

# Spectral Characterization and Mitigation of Sequential Knowledge Editing Collapse

Chi Zhang<sup>1</sup>, Mengqi Zhang<sup>1,4\*</sup>, Xiaotian Ye<sup>2</sup>, Runxi Cheng<sup>3</sup>,  
Zisheng Zhou<sup>1</sup>, Ying Zhou<sup>1</sup>, Pengjie Ren<sup>1</sup>, Zhumin Chen<sup>1</sup>

<sup>1</sup>Shandong University

<sup>2</sup>School of Computer Science, Beijing University of Posts and Telecommunications

<sup>3</sup>Tsinghua Shenzhen International Graduate School, Tsinghua University

<sup>4</sup>Suzhou Research Institute of Shandong University

zhangchi2001@sdu.edu.cn, mengqi.zhang@sdu.edu.cn

## Abstract

Sequential knowledge editing in large language models often causes catastrophic collapse of the model’s general abilities, especially for parameter-modifying methods. Existing approaches mitigate this issue through heuristic constraints on parameter updates, yet the mechanisms underlying such degradation remain insufficiently understood. In this work, we present a spectral analysis of sequential knowledge editing and show that a model’s general abilities are closely associated with dominant singular directions of pretrained weight matrices. These directions are highly sensitive to perturbations and are progressively disrupted by repeated edits, closely tracking the collapse in both editing efficacy and general performance. Building on this insight, we propose REVIVE, a plug-and-play framework that stabilizes sequential editing by explicitly preserving the dominant singular subspace. REVIVE represents parameter updates in the spectral basis of the original weights and filters components that would interfere with the protected region. Extensive experiments across multiple models and benchmarks show that REVIVE consistently improves editing efficacy while substantially preserving general abilities under long-horizon sequential editing, including extreme settings with up to 20,000 edits.

## 1 Introduction

Large language models (LLMs) often generate outdated or incorrect information due to limitations in pretraining data or evolving real-world knowledge (Cao et al., 2021). Knowledge editing (Meng et al., 2022b; Zhang et al., 2026, 2025b,a) provides a lightweight way to correct specific facts without retraining models from scratch. In practice, however, edits are rarely isolated: models are updated repeatedly over time, leading to the setting of **sequential knowledge editing** (Fang et al.,

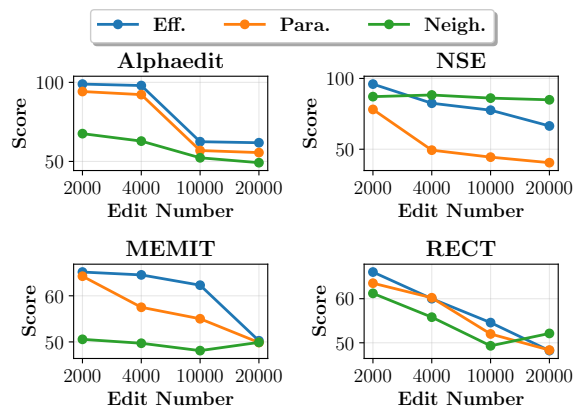


Figure 1: Performance of representative sequential knowledge editing methods on COUNTERFACT using LLaMA3.

2024), which requires both accurate edits and the preservation of general abilities (Gu et al., 2024). These challenges are particularly pronounced for *parameter-modifying* methods that directly update model weights, which are the focus of this work.

Despite the development of several sequential knowledge editing methods, including RECT (Gu et al., 2024), PRUNE (Ma et al., 2024), and ALPHAEDIT (Fang et al., 2024), long-horizon degradation remains largely unresolved. As illustrated in Figure 1, although these approaches appear to be effective in short or small-scale setting, they still suffer pronounced drops in both editing efficacy and general abilities as the number of edits increases. A key limitation is that most existing methods primarily mitigate interference across updates through heuristic constraints or regularization, such as bounding update magnitudes, reducing interference across edits, or minimizing disruption to unrelated knowledge, without explicitly modeling how parameter updates interact with the intrinsic structure of pretrained weights. In particular, they do not clarify which parts of the pretrained weights are most relevant to general abilities, or how repeated updates accumulate and gradually degrade

\*Corresponding author.

them. As a result, these methods provide limited mechanistic insight into why performance progressively degrades as edits accumulate.

In this work, we conduct a systematic analysis of sequential knowledge editing from a spectral perspective (§2). Through a detailed study of the singular value decomposition (SVD) of weight matrices, we obtain three key findings: (i) general abilities are concentrated in a low-rank spectral structure spanned by dominant singular directions (§2.2); (ii) this structure is unusually fragile, such that even small perturbations aligned with dominant directions can disproportionately degrade general performance (§2.3); and (iii) repeated edits progressively distort these dominant directions in practice (§2.4). Together, these findings provide strong empirical evidence that **interference with the dominant singular subspace is closely tied to failure in sequential editing, suggesting a primary mechanism underlying this collapse**.

Despite this insight, preserving dominant spectral structure during editing is nontrivial: editing objectives are local and offer little direct control over global spectral properties, and parameter-modifying editors differ substantially across methods. To address these challenges, we develop a **Robust sEquential editing Via domInant subspace preserVation framEwork** (REVIVE) (§3), a plug-and-play framework that stabilizes sequential knowledge editing by explicitly protecting dominant spectral structure. The core idea of REVIVE is to analyze parameter updates in the singular vector basis of the original weight matrix, which enables a fine-grained characterization of how edits interact with intrinsic functional directions. Based on this representation, REVIVE identifies the dominant subspace using a spectral energy criterion and constructs constrained updates by filtering out components that interfere with this critical region. This design allows factual knowledge to be incorporated through lower-energy directions while preserving the dominant spectral structure associated with general abilities.

Our contributions can be summarized as follows:

- We present a systematic spectral analysis of sequential knowledge editing, revealing that interference with the dominant singular subspace of pretrained weight matrices is a key mechanism underlying model collapse.
- Based on this analysis, we propose REVIVE, a simple and plug-and-play framework that stabi-

lizes sequential knowledge editing by explicitly preserving the dominant singular subspace during parameter updates.

- We conduct extensive experiments on multiple models and benchmarks, demonstrating that REVIVE consistently outperforms state-of-the-art methods in both editing efficacy and the preservation of general abilities under sequential editing.

## 2 Why Sequential Editing Collapses: A Spectral Perspective

To understand why sequential editing leads to model collapse, this section presents a spectral analysis of parameter matrices. We begin by viewing each weight matrix as a composition of independent input-output mappings derived from its Singular Value Decomposition (SVD). Based on this perspective, we empirically investigate three key questions: (i) *where a model’s general abilities are concentrated within these spectral components*, (ii) *how sensitive these components are to perturbations*, and (iii) *how sequential editing progressively distorts these components in practice*.

Together, these analyses lead to a unified insight: **sequential editing fails when cumulative update progressively corrupt the dominant singular directions of weight matrices, which are strongly associated with general abilities**. As mainstream editing methods primarily target feed-forward network (FFN) layers for modification (Meng et al., 2022b,a), we ground our investigation in the FFN matrices of LLaMA3-8B.

### 2.1 Spectral View of Weight Matrices

From a spectral perspective, a parameter matrix  $\mathbf{W} \in \mathbb{R}^{m \times n}$  can be decomposed into a set of independent input-output mappings using Singular Value Decomposition (SVD):

$$\mathbf{W} = \mathbf{U}\Sigma\mathbf{V}^\top = \sum_{i=1}^r \sigma_i \mathbf{u}_i \mathbf{v}_i^\top, \quad (1)$$

where  $\mathbf{U} \in \mathbb{R}^{m \times m}$  and  $\mathbf{V} \in \mathbb{R}^{n \times n}$  contain orthogonal left and right singular vectors, respectively, and  $\Sigma \in \mathbb{R}^{m \times n}$  contains the singular values  $\sigma_1 \geq \sigma_2 \geq \dots \geq \sigma_r \geq 0$ .

Each rank-one component  $\sigma_i \mathbf{u}_i \mathbf{v}_i^\top$  acts as a distinct input-output mapping: an input vector  $\mathbf{x} \in \mathbb{R}^n$  is projected onto  $\mathbf{v}_i$ , scaled by  $\sigma_i$ , and expanded along  $\mathbf{u}_i$  to produce the output. The orthogonality

of the singular vectors ensures these mappings operate independently. Pretraining learns this highly structured functional decomposition, making it a critical component of the model’s general abilities (Wang et al., 2024b; Hao et al., 2025). Crucially, this spectral view provides a unified lens for analyzing the failure of sequential editing. In the remainder of this section, we use this perspective to systematically investigate the three questions outlined above.

## 2.2 Where Are General Abilities Concentrated?

To identify where general abilities are concentrated, we evaluate model performance on GLUE tasks (MRPC, CoLA, RTE, NLI) (Wang et al., 2019) using weight matrices reconstructed from subsets of their singular components.

Specifically, we define the singular value energy of an index set  $\mathcal{I}$  as  $E_{\mathcal{I}} = \sum_{i \in \mathcal{I}} \sigma_i$ , and the total energy as  $E_{\text{total}} = \sum_{i=1}^r \sigma_i$ . We then reconstruct the weight matrix using the smallest set of singular components that captures  $n\%$  of the total energy,  $\tilde{\mathbf{W}}^n = \sum_{i \in \mathcal{I}} \sigma_i \mathbf{u}_i \mathbf{v}_i^\top$ , and evaluate model performance using the reconstructed matrices.

**Finding: General abilities are highly concentrated in the dominant singular subspace.** As shown in Figure 2, reconstructing weight matrices with only the top 5% of singular components (by energy) recovers approximately 62.6% of the model’s original performance. This result indicates that a substantial portion of a model’s general abilities is concentrated in a small, low-rank subspace spanned by the singular vectors associated with the largest singular values.

## 2.3 How Sensitive Are These Components to Perturbations?

To evaluate the robustness of different singular components, we partition the spectrum into ten non-overlapping groups by cumulative energy (0–10%, ..., 90–100%). For each group  $\mathcal{G}$ , we inject a structured perturbation that selectively acts on the corresponding input singular directions. Concretely, we first generate

$$\Delta = \sum_{j \in \mathcal{G}} \sum_{i=1}^r \alpha_{i,j} \mathbf{u}_i \mathbf{v}_j^\top, \alpha_{i,j} \sim \mathcal{N}(0, 1), \quad (2)$$

which randomly remaps the input directions  $\{\mathbf{v}_j\}_{j \in \mathcal{G}}$  to output directions. The perturbation is then normalized and scaled to a fixed strength:  $\varepsilon$ :




Figure 2: General ability recovery with increasing proportion of retained singular components.




Figure 3: Sensitivity of model’s general ability to perturbations across different spectral groups.

$\tilde{\Delta} = \varepsilon \cdot \frac{\Delta}{\|\Delta\|_F}$ . This ensures identical Frobenius norms across groups, thereby enabling a fair comparison. We evaluate the impact of the perturbed matrix  $\mathbf{W}' = \mathbf{W} + \tilde{\Delta}$  on general performance. A symmetric analysis on output directions is provided in Appendix F.1 and exhibits similar trends.

**Finding: The dominant singular subspace most strongly associated with general abilities is highly sensitive to perturbations.** As shown in Figure 3, perturbations applied to high-energy singular components (e.g., 0–20%), which span the dominant singular subspace, lead to sharp and consistent degradation in performance. In contrast, perturbations to low-energy components (70–100%) produce only minor or negligible effects. These results reveal a paradoxical property: **the spectral directions most associated with general abilities are also the most fragile under perturbations.**

## 2.4 How Does Sequential Editing Distort These Components in Practice?

Having established that general abilities are concentrated in a fragile dominant singular subspace, we now analyze how these critical subspaces degrade during sequential editing progress. We apply 2,000 edits from COUNTERFACT dataset to LLaMA3 using MEMIT, organized into 20 rounds with 100 edits per round. After each editing round, we evaluate (i) editing performance using Efficacy Score and Paraphrase Score, (ii) general abilities using the GLUE benchmark, and (iii) the spectral metrics introduced below. This allows us to analyze the relationship between behavioral degradation and structural changes in model parameters. We conduct the same analytical experiments on AlphaEdit; due to space constraints, these results are reported in Appendix F.2.

**Spectrum-based metrics.** To characterize the stability of the dominant singular subspace (top 10% components by singular value energy), we in-

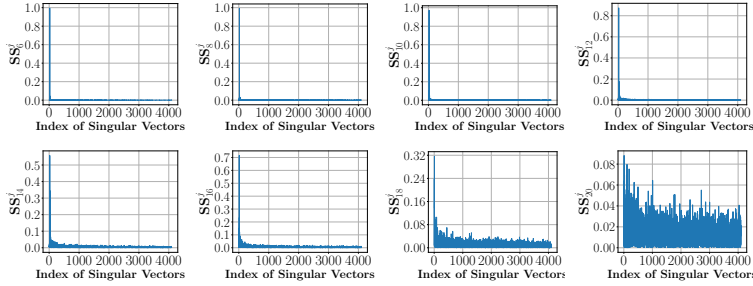


Figure 4: Singular vector similarity (SS) under sequential editing from rounds 6 to 20.




Figure 5:  $LS_t$  from rounds 1 to 20.




Figure 6: The Efficacy Score and Paraphrase Score vs. Edit Number.




Figure 7: General Ability Performance on GLUE vs. Edit Number.

Introduce two complementary metrics that capture degradation at different granularities.

- **Low-rank Subspace Similarity (LS)** measures the macroscopic drift of the dominant singular subspace. It is defined as the Frobenius cosine similarity between the low-rank reconstructions of the original matrix  $\hat{\mathbf{W}}_0$  and the edited matrix  $\hat{\mathbf{W}}_t$  at editing round  $t$ :

$$LS_t = \frac{\langle \hat{\mathbf{W}}_t, \hat{\mathbf{W}}_0 \rangle_F}{\|\hat{\mathbf{W}}_t\|_F \cdot \|\hat{\mathbf{W}}_0\|_F}. \quad (3)$$

- **Singular Vector Similarity (SS)** provides a microscopic view by quantifying how individual dominant singular vectors rotate during editing. We compute the cosine similarity between a dominant singular vector  $\mathbf{v}_t$  of the edited matrix  $\mathbf{W}_t$  and each original singular vector  $\mathbf{v}_j$  of  $\mathbf{W}_0$  (a vector basis):  $SS_t^j = \langle \mathbf{v}_t, \mathbf{v}_j \rangle$ . Results for left singular vectors show the same trend and are reported in Appendix F.3.

**Finding: Sequential editing progressively degrades the dominant singular subspace, accompanying model collapse.** This finding is supported by consistent evidence across behavioral, macroscopic, and microscopic analyses. At the *behavioral level*, both editing performance and general ability remain stable in early rounds but deteriorate rapidly after round 10, collapsing almost entirely

by round 20 (Figure 6 and 7). At the *macroscopic level*, this behavioral degradation is closely mirrored by a marked decline in the Low-rank Subspace Similarity (LS) metric (Figure 4). LS remains near its initial value in early rounds, indicating that the dominant singular subspace is largely preserved, before drifting and dropping sharply after roughly round 15. At the *microscopic level*, the Singular Vector Similarity (SS) metric (Figure 5) reveals that individual dominant singular directions gradually rotate away from their original directions, becoming nearly orthogonal by round 20. This gradual rotation reflects a systematic corruption of the learned input–output mappings represented by the dominant singular subspace. Taken together, these observations provide strong empirical evidence that collapse under sequential editing is structurally associated with the progressive degradation of the dominant singular subspace.

### 3 Methodology

Our analysis in Section 2 show that collapse under sequential editing is closely associated with the cumulative degradation of the dominant singular subspace, which is strongly linked to a model’s general performance. This motivates explicitly preserving this subspace to stabilize long-horizon sequential editing.

To this end, we introduce a **Robust sEquential editing Via dominant subspace preserVation framEwork** (REVIVE), a plug-and-play framework that operates directly on parameter updates produced by existing editing methods. Rather than modifying editing objectives or imposing method-specific constraints, REVIVE analyzes each update in the singular-vector basis of the original weight matrix, providing a unified spectral coordinate system for characterizing how edits interact with intrinsic functional directions. Within this representation,

REVIVE identifies update components that interfere with dominant singular directions and removes them before applying the update. An overview of the framework is shown in Figure 8, and the complete algorithm is provided in Appendix A.

### 3.1 Spectral Decomposition of Parameter Updates

To explicitly characterize how a parameter update interacts with different spectral components of the original weight matrix, REVIVE analyzes parameter updates through spectral decomposition.

Recall from Section 2.1 that the SVD of the original weight matrix  $\mathbf{W}$  defines sets of left and right singular vectors  $\{\mathbf{u}_i\}$  and  $\{\mathbf{v}_j\}$ . The rank-one outer products  $\{\mathbf{u}_i \mathbf{v}_j^\top\}_{i,j}$  constitute an orthogonal basis for the matrix space  $\mathbb{R}^{m \times n}$ ; a formal proof is provided in Appendix B. Using this basis, any update matrix  $\Delta \mathbf{W}$  produced by an editing method can be expressed as

$$\Delta \mathbf{W} = \sum_{i=1}^m \sum_{j=1}^n \alpha_{ij} \mathbf{u}_i \mathbf{v}_j^\top, \quad (4)$$

where the coefficients are given by  $\alpha_{ij} = \langle \Delta \mathbf{W}, \mathbf{u}_i \mathbf{v}_j^\top \rangle_F$ . This spectral decomposition resolves a parameter update into its effects on individual input–output mappings of the original weight matrix. Each coefficient  $\alpha_{ij}$  quantifies the contribution of the update along the mapping from input singular direction  $\mathbf{v}_j$  to output singular direction  $\mathbf{u}_i$ . As a result, update components aligned with dominant spectral directions can be distinguished from those operating in lower-energy regions.

### 3.2 Dominant Subspace Protection

Building on the spectral decomposition above, REVIVE introduces a Dominant Subspace Protection (DSP) mechanism that constrains parameter updates to avoid interference with the dominant singular subspace. DSP consists of two components: Dominant Subspace Identification and Safe Update Construction.

**Dominant Subspace Identification.** To identify the critical components for preservation, we adopt an energy-based criterion. Given a **singular-value energy threshold**  $\tau \in (0, 1)$  (analyzed in Section 4.2), we select the smallest index  $k$  such that the cumulative energy of the top- $k$  singular values exceeds  $\tau$ :

$$\frac{\sum_{i=1}^k \sigma_i}{\sum_{i=1}^r \sigma_i} \geq \tau. \quad (5)$$

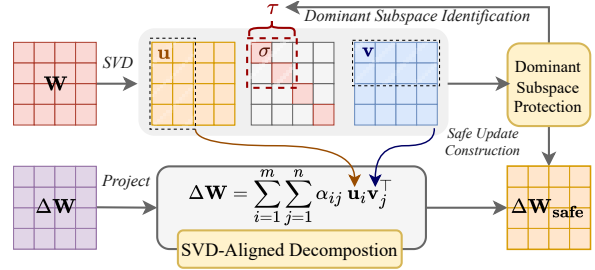


Figure 8: Overview of the REVIVE. An update matrix  $\Delta \mathbf{W}$  is represented in the singular vector basis of the original weight matrix. The dominant subspace, identified by an energy threshold  $\tau$ , is then used to filter out update components that would interfere with this protected region. The resulting safe update  $\Delta \mathbf{W}_{\text{safe}}$  preserves the dominant subspace while allowing effective knowledge modification.

The corresponding singular vectors  $\{\mathbf{u}_i\}_{i=1}^k$  and  $\{\mathbf{v}_i\}_{i=1}^k$  span the dominant input and output subspaces, respectively.

**Safe Update Construction.** Once the dominant subspace is identified, we construct a safe update by removing update components associated with this protected region. Concretely, the coefficient  $\alpha_{ij}$  is set to zero whenever the corresponding input direction  $\mathbf{v}_j$  or output direction  $\mathbf{u}_i$  lies in the dominant subspace, i.e.,  $\alpha_{ij} = 0$  if  $i \leq k$  or  $j \leq k$ . The resulting safe update is given by

$$\Delta \mathbf{W}_{\text{safe}} = \sum_{i>k} \sum_{j>k} \alpha_{ij} \mathbf{u}_i \mathbf{v}_j^\top. \quad (6)$$

By construction,  $\Delta \mathbf{W}_{\text{safe}}$  operates exclusively in lower-energy spectral directions, thereby limiting interference with dominant components that are critical for preserving general abilities.

Importantly, the protection mechanism of REVIVE operates as a plug-and-play wrapper that can be applied to existing parameter-modifying knowledge editing methods, without altering how edits are generated.

## 4 Experiments

### 4.1 Experiment Setup

**Base Models.** We conduct experiments on three LLMs that are widely used in prior knowledge editing work: GPT2-XL (1.5B) (Radford et al., 2019), GPT-J (6B) (Wang and Komatsuzaki, 2021), and LLaMA3 (8B) (Grattafiori et al., 2024). Due to space constraints, results on GPT2-XL are reported in Appendix G.1.

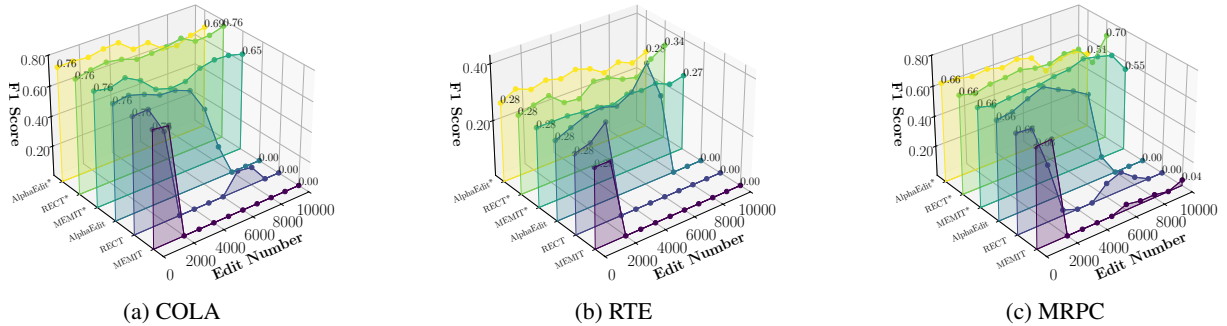


Figure 9: Performance of baselines and their REVIVE-enhanced versions (\*) on GLUE.

**Baselines.** We compare REVIVE against a suite of strong baselines, including the canonical MEMIT (Meng et al., 2022b), as well as four state-of-the-art methods designed for sequential editing: PRUNE (Ma et al., 2024), RECT (Gu et al., 2024), ALPHAEDIT (Fang et al., 2024), DELTAEDIT (Cao et al., 2025), and NSE (Jiang et al., 2025b). Due to differences in experimental settings, comparison with DeltaEdit can be found in Appendix G.2. Additional implementation details and baseline comparisons are provided in Appendix C.1.

**Datasets and Metrics.** We evaluate all methods on two standard factual knowledge editing benchmarks, COUNTERFACT (Meng et al., 2022b) and ZSRE (Levy et al., 2017). For ZSRE, we report Efficacy, Paraphrase, and Neighborhood Scores. For COUNTERFACT, we additionally report Fluency and Consistency metrics. Detailed definitions of all datasets and evaluation metrics are provided in Appendix C.2, C.3, and C.4.

## 4.2 Results and Analysis

This section presents a comprehensive evaluation of REVIVE. We first assess its performance under extended sequential editing, measuring both editing success and the preservation of general abilities. We then analyze robustness properties, including scalability to 20,000 edits, sensitivity to hyperparameters, the stability of internal representations using t-SNE visualizations, and matrix norm growth. Additionally, Appendix G.6 provides a direct verification that REVIVE preserves the dominant singular directions during long editing sequences. Further efficiency analyses are reported in Appendix G.8.

**Sequential Editing Performance.** We evaluate the effectiveness of REVIVE over an extended sequence of 10,000 edits, applied in 100 rounds of 100 edits each, on the COUNTERFACT and ZSRE

benchmarks. As shown in Table 1, incorporating REVIVE as a plug-and-play module consistently improves performance across all base methods and models. The improvements are particularly pronounced on the more challenging ZSRE benchmark, where several baselines, including MEMIT and RECT, degrade rapidly when applied sequentially without protection. When augmented with REVIVE, these methods remain stable and achieve substantially higher efficacy (e.g., 83.45% for MEMIT+REVIVE on LLaMA3). Notably, the standard MEMIT+REVIVE combination consistently outperforms specialized sequential baselines such as PRUNE and RECT, suggesting that proactively constraining harmful update directions is more effective than post-hoc regularization. For certain methods such as NSE, applying REVIVE leads to a moderate decrease in Neighborhood Success. This behavior is likely attributable to the fact that high neighborhood scores can arise trivially when edits fail to modify the model. In contrast, REVIVE achieves large gains in Efficacy and Paraphrase while maintaining strong Neighborhood performance, reflecting a more meaningful form of successful editing.

**Preservation of General Abilities.** To evaluate the preservation of general abilities, we assess edited LLaMA3 models on the GLUE benchmark after 10,000 sequential edits. For clarity, Figure 9 reports results on three representative tasks, with full results provided in Appendix G.3. All tasks exhibit consistent trends. As shown in Figure 9, baseline methods without protection experience rapid degradation in general performance. MEMIT and RECT collapse to near-zero performance after only 3,000 edits, and even the more robust ALPHAEDIT eventually suffers a complete collapse after 8,000 edits. In contrast, REVIVE enhanced methods maintains an overall average 86.34% of

Table 1: Performance on sequential editing over 10,000 Samples. The abbreviations *Eff.* (Efficacy Success), *Para.* (Paraphrase Success), *Neigh.* (Neighborhood Success), *Flu.* (Generation Entropy), and *Consis.* (Reference Score) denote respective evaluation metrics. Relative improvements (%) are shown in blue and decreases in orange.  $\uparrow\uparrow$  indicates a large improvement where the baseline score was near zero.

Method	Counterfact					ZsRE		
	Eff. $\uparrow$	Para. $\uparrow$	Neigh. $\uparrow$	Flu. $\uparrow$	Consis. $\uparrow$	Eff. $\uparrow$	Para. $\uparrow$	Neigh. $\uparrow$
<b>LLaMA3</b>	7.02	9.44	89.73	635.47	24.24	35.67	34.81	31.83
MEMIT	62.3	55.02	48.11	522.1	4.4	0.08	0.08	1.36
+REVIVE	95.62 $\uparrow$ 53.5%	84.60 $\uparrow$ 53.8%	62.17 $\uparrow$ 29.2%	603.22 $\uparrow$ 15.5%	29.39 $\uparrow$ 568.0%	83.45 $\uparrow\uparrow$	79.90 $\uparrow\uparrow$	32.01 $\uparrow\uparrow$
PRUNE	59.98	55.72	48.56	571.27	1.89	0.00	0.00	0.08
+REVIVE	80.57 $\uparrow$ 34.3%	69.54 $\uparrow$ 24.7%	54.76 $\uparrow$ 12.8%	570.85 $\downarrow$ 0.1%	28.49 $\uparrow\uparrow$	56.61 $\uparrow\uparrow$	53.30 $\uparrow\uparrow$	27.74 $\uparrow\uparrow$
RECT	60.23	54.9	50.56	441.61	5.08	0.00	0.00	0.00
+REVIVE	92.69 $\uparrow$ 53.9%	79.95 $\uparrow$ 45.6%	63.09 $\uparrow$ 24.8%	600.13 $\uparrow$ 35.9%	29.28 $\uparrow$ 476.8%	84.20 $\uparrow\uparrow$	80.27 $\uparrow\uparrow$	31.92 $\uparrow\uparrow$
AlphaEdit	62.48	56.9	52.31	505.5	4.25	90.57	85.66	30.5
+REVIVE	98.74 $\uparrow$ 58.0%	90.08 $\uparrow$ 58.4%	60.19 $\uparrow$ 15.1%	615.97 $\uparrow$ 21.9%	32.66 $\uparrow$ 668.5%	93.40 $\uparrow$ 3.1%	89.31 $\uparrow$ 4.3%	31.72 $\uparrow$ 4.0%
NSE	77.59	44.42	86.12	607.86	23.31	45.61	45.04	31.27
+REVIVE	98.89 $\uparrow$ 27.4%	92.28 $\uparrow$ 107.8%	65.72 $\downarrow$ 23.6%	618.66 $\uparrow$ 1.8%	32.74 $\uparrow$ 40.5%	94.37 $\uparrow$ 107.0%	90.57 $\uparrow$ 101.2%	32.17 $\uparrow$ 2.9%
<b>GPT-J</b>	15.22	17.65	83.50	622.01	29.61	26.45	25.74	27.04
MEMIT	54.03	52.66	53.63	594.16	5.17	0.10	0.10	0.17
+REVIVE	97.63 $\uparrow$ 80.7%	87.76 $\uparrow$ 66.6%	66.52 $\uparrow$ 24.1%	616.47 $\uparrow$ 3.8%	40.69 $\uparrow$ 687.4%	88.88 $\uparrow\uparrow$	83.22 $\uparrow\uparrow$	27.87 $\uparrow\uparrow$
PRUNE	52.92	51.47	53.91	576.95	5.14	0.03	0.02	0.05
+REVIVE	86.95 $\uparrow$ 64.3%	81.03 $\uparrow$ 57.5%	64.21 $\uparrow$ 19.1%	583.05 $\uparrow$ 1.1%	35.73 $\uparrow$ 595.7%	63.08 $\uparrow\uparrow$	58.90 $\uparrow\uparrow$	26.03 $\uparrow\uparrow$
RECT	63.60	55.33	56.69	404.13	4.49	23.60	22.02	12.44
+REVIVE	94.96 $\uparrow$ 49.4%	77.27 $\uparrow$ 39.6%	67.78 $\uparrow$ 19.6%	612.76 $\uparrow$ 51.62%	38.69 $\uparrow$ 761.7%	81.28 $\uparrow$ 244.4%	74.78 $\uparrow$ 239.6%	28.20 $\uparrow$ 126.7%
AlphaEdit	96.51	86.76	60.80	544.18	19.33	87.84	78.65	22.31
+REVIVE	99.50 $\uparrow$ 3.1%	93.92 $\uparrow$ 8.3%	67.35 $\uparrow$ 10.8%	600.64 $\uparrow$ 10.4%	40.63 $\uparrow$ 110.3%	97.53 $\uparrow$ 11.0%	91.33 $\uparrow$ 16.1%	23.40 $\uparrow$ 4.9%
NSE	88.95	69.69	75.46	611.35	33.31	44.03	42.39	24.86
+REVIVE	94.88 $\uparrow$ 6.7%	89.49 $\uparrow$ 28.4%	64.06 $\downarrow$ 15.1%	608.12 $\downarrow$ 0.5%	40.18 $\uparrow$ 20.6%	97.80 $\uparrow$ 122.1%	91.75 $\uparrow$ 116.4%	26.84 $\uparrow$ 8.0%

its performance across all tasks after 10,000 edits. These results clearly demonstrate that shielding the dominant singular subspace is a highly effective strategy for preserving a model’s general abilities during sequential editing.

### Scalability under Extreme Sequential Editing.

We further evaluate the scalability of REVIVE under substantially extended sequential editing on LLaMA3. Specifically, we apply 20,000 sequential edits on COUNTERFACT, organized into 200 rounds of 100 edits each, as well as the full set of 19,086 edits on ZSRE. As shown in Figure 10, REVIVE continues to deliver substantial gains over the original base methods. On COUNTERFACT, REVIVE achieves average improvements of +75.1% in Efficacy and +53.1% in Fluency. Complete results for all metrics and methods, provided in Appendix G.4. These results indicate that REVIVE effectively maintains editing performance even when the number of edits is significantly scaled up.

**Sensitivity Analysis.** We evaluate the robustness of REVIVE with respect to its only intrinsic hyperparameter, the singular value energy threshold  $\tau$ , defined in Section 3.2. This threshold controls the size of the dominant singular subspace that is protected from modification: larger values of  $\tau$  emphasize stronger preservation of general abilities, while smaller values permit more flexible updates by exposing a larger portion of the parameter space to editing. Figure 11 shows that REVIVE exhibits strong robustness, maintaining high performance across a wide range of  $\tau$  values. This behavior indicates that the effectiveness of REVIVE is not sensitive to the precise choice of the dominant subspace boundary, and that stable performance can be achieved without careful hyperparameter tuning. Further hyperparameter results for all models are reported in Appendix G.5.

**Visualization of Representational Stability.** To qualitatively examine how extended sequential editing affects internal representations, we apply




Figure 10: REVIVE vs. original methods under 20,000 sequential edits on COUNTERFACT.

t-SNE (Maaten and Hinton, 2008) to visualize the representations of 1,000 factual prompts from LLaMA3, both before and after applying 20,000 edits. As illustrated in Figure 12, a strong baseline like ALPHAEDIT causes a noticeable distributional shift, where post-edit representations drift away from their original positions. In contrast, when combined with REVIVE, MEMIT exhibits considerably smaller shifts, with post-edit representations remaining closely aligned with their original clusters. This visualization is consistent with our quantitative findings and suggests that preserving the dominant singular subspace helps maintain the overall representational geometry of the model under extended sequential editing.




Figure 11: Performance of MEMIT-REVIVE on LLaMA3 (CounterFact) under different thresholds. Figure 12: t-SNE visualization of LLaMA3 representations before and after 20,000 sequential edits.

**Matrix Norm Growth Analysis** Prior studies (Gupta et al., 2025; Xu et al., 2025) have consistently observed that parameter matrix norms tend to exhibit abnormal growth during sequential editing, a phenomenon often linked to model instability and catastrophic collapse. To verify whether preserving the dominant singular subspace inherently mitigates this issue, we examine the evolution of weight magnitudes during long-horizon editing.

Specifically, we track the  $L_1$  and  $L_2$  norms of the parameter matrix at Layer 3 of the GPT-J (6B) model over a sequence of 10,000 edits on the COUNTERFACT, utilizing MEMIT as the base edi-

tor. The results are summarized in Table 2. As edits accumulate, the unprotected baseline exhibits severe norm inflation, with the  $L_2$  norm surging from 105.51 to 20946.66. In contrast, REVIVE substantially constrains and stabilizes this growth. This empirical evidence confirms that by addressing the interference in the dominant subspace, REVIVE naturally suppresses the abnormal weight inflation associated with long-horizon degradation.

Table 2: Evolution of parameter matrix norms in GPT-J (Layer 3) over 10,000 sequential edits.

Edits	$L_1$ Norm		$L_2$ Norm	
	MEMIT	+REVIVE	MEMIT	+REVIVE
0	$6.88 \times 10^5$	$6.88 \times 10^5$	105.51	105.51
2000	$1.37 \times 10^6$	$7.47 \times 10^5$	5590.36	153.57
4000	$1.75 \times 10^6$	$7.71 \times 10^5$	8649.95	156.45
6000	$2.29 \times 10^6$	$7.96 \times 10^5$	13012.32	159.32
8000	$2.51 \times 10^6$	$8.13 \times 10^5$	13871.72	161.41
10000	$3.92 \times 10^6$	$8.30 \times 10^5$	20946.66	163.47

## 5 Related Work

Our work focuses on *parameter-modifying methods* for knowledge editing. Early methods such as MEND and MEMIT (Mitchell et al., 2022a; Meng et al., 2022b) demonstrate strong performance for isolated edits, but often degrade under sequential editing due to the accumulation of interfering parameter updates.

To mitigate interference across edits, several sequential editing methods introduce mechanisms that restrict or regularize parameter updates during editing (Fang et al., 2024; Ma et al., 2024; Gu et al., 2024; Cao et al., 2025). A common strategy among these approaches is to construct *protection subspaces* that constrain how updates are applied. Notably, **these subspaces are typically derived from empirical statistics over external data, such as covariance estimates computed from sampled factual triples (Fang et al., 2024), or from historical editing directions accumulated during**

**previous updates (Cao et al., 2025).** While effective in reducing empirically observed conflicts, such approaches are largely driven by external data distributions or past editing trajectories, and are not explicitly grounded in the intrinsic organization of the model’s parameter space.

By contrast, our approach directly analyzes the spectral structure of the model’s weight matrices, identifies degradation of dominant functional subspaces as the primary failure mechanism under sequential editing, and introduces a plug-and-play intervention to preserve these subspaces during updates. A more comprehensive discussion of related work, including detailed comparisons with SVD-based editing methods, is provided in Appendix D and Appendix E.

## 6 Conclusion

In this work, we investigated the mechanisms behind model collapse in sequential knowledge editing from a spectral perspective. Our analysis reveals that repeated parameter updates progressively distort dominant singular directions of pretrained weight matrices, which are crucial for maintaining a model’s general abilities. Based on this insight, we propose REVIVE, a plug-and-play framework that protects this intrinsic structure by filtering update components that interfere with dominant singular directions. Extensive experiments across multiple models and benchmarks show that REVIVE substantially improves editing performance while preserving general abilities, even under large numbers of sequential edits. Overall, our work provides a structural explanation for sequential editing failure and offer a principled approach for maintaining the long-term stability of edited LLMs.

## Acknowledgements

This work was supported by the National Natural Science Foundation of China (Nos. 62502286, 62472261, 62576339, 62372275). It was also Funded by Basic Research Program of Jiangsu (BK20250429), the Shandong Provincial Natural Science Foundation (ZR2024QF203), the Technology Innovation Guidance Program of Shandong Province (YDZX2024088), and the Provincial Key R&D Program of Shandong Province (Nos. 2024CXGC010108, 2025CXGC010112).

## Limitations

While our results demonstrate the effectiveness of preserving the dominant singular subspace for stabilizing sequential knowledge editing, several limitations remain and suggest directions for future work.

First, the dominant subspace in REVIVE is identified using a singular-value energy criterion. While this choice is empirically stable and effective across models and tasks, it is not theoretically optimal in a strict sense. Exploring alternative criteria for defining task-relevant or functionally critical spectral components remains an open direction.

Second, our analysis and experiments primarily focus on the feed-forward network (FFN) layers, which are known to play a central role in factual knowledge storage. However, other architectural components, such as attention mechanisms, may also contribute to knowledge representation and could exhibit different spectral behaviors under sequential editing. Extending our spectral analysis and protection strategy to these components remains an important avenue for future work.

Third, while recent inspiring works (Liu et al., 2025, 2026) have questioned the adequacy of existing evaluation protocols, this study primarily focuses on mitigating model collapse. Consequently, we did not benchmark our method against these emerging metrics at this stage. We are committed to investigating these evaluation-related issues in our subsequent research.

## Usage of large language models

In the preparation of this work, large language models (LLMs) have been utilized to assist in several stages of writing. In particular, LLMs played a significant role in polishing the manuscript by improving readability and correcting grammatical issues. Moreover, they provided valuable assistance in certain aspects of data visualization, such as generating and refining plotting scripts, which streamlined the experimental analysis process.

## Ethics Statement

This work investigates knowledge editing techniques for large language models to enable targeted updates of factual behaviors. All experiments are conducted on publicly available datasets and do not involve the collection of personal data or interaction with human subjects. We emphasize that knowledge editing methods should be applied

with caution, as they could potentially be misused to inject harmful, misleading, or inappropriate information into model parameters. Our approach is intended solely for controlled research settings, such as correcting erroneous knowledge and analyzing model behavior, and should not be deployed without appropriate safeguards and oversight. We acknowledge that edited models may still inherit biases present in the original training data.

## References

- Luisa Bentivogli, Bernardo Magnini, Ido Dagan, Hoa Trang Dang, and Danilo Giampiccolo. 2009. The fifth PASCAL recognizing textual entailment challenge. In *TAC. NIST*.
- Ding Cao, Yuchen Cai, Rongxi Guo, Xuesong He, and Guiquan Liu. 2025. Deltaedit: Enhancing sequential editing in large language models by controlling superimposed noise. *arXiv preprint arXiv:2505.07899*.
- 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.
- William B. Dolan and Chris Brockett. 2005. Automatically constructing a corpus of sentential paraphrases. In *IWP@IJCNLP. Asian Federation of Natural Language Processing*.
- Qingxiu Dong, Damai Dai, Yifan Song, Jingjing Xu, Zhifang Sui, and Lei Li. 2022. Calibrating factual knowledge in pretrained language models. In *Findings of the Association for Computational Linguistics: EMNLP 2022*, pages 5937–5947.
- Yanai Elazar, Nora Kassner, Shauli Ravfogel, Abhishava Ravichander, Eduard Hovy, Hinrich Schütze, and Yoav Goldberg. 2021. Measuring and improving consistency in pretrained language models. *Transactions of the Association for Computational Linguistics*, 9:1012–1031.
- Junfeng Fang, Houcheng Jiang, Kun Wang, Yunshan Ma, Jie Shi, Xiang Wang, Xiangnan He, and Tat-Seng Chua. 2024. Alphaedit: Null-space constrained knowledge editing for language models. In *The Thirteenth International Conference on Learning Representations*.
- Mor Geva, Roei Schuster, Jonathan Berant, and Omer Levy. 2021. Transformer feed-forward layers are key-value memories. In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 5484–5495.
- Aaron Grattafiori, Abhimanyu Dubey, Abhinav Jauhri, Abhinav Pandey, Abhishek Kadian, Ahmad Al-Dahle, Aiesha Letman, Akhil Mathur, Alan Schelten, Alex Vaughan, and 1 others. 2024. The llama 3 herd of models. *arXiv preprint arXiv:2407.21783*.
- Jia-Chen Gu, Hao-Xiang Xu, Jun-Yu Ma, Pan Lu, Zhen-Hua Ling, Kai-Wei Chang, and Nanyun Peng. 2024. Model editing harms general abilities of large language models: Regularization to the rescue. In *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing, EMNLP 2024*, pages 16801–16819.
- Akshat Gupta, Phudish Prateepamornkul, Maochuan Lu, Ahmed M. Alaa, Thomas Hartvigsen, and Gopala Anumanchipalli. 2025. Lifelong knowledge editing requires better regularization. In *Findings of the Association for Computational Linguistics: EMNLP 2025, Suzhou, China, November 4-9, 2025*, pages 22653–22675. Association for Computational Linguistics.
- Jitai Hao, Qiang Huang, Hao Liu, Xinyan Xiao, Zhaochun Ren, and Jun Yu. 2025. A token is worth over 1,000 tokens: Efficient knowledge distillation through low-rank clone. In *The Thirty-ninth Annual Conference on Neural Information Processing Systems*.
- Tom Hartvigsen, Swami Sankaranarayanan, Hamid Palangi, Yoon Kim, and Marzyeh Ghassemi. 2023. Aging with grace: Lifelong model editing with discrete key-value adaptors. *Advances in Neural Information Processing Systems*, 36:47934–47959.
- Dan Hendrycks, Collin Burns, Steven Basart, Andy Zou, Mantas Mazeika, Dawn Song, and Jacob Steinhardt. 2021. Measuring massive multitask language understanding. In *The Ninth International Conference on Learning Representations*.
- Zeyu Huang, Yikang Shen, Xiaofeng Zhang, Jie Zhou, Wenge Rong, and Zhang Xiong. 2023. Transformer-patcher: One mistake worth one neuron. In *The Eleventh International Conference on Learning Representations*.
- Houcheng Jiang, Junfeng Fang, Ningyu Zhang, Mingyang Wan, Guojun Ma, Xiang Wang, Xiangnan He, and Tat-Seng Chua. 2025a. Anyedit: Edit any knowledge encoded in language models. In *Forty-second International Conference on Machine Learning*.
- Houcheng Jiang, Junfeng Fang, Tianyu Zhang, Baolong Bi, An Zhang, Ruipeng Wang, Tao Liang, and Xiang Wang. 2025b. Neuron-level sequential editing for large language models. In *Proceedings of the 63rd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers), 2025*, pages 16678–16702.
- Omer Levy, Minjoon Seo, Eunsol Choi, and Luke Zettlemoyer. 2017. Zero-shot relation extraction via reading comprehension. In *21st Conference on Computational Natural Language Learning, CoNLL 2017*, pages 333–342. Association for Computational Linguistics (ACL).

- Wei Liu, Haomei Xu, Bingqing Liu, Zhiying Deng, Haozhao Wang, Jun Wang, Ruixuan Li, Yee Whye Teh, and Wee Sun Lee. 2025. Is model editing built on sand? revealing its illusory success and fragile foundation. *arXiv preprint arXiv:2510.00625*.
- Wei Liu, Haomei Xu, Hongkai Liu, Zhiying Deng, Ruixuan Li, Heng Huang, Yee Whye Teh, and Wee Sun Lee. 2026. Are we evaluating the edit locality of llm model editing properly? *arXiv preprint arXiv:2601.17343*.
- Jun-Yu Ma, Hong Wang, Hao-Xiang Xu, Zhen-Hua Ling, and Jia-Chen Gu. 2024. Perturbation-restrained sequential model editing. In *The Thirteenth International Conference on Learning Representations*.
- Laurens van der Maaten and Geoffrey Hinton. 2008. Visualizing data using t-sne. *Journal of machine learning research*, 9(Nov):2579–2605.
- Aman Madaan, Niket Tandon, Peter Clark, and Yiming Yang. 2022. Memory-assisted prompt editing to improve GPT-3 after deployment. In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing, EMNLP 2022*, pages 2833–2861.
- Kevin Meng, David Bau, Alex Andonian, and Yonatan Belinkov. 2022a. Locating and editing factual associations in gpt. *Advances in neural information processing systems*, 35:17359–17372.
- Kevin Meng, Arnab Sen Sharma, Alex J Andonian, Yonatan Belinkov, and David Bau. 2022b. Mass-editing memory in a transformer. In *The Eleventh International Conference on Learning Representations*.
- Eric Mitchell, Charles Lin, Antoine Bosselut, Chelsea Finn, and Christopher D. Manning. 2022a. Fast model editing at scale. In *The Tenth 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 *International Conference on Machine Learning*, pages 15817–15831. PMLR.
- Alec Radford, Jeffrey Wu, Rewon Child, David Luan, Dario Amodei, Ilya Sutskever, and 1 others. 2019. Language models are unsupervised multitask learners. *OpenAI blog*, 1(8):9.
- Richard Socher, Alex Perelygin, Jean Wu, Jason Chuang, Christopher D. Manning, Andrew Y. Ng, and Christopher Potts. 2013. Recursive deep models for semantic compositionality over a sentiment treebank. In *Conference on Empirical Methods in Natural Language Processing.*, pages 1631–1642.
- Chenmien Tan, Ge Zhang, and Jie Fu. 2024. Massive editing for large language models via meta learning. In *The Twelfth International Conference on Learning Representations*.
- Alex Wang, Amanpreet Singh, Julian Michael, Felix Hill, Omer Levy, and Samuel R. Bowman. 2019. GLUE: A multi-task benchmark and analysis platform for natural language understanding. In *The Seventh International Conference on Learning Representations*.
- Ben Wang and Aran Komatsuzaki. 2021. Gpt-j-6b: A 6 billion parameter autoregressive language model.
- Peng Wang, Zexi Li, Ningyu Zhang, Ziwen Xu, Yunzhi Yao, Yong Jiang, Pengjun Xie, Fei Huang, and Hua-jun Chen. 2024a. Wise: Rethinking the knowledge memory for lifelong model editing of large language models. *Advances in Neural Information Processing Systems*, 37:53764–53797.
- Xin Wang, Yu Zheng, Zhongwei Wan, and Mi Zhang. 2024b. Svd-llm: Truncation-aware singular value decomposition for large language model compression. In *The Thirteenth International Conference on Learning Representations*.
- Alex Warstadt, Amanpreet Singh, and Samuel R. Bowman. 2019. Neural network acceptability judgments. *Trans. Assoc. Comput. Linguistics*, 7:625–641.
- Adina Williams, Nikita Nangia, and Samuel R. Bowman. 2018. A broad-coverage challenge corpus for sentence understanding through inference. In *NAACL-HLT*, pages 1112–1122. Association for Computational Linguistics.
- Hao-Xiang Xu, Jun-Yu Ma, Zhen-Hua Ling, Ningyu Zhang, and Jia-Chen Gu. 2025. Constraining sequential model editing with editing anchor compression. In *Findings of the Association for Computational Linguistics: NAACL 2025*, pages 5499–5515.
- Lang Yu, Qin Chen, Jie Zhou, and Liang He. 2024. Melo: Enhancing model editing with neuron-indexed dynamic lora. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 38, pages 19449–19457.
- Mengqi Zhang, Bowen Fang, Qiang Liu, Xiaotian Ye, Shu Wu, Pengjie Ren, Zhumin Chen, and Liang Wang. 2025a. Kele: Residual knowledge erasure for enhanced multi-hop reasoning in knowledge editing. In *Findings of the Association for Computational Linguistics: EMNLP 2025*, pages 24537–24552.
- Mengqi Zhang, Xiaotian Ye, Qiang Liu, Pengjie Ren, Shu Wu, and Zhumin Chen. 2024a. Knowledge graph enhanced large language model editing. In *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing, EMNLP 2024*, pages 22647–22662.
- Mengqi Zhang, Xiaotian Ye, Qiang Liu, Shu Wu, Pengjie Ren, and Zhumin Chen. 2025b. Uncovering overfitting in large language model editing. In *The Thirteenth International Conference on Learning Representations*.

Mengqi Zhang, Zisheng Zhou, Xiaotian Ye, Qiang Liu, Zhaochun Ren, Zhumin Chen, and Pengjie Ren. 2026. [Disentangling knowledge representations for large language model editing](#). In *The Fourteenth International Conference on Learning Representations*.

Ningyu Zhang, Bozhong Tian, Siyuan Cheng, Xiaozhuan Liang, Yi Hu, Kouying Xue, Yanjie Gou, Xi Chen, and Huajun Chen. 2024b. [Instructedit: Instruction-based knowledge editing for large language models](#). In *Proceedings of the Thirty-Third International Joint Conference on Artificial Intelligence, 2024*, pages 6633–6641.

Ningyu Zhang, Zekun Xi, Yujie Luo, Peng Wang, Bozhong Tian, Yunzhi Yao, Jintian Zhang, Shumin Deng, Mengshu Sun, Lei Liang, and 1 others. 2024c. [Oneedit: A neural-symbolic collaboratively knowledge editing system](#). *Proceedings of the VLDB Endowment*. ISSN, 2150:8097.

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, EMNLP 2023*, pages 4862–4876.

## A Algorithm Details

---

### Algorithm 1 REVIVE

---

**Require:** Current weight matrix  $\mathbf{W} \in \mathbb{R}^{m \times n}$ ; update matrix  $\Delta \mathbf{W}$ ; singular-value energy threshold  $\tau \in (0, 1)$

**Ensure:** Safe update  $\Delta \mathbf{W}_{\text{safe}}$

1: **SVD-Aligned Decomposition:**

2:  $\{\mathbf{u}_i\}_{i=1}^m, \{\sigma_i\}_{i=1}^r, \{\mathbf{v}_i\}_{i=1}^n = \text{SVD}(\mathbf{W})$

3: Construct orthogonal basis  $\{\mathbf{u}_i \mathbf{v}_j^\top \mid i = 1, \dots, m; j = 1, \dots, n\}$

4: **for**  $i = 1$  to  $m$  **do**

5:     **for**  $j = 1$  to  $n$  **do**

6:          $\alpha_{ij} \leftarrow \langle \Delta \mathbf{W}, \mathbf{u}_i \mathbf{v}_j^\top \rangle_F$

7:     **end for**

8: **end for**

9: Represent update as  $\Delta \mathbf{W} = \sum_{i=1}^m \sum_{j=1}^n \alpha_{ij} \mathbf{u}_i \mathbf{v}_j^\top$

10: **Dominant Subspace Identification:**

11: Find smallest  $k$  s.t.  $\frac{\sum_{i=1}^k \sigma_i}{\sum_{i=1}^n \sigma_i} \geq \tau$

12: Define dominant subspace  $\{\mathbf{u}_1, \dots, \mathbf{u}_k\}, \{\mathbf{v}_1, \dots, \mathbf{v}_k\}$

13: **Safe Update Construction:**

14: Initialize  $\Delta \mathbf{W}_{\text{safe}} \leftarrow 0$

15: **for**  $i = 1$  to  $m$  **do**

16:     **for**  $j = 1$  to  $n$  **do**

17:         **if**  $i > k$  **and**  $j > k$  **then**

18:              $\Delta \mathbf{W}_{\text{safe}} \leftarrow \Delta \mathbf{W}_{\text{safe}} + \alpha_{ij} \mathbf{u}_i \mathbf{v}_j^\top$

19:         **end if**

20:     **end for**

21: **end for**

22: **return**  $\Delta \mathbf{W}_{\text{safe}}$

---

## B Proof of the SVD-aligned Matrix Basis

This section provides the formal proof for the claim made in Section 3.1 that the set of rank-one outer products  $\{\mathbf{u}_i \mathbf{v}_j^\top\}_{i,j}$  derived from the singular vectors of a matrix  $\mathbf{W}$ , forms an orthonormal basis for the space of matrices  $\mathbb{R}^{m \times n}$ .

**Theorem 1** (Outer-product bases from two orthonormal vector bases). *Let  $\{\mathbf{u}_1, \dots, \mathbf{u}_m\} \subset \mathbb{R}^m$  and  $\{\mathbf{v}_1, \dots, \mathbf{v}_n\} \subset \mathbb{R}^n$  be orthonormal bases of  $\mathbb{R}^m$  and  $\mathbb{R}^n$  respectively. Consider the set of  $mn$  matrices*

$$\mathcal{B} = \left\{ \mathbf{u}_p \mathbf{v}_q^\top : p = 1, \dots, m, q = 1, \dots, n \right\}. \quad (7)$$

*Then  $\mathcal{B}$  forms an orthonormal basis of the real vector space  $\mathbb{R}^{m \times n}$  with respect to the Frobenius inner product  $\langle \mathbf{X}, \mathbf{Y} \rangle_F = \text{tr}(\mathbf{X}^\top \mathbf{Y})$ .*

In particular, every  $\mathbf{X} \in \mathbb{R}^{m \times n}$  admits the unique expansion

$$\mathbf{X} = \sum_{p=1}^m \sum_{q=1}^n c_{pq} \mathbf{u}_p \mathbf{v}_q^\top, \quad (8)$$

where the coefficients are given by

$$c_{pq} = \langle \mathbf{X}, \mathbf{u}_p \mathbf{v}_q^\top \rangle_F = \mathbf{u}_p^\top \mathbf{X} \mathbf{v}_q. \quad (9)$$

*Proof.* We split the proof into three parts: (i) orthogonality, (ii) spanning (completeness), and (iii) uniqueness / coefficient formula.

**(i) Orthogonality.** Take two generic elements  $\mathbf{u}_p \mathbf{v}_q^\top$  and  $\mathbf{u}_{p'} \mathbf{v}_{q'}^\top$  from  $\mathcal{B}$ . Their Frobenius inner product is

$$\langle \mathbf{u}_p \mathbf{v}_q^\top, \mathbf{u}_{p'} \mathbf{v}_{q'}^\top \rangle_F = (\mathbf{u}_p^\top \mathbf{u}_{p'}) (\mathbf{v}_q^\top \mathbf{v}_{q'}). \quad (10)$$

Since  $\{\mathbf{u}_p\}$  and  $\{\mathbf{v}_q\}$  are orthonormal bases, we have

$$\mathbf{u}_p^\top \mathbf{u}_{p'} = \delta_{pp'}, \quad \mathbf{v}_q^\top \mathbf{v}_{q'} = \delta_{qq'}, \quad (11)$$

where  $\delta_{ij}$  is the Kronecker delta:

$$\delta_{ij} = \begin{cases} 1, & \text{if } i = j, \\ 0, & \text{if } i \neq j. \end{cases} \quad (12)$$

Therefore,

$$\langle \mathbf{u}_p \mathbf{v}_q^\top, \mathbf{u}_{p'} \mathbf{v}_{q'}^\top \rangle_F = \delta_{pp'} \delta_{qq'}. \quad (13)$$

In particular, if either  $p \neq p'$  or  $q \neq q'$ , then one of the Kronecker deltas vanishes, making the inner product equal to 0. This proves that distinct basis elements are orthogonal.

**(ii) Completeness.** The vector space  $\mathbb{R}^{m \times n}$  has dimension  $mn$ . We have produced  $mn$  elements in  $\mathcal{B}$ , which are mutually orthonormal. Mutual orthonormality implies linear independence. Since we have exactly  $mn$  linearly independent matrices in an  $mn$ -dimensional space,  $\mathcal{B}$  must span  $\mathbb{R}^{m \times n}$ , and therefore forms a basis.

For a constructive argument, let  $\mathbf{X} \in \mathbb{R}^{m \times n}$  be arbitrary. Define coefficients

$$c_{pq} = \langle \mathbf{X}, \mathbf{u}_p \mathbf{v}_q^\top \rangle_F = \mathbf{u}_p^\top \mathbf{X} \mathbf{v}_q. \quad (14)$$

Form the matrix

$$\widehat{\mathbf{X}} = \sum_{p=1}^m \sum_{q=1}^n c_{pq} \mathbf{u}_p \mathbf{v}_q^\top. \quad (15)$$

For any fixed indices  $(p', q')$ , we compute

$$\begin{aligned} \langle \widehat{\mathbf{X}}, \mathbf{u}_{p'} \mathbf{v}_{q'}^\top \rangle_F &= \sum_{p,q} c_{pq} \langle \mathbf{u}_p \mathbf{v}_q^\top, \mathbf{u}_{p'} \mathbf{v}_{q'}^\top \rangle_F \\ &= \sum_{p,q} c_{pq} \delta_{pp'} \delta_{qq'} \\ &= c_{p'q'}. \end{aligned} \quad (16)$$

But by definition,  $c_{p'q'} = \langle \mathbf{X}, \mathbf{u}_{p'} \mathbf{v}_{q'}^\top \rangle_F$ , hence

$$\langle \widehat{\mathbf{X}} - \mathbf{X}, \mathbf{u}_{p'} \mathbf{v}_{q'}^\top \rangle_F = 0 \quad \text{for all } p', q'. \quad (17)$$

Since  $\mathcal{B}$  spans the space, the only matrix orthogonal to every basis element is the zero matrix. Therefore,  $\widehat{\mathbf{X}} - \mathbf{X} = \mathbf{0}$ , proving  $\mathbf{X} = \widehat{\mathbf{X}}$ . This provides an explicit expansion of any matrix in the basis  $\mathcal{B}$ , proving completeness.

**(iii) Uniqueness.** Orthogonality immediately gives that the coefficients are unique and equal to the Frobenius inner products:

$$c_{pq} = \langle \mathbf{X}, \mathbf{u}_p \mathbf{v}_q^\top \rangle_F = \mathbf{u}_p^\top \mathbf{X} \mathbf{v}_q. \quad (18)$$

This completes the proof. Therefore, using the  $\mathbf{u}$  and  $\mathbf{v}$  matrices obtained from the SVD of a matrix to construct such outer-product basis matrices is valid and well-founded.  $\square$

## C Experimental Details

### C.1 Baselines

- **MEMIT** (Meng et al., 2022b) is the first method to support large-scale knowledge injection across multiple layers. It exploits the key–value structure (Geva et al., 2021) of FFNs and improves upon ROME by restricting updates to a contiguous set of layers, allowing thousands of new facts to be inserted in one pass. However, MEMIT does not consider sequential editing, leaving space for later improvements.
- **RECT** (Gu et al., 2024) RECT is designed to mitigate the degradation of general abilities during sequential editing. It observes that general ability performance declines as more edits are applied, and addresses this by updating parameters based on the magnitude of change in individual weights. However, as our earlier analysis suggests, general abilities are governed by mappings between directions rather than individual parameters. Consequently, RECT remains too **localized**

- at the parameter level and fails to effectively preserve general abilities in long-horizon sequential editing.
- **PRUNE** (Ma et al., 2024) is specifically designed for sequential editing with the goal of protecting the general abilities of LLMs. From the perspective of matrix conditioning, it constrains the maximum singular value of the update matrix so that it does not exceed that of the original parameter matrix, thereby reducing the risk of collapse. However, unlike our method, PRUNE does not filter the directions associated with large singular values, which may weaken knowledge retention. Moreover, its constraint only limits singular values to remain below a threshold, effectively attenuating but not eliminating the influence of noise.
  - **NSE** (Jiang et al., 2025b) is a method specifically designed for sequential knowledge editing. It preserves the original parameters during update computation, ensuring that each new edit does not interfere with previously injected knowledge. Inspired by the key–value view of FFN layers (Geva et al., 2021), NSE uses activation values to identify those neurons most relevant to the current update, restricting parameter changes within this subset. While this reduces unnecessary disturbance to the model, neuron-level selection alone cannot fully protect general abilities due to the problem of *superposition*, where a single neuron may encode multiple orthogonal directions. As a result, NSE still fails to maintain general abilities under long-horizon sequential editing.
  - **AlphaEdit** (Fang et al., 2024) is a method specifically designed for sequential knowledge editing. It constructs a protection subspace for previously stored knowledge by collecting 100K (subject, relation, object) triples from Wikipedia. During subsequent edits, parameter updates are projected onto the null space of this protection subspace to prevent interference with existing knowledge. However, based on our earlier analysis, sequential editing primarily perturbs the subspace associated with general abilities rather than factual knowledge alone. Thus, the choice of protection subspace in AlphaEdit is not sufficiently precise. As shown in our experiments, AlphaEdit can withstand more editing steps compared to other baselines, but eventually still suffers from a collapse of general abilities.
  - **DeltaEdit** (Cao et al., 2025) is a sequential knowledge editing method that addresses the superimposed noise accumulation problem in LLMs. It mitigates performance degradation by identifying and reducing conflicts between consecutive edits. Specifically, it employs a dynamic orthogonal constraint strategy to ensure that new knowledge updates remain orthogonal to previous ones, thereby minimizing interference between different editing instances and maintaining model stability over long-term sequences.
  - **MEND** (Mitchell et al., 2022a) is a hypernetwork-based knowledge editing method that addresses the scalability and efficiency challenges of updating massive LLMs. It enables fast, localized modifications by learning to transform standard fine-tuning gradients into targeted parameter updates. Specifically, it employs a low-rank gradient decomposition strategy to parameterize the editor network, thereby reducing the immense computational overhead and allowing rapid fact corrections without degrading the model’s general capabilities.

## C.2 Datasets

- **ZsRE** (Levy et al., 2017) is a question-answering (QA) dataset. Each sample contains a subject string and a corresponding answer, which serve as the editing target to assess **Efficacy**. To evaluate **Paraphrase**, it utilizes rephrased questions generated through back-translation. Following prior work, it employs out-of-scope Natural Questions to measure **Neighborhood** (also referred to as **Locality**).
- **Counterfact** (Meng et al., 2022b) is a more challenging dataset that contrasts Counterfactual with factual statements. Each record is derived from an entry in the PARAREL dataset (Elazar et al., 2021), with all entities originating from WikiData. It uses metrics similar to ZsRE to evaluate **Efficacy Score**, **Paraphrase Score**, and **Neighborhood Score**. For its out-of-scope data, it replaces the subject entity with an approximate entity that shares the same predicate. Furthermore, Counterfact uniquely includes multiple generation prompts with the same meaning to test the **Fluency**(Generation Entropy) and **Consistency**(Reference Score) of the generated text.

### C.3 ZsRE Metrics

In line with prior work (Meng et al., 2022a,b), we define the evaluation metrics on the ZSRE dataset. Given a language model  $f_\theta$ , a factual prompt  $(s_i, r_i)$ , its target output  $o_i$ , and the model’s pre-edit prediction  $o_i^c$ , the following metrics are used:

- **Efficacy:** This metric reflects the model’s accuracy on the edited samples, computed as the average top-1 success rate:

$$\mathbb{E}_i \left\{ o_i = \arg \max_o \mathbb{P}_{f_\theta}(o \mid (s_i, r_i)) \right\}. \quad (19)$$

- **Paraphrase:** This measures how well the model transfers the edit to paraphrased forms of  $(s_i, r_i)$ , denoted as  $N((s_i, r_i))$ . It is defined as the average top-1 accuracy over these rephrasings:

$$\mathbb{E}_i \left\{ o_i = \arg \max_o \mathbb{P}_{f_\theta}(o \mid N((s_i, r_i))) \right\}. \quad (20)$$

- **Neighborhood:** This evaluates whether unrelated prompts  $O(s_i, r_i)$  remain unaffected by the edit. It is measured as the proportion of cases where predictions for such prompts stay consistent:

$$\mathbb{E}_i \left\{ o_i^c = \arg \max_o \mathbb{P}_{f_\theta}(o \mid O((s_i, r_i))) \right\}. \quad (21)$$

### C.4 Counterfact Metrics

In line with prior work (Meng et al., 2022a,b), we further define the metrics used in the Counterfact benchmark. Given a language model  $f_\theta$ , a factual prompt  $(s_i, r_i)$ , a target output  $o_i$ , and the model’s pre-edit prediction  $o_i^c$ , we define:

- **Efficacy Score:** The fraction of cases where, for the prompt  $(s_i, r_i)$ , the target  $o_i$  receives higher probability than the original output  $o_i^c$ :

$$\mathbb{E}_i \left[ \mathbb{P}_{f_\theta}[o_i \mid (s_i, r_i)] > \mathbb{P}_{f_\theta}[o_i^c \mid (s_i, r_i)] \right]. \quad (22)$$

- **Paraphrase Score:** The proportion of paraphrased prompts  $N((s_i, r_i))$  where the edited output  $o_i$  is ranked higher than the original response  $o_i^c$ :

$$\mathbb{E}_i \left[ \mathbb{P}_{f_\theta}[o_i \mid N((s_i, r_i))] > \mathbb{P}_{f_\theta}[o_i^c \mid N((s_i, r_i))] \right]. \quad (23)$$

- **Neighborhood Score:** The proportion of semantically related but distinct prompts  $O((s_i, r_i))$  where the model maintains correct predictions, assigning higher probability to  $o_i$  over  $o_i^c$ :

$$\mathbb{E}_i \left[ \mathbb{P}_{f_\theta}[o_i \mid O((s_i, r_i))] > \mathbb{P}_{f_\theta}[o_i^c \mid O((s_i, r_i))] \right]. \quad (24)$$

- **Fluency:** A measure of output repetition, defined using the entropy of the  $n$ -gram distribution:

$$-\frac{2}{3} \sum_k g_2(k) \log_2 g_2(k) + \frac{4}{3} \sum_k g_3(k) \log_2 g_3(k), \quad (25)$$

where  $g_n(\cdot)$  denotes the frequency distribution over  $n$ -grams.

- **Consistency:** This evaluates how consistent the model’s generations are with external references. Given a subject  $s$ , we compute the cosine similarity between TF-IDF embeddings of the model’s text and the corresponding Wikipedia article about  $o$ .

### C.5 Details of GLUE

GLUE (Wang et al., 2019) is a comprehensive benchmark, and this paper leverages the following six subtasks:

- **SST (Socher et al., 2013)** (Stanford Sentiment Treebank): A single-sentence classification task based on movie reviews, where the goal is to predict binary sentiment labels.
- **MRPC (Dolan and Brockett, 2005)** (Microsoft Research Paraphrase Corpus): A sentence-pair task aiming to identify whether two sentences are semantically equivalent.
- **MMLU (Hendrycks et al., 2021)** (Massive Multi-task Language Understanding): A broad benchmark covering diverse subjects, designed to evaluate models under zero-shot and few-shot conditions.
- **RTE (Bentivogli et al., 2009)** (Recognizing Textual Entailment): A natural language inference task where the objective is to determine if a premise entails its corresponding hypothesis.
- **CoLA (Warstadt et al., 2019)** (Corpus of Linguistic Acceptability): A single-sentence classification benchmark that tests whether sentences are grammatically acceptable.

- **NLI** (Williams et al., 2018) (Natural Language Inference): A task requiring models to infer the logical relationship between a premise and a hypothesis.

### C.6 Method Implementation Details

All experiments based on GPT-J and LLaMA3 are conducted on NVIDIA A800 GPUs (80GB), while experiments on GPT2-XL are performed on NVIDIA RTX 4090 GPUs (24GB). For baselines, we adopt the official implementations of ALPHAEEDIT and NSE and keep their default hyperparameter configurations unchanged. The only additional hyperparameter introduced by our method is the singular value projection threshold. Following Appendix G.5, we set this threshold to preserve the top 10% dominant singular components for all models. We use the same threshold across baselines because the protected structure is model-dependent and determined solely by the spectrum of the underlying pretrained weights. For SVD computation, we use `torch.linalg.svd`. Before each update, we perform SVD on the current parameter matrix to construct the spectral basis used by our projection.

### D Related Work (Full Version)

Knowledge editing methods for large language models can be broadly categorized into *parameter-preserving* and *parameter-modifying* approaches, depending on whether the original model parameters are directly updated. These two lines of work differ fundamentally in their design philosophy, flexibility, and susceptibility to interference under sequential editing. Below we review representative methods in each category, with an emphasis on techniques relevant to sequential knowledge editing.

**Parameter-Preserving Methods.** Parameter-preserving approaches maintain the base model’s parameters unchanged and instead incorporate external mechanisms to realize edits. A common direction is to attach additional modules that store and retrieve edited knowledge. For example, SERAC (Mitchell et al., 2022b) introduces an auxiliary memory with a Counterfactual model, CaliNet (Dong et al., 2022) and T-Patcher (Huang et al., 2023) insert neuron-based units, and GRACE (Hartvigsen et al., 2023) organizes edits in a dynamic codebook. MELO (Yu et al., 2024) uses additional LoRA-style adapters to preserve original parameters, while WISE (Wang

et al., 2024a) improves stability and general ability with dual-memory and conflict-free sharding. Another line of work performs edits through prompting: MemPrompt (Madaan et al., 2022) and IKE (Zheng et al., 2023) rely on injecting new facts into the input context. More recent efforts combine symbolic structures with neural editing, such as OneEdit (Zhang et al., 2024c), which leverages knowledge graphs for collaborative knowledge updates.

**Parameter-Modifying Methods.** Parameter-modifying methods directly update the model’s weights to encode new knowledge. Meta-learning based techniques predict parameter shifts through hypernetworks, including MEND (Mitchell et al., 2022a), MALMEN (Tan et al., 2024), and InstructEdit (Zhang et al., 2024b). Locate-then-edit methods first determine the locations where knowledge is stored and then apply targeted modifications. Typical examples are ROME (Meng et al., 2022a), which computes updates using closed-form equations, MEMIT (Meng et al., 2022b), which scales editing to batches, GLAME (Zhang et al., 2024a), which integrates knowledge graphs, and AnyEdit (Jiang et al., 2025a), which recursively edits knowledge of arbitrary structure.

When parameter-modifying edits are applied sequentially, additional difficulties arise due to the accumulation of interfering updates. To address these issues, several methods introduce explicit mechanisms to regularize or constrain updates during editing. RECT (Gu et al., 2024) enforces sparsity at the level of individual parameters, PRUNE (Ma et al., 2024) controls the conditioning of update matrices, AlphaEdit (Fang et al., 2024) restricts edits to a null space derived from previously stored knowledge, and DeltaEdit (Cao et al., 2025) mitigates edit–edit interference by projecting new updates onto the orthogonal complement of historical update directions. NSE (Jiang et al., 2025b) selects editing locations based on neuron activation patterns that are deemed most relevant for knowledge storage. While these methods alleviate certain forms of interference, they do not explicitly account for how sequential updates interact with the intrinsic structure of pretrained model parameters, which is the focus of our work.

Table 3: Comparison between REVIVE and representative SVD-based or subspace-constrained editing methods.

Method	Protected Subspace	Failure Mode Addressed
PRUNE	Singular value magnitude constraint	Large-magnitude updates
AlphaEdit	Feature covariance from external factual data	Interference across factual activations
Delta-Edit	Historical update directions	Edit–edit interference
<b>REVIVE</b>	<b>Dominant singular vectors of weight matrices</b>	<b>Dominant subspace corruption</b>

## E Comparison with SVD-based Editing Methods

This section clarifies the conceptual and methodological differences between REVIVE and prior knowledge editing methods that employ SVD or subspace-based constraints, including PRUNE (Ma et al., 2024), AlphaEdit (Fang et al., 2024), and DeltaEdit (Cao et al., 2025). Although these approaches share a common mathematical tool (e.g., singular value decomposition or subspace projection), **they differ fundamentally in how protected subspaces are defined, which failure modes they are intended to mitigate, and what aspects of model degradation they address.**

**PRUNE.** PRUNE constrains parameter updates by limiting the singular values of the update matrix relative to those of the original weights. This design primarily regulates the *magnitude* of updates and provides a form of spectral regularization. However, it does not distinguish between different singular directions and therefore does not explicitly preserve functionally meaningful subspaces of the pretrained model. Consequently, PRUNE may suppress useful update components while failing to preserve the functional subspace that supports general model abilities.

**AlphaEdit.** AlphaEdit constructs a protected subspace based on the covariance of feature activations induced by a large external factual dataset (e.g., 100k triples from wikipedia). This approach aims to reduce interference between factual representations during sequential edits. As a result, the protected subspace is *extrinsic* to the model and reflects correlations arising from the sampled data distribution rather than the intrinsic structure of the pretrained weight matrices. While effective at mitigating interference between edits, this design does not explicitly safeguard the dominant spectral components of the model parameters. Empirically, AlphaEdit exhibits severe degradation after long edit sequences, with both editing efficacy and general abilities collapsing after approximately 8k edits.

**DeltaEdit.** DeltaEdit focuses on preventing *edit–edit interference* by tracking historical update directions and projecting new edits to be orthogonal to them. The protected subspace is thus derived entirely from past edits, without reference to the pretrained model’s internal parameter structure. Consequently, DeltaEdit addresses interference between edits but does not target model-level spectral drift, and cannot prevent degradation of general abilities that may arise even when consecutive edits are semantically independent.

**REVIVE (ours).** REVIVE is motivated by a failure mode that has received limited attention in prior sequential editing work: the progressive degradation of the *dominant singular subspace* of weight matrices over long edit sequences. Our spectral analyses suggest that high-energy singular directions are strongly associated with general abilities. REVIVE therefore protects an *intrinsic*, model-derived subspace defined by the dominant singular vectors of the pretrained weights. By selectively filtering update components that would interfere with this subspace, REVIVE mitigates spectral degradation and supports stable editing over long horizons (up to 20k sequential edits in our experiments).

**Summary.** As summarized in Table 3, existing methods employ subspace constraints to mitigate different forms of interference, they are primarily designed to control interactions between edits or between data instances. In contrast, REVIVE identifies degradation of the model’s intrinsic dominant singular subspace as a fundamental cause of failure in sequential editing and intervenes by selectively filtering update components that interfere with the dominant singular subspace.

## F Supplementary Spectral Analyses for Section 2

### F.1 Robustness to Output-Side Perturbations

To complement the input-side perturbation analysis presented in Section 2.3, we conduct a symmetric study on *output-side perturbations*, examining how




Figure 13: Sensitivity of general ability to perturbations across different output side (left singular vector) spectral groups.

perturbations aligned with different groups of *left* singular vectors affect model robustness.

Specifically, we partition the left singular vectors  $\{\mathbf{u}_i\}$  of the weight matrix  $\mathbf{W}$  into groups based on the cumulative energy of their corresponding singular values, using the same grouping strategy as in the main text. For a selected group  $\mathcal{H}$ , we construct structured perturbations that modify the contribution of these output directions while mixing them with all input directions:

$$\tilde{\Delta} = \sum_{i \in \mathcal{H}} \sum_{j=1}^r \beta_{i,j} \mathbf{u}_i \mathbf{v}_j^\top, \beta_{i,j} \sim \mathcal{N}(0, 1).$$

To ensure comparability across groups, the perturbation is normalized and rescaled to a fixed magnitude:

$$\tilde{\Delta} = \varepsilon \cdot \frac{\Delta}{\|\Delta\|_F}.$$

The perturbed weight matrix is given by  $\mathbf{W}' = \mathbf{W} + \tilde{\Delta}$ . This construction can be interpreted as altering the *input representation* of selected outputs  $\{\mathbf{u}_i\}_{i \in \mathcal{H}}$  (left singular vector) into random mixtures of all inputs  $\{\mathbf{v}_j\}_{j=1}^r$  (right singular vector).

Figure 13 reports the robustness curves under output-side perturbations for different singular value groups. The observed trends closely mirror those of the input-side experiments: perturbations aligned with dominant singular directions lead to substantially larger degradation in model performance than perturbations confined to lower-energy components.

## F.2 Spectral Dynamics of AlphaEdit under Sequential Editing

This section supplements the analysis in Section 2.4 by examining the spectral dynamics of

the strongest existing baseline, ALPHAEDIT, under long-sequence editing. We consider the same experimental setting as in the main text: 10,000 sequential edits (100 edits per round for 100 rounds) on the COUNTERFACT dataset using the LLAMA3 model. We track the evolution of the Singular Vector Similarity (SS) metric for dominant singular vectors throughout the editing process.

As shown in Figure 14, ALPHAEDIT exhibits relatively small deviations in dominant singular vectors during the early editing rounds. However, as edits accumulate, the maximum SS value steadily decreases, indicating progressive rotation of critical singular directions. By the end of the editing sequence, the maximum SS drops below 0.3, suggesting substantial distortion of the dominant singular subspace. This spectral degradation closely mirrors the collapse in general abilities observed on the GLUE benchmark (Figure 9). These results further support the conclusion that corruption of the dominant singular subspace is a general failure mode of long-sequence knowledge editing, rather than a method-specific artifact.

## F.3 Dynamics of Left Singular Vectors

In Section 2.4, we focus on the evolution of right singular vectors to characterize microscopic degradation of the dominant subspace during sequential editing. Here, we provide a complementary analysis on the corresponding *left* singular vectors. As shown in Figure 15, left singular vectors exhibit the same qualitative behavior as their right counterparts: dominant singular directions progressively rotate away from their original orientations as edits accumulate. This consistency confirms that sequential editing induces a symmetric degradation of input–output mappings, rather than an artifact specific to one side of the decomposition.

## G Supplementary Experimental Results

### G.1 Performance on GPT2-XL under Sequential Editing

In this subsection, we additionally report the results on GPT2-XL (1.5B). Table 4 summarizes the performance of REVIVE under sequential editing on this model. Across all evaluated baselines, incorporating REVIVE consistently improves editing performance and stability. These results indicate that the effectiveness of REVIVE generalizes well across different model sizes.

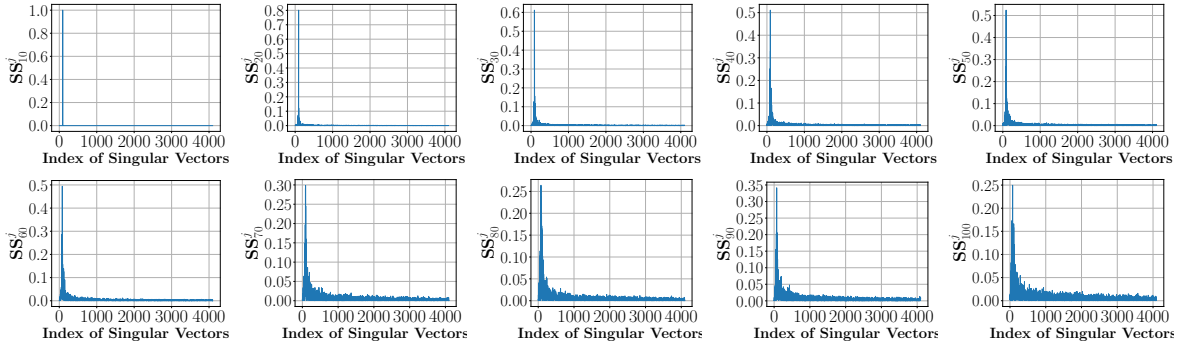


Figure 14: Evolution of SS of AlphaEdit over sequential editing, from  $SS_{10}$  to  $SS_{100}$  with step size 10.




Figure 15: Evolution of Left Singular Vector Similarity (SS) over sequential editing, from  $SS_2$  to  $SS_{20}$  with step size 2.

## G.2 Performance Comparison with DeltaEdit

DeltaEdit (Cao et al., 2025) is evaluated under an experimental protocol that differs substantially from our main setting. Specifically, DeltaEdit performs one edit per sample and reports results over 3,000 edited samples, whereas our primary experiments apply batched updates (100 edits per round) over a total of 10,000 edits. To ensure a fair and meaningful comparison, we therefore conduct a separate evaluation of DeltaEdit under its original experimental setting. Moreover, since DeltaEdit is proposed as an extension of AlphaEdit, we compare it against ALPHAEDIT+REVIVE to ensure a fair and meaningful evaluation.

Table 5 reports the results. Both DeltaEdit and the REVIVE-enhanced model substantially improve over the AlphaEdit baseline across most metrics. Notably, ALPHAEDIT+REVIVE achieves the strongest overall performance, outperforming DeltaEdit on all metrics except for *Neighborhood Score* on COUNTERFACT. We further note that REVIVE and DeltaEdit address long-horizon degradation from different perspectives. REVIVE prevents collapse by protecting a model-derived dominant structure during parameter updates, which yields broad and consistent improvements across metrics.

In contrast, DeltaEdit primarily improves edit quality by reducing interference between edits. Overall, REVIVE delivers more pronounced gains in our evaluation.

## G.3 Full Glue Results

This section reports the remaining GLUE evaluation results that were omitted from Section 4.2 due to space constraints. Results for the remaining three GLUE tasks are shown in Figure 16.

## G.4 REVIVE Enhanced Baselines Under Extreme Settings

This subsection reports the complete results of REVIVE-enhanced methods under extreme sequential editing settings. Table 6 presents detailed performance metrics, including results on the ZSRE dataset that were omitted from the main text due to space constraints.

## G.5 Full Threshold Experiments Results

This section provides detailed results that complement the sensitivity analysis reported in Section 4.2. We present extended experiments on the singular value energy threshold  $\tau$ , which determines the size of the dominant singular subspace protected by

Table 4: Performance on sequential editing over 10,000 Samples. The abbreviations *Eff.* (Efficacy Success), *Para.* (Paraphrase Success), *Neigh.* (Neighborhood Success), *Flu.* (Generation Entropy), and *Consis.* (Reference Score) denote respective evaluation metrics. Relative improvements (%) are shown in blue and decreases in orange.  $\uparrow\uparrow$  indicates a large improvement where the baseline score was near zero.

Method	Counterfact					ZsRE		
	Eff.↑	Para.↑	Neigh.↑	Flu.↑	Consis.↑	Eff.↑	Para.↑	Neigh.↑
<b>GPT2-XL</b>	21.82	24.16	78.32	626.69	31.34	22.17	21.28	24.20
MEMIT	70.56	62.42	55.94	516.26	8.74	53.00	46.27	12.76
+REVIVE	90.46 $\uparrow$ 28.2%	75.88 $\uparrow$ 21.5%	63.83 $\uparrow$ 14.1%	598.21 $\uparrow$ 15.9%	34.32 $\uparrow$ 292.7%	68.20 $\uparrow$ 28.7%	60.80 $\uparrow$ 31.4%	27.09 $\uparrow$ 112.4%
PRUNE	57.61	54.01	52.87	596.56	6.93	0.21	0.19	2.06
+REVIVE	82.00 $\uparrow$ 42.4%	70.90 $\uparrow$ 31.3%	62.82 $\uparrow$ 18.8%	600.99 $\uparrow$ 0.7%	34.55 $\uparrow$ 398.1%	40.92 $\uparrow\uparrow$	37.61 $\uparrow\uparrow$	25.29 $\uparrow$ 1127.2%
RECT	86.52	69.50	55.71	499.64	11.41	29.80	27.17	6.94
+REVIVE	82.99 $\downarrow$ 4.1%	69.20 $\downarrow$ 0.4%	65.60 $\uparrow$ 17.7%	595.69 $\uparrow$ 19.2%	34.05 $\uparrow$ 198.3%	62.45 $\uparrow$ 109.6%	55.17 $\uparrow$ 103.0%	26.20 $\uparrow$ 278.0%
AlphaEdit	92.13	76.80	56.85	581.49	31.72	53.00	46.27	12.76
+REVIVE	94.48 $\uparrow$ 2.6%	78.70 $\uparrow$ 2.5%	62.87 $\uparrow$ 10.6%	587.94 $\uparrow$ 1.1%	38.51 $\uparrow$ 21.5%	68.10 $\uparrow$ 28.5%	57.17 $\uparrow$ 23.5%	20.35 $\uparrow$ 59.4%
NSE	69.22	54.54	69.26	596.41	28.87	33.71	32.31	22.70
+REVIVE	96.12 $\uparrow$ 38.8%	84.49 $\uparrow$ 54.9%	64.17 $\downarrow$ 7.4%	592.71 $\downarrow$ 0.6%	37.74 $\uparrow$ 30.9%	77.83 $\uparrow$ 131.0%	70.55 $\uparrow$ 118.4%	24.84 $\uparrow$ 9.4%

Table 5: Comparison between AlphaEdit-REVIVE and DeltaEdit.

Model	Method	CounterFact			ZsRE		
		Eff.↑	Para.↑	Neigh.↑	Eff.↑	Para.↑	Neigh.↑
<b>Llama3-8B</b>	MEMIT	50.54	50.50	49.46	0.05	0.05	1.41
	PRUNE	50.33	50.40	49.63	0	0	0
	RECT	49.02	49.23	51.00	0	0	0
	DeltaEdit	98.63	83.20	<b>78.67</b>	94.94	91.19	30.08
	AlphaEdit	93.83	84.75	64.44	94.49	91.52	29.77
	+REVIVE	<b>99.37</b>	<b>93.50</b>	64.40	<b>95.19</b>	<b>91.58</b>	<b>32.17</b>

Table 6: Performance results of REVIVE enhanced Baselines under extreme sequential editing (20,000 edits).

Model	Method	Counterfact					ZsRE		
		Eff.↑	Para.↑	Neigh.↑	Flu.↑	Consis.↑	Eff.↑	Para.↑	Neigh.↑
LLaMA3	MEMIT-REVIVE	91.94	79.67	56.90	557.61	26.44	84.11	79.85	32.92
	RECT-REVIVE	89.00	76.78	60.54	594.38	27.93	79.35	76.35	30.24
	AlphaEdit-REVIVE	97.50	87.24	57.65	613.22	32.77	92.62	88.25	31.31
	NSE-REVIVE	98.50	90.38	61.78	615.65	33.23	93.91	89.67	31.58

REVIVE. We report results for MEMIT+REVIVE across three base models. Table 7 summarizes the variations of editing-related metrics under different threshold values. In addition, Figure 17 reports the corresponding GLUE performance on LLaMA3 for different choices of  $\tau$ . Consistent with the observations in the main text, the results indicate that REVIVE maintains stable performance across a broad range of threshold values. Results on GPT-J and GPT2-XL for GLUE are omitted, as these

models exhibit limited GLUE performance prior to editing, making threshold-dependent variations less informative in this setting.

## G.6 Preserving Dominant Singular Directions with REVIVE

We now investigate whether the spectral degradation observed above can be mitigated by explicitly protecting the dominant singular subspace. To this end, we compare the singular vector dynamics of




Figure 16: Baseline and corresponding REVIVE version(\*) performance on GLUE across datasets.

MEMIT with those of MEMIT-REVIVE. All experiments follow the same setup as in Section 2.4. Figure 18 reports the evolution of the maximum SS value for MEMIT-REVIVE throughout the entire editing sequence.

In contrast to MEMIT, MEMIT-REVIVE consistently maintains a maximum SS value of 1 across all rounds, indicating that the dominant singular vectors remain aligned with their original directions throughout the editing process. This absence of spectral drift demonstrates that REVIVE effectively shields the critical subspace from cumulative perturbations. Importantly, this spectral stability directly corresponds to the preserved general performance observed on the GLUE benchmark. Together with the results in Section 2.4, these findings provide strong evidence that explicitly protecting the dominant singular subspace is sufficient to prevent the collapse of general abilities under long-sequence knowledge editing.

## G.7 Generalization to Meta-Learning Methods

To evaluate whether REVIVE generalizes beyond locate-then-edit methods to meta-learning-based editors, we conduct additional experiments using MEND (Mitchell et al., 2022a) on the GPT-J model. Because meta-learning methods compute parameter updates independently for each sample, we evaluate sequential editing with a batch size of 1 and scale the sequence up to 1,000 edits.

As shown in Table 8, the standard MEND model suffers from rapid catastrophic collapse under sequential updates. For instance, on the ZSRE dataset, its Efficacy drops to 2.40% after just 50 edits and completely collapses to 0.00% by 1,000 edits. In sharp contrast, applying REVIVE as a plug-

and-play module effectively stabilizes the process. After 1,000 sequential edits, MEND equipped with REVIVE maintains 81.20% Efficacy on ZSRE and 86.60% on COUNTERFACT, while generative metrics such as Fluency and Consistency remain stable.

This strong generalization stems from REVIVE’s structural design. REVIVE operates purely on the resulting update matrix within the spectral basis of the pretrained weights, making no assumptions about the underlying algorithm or hypernetwork generating the update. Consequently, whether the update is derived from a closed-form editor or a meta-learning method, REVIVE can consistently filter out harmful components aligned with the dominant singular subspace, thereby preserving the model’s general abilities.

## G.8 Computational Overhead Analysis

REVIVE introduces additional spectral projections based on singular value decomposition, which may incur extra computational cost. To quantify this overhead, we measure the runtime of a baseline editing method and its REVIVE-augmented variant when editing 100 samples. Since the projection operation is independent of the underlying editing algorithm, we conduct this analysis using MEMIT as a representative baseline. As can be seen from Table 9, since knowledge editing is applied only to the FFN weight matrices in a subset of model layers, the number of required SVD computations is limited, making the additional projection overhead negligible.

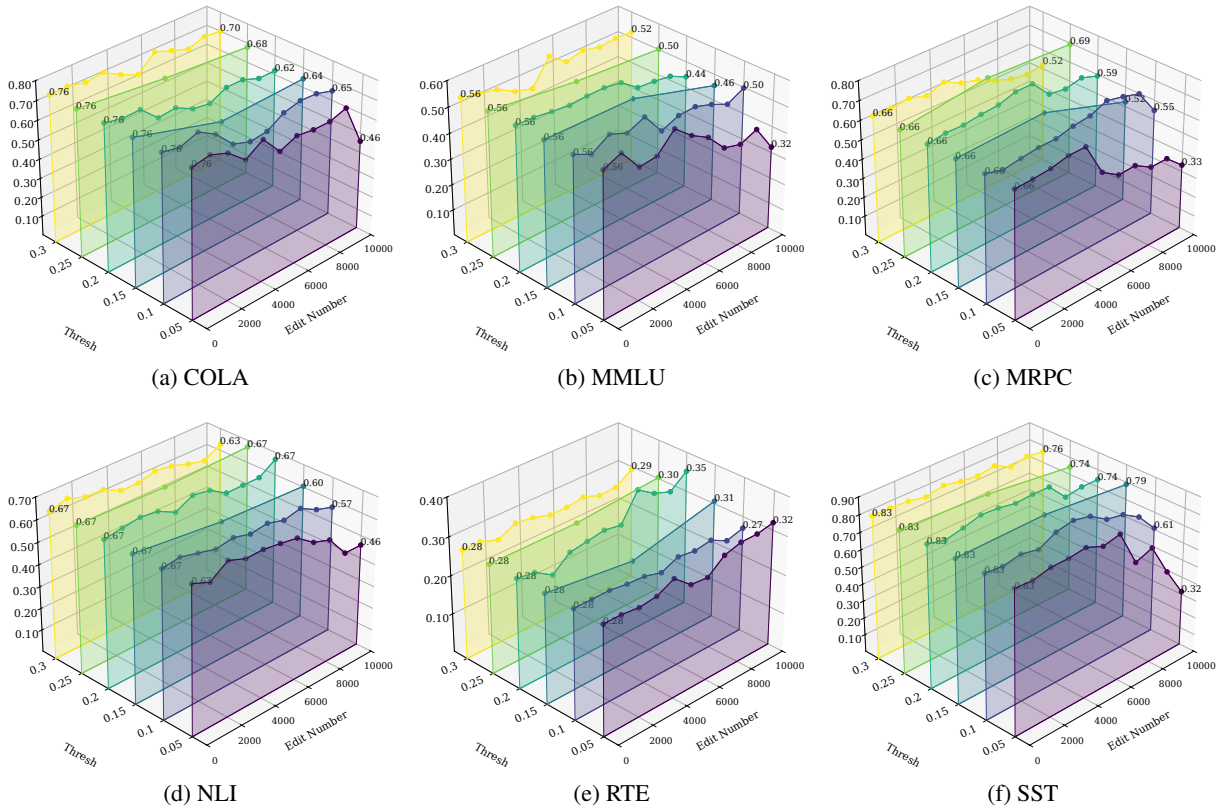


Figure 17: GLUE evaluation results on LLaMA3 after 10,000 edits on the CounterFact dataset using MEMIT-REVIVE with different protection thresholds.

Table 9: Runtime comparison under the setting of editing 100 samples.

Method	GPT-J	LLaMA3-8B
MEMIT	319.21s	813.86s
+REVIVE	332.14s	829.39s

Table 7: Performance results of MEMIT-REVIVE on sequential editing task under different singular value energy thresholds (10,000 Samples from CounterFact).

Model	Thresh	Counterfact					ZsRE		
		Eff.↑	Para.↑	Neigh.↑	Flu.↑	Consis.↑	Eff.↑	Para.↑	Neigh.↑
LLaMA3	0.05	94.46	86.03	59.70	587.30	29.83	78.66	75.89	29.29
	0.10	95.62	84.60	62.17	603.22	29.39	86.56	83.07	31.88
	0.15	95.03	80.60	64.49	613.66	29.19	87.10	83.36	32.41
	0.20	94.58	78.38	66.19	621.15	29.63	86.85	83.46	32.75
	0.25	92.96	73.94	68.94	624.63	29.49	83.85	80.23	33.27
	0.30	88.94	67.56	71.86	625.68	28.84	81.18	77.81	33.02
GPT-J	0.05	91.23	83.72	57.26	596.20	33.29	78.50	73.19	27.44
	0.10	97.09	87.01	67.10	616.15	40.00	83.87	77.28	29.77
	0.15	96.74	81.20	69.98	617.42	39.63	88.57	82.87	29.15
	0.20	94.95	76.59	72.13	621.01	38.36	85.83	79.97	29.27
	0.25	92.84	69.42	74.15	621.61	37.19	81.67	74.67	27.59
	0.30	88.49	64.30	74.94	623.53	36.66	81.32	73.54	28.56
GPT2-XL	0.05	91.89	80.72	61.13	575.14	32.12	62.13	55.40	25.90
	0.10	90.82	77.24	63.73	595.36	34.28	63.34	55.29	25.93
	0.15	87.82	73.39	65.89	607.06	35.33	66.19	58.40	27.13
	0.20	83.10	66.95	68.44	615.17	35.46	64.53	57.45	26.60
	0.25	78.77	61.82	69.66	618.28	35.17	58.11	51.80	26.89
	0.30	73.03	57.28	71.12	621.46	34.60	57.05	51.15	26.42

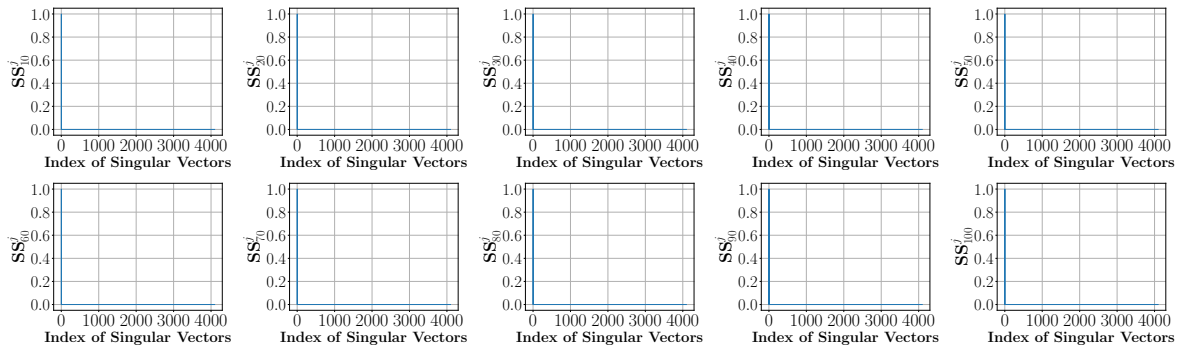


Figure 18: Evolution of SS of MEMIT-REVIVE over sequential editing, from  $SS_{10}$  to  $SS_{100}$  with step size 10.

Table 8: Performance comparison of MEND and MEND + REVIVE on the GPT-J model under sequential editing. Results demonstrate that REVIVE successfully prevents catastrophic collapse for meta-learning methods.

Edits	Method	COUNTERFACT					ZSRE		
		Eff. $\uparrow$	Para. $\uparrow$	Neigh. $\uparrow$	Flu. $\uparrow$	Consis. $\uparrow$	Eff. $\uparrow$	Para. $\uparrow$	Neigh. $\uparrow$
1	MEND	100.00	50.00	30.00	627.95	31.78	75.00	75.00	66.67
	+ REVIVE	100.00	50.00	30.00	626.93	31.08	75.00	75.00	66.67
50	MEND	54.00	48.00	54.00	488.27	6.18	2.40	1.60	93.11
	+ REVIVE	86.00	41.00	65.60	617.13	28.19	95.91	91.25	30.27
100	MEND	43.00	42.50	58.50	524.40	1.19	0.20	0.60	0.02
	+ REVIVE	89.00	45.50	60.26	618.11	28.98	97.41	89.98	27.85
500	MEND	45.40	45.90	54.10	498.58	2.92	0.34	0.34	0.64
	+ REVIVE	85.20	45.50	60.28	615.47	28.85	86.01	78.96	27.09
1000	MEND	46.10	46.40	53.60	488.05	2.41	0.00	0.00	0.08
	+ REVIVE	86.60	50.80	56.77	619.08	29.38	81.20	74.38	25.97