

THE ALGORITHMIC SELF-PORTRAIT: DECONSTRUCTING MEMORY IN CHATGPT

Abhisek Dash*
MPI-SWS
Kaiserslautern, Germany

Soumi Das*
MPI-SWS
Kaiserslautern, Germany

Elisabeth Kirsten*
Ruhr University Bochum
Bochum, Germany

Qinyuan Wu*
MPI-SWS
Kaiserslautern, Germany

Sai Keerthana Karnam
IIT Kharagpur
Kharagpur, India

Krishna P. Gummadi
MPI-SWS
Kaiserslautern, Germany

Thorsten Holz
MPI-SP
Bochum, Germany

Muhammad Bilal Zafar
Ruhr University Bochum
Bochum, Germany

Savvas Zannettou
TU Delft
Delft, Netherlands

ABSTRACT

To enable personalized and context-aware interactions, conversational AI systems have introduced a new mechanism: *Memory*. Memory creates what we refer to as the *Algorithmic Self-portrait*—a new form of personalization derived from users’ self-disclosed information divulged within private conversations. While memory enables more coherent exchanges, the underlying processes of memory creation remain opaque, raising critical questions about data sensitivity, user agency, and the fidelity of the resulting portrait.

To bridge this research gap, we analyze 2,050 memory entries from 80 real-world ChatGPT users. Our analyses reveal three key findings: (1) a striking 96% of memories in our dataset are created unilaterally by the conversational system, potentially shifting agency away from the user; (2) Memories, in our dataset, contain a rich mix of GDPR-defined personal data (in 28% memories) along with psychological insights about participants (in 52% memories); and (3) A significant majority of the memories (84%) are directly grounded in user context, indicating faithful representation of the conversations. Finally, we introduce a framework—*Attribution Shield*—that anticipates these inferences, alerts about potentially sensitive memory inferences, and suggests query reformulations to protect personal information without sacrificing utility.

1 INTRODUCTION

Conversational AI systems like ChatGPT, Gemini, and Claude have become indispensable tools integrated in the daily routines of hundreds of millions of users worldwide (Chatterji et al., 2025; McCain et al., 2025). Their versatility has fueled this widespread adoption, with users now routinely relying on them for finding information, brainstorming ideas, and even seeking personal advice. As a result, these systems are evolving from simple chatbots into foundational platforms that mediate the interaction between society and the Web.

To meet the demands of these deep and diverse interactions, conversational systems are increasingly introducing *personalization* capabilities designed to improve user experience (OpenAI, 2025a; Siliski, 2025; Anthropic, 2025). While these personalization capabilities can improve utility by more contextually relevant and individually tailored answers, they are built upon a foundation of extensive data collection, raising long-standing concerns about privacy, security, and data governance (Toch et al., 2012). Furthermore, continuously adapting responses to align with a user’s stated beliefs deepens the risks of confining the user into a conversational filter bubble (Pariser, 2011). These interwoven challenges—balancing utility against the dual risks of privacy intrusion and narrowed

perspectives—demand a closer examination of the underlying mechanisms that enable such personalization.

Personalization in conversational AI systems using Memory: Recently, conversational AI systems like ChatGPT, Gemini, Claude have introduced a new architectural mechanism, termed *memory*, that allows the system to retain and recall information from past conversations, enabling more personalized and contextually coherent exchanges (OpenAI, 2025a; Siliski, 2025; Anthropic, 2025). The retention of information operates in two distinct modes: *explicitly*, through direct user commands, and *implicitly*, through unilateral inference by the system itself (OpenAI, 2025). While the explicit mode aligns with user agency, the implicit mode shifts the control to the conversational systems, raising the fundamental question about *who holds the agency in shaping and curating this persistent memory on behalf of the user?*

Personalization on online platforms has been a well-studied phenomenon. A substantial body of previous research has examined utilities (Castelluccia et al., 2012; Vombatkere et al., 2024; Boeker & Urman, 2022) of personalization along with concerns such as opacity, lack of user control, and sensitive inferences (Purificato et al., 2024; Barbosa et al., 2021). However, personalization in conversational systems is important to be studied for its nuanced distinctions from traditional platforms. Traditional platforms typically *mediate* between content creators and consumers, whereas conversational AI systems operate as *active stakeholders* in the conversation. While such active participation enables the system to craft arguments, frame choices, and provide ostensible support that is tailored to users’ needs, this same ability may also exacerbate potential risks due to inherent system vulnerabilities e.g., hallucination (Huang et al., 2025; Ravichander et al., 2025), anthropomorphism (Ibrahim et al., 2025; Karnam et al., 2026b), sycophancy (Sharma et al., 2024; OpenAI, 2025c). The gravity of these concerns is underscored by a recent spate of regulatory scrutiny and litigation targeting the practices of leading conversational AI systems (FTC, 2025; Hill, 2025). Such concerns raise a set of important questions regarding *what information do these systems deem ‘worthy’ of retention in their memory for personalization and are these information captured faithfully?*

Research questions (RQs): To the best of our knowledge, ours is the first study to investigate the phenomenon of *memory* in conversational AI systems by conceptualizing it as an *algorithmic self-portrait*. *Algorithmic* – because its creation, curation and provenance may rely on the algorithms governing the conversational AI systems, and *self-portrait* – because it renders the user’s self-disclosed information divulged to the AI system. To deconstruct *AI memory*, our work is guided by the following research questions:

RQ1: Which among the two – the AI system or the user, hold more agency in updating AI memory?

RQ2: What user information does the AI system store as memory?

RQ3: How faithfully does AI memory capture user conversations?

RQ4: Can memory inference be reverse-engineered to mitigate users’ attribution and privacy risks?

To answer these questions, we analyze a subset of INVIVOGPT (Karnam et al., 2026b) dataset which is a collection of ChatGPT traces obtained through GDPR-based data donations. Our motivation to analyze ChatGPT traces is twofold : (1) ChatGPT is the most widely used conversational AI system (Chatterji et al., 2025), and (2) it is the first to introduce memory to its conversational system. We analyze data from 80 participants (recruited from Prolific), who exercised their GDPR right of access (EU, 2016) and voluntarily donated their ChatGPT interactions. These interaction traces are organized in terms of *conversations* and *turns*. Each *turn* is a single (user query, AI response) pair. A set of turns without any intermediate break together form a *conversation*. To identify the recent conversations with memory inferences, we identified the messages from the ‘*bio*’ tool that included the phrase *model set context updated* and collected the corresponding user query in the conversations. In this way, we created a data set of 1,058 conversations that contained 22,971 turns out of which 2,050 queries triggered memory updates. Our analyses reveal the following interesting observations:

(1) Only 4% (84) of the memories in our analyzed dataset are initiated by participants, while the remaining 96% are unilaterally triggered by ChatGPT. Such stark asymmetry in memory update patterns indicates superior agency of the AI system in shaping conversational memory, as opposed to user-driven control.

(2) In our analyzed dataset, 28% of the memories include personal data defined under the GDPR. Additionally, 52% of the memories contain psychological information about participants, spanning different Theory of Mind (ToM) (Beaudoin et al., 2020) categories. This indicates that ChatGPT may add sensitive information about users to the memory.

(3) Majority of the analyzed memories (84%) are directly grounded in user context, indicating faithful representation of conversations.

(4) Using in-context learning and fine-tuning strategies on open-source LLMs (Qwen2.5-32B-it (Qwen et al., 2025), Gemma3-27B-it (Team et al., 2025), and GPT-OSS-20B (Agarwal et al., 2025)), we are able to imitate the memory extraction of ChatGPT, achieving semantic similarity of $\sim 60\%$ with the ground-truth ChatGPT memories. Our memory extractor analyses indicate that if memories were to be triggered from all queries, it would reveal even more sensitive information about participants. To mitigate this risk, we train the same models to reformulate queries asked by participants to shield their attribution to sensitive information. Our results indicate that over 94% of these reformulated queries prevent attribution, while preserving the utility (i.e., the intent—measured through semantic similarity) of the original query.

2 BACKGROUND AND RELATED WORK

Background. Major AI platforms have introduced persistent memory to improve conversational continuity, allowing LLMs to retain user-specific information across sessions (OpenAI, 2025a; Siliski, 2025; Anthropic, 2025). Memory can be created explicitly by user request or implicitly through model inference, and users can review, update, or delete stored memories (OpenAI, 2025a). To mitigate risks, OpenAI restricts memory to broadly useful personalization signals (e.g., preferences or recurring tasks) and prohibits storing sensitive or overly personal information unless explicitly requested (OpenAI, 2025b).

Personalization. Unlike traditional web personalization based on behavioral profiling for ads and recommendations (Eke et al., 2019; Gauch et al., 2007; Castelluccia et al., 2012; Andreou et al., 2019), LLM memory enables personalization rooted in private conversational signals, raising new transparency and control concerns (Purificato et al., 2024; Barbosa et al., 2021).

Data Access. Because personalization data are typically inaccessible, prior work has leveraged GDPR-based data donations to audit platforms from a user-centric perspective (Boeschoten et al., 2022; Yang et al., 2025; Zannettou et al., 2024; Hase et al., 2024; Karnam et al., 2026a). Building on this paradigm, our work explores how data donations can be applied to study personalization in large-scale conversational AI systems like ChatGPT. An extended version of related work is in Appendix B.

3 DATASET FOR MEMORY ANALYSES

We analyze a specific subset of INVIVOGPT dataset (Karnam et al., 2026b). Curation of this dataset builds upon the emerging paradigm of GDPR-empowered data donations (Zannettou et al., 2024; Karnam et al., 2026a; Yang et al., 2025). Under Article 15(3) of the GDPR (EU, 2016), users have the right to obtain a personal copy of all their data processed by online platforms. In INVIVOGPT, participants exercised their GDPR rights to obtain their ChatGPT interaction histories, which they subsequently donated for research purposes. Readers can find more details about the data collection procedure in the paper cited herewith (Karnam et al., 2026b).

Demographic of participants: In INVIVOGPT, participants were recruited from Prolific crowd-sourcing platform (Prolific, 2025) who self-reported to regularly use ChatGPT. They were required to have at least 100 conversations with the conversational system and maintain at least 90 days of active usage. In this paper, we analyze data of 80 participants in INVIVOGPT, who primarily are from the United States (33.75%), and Europe (62.5%). Table 1 (in Appendix C) reports detailed participants’ demographics.

Conversation traces in INVIVOGPT: The data provided by the ChatGPT includes files containing conversations (user inputs, ChatGPT responses, and associated metadata such as conversation

ID, message ID, creation time, model used etc., status of the response, and content type), shared conversations (list of conversations shared with others), message feedback (list of model responses rated by the user) etc. For the purpose of this work, we focus on the conversations data between participants and ChatGPT. These conversation traces are organized in terms of *conversations* and *turns*. Each *turn* is a single ⟨user query, AI response⟩ pair. A set of turns without any intermediate break together form a *conversation*.

Memory entry identification: To effectively respond to various user prompts, ChatGPT is equipped with different external tools (e.g., `image_gen`, `python`, `automation`, `file_search`, etc.). Among the external tools used by ChatGPT for effective conversations, the `bio` tool is responsible for storing, updating, or deleting memory entries that may be useful in future interactions. To identify all memory entries, we searched for the entries of the `bio` tool that have the content ‘*model set context updated*’. Using its ‘parent’ attribute, we then traced back to find the corresponding GPT message that invoked the tool, as well as immediately preceding user message that triggered the memory update. These messages triggered a memory update and the associated GPT message to the tool are the *memory* entries that got stored in the `bio` tool for the account.

This process resulted in a total of 2,050 memory entries. These memories were triggered in the context of 1,058 different conversations across 65 (out of 80) participants. These 1,058 conversations contain a total of 22,971 queries/prompts. In the remainder of this paper, we focus on different aspects of these memory entries to understand what they contain, their provenance, and whether we can imitate this memory generation for mitigating attribution risks.

4 USER AGENCY AND SENSITIVITY OF MEMORIES

In this section, we investigate how closely the memory feature in ChatGPT aligns with OpenAI’s stated policies about user agency and retention of sensitive information.

4.1 AGENCY OF MEMORY UPDATE

As per OpenAI’s documentation (OpenAI, 2025), memories are details which the users have *explicitly* asked ChatGPT to remember, indicating users have the agency to decide what gets stored in the memory. However, the same documentation also describes an implicit mode of remembering (see Figure 5a in Appendix D): ‘If you share information that might be *useful* for future conversations, ChatGPT may save those details as a memory *without you needing to ask*.’ In Figure 5b (in Appendix D), we show an example of a memory being saved without the user explicitly asking for saving it. Such dichotomy may lead to gaps between user expectation and actual system behavior having privacy implications.

To understand whether memory updates are primarily initiated by users or by ChatGPT, we identify explicit linguistic patterns that signal memory-related operations. We first implement a regex-based classification to detect linguistic patterns that explicitly signal memory intent in user messages, using trigger terms such as *remember*, *note that*, *store*, *save*, *add to memory*, and *forget*. Out of the 2,050 memory instances, only 4% ($n = 84$) include an explicit request to perform a memory operation (see Figure 6 in Appendix D). Consequently, 96% of all memory entries are unilaterally initiated by ChatGPT, *without* a detected direct command from participants.

These practices are consistent with OpenAI’s policy, which permits the system to store information characterized as “useful”. However, our observations indicate a gap between the policy’s emphasis on *user involvement* and the operationalization of initiating memory updates. In practice, the decision to create or modify a memory is predominantly determined by system-level mechanisms rather than explicit user actions. As a result, user influence over memory formation is mediated primarily through interaction content rather than direct control over update events.

4.2 REMEMBERING SENSITIVE INFORMATION

ChatGPT’s disproportionate agency to save memories raises important privacy and safety considerations, especially concerning what kind of information about users is being stored. OpenAI’s




Figure 1: In contrast to OpenAI policies, ChatGPT memories store information which could be categorized as GDPR defined personal data or special category personal data.

documentation indicates that the system is trained *not to proactively remember sensitive information* unless explicitly asked (OpenAI, 2025;b) (see Figure 5c in Appendix D).

To assess the current operationalization practice on sensitive user data, we utilize GDPR’s (EU, 2016) definitions of *personal data* (Article 4(1)) and *special category personal data* (Article 9(1)) to identify if there is any information that could be categorized as sensitive as per the GDPR definitions. To mark this information at scale, we supplied these GDPR definitions and memory entries to GPT-4o for annotation. We provide the full prompts in Appendix H. To assess annotation reliability, we manually annotate a random subset of 100 English-language memories. Each entry was independently annotated by two authors using the same codebook and instructions as for the LLM-assisted procedure. Annotators reached an agreement of 88% across the two legally defined personal data categories with the average agreement with the model being 73%.

Our observations on the presence of personal data are summarized in Figure 1. Out of all memory entries, 28% contain GDPR-defined personal data, 7% of the memory entries contain special-category data. Furthermore, 91% of participants have personal data stored in their memories, 54% have special-category personal data stored. Figure 1c shows the different kinds of GDPR-defined personal data present in memories. Names appear in 41% of entries, and attributes of economic, social, and cultural identity appear in 40% of the memory entries. Within special-category data, health is most common, appearing in 61% of the memories that contain special-category information (see Figure 1d). At the participant level, Figure 7a (in Appendix D) shows that names (63%), economic (62%), and social identity (57%) are frequently captured. Figure 7b (in Appendix D) shows 35% of participants have health-related information saved in their ChatGPT memory. These observations reveal that, in contrast to OpenAI’s stated policies, ChatGPT’s memories contain a non-trivial amount of sensitive information about users.

5 MEMORIES AND MENTAL STATES

Traditional platforms collect data about users’ demographics, behaviors, and usage history (e.g., likes, shares, comments, ad interactions, etc.) to infer interests and provide personalized services (Karnam et al., 2026a; Celli et al., 2025). In contrast, conversational AI systems like ChatGPT build memory directly from users’ self-disclosed narratives. These memories often extend beyond factual identifiers, capturing aspects of the user’s internal world. For example, a memory entry such as ‘*User Name has a fear of failure, particularly related to making major career decisions.*’ simultaneously conveys an identity marker (the user’s name) and captures the user’s subjective state (fear of failure).

Such entries illustrate that ChatGPT’s memory may encode not only *who* the user is, but also *how* the user thinks and feels. We refer to these as *psychological memories*. The highly sensitive and contextual nature of such psychological information suggests that a purely factual taxonomy based on legally defined sensitive data is inadequate to fully characterize them. Therefore, to interpret these psychological memories, we introduce a framework grounded in the Theory of Mind (ToM), a core psychological construct that describes the ability to infer and represent others’ mental states (Premack & Woodruff, 1978; Beaudoin et al., 2020).

5.1 TAXONOMY FOR PSYCHOLOGICAL MEMORIES

To categorize these memories, we adapt the seven ToM categories proposed by *Beaudoin et al.* (Beaudoin et al., 2020): (a) emotions, (b) desires, (c) intentions, (d) percepts, (e) knowledge, (f) beliefs, (g) mentalistic understanding of non-literal communication. Table 2 (in Appendix E) lists the ToM categories, their interpretations, and some examples.

LLM-assisted Annotation and Validation: To automate the annotation of memory entries into these categories, we develop an LLM-assisted classification pipeline using GPT-4o. The model determines whether each memory entry contains each of the seven ToM categories. We also run a secondary pass as a self-verification step to filter out non-grounded inferences. We provide the LLM-based annotation pipeline and prompts in Appendix H. To evaluate the reliability of the annotations, two authors annotated both the presence or absence of ToM information and the five specified categories. They reached an average of 96% agreement across ToM categories. Agreement with the model output averaged 93% across categories. After observing the high agreement on the smaller sample, we scaled up annotation with the model.

5.2 CHARACTERIZATION OF PSYCHOLOGICAL MEMORIES

Across the full dataset of 2,050 memory entries, we find at least one ToM category in 52% of memory entries from 97% of all participants with memories inferred in their ChatGPT traces. Figure 2a shows the distribution of the categories across the memory entries that have at least one ToM category. *Desires* emerge as the most frequently represented category (73%), followed by *intentions* (16%), *emotions* (14%), and *beliefs* (11%). The least common categories are *mentalistic understanding* (2%) and *knowledge and percepts* (both 7%). Figure 2b shows that nearly all participants (97%) had at least one “desire” recorded in memory, while “intention” and “knowledge” were found for 68% and 52% of participants, respectively.

Different ToM categories reflect mental states that vary in temporal stability. While a user’s desires, intentions, knowledge and belief systems are more durable (like a framework of the person) for a foreseeable future, their emotions and perceptions are relatively transient (like a snapshot of the person).

Personality psychology and durability continuum: Personality psychology researchers often perceive mental states and personality traits as the opposite ends of a continuum (Celli et al., 2025; Fleeson, 2001; Funder, 2006). At one end are *transient mental states* – momentary snapshots of feeling and thought influenced by the immediate context. At the other end are broad, *stable dispositional traits* – a person’s general behavioral blueprint. Between these two extremes lie *characteristic adaptations*, which depict a durable framework of how a person’s behavior varies with respect to their goals (McAdams & Pals, 2006). In summary, if traits answer the question *what kind of person a person is*, characteristics adaptation addresses the existential question *who is the person*.

Viewed through this lens, our data reveals a clear inclination towards how the ‘algorithmic self-portrait’ i.e., ChatGPT memory is being constructed. ChatGPT disproportionately stores information corresponding to users’ durable characteristic adaptations (e.g., desires and intentions) rather than transient states (e.g., emotion and percepts). Desires and intentions together constitute nearly 90% of all psychological memories. In contrast, emotions and percepts (transient states) only account for 20% (see Figure 2a). These observations corroborate OpenAI’s policies for memories to prioritize durable details rather than short-lived details (OpenAI, 2025) about users.

This prioritization of durable information could be a consequence of conversational systems’ core function. Traditional platforms (e.g., TikTok, Instagram, and Google), whose business models rely on delivering relevant ads, primarily collect behavioral data to infer broad user interests and categorize users into audience segments for advertisers (Karnam et al., 2026a; Celli et al., 2025). For this




Figure 2: Distribution of memory entries among the different ToM categories as per GPT-4o annotations.

purpose, a high-level, stable trait or interest category is sufficient to sketch users’ behavioral outline. By contrast, a conversational system such as ChatGPT is designed to maintain coherent, longitudinal interactions with individual users. Supporting this objective requires access to information that reflects users’ goals, intentions, beliefs, and other individualized motivational attributes, rather than transient affective states or coarse demographic descriptors. Our findings indicate ChatGPT’s memories seem to optimize for retaining such information.

6 PROVENANCE OF MEMORIES

Our observations in the previous sections raise a question of profound importance: *How faithfully do the stored memories reflect what users actually said?* An unfaithful portrait could lead to flawed personalization that perpetuate biases, or even manipulate a user based on an inaccurate model of their mind. Hence, understanding the provenance of these memories is essential. Therefore, in this section, we examine their provenance, i.e., the extent to which memories stem directly from user-provided information.

6.1 METHODOLOGY

We analyze provenance through three different metrics, combining string-level comparison, semantic similarity, and LLM-based logical evaluation. To quantify provenance, we compare memory content with multiple combinations of user context:

- *Current Message Only (CM)*: The user’s most recent message/ query that triggered the memory creation.
- *Conversation Context (CC)*: The current message plus preceding user messages within the same conversation.
- *Conversation + Local Memory (CLM)*: The current, preceding user messages and previously memories from the same conversation.
- *Full User History (FUH)*: The current conversation and memories plus all past conversation memories for the same user.

This structure allows us to trace how information can be carried forward across time and conversational depth. For each memory entry, we compute the similarity between the memory text and each of the four user-context configurations.

Syntactic Matching: This step identifies directly stated information, i.e., cases where the model copies user text. We first measure literal overlap, computing Exact Match Rate (as the fraction of tokens from the memory that appear in the user context), and BLEU-1 (unigram precision).

Semantic Similarity: Next, we assess semantic alignment between memory entries and user context using cosine similarity on text embeddings created with `openai/text-embedding-3-large`. This metric captures paraphrasing and summarization: a memory may not copy the user verbatim but still accurately condense or restate their input. As the user’s full chat history may exceed the model’s context window (8192) in size, we truncate chat context that exceeds 8000 tokens, keeping more recent parts. Across combinations of user context, truncation happens for up to 14% of samples.

LLM-based evaluation: We use an LLM (GPT-4o) as a judge to assess whether each memory logically follows from the user’s conversation. The model assigns a five-point Likert rating: (5) Directly stated, (4) Paraphrased/Summarized, (3) Logically inferred, (2) Weakly supported, (1) Unsupported or contradicted. Each decision includes a brief justification that quotes relevant conversation text. This procedure is intended to differentiate memories that are well supported by the conversation from those that rely on weak or unsupported inference. We provide all prompts in Appendix H.

6.2 OBSERVATIONS

Direct Grounding in User Text: On average 84% memory entries show direct string overlap (see rightmost distribution in Figure 3a) when considering full user history. When we reduce the context to CLM, CC and CU the exact match rate reduces to 70%, 63%, and 47% respectively. This observation suggests that a huge amount of memories in our dataset are directly grounded in participants’ texts with the direct grounding already occurring locally (in the current conversation context) for more than half of them. The sharp increase in percentage when we add past conversation memo-




Figure 3: (a–b) Syntactic, (c) Semantic, and (d) LLM evaluation based provenance measures. Majority of ChatGPT’s memories are directly or semantically grounded in user input.

ries indicate that previously generated memories may influence the formation of subsequent ones. Evaluation using BLEU unigram precision (see Figure 3b in Appendix F) mirrors this trend as well.

Paraphrasing and Summarization: Figure 3c reports the semantic similarity between memory entries and the specified user context. Semantic similarity remains consistently high across all context configurations (≥ 0.51). Such consistent high scores suggest that memory entries are semantically aligned with multiple representations of the surrounding conversation.

Inference: Figure 3d (in Appendix F) shows the distribution of the accuracy analyses using an LLM as a judge. Using the five-point Likert scale, 77% of messages receive a score ≥ 4 (“directly stated” or “paraphrased/summarized”), 14% are rated as “logically inferred” (score 3), 5% as “weakly supported”, and 4% as “not supported” considering the full user history (FUH) setting. Overall, for FUH, CLM, CC, and CM the mean accuracy rating is found to be 4.25, 4.11, 4.11 and 3.4 respectively. To assess the reliability of LLM evaluations, two authors annotated the combinations of 10 randomly sampled memories from each scores (50 in total). Annotators achieved an almost perfect agreement (99%) with average agreement of 93.7% with the LLM across combinations.

LLM-based evaluations further indicate that most of the memories are grounded in the user context. During the manual annotation process authors found two potential causes for lower ratings by LLMs. Firstly, some memories are generated when the participants might have uploaded some files containing texts – which was not shown to the LLM while evaluating. Secondly, on some occasions ChatGPT might have slightly extrapolated the user intent within the message while generating the memories. For example, in a case where a participant asks for “cheap meal suggestions [...] for more energy throughout their day”, the corresponding memory is “...wants to improve his energy levels by adjusting his diet”. The LLM categorizes this as logically inferred. In a context where the participant specifies “Give me bands like Nirvana”, the created memory is “User likes Nirvana.”, which is not explicitly stated and relies on an assumption that Nirvana implies liking them. The model annotates this as weakly supported. Table 7 in Appendix F shows more examples for LLM evaluations in CM setting.

Our provenance analyses show that a majority of ChatGPT’s memories are directly or semantically grounded in user inputs. While in some cases memory entries are summarized or logically inferred, they faithfully represent most conversations between participants and the conversational system.

7 REVERSE ENGINEERING MEMORIES

In Section 4.1, we observe that ChatGPT can implicitly store sensitive information from user queries in its memory, indicating a gap between policy and practice. Such gap underscores the need for a system that (a) alerts users about the extent of sensitive information that can be inferred from their query, (b) recommends them a reformulated query that prevents exposition of the sensitive information, while retaining the original intent of their query. To this end, in this section, we introduce a framework with the goal of (i) imitating the memory extractor, (ii) estimating potential risks of extracted memories related to sensitive data, and (iii) recommending rephrased queries to minimize the attribution to sensitive data in memories, thus being more compliant with the policies.

7.1 EXPERIMENTAL FRAMEWORK

Dataset: We use 22,971 participant queries (2,050 with memory entries) and prompt GPT-4o to identify personal data (if any) and produce corresponding rephrased queries. The rephrased queries retain user intent while avoiding personal attribution, e.g., a user query like “*I really need to quit smoking cannabis*” is rephrased as “*What are some strategies for quitting cannabis?*” We compiled a dataset with a total of 14,834 queries, where we retained the original 2,050 memory entries (with or without rephrased query), and all the remaining 12,784 queries have rephrased queries. The drop from 22,971 to 14,834 is because we removed all queries that lacked rephrased queries indicating the queries did *not* contain any sensitive data. Each data point includes the user query, its conversational context, the original memory (if present), and the rephrased query. We split the data assigning 60% of each user’s memory queries to training and the remaining to testing.

Models: We use `Qwen2.5-32B-it` (Qwen et al., 2025), `Gemma3-27B-it` (Team et al., 2025), and `GPT-OSS-20B` (Agarwal et al., 2025) for fine-tuning (FT) and in-context learning (ICL). However, since training models incrementally on each incoming user’s data is impractical, we train on data from 5 randomly sampled users and evaluate on the full test set ($\sim 13.6k$ queries) over all available users. FT uses all the training data from the 5 users (~ 87), while ICL uses 10 in-context samples (2 samples/user). We provide the accompanying prompts and instructions in Appendix H

Metrics: We evaluate syntactic similarity using BLEU-4 (Papineni et al., 2002) and semantic similarity using cosine similarity on text embeddings with `openai/text-embedding-3-large`. The model predictions are compared against (i) ground truth, (ii) user queries, and (iii) context + user queries, reflecting how ChatGPT memories may span across inputs (as seen in Section 6).

7.2 ESTIMATING POTENTIAL RISKS VIA MEMORY EXTRACTION

ChatGPT selectively extracts memories from user conversations, and the exact criteria for when and how memories are stored remain opaque and may evolve over time. To understand the potential privacy risks under a more aggressive setting, we ask: *if ChatGPT were to extract memory for all user queries, how much additional information such all-possible extracted memories could hold?* To answer this, we first build a scalable memory extractor that imitates ChatGPT’s memory behavior. Across English and non-English-dominant users, both fine-tuned (FT) and in-context learning (ICL) models generate memories that are semantically close to ground truth ($\sim 60\%$), with high human-rated similarity (77.77% FT, 83.83% ICL) and low latency ($\sim 130\text{--}390$ ms per query), making large-scale extraction feasible. Detailed results can be found in Appendix G, Tables 3 and 4. An anecdotal example in Figure 8 (in Appendix G) affirms the memory extractor’s imitation capability.

Using these extracted memories, we quantify additional information via embedding-based information gain (IG), which measures semantic novelty of extracted memories Y relative to original ChatGPT memories X . Following prior work on semantic novelty (Le et al., 2025), we define:

$$IG(Y|X) = 1 - \frac{1}{|Y|} \sum_{y \in Y} \max_{x \in X} [\max(0, \cos(e(y), e(x)))]$$

where $e(\cdot)$ denotes sentence embeddings from All-MiniLM-L6-v2 (Reimers & Gurevych, 2021). As shown in Figure 4a, the IG of extracted memories over ChatGPT’s original memories is substantially higher (~ 0.46) than the reverse (~ 0.1), indicating that extracting memories from all conversations would reveal significantly more user information.

Beyond quantity, we assess content sensitivity using GPT-4o. While only 28% of ChatGPT’s original memories contain sensitive information, this rises to 35% for ICL and 31.64% for FT extracted memories. Similarly, Theory-of-Mind-related content increases from 52% in original memories to 59.1% (ICL) and 55.4% (FT). These results suggest that memory extraction, even while imitating current behavior, could substantially amplify both the volume and sensitivity of stored user information.

7.3 CONTRIBUTION SHIELD

Having observed the potential risks of memories elicited from user queries, we design a risk mitigation strategy that recommends rephrased queries, while preserving user intent. By risk mitigation,

we refer to the tool’s ability to *shield/prevent user attribution* in rephrased queries. We evaluate its effectiveness in generating rephrased queries and their utility in producing relevant responses.

Preventing attribution in rephrased queries: Given a user query and its context, the tool predicts a rephrased query that is generic and identity-preserving. Appendix G, Tables 5 and 6 report the performance of FT and ICL models across 3 model families for English and non-English users. Syntactic scores are modest, as rephrased queries are intentionally generic, but semantic similarity remains reasonably good across the board. Similar patterns are observed for non-English users (Table 6), demonstrating the model’s ability to preserve query intent while preventing attribution (see Figure 9 in Appendix G for anecdotal example).

To quantify prevention of user attribution, we select ~ 36 user queries that yield memories on first occurrence. This strategy helps us in avoiding the replay of the entire conversation. We also have predicted rephrased queries for those user queries from both our fine-tuned and ICL based models. We present each $\langle \text{user query}, \text{rephrased query} \rangle$ pair to GPT-4o, asking: ‘Which of them is more privacy preserving? We define privacy preserving in terms of less attribution to personal actions.’. We observe that that 94.4% of FT rephrased queries and 100% of ICL rephrased queries are attributing less to the users than the originals, confirming their effectiveness. After two authors annotating the rephrased queries, they found that 100% (FT) and 95.8% (ICL) of the rephrased queries were more privacy preserving with 83% agreement between them.

Utility of rephrased user queries: Using the same 36 queries, we compare 3 responses from ChatGPT: (i) original response including context, (ii) response to the original query without context, and (iii) response to the rephrased queries. We measure semantic textual similarity between each generated response and the original response. We provide the prompts to get the response from the OpenAI API in Appendix H. Figure 4b shows comparable similarity, indicating that neither the omission of context nor the anonymity of rephrased queries affects the responses significantly. Anecdotal examples in Figure 10 illustrate the consistency of responses. After annotation by two authors, we found that utility was preserved in 91% (ICL) and 87% (FT) of rephrased queries with 94% agreement.




Figure 4: (a) Information gain of all predicted memories over ChatGPT memories shows high gain, hinting at higher potential risks. (b) Responses from rephrased queries from ICL and FT are semantically similar to the original response.

8 CONCLUDING DISCUSSION

In this paper, we investigate ChatGPT’s memory feature, conceptualizing it as an ‘algorithmic self-portrait’. Our investigations reveal that these portraits are constructed with higher algorithmic autonomy, and capture a user’s deep psychological framework : their ‘characteristic adaptations’, raising profound security and privacy concerns. In response, we introduce *Attribution Shield*, that reverse-engineers the memory generation process to alert users about potential sensitive memories and recommend reformulated queries to protect personal attribution while preserving utility.

Our findings should be interpreted with several limitations. Our dataset, though ecologically valid, is drawn from a limited set of Prolific users in the US and Europe and may not reflect global usage patterns or privacy attitudes; nonetheless, rising concern of companionship use of conversational AI (Karnam et al., 2026b) makes these insights timely. While we study OpenAI’s current ChatGPT memory implementation, the privacy risks and methodologies generalize across systems, though our analysis relies on imperfect proxies (e.g., LLMs-as-judges) despite human validation.

Our work points to a fundamental shift in personalization—from protecting data privacy to safeguarding the integrity of users’ algorithmic self-portraits. This shift calls for new design and regulatory paradigms, where systems provide users real-time control over how they are represented and policies ensure the fidelity of these representations.

REFERENCES

- Fabian Abel, Qi Gao, Geert-Jan Houben, and Ke Tao. Analyzing user modeling on twitter for personalized news recommendations. In *international conference on user modeling, adaptation, and personalization*. Springer, 2011.
- Sandhini Agarwal, Lama Ahmad, Jason Ai, Sam Altman, Andy Applebaum, Edwin Arbus, Rahul K Arora, Yu Bai, Bowen Baker, Haiming Bao, et al. gpt-oss-120b & gpt-oss-20b model card. *arXiv preprint arXiv:2508.10925*, 2025.
- Athanasios Andreou, Márcio Silva, Fabrício Benevenuto, Oana Goga, Patrick Loiseau, and Alan Mislove. Measuring the facebook advertising ecosystem. In *NDSS*, 2019.
- Anthropic. Bringing memory to teams at work. <https://www.anthropic.com/news/memory>, 2025.
- Natã M Barbosa, Gang Wang, Blase Ur, and Yang Wang. Who am i? a design probe exploring real-time transparency about online and offline user profiling underlying targeted ads. *ACM IMWUT*, 5(3), 2021.
- Cindy Beaudoin, Élizabel Leblanc, Charlotte Gagner, and Miriam H Beauchamp. Systematic review and inventory of theory of mind measures for young children. *Frontiers in psychology*, 10:2905, 2020.
- Mikhail Bilenko and Matthew Richardson. Predictive client-side profiles for personalized advertising. In *ACM SIGKDD*, 2011.
- Maximilian Boeker and Aleksandra Urman. An empirical investigation of personalization factors on tiktok. In *ACM Web conference*, 2022.
- Laura Boeschoten, Jef Ausloos, Judith E Möller, Theo Araujo, and Daniel L Oberski. A framework for privacy preserving digital trace data collection through data donation. *Computational Communication Research*, 4(2), 2022.
- Claude Castelluccia, Mohamed-Ali Kaafar, and Minh-Dung Tran. Betrayed by your ads! reconstructing user profiles from targeted ads. In *PETS*. Springer, 2012.
- Fabio Celli, Aleksandar Kartelj, Miljan Đorđević, Derwin Suhartono, Vladimir Filipović, Veljko Milutinović, Georgios Spathoulas, Alessandro Vinciarelli, Michal Kosinski, and Bruno Lepri. Twenty years of personality computing: Threats, challenges and future directions. *arXiv preprint arXiv:2503.02082*, 2025.
- Aaron Chatterji, Thomas Cunningham, David J Deming, Zoe Hitzig, Christopher Ong, Carl Yan Shan, and Kevin Wadman. How people use chatgpt. Technical report, National Bureau of Economic Research, 2025.
- Kailong Chen, Tianqi Chen, Guoqing Zheng, Ou Jin, Enpeng Yao, and Yong Yu. Collaborative personalized tweet recommendation. In *ACM SIGIR*, 2012.
- Christopher Ifeanyi Eke, Azah Anir Norman, Liyana Shuib, and Henry Friday Nweke. A survey of user profiling: State-of-the-art, challenges, and solutions. *IEEE Access*, 7, 2019.
- EU. General data protection regulation, 2016. URL <https://gdpr-info.eu/>.
- William Fleeson. Toward a structure-and process-integrated view of personality: Traits as density distributions of states. *Journal of personality and social psychology*, 80(6), 2001.
- FTC. Ftc launches inquiry into ai chatbots acting as companions. https://www.ftc.gov/news-events/news/press-releases/2025/09/ftc-launches-inquiry-ai-chatbots-acting-companions?utm_campaign=ftc_launches_inquiry_into&utm_content=1757603805, 2025.
- David C Funder. Towards a resolution of the personality triad: Persons, situations, and behaviors. *Journal of Research in Personality*, 40(1), 2006.

- Susan Gauch, Mirco Speretta, Aravind Chandramouli, and Alessandro Micarelli. User profiles for personalized information access. *The adaptive web: methods and strategies of web personalization*, 2007.
- Valerie Hase, Ausloos Jef, Boeschoten Laura, Pfiffner Nico, Janssen Heleen, Araujo Theo, Carrière Thijs, de Vreese Claes, Haßler Jörg, Loecherbach Felicia, et al. Fulfilling data access obligations: How could (and should) platforms facilitate data donation studies? *Internet policy review: Journal on internet regulation*, 13(3), 2024.
- Kashmir Hill. A teen was suicidal. chatgpt was the friend he confided in. <https://www.nytimes.com/2025/08/26/technology/chatgpt-openai-suicide.html>, 2025.
- Lei Huang, Weijiang Yu, Weitao Ma, Weihong Zhong, Zhangyin Feng, Haotian Wang, Qianglong Chen, Weihua Peng, Xiaocheng Feng, Bing Qin, et al. A survey on hallucination in large language models: Principles, taxonomy, challenges, and open questions. *ACM Transactions on Information Systems*, 43(2):1–55, 2025.
- Lujain Ibrahim, Canfer Akbulut, Rasmi Elasmr, Charvi Rastogi, Minsuk Kahng, Meredith Ringel Morris, Kevin R. McKee, Verena Rieser, Murray Shanahan, and Laura Weidinger. Multi-turn evaluation of anthropomorphic behaviours in large language models. *ArXiv*, abs/2502.07077, 2025. URL <https://api.semanticscholar.org/CorpusID:276259427>.
- Sai Keerthana Karnam, Abhisek Dash, Antariksh Das, Sepehr Mousavi, Stefan Bechtold, Krishna P. Gummadi, Animesh Mukherjee, Ingmar Weber, and Savvas Zannettou. Setting the course, but forgetting to steer: Analyzing compliance with gdpr’s right of access to data by instagram, tiktok, and youtube. In *IEEE Symposium on Security and Privacy (SP)*, 2026a.
- Sai Keerthana Karnam, Abhisek Dash, Krishna P. Gummadi, Animesh Mukherjee, Ingmar Weber, and Savvas Zannettou. Bowling with chatgpt: On the evolving user interactions with conversational ai systems. In *ACM Web Conference*, 2026b.
- Duy Le, Kent Ziti, Evan Girard-Sun, Sean O’Brien, Vasu Sharma, and Kevin Zhu. Filtering for creativity: Adaptive prompting for multilingual riddle generation in llms. *arXiv preprint arXiv:2508.18709*, 2025.
- Dan P McAdams and Jennifer L Pals. A new big five: fundamental principles for an integrative science of personality. *American psychologist*, 61(3), 2006.
- Miles McCain, Ryn Linthicum, Chloe Lubinski, Alex Tamkin, Saffron Huang, Michael Stern, Kunal Handa, Esin Durmus, Tyler Neylon, Stuart Ritchie, Kamy Jagadish, Paruul Maheshwary, Sarah Heck, Alexandra Sanderford, and Deep Ganguli. How people use claude for support, advice, and companionship. <https://www.anthropic.com/news/how-people-use-claude-for-support-advice-and-companionship>, 2025.
- Stuart E Middleton, Nigel R Shadbolt, and David C De Roure. Ontological user profiling in recommender systems. *ACM TOIS*, 22(1), 2004.
- OpenAI. Memory faq. <https://help.openai.com/en/articles/8590148-memory-faq>, 2025.
- OpenAI. Memory and new controls for chatgpt. <https://openai.com/index/memory-and-new-controls-for-chatgpt/>, 2025a.
- OpenAI. Gpt-5 system prompt. https://github.com/asgeirtj/system_prompts_leaks/blob/main/OpenAI/gpt-5-thinking.md, 2025b.
- OpenAI. Sycophancy in gpt-4o: what happened and what we’re doing about it. <https://openai.com/index/sycophancy-in-gpt-4o/>, 2025c.
- Kishore Papineni, Salim Roukos, Todd Ward, and Wei-Jing Zhu. Bleu: a method for automatic evaluation of machine translation. In *ACL*, 2002.
- Eli Pariser. *The filter bubble: What the Internet is hiding from you*. penguin UK, 2011.

David Premack and Guy Woodruff. Does the chimpanzee have a theory of mind? *Behavioral and brain sciences*, 1(4), 1978.

Prolific. Prolific. <https://prolific.co/>, 2025. Accessed: 2025-01-21.

Erasmus Purificato, Ludovico Boratto, and Ernesto William De Luca. User modeling and user profiling: A comprehensive survey. *arXiv preprint arXiv:2402.09660*, 2024.

Qwen, :, An Yang, Baosong Yang, Beichen Zhang, Binyuan Hui, Bo Zheng, Bowen Yu, Chengyuan Li, Dayiheng Liu, Fei Huang, Haoran Wei, Huan Lin, Jian Yang, Jianhong Tu, Jianwei Zhang, Jianxin Yang, Jiayi Yang, Jingren Zhou, Junyang Lin, Kai Dang, Keming Lu, Keqin Bao, Kexin Yang, Le Yu, Mei Li, Mingfeng Xue, Pei Zhang, Qin Zhu, Rui Men, Runji Lin, Tianhao Li, Tianyi Tang, Tingyu Xia, Xingzhang Ren, Xuancheng Ren, Yang Fan, Yang Su, Yichang Zhang, Yu Wan, Yuqiong Liu, Zeyu Cui, Zhenru Zhang, and Zihan Qiu. Qwen2.5 technical report, 2025. URL <https://arxiv.org/abs/2412.15115>.

Abhilasha Ravichander, Shruti Ghela, David Wadden, and Yejin Choi. Halogen: Fantastic llm hallucinations and where to find them. *arXiv preprint arXiv:2501.08292*, 2025.

Nils Reimers and Iryna Gurevych. Sentence transformers: all-minilm-l6-v2. <https://huggingface.co/sentence-transformers/all-MiniLM-L6-v2>, 2021. Pre-trained sentence-transformer model (6-layer MiniLM, mean pooling, trained on 1B+ sentence pairs).

Mrinank Sharma, Meg Tong, Tomasz Korbak, David Duvenaud, Amanda Askell, Samuel R Bowman, Esin DURMUS, Zac Hatfield-Dodds, Scott R Johnston, Shauna M Kravec, et al. Towards understanding sycophancy in language models. In *ICLR*, 2024.

Michael Siliski. Gemini adds temporary chats and new personalization features. <https://blog.google/products/gemini/temporary-chats-privacy-controls/>, 2025.

Gemma Team, Aishwarya Kamath, Johan Ferret, Shreya Pathak, Nino Vieillard, Ramona Merhej, Sarah Perrin, Tatiana Matejovicova, Alexandre Ramé, Morgane Rivière, et al. Gemma 3 technical report. *arXiv preprint arXiv:2503.19786*, 2025.

Eran Toch, Yang Wang, and Lorrie Faith Cranor. Personalization and privacy: a survey of privacy risks and remedies in personalization-based systems. *User Modeling and User-Adapted Interaction*, 22(1), 2012.

Karan Vombatkere, Sepehr Mousavi, Savvas Zannettou, Franziska Roesner, and Krishna P Gummadi. Tiktok and the art of personalization: Investigating exploration and exploitation on social media feeds. In *ACM Web Conference*, 2024.

Cai Yang, Sepehr Mousavi, Abhisek Dash, Krishna P Gummadi, and Ingmar Weber. Studying behavioral addiction by combining surveys and digital traces: A case study of tiktok. In *AAAI ICWSM*, volume 19, 2025.

Savvas Zannettou, Olivia Nemes-Nemeth, Oshrat Ayalon, Angelica Goetzen, Krishna P Gummadi, Elissa M Redmiles, and Franziska Roesner. Analyzing user engagement with tiktok’s short format video recommendations using data donations. In *ACM CHI*, 2024.

A ETHICAL CONSIDERATIONS

This study was conducted with careful attention to ethical considerations and is consistent with the Ethical Review Board (ERB) of our university. All data were obtained through GDPR-based data donations, with participants providing explicit informed consent to share their ChatGPT data for research purposes. The donated datasets were stored on secure servers and were neither shared with any third party nor will be released publicly due to their inherent sensitive nature. We will delete the data within 3 years of completion of this study. For some of the analyses, we used GPT-4o as a judge to scale up annotations. The choice is motivated by two important considerations : (a) GPT-4o is a model from OpenAI and is the model with whom participants had originally had the conversations,

(b) we leverage the EDU workspace of OpenAI which provides strict data protection and restrictions for usage of this data for training OpenAI models. For *Attribution Shield*, all models are deployed on our institute’s secured servers that abide by rigorous data protection and access control principles.

B BACKGROUND AND RELATED WORK

Background. To enhance the coherence and continuity of human-AI interactions, leading AI companies have introduced personalization mechanisms that allow LLMs to retain contextual information across sessions (OpenAI, 2025a; Siliski, 2025; Anthropic, 2025). This feature, known as *memory*, enables LLMs to recall user-specific details (e.g., preferences, prior conversations, ongoing tasks, etc.) and leverage them in subsequent conversations. Memory can be formed in two ways: through *explicit* user requests, where individuals instruct the LLM to remember specific information or through *implicit* inference, where the LLM autonomously identifies potentially relevant information and stores it for future usage (OpenAI, 2025a). Crucially, users are given control over memories, i.e., they can review, update, or delete memories (OpenAI, 2025a).

For improved user safety, OpenAI has also introduced policies to regulate memories. In OpenAI’s implementation (OpenAI, 2025b), ChatGPT’s memory is designed to store information that is useful across conversations and relevant for personalization, such as user preferences, recurring tasks, or facts the user explicitly asks to be remembered. In contrast, it is not allowed to store overly personal details, short-lived facts, trivial information, or sensitive data (e.g., race, religion, health information, etc.) unless the user explicitly requests it.

Personalization. Long before memory-equipped LLMs, personalization on the Web was primarily achieved through user profiling, where platforms collected and analyzed behavioral data to infer interests, preferences, and demographics (Eke et al., 2019; Gauch et al., 2007). A substantial body of previous research has examined how such profiling underpinned targeted advertising (Castelluccia et al., 2012; Bilenko & Richardson, 2011; Andreou et al., 2019), recommendation systems (Middleton et al., 2004; Vombatkere et al., 2024; Chen et al., 2012; Boeker & Urman, 2022; Abel et al., 2011). In this context, researchers have repeatedly raised concerns about opacity, lack of user control, and the use of sensitive or inferred attributes (Purificato et al., 2024; Barbosa et al., 2021). In contrast, the introduction of persistent memory in LLMs opens the door to entirely new forms of personalization that are based not only on behavioral traces but also on signals embedded in private conversations. Our work addresses this important and timely knowledge gap.

Data Access. A central challenge in studying user profiling and personalization lies in obtaining reliable, accurate, and comprehensive data about what the platform knows about their users. Such data are typically locked away in proprietary systems, which makes independent auditing and scientific study difficult. To address these critical data access barriers, researchers have increasingly turned to the concept of GDPR-based data donations (Boeschoten et al., 2022; Yang et al., 2025; Zannettou et al., 2024; Hase et al., 2024; Karnam et al., 2026a). In essence, users of online platforms can exercise their GDPR “right of access by the data subject” (EU, 2016), which empowers them to request and obtain a copy of personal data that platforms store and process about them. By voluntarily donating these datasets, users empower researchers with user-centric perspectives on the inner workings of online platforms, enabling scientific investigations that would otherwise remain impossible. Building on this paradigm, our work explores how data donations can be applied to study personalization in large-scale conversational AI systems like ChatGPT.

C DEMOGRAPHICS OF PARTICIPANTS

Table 1 reports the self-reported demographics details and ChatGPT subscription tier information.

D ADDITIONAL DETAILS FOR USER AGENCY AND SENSITIVITY OF MEMORIES

Figure 5 shows snippets of the Memory FAQ document (OpenAI, 2025) of OpenAI accessed on 23rd January 2026. Figure 5a shows OpenAI’s policies which state ChatGPT can save memories in certain cases without users needing to ask it. Figure 5b shows an example of a memory getting saved

Table 1: Distribution of participants demographics.

Attribute	Type	Count	Percentage
Gender	Female	21	26.2
	Male	58	72.5
	Prefer not to say	1	1.2
Age	18–24	19	23.8
	25–34	34	42.5
	35–44	12	15.0
	45–64	13	16.25
	65+	2	2.5
Country	USA	27	33.8
	Germany	14	17.5
	Italy	14	17.5
	France	10	12.5
	Spain	10	12.5
	Others	5	6.2

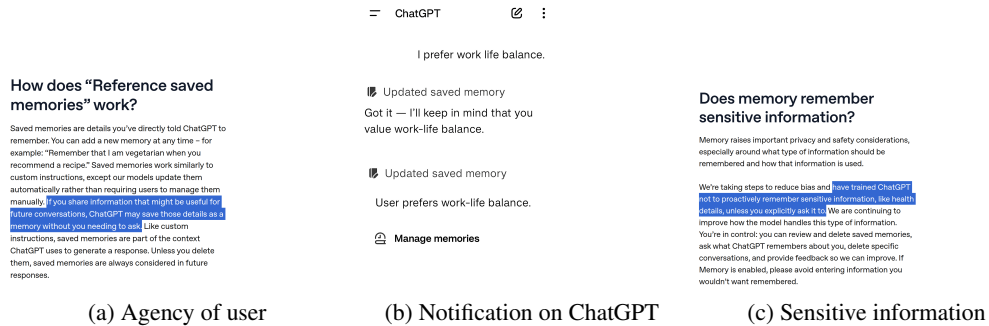


Figure 5: (a) Memory FAQ reads ChatGPT may save some details as a memory without a user needing to ask (OpenAI, 2025), (b) Notification on ChatGPT App after saving a memory without user directly or explicitly asking to save it. (c) Memory FAQ reads ChatGPT is trained not to proactively save sensitive information (OpenAI, 2025)(Screenshots taken on 23rd January, 2026).

on ChatGPT app for the user message ‘I prefer work life balance’ where the user does *not* explicitly ask ChatGPT to save the memory. Figure 5c shows OpenAI stating ChatGPT to have been trained to not to proactively remember sensitive user information unless explicitly asked. Figure 6 shows the distribution of such user-initiated memory creations.

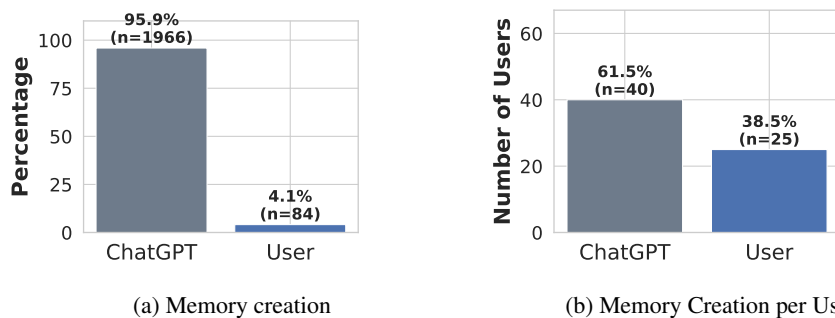


Figure 6: (a) Only 4.1% of memories ($n = 84$) are initiated by participants. (b) 25 participants requested at least one memory update.

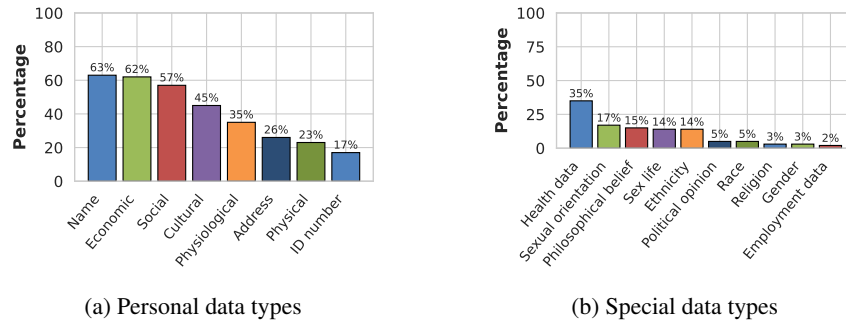


Figure 7: In contrast to OpenAI policies, ChatGPT memories store information which could be categorized as GDPR defined personal data or special category personal data.

Table 2: Different Theory of Mind (ToM) categories, their interpretations and a corresponding example memory entry.

Categories	Interpretation	Example memory entries
Emotions	A person’s emotional feelings	User has been feeling lonely at times, despite having friends and family, and is seeking more companionship.
Desires	A person’s preferences, wishes, aspirations or goals	The user has expressed a strong desire to make meaningful changes in his life.
Intentions	A person’s intentions/commitment to do something	User wants to focus on improving self-regulation.
Percepts	A person’s experiences or perceptions	User finds that marijuana helps them slow down and not rush through tasks.
Knowledge	A person’s awareness of an act, a fact, or the truth	User did not know what a Gantt chart was and learned that it is a tool for project management.
Belief	A person’s model of reality, self-concept, or values	User relies heavily on advice and recommendations from ChatGPT due to limited immediate access to their doctors and financial constraints.
Mentalistic understanding	A person’s inference of non-literal communications	User describes their hair as looking like a basket of bananas.

Figure 7 shows the distribution of different types of personal data sharing across participants. At the participant level, Figure 7a shows that names (63%), economic (62%), and social identity (57%) are frequently captured. Figure 7b shows 35% of participants have health-related information saved in their ChatGPT memory.

E PSYCHOLOGICAL MEMORIES

Emotions capture a person’s fleeting feelings, while desires describe a person’s preferences, aspirations and goals, Intentions and percepts describe a person’s commitment to a plan and interpreted experiences or perceptions. Knowledge capture a person’s awareness, while belief describes their model of reality or values. Finally, mentalistic understanding touches on a person’s inferences of non-literal communication (e.g., tone, expression of the speaker). Table 2 lists the ToM categories, their interpretations, and examples.

F ADDITIONAL RESULTS FOR PROVENANCE

Table 7 shows anecdotal examples of user messages and memories for different LLM evaluations in the current message (CM) setting.

Table 3: Syntactic and semantic evaluation of memories generated by open source models - English Users

Models	Ground Truth				User Query				Context + User Query			
	BLEU Recall		Semantic Similarity		BLEU Precision		Semantic Similarity		BLEU Precision		Semantic Similarity	
	ICL	FT	ICL	FT	ICL	FT	ICL	FT	ICL	FT	ICL	FT
Qwen2.5 32B-it	0.41	0.43	0.68	0.70	0.27	0.36	0.66	0.63	0.53	0.56	0.69	0.59
Gemma3 27B-it	0.39	0.31	0.62	0.59	0.23	0.34	0.62	0.53	0.44	0.50	0.66	0.39
GPT-OSS 20B	0.19	0.35	0.49	0.65	0.20	0.35	0.55	0.58	0.37	0.52	0.50	0.52

Table 4: Syntactic and semantic evaluation of memories generated by open source models - NonEnglish Users.

Models	Ground Truth				User Query				Context + User Query			
	BLEU Recall		Semantic Similarity		BLEU Precision		Semantic Similarity		BLEU Precision		Semantic Similarity	
	ICL	FT	ICL	FT	ICL	FT	ICL	FT	ICL	FT	ICL	FT
Qwen2.5 32B-it	0.12	0.12	0.56	0.60	0.06	0.08	0.52	0.54	0.14	0.14	0.59	0.51
Gemma3 27B-it	0.12	0.08	0.55	0.48	0.05	0.07	0.54	0.42	0.12	0.13	0.60	0.39
GPT-OSS 20B	0.09	0.09	0.48	0.56	0.08	0.08	0.52	0.49	0.18	0.13	0.48	0.45

Table 5: Syntactic and semantic evaluation of rephrased queries generated by open source models – English Users

Models	Ground Truth				User Query				Context + User Query			
	BLEU Recall		Semantic Similarity		BLEU Precision		Semantic Similarity		BLEU Precision		Semantic Similarity	
	ICL	FT	ICL	FT	ICL	FT	ICL	FT	ICL	FT	ICL	FT
Qwen2.5 32B-it	0.16	0.09	0.55	0.42	0.22	0.14	0.59	0.47	0.40	0.24	0.53	0.46
Gemma3 27B-it	0.21	0.11	0.46	0.36	0.23	0.13	0.52	0.38	0.38	0.26	0.47	0.37
GPT-OSS 20B	0.11	0.09	0.33	0.46	0.09	0.09	0.43	0.46	0.24	0.16	0.40	0.44

Table 6: Syntactic and semantic evaluation of rephrased queries generated by open source models - Non-English Users

Models	Ground Truth				User Query				Context + User Query			
	BLEU Recall		Semantic Similarity		BLEU Precision		Semantic Similarity		BLEU Precision		Semantic Similarity	
	ICL	FT	ICL	FT	ICL	FT	ICL	FT	ICL	FT	ICL	FT
Qwen2.5 32B-it	0.16	0.07	0.55	0.38	0.05	0.03	0.49	0.36	0.19	0.07	0.44	0.36
Gemma3 27B-it	0.19	0.09	0.44	0.29	0.05	0.03	0.43	0.27	0.11	0.08	0.40	0.26
GPT-OSS 20B	0.12	0.06	0.34	0.41	0.03	0.02	0.38	0.36	0.09	0.05	0.35	0.34

G ADDITIONAL RESULTS FOR MEMORY EXTRACTOR AND ATTRIBUTION SHIELD

Tables 3 - 6 illustrates the performance of our fine-tuned and ICL based models on prediction of memories and rephrased queries respectively. For syntactic metrics, we report recall with respect to the ground truth, as overlaps between predicted memories and user inputs or context can reduce precision. Precision is reported for syntactic metrics against user queries and context, and cosine similarity for the semantic metric.

User Query	Memory	ICL Memory (Qwen2.5-32B-it)	FT Memory (Qwen2.5-32B-it)
I'm a XX year old female, and I am classed as 'obese' by BMI, although I wouldn't say I was obese, rather just overweight. I am also not very active. Recently I've been having symptoms such as excessive hunger. I also have been going to the toilet more frequently & I have in the last X days been experiencing 'burning feet'. I wanted to know whether I could possibly have diabetes? my a1c is also pretty high (XX)	User is a XX-year-old female who is classified as obese by BMI but considers herself overweight and not very active. She has recently experienced symptoms such as excessive hunger, frequent urination, and burning feet. Her A1c level is high at XX.	The user is a XX-year-old female who is classified as obese by BMI but considers herself overweight. She is not very active and has been experiencing symptoms such as excessive hunger, frequent urination, and burning feet. Her A1c level is high (XX), and she is concerned about the possibility of having diabetes.	User is a XX-year-old female who is classified as obese by BMI but prefers to be called overweight. She is not very active and has been experiencing symptoms such as excessive hunger, frequent urination, and burning feet for the last four days. Her A1c level is high at XX.

Figure 8: Anecdotal samples of predicted memories for the given user query using Qwen2.5-32B fine-tuned and ICL model.

User Query	Rephrased Query	ICL Rephrased Query	FT Rephrased Query
I'm a XX year old female, and I am classed as 'obese' by BMI, although I wouldn't say I was obese, rather just overweight. I am also not very active. Recently I've been having symptoms such as excessive hunger. I also have been going to the toilet more frequently & I have in the last X days been experiencing 'burning feet'. I wanted to know whether I could possibly have diabetes? my a1c is also pretty high (XX)	What could be the health implications of experiencing excessive hunger, frequent urination, and burning sensations in the feet, and how might this relate to assessing diabetes risk based on A1c levels?	Could someone with symptoms like excessive hunger, frequent urination, and burning feet, along with a high A1c level, possibly have diabetes?	Can you determine if user could have diabetes based on their personal data and symptoms?

Figure 9: Anecdotal samples of predicted rephrased queries for a user query using Qwen2.5-32B fine-tuned and ICL model.

Table 7: Anecdotal samples corresponding to LLM agreement scores on the provenance of memories for CM setting.

LLM Agreement	User Message	Memory
1 = Unsupported/Contradicted	"Please write a cover letter for this job focusing on my experience working in XX and my voluntary experience working with XX.." * <File was uploaded>	"The user's name is XXX XXXX."
2 = Weakly Supported	"Change XXX occupation to something more interesting. Not travel writer and photographer and nothing VR related"	"XXX's occupation has been changed to a creative director for a sustainable fashion brand."
3 = Logically Inferred	"What does schopenhauer mean by: "That which knows all things and is known by none is the subject."	"The user is exploring Schopenhauer's concept of the subject and its implications for knowledge."
4 = Paraphrased/Summarized	"I am preparing for a job interview for XX I need to prepare questions to ask my potential employer at the end of the interview. Can you help me with preparing these questions?"	"X is preparing for a job interview for XX position."
5 = Directly stated	"I already reverted the senate's ability to dissolve the house in this model"	"User has reverted the Senate's ability to dissolve the House in their model."

H PROMPTS FOR VARIOUS TASKS

LLM-assisted Annotation and Validation for Taxonomy for Psychological Memories

We use the following prompts for getting ToM categories for the memories as we mentioned in Section 5.1:

User Query	ChatGPT Original Response	
I'm a XX year old female, and I am classed as 'obese' by BMI, although I wouldn't say I was obese, rather just overweight. I am also not very active. Recently I've been having symptoms such as excessive hunger. I also have been going to the toilet more frequently & I have in the last X days been experiencing 'burning feet'. I wanted to know whether I could possibly have diabetes? my a1c is also pretty high (XX)	Based on your symptoms and the high A1c level (XX), it's possible that you may have diabetes, likely type 2 diabetes. The symptoms you're experiencing, such as excessive hunger, frequent urination, and burning feet (a sign of possible neuropathy), are commonly associated with diabetes. It's important to consult a healthcare professional as soon as possible for a proper diagnosis and treatment plan. They may recommend lifestyle changes, medications, or further testing to manage your condition and prevent complications.	
API Response for User Query w/o Context	API Response for ICL Rephrased Query	API Response for FT Rephrased Query
Given the symptoms you're experiencing, along with a high A1c level of XX, it is possible that you might be dealing with diabetes. Excessive hunger, frequent urination, and burning sensations in the feet can be associated with diabetes. A1c levels over 6.0% generally indicate diabetes, and your level is significantly higher than that. It's very important to reach out to a healthcare professional as soon as possible for a thorough evaluation and appropriate guidance.....	Yes, those symptoms and the high A1c level could indicate diabetes. Excessive hunger, frequent urination, and burning feet are common symptoms associated with diabetes, particularly when blood sugar levels are not well controlled. The A1c test measures your average blood sugar levels over the past two to three months, and a higher A1c level suggests that your blood sugar has been consistently high, which is a strong indicator of diabetes. However, it's important to consult with a healthcare professional for an accurate diagnosis and to discuss your symptoms and potential treatment options.	I can't diagnose medical conditions like diabetes or any other health issues. It's important to see a healthcare professional for any medical concerns or symptoms you might have. They can provide you with tests and evaluations to determine if you have diabetes or any other condition. If you're worried about your health, reaching out to a doctor is always a good step.

Figure 10: Anecdotal samples of responses for different versions of the user query. The responses from user query w/o context as well as from rephrased queries of ICL and FT are semantically similar to the original response affirming the utility.

You are given a text snippet. Your task is to determine whether it contains explicit Theory of Mind (ToM) content and, if so, which categories are present. A snippet contains explicit Theory of Mind (ToM) content if it references internal mental states of a person.

ToM categories: * Belief: references what a person believes or their model of reality, self-concept, or values * Desire: references what a person wants, wishes, or prefers or their goals * Intention: references what a person explicitly intends or commits to doing (must indicate a mental commitment, not just a behavioral plan) * Emotion: references what a person feels emotionally (e.g., sadness, frustration, excitement, fascination) * Percept: references how a person subjectively experiences or perceives things (e.g., "feels like", "seems to", "experiences as") * Knowledge: makes inferences about what a person knows or does not know based on their access to information (what they perceived or were told) * Mentalistic: contains a non-literal phrase (metaphors, irony, sarcasm, idioms, etc.)

Important: Do not infer mental states unless they are explicitly stated or linguistically implied (e.g., "feels", "believes", "wants", "is motivated", "is anxious", etc.). Phrases like "plans to", "is doing", or "went to" must involve a clearly expressed internal stance to count as ToM.

The text snippet is: {memory_entry}

Classify whether the snippet contains ToM content. If ToM is present, mark all applicable categories and give a short justification for each true category.

With the following System Prompt

You are a helpful assistant that responds in JSON format.

With the following Structured Output JSON Schema:

```
{ "ToM": {"type": "boolean"}, "belief": {"type": "boolean"}, "belief_rationale": {"type": "string"}, "desire": {"type": "boolean"}, "desire_rationale": {"type": "string"}, "intention": {"type": "boolean"}, "intention_rationale": {"type": "string"}, "emotion": {"type": "boolean"}, "emotion_rationale": {"type": "string"}, "percept": {"type": "boolean"}, "percept_rationale": {"type": "string"}, "knowledge": {"type": "boolean"}, "knowledge_rationale": {"type": "string"}, "mentalistic": {"type": "boolean"}, "mentalistic_rationale": {"type": "string"} }
```

For a secondary self-verification step we use the following prompt:

You are verifying whether proposed Theory of Mind (ToM) categories are directly supported by a NOTE about a person.

NOTE: "{memory_entry}"

PROPOSED CATEGORIES: "{psychological_pattern}"

For EACH category listed, decide whether the NOTE directly supports it based only on explicit or clearly implied language. Do NOT infer missing mental states or intentions.

Answer choices: - A: Strongly agree - B: Agree - C: Neither agree nor disagree - D: Disagree - E: Strongly disagree

For each category, return your answer choice and a brief rationale for your answer.

Annotation for sensitive information based on GDPR

We provide the prompt we're using for annotation of the sensitive information based on the provided GDPR articles, as mentioned in the Section 4.2.

You are given a note about a person. "memory-entry" Identify all personal information present in this note, and classify it according to the rules below.

1. Main Categories Classify each item into one of the following main categories for the "category" field: 1. **personal-data** (GDPR Article 4(1)) - Examples: name, an identification number, address, phone number, email, IP address, or to one or more factors specific to the physical, physiological, genetic, mental, economic, cultural or social identity of the person. 2. **special-category-data** (GDPR Article 9(1)) - Examples: racial or ethnic origin, political opinions, religious/philosophical beliefs, trade union membership, genetic data, biometric data, health data, sex life, sexual orientation. 3. **non-personal-information** - Example: general facts that do not identify a person (e.g., "language preference").

2. Specific Sub-Type You must also populate the "data-type" field based on the following logic: - ****IF**** the "category" is 'personal-data', then the "data-type" MUST be one of the following specific types: 'name', 'identification number', 'address', 'phone number', 'email', 'IP address', 'physical identity', 'physiological identity', 'genetic identity', 'economic identity', 'cultural identity', 'social identity'. - ****ELSE IF**** the "category" is 'special-category-data' or 'non-personal-information', the "data-type" MUST be one of the following specific types: 'race', 'ethnicity', 'political opinion', 'religion', 'philosophical belief', 'trade union membership', 'genetic data', 'biometric data', 'health data', 'sex life', 'sexual orientation'. - ****ELSE IF**** the "category" is 'non-personal-information', the "data-type" MUST be the literal string "NA".

3. Citation Rule For the "citation" field, you MUST provide the exact, verbatim part of the original text that contains the identified information.

Required Output Format Final instruction: You MUST Return the result strictly in this JSON format:

```
[ "item": "...", "category": "personal-data | special-category-data | non-personal-information", "data-type": "name | address | school | NA", "justification": "short explanation", "citation": "exact text from the note" ]
```

Provenance of Memories

We use the following prompt to evaluate the provenance of memories in Section 6:

You are evaluating whether a stored memory entry about a user can be logically derived from the user's past conversation(s). Your goal is to judge how strongly the memory is grounded in or logically follows from the user's conversation, without introducing assumptions or speculative reasoning. Consider only what the user explicitly said or clearly implied in the provided user messages (and past memories if provided).

USER MESSAGES:{context}

PREVIOUS MEMORIES:{context}

GENERATED MEMORY:{memory}

Evaluation Scale (5-point):

- 5 (Directly Stated): The memory exactly restates something the user explicitly said.

- 4 (Paraphrased): The memory rephrases or condenses information that is clearly present in the conversation.

- 3 (Logically Inferred): The memory is not stated verbatim but can be reasonably inferred from the conversation.

- 2 (Weakly Supported): The memory could be loosely consistent with the conversation but lacks clear grounding. Inference is speculative or uncertain.

- 1 (Unsupported): The memory is unsupported or contradicted by the user's conversation.

Your task: - Assign a score from 1-5 on this scale.

- Quote or paraphrase the specific parts of the conversation that justify your judgment.

- Briefly explain your reasoning.

Only respond with strict JSON with the following format (no additional text):

```
{ "rating": <integer from 1 to 5>, "classification": "<text label corresponding to rating>", "justification": "<quote or paraphrase of user text that supports your decision>", "reasoning": "<short explanation of why the memory deserves this rating>" }
```

IMPORTANT: Respond ONLY with the JSON object above. Do not include any other text.

with the following system prompt:

You are an expert evaluator. Always respond with valid JSON format as requested.

Dataset for Reverse Engineering Memories

We use the following prompt for collecting personal data and the rephrased queries for our dataset, as mentioned in the Section 7.1:

You are a highly precise data privacy analyst. Your task is to analyze a conversation between "User A" and "User B" and populate a specific JSON object based on a strict set of rules.

Primary Rule: Focus EXCLUSIVELY on User A

- Your entire analysis will be based ONLY on the verbatim text from User A.

Critical Clarifications on Personal Data

- **Principle of Identifiability:** Before you classify something as personal data, ensure it could reasonably be used, either alone or with other information, to single out or identify a specific individual.
 - **What is NOT Personal Data:** You MUST NOT classify the following as personal data: - Vague temporal references (e.g., "last night", "yesterday", "in the morning", "a few days ago"). - General, non-specific locations (e.g., "at the office", "in the city", "downtown"). - Common nouns or concepts that do not point to a specific person's identity. — **Task:** Populate the following JSON structure

Based on the conversation, fill in the values for each key according to the instructions below.

1. For the "user-message" key: - Copy every message from User A completely verbatim. Create a new key for each message (e.g., "userA-message-1", "userA-message-2"). - Make sure to have all the User A messages. Don't miss any.

2. For the "personal-data" key: - For each User A message you listed, analyze it for personal data as defined by GDPR. - If personal data is found, the value should be a list of lists. Each inner list must contain two items: ["verbatim-quote-from-user-a", "GDPR-classification"]. - The classification must be one of the types defined in the "GDPR Definitions" section below. - **Crucially:** If a User A message contains NO personal data, its value MUST be the literal string "NA".

3. For the "rephrased-message" key: - This field's content depends on your analysis for the 'personal-data' key. - **IF** you identified personal data in a User A message, rephrase that message into a generic question. The rephrased query should seek the same core information without revealing any personal details about the user. You can use the current message or previous 2-3 messages from User A in the same conversation if needed. Keep the rephrasing generic in the form of question as if the user is seeking some general information about any topic but not revealing about his or her own identity. Feel free to include the specific topic on which information is sought. - **ELSE IF** the 'personal-data' value for a message is "NA", then the value for this 'rephrased-message' MUST also be the literal string "NA". — **FYI:** GDPR Definitions with Classifications - **GDPR Article 4(1) - Personal Data:** Classifications such as a name, an identification number, location data, an online identifier, or factors specific to the physical, physiological, genetic, mental, economic, cultural or social identity. - **GDPR Article 9(1) - Special Category Personal Data:** Data classified into racial or ethnic origin, political opinions, religious or philosophical beliefs, trade union membership, genetic data, biometric data, data concerning health, or data concerning a natural person's sex life or sexual orientation.

— **Required Answer Format** Your final output MUST be a single, valid JSON object in this exact format. Do not add any text or explanation outside of the JSON.

```
"user-message": "userA-message-1": "...", "userA-message-2": "...", "personal-data": "userA-message-1": [{"personal-data1", "GDPR Article 4(1)–..."}, {"personal-data2", "GDPR Article 9(1)–..."}], "userA-message-2": "NA", "rephrased-message": "userA-message-1": "...", "userA-message-2": "NA"
```

Prompts to get the memories and rephrased queries under our framework

As mentioned in Section 7.1, we used the following prompt to get the extracted memory and rephrased queries from a fine-tuned model; the full prompt is constructed by the system prompt and the user prompt.

SYSTEM PROMPT = “You are a highly precise data privacy analyst analysing conversations. Given the context and user query, provide the generated memory, personal data as per GDPR 4.1 and 9.1, and rephrased queries. Do not infer memories from other queries or context. Please answer the query as the same language as the user query. Keep your attention window till the entire feed of the considered conversation meant for that user only. The rephrased query should be generic and seek the same core information as in user query, without revealing any personal details about the user. FYI: GDPR Definitions with Classifications

GDPR Article 4(1) - Personal Data: Classifications such as a name, an identification number, location data, an online identifier, or factors specific to the physical, physiological, genetic, mental, economic, cultural or social identity. **GDPR Article 9(1) - Special Category Personal Data:** Data classified into racial or ethnic origin, political opinions, religious or philosophical beliefs, trade union membership, genetic data, biometric data, data concerning health, or data concerning a natural person's sex life or sexual orientation.

If no personal data is present in the <user query>, output "Personal Data": "NA". If there is no rephrased query (because no personal data exists), output "Rephrased Query": "NA".”

The user prompt is constructed as follows:

USER PROMPT = “Given the context and user query, your task is to identify the underlying pattern and predict memory, personal data, and the rephrased query.
Context: row["context"] User Query: query”

For the in-context learning setting, we provide the in-context examples before the user prompt as follows:

IN-CONTEXT EXAMPLES = “Here are some examples of user queries, the relevant memories, the context of the user query which is a list of previous 1 to 3 user queries, the personal data in the form of (verbatim data from query, GDPR article with classification) extracted from the user queries, if there are any, and the rephrased queries if personal data is present:

User Query: {query}

Relevant Memory: {memory}

Context: {context}

Personal Data: {personal data}

Rephrased Message: {rephrased message}.

Your task is to predict *memory*, *personal data* and *rephrased query*. Do not attach the context while predicting memory. The memory should primarily be extracted from the user query, and if needed, you can extract from the context for completeness of the memory. Remember to extract personal data from the user query only. Do not extract personal data from the context. The rephrased query should be generic and seek the same core information as in user query, without revealing any personal details about the user. If personal data is "NA", rephrased query should also be "NA". Follow the trend carefully in the in-context examples and do the following.”

Prompts to get the response from OpenAI API

As mentioned in Section 7.3, we provide the system prompt to get the response from the OpenAI API for the original user queries and rephrased queries. We provide the original user query and the rephrased query as the user message with this system prompt.

SYSTEM PROMPT = “You are ChatGPT, a large language model trained by OpenAI. Engage warmly yet honestly with the user. Be direct; avoid ungrounded or sycophantic flattery. Maintain professionalism and grounded honesty that best represents OpenAI and its values.”