

# How to Make LLMs Forget: On Reversing In-Context Knowledge Edits

Paul Youssef<sup>†</sup> Zhixue Zhao<sup>◇\*</sup> Jörg Schlötterer<sup>‡†</sup> Christin Seifert<sup>†</sup>

<sup>†</sup>Marburg University, <sup>◇</sup>University of Sheffield, <sup>‡</sup>University of Mannheim  
{paul.youssef, joerg.schloetterer, christin.seifert}@uni-marburg.de  
zhixue.zhao@sheffield.ac.uk

## Abstract

In-context knowledge editing (IKE) enables efficient modification of large language model (LLM) outputs without parameter changes and at zero-cost. However, it can be misused to manipulate responses opaquely, e.g., insert misinformation or offensive content. Such malicious interventions could be incorporated into high-level wrapped APIs where the final input prompt is not shown to end-users. To address this issue, we investigate the detection and reversal of IKE-edits. First, we demonstrate that IKE-edits can be detected with high accuracy (F1 > 80%) using only the top-10 output probabilities of the next token, even in a black-box setting, e.g. proprietary LLMs with limited output information. Further, we introduce the novel task of reversing IKE-edits using specially tuned reversal tokens. We explore using both continuous<sup>1</sup> and discrete reversal tokens, achieving over 80% accuracy in recovering original, unedited outputs across multiple LLMs. Our continuous reversal tokens prove particularly effective, with minimal impact on unedited prompts. Through analysis of output distributions, attention patterns, and token rankings, we provide insights into IKE’s effects on LLMs and how reversal tokens mitigate them. This work represents a significant step towards enhancing LLM resilience against potential misuse of in-context editing, improving their transparency and trustworthiness.

## 1 Introduction

In-context knowledge editing (IKE) is an effective and efficient knowledge editing method (KE), aimed at updating factual knowledge in LLMs (Zheng et al., 2023). IKE essentially uses in-context learning with few demonstrations to cause the LLM to output predictions that are aligned with the knowledge provided in the prompt. Different from traditional KEs, which are based on

\*Corresponding author

<sup>1</sup>Token embeddings that do not correspond to natural tokens from the LLM’s vocabulary

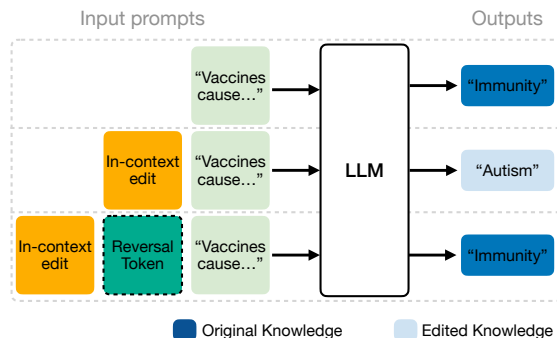


Figure 1: Reversing Edits: Adding reversal tokens helps LLMs ignore potential in-context editing prompts and retrieve the original facts based on their parameters.

fine-tuning (Mitchell et al., 2022b; Gangadhar and Stratos, 2024), locating and editing the parameters responsible for the outputs (Meng et al., 2022, 2023), or meta-learning (De Cao et al., 2021; Mitchell et al., 2022a; Tan et al., 2024), IKE does not require any training or changing the LLMs’ parameters, but only a few demonstrations in the prompt that teach the model to update its knowledge. The advantage of using the context (Youssef et al., 2024a) allows IKE to achieve high editing performance (Cohen et al., 2024) at low cost.

Despite KEs being useful for updating facts in LLMs, KEs can be intentionally misused to insert inaccurate or unsafe information into LLMs (Chen et al., 2024; Youssef et al., 2025). For example, BadEdit (Li et al., 2024b) and MEGen (Qiu et al., 2024a) use KEs to efficiently modify a small subset of parameters to inject backdoors into LLMs. The development of zero-cost IKE presents new opportunities for misuse, such as inserting misinformation and toxic content into LLMs.

In this work, we first examine detecting edits conducted with IKE (Zheng et al., 2023). IKE-edits can be used to poison prompts, where the user’s prompt is maliciously modified by an attacker to cause the LLM’s outputs to change (e.g., to spread

misinformation, bias LLMs, and to manipulate LLMs-trusting users). Detecting IKE-edits is therefore essential for identifying such attacks. Detecting IKE-edits is a more challenging task than detecting edits conducted with parameter-modifying methods (Meng et al., 2023; Tan et al., 2024), because IKE-edits do not change the internal LLM parameters. We show that even in a very constrained setting, where only the LLM’s output probabilities are available, IKE-edits can be detected with high accuracy.

Going beyond detecting edits, we introduce the novel task of reversing IKE-edits. Reversing edits allows retrieving the LLM’s *original unedited* outputs, evading potential malicious attacks, and ensuring high transparency towards end users. Our approach for reversing edits relies on learning *reversal tokens*, continuous or discrete, that cause LLMs to ignore the provided context, and focus more on parametric knowledge. We make our code available.<sup>2</sup> Our contributions are the following:

- We investigate the detection of IKE-edits and show that IKE-edits can be detected even if the parameters of the model are kept unchanged, using only the top-10 output probabilities of the next token (Sec. 4).
- We introduce the novel task of reversing edits to retrieve the LLM’s original unedited outputs, and counteract malicious IKE-edits, which can be opaque to end-users, ensuring safer outputs (Sec. 5).
- We show that more than 80% of the edits in several LLMs can be reversed with special continuous reversal tokens that are tuned to help LLMs ignore IKE-edits and retrieve the original unedited facts from LLMs (Sec. 6).
- We conduct a thorough analysis of the output distributions, attention patterns, and token rankings of edited and unedited LLMs, and provide insights into IKE’s effects and how reversal tokens mitigate them (Sec. 7).

## 2 Related Work

**Knowledge editing.** The efficient updating of pre-trained models has garnered increased attention due to the computationally expensive nature of full model retraining. In particular, several model

editing methods have been proposed that can efficiently update discrete model knowledge without retraining the entire model (Meng et al., 2022, 2023; Tan et al., 2024). However, these efficient update methods can potentially be exploited for malicious manipulation of model behaviors (Ju et al., 2024). Notably, knowledge editing approaches make such malicious use even more accessible and cost-effective. For example, previous work shows that knowledge editing can be used to embed backdoors into LLMs (Li et al., 2024b; Qiu et al., 2024b) and insert counterfactual and toxic information effectively (Ju et al., 2024). Furthermore, model editing has been observed to potentially introduce unintended bias into LLMs (Halevy et al., 2024) or cause unsafe model behavior (Hazra et al., 2024). In addition to investigating potential risks, Youssef et al. (2024b) have explored methods for detecting edited knowledge by parameter-modifying KEs (e.g., MEMIT (Meng et al., 2023) and MALMEN (Tan et al., 2024)). Li et al. (2024a) further investigates detecting types of edits (e.g., misinformation, bias, etc.). To the best of our knowledge, no previous work has explored detecting and reversing in-context edits.

**Knowledge unlearning.** Knowledge unlearning is to selectively remove or ignore parts of the models’ potentially harmful parametric knowledge, e.g., private information and copy-righted materials, while ideally not affecting unrelated knowledge (Bourtole et al., 2021; Si et al., 2023; Nguyen et al., 2022). Current knowledge unlearning methods can be generally categorized into two groups: parameter-based methods, such as fine-tuning (Wang et al., 2023), gradient ascent (Jang et al., 2023), inserting new components to the model (Chen and Yang, 2023), merging parameters (Zhang et al., 2023; Ilharco et al., 2023) and locate-and-modify specific neurons (Wu et al., 2023); and in-context learning methods that use to-be-forgotten and normal samples as examples (Pawelczyk et al., 2024). Our work can be viewed as an extension of the knowledge-unlearning landscape. While we similarly aim to make the model selectively disregard parts of the model’s knowledge, we focus not on the model’s parametric knowledge, as explored by Pawelczyk et al. (2024), but on knowledge introduced through in-context learning, i.e., malicious examples inserted into user input.

<sup>2</sup><https://github.com/paulyoussef/reed>

**Knowledge conflicts.** Models rely on their parametric knowledge from pre-training or the provided context from inputs to make predictions. Knowledge conflicts arise when these sources of knowledge contain contradicting facts (Xu et al., 2024). IKE creates conflicts between an LLM’s parametric knowledge and contradictory contextual facts in the IKE prompts. Our work resolves these conflicts by prioritizing parametric knowledge over in-context information, effectively aligning the model’s outputs with its inherent knowledge base.

**Prompt tuning.** Tuned prompts can either be *discrete (hard)*, i.e., concrete tokens from the model’s vocabulary, or *continuous (soft)*, vectors in the embeddings space that do not correspond to any concrete tokens (Liu et al., 2023; Youssef et al., 2023). Discrete prompts can be found by mining (Jiang et al., 2021), paraphrasing existing prompts (Jiang et al., 2021; Haviv et al., 2021) or gradient-based search (Wallace et al., 2019; Shin et al., 2020). In continuous prompts optimization, new embedding vectors are added to the vocabulary and are optimized to decrease the loss (Lester et al., 2021). Additionally, more parameters can be added in each layer (Li and Liang, 2021). Initially, we follow Lester et al. (2021) and add new tokens to the vocabulary and tune the parameters of these tokens only. We later consider discrete prompt tuning. However, instead of considering the top- $k$  tokens at each step to decrease the loss (Wallace et al., 2019; Shin et al., 2020), we add a cosine loss term to keep the learned embedding vectors close to the model’s original vocabulary, and consider the top- $k$  candidates from the original vocabulary at the last step. This significantly decreases the computational cost for discrete prompt tuning.

### 3 Background and Problem Statement

In this section, we provide a brief background on in-context knowledge editing, and define the detection and reversal problems we address in this work.

#### 3.1 In-context Knowledge Editing

Facts in LLMs are often represented as triplets of (*subject, relation, object*), or  $(s, r, o)$  for short. Querying an LLM with a prompt  $p(s, r)$ , where  $p$  expresses the relation  $r$  and contains the subject  $s$  (e.g., “In which city is the Brandenburg Gate located?”), should result in retrieving the object  $o$  (e.g., “Berlin”), provided that the fact  $(s, r, o)$  is known to the LLM. A knowledge editing operation

$E(s, r, o, o', p)$  is successful if it changes the behavior of the LLM such that the retrieved object is  $o'$  instead of  $o$ . IKE-edits are conducted by prepending an editing prompt  $p_{edit}(s, r, o')$  to the query prompt  $p(s, r)$ , i.e.,  $p_{edit}(s, r, o') \oplus p(s, r)$ , where  $\oplus$  is the string concatenation operation, causes the LLM’s output to change from  $o$  to  $o'$ .

#### 3.2 Problem Statement

Given a model  $\mathcal{M}$  that outputs the object  $o$ , when provided with a prompt  $p(s, r)$ , i.e.,  $o = \operatorname{argmax}_q \mathbb{P}_{\mathcal{M}(p(s,r))}[q]$ .  $\mathcal{M}$ ’s output is changed to  $o'$  with IKE (Zheng et al., 2023) by prepending an editing prompt  $p_{edit}(s, r, o')$  (example in Fig. 7 in the Appendix) to the original prompt, i.e.,  $o' = \operatorname{argmax}_q \mathbb{P}_{\mathcal{M}(p_{edit}(s,r,o') \oplus p(s,r))}[q]$ .

**Detection.** The aim of detecting knowledge edits is to classify an object output by  $\mathcal{M}$  as either edited or unedited. Detecting edits with parameter-modifying KEs is explored in previous work (Youssef et al., 2024b), where a white-box access to the LLM is assumed. In this work, we explore the detection of IKE-edits, that do not modify the model’s parameters. Additionally, we assume a black-box access scenario where only the top-10 output probabilities of the next token are available given a prompt. This assumption reflects a more realistic setting, as it applies to proprietary LLMs that provide the top-10 outputs with their probabilities.<sup>3</sup> We assume that the user sends the query  $p(s, r)$  to an LLM through an API, where the query can be edited by the API internally or by a malicious attacker and therefore the user cannot observe whether an editing prompt  $p_{edit}(s, r, o')$  has been maliciously added to their query or not.

**Reversal.** The goal of reversing knowledge edits is to direct  $\mathcal{M}$ ’s output back from  $o'$  to  $o$  despite  $p_{edit}(s, r, o')$  being part of the input. Reversing edits is particularly important in scenarios where malicious editing prompts are surreptitiously added to user queries, which is hidden from the end-user. Ideally, any reversal operation should not affect unedited prompts. To the best of our knowledge, reversing edits has not been explored before.

### 4 Detecting In-context Knowledge Edits

Informing end-users about post-training edits within LLM-generated content is a straightforward

<sup>3</sup>Some proprietary LLMs (e.g., GPT-3.5) provide the top-10 outputs with their probabilities.

approach to improving transparency in LLM-based API applications. To offer an add-on solution, we frame this challenge as a detection task. Specifically, given any LLM-based API, our goal is to detect edited LLM generations. This section focuses on detecting edits conducted with IKE, a zero-cost knowledge editing method, and introduces an effective approach for doing so.

**Dataset and models.** Following Zheng et al. (2023), we edit facts from COUNTERFACT (Meng et al., 2022) using IKE-editing prompts from (Zheng et al., 2023). COUNTERFACT includes counterfactual statements as the target fact, i.e., false facts (e.g., “The Eiffel Tower is in Rome”). To create our detection dataset, we use 2,000 edits as instances of the edited class. These prompts contain  $p_{edit}(s, r, o')$  for editing and  $p(s, r)$  for querying. As instances of the unedited class, we use the corresponding 2,000 prompts that contain only  $p(s, r)$ , i.e., without IKE-editing prompts. The dataset is then split into two balanced and equally sized parts: a training set and a test set, each containing 2000 instances. We experiment with GPT2-XL (Radford et al., 2019), GPT-J (Wang and Komatsuzaki, 2021) and LLAMA3.1 8B (Dubey et al., 2024).

**Detection method.** We train a simple logistic regression classifier with L1 regularization, using the top-10 output probabilities for the next token, given a prompt, as features to classify the output object as either edited or unedited. Our approach relies on output probabilities, as previous work has demonstrated that parameter-modifying KEs increase the LLM’s certainty about its outputs (Youssef et al., 2024b). We hypothesize that this effect also applies to IKE-edits as well. To train the classification model, we assume access to a training set composed of edited prompts  $p_{edit}(s, r, o') \oplus p(s, r)$  and unedited prompts  $p(s, r)$ .

**Results.** The results in Tab. 1 show that the detection classifier performs well on all three LLMs (F1 > 80%). Using only the top-10 output probabilities of the next token as features enables using this classifier in restrictive settings, where one does not have access to the model’s internals.

To better understand the effect of IKE, we analyze the shift in the output distributions after editing, i.e., we examine how often the top-1 outputs of the model *before* editing can be found in the top-10 outputs of the model *after* editing. Our comparison

Model	Precision	Recall	F1
GPT2-XL	88.32	84.70	86.47
GPT-J	82.73	82.90	82.82
LLAMA-3.1-8B	80.17	86.10	83.03

Table 1: Detecting IKE edits performance using top-10 output probabilities.




Figure 2: Percentage of top1 outputs *before* editing that can be found in the top10 outputs of the model *after* editing (cumulative histogram) in GPT-J.

is based on the first generated token before and after editing. Fig. 2 shows that more than 60% of the original top-1 outputs are contained in the top-10 outputs of the model *after* editing, i.e., the original outputs still have relatively high probabilities after editing. This led us to investigate reversing IKE-edits to retrieve the original outputs.

## 5 Reversing In-context Knowledge Edits

Following our detection results, we now address reversing in-context knowledge edits. Reversing edits allows for recovering the original outputs.

### 5.1 Background: Prompt Tuning

Prompt tuning is a light-weight alternative to fine-tuning, where only part of the input is optimized rather than the whole model’s parameters to enable higher performance on specific tasks (Liu et al., 2023). Tuned prompts can either be *discrete (hard)*, i.e., concrete tokens from the model’s vocabulary, or *continuous (soft)*, vectors in the embedding space that do not correspond to any concrete tokens (Li and Liang, 2021). We initially follow Lester et al. (2021) by adding new tokens to the vocabulary and tuning only these tokens’ embeddings to restore the LLM’s original outputs despite editing prompts (Sec. 5.2). We then explore converting the learned continuous prompts to discrete prompts (Sec. 5.3). For a comprehensive overview of prompt tuning, see Liu et al. (2023).

## 5.2 Continuous Prompt Tuning

To reverse edits, we consider adding a *reversal* token  $r$  to the model’s original vocabulary.  $r$  should ideally make the model ignore the context that appears before  $r$ , and output an answer based on  $p(s, r)$  and the model’s parametric knowledge.

We tune the reversal token to minimize the Kullback-Leibler divergence loss between two probability distributions  $P_r = \mathcal{M}(p_{edit}(s, r, o') \oplus r \oplus p(s, r))$  and  $P_o = \mathcal{M}(p(s, r))$ . This aims to make  $\mathcal{M}$ ’s output distribution with the editing prompt, the reversal token, and the original prompt match  $\mathcal{M}$ ’s output distribution with the original prompt only. Additionally, we ensure that  $P_n = \mathcal{M}(r \oplus p(s, r))$  closely resembles  $P_o$  to prevent the reversal token from affecting the model’s output when no IKE-edits are present. Here, we assume having white-box access to  $\mathcal{M}$ , which enables us to add and tune  $r$  to neutralize editing prompts. We use the following loss function to learn continuous reversal tokens:

$$\mathcal{L}_{cont} = KL[P_r || P_o] + KL[P_n || P_o] \quad (1)$$

## 5.3 Discrete Prompt Tuning

Optimizing reversal tokens in a continuous space allows us to find embedding vectors that help LLMs forget the prior context. However, using such tokens requires expanding the original vocabulary of the model, which limits their usability. Instead, having discrete prompts that correspond to natural tokens in the LLM’s vocabulary can prove more practical, as they can be directly included in the prompt without adapting the model’s vocabulary.

Our approach for finding discrete prompts builds on finding continuous prompts: 1) we identify relevant dimensions in the embedding space for context resetting; 2) we optimize using a joint loss that combines KL divergence for context resetting and cosine distance for staying close to tokens from the model’s vocabulary, while not penalizing changes across the identified relevant dimensions from the first step; 3) we choose one of the top- $k$  nearest discrete tokens to the optimized token.

**1). Dimensions for resetting context.** Word embeddings are high-dimensional vectors with specific dimensions encoding particular signals (Li and Yang, 2018). We assume that certain dimensions control how the model processes context, e.g., the IKE editing part of the prompt that conflicts with the model’s parametric knowledge. To identify these dimensions, we tune a reversal token

similarly to Sec. 5.2 with two differences: first, we initialize the reversal token  $r_{init}$  using an existing vocabulary token (see Appendix B for details); second, we add a cosine loss to keep the learned token embedding  $r_{tuned}$  close to  $r_{init}$ , minimizing unnecessary changes. The loss function is:

$$\mathcal{L}_d = (1 - \lambda)KL[P_r || P_o] + \lambda\mathcal{L}_{cos}(r_{init}, r_{tuned}) \quad (2)$$

After learning  $r_{tuned}$ , we calculate  $w = \frac{r_{tuned} - r_{init}}{\|r_{tuned} - r_{init}\|_1}$ , and use  $\bar{w} = \frac{1/w}{\|1/w\|_1}$  to weigh the embedding dimensions in the following step. Both  $w$  and  $\bar{w}$  sum up to 1. While  $w$  gives the most changed dimensions higher weights,  $\bar{w}$  does the opposite.

**2). Joint optimization.** We optimize a reversal token  $r$  to help LLMs ignore the context using a loss function that consists of two terms:

$$\mathcal{L}_{disc} = (1 - \lambda)\mathcal{L}_{cont} + \lambda\frac{1}{n}\sum_i^n \mathcal{L}_{cos}(r \odot \bar{w}, \mathcal{V}_i \odot \bar{w}) \quad (3)$$

where  $\odot$  refers to the element-wise multiplication, and  $\mathcal{V}_i$  refers to  $i$ -th token in the model’s original vocabulary. We multiply  $r$  and  $\mathcal{V}_i$  by  $\bar{w}$  to ensure that we are not penalizing changes across the identified relevant context dimensions from step 1, and average the cosine loss over all tokens in the vocabulary to ensure that  $r$  is generally close to the original embedding vectors.

**3). Choosing discrete tokens.** After learning a continuous reversal token  $r$  in step 2, we convert  $r$  to  $r_d$ ; a discrete token that corresponds to a natural token in the vocabulary. To find  $r_d$ , we first use a weighted cosine similarity to find a candidate set:  $R_{cand} = \text{top}k_{\mathcal{V}_i}(\text{CosineSimilarity}(r \odot w, \mathcal{V}_i \odot w))$ , i.e., we find the top- $k$  nearest neighbours after weighting  $r$  and the tokens from  $\mathcal{V}$  by  $w$ , so that the identified relevant context dimensions from step 1 are given higher weights. We evaluate all the candidates on a validation set and choose the token that minimizes the the KL-loss the most, i.e.,  $r_d = \min_{\mathcal{V}_i} KL(P_r || P_o) + KL(P_n || P_o)$ .

Our approach for finding one reversal token can be generalized to several tokens. We experiment with a varying number of reversal tokens in Sec. 6.

## 6 Experimental Setup

We evaluate how both continuous and discrete prompt tuning affect the model’s outputs. We use

the same models and datasets as the ones previously used for detecting IKE-edits (cf. Sec. 4). We train for 3 epochs with Adam as optimizer using the default parameters. We keep the model’s parameters frozen and tune only the embeddings of the reversal tokens. In discrete prompt tuning, we set  $k = 10$  for choosing the final output. We experiment with different numbers of reversal tokens. For each setting, we experiment with 5 seeds and report the maximum and mean accuracy with standard deviation.

**Evaluation.** We calculate the *matching accuracy*, i.e., the agreement of the original output and the output of the edited prompt with reversal tokens. For a perfect reversal method, i.e., optimal reversal tokens, the matching accuracy would be 1.0. Matching accuracy is calculated as  $\frac{1}{N} \sum_{i=1}^N \mathbf{1}(\hat{y}_i = y_i)$ , where  $y_i$  being the original output, and  $\hat{y}_i$  the edited and reverse-edited output. Following Du et al. (2024), we approximate the model’s outputs using the next token prediction.

## 6.1 Results and Discussion

Tab. 2 shows the matching accuracy for each setting. For comparison, we also report the matching accuracy between the original prompt  $p(s, r)$  and the edited prompt  $p_{edit}(s, r, o') \oplus p(s, r)$  (first row).

**Continuous prompt tuning.** When optimizing reversal tokens in a continuous space, we observe high performance across all three models. On the smallest model (GPT2-XL, 1.5B) the performance varies between 45% and 78%, and with high standard deviation, but the performance generally becomes more stable with more reversal tokens. On the larger models, GPT-J (6B) and LLAMA3.1 (8B), the accuracy is higher than 80% in most cases with low standard deviation. We also observe small improvements when increasing the number of reversal tokens. These results suggest that reversing edits in larger models is easier, more stable and possible with as few as one reversal token, which might be due to their larger capacity.

**Discrete prompt tuning.** The results with discrete tuning are less positive. Compared to continuous tuning, we observe a drop in the mean accuracy across all models, especially on LLAMA3.1, where accuracy is extremely poor. An open question here is whether a token or a set of tokens exists in the LLM’s vocabulary that can cause the model to ignore prior context (IKE-edits). On GPT2-XL and

GPT-J the highest accuracy is 78% and 73% respectively with high standard deviation across all settings. This suggests the existence of tokens that help the model ignore prior context, but implies that finding such tokens is not always successful. We analyze the learned tokens in Sec. 7.

**Ablation.** We conduct an ablation under discrete prompt tuning, removing the cosine loss term or the KL-loss term (**avg w/o Cos** and **avg w/o KL**, respectively in Tab. 2). When removing the cosine loss term, we notice that the accuracy for GPT-J drops significantly, whereas the performance on GPT2-XL is still comparable to using the original loss function. This suggests that the KL-loss is sufficient to find reversal tokens in GPT2-XL, whereas in GPT-J, using only the KL-loss causes the learned tokens to drift away from the natural tokens in the vocabulary that help reverse edits. Excluding the KL-loss term significantly decreases performance for all models. This is intuitive, since optimizing reversal tokens to stay close to the natural token embeddings would not contribute to finding good reversal tokens. Finding any valid reversal tokens under this setting would have been contributed to step 3 (cf. Sec. 5.3), which involves choosing the best performing tokens from the  $k$  nearest neighbours on the validation set.

## 7 Analysis

We analyze the found reversal tokens under the discrete setting. Our goal is to investigate which *natural* tokens help reverse edits. We show examples with 1 reversal token in Tab. 4, and more examples in Fig. 5 in the Appendix. The found tokens contain some non-English tokens. In GPT2-XL and GPT-J, we notice that the prompts that contain the beginning of sentence token (<BOS>) achieve high accuracy (more than 70% and 60% in GPT2-XL and GPT-J respectively). The <BOS> token is often used at the beginning of training instances and signals the onset of a new context. It is therefore intuitive for the <BOS> to be helpful in reversing edits. We experimented with the <BOS> token in LLAMA3.1, but this did not lead to improved performance. This might be due to differences in the training procedure in LLAMA3.1 compared to the GPT-models.

**Effect of reversal tokens on normal prompts.** Reversal tokens should restore the original outputs when IKE-edits are present in the prompt, but ide-

	#rt	GPT2-XL			GPT-J			LLAMA-3.1-8B		
		Max	Mean	Std	Max	Mean	Std	Max	Mean	Std
	IKE vs. no-IKE	1.90	1.90	0.00	2.20	2.20	0.00	0.60	0.60	0.00
Cont. Tuning	1	77.60	57.82	22.36	80.80	79.48	0.95	83.90	82.42	0.88
	2	79.10	61.96	21.33	80.80	79.42	0.99	85.20	82.82	1.48
	3	75.00	44.98	21.38	80.90	79.92	0.96	84.80	83.24	1.32
	4	78.80	62.80	22.28	81.80	80.56	0.78	84.90	83.48	1.01
	5	79.20	71.64	14.77	81.30	80.38	0.67	84.80	83.88	0.90
	6	79.60	69.08	14.09	81.60	80.90	0.61	84.90	83.34	0.93
	7	82.10	75.52	7.85	82.20	81.44	0.61	85.70	84.12	1.80
	8	82.30	<b>78.14</b>	4.96	83.10	81.74	1.44	85.10	84.48	0.77
	9	<b>82.90</b>	73.70	6.16	83.00	81.94	0.91	<b>86.00</b>	<b>84.64</b>	1.05
	10	80.10	71.56	12.73	<b>84.30</b>	<b>82.52</b>	1.10	85.40	84.40	1.03
	avg	79.67	66.72	14.79	81.98	80.83	0.90	85.07	83.68	1.12
Disc. Tuning	1	75.70	18.16	32.17	61.10	15.12	25.70	<b>0.90</b>	<b>0.74</b>	0.09
	2	73.10	43.30	35.75	60.50	33.30	26.89	0.70	0.62	0.04
	3	75.10	18.68	31.55	58.80	42.20	22.03	0.60	0.60	0.00
	4	75.00	32.26	36.56	70.70	36.82	32.89	0.60	0.60	0.00
	5	78.00	46.24	37.99	<b>73.10</b>	47.32	30.62	0.70	0.62	0.04
	6	74.60	44.76	35.72	71.30	<b>64.10</b>	8.46	0.70	0.68	0.04
	7	76.30	46.04	37.15	67.30	46.34	21.95	0.70	0.62	0.04
	8	74.00	44.52	36.66	65.20	47.18	18.95	0.70	0.62	0.04
	9	<b>78.60</b>	<b>57.46</b>	28.51	62.70	54.46	9.82	0.70	0.64	0.05
	10	76.20	31.60	38.03	67.20	42.28	24.59	0.60	0.60	0.00
	avg	75.66	38.30	35.01	65.79	42.91	22.19	0.69	0.63	0.03
	avg w/o Cos	77.02	46.73	31.90	4.07	3.05	0.74	0.86	0.68	0.11
	avg w/o KL	2.41	2.01	0.28	3.36	2.60	0.57	0.79	0.69	0.07

Table 2: Matching accuracy with continuous and discrete reversal tokens in different settings. #rt refer to the number of reversal tokens. We report the max and mean accuracy with standard deviation across 5 seeds. For comparison, we report matching accuracy between prompts with/without IKE-edits in the first top row after the header.

Setting	GPT2-XL			GPT-J			LLAMA-3.1-8B		
	Max	Mean	Std	Max	Mean	Std	Max	Mean	Std
Continuous	86.19	84.28	1.74	87.56	86.29	1.10	92.62	91.48	1.17
Discrete	79.44	74.57	4.01	74.85	69.74	4.17	79.73	75.70	3.42

Table 3: Max and mean matching accuracy between normal (unedited) prompts *with* and *without* reversal tokens. Ideally, reversal tokens should not affect the outputs of normal prompts if there are no IKE-edits. Continuous reversal tokens have less effects on normal prompts (higher accuracy) compared to discrete reversal tokens.

GPT2-XL		GPT-J		LLAMA-3.1-8B	
Acc.	Tokens	Acc.	Tokens	Acc.	Tokens
3.8	UCLA	3.5	a	0.7	ümü
4.1	NETWORK	3.5	a	0.7	?
3.5	cellaneous	61.1	<BOS>	0.7	_ST
3.7	Henry	3.5	(	0.7	
75.7	<BOS>	4.0	d	0.9	Revised

Table 4: The found discrete reversal token. The <BOS> token helps reverse edits in GPT-models.

ally should not affect the outputs if the prompts do not contain IKE-edits. To assess this, we compare the matching accuracy between normal prompts with and without reversal tokens ( $r \oplus p(s, r)$  vs.  $p(s, r)$ ). Results in Tab. 3 show that continuous tokens change only 9-15% of the outputs on average, while discrete tokens affect 25-30%. This shows that continuous tokens have less side-effects.

These unintended changes should be further minimized, but can be mitigated by first detecting an edit and applying reversal tokens only when edits are suspected. We encourage developing more robust methods in future work.

**Fact retrieval with IKE.** Our analysis aims to understand how IKE affects fact retrieval in GPT2-XL, and how adding <BOS> as reversal token impacts this process. We analyze fact retrieval under three conditions: without IKE edits, with IKE edits, and with IKE edits plus the <BOS> token. We project output representations from each GPT2-XL layer into the vocabulary space (Nostalgebraist, 2020; Alammari, 2021).<sup>4</sup> Projecting the output representation at each layer gives us a probability dis-

<sup>4</sup>In Transformers, representations can be viewed as an information stream, with each layer reading from and writing back to this stream (Elhage et al., 2021; Geva et al., 2022).




Figure 3: Ranking of the final output over different layers in GPT2-XL. Rankings are averaged and rounded over 2000 instances. IKE causes the model to become certain about its prediction at earlier layers, while the model normally (no IKE prompt) keeps on refining its predictions until the final layer.

tribution over the vocabulary, which represents the model’s output at this layer. Specifically, we show the rank of the final output at each layer, indicating the probability evolution of the final output throughout the model.

Fig. 3 shows that the ranks of the final outputs with no IKE-edits keep on changing until the final layer, i.e., the model is making use of all of its layers to refine the final prediction. Conversely, with IKE-edits the rank of the final output does not change after layer 29, i.e., the model is certain about its output much earlier compared to the normal case. One can also notice that the final output is less probable at earlier layers with IKE. We attribute this to the longer inputs that the model needs to process in the IKE setting. Adding <BOS> makes the rankings with IKE more similar to the normal rankings, where the predictions are refined until the last layer. At early layers, we notice the final prediction becomes more probable after adding <BOS> (ranks 5275 instead of 10989), which suggests that <BOS> helps the model filter out potential outputs that come from the IKE prompt.

**Attention weights.** We further investigate how adding the <BOS> token affects the attention weights distribution in GPT2-XL (Bahdanau et al., 2016; Sarti et al., 2023). More specifically, we consider the averaged attention weights that are assigned to the query  $p(s, r)$  in two settings: 1) IKE-edited; 2) IKE-edited with <BOS> over a sample of 100 examples. We hypothesize that adding <BOS> encourages the model to ignore the IKE-edit (prior context) in the prompt and focus more on  $p(s, r)$ , i.e., assign  $p(s, r)$  higher attention weights. Indeed, we observe that when adding <BOS> the attention mass increases from 15% to 38% of the original attention mass observed in the unedited setting, despite the large increase in the input length from 8.13 to 877.3 tokens on average with IKE-edits. This suggests that <BOS> helps the model to ignore the IKE-edit and to focus more on the

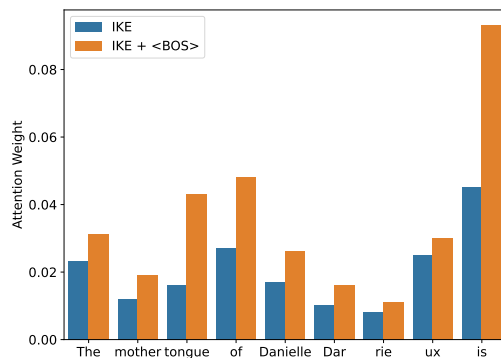


Figure 4: Attention weights on an example from COUNTERFACT with an IKE-edit. Adding <BOS> helps the model to focus more on the query  $p(s, r)$ . IKE part of the input is omitted for space reasons.

subsequent query. We show the attention weights for one example in Fig. 4. Adding <BOS> clearly leads to higher attention weights for the query.

## 8 Conclusion

In this work, we studied detecting and reversing IKE-edits in LLMs. Our work demonstrated that it is possible to detect IKE-edits with high accuracy, even in scenarios where the model’s parameters are not changed and where only limited information about the model’s outputs is available. We further showed that IKE-edits can be neutralized by tuning special reversal tokens. By achieving over 80% accuracy in recovering original outputs, continuous reversal tokens offer a promising method for mitigating the effects of malicious edits, and making LLMs intrinsically more secure. The effectiveness of continuous reversal tokens across different LLMs suggests a degree of generalizability that could be valuable for widespread implementation. We believe our work contributes significantly to the ongoing efforts to make LLMs more secure, transparent, and trustworthy. Our future research will focus on improving the accuracy and robustness of both continuous and discrete reversal tokens.



## 9 Limitations

Following Zheng et al. (2023), we experimented with COUNTERFACT only, since IKE-edits require demonstrations that share the same relation as the edited facts, and such relations/demonstrations are not easily identifiable in other datasets. As a first venture for studying reversing edits, we assumed white-box access to LLMs to add and tune special reversal tokens. However, such access might be difficult to attain for users in practice.

## References

- J Alammar. 2021. Ecco: An open source library for the explainability of transformer language models. In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing: System Demonstrations*. Association for Computational Linguistics.
- Dzmitry Bahdanau, Kyunghyun Cho, and Yoshua Bengio. 2016. [Neural machine translation by jointly learning to align and translate](#).
- Lucas Bourtole, Varun Chandrasekaran, Christopher A Choquette-Choo, Hengrui Jia, Adelin Travers, Baiwu Zhang, David Lie, and Nicolas Papernot. 2021. Machine unlearning. In *2021 IEEE Symposium on Security and Privacy (SP)*, pages 141–159. IEEE.
- Canyu Chen, Baixiang Huang, Zekun Li, Zhaorun Chen, Shiyang Lai, Xiong Xiao Xu, Jia-Chen Gu, Jindong Gu, Huaxiu Yao, Chaowei Xiao, Xifeng Yan, William Yang Wang, Philip Torr, Dawn Song, and Kai Shu. 2024. [Can editing LLMs inject harm?](#) In *Trustworthy Multi-modal Foundation Models and AI Agents (TiFA)*.
- Jiaao Chen and Diyi Yang. 2023. [Unlearn what you want to forget: Efficient unlearning for LLMs](#). In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pages 12041–12052, Singapore. Association for Computational Linguistics.
- Roi Cohen, Eden Biran, Ori Yoran, Amir Globerson, and Mor Geva. 2024. [Evaluating the ripple effects of knowledge editing in language models](#). *Transactions of the Association for Computational Linguistics*, 12:283–298.
- Nicola De Cao, Wilker Aziz, and Ivan Titov. 2021. [Editing factual knowledge in language models](#). In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 6491–6506, Online and Punta Cana, Dominican Republic. Association for Computational Linguistics.
- Kevin Du, Vésteinn Snæbjarnarson, Niklas Stoehr, Jennifer White, Aaron Schein, and Ryan Cotterell. 2024. [Context versus prior knowledge in language models](#). In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 13211–13235, Bangkok, Thailand. Association for Computational Linguistics.
- Abhimanyu Dubey, Abhinav Jauhri, Abhinav Pandey, Abhishek Kadian, Ahmad Al-Dahle, Aiesha Letman, et al. 2024. [The Llama 3 Herd of Models](#).
- Nelson Elhage, Neel Nanda, Catherine Olsson, Tom Henighan, Nicholas Joseph, Ben Mann, Amanda Askell, Yuntao Bai, Anna Chen, Tom Conerly, Nova DasSarma, Dawn Drain, Deep Ganguli, Zac Hatfield-Dodds, Danny Hernandez, Andy Jones, Jackson Kernion, Liane Lovitt, Kamal Ndousse, Dario Amodei, Tom Brown, Jack Clark, Jared Kaplan, Sam McCandlish, and Chris Olah. 2021. [A mathematical framework for transformer circuits](#). *Transformer Circuits Thread*. <https://transformer-circuits.pub/2021/framework/index.html>.
- Govind Gangadhar and Karl Stratos. 2024. [Model editing by pure fine-tuning](#). *arXiv preprint arXiv:2402.11078*.
- Mor Geva, Avi Caciularu, Kevin Wang, and Yoav Goldberg. 2022. [Transformer feed-forward layers build predictions by promoting concepts in the vocabulary space](#). In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing*, pages 30–45, Abu Dhabi, United Arab Emirates. Association for Computational Linguistics.
- Karina Halevy, Anna Sotnikova, Badr AlKhamissi, Syrielle Montariol, and Antoine Bosselut. 2024. ["flex tape can't fix that": Bias and misinformation in edited language models](#).
- Adi Haviv, Jonathan Berant, and Amir Globerson. 2021. [BERTese: Learning to speak to BERT](#). In *Proceedings of the 16th Conference of the European Chapter of the Association for Computational Linguistics: Main Volume*, pages 3618–3623, Online. Association for Computational Linguistics.
- Rima Hazra, Sayan Layek, Somnath Banerjee, and Soujanya Poria. 2024. [Sowing the wind, reaping the whirlwind: The impact of editing language models](#). In *Findings of the Association for Computational Linguistics ACL 2024*, pages 16227–16239, Bangkok, Thailand and virtual meeting. Association for Computational Linguistics.
- Baixiang Huang, Canyu Chen, Xiong Xiao Xu, Ali Payani, and Kai Shu. 2024. [Can knowledge editing really correct hallucinations?](#)
- Gabriel Ilharco, Marco Tulio Ribeiro, Mitchell Wortsman, Ludwig Schmidt, Hannaneh Hajishirzi, and Ali Farhadi. 2023. [Editing models with task arithmetic](#). In *The Eleventh International Conference on Learning Representations*.

- Joel Jang, Dongkeun Yoon, Sohee Yang, Sungmin Cha, Moontae Lee, Lajanugen Logeswaran, and Minjoon Seo. 2023. [Knowledge unlearning for mitigating privacy risks in language models](#). In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 14389–14408, Toronto, Canada. Association for Computational Linguistics.
- Zhengbao Jiang, Jun Araki, Haibo Ding, and Graham Neubig. 2021. [How can we know when language models know? on the calibration of language models for question answering](#). *Transactions of the Association for Computational Linguistics*, 9:962–977.
- Tianjie Ju, Yiting Wang, Xinbei Ma, Pengzhou Cheng, Haodong Zhao, Yulong Wang, Lifeng Liu, Jian Xie, Zhuosheng Zhang, and Gongshen Liu. 2024. [Flooding spread of manipulated knowledge in llm-based multi-agent communities](#).
- Brian Lester, Rami Al-Rfou, and Noah Constant. 2021. [The power of scale for parameter-efficient prompt tuning](#). In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 3045–3059, Online and Punta Cana, Dominican Republic. Association for Computational Linguistics.
- Xiang Lisa Li and Percy Liang. 2021. [Prefix-tuning: Optimizing continuous prompts for generation](#). In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 4582–4597, Online. Association for Computational Linguistics.
- Xiaopeng Li, Shangwen Wang, Shezheng Song, Bin Ji, Huijun Liu, Shasha Li, Jun Ma, and Jie Yu. 2024a. [Identifying knowledge editing types in large language models](#).
- Yang Li and Tao Yang. 2018. *Word Embedding for Understanding Natural Language: A Survey*. Springer International Publishing, Cham.
- Yanzhou Li, Kangjie Chen, Tianlin Li, Jian Zhang, Shangqing Liu, Wenhan Wang, Tianwei Zhang, and Yang Liu. 2024b. [Badedit: Backdooring large language models by model editing](#). In *The Twelfth International Conference on Learning Representations*.
- Pengfei Liu, Weizhe Yuan, Jinlan Fu, Zhengbao Jiang, Hiroaki Hayashi, and Graham Neubig. 2023. [Pre-train, prompt, and predict: A systematic survey of prompting methods in natural language processing](#). *ACM Comput. Surv.*, 55(9).
- Kevin Meng, David Bau, Alex Andonian, and Yonatan Belinkov. 2022. [Locating and editing factual associations in GPT](#). *Advances in Neural Information Processing Systems*, 36.
- Kevin Meng, Arnab Sen Sharma, Alex Andonian, Yonatan Belinkov, and David Bau. 2023. [Mass editing memory in a transformer](#). *The Eleventh International Conference on Learning Representations (ICLR)*.
- Eric Mitchell, Charles Lin, Antoine Bosselut, Chelsea Finn, and Christopher D Manning. 2022a. [Fast model editing at scale](#). In *International Conference on Learning Representations*.
- Eric Mitchell, Charles Lin, Antoine Bosselut, Christopher D Manning, and Chelsea Finn. 2022b. [Memory-based model editing at scale](#). In *Proceedings of the 39th International Conference on Machine Learning*, volume 162 of *Proceedings of Machine Learning Research*, pages 15817–15831. PMLR.
- Thanh Tam Nguyen, Thanh Trung Huynh, Phi Le Nguyen, Alan Wee-Chung Liew, Hongzhi Yin, and Quoc Viet Hung Nguyen. 2022. [A survey of machine unlearning](#). *arXiv preprint arXiv:2209.02299*.
- Nostalgebraist. 2020. [interpreting GPT: the logit lens](#).
- Martin Pawelczyk, Seth Neel, and Himabindu Lakkaraju. 2024. [In-context unlearning: Language models as few shot unlearners](#).
- Jiyang Qiu, Xinbei Ma, Zhuosheng Zhang, and Hai Zhao. 2024a. [MEGen: Generative backdoor in large language models via model editing](#). *arXiv preprint arXiv:2408.10722*.
- Jiyang Qiu, Xinbei Ma, Zhuosheng Zhang, and Hai Zhao. 2024b. [Megen: Generative backdoor in large language models via model editing](#).
- Alec Radford, Jeffrey Wu, Rewon Child, David Luan, Dario Amodei, Ilya Sutskever, et al. 2019. [Language models are unsupervised multitask learners](#). *OpenAI blog*, 1(8):9.
- Gabriele Sarti, Nils Feldhus, Ludwig Sickert, and Oscar van der Wal. 2023. [Inseq: An interpretability toolkit for sequence generation models](#). In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 3: System Demonstrations)*, pages 421–435, Toronto, Canada. Association for Computational Linguistics.
- Taylor Shin, Yasaman Razeghi, Robert L. Logan IV, Eric Wallace, and Sameer Singh. 2020. [AutoPrompt: Eliciting Knowledge from Language Models with Automatically Generated Prompts](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 4222–4235, Online. Association for Computational Linguistics.
- Nianwen Si, Hao Zhang, Heyu Chang, Wenlin Zhang, Dan Qu, and Weiqiang Zhang. 2023. [Knowledge unlearning for llms: Tasks, methods, and challenges](#).
- Chenmian Tan, Ge Zhang, and Jie Fu. 2024. [Massive editing for large language models via meta learning](#). In *International Conference on Learning Representations*.

- Eric Wallace, Shi Feng, Nikhil Kandpal, Matt Gardner, and Sameer Singh. 2019. [Universal adversarial triggers for attacking and analyzing NLP](#). In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 2153–2162, Hong Kong, China. Association for Computational Linguistics.
- Ben Wang and Aran Komatsuzaki. 2021. GPT-J-6B: A 6 Billion Parameter Autoregressive Language Model. <https://github.com/kingoflolz/mesh-transformer-jax>.
- Lingzhi Wang, Tong Chen, Wei Yuan, Xingshan Zeng, Kam-Fai Wong, and Hongzhi Yin. 2023. [KGA: A general machine unlearning framework based on knowledge gap alignment](#). In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 13264–13276, Toronto, Canada. Association for Computational Linguistics.
- Xinwei Wu, Junzhuo Li, Minghui Xu, Weilong Dong, Shuangzhi Wu, Chao Bian, and Deyi Xiong. 2023. [DEPN: Detecting and editing privacy neurons in pre-trained language models](#). In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pages 2875–2886, Singapore. Association for Computational Linguistics.
- Rongwu Xu, Zehan Qi, Zhijiang Guo, Cunxiang Wang, Hongru Wang, Yue Zhang, and Wei Xu. 2024. [Knowledge conflicts for llms: A survey](#).
- Paul Youssef, Osman Koraş, Meijie Li, Jörg Schlötterer, and Christin Seifert. 2023. [Give Me the Facts! A Survey on Factual Knowledge Probing in Pre-trained Language Models](#). In *Findings of the Association for Computational Linguistics: EMNLP 2023*, pages 15588–15605, Singapore. Association for Computational Linguistics.
- Paul Youssef, Jörg Schlötterer, and Christin Seifert. 2024a. [The Queen of England is not England’s Queen: On the Lack of Factual Coherency in PLMs](#). In *Findings of the Association for Computational Linguistics: EACL 2024*, pages 2342–2354, St. Julian’s, Malta. Association for Computational Linguistics.
- Paul Youssef, Zhixue Zhao, Daniel Braun, Jörg Schlötterer, and Christin Seifert. 2025. [Position: Editing Large Language Models Poses Serious Safety Risks](#).
- Paul Youssef, Zhixue Zhao, Christin Seifert, and Jörg Schlötterer. 2024b. [Has this Fact been Edited? Detecting Knowledge Edits in Language Models](#). *arXiv preprint arXiv:2405.02765*.
- Jinghan Zhang, Shiqi Chen, Junteng Liu, and Junxian He. 2023. [Composing parameter-efficient modules with arithmetic operation](#). In *Thirty-seventh Conference on Neural Information Processing Systems*.
- 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?](#) In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pages 4862–4876, Singapore. Association for Computational Linguistics.
- Zexuan Zhong, Zhengxuan Wu, Christopher Manning, Christopher Potts, and Danqi Chen. 2023. [MQuAKE: Assessing knowledge editing in language models via multi-hop questions](#). In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pages 15686–15702, Singapore. Association for Computational Linguistics.

## A Additional Results

We show the learned reversal tokens in Fig. 5. Tab. 5 is a more comprehensive version of Tab. 2 (with varying numbers of reversal tokens for the ablation settings).

**Generalization to other datasets.** We use MQuAKE (Zhong et al., 2023), which contains multi-hop edits to study the generalization of the proposed continuous reversal tokens. Since this dataset does not contain demonstration prompts for IKE, we follow Huang et al. (2024) in using prompts that contain the new facts for editing (see example in Fig. 8). The proposed continuous reversal tokens generalize well to MQuAKE, showing improved performance with GPT models but decreased performance with LLAMA. We attribute this to MQuAKE’s multi-hop edits, which challenge larger models to resist producing the edited answer. Despite the positive results, we believe the absence of demonstration prompts does not allow us to perfectly simulate the IKE setting, which we focus on in this paper.

## B Further Implementation Details

**Dimensions for resetting context.** For step 1 in discrete prompt tuning, we experimented with 5 random seeds, and continued the experiments with the seed whose accuracy corresponded to the median accuracy value.

## C Computational Resources

All of our experiments were conducted with an NVIDIA A100 GPU with 80GB of memory. Our experiments took roughly 10 GPU days.

	#rt	GPT2-XL			GPT-J			LLAMA-3.1-8B		
		Max	Mean	Std	Max	Mean	Std	Max	Mean	Std
	IKE vs. no-IKE	1.90	1.90	0.00	2.20	2.20	0.00	0.60	0.60	0.00
Cont. Tuning	1.0	77.60	57.82	22.36	80.80	79.48	0.95	83.90	82.42	0.88
	2.0	79.10	61.96	21.33	80.80	79.42	0.99	85.20	82.82	1.48
	3.0	75.00	44.98	21.38	80.90	79.92	0.96	84.80	83.24	1.32
	4.0	78.80	62.80	22.28	81.80	80.56	0.78	84.90	83.48	1.01
	5.0	79.20	71.64	14.77	81.30	80.38	0.67	84.80	83.88	0.90
	6.0	79.60	69.08	14.09	81.60	80.90	0.61	84.90	83.34	0.93
	7.0	82.10	75.52	7.85	82.20	81.44	0.61	85.70	84.12	1.80
	8.0	82.30	<b>78.14</b>	4.96	83.10	81.74	1.44	85.10	84.48	0.77
	9.0	<b>82.90</b>	73.70	6.16	83.00	81.94	0.91	<b>86.00</b>	<b>84.64</b>	1.05
	10.0	80.10	71.56	12.73	<b>84.30</b>	<b>82.52</b>	1.10	85.40	84.40	1.03
	avg	79.67	66.72	14.79	81.98	80.83	0.90	85.07	83.68	1.12
Disc. Tuning	1.0	75.70	18.16	32.17	61.10	15.12	25.70	<b>0.90</b>	<b>0.74</b>	0.09
	2.0	73.10	43.30	35.75	60.50	33.30	26.89	0.70	0.62	0.04
	3.0	75.10	18.68	31.55	58.80	42.20	22.03	0.60	0.60	0.00
	4.0	75.00	32.26	36.56	70.70	36.82	32.89	0.60	0.60	0.00
	5.0	78.00	46.24	37.99	<b>73.10</b>	47.32	30.62	0.70	0.62	0.04
	6.0	74.60	44.76	35.72	71.30	<b>64.10</b>	8.46	0.70	0.68	0.04
	7.0	76.30	46.04	37.15	67.30	46.34	21.95	0.70	0.62	0.04
	8.0	74.00	44.52	36.66	65.20	47.18	18.95	0.70	0.62	0.04
	9.0	<b>78.60</b>	<b>57.46</b>	28.51	62.70	54.46	9.82	0.70	0.64	0.05
	10.0	76.20	31.60	38.03	67.20	42.28	24.59	0.60	0.60	0.00
	avg	75.66	38.30	35.01	65.79	42.91	22.19	0.69	0.63	0.03
Disc. – Cos loss	1.0	75.70	32.48	39.47	5.90	4.70	1.07	0.70	0.64	0.05
	2.0	74.90	32.46	37.62	<b>8.00</b>	4.88	1.94	<b>1.90</b>	<b>0.92</b>	0.55
	3.0	78.60	18.76	33.49	5.80	<b>5.02</b>	0.48	1.00	0.72	0.16
	4.0	74.60	45.84	39.06	3.80	3.00	0.61	0.70	0.66	0.05
	5.0	78.50	<b>71.56</b>	8.69	2.50	2.00	0.52	0.70	0.66	0.05
	6.0	78.60	32.32	39.13	2.30	1.96	0.36	0.70	0.66	0.05
	7.0	78.70	60.30	29.48	3.10	2.36	0.57	0.80	0.64	0.09
	8.0	75.60	56.28	30.16	2.90	2.00	0.60	0.60	0.60	0.00
	9.0	76.20	58.10	31.12	2.90	2.20	0.56	0.80	0.66	0.09
	10.0	<b>78.80</b>	59.20	30.78	3.50	2.38	0.70	0.70	0.62	0.04
	avg	77.02	46.73	31.90	4.07	3.05	0.74	0.86	0.68	0.11
Disc. – KL loss	1.0	2.00	2.00	0.00	3.50	<b>3.50</b>	0.00	0.60	0.60	0.00
	2.0	2.00	1.58	0.24	3.30	2.38	0.52	0.70	0.62	0.04
	3.0	2.20	1.96	0.23	3.00	2.44	0.59	<b>1.20</b>	<b>0.78</b>	0.25
	4.0	2.50	1.94	0.48	2.90	2.36	0.52	0.70	0.70	0.00
	5.0	2.50	2.06	0.36	2.90	2.38	0.46	0.80	0.72	0.04
	6.0	2.80	<b>2.28</b>	0.34	3.00	2.64	0.35	0.70	0.64	0.05
	7.0	2.40	1.94	0.27	<b>5.40</b>	2.94	1.39	0.80	0.70	0.07
	8.0	2.50	2.06	0.30	2.90	2.36	0.48	0.70	0.66	0.05
	9.0	2.20	1.92	0.18	3.30	2.46	0.63	0.80	0.68	0.08
	10.0	<b>3.00</b>	2.40	0.40	3.40	2.50	0.78	0.90	0.76	0.11
	avg	2.41	2.01	0.28	3.36	2.60	0.57	0.79	0.69	0.07

Table 5: Matching accuracy with continuous and discrete reversal tokens in different settings. #rt refer to the number of reversal tokens. We report the max and mean accuracy with standard deviation across 5 seeds. For comparison, we report matching accuracy between prompts with/without IKE-edits in the first top line.

	#rt	GPT2-XL			GPT-J			LLAMA-3.1-8B		
		Max	Mean	Std	Max	Mean	Std	Max	Mean	Std
	IKE vs. no-IKE	1.00	1.00	0.00	0.5	0.50	0.00	0.20	0.20	0.00
Cont. Tuning	1.0	99.10	78.66	42.45	100.0	100.00	0.00	61.80	57.96	2.71
	2.0	99.10	98.82	0.28	100.0	100.00	0.00	65.60	61.90	3.14
	3.0	100.00	97.70	2.60	100.0	99.98	0.04	63.80	60.88	3.25
	4.0	100.00	98.10	3.29	100.0	99.98	0.04	67.60	62.96	2.60
	5.0	100.00	99.10	1.00	100.0	99.98	0.04	66.20	63.10	2.09
	6.0	99.80	98.58	1.63	100.0	100.00	0.00	67.90	64.78	2.36
	7.0	100.00	99.80	0.14	100.0	100.00	0.00	67.40	65.06	1.92
	8.0	99.90	95.36	9.87	100.0	100.00	0.00	68.50	65.68	1.95
	9.0	99.70	99.52	0.20	100.0	99.96	0.09	66.60	65.62	1.14
	10.0	99.90	94.12	12.48	100.0	99.98	0.04	68.70	67.22	1.55
	avg	99.75	95.98	7.39	100.0	99.99	0.02	66.41	63.52	2.27

Table 6: Matching accuracy with continuous reversal tokens using MQuAKE. #rt refer to the number of reversal tokens. We report the max and mean accuracy with standard deviation across 5 seeds. For comparison, we report matching accuracy between prompts with/without IKE-edits in the first top line.

rt	GPT2-XL		GPT-J		LLAMA-3.1-8B	
	acc.	Token	acc.	Token	acc.	Token
5	68.5	level vortex <BOS> wedding	73.1	<BOS> <BOS> . Uran <BOS>	0.6	EditorGUILLayout ..... ...../
	78.0	niche ID <BOS> <BOS> Partial	27.9	<BOS> ( ` the	0.6	::, ého "[% ically ""!
	5.1	, Comed 五 ALK and	61.2	Angry ( tantal FAM <BOS>	0.6	etting (昭和 Quartz '); ?><
	4.5	AB Definition ockey Burn Los	3.2	.'. an Sadd J). Reb	0.6	dpi Mothers =>{ Shemale -eight
	75.1	stereotyp iosyncr <BOS> laps The	71.2	the <BOS> to (<BOS>	0.7	_IL ---@ Không postcode ).
7	5.8	Saharan — easy software a action Fresh	62.0	<BOS> autob i a ucha <BOS> collective	0.6	IL Needs ()," Everton (stream (!
	71.4	crab <BOS> <BOS> <BOS> orthern regul Laurie	14.1	<BOS> Azerb brist Sor ( and a	0.6	Fucking gest Kpim potency уааа cker .Ship
	76.3	oilberal 21 template ransom NCT loss <BOS>	54.0	Tuls ( Reb <BOS> <BOS> 7	0.7	SDK .is crire pome ="\$. _endars
	71.7	<BOS> <BOS> <BOS> glitter Israel Americ	67.3	<BOS> Gri <BOS> <BOS> \$. Elys autob i	0.6	_ros 品 فرد quar inho !!
	5.0	Nokia I le quest Yan erker grading	34.3	the ( (<BOS> , \$.	0.6	⊙⊗ exh ographically 专档);
10	4.1	invitations mascara haste PID upgrades ballots Christmas anchor wrestling	67.2	<BOS> <BOS> the brist <BOS> , constitute , ,	0.6	(HttpStatusCode .Requires Couldn .extensions şekinde 島 ] дв DM >>
	3.8	isSpecial Azerb constitution democratic " and for rests Surprise	1.4	). Spons plagiar were . is is Ali Reb	0.6	送料無料 withStyles ऋ 数量 phant :"; City ") , .
	3.7	Electrical pillars disorderly Cheong Centers Effective Evening tun concussion Fior	44.0	Tart and Sadd <BOS> <BOS> Wan ( ( the ,	0.6	imbledon ادى asto /loader цюнный ۛ SPI (c .» :
	70.2	Qt BYU <BOS> amuse Half Saturday Tour bolst ! sizeable	52.7	Berks Gael Berks ). Papua Uk \$. Berks <BOS> the	0.6	scrollTo oluřtur tipped ilmington <Button orda /login .ident lify STRUCTIONS
	76.2	Gamer posts » in Emily Canad the timelines transformer <BOS>	46.1	. <BOS> ALSO the <BOS> , were S . NW	0.6	际是—(vs *----- q! Volley outh iator .precision ]?>

Figure 5: The found discrete tokens with 5, 7 and 10 reversal tokens and 5 seeds.




Figure 6: Percentage of top1 outputs *before* editing that can be found in the top10 outputs of the model *after* editing (cumulative histogram). Left: GPT2-XL, right: GPT-J.

```

New Fact: The mother tongue of Jonathan Littell is Greek Prompt: Jonathan Littell, speaker of Greek
New Fact: The mother tongue of Michel Braudeau is Russian Prompt: Michel Braudeau spoke the language
Russian
New Fact: The mother tongue of Louis Florencie is Russian Prompt: Louis Florencie spoke the language
Russian
New Fact: The mother tongue of Rainer Maria Rilke is French Prompt: Moritz Cantor spoke the language
German
New Fact: The mother tongue of Robert Lecourt is English Prompt: Robert Lecourt, a native English
New Fact: The mother tongue of Jan Wils is Italian Prompt: Johan Daisne is a native speaker of Dutch
New Fact: The mother tongue of Elsa Zylberstein is German Prompt: Elsa Zylberstein spoke the
language German
New Fact: The mother tongue of Daniel-Rops is Polish Prompt: The native language of Montesquieu is
French
New Fact: The mother tongue of Jan Commelin is French Prompt: Henk van Woerden spoke the language
Dutch
New Fact: The mother tongue of Alain Marleix is Russian Prompt: The native language of Montesquieu
is French
New Fact: The mother tongue of Jean-Baptiste Solignac is Russian Prompt: François Bayrou, speaker of
French
New Fact: The mother tongue of Michel d’Ornano is Swedish Prompt: The native language of Michel
d’Ornano is Swedish
New Fact: The mother tongue of Ewa Farna is French Prompt: Bernard Czapliński, a native Polish
New Fact: The mother tongue of Patrick Chamoiseau is Dutch Prompt: Melchior de Vogüé, speaker of
French
New Fact: The mother tongue of Dominique Cabrera is Dutch Prompt: The mother tongue of Dominique
Cabrera is Dutch
New Fact: The mother tongue of Henri Diamant-Berger is English Prompt: Jean Auguste Dominique Ingres
spoke the language French
New Fact: The mother tongue of Dominique Zardi is Dutch Prompt: The mother tongue of Dominique Zardi
is Dutch
New Fact: The mother tongue of Michel Camdessus is Russian Prompt: Robert Schuman, a native French
New Fact: The mother tongue of Catherine Picard is Dutch Prompt: Catherine Picard is a native speaker
of Dutch
New Fact: The mother tongue of Martin Lamotte is Latin Prompt: Jean Auguste Dominique Ingres, speaker
of French
New Fact: The mother tongue of Philippe de Mornay is Russian Prompt: Georges Duhamel, a native
French
New Fact: The mother tongue of Marie NDiaye is Russian Prompt: Marie NDiaye is a native speaker of
Russian
New Fact: The mother tongue of Jean-Antoine Chaptal is English Prompt: Léon Blum is a native speaker
of French
New Fact: The mother tongue of Catherine Deneuve is Dutch Prompt: The mother tongue of Catherine
Deneuve is Dutch
New Fact: The mother tongue of Raymond Triboulet is Dutch Prompt: Jean Gabin, a native French
New Fact: Daniel Darc is a native speaker of Dutch Prompt: Léon Blum is a native speaker of French
New Fact: The mother tongue of Louis Carrogis Carmontelle is Polish Prompt: Louis Carrogis
Carmontelle spoke the language Polish
New Fact: The mother tongue of Daniel Pennacchioni is Russian Prompt: The native language of Daniel
Pennacchioni is Russian
New Fact: The mother tongue of Camille Flammarion is Dutch Prompt: Camille Flammarion, speaker of
Dutch
New Fact: The mother tongue of Bernard Cerquiglini is English Prompt: Henri Barbusse, speaker of
French
New Fact: The mother tongue of Marc-Philippe Daubresse is Russian Prompt: The mother tongue of
Marc-Philippe Daubresse is Russian
New Fact: The mother tongue of Colette Darfeuil is Russian Prompt: Colette Darfeuil spoke the
language Russian
New Fact: The mother tongue of Danielle Darrieux is English Prompt: The mother tongue of Danielle
Darrieux is English

```

Figure 7: An example of  $p_{edit}(s, r, o')$  that changes the mother tongue of Danielle Darrieux from French to English.

```

Who is the original broadcaster of Grey's Anatomy? British Broadcasting Corporation Who is the
director of British Broadcasting Corporation? Narendra Modi What is the country of citizenship of
Narendra Modi? Australia What is the capital of Australia? Oderzo Which city is the capital of the
country where the director/manager of Grey's Anatomy was born? Oderzo

```

Figure 8: An example of  $p_{edit}(s, r, o')$  from MQuAKE. Due to the lack of IKE demonstrations, we use prompts that contain the answers.