Title: Flexible structure learning under uncertainty

Rui Wang[1,2], Vael Gates[3], Yuan Shen[4], Peter Tino[5], Zoe Kourtzi[6]

1. State Key Laboratory of Brain and Cognitive Science, CAS Center for Excellence in Brain Science and Intelligence Technology, Institute of Psychology, Chinese Academy of Sciences, Beijing, 100101, China

2. Department of Psychology, University of Chinese Academy of Sciences, Beijing, 100049, China

3. Institute for Human-Centered AI, Stanford University, Stanford, CA 94305, USA

4. School of Science and Technology, Nottingham Trent University, Nottingham, NG11 8NF, UK

5. School of Computer Science, University of Birmingham, Birmingham, B15 2TT, UK

6. Department of Psychology, University of Cambridge, Cambridge, CB2 3EB, UK

Corresponding author:

Zoe Kourtzi

Department of Psychology

University of Cambridge

Downing Street, CB2 3EB

Cambridge, UK

Email: zk240@cam.ac.uk

Electronic copy available at: https://ssrn.com/abstract=4014051

Title: Flexible structure learning under uncertainty

Rui Wang[1,2], Vael Gates[3], Yuan Shen[4], Peter Tino[5], Zoe Kourtzi[6]

1. State Key Laboratory of Brain and Cognitive Science, CAS Center for Excellence in Brain Science and Intelligence Technology, Institute of Psychology, Chinese Academy of Sciences, Beijing, 100101, China

2. Department of Psychology, University of Chinese Academy of Sciences, Beijing, 100049, China

3. Institute for Human-Centered AI, Stanford University, Stanford, CA 94305, USA

4. School of Science and Technology, Nottingham Trent University, Nottingham, NG11 8NF, UK

5. School of Computer Science, University of Birmingham, Birmingham, B15 2TT, UK

6. Department of Psychology, University of Cambridge, Cambridge, CB2 3EB, UK

Corresponding author:

Zoe Kourtzi

Department of Psychology

University of Cambridge

Downing Street, CB2 3EB

Cambridge, UK

Email: zk240@cam.ac.uk

Electronic copy available at: https://ssrn.com/abstract=4014051

Acknowledgements

We would like to thank Ignacio Perez-Pozuelo, Andre Zylstra for help with data collection; Vasilis M. Karlaftis, Peng Cai and Yi Jiang for helpful discussions.

Funding

This work was supported by grants to (i) ZK from the Biotechnology and Biological Sciences Research Council (H012508, BB/P021255/1), the Wellcome Trust (205067/Z/16/Z), the European Community's Seventh Framework Programme (FP7/2007-2013) under agreement PITN-GA-2011-290011, and the National Research Foundation, Prime Minister's Office, Singapore, under its Campus for Research Excellence and Technological Enterprise (CREATE) programme; (ii) RW from Ministry of Science and Technology of China (2021ZD0203801), National Natural Science Foundation of China (31701003), the Youth Innovation Promotion Association of Chinese Academy of Sciences (2018115). For the purpose of open access, the author has applied for a CC BY public copyright license to any Author Accepted Manuscript version arising from this submission.

Electronic copy available at: https://ssrn.com/abstract=4014051

Highlights

1. Manipulating probabilistic contingencies, but not sequence rhythm, impairs structure learning.
2. Trial-by-trial feedback facilitates maximization and structure learning.
3. Adopting a decision strategy closer to maximization relates to better learning.
4. Selective attention and working memory account for individual variability in strategy and learning.

Key words: structure learning, uncertainty, perceptual decisions

Electronic copy available at: https://ssrn.com/abstract=4014051

Abstract

Experience is known to facilitate our ability to interpret sequences of events and make predictions about the future by extracting temporal regularities in our environments. Here, we ask whether uncertainty in dynamic environments affects our ability to learn predictive structures. We exposed participants to sequences of symbols determined by first-order Markov models and asked them to indicate which symbol they expected to follow each sequence. We introduced uncertainty in this prediction task by manipulating the: a) probability of symbol co-occurrence, b) stimulus presentation rate, c) feedback. Our results demonstrate that increasing the similarity in the probabilities of symbol co-occurrence impaired performance on the prediction task. In contrast, increasing uncertainty in stimulus presentation rate by introducing temporal jitter resulted in participants adopting a strategy closer to probability maximization than matching and improving in the prediction tasks. Next, we show that feedback plays a key role in learning predictive statistics. Trial-by-trial feedback yielded stronger improvement than block feedback or no feedback; that is, participants adopted a strategy closer to probability maximization and showed stronger improvement when trained with trial-by-trial feedback. Further, correlating individual strategy with learning performance showed higher improvement in structure learning for observers who adopted a strategy closer to maximization. Our results indicate that executive cognitive functions (i.e. selective attention, working memory) may account for this individual variability in strategy and structure learning ability. Taken together, our results provide evidence for flexible structure learning; individuals adapt their decision strategy closer to probability maximization, reducing uncertainty in temporal sequences and improving their ability to learn predictive statistics in variable environments.

Electronic copy available at: https://ssrn.com/abstract=4014051

Introduction

Successful everyday interactions entail that we identify spatiotemporal regularities (i.e. patterns that repeat frequently) in our cluttered and dynamic environments and exploit them to predict future events. Learning and experience are known to facilitate our ability to extract the environment's statistics (Aslin & Newport, 2012; Perruchet & Pacton, 2006). For example, humans become sensitive to stimuli (shapes, tones or syllables) that co-occur following a spatial or temporal pattern through repetitive exposure (Chun, 2000; Fiser & Aslin, 2002; J. R. Saffran, Aslin, & Newport, 1996; Jenny R Saffran, Johnson, Aslin, & Newport, 1999; N. B. Turk-Browne, Junge, & Scholl, 2005).

Our recent work demonstrates that individuals extract the statistics that govern the temporal structure of events and exploit them to make predictions about future events (R. Wang, Y. Shen, P. Tino, A. E. Welchman, & Z. Kourtzi, 2017). Further, we show that this learning of predictive structures relates to the decision strategy of individuals. In particular, previous work has highlighted the role of strategies in probabilistic learning and decision making (Acerbi, Vijayakumar, & Wolpert, 2014; Erev & Barron, 2005; Murray, Patel, & Yee, 2015; Schulze, van Ravenzwaaij, & Newell, 2015; Shanks, Tunney, & McCarthy, 2002). Humans and animals are known to engage in probability matching (match their choices probabilistically according to the underlying input statistics) or probability maximization (maximize their success by selecting the most probable outcomes) when making choices. Yet, the factors that determine individual decision strategies and influence learning ability remain largely unknown.

Previous work provides evidence for the role of uncertainty in perceptual decision making (Bach & Dolan, 2012). In particular, noisy sensory signals (Daikoku, 2018;

Electronic copy available at: https://ssrn.com/abstract=4014051

Dayan & Daw, 2008; Dosher & Lu, 1998; Hasson, 2017) or increased stochasticity in temporal sequences (A. C. Nobre & van Ede, 2018; Rolke & Hofmann, 2007) impact the difficulty of perceptual tasks. Further, feedback is known to play a key role in resolving uncertainty and facilitating perceptual decisions (Dayan & Abbott, 2001; Kluger & DeNisi, 1996; Petrov, Dosher, & Lu, 2005).

Here, we test whether uncertainty in sensory processing and feedback affect decision strategy and structure learning in the context of a sequence prediction task. In particular, we trained participants with temporal sequences comprising unfamiliar symbols and determined by first-order Markov models. Participants were exposed to these context-based statistics (i.e. symbol probability is contingent on previous symbols) and they were asked to judge whether a test symbol that followed the sequence presentation matched the expected symbol based on the preceding sequence. This sequence prediction task allows us to track participant responses over time and interrogate the decision strategy that individuals adopted during learning. We introduced uncertainty in the task by manipulating: a) the probability of symbol co-occurrence, b) the stimulus presentation rate, c) feedback. We reasoned that during training individuals will adapt their decision strategies and performance in the sequence prediction task. Our results demonstrate that: (1) increasing the similarity in the probabilities with which symbol contingencies appear in the sequence impaired performance on the prediction task; (2) increasing uncertainty in stimulus presentation rate by temporal jittering facilitated probability maximization and performance; (3) trial-by-trial feedback enhanced performance compared to block feedback or no feedback and facilitated probability maximization, while uncorrelated feedback resulted in limited improvement. Correlating individual strategies with learning performance showed that observers that adopted a strategy closer to maximization showed stronger performance improvement.

Electronic copy available at: https://ssrn.com/abstract=4014051

Finally, we show that attentional and working memory ability may account for individual differences in decision strategy and structure learning ability.

Materials and Methods

Observers

105 observers (40 males and 65 females, mean age = 22.1±0.3 years) participated in this study and they were randomly allocated into different experimental groups. All observers were naive to the aim of the study, had normal or corrected-to-normal vision and gave written informed consent. This study was approved by the University of Cambridge Ethics Committee and the institutional review board of the Institute of Psychology, Chinese Academy of Sciences.

Stimuli

Stimuli comprised 4 symbols chosen from Sabaean alphabet and Ndjuká syllabary (**Figure 1a**). These symbols were highly discriminable from each other and were unfamiliar to the observers. Each symbol was presented at 6.5° of visual angle in black on mid-grey background. Experiments were controlled using Matlab and the Psychophysics toolbox 3 (Brainard, 1997; Pelli, 1997). Stimuli were presented on a 21-inch CRT monitor (ViewSonic P225f , 1024×768 pixel, 85 Hz frame rate) at a distance of 60 cm.

Sequence design

Electronic copy available at: https://ssrn.com/abstract=4014051

We employed first-order Markov model (i.e. level-1) to generate probabilistic sequences (R. Wang et al., 2017). The level-1 Markov model produces a sequence of symbols, where the symbol at time i is determined probabilistically by the immediately preceding symbol. We refer to the symbol presented at time i, $s(i)$, as the *target* and to the previous symbol $s(i\text{-}1)$ as the *context*:

$$P(s(i)|s(i\text{-}1),s(i\text{-}2),\ldots,s(1))=P(s(i)|s(i\text{-}1)).$$

At each time point in the sequence, the symbol that follows a given context is determined probabilistically. The underlying Markov model can be represented through the associated context-conditional target probabilities. We used 4 symbols that we refer to as stimuli A, B, C and D. The correspondence between stimuli and symbols was counterbalanced across participants. Specifically, for level-1, the target depended on the item that immediately preceded it. Given a context (the last seen symbol), only one of two targets could follow (**Figure 1b**): one had a high probability of being presented (80% of occurrence in most groups, 60% in Group 2) and the other a low probability (20% of occurrence in most groups, 40% in Group 2). For example, when Symbol A was presented, only symbols B or C were allowed to follow, and B had a higher probability of occurrence than C.

Experimental Design:

We tested six groups of participants. Experiment 1 (Group 1: N=18) aimed to replicate our previous findings (R. Wang et al., 2017) and test whether learning is maintained over time. In Experiments 2-4, we manipulated: (1) symbol occurrence in sequences; (2) stimulus presentation rate; (3) feedback. In Experiment 2, we modified the

Electronic copy available at: https://ssrn.com/abstract=4014051

probability with which symbols appeared in the sequence. A new group of participants (Group 2: N = 18) were exposed to the sequences with probabilities of context-target contingencies at 60% vs. 40%, in contrast to 80% vs. 20% in Group 1. In Experiment 3 (Group 3: N = 18), visual stimuli appeared in a stream separated by jittered blank intervals, in contrast to Group 1 in which the stimuli were presented at a fixed interval. In Experiment 4, we examined the role of feedback in learning predictive structures: In contrast to participants in Group 1 who were trained with block feedback, participants in Group 4 (N = 18) were trained with trial-by-trial feedback), participants in Group 5 (N = 15) were trained without any feedback, and participant in Group 6 (N = 18) were trained with uncorrelated feedback.

All participants underwent six sessions: one session involved testing on cognitive tasks (i.e. working memory and selective attention), the remaining five sessions involved testing and training on the sequence prediction task using first-order Markov sequences. Before and after training (pre- and post-training sessions), participants were tested with structured sequences and random sequences (i.e., all four symbols were presented with equal probability 25% in a random order). To investigate whether the learning effect was maintained over time, ten observers in Group 1 were re-tested four weeks after training.

Training sessions: Training comprised 23 blocks of structured sequences (60 trials per block) that were conducted on four consecutive days. For each trial (**Figure 1a)**, a sequence of 9–13 stimuli appeared in the center of the screen, one at a time in a continuous stream, for 100 ms each followed by a central white fixation dot (interstimulus interval, ISI) for 400 ms on average. The ISI was fixed at 400 ms, except for Group 3 in which the ISI was jittered; that is the ISI in a given trial was chosen randomly from a uniform distribution of values ranging between 100 and 700 ms and

Electronic copy available at: https://ssrn.com/abstract=4014051

binned in temporal windows of 20 ms (i.e.100, 120, 140 ms etc). The end of each trial was indicated by a red-dot cue that was presented for 400 ms. Following this, all four symbols appeared in the center (2 x 2 grid) of the screen. Observers were asked to indicate which symbol they expected to appear following the preceding sequence by pressing a key corresponding to the location of the predicted symbol. Following the observer's response, a circle appeared on the selected stimulus for 300 ms to highlight the observer's choice. For Group 4 and Group 6, trial-by-trial feedback was provided by coloring this circle (green vs. red signified correct vs. incorrect responses, respectively). For Group 4 feedback matched the presented sequence (Group 4: trial feedback), while for Group 6 feedback was uncorrelated to observers' responses (Group 6: uncorrelated feedback). For other groups (Group 1-3,5), the color of the circle was always white, simply indicating the observer's choice rather than providing feedback. If no response was made within 2 s, a null response was recorded and the next trial started. For Group 1-3, observers were given feedback (i.e. score in the form of performance index (PI), see "Behavioural analysis") at the end of each block (block feedback). For Group 5 (no feedback), neither block feedback nor informative trial feedback was provided.

Test sessions: To compare performance before and after training, the pre- and post-training sessions included three blocks, that is, two blocks of structured sequences interleaved with one block of random sequences (i.e., all four symbols were presented with equal probability 25% in a random order). Participants were trained with structured sequences and tested with both structured and random sequences to ensure that training was specific to the trained sequences. Each block comprised 40 trials, during which participants performed the same sequence prediction task as in the training sessions.

Electronic copy available at: https://ssrn.com/abstract=4014051

The stimuli and procedure were identical to the training sessions but no feedback was given during test sessions.

Figure 1

Figure 1

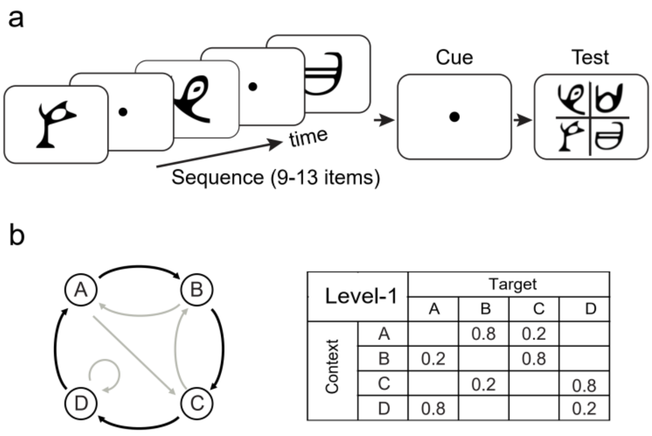

Level-1		Target			
		A	B	C	D
Context	A		0.8	0.2	
	B	0.2		0.8	
	C		0.2		0.8
	D	0.8			0.2

Figure 1. Trial and sequence design. (a) 9 to 13 symbols were presented one at a time in a continuous stream followed by a cue and the test display. (b) Sequence design. For the first-order Markov model (Level1), a diagram indicates states (circles) and transitional probabilities (black arrow: high probability, e.g., 80%; gray arrow: low probability, e.g., 20%). Transitional probabilities are shown in a four-by-four conditional probability matrix, with rows indicating temporal contexts and columns indicating the corresponding targets.

Cognitive testing

Memory: visual short-term memory

The working memory task was designed based on the sequential working memory task by Luck and Vogel (Luck & Vogel, 1997). Colored dots were displayed on a grey background for 500ms, followed by a 1000ms delay. After the delay, the dot display re-appeared with one of the dots highlighted by a white square. Participants reported

Electronic copy available at: https://ssrn.com/abstract=4014051

whether the highlighted dot had remained the same color on the second presentation. An initial display of two dots was used. We manipulated the number of dots in the display using a two-down one-up staircase, resulting in 70.7% performance. Working memory thresholds (i.e. number of dots in the display) were calculated by averaging the last two-third reversals in each staircase. For each trial, each dot was randomly assigned a color, and one dot was randomly chosen as the target. Each dot had a radius of 12 pixels and dots were displayed in random locations within a 10x10 grid (jittered ± 10 pixels). Each run consisted of 10 staircase reversals, and participants completed 3 runs, after which we computed the average threshold as their working memory score. In this task, a higher score (greater number of items in display) denotes better performance.

Attention: useful field of view

We used the Useful Field of View (UFOV; Visual Awareness Inc.) task to assess selective attention (Edwards et al., 2006). Each trial started with a fixation bounding box (1-s duration), followed by the test stimuli (variable duration between 16.7 and 500 ms), a white noise visual mask to control for after images (1-s duration) and the response screen (displayed until a response was made). The central stimulus (a silhouette of 2 cm × 1.5 cm of a car or a truck) was presented on a black background inside a white bounding box, with a simultaneously presented peripheral stimulus (2 cm × 1.5 cm silhouette of a vehicle) which was fixed at 11 cm from the central stimulus at one of the eight radial locations. The target stimuli were embedded in the context of distractors (47 triangles of the same size and luminance as the targets). Participants were asked to ignore the triangles and point out whether the central stimulus comprised a car or a truck, as well as the location of the peripheral target. Using a double-staircase method, the duration of the display within each task varied between 16.7 and 500 ms.

Electronic copy available at: https://ssrn.com/abstract=4014051

This allowed us to establish the minimal display duration at which the participant could correctly perform the tests 75 % of the time. Participants completed three runs, after which we computed the average threshold as their selective attention score. Thus, a lower score (shorter display duration) indicates better performance in this task.

Behavioural analysis

Performance index (PI): We assessed participant responses in a probabilistic manner, following our previous work (R. Wang et al., 2017). We computed a performance index per context that quantifies the minimum overlap (min: minimum) between the distribution of participant responses (Presp) and the distribution of presented targets (Ppres) estimated across 60 trials per block by:

PI(context) = $\sum_{target} min\ (Presp(target|context), Ppres(target|context))$

The overall performance index is then computed as the average of the performance indices across contexts, PI(context), weighted by the corresponding context probabilities P(context):

PI = $\sum_{context} PI(context) \cdot P(context)$

To compare across different conditions, we defined a normalized PI measure that quantifies participant performance relative to random guessing. We computed a random guess baseline; i.e. performance index PI_{rand} that reflects participant responses to targets with equal probability for each target for a given context for level-1(PI_{rand} = 0.45 for probability of 80% vs. 20%, PI_{rand} = 0.50 for probability of 60% vs. 40%). To correct for differences in random-guess baselines, we subtracted the random guess baseline

Electronic copy available at: https://ssrn.com/abstract=4014051

from the performance index ($PI_{normalized} = PI - PI_{rand}$). PI improvement was the difference in normalized PI between pre- and post-training sessions.

Strategy choice and strategy index

Following our previous work (R. Wang et al., 2017), we quantified each participant's strategy, by comparing individual participant response distributions (response-based model) to two baseline models: (i) probability matching, where probabilistic distributions are derived from the Markov models that generated the presented sequences (Model-matching) and (ii) a probability maximization model, where only the single most likely outcome is allowed for each context (Model-maximization). We used Kullback-Leiber (KL) divergence to compare the response distribution to each of these two models. KL is defined as follows:

$$KL = \sum_{context} M(context) \sum_{target} M(target|context) \, log \, (\frac{M(target|context)}{R(target)|context})$$

for level-1 model where R( ) and M( ) denote the probability distribution or conditional probability distribution derived from the human responses and the models (i.e. probability matching or maximization) respectively, across all the conditions.

We quantified the difference between the KL divergence from the response-based model to Model-matching and the KL divergence from the response-based model to Model-maximization. We refer to this quantity as strategy choice indicated by ΔKL(Model-maximization, Model-matching). We computed strategy choice per training block, resulting in a strategy curve across training for each individual participant. We then derived an individual strategy index by calculating the integral of each participant's strategy curve and subtracting it from the integral of the exact matching curve, as defined by Model-matching across training. We defined the integral

Electronic copy available at: https://ssrn.com/abstract=4014051

curve difference (ICD) between individual strategy and exact matching as the individual strategy index. Negative strategy index indicates a strategy closer to matching, while positive index indicates a strategy closer to maximization.

Results

Experiment 1: Learning temporal statistics with block feedback

To test whether individuals adapt to the environment's statistics, we trained eighteen participants (Group 1) on multiple training blocks over four sessions, during which they were presented with structured sequences of symbols that were determined by the first-order Markov model and were asked to perform a prediction task; that is, participants indicated the symbol they expected to appear following the preceding sequence. During the training phase, the visual stimuli were presented one after another at a fixed rate of 2Hz. Participants were given block feedback; that is, the Performance Index (PI) score (indicating how closely the probability distribution of participant responses matches the probability distribution of the presented symbols) was shown to the participants at the end of each block (i.e. 60 trials). To quantify the learning effect, we compared the normalized PI (i.e. after subtracting performance based on random guessing) before and after training. A repeated measures ANOVA showed a significant session effect ($F(1,17) = 36.72$, $p < 0.001$), indicating that repeated exposure facilitates learning of structured sequences (**Figure 2**). Specifically, most observers (13/18) improved in the prediction task with an average PI improvement of 30.7±2.8%; only five participants showed performance less than 10% above random guessing. The learning curves in **Figure 2a** indicate that performance improves throughout training. These results corroborated our

Electronic copy available at: https://ssrn.com/abstract=4014051

previous findings (R. Wang et al., 2017) showing that participants succeed in extracting regularities and making predictions about upcoming events.

Figure 2

Figure 2

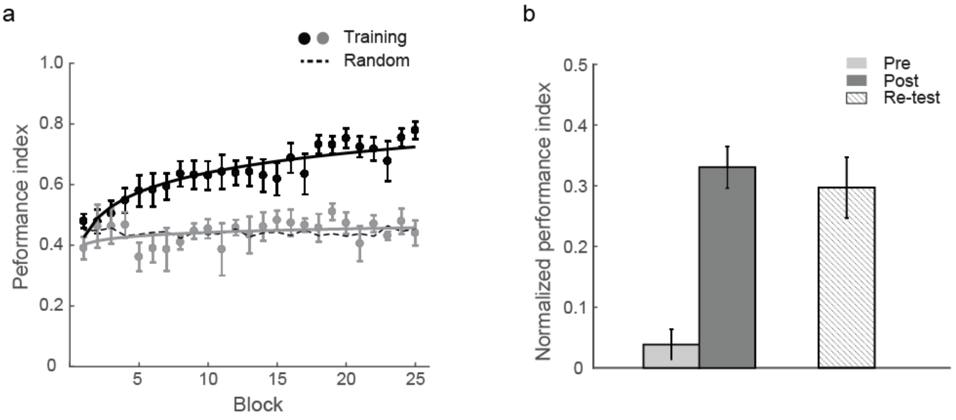

Figure 2. Experiment 1. (a) Mean Performance index across test and training blocks. Data are fitted separately for participants who improved during training (black symbols, N = 13) and those who did not improve (grey symbols, N = 5). Random guess baseline is indicated by dotted lines. (b) Normalized PI for test sessions. Data are shown before (gray bars), immediately after (black bars) and four weeks after training (dotted bars). Error bars indicate standard error of the mean.

To examine whether the learning effect we observed was maintained over time, ten participants were called back for an additional test session four weeks after training (spaced by 27.2±4.8 days). Performance in this test session was significantly higher than the pre-training test (F(1,9) = 38.54, p < 0.001) (**Figure 2b**). Mean PI improvement immediately after and four weeks after training was 29.2±3.4% and 25.9±4.2%,

Electronic copy available at: https://ssrn.com/abstract=4014051

respectively and did not differ significantly between these post-tests (t(9) = 1.06, p = 0.319), suggesting that the training-dependent improvement we observed was sustained for a prolonged time.

Experiment 2: Manipulating context-conditional probability of symbol occurrence

In this experiment we asked whether increasing uncertainty during the training by manipulating the context-conditional probability of symbol occurrence affects learning in the context of the prediction task. In particular we changed the symbol transitional probability from 80% vs. 20% (Experiment 1) to 60% vs.40%, while keeping the context and targets identical to the model used in Experiment 1. We hypothesized that decreasing the discriminability of contingency probabilities would impair learning. We trained a new group of eighteen observers (Group 2) on the prediction task using the less discriminable contingency probabilities (60% vs. 40%).

Figures 3a, 3c show that this manipulation resulted in low PI improvement (4.6 ± 2.2%) A two-way repeated measures ANOVA comparing performance before and after training between Group 1 (Experiment 1) and Group 2 (Experiment 2) showed a significant interaction of session and group (F(1,34) = 18.12, $p < 0.001$). There was no significant difference between groups for the pre-training performance (F(1,34) = 0.11, p = 0.738). In contrast, performance after training for Group 1 was significantly higher than Group 2 (F(1,34) = 16.58, $p < 0.001$). These results suggest that probability of context-target contingencies affects learning of temporal statistics; that is, making the probabilities of symbol co-occurrence less discriminable compromises performance and learning in the prediction task.

Electronic copy available at: https://ssrn.com/abstract=4014051

Experiment 3: Manipulating uncertainty in stimulus presentation rate by temporal jitter

Previous studies have shown that rhythmic stimulation that induces strong temporal expectation facilitates processing of events (Jones, Moynihan, MacKenzie, & Puente, 2002; Rohenkohl, Cravo, Wyart, & Nobre, 2012; Schroeder & Lakatos, 2009). Here, we tested whether disrupting rhythmic stimulation by introducing temporal jitter disrupts learning in the prediction task. In particular, we varied the ISI between successive stimuli in a trial. That is, in Group 1, the stimuli were presented at a fixed rate of 2Hz and the ISI was fixed at 400ms. In contrast, in Experiment 3 (N=18, Group 3) the ISI was jittered, ranging from 100 to 700ms. We reasoned that jittering the ISI would prevent temporal expectation and may impair learning of temporal statistics.

In contrast to this prediction, our results show that training resulted in a significant improvement for most participants in Group 3 (13/18, mean PI improvement: 25.6±3.8%) except five participants who showed only small improvement (i.e. performance after training less than 10% above random guessing) (**Figure 3b**). A mixed ANOVA comparing across Experiment 1 and Experiment 3 with Session (Pre vs. Post) and Group (Group 1 vs. Group 3) showed a significant main effect of session (F(1,34) = 72.81, $p < 0.001$), consistent with enhanced performance after training. There was no significant main effect of group (F(1,34) = 1.79, p = 0.190) nor interaction between session and group (F(1,34) = 0.22, p = 0.639), indicating similar improvement across groups despite temporal jitter in Group 3.

Decision strategies for learning temporal statistics

We next asked whether increasing uncertainty in the prediction task by manipulating the probability of context-target contingencies or the stimulus presentation rate affects participant decision strategies when making predictions. Previous work on probabilistic

Electronic copy available at: https://ssrn.com/abstract=4014051

learning and decision making has proposed that individuals adopt decision strategies ranging from matching to maximization when making probabilistic choices (Acerbi et al., 2014; Erev & Barron, 2005; Murray et al., 2015; Schulze et al., 2015; Shanks et al., 2002). We have previously shown that in the context of our prediction task, participants are exposed to stochastic sequences and use these strategies when learning the probabilities of different outcomes (R. Wang et al., 2017). Modelling the participants' responses allows us to quantify their decision strategy. Specifically, participants may adopt (1) probability matching (that is, match their choices to the relative probabilities of the context-target contingencies presented in the sequences); or (2) deviate from matching towards maximization (that is, choose the most likely outcome in a given context). To quantify these strategies, we computed a strategy index that indicates participant's preference (on a continuous scale) for responding using probability matching versus maximization. **Figure 3d** illustrates variability of strategy index for learners from Group 1 and Group 3. The strategy index for Group 1 was not significantly different from matching (that is, zero strategy index; $t(12) = 1.32$; $P = 0.213$), while the strategy index for Group 3 was significantly higher than zero ($t(12) = 8.93$; $P < 0.001$). Comparing individual strategy across groups showed significantly higher strategy index for Group 3 than Group 1 ($t(24) = 3.14$, $p = 0.004$), suggesting that disrupting rhythmic stimulation by temporal jitter results to a decision strategy closer to maximization.

Correlating PI with strategy index showed a significant positive relationship ($r = 0.312$, $p=0.022$ for Group 1-3, N=54), suggesting that maximization strategy relates to improved performance in the prediction task. This relationship may explain the surprising result we observed for Group 3; that is, performance in the prediction task improves despite temporal jitter. That is adopting a strategy closer to maximization may

Electronic copy available at: https://ssrn.com/abstract=4014051

facilitate learning when uncertainty in stimulus presentation rate is increased due to temporal jitter.

Figure 3

Figure 3

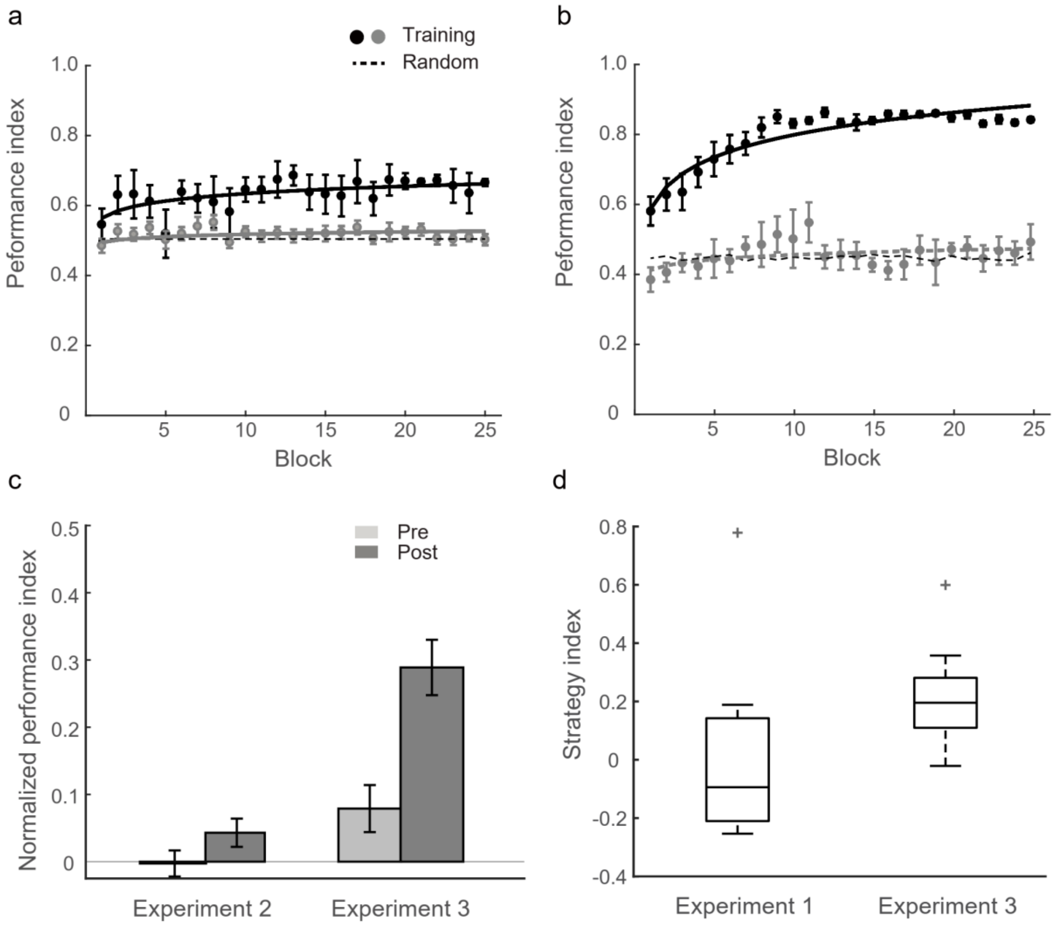

Figure 3. Experiment 2 and Experiment 3. (a) Mean Performance index across test and training blocks for Group 2 (Experiment 2). Data are fitted separately for participants who improved during training (black symbols, N = 5) and those who did not improve (grey symbols, N = 13). (b) Mean Performance index across test and training blocks for Group 3 (Experiment 3). Data are fitted separately for participants who improved during training (black symbols, N = 13) and those who did not improve (grey symbols, N = 5). (c) Normalized PI pre- and post-training for Group 2 and Group 3. Error bars indicate standard error of the mean. (d) Box plots of strategy

Electronic copy available at: https://ssrn.com/abstract=4014051

index show individual variability for learners in Group 3 and Group 1 (Experiment 1). The upper and lower error bars display the minimum and maximum data values, and the central boxes represent the interquartile range (25th–75th percentiles). The thick line in the central boxes represents the median. Crosses denote outliers.

Experiment 4: Manipulating feedback

Theoretical work has suggested that supervised, error-correcting learning mechanisms rely on external feedback (Dayan & Abbott, 2001). To understand the role of feedback in learning temporal statistics, we trained 3 additional groups of participants with a) trial-by-trial feedback based on the symbol expected by the pre-defined sequences (Group 4, **Figure 4a**); b) no feedback (Group 5, **Figure 4b**); c) uncorrelated feedback, that is, random trial-by-trial feedback that was uncorrelated to the observers' responses (Group 6, **Figure 4c**).

Figure 4d shows mean performances before and after training per group. A mixed two-way ANOVA showed a significant interaction of session and group (F(2,48) = 12.69, $p < 0.001$), suggesting that performance improvement differed across groups. PI improvement for Group 4, Group 5 and Group 6 was 33.2±3.6%, 19.0±1.1%, and 8.6±0.5% respectively. Specifically, trial-by-trial feedback (Group 4) resulted in most participants (17/18 learners) showing improvement in the task that was on average higher than the improvement observed for the other groups (Group 4 vs. Group 5: t(31) = 2.697, p =0.011; Group 4 vs. Group 6: t(34) = 5.043, $p < 0.001$). Most participants improved in the task even without feedback (12/15 learners) and there was no significant difference in performance between block feedback and no feedback (Group 1 vs. Group 5: t(31) = 0.805, p =0.427). However, providing uncorrelated feedback

Electronic copy available at: https://ssrn.com/abstract=4014051

resulted in limited improvement (i.e. two thirds participants showed < 10% improvement; Group 5 vs Group 6: t(31) = 2.064, p =0.047), and nearly half participants (8/18) showed performance less than 10% above random guessing in the prediction task after training.

We then compared participant decision strategies across groups to test whether feedback modulates decision strategy (**Figure 4e**). A one-way ANOVA on strategy index showed a significant effect of group (F(2,48) = 15.70, p < 0.001). For participants who trained with trial-by-trial feedback (Group 4), the strategy index was significantly higher than zero (t(17) = 5.225, p < 0.001) and higher than the strategy index for groups that trained with no feedback or uncorrelated feedback (Group 4 vs. Group 5: t(31) = 2.936, p = 0.006; Group 4 vs. Group 6: t(34) = 5.913, p < 0.001), suggesting that participants adopted a strategy closer to maximization. In contrast, participants who trained without feedback (Group 5) showed strategy index that did not differ significantly from matching (that is, zero strategy index, t(14) = 0.802, p = 0.436), suggesting that participants learned by matching the probability distribution of the presented context-target contingencies. Due to the lower number of participants who improved when trained with uncorrelated feedback (Group 6), the strategy index in this group was lower than zero (t(17) = -2.783, p = 0.013) and significantly lower than Group 5(t(31) = 2.129, p = 0.041). These results suggest that decreasing feedback uncertainty by providing trial-by-trial feedback facilitates maximization and learning of temporal statistics.

Figure 4

Electronic copy available at: https://ssrn.com/abstract=4014051

Figure 4

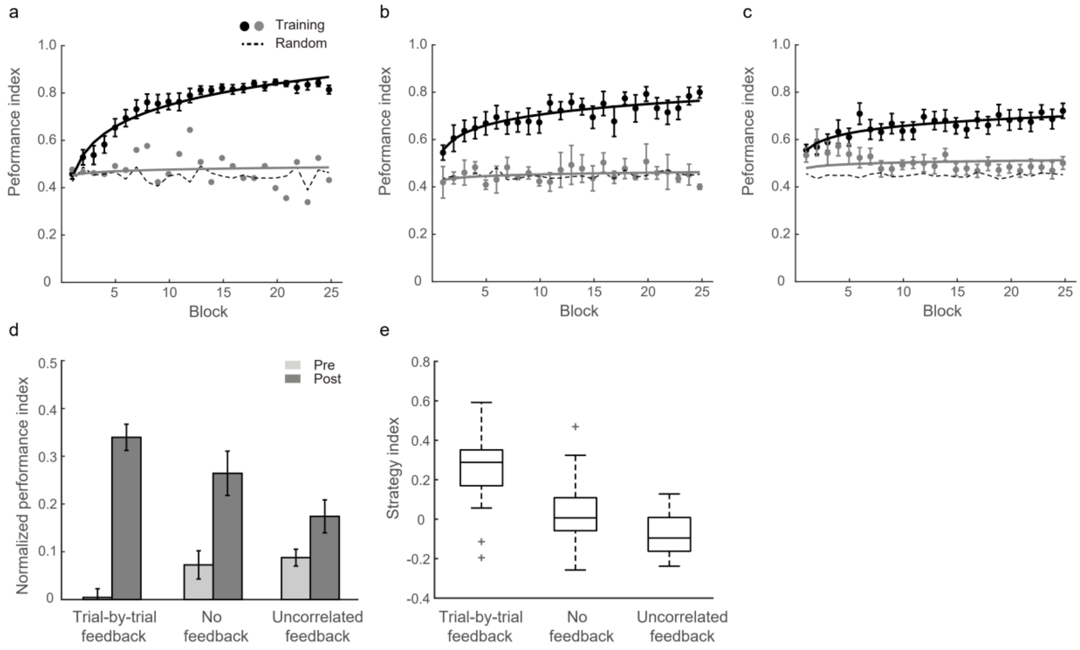

Figure 4. Experiment 4. (a) Mean Performance index across blocks for Group 4 (trial-by-trial feedback). Data are fitted separately for participants who improved during training (black symbols, N = 17) and one participant who did not improve (grey symbols). (b) Mean Performance index across blocks for Group 5 (no feedback). Data are fitted separately for participants who improved during training (N = 12) and those who did not improve (N = 3). (c) Mean Performance index across blocks for Group 6 (uncorrelated feedback). Data are fitted separately for participants who improved during training (N = 10) and those who did not improve (N = 8). (d) Normalized PI for test sessions in Groups 4,5,6. Error bars indicate standard error of the mean. (e) Box plots of strategy index show individual variability per group.

Correlating learning performance to strategies

We further tested whether individual strategies relate to learning performance. Combing data across experiments (N = 105), there was a significant correlation (**Figure 5a**) between participants' strategy index and behavioral improvement (i.e. difference in normalized PI between pre- and post-training sessions) (r = 0.471, p<0.001), indicating

Electronic copy available at: https://ssrn.com/abstract=4014051

that participants who adopt a strategy closer to maximization show higher improvement. Despite the fact that improvement varied across groups due to the experiment manipulation (a linear regression model showed group as a significant predictor on PI improvement, $R^2 = 0.325$, $F = 9.551$, $p<0.001$), including strategy index as an additional regressor significantly explained 7.9% more variance in performance improvement ($p = 0.001$, $\Delta R^2 = 0.079$). At first glance, this result may appear surprising, as exact matching is expected to result to 100% performance. Interrogating the response distribution across participants showed that most learners with high PI adopted a strategy towards maximization (i.e. responding more than 80% to high probability contingencies) over training. This result is in line with previous findings suggesting that probability maximization is favored when learning complex probabilistic tasks (Lagnado, Newell, Kahan, & Shanks, 2006; R. Wang et al., 2017).

Correlating cognitive abilities to learning temporal statistics

Finally, we asked whether cognitive control abilities (i.e. attention, working memory) relate to learning performance. Selective attention and working memory were assessed before training on the prediction task using the Useful Field of View task and visual short-term memory tasks respectively. We observed individual variability in cognitive tasks across participants. Performance in selective attention—as measured by SOA duration needed for separating targets from cluttered distractors —ranged from 16 to 165 ms, and performance in working memory—as measured by number of items which were correctly memorized —ranged from 3.43 to 9.5 number of dots. There was a significant correlation between selective attention and working memory scores ($r = -0.304$, $p = 0.002$) across participants in all groups. Further, we found: a) significant

Electronic copy available at: https://ssrn.com/abstract=4014051

correlation between selective attention scores ($r = -0.327$, $p = 0.001$) and working memory scores ($r = 0.234$, $p = 0.016$) with strategy index (**Figure 5b**) in the prediction task across all participants ($N = 105$), b) significant correlation between selective attention scores ($r = -0.280$, $p = 0.004$) and working memory scores ($r = 0.192$, $p = 0.050$) with PI improvement. These results suggest that participants with better attentional and working memory skills are more likely to adopt a strategy closer to maximization and improve in statistical learning.

Further, multiple regression analysis showed that group and cognitive abilities (selective attention, working memory) explained significantly ($F(7,97) = 5.31$, $p < 0.001$) 27.7 % of the variance in strategy index ($R = 0.526$), and group had the strongest impact on decision strategy ($p < 0.001$, $R^2 = 0.234$). Interestingly, we found that selective attention rather than working memory was better at predicting the strategy adopted during training; that is, excluding the variation accounted for by group (i.e. experimental manipulation), a model with selective attention as an additional regressor significantly explained 3.3% more of the variance in strategy index ($p = 0.038$). However, we did not observe a significant impact of working memory in the model ($p = 0.101$). These results suggest that selective attention is a key predictor of decision strategy; that is, it is likely that selecting the most probable outcomes when maximizing facilitates learning of temporal statistics.

Figure 5

Electronic copy available at: https://ssrn.com/abstract=4014051

Figure 5

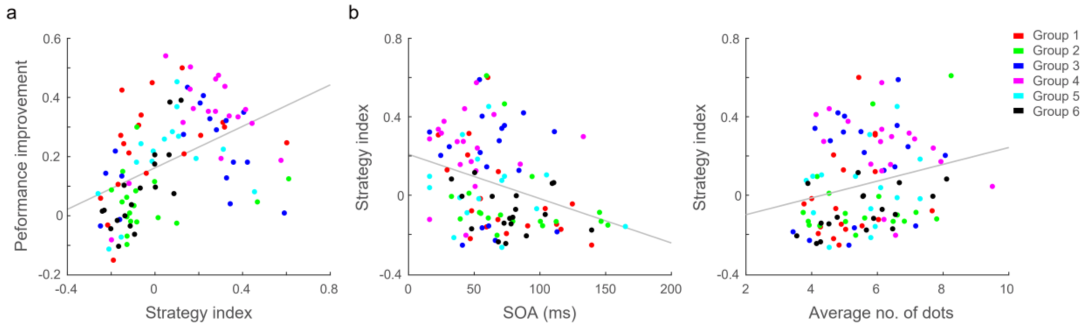

Figure 5. Correlations across participants in all groups (N=105). (a) Significant correlations of individual decision strategy and performance improvement (r = 0.471， CI= [0.331， 0.603]). A skipped Pearson correlation analysis using the Robust correlation toolbox (Pernet, Wilcox, & Rousselet, 2012) replicated this significant positive correlation following exclusion of six bivariate outliers (r = 0.623， CI= [0.508， 0.724]). Negative strategy-index values indicate a strategy closer to matching, while positive values indicate a strategy closer to maximization. (b) Correlating cognitive skills with decision strategy. Left, correlation of selective attention scores with strategy index. A lower score (SOA: stimulus onset asynchrony; i.e. shorter display duration) indicates better performance in the selective attention task that relates to decision strategy closer to maximization. Right, correlation of working memory scores with strategy index. A higher score (larger number of dots in the display) indicates better performance in the working memory task that relates positively with decision strategy closer to maximization. The color of the dots indicates participant group.

Discussion

Extracting the statistics governing event streams is critical for adaptive behavior in rapidly changing environments. Our findings demonstrate that exposure to temporal sequences facilitates our ability to extract their structure and predict upcoming events; an improvement that lasts for a prolonged period following training (up to 4 weeks).

Electronic copy available at: https://ssrn.com/abstract=4014051

We show that this learning of predictive structures is maintained under uncertainty that relates to the characteristics of the temporal sequences and task feedback. In particular, learners adapt their behavior to changes in the sequence design, rate of stimulus presentation and feedback. Further, attentional and working memory skills account for variability in decision strategy that strongly relates to individual structure learning ability. Our findings advance our understanding of structure learning in four main respects.

First, we show that less discriminable contingency probabilities compromised learning performance. This is consistent with previous work (Hasson, 2017; Okano, Daikoku, Ugawa, Kanai, & Yumoto, 2021; Thiessen, Kronstein, & Hufnagle, 2013), showing that probability of stimulus occurrence is key for extracting spatiotemporal structures. For example, previous studies have shown that sequences of syllables with high conditional probabilities are perceived to correspond to words, while syllable transitions with low predictability are more likely to be perceived as word-boundaries (Jenny R. Saffran, Newport, & Aslin, 1996). Our analysis showed that participants failed to extract the underlying first order Markov structure (i.e. identify the correct context-target contingencies) in sequences with highly similar contingency probabilities. In this model, the target depended on the immediately preceding stimulus. Specifically, given a context, only one of two targets (among all four symbols) could follow: one with a high probability and the other with a low probability. For example, for context A, the probability of AB, AC, AD, AA were 0.8, 0.2, 0 and 0, respectively. Tracking observers' responses across trials showed that decreasing the probability difference (i.e. changing the context-conditional probabilities from 80% vs. 20% to 60% vs.40%) resulted in impaired performance. That is, most participants were not able to predict the corresponding conditional probabilities nor extract the appropriate context-

Electronic copy available at: https://ssrn.com/abstract=4014051

target contingencies. These results indicate that probability of context-target contingencies plays a key role in learning temporal structure.

Second, we asked whether temporal uncertainty influences learning of predictive structure. Previous studies have shown that temporal variability and uncertainty disrupts temporal expectation and impairs performance (A. Nobre, Correa, & Coull, 2007). In contrast, presenting stimuli at a regular rhythm or at the expected time has been shown to facilitate action preparation and execution (e.g. reduced reaction times and saccade latencies) (Niemi & Näätänen, 1981) and enhance perceptual judgements (Lasley & Cohn, 1981; Rohenkohl et al., 2012; Rolke & Hofmann, 2007; Westheimer & Ley, 1996). In contrast to this previous work, we found that disrupting the rhythmic presentation of the sequence by introducing temporal jitter did not disrupt learning of predictive structures. Interestingly, learners presented with temporally jittered sequences adopted a strategy closer to maximization, suggesting that maximizing may facilitate learning of temporal structure under temporal uncertainty. It is likely that our participants focused on the probabilistic associations between symbols rather than the sequence rhythm, as our prediction task requires the participants to make an explicit judgement about the expected stimulus. This is consistent with previous work suggesting that humans are rational probabilistic learners and able to extract organized structures from ambiguous information (e.g. feature correlations in multidimensional sequences (N. Turk-Browne, Isola, Scholl, & Treat, 2008)) in a flexible manner (Aslin & Newport, 2012). Further, temporal jitter may result in increased cognitive load. Our results suggest that adopting a strategy closer to maximization facilitates structure learning under conditions of higher task demands, consistent with previous work showing that participants adopt a strategy closer to maximization when learning more

Electronic copy available at: https://ssrn.com/abstract=4014051

complex probabilistic sequences and tasks after training (Lagnado, Newell, Kahan, & Shanks, 2006; R. Wang et al., 2017).

Previous behavioral and neurophysiological studies have suggested that temporal uncertainty induced by varying the regularity of rhythmic stimulus streams influences stimulus processing at a perceptual level (Rohenkohl et al., 2012; Rolke & Hofmann, 2007; Schroeder, Wilson, Radman, Scharfman, & Lakatos, 2010). Entraining the brain to rhythmic events has been shown to facilitate sensory processing (Lakatos, Karmos, Mehta, Ulbert, & Schroeder, 2008; A. C. Nobre & van Ede, 2018; Schroeder & Lakatos, 2009). It is likely that irregular stimulus presentation disrupts perceptual processing in visual cortico-striatal circuits that have been shown to relate to structure learning based on probability matching. In particular, we have previously shown that distinct brain circuits relate to individual strategies for learning temporal statistics (Karlaftis et al., 2019; Rui Wang, Yuan Shen, Peter Tino, Andrew E. Welchman, & Zoe Kourtzi, 2017): probability matching engages occipitotemporal and ventral caudate regions, whereas maximization engages fronto-striatal circuits (i.e. dorsolateral prefrontal cortex, cingulate, sensory–motor regions, and putamen). Thus, it is possible that structure learning under temporal uncertainty recruits fronto-striatal circuits that support learning by maximization rather than matching facilitating learners to flexibly adapt their decision strategy and learn the environments statistics.

Third, we test whether feedback modulates decision strategy and learning of predictive statistics. Feedback is known to play a key role in learning new skills from simple feature processing to complex social interactions (Kluger & DeNisi, 1996). Theoretical work has proposed that supervised, error-correcting learning mechanisms rely on external feedback (Dayan & Abbott, 2001; Liu, Dosher, & Lu, 2014; Petrov et al., 2005). Yet, previous work on statistical learning has shown that learning of spatiotemporal

Electronic copy available at: https://ssrn.com/abstract=4014051

regularities may occur implicitly (i.e., by mere exposure rather than external feedback) (Perruchet & Pacton, 2006). Our results demonstrate that participants were able to extract the underlying sequence structure without any feedback; that is, participant who received no feedback or sparse performance feedback (i.e. mean performance feedback quantitatively across a block of sixty trials) performed similarly in the prediction task. However, trial-by-trial correct feedback enhanced task performance and resulted in learners adopting a strategy closer to maximization than matching. In contrast, random feedback that was uncorrelated to the participants responses compromised learning substantially. Our results are consistent with previous work showing that informative feedback shifts decision strategy towards maximization in probabilistic choice tasks (Shanks et al., 2002). That is, trial-by-trial feedback supports error correction, reducing uncertainty and facilitating a decision strategy that delivers best outcomes and increased reward.

Finally, we show that individual strategies and performance in learning temporal statistics correlate with attentional and working memory skills. The role of attention and working memory in statistical learning remains debated (Conway, 2020). It is possible that working memory is involved in the encoding of multiple sequence items, facilitating learning of temporal statistics; yet the role of working memory in sequence learning remains controversial (Janacsek & Nemeth, 2013). Further, some studies propose that selective attention may gate learning of statistics; that is, regularities are only learned when the stimuli are attended (N. B. Turk-Browne et al., 2005), while others argue that extracting regularities is a consequence of attentional processing (Pacton & Perruchet, 2008). Statistical learning has been proposed to involve a multicomponent learning system that relates to stimulus encoding, retention and abstraction, with each component of this system depending on attention or working memory to a different degree (Arciuli, 2017). Although attention and working memory have traditionally been considered to be distinct cognitive processes, recent studies propose an overlap between the brain systems that support these processes (Awh, Vogel, & Oh, 2006). Here, we provide evidence that participants with better attentional and working memory skills adopt a strategy closer to maximization and show improved structure learning. Interestingly, we demonstrate that selective attention is a stronger predictor of maximization strategy than working memory. This is likely due to the fact that extracting probabilistic conjunctions is more relevant in the context of our prediction task than memorizing sequences. Consistent with previous computational work proposing a key role of attentional selection in learning (Dayan, Kakade, & Montague, 2000; Yu & Dayan, 2005), our findings suggest that selective attention may facilitate the selection of the most probable outcomes (i.e. adopting a maximization strategy), reducing uncertainty and supporting learning of temporal statistics.

In sum, our findings provide evidence for flexible learning of predictive statistics; that is, individuals adapt their decision strategy to learn the underlying structure of events in the face of sensory or feedback-related uncertainty and predict upcoming events. Future work may investigate the brain plasticity mechanisms that mediate our ability for this flexible structure learning under uncertainty.

References

Acerbi, L., Vijayakumar, S., & Wolpert, D. M. (2014). On the origins of suboptimality in human probabilistic inference. *PLoS Comput Biol, 10*(6), e1003661. doi:10.1371/journal.pcbi.1003661

Arciuli, J. (2017). The multi-component nature of statistical learning. *Philos Trans R Soc Lond B Biol Sci, 372*(1711). doi:10.1098/rstb.2016.0058

Aslin, R. N., & Newport, E. L. (2012). Statistical learning from acquiring specific items to forming general rules. *Current directions in psychological science, 21*(3), 170-176.

Electronic copy available at: https://ssrn.com/abstract=4014051

Awh, E., Vogel, E. K., & Oh, S. H. (2006). Interactions between attention and working memory. *Neuroscience, 139*(1), 201-208. doi:https://doi.org/10.1016/j.neuroscience.2005.08.023

Bach, D. R., & Dolan, R. J. (2012). Knowing how much you don't know: a neural organization of uncertainty estimates. *Nature Reviews Neuroscience, 13*(8), 572-586. doi:10.1038/nrn3289

Brainard, D. H. (1997). The psychophysics toolbox. *Spatial vision, 10*, 433-436.

Chun, M. M. (2000). Contextual cueing of visual attention. *Trends in cognitive sciences, 4*(5), 170-178. doi:10.1016/s1364-6613(00)01476-5

Conway, C. M. (2020). How does the brain learn environmental structure? Ten core principles for understanding the neurocognitive mechanisms of statistical learning. *Neuroscience & Biobehavioral Reviews, 112*, 279-299. doi:https://doi.org/10.1016/j.neubiorev.2020.01.032

Daikoku, T. (2018). Neurophysiological Markers of Statistical Learning in Music and Language: Hierarchy, Entropy and Uncertainty. *Brain Sciences, 8*(6), 114. Retrieved from https://www.mdpi.com/2076-3425/8/6/114

Dayan, P., & Abbott, L. (2001). *Theoretical Neuroscience: Computational and Mathematical Modeling of Neural Systems* (Vol. 15): MIT Press.

Dayan, P., & Daw, N. D. (2008). Decision theory, reinforcement learning, and the brain. *Cognitive, Affective and Behavioral Neuroscience, 8*(4), 429-453. doi:10.3758/CABN.8.4.429

Dayan, P., Kakade, S., & Montague, P. R. (2000). Learning and selective attention. *Nat Neurosci, 3 Suppl*, 1218-1223. doi:10.1038/81504

Dosher, B. A., & Lu, Z.-L. (1998). Perceptual learning reflects external noise filtering and internal noise reduction through channel reweighting. *Proceedings of the National Academy of Sciences, 95*(23), 13988-13993. doi:10.1073/pnas.95.23.13988

Edwards, J. D., Ross, L. A., Wadley, V. G., Clay, O. J., Crowe, M., Roenker, D. L., & Ball, K. K. (2006). The useful field of view test: Normative data for older adults. *Archives of Clinical Neuropsychology, 21*(4), 275-286. doi:https://doi.org/10.1016/j.acn.2006.03.001

Erev, I., & Barron, G. (2005). On adaptation, maximization, and reinforcement learning among cognitive strategies. *Psychol Rev, 112*(4), 912-931. doi:10.1037/0033-295x.112.4.912

Fiser, J., & Aslin, R. N. (2002). Statistical learning of higher-order temporal structure from visual shape sequences. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 28*(3), 458.

Hasson, U. (2017). The neurobiology of uncertainty: implications for statistical learning. *Philos Trans R Soc Lond B Biol Sci, 372*(1711). doi:10.1098/rstb.2016.0048

Janacsek, K., & Nemeth, D. (2013). Implicit sequence learning and working memory: Correlated or complicated? *Cortex, 49*(8), 2001-2006. doi:https://doi.org/10.1016/j.cortex.2013.02.012

Jones, M. R., Moynihan, H., MacKenzie, N., & Puente, J. (2002). Temporal aspects of stimulus-driven attending in dynamic arrays. *Psychol Sci, 13*(4), 313-319. doi:10.1111/1467-9280.00458

Karlaftis, V. M., Giorgio, J., Vértes, P. E., Wang, R., Shen, Y., Tino, P., . . . Kourtzi, Z. (2019). Multimodal imaging of brain connectivity reveals predictors of individual decision strategy in statistical learning. *Nature Human Behaviour, 3*(3), 297-307. doi:10.1038/s41562-018-0503-4

Kluger, A., & DeNisi, A. (1996). The Effects of Feedback Interventions on Performance: A Historical Review, a Meta-Analysis, and a Preliminary Feedback Intervention Theory. *Psychological bulletin, 119*, 254-284. doi:10.1037/0033-2909.119.2.254

Electronic copy available at: https://ssrn.com/abstract=4014051

Lagnado, D., Newell, B., Kahan, S., & Shanks, D. (2006). Insight and strategy in multiple-cue learning. *Journal of experimental psychology. General, 135*, 162-183. doi:10.1037/0096-3445.135.2.162

Lakatos, P., Karmos, G., Mehta, A. D., Ulbert, I., & Schroeder, C. E. (2008). Entrainment of neuronal oscillations as a mechanism of attentional selection. *Science, 320*(5872), 110-113. doi:10.1126/science.1154735

Lasley, D. J., & Cohn, T. (1981). Detection of a luminance increment: effect of temporal uncertainty. *Journal of the Optical Society of America, 71*(7), 845-850. doi:10.1364/JOSA.71.000845

Liu, J., Dosher, B., & Lu, Z.-L. (2014). Modeling trial by trial and block feedback in perceptual learning. *Vision Research, 99*, 46-56. doi:https://doi.org/10.1016/j.visres.2014.01.001

Luck, S. J., & Vogel, E. K. (1997). The capacity of visual working memory for features and conjunctions. *Nature, 390*(6657), 279-281. doi:10.1038/36846

Murray, R. F., Patel, K., & Yee, A. (2015). Posterior Probability Matching and Human Perceptual Decision Making. *PLoS Comput Biol, 11*(6), e1004342. doi:10.1371/journal.pcbi.1004342

Niemi, P., & Näätänen, R. (1981). Foreperiod and simple reaction time. *Psychological bulletin, 89*(1), 133-162. doi:10.1037/0033-2909.89.1.133

Nobre, A., Correa, A., & Coull, J. (2007). The hazards of time. *Curr Opin Neurobiol, 17*(4), 465-470. doi:10.1016/j.conb.2007.07.006

Nobre, A. C., & van Ede, F. (2018). Anticipated moments: temporal structure in attention. *Nature Reviews Neuroscience, 19*(1), 34-48. doi:10.1038/nrn.2017.141

Okano, T., Daikoku, T., Ugawa, Y., Kanai, K., & Yumoto, M. (2021). Perceptual uncertainty modulates auditory statistical learning: A magnetoencephalography study. *International Journal of Psychophysiology, 168*, 65-71. doi:https://doi.org/10.1016/j.ijpsycho.2021.08.002

Pacton, S., & Perruchet, P. (2008). An attention-based associative account of adjacent and nonadjacent dependency learning. *J Exp Psychol Learn Mem Cogn, 34*(1), 80-96. doi:10.1037/0278-7393.34.1.80

Pelli, D. G. (1997). The VideoToolbox software for visual psychophysics: Transforming numbers into movies. *Spatial vision, 10*(4), 437-442.

Pernet, C. R., Wilcox, R., & Rousselet, G. A. (2012). Robust correlation analyses: false positive and power validation using a new open source matlab toolbox. *Front Psychol, 3*, 606. doi:10.3389/fpsyg.2012.00606

Perruchet, P., & Pacton, S. (2006). Implicit learning and statistical learning: One phenomenon, two approaches. *Trends in cognitive sciences, 10*(5), 233-238.

Petrov, A. A., Dosher, B. A., & Lu, Z. L. (2005). The dynamics of perceptual learning: an incremental reweighting model. *Psychol Rev, 112*(4), 715-743. doi:10.1037/0033-295x.112.4.715

Rohenkohl, G., Cravo, A. M., Wyart, V., & Nobre, A. C. (2012). Temporal expectation improves the quality of sensory information. *J Neurosci, 32*(24), 8424-8428. doi:10.1523/jneurosci.0804-12.2012

Rolke, B., & Hofmann, P. (2007). Temporal uncertainty degrades perceptual processing. *Psychonomic Bulletin & Review, 14*(3), 522-526. doi:10.3758/BF03194101

Saffran, J. R., Aslin, R. N., & Newport, E. L. (1996). Statistical learning by 8-month-old infants. *Science, 274*(5294), 1926-1928. Retrieved from http://www.ncbi.nlm.nih.gov/entrez/query.fcgi?cmd=Retrieve&db=PubMed&dopt=Citation&list_uids=8943209

Saffran, J. R., Johnson, E. K., Aslin, R. N., & Newport, E. L. (1999). Statistical learning of tone sequences by human infants and adults. *Cognition, 70*(1), 27-52.

Electronic copy available at: https://ssrn.com/abstract=4014051

Saffran, J. R., Newport, E. L., & Aslin, R. N. (1996). Word Segmentation: The Role of Distributional Cues. *Journal of Memory and Language, 35*(4), 606-621. Retrieved from http://www.sciencedirect.com/science/article/pii/S0749596X96900327

Schroeder, C. E., & Lakatos, P. (2009). Low-frequency neuronal oscillations as instruments of sensory selection. *Trends Neurosci, 32*(1), 9-18. doi:10.1016/j.tins.2008.09.012

Schroeder, C. E., Wilson, D. A., Radman, T., Scharfman, H., & Lakatos, P. (2010). Dynamics of Active Sensing and perceptual selection. *Curr Opin Neurobiol, 20*(2), 172-176. doi:10.1016/j.conb.2010.02.010

Schulze, C., van Ravenzwaaij, D., & Newell, B. R. (2015). Of matchers and maximizers: How competition shapes choice under risk and uncertainty. *Cogn Psychol, 78*, 78-98. doi:10.1016/j.cogpsych.2015.03.002

Shanks, D. R., Tunney, R. J., & McCarthy, J. D. (2002). A Re-examination of Probability Matching and Rational Choice. *J. Behav. Dec. Making, 15*, 233–250.

Thiessen, E. D., Kronstein, A. T., & Hufnagle, D. G. (2013). The extraction and integration framework: a two-process account of statistical learning. *Psychol Bull, 139*(4), 792-814. doi:10.1037/a0030801

Turk-Browne, N., Isola, P., Scholl, B., & Treat, T. (2008). Multidimensional Visual Statistical Learning. *J Exp Psychol Learn Mem Cogn, 34*, 399-407. doi:10.1037/0278-7393.34.2.399

Turk-Browne, N. B., Junge, J. A., & Scholl, B. J. (2005). The automaticity of visual statistical learning. *Journal of Experimental Psychology: General, 134*(4), 552-564. doi:10.1037/0096-3445.134.4.552

Wang, R., Shen, Y., Tino, P., Welchman, A. E., & Kourtzi, Z. (2017). Learning predictive statistics from temporal sequences: Dynamics and strategies. *J Vis, 17*(12), 1. doi:10.1167/17.12.1

Wang, R., Shen, Y., Tino, P., Welchman, A. E., & Kourtzi, Z. (2017). Learning Predictive Statistics: Strategies and Brain Mechanisms. *The Journal of neuroscience : the official journal of the Society for Neuroscience, 37*(35), 8412-8427. doi:10.1523/JNEUROSCI.0144-17.2017

Westheimer, G., & Ley, E. (1996). Temporal uncertainty effects on orientation discrimination and stereoscopic thresholds. *J Opt Soc Am A Opt Image Sci Vis, 13*(4), 884-886. doi:10.1364/josaa.13.000884

Yu, A. J., & Dayan, P. (2005). Uncertainty, neuromodulation, and attention. *Neuron, 46*(4), 681-692. doi:10.1016/j.neuron.2005.04.026

Electronic copy available at: https://ssrn.com/abstract=4014051