

Latent structure in random sequences drives neural learning toward a rational bias

Yanlong Sun^{a,1}, Randall C. O'Reilly^b, Rajan Bhattacharyya^c, Jack W. Smith^a, Xun Liu^d, and Hongbin Wang^{a,1}

^aCenter for Biomedical Informatics, Texas A&M University Health Science Center, Houston, TX 77030; ^bDepartment of Psychology and Neuroscience, University of Colorado, Boulder, CO 80309; ^cCenter for Neural and Emergent Systems, Information and System Sciences Lab, HRL Laboratories LLC, Malibu, CA 90265; and ^dKey Laboratory of Behavioral Science, Institute of Psychology, Chinese Academy of Sciences, Beijing, China 100101

Edited by Michael I. Posner, University of Oregon, Eugene, OR, and approved February 4, 2015 (received for review November 18, 2014)

People generally fail to produce random sequences by overusing alternating patterns and avoiding repeating ones—the gambler's fallacy bias. We can explain the neural basis of this bias in terms of a biologically motivated neural model that learns from errors in predicting what will happen next. Through mere exposure to random sequences over time, the model naturally develops a representation that is biased toward alternation, because of its sensitivity to some surprisingly rich statistical structure that emerges in these random sequences. Furthermore, the model directly produces the best-fitting bias-gain parameter for an existing Bayesian model, by which we obtain an accurate fit to the human data in random sequence production. These results show that our seemingly irrational, biased view of randomness can be understood instead as the perfectly reasonable response of an effective learning mechanism to subtle statistical structure embedded in random sequences.

gambler's fallacy | waiting time | neural network | temporal integration | Bayesian inference

People are prone to search for patterns in sequences of events, even when the sequences are completely random. In a famous game of roulette at the Monte Carlo casino in 1913, black repeated a record 26 times—people began extreme betting on red after about 15 repetitions (1). The gambler's fallacy—a belief that chance is a self-correcting process where a deviation in one direction would induce a deviation in the opposite direction—has been deemed a misperception of random sequences (2). For decades, this fallacy is thought to have originated from the “representativeness bias,” in which a sequence of events generated by a random process is expected to represent the essential characteristics of that process even when the sequence is short (3).

However, there is a surprising amount of systematic structure lurking within random sequences. For example, in the classic case of tossing a fair coin, where the probability of each outcome (heads or tails) is exactly 0.5 on every single trial, one would naturally assume that there is no possibility for some kind of interesting structure to emerge, given such a simple and desolate form of randomness. And yet, if one records the average amount of time for a pattern to first occur in a sequence (i.e., the waiting time statistic), it is significantly longer for a repetition (head–head HH or tail–tail TT, six tosses) than for an alternation (HT or TH, four tosses). This is despite the fact that on average, repetitions and alternations are equally probable (occurring once in every four tosses, i.e., the same mean time statistic). For both of these facts to be true, it must be that repetitions are more bunched together over time—they come in bursts, with greater spacing between, compared with alternations. Intuitively, this difference comes from the fact that repetitions can build upon each other (e.g., sequence HHH contains two instances of HH), whereas alternations cannot. Statistically, the mean time and waiting time delineate the mean and variance in the distribution of the interarrival times of patterns, respectively (4). Despite the same frequency of occurrence (i.e., the same mean), alternations are more evenly distributed over time than repetitions (i.e., different variances). Another source of insight comes from the transition graph (Fig. 14), which reveals a structural asymmetry in the process

of fair coin tossing. For example, when the process has the same chance to visit any of the states, the minimum number of transitions it takes to leave and then revisit a repetition state is longer than that for an alternation state. Let p_A denote the probability of alternation between any two consecutive trials; despite the same mean time at $p_A = 1/2$, repetitions will have longer waiting times than alternations as long as $p_A > 1/3$ (Fig. 1B). (See *SI Text* for the calculation of mean time and waiting time statistics.)

Is this latent structure of waiting time just a strange mathematical curiosity or could it possibly have deep implications for our cognitive-level perceptions of randomness? It has been speculated that the systematic bias in human randomness perception such as the gambler's fallacy might be due to the greater variance in the interarrival times or the “delayed” waiting time for repetition patterns (4, 5). Here, we show that a neural model based on a detailed biological understanding of the way the neocortex integrates information over time when processing sequences of events (6, 7) is naturally sensitive to both the mean time and waiting time statistics. Indeed, its behavior is explained by a simple averaging of the influences of both of these statistics, and this behavior emerges in the model over a wide range of parameters. Furthermore, this averaging dynamic directly produces the best-fitting bias-gain parameter for an existing Bayesian model of randomness judgments (8), which was previously an unexplained free parameter and obtained only through parameter fitting. We also show that we can extend this Bayesian model to better fit the full range of human data by including a higher-order pattern statistic, and the neurally derived bias-gain parameter still provides the best fit to the human data in the augmented model. Overall, our model provides a neural

Significance

The human mind has a unique capacity to find order in chaos. The way the neocortex integrates information over time enables the mind to capture rich statistical structures embedded in random sequences. We show that a biologically motivated neural network model reacts to not only how often a pattern occurs (mean time) but also when a pattern is first encountered (waiting time). This behavior naturally produces the alternation bias in the gambler's fallacy and provides a neural grounding for the Bayesian models of human behavior in randomness judgments. Our findings support a rational account for human probabilistic reasoning and a unifying perspective that connects the implicit learning without instruction with the generalization under structured and expressive rules.

Author contributions: Y.S., R.C.O., and H.W. designed research; Y.S. and H.W. performed research; Y.S., R.C.O., J.W.S., and H.W. contributed new reagents/analytic tools; Y.S. and H.W. analyzed data; and Y.S., R.C.O., R.B., J.W.S., X.L., and H.W. wrote the paper.

The authors declare no conflict of interest.

This article is a PNAS Direct Submission.

Freely available online through the PNAS open access option.

¹To whom correspondence may be addressed. Email: ysun@tamhsc.edu or hwang@tamhsc.edu.

This article contains supporting information online at www.pnas.org/lookup/suppl/doi:10.1073/pnas.1422036112/-DCSupplemental.




Fig. 1. Time of patterns described by the probability of alternation between consecutive trials (p_A). (A) Transitions between patterns of length 2. At $p_A = 1/2$, the process has the same chance to visit either a repetition state (HH or TT) or an alternation state (HT or TH). However, it takes a minimum of three transitions for the process to leave and then revisit a repetition state (e.g., HH \rightarrow HT \rightarrow TH \rightarrow HH), but only two for an alternation state (e.g., HT \rightarrow TH \rightarrow HT). (B) Equilibria by p_A values. A repetition (R) and an alternation (A) have the same mean time $E[T_R] = E[T_A]$ at $p_A = 1/2$, the same waiting time $E[T_R^*] = E[T_A^*]$ at $p_A = 1/3$, and the same sum $E[T_R] + E[T_R^*] = E[T_A] + E[T_A^*]$ at $p_A = 3/7$.

grounding for the pervasive gambler's fallacy bias in human judgments of random processes, where people systematically discount repetitions and emphasize alternations (9, 10).

Neural Model of Temporal Integration

Our neural model is extremely simple (Fig. 2A). It consists of a sensory input layer that scans nonoverlapping binary digits of H vs. T and an internal prediction layer that attempts to predict the next input, while the prior inputs in the sequence are encoded in the temporal context. This model is based on a biologically motivated computational framework that has been used to explain the neural basis of cognition in a wide range of different domains (6), with the benefit of integrating prior temporal context information according to the properties of the deep neocortical neurons (layers 5b and 6) (7).

Our main hypothesis is that the predictive learning and temporal integration properties of this model, which reflect corresponding features of the neocortex, will produce representations that incorporate both the waiting time and mean time statistics of the input sequences (despite the inability of the model to accurately predict the next input in these random sequences). In other words, we predict a systematic interaction between these basic learning mechanisms and the surprisingly rich statistical structure of the input. This is a principled prediction based on the well-established sensitivity of these kinds of neural learning mechanisms to the statistical structure of inputs (e.g., ref. 11), and extensive parameter exploration demonstrates that our results hold across a wide range of parameters and that the model's behavior is systematically affected by certain parameters in sensible ways (SI Text). Thus, despite the emergent nature of learning in our model, it nevertheless reflects systematic behavior and is not the result of a random parameter-fitting exercise. Moreover, we show below that the model's behavior can be largely explained by a simple equation as a function of the mean time and waiting time statistics, further establishing the systematicity of the model's behavior and also establishing a direct connection to more abstract Bayesian models as elaborated subsequently.

The model was trained with binary sequences generated at various levels of the probability of alternation (p_A), each sequence consisting of 10,000 coin tosses (although learning occurred quickly within a few hundred trials). Crucially, learning was concerned with only reconstructing the input sequence but not pattern discrimination, as no teaching signals were provided regarding the underlying p_A values and pattern time statistics. After training, the model was tested with a sequence of 1,000 tosses generated at the same p_A level as in the training sequence.

We then decoded these sequence representations through a reverse correlation technique. Based on the sensitivity of the unit activations to the temporal patterns of length 2, we classified and then counted the number of the units on the internal prediction layer as either repetition detectors (R, being sensitive to either HH or TT) or alternation detectors (A, being sensitive to either HT or TH). (See SI Text for the method of detector classification.)

Most intriguingly, at $p_A = 1/2$ (i.e., independent tosses of a fair coin), the model produced a ratio of $R/A \approx 0.70$ —despite the equal frequency of pattern occurrences, repetition detectors were significantly less likely than alternation detectors. Such alternation bias is in the same direction of the representativeness bias, where people perceive alternation patterns as more representative of a random process than repetition patterns (2, 3). Effectively, this result demonstrates the gambler's fallacy emerging naturally as a consequence of the alternation bias, due to the model's sensitivity to the waiting time advantage of alternations compared with repetitions.

We then used the R/A ratio to compute the subjective probability of alternation, p'_A , as the model's internal representation of its actually experienced p_A . With $R/A \approx 0.70$, we have

$$p'_A = \frac{A}{R+A} = \frac{1}{1+R/A} \approx 0.59. \quad [1]$$

This p'_A value is consistent with the empirical findings on subjective randomness. From a comprehensive review of the studies on random sequence perception and generation, it was found that the subjective probability of alternation was around $0.58 \sim 0.63$ (9).

To further characterize the nature of the alternation bias, we systematically varied the probability of alternation (p_A) in generating the training sequence (i.e., departures from tossing a fair coin independently) and then measured the effects on the R/A ratio. We found a smooth curve, where the R/A ratio increased (more repetition detectors) as p_A decreased (less frequent occurrences of alternations). At $p_A = 3/7$, the model reached an equilibrium point with equal numbers of repetition and alternation detectors, $R/A = 1$ (Fig. 2B). That is, alternations have to be this much less frequent (i.e., greater mean time) to cancel out their waiting time advantage. This corresponds exactly to the equilibrium point where repetitions and alternations have the same sum of mean and waiting times (Fig. 1B).

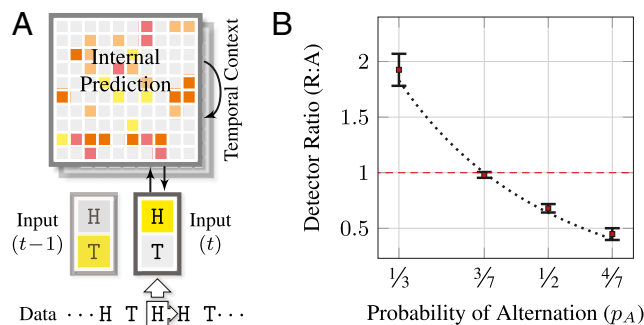


Fig. 2. Neural model of temporal integration to capture the statistics of pattern times in random sequences. (A) Architecture of the neural model. A single sensory input layer scans through a sequence of binary digits one digit at a time (input at $t - 1$ is for illustration only). An internal prediction layer, with bidirectional connections from the input layer and its own temporal context representation, attempts to predict the next input. (B) Neural model behavior depicted by the ratio between repetition and alternation detectors in response to the actual probability of alternation (p_A) in the input sequences. At $p_A = 1/2$, the model showed $R/A \approx 0.70$ (i.e., fewer repetition detectors than alternation detectors). Error bars (\pm SEM) represent the variability of model predictions. The dotted line is the squared total time ratio between alternation and repetition patterns (Eq. 2).

Overall, the model's behavior can be mostly replicated by a simple equation that averages the effects of the mean time and waiting time statistics (dotted line in Fig. 2B),

$$\frac{R}{A} \approx \left(\frac{E[T_A] + E[T_A^*]}{E[T_R] + E[T_R^*]} \right)^2, \quad [2]$$

where $E[T]$ is the mean time, $E[T^*]$ is the waiting time, and subscripts R and A represent repetitions and alternations, respectively. This establishes a clear higher-level explanation for the emergent behavior of the model, allowing us to summarize its behavior as simply averaging the effects of these two relevant statistics over the random sequences.

Bayesian Models of Random Sequence Production

A unifying perspective on human statistical learning requires bridging the gap between the implicit learning without instruction and the generalization of the learned patterns under structured and expressive rules (12). On the one hand, our neural model shows that through mere exposure to a set of input stimuli, a systematic bias was developed toward the alternation pattern in random sequences. On the other hand, recent Bayesian accounts for probabilistic learning suggest that the human mind performs rational inferences at both neural and behavioral levels (13, 14). Thus, we asked whether it was possible to relate the emergent behavior of the neural model to an existing Bayesian model of randomness judgments, specifically whether we could demonstrate a quantitative connection between the bias for local patterns at the neural level and the behavior of generating longer random sequences governed by the rules of Bayesian inference.

Let $f(H, T)$ denote the degree of the belief that a sequence of coin tosses consisting of H heads and T tails is generated by a fair coin, where the probability of heads in any single toss is $p = 1/2$. By Bayes' theorem, assuming a uniform prior distribution $p \in [0, 1]$, $f(H, T)$ can be formulated as the posterior probability density,

$$f(H, T) = \frac{2^{-(H+T)}}{\int_0^1 p^H (1-p)^T dp} = 2^{-(H+T)} (H+T+1) \binom{H+T}{H}. \quad [3]$$

Because of the binomial coefficient $\binom{H+T}{H}$, Eq. 3 is maximized when $H = T$. That is, governed by the belief function $f(H, T)$, the optimal solution to generating a random sequence is to always seek a balance between the numbers of heads and tails (10). Based on this belief function, Griffiths and Tenenbaum (8) proposed a Bayesian model of random sequence production. They first defined a likelihood function, L_k , to represent the local representativeness that choosing a head instead of a tail as the outcome of the k th toss would result in a more random sequence,

$$L_k = \sum_{i=1}^{k-1} \log f(H_i + 1, T_i) - \log f(H_i, T_i + 1) = \log \prod_{i=1}^{k-1} \left(\frac{T_i + 1}{H_i + 1} \right), \quad k \geq 2, \quad [4]$$

where H_i and T_i were, respectively, the numbers of heads and tails counting back i steps in the sequence. Then, with a free parameter (λ) to scale the contribution of L_k , the probability of choosing a head at each response (R_k) was obtained by a logistic function:

$$P(R_k = H) = \frac{1}{1 + e^{-\lambda L_k}}. \quad [5]$$

This Bayesian model was then used to fit a massive database from the "Zenith radio experiment," where 20,099 participants attempted to produce five random binary symbols one at a time (15). It was found that a λ value of around 0.60 produced the optimal fit to 15 of the 16 data points from the human data (Fig. 3A).

The data point that Griffiths and Tenenbaum (8) did not predict well was the sequence HTHTH, which was judged by people to not be a very good random sequence, but Eq. 5 ranked it highly. It seems that in generating random sequences, people were seeking a balance not only between the heads and tails but also between higher-order pattern events (e.g., alternation itself is repeated four times in HTHTH). We can add this mechanism into the Bayesian model with an additional term, M_k ,

$$M_k = \log \left(\frac{O_T + 1}{O_H + 1} \right), \quad k \geq 3, \quad [6]$$

where M_k performs a similar function to that of L_k , except being based on the numbers of the second-order pattern events, O_H and O_T (either alternation or repetition, depending the choice at R_{k-1}).

Applying the same scaling factor λ to both L_k and M_k , Eq. 5 becomes

$$P(R_k = H) = \frac{1}{1 + e^{-\lambda(L_k + M_k)}}. \quad [7]$$

This augmented model now produces an excellent fit to the full set of sequence data points (Pearson's $R^2 \approx 0.86$), with $\lambda \approx 0.51$ as the optimal parameter value (Fig. 3A). In addition to improving the prediction on the sequence HTHTH, Eq. 7 also consistently makes better predictions than Eq. 5 on other data points, and the same value $\lambda \approx 0.51$ also produces the best fit to the partial dataset excluding the sequence HTHTH (Pearson's $R^2 \approx 0.89$) (Fig. 3B).

The Bayesian model provides a formalization of the representativeness heuristic (2, 3, 8, 14). It captures the idea that when generating random sequences, people are seeking a balance between heads and tails and between repetitions and alternations not only in the global sequence, but also in the local subsequences (10). Apparently, the extent to which this balance is adjusted is determined by the free scaling parameter λ in Eqs. 5 and 7. However, beyond parameter fitting, the Bayesian model does not have any independent basis for specifying this parameter.

In the light of the neural model's behavior (i.e., the alternation bias in Eq. 1), we predict that the scaling parameter λ should have originated naturally from people's actual experiences of random sequences in the learning environment. Specifically, we can deduce from either Eq. 5 or Eq. 7 that λ actually serves as a bias-gain parameter that modulates the strength of the alternation bias:

$$p'_A = P(R_2 = T | R_1 = H) = \frac{1}{1 + 2^{-\lambda}}. \quad [8]$$

When $\lambda = 0$, Eq. 8 produces unbiased responses with the subjective probability of alternation $p'_A = 1/2$, corresponding to the process of independent coin tossing; and higher values $\lambda > 0$ produce an increasing alternation bias with $p'_A > 1/2$. In other words, $\lambda > 0$ corresponds to the tendency of avoiding repetitions, which applies to both the first- and second-order events (Eq. 7).

