Optimal Timing for Episodic Retrieval and Encoding for Event Understanding

Qihong Lu[1], Zi Ying Fan[2], Uri Hasson[1], Kenneth A. Norman[1]

1. Department of Psychology and Princeton Neuroscience Institute, Princeton University; 2. Department of Computer Science, Princeton University

Summary

- We present a neural network that learns to use episodic memory for event prediction.
- The learned memory-retrieval policy shows a speed-accuracy trade-off that is sensitive to the cost of incorrect predictions.
- Models that selectively encode at event boundaries had fewer retrieval errors.
- Collectively, this model provides insights about why, in real data, episodic retrieval seems to be sensitive to prediction demand and uncertainty [1] and why episodic encoding seems to happen selectively at event boundaries [3, 4].

Model detail

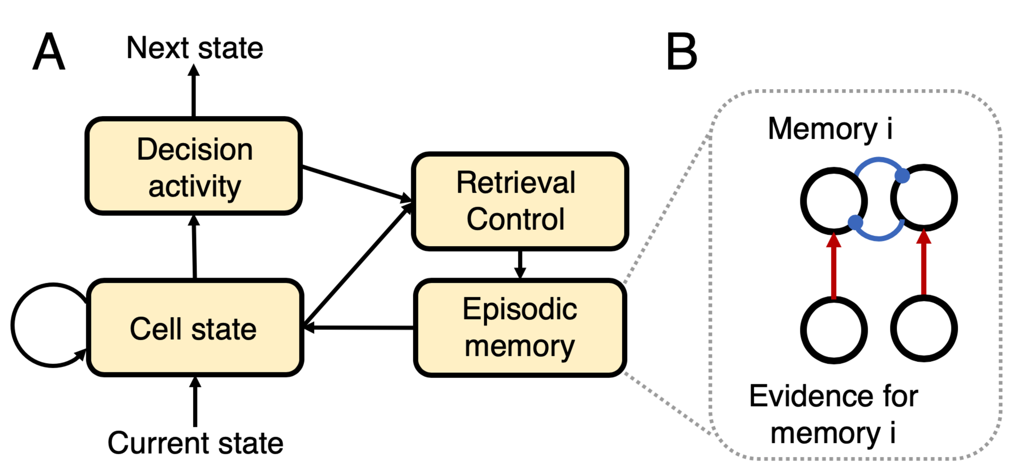

Figure 1: A) The model architecture; B) Memories are a set of laterally competing evidence accumulators.

Cortex is a recurrent neural network (LSTM) that predicts the upcoming state.

Hippocampus represents memories as a set of evidence accumulators. Each memory is a previously saved cortical pattern.

- **Retrieval** is an evidence-accumulation process that determines which memory to retrieve. Evidence is proportional to the similarity between the current state and stored memory.
- The **retrieval control** layer controls the feed-forward weights and the level of lateral competition between different memories.
- **Encoding** a new memory corresponds to adding a new accumulator.

The model is trained with reinforcement learning. The reward is positive/negative if the prediction is correct/incorrect. The model can say "don't know", in which case the reward is zero.

A context-dependent event prediction task

An **event sequence** is a sample path from a event schema conditioned on a situation (fig 2). An **event schema** is a graph, where each transition is controlled by a particular feature of the situation (2 A). Thus, knowing the features of the situation is useful for event prediction.

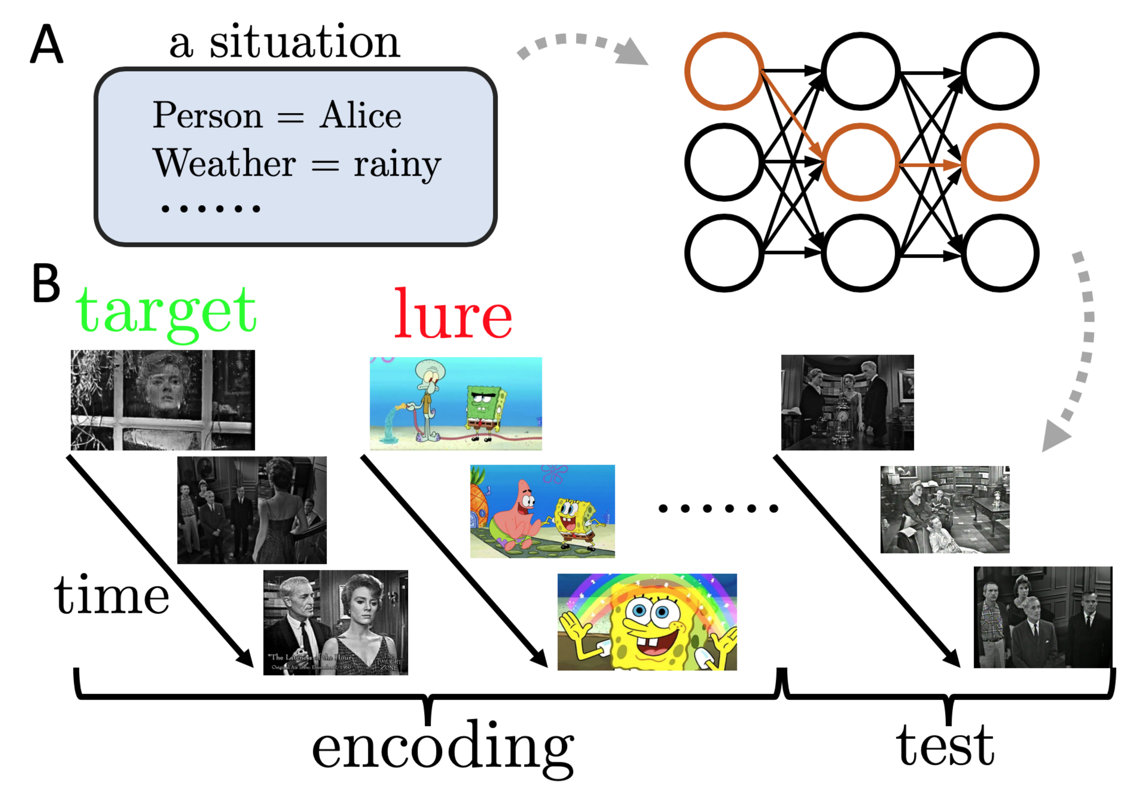

Figure 2: A) Sample an event sequence from the schema; B) An example trial.

Use episodic memory to predict upcoming events

There are three conditions (inspired by [1]): during test, the ongoing situation is ...

- recently observed (RM; recent memory)
- observed in distant past (DM; distant memory)
- new (NM; no memory)

In the DM condition, relevant information is not in working memory, so it needs to be recalled from episodic memory. DM prediction accuracy starts low, and increases after episodic recall takes place.

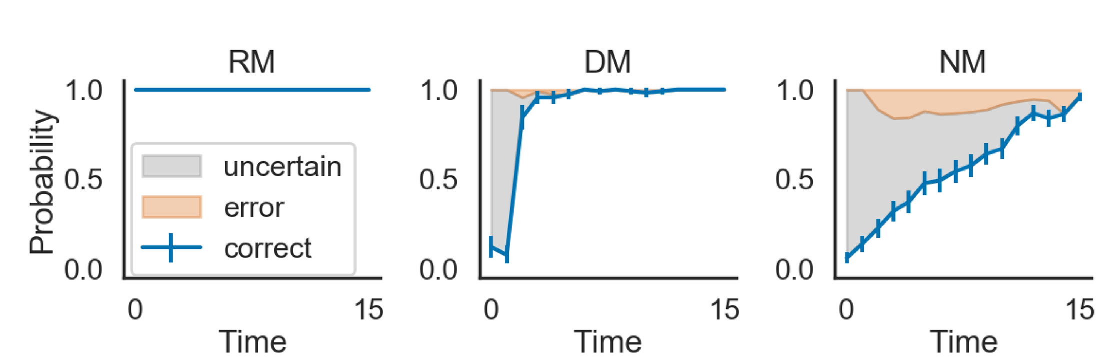

Figure 3: Event prediction performance across the three conditions.

Properties of the learned recall policy

Recall is sensitive to whether information is already in working memory (fig 4: lower recall in RM condition than DM; see also fig 6). We also manipulated the penalty associated with incorrect predictions. When penalty is high, recall is delayed (fig 4 A vs. B) and false recall is low (fig 5), similar to a well-established model of hippocampus [2].

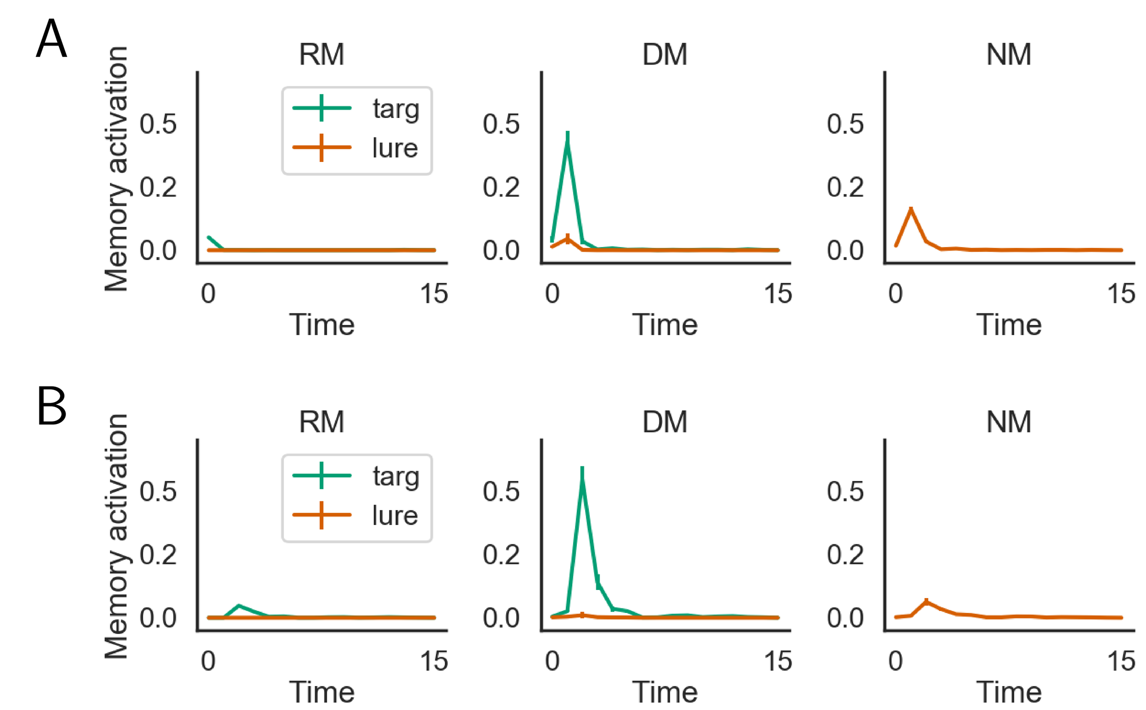

Figure 4: Memory activation in the A) low; and B) high penalty environment.

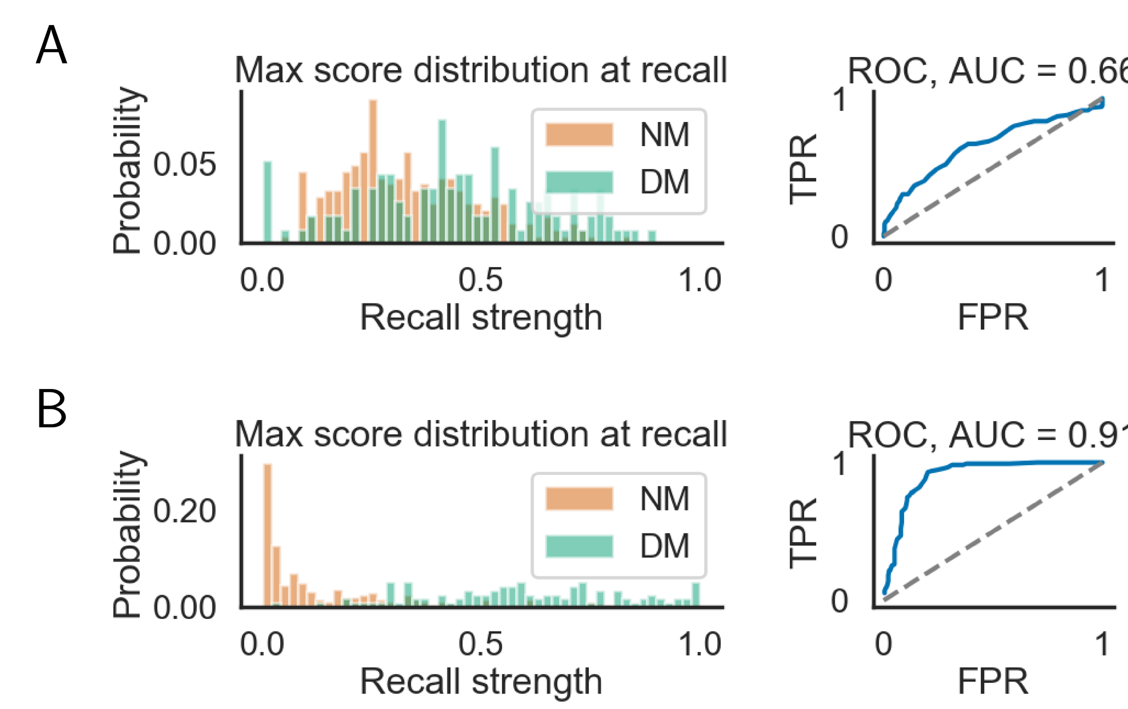

Figure 5: ROC analysis for the memory activation in the A) low; and B) high penalty environment.

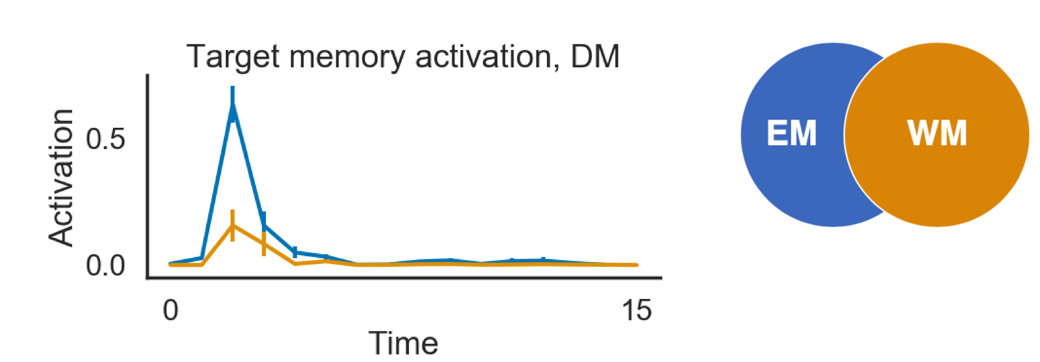

Figure 6: Recall is suppressed when the queried information is in working memory (i.e. when uncertainty is low).

Encoding at event boundaries reduces subsequent memory errors

We found models that **encode at event boundaries** performed better at subsequent recall (fig 7), compared to models that also encode episodic memories within an event sequence (i.e. **cumulative encoding**), because encoding within an event sequence leads to more confusable memories (fig 8).

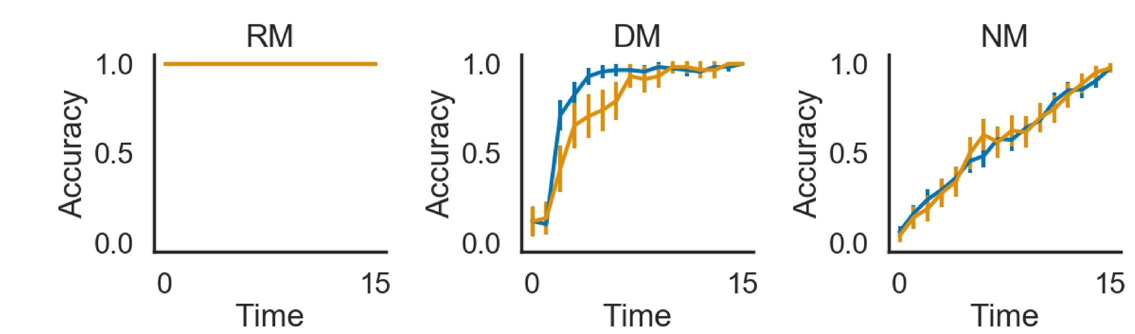

Figure 7: Event prediction accuracy for models that encode at event boundaries ; vs. models that also encode within an event sequence.

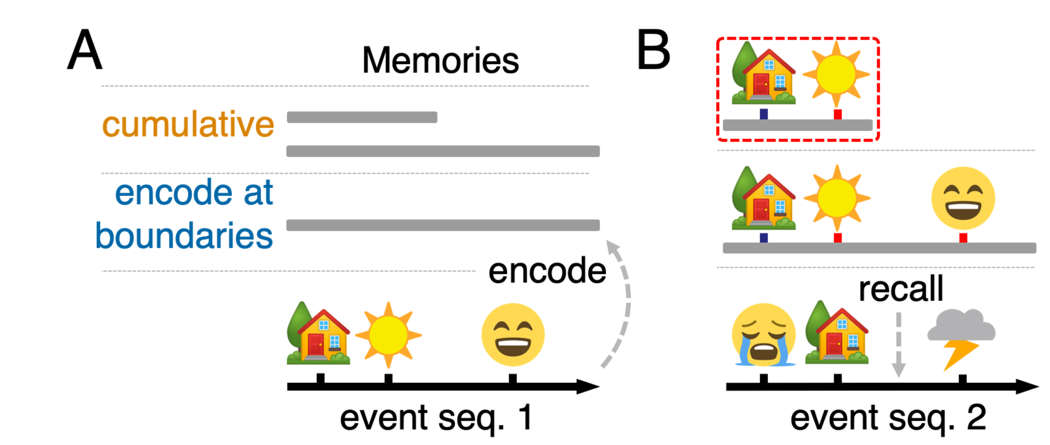

Figure 8: A) The resulting memory chunks under the two encoding regimes; B) Encoding within an event sequence might cause subsequent false recall. Connecting all information make lures easier to reject.

References & Acknowledgement

[1] Chen, J. et al. (2016) Cereb Cortex.
[2] Norman, K. (2010) Hippocampus.
[2] Baldassano, C., et al. (2017) Neuron.
[3] Ben-Yakov A., & Henson, N. (2018). J Neuro.

Acknowledgement: This work was supported by a Multi-University Research Initiative grant to KAN and UH (ONR/DoD N00014-17-1-2961).

Download the poster: `https://tinyurl.com/CCN19-QL`