

A Minimal Neural Network Model of The Gambler's Fallacy

Yanlong Sun (ysun@tamhsc.edu)

Hongbin Wang (hwang@tamhsc.edu)

Center for Biomedical Informatics, Texas A&M University Health Science Center
Houston, TX 77030 USA

Abstract

The gambler's fallacy has been a notorious showcase of human irrationality in probabilistic reasoning. Recent studies suggest the neural basis of this fallacy might have originated from the predictive learning by neuron populations over the latent temporal structures of random sequences, particularly due to the statistics of pattern times and the precedence odds between patterns. Here we present a biologically-motivated minimal neural network model with only eight neurons. Through unsupervised training, the model naturally develops a bias toward alternation patterns over repetition patterns, even when both patterns are equally likely presented to the model. Our analyses suggest that the way the neocortex integrates information over time makes the neuron populations not only sensitive to the frequency signals but also relational structures embedded over time. Moreover, we offer an explanation for how higher-level cognitive biases may have an early start at the level of sensory processing.

Keywords: gambler's fallacy; alternation bias; waiting time; temporal integration; predictive learning.

Introduction

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 a notorious showcase of human irrationality in probabilistic reasoning (Tversky & Kahneman, 1974). For decades, this fallacy is thought to have originated from a cognitive bias called the “representativeness heuristic”, which is attributed to the belief of the “law of small numbers” that small samples are highly representative of the populations from which they are drawn (Gilovich, Vallone, & Tversky, 1985; Tversky & Kahneman, 1974).

Recent development in neuroscience and computational models suggests that the human mind develops structured probabilistic representations about the world and performs near-optimal Bayesian inferences (Pouget, Beck, Ma, & Latham, 2013; Tenenbaum, Kemp, Griffiths, & Goodman, 2011). For example, representativeness has been defined with a Bayesian belief-updating structure in which different hypotheses are evaluated based on different sets of the input data (Gigerenzer & Hoffrage, 1995; Griffiths & Tenenbaum, 2001). However, it remains elusive how the structured hypothesis space has originated in the first place, and how cognitive biases can arise from normative probabilistic models.

On the topic of randomness perception, there has been a growing speculation that people's intuition about random process, also known as the subjective randomness, is biased by the statistical structures in the learning environment (Budescu, 1987; Falk & Konold, 1997; Hahn & Warren, 2009; Lopes & Oden, 1987; Nickerson, 2002; Oppenheimer & Monin, 2009; Oskarsson, Van Boven, McClelland, & Hastie, 2009; Sun,

Tweney, & Wang, 2010). Particularly, we have argued that without a predefined hypothesis structure, biases underpinning the gambler's fallacy can emerge by simply capturing the temporal relations between patterns as a random process unfolds over time (Sun & Wang, 2010a, 2010b, 2012, 2015). With a biologically realistic simple recurrent model that learns to re-encode sequential binary data through unsupervised learning, we show that dissociation of random patterns can naturally emerge as the consequence of inhibitory competition between overlapped representations (Sun et al., 2015). Our findings indicate that cognitive biases in overt behavior can emerge early and locally at the level of sensory processing, and neurons' sensitivity to the temporal structures in the learning environment is the key in bridging the gap between neurons and behavior.

In the following, we first introduce some basic normative measures on the time of random patterns. Then, based on the neural model we reported earlier, we present a minimal neural network model with only eight units. We will show that this minimal model can mostly replicate our previous findings and provide new insights regarding the neural encodings of sequential patterns.

Temporal Distance between Patterns

In sequences generated by a random process, there can be fundamentally different types of statistical structures regarding *how often* a pattern occurs and *when* a pattern is to occur. Our previous works have been focusing on the distinction between the mean time statistic that measures how often the pattern occurs in a global sequence, and the waiting time statistic that measures when a pattern will first occur since the beginning of the observation. Here we introduce a more compact yet more comprehensive framework that incorporates not only both types of statistics for individual patterns but also the statistics depicting the relational structures between different patterns.

To compute the temporal distances between different patterns with different initial states, we use the first-order dependent Markov chains parameterized by the *probability of alternation* (p_A) between consecutive trials and the corresponding generating functions (Figure 1).¹ Define $E[T_{ji}]$ as the expected number of transitions from the initial state i until the

¹The method of generating functions by Markov chains also applies to independent Bernoulli trials parameterized the probabilities of single elements (e.g., the probability of heads or tails), and it also generates higher-moment statistics such as variance (e.g., Sun & Wang, 2015). Here we only present the main results and the exact generating functions are omitted.




Figure 1: Markov chains for generating the *waiting times* $E[T_{HH|\emptyset}]$ and $E[T_{HT|\emptyset}]$ given the initial empty state $\emptyset$ (Figure A), and the *additional times* $E[T_{HH|H}]$, $E[T_{HT|H}]$, and $E[T_{T|H}]$, given the same initial state H (Figure B). In each chain, states $S_{j|i}$ represent all possible sequences that start from the pattern i ($i = \emptyset$ means starting anew) and end with the first arrival of the pattern j . States M_k represent all possible sequences that end with the pattern k but do not contain the expected pattern j . Transitions between nonempty states are labeled as either repetition (R) or alternation (A). Figure C: Pairwise precedence odds between patterns when the probability of alternation $p_A = 1/2$, for example, the odds are 3 to 1 that one is to first encounter TH than to first encounter HH.

first arrival of the pattern j . When the initial state is empty $i = \emptyset$ (i.e., the counting process starts anew), $E[T_{j|\emptyset}]$ is referred to as the *waiting time* of pattern j . For example, from Figure 1A, the waiting times for the patterns HH and HT are respectively,

$$\begin{aligned} E[T_{HH|\emptyset}] &= 1 + \frac{1}{2p_A} + \frac{2}{1-p_A}, \\ E[T_{HT|\emptyset}] &= 1 + \frac{1}{2p_A} + \frac{1}{p_A}. \end{aligned} \quad (1)$$

When $p_A = 1/2$ (namely, independent Bernoulli trials with a fair coin where repetitions and alternations are equally likely), we have $E[T_{HH|\emptyset}] = 6$ and $E[T_{HT|\emptyset}] = 4$.

When the initial state is not empty, $E[T_{j|i}]$ is referred to as the *additional time* for pattern j given the initial state i . For example, from Figure 1B, we have

$$E[T_{HH|H}] = \frac{2}{1-p_A}, \quad E[T_{HT|H}] = E[T_{T|H}] = \frac{1}{p_A}. \quad (2)$$

At $p_A = 1/2$, we have $E[T_{HH|H}] = 4$ and $E[T_{HT|H}] = E[T_{T|H}] = 2$.

When the initial state i is exactly the desired pattern j , $E[T_{j|i}]$ denotes the expected number of transitions between any two consecutive occurrences of the pattern j , and is referred to as the *mean time* of pattern j . Since the first-order Markov chain is memoryless between consecutive transitions, we have relations such as $E[T_{HH|HH}] = E[T_{HH|H}]$, and $E[T_{HT|HT}] = E[T_{HT|T}]$. Therefore, the mean times for the patterns HH and HT are respectively,

$$E[T_{HH|H}] = \frac{2}{1-p_A}, \quad E[T_{HT|T}] = \frac{2}{p_A}. \quad (3)$$

At $p_A = 1/2$, we have $E[T_{HH|H}] = E[T_{HT|T}] = 4$. The inverse of mean time is *frequency*. For example, $E[T_{HH|H}] = 4$ means that we expect to see the pattern HH once in every 4 tosses of a fair coin.

Among these different measures, the most striking distinction is that at $p_A = 1/2$, we have $E[T_{HH|H}] = 4$ but $E[T_{HT|H}] = 2$, in spite of the fact that given an H, the next digit is equally likely to be either an H or a T. This is because a reoccurrence of the pattern HH can “reuse” the ending elements from its previous occurrence, but a reoccurrence of the pattern HT must always start anew. This statistical property of faster transition times when starting anew is known as *new better than used* (NBU) (Ross, 2007). Essentially, the NBU property is due to the *overlap* between a pattern and a shifted copy of itself or between different patterns. For example, as shown in Figure 1B, pattern HH overlaps with its shifted copy by one element H, but pattern HT does not. Then, towards the destination state $S_{HH|H}$, anytime things go astray (i.e., the process ends in state M_T), the waiting for HH has to start all over. In contrast, the waiting for HT is always on average two flip away from the state M_H . As a result, the transition to $S_{HH|H}$ is “delayed” than the transition to $S_{HT|H}$. This overlap also explains the pairwise precedence relation shown in Figure 1C. For example, in the competition between the patterns HH and TH, the former reuses the last element of the latter but the latter starts anew. As a result, if we toss a fair coin repeatedly (i.e., $p_A = 1/2$), the odds are 3 to 1 that we first encounter TH than first encounter HH.

A Neural Model of Temporal Integration

The NBU property of pattern times has fundamental implications in neural encoding of pattern events. As shown in Figure 1, different measures of waiting time, additional time, mean time and pairwise precedence odds are all due to the overlap between temporal patterns. Recent developments in neuroscience and computational models suggest that neural encodings of events and values are always overlapped, and it is the encoding of neural populations that give rise to higher-level and more abstract representations (Adolphs, 2015; Dehaene & Brannon, 2010; Pouget et al., 2013). In the domain of temporal integration, probabilistic encoding must consider the overlap between representations at different times, namely, recurrent processing (Elman, 1990). Then, we would immediately conjecture that via merely encoding the random sequences that unfold over time, populations of neurons would naturally capture the temporal structures depicted by the pattern times statistics.

prediction layer has only four units without explicit temporal context representation. Also different from the bigger model where the initial status of detectors was set by random weights, the detectors in the eight-unit model are initially set by distinctive projection weights from the input units. For example, detectors HH and HT receive the same projection weights from the context input units, but detector HH receives a stronger weight from the online input unit H, and detector HT receives a stronger weight from the online input unit T. In other words, given the same initial state H, detector HH tends to predict a repetition and detector HT tends to predict an alternation.

Crucially, the controlled weights allow a more precisely controlled experiment by eliminating variations produced by random weights. In the bigger model, different detectors would “naturally” emerge by a random initialization of weights. However, this may produce a disparity in the number of detector types at the initial stage (e.g., more HT detectors than HH detectors), and such a disparity has to be accounted for by averaging multiple simulations with different random initializations. This disparity is eliminated in the eight-unit model, such that the model as a whole is initially unbiased toward any of the four patterns. (We have implemented different sets of controlled weights, e.g., 0.2 versus 0.8, or 0.4 versus 0.6, and find that they all produce the same results.)




Figure 4: Pattern dissociation at different levels of the probability of alternation p_A after training. The repetition detectors HH and TT only showed higher activations when the model was trained with sequences generated by $p_A < 1/3$. Box plots represent distribution quantile.

The eight-unit model was trained and tested in the same way as the bigger model, and the only difference is the analyses of the test results. Instead of counting the number of detectors, we directly measure the activations of each detector given different input patterns. Figure 4 shows the main result. We first notice that after being trained with truly random sequences (i.e., $p_A = 1/2$ in independent fair coin tossing), the aver-

aged activations of detectors were significantly lower when the current inputs were repetition patterns (HH or TT) than alternation patterns (HT or TH). That is, in spite of the initially unbiased representations of all patterns and the equal training frequency (i.e., the mean time is the same for all patterns at $p_A = 1/2$), repetition patterns eventually were significantly under-represented than alternation patterns.




Figure 5: The updating trajectories of projection weights from context input units to detector units HH and TH during the training phase. The initial weight values are marked by the circles at the first trial. For example, at $p_A = 1/2$, the unit initially designated as the HH detector became an TH detector after approximately 200 trials, as all its weight values from the context units switched to the opposite side of 0.5, whereas the unit initially designated as the TH detector remained stable. The same trend was also observed between TT and HT detectors. The projection weights from the online input units remained about the same thus are not plotted here.

To locate the source of the alternation bias, we found that during the training phase, the projection weights from the context input units to each detector underwent a dramatic remapping (Figure 5), whereas the projection weights from the online input units remained about the same.

Specifically, at $p_A = 1/2$, the detector unit HH initially received a weight of 0.7 from the context input unit H and a weight of 0.3 from the context input unit T (hence its initially designated detecting status). After about 200 trials, these two weights switched to the opposite sides of 0.5, effectively “switching” the HH detector into a TH detector (but not a HT detector). Similarly, the TT detector switched to an HT detector (but not a TH detector). The directions of these switches corresponded exactly to the pairwise precedence relationship depicted in Figure 1C. At $p_A = 2/3$, the switch was even more obvious. At $p_A = 1/6$, it was the alternation detectors’ turn to

be under-represented and switched to the repetition detectors. Finally at $p_A = 1/3$, all projection weights from the context units approached then stabilized around 0.5, so that the model eventually learned to be indifferent to the contextual information, resulting in unbiased activations for all patterns as shown in Figure 4.

Moreover, we also tested models with only a subset of particular detectors (e.g., HH versus TH only). When $p_A = 1/2$, regardless of the initial pattern preference set by the projection weights, the model would eventually react indifferently to all patterns. This result indicates that in order to capture different pattern time statistics or the pairwise precedence odds between patterns of length two (Figure 1), the inhibitory competition between at least four types of detectors is required. When we tested models with only online input units, the model showed the same indifference at $p_A = 1/2$. This indicates that predictive learning, namely, predicting what will happen next based on the historical context, is critical in producing the alternation bias. Together, these observations confirmed our hypothesis that this eight-unit model is a minimal model to produce the alternation bias in the gambler's fallacy.

In comparison, the alternation bias exhibited by the minimal eight-unit model in Figure 3 is in the same direction as that exhibited by the bigger model in Figure 2. However, the equilibrium point (the p_A level where the model was indifferent to all patterns) was different (compare Figure 2 with Figure 4). Specifically, the equilibrium point was $p_A = 1/3$ for the eight-unit model but $p_A = 3/7$ for the bigger model. This indicates that the minimal model was more sensitive to the waiting time (delay) than to the mean time (frequency) of pattern occurrences, because by Equations 1 and 3, all patterns have the same waiting time but different mean times at $p_A = 1/3$,

$$E[T_{HH|\emptyset}] = E[T_{HT|\emptyset}] = 11/2, \quad E[T_{HH|H}] = 3, \quad E[T_{HT|T}] = 6.$$

In contrast, the bigger model was more “balanced” toward both statistics, because at $p_A = 3/7$,

$$E[T_{HH|\emptyset}] + E[T_{HH|H}] = E[T_{HT|\emptyset}] + E[T_{HT|T}] = 55/6.$$

One particular reason for such difference is that the competition would be stronger among fewer detectors due to the homeostatic mechanism implemented in the network. This mechanism keeps individual neurons from firing too much or too little over time, which is essentially a normalization mechanism in self-organizing learning at long time scales (Bienenstock, Cooper, & Munro, 1982; Cooper, 2000; Hebb, 1949; O'Reilly, Munakata, Frank, Hazy, & Contributors, 2012). As a consequence, the bigger model with more neurons would be more likely to maintain diversity in the specialization of neurons thus its equilibrium point could be determined by both waiting time and mean time statistics.

Conclusion

Overall, our results from both models suggest that pattern dissociation can naturally emerge from temporal reconstructions of the input data. Particularly with the minimal eight-neuron

model, detector neurons “reoriented attention” to the past information (i.e., remapping the projection weights from the context units), and the driving force behind such reorientation was more of the waiting time rather than of the mean time of patterns. As for the model with more neurons we reported earlier, the specialization of neurons would be more diversified thus would enable the model to develop sensitivity to both types of pattern time statistics.

The observation that both models exhibited the alternation bias is consistent with the representativeness bias underpinning the gambler's fallacy (Gilovich et al., 1985; Tversky & Kahneman, 1974). For example, Figure 4 show that at $p_A = 1/2$, the model had higher activations for alternation patterns than repetition patterns, in spite of the sequential independence of events. Critically, such bias emerged through unsupervised training without any pre-defined hypothesis structures, since both models were initially symmetrically structured, and were not provided with any prior knowledge on how different p_A levels would affect the occurrences of different patterns.

Given the simplicity of our models, one far-reaching implication is that cognitive biases and structured abstractions can emerge early and locally at the level of sensory processing. Nevertheless, it should be noted that our models only address purely bottom-up learning mechanisms without implementing any top-down learning or higher-level representations such as beliefs or goals. For the early and locally developed biases to be maintained and utilized in later and global processes, higher-level representations and top-down structures must also be involved through a hierarchical structure of abstractions (Munakata et al., 2011; Tenenbaum et al., 2011).

Lastly, probabilistic thinking has to consider the consequence of time (Buchanan, 2013; Hawkins & Blakeslee, 2004). Our findings suggest that rich semantics in the learning environment can be extracted by neuron populations in predictive learning through temporal integration. This learning over time would lead to the structured hypothesis spaces such as those required by Bayesian inference thus provide essential building blocks in bridging the gap between neural computations and overt behavior.

Acknowledgments

This work was supported by the Air Force Office of Scientific Research (AFOSR) grant number FA9550-12-1-0457, the Office of Naval Research (ONR) grant number N00014-16-1-2111.

References

- Adolphs, R. (2015). The unsolved problems of neuroscience. *Trends in Cognitive Sciences*, 19(4), 173–175. doi: 10.1016/j.tics.2015.01.007
- Bienenstock, E. L., Cooper, L. N., & Munro, P. W. (1982). Theory for the development of neuron selectivity: Orientation specificity and binocular interaction in visual cortex. *The Journal of Neuroscience*, 2(1), 32–48.
- Buchanan, M. (2013). Gamble with time. *Nature Physics*, 9(1), 3. doi: 10.1038/nphys2520

- Budescu, D. V. (1987). A Markov model for generation of random binary sequences. *Journal of Experimental Psychology: Human Perception and Performance*, 13(1), 25–39. doi: 10.1037/0096-1523.13.1.25
- Cooper, L. N. (2000). Memories and memory: A physicist's approach to the brain. *International Journal of Modern Physics A*, 15(26), 4069–4082. doi: 10.1142/S0217751X0000272X
- Dehaene, S., & Brannon, E. M. (2010). Space, time, and number: A Kantian research program. *Trends in Cognitive Sciences*, 14(12), 517–519. doi: 10.1016/j.tics.2010.09.009
- Elman, J. L. (1990). Finding structure in time. *Cognitive Science*, 14(2), 179–211. doi: 10.1207/s15516709cog1402_1
- Falk, R., & Konold, C. (1997). Making sense of randomness: Implicit encoding as a basis for judgment. *Psychological Review*, 104(2), 301–318. doi: 10.1037/0033-295x.104.2.301
- Gigerenzer, G., & Hoffrage, U. (1995). How to improve Bayesian reasoning without instruction: Frequency formats. *Psychological Review*, 102(4), 684–704. doi: 10.1037/0033-295x.102.4.684
- Gilovich, T., Vallone, R., & Tversky, A. (1985). The hot hand in basketball: On the misperception of random sequences. *Cognitive Psychology*, 17(3), 295–314. doi: 10.1016/0010-0285(85)90010-6
- Goodfellow, L. D. (1938). A psychological interpretation of the results of the zenith radio experiments in telepathy. *Journal of Experimental Psychology*, 23(6), 601–632. doi: 10.1037/h0058392
- Griffiths, T. L., & Tenenbaum, J. B. (2001). Randomness and coincidences: Reconciling intuition and probability theory. In J. D. Moore & K. Stenning (Eds.), *Proceedings of the 23rd annual conference of the cognitive science society* (pp. 370–375). Mahwah, NJ: Lawrence Erlbaum Associates.
- Hahn, U., & Warren, P. A. (2009). Perceptions of randomness: Why three heads are better than four. *Psychological Review*, 116(2), 454–461. doi: 10.1037/a0015241
- Hawkins, J., & Blakeslee, S. (2004). *On intelligence*. New York: Henry Holt.
- Hebb, D. O. (1949). *The organization of behavior*. New York: Wiley.
- Lopes, L. L., & Oden, G. C. (1987). Distinguishing between random and nonrandom events. *Journal of Experimental Psychology: Learning Memory and Cognition*, 13(3), 392–400. doi: 10.1037/0278-7393.13.3.392
- Munakata, Y., Herd, S. A., Chatham, C. H., Depue, B. E., Banich, M. T., & O'Reilly, R. C. (2011). A unified framework for inhibitory control. *Trends in Cognitive Sciences*, 15(10), 453–459. doi: 10.1016/j.tics.2011.07.011
- Nickerson, R. S. (2002). The production and perception of randomness. *Psychological Review*, 109(2), 330–357. doi: 10.1037//0033-295X.109.2.330
- Oppenheimer, D. M., & Monin, B. (2009). The retrospective gambler's fallacy: Unlikely events, constructing the past, and multiple universes. *Judgment and Decision Making*, 4(5), 326–334.
- O'Reilly, R. C., Munakata, Y., Frank, M. J., Hazy, T. E., & Contributors. (2012). *Computational cognitive neuroscience*. Wiki Book, 1st Edition, URL: <http://ccnbook.colorado.edu>.
- O'Reilly, R. C., Wyatte, D., & Rohrlich, J. (2014). Learning through time in the thalamocortical loops. *Preprint at: http://arxiv.org/abs/1407.3432*.
- Oskarsson, A. T., Van Boven, L., McClelland, G. H., & Hastie, R. (2009). What's next? Judging sequences of binary events. *Psychological Bulletin*, 135(2), 262–285. doi: 10.1037/a0014821
- Pouget, A., Beck, J. M., Ma, W. J., & Latham, P. E. (2013). Probabilistic brains: Knowns and unknowns. *Nature Neuroscience*, 16(9), 1170–1178. doi: 10.1038/nn.3495
- Ross, S. M. (2007). *Introduction to probability models* (9th ed.). San Diego, CA: Academic Press.
- Sun, Y., O'Reilly, R. C., Bhattacharyya, R., Smith, J. W., Liu, X., & Wang, H. (2015). Latent structure in random sequences drives neural learning toward a rational bias. *Proceedings of the National Academy of Sciences*, 112(12), 3788–3792. doi: 10.1073/pnas.1422036112
- Sun, Y., Tweney, R. D., & Wang, H. (2010). Occurrence and nonoccurrence of random sequences: Comment on Hahn and Warren (2009). *Psychological Review*, 117(2), 697–703. doi: 10.1037/a0018994
- Sun, Y., & Wang, H. (2010a). Gambler's fallacy, hot hand belief, and time of patterns. *Judgment and Decision Making*, 5(2), 124–132.
- Sun, Y., & Wang, H. (2010b). Perception of randomness: On the time of streaks. *Cognitive Psychology*, 61(4), 333–342. doi: 10.1016/j.cogpsych.2010.07.001
- Sun, Y., & Wang, H. (2012). Perception of randomness: Subjective probability of alternation. In N. Miyake, D. Peebles, & R. P. Cooper (Eds.), *Proceedings of the 34th annual conference of the cognitive science society* (pp. 1024–1029). Austin, TX: Cognitive Science Society.
- Sun, Y., & Wang, H. (2015). Generating functions in neural learning of sequential structures. In D. C. Noelle et al. (Eds.), *Proceedings of the 37th annual conference of the cognitive science society* (pp. 2302–2307). Austin, TX: Cognitive Science Society.
- Tenenbaum, J. B., Kemp, C., Griffiths, T. L., & Goodman, N. D. (2011). How to grow a mind: Statistics, structure, and abstraction. *Science*, 331(6022), 1279–1285. doi: 10.1126/science.1192788
- Tversky, A., & Kahneman, D. (1974). Judgment under uncertainty: Heuristics and biases. *Science*, 185(4157), 1124–1131. doi: 10.1126/science.185.4157.1124