

Editing Language Model-Based Knowledge Graph Embeddings

Siyuan Cheng^{1, 2, 4*}, Ningyu Zhang^{1, 2†},
Bozhong Tian^{1, 2*}, Xi Chen^{4*}, Qingbin Liu⁴, Huajun Chen^{1, 2, 3†}

¹Zhejiang University

²Zhejiang University - Ant Group Joint Laboratory of Knowledge Graph

³Donghai Laboratory

⁴Platform and Content Group, Tencent

{sycheng, zhangningyu, tbozhong, huajunsir}@zju.edu.cn

{jasonxchen, qingbinliu}@tencent.com

Abstract

Recently decades have witnessed the empirical success of framing Knowledge Graph (KG) embeddings via language models. However, language model-based KG embeddings are usually deployed as static artifacts, making them difficult to modify post-deployment without re-training after deployment. To address this issue, we propose a new task of editing language model-based KG embeddings in this paper. This task is designed to facilitate rapid, data-efficient updates to KG embeddings without compromising the performance of other aspects. We build four new datasets: E-FB15k237, A-FB15k237, E-WN18RR, and A-WN18RR, and evaluate several knowledge editing baselines demonstrating the limited ability of previous models to handle the proposed challenging task. We further propose a simple yet strong baseline dubbed KGEitor, which utilizes additional parametric layers of the hypernetwork to edit/add facts. Our comprehensive experimental results reveal that KGEitor excels in updating specific facts without impacting the overall performance, even when faced with limited training resources. Code and datasets will be available at <https://github.com/AnonymousForPapers/DeltaKG>.

Introduction

Knowledge Graphs (KGs) represent large-scale, multi-relational graphs containing a wealth of symbolic facts. These structures offer invaluable back-end support for various knowledge-intensive tasks, such as information retrieval, question-answering, and recommender systems. (Yang et al. 2021, 2022; Wu et al. 2022b,a). To better utilize that symbolic knowledge in KGs for machine learning models, many KG embedding approaches have been devoted to representing KGs in low-dimension vector spaces (Wang et al. 2017; Zhang et al. 2022a). Traditional KG embedding models, e.g., TransE (Bordes et al. 2013), RotatE (Sun et al. 2019), are naturally taxonomized as structure-based methods (Xie et al. 2016; Zhang et al. 2020b,a; Wang et al. 2022b). These approaches employ supervised machine

*These authors contributed equally.

†Corresponding author.

Copyright © 2024, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

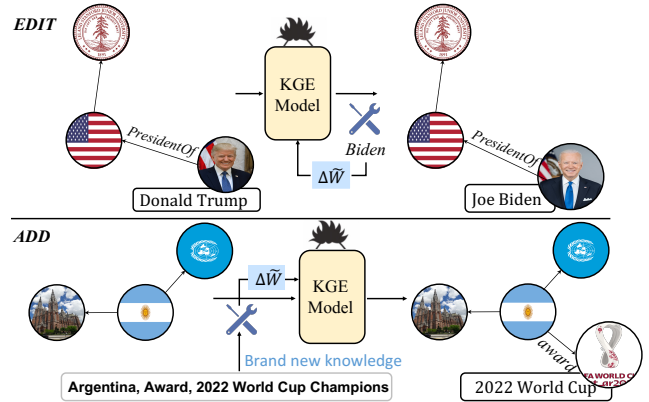


Figure 1: Top: Illustration of the EDIT task. We edit the wrong fact knowledge stored in the KG embeddings. Bottom: Illustration of the ADD task. We add brand-new knowledge into the model without re-training.

learning to optimize target objectives by utilizing scoring functions that preserve the inherent structure of KGs.

However, a recent shift in KG embedding methodologies has emerged, moving away from the explicit modeling of structure (Yao, Mao, and Luo 2019; Zhang et al. 2020c; Wang et al. 2022b,a). Instead, contemporary techniques focus on incorporating text descriptions through the use of expressive black-box models, e.g., pre-trained language models. This new paradigm operates under the assumption that the model will inherently capture the underlying structure without requiring explicit instruction. Leveraging language models to frame KG embeddings has emerged as a highly promising approach, yielding considerable empirical success. This technique offers the potential to generate informative representations for long-tail entities and those on-the-fly emerging entities. However, KG embeddings with language models are usually deployed as static artifacts, which are challenging to modify without re-training. To respond to changes (e.g., emerging new facts) or a correction for facts of existing KGs, the ability to conduct flexible knowledge updates to KG embeddings **after deployment** is desirable.

To address this need, we introduce a new task of editing

language model-based KG embeddings, which aims to enable data-efficient and fast updates to KG embeddings for a small region of parametric space without influencing the performance of the rest. It's important to note that while editing KG embeddings primarily centers on link prediction tasks, the editing of standard language models mainly addresses tasks such as question answering, each presenting its unique set of challenges. For instance, KG embeddings contend with the complex issue of handling numerous many-to-many (N-M) triples. Intuitively, we define two tasks for pre-trained KG embeddings, namely, EDIT and ADD, to support editing incorrect or adding new facts without re-training the whole model during deployment as shown in Figure 1. We treat editing the memories of KG embeddings as a *learning-to-update* problem, which can ensure that knowledge represented in KG embeddings remains accurate and up-to-date. To evaluate the performance of the proposed task, we establish three guiding principles: **Knowledge Reliability**, which indicates that those edited or newly added knowledge should correctly be inferred through link prediction; **Knowledge Locality**, which means that editing KG embeddings will not affect the rest of the acquired knowledge when successfully updating specific facts; **Knowledge Efficiency** indicates being able to modify a model with low training resources. Specifically, we build four new datasets based on FB15k237 and WN18RR for EDIT and ADD tasks. We leverage several approaches, including Knowledge Editor (KE) (Cao, Aziz, and Titov 2021), MEND (Mitchell et al. 2022), and CALINET (Dong et al. 2022) as baselines. Editing KGE proves more challenging than refining language models. Experimental results demonstrate that current methods struggle with efficiently modifying KG embeddings. We further propose a simple yet effective strong baseline dubbed **Knowledge Graph Embeddings Editor (KGEditor)**, which can efficiently manipulate knowledge in embeddings by editing additional parametric layers. Our experiments demonstrate that **KGEditor** can effectively modify incorrect knowledge or add new knowledge while preserving the integrity of other information. We summarize our contributions as follows:

- We propose a new task of editing language model-based KG embeddings. The proposed task with datasets may open new avenues for improving KG embedding via knowledge editing.
- We introduce **KGEditor** that can efficiently modify incorrect knowledge or add new knowledge without affecting the rest of the acquired knowledge.
- We conduct extensive comparisons with in-depth analysis of four datasets and report empirical results with insightful findings, demonstrating the effectiveness of the proposed approach.

Related Work

Knowledge Graph Embedding

Early KG embedding approaches primarily focused on deriving embeddings from structured information alone. Most existing KG embedding methods utilize the translation-based paradigm such as TransE (Bordes et al. 2013), TransR

(Lin et al. 2015), TransH (Wang et al. 2014), and RotatE (Sun et al. 2019) or semantic matching paradigms, including DistMult (Yang et al. 2015), RESCAL (Nickel, Tresp, and Kriegel 2011), and HolE (Nickel, Rosasco, and Poggio 2016). Additionally, there has been significant interest in explicitly utilizing structural information through graph convolution networks (Kipf and Welling 2017; Velickovic et al. 2018; Vashishth et al. 2020; Liu et al. 2021b; Zhang et al. 2022b). Since pre-trained language models (PLMs) (Zhao et al. 2023) have made waves for widespread tasks, framing KG embedding via language models is an increasing technique that has led to empirical success (Pan et al. 2023). On the one hand, several works have emerged to leverage PLMs for KG embeddings. Some studies (Yao, Mao, and Luo 2019; Zhang et al. 2020c; Wang et al. 2021c, 2022a) leverage finetuning the PLMs, for example, KG-BERT (Yao, Mao, and Luo 2019), which takes the first step to utilize BERT (Devlin et al. 2019) for KG embeddings by regarding a triple as a sequence and turning link prediction into a sequence classification task. StAR (Wang et al. 2021a) proposes a structure-augmented text representation approach that employs both spatial measurement and deterministic classifier for KG embeddings, respectively. LMKE (Wang et al. 2022b) adopts language models which utilize description-based KG embedding learning with a contrastive learning framework. On the other hand, some studies (Lv et al. 2022) adopt the prompt tuning (Liu et al. 2022; Chen et al. 2022b; Zhang, Li, and et al. 2022) with language models. PKGC (Lv et al. 2022) converts triples into natural prompt sentences for KG embedding learning. There are also some other studies (Xie et al. 2022; Saxena, Kochsiek, and Gemulla 2022; Chen et al. 2022a) formulate KG embedding learning as sequence-to-sequence generation with language models. Nonetheless, KG embeddings in previous studies are usually deployed as static artifacts whose behavior is challenging to modify after deployment. We propose a novel task: editing language model-based KG embeddings to correct or modify facts in existing KGs, marking the first initiative in this field to our knowledge.

Editing Factual Knowledge

Editable training (Sinitsin et al. 2020; Yao et al. 2023) represents an early, model-agnostic attempt to facilitate rapid editing of trained models. Subsequently, numerous reliable and effective approaches have been proposed to enable model editing without the necessity for resource-intensive re-training (Cohen et al. 2023; Geva et al. 2023; Hase et al. 2023; Han et al. 2023). These methods can be broadly classified into two distinct categories: external model-based editor and additional parameter-based editor. The first category, external model-based editors, employs extrinsic model editors to manipulate model parameters. For instance, KnowledgeEditor (KE) (Cao, Aziz, and Titov 2021) adopts a hyper-network-based approach to edit knowledge within language models. MEND (Mitchell et al. 2022) utilizes a hypernetwork (MLP) to predict the rank-1 decomposition of fine-tuning gradients. The predicted gradients are used to update a subset of the parameters of PLMs. The second category, additional parameter-based editors,

involves using supplementary parameters to adjust the final output of the model, thereby achieving model editing. Building upon prior research (Dai et al. 2022a), which suggests that Feed-Forward Networks (FFNs) in PLMs may store factual knowledge, CALINET (Dong et al. 2022) enhances a specific FFN within the PLM by incorporating additional parameters composed of multiple calibration memory slots. Existing approaches predominantly focus on modifying knowledge within pre-trained language models, constraining the manipulation and interpretation of facts. In contrast, by editing knowledge in KG embeddings, our proposed tasks enable adding and modifying knowledge, thereby extending the applicability of KG embeddings across various downstream tasks. Previous methods for updating knowledge graph embeddings (KGE), such as OUK (Fei, Wu, and Khan 2021) and RotatH (Wei et al. 2021), mainly focused on modifying the entity and relation vectors in score function-based KGE models. These methods employed different dataset versions as snapshots to represent changes within the knowledge graph and validated these changes directly using hit@k. However, this approach falls short of effectively measuring the efficacy of edits. Our paper improves these methods by creating datasets and introducing metrics for a more accurate evaluation of KGE updates.

Methodology

Task Definition

A KG can be represented as $\mathcal{G} = (\mathcal{E}, \mathcal{R}, \mathcal{T})$, where \mathcal{E} , \mathcal{R} and \mathcal{T} are sets of entities, relation types, and triples. Each triple in \mathcal{T} takes the form (h, r, t) , where $h, t \in \mathcal{E}$ are the head and tail entities. For the EDIT task, knowledge requiring edits is defined as (h, r, y, a) or (y, r, t, a) , where y is an incorrect/outdated entity. We denote the original input entity as x and the predicted target as y . a points to the desired edited entity (change from y to a). E.g., $\langle \text{Donald Trump, president_of, U.S.A.} \rangle$ is an outdated triple, which should be changed to $\langle \text{Joe Biden, president_of, U.S.A.} \rangle$ to update the KG embeddings. Apart from outdated information, numerous new facts may be absent in current KGs. Consequently, it becomes essential to flexibly incorporate these new triples into the KG embeddings. Thus, we introduce the ADD task to integrate new knowledge seamlessly. Note that the ADD task is similar to the inductive setting in KG completion but without re-training the model. Formally, we define the task of editing language model-based KG embeddings as follows:

$$p(a \mid \tilde{\mathcal{W}}, \phi, x) \leftarrow p(y \mid \mathcal{W}, \phi, x) \quad (1)$$

where \mathcal{W} and $\tilde{\mathcal{W}}$ denote the original and edited parameters of KG embeddings, respectively. In this paper, x refers to the original input (e.g., $(?, h, r)$), output entity (y), and desired edited entity (a). ϕ represents the external parameters of the editor network. At the same time, we need to strive to maintain the stability of the model for other correct knowledge, as described below:

$$p(y' \mid \tilde{\mathcal{W}}, \phi, x') \leftarrow p(y' \mid \mathcal{W}, \phi, x') \quad (2)$$

where x' and y' represent the input and label of the factual knowledge stored in the model, respectively, with $x' \neq x$.

Evaluation Metrics We set three principles to measure our proposed tasks' efficacy.

Knowledge Reliability evaluates whether edited or new knowledge is correctly inferred through link prediction. Changes in KG embeddings should align with the intended edits without bringing unintended biases/errors. For editing effect validation, we adopt the KG completion setting. By ranking candidate entity scores, we generate an entity list. We define the **Success@1 metric (Succ@k)** by counting the number of correct triples that appear at position k .

Knowledge Locality seeks to evaluate whether editing KG embeddings will impact the rest of the acquired knowledge \mathcal{O}^x when successfully updating specific facts. We introduce the **Retain Knowledge (RK@k)** metric, indicating that the entities predicted by the original model are *still correctly inferred by the edited model*. Specifically, we sample the rank k triples predicted by the original model as a reference dataset (L-test in Table 1). We posit that if the predicted triples remain unchanged after editing, then the model adheres to knowledge locality and does not affect the rest of the facts. We calculate the proportion of retaining knowledge as a measure of the stability of the model editing:

$$RK@k = \frac{\sum f(x'; \tilde{\mathcal{W}})_{\leq k}}{\sum f(x'; \mathcal{W})_{\leq k}} \quad (3)$$

where x' are randomly sampled from the reference dataset. f is the function with input x and parameterized by \mathcal{W} , computing the rank of x . $f(x'; \mathcal{W})_{\leq k}$ signifies the count of rank value under k predicted by the model with \mathcal{W} parameters.

Besides, to better present the effect of models, we introduce two additional metrics. **Edited Knowledge Rate of Change** ER_{roc} and **Retaining Knowledge Rate of Change** RK_{roc} are denoted as:

$$ER_{roc} = \frac{|R_{edit} - R_{origin}|}{R_{origin}}, \quad (4)$$

$$RK_{roc} = \frac{|R_{s_edit} - R_{s_origin}|}{R_{s_edit}}$$

where R_{edit} and R_{origin} represent the mean rank on the test set before and after editing. R_{s_edit} and R_{s_origin} denote the reference L-test's mean rank before and after editing.

Knowledge Efficiency aims to evaluate the efficiency of the editors. We leverage the tuning parameters (including external model or additional parameters; see Params in Table 2, 3) as a metric for evaluation.

Datasets Construction

We construct four datasets for EDIT and ADD tasks, based on two benchmarks FB15k-237 (Toutanova et al. 2015), and WN18RR (Dettmers et al. 2018). After initially training KG embeddings with language models, we sample challenging triples as candidates, as detailed below. Recent research indicates that pre-trained language models can learn significant factual knowledge (Petroni et al. 2019; Cao et al. 2021), and predict the correct tail entity with a strong bias. This may hinder the proper evaluation of the editing task. Thus, we exclude relatively simple triples inferred from language

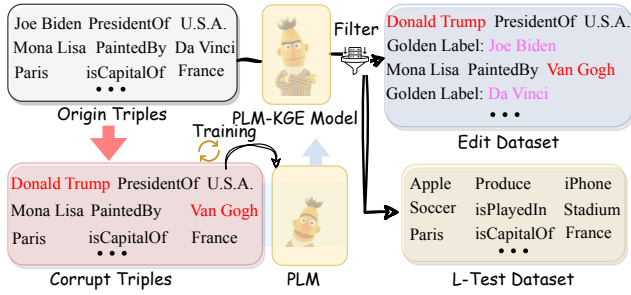


Figure 2: Data Construction Process. Step 1: Randomly disrupt the triples from the existing knowledge graph to produce a corrupted dataset. Step 2: Employ this dataset to fine-tune a pre-existing pre-trained KGE model, yielding a model needing editing. Step 3: Filtering, we reassess the data with the pre-trained KGE model, accurately sorting the correctly labeled data into the L-Test dataset and allocating the mislabeled data to a designated dataset for correction.

models’ **internal knowledge** for precise assessment. Specifically, we select data with link prediction ranks $>2,500$. For EDIT, the Top-1 facts from the origin model replace the entities in the training set of FB15k-237 to build the pre-train dataset (original pre-train data), while golden facts serve as target edited data. For ADD, we leverage the original training set of FB15k-237 to build the pre-train dataset (original pre-train data) and the data from the inductive setting as they are not seen before. Unlike EDIT, for ADD, since new knowledge is not present during training, we **directly use the training set to evaluate the ability to incorporate new facts**. We adopt the same strategy to construct two datasets from WN18RR. To assess knowledge locality, we further create a reference dataset (L-Test) based on the link prediction performance with a rank value below k (the same k as in the RK metric). Five human experts **scrutinize all datasets** for ethical concerns. Ultimately, we establish four datasets: E-FB15k237, A-FB15k237, E-WN18RR, and A-WN18RR. Table 1 provides details of the datasets. The creation of the datasets for the EDIT task is shown in Figure 2.

Language Model-based KGE

This section introduces the technical background of language model-based KG embeddings. Specifically, we illustrate two kinds of methodologies to leverage language models, namely: finetuning (**FT-KGE**) and prompt tuning (**PT-KGE**). Note that our task is orthogonal to language-based KG embeddings, performance differences between **FT-KGE** and **PT-KGE** are discussed in the next section.

FT-KGE methods, such as KG-BERT (Yao, Mao, and Luo 2019), which regard triples in KGs as textual sequences. Specifically, they are trained with the description of triples that represent relations and entities connected by the [SEP] and [CLS] tokens and then take the description sequences as the input for finetuning. Normally, they use the presentation of [CLS] token to conduct a binary classification, note as:

$$p(\mathbf{y}_\tau | x) = p([\text{CLS}] | x; \mathcal{W}) \quad (5)$$

where $\mathbf{y}_\tau \in [0, 1]$ is the label of whether the triple is positive.

TASK	FB15k-237				WN18RR			
	Pre-train	Train	Test	L-Test	Pre-train	Train	Test	L-Test
ADD	215,082	2,000	-	16,872	69,721	2,000	-	10,000
EDIT	310,117	3,087	3,087	7,051	93,003	1,491	830	5,003

Table 1: The statistics of datasets for the EDIT and ADD tasks. L-Test is the test set for knowledge locality to evaluate the rest of the knowledge in KGE.

PT-KGE methods like PKGC (Lv et al. 2022), utilize natural language prompts to elicit knowledge from pre-trained models for KG embeddings. Prompt tuning uses PLMs as direct predictors to complete a cloze task, linking pre-training and finetuning. In this paper, we implement PT-KGE with entity vocabulary expansion, treating each unique entity as common knowledge embeddings. Specifically, we consider entities $e \in \mathcal{E}$ as special tokens in the language model, turning link prediction into a masked entity prediction. Formally, we can obtain the correct entity by ranking the probability of each entity in the knowledge graph as follows:

$$p(\mathbf{y}|x) = p([\text{MASK}] = y|x; \mathcal{W}) \quad (6)$$

where $p(\mathbf{y}|x)$ is denoted as the probability distribution of the entity in the KG.

Editing KGE Baselines

With the pre-trained KG embeddings via the sections mentioned above, we introduce several baselines to edit KG embeddings. Concretely, we divide baseline models into two categories. The first is the **external model-based editor**, which uses an extrinsic model editor to control the model parameters such as KE and MEND. Another is the **additional parameter-based editor**, which introduces additional parameters to adjust the final output of the model to achieve the model editing. There are also some non-editing methods in the baseline: Finetune and K-Adapter (Wang et al. 2021b).

External Model-based Editor This method employs a hypernetwork to learn the weight update Δ for editing the language model. KE utilizes a hypernetwork (specifically, a bidirectional-LSTM) with constrained optimization, which is used to predict the weight update during inference. Thus, KE can modify facts without affecting the other knowledge in parametric space. MEND conducts efficient local edits to language models with a single input-output pair, which learns to transform the gradient of finetuned language models, which utilizes a low-rank decomposition of gradients. Hence, MEND can edit the parameters of large models, such as BART (Lewis et al. 2020), GPT-3 (Brown and et al. 2020), and T5 (Raffel et al. 2020), with little resource cost. To apply KE and MEND to our task setting, we remove the semantically equivalent set¹ of the original task and only utilize the triples that need to be edited as input x .

¹Semantic equivalence refers to a declaration that two data elements from different vocabularies contain data having a similar meaning (Cao, Aziz, and Titov 2021).

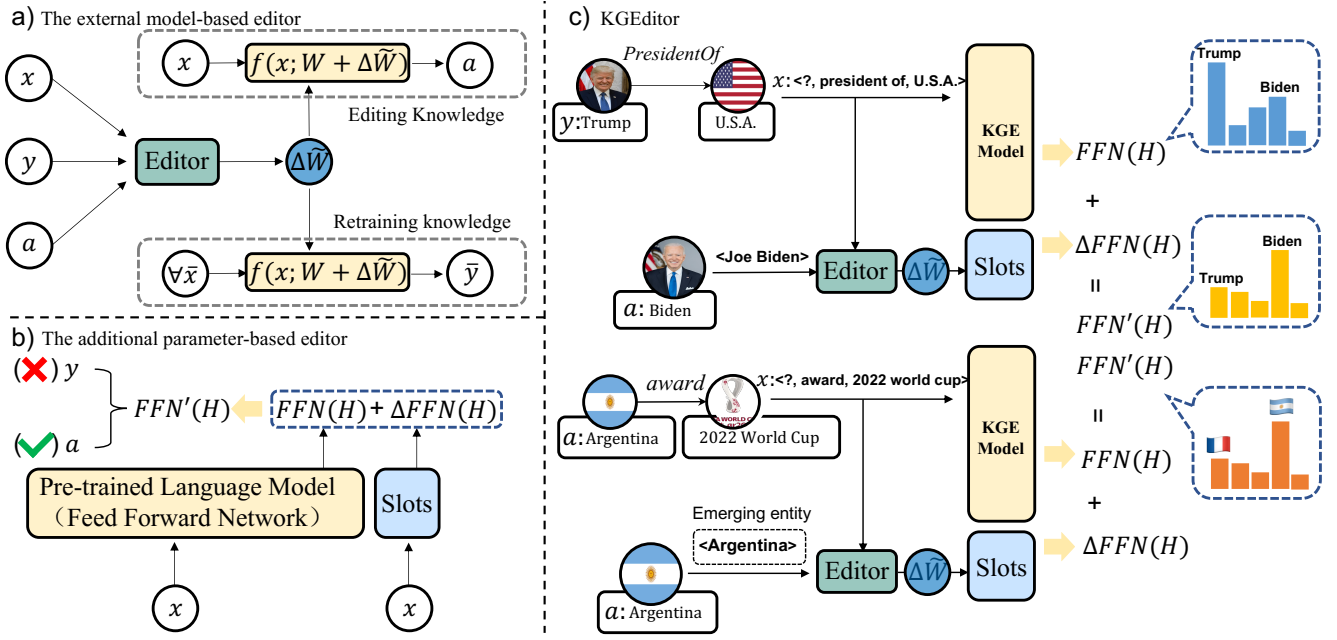


Figure 3: The introduction of baselines and KGEitor. The external model-based editors (a) utilize a hyper external network to obtain the parameters’ shift and add to the original model parameters for editing (replacing the origin entity y with the alternative entity a). The additional parameter-based editors (b) rectify the erroneous knowledge stored in FFN by adjusting its predicted distributions from $FFN(H; W)$ to $FFN'(H)$. KGEitor (c) utilizes a hyper external network to update the knowledge in FFN (Top right: EDIT, Bottom right: ADD).

Additional Parameter-based Editor Since previous study (Dai et al. 2022a) illustrates that the FFNs in PLMs may store factual knowledge, it is intuitive to edit models by modifying the FFNs. This paradigm introduces extra trainable parameters within the language models. CALINET extends the FFN with additional knowledge editing parameters, consisting of several calibration memory slots. To apply CALINET to the proposed task, we leverage the same architecture with FFN but with a smaller intermediate dimension d and add its output to the original FFN output as an adjustment term to edit knowledge.

The Proposed Strong Baseline: KGEitor

Note that external model-based edit methods are flexible but have to optimize many editor parameters. In contrast, the additional parameter-based method only tunes a small number of additional parameters (0.9M) but encounters poor empirical performance due to the challenge of manipulating knowledge in PLMs. Intuitively, we capitalize on the advantages of both approaches and propose a simple yet robust baseline **KGEitor**, as shown in Figure 3, which employs additional parameters through hypernetwork. We construct an additional parameter layer with the same architecture of FFN and leverage its parameters for knowledge editing. Moreover, we leverage an additional hypernetwork to generate the additional layer. With **KGEitor**, we can optimize fewer parameters while keeping the performance of editing KG embeddings. Specifically, when computing the output of FFN, we add the output of the additional FFNs to the original FFN output as

an adjustment term for editing:

$$FFN'(H) = FFN(H; W) + \Delta FFN(H; \tilde{W} + \Delta \tilde{W}) \quad (7)$$

where W is denoted as the origin FFNs’ weight, and \tilde{W} is the weight of additional FFNs. $\Delta \tilde{W}$ is the parameters’ shift generated by the hypernetwork. Inspired by KE (Cao, Aziz, and Titov 2021), we build the hypernetworks with a bidirectional LSTM. We encode $\langle x, y, a \rangle$ and then concatenate them with a special separator to feed the bidirectional LSTM. Then, we feed the last hidden states of bidirectional LSTM into the FNN to generate a single vector h for knowledge editing. To predict the shift for the weight matrix $\tilde{W}^{n \times m}$, we leverage FNN conditioned on h which predict vectors $\alpha, \beta \in \mathbb{R}^m, \gamma, \delta \in \mathbb{R}^n$ and a scalar $\eta \in \mathbb{R}$. Formally, we have the following:

$$\Delta \tilde{W} = \sigma(\eta) \cdot \left(\hat{\alpha} \odot \nabla_W \mathcal{L}(W; x, a) + \hat{\beta} \right), \quad (8)$$

with $\hat{\alpha} = \hat{\sigma}(\alpha) \gamma^\top$ and $\hat{\beta} = \hat{\sigma}(\beta) \delta^\top$,

where $\hat{\sigma}$ refers to the Softmax function (i.e., $x \mapsto \exp(x) / \sum_i \exp(x_i)$) and σ refers to the Sigmoid function (i.e., $x \mapsto (1 + \exp(-x))^{-1}$). $\nabla_W \mathcal{L}(W; x, a)$ presents the gradient, which contains rich information regarding how to access the knowledge in W . Note that the hypernetwork ϕ parameters can linearly scale with the size of W , making the **KGEitor** efficient for editing KG embeddings in terms of computational resources and time.

Method	Params	E-FB15k237						E-WN18RR					
		Time ↓	Succ@1 ↑	Succ@3 ↑	ER_{roc} ↑	RK@3 ↑	RK_{roc} ↓	Time ↓	Succ@1 ↑	Succ@3 ↑	ER_{roc} ↑	RK@3 ↑	RK_{roc} ↓
No Model Edit													
KGE_FT	121M	0.103	0.472	0.746	0.998	0.543	0.977	0.109	0.758	0.863	0.998	0.847	0.746
KGE_ZSL	0M	0.000	0.000	0.000	-	1.000	0.000	0.000	0.000	0.000	-	1.000	0.000
K-Adapter	32.5M	0.056	0.329	0.348	0.926	0.001	0.999	0.061	0.638	0.752	0.992	0.009	0.999
Model Edit Method													
CALINET	0.9M	<u>0.257</u>	0.328	0.348	0.937	0.353	0.997	<u>0.238</u>	0.538	0.649	0.991	0.446	0.994
KE	88.9M	0.368	0.702	<u>0.969</u>	0.999	0.912	<u>0.685</u>	0.386	0.599	0.682	<u>0.978</u>	0.935	0.041
MEND	59.1M	0.280	<u>0.828</u>	0.950	<u>0.954</u>	0.750	0.993	0.260	<u>0.815</u>	<u>0.827</u>	0.948	0.957	0.772
KGEditor	38.9M	0.226	0.866	0.986	0.999	<u>0.874</u>	0.635	0.232	0.833	0.844	0.991	<u>0.956</u>	<u>0.256</u>

Table 2: The main result of the EDIT task on E-FB15k237 and E-WN18RR. Bold indicates the best result among *Model Edit Method*, while underline represents the second-best result. The same format applies to Table 3.

Method	Params	A-FB15k237						A-WN18RR					
		Time ↓	Succ@1 ↑	Succ@3 ↑	ER_{roc} ↑	RK@3 ↑	RK_{roc} ↓	Time ↓	Succ@1 ↑	Succ@3 ↑	ER_{roc} ↑	RK@3 ↑	RK_{roc} ↓
No Model Edit													
KGE_FT	121M	0.100	0.906	0.976	0.999	0.223	0.997	0.108	0.997	0.999	0.999	0.554	0.996
KGE_ZSL	0M	0.000	0.000	0.000	-	1.000	0.000	0.000	0.000	0.000	-	1.000	0.000
K-Adapter	32.5M	0.055	0.871	0.981	0.999	0.000	0.999	0.061	0.898	0.978	0.999	0.002	0.999
Model Edit Method													
CALINET	0.9M	<u>0.261</u>	<u>0.714</u>	0.870	<u>0.997</u>	0.034	0.999	<u>0.275</u>	0.832	0.913	<u>0.995</u>	0.511	0.989
KE	88.9M	0.362	0.648	0.884	<u>0.997</u>	0.926	<u>0.971</u>	0.384	0.986	<u>0.996</u>	0.999	0.975	0.090
MEND	59.1M	0.400	0.517	0.745	0.991	0.499	0.977	0.350	0.999	1.0	0.999	0.810	0.987
KGEditor	58.7M	0.203	0.796	0.923	0.998	<u>0.899</u>	0.920	0.203	<u>0.998</u>	1.0	0.999	<u>0.956</u>	<u>0.300</u>

Table 3: The main result of the ADD task on A-FB15k237 and A-WN18RR.

Evaluation

Settings

We employ the constructed datasets, E-FB15k237, A-FB15k237, E-WN18RR, and A-WN18RR for evaluations. Initially, we utilize the pre-train datasets (the identical training set for the ADD task and the corrupted training set for the EDIT task) to initialize the KG embeddings with BERT. We then train the editor using the training dataset and evaluate using the testing dataset. Regarding the ADD task, the evaluation is conducted based on the performance of the training set, as they are not present in the pre-training dataset. We adopt PT-KGE as the default setting for the main experiments. Additionally, We supply the datasets, pre-trained KG embeddings, and a leaderboard for reference.

Main Results

Our comparison of **KGEditor** spans multiple baselines like MEND, KE, CALINET, and variants: KGE_FT involving direct fine-tuning of all parameters, KGE_ZSL for inferring triples without parameter adjustments, and K-Adapter using an adaptor to fine-tune a subset of parameters. Our experiments aim to 1) evaluate the efficacy of various approaches for editing KG embeddings (Table 2,3), 2) investigate the performance when varying the number (n) of target triples (Figure 5), and 3) examine the impact of distinct KGE initialization approaches on performance (Figure 4).




Figure 4: Left: ER_{roc} and Right: RK_{roc} . The performance of different KGE initializations with different models.

EDIT&ADD Task As Table 2 shows, KGE_ZSL fails to infer any facts on both datasets for the EDIT task. Intriguingly, we also find that fine-tuning all (KGE_FT) or a subset (K-Adapter) of parameters does not result in satisfactory performance. We think this is due to the numerous N - M facts, and fine-tuning with more data may lead to confusion about the knowledge. We find that well-crafted knowledge editing baselines surpass KGE_FT and K-Adapter, indicating the need for specialized architectures for editing knowledge to more effectively identify accurate facts. Furthermore, we notice that CALIET exhibits poor knowledge

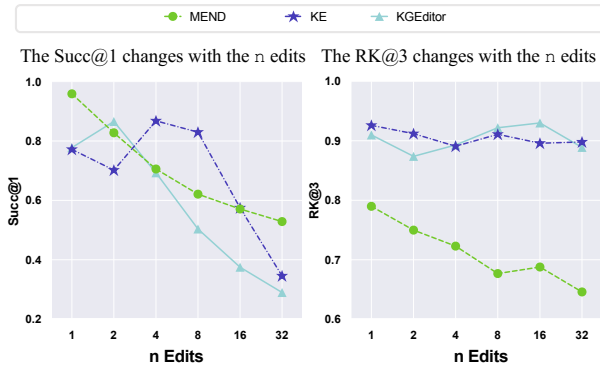


Figure 5: Left: the variation of Succ@1 value with n edits. All models drop when n edits become larger. Right: the RK@3 score drops when n changes, KE and KGEitor can remain stable performance, while MEND decreases.

locality (RK@3 and RK_{roc}), indicating that merely modifying a few parameters in FFNs cannot guarantee precise knowledge editing. Our model **KGEitor** surpasses nearly all baselines, utilizing fewer tunable parameters than KE and MEND. Regarding time and resource expenditure, **KGEitor** exhibits markedly enhanced editing efficiency compared to KGE_FT. Note that our model employs a hypernetwork to guide the parameter editing of the FFN, resulting in greater parameter efficiency than KE and MEND while simultaneously achieving superior performance to CALINET.

For the ADD task, as shown in Table 3, we observe that KGE_ZSL continues to fail in inferring any facts on both datasets (We filter all facts that can be directly inferred by probing the PLM). Finetuning either all parameters (KGE_FT) or a subset (K-Adaptor) maximizes the performance of reliability, but compromises previously acquired knowledge, reflecting poor knowledge locality. Given that the ADD task introduces unseen facts during KG embedding initialization, finetuning parameters naturally capture new knowledge while sacrificing prior information.

FT-KGE&PT-KGE We evaluate various KGE initialization methods with different models on the EDIT task. Figure 4 contrasts the performance of FT-KGE and PT-KGE in terms of knowledge reliability and locality. We observe that the editing performance of PT-KGE is superior to that of FT-KGE, suggesting that employing prompt-based models is more suitable for editing KG embeddings. Prior research (Liu et al. 2021a; Meng et al. 2022) has demonstrated that the prompt-based approaches are more effective in harnessing the “modeledge” (Han and et al. 2021) in PLMs, making it reasonable to utilize PT-KGE for editing purposes.

Number of Edited Knowledge We further analyze the impact of varying the number of edited facts. We employ different numbers of edits in the E-FB15k237, with the number of edits $n \in \{1, 2, 4, 8, 16, 32\}$. Figure 5 illustrates a pronounced impact of n edits on all models’ knowledge reliability, measured by Succ@1. Furthermore, we observe that KE and **KGEitor** show stable performance in knowledge




Figure 6: We edit the fact (Mila Kunis, type of union, Domestic partnership) and replace the head entity by Marc Shaiman (Golden Label). The points close to the center refer to the entities the model prefers. Left: predictions of the original KG embeddings. Right: predictions after editing.

locality, unlike MEND, which sees a notable decline. We attribute this to KE and **KGEitor** employing FFN (termed as knowledge neurons (Dai et al. 2022b)), which can be scaled to accommodate a higher number of edited facts.

We visualize entities before and after editing for clearer model insights. Figure 6 illustrates a substantial shift in predicted entities’ positions before and after the model editor’s application. These points signify that the inferred entity is provided with a head entity and relation, with the right center representing the golden standard. Upon editing the model, the correct entity distinctly emerges close to the circle’s center (knowledge reliability), while the relevant distances of other entities remain largely unaltered, thereby showcasing the effectiveness of editing KG embeddings.

Discussion and Conclusion

Our proposed task of editing the KGE model allows for direct modification of knowledge to suit specific tasks, thereby improving the efficiency and accuracy of the editing process. Contrary to earlier pre-training language model tasks for editing, our approach relies on KG facts to modify knowledge, without using pre-trained model knowledge. These methods enhance performance and offer important insights for research in knowledge representation and reasoning.

Future Work

Editing KGE models presents ongoing issues, notably handling intricate knowledge and many-to-many relations. Edited facts, arising from such relations, can bias the model to the edited entity, overlooking other valid entities. Besides, the experimental KGE models are small, all leveraging the standard BERT. Yet, with rising large-scale generative models (LLMs) like LLaMA (Touvron et al. 2023), ChatGLM (Zeng et al. 2022), and Alpaca (Taori et al. 2023), the demand to edit LLMs grows. In the future, we aim to design models to edit knowledge with many-to-many relations and integrate LLMs editing techniques.

Acknowledgments

We thank the anonymous reviewers for their kind comments. This work was supported by the National Natural Science Foundation of China (No. 62206246), the Fundamental Research Funds for the Central Universities (226-2023-00138), Zhejiang Provincial Natural Science Foundation of China (No. LGG22F030011), Ningbo Natural Science Foundation (2021J190), Yongjiang Talent Introduction Programme (2021A-156-G), CCF-Tencent Rhino-Bird Open Research Fund, and Information Technology Center and State Key Lab of CAD&CG, Zhejiang University.

References

- Bordes, A.; Usunier, N.; García-Durán, A.; Weston, J.; and Yakhnenko, O. 2013. Translating Embeddings for Modeling Multi-relational Data. In *NeurIPS*.
- Brown, T. B.; and et al. 2020. Language Models are Few-Shot Learners. In *NeurIPS*.
- Cao, B.; Lin, H.; Han, X.; Sun, L.; Yan, L.; Liao, M.; Xue, T.; and Xu, J. 2021. Knowledgeable or Educated Guess? Revisiting Language Models as Knowledge Bases. In *ACL*.
- Cao, N. D.; Aziz, W.; and Titov, I. 2021. Editing Factual Knowledge in Language Models. In *EMNLP*.
- Chen, C.; Wang, Y.; Li, B.; and Lam, K. 2022a. Knowledge Is Flat: A Seq2Seq Generative Framework for Various Knowledge Graph Completion. In *COLING*.
- Chen, X.; Zhang, N.; Xie, X.; Deng, S.; Yao, Y.; Tan, C.; Huang, F.; Si, L.; and Chen, H. 2022b. Know-Prompt: Knowledge-aware Prompt-tuning with Synergistic Optimization for Relation Extraction. In *WWW*.
- Cohen, R.; Biran, E.; Yoran, O.; Globerson, A.; and Geva, M. 2023. Evaluating the Ripple Effects of Knowledge Editing in Language Models. *CoRR*, abs/2307.12976.
- Dai, D.; Dong, L.; Hao, Y.; Sui, Z.; Chang, B.; and Wei, F. 2022a. Knowledge Neurons in Pretrained Transformers. In *ACL*.
- Dai, D.; Dong, L.; Hao, Y.; Sui, Z.; Chang, B.; and Wei, F. 2022b. Knowledge Neurons in Pretrained Transformers. In *ACL*.
- Dettmers, T.; Minervini, P.; Stenetorp, P.; and Riedel, S. 2018. Convolutional 2D Knowledge Graph Embeddings. In *AAAI*.
- Devlin, J.; Chang, M.; Lee, K.; and Toutanova, K. 2019. BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding. In Burstein, J.; Doran, C.; and Solorio, T., eds., *NAACL*.
- Dong, Q.; Dai, D.; Song, Y.; Xu, J.; Sui, Z.; and Li, L. 2022. Calibrating Factual Knowledge in Pretrained Language Models. In *EMNLP*, Findings of EMNLP.
- Fei, L.; Wu, T.; and Khan, A. 2021. Online Updates of Knowledge Graph Embedding. In *Complex Networks 2021*, volume 1016 of *Studies in Computational Intelligence*, 523–535. Springer.
- Geva, M.; Bastings, J.; Filippova, K.; and Globerson, A. 2023. Dissecting Recall of Factual Associations in Auto-Regressive Language Models. *CoRR*, abs/2304.14767.
- Han, X.; and et al. 2021. Pre-trained models: Past, present and future. *AI Open*.
- Han, X.; Li, R.; Li, X.; and Pan, J. Z. 2023. A divide and conquer framework for Knowledge Editing. *Knowledge-Based Systems*, 110826.
- Hase, P.; Bansal, M.; Kim, B.; and Ghandeharioun, A. 2023. Does Localization Inform Editing? Surprising Differences in Causality-Based Localization vs. Knowledge Editing in Language Models. *CoRR*, abs/2301.04213.
- Kipf, T. N.; and Welling, M. 2017. Semi-Supervised Classification with Graph Convolutional Networks. In *ICLR*.
- Lewis, M.; Liu, Y.; Goyal, N.; Ghazvininejad, M.; Mohamed, A.; Levy, O.; Stoyanov, V.; and Zettlemoyer, L. 2020. BART: Denoising Sequence-to-Sequence Pre-training for Natural Language Generation, Translation, and Comprehension. In *ACL*.
- Lin, Y.; Liu, Z.; Sun, M.; Liu, Y.; and Zhu, X. 2015. Learning Entity and Relation Embeddings for Knowledge Graph Completion. In Bonet, B.; and Koenig, S., eds., *AAAI*.
- Liu, P.; Yuan, W.; Fu, J.; Jiang, Z.; Hayashi, H.; and Neubig, G. 2021a. Pre-train, prompt, and predict: A systematic survey of prompting methods in natural language processing. *ACM Computing Surveys*.
- Liu, S.; Grau, B. C.; Horrocks, I.; and Kostylev, E. V. 2021b. INDIGO: GNN-Based Inductive Knowledge Graph Completion Using Pair-Wise Encoding. In *NeurIPS*, 2034–2045.
- Liu, X.; Ji, K.; Fu, Y.; Tam, W.; Du, Z.; Yang, Z.; and Tang, J. 2022. P-Tuning: Prompt Tuning Can Be Comparable to Fine-tuning Across Scales and Tasks. In *ACL*.
- Lv, X.; Lin, Y.; Cao, Y.; Hou, L.; Li, J.; Liu, Z.; Li, P.; and Zhou, J. 2022. Do Pre-trained Models Benefit Knowledge Graph Completion? A Reliable Evaluation and a Reasonable Approach. In *Findings of ACL*.
- Meng, K.; Bau, D.; Andonian, A.; and Belinkov, Y. 2022. Locating and Editing Factual Knowledge in GPT. In *NeurIPS*.
- Mitchell, E.; Lin, C.; Bosselut, A.; Finn, C.; and Manning, C. D. 2022. Fast Model Editing at Scale. In *ICLR*.
- Nickel, M.; Rosasco, L.; and Poggio, T. A. 2016. Holographic Embeddings of Knowledge Graphs. In *AAAI*.
- Nickel, M.; Tresp, V.; and Kriegel, H. 2011. A Three-Way Model for Collective Learning on Multi-Relational Data. In *ICML*.
- Pan, S.; Luo, L.; Wang, Y.; Chen, C.; Wang, J.; and Wu, X. 2023. Unifying Large Language Models and Knowledge Graphs: A Roadmap. *CoRR*, abs/2306.08302.
- Petroni, F.; Rocktäschel, T.; Riedel, S.; Lewis, P. S. H.; Bakhtin, A.; Wu, Y.; and Miller, A. H. 2019. Language Models as Knowledge Bases? In *EMNLP*.
- Raffel, C.; Shazeer, N.; Roberts, A.; Lee, K.; Narang, S.; Matena, M.; Zhou, Y.; Li, W.; and Liu, P. J. 2020. Exploring the Limits of Transfer Learning with a Unified Text-to-Text Transformer. *J. Mach. Learn. Res.*, 21: 140:1–140:67.
- Saxena, A.; Kochsiek, A.; and Gemulla, R. 2022. Sequence-to-Sequence Knowledge Graph Completion and Question Answering. In *ACL*.

- Sinitin, A.; Plokhotnyuk, V.; Pyrkov, D. V.; Popov, S.; and Babenko, A. 2020. Editable Neural Networks. In *ICLR*.
- Sun, Z.; Deng, Z.; Nie, J.; and Tang, J. 2019. RotatE: Knowledge Graph Embedding by Relational Rotation in Complex Space. In *ICLR*.
- Taori, R.; Gulrajani, I.; Zhang, T.; Dubois, Y.; Li, X.; Guestrin, C.; Liang, P.; and Hashimoto, T. B. 2023. Stanford Alpaca: An Instruction-following LLaMA model. https://github.com/tatsu-lab/stanford_alpaca. Accessed: 2023-03-13.
- Toutanova, K.; Chen, D.; Pantel, P.; Poon, H.; Choudhury, P.; and Gamon, M. 2015. Representing Text for Joint Embedding of Text and Knowledge Bases. In *EMNLP*.
- Touvron, H.; Lavril, T.; Izacard, G.; Martinet, X.; Lachaux, M.; Lacroix, T.; Rozière, B.; Goyal, N.; Hambro, E.; Azhar, F.; Rodriguez, A.; Joulin, A.; Grave, E.; and Lample, G. 2023. LLaMA: Open and Efficient Foundation Language Models. *CoRR*, abs/2302.13971.
- Vashishth, S.; Sanyal, S.; Nitin, V.; and Talukdar, P. P. 2020. Composition-based Multi-Relational Graph Convolutional Networks. In *ICLR*.
- Velickovic, P.; Cucurull, G.; Casanova, A.; Romero, A.; Liò, P.; and Bengio, Y. 2018. Graph Attention Networks. In *ICLR*.
- Wang, B.; Shen, T.; Long, G.; Zhou, T.; Wang, Y.; and Chang, Y. 2021a. Structure-Augmented Text Representation Learning for Efficient Knowledge Graph Completion. In *WWW*.
- Wang, L.; Zhao, W.; Wei, Z.; and Liu, J. 2022a. SimKGC: Simple Contrastive Knowledge Graph Completion with Pre-trained Language Models. In *ACL*.
- Wang, Q.; Mao, Z.; Wang, B.; and Guo, L. 2017. Knowledge Graph Embedding: A Survey of Approaches and Applications. *TKDE*.
- Wang, R.; Tang, D.; Duan, N.; Wei, Z.; Huang, X.; Ji, J.; Cao, G.; Jiang, D.; and Zhou, M. 2021b. K-Adapter: Infusing Knowledge into Pre-Trained Models with Adapters. In *Findings of ACL*.
- Wang, X.; Gao, T.; Zhu, Z.; Zhang, Z.; Liu, Z.; Li, J.; and Tang, J. 2021c. KEPLER: A Unified Model for Knowledge Embedding and Pre-trained Language Representation. *Trans. Assoc. Comput. Linguistics*, 9: 176–194.
- Wang, X.; He, Q.; Liang, J.; and Xiao, Y. 2022b. Language Models as Knowledge Embeddings. In *IJCAI*.
- Wang, Z.; Zhang, J.; Feng, J.; and Chen, Z. 2014. Knowledge Graph Embedding by Translating on Hyperplanes. In *AAAI*.
- Wei, Y.; Chen, W.; Li, Z.; and Zhao, L. 2021. Incremental Update of Knowledge Graph Embedding by Rotating on Hyperplanes. In *ICWS 2021*, 516–524. IEEE.
- Wu, T.; Shiri, F.; Kang, J.; Qi, G.; Haffari, G.; and Li, Y.-F. 2022a. KC-GEE: Knowledge-based Conditioning for Generative Event Extraction.
- Wu, T.; Wang, G.; Zhao, J.; Liu, Z.; Qi, G.; Li, Y.; and Haffari, G. 2022b. Towards relation extraction from speech. In Goldberg, Y.; Kozareva, Z.; and Zhang, Y., eds., *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing, EMNLP 2022, Abu Dhabi, United Arab Emirates, December 7-11, 2022*, 10751–10762. Association for Computational Linguistics.
- Xie, R.; Liu, Z.; Jia, J.; Luan, H.; and Sun, M. 2016. Representation Learning of Knowledge Graphs with Entity Descriptions. In Schuurmans, D.; and Wellman, M. P., eds., *AAAI*.
- Xie, X.; Zhang, N.; Li, Z.; Deng, S.; Chen, H.; Xiong, F.; Chen, M.; and Chen, H. 2022. From Discrimination to Generation: Knowledge Graph Completion with Generative Transformer. In *WWW*.
- Yang, B.; Yih, W.; He, X.; Gao, J.; and Deng, L. 2015. Embedding Entities and Relations for Learning and Inference in Knowledge Bases. In *ICLR*.
- Yang, Z.; Ma, J.; Chen, H.; Zhang, J.; and Chang, Y. 2022. Context-Aware Attentive Multilevel Feature Fusion for Named Entity Recognition. *IEEE Transactions on Neural Networks and Learning Systems*.
- Yang, Z.; Ma, J.; Chen, H.; Zhang, Y.; and Chang, Y. 2021. HiTRANS: A Hierarchical Transformer Network for Nested Named Entity Recognition. In *Findings of ACL*.
- Yao, L.; Mao, C.; and Luo, Y. 2019. KG-BERT: BERT for Knowledge Graph Completion. *CoRR*, abs/1909.03193.
- Yao, Y.; Wang, P.; Tian, B.; Cheng, S.; Li, Z.; Deng, S.; Chen, H.; and Zhang, N. 2023. Editing Large Language Models: Problems, Methods, and Opportunities. *CoRR*, abs/2305.13172.
- Zeng, A.; Liu, X.; Du, Z.; Wang, Z.; Lai, H.; Ding, M.; Yang, Z.; Xu, Y.; Zheng, W.; Xia, X.; Tam, W. L.; Ma, Z.; Xue, Y.; Zhai, J.; Chen, W.; Zhang, P.; Dong, Y.; and Tang, J. 2022. GLM-130B: An Open Bilingual Pre-trained Model. *CoRR*, abs/2210.02414.
- Zhang, N.; Bi, Z.; Liang, X.; Cheng, S.; Hong, H.; Deng, S.; Zhang, Q.; Lian, J.; and Chen, H. 2022a. OntoProtein: Protein Pretraining With Gene Ontology Embedding. In *ICLR*.
- Zhang, N.; Deng, S.; Sun, Z.; Chen, J.; Zhang, W.; and Chen, H. 2020a. Relation Adversarial Network for Low Resource Knowledge Graph Completion. In *WWW*.
- Zhang, N.; Li, L.; and et al. 2022. Differentiable Prompt Makes Pre-trained Language Models Better Few-shot Learners. In *ICLR*.
- Zhang, Z.; Cai, J.; Zhang, Y.; and Wang, J. 2020b. Learning Hierarchy-Aware Knowledge Graph Embeddings for Link Prediction. In *AAAI*.
- Zhang, Z.; Liu, X.; Zhang, Y.; Su, Q.; Sun, X.; and He, B. 2020c. Pretrain-KGE: Learning Knowledge Representation from Pretrained Language Models. In *Findings of EMNLP*.
- Zhang, Z.; Wang, J.; Ye, J.; and Wu, F. 2022b. Rethinking Graph Convolutional Networks in Knowledge Graph Completion. In *WWW*.
- Zhao, W. X.; Zhou, K.; Li, J.; and et al. 2023. A Survey of Large Language Models. *CoRR*, abs/2303.18223.