

SynapticRAG: Enhancing Temporal Memory Retrieval in Large Language Models through Synaptic Mechanisms

Yuki Hou¹, Haruki Tamoto², Qinghua Zhao³, Homei Miyashita¹

¹Meiji University, ²Kyoto University, ³Hefei University

houhoutime@gmail.com, harukiririwuru@gmail.com, zhaoqh@buaa.edu.cn

Abstract

Existing retrieval methods in Large Language Models show degradation in accuracy when handling temporally distributed conversations, primarily due to their reliance on simple similarity-based retrieval. Unlike existing memory retrieval methods that rely solely on semantic similarity, we propose SynapticRAG, which uniquely combines temporal association triggers with biologically-inspired synaptic propagation mechanisms. Our approach uses temporal association triggers and synaptic-like stimulus propagation to identify relevant dialogue histories. A dynamic leaky integrate-and-fire mechanism then selects the most contextually appropriate memories. Experiments on four datasets of English, Chinese and Japanese show that compared to state-of-the-art memory retrieval methods, SynapticRAG achieves consistent improvements across multiple metrics up to 14.66% points. This work bridges the gap between cognitive science and language model development, providing a new framework for memory management in conversational systems.¹

1 Introduction

As AI systems become increasingly integrated into daily communications, the ability to maintain natural conversations has become crucial (Dai et al., 2019; Trivedi et al., 2023; Deng et al., 2024). While Large Language Models (LLMs) have achieved remarkable progress in dialogue generation, they face a fundamental challenge: maintaining coherent long-term memory (Zaheer et al., 2020; Brown et al., 2020; Beltagy et al., 2020). This limitation becomes particularly evident in their struggle to retrieve and integrate contextually relevant information from temporally distributed dialogue histories (Sukhbaatar et al., 2015; Gao et al., 2018; Olabiyi et al., 2019; Gutiérrez et al., 2025). For

¹The code is available at: <https://github.com/tamoharu/SynapticRAG-Benchmark>

instance, current state-of-the-art memory retrieval methods show a 17% degradation in response relevance when conversation context spans multiple sessions, significantly impacting user experience in real-world applications.

Inspired by cognitive science research, which shows that human memory retrieval relies heavily on temporal triggers (Tulving, 1985; McDaniel and Einstein, 2000) and synaptic connections between memory nodes (Hebb, 1949; O’Reilly and McClelland, 1994; Norman and O’Reilly, 2003; Neves et al., 2008), rather than pure memory strength. In contrast, current methods primarily use simple similarity-based retrieval, failing to capture these temporal-associative aspects.

Inspired by these cognitive principles, we propose SynapticRAG, a neuroscience-inspired memory retrieval framework. Our approach integrates temporal association triggers, synaptic-like memory connections, and dynamic activation control to achieve more human-like memory retrieval in dialogue systems. This approach is theoretically motivated by the proven effectiveness of biological memory systems in handling temporal dependencies and context integration. Our work makes three key contributions: (1) A novel biologically-inspired memory retrieval mechanism, (2) Integration of temporal and associative memory aspects in memory retrieval systems, and (3) Empirical validation across three languages and datasets, showing consistent improvements up to 14.66% points over state-of-the-art memory retrieval methods on English, Chinese and Japanese benchmarks.

2 Related Works

2.1 Memory Retrieval

Traditional retrieval-augmented generation (RAG) approaches focus primarily on semantic similarity for knowledge retrieval (Lewis et al., 2020; Borgeaud et al., 2022; Guu et al., 2020). While

effective for static knowledge bases, these methods often struggle with maintaining dialogue coherence over extended interactions. Recent studies have explored various temporal enhancement strategies: [Zhong et al. \(2023\)](#) and [Alonso et al. \(2024\)](#) introduced time-based memory updates, while [Hou et al. \(2024\)](#) demonstrated the benefits of dynamic memory consolidation. Despite these advances, significant challenges remain in temporal aspects of task-oriented dialogues ([Wu, 2019](#); [Zhang et al., 2019](#)). The latest developments in self-updatable models ([Wang et al., 2024](#)) and behavioral memory ([Lee et al., 2024](#)) underscore the importance of integrating semantic and temporal information, which our work directly addresses.

2.2 Biological Inspiration

Neuroscience research has revealed that synaptic mechanisms play crucial roles in memory formation and retrieval ([Gerstner and Kistler, 2002](#)). Particularly relevant to our work is the phenomenon of synaptic reverberation in working memory maintenance ([Compte et al., 2000](#); [Wang, 2001](#)). While existing dialogue systems often employ simplified decay mechanisms, biological studies suggest the importance of diverse temporal dynamics in memory processes ([Fell and Axmacher, 2011](#); [Howard et al., 2024](#)). Our SynapticRAG draws upon these insights by implementing biologically-motivated synaptic propagation for temporal-aware memory retrieval, addressing the limitations of conventional approaches that lack biological grounding.

3 Approach

Our approach enhances dialogue memory recall by integrating temporal dynamics into retrieval mechanisms. The model consists of four components, as described in Figure 1.

3.1 Memory Construction and Storage

We embed each dialogue turn into a vector representation and store them as memory nodes in a graph structure. The edges between nodes are weighted by their cosine similarities, reflecting the semantic relationships in the dialogue history. Each node maintains a weighted spike train $[(t_i, s_i)]_{i=1}^n$, where t_i represents the dialogue turn number, s_i denotes the stimulus strength (initialized to 1), and n counts the received stimuli. This temporal information enables our model to capture both the semantic relevance and temporal patterns of dialogue interactions.

3.2 Selecting Memory Nodes

Our selection process combines semantic and temporal filtering to identify relevant memory nodes. Given a new query, we first select candidate nodes whose semantic similarity exceeds threshold \cos_{th} . These candidates are then filtered based on temporal association strength, leveraging the observation that temporally proximate memories often form coherent dialogue chains ([Howard and Kahana, 2002](#); [Dubrow and Davachi, 2013](#)).

Temporal Association Computation We quantify temporal associations between memory nodes using a framework based on Dynamic Time Warping (DTW) ([Sakoe and Chiba, 1978](#)), which enables robust alignment of memory spike trains across time and supports flexible modeling of non-linear temporal shifts. For two memory nodes with spike trains $A = (t_1^a, \dots, t_n^a)$ and $B = (t_1^b, \dots, t_m^b)$, where m, n denotes the largest turn number, we construct a distance matrix $D \in \mathbb{R}^{n \times m}$ where each element $D(i, j) = |t_i^a - t_j^b|$, $1 \leq i \leq n, 1 \leq j \leq m$ represents the absolute temporal difference between stimulus reception timestamps.

To transform temporal distances into association strengths, we apply an exponential decay using the average time constant $\bar{\tau} = 0.5(\tau_A + \tau_B)$, where τ_A and τ_B are node-specific time constants, we construct the association strength matrix $W \in \mathbb{R}^{n \times m}$ as:

$$W(i, j) = \exp\left(\frac{-D(i, j)}{\bar{\tau}}\right) \in (0, 1], \quad (1)$$

where larger $\bar{\tau}$ values maintain robust memory connections with sustained information retention, while smaller values lead to rapid decay. The exponential transformation bounds association strengths within $(0, 1]$, with stronger temporal proximity reflected in higher scores.

Given W , we maximize the cumulative association strength by defining cumulative matrix L as:

$$L(i, 1) = \sum_{k=1}^i W(k, 1), \quad i \in \{1, \dots, n\} \quad (2)$$

$$L(1, j) = \sum_{k=1}^j W(1, k), \quad j \in \{1, \dots, m\} \quad (3)$$

$$L(i, j) = W(i, j) + \max\{L(i-1, j), L(i, j-1), L(i-1, j-1)\} \quad (4)$$

The cumulative score $L(n, m)$ uniquely captures the temporal association strength between sequences A and B .

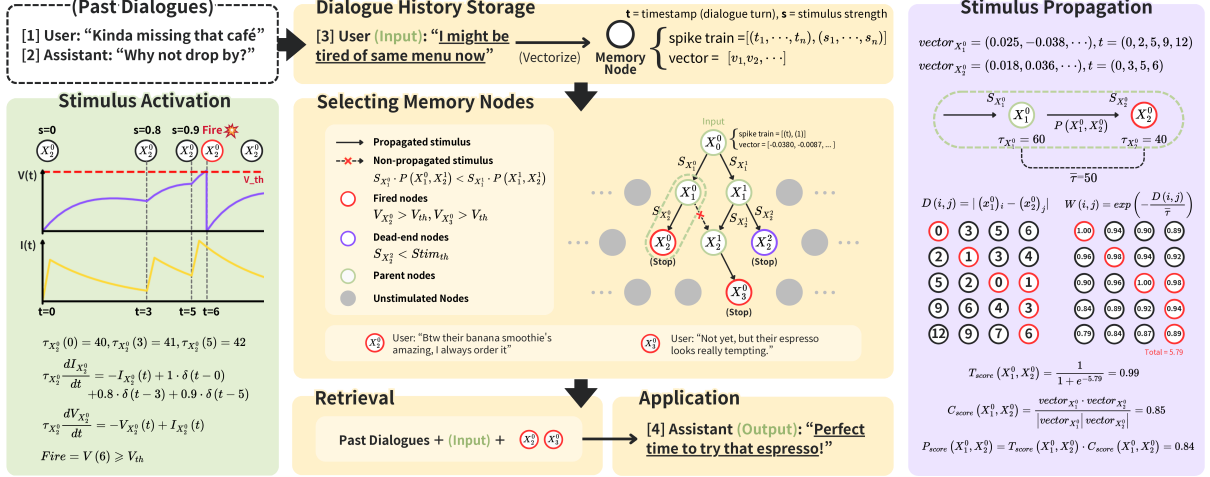


Figure 1: The framework of SynapticRAG. First, past dialogues and current input are stored as memory nodes and formalized into an attributed graph; Second, the stimulus strength is propagated among the graph following propagation mechanisms; Third, nodes are activated through temporal-semantic connections following neuron-inspired dynamics; Finally, responses are generated by integrating selected memories with the input query.

Besides, as $L(n, m)$ may exceed 1, we normalize this cumulative score through sigmoid transformation $\sigma(x) = \frac{1}{1 + \exp(-x)}$:

$$T_{score}(A, B) = \sigma(L(n, m)) \quad (5)$$

Score Integration T_{score} captures the temporal association strength between memory nodes. To integrate temporal and semantic similarities into a unified association measure, we define the propagation score P_{score} as:

$$P_{score}(A, B) = T_{score}(A, B) \cdot C_{score}(A, B) \quad (6)$$

where $C_{score}(A, B) = \frac{A \cdot B}{|A| \cdot |B|}$ quantifies semantic similarity through cosine distance between node vectors. This integration provides a comprehensive measure for subsequent stimulus propagation.

3.3 Stimulus Propagation Mechanism

We organize memory nodes into a layered graph for efficient stimulus propagation. In layer i , nodes that accumulate sufficient stimulus become active and can propagate to connected nodes in layer $i+1$. The propagation strength between nodes follows P_{score} (Eq. (6)), forming contextually coherent memory paths. The propagation strength between parent node X_i and its child node X_{i+1} follows:

$$S_{X_{i+1}} = P_{score}(X_i, X_{i+1}) \cdot S_{X_i}, \quad S_{X_0} = 1 \quad (7)$$

where $S_{X_{i+1}}$ denotes the stimulus strength received at node X_{i+1} , $P_{score}(X_i, X_{i+1})$ quantifies

the temporal and semantic association strengths between parent nodes, S_{X_i} denotes the current stimulus strength at node X_i , and S_{X_0} establishes the baseline stimulus strength for the initial input node.

Memory nodes receive activation signals that propagate along paths in the layered graph. For a propagation path X_0, X_1, \dots, X_n , the stimulus strength at node X_n decays according to:

$$S_{X_n} = S_{X_0} \prod_{i=0}^{n-1} P_{score}(X_i, X_{i+1}) \quad (8)$$

where S_{X_0} is the initial stimulus strength and P_{score} captures the decay along each step.

Stimulus Propagation Path Construction We construct propagation paths through three steps: parent node qualification, child node selection, and path formation. Let \mathcal{N} denote the universal node set, with thresholds $stim_{th}$ and $cost_{th}$ controlling activation and connection sensitivity respectively.

The propagation process follows a layer-wise approach, where qualified parent nodes \mathcal{P}_i at layer i are determined by:

$$\mathcal{P}_i = \{p \in \mathcal{N}_i \mid (S_p > stim_{th}) \wedge (p \notin Fire)\} \quad (9)$$

where S_p denotes the stimulus strength. The thresholds are optimized as detailed in Appendix B, with Fire Mechanism described in Section 3.4.

For each parent node $p \in \mathcal{P}_i$, we define its candidate child node set $\mathcal{C}_i(p)$ for layer $i+1$. $\mathcal{C}_i(p)$ is

constructed as the set of nodes whose cosine similarity with parent node p exceeds threshold \cos_{th} , while excluding all nodes from layers 0 through i :

$$\mathcal{C}_i(p) = \left\{ c \in \mathcal{N} \setminus \bigcup_{j=0}^i \mathcal{N}_j \mid \frac{p \cdot c}{\|p\| \|c\|} > \cos_{\text{th}} \right\} \quad (10)$$

Based on the candidate child node set $\mathcal{C}_i(p)$, we determine the actual stimulus-receiving child nodes. For cases where a child node c receives stimuli from multiple parent nodes, we select the parent node providing the maximum stimulus strength and utilize its stimulus value.

We first formalize a candidate parent node set $\mathcal{P}_i(c)$ for each child node as:

$$\mathcal{P}_i(c) = \left\{ p \in \mathcal{P}_i \mid \frac{p \cdot c}{\|p\| \|c\|} > \cos_{\text{th}} \right\} \quad (11)$$

The optimal parent node p^* that transmits stimulus to child node c is uniquely determined from $\mathcal{P}_i(c)$ through:

$$p^* = \arg \max_{p \in \mathcal{P}_i(c)} S_p \cdot P_{\text{score}}(p, c) \quad (12)$$

We formalize the actual child node set \mathcal{C}_i^* as:

$$\begin{aligned} \mathcal{C}_i^*(p) &= \left\{ c \in \mathcal{C}_i(p) \mid S_p \cdot P_{\text{score}}(c, p) \right. \\ &\quad \left. = \max_{p' \in \mathcal{P}_i} S_{p'} \cdot P_{\text{score}}(p', c) \right\} \end{aligned} \quad (13)$$

The node set \mathcal{N}_{i+1} at layer $i+1$ is constructed as the union of all actual child node sets $\mathcal{C}_i^*(p)$, with the constraint that these child node sets must be mutually disjoint for different parent nodes:

$$\begin{aligned} \mathcal{N}_{i+1} &= \bigcup_{p \in \mathcal{P}_i} \mathcal{C}_i^*(p), \mathcal{C}_i^*(p) \cap \mathcal{C}_i^*(q) = \emptyset \\ &\quad (\forall p, q \in \mathcal{P}_i, p \neq q) \end{aligned} \quad (14)$$

With \mathcal{N}_{i+1} established for layer $i+1$, we iteratively execute this process for subsequent layers. The propagation terminates when there are no further qualified parent nodes, specifically when $\mathcal{P}_i = \emptyset$. Following this propagation framework, we can define the pathway to node X_n using the method for determining parent nodes mentioned earlier, as expressed by the following equation:

$$X_i = \arg \max_{x \in \mathcal{P}_i(X_{i+1})} S_p \cdot P_{\text{score}}(x, X_{i+1}) \quad (15)$$

3.4 Stimulus Activation

Based on the stimulus propagation mechanism (Section 3.3), we formalize how nodes process received stimuli through membrane potential: a dynamic variable that accumulates input stimuli and governs node activation. For a node X , the processing pipeline consists of three steps: (1) converting discrete stimuli into continuous input current, (2) evolving membrane potential with dynamic time constants, and (3) determining activation through threshold-based firing. This framework enables adaptive memory selection in dialogue systems by emulating neurobiological principles (Dayan and Abbott, 2005).

Input Current Generation The Leaky Integrate-and-Fire (LIF) model (Brunel and Rossum, 2007) provides a biologically inspired abstraction for simulating how stimulus input accumulates and decays over time within memory nodes, forms the basis for processing temporal signals in dialogue systems, however, its fixed membrane time constant τ limits its ability to balance recent inputs with long-term memories in dialogue systems. To alleviate this, we propose an enhanced LIF model with a dynamic time constant τ_X and a spike train-based input accumulation mechanism. The discrete input stimuli $S_X(t)$ are converted into continuous input current $I_X(t)$ through:

$$\tau_X \frac{dI_X}{dt} = -I_X(t) + \sum_{t_s \in \Gamma} S_X(t) \delta(t - t_s) \quad (16)$$

where Γ denotes input pulse timestamps, t_s represents stimulus reception, and $\delta(t - t_s)$ is the Dirac delta function for discrete pulse modeling (Bracewell, 2000). The leakage term $-I_X(t)$ simulates neuronal ion channels, enabling natural decay of past stimuli while maintaining responsiveness to new inputs. When no stimulus is present, the summation term becomes zero, allowing gradual memory decay analogous to biological systems.

Membrane Potential Evolution Node state is represented by membrane potential $V_X(t)$, reflecting both input stimulus accumulation and natural memory decay, which evolves according to:

$$\tau_X \frac{dV_X}{dt} = -V_X(t) + I_X(t) \quad (17)$$

where the decay term $-V_X(t)$ prevents unbounded accumulation while $I_X(t)$ incorporates

new inputs. This evolution enables selective activation as nodes fire only when $V_X(t)$ exceeds a predefined threshold. To balance memory retention with input sensitivity, we introduce a dynamic membrane time constant τ_X that updates based on input intervals Δt :

$$\tau_X(t + \Delta t) = \tau_X(t) + \frac{1 - \exp(-\Delta t)}{1 + \exp(-\Delta t)} \quad (18)$$

This sigmoid mapping ensures that τ_X remains stable (approaches 0) for frequent inputs to maintain sensitivity, while increasing (approaches 1) for sparse inputs to preserve long-term memories. τ_X starts with its optimized initial value τ_{init} , which is fixed for all nodes. This mechanism effectively emulates synaptic plasticity in biological neurons.

Node Activation and Fire Mechanism A node enters the Fired state when its membrane potential exceeds threshold V_{th} , with the firing nodes at time t defined as $\text{Fire}(t) = \{X \in \mathcal{N} \mid V_X(t) \geq V_{\text{th}}\}$. Post-firing, nodes reset to optimized resting states ($V_{\text{rest}}, I_{\text{rest}}$) (See Appendix C), preventing continuous activation and enabling contextually relevant memory selection. The dialogue text of nodes in fired state is incorporated into the prompt for response generation, thus contributing to the system’s contextual memory retrieval.

4 Experiments

4.1 Dataset

Temporal information is critical for dialogue response generation, requiring access to historical context for time-dependent interactions. Effective evaluation of temporal-aware dialogue capabilities requires datasets with three key properties: (1) multi-turn dialogues that encompass complex temporal contexts spanning multiple interactions across distinct temporal segments, (2) complex temporal dependencies where correct memory retrieval depends critically on temporal context rather than solely semantic similarity, and (3) explicit annotations identifying correct memory retrievals, providing ground-truth labels for quantitatively evaluating retrieval accuracy. Currently, only **PerLTQA** (Du et al., 2024) partially addresses these requirements through its format with temporal dependencies, containing 625 English and Chinese dialogues.

To fill this gap, we introduce Synaptic Memory Retrieval Conversations **SMRCs**, a bilingual dataset in English and Japanese consisting of

101 dialogues with 456 tasks, constructed using Claude 3.5 Sonnet (Anthropic, 2024) (See generation guidelines in Appendix A). Each task consists of two segments: a past and a present dialogue, with 10 turns of user-assistant interactions each.

For quantitative evaluation, we annotate each task with multiple trigger-memory pairs, where a trigger is an utterance in the present dialogue referencing the past context, and memories are relevant utterances from the past dialogue. To maintain natural conversational flow and consistent information density, utterances are controlled at lengths between 10 and 20 words. Temporal relationships between triggers and memories typically span multiple turns, requiring models to comprehend dialogue contexts beyond superficial semantic matching. Figure 4 provides a representative example.

4.1.1 Annotation Reliability Evaluation

To ensure semantic authenticity and annotation validity, we implemented a rigorous filtering method for dataset construction. Specifically, annotations were independently performed by four annotators, and only annotations meeting one of two strict conditions were retained: (1) annotations were exactly identical, with at least two memory references; or (2) annotations partially overlapped, and the intersection contained at least two memory references, retaining only the overlapping annotations.

Table 1: Annotation results from four annotators (A1–A4) across 10 tasks. Each cell indicates the indices of selected utterances.

Task ID	A1	A2	A3	A4
1	[4,5]	[4,5]	[5,6]	[4,5]
2	[6,7]	[6,7]	[4,6]	[6,7]
3	[2,4]	[2,4]	[2,4]	[2,4]
4	[8,9]	[8,9]	[8,9]	[6,8]
5	[0,2]	[0,2]	[0,2]	[0,2]
6	[5,7]	[5,7]	[5,7]	[5,7]
7	[2,7]	[2,7]	[2,7]	[2,6]
8	[6,8]	[6,8]	[8,9]	[8,9]
9	[5,6]	[5,6]	[5,6]	[5,6]
10	[0,7]	[0,7]	[0,6]	[6,7]

Furthermore, we conducted an inter-annotator agreement evaluation using Krippendorff’s Alpha (Krippendorff, 2011). We selected 10 representative tasks from the dataset, covering diverse dialogue scenarios. Annotators independently followed clearly defined annotation guidelines: each annotator was given a set of 10 past dialogue utter-

ances and one trigger utterance from the present dialogue and instructed to select 2 past utterances relevant for accurately comprehending and responding to the trigger. For calculating agreement, we constructed a binary selection matrix representing each of the 10 past utterances per task, assigning a value of 1 if selected by an annotator and 0 otherwise. This matrix was analyzed using a nominal scale to compute inter-annotator agreement. The annotation results are summarized in Table 1.

Based on these annotations, we computed a Krippendorff’s Alpha agreement value of 0.782. Typically, a Krippendorff’s Alpha value ≥ 0.8 indicates excellent reliability, while a value ≥ 0.67 is considered an acceptable minimum standard (Krippendorff, 2004). This result demonstrates robust consistency and strong consensus among annotators, with limited subjectivity in the annotations.

4.2 Baselines and Model Adaptations

We compare SynapticRAG with five baseline models. The three base models are standard RAG, MemoryBank (Zhong et al., 2023), MyAgent (Hou et al., 2024), and our SynapticRAG. For fair comparison against SynapticRAG’s temporal-aware retrieval mechanism, we also developed adapted versions of MemoryBank and MyAgent.

RAG serves as our fundamental baseline, as its vector similarity search forms the foundation of retrieval process and provides a robust benchmark for evaluating basic retrieval capabilities.

MemoryBank incorporates Ebbinghaus Forgetting Curve-inspired (Graves et al., 2016) mechanisms to dynamically adjust memory based on temporal factors, representing an improved memory management approach.

MyAgent represents a cognitive-inspired approach to memory retrieval, adopting human memory cue recall as a trigger and developing a mathematical model that dynamically quantifies memory consolidation by considering contextual relevance r , elapsed time t , and recall probability p .

MemoryBank (Adt) enhances the original MemoryBank’s temporal awareness through additional parameters including top- K , Θ , and scaling factors (t_{scale} , s_{scale} , s_{init}) to enable fair comparison with temporal-aware retrieval mechanisms.

MyAgent (Adt) modifies the original retrieval process to select the same number of memories as SynapticRAG from those exceeding the cosine similarity threshold, ranked by recall probability p , ensuring consistent comparison conditions.

4.3 Metrics

We evaluate SynapticRAG using both traditional and dynamic retrieval metrics. Traditional metrics like $\text{recall}@k$ and $\text{precision}@k$ measure retrieval performance with a fixed k number of memories (Manning et al., 2008; Hambarde and Proença, 2023; Croft et al., 2009). Unlike the baseline methods, SynapticRAG does not inherently retrieve a fixed number of memories. To fairly compute and compare Precision and Recall metrics, we implemented the following procedure:

For each retrieval scenario, we ranked memory nodes according to their membrane potentials (V) computed by SynapticRAG. Subsequently, we selected the top- k memory nodes with the highest V values to match the number of memories retrieved by the baseline methods. Precision and Recall were then calculated based on these top- k selected memories. Since V in SynapticRAG directly determines whether a node fires (i.e., is retrieved), ranking memories by membrane potential reflects the model’s intrinsic memory-selection mechanism, ensuring that our evaluation is both methodologically sound and consistent with the biological inspiration underlying SynapticRAG.

Equal Retrieval Count (ERC) However, fixed- k retrieval becomes problematic as memory needs vary by context - from multiple memories for complex past events to few for simple queries. To address this, we propose ERC where retrieval count adapts dynamically: when our method retrieves k memories for a query, other methods retrieve their top- k ranked results, with accuracy calculated as the ratio of correct matches to total correct labels. Additionally, for queries requiring at least m correct memories, we allow baselines to retrieve up to m memories when the evaluated method retrieves fewer, denoted as Equal Retrieval Count with Minimum Guarantee, *i.e.*, ERC-MG. The ERC-MG metric provides a stricter benchmark than ERC by ensuring baseline methods retrieve at least the annotated number of memories whenever SynapticRAG retrieves fewer. This setup rigorously evaluates SynapticRAG’s robustness under conditions favorable to baseline methods.

Notably, SynapticRAG achieves consistently high performance under the ERC-MG metric across all datasets, surpassing the strongest baseline by 7.76%-14.66%. Given that ERC-MG imposes more stringent retrieval requirements by allowing baseline methods the opportunity to retrieve

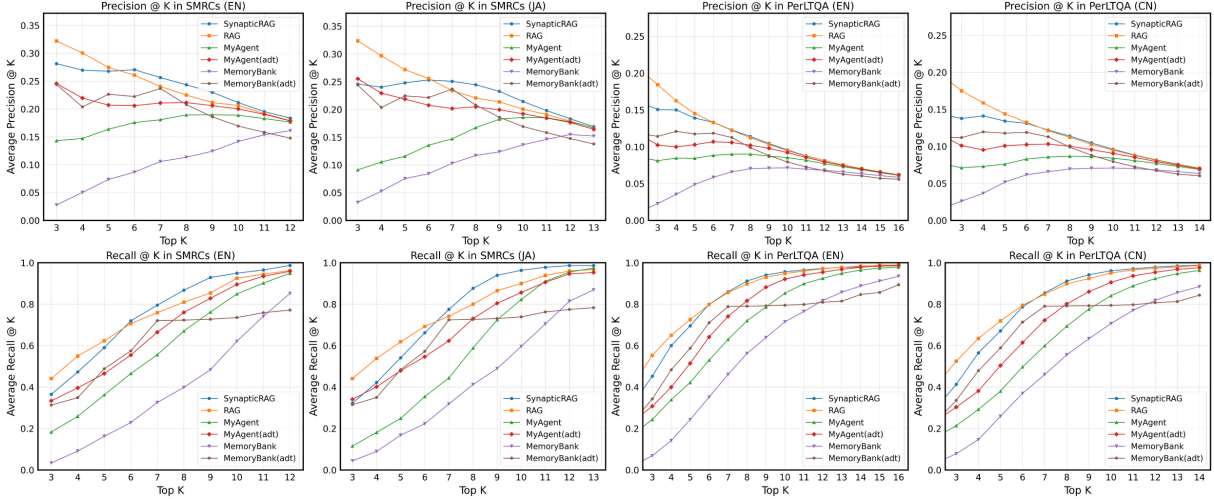


Figure 2: Traditional memory retrieval metrics values (Precision@ k and Recall@ k , $k = [3, \dots, 12]$) across four tested datasets and five baselines.

Table 2: Dynamic memory retrieval metrics values (ERC and ERC-MG). Δ denotes the improvement of SynapticRAG over the best baseline model.

	SMRCs ^{EN}		SMRCs ^{JA}		PerLTQA ^{EN}		PerLTQA ^{CN}		Avg.
	ERC	ERC-MG	ERC	ERC-MG	ERC	ERC-MG	ERC	ERC-MG	
RAG	71.55	71.55	81.58	81.58	82.14	82.51	84.95	85.36	80.15
MemoryBank	32.60	32.82	49.34	49.34	33.62	33.62	36.75	36.75	38.11
MemoryBank _(Adt)	62.36	62.58	72.15	72.15	67.24	67.54	71.31	71.61	68.37
MyAgent	62.14	62.14	70.61	70.61	61.99	62.29	65.14	65.47	65.05
MyAgent _(Adt)	63.24	63.46	77.19	77.19	66.49	66.64	69.54	69.76	69.19
SynapticRAG	86.21	86.21	92.32	92.32	90.47	90.47	93.12	93.12	90.53
$\Delta\%$	+14.66	+14.66	+10.74	+10.74	+8.33	+7.96	+8.17	+7.76	+10.38

all annotated memories, this result is particularly significant.

4.4 Implementation Details

Our implementation uses text-embedding-3-large (OpenAI, 2024a) for text vectorization and GPT-4o (OpenAI, 2024b) for generation. The memory management parameters include cosine similarity threshold $\cos_{th} = 0.262$ and temporal decay factor $\tau_{init} = 43.07$. For the LIF model, we set membrane potential threshold $V_{th} = 0.099$, stimulus threshold $\text{stim}_{th} = 0.037$, and post-firing reset values $V_{rest} = 2.903$, $I_{rest} = -7.128$. Complete hyperparameter settings are available in Table 4.

5 Results

5.1 Main Results

Experimental results shown in Figure 2 and Table 2 demonstrate that SynapticRAG achieves consistent performance improvements across four datasets:

SMRCs (EN), SMRCs (JA), PerLTQA (EN) and PerLTQA (CN). Compared to the strongest baseline, our method shows an average improvement of 10.38% in memory retrieval accuracy.

Specifically, from a cross-dataset perspective, SynapticRAG maintains stable advantages on both SMRCs with complex temporal dependencies and PerLTQA with natural dialogue flow. This consistency extends across languages (English, Japanese, and Chinese), validating our model’s language-agnostic nature and generalization capability.

In terms of performance metrics, SynapticRAG successfully balances precision and recall, particularly excelling when $k \geq 6$, indicating superior handling of complex dialogues requiring multiple related memories. On SMRCs (EN), it achieves 86.21% accuracy under ERC evaluation, surpassing the strongest baseline by 14.66%. Similar improvements are observed across other datasets, with PerLTQA (CN) reaching 93.12% accuracy.

The performance curves reveal that while preci-

Table 3: Performance improvement of SynapticRAG and its modules ($\Delta\%$).

Model Ablations	SMRCs ^{EN}	SMRCs ^{JA}	PerLTQA ^{EN}	PerLTQA ^{CN}	Avg.
<i>full</i> SynapticRAG	14.66	10.74	8.33	8.17	10.41
<i>w/o</i> Temporal Integration	-11.82	-8.55	-9.68	-9.54	-9.90
<i>w/o</i> Stimulus Propagation	-8.31	-2.41	-3.60	-3.36	-4.35
<i>w/o</i> Dynamic Time Constant	-14.66	-10.96	-8.29	-8.24	-10.54

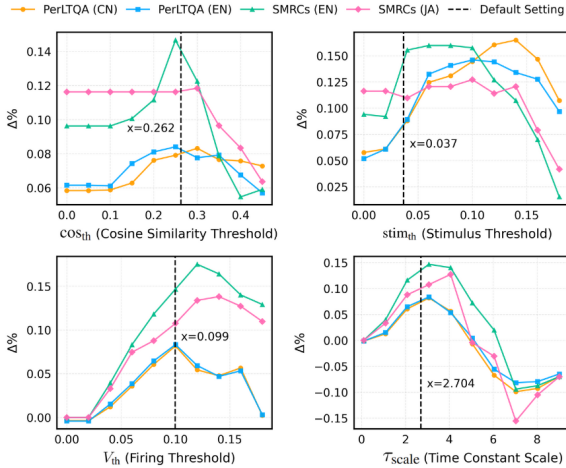


Figure 3: Effect of key parameters on SynapticRAG’s performance gain over RAG. $\Delta\%$ indicates the accuracy improvement relative to RAG, and Default Setting denotes hyperparameter-optimal values.

sion decreases with increasing k for all methods, SynapticRAG maintains a more gradual decline. Its recall curves consistently dominate other methods, demonstrating effective identification of relevant memories across temporal spans.

These results validate SynapticRAG’s effectiveness in enhancing temporal memory retrieval, particularly in maintaining high recall while preserving precision. This balanced performance is especially valuable for applications requiring accurate memory retrieval across dialogue histories.

5.2 Ablation Study

Our method comprises three key components: a) Temporal Integration capturing temporal dependencies through DTW; b) Stimulus Propagation establishing synaptic-like memory connections through optimal propagation pathways; and c) Stimulus Activation controlling memory firing through dynamic LIF model, where decay factor τ balances recent and long-term memories. We evaluate each component’s contribution through individual removal. Specifically, we set $T_{\text{score}}=1$ to remove tem-

poral association (*w/o* Temporal Integration) for examining pure semantic similarity-based retrieval; disable stimulus propagation by setting $\text{stim}_{\text{th}}=1$ (*w/o* Stimulus Propagation) to observe performance when memories are independent and unaffected by each other’s stimulation; and fix τ to its initial value τ_{init} (*w/o* Dynamic Time Constant) to examine the effect of prioritizing recent inputs without balancing long-term memory consideration.

Results in Table 3 demonstrate that each module is crucial, with their removal leading to significant performance degradation across all datasets. Specifically, removing Temporal Integration causes an average performance drop of 9.90%, indicating that semantic similarity alone struggles to capture complex temporal dependencies in dialogues; disabling Stimulus Propagation results in an average drop of 4.35%, validating the necessity of memory interconnections; removing Dynamic Time Constant leads to substantial degradation with an average drop of 10.54%, demonstrating the critical role of balancing recent inputs and long-term memories. These findings validate both our component design and biologically-inspired principles, particularly on SMRCs with complex temporal dependencies.

5.3 Parameter Analysis

Building upon our validation of individual components, we analyze SynapticRAG’s parameter sensitivity through four fundamental parameters: cos_{th} for cosine similarity threshold, stim_{th} for stimulus propagation threshold, V_{th} for firing threshold, and τ_{scale} for scaling dynamic time constant. While Optuna optimization identified their optimal values, understanding their impacts remains crucial for both theoretical and practical purposes.

As shown in Figure 3, our analysis reveals distinct parameter sensitivities. The filtering threshold cos_{th} achieves optimal performance at 0.262 within $[0.0, 0.4]$, with SMRCs showing higher sensitivity than PerLTQA due to their complex semantic relationships. The stimulus threshold stim_{th} exhibits

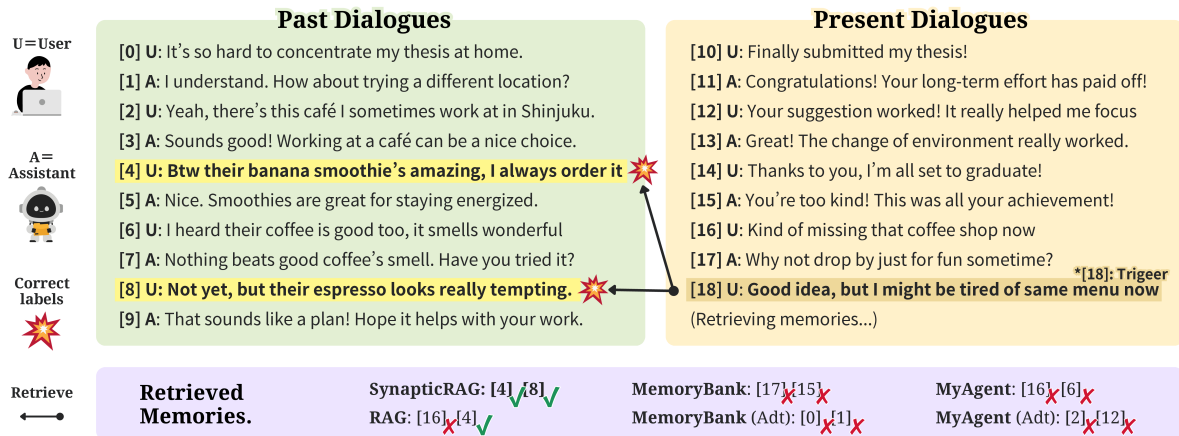


Figure 4: A case study illustrating how different models retrieve memories.

stable behavior around 0.037 in $[0.0, 0.17]$, where higher values block essential connections while lower values cause over-propagation. The firing threshold V_{th} shows moderate sensitivity at 0.099 within $[0.0, 0.2]$, balancing between over-activation and insufficient utilization.

Most critically, τ_{scale} demonstrates dramatic performance variations, peaking at 2.704 with a sharp decline beyond this value (negative $\Delta\%$). These findings suggest task-specific tuning strategies: precise adjustment for semantic-heavy tasks, stable propagation thresholds for general applications, and potentially adaptive mechanisms for temporal scaling. The identified robust regions can guide efficient parameter tuning in practical deployments.

5.4 Case Study

To demonstrate SynapticRAG's temporal memory capabilities, we present a case from SMRCs where a user discusses their thesis writing experience at a café (Figure 4). When the user expresses potential menu fatigue at turn [18], SynapticRAG successfully retrieves the correct label memories about beverage preferences ([4] and [8]) from past dialogue, where they mentioned a banana smoothie and interest in espresso. This temporal association enables a contextually appropriate response that acknowledges both their past interest and current desire for change. RAG retrieves one correct memory ([4]) but misses the temporal connection ([16]), MemoryBank and MyAgent fail to retrieve either correct memory, and their adapted versions show similar temporal association failures. This example illustrates SynapticRAG's ability to form coherent memory chains across temporal gaps while maintaining contextual relevance, a key advantage of

our synaptic-inspired propagation mechanism.

6 Conclusion

In this paper, we introduced SynapticRAG, a biologically-inspired memory retrieval model that enhances temporal memory management in dialogue systems. By integrating temporal association triggers, synaptic-like stimulus propagation, and dynamic memory activation mechanisms, our approach effectively addresses the limitations of simple similarity-based retrieval in handling temporally distributed conversations. Our experiments demonstrate while RAG excels in simpler straight-forward retrieval tasks, SynapticRAG demonstrates superior performance in temporal-aware dialogue contexts where understanding dialogue history is critical, achieving consistent improvements across multiple languages and dialogue scenarios.

7 Limitations

Over-retrieval Risk. The stimulus propagation mechanism may trigger excessive memory node activations through synaptic connections, particularly in complex multi-topic dialogues. Such over-retrieval leads to two main issues: lengthy generation prompts as all retrieved memory texts are incorporated for response generation, and increased computational overhead due to linear growth of memory storage with dialogue length. To mitigate these issues, early stopping criteria and pruning strategies could be implemented to terminate unnecessary activation paths while preserving essential temporal connections, and importance sampling could be used to maintain only the most crucial historical information.

Computational Complexity. The original SynapticRAG architecture incurs computational complexity on the order of $O(n^2)$ in memory retrieval, where n represents the total number of memory nodes. This quadratic complexity arises from the cosine-similarity filtering process, as each propagation step involves multiple parent nodes independently performing full scans through the entire memory database. To address this limitation, we have developed an optimized architecture that computes a single representative centroid vector based on all parent nodes, reducing the complexity from $O(n^2)$ to $O(n)$. See details in Appendix D.

Parameter Sensitivity. The optimization of model parameters requires more extensive tuning compared to simpler models, and the variable number of recalled memories presents challenges when strict control is needed. Different language contexts may require separate parameter tuning. This could be mitigated through the development of automated parameter adaptation mechanisms that dynamically adjust based on input characteristics and language-specific features.

Model Interpretability. The complex stimulus propagation and firing mechanisms reduce interpretability compared to simpler similarity-based approaches. Tracing memory selection reasoning becomes challenging due to dynamic synaptic connections. However, this could potentially be solved by introducing attention-like mechanisms to quantify the contribution of each synaptic connection to the final memory selection, providing insights into the model's decision process.

Acknowledgments

The work of Qinghua Zhao was partially supported by the Hefei College Talent Research Fund Project (No. 24RC20, 21-22RC13), and Natural Science Foundation of the Anhui Higher Education Institutions of China (No. 2022AH051779).

References

- Nick Alonso, Tomás Figliolia, Anthony Ndirango, and Beren Millidge. 2024. [Toward conversational agents with context and time sensitive long-term memory](#). *Preprint*, arXiv:2406.00057.
- Anthropic. 2024. Claude 3.5 sonnet. <https://www.anthropic.com>. Large language model, accessed February 3, 2025.

Iz Beltagy, Matthew E. Peters, and Arman Cohan. 2020. [Longformer: The long-document transformer](#). *Preprint*, arXiv:2004.05150.

Sebastian Borgeaud, Arthur Mensch, Jordan Hoffmann, Trevor Cai, Eliza Rutherford, Katie Millican, George van den Driessche, Jean-Baptiste Lespiau, Bogdan Damoc, Aidan Clark, Diego de Las Casas, Aurelia Guy, Jacob Menick, Roman Ring, Tom Hennigan, Saffron Huang, Loren Maggiore, Chris Jones, Albin Cassirer, Andy Brock, Michela Paganini, Geoffrey Irving, Oriol Vinyals, Simon Osindero, Karen Simonyan, Jack W. Rae, Erich Elsen, and Laurent Sifre. 2022. [Improving language models by retrieving from trillions of tokens](#). *Preprint*, arXiv:2112.04426.

R.N. Bracewell. 2000. *The Fourier Transform and Its Applications*. Circuits and systems. McGraw Hill.

Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, Sandhini Agarwal, Ariel Herbert-Voss, Gretchen Krueger, Tom Henighan, Rewon Child, Aditya Ramesh, Daniel Ziegler, Jeffrey Wu, Clemens Winter, Chris Hesse, Mark Chen, Eric Sigler, Mateusz Litwin, Scott Gray, Benjamin Chess, Jack Clark, Christopher Berner, Sam McCandlish, Alec Radford, Ilya Sutskever, and Dario Amodei. 2020. [Language models are few-shot learners](#). In *Advances in Neural Information Processing Systems*, volume 33, pages 1877–1901. Curran Associates, Inc.

Nicolas Brunel and Mark C. Rossum. 2007. Lapicque's 1907 paper: from frogs to integrate-and-fire. *Biol. Cybern.*, 97(5–6):337–339.

Albert Compte, Nicolas Brunel, Patricia S. Goldman-Rakic, and Xiao-Jing Wang. 2000. [Synaptic mechanisms and network dynamics underlying spatial working memory in a cortical network model](#). *Cerebral Cortex*, 10(9):910–923.

Bruce Croft, Donald Metzler, and Trevor Strohman. 2009. *Search Engines: Information Retrieval in Practice*, 1st edition. Addison-Wesley Publishing Company, USA.

Zihang Dai, Zhilin Yang, Yiming Yang, Jaime G. Carbonell, Quoc V. Le, and Ruslan Salakhutdinov. 2019. [Transformer-xl: Attentive language models beyond a fixed-length context](#). In *Annual Meeting of the Association for Computational Linguistics*.

Peter Dayan and L. F. Abbott. 2005. *Theoretical Neuroscience: Computational and Mathematical Modeling of Neural Systems*. The MIT Press.

Xi Deng, Shasha Li, Jie Yu, and Jun Ma. 2024. [Hm-transformer: Hierarchical multi-modal transformer for long document image understanding](#). In *Web and Big Data: 7th International Joint Conference, APWeb-WAIM 2023, Wuhan, China, October 6–8, 2023, Proceedings, Part IV*, page 232–245, Berlin, Heidelberg. Springer-Verlag.

- Yiming Du, Hongru Wang, Zhengyi Zhao, Bin Liang, Baojun Wang, Wanjun Zhong, Zezhong Wang, and Kam-Fai Wong. 2024. [PerLTQA: A personal long-term memory dataset for memory classification, retrieval, and fusion in question answering](#). In *Proceedings of the 10th SIGHAN Workshop on Chinese Language Processing (SIGHAN-10)*, pages 152–164, Bangkok, Thailand. Association for Computational Linguistics.
- Sarah Dubrow and Lila Davachi. 2013. [The influence of context boundaries on memory for the sequential order of events](#). *Journal of experimental psychology. General*, 142 4:1277–86.
- Juergen Fell and Nikolai Axmacher. 2011. [The role of phase synchronization in memory processes](#). *Nature reviews. Neuroscience*, 12:105–18.
- Jianfeng Gao, Michel Galley, and Lihong Li. 2018. [Neural approaches to conversational AI](#). In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics: Tutorial Abstracts*, pages 2–7, Melbourne, Australia. Association for Computational Linguistics.
- Wulfram Gerstner and Werner M. Kistler. 2002. *Spiking Neuron Models: Single Neurons, Populations, Plasticity*. Cambridge University Press.
- Alex Graves, Greg Wayne, Malcolm Reynolds, Tim Harley, Ivo Danihelka, Agnieszka Grabska-Barwińska, Sergio Gómez, Edward Grefenstette, Tiago Ramalho, John Agapiou, Adrià Badia, Karl Hermann, Yori Zwols, Georg Ostrovski, Adam Cain, Helen King, Christopher Summerfield, Phil Blunsom, Koray Kavukcuoglu, and Demis Hassabis. 2016. [Hybrid computing using a neural network with dynamic external memory](#). *Nature*, 538.
- Bernal Jiménez Gutiérrez, Yiheng Shu, Yu Gu, Michihiro Yasunaga, and Yu Su. 2025. [Hipporag: Neurobiologically inspired long-term memory for large language models](#). *Preprint*, arXiv:2405.14831.
- Kelvin Guu, Kenton Lee, Zora Tung, Panupong Pasupat, and Ming-Wei Chang. 2020. [Realm: Retrieval-augmented language model pre-training](#). *Preprint*, arXiv:2002.08909.
- Kailash A. Hambarde and Hugo Proença. 2023. [Information retrieval: Recent advances and beyond](#). *IEEE Access*, 11:76581–76604.
- Donald Olding Hebb. 1949. *The organization of behavior; a neuropsychological theory*. Wiley.
- Yuki Hou, Haruki Tamoto, and Homei Miyashita. 2024. [My agent understands me better: Integrating dynamic human-like memory recall and consolidation in llm-based agents](#). In *Extended Abstracts of the CHI Conference on Human Factors in Computing Systems, CHI EA '24*, New York, NY, USA. Association for Computing Machinery.
- Marc W. Howard, Zahra G. Esfahani, Bao Le, and Per B. Sederberg. 2024. [Learning temporal relationships between symbols with laplace neural manifolds](#). *Preprint*, arXiv:2302.10163.
- Marc W. Howard and Michael J. Kahana. 2002. [A distributed representation of temporal context](#). *Journal of Mathematical Psychology*, 46(3):269–299.
- Klaus Krippendorff. 2004. *Content Analysis: An Introduction to Its Methodology (second edition)*. Sage Publications.
- Klaus Krippendorff. 2011. [Computing krippendorff's alpha-reliability](#).
- Janghwan Lee, Seongmin Park, Sukjin Hong, Minsoo Kim, Du-Seong Chang, and Jungwook Choi. 2024. [Improving conversational abilities of quantized large language models via direct preference alignment](#). In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 11346–11364, Bangkok, Thailand. Association for Computational Linguistics.
- Patrick Lewis, Ethan Perez, Aleksandra Piktus, Fabio Petroni, Vladimir Karpukhin, Naman Goyal, Heinrich Küttler, Mike Lewis, Wen-tau Yih, Tim Rocktäschel, Sebastian Riedel, and Douwe Kiela. 2020. [Retrieval-augmented generation for knowledge-intensive nlp tasks](#). In *Proceedings of the 34th International Conference on Neural Information Processing Systems, NIPS '20*. Curran Associates Inc.
- Christopher D. Manning, Prabhakar Raghavan, and Hinrich Schütze. 2008. *Introduction to Information Retrieval*. Cambridge University Press.
- Mark A McDaniel and Gilles O Einstein. 2000. [Strategic and automatic processes in prospective memory retrieval: a multiprocess framework](#). *Applied Cognitive Psychology*, 14.
- Guilherme Neves, Sam Cooke, and Tim Bliss. 2008. [Synaptic plasticity, memory and the hippocampus: a neural network approach to causality](#). *Nature Reviews Neuroscience*, 9:65–75.
- Kenneth Norman and Randall O'Reilly. 2003. [Modeling hippocampal and neocortical contributions to recognition memory: A complementary learning systems approach](#). *Psychological review*, 110:611–46.
- Oluwatobi Olabiyi, Alan O Salimov, Anish Khazane, and Erik Mueller. 2019. [Multi-turn dialogue response generation in an adversarial learning framework](#). In *Proceedings of the First Workshop on NLP for Conversational AI*, pages 121–132, Florence, Italy. Association for Computational Linguistics.
- OpenAI. 2024a. [Embeddings - openai api](#). <https://platform.openai.com/docs/guides/embeddings>. (Accessed on 05/21/2024).
- OpenAI. 2024b. [Gpt-4o system card](#). *Preprint*, arXiv:2410.21276.

- Randall C. O'Reilly and James L. McClelland. 1994. [Hippocampal conjunctive encoding, storage, and recall: Avoiding a trade-off](#). *Hippocampus*, 4.
- H Sakoe and S Chiba. 1978. [Dynamic programming algorithm optimization for spoken word recognition](#). *IEEE Transactions on Acoustics, Speech, and Signal Processing*, 26(1):43–49.
- Sainbayar Sukhbaatar, Arthur Szlam, Jason Weston, and Rob Fergus. 2015. End-to-end memory networks. In *Proceedings of the 29th International Conference on Neural Information Processing Systems - Volume 2, NIPS'15*, page 2440–2448, Cambridge, MA, USA. MIT Press.
- Harsh Trivedi, Niranjan Balasubramanian, Tushar Khot, and Ashish Sabharwal. 2023. [Interleaving retrieval with chain-of-thought reasoning for knowledge-intensive multi-step questions](#). In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 10014–10037, Toronto, Canada. Association for Computational Linguistics.
- Endel Tulving. 1985. *Elements of episodic memory*. Number 2 in Oxford psychology series. Clarendon Press, Oxford University Press.
- Xiao-Jing Wang. 2001. [Synaptic reverberation underlying mnemonic persistent activity](#). *Trends in Neurosciences*, 24:455–463.
- Yu Wang, Xinshuang Liu, Xiusi Chen, Sean O'Brien, Junda Wu, and Julian J. McAuley. 2024. [Self-updatable large language models with parameter integration](#). *ArXiv*, abs/2410.00487.
- Chien-Sheng Wu. 2019. [Learning to memorize in neural task-oriented dialogue systems](#). *ArXiv*, abs/1905.07687.
- Manzil Zaheer, Guru Guruganesh, Kumar Avinava Dubey, Joshua Ainslie, Chris Alberti, Santiago Ontanon, Philip Pham, Anirudh Ravula, Qifan Wang, Li Yang, and Amr Ahmed. 2020. [Big bird: Transformers for longer sequences](#). In *Advances in Neural Information Processing Systems*, volume 33, pages 17283–17297. Curran Associates, Inc.
- Zheng Zhang, Minlie Huang, Zhongzhou Zhao, Feng Ji, Haiqing Chen, and Xiaoyan Zhu. 2019. [Memory-augmented dialogue management for task-oriented dialogue systems](#). *ACM Trans. Inf. Syst.*, 37(3).
- Wanjun Zhong, Lianghong Guo, Qiqi Gao, He Ye, and Yanlin Wang. 2023. [Memorybank: Enhancing large language models with long-term memory](#). *Preprint*, arXiv:2305.10250.

A SMRCs Dataset Generation Guidelines

We used the following prompt to generate the SMRCs dataset:

Please generate a dialogue dataset following these guidelines:

1. Temporal Structure:

- Create conversations at two time points: past (1-6 months ago) and present
- Each conversation set should contain approximately 10 turns of dialogue

2. Conversation Design:

- Topics should cover a broad range including daily life, work, school, and personal consultations
- Focus on personalized conversations rather than generic exchanges
- Keep each turn concise, limiting to 10-20 phrases per interaction

3. Memory Retrieval Design:

- The assistant should naturally recall past conversations based on the user's current statements
- Assistant's responses should reference specific recalled utterances from either the user or assistant
- Responses must maintain logical coherence with the recalled memories

4. Memory Trigger Requirements:

- Each trigger should be a single user utterance from the present conversation
- Recalled memories should be complete utterances from past conversations (either user or assistant)
- Multiple trigger-memory pairs within a single conversation are preferred

5. Quality Criteria:

- Ensure natural conversation flow with appropriate memory recalls
- Prioritize conceptual relevance over simple keyword matching
- Maintain clear and appropriate relationships between triggers and recalled memories

6. Documentation Requirements:

- For each conversation set, explicitly document:
- Trigger: The specific user utterance in the present conversation that prompts recall
 - Recalled Memories: The complete related utterances from past conversations
 - All memories should be documented without abbreviation

7. Output Format:

- Provide results in JSON format only

Note: The focus should be on creating natural dialogue flows where memory recalls are contextually appropriate and based on potential conceptual relationships rather than simple pattern matching.

B Hyperparameter Optimization

We optimized hyperparameters using Optuna, ensuring the evaluation metric was independent of comparative models. Optuna iteratively samples parameters, evaluates performance, and refines sampling to efficiently explore the parameter space and identify optimal configurations, eliminating the need for exhaustive manual tuning. When multiple trials achieved the same best score, we selected the parameter set with the highest ERC score to ensure robust performance. The number of hyperparameters varied by model: MemoryBank (0), MemoryBank(Adapted) (5), MyAgent (1), MyA-

gent(Adapted) (4), and SynapticRAG (9).

B.1 Optimization for Adapted Models

For adapted models, where the number of retrieved memories cannot be predetermined, we optimized to maximize the scores of correct label memories. Each memory score s_i was normalized to $[0, 1]$ using the total score sum $Sum_n = \sum_{i=1}^{M_n} s_i$, with M_n being the number of memories in the n -th task. The normalized scores of correct label memories, indexed by y_n , were summed as $Score_n = \sum_{j \in y_n} s_j / Sum_n$. The final score across all N tasks, $Score = \sum_{n=1}^N Score_n$, was maximized by optimizing hyperparameters.

B.2 Optimization for SynapticRAG

For SynapticRAG, which uses the score v for threshold judgment, optimization was based on conditions related to fired memories. Two conditions were imposed. The first penalized insufficient correct label memories, defined as $A_n = \frac{|y_n| - |f_n \cap y_n|}{|y_n|}$ if $|y_n| > |f_n \cap y_n|$, and 0 otherwise, where f_n is the set of fired memory indices and y_n is the set of correct label indices. The second penalized excessive memory retrieval, expressed as $B_n = \frac{|f_n|}{|y_n|}$. The final score was calculated as $Score = \sum_{n=1}^N -(A_n + 0.15 * B_n)$. The factor 0.15 balanced the firing of correct label memories and excessive retrieval, determined through extensive experimentation. A larger value prioritizes precision, while a smaller value emphasizes recall. The negative sum avoided division by zero issues. Hyperparameters were optimized to maximize this score.

C Hyperparameter Settings

All hyperparameters are listed in Table 4.

C.1 SynapticRAG

In our model, in addition to \cos_{th} , V_{th} , stim_{th} , τ_{init} , V_{rest} , I_{rest} mentioned in the text, we include scaling constants τ_{scale} for adjusting τ , t_{scale} for adjusting t and P_{scale} for adjusting P_{score} as targets for hyperparameter optimization.

Using these scaling constants, τ , t , and P_{score} are rewritten as follows:

$$\begin{aligned}\tau &\rightarrow \tau_{scale} \tau \\ t &\rightarrow t_{scale} t \\ P_{score} &\rightarrow P_{scale} P_{score}\end{aligned}$$

In summary, the following nine parameters are subject to hyperparameter optimization in our model: \cos_{th} , V_{th} , stim_{th} , τ_{init} , V_{rest} , I_{rest} , τ_{scale} , t_{scale} , P_{scale} .

C.2 MemoryBank (Adapted)

The original MemoryBank model calculates memory scores as follows:

First, for each memory:

$$\text{Score}^{(i)} = \exp\left(\frac{\Delta t^{(i)}}{5s^{(i)}}\right)$$

where $\Delta t^{(i)}$ is the time difference between when the memory occurred and the current time, and $s^{(i)}$ is the memory strength, with an initial value of 1.

Subsequently, the Score is cut off by a randomly set threshold θ , and memories below this are considered forgotten. We consider the set of memories that exceed the threshold:

$$\text{Retain} = \{i \mid \text{Score}^{(i)} \geq \theta\}$$

Based on the cosine similarity with the input vector, the top 6 memories are selected from this set:

$$\text{Top}_6 = \text{argsort}_{i \in \text{Retain}}(\cos_{\text{sim}}(i))[:6]$$

The strength s of the selected memories is updated by adding 1:

$$s_0^{(i)} = 1, \quad s_n^{(i)} = s_{n-1}^{(i)} + 1 \quad \text{for } i \in \text{Top}_6$$

The memory strength is updated in this way at each turn. In the Adapted model, we introduced five parameters to this model: Top_k , Θ , t_{scale} , s_{scale} , s_{init} .

The equations are as follows:

$$\text{Score}^{(i)} = \exp\left(\frac{t_{scale} \Delta t^{(i)}}{s_{scale} s_i}\right)$$

$$\text{Retain} = \{i \mid \text{Score}^{(i)} \geq \Theta\}$$

$$\text{Top}_k = \text{argsort}_{i \in \text{Retain}}(\cos_{\text{sim}}(i))[:k]$$

$$s_0^{(i)} = s_{init}, \quad s_n^{(i)} = s_{n-1}^{(i)} + 1 \quad \text{for } i \in \text{Top}_k$$

The optimization was performed based on these set parameters.

Model	Parameter	Value
MemoryBank (Adt)	Top _k	7
	Θ	0.12
	t _{scale}	13.82
	s _{scale}	9.06
	s _{init}	0.05
MyAgent	cos _{th}	0.30
MyAgent (Adt)	cos _{th}	0.49
	r _{scale}	10.07
	t _{scale}	-1.72
	g _{scale}	-2.05
SynapticRAG	cos _{th}	0.26
	V _{th}	0.01
	stim _{th}	0.04
	τ _{init}	43.07
	τ _{scale}	2.70
	t _{scale}	7.93
	P _{scale}	0.46
	V _{rest}	2.90
	I _{rest}	-7.13

Table 4: Hyperparameter Settings

C.3 MyAgent (Adapted)

The original MyAgent model is defined as follows: First, the recall probability p is evaluated for memories exceeding a cosine similarity threshold with the input vector:

$$\text{Retain} = \{i \mid \text{cos}_{\text{sim}}(i) \geq \text{cos}_{\text{th}}\}$$

p is calculated as:

$$p^{(i)}(t) = \frac{1 - \exp\left(-\frac{r^{(i)}e^{-\Delta t^{(i)}}}{g^{(i)}}\right)}{1 - e^{-1}}$$

for $i \in \text{Retain}$

where Δt is the time difference between memory occurrence and current time, and g is memory strength, initially set to 1. The memory with the highest p is retrieved, and its strength g is updated:

$$j = \text{argmax}_{i \in \text{Retain}}(p^{(i)}(t))$$

$$g_0^{(j)} = 1, \quad g_n^{(j)} = g_{n-1}^{(j)} + \frac{1 - e^{-\Delta t^{(j)}}}{1 + e^{-\Delta t^{(j)}}}$$

The memory strength is thus updated in each turn. The Adapted model introduces four parameters: cos_{th} , r_{scale} , t_{scale} , g_{scale} . cos_{th} is the cosine similarity threshold from the original model.

The equations for other parameters are as follows:

$$\text{Retain} = \{i \mid \text{cos}_{\text{sim}}(i) \geq \text{cos}_{\text{th}}\}$$

$$p^{(i)}(t) = \frac{1 - \exp(-r_{\text{scale}}r^{(i)}e^{-t_{\text{scale}}\Delta t^{(i)}}/g_{\text{scale}}g^{(i)})}{1 - e^{-1}}$$

$$j = \text{argmax}_{i \in \text{Retain}}(p^{(i)}(t))$$

$$g_0^{(j)} = 1, \quad g_n^{(j)} = g_{n-1}^{(j)} + \frac{1 - e^{-t_{\text{scale}}\Delta t^{(j)}}}{1 + e^{-t_{\text{scale}}\Delta t^{(j)}}}$$

Optimization was performed based on these parameters.

D Computational Complexity Optimization

As we noted in section Limitations, the original SynapticRAG architecture indeed incurs computational complexity on the order of $O(n^2)$ in the worst case. The main reason for this quadratic complexity arises from the cosine-similarity filtering process used during memory retrieval. Specifically, in the original method, each propagation step involves selecting parent nodes (whose number we denote as $(1/a)n$), proportional to the total number of memory nodes n). Each of these parent nodes independently performs a full scan through the entire memory database, resulting in a total computational complexity of approximately $(1/a)n \times n = (1/a)n^2$.

To directly address this issue, we have developed and experimentally validated an optimized alternative architecture. Rather than allowing each parent node to individually scan the entire database, our optimized method computes a **single representative centroid vector** based on **all parent nodes**. This centroid vector is then used to perform only one cosine-similarity scan through the entire dataset.

The computational complexity of this optimized approach can be explained as follows: calculating the centroid vector involves processing $(1/a)n$ parent nodes, thus incurring a computational cost proportional to $(1/a)n$. Subsequently, performing a single cosine-similarity scan using this centroid vector across all memory nodes incurs an additional cost of n . Hence, the total computational complexity of the optimized architecture becomes $(1/a)n + n = (1 + 1/a)n$, effectively **reducing the complexity from $O(n^2)$ to $O(n)$** .

We further analyzed the execution time differences on the SMRCs(EN) dataset. The execution time reported here represents average CPU processing time per retrieval task. Slight variations observed in baseline methods (except SynapticRAG) are mainly attributable to minor differences in the

Table 5: Retrieval Accuracy Comparison (ERC & ERC-MG scores)

Model	SMRCs(EN)		SMRCs(JA)		PerLTQA(EN)		PerLTQA(CN)		Avg.
	ERC	ERC-MG	ERC	ERC-MG	ERC	ERC-MG	ERC	ERC-MG	
RAG	71.55	71.55	81.58	81.58	82.14	82.51	84.95	85.36	80.15
MemoryBank (Adt)	62.36	62.58	72.15	72.15	67.24	67.54	71.31	71.61	68.37
MyAgent (Adt)	63.24	63.46	77.19	77.19	66.49	66.64	69.54	69.76	69.19
SynapticRAG (Original)	86.21	86.21	92.32	92.32	90.47	90.47	93.12	93.12	90.53

Table 6: Optimized Architecture Results

Model	SMRCs(EN)		SMRCs(JA)		PerLTQA(EN)		PerLTQA(CN)		Avg.
	ERC	ERC-MG	ERC	ERC-MG	ERC	ERC-MG	ERC	ERC-MG	
RAG	66.30	66.30	69.52	69.52	73.77	74.15	74.68	75.08	71.07
MemoryBank (Adt)	57.77	58.21	62.50	62.50	64.13	64.43	66.91	67.21	62.96
MyAgent (Adt)	55.14	55.14	58.11	58.11	53.62	53.77	54.05	54.27	55.28
SynapticRAG (Optimized)	79.43	79.43	83.11	83.11	85.74	85.74	86.95	86.95	83.81

Table 7: Execution Time Comparison

Model	Original (s)	Optimized (s)
RAG	0.00147	0.00137
MemoryBank (Adt)	0.08099	0.07916
MyAgent (Adt)	0.04962	0.04948
SynapticRAG	0.14534	0.10801

number of retrieved memories and computational overhead, but these variations remain within a negligible margin of error. Moreover, the optimized architecture achieves approximately a $1.3 \times$ fold-increase in execution speed compared to the original SynapticRAG. Given the difference in computational complexity (from $O(n^2)$ to $O(n)$), we expect this performance improvement to become significantly more pronounced as the dataset size increases.