

Sequence learning in noise: Modelling the behavior of non-human primates

Jeremy D. Yeaton (jyeaton@uci.edu)

Department of Language Science, 2243 Social Sciences Plaza
Irvine, CA 92697 USA

Abstract

When human and non-human animals learn sequences, they develop a serial expertise about the order in which elements are organized in the sequence. This expertise can originate both from associating stimuli to each other and associating each stimulus to its ordinal position. However, when a simple regularity of two elements is inserted in a random sequence at different positions, ordinal position and adjacent associations do not have the same reliability, and it is unclear how associative mechanisms will constraint ordinal knowledge about the sequence. Behavioral results from baboons indicate that ordinal position information does not play a role in the acquisition of these regularities. Here, we have implemented a simple, cognitively-inspired computational model to gain insight into the mechanisms underlying statistical learning of regularities in noise of non-human primates. We find crucial similarities and differences between the model behavior and the observed primate behavior.

Keywords: statistical learning; associative learning; animal cognition

Introduction

As a task is performed, predictions are constantly being made about what will come next so as to respond to the stimuli as quickly and efficiently as possible. In order to do this, patterns must be extracted from the environment, represented, and subsequently implemented as motor behaviors. The processes by which these steps take place, however, are not well understood nor are variables that impact the extractability of a pattern—or regularity—from the input (Perruchet & Vinter, 1998). Here we implement a cognitively plausible model of regularity extraction and compare it to behavior of non-human primates based on data from a Hebbian learning task (Hebb, 1961).

Background

Statistical Learning

Statistical learning (SL) refers to the ability to extract regularities from the environment exclusively through exposure without the use of explicit scaffolding or instruction. In many SL experiments, participants must extract regularities in a stream of stimuli which is free from random elements. This fact represents a non-ecological design, as a natural environment is replete with random, unpredictable, noisy input, which

must be perceived and processed in order to extract patterns. More recent studies have examined the extraction of regularities in noise, and manipulate the number of elements comprising the regularity, how many random elements appear between repetition of the regularity, and the position in sequence of the regularity. Such studies test for serial position effects, i.e.: primacy and recency effects wherein the first and last elements of a sequence are more readily recalled than elements in the middle (Ebbinghaus, 1885). In essence, a relationship is learned between the element and its position in the sequence. In contrast to this would be the learning of associations between internal elements in the sequence (e.g.: element j always follows element i ; Young, 1962). A recent study pitted these competing hypotheses against one another by manipulating the position in sequence of simple regularities in short visuo-motor sequences with noise (Yeaton, Tosatto, Fagot, Grainger, & Rey, 2021).

Yeaton, et al. (2021)

In two experiments with a group of Guinea baboons, Yeaton and colleagues manipulated 1) the position of a two-element regularity (AB) within a noisy sequence, 2) whether it had a fixed position, and 3) the length of the sequence it was embedded in, to determine whether these factors impact its learnability. Their results showed no difference in learning rate regardless of length or position manipulations.

Yeaton et al. used a sequence pointing task in which baboons must point to a series of circles that appear at different fixed points on a screen (Minier, Fagot, & Rey, 2016). In a trial, a red circle would appear over a point on the grid (Fig. 1). When touched, the circle disappears and reappears immediately at a different point. They operationalized learning rate as the decrease in response time (RT) for a given transition over time, i.e.: if the RT for the transition from point A to point B decreased relative to a baseline condition, then it was inferred that the baboons learned the relationship between those points. In the task, accelerated learning in experimental conditions would be attributed to the use of top-down attentional or cognitive resources related to pattern identification and extraction. By contrast, a constant learning rate would mean a strictly bottom-up Hebbian process driven only by the frequency of co-occurrence (Hebb,

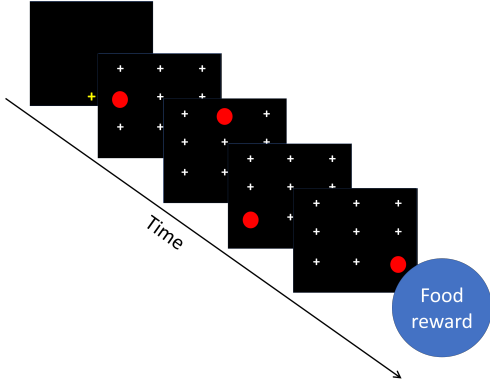


Figure 1: Schematic of a trial of length 4. The baboon first touches the yellow cross, then each of the series of red circles, with response time measured for each. After successfully touching all 4 circles, a reward is automatically distributed (grain of dry wheat).

1961).

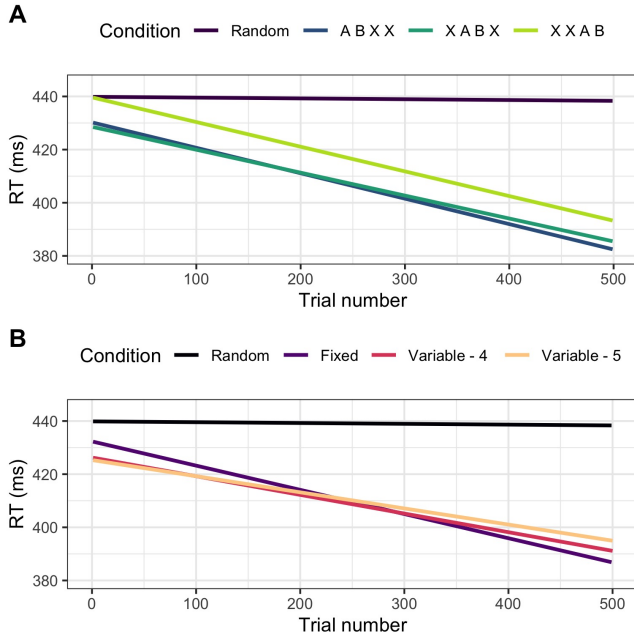


Figure 2: Results from Yeaton et al. (2021). Regression lines fit to RTs aggregated across participants in Experiment 1 (A) and Experiment 2 (B). The slope of the line fit to the RTs is the operational definition of learning rate—if the slope is near 0 (as in Random condition) then no learning took place, whereas if the slope is more negative then learning can be inferred as improvement was observed.

In the experiments, an *AB* regularity appeared either in a fixed position at the beginning, middle, or end of a 4-element sequence (Exp. 1), or in a variable position in

a 4- or 5-element sequence (Exp. 2). The fixed-position conditions are denoted as *ABXX*, *XABX*, and *XXAB* where *X* is some other randomly drawn element that is not *A*, *B*, or the other *X*. Experiment 2 was basically a mix of these three conditions (Variable - 4) or a mix of these with an additional *X* added (Variable - 5). The different conditions allow for varying amounts of context to inform predictions, however no significant differences between these conditions were found (Fig. 2), meaning that the system extracting the regularities is only weakly sensitive to context, if at all. In other words, the baboons extract the inter-element relationship and this extraction is neither helped nor hindered by serial position information.

The Present Study

In this work, we set out to test a potential system architecture underlying regularity extraction in non-human primates. By implementing a simple two-layer neural network¹ crudely corresponding to perceptual and motor systems we can examine whether the learning dynamics of the model correspond to those observed in the baboons.

Architecture & Dynamics

Our proposed model is comprised of two main layers: Perceptual/Attentional and Motor. Each layer has a matrix of activation values (*a* and *m* respectively), and a matrix of weights (*A* and *M* respectively). In each trial *t*, some visual stimulus V_t with values $v \in \{0, 1\}$ is presented. At each trial *t*, the attentional activation a_t is updated according to the visual input V_t , the attentional weights A_t , and the previous activation a_{t-1} :

$$a_t = A_t \cdot V_t + \lambda(a_{t-1}) \quad (1)$$

where λ is a scalar modulating the influence of the previous activation state. The activations a_t are then passed to the motor layer as the input. The motor activations m_t are then updated according to that input, the motor weights M_t , and the previous motor activation state m_{t-1} :

$$m_t = M_t \cdot a_t + \lambda(m_{t-1}) \quad (2)$$

The loss at each trial is calculated as the cross-entropy loss between the visual input for that trial V_t and the output of the softmax transformation of the motor activations m_t :

$$J_t = H(V_t, \text{softmax}(m_t)) \quad (3)$$

The weights for both layers *A* and *M* are then updated using gradient descent after each pair of touches.

¹I think this might be an Elman SRN but I'm not sure.

Training

The input for the model is the same stimuli used in the behavioral studies described above. Instead of being divided by trial, however, the trials were appended one after another and were presented as a continuous string of touches. Here we define a **trial** as a sequence of four **touches**. Each “touch” was represented as an 9-dimensional vector: one for each of the possible targets². The touches were fed to the model two at a time³.

Our primary manipulation was of the weight of the activation at time $t-1$ on learning of the AB regularity, as well as overall accuracy. We used values of λ from 0 to 1.0 in steps of 0.1. The Adam algorithm in PyTorch was used with an initial learning rate of 0.001.

Overall, 880 different models were fit (11 λ -values \times 4 conditions \times 20 baboons). Each model was trained on the first 1,000 touches (250 4-touch trials) and tested on the next 200 touches (50 4-touch trials). Accuracy for B in non-ceiling models did not significantly improve with additional training. Training was conducted in a single pass through all of the touches.




Figure 3: Loss by trial and position-in-sequence. Note that while the loss on the two unpredictable positions (denoted by X) does not meaningfully decrease.

Evaluation

We used two different accuracy metrics to assess model fit. In both cases, accuracy was measured as how often the highest probability value from the softmax of the motor activations m_t corresponded to the stimulus in V_t . This was calculated for the test set for both overall accuracy, and the accuracy on B where B is the predictable element in the regularity.

²A different version of the model used an 11-dimensional vector which included `START` and `END` states and returned somewhat different results which we do not report here.

³When fed one at a time, the model did not learn, and I don’t know why.

Results

The model is quite effective at learning the regularity, and learns very quickly with less input data than the baboons saw, minimizing loss on B in some 250 trials (Fig. 3). In the same way, we see more errors being made on earlier touches than on later ones (Fig. 4). Qualitatively this is comparable to baboon behavior RTs got faster over the course of training. An important contrast with the baboons is observable in Fig. 3 where reduced loss is observed on the A element in the regularity – not as much as on B but still less than the X elements. In contrast, Yeaton et al. did not observe significant improvement in RTs on A .




Figure 4: Touch-by-touch performance at different stages of training. Each vertical column represents a touch. Yellow squares are false alarms (targeted the wrong state), and dark squares are the corresponding missed targets. Note that the x-axis is in *touches*, not *trials*.

Accuracy on Regularity

In terms of accuracy in the learned regularity, we observe quite different patterns between conditions and their interaction with λ -values (Fig. 5 - top panel). While the fixed-position conditions ($ABXX$, $XABX$, and $XXAB$) have ceiling-level accuracy on B at lower values of λ up to about 0.6, performance drops off rather steeply after that and decreases as λ approaches 1. The three fixed position conditions do not vary much from each other. Some differences are observed, but the amplitude and direction of these differences varies between runs.

In stark contrast, accuracy in the variable position condition (Var-4) starts relatively low, increases in accuracy up to $\lambda = 0.6$, then drops back off again. This contrast is important for two reasons: 1) the fixed position conditions behave differently from the variable po-

sition condition, and 2) the fixed and variable conditions interact differently with λ .




Figure 5: Accuracy by condition and λ -value for learning with no START and END states. Top: Accuracy on the B element in the AB regularity. Bottom: Overall accuracy.

It is possible that different results would be observed if trials were separated by START and END. In order to test this, we used a slightly different set of training data wherein additional states were included in each trial – one for the start of the trial, and one for the end. Thus, instead of $XXABXXABXXAB\dots$ as the sequence of states, it was $SXXABTSXXABT\dots$ where S and T are the START and END states respectively.

We found a qualitatively different pattern of results when these START and END states were used (Fig. 6 - top panel). Whereas without them, the three fixed position conditions remained quite similar in accuracy, with the additional states, then the $ABXX$ and $XXAB$ conditions significantly outperformed the $XABX$ condition with $0.5 < \lambda < 0.9$.

Overall Accuracy

The results of the overall accuracy are much less interesting (Fig. 5 - bottom panel). No condition approaches ceiling-level accuracy, and the dropoff observed in the




Figure 6: Accuracy by condition and λ -value for learning with START and END states. Top: Accuracy on the B element in the AB regularity. Bottom: Overall accuracy.

fixed position conditions with increasing values of λ is not as dramatic. We do not observe much meaningful impact of λ in the variable position condition, as accuracy remains more or less stagnant until $\lambda = 1$. So again, we see different patterns of behavior in the fixed and variable position conditions, and different interactions with λ .

With the START and END states, we see a significant increase in overall accuracy across the board (Fig. 6 - bottom panel). Some of this is driven by the transition from END to START which has a transitional probability of 1. Nonetheless, the START and END states drive a significant increase in accuracy.

Discussion

The models' behavior bears both important similarities and differences from the baboons'. In terms of similarities, the accuracy in the fixed position conditions does not appear to significantly differ amongst themselves which is consistent with the results of Yeaton et al.'s Experiment 1. It is important to note, however, that the stimuli were presented to the model as a continuous

string of touches whereas the baboons were presented with 4-touch trials with distinct start and end points.

In terms of differences, the model diverges in two crucial ways from what would have been predicted if it were using the same extraction mechanisms as the baboons. The first is that the fixed position conditions have quantitatively and qualitatively different patterns of behavior from each the variable position conditions. While the baboons did not perform differently between these conditions, the model has produced drastically different results. The second is that the fixed and variable conditions interact differently with the λ parameter. Because no differences were observed between conditions in the baboon behavior, one would predict that they would interact comparably with λ , but we find that the fixed position conditions begin at ceiling-level accuracy and drop off with increasing values of λ . By contrast, the variable position conditions begin with relatively low accuracy, increase to a maximum at $\lambda = 0.6$, and drop off again without ever reaching ceiling. Thus, while it seems to be the case that the *AB* regularity is learned in both the fixed and variable position conditions the weight of previous touch information plays a different role in each.

Furthermore, in the models trained with START and END states, we see that the regularity is extracted more readily adjacent to a boundary (*ABXX* and *XXAB* conditions). This is in contrast to what is observed in the baboon behavior where we saw no such advantage in these conditions.

Circling back to the question of what is being learned from Young (1962), it seems that in the fixed position conditions, the serial position information provides a significant boost in learning, while the variable position condition suffers for a lack thereof. It seems that the models trained on the fixed position stimuli learn both the serial position as well as the regularity, while the variable position models can only learn the regularity as the serial position information is unreliable.

While more simulations are necessary, this work provides crucial first step towards modelling the dynamics of regularity extraction in from noisy input which can help us to better understand non-human primate cognition, as well as human cognition.

The code to run the simulations and produce the figures is available here: https://colab.research.google.com/drive/1DIyxJjnBBoqF1UdHrmb1h_Jek9zS1nK0?usp=sharing

References

Ebbinghaus, H. (1885). *Über das gedächtnis: untersuchungen zur experimentellen psychologie*. Duncker & Humblot.

- Hebb, D. O. (1961). Distinctive features of learning in the higher animal. *Brain mechanisms and learning*, 37, 46.
- Minier, L., Fagot, J., & Rey, A. (2016). The temporal dynamics of regularity extraction in non-human primates. *Cognitive science*, 40(4), 1019–1030.
- Perruchet, P., & Vinter, A. (1998). Parser: A model for word segmentation. *Journal of memory and language*, 39(2), 246–263.
- Yeaton, J. D., Tosatto, L., Fagot, J., Grainger, J., & Rey, A. (2021). *Simple questions on simple associations: Regularity extraction in Papio papio*. (62nd Annual Meeting of the Psychonomic Society (Psychonomics 2021))
- Young, R. K. (1962). Tests of three hypotheses about the effective stimulus in serial learning. *Journal of Experimental Psychology*, 63(3), 307.