

Human Learning of Non-Markov Structures

Juliana E. Trach (juliana.trach@yale.edu)

Department of Psychology, Yale University,
100 College St. New Haven, CT

Samuel D. McDougle

Department of Psychology, Yale University,
100 College St. New Haven, CT

David P. Carcamo

Department of Physics, Yale University,
217 Prospect St. New Haven, CT

Christopher W. Lynn

Department of Physics, Yale University,
217 Prospect St. New Haven, CT

Abstract

From comprehending language to learning new dance moves, extracting complex relationships between sequences of input is a key feature of human cognition. Prior studies have predominantly explored the cognitive mechanisms of structure learning using Markov sequences, where each element depends only on the previous one. Real-world experience, however, is rife with complex dependencies beyond Markov processes. Here, we study the effects of non-Markov dependencies on sequence learning by leveraging graph learning approaches. We introduce a motor sequence task in which transitional probabilities between pairs of stimuli are identical from a Markov perspective, but differ on higher-order non-Markov dependencies. We find that participants are better able to anticipate stimuli with higher non-Markov probabilities, providing corroboratory evidence that humans are sensitive to statistical structure beyond Markov dependencies. Further, behavior differed from other participants trained only on Markov sequences. Overall, this work demonstrates that humans can rapidly learn and represent statistical dependencies beyond the Markov regime.

Keywords: statistical learning; graph learning; non-Markov processes; motor sequences

Introduction

Humans are constantly parsing continuous experience into discrete units. For example, we might parse a piece of music into phrases or motifs or divide the process of tying your shoe into discrete steps. Our ability to uncover structure in sequential inputs relies on statistical learning—the automatic extraction of statistical regularities in the environment (Aslin & Newport, 2012; Saffran, 2020). This ability is active across the lifespan (Ellis et al., 2021; Saffran & Kirkham, 2018) and is fundamental to core cognitive processes like language acquisition (Saffran et al., 1996; Erickson & Thiessen, 2015), event perception (Levine et al., 2019; Zacks & Swallow, 2007), and motor learning (Nissen & Bullemer, 1987; Hunt & Aslin, 2001).

Foundational work in statistical learning has established that people readily learn item-to-item regularities from sequences of inputs (e.g., Saffran et al., 1996; Nissen & Bullemer, 1987). For example, infants can use transitional probabilities between speech sounds to segment continuous speech into words (Saffran et al., 1996). In another example, Nissen and Bullemer (1987) found that people were faster and more accurate at a cued motor sequence task after introducing statistical structure into the sequence of cued key presses. In other words, participants responded faster when they could

predict the next key press from the current one. Such sequences, where the probability of the next item depends solely on the current item, are known as *Markov* sequences.

This work prompted questions about the limits of our powerful statistical learning abilities. As such, the field of “graph learning” emerged. In this domain, researchers study sequences where the first-order transitional probabilities are uniform (i.e., a uniform first-order Markov sequence), but where larger-scale properties of the network structure affect behavior. In one recent example, Kahn et al. (2018) designed a cued-motor task where sequences of key presses were drawn from an underlying graph structure. Crucially, the graphs were designed such that the transitional probability from each action to another was uniform, however there was larger-scale structure in the graph that made certain transitions more surprising to the participants than others. Despite never being exposed to the whole graph structure simultaneously and the uniform transitional probabilities across actions, participants were slower to respond to more (versus less) surprising transitions, demonstrating that they had indeed learned the structure of the graph. Similar results have been found in the domain of visual statistical learning (Karuza et al., 2017; Schapiro et al., 2015; Rmus et al., 2022), and there has been a great deal of interest in formalizing the computations and mechanisms by which people extract latent structure from continuous experience (Tang et al., 2023; Lynn et al., 2020). Notably, these studies of graph learning primarily leverage first-order Markov sequences.

Pure Markov sequences are rare, if not completely absent from real-world settings. For example, consider playing a phrase of music on the piano—while the next note does depend on the current note, it also depends on the entire sequence of notes that preceded. Thus, examining how people learn and represent statistical dependencies that exceed Markov processes is crucial to a thorough understanding of this fundamental capacity. While Markov sequences represent the status quo in recent graph learning research, there are a number of earlier motor sequence learning studies that tested the effects of higher-order dependencies on motor sequence performance (Remillard & Clark, 2001; Remillard, 2010; Gureckis & Love, 2010). Perhaps unsurprisingly, this work generally indicates that longer-range dependencies do indeed influence motor sequence execution. Still, extant work largely relies on a

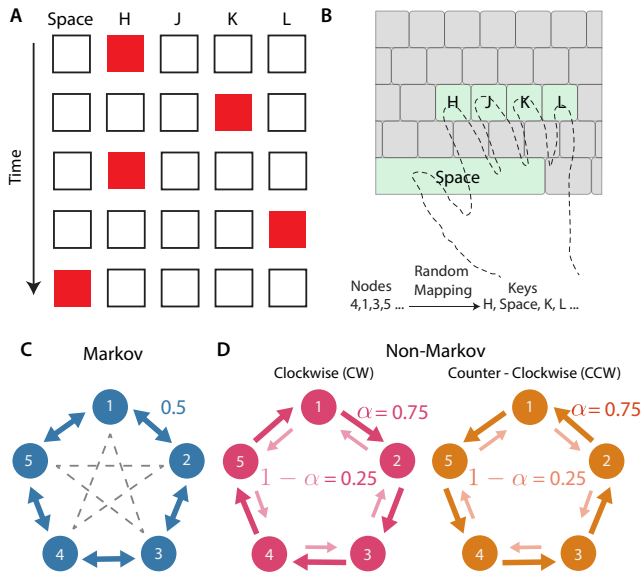


Figure 1: A) Illustration of task stimuli. Participants were shown a sequence of 1000 stimuli, where each stimulus consisted of four white squares and one red square. Each row represents the stimulus for one trial. B) Participants responded to each stimulus by pressing the key corresponding to the red square. Their right hand was positioned over the space, H, J, K, and L keys. C) Illustration of ring network used to generate motor sequences. Nodes correspond to a specific key press (and thus also a specific stimulus). Connections between nodes represent possible transitions between actions. Markov sequences were generated using a ring network with uniform transition probabilities (.5) between neighbors. The gray dashed edges indicate novel transitions that only occurred during the Probe Phase at the end of the experimental session. D) Illustration of transitional structure for non-Markov sequences. Non-Markov sequences were generated using a second-order Markov process, wherein the next stimulus depends not only on the current stimulus, but also on the previous stimulus. When $\alpha = 0.75$ the sequence prefers to continue moving around the network in the same direction.

limited number carefully constructed sequences comprised of transitional rules that might limit the generalizability of these results. Further, controlling for potential motor confounds in these studies presents a significant challenge. In the following, we leverage ideas from the field of graph learning to study the human learning of non-Markov dependencies and expand on previous work examining higher-order statistical associations in motor sequence execution.

To that end, we developed a paradigm to examine how humans learn non-Markov statistics while explicitly controlling for differences in Markov statistics. Importantly, our approach allows for a concise mathematical form that can be used to generate a diversity of specific sequences that

arise from the same underlying structure. The general logic of the experiment is that variation in a participant's reaction time (RT) and accuracy across transitions is related to their expectations: If they have a strong expectation of what will happen next, they should be able to respond quickly and accurately; if, instead, the next stimulus is surprising (i.e., violates their expectations), then they should respond more slowly and be more likely to respond incorrectly (Hyman, 1953). We designed sequences such that the first-order structure was a uniform Markov network (Figure 1C), but the second-order structure of the sequences made certain transitions more likely than others (.75 probability versus .25 probability; Figure 1D). Thus, if participants learned only the first-order structure, we would not expect significant differences in performance across transitions (reflecting the uniform Markov network). However, if they did learn the higher-order structure, the higher-probability transitions in this case should be performed faster than the more surprising transitions. As a control, we compared performance between the non-Markov learners to a separate group of participants exposed only to the first-order Markov network structure (i.e., without higher-order structure; Figure 1C).

To preview our results, we found that participant behavior was influenced by the latent graph structure in both groups of participants. While we found strong evidence that participants had learned the longer-range statistical dependencies embedded in the sequences, we did not find gross differences in learning trajectories across groups, demonstrating that participants did not necessarily find these more complex structures more difficult to learn. This pattern of results implies that learning non-Markovian structures is an efficient, and perhaps automatic, cognitive process.

Methods

Participants

We used the online platform Prolific to recruit 100 participants for each type of structure (non-Markov group: $N = 47$ female, mean age = 27.8, range = 18-35; Markov group: $N = 46$ female, mean age = 28.8, range = 18-35). All participants were from the US or UK and had normal or corrected-to-normal vision. Participants were paid \$2.50 to complete the 15 minute study. We planned on excluding participants who were very inaccurate (<50% accuracy) or excessively slow (average RT > 2s) in their responses. No participants met this exclusion threshold.

Task Design

The basic task was a cued-response motor task where participants were shown a stimulus on each trial that indicated which of five keys to press on a keyboard (Figure 1A&B). During the task, participants positioned their right hand on the space (thumb), H (index), J (middle), K (ring), and L (pinky) keys. On each trial, participants were shown five squares that were spatially aligned with the fingers of their right hand. One of the squares was colored red on

each trial while the others remained white. Participants were instructed to press the key that corresponded to the red square as quickly and accurately as possible. Once the participant responded correctly, a new square would turn red, thus beginning the subsequent trial. The next trial would not start until the participant responded correctly, and participants were instructed to correct their response if they made an error. If the participant responded incorrectly four times in a row, the task paused for 8 s and displayed a message to the participant reminding them to respond accurately.

The experimental session began with instructions and a short practice phase (15 trials). During this phase, the key name associated with the target square (e.g., ‘space’ or ‘J’) was displayed above the red square on each trial in order to help orient participants to the task rules. After the practice phase, participants completed approximately 1,000 trials of the cued-response task (200 iterations per key, 10 min). Trials with especially long (> 1500 ms) or short (< 100 ms) reaction times (RTs) were excluded from analysis. Additionally, for RT analyses, only trials where the participant responded correctly on the first attempt were included. We use median RTs for analysis to attenuate the influence of outlier RTs.

Trial Sequences

Structured Sequences. Our objective was to create sequences that had a uniform first-order Markov structure, but also discoverable second-order structure that could influence participant behavior. To do this, we started with a five-node ring structure with uniform first-order Markov structure (Figure 1C). In this structure, each node represents a stimulus and corresponding key press, and the connections between nodes represent possible transitions between stimuli. We then imposed second-order structure onto that network by conditioning the probability of the next node, x_{t+1} , on the current, x_t , and previous, x_{t-1} , nodes (Figure 1D), where x is the node identity and t is the trial.

Participants in the *Markov* group were exposed to sequences drawn from the first-order Markov network (Figure 1C). In this case, there was a .5 probability of transitioning from the current node (x_t) to either neighboring node on the next trial (x_{t+1} ; i.e., $P(x_{t+1}|x_t) = .5$; Figure 1C). For example, if the current stimulus corresponded to node 2, then there was an equal probability of transitioning to nodes 1 or 3 on the next trial, but transitions to nodes 4 and 5 were not allowed.

For the *non-Markov* sequences, we conditioned the probability of the next node (x_{t+1}), on the current (x_t) and previous node (x_{t-1}) that the participant had seen (i.e., embedded second-order structure into the sequences). Specifically, we biased the probabilities such that x_{t+1} was more likely to continue in the same direction around the ring, than it was to reverse directions. Mathematically, this process is described by a second-order Markov model:

$$P(x_{t+1}|x_t, x_{t-1}) = \begin{cases} \alpha & x_{t+1} - x_t \equiv x_t - x_{t-1} \pmod{5} \\ 1 - \alpha & \text{else.} \end{cases} \quad (1)$$

Here, α is a parameter between 0 and 1, where the next item in the sequence will continue in the same direction around the structure with probability α and reverse direction with probability $1 - \alpha$. In our design, we set $\alpha = .75$ (Figure 1D). Thus, if participants visited node 1 and then node 2 in sequence, they then had .75 probability of transitioning to node 3 and a .25 probability of transitioning back to node 1. Importantly, the first-order statistics were identical to those of the Markov sequences – across all trials, the probability of transitioning from node 2 to node 1 or 3 was equal ($\sum_{x_{t-1}} P(x_{t+1}|x_t, x_{t-1})P(x_{t-1}|x_t) = P(x_{t+1}|x_t) = .5$). Therefore, any differences in the reaction times between these two example sequences (e.g., 1-2-3 versus 1-2-1) must arise from learning the second-order structure. Importantly, we shuffled the assignment of keys to nodes for each participant such that adjacent fingers were not necessarily assigned to adjacent nodes in the ring structure. Participants were never informed of the structure of the trial sequences.

Probe Phase. We added a brief “probe” phase (30 trials) after the 1,000 trials of training on the structured sequences. In this phase, participants responded to novel transitions that violated the sequence structure of the first 1,000 trials. For example, node 2 could transition to node 1 or node 3 during the first 1,000 trials, so during the probe phase, node 2 transitioned only to node 4 or node 5 (Figure 1C, grey lines). If participants learned the sequence structure in the structured sequence phase, they should be slower to respond during the probe trials, since those transitions violated their expectations. There was no break or any task feature (other than the new transitions) that informed participants about the beginning of the probe phase. We opted to keep this phase extremely short to avoid participants learning a new structure.

Results

Learning of Non-Markov Statistics from Sequential Input

Our primary question was whether participants trained on the non-Markov transitions would learn the second-order transitional structure that we imposed on the motor sequences. In general, we expected participants to be more accurate and faster in their responses when they have stronger expectations about what will happen next. That is, if the participant had a strong expectation that the next key press would be with their thumb, they should be able to efficiently plan and execute that response. In contrast, if the next stimulus violated their expectations, they should be slower to make the corresponding response and they may make an error.

We thus compared performance between trials where the sequence continued in the same direction around the ring (“continue trials”) to trials where the sequence reversed direction (“reverse trials”). In the non-Markov sequences, the next node is more likely to continue in the same direction around the ring (.75 probability), than it is to reverse directions (.25 probability). Thus, if participants learned

the non-Markov structure of the sequences, they should be more accurate and faster on continue versus reverse trials. In contrast, if participants attended only to the first-order Markov structure of the sequences (uniform .5 transitional probability to neighboring nodes), then we would not predict differences in accuracy or RT across these trials types. We excluded the first 250 trials from this analysis to focus on trials where participants had enough exposure to learn the latent structure.

We found clear evidence that participants trained on the non-Markov sequences did indeed learn the second-order structure from this analysis, consistent with previous results (Remillard & Clark, 2001). Participants were less accurate and slower when the next trial reversed direction (i.e., the less likely transition) versus when it continued in the same direction (Figure 2A; paired sample t test: accuracy: $t(99) = 4.59, p < .001$; RT: $t(99) = -4.87, p < .001$). This result demonstrates that individuals extracted the non-Markov statistical structure within the motor sequences, expanding on previous work on graph learning.

Comparison to Participants Trained on Markov Sequences

We compared the behavior of the non-Markov group to the behavior of a separate group of participants trained on sequences drawn from the first-order Markov network as a control. We performed a mixed-factor ANOVA (Group [non-Markov versus Markov] x Trial type [continue versus reverse]) to examine differences in behavior between the two groups (excluding the first 250 trials). We found that while there were no overall differences in behavior between the two groups (Figure 2A&B; main effect of Group: accuracy: $F(1,198) = 0.82, p = .365$; RT: $F(1,198) = 0.06, p = .802$), there were indeed significant differences in RT between trial types (main effect of Trial type: RT: $F(1,198) = 15.03, p < .001, \eta_p^2 = 0.07$; although not in accuracy: $F(1,198) = 0.31, p = .577$). Crucially, we found that the effect of continuing versus reversing direction was different between groups (interaction Group x Trial type: accuracy: $F(1,198) = 60.59, p < .001, \eta_p^2 = 0.23$; RT: $F(1,198) = 125.7, p < .001, \eta_p^2 = 0.39$). Thus, we can conclude that participants trained on Markov versus non-Markov sequences extracted different underlying structures based on the sequences they saw.

Intriguingly, while the non-Markov group performed worse on reverse (i.e., the less likely transition) versus continue trials, participants trained on the Markov networks showed the *opposite* pattern of behavior: They were more accurate and faster when the next trial reversed direction around the ring versus continued in the same direction (paired sample t test: accuracy: $t(99) = -6.77, p < .001$; RT: $t(99) = 11.46, p < .001$). While we expected equivalent performance across trial types in this group (because the transitional probabilities are uniform), there is one rather straightforward explanation for this effect that is wholly unrelated to the learning of the statistical structure: Due to the ring structure

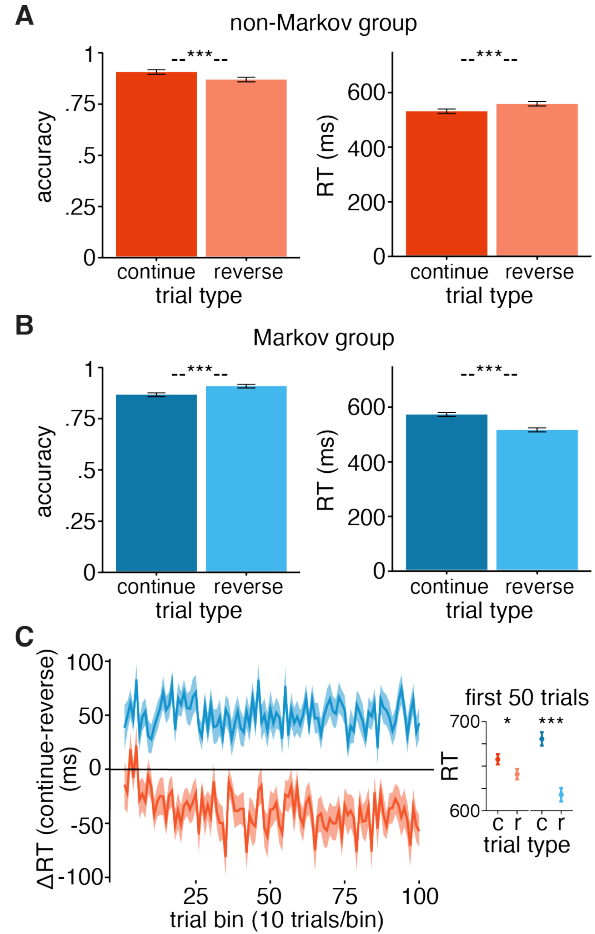


Figure 2: A) non-Markov group: Comparison of accuracy (left) and median RT (right) on "continue" trials (.75 probability) versus "reverse" trials (.25 probability). Error bars = 1 SEM. B) Markov group: Comparison of accuracy (left) and median RT (right) on "continue" trials (.5 probability) versus "reverse" trials (.5 probability). Error bars = 1 SEM. C) Difference in RT between continue and reverse trials (continue trial median RT - reverse trial median RT) as a function of trial bin. Positive values indicate faster performance on reverse trials, whereas negative values indicate faster performance on continue trials. Orange = non-Markov group, Blue = Markov group. Inset depicts RT on continue (c) versus reverse (r) trials during the first 50 trials. * $p < .05$, *** $p < .001$

of the transitions, reversing direction within this network necessarily means repeating the same action from two trials back (e.g., $trial_{t-2} = \text{thumb}$, $trial_{t-1} = \text{pinky}$, $trial_t = \text{thumb}$). Thus, participant responses might be facilitated by returning to a recently prepared motor response. We note that this same effect would influence the reversal trials in the non-Markov group as well, but, critically, in a direction that goes *against* the predicted effect of the non-Markov regularities.

Thus, we reasoned that if the decrease in RTs for reverse

trials in the Markov group was due to a recency bias (rather than learning), then we should see the same facilitation on reverse trials in the non-Markov group early in the session before they had learned the statistical structure. To test this, we compared accuracy and RT on continue versus reverse trials in the non-Markov group during the first 50 trials of the task. We chose the first 50 trials of the task (rather than a smaller or larger number) to limit the amount of exposure that participants had to the structure while also ensuring enough reverse trials for analysis, since reversals were infrequent occurrences (.25 probability). Consistent with our hypothesis, we found that participants were indeed faster on reverse trials early in the task in the non-Markov condition (Figure 2C, inset; $t(99) = 2.24$, $p = .028$, $d = 0.1$), although this effect was not reliable in accuracy ($t(99) = 1.04$, $p = .302$). This analysis suggests that RTs on reverse trials were indeed facilitated prior to any learning, consistent with a simple motor facilitation effect. Exposure to the latent statistical structure can then overcome and reverse this prior.

Further corroborating this account, we found that the facilitation of continue (relative to reverse) trials in the non-Markov group emerged gradually over the course of the session. To quantify this, we divided the task into ten-trial bins (100 bins total) and calculated a difference score for RTs for continue versus reverse trials (continue trial RTs - reverse trial RTs). We then entered these difference scores into a linear mixed effects model with Bin Number and Group as fixed effects and random slopes and intercepts for each participant. In this case, we found no significant effect of Bin number (Figure 2C; $\beta = -0.06$, $p = .219$), however there was a significant effect of Group ($\beta = -0.63$, $p < .001$), such that RT difference scores were more negative in the non-Markov group, reflecting better performance on continue versus reverse trials. Crucially, there was a significant interaction between Group and Bin number ($\beta = -0.19$, $p = .0057$). We conducted additional linear mixed effects analyses on each experiment separately to further investigate this interaction.

As predicted, we found a significant effect of Bin number in the non-Markov group ($\beta = -0.25$, $p < .001$) reflecting the emergence of this effect with practice. This was not the case in the Markov group ($\beta = -0.06$, $p = .194$). In this case, the facilitation of reverse trial persisted at the same level throughout the task, consistent again with an inherent (rather than learned) bias. Taken together, these results establish that participants in the non-Markov group learned the second-order statistics over time, even overcoming the competing effect of recency.

Comparable Learning of Non-Markov and Markov Structures

Finally, our design allowed us to examine differences in learning outcomes and trajectories across groups. There were no differences in overall accuracy (Figure 3A; non-Markov group: $M = 89.8\%$; Markov group: $M = 89.5\%$; $t(198) = 0.28$, $p = .776$) nor median reaction time (RT; Figure 3B;

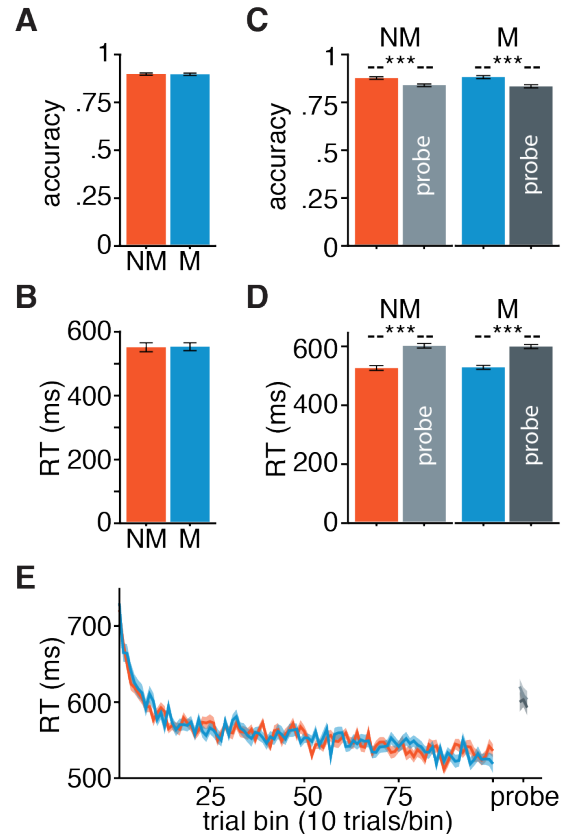


Figure 3: A) Average accuracy in the non-Markov (NM) versus Markov (M) groups. B) Median RT in the non-Markov versus Markov groups. C) Accuracy during the last 30 trials of the structured sequence phase versus probe phase for non-Markov (left) and Markov (right) groups. D) Median RT during the last 30 trials of the structured sequence phase versus probe phase for non-Markov (left) and Markov (right) groups. E) RT learning curve over trial bins and probe phase. *** $p < .001$

non-Markov group: average median RT = 551ms; Markov group: average median RT = 553ms; $t(194.08) = 0.07$, $p = .945$) between the non-Markov and Markov versions of the task. We also compared accuracy and median RT at the end of structured sequences (last 30 trials) to accuracy and median RT during the probe phase (30 trials at the end of the task where the learned graph structure was violated) across the two groups. We used a mixed-factor ANOVA to do this with a between-subjects factor for Group (Markov versus non-Markov) and a within-subject factor for Phase (structured versus probe). If participants learned better in one group or the other, we would expect an interaction between Phase and Group, reflecting differential performance reductions during the probe phase. As expected, participants were slower and less accurate during the probe phase as compared to the end of the structured sequences (Figure 3C&D; Accuracy: main effect of Phase: $F(1,198) = 33.39$,

$p < .001$, $\eta_p^2 = 0.14$; non-Markov group, paired t-test: $t(99) = 3.91$, $p < .001$; Markov group, paired t-test: $t(99) = 4.26$, $p < .001$; Median RT: main effect of Phase: $F(1,198) = 100.49$, $p < .001$, $\eta_p^2 = 0.34$; non-Markov group, paired t-test: $t(99) = 6.83$, $p < .001$; Markov group, paired t-test: $t(99) = 7.43$, $p < .001$). This effect did not differ across groups (Accuracy: main effect of Group: $F(1,198) = 0.08$, $p = .782$; Group x Phase: $F(1,198) = 0.53$, $p = .467$; Median RT: main effect of Group: $F(1,198) = 0.11$, $p = .736$; Group x Phase: $F(1,198) = 0.008$, $p = .929$), further supporting the idea that learning was comparable across graph structures.

Additionally, we found that the learning trajectories (measured with median RT) were similar across the two groups (Figure 3E). To test this statistically, we divided the task into bins of 10 trials (100 bins), calculated the average median RT in each of these bins, and performed a one-sample t-test to assess whether the difference between the two RT learning curves (i.e., Markov versus non-Markov) was reliably different than 0. We also repeated this analysis for only the first quarter of trials (25 bins), where one might expect differences in learning to be most prominent. Again, we found no evidence of learning differences across the two versions of the task (All data: $t(99) = 1.46$, $p = 0.147$; first 250 trials: $t(24) = 1.4$, $p = 0.175$). Taken together, these results show that participants performed at a similar level overall regardless of whether they were trained on non-Markov or Markov sequences and despite the increased complexity of the non-Markov sequences.

Discussion

Here, we developed a motor sequence task to examine statistical learning of non-Markov dependencies in sequential experience. We drew on ideas from graph learning to generate sequences with non-Markov structure, but with uniform Markov statistics (Remillard & Clark, 2001), thus eliminating confounding Markov dependencies. We found robust behavioral evidence that participants learned the non-Markov structure. Moreover, comparing against participants trained on Markov sequences with the same first-order structure, non-Markov learners exhibited distinct patterns of behavior. Taken together, these results show that people rapidly, and perhaps automatically, extract latent statistical structures within sequences beyond simple Markov dependencies. Further, our generalizable approach allows for rigorous investigation of the limits of these abilities in future work.

We note that this work sits at the intersection of two larger literatures. On one side, previous motor sequence learning work provided initial evidence that long-range dependencies in sequential input impact motor sequence execution (Remillard & Clark, 2001; Remillard, 2010; Gureckis & Love, 2010). On the other, research on graph learning has suggested that people build structured cognitive representations from sequential input that impact prediction and behavior (Lynn et al., 2020; Kahn et al., 2018; Karuza

et al., 2017). Thus, one intriguing question is about the structure of the representations that people construct during learning. While our approach implies that people are extracting the latent structure, there is some evidence that similar behavioral effects could arise from simpler association learning (Gureckis & Love, 2010), instead of constructing the whole latent structure. Our approach makes this question more tractable to test, as the underlying graph structure is likely easier learned than the transitional rules employed in previous motor sequence learning work.

One surprising feature of our results was that participants in the Markov group performed significantly better on reverse trials relative to continue trials. This is in stark contrast to the non-Markov group, which exhibited facilitated performance on continue trials due to the second-order statistics. We hypothesized that this effect was driven by the fact that reverse trials necessarily involved repeating a recent response, thus giving rise to a recency bias. Indeed, we found statistical evidence to support this account, which further strengthens our central findings. Specifically, the effects of the non-Markov structure were so strong that they overcame the natural bias to respond more quickly (and accurately) to reverse trials.

A potentially exciting result from this work is that there were no overall differences in performance or learning trajectory between the two groups of participants. From one perspective, one might expect better overall performance in the non-Markov condition, since participants have more information to predict upcoming stimuli. Yet from another perspective, one might expect slower learning in the non-Markov condition (even if performance is facilitated overall) since these dependencies are more complex to learn and internally represent. However, we did not observe either of these patterns of behavior. Our results could be interpreted as further evidence of the power of statistical learning—participants readily acquired the second-order dependencies at the same rate as the Markov statistics. At this point, we hesitate to take a strong interpretation of these results, as the design of the networks meant that the second-order dependencies were working against inherent motor biases that we observed in the Markov group. A productive next step would be to vary α in the non-Markov sequences and train participants on sequences that align with the motor bias to better quantify the effect of statistical learning versus motor processing.

Statistical learning from sequential input is a fundamental component of human cognition and underpins a wide range of cognitive capacities. This work presents a method for examining how people extract and represent higher-order statistical regularities in the environment. In ongoing work, we are using computational modeling to characterize learning trajectories and internal representations of non-Markov structure learning. Examining our ability to uncover complex statistical relationships from continuous experience is central to understanding many facets of behavior.

Acknowledgments

We would like to thank members of the ACT Lab for helpful discussions about this project. J.E.T. was funded by NSF GRFP.

References

- Aslin, R. N., & Newport, E. L. (2012, June). Statistical Learning: From Acquiring Specific Items to Forming General Rules. *Current Directions in Psychological Science*, 21(3), 170–176. doi: 10.1177/0963721412436806
- Ellis, C. T., Skalaban, L. J., Yates, T. S., Bejjanki, V. R., Córdova, N. I., & Turk-Browne, N. B. (2021, August). Evidence of hippocampal learning in human infants. *Current Biology*, 31(15), 3358–3364.e4. doi: 10.1016/j.cub.2021.04.072
- Erickson, L. C., & Thiessen, E. D. (2015). Statistical learning of language: Theory, validity, and predictions of a statistical learning account of language acquisition. *Developmental Review*. doi: 10.1016/j.dr.2015.05.002
- Gureckis, T. M., & Love, B. C. (2010). Direct associations or internal transformations? exploring the mechanisms underlying sequential learning behavior. , 34(1), 10–50. Retrieved 2025-05-08, from doi: 10.1111/j.1551-6709.2009.01076.x
- Hunt, R. H., & Aslin, R. N. (2001, December). Statistical learning in a serial reaction time task: access to separable statistical cues by individual learners. *Journal of experimental psychology. General*, 130(4), 658–80.
- Hyman, R. (1953). Stimulus information as a determinant of reaction time. *Journal of Experimental Psychology*, 45(3), 188–196. Retrieved 2025-01-24, from doi: 10.1037/h0056940
- Kahn, A. E., Karuza, E. A., Vettel, J. M., & Bassett, D. S. (2018). Network constraints on learnability of probabilistic motor sequences. *Nature Human Behaviour*, 2, 936–947. doi: 10.1038/s41562-018-0463-8
- Karuza, E. A., Kahn, A. E., Thompson-Schill, S. L., & Bassett, D. S. (2017, December). Process reveals structure: How a network is traversed mediates expectations about its architecture. *Scientific Reports*, 7(1), 1–9. (Publisher: Nature Publishing Group) doi: 10.1038/s41598-017-12876-5
- Levine, D., Buchsbaum, D., Hirsh-Pasek, K., & Golinkoff, R. M. (2019). Finding events in a continuous world: A developmental account. *Developmental Psychobiology*, 61(3), 376–389. doi: 10.1002/dev.21804
- Lynn, C. W., Kahn, A. E., Nyema, N., & Bassett, D. S. (2020, May). Abstract representations of events arise from mental errors in learning and memory. *Nature Communications*, 11(1), 2313. doi: 10.1038/s41467-020-15146-7
- Nissen, M. J., & Bullemer, P. (1987, January). Attentional requirements of learning: Evidence from performance measures. *Cognitive Psychology*, 19(1), 1–32. Retrieved 2019-02-10, from doi: 10.1016/0010-0285(87)90002-8
- Remillard, G. (2010). Implicit learning of fifth- and sixth-order sequential probabilities. , 38(7), 905–915. Retrieved 2025-05-08, from doi: 10.3758/MC.38.7.905
- Remillard, G., & Clark, J. M. (2001). Implicit learning of first-, second-, and third-order transition probabilities. , 27(2), 483–498. Retrieved 2025-05-08, from doi: 10.1037/0278-7393.27.2.483
- Rmus, M., Ritz, H., Hunter, L. E., Bornstein, A. M., & Shenhav, A. (2022, August). Humans can navigate complex graph structures acquired during latent learning. *Cognition*, 225, 105103. doi: 10.1016/j.cognition.2022.105103
- Saffran, J. R. (2020). Statistical Language Learning in Infancy. *Child Development Perspectives*, 14(1), 49–54. doi: 10.1111/cdep.12355
- Saffran, J. R., Aslin, R. N., & Newport, E. L. (1996). Statistical learning by 8-month-old infants. *Science*, 274, 1926–1928.
- Saffran, J. R., & Kirkham, N. Z. (2018). Infant Statistical Learning. *Annual Review of Psychology*, 69(1), 181–203. doi: 10.1146/annurev-psych-122216-011805
- Schapiro, A. C., Turk-Browne, N. B., Norman, K. A., & Botvinick, M. M. (2015). Statistical Learning of Temporal Community Structure in the Hippocampus. *Hippocampus*. doi: 10.1002/hipo.22523
- Tang, W., Shin, J. D., & Jadhav, S. P. (2023, March). Geometric transformation of cognitive maps for generalization across hippocampal-prefrontal circuits. *Cell Reports*, 42(3), 112246. doi: 10.1016/j.celrep.2023.112246
- Zacks, J. M., & Swallow, K. M. (2007, April). Event Segmentation. *Current Directions in Psychological Science*, 16(2), 80–84. (Publisher: SAGE Publications Inc) doi: 10.1111/j.1467-8721.2007.00480.x