

A neural circuit model for a contextual association task inspired by recommender systems

Henghui Zhu¹ | Ioannis Ch. Paschalidis²  | Allen Chang³  | Chantal E. Stern³ | Michael E. Hasselmo³ 

¹Division of Systems Engineering, Boston University, Boston, Massachusetts

²Department of Electrical and Computer Engineering, Division of Systems Engineering, and Department of Biomedical Engineering, Boston University, Boston, Massachusetts

³Department of Psychological and Brain Sciences, and Center for Systems Neuroscience, Boston University, Boston, Massachusetts

Correspondence

Ioannis Ch. Paschalidis, Department of Electrical and Computer Engineering, Division of Systems Engineering, and Department of Biomedical Engineering, Boston University, 8 Mary's St., Boston, MA 02215.
Email: yannis@bu.edu

Funding information

National Institute of General Medical Sciences, Grant/Award Number: 1R01GM135930; National Science Foundation, Grant/Award Numbers: CNS-1645681, DMS-1664644, IIS-1914792; Office of Naval Research, Grant/Award Numbers: N00014-16-1-2832, N00014-19-1-2571

Abstract

Behavioral data shows that humans and animals have the capacity to learn rules of associations applied to specific examples, and generalize these rules to a broad variety of contexts. This article focuses on neural circuit mechanisms to perform a context-dependent association task that requires linking sensory stimuli to behavioral responses and generalizing to multiple other symmetrical contexts. The model uses neural gating units that regulate the pattern of physiological connectivity within the circuit. These neural gating units can be used in a learning framework that performs low-rank matrix factorization analogous to recommender systems, allowing generalization with high accuracy to a wide range of additional symmetrical contexts. The neural gating units are trained with a biologically inspired framework involving traces of Hebbian modification that are updated based on the correct behavioral output of the network. This modeling demonstrates potential neural mechanisms for learning context-dependent association rules and for the change in selectivity of neurophysiological responses in the hippocampus. The proposed computational model is evaluated using simulations of the learning process and the application of the model to new stimuli. Further, human subject behavioral experiments were performed and the results validate the key observation of a low-rank synaptic matrix structure linking stimuli to responses.

KEYWORDS

context association task, matrix factorization, neural circuit model

1 | INTRODUCTION

Behavioral data from a range of cognitive tasks indicate that humans and animals can learn rules based on specific examples and generalize these rules to a broader range of different contexts (Aminoff, Gronau, & Bar, 2006; Badre & Frank, 2012; Bar, Aminoff, & Ishai, 2007; Bhandari & Badre, 2018; Carpenter, Just, & Shell, 1990; Chatham, Frank, & Badre, 2014; Eliasmith et al., 2012; Hummel & Holyoak, 1997; Miller & Cohen, 2001; Rasmussen & Eliasmith, 2011; Raudies & Hasselmo, 2017; Santoro, Hill, Barrett, Morcos, & Lillicrap, 2018; Wallis, Anderson, & Miller, 2001). This learning of context-dependent

rules is consistent with the interaction of general roles and specific fillers in symbolic processing, in which a rule learned with specific instances of stimuli and contexts can be generalized to apply to previously unseen combinations of stimuli and contexts (Badre & Frank, 2012; Chatham et al., 2014; Hasselmo & Stern, 2018; Hummel & Holyoak, 1997). This process enables agents to generalize well from previous experiences and interpret previously unseen sensory input according to a learned context-dependent set of rules. For example, humans learn that a red light means to stop when driving, and can generalize this rule to multiple locations, but they also learn more complex, location-dependent rules such as the fact that one can

turn right after stopping at a red light, except in certain cities and countries.

Experimental data from the rodent hippocampus addresses potential neurophysiological changes associated with context-dependent learning and generalization. Data shows that neural responses increase in selectivity dependent upon context in a behavioral task that provides reward for different stimulus items in different contexts (Komorowski et al., 2013; Komorowski, Manns, & Eichenbaum, 2009). In this task, the rat is in one of two visually distinct environments, which define the context for the rat. Each of the two environments has two stimulus pots, with distinct features such as filling material. The stimulus pots can appear in different locations, but only one contains reward dependent upon context. Neurophysiological recordings show that hippocampal neurons develop specificity toward specific pairings of a stimulus item in a specific context during learning (Komorowski et al., 2009). In addition, data from the hippocampus show the replay of memory representations during sharp-wave ripple events during quiet waking and sleep that can occur in forward or backward order (Carr, Jadhav, & Frank, 2011; Diba & Buzsáki, 2007; Karlsson & Frank, 2009; Lee & Wilson, 2002). These replay events provide an opportunity for the development of rule representations for generalization. In this light, our work can be seen as modeling the effect of hippocampal circuit mechanisms involving repeated interleaved reactivation of learning examples in different sequential orders referred to here as primal (stimulus followed by context) and dual (context followed by stimulus). This interleaved reactivation could generate the context-item selectivity seen experimentally in this task, which could then generalize to the selectivity for different contexts. The replay could alter context-item selectivity by updating gating mechanisms in the hippocampus in which entorhinal input to region CA1 gates the influence of synaptic input from region CA3 to region CA1 as shown in previous models (Hasselmo & Eichenbaum, 2005; Katz, Kath, Spruston, & Hasselmo, 2007).

Understanding how humans learn to make flexible decisions has motivated considerable research on biologically plausible neural circuit models in the brain (Badre & Frank, 2012; Badre, Kayser, & D'Esposito, 2010; Bhandari & Badre, 2018; Chang, Johnson, Whiteman, & Stern, 2019; Chang, Ren, Whiteman, & Stern, 2017; Chatham et al., 2014; Eliasmith et al., 2012; Hasselmo, 2005; Hasselmo & Stern, 2018; Melrose, Poulin, & Stern, 2007; O'Reilly, 1998; O'Reilly, Hazy, & Herd, 2016; Zhu, Paschalidis, & Hasselmo, 2018). This work includes models of cognitive function in cortical circuits based on theoretical frameworks such as the Semantic Pointer architecture (Eliasmith et al., 2012) and the LEABRA cognitive architecture (O'Reilly, 1998; O'Reilly et al., 2016). In particular, many models include mechanisms for gating the spread of neural activity between regions to recruit different neural circuits for flexible application of different context-dependent rules (Badre et al., 2010; Badre & Frank, 2012; Bhandari & Badre, 2018; Chang et al., 2017, 2019; Chatham et al., 2014; Eliasmith et al., 2012; Hasselmo, 2005; Hasselmo & Stern, 2018; Miller & Cohen, 2001; O'Reilly, 1998; O'Reilly et al., 2016; Zhu et al., 2018). On a circuit level, this could involve the role of gating neurons that regulate the response of other neurons to synaptic input due to the nonlinear interaction of adjacent synapses on the dendritic tree (Mel, 1993; Poirazi, Brannon, &

Mel, 2003), or circuit-level multiplicative interactions (Nezis & Van Rossum, 2011; Sherfey, Ardid, Hass, Hasselmo, & Kopell, 2018) that can be mediated by populations of neurons (Eliasmith & Anderson, 2004), or oscillatory dynamics of cortical circuits (Buschman, Denovellis, Diogo, Bullock, & Miller, 2012; Lundqvist et al. 2018,b; Sherfey et al., 2018).

In this article, we present two different levels of analysis for mechanisms of learning a context-dependent association task and generalizations to additional instances of stimuli-context pairs that were not previously seen. We will describe a mechanism of matrix factorization that is used to learn and generalize with accuracy of 100% for this task. We will also describe a neurally plausible learning rule that can achieve the same performance on a symmetrical version of the task.

Our computational model significantly outperforms earlier models because it has the ability, through low-rank matrix factorization, to discover the proper amount of internal memory needed for the task and use additional available memory for redundancy. We analyze results from human subject behavioral experiments that validate the key observation of a low-rank synaptic connection structure linking stimuli to behavioral responses.

2 | METHODS

2.1 | Task overview

We model a behavioral task that requires learning of rules guiding the association of specific sets of stimuli with the behavioral responses under different location contexts (Chang et al., 2017, 2019; Hasselmo & Stern, 2018). This task represents a generalization from the context-dependent learning task used in the rodent hippocampus (Komorowski et al., 2009), which had only two spatial contexts compared to the four used in the more recent task and modeled here. The association rule is summarized in Figure 1. There are four location contexts 1, 2, 3, and 4, represented by different quadrants on a computer monitor screen. Under different contexts, four stimuli, A, B, C, and D are associated with two responses X and Y in different ways. We will use the term "response" for X and Y because they can be seen as behavioral response to an input consisting of a stimulus-context pair, for example, A2. As we will see later in the Experimental Results, X and Y are represented by a second stimulus and human participants are asked to associate the first stimulus (A–D) with the second (X,Y) in different contexts (1–4). Contexts 1 and 4 share the same association rule while Contexts 2 and 3 share a different association rule, as shown in Figure 1. We aim at investigating how a neural circuit model can be constructed to learn these rules for association using interleaved reactivation of the examples, consistent with the replay phenomenon described for the hippocampus (Carr et al., 2011; Diba & Buzsáki, 2007; Karlsson & Frank, 2009; Lee & Wilson, 2002). These replay phenomena can provide the interleaved dynamics for learning involving bidirectional reactivations of stimulus representations followed by context (primal) or context followed by stimulus (dual)

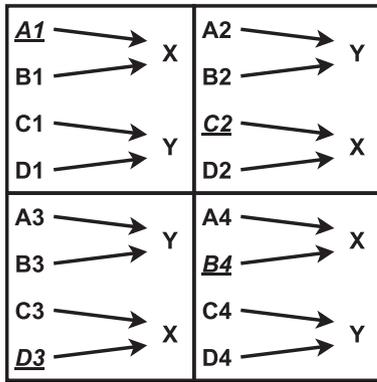


FIGURE 1 The association rules for the context association task with four spatial location contexts indicated by the numbers 1, 2, 3, or 4. Different stimuli indicated by A, B, C, or D can be presented in each context. The correct association is indicated by arrows. For example, stimulus B in Location 1 is associated with response X, but stimulus B in Location 2 is associated with response Y. The underlined stimulus-context pairs are hidden during training. The task is described in more detail in another article in this special issue (Chang et al., 2019)

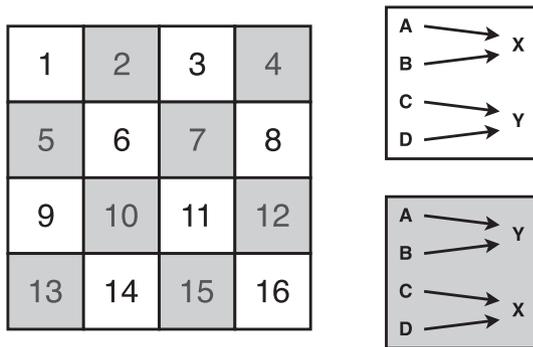


FIGURE 2 A different variant of the task testing the capacity for broader generalization of the context association task to 16 contexts. Different context-dependent rules are shown in the white and shaded boxes on the right. These rules are applied to a larger number of 16 different contexts on the left, with the (white or shaded) background indicating the rule for each context

that prove most effective for learning as described in Section 2.3. Also, we consider the ability of the neural circuit model to generalize to other instances of the task that were not previously encountered. In this generalization test, we hide some stimuli in the task during training (corresponding to the underlined stimulus-context pairs in Figure 1). We train a neural circuit with all the stimulus-context pairs without ever presenting the hidden ones during training, and then test its accuracy on all stimulus-context pairs.

In addition to the context association task defined earlier, we also consider a larger context association task with 16 different contexts, shown in Figure 2.

On a notational remark, all vectors are assumed to be column vectors and denoted by bold lowercase letters; for economy of space, we will write $\mathbf{x} = (x_1, \dots, x_n)$ for $\mathbf{x} \in \mathbb{R}^n$. Matrices will be denoted by bold uppercase letters.

2.2 | Response matrix factorization

We first define the response matrix of the task, whose columns correspond to stimuli and rows to contexts. The elements of the matrix represent the likelihood of the response being Y for the corresponding stimulus-context pair. The likelihood starts with a value between 0 and 1 and can change in a graded manner during learning. The correct response matrix for the task of Figure 1 is shown below:

$$\mathbf{R} = \begin{matrix} & \begin{matrix} A & B & C & D \end{matrix} \\ \begin{matrix} 1 \\ 2 \\ 3 \\ 4 \end{matrix} & \begin{bmatrix} 0 & 0 & 1 & 1 \\ 1 & 1 & 0 & 0 \\ 1 & 1 & 0 & 0 \\ 0 & 0 & 1 & 1 \end{bmatrix} \end{matrix}. \quad (1)$$

Notice that \mathbf{R} can be factorized as the product of two low-rank matrices $\mathbf{M} = (M_{i,j}) \in \mathbb{R}^{2 \times 4}$ and $\mathbf{G} = (G_{i,j}) \in \mathbb{R}^{2 \times 4}$:

$$\mathbf{R} = \begin{matrix} & \begin{matrix} t_1 & t_2 \end{matrix} \\ \begin{matrix} 1 \\ 2 \\ 3 \\ 4 \end{matrix} & \begin{bmatrix} 0 & 1 \\ 1 & 0 \\ 1 & 0 \\ 0 & 1 \end{bmatrix} \end{matrix} \cdot \begin{matrix} & \begin{matrix} A & B & C & D \end{matrix} \\ \begin{matrix} t_1 \\ t_2 \end{matrix} & \begin{bmatrix} 1 & 1 & 0 & 0 \\ 0 & 0 & 1 & 1 \end{bmatrix} \end{matrix} = \mathbf{G}^T \mathbf{M}, \quad (2)$$

where superscript T denotes transpose and t_1 and t_2 can be interpreted as two types of association (as in the task).

Let us encode each of the four stimuli A, B, C, and D, by a vector \mathbf{s} equal to one of the four unit vectors in \mathbb{R}^4 , respectively (e.g., B is encoded by $(0, 1, 0, 0)$). This vector represents activity across a population of neurons representing the input stimuli, but the model uses only simple connectionist threshold units and does not include the intrinsic dynamics of real neurons. Similarly, each of the contexts 1, 2, 3, and 4, are encoded by one vector \mathbf{c} equal to one of the four unit vectors in \mathbb{R}^4 , respectively. Then, the correct response can be computed from \mathbf{R} as $\mathbf{c}^T \mathbf{R} \mathbf{s} = (\mathbf{Gc})^T \mathbf{M} \mathbf{s}$ and is 1 for Y and 0 for X.

To find the low-rank factorization of \mathbf{R} , one can refer to the collaborative filter algorithm in recommender systems (Murphy, 2012). As we will see, the low-rank property of the solution leads to a more explainable solution and better generalization ability. To provide some intuition through an example, the recommender system can start with a general matrix mapping a set of movies to a set of individual viewer preferences, and then factorize this general matrix into two separate low-rank matrices. One of these low-rank matrices maps individual viewers to preferred movie categories, while the other low-rank matrix allows mapping of individual movies to movie categories and can thus provide a prediction of viewer preferences. In our application, the low-rank matrix factorization allows each context to be associated with a specific set of weights representing a rule which then allows each stimulus to be mapped to a response. Similar to the prediction of viewer preferences based on mapping movies to previously learned movie categories, the correct response to a stimulus can be

learned by mapping the context of that stimulus to a specific context rule, and then the mapping can be generalized to other stimuli from that same context. In the work presented here, matrix factorization similar to the recommender system factorization is obtained by interleaved training of different orders of presentation, proposed to result from replay of neural representations in the hippocampus as described next.

2.3 | Model and learning

We next develop a neural circuit model inspired by the factorization of the response matrix R . The model employs two neural circuit models working together. The interaction of these circuits could occur during learning when stimuli are present, but would be further enhanced by

the forward and backward replay of neural representations during quiet waking and sleep (Carr et al., 2011; Diba & Buzsáki, 2007; Karlsson & Frank, 2009; Lee & Wilson, 2002). We denote the weight matrices of the two circuits by M and G (cf. Equation (2)). Circuit M processes the stimulus s and circuit G the context c .

2.3.1 | Primal circuit

We first consider a neural circuit which first processes the neural representation of the stimulus and then the neural representation of the context. This could be considered as one direction of replay of examples of the component stimulus and context events in the hippocampus. The computation $(Gc)^T Ms$ we referred to earlier is organized as follows. First, Ms is performed as:

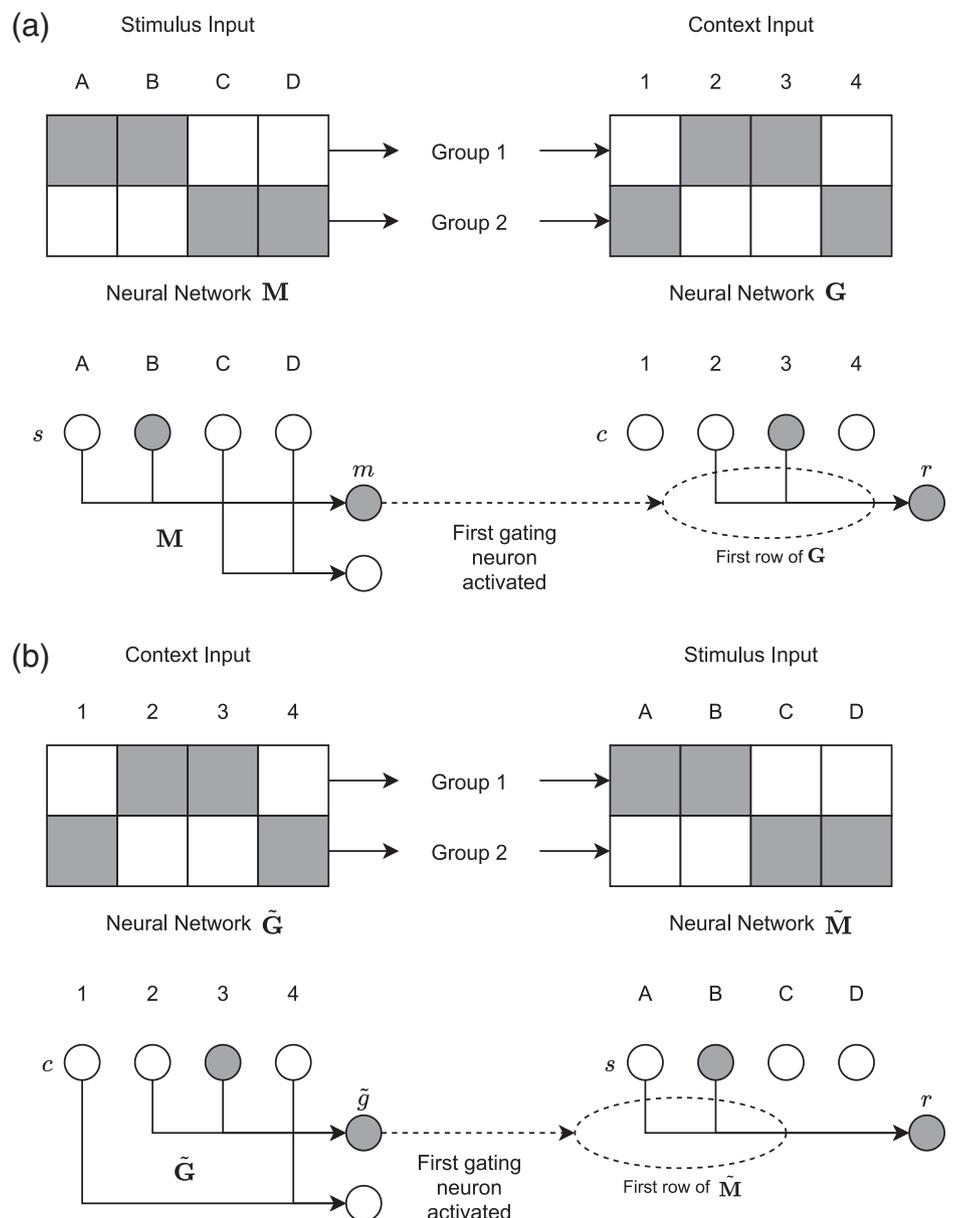


FIGURE 3 The neural circuit models under the two different update rules. Dark cells indicate activation. (a) Neural circuit model with the stimulus input first followed by context (primal order). For example, consider input B3. Stimulus B activates the first (top) cell of the B column in the left network M , which selectively gates the synaptic connectivity in network G to activate one set of synapses (Group 1 = Row 1) that mediate the influence of context on associations. Then, the input of Context 3 in the right network G has been gated by Group 1, so that Context 3 activates the first (top) cell of the third column in the right circuit and generates the output response $r = 1$, that is, Y. (b) Neural circuit model with the context input first followed by stimulus input (dual order). In the dual order, the context is presented first and processed by network G . This selects an appropriate group that gates the processing of the stimulus by network M

$$m_i = \sum_{j=1}^4 M_{ij}s_j + \mu\sigma_i, \quad i = 1, \dots, p, \quad (3)$$

where m_i represents the activation of a postsynaptic neuron, p is the number of hidden neurons, and σ_i represents environmental noise assumed uniform in $[0, 1]$, with noise gain $0 < \mu < 1$. This population of hidden neurons essentially represents the different elements of the generalized rule. Then, the neuron $k = \operatorname{argmax}_i m_i$ with the highest activation becomes the selected neuron. The activity of this neuron acts as a gate for the circuit corresponding to \mathbf{G} to process the context \mathbf{c} as $\hat{r} = \sum_{j=1}^4 \tilde{G}_{kj}c_j$, where \hat{r} denotes the activation of a postsynaptic neuron. The activation level \hat{r} indicates the probability of selecting response Y. Specifically, action X is selected if $\hat{r} < 0.5$, and action Y otherwise. A diagram of this kind of stimulus-first forward propagation rule, to which we refer as the *primal circuit*, is shown in Figure 3a.

A Hebbian learning rule is applied to the neural circuit model for obtaining the best weight matrices \mathbf{M} and \mathbf{G} after observing some training examples. The Hebbian learning rule is applied after each individual instance, which includes sequential presentation of a stimulus input \mathbf{s} , a context input \mathbf{c} and the correct response r . The update of the weights in \mathbf{M} and \mathbf{G} is performed as follows. Suppose the gating neuron k is activated after applying Equation (3). Let $i_s = \operatorname{argmax}_{j=1, \dots, 4} s_j$ be the index of the input stimulus and $i_c = \operatorname{argmax}_{j=1, \dots, 4} c_j$ the input context index. If the final response is correct, then the following learning rule is applied, analogous to Hebbian *Long-Term Potentiation (LTP)*:

$$M_{k,i_s} = M_{k,i_s} + \alpha, \quad G_{k,i_c} = G_{k,i_c} + \alpha, \quad (4)$$

where $\alpha \in (0, 1)$ is a learning rate. Otherwise, we perform an update analogous to *Long-Term Depression (LTD)*:

$$M_{k,i_s} = M_{k,i_s} - \alpha, \quad G_{k,i_c} = G_{k,i_c} - \alpha. \quad (5)$$

Finally, we project the elements of \mathbf{M} and \mathbf{G} onto $[0, 1]$, that is, any element larger than 1 is set to 1 and any element smaller than 0 is set to 0; the matrix elements with value greater than 0 and smaller than 1 are not modified.

The above algorithm can be viewed as a process of either LTD or LTP modulated by the error signal, based on whether the output response r on a given trial is correct or not. These effects could correspond to neuromodulation of the mechanisms of LTD or LTP (Adams, Winterer, & Müller, 2004; Blitzer, Gil, & Landau, 1990; Bröcher, Artola, & Singer, 1992; Hasselmo, Schnell, & Barkai, 1995). We can establish the following convergence guarantee for the learning process.

Theorem 1 *The neural circuit model converges to an optimal state under the primal-only update if each input (stimulus-context pair) is sampled uniformly, noise is independent across time, and all examples are provided during training.*

The proof to the above theorem is almost the same as the proof to Theorem 12 in our previous work (Zhu, Paschalidis, & Hasselmo, 2019); hence, we omit the proof.

2.3.2 | Dual circuit and primal-dual learning

Next we consider a symmetrical way of computing the response. This could be considered to arise from an alternate order of replay of events in the hippocampus, with context input occurring before stimulus input. Specifically, the response can be written as $r = (\mathbf{Gc})^T \mathbf{Ms} = (\mathbf{Ms})^T (\mathbf{Gc})$, which suggests the following circuit we call *dual*. In the dual circuit, we will use $\tilde{\mathbf{G}}$ and $\tilde{\mathbf{M}}$ to denote the weight matrices. The context is processed first to determine activation levels of gating units, by computing $\tilde{\mathbf{Gc}}$. The update is:

$$\tilde{g}_i = \sum_{j=1}^4 \tilde{G}_{ij}c_j + \mu\sigma_i, \quad (6)$$

where \tilde{g}_i represents the activation level of a gating neuron. The gating neuron $l = \operatorname{argmax}_i \tilde{g}_i$ with the highest activation potential becomes the selected neuron, gating the neural circuit corresponding to $\tilde{\mathbf{M}}$, which gates the stimulus input and produces the response $\hat{r} = \sum_{j=1}^4 \tilde{M}_{lj}s_j$, interpreted as the probability of selecting Y (see the diagram in Figure 3b).

To simultaneously learn the elements of \mathbf{G} , \mathbf{M} in the primal circuit and $\tilde{\mathbf{G}}$, $\tilde{\mathbf{M}}$ in the dual circuit (all initialized with elements in $[0, 1]$), we can use Hebbian updates analogous to Equations (4) and (5). In the following, we combine the primal and dual circuit computations in what we call a Hebbian *primal-dual* learning. Each stimulus-context pair (\mathbf{s}, \mathbf{c}) , with corresponding indices i_s , i_c (defined earlier) and correct response r , is first processed by the primal circuit (cf. Equation (3)). Using k as the index of the activated gating neuron, we update:

$$G_{k,i_c} = G_{k,i_c} + \alpha \operatorname{sgn}(r - G_{k,i_c}), \quad (7)$$

where $\operatorname{sgn}(x)$ is the sign of x . The input is also processed using the dual circuit (cf. Equation (6)) which activates gating neuron l . We update:

$$\tilde{M}_{l,i_s} = \tilde{M}_{l,i_s} + \alpha \operatorname{sgn}(r - \tilde{M}_{l,i_s}). \quad (8)$$

In both Equations (7) and (8), the argument of the sgn function is $r - \hat{r}$, the difference between the correct and predicted response. After each primal-dual (Equations (7)–(8)) update, we project the elements in $\tilde{\mathbf{M}}$ and \mathbf{G} onto $[0, 1]$. The algorithm keeps performing one primal and one dual update for each input presented. After K inputs processed, the primal and dual weights are synchronized:

$$\mathbf{M} = \tilde{\mathbf{M}}, \quad \tilde{\mathbf{G}} = \mathbf{G}. \quad (9)$$

We call K the *synchronization period* of the update.

The following result (shown in Data S1) provides a guarantee of convergence.

Theorem 2 *The neural circuit model converges to an optimal state under the primal-dual update with $K = 1$ if each stimulus-context pair is sampled uniformly, noise is independent across time, and all examples are provided during training.*

3 | RESULTS

3.1 | Simulation results

We use simulations to evaluate the performance of the proposed models. In a first simple example we consider the behavioral task shown in Figure 1, and set the number of the gating neurons to $p = 2$ and the synchronization period $K = 1$. The accuracy of the neural circuit models, measured by the number of stimulus-context pairs correctly processed in testing, is shown in Figure 4 as a function of the number of learning iterations. The accuracy increases during the training process and converges to 16 accurately identified pairings.

We next investigate the capacity of the neural circuits to generalize to stimulus-context pairs not seen during training. For both tasks shown in Figures 1 and 2, we fix the learning rate to $\alpha = 0.2$, the noise gain to $\mu = 0.99$, and vary the number of hidden neurons p . The neural circuit models are initialized with the value 0.5 in all elements of the two matrices \mathbf{M} and \mathbf{G} , and are trained until convergence. For each instance m , we perform 1,000 trials and report the average accuracy (normalized to 1) of the model for all the stimulus-context pairs, including the hidden ones (which were not used for training). For the behavioral task shown in Figure 1, the underlined inputs are hidden during training. For this task, the generalization results are shown in Figure 5a. In light of Thm. 2, which requires that all inputs are used during training, it is interesting that primal-dual learning yields 100% accuracy. Apparently, even though it is possible for the neural circuit to converge to a non-optimal state, the corresponding probability is very low as quantified by the simulation results in Figure 5a (square symbols) that show no instances of non-optimal states. On the other hand, primal-only learning does not lead to good generalization except when there are only two gating units (the minimum required for the task).

Figure 3 shows how the network generates a correct response for a previously seen stimulus-context pair B3 (stimulus B, Context 3).

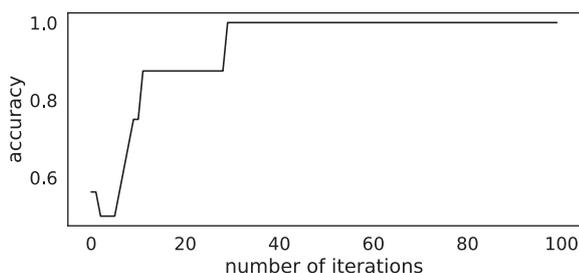


FIGURE 4 Accuracy of a primal-dual neural circuit

The network also generalizes to hidden stimulus-context pairs, such as A1, which is underlined in Figure 1. A1 is not seen during learning, but the presentation of other stimuli have already created the connectivity shown in Figure 3. For example, in the primal model, learning of A2, A3, and A4 results in stimulus A activating Group 1, and learning of B1 results in Group 1 gating stimulus 1 to generate output X ($r = 0$, white). In the dual model, learning of B1, C1, and D1 results in Context 1 leading to Group 1, and learning of A4 results in gating by Group 1 to include A going to X ($r = 0$, white). Thus, when A1 is presented for the first time after learning, in the primal model, input of stimulus A activates the gating unit for Group 1, and this gating of row one of connectivity ensures that input of Context 1 results in output X ($r = 0$, white). Thus, the gating process allows stimuli and context input to gate activity that allows generalization for correct responses to stimulus-context pairs not seen during learning.

To better understand why increasing the number of hidden neurons does not impact performance for primal-dual learning, we run the algorithm for $p = 6$ and $K = 1$. The algorithm converges to the following matrices:

$$\mathbf{M} = \begin{bmatrix} 1 & 1 & 0 & 0 \\ 1 & 1 & 0 & 0 \\ 0 & 0 & 1 & 1 \\ 1 & 1 & 0 & 0 \\ 0 & 0 & 1 & 1 \\ 0 & 0 & 1 & 1 \end{bmatrix}, \quad \mathbf{G} = \begin{bmatrix} 0 & 1 & 1 & 0 \\ 0 & 1 & 1 & 0 \\ 1 & 0 & 0 & 1 \\ 0 & 1 & 1 & 0 \\ 1 & 0 & 0 & 1 \\ 1 & 0 & 0 & 1 \end{bmatrix}, \quad (10)$$

where Rows 1, 2, and 4 are identical and correspond to Group 1 (cf. Figure 3a) and Rows 3, 5, and 6 correspond to Group 2. It is evident that primal-dual learning enforces a low-rank property and simply replicates the appropriate entries when the number of hidden neurons exceeds the minimum necessary.

We separately tested broader generalization in a different variant of the task with a larger number of different contexts as shown in Figure 2, which has 16 contexts rather than the four contexts shown in the task in Figure 1. For testing this broader generalization, there were two stages of training, one stage to learn the association rules and the second stage to determine which rule should be applied to each context. During the initial stage of training to learn the rules, only examples in Contexts 1–4 are provided during training. As with the smaller task of Figure 1, one context-stimulus pair for Contexts 1–4 is hidden during training. After the neural circuit converges on the four contexts, we implemented a second stage of training in which we randomly choose one context-stimulus example from each of Contexts 5–16 and present them 10 times repeatedly so that the model can learn from them. We then evaluate the generalization performance using all inputs from the 16 contexts. Thus, in this different variant of the task, the network has been trained on less than 1/3 of the inputs across the 16 contexts. The result is shown in Figure 5b. Again, primal-dual learning exhibits a very high generalization ability, whereas primal-only learning merely remembers what was previously learned.

To explore the convergence properties of the two types of Hebbian learning rules with respect to the hyperparameters, we trained our neural

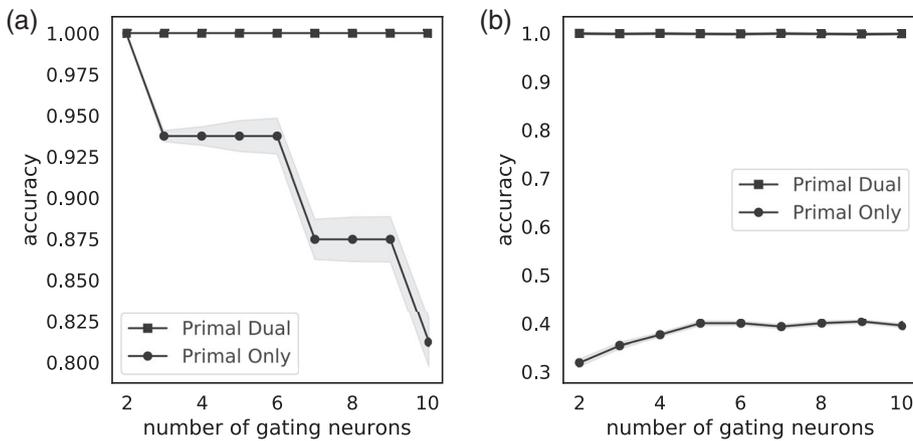


FIGURE 5 Comparison of generalization accuracy (normalized to 1). (a) The primal-dual model (top, square symbols) shows perfect generalization accuracy for the context-association task with four contexts as shown in Figure 1. (b) The primal-dual model (top, square symbols) shows perfect accuracy for the task variant testing broader generalization to 16 contexts as shown in Figure 2

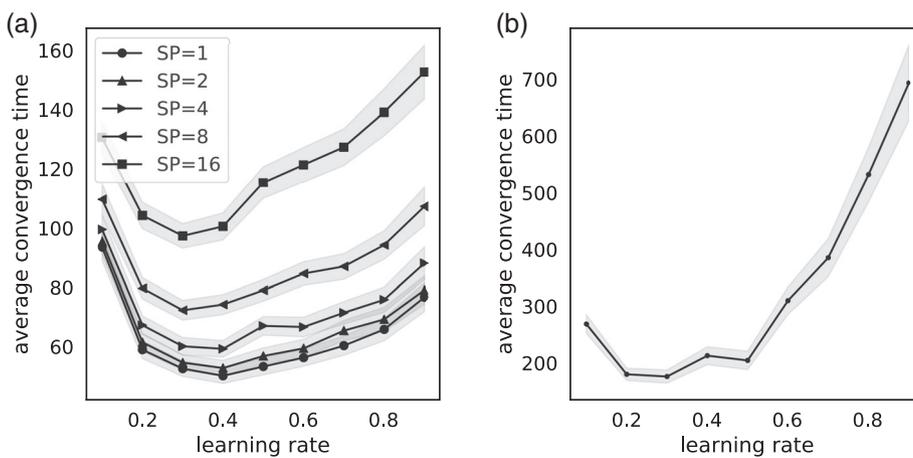


FIGURE 6 Average convergence time as a function of learning rate for primal-dual training (a) and primal-only training (b)

circuit models with different hyperparameters. Under each hyperparameter, we run the simulation for 1,000 trials and report the average convergence time of the model. The convergence time is defined as the number of learning iterations before the neural circuit reaches the state that produces all correct actions on all training input pairs.

We considered the task of Figure 1 and first evaluated the effect of the learning rate α . We fixed the noise level to $\mu = 0.99$ and the number of hidden neurons to $p = 2$. We also varied the synchronization period K (abbreviated as SP in Figure 6a) in the primal-dual learning. The results are shown in Figures 6a,b. For both types of learning, there is an optimal learning rate (not too small and not too large) minimizing convergence time. The convergence time also increases with K . Finally, we observed that primal-dual training requires a much smaller amount of time to converge compared with primal-only training.

Then, we explored the dependence on the noise gain μ . We set $\alpha = 0.2$ and $p = 2$ during training. Figures 7a,b depict average convergence time as a function of noise gain. When the noise gain is high, the neural circuit models can be easily misled by noise and oscillate near the optimal states. But when the noise level is low, it is possible for a neural circuit to too strongly trust what it has learned and become “stuck” in a local minimum. (Notice that convergence time can increase as μ approaches zero for both learning modes.) Again a huge difference in the convergence times is shown for primal-dual versus primal-only training.

Finally, we evaluate the convergence speed with respect to the number of hidden neurons. We set $\alpha = 0.2$ and $\mu = 0.99$ during training. The results are shown in Figure 8a,b. When the number of hidden neurons increases, the model complexity increases. Therefore, it takes more time for the neural circuits to converge.

3.2 | Experimental results

Based on our analysis and simulation results, we hypothesize that humans who do well in the context-association tasks are able to learn a low-rank synaptic matrix structure relevant to storing information in the human cortex that could thereby help achieve better generalization to a wider range of contexts. In this section, we present an analysis of behavioral data from human subjects performing the task of Figure 1; testing the hypothesis of low-rank response matrix structure.

3.2.1 | Participants

The experiments involved collection of behavioral data from the work in (Chang et al., 2017, 2019). There were 70 valid healthy young participants in the experiments. These participants were recruited from

FIGURE 7 Average convergence time as a function of noise gain for primal-dual training (a) and primal-only training (b)

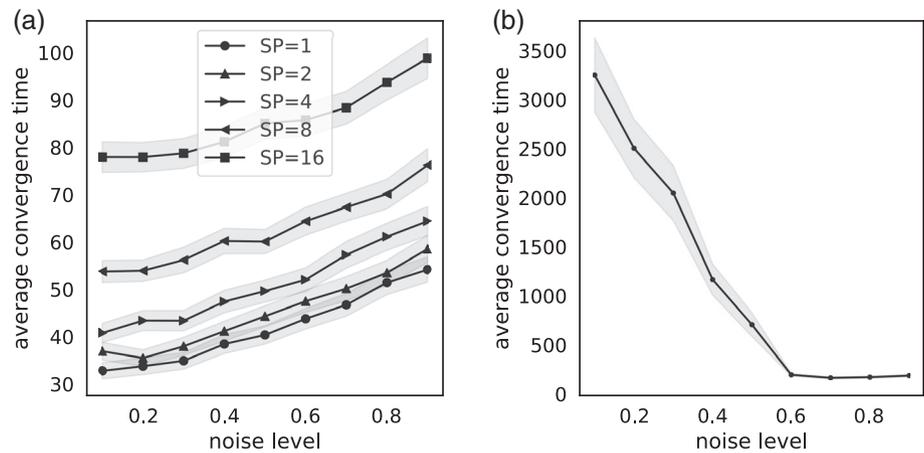
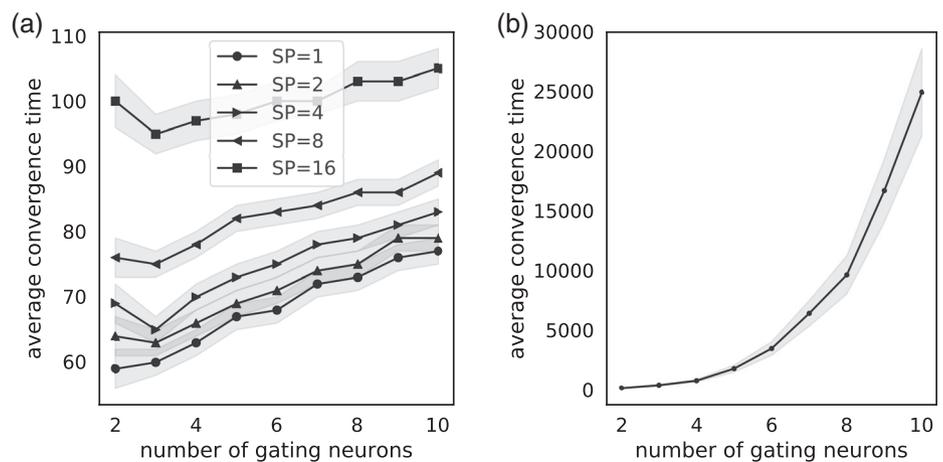


FIGURE 8 Average convergence time as a function of the number of hidden neurons for primal-dual training (a) and primal-only training (b)



Boston University and received monetary compensation for their participation. All experimental procedures were approved by the Boston University Institutional Review Board. All participants were cognitively intact, provided written informed consent, and at the conclusion of the experiment were debriefed.

3.2.2 | Task procedure

For the human participants, the task stimuli were realistic visual objects selected from a public database (Brady, Konkle, Alvarez, & Oliva, 2008). Among the visual objects, five of them are used as the first (cue) stimuli (A, B, C, D, E) and three of them are used as the second stimuli (X, Y, Z) such that the association of the two stimuli guides the response of the participant. The association rule for cue stimuli A, B, C, and D, in four contexts, and second stimuli X and Y that guide responses, was as in Figure 3. Stimulus E was always associated with second stimulus Z (the “easy rule”).

During the task, cue and associate stimulus in one of four quadrants on a screen were sequentially presented to the participants. The four quadrants of the screen represent the four contexts. Each input was presented on the screen for 1,500 ms and another 1,500 ms was given for participants to determine if the presented stimulus was

associated with the presented second stimulus (“match” or “no match”). After the participant's response, the correct response was provided for 1,000 ms. Finally, there was a 500 ms interval between repeated trials.

3.2.3 | Training condition

Two types of training sequences were used for the human subjects. The first one was a *context-spaced* training condition. In this condition, one cue (e.g., A) was presented with corresponding second stimulus (X or Y) in *all* spatial quadrants (1, 2, 3, 4) before a new cue (e.g., B) and its permutations of quadrants and second stimuli were presented. The order of cue objects was randomly generated for each block, and trials within each chunk were randomly generated as well. The four simple rule trials were randomly distributed throughout the block of 36 trials.

The second training condition was a *context-massed* training condition. In this condition, all cues (A, B, C, D) and associate stimuli (X, Y) permutations were presented in a single quadrant (e.g., 1) before moving to a new quadrant (e.g., 4) and its permutations of cues and associate stimuli. Again, the order of spatial quadrants was randomly generated for each block, and trials within each chunk were randomly generated as well.

3.2.4 | Response matrix

Each subject completed 320 trials. One trial in both training conditions consists of 10 loops, in which a complete set of context-stimulus-second stimulus triples are enumerated using the two training conditions. For each participant, a response matrix is constructed using all responses in a single loop as follows. If the subject reports a match then the response of the (stimulus-context-second stimulus) triplet shown to the subject is recorded as the response, otherwise (if no-match is reported), then the complementary response is

recorded. We sum across all trials to form the response matrix. For example, if (A-1-Y-match) and (A-1-X-no match) are in a loop, then the response matrix \mathbf{R} has $R_{1,1} = 2$ (response Y counts for 1 and response X for 0). If (A-1-Y-match) and (A-1-X-match) is observed, then $R_{1,1} = 1$. If (A-1-Y-no match) and (A-1-X-match) is observed, then $R_{1,1} = 0$.

The rank of the response matrices and the corresponding accuracy is shown in Figure 9. During the training process, the average accuracy increases while the rank of the response matrix decreases. The relationship between the matrix rank and the accuracy is shown in Figure 9a. It is clear that the lower the rank of the response matrices is, the higher the accuracy on the behavioral task was.

We also segregated the subjects into two cohorts—the good learners and the less effective learners—and computed the average rank of the response matrices in each cohort. We consider a subject as a good learner if his/her accuracy during the last three loops is above the median, and a less effective learner if this accuracy is below the median. We calculated the mean of the response matrices of good and less effective learners in their last three loops. The good learners have an average rank of 2.92 and the less effective learners have an average rank of 3.59, with a p -value of less 1% (using a t test across the two sets of samples). This suggests that the rank of the response matrix is the critical factor for good performance in the context association task and our primal-dual learning algorithm, which efficiently discovers the low-rank property of the response matrix, is useful for understanding the performance in the task.

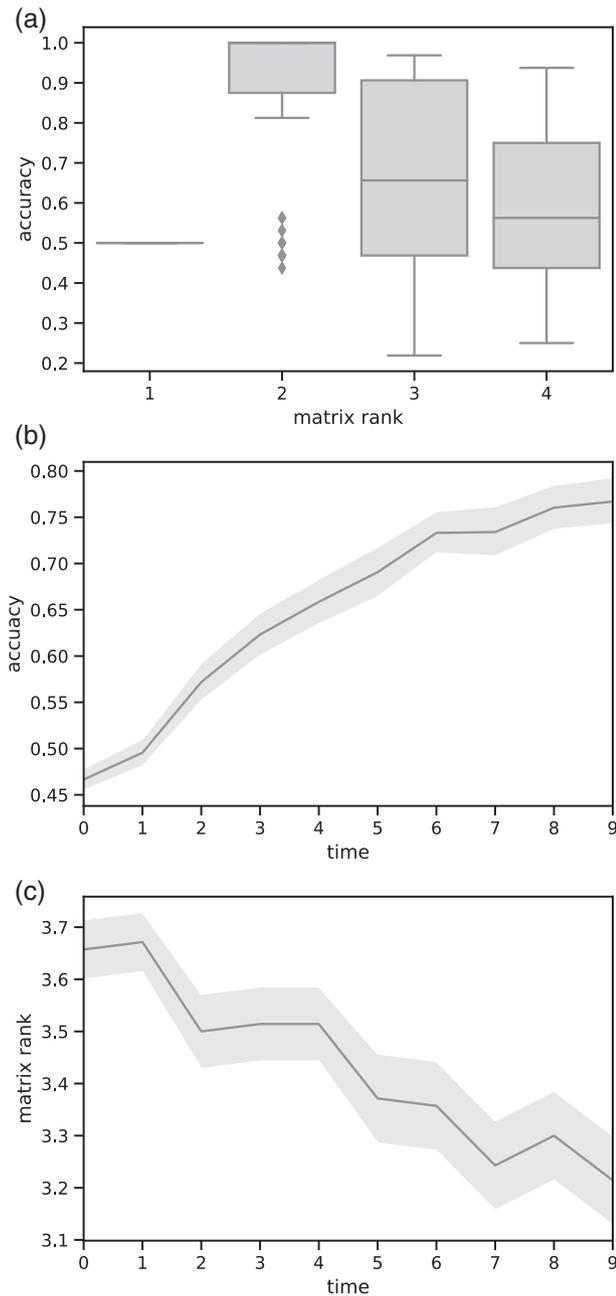


FIGURE 9 Experimental results. (a) Relationship between response matrix rank and accuracy for subjects trained in the context-space mode. (b) Average accuracy of the subjects through time. (c) Average rank of the response matrix of the subjects through time

4 | DISCUSSION

In this article, we address potential mechanisms for learning a set of context-dependent rules involving a symmetrical set of associations between specific cues and specific responses, and generalizing to a number of hidden stimulus-context pairs. In addition, in a different variant of the task, we demonstrated that the model can perform with high accuracy when tested on broader generalization of these context-dependent rules to a large number of additional contexts when an example association is presented in each new context. A simplified neural gating model presented here uses the recommender system framework to perform matrix factorization that maps each example of a contextual association to the correct response and generalizes with perfect accuracy. This illustrates how gating units that regulate network connections can be used to learn the initial set of rules concerning associations in a subset of contexts, and to effectively generalize to a larger set of contexts. Behavioral data from this task were analyzed in terms of the rank of the matrix that describes the responses of the participants. The analysis shows that the formation of a lower rank matrix was associated with better performance in the behavioral task.

One may postulate that the symmetric nature of the tasks in both Figures 1 and 2 contributes to the generalization ability of primal-dual learning (yielding 100% accuracy as we discussed in Section 3.1). We considered additional tasks with no such symmetry. Specifically, we

considered a task with five contexts and two association rules (three contexts using Rule 1 and the remaining two contexts using Rule 2) and a task with four contexts and five stimuli. In both these tasks, primal-dual learning achieved 100% accuracy in the generalization task. This suggests that the generalization ability of primal-dual learning is rooted in the existence of a low-rank factorization of the response matrix (i.e., some association rules get repeated in some contexts) and not on the symmetry of the task.

The pattern of synaptic connectivity learned by this network (Figure 3) shows a potential mechanism for the learning of selective responses to item-context pairings, as observed in neurophysiological data (Komorowski et al., 2009). For modeling the hippocampal response, the input representing context and stimulus item in this model could be considered to represent the sensory input for context and item arriving in the entorhinal cortex after processing through multiple cortical regions. These representations then interact via matrices representing synaptic connectivity from the entorhinal cortex to hippocampal regions CA3 and CA1 and from region CA3 to CA1. The arrival of contextual information into the entorhinal cortex could activate a contextual representation in region CA3 that interacts with the synaptic input representing items. For example, a contextual representation in CA3 for Group 1 or Group 2 (Figure 3) could interact with the connectivity from the entorhinal cortex to region CA1 via multiplicative influences on synaptic input (e.g., Hasselmo and Eichenbaum (2005); Katz et al. (2007)) to activate neurons coding a conjunction of context and item stimulus that could further influence the selection of a behavioral response. In this manner, the interaction of a context input (e.g., represented by input from CA3 to CA1) with a stimulus input (represented by inputs from the entorhinal cortex to CA1) can regulate the selective firing for a specific stimulus in a specific context, as shown in neurophysiological data (Komorowski et al., 2009). The conjunctive interaction of CA3 and entorhinal input to CA1 has been used previously (Hasselmo & Eichenbaum, 2005; Katz et al., 2007) to model the mechanisms for context-dependent firing in the hippocampus during spatial alternation, but these previous simulations did not address the potential mechanism for generalizing this response to a wider range of different contexts.

The mechanisms of gating were previously explored in a neural network simulation of this task (Hasselmo & Stern, 2018) in which gating units activated randomly during learning would regulate the pattern of connectivity for subsequent sensory input. That network was able to generalize when only a small number of gating units were used, because this forced each unit during learning to cover the full range of associations within a context, allowing one association to reactivate that same gating unit to provide the full range of associations within that context. However, when larger numbers of gating units were used, this prevented accurate generalization, because subsets of associations could be coded on different gating units, preventing full generalization by reactivation of an individual gating unit. The recommender system framework presented here avoids this problem by ensuring that each gating unit (row of the matrix) is associated with a second matrix representing the full set of associations.

Future studies must address how the learning of the symmetrical associations of stimulus-response pairs in different contexts could be combined with the learning of the baseline association between stimulus E and response Z that does not change between the different contexts.

These networks rely on a framework in which the activity of individual units gate a full matrix of connectivity at the subsequent step. This gating could arise from different physiological mechanisms. In one potential mechanism, the gating unit could be an axoaxonic interneuron that directly regulates the output from a subset of pyramidal cells within the circuit (Baude, Bleasdale, Dalezios, Somogyi, & Klausberger, 2006; Cutsuridis & Hasselmo, 2012). Another potential mechanism could involve the nonlinear interaction of adjacent synapses on the dendritic tree of pyramidal cells. In this framework, the activity of a first set of gating neurons may cause synaptic currents in the dendritic tree of postsynaptic neurons that are adjacent to the synaptic inputs from a second set of neurons. The nonlinear interaction of synapses can occur due to voltage-sensitive postsynaptic channels known as N-Methyl-D-Aspartate (NMDA) receptors, as modeled extensively in previous work (Koch & Poggio, 1992; Mel, 1993; Poirazi et al., 2003). NMDA receptor channels only allow current to pass through when postsynaptic depolarization of voltage causes release of magnesium blockade of the channel that then allows presynaptic glutamate release to activate excitatory ionic current through the channel. The synapses arising from gating units could provide sufficient postsynaptic depolarization to release magnesium blockade, thereby allowing the glutamate released from the second set of synaptic inputs to cause suprathreshold excitatory currents in the postsynaptic neurons. Thus, the synapses from the second set of input neurons on the postsynaptic neuron have strength zero unless they are adjacent to an active gating synapse. In this manner, the gating unit can set the pattern of active connectivity within the circuit. Other potential mechanisms for gating could involve circuit level multiplicative interactions (Nezis & Van Rossum, 2011; Sherfey et al., 2018) that can be implemented by interacting populations of neurons (Eliasmith & Anderson, 2004) or gating by the interaction of local oscillatory dynamics of cortical circuits with the dominant frequency of input from other regions (Buschman et al., 2012; Lundqvist et al. 2018,b; Sherfey et al., 2018).

This framework for gating focuses on local interactions within neocortical circuits. In contrast, many other gating mechanisms have been proposed to involve interaction of different regions, such as the interaction of the basal ganglia with cortical circuits. These are not mutually exclusive theories. These other types of gating models could potentially be utilized to learn and generalize the contextual association task as well.

ACKNOWLEDGMENTS

Research partially supported by the Office of Naval Research under MURI grants N00014-16-1-2832 and N00014-19-1-2571, by the NSF under grants IIS-1914792, DMS-1664644, and CNS-1645681, by the NIH/NIGMS under grant 1R01GM135930, and by Philips Healthcare.

DATA AVAILABILITY STATEMENT

The data that support the findings of this study are available from C.E. Stern upon reasonable request.

ORCID

Ioannis Ch. Paschalidis  <https://orcid.org/0000-0002-3343-2913>

Allen Chang  <https://orcid.org/0000-0002-4500-4916>

Michael E. Hasselmo  <https://orcid.org/0000-0002-9925-6377>

REFERENCES

- Adams, S. V., Winterer, J., & Müller, W. (2004). Muscarinic signaling is required for spike-pairing induction of long-term potentiation at rat schaffer collateral-CA1 synapses. *Hippocampus*, 14(4), 413–416.
- Aminoff, E., Gronau, N., & Bar, M. (2006). The parahippocampal cortex mediates spatial and nonspatial associations. *Cerebral Cortex*, 17(7), 1493–1503.
- Badre, D., & Frank, M. J. (2012). Mechanisms of hierarchical reinforcement learning in cortico-striatal circuits 2: Evidence from fMRI. *Cerebral Cortex*, 22(3), 527–536.
- Badre, D., Kayser, A. S., & D'Esposito, M. (2010). Frontal cortex and the discovery of abstract action rules. *Neuron*, 66(2), 315–326.
- Bar, M., Aminoff, E., & Ishai, A. (2007). Famous faces activate contextual associations in the parahippocampal cortex. *Cerebral Cortex*, 18(6), 1233–1238.
- Baude, A., Bleasdale, C., Dalezios, Y., Somogyi, P., & Klausberger, T. (2006). Immunoreactivity for the gabaa receptor $\alpha 1$ subunit, somatostatin and connexin36 distinguishes axoaxonic, basket, and bistratified interneurons of the rat hippocampus. *Cerebral Cortex*, 17(9), 2094–2107.
- Bhandari, A., & Badre, D. (2018). Learning and transfer of working memory gating policies. *Cognition*, 172, 89–100.
- Blitzer, R. D., Gil, O., & Landau, E. M. (1990). Cholinergic stimulation enhances long-term potentiation in the CA1 region of rat hippocampus. *Neuroscience Letters*, 119(2), 207–210.
- Brady, T. F., Konkle, T., Alvarez, G. A., & Oliva, A. (2008). Visual long-term memory has a massive storage capacity for object details. *Proceedings of the National Academy of Sciences*, 105(38), 14325–14329.
- Bröcher, S., Artola, A., & Singer, W. (1992). Agonists of cholinergic and noradrenergic receptors facilitate synergistically the induction of long-term potentiation in slices of rat visual cortex. *Brain Research*, 573(1), 27–36.
- Buschman, T. J., Denovelli, E. L., Diogo, C., Bullock, D., & Miller, E. K. (2012). Synchronous oscillatory neural ensembles for rules in the prefrontal cortex. *Neuron*, 76(4), 838–846.
- Carpenter, P. A., Just, M. A., & Shell, P. (1990). What one intelligence test measures: A theoretical account of the processing in the raven progressive matrices test. *Psychological Review*, 97(3), 404.
- Carr, M. F., Jadhav, S. P., & Frank, L. M. (2011). Hippocampal replay in the awake state: A potential substrate for memory consolidation and retrieval. *Nature Neuroscience*, 14(2), 147.
- Chang, A., Johnson, A., Whiteman, A. S., & Stern, C. E. (2019). A latent structure analysis for understanding individual differences in context-dependent rule learning in humans. *Hippocampus in review*.
- Chang, A., Ren, Y., Whiteman, A. S., & Stern, C. E. (2017). Evaluating the relationship between activity in frontostriatal regions and successful and unsuccessful performance of a context-dependent rule learning task. In *Society for Neuroscience Abstracts*. Society for Neuroscience.
- Chatham, C. H., Frank, M. J., & Badre, D. (2014). Corticostriatal output gating during selection from working memory. *Neuron*, 81(4), 930–942.
- Cutsuridis, V., & Hasselmo, M. (2012). Gabaergic contributions to gating, timing, and phase precession of hippocampal neuronal activity during theta oscillations. *Hippocampus*, 22(7), 1597–1621.
- Diba, K., & Buzsáki, G. (2007). Forward and reverse hippocampal place-cell sequences during ripples. *Nature Neuroscience*, 10(10), 1241.
- Eliasmith, C., & Anderson, C. H. (2004). *Neural engineering: Computation, representation, and dynamics in neurobiological systems*. Cambridge, Massachusetts: MIT press.
- Eliasmith, C., Stewart, T. C., Choo, X., Bekolay, T., DeWolf, T., Tang, Y., & Rasmussen, D. (2012). A large-scale model of the functioning brain. *Science*, 338(6111), 1202–1205.
- Hasselmo, M. E. (2005). ISSN 0898-929X). A model of prefrontal cortical mechanisms for goal-directed behavior. *Journal of Cognitive Neuroscience*, 17(7), 1115–1129. <https://doi.org/10.1162/0898929054475190>
- Hasselmo, M. E., & Eichenbaum, H. (2005). Hippocampal mechanisms for the context-dependent retrieval of episodes. *Neural Networks*, 18(9), 1172–1190.
- Hasselmo, M. E., Schnell, E., & Barkai, E. (1995). Dynamics of learning and recall at excitatory recurrent synapses and cholinergic modulation in rat hippocampal region CA3. *Journal of Neuroscience*, 15(7), 5249–5262.
- Hasselmo, M. E., & Stern, C. E. (2018). A network model of behavioural performance in a rule learning task. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 373(1744), 1–10. <https://doi.org/10.1098/rstb.2017.0275>
- Hummel, J. E., & Holyoak, K. J. (1997). Distributed representations of structure: A theory of analogical access and mapping. *Psychological Review*, 104(3), 427.
- Karlsson, M. P., & Frank, L. M. (2009). Awake replay of remote experiences in the hippocampus. *Nature Neuroscience*, 12(7), 913.
- Katz, Y., Kath, W. L., Spruston, N., & Hasselmo, M. E. (2007). Coincidence detection of place and temporal context in a network model of spiking hippocampal neurons. *PLoS Computational Biology*, 3(12), e234.
- Koch, C., & Poggio, T. (1992). Multiplying with synapses and neurons. In: T. M. McKenna, J. L. Davis, & S. F. Zornetzer (eds.) *Single neuron computation* (pp. 315–345). San Diego, CA: Academic Press Inc.
- Komorowski, R. W., Garcia, C. G., Wilson, A., Hattori, S., Howard, M. W., & Eichenbaum, H. (2013). Ventral hippocampal neurons are shaped by experience to represent behaviorally relevant contexts. *Journal of Neuroscience*, 33(18), 8079–8087.
- Komorowski, R. W., Manns, J. R., & Eichenbaum, H. (2009). Robust conjunctive item–place coding by hippocampal neurons parallels learning what happens where. *Journal of Neuroscience*, 29(31), 9918–9929.
- Lee, A. K., & Wilson, M. A. (2002). Memory of sequential experience in the hippocampus during slow wave sleep. *Neuron*, 36(6), 1183–1194.
- Lundqvist, M., Herman, P., & Miller, E. K. (2018). Working memory: Delay activity, yes! Persistent activity? Maybe not. *Journal of Neuroscience*, 38(32), 7013–7019.
- Lundqvist, M., Herman, P., Warden, M. R., Brincat, S. L., & Miller, E. K. (2018). Gamma and beta bursts during working memory readout suggest roles in its volitional control. *Nature Communications*, 9(1), 394.
- Mel, B. W. (1993). Synaptic integration in an excitable dendritic tree. *Journal of Neurophysiology*, 70(3), 1086–1101.
- Melrose, R. J., Poulin, R. M., & Stern, C. E. (2007). An fMRI investigation of the role of the basal ganglia in reasoning. *Brain Research*, 1142, 146–158.
- Miller, E. K., & Cohen, J. D. (2001). An integrative theory of prefrontal cortex function. *Annual Review of Neuroscience*, 24(1), 167–202.
- Murphy, K. P. (2012). *Machine learning: A probabilistic perspective*. Cambridge, Massachusetts: MIT Press.
- Nezis, P., & Van Rossum, M. C. (2011). Accurate multiplication with noisy spiking neurons. *Journal of Neural Engineering*, 8(3), 034005.
- O'Reilly, R. C. (1998). Six principles for biologically based computational models of cortical cognition. *Trends in Cognitive Sciences*, 2(11), 455–462.
- O'Reilly, R. C., Hazy, T. E., & Herd, S. A. (2016). The leabra cognitive architecture: How to play 20 principles with nature. *The Oxford Handbook of Cognitive Science*, 91, 1–31.
- Poirazi, P., Brannon, T., & Mel, B. W. (2003). Arithmetic of subthreshold synaptic summation in a model CA1 pyramidal cell. *Neuron*, 37(6), 977–987.

- Rasmussen, D., & Eliasmith, C. (2011). A neural model of rule generation in inductive reasoning. *Topics in Cognitive Science*, 3(1), 140–153.
- Raudies, F., & Hasselmo, M. E. (2017). A model of symbolic processing in Raven's progressive matrices. *Biologically Inspired Cognitive Architectures*, 21, 47–58.
- Santoro, A., Hill, F., Barrett, D., Morcos, A., & Lillicrap, T. (2018). Measuring abstract reasoning in neural networks. *International Conference on Machine Learning*, 4477–4486.
- Sherfey, J. S., Ardid, S., Hass, J., Hasselmo, M. E., & Kopell, N. J. (2018). Flexible resonance in prefrontal networks with strong feedback inhibition. *PLoS Computational Biology*, 14(8), e1006357.
- Wallis, J. D., Anderson, K. C., & Miller, E. K. (2001). Single neurons in prefrontal cortex encode abstract rules. *Nature*, 411(6840), 953.
- Zhu, H., Paschalidis, I. C., & Hasselmo, M. E. (2018). Neural circuits for learning context-dependent associations of stimuli. *Neural Networks*, 107, 48–60. <https://doi.org/10.1016/j.neunet.2018.07.018>, <http://www.sciencedirect.com/science/article/pii/S089360801830217X>
- Zhu, H., Paschalidis, I. C., & Hasselmo, M. E. (2019). A Hebbian learning algorithm for training a neural circuit to perform context-dependent associations of stimuli. In *2019 American Control Conference (ACC)*, 848–853. IEEE

SUPPORTING INFORMATION

Additional supporting information may be found online in the Supporting Information section at the end of this article.

How to cite this article: Zhu H, Paschalidis IC, Chang A, Stern CE, Hasselmo ME. A neural circuit model for a contextual association task inspired by recommender systems. *Hippocampus*. 2020;30:384–395. <https://doi.org/10.1002/hipo.23194>