
Probabilistic Information Processing in Humans and Recurrent Neural Networks

Anonymous Author(s)

Affiliation

Address

email

Abstract

1 In nature, sensory inputs are often highly structured, and statistical regularities
2 of these signals can be extracted to form expectation about future sensorimotor
3 associations, thereby facilitating optimal behavior. To date, the circuit mechanisms
4 that underlie these probabilistic computations are not well understood. Through
5 a human electrophysiological experiment and a recurrent neural network (RNN)
6 model, the present study investigates how the brain extracts, processes, and utilizes
7 probabilistic structures of sensory signals to guide behavior. To achieve this goal,
8 we first constructed and trained a biophysically constrained RNN model to perform
9 a probabilistic decision making task similar to task paradigms designed for humans.
10 Specifically, the training environment was probabilistic such that one stimulus
11 was more probable than the others. We show that both humans and the RNN
12 model successfully extract information about stimulus probability and integrate this
13 knowledge into their decisions and task strategy in a new environment. Specifically,
14 performance of both humans and the RNN model varied with the degree to which
15 the stimulus probability of the new environment matched the formed expectation.
16 In both humans and RNNs, this expectation effect was more prominent when the
17 strength of sensory evidence was low. These findings suggest that both humans
18 and our RNN model placed more emphasis on prior expectation (top-down signals)
19 when the available sensory information (bottom-up signals) was limited. Finally,
20 by dissecting the trained RNN model, we demonstrate how competitive inhibition
21 and recurrent excitation form the basis for neural circuitry optimized to perform
22 probabilistic information processing.

23 1 Introduction

24 The brain contains billions of neurons, each connecting with up to 10,000 other neurons. Together,
25 these neurons continually produce electrical signals to represent and relay information about the
26 sensory environments. Notably, this task is made more challenging by the fact that sensory inputs are
27 highly dynamic and often ambiguous, especially as animals traverse novel environments. In addition,
28 naturalistic sensory environments are inherently probabilistic and certain stimuli are subsequently
29 encountered more frequently than others. To optimize sensory processing, the brain needs to extract
30 the statistics of the sensory environments to form expectation against which incoming signals
31 are compared [1]. This expectation captures the inherent probabilistic structures of the sensory
32 environments and can be used to fine-tune and adapt behavioral responses as the animals encounter
33 novel environments. For example, knowledge of statistical regularities in the environment can lead
34 to faster recognition of objects when they are encountered in an expected context (e.g., a bird in a
35 backyard) than when they are encountered in an unlikely context (e.g., a bird in a washing machine;
36 [2, 3]). While certain aspects of probabilistic information processing have been investigated through

37 human psychophysics, neuroimaging, and lesion studies, the circuit mechanisms that underlie this
38 complex cognitive function is not well understood.

39 Recent work in patients with bilateral damage to the hippocampus demonstrates that an ability
40 to learn stimulus probability and use this knowledge to perform cognitive tasks circumvents the
41 main hippocampal pathway which mediates various other kinds of learning [4]. In addition, human
42 electrophysiological and psychophysical studies have shown that learned expectation based on
43 stimulus probability improved behavior by primarily modulating post-perceptual processes including
44 the selection and execution of motor responses rather than directly sharpening sensory signals in
45 the primary visual areas [5, 6]. To investigate the circuitry that underlies probabilistic information
46 processing, the present study examines how humans and recurrent neural network (RNN) model
47 process and utilize information about stimulus probability in a perceptual decision making task.
48 We show that both humans and the RNN model extract and use probabilistic information of the
49 sensory stimuli to improve their performance in an environment where the stimulus probability
50 closely aligns with the learned expectation. In addition, we show that the enhanced processing of the
51 highly probable (expected) stimulus can be attributed to a combination of decreased inhibition and
52 increased recurrent excitation within a neuronal subgroup selective for the expected stimulus. This
53 unique circuit mechanism gives rise to a sub-population of neurons with enhanced neural activity
54 which has been previously observed in experimental studies.

55 **2 Human behavioral data**

56 Expectation based on stimulus regularities exert a powerful influence on human perception and
57 decision making. To examine how humans process and use probabilistic information to guide
58 behavior, we tested human participants on a perceptual decision making task where the stimulus
59 probability and the amount of sensory strength were manipulated. On each trial, participants
60 monitored a display of flickering randomly-oriented red and blue bars and reported a brief target
61 display which was either red or blue bars coherently oriented at 0° (horizontal) or 90° (vertical;
62 Figure 1a). Electroencephalography (EEG) was concurrently measured as the participants performed
63 this task. Stimulus probability was manipulated on a block-to-block basis such that within each block
64 of trials, one target feature (e.g., horizontal orientation) was more prevalent (70%) than the other
65 feature (i.e., vertical orientation; 30%). In addition, experimental blocks where the target features
66 (e.g., orientation) was counterbalanced (50% each) were also interleaved throughout the experiment
67 and were used as a neutral testing context to evaluate if stimulus probability was learned on the
68 previous (training) block. Finally, sensory strength was manipulated through the frequency at which
69 the red and blue bars were being rendered. In this study, we considered two strength values: low
70 flicker frequency at 33 Hz and high frequency at 50 Hz.

71 Together, this behavioral paradigm incorporates three comparison conditions by varying the stimulus
72 probability on a block-to-block basis: (1) ‘expected’ condition where two adjacent blocks shared
73 the likely target stimulus (e.g., coherently horizontal bars presented as targets more often in two
74 adjacent blocks), (2) ‘unexpected’ condition where two adjacent blocks did not share the likely target
75 stimulus (e.g., coherently horizontal bars presented as targets more likely in one block and coherently
76 vertical bars as targets more likely in the next block), and (3) ‘neutral’ condition where the coherently
77 horizontal and coherently vertical bars were presented as targets equally likely in each block. Finally,
78 the manipulation of the flicker frequency allowed us to investigate the probabilistic information
79 processing under different levels of sensory strength.

80 **3 RNN model**

81 Recurrent neural network (RNN) models have recently been used to investigate various components
82 of neural computations such as perceptual inference and working memory [7, 8]. Because RNN
83 models allow an experimenter to manipulate many aspects of the sensory environment and learning
84 processes—including the network architecture, the task, and the stimulus set—these models offer a
85 complementary approach to in-vivo methods. To date, RNN models have been primarily trained in an
86 environment where a stimulus space is counterbalanced and unbiased to ensure that each stimulus
87 is equally represented and accounted for by the neural dynamics of the trained models. However,
88 naturalistic sensory environments are inherently probabilistic and certain stimuli are subsequently
89 encountered more frequently than others.

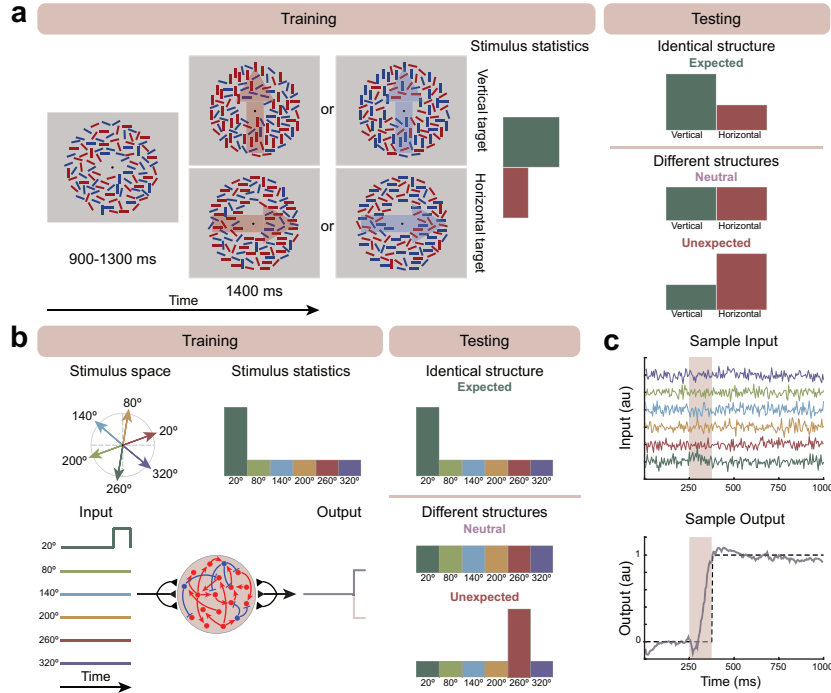


Figure 1: Task paradigms for humans and RNN model. **a**. Probabilistic decision making task. Human participants reported an orientation (horizontal or vertical) of coherently oriented red or blue bars. In this task paradigm, vertically oriented bars (green bar in the histogram) were more likely to be the targets as compared to horizontally oriented bars (red bar in the histogram) during training. After training, the participants were tested on three different ‘environments’ (right panel). **b**. Probabilistic decision making task used to train and test the RNN model. **c**. Network input and output from a sample RNN model on a trial where the presented stimulus was highly probable (i.e., Stimulus 1).

90 In order to model more realistic and biased sensory environments, we developed a simple training
 91 paradigm that resembled the one used in human studies (Figure 1b). In this paradigm, one input
 92 stimulus (out of 6 possible stimuli) was over-represented and presented to our RNN model more often
 93 than the rest of the stimuli during training. More specifically, the ‘expected’ stimulus was present in
 94 80% of the training trials, while the other five ‘unexpected’ stimuli were equally represented (4%
 95 each). On each trial, one out of 6 possible stimuli was presented for 125 ms. The stimulus signals
 96 were modeled as white-noise signals (drawn from the standard normal distribution) with a constant
 97 offset value added during the stimulus window (Figure 1c, top panel). The offset value was varied to
 98 model the sensory strength and coherence (i.e., flicker frequency in the human data). The offset value
 99 of 0.6 and 0.7 were used to simulate the low and high flicker frequency, respectively. In total, we
 100 trained 30 RNNs of 200 units (80% excitatory and 20% inhibitory units) to perform this task.

101 Using this paradigm, we trained a continuous rate RNN model to produce an output signal approaching
 102 +1 when the expected stimulus was shown and 0 when an unexpected stimulus was given (Figure 1c).
 103 The RNN model employed in this study was similar to the one used in [9], and we used a gradient-
 104 descent method to train the model.

105 4 Results

106 In the human dataset, stimulus probability influenced behavioral performance such that accuracy was
 107 higher in the expected context than in the neutral and unexpected context in both low and high flicker
 108 frequency conditions (Figure 2a). Accuracy was also higher on trials where stimuli were rendered at
 109 a high flicker frequency (strong sensory strength) compared to the trials with a low flicker frequency.
 110 In addition, improvement in accuracy from the unexpected to expected context was more prominent
 111 in the low flicker frequency condition, suggesting that participants relied on their formed expectation
 112 more when the available sensory information was limited. This finding illustrates an interplay of

113 top-down (e.g., expectation) and bottom-up (e.g., sensory strength) signals and is in line with several
 114 theoretical and empirical work [10–13].

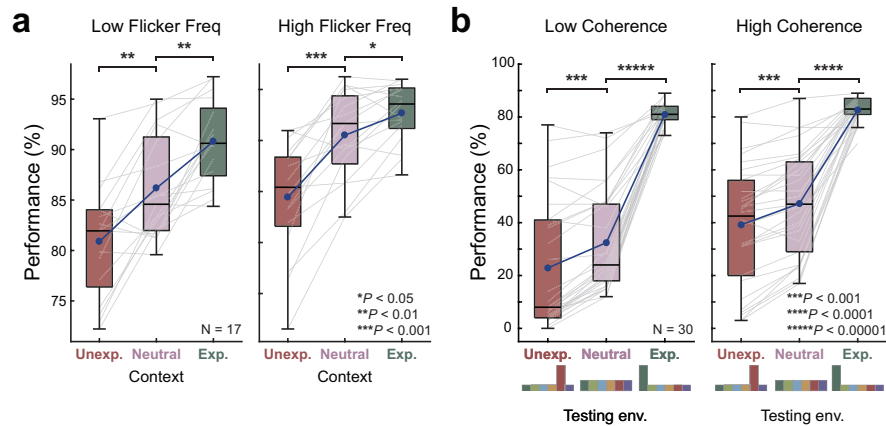


Figure 2: Performance on the probabilistic decision making task across contexts and testing environments as a function of the strength of sensory evidence. **a.** In humans, accuracy was higher in the expected context than in the neutral and unexpected context in both low and high flicker frequency conditions. Accuracy was also higher on trials where stimuli were presented at a high compared to a low flickering frequency. **b.** The RNN models exhibited similar performance trends as the human participants where performance was highest when the stimulus probability of the testing environment matched the learned expectation for both levels of sensory strength.

115 We observed similar task performance trends in our RNN model. When the stimulus probability
 116 was identical during the testing phase (i.e., Stimulus 1 was over-represented), our RNN model
 117 performed the task with high accuracy (Figure 2b) in both coherence conditions. When the stimulus
 118 environment was dramatically changed (i.e., one of the unexpected stimuli was over-represented),
 119 the model task performance significantly decreased (Figure 2b). The model performance for the
 120 ‘neutral’ environment, where all 6 stimuli were equally represented, was significantly higher than the
 121 performance from the ‘unexpected’ testing environment but lower than the one from the ‘expected’
 122 condition. Overall, our RNN model findings are closely aligned with the results from the human
 123 behavioral data.

124 Although the exact circuit mechanisms underlying probabilistic information processing are not
 125 known, a recent experimental study revealed that probabilistic learning led to a group of neurons
 126 in the mouse primary visual cortex responding more robustly to expected or likely stimuli [14]. In
 127 order to investigate if such subgroup of neurons also exists in our RNN model, we first classified all
 128 the neurons in each trained RNN model based on their firing patterns in response to the six stimuli.
 129 For example, if a neuron fired more often when Stimulus 1 was presented, the neuron was assigned
 130 to the ‘Stimulus 1’ cluster. Using this method, we identified 6 subgroups of neurons in each RNN.
 131 For each subgroup, we then characterized its connectivity patterns to the rest of the subgroups. The
 132 resulting connectivity patterns revealed that the subgroup corresponding to the expected stimulus
 133 (i.e., Stimulus 1) had significantly stronger within-group excitation than across-group excitation
 134 for both coherence levels (Figure 3). In addition, the subgroup also exhibited significantly weaker
 135 within-group inhibition than across-group inhibition (Figure 3). Interestingly, the low coherence
 136 condition accentuated these differences in excitation and inhibition (Figure 3 d–f). Therefore, training
 137 our RNN model using a biased set of stimuli led to emergence of a subgroup of neurons that were
 138 more likely to stimulate themselves and inhibit other neurons resulting in enhanced neural responses
 139 at baseline.

140 5 Conclusions and future directions

141 In the present study, we employed a biophysically constrained RNN model to investigate the circuit
 142 mechanisms that underlie probabilistic information processing. By devising an RNN paradigm that
 143 resembles the perceptual tasks used to study probabilistic learning in humans, we validated and
 144 compared model performance to that of humans. Specifically, the model and humans both displayed

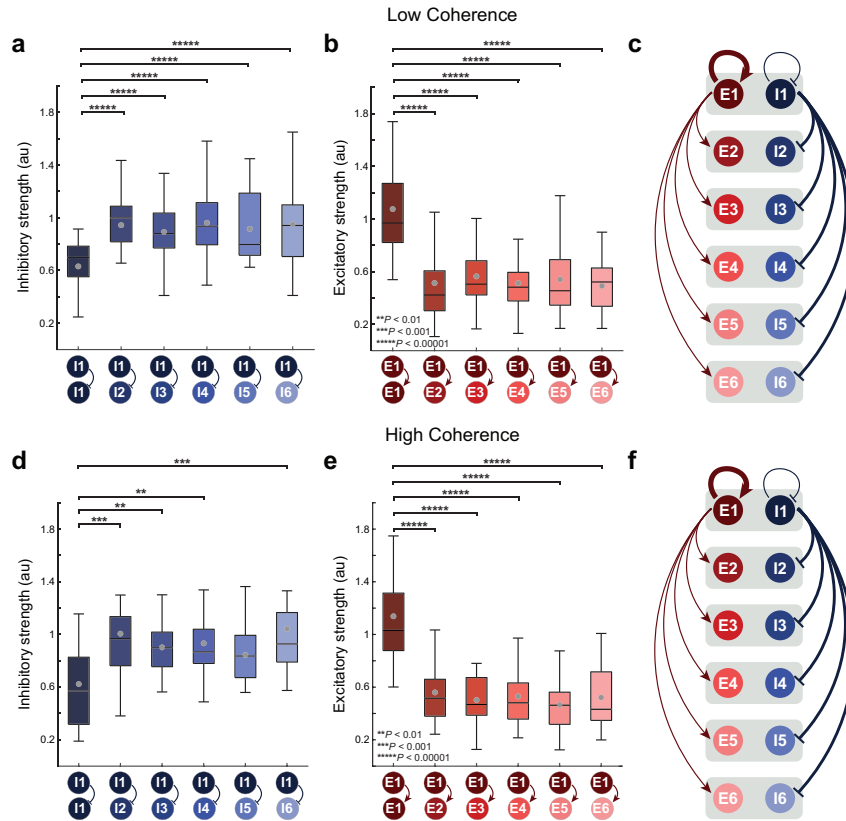


Figure 3: Network connectivity of the RNN models trained on an environment where stimulus 1 was highly probable (i.e., expected). Strength of pairwise inhibitory (left panel) and excitatory signals (middle panel) between neuron subgroups selective for each of the 6 stimuli were plotted. I_i and E_i represent the inhibitory and excitatory subgroup selective for stimulus i , respectively. **a.** In the low coherence condition, the neuron subgroup selective for the expected stimulus displayed a weaker within-group inhibition than across-group inhibition. **b.** The subgroup selective for the expected stimulus also exhibited a stronger within-group excitation than across-group excitation. **c.** A summary diagram of the pairwise inhibitory and excitatory strength in the low coherence condition. **d-f.** In the high coherence condition, the RNN models exhibited similar network connectivity patterns as observed in the low coherence condition. Boxplot central lines, median; gray circles, mean; bottom and top edges, lower and upper quartiles; whiskers, 1.5*interquartile range; outliers not plotted. ** $P < 0.01$, *** $P < 0.001$, ***** $P < 0.00001$ by Kruskal-Wallis, Dunn's multiple comparisons post hoc.

145 similar performance trends such that task accuracy increased with the degree of similarity between the
 146 stimulus probability of the testing and training environment or context. In addition, this expectation
 147 effect on behavior was accompanied by a combination of decreased within-group inhibition and
 148 increased within-group excitation of the neuronal subgroup selective for the expected stimulus.

149 Future works include comparing the RNN model connectivity results with animal electrophysiological
 150 data such as EEG and local field potentials (LFP). In particular, the behavioral improvement in the
 151 expected context reported in our human study was accompanied by changes in the pattern of parietal
 152 alpha and frontal theta oscillatory signals, previously used to index task effort and cognitive control.
 153 Specifically, parietal alpha power showed a sustained decrease in the unexpected context condition.
 154 This is consistent with the notion that violations of expectation require greater and prolonged task
 155 engagement [15]. Further, we observed an increase in frontal theta power when the target appeared
 156 in the unexpected context. This pattern is consistent with the hypothesis that unexpected contexts
 157 require greater executive control during later stages of decision-making including preparation and
 158 execution of motor responses [16]. Further analyses of the RNN network structures could shed light
 159 on the circuitry that gives rise to these oscillatory differences as a basis for probabilistic information
 160 processing in the human cortex.

161 **Broader Impact**

162 By analyzing both human and model data, the present study proposes an experimentally testable
163 neural circuit mechanism important for decision making. Elucidating circuit mechanisms required for
164 decision making will help better understand how such mechanisms are disrupted in neuropsychiatric
165 disorders such as Alzheimer’s disease and schizophrenia. We believe that the present study poses no
166 negative ethical or societal issues.

167 **References**

- 168 [1] Christopher Summerfield and Floris P De Lange. Expectation in perceptual decision making: neural and
169 computational mechanisms. *Nature Reviews Neuroscience*, 15(11):745–756, 2014.
- 170 [2] Irving Biederman, Arnold L Glass, and E Webb Stacy. Searching for objects in real-world scenes. *Journal*
171 *of experimental psychology*, 97(1):22, 1973.
- 172 [3] Floris P De Lange, Micha Heilbron, and Peter Kok. How do expectations shape perception? *Trends in*
173 *cognitive sciences*, 22(9):764–779, 2018.
- 174 [4] Nuttida Rungratsameetaweemana, Larry R Squire, and John T Serences. Preserved capacity for learning
175 statistical regularities and directing selective attention after hippocampal lesions. *Proceedings of the*
176 *National Academy of Sciences*, 116(39):19705–19710, 2019.
- 177 [5] Ji Won Bang and Dobromir Rahnev. Stimulus expectation alters decision criterion but not sensory signal
178 in perceptual decision making. *Scientific reports*, 7(1):1–12, 2017.
- 179 [6] Nuttida Rungratsameetaweemana, Sirawaj Itthipuripat, Annalisa Salazar, and John T Serences. Expecta-
180 tions do not alter early sensory processing during perceptual decision-making. *Journal of Neuroscience*,
181 38(24):5632–5648, 2018.
- 182 [7] Guangyu Robert Yang, Madhura R. Joglekar, H. Francis Song, William T. Newsome, and Xiao-Jing Wang.
183 Task representations in neural networks trained to perform many cognitive tasks. *Nature Neuroscience*, 22
184 (2):297–306, 2019.
- 185 [8] Robert Kim and Terrence J. Sejnowski. Strong inhibitory signaling underlies stable temporal dynamics
186 and working memory in spiking neural networks. *bioRxiv*, 2020.
- 187 [9] Robert Kim, Yinghao Li, and Terrence J Sejnowski. Simple framework for constructing functional spiking
188 recurrent neural networks. *Proceedings of the national academy of sciences*, 116(45):22811–22820, 2019.
- 189 [10] Andre M Bastos, W Martin Usrey, Rick A Adams, George R Mangun, Pascal Fries, and Karl J Friston.
190 Canonical microcircuits for predictive coding. *Neuron*, 76(4):695–711, 2012.
- 191 [11] Hendrikje Nienborg and Bruce G Cumming. Decision-related activity in sensory neurons reflects more
192 than a neuron’s causal effect. *Nature*, 459(7243):89–92, 2009.
- 193 [12] Karl Friston. A theory of cortical responses. *Philosophical transactions of the Royal Society B: Biological*
194 *sciences*, 360(1456):815–836, 2005.
- 195 [13] Michael W Spratling. Reconciling predictive coding and biased competition models of cortical function.
196 *Frontiers in computational neuroscience*, 2:4, 2008.
- 197 [14] Anderson Speed, Joseph Del Rosario, Navid Mikail, and Bilal Haider. Spatial attention enhances network,
198 cellular and subthreshold responses in mouse visual cortex. *Nature communications*, 11(1):1–11, 2020.
- 199 [15] Jan R Wessel and Adam R Aron. On the globality of motor suppression: unexpected events and their
200 influence on behavior and cognition. *Neuron*, 93(2):259–280, 2017.
- 201 [16] James F Cavanagh and Michael J Frank. Frontal theta as a mechanism for cognitive control. *Trends in*
202 *cognitive sciences*, 18(8):414–421, 2014.