

InfuserKI: Enhancing Large Language Models with Knowledge Graphs via Infuser-Guided Knowledge Integration

Fali Wang
Pennsylvania State University
University Park, USA
fqw5095@psu.edu

Runxue Bao
GE Healthcare
Bellevue, USA
runxue.bao@gehealthcare.com

Suhang Wang
Pennsylvania State University
University Park, USA
szw494@psu.edu

Wenchao Yu, Yanchi Liu, Wei Cheng, Haifeng Chen
NEC Laboratories America, Princeton, USA
{wyu, yanchi, weicheng, haifeng}@nec-labs.com

Abstract

Large Language Models (LLMs) have achieved exceptional capabilities in open generation across various domains, yet they encounter difficulties with tasks that require intensive knowledge. To address these challenges, methods for integrating knowledge have been developed, which augment LLMs with domain-specific knowledge graphs through external modules. These approaches, however, face data inefficiency issues as they necessitate the processing of both known and unknown knowledge for fine-tuning. Thus, our research focuses on a novel problem: efficiently integrating unknown knowledge into LLMs without unnecessary overlap of known knowledge. A risk of introducing new knowledge is the potential forgetting of existing knowledge. To mitigate this risk, we propose the innovative InfuserKI framework. This framework employs transformer internal states to determine when to enrich LLM outputs with additional information, effectively preventing knowledge forgetting. Performance evaluations using the UMLS-2.5k and MetaQA domain knowledge graphs reveal that InfuserKI not only successfully integrates new knowledge but also outperforms state-of-the-art baselines, reducing knowledge forgetting by 9% and 6%, respectively.

1 Introduction

Large Language Models (LLMs) have significantly advanced the capabilities of various language tasks, including Question Answering (QA), dialogue, and information retrieval, showcasing impressive performance across different fields (Touvron et al., 2023a,b). However, in knowledge-intensive tasks like open-domain QA, LLMs can produce texts that are misleading or inaccurate due to a lack of domain knowledge and the phenomenon of catastrophic forgetting post-fine-tuning (Kwiatkowski et al., 2019; Zhai et al., 2024; Li et al., 2022). The step of updating and customizing LLMs with *domain knowledge integration* is thus highly valued

for enhancing their application. This could involve companies customizing models with specialized product knowledge, or hospitals adapting models to reflect specific case data.

Knowledge Graphs (KGs) are ideal sources for bolstering domain-specific knowledge, thanks to their structured and measurable knowledge units. Various strategies have been devised to utilize this knowledge effectively. Typically, these strategies encompass instruction tuning of LLMs using explanations of knowledge entities (Wu et al., 2023), developing triplet-based pre-training tasks (Zhang et al., 2022; Qin et al., 2021; Wang et al., 2021), using KGs as external sources in retrieval tasks (Sridhar and Yang, 2022; Yu et al., 2022), and applying parameter-efficient fine-tuning (PEFT) techniques such as LoRA (Hu et al., 2021) and adapters (Houlsby et al., 2019), or model editing (ME) methods like T-Patcher (Huang et al., 2023) to implement knowledge in a triplet-to-text format (Meng et al., 2021; Emelin et al., 2022; Dong et al., 2022). However, pre-training or fine-tuning LLMs with the entire KGs is not only time-consuming but also leads to data inefficiencies, especially when models relearn knowledge they already have. To address this issue, we focus on integrating new, previously unknown knowledge only. This precise focus, however, introduces the risk of catastrophic forgetting, where the addition of new knowledge may affect existing knowledge. Fig. 1 illustrates a comparison between a standard LLM and its fine-tuned variant by visualizing the internal states of the 10th transformer layer from the training data using the TSNE tool, where each UMLS knowledge unit sample is processed to obtain these states and then mapped to two dimensions for display. Fig. 1 (a) and (b) demonstrate how direct fine-tuning can lead to the loss of previously known data, while Fig. 1 (c) illustrates the ideal integration of new knowledge without compromising existing information. Thus, we pose a novel research question: **How**




Figure 1: An illustrative comparison among (a) Vanilla LLM, (b) Fine-Tuned LLM, and (c) our Knowledge-Infused LLM.

can we efficiently integrate new knowledge from domain-specific KGs into LLMs while preventing catastrophic forgetting?

In this work, we introduce the Infuser-guided Knowledge Integration (**InfuserKI**) framework, meticulously designed to integrate domain-specific knowledge from KGs into LLMs. Drawing inspiration from [Azaria and Mitchell \(2023\)](#), which reveals that an LLM’s internal states can reflect the truthfulness of its generated texts, our framework incorporates an infusing mechanism that verifies the presence of current knowledge in LLMs. This mechanism facilitates the adaptive selection of additional information for both known and unknown knowledge, effectively minimizing the impact on existing knowledge and preventing knowledge forgetting. Additionally, InfuserKI employs knowledge adapters to embed new knowledge while maintaining the integrity of the original model parameters. The process within the InfuserKI framework initiates by identifying knowledge that LLMs do not yet know. Following methodologies from [Zhao et al. \(2023\)](#) and [Seyler et al. \(2017\)](#), we craft a *knowledge statement* and *multiple-choice questions* for a knowledge triplet $\langle h, r, t \rangle$ using established relational templates, as illustrated in Fig. 3. Furthermore, to broaden the generality of the integrated knowledge, InfuserKI implements a relation classification task. This task is designed to refine the linguistic representations developed by the adapters, enabling the prediction of relationships within knowledge statements based on the adapter outputs for head and tail entities. This approach not only ensures a solid integration of new knowledge but also bolsters the framework’s ability to generalize this knowledge to unseen scenarios.

Our main contributions are:

- We explore a novel problem: effectively integrating unknown knowledge from KGs into LLMs without impacting existing knowledge.
- We introduce a new knowledge integration framework, InfuserKI, which facilitates the adaptive se-

lection of known and unknown knowledge for integration into LLMs, effectively reducing knowledge forgetting.

- Comprehensive evaluations on the UMLS and MetaQA datasets demonstrate that InfuserKI achieves effective knowledge integration with less forgetting, maintains performance on large-scale data, and offers enhanced generality across unseen templates and downstream tasks.

2 Related Work

Knowledge Integration LLMs often produce seemingly accurate but incorrect answers due to missing knowledge. Addressing this, knowledge integration (KI) into LLMs has become popular. KGs, which capture wide or domain-specific knowledge, serve as an ideal option due to their structured and quantifiable knowledge units. KI from KGs usually occurs during pre-training or fine-tuning. For example, ERNIE ([Sun et al., 2019](#)) injects KG’s embeddings, such as TransE ([Fan et al., 2014](#)), into models using an entity-token alignment masking loss. However, retraining is time-consuming. In fine-tuning, methods including JointLK ([Sun et al., 2022](#)) and GreaseLM ([Zhang et al., 2021](#)) apply graph neural networks to model knowledge subgraphs, relying on KGs until inference. Fully fine-tuning models such as PMC-LLaMa ([Wu et al., 2023](#)) is computationally costly; therefore PEFT methods, for instance, LoRA ([Hu et al., 2021](#)) and Adapters ([Houlsby et al., 2019](#)), are more feasible. Based on these works, MoP ([Meng et al., 2021](#)), K-Adapter ([Wang et al., 2021](#)), and KB-adapters ([Emelin et al., 2022](#)) inject knowledge directly into model parameters but risk catastrophic forgetting of unrelated knowledge ([Meng et al., 2022b](#)). Thus, we focus on adapter-based integration that minimizes the impact on unrelated knowledge.

Model Editing Model Editing (ME) for LLMs falls into two categories: gradient-based and extension-based. Gradient-based methods, as described by [Dai et al. \(2022\)](#), modify specific weights related to knowledge edits. ROME ([Meng et al., 2022a](#)) and MEMIT ([Meng et al., 2022b](#)) take this further by updating entire Feedforward Network (FFN) layers to enhance model editing. These methods, however, are limited in the number of edits or may require considerable time for execution. On the other hand, extension-based methods add new parameters to correct inaccurate

information. CALINET (Dong et al., 2022) and T-Patcher (Huang et al., 2023) incorporate memory slots or trainable "patches" into final FFN outputs. GRACE (Hartvigsen et al., 2023) employs a key-value adapter with a deferral mechanism for the selective use of knowledge based on input. However, the adapter-based modules positioned in top transformer layers are designed to calibrate false facts. Instead, our method aims to infuse new knowledge by placing adapters throughout transformer layers.

Catastrophic Forgetting Catastrophic forgetting occurs when learning new information causes a drastic loss of previously learned knowledge (Ratcliff, 1990). This phenomenon is particularly evident in sequential inter-task learning, where acquiring new task knowledge can lead to forgetting older task knowledge (McCloskey and Cohen, 1989). To address this, various strategies have been developed. Xuhong et al. (2018) applied constraint to minimize parameter changes during new task learning. Elastic Weight Consolidation (EWC) incorporates the Hessian matrix into parameter regularization to reduce forgetting (Kirkpatrick et al., 2017). Replay-based methods, including sampling strategies that retain original training samples in a memory buffer (Lopez-Paz and Ranzato, 2017). Knowledge Distillation aligns the predictions of a fine-tuned model with the pre-fine-tuning model (Buzzega et al., 2020). Parameter-Efficient Fine-Tuning can also mitigate forgetting, represented by LoRA (Hu et al., 2021), which uses low-rank matrices for weight modifications while maintaining pre-trained parameters frozen, and achieves results akin to full fine-tuning. However, these studies emphasize sequential inter-task transfer learning. Our focus shifts to intra-task knowledge forgetting, where integrating new knowledge leads to the potential loss of previously existing knowledge.

3 Proposed Framework - InfuserKI

The objective of our method is to leverage domain knowledge from KGs to enhance LLMs for knowledge-intensive tasks. Specifically, given an LLM $p_\theta \in \mathbb{P}$ and a set of knowledge triplets $\mathcal{T} \in \mathbb{T}$, our goal is to fine-tune the LLM p_θ into p'_θ , incorporating previously unknown knowledge \mathcal{T}_{unk} without affecting existing knowledge \mathcal{T}_{known} . For efficiency, we only inject knowledge that is unknown to the LLM as:

$$\mathbb{F}_{KI} : \mathbb{P} \times \mathbb{T} \rightarrow \mathbb{P} \quad p'_\theta = f_{KI}(p_\theta, \mathcal{T}_{unk})$$




Figure 2: Knowledge Areas in LLMs: Original ($\mathcal{N}_1+\mathcal{N}_2$), Post-Fine-Tuning ($\mathcal{N}_1+\mathcal{N}_3$), Forgotten (\mathcal{N}_2), and Failed Integration (\mathcal{N}_4).

The core design of our InfuserKI framework comprises two steps: knowledge detection and knowledge integration, as illustrated in Fig. 3. To be specific, we first detect previously unknown knowledge by feeding questions derived from knowledge triplets to the LLMs. Upon identifying a set of unknown knowledge, we employ the knowledge adapter, which is parallel to the original transformer layer and trained to store new knowledge. The core of our framework, the *knowledge Infuser*, is designed to strategically determine whether new knowledge from the knowledge adapter should be engaged. Throughout this process, we only fine-tune the knowledge adapter and the Infuser while keeping the original transformer parameters fixed.

3.1 Knowledge Detection

Given the inefficiency of fine-tuning LLMs on entire graphs, we aim to identify and integrate only the LLMs' unknown knowledge. To overcome the difficulty of evaluating open-ended questions, we convert triplets into multiple-choice questions (Manakul et al., 2023), allowing for a precise assessment of LLMs' initial unknown knowledge ($\mathcal{N}_3 + \mathcal{N}_4$ in Fig. 2). This strategy enables efficient knowledge integration, using multiple-choice training data to enhance domain-specific performance.

Multiple-choice Question Generation Given a knowledge triplet, it is transformed into multiple-choice questions and a knowledge statement using relation templates generated by GPT-4. For instance, the triplet $\langle \textit{Sutura cranii}, \textit{has finding site}, \textit{Acrocephalosyndactyly type 5} \rangle$ is rephrased into the question with golden answer as "What diagnosis is associated with the finding site of *Sutura cranii*? Answer: *Acrocephalosyndactyly type 5*," along with a knowledge statement as "The finding site for *Sutura cranii* is associated with *Acrocephalosyndactyly type 5*." The prompt for generating templates and knowledge evaluation method are detailed in Appendix A.1.

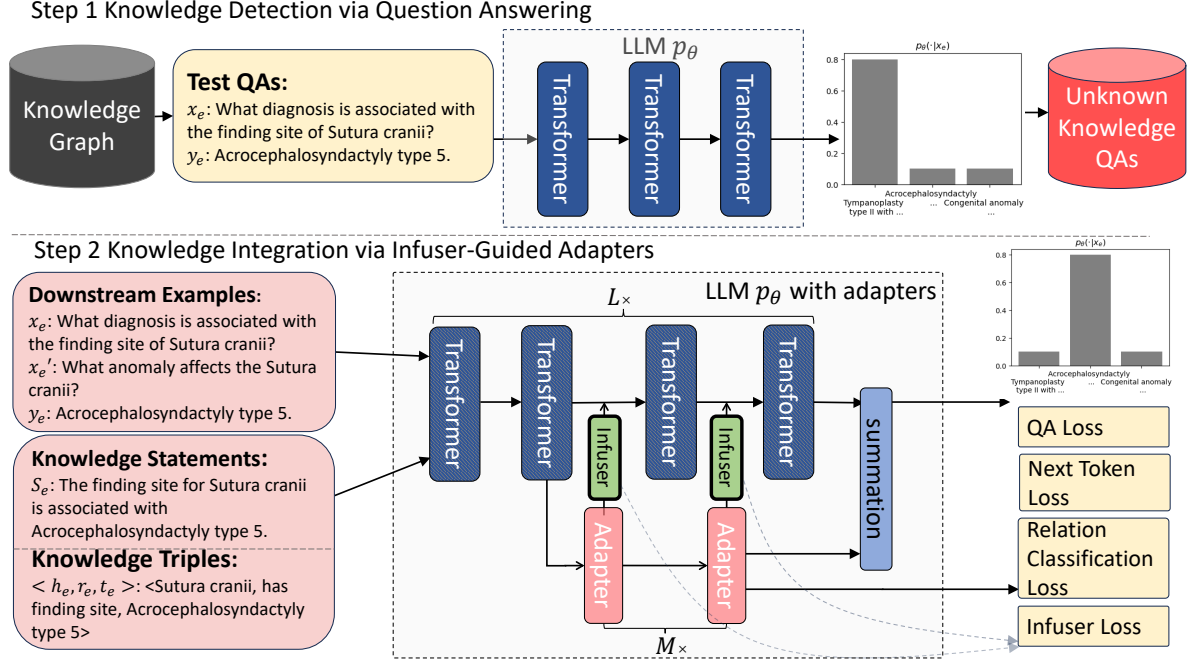


Figure 3: Infuser-Guided Knowledge Integration Framework.

Unknown Knowledge Detection With multiple-choice questions, we input them into LLMs. The testing prompts are in Table 8 in Appendix. We use regular expressions to extract the chosen options from the output of LLMs, treating the response as incorrect if no options can be extracted. This helps us detect the LLMs’ known and unknown knowledge. As shown in Fig. 2, the regions labeled \mathcal{N}_1 and \mathcal{N}_2 represent the set of known knowledge, denoted as \mathcal{T}_{known} , while the regions labeled \mathcal{N}_3 and \mathcal{N}_4 represent the set of unknown knowledge, as \mathcal{T}_{unk} . We then develop a new method to integrate these unknown knowledge into the LLMs without affecting existing knowledge.

3.2 Infuser-Guided Knowledge Integration

Next, we detail our Infuser-guided Knowledge Integration method that effectively and efficiently injects unknown knowledge of LLMs.

Knowledge Adapter To improve parameter efficiency, we use parallel adapters as extra modules to learn new knowledge, keeping the original LLM parameters unchanged, as shown in Fig. 4. Existing works (Dai et al., 2022; Geva et al., 2021) show that Feed-Forward Network (FFN) layers in transformer-based language models store knowledge effectively. Thus, we add adapters parallel to the last M FFN layers for the entire L layers. For the l -th selected adapter layer where $l \in [L - M + 1, L]$, we combine the FFN input

$\mathbf{H}_P^l \in \mathbb{R}^{n \times d}$ with the output \mathbf{H}_A^{l-1} from the previous adapter layer as:

$$\tilde{\mathbf{H}}_A^l = \mathbf{H}_A^{l-1} + \mathbf{H}_P^l \quad (1)$$

where n is the length of the LLM input sequence, and d is the hidden dimension. The initial \mathbf{H}_A^{L-M} is set to a vector of all zeros. Following He et al. (2022), the adapter layer utilizes a down-projection with $\mathbf{W}_{\text{down}} \in \mathbb{R}^{d \times d'}$ to transform the combined input $\tilde{\mathbf{H}}_A^l$ into a lower-dimensional space specified by the bottleneck dimension d' so as to facilitate the learning of new patterns with minimal extra space. This is followed by a nonlinear activation function σ , and subsequently, an up-projection is applied with $\mathbf{W}_{\text{up}} \in \mathbb{R}^{d' \times d}$ as:

$$\mathbf{H}_A^l = \sigma(\tilde{\mathbf{H}}_A^l \mathbf{W}_{\text{down}}) \mathbf{W}_{\text{up}} \quad (2)$$

Typically, the adapter output directly merges with the original output from the FFN as follows:

$$\mathbf{H}_O^l = \mathbf{H}_A^l + \text{FFN}(\mathbf{H}_P^l) \quad (3)$$

\mathbf{H}_O^l is then fed into either the next transformer attention layer or the final linear and softmax layer. However, this approach can *overload the LLM with unnecessary information about knowledge it already knows*, causing the forgetting issue.

Knowledge Infuser To ensure that these extra modules do not confuse the LLM about its existing knowledge, we propose an Infuser model to more

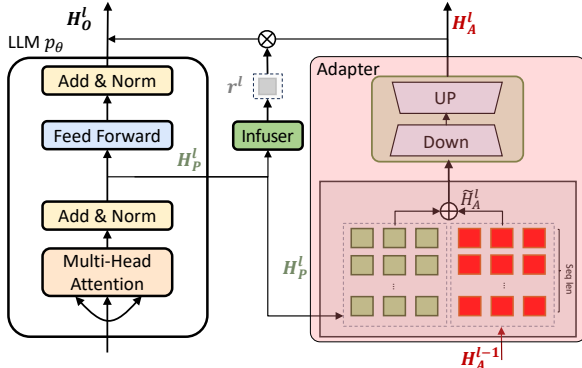


Figure 4: Infuser-Guided Knowledge Adapters.

effectively infuse the knowledge from the knowledge adapter to the LLM. Intuitively, for a given question, the Infuser assesses if the LLM knows the knowledge at hand. If not, the Infuser can fuse more knowledge from \mathbf{H}_A^l to LLM to provide extra information. If the LLM already knows, \mathbf{H}_A^l should have less impact. Recent work (Azaria and Mitchell, 2023) indicates that checking the LLM’s internal states can determine if it knows the current question, which paves us a way to design the Infuser. Specifically, we derive an infusing score from the input of an FFN sublayer as follows:

$$r^l = f_{In}(\text{Mean}(\mathbf{H}_P^l)) \quad (4)$$

where f_{In} denotes the Infuser module implemented as a multilayer perceptron (MLP) with a sigmoid activation function and the Mean function averages the vector along the sequence length. This allows infusing score r^l to be mapped to the range $[0, 1]$, indicating how well the LLMs know about the knowledge based on their intermediate states in the l -th FFN layer (\mathbf{H}_P^l). As a result, the infusing mechanism helps LLMs learn new knowledge without forgetting what they already know. However, it is difficult for the Infuser to recognize existing knowledge if it only encounters new knowledge during fine-tuning. To fix this, we also include a modest quantity of samples representing knowledge the LLMs already have. Before fine-tuning, we first pre-train the Infuser on a binary infusing task with a balanced mix of known and unknown samples. The Infuser loss is a binary cross-entropy loss function as:

$$\mathcal{L}_{In} = \mathbb{E}_{x, y_{In}} [\text{BCE}(f_{In}(\mathbf{H}_P^l), y_{In})] \quad (5)$$

where x is the sample and the infusing label y_{In} is 1 for new knowledge and 0 for previously acquired knowledge. Finally, we obtain an additive filtered

adapter vector, which is integrated with the original FFN output:

$$\mathbf{H}_O^l = r^l \mathbf{H}_A^l + \text{FFN}(\mathbf{H}_P^l), \quad (6)$$

which can selectively incorporate knowledge from the adapter into the fixed base model.

Objective Function of InfuserKI We employ unknown knowledge identified during the knowledge detection phase to fine-tune both the knowledge adapter and the Infuser. The InfuserKI framework is divided into three phases: Infuser tuning, QA (Question Answering) training, and RC (Relation Classification) training, as illustrated by the following objective function:

$$\mathcal{L} = \begin{cases} \mathcal{L}_{In}, & \text{Infuser Tuning} \\ \mathcal{L}_{QA}, & \text{QA Training} \\ \mathcal{L}_{NTL} + \lambda_{RC} \mathcal{L}_{RC}, & \text{RC Training.} \end{cases} \quad (7)$$

In terms of QA training, we use question-based instructions with standard answers as golden responses. The QA loss is akin to the conventional training loss used in transformer-based language models, tailored to adapt instructions within a specific domain:

$$\mathcal{L}_{QA} = \mathbb{E}_{x, y} \left[\frac{1}{|y|} \sum_{i=1}^{|y|} \text{CE}(p_\theta(\cdot | x, y_1, \dots, y_{i-1}), y_i) \right] \quad (8)$$

where $\text{CE}(\cdot, \cdot)$ denotes the cross-entropy loss function, $y = y_1, \dots$, is the golden output, and $p_\theta(\cdot | x, y_1, \dots, y_{i-1})$ is the prediction of an LLM. Note that we also incorporate a small set of yes/no QA samples to enhance the model generality to various question types.

To boost generality of InfuserKI, we adopt a relation classification task, following Zhao et al. (2023), to enhance our knowledge adapters’ understanding of relational facts. For a given knowledge statement k and its triplet $\langle h, r, t \rangle$, we perform mean pooling on the adapter output \mathbf{H}_A^L for the entity mentions, obtaining representations v^h and v^t . Following Qin et al. (2021), we form a relational representation $v^r = [v^h, v^t]$, treating r as a positive sample and other relations as negatives. The relation classification (RC) loss, employing the InfoNCE loss (Oord et al., 2018), aims to distinguish positive relations from negatives, as shown below:

$$\mathcal{L}_{RC} = \mathbb{E}_k \left[-\log \frac{\exp(f_1^R(v^r) \cdot f_2^R(r)/\tau)}{\sum_{r' \in \mathcal{E}} \exp(f_1^R(v^r) \cdot f_2^R(r')/\tau)} \right] \quad (9)$$

where τ acts as a temperature hyperparameter. The functions f_1^R and f_2^R align entity and relation embeddings into a unified dimensional space, respectively, with \mathcal{E} denoting the complete set of relations. Besides that, we also adopt the conventional training loss (i.e. next token loss) used in transformer models:

$$\mathcal{L}_{NTL} = \mathbb{E}_k \left[\frac{1}{|k|} \sum_{i=1}^{|k|} \text{CE}(P_{\theta}(k_i | k_1, \dots, k_{i-1})) \right] \quad (10)$$

The training algorithm is detailed in Appendix A.2. To be specific, given an LLM p_{θ} and a KG with knowledge triplets $\langle h, r, t \rangle$, we generate question-based instructions q , standard answers y , and knowledge statements k . The training is divided into three stages. Initially, we tune the Infuser using a small set of balanced samples of known and unknown, as per Eq. 5. In the second stage, we fine-tune the model using a QA loss to integrate unknown knowledge, following Eq. 8. In the final stage, we use knowledge statements and triplets to enhance the model generality, according to Eq. 9 and 10.

4 Experiments

In this section, we evaluate the proposed framework by conducting experiments on two knowledge graphs across different data scales, comparing against PEFT and ME baselines.

4.1 Experimental Setup

We evaluate our InfuserKI framework with competitive baselines on two domain KGs and their corresponding downstream tasks in terms of reliability, locality, and generality.

Datasets We conduct experiments on a medical KG UMLS (Bodenreider, 2004) with PubMedQA (Jin et al., 2019) and a movie KG MetaQA (Zhang et al., 2018) with MetaQA-1HopQA as the downstream task respectively. The detailed description is in Appendix A.3.

Metrics Following Huang et al. (2023) (see Appendix A.4), as shown in Fig. 2 with areas for various knowledge dynamics, we use the following metrics: (1) **Newly-learned Rate (NR)** for reliability, calculated by $NR = \mathbb{E}_{x \in \mathcal{N}_3 + \mathcal{N}_4} [p_{known}(x)]$ with $p_{known}(x) = 1$ for correct answers and 0 for incorrect ones; (2) **Remembering Rate (RR)** for locality, defined as $RR = \mathbb{E}_{x \in \mathcal{N}_1 + \mathcal{N}_2} [p_{known}(x)]$;

(3) **F1_T1 and F1_T2** for seen templates to assess reliability and locality and **F1_T3 to F1_T5** for unseen templates, with their average, denoted as **F1_Unseen**, serving to assess generality; and (4) **Downstream-Task F1** for the effectiveness of knowledge integration on downstream tasks.

Baselines We compare InfuserKI against both PEFT methods and ME techniques. The PEFT baselines include: (i) **Prefix Tuning** (Li and Liang, 2021) employs learnable prompts in input or intermediate layers; (ii) **LoRA** (Hu et al., 2021) uses trainable low-rank matrices for self-attention weights while freezing other parameters; (iii) **QLoRA** (Dettmers et al., 2023) quantizes pre-trained models to 4 bits based on LoRA. All PEFT methods are tested with the same mix of unknown and known samples to ensure fairness. The adopted **Knowledge Model Editing Methods** are: (i) **CALINET** (Dong et al., 2022) corrects false knowledge by fine-tuning an adapter in a specific FFN layer while keeping original model parameters intact; (ii) **T-Patcher** (Huang et al., 2023) adds a few trainable neurons to the last FFN layer for error correction.

Experimental Details We use LLaMa-2-7B (Touvron et al., 2023a) as our base LLM. Following MoP (Meng et al., 2021), we sample parts of the KG (2,500 and 25,000 triplets for UMLS, and 2,900 for MetaQA) in our experiments. During fine-tuning, we set the dimensionality d' to 10, and positioned the adapters in the last 30 layers out of 32. The RC loss temperature is set at $\tau = 0.7$. Our approach adds approximately 2.5M extra parameters. Using the AdamW optimizer (Loshchilov and Hutter, 2018) with a batch size of 8 and a learning rate of $1 \times e^{-4}$, training takes about 30 minutes per epoch for UMLS 2.5k and MetaQA, and 4 hours for UMLS 25k on $4 \times A100$ GPU servers. We adjust loss weights with $\lambda_{RC} = 10$. The PEFT baselines are implemented following LLaMa-Adapter (Zhang et al., 2023) and PEFT (Mangrulkar et al., 2022).

4.2 Results and Analysis

Table 1 and 2 show a comparison of our InfuserKI against existing PEFT and ME methods on the UMLS and MetaQA with 2,500 and 2,900 triplets respectively. We can observe: (1) The performance of Vanilla LLaMa-2-7B underscores a lack of domain-specific knowledge, highlighting its knowledge limitations in specialized do-

Methods	Reliability	Locality	Generality						
	NR	RR	F1_T1	F1_T2	F1_T3	F1_T4	F1_T5	F1_Unseen	PubMedQA
LLaMa-2-7B	-	-	0.41	0.53	0.42	0.50	0.39	0.44	0.38
CALINET	1.00	0.52	0.81	0.75	0.50	0.68	0.46	0.55	0.46
T-Patcher	0.73	0.06	0.45	0.71	0.30	0.65	0.32	0.42	0.40
Prefix Tuning	0.70	0.90	0.78	0.71	0.63	0.54	0.60	0.59	0.44
LoRA	0.92	0.80	0.87	0.74	0.82	0.72	0.78	0.77	0.47
QLoRA	0.97	0.88	0.93	0.78	0.79	0.64	0.81	0.75	0.49
Ours	0.99	0.99	0.99	0.89	0.91	0.82	0.92	0.88	0.58

Table 1: Comparative results of InfuserKI with PEFT and ME methods on the UMLS 2.5k triplets.

Methods	Reliability	Locality	Generality						
	NR	RR	F1_T1	F1_T2	F1_T3	F1_T4	F1_T5	F1_Unseen	1HopQA
LLaMa-2-7B	-	-	0.57	0.45	0.53	0.42	0.52	0.49	0.47
CALINET	0.97	0.84	0.90	0.74	0.85	0.68	0.85	0.79	0.44
T-Patcher	0.39	0.75	0.60	0.69	0.57	0.62	0.61	0.81	0.36
Prefix Tuning	0.12	0.88	0.56	0.53	0.53	0.51	0.53	0.52	0.45
LoRA	0.90	0.80	0.84	0.79	0.81	0.76	0.82	0.80	0.62
QLoRA	0.93	0.90	0.91	0.82	0.89	0.80	0.90	0.86	0.69
Ours	0.99	0.96	0.97	0.88	0.97	0.86	0.94	0.92	0.67

Table 2: Comparative results of InfuserKI with PEFT and ME methods on the MetaQA KG.

mains. (2) Our method outperforms ME baselines such as CALINET and T-Patcher, which focus on correcting existing knowledge by positioning adapters in earlier transformer layers. This emphasis makes them less suited for integrating new knowledge compared to our approach. (3) Compared to PEFT methods such as Prefix Tuning, LoRA, and QLoRA, our method achieves superior locality (RR). This improvement stems from our infusing mechanism’s adaptive selection of supplementary information, which effectively prevents adapters from interfering with previously acquired knowledge. (4) Our method outperforms the T-Patcher across all metrics. Although T-Patcher reduces the impact on a minimal number of unrelated samples, it lacks robustness in locality, which our infusing mechanism effectively addresses. (5) Our approach demonstrates better generality on unseen templates and in the downstream tasks PubMedQA/1-HopQA, benefiting from our well-designed relation classification task.

Besides, Table 3 reveals our method maintains excellent performance in reliability, locality, and generality when scaling from 2,500 to 25,000 triplets on the UMLS KG, proving its capability in large-scale knowledge integration. In contrast, traditional ME methods show a performance decline at a larger scale, indicating their limitation to small-scale editing. For additional results on more datasets and with more baselines, please refer to Appendices A.5 and 4.8. Besides, despite the significant increase in triplets, we observe the unchanged performance on PubMedQA due to the

nature of PubMedQA as a new downstream task in the same domain with limited knowledge overlap. One primary benefit of knowledge injection via fine-tuning is to stimulate domain-specific knowledge. Therefore, injecting 2.5k pieces of knowledge may have already reached the saturation point for PubMedQA, beyond which no additional performance gains from 25k pieces are observed.

4.3 Ablation Study

To assess the impact of each component in InfuserKI, we compare it against variants without certain parts: (1) InfuserKI-w/o-RL, a variant without the Infuser loss; (2) InfuserKI-w/o-Ro, a variant without the Infuser module; (3) InfuserKI-w/o-RC, which excludes the relationship classification task. In Table 4, we notice: (1) Removing Infuser loss diminishes NR by 10%, indicating the role of infusing loss in distinguishing known from unknown information for effective integration. (2) Excluding the Infuser lowers RR by 7%, emphasizing its importance in minimizing knowledge forgetting. (3) Without the relation classification task, F1_Unseen decreases by 5%, showing its effectiveness in leveraging knowledge triplets to generalize new knowledge integration.

4.4 Impact of Adapter Position

To explore the benefits of adapter positions within the transformer architecture, we position adapters in the 3rd to 12th (bottom), 13th to 22nd (middle), and 23rd to 32nd (top) FFN layers, as well as across the 3rd to 32nd attention layers. Fig. 5 shows that

Methods	Reliability	Locality	Generality						
	NR	RR	F1_T1	F1_T2	F1_T3	F1_T4	F1_T5	F1_Unseen	PubMedQA
LLaMa-2-7B	-	-	0.35	0.47	0.36	0.50	0.36	0.41	0.38
CALINET	0.86	0.44	0.69	0.57	0.66	0.55	0.68	0.63	0.45
T-Patcher	0.63	0.20	0.45	0.55	0.38	0.53	0.37	0.43	0.43
Prefix-Tuning	0.82	0.80	0.82	0.59	0.79	0.61	0.77	0.72	0.47
LoRA	0.96	0.90	0.95	0.62	0.94	0.58	0.91	0.81	0.40
QLoRA	0.94	0.91	0.93	0.70	0.90	0.69	0.87	0.82	0.45
Ours	0.99	0.99	0.99	0.83	0.94	0.80	0.96	0.90	0.58

Table 3: Comparative results of InfuserKI with PEFT and ME methods on the UMLS 25k triplets.

Methods	NR	RR	F1_Unseen
InfuserKI	0.99	0.99	0.88
InfuserKI-w/o-RL	0.89	0.97	0.77
InfuserKI-w/o-Ro	0.97	0.92	0.87
InfuserKI-w/o-RC	0.96	0.97	0.83

Table 4: Ablation study on UMLS-2.5k.

(1) NR diminishes from the bottom to the top layers, indicating that top-layer adapters are less effective for knowledge integration. This could be attributed to the fact that knowledge representations in the upper layers depend on information from the lower layers and any deficiencies in the lower layers can impact the integration of knowledge. This observation aligns with prior studies (Huang et al., 2023; Dong et al., 2022), suggesting that while top layers are better for refining abstract concepts and knowledge correction, bottom layers are more suited for injecting new information; and (2) placing adapters in attention layers proves less effective for new knowledge integration, confirming that FFN layers act as storage for factual knowledge, which also agrees to the findings in previous studies (Dai et al., 2022; Geva et al., 2021).




Figure 5: Impact of Adapter Positions on InfuserKI.

4.5 Infuser Analysis

To delve deeper into the infusing mechanism, we visualize its values on the test set. As shown in Fig. 6, we display the infusing scores for both original known and unknown samples. Our observation is that infusing scores are lower on known samples, helping to block interfering information and thus mitigating knowledge forgetting.




Figure 6: Infusing Scores for Known vs. Unknown Samples.

4.6 Resource Requirements

To analyze our resource requirements, we compare various techniques, focusing on latency and parameter demands. All methods show similar latencies, due to providing short answers after fine-tuning. We examine memory usage by comparing additional parameter sizes for 2.5K and 25K scenarios using the LLaMa-2-7b model, as detailed in Table 5. Currently, both the 2.5K and 25K scenarios use the same parameter sizes. Both CALINET and our method use adapters of the same size, noted as 10. However, our InfuserKI framework perform better by incorporating the Infuser module.

4.7 Case Study

To intuitively understand the effectiveness of our framework, we compare the prediction score distributions over candidate choices from the vanilla

Methods	Parameter Demands (2.5K/25K)
CALINET	3.7M / 3.7M
T-Patcher	9.2M / 92M
Ours	3.7M / 3.7M

Table 5: Comparison of parameter amounts for different methods

LLaMa-2, LoRA, and our InfuserKI in two cases. Fig. 7 (a) shows that LLaMa-2, which initially gives incorrect answers, can provide correct answers after applying our InfuserKI and LoRA. However, LoRA induces forgetting for the second case, as depicted in Fig. 7 (b) while InfuserKI retains the knowledge.

What diagnosis is associated with the finding site of Sutura cranii? (A) Discharging mastoid cavity (finding) (B) Congenital anomaly of basioccipital bone (disorder) (C) Swelling over mastoid (finding) (D) Overlapping cranial sutures (finding)		
LLaMa: A 3e-7 B 0.999 C 3e-7 *D 3e-7	LoRA: A 3e-5 B 2e-5 C 2e-5 *D 0.999	InfuserKI: A 1e-6 B 6e-7 C 3e-6 *D 0.999
(a) LoRA and InfuserKI successfully inject knowledge into LLaMa.		
What procedure is performed on the Process Mastoideus? (A) Epithelial debris of mastoid cavity (finding) (B) Congenital anomaly of basioccipital bone (disorder) (C) Repair of mastoid antrum or cavity (procedure) (D) Finding of moistness of mastoid cavity (finding)		
LLaMa: A 5e-5 B 1e-4 *C 0.999 D 1e-4	LoRA: A 0.02 B 1e-5 *C 0.007 D 0.97	InfuserKI: A 3e-6 B 1e-7 *C 0.998 D 0.001
(b) LoRA forgets knowledge LLaMa knows and InfuserKI remembers.		

Figure 7: Illustration of Infuser-Guided Knowledge Integration with less forgetting.

4.8 Comparison with RAG Baselines

Both the Retrieval-Augmented Generation (RAG) method and our approach aim to enhance LLMs using external knowledge, necessitating a comparative analysis using the UMLS dataset. We have designed experiments to inject and assess knowledge specific to certain relation types, developing two RAG variants: RAG-TKS, which uses a BM25 retriever to utilize knowledge statements from the training set for context, and RAG-Google, which retrieves top-ranked content using Google. The results in Table 6 demonstrate that our method, which integrates knowledge directly into the model parameters, significantly outperforms both RAG variants. This enhanced performance may be attributable to the direct integration of knowledge into the parameters, which effectively stimulates the model capabil-

Methods	F1	Latency (ms)
LLaMa-2-7B	0.40	933
RAG-Google	0.37	2027
RAG-TKS	0.42	1113
Ours	0.66	860

Table 6: Comparative results of InfuserKI with RAG methods on the UMLS KG.

ity within specific domains. Moreover, our method exhibits lower inference latency than RAG, as it eliminates the need for external searches, and outperforms LLaMa-2-7B by delivering precise and concise answers without long explanatory texts.

5 Conclusion

In this study, we tackle a novel problem of integrating new knowledge from KGs into LLMs without affecting existing knowledge. We introduce the Infuser-guided Knowledge Integration framework, designed to selectively add new information to LLMs, minimizing the impact on prior knowledge and preventing catastrophic forgetting. A relation classification task further enhances the model’s generality. Evaluations on UMLS and MetaQA demonstrate InfuserKI’s effectiveness in integrating knowledge with less forgetting, maintaining sustained performance with large-scale data, and exhibiting exceptional generality on unseen templates and downstream tasks. Future work will study methods to test and integrate knowledge into LLMs with multi-hop knowledge triplets.

6 Limitations

We note that the effectiveness of our method is contingent upon the base language model’s ability to follow instructions accurately. In scenarios where the underlying model exhibits suboptimal instruction-following capabilities, the integration of knowledge, regardless of its quality, may not significantly improve performance on downstream tasks. Consequently, applying our knowledge integration framework to models with limited instruction-following proficiency presents a considerable challenge.

Acknowledgements

This work is supported by, or in part by, NEC Labs America gift funding.

References

- Amos Azaria and Tom Mitchell. 2023. The internal state of an llm knows when its lying. *arXiv preprint arXiv:2304.13734*.
- Olivier Bodenreider. 2004. The unified medical language system (umls): integrating biomedical terminology. *Nucleic acids research*, 32(suppl_1):D267–D270.
- Pietro Buzzega, Matteo Boschini, Angelo Porrello, Davide Abati, and Simone Calderara. 2020. Dark experience for general continual learning: a strong, simple baseline. *Advances in neural information processing systems*, 33:15920–15930.
- Damai Dai, Li Dong, Yaru Hao, Zhifang Sui, Baobao Chang, and Furu Wei. 2022. Knowledge neurons in pretrained transformers. In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 8493–8502.
- Tim Dettmers, Artidoro Pagnoni, Ari Holtzman, and Luke Zettlemoyer. 2023. Qlora: Efficient finetuning of quantized llms. *arXiv preprint arXiv:2305.14314*.
- 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.
- Hady Elsahar, Pavlos Vougiouklis, Arslan Remaci, Christophe Gravier, Jonathon Hare, Frederique Laforest, and Elena Simperl. 2018. T-rex: A large scale alignment of natural language with knowledge base triples. In *Proceedings of the Eleventh International Conference on Language Resources and Evaluation (LREC 2018)*.
- Denis Emelin, Daniele Bonadiman, Sawsan Alqahtani, Yi Zhang, and Saab Mansour. 2022. Injecting domain knowledge in language models for task-oriented dialogue systems. In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing*, pages 11962–11974.
- Miao Fan, Qiang Zhou, Emily Chang, and Fang Zheng. 2014. Transition-based knowledge graph embedding with relational mapping properties. In *Proceedings of the 28th Pacific Asia conference on language, information and computing*, pages 328–337.
- 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.
- Thomas Hartvigsen, Swami Sankaranarayanan, Hamid Palangi, Yoon Kim, and Marzyeh Ghassemi. 2023. Aging with grace: Lifelong model editing with discrete key-value adapters. In *NeurIPS Workshop on Robustness in Sequence Modeling*.
- Junxian He, Chunting Zhou, Xuezhe Ma, Taylor Berg-Kirkpatrick, and Graham Neubig. 2022. [Towards a unified view of parameter-efficient transfer learning](#). In *International Conference on Learning Representations*.
- Neil Houlsby, Andrei Giurgiu, Stanislaw Jastrzebski, Bruna Morrone, Quentin De Laroussilhe, Andrea Gesmundo, Mona Attariyan, and Sylvain Gelly. 2019. Parameter-efficient transfer learning for nlp. In *International Conference on Machine Learning*, pages 2790–2799. PMLR.
- Edward J Hu, Phillip Wallis, Zeyuan Allen-Zhu, Yuanzhi Li, Shean Wang, Lu Wang, Weizhu Chen, et al. 2021. Lora: Low-rank adaptation of large language models. In *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*.
- Qiao Jin, Bhuwan Dhingra, Zhengping Liu, William Cohen, and Xinghua Lu. 2019. Pubmedqa: A dataset for biomedical research question answering. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 2567–2577.
- James Kirkpatrick, Razvan Pascanu, Neil Rabinowitz, Joel Veness, Guillaume Desjardins, Andrei A Rusu, Kieran Milan, John Quan, Tiago Ramalho, Agnieszka Grabska-Barwinska, et al. 2017. Overcoming catastrophic forgetting in neural networks. *Proceedings of the national academy of sciences*, 114(13):3521–3526.
- Tom Kwiatkowski, Jennimaria Palomaki, Olivia Redfield, Michael Collins, Ankur Parikh, Chris Alberti, Danielle Epstein, Illia Polosukhin, Jacob Devlin, Kenton Lee, et al. 2019. Natural questions: A benchmark for question answering research. *Transactions of the Association for Computational Linguistics*, 7:452–466.
- 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)*, pages 333–342, Vancouver, Canada. Association for Computational Linguistics.
- Shaobo Li, Xiaoguang Li, Lifeng Shang, Zhenhua Dong, Chengjie Sun, Bingquan Liu, Zhenzhou Ji, Xin Jiang, and Qun Liu. 2022. [How pre-trained language models capture factual knowledge? a causal-inspired analysis](#). In *Findings of the Association for Computational Linguistics: ACL 2022*, pages 1720–1732, Dublin, Ireland. Association for Computational Linguistics.

- Xiang Lisa Li and Percy Liang. 2021. Prefix-tuning: Optimizing continuous prompts for generation. In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 4582–4597.
- David Lopez-Paz and Marc’Aurelio Ranzato. 2017. Gradient episodic memory for continual learning. *Advances in neural information processing systems*, 30.
- Ilya Loshchilov and Frank Hutter. 2018. Decoupled weight decay regularization. In *International Conference on Learning Representations*.
- Potsawee Manakul, Adian Liusie, and Mark JF Gales. 2023. Selfcheckgpt: Zero-resource black-box hallucination detection for generative large language models. *arXiv preprint arXiv:2303.08896*.
- Sourab Mangrulkar, Sylvain Gugger, Lysandre Debut, Younes Belkada, Sayak Paul, and Benjamin Bossan. 2022. Peft: State-of-the-art parameter-efficient fine-tuning methods. <https://github.com/huggingface/peft>.
- Michael McCloskey and Neal J Cohen. 1989. Catastrophic interference in connectionist networks: The sequential learning problem. In *Psychology of learning and motivation*, volume 24, pages 109–165. Elsevier.
- 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*.
- Zaiqiao Meng, Fangyu Liu, Thomas Clark, Ehsan Shareghi, and Nigel Collier. 2021. Mixture-of-partitions: Infusing large biomedical knowledge graphs into bert. In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 4672–4681.
- Aaron van den Oord, Yazhe Li, and Oriol Vinyals. 2018. Representation learning with contrastive predictive coding. *arXiv preprint arXiv:1807.03748*.
- Yujia Qin, Yankai Lin, Ryuichi Takano, Zhiyuan Liu, Peng Li, Heng Ji, Minlie Huang, Maosong Sun, and Jie Zhou. 2021. ERICA: Improving entity and relation understanding for pre-trained language models via contrastive learning. In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 3350–3363, Online. Association for Computational Linguistics.
- Roger Ratcliff. 1990. Connectionist models of recognition memory: constraints imposed by learning and forgetting functions. *Psychological review*, 97(2):285.
- Dominic Seyler, Mohamed Yahya, and Klaus Berberich. 2017. Knowledge questions from knowledge graphs. In *Proceedings of the ACM SIGIR International Conference on Theory of Information Retrieval*, pages 11–18.
- Rohit Sridhar and Diyi Yang. 2022. Explaining toxic text via knowledge enhanced text generation. In *Proceedings of the 2022 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 811–826.
- Yu Sun, Shuohuan Wang, Yukun Li, Shikun Feng, Xuyi Chen, Han Zhang, Xin Tian, Danxiang Zhu, Hao Tian, and Hua Wu. 2019. Ernie: Enhanced representation through knowledge integration. *arXiv preprint arXiv:1904.09223*.
- Yueqing Sun, Qi Shi, Le Qi, and Yu Zhang. 2022. Jointlk: Joint reasoning with language models and knowledge graphs for commonsense question answering. In *Proceedings of the 2022 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 5049–5060.
- Hugo Touvron, Thibaut Lavril, Gautier Izacard, Xavier Martinet, Marie-Anne Lachaux, Timothée Lacroix, Baptiste Rozière, Naman Goyal, Eric Hambro, Faisal Azhar, et al. 2023a. Llama: Open and efficient foundation language models. *arXiv preprint arXiv:2302.13971*.
- Hugo Touvron, Louis Martin, Kevin Stone, Peter Albert, Amjad Almahairi, Yasmine Babaei, Nikolay Bashlykov, Soumya Batra, Prajjwal Bhargava, Shrutu Bhosale, et al. 2023b. Llama 2: Open foundation and fine-tuned chat models. *arXiv preprint arXiv:2307.09288*.
- Ruize Wang, Duyu Tang, Nan Duan, Zhongyu Wei, Xuan-Jing Huang, Jianshu Ji, Guihong Cao, Daxin Jiang, and Ming Zhou. 2021. K-adapter: Infusing knowledge into pre-trained models with adapters. In *Findings of the Association for Computational Linguistics: ACL-IJCNLP 2021*, pages 1405–1418.
- Chaoyi Wu, Xiaoman Zhang, Ya Zhang, Yanfeng Wang, and Weidi Xie. 2023. Pmc-llama: Further fine-tuning llama on medical papers. *arXiv preprint arXiv:2304.14454*.
- LI Xuhong, Yves Grandvalet, and Franck Davoine. 2018. Explicit inductive bias for transfer learning with convolutional networks. In *International Conference on Machine Learning*, pages 2825–2834. PMLR.

Wenhao Yu, Chenguang Zhu, Lianhui Qin, Zhihan Zhang, Tong Zhao, and Meng Jiang. 2022. Diversifying content generation for commonsense reasoning with mixture of knowledge graph experts. In *Proceedings of the 2nd Workshop on Deep Learning on Graphs for Natural Language Processing (DLG4NLP 2022)*, pages 1–11.

Yuexiang Zhai, Shengbang Tong, Xiao Li, Mu Cai, Qing Qu, Yong Jae Lee, and Yi Ma. 2024. Investigating the catastrophic forgetting in multimodal large language model fine-tuning. In *Conference on Parsimony and Learning*, pages 202–227. PMLR.

Renrui Zhang, Jiaming Han, Aojun Zhou, Xiangfei Hu, Shilin Yan, Pan Lu, Hongsheng Li, Peng Gao, and Yu Qiao. 2023. Llama-adapter: Efficient fine-tuning of language models with zero-init attention. *arXiv preprint arXiv:2303.16199*.

Taolin Zhang, Chengyu Wang, Nan Hu, Minghui Qiu, Chengguang Tang, Xiaofeng He, and Jun Huang. 2022. Dkplm: decomposable knowledge-enhanced pre-trained language model for natural language understanding. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 36, pages 11703–11711.

Xikun Zhang, Antoine Bosselut, Michihiro Yasunaga, Hongyu Ren, Percy Liang, Christopher D Manning, and Jure Leskovec. 2021. Greaselm: Graph reasoning enhanced language models. In *International conference on learning representations*.

Yuyu Zhang, Hanjun Dai, Zornitsa Kozareva, Alexander Smola, and Le Song. 2018. Variational reasoning for question answering with knowledge graph. In *Proceedings of the AAAI conference on artificial intelligence*, volume 32.

Ziwan Zhao, Linmei Hu, Hanyu Zhao, Yingxia Shao, and Yequan Wang. 2023. Knowledgeable parameter efficient tuning network for commonsense question answering. In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 9051–9063.

A Appendix

A.1 Template Prompts and MCQA Construction

To facilitate an effective comparison between long-form answers from LLMs and standard answers for open-ended questions, we utilize a multiple-choice format, as detailed in Table 7. This format comprises a correct answer alongside three distractors. The first distractor is chosen for its minimal edit distance to the head entity, while the remaining two are randomly selected from a set of ten candidates based on their edit distance to the correct answer. Subsequently, these choices are randomized and

I need five question-answer templates and a knowledge statement to analyze relationships in triplets formatted as <SUBJECT, RELATION, OBJECT>, focusing on the relation {RELATION}. Answers should be either the [OBJECT] entity or a yes/no response. Use placeholders [SUBJECT] and [OBJECT] to denote where the subject and object entities will be inserted. The knowledge statement should be a VERY brief, declarative sentence illustrating the RELATION between [SUBJECT] and [OBJECT], incorporating the original relation words ‘possibly equivalent to’.

Context is provided by the following examples: {EXAMPLE TRIPLETS}

Please create five unique question-answer templates and one knowledge statement, formatted as a JSON string. For clarity, the output should follow this format:

```
{ 'rel': { RELATION },
  'template#1': '[Question-answer template 1]',
  'template#2': '[Question-answer template 2]',
  'template#3': '[Question-answer template 3]',
  'template#4': '[Question-answer template 4]',
  'template#5': '[Question-answer template 5]',
  'knowledge_statement': '[Knowledge statement]',
  'memo': '[Additional memo or notes]' }
```

Note: ONLY OUTPUT A JSON STRING, NO ANY OTHER CONTENT.

Output: <Your generated JSON string>

Table 7: Prompt to GPT-4 to generate QA templates

presented as options (A), (B), (C), and (D) alongside the question, allowing for a precise assessment of LLMs’ knowledge in specific domains.

A.2 Algorithm

The algorithm is described in Algorithm 1.

A.3 Knowledge Graphs and Datasets

UMLS (Bodenreider, 2004): The Unified Medical Language System (UMLS) knowledge graph, developed by the US National Library of Medicine, integrates over 2 million terms for nearly 900,000 concepts from more than 60 biomedical vocabularies. These include the NCBI taxonomy, Gene Ontology, and Medical Subject Headings (MeSH), along with 12 million concept relations. For testing, we employ the PubMedQA dataset (Jin et al., 2019), a biomedical QA dataset derived from PubMed ab-

<p>Below is an instruction that describes a task. Write a response that appropriately completes the request. ### Instruction: {instruction} ### Response:</p>
--

Table 8: Prompt to LLMs to answer MCQA

Algorithm 1 Infuser-Guided Knowledge Integration.

- 1: **procedure** ROUTERKI(p_θ, \mathcal{G}) ▷ Target LLM p_θ and KG \mathcal{G} with triplets $\langle h, r, t \rangle$
- 2: **# Step 1: Knowledge Detection**
- 3: Convert triplets into MCQs q , with correct answers y and knowledge statements k , using relational templates.
- 4: Input MCQs into p_θ to identify unknown knowledge.
- 5: **# Step 2: Knowledge Integration**
- 6: Tune Infuser on a balanced mix of known and unknown samples as per Eq. 5.
- 7: Fine-tune adapters for templates #1 and #2 using QA loss in Eq. 8.
- 8: Apply relation classification to unknown statements, following Eq. 9 and Eq. 10.

stracts, featuring Yes/No/Maybe questions alongside context, as highlighted in Wu et al. (2023).

MetaQA (Zhang et al., 2018) serves as a multi-hop KGQA benchmark in the movie domain, presenting a knowledge graph with 135,000 triplets, 43,000 entities, and 9 relations. It organizes over 400,000 questions into 1-hop, 2-hop, and 3-hop categories, each annotated with head entities, answers, and reasoning paths. Our analysis concentrates on the 1-hop version for downstream testing.

A.4 Three Evaluation Properties

Following Huang et al. (2023), the enhanced LLM should meet these properties:

Property 1, Reliability: The enhanced model p'_θ incorporates knowledge previously unknown to p_θ as

$$p'_\theta(x) = y \text{ if } p_\theta(x) \neq y. \quad (11)$$

Reliability is quantified using the Newly-learned Rate (NR) in our work.

Property 2, Locality: Knowledge integration should be localized and precise, ensuring the fine-tuned model p'_θ retains accuracy on \mathcal{T}_{known} , the

knowledge previously known to p_θ as

$$p'_\theta(x) = y \text{ if } p_\theta(x) = y. \quad (12)$$

Here, this property is measured by the Remembering Rate (RR), which indicates the accuracy on the previously acquired knowledge.

Property 3, Generality: For any unknown sample x , let $\mathbb{E}_x = \{x' | y_{x'} = y_x\}$ denote a set of equivalent inputs. The model p'_θ should correctly answer all instances $x' \in \mathbb{E}_x$ as

$$\forall x' \in \mathbb{E}_x, p'_\theta(x') = y. \quad (13)$$

In this study, generality is assessed by averaging F1 scores (F1_Unseen) across three unseen templates during training as well as performance on downstream tasks.

A.5 Results on ME Datasets and YAGO

We conduct experiments on two Wikipedia-sourced datasets used in Model Editing (ME) methods: the Zero-Shot Relation Extraction (zsRE) dataset (Levy et al., 2017) and the T-REx dataset (Elsahar et al., 2018). We also perform comparative experiments using sampled knowledge graphs from YAGO. The results in Table 9, 10, and 11 show that the LLM backbone has deficiencies in handling world knowledge across three datasets, but performance improves with our knowledge injection method, achieving optimal specificity, locality, and generality.

Methods	NR	RR	F1_T1	F1_T2	F1_T3	F1_T4	F1_T5	F1_Unseen
LLaMa-2-7B	-	-	0.51	0.59	0.48	0.59	0.49	0.52
CALINET	0.61	0.49	0.54	0.66	0.53	0.63	0.49	0.55
LoRA	0.55	0.55	0.55	0.54	0.57	0.52	0.51	0.53
Ours	0.84	0.95	0.91	0.80	0.82	0.65	0.81	0.76

Table 9: Comparative results of InfuserKI with PEFT and ME methods on the zsRE-1k.

Methods	NR	RR	F1_T1	F1_T2	F1_T3	F1_T4	F1_T5	F1_Unseen
LLaMa-2-7B	-	-	0.64	0.62	0.66	0.63	0.62	0.64
CALINET	0.94	0.72	0.80	0.65	0.68	0.62	0.72	0.67
LoRA	0.66	0.63	0.64	0.64	0.68	0.57	0.68	0.64
Ours	1.00	0.98	0.99	0.89	0.97	0.79	0.97	0.84

Table 10: Comparative results of InfuserKI with PEFT and ME methods on the TREx-1k.

Methods	NR	RR	F1_T1	F1_T2	F1_T3	F1_T4	F1_T5	F1_Unseen
LLaMa-2-7B	-	-	0.63	0.58	0.61	0.61	0.60	0.61
CALINET	0.65	0.60	0.61	0.71	0.71	0.68	0.64	0.68
LoRA	0.81	0.79	0.80	0.83	0.80	0.62	0.57	0.66
Ours	1.00	0.90	0.94	0.95	0.95	0.79	0.79	0.84

Table 11: Comparative results of InfuserKI with PEFT and ME methods on the YAGO-1k KG.