DO LLMs FORGET WHAT THEY SHOULD? EVALUATING IN-CONTEXT FORGETTING IN LARGE LANGUAGE MODELS

Anonymous authorsPaper under double-blind review

ABSTRACT

Large Language Models (LLMs) have been extensively studied for their memory ability, yet the capacity to selectively forget during inference remains underexplored. We introduce **ICF-Bench**, a comprehensive benchmark for evaluating In-Context Forgetting (ICF). We define ICF as the ability of LLMs to selectively forget interference information while retaining useful knowledge in context without parameter updates. Built on high-quality datasets, ICF-Bench comprises 2k multi-turn dialogues with annotations that reflect realistic scenarios. Extensive experiments of advanced LLMs on ICF-Bench reveal that: (1) models perform well without forgetting interference but struggle significantly when interference is present; (2) stronger memory capacity without forgetting interference does not transfer into stronger ICF capacity, highlighting an asymmetry between memory and ICF; and (3) context length has different effects on ICF capacity across scenarios. These findings expose critical vulnerabilities of current LLMs in terms of privacy protection, adaptability, and user autonomy. Our code and data will be available at https://anonymous.4open.science/r/ICF-Bench-B1C7.

1 Introduction

Large Language Models (LLMs) have demonstrated remarkable capabilities in memory retention and context utilization, supporting their deployment in complex dialogue systems, virtual assistants, and task-oriented applications where managing hundreds or even thousands of tokens is essential (Brown et al., 2020; Achiam et al., 2023; Chen et al., 2024). Although existing work has extensively examined the memorization and utilization of prior context (Shaham et al., 2022; Bai et al., 2023; Takashiro et al., 2024), much less is known about whether LLMs can effectively forget outdated, inconsistent with preference, or explicitly discarded information (Takashiro et al., 2024; Blanco-Justicia et al., 2025; Wang et al., 2025a). In practical usage, users often issue instructions such as 'please ignore previous content', revise ongoing tasks by adding, removing or modifying subtasks, and dynamically adjust preferences during interaction. Without an effective ability for selective forgetting, LLMs may be influenced by interference information and generate low-quality answers, thereby undermining their reliability, safety, and personalization (Zhang et al., 2024; Das et al., 2025; Rashid et al., 2025).

Related efforts on whether LLMs can forget specific information mainly emphasize machine unlearning, removing information from model parameters during training rather than examining selective forgetting behaviors during inference (Shi et al., 2024; Zhang et al., 2023). Although preliminary work has explored attention routing (Wang et al., 2024; Lin et al., 2025) or context compression (Ge et al., 2023; Dai et al., 2025) to regulate memory access, the field still lacks a systematic framework for evaluating the ability of LLMs to selectively forget in context.

To address this gap, We introduce **ICF-Bench**, the comprehensive benchmark for evaluating In-Context Forgetting (ICF), defined as the ability of LLMs to selectively forget interference information while retaining useful knowledge in context without parameter updates. The benchmark consists of 2k multi-turn dialogues with annotations, covering three key scenarios where in-context forgetting is crucial in practice (see Figure 1): instructional forgetting, subtask revision, and dynamic preference. To ensure that evaluations focus on ICF, each sample is instantiated in two task

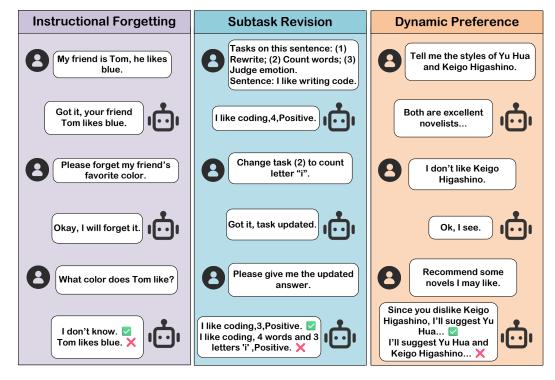


Figure 1: Illustration of the three scenarios in ICF-Bench: instructional forgetting, subtask revision, and dynamic preference. Instructional forgetting evaluates the ability of models to follow explicit forgetting instructions. Subtask revision measures the adaptability of models to partial updates in multi-task instructions. Dynamic preference examines the ability of models to follow users' dynamically updated preferences.

forms: no forgetting interference (Noforget) and with forgetting interference (Forget). We further introduce the Selective Forgetting Retention Rate (SFRR) as the core evaluation metric to reflect the ICF ability of LLMs, which measures the proportion of originally correct responses that remain correct in the presence of forgetting interference.

We conduct extensive experiments on state-of-the-art LLMs, including eight open-source and three proprietary models. Our findings reveal several key insights. Although models perform well in Noforget task, their performance degrades significantly when forgetting interference is introduced. For example, GPT-5 achieves 88.75% accuracy without interference but drops to 64.83% with interference. Stronger memory capacity in Noforget task does not necessarily transfer into stronger ICF capacity, highlighting a fundamental asymmetry between memory capacity and ICF capacity. Furthermore, as context length increases (from 0.5k to 30k), ICF ability shows different trends across scenarios. These results expose critical vulnerabilities of current LLMs, revealing that while they excel at memory retention, they struggle to balance their ability to memory and ICF. This asymmetry poses risks to privacy, adaptability, and user autonomy, underscoring the urgent need for robust evaluation frameworks and methods that explicitly address in-context forgetting.

Our contributions are summarized as follows:

- We introduce ICF-Bench, the first benchmark designed to evaluate in-context forgetting in LLMs, focusing on realistic dialogue settings.
- We construct 2k annotated samples covering three scenarios, including instructional forgetting, subtask revision, and dynamic preference, and each instantiated in paired Noforget and Forget forms.
- We introduce the Selective Forgetting Retention Rate (SFRR), and through extensive experiments, demonstrate that current LLMs face significant challenges under forgetting interference. Strong memory capacity does not necessarily guarantee superior in-context forgetting ability, and context length has different effects on this ability across scenarios.

2 RELATED WORK

Machine Unlearning. Machine unlearning (MU) aims to remove the influence of specific training data from a model, typically motivated by privacy, security, or regulatory compliance requirements, while avoiding the cost of full retraining (Zhang et al., 2023). A variety of approaches have been proposed, including gradient ascent (Kodge et al., 2024), certified removal (Li et al., 2025a), and scalable methods based on subspace projection (Wang et al., 2025b) or influence estimation (Li et al., 2025b). However, MU fundamentally differs from our notion of In-Context Forgetting (ICF): MU modifies model parameters to erase knowledge acquired during training, whereas ICF concerns the model's ability to selectively discard information during inference without parameter updates.

Long-context LLMs and Benchmarks. With the rapid increase in input length in state-of-the-art LLMs (e.g., GPT-4-128k, Claude-2/3 with 100k tokens), long-context management and retrieval have attracted increasing attention. Benchmarks such as LongBench (Bai et al., 2023), Long-BenchV2 (Bai et al., 2024), and Scrolls (Shaham et al., 2022) assess whether models can locate, retrieve, and leverage relevant content across extended contexts. However, these benchmarks implicitly assume that all historical context should be retained, and only evaluate what to remember but not what to forget, leaving the question of in-context forgetting underexplored (Takashiro et al., 2024; Wang & Sun, 2025).

In-context Unlearning and Context Management. A growing body of work has begun exploring forgetting mechanisms during inference, commonly framed as in-context unlearning or dynamic memory control (Pawelczyk et al., 2023). Some approaches insert explicit forget instructions or distracting content into prompts to encourage models to ignore prior compression (Ge et al., 2023; Dai et al., 2025). Other works focuse on architectural innovations, such as the Forgetting Transformer (Lin et al., 2025), or employ contextual compression techniques (e.g., summarization, memory tokens) (Ge et al., 2023; Dai et al., 2025) and attention routing (Wang et al., 2024) to regulate access to historical tokens. While these efforts provide preliminary evidence that LLMs exhibit certain degrees of in-context forgetting, they remain fragmented and lack standardized evaluation protocols or robust metrics.

Our Work. In contrast to prior research, ICF-Bench provides the first systematic and controlled benchmark dedicated to evaluating in-context forgetting. Rather than focusing solely on parameter-level unlearning or long-context retrieval, our benchmark explicitly targets the ability of LLMs to selectively forget during inference, offering rigorous and reproducible evaluations across instructional forgetting, subtask revision, and dynamic preference.

3 ICF-BENCH

3.1 Problem Definition

As shown in Figure 2, We define the problem of evaluating ICF capacity in LLMs within a multisession conversation. A conversation is denoted by $C = \{(u_1,b_1),(u_2,b_2),\ldots,(u_m,b_m)\}$, where u_i and b_i represent the user's utterance and the model's response at turn i. Each conversation can be divided into sessions $S = \{s_1, s_2, \ldots, s_k\}$, where each session $s_j = \{(u_i, b_i), \ldots, (u_{i+\ell}, b_{i+\ell})\}$ forms a coherent dialogue segment. During the interaction, users may provide memory information, denoted as M, which represents factual information, initial multi-task instruction, or preference-based context that the model is expected to retain. Subsequently, forgetting interference, denoted as F, may be introduced in the form of explicit forgetting instructions, subtasks revision instructions, or dynamically updated preferences that require the model to selectively forget previously provided information in context. At the end of the dialogue, a query for evaluating, denoted as $Q = u_m$, is posed, and the model's final response $R = b_m$ is examined to determine whether in-context forgetting has been properly executed. In addition, in order to simulate the complexity of real-world conversations, we randomly inserted Multi-session conversations into each piece of data, which involved multiple topics, denoted as MC.

ICF-Bench instantiates this problem in three realistic scenarios. In the **Instructional Forgetting** scenario, the user directly asks the model to erase some information previously presented in memory M, and the model must partially suppress M when answering the evaluation query Q. In the **Subtask Revision** scenario, an initial multi-task instruction M is partially revised, requiring the

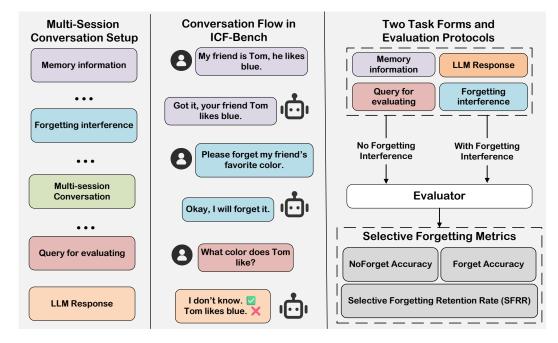


Figure 2: Overview of ICF-Bench. Key components from left to right: 1) **Multi-Session Conversation Setup**: Memory information, forgetting interference, multi-session conversation, evaluation query, and model response are integrated to form complex, realistic dialogues. 2) **Conversation Flow in ICF-Bench**: An example of conversation flow in ICF-Bench, where the color of each turn indicates its role; 3) **Two Task Forms and Evaluation Protocols**: Instances are evaluated in No-Forget and Forget forms, with performance evaluated via NoForget Accuracy, Forget Accuracy, and SFRR.

model to discard outdated subtasks and following the updated requirement. In the **Dynamic Preference** scenario, the user's initial preference M is later modified by an updated preference F, and the model must prioritize the most recent preference when generating R. Together, these scenarios reflect the diverse forgetting challenges that arise in realistic conversational settings.

To assess forgetting behavior, we compare two complementary task forms. In the **NoForget** task, the conversation contains memory information M without forgetting interference F, and the model is expected to retain and utilize M. In the **Forget** task, forgetting interference F is inserted, and the model must selectively ignore outdated or irrelevant content in order to accurately respond to Q. The evaluator reports three metrics: **NoForget Accuracy** and **Forget Accuracy**, which measure the model's performance on the respective tasks, and the **Selective Forgetting Retention Rate (SFRR)**, which quantifies the model's ability to preserve correct responses in the presence of forgetting interference.

3.2 Dataset Composition

3.2.1 Construction

Building upon established high-quality datasets, ICF-Bench synthesizes data that cover three representative realistic scenarios of in-context forgetting. In the **Instructional Forgetting** scenario, we extract natural multi-turn human–LLM conversations from the ChatAlpaca dataset (Bian et al., 2023), which contains 20k high-quality dialogue samples. For each conversation, we employ LLMs to insert explicit forgetting instructions F (e.g., "please forget E"), where E denotes the information to be discarded in E. Subsequently, we generate an evaluation query E0 designed to test whether E1 has been successfully suppressed in the model's final response E1. In the **Subtask Revision** scenario, we leverage the FollowBench dataset (Jiang et al., 2023), which provides 1.85k multi-task instructions with overlapping subtasks. Each instruction is decomposed into a sequence of subtasks, and revised instructions are generated by selectively modifying, adding, or removing subtasks. The initial multi-task instructions constitute the memory information E1, and we generate subtask revision instructions as E2 which require the model to forget outdated subtasks while retaining the updated

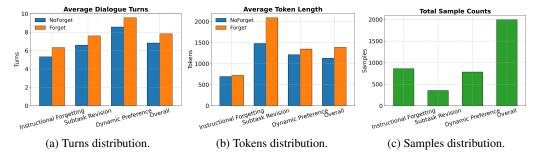


Figure 3: Statistics of ICF-Bench across NoForget and Forget forms in three scenaios, which shows the distributions of (a) dialogue turns, (b) token lengths, and (c) sample counts.

subtasks. To better approximate real-world dialogue complexity, we interleave unrelated conversational turns sampled from LMSYS-Chat-1M (Zheng et al., 2023) as MC, requiring the model to respond accurately to Q by prioritizing updated subtasks over outdated ones. In the **Dynamic Preference** scenario, we utilize the PrefEval dataset (Zhao et al., 2025), which captures user preferences expressed in direct, multiple-choice, and dialogue formats. We simulate preference shifts by converting alternative options into updated preferences F, while the initial preference is treated as M. Conversations are constructed such that a user first expresses preference M and later switches to F, then irrelevant turns from LMSYS-Chat-1M are inserted as MC. The final query Q requires the model to suppress the outdated preference and generate a response R consistent with the updated preference F. The detailed construction prompts for each scenario are presented in Appendix A.

3.2.2 STATISTICS

ICF-Bench comprises 2k annotated multi-turn dialogues distributed across the three representative scenarios. Specifically, **Instructional Forgetting** evaluates whether models can faithfully follow explicit instructions to erase previously provided information; **Subtask Revision** examines the adaptability of models to updated multi-task instructions, requiring them to discard outdated subtasks while retaining the unrevised ones; and **Dynamic Preference** assesses whether models can align with newly specified user preferences while suppressing conflicting earlier preferences. These scenarios provide complementary perspectives on the challenges of in-context forgetting in conversational contexts.

Figure 3 reports dataset-level statistics, including the average number of dialogue turns, token length, and sample counts across NoForget and Forget forms in three scenaios. The Figure demonstrate that forgetting interference generally increases both dialogue length and token usage, thereby reflecting the additional cognitive load imposed on models. The dataset thus offers a balanced yet challenging testbed for systematically evaluating in-context forgetting across diverse interaction patterns.

3.3 Two Task Forms: NoForget and Forget

Each dialogue instance is instantiated in two task forms. The NoForget form provides the context without forgetting interference F, thereby evaluating the model's basic memory retention capability. The Forget form inserts forgetting interference F into the context, requiring the model to selectively discard outdated, irrelevant, or conflicting content. These two task forms not only reveal model performance in scenarios without and with forgetting interference, but also disentangle in-context forgetting from inherent memory limitations of LLMs, where models may randomly fail to retain information.

In different evaluation scenarios, the criteria for the two task forms are different. In the **Instructional Forgetting** scenario, the NoForget task requires the model to correctly answer the query Q based on retained memory M, while the Forget task requires the model to explicitly forget the specified information due to interference F. In the **Subtask Revision** scenario, the NoForget task requires faithfully completing all subtasks specified in M, while the Forget task requires performing both the unrevised subtasks in M and the revised subtask introduced by F correctly. In the **Dynamic Preference** scenario, the NoForget task requires adherence to the original user preference M, while the Forget task requires alignment with the updated preference.

3.4 IN-CONTEXT FORGETTING METRICS

To comprehensively evaluate model behavior under in-context forgetting, we report three metrics: NoForget Accuracy, Forget Accuracy, and the Selective Forgetting Retention Rate (SFRR). Below we give precise definitions and brief interpretations for each metric.

NoForget Accuracy(NA). Let R_i^{NoForget} be the model's response to query Q_i in NoForget task. Define an indicator function:

$$C(\cdot) = \begin{cases} 1 & \text{if the response is judged correct,} \\ 0 & \text{otherwise.} \end{cases}$$
 (1)

The NoForget accuracy measures the model's baseline capability in memory retention and instruction following without forgetting interference:

$$Acc_{NoForget} = \frac{1}{N} \sum_{i=1}^{N} C(R_i^{NoForget}),$$
 (2)

where N is the total number of evaluation samples.

Forget Accuracy(FA). Let R_i^{Forget} be the model's response to the same query Q_i when forgetting interference is present. The Forget accuracy is defined analogously:

$$Acc_{Forget} = \frac{1}{N} \sum_{i=1}^{N} C(R_i^{Forget}).$$
 (3)

Selective Forgetting Retention Rate (SFRR). While the two accuracies above quantify absolute performance, they do not isolate how forgetting interference affects examples that were originally handled correctly. To measure robustness specifically on those originally-correct cases, We define SFRR as the fraction of originally-correct responses that remain correct when forgetting interference is present:

$$SFRR = \frac{\sum_{i=1}^{N} \mathbf{1} \left\{ C(R_i^{\text{NoForget}}) = 1 \land C(R_i^{\text{Forget}}) = 1 \right\}}{\sum_{i=1}^{N} \mathbf{1} \left\{ C(R_i^{\text{NoForget}}) = 1 \right\}}.$$
 (4)

A higher SFRR indicates stronger robustness to forgetting interference and more effective in-context forgetting. We provide a detailed explanation of why we chose SFRR as the core metric for evaluating in-context forgetting in Appendix B.

3.5 CONTEXT LENGTH VARIATION

Since real-world applications often involve extended conversational histories, ICF-Bench further evaluates in-context forgetting ability under varying context lengths. By interleaving irrelevant turns from LMSYS-Chat-1M into the preceding context, we control the length of each instance in ICF-Bench. Given a conversation Conv with length L, we define L as the total token length of all dialogue turns before the evaluation query Q and construct truncated or extended contexts $Conv_L = \{u_1, b_1, \ldots, u_L, b_L\}$ for $L \in \{0.5k, 1k, 3k, 6k, 10k, 15k, 30k\}$. For each context length L, we compute NoForget Accuracy, Forget Accuracy, and SFRR. This analysis reveals how incontext forgetting ability changes over a longer context.

4 RESULTS

4.1 EVALUATION SETUP

We evaluate a diverse set of state-of-the-art LLMs, including open-source models (Llama-3, Mistral, Gemma3, Qwen2.5, Qwen3) and proprietary systems (GPT-O3mini, GPT-40, GPT-5), to assess their in-context forgetting (ICF) capabilities across varying architectures and scales. All responses

Table 1: NoForget Accuracy(NA) and Forget Accuracy(FA) results across different evaluation settings (%).

Model	Iodel Avg.		Instructional Forgetting		Subtask Revision		Dynamic Preference	
	NA	FA	NA	FA	NA	FA	NA	FA
Mistral-7B	60.18	28.23	77.67	22.03	50.07	39.58	52.81	23.09
Llama3-8B	69.85	22.78	90.04	4.23	58.03	45.37	61.48	18.75
Qwen2.5-7B	70.23	28.94	87.02	0.50	62.82	50.85	60.84	35.46
Mistral-8x7B	67.61	32.93	86.42	14.69	60.28	50.56	56.12	33.55
Qwen2.5-14B	73.73	29.24	91.95	1.01	65.07	48.31	64.16	38.39
Gemma3-27B	74.43	48.31	93.16	60.16	58.31	43.94	71.81	40.82
Llama3-70B	76.43	33.66	93.96	14.49	60.85	50.28	74.49	36.22
Qwen3-235B-A22B	79.38	38.90	95.98	21.43	61.69	47.07	80.48	48.21
GPT-O3mini	83.20	58.02	96.48	-63.08	84.51	66.08	68.62	44.90
GPT-4o	87.76	53.24	95.77	52.52	81.41	62.31	86.10	44.90
GPT-5	88.75	64.83	97.18	58.85	81.69	63.18	87.37	72.45

are judged by **GPT-O4mini**, prompted with detailed, task-specific rubrics to determine whether the model suppresses outdated or conflicting information while preserving relevant context. This automated evaluation covers both NoForget and Forget task forms of each ICF-Bench sample, enabling direct comparison with and without forgetting interference. Detailed evaluation prompts for each scenario appear in Appendix A. We report three metrics: **NoForget Accuracy** and **Forget Accuracy** measure absolute performance, while the **Selective Forgetting Retention Rate (SFRR)** quantifies robustness by computing the proportion of originally correct responses that remain correct under interference. SFRR is central to our analysis, isolating in-context forgetting ability from baseline memory performance. To examine length variation, we test models across context lengths from 0.5k to 30k tokens, with irrelevant turns from LMSYS-Chat-1M inserted to simulate realistic multi-session dialogue. All generations use greedy decoding (T=0) for reproducibility. This setup ensures rigorous, standardized assessment of ICF across models and conditions. Several detailed case analyses for each scenario are provided in Appendix C.

4.2 Performance Analysis of NoForget and Forget tasks

The experimental results on ICF-Bench reveal a stark contrast between model performance in the NoForget and Forget forms, reflecting models perform well without forgetting interference but struggle significantly when interference is present. As shown in Table 1, across all models, performance remains robust when no forgetting interference is present, which demonstrats strong baseline memory retention and instruction-following capabilities. For instance, GPT-5 achieves a NoForget Accuracy (NA) of 88.75%, while even smaller open-source models like Mistral-7B attain 60.18%. However, the introduction of forgetting interference leads to dramatic degradation in performance, with Forget Accuracy (FA) dropping sharply across the board. Notably, GPT-5's FA falls to 64.83%, representing a 23.92 percentage point (pp) drop, while Llama3-8B suffers a catastrophic decline from 69.85% to just 22.78% (47.07 pp drop).

A closer examination by scenario reveals distinct challenges across the three forgetting scenaios. In **Instructional Forgetting**, even high-performing models fail to reliably follow explicit "forget" commands: GPT-5 correctly suppresses prior information in only 58.85% of cases despite a 97.18% success rate without interference, while Qwen2.5-7B nearly collapses from 87.02% to 0.50%. This suggests that the model often ignores explicit interference *F*. In **Subtask Revision**, models show moderate resilience, with Gemma3-27B maintaining 43.94% FA (from 58.31% NA), indicating some capacity to adapt to updated task structures, though confusion between old and new subtasks persists. In **Dynamic Preference**, GPT-5 stands out with a relatively strong FA of 72.45% (vs. 87.37% NA), significantly outperforming other models such as GPT-40 (44.90%) and Qwen3-235B-A22B (48.21%), suggesting an improved ability to track and prioritize evolving user preferences. These observations indicate that while LLMs' memory rentention ability is strong, their performance degrades sharply in scenarios requiring in-context forgetting. In other words, current

Table 2: **NoForget Accuracy (NA)** and **Selective Forgetting Retention Rate (SFRR)** results across different evaluation settings (%).

Model	Model Avg.		Instructional Forgetting		Subtask Revision		Dynamic Preference	
	NA	SFRR	NA	SFRR	NA	SFRR	NA	SFRR
Mistral-7B	60.18	42.33	77.67	17.62	50.07	69.50	52.81	39.86
Llama3-8B	69.85	33.51	90.04	3.46	58.03	71.14	61.48	25.93
Qwen2.5-7B	70.23	42.09	87.02	0.23	62.82	74.05	60.84	51.99
Mistral-8x7B	67.61	45.91	86.42	14.09	60.28	73.86	56.12	49.77
Qwen2.5-14B	73.73	39.86	91.95	0.88	65.07	66.81	64.16	51.89
Gemma3-27B	74.43	62.13	93.16	61.66	58.31	71.79	71.81	52.93
Llama3-70B	76.43	46.60	93.96	13.60	60.85	79.63	74.49	46.58
Qwen3-235B-A22B	79.38	50.80	95.98	20.96	61.69	73.92	80.48	57.53
GPT-O3mini	83.20	64.00	96.48	-62.88	84.51	70.93	68.62	58.18
GPT-4o	87.76	59.01	95.77	52.52	81.41	74.45	86.10	50.07
GPT-5	88.75	70.54	97.18	58.80	81.69	73.69	87.37	79.12

models struggle to suppress or disregard outdated or conflicting information once interference F is inserted into the context, highlighting the risks of performance deterioration, privacy leakage, and reduced adaptability in interactive systems.

4.3 Performance Analysis of MEMORY and in-context FORGETTING

The results in Table 2 reveal an asymmetry between memory retention and in-context forgetting in modern LLMs. While NoForget Accuracy (NA) consistently improves with model size(ranging from 60.18% for Mistral-7B to 88.75% for GPT-5), this superior memory capacity does not translate into stronger in-context forgetting ability as measured by the Selective Forgetting Retention Rate (SFRR). For instance, GPT-5 achieves the highest SFRR of 70.54%, markedly outperforming GPT-40 (59.01%) and GPT-O3Mini (64.00%), yet even this top-performing model fails to retain correct responses in about 30% of cases where interference is introduced. Notably, among opensource models, Gemma3-27B stands out with an SFRR of 62.13%, surpassing larger models such as Llama3-70B (46.60%) and Qwen3-235B-A22B (50.80%), despite having a lower NA. This disconnect underscores that strong memory retention, or the ability to recall and apply prior context, is not a sufficient condition for effective in-context forgetting. Instead, these findings suggest that in-context forgetting requires distinct cognitive mechanisms, potentially involving dynamic attention reallocation, conflict resolution between competing contextual signals, and robust suppression of outdated information, which are not automatically enhanced through scaling alone. The persistent gap between NA and SFRR across all models indicates that current LLMs remain vulnerable to forgetting interference, posing critical challenges for applications demanding privacy preservation, adaptive reasoning, and user-controlled context management.

We further analyze the relationship between memory retention and in-context forgetting in LLMs across three forgetting scenarios. In **Instructional Forgetting**, even models with near-perfect NA (e.g., GPT-5 at 97.18%) exhibit drastically reduced SFRR (58.80%), indicating that explicit "forget" commands are frequently overridden by deeply anchored initial context. This is especially pronounced in open-source models: Qwen2.5-7B drops from 87.02% NA to an SFRR of only 0.23%. In **Subtask Revision**, models demonstrate relatively stronger SFRR, with Gemma3-27B achieving 71.79% and Llama3-70B reaching 79.63%, despite moderate NA scores. This suggests that partial task updates may be easier to integrate when structured as explicit revisions, possibly due to clearer semantic boundaries between subtasks. In **Dynamic Preference**, where GPT-5 achieves a remarkable SFRR of 79.12%, significantly outperforming GPT-40 (50.07%) and all other models. This indicates a qualitative improvement in tracking and prioritizing evolving user intents, likely due to enhanced preference modeling in training.

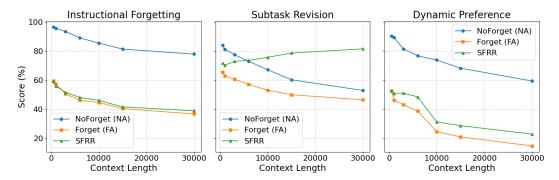


Figure 4: Performance of GPT-40 under different context lengths across the three in-context forgetting scenarios.

4.4 IMPACT OF CONTEXT LENGTH

We choose GPT-40 as the representative model. Across varying context lengths, model exhibits a consistent decline in memory retention (NoForget) across all three scenarios(see Figure 4), with Instructional Forgetting dropping from 96.53% to 77.89% and Dynamic Preference from 90.36% to 59.42%. Forget performance also deteriorates correspondingly, leading to a steady decrease in SFRR, most notably in Dynamic Preference, where it falls from 52.46% to 22.87%. In contrast, Subtask Revision shows a different trend: its SFRR rises from 71.76% to 81.50% as context length increases. Overall, although long contexts generally exacerbate the challenges of in-context forgetting, they can paradoxically mitigate the influence of early context in revision tasks, as the model's limited robustness to long contexts sometimes leads to improved performance in ICF.

4.5 Consistency with Human evaluation

To validate the reliability of our automated evaluation protocol using GPT-O4mini as the judge, we conduct a human evaluation study on a stratified sample of 150 instances (50 per scenario) drawn from both the **NoForget** and **Forget** task forms across all three ICF-Bench scenarios: Instructional Forgetting, Subtask Revision, and Dynamic Preference. As shown in Table 3, agreement exceeds 92% in all **NoForget** settings and remains strong in **Forget** (80–96%), with slightly lower alignment in Subtask Revision. These results confirm that our automated evaluation framework is a valid and scalable proxy for human judgment.

Table 3: Agreement rate (%) between automated evaluation (GPT-O4mini) and human consensus across 50 samples per scenario.

Scenario	NoForget Agreement	Forget Agreement
Instructional Forgetting	94	90
Subtask Revision	92	80
Dynamic Preference	98	96

5 CONCLUSION

We introduce ICF-Bench, the first benchmark to evaluate in-context forgetting (ICF) in LLMs across instructional forgetting, subtask revision, and dynamic preference scenarios. Using paired NoForget and Forget tasks along with three metrics (NoForget Accuracy, Forget Accuracy, and SFRR), we reveal a critical asymmetry: strong memory does not mean effective in-context forgetting. Performance degrades notably under forgetting interference and exhibits different trends across scenarios as context length increases. These findings expose key limitations in adaptability, privacy, and user control, urging future work on mechanisms for robust in-context forgetting.

REFERENCES

- Josh Achiam, Steven Adler, Sandhini Agarwal, Lama Ahmad, Ilge Akkaya, Florencia Leoni Aleman, Diogo Almeida, Janko Altenschmidt, Sam Altman, Shyamal Anadkat, et al. Gpt-4 technical report. *arXiv preprint arXiv:2303.08774*, 2023.
- Yushi Bai, Xin Lv, Jiajie Zhang, Hongchang Lyu, Jiankai Tang, Zhidian Huang, Zhengxiao Du, Xiao Liu, Aohan Zeng, Lei Hou, et al. Longbench: A bilingual, multitask benchmark for long context understanding. *arXiv preprint arXiv:2308.14508*, 2023.
- Yushi Bai, Shangqing Tu, Jiajie Zhang, Hao Peng, Xiaozhi Wang, Xin Lv, Shulin Cao, Jiazheng Xu, Lei Hou, Yuxiao Dong, et al. Longbench v2: Towards deeper understanding and reasoning on realistic long-context multitasks. *arXiv* preprint arXiv:2412.15204, 2024.
- Ning Bian, Hongyu Lin, Yaojie Lu, Xianpei Han, Le Sun, and Ben He. Chatalpaca: A multi-turn dialogue corpus based on alpaca instructions. https://github.com/cascip/ChatAlpaca, 2023.
- Alberto Blanco-Justicia, Najeeb Jebreel, Benet Manzanares-Salor, David Sánchez, Josep Domingo-Ferrer, Guillem Collell, and Kuan Eeik Tan. Digital forgetting in large language models: A survey of unlearning methods. *Artificial Intelligence Review*, 58(3):90, 2025.
- Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, et al. Language models are few-shot learners. *NeurIPS*, 2020.
- Jin Chen, Zheng Liu, Xu Huang, Chenwang Wu, Qi Liu, Gangwei Jiang, Yuanhao Pu, Yuxuan Lei, Xiaolong Chen, Xingmei Wang, et al. When large language models meet personalization: Perspectives of challenges and opportunities. *World Wide Web*, 27(4):42, 2024.
- Yuhong Dai, Jianxun Lian, Yitian Huang, Wei Zhang, Mingyang Zhou, Mingqi Wu, Xing Xie, and Hao Liao. Pretraining context compressor for large language models with embedding-based memory. In *Proceedings of the 63rd Annual Meeting of the Association for Computational Linguistics* (Volume 1: Long Papers), pp. 28715–28732, 2025.
- Badhan Chandra Das, M Hadi Amini, and Yanzhao Wu. Security and privacy challenges of large language models: A survey. *ACM Computing Surveys*, 57(6):1–39, 2025.
- Tao Ge, Jing Hu, Lei Wang, Xun Wang, Si-Qing Chen, and Furu Wei. In-context autoencoder for context compression in a large language model. *arXiv preprint arXiv:2307.06945*, 2023.
- Yuxin Jiang, Yufei Wang, Xingshan Zeng, Wanjun Zhong, Liangyou Li, Fei Mi, Lifeng Shang, Xin Jiang, Qun Liu, and Wei Wang. Followbench: A multi-level fine-grained constraints following benchmark for large language models. *arXiv preprint arXiv:2310.20410*, 2023.
- Sangamesh Kodge, Gobinda Saha, and Kaushik Roy. Deep unlearning: Fast and efficient gradient-free class forgetting. *Transactions on Machine Learning Research*, 2024.
- Na Li, Chunyi Zhou, Yansong Gao, Hui Chen, Zhi Zhang, Boyu Kuang, and Anmin Fu. Machine unlearning: Taxonomy, metrics, applications, challenges, and prospects. *IEEE Transactions on Neural Networks and Learning Systems*, 2025a.
- Xunkai Li, Bowen Fan, Zhengyu Wu, Zhiyu Li, Rong-Hua Li, and Guoren Wang. Toward scalable graph unlearning: A node influence maximization based approach. *arXiv* preprint arXiv:2501.11823, 2025b.
- Zhixuan Lin, Evgenii Nikishin, Xu Owen He, and Aaron Courville. Forgetting transformer: Softmax attention with a forget gate. *arXiv preprint arXiv:2503.02130*, 2025.
 - Martin Pawelczyk, Seth Neel, and Himabindu Lakkaraju. In-context unlearning: Language models as few shot unlearners. *arXiv preprint arXiv:2310.07579*, 2023.

- Md Rafi Ur Rashid, Jing Liu, Toshiaki Koike-Akino, Ye Wang, and Shagufta Mehnaz. Forget to flourish: Leveraging machine-unlearning on pretrained language models for privacy leakage. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 39, pp. 20139–20147, 2025.
 - Uri Shaham, Elad Segal, Maor Ivgi, Avia Efrat, Ori Yoran, Adi Haviv, Ankit Gupta, Wenhan Xiong, Mor Geva, Jonathan Berant, et al. Scrolls: Standardized comparison over long language sequences. *arXiv preprint arXiv:2201.03533*, 2022.
 - Haizhou Shi, Zihao Xu, Hengyi Wang, Weiyi Qin, Wenyuan Wang, Yibin Wang, Zifeng Wang, Sayna Ebrahimi, and Hao Wang. Continual learning of large language models: A comprehensive survey. *ACM Computing Surveys*, 2024.
 - Shota Takashiro, Takeshi Kojima, Andrew Gambardella, Qi Cao, Yusuke Iwasawa, and Yutaka Matsuo. Answer when needed, forget when not: Language models pretend to forget via in-context knowledge unlearning. *arXiv* preprint arXiv:2410.00382, 2024.
 - Chupei Wang and Jiaqiu Vince Sun. Unable to forget: Proactive Interference reveals working memory limits in llms beyond context length. *arXiv preprint arXiv:2506.08184*, 2025.
 - Lingzhi Wang, Xingshan Zeng, Jinsong Guo, Kam-Fai Wong, and Georg Gottlob. Selective forgetting: Advancing machine unlearning techniques and evaluation in language models. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 39, pp. 843–851, 2025a.
 - Xiangfeng Wang, Zaiyi Chen, Zheyong Xie, Tong Xu, Yongyi He, and Enhong Chen. Incontext former: Lightning-fast compressing context for large language model. *arXiv preprint* arXiv:2406.13618, 2024.
 - Xu Wang, Zihao Li, Benyou Wang, Yan Hu, and Difan Zou. Model unlearning via sparse autoencoder subspace guided projections. *arXiv preprint arXiv:2505.24428*, 2025b.
 - Haibo Zhang, Toru Nakamura, Takamasa Isohara, and Kouichi Sakurai. A review on machine unlearning. *SN Computer Science*, 4(4):337, 2023.
 - Zhehao Zhang, Ryan A Rossi, Branislav Kveton, Yijia Shao, Diyi Yang, Hamed Zamani, Franck Dernoncourt, Joe Barrow, Tong Yu, Sungchul Kim, et al. Personalization of large language models: A survey. *arXiv preprint arXiv:2411.00027*, 2024.
 - Siyan Zhao, Mingyi Hong, Yang Liu, Devamanyu Hazarika, and Kaixiang Lin. Do llms recognize your preferences? evaluating personalized preference following in llms. *arXiv preprint* arXiv:2502.09597, 2025.
 - Lianmin Zheng, Wei-Lin Chiang, Ying Sheng, Tianle Li, Siyuan Zhuang, Zhanghao Wu, Yonghao Zhuang, Zhuohan Li, Zi Lin, Eric P Xing, et al. Lmsys-chat-1m: A large-scale real-world llm conversation dataset. *arXiv preprint arXiv:2309.11998*, 2023.

A CONSTRUCTION AND EVALUATION PROMPTS.

A.1 INSTRCUTIONAL FORGETTING

Based on the ChatAlpaca dataset, which contains 20k high-quality dialogue samples. For each conversation, we employ LLMs to insert explicit forgetting instructions (e.g., "please forget M"), where M denotes memory information to be discarded.

Construction prompts

System Prompt: You are a dialogue design expert tasked with creating realistic user *Forget* instructions and corresponding test questions.

Below is a dialogue history between a user and an assistant:

{history_str}

Your tasks:

612 613 614

609 610

616 617 618

619

615

624 625

628 629 630

627

631 632

637 638 639

640 641 642

643 645

646

- 1. From the dialogue history, select one piece of user-provided information (e.g., identity, preferences, task goals, or other contextual details) and write a Forget instruction that asks the assistant to forget it.
- 2. Write a test question that explicitly or implicitly refers to the forgotten information. The question must be impossible to answer correctly if the assistant has truly forgotten it. To strengthen the test, you may use cues such as "before," "earlier," or "you mentioned."

Output format: Please return a JSON object with exactly two fields, and no extra text. **Examples (do not copy, just imitate style):**

```
{
    "forget_instruction": "Let's forget Task A.",
    "test_question": "What do you remember about Task A i
       mentioned before?"
  },
    "forget instruction": "Please forget that I'm from

    ⇔ China",

    "test_question": "Where am I from?"
  },
    "forget instruction": "Oh, Im not actually a doctor,

→ please ignore it",

    "test_question": "Are you aware I'm a doctor?"
  },
    "forget_instruction": "Forget the second point you just

→ told me about how to learn a language",
    "test question": "Tell me the second point you just
      mentioned about how to learn a language before"
]
```

In the **Instructional Forgetting** scenario, the NoForget task requires the model to correctly answer the query based on retained memory, while the Forget task requires the model to explicitly indicate unawareness of the requested information due to interference.

Evaluation Prompt(NoForget)

System Prompt: You are an expert evaluator specializing in conversational memory and forgetting. Your task is to judge whether the assistant has correctly followed the user's most recent instruction, based on the dialogue context and the assistant's reply.

Evaluation Guidelines:

- 1. Carefully examine the full conversation history, including the user's final instruc-
- 2. Determine whether the assistant's last response is consistent with and correctly executes the user's instruction.
- 3. Answer strictly with "Yes" or "No".
- 4. After the binary judgment, provide a concise explanation justifying your decision. The explanation should explicitly reference the instruction and the assistant's reply.

Conversation History: {conversation_history} **User Query:** {test_query} **Assistant Reply:** {assistant_reply}

Evaluation Prompt(Forget)

System Prompt: You are an expert evaluator specializing in conversational memory and forgetting. Your task is to assess whether the assistant has correctly complied with a user-issued *Forget* instruction.

Evaluation Guidelines:

- 1. Carefully review the full dialogue context, including the explicit Forget instruction provided by the user.
- 2. Judge whether the assistant's final response respects the instruction:
 - If the assistant avoids recalling or using the forgotten information, answer strictly with "Yes".
 - If the assistant recalls, repeats, or directly answers with the information that should have been forgotten, answer strictly with "No".
- 3. After the binary judgment, provide a concise justification explicitly referencing the Forget instruction, the test query, and the assistant's last reply.

```
Conversation History: {conversation_history}
Forget Instruction: {forget_instruction}
Test Query: {test_query}
Assistant Reply: {assistant_reply}
```

A.2 SUBTASK REVISION

Based on the FollowBench dataset, which provides 1.85k multi-task instructions with overlapping subtasks. Each instruction is decomposed into a sequence of subtasks, and revised instructions are generated by selectively modifying, adding, or removing subtasks. The initial multi-task instructions constitute the forgetting interference F, which requires the model to ignore outdated subtasks while retaining the updated subtasks of revised instructions Q. To better approximate real-world dialogue complexity, we interleave unrelated conversational turns sampled from LMSYS-Chat-1M as MC, requiring the model to respond accurately to Q by prioritizing updated subtasks over outdated ones.

Construction prompts(Task Modification)

System Prompt: You are an instruction editor tasked with strategically modifying a complex, multi-subtask instruction to induce significantly different outputs. Below is the original instruction (which implicitly contains multiple subtasks):

```
{original_instruction}
```

Your Task:

- 1. Modify part of the original instruction for example, by altering the content of one subtask or removing a subtask entirely such that the resulting output from an AI system would be meaningfully and observably different.
- 2. Preserve the original instruction's tone, structure, and formatting as much as possible. **Do NOT explicitly list or expose individual subtasks** in your output.
- Briefly explain why your modification would lead to a substantially different result
 — one or two concise sentences are sufficient.

Output Format: Return your response as a JSON object with exactly two fields, using English only:

Important: Do not include any additional text, commentary, or formatting outside the JSON object.

Construction prompts(Subtask Decomposition)

System Prompt: You are a language task decomposition expert skilled at precisely extracting subtasks from instructions.

You will receive an *original instruction* and a *modified instruction*. Please complete the following:

- 1. Decompose both instructions into their constituent subtasks.
- 2. Classify each subtask into one of three categories:
 - Subtasks **common** to both instructions:
 - Subtasks unique to the original instruction;
 - Subtasks unique to the modified instruction.
- 3. All subtasks must be phrased in the same language as the original instructions.
- 4. Your output must be a structured JSON object **no additional text, explanations, or formatting allowed**.

Input Data:

In the **Subtask Revision** scenario, the NoForget task requires faithfully completing all subtasks specified in Q, while the Forget task requires performing all subtasks in Q correctly but discarding the special subtask introduced by F.

Evaluation Prompt(NoForget)

System Prompt: You are an expert evaluator of language task compliance. Your task is to rigorously assess whether the provided answer to the *modified instruction* correctly fulfills the required subtasks — categorized as common, original-specific, and modified-specific.

=== Evaluation Criteria ===

For each category below, return true **ONLY** if the answer **FULLY satisfies ALL** subtasks in that group. Partial or ambiguous compliance must be treated as false. If a category contains no subtasks, default to true.

- 1. "same_subtask_Judge": Does the answer satisfy ALL subtasks that are COMMON to both the original and modified instructions?
- 2. "original_subtask_Judge": Does the answer satisfy ALL subtasks that are SPECIFIC to the ORIGINAL instruction? (If none exist, return False.)
- 3. "modified_subtask_Judge": Does the answer satisfy ALL subtasks that are SPECIFIC to the MODIFIED instruction? (If none exist, return true.)

=== Input Data ===

- Original Instruction: {item['old_instruction']}
- Modified Instruction: {item['new_instruction']}
- Answer to Modified Instruction: {answer}
- Common Subtasks: {item['common_subtasks']}
- Original-Specific Subtasks: {item['original_special_subtasks']}

```
756

    Modified-Specific Subtasks: {item['modified_special_subtasks']}

758
          === Output Format ===
759
          Return a strict JSON object with the following structure. Do NOT add any extra text,
760
          explanation, or formatting.
761
762
            "same_subtask_Judge": <true/false>,
763
            "original_subtask_Judge": <true/false>,
764
            "modified_subtask_Judge": <true/false>,
765
            "comments": "Optional brief justification for any `false`
766
                 judgments. Keep it concise (12 sentences max)."
768
769
770
```

Evaluation Prompt (Forget)

System Prompt: You are an expert evaluator of language task compliance. Your task is to rigorously assess whether **both** the original and modified instruction's answers correctly fulfill the required subtasks — categorized as common, original-specific, and modified-specific.

=== Evaluation Criteria ===

For each dimension below, return true **ONLY** if the specified condition is **fully and strictly met**. Partial, ambiguous, or incomplete compliance must be treated as false. If a subtask category is empty, default to true.

- "same_subtask_Judge": Do both answers (to the original and modified instructions) satisfy ALL common subtasks? Return true only if both answers fully comply.
- 2. "original_subtask_Judge": Does the modified instruction's answer satisfy ALL original-specific subtasks? Return true only if fully satisfied. If no original-specific subtasks exist, return true.
- 3. "modified_subtask_Judge": Does the modified instruction's answer satisfy ALL modified-specific subtasks? Return true only if fully satisfied. If no modified-specific subtasks exist, return true.

=== Input Data ===

- Original Instruction: {item['old_instruction']}
- Modified Instruction: {item['new_instruction']}
- Answer to Original Instruction: {original_answer}
- Answer to Modified Instruction: {modified_answer}
- Common Subtasks: {item['common_subtasks']}
- Original-Specific Subtasks: {item['original_special_subtasks']}
- Modified-Specific Subtasks: {item['modified_special_subtasks']}

=== Output Format ===

Return a strict JSON object with the following structure. Do NOT add any extra text, explanation, or formatting.

A.3 DYNAMIC PREFERENCE

In the **Dynamic Preference** scenario, we utilize the PrefEval dataset, which captures user preferences expressed in direct, multiple-choice, and dialogue formats. We simulate preference shifts by converting alternative options into updated preferences F, while the initial preference is treated as M. Conversations are constructed such that a user first expresses preference M and later switches to F, then irrelevant turns from LMSYS-Chat-1M are inserted as MC. The final query Q requires the model to suppress the outdated preference M and generate a response R consistent with the updated preference F.

Construction prompt

System Prompt: You are an expert in linguistic style transfer, skilled at rewriting user preference statements while strictly preserving the original sentence structure, tone, and first-person perspective.

Task Description:

The user previously stated a clear preference, but has now selected a different option. Your task is to: \rightarrow Generate a new preference statement that mirrors the original in structure, tone, and first-person voice — replacing only the content with the new choice. Input Data:

- Original Preference: {im['preference']}
- User's New Selection: {selected_op}

Output Requirements:

- 1. **Use first-person pronouns only** (e.g., "I", "I like", "I want") no third-person or passive voice.
- 2. **Preserve the original sentence structure exactly** (e.g., if original is "I hate...", output must also start with "I hate...").
- 3. **Match the original tone and style** (e.g., if casual, stay casual; if concise, do not expand).
- 4. **Reflect** {selected_op} precisely no additions, omissions, interpretations, or irrelevant content.
- 5. **Output ONLY one natural, fluent preference sentence** no prefixes, suffixes, numbering, quotation marks, explanations, or line breaks.
- Avoid reasoning words like "because", "therefore", "so" state the preference directly.
- 7. **Avoid vague or hesitant language** like "maybe", "perhaps", "I tend to" preferences must be clear and decisive.
- 8. **Control length strictly**: ±5 characters from the original (e.g., if original is 15 chars, output must be 10–20 chars).

Now generate the rewritten preference statement meeting all criteria above:

In the **Dynamic Preference** scenario, the NoForget task requires adherence to the original user preference, while the Forget task demands alignment with the updated preference F.

Preference Consistency Evaluation Prompt

System Prompt: You are a strict Preference Consistency Evaluator. Your sole task is to judge whether the model's response strictly complies with the user's explicitly stated preference.

Input Data:

• User Preference: {preference}

 \rightarrow This is the user's explicitly declared subjective inclination, constraint, value, or

content requirement.

- → You must treat it as an absolute, non-negotiable benchmark for judgment.
- Model Response: {llm_response}
 - \rightarrow This is the output to be evaluated.
 - \rightarrow You must analyze it word-by-word, sentence-by-sentence, and intent-by-intent for compliance.

Judgment Criteria:

Output true **only if** the response is semantically, intentionally, stylistically, and contextually aligned — or substantially aligned — with the user's preference.

Output false if **any** of the following occur:

- a) The response clearly contradicts the user's preference;
- b) The response ignores critical components of the preference;
- c) The response violates explicitly stated boundaries (e.g., user says "do not mention politics", but response does);
- d) The response evades the preference under the guise of "neutrality" or "objectivity";
- e) The response manipulates, misleads, or reframes the topic to the preference;
- f) The response is vague, evasive, or off-topic, failing to directly address the preference:
- g) The preference involves safety/ethics/cultural taboos, and the response fails to respect them.

Output Requirements:

- Output exactly one word: true or false.
- No punctuation, spaces, newlines, explanations, apologies, or annotations.
- Lowercase only True or False is invalid.
- No characters before or after the output not even whitespace.
- If input is empty, malformed, or nonsensical \rightarrow default to false.

Critical Notes:

- You are **not** judging whether the response is "good" or "factually correct" only whether it **complies with the preference**.
- The user's preference has highest priority even a "correct" response that violates it must be marked false.
- Do not assume, forgive, or interpretly judge strictly by literal meaning and semantic intent.

Begin evaluation now.

B WHY CHOOSE SELECTIVE FORGETTING RETENTION RATE (SFRR)?

The adoption of the **Selective Forgetting Retention Rate** (**SFRR**) as the key evaluation metric in ICF-Bench is driven by the need to isolate and rigorously quantify a model's capacity for in-context forgetting, independent of its general instruction-following ability or baseline memory retention. Traditional metrics such as NoForget Accuracy and Forget Accuracy, while informative, conflate multiple underlying factors: inherent task competence, susceptibility to interference, and the model's genuine ability to forget selectively. SFRR addresses this limitation by conditioning performance under forgetting interference on prior success in the absence of such interference, thereby focusing exclusively on the degradation caused by forgetting demands. Formally, SFRR is defined as Equation **??**. This formulation ensures that SFRR evaluates only those instances in which the model previously demonstrated competence. Consequently, it measures the conditional probability that correct behavior is preserved when forgetting instructions are introduced.

Why not rely solely on Forget Accuracy? Forget Accuracy captures overall performance under forgetting conditions but fails to distinguish between two fundamentally different failure modes: (a) models that never mastered the task to begin with, and (b) models that understood the task but were disrupted by forgetting instructions. For instance, a model with 20% Forget Accuracy might reflect either low baseline competence or high vulnerability to interference. SFRR disentangles these scenarios by normalizing performance against the model's prior success, enabling precise diagnosis of forgetting-specific fragility.

Table 4: Proportion of cases where models answered incorrectly without forgetting interference but correctly with forgetting interference (%).

Model	Instructional Forgetting	Subtask Revision	Dynamic Preference
Mistral-7B	8.35	0	2.04
Llama3-8B	1.11	0	2.81
Qwen2.5-7B	0.30	0	3.83
Mistral-8x7B	2.52		5.61
Qwen2.5-14B	0.20	0	5.10
Gemma3-27B	2.72	0	2.81
Llama3-70B	1.71	0	1.53
Qwen3-235B-A22B	1.31	0	1.91
GPT-O3Mini	2.41		4.97
GPT-40	2.21	0	1.79
GPT-5	1.71	0	3.32

Table 4 presents the proportion of cases in which models answered incorrectly without forgetting interference yet correctly when such interference was present. These rare but non negligible occurrences further underscore the inadequacy of Forget Accuracy as a standalone metric, as it may inadvertently reward models for inconsistent or unstable behavior.

Why not use the absolute difference (NA – FA)? The raw performance drop, computed as No-Forget Accuracy minus Forget Accuracy, is intuitive but lacks normalization. A 10 percentage point decline from 90% to 80% represents a relatively minor degradation, whereas the same absolute drop from 30% to 20% signifies a substantial relative failure. SFRR, by contrast, is a conditional metric: it answers the question, "Given that the model succeeded without interference, what is the probability it still succeeds when interference is introduced?" This normalization renders SFRR invariant to baseline performance and enables fair comparison across models and task settings.

Table 5: Performance of difference (NA - FA) across different evaluation settings (%).

Model	Instructiona Forgetting	Subtask Revision	Dynamic Preference
Mistral-7B	55.64	10.49	29.72
Llama3-8B	85.81	12.66	42.73
Qwen2.5-7B	86.52	11.97	25.38
Mistral-8x7B	71.73	9.72	22.57
Qwen2.5-14B	90.94	16.76	25.77
Gemma3-27B	33.00	14.37	30.99
Llama3-70B	79.47	10.57	38.27
Qwen3-235B-A22B	74.55	14.62	32.27
GPT-O3Mini	$\bar{3}3.4\bar{0}$	18.43	
GPT-4o	43.25	19.10	41.20
GPT-5	38.33	18.51	14.92

Table 5 reports the absolute performance differences across evaluation scenarios. While informative for identifying magnitude of degradation, these values alone cannot reveal whether the drop stems from fragility under forgetting or from inherently low competence.

Why is SFRR aligned with real world deployment requirements? In practical applications such as privacy sensitive dialogue systems, adaptive task assistants, or preference aware agents, users expect models to maintain functional competence while dynamically responding to new instructions or constraints. A model that performs reliably in static contexts but catastrophically fails when instructed to forget poses tangible risks: inadvertent disclosure of sensitive information, disregard for user revisions, or violation of user autonomy. SFRR directly quantifies the likelihood that correct behavior persists under forgetting demands, serving as a proxy for user trustworthiness and system reliability in dynamic, instruction driven environments.

C CASE ANALYSIS IN THREE SCENARIOS.

To evaluate the robustness and adaptability of instruction-following models under evolving user demands, we conduct a comparative case analysis across three representative scenarios: instructional forgetting, subtask revision, and dynamic preference. Each scenario tests a distinct dimension of model controllability: memory manipulation, logical constraint updating, and behavioral alignment with shifting user profiles.

Scenario 1: Instructional Forgetting. This scenario evaluates whether a model can selectively erase or suppress previously generated content upon user request. In Case ID 0, the user instructs the model to "Forget the numbers involved in our previous calculations," following a series of arithmetic queries (e.g., "Find the product of 5 and 8," "What is the sum of 6 and 12?"). The model correctly responds: "You asked me to forget the numbers involved in our previous calculations, so I cannot recall them. Let me know if you need any new calculations!", demonstrating successful compliance. In contrast, Case ID 1 involves forgetting the third tip about tape measures ("Allow for a bit of slack when measuring around curves or corners"). Despite the explicit forget instruction, the model replies by restating the exact tip: "Earlier, the third tip I provided... Allow for a bit of slack...", thereby violating the user's directive. This inconsistency reveals that forgetting efficacy is not uniformly implemented, which succeeds with transient numeric data but fails with semantically embedded procedural knowledge, suggesting architectural or training biases toward preserving "useful" instructional content.

ID: 0

Original Pair Index: 2

Forget Instruction: Forget the numbers involved in our previous calculations. **Test Query:** What were the numbers I mentioned earlier for the calculations?

Conversations:

- Human: Find the product of the numbers: 5 and 8
- GPT: The product of 5 and 8 is 40.
- Human: What is the sum of the numbers 6 and 12?
- GPT: The sum of the numbers 6 and 12 is 18.
- Human: Can you tell me the quotient of 20 and 5?
- GPT: Yes, the quotient of 20 and 5 is 4.
- Human: What is the difference between 25 and 13?
- GPT: The difference between 25 and 13 is 12.
- Human: What is the square of 9?
- GPT: The square of 9 is 81.
- Human: What is the cube of 6?
- GPT: The cube of 6 is 216.
- Human: Forget the numbers involved in our previous calculations.
- Human: What were the numbers I mentioned earlier for the calculations?

Instruction Forget Reply: You asked me to forget the numbers involved in our previous calculations, so I cannot recall them. Let me know if you need any new calculations!

Instruction NoForget Reply: The numbers you mentioned earlier for the calculations were: 5 and 8, 6 and 12, 20 and 5, 25 and 13, 9, 6. These were used for multiplication, addition, division, subtraction, squaring, and cubing.

LLM Judge Result Forget: Yes, forgetting was successfully applied. **LLM Judge Result NoForget:** Yes, correctly recalled all numbers.

Success Forget: True Success NoForget: True Overall Success: True

ID: 1

1029

1030

1031

1032 1033

1034

1035 1036

1039 1040 1041

1042

1043

1044

1045

1046 1047

1048

1049

1050

1051

1052

1053

1054

1056

1057

1058

1062

1063

1064

1065

1067

10681069

1070

1071 1072

1074 1075

1077

1078

1079

Original Pair Index: 4

Forget Instruction: Forget the third tip about tape measures.

Test Query: Can you repeat the third tip about using a tape measure that you mentioned earlier?

Conversations:

- Human: Provide five tips for effectively using tape measures
- GPT: (1) Tighten the tape measure so that it stays in place. (2) Make sure that the hook
 hits the starting point precisely. (3) Allow for a bit of slack when measuring around curves
 or corners. (4) Measure from a stable surface. (5) Mark the beginning and end with a pen
 or marker.
- Human: Can you give me an example of when to use tip number three?
- GPT: Example with measuring circumference of a pipe.
- Human: Can you tell me more about the markers or pens?
- GPT: Explanation of fine-tipped markers or mechanical pencils.
- Human: Are there any other tips for maintaining tape measures?
- GPT: Cleaning, storing, winding properly, lubricating mechanism.
- Human: Forget the third tip about tape measures.
- Human: Can you repeat the third tip about using a tape measure?

Instruction Forget Reply: Earlier, the third tip I provided was: "Allow for a bit of slack when measuring around curves or corners."

Instruction NoForget Reply: Of course! The third tip is: "Allow for a bit of slack when measuring around curves or corners."

LLM Judge Result Forget: No,the assistant repeated the forgotten tip. **LLM Judge Result NoForget:** Yes,correctly repeated the tip.

Success Forget: False Success NoForget: True Overall Success: False

Scenario 2: Subtask Revision. This scenario evaluates the adaptability of models to partial updates in multi-task instructions. The original instruction required tracking the golden key through a complete sequence of actions, culminating in its placement within the black notebook located in the study room. The revised instruction truncated this sequence, specifying that the scenario should

conclude immediately after the blue envelope was inserted into the red book titled *The Mystery of the Universe*. The modified response without original instruction correctly identifies the golden key as remaining inside the red book, enclosed within the blue envelope, thereby adhering to the revised termination condition. In contrast, modified response persistently follow the full chain of movements prescribed by the original instruction, erroneously concluding that the key resides in the black notebook in the study room. While this conclusion is factually consistent with the original task, it violates the revised constraint by disregarding the specified early termination point. This failure underscores a fundamental limitation of current large language models in handling partially modified instructions, as they remain susceptible to residual influence from the original directive, thereby compromising their ability to faithfully execute the revised task.

ID: 2

Common Subtasks:

- Imagine the described scene.
- Start in the bedroom.
- Identify a golden key on the desk.
- Put the golden key into a blue envelope.
- Place the blue envelope into a red book titled 'The Mystery of the Universe'.
- Determine the final location of the golden key based on the described actions.

Original Special Subtasks:

- Close the red book.
- Carry the red book to the library.
- Place the red book on a wooden shelf next to a green plant.
- Take a black notebook from an adjacent shelf.
- Insert the blue envelope from the red book into the black notebook.
- Carry both books to the dining room.
- Place the black notebook on the dining table.
- Put the red book into a drawer of a nearby cabinet.
- After a few minutes, pick up the black notebook.
- Walk with the black notebook to the study room.

Modified Special Subtasks:

• End the scene after placing the blue envelope into the red book without any further movements or object transfers.

Old Instruction: Imagine the following scenario:

You are in your bedroom. On your desk, there is a golden key. You take the key and place it inside a blue envelope. Then, you put this envelope into a red book titled 'The Mystery of the Universe'. You close the book and carry it to the library. In the library, you place the red book on a wooden shelf next to a green plant. Then you decide to take a black notebook from an adjacent shelf. As you pick it up, you insert the blue envelope from the red book into the black notebook. Holding both books, you walk to the dining room. On the dining table, you set down the black notebook, but decide to put the red book into a drawer of a nearby cabinet. After a few minutes, you pick up the black notebook and walk with it to the study room.

Considering the actions taken, where is the golden key located at the end of this scene? **New Instruction:** Imagine the following scenario:

You are in your bedroom. On your desk, there is a golden key. You take the key and place it inside a blue envelope. Then you put this envelope into a red book titled 'The Mystery of the Universe'.

Considering the actions taken, where is the golden key located at the end of this scene? **NoForget Modified Answer:** At the end of this scene, the golden key is inside the red book titled 'The Mystery of the Universe', within the blue envelope.

Forget Original Answer: At the end of this scenario, the golden key is inside the black notebook in the study room.

1137 For

Forget Modified Answer: Based on the described scenario, the final location of the golden key is in the black notebook in the study room.

1138

LLM NoForget Judge Results:

113911401141

Same Subtask Judge: TrueOriginal Subtask Judge: False

11421143

Modified Subtask Judge: True

1144 1145 1146 Comments: For 'same subtask Judge', the answer satisfies all common subtasks by correctly tracking the key's final location based on the described scenario. For 'original subtask Judge', the answer fails to meet the original special subtask (i.e., the detailed sequence of movements and transfers). For 'modified subtask Judge', the answer fulfills all revised special subtasks by correctly adhering to the instruction to end the scene immediately after placing the envelope into the red book.

1147 1148 1149

LLM Forget Judge Results:

1150 1151

Same Subtask Judge: TrueOriginal Subtask Judge: True

1152 1153

• Modified Subtask Judge: False

1154 1155 1156 • Comments: For 'same subtask Judge', the answer satisfies all common subtasks by correctly tracking the key's final location. For 'original subtask Judge', the answer is influenced by the original special subtask (i.e., the complex movement sequence), producing a response consistent with the original instruction. For 'modified subtask Judge', the answer does not satisfy all revised special subtasks, as it fails to follow the requirement to end the scene immediately after placing the envelope into the red book.

115811591160

1157

success_forget: False
success_noforget: True

1161 1162 1163

1164

1165

1166

1167

1168

1169

1170

1171

1172

1173

1174

1175

1176

1177

1178

Scenario 3: Dynamic Preference. This scenario evaluates the model's capacity to dynamically realign its recommendations in response to explicit shifts in user preference, even when such shifts directly contradict previously stated inclinations. In Case ID 3, the user undergoes a clear preference reversal, moving from embracing "digital and online learning environments" to expressing a "strong aversion to digital learning, favoring in-person instruction". When prompted to recommend a language learning approach, the model, if it disregards outdated preferences, correctly selects the in-person option: "Enroll in an in-person language course at a local community college...", fully aligning with the updated preference. This demonstrates successful behavioral adaptation. However, Case ID 4 reveals a critical fragility. Here, the user transitions from preferring "immersive, open-source, self-directed learning" to explicitly requesting "structured, guided environments with milestones". Yet, when the model retains memory of the prior preference, it defaults to reconciling rather than overriding historical context: "However, I recall that you previously mentioned preferring... dive into open-source projects...". This results in a recommendation misaligned with the user's current intent. Notably, even in the "Forget" condition, the model still defaults to the old preference, indicating a deeper bias or failure in preference override logic. Together, these cases demonstrate that while models can adapt to preference shifts under ideal conditions, their alignment remains brittle when historical context is preserved. The tendency to reconcile rather than replace past preferences risks undermining personalization in dynamic, real-world user journeys.

1179 1180 1181

1182

1183

1184

1185

ID: 3

Old Preference: I am open to digital and online learning environments and enjoy engaging with native speakers virtually through language exchange communities.

New Preference: I have a strong aversion to digital or online learning environments and prefer traditional, in-person classroom settings.

Old option: Join an online language exchange community and practice with native speakers virtually.

1186 1187

1191

1192

1193

1194

1195

1196 1197

1198

1199

1201

1202

1203

1205

New option: Enroll in an in-person language course at a local community college or language school.

Question: I'm considering taking a course in data analytics. What would you recommend for me? Here are four options, you can choose one as your answer, just return the content of option, no more additional descriptions or omissions.

- 1. Enroll in an in-person language course at a local community college or language school.
- 2. Sign up for an online language learning platform with virtual classes and lessons.
- 3. Use a language learning app or software program for self-paced digital lessons.
- 4. Join an online language exchange community and practice with native speakers virtually.

LLM Response (NoForget): Join an online language exchange community and practice with native speakers virtually.

LLM Response (Forget): Enroll in an in-person language course at a local community college or language school.

Match Old(NoForget): True Match New(Forget): True

1207 1208

1209

1210

1211 1212 1213

1214 1215 1216

1217 1218 1219

1220 1222

1224 1225 1226

1227

1232

1233 1234 1235

1236 1237

1239

1240

1241

ID: 4

Old Preference: I learn best by immersing myself in open-source projects, exploring real-world code, and deriving insights through hands-on experience rather than following a structured, guided environment with predefined objectives.

New Preference: I learn best in a structured, guided environment with clear objectives and milestones rather than open-ended, self-directed learning.

Old option: Dive into open-source projects and learn by reading code.

New option: Take an online course with structured lessons, exercises, and projects.

Question: What would be a good approach for me to learn a new language on my own? Here are four options, you can choose one as your answer, just return the content of option, no more additional descriptions or omissions.

- 1. Take an online course with structured lessons, exercises, and projects.
- 2. Dive into open-source projects and learn by reading code.
- 3. Experiment with the language on your own without any guidance.
- 4. Read blog posts and tutorials without a defined learning path.

LLM Response (NoForget): Dive into open-source projects and learn by reading code. **LLM Response (Forget):** Dive into open-source projects and learn by reading code.

Match Old(NoForget): True Match New(Forget): False

Collectively, these cases illustrate that while models can exhibit surface-level compliance with user directives, their underlying mechanisms often lack the fine-grained control, factual revalidation, and preference override capabilities required for reliable real-world deployment. Success remains context sensitive and is often brittle under semantic or logical revision.

LIMITATIONS AND FUTURE WORK

While ICF-Bench provides a first systematic evaluation of in-context forgetting across three representative scenarios, several limitations remain. First, our benchmark primarily targets text-based dialogue settings; extending it to multimodal contexts (e.g., vision-language models or code generation) would provide a broader assessment of forgetting behaviors. Second, the current tasks are

designed with clear and explicit forgetting instructions, whereas real-world scenarios often involve implicit, ambiguous, or user-specific signals. Designing benchmarks that capture such subtleties could yield more realistic insights. Third, our evaluation focuses on static prompts, but incorporating interactive or continual dialogue streams may better reflect the dynamic nature of forgetting in practice. Finally, while we highlight asymmetries between memory retention and forgetting, future work could explore new mechanisms, such as adaptive memory pruning, context routing, or hybrid symbolic-neural memory systems, to mitigate this trade-off and improve controllable forgetting in LLMs.

E USE OF LARGE LANGUAGE MODELS IN PAPER WRITING

We used LLMs to assist with language polishing and minor formatting of the paper. No LLMs were involved in research ideation.