

Aligning Language Models with Real-time Knowledge Editing

Chenming Tang* Yutong Yang* Kexue Wang Yunfang Wu†

National Key Laboratory for Multimedia Information Processing, Peking University
School of Computer Science, Peking University

{tangchenming, 2300013219}@stu.pku.edu.cn {yytpku, wuyf}@pku.edu.cn

*Equal contribution †Corresponding author

Abstract

Knowledge editing aims to modify outdated knowledge in language models efficiently while retaining their original capabilities. Mainstream datasets for knowledge editing are predominantly static and fail to keep in pace with the evolving real-world knowledge. In this work, we introduce CRAFT, an ever-evolving real-world dataset for knowledge editing. It evaluates models on temporal locality, common-sense locality, composite portability and alias portability, providing a comprehensive and challenging evaluation for knowledge editing, on which previous methods hardly achieve balanced performance. Towards flexible real-time knowledge editing, we propose KEDAS, a novel paradigm of knowledge editing alignment featuring diverse edit augmentation and self-adaptive post-alignment inference, exhibiting significant performance gain on both CRAFT and traditional datasets compared to previous methods. We hope this work may serve as a catalyst for shifting the focus of knowledge editing from static update to dynamic evolution.¹

1 Introduction

Once language models (LMs) are pre-trained, knowledge is injected into them and becomes their static internal capability (Petroni et al., 2019). As time goes by, some knowledge in LMs inevitably becomes incorrect or out of date, and it is vital to update the outdated knowledge. However, re-training a model frequently is highly costly and usually infeasible. To this end, the techniques of knowledge editing have been developed (Zhang et al., 2024) to edit specific knowledge in LMs efficiently without heavy re-training.

There have been a large number of datasets to evaluate the effectiveness of knowledge editing approaches. However, most traditional ones like

¹All of the code and data of this work are publicly available at <https://github.com/JamyDon/CRAFT-KEDAS>.

Dataset	RT	RW	RP	CRP
ZsRE	✗	✓	✗	✗
MQuAKE-T	✗	✓	✓	✗
MQuAKE-CF	✗	✗	✓	✓
RippleEdits	✗	✓	✓	✗
WikiBigEdit	✓	✓	✗	✗
EvoWiki	△	✓	✓	✗
CRAFT (Ours)	✓	✓	✓	✓

Table 1: Comparison of knowledge editing datasets. **RT** refers to “real-time” and **RW** denotes “real-world”. **RP** (reasoning portability) tests reasoning chains derived from a single edit, whereas **CRP** (composite reasoning portability) evaluates a model’s ability to integrate multiple distinct edits into a single reasoning query.

ZsRE (Levy et al., 2017), MQuAKE (Zhong et al., 2023) and RippleEdits (Cohen et al., 2024a) are static and not real-time. Once the datasets are constructed and released, they remain fixed and can not be updated anymore, which is far from real-world knowledge editing applications. Recently, there have been a few real-time datasets released. For example, WikiBigEdit (Thede et al., 2025) collects real-world changes in Wikipedia, but it requires processing massive Wiki data (hundreds of gigabytes) from the Internet, which is costly and inconvenient in practice. Meanwhile, the automatically collected dataset lacks necessary filtering and thus suffers from severe sparsity (for instance, only 3.6% of the collected data have an entry of portability evaluation). EvoWiki (Tang et al., 2025) is another real-time dataset based on Wikipedia, but it focuses on retrieval-augmented generation (RAG) and continual learning and does not support comprehensive evaluation for knowledge editing due to lack of crucial aspects like locality. Overall, there still lacks a high-quality, easily-updated dataset for real-world and real-time knowledge editing.

In this work, we introduce CRAFT (**C**hinese **R**eal-time statistics **A**nd **F**inance knowledge editing dataset**T**), an ever-evolving real-world dataset

for knowledge editing in Chinese. It leverages publicly available official data that are continuously updated, ensuring both temporal freshness and real-world applications. CRAFT is organized in well-designed paired edits, each serving as a composite reasoning test to evaluate models’ ability to integrate multiple related factual updates. Meanwhile, it supports evaluations on alias portability and temporal and common-sense locality, providing a comprehensive assessment of LMs’ adaptability and robustness under dynamic knowledge changes. The differences between CRAFT and previous datasets are summarized in Table 1.

We perform exposure analysis on CRAFT and other evaluation datasets with five different LMs. The results indicate that a large portion of knowledge in existing datasets has been known to LMs while the exposure rate of CRAFT is nearly zero. This validates the real-time property of CRAFT, which makes it challenging and fair.

We evaluate a suite of representative knowledge editing methods on CRAFT and demonstrate current approaches exhibit inherent limitations. Parameter-based approaches like ROME (Meng et al., 2022) and WISE (Wang et al., 2024a) suffer from gradual model degradation and scarcely edit successfully when new knowledge comes sequentially and continually. Simple retrieval-based approaches like IKE (Zheng et al., 2023) and EREN (Chen et al., 2024b) fail to achieve consistent balanced performance due to lack of alignment. LTE (Jiang et al., 2024), which aligns LMs with knowledge editing via finetuning, shows weak locality due to overfitting on irrelevant queries.

Towards flexible real-time knowledge editing, we propose KEDAS², namely **K**nowledge **E**dit alignment with **D**iverse **A**ugmentation and **S**elf-adaptive inference, an advanced knowledge editing framework featuring diverse representations of edits and flexible inference paths. The one-time offline alignment stage first finetunes the LM with low-rank adaptation (LoRA) (Hu et al., 2022). Once the alignment is finished, users may edit any knowledge without modifying any parameters. During editing, we introduce diverse edit augmentation to store different forms of each edit in the memory. In the inference phase, a filter-enhanced smart retriever is employed to adaptively select the base model or the aligned model. If any edits are recalled, the prompt will go through the

LM with LoRA activated. Otherwise, the prompt just passes to the original LM to guarantee locality and avoid over-fitting. On both CRAFT and traditional datasets, KEDAS exhibits outstanding performance in all metrics and significantly outperforms previous methods, illustrating an ideal paradigm of knowledge editing alignment.

Our contributions are as follows:

- We reveal the leakage issue of existing datasets and introduce CRAFT, an open-source, contamination-free and easily-updated dataset for real-time knowledge editing.
- We evaluate existing knowledge editing methods on CRAFT and reveal the limitations of these methods in real-time knowledge editing, especially the tradeoff among edit success, locality and portability.
- We propose KEDAS, a novel paradigm of aligning LMs with real-time knowledge editing which significantly outperforms existing methods with balanced performance.

This paper is organized as follows. We introduce the preliminary of knowledge editing in § 2. Then, we present our CRAFT dataset in § 3 and propose KEDAS for real-time knowledge editing in § 4. After that, we conduct empirical experiments in § 5. Finally, we discuss related work in § 6 and conclude our study in § 7.

2 Preliminary

2.1 Task Formulation

An LM system f can be regarded as a function $f : \mathcal{Q} \mapsto \mathcal{A}$ mapping an query $q \in \mathcal{Q}$ to its output answer $a = f(q) \in \mathcal{A}$.

Knowledge editing aims to change the behavior of the LM after modifying some knowledge. Edits of knowledge are usually in the form of query-answer (QA) pairs:

$$\mathcal{E} = \{e^t\}_{t=1}^N = \{(q_e^t, a_e^t)\}_{t=1}^N, \quad (1)$$

where q_e^t is an input query triggering knowledge (*e.g.*, The current US president is), a_e^t is the corresponding target of editing (*e.g.*, Donald Trump), t is the index of each edit, and N denotes the total number of edits.

To assess the efficacy of knowledge editing, the post-edit LM f^* is evaluated via the following criteria (Zhang et al., 2024):

²/ki:dəs/, pronounced as “kee-das”.

Edit success measures the accuracy of editing, requiring f^* to correctly recall the edits:

$$\mathbb{E}_{(q_e, a_e) \in \mathcal{E}} \mathbb{1}\{f^*(q_e) = a_e\}. \quad (2)$$

Locality measures the precision of editing, requiring f^* not to change its behavior out of the scope of edits:

$$\mathbb{E}_{(q_l, a_l) \in \mathcal{E}_l} \mathbb{1}\{f^*(q_l) = f(q_l)\}, \quad (3)$$

where \mathcal{E}_l are QA pairs of unrelated queries. Note that the target of locality is unchanged answer rather than exactly the ground truth answer.

Portability measures how well edits transfer to related queries, requiring f^* to correctly answer such queries:

$$\mathbb{E}_{(q_p, a_p) \in \mathcal{E}_p} \mathbb{1}\{f^*(q_p) = a_p\}, \quad (4)$$

where \mathcal{E}_p are QA pairs related to existing edits.

2.2 Settings of Knowledge Editing

Single Editing In single editing, each edit is evaluated directly after it is applied to the original model f :

$$f^t = \mathbf{Edit}(f, q_e^t, a_e^t), \quad 1 \leq t \leq N. \quad (5)$$

This traditional setting is widely employed by earlier work. It is far from real-world applications, since only one edit can be applied.

Sequential Editing In sequential editing, edits are applied step by step in a continual manner (f^0 denotes f for consistency):

$$f^t = \mathbf{Edit}(f^{t-1}, q_e^t, a_e^t), \quad 1 \leq t \leq N. \quad (6)$$

The final model f^N after applying all N edits is evaluated. This setting is frequently used to evaluate the lifelong performance and scalability of knowledge editing. We focus on this setting in this work because it is closer to practical applications, especially the real-time knowledge editing setting.

2.3 Knowledge Editing Alignment

LTE (Jiang et al., 2024) relates knowledge editing to LM alignment, aligning LMs with ever-changing real-time knowledge edits. During alignment, LMs’ capabilities of knowledge updating are enlightened by an editing prompt:

```
[Updated Information] {edit}
[Query] {query}
```

The training data consist of in-scope and out-of-scope queries with or without the specific edit to improve edit efficacy while keeping locality.

In the inference phase, LMs are required to conduct on-the-fly and streaming knowledge editing by retrieving relevant edits to the query from the stored memory.

3 The CRAFT Dataset

We introduce CRAFT, a real-time and real-world dataset designed to evaluate knowledge editing under dynamically evolving factual contexts. Unlike static datasets that rely on outdated or widely exposed Wikipedia facts, CRAFT continuously updates with publicly available official data sources, ensuring fairness and real-world relevance.

3.1 Data Sources

CRAFT consists of two complementary subsets, namely **CRAFT-Statistics** and **CRAFT-Finance**, representing two highly correlated real-world domains that evolve periodically. Example instances from the two subsets are shown in Appendix G and H respectively.

CRAFT-Statistics We collect monthly statistical reports from the National Bureau of Statistics of China³, covering July 2024 to June 2025 in the current version used in this work. The dataset spans four major domains – finance, fiscal affairs, telecommunications, and transportation – containing 221 indicators and 4,468 data points. Our data collection pipeline, leveraging the `cn-stats` API⁴, supports automatic temporal updates, enabling effortless collection of the most recent monthly data.

CRAFT-Finance CRAFT-Finance includes annual financial statements of 390 publicly listed Chinese companies (2023–2024), collected via the `AKShare` API⁵. It features financial indicators such as debt ratio, shareholder equity, and operating cash flow, supporting yearly updates to reflect the latest financial disclosures.

3.2 Evaluation Metrics

To comprehensively assess model adaptability, CRAFT provides four evaluation dimensions besides the conventional edit success (ES) as described in § 2.1:

³<https://data.stats.gov.cn/>

⁴<https://github.com/songjian/cnstats>

⁵<https://github.com/akfamily/akshare>

Dataset	GPT-3.5 (OpenAI, 2023)	Mistral-7B (Jiang et al., 2023)	Llama-3.1-8B (Llama Team, 2024)	DeepSeek-chat (DeepSeek-AI, 2025)	Doubao-1.5-pro-32k (ByteDance, 2025)	Size
ZsRE (Levy et al., 2017)	43.98%	23.48%	22.45%	52.47%	26.77%	6,000
MQuAKE-T (Zhong et al., 2023)	35.46%	76.21%	46.11%	42.58%	76.09%	5,604
EvoWiki ^{evolved} (Tang et al., 2025)	8.07%	16.22%	6.43%	10.24%	5.08%	1,338
EvoWiki ^{stable} (Tang et al., 2025)	28.45%	42.37%	25.90%	39.76%	12.32%	1,494
EvoWiki ^{uncharted} (Tang et al., 2025)	8.19%	22.62%	7.39%	16.02%	3.43%	1,136
CRAFT-Statistics (Ours)	0.01%	0.00%	0.30%	1.80%	0.99%	4,468
CRAFT-Finance (Ours)	0.00%	0.00%	0.14%	0.42%	0.00%	7,800

Table 2: Exposure rate (%) of different LLMs to benchmark datasets. Model release years are indicated in citations. CRAFT exhibits minimal exposure due to real-time data collection and domain freshness.

Composite Reasoning Portability ($P_{\text{reasoning}}$)

We propose *composite reasoning portability*, a new evaluation designed specifically for CRAFT. Unlike traditional *multi-hop reasoning portability*, which constructs reasoning chains from a single edit, composite portability evaluates a model’s ability to integrate multiple distinct edits into a single reasoning query.

In CRAFT, all data are organized in *paired format*: each test instance consists of *two edits* and one corresponding composite reasoning query. Formally, given a pair of edits $\{e_1, e_2\}$, the composite portability query q_p requires reasoning over both edits simultaneously. For example, consider a CRAFT-Finance instance with a pair of edits:

- e_1 : 2024 Company A’s total assets = 2M
- e_2 : 2024 Company A’s total liabilities = 1M

The corresponding composite reasoning query and answer are:

- q_p : 2024 Company A’s debt ratio
- a_p : 0.5

CRAFT-Finance includes 10 types of such combinations, demonstrating the generality and practical relevance of composite reasoning portability. For CRAFT-Statistics, composite portability is constructed by pairing two edits of the same indicator from adjacent years.

Alias Portability (P_{alias}) We propose *alias portability* to evaluate a model’s robustness to synonymous indicator names by replacing key terms with domain-specific aliases constructed from a manually curated alias list.

Temporal Locality (L_{temporal}) We propose *temporal locality* to how well a model’s unrelated general knowledge and reasoning capability are preserved after knowledge editing.

Common-sense Locality (L_{common}) We propose *common-sense locality* to measure whether a model’s irrelevant general knowledge and reasoning capability remain consistent after knowledge editing. For CRAFT, the questions for this evaluation are directly sampled from the C^3 dataset (Sun et al., 2020), which contains 19,577 Chinese multiple-choice comprehension questions.

3.3 Exposure Analysis

A perfect evaluation dataset should be unlearned, with a low portion of knowledge has been exposed to LMs (*i.e.*, a low exposure rate). If a large portion of knowledge has already been known by LMs, the evaluation results on it can be inaccurate, unfair, or even misleading.

We evaluate the exposure rates of CRAFT and existing benchmarks across five representative LMs by directly prompting with the questions in the datasets and counting how many questions can be correctly answered. As shown in Table 2, Due to its temporal freshness and domain specificity, CRAFT exhibits extremely low exposure rates (nearly 0%) across all the five LMs, mitigating data contamination and ensuring fair evaluation. Meanwhile, for other datasets, a great portion of the knowledge has been known to the LMs, making them less challenging and less fair in practice. This reveals the advantage of CRAFT as a truly real-time dataset.

3.4 Dataset Statistics

Subset	Edit	$P_{\text{reasoning}}$	P_{alias}	L_{temporal}	L_{common}
CRAFT-Statistics	4,468	2,234	4,468	4,468	4,468
CRAFT-Finance	7,800	3,900	7,800	7,800	7,800

Table 3: The numbers of instances in different evaluation types of CRAFT. Composite reasoning uses edit pairs; other metrics are per edit.

Table 3 shows the numbers of instances in

CRAFT’s subsets. Composite reasoning portability is evaluated per edit pair, and other metrics per edit. For experiments, we split each subset of CRAFT into training and test sets. The test set contains 500 instances from the two subsets each (1,000 instances in total) and the training set contains the remaining instances.

3.5 Discussion

CRAFT is a scalable and continuously evolving dataset reflecting real-world knowledge updates. Its dynamic nature, factual authenticity, and comprehensive reasoning evaluation make it a robust foundation for future research on reliable knowledge editing. Anyone who wants to construct CRAFT within an arbitrary time period can easily obtain a customized dataset with our provided open-source scripts. Through CRAFT, we also demonstrate that researchers can make full use of public available data sources (like public APIs) besides the widely used Wikipedia data to create datasets with real-world and real-time properties and scalability that can help evaluate, train or develop more scalable intelligent systems without much human labor.

4 KEDAS: A Proposal for Real-time Knowledge Editing Alignment

As will be shown in the experiments (§ 5), previous methods show certain limitations either in recall (*i.e.*, edit success and portability) or precision (*i.e.*, locality) on CRAFT. To improve edit success and generalization while keeping a good locality in real-time knowledge editing, we propose KEDAS, an all-round approach consisting of several complementary designs.

The overall framework of KEDAS is illustrated in Figure 1, including alignment, editing and inference. In alignment, the LM is aligned with knowledge editing using LoRA. In editing, each edit is converted into diverse forms and stored to the memory module. In the inference phase, edits are retrieved from the memory based on the user query and then filtered by a trained classifier. If there exist relevant edits, the model with trained LoRA adapters will be called for inference. Otherwise, the query will go through the frozen base model. The order of editing and inference is flexible based on practical needs (“edit as you go”), as shown by the bidirectional arrow. For example, we may apply several edits and employ the LM for

inference immediately. If there come further new edits, we can return to the editing phase to perform these edits and then resume inference.

4.1 Alignment Based on LoRA

LTE (Jiang et al., 2024) suffers from a limited locality because it employs the post-alignment model, whose behavior can be significantly different from the original one, to all incoming queries. To address this, we propose the idea of self-adaptive post-alignment inference. To achieve this elegantly, we align the base model via parameter-efficient fine-tuning with LoRA (Hu et al., 2022), which tunes additional adapters during alignment. This costs low computational resources for both training and inference, making KEDAS efficient. We follow LTE to construct alignment data using the training set of CRAFT and details are in Appendix A.

4.2 Diverse Edit Augmentation

Form	Content
QA	What are Microsoft’s total assets? \$619B.
<i>Declarative</i>	Microsoft’s total assets are \$619B.
<i>Aliased</i>	MSFT’s total assets are \$619B.

Table 4: A demonstration of diverse edit expressions.

In previous work, each piece of edited knowledge is directly applied to the system as it is, limiting generalization. Since each knowledge can be expressed in various ways, leveraging just one fixed form is inflexible, especially for challenging portability evaluations like that in CRAFT.

In this work, we propose diverse edit augmentation, a novel technique that augments each edit by converting it into multiple expressions. The *de facto* standard form of edits in knowledge editing is a QA pair, typically describing the relation between entities or attributes of them. Based on this observation, we manually design several ways of form augmentation, including *declarative* and *aliased* forms, as demonstrated in Table 4.

Concretely, for an edit $e^t = (q_e^t, a_e^t)$ at step t^6 , its conventional QA form is:

$$e_{\text{qa}}^t = q_e^t \oplus a_e^t, \quad (7)$$

where \oplus denotes concatenation. Then, it is aug-

⁶For simplicity, we deem the two edits in each CRAFT instance as two separate time steps.

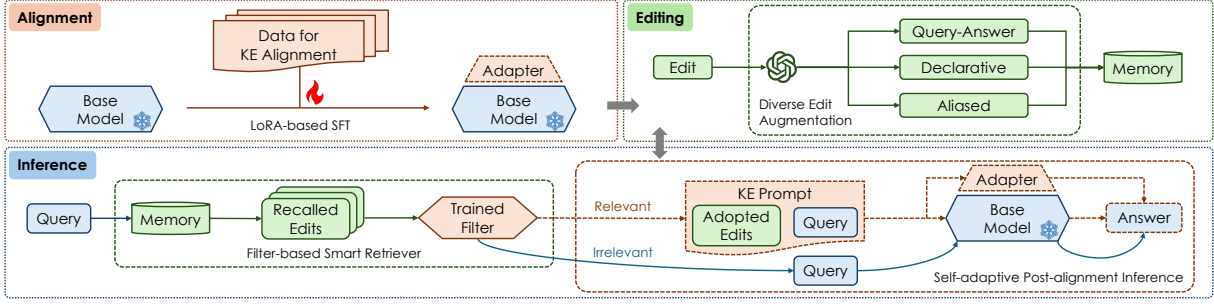


Figure 1: Illustration of our proposed KEDAS framework.

mented as:

$$e_{\text{dec}}^t = \mathbf{Declarative}(q_e^t, a_e^t), \quad (8)$$

$$e_{\text{als}}^t = \mathbf{Aliased}(q_e^t, a_e^t). \quad (9)$$

Specifically, *declarative* is to convert the QA pair into a declaration while *aliased* is to replace specialized terms in the declaration with their aliases for more flexible retrieval. In experiments, we conduct these augmentations by prompting an API-based off-the-shelf LM.

We store the augmented diverse edits into memory for future retrieval:

$$\mathcal{M}^0 = \emptyset, \quad \mathcal{M}^t = \mathcal{M}^{t-1} \cup \{e_{\text{qa}}^t, e_{\text{dec}}^t, e_{\text{als}}^t\}, \quad (10)$$

where \mathcal{M}^t denotes the memory at time step t .

Note that our augmentation methods are specifically designed for the CRAFT dataset. There can be more diverse augmentations like paraphrasing in practice for broader use.

4.3 Filter-based Smart Retriever

During inference, the common practice regarding retrieval in previous knowledge editing work is to directly employ an embedder and perform single-stage retrieval to obtain the top- k hits. This suffers from a poor locality due to the lack of a filter measuring the relevance of queries. In this work, we propose a filter-based smart retriever that pairs the high-recall retrieval with a high-precision filter by precisely identifying relevant queries. In this way, on challenging datasets like CRAFT, the post-edit system can keep consistent behaviors for neighboring locality queries and preserve its original answers to general queries.

First, we retrieve top- n candidates from the memory using a normal embedding-based retriever. Thanks to our diverse edit augmentation, the memory contains various forms of edits, improving the recall rate. At the time step t , for an input query q , a set of top- n candidates $\mathcal{C}^{\text{top-}n}$ is retrieved as:

$$\mathcal{C}^{\text{top-}n} = \mathbf{Retriever}^{\text{top-}n}(q, \mathcal{M}^t). \quad (11)$$

Then, a binary classifier θ is exploited to predict the relevance between candidates and a query:

$$\mathcal{C}^{\text{filtered}} = \{e \mid \mathbf{Filter}_{\theta}(e) = 1, e \in \mathcal{C}^{\text{top-}n}\}. \quad (12)$$

Then we only adopt the top- k edits with the highest similarity to the query if the filtered candidate set contains too many edits:

$$\mathcal{E}^* = \arg \text{top}_k(\mathbf{Sim}(e, q)) \text{ for } e \in \mathcal{C}^{\text{filtered}}. \quad (13)$$

4.4 Self-adaptive Post-alignment Inference

In the inference phase, we employ both the original pre-alignment model and the post-alignment adapter, which does not consume extra memory while enabling two inference paths. Each time a query q is requested, the filter-based smart retriever is utilized to recall possible relevant edits in the memory. If there exists a set of relevant edits \mathcal{E}^* , indicating that the query may involve edited knowledge, we fill both \mathcal{E}^* and q into the knowledge editing prompt (see § 2.3) and then pass the prompt to the model with the post-alignment adapters activated. Otherwise, the query directly goes through the pre-alignment model without any processing. The inference strategy is formularized below:

$$a = \begin{cases} f_{\Phi}(q) & \text{irrelevant} \\ f_{\Phi+\Delta\Phi(\Theta)}(\mathbf{KEPrompt}(\mathcal{E}^*, q)) & \text{relevant} \end{cases}, \quad (14)$$

where a denotes the final answer, f_{Φ} denotes the original model parametrized by Φ , $f_{\Phi+\Delta\Phi(\Theta)}$ denotes the model with post-alignment adapters $\Delta\Phi(\Theta)$, and $\mathbf{KEPrompt}(\cdot, \cdot)$ denotes the knowledge editing prompt template.

4.5 Discussion

KEDAS borrows the idea of inference-time retrieval from EREN (Chen et al., 2024b) and the idea

of knowledge editing alignment from LTE (Jiang et al., 2024). The novelty of KEDAS lies in the two unique designs: diverse edit augmentation and filter-based smart retriever paired with self-adaptive post-alignment inference. The former follows the idea that edits can go beyond their initial forms, while the latter demonstrates the insight that we do not have to perform knowledge editing for irrelevant queries, and a lightweight mechanism leads to more balanced performance.

5 Experiments

5.1 Experimental Setup

5.1.1 Evaluation

We evaluate representative knowledge editing methods and our proposed KEDAS with CRAFT. We use the metrics described in § 3.2, including ES, L_{temporal} , L_{common} , $P_{\text{reasoning}}$ and P_{alias} .

We also validate the effectiveness of KEDAS in traditional knowledge editing with four widely used datasets from KnowEdit (Zhang et al., 2024), including ZsRE (Levy et al., 2017), WikiBio (Manakul et al., 2023), WikiData_{recent} and WikiData_{counterfact} (Cohen et al., 2024b). We use the metrics described in § 2.1 for evaluation. See Appendix E for details.

In the main experiments, we focus on sequential editing for both evaluation datasets.

5.1.2 Language Models

For CRAFT, we use Llama-3.1-8B-Instruct (Llama Team, 2024) and Qwen2.5-3B-Instruct (Qwen, 2025), both from widely used herds of open-source LMs (LLAMA-3.1 and QWEN2.5) and with different scales. For KnowEdit, we focus on the same LLAMA-3.1 model and provide the results on two other LMs in Appendix I.

5.1.3 Evaluated Methods

We evaluate the following knowledge editing methods: (1) **LoRA** (Zhang et al., 2024) directly updates knowledge by LoRA-based finetuning. (2) **ROME** (Meng et al., 2022) locates multilayer perceptron weights of related knowledge and writes in new key-value pairs. (3) **IKE** (Zheng et al., 2023) leverages in-context examples of copying, updating and retaining knowledge to edit knowledge via prompting. (4) **EREN** (Chen et al., 2024b) prompts LMs to decide the relevance of retrieved edits and then queries LMs with or without edits. (5) **WISE** (Wang et al., 2024a) stores

edits in a parametric side memory module and utilized a router to decide which memory to go through. (6) **LTE** (Jiang et al., 2024) aligns LMs with the knowledge editing task by post-training. (7) **KEDAS** is our proposed method.

Due to compatibility, cost, and performance issues, we do not include some other baselines as discussed in Appendix F.

5.1.4 Implementation Details

Our experiments are based on the framework of EasyEdit (Wang et al., 2024b). For most of the baselines, we follow the default hyper-parameter settings provided by EasyEdit.

For IKE, we only evaluate the 8-shot setting because of limited resources. For both LTE and KEDAS, we use paraphrase-multilingual-MiniLM-L12-v2 (Reimers and Gurevych, 2019) as the retriever and set the number of adopted edits k as 3, the default setting of LTE. For KEDAS, we adopt bert-base-chinese (Devlin et al., 2019) as the binary filter, which is trained based on the training set of CRAFT with details specified in Appendix D. We set the number of retrieved candidates n as 8 and that of adopted edits k as 3, in line with LTE. The diverse edit augmentation of KEDAS is done by prompting `gpt-4o-mini`⁷ and the prompts are in Appendix C.

5.2 Results on CRAFT

The results on CRAFT are presented in Table 5.

Previous methods of knowledge editing show limitations in certain aspects. LoRA and ROME, as traditional methods primarily designed for single editing or a limited number of edits, show no advantage when applied to the sequential editing setting on CRAFT. IKE, the retrieval-based approach, suffers from poor locality and mediocre portability, albeit achieves outstanding edit success. The prompting-based EREN requires specific prompt optimization for different LMs and can scarcely perform editing with its default prompt, as can be seen from the low edit success. WISE is parameter-based and suffers from downgrading performance under the sequential setting, with weak performance in most metrics. LTE, with an alignment stage, achieves relatively robust performance in edit success and portability but struggles to achieve locality due to its fixed inference path. These indicate that CRAFT is challenging for existing knowl-

⁷<https://platform.openai.com>

Model	Method	CRAFT-Statistics					CRAFT-Finance				
		ES	L _{temporal}	L _{common}	P _{reasoning}	P _{alias}	ES	L _{temporal}	L _{common}	P _{reasoning}	P _{alias}
LLAMA-3.1	LoRA	47.93	0.71	0.00	41.56	47.62	44.19	13.97	0.00	9.81	7.43
	ROME	21.59	0.00	12.20	13.67	21.59	4.25	4.00	14.00	7.62	4.40
	IKE	99.82	5.49	11.20	43.03	<u>55.73</u>	99.98	27.50	11.60	40.71	57.86
	EREN	22.12	<u>91.68</u>	63.20	24.99	22.29	24.37	70.93	65.80	25.48	22.27
	WISE	47.54	29.74	<u>81.00</u>	41.26	47.30	32.66	49.62	<u>79.80</u>	34.41	32.48
	LTE	97.90	5.44	12.60	<u>68.58</u>	54.14	80.68	27.94	8.40	<u>50.21</u>	60.98
	KEDAS	<u>99.80</u>	100.00	100.00	81.20	62.17	<u>99.85</u>	<u>62.50</u>	100.00	55.95	70.61
QWEN2.5	LoRA	43.81	18.86	0.00	39.48	35.91	24.73	11.80	0.00	31.59	32.89
	ROME	35.06	32.79	0.00	29.00	26.90	9.75	2.95	0.00	14.39	9.34
	IKE	<u>96.82</u>	19.44	82.40	48.48	42.02	97.96	28.78	81.80	39.27	46.16
	EREN	28.32	<u>61.08</u>	92.60	34.55	30.67	31.25	42.43	93.00	33.69	29.53
	WISE	40.72	43.46	<u>94.40</u>	34.57	38.39	28.07	<u>51.43</u>	<u>95.00</u>	28.67	21.25
	LTE	99.95	8.38	6.60	<u>77.67</u>	58.92	99.96	31.56	5.20	51.86	<u>60.18</u>
	KEDAS	99.95	100.00	100.00	80.88	<u>58.59</u>	<u>99.88</u>	66.47	100.00	<u>51.55</u>	65.97

Table 5: Main results on CRAFT. The highest and second-highest scores are **bolded** and underlined, respectively.

edge editing methods, making it hard to achieve a consistently balanced performance.

Meanwhile, KEDAS exhibits outstanding performance on CRAFT, securing the highest or second-highest scores in all the metrics. The edit success of KEDAS is nearly 100% while all of its locality scores are 100% except for L_{temporal} on the Finance subset possibly because the BERT-based filter suffers from limited generalization in certain cases. Meanwhile, the portability of KEDAS is significantly outstanding on LLAMA-3.1 and is better or on par with LTE on QWEN2.5. Overall, KEDAS outperforms WISE, the representative parameter-based approach, by **+35.63** and **+34.73** averaged points on the two LMs respectively, and surpasses LTE, the strong alignment-based baseline, by **+36.52** and **+32.30** averaged points. This confirms that KEDAS achieves the most balanced performance on CRAFT and maintains competitive performance across all dimensions.

5.3 Results on KnowEdit

The results on KnowEdit are shown in Table 6. KEDAS exhibits outstanding performance in all the metrics and secures the highest harmonic mean (HM) scores across all the datasets. On average, its HM outperforms WISE by **+26.9** and surpasses LTE by **+18.2**. As an alignment-based approach, KEDAS is superior especially in locality compared with LTE, with an improvement of **+31.7**.

On KnowEdit, we also further experiment with other LMs and under the traditional setting of single editing. These results can be found in Appendix I, where KEDAS also secures consistent advantage with balanced performance.

5.4 Ablation Study

In this section, we assess the indispensability of components of KEDAS, including diverse edit augmentation (DEA), filter (FLT) and self-adaptive post-alignment inference (SPI). The results are presented in Table 7. The exclusion of diverse edit augmentation leads to a decline in edit success and portability, indicating that it contributes to the recall of the retriever. Meanwhile, both the filter and the paradigm of self-adaptive post-alignment inference show great contribution to locality, demonstrating that they are playing important roles in KEDAS. These results confirm that all the core components of KEDAS make contributions to the final performance and are indispensable.

6 Related Work

6.1 Related Evaluation Datasets

Conventionally, knowledge editing datasets are created based on temporal changes or counter-fact modifications of world knowledge either from existing QA datasets (Levy et al., 2017; Yao et al., 2023) or Wikipedia (Cohen et al., 2024a; Zhong et al., 2023), in which Zhong et al. (2023) first propose multi-hop composite portability. However, these are all static and suffer from data contamination as time goes by.

Thede et al. (2025) propose WikiBigEdit, a dynamic dataset based on periodical changes in Wikipedia. However, only a small fraction (3.6%) of its samples support reasoning portability evaluation, limiting its comprehensiveness. Similarly, Tang et al. (2025) propose EvoWiki, a Wiki-based real-time dataset primarily designed for RAG.

Method	ZsRE				WikiBio			WikiData _{recent}				WikiData _{counteract}			
	ES	L	P	HM	ES	L	HM	ES	L	P	HM	ES	L	P	HM
IKE	98.5	43.5	64.4	61.7	96.9	37.9	54.5	97.4	48.8	67.2	65.7	91.3	53.7	61.4	65.4
EREN	32.3	70.7	44.9	44.5	59.4	64.3	61.8	44.4	73.1	37.5	47.7	23.5	75.4	22.9	30.2
WISE	60.0	51.3	40.4	49.3	84.5	95.9	89.8	68.8	94.5	43.1	62.1	47.1	45.4	35.0	41.8
LTE	99.6	55.5	66.7	69.7	97.7	59.7	74.1	99.8	52.9	74.7	70.9	98.3	59.3	73.4	73.8
KEDAS	99.6	90.5	71.4	85.5	97.7	100.0	98.9	99.8	73.9	76.1	81.7	98.3	69.2	77.4	79.9

Table 6: Results on KnowEdit. Metrics include edit success (ES), locality (L), portability (P) and harmonic mean (HM) of the previous metrics. The highest scores are **bolded**.

Method	CRAFT-Statistics					CRAFT-Finance				
	ES	L _{temporal}	L _{common}	P _{reasoning}	P _{alias}	ES	L _{temporal}	L _{common}	P _{reasoning}	P _{alias}
KEDAS	99.80	100.00	100.00	81.20	62.17	99.85	62.50	100.00	55.95	70.61
w/o DEA	96.51	100.00	100.00	71.02	59.89	98.10	62.50	100.00	53.90	63.35
w/o FLT	99.80	5.44	11.40	81.20	62.17	100.00	27.48	8.40	55.95	70.61
w/o SPI	99.80	5.68	10.80	81.20	62.17	99.85	27.42	8.40	55.95	70.61

Table 7: Ablation results on CRAFT. The highest scores are **bolded**.

Although EvoWiki claims to automatically update with new edits, its implementation and data pipeline remain unpublished.

Besides conventional knowledge editing datasets, there have also been some designed for special use. [Mitchell et al. \(2022b\)](#) propose ConvSent to edit LMs’ sentiments on given topics. [Ishibashi and Shimodaira \(2024\)](#) propose Sanitation to forget specific information stored in LMs and thus address privacy concerns. [Huang et al. \(2025\)](#) propose HalluEditBench to evaluate whether knowledge editing correct hallucinations.

Our proposed CRAFT leverages statistical and financial data, which are both real-time and inherently dynamic. The strong interdependence among indicators allows each pair of edits to correspond to a meaningful reasoning portability test. This enables comprehensive evaluation across five dimensions: data type, real-time property, real-world grounding, reasoning portability, and composite reasoning portability.

6.2 Knowledge Editing Methods

Parameter-based Editing This line of work edits model parameters or adds extra parameters for knowledge editing. [Meng et al. \(2022\)](#), [Meng et al. \(2023\)](#) and [Fang et al. \(2025\)](#) adopt a locate-then-edit manner and aim to precisely edit MLP weights. [Mitchell et al. \(2022a\)](#) adopt meta-learning ([Hospedales et al., 2022](#)) to predict the change of model parameters. [Dong et al. \(2022\)](#), [Huang et al. \(2023\)](#) and [Wang et al. \(2024a\)](#) add extra parameters to store edits in a parametric form.

[Mitchell et al. \(2022b\)](#) train a separate counter-fact model and does not involve retrieval.

Retrieval-based Editing This line of work stores edits in a memory module and retrieve from it when needed. [Zheng et al. \(2023\)](#) leverage in-context examples of copying, updating and retaining knowledge to edit knowledge via prompting. [Chen et al. \(2024b\)](#) first prompt LMs to decide the relevance of retrieved edits and then query LMs with or without in-context edits, which may seem conceptually related to KEDAS but is essentially different in that it relies merely on prompting. [Chen et al. \(2024a\)](#) convert edits into soft prompt tokens with trained models and then concatenate user prompts to modify the behavior of frozen LMs. [Jiang et al. \(2024\)](#) align LMs with knowledge editing by post-training with in-context edits. Our proposed KEDAS can also be categorized as retrieval-based.

7 Conclusion

We have introduced CRAFT, a novel automatically updatable real-time dataset for knowledge editing that focuses on national statistics and finance of China. We evaluate representative knowledge editing methods on CRAFT and reveal that existing methods have certain limitations. We then propose a simple yet effective approach, KEDAS, to better align LMs with real-time knowledge editing. Experiments validate the effectiveness of KEDAS on CRAFT in all aspects including edit success, locality and portability. We hope this work may serve as a catalyst for shifting the focus of knowledge editing from static update to dynamic evolution.

Acknowledgments

This work is supported by the National Natural Science Foundation of China (62076008).

Limitations

Domain CRAFT focuses on the Chinese language only and is limited to the domains of national statistics and finance. Nevertheless, our methodology is universal and can be applied to any other languages and CRAFT is easily expandable to domains that data are periodically updated and automatically accessible.

Portability The composite reasoning portability of CRAFT is merely 2-hop. Nevertheless, it can be expanded to multiple hops via combination of time and indicators.

Baselines Only a limited set of existing methods is evaluated in our experiments due to limited resources and compatibility issues.

Filter The performance of KEDAS heavily depends on the filter. If the filter is not sufficiently robust, the overall performance of KEDAS will also be less competitive.

Ethical Considerations

Computational Budget All our experiments are conducted on a machine with Ubuntu 20.04.6 LTS, Intel[®] Xeon[®] Silver 4310 CPU and 256G memory. We use one NVIDIA A40 48G GPU for all the experiments. The training of KEDAS takes about 2 hours. The editing of KEDAS takes about 50 minutes on both subsets of CRAFT.

Reproducibility Our work is fully reproducible since we have provided our implementation details and scripts for experiments.

Potential Risks To the best of our knowledge, there are no potential risks concerning our work.

Scientific Artifacts We cite all the creators of scientific artifacts we use in this paper. Licenses of these scientific artifacts are shown in Table 8. Our use of these artifacts is consistent with their intended use.

Privacy and Offense Concerns Our data source ensures that our CRAFT dataset contains no information that names or uniquely identifies individual people or offensive content.

Artifact	License
C ³	Non-commercial Purpose
cn-stats	MIT
AKShare	MIT
BERT	Apache
MiniLM	Apache
EasyEdit & KnowEdit	MIT
LLAMA-3	Llama 3 Community License
QWEN2.5	Apache
LLAMA-2	Llama 2 Community License

Table 8: Licenses of scientific artifacts we use.

References

- ByteDance. 2025. [Doubao-1.5-pro-32k model](#).
- Qizhou Chen, Taolin Zhang, Xiaofeng He, Dongyang Li, Chengyu Wang, Longtao Huang, and Hui Xue'. 2024a. Lifelong knowledge editing for LLMs with retrieval-augmented continuous prompt learning. In *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing*.
- Yingfa Chen, Zhengyan Zhang, Xu Han, Chaojun Xiao, Zhiyuan Liu, Chen Chen, Kuai Li, Tao Yang, and Maosong Sun. 2024b. Robust and scalable model editing for large language models. In *Proceedings of the 2024 Joint International Conference on Computational Linguistics, Language Resources and Evaluation (LREC-COLING 2024)*.
- Roi Cohen, Eden Biran, Ori Yoran, Amir Globerson, and Mor Geva. 2024a. [Evaluating the ripple effects of knowledge editing in language models](#). *Transactions of the Association for Computational Linguistics*.
- Roi Cohen, Eden Biran, Ori Yoran, Amir Globerson, and Mor Geva. 2024b. Evaluating the ripple effects of knowledge editing in language models. *Transactions of the Association for Computational Linguistics*.
- DeepSeek-AI. 2025. [Deepseek-v3 technical report](#). Preprint, arXiv:2412.19437.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. BERT: Pre-training of deep bidirectional transformers for language understanding. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*.
- 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*.
- Junfeng Fang, Houcheng Jiang, Kun Wang, Yunshan Ma, Jie Shi, Xiang Wang, Xiangnan He, and Tat-Seng Chua. 2025. Alphaedit: Null-space constrained model editing for language models. In *The Thirteenth*

- International Conference on Learning Representations*.
- Timothy Hospedales, Antreas Antoniou, Paul Micaelli, and Amos Storkey. 2022. Meta-learning in neural networks: A survey. *IEEE Transactions on Pattern Analysis and Machine Intelligence*.
- Edward J Hu, yelong shen, Phillip Wallis, Zeyuan Allen-Zhu, Yuanzhi Li, Shean Wang, Lu Wang, and Weizhu Chen. 2022. LoRA: Low-rank adaptation of large language models. In *International Conference on Learning Representations*.
- Baixiang Huang, Canyu Chen, Xiong Xiao Xu, Ali Payani, and Kai Shu. 2025. [Can knowledge editing really correct hallucinations?](#) In *The Thirteenth 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*.
- Yoichi Ishibashi and Hidetoshi Shimodaira. 2024. [Knowledge sanitization of large language models](#). Preprint, arXiv:2309.11852.
- Albert Q. Jiang, Alexandre Sablayrolles, Arthur Mensch, Chris Bamford, Devendra Singh Chaplot, Diego de las Casas, Florian Bressand, Gianna Lengyel, Guillaume Lample, Lucile Saulnier, L  lio Renard Lavaud, Marie-Anne Lachaux, Pierre Stock, Teven Le Scao, Thibaut Lavril, Thomas Wang, Timoth  e Lacroix, and William El Sayed. 2023. [Mistral 7b](#). Preprint, arXiv:2310.06825.
- Yuxin Jiang, Yufei Wang, Chuhan Wu, Wanjuan Zhong, Xingshan Zeng, Jiahui Gao, Liangyou Li, Xin Jiang, Lifeng Shang, Ruiming Tang, Qun Liu, and Wei Wang. 2024. Learning to edit: Aligning LLMs with knowledge editing. In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*.
- Omer Levy, Minjoon Seo, Eunsol Choi, and Luke Zettlemoyer. 2017. Zero-shot relation extraction via reading comprehension. In *Proceedings of the 21st Conference on Computational Natural Language Learning (CoNLL 2017)*.
- Llama Team. 2024. [The llama 3 herd of models](#). Preprint, arXiv:2407.21783.
- Potsawee Manakul, Adian Liusie, and Mark Gales. 2023. SelfCheckGPT: Zero-resource black-box hallucination detection for generative large language models. In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*.
- Kevin Meng, David Bau, Alex Andonian, and Yonatan Belinkov. 2022. Locating and editing factual associations in gpt. In *Advances in Neural Information Processing Systems*.
- Kevin Meng, Arnab Sen Sharma, Alex J Andonian, Yonatan Belinkov, and David Bau. 2023. Mass-editing memory in a transformer. In *The Eleventh International Conference on Learning Representations*.
- Eric Mitchell, Charles Lin, Antoine Bosselut, Chelsea Finn, and Christopher D Manning. 2022a. Fast model editing at scale. In *International Conference on Learning Representations*.
- Eric Mitchell, Charles Lin, Antoine Bosselut, Christopher D Manning, and Chelsea Finn. 2022b. Memory-based model editing at scale. In *Proceedings of the 39th International Conference on Machine Learning*.
- OpenAI. 2023. Gpt-3.5 technical report. <https://platform.openai.com/docs/models/gpt-3-5>.
- Fabio Petroni, Tim Rockt  schel, Sebastian Riedel, Patrick Lewis, Anton Bakhtin, Yuxiang Wu, and Alexander Miller. 2019. Language models as knowledge bases? In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*.
- Qwen. 2025. [Qwen2.5 technical report](#). Preprint, arXiv:2412.15115.
- Nils Reimers and Iryna Gurevych. 2019. Sentence-bert: Sentence embeddings using siamese bert-networks. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing*.
- Kai Sun, Dian Yu, Dong Yu, and Claire Cardie. 2020. Investigating prior knowledge for challenging chinese machine reading comprehension. *Transactions of the Association for Computational Linguistics*.
- Wei Tang, Yixin Cao, Yang Deng, Jiahao Ying, Bo Wang, Yizhe Yang, Yuyue Zhao, Qi Zhang, Xuanjing Huang, Yu-Gang Jiang, and Yong Liao. 2025. [EvoWiki: Evaluating LLMs on evolving knowledge](#). In *Proceedings of the 63rd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*.
- Lukas Thede, Karsten Roth, Matthias Bethge, Zeynep Akata, and Tom Hartvigsen. 2025. [Wikibigedit: Understanding the limits of lifelong knowledge editing in llms](#). Preprint, arXiv:2503.05683.
- Hugo Touvron, Louis Martin, Kevin Stone, Peter Albert, Amjad Almahairi, Yasmine Babaei, Nikolay Bashlykov, Soumya Batra, Prajjwal Bhargava, Shrutit Bhoale, Dan Bikel, Lukas Blecher, Cristian Canton Ferrer, Moya Chen, Guillem Cucurull, David Esiobu, Jude Fernandes, Jeremy Fu, Wenyin Fu, and 49 others. 2023. [Llama 2: Open foundation and fine-tuned chat models](#). Preprint, arXiv:2307.09288.
- 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. In *Advances in Neural Information Processing Systems*.
- Peng Wang, Ningyu Zhang, Bozhong Tian, Zekun Xi, Yunzhi Yao, Ziwen Xu, Mengru Wang, Shengyu Mao, Xiaohan Wang, Siyuan Cheng, Kangwei Liu, Yuansheng Ni, Guozhou Zheng, and Huajun Chen. 2024b. EasyEdit: An easy-to-use knowledge editing framework for large language models. In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 3: System Demonstrations)*.
- Yunzhi Yao, Peng Wang, Bozhong Tian, Siyuan Cheng, Zhoubo Li, Shumin Deng, Huajun Chen, and Ningyu Zhang. 2023. [Editing large language models: Problems, methods, and opportunities](#). In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*.
- Ningyu Zhang, Yunzhi Yao, Bozhong Tian, Peng Wang, Shumin Deng, Mengru Wang, Zekun Xi, Shengyu Mao, Jintian Zhang, Yuansheng Ni, Siyuan Cheng, Ziwen Xu, Xin Xu, Jia-Chen Gu, Yong Jiang, Pengjun Xie, Fei Huang, Lei Liang, Zhiqiang Zhang, and 3 others. 2024. [A comprehensive study of knowledge editing for large language models](#). *Preprint*, arXiv:2401.01286.
- 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*.
- Yaowei Zheng, Richong Zhang, Junhao Zhang, Yanhan Ye, and Zheyang Luo. 2024. LlamaFactory: Unified efficient fine-tuning of 100+ language models. In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 3: System Demonstrations)*.
- Zexuan Zhong, Zhengxuan Wu, Christopher Manning, Christopher Potts, and Danqi Chen. 2023. MQuAKE: Assessing knowledge editing in language models via multi-hop questions. In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*. Association for Computational Linguistics.

A Details of Alignment Data for KEDAS

For each instance of CRAFT’s training set \mathcal{E}^t , we construct the edit candidate set \mathcal{E}^* by keeping the two edits e_1 and e_2 of the instance and adding 0-3 (randomly decided) extra retrieved edits with the same retriever as LTE that are different from the two to promote model’s robustness to incorrect retrieved results.

To promote edit success, for each edit query q_e and answer a_e of either e_1 or e_2 , the input is **KEPrompt**(\mathcal{E}^* , q_e) and the target is a_e .

To promote portability, for each portability query q_p and answer a_p in \mathcal{E}_p^t , the input is **KEPrompt**(\mathcal{E}^* , q_p) and the target is a_p .

To promote locality, for each locality query q_l and answer a_l in \mathcal{E}_l^t , the in-scope input is **KEPrompt**(\mathcal{E}^* , q_l) with the target a_l . We also include an out-of-scope input q_l with the target a_l that is without any retrieved edits.

In this way, the LM learns to leverage retrieved edits for relevant queries while keeping the answer unchanged for locality queries.

B Details of LM Alignment

Hyperparameter	Value
Batch size	1
Gradient accumulation steps	8
Learning rate	1e-4
Epoches	1
Max length	2048
Optimizer	AdamW
Scheduler	cosine
Warmup ratio	0.1
LoRA rank	8
LoRA alpha	16
LoRA dropout	0

Table 9: Hyper-parameters for LM finetuning.

We adopt LLaMA-Factory (Zheng et al., 2024) to finetune the LM. The alignment is done on an NVIDIA A40 48GB GPU and takes 2 hours.

The main training hyperparameters are presented in Table 9, most remain unchanged as the default values. Note that we train the LM for only one epoch because finetuning for multiple epoches will cause over-fitting.

C Details of Diverse Edit Augmentation

The prompt templates of diverse edit augmentation for gpt-4o-mini are presented in Figure 2 and 3. We include manually written in-context examples in the templates to better instruct the model.

Prompt Template for Declarative

You are a helpful assistant. You are given a query and a target new. Please generate the declaration form of the query and target new.

Here is an example:

Query: 2024年华大智造的总资产（亿元）是多少？

Target new: 103.15

Declaration: 2024年华大智造的总资产（亿元）是103.15。

Please generate the declaration form of the query and target new below:

Query: {query}

Target new: {target_new}

Declaration:

Figure 2: Prompt template for the declarative form.

D Details of Training Data for the Filter

For each instance of CRAFT’s training set \mathcal{E}^t with \mathcal{E}_e^t , \mathcal{E}_p^t and \mathcal{E}_l^t as the set of edits, portability QAs and locality QAs respectively, we simply construct the training data for the filter based on relevance. For any two edits $e = (q_e, a_e)$, $e' = (q_{e'}, a_{e'})$ (may be identical) in \mathcal{E}_e^t , the input is $q_{e'} [\text{sep}] e$ and the target is 1 (relevant). For any edit $e \in \mathcal{E}_e^t$ and any portability QA pair $(q_p, a_p) \in \mathcal{E}_p^t$, the input is $q_p [\text{sep}] e$ and the target is 1 (relevant). For any edit $e \in \mathcal{E}_e^t$ and any locality QA pair $(q_l, a_l) \in \mathcal{E}_l^t$, the input is $q_l [\text{sep}] e$ and the target is 0 (irrelevant). In this way, the filter learns to decide the relevance between the query and the edit.

E Details Regarding KnowEdit

Dataset	ZsRE	WikiBio	WikiData _{recent}	WikiData _{counterfact}
# Train	10000	464	570	1428
# Test	1304	306	1266	885

Table 10: Data statistics of KnowEdit.

Table 10 presents the data statistics of KnowEdit.

Following Zhang et al. (2024) and Jiang et al. (2024), our evaluation metrics for KnowEdit include edit success (**ES**), locality (**L**) and portability (**P**). We also calculate the harmonic mean (**HM**) of the three scores as the edit quality. Note that WikiBio does not contain portability queries and we thus calculate the harmonic mean based on the first two metrics.

F Other Baselines

In our primary experiments, we also tried to evaluate many other modern baselines but encountered different sorts of issues (including compatibility, availability, cost, performance, etc.).

Prompt Template for *Aliased*

You are a helpful assistant. You are given a query, a target new, and a declaration. Please generate a paraphrased sentence of the declaration with the central term translated to English.

Here is an example:

Query: 2024年华大智造的总资产(亿元)是多少?

Target new: 103.15

Declaration: 2024年华大智造的总资产(亿元)是103.15。

Paraphrased sentence: 2024年华大智造的Total Assets(100 million yuan)是103.15。

Here is another example:

Query: 2024年9月中国的订销报纸份数当期值(万份)是多少?

Target new: 137885.2

Declaration: 2024年9月中国的订销报纸份数当期值(万份)是137885.2。

Paraphrased sentence: 2024年9月中国的Issue of Newspapers, Current Period Value(10000 pieces)是137885.2。

Please generate a paraphrased sentence of the declaration below:

Query: {query}

Target new: {target_new}

Declaration: {declaration}

Paraphrased sentence:

Figure 3: Prompt template for the *aliased* form.

When running MEMIT (Meng et al., 2023), our GPU ran out of memory. Meanwhile, it performed poorly in all the metrics on some datasets (e.g., merely 1.91 and 0.75 edit success on WikiData_{recent} and WikiData_{counterfact}). Since MEMIT is primarily designed for single editing, we attribute the poor performance to its incompatibility with the sequential editing setting.

AlphaEdit (Fang et al., 2025), a more advanced method designed on the basis of MEMIT, also exhibited relatively poor performance (for instance, its edit success scores on WikiData_{recent} and WikiData_{counterfact} were merely 20.60 and 5.91 respectively) so that we do not include it as a baseline for comparison. We believe this is due to its incompatibility with the sequential editing setting or the LMs we use for experiments.

MEND (Mitchell et al., 2022a), SERAC (Mitchell et al., 2022b) and RECIPE (Chen et al., 2024a) require a lot of resources that we cannot afford currently for training and do not provide trained checkpoints for the LLAMA-3 herd of LMs used in our experiments. Therefore, we do not include these in our main experiments.

Nevertheless, for a comprehensive comparison, we experiment on KnowEdit with all the mentioned baselines under the single editing setting using Llama-2-7B-Chat (Touvron et al., 2023) due to that all the methods are compatible with single editing and the LLAMA-2 herd of LMs. The results are presented in Appendix I.

G Example from CRAFT-Statistics

- **Prompt 1:** 2024年7月中国的货币和准货币(M2)供应量期末值(亿元)是多少? (English translation: What is the period-end value (100 million yuan) of China's Money and Quasi-Money (M2) supply in July 2024?)
Target New: 3033060.78
- **Prompt 2:** 2023年7月中国的货币和准货币(M2)供应量期末值(亿元)是多少? (English translation: What is the period-end value (100 million yuan) of China's Money and Quasi-Money (M2) supply in July 2023?)
Target New: 2854031.56

Portability

- **Alias:**
 - Prompt: 2024年7月中国的Money and Quasi-Money (M2) Supply, period-end(100 million yuan)是多少? (English translation: What is the Money and Quasi-Money (M2) Supply, period-end (100 million yuan), in July 2024 in China?) Answer: 3033060.78
 - Prompt: 2023年7月中国的Money and Quasi-Money (M2) Supply, period-end(100 million yuan)是多少? (English translation: What is the Money and Quasi-Money (M2) Supply, period-end (100 million yuan), in July 2023 in China?) Answer: 2854031.56
- **Composite Reasoning:** Prompt: 2024年7月中国的货币和准货币(M2)供应量期末值(亿元)比2023年7月的高多少? (English translation: How much higher was China's M2 supply (100 million yuan) in July 2024 than in July 2023?) Answer: 179029.22

Locality

- **Temporal:**
 - Prompt: 2022年7月中国的货币和准货币(M2)供应量期末值(亿元)是? (English translation: What was the M2 supply (100 million yuan) at the end of July 2022 in China?)
Answer: 2578078.57
 - Prompt: 2025年7月中国的货币和准货币(M2)供应量期末值(亿元)是? (English translation: What is the M2 supply

(100 million yuan) at the end of July 2025 in China?) Answer: 3299429.06

- **Common-sense:**

- Prompt: 请根据以下材料回答问题。(略) 问题: 动物的器官感觉与人的相比有什么不同? (English translation: Read the following passage and answer the question. Question: How do animals' sensory organs differ from those of humans?) Answer: D
- Prompt: 请根据以下材料回答问题。(略) 问题: 录音中提到能预报风暴的动物是什么? (English translation: Read the following passage and answer the question. Question: Which animal mentioned in the passage can predict storms?) Answer: C

H Example from CRAFT-Finance

- **Prompt 1:** 2024年三友医疗的总资产(亿元)是多少? (English translation: What is the total assets (100 million yuan) of Sanyou Medical in 2024?) **Target New:** 23.07
- **Prompt 2:** 2024年三友医疗的总负债(亿元)是多少? (English translation: What is the total liabilities (100 million yuan) of Sanyou Medical in 2024?) **Target New:** 2.58

Portability

- **Alias:**

- Prompt: 2024年三友医疗的Total Assets(100 million yuan)是多少? (English translation: What is the total assets (100 million yuan) of Sanyou Medical in 2024?) Answer: 23.07
- Prompt: 2024年三友医疗的Total Liabilities(100 million yuan)是多少? (English translation: What is the total liabilities (100 million yuan) of Sanyou Medical in 2024?) Answer: 2.58

- **Composite Reasoning:** Prompt: 2024年三友医疗的资产负债率(%)是多少? (English translation: What is the debt-to-asset ratio (%) of Sanyou Medical in 2024?)

Answer: 11.2

Locality

- **Temporal:**

- Prompt: 2023年三友医疗的总资产(亿元)是多少? (English translation: What was the total assets (100 million yuan) of Sanyou Medical in 2023?) Answer: 22.61
- Prompt: 2023年三友医疗的总负债(亿元)是多少? (English translation: What was the total liabilities (100 million yuan) of Sanyou Medical in 2023?) Answer: 2.19

- **Common-sense:**

- Prompt: 请根据以下材料回答问题。(略) 问题: 小羊的结局是什么? (English translation: Read the following passage and answer the question. Question: What happened to the little lamb?) Answer: C
- Prompt: 请根据以下材料回答问题。(略) 问题: 这个故事告诉我们什么道理? (English translation: Read the following passage and answer the question. Question: What moral lesson does this story convey?) Answer: A

I More Results on KnowEdit

The sequential editing results with Llama-2-7B-Chat (Touvron et al., 2023) and Qwen2.5-7B-Instruct (Qwen, 2025) are presented in Table 11. KEDAS secures the highest harmonic mean scores on all the four datasets, outperforming previous methods significantly.

We also experiment with Llama-2-7B-Chat under the traditional single editing setting and the results are presented in Table 12. KEDAS exhibits the most balanced performance among all the 11 evaluated methods.

Model	Method	ZsRE				WikiBio			WikiData _{recent}				WikiData _{counteract}			
		ES	L	P	HM	ES	L	HM	ES	L	P	HM	ES	L	P	HM
LLAMA-2	SERAC	94.9	44.9	40.6	52.2	98.4	70.1	81.9	92.0	33.6	47.0	48.4	78.0	29.3	42.7	42.6
	IKE	99.5	46.9	67.6	65.0	99.7	41.6	58.7	85.0	45.8	57.3	58.8	85.7	51.5	60.6	63.0
	EREN	36.7	67.1	47.8	47.6	60.8	64.0	62.4	44.4	59.8	38.5	46.0	25.6	67.8	25.4	32.2
	WISE	71.1	67.5	51.1	61.9	95.5	99.2	97.3	80.1	48.9	53.2	58.0	72.9	56.6	50.7	58.7
	RECIPE	96.4	60.0	52.6	65.1	63.8	61.9	62.8	80.8	61.2	51.4	62.2	55.6	66.8	36.8	49.9
	LTE	99.6	49.2	69.2	66.9	100.0	47.2	64.1	100.0	55.0	74.6	72.1	98.4	63.6	73.6	76.0
	KEDAS	99.6	91.0	73.6	86.7	100.0	100.0	100.0	100.0	77.6	75.9	83.1	98.4	76.3	77.1	82.8
QWEN2.5	IKE	96.7	42.2	65.1	60.7	99.3	33.4	50.0	93.6	40.7	61.7	58.3	86.2	48.4	59.3	61.1
	EREN	34.4	60.7	46.0	44.6	58.1	53.0	55.4	42.1	55.6	37.4	43.8	21.5	62.4	22.0	27.8
	WISE	53.0	14.7	25.1	23.7	86.6	100.0	92.8	54.0	96.8	33.0	50.7	45.4	17.9	30.3	27.1
	LTE	99.1	48.0	65.4	64.9	98.4	38.7	55.5	99.7	42.5	68.9	62.4	98.0	49.4	69.9	67.0
	KEDAS	99.1	89.6	70.2	84.5	98.4	100.0	99.2	99.7	67.1	70.5	76.7	98.0	66.6	73.4	77.2

Table 11: Main results of sequential editing. Metrics include edit success (ES), locality (L), portability (P) and harmonic mean (HM). The highest scores are **bolded**.

Method	ZsRE				WikiBio			WikiData _{recent}				WikiData _{counteract}			
	ES	L	P	HM	ES	L	HM	ES	L	P	HM	ES	L	P	HM
LoRA	100.0	75.7	56.5	73.3	100.0	81.7	89.9	100.0	56.3	64.5	69.3	100.0	70.7	68.9	77.6
MEND	97.9	96.4	61.1	81.2	58.3	90.8	71.0	94.4	77.9	43.3	64.5	-	-	-	-
ROME	96.7	53.7	52.7	62.5	96.0	62.7	75.8	97.0	54.7	55.7	64.5	98.5	50.4	55.8	62.6
SERAC	97.0	60.0	61.3	69.3	98.4	72.8	83.7	99.0	36.4	64.9	56.6	99.0	32.4	76.7	55.6
IKE	99.8	47.1	78.1	68.1	99.7	41.3	58.4	85.0	45.9	62.6	60.6	87.3	52.0	72.2	67.4
EREN	36.7	67.1	47.8	47.6	60.8	64.0	62.4	44.4	59.8	38.5	46.0	25.6	67.8	25.4	32.2
WISE	99.9	99.7	47.4	72.9	100.0	100.0	100.0	98.3	82.7	62.2	78.2	99.6	82.3	71.1	82.8
AlphaEdit	97.5	66.9	51.2	67.0	99.0	71.8	83.2	97.8	60.8	52.3	65.5	98.9	64.1	47.7	64.3
RECIPE	96.6	78.3	55.9	73.1	63.8	94.6	76.2	81.0	71.1	53.2	66.3	55.7	76.9	38.3	52.6
LTE	99.9	49.2	77.6	69.4	100.0	47.2	64.1	100.0	55.0	82.0	74.3	100.0	65.0	88.0	81.6
KEDAS	99.6	96.3	77.3	89.9	100.0	100.0	100.0	100.0	98.1	81.1	92.2	98.4	99.4	85.9	94.1

Table 12: Results of single editing on KnowEdit with Llama-2-7B-Chat. The highest scores are **bolded**. “-” means the method encounters out-of-memory (OOM) or tensor overflow on our GPU.