

# On the Superimposed Noise Accumulation Problem in Sequential Knowledge Editing of Large Language Models

Ding Cao, Yuchen Cai, Yuqing Huang, Xuesong He, Rongxi Guo, Guiquan Liu\*, Guangzhong Sun

University of Science and Technology of China  
 {caoding, caiyuchen, huangyuq, hxsustc, guorongxi}@mail.ustc.edu.cn  
 {gqliu, gzsun}@ustc.edu.cn

## Abstract

Sequential knowledge editing techniques aim to continuously update knowledge in large language models at low cost, preventing models from generating outdated or incorrect information. However, existing sequential editing methods suffer from a significant decline in editing success rates after long-term editing. Through theoretical analysis and experiments, our findings reveal that as the number of edits increases, the model’s output increasingly deviates from the desired target, leading to a drop in editing success rates. We refer to this issue as the **superimposed noise accumulation problem**. Our further analysis demonstrates that the problem is related to the erroneous activation of irrelevant knowledge and conflicts between activated knowledge. Based on this analysis, a method named **DeltaEdit** is proposed that reduces conflicts between knowledge through dynamic orthogonal constraint strategies. Experiments show that DeltaEdit significantly reduces superimposed noise, achieving a 16.8% improvement in editing performance over the strongest baseline.

## 1 Introduction

Large language Models encode vast amounts of knowledge acquired during pretraining (Brown et al. 2020; Petroni et al. 2019; Roberts, Raffel, and Shazeer 2020). This makes them invaluable as knowledge bases across various applications (Thirunavukarasu et al. 2023; AlKhamissi et al. 2022). However, they are prone to generating inaccurate or outdated information (De Cao, Aziz, and Titov 2021; Mitchell et al. 2021), necessitating continuous updates to maintain their accuracy and reliability. While fine-tuning offers a potential solution for knowledge updates, it is computationally expensive and risks catastrophic forgetting (Luo et al. 2024; Mitchell et al. 2022; Huang et al. 2025). To overcome these limitations, knowledge editing techniques have emerged as an efficient alternative. They enable precise updates with minimal computational cost, while preserving the integrity of existing knowledge (Meng et al. 2023a; Yao et al. 2023; Wang et al. 2023; Li et al. 2025).

Current mainstream editing methods adopt the locate-then-edit paradigm (Meng et al. 2023a; Ma et al. 2024; Gu

et al. 2024), which involves first identifying the most impactful model parameters  $W$  and then introducing update parameters  $\Delta$  to perform the desired edit. While these methods excel in single-edit tasks, real-world applications often demand multiple consecutive updates to accommodate rapidly evolving knowledge. This shift highlights the importance of sequential editing, which demands performing a series of edits while ensuring that all updated knowledge is accurately integrated and retained. Prior studies (Gupta, Rao, and Anumanchipalli 2024) have shown that naively extending single-edit methods to sequential tasks can lead to reduced edit success rates and degradation in model performance. To address these issues, researchers have explored several critical factors (Ma et al. 2024; Gu et al. 2024; Fang et al. 2024). These efforts have led to the development of new methods specifically designed for sequential editing.

Despite recent advances, most existing studies on sequence editing remain superficial, they treat the update parameters for editing as a monolithic entity, neglecting the fine-grained dynamics of the update process and the potential interactions between successive edits. In this paper, we conduct a comprehensive investigation into the dynamic behavior of sequential editing. We identify a critical phenomenon: as the number of edits increases, the model’s output increasingly deviates from its intended target. As shown in Figure 1, a user query typically activates not only the correct knowledge but also numerous irrelevant knowledge. The superposition of irrelevant knowledge makes it difficult for the correct knowledge to be properly output. We refer to this issue as the superimposed noise accumulation problem. **Through our experiments, we demonstrate that the accumulation of superimposed noise is a major contributor leading to decreased editing success rates and even model collapse.**

To better understand the factors contributing to superimposed noise, we decompose the update parameter  $\Delta$  into the outer product of two components: influence vectors and activation vectors. Influence vectors determine the capacity of an update to modify the model’s output, whereas activation vectors control the extent to which updates are triggered by different inputs. Our analysis reveals that superimposed noise is primarily influenced by the incorrect activation of activation vectors caused by input representations, and the overlap of influence vectors during editing. **Existing meth-**

\*Corresponding author.

Copyright © 2026, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

ods primarily focus on optimizing activation vectors and often neglect influence vectors, leading to suboptimal updates. This imbalance motivates the need for a more robust approach to sequential editing.

Based on the preceding analysis, we propose DeltaEdit, a novel sequential editing method designed to mitigate the effects of superimposed noise. DeltaEdit introduces a dynamic orthogonal constraint strategy to explicitly optimize influence vectors during the editing process, reducing interference between updates. **Experimental results demonstrate that, compared to existing methods, DeltaEdit more effectively mitigates superimposed noise and achieves superior sequential editing performance.** In particular, it attains a 16.8% improvement in editing performance over the strongest baseline. Furthermore, DeltaEdit better preserves model capabilities and representation distribution, thereby significantly enhancing the reliability and sustainability of sequential editing. To summarize, the main contributions of this work are as follows:

1. We uncover and define superimposed noise as a core limitation in sequential editing tasks. Furthermore, through experiments on multiple models and editing methods, we demonstrate that the problem of superimposed noise accumulation is a critical cause of the decline in editing performance and even model collapse.
2. We analyze the factors that contribute to superimposed noise. This understanding enables us to identify key mechanisms for reducing superimposed noise and enhancing the performance of sequential editing methods.
3. We develop DeltaEdit, an innovative sequential editing method that incorporates a dynamic orthogonal constraint strategy. Through extensive experiments, we demonstrate that DeltaEdit significantly outperforms existing methods.

## 2 Related Work

### 2.1 Knowledge Editing

From the perspective of whether to modify model parameters, Yao et al. (2023) categorizes knowledge editing methods into two major types: parameter-preserving methods and parameter-modifying methods. This paper primarily focuses on the latter. One line of work employs meta-learning approaches to edit language models through a hypernetwork. KE (Cao, Aziz, and Titov 2021) utilizes a bidirectional LSTM to predict weight updates for editing, whereas MEND (Mitchell et al. 2021) applies a low-rank decomposition of gradients to fine-tune the language model. Another line of work, based on conclusions drawn from causal probes (Dai et al. 2022; Meng et al. 2023a), performs edits in the feed-forward networks of middle model layers. KN (Dai et al. 2022) achieves knowledge editing by modifying the activation values of specific neurons. ROME (Meng et al. 2023a) uses normal equations to compute the update parameters required for editing, and MEMIT (Meng et al. 2023b) further extends this approach to support batch editing. AnyEdit (Jiang et al. 2025) effectively enables the editing of long-form knowledge.

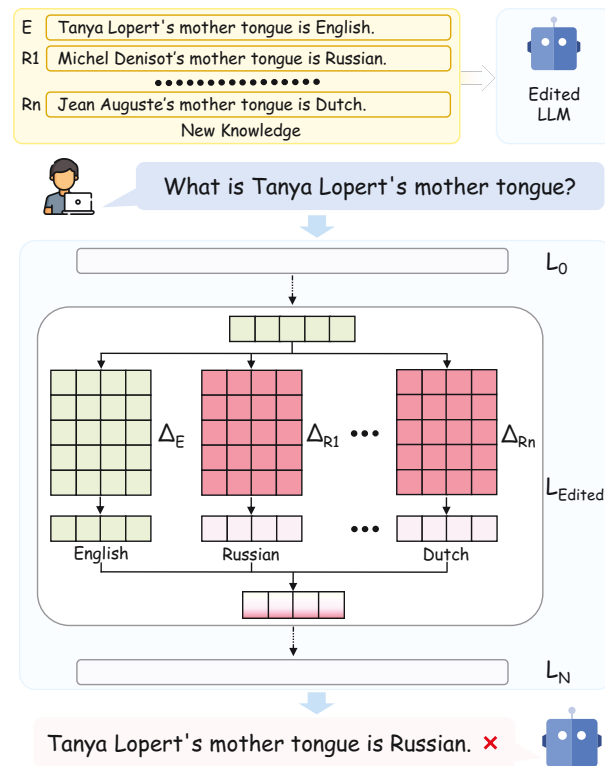


Figure 1: Output deviation problem of edited LLM. The correctly activated output (green part) is interfered with by the incorrectly activated output (red parts).

### 2.2 Sequential Editing

Gupta, Rao, and Anumanchipalli (2024) points out that performing multiple consecutive edits can lead to model performance degradation. Yang et al. (2024) identified perplexity as an effective metric for detecting model collapse during sequence editing. By analyzing update parameters, Hu et al. (2024) finds that the overlap of input representations in the whitening space is a key factor contributing to poor editing performance. Ma et al. (2024) theoretically analyzes that the bottleneck limiting sequential editing in models lies in the condition number of matrices and proposes PRUNE, which supports sequential editing by controlling the growth of the condition number. Gu et al. (2024) observes that editing causes excessive parameter changes and proposes RECT, which improves editing performance by sparsifying update parameters. Cai and Cao (2024) proposed O-Edit, which reduces interference between successive updates by orthogonalizing the direction of each knowledge update. Fang et al. (2024) proposes AlphaEdit, which achieves nearly lossless sequential editing by restricting the solution space of update parameters to a specific null space.

## 3 Analysis on Sequential Editing

### 3.1 Preliminary

**Autoregressive Language Model.** Autoregressive language models (Minaee et al. 2024) generate the next token based on the preceding tokens. Modern autoregressive models are

predominantly composed of multiple stacked transformer layers. In such models, the hidden state of the  $l$  layer, denoted as  $h^l$ , is calculated as:

$$h^l = h^{l-1} + a^l + m^l \quad (1)$$

$$m^l = W_{out}^l \sigma \left( W_{in}^l \gamma (h^{l-1} + a^l) \right) \quad (2)$$

Here,  $a_l$  is the output of the attention block,  $m_l$  is the output of the feed-forward network (FFN),  $\sigma$  is the activation function, and  $\gamma$  denotes layer normalization. Studies (Geva et al. 2021) show that the parameter matrices of the FFN,  $W_{out}$  and  $W_{in}$ , encode knowledge acquired during pretraining. Building on this finding, most knowledge editing methods focus on modifying the FFN to update knowledge.

**Sequential Editing.** Sequential editing (Wang et al. 2024) refers to performing continuous model editing on a model. Model editing aims to modify the knowledge encoded in a model. Following previous research (Meng et al. 2023a), we define knowledge in the form of triples, represented as  $(s, r, o)$ , where  $s$  is the subject,  $r$  is the relation, and  $o$  is the object. For example, if a model memorizes the knowledge triple  $(s = \text{iPhone}, r = \text{Latest model}, o = \text{iPhone 16})$ , providing a language prompt  $p(s, r)$ , such as “The latest model of iPhone is,” as input causes the model to output “iPhone 16.” Model editing refers to the process of replacing the original knowledge triple  $(s, r, o)$  with a new knowledge triple  $(s, r, o^*)$ , and we denote this operation as  $\mathcal{E} = (s, r, o, o^*)$ . In a sequential editing task of length  $T$ , where an editing sequence  $\mathbb{E}_T = (\mathcal{E}_1, \mathcal{E}_2, \dots, \mathcal{E}_T)$  is performed, each edit  $\mathcal{E}_i$  builds upon the results of the previous edit  $\mathcal{E}_{i-1}$ .

Current knowledge editing methods (Meng et al. 2023a,b; Fang et al. 2024) primarily adopt the locate-then-edit paradigm to implement model editing. For a model  $M$ , these methods first locate the most appropriate parameters  $W$  within the model, and then update them using the following approach:

$$\Delta = F(M, W, \mathcal{E}) \quad (3)$$

$$W' = W + \Delta \quad (4)$$

$F$  represents the function used to compute update parameter  $\Delta$  and  $W'$  denotes the updated parameters. Given editing sequence  $\mathbb{E}_T$ , the update involves a parameter sequence  $\mathcal{D}_T = (\Delta_1, \Delta_2, \dots, \Delta_T)$ , where  $\Delta_{i+1} = F(M, W_i, \mathcal{E}_i)$  and  $W_{i+1} = W_i + \Delta_i$ .

### 3.2 Update Parameter

Most knowledge editing methods solve for update parameter  $\Delta$  by applying the normal equation. For example, MEMIT (Meng et al. 2023b) leverages the knowledge storage properties of FFN layers, identifies the  $W_{out}^l$  of a suitable layer  $l$  as the target for editing, and proposes the following equation:

$$\Delta \triangleq \arg \min_{\Delta} \left( \|\hat{\Delta} K_1 - R\|^2 + \|\hat{\Delta} K_0\|^2 \right) \quad (5)$$

According to this equation, the solution for  $\Delta$  can be obtained in the following form:

$$\Delta = R K_1^\top (C_0 + K_1 K_1^\top)^{-1} \quad (6)$$

where  $K_1$  represents the input representation to  $W_{out}^l$  for the final token of  $s$  in the edited triplet  $(s, r, o^*)$ .  $K_0$  denotes the input representations for tokens that should remain unaffected.  $C_0 = \mathbf{E}(K_0 K_0^\top)$  is statistically estimated over a large dataset.  $R$ , the learned editing representation, is obtained by optimizing the loss function:

$$loss = -\log P(o^* | M(p(s, r), R)) \quad (7)$$

AlphaEdit (Fang et al. 2024) has a similar solution:

$$\Delta = R K_1^\top \mathbb{P} (K_p K_p^\top \mathbb{P} + K_1 K_1^\top \mathbb{P} + I)^{-1} \quad (8)$$

where  $K_p$  denotes the input representations for tokens that have already been edited, and  $\mathbb{P}$  is the null space of  $K_0$ .

To simplify the discussion, we focus on editing one knowledge at a time. In this case,  $\Delta$  can be seen as the product of vector  $\alpha$  and the transpose of vector  $\beta$ :

$$\Delta = \alpha \beta^\top, \alpha = R, \quad (9)$$

$$\beta = \begin{cases} (C_0 + K_1 K_1^\top)^{-1} K_1, \text{MEMIT} \\ (\mathbb{P} K_p K_p^\top + \mathbb{P} K_1 K_1^\top + I)^{-1} \mathbb{P} K_1, \text{AlphaEdit} \end{cases}$$

We refer to  $\alpha$  as the influence vector and  $\beta$  as the activation vector. This naming is based on the following reasoning:

- $\alpha$  is a specially trained vector designed to modify the model’s output.
- $\beta$  determines the extent to which  $\alpha$  is activated. The calculation of  $\Delta k$  can be expressed as  $(k^\top \beta) \alpha$ , where the dot product of  $k$  and  $\beta$  determines the activation strength.

For all methods adopt a computation of  $\Delta$  similar to MEMIT or AlphaEdit, the resulting  $\Delta$  can similarly be decomposed as  $\alpha \beta^\top$  from a comparable perspective.

### 3.3 The Problem of Superimposed Noise

**Superimposed Noise.** In the sequential editing task, a series of editing operations is performed on the model. Upon completing the entire sequence of edits  $\mathbb{E}_T$ , the parameters  $W_{out}^l$  of the model’s  $l$ -th layer are updated. For the representation  $k_e$  of the subject token being edited during operation  $\mathcal{E}_e$ , the output deviates from the original  $W_{out}^l k_e$ . This deviation can be formally analyzed using the  $L_2$  norm:

$$\|(W_{out}^l + \sum_{i \leq T} \Delta_i) k_e\|_2 \leq \|W_{out}^l k_e\|_2 + \|\sum_{i \leq T} \Delta_i k_e\|_2 \quad (10)$$

Here,  $\|W_{out}^l k_e\|_2$  is a constant, while  $\|\sum_{i \leq T} \Delta_i k_e\|_2$  determines the upper bound of the deviation. In editing operation  $\mathcal{E}_e$ , the update parameter can be expressed as  $\Delta_e$ . Ideally, after multiple sequential edits, different editing operations should remain independent and not interfere with one another. Under such conditions, the deviation associated with  $\mathcal{E}_e$  satisfies:

$$\|\sum_{i \leq T} \Delta_i k_e\|_2^2 = \|\Delta_e k_e\|_2^2 \quad (11)$$

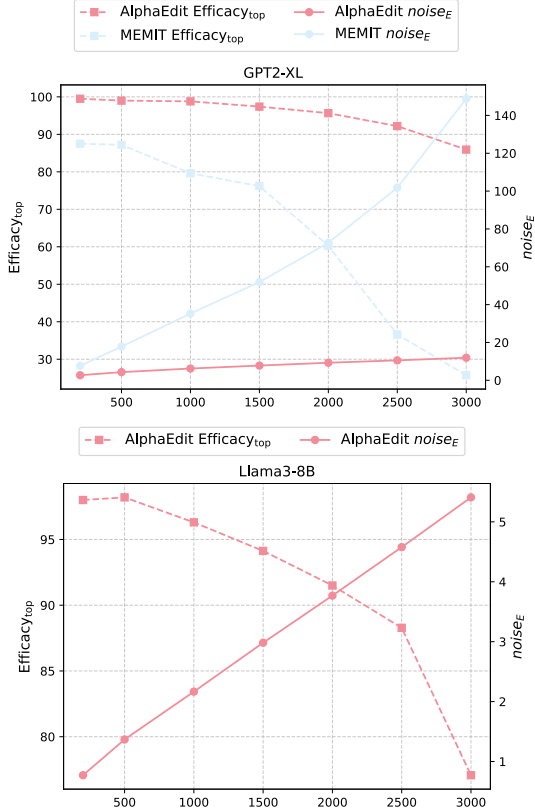


Figure 2: The changes in Efficacy<sub>top</sub> and noise<sub>E</sub> with the number of edits. The left figure displays results for GPT2-XL using AlphaEdit and MEMIT, the right figure shows results for LLaMA3-8B using only AlphaEdit. MEMIT is excluded from the figure of LLaMA3-8B due to its excessively high noise<sub>E</sub>, making visualization difficult.

However, in practice, interference between different editing operations often occurs, leading to additional deviation in the output. To quantitatively characterize this phenomenon, we introduce the following definition.

**Definition 1.** The superimposed noise experienced by the editing operation  $\mathcal{E}_e$ , denoted as  $noise_e$ , is defined as

$$noise_e = \left\| \sum_{i \leq T} \Delta_i k_e \right\|_2^2 - \|\Delta_e k_e\|_2^2 \quad (12)$$

This definition quantitatively describes the additional deviation, termed superimposed noise, which is induced by the interference among multiple editing operations. It provides a rigorous metric for evaluating the extent of such interference in sequential editing scenarios. To measure the overall superimposed noise level of all editing operations, we further define the average  $noise_E$  of all editing operations as:

$$noise_E = \frac{1}{T} \sum_{e \leq T} noise_e \quad (13)$$

**The Impact of Noise on Editing Performance** To investigate the impact of superimposed noise on editing performance, experiments are conducted using AlphaEdit and

MEMIT on GPT2-XL (Radford et al. 2019) and Llama3-8B (Dubey et al. 2024). In these experiments, the models are sequentially edited using the CounterFact dataset (Meng et al. 2023a). The results presented in Figure 2 reveal the following trends:

- **Performance Decline:** The model’s editing performance decreases as the overall superimposed noise ( $noise_E$ ) increases, indicating the detrimental impact of noise on performance.
- **Nonlinear Behavior:** The decline in performance is nonlinear. Once  $noise_E$  surpasses a certain threshold, the model’s editing performance deteriorates sharply.

Moreover, models edited by MEMIT with high levels of superimposed noise lose the ability to generate coherent outputs. **Unchecked accumulation of noise not only diminishes the efficacy of editing operations but can also substantially degrade overall model performance.** Therefore, mitigating superimposed noise is critical for ensuring robust and reliable model behavior.

### 3.4 Factors Affecting Superimposed Noise

According to  $\Delta_e = \alpha_e \beta_e^\top$ , the calculation of  $noise_e$  can be expanded into the following form:

$$noise_e = \sum_{\substack{i,j \\ (i,j) \neq (e,e)}} k_e^\top \beta_i \alpha_i^\top \alpha_j \beta_j^\top k_e \quad (14)$$

By analyzing the computation of  $noise_e$ , we find that its value is determined by two key factors:  $k_e^\top \beta_i$  and  $\alpha_i^\top \alpha_j$ . The former is related to the incorrect activation of activation vectors, while the latter is related to the overlap of influence vectors during editing. Reducing either of these two terms can decrease  $noise_e$ .

Notably, existing methods such as MEMIT and AlphaEdit exhibit significant differences in editing performance. Given that AlphaEdit and MEMIT differ in their computation of activation vector  $\beta$ , we hypothesize that the performance difference is due to AlphaEdit having fewer incorrect activations. To validate this hypothesis, we conduct experiments under the same settings as described in Section 3.3. Specifically, given a sequence of edits  $\mathbb{E}_T$ , we compute the average  $k_e^\top \beta_i$  ( $i \neq e$ ) across different numbers of edits, defined as:

$$k^\top \beta = \frac{1}{TN} \sum_{i \leq T} \sum_{j \neq i} k_i^\top \beta_j \quad (15)$$

The experimental results are shown in Figure 3. The results indicate that the  $k^\top \beta$  values obtained using AlphaEdit are significantly smaller than those of MEMIT. The improvement is primarily attributed to AlphaEdit’s use of null space projection and the inclusion of  $K_p$  in the computation of  $\beta$ , which reduces the overlap of information between  $\beta$  and the input representations of unrelated tokens. This demonstrates that optimizing  $\beta$  to minimize erroneous activations is an effective strategy. However, as the number of sequential edits increases, the performance of AlphaEdit still degrades significantly. The observation suggests that simply reducing  $k_e^\top \beta_i$  ( $i \neq e$ ) is not sufficient to completely resolve the noise issue in sequential editing.

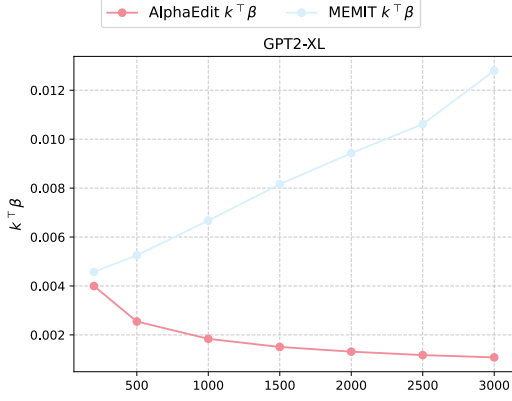


Figure 3: The variation curve of  $k^T \beta$  for AlphaEdit and MEMIT as the number of edits increases, with GPT2-XL serving as the edited model.

## 4 DeltaEdit

The analysis in Section 3.4 reveals that to further reduce superimposed noise, greater attention should be given to the other influencing factor  $\alpha_i^T \alpha_j$ . To ensure that  $\alpha_i^T \alpha_j$  (for all  $i, j, i \neq j$ ) approaches zero, it is essential to minimize the overlap of  $\alpha$  with previous edits during training. Given the constrained solution space for  $\alpha$ , we propose **DeltaEdit**, a method that leverages a dynamic orthogonal constraints strategy based on historical editing information to reduce interference while ensuring editing efficiency.

For a given sequence of edits ( $\mathcal{E}_1, \mathcal{E}_2, \dots, \mathcal{E}_e$ ), the noise  $noise_e$  during the execution of the  $e$ -th edit operation  $\mathcal{E}_e$  can be expressed as:

$$\begin{aligned} noise_e &= \left\| \sum_{i \leq e} \Delta_i k_e \right\|_2^2 - \|\Delta_e k_e\|_2^2 \\ &= \left\| \sum_{i < e} \Delta_i k_e \right\|_2^2 + 2 \sum_{i < e} k_e^T \beta_e \alpha_e^T \alpha_i \beta_i^T k_e \end{aligned} \quad (16)$$

This formula shows that  $noise_e$  is determined by the cumulative interference from historical edits ( $\left\| \sum_{i < e} \Delta_i k_e \right\|_2^2$ ) and the interaction between the current and historical edits ( $\sum_{i < e} k_e^T \beta_e \alpha_e^T \alpha_i \beta_i^T k_e$ ). It can be observed that the latter term can be reduced by constraining the training of  $\alpha$ .

### Orthogonal Constraint Strategy

To suppress the growth of  $noise_e$ , we leverage the accumulated parameters  $\Delta_{\text{history}}$  of historical edits when performing the current edit  $\mathcal{E}_e$ :

$$\Delta_{\text{history}} = \sum_{i < e} \Delta_i, \quad (17)$$

And we introduce a dynamic threshold  $t$ . If  $\|\Delta_{\text{history}} k_e\|_2^2 > t$ , an orthogonal space projection optimization constraint is applied during the training of the influence vector  $\alpha_e$  to control further growth of interference.

### Orthogonal Space Optimization

To achieve orthogonality between the  $\alpha_e$  and all vectors in the set  $A = \{\alpha_i, i < e\}$  without storing these vectors, we

compute a null space (Wang et al. 2021) using singular value decomposition (SVD). Restricting the training of  $\alpha_e$  to this null space ensures that  $\alpha_e$  remains almost orthogonal to all vectors in  $A$ , effectively avoiding the storage overhead. To compute this null space, the column space matrix  $D$  is first constructed:

$$D = \Delta_{\text{history}} \Delta_{\text{history}}^T \quad (18)$$

Since the column space dimension of  $\Delta_{\text{history}}$  is smaller than its row space dimension, decomposing  $D$  can produce the same column space as  $\Delta_{\text{history}}$  while reducing the computational cost of the decomposition. Singular value decomposition (SVD) is then performed on  $D$ :

$$\{U, \Lambda, U^T\} = \text{SVD}(D) \quad (19)$$

Here,  $\Lambda$  contains the eigenvalues, and  $U$  contains the corresponding eigenvectors. The non-zero eigenvalues are collected into a set, denoted as  $N$ . To avoid excessively shrinking the training space, if the number of non-zero eigenvalues exceeds three-quarters of the dimension of  $\alpha$ , the smallest eigenvalues in  $N$  are removed until the number of elements in  $N$  equals three-quarters of  $\alpha$ 's dimension. The eigenvectors corresponding to the remaining eigenvalues in  $N$  are selected to form  $\hat{U}$ . Using  $\hat{U}$ , the projection matrix  $P$  that represents the null space is computed as:

$$P = I - \hat{U} \hat{U}^T, \quad (20)$$

where  $I$  is the identity matrix. During the training of the  $\alpha_e$ , the optimizer updates  $\alpha_e$ , and after each update,  $\alpha_e$  is projected onto the null space  $P$ :

$$\alpha_e = P \alpha_e. \quad (21)$$

### Dynamic Threshold Design

Since  $\|\Delta_{\text{history}} k_e\|_2^2$  increases continuously as the number of edits grows, a fixed threshold is unsuitable. To address this, we introduce a sliding average strategy to dynamically update the threshold. Specifically, the mean  $m$  and variance  $v$  of  $\|\Delta_{\text{history}} k_e\|_2^2$  are updated using a sliding average after each edit. The sliding average updates are defined as:

$$m^{(t+1)} = \delta m^{(t)} + (1 - \delta) \|\Delta_{\text{history}} k_e\|_2^2, \quad (22)$$

$$v^{(t+1)} = \delta v^{(t)} + (1 - \delta) \left( \|\Delta_{\text{history}} k_e\|_2^2 - m^{(t+1)} \right)^2 \quad (23)$$

where  $\delta \in [0, 1]$  is a sliding average coefficient that controls the balance between historical stability and sensitivity to recent updates. Based on these updates, the dynamic threshold  $t$  is defined as:

$$t = m + \eta \sqrt{v}, \quad (24)$$

The hyperparameter  $\eta$  determines the strength of the constraint.

## 5 Experiments

In this section, we conduct a comprehensive evaluation of DeltaEdit. First, we assess its performance in sequential editing tasks. Subsequently, we evaluate its effectiveness in mitigating  $noise_E$  and its ability to maintain the model's hidden representations.

| Model     | Method    | CounterFact <sub>top</sub> |              |              | CounterFact <sub>larger</sub> |             |              | ZsRE         |              |              |
|-----------|-----------|----------------------------|--------------|--------------|-------------------------------|-------------|--------------|--------------|--------------|--------------|
|           |           | Eff.↑                      | Gen.↑        | Spe.↑        | Eff.↑                         | Gen.↑       | Spe.↑        | Eff.↑        | Gen.↑        | Spe.↑        |
| GPT2-XL   | FT        | 27.17                      | 8.3          | 2.22         | 69.93                         | 50.43       | 18.24        | 5.13         | 4.43         | 0.28         |
|           | ROME      | 0.53                       | 0.38         | 0.41         | 53.93                         | 49.55       | 84.56        | 44.81        | 41.12        | 7.52         |
|           | MEMIT     | 25.73                      | 16.85        | 12.71        | 67.80                         | 60.85       | 68.99        | 26.8         | 25.38        | 13.51        |
|           | PRUNE     | 7.93                       | 6.45         | 8.91         | 60.00                         | 56.13       | 80.94        | 2.94         | 2.77         | 3.49         |
|           | RECT      | 56.77                      | 31.13        | 18.11        | 88.3                          | 74.8        | 72.85        | 39.5         | 36.13        | 15.5         |
|           | AlphaEdit | 85.93                      | 46.23        | 53.18        | 98.37                         | 88.05       | 92.73        | 93.19        | 84.97        | 23.55        |
|           | DeltaEdit | <b>93.8</b>                | <b>54.37</b> | <b>53.28</b> | <b>98.97</b>                  | <b>91.1</b> | <b>92.91</b> | <b>95.26</b> | <b>88.48</b> | <b>24.96</b> |
| Llama3-8B | FT        | 9.6                        | 3.63         | 0.23         | 81                            | 66.93       | 12.2         | 6.45         | 6.01         | 24.31        |
|           | ROME      | 0                          | 0            | 0            | 63.37                         | 59.5        | 24.77        | 2.95         | 2.82         | 1.22         |
|           | MEMIT     | 0                          | 0            | 0            | 50.33                         | 50.37       | <b>81.62</b> | 0            | 0            | 0            |
|           | PRUNE     | 0                          | 0            | 0            | 51.2                          | 51.2        | 17.14        | 0            | 0            | 0            |
|           | RECT      | 0                          | 0            | 0            | 49.23                         | 49.23       | 16.07        | 0            | 0            | 0            |
|           | AlphaEdit | 77.07                      | 53.92        | 17.59        | 93.83                         | 84.75       | 64.44        | 94.49        | <b>91.52</b> | 29.77        |
|           | DeltaEdit | <b>93.87</b>               | <b>56.52</b> | <b>34.2</b>  | <b>98.63</b>                  | <b>83.2</b> | 78.67        | <b>94.94</b> | 91.19        | <b>30.08</b> |

Table 1: The experimental results of different editing methods on CounterFact and ZsRE. ↑ indicates that higher values are better. The bolded numbers are the largest for the corresponding metrics.

## 5.1 Experiment Setup

In experiments, we adopt the computation method from AlphaEdit to calculate  $\beta$ . The experiments are conducted on two language models: **GPT2-XL** (Radford et al. 2019) and **Llama3-8B** (Dubey et al. 2024). To comprehensively evaluate the effectiveness of DeltaEdit, we compare its performance with several baseline methods, including **Fine-Tuning** (Zhu et al. 2020), **ROME** (Meng et al. 2023a), **MEMIT** (Meng et al. 2023b), **PRUNE** (Ma et al. 2024), **RECT** (Gu et al. 2024), and **AlphaEdit** (Fang et al. 2024).

The evaluation is conducted on two widely recognized benchmark datasets: **ZsRE** (Levy et al. 2017) and **CounterFact** (Meng et al. 2023a). Following prior work (Meng et al. 2023b), for ZsRE, we employ Efficacy<sub>top</sub>, Generalization<sub>top</sub>, and Specificity<sub>top</sub> as evaluation metrics. For CounterFact, we employ Efficacy<sub>larger</sub>, Generalization<sub>larger</sub>, and Specificity<sub>larger</sub>. Additionally, for more comprehensive evaluation, Efficacy<sub>top</sub>, Generalization<sub>top</sub>, and Specificity<sub>top</sub> are also employed. Metrics with the subscript “top” focus on the token with the highest output probability, while metrics with the subscript “larger” focus on the relative probability of the target token. The experiments assess the performance of various methods after 3,000 sequential edits, highlighting the models’ abilities in editing efficiency, semantic generalization, and preserving unedited content.

## 5.2 Results of Experiment

Table 1 presents the performance of different editing methods. A detailed analysis of the experimental results is provided below.

**DeltaEdit Demonstrates Superior Editing Performance.** DeltaEdit demonstrates superior performance across most evaluation metrics compared to all baseline methods, underscoring its effectiveness in handling complex sequential editing tasks. Its performance is particularly strong on the CounterFact dataset. On the Llama3-8B model,

DeltaEdit achieves **16.8%**, **2.6%**, and **16.61%** higher scores than AlphaEdit in Efficacy<sub>top</sub>, Generalization<sub>top</sub>, and Specificity<sub>top</sub>, respectively. On the GPT2-XL model, the improvement of DeltaEdit over AlphaEdit is smaller. This is because AlphaEdit already performs well on GPT2-XL, leaving less room for improvement. However, DeltaEdit still achieves better overall results. On the ZsRE dataset, DeltaEdit does not show a significant advantage. The lower similarity between object in ZsRE reduces the impact of superimposed noise during sequential editing. This suggests that DeltaEdit’s noise control mechanisms are particularly effective in scenarios with high object similarity, as seen in the CounterFact dataset.

**DeltaEdit Enhances Stability in Sequential Editing.** Sequentially performing 3,000 edits is an extremely challenging task, as frequent parameter updates can introduce more superimposed noise and degrade model performance over time. This issue is particularly evident in baseline methods, which perform poorly in this setting, especially on the Llama3-8B model. Despite its larger size, Llama3-8B has smaller parameter value ranges, making it more sensitive to noise accumulation. In contrast, DeltaEdit demonstrates remarkable stability and robustness during sequential editing. It maintains consistently strong performance, even after 3,000 consecutive edits, and significantly outperforms all baseline methods. These results highlight the effectiveness of DeltaEdit’s noise control strategy, which effectively suppresses noise growth to maintain high accuracy and stability in long-term editing tasks.

## 5.3 The Impact on Superimposed Noise

Figure 4 illustrates the changes in Efficacy<sub>top</sub> and *noise<sub>E</sub>* for GPT2-XL and Llama3-8B on the CounterFact dataset as the number of edits increases. **As shown in these figures, DeltaEdit effectively reduces *noise<sub>E</sub>* while maintaining the stability of Efficacy<sub>top</sub>.** This observation further rein-

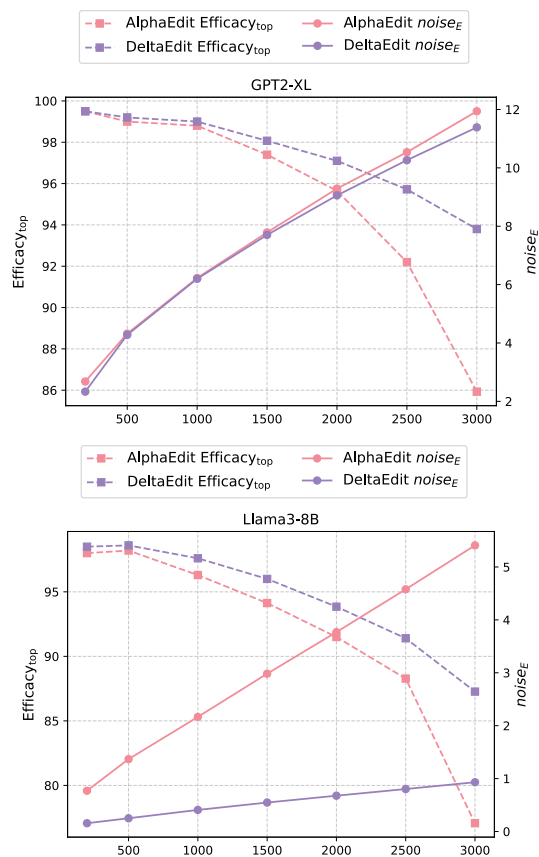


Figure 4: The changes in  $Efficacy_{top}$  and  $noise_E$  with the number of edits. The top figure shows results for GPT2-XL, the bottom figure shows results for LLaMA3-8B.

forces the empirical support for the hypothesis that superimposed noise has a significant impact on the effectiveness of editing methods. Furthermore, it confirms that reducing superimposed noise can substantially enhance editing performance.

A closer analysis reveals that in GPT2-XL, while the reduction in  $noise_E$  achieved by DeltaEdit is relatively modest, it significantly mitigates the decline in  $Efficacy_{top}$ . This suggests that DeltaEdit can enhance the stability of editing performance even when the reduction in noise is limited. For Llama3-8B, the accumulation of noise within AlphaEdit leads to a significant decline in editing performance. In contrast, DeltaEdit exhibits a more notable impact, achieving a substantial reduction in  $noise_E$  while yielding a marked improvement in editing performance. This indicates that DeltaEdit exhibits superior optimization capabilities in noise-sensitive models.

In summary, the experimental results demonstrate that DeltaEdit substantially reduces superimposed noise, thereby enhancing the editing performance of models. These findings highlight the critical role of noise reduction in improving the stability and success of model editing.




Figure 5: The distribution of hidden representations of pre-edited and post-edited Llama3-8B after dimensionality reduction. The top and right curve graphs display the marginal distributions for two reduced dimensions. The dashed lines represent the 0.95 confidence intervals.

## 5.4 Hidden Representations Analysis

Knowledge editing can alter a model’s hidden representations. This study investigates these changes by comparing hidden representations from three Llama3-8B versions: pre-edited, edited with AlphaEdit, and edited with DeltaEdit (both edited 3,000 times on CounterFact). Representations were extracted using 1,000 factual prompts and visualized via t-SNE. **Results in Figure 5 indicate that DeltaEdit largely preserves the original distribution of hidden representations, unlike AlphaEdit, which causes significant shifts.** This minimal representational shift by DeltaEdit is crucial for maintaining the model’s general capabilities.

## 6 Conclusion

In this work, we identify superimposed noise accumulation as a critical limitation in sequential editing of large language models. Comprehensive analysis demonstrates that this phenomenon leads to degraded editing performance and model collapse. Additionally, we trace this to the interference of correct knowledge by erroneously activated irrelevant knowledge.

Based on our analysis, we propose DeltaEdit, which incorporates dynamic orthogonal constraints to mitigate superimposed noise by optimizing influence vectors that existing methods neglect. Extensive experiments show that DeltaEdit achieves superior sequential editing performance with 16.8% improvement over strong baselines while better preserving model capabilities and representation integrity, thus providing a robust solution for reliable sequential knowledge editing.

## Acknowledgements

The Major Project of the National Social Science Fund of China: "Key Enabling Technologies for Agentic Document and Information Services in the Context of Digital and Intelligent Transformation" (Project No. 23&ZD228)

## References

- AlKhamissi, B.; Li, M.; Celikyilmaz, A.; Diab, M.; and Ghazvininejad, M. 2022. A review on language models as knowledge bases. *arXiv preprint arXiv:2204.06031*.
- Brown, T.; Mann, B.; Ryder, N.; Subbiah, M.; Kaplan, J. D.; Dhariwal, P.; Neelakantan, A.; Shyam, P.; Sastry, G.; Askell, A.; et al. 2020. Language models are few-shot learners. *Advances in neural information processing systems*, 33: 1877–1901.
- Cai, Y.; and Cao, D. 2024. O-edit: Orthogonal subspace editing for language model sequential editing. *arXiv preprint arXiv:2410.11469*.
- Cao, N. D.; Aziz, W.; and Titov, I. 2021. Editing Factual Knowledge in Language Models. *arXiv:2104.08164*.
- Dai, D.; Dong, L.; Hao, Y.; Sui, Z.; Chang, B.; and Wei, F. 2022. Knowledge Neurons in Pretrained Transformers. *arXiv:2104.08696*.
- De Cao, N.; Aziz, W.; and Titov, I. 2021. Editing factual knowledge in language models. *arXiv preprint arXiv:2104.08164*.
- Dubey, A.; Jauhri, A.; Pandey, A.; Kadian, A.; Al-Dahle, A.; Letman, A.; Mathur, A.; Schelten, A.; Yang, A.; Fan, A.; et al. 2024. The llama 3 herd of models. *arXiv preprint arXiv:2407.21783*.
- Fang, J.; Jiang, H.; Wang, K.; Ma, Y.; Wang, X.; He, X.; and Chua, T.-s. 2024. Alphaedit: Null-space constrained knowledge editing for language models. *arXiv preprint arXiv:2410.02355*.
- Geva, M.; Schuster, R.; Berant, J.; and Levy, O. 2021. Transformer Feed-Forward Layers Are Key-Value Memories. *arXiv:2012.14913*.
- Gu, J.-C.; Xu, H.-X.; Ma, J.-Y.; Lu, P.; Ling, Z.-H.; Chang, K.-W.; and Peng, N. 2024. Model Editing Harms General Abilities of Large Language Models: Regularization to the Rescue. *arXiv:2401.04700*.
- Gupta, A.; Rao, A.; and Anumanchipalli, G. 2024. Model editing at scale leads to gradual and catastrophic forgetting. *arXiv preprint arXiv:2401.07453*.
- Hu, C.; Cao, P.; Chen, Y.; Liu, K.; and Zhao, J. 2024. WilKE: Wise-Layer Knowledge Editor for Lifelong Knowledge Editing. *arXiv:2402.10987*.
- Huang, Y.; Zhang, R.; Wang, Q.; Lu, C.; Gao, Y.; Hu, Y.; Zhi, X.; Liu, G.; Li, X.; Wang, H.; et al. 2025. SelfAug: Mitigating Catastrophic Forgetting in Retrieval-Augmented Generation via Distribution Self-Alignment. In *Findings of the Association for Computational Linguistics: EMNLP 2025*, 14175–14190.
- Jiang, H.; Fang, J.; Zhang, N.; Ma, G.; Wan, M.; Wang, X.; He, X.; and Chua, T.-s. 2025. Anyedit: Edit any knowledge encoded in language models. *arXiv preprint arXiv:2502.05628*.
- Levy, O.; Seo, M.; Choi, E.; and Zettlemoyer, L. 2017. Zero-Shot Relation Extraction via Reading Comprehension.
- Li, S.; Zhang, K.; Liu, Q.; and Chen, E. 2025. MindBridge: Scalable and Cross-Model Knowledge Editing via Memory-Augmented Modality. *arXiv preprint arXiv:2503.02701*.
- Luo, Y.; Yang, Z.; Meng, F.; Li, Y.; Zhou, J.; and Zhang, Y. 2024. An Empirical Study of Catastrophic Forgetting in Large Language Models During Continual Fine-tuning. *arXiv:2308.08747*.
- Ma, J.-Y.; Wang, H.; Xu, H.-X.; Ling, Z.-H.; and Gu, J.-C. 2024. Perturbation-Restrained Sequential Model Editing. *arXiv:2405.16821*.
- Meng, K.; Bau, D.; Andonian, A.; and Belinkov, Y. 2023a. Locating and Editing Factual Associations in GPT. *arXiv:2202.05262*.
- Meng, K.; Sharma, A. S.; Andonian, A.; Belinkov, Y.; and Bau, D. 2023b. Mass-Editing Memory in a Transformer. *arXiv:2210.07229*.
- Minaee, S.; Mikolov, T.; Nikzad, N.; Chenaghlu, M.; Socher, R.; Amatriain, X.; and Gao, J. 2024. Large language models: A survey. *arXiv preprint arXiv:2402.06196*.
- Mitchell, E.; Lin, C.; Bosselut, A.; Finn, C.; and Manning, C. D. 2021. Fast model editing at scale. *arXiv preprint arXiv:2110.11309*.
- Mitchell, E.; Lin, C.; Bosselut, A.; Manning, C. D.; and Finn, C. 2022. Memory-based model editing at scale. In *International Conference on Machine Learning*, 15817–15831. PMLR.
- Petroni, F.; Rocktäschel, T.; Lewis, P.; Bakhtin, A.; Wu, Y.; Miller, A. H.; and Riedel, S. 2019. Language models as knowledge bases? *arXiv preprint arXiv:1909.01066*.
- Radford, A.; Wu, J.; Child, R.; Luan, D.; Amodei, D.; and Sutskever, I. 2019. Language Models are Unsupervised Multitask Learners.
- Roberts, A.; Raffel, C.; and Shazeer, N. 2020. How much knowledge can you pack into the parameters of a language model? *arXiv preprint arXiv:2002.08910*.
- Thirunavukarasu, A. J.; Ting, D. S. J.; Elangovan, K.; Gutierrez, L.; Tan, T. F.; and Ting, D. S. W. 2023. Large language models in medicine. *Nature medicine*, 29(8): 1930–1940.
- Wang, P.; Li, Z.; Zhang, N.; Xu, Z.; Yao, Y.; Jiang, Y.; Xie, P.; Huang, F.; and Chen, H. 2024. WISE: Rethinking the Knowledge Memory for Lifelong Model Editing of Large Language Models. *arXiv:2405.14768*.
- Wang, S.; Li, X.; Sun, J.; and Xu, Z. 2021. Training networks in null space of feature covariance for continual learning. In *Proceedings of the IEEE/CVF conference on Computer Vision and Pattern Recognition*, 184–193.
- Wang, S.; Zhu, Y.; Liu, H.; Zheng, Z.; Chen, C.; and Li, J. 2023. Knowledge Editing for Large Language Models: A Survey. *arXiv:2310.16218*.

Yang, W.; Sun, F.; Ma, X.; Liu, X.; Yin, D.; and Cheng, X. 2024. The butterfly effect of model editing: Few edits can trigger large language models collapse. *arXiv preprint arXiv:2402.09656*.

Yao, Y.; Wang, P.; Tian, B.; Cheng, S.; Li, Z.; Deng, S.; Chen, H.; and Zhang, N. 2023. Editing Large Language Models: Problems, Methods, and Opportunities. *arXiv:2305.13172*.

Zhu, C.; Rawat, A. S.; Zaheer, M.; Bhojanapalli, S.; Li, D.; Yu, F.; and Kumar, S. 2020. Modifying memories in transformer models. *arXiv preprint arXiv:2012.00363*.