Acta Psychologica 266 (2026) 106796

Contents lists available at ScienceDirect

Acta Psychologica

journal homepage: www.elsevier.com/locate/actpsy

Modeling human synchronization to rhythmic patterns with varying statistical regularities

Dunia Giomo [a,b,*], Federico Mancinelli [a,c], Andrea Ravignani [b,d,e,f], Domenica Bueti [a]

[a] International School for Advanced Studies (SISSA), Trieste, Italy
[b] Department of Human Neurosciences, Sapienza University of Rome, Rome, Italy
[c] University of Bonn, Transdisciplinary Research Area "Life & Health", Hertz Chair for Artificial Intelligence and Neuroscience, Bonn, Germany
[d] Centre for Music in the Brain, Department of Clinical Medicine, Aarhus University & The Royal Academy of Music, Aarhus/Aalborg, Denmark
[e] Research Center of Neuroscience "CRiN-Daniel Bovet", Sapienza University of Rome, Rome, Italy
[f] Institute of Cognitive Sciences and Technologies, National Research Council, Rome, Italy

ARTICLE INFO

Keywords:
Timing
Rhythm
Synchronization
Finger tapping
Bayesian inference
Rescorla-Wagner
Associative learning

ABSTRACT

The ability to learn and adapt to new temporal regularities is pivotal to successfully interact with our environment. Research on rhythm processing has identified which features boost the learning of new rhythmic patterns and facilitate predictive behaviors, such as motor synchronization to them. How do humans learn to synchronize with new temporal patterns when these lack strong rhythmicity and predictability? Here we hypothesized that optimal synchronization with such patterns would require a probabilistic form of associative learning. In our experiment participants tapped a finger in synchrony with an auditory stream in which a familiar and a novel pattern alternated. The two patterns were made of the same intervals but arranged in reversed order, thus having different structures. Within the stream, the two patterns occurred with complementary probabilities, following a gradual transition from familiar to novel pattern across trials. We modelled synchronization performances as the result of a Bayesian inference process parameterized by a Rescorla-Wagner learning rule. Results showed that participants learned to synchronize to the novel pattern only when its predictability clearly exceeded that of the old one, i.e., a few trials after the actual transition in the patterns' probabilities. Parameter estimation makes us hypothesize two learning strategies to be tested in future work. On the one hand, "fast learners" may quickly learn to synchronize to each pattern and adapt to changes between them across trials. On the other hand, "slow learners" may slowly learn to synchronize to a pattern and fail to fully adjust to a novel one. Overall, our results provide preliminary evidence that probabilistic associative learning may offer a parsimonious account for how individuals learn to synchronize with new, variably predictable temporal patterns. This learning could be modulated by subject-specific strategies to adapt to rhythmic changes.

1. Introduction

Countless sensory and motor experiences engage our remarkable abilities in detecting, learning and producing complex temporal regularities. We are able to enjoy a rhythmic change while listening to a jazz performance, and likewise easily switch from one motor sequence to another while playing sports. Such kinds of activities also require us to promptly adapt to changes in the temporal structure of the events, often at multiple timescales. This entails learning new temporal associations, deciding whether they are stable and adjusting our actions accordingly.

Humans' ability to track temporal regularities has its finest expression in the perception and production of rhythms. Extensive research on the topic (for a review see Grahn, 2012) has identified the combination of features that allow us to detect a rhythmic structure in external stimuli. These features include, among others, the presence of hierarchical levels of regularities (i.e., meter), an underlying regular pulse (i.e., beat), and durations related by small integer ratios (Essens & Povel, 1985; Fraisse, 1982; Grahn & Rowe, 2013; Honing, 2013). Humans quickly exploit these features and their predictability when synchronizing to rhythmic patterns. In particular, sensorimotor synchronization (SMS) studies investigate human and non-human animals' abilities to synchronize a motor output with an external rhythmic input (Iversen & Balasubramaniam, 2016; Repp, 2005; Repp & Su, 2013). Leveraging the regular and highly predictable structure of the input, synchronization

* Corresponding author at: Department of Human Neurosciences, University of Rome La Sapienza, 00185 Piazzale Aldo Moro 5, Rome, Italy.
E-mail address: dunia.giomo@uniroma1.it (D. Giomo).

https://doi.org/10.1016/j.actpsy.2026.106796
Received 11 August 2025; Received in revised form 3 April 2026; Accepted 6 April 2026
Available online 23 April 2026

behaviors are typically anticipatory (Repp & Su, 2013), i.e., the motor output slightly predates the onset of the stimulus based on temporal predictions. Several SMS studies have tested how unpredictable changes in the timing of the stimulus affect tapping performance. Most paradigms intelligently employ local perturbations among the inter onset intervals (IOI, i.e., the intervals between adjacent cues) of a sequence. Typically, the disrupted sequence is an isochronous one (i.e., one where all IOIs are equal, hence maximally predictable). Disruption causes sudden phase or period variations (e.g., Repp & Keller, 2004; Repp, Keller, & Jacoby, 2012; Schulze, Cordes, & Vorberg, 2005); it lowers the sequence predictability and spurs participants to adjust their tapping (i.e., error correction; Repp, 2005) to re-synchronize to the new established regularity. Computational approaches to SMS account for re-synchronization behaviors mainly in three ways: as a re-adjustment between an internal time-keeper and the new input (linear, event-based models; e.g., Schulze & Vorberg, 2002; Jacoby, Tishby, Repp, Ahissar, & Keller, 2015), as a re-coupling between internal oscillatory dynamics and the external stimulus' periodicity (non-linear, oscillator-based models; e.g., Large, Fink, & Kelso, 2002), or as a dynamic update of internal, periodicity-based predictions about the incoming stimulus based on sensory evidence (Bayesian inference-based models; e.g., Cannon, 2021).

Synchronization perturbations in ecological contexts are not only limited to abrupt local IOI or global beat/meter changes in highly regular rhythmic stimuli, rather they more often come about as fluctuations in intrinsically variable rhythms (Torre, Varlet, & Marmelat, 2013). What happens instead if the perturbation does not involve the durations of IOIs, but their overarching sequential structure, like scrambling letters in the game Scrabble? This is the case of changes in the structure of complex temporally structured stimuli like speech or motor sequences, which feature neither isochronous sequences, nor global periodicities (e.g., they lack a clear underlying beat). Yet, these complex stimuli still retain a temporal structure that identifies them and can elicit synchronized behaviors, like joint actions (Keller, Novembre, & Hove, 2014) or synchronous speech production (Cummins, 2009, 2011). Moreover, perturbations in the temporal structure of sequences might introduce a stable change, equivalent to a new, fully predictable temporal regularity. Alternatively, changes might become stable over time, thus requiring a more flexible prediction update. In other words, how do humans learn to synchronize to a new temporal structure which introduces statistical, hence predictability, changes over the structure of a familiar one?

We can hypothesize that efficient synchronization to a new, variably predictable temporal structure would require multiple computations: first, learning the novel temporal structure, based on a re-mapping of the association between external cues and motor actions; second, dynamically tracking how much the learned structural change is stable or volatile in the current context; third, based on this, updating through Bayesian inference what is the most likely temporal structure to come next and adjust synchronization behavior accordingly. Canonical computational models of SMS, apt at explaining either changes in local IOI or global periodicities, cannot fully account for such a probabilistic kind of synchronization (but see Elliott, Wing, & Welchman, 2014, for a Bayesian account of SMS in the presence of conflicting periodic acoustic cues). A probabilistic form of associative learning (Gershman, 2015; Kruschke, 2008), instead, may support this type of synchronization. Both Bayesian inference and associative learning rules represent simple and powerful computations through which human and non-human animals acquire information from the environment, make predictions about the world and guide their actions. In fact, Bayesian inference and associative learning have been repeatedly identified, either in combination or individually, as the underlying principles guiding a multitude of cognitive and behavioral processes (Kang, Tobler, & Dayan, 2024; Soto, Vogel, Uribe-Bahamonde, & Perez, 2023). The combination of associative principles and Bayesian inference has been extensively applied to reward-based learning and decision-making under uncertainty (Behrens, Hunt, Woolrich, & Rushworth, 2008; Gershman, 2015; Kang et al., 2024). In many of these studies participants are instructed to learn stimulus-reward contingencies that feature a certain degree of uncertainty, for example due to context-dependent (e.g. Behrens, Woolrich, Walton, & Rushworth, 2007) or time-dependent (e.g. Piray & Daw, 2020) variations. Thus an optimal learning strategy in such tasks is to rely on probabilistic inference to weigh different stimulus-response associations simultaneously and flexibly adjust these weights (Kruschke, 2008).

Theories of associative learning traditionally posit that whenever two events (e.g. a stimulus and an action) are repeatedly experienced together, an association between them is learned and stored in memory as a direct connection between the two events (Pearce & Bouton, 2001). This is true also when learning sequences of multiple stimulus-response events, which are represented in memory as a distributed network of direct association nodes (Eichenbaum, 2017). Studies on sequential learning have indeed shown that associative principles can account for the learning of complex sequence material (i.e., sequences featuring higher-order relations between events) too. For example, Gureckis and Love (2010) employed a serial reaction time task using sequences with different statistical and transformational complexity. They showed that a simple associative model, based on direct associations, could better predict participants' performance compared to a transformations-based recurrent network. This finding is particularly relevant to the focus of our study because temporal patterns are indeed a type of complex sequence and, when lacking strong periodicity, their encoding and memorization can require sequential learning (Bouwer, Werner, Knetemann, & Honing, 2016).

In the current study we asked whether learning to synchronize to a novel, weakly predictable temporal pattern is supported by probabilistic associative learning. Specifically, we tested how humans learn to synchronize to a new temporal pattern, made of the same intervals of a previously learned one, which gradually becomes more predictable across trials. To this aim we parameterized a gradual transition between two temporal patterns with symmetric structures, i.e., composed of the same intervals but arranged in reversed order. We manipulated the probability of occurrence of each of the two patterns across trials to investigate how their predictability affected the learning transition experienced by our participants. We presented the two patterns embedded in an auditory stream in a synchronization task, where participants were instructed to tap their finger in synchrony with the stimuli. We conducted model-agnostic and model-based analyses on tapping data, with the goal of defining and retrieving the parameters underlying the learning transition expressed in participants' synchronization performance. More specifically, our models draw upon computational accounts of probabilistic associative learning and define an optimal synchronizer who solves the task using Bayesian inference, following a Rescorla-Wagner learning rule (Rescorla & Wagner, 1972). We hypothesized that our participants would adopt this strategy to learn and successfully synchronize to the new temporal pattern, while taking into account the patterns' volatile predictability across trials.

2. Methods

2.1. Participants

Fourteen participants (10 females; 3 left-handed; mean age 25 ± 3.2; age range 21–32) were recruited for the present study. Three participants reported some years of formal musical training (range 3–8 years). Before starting the experiment, each subject gave written informed consent to take part in the study. The experimental protocol was approved by the SISSA Ethics Committee (protocol number 11175-III/13) and complied with the Declaration of Helsinki. All recruited participants were naïve to the purposes of the study and received fixed monetary compensation after completing the task.

2.2. Experimental design and stimuli

Our study focuses on how human participants learn to synchronize to a novel temporal pattern, which becomes gradually more predictable than a familiar one across time. To this aim, we designed a gradual transition between two temporal patterns embedded in an auditory stream. Each pattern was created as a series of beeps. We instructed our participants to tap the index finger in synchrony with each beep. We employed two temporal patterns composed of the same intervals, but arranged in reversed order, such that only their overall structure changed while their information content remained identical. The transition between the two patterns was parameterized following a sigmoid function, with the probability of occurrence of the familiar and novel patterns gradually decreasing and increasing across trials, respectively.

2.2.1. Stimuli

We employed two patterns (A and B, Fig. 1b) made of brief beeps separated by silent pauses (i.e., IOI). Beeps consisted of 33 ms pure tones of 1000 Hz with 2-ms rise and fall time filter, created using a MATLAB (R2019b, The MathWorks) custom code. The IOI used were identical across both patterns (250 ms, 500 ms, 750 ms), related by small integer ratios (1:2:3), with only the order of the IOI differing between patterns. The structures of the two patterns mirrored each other: pattern A contained IOI of 500–250–250-750 ms (with ratios 2:1:1:3) while pattern B contained IOI of 750–250–250-500 ms (with ratios 3:1:1:2). Using the same set of IOI arranged in reversed order ensured the two patterns had identical information content (i.e., same levels of overall and conditional entropy), but still differed in their global structures. Additionally, these specific structures were selected to ensure comparable rhythmic complexity: neither could metrically fit a simple 4/4 bar while both were easy enough to be learned. The two patterns were presented multiple times per trial as a continuous auditory stream (see next paragraph for a detailed description). Hence the selected structures also ensured that the resulting auditory stream would not feature a highly regular overarching metrical grid, with accents perceived always on-beat (Grahn & Brett, 2007; Povel & Essens, 1985).

2.2.2. Transition

Within a trial the two patterns were presented several times in random order, resulting in a continuous auditory stream. The relative frequency of presentation of the two patterns within the stream changed parametrically across trials following a logistic function (see Fig. 1a). The function's midpoint was set at the 16th trial (i.e., halfway the total number of trials) and its slope at 0.1. This manipulation entailed a gradual transition whereby pattern A was mostly experienced in the first half of the trials and its presence gradually decreased as the trial number increased. Conversely, pattern B progressively replaced pattern A as the trial number increased and was mostly present in the second half of the trials. Both the probability of the patterns and the random order in which they appeared within a trial of the first half (1–15) were mirrored in an equivalent trial of the second half (17–31). This meant, for example, that in the 17th trial the proportions of pattern A and B within the auditory stream were reversed compared to the 15th trial, and the order in which they alternated was specular to that of the 15th trial (an example of trials' mirror-randomization is shown in Fig. S1a).

2.3. Apparatus

The same custom-made set up (www.cynexo.com) was used to deliver the auditory stimuli series and collect participants' tapping. The system included a microcontroller Arduino Due (www.arduino.cc), connected to a response pad and a photodiode, and a microcontroller Arduino Uno paired to a shield for sound delivery (Adafruit Wave Shield for Arduino, www.adafruit.com), into which noise-cancelling headphones were directly plugged.

Each trial's stimuli series was generated using MATLAB (R2019b, The MathWorks) and Psychtoolbox (Brainard, 1997; Kleiner, Brainard, & Pelli, 2007; Pelli, 1997). Each stimulus onset was coded as a visual trigger appearing on the screen (hidden from the participant) and read by the photodiode. The instantaneous change of status of the photodiode triggered the two microcontrollers serially to play the beep sound, marking the onset of each auditory stimulus. The response pad was equipped with a pressure sensor (SingleTact CS15–4.5 N, www.singletact.com), attached underneath a resin-made button. The analogue signal was sampled at 100 Hz and reflected the force exerted on the button (range of sensitivity 0–4.5 Newton, represented in an arbitrary integer scale spanning from 0 to 1022). The resin-made button prevented any acoustic feedback generated by its press, which would otherwise be an unwanted byproduct of pressing a plastic-made button. Benchmarking tests before the experiment showed a constant delay of 23 ms in sound delivery, from the onset of the visual trigger to the headphones, with a negligible standard deviation (in the microsecond scale). This custom-made set up allowed a single common recording and timescale for all trial events (i.e., stimuli and responses).

2.4. General procedure and task

Participants sat in front of a computer screen and wore noise-

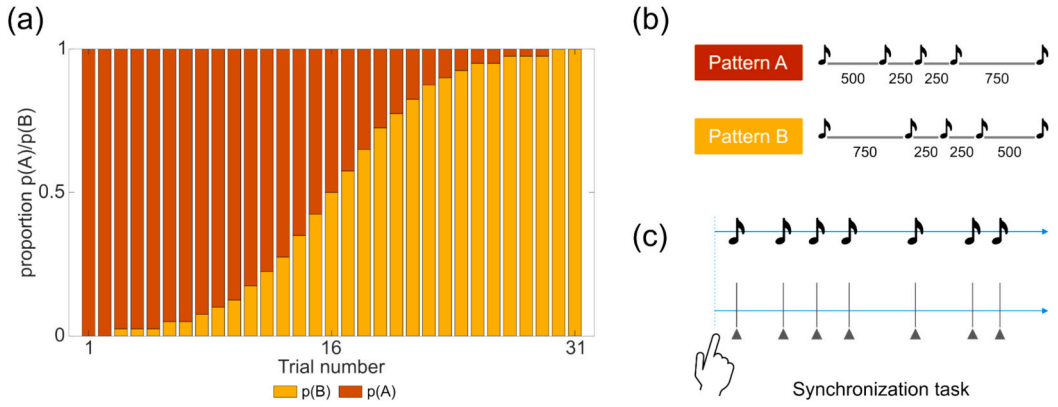

Fig. 1. Experimental design and task. (a) The bar plot shows the transition between pattern A and B across trials. The transition is parameterized as a gradual change in the probability of occurrence of the two patterns within each trial. The parameterization follows a logistic function (color code: pattern A = red, pattern B = yellow). (b) Pattern structure: pattern A and B are made of the same set of empty intervals, marked by brief tones, arranged in different order. The intervals are related by small integer ratios (1:2:3). The two structures are symmetrical, with the same level of information content. (c) Synchronization task: participants were asked to tap their finger in synchrony with an auditory stream made of several presentations of the two patterns. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

cancelling headphones. After adjusting the sound volume to a comfortable level and receiving general instructions on the experimental procedure, they performed a short training to familiarize with the synchronization task (Fig. 1c) and the response pad. The training consisted of two trials. Each trial was an auditory stream of 44 events (i.e., event = beep + empty interval), corresponding to 11 repetitions of the same temporal pattern. Participants were instructed to tap their finger on the response pad in synchrony with each beep. The training pattern was made of three intervals (270, 540 and 810 ms), different from those used in the main experiment. The training intervals were related by small integer ratio (1:2:3) with the following structure: 810–540–270-540 ms (3:2:1:2).

After the training participants started the main experiment, during which they performed 31 trials. Each trial was an auditory stream of 164 events, lasting approximately 80 s. The 164 events corresponded to 41 presentations of pattern A, B or both, following the parameterized trial-specific proportions. The beginning of each trial was preceded by a brief interval (randomly selected between 0.9 and 1.7 s), after which a message appeared on the screen concurrently with the onset of the first auditory stimulus of the first pattern's presentation. The message cued the participant to start tapping the finger on the response pad button in synchrony with the beep. The first pattern presented within each trial stream was a copy of the second and served to prompt participants to synchronize their tapping with the stream. The following 40 patterns were instead experimental presentations. Within a trial the probability of experiencing pattern A or pattern B was fixed, based on the parameterization defined in the experimental design. The order in which the patterns' repetitions occurred within a trial stream was instead randomized across subjects, with the constraint of the mirror-like configuration described above (see paragraph 2.2.2). The start of a new trial was self-paced to allow short breaks if necessary.

2.5. Data analysis and modelling

2.5.1. Pre-processing

Raw signals recorded by our delivery-response apparatus were pre-processed to extract stimuli and tap onsets. All pre-processing steps were performed using custom codes written in MATLAB and applied to each trial of each participant separately.

A stimulus onset was identified as a change of status of the photodiode, whose signal was recorded as a time series of 0 and 1, with 1 corresponding to the appearance of the visual trigger on the screen. Each positive derivative of the trigger's appearance was classified as an onset and stored with its timestamp, to which a delay of 23 ms was added (as estimated by previous technical tests on sound delivery timing, see paragraph 2.3).

The tapping signal obtained from each trial was a time series of force values recorded from the beginning to the end of it (an example of tapping signal from one trial is shown in Fig. S1b). Tap onset detection was performed through hierarchical steps. First, in order to remove mechanical noise-related fluctuations, the raw signal from each trial was smoothed using a 3rd order polynomial Savitzky-Golay filter (Savitzky & Golay, 1964), with a window of 15 samples (i.e., corresponding to roughly 150 ms). The standard deviation of the first 50 samples of each trial was calculated to obtain a baseline-noise standard deviation. As a second step, we calculated the mean force of the 150 highest values of the signal. Based on this average force, we categorized the signal either as “Strong”, “Medium” or “Weak”, and assigned a trial-specific derivative-threshold (i.e., a positive difference between consecutive force samples representing a “jump” in the signal, large enough to be caused by a tap). Based on this categorization, we also defined a trial-specific standard deviation threshold as a multiple of the baseline standard deviation. The standard deviation threshold represented the minimum amount of increase in the standard deviation of neighboring samples indicative of a “tap” event. These tailored thresholds allowed us to take into account the great variability observed between participants in the amount of force they exerted on the response pad. Based on these threshold values, the third step for the detection of tap onsets comprised the following inclusion/exclusion criteria. First, samples were filtered based on the presence of either a suprathreshold jump in the signal (i.e., a signal derivative greater than the derivative-threshold), or the cumulative presence of around-threshold derivatives in neighboring samples. For this second subgroup of potential onsets, the standard deviation threshold was used as additional criterion to filter the candidate onsets. Second, all samples within a 100 ms window past a potential onset were excluded (i.e., they were considered either as signals' artifacts or “false” onsets, e.g., double taps). The final criterion was the presence of large enough negative derivatives in the tapping signal between two candidate tap onsets, which indicated a tap offset (i.e., return to baseline of force values). As a result of the pre-processing stage, four participants had to be excluded due to technical issues in filtering their signal from baseline noise and reliably detecting tap onsets in more than 25% of their trials. Hence the resultant sample used for the analysis is $N = 10$ (7 females; 3 left-handed; mean age 25 ± 3.5; two participants with 3 and 5 years of musical training, respectively).

From the pre-processed series, all stimuli and tap onsets detected earlier than 350 ms before the 5th stimulus onset of the trial stream were removed. This was done because the first pattern of each trial, a copy of the actual first experimental pattern, was added to the trial stream only to prompt the synchronization of the tapping. Stimuli and tap onset series were then aligned by means of Dynamic Time Warping (DTW); this pattern matching algorithm allows to pair time series of unequal length by computing the best warping path to minimize their distance (Sakoe & Chiba, 1978). Being a synchronization task, this alignment allowed us to obtain an optimal one-to-one pairing between taps and stimuli and, consequently, detect and discard possible double taps (two taps matched with the same stimulus) or missing taps (stimuli with no match). The paired series were used to compute Inter Onset Intervals (IOI), Inter Tap Intervals (ITI), ITI error and asynchronies. IOI and ITI were calculated as the intervals between two adjacent stimuli onsets or tap onsets, respectively. IOI and ITI both measure intervals and are similar in logic; the former applies to what participants heard, the latter to what they produced. ITI error was calculated as the difference between matched ITI–IOI, with positive values indicating over-reproduction and negative values under-reproduction. Asynchrony was calculated as the positive or negative difference between paired tap–stimulus onsets, with negative values indicating anticipated taps and positive values delayed taps (Repp & Su, 2013).

2.5.2. Model-agnostic analysis

We analyzed tapping data focusing on the first interval of all occurrences of the patterns within a trial's stream. The 1st interval is in fact the most informative one from a learner perspective, as it allows to discriminate between the two temporal structures. More specifically, its identity cannot be inferred from the intervals that preceded it. Therefore its reproduction should be mostly influenced by internal predictions and more susceptible to the level of uncertainty of the trial. The reproduction of the 1st interval can thus represent a proxy of the ongoing learning and prediction update. To this purpose, we analyzed the error of the 1st ITI and the asynchrony of the 2nd tap onset, as they both can be referred to the reproduction of the first interval of the pattern (e.g., a positive asynchrony in the 2nd tap is indicative of an over-reproduction of the 1st ITI). We investigated how these measures differed between pattern A and B within and across trials. To overcome the unbalanced occurrences of the two patterns within a single trial, we binned the observations according to the number of times a given pattern was presented, irrespective of the trial number (one bin = 31 presentations of a given pattern; e.g. bin nr. 2 includes from the 32nd to 62nd presentations of pattern A and B). We then calculated bin-wise mean asynchronies and mean ITI errors for each participant and pattern.

We tested the effects of pattern type and bin history on mean asynchronies and mean ITI by fitting two separate Linear Mixed Effects

(LME) models. We included pattern identity and bin number as categorical and continuous predictors respectively, and participants as random intercept. LME models were fitted in R (R Core Team, 2025), using the packages *lme4* (Bates, Mächler, Bolker, & Walker, 2015) and *lmerTest* (Kuznetsova, Brockhoff, & Christensen, 2017), and *MuMIn* package (Bartoń, 2020) for calculating R^2 statistics. Overall effects were tested using Type III Analysis of Variance with Satterthwaite's method for degrees of freedom. Interaction effects were tested using the *emmeans* package (Lenth, 2020), using Kenward-Roger's method for degrees of freedom. Alpha level was always set at 0.05.

2.5.3. Model-based analysis

Our model-based analysis defines and estimates the parameters governing our participants' learning transition from a familiar to a novel temporal pattern, modulated by the change in the patterns' predictability across trials. We hypothesized that a successful synchronization to our target patterns and their evolving predictability would result from an optimal Bayesian inference process, parameterized by a Rescorla-Wagner rule (Rescorla & Wagner, 1972). We first modelled participants' behavior as the result of Bayesian causal inference on the probabilistic structure of each trial, thus obtaining a macroscopic description of their performances. In a second model we increased the granularity of this estimation to the single tap level and defined participants' inference process as governed by a simple Rescorla-Wagner rule, and subject to a memory decay factor. This allowed us to go beyond the macroscopic picture obtained with the first model, and describe in greater depth participants' learning transition from the familiar to the novel pattern. In both models we computed the likelihood of the 1st ITI of the patterns, as this measure directly refers to the most informative interval of both temporal patterns.

2.5.3.1. Mixture model. For each subject and each trial our first model estimates the parameter θ, representing the proportions of pattern A and B in the tapping series. The model is a simple mixture of two gaussians, with θ_A and θ_B (i.e., $\theta_B = 1 - \theta_A$) as mixture weight parameters. The likelihood is defined as:

$$p(y) = \theta_A N(y \mid \mu_A, \sigma_A) + \theta_B N(y \mid \mu_B, \sigma_B) \quad (1)$$

where y is the series of ITI, μ and σ are the experimentally defined means and standard deviations of the two stimuli distributions (i.e., only the 1st IOI). Each θ represents the mixture proportion of the distribution with mean μ_i and standard deviation σ_i (i.e., the proportion of pattern i). For the estimation of θ_A and θ_B we set as prior two beta distributions with scale hyperparameters a and b. The beta distributions account for the uncertainty in the posterior over each θ, which otherwise would be a point estimate. This model was also run on all IOI series as a sanity check to test whether the parameterized transition was accurately retrieved.

2.5.3.2. Static learning rate model. The second model quantifies the proportions of pattern A and B at the single tap level. Additionally, it estimates two learning rate parameters (one for each pattern) governing the evolution of the estimated patterns' proportions across trials. A mixture of two gaussians defines each tap's likelihood:

$$p(y) = p(A)N(y \mid \mu_A, \sigma_A) + (1 - p(A))N(y \mid \mu_B, \sigma_B) \quad (2)$$

where y is the n^{th} tap, μ_A and μ_B are the experimentally defined means of the two stimuli distributions (i.e., 1st IOI of the patterns), σ_A and σ_B are their unknown standard deviations. For the estimation of σ_A and σ_B we set a normal prior centered on 0 with variance defined by the hyperparameter λ. The mixture weight parameter $p(A)$ represents the probability of the n^{th} tap being pattern A, whereas $\{1 - p(A)\}$ corresponds to its counterpart for pattern B. The posterior over each $p(A)_n$ is estimated following a Rescorla-Wagner rule:

$$p(A)_n = \begin{cases} p(A)_{n-1} + \alpha_A(1 - p(A)_{n-1}) & \text{if pattern}_{n-1} = A \\ p(A)_{n-1} - \alpha_B\, p(A)_{n-1} & \text{if pattern}_{n-1} = B \end{cases} \quad (3)$$

where $p(A)_n$ is the probability of the n^{th} ITI of being pattern A; $p(A)_{n-1}$ is the previous ITI estimated probability; α_A and α_B are the learning rates ruling $p(A)_n$ update. The lower their values, the more static $p(A)$ estimation will be over time; the higher their values, the more frequently $p(A)$ will be updated, based on the identity of the previous pattern. While the posterior over $p(A)$ is estimated at the tap level, the two learning rates are assumed as fixed across trials. Hence the model estimates only one α_A and one α_B for each participant.

Following eq. [3], the posterior over each $p(A)_n$ is weighted according to a decay parameter, accounting for memory persistence (vs. fading) of the two patterns across trials. This forgetting process is based on F-Q variants of Q-learning models (Ito & Doya, 2009; Xia et al., 2021) and is implemented at each n^{th} tap as:

$$p(A)_n = \begin{cases} (1 - decay_1)^* p(A)_n + decay_1 ^* 0.5 & \text{if pattern}_{n-1} = A \\ (1 - decay_2)^* p(A)_n + decay_2 ^* 0.5 & \text{if pattern}_{n-1} = B \end{cases} \quad (4)$$

where $p(A)_n$ is the probability of the n^{th} ITI of being pattern A; $p(A)_{n-1}$ is the previous ITI estimated probability; $decay_1$ and $decay_2$ are the forgetting rates for previous choices ($p(A)_{1:n}$) of tapping pattern A and B respectively. The forgetting process is assumed to occur on every tap and the overall estimation represents an average decay factor for each of the two patterns. Similarly to the learning rates α_A and α_B, also $decay_1$ and $decay_2$ are bounded between 0 and 1, so that a maximal decay ($decay_i = 1$) would mean equal probability of performing pattern A or B at the n^{th} tap (i.e., $p(A) = 0.5$). A maximal decay is indicative of the absence of history effects on the probability of producing one pattern or the other.

We performed a subject-wise parameter recovery to evaluate the goodness of fit of the second model. Specifically, we used each participant's original parameter estimates to generate 20 synthetic datasets, according to eq. [2] with the inclusion of a gaussian noise component. The synthetic datasets were fed back to the model, thus obtaining 20 new sets of parameter estimates per participant. The new retrieved parameters were compared with the generating parameters.

Both models were implemented in Stan using RStan as interface (Stan Development Team, 2023; Carpenter et al., 2017). Stan uses a Markov Chain Monte Carlo sampling method to estimate posterior distributions over model parameters. For each participant (and each trial, in the case of the first model) we ran the models with four independent chains and 5000 iterations (2500 warm-up). Convergence of sampling chains was estimated through the Gelman-Rubin $\hat{R}$ statistic (Gelman & Rubin, 1992).

3. Results

3.1. Model-agnostic analysis

A regression analysis on ITI errors and asynchronies tested how our experimental manipulations affected participants' performances. Mean ITI errors and asynchronies were computed by binning the number of patterns' occurrences across trials (see paragraph 2.5.2) and were entered as dependent variables in two LME models. Both LME models included pattern identity and bin history as fixed effects (categorical and continuous predictors, respectively), and participants as random intercept (model formula: measure ~ pattern identity * bin number + (1 | subject)). Goodness of fit of both models was evaluated using R^2 statistic ($R^2_{marginal} = 0.70$, $R^2_{conditional} = 0.76$ for ITI error fit; $R^2_{marginal} = 0.33$, $R^2_{conditional} = 0.75$ for asynchrony fit). The LME model on mean ITI errors revealed a main effect of pattern identity (model estimate: $\beta = -0.2772$, SE = 0.012; type III ANOVA on model estimate: $F(1, 387) = 526.21$, $p < 0.0001$) and of bin history ($\beta = 0.0006$, SE = 0.0007; ANOVA: $F(1, 387) = 92.68$, $p < 0.0001$), meaning that ITI error values differed significantly between pattern A and B and across bin history. In terms of

tapping behavior, participants over-reproduced the 1st interval of pattern A and under-reproduced the 1st of pattern B. For both patterns, ITI errors gradually decreased as the number of patterns' presentations increased (see Fig. 2). The model on ITI errors also showed a significant interaction between pattern identity and bin history ($\beta = 0.0085$, SE = 0.001; ANOVA: $F(1, 387) = 71.43$, $p < 0.0001$). This was due to a significantly steeper reduction of the ITI error in pattern B compared to pattern A across bin history (A x bin: $\beta = 0.0006$; B x bin: $\beta = 0.0091$). Similarly, the LME model on mean asynchrony values showed a main effect of pattern identity ($\beta = -0.1564$, SE = 0.01; ANOVA: $F(1, 387) = 236.46$, $p < 0.0001$) and of bin history ($\beta = 0.0014$, SE = 0.0006; ANOVA: $F(1, 387) = 103.93$, $p < 0.0001$). These effects were due to a significantly greater negative asynchrony for pattern B compared to A and a gradual reduction of the asynchrony for both patterns across bin history (Fig. S2). This model also showed a significant interaction between pattern identity and bin history ($\beta = 0.0057$, SE = 0.0008; ANOVA: $F(1, 387) = 45.9$, $p < 0.0001$). The reduction in the amount of asynchrony across bins was, in fact, significantly greater for pattern B than A (A x bin: $\beta = 0.00145$; B x bin: $\beta = 0.0072$).

Overall, these results reveal that participants quickly learned to synchronize with pattern A, as suggested by smaller positive ITI errors (i.e., over-reproduction) and smaller negative asynchronies compared to pattern B throughout the task. On the contrary, synchronization with pattern B was characterized by greater ITI error in the opposite direction (i.e., under-reproduction) and larger negative asynchronies. The latter were large not only in the trials where pattern B's presence was unlikely, but also in subsequent trials where its probability of appearing was the same as pattern A. Nevertheless, pattern B showed a constant steep improvement in error and asynchrony throughout all trials, and reached comparable levels of accuracy of pattern A in the second part of the task. This improvement was likely due to participants learning the structure of pattern B and increasingly predicting its appearance towards the last part of the task.

3.2. Model-based analysis

The models aim at providing a computational account of a) how humans learn to synchronize to a new temporal pattern gradually replacing a familiar one, and b) how patterns' predictability influences the transition from one to the other. In the current study this process is defined as the transition from a known to a novel temporal pattern, made of the same intervals, which increasingly becomes more predictable across trials.

3.2.1. Mixture model

Our first model provides a coarse-grained picture of participants' performance: it estimates, trial by trial, the relative abundance (proportions) of the two patterns. Pattern A and B proportions are modelled as mixture weight parameters θ_A and θ_B of a mixture of two gaussians. The model includes a beta prior, with scale hyperparameters a and b, to estimate θ_A and θ_B. The beta prior quantifies the uncertainty in each θ estimation. We performed a sanity check of the estimation accuracy by feeding the model with each trial's stimuli series (i.e., 1st IOI of the patterns). Results showed that the parameterized transition embedded in the target patterns was accurately retrieved by the model. The trial-specific level of predictability was clearly captured by the evolution of the estimated beta distributions across trials (Fig. S3a and S3b). Concerning tapping performances, parameters estimation was done subject- and trial-wise using the 1st ITI of each occurrence of the patterns. A representative example of the model estimates for one participant is shown in Fig. 3a for the θ parameter and in Fig. 3b for the beta priors over θ. The evolution of estimated θ_A and θ_B reveals that this participant's tapping was reflecting quite accurately the real distributions of the two patterns across trials. However, the switch between the two θ happened later compared to the target patterns. This indicates that this participant was more likely to synchronize their tapping with pattern A even when its probability of occurrence was equal to that of pattern B. The estimation of the beta priors over θ further emphasizes the delayed switch in behavior. The estimated beta distributions are indeed characterized by greater uncertainty compared to those embedded in the target patterns, as evidenced by the large overlap between the two distributions persisting until the final trials. Results at the group level confirm those observed for this participant (posterior mean, std. and quantiles of θ_A and θ_B for each trial and participant are reported in Table S1). Fig. 4 shows a group summary plot of θ_A and θ_B estimation, where the delayed transition to pattern B is very clear. This delay may reflect either a carryover effect of pattern A on tapping performance and/or the influence of low predictability during the middle trials which may have biased the tapping towards the more familiar pattern. The prolonged overlap in the estimated beta distributions is also present at the group level (Fig. S3c), thus confirming the greater uncertainty compared to the distributions estimated from the stimuli series.

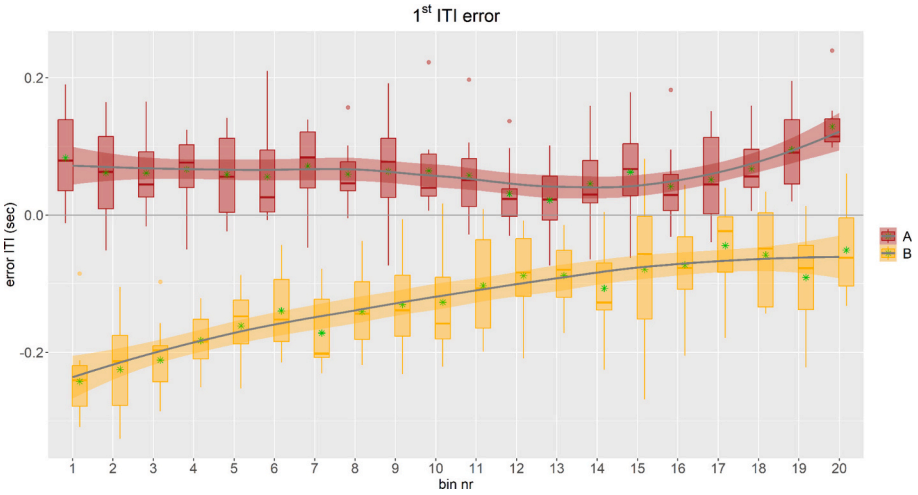

Fig. 2. Evolution of mean ITI error of the 1st interval of pattern A (intervals: 500–250–250-750 ms; color coded in red) and of pattern B (intervals: 750–250–250-500 ms; color coded in yellow) across presentation-bin history (see paragraph 3.2 for details on mean ITI error calculation). ITI errors are plotted as a function of the binned number of occurrences of the patterns across trials. Pattern A is characterized by a mild over-reproduction, stable across trials, whereas pattern B shows a large under-reproduction that progressively reduces. Statistical analysis confirmed the presence of a trial history effect for pattern B as a steeper slope coefficient (here graphically represented using LOESS method). Boxplots' borders represent interquartile range. Solid lines and green dots inside the boxplots represent median and mean ITI errors respectively. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

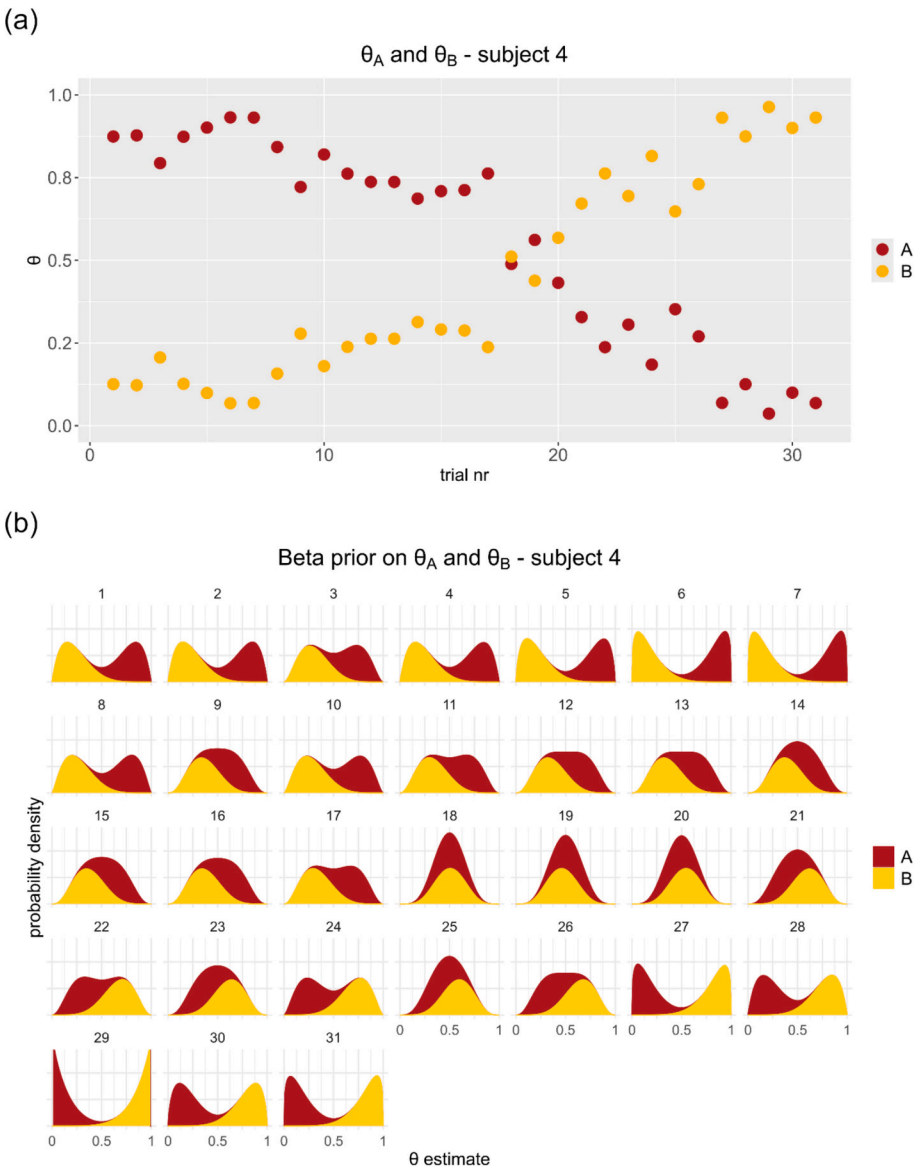

Fig. 3. Mixture model parameters estimation. (a) Estimated θ_A and θ_B (color code: θ_A = red, θ_B = yellow) across trials for participant n.4: the values reflect the proportions of pattern A and B across trials in the reproduction of this participant, as estimated by the mixture model. θ estimation reveals that participant n.4 was more likely to synchronize to pattern A even when its trial-specific probability was lower than pattern B, thus experiencing a delayed transition compared to the one designed in our target patterns (see Fig. 1a). (b) Estimated beta priors over θ_A and θ_B across trials for participant n.4, color-coded for pattern A and B as in (a). Beta distributions result from the estimation of the scale hyperparameters a and b, and act as priors over θ_A and θ_B. The overlap between each pair of beta distributions (one pair per trial) is a quantification of the uncertainty in the estimation of θ. As the plot reveals, the overlap between beta distributions is high until the 26th trial for participant n.4, which is indicative of greater uncertainty in the corresponding θ_A and θ_B plotted in panel (a), compared to the initial and last part of trials. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

3.2.2. *Static learning rate model*

The second model provides a finer-grained estimation of the proportions of pattern A and B in participants' performance; in fact, it has better temporal resolution and targets the underlying process governing the evolution patterns' proportion in time. The model, in brief, defines the probability of each tap belonging to pattern A ($p(A)$) or B ($\{1 - p(A)\}$) as a mixture of two gaussians. The mixture weight parameters $p(A)$ and $\{1 - p(A)\}$ are updated based on two static learning rates α_A and α_B and are each subject to a decay parameter, accounting for the persistence (vs. fading) of patterns' memory across trials. Parameters estimation was done subject-wise, using the 1st ITI of each pattern occurrence throughout the trials. Mixture weight parameters were estimated tap-wise, whereas a unique value was estimated for their respective learning rates, decay parameters and standard deviations of the two gaussians (posterior mean, std. and quantiles of learning rates, decay parameters and standard deviations for each participant are reported in Table S2). Results for parameter $p(A)$ estimation are shown for each participant in Fig. 5. In general, the estimation follows the delayed-switch scenario already seen in the estimation of θ from the first model. The estimated values reflect participants' tendency to produce pattern A more often, even after the point of transition designed in the target patterns (i.e., at trial 16). As indicated by the drop in $p(A)$ estimates, participants produced pattern B more frequently only in the final group of trials (approximately after the 20th), when its probability of occurrence was substantially higher than that of pattern A. These results appear broadly consistent across participants.

Participants, however, do not behave homogeneously on all fronts. Based on $p(A)$ estimates we may hypothesize two subgroups of performers: “fast learners” and “slow learners”. In the “fast learners” group (e.g. subject 7) $p(A)$ estimates relatively quickly saturate to 1 (i.e., a

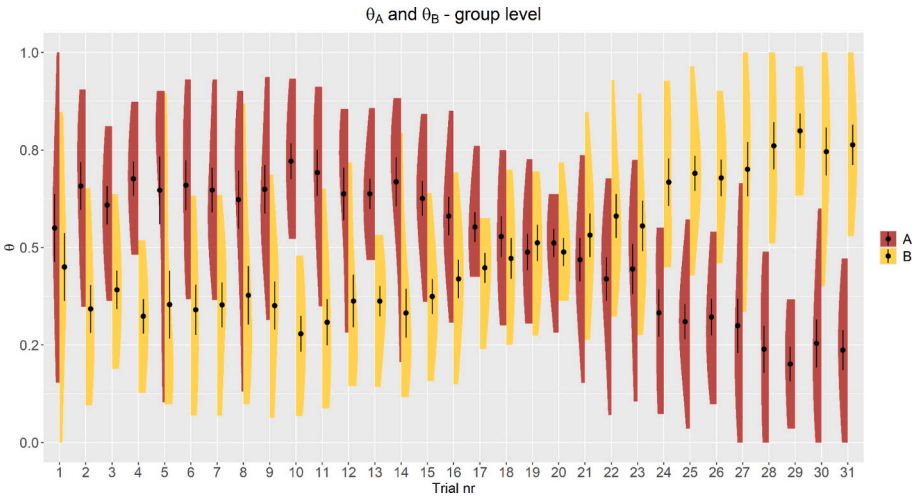

Fig. 4. Group-summary plot on estimated parameters θ_A and θ_B. Distribution of θ_A and θ_B values estimated for all participants using the mixture model. Theta values are plotted as a function of trial number and grouped by pattern identity (color code: θ_A = red, θ_B = yellow). θ estimation reveals that overall participants experienced a delayed transition from pattern A to pattern B. In terms of tapping behavior, this means that they were more likely to synchronize to pattern A even when its trial-specific probability was lower than pattern B. Black dots and lines represent mean and SE of estimates' distribution respectively. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

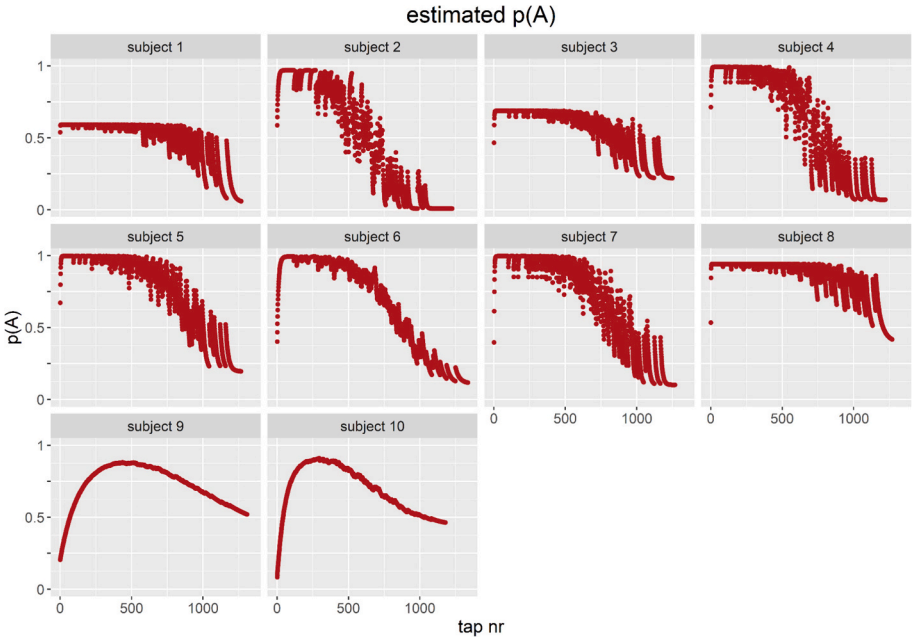

Fig. 5. Static learning rate model: parameter *p(A)* estimation. Each plot shows the subject-wise estimation of the parameter *p(A)*, representing the probability of reproducing pattern A tap by tap (i.e. for pattern B: {1 - *p(A)*}).

maximal probability of pattern A in the mixture) and remain stable over multiple taps, except for transient perturbations likely triggered by the appearance of the alternative pattern. Values of *p(A)* begin to decline shortly after the point of transition in the target patterns' proportions. Towards the end of the tap series *p(A)* reaches very low probability levels, indicating a full transition to pattern B. In the "slow learners" group (e.g. subject 9), instead, *p(A)* estimates take quite a large number of samples to reach high proportion values, if not remaining slightly above chance (e.g. subject 1) for the first half of trials. This latter case could suggest that during the first trials, the tapping was either very noisy or synchronized with an internal temporal reference. Additionally, *p(A)* estimates of "slow learners" tend to decline more gradually in the second half of trials, compared to the other group. In general, *p(A)* estimates for "slow learners" may reflect lower sensitivity to the actual proportions of pattern A and B defined in the target patterns.

The potential presence of two groups of performers is also suggested by specific asymmetries in estimated learning rates and decay rates across participants (Fig. 6a-b). A group of subjects (the same we termed "slow learners") shows either low learning rates for both patterns (e.g. subject 9) or a high decay specifically for pattern A (e.g. subject 1). Very low learning rates signal that *p(A)* estimates relied primarily on previous taps, whereas high decay for pattern A indicates a consistent decline in the memory of previous occurrences of pattern A. A second group of subjects (matching the "fast learners" one), instead, show higher learning rates and low decay rates for both patterns, compared to the first group: this suggests that *p(A)* estimates were more influenced by

the ongoing changes in patterns' occurrences across trials.

3.2.3. Parameter recovery

A subject-wise parameter recovery was performed to assess whether the static learning rate model could accurately simulate participants' behavior and to evaluate the consistency between parameters estimated from the original and the simulated data. All simulated datasets showed a positive correlation with the original data (Fig. 7a and S4a-b), suggesting the model was overall able to reproduce each participant's tapping performance. Based on correlation coefficients, participants seem to cluster around two groups that, interestingly, match the "fast vs slow learners" hypothesis: simulated data from fast learners better correlate with original data, compared to slow learners. Concerning tapping variability, the inclusion of a gaussian noise component at the stage of data generation allowed to mimic the inherent noise in participants' tapping (Fig. S4a). Additionally, the recovered estimates of the mixtures' standard deviations σ_A and σ_B had a good overlap with the original standard deviations for almost all participants (Fig. 7b). Similarly, results from simulation-based estimations of *p(A)* showed overall a high consistency with the original parameters (Fig. S5).

The estimation of learning rates α_A and α_B from the synthetic data led on average to coherent values with the original parameter estimates. In particular, the model consistently retrieved small values for α_B clustered around the original estimates, and high values of α_A. However, the results also revealed some non-negligible differences between simulation-based and original estimates of α_A for nearly half of the participants (Fig. 6a). Specifically, recovered values of α_A were highly variable, and in some cases they diverged substantially from the original parameters (e.g. subject 3). The recovered decay parameters, instead, mostly overlapped with the small values estimated from the original data (Fig. 6b). As for the two participants with high decay rates for pattern A, in all simulations the model correctly retrieved high estimates, albeit with some variability. The estimation differences found for the learning rates might be related to their static formalization in the model. We further examine this potential model optimization in the discussion section.

4. Discussion

In this work we investigated how humans learn to synchronize to a novel temporal pattern, which gradually replaces, hence becomes more predictable than, a familiar pattern composed of the same intervals. To this aim we embedded a gradual transition between two symmetrical auditory temporal patterns, by manipulating their probability of occurrence across trials, and presented them to participants in a synchronization task. We hypothesized that efficient synchronization would require a probabilistic form of associative learning and modelled tapping performances as the result of an optimal Bayesian inference process, guided by a Rescorla-Wagner rule. Model-based and model-agnostic analyses of synchronization performances revealed a delayed adaptation from the old to the novel pattern, compared to the time course of the transition designed in our target patterns. Model-based results identified the parameters describing the learning process underlying this delayed adaptation. Additionally, our model was able to reproduce participants' synchronization performance and consistently retrieve the subject-specific parameters based on simulated data. The estimated parameters seem to point, speculatively, to two potential subgroups of participants ("fast" and "slow learners"), based on the pace of the transition from one pattern to the other, as well as on the sensitivity to the predictability of patterns' occurrences across trials.

Despite some limitations, model-based results offer preliminary evidence in support of the hypothesis that participants would adopt a "Bayesian synchronizer" strategy to perform the task. Consistent with model-agnostic findings, model-based results revealed a more nuanced scenario of learning pace among participants. In particular, both analyses showed that almost all participants relatively quickly learned to synchronize to pattern A. This was expressed as lower ITI error and better asynchrony for pattern A compared to pattern B throughout the task, as well as in the estimation of the mixture weight parameter θ_A across trials (and of *p(A)* across taps). Pattern B was instead performed with lower accuracy than pattern A even beyond the middle trials. Accordingly, mixture weights' estimates reflected a delayed adjustment to the change in patterns' proportions relative to the target patterns, consistent with a carryover effect and with low decay rates estimated in the second model. Taken together, these results seem to suggest that pattern A served as a prior expectation to guide synchronization behavior, particularly when the stimuli were too unpredictable. Once the presence of pattern B became frequent and stable enough within the trial, mixture weights' estimates showed the transition in patterns' proportions. This was reflected in the reduction of ITI error and asynchronies for pattern B. With the second model, parameter estimation revealed some notable differences among our subjects. These differences led us to hypothesize the possible presence of two synchronization behaviors and underlying learning strategies. A group of participants ("fast learners") quickly learned to synchronize to pattern A, as indicated by their *p(A)* estimates rapidly saturating to the maximum proportion value. These participants appeared more sensitive to the ongoing changes in patterns' proportions, i.e., high learning rates for both patterns. Also, they showed a sharper transition to pattern B, i.e., a clear

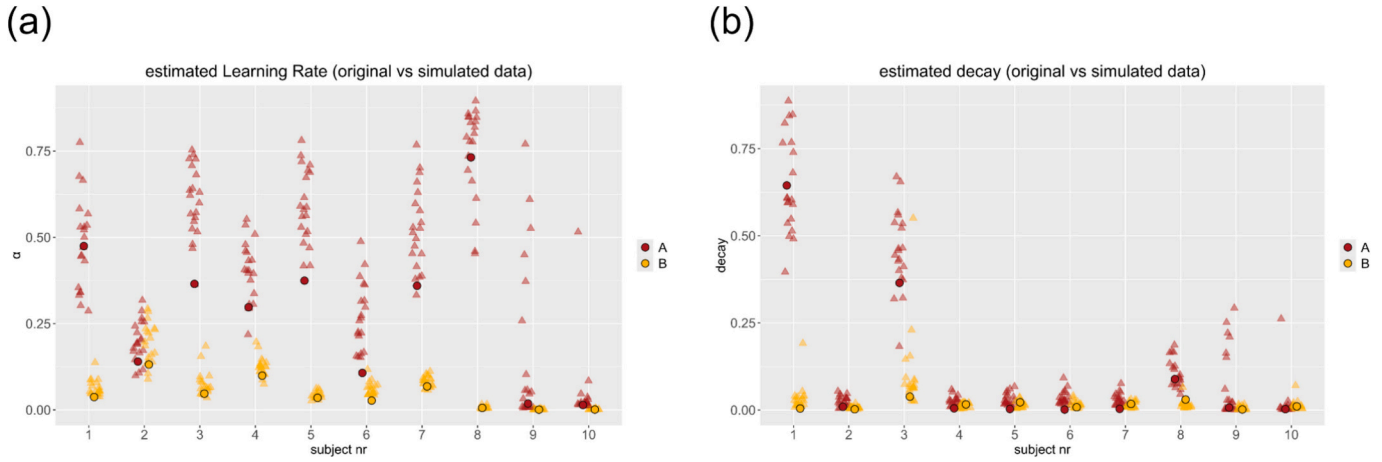

Fig. 6. Static learning rate model: learning rates and decay parameters. (a) Estimated learning rates α_A and α_B from original (circles) and simulated data (triangles), plotted by subject (x-axis). Color code: σ_A = red; σ_B = yellow). (b) Estimated decay rates for pattern A and B from original (circles) and simulated data (triangles), plotted by subject and color-coded as in (b). (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

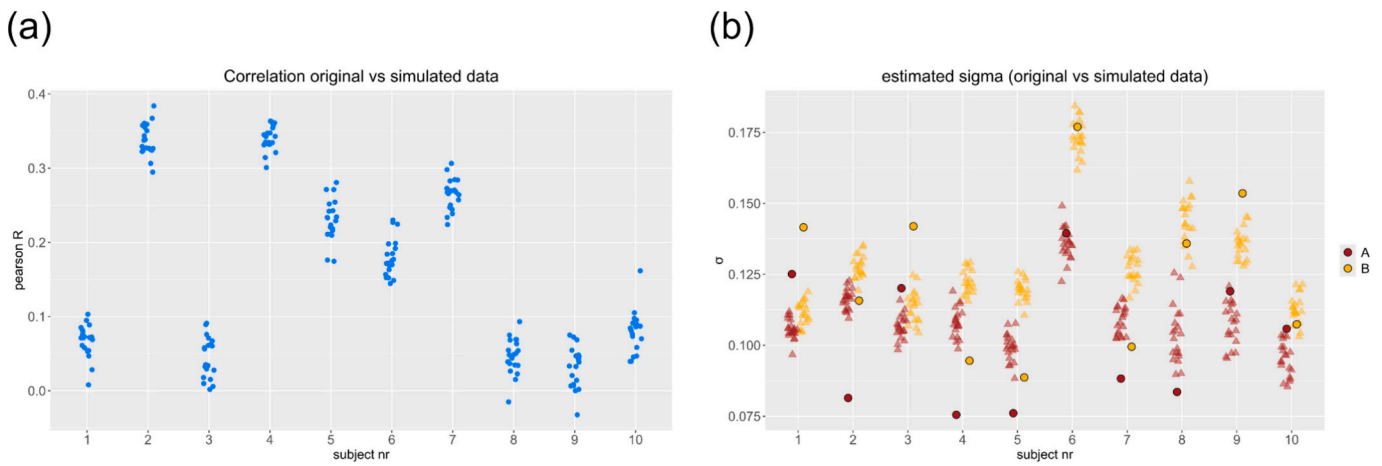

Fig. 7. Static learning rate model: parameter recovery. (a) Subject-wise correlation between original data and simulated data. Each dot represents the correlation (Pearson's r) between the subject's original dataset and each simulated datasets (20 per participant). Each column of dots refers to one participant (ordered on the x-axis). Simulated data were generated from the model, fed with each subject's estimated parameters. (b) Estimation of mixture components' standard deviations σ_A and σ_B from original (circles) and simulated data (triangles), plotted by subject (x-axis). Color code: σ_A = red; σ_B = yellow). (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

drop in *p(A)* estimates, albeit slightly delayed compared to the one defined in the target patterns. A second group ("slow learners") seemed to need on average more time to learn to synchronize to pattern A (i.e., *p(A)* estimates saturated at a slower pace, or not saturated at all). This group was less sensitive to trials' variability and, in some cases, stuck to pattern A longer than the other group (i.e., slower decrease in *p(A)* estimates). Some of these participants seemed to rely more on previous taps (i.e., very low learning rates), rather than on the actual changes in patterns' proportions across trials. Others, instead, showed a consistent forgetfulness for the history of occurrences of pattern A (i.e., high decay rate). Both behaviors thus resulted in a less clear and, on average, slower transition from pattern A to B. It is important to emphasize here that the sample size used for the analysis was relatively small. Thus both the general trend and the hypothetical learning profiles we observed should be put to the test in future studies using larger samples.

Our simulation-based results provide preliminary support to the hypothesis that participants may employ a probabilistic form of associative learning to synchronize with novel, variably predictable, temporal structures. Specifically, we were able to retrieve the evolution of the mixture proportions of pattern A and B estimated from participants' data, as well as the learning and memory decay parameters ruling the proportions update. These results suggest that simple associative rules combined with Bayesian inference may be sufficient (though not exhaustive) to explain the behavior observed. Our computational account is consistent with several studies showing how associative learning principles (Pearce & Bouton, 2001; Soto et al., 2023), Bayesian inference (Knill & Pouget, 2004) or a combination of both (Behrens et al., 2007; Behrens et al., 2008; Blanco & Moris, 2018) underlie many aspects of behavior and cognition. In contexts where new associations between events with variable degree of predictability (uncertainty) are introduced, similarly to our task, individuals may rely on a Bayesian form of associative learning (Gershman, 2015; Kruschke, 2008). This form of learning would allow to create and hold memory representations of multiple associations simultaneously and in a probabilistic way, i.e. each differentially weighted according to contextual uncertainty. When specific associations between events become more likely than others, as in our task, learning rates reflect the sensitivity to such changes in the environment and, in turn, the frequency at which predictions about future events are updated. In our model learning rates are assumed to remain static across trials, thus they reflect each participant's overall sensitivity to changes in patterns' occurrences throughout the task. Given the shape of the transition embedded in our target patterns and the different level of predictability across trials, we can reasonably hypothesize that learning rates could evolve across the task: higher rates when predictability is low, lower rates when predictability is high. A dynamic formalization of α_A and α_B that tracks trial uncertainty, for example in the form of a Kalman filter (Kang et al., 2024; Piray & Daw, 2020), could improve the explanatory power of our model. These dynamic learning rates would reflect how much participants are sensitive to changes in trials' volatility (Behrens et al., 2007).

Our model includes two decay parameters accounting for the forgetting of pattern A and B's past occurrences. In F-Q-learning models, which our formalization of decay parameters draws upon, such decay factors reflect a forgetting process affecting the action/choice *not taken* in the current trial (Ito & Doya, 2009; Master et al., 2020; Xia et al., 2021). In our model the forgetting process refers also to the pattern *not experienced* in the current trial. Our model, however, cannot disentangle whether the forgetting affects the memory representation of the pattern structure per se, or its retrieval process. We can speculate that the latter option is more likely. In fact, several studies on context-dependent learning and extinction (for reviews, see Bouton, 2004; Delamater, 2004) have shown that new associations between events do not necessarily overwrite the memory trace of previously learned associations, but coexist with them (Dunsmoor, Niv, Daw, & Phelps, 2015). Old associations are retained in memory and can be reactivated after exposure to the same context in which they were originally experienced (Bouton & Peck, 1989). However, when old and new associations are thought to be generated from the same latent cause, the old memory representation may be overwritten by the new one (Gershman, Monfils, Norman, & Niv, 2017). Future experimental and modelling works, that are based on our study, should investigate whether the forgetting process affects the memory representation of the patterns, their retrieval or both.

Relevant to our study, associative learning has been shown to play a key role also in context-dependent sensorimotor adaptation (Avraham, Taylor, Breska, Ivry, & McDougle, 2022). Our synchronization task shares some similarities with sensorimotor adaptation experiments. In fact, our participants needed to learn an association between an external cue (i.e., a temporal structure) and an action (i.e., the synchronized movement). The synchrony between cue and movement acted as a positive feedback reinforcing the learned association. The introduction of a novel temporal structure required learning a new association between cue and movement and adapting the latter to changes in task contingencies. However, the novel structure was introduced with varying degrees of predictability, thus the individual had to take into account uncertainty to guide his actions. This is where Bayesian inference comes into play: based on prior expectations, the individual had to

decide the temporal structure most likely to come, move in synchrony with it and update his expectations based on the outcome of his action (the match or mismatch between movement and cue).

Our basic computational account somehow implicitly assumes that a "controller" mechanism oversees both the learning and the prediction components involved in our task. In other words, there should be a form of integration or supervision over mapping, memorizing, predicting and updating the association cue-movement. The biological plausibility of this assumption is supported by an incredibly vast research on the central role of the cerebellum in sensorimotor adaptation and control (for extensive reviews, see Manto et al., 2012; Popa, Streng, Hewitt, & Ebner, 2016). Most importantly, the fundamental role of cerebellar structures in supporting temporal predictions (Breska & Ivry, 2016, 2018) and associative learning (Nguyen & Person, 2025) makes it the most likely "controller" mechanism potentially involved in our task (and model). Future works could test this idea both experimentally (e.g. with neurostimulation techniques) and computationally (e.g. model optimization by including a cerebellar controller).

Bayesian inference is at the core of recent computational accounts of synchronization behavior (e.g., Cannon, 2021; Elliott et al., 2014; Pérez, Delle Monache, Lacquaniti, Bosco, & Merchant, 2023). The study by Elliott et al. (2014), in particular, shares some similarities with ours. In their task participants had to tap along with two simultaneously presented metronomes differing either in phase or in standard deviation of the period. Essentially, to "solve" the task participants had to either separate the cues or integrate them. Synchronization asynchronies were statistically explained by a Bayesian inference-based mixture model that switches between separation (two mixture components) and integration (one component) of cues according to the degree of phase-offset between metronomes. The authors also showed that the amount of phase-offset prompting the switch between separation and integration was subject-specific. In our study the mixture components represent the two patterns. Rather than one or two mixtures, our participants had to "choose" in which proportions the mixtures were present in a given trial, which reflected their online predictions. In this sense, an optimal way to solve the task required a flexible, trial-by-trial switching between external stimuli and internal predictions. More specifically, on each trial our participants could choose to synchronize with the external stimuli (i.e., the pattern heard), hence frequently update their predictions. Alternatively, they could synchronize with the most likely pattern to occur, hence rely on relatively stable, stimuli history-based predictions. According to each trial's predictability, participants' synchronization would benefit more from one strategy than the other. These alternative strategies may explain part of the inter-individual differences we observed: some participants ("fast learners") seemed to update more frequently their predictions (i.e., high learning rates for both patterns), whereas others (a portion of "slow learners") seemed to rely more on previous taps and/or stimuli history (i.e., low learning rates).

Within a Bayesian perspective, a growing number of experimental works and theoretical accounts have proposed that rhythm processing follows the principles of predictive coding (Vuust, Ostergaard, Pallesen, Bailey, & Roepstorff, 2009, 2018; Vuust & Witek, 2014; Koelsch, Vuust, & Friston, 2019). According to this view, tracking and synchronizing to rhythms is a highly predictive and hierarchical process; it is described in terms of optimal Bayesian inference based on a continuous comparison between prior expectations (i.e., driven by metrical predictions, top-down) and sensory experience (i.e., driven by the experienced rhythmic patterns, bottom-up). Our results may indirectly support this framework. Specifically, the predictive coding view of rhythm processing posits that top-down predictions are mainly built upon specific rhythmic features, i.e., beat and meter. Even if our stimuli intentionally lacked a clear underlying beat and metrical hierarchy (see Method section), it is possible that some of our participants may have imposed some metrical grid in their tapping performance. This might explain the behavior of the hypothesized "slow learners": these participants appeared to be poorly sensitive to the external changes in the stimuli, but actually they might have performed the task following an internal metrical reference. More in general, we can speculate that these participants may have paradoxically used a meter-based strategy to guide their predictions, supporting the predictive coding framework.

4.1. Limitations and model optimization

The interpretation of our results is limited by the small sample size of participants we tested, thus it is based on preliminary evidence that must necessarily be confirmed by a replication study. Additionally, our parameter recovery analysis shows some limitations of our models, which could be addressed in future work. We will focus here on two of them which link to the extensive sensorimotor synchronization (SMS) literature. First, our formalization of tapping variance could be improved by incorporating in the current model classic accounts of motor timing variability (Getty, 1975; Ivry & Hazeltine, 1995; Wing & Kristofferson, 1973), their more recent developments (García-Garibay, Cadena-Valencia, Merchant, & De Lafuente, 2016; Laje, Cheng, & Buonomano, 2011), or their SMS-specific implementations (e.g., Vorberg & Schulze, 2002). These models take into account both motor-dependent and time-dependent components of the variance associated with a reproduced temporal interval. Second, including in the model an error correction mechanism, typically observed in tapping tasks, could contribute to explain part of the variability we observed in the data. When synchronizing with an external rhythm, in fact, humans typically adopt error correction mechanisms to re-adjust to it and compensate for the inherent variability of their motor actions (Repp, 2005). It is possible that such mechanisms were at play also in our task. Error correction, as probed in perturbations experimental paradigms, is a defining component of event-based models of SMS (Repp & Su, 2013), and is also incorporated in some oscillator-based (e.g., Bose, Byrne, & Rinzel, 2019) or Bayesian (e.g., Castillo-Almazán, Pérez, Prado, Jacoby, & Merchant, 2025) models of SMS. All these models, however, account for error correction when synchronizing to highly regular rhythmic stimuli, as those employed in standard SMS tasks (e.g., Repp et al., 2012). Our synchronization task is quite atypical compared to standard SMS tasks for two reasons: the peculiar experimental manipulation (i.e., the gradual change in the two patterns' predictability across trials), and the patterns used (i.e., complex rhythms without a clearly detectable underlying beat). Therefore, if any error correction mechanism influenced our participants' tapping performance, it is difficult to disentangle what triggered it: whether it was violations of the immediately expected IOI (i.e., local perturbation), internally driven metrical predictions, or predictions based on the statistical properties of the patterns. A possible way to overcome this limitation is to test different classes of models, in particular extensions of the ACT-R (adaptive control of thought - rational) architecture for skill acquisition (Anderson, Betts, Bothell, Hope, & Lebiere, 2019). In particular, the variant proposed by Gianferrara, Betts, and Anderson (2024) includes a double mechanism engaged in motor skill learning: a proper ACT-R model operating cognitive control and motor adaptation, paired with a simple periodic module which automatically executes the tapping based on a periodic motor clock. In our task this double mechanism could explain the initial learning stages for both patterns, the relatively automatic synchronization to them at later stages, as well as potential error correction mechanisms at play.

5. Conclusions

In the current study we investigated whether learning to synchronize to new, weakly predictable temporal patterns is supported by simple associative learning and Bayesian inference mechanisms. In our synchronization task a new temporal pattern, made of the same constitutive intervals of a familiar one, became increasingly more predictable across trials. Results from model-based and model-agnostic analyses showed that our participants experienced a delayed transition from the known to

the novel temporal pattern, relative to the transition designed in the target patterns. Model-based results offered preliminary support to the hypothesis that participants' behavior could be explained as the result of a Bayesian inference process, guided by a Rescorla-Wagner learning rule, though with important subject-specific nuances. Specifically, we observed individual differences in the pace of transition to the novel pattern, and in the sensitivity to the predictability of patterns' occurrences across trials. This led us to hypothesize the potential presence of individual learning strategies interacting with the more general learning mechanisms accounted for by our models. Overall the probabilistic associative learning account we proposed offers insights on how humans may learn and adapt to new, weakly rhythmic and variably predictable temporal regularities. Our study provides preliminary evidence in support of this general idea, thus any future research direction should start from a replication of our results. We suggest three possible directions to be pursued in future work. First, one could test if and how our modelling framework generalizes to different timing tasks involving similar probabilistic temporal patterns without overt motor synchronization. Second, our model formalization can be refined at multiple levels, as suggested throughout the discussion (e.g., introduction of dynamic learning rates, more complex memory components, or additional periodic mechanisms). Third, one could search for neural correlates of the computational mechanisms which explain our behavioral data.

CRediT authorship contribution statement

Dunia Giomo: Writing – review & editing, Writing – original draft, Visualization, Software, Methodology, Investigation, Formal analysis, Data curation, Conceptualization. **Federico Mancinelli:** Writing – review & editing, Methodology, Formal analysis, Conceptualization. **Andrea Ravignani:** Writing – review & editing, Funding acquisition. **Domenica Bueti:** Writing – review & editing, Supervision, Resources, Project administration, Funding acquisition, Conceptualization.

Ethics approval statement

The experimental protocol was approved by the SISSA Ethics Committee (protocol number 11175-III/13) and complied with the Declaration of Helsinki.

Funding statement

D.B. acknowledges the financial support of the European Research Council (ERC) under the European Union's Horizon 2020 research and innovation program (grant agreement no. 682117 BIT-ERC-2015-CoG), and the Italian Ministry of University and Research under the call PRIN2022 (project ID:2022CCPJ3J). A.R. acknowledges the financial support of the European Union (ERC, TOHR, 101041885) and of the HFSP research grant RGP0019/2022. The views and opinions expressed are those of the authors only and do not necessarily reflect those of the European Union or the European Research Council Executive Agency. Neither the European Union nor the granting authority can be held responsible for them. The Centre for Music in the Brain (MIB) is funded by The Lundbeck Foundation (R469-2024-1573) and Købmand Herman Sallings Fond.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Acknowledgments

We thank Gianfranco Fortunato for helpful discussions on the experimental design, setup and data analysis, and Fabrizio Manzino for technical support in the experimental setup.

Appendix A. Supplementary data

Supplementary data to this article can be found online at https://doi.org/10.1016/j.actpsy.2026.106796.

Data availability

Original scripts and data are available at the following link: https://osf.io/uqwa5/overview

References

Anderson, J. R., Betts, S., Bothell, D., Hope, R., & Lebiere, C. (2019). Learning rapid and precise skills. *Psychological Review, 126*(5), 727–760.
Avraham, G., Taylor, J. A., Breska, A., Ivry, R. B., & McDougle, S. D. (2022). Contextual effects in sensorimotor adaptation adhere to associative learning rules. *Elife, 11*, Article e75801.
Bartoń, K. (2020). MuMIn: Multi-Model Inference. R package version 1.43.17 https://CRAN.R-project.org/package=MuMIn.
Bates, D., Mächler, M., Bolker, B., & Walker, S. (2015). Fitting linear mixed-effects models using lme4. *Journal of Statistical Software, 67*, 1–48.
Behrens, T. E., Hunt, L. T., Woolrich, M. W., & Rushworth, M. F. (2008). Associative learning of social value. *Nature, 456*(7219), 245–249.
Behrens, T. E., Woolrich, M. W., Walton, M. E., & Rushworth, M. F. (2007). Learning the value of information in an uncertain world. *Nature Neuroscience, 10*(9), 1214–1221.
Blanco, F., & Moris, J. (2018). Bayesian methods for addressing long-standing problems in associative learning: The case of PREE. *Quarterly Journal of Experimental Psychology, 71*(9), 1844–1859.
Bose, A., Byrne, Á., & Rinzel, J. (2019). A neuromechanistic model for rhythmic beat generation. *PLoS Computational Biology, 15*(5), Article e1006450.
Bouton, M. E. (2004). Context and behavioral processes in extinction. *Learning & Memory, 11*(5), 485–494.
Bouton, M. E., & Peck, C. A. (1989). Context effects on conditioning, extinction, and reinstatement in an appetitive conditioning preparation. *Animal Learning & Behavior, 17*(2), 188–198.
Bouwer, F. L., Werner, C. M., Knetemann, M., & Honing, H. (2016). Disentangling beat perception from sequential learning and examining the influence of attention and musical abilities on ERP responses to rhythm. *Neuropsychologia, 85*, 80–90.
Brainard, D. H. (1997). The psychophysics toolbox. *Spatial Vision, 10*, 433–436.
Breska, A., & Ivry, R. B. (2016). Taxonomies of timing: Where does the cerebellum fit in? *Current Opinion in Behavioral Sciences, 8*, 282–288.
Breska, A., & Ivry, R. B. (2018). Double dissociation of single-interval and rhythmic temporal prediction in cerebellar degeneration and Parkinson's disease. *Proceedings of the National Academy of Sciences, 115*(48), 12283–12288.
Cannon, J. (2021). Expectancy-based rhythmic entrainment as continuous Bayesian inference. *PLoS Computational Biology, 17*(6), Article e1009025.
Carpenter, B., Gelman, A., Hoffman, M. D., Lee, D., Goodrich, B., Betancourt, M., … Riddell, A. (2017). Stan: A probabilistic programming language. *Journal of Statistical Software, 76*(1).
Castillo-Almazán, A., Pérez, O., Prado, L., Jacoby, N., & Merchant, H. (2025). Flexible tapping synchronization in macaques: Dynamic switching of timing strategies within rhythmic sequences. *Journal of Neurophysiology, 134*(2), 580–590.
Cummins, F. (2009). Rhythm as entrainment: The case of synchronous speech. *Journal of Phonetics, 37*(1), 16–28.
Cummins, F. (2011). Periodic and aperiodic synchronization in skilled action. *Frontiers in Human Neuroscience, 5*, 170.
Delamater, A. R. (2004). Experimental extinction in Pavlovian conditioning: Behavioural and neuroscience perspectives. *Quarterly Journal of Experimental Psychology Section B, 57*(2), 97–132.
Dunsmoor, J. E., Niv, Y., Daw, N., & Phelps, E. A. (2015). Rethinking extinction. *Neuron, 88*(1), 47–63.
Eichenbaum, H. (2017). Memory: Organization and control. *Annual Review of Psychology, 68*, 19–45.
Elliott, M. T., Wing, A. M., & Welchman, A. E. (2014). Moving in time: Bayesian causal inference explains movement coordination to auditory beats. *Proceedings of the Royal Society B: Biological Sciences, 281*(1786), 20140751.
Essens, P. J., & Povel, D. J. (1985). Metrical and nonmetrical representations of temporal patterns. *Perception & Psychophysics, 37*(1), 1–7.
Fraisse, P. (1982). Rhythm and tempo. In D. Deutsch (Ed.), *Psychology of music* (pp. 149–180). New York, NY: Academic Press.
García-Garibay, O., Cadena-Valencia, J., Merchant, H., & De Lafuente, V. (2016). Monkeys share the human ability to internally maintain a temporal rhythm. *Frontiers in Psychology, 7*, 1971.
Gelman, A., & Rubin, D. B. (1992). Inference from iterative simulation using multiple sequences. *Statistical Science*, 457–472.
Gershman, S. J. (2015). A unifying probabilistic view of associative learning. *PLoS Computational Biology, 11*(11), Article e1004567.
Gershman, S. J., Monfils, M. H., Norman, K. A., & Niv, Y. (2017). The computational nature of memory modification. *Elife, 6*, Article e23763.

Getty, D. J. (1975). Discrimination of short temporal intervals: A comparison of two models. *Perception & Psychophysics, 18*(1), 1–8.
Gianferrara, P. G., Betts, S. A., & Anderson, J. R. (2024). Periodic tapping mechanisms of skill learning in a fast-paced video game. *Journal of Experimental Psychology: Human Perception and Performance, 50*(1), 39.
Grahn, J. A. (2012). Neural mechanisms of rhythm perception: Current findings and future perspectives. *Topics in Cognitive Science, 4*(4), 585–606.
Grahn, J. A., & Brett, M. (2007). Rhythm and beat perception in motor areas of the brain. *Journal of Cognitive Neuroscience, 19*(5), 893–906.
Grahn, J. A., & Rowe, J. B. (2013). Finding and feeling the musical beat: Striatal dissociations between detection and prediction of regularity. *Cerebral Cortex, 23*(4), 913–921.
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.
Honing, H. (2013). Structure and interpretation of rhythm in music. In D. Deutsch (Ed.), *The psychology of music* (3rd ed., pp. 369–404). Elsevier Academic Press.
Ito, M., & Doya, K. (2009). Validation of decision-making models and analysis of decision variables in the rat basal ganglia. *Journal of Neuroscience, 29*(31), 9861–9874.
Iversen, J. R., & Balasubramaniam, R. (2016). Synchronization and temporal processing. *Current Opinion in Behavioral Sciences, 8*, 175–180.
Ivry, R. B., & Hazeltine, R. E. (1995). Perception and production of temporal intervals across a range of durations: Evidence for a common timing mechanism. *Journal of Experimental Psychology: Human Perception and Performance, 21*(1), 3.
Jacoby, N., Tishby, N., Repp, B. H., Ahissar, M., & Keller, P. E. (2015). Parameter estimation of linear sensorimotor synchronization models: Phase correction, period correction, and ensemble synchronization. *Timing & Time Perception, 3*(1–2), 52–87.
Kang, P., Tobler, P. N., & Dayan, P. (2024). Bayesian reinforcement learning: A basic overview. *Neurobiology of Learning and Memory, 211*, Article 107924.
Keller, P. E., Novembre, G., & Hove, M. J. (2014). Rhythm in joint action: Psychological and neurophysiological mechanisms for real-time interpersonal coordination. *Philosophical Transactions of the Royal Society, B: Biological Sciences, 369*(1658), 20130394.
Kleiner, M. H., Brainard, D., & Pelli, D. (2007). *"What's new in Psychtoolbox-3?" perception, 36 ECVP abstract supplement*.
Knill, D. C., & Pouget, A. (2004). The Bayesian brain: The role of uncertainty in neural coding and computation. *Trends in Neurosciences, 27*(12), 712–719.
Koelsch, S., Vuust, P., & Friston, K. (2019). Predictive processes and the peculiar case of music. *Trends in Cognitive Sciences, 23*(1), 63–77.
Kruschke, J. K. (2008). Bayesian approaches to associative learning: From passive to active learning. *Learning & Behavior, 36*(3), 210–226.
Kuznetsova, A., Brockhoff, P. B., & Christensen, R. H. (2017). lmerTest package: Tests in linear mixed effects models. *Journal of Statistical Software, 82*, 1–26.
Laje, R., Cheng, K., & Buonomano, D. V. (2011). Learning of temporal motor patterns: An analysis of continuous versus reset timing. *Frontiers in Integrative Neuroscience, 5*, 61.
Large, E. W., Fink, P., & Kelso, S. J. (2002). Tracking simple and complex sequences. *Psychological Research, 66*, 3–17.
Lenth, R. (2020). Emmeans: Estimated marginal means, aka least-squares means. R package version 1.5.0.. https://CRAN.R-project.org/package=emmeans.
Manto, M., Bower, J. M., Conforto, A. B., Delgado-García, J. M., Da Guarda, S. N. F., Gerwig, M., & Timmann, D. (2012). Consensus paper: Roles of the cerebellum in motor control—the diversity of ideas on cerebellar involvement in movement. *The Cerebellum, 11*(2), 457–487.
Master, S. L., Eckstein, M. K., Gotlieb, N., Dahl, R., Wilbrecht, L., & Collins, A. G. (2020). Disentangling the systems contributing to changes in learning during adolescence. *Developmental Cognitive Neuroscience, 41*, Article 100732.
Nguyen, K. P., & Person, A. L. (2025). Cerebellar circuit computations for predictive motor control. *Nature Reviews Neuroscience*, 1–16.
Pearce, J. M., & Bouton, M. E. (2001). Theories of associative learning in animals. *Annual Review of Psychology, 52*(1), 111–139.
Pelli, D. G. (1997). The VideoToolbox software for visual psychophysics: Transforming numbers into movies. *Spatial Vision, 10*, 437–442.
Pérez, O., Delle Monache, S., Lacquaniti, F., Bosco, G., & Merchant, H. (2023). Rhythmic tapping to a moving beat motion kinematics overrules natural gravity. *Iscience, 26*(9).
Piray, P., & Daw, N. D. (2020). A simple model for learning in volatile environments. *PLoS Computational Biology, 16*(7), Article e1007963.
Popa, L. S., Streng, M. L., Hewitt, A. L., & Ebner, T. J. (2016). The errors of our ways: Understanding error representations in cerebellar-dependent motor learning. *The Cerebellum, 15*(2), 93–103.
Povel, D. J., & Essens, P. (1985). Perception of temporal patterns. *Music Perception, 2*(4), 411–440.
R Core Team. (2025). *R: A language and environment for statistical computing*. Vienna, Austria: R Foundation for Statistical Computing. https://www.R-project.org/.
Repp, B. H. (2005). Sensorimotor synchronization: A review of the tapping literature. *Psychonomic Bulletin & Review, 12*(6), 969–992.
Repp, B. H., & Keller, P. E. (2004). Adaptation to tempo changes in sensorimotor synchronization: Effects of intention, attention, and awareness. *The Quarterly Journal of Experimental Psychology Section A, 57*(3), 499–521.
Repp, B. H., Keller, P. E., & Jacoby, N. (2012). Quantifying phase correction in sensorimotor synchronization: Empirical comparison of three paradigms. *Acta Psychologica, 139*(2), 281–290.
Repp, B. H., & Su, Y. H. (2013). Sensorimotor synchronization: A review of recent research (2006–2012). *Psychonomic Bulletin & Review, 20*(3), 403–452.
Rescorla, R. A., & Wagner, A. R. (1972). A theory of Pavlovian conditioning: Variations in the effectiveness of reinforcement and nonreinforcement. In *, 2. Classical Conditioning II: Current Research and Theory* (pp. 64–99). New York, NY: Appleton-Century-Crofts.
Sakoe, H., & Chiba, S. (1978). Dynamic programming algorithm optimization for spoken word recognition. *IEEE Transactions on Acoustics, Speech, and Signal Processing, 26*(1), 43–49.
Savitzky, A., & Golay, M. J. (1964). Smoothing and differentiation of data by simplified least squares procedures. *Analytical Chemistry, 36*(8), 1627–1639.
Schulze, H. H., Cordes, A., & Vorberg, D. (2005). Keeping synchrony while tempo changes: Accelerando and ritardando. *Music Perception, 22*(3), 461–477.
Schulze, H. H., & Vorberg, D. (2002). Linear phase correction models for synchronization: Parameter identification and estimation of parameters. *Brain and Cognition, 48*(1), 80–97.
Soto, F. A., Vogel, E. H., Uribe-Bahamonde, Y. E., & Perez, O. D. (2023). Why is the Rescorla-Wagner model so influential? *Neurobiology of Learning and Memory, 204*, Article 107794.
Stan Development Team. (2023). Stan Modeling Language Users Guide and Reference Manual, version 2.32.2. https://mc-stan.org.
Torre, K., Varlet, M., & Marmelat, V. (2013). Predicting the biological variability of environmental rhythms: Weak or strong anticipation for sensorimotor synchronization? *Brain and Cognition, 83*(3), 342–350.
Vorberg, D., & Schulze, H. H. (2002). Linear phase-correction in synchronization: Predictions, parameter estimation, and simulations. *Journal of Mathematical Psychology, 46*(1), 56–87.
Vuust, P., Ostergaard, L., Pallesen, K. J., Bailey, C., & Roepstorff, A. (2009). Predictive coding of music–brain responses to rhythmic incongruity. *cortex, 45*(1), 80–92.
Vuust, P., & Witek, M. A. (2014). Rhythmic complexity and predictive coding: A novel approach to modeling rhythm and meter perception in music. *Frontiers in Psychology, 5*, 1111.
Wing, A. M., & Kristofferson, A. B. (1973). Response delays and the timing of discrete motor responses. *Perception & Psychophysics, 14*(1), 5–12.
Xia, L., Master, S. L., Eckstein, M. K., Baribault, B., Dahl, R. E., Wilbrecht, L., & Collins, A. G. E. (2021). Modeling changes in probabilistic reinforcement learning during adolescence. *PLoS Computational Biology, 17*(7), Article e1008524.