Modeling Multiplicity of Strategies in Free Recall with Neural Networks

Moufan Li (moufan.li@nyu.edu)
Department of Psychology, New York University
6 Washington Place, New York, NY 10009, USA

Kristopher T. Jensen (kris.jensen@ucl.ac.uk)
Sainsbury Wellcome Centre, University College London
25 Howland Street, London W1T 4JG, UK

Qihong Lu (qihong.lu@columbia.edu)
The Mortimer B. Zuckerman Mind Brain Behavior Institute
Center for Theoretical Neuroscience, Columbia University
3227 Broadway, New York, NY 10027, USA

Qiong Zhang (qiong.z@rutgers.edu)
Department of Psychology, Rutgers University
152 Frelinghuysen Road, Piscataway, NJ 08854, USA

Marcelo G. Mattar (marcelo.mattar@nyu.edu)
Department of Psychology, New York University
6 Washington Place, New York, NY 10009, USA

Abstract

Humans preferentially recall items that are presented in close temporal proximity together – a phenomenon known as the 'temporal contiguity effect'. In this study, we investigated how this phenomenon emerges naturally when training a recurrent neural network with episodic memory on free recall tasks, and the neural mechanisms underlying this process. The model managed to produce the temporal contiguity effect, and we found individual differences in neural mechanisms for different models. Some models learned an item index code that matches the 'memory palace' technique and recalled in a forward order, while the other models learned to recall in a backward order and relied more on item-related temporal context. We found that the extent to which the model changes the context between encoding and recalling memories affects the learned recall strategy. Our findings provide insights into how different memory strategies may arise in human free recall.

Keywords: neural network; episodic memory; free recall; temporal contiguity effect

Introduction

Humans exhibit the temporal contiguity effect in free recall tasks, where recalled items tend to be at nearby serial positions from the previously recalled item (Kahana, 1996). The temporal context model (TCM) (Howard & Kahana, 2002; Polyn, Norman, & Kahana, 2009) explains this by proposing a slowly drifting temporal context that retains information about previously studied items. This can lead to temporal contiguity as items near the previously recalled item are more likely to be recalled due to the similarity between their context. While TCM accounts for the recall patterns in most participants, some individuals use techniques like the memory palace or 'method of loci' (Yates, 2013) to enhance their memory performance. This involves encoding items at specific locations within a mentally constructed environment and recalling memories by following a stereotypical route. This strategy facilitates forward recall, which has been found to be an optimal solution for free recall (Zhang, Griffiths, & Norman, 2023).

In this work, we trained neural networks to study possible neural mechanisms underlying the temporal contiguity effect. The models exhibit the temporal contiguity effect, consistent with human behavioral data, while the training process does not explicitly encourage this behavior. We found different memory strategies in different individual models, with some models exhibiting the memory palace technique while others relying more on a slowly drifting temporal context. We then investigated factors that affect the learned strategies and found that the noise injected between memory encoding and recall can make the model switch between the two strategies. Our results provide insights into how different memory strategies may arise in human free recall.

Methods

We used a free-recall task to train the models. There was an encoding phase and a recall phase in each trial. During the encoding phase, the model received a list of items as one-hot vectors. In the recall phase, it had to recall all items in any order. Correct recalls were rewarded (+1), while incorrect or repeated items were penalized (-1).

The model was composed of a 'context module' with 128 gated recurrent units (GRUs) (Cho et al., 2014) and a 'memory module' with a list of slots for storing memories. During encoding, the GRU received a one-hot input as an item, updated its hidden state, and appended this updated hidden state to the memory module as an episodic memory at each time step (Figure 1a). During recall, the model retrieved a weighted average of memories from the memory module based on cosine similarity between the memories and the current hidden state, then used the retrieved memory as an input to update the GRU and produce an output policy over all possible items (Figure 1b). This weighted averaged memory was used instead of the single most similar memory to make the memory retrieval process differentiable during training (Graves, Wayne, & Danihelka, 2014). The memory module was cleared at the start of each trial to allow the storage of new memories.

We trained the model with the advantage actor-critic (A2C) reinforcement learning algorithm (Mnih et al., 2016; Jensen, 2023). The last reward and last action (recalled item) sampled from the policy were returned to the model as inputs at the next time step (Wang et al., 2016). To encourage the model to recall only one memory at a time and inhibit the retrieval of other memories, we added a negative entropy regularization term on the memory similarity vector in the loss function.

We trained 50 models with different random seeds for different amounts of noise injected into the hidden state before the start of recall. Adding noise reduces the information about the last items in the hidden state, allowing the model to more flexibly choose where to start recalling instead of starting from the end of the list.

To characterize the recall strategy of the models, we used two metrics to quantify the contiguity effect of a model: *forward asymmetry* evaluates the tendency of a model to recall an item presented after the previously recalled item instead of before the previously recalled item, and *temporal factor* quantifies the tendency of an agent to recall an item that is in close temporal proximity to the previously recalled item.

Results

We used a k-means algorithm to cluster the models based on the forward asymmetry and the temporal factor, as well as the decoding accuracy of item index (position of items in the list) and item identity. We observed two clusters of strategies in models with different random seeds and the amount of noise added to the hidden state at the start of recall (Figure 1c,d). The main factor that separates these two groups was the forward asymmetry, with models that tend to recall more in a forward order (forward asymmetry > 0.5) falling into the same

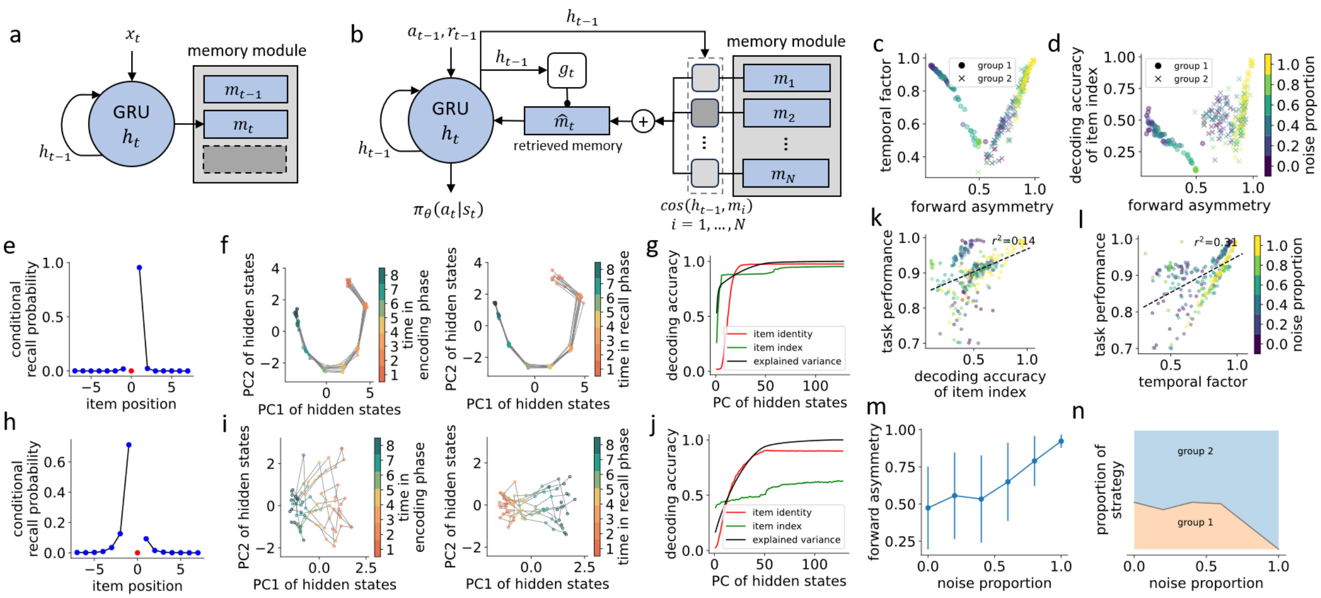

Figure 1: (a) Model structure during memory encoding. (b) Model structure during memory retrieval. (c) Clustering of models in the space of forward asymmetry and temporal factors. (d) Decoding accuracy of item index for the two clusters of models. (e-g) The first cluster of models. (e) Conditional recall probability as a function of the relative position of items in the list. (f) The first two principal components (PCs) of the hidden state across 10 trials during the encoding (left) and recall (right) phases. (g) Decoding accuracy of item identity (red) and item index (green) from an increasing number of PCs of the hidden states, together with the cumulative explained variance by the PCs (black). (h-j) Same as (e-g) for the second cluster of models. (k) Relation between task performance and decoding accuracy of item index. (l) Relation of task performance with temporal factor. (m) Forward asymmetry as a function of the proportion of noise injected between the encoding and recall phases, with error bars representing the standard deviation of models with different random seeds. (n) The proportion of models that fall in each cluster as a function of the noise proportion.

group and models that tend to do more backward recall (forward asymmetry < 0.5) falling into another.

The first group of models learned to mostly recall in a forward order by encoding the item index information in the hidden state of the recurrent layer (Figure 1e). The hidden state of this group of models followed very similar trajectories across trials in both the encoding and recall phases, regardless of what specific items were presented in the input list (Figure 1f). This allowed the model to sequentially reinstate each hidden state from the encoding phase to retrieve the correct memories. We also found that item index was encoded in a low-dimensional hidden state space, and that item identity was encoded in a higher-dimensional space (Figure 1g). The behavior of this group of models resembles the memory palace technique, where subjects mentally traverse the same set of locations along a stereotypical path during encoding and recall, in a pre-constructed cognitive map. Conversely, the second group of models had a higher tendency to recall in a backward order (Figure 1h). There was a significant difference in the hidden state trajectories across trials with different item content (Figure 1i), and we can decode the item identity much better than the item index (Figure 1j). This group of models exhibits a more TCM-like recall pattern.

For each group of models, the temporal factor increased as the tendency of recalling in a particular order increased (Figure 1c). The decoding accuracy of item index information also increased as the tendency to recall in either forward or backward order increased (Figure 1d), though the decoding accuracy was overall lower for models that recall in a backward order. We found that the decoding accuracy of item index and the temporal factor were both positively correlated with the free recall performance of the models (Figure 1k, l). These results suggest that the model can improve its recall performance by learning to recall with high temporal contiguity, which matches findings in human experiments (Sederberg, Miller, Howard, & Kahana, 2010). The use of item index information can increase recall performance, consistent with the improvement in free recall performance of humans from using the memory palace technique.

We were able to influence the strategy used by the models by changing the proportion of noise injected into the hidden state before the start of recall. We found that the forward asymmetry increased as the proportion of noise increased (Figure 1m), and there was a larger number of models that learned to recall in a forward order when the noise was high (Figure 1n). This suggests that models with more context about the last few items in the list tend to recall in a backward order and rely less on item index information. When the model has little information about the context of the last few items, it can avoid recalling from the end of the list, thus managing to begin from the start and recall in a forward order by relying on item index information to obtain optimal performance.

References

Cho, K., Van Merriënboer, B., Gulcehre, C., Bahdanau, D., Bougares, F., Schwenk, H., & Bengio, Y. (2014). Learning phrase representations using rnn encoder-decoder for statistical machine translation. *arXiv preprint arXiv:1406.1078*.

Graves, A., Wayne, G., & Danihelka, I. (2014). Neural turing machines. *arXiv preprint arXiv:1410.5401*.

Howard, M. W., & Kahana, M. J. (2002). A distributed representation of temporal context. *Journal of mathematical psychology*, *46*(3), 269–299.

Jensen, K. T. (2023). An introduction to reinforcement learning for neuroscience. *arXiv preprint arXiv:2311.07315*.

Kahana, M. J. (1996). Associative retrieval processes in free recall. *Memory & cognition*, *24*(1), 103–109.

Mnih, V., Badia, A. P., Mirza, M., Graves, A., Lillicrap, T., Harley, T., ... Kavukcuoglu, K. (2016). Asynchronous methods for deep reinforcement learning. In *International conference on machine learning* (pp. 1928–1937).

Polyn, S. M., Norman, K. A., & Kahana, M. J. (2009). A context maintenance and retrieval model of organizational processes in free recall. *Psychological review*, *116*(1), 129.

Sederberg, P. B., Miller, J. F., Howard, M. W., & Kahana, M. J. (2010). The temporal contiguity effect predicts episodic memory performance. *Memory & cognition*, *38*, 689–699.

Wang, J. X., Kurth-Nelson, Z., Tirumala, D., Soyer, H., Leibo, J. Z., Munos, R., ... Botvinick, M. (2016). Learning to reinforcement learn. *arXiv preprint arXiv:1611.05763*.

Yates, F. A. (2013). *Art of memory*. Routledge.

Zhang, Q., Griffiths, T. L., & Norman, K. A. (2023). Optimal policies for free recall. *Psychological Review*, *130*(4), 1104.