



HAL
open science

Learning Higher-Order Transitional Probabilities in Nonhuman Primates

Arnaud Rey, Joël Fagot, Fabien Mathy, Laura Lazartigues, Laure Tosatto, Guillem Bonafos, Jean-marc Freyermuth, Frédéric Lavigne

► **To cite this version:**

Arnaud Rey, Joël Fagot, Fabien Mathy, Laura Lazartigues, Laure Tosatto, et al.. Learning Higher-Order Transitional Probabilities in Nonhuman Primates. *Cognitive Science*, 2022, 46 (4), <10.1111/cogs.13121>. <hal-03631173>

HAL Id: hal-03631173

<https://hal.science/hal-03631173v1>

Submitted on 5 Apr 2022

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

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



HAL Authorization

Learning higher-order transitional probabilities in nonhuman primates

Arnaud Rey^a, Joël Fagot^{a,b}, Fabien Mathy^c, Laura Lazartigues^c,

Laure Tosatto^a, Guillem Bonafos^{a,d}, Jean-Marc Freyermuth^d, and Frédéric Lavigne^c

- a. Laboratoire de Psychologie Cognitive, CNRS & Aix-Marseille Université,
Marseille, France
- b. Station de Primatologie - Celphedia, CNRS UAR846, Rousset, France
- c. Bases, Corpus, Langage, CNRS & Université Côte d'Azur, Nice, France
- d. Institut de Mathématiques de Marseille, CNRS & Aix-Marseille Université,
Marseille, France

RUNNING TITLE: HIGHER-ORDER TPS

Corresponding author:

Arnaud Rey
Laboratoire de Psychologie Cognitive
CNRS & Aix-Marseille Université
3, place Victor Hugo - Case D
13331 Marseille Cedex 03 – France
arnaud.rey@cnrs.fr

ORCID: 0000-0001-8204-483X (A.R.)

Abstract

The extraction of cooccurrences between two events, A and B, is a central learning mechanism shared by all species capable of associative learning. Formally, the cooccurrence of events A and B appearing in a sequence is measured by the transitional probability (TP) between these events, and it corresponds to the probability of the second stimulus given the first (i.e., $p(B|A)$). In the present study, nonhuman primates (Guinea baboons, *Papio papio*) were exposed to a serial version of the XOR (i.e., exclusive-OR), in which they had to process sequences of three stimuli: A, B, and C. In this manipulation, first-order TPs (i.e., AB and BC) were uninformative due to their transitional probabilities being equal to .5 (i.e., $p(B|A) = p(C|B) = .5$), while second-order TPs were fully predictive of the upcoming stimulus (i.e., $p(C|AB) = 1$). In Experiment 1, we found that baboons were able to learn second-order TPs, while no learning occurred on first-order TPs. In Experiment 2, this pattern of results was replicated, and a final test ruled out an alternative interpretation in terms of proximity to the reward. These results indicate that a non-human primate species can learn a nonlinearly separable problem such as the XOR. They also provide fine-grained empirical data to test models of statistical learning on the interaction between the learning of different orders of TPs. Recent bioinspired models of associative learning are also introduced as promising alternatives to the modeling of statistical learning mechanisms.

Keywords: Associative learning, predictive coding, transitional probabilities, XOR

Introduction

The prediction of future stimuli requires the learning of statistical relations between stimuli in sequences (Erickson & Desimone, 1999; Rey et al., 2019). The prediction of a stimulus B based on a preceding stimulus A requires knowledge of the transitional probability (TP) of the second given the first ($p(B|A)$), which can be learned as an associative strength between the two stimuli. Such paired learning is thought to occur at the synaptic level by standard Hebbian learning mechanisms between two populations of neurons, with each coding for one stimulus (Mongillo et al., 2003).

The prediction of a stimulus can also depend on several preceding stimuli in a temporal sequence and on higher-order TPs (Gureckis & Love, 2010). Higher-order TPs are characterized by one stimulus depending on more than one previous element in a context. The order of a relationship is defined as the number of previous elements upon which the prediction depends. For instance, a pattern of input stimuli such as $A \Rightarrow B$ is a first-order statistical relationship, while $AB \Rightarrow C$ is a second-order relationship (here, A , B and C are considered events or stimuli). Although many studies have focused on first-order TPs, notably in the field of implicit statistical learning (Christiansen, 2019), less is known about higher-order TPs and the related brain mechanisms supporting such fundamental predictive abilities.

A paradigmatic case regarding the relation between first- and second-order TPs is the serial version of the exclusive disjunction (also called, exclusive-OR, or more commonly XOR; see Elman, 1990; Minsky & Papert, 1969). In this specific case involving a sequence of three stimuli, A , B , and C , first-order TPs do not predict the occurrence of the next item given the previous item (i.e., $p(B|A) = p(C|B) = .5$), while second-order TPs are deterministic (i.e., $p(C|AB) = 1$) and fully predict the occurrence

of the third stimulus given the previous two stimuli. To more concretely illustrate the serial XOR, suppose six distinct stimuli (A , B , C , D , E , and F) are displayed throughout four sequences composed of specific serial combinations of these stimuli: ABC , ADE , FBE , and FDC . In this example, the TP between the first stimulus (i.e., A or F) and the second stimulus (i.e., B or D) is equal to .5 (indeed, A is either followed by B in the first sequence ABC or by D in the second sequence ADE). Similarly, the TP between the second stimulus (i.e., B or D) and the third stimulus (i.e., E or C) is also equal to .5. In a serial version of the XOR, each combination of the first and second stimuli (i.e., AB , AD , FB , and FD) fully predicts the occurrence of the next stimulus (i.e., C , E , E , and C , respectively), with probabilities of $p(C|AB)=1$, $p(E|AD)=1$, $p(E|FB)=1$, and $p(E|FD)=1$. An interesting component of the XOR as a nonlinear separable case is that it proves to be challenging for learning because the second-order TP (e.g., $p(C|AB)$) cannot be computed from the first-order probabilities (i.e., $p(B|A)$ and $p(C|B)$) and, thus, can be learned only by taking into account the combination of the first two stimuli (i.e., AB).

A recent study (Minier et al., 2016) reported that nonhuman primates (Guinea baboons, *Papio papio*) can use conjointly first- and second-order TPs in a serial response-time task (Nissen & Bullemer, 1987). Over repeated exposures to ABC sequences, it was found that response times for C decreased faster than those for B , indicating that C benefited from the richer context of both A and B (and therefore, from both first- and second-order TPs). In this previous study, however, contrary to the logic of the XOR, both first-order TPs (i.e., $p(B|A)$ and $p(C|B)$) and the second-order TP (i.e., $p(C|AB)$) were equal to 1. In that case, C can be predicted on the sole basis of first-order TPs, such as $p(C|A)$ or $p(C|B)$, leaving the question of whether C was predicted based on the combination of A and B unanswered.

In the present study, we tested whether nonhuman primates can extract second-order TPs on the basis of a combination of two preceding stimuli when first-order TPs are noninformative. Using the serial response-time task from Minier et al. (2016), baboons were trained to touch a red circle that would appear at nine possible locations on a touch screen (see Figure 1A). Each trial was composed of three presentations of the red circle and three touch responses by the baboons. A food reward was provided at the end of each completed trial.

Studying predictive (or regularity extraction) abilities in humans often raises the question of implicit versus explicit learning (e.g., Cleeremans, Destrebecqz, & Boyer, 1998; Dienes & Perner, 1999). Disentangling these two sources of learning is indeed problematic in humans because of our verbal recoding and reasoning skills, allowing us to recode regularities into explicit rules or hypotheses about the statistical structure of the experimental material. One way to get around this bias is to test species that do not have linguistic recoding skills, such as the group of Guinea baboons (*Papio papio*) which is participating in the present experiments. We can be confident in this case that participants will not use language recoding skills to explicitly code the regularities and that their performance will mainly reflect the properties and the dynamics of their implicit and associative learning processes.

Testing baboons with a serial response time task also provides an online measure of the extraction dynamics of statistical regularities. Used conjointly with offline measures (Batterink, Reber, Neville, & Paller, 2015), online measures can strongly constrain the elaboration of current computational models of statistical learning (e.g., Elman, 1990; Endress & Johnson, 2021; Frank, Goldwater, Griffiths, & Tenenbaum, 2010; French et al., 2011; Giroux & Rey, 2009; Perruchet & Vinter, 1998;

Pothos, 2007; Tovar, Westermann, & Torres, 2018). They indeed provide a more direct measure of the dynamics of implicit statistical learning processes. More specifically, by monitoring the online ability of a non-human primate species to learn second-order TPS when first-order TPS are not predictive, the present experiment will inform us about the ability of this non-human species to learn a nonlinearly separable situation such as the XOR problem. It should also help the first generation of statistical learning models (e.g., Elman, 1990; Perruchet & Vinter, 1998) and the new generation of connectionist models (e.g., Endress & Johnson, 2021; Tovar et al., 2018) to take a step towards more precise predictions on the interaction between the learning of different orders of TPS.

Experiment 1

In Experiment 1, the serial XOR was implemented by randomly presenting to each baboon one triplet of positions out of four possible triplets. For example, one set of four triplets could be: 1-2-4, 7-2-9, 1-8-9 and 7-8-4 (see Figure 1B). In this case, the first possible positions (1 and 7) can be followed by two possible second positions (2 and 8). Similarly, second positions (2 and 8) can be followed by two possible third positions (4 and 9). Therefore, first-order TPS (i.e., between Positions 1 and 2, or between Positions 2 and 3) are equal to .5. However, the third stimulus is fully predictable based on the combination of the first and second stimuli (i.e., second-order TPS = 1). To test whether monkeys were able learn second-order TPS, we compared the evolution of the first transition time (i.e., TT1, from the first to the second position) to the second transition time (i.e., TT2, from the second to the third position). Theoretically, no learning can occur on TT1 (i.e., we predict no decrease in response time over time) given that the second position cannot be predicted from the first.

Conversely, if monkeys can extract the second-order TP, then we should observe a progressive decrease in response time on TT2.

< insert Figure 1 here >

Method

Participants

Fourteen adult females and six males (age range 2.4-20.7 years) participated in Experiments 1. The monkeys were members of a social group of 25 individuals living in a 700-m² outdoor enclosure containing climbing structures connected to two indoor experimental areas containing the test equipment (see below). Water was provided *ad libitum* during the test, and the monkeys received their normal food ration of fruits every day at 5 PM.

Apparatus

This experiment was conducted using a computer-learning device based on the voluntary participation of baboons (Fagot & Bonté, 2010). The baboons were implanted with RFID microchips and had free access to 10 automatic operant conditioning learning devices. Whenever a monkey entered a test chamber, it was identified by its microchip, and the system was prompted to resume the trial list at the place at which the subject left it at its previous visit. The experiment was controlled by a software testing program written by JF using E-prime (Version 2.0 professional, Psychology Software Tools, Pittsburgh, PA, USA).

Procedure

The general procedure was identical to that used by Minier et al. (2016). The screen was divided into nine equidistant positions represented by white crosses on a black background. A trial began with the presentation of a yellow fixation cross at the

bottom of the screen. After the baboon touched it, the fixation cross disappeared, and the nine crosses were displayed, with one of them being replaced by the target, namely, a red circle. When the first target circle was touched, it disappeared and was immediately replaced by the cross. The next position in the sequence was then replaced by the second red target circle until the end of the sequence was reached. A reward (a drop of dry wheat) was provided at the end of a sequence of three touch responses. To learn the task, the baboons initially received 1-item trials that were rewarded after one touch, after which the number of touches in a trial was progressively increased to three. If the baboon touched an inappropriate location (incorrect trial) or failed to touch the screen within 5 seconds after the red circle appeared (aborted trial), a green screen was displayed for 3 seconds as a marker of failure. Aborted trials were not counted as trials and were therefore presented again, while incorrect trials were not. The elapsed time between the appearance of the red circle and the amount of time before the baboon responded by touching the circle was recorded as the transition time between the first and second positions of the circles (i.e., TT1) and between the second and third positions of the circles in the sequence (i.e., TT2).

To control and match the motor difficulty of TT1 and TT2, each baboon was first tested on 1000 random triplets. On the basis of these random trials, a baseline measure for all possible transitions from one position to another was computed by calculating the mean TT for each transition (e.g., from Position 1 to 9) and for each monkey, yielding a 9×9 matrix of mean TT (calculated over the entire group of monkeys, see Appendix 1).

After these random trials, each monkey was exposed to 4000 trials, each involving one of the four possible regular sequences. These four 3-item regular

sequences were carefully constructed so that the mean TTs of their first and second transitions would not be statistically different compared with the baseline measurements obtained for these transitions during the random trials. Because we were interested in the evolution of TTs for the first and second transitions in the triplets within the regular sequences, a computer program was developed to find the smallest TT differences between these transitions within the random trials. Following these constraints, each monkey received a different set of four sequences (see Appendix 2).

Results

We analyzed the evolution of TT1 and TT2 by dividing the 4000 trials into 10 successive blocks of 400 trials (one trial being composed of three screen touches). Trials for which one of the three responses was incorrect were discarded. Inspection of the response times distribution revealed that a majority of responses were produced around 500ms. A smaller group of RTs appeared around 1,000 ms and was likely due to situations in which baboon's response was not recorded by the computer, because their hands were dirty. In this situation, they had to touch the screen again, and longer RTs were recorded (that are on average **twice as long** compared to the first responses). This is why we have adopted the following two-steps trimming procedure. We first excluded RTs higher than 800 ms. Second, a recursive trimming procedure excluded response times greater than 2.5 SDs in either direction of the participant's mean for each of the three possible positions in a trial. Mean TTs for TT1 and TT2 were then computed per block and per monkey, and the results are reported in Figure 2.

< insert Figure 2 here >

The RT analysis for Experiment 1 (see Table 1) revealed no main effect of Block ($\beta = -0.56$, $t = -1.07$, 95% CI [-1.6, .49]), a main effect of Transition, with TT2 being

faster than TT1 ($\beta = -6.27$, $t = -10.7$, 95% CI [-7.4, -5.11]), and an interaction effect between Block and Transition ($\beta = -1.59$, $t = -16.8$, 95% CI [-1.77, -1.4]) revealing that the difference between TT1 and TT2 increased across Blocks. We used R (R Core Team, 2012) and *lme4* (Bates et al., 2012) to perform a linear mixed effects analysis of the relationship between Block and Transition. As fixed effects, we entered the interaction between Block and Transition into the model. As random effects, we modeled the by-subject random slopes for the effect of Block. Visual inspection of residual plots did not reveal any obvious deviations from homoscedasticity or normality.

< insert Table 1 here >

Discussion

In Experiment 1, baboons could not predict the second position on the basis of the first because the transitional probability of the first transition was equal to .5. Although the third position could not be predicted on the basis of the second position (because the TP was also .5), baboons could improve their performance on the third position by taking into account both Position 1 and 2, i.e., by learning the second-order TP. Results clearly showed that they indeed managed to produce faster RTs on the third position compared to the second, and that this difference increased across blocks.

However, one possible limitation of Experiment 1 is the proximity between the last element of the sequence (i.e., C in ABC) and the reward provided to the monkey at the end of a series of correct touches. The decrease in RTs could indeed be due to the monkey accelerating its response to get the reward faster and not to the learning of the statistical structure of the triplet. Another possibility is that animals learn to better time their responses in the task, improving across trials, and that this timing is

facilitated by the two preceding responses more than just one. We tested these alternative interpretations in the next experiment.

Experiment 2

Experiment 2 had two objectives. First, to test the robustness of the results obtained in Experiment 1, we replicated the same design with a different set of XOR sequences for each monkey. Second, to test the alternative interpretations of the acceleration of TT2 in terms of proximity to the reward or of **a quicker timing** of their responses, we added a final block of 400 trials in which we switched the final element of each sequence. Baboons were then trained during 10 blocks on the four XOR sequences (*ABC*, *ADE*, *FBE*, and *FDC*) and the last element of each sequence was switched during the 11th block (i.e., *ABE*, *ADC*, *FBC*, and *FDE*). The switch did not change the value of first-order TPs (i.e., $p=.5$) but changed the value of second-order TPs (i.e., from $p=1$ to $p=0$). Because we simply switched the two possible final elements of the four initial sequences, if monkeys responded faster to TT2 due to its proximity to the reward or to **a quicker timing** of their response, the switch should not produce any difference in RTs. Conversely, if baboons had learned the second-order TPs (e.g., $p(C/AB) = 1$), they should be in trouble with the switch and longer RTs on TT2 should be observed in Block 11.

Method

Participants

Thirteen adult females and four males (age range 2.75-24.8 years) participated in Experiments 2. The monkeys were members of the same social group of baboons.

Apparatus

This experiment was conducted with the same computer-learning device as the one used in Experiment 1.

Procedure

The general procedure was identical to the one used in Experiment 1. Each monkey was exposed to 4000 trials, each involving one of the four possible regular sequences. These four 3-item regular sequences had the same properties as the ones used in Experiment 1. Each monkey simply received a different set of four sequences (see Appendix 3). In addition, after Block 10, monkeys had to perform a final block of 400 trials in which the last element of each sequence (e.g., C) was switched with the other possible last element (i.e., E).

Results

We analyzed the evolution of TT1 and TT2 as we did in Experiment 1. Mean TTs for TT1 and TT2 were then computed per block and per monkey, and the results are reported in Figure 3.

< insert Figure 3 here >

The RT analysis for Experiment 2 (see Table 2) revealed a main effect of Block ($\beta = -1.03$, $t = -2.97$, 95% CI [-1.73, -.34]), indicating that RTs decreased from Block 1 to 10, a main effect of Transition, with TT2 being faster than TT1 ($\beta = -8.37$, $t = -11.4$, 95% CI [-9.8, -6.9]), and an interaction effect between Block and Transition ($\beta = -1.82$, $t = -15.4$, 95% CI [-2.05, -1.59]), revealing that the difference between TT1 and TT2 increased from Blocks 1 to 10. Critically, mean RTs in Block 11 were significantly slower than in Block 10 ($t(16) = -2.58$, $p = .02$), indicating that the switch in the final position of the sequences had an effect on RTs.

< insert Table 2 here >

Discussion

Experiment 2 replicated the main results of Experiment 1 and demonstrated that baboons are able to learn second-order TPs. Thanks to the switch conducted in Block 11 on the last element of the sequences, we found that the baboons' decrease in RTs for TT2 was not due to the proximity with the reward or to a quicker timing of their response but can be attributed to their learning of the second-order TPs.

General discussion

The results of Experiment 1 showed that the difference between TT1 and TT2 increased from Block 1 to 10, indicating that nonhuman primates were able to learn fully predictive second-order TPs ($p = 1$) when first-order TPs were not fully predictive ($p = .5$), a situation corresponding to a serial version of the XOR. The same pattern of results was replicated in Experiment 2 and an additional control allowed us to rule out an alternative interpretation in terms of proximity with the reward or to a quicker timing of their response. These findings suggest that monkeys were capable of learning second-order TPs by taking into account a combination of the first and second stimuli to predict the third. More crucially, they show that a nonhuman primate species can learn a nonlinearly separable situation such as the XOR problem.

These data are consistent with previous findings obtained on humans suggesting that statistical learning processes can operate beyond first-order transitional probabilities (Perruchet & Poulin-Charronnat, 2012). Indeed, as mentioned by Perruchet and Poulin-Charronnat (2012, p. 808), statistical learning is often conceptually reduced to cases where first-order transitional probabilities seem to prevail (Aslin et al., 1998). While the present findings do not underestimate the

prominence of first-order TPS for statistical learning, they clearly indicate that learning can occur at higher orders and *independently* of first-order TPS.

Accounting for the learnability of first-order and higher-order TPS is certainly an important challenge for models of statistical learning. Let us consider, for example, two of the most influential models in this field, the Parser model (Perruchet & Vinter, 1998) and the Simple Recurrent Network model (hereafter: SRN; Elman, 1990). A critical feature of Parser is the assumption that repeated sequential patterns will progressively become chunks of information. In the present situation, each of the four triplets in the XOR situation may therefore be coded as chunks according to Parser. However, if the triplets become chunks, Parser would certainly predict that RTs on the third position of a triplet should become faster with practice but also RTs on the second position, which does not fit with the present data. On the other side, SRN is typically designed for predicting first-order TPS but the presence of context units may also allow the model to take into account higher-order TPS. Context units are supposed to receive a copy of the activation hidden units had one time-step before. Although one might question the plausibility of this computational trick, it certainly implements a general feature of the brain that can keep activated just-processed information in order to combine it with currently-processed information. This way, the SRN model could certainly account for the baboons learning a XOR situation.

New generations of computational models based on the assumptions that each item in a sequence is coded by a specific population of neurons and that learning between these populations is following Hebbian learning principles, may provide promising and more plausible descriptions of these statistical learning mechanisms (e.g., Endress & Johnson, 2021; Tovar et al., 2018). These models may use the

present set of results to test if they can predict that second-order TPS can be learnt when first-order TPS are unpredictable.

Similarly, bioinspired models have also addressed the question of associative learning of second-order TPS in cortical network models (Bourjaily & Miller, 2011; Lavigne et al., 2014; Rigotti et al., 2010). These types of bioinspired models also assume that each element of an *ABC* sequence is coded by a specific population of neurons. Learning between these neural populations occurs at the level of mixed-coding neurons (Bourjaily & Miller, 2010; Rigotti et al., 2010) and dendrites (Lavigne et al., 2014). They provide a simple and biologically grounded framework accounting for the learning of first- and second-order statistical relationships in terms of synaptic learning (see also Lavigne et al., 2016). Learning second-order TPS is possible in an initially unstructured network that does not require specific assumptions about the coding of stimuli in various prewired hierarchical layers of neurons. Using a generalization of the classical Hebbian learning rule, the efficacy of single synapses is influenced not only by the two neural populations coding for the two stimuli (as in standard Hebbian learning) but also by other neural populations coding for the whole sequence of three stimuli (Lavigne et al., 2014).

In these models, during the online presentation of stimulus sequences, long-term potentiation (LTP) and long-term depression (LTD) determine the efficacy values between populations of neurons coding for the different stimuli. The synaptic matrix generated by learning encodes the probabilistic relations between stimuli (i.e., first-, second-, and probably higher-order TPS). After learning, the presentation of a stimulus activates the neuronal population coding for that stimulus (i.e., retrospective activity). The activated population, in turn, activates associated populations (i.e., prospective

activity) according to the learned synaptic matrix. The level of activation of a given population can be used as a predictor of the transition time from one stimulus to the next. Therefore, the present framework provides theoretical tools for understanding the extraction of statistical regularities during the processing of sequences and according to the various TPs between groups of stimuli (Baker et al., 2002; Wallis et al., 2001; Wallis & Miller, 2003).

Compared to connectionist models of statistical learning, such as the SRN model (Elman, 1990), this type of generalized Hebbian learning provides a simple alternative account of the data without making prior assumptions regarding the structure of the network in layers or the coding of stimuli in input/output layers. Assumptions regarding the learning mechanisms are also compatible with our current neurobiological knowledge of synaptic learning mechanisms (Govindarajan et al., 2011) and the mixed selectivity of neurons (Rigotti et al., 2013). Finally, it has an explicit computational format that allows for precise simulations and predictions. Future empirical work using longer repeated sequences (for example, with four successively repeated elements, such as *ABCD*) will be critical for testing if learning can occur between more than three populations of activated neurons, revealing a possible limitation in the range of neuronal populations that could be associated through generalized Hebbian learning mechanisms.

Acknowledgments

We are grateful to Gianluigi Mongillo for his comments and discussion regarding a previous version of the manuscript. This work was supported by the UCAJEDI Investments in the Future project (ANR-15-IDEX-01), the Convergence Institute ILCB (ANR-16-CONV-0002), the Premilang2 ANR-project (#ANR-13-BSH2-0002) and the Chunked ANR-project (#ANR-17-CE28-0013-02).

Open Practices Statements

Data from the experiment are available on Open Science Framework at https://osf.io/dg8nw/?view_only=547ae6c485ba48139c1d8713232657e9

Animal rights

This research adhered to the applicable French rules for the ethical treatment of research animals and received ethical approval from the French Ministry of Education (approval APAFIS#2717-2015111708173794 10 v3).

Author Contributions

A.R., J.F., F.M., F.L. designed research; A.R., J.F., L.L., L.T., performed research; A.R., G.B., J.-M.F. analyzed data; A.R., J.F., F.M., F.L. wrote the paper.

Conflict of interest

The authors declare no conflict of interest.

References

- Aslin, R. N., Saffran, J. R., & Newport, E. L. (1998). Computation of Conditional Probability Statistics by 8-Month-Old Infants. *Psychological Science*, 9(4), 321–324. <https://doi.org/10.1111/1467-9280.00063>
- Baker, C. I., Behrmann, M., & Olson, C. R. (2002). Impact of learning on representation of parts and wholes in monkey inferotemporal cortex. *Nature Neuroscience*, 5(11), 1210–1216. <https://doi.org/10.1038/nn960>
- Bates, D., Maechler, M., & Bolker, B. (2012). *lme4: Linear mixed-effects models using Eigen and syntax*. R package version 0.999999-0.
- Bourjaily, M. A., & Miller, P. (2011). Synaptic plasticity and connectivity requirements to produce stimulus-pair specific responses in recurrent networks of spiking neurons. *PLoS Computational Biology*, 7(2), e1001091.
- Christiansen, M. H. (2019). Implicit Statistical Learning: A Tale of Two Literatures. *Topics in Cognitive Science*, 11(3), 468–481. <https://doi.org/10.1111/tops.12332>
- Elman, J. L. (1990). Finding Structure in Time. *Cognitive Science*, 14(2), 179–211. https://doi.org/10.1207/s15516709cog1402_1
- Endress, A. D., & Johnson, S. P. (2021). When forgetting fosters learning: A neural network model for statistical learning. *Cognition*, 213, 104621.
- Erickson, C. A., & Desimone, R. (1999). Responses of Macaque Perirhinal Neurons during and after Visual Stimulus Association Learning. *Journal of Neuroscience*, 19(23), 10404–10416. <https://doi.org/10.1523/JNEUROSCI.19-23-10404.1999>
- Fagot, J., & Bonté, E. (2010). Automated testing of cognitive performance in monkeys: Use of a battery of computerized test systems by a troop of semi-free-ranging baboons (*Papio papio*). *Behavior Research Methods*, 42(2), 507–516. <https://doi.org/10.3758/BRM.42.2.507>

- Frank, M. C., Goldwater, S., Griffiths, T. L., & Tenenbaum, J. B. (2010). Modeling human performance in statistical word segmentation. *Cognition*, *117*(2), 107–125. <https://doi.org/10.1016/j.cognition.2010.07.005>
- Govindarajan, A., Israely, I., Huang, S.-Y., & Tonegawa, S. (2011). The Dendritic Branch Is the Preferred Integrative Unit for Protein Synthesis-Dependent LTP. *Neuron*, *69*(1), 132–146. <https://doi.org/10.1016/j.neuron.2010.12.008>
- Gureckis, T. M., & Love, B. C. (2010). Direct Associations or Internal Transformations? Exploring the Mechanisms Underlying Sequential Learning Behavior. *Cognitive Science*, *34*(1), 10–50. <https://doi.org/10.1111/j.1551-6709.2009.01076.x>
- Lavigne, F., Avnaïm, F., & Dumercy, L. (2014). Inter-synaptic learning of combination rules in a cortical network model. *Frontiers in Psychology*, *5*. <https://doi.org/10.3389/fpsyg.2014.00842>
- Lavigne, F., Longrée, D., Mayaffre, D., & Mellet, S. (2016). Semantic integration by pattern priming: Experiment and cortical network model. *Cognitive Neurodynamics*, *10*(6), 513–533. <https://doi.org/10.1007/s11571-016-9410-4>
- Minier, L., Fagot, J., & Rey, A. (2016). The Temporal Dynamics of Regularity Extraction in Non-Human Primates. *Cognitive Science*, *40*(4), 1019–1030. <https://doi.org/10.1111/cogs.12279>
- Minsky, M., & Papert, S. (1969). *Perceptron: An introduction to computational geometry*. Cambridge tiass., MIT.
- Mongillo, G., Amit, D. J., & Brunel, N. (2003). Retrospective and prospective persistent activity induced by Hebbian learning in a recurrent cortical network. *European Journal of Neuroscience*, *18*(7), 2011–2024. <https://doi.org/10.1046/j.1460-9568.2003.02908.x>
- Nissen, M. J., & Bullemer, P. (1987). Attentional requirements of learning: Evidence from performance measures. *Cognitive Psychology*, *19*(1), 1–32. [https://doi.org/10.1016/0010-0285\(87\)90002-8](https://doi.org/10.1016/0010-0285(87)90002-8)

- Perruchet, P., & Poulin-Charronnat, B. (2012). Beyond transitional probability computations: Extracting word-like units when only statistical information is available. *Journal of Memory and Language*, 66(4), 807–818. <https://doi.org/10.1016/j.jml.2012.02.010>
- Perruchet, P., & Vinter, A. (1998). PARSER: A Model for Word Segmentation. *Journal of Memory and Language*, 39(2), 246–263. <https://doi.org/10.1006/jmla.1998.2576>
- R Core Team. (2012). *R: A language and environment for statistical computing*. R Foundation for Statistical Computing, Vienna, Austria.
- Rey, A., Minier, L., Malassis, R., Bogaerts, L., & Fagot, J. (2019). Regularity Extraction Across Species: Associative Learning Mechanisms Shared by Human and Non-Human Primates. *Topics in Cognitive Science*, 11(3), 573–586. <https://doi.org/10.1111/tops.12343>
- Rigotti, M., Barak, O., Warden, M. R., Wang, X. J., Daw, N. D., Miller, E. K., & Fusi, S. (2013). The importance of mixed selectivity in complex cognitive tasks. *Nature*, 497(7451), 585-590.
- Rigotti, M., Ben Dayan Rubin, D. D., Wang, X. J., & Fusi, S. (2010). Internal representation of task rules by recurrent dynamics: the importance of the diversity of neural responses. *Frontiers in computational neuroscience*, 4, 24.
- Tovar, Á. E., Westermann, G., & Torres, A. (2018). From altered synaptic plasticity to atypical learning: A computational model of Down syndrome. *Cognition*, 171, 15-24.
- Wallis, J. D., Anderson, K. C., & Miller, E. K. (2001). Single neurons in prefrontal cortex encode abstract rules. *Nature*, 411(6840), 953–956. <https://doi.org/10.1038/35082081>
- Wallis, J. D., & Miller, E. K. (2003). From Rule to Response: Neuronal Processes in the Premotor and Prefrontal Cortex. *Journal of Neurophysiology*, 90(3), 1790–1806. <https://doi.org/10.1152/jn.00086.2003>

Table 1: Mixed model regression results for response times in Experiments 1.

<i>Fixed effects</i>	<i>Coefficient</i>	<i>SE</i>	<i>t</i>	<i>95% CI</i>		<i>p</i>
Intercept	424.9	6.23	68.1	412	437	< 2e-16
Block (1-10)	-0.56	0.52	-1.07	-1.6	0.48	0.299
Transition (TT1/TT2)	-6.27	0.59	-11.9	-7.4	-5.11	< 2e-16
Block*Transition	-1.59	0.09	-16.8	-1.8	-1.4	< 2e-16
<i>Random effects</i>	<i>Variance</i>					
Participant intercept	774.7					
Block participant	5.34					

Table 2: *Mixed model regression results for response times in Experiments 2.*

<i>Fixed effects</i>	<i>Coefficient</i>	<i>SE</i>	<i>t</i>	<i>95% CI</i>		<i>p</i>
Intercept	447.8	8.4	53.5	431	464.7	< 2e-16
Block (1-10)	-1.03	0.35	-2.98	-1.73	-.34	.008
Transition (TT1/TT2)	-8.37	.74	-11.4	-9.8	-6.9	< 2e-16
Block*Transition	-1.82	0.12	-15.4	-2.05	-1.59	< 2e-16
<i>Random effects</i>	<i>Variance</i>					
Participant intercept	1187.6					
Block participant	1.92					

Figure captions

Figure 1: The serial response-time task and the serial XOR. A: Representation of the nine possible positions on the touch screen, Position 6 being replaced by the target red circle. B: Representation of four triplets implementing the XOR.

Figure 2: Mean response times as a function of transition type and block number in Experiment 1. Each block corresponds to 400 successive trials (100 for each sequence). Error bars represent confidence intervals.

Figure 3: Mean response times as a function of transition type and block number in Experiment 2. Each block corresponds to 400 successive trials (100 for each sequence). Error bars represent confidence intervals.

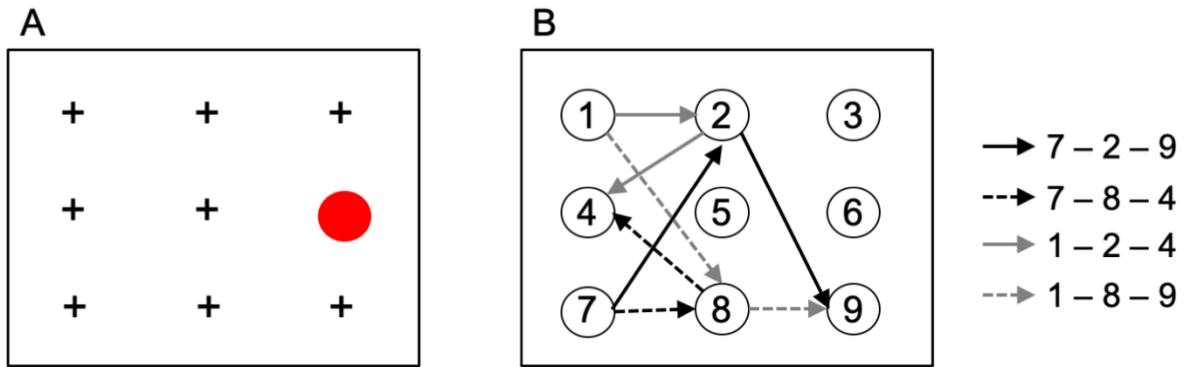


Figure 1

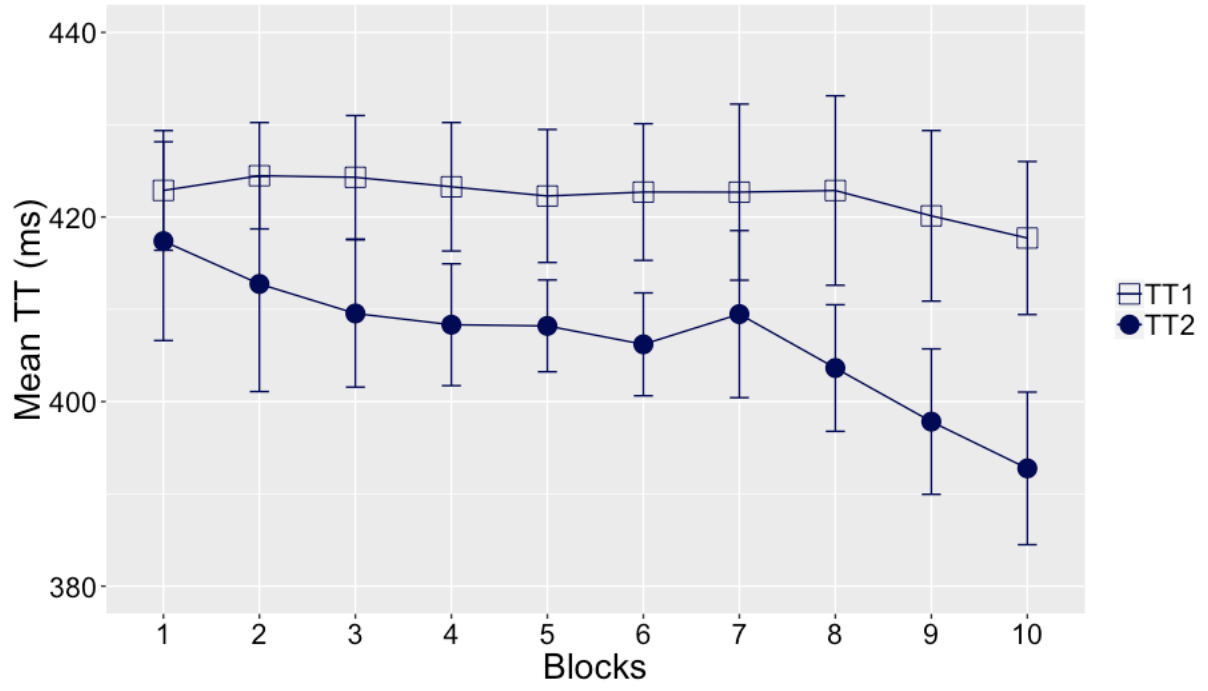


Figure 2

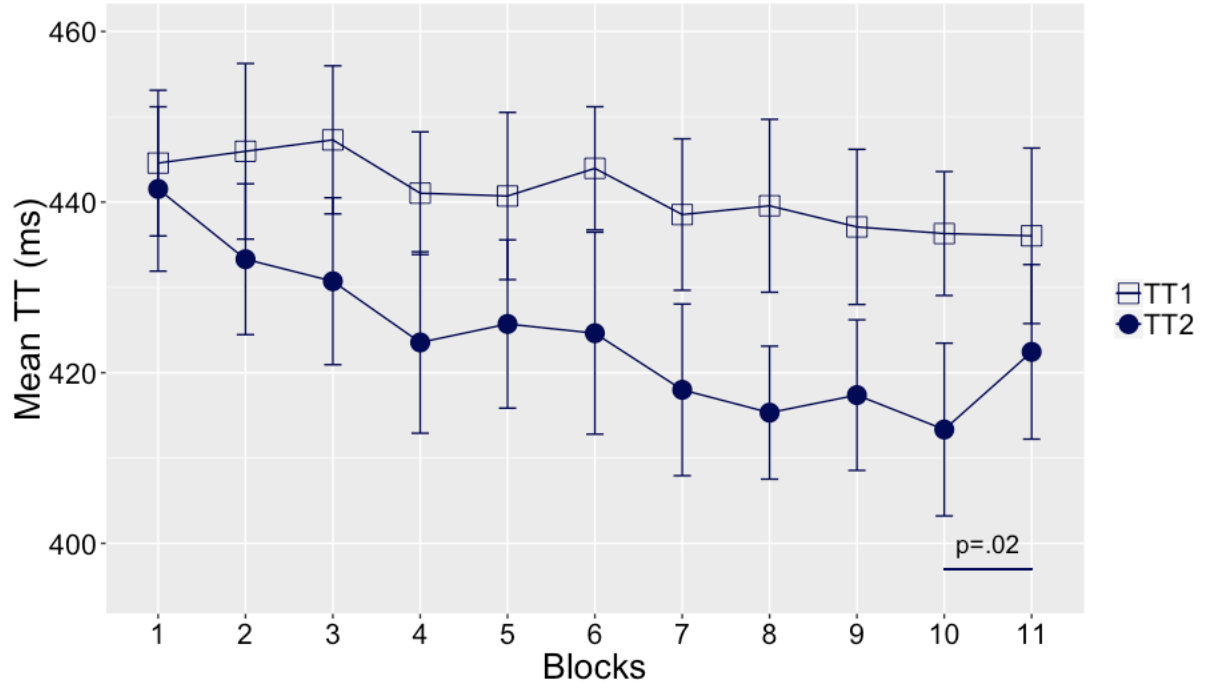


Figure 3

Appendix 1

Mean response times for each of the 72 possible transitions calculated from the 1000 random trials, over the entire group of baboons. These data were collected just before Experiment 1.

1 st Element in Transition	2 nd Element in Transition								
	1	2	3	4	5	6	7	8	9
1	-	426	421	438	365	360	447	359	371
2	506	-	457	411	377	393	391	365	393
3	502	435	-	443	368	353	439	372	365
4	486	423	448	-	366	374	434	339	358
5	485	408	378	444	-	345	449	392	380
6	477	383	379	426	344	-	448	384	418
7	472	424	435	423	370	381	-	374	371
8	445	388	401	396	342	367	443	-	396
9	487	403	410	425	334	361	437	362	-

Appendix 2

The four XOR triplets used for each monkey in Experiment 1.

Baboon	Triplet			
	1	2	3	4
1	213	718	248	743
2	218	713	243	748
3	243	713	218	748
4	718	213	248	743
5	218	748	243	713
6	743	718	248	213
7	743	718	248	213
8	568	364	328	524
9	368	324	528	564
10	328	364	524	568
11	368	324	528	564
12	324	368	564	528
13	568	328	524	364
14	564	368	528	324
15	129	526	186	589
16	129	589	186	526
17	186	526	589	129
18	529	586	189	126
19	529	586	126	189
20	189	586	126	529

Appendix 3

The four XOR triplets used for each monkey in Experiment 2.

Baboon	Triplet			
	1	2	3	4
1	748	713	218	243
2	368	324	564	528
3	364	524	328	568
4	564	528	368	324
5	189	586	126	529
6	526	186	589	129
7	589	186	129	526
8	529	568	328	369
9	568	529	328	369
10	962	914	864	812
11	964	814	912	862
12	358	394	698	654
13	354	398	658	694
14	914	962	864	812
15	814	862	912	964
16	914	213	284	983
17	214	913	984	283