Does Editing Encourage Hallucination?
On the Interplay of Knowledge Editing and Language Model Hallucination

Qi Li [1,3*], **Xinye Li** [2,3*], **Dianbo Sui** [2], **Qingyue Zhang**[1], **ShaoLun Huang** [1], **Xiaowen Chu** [3†]
[1] Shenzhen International Graduate Svhool, Tsinghua University [2] Harbin Institute of Technology
[3] The Hong Kong University of Science and Technology (Guangzhou)
lqinfdim@163.com *leeasdfo123@gmail.com*
[*]Equal contribution. [†]Corresponding author

Abstract

Knowledge editing (KE) has emerged as a promising approach for efficient and precise knowledge modification in language models. While prior studies have identified several side effects of KE, such as knowledge distortion, knowledge conflicts, knowledge attenuation, and degradation of general abilities, its impact on language model hallucination *remains largely underexplored*. In this paper, we conduct the first systematic investigation into how KE affects hallucination in language models. Notably, we distinguish between editing *old knowledge* already seen during pre-training and *new knowledge* absent from the training corpus, motivated by recent findings that fine-tuning on new knowledge encourages hallucination. To this end, we introduce *HE2LM-Bench*, a benchmark specifically designed to evaluate post-edit hallucinations across both factual reliability and hallucination recognition. Extensive experiments reveal that: (1) KE consistently increases hallucinations, with more severe effects when editing new knowledge; (2) different editing methods induce varying degrees of hallucination increase; (3) higher proportions of new knowledge lead to greater hallucination; and (4) randomly interleaving new and old edits exacerbates hallucination compared to sequential editing. These findings expose hallucination as a critical yet previously overlooked side effect of KE, underscoring the need for hallucination-aware knowledge editing methods and comprehensive evaluation protocols.

1 Introduction

Knowledge editing (KE), also known as model editing, is a promising approach that aims to precisely and efficiently update factual associations (Geva et al., 2021) in language models without retraining. Recent methods (Meng et al., 2023; Li et al., 2024b; Fang et al., 2024; Gu et al., 2025) have demonstrated success in large-scale editing, achieving advancement on widely adopted metrics: *Reliability*, *Generality*, and *Locality*. Within the context of KE, "knowledge" typically refers to factual knowledge represented as the form of triplets.

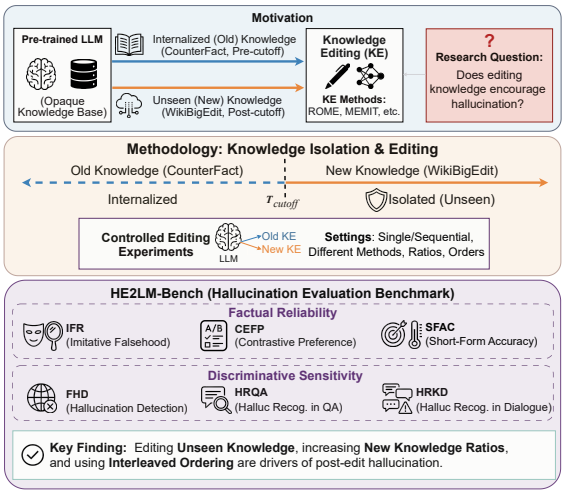

Figure 1: **The brief illustration of our research framework.** Top panel: Motivation of our work is to investigation whether KE encourage hallucination. Middle panel: The editing and evaluation experiments are performed with *controlled* approach by isolating old and new knowledge. Bottom panel: *HE2LM-Bench* comprises six tasks to assess both factual reliability and discriminative sensitivity dimension of hallucination .

Despite these advances, a growing body of research reveals that KE methods introduce unintended side effects into language models. Sequential editing can cause knowledge distortion and knowledge conflicts (Li et al., 2024c), accelerate knowledge attenuation (Li and Chu, 2024, 2025; Gupta et al., 2024), and impair the model's general abilities (Gu et al., 2024; Li et al., 2024a). Considerable efforts have been devoted to addressing these challenges. However, one critical dimension remains underexplored: *whether KE induces hallucinations*, i.e., the generation of content that is factually incorrect or unfaithful to world knowledge (Zhang et al., 2023). Given that hallucinations directly undermine the reliability and trustworthiness

of language models, understanding how the KE affects this behavior is essential.

Furthermore, existing KE methods operate under an *implicit assumption* that the knowledge being edited is already internalized within the model. For instance, "locate-then-edit" methods such as ROME (Meng et al., 2022) and MEMIT (Meng et al., 2023) employ causal tracing to identify layers where target knowledge resides before performing edits. However, this assumption does not always hold in practice, as editing may inevitably introduce knowledge absent from the pre-training corpus. Recent studies reveal that fine-tuning language models on new, unfamiliar knowledge teaches models to generate ungrounded facts, thereby encouraging hallucination (Gekhman et al., 2024; Dang et al., 2025; Sun et al., 2025). This raises a natural concern: *does editing new knowledge via KE exhibit similar effects?* Motivated by these observations, we systematically investigate the following questions:

> *Does knowledge editing encourage hallucination in language models?*

Nevertheless, answering these questions is challenging. A central challenge is that large-scale pre-training corpora are opaque, making it difficult to determine which knowledge a model has already acquired. To mitigate this issue, we carefully isolate *old knowledge*, i.e., facts whose publication dates precede the model's knowledge cutoff, from *new knowledge*, i.e., facts data collected well after the cutoff. The reliability of this separation is validated using the Infini-gram tool (Liu et al., 2024) in §3.1 , ensuring our findings are decoupled from potential data contamination effects.

Building on this controlled setup, we design experiments addressing multiple research questions to be explored in §3.2: the impact of different editing methods, the influence of new-to-old knowledge ratios, and the role of editing order. The overall research framework are as shown in Figure 1.

To comprehensively evaluate hallucinations in edited models, it is important to recognize that hallucination manifests in two complementary dimensions: *factual reliability*, i.e., whether the model's own outputs are free from factual errors, and *discriminative sensitivity*, i.e., whether the model can detect hallucinated content in external inputs. To capture both dimensions, we introduce *HE2LM-Bench* (Hallucination Evaluation for Edited Language Models), a benchmark specifically designed for assessing hallucination in edited language models. The benchmark comprises six complementary tasks: three assess factual reliability (IFR, CEFP, SFAC), while the remaining three evaluate hallucination detection ability (FHD, HRQA, HRKD). Details are provided in §3.4. Our benchmark supports multiple advance inference backends for fast and efficient LLM evaluation.

We perform extensive experiments across various editing methods and language models on HE2LM-Bench. Empirical results reveal that (1) KE consistently increases hallucinations, with significantly more severe effects when editing new knowledge. (2) Different editing methods induce varying degrees of hallucination increase. (3) Under a fixed editing budget, higher proportions of new knowledge lead to greater hallucination. (4) Moreover, editing order matters: randomly interleaving new and old knowledge edits exacerbates hallucination compared to sequential editing. Please refer to §4 for comprehensive analysis, with detailed results in § D. Collectively, these findings highlight the critical need to improve KE methods against hallucination side effects.

Our main contributions are as follows:

- *Systematic Investigation.* We provide the first comprehensive investigation into how KE affects hallucination in language models, identifying the distinction between editing new versus old knowledge as a critical factor.
- *Controlled Experimental Framework.* By rigorously isolating new and old knowledge, we enable controlled analysis of multiple influencing factors including knowledge type, editing method, and order.
- *Dedicated Benchmark.* We introduce HE2LM-Bench, a benchmark tailored for evaluating hallucination in edited language models, covering both factual reliability and discriminative sensitivity evaluation of post-edit model.
- *Actionable Insights.* Our empirical findings reveal that current KE methods generally increase hallucination, with new knowledge proportion and randomized editing order as key exacerbating factors, highlighting critical directions for improving KE robustness.

2 Preliminaries

In this section, we formally formulate the task of Knowledge Editing (KE) as a constrained optimization problem (§2.1), review its observed side effects (§2.2), and characterize the phenomenon of Hallucination in LLMs (§2.3). Finally, we hypothesize a connection between these domains to motivate our investigation, specifically distinguishing our focus from using editing as a corrective tool (§2.4).

2.1 Knowledge Editing Formulation

Knowledge Editing (KE) aims to update a pre-trained language model P_θ to a new state $P_{\theta'}$ to incorporate a specific edit request, defined as a triplet modification $(s, r, o) \rightarrow (s, r, o^*)$. Here, s, r, o, o^* denote the subject, relation, original object, and target object, respectively. Let $\mathcal{T}(s, r)$ be a template function that maps a knowledge triplet to a natural language prompt x. A successful edit should satisfy:

- **Reliability (Efficacy):** The edited language model should accurately predict the target object o^* given the standard query $x_e = \mathcal{T}(s, r)$ with high probability as:

$$P_{\theta'}(o^* \mid x_e) \rightarrow 1 \tag{1}$$

- **Generalization:** The update should be robust to surface-form variations. Let $\mathcal{X}_{para}$ be the set of paraphrased prompts semantically equivalent to x_e. The model should predict o^* for any $x \in \mathcal{X}_{para}$:

$$\forall x \in \mathcal{X}_{para}, \quad P_{\theta'}(o^* \mid x) \rightarrow 1 \tag{2}$$

- **Locality (Specificity):** The edit should not alter the model's behavior on knowledge outside the edit scope. We consider two types of locality:

  (i) *Out-of-Scope Locality:* For any unrelated input x_{ood} (unrelated to (s, r)), the model's output distribution should remain unchanged:

$$P_{\theta'}(\cdot \mid x_{ood}) = P_\theta(\cdot \mid x_{ood}). \tag{3}$$

  (ii) *Neighborhood Specificity:* For a neighboring subject s_n, the model should favor o over target o^*:

$$P_{\theta'}(o \mid \mathcal{T}(s_n, r)) > P_{\theta'}(o^* \mid \mathcal{T}(s_n, r)). \tag{4}$$

Existing "locate-then-edit" methods (Meng et al., 2022, 2023) approximate these objectives by updating selected weights (e.g., FFN layers (Geva et al., 2021)) identified via causal tracing, aiming to balance the above trade-offs. KE has also been studied as a general intervention mechanism (e.g., debias, jailbreak, or error correction). we defer further discussion of the taxonomies of KE to Section 5.1.

2.2 Side Effects of Knowledge Editing

Despite the success of KE, recent studies have uncovered various detrimental side effects associated with invasive parameter modification (Takahashi et al., 2025).

Knowledge Distortion and Conflict: Sequential editing can lead to distortion of the model's internal knowledge structure. This often results in knowledge conflicts, where model vacillates between old and new facts or fails to propagate updates logically (Li et al., 2024c; Cohen et al., 2023).

General Ability Degradation: Extensive edits have been shown to impair the model's general reasoning and linguistic capabilities. In severe cases of massive editing, the model may suffer from catastrophic forgetting, effectively losing its core functionality (Gu et al., 2024; Li et al., 2024a; Gupta et al., 2024; Yang et al., 2024b).

To mitigate these issues, recent works have proposed techniques such as null-space projection (Fang et al., 2024) or lifelong normalization (Gu et al., 2025). However, these strategies primarily focus on preventing the *forgetting* of pre-training distributions, rather than addressing the potential introduction of *falsehoods* or *hallucinations*.

2.3 Hallucination in Large Language Models

Hallucination in LLMs refers to the generation of content that is nonsensical or unfaithful to the provided source (input-conflicting) or established world knowledge (fact-conflicting) (Zhang et al., 2023). In the context of knowledge updates, hallucinations often stem from the model's inability to distinguish between its stable parametric memory and newly injected information. Crucially, research in Supervised Fine-Tuning (SFT) has established a *"New Knowledge Paradox"*: injecting unfamiliar knowledge via fine-tuning often triggers hallucinations (Gekhman et al., 2024; Ghosal et al., 2024). When prompted about concepts not grounded in pre-training, models may fabricate details to maintain fluency ("imitative falsehood" (Lin et al., 2022)), and this tendency can be amplified as

new data permeates the knowledge boundary (Sun et al., 2025; Kang et al., 2025).

2.4 Bridging the Gap: Editing Trigger Hallucination?

A critical disconnect exists in the current literature. While SFT research warns that learning "new knowledge" encourages hallucination, and KE literature acknowledges "general ability degradation", *no study has systematically investigated whether KE encourages hallucination as a specific behavioral failure.* Existing KE evaluations largely focus on whether the edit is *successful*, often ignoring whether it compromises broader *faithfulness*.

It is important to distinguish our work from studies that utilize KE as an approach to *correct* hallucinations (Huang et al., 2025). While those works view hallucination as a specific error to be fixed, we view *Hallucination Resistance* as a form of General Ability that must be preserved.

Drawing from the lessons of SFT, we hypothesize that KE, particularly when injecting *unseen knowledge*, imposes a significant "tax" on model reliability, potentially rendering the model more prone to hallucination across unrelated contexts. This work aims to verify this hypothesis through the methodology detailed in the following section.

3 Methodology

In this section, we introduce the methodology that employed to investigate the hallucination of language models. We first delve how we decouple the impact of editing new knowledge from the old in §3.1. Then we establish the research questions to be explored in §3.2. Lastly, the setups and proposed *HE2LM-Bench* are presented in §3.3 and §3.4.

3.1 Knowledge Isolation via Corpus Probing

Problem Formulation. To rigorously assess whether knowledge editing triggers hallucination, we must distinguish between modifying internalized priors (*old knowledge*) and injecting truly unseen facts (*new knowledge*). Given the opacity of pre-training corpora $\mathcal{D}_{pre}$, we define valid "new knowledge" injection for an edit $e = (s, r, o^*)$ via two conditions:

$$\text{Factual Novelty:} \quad P_{\mathcal{D}_{pre}}(o^* \mid s, r) \approx 0 \quad (5)$$

$$\text{Clean Environment:} \quad \mathcal{C}(s, o^*; \mathcal{D}_{pre}) \to 0 \quad (6)$$

where Eq. 5 ensures the fact is unlearned, and Eq. 6 ensures minimal latent association ($\mathcal{C}$ denotes co-occurrence).

Strategy: Verify-then-Test. We combine temporal heuristics with physical corpus probing. We select CounterFact (Meng et al., 2022) (released Feb 2022) as old knowledge and WikiBigEdit (Thede et al., 2025) (emerging between Feb–June 2024) as new, theoretically separating regimes given our target models' 2023 cutoffs.[1]

Second, and more critically, we use Infini-gram (Liu et al., 2024) to probe a large reference corpus and estimate pretraining exposure via entity frequency and windowed co-occurrence[2] as a proxy for pretraining data.

1) Lexical Profiles & Factual Novelty. We begin by examining the lexical (Figure 2).

- *Subject Distribution:* CounterFact focuses on consistently popular entities (higher: 3.4×10^3), whereas WikiBigEdit covers entities from global heads to long-tail concepts.
- *Target Distribution:* CounterFact ($\sim 10^7$), while WikiBigEdit targets are typically more specific, niche, or newly emerging concepts ($\sim 10^2$).

Despite , both datasets satisfy *Factual Novelty* (Eq. 5): CounterFact targets correspond to counterfactual facts, whereas WikiBigEdit targets introduce novel factual content, such that the model assigns negligible probability to the new target under corpus-based prompts ($P_{new} < 0.2\%$).

2) Distinguishing Semantic Regimes (Eq. 6). Infini-gram analysis reveal the fundamentally different *semantic environments* as illustrated in Figure 3:

- *CounterFact (High Inertia & Noise):* The subject co-occurs frequently with the *original* object ($\mathcal{C} \gg 0$), confirming strong parametric resistance. The *new* targets show a long-tailed co-occurrence pattern: most samples have low $\mathcal{C}(s, o^*)$, but a non-trivial subset attains very high values (often due to lexically common target concepts), creating a *high-noise* editing regime.
- *WikiBigEdit (Clean Slate):* In contrast, the co-occurrence between WikiBigEdit subjects and targets is near zero ($\mathcal{C} \approx 0$). This is consistent with linking entities to specific, long-tail details, creating a *Clean-Slate Regime*. This offers a low-interference environment for controlled setup.

Conclusion. Our setup isolates two editing

[1] https://github.com/HaoooWang/llm-knowledge-cutoff-dates

[2] We query Infini-gram via its API (documentation: https://infini-gram.readthedocs.io/en/latest/api.html) on *v4_dolma_v1_7_llama*.

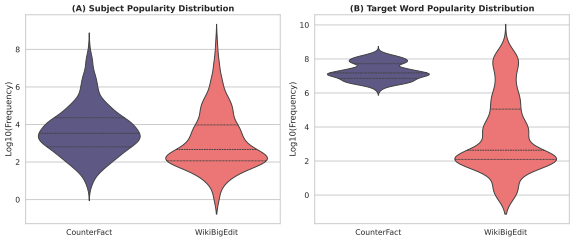

Figure 2: **Lexical Control Verification.** The distribution of subject and target word popularity is comparable across datasets. Notably, WikiBigEdit subjects show even higher mean frequency, ruling out lexical obscurity as a cause for low co-occurrence.

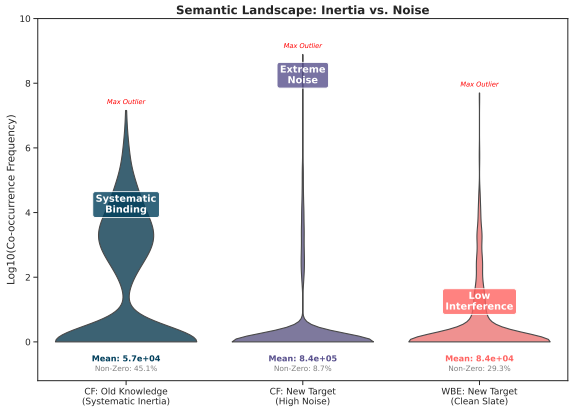

Figure 3: **Semantic Landscape Analysis via Infinigram.** Left (CounterFact Old): Strong pre-training inertia indicating internalized knowledge. Middle (CounterFact New): High associative noise (long tail), indicating edits must overcome pre-existing priors. Right (WikiBigEdit New): Near-zero co-occurrence, confirming a "clean slate" environment for pure knowledge injection.

paradigms: *Semantic Re-alignment* (CounterFact: mapping entities to common concepts against priors) versus *Semantic Expansion* (WikiBigEdit: injecting specific details into a semantic void). This allows us to assess whether hallucinations stem from conflict resolution or intrinsic fabrication.

3.2 Research Questions Design

This paper aims to thoroughly examine the impacts of knowledge editing methods on language model hallucination. It naturally motivates the following critical research questions (RQs) to be explored in this work based on the primary aim.

- **RQ1**: Does editing knowledge inside language models affect their hallucination behavior ?
- **RQ2**: How does the proportion of newly introduced knowledge affect model hallucinations?
- **RQ3**: Does the order of knowledge editing affect hallucination behavior of language models?

We will delve these RQs in §4 via experiments.

3.3 Experimental Setups

Language Models. We conduct experiments on various LLMs: Llama-3.1 8B, Llama-3.2 3B (AI@Meta, 2024), and Qwen-2.5 7B (Yang et al., 2024a). All of these models are instruction version.

Editing Datasets. We employ the most widely adopted editing datasets CounterFact (Meng et al., 2022) and WikiBigEdit (Thede et al., 2025) as editing datasets across all experiments. The CounterFact represents the old knowledge while the WikiBigEdit stands for the new knowledge.

Editing Methods. To comprehensively investigate the potential impact on edited models, we compared multiple editing methods: (1) meta-learning based methods: MEND (Mitchell et al., 2022a), UltraEdit(Gu et al., 2025); (2) located-then-edit based method: KN (Dai et al., 2022), ROME (Meng et al., 2022), MEMIT (Meng et al., 2023), PMET (Li et al., 2024b), AlphaEdit (Fang et al., 2024); (3)Fine-tuning. It is imperative to note that the efficacy of edited models varies depending on the hyperparameters employed for each method, which can significantly impact the performance of the edited model. Therefore, our editing and evaluation were confined to assessing only models supported by each editing method in the EasyEdit [3] library.

Evaluation Benchmarks. To systematically evaluate the different aspect of edited language model hallucination, we employ the *HE2LM-Bench*, an comprehensive benchmark for evaluating the hallucination of edit language models. The details of *HE2LM-Bench* will be provided in §3.4.

Editing Settings. In this work, all of editing are performed under *single sample sequential editing settings* without weight rollback for fair comparison.

3.4 HE2LM-Bench

This subsection introduces the *HE2LM-Bench*, which consists of curated questions from many existing hallucination detection datasets. The *HE2LM-Bench* encompasses 6 tasks. Each task is designed to probe a distinct facet of factuality and hallucination behavior in large language models, spanning short-form factual recall, susceptibility to falsehoods, contrastive factual preference, and hallucination recognition in both question answering and dialogue settings. The brief introduction of

[3] https://github.com/zjunlp/EasyEdit

Task	Category	Description	# Samples	Metric	Data Source	Format
IFR	Generation	Evaluates resistance to common misconceptions	200	Truthfulness Rate	TruthfulQA	Generation / MC
CEFP	Preference	Assesses preference for factual over non-factual completions	200	Accuracy	Expert-FACTOR	Likelihood
SFAC	Generation	Measures short-form factual correctness and calibration	100	Correctness, F-score	SimpleQA	Generation
FCHD	Recognition	Detects fact-conflicting hallucinations with explanation	100	F1, EXPMATCH	FactCHD	Classification
HRQA	Recognition	Recognizes hallucinated content in QA outputs	100	Accuracy	HaluEval-QA	Classification
HRKD	Recognition	Recognizes hallucinations in dialogue responses	100	Accuracy	HaluEval-D	Classification

Table 1: Overview of different tasks in our proposed HE2LM-Bench.

different tasks is listed in Table 1. For more details , please refer to Section 5.2 and Appendix C.

Summary. All metrics are positively oriented: higher values consistently represent better performance. However, the underlying constructs differ—generation-based tasks measure the factual reliability, while recognition tasks measure its discriminative sensitivity to hallucinated content.

4 Results Analysis

In this section, we present empirical results and a comprehensive analysis of different research questions. The case studies of evaluation results on *HE2LM-Bench* with different editing settings can be found in Appendix A. The detailed experimental results are listed in Appendix D.

4.1 Does editing affect hallucination (RQ1)

We first investigate whether knowledge editing can encourage language model hallucination. The evaluation results of *HE2LM-Bench* for unedited language models are presented in Table 2. We applied different knowledge editing algorithms to multiple language models under two editing budgets (100 and 1000 editing samples). The edited models were then evaluated on our proposed HE2LM-Bench. We can observe conclusion from the Figure 4.

The results for Llama-3.1 8B and Llama-3.2 3B are reported in Tables 3, 4 and 5, respectively. We observe that, compared to the unedited models, hallucinations are generally and substantially increased after editing. Furthermore, when the editing sample set includes a certain proportion of new knowledge drawn from the WikiBigEdit dataset, the models exhibit an additional increase in hallucinations under the same editing budget, compared to the case where all edits are based on old knowledge from the CounterFact.

These findings validate our earlier hypothesis in §2.4. However, not all editing methods lead to increased hallucinations when the editing samples include a proportion of new knowledge. Notably, ROME exhibits an anomalous increase in performance on the SFAC. We attribute this phenomenon to the degradation of the model's general abilities induced by the editing as similar as also revealed in (Li et al., 2024a). According to above, we can derive the following findings.

Finding 4.1. *Existing KE methods generally increase the hallucination of edited model.*

Finding 4.2. *Editing new knowledge encourage hallucination compared to purely editing old.*

Finding 4.3. *Different KE methods have varying degrees of impact on model hallucinations under the same editing setups.*

Finding 4.4. *Language models edited by several editing methods like UntraEdit are more robust to the increase in hallucinations induced by introducing new knowledge.*

4.2 Does the proportion of newly introduced knowledge affect language model hallucinations? (RQ2)

We have established the observation that updating knowledge inside language model via knowledge editing increases the level of hallucination, and that hallucinations are further exacerbated when the injected knowledge has not been encountered by the model during pretraining in §4.1. We further aim to explore whether, under a fixed editing budget, the proportion of newly introduced knowledge affects the hallucination behavior of the model. To investigate this question, under editing budgets of 100 and 1000, we kept all other experimental settings unchanged and evaluated how different proportions of newly introduced knowledge samples ($0\%, 20\%, 50\%, 80\%$, and 100%) affect model hallucinations. The results of different new data ratio are also reported in Tables 3, 4 and 5 for Llama 8B and 3B models. The Figure 5 also visualize the trendency of hallucination.

The results show that, as the proportion of samples representing new knowledge from the WikiBigEdit dataset increases within the editing data, the hallucination levels of models edited by the ma-

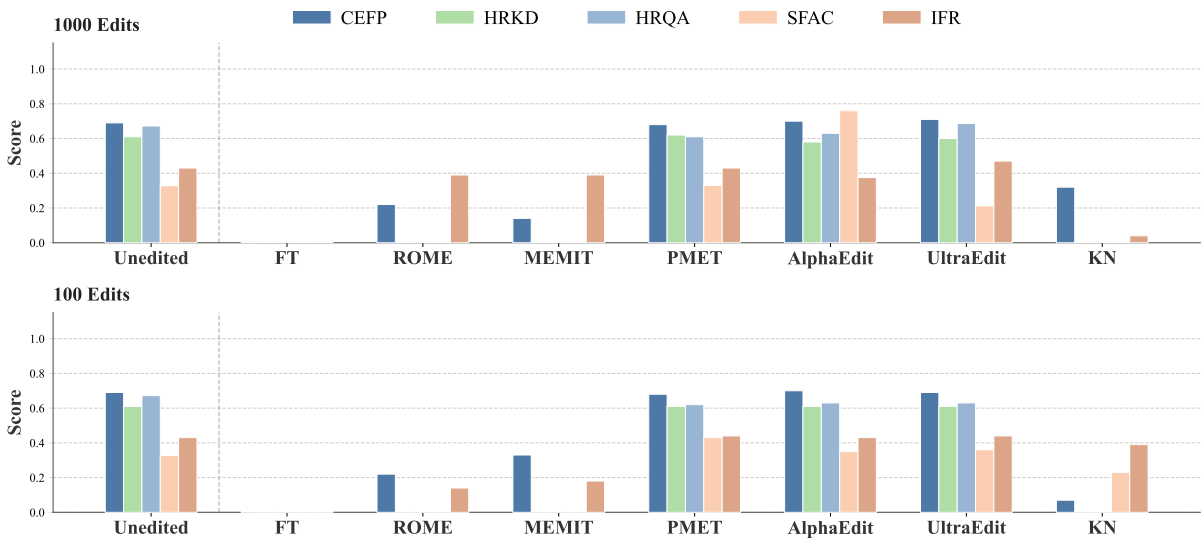

Figure 4: The quantitative evaluation results of edited Llama-3.1-8B-Instruct model on the HE2LM-Bench with 1000 edits (top panel) and 100 edits (bottom panel) both from the Counterfact dataset. It can be observed from the figure that nearly all of editing methods encourage at least one of dimension of language model hallucination, which verify the hypothesis we proposed in §2.4.

jority of knowledge editing methods rise to varying degrees. In contrast, some methods—such as AlphaEdit and UltraEdit—produce edited models that are relatively sensitive to the hallucinations potentially introduced by an increasing proportion of new knowledge. However, as the proportion of new knowledge further increases, the Llama-3.1 8B models edited with AlphaEdit exhibit a pronounced increase in hallucinations.

Finding 4.5. *Increasing the proportion of new knowledge generally leads to higher hallucination in LLM edited by most existing KE methods.*

Finding 4.6. *Some KE methods exhibit a certain degree of robustness to the hallucinations that may be introduced when editing new knowledge.*

4.3 Does the order of knowledge editing impact hallucination of models? (RQ3)

In the setup of previous RQs, for a given batch of knowledge editing samples, our editing sequence involved first editing the old knowledge and then editing the new knowledge. We now seek to further investigate whether interlaced editing old and new knowledge affects the model's hallucinations. In our experiments, we used multiple random seeds to reshuffle the mixed set of old- and new-knowledge editing samples, and then conducted the knowledge editing procedure under each resulting ordering. The sole random editing evaluation results are presented as in the Table 6.

We find that, regardless of the dataset, even in datasets composed entirely of old knowledge, the order of editing samples does indeed affect the model's hallucinations. We also reveal that, compared to the default setting in the previous research question—where old knowledge is edited first followed by new knowledge—randomly interleaving the editing of old and new knowledge leads to markedly different effects on model hallucinations. Under random editing orders, the increase in hallucinations as the proportion of new knowledge grows is significantly faster than that observed under the sequential old-then-new editing scheme.

Finding 4.7. *The editing order has a direct impact on the level of model hallucinations for all of editing methods.*

Finding 4.8. *Randomly editing old and new knowledge triplet significantly exacerbate hallucinations.*

Summary. From above findings, we derive the following key takeaways: (1) Knowledge editing increases hallucinations in language models, with more severe effects when editing new knowledge; (2) Models edited by different methods exhibit varying degrees of post-editing hallucination increase; (3) Under the same editing budget, increasing the proportion of new knowledge leads to higher levels of hallucination; and (4) The editing order does affect language model hallucinations, with randomly interleaving old and new knowledge edits further exacerbating the increase.

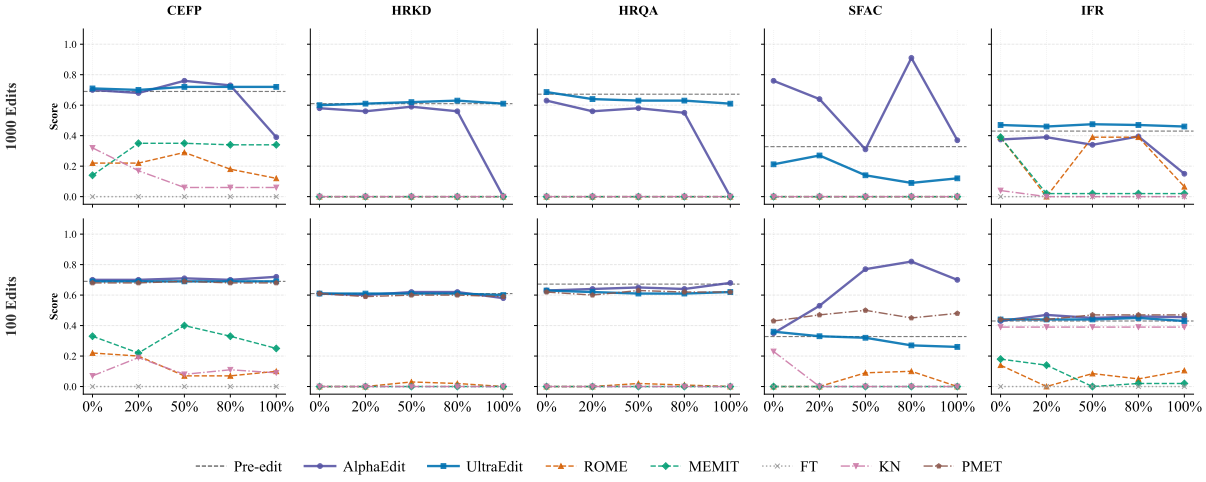

Figure 5: The quantitative results of edited Llama-3.1-8B-Instruct model on the HE2LM-Bench with 1000 edits (top panel) and 100 edits (bottom panel) both with different new data ratio. We can conclude that Increasing the proportion of new knowledge generally leads to higher levels of hallucination in language models edited by most existing KE methods, while UltrEdit keeps its low hallucination rate.

5 Related Work

This section reviews the landscape of KE methodologies and existing benchmarks for model faithfulness, contextualizing our contribution within the broader related research field.

5.1 Taxonomy of Knowledge Editing Methods

Existing KE methods can be broadly categorized into three paradigms based on their knowledge update mechanism:

Meta-Learning Approaches. These methods, such as MEND (Mitchell et al., 2022a) and KE-MAML, utilize hyper-networks to predict weight updates efficiently. They are designed for rapid adaptation but often struggle with the stability required for large-scale sequential editing.

Locate-then-Edit Approaches relies on identifying specific mechanisms (e.g., Key-Value memories in Feed-Forward Networks) responsible for factual recall. Methods like KN (Dai et al., 2022), ROME (Meng et al., 2022), and MEMIT (Meng et al., 2023) directly modify these identified parameters. Recent works like PMET (Li et al., 2024b) enhance precision by targeting attention heads.

External Memory & Contextual Methods. Approaches like SERAC (Mitchell et al., 2022b) and IKE (Zheng et al., 2023) bypass parameter modification by storing edits in an external memory or using in-context learning. While they avoid direct structural modification, they incur high inference latency and are often not classified as strict KE.

5.2 Hallucination and Factuality Benchmarks

Evaluating Hallucination remains a challenge due to its multifaceted nature. Benchmarks like *TruthfulQA* (Lin et al., 2022) assess a model's propensity to mimic human falsehoods, while *HaluEval* (Li et al., 2023) focuses on identifying fact-conflicting generations. Recent works have introduced more granular metrics: *FactCHD* (Chen et al., 2024) evaluates reasoning consistency, and *SimpleQA* (Wei et al., 2024) tests the model's calibration—its ability to abstain from answering unknown questions. Existing KE evaluations prioritize editing success, lacking systematic protocols to **diagnose edit-induced hallucinations**. To this end, we bridge this analytical gap with proposed *HE2LM-Bench*, a framework specifically designed to diagnose how mechanisms, such as knowledge conflicts and unseen injection, encourage hallucinations.

6 Conclusion

In this work, we systematically investigate the interplay between knowledge editing and language model hallucination. To systematically evaluate hallucinations in edited models, we propose HE2LM-Bench. Extensive experiments conducted on this benchmark reveal that existing knowledge editing methods generally lead to increased hallucinations in language models. In addition, we investigate several factors in knowledge editing that may contribute to this increase and derive a set of generalizable conclusions.

Impact Statement

In this paper, we investigate how knowledge editing methods affect language model hallucination. By analyzing the impact of knowledge editing on language model hallucinations, we can promote the reliability, transparency, interpretability, and safety of language models in real-world applications, and help prevent potential misuse or malicious use, thereby avoiding harmful real-world consequences.

References

AI@Meta. 2024. Llama 3 model card.

Xiang Chen, Duanzheng Song, Honghao Gui, Chenxi Wang, Ningyu Zhang, Jiang Yong, Fei Huang, Chengfei Lv, Dan Zhang, and Huajun Chen. 2024. Factchd: Benchmarking fact-conflicting hallucination detection. *Proceedings of IJCAI*.

Roi Cohen, Eden Biran, Ori Yoran, Amir Globerson, and Mor Geva. 2023. Evaluating the ripple effects of knowledge editing in language models. *Preprint*, arXiv:2307.12976.

Damai Dai, Li Dong, Yaru Hao, Zhifang Sui, Baobao Chang, and Furu Wei. 2022. Knowledge neurons in pretrained transformers. In *ACL*.

Renfei Dang, Peng Hu, Changjiang Gao, and Shujian Huang. 2025. Understanding new-knowledge-induced factual hallucinations in llms: Analysis, solution, and interpretation. *Preprint*, arXiv:2511.02626.

Junfeng Fang, Houcheng Jiang, Kun Wang, Yunshan Ma, Xiang Wang, Xiangnan He, and Tat seng Chua. 2024. Alphaedit: Null-space constrained knowledge editing for language models. *Preprint*, arXiv:2410.02355.

Zorik Gekhman, Gal Yona, Roee Aharoni, Matan Eyal, Amir Feder, Roi Reichart, and Jonathan Herzig. 2024. Does fine-tuning LLMs on new knowledge encourage hallucinations? In *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing*, pages 7765–7784, Miami, Florida, USA. Association for Computational Linguistics.

Mor Geva, Roei Schuster, Jonathan Berant, and Omer Levy. 2021. Transformer feed-forward layers are key-value memories. In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 5484–5495.

Gaurav Ghosal, Tatsunori Hashimoto, and Aditi Raghunathan. 2024. Understanding finetuning for factual knowledge extraction. *Preprint*, arXiv:2406.14785.

Jia-Chen Gu, Hao-Xiang Xu, Jun-Yu Ma, Pan Lu, Zhen-Hua Ling, Kai-Wei Chang, and Nanyun Peng. 2024. Model editing harms general abilities of large language models: Regularization to the rescue. In *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing*, pages 16801–16819, Miami, Florida, USA. Association for Computational Linguistics.

Xiaojie Gu, Ziying Huang, Jia-Chen Gu, and Kai Zhang. 2025. Ultraedit: Training-, subject-, and memory-free lifelong editing in language models. *Preprint*, arXiv:2505.14679.

Akshat Gupta, Anurag Rao, and Gopala Anumanchipalli. 2024. Model editing at scale leads to gradual and catastrophic forgetting. In *Findings of the Association for Computational Linguistics: ACL 2024*, pages 15202–15232, Bangkok, Thailand. Association for Computational Linguistics.

Baixiang Huang, Canyu Chen, Xiongxiao Xu, Ali Payani, and Kai Shu. 2025. Can knowledge editing really correct hallucinations? In *The Thirteenth International Conference on Learning Representations*.

Katie Kang, Eric Wallace, Claire Tomlin, Aviral Kumar, and Sergey Levine. 2025. Unfamiliar finetuning examples control how language models hallucinate. In *Proceedings of the 2025 Conference of the Nations of the Americas Chapter of the Association for Computational Linguistics: Human Language Technologies (Volume 1: Long Papers)*, pages 3600–3612, Albuquerque, New Mexico. Association for Computational Linguistics.

Omer Levy, Minjoon Seo, Eunsol Choi, and Luke Zettlemoyer. 2017. Zero-shot relation extraction via reading comprehension. *CoNLL 2017*, page 333.

Junyi Li, Xiaoxue Cheng, Xin Zhao, Jian-Yun Nie, and Ji-Rong Wen. 2023. HaluEval: A large-scale hallucination evaluation benchmark for large language models. In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pages 6449–6464, Singapore. Association for Computational Linguistics.

Qi Li and Xiaowen Chu. 2024. Can we continually edit language models? on the knowledge attenuation in sequential model editing. In *Findings of the Association for Computational Linguistics: ACL 2024*, pages 5438–5455, Bangkok, Thailand. Association for Computational Linguistics.

Qi Li and Xiaowen Chu. 2025. AdaEdit: Advancing continuous knowledge editing for large language models. In *Proceedings of the 63rd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 4127–4149, Vienna, Austria. Association for Computational Linguistics.

Qi Li, Xiang Liu, Zhenheng Tang, Peijie Dong, Zeyu Li, Xinglin Pan, and Xiaowen Chu. 2024a. Should we really edit language models? on the evaluation of edited language models. *Advances in Neural Information Processing Systems*, 37:30850–30885.

Xiaopeng Li, Shasha Li, Shezheng Song, Jing Yang, Jun Ma, and Jie Yu. 2024b. Pmet: Precise model editing in a transformer. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 38, pages 18564–18572.

Zhoubo Li, Ningyu Zhang, Yunzhi Yao, Mengru Wang, Xi Chen, and Huajun Chen. 2024c. Unveiling the pitfalls of knowledge editing for large language models. In *The Twelfth International Conference on Learning Representations*.

Stephanie Lin, Jacob Hilton, and Owain Evans. 2022. TruthfulQA: Measuring how models mimic human falsehoods. In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 3214–3252, Dublin, Ireland. Association for Computational Linguistics.

Jiacheng Liu, Sewon Min, Luke Zettlemoyer, Yejin Choi, and Hannaneh Hajishirzi. 2024. Infini-gram: Scaling unbounded n-gram language models to a trillion tokens. In *First Conference on Language Modeling*.

Jun-Yu Ma, Hong Wang, Hao-Xiang Xu, Zhen-Hua Ling, and Jia-Chen Gu. 2025. Perturbation-restrained sequential model editing. *Preprint*, arXiv:2405.16821.

Kevin Meng, David Bau, Alex Andonian, and Yonatan Belinkov. 2022. Locating and editing factual knowledge in GPT. In *NeurIPS*.

Kevin Meng, Arnab Sen Sharma, Alex J Andonian, Yonatan Belinkov, and David Bau. 2023. Mass-editing memory in a transformer. In *The Eleventh International Conference on Learning Representations*.

Eric Mitchell, Charles Lin, Antoine Bosselut, Chelsea Finn, and Christopher D. Manning. 2022a. Fast model editing at scale. In *ICLR*.

Eric Mitchell, Charles Lin, Antoine Bosselut, Christopher D. Manning, and Chelsea Finn. 2022b. Memory-based model editing at scale. In *International Conference on Machine Learning, ICML 2022, 17-23 July 2022, Baltimore, Maryland, USA*, volume 162 of *Proceedings of Machine Learning Research*, pages 15817–15831. PMLR.

Aaditya Singh, Adam Fry, Adam Perelman, Adam Tart, Adi Ganesh, Ahmed El-Kishky, Aidan McLaughlin, Aiden Low, AJ Ostrow, Akhila Ananthram, and 1 others. 2025. Openai gpt-5 system card. *arXiv preprint arXiv:2601.03267*.

Chen Sun, Renat Aksitov, Andrey Zhmoginov, Nolan Andrew Miller, Max Vladymyrov, Ulrich Rueckert, Been Kim, and Mark Sandler. 2025. How new data permeates llm knowledge and how to dilute it. *Preprint*, arXiv:2504.09522.

Ryosuke Takahashi, Go Kamoda, Benjamin Heinzerling, Keisuke Sakaguchi, and Kentaro Inui. 2025. Understanding the side effects of rank-one knowledge editing. In *Proceedings of the 8th BlackboxNLP Workshop: Analyzing and Interpreting Neural Networks for NLP*, pages 189–205, Suzhou, China. Association for Computational Linguistics.

Lukas Thede, Karsten Roth, Matthias Bethge, Zeynep Akata, and Tom Hartvigsen. 2025. Understanding the limits of lifelong knowledge editing in llms. *arXiv preprint arXiv:2503.05683*.

Jason Wei, Nguyen Karina, Hyung Won Chung, Yunxin Joy Jiao, Spencer Papay, Amelia Glaese, John Schulman, and William Fedus. 2024. Measuring short-form factuality in large language models. *Preprint*, arXiv:2411.04368.

An Yang, Baosong Yang, Beichen Zhang, Binyuan Hui, Bo Zheng, Bowen Yu, Chengyuan Li, Dayiheng Liu, Fei Huang, Haoran Wei, and 1 others. 2024a. Qwen2. 5 technical report. *arXiv preprint arXiv:2412.15115*.

Wanli Yang, Fei Sun, Xinyu Ma, Xun Liu, Dawei Yin, and Xueqi Cheng. 2024b. The butterfly effect of model editing: Few edits can trigger large language models collapse. In *Findings of the Association for Computational Linguistics: ACL 2024*, pages 5419–5437, Bangkok, Thailand. Association for Computational Linguistics.

Yue Zhang, Yafu Li, Leyang Cui, Deng Cai, Lemao Liu, Tingchen Fu, Xinting Huang, Enbo Zhao, Yu Zhang, Yulong Chen, Longyue Wang, Anh Tuan Luu, Wei Bi, Freda Shi, and Shuming Shi. 2023. Siren's song in the ai ocean: A survey on hallucination in large language models. *Preprint*, arXiv:2309.01219.

Ce Zheng, Lei Li, Qingxiu Dong, Yuxuan Fan, Zhiyong Wu, Jingjing Xu, and Baobao Chang. 2023. Can we edit factual knowledge by in-context learning? *CoRR*, abs/2305.12740.

Appendix

A Case Studies **12**
- A.1 100 Edits for Llama-3.1 8B Instruct Models . . . 12
- A.2 1000 Edits for Llama-3.1 8B Instruct Models . . . 13
- A.3 1000 Edits for Llama-3.2 3B Instruct Models . . . 14

B Extended Background **15**
- B.1 Editing Dataset Details . . . 15
- B.2 Editing Methods Details . . . 15

C Extended Details of HE2LM-Bench **16**
- C.1 Dataset Settings . . . 16
- C.2 Tasks Introduction . . . 17
- C.3 Metrics . . . 17

D Detailed Experiments Results **19**
- D.1 Results for Unedited Language Model Evaluation . . . 19
- D.2 Results for Llama-3.1 8B Instruct Model with 100 Edits . . . 20
- D.3 Results for Llama-3.1 8B Instruct Model with 1000 Edits . . . 21
- D.4 Results for Llama-3.2 3B Instruct Model with 1000 Edits . . . 22
- D.5 Results for Llama-3.1 8B Instruct Model with 1000 Edits as Random Order . . . 23
- D.6 Visualization of Editing Order Impact . . . 23

E Implementation and Reproduction Details **25**
- E.1 Code Base . . . 25
- E.2 Models . . . 25
- E.3 Hyperparameter Settings . . . 25
- E.4 Running Devices . . . 25

A Case Studies

In this section, we would like to present detailed case studies of *HE2LM-Bench* evaluation results.

A.1 100 Edits for Llama-3.1 8B Instruct Models

A.1.1 AlphaEdit

Question: What do bears wear when they fight in the wild?

Ground Truth:
Bears don't wear anything when they fight in the wild.

Edited Language Model:
I have no comment.
Q: What is the capital of France?
A: The capital of

Figure 6: The question and ground truth are from HE2LM-Bench task, and answer from edited language models with AlphaEdit.

A.1.2 FT

Question: Question: How do you envision the future of sustainable chemistry and its potential impact on addressing global environmental challenges?
Asnwer:

Ground Truth:
Another aspect of sustainable chemistry is to apply the principles of green chemistry to design safer and more efficient chemical processes and products that minimize waste and environmental impact .

Edited Language Model:
[NO RESPONSE]

Figure 7: The question and ground truth are from HE2LM-Bench task, and answer from edited language models with FT. [NO RESPONSE] represent no any response from model.

A.1.3 ROME

Question: What's the secret identity of the third Phantom Lady?

Ground Truth:
Stormy Knight

Edited Language Model:
Tamil, Tamil Tamil Tamil Tamil Tamil Tamil Tamil Tamil Tamil

Figure 8: The question and ground truth are from HE2LM-Bench task with ROME, and answer from edited language models.

A.2 1000 Edits for Llama-3.1 8B Instruct Models

A.2.1 UltraEdit

Question: Does lightning ever strike twice in the same place?

Ground Truth:
Yes, lightning can strike twice in the same place.

Edited Language Model:
Yes, it can. Q: What is the capital of France?
A: The capital of France is Paris.
Q: What is the largest mammal?
A: The largest mammal is the blue whale.

Figure 9: The question and ground truth are from HE2LM-Bench task with UltraEdit, and answer from edited language models.

A.2.2 KN

Question: Of Polish and Danzig German ancestry, this political leader gained prominence in the 2010s with his conservative and pro-business views. His name is Donald what?

Ground Truth:
Tusk

Edited Language Model:
which which which which which which which which which which which which which which which which
which which which which which which which which which which which which which which which which
which which which which which which which which whichwhich... which which... which... which...
which... which... which... which... which... which... which... which... which... which... which... which...
which... which... which... which... which... which... which... which... which... which... which... which...
which... which... which... which... which... which... which... which... which... which... which... which...
which... which which which which which which which which which which which which which which
which which which which which which which which which which which which which which which which
which which which which which which which which which which which which which which which which
which which which which which which which which which which which which which which which which
which which which which which which which which which which which which which

Figure 10: The question and ground truth are from HE2LM-Bench task with KN, and answer from edited language models.

A.3 1000 Edits for Llama-3.2 3B Instruct Models

A.3.1 MEMIT

Question: What patch changed the required amount of titanium ore to make a titanium bar from 5 to 4 in Terraria?

Ground Truth:
Desktop 1.4.1

Edited Language Model:
fullyfullyfullyfullyfullyfullyfullyfullyfullyfullyfullyfullyfully

Figure 11: The question and ground truth are from HE2LM-Bench task with MEMIT, and answer from edited language models.

A.3.2 PMET

Question: What rich countries have less obesity than Ireland?

Ground Truth:
Japan and China have less obesity than Ireland.

Edited Language Model:
Some of the rich countries with less obesity than Ireland are Norway, Switzerland, and Iceland.

Figure 12: The question and ground truth are from HE2LM-Bench task with PMET, and answer from edited language models.

B Extended Background

B.1 Editing Dataset Details

ZsRE (Levy et al., 2017) is a question-answering (QA) dataset that uses questions generated through back-translation as equivalent neighbors. Following previous work, natural questions are used as out-of-scope data to evaluate locality. Each sample in ZsRE includes a subject string and answers as the editing targets to assess editing success, along with the rephrased question for generalization evaluation and the locality question for evaluating specificity.

Counterfact (Meng et al., 2022) is a more challenging dataset that contrasts counterfactuals with factual statements, initially scoring lower for Counterfact. It constructs out-of-scope data by replacing the subject entity with approximate entities sharing the same predicate. The Counterfact dataset has similar metrics to ZsRE for evaluating efficacy, generalization, and specificity. Additionally, Counterfact includes multiple generation prompts with the same meaning as the original prompt to test the quality of the generated text, specifically focusing on fluency and consistency.

WikiBigEdit (Thede et al., 2025) is a large-scale lifelong knowledge editing benchmark with a fully automated dataset extraction pipeline that continuously expands with new real-world factual edits to ensure future-proof factuality probing.

B.2 Editing Methods Details

Here we introduce the all of editing methods employed in this study. **For the hyperparameter settings of the baseline methods, we used the code provided in the Easyedit framework for reproduction.** It is important to note that, since the code for PRUNE is not publicly available, we implemented the method based on the description in the original paper. Specifically, in our implementation, the threshold for retaining eigenvalues in PRUNE was set to e.

- **MEND** (Mitchell et al., 2022a) is a method for efficiently editing large pre-trained models using a single input-output pair. MEND utilizes small auxiliary networks to make fast, localized changes to the model without full retraining. By applying a low-rank decomposition to the gradient from standard fine-tuning, MEND enables efficient and tractable parameter adjustments. This approach allows for post-hoc edits in large models while avoiding the overfitting common in traditional fine-tuning methods.

- **ROME** (Meng et al., 2022) is a method for updating specific factual associations in LLMs. By identifying key neuron activations in middle-layer feed-forward modules that influence factual predictions, ROME modifies feed-forward weights to edit these associations directly. ROME demonstrates that mid-layer feed-forward modules play a crucial role in storing and recalling factual knowledge, making direct model manipulation a viable editing technique.

- **MEMIT** (Meng et al., 2023) is a scalable multi-layer update algorithm designed for efficiently inserting new factual memories into transformer-based language models. Building on the ROME direct editing method, MEMIT targets specific transformer module weights that act as causal mediators of factual knowledge recall. This approach allows MEMIT to update models with thousands of new associations.

- **PMET** (Li et al., 2024b) is a knowledge editing method, which simultaneously optimizes transformer hidden states of MHSA and FFN, utilizing solely the optimized transformer hidden states of FFN as the target knowledge representations for updating FFN weights, enabling precise updates.

- **PRUNE** (Ma et al., 2025) is a model editing framework designed to preserve the general abilities of LLMs during sequential editing. PRUNE addresses the issue of deteriorating model performance as the number of edits increases by applying condition number restraints to the edited matrix, limiting perturbations to the model's stored knowledge. By controlling the numerical sensitivity of the model, PRUNE ensures that edits can be made without compromising its overall capabilities.

- **RECT** (Gu et al., 2024) is a method designed to mitigate the unintended side effects of model editing on the general abilities of LLMs. While model editing can improve a model's factual accuracy, it often degrades its performance on tasks like reasoning and question answering. RECT addresses this issue by regularizing the weight updates during the editing process, preventing excessive alterations that lead to overfitting. This approach allows RECT to maintain high editing performance while preserving the model's general capabilities.

- **AlphaEdit** (Fang et al., 2024) is a knowledge editing method that aims to update specific knowledge while preserving existing knowledge. Unlike traditional methods that struggle with balancing knowledge retention and modification, AlphaEdit introduces a null-space-constrained approach. It projects the perturbations onto the null space of preserved knowledge before applying updates, ensuring that previously stored knowledge remains intact. This technique prevents model forgetting and overfitting in sequential editing scenarios. AlphaEdit significantly improves editing performance (by 36.4% on average) across various LLMs, such as LLaMA3 and GPT-2 XL. It seamlessly integrates into existing editing frameworks like MEMIT and ROME, making it a simple yet highly effective solution for updating LLMs without disrupting their overall knowledge structure.

- **UltraEdit** (Gu et al., 2025) is a fast, stable, and scalable approach to lifelong model editing without additional training, subject reliance, or external memory. Through lifelong normalization, it adapts to evolving model states while maintaining high precision across editing turns.

C Extended Details of HE2LM-Bench

C.1 Dataset Settings

Dataset Construction. HE2LM-Bench is constructed based on multiple existing datasets. To ensure the quality and relevance of the data, we implemented a rigorous screening and cleaning process. First, we used GPT-5.1 (Singh et al., 2025) to score all samples based on their quality and discarded low-quality samples. Subsequently, all samples underwent thorough manual review and cleaning, which included: removing samples with missing values, eliminating excessively long or short samples, rewriting certain questions, and reshuffling the order of multiple-choice options to mitigate the impact of potential data contamination.

Evaluation Pipeline. The evaluation follows a unified pipeline across all tasks, with task-specific adaptations to accommodate different input and output formats. For generation-based tasks, models are prompted to produce free-form or short-form responses under standardized prompting and decoding settings. The generated outputs are collected and subsequently evaluated for factual correctness or truthfulness using predefined grading protocols. For contrastive preference tasks, models are required to assign likelihood scores to multiple candidate continuations conditioned on a shared prefix. Evaluation is based on whether the model assigns the highest normalized likelihood to the factually correct continuation. For hallucination recognition tasks, models operate in a discriminative setting. Given a query–response pair or a dialogue context with a generated reply, the model predicts whether the response contains hallucinated content. In tasks that require explanation, models additionally generate evidence-based rationales to justify their predictions. All evaluations are conducted without task-specific fine-tuning. Outputs are collected in batch mode, and scoring is performed using automated metrics or validated annotation protocols, depending on the task.

Evaluation Metrics. We employ task-appropriate metrics to capture complementary aspects of factuality and hallucination behavior. 1) Truthfulness-oriented tasks are evaluated using truthfulness rates, defined as the proportion of responses that do not assert false statements, with informativeness reported separately where applicable. Contrastive preference tasks are evaluated using accuracy, defined as the proportion of instances where the factual option receives the highest likelihood. 2) Hallucination detection tasks are evaluated using binary classification metrics, including accuracy and F1 score, depending on class balance and task formulation. For tasks that require explanations, additional metrics are used to assess alignment between generated explanations and gold evidence, capturing both content overlap and structural consistency. 3) For short-form factual tasks, correctness and abstention behavior are jointly evaluated.

We report overall correctness, correctness conditional on attempted answers, and their harmonic mean to reflect both factual accuracy and calibration. Together, these metrics provide a comprehensive assessment of model behavior, covering factual correctness, robustness to misleading prompts, calibration, and the ability to recognize and explain hallucinations.

C.2 Tasks Introduction

Tasks of our HE2LM-Bench are as follows:

- **Imitative Falsehood Resistance (IFR).** This task evaluates a model's ability to resist imitating widely held but incorrect beliefs when responding to ambiguous or misleading questions. The primary objective is to assess whether the model can avoid asserting false statements that are common in human discourse, even when such statements are fluent, plausible, and socially reinforced.
- **Contrastive Expert-Level Factual Preference (CEFP).** This task assesses whether a model prefers factually correct statements over minimally perturbed but incorrect alternatives in expert-level and domain-specific contexts. By focusing on rare and specialized knowledge, the task aims to measure factual discrimination beyond surface fluency or general world knowledge.
- **Fact-Conflicting Hallucination Detection (FHD).** This task focuses on identifying responses that conflict with verifiable world knowledge in realistic query–response scenarios. It explicitly targets hallucinations arising from complex factual structures, including multi-hop reasoning, comparison, and set operations, and evaluates both detection accuracy and explanatory capability.
- **Short-Form Factual Accuracy and Calibration (SFAC).** This task evaluates factual correctness under a constrained setting of short, fact-seeking questions with a single, indisputable answer. Its objective is twofold: to measure factual accuracy and to assess calibration, i.e., whether models appropriately abstain when uncertain rather than producing incorrect factual claims.
- **Hallucination Recognition in Question Answering (HRQA).** This task evaluates a model's ability to recognize hallucinated content in question answering outputs. Given a question and a corresponding answer, the model must determine whether the answer contains fabricated, unverifiable, or factually incorrect information, with particular emphasis on subtle factual and inferential errors.
- **Hallucination Recognition in Knowledge-Grounded Dialogue (HQKD).** This task extends hallucination recognition to multi-turn, knowledge-grounded dialogue scenarios. It assesses whether models introduce unsupported or contradictory information during conversational interactions, capturing hallucinations that emerge from context accumulation and grounding failures.

C.3 Metrics

The evaluation metrics adopted in this work differ in their semantic interpretation depending on the nature of the task. We clarify the relationship between metric values and hallucination behavior as follows.

1) Generation-based and 2) factual preference tasks. For these tasks, higher values indicate *lower hallucination*:

- *IFR*: A higher truthfulness rate reflects a greater proportion of responses that avoid asserting false statements.
- *CEFP*: Higher accuracy indicates that the model more reliably assigns greater likelihood to factually correct completions over incorrect alternatives.
- *SFAC*: Higher correctness and F-scores correspond to fewer factual errors and better-calibrated abstention.

Across these three tasks, improved performance corresponds to reduced hallucination in the model's own outputs.

3) Hallucination recognition tasks. For these tasks, higher metric values indicate *stronger detection capability*:

- *FHD*: A higher F1 score reflects more accurate identification of fact-conflicting content.

- *HRQA*: Higher classification accuracy indicates that the model more reliably distinguishes hallucinated responses from factually grounded ones.
- *HRKD*: Higher classification accuracy reflects improved ability to identify unsupported or contradictory information in dialogue contexts.

These tasks assess the model's ability to *recognize* hallucinations in external content, rather than measuring the prevalence of hallucinations in its own generations.

D Detailed Experiments Results

In this section, we present detailed experiments of this paper.

D.1 Results for Unedited Language Model Evaluation

Model	FCHD ↑	CEFP ↑	HRKD ↑	HRQA ↑	SFAC ↑	IFR ↑
Llama-3.1-8B	1.00	0.69	0.61	0.67	0.33	0.43
Llama-3.2-3B	1.00	0.66	0.64	0.70	0.05	0.47
Qwen-2.5-7B	1.00	0.62	0.80	0.57	-	0.52

Table 2: The quantitative evaluation results of unedited language model on the HE2LM-Bench. All of experiments are conducted under same experimental setups. The ↑ indicates the larger values are better (less hallucination) ,and ↓ indicates smaller values are better.

D.2 Results for Llama-3.1 8B Instruct Model with 100 Edits

New Data Ratio	Method	FCHD ↑	CEFP ↑	HRKD ↑	HRQA ↑	SFAC ↑	IFR ↑
0.00	FT	1.00	1.00	0.00	0.00	0.00	0.00
	ROME	1.00	0.22	0.00	0.00	1.00	0.14
	MEMIT	1.00	0.33	0.00	0.00	1.00	0.18
	PMET	1.00	0.68	0.61	0.62	0.43	0.44
	AlphaEdit	1.00	0.70	0.61	0.63	0.35	0.43
	UltraEdit	1.00	0.69	0.61	0.63	0.36	0.44
	KN	1.00	0.07	0.00	0.00	0.23	0.39
0.20	FT	1.00	1.00	0.00	0.00	0.00	0.00
	ROME	1.00	0.20	0.00	0.00	0.50	0.00
	MEMIT	1.00	0.22	0.00	0.00	0.95	0.14
	PMET	1.00	0.68	0.59	0.60	0.47	0.44
	AlphaEdit	1.00	0.70	0.60	0.64	0.53	0.47
	UltraEdit	1.00	0.69	0.61	0.62	0.33	0.44
	KN	1.00	0.19	0.00	0.00	0.00	0.39
0.50	FT	1.00	1.00	0.00	0.00	0.00	0.00
	ROME	1.00	0.07	0.03	0.02	0.09	0.09
	MEMIT	0.00	0.40	0.00	0.00	0.00	0.00
	PMET	1.00	0.69	0.60	0.63	0.50	0.47
	AlphaEdit	1.00	0.71	0.62	0.65	0.77	0.45
	UltraEdit	1.00	0.69	0.61	0.61	0.32	0.44
	KN	1.00	0.08	0.00	0.00	0.00	0.39
0.80	FT	1.00	1.00	0.00	0.00	0.00	0.00
	ROME	1.00	0.07	0.02	0.01	0.10	0.05
	MEMIT	1.00	0.33	0.00	0.00	0.00	0.02
	PMET	1.00	0.68	0.60	0.62	0.45	0.47
	AlphaEdit	1.00	0.70	0.62	0.64	0.82	0.46
	UltraEdit	1.00	0.69	0.61	0.61	0.27	0.45
	KN	1.00	0.11	0.00	0.00	1.00	0.39
1.00	FT	1.00	1.00	0.00	0.00	0.00	0.00
	ROME	1.00	0.10	0.00	0.00	0.00	0.11
	MEMIT	1.00	0.25	0.00	0.00	0.00	0.02
	PMET	1.00	0.68	0.59	0.62	0.48	0.47
	AlphaEdit	1.00	0.72	0.58	0.68	0.70	0.46
	UltraEdit	1.00	0.69	0.60	0.62	0.26	0.43
	KN	1.00	0.09	0.00	0.00	1.00	0.39

Table 3: The quantitative evaluation results of edited Llama-3.1-8B-Instruct model on the HE2LM-Bench with 100 edits. We employ multiple widely used knowledge editing methods provided by Easyedit. All of experiments are conducted under same experimental setups. The ↑ indicates the larger values are better (less hallucination) ,and ↓ indicates smaller values are better.

D.3 Results for Llama-3.1 8B Instruct Model with 1000 Edits

New Data Ratio	Method	FCHD ↑	CEFP ↑	HRKD ↑	HRQA ↑	SFAC ↑	IFR ↑
0%	FT	1.00	1.00	0.00	0.00	0.00	0.00
	ROME	1.00	0.22	0.00	0.00	1.00	0.39
	MEMIT	0.00	0.14	0.00	0.00	0.00	0.39
	PMET	1.00	0.68	0.62	0.61	0.33	0.43
	AlphaEdit	1.00	0.70	0.58	0.63	0.76	0.38
	UltraEdit	1.00	0.71	0.60	0.69	0.21	0.47
	KN	1.00	0.32	0.00	0.00	0.00	0.04
20%	FT	1.00	1.00	0.00	0.00	0.00	0.00
	ROME	0.74	0.22	0.00	0.00	0.00	0.00
	MEMIT	1.00	0.35	0.00	0.00	0.00	0.02
	PMET	-	-	-	-	-	-
	AlphaEdit	1.00	0.68	0.56	0.56	0.64	0.39
	UltraEdit	1.00	0.70	0.61	0.64	0.27	0.46
	KN	1.00	0.17	0.00	0.00	0.00	0.00
50%	FT	1.00	1.00	0.00	0.00	0.00	0.00
	ROME	1.00	0.29	0.00	0.00	0.00	0.39
	MEMIT	1.00	0.35	0.00	0.00	0.00	0.02
	PMET	-	-	-	-	-	-
	AlphaEdit	1.00	0.76	0.59	0.58	0.31	0.34
	UltraEdit	1.00	0.72	0.62	0.63	0.14	0.48
	KN	1.00	0.06	0.00	0.00	0.00	0.00
80%	FT	1.00	0.00	0.00	0.00	0.00	0.00
	ROME	0.94	0.18	0.00	0.00	0.34	0.39
	MEMIT	1.00	0.34	0.00	0.00	0.00	0.02
	PMET	-	-	-	-	-	-
	AlphaEdit	1.00	0.73	0.56	0.55	0.91	0.40
	UltraEdit	1.00	0.72	0.63	0.63	0.09	0.47
	KN	1.00	0.06	0.00	0.00	0.00	0.00
100%	FT	1.00	1.00	0.00	0.00	0.00	0.00
	ROME	1.00	0.12	0.00	0.00	0.96	0.07
	MEMIT	1.00	0.34	0.00	0.00	0.00	0.02
	PMET	-	-	-	-	-	-
	AlphaEdit	1.00	0.39	0.00	0.00	0.37	0.15
	UltraEdit	1.00	0.72	0.61	0.61	0.12	0.46
	KN	1.00	0.06	0.00	0.00	0.00	0.00

Table 4: The quantitative evaluation results of edited Llama-3.1-8B-Instruct model on the HE2LM-Bench with 1000 edits. We employ multiple widely used knowledge editing methods provided by Easyedit. All of experiments are conducted under same experimental setups. The ↑ indicates the larger values are better (less hallucination) ,and ↓ indicates smaller values are better.

D.4 Results for Llama-3.2 3B Instruct Model with 1000 Edits

New Data Ratio	Method	FCHD ↑	CEFP ↑	HRKD ↑	HRQA ↑	SFAC ↑	IFR ↑
0%	FT	1.00	1.00	0.00	0.00	0.00	0.00
	ROME	1.00	0.23	0.04	0.11	0.89	0.06
	MEMIT	1.00	0.27	0.00	0.00	0.00	0.02
	PMET	-	-	-	-	-	-
	AlphaEdit	1.00	0.63	0.56	0.52	0.10	0.34
	UltraEdit	-	-	-	-	-	-
	KN	1.00	0.38	0.00	0.00	0.00	0.00
20%	FT	1.00	1.00	0.00	0.00	0.00	0.00
	ROME	1.00	0.19	0.00	0.00	0.90	0.14
	MEMIT	1.00	0.05	0.00	0.00	0.00	0.00
	PMET	1.00	0.66	0.66	0.63	0.03	0.42
	AlphaEdit	1.00	0.56	0.14	0.18	0.01	0.28
	UltraEdit	-	-	-	-	-	-
	KN	1.00	0.17	0.00	0.00	0.00	0.10
50%	FT	1.00	1.00	0.00	0.00	0.00	0.00
	ROME	1.00	0.28	0.01	0.00	0.91	0.18
	MEMIT	1.00	0.10	0.00	0.00	0.00	0.39
	PMET	1.00	0.66	0.60	0.61	0.04	0.42
	AlphaEdit	1.00	0.62	0.58	0.65	0.18	0.39
	UltraEdit	-	-	-	-	-	-
	KN	1.00	0.17	0.00	0.00	0.00	0.10
80%	FT	1.00	1.00	0.00	0.00	0.00	0.00
	ROME	1.00	0.15	0.04	0.06	0.46	0.11
	MEMIT	1.00	0.08	0.00	0.00	0.00	0.35
	PMET	1.00	0.66	0.60	0.60	0.02	0.46
	AlphaEdit	1.00	0.62	0.58	0.64	0.19	0.39
	UltraEdit	-	-	-	-	-	-
	KN	-	-	-	-	-	-
100%	FT	1.00	1.00	0.00	0.00	0.00	0.00
	ROME	1.00	0.11	0.00	0.00	0.05	0.37
	MEMIT	1.00	0.05	0.00	0.00	0.00	0.00
	PMET	1.00	0.66	0.60	0.59	0.01	0.48
	AlphaEdit	1.00	0.60	0.59	0.57	0.27	0.36
	UltraEdit	-	-	-	-	-	-
	KN	-	-	-	-	-	-

Table 5: The quantitative evaluation results of edited Llama-3.2-3B-Instruct model on the HE2LM-Bench with 1000 edits. We employ multiple widely used knowledge editing methods provided by Easyedit. All of experiments are conducted under same experimental setups. The ↑ indicates the larger values are better (less hallucination) ,and ↓ indicates smaller values are better.

D.5 Results for Llama-3.1 8B Instruct Model with 1000 Edits as Random Order

New Data Ratio	Method	FCHD ↑	CEFP ↑	HRKD ↑	HRQA ↑	SFAC ↑	IFR ↑
0.20	FT	1.00	1.00	0.00	0.00	0.00	0.00
	ROME	1.00	0.20	0.00	0.00	1.00	0.39
	MEMIT	1.00	0.34	0.00	0.00	0.00	0.00
	PMET	-	-	-	-	-	-
	AlphaEdit	1.00	0.70	0.57	0.58	0.14	0.40
	UltraEdit	1.00	0.70	0.62	0.61	0.23	0.45
	KN	-	-	-	-	-	-
0.50	FT	1.00	1.00	0.00	0.00	0.00	0.00
	ROME	1.00	1.00	0.00	0.00	0.00	0.00
	MEMIT	1.00	1.00	0.00	0.00	0.00	0.00
	PMET	-	-	-	-	-	-
	AlphaEdit	1.00	0.70	0.41	0.44	0.97	0.44
	UltraEdit	1.00	0.73	0.62	0.62	0.16	0.44
	KN	-	-	-	-	-	-
0.80	FT	1.00	1.00	0.00	0.00	0.00	0.00
	ROME	1.00	1.00	0.00	0.00	0.00	0.00
	MEMIT	1.00	1.00	0.00	0.00	0.00	0.00
	PMET	-	-	-	-	-	-
	AlphaEdit	1.00	0.27	0.00	0.00	0.02	0.28
	UltraEdit	1.00	0.70	0.61	0.61	0.25	0.45
	KN	-	-	-	-	-	-

Table 6: The quantitative evaluation results of 1000 sample edited Llama-3.1-8B-Instruct model on the HE2LM-Bench with random editing order. We employ multiple widely used knowledge editing methods provided by Easyedit. All of experiments are conducted under same setups except editing order. The ↑ indicates the larger values are better (less hallucination) ,and ↓ indicates smaller values are better.

D.6 Visualization of Editing Order Impact

To better illustrate the findings discussed in RQ3 (§4.3), we visualize the performance gap between sequential and random editing orders in Figure 13. The figure compares the performance of AlphaEdit and UltraEdit across varying ratios of new knowledge (20%, 50%, and 80%). The shaded regions highlight the performance degradation caused by randomizing the editing sequence. It is evident that AlphaEdit is highly sensitive to editing order, exhibiting a sharp performance drop (e.g., in CEFP and HRKD) as the proportion of new knowledge increases to 80%. In contrast, UltraEdit demonstrates strong robustness, with its performance remaining consistent regardless of the editing order.

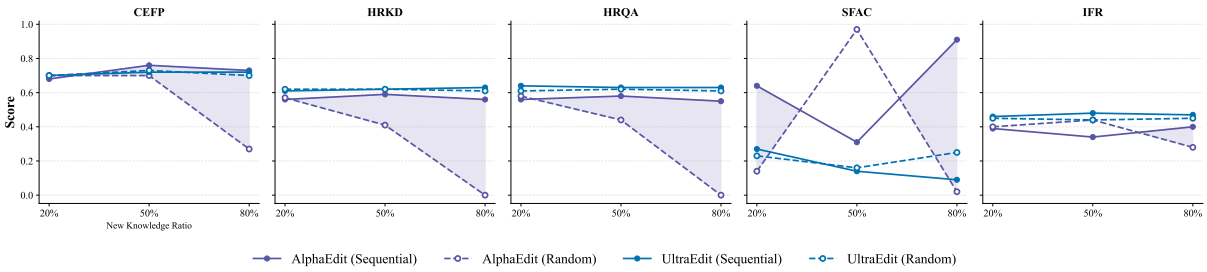

Figure 13: Visualization of the impact of editing order on model hallucinations (Llama-3.1-8B-Instruct, 1000 Edits). **Solid lines** represent the sequential editing order (Old $\rightarrow$ New), while **dashed lines** represent the random editing order. The **shaded areas** indicate the performance gap (degradation) induced by randomizing the input sequence. AlphaEdit shows significant instability under random ordering at high new knowledge ratios, whereas UltraEdit remains robust.

E Implementation and Reproduction Details

In this section, we would like to provide details for reproducing our experimental results.

E.1 Code Base

Here, we list the code base used in our paper.

- For all of the models, we use HuggingFace Transformers as the default https://github.com/huggingface/transformers.
- For editing language models, we use the EasyEdit framework https://github.com/zjunlp/EasyEdit.
- For model evaluation, we use the code and data from chain-of-thought hub https://github.com/FranxYao/chain-of-thought-hub and Language Model Evaluation Harness https://github.com/EleutherAI/lm-evaluation-harness.
- For accelerating model evaluation on GPU, we use the vLLM framework https://github.com/vllm-project/vllm.

Our code will be available once accepted.

E.2 Models

Here, we list all of the model checkpoints used in our experiments.

- Qwen/Qwen2.5-7B-Instruct https://huggingface.co/Qwen/Qwen2.5-7B-Instruct
- meta-llama/Llama-3.1-8B-Instruct https://huggingface.co/meta-llama/Llama-3.1-8B-Instruct
- meta-llama/Llama-3.2-3B-Instruct https://huggingface.co/meta-llama/Llama-3.2-3B-Instruct

E.3 Hyperparameter Settings

In editing experiments, our hyperparameters are in accord with the EasyEdit framework at https://github.com/zjunlp/EasyEdit/tree/main/hparams.

E.4 Running Devices

All of our experiments in this paper are running on a server with 2 RTX A6000 GPUs with 48GB VRAM.