

Learning predictive statistics from temporal sequences: Dynamics and strategies

Rui Wang*

Key Laboratory of Mental Health, Institute of Psychology,
Chinese Academy of Sciences, Beijing, China,
Department of Psychology, University of Cambridge,
Cambridge, UK



Yuan Shen*

Department of Mathematical Sciences,
Xi'an Jiaotong-Liverpool University, Suzhou, China
School of Computer Science, University of Birmingham,
Birmingham, UK



Peter Tino

School of Computer Science, University of Birmingham,
Birmingham, UK



Andrew E. Welchman

Department of Psychology, University of Cambridge,
Cambridge, UK



Zoe Kourtzi*

Department of Psychology, University of Cambridge,
Cambridge, UK



Human behavior is guided by our expectations about the future. Often, we make predictions by monitoring how event sequences unfold, even though such sequences may appear incomprehensible. Event structures in the natural environment typically vary in complexity, from simple repetition to complex probabilistic combinations. How do we learn these structures? Here we investigate the dynamics of structure learning by tracking human responses to temporal sequences that change in structure unbeknownst to the participants. Participants were asked to predict the upcoming item following a probabilistic sequence of symbols. Using a Markov process, we created a family of sequences, from simple frequency statistics (e.g., some symbols are more probable than others) to context-based statistics (e.g., symbol probability is contingent on preceding symbols). We demonstrate the dynamics with which individuals adapt to changes in the environment's statistics—that is, they extract the behaviorally relevant structures to make predictions about upcoming events. Further, we show that this structure learning relates to individual decision strategy; faster learning of complex structures relates to selection of the most probable outcome in a given context (maximizing) rather than matching of the exact sequence statistics. Our findings provide evidence for alternate routes to learning of behaviorally relevant statistics that facilitate our ability to predict future

events in variable environments.

Introduction

Extracting structure from initially incomprehensible streams of events is fundamental to a range of human abilities, from navigating in a new environment to learning a language. These skills rely on identifying spatial and temporal regularities, often with minimal explicit feedback (Aslin & Newport, 2012; Perruchet & Pacton, 2006). The human brain appears expert at learning contingencies between co-occurring stimuli on the basis of mere exposure. For instance, structured patterns become familiar after simple exposure to items (shapes, tones, or syllables) that co-occur spatially or follow in a temporal sequence (Chun, 2000; Fiser & Aslin, 2002a; Saffran, Aslin, & Newport, 1996; Saffran, Johnson, Aslin, & Newport, 1999; Turk-Browne, Junge, & Scholl, 2005).

Previous work on human statistical learning has focused on repetitive patterns and associative pairings. However, event structures in the natural environment typically comprise regularities of variable complexity, from simple repetition to complex probabilistic com-

Citation: Wang, R., Shen, Y., Tino, P., Welchman, A. E., & Kourtzi, Z. (2017). Learning predictive statistics from temporal sequences: Dynamics and strategies. *Journal of Vision*, 17(12):1, 1–16, doi:10.1167/17.12.1.

doi: 10.1167/17.12.1

Received June 13, 2017; published October 2, 2017

ISSN 1534-7362 Copyright 2017 The Authors



binations. For instance, when learning a new piece of music, we can make use of the intrinsic structure that ranges from tones to melodies (Fitch & Martins, 2014). Or, when reading about a new topic, we first extract information about the key components and then their interdependencies that together explain particular phenomena. To account for the full range of learning behaviors, we therefore need to understand the processes involved in extracting information of variable complexity.

Here we investigate the dynamics of learning predictive structures by modeling whether and when participants extract the structure that governs temporal sequences that change in their complexity. To allow them to do so unencumbered by past experience, we tested participants with sequences of unfamiliar symbols, where the sequence structure changed unbeknownst to the participants (Figure 1). We increased sequence complexity by manipulating the memory order (i.e., context length) of the Markov model used to generate the sequences. In particular, we presented participants with sequences that were determined first by frequency statistics (i.e., occurrence probability per symbol) and then by more complex context-based statistics (i.e., the probability of a given symbol appearing depends on the preceding symbols). Participants performed a prediction task in which they indicated which symbol they expected to appear following exposure to a sequence of variable length. Following previous statistical learning paradigms, participants were exposed to the sequences without trial-by-trial feedback.

Our results show that the participants' ability to extract behaviorally relevant temporal statistics improved with training. To understand the dynamics of this structure learning, we track human predictions as they evolve over time (i.e., during exposure to the sequences). In particular, we compare the sequence of button presses made by the participants to the presented temporal sequences and test whether and when participants' judgments approximate the Markov model that produced the presented sequences. Using this approach, we show that participants adapt to changes in the environment's statistics and exploit previous knowledge of similar, but simpler, statistics when learning higher order structures. Further, we show that learning of predictive structures relates to individual decision strategy; some individuals used a probability-maximization strategy (i.e., extracting the most probable outcome in a given context), while others chose to match the exact sequence statistics. We demonstrate that faster learning of complex structures is associated with selecting the most probable outcomes in a given context, suggesting that attempting to learn all possible statistical contingencies may limit the ability to learn higher order structures.

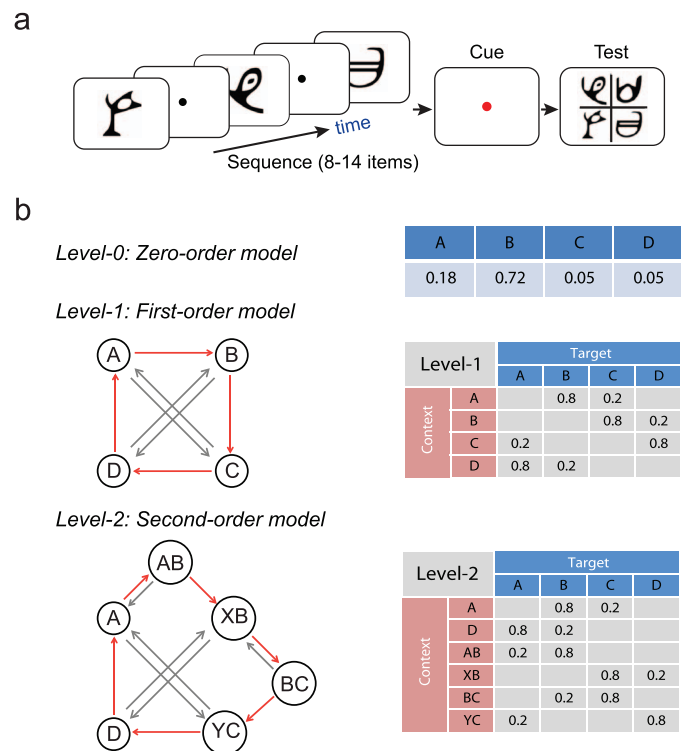


Figure 1. Trial and sequence design. (a) Eight to 14 symbols were presented one at a time in a continuous stream followed by a cue and the test display. (b) Sequence design. For the zero-order model (Level 0): Different states (A, B, C, D) are assigned to four symbols with different probabilities. For first- (Level 1) and second- (Level 2) order models, diagrams indicate states (circles) and conditional probabilities (red arrow: high; gray arrow: low). Transitional probabilities were arranged in a 4×4 (Level 1) or 4×6 (Level 2) conditional-probability matrix.

Material and methods

Observers

We tested 50 participants (mean age = 22.9 years)—Experiment 1, Group 0: $n = 19$; Experiment 2, Group 1: $n = 8$; Group 2: $n = 12$; no-training control experiment, $n = 11$. All observers were unaware of 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 Birmingham Ethics Committee.

Stimuli

Stimuli comprised four symbols chosen from the Ndjuká syllabary (Figure 1a; Turk-Browne, Scholl, Chun, & Johnson, 2009). The symbols (black on a mid-gray background; size = 8.5°) were highly discriminable

and were unfamiliar to the participants. Experiments were controlled using MATLAB and the Psychophysics Toolbox 3 (Brainard, 1997; Pelli, 1997). Stimuli were presented on a gamma-corrected 21-in. ViewSonic P225f monitor (1,280 × 1,024 pixel; 0.3 × 0.3 mm/pixel; 85-Hz refresh rate). Experiments were conducted in a dark room and the viewing distance was 45 cm.

Sequence design

To generate probabilistic sequences that differed in complexity, we used a temporal Markov model and manipulated the memory order of the sequence, which we refer to as the context length. The Markov model determines a sequence of symbols, where the symbol at time i is determined probabilistically by the previous k symbols. We refer to the symbol $s(i)$ presented at time i as the *target* and to the preceding k -tuple of symbols ($s(i-1)$, $s(i-2)$, ..., $s(i-k)$) as the *context*. The value of k is the order or level of the model

$$\begin{aligned} P(s(i)|s(i-1), s(i-2), \dots, s(1)) \\ = P(s(i)|s(i-1), s(i-2), \dots, s(i-k)), k < i. \end{aligned}$$

The simplest $k = 0^{\text{th}}$ order model is a random memoryless source. This generates, at each time point i , a symbol according to symbol probability $P(s)$, without taking account of the previously generated symbols.

The order $k = 1$ Markov model generates symbol $s(i)$ at each time i conditional on the previous symbol $s(i-1)$. This introduces a memory in the sequence; that is, the probability of a particular symbol at time i depends on the preceding symbol $s(i-1)$. Unconditional symbol probabilities $P(s(i))$ for the case $k = 0$ are replaced with conditional ones $P(s(i)|s(i-1))$.

We applied the same logic to higher memory orders: When $k = 2$, the probability of a symbol at time i depends on the two preceding symbols $s(i-1)$, $s(i-2)$: $P(s(i)|s(i-1), s(i-2))$. That is, the memory in the sequence is deeper and the number of conditioning contexts increases with memory depth k .

At each time point in the sequence, the symbol that follows a given context is determined probabilistically, making the Markov sequences stochastic. The underlying Markov model can be represented through the associated context-conditional target probabilities. We used four symbols that we refer to as stimuli A, B, C and D. The correspondence between stimuli and symbols was counterbalanced across participants.

For Level 0, the Markov model was based on the probability of symbol occurrence: One symbol had a high probability of occurrence and one low, while the remaining two symbols appeared rarely (Figure 1b). For example, the probabilities of occurrence for the

four symbols A, B, C, and D were 0.18, 0.72, 0.05, and 0.05, respectively. Presentation of a given symbol was independent of the stimuli that preceded it.

For Level 1, the target depended on the immediately preceding stimulus (Figure 1b). Given a context (the last-seen symbol), only one of two targets could follow: One had a high probability of being presented and the other a low probability (e.g., 80% vs. 20%). 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.

For Level 2, the Markov model contained temporal contexts of variable length (Figure 1b), extending the Level 1 model. That is, the Markov model included both first- and second-order contexts (i.e., the target symbols depended on the preceding two symbols). As with the Level 1 model, given a specific context, only two symbols were allowed to follow, one with a high and one with a low probability (e.g., 80% vs. 20%). The target probabilities for contexts with B as the last symbol (i.e., AB, BB, CB, DB) were constrained by allowing only two sets of conditional target probabilities, namely $P(s|AB)$ and $P(s|XB)$, where s is the target symbol (A, B, C, or D) and X stands for any other symbol apart from A (i.e., B, C, or D). The same structure was imposed for second-order contexts with C as the last symbol. In this case, the two sets of conditional target probabilities were $P(s|BC)$ and $P(s|YB)$, where Y stands for any other symbol apart from B (i.e., A, C, or D). To discriminate between contexts that shared the same last symbol (i.e., XB vs. AB, and YC vs. BC), different targets were assigned to each context (one with high and one with low probability). For example, the allowed targets following XB were C and D, while the targets for context AB were B and A. To ensure that learning was not biased by differences in context probability, the four Level 1 contexts (A, B, C, D) appeared at an equal 25% probability, and the six Level 2 contexts (A, AB, XB, BC, YC, D) appeared at similar probabilities (0.19, 0.19, 0.16, 0.16, 0.15, and 0.15, respectively).

To test whether participants adapt to changes in the temporal structure, we ensured that the sequences across levels were matched for properties (i.e., number or identity of symbols) other than context length. Further, we designed the stochastic sources from which the sequences were generated so that the context-conditional uncertainty remained highly similar across levels. In particular, for the zero-order source only two symbols were likely to occur most of the time; the remaining two symbols had very low probability (0.05). This was introduced to ensure that there was no difference in the number of symbols presented across levels. Of the two dominant symbols, one was more probable (probability 0.72) than the

other (probability 0.18). This structure is preserved in the Markov chain of order 1 or 2, where conditional on the previous symbols, only two symbols were allowed to follow, one with higher probability (0.80) than the other (0.20). This ensures that the structure of the generated sequences across levels differed predominantly in memory order (i.e., context length) rather than context-conditional probability.

Experimental design

We generated probabilistic sequences of symbols that differed in their complexity using temporal Markov models—that is, sequences determined by simple frequency statistics (Level 0) and more complex sequences defined by context-based statistics (Levels 1 and 2). Manipulating the memory depth of the stochastic source that generated the sequences allowed us to systematically vary the context length of the sequences (Figure 1, Supplementary Material). In Experiment 1 (Group 0), we trained participants with sequences that changed in their complexity starting from Level 0 and then moving to Level 1 and Level 2 sequences. In Experiment 2, we tested two additional groups of participants: Group 1 trained first at Level 1 and then Level 2; Group 2 trained only at Level 2. For each level, observers completed a minimum of three and a maximum of five training sessions (840–1,400 trials). Training at each level ended when participants reached plateau performance (i.e., performance did not change significantly for two sessions). A posttraining test session followed training per level, during which observers were presented with sequences determined by the statistics of the trained level (90 trials). Before and after training (pre- and posttraining sessions), participants were tested with sequences from all three levels (30 trials per level). Overall, Group 0 completed 13–15 training sessions and five test sessions (on average 23.3 days); Group 1 completed 8–10 training sessions and four test sessions (on average 15.6 days); Group 2 completed four or five training sessions and three test sessions (on average 9.5 days). Further, to ensure that any changes observed across time were a result of active training, we performed a no-training control experiment. Specifically, participants were tested on all three levels in two behavioral sessions that were separated by a period (27.9 days on average) comparable to that between the pre- and posttraining sessions for Group 0. The stimuli, sequences, and procedure matched the first and last test sessions in Experiment 1, but no training took place between these two sessions.

Training sessions

Each training session comprised five blocks of structured sequences (56 trials per block) and lasted 1 hr. To ensure that sequences in each block were representative of the Markov-model order per level, we generated 10,000 Markov sequences per level comprising 672 stimuli per sequence. We then estimated the Kullback–Leibler divergence (KL divergence) between each example sequence and the generating source. In particular, for Level 0 sequences this was defined as

$$KL = \sum_{\text{target}} Q(\text{target}) \log \frac{Q(\text{target})}{P(\text{target})},$$

and for Level 1 and 2 sequences it was defined as

$$KL = \sum_{\text{context}} \left(Q(\text{context}) \cdot \sum_{\text{target}} Q(\text{target}|\text{context}) \times \log \frac{Q(\text{target}|\text{context})}{P(\text{target}|\text{context})} \right),$$

where $P()$ refers to probabilities or conditional probabilities derived from the presented sequences and $Q()$ refers to those specified by the source. We selected 50 sequences with the lowest KL divergence (i.e., these sequences closely matched the Markov model per level). The sequences presented to the participants during the experiments were selected randomly from this sequence set.

For each trial, a sequence of 8–14 stimuli appeared in the center of the screen, one at a time in a continuous stream, for 300 ms each followed by a central white fixation dot (interstimulus interval) for 500 ms (Figure 1a). This variable trial length ensured that observers maintained attention during the whole trial. Each block comprised an equal number of trials with the same number of stimuli. The end of each trial was indicated by a red-dot cue that was presented for 500 ms. Following this, all four symbols were shown in a 2×2 grid. The positions of test stimuli were randomized from trial to trial. 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. Observers learned a stimulus–key mapping during the familiarization phase: 8, 9, 5, and 6 on the number pad corresponded to the four positions of the test stimuli—upper left, upper right, lower left, and lower right, respectively. After the observer’s response, a white circle appeared on the selected stimulus for 300 ms to indicate the observer’s choice, followed by a fixation dot for 150 ms (intertrial interval) before the start of the next trial. If no response was made within 2 s, a null response was recorded and the next trial started. Participants were

given feedback (i.e., score in the form of a performance index; see Data analysis) at the end of each block—rather than per-trial error feedback—that motivated them to continue with training.

Test sessions

Test sessions were conducted at the beginning and end of Experiments 1 and 2. Pre- and posttraining test sessions comprised nine runs (i.e., three runs per level). Intermediate test sessions (i.e., test sessions after training per level) included nine runs with sequences from the trained level. Each run comprised five blocks of structured and five blocks of random sequences presented in random order (two trials per block; a total of 10 structured and 10 random trials per run). For random sequences the four symbols were presented with equal probability in a random order. Each trial comprised a sequence of 10 symbols that were presented for 250 ms each, separated by a blank interval during which a white fixation dot was presented for 250 ms. Following the sequence, a response cue (central red dot) appeared on the screen before the four test stimuli were displayed for 1.5 s. No feedback was given during the test sessions.

Data analysis

Performance index

We assessed participant responses in a probabilistic manner. For each context, we computed the absolute Euclidean distance between the distribution of participant responses and the distribution of presented targets estimated across 56 trials per block:

$$\text{AbDist}(\text{context}) = \sum_{\text{target}} |P_{\text{resp}}(\text{target}|\text{context}) - P_{\text{pres}}(\text{target}|\text{context})|,$$

where the sum is over targets from the symbol set A, B, C, and D. We estimate AbDist per context for each block. We quantified the minimum overlap between these two distributions by computing a Performance Index (PI) per context:

$$\text{PI}(\text{context}) = \sum_{\text{target}} \min(P_{\text{resp}}(\text{target}|\text{context}), P_{\text{pres}}(\text{target}|\text{context})).$$

Note that $\text{PI}(\text{context}) = 1 - \text{AbDist}(\text{context})/2$. The overall performance index is then computed as the average of the performance indices across contexts, $\text{PI}(\text{context})$, weighted by the corresponding stationary context probabilities:

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

To compare across different levels, we defined a normalized PI measure that quantifies participant performance relative to random guessing. We computed a random-guess baseline—that is, performance index PI_{rand} —that reflects participant responses to targets with equal probability of 25% for each target per trial for Level 0 ($\text{PI}_{\text{rand}} = 0.53$) and equal probability for each target for a given context for Levels 1 and 2 ($\text{PI}_{\text{rand}} = 0.45$). To correct for differences in random-guess baselines across levels, we subtracted the random-guess baseline from the performance index ($\text{PI}_{\text{normalized}} = \text{PI} - \text{PI}_{\text{rand}}$).

Strategy choice and strategy index

To quantify each observer's strategy, we compared individual participant responses to probability matching, where probability distributions are derived from the Markov models that generated the presented sequences (matching), and probability maximization, where only the single most likely outcome is allowed for each context (maximization). We used KL divergence to compare the response distribution to matching versus maximization. KL is defined as follows:

$$KL = \sum_{\text{target}} M(\text{target}) \log \left(\frac{M(\text{target})}{R(\text{target})} \right)$$

for the Level 0 model and

$$KL = \sum_{\text{context}} M(\text{context}) \sum_{\text{target}} M(\text{target}|\text{context}) \times \log \left(\frac{M(\text{target}|\text{context})}{R(\text{target}|\text{context})} \right)$$

for the Levels 1 and 2 model, where $R(\cdot)$ and $M(\cdot)$ denote the probability distribution or conditional probability distribution derived from the human responses and probability matching versus maximization respectively, across all the conditions.

We quantified the difference between the KL divergence from maximization and matching to the response-based distribution, respectively. We refer to this quantity as *strategy choice*, indicated by $\Delta\text{KL}(\text{maximization}, \text{matching})$. We updated the strategy choice per trial and averaged across blocks, 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 matching across training. We defined the integral-curve difference between individual strategy and exact matching as the individual strategy index.

Results

Experiment 1: Behavioral performance

Previous studies have compared learning of different spatiotemporal contingencies in separate experiments across different participant groups (Fiser & Aslin, 2002a, 2005). Here, to investigate whether individuals extract changes in structure, we presented the same participants with sequences that changed in structure unbeknownst to them (Figure 1a). We parameterized structure complexity based on the memory order of the Markov models used to generate the sequences—that is, the degree to which the presentation of a symbol depended on the history of previously presented symbols (Figure 1b). We first presented participants with simple zero-order sequences (Level 0), followed by more complex first- and second-order sequences (Level 1, Level 2), as previous work has shown that temporal dependencies are more difficult to learn as their length increases (van den Bos & Poletiek, 2008) and training with simple dependencies may facilitate learning of more complex contingencies (Antoniou, Ettliger, & Wong, 2016). Zero-order sequences (Level 0) were contextless—that is, the presentation of each symbol depended only on the probability of occurrence of each symbol. First- and second-order sequences were governed by context-based statistics—that is, the presentation of a particular symbol was conditionally dependent on the previously presented symbols (i.e., context length of 1 or 2). Participants were presented with first-order (Level 1: context length of one stimulus) followed by variable-order (Level 2: context length of one or two stimuli) context–target contingencies. We measured participant performance in the prediction task before and after training.

As the sequences we employed were stochastic, we developed a probabilistic measure to assess participants' performance in the prediction task. Specifically, we computed a performance index (PI) that indicates how closely the distribution of participant responses matched the probability distribution of the presented symbols. This is preferable to a simple measure of accuracy because the probabilistic nature of the sequences means that the correct upcoming symbol is not uniquely specified; thus, designating a particular choice as correct or incorrect is often arbitrary.

Our results showed fast learning initially (i.e., enhanced performance in the first two training blocks compared to the pretraining test) that was followed by further improvement during the rest of the training (Figure 2a). This is consistent with the time course demonstrated by previous perceptual-learning studies (Karni & Sagi, 1993). Comparing normalized perfor-

mance (i.e., after subtracting random guessing) before and after training showed that participants were able to learn the presented sequences (only one participant showed less than 10% improvement after four training sessions for Level 2). A repeated-measures ANOVA with session (pre-, posttest) and complexity level (0, 1, 2) as factors showed significant main effects of session, $F(1, 18) = 145.8, p < 0.001$, and level, $F(1, 18) = 19.0, p < 0.001$, consistent with enhanced performance after training and increasing task difficulty for higher order sequences. Further, the lack of a significant interaction between session and level, $F(2, 36) = 2.40, p = 0.106$, suggests similar improvement across levels.

The learning functions in Figure 2a highlight that performance improves through training. Next we directly assessed how well participants were able to extract structures that were predictive of upcoming events. Figure 2b shows that the participants' ability to extract the most frequently presented symbols (Level 0) or context–target contingencies (Levels 1 and 2) improved with training across levels. When participants were presented with sequences of variable context length (Level 2), they maintained good performance for the first-order contingencies and also improved in extracting second-order contingencies.

Finally, we asked whether these learning effects were specific to the trained sequences. First, we contrasted performance accuracy on structured versus random sequences before and after training sessions. We found significant interactions between session and sequence, indicative of effects specific to the structured sequences—Level 0: $F(1, 18) = 9.17, p = 0.007$; Level 1: $F(1, 18) = 83.8, p < 0.001$; Level 2: $F(1, 18) = 61.7, p < 0.001$. Second, we conducted a no-training control experiment. Participants ($n = 11$) were tested with structured sequences in two sessions, but they did not receive training between sessions. Our results showed no significant main effect of session, $F(1, 10) = 0.12, p = 0.736$, or level, $F(1, 10) = 1.84, p = 0.205$, nor a significant interaction between session and level, $F(1, 10) = 1.16, p = 0.308$, indicating that improvements were specific to trained sequences rather than a result of repeated exposure during the pre- and posttraining sessions.

Response tracking

To quantify our results, we tracked the participants' responses across trials using a weighted combination (i.e., mixture) of Markov processes (i.e., zero-, first-, second-order). Previous work has used a Hebbian process to account for perceptual learning without explicit feedback (Liu, Lu, & Doshier, 2010; Petrov, Doshier, & Lu, 2005, 2006). For our purposes, however, capturing the dynamics of participants'




Figure 2. Experiment 1: Behavioral performance. (a) Performance index for Group 0 ($n = 19$) across training (solid circles) blocks, pretraining test (Pre: open squares), and posttraining test (Post: open squares). The performance index expresses the absolute distance (proportion overlap) between the distribution of participant responses and the distribution of presented targets. Overall performance index is calculated as the weighted average across context probabilities. Data are fitted for participants who improved during training (black circles). Data are also shown for one participant who did not improve during training (Level 2, gray symbols). Error bars show standard error of the mean. (b) Response probabilities for individual targets (Level 0) or conditional probabilities of context–target contingencies (Levels 1 and 2) across training blocks. Red lines indicate targets or context–target contingencies with the highest (conditional) probability (i.e., 0.72 for Level 0 and 0.8 for Levels 1 and 2), blue lines indicate the second-highest (conditional) probabilities (i.e., 0.18 for Level 0 and 0.2 for Levels 1 and 2), and green lines indicate targets or context–target contingencies that appear rarely (i.e., 0.05) or not at all. For Level 2, first- and second-order contexts are presented separately (dashed vs. solid lines).

responses as they learn to condition their responses on higher order statistics is difficult for a Hebbian process, due to the limited discrete data (i.e., one response per trial) during the learning process.

Following previous work on the learning of visual statistics (Droll, Abbey, & Eckstein, 2009; Eckstein, Abbey, Pham, & Shimozaki, 2004), we used a Bayesian process to adjust the mixture coefficient weights assigned to these component Markov processes during training (Supplementary Material). In particular, we extracted changes in participants' responses over time that relate to the rule used to generate the sequences—that is, memory or context length (e.g., the current target depends on the last symbol or the last two symbols)—and to the contingencies between individual

stimuli in the sequence (e.g., last stimulus was A, so next is likely to be B).

Extracting context length from participants' responses

First, we asked whether participants were able to extract the correct context length during training. In particular, a significant increase in the mixture coefficient for a given Markov order (e.g., Level 1) provides an indication that participants use a given memory length (e.g., context length 1) when responding. As the participants learned, we dynamically tracked whether and when the memory (context length) in their

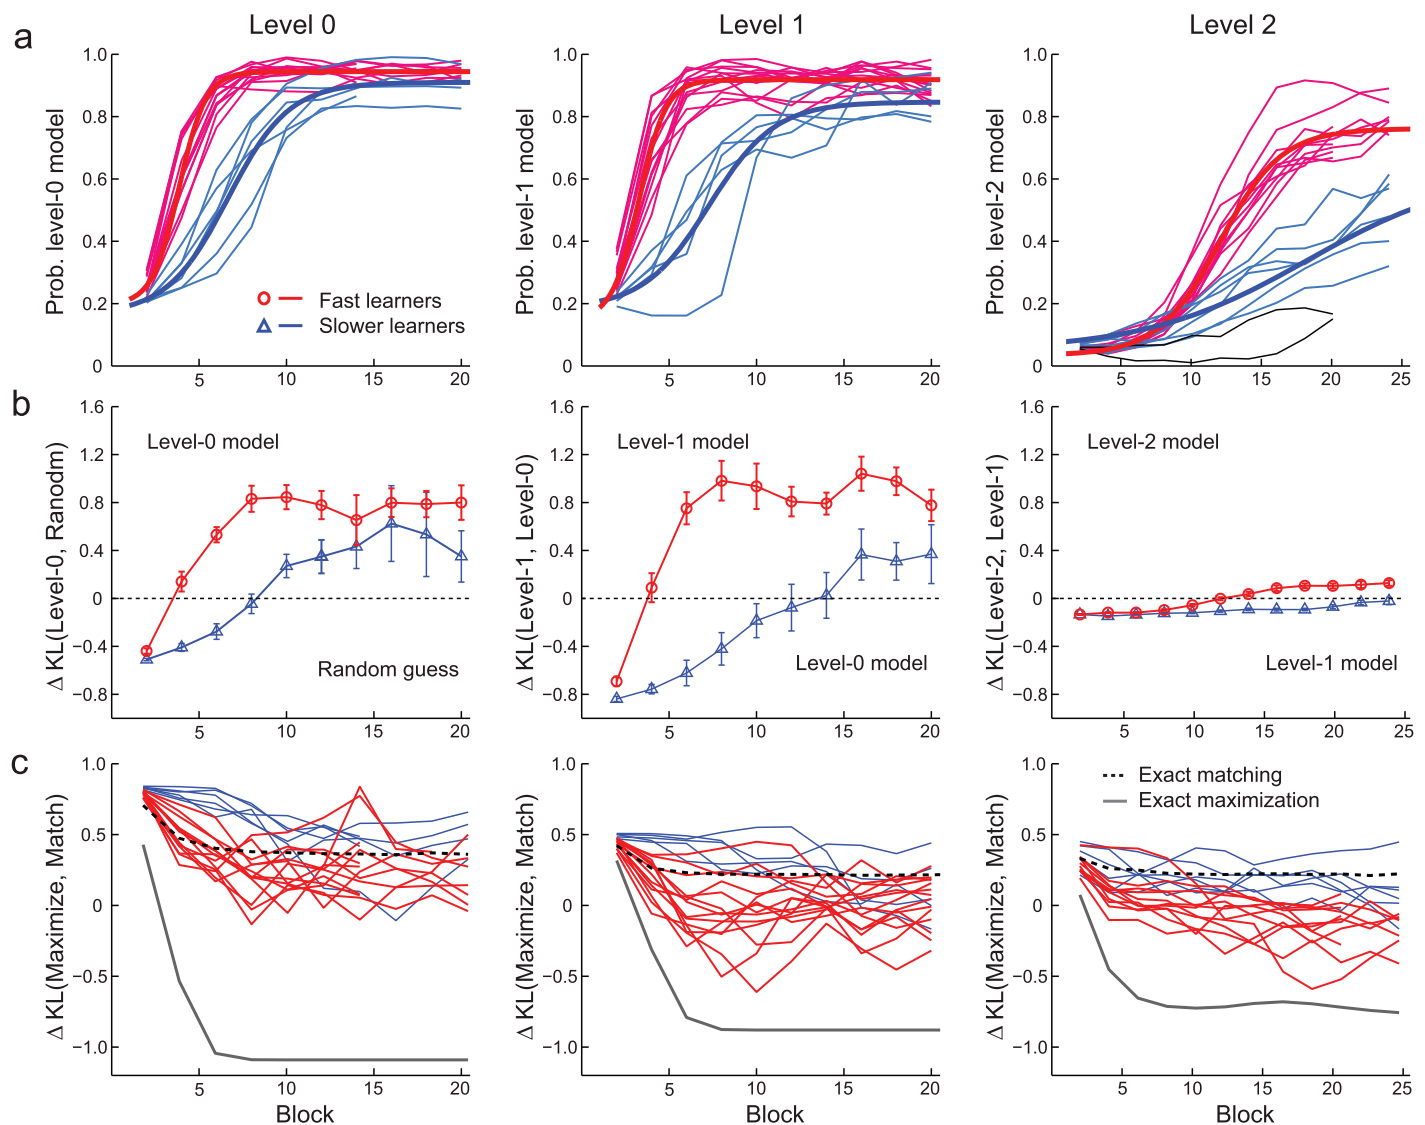


Figure 3. Experiment 1: Response tracking. (a) Functional clustering analysis (Group 0) showed two data clusters, indicated in red (Level 0: $n = 13$, Level 1: $n = 14$, Level 2: $n = 11$) versus blue (Level 0: $n = 6$, Level 1: $n = 5$, Level 2: $n = 6$). Mixture coefficient curves are shown for each individual participant; bold curves indicate sigmoid fits to each cluster. Data are also shown for two participants (black lines) who showed less than a 25% probability of extracting the correct context length at the end of training. (b) Learning predictive probabilities. ΔKL curves between the predictive mixture model for each level and baseline models across training blocks. ΔKL values above zero indicate that the participant responses approximated the Markov model that generated the sequences. Average data are shown per participant cluster (i.e., red vs. blue). Note: The smaller ΔKL values and error bars for Level 2 reflect small differences between Level 1 and Level 2 models; yet fast learners show higher values than zero, indicating that they are able to learn second-order context–target contingencies. Error bars show the standard error of the mean. (c) Strategy choice, as indicated by comparing (ΔKL) matching versus maximization for each participant per cluster (i.e., red vs. blue).

responses changed. In particular, we traced the evolution of the coefficients of the individual mixture components across training blocks. Mixture coefficient curves for individual participants followed a sigmoid shape, indicating changes in the context length extracted by the observers during training; we refer to these curves as *learning curves*. This analysis (Figure 3a) revealed that most participants became better at extracting the correct context length during training,

except two participants (gray lines for Level 2 in Figure 3a) who showed less than a 25% probability of selecting the correct context length. Further, comparing learning rate—as determined by the sigmoid mixture coefficient curves—across levels (0, 1, and 2) showed significantly slower learning rates for higher order than simpler sequences, $F(2, 49) = 23.7$, $p < 0.001$.

A notable feature of the learning curves in Figure 3a is the variability in learning rates between different

participants: Some individuals extracted the correct context length earlier in the training than others. To characterize prototypical learning profiles, we performed a functional clustering analysis of the learning curves (Supplementary Material). We found that two clusters were adequate to capture the individual variability in the data (Supplementary Figure S1). Given the apparent difference between participant groups in the speed of extracting the correct context length, we refer to these clusters as fast and slower learners. Supplementary Figure S2 shows differences in the learning rate of the more probable contingencies between the two clusters, confirming that some learners extracted the behaviorally relevant statistics faster than others.

We took a number of steps to validate our response-tracking analysis in a controlled manner. As a first step, we applied this analysis to random responses. We found no evolution of the coefficients of the individual mixture components, suggesting that the changes revealed using the participants' data do not simply reflect the dynamics of parameter initialization. We also tested our response-tracking analysis on responses generated by a synthetic learner (Supplementary Material), controlling for key parameters (learning rate and memory-order transition point). We varied the synthetic learner's parameters and recorded the sequence of predictions it made. This test showed that we could recover the key parameters that determined the synthetic learner's predictions (Supplementary Figure S3).

Extracting predictive contingencies from participants' responses

For individuals to succeed in the prediction task, they needed to extract not only the appropriate context length but also the correct conditional probabilities (i.e., context–target contingencies). To capture the dynamics of learning predictive contingencies, we sought to quantify the relationship between the participants' responses and the Markov models used to generate stimulus sequences. For each Markov order level, we considered two alternative models: the correct model order (e.g., Level 1 choices for Level 1 sequences) or a lower order approximation based on the previously trained sequence level (e.g., Level 0 choices for Level 1 sequences). We initially favored the lower order approximation to prevent emulating lower order structure using a higher order model. Using a Bayesian updating process, we obtained evidence that allowed us to discern whether responses were governed by a lower or a higher order process. We quantified how close participants' behavior was to a particular model using the Kullback–Leibler (KL) divergence statistic. We then contrasted KL statistics (i.e., slope of

Δ KL learning curves) to test which model the participants' responses approximated (Figure 3b). A two-way ANOVA showed a significant interaction between complexity level (0, 1, 2) and cluster (fast vs. slower learners), $F(2, 49) = 3.90$, $p = 0.027$, suggesting that individuals who extracted the correct context length early in the training also learned the appropriate context–target contingencies. Further, we observed a main effect of level—fast learners: $F(2, 49) = 39.0$, $p < 0.001$; slower learners: $F(2, 49) = 4.90$, $p = 0.012$ —suggesting that learning the correct predictive contingencies was more difficult for higher order sequences.

Previous work (Jensen, Boley, Gini, & Schrater, 2005) has demonstrated that temporal structure can be extracted without an explicit representation of the underlying model based on computing the entropy of excerpts from temporal sequences. We implemented an entropy-based approach and showed that it could recover first- and second-order contexts from the participant responses (Supplementary Figure S4). However, we found that this approach was limited in tracking the learning dynamics, as it required more trials to extract learning strategies from participant responses (i.e., there were insufficient participant responses to reliably estimate entropy in the first 10 blocks of trials).

Strategies for probability learning: Matching versus maximization

As the Markov models that generated stimulus sequences were stochastic, participants needed to learn the probabilities of different outcomes to succeed in the prediction task. Motivated by previous work on decision making in the context of cognitive (Shanks, Tunney, & McCarthy, 2002) and sensorimotor tasks (Acerbi, Vijayakumar, & Wolpert, 2014; Eckstein et al., 2013; Murray, Patel, & Yee, 2015), we formulated two possible strategies for making predictions. First, participants might use *probability maximization*, whereby they would always select the most probable outcome in a particular context. Alternatively, they might learn the relative probabilities of each symbol—for example, $p(A) = 0.18$, $p(B) = 0.72$, $p(C) = 0.05$, $p(D) = 0.05$ —and respond so as to reproduce this distribution, a strategy referred to as *probability matching*.

To quantify participants' strategies across training, we computed a strategy index that indicates each participant's preference (on a continuous scale) for responding using probability matching versus maximization (Figure 3c). We found that for Level 0 sequences, participants adopted a strategy that was closer to probability matching than maximization, suggesting that they solved the task by memorizing the frequency with which each symbol occurred. However, for Levels 1 and 2 they

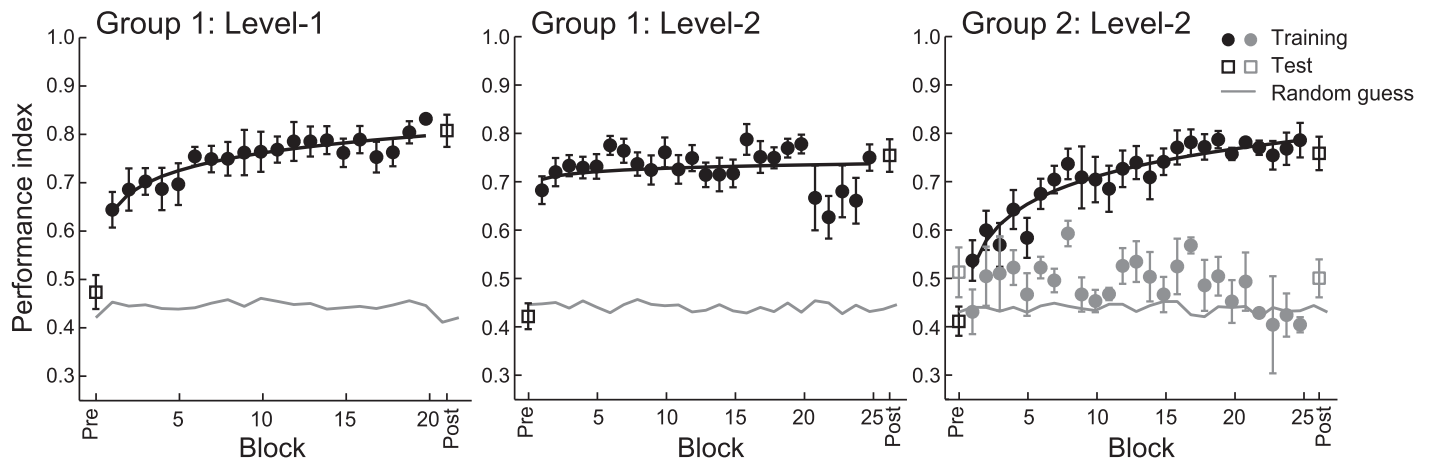


Figure 4. Experiment 2: Behavioral performance. Data for Group 1 ($n = 8$; Levels 1 and 2) and Group 2 ($n = 12$; Level 2). Performance index is shown across training (solid circles) blocks, pretraining test (Pre: open squares), and posttraining test (Post: open squares). Fitted data are shown for participants who improved during training (black circles). Data are also shown for participants ($n = 4$) in Group 2 who did not improve during training (Level 2, gray symbols). Error bars show standard error of the mean.

shifted toward maximization. Comparing individual strategy across levels and participant clusters showed a significant main effect of complexity level, $F(2, 49) = 12.2$, $p < 0.001$, suggesting that participants' strategies shifted closer to maximization for higher order sequences. Further, a significant main effect of cluster, $F(1, 49) = 60.9$, $p < 0.001$, indicates that fast learners who extracted the correct context length early in training deviated from matching and adopted a strategy closer to maximization. The lack of a significant interaction between cluster and level, $F(2, 18) = 0.025$, $p = 0.915$, suggests that each cluster of participants adopted a similar strategy across levels (i.e., closer to maximization for fast than for slower learners).

Despite greater maximization at higher complexities, we note that participants did not achieve optimal maximization performance (Figure 3c). Maximization is typically observed under supervised or reinforcement learning paradigms (Shanks et al., 2002), so it is perhaps not surprising that our participants did not achieve exact maximization, as trial-by-trial feedback was not provided. Moreover, the tendency for participants to respond using probability matching may be higher when individual elements are clearly discriminable (i.e., our symbols) but nevertheless ambiguous because different processes can give rise to similar sequences of symbols (as in our sequence-generation process; Murray et al., 2015). Our findings are consistent with previous studies showing that participants adopt a strategy closer to matching when learning a simple probabilistic task in the absence of trial-by-trial feedback (Shanks et al., 2002). However, for more complex probabilistic tasks, participants weight their responses toward the most-likely outcome (i.e., adopt a strategy closer to maximization) after training (Lagnado, Newell, Kahan, & Shanks, 2006).

Experiment 2: Behavioral performance

We next asked whether learning of simple structures facilitates subsequent learning of complex structures. In Experiment 2, we tested two additional participant groups who started training from Level 1 (Group 1) or Level 2 (Group 2) rather than Level 0. We then compared performance in Groups 1 and 2 with performance by participants who trained on all three levels (i.e., Experiment 1, Group 0).

Group 1 participants ($n = 8$) were first trained on Level 1 and then Level 2, but not Level 0. The results from this group (Figure 4) were similar to the results from Experiment 1. In particular, comparing performance between Group 0 and Group 1 (three-way mixed ANOVA) showed significant effects of session (pre vs. post), $F(1, 25) = 191.3$, $p < 0.001$, and complexity level (1 vs. 2), $F(1, 25) = 25.9$, $p < 0.001$, but no significant effect of group, $F(1, 25) = 0.253$, $p = 0.619$, nor any significant interactions: session, level, and group, $F(1, 25) = 0.311$, $p = 0.582$; session and group, $F(1, 25) = 2.22$, $p = 0.149$; level and group, $F(1, 25) = 1.15$, $p = 0.293$. Further, comparing initial training performance (mean of first two training blocks) between the two groups did not show a significant group effect, $F(1, 25) = 0.106$, $p = 0.747$, suggesting that training with zero-order sequences does not facilitate the learning of higher order sequences.

In contrast, extracting higher order structures proved to be more difficult for Group 2 participants ($n = 12$), who did not have prior experience with zero- or first-order sequences (Figure 4). In particular, eight of 12 participants improved significantly in the task during training, while the rest of the participants showed less than 10% improvement. A mixed ANOVA comparing training session (start, end of training) and group (0, 1,




Figure 5. Strategies for learning context-based statistics. (a) Correlations of individual strategy index and learning rate for participants who improved at both Levels 1 and 2 during training in Group 0 and Group 1. (b) Correlation of individual strategy index between Level 1 and Level 2 for participants trained in Group 0 and Group 1. Negative strategy-index values indicate a strategy closer to matching, while positive values indicate a strategy closer to maximization.

2) showed a significant interaction between the two, $F(2, 31) = 4.41$, $p = 0.021$. In particular, there was a significant difference between groups in performance at the start, $F(2, 31) = 5.14$, $p = 0.012$, but not the end of training, $F(2, 31) = 0.893$, $p = 0.420$. To investigate this difference further, we compared performance on the second-order contexts only (i.e., excluding first-order contexts in Level 2) between groups. There was a significant interaction between session and group, $F(2, 31) = 10.52$, $p < 0.001$, and a significant difference between groups in performance at the start, $F(2, 32) = 5.05$, $p = 0.013$, but not at the end of training, $F(2, 32) = 1.75$, $p = 0.191$. Post hoc comparisons showed significantly higher performance indices in the prediction task for second-order contexts in Group 0 and Group 1 than in Group 2—Group 0 versus Group 2: $p = 0.023$; Group 1 versus Group 2: $p = 0.009$.

Taken together, these results suggest that learning first-order sequences facilitates learning of higher order sequences. In contrast, learning frequency statistics does not facilitate performance in learning higher order sequences. Further, fast learners in Experiment 2 extracted the correct context length and context–target contingencies early in training and deviated from matching toward maximization (Supplementary Figure S5). In particular, fast learners extracted second-order contexts earlier than slower learners, who continued to rely on first-order contexts (Supplementary Figure S6).

Tracking individual strategy across levels

Combining data across experiments, we asked how individual strategy relates to learning performance (i.e., learning rate). Significant correlations (Figure 5a) between participants' learning rate and strategy in-

dex—Level 1 ($n = 27$): $R = 0.461$, $p = 0.016$; Level 2 ($n = 33$): $R = 0.519$, $p = 0.002$ —indicate that participants who extracted the correct context length early in the training adopt a strategy closer to maximization. These results suggest that fast learning relates to selecting the most probable outcome when learning context–target contingencies. We then asked how the participants' strategies developed during training across levels. Correlating individual strategy index across Levels 1 and 2 (Figure 5b) showed that participants' strategy was highly correlated ($R = 0.489$, $p = 0.0131$) across Levels 1 and 2 ($n = 25$ from Groups 0 and 1). These results suggest that participants mostly retained the same strategy across levels of complexity (i.e., from first- to second-order sequences).

Discussion

Here we ask how individuals adapt to changes in the environment's statistics to make predictions about future events. In particular, we sought to characterize the dynamics of learning temporal structures that change in their complexity. We tracked each participant's responses across trials and tested whether and when participants extract the structure that governs sequences of unfamiliar symbols. This enabled us to provide the following four main advances in understanding the dynamics of human statistical learning.

First, we show that participants adapt to the environment's statistics: They extract behaviorally relevant structures from temporal sequences that change in their complexity to make predictions about upcoming events. Further, they benefit from previous exposure to lower order statistics (i.e., first-order

sequences) when learning higher order structures. Previous studies (Fiser & Aslin, 2002a, 2005) have shown that humans are able to extract complex spatiotemporal statistics (e.g., joint vs. conditional probability statistics). These statistics are typically manipulated in separate short-lasting experiments and tested across separate groups of individuals. Here, we test how the same individuals extract structures that change in their complexity, simulating more naturalistic situations that require extracting a range of patterns from simple repetition to probabilistic combinations. Our response-tracking approach allows us to monitor whether and when individuals shift from learning simple to complex structures, as the complexity of the presented sequences changed unbeknownst to them. Our findings demonstrate that individuals extract the behaviorally relevant context length and context–target contingencies that correspond to the structure of the presented sequences.

Second, our response-tracking approach allowed us to extract prototypical patterns of learning dynamics. We demonstrate that fast learners succeeded in identifying the correct statistical structure early in the training. Interestingly, when learning complex structures, fast learners extracted higher order contexts and adopted a learning strategy closer to maximization (i.e., extracted the most probable target per context) earlier in the training. Previous work has tested the role of matching versus maximization strategies in perceptual decision making (Acerbi et al., 2014; Eckstein et al., 2013; Murray et al., 2015) and reward-based learning (Shanks et al., 2002): Observers may distribute their choice responses so as to match the underlying input statistics versus maximize their reward by selecting the most frequently rewarded outcome in each trial. Here, we test these strategies in the context of statistical learning. We show that fast learners tend to use a strategy closer to maximization, suggesting that there may be a benefit to extracting the most probable target per context rather than attempting to learn all statistical dependencies. Further, our findings are consistent with studies suggesting that previous experience shapes the selection of decision strategies (Fulvio, Green, & Schrater, 2014; Rieskamp & Otto, 2006).

Third, we ask whether learning temporal structures occurs in an incidental manner through exposure to regularities or whether it involves explicit knowledge of the underlying sequence structure. Previous studies have suggested that learning of regularities may occur implicitly (i.e., by mere exposure rather than external feedback) in a range of tasks: visuomotor sequence learning (Nissen & Bullemer, 1987), artificial grammar learning (Reber, 1967), probabilistic category learning (Knowlton, Squire, & Gluck, 1994), and contextual cue learning (Chun & Jiang, 1998). Most studies have

focused on implicit measures of sequence learning, such as familiarity judgments or reaction times (for a review, see Schwarb & Schumacher, 2012). In contrast, our paradigm allows us to directly test whether exposure to temporal sequences facilitates observers' ability to explicitly predict the identity of the next stimulus in a sequence. Our experimental design makes it unlikely that the participants memorized specific stimulus positions or the full sequences. Further, participants were exposed to the sequences without trial-by-trial feedback, but were given block feedback about their performance that motivated them to continue with training. A control experiment during which the participants were not given any feedback showed similar results to our main experiment (Supplementary Figure S7), suggesting that it is unlikely that the block feedback facilitated explicit sequence memorization. Yet it is possible that making an explicit prediction about the identity of the test stimulus made the participants aware of the dependencies between the stimuli presented in the sequence. During debriefing, most participants reported some predictive sequence structures (i.e., high-probability symbols or context–target combinations). Thus, it is possible that prolonged exposure to probabilistic structures (i.e., multiple sessions in contrast to single-exposure sessions typically used in statistical-learning studies) in combination with prediction judgments (Dale, Duran, & Morehead, 2012) may evoke some explicit knowledge of temporal structures, in contrast to implicit measures of anticipation typically used in statistical-learning studies.

Finally, previous work has discussed a range of possible representations that are formed during statistical learning. This has mainly focused on deriving generative structure from the stimulus space (for a review, see Dehaene, Meyniel, Wacongne, Wang, & Pallier, 2015) and implicated a range of representations from learning stimulus associations and transitional probabilities to sequence chunks (i.e., statistical contingencies) and abstract rules (Aslin & Newport, 2012; Fiser, Berkes, Orbán, & Lengyel, 2010; Opitz, 2010; Orbán, Fiser, Aslin, & Lengyel, 2008; Reber, 1967). In the context of our task, extracting the sequence context length may relate to rule-based learning, while learning behaviorally relevant contingencies may relate to chunk learning. Further, this range of processes parallels the distinctions between model-free and model-based learning by exploring new strategies versus exploiting previously learned associations in the context of reward-based learning (Dayan & Niv, 2008; Koehlin, 2014). However, distinguishing between these accounts in the context of statistical learning is complicated by task setting and complexity (Franco & Destrebecqz, 2012; Pothos, 2007). Here we take a different perspective: To understand the dynamics of human behavior,

we track human responses during mere exposure to temporal sequences that change in their structure, simulating interactions in naturalistic settings that vary in context and complexity. We show that learning predictive statistics proceeds without explicit trial-by-trial feedback and relates to individual strategy in extracting behaviorally relevant structure from sequences of events.

In sum, our findings provide evidence that successful learning of complex structures relies on extracting behaviorally relevant statistics that are predictive of upcoming events. This learning of predictive structures relates to individual decision strategy: Faster learning of complex structures relates to selecting the most probable outcomes in a given context rather than learning the exact sequence statistics, providing evidence for an alternate route to learning. In future work, it would be interesting to investigate whether these strategies are specific to the sensory input modality or mediate domain-general learning of temporal structure (Nastase, Iacovella, & Hasson, 2014). Recent work has provided evidence for statistical learning within and across different sensory modalities (vision, audition, touch; Conway & Christiansen, 2005; Mitchel & Weiss, 2011), suggesting that statistical learning is implemented by domain-general principles that are subject to modality-specific constraints (Frost, Armstrong, Siegelman, & Christiansen, 2015). For example, in vision statistical learning has been mainly demonstrated by extracting spatial relations, while in audition by extracting temporal regularities. Learning predictive statistics across modalities is critical not only for sensorimotor interactions with the environment but also higher cognitive functions that involve complex structures, such as action organization, music comprehension, and language learning (Conway & Christiansen, 2001; Dehaene et al., 2015; Fitch & Martins, 2014; Frost et al., 2015). Finally, it would be interesting to investigate the developmental time course of learning predictive statistics. Previous work has provided evidence for statistical learning from infancy to older age (for a review, see Krogh, Vlach, & Johnson, 2012) in both vision (e.g., Bulf, Johnson, & Valenza, 2011; Fiser & Aslin, 2001, 2002a, 2002b; Kirkham, Slemmer, & Johnson, 2002; Kirkham, Slemmer, Richardson, & Johnson, 2007) and audition (e.g., Pelucchi, Hay, & Saffran, 2009; Saffran et al., 1999; Saffran, Aslin, & Newport, 1996; Saffran, Newport, & Aslin, 1996). Further, it has been suggested that while learning probabilities is achieved early in life, learning meaningful statistical patterns develops later in adolescence (Amso & Davidow, 2012; Janacek, Fiser, & Nemeth, 2012). This may relate to the suggestion that young children maximize, while matching develops later in life (Kam & Newport, 2009; Stevenson & Weir, 1959; Weir, 1964). Future work on the brain mechanisms of

learning predictive statistics may explore the development of common brain routes to structure learning across domains of perceptual and cognitive expertise.

Keywords: learning, behavior, vision

Acknowledgments

This work was supported by grants to PT from the Engineering and Physical Sciences Research Council (EP/L000296/1); to ZK from the Biotechnology and Biological Sciences Research Council (H012508), the Leverhulme Trust (RF-2011-378), and the European Community's Seventh Framework Programme (FP7/2007-2013) under agreement PITN-GA-2011-290011; and to AW from the Wellcome Trust (095183/Z/10/Z).

*RW and YS contributed equally to this article.

Commercial relationships: none.

Corresponding author: Zoe Kourtzi.

Email: zk240@cam.ac.uk.

Address: Department of Psychology, University of Cambridge, Cambridge, UK.

References

- Acerbi, L., Vijayakumar, S., & Wolpert, D. M. (2014). On the origins of suboptimality in human probabilistic inference. *PLoS Computational Biology*, *10*(6), e1003661, doi:10.1371/journal.pcbi.1003661.
- Amso, D., & Davidow, J. (2012). The development of implicit learning from infancy to adulthood: Item frequencies, relations, and cognitive flexibility. *Developmental Psychobiology*, *54*(6), 664–673, doi: 10.1002/dev.20587.
- Antoniou, M., Ettliger, M., & Wong, P. C. M. (2016). Complexity, training paradigm design, and the contribution of memory subsystems to grammar learning. *PLOS ONE*, *11*(7), e0158812, doi:10.1371/journal.pone.0158812.
- 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.
- Brainard, D. H. (1997). The Psychophysics Toolbox. *Spatial Vision*, *10*, 433–436.
- Bulf, H., Johnson, S. P., & Valenza, E. (2011). Visual statistical learning in the newborn infant. *Cognition*, *121*(1), 127–132, doi:10.1016/j.cognition.2011.06.010.
- 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.
- Chun, M. M., & Jiang, Y. H. (1998). Contextual cueing: Implicit learning and memory of visual context guides spatial attention. *Cognitive Psychology*, 36(1), 28–71, doi:10.1006/cogp.1998.0681.
- Conway, C. M., & Christiansen, M. H. (2001). Sequential learning in non-human primates. *Trends in Cognitive Sciences*, 5(12), 539–546.
- Conway, C. M., & Christiansen, M. H. (2005). Modality-constrained statistical learning of tactile, visual, and auditory sequences. *Journal of Experimental Psychology: Learning, Memory, & Cognition*, 31(1), 24–39, doi:10.1037/0278-7393.31.1.24.
- Dale, R., Duran, N. D., & Morehead, J. R. (2012). Prediction during statistical learning, and implications for the implicit/explicit divide. *Advances in Cognitive Psychology*, 8(2), 196–209.
- Dayan, P., & Niv, Y. (2008). Reinforcement learning: The good, the bad and the ugly. *Current Opinion in Neurobiology*, 18(2), 185–196.
- Dehaene, S., Meyniel, F., Wacongne, C., Wang, L., & Pallier, C. (2015). The neural representation of sequences: From transition probabilities to algebraic patterns and linguistic trees. *Neuron*, 88(1), 2–19, doi:10.1016/j.neuron.2015.09.019.
- Droll, J. A., Abbey, C. K., & Eckstein, M. P. (2009). Learning cue validity through performance feedback. *Journal of Vision*, 9(2):18, 1–23, doi:10.1167/9.2.18. [PubMed] [Article]
- Eckstein, M. P., Abbey, C. K., Pham, B. T., & Shimozaki, S. S. (2004). Perceptual learning through optimization of attentional weighting: Human versus optimal Bayesian learner. *Journal of Vision*, 4(12):3, 1006–1019, doi:10.1167/4.12.3. [PubMed] [Article]
- Eckstein, M. P., Mack, S. C., Liston, D. B., Bogush, L., Menzel, R., & Krauzlis, R. J. (2013). Rethinking human visual attention: Spatial cueing effects and optimality of decisions by honeybees, monkeys and humans. *Vision Research*, 85, 5–19.
- Fiser, J., & Aslin, R. N. (2001). Unsupervised statistical learning of higher-order spatial structures from visual scenes. *Psychological Science*, 12(6), 499–504.
- Fiser, J., & Aslin, R. N. (2002a). Statistical learning of higher-order temporal structure from visual shape sequences. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 28(3), 458–467.
- Fiser, J., & Aslin, R. N. (2002b). Statistical learning of new visual feature combinations by infants. *Proceedings of the National Academy of Sciences, USA*, 99(24), 15822–15826.
- Fiser, J., & Aslin, R. N. (2005). Encoding multielement scenes: Statistical learning of visual feature hierarchies. *Journal of Experimental Psychology: General*, 134(4), 521–537.
- 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.
- Fitch, W., & Martins, M. D. (2014). Hierarchical processing in music, language, and action: Lashley revisited. *Annals of the New York Academy of Sciences*, 1316(1), 87–104.
- Franco, A., & Destrebecqz, A. (2012). Chunking or not chunking? How do we find words in artificial language learning? *Advances in Cognitive Psychology*, 8(2), 144–154.
- Frost, R., Armstrong, B. C., Siegelman, N., & Christiansen, M. H. (2015). Domain generality versus modality specificity: The paradox of statistical learning. *Trends in Cognitive Science*, 19(3), 117–125, doi:10.1016/j.tics.2014.12.010.
- Fulvio, J. M., Green, C. S., & Schrater, P. R. (2014). Task-specific response strategy selection on the basis of recent training experience. *PLoS Computational Biology*, 10(1), e1003425.
- Janacek, K., Fiser, J., & Nemeth, D. (2012). The best time to acquire new skills: Age-related differences in implicit sequence learning across the human lifespan. *Developmental Science*, 15(4), 496–505, doi:10.1111/j.1467-7687.2012.01150.x.
- Jensen, S., Boley, D., Gini, M., & Schrater, P. (2005, Month). *Rapid on-line temporal sequence prediction by an adaptive agent*. Paper presented at the Fourth International Joint Conference on Autonomous Agents and Multiagent Systems, July 25–29, Utrecht, the Netherlands.
- Kam, C. L., & Newport, E. L. (2009). Getting it right by getting it wrong: When learners change languages. *Cognitive Psychology*, 59(1), 30–66, doi:10.1016/j.cogpsych.2009.01.001.
- Karni, A., & Sagi, D. (1993). The time course of learning a visual skill. *Nature*, 365(6443), 250–252.
- Kirkham, N. Z., Slemmer, J. A., & Johnson, S. P. (2002). Visual statistical learning in infancy: Evidence for a domain general learning mechanism. *Cognition*, 83(2), B35–B42.
- Kirkham, N. Z., Slemmer, J. A., Richardson, D. C., & Johnson, S. P. (2007). Location, location, location: Development of spatiotemporal sequence learning in infancy. *Child Development*, 78(5), 1559–1571, doi:10.1111/j.1467-8624.2007.01083.x.
- Knowlton, B. J., Squire, L. R., & Gluck, M. A. (1994).

- Probabilistic classification learning in amnesia. *Learning & Memory*, 1(2), 106–120.
- Koechlin, E. (2014). An evolutionary computational theory of prefrontal executive function in decision-making. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 369 (1655), doi:10.1098/rstb.2013.0474.
- Krogh, L., Vlach, H. A., & Johnson, S. P. (2012). Statistical learning across development: Flexible yet constrained. *Frontiers in Psychology*, 3, 598, doi:10.3389/fpsyg.2012.00598.
- Lagnado, D. A., Newell, B. R., Kahan, S., & Shanks, D. R. (2006). Insight and strategy in multiple-cue learning. *Journal of Experimental Psychology: General*, 135(2), 162–183.
- Liu, J., Lu, Z. L., & Doshier, B. A. (2010). Augmented Hebbian reweighting: Interactions between feedback and training accuracy in perceptual learning. *Journal of Vision*, 10(10):29, 1–14, doi:10.1167/10.10.29. [PubMed] [Article]
- Mitchel, A. D., & Weiss, D. J. (2011). Learning across senses: Cross-modal effects in multisensory statistical learning. *Journal of Experimental Psychology: Learning, Memory, & Cognition*, 37(5), 1081–1091, doi:10.1037/a0023700.
- Murray, R. F., Patel, K., & Yee, A. (2015). Posterior probability matching and human perceptual decision making. *PLoS Computational Biology*, 11(6), e1004342, doi:10.1371/journal.pcbi.1004342.
- Nastase, S., Iacovella, V., & Hasson, U. (2014). Uncertainty in visual and auditory series is coded by modality-general and modality-specific neural systems. *Human Brain Mapping*, 35(4), 1111–1128, doi:10.1002/hbm.22238.
- Nissen, M. J., & Bullemer, P. (1987). Attentional requirements of learning: Evidence from performance measures. *Cognitive Psychology*, 19(1), 1–32, doi:10.1016/0010-0285(87)90002-8.
- Opitz, B. (2010). Neural binding mechanisms in learning and memory. *Neuroscience & Biobehavioral Reviews*, 34(7), 1036–1046.
- Orbán, G., Fiser, J., Aslin, R. N., & Lengyel, M. (2008). Bayesian learning of visual chunks by human observers. *Proceedings of the National Academy of Sciences, USA*, 105(7), 2745–2750.
- Pelli, D. G. (1997). The VideoToolbox software for visual psychophysics: Transforming numbers into movies. *Spatial Vision*, 10(4), 437–442.
- Pelucchi, B., Hay, J. F., & Saffran, J. R. (2009). Learning in reverse: Eight-month-old infants track backward transitional probabilities. *Cognition*, 113(2), 244–247, doi:10.1016/j.cognition.2009.07.011.
- 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., Doshier, B. A., & Lu, Z. L. (2005). The dynamics of perceptual learning: An incremental reweighting model. *Psychological Review*, 112(4), 715–743, doi:10.1037/0033-295x.112.4.715.
- Petrov, A. A., Doshier, B. A., & Lu, Z. L. (2006). Perceptual learning without feedback in non-stationary contexts: Data and model. *Vision Research*, 46(19), 3177–3197, doi:10.1016/j.visres.2006.03.022.
- Pothos, E. M. (2007). Theories of artificial grammar learning. *Psychological Bulletin*, 133(2), 227–244.
- Reber, A. S. (1967). Implicit learning of artificial grammars. *Journal of Verbal Learning and Verbal Behavior*, 6(6), 855–863.
- Rieskamp, J., & Otto, P. E. (2006). SSL: A theory of how people learn to select strategies. *Journal of Experimental Psychology: General*, 135(2), 207–236.
- Saffran, J. R., Aslin, R. N., & Newport, E. L. (1996, Dec 13). Statistical learning by 8-month-old infants. *Science*, 274(5294), 1926–1928.
- 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.
- 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.
- Schwarb, H., & Schumacher, E. H. (2012). Generalized lessons about sequence learning from the study of the serial reaction time task. *Advances in Cognitive Psychology*, 8(2), 165–178.
- Shanks, D. R., Tunney, R. J., & McCarthy, J. D. (2002). A re-examination of probability matching and rational choice. *Journal of Behavioral Decision Making*, 15, 233–250.
- Stevenson, H. W., & Weir, M. W. (1959). Variables affecting children's performance in a probability learning task. *Journal of Experimental Psychology*, 57(6), 403–412.
- 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.
- Turk-Browne, N. B., Scholl, B. J., Chun, M. M., & Johnson, M. K. (2009). Neural evidence of

- statistical learning: Efficient detection of visual regularities without awareness. *Journal of Cognitive Neuroscience*, 21(10), 1934–1945.
- van den Bos, E., & Poletiek, F. H. (2008). Effects of grammar complexity on artificial grammar learning. *Memory & Cognition*, 36(6), 1122–1131, doi:10.3758/mc.36.6.1122.
- Weir, M. W. (1964). Developmental changes in problem-solving strategies. *Psychological Review*, 71, 473–490.