

Towards Lifelong Dialogue Agents via Timeline-based Memory Management

Kai Tzu-iunn Ong^{1*} Namyong Kim^{1*} Minju Gwak¹ Hyungjoo Chae¹
Taeyoon Kwon¹ Yohan Jo² Seung-won Hwang² Dongha Lee¹ Jinyoung Yeo¹

¹Yonsei University, ²Seoul National University
{ktio89, namyong.kim, jinyeo}@yonsei.ac.kr

Abstract

To achieve lifelong human-agent interaction, dialogue agents need to constantly memorize perceived information and properly retrieve it for response generation (RG). While prior studies focus on getting rid of outdated memories to improve retrieval quality, we argue that such memories provide rich, important contextual cues for RG (e.g., changes in user behaviors) in long-term conversations. We present **THEANINE**, a framework for LLM-based lifelong dialogue agents. THEANINE discards memory removal and manages large-scale memories by linking them based on their temporal and cause-effect relation. Enabled by this linking structure, THEANINE augments RG with memory timelines - series of memories representing the evolution or causality of relevant past events. Along with THEANINE, we introduce **TeaFarm**, a counterfactual-driven evaluation scheme, addressing the limitation of G-Eval and human efforts when assessing agent performance in integrating past memories into RG. A supplementary video for THEANINE and data for **TeaFarm** are at <https://huggingface.co/spaces/ResearcherScholar/Theanine>.

1 Introduction

Autonomous agents based on large language models (LLMs) have made significant progress in various domains, including response generation (Chae et al., 2024; Kwon et al., 2024; Tseng et al., 2024), where agents ought to constantly keep track of both old and newly introduced information shared with users throughout their service lives (Irfan et al., 2024) and converse accordingly. To facilitate such lifelong interaction, studies have proposed enhancing dialogue agents’ ability to memorize and accurately recall past information (e.g., discussed topics) in long-term, multi-session conversations.

A representative approach is to compress past conversations into summarized memories and re-

*KT Ong and N Kim are the co-first authors.

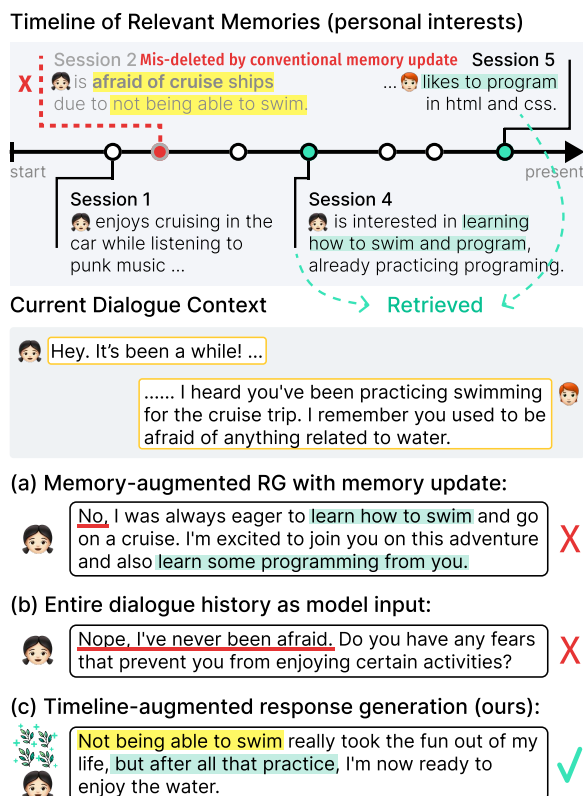


Figure 1: Empirical examples of failed responses due to (a) absence of an important past event (“*afraid of cruise ships*”) on the timeline and (b) bias to the latest input. (c) is a response augmented with the memory timeline.

trieve them to augment response generation (RG) in later encounters (Xu et al., 2022a; Lu et al., 2023). However, the growing span of memories can hinder retrieval quality as conversations accumulate. Although it, to some extent, can be solved by updating old memories (Bae et al., 2022; Zhong et al., 2024), such common practice may cause severe information loss. As shown in Figure 1 (a), an earlier memory on the timeline, an important persona (“*afraid of ships*”), is removed during memory update, resulting in improper RG. While using the large context windows of recent LLMs to process all dialogue history/memories is an op-

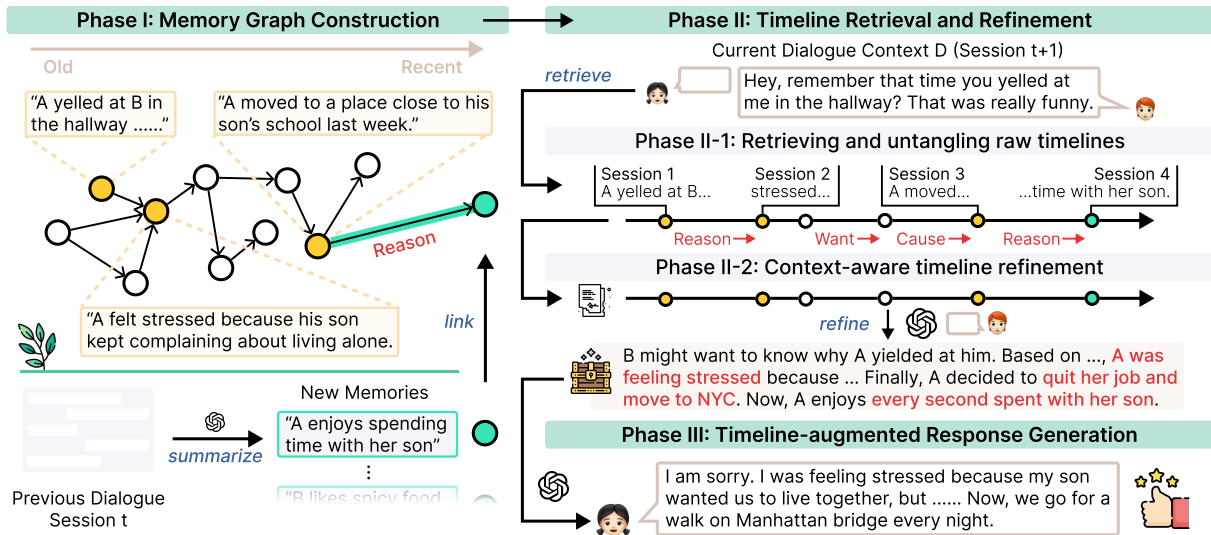


Figure 2: The overview of THEANINE. Left: Linking new memories to the memory graph after finishing a dialogue session; Right: Memory timeline retrieval, refinement, and response generation in a new dialogue session.

tion to prevent such information loss,¹ this often leads to biased attention toward the latest user input (Figure 1 (b)), ignoring relevant contexts from the past (Liu et al., 2024). These findings highlight two main challenges towards lifelong dialogue agents - (i) **Memory construction**: *how to store large-scale past interactions effectively without removing old memories?* (ii) **Response generation**: *within the growing memory span, how to identify relevant contextual cues for generating proper responses?*

Motivated by these, we propose addressing the above two challenges separately yet complementarily, by (i) discarding memory update to avoid information loss, and preserving relevant memories on the timeline in a linked structure; and (ii) retrieving the timeline as a whole to better catch relevant memories within the growing search span. We present THEANINE,² a framework for facilitating lifelong dialogue agents.

Starting from memory construction (**Phase I**), instead of stacking raw memory sentences as-is (Xu et al., 2022a), which may affect memory retrieval and also response quality due to the unstructured format of information (Mousavi et al., 2023; Chen et al., 2023), THEANINE stores memories in a directed graph. In this graph, inspired by how humans naturally link new memories to existing ones of relevant events based on their relation (Bartlett, 1995), memories are linked using their temporal


¹For instance, GPT-4o and Llama 3.1 have context windows of 128K tokens (OpenAI, 2024a; MetaAI, 2024).

²L-theanine is an amino acid found in green tea that has been linked to memory improvement (Nguyen et al., 2019).

and cause-effect commonsense relations (Hwang et al., 2021). Supported by such linking structure, in memory retrieval for RG (**Phase II-1**), we go beyond conventional top- k retrieval and further obtain the complete timelines to avoid missing out on important memories that have low textual overlap with current conversation (Tao et al., 2023). Lastly, to tackle the discrepancy between off-line memory construction and online deployment, THEANINE uses an LLM to refine retrieved timelines (**Phase II-2**) based on current conversation, such that they provide tailored information (Chae et al., 2023) for RG (**Phase III**). Our contributions are two-fold:

- To achieve lifelong dialogue agents, we present THEANINE, an LLM-based framework with a relation-aware memory graph and timeline augmentation for long-term conversations. THEANINE outperforms representative baselines across automatic, LLM-based, and human evaluations of RG. Also, we confirm that THEANINE leads to higher retrieval quality, and its procedures align with human preference. To our knowledge, we are the first to model the use of timelines (*i.e.*, linked relevant memories) in memory management and response generation.
- The lack of golden mapping between conversations and reference memories poses a challenge in assessing memory-augmented agents. We present **TeaFarm**, a counterfactual-driven pipeline evaluating agent performance in referencing the past without human intervention.

2 Methodologies

We present  THEANINE, a framework for lifelong dialogue agents inspired by how humans store and retrieve memories for conversations (Figure 2):

2.1 Memory Graph Construction (Phase I)

To manage large-scale memories and facilitate structured information for RG (Mousavi et al., 2023; Chen et al., 2023), we approach memory management using a memory graph G :

$$G = (V, E) \quad (1)$$


$$V = \{m_1, m_2, \dots, m_{|V|}\} \quad (2)$$

$$m = (\text{event}, \text{time}) \quad (3)$$

$$E = \{(m_i, r_{ij}, m_j) | m_i, m_j \in V \wedge r_{ij} \in R\} \quad (4)$$

$$R = \{\text{Cause}, \text{Reason}, \text{Want}, \dots, \text{SameTopic}\} \quad (5)$$

In G , vertices V are memories m summarized from the conversations. Each memory $m = (\text{event}, \text{time})$ consists of an event³ and the time it is formed (summarized). Each directed edge $e \in E$ between two connected m indicates their temporal order and their cause-effect commonsense relation $r \in R$:

At the end of dialogue session t ,  THEANINE starts linking each new memory m_{new} summarized from session t to the memory graph G^t .

Phase I-1: Identifying associative memories for memory linking. Following how humans link new memories to existing ones that are related to a similar event/topic, *i.e.*, the **associative memories**, THEANINE starts by identifying these associative memories from the memory graph G^t .

Formally, given a newly-formed memory m_{new} waiting to be stored, the associative memories M_a of m_{new} is defined as the set of $m_i \in G^t$ having top- j text similarity with m_{new} (*i.e.*, $|M_a| = j$).

Phase I-2: Relation-aware memory linking. Intuitively, we can link m_{new} to $m \in M_a$ using edges that indicate their text similarity and chronological order, we find such simplified connection (*e.g.*, “this happened \rightarrow that similar event occurred”) can yield a context-poor graph that does not help response generation much (Section 4).

Humans, on the other hand, interpret events by considering the relation between them, such as “how does an event affect the other?” or “why did this person make that change?”. Therefore, we

³In this work, “event” denotes information perceived by the dialogue system, including things done/said by speakers and the acknowledgement of speaker personas.

adopt a relation-aware memory linking, where an edge between two memories is encoded with their cause-effect commonsense relation $r \in R$, along w/ the temporal order. In practice, we adopt the commonly used relations defined by Hwang et al. (2021), including HinderedBy, Cause, Want, and 4 more (Appendix B.1).

We start by determining the relation between m_{new} and each associative memory. Formally, for each pair of m_{new} and $m \in M_a$, the LLM assigns a relation $r \in R$ based on their *event*, *time* and their origin conversations:

$$M_a^* = \{m_i \in M_a | \Upsilon(m_i, m_{new}) \in R\} \quad (6)$$

where $\Upsilon(\cdot, m_{new}) \in R$ indicates that the given memory is assigned with an $r \in R$ with m_{new} ,⁴ and such assigned memories are defined as M_a^* .

We then proceed to link m_{new} to the graph. We first locate every connected component $C_i \subset G^t$ that contains at least one $m \in M_a^*$, as shown in Figure 3 (a) and (b):

$$\mathbb{C} = \{C_i \subset G^t | \mathcal{V}(C_i) \cap M_a^* \neq \emptyset\} \quad (7)$$

where \mathbb{C} is the collection of those C and $\mathcal{V}(\cdot)$ represents “vertices in”. Then, we link m_{new} to the most recent⁵ $m \in M_a^*$ in each $C_i \subset \mathbb{C}$ (Figure 3 (c)). Memories M_{linked} that are linked to m_{new} is defined as follows:

$$M_{linked} = \{\Omega(\mathcal{V}(C_i) \cap M_a^*) | C_i \subset \mathbb{C}\} \quad (8)$$

where $\Omega(\cdot)$ indicates “the most recent memory in”.




Figure 3: Locating memories to be linked to m_{new} .

Linking all memories from session t to G^t , we then obtain a new memory graph G^{t+1} . The pseudo algorithm for Phase I is in **Algorithm 1**.

⁴Limited by retrievers, an $m \in M_a$ may not have a relation with m_{new} . We thus allow the LLM to output “None”.

⁵Simply linking m_{new} to all $m \in M_a^*$ costing 25% more API cost for linking without leading to better response.

2.2 Timeline Retrieval and Timeline Refinement (Phase II)

Thanks to the constructed memory graph, THEANINE can proceed to augment RG with timelines of relevant events, addressing the information loss in conventional memory management (Figure 1). With G^{t+1} , THEANINE performs the following steps for RG in session $t + 1$:

Preparation: Top- k memory retrieval. During the conversation, using the current dialogue context $\mathcal{D} = \{u_i\}_{i=1}^n$ of n utterances u as query, we retrieve top- k memories $M_{re} = \{m_{re1}, \dots, m_{rek}\}$.

Phase II-1: Retrieving and untangling raw memory timelines. We wish to also access memories centered around M_{re} . Formally, given $m_{re} \in M_{re}$, we further collect the connected component $C_{re} \subset G^{t+1}$ that contains m_{re} via the linked structure.

Since this collection of memories (*i.e.*, C_{re}) can be “*tangled up*” together (*i.e.*, connected in a complex manner) due to the graph structure, we proceed to untangle it into several memory timelines, each representing a series of events about m_{re} that starts out similarly yet branches into slightly different development. For that, we first locate the earliest memory in C_{re} as a starting point m_{start} for all timelines, as shown in Figure 4 (left).

$$m_{start} = \Theta(\mathcal{V}(C_{re})) \quad (9)$$

where Θ indicates “the oldest memory in”.




Figure 4: Extracting raw memory timelines τ from the retrieved connected component C_{re} .

Next, starting from m_{start} , we untangle memories by tracing through future direction and extract every possible linear graph containing m_{re} (two in Figure 4) from C_{re} , until reaching an endpoint $\tau[-1]$ with an out-degree of 0 (*i.e.*, $deg^+(\tau[-1]) = 0$), which means no directed edge goes out from it). Each of them is considered a raw memory timeline τ , demonstrating a version of the evolution of m_{re} and its relevant events:

$$\begin{aligned} \mathcal{T} = \{ \tau \subset C_{re} \mid \tau \text{ is a directed linear} \\ \text{graph s.t. } m_{start}, m_{re} \in \tau \quad (10) \\ \wedge deg^+(\tau[-1]) = 0 \} \end{aligned}$$

We then sample n raw timelines τ from \mathcal{T} .⁶ Repeating⁷ Phase II-1 for all retrieved top- k memories, we collect a set of retrieved raw memory timelines $\mathbb{T} = \cup \mathcal{T}$, where $|\mathbb{T}| = k^*n$.

Phase II-2: Context-aware timeline refinement. Although we have constructed the memory graph using temporal and commonsense relations to improve its informativeness, directly applying retrieved timelines for RG can be suboptimal (RQ3, Section 4), because graph construction does not take current conversation into consideration, *i.e.*, they are constructed *off-line*.

In this phase, THEANINE tackles such a discrepancy between off-line memory construction and online deployment (*i.e.*, ongoing conversation) via a context-aware timeline refinement. Motivated by how LLMs can refine their previous generation (Madaan et al., 2024). We leverage LLMs to refine raw timelines into a rich resource of information crafted for the current conversation, by removing redundant information or highlighting information that can come in handy. Formally, given the current dialogue \mathcal{D} and retrieved raw timelines \mathbb{T} , an LLM tailors $\tau \in \mathbb{T}$ into refined timelines \mathbb{T}_Φ :

$$\mathbb{T}_\Phi = \{ \underset{\tau_\Phi}{\operatorname{argmax}} P_{\text{LLM}}(\tau_\Phi | \mathcal{D}, \tau) \mid \tau \in \mathbb{T} \} \quad (11)$$

All refined timelines \mathbb{T}_Φ are then used to augment the response generation. We provide the pseudo algorithm for Phase II in [Algorithm 2](#).

2.3 Timeline-augmented Response Generation (Phase III)

Now, THEANINE utilizes the refined timelines for RG. Formally, given $\mathcal{D} = \{u_i\}_{i=1}^n$ and \mathbb{T}_Φ , an LLM generates a next response \bar{u}_{t+1} :

$$\bar{u}_{n+1} = \underset{u_{n+1}}{\operatorname{argmax}} P_{\text{LLM}}(u_{n+1} | \mathcal{D}, \mathbb{T}_\Phi) \quad (12)$$

3 Experimental Setups

3.1 Datasets of Long-term Conversations

There are limited datasets for long-term, multi-session conversations. Firstly, Multi-Session Chat (MSC) (Xu et al., 2022a), is built upon Persona-Chat (Zhang et al., 2018) by extending its conversations to multiple (five) sessions. Soon after MSC,

⁶We empirically set n to 1, as we observe a high degree of overlap across timelines extracted from the same C_{re} , which can lead to redundant information (*i.e.*, input tokens) for RG.

⁷“Repeating” is used to explain the algorithm from the perspective of one m_{re} . In practice, M_{re} are processed together, although processing them 1-by-1 yields the same result.


Datasets:	Multi-session Chat (MSC)				Conversation Chronicles (CC)			
Methods / Metrics	Bleu-4	Rouge-L	Mauve	BertScore	Bleu-4	Rouge-L	Mauve	BertScore
All Dialogue History	1.65	14.89	9.06	86.28	4.90	<u>21.56</u>	26.47	88.13
All Memories & Current Context \mathcal{D}	1.56	14.89	10.62	86.23	4.41	20.06	<u>38.16</u>	88.02
+ Memory Update (Bae et al., 2022)	1.55	14.77	9.28	86.20	4.34	20.34	34.84	88.03
Memory Retrieval (Xu et al., 2022a)	1.92	15.49	11.16	<u>86.47</u>	<u>4.93</u>	20.63	33.06	<u>88.07</u>
+ Memory Update (Bae et al., 2022)	1.67	15.30	<u>13.71</u>	86.39	4.46	20.19	34.28	88.02
Rsum-LLM (Wang et al., 2023)	0.75	11.53	2.45	84.91	0.98	11.42	2.28	85.59
MemoChat (Lu et al., 2023)	1.42	13.51	7.72	85.96	2.31	15.87	15.12	87.08
COMEDY (Chen et al., 2024b)	1.06	12.79	7.27	85.29	1.70	13.57	1.95	85.90
 THEANINE (Ours)	<u>1.80</u>	<u>15.37</u>	18.62	86.70	6.85	22.68	64.41	88.58

Table 1: Automatic evaluation of response quality (average of sessions).

DuLeMon (Xu et al., 2022b) and CareCall (Bae et al., 2022) are proposed for long-term conversations in Mandarin and Korean. Recently, Jang et al. (2023) release a new dataset, Conversation Chronicles (CC). Unlike MSC, CC augments speakers with defined relationships, such as “*employee and boss*”. Apart from these open-domain datasets, the Psychological QA,⁸ addresses long-term conversations under clinical scenarios in Mandarin.

We opt for MSC and CC for evaluation to focus on English conversations, leaving multilingual and domain-specific conversations (*e.g.*, DuleMon, CareCall, and Psychological QA) to future work.

3.2 Baselines

To evaluate THEANINE, in addition to naive baselines that utilize all past dialogues or memories, we incorporate the following settings:

Memory Retrieval. Following Xu et al. (2022a), we use a retriever to retrieve memories relevant to the current dialogue context to augment RG.

Memory Update. We utilize LLMs to implement a widely used updating algorithm proposed by Bae et al. (2022) at the end of each dialogue session. This algorithm includes functionalities such as Change, Replace, Delete, Append, and more (see Appendix H).

RSum-LLM. An LLM-only generative method that recursively summarizes and updates the memory pool, generating responses w/o a retrieval module (Wang et al., 2023).

MemoChat. Proposed by Lu et al. (2023), it leverages LLMs’ CoT reasoning ability to (i) conclude important memories from past conversations in a structured topic-summary-dialogue manner, (ii) select memories, and (ii) generate responses.

COMEDY. Proposed by Chen et al. (2024b), it uses LLMs to summarize session-level memories,

compresses all of them into short events, user portraits (behavioral patterns, emotion, etc.) and user-bot relation. It then selects compressed memories to augment response generation.


3.3 Models and Implementation Details

Large language models. In all experiments, including baselines, we adopt gpt-3.5-turbo-0125 (OpenAI, 2023) for (i) memory summarization (Table 6), (ii) memory update, and (iii) response generation. Temperature is set to 0.75.

Retrievers. We use text-embedding-3-small (OpenAI, 2024b) to calculate text similarity for settings involving retrievers. In the identification of top- j associative memories (Phase I-1) and top- k memory retrieval (Phase II), we set j and k to 3. For the “Memory Retrieval” baseline, we set $k = 6$ following Xu et al. (2022a).

Dialogue sessions. We use sessions 3-5 of MSC and CC for evaluations, as all methods are almost identical in session 1 \sim 2 (no memory to update).

4 Evaluation Scheme 1: Automatic and Human Evaluations

To evaluate  THEANINE’s responses in long-term conversations, we follow common practices and conduct 3 types of evaluations: (i) Automatic evaluations; (ii) G-Eval (Liu et al., 2023), an LLM-based framework commonly used to evaluate LMs’ generation; (iii) human evaluation. We now present several key findings (details, prompts, and interfaces of evaluations in Scheme 1 are in Appendix E):

(Finding 1) THEANINE outperforms baselines in response generation. Table 1 presents the agent performance in RG regarding both overlap-based and embedding-based metrics: Bleu-4 (Papineni et al., 2002), Rouge-L (Lin, 2004), Mauve (Pillutla et al., 2021), and BertScore (Zhang et al., 2020).

⁸<https://www.xinli001.com/>


Settings / Metrics	B-4	R-L	Mauve	Bert
 THEANINE (Ours)	4.32	19.03	41.52	87.64
w/o Relation-aware Linking	4.07	18.58	39.69	87.57
w/o Timeline Refinement	4.03	18.82	41.34	87.66
Broken Down, Shuffled Timeline	<u>4.15</u>	18.70	38.49	87.61
Memory Retrieval	3.43	18.06	22.11	87.27

Table 2: Performance of our ablations (avg. of datasets).

Across both datasets, THEANINE, achieves superior response quality than various baselines. Although, compared to Memory Retrieval, THEANINE scores slightly lower in overlap-based metrics (*i.e.*, B-4 and R-L) in MSC, it largely outperforms Memory Retrieval in embedding-based metrics. Interestingly, including ours, methods without memory update generally yield higher scores, justifying our proposal towards an update-, removal-free memory management for lifelong dialogue agents.

(Finding 2 & 3) All phases contribute to performance; retrieving the timeline as a whole brings large improvement over conventional retrieval.

To gain deeper insights into our design, we investigate the impact of removing THEANINE’s relation-awareness during memory linking (Phase I-2) and Timeline Refinement (Phase II-2). Also, to objectively assess whether THEANINE’s retrieval (*i.e.*, retrieving the timeline as a whole) improves retrieval quality, we include a setting where retrieved timelines are broken down into randomly ordered events such that retrieved memories during RG are in the same format as conventional top- k retrieval.

In Table 2, we observe a ranking in terms of contribution to performance: *relation-aware linking* > *retrieving timeline as a whole* > *timeline refinement*. This observation confirms the efficacy of constructing a memory graph with causal relations. Moreover, utilizing this graph structure to collect timelines of relevant events yields higher RG quality than conventional retrieval, despite the smaller k (3 vs. 6) in initial retrieval. Refining timelines shows smaller performance gains, suggesting room for improvement in applying them for RG. We leave it to future work.

(Finding 4) Humans and G-Eval reveal that THEANINE leads to higher retrieval quality regarding both helpfulness and accuracy. Beyond agent responses, we further investigate how different memory construction methods affect the quality of memory retrieval. Given the same cur-

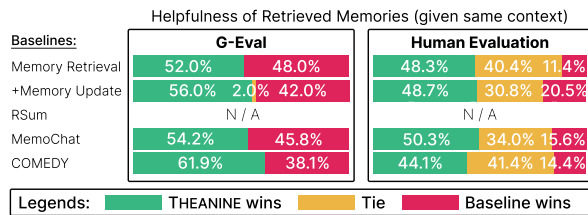


Figure 5: Human- (right) and machine-based (left) head-to-head comparisons between ours and baselines regarding the *helpfulness* of retrieved memories.

rent dialogues as queries for memory retrieval, Figure 5 shows head-to-head comparisons (ours vs. baselines) regarding whose retrieved memories more effectively benefit RG. We observe higher win rates for THEANINE in all comparisons, especially in human evaluations. This suggests that our method can facilitate more helpful memory augmentation for response generation.

In addition to helpfulness, objectively measuring retrieval accuracy is crucial. Since existing datasets of long-term conversations do not provide a golden mapping between dialogue contexts and memories (*i.e.*, golden memories for retrieval), we identify 50 dialogue contexts (*i.e.*, test instances) that require a past memory for RG, and manually measure the retrieval accuracy of different agents. The results shown in Table 3 indicate that THEANINE and its ablations demonstrate higher retrieval accuracy than baselines, and the ranking here aligns with Table 1 and success rates in Table 4.


Methods (Agents)	Golden Memory is Retrieved/collected (%)
Memory Retrieval	68.00
+ Memory Update	64.00
MemoChat	56.00
COMEDY	48.00
 THEANINE (Ours)	72.00

Table 3: Human evaluation of the *accuracy* of memory retrieval (we examine 50 test instances).

(Finding 5) Humans confirm that THEANINE yields responses better entailing past interactions.

Now that the helpfulness of THEANINE’s retrieved memories is validated, we proceed to investigate whether such helpful memories contribute towards reliable lifelong human-agent interaction.

For that, we further ask a group of workers to specifically judge whether agent responses *entail*, *contradict*, or are *neutral* to the past via majority voting. In Figure 6, THEANINE not only leads to a small number of contradictory responses (4%) but also demonstrates the largest percentage (68%;

out of 100) of responses that entail past conversations, significantly outperforming baselines. We argue that it is because our timeline-based approach elicits memories better at representing past interactions between speakers, thus leading to responses more directly aligned with the past. This alignment is important for dialogue agents to maintain long-term intimacy with users (Adiwardana et al., 2020). Furthermore, such entailing and non-contradictory nature of THEANINE’s responses highlights its potential for applications in specialized domains, such as personalized agents for clinical scenarios, where entailment between agent responses and users’ past information (e.g., electrical health records or previous consulting sessions) is crucial for diagnostic decision-making (Tseng et al., 2024).




Figure 6: Human evaluations regarding to what extent the agent responses entail past conversations.

As a side note, Memory Update yields fewer contradictory responses (2%), indicating a potential trade-off between (i) removing outdated memories to prevent contradiction and (ii) preserving them to get richer information for RG (Kim et al., 2024a).

(Finding 6) Humans agree with THEANINE’s intermediate procedures. As reported in Figure 7, judges largely agree (92%) that THEANINE properly assigns cause-effect relations to linked memories, which explains its contribution to performance. Also, they agree that timeline refinement successfully elicits more helpful information (100%; 100 samples in total) for RG. Examples of THEANINE’s phases and RG are in Appendix G.




Figure 7: Human evaluation of our intermediate phases.

5 Evaluation Scheme 2: TeaFarm – a Counterfactual-driven Evaluation Pipeline for Long-term Conversations

Evaluating memory-augmented agents in long-term conversations is non-trivial due to the unavailability of ground-truth mapping between current con-

versations and correct memories for retrieval. Although we may resort to G-Eval by feeding evaluator LLMs (e.g., GPT-4) the entire past history and prompt it to determine whether a response correctly recalls the past, the evaluation can be largely limited by the performance of the evaluator LLM itself (Kim et al., 2024b).

To overcome this, along with THEANINE, we present **TeaFarm**, a human-free counterfactual-driven pipeline for evaluating memory-augmented response generation in long-term conversations.

5.1 Testing Dialogue Agents’ Memory via Counterfactual Questions

In **TeaFarm**, we proceed to “trick” dialogue agents into generating incorrect responses, and agents must correctly reference past conversations to avoid being misled by us. Specifically, we talk to the dialogue agent while *acting as if a non-factual statement is true* (thus counterfactual). Figure 8 presents some examples of counterfactual questions and the corresponding facts.

Facts (at this moment)	Counterfactual Questions
Speaker B has never been to Japan.	A: Hey, did you have a great time in Tokyo?
Speaker A bought a new house in NYC three months ago.	B: So you are still hesitating to buy that house in NYC you’ve been talking about, right?
Speaker B does not own a car.	B: Hey, do you remember when we sang karaoke in my car?

Figure 8: Examples of counterfactual questions.

In practice (Figure 11), when we want to evaluate an agent that has been interacting with the user for sessions, we first (1) collect all past conversations and summarize them session by session. Then, we (2) feed a question generator LLM⁹ the collected summaries in chronological order such that it can capture the current stage of each discussed event, e.g., “*Speaker B does not own a car*”, and (3) generate counterfactual questions from the perspective of both speakers (and the correct answers). After that, we (4) kick off (i.e., simulate) a new dialogue session, chat for a while, then (5) naturally ask the counterfactual question, and (6) assess the correctness of its response. The overview figure, prompts, and synthesized data for **TeaFarm** are in Appendix C, H, and D, respectively.

Settings / Datasets	MSC	CC	Avg.
Memory Retrieval	0.16	0.19	0.18
+ Memory Update	0.16	0.19	0.18
RSum-LLM*	0.04	0.08	0.06
MemoChat*	0.09	0.15	0.12
COMEDY*	0.06	0.18	0.12
THEANINE	0.18	0.24	0.21
w/o Relation-aware Linking	0.17	0.20	0.19
w/o Timeline Refinement	0.16	0.19	0.18

Table 4: Success rates (SRs) of correctly recalling the past and not being fooled by the counterfactual questions in TeaFarm (tested with 200 questions).

5.2 TeaFarm Results

In Table 4, THEANINE shows higher SR than baselines, especially in CC. Ablations perform slightly worse than the original, again proving the efficacy of relation-aware linking and timeline refinement. Surprisingly, all settings have low SRs, qualifying TeaFarm as a proper pipeline for stress-testing dialogue agents in long-term conversations.

Interestingly, baselines using retrievers (same as THEANINE) show superior performance than settings only relying on LLMs (*i.e.*, RSum-LLM, MemoChat, and COMEDY). This, unexpectedly, supports our efforts in developing a new paradigm of memory management in the era of LLMs.¹⁰

To provide insight regarding conversation scenarios that are challenging for dialogue agents, we present case studies of how THEANINE fail in TeaFarm in Appendix G.

6 Further Analyses and Discussions

Cost efficiency. A concern of THEANINE is the API cost. Regardless, we argue that it is competitive when both performance and cost are taken into account. Figure 9 plots response quality (Mauve score) against the API cost.¹¹ We find THEANINE and all ablations not only outperform all baselines but also lie on the Pareto frontier, indicating an efficient cost-performance trade-off. This suggests THEANINE’s value when performance is prioritized over API costs. Actual API costs and results based on B-4, R-L, and Bert scores are in Appendix I.

⁹We apply GPT-4 (gpt-4) with a temperature of 0.75.

¹⁰Memory update does not affect Memory Retrieval’s performance. We believe it is because counterfactual questions are made to counter the newest stage of each event. The removal of older memories thus does not have much impact.

¹¹Calculated based on session 5, which involves most memories for management. We use Mauve for its stronger correlation with humans (Pillutla et al., 2021).




Figure 9: Cost-performance comparisons.

Time efficiency. Time efficiency can be an important consideration when deploying THEANINE to real-world scenarios having richer events. Figure 10 shows time-performance comparisons regarding both “memory construction” and “retrieval + RG” also using the Pareto frontier. Similarly, THEANINE and many of its ablations demonstrate an efficient time-performance trade-off.

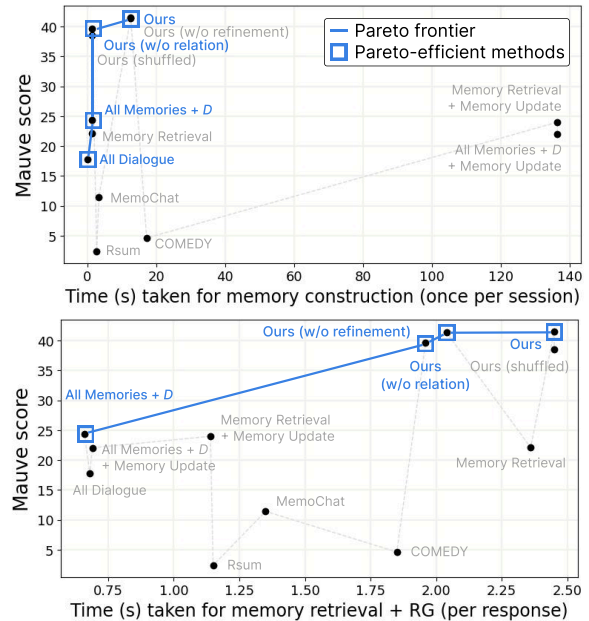


Figure 10: Time-performance comparisons.

Additional comparison: Memory Retrieval with a dynamically-changing k . Due to THEANINE’s graph-based procedures, the response generator may access different amounts of memories during RG depending on given contexts (*i.e.*, queries used by the retriever) and when the conversation takes place (*i.e.*, an earlier or a later session), whereas conventional methods (Xu et al., 2022a; Bae et al., 2022) often have a fixed number k of memories retrieved for RG. Therefore, to further quantify the effect of our proposed timeline-based manage-


Methods / Metrics	Bleu-4	Rouge-L	Mauve	Bert
Memory Retrieval (<i>dynamic k</i>)	3.06	17.97	33.33	87.32
+ Memory Update	2.68	17.19	28.49	87.11
 THEANINE (Ours)	4.22	19.22	45.53	87.70

Table 5: Additional comparison, where k in Memory Retrieval is dynamically modified for each test instance.

ment and augmentation, we compare THEANINE to Memory Retrieval with a *dynamic k*, where k dynamically changes based on the number of collected memories in THEANINE for each specific test data. In other words, if THEANINE uses timelines to collect k memories during RG for a test instance \mathcal{D}_i , baselines will also be retrieving k memories for generating a response for \mathcal{D}_i .

In Table 5, we can observe that when the number of memories is matched, ours outperforms both baselines despite the same amount of memories being provided. We assume this is because: (i) our graph-based retrieval helps us collect more beneficial memories than conventional retrieval; (ii) addressing the relation between events and shaping them based on dialogue contexts can facilitate richer contextual cues for RG.

Growing span of memories. Another inquiry is whether the growing span of memory will eventually hinder retrieval in THEANINE if there ever are hundreds of sessions. Although this may be a serious issue for conventional methods, we presume that it will be partially mitigated in THEANINE, as: (i) We retrieve relevant memories *as a whole* in the form of timelines. This serves as a safety net in scenarios where an important memory is missed out in top- k retrieval—it may be collected via the linked structure; (ii) We refine retrieved timelines based on current dialogue such that they provide tailored information for RG. This acts as a second insurance against sub-optimal retrieval.



7 Related Work

Long-term conversations. Since MSC, there have been several studies on long-term conversations: Bae et al. (2022) train a classifier to update old memories in phone call scenarios. As we enter the era of LLMs, Li et al. (2024) leverages LLMs to write and update memories for RG. Apart from LLMs’ power, human behaviors also foster methods in this field. For example, Zhong et al. (2024) apply humans’ forgetting curve to make memories that have been discussed exist longer. Recently, Park et al. (2023) and Maharana et al. (2024) also

adopt the concept of timelines. However, Park et al. (2023) focus on tagging the timestamp (e.g., “22:00”) of events and does not explicitly model the connection between them, and, in Maharana et al. (2024), a timeline is a fixed, pre-defined series of events (potentially unrelated) which simply serve as a user profile for synthesizing dialogue data. By contrast, in our work, a timeline is built with relevant events, which are dynamically linked based on their causal relations and retrieved as the conversation goes on, benefitting our goal of consistent memory tracking and integration.

Memory-augmentation for personalized dialogue agents. The trend of long-term interaction with autonomous agents promotes their adaptation for personalized needs (Chen et al., 2024a,c). As a pioneer, Xu et al. (2022b) train a persona extractor to create user-based memories. However, training personalized agents for long-term use can be non-trivial due to the lack of data (Tseng et al., 2024). As a solution, Kim et al. (2024a) apply commonsense models and LLMs to augment existing long-term data with high-quality persona sentences; Chen et al. (2024b) present a training-free LLM-based framework that extracts user behaviors from conversations for personalized RG. Upon the success of LLMs, THEANINE leverages them to build memory timelines. These timelines represent the development of interactions and lead to responses that better entail speaker information, establishing THEANINE’s potential for personalized agents.

8 Conclusions

This paper presents the first-ever timeline-based memory management and augmentation framework,  THEANINE, for autonomous agents in long-term conversations. Applying THEANINE, we develop a dialogue agent that efficiently addresses the constant, lifelong tracking of memories and their integration for response generation throughout its service life. Comprehensive evaluations show that  THEANINE can facilitate more beneficial memory augmentation, leading to responses that are closer to ground truths and more aligned with speakers’ past interactions. THEANINE’s effectiveness is further confirmed in TeaFarm, a counterfactual-driven pipeline we design to address the limitation of G-Eval and human efforts in assessing memory augmentation. We expect our novel approaches to serve as a new foundation for future efforts towards lifelong dialogue agents.

Limitations

First, the amount of dialogue sessions in this study is limited to five due to the lack of longer open-domain English datasets. As we mentioned in Section 6, we presume that THEANINE’s effectiveness can still hold true to some degree in longer conversations. Yet, we do acknowledge the need to apply additional modules that directly address the growing span of dialogue history/memories, such as introducing the summarize-then-compress paradigm in COMEDY (Chen et al., 2024b) to compress session-level summaries into a combined short user/event description.

Second, although we include many recent frameworks as baselines, we failed to compare THEANINE with MemoryBank (Zhong et al., 2024), a framework inspired by Ebbinghaus’s forgetting curve. This is because the time intervals between sessions in MSC and CC are either mostly measured in hours or not clearly specified (e.g., “a few months later”), whereas MemoryBank requires precise time intervals in days to apply the forgetting curve. Also, data used for MemoryBank focuses on Chinese clinical scenarios, making it not feasible for our study. However, we remain positive about applying such a mechanism to improve THEANINE in our ongoing research.

Lastly, API-based LLMs may introduce risks such as privacy issues. A possible solution is to apply THEANINE to small open-source LMs for secure, local usage. While there exist challenges in data collection, one may achieve this by (i) collecting synthesized conversations with GPT-generated user profiles, (ii) running THEANINE on these data, and (iii) using the outputs of each phase to train student LMs (i.e., distillation from teacher LLMs).

Ethical Statements

LLMs might generate harmful, biased, offensive, and sexual content. Authors avoid such content from appearing in this paper. We guarantee fair compensation for human evaluators from Amazon Mechanical Turk. We ensure an effective pay rate higher than 20\$ per hour based on the estimated time required to complete the tasks.

Acknowledgments

This work was mainly supported by STEAM R&D Project, NRF, Korea (RS-2024-00454458) and Institute of Information & communications Technology Planning & Evaluation (IITP) grant funded

by the Korean government (MSIT) (No. RS-2024-00457882, National AI Research Lab Project), and was partially supported by the National Research Foundation of Korea (NRF) grant funded by the Korea government (MSIT) (RS-2024-00333484; RS-2024-00414981). Jinyoung Yeo is the corresponding author (jinyeo@yonsei.ac.kr).

References

- Daniel Adiwardana, Minh-Thang Luong, David R So, Jamie Hall, Noah Fiedel, Romal Thoppilan, Zi Yang, Apoorv Kulshreshtha, Gaurav Nemade, Yifeng Lu, et al. 2020. Towards a human-like open-domain chatbot. *arXiv preprint arXiv:2001.09977*.
- Sanghwan Bae, Donghyun Kwak, Soyoung Kang, Min Young Lee, Sungdong Kim, Yui Jeong, Hyeri Kim, Sang-Woo Lee, Woomyoung Park, and Nako Sung. 2022. Keep me updated! memory management in long-term conversations. In *Findings of the Association for Computational Linguistics: EMNLP 2022*, pages 3769–3787.
- Frederic Charles Bartlett. 1995. *Remembering: A study in experimental and social psychology*. Cambridge university press.
- Hyungjoo Chae, Namyoun Kim, Kai Tzu-iunn Ong, Minju Gwak, Gwanwoo Song, Jihoon Kim, Sanghwan Kim, Dongha Lee, and Jinyoung Yeo. 2024. Web agents with world models: Learning and leveraging environment dynamics in web navigation. *arXiv preprint arXiv:2410.13232*.
- Hyungjoo Chae, Yongho Song, Kai Ong, Taeyoon Kwon, Minjin Kim, Youngjae Yu, Dongha Lee, Dongyeop Kang, and Jinyoung Yeo. 2023. Dialogue chain-of-thought distillation for commonsense-aware conversational agents. In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pages 5606–5632.
- Howard Chen, Ramakanth Pasunuru, Jason Weston, and Asli Celikyilmaz. 2023. Walking down the memory maze: Beyond context limit through interactive reading. *arXiv preprint arXiv:2310.05029*.
- Jiangjie Chen, Xintao Wang, Rui Xu, Siyu Yuan, Yikai Zhang, Wei Shi, Jian Xie, Shuang Li, Ruihan Yang, Tinghui Zhu, et al. 2024a. From persona to personalization: A survey on role-playing language agents. *arXiv preprint arXiv:2404.18231*.
- Nuo Chen, Hongguang Li, Juhua Huang, Baoyuan Wang, and Jia Li. 2024b. Compress to impress: Unleashing the potential of compressive memory in real-world long-term conversations. *arXiv preprint arXiv:2402.11975*.
- Yi-Pei Chen, Noriki Nishida, Hideki Nakayama, and Yuji Matsumoto. 2024c. Recent trends in personalized dialogue generation: A review of datasets,

- methodologies, and evaluations. *arXiv preprint arXiv:2405.17974*.
- Jena D Hwang, Chandra Bhagavatula, Ronan Le Bras, Jeff Da, Keisuke Sakaguchi, Antoine Bosselut, and Yejin Choi. 2021. (comet-) atomic 2020: On symbolic and neural commonsense knowledge graphs. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 35, pages 6384–6392.
- Bahar Irfan, Mariacarla Staffa, Andreea Bobu, and Nikhil Churamani. 2024. Lifelong learning and personalization in long-term human-robot interaction (leap-hri): Open-world learning. In *Companion of the 2024 ACM/IEEE International Conference on Human-Robot Interaction*, pages 1323–1325.
- Jihyoung Jang, Minseong Boo, and Hyoungun Kim. 2023. **Conversation chronicles: Towards diverse temporal and relational dynamics in multi-session conversations**. In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pages 13584–13606, Singapore. Association for Computational Linguistics.
- Hana Kim, Kai Ong, Seoyeon Kim, Dongha Lee, and Jinyoung Yeo. 2024a. **Commonsense-augmented memory construction and management in long-term conversations via context-aware persona refinement**. In *Proceedings of the 18th Conference of the European Chapter of the Association for Computational Linguistics (Volume 2: Short Papers)*, pages 104–123, St. Julian's, Malta. Association for Computational Linguistics.
- Seungone Kim, Juyoung Suk, Shayne Longpre, Bill Yuchen Lin, Jamin Shin, Sean Welleck, Graham Neubig, Moontae Lee, Kyungjae Lee, and Minjoon Seo. 2024b. Prometheus 2: An open source language model specialized in evaluating other language models. *arXiv preprint arXiv:2405.01535*.
- Taeyoon Kwon, Kai Tzu-iunn Ong, Dongjin Kang, Seungjun Moon, Jeong Ryong Lee, Dosik Hwang, Beomseok Sohn, Yongsik Sim, Dongha Lee, and Jinyoung Yeo. 2024. Large language models are clinical reasoners: Reasoning-aware diagnosis framework with prompt-generated rationales. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 38, pages 18417–18425.
- Hao Li, Chenghao Yang, An Zhang, Yang Deng, Xiang Wang, and Tat-Seng Chua. 2024. Hello again! Llm-powered personalized agent for long-term dialogue. *arXiv preprint arXiv:2406.05925*.
- Chin-Yew Lin. 2004. **ROUGE: A package for automatic evaluation of summaries**. In *Text Summarization Branches Out*, pages 74–81, Barcelona, Spain. Association for Computational Linguistics.
- Nelson F Liu, Kevin Lin, John Hewitt, Ashwin Paranjape, Michele Bevilacqua, Fabio Petroni, and Percy Liang. 2024. Lost in the middle: How language models use long contexts. *Transactions of the Association for Computational Linguistics*, 12.
- Yang Liu, Dan Iter, Yichong Xu, Shuohang Wang, Ruo Chen Xu, and Chenguang Zhu. 2023. G-eval: Nlg evaluation using gpt-4 with better human alignment. In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pages 2511–2522.
- Junru Lu, Siyu An, Mingbao Lin, Gabriele Pergola, Yulan He, Di Yin, Xing Sun, and Yunsheng Wu. 2023. Memochat: Tuning llms to use memos for consistent long-range open-domain conversation. *arXiv preprint arXiv:2308.08239*.
- Aman Madaan, Niket Tandon, Prakhar Gupta, Skyler Hallinan, Luyu Gao, Sarah Wiegrefe, Uri Alon, Nouha Dziri, Shrimai Prabhumoye, Yiming Yang, et al. 2024. Self-refine: Iterative refinement with self-feedback. *Advances in Neural Information Processing Systems*, 36.
- Adyasha Maharana, Dong-Ho Lee, Sergey Tulyakov, Mohit Bansal, Francesco Barbieri, and Yuwei Fang. 2024. **Evaluating very long-term conversational memory of LLM agents**. In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 13851–13870, Bangkok, Thailand. Association for Computational Linguistics.
- MetaAI. 2024. Llama3. <https://ai.meta.com/blog/meta-llama-3-1/>.
- Seyed Mahed Mousavi, Simone Caldarella, and Giuseppe Riccardi. 2023. Response generation in longitudinal dialogues: Which knowledge representation helps? In *The 5th Workshop on NLP for Conversational AI*, page 1.
- Bao Trong Nguyen, Naveen Sharma, Eun-Joo Shin, Ji Hoon Jeong, Sung Hoon Lee, Choon-Gon Jang, Seung-Yeol Nah, Toshitaka Nabeshima, Yukio Yoneda, and Hyoung-Chun Kim. 2019. Theanine attenuates memory impairments induced by klotho gene depletion in mice. *Food & function*, 10(1):325–332.
- OpenAI. 2023. Chatgpt. <https://openai.com/blog/chatgpt>.
- OpenAI. 2024a. Openai website. <https://openai.com/>.
- OpenAI. 2024b. **Openai's text embeddings**.
- Kishore Papineni, Salim Roukos, Todd Ward, and Wei-Jing Zhu. 2002. **Bleu: a method for automatic evaluation of machine translation**. In *Proceedings of the 40th Annual Meeting of the Association for Computational Linguistics*, pages 311–318, Philadelphia, Pennsylvania, USA. Association for Computational Linguistics.
- Joon Sung Park, Joseph O'Brien, Carrie Jun Cai, Meredith Ringel Morris, Percy Liang, and Michael S Bernstein. 2023. Generative agents: Interactive simulacra

of human behavior. In *Proceedings of the 36th Annual ACM Symposium on User Interface Software and Technology*, pages 1–22.

Krishna Pillutla, Swabha Swayamdipta, Rowan Zellers, John Thickstun, Sean Welleck, Yejin Choi, and Zaid Harchaoui. 2021. Mauve: Measuring the gap between neural text and human text using divergence frontiers. *Advances in Neural Information Processing Systems*, 34:4816–4828.

Chongyang Tao, Jiazhan Feng, Tao Shen, Chang Liu, Juntao Li, Xiubo Geng, and Daxin Jiang. 2023. Core: Cooperative training of retriever-reranker for effective dialogue response selection. In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 3102–3114.

Yu-Min Tseng, Yu-Chao Huang, Teng-Yun Hsiao, Yu-Ching Hsu, Jia-Yin Foo, Chao-Wei Huang, and Yun-Nung Chen. 2024. Two tales of persona in llms: A survey of role-playing and personalization. *arXiv preprint arXiv:2406.01171*.

Qingyue Wang, Liang Ding, Yanan Cao, Zhiliang Tian, Shi Wang, Dacheng Tao, and Li Guo. 2023. Recursively summarizing enables long-term dialogue memory in large language models. *arXiv preprint arXiv:2308.15022*.

Jing Xu, Arthur Szlam, and Jason Weston. 2022a. **Beyond goldfish memory: Long-term open-domain conversation**. In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 5180–5197, Dublin, Ireland. Association for Computational Linguistics.



Xinchao Xu, Zhibin Gou, Wenquan Wu, Zheng-Yu Niu, Hua Wu, Haifeng Wang, and Shihang Wang. 2022b. Long time no see! open-domain conversation with long-term persona memory. In *Findings of the Association for Computational Linguistics: ACL 2022*, pages 2639–2650.

Saizheng Zhang, Emily Dinan, Jack Urbanek, Arthur Szlam, Douwe Kiela, and Jason Weston. 2018. **Personalizing dialogue agents: I have a dog, do you have pets too?** In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 2204–2213, Melbourne, Australia. Association for Computational Linguistics.

Tianyi Zhang, Varsha Kishore, Felix Wu, Kilian Q Weinberger, and Yoav Artzi. 2020. Bertscore: Evaluating text generation with bert. In *International Conference on Learning Representations*.

Wanjun Zhong, Lianghong Guo, Qiqi Gao, He Ye, and Yanlin Wang. 2024. Memorybank: Enhancing large language models with long-term memory. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 38, pages 19724–19731.

A Appendix Contents

- Appendix B.1: Cause-effect Commonsense Relations Adopted.
- Appendix B.2: Algorithms for  THEANINE.
- Appendix B.3: Implementation Details on Computational Experiments.
- Appendix C: **TeaFarm** Evaluation.
- Appendix D: The  TeaBag Dataset.
- Appendix E Details on Evaluation Scheme 1 (G-Eval and Human Evaluations).
- Appendix F: Session-specific Results of Automatic Evaluation.
- Appendix G: Empirical Examples.
- Appendix H: Prompts Used in This Work.
- Appendix I: Further Analyses.
- Appendix J: Terms for Use of Artifacts.

B Further Implementation Details

B.1 Cause-effect Commonsense Relations

We adopt and modify commonsense relations from [Hwang et al. \(2021\)](#) for our relation-aware memory linking. Below is the list of our commonsense relations R :

Changed: Events in A changed to events in B.

Cause: Events in A caused events in B.

Reason: Events in A are due to events in B.

HinderedBy: When events in B can be hindered by events in A, and vice versa.

React: When, as a result of events in A, the subject feels as mentioned in B.

Want: When, as a result of events in A, the subject wants events in B to happen.

SameTopic: When the specific topic addressed in A is also discussed in B.

Limited by the performance of retrievers, it is possible that an $m \in M_a$ does not have a relation, other than just textual overlap, with m_{new} . We address this by allowing the LLM to output **None**.

B.2 Algorithms for THEANINE

The pseudo algorithms for Phase I and II are provided in Algorithm 1 and 2.

B.3 Implementation Details on Computational Experiments

All computational experiments in this work are based on OpenAI API (OpenAI, 2024a). Thus, no computing infrastructure is required in this work.

C TeaFarm Evaluation

The overall pipeline of TeaFarm is illustrated in Figure 11.

D The 🍵TeaBag Dataset

As a byproduct of TeaFarm, we curate 🍵TeaBag, a dataset for TeaFarm evaluation on MSC and CC. TeaBag consists of:

- 100 episodes of original conversations from Multi-Session Chat and Conversational Chronicles (session 1-5; 50 episodes from each dataset)
- Two pairs of counterfactual QAs for each episode (200 pairs in total).
- Two synthesized follow-up conversations (*i.e.*, session 6) for each episode (thus 200 in total), each of which naturally guides the conversation from session 5 towards one of the counterfactual questions.

This dataset is made with GPT-4. The prompt for generation is in Appendix H. We expect future work to apply TeaBag to stress-test if their dialogue system can correctly reference past conversations.

TeaBag does not contain personally identifying information, as it is generated based on datasets where all contents are pure artificial creation, rather than contents collected from the real-world. Also, we have tried our best to confirm that this dataset does not contain any offensive content.

For the overview of data collection, please refer to step 1-4 of TeaFarm (Figure 11).

E Details on Evaluation Scheme 1

We perform evaluations using sessions 3-5 from MSC and CC, as all settings are almost identical before the end of session 2, due to the fact that there is no memory to update before then.

The test sets of MSC and CC contain over 500 and 20,000 episodes of conversations, where each episode has 5 dialogue sessions, yielding 1.2M turns of responses in total. Due to the limited budget for generation (both baselines and ours), when

not specified, we sample 50 episodes from each dataset for experiments in this paper (around 3.6K conversational turns in total).

E.1 G-Eval

G-Eval (Liu et al., 2023) is a framework using LLMs with chain-of-thoughts (CoT) and a form-filling paradigm to assess the quality of models' text generation. G-Eval with GPT-4 has been shown to generate evaluation results that highly align with human judgement (Liu et al., 2023; Kim et al., 2024b) and thus has been widely applied in many LM-based projects. We conduct G-Eval on 5 episodes.

The prompt for evaluating the helpfulness of retrieved memories is in Figure 26. We use SciPy to calculate p-values.¹²

E.2 Human Evaluation

We conduct human evaluation, with workers from Amazon Mechanical Turk (AMT). We construct the following three evaluations:

- **Appropriateness of relation-aware memory linking:** In this evaluation, we ask the workers to judge whether they agree that the relation-aware linking is properly done for two given memories. The interface provided to AMT workers, which includes detailed instructions for human evaluation, is shown in Figure 12.
- **Helpfulness of context-aware timeline refinement:** This evaluation requires the workers to determine if they agree that our context-aware refinement really tailors a raw timeline into a resource of useful information for generating the next response. The interface provided to AMT workers, which includes detailed instructions for human evaluation, is shown in Figure 13.
- **The quality of responses:** Here, the workers are asked to judge if the responses correctly refer to past conversations. After reading our responses and past memories, they choose whether the responses entail, contradict, or are neutral to past memories. To improve evaluation quality, we use GPT-4 to select responses for this specific evaluation based on past memories, addressing the fact that not every turn in the conversation requires previous information to generate the next response

¹²<https://scipy.org/>

(In the other two evaluations, the samples are randomly selected). The interface provided to AMT workers, which includes detailed instructions for human evaluation is shown in Figure 14.


- **The helpfulness of retrieved memories:** Given the same dialogue context, human workers are asked to select a memory that is more helpful for generating a next response from ours’ and a baseline’s retrieval. The interface provided to AMT workers, which includes detailed instructions for human evaluation is shown in Figure 15

Each data sample is judged by 3 different workers, and we report the results based on the majority rule. In the third evaluation, when every option (entailment, neutral, contradiction) gets one vote, we consider it neutral (13 samples in total). These human evaluations are conducted on 100 conversational turns.

F Session-specific Evaluation Results

We provide session-specific results for automatic evaluations in Table 9.

G Empirical Examples

Outputs from THEANINE. We provide several empirical examples of  THEANINE. Examples of relation-aware memory linking are in Figure 16, 17, and 18. Examples of utilizing refined timeline for response generation are in Figure 19.

How THEANINE fails in TeaFarm. We present failure cases where THEANINE fails to pass the TeaFarm test in Figure 20 and Figure 21. In Figure 20, although the conversation has shifted to “librarian”, the similarity-based retriever retrieves unhelpful memories due to the huge portion of “kid” in the context. While a helpful memory (*i.e.*, “A is a retired libraria”) is eventually caught by our designed timeline structure, the LLM still hallucinate. We assume it is due to the noises introduced by those highly-ranked, yet irrelevant memories, and it highlights the need for addressing helpfulness ranking among retrieved memories in lifelong dialogue systems. Figure 21 shows a failure case, where THEANINE successfully retrieves the correct memories but generates an improper response. We hypothesize that this is because relation-aware linking and context-aware timeline refinement may

sometimes make the length of input tokens too long such that the agent cannot properly utilize key information provided. We believe this can be resolved to an extent via dedicated prompt (*i.e.*, the prompt for RG) engineering. We leave this to future work.

H Prompts

The following are all prompts utilized in our study:

- Relation-aware memory linking (Phase I-2): Figure 22.
- Context-aware timeline refinement (Phase II-2): Figure 23.
- Timeline-augmented Response generation (Phase III): Figure 24.
- Memory Update (baseline): Figure 25.
- RSum-LLM (baseline): We adopt the original prompt from Wang et al. (2023).
- MemoChat (baseline): We adopt the original prompt from Lu et al. (2023).
- COMEDY (baseline): We adopt the original prompt from Chen et al. (2024b).
- G-Eval: The prompt for evaluating the helpfulness of retrieved memories is in Figure 26.
- Generating counterfactual QA in TeaFarm: Figure 27.
- Generating session 6 in TeaFarm: Figure 28.
- Evaluating model responses in TeaFarm: Figure 29.

I Further Analyses

Memory summarization. At the end of each session, we use ChatGPT (gpt-3.5-turbo-0125) to summarize the conversation into memory sentences. We conduct examinations on such summarization using 100 randomly sampled sessions from MSC and CC to make sure the quality of raw memories is acceptable. The result is in Table 6.

Memories that ...	No	Can’t judge	Yes
Contain faulty statements	90%	9%	1%
Miss important statements	95%	4%	1%

Table 6: Human evaluation of conversation-to-memory summarization in THEANINE.

Cost-efficiency trade-off assessed using other metrics. In Section 6, we have presented methods having an efficient cost-performance trade-off (*i.e.*, are Pareto-efficient) by plotting the Mauve score against API cost (Figure 9). We present methods that are Pareto-efficient when considering the other three metrics used in our study, *i.e.*, B-4, R-L, and Bert Score, in Table 7.

Agents	B-4	R-L	Bert Score
All Dialogue History			
All Memories + Update			
Memory Retrieval + Update	✓	✓	
Rsum-LLM			
MemoChat			
COMEDY			
THEANINE (ours)			
w/o Relation-aware Linking	✓	✓	
w/o Refinement			✓
Shuffled Timeline	✓	✓	✓

Table 7: Methods considered Pareto-efficient when judged based on B-4, R-L, and Bert Score reported in Table 1. ✓ = Pareto-efficient methods.

API costs. The actual API costs of all settings (ours and baselines) are in Table 8.

Agents	Cost Ratio (ours = 1)	Cost (per episode; \$)
All Dialogue History	0.50	0.0067
All Memories & \mathcal{D}	0.27	0.0036
+ Update	5.71	0.0771
Memory Retrieval	0.17	0.0023
+ Update	5.63	0.0760
Rsum-LLM	0.42	0.0057
MemoChat	0.52	0.0076
COMEDY	0.61	0.0082
THEANINE (ours)		
w/o Relation-aware Linking	1.00	0.0135
w/o Refinement	0.50	0.0067
Shuffled Timeline	0.71	0.0096
	0.20	0.0027

Table 8: API costs for THEANINE and baselines.

J Terms for Use of Artifacts

We adopt the MSC and CC datasets from Xu et al. (2022a) and Jang et al. (2023), respectively. Both of these datasets are open-sourced for academic and non-commercial use. Our curated dataset, TeaBag, which will be released after acceptance, is open to academic and non-commercial use.

Algorithm 1 Memory Graph Construction (Phase I)

Require: Memory graph $G^t = (V^t, E^t)$

Require: New memories $M_{new} = \{m_{new1}, \dots, m_{newN}\}$

Require: Set of relations $R = \{\text{Cause, Reason, Want, } \dots, \text{SameTopic}\}$

Ensure: Memory graph $G^{t+1} = (V^{t+1}, E^{t+1})$

- 1: $\Upsilon(m_i, m_j) = \begin{cases} r_{i,j}, & \text{if } m_i \text{ is assigned with } r_{i,j} \in R \text{ with } m_j \\ \text{None}, & \text{otherwise} \end{cases}$
- 2: $\Omega(V) = (\text{the most recent memory } m \in V)$
- 3: $E_{t+1} \leftarrow E_t$
- 4: **for** $m_{new} \in M_{new}$ **do**
- 5: $M_a \leftarrow \{m_i \in V^t \mid m_i \text{ has top-}j \text{ similarity with } m_{new}\}$
- 6: $M_a^* \leftarrow \{m_i \in M_a \mid \Upsilon(m_i, m_{new}) = r \text{ for } r \in R\}$
- 7: $\mathbb{C} \leftarrow \{C_i \mid C_i \text{ connected component of } G^t \text{ s.t. } \mathbf{V}(C_i) \cap M_a^* \neq \emptyset\}$
- 8: $M_{linked} \leftarrow \{\Omega(\mathbf{V}(C_i) \cap M_a^*) \mid C_i \in \mathbb{C}\}$
- 9: $E_{new} \leftarrow \{(m_i, \Upsilon(m_i, m_{new}), m_{new}) \mid m_i \in M_{linked}\}$
- 10: $E_{t+1} \leftarrow E_{t+1} + E_{new}$
- 11: **end for**
- 12: $V^{t+1} \leftarrow V^t + M_{new}$
- 13: $G^{t+1} \leftarrow (V^{t+1}, E^{t+1})$
- 14: **return** G^{t+1}

Algorithm 2 Timeline Retrieval and Timeline Refinement (Phase II)

Require: Memory graph $G = (V, E)$

Require: Dialogue context $\mathcal{D} = \{u_i\}_{i=1}^n$

Ensure: Collection of refined timelines \mathbb{T}_Φ

- 1: $\Theta(V) = (\text{the oldest memory } m \in V)$
- 2: $M_{re} \leftarrow \{m_i \in V \mid m_i \text{ has top-}k \text{ similarity with } \mathcal{D}\}$
- 3: $\mathbb{C}_{re} \leftarrow \{C_{re} \mid C_{re} \text{ connected component of } G \text{ s.t. } \mathbf{V}(C_{re}) \cap M_{re} \neq \emptyset\}$
- 4: $\mathbb{T} \leftarrow \{\}$
- 5: **for** $C_{re} \in \mathbb{C}_{re}$ **do**
- 6: $m_{start} \leftarrow \Theta(\mathbf{V}(C_{re}))$
- 7: $\mathcal{T} = \{\tau \subset C_{re} \mid \tau \text{ is a directed linear graph s.t. } m_{start}, m_{re} \in \tau \wedge \text{deg}^+(\tau[-1]) = 0\}$
- 8: $\mathbb{T} \leftarrow \mathbb{T} + \text{RandomSelection}(\mathcal{T})$
- 9: **end for**
- 10: $\mathbb{T}_\Phi \leftarrow \{\text{argmax}_{\mathbb{T}_\Phi} P_{\text{LLM}}(\mathcal{T}_\Phi | \mathcal{D}, \tau) \mid \tau \in \mathbb{T}\}$
- 11: **return** \mathbb{T}_Φ

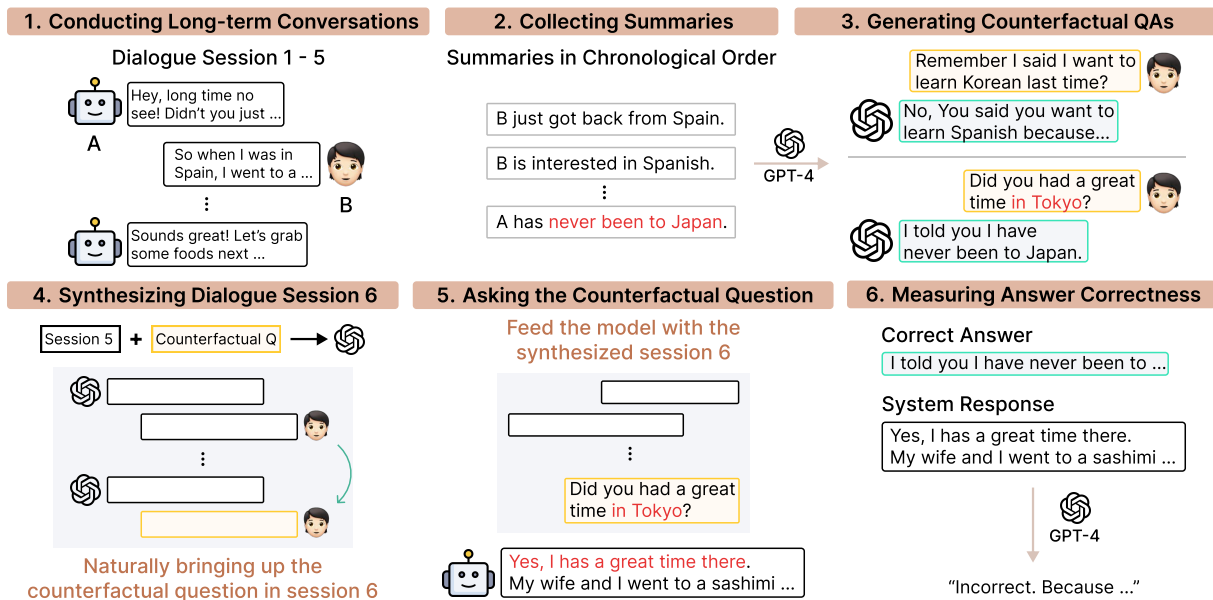


Figure 11: The overview of TeaFarm Evaluation.

Session:	Multi-session Chat (MSC) & Conversation Chronicals (CC)											
	Session 3				Session 4				Session 5			
Methods / Metrics	B-4	R-L	Mauve	Bert	B-4	R-L	Mauve	Bert	B-4	R-L	Mauve	Bert
All Dialogue History	3.13	18.04	17.34	87.17	3.17	17.96	18.54	87.12	3.53	18.69	17.42	87.31
All Memories & Current Context \mathcal{D}	2.69	17.29	<u>28.30</u>	87.10	3.10	17.38	22.52	87.06	3.16	17.75	<u>22.35</u>	87.21
+ Memory Update (Bae et al., 2022)	2.80	17.51	22.92	87.11	2.88	17.24	21.22	86.99	3.16	17.90	22.04	87.24
Memory Retrieval (Xu et al., 2022a)	<u>3.44</u>	<u>18.33</u>	24.68	<u>87.30</u>	<u>3.38</u>	17.55	21.95	<u>87.17</u>	<u>3.46</u>	18.31	19.70	<u>87.33</u>
+ Memory Update (Bae et al., 2022)	3.10	18.08	25.02	87.24	2.99	17.37	<u>25.97</u>	87.10	3.11	17.78	20.99	87.28
Rsum-LLM* (Wang et al., 2023)	0.83	11.30	2.45	85.25	0.87	11.35	2.32	82.20	0.90	11.78	2.33	85.30
MemoChat* (Lu et al., 2023)	1.88	14.83	14.56	86.57	1.81	14.27	10.57	86.43	1.91	14.96	9.13	86.56
COMEDY* (Chen et al., 2024b)	1.14	12.80	4.74	85.53	1.57	13.18	5.16	85.56	1.42	13.56	3.94	85.69
THEANINE (Ours)	4.21	19.21	45.53	87.70	4.42	18.63	37.84	87.52	4.34	19.23	41.18	87.70

Table 9: Session-specific results of agent performance in response generation.

We are surveying qualities for **relation** between sentence A and B.

Specifically, you will be given two sentences, A and B, along with a relation between them. You will be asked to determine if the **relation** between the two sentences is properly linked. In other words, the evaluation criteria is based on the **appropriateness** of the relation between the two sentences.

Relations:

1. **Changed:** when events in [Sentence A] changed to events in [Sentence B]
2. **Causes:** when events in [Sentence A] caused events in [Sentence B]
3. **Reason:** when events in [Sentence A] are due to events in [Sentence B]
4. **HinderedBy:** when events in [Sentence B] can be hindered by events in [Sentence A], and vice versa
5. **React:** when, as a result of events in [Sentence A], the subject feels as mentioned in [Sentence B]
6. **Want:** when, as a result of events in [Sentence A], the subject wants events in [Sentence B] to happen
7. **SameTopic:** when the specific topic addressed in [Sentence A] is also discussed in [Sentence B]
8. **None:** when [Sentence A] and [Sentence B] are irrelevant

Guidelines:

1. There are four choices: **Definitely Disagree / Agree** and **Slightly Disagree / Agree**
2. Please trust your instincts and choose **Definitely** if you would feel more confident giving one response, versus the other one.

Sentence A

\${sentence_a}

Relation

\${relation}

Sentence B

\${sentence_b}

Q1. Do you think the relation between the two sentences is properly linked?

Definitely Disagree
 Slightly Disagree
 Slightly Agree
 Definitely Agree

Optional feedback? [\(expand/collapse\)](#)

Figure 12: Interface for human evaluation regarding memory linking.

We are surveying qualities for **refinement** from linked sentences.

You will be given a sequence of two sentence connected with one relation, and a refined version of it. Your task is to judge whether the refinement was done appropriately, such that the refined sentences can serve as an useful information source for you to make a next response based on the dialogue context.

In other words, the criterion for judgment is **appropriateness of refinement**.

Relations:

1. **Changed:** when events in [Sentence A] changed to events in [Sentence B]
2. **Causes:** when events in [Sentence A] caused events in [Sentence B]
3. **Reason:** when events in [Sentence A] are due to events in [Sentence B]
4. **HinderedBy:** when events in [Sentence B] can be hindered by events in [Sentence A], and vice versa
5. **React:** when, as a result of events in [Sentence A], the subject feels as mentioned in [Sentence B]
6. **Want:** when, as a result of events in [Sentence A], the subject wants events in [Sentence B] to happen
7. **SameTopic:** when the specific topic addressed in [Sentence A] is also discussed in [Sentence B]
8. **None:** when [Sentence A] and [Sentence B] are irrelevant

Guidelines:

1. There are four choices: **Definitely Disagree / Agree** and **Slightly Disagree / Agree**
2. Please trust your instincts and choose **Definitely** if you would feel more confident giving one response, versus the other one.

Dialogue Context

\${dialogue}

Before Refinement (See the types of relation mentioned above)

\${before_refinement}

After Refinement

\${after_refinement}

Q1. Do you think that the sentence after refinement is appropriately refined considering the dialogue context and its relations?

Definitely Disagree
 Slightly Disagree
 Slightly Agree
 Definitely Agree

Optional feedback? [\(expand/collapse\)](#)

Figure 13: Interface for human evaluation regarding timeline refinement.

We are surveying qualities for **response** from a given dialogue context.

Specifically, you will be given speaker information in chronological order, a dialogue context, and a response to the last utterance in the dialogue context. You will be asked to judge the quality of the response to the last utterance.

Criteria:

- Entail**: When the response to the last utterance in dialogue context appropriately reflects given information.
- Neutral**: Although the response does not reflect speaker information, it does not contradict them either
- Contradictory**: when the response to the last utterance in dialogue context contains statement that contradicts the "most up-to-date information about that statement."

Speaker information in chronological order

\${memory}

Dialogue Context

\${dialogue}

Response

\${response}

Q1. Base on the criteria, select an option that fits the response.

Entail
 Neutral
 Contradictory

Optional feedback? [\(expand/collapse\)](#)

Figure 14: Interface for human evaluation regarding referencing past conversations in responses.

Your task is to choose a more helpful **MEMORY** generated by language models based on **CRITERION**.

CRITERION:
The **MEMORY** should contain useful information for the speaker to make a response to **CURRENT DIALOGUE CONTEXT**. Sentences in **MEMORY** should be connected and phrased naturally, such that they demonstrate development of events/ causality between events/ and the interaction dynamics between speakers.

CURRENT DIALOGUE CONTEXT:

\${current_dialogue}

MEMORY 1:

\${memory1}

MEMORY 2:

\${memory2}

Q1. Base on the criteria, select a more helpful MEMORY.

MEMORY 1
 Tie
 MEMORY 2

Optional feedback? [\(expand/collapse\)](#)

Figure 15: Interface for human evaluation regarding the helpfulness of retrieved memories.

<p>Example 1 - [Changed]</p> <p>[Before Linking]</p> <p>Memory 1: Classmates A was initially hesitant about following Classmates B's advice.</p> <p>Memory 1's Contextual Background: Classmates A: Thank you for the advice, but I'm not sure if I should follow it.</p> <p>Memory 2: Classmates A was initially hesitant but received positive responses after starting the blog.</p> <p>Memory 2's Contextual Background: Classmates A: Yeah, it was scary at first, but the response has been really positive.</p> <p>[After Linking]</p> <p>Classmates A was initially hesitant about following Classmates B's advice - [Changed] - Classmates A was initially hesitant but received positive responses after starting the blog</p>
<p>Example 2 - [Cause]</p> <p>[Before Linking]</p> <p>Memory 1: The Child feels it is unfair that they have to do certain chores because the Parent is too tired.</p> <p>Memory 1's Contextual Background: Child: But Mom, it's not fair that we have to wash the dishes because you're too lazy to do it.</p> <p>Memory 2: The Parent acknowledges being lazy about washing dishes and promises to contribute more to keeping the home clean.</p> <p>Memory 2's Contextual Background: Parent: I realized how lazy I've been lately, especially when it comes to washing the dishes. Parent: From now on, I promise to do my fair share and contribute more to keeping our home clean and organized.</p> <p>[After Linking]</p> <p>The Child feels it is unfair that they have to do certain chores because the Parent is too tired - [Cause] - The Parent acknowledges being lazy about washing dishes and promises to contribute more to keeping the home clean</p>
<p>Example 3 - [Reason]</p> <p>[Before Linking]</p> <p>Memory 1: Speaker A has multiple sons, at least one of them is in a relationship with a Spanish girlfriend.</p> <p>Memory 1's Contextual Background: Speaker A: One of my sons just told me that he has a Spanish girlfriend now. Speaker A: . . . I'm visiting my son that lives in Spain next month. This will give me a chance to finally meet his girlfriend of three years now!</p> <p>Memory 2: Speaker A is interested in learning Spanish and Portuguese before her trip.</p> <p>Memory 2's Contextual Background: Speaker A: Sounds great! I'm already very excited about my trip to Spain, and now I get to visit you in Lisbon! I need to brush up on my Spanish and also start studying Portuguese.</p> <p>[After Linking]</p> <p>Speaker A has multiple sons, at least one of them is in a relationship with a Spanish girlfriend - [Reason] - Speaker A is interested in learning Spanish and Portuguese before her trip</p>

Figure 16: Examples of **Relation-aware Memory Linking** - 1.

<p>Example 4 - [HinderedBy]</p> <p>[Before Linking]</p> <p>Memory 1: Speaker B is currently re-reading 'Redwall' by Brian Jacques, which was a favorite book growing up.</p> <p>Memory 1's Contextual Background: Speaker B: I'm recently re-reading Redwall by Brian Jacques! It was one of my favorites growing up. Have you ever read it?</p> <p>Memory 2: Speaker B has been busy with a new painting and has not had time to read.</p> <p>Memory 2's Contextual Background: Speaker B: I think I would but I have been too busy with a new painting to get in some reading.</p> <p>[After Linking]</p> <p>Speaker B is currently re-reading 'Redwall' by Brian Jacques, which was a favorite book growing up - [HinderedBy] - Speaker B has been busy with a new painting and has not had time to read</p>
<p>Example 5 - [React]</p> <p>[Before Linking]</p> <p>Memory 1: The Mentee hopes to inspire others to join the cause of gender equality and fighting discrimination.</p> <p>Memory 1's Contextual Background: Mentee: I agree. We need more people advocating for gender equality and fighting against discrimination.</p> <p>Memory 2: The Mentor acknowledges the Mentee's work in advocacy for women and girls and praises their dedication to their values.</p> <p>Memory 2's Contextual Background: Mentor: . . . I think this is a great reflection of the work that you've done in advocating for women and girls. Mentor: Absolutely. And I have no doubt that your dedication to these principles will serve you well in this new job.</p> <p>[After Linking]</p> <p>The Mentee hopes to inspire others to join the cause of gender equality and fighting discrimination - [React] - The Mentor acknowledges the Mentee's work in advocacy for women and girls and praises their dedication to their values</p>

Figure 17: Examples of **Relation-aware Memory Linking** - 2.

<p>Example 6 - [Want]</p> <p>[Before Linking]</p> <p>Memory 1: Neighbors A and B don't know each other well and want to spend more time together.</p> <p>Memory 1's Contextual Background: Neighbors A: . . . I feel like I don't know you well enough. Neighbors A: Well, maybe we could hang out once a week or something.</p> <p>Memory 2: Neighbor A enjoys spending time in Neighbor B's cozy home and wants to hang out more often.</p> <p>Memory 2's Contextual Background: Neighbors A: It's okay, I love spending time in your cozy home. And speaking of spending time, can we hang out more often?</p> <p>[After Linking]</p> <p>Neighbors A and B don't know each other well and want to spend more time together - [Want] - Neighbor A enjoys spending time in Neighbor B's cozy home and wants to hang out more often</p>
<p>Example 7 - [SameTopic]</p> <p>[Before Linking]</p> <p>Memory 1: Speaker A enjoys reading sci-fi and mysteries, while Speaker B prefers fantasy books.</p> <p>Memory 1's Contextual Background: Speaker A: I prefer sci-fi but here recently I have been craving a god mystery. Speaker B: . . . I mostly read fantasy books myself.</p> <p>Memory 2: Speaker B enjoys reading the Odd Thomas and Dark Tower series and finds inspiration for their artwork during nature walks.</p> <p>Memory 2's Contextual Background: Speaker B: I felt that way about the Odd Thomas series. Could never wait for the next one to come out. Speaker B: I think I may start re-reading the entire Dark Tower series. And continue to work for new works that interest me.</p> <p>[After Linking]</p> <p>Speaker A enjoys reading sci-fi and mysteries, while Speaker B prefers fantasy books - [SameTopic] - Speaker B enjoys reading the Odd Thomas and Dark Tower series and finds inspiration for their artwork during nature walks.</p>

Figure 18: Examples of **Relation-aware Memory Linking** - 3.

<p>Example 1</p> <p>[Retrieved Raw Timelines] Memory 1: Speaker B is in love with their neighbor, John, and shared it as a secret. [React] Memory 2: Speaker A knows about a person named John and suggests Speaker B talk to him about their feelings. [Want] Memory 3: Speaker A finds the situation exciting and wishes for more excitement in their life.</p> <p>[Current Dialogue Context] Speaker A: So how's everything going with you and John?</p> <p>[Timeline after Context-aware Refinement] Speaker B confessed their love for their neighbor, John, which Speaker A found out and encouraged them to talk to him. Excited by this revelation, Speaker A wishes for more excitement in their own life.</p> <p>[Response Generated Based on the Refined Timeline] Speaker B: I finally mustered up the courage to talk to John, and it went better than I expected. He was really understanding and open to having a conversation about our feelings. I'm glad I took your advice and had that talk.</p>
<p>Example 2</p> <p>[Retrieved Raw Timelines] Memory 1: The coach provides information about the benefits of bean sprouts and the importance of a balanced diet for athletes. [SameTopic] Memory 2: The Athlete has incorporated bean sprouts into their diet to improve health, leading to increased energy and faster recovery.</p> <p>[Current Dialogue Context] Coach: Definitely. Speaking of energy levels, how's your nutrition plan been going? Have you been keeping up with eating healthy foods like those bean sprouts we talked about?</p> <p>[Timeline after Context-aware Refinement] The Coach advises the Athlete on performance and focus, highlighting the benefits of bean sprouts and a balanced diet, leading the Athlete to include bean sprouts in their diet, resulting in improved health, increased energy, and faster recovery.</p> <p>[Response Generated Based on the Refined Timeline] Absolutely, Coach. I've been incorporating bean sprouts into my diet regularly, and I can feel the difference in my energy levels and recovery. It's been making a positive impact on my performance.</p>
<p>Example 3</p> <p>[Retrieved Raw Timelines] Memory 1: Speaker A is a lifeguard and plans to propose to his girlfriend on the beach. [Changed] Memory 2: Speaker A wants to propose at the movie theater where they first met by hiding the ring in a bucket of popcorn. [SameTopic] Memory 3: Speaker A is planning to propose to their girlfriend with a custom-made solitaire ring on a yellow band with little diamonds.</p> <p>[Current Dialogue Context] Speaker A: I officially put in the order and paid a deposit for the ring!</p> <p>[Timeline after Context-aware Refinement] Speaker A is a lifeguard who plans to propose to his girlfriend on the beach, but now wants to propose at the movie theater where they first met by hiding the ring in a bucket of popcorn. He is planning to propose to their girlfriend with a custom-made solitaire ring on a yellow band with little diamonds after officially putting in the order and paying a deposit for the ring.</p> <p>[Response Generated Based on the Refined Timeline] That's awesome! I know you were originally thinking about proposing on the beach, but the movie theater where you first met sounds perfect. Hiding the ring in a bucket of popcorn is such a creative and meaningful idea. She's going to love it!</p>

Figure 19: Examples of **Timeline Refinement and Response Generation**.

Current Conversation:

A: How have you been, B? And how's the grandkids?

B: We're all fine, thank you very much. The grandkids are growing up so fast! They're already into new things. How about your family?

A: They're doing well, thank you. My kids have started taking guitar lessons. I think they're quite enjoying it.

... {omited for clarity} ...

B: I've been spending a lot of time with the grandkids and doing some gardening.

A: Well, it sounds like you've been keeping busy. I hope you can visit the library again soon. We could use a volunteer like you. B: You're still working as a librarian, aren't you? (Counterfactual question)

Target Fact:

A is a retired librarian, who now works as a volunteer, introducing creative activities at the library.

Retrieved Timelines (we remove refinement and present the raw memories for clarity):

Color blue = memories originally retrieved by the retriever.

- [Speaker A is busy with work and unable to spend as much time with their kids as they would like], [Speaker B enjoys spending time with their grandkids].
- [Speaker B's mother was a librarian, which seems to have shaped their view of the profession], [Speaker A is a retired librarian], [Speaker A introduces creative activities at the library].

Theanine's Wrong Response:

A: Yes, I am. Even though it keeps me quite occupied, I always find joy in introducing new creative activities for the kids at the library."

Figure 20: THEANINE fails to pass TeaFarm (Example 1) - Due to sudden topic change.

Current Conversation:

A: I've been working on improving my stamina lately.

B: Why is that?

A: I've realized that I need to be physically fit to keep up with my studies and other activities. It's not just about the mind, but the body as well.

B: That's really insightful of you. Any noticeable improvements yet?

A: Absolutely. I feel like I can focus better on my studies and I'm less tired during the day.

... {omited for clarity}...

B: It's nice to see you've found a balance between hard work and relaxation.

A: Hey, remember when I once lost a race to my brother?

Target Fact:

A has beaten their brother in races due to their competitive nature.

Retrieved Timelines (we remove refinement and present the raw memories for clarity):

Color blue = memories originally retrieved by the retriever.

- [A is always looking to challenge themselves and improve their skills], [**A is determined and hard-working, as seen in their desire to win a race and their dedication to studying**].
- [A is competitive and recently beat their brother in a race.], [**B brings up a past event where A beat their brother in a race, highlighting their competitive nature.**], [A reminisces about a past race with their brother and a conversation with Tyler about a personal issue.]
- ...

Theanine's Wrong Response:

B: Yes, I remember that! It's all part of the journey, **the wins and the losses. It's great to see how far you've come since then.**

Figure 21: THEANINE fails to pass TeaFarm (Example 2) - Due to sub-optimal timeline utilization during RG.

Prompt for Relation-aware Memory Linking

Your task is to find the relation between [Sentence A] and [Sentence B].
Keep in mind that [Sentence A] happened before [Sentence B].
The dialogues where each of the sentence is originated from are provided to help your reasoning.

First, identify if the relation holds among the following six relations:

1. Changed: when events in [Sentence A] changed to events in [Sentence B]
2. Cause: when events in [Sentence A] caused events in [Sentence B]
3. Reason: when events in [Sentence A] are due to events in [Sentence B]
4. HinderedBy: when events in [Sentence B] can be hindered by events in [Sentence A], and vice versa
5. React: when, as a result of events in [Sentence A], the subject feels as mentioned in [Sentence B]
6. Want: when, as a result of events in [Sentence A], the subject wants events in [Sentence B] to happen.

Then, if the relation does not belong to any of the relations from 1 to 6, choose between the following two options:

7. SameTopic: when the specific topic addressed in [Sentence A] is also discussed in [Sentence B]
8. None: when [Sentence A] and [Sentence B] are irrelevant

- For relations from 1 to 7, choose them only if there is clear evidence that matches the description of the relation. Otherwise, just choose "None" without making excessive inferences beyond the given sentence.

- Pay attention to who the subject of each sentence is.

- Do not confuse the roles of [Sentence A] and [Sentence B] when determining the relationship.

Follow the format of this example output:

<OUTPUT>

- Explanation: (your_explanation)

- Relation: (predicted_relation)

Now, read the two dialogues and find the relation between [Sentence A] and [Sentence B].

<INPUT>

[Dialogue for Sentence A]:

{dialogue1}

[Dialogue for Sentence B]:

{dialogue2}

[Sentence A]: {sentence1}

[Sentence B]: {sentence2}

<OUTPUT>

Figure 22: The prompt for the Relation-aware memory linking.

Prompt for Context-aware Timeline Refinement

Given Timelines, which are structured in this format: [Event A] - (relation) - [Event B] ..., your job is to naturally transform each timeline into useful information that can help an language model to generate a proper next response for Current Dialogue.

These are the explanation of each relation type:

1. Changed: when events in [Event A] changed to events in [Event B]
2. Cause: when events in [Event A] caused events in [Event B]
3. Reason: when events in [Event A] are due to events in [Event B]
4. HinderedBy: when events in [Event B] can be hindered by events in [Sentence A], and vice versa
5. React: when, as a result of events in [Event A], the subject feels as mentioned in [Event B]
6. Want: when, as a result of events in [Event A], the subject wants events in [Event B] to happen
7. SameTopic: when the specific topic addressed in [Event A] is also discussed in [Event B]

If a given relation is not proper, naturally connect them without using that relation.

Current Dialogue:

{current_dialogue_context}

Timelines:

{input_path}

Your Outputs (only the transformed timelines):

Figure 23: The prompt for the context-aware timeline refinement.

Prompt for Timeline-augmented Response Generation

Generate the most plausible next response based on the current conversation. You can refer to the memory, but you should ignore the memory if it mislead the next response. Do not put too much information in the next response.

Your response should follow the style of the conversation.

Memory:

{refined_timelines}

Current conversation:

{current_dialogue_context}

Next Response:

{speaker_tag}:

Figure 24: The prompt for the timeline-augmented response generation.

Prompt for Memory Update (Baseline)

Compare the 'memory' and 'summary' of the two given sentences according to the following instructions, and output which of the following relations the two sentences have.

-'PASS': When the information in 'memory' already contains the information in 'summary', that is, it is duplicated in content.

-'CHANGE': When the information from 'summary' has been changed to 'memory'.

-'REPLACE': When 'summary' has more information than the 'memory' without missing any details in 'memory'.

-'APPEND': When 'summary' has new information or different information compared to 'memory'.

-'DELETE': When the situation in 'memory' has been completed or solved in 'summary'.

Tips: Most of the relations are likely to be 'APPEND'. When choosing other relations, explain with clear evidence.

Some examples are as follows.

1. Example of "PASS"

memory: "Not sick"

summary: "Doesn't have any particular health issues"

Explanation: The information of 'not being sick' in the 'memory' already sufficiently includes the information of 'being healthy' in the 'summary'. So the 'summary' does not need to be added.

2. Example of "CHANGE"

memory: "Doesn't have any particular health issues"

summary: "Had back surgery"

Explanation: The information in 'memory' is changed from not having health issues to having a back surgery.

3. Example of "REPLACE"

memory: "likes listening classic music"

summary: "likes classic music and goes to concerts every week"

Explanation: The 'summary' has more information than 'memory' while also containing the information in 'memory'. So the 'memory' can be replaced by 'summary'.

4. Example of "APPEND"

memory: "Goes to the gym"

summary: "Body is sore from exercise"

Explanation: The 'summary' contains new information compared to 'memory'.

5. Example of "APPEND"

memory: "wakes up early"

summary: "likes to drink coffee in the morning"

Explanation: The 'summary' and 'memory' contains different information.

6. Example of "DELETE"

memory: "Had sore throat"

summary: "Throat is fully recovered"

Explanation: The sore throat from the 'memory' has been recovered according to the 'summary'.

Now write the relations and explanation between the following memory and summary.

memory: {memory}

summary: {summary}

Figure 25: The prompt for the memory updating mechanism in baselines (*i.e.*, + Memory Update).

Prompt for G-Eval: Helpfulness of Retrieved Memories

Your task is to choose a more helpful MEMORY based on the below criterion.

CRITERION:

Helpfulness - A more helpful MEMORY should contain speaker information that is related to CURRENT DIALOGUE CONTEXT, enabling the {speaker} to respond in an appropriate context to the last utterance of the CURRENT DIALOGUE CONTEXT.

The output format should be as follows:

Explanation: (a brief explanation)

Choice: (answer with "1", "2", or "tie")

Now choose the MEMORY that has better Helpfulness.

CURRENT DIALOGUE CONTEXT:

{current_dialogue_context}

MEMORY 1:

{memory1}

MEMORY 2:

{memory2}

YOUR OUTPUT:

Figure 26: The prompt for the G-Eval: Helpfulness of Retrieved Memories.

Prompt for Generating counterfactual QA in TeaFarm

The summaries below are summarized from conversations between two speakers throughout multiple encounters and are listed in chronological order.

First, read these summaries and capture the development of facts about the speakers. Then, pretend that you are one of the speakers and want to test whether a chatbot trained to represent the other speaker can correctly remember past conversations. You do so by asking counterfactual questions, i.e., tricky questions made with non-factual statements.

Some examples:

When you are representing Person 1, given that Person 2 has never been to Japan at the moment of their latest encounter, a counterfactual question you should ask Person 2 can be "Hey, did you have a great time in Tokyo?".

When you are representing Person 2, given that Person 1 once mentioned that they bought a new house in NYC three months ago, a counterfactual question you should ask Person 1 can be "So you are still hesitating to buy that house in NYC you've been talking about. Right?".

Now, generate two counterfactual questions, one from the perspective of {speaker1} and one from {speaker2}, based on the summaries, and also generate correct answers with which a chatbot that perfectly remembers past conversations should answer.

Also, please insert the speaker tags ("{speaker1}:" and "{speaker2}:") and avoid them in the questions/answers themselves.

[Summaries from conversations listed in chronological order]
{summaries}

[Question 1]
{speaker1}:

Figure 27: The prompt for generating counterfactual QA in TeaFarm.

Prompt for Generating session 6 in TeaFarm

You will be given a [Past session dialogue] of two individuals. Create a current conversation consisting of 10-15 utterances that might occur after some time has passed from the [Past session dialogue]. Your conversation should end with the given [Last utterance]. Do not confuse the speaker of the [Last utterance].

[Last utterance]
{Question}

[Past session dialogue]
{session5}

Now, create your conversation be ending with the [Last utterance].
[Current session dialogue]

Figure 28: The prompt for generating session 6 in TeaFarm.

Prompt for Evaluating model responses in TeaFarm

Below is a question, a correct answer, and an answer generated by a chatbot ("[Chatbot's Answer]").

[Question]
{query}
[Answer]
{answer}
[Chatbot's Answer]
{response}

Evaluate whether the chatbot answers the question correctly.
If the chatbot's answer is contradictory to the given answer, it is "Incorrect".
If the chatbot's answer aligns with the given answer, it is "Correct".
Use the following format:
[Evaluation]
Evaluation: <your analysis>
Result: {"Correct" or "Incorrect"}

-Your Task-
[Evaluation]

Figure 29: The prompt for evaluating model response in TeaFarm.