, specifically, the spontaneous recovery of long-term associations in Experiment 2. As described in the Methods section, the test phase was identical in the two experiments, where targets could be predicted according to any of the two possible rules with equal probability. Results showed that participants rapidly learned about these contextual changes, as indicated by the immediate cancellation of the differences between rules in P3 amplitude and by the progressive fading and eventual suppression of RT differences. The existence of differences in RT at the beginning of this phase in the two experiments would indicate that the associations learned in the immediately preceding short-term learning phase were still being used by participants to anticipate the targets, until enough information about the new rules' relative probabilities was collected. Importantly, while in Experiment 1 differences in both variables disappeared after a number of trials, indicating that the new relative probabilities were learned, in Experiment 2 results reverted to a pattern similar to that observed in the long-term learning phase, i.e., faster RTs and reduced P3 amplitude in response to targets presented according the rule that most likely predicted them in that phase. As in the Short-term learning phase, this result can only be explained by the previous exposure to the long-term learning phase of participants in Experiment 2. More specifically, it would reflect the system's attempt to interpret contextual information and predict upcoming events according to the rule that dominated in the Long-term learning phase once the alternative rule, valid in the Short-term learning phase, could not effectively predict targets anymore. As said above, gathering evidence to uncover the relative probabilities of events and to adjust predictions accordingly takes time, and while this was achieved the system might have tried to optimize this process by exploiting all sources of information available. This would include not only the sensory information extracted from the environment while performing the test phase, but also the knowledge built upon similar past experiences (Domenech & Dreher, 2010), i.e., during the Long-term learning phase. This look back into previous knowledge would have been additionally impelled by the significantly larger number of trials

included in the Long-term learning phase: from a participant's perspective the associations learned during the Short-term learning phase might have been perceived as transiently valid only, leading participants to rely on those learned in the Long-term learning phase until enough new information was available. Consequently, if the test phase extended in time, we would expect participants to eventually learn about the new relative probabilities and differences between the two possible targets to disappear, showing a similar pattern to that observed in Experiment 1.

The spontaneous recovery of the associations learned in the Long-term learning phase indicates that these were not permanently replaced by short-term associations, but remained accessible to be used when the transiently useful short-term associations were no longer valid. This suggests independent control of short- and long-term associations. The question at this point would be whether this independent control requires or not the operation of independent mechanisms. There is increasing evidence that learning invokes multiple processes working at different timescales. In Pavlovian learning literature it has been shown that long-term associations between stimulus and behaviour spontaneously recover after being temporarily replaced by alternative stimulus-behaviour associations conditioned on a shorter time-scale (see for instance Bouton, 1993, for a review), which would indicate the participation of at least two mechanisms, one implementing the short-term association for transient adaptation and the other preserving the long-term association for being reused when conditions revert to the most common state. Working on perceptual learning, Bao and colleagues (Bao & Engel, 2012a, 2012b) ran a series of experiments in which participants were adapted to high contrast for a relatively long time to be later deadapted to a lower contrast for a shorter period. They found that although the short phase initially produced deadaptation, adaptation effects spontaneously recovered when participants were tested in a neutral environment. They concluded that a single controlling mechanism could not account for the recovery effects, and that perceptual learning was possibly controlled by a continuum of mechanisms acting over a large range of timescales (Bao & Engel, 2012a). Along the same lines, Vul and co-workers (Vul et al., 2008) had previously demonstrated that the McCollough effect, an orientation contingent colour aftereffect, is a product of two distinct and separable timescales of learning in early visual cortex. Similar examples can also be found in motor learning research, where Smith and colleagues showed that in motor learning of reach, adaptation depends on two distinct processes, with different characteristics, that, importantly, operate at different learning rates, favouring motor adaptation to disturbances in the environment or within the motor system (Smith et al., 2006).

Taken together, those works suggest that the participation of multiple operators subserving learning at distinct timescales might be a widely implemented resource, providing humans with the necessary flexibility to successfully adapt to ever-changing contexts. However, while hypotheses proposing different independent operators control learning and adaptation at different timescales fit well in simple scenarios like those described above, they are problematic when explaining more complex situations. More specifically, in our study such an explanation would imply the contribution of as many operators as sets of associations learned, i.e., the associations participants would eventually learn in the test phase of Experiment 2 would call for a third operator, a hypothetical additional phase with different associations presented after the test phase would require the operation of a fourth mechanism, and so on. Alternatively, we think that the pattern of results found in our study needs for the associations learned at different timescales being stored and controlled separately but not, strictly speaking, for the operation of separate mechanisms. In this regard, although the present study does not allow us to draw conclusions about the neural mechanisms underlying the operators supporting predictions at different time scales, we believe that our data fit quite well into previous findings about the biological bases of implicit statistical learning (Christiansen, 2019). There is strong evidence that implicit statistical learning activates the declarative and nondeclarative memory systems of the brain, essentially dependent on the medial temporal lobe (MTL) and the striatum, respectively (Batterink et al., 2019; Reber, 2008, 2013; Squire, 1992). Specifically, neuroimaging results suggest a competitive interaction of the two systems in implicit statistical learning paradigms. The MTL system, supporting more flexible and abstract learning that can be generalized to new retrieval contexts, would play a role in learning during early stages of training, supporting rapid initial acquisition of higher order associations in complex sequences and making predictions about possible outcomes (Bornstein & Daw, 2012). As learning progresses and the statistical structure of the context is unravelled, the associations between stimuli would become more predictable, which would lead to the progressive disengagement of the MTL system and to the striatal system, which supports more specific and context-dependent learning, subsequently taking control (Poldrack et al., 2001).

The interaction between the MTL and the striatum systems could hypothetically support the operation of predictions at different timescales observed in our study. Learning in each phase would be initially supported by the MTL system extracting the rules from the predictive sequences. Once the rules are learned, the activity in the MTL system would decrease and the striate system would take control of the predictive activity.

Importantly, the information learned in every phase would be stored by the MTL system. When there are no

previously learned associations available, statistical learning would simply proceed as explained above. In the test phase of both experiments, however, the initial involvement of the MTL system would not only support the extraction of statistical regularities, but also activate previously learned sets of associations, preferentially the most recently valid for target prediction (i.e., those learned in the short-term phase), as indicated by the RT and P3 amplitude pattern of results. As these associations fail to accurately predict targets, in Experiment 1 participants would eventually learn that both rules predict targets with equal probability and consequently the RT and P3 amplitude differences between rules would disappear. In Experiment 2, however, the long-term set of associations would be retrieved. This retrieval might be responsible for the spontaneous recovery observed in RTs and P3 amplitudes. Since these rules do not efficiently predict targets either, we speculate that participants would finally generate a new set of rules based on what they learn about the environment, so that RT and amplitude differences between rules would eventually disappear, as observed in Experiment 1. Unfortunately, the limited length of the test phase does not allow us to check this hypothesis.

In sum, while our behavioural and EEG data present obvious limitations to gain insight into the neural mechanisms supporting predictions at different timescales, we consider that the pattern of results obtained in our study could be better explained on the basis of the competitive interaction between the MTL and striatal systems, rather than on the contribution of separate, independent mechanisms operating at different timescales. This interaction would provide the system with the necessary flexibility to allow multiple sets of rules to be latently available to respond to changes in the environment when needed. The long-term / short-term results presented here would thus represent just a particular sample of such dynamics.

Finally, as explained above, our results suggest that in the test phases of both experiments, instead of immediately learning about the regularities of the new context and applying that knowledge to anticipate the targets, at first participants used the associations that had been proven to correctly predict sensory inputs in the preceding experimental phases. This result is coherent with well-established findings on associative learning, particularly from interference paradigms (Spear, 1978, 1981; Bouton, 1991; Miller et al., 1986; Thomas, 1981) according to which different sets of associations can be learned and maintained through changing conditions, remaining more or less available for future retrieval (Bouton, 1993). However, the ensuing question would be what determines which one, among the different sets of associations stored, would be preferentially retrieved. Data obtained in previous experiments suggest that the relative availability of the different sets depends on

variables such as the degree of match between the context in which the sets were learned and the current context, the time that has passed since the associations have been learned, and the relative strength of those associations, which in turn relies on several factors: associative learning is incremental, so that with practice and consistency associations strengthen and performance becomes faster and more accurate (Schacter & Wagner, 2013). Hence, the pattern of results obtained in the test phase, and particularly the spontaneous recovery of long-term associations in Experiment 2, would be explained by participants initially applying the set of associations learned in the immediately preceding short-term learning phase, as those are the most recently acquired and, once these failed to consistently predict targets, either generating a new set of associations, as results from Experiment 1 suggest, or resorting before that to the more extensively practiced and consolidated long-term associations, as shown by results in Experiment 2. As we conjectured before, in Experiment 2 participants would also finally learn the probabilistic structure of the test phase, eventually causing RT and P3 amplitude differences between rules to disappear, as in Experiment 1, if the phase were prolonged long enough.

Conclusions

In order to successfully operate in an ever-changing environment, the brain needs to be able to flexibly adapt its predictions on the basis of the information it gathers over time. This is particularly important when a stable context is transiently modified: the brain needs to adapt to changes quickly, but without weakening the previous representations, as the new associations are only transient. Accordingly, our data show that the long-term associations were not simply overwritten by the short-term ones, since they spontaneously recovered in a subsequent neutral context different from that in which they were learned, as the pattern of RT and P3 amplitude modulations in the test phases suggests. The results of our experiments indicate that when the context changes and the current associations turn out not to be valid anymore, participants retrieve and use alternative associations they have previously learned in order to adapt their predictions, at least until enough contextual information is available to generate a new set of associations. The fact that the long-term associations recovered spontaneously, without needing to be relearned, indicates that they remained intact while transiently replaced by the short-term ones, and therefore supports the notion that they are stored and controlled independently from each other. Future research should attempt to further explore the generation, maintenance and possible interaction of predictions through different successive contexts by systematically manipulating their temporal and probabilistic relationships and, particularly, to determine the underlying neural bases supporting those processes.

Acknowledgements

This work was supported by the French Agence Nationale de la Recherche (ANR) within the *Programme franco-allemand en Sciences humaines et sociales (FRAL)* 2016 (PROJECT ID: ANR-16-FRAL-0008 to F.W.)

This research was partly supported by IdEx Université de Paris ANR-18-IDEX-0001

Appendix I: Priors used in each model

Parameter	Parameter code brms	Prior	Prior decision
Fixed effects	b	Normal (0,1)	Weakly informative prior of null-effect
Intercept	Intercept	student_t (3,0,2.5)	Default
Cholesky factors of correlations random effects	L	lkj_corr_cholesky (1)	Default
Random Intercept and Random slopes	sd, sigma	student_t (3,0,2.5)	Default

References

- Acharya, J. N., Hani, A., Cheek, J., Thirumala, P., & Tsuchida, T. N. (2016). American Clinical Neurophysiology Society Guideline 2: Guidelines for Standard Electrode Position Nomenclature. *Journal of Clinical Neurophysiology: Official Publication of the American Electroencephalographic Society*, 33(4), 308–311. <https://doi.org/10.1097/WNP.0000000000000316>
- Adams, R. A., Shipp, S., & Friston, K. J. (2013). Predictions not commands: Active inference in the motor system. *Brain Structure & Function*, 218(3), 611–643. <https://doi.org/10.1007/s00429-012-0475-5>
- Baayen, R. H. (2008). *Analyzing Linguistic Data: A Practical Introduction to Statistics using R*. Cambridge University Press. <https://doi.org/10.1017/CBO9780511801686>
- Baayen, R. Harald, & Milin, P. (2010). *Analyzing reaction times*. <https://doi.org/10.21500/20112084.807>
- Bao, M., & Engel, S. (2012a). Distinct mechanisms control contrast adaptation over different timescales. *Journal of Vision*, 12(9), 698–698. <https://doi.org/10.1167/12.9.698>
- Bao, M., & Engel, S. A. (2012b). Distinct mechanism for long-term contrast adaptation. *Proceedings of the National Academy of Sciences*, 109(15), 5898–5903. <https://doi.org/10.1073/pnas.1113503109>

- Barr, D. J., Levy, R., Scheepers, C., & Tily, H. J. (2013). Random effects structure for confirmatory hypothesis testing: Keep it maximal. *Journal of Memory and Language*, *68*(3), 255–278.
<https://doi.org/10.1016/j.jml.2012.11.001>
- Bastos, A. M., Usrey, W. M., Adams, R. A., Mangun, G. R., Fries, P., & Friston, K. J. (2012). Canonical microcircuits for predictive coding. *Neuron*, *76*(4), 695–711.
<https://doi.org/10.1016/j.neuron.2012.10.038>
- Batterink, L. J., Paller, K. A., & Reber, P. J. (2019). Understanding the Neural Bases of Implicit and Statistical Learning. *Topics in Cognitive Science*, *11*(3), 482–503. <https://doi.org/10.1111/tops.12420>
- Bekinschtein, T. A., Dehaene, S., Rohaut, B., Tadel, F., Cohen, L., & Naccache, L. (2009). Neural signature of the conscious processing of auditory regularities. *Proceedings of the National Academy of Sciences of the United States of America*, *106*(5), 1672–1677. <https://doi.org/10.1073/pnas.0809667106>
- Bell, A. J., & Sejnowski, T. J. (1995). An information-maximization approach to blind separation and blind deconvolution. *Neural Computation*, *7*(6), 1129–1159. <https://doi.org/10.1162/neco.1995.7.6.1129>
- Bidet-Caulet, A., Barbe, P.-G., Roux, S., Viswanath, H., Barthélémy, C., Bruneau, N., Knight, R. T., & Bonnet-Brilhault, F. (2012). Dynamics of anticipatory mechanisms during predictive context processing. *European Journal of Neuroscience*, *36*(7), 2996–3004. <https://doi.org/10.1111/j.1460-9568.2012.08223.x>
- Bornstein, A. M., & Daw, N. D. (2012). Dissociating hippocampal and striatal contributions to sequential prediction learning. *European Journal of Neuroscience*, *35*(7), 1011–1023.
<https://doi.org/10.1111/j.1460-9568.2011.07920.x>
- Bouton, M. E. (1993). Context, time, and memory retrieval in the interference paradigms of Pavlovian learning. *Psychological Bulletin*, *114*(1), 80–99. <https://doi.org/10.1037/0033-2909.114.1.80>
- Brainard, D. H. (1997). The Psychophysics Toolbox. *Spatial Vision*, *10*(4), 433–436.
<https://doi.org/10.1163/156856897X00357>
- Bürkner, P.-C. (2017). brms: An R Package for Bayesian Multilevel Models Using Stan. *Journal of Statistical Software*, *80*(1), 1–28. <https://doi.org/10.18637/jss.v080.i01>
- Carpenter, B., Gelman, A., Hoffman, M. D., Lee, D., Goodrich, B., Betancourt, M., Brubaker, M., Guo, J., Li, P., & Riddell, A. (2017). Stan: A Probabilistic Programming Language. *Journal of Statistical Software*, *76*(1), 1–32. <https://doi.org/10.18637/jss.v076.i01>

- Christiansen, M. H. (2019). Implicit Statistical Learning: A Tale of Two Literatures. *Topics in Cognitive Science*, 11(3), 468–481. <https://doi.org/10.1111/tops.12332>
- Dayan, P., Hinton, G. E., Neal, R. M., & Zemel, R. S. (1995). The Helmholtz Machine. *Neural Computation*, 7(5), 889–904. <https://doi.org/10.1162/neco.1995.7.5.889>
- Domenech, P., & Dreher, J.-C. (2010). Decision Threshold Modulation in the Human Brain. *Journal of Neuroscience*, 30(43), 14305–14317. <https://doi.org/10.1523/JNEUROSCI.2371-10.2010>
- Donchin, E., & Coles, M. G. H. (1988). Is the P300 component a manifestation of context updating? *Behavioral and Brain Sciences*, 11(3), 357–374. <https://doi.org/10.1017/S0140525X00058027>
- Ehinger, B. V., König, P., & Ossandón, J. P. (2015). Predictions of Visual Content across Eye Movements and Their Modulation by Inferred Information. *Journal of Neuroscience*, 35(19), 7403–7413. <https://doi.org/10.1523/JNEUROSCI.5114-14.2015>
- Feldman, H., & Friston, K. J. (2010). Attention, Uncertainty, and Free-Energy. *Frontiers in Human Neuroscience*, 4. <https://doi.org/10.3389/fnhum.2010.00215>
- Fischer, A. G., & Ullsperger, M. (2013). Real and fictive outcomes are processed differently but converge on a common adaptive mechanism. *Neuron*, 79(6), 1243–1255. <https://doi.org/10.1016/j.neuron.2013.07.006>
- Fiser, J., Berkes, P., Orbán, G., & Lengyel, M. (2010). Statistically optimal perception and learning: From behavior to neural representations. *Trends in Cognitive Sciences*, 14(3), 119–130. <https://doi.org/10.1016/j.tics.2010.01.003>
- Friston, K. (2005). A theory of cortical responses. *Philosophical Transactions of the Royal Society of London. Series B, Biological Sciences*, 360(1456), 815–836. <https://doi.org/10.1098/rstb.2005.1622>
- Garrido, M. I., Kilner, J. M., Kiebel, S. J., Stephan, K. E., Baldeweg, T., & Friston, K. J. (2009). Repetition suppression and plasticity in the human brain. *NeuroImage*, 48(1), 269–279. <https://doi.org/10.1016/j.neuroimage.2009.06.034>
- Helmholtz, H. von. (2013). *Treatise on Physiological Optics*. Courier Corporation.
- Hsu, Y.-F., Bars, S. L., Hämäläinen, J. A., & Waszak, F. (2015). Distinctive Representation of Mispredicted and Unpredicted Prediction Errors in Human Electroencephalography. *Journal of Neuroscience*, 35(43), 14653–14660. <https://doi.org/10.1523/JNEUROSCI.2204-15.2015>
- Huettel, S. A., Mack, P. B., & McCarthy, G. (2002). Perceiving patterns in random series: Dynamic processing of sequence in prefrontal cortex. *Nature Neuroscience*, 5(5), 485–490. <https://doi.org/10.1038/nn841>

- Hyman, R. (1953). Stimulus information as a determinant of reaction time. *Journal of Experimental Psychology*, 45(3), 188–196. <https://doi.org/10.1037/h0056940>
- Kelly, S. P., & O’Connell, R. G. (2013). Internal and External Influences on the Rate of Sensory Evidence Accumulation in the Human Brain. *Journal of Neuroscience*, 33(50), 19434–19441. <https://doi.org/10.1523/JNEUROSCI.3355-13.2013>
- Kleiner, M., Brainard, D., Pelli, D., Ingling, A., Murray, R., & Broussard, C. (2007). What’s new in psychtoolbox-3. *Perception*, 36(14), 1–16.
- Kolossa, A., Kopp, B., & Fingscheidt, T. (2015). A computational analysis of the neural bases of Bayesian inference. *NeuroImage*, 106, 222–237. <https://doi.org/10.1016/j.neuroimage.2014.11.007>
- Kording, K. P., Tenenbaum, J. B., & Shadmehr, R. (2007). The dynamics of memory as a consequence of optimal adaptation to a changing body. *Nature Neuroscience*, 10(6), 779–786. <https://doi.org/10.1038/nn1901>
- Lemoine, N. P. (2019). Moving beyond noninformative priors: Why and how to choose weakly informative priors in Bayesian analyses. *Oikos*, 128(7), 912–928. <https://doi.org/10.1111/oik.05985>
- Mars, R. B., Debener, S., Gladwin, T. E., Harrison, L. M., Haggard, P., Rothwell, J. C., & Bestmann, S. (2008). Trial-by-Trial Fluctuations in the Event-Related Electroencephalogram Reflect Dynamic Changes in the Degree of Surprise. *Journal of Neuroscience*, 28(47), 12539–12545. <https://doi.org/10.1523/JNEUROSCI.2925-08.2008>
- Meyniel, F., Maheu, M., & Dehaene, S. (2016). Human Inferences about Sequences: A Minimal Transition Probability Model. *PLOS Computational Biology*, 12(12), e1005260. <https://doi.org/10.1371/journal.pcbi.1005260>
- O’Connell, R. G., Dockree, P. M., & Kelly, S. P. (2012). A supramodal accumulation-to-bound signal that determines perceptual decisions in humans. *Nature Neuroscience*, 15(12), 1729–1735. <https://doi.org/10.1038/nn.3248>
- Perrin, F., Pernier, J., Bertrand, O., Giard, M. H., & Echallier, J. F. (1987). Mapping of scalp potentials by surface spline interpolation. *Electroencephalography and Clinical Neurophysiology*, 66(1), 75–81. [https://doi.org/10.1016/0013-4694\(87\)90141-6](https://doi.org/10.1016/0013-4694(87)90141-6)
- Poldrack, R. A., Clark, J., Paré-Blagoev, E. J., Shohamy, D., Creso Moyano, J., Myers, C., & Gluck, M. A. (2001). Interactive memory systems in the human brain. *Nature*, 414(6863), 546–550. <https://doi.org/10.1038/35107080>

- Polich, J. (2007). Updating P300: An Integrative Theory of P3a and P3b. *Clinical Neurophysiology: Official Journal of the International Federation of Clinical Neurophysiology*, 118(10), 2128–2148.
<https://doi.org/10.1016/j.clinph.2007.04.019>
- Rao, R. P. N., & Ballard, D. H. (1999). Predictive coding in the visual cortex: A functional interpretation of some extra-classical receptive-field effects. *Nature Neuroscience*, 2(1), 79–87.
<https://doi.org/10.1038/4580>
- RCore, T. (2016). *R: A language and environment for statistical computing*. R Foundation for Statistical Computing, Vienna, Austria.
- Reber, P. J. (2008). Cognitive Neuroscience of Declarative and Nondeclarative Memory. In A. S. Benjamin, J. S. De Belle, B. Etnyre, & T. A. Polk (Eds.), *Advances in Psychology* (Vol. 139, pp. 113–123). North-Holland. [https://doi.org/10.1016/S0166-4115\(08\)10010-3](https://doi.org/10.1016/S0166-4115(08)10010-3)
- Reber, P. J. (2013). The neural basis of implicit learning and memory: A review of neuropsychological and neuroimaging research. *Neuropsychologia*, 51(10), 2026–2042.
<https://doi.org/10.1016/j.neuropsychologia.2013.06.019>
- Sawaki, R., & Katayama, J. (2006). Stimulus context determines whether non-target stimuli are processed as task-relevant or distractor information. *Clinical Neurophysiology: Official Journal of the International Federation of Clinical Neurophysiology*, 117(11), 2532–2539.
<https://doi.org/10.1016/j.clinph.2006.06.755>
- Schacter, D.L & Wagner, A. D. (2013). *Learning and memory*. In: Kandel, E. R., Schwartz, J. R., Jessell, T. M. *Principles of neural science (5th Ed.)*. New York: McGraw-Hill.
- Schadow, J., Lenz, D., Dettler, N., Fründ, I., & Herrmann, C. S. (2009). Early gamma-band responses reflect anticipatory top-down modulation in the auditory cortex. *NeuroImage*, 47(2), 651–658.
<https://doi.org/10.1016/j.neuroimage.2009.04.074>
- Smith, M. A., Ghazizadeh, A., & Shadmehr, R. (2006). Interacting Adaptive Processes with Different Timescales Underlie Short-Term Motor Learning. *PLOS Biology*, 4(6), e179.
<https://doi.org/10.1371/journal.pbio.0040179>
- Sohoglu, E., & Davis, M. H. (2016). Perceptual learning of degraded speech by minimizing prediction error. *Proceedings of the National Academy of Sciences*, 113(12), E1747–E1756.
<https://doi.org/10.1073/pnas.1523266113>

- Squire, L. R. (1992). Declarative and Nondeclarative Memory: Multiple Brain Systems Supporting Learning and Memory. *Journal of Cognitive Neuroscience*, 4(3), 232–243.
<https://doi.org/10.1162/jocn.1992.4.3.232>
- Stefanics, G., & Czigler, I. (2012). Automatic prediction error responses to hands with unexpected laterality: An electrophysiological study. *NeuroImage*, 63(1), 253–261.
<https://doi.org/10.1016/j.neuroimage.2012.06.068>
- Strange, B. A., Duggins, A., Penny, W., Dolan, R. J., & Friston, K. J. (2005). Information theory, novelty and hippocampal responses: Unpredicted or unpredictable? *Neural Networks*, 18(3), 225–230.
<https://doi.org/10.1016/j.neunet.2004.12.004>
- Sutton, S., Braren, M., Zubin, J., & John, E. R. (1965). Evoked-Potential Correlates of Stimulus Uncertainty. *Science*, 150(3700), 1187–1188. <https://doi.org/10.1126/science.150.3700.1187>
- Swanson, L. R. (2016). The Predictive Processing Paradigm Has Roots in Kant. *Frontiers in Systems Neuroscience*, 10, 79. <https://doi.org/10.3389/fnsys.2016.00079>
- Thornton, C. (2017). Predictive processing simplified: The infotopic machine. *Brain and Cognition*, 112, 13–24. <https://doi.org/10.1016/j.bandc.2016.03.004>
- Todorovic, A., Ede, F. van, Maris, E., & Lange, F. P. de. (2011). Prior Expectation Mediates Neural Adaptation to Repeated Sounds in the Auditory Cortex: An MEG Study. *Journal of Neuroscience*, 31(25), 9118–9123. <https://doi.org/10.1523/JNEUROSCI.1425-11.2011>
- Todorovic, A., & Lange, F. P. de. (2012). Repetition Suppression and Expectation Suppression Are Dissociable in Time in Early Auditory Evoked Fields. *Journal of Neuroscience*, 32(39), 13389–13395.
<https://doi.org/10.1523/JNEUROSCI.2227-12.2012>
- Vul, E., Krizay, E., & MacLeod, D. I. A. (2008). The McCollough effect reflects permanent and transient adaptation in early visual cortex. *Journal of Vision*, 8(12), 4–4. <https://doi.org/10.1167/8.12.4>
- Wacongne, C., Labyt, E., Wassenhove, V. van, Bekinschtein, T., Naccache, L., & Dehaene, S. (2011). Evidence for a hierarchy of predictions and prediction errors in human cortex. *Proceedings of the National Academy of Sciences*, 108(51), 20754–20759. <https://doi.org/10.1073/pnas.1117807108>
- Waszak, F., & Herwig, A. (2007). Effect anticipation modulates deviance processing in the brain. *Brain Research*, 1183, 74–82. <https://doi.org/10.1016/j.brainres.2007.08.082>
- Wickham, H. (2016). *ggplot2: Elegant Graphics for Data Analysis*. Springer.

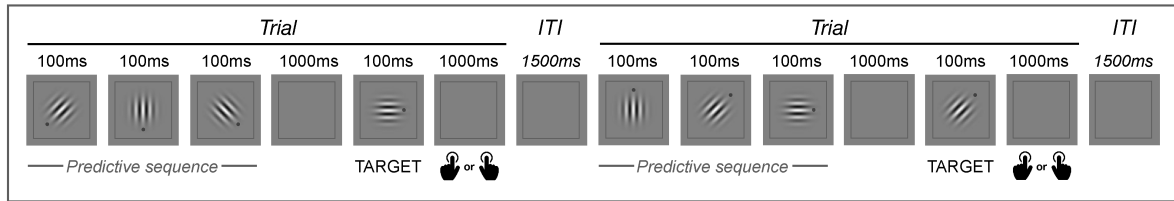


Figure 1. Schematic representation of events in two consecutive trials. All the phases in both Experiment 1 and Experiment 2 contained the same type of sequences. Phases only differed in the probability of targets following or not the rotation direction indicated by the first three stimuli in the sequence.

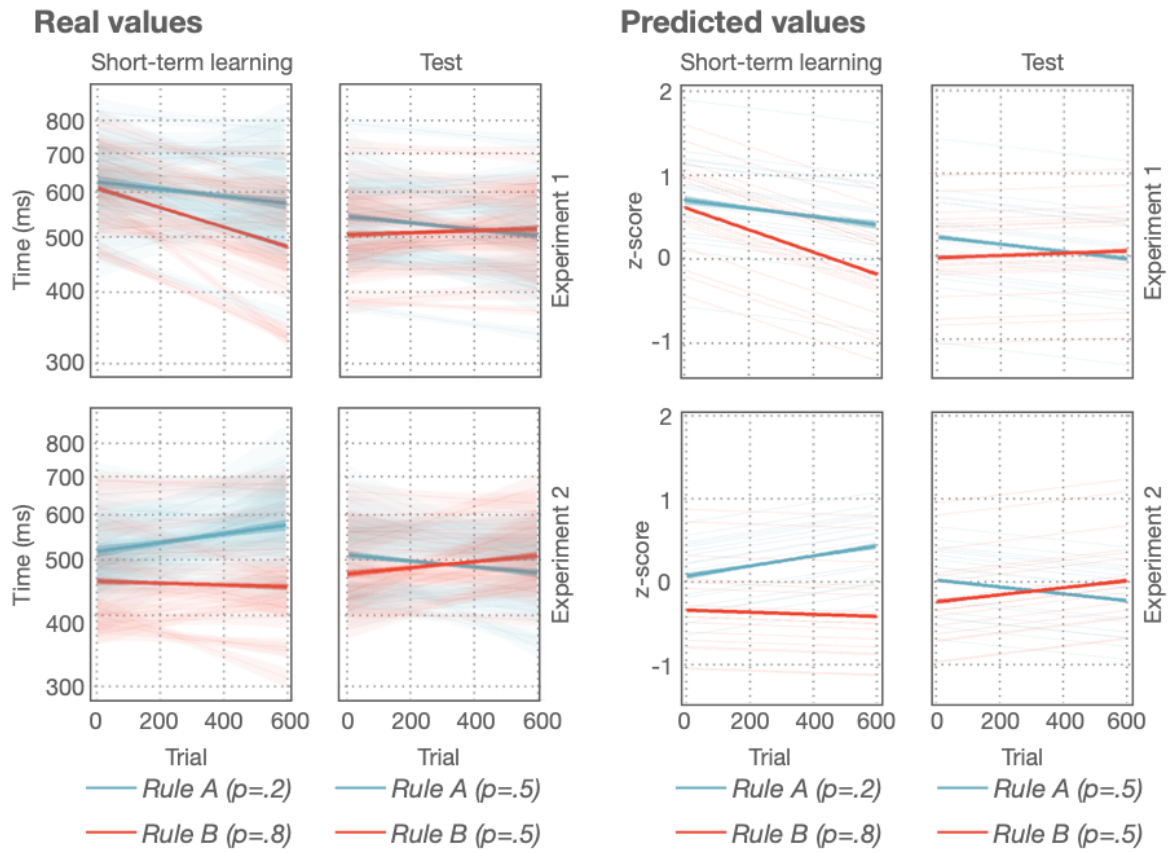


Figure 2. Behavioural results obtained in the Short-term Learning and Test phases in Experiments 1 (upper panels) and 2 (lower panels). Overall and per-participant results are given for the data (left panel, RTs on semi-log scale) and for the model predictions (right panel, z-scores).

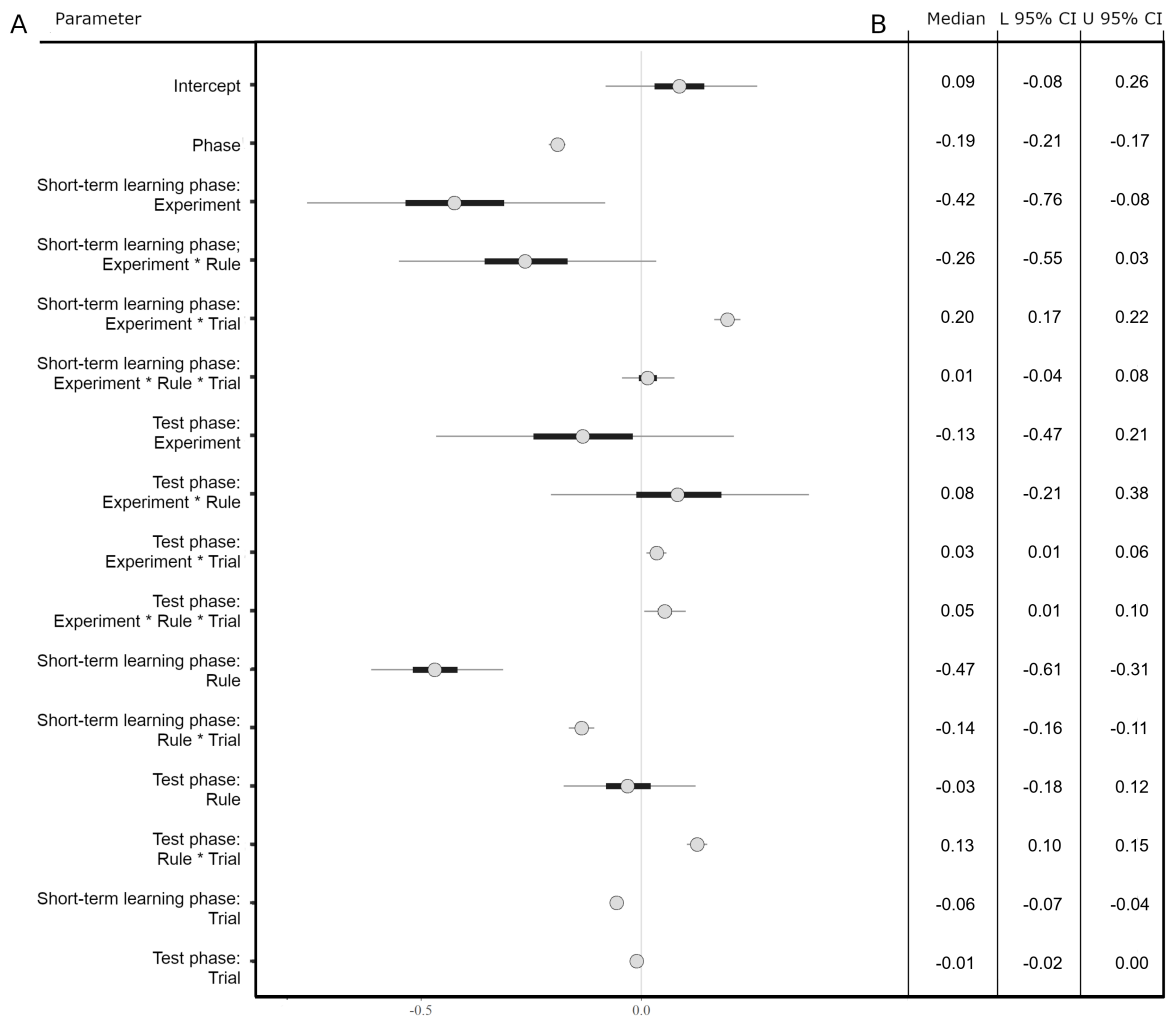


Figure 3. Behavioural data, short-term learning and test phases, general model. We grouped the interaction of Experiment, Rule and Trial under the short-term learning and test phase. (A) Medians and credible intervals of parameter values in the behavioural data (50% probability, thick line; 95% probability, thin line). Values on the x axis represent effect size. Fixed-effects parameters included in the model are shown on the y-axis. Credible intervals that do not cross the 0 line represent reliable effects at 95%. (B) values of the medians (left column) and lower and upper boundaries (central and right columns, respectively) of the 95% credible interval. The first level of each categorical variable is contrast-coded in negative, for example: RTs in the short-term learning phase are longer (-.5*-.19) than the RTs in the test phase (.5*-.19).

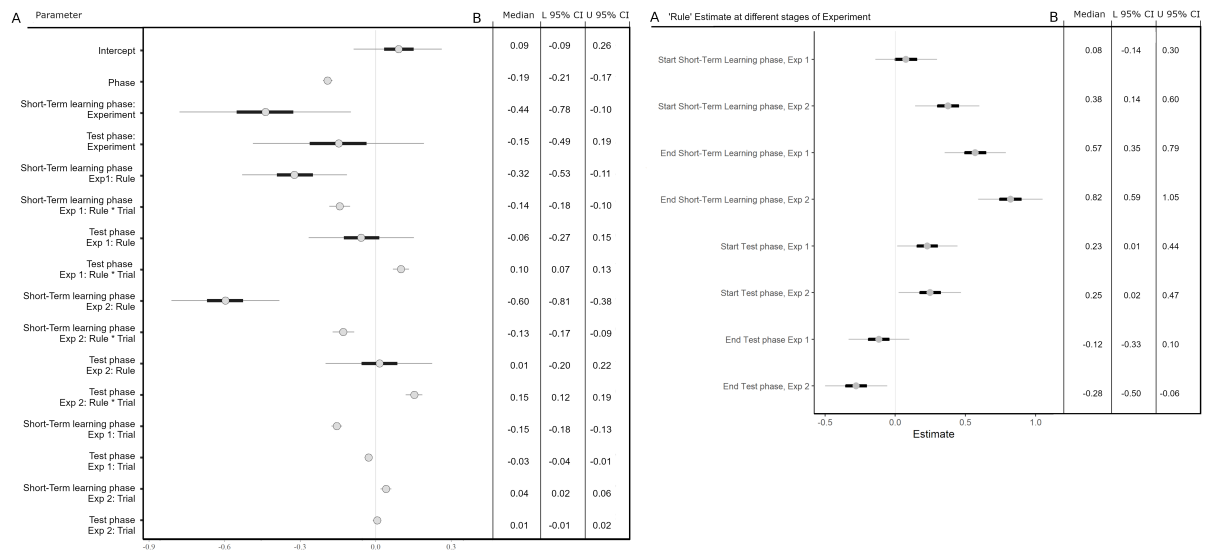


Figure 4. Behavioural data, short-term learning and test phases, model grouped by Experiment. We grouped the interaction of Rule and Trial by Short-Term Learning and Test phase in Experiment 1 and 2. (A) Medians and credible intervals of parameter values in the RT data (50% probability, thick line; 95% probability, thin line). Values on the x axis represent effect size. Credible intervals that do not cross the 0 line represent reliable effects at 95%. (B) Numerical values of the medians and lower and upper boundaries of the 95% credible interval. The first level of each categorical variable is contrast-coded in negative. Left panel: Fixed-effects parameters included in the model are shown on the y axis. Right panel: Credible intervals of the Rule parameter at different stages of each experiment, obtained with simple slope comparisons.

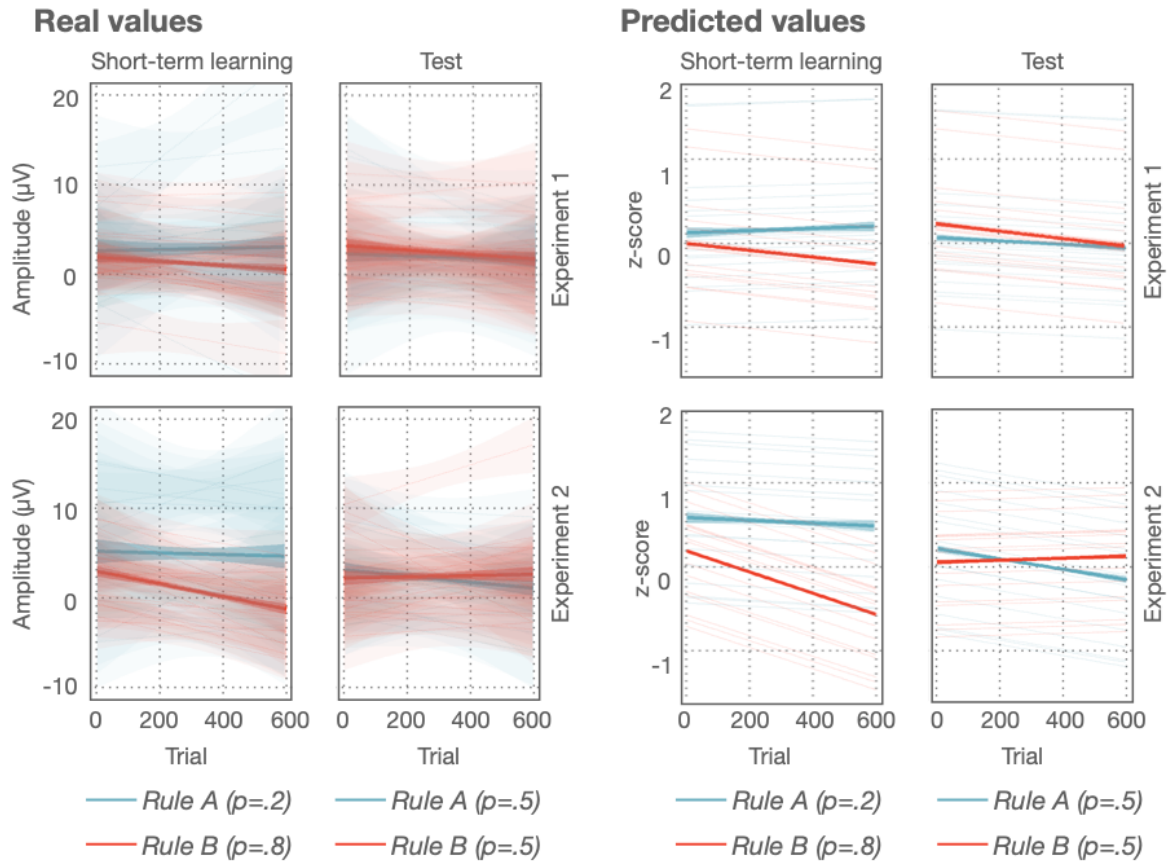


Figure 5. Variations of P3 amplitude throughout the Short-Term Learning and Test phases in Experiment 1 (upper panels) and 2 (lower panels). Overall and per-participant results are given for the data (left panel, μV) and for the model predictions (right panel, z-scores).

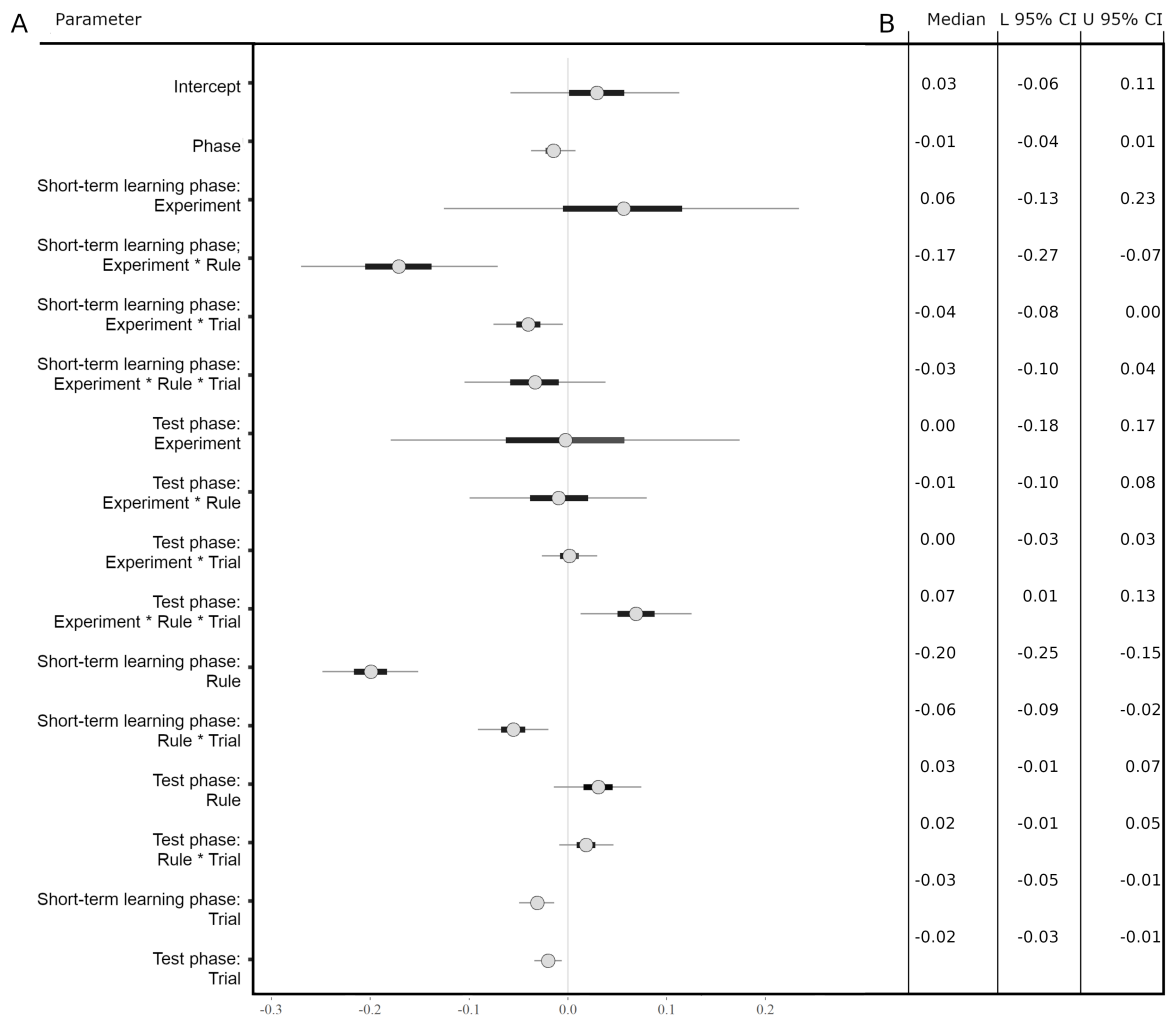


Figure 6. EEG data, short-term learning and test phases, general model. We grouped the interaction of Experiment, Rule and Trial under the short-term learning and test phase. (A) Medians and credible intervals of parameter values in the EEG data (50% probability, thick line; 95% probability, thin line). Values on the x axis represent effect size. Fixed-effects parameters included in the model are shown on the y-axis. Credible intervals that do not cross the 0 line represent reliable effects at 95%. (B) values of the medians (left column) and lower and upper boundaries (central and right columns, respectively) of the 95% credible interval. The first level of each categorical variable is contrast-coded in negative, for example: amplitudes in the short-term learning phase are longer (-.5*-.19) than the amplitudes in the test phase (.5*-.19).

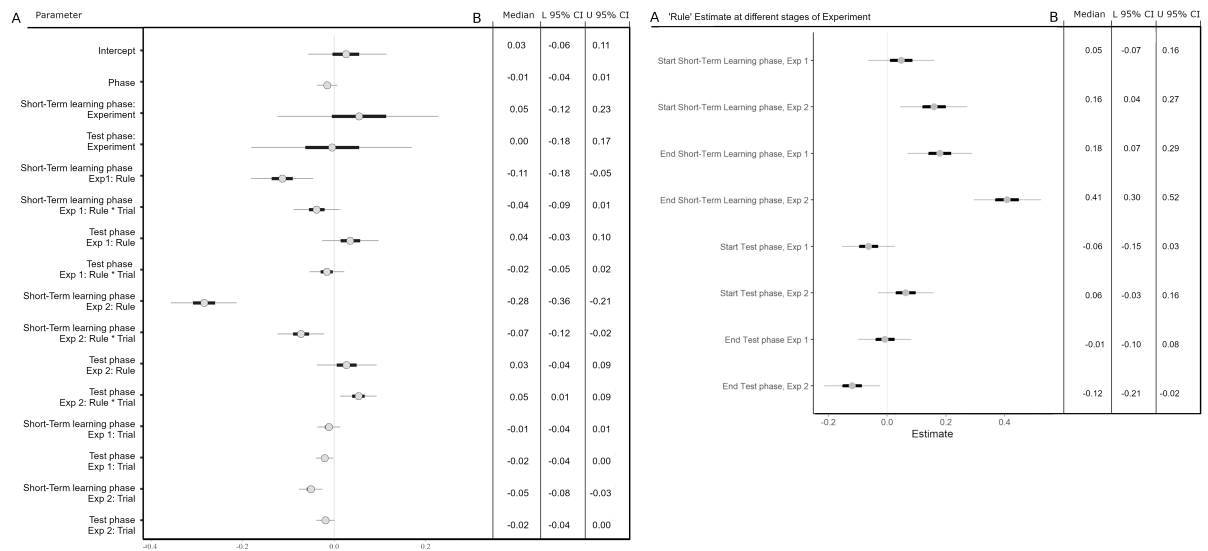


Figure 7. EEG data, short-term learning and test phases, model grouped by Experiment. We grouped the interaction of Rule and Trial by Short-Term Learning and Test phase in Experiment 1 and 2. (A) Medians and credible intervals of parameter values in the EEG data (50% probability, thick line; 95% probability, thin line). Values on the x axis represent effect size. Credible intervals that do not cross the 0 line represent reliable effects at 95%. (B) Numerical values of the medians and lower and upper boundaries of the 95% credible interval for trials referenced to the middle of the experiment (left panel), the start (central panel) and the end of the experiment (right panel). The first level of each categorical variable is contrast-coded in negative. Left panel: Fixed-effects parameters included in the model are shown on the y axis. Right panel: Credible intervals of the Rule parameter at different stages of each experiment, obtained with simple slope comparisons.

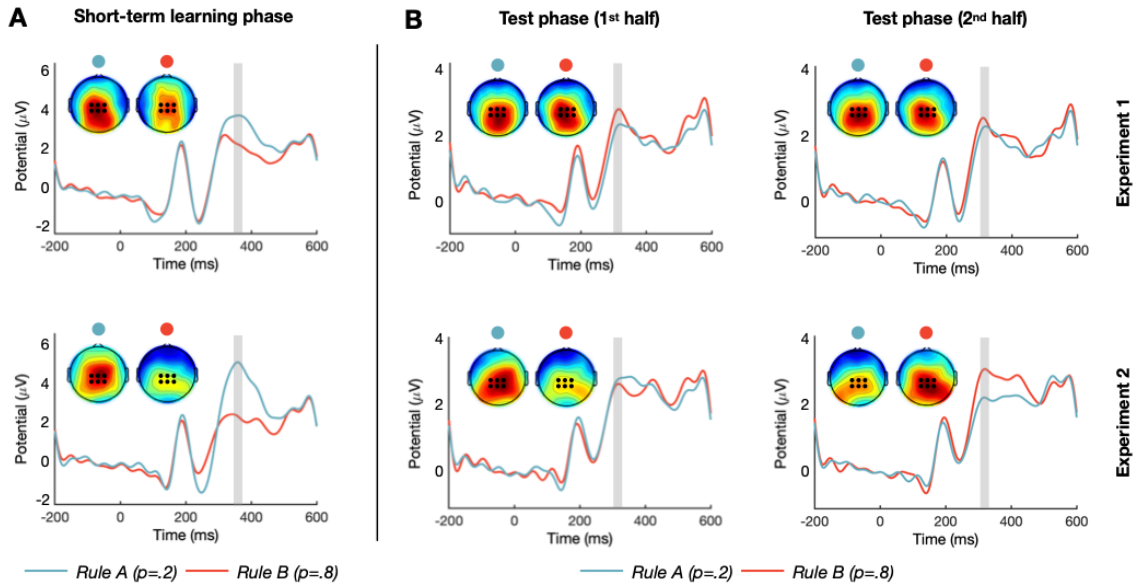


Figure 8. (A) ERP waveforms corresponding to the Short-term Learning phase for Experiments 1 (upper panel) and 2 (lower panel) at the central electrode cluster. Topographic maps show the distribution of voltage of P3 in response to targets presented according to rules A and B. (B) ERP waveforms corresponding to the Test phase for Experiments 1 (upper row) and 2 (lower row) at the central electrode cluster, together with voltage scalp distribution maps for n targets presented according to rules A and B. Average waveforms for the first and second halves of the Test phase are presented in the left and right columns, respectively.

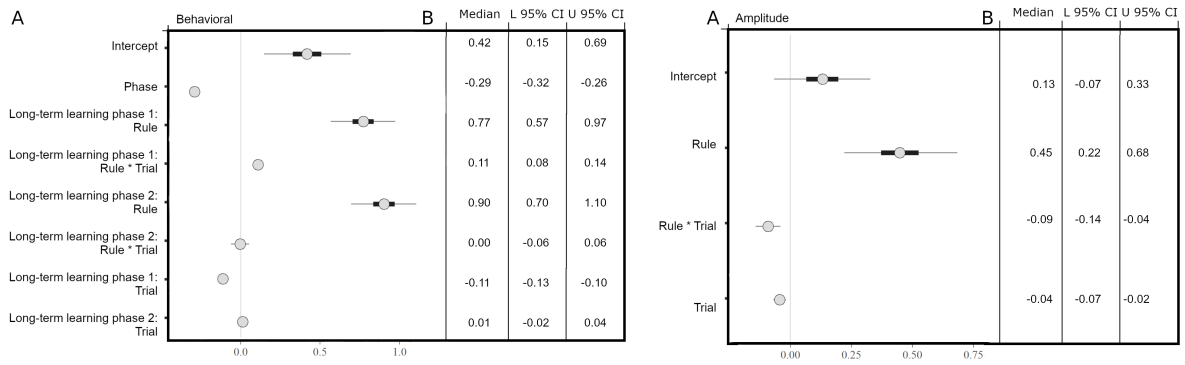


Figure 9. Model of Long-term Learning Phases. We modelled the interaction of rule and trial, grouped by Long-Term Learning phase I and II in the case of behavioural data. Medians and credible intervals of parameter values in the behavioural (left panel) and EEG (right panel) data (50% probability, thick line; 95% the model probability, thin line). Values on the x axis represent effect size. Fixed-effects parameters included in the model are shown on the y axis. Credible intervals that do not cross the 0 line represent reliable effects at 95%. Each panel includes the numerical values of the medians and lower and upper boundaries of the 95% credible interval. The first level of each categorical variable is contrast-coded in negative.

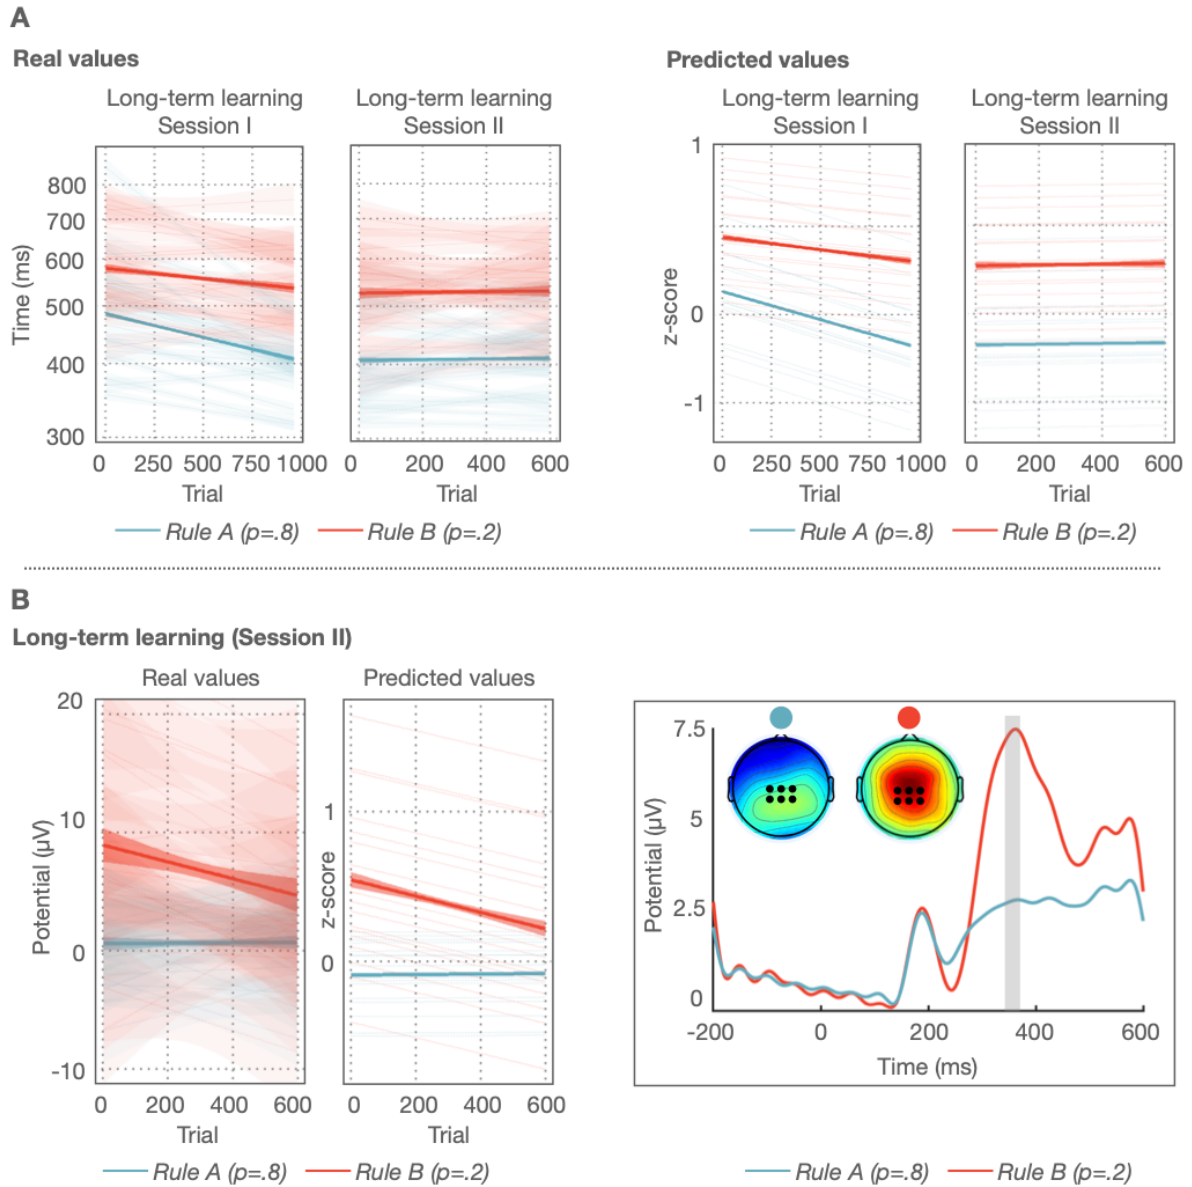


Figure 10. (A) Behavioural results obtained in the Long-term Learning phase Session I (without EEG recording) and Session II (with EEG recording). Overall and per-participant results are given for the data (left panel, RTs on semi-log scale) and for the model predictions (right panel, z-scores). (B) Variation of P3 amplitude throughout the Long-term Learning (Session II) phase with overall and per participant results for data and model predictions (left panel) at the central electrode cluster. ERP waveform at the central electrode cluster and scalp voltage distribution maps in the Long-term Learning (Session II) phase for targets presented according to rules A and B (right panel).