UC Merced

Proceedings of the Annual Meeting of the Cognitive Science Society

Title

Parallels between Neural Machine Translation and Human Memory Search: A Cognitive Modeling Approach

Permalink

https://escholarship.org/uc/item/43h3w8cv

Journal

Proceedings of the Annual Meeting of the Cognitive Science Society, 46(0)

Authors

Salvatore, Nikolaus

Zhang, Qiong

Publication Date

2024

Copyright Information

This work is made available under the terms of a Creative Commons Attribution License, available at https://creativecommons.org/licenses/by/4.0/

Peer reviewed

eScholarship.org Powered by the California Digital Library
University of California

Parallels between Neural Machine Translation and Human Memory Search: A Cognitive Modeling Approach

Nikolaus Salvatore [a], **Qiong Zhang** [a,b,c]
{nikolaus.salvatore, qiong.z}@rutgers.edu

[a]Department of Computer Science, [b]Department of Psychology, [c]Rutgers Center for Cognitive Science
Rutgers University–New Brunswick, Piscataway, NJ, USA 08854

Abstract

In this work, we propose a neural network model for free recall that draws direct parallels between neural machine translation (NMT) and cognitive models of memory search, specifically the Context Maintenance and Retrieval (CMR) model. We hypothesize that NMT advancements such as attention mechanisms (Luong et al., 2015) closely resemble how humans reactivate prior contexts ("mental time travel"; Tulving, 1985). To demonstrate these parallels, we train a seq2seq model with attention as a cognitive model of memory search, and evaluate behavior against available human free recall data. We find that at intermediate levels of training, the model can capture several phenomena observed in human free recall experiments (Kahana et al., 2022); and after optimization, the model demonstrates the same optimal behavior as previously derived by the CMR model (Zhang, Griffiths, & Norman, 2023). Performing an ablation study, we demonstrate that behavioral differences between models with and without attention align with impaired behavior observed in hippocampal amnesia patients (Palombo, Di Lascio, Howard, & Verfaellie, 2019).

Introduction

Humans and artificial agents frequently encounter similar computational problems (Griffiths, Steyvers, & Firl, 2007), which has often led to a synergy between machine learning and cognitive science research in the past (Griffiths et al., 2007; Callaway, Griffiths, Norman, & Zhang, 2023; Lampinen, Chan, Banino, & Hill, 2021; Van de Ven, Siegelmann, & Tolias, 2020). Over the past decade, neural machine translation (NMT) has revolutionized the field of automated translation by leveraging deep learning techniques to translate text from one language to another (Luong et al., 2015; Bahdanau, Cho, & Bengio, 2014; Cho, Van Merriënboer, et al., 2014). In this work, we demonstrate that neural network models for machine translation and a cognitive model of human memory search share strikingly similar architectural components, despite being developed in separate communities and with distinct applications.

First, we draw a parallel between the sequence-to-sequence (seq2seq) model and the *Context Updating* mechanism critical to context-based models in human memory. It has been long recognized that context plays an important role in supporting episodic memory recall in humans (Estes, 1955; Bower, 1967). The Context Maintenance and Retrieval model (CMR; Polyn, Norman, & Kahana, 2009; Lohnas, Polyn, & Kahana, 2015), a successor of the Temporal Context Model (Howard & Kahana, 2002a), posits that context gradually evolves over time and binds with to-be-remembered items. During testing, the initial recall context is equivalent to the final context state during encoding, which serves as a cue to drive further recalls (Howard & Kahana, 2002b; Polyn et al., 2009; Lohnas et al., 2015). These models have been supported by a range of behavioral findings in human episodic memory, particularly in free recall tasks, where participants are asked to recall items in any order from a list just studied (Murdock, 1962). We propose that the context updating mechanism in CMR closely resembles the encoding and decoding process found in the classic sequence-to-sequence (seq2seq) models (Sutskever, Vinyals, & Le, 2014; Cho, Van Merriënboer, et al., 2014), such as Long Short-Term Memory (LSTM; Hochreiter & Schmidhuber, 1997) units and Gated Recurrent Units (GRUs; Cho, van Merriënboer, Bahdanau, & Bengio, 2014). By maintaining and updating a latent hidden state, the input encoding process in seq2seq models is analogous to how context in human memory evolves and is updated over time, whereas the decoding process is similar to how shifting context in humans drives further recall.

Next, we establish the similarity between the attention mechanism in NMT and the *Context Reinstatement* mechanism (or "jump back in time") in CMR. Long-range temporal dependencies have been a significant obstacle for classic seq2seq models. With regard to translation, long-range dependencies from an input sentence must be encoded entirely into a fixed-length vector (hidden state) before it is passed to the decoder. The introduction of attention mechanisms (Bahdanau et al., 2014; Luong et al., 2015), allows the model to focus on different parts of the input sequence while generating each word in the output, giving the decoder direct access to previous encoder states during the decoding process. We argue that the attention mechanism closely resembles the context reinstatement mechanism in CMR, as recently recalled items create current contexts that partially reactivate previous contexts in which the item was originally studied. As one might have exhausted items associated with the current context, this additional mechanism of context reinstatement helps one "jump back in time" toward the original study contexts (similar to Tulving's notion of mental time travel; Tulving, 1985), and serves as a retrieval cue for the recall of remaining items. The method by which CMR reactivates previous context states is analogous to how an attention mechanism can

In L. K. Samuelson, S. L. Frank, M. Toneva, A. Mackey, & E. Hazeltine (Eds.), *Proceedings of the 46th Annual Conference of the Cognitive Science Society.* ©2024 The Author(s). This work is licensed under a Creative Commons Attribution 4.0 International License (CC BY).

reactivate previous encoding hidden states.

To demonstrate that the parallels between seq2seq models and CMR go beyond the conceptual level, we implement a basic seq2seq model (the cornerstone of NMT) with an attention mechanism (Luong et al., 2015) as a cognitive model of human memory search. To do this, we train our model in a reinforcement learning framework to complete a typical human free recall task involving recalling words from a presented list. We show that the fully optimized seq2seq model with attention converges to the same behavioral patterns as optimal free recall, previously derived under a rational analysis of the Context Maintenance and Retrieval model (rational-CMR; Zhang et al., 2023). Furthermore, model evaluations conducted intermittently throughout the model's training exhibit similar recall characteristics as human patients, in terms of the primacy, recency, and temporal contiguity effects (Murdock, 1962; Kahana, 1996; Howard & Kahana, 1999).

In addition to demonstrating that our seq2seq model can serve as a cognitive model of human memory search, we provide additional novel modeling analysis to understand memory deficiencies in hippocampal amnesia. Medial temporal lobe (MTL) lesions have been previously associated with the inability to reinstate prior experienced contexts (Reed & Squire, 1998; Palombo et al., 2019; Howard, Fotedar, Datey, & Hasselmo, 2005; Scoville & Milner, 1957). If the attention mechanism in our seq2seq model is analogous to the context reinstatement mechanism in CMR, an ablation of the attention mechanism should capture similar recall characteristics as hippocampal amnesic patients compared with healthy controls. To foreshadow our results, our ablation study, consisting of the model without the attention mechanism, reveals a similar impairment as that of amnesiac patients.

In the following sections of the paper, we give an overview of the CMR model and draw parallels with the NMT mechanisms found in the seq2seq model. We describe the overall structure of our seq2seq model and our free recall training procedure. Finally, we detail evaluation results across a range of model configurations, demonstrating alignment in behavior characteristics between our model and human subjects.

Context Maintenance and Retrieval Model (CMR)

Context plays an important role in the encoding and retrieval of information (Estes, 1955; Bower, 1967; Anderson & Bower, 1972). Computational models such as TCM and CMR (Howard & Kahana, 2002b; Polyn et al., 2009; Lohnas et al., 2015) posit that context slowly drifts over time and binds with to-be-remembered external or internal experiences. Context at time t, denoted as c_t, follows the process:

$$c_t = \rho c_{t-1} + \beta_{enc} c^{IN}_{enc} \tag{1}$$

where c^{IN} is the retrieved context induced by an encountered experience, $\beta_{enc} \in [0,1]$ is a parameter determining the rate at which context drifts toward the new experience, and ρ is a scalar ensuring $||c_t|| = 1$. When an item is presented in the study list, it activates its pre-experimental context $c^{IN}_{enc} = M^{FC}_{pre} f_t$, where M^{FC}_{pre} represents item-to-context associations that existed prior to the experiment (initialized as an identity matrix, under the simplifying assumption that an item is only associated with its own context; see Polyn et al., 2009), and f_t is a binary vector that is all zeros except at the presented item's position. Therefore, $M^{FC}_{pre} f_t$ is the context previously associated with the presented item. In addition to these fixed pre-experimental item-to-context associations held in M^{FC}_{pre}, there are also experimental item-to-context and context-to-item associations held in M^{FC}_{exp} and M^{CF}_{exp} that capture new learning in the experiment. These matrices are initialized to zero and are updated during the study phase. Specifically, when an item is presented, a new association is formed between the presented item and the current context state via the Hebbian outer-product learning rule: $\Delta M^{FC}_{exp} = \Delta M^{CF}_{exp} = f_t c^T_{t-1}$.

During recall, the memory search process is driven by the current state of the context representation c_t. The support for recalling each item depends on how much the current context matches the items' study context. The starting context at recall is close in time to (and thus similar to) the end-of-list context during encoding, giving rise to better recalls for items studied at the end of the list (i.e., recency effects; Murdock, 1962). As recall continues, context evolves under the same process as it did during the encoding phase, $c_t = \rho c_{t-1} + \beta_{rec} c^{IN}_{rec}$ as in Equation (1), but with the retrieved context c^{IN}_{rec} introduced differently. The key difference is that items are encountered for the first time during the experiment in the encoding stage (which only activates its pre-experimental context $c^{IN}_{enc} = M^{FC}_{pre} f_i$), whereas – in the recall phase – items are encountered for the second time when they are recalled. The retrieved context c^{IN}_{rec} can come from both the pre-experimental context associated with the item $M^{FC}_{pre} f_j$ and the experimental context associated with the item $M^{FC}_{exp} f_j$, which has been acquired through Hebbian learning during the encoding phase. The extent of retrieving the pre-experimental context versus retrieving the experimental context is regulated by the parameter $\gamma_{FC} \in [0,1]$,

$$c^{IN}_{rec} = (1-\gamma_{FC}) M^{FC}_{pre} f_j + \gamma_{FC} M^{FC}_{exp} f_j = (1-\gamma_{FC}) c^{IN}_{enc} + \gamma_{FC} c_{i-1}. \tag{2}$$

The value of γ_{FC} has important implications for the recall transition patterns. When $\gamma_{FC} = 0$, the retrieved context entirely consists of pre-experimental context associated with the item $M^{FC}_{pre} f_j$, which is identical to the retrieved context c^{IN}_{enc} when the item was first encountered during encoding. When $\gamma_{FC} = 1$, the retrieved context entirely consists of experimental context associated with the item $M^{FC}_{exp} f_j$, which is essentially the context c_{i-1} during encoding.

Both c^{IN}_{enc} and c_{i-1} are part of the original study context, so reinstating them into the current context has the effect of mentally "jumping back in time". As a consequence of this "jumping back in time", items that were studied close in time

to the just-recalled item have a higher chance of being recalled next, because of the similarity between their study context and the current context. This contributes to the temporal contiguity effect commonly observed in free recall (Kahana, 1996; Howard & Kahana, 1999). Though both contribute to temporal contiguity, reinstating c^{IN}_{enc} and reinstating c_{i-1} bias forward recalls and backward recalls differently. During encoding, at $t = i$, the drifting part of the study context ρc_{i-1} is similar to contexts both before and after $t = i$, and has the chance to be associated with items that come before and after $t = i$. Therefore, reinstating c_{i-1} gives rise to both forward recalls and backward recalls. However, since an item's pre-experimental context is incorporated into temporal context after the item is presented, the retrieved context c^{IN}_{enc} does not share any similarity with contexts before $t = i$, and only has an opportunity to be associated with items that come after. Therefore, reinstating c^{IN}_{enc} gives rise to only forward recalls, which accounts for the forward asymmetry commonly seen in free recall experiments. Medial temporal lobe amnesia has been associated with the impairment in backward contiguity, with the lack of access to the experimental context c_{i-1} but intact access to the pre-experimental or semantic context of an item (Palombo et al., 2019). Finally, to be able to fully simulate recall patterns, variants of CMR are equipped with different retrieval rules and recall termination rules (see more details in Howard & Kahana, 2002b; Polyn et al., 2009; Lohnas et al., 2015; Zhang et al., 2023).

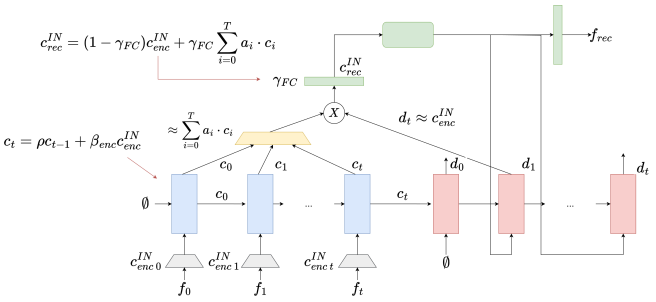

Figure 1: A graphical illustration of NMT model to highlight its alignment with CMR

Neural Machine Translation Figure 1 represents a seq2seq neural network model framed using the mathematical notation commonly associated with the context maintenance and retrieval (CMR) model. The motivation behind this graphical representation is to conceptually bridge the gap between NMT and CMR. In this model architecture, both the encoder and decoder blocks represent GRUs. In this diagram, the encoder (blue boxes) processes the input sequence of words (f_{t-1}, f_t, f_{t+1}), which is embedded into a lower-dimensional semantic space (represented by c^{IN}_{enc}). The sequence of encoded vectors (c_{t-1}, c_t, c_{t+1}) is generated in a unidirectional manner, with each vector being a function of the previous context vector and the input at that time step. Internally, the GRU computes its output and updates its internal hidden state according to a learned update rule z_t, such that $h_t = (1 - z_t) \cdot h_{t-1} + z_t \cdot \hat{h}_t$. $\hat{h}_t$ represents the cell's current, potential memory content, where a certain amount of information from the last timestep is parametrically removed by a reset gate, r_t, and new information is incorporated. h_t is the final memory state at the end of the timestep that controls the relative change from the hidden state at the previous timestep. The degree by which the current hidden state is modified by new inputs is analogous to the drifting context in the CMR model, where newly encoded items update the current recall context, but do not completely override it. We draw a direct parallel with CMR's context updating mechanism with the equation for c_t (shown in the diagram).

In our comparative model, the attention mechanism (yellow box) parallels the "jump back in time" mechanism in human memory search, giving the decoder access to the previously encoded contexts, c_t. The form of attention used in this comparative model is global Luong (or multiplicative) attention, in which the current hidden state is updated using a sum of encoder hidden states weighted by similarity. The equation $a_i = softmax(d_t^T \cdot c_i)$ represents the similarity scores calculated between the current decoder hidden state d_t and each of the encoder hidden states, c_i. The actual update of the context is performed according to $\sum_{i=1}^{T} a_i \cdot c_i$, which is used to update recall context, c^{IN}_{rec} (Figure 1). In this manner, the attention mechanism can trigger a shift in the decoder's context that may ordinarily not be possible when relying on the internal hidden state alone.

The decoder (red boxes) receives the final hidden state from the encoder, c_t. Once the decoder's hidden state and previous encoder states are passed through the attention mechanism, the resulting recall context, c^{IN}_{rec}, is used to generate an output for a given time step. The decoder's hidden state is then updated by the feedback of the output embedding before it proceeds to the next stage of the decoding process. The parameters $(\gamma_{FC}, \beta, \rho)$ in the encoder's and the decoder's update equations mirror the CMR parameters for context updating and recall, but are implicitly learned. Specifically, γ_{FC} represents the relative weight given to the previous contexts reactivated through the attention mechanism, which necessarily becomes 0 with the removal of the attention mechanism.

Methods

The model training procedure is posed as a reinforcement learning problem, in which an agent is presented with a sequence of words and then expected to recall this sequence for words when prompted with a start-of-sequence token. We use the Proximal Policy Optimization (PPO) algorithm (Schulman, Wolski, Dhariwal, Radford, & Klimov, 2017) and allow the agent to terminate its own recall by the prediction of an end-of-sequence token. Following techniques in Bahdanau et al. (2014), the seq2seq model (serving as the actor in the PPO agent) is pre-trained for a brief period using supervised training before transitioning to reinforcement learning.

For this experiment, we opt to use GRUs due to their simplified internal architecture, without sacrificing performance.

During decoding (see Figure 1), the decoder's hidden state (after passing through the attention mechanism) is compared via cosine similarity to each encoder hidden state (stored in a memory table as keys). The output embedding is a weighted combination of word embeddings (stored in the same memory table as values) based on cosine similarity with the corresponding encoding state, which is then transformed into an original list item for recall through a feed-forward layer. This implements a similar retrieval rule as in CMR, where the probability of recalling each item depends on how much the current context matches the items' study contexts (Polyn et al., 2009; Lohnas et al., 2015).

Given that free recall is inherently an unordered task, unlike translation, conventional cross-entropy loss is unsuitable due to the importance of item order in loss calculation. In order to train the seq2seq model independent of the order of recall, we treat the problem as a set prediction task and opt to use an iterative Sinkhorn algorithm approximation method that impacts training time, but ensures that network learning is not impacted by token order (Brun, Gaüzère, Renton, Bougleux, & Yger, 2022). For data, we use the PEERS free recall dataset as a basis of comparison with human subjects, extracting an experimental vocabulary from all words appearing in human trials (Kahana et al., 2022). All model configurations were trained using 50,000 randomly generated sequences of 14 words sampled from the PEERS vocabulary. The embedding layer of the encoder model is instantiated using pre-trained GloVe embeddings (Pennington et al., 2014), which is then frozen to prevent the model from modifying these embeddings during training. This frozen embedding layer is analogous to a human's long-term semantic memory and creates an initial embedding space in which more similar words should appear more closely together.

To test the ability of our model to predict human data in medial temporal lobe amnesia patients and healthy controls (Palombo et al., 2019), we investigate two distinct configurations of the free recall model: the standard model as described and a model with the attention mechanism removed. In addition, each configuration is evaluated using several recurrent network hidden dimension sizes (32, 64, 128).

Results and Discussion

Figures 2-4 illustrate the behavioral patterns for each configuration trained on the free recall task. Each configuration was trained based on the stimuli and trial structures in the PEERS free recall dataset (Kahana et al., 2022). For direct comparison with the CMR model and human subjects, we analyze model simulation results describing three sets of recall patterns: 1) how well on average model retrieves items for each position in the study list (serial position curve), 2) where the model initiates its recall from (probability of the first recall curve), and 3) how likely it is to recall items studied consecutively in the study list (conditional response probability). For the final analysis, the probability is computed by dividing the number of times a transition of that lag is actually made by the number of times it could have been made (Kahana, 1996). Following how the above patterns are analyzed in human data (Kahana, 1996; Polyn et al., 2009), we removed repetitions in the model when performing our analyses.

The optimal attention model demonstrates the same recall behavior as the optimal policy of the cognitive model CMR. Past cognitive modeling work with CMR has demonstrated that not only could it capture averaged human behavioral patterns (through fitting model parameters to human data; Polyn et al., 2009; Lohnas et al., 2015), but it can explain why some individuals achieve better memory performance than others in the free recall task by analyzing how well their behavior aligns with an optimal policy of CMR (rational-CMR; Zhang et al., 2023). When the free recall behavior is optimized under the architectural constraints of the CMR model, the optimal policy always starts by recalling from the beginning of the list (Figure 2b) and sequentially recalls forwards (Figure 2c) despite no constraint placed on the order of recall. This optimal behavior is non-trivial, as one already has access to the end-of-list context at the start of recall, whereas reactivating the beginning-of-list context requires an extra jump back to early items of the list. Figure 2 shows that our trained seq2seq model with attention (blue) demonstrated the same behavior as the optimal policy of CMR (orange; reproduced from Figure 3A in Zhang et al., 2023). Both models achieve near-optimal performance (Figure 2a), with a near certainty of starting recall at the beginning of the sequence (Figure 2b) and making a forward recall transition to the next serial position at lag +1 (Figure 2c). In 2c, the optimal CMR model shows a small amount of backward recall in contrast to the NMT model's complete lack of it. CMR, regardless of optimality, includes some stochasticity in its recall of items whereas the NMT model does not.

The intermediate training evaluations of the attention model exhibit similar qualitative patterns as is typically observed in human participants. Only a small proportion of top-performing human participants are able to demonstrate the exact behavior of the optimal policy (Zhang et al., 2023). Averaged human recall behavior in free recall experiments (blue line in Figure 3a-3c; reproduced from Kahana et al., 2022) also exhibits recency (enhanced end-of-list recall; Figure 3b) and backward contiguity (tendency toward shorter, backward lags; Figure 3c), in addition to what is amplified in the optimal behavior with primacy (enhanced beginning-of-list recall; Figure 3b) and forward contiguity (tendency toward shorter, forward lags; Figure 3c). We show that while the trained NMT model with attention demonstrates the same behavior as the optimal policy of CMR, exhibiting primacy and forward contiguity, intermediate training evaluations of the model (lighter yellow lines) exhibit recency (Figure 3b) and backward contiguity (Figure 3c), though to a smaller degree compared to averaged human recall behavior. Recency and backward contiguity quickly disappear as the model is optimized to further rely on forward recalls initiated from the beginning of the sequence. Note that in some cases (i.e.

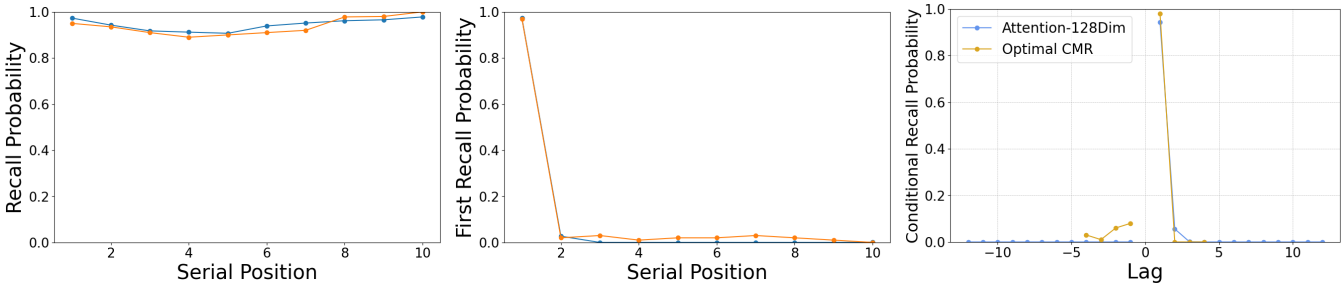

(a) (b) (c)

Figure 2: Behavioral recall patterns for the fully optimized, 128 dimension model with attention compared to the expected behavioral results derived by the optimal policy CMR model. Optimal CMR results are reproduced from Zhang et al. (2023).

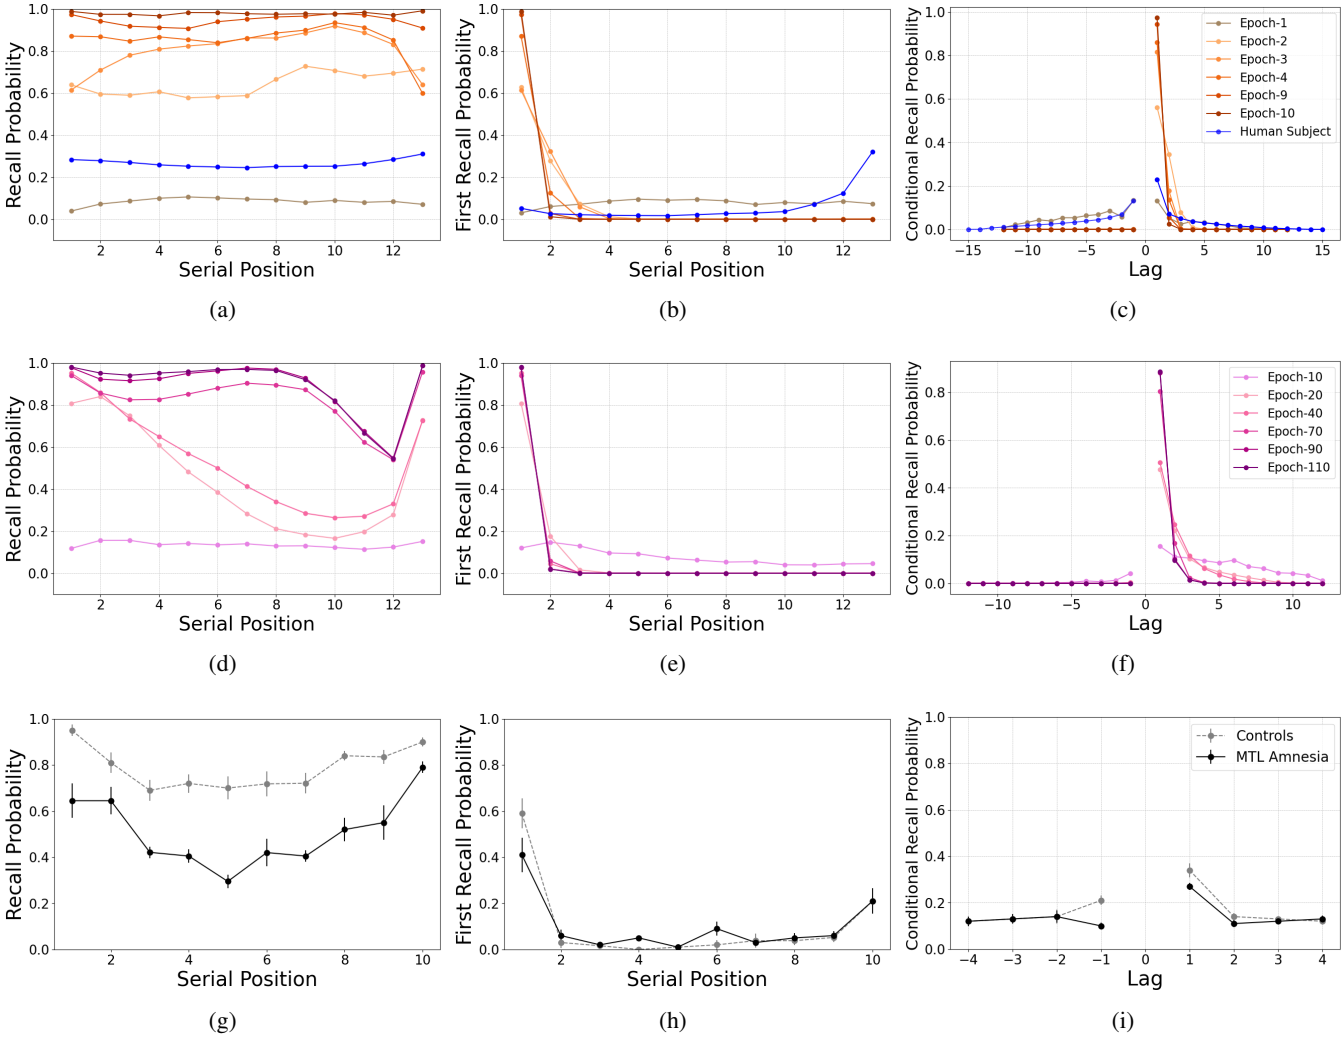

(a) (b) (c)

(d) (e) (f)

(g) (h) (i)

Figure 3: Behavioral recall patterns under different intermediate states of model training with attention (a-c), without attention (d-f), and human subjects of medial temporal lobe amnesia and healthy controls (g-i), reproduced from Palombo et al. (2019). Model with attention was overlaid with human results reproduced from Kahana et al. (2022).

Epochs 3 and 4), unlike human participants, the recall curve in the model shows an inverted U-shape.

An ablation study of the attention mechanism demonstrates alignment between model results and human data with MTL amnesia. Medial temporal lobe (MTL) lesions have been previously associated with the inability to re-

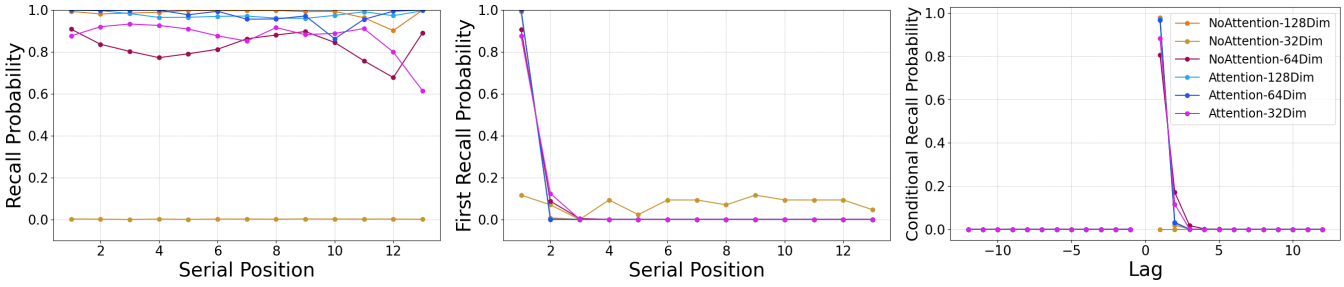

(a) (b) (c)

Figure 4: The effect of different hidden dimension sizes (32, 64, 128) on the behavioral recall patterns for attention and no attention models.

instate prior experienced contexts (Reed & Squire, 1998; Palombo et al., 2019; Howard et al., 2005; Scoville & Milner, 1957), mathematically implemented as the context reinstatement mechanism in CMR. As we hypothesized that the attention mechanism in our seq2seq model is analogous to context reinstatement in CMR, we carried out an ablation study of the attention mechanism to examine if model differences could capture the key behavioral difference in the amnesic patients and the healthy controls. Results of this ablation are displayed in figures 3d-3f. Models with attention converge much more quickly to optimal behavior and thus require fewer training epochs. Figures 3g-3i depict recall patterns for patients with medial temporal lobe (MTL) amnesia, reproduced from (Palombo et al., 2019). Notably, amnesiac patients lack the ability to "jump back in time" and therefore display a greatly reduced backward contiguity relative to healthy controls (lower -1 lag in 3i), consistent with CMR predictions when the ability to reinstate the original study contexts is impaired. Similarly, the lack of the attention mechanism in Figure 3f also largely eliminates the model's capacity for backward contiguity (at any stage of training) compared to with attention, where backward contiguity is observed in early iterations of training (Figure 3c).

An ablation study of the attention mechanism provides insight into the performance difference between amnesia patients and healthy controls. Figure 4 shows evaluation results for each configuration with respect to hidden dimension size and the presence of the attention mechanism after full optimization. The attention model consistently exhibits higher recall probability than the no-attention model (Figure 4a), capturing the same memory performance patterns observed in amnesia patients and healthy control (Figure 3g). Since the final model iteration of both the attention model and the no-attention model demonstrate optimal free recall behavior (with primacy and forward contiguity), one might ask why the attention model has better memory performance when it must learn to ignore backward contiguity, while the no-attention model has forward contiguity by default (i.e. is incapable of backward contiguity). It has been demonstrated in the optimal CMR model that forward contiguity only contributes to performance in combination with primacy (Zhang et al., 2023). We hypothesize that the lack of the attention mechanism (i.e. the ability to reinstate previous study contexts) removes the ability to reinstate the beginning-of-list context, which is necessary to establish primacy. We suggest that this lack of primacy gives rise to reduced performance in the no-attention case. An alternative, and more difficult way, to obtain primacy in the no-attention model is to maintain all study items in the working memory (i.e. hidden states); and this difficulty is greater when the working memory capacity is smaller. Consistent with our hypothesis, we observed higher performance difference between the attention model and the no-attention model when the models' hidden dimensions are small. A hidden dimension size of 32 is inadequate to observe any recall performance in the no attention case, whereas the attention configuration shows considerably higher recall performance regardless of hidden dimension size (Figure 4a).

Conclusion

We have demonstrated that seq2seq models, originally introduced for the purpose of machine translation, can also serve as a neural network representation of the CMR model of human memory search. Various components of the seq2seq model have direct comparisons with mechanisms of the CMR model, and an added attention mechanism grants the model a form of mental time travel, allowing it to explicitly reactive past contexts in a similar manner to CMR. Comparisons to the optimal CMR condition illustrate the model's ability to optimally recall, while comparisons between human subject data and incompletely optimized models reflect the model's potential for capturing sub-optimal recall behavior as well.

In addition to establishing this seq2seq model as an alternative memory search model (capable of capturing optimal and sub-optimal human behavior as CMR does), we additionally present the model as a potential framework for understanding memory deficits. By removing the attention mechanism and hampering the model's ability to reactivate previous encoding contexts, we eliminate the model's capacity for backward contiguity and hamper the model's recall performance in a manner similar to patients with MTL amnesia.

References

Anderson, J. R., & Bower, G. H. (1972). Recognition and retrieval processes in free recall. *Psychological Review*, *79*(2), 97–123.

Bahdanau, D., Cho, K., & Bengio, Y. (2014). Neural machine translation by jointly learning to align and translate. *arXiv preprint arXiv:1409.0473*.

Bower, G. (1967). A multicomponent theory of the memory trace. *Psychology of Learning and Motivation*, *1*, 229–325.

Brun, L., Gaüzère, B., Renton, G., Bougleux, S., & Yger, F. (2022). A differentiable approximation for the linear sum assignment problem with edition. *International Conference on Pattern Recognition*. doi: 10.1109/icpr56361.2022.9956203

Callaway, F., Griffiths, T. L., Norman, K. A., & Zhang, Q. (2023). Optimal metacognitive control of memory recall. *Psychological Review*.

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*.

Cho, K., van Merriënboer, B., Bahdanau, D., & Bengio, Y. (2014). On the properties of neural machine translation: Encoder–decoder approaches. *SSST@EMNLP*. doi: 10.3115/v1/w14-4012

Estes, W. K. (1955). Statistical theory of spontaneous recovery and regression. *Psychological Review*, *62*(3), 145–154.

Griffiths, T. L., Steyvers, M., & Firl, A. (2007). Google and the mind: Predicting fluency with PageRank. *Psychological science*, *18*(12), 1069–1076.

Hochreiter, S., & Schmidhuber, J. (1997). Long short-term memory. *Neural Computation*. doi: 10.1162/neco.1997.9.8.1735

Howard, M. W., Fotedar, M. S., Datey, A. V., & Hasselmo, M. E. (2005). The temporal context model in spatial navigation and relational learning: toward a common explanation of medial temporal lobe function across domains. *Psychological Review*, *112*(1), 75.

Howard, M. W., & Kahana, M. J. (1999). Contextual variability and serial position effects in free recall. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *25*(4), 923–941. doi: 10.1037/0278-7393.25.4.923

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

Howard, M. W., & Kahana, M. J. (2002b, Jun). A distributed representation of temporal context. *Journal of Mathematical Psychology*, *46*(3), 269–299. doi: 10.1006/jmps.2001.1388

Kahana, M. J. (1996, Jan). Associative retrieval processes in free recall. *Memory & Cognition*, *24*(1), 103–109. doi: 10.3758/bf03197276

Kahana, M. J., Lohnas, L. J., Healey, K., Aka, A., Broitman, A., Crutchley, E., ... others (2022). The Penn electrophysiology of encoding and retrieval study.

Lampinen, A., Chan, S., Banino, A., & Hill, F. (2021). Towards mental time travel: a hierarchical memory for reinforcement learning agents. *Advances in Neural Information Processing Systems*, *34*, 28182–28195.

Lohnas, L. J., Polyn, S. M., & Kahana, M. J. (2015). Expanding the scope of memory search: Modeling intralist and interlist effects in free recall. *Psychological Review*, *122*(2), 337–363.

Luong, M.-T., Luong, M.-T., Pham, H., Pham, H., Manning, C. D., & Manning, C. D. (2015). Effective approaches to attention-based neural machine translation. *Conference on Empirical Methods in Natural Language Processing*. doi: 10.18653/v1/d15-1166

Murdock, B. B. (1962, Nov). The serial position effect of free recall. *Journal of Experimental Psychology*, *64*(5), 482–488. doi: 10.1037/h0045106

Palombo, D. J., Di Lascio, J. M., Howard, M. W., & Verfaellie, M. (2019). Medial temporal lobe amnesia is associated with a deficit in recovering temporal context. *Journal of cognitive neuroscience*, *31*(2), 236–248.

Pennington, J., Pennington, J., Socher, R., Socher, R., Manning, C., & Manning, C. D. (2014). Glove: Global vectors for word representation. *Conference on Empirical Methods in Natural Language Processing*. doi: 10.3115/v1/d14-1162

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), 129–156.

Reed, J. M., & Squire, L. R. (1998). Retrograde amnesia for facts and events: findings from four new cases. *Journal of Neuroscience*, *18*(10), 3943–3954.

Schulman, J., Wolski, F., Dhariwal, P., Radford, A., & Klimov, O. (2017). Proximal policy optimization algorithms. *arXiv preprint arXiv:1707.06347*.

Scoville, W. B., & Milner, B. (1957). Loss of recent memory after bilateral hippocampal lesions. *Journal of neurology, neurosurgery, and psychiatry*, *20*(1), 11.

Sutskever, I., Vinyals, O., & Le, Q. V. (2014). Sequence to sequence learning with neural networks. *Advances in neural information processing systems*, *27*.

Tulving, E. (1985). Memory and consciousness. *Canadian Psychology/Psychologie canadienne*, *26*(1), 1.

Van de Ven, G. M., Siegelmann, H. T., & Tolias, A. S. (2020). Brain-inspired replay for continual learning with artificial neural networks. *Nature communications*, *11*(1), 4069.

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