

Filter-then-Generate: Large Language Models with Structure-Text Adapter for Knowledge Graph Completion

Ben Liu^{1†*}, Jihai Zhang^{2†}, Fangquan Lin², Cheng Yang², Min Peng^{1‡}

¹School of Computer Science, Wuhan University, China

²DAMO Academy, Alibaba Group, Hangzhou, 310023, China

{liuben123, pengm}@whu.edu.cn

{jihai.zjh, fangquan.linfq, charis.yangc}@alibaba-inc.com

Abstract

Large Language Models (LLMs) present massive inherent knowledge and superior semantic comprehension capability, which have revolutionized various tasks in natural language processing. Despite their success, a critical gap remains in enabling LLMs to perform knowledge graph completion (KGC). Empirical evidence suggests that LLMs consistently perform worse than conventional KGC approaches, even through sophisticated prompt design or tailored instruction-tuning. Fundamentally, applying LLMs on KGC introduces several critical challenges, including a vast set of entity candidates, hallucination issue of LLMs, and under-exploitation of the graph structure. To address these challenges, we propose a novel instruction-tuning-based method, namely FtG. Specifically, we present a *filter-then-generate* paradigm and formulate the KGC task into a multiple-choice question format. In this way, we can harness the capability of LLMs while mitigating the issue caused by hallucinations. Moreover, we devise a flexible ego-graph serialization prompt and employ a structure-text adapter to couple structure and text information in a contextualized manner. Experimental results demonstrate that FtG achieves substantial performance gain compared to existing state-of-the-art methods. The instruction dataset and code are available at <https://github.com/LB0828/FtG>.

1 Introduction

Knowledge graphs (KGs) encode and store abundant factual knowledge in the format of triples like (*head entity, relation, tail entity*), which provide faithful knowledge source for downstream knowledge-intensive tasks (Pan et al., 2023; Luo et al., 2023b). However, due to the evolving nature

*This work is conducted when Ben Liu was interning at DAMO Academy, Alibaba Group

† Equal contribution.

‡ Corresponding author

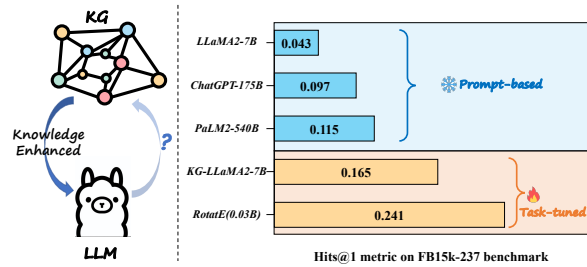


Figure 1: Unsatisfactory performance (Hits@1 metric) of LLMs on the test set of FB15k-237 compared to conventional KGC method RotatE (Sun et al., 2019).

of knowledge, real-world KGs often suffer from incompleteness, urging auto-completion of them. Therefore, knowledge graph completion (KGC), which aims to infer missing triples from the existing KG, has been a fundamental and challenging problem in artificial intelligence research.

Recently, large language models (LLMs) have demonstrated impressive memorization and reasoning abilities through pre-training on massive text corpora (Zhang et al., 2024; He et al., 2024). Despite their success, LLMs are limited in insufficient knowledge and prone to generate hallucinations (Huang et al., 2023; Li et al., 2023). Multiple attempts that exploit KGs to address the hallucinations issues of LLMs have demonstrated promising achievements on various natural language processing tasks (Luo et al., 2023b; Xu et al., 2024). However, the integrating LLMs with KGs for KGC remains underexplored, which present a significant opportunity to revolutionize the approach to KGC by leveraging the contextual understanding and reasoning abilities of LLMs.

Zhang et al. (2023) and Shu et al. (2024) apply LLMs to relatively simple KG triple classification (i.e., determining if it is true) has achieved promising results. For more practical and challenging KGC task, Zhu et al. (2023) exploit ChatGPT (OpenAI, 2023) with in-context learning (Wei et al.,

2022b) to transform the KGC task into a text-based prediction. However, contrary to expectations, ChatGPT, equipped with hundreds of billions of parameters, lags behind the performance achieved by conventional structure-based methods (e.g., RotatE (Sun et al., 2019)). As shown in Figure 1, on widely-used dataset FB15k-237 (Toutanova and Chen, 2015), ChatGPT only achieves 0.097 on the Hits@1 metric while RotatE achieves 0.241. Furthermore, KG-LLaMA (Yao et al., 2023) utilizes instruction tuning (Wei et al., 2022a) to adapt LLaMA2-7b (Touvron et al., 2023) to KGC task, the performance achieved still is not satisfactory (0.165 Hits@1 on FB15k-237). Considering the powerful abilities of LLMs, yet they perform worse than conventional methods on KGC task, a pertinent question arises: *Why LLMs cannot present satisfactory performance on KGC task?*

We speculate LLMs struggle in the KGC task from two aspects: (1) **Large entity candidate set.** The KGC task can essentially be regarded as a classification task, where the label is all entities in KG. Recent works (Ma et al., 2023) evaluating the performance of LLMs on classification tasks have shown that LLMs struggle in datasets with a large number of labels. The result is consistent with Yao et al. (2023), where LLMs perform relatively better on the KG with a smaller entity set. The enormity of the entity candidate set in the KG poses a challenge for LLMs. Furthermore, LLMs are known for their propensity to hallucinate content, generating information not grounded by world knowledge (Ji et al., 2023). Directly applying LLMs to elicit missing entities from the entire token space often leads to generating invalid entities outside the entity set. (2) **Inherent Graph Structure of KGs.** Different from text-based NLP tasks, KGs demonstrate unorganized and complicated graph structures. Current efforts (Fan et al., 2024) have revealed that performance of LLMs on fundamental graph structural understanding tasks is subpar. Effectively guiding LLMs to comprehend the structural information of KGs remains a considerable challenge.

Building on these findings, we propose a novel instruction-tuning based method, namely FtG, which harnesses and unleashes the capability of LLMs for KGC task. To address the issues of enormous entity candidate set, we present a *filter-then-generate* paradigm, where FtG first employs a conventional KGC method as the filter to eliminate unlikely entities and retain only the top- k candidates. Then we formulate the KGC task into

a multiple-choice question format and construct instruction template to prompt LLMs to generate target entity from the top- k candidates. Essentially, this paradigm mirrors human behavior. For example, when answering a question, humans would eliminate obviously wrong answers and find the answer from few remaining candidates. In this way, we narrow the candidate set significantly, and the multiple-choice question format effectively avoids LLMs to output uncontrollable text. Moreover, to incorporate graph structure information into LLMs, we devise a structure-aware ego-graph serialization prompt and propose a light-weight structure-text adapter to map the graph features into the text space. Comparison experiment results show that our proposed FtG greatly improves performance of LLMs on KGC task. In summary, our contributions are:

- We propose a novel instruction-tuning based method FtG to enhance the performance of LLMs in the KGC task. In which, our proposed *filter-then-generate* paradigm can effectively harness the powerful capability of LLMs while mitigating the issues caused by hallucinations.
- To bridge the gap between the graph structure and LLMs, we introduce a structure-aware ego-graph prompt and devise a structure-text adapter in a contextualized way.
- Our FtG outperforms the state-of-the-art methods over three widely used benchmarks, which demonstrates our model’s superiority. Further results reveal that our FtG paradigm can enhance existing KGC methods in a plug-and-play manner.

2 Preliminaries

Knowledge Graph Completion (KGC). Knowledge graph (KG) is commonly composed of a collection of fact triples. Let $\mathcal{G} = (\mathcal{E}, \mathcal{R}, \mathcal{T})$ be a KG instance, where \mathcal{E} , \mathcal{R} , and \mathcal{T} represent the set of entities, relations, and triples, respectively. Each triple $(h, r, t) \in \mathcal{E} \times \mathcal{R} \times \mathcal{E}$ describes the fact that a relation r exists between head entity h and tail entity t . Given an incomplete triple $(h, r, ?)$ or $(?, r, t)$ as query q , knowledge graph completion aims to predict the missing tail or head entity. In conventional KGC models, they learn specific structural embeddings for KGs, and the missing entity is predicted by finding the highest score $f(h, r, e)$

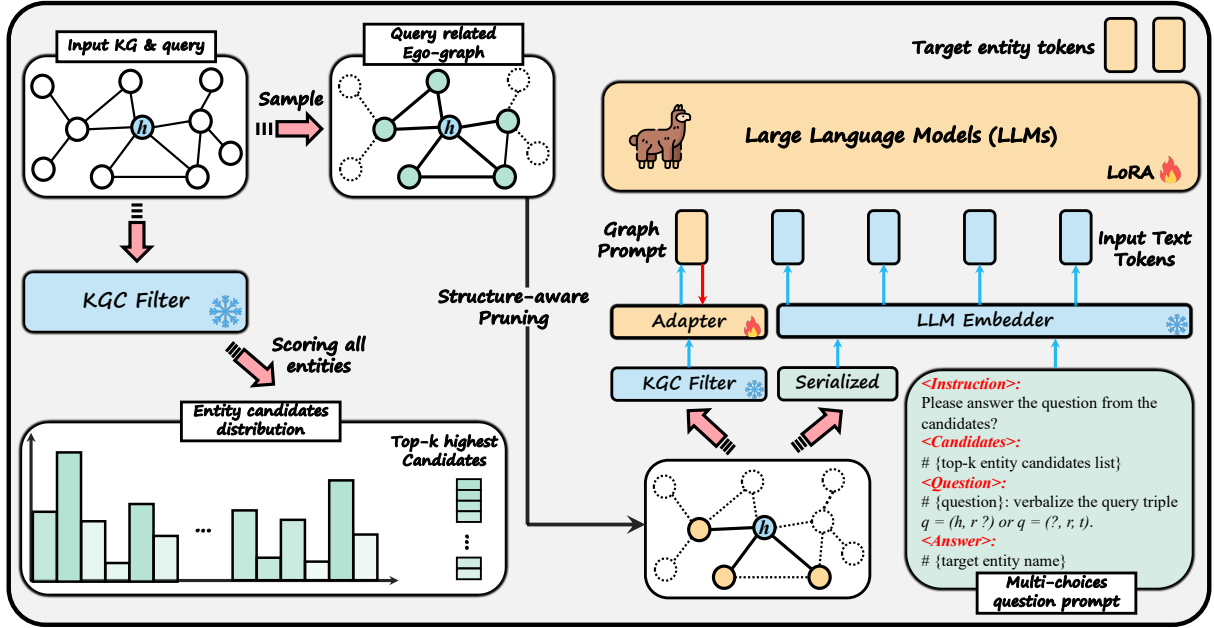


Figure 2: The overall framework of FtG. For a query triple, we employ a KGC filter to obtain top- k candidates and construct corresponding multiple-choice question instruction. Then we sample the ego-graph of query entity and prune irrelevant neighbors. To bridge the gap between graph and text, we encode the pruned ego-graph into a soft graph token and map the graph token into text embedding space with a lightweight adapter. The target entity name is generated with soft graph token, textual serialization of pruned ego-graph, and instruction prompt.

or $f(e, r, t), \forall e \in \mathcal{E}$, where f is the model-specific scoring function.

Instruction Tuning for KGC. Instruction tuning refers to fine-tuning LLMs to follow human-curated instructions, enabling adaptation of LLMs to specific tasks. When applying LLMs to the KGC task, an instruction tuning sample comprises an instruction prompt and input-output pair. The instruction prompt \mathcal{I} (e.g., "Predict the missing tail entity") is definition of KGC task for LLMs to comprehend and execute. The input X is the verbalization of the query q described in natural language. The instruction tuning process aims to strictly generate the missing entity in natural language given the instruction and the query input: $Y = \text{LLM}_\theta(\mathcal{I}, X)$, where θ are the parameters of LLM. The prevalent Negative Log-Likelihood Loss in language modeling is selected as training objective, which can be formed as:

$$\mathcal{L}(\theta) = - \sum_{i=1}^L \log P_\theta(Y_i | \mathcal{I}, X, Y_{<i}), \quad (1)$$

where $Y_{<i}$ represents the prefix of missing entity name sequence Y up to position $i - 1$, $P_\theta(Y_i | \mathcal{I}, X, Y_{<i})$ represents the probability of generating token Y_i , and L is the sequence length of Y .

3 Methodology

In this section, we first provide an in-depth description of our *filter-then-generate* paradigm in Sec. 3.1, specifically designed for KGC task. Based on such paradigm, to bridge the gap between graph structure and text, we further introduce two novel modules: 1) a flexible ego-graph serialization prompt in Sec. 3.2, which can effectively convey the structural information around the query triple, and 2) a structure-text prefix adapter in Sec. 3.3, to map graph structure features into the text embedding space. Finally, we detail the instruction tuning strategy in Sec. 3.4, focusing on efficient adaptation to KGC task. The overall architecture of our proposed model is illustrated in Figure 2.

3.1 Filter-then-Generate Paradigm

To address the challenge of large entity candidate set, we propose a novel *filter-then-generate* paradigm for LLMs on KGC task. As its name implies, we utilize a filter to eliminate unlikely entities and retain only the top- k candidates. LLMs then generate the target entities conditioned on the query and candidates list. Within our paradigm, given a query $q = (h, r, ?)$ or $q = (?, r, t)$, we employ a conventional structure-based KGC method as filter to score each entity e in KG \mathcal{G} and re-

tain top- k highest scoring entities as candidate set $C_k = [e_1, e_2, e_3, \dots, e_k]$. The idea behind this paradigm is that conventional shallow embedding based models are good at easy samples but fail to discriminate target entities from a set of hard samples. Then we reformulate the KGC task into the form of multiple-choice question and design a simple instruction template that prompts LLMs to generate answer:

Instruction: Please answer the following question and select only one answer from the candidates that is most relevant to the question.

Question: <Verbalization of query q >

[Optional] Context: <Ego-graph serialization prompts>

Candidates: <Candidate set C_k >

Answer: <Target entity name>

In above instruction template, we adopt the same verbalization as Yao et al. (2023) to transform the query triple into a simple question, and the context is optional, which conveys the structural information around the query described in following section. Essentially, our *filter-then-generate* pipeline demonstrates following advantages. Firstly, it can rescue LLMs from a large number of candidate entities, enabling LLMs to leverage their own knowledge and reasoning abilities to identify the target entity from a group of candidate entities, which go beyond the abilities of conventional methods but could be well solved by LLMs. Additionally, the multiple-choice question format effectively avoids LLMs to output uncontrollable text.

3.2 Ego-graph Serialization Prompt

In our paradigm, we aim to exploit semantic comprehension and reasoning capability of LLMs for KGC task. Nevertheless, transforming the query triple into a text-based prediction inevitably neglects the structural information of KG, which is an important feature for KGC task. Moreover, understanding graph structures using LLM continues to be a challenge, and although there has been some exploration into designing prompts to convey structural information, a comprehensive solution is still lacking.

To incorporate the structural information of KG into LLMs, we design an ego-graph serialization prompt. Instead of accessing to the entire KG, we

extract the 1-hop ego-graph (Wang et al., 2019) around the query entity, which characterizes the first order connectivity structure of entity. Considering that not all neighborhoods are useful for query, and some of them even introduce additional noise, we employ structure embeddings of KGs to sample more informative neighbors. Specifically, given a KG \mathcal{G} and the query triple $q = (h, r, ?)$ under the tail entity prediction setting (same in head entity prediction), we first sample the both incoming and outgoing triples of h as the 1-hop ego-graph $\mathcal{N}_h = \{(h, r', e) \in \mathcal{G}\} \cup \{(e, r', h) \in \mathcal{G}\}$. Then let $\mathbf{E} \in \mathbb{R}^{|\mathcal{E}| \times d_s}$ and $\mathbf{R} \in \mathbb{R}^{|\mathcal{R}| \times d_s}$ denote the structural entity embedding matrix and relation embedding matrix, respectively, and d_s is structural embedding dimension. The structure embedding matrices are provided by the KGC model adopted as the filter of FtG. We take in query $q = (h, r, ?)$ and the 1-hop ego-graph \mathcal{N}_h to extract the most relevant neighbors $\widetilde{\mathcal{N}}_h$ as follows:

$$\begin{aligned} \widetilde{\mathcal{N}}_h &= \{(h', r', t') \mid (h', r', t') \in \mathcal{N}_h \\ &\quad \text{and } \cos(\mathbf{h}' \parallel \mathbf{r}', \mathbf{h} \parallel \mathbf{r}) > \epsilon\} \end{aligned} \quad (2)$$

where $\mathbf{h} \in \mathbb{R}^{1 \times d_s}$ and $\mathbf{r} \in \mathbb{R}^{1 \times d_s}$ represent the structural embedding of h and r , $\cos(\cdot, \cdot)$ is the cosine similarity, ϵ is the threshold, and \parallel is the concatenation operation.

After obtaining the extracted ego-graph $\widetilde{\mathcal{N}}_h$, we follow existing work (Jiang et al., 2023a) perform breadth-first search (BFS) serialization to linearize it into a textual sentence. Specifically, starting from the entity h , we perform BFS to visit all entities in $\widetilde{\mathcal{N}}_h$. We then concatenate all the visited triples in the order of their traversal during the BFS process and remove duplicate entities, resulting in a long sequence, denotes as:

$$S_{\widetilde{\mathcal{N}}_h} = \{h, r_1, e_1, r_2, e_2, \dots, r_m, e_m\}, \quad (3)$$

where m is the number of triples in $\widetilde{\mathcal{N}}_h$.

3.3 Structure-Text Adapter

Graph Encoding and Adaption. While our ego-graph serialization prompt has captured the local structure information around the query, the linearization process inevitably loses the connective pattern of the graph. Therefore, we propose a soft prompt strategy to couple the KG structure and text information in a contextualized way. Given the pruned ego-graph $\widetilde{\mathcal{N}}_h$, we obtain the ego-graph representation through *parameter-free*

message passing on encoded structure features, and map the graph representation into the embedding space of LLM via a trainable projection matrix $\mathbf{W}_p \in \mathbb{R}^{d_s \times d_x}$:

$$\mathbf{S}_{\tilde{\mathcal{N}}_h} = \frac{1}{|\tilde{\mathcal{N}}_h|} \sum_{e' \in \tilde{\mathcal{N}}_h} e', \quad \mathbf{S}'_{\tilde{\mathcal{N}}_h} = \mathbf{W}_p \cdot \mathbf{S}_{\tilde{\mathcal{N}}_h}, \quad (4)$$

where $\mathbf{S}'_{\tilde{\mathcal{N}}_h}$ is the projected ego-graph representation, $e' \in \mathbb{R}^{1 \times d_s}$ is corresponding entity structural embedding, and d_x denotes the dimension of embedding space of LLMs. We do not explore more complex adaptation schemes (e.g., cross-attention) because they require extra graph-text pairs for pre-training. Moreover, such straightforward linear projection allows us to iterate data-centric experiments quickly, which has been proven effective in visual-text alignment (Liu et al., 2023).

Target Entity Generation. Given a query $q = (h, r, ?)$ and ego-graph serialization sequence $S_{\tilde{\mathcal{N}}_h}$, we formulate them to corresponding textual version and obtain the input of the LLM, denoted as $X = X_q + X_{S_{\tilde{\mathcal{N}}_h}}$. Let $\mathbf{X} \in \mathbb{R}^{|X| \times d_x}$ denote the textual content embeddings of input, where $|X|$ is the token length of X . We concatenate the soft graph token and input embeddings as final input of LLMs, i.e., $\mathbf{X}' = \mathbf{S}'_{\tilde{\mathcal{N}}_h} \parallel \mathbf{X}$. In this way, the structure information can interact frequently with the textual information, enabling LLMs to leverage the underlying graph structure. Finally, our optimization objective is to maximize the probability of generating the target entity name Y_t by maximizing:

$$P(\mathbf{Y}_t | \mathbf{X}', \mathbf{X}_{\mathcal{I}}) = \prod_{i=1}^L P_{\theta}(y_i | \mathbf{S}'_{\tilde{\mathcal{N}}_h} \parallel \mathbf{X}, \mathbf{X}_{\mathcal{I}}, \mathbf{Y}_{t, < i}), \quad (5)$$

where $\mathbf{X}_{\mathcal{I}}$ denotes the representation of instruction tokens, and L is the token length of target entity.

Connection to Graph Neural Networks. Our model shares essential mechanism similarities with existing GNNs, thus covering their advantages. First, due to the input length limitation of LLMs, our ego-graph serialization prompt for the query entity is aligned with GraphSAGE (Hamilton et al., 2017). And our similarity-based extraction module resembles graph regularization techniques like DropEdge (Rong et al., 2020). Additionally, our structure-text adapter carries structure features that can interact with text semantic features deeply in the encoding phase. Causal attention in LLMs can

be regarded as an advanced weighted average aggregation mechanism of GAT (Velickovic et al., 2018), facilitating our model to effectively model the varying importance of different neighbors to the central entity. Therefore, our framework integrates inductive bias required for graph tasks and enhances the graph structure understanding capability of LLMs.

3.4 KGC-Specific Instruction Tuning Strategy

The instruction tuning process aims to customize the reasoning behavior of LLM to meet the specific constraints and requirements of KGC task. An example of our instruction data can be seen in Appendix D.2. During the training process, we always keep the parameters of KGC filter frozen, and update both the weights of the projection layer and LLM. Considering the computation overhead of full-parameters updates for LLM, we employ low-rank adaptation (i.e., LoRA (Hu et al., 2022)) due to its simple implementation and promising performances (Liu et al., 2022). This approach freezes the pre-trained model parameters and updates the parameters of additional trainable weight matrix $\mathbf{W} \in \mathbb{R}^{d_1 \times d_2}$ by decomposing it into a product of two low-rank matrices: $\mathbf{W} = \mathbf{B}\mathbf{A}$, where $\mathbf{B} \in \mathbb{R}^{d_1 \times r}$, $\mathbf{A} \in \mathbb{R}^{d_2 \times r}$, and $r \ll \min(d_1, d_2)$. Hence, LoRA can effectively adapt the LLM to KGC task while requiring little memory overhead for storing gradients.

4 Experiment

4.1 Experimental Setup

We employ three widely-used KG datasets for our evaluation: FB15k-237 (Toutanova and Chen, 2015), CoDEX-M (Safavi and Koutra, 2020), and NELL-995 (Xiong et al., 2017). Detailed dataset statistics are shown in Appendix A. And the baselines adopted in our experiments are shown in Appendix B. In our implementation, we use LLaMA2-7B (Touvron et al., 2023) as the LLM backbone. We employ RotatE (Sun et al., 2019) as our filter for its simplicity and lightweight nature. The effect of different filters is discussed in Sec. 4.4. We report Mean Reciprocal Rank (MRR) and Hits@N (N=1,3,10) metric following the previous methods. Specific implementation details please refer to Appendix C.

Model	FB15k-237				CoDEX-M				NELL-995			
	MRR	H@1	H@3	H@10	MRR	H@1	H@3	H@10	MRR	H@1	H@3	H@10
Structure-Based Methods												
TransE (Bordes et al., 2013)	.279	.198	.376	.441	.303	.223	.336	.454	.401	.344	.472	.501
DistMult (Kazemi and Poole, 2018)	.281	.199	.301	.446	.223	.145	.245	.383	.485	.401	.524	.610
ComplEx (Trouillon et al., 2016)	.278	.194	.297	.450	<u>.337</u>	<u>.262</u>	<u>.370</u>	.476	.482	.399	.528	.606
ConvE (Dettmers et al., 2018)	.312	.225	.341	.497	.318	.239	.355	.464	.491	.403	.531	.613
RotatE (Sun et al., 2019)	.338	.241	.375	.533	.302	.219	.341	.461	.483	.435	.514	.565
KG-Mixup (Shomer et al., 2023)	<u>.358</u>	.264	-	<u>.548</u>	.319	.242	-	.465	<u>.522</u>	<u>.458</u>	-	<u>.621</u>
PLM-Based Methods												
GenKGC (Xie et al., 2022)	-	.192	.355	.439	-	-	-	-	-	-	-	-
KG-S2S (Chen et al., 2022)	.336	.257	.373	.498	.246	.186	.268	.372	.392	.324	.438	.511
CSProm-KG (Chen et al., 2023)	<u>.358</u>	<u>.269</u>	<u>.393</u>	.538	.320	.243	.355	.464	.508	.438	.548	.626
ChatGPT (Zhu et al., 2023) ♣	-	.097	.112	.124	-	-	-	-	-	-	-	-
PaLM2-540B (Anil et al., 2023) ♣	-	.115	.166	.217	-	-	-	-	-	-	-	-
KG-LLaMA-7B (Yao et al., 2023) ◇	.238	.165	.272	.423	.179	.159	.200	.204	.397	.388	.405	.406
FtG(Ours)	.392	.321	.413	.542	.395	.352	.407	<u>.473</u>	.538	.479	.573	.626

Table 1: The performance of FtG and baselines on three KG datasets. ♣ denotes results are obtained through evaluating all test triples from the Li et al. (2024). ◇ means that partial unreported results are obtained through the implementations as Yao et al. (2023).

4.2 Main Results

Table 1 displays the results of our experiments. Overall, we can observe that FtG achieves consistent and significant improvement on both datasets across most metrics, which demonstrates the effectiveness of our proposed FtG.

Compared to structure-based baselines, FtG showcases remarkable performance, which demonstrates the capability of FtG to understand and leverage graph structure. Additionally, compared to RotatE, which is employed as the filter in our model, FtG achieves Hits@1 improvements of 33.2%, 60.7%, and 10.11% across three datasets, indicating that *filter-then-generate* paradigm can effectively incorporate the strength of RotatE and LLMs, enabling FtG to leverage knowledge memorization and reasoning ability of LLM to address indistinguishable entity candidates. Compared to the sparse NELL-995, FtG improves more in the remaining two datasets, suggesting that FtG can fully utilize the structural information of the KG.

For the PLM-based baselines, FtG outperforms the SOTA method CSProm-KG by a substantial margin, indicating the superiority of our method. Focusing on the LLMs-based methods, we can find that even with instruction fine-tuning on KGs, LLMs still yield inferior performance. The reason is that directly eliciting LLMs to generate answers is prone to be influenced by hallucination of LLMs, leading to uncontrollable responses. In comparison, FtG achieves substantial improvements across three datasets, indicating that FtG can effectively

harness and unleash the capability of LLMs.

4.3 Ablation Study

In this subsection, we conduct an ablation study to investigate the individual contributions of different components in FtG. The results and meanings of various variants are reported in Table 2. The results reveal that all modules are essential because their absence has a detrimental effect on performance.

FGP	ESP	STA	FB15k-237		NELL995	
			MRR	H@1	MRR	H@1
✓	✓	✓	.392	.321	.538	.479
✓	✗	✗	.363(↓.029)	.279(↓.042)	.534(↓.004)	.469(↓.010)
✓	✓	✗	.374(↓.018)	.295(↓.026)	.535(↓.003)	.471(↓.008)
✓	✗	✓	.382(↓.010)	.306(↓.015)	.532(↓.006)	.472(↓.007)
✗	✗	✗	.238(↓.154)	.165(↓.156)	.397(↓.141)	.388(↓.091)

Table 2: Ablation for the FtG in FB15k-237 and NELL995. FGP denotes *filter-then-generate* paradigm. ESP denotes ego-graph serialization prompt. STA denotes structure-to-text adapter.

Specifically, to demonstrate the effectiveness of *filter-then-generate* paradigm, we directly adopt multiple-choice question instruction to fine-tune LLaMA-7b, and the results are shown in Table 2 Line 2. We observe that our paradigm significantly outperforms the base model that directly adopts instruction tuning (the last line). This demonstrates *filter-then-generate* paradigm can harness and unleash the capability of LLMs. Moreover, by comparing the variant with ego-graph serialization prompt (Line 3) and Line 2, we find that neighborhood information surrounding the query entity

can facilitate LLMs to perform KGC task. And the effect of ego-graph serialization prompt is more significant in FB15k-237, which is due to the fact that NELL-995 is a more sparse KG (the average node degree is lower than 2).

Besides, comparing the Line 2 and Line 4, we can find that our proposed graph soft prompt can achieve impressive Hits@1 improvements of 9.7% on FB15k-237, which demonstrates the soft graph prompt with lightweight adapter can deeply encode the inherent structural characteristics of the KG.

4.4 Discussion

In this section, we conduct a multifaceted performance analysis of FtG by answering the following questions. More analysis please refer to Appendix D.

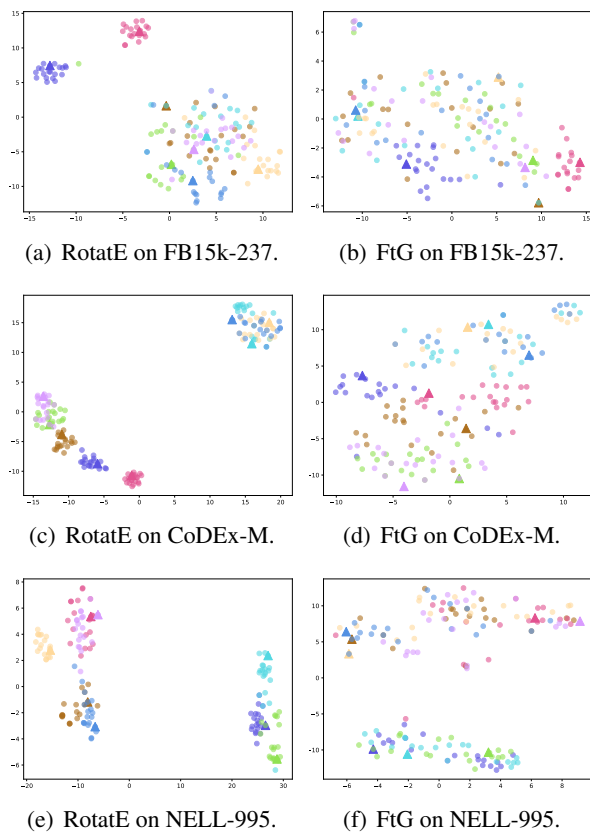


Figure 3: Visualization of candidate entity embeddings in CoDEX-M and NELL-995. Each color denotes a query, and candidate entities of the same color belong to the same query. \triangle denotes the target entity.

Q1: Why can FtG make impressive gains? Our motivation for proposing FtG is that we argue that *LLMs are not good KGC reasoner, but strong discriminator of hard samples*. We select several test samples where RotatE prediction failed in

FB15k-237, and visualize the top-20 highest scoring entities embeddings of each query in Figure 3 (a). Dots of the same color represent entities that are candidates under the same query, and \triangle is the target entity. Obviously, we can find that these candidate entities are indistinguishable for RotatE, and they overlap in the embedding space. As shown in Table 7, these hard samples require additional contextual knowledge to be distinguished. In contrast, we visualize the candidate entity’s hidden states in the last transformer-layer of FtG in Figure 3 (b). The figure demonstrates that our FtG can well distinguish the target entity from the hard samples, which is attributed to the inherent contextual knowledge of LLMs. Besides, we observe similar visualization results on the CoDEX-M and NELL-995 datasets.

Q2: Is FtG compatible with existing KGC methods? Here, we further evaluate whether our proposed FtG is robust enough when equipped with various KGC methods as a filter. Taking on a more challenging setting, we do not resort to retrain with different KGC methods. Instead, we load the trained LoRA weights directly and then switch different KGC filters for evaluation. Specifically, we employ a range of prevalent KGC methods, including structure-based TransE, ComplEx, RotatE, as well as PLM-based method CSProm-KG. Our results, as shown in Figure 4, demonstrate that existing KGC methods achieve significant improvements when integrated with FtG, across both FB15k-237 and NELL-995 datasets. This suggests that FtG can effectively incorporate the strength of conventional KGC methods and LLMs, enabling FtG to leverage the reasoning ability of LLMs to address indistinguishable candidates. Furthermore, this also underscores our method’s capacity to **enhance existing KGC methods in a plug-and-play manner**, demonstrating the versatility and effectiveness of FtG.

Q3: Is structure-aware pruning necessary in ego-graph serialization prompt? We replace the ego-graph serialization prompt in FtG with three other common heuristics, including random walk of the query entity, the entire one-hop ego-graph, and two-hop ego-graph. Empirical results are shown in Table 3, their final prediction are outperformed by FtG. Notably, comparing the results of the entire one-hop ego-graph (Line 3) with FtG, we can see that structure-aware pruning plays a crucial role, especially on the FB15k-237 dataset, which is rich in graph structure information (average degree of

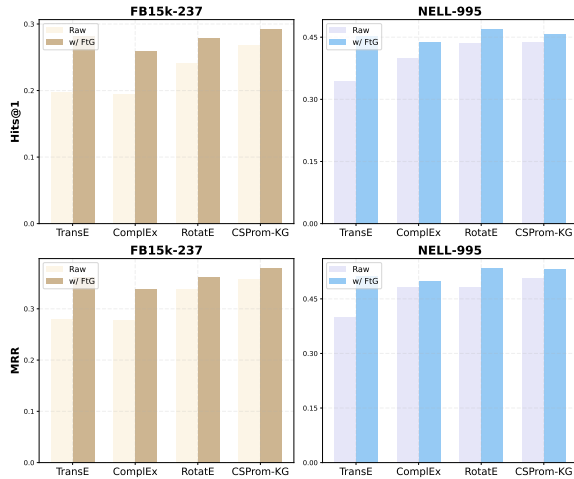


Figure 4: Performance (Hits@1&MRR) of FtG with various KGC filter on FB15k-237 and NELL-995 dataset.

entity is 37.4). These results demonstrate that our structure-aware pruning strategy can retain only the relevant information while filtering out the irrelevant neighbors. And it shares essential mechanism similarities with GAT (Velickovic et al., 2018), thus covering the advantage of GAT.

Heuristics	FB15k-237			NELL995		
	MRR	H@1	H@3	MRR	H@1	H@3
Structure-aware pruning	.392	.321	.413	.538	.479	.573
Random Walk	.356	.286	.402	.531	.468	.567
Entire 1-hop Ego-graph	.368	.313	.405	.536	.476	.568
2-hop Ego-graph	.382	.308	.404	.528	.463	.565

Table 3: Comparison of prediction performance with different heuristics.

Q4: How does the number of candidate entities retained in the filtering stage affect the performance? Here, we analyze the connection between the size of candidate set and performance. Our results, as presented in Figure 5. We observe that when increasing the size of the candidate set, the variation in model performance is not significant. The impact of this hyperparameter on results is akin to the trade-off between accuracy and recall. A larger candidate set implies a higher likelihood of containing the correct entity but also means that the LLM needs to comprehend more entities. In this paper, We prefer to enable LLMs to focus on hard samples that conventional KGC can not solve them well with limited model capacity and data amount. Therefore, we finally chose to set the size of the candidate entities to 20.

Q5: What are effects of *filter-then-generate* paradigm? We devise several variants to fully




Figure 5: Performance of FtG with different sizes of candidate entities.

analyze the impact of the FtG paradigm LLMs. Specifically, these variants are:

- LLaMA2-7B: directly prompt the LLM for KGC, the prompt format can refer to Table 4.
- LLaMA2-7B-FtG: we do not fine-tune the LLM and only adopt the FtG prompt.
- ChatGPT: we utilize the same prompt (refer to Figure 7) as Li et al. (2024) to evaluate.
- KG-LLaMA2-7B: adopt instruction-tuning to adapt LLaMA2-7B for KGC.
- FtG*: we only adopt *filter-then-generate* paradigm to fine-tune the LLM, it can be regarded as an ablation version.

The results, as illustrated in Figure 6, indicate that the *filter-then-generate* paradigm serves as an effective strategy for leveraging the capabilities of LLM in KGC. Especially, LLaMA2-7B with multiple-choice question prompt can outperform ChatGPT across some metrics in FB15k-237 dataset.

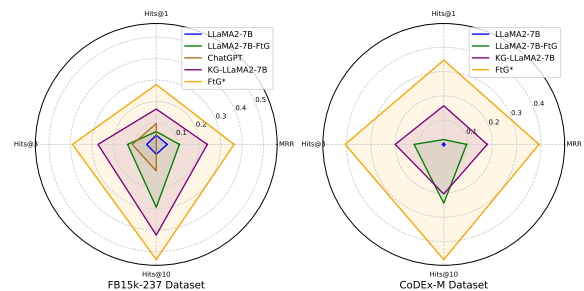


Figure 6: Effect of *filter-then-generate* paradigm. The MRR metric of ChatGPT in FB15k-237 is unreported, to include it in our figure without causing distraction, we’ve set its value to 0 by default.

5 Related Works

Structure-based KGC methods. Early methods typically define a score function to evaluate the

scores of triples through spatial measurement or latent matching. TransE (Bordes et al., 2013) defines each relation as translation from the head entity to tail entity. RotatE (Sun et al., 2019) further extends this idea in a complex space, enabling to model the symmetry relation pattern. Semantic matching methods, such as DistMult (Kazemi and Poole, 2018) and ComplEx (Trouillon et al., 2016), leverage the semantic similarity to capture complex interactions among entities and relations. Additionally, RGCN (Schlichtkrull et al., 2018), CompGCN (Vashishth et al., 2020), and SMiLE (Peng et al., 2022) employ graph neural networks to model the graph structure patterns inherent in KGs. KG-Mixup (Shomer et al., 2023) address the degree bias in KG, achieving promising results.

PLM-based KGC methods. KG-BERT (Yao et al., 2019) and StAR (Wang et al., 2021) utilize cross-entropy object and fine-tune PLM to produce entity embeddings. SimKGC (Wang et al., 2022) convert the KGC task into a semantic matching task and introduce contrastive learning to model fine-grained semantics. However, these methods suffer from unstable negative sampling. KGT5 (Saxena et al., 2022), KG-S2S (Chen et al., 2022), and UniLP (Liu et al., 2024) further exploit T5 (Raffel et al., 2020) with soft prompt to improve performance of generative KGC. Jiang et al. (2023b) propose a retrieve-based method to probe knowledge from PLM for open KG completion. CSProm-KG (Chen et al., 2023) integrates PLM with a structure-based method to bridge the structure and text information, achieving SOTA performance.

Moreover, Large Language Models (LLMs) have revolutionized various tasks in natural language processing. Focusing on KGC task, Zhu et al. (2023) construct prompts and evaluate the performance of ChatGPT with in-context learning on KGC task. ChatRule (Luo et al., 2023a) leverages ChatGPT to mine logical rules in KGs and applies these rules to make predictions. With the development of techniques of LLMs in the open source community, KG-LLaMA (Yao et al., 2023) makes the first step by applying instruction tuning to adapt LLaMA on KGC task. Additionally, some methods (Li et al., 2024) distill contextual knowledge from LLMs to improve the quality of entity texts, thus benefiting existing PLM-based approaches. Nevertheless, existing approaches that utilize LLMs for KGC task have not demonstrated satisfactory performance, and it remains a challenge to apply LLMs in the KGC task. The detailed comparison

between our FtG and existing LLM-based methods refer to Appendix F.

6 Conclusions

In this paper, we propose FtG, a instruction-tuning based method to enhance the performance of LLMs in KGC task. Our proposed *filter-then-generate* paradigm can effectively harness the capabilities of LLMs. To further incorporate the structural information into LLMs, we devise an ego-graph prompt and introduce a structure-text adapter. Extensive experiments demonstrate the effectiveness of FtG. In the future, we plan to adapt our method to other relevant downstream tasks, such as recommendation and open question answering.

Limitations

FtG can effectively harness the reasoning ability of LLMs and successfully incorporate the graph structural information into the LLMs, achieving substantial performance improvement on KGC task. However, the extremely large number of parameters in LLMs makes fine-tuning them resource-intensive. At the same time, LLMs are notoriously slow at decoding during inference. In our experiment, we use DeepSpeed (Rajbhandari et al., 2020) to accelerate training and inference, but FtG remain slower than traditional methods due to its inherent scale. Besides, if the KGC filter is not able to recall the target entity within the top- k candidates, FtG cannot make correct prediction. Therefore, a potential way to improve the effectiveness of FtG is to improve the success rate of target entity recall, and our FtG is more of a general framework to adapt LLM to KGC task.

Acknowledgements

We would like to thank all the anonymous reviewers and area chairs for their comments. This research is supported by National Natural Science Foundation of China (U23A20316), General Program of Natural Science Foundation of China (NSFC) (Grant No.62072346), and founded by Joint&Laboratory on Credit Technology.

References

Rohan Anil, Andrew M. Dai, Orhan Firat, Melvin Johnson, Dmitry Lepikhin, Alexandre Passos, Siamak Shakeri, Emanuel Taropa, Paige Bailey, Zhifeng Chen, Eric Chu, Jonathan H. Clark, Laurent El

- Shafey, Yanping Huang, Kathy Meier-Hellstern, Gaurav Mishra, Erica Moreira, Mark Omernick, Kevin Robinson, Sebastian Ruder, Yi Tay, Kefan Xiao, Yuanzhong Xu, Yujing Zhang, Gustavo Hernández Ábrego, Junwhan Ahn, Jacob Austin, Paul Barham, Jan A. Botha, James Bradbury, Siddhartha Brahma, Kevin Brooks, Michele Catasta, Yong Cheng, Colin Cherry, Christopher A. Choquette-Choo, Aakanksha Chowdhery, Clément Crepy, Shachi Dave, Mostafa Dehghani, Sunipa Dev, Jacob Devlin, Mark Díaz, Nan Du, Ethan Dyer, Vladimir Feinberg, Fangxiaoyu Feng, Vlad Fienber, Markus Freitag, Xavier Garcia, Sebastian Gehrmann, Lucas Gonzalez, and et al. 2023. [Palm 2 technical report](#). *CoRR*, abs/2305.10403.
- Kurt D. Bollacker, Colin Evans, Praveen K. Paritosh, Tim Sturge, and Jamie Taylor. 2008. [Freebase: a collaboratively created graph database for structuring human knowledge](#). In *Proceedings of the ACM SIGMOD International Conference on Management of Data, SIGMOD 2008, Vancouver, BC, Canada, June 10-12, 2008*, pages 1247–1250. ACM.
- Antoine Bordes, Nicolas Usunier, Alberto García-Durán, Jason Weston, and Oksana Yakhnenko. 2013. [Translating embeddings for modeling multi-relational data](#). In *Advances in Neural Information Processing Systems 26: 27th Annual Conference on Neural Information Processing Systems 2013. Proceedings of a meeting held December 5-8, 2013, Lake Tahoe, Nevada, United States*, pages 2787–2795.
- Chen Chen, Yufei Wang, Bing Li, and Kwok-Yan Lam. 2022. [Knowledge is flat: A seq2seq generative framework for various knowledge graph completion](#). In *Proceedings of the 29th International Conference on Computational Linguistics, COLING 2022, Gyeongju, Republic of Korea, October 12-17, 2022*, pages 4005–4017. International Committee on Computational Linguistics.
- Chen Chen, Yufei Wang, Aixin Sun, Bing Li, and Kwok-Yan Lam. 2023. [Dipping plms sauce: Bridging structure and text for effective knowledge graph completion via conditional soft prompting](#). In *Findings of the Association for Computational Linguistics: ACL 2023, Toronto, Canada, July 9-14, 2023*, pages 11489–11503.
- Tim Dettmers, Pasquale Minervini, Pontus Stenetorp, and Sebastian Riedel. 2018. [Convolutional 2d knowledge graph embeddings](#). In *Proceedings of the Thirty-Second AAAI Conference on Artificial Intelligence, (AAAI-18), New Orleans, Louisiana, USA, February 2-7, 2018*, pages 1811–1818. AAAI Press.
- Wenqi Fan, Shijie Wang, Jiani Huang, Zhikai Chen, Yu Song, Wenzhuo Tang, Haitao Mao, Hui Liu, Xiaorui Liu, Dawei Yin, and Qing Li. 2024. [Graph machine learning in the era of large language models \(llms\)](#). *CoRR*, abs/2404.14928.
- William L. Hamilton, Zitao Ying, and Jure Leskovec. 2017. [Inductive representation learning on large graphs](#). In *Advances in Neural Information Processing Systems 30: Annual Conference on Neural Information Processing Systems 2017, December 4-9, 2017, Long Beach, CA, USA*, pages 1024–1034.
- Yuhang He, Jihai Zhang, Jianzhu Bao, Fangquan Lin, Cheng Yang, Bing Qin, Ruifeng Xu, and Wotao Yin. 2024. [Bc-prover: Backward chaining prover for formal theorem proving](#). In *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing, EMNLP 2024, Miami, FL, USA, November 12-16, 2024*, pages 3059–3077.
- Edward J. Hu, Yelong Shen, Phillip Wallis, Zeyuan Allen-Zhu, Yuanzhi Li, Shean Wang, Lu Wang, and Weizhu Chen. 2022. [Lora: Low-rank adaptation of large language models](#). In *The Tenth International Conference on Learning Representations, ICLR 2022, Virtual Event, April 25-29, 2022*. OpenReview.net.
- Lei Huang, Weijiang Yu, Weitao Ma, Weihong Zhong, Zhangyin Feng, Haotian Wang, Qianglong Chen, Weihua Peng, Xiaocheng Feng, Bing Qin, and Ting Liu. 2023. [A survey on hallucination in large language models: Principles, taxonomy, challenges, and open questions](#). *CoRR*, abs/2311.05232.
- Ziwei Ji, Nayeon Lee, Rita Frieske, Tiezheng Yu, Dan Su, Yan Xu, Etsuko Ishii, Yejin Bang, Andrea Madotto, and Pascale Fung. 2023. [Survey of hallucination in natural language generation](#). *ACM Comput. Surv.*, 55(12):248:1–248:38.
- Jinhao Jiang, Kun Zhou, Wayne Xin Zhao, Yaliang Li, and Ji-Rong Wen. 2023a. [Reasoninglm: Enabling structural subgraph reasoning in pre-trained language models for question answering over knowledge graph](#). In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing, EMNLP 2023, Singapore, December 6-10, 2023*, pages 3721–3735.
- Pengcheng Jiang, Shivam Agarwal, Bowen Jin, Xuan Wang, Jimeng Sun, and Jiawei Han. 2023b. [Text augmented open knowledge graph completion via pre-trained language models](#). In *Findings of the Association for Computational Linguistics: ACL 2023, Toronto, Canada, July 9-14, 2023*, pages 11161–11180.
- Seyed Mehran Kazemi and David Poole. 2018. [Simple embedding for link prediction in knowledge graphs](#). In *Advances in Neural Information Processing Systems 31: Annual Conference on Neural Information Processing Systems 2018, NeurIPS 2018, December 3-8, 2018, Montréal, Canada*, pages 4289–4300.
- Dawei Li, Zhen Tan, Tianlong Chen, and Huan Liu. 2024. [Contextualization distillation from large language model for knowledge graph completion](#). *CoRR*, abs/2402.01729.
- Junyi Li, Xiaoxue Cheng, Xin Zhao, Jian-Yun Nie, and Ji-Rong Wen. 2023. [Halueval: A large-scale hallucination evaluation benchmark for large language models](#). In *Proceedings of the 2023 Conference on*

- Empirical Methods in Natural Language Processing, EMNLP 2023, Singapore, December 6-10, 2023*, pages 6449–6464.
- Ben Liu, Miao Peng, Wenjie Xu, Xu Jia, and Min Peng. 2024. [Unilp: Unified topology-aware generative framework for link prediction in knowledge graph](#). In *Proceedings of the ACM on Web Conference 2024, WWW 2024, Singapore, May 13-17, 2024*, pages 2170–2180. ACM.
- Haokun Liu, Derek Tam, Mohammed Muqeeth, Jay Mohata, Tenghao Huang, Mohit Bansal, and Colin Raffel. 2022. [Few-shot parameter-efficient fine-tuning is better and cheaper than in-context learning](#). In *Advances in Neural Information Processing Systems 35: Annual Conference on Neural Information Processing Systems 2022, NeurIPS 2022, New Orleans, LA, USA, November 28 - December 9, 2022*.
- Haotian Liu, Chunyuan Li, Qingyang Wu, and Yong Jae Lee. 2023. [Visual instruction tuning](#). *CoRR*, abs/2304.08485.
- Linhao Luo, Jiaxin Ju, Bo Xiong, Yuan-Fang Li, Gholamreza Haffari, and Shirui Pan. 2023a. [Chatrule: Mining logical rules with large language models for knowledge graph reasoning](#). *CoRR*, abs/2309.01538.
- Linhao Luo, Yuan-Fang Li, Gholamreza Haffari, and Shirui Pan. 2023b. [Reasoning on graphs: Faithful and interpretable large language model reasoning](#). *CoRR*, abs/2310.01061.
- Yubo Ma, Yixin Cao, Yong Hong, and Aixin Sun. 2023. [Large language model is not a good few-shot information extractor, but a good reranker for hard samples!](#) In *Findings of the Association for Computational Linguistics: EMNLP 2023, Singapore, December 6-10, 2023*, pages 10572–10601.
- OpenAI. 2023. [GPT-4 technical report](#). *CoRR*, abs/2303.08774.
- Shirui Pan, Linhao Luo, Yufei Wang, Chen Chen, Jiapu Wang, and Xindong Wu. 2023. [Unifying large language models and knowledge graphs: A roadmap](#). *CoRR*.
- Miao Peng, Ben Liu, Qianqian Xie, Wenjie Xu, Hua Wang, and Min Peng. 2022. [Smile: Schema-augmented multi-level contrastive learning for knowledge graph link prediction](#). In *Findings of the Association for Computational Linguistics: EMNLP 2022, Abu Dhabi, United Arab Emirates, December 7-11, 2022*, pages 4165–4177.
- Colin Raffel, Noam Shazeer, Adam Roberts, Katherine Lee, Sharan Narang, Michael Matena, Yanqi Zhou, Wei Li, and Peter J. Liu. 2020. [Exploring the limits of transfer learning with a unified text-to-text transformer](#). *J. Mach. Learn. Res.*, 21:140:1–140:67.
- Samyam Rajbhandari, Jeff Rasley, Olatunji Ruwase, and Yuxiong He. 2020. [Zero: memory optimizations toward training trillion parameter models](#). In *Proceedings of the International Conference for High Performance Computing, Networking, Storage and Analysis, SC 2020, Virtual Event / Atlanta, Georgia, USA, November 9-19, 2020*, page 20. IEEE/ACM.
- Yu Rong, Wenbing Huang, Tingyang Xu, and Junzhou Huang. 2020. [Dropedge: Towards deep graph convolutional networks on node classification](#). In *8th International Conference on Learning Representations, ICLR 2020, Addis Ababa, Ethiopia, April 26-30, 2020*.
- Tara Safavi and Danai Koutra. 2020. [Codex: A comprehensive knowledge graph completion benchmark](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing, EMNLP 2020, Online, November 16-20, 2020*, pages 8328–8350.
- Apoorv Saxena, Adrian Kochsiek, and Rainer Gemulla. 2022. [Sequence-to-sequence knowledge graph completion and question answering](#). In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers), ACL 2022, Dublin, Ireland, May 22-27, 2022*, pages 2814–2828.
- Michael Sejr Schlichtkrull, Thomas N. Kipf, Peter Bloem, Rianne van den Berg, Ivan Titov, and Max Welling. 2018. [Modeling relational data with graph convolutional networks](#). In *The Semantic Web - 15th International Conference, ESWC 2018, Heraklion, Crete, Greece, June 3-7, 2018, Proceedings*, volume 10843 of *Lecture Notes in Computer Science*, pages 593–607. Springer.
- Harry Shomer, Wei Jin, Wentao Wang, and Jiliang Tang. 2023. [Toward degree bias in embedding-based knowledge graph completion](#). In *Proceedings of the ACM Web Conference 2023, WWW 2023, Austin, TX, USA, 30 April 2023 - 4 May 2023*, pages 705–715. ACM.
- Dong Shu, Tianle Chen, Mingyu Jin, Yiting Zhang, Chong Zhang, Mengnan Du, and Yongfeng Zhang. 2024. [Knowledge graph large language model \(KG-LLM\) for link prediction](#). *CoRR*, abs/2403.07311.
- Zhiqing Sun, Zhi-Hong Deng, Jian-Yun Nie, and Jian Tang. 2019. [Rotate: Knowledge graph embedding by relational rotation in complex space](#). In *7th International Conference on Learning Representations, ICLR 2019, New Orleans, LA, USA, May 6-9, 2019*. OpenReview.net.
- Kristina Toutanova and Danqi Chen. 2015. [Observed versus latent features for knowledge base and text inference](#). In *Proceedings of the 3rd Workshop on Continuous Vector Space Models and their Compositionality, CVSC 2015, Beijing, China, July 26-31, 2015*, pages 57–66.
- Hugo Touvron, Louis Martin, Kevin Stone, Peter Albert, Amjad Almahairi, Yasmine Babaei, Nikolay Bashlykov, Soumya Batra, Prajjwal Bhargava, Shruti

- Bhosale, Dan Bikel, Lukas Blecher, Cristian Canton-Ferrer, Moya Chen, Guillem Cucurull, David Esiobu, Jude Fernandes, Jeremy Fu, Wenyin Fu, Brian Fuller, Cynthia Gao, Vedanuj Goswami, Naman Goyal, Anthony Hartshorn, Saghar Hosseini, Rui Hou, Hakan Inan, Marcin Kardas, Viktor Kerkez, Madian Khabsa, Isabel Kloumann, Artem Korenev, Punit Singh Koura, Marie-Anne Lachaux, Thibaut Lavril, Jenya Lee, Diana Liskovich, Yinghai Lu, Yuning Mao, Xavier Martinet, Todor Mihaylov, Pushkar Mishra, Igor Molybog, Yixin Nie, Andrew Poulton, Jeremy Reizenstein, Rashi Rungta, Kalyan Saladi, Alan Schelten, Ruan Silva, Eric Michael Smith, Ranjan Subramanian, Xiaoqing Ellen Tan, Binh Tang, Ross Taylor, Adina Williams, Jian Xiang Kuan, Puxin Xu, Zheng Yan, Iliyan Zarov, Yuchen Zhang, Angela Fan, Melanie Kambadur, Sharan Narang, Aurélien Rodriguez, Robert Stojnic, Sergey Edunov, and Thomas Scialom. 2023. [Llama 2: Open foundation and finetuned chat models](#). *CoRR*, abs/2307.09288.
- Théo Trouillon, Johannes Welbl, Sebastian Riedel, Éric Gaussier, and Guillaume Bouchard. 2016. [Complex embeddings for simple link prediction](#). In *Proceedings of the 33rd International Conference on Machine Learning, ICML 2016, New York City, NY, USA, June 19-24, 2016*, volume 48 of *JMLR Workshop and Conference Proceedings*, pages 2071–2080. JMLR.org.
- Shikhar Vashishth, Soumya Sanyal, Vikram Nitin, and Partha P. Talukdar. 2020. [Composition-based multi-relational graph convolutional networks](#). In *8th International Conference on Learning Representations, ICLR 2020, Addis Ababa, Ethiopia, April 26-30, 2020*. OpenReview.net.
- Petar Velickovic, Guillem Cucurull, Arantxa Casanova, Adriana Romero, Pietro Liò, and Yoshua Bengio. 2018. [Graph attention networks](#). In *6th International Conference on Learning Representations, ICLR 2018, Vancouver, BC, Canada, April 30 - May 3, 2018, Conference Track Proceedings*.
- Bo Wang, Tao Shen, Guodong Long, Tianyi Zhou, Ying Wang, and Yi Chang. 2021. [Structure-augmented text representation learning for efficient knowledge graph completion](#). In *WWW '21: The Web Conference 2021, Virtual Event / Ljubljana, Slovenia, April 19-23, 2021*, pages 1737–1748. ACM / IW3C2.
- Liang Wang, Wei Zhao, Zhuoyu Wei, and Jingming Liu. 2022. [Simkgc: Simple contrastive knowledge graph completion with pre-trained language models](#). In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers), ACL 2022, Dublin, Ireland, May 22-27, 2022*, pages 4281–4294.
- Xiang Wang, Xiangnan He, Yixin Cao, Meng Liu, and Tat-Seng Chua. 2019. [KGAT: knowledge graph attention network for recommendation](#). In *Proceedings of the 25th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining, KDD 2019, Anchorage, AK, USA, August 4-8, 2019*, pages 950–958. ACM.
- Jason Wei, Maarten Bosma, Vincent Y. Zhao, Kelvin Guu, Adams Wei Yu, Brian Lester, Nan Du, Andrew M. Dai, and Quoc V. Le. 2022a. [Finetuned language models are zero-shot learners](#). In *The Tenth International Conference on Learning Representations, ICLR 2022, Virtual Event, April 25-29, 2022*. OpenReview.net.
- Jason Wei, Yi Tay, Rishi Bommasani, Colin Raffel, Barret Zoph, Sebastian Borgeaud, Dani Yogatama, Maarten Bosma, Denny Zhou, Donald Metzler, Ed H. Chi, Tatsunori Hashimoto, Oriol Vinyals, Percy Liang, Jeff Dean, and William Fedus. 2022b. [Emergent abilities of large language models](#). *Trans. Mach. Learn. Res.*, 2022.
- Xin Xie, Ningyu Zhang, Zhoubo Li, Shumin Deng, Hui Chen, Feiyu Xiong, Mosha Chen, and Huajun Chen. 2022. [From discrimination to generation: Knowledge graph completion with generative transformer](#). In *Companion of The Web Conference 2022, Virtual Event / Lyon, France, April 25 - 29, 2022*, pages 162–165. ACM.
- Wenhan Xiong, Thien Hoang, and William Yang Wang. 2017. [DeepPath: A reinforcement learning method for knowledge graph reasoning](#). In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing, EMNLP 2017, Copenhagen, Denmark, September 9-11, 2017*, pages 564–573.
- Zhentao Xu, Mark Jerome Cruz, Matthew Guevara, Tie Wang, Manasi Deshpande, Xiaofeng Wang, and Zheng Li. 2024. [Retrieval-augmented generation with knowledge graphs for customer service question answering](#). *CoRR*, abs/2404.17723.
- Liang Yao, Chengsheng Mao, and Yuan Luo. 2019. [KG-BERT: BERT for knowledge graph completion](#). *CoRR*, abs/1909.03193.
- Liang Yao, Jiazhen Peng, Chengsheng Mao, and Yuan Luo. 2023. [Exploring large language models for knowledge graph completion](#). *CoRR*, abs/2308.13916.
- Jihai Zhang, Wei Wang, Siyan Guo, Li Wang, Fangquan Lin, Cheng Yang, and Wotao Yin. 2024. [Solving general natural-language-description optimization problems with large language models](#). In *Proceedings of the 2024 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies: Industry Track, NAACL 2024, Mexico City, Mexico, June 16-21, 2024*, pages 483–490.
- Yichi Zhang, Zhuo Chen, Wen Zhang, and Huajun Chen. 2023. [Making large language models perform better in knowledge graph completion](#). *CoRR*, abs/2310.06671.
- Yuqi Zhu, Xiaohan Wang, Jing Chen, Shuofei Qiao, Yixin Ou, Yunzhi Yao, Shumin Deng, Huajun Chen, and Ningyu Zhang. 2023. [Llms for knowledge graph construction and reasoning: Recent capabilities and future opportunities](#). *CoRR*, abs/2305.13168.

A Dataset

We use FB15k-237 (Toutanova and Chen, 2015), CoDEX-M (Safavi and Koutra, 2020), and NELL-995 (Xiong et al., 2017) for evaluation. FB15k-237 is a subset extracted from the Freebase (Bollacker et al., 2008), which includes commonsense knowledge about movies, sports, locations, etc. CoDEX-M is extracted from Wikipedia, which contains tens of thousands of hard negative triples, making it a more challenging KGC benchmark. NELL-995 is taken from the Never Ending Language Learner (NELL) system and covers many domains. Detailed statistics of all these datasets are shown in Table 5.

B Baseline Details

B.1 Baselines

We compare our FtG against three types of baselines: (1) *structure-based methods*, including TransE (Bordes et al., 2013), DistMult (Kazemi and Poole, 2018), ComplEx (Trouillon et al., 2016), ConvE (Dettmers et al., 2018), RotatE (Sun et al., 2019), and KG-Mixup (Shomer et al., 2023). (2) *PLM-based methods*, including GenKGC (Xie et al., 2022), KG-S2S (Chen et al., 2022), and CSProm-KG (Chen et al., 2023). (3) *LLM-based methods*, including ChatGPT (Zhu et al., 2023), PaLM2-540B (Anil et al., 2023), and KG-LLaMA-7B (Yao et al., 2023). Both ChatGPT and PaLM2-540B employ LLMs as the backbone and focus on prompt engineer, enabling LLMs to understand the KGC task. KG-LLaMA-7B is an instruction fine-tuned LLaMA2-7B based on KG datasets.

B.2 Implementation of Baselines

Since some baselines miss results on some metrics, we implement these baselines based on their released code. For structure-based baselines, we use the toolkit provided in RotatE¹, which gives state-of-the-art performance of existing KGC models in a unified framework. We adopt the optimal hyperparameter configurations reported in their papers.

For KG-S2S, on the CoDEX-M dataset, we utilize the official code provided² and set the length of the entity description text to 30, epochs to 50, batch size to 32, and beam width to 40. For experiments on the NELL-995 dataset, the epoch is increased to 100, batch size is determined to be

¹<https://github.com/DeepGraphLearning/KnowledgeGraphEmbedding>

²<https://github.com/chenchens190009/KG-S2S>

64, and the learning rate is fixed at 0.001. It is worth noting that the NELL-995 dataset inherently lacks entity description texts, the length of entity description text is set to 0. For CSProm-KG³, following the original paper, we choose ConvE as graph model owing to its superior performance metrics, alongside the BERT-base model serving as the foundational PLM. Specifically, on the CoDEX-M dataset, we set the batch size to 128, epoch to 60, length of entity description text to 30, learning rate to 0.0005, prompt length to 30, label smoothing to 0.1, and embedding dimension to 156. On the NELL-995 dataset, we adopt same configurations expect the length of entity description is set 0. For the KG-LLaMA-7B, we adopt same LoRA configurations in Table 6 as our FtG for fair comparison.

C Implementation Details

In our implementation, the quantity k of candidates retained is selected from $\{10, 20, 30, 40\}$. During training, we keep the RotatE frozen and employ LoRA to fine-tune the model. The detailed hyperparameters we use during training and inference are shown in Table 6. We employ identical hyperparameters in different datasets. DeepSpeed ZeRO stage³⁴ is enabled for optimization. All models are trained using 2 Nvidia A800 GPUs, each with 80GB of memory. For all datasets, we report Mean Reciprocal Rank (MRR) and Hits@N (N=1,3,10) metric following the previous works.

D Instruction Template

D.1 Prompt for ChatGPT and PaLM2

Zhu et al. (2023) construct few-shot demonstrations to assess the performance of LLM in KGC. Figure 7 shows a example of the input to LLMs, and (Li et al., 2024) utilize the API parameter to obtain multiple candidates, enabling the calculation of Hits@1, Hits@3, and Hits@10 metrics.

D.2 Prompt for FtG

In our framework, we use a simple template in Table 4 to convert the query triple to text format like Yao et al. (2023). Then we formulate the KGC task into a multiple-choice question format. A specific example is shown in Figure 8.

³<https://github.com/chenchens190009/CSProm-KG>

⁴<https://github.com/microsoft/Megatron-DeepSpeed>

Triple	(War on Terrorism, military/military conflict/combatants./military/military combatant group/combatants, Canada)
Tail Prompt	War on Terrorism, military/military conflict/combatants./military/military combatant group/combatants?
Head Prompt	What/Who/When/Where/Why military/military conflict/combatants./military/military combatant group/combatants Canada?

Table 4: The prompt we use to verbalize the query triple.

Dataset	\mathcal{E}	\mathcal{R}	Train	Valid	Test
FB15k-237	14,541	237	272,115	17,535	20,466
CoDex-M	17,050	51	185,584	10,310	10,311
NELL-995	74,536	200	149,678	543	2,818

Table 5: Statistics of the Datasets.

Name	Value
lora r	16
lora alpha	32
lora dropout	0.05
lora target modules	(q, k, v, o, down, up, gate) proj
cutoff len	1024
epochs	2
per device batch size	64
gradient accumulation steps	1
learning rate	$3e - 4$
weight decay	$1e - 5$
warm ratio	0.01
lr scheduler type	cosine
num return sequences	10
projection layers	1

Table 6: Detailed hyperparameters used in our paper.

E Case Study

In Table 8, we demonstrate some cases to illustrate the differences in responses between existing LLMs-based methods and FtG. From the cases, our FtG can effectively leverage the capabilities of LLMs while avoiding the generation of uncontrollable text.

F Comparison with existing LLMs-based methods

In this subsection, we provide a detailed introduction to existing LLMs-based methods and further discuss the potential application of our FtG. The existing LLM-based methods mainly include:

- **KoPA** (Zhang et al., 2023): proposes an instruction-tuning method based on LLaMA2-

Hard negative candidates	Explanation
(Senegal, part of, Middle East)	Senegal is part of West Africa.
(Lesotho, official language, American English)	English, not American English, is an official language of Lesotho.
(Vatican City, member of, UNESCO)	Vatican City is a UNESCO World Heritage Site but not a member state.

Table 7: Selected examples of hard samples in Codex.

Prompt for triple: (Friedrich Gundolf, employer, Heidelberg University)	
Tail Prediction:	Predict the tail entity [MASK] from the given (Erich Brandenburg, employer, [MASK]) by completing the sentence "who is the employer of Erich Brandenburg? The answer is ". The answer is Leipzig University, so the [MASK] is Leipzig University. Predict the tail entity [MASK] from the given (Friedrich Gundolf, employer, [MASK]) by completing the sentence "who is the employer of Friedrich Gundolf? The answer is ". The answer is
Head Prediction:	Predict the head entity [MASK] from the given ([MASK], employer, Leipzig University) by completing the sentence "Leipzig University is the employer of who? The answer is ". The answer is Erich Brandenburg, so the [MASK] is Erich Brandenburg. Predict the head entity [MASK] from the given ([MASK], employer, Heidelberg University) by completing the sentence "Heidelberg University is the employer of who? The answer is ". The answer is

Figure 7: The prompt that directly leverage LLMs to perform KGC. Tail Prompt and Head Prompt mean the input to predict the tail and head entity respectively.

FtG Prompt for triple: (Friedrich Gundolf, employer, Heidelberg University)	
Instruction:	Please answer the following commonsense question and select only one answer from the candidates that is most relevant to the question.
Tail Prediction	
Context:	Friedrich Gundolf is German poet.
Question:	Friedrich Gundolf employer?
Candidate Answers:	Ludwig Maximilian University of Munich, University of Giessen, Heidelberg University, Humboldt University of Berlin, Karlsruhe Institute of Technology, University of Marburg, University of Bonn, University of Freiburg, Leipzig University, University of Kiel, University of Tu00fcbingen, University of Gu00f6ttingen, Goethe University Frankfurt, University of Wu00fcrzburg, Freie Universit00e4t Berlin, Technische Universit00e4t Darmstadt, Technical University of Berlin, University of Halle-Wittenberg, University of Wroc0142aw, University of Stuttgart.
Head Prediction	
Context:	Hubertus Strughold employer Heidelberg University Robert Spaemann employer Heidelberg University Michael von Albrecht employer Heidelberg University.
Question:	What/Who/When/Where/Why employer Heidelberg University?
Candidate Answers:	Friedrich Julius Stahl, Wolfgang Schadewaldt, Theodor Mommsen, Wilhelm Kahl, Ludwig Bergstr00e4sser, Carl Engler, Karl von Gareis, Adolf Butenandt, Heinrich von Sybel, Friedrich Gundolf, Hermann Usener, Ernst Zitelmann, Walther Sch00fcking, Hermann Dietrich, Albert Hu00e4nel, Johann Michael Franz Birnbaum, Ludwig Bamberger, Carl Bosch, Rudolf Otto, Friedrich Carl von Savigny.

Figure 8: An example prompt that we adopt in FtG. █ denotes target entity.

7B for KG triple classification task. Although the authors claim their focus is on KGC, their work is strictly speaking a triple classification task. This means given a true triple, they randomly replace the head entity or tail entity to construct a negative sample, and then let the model perform binary classification, i.e., simply outputting True or False.

- **KG-LLM** (Shu et al., 2024): similarly focuses on the triple classification task. It constructs chain-of-thought prompts via random walks on the KG and then fine-tunes LLMs to

Query head: (? , employer, Heidelberg University)		
Models	Labels	Output
LLaMA2-7B	✗	What/Who/When/Where/Why employer University of Cincinnati end-code
ChatGPT	✗	[MASK] is not able to provide a specific answer for the second sentence without more information.
ChatGPT+FtG Prompt	✗	Hubertus Strughold, Robert Spaemann, and Michael von Albrecht are employed by Heidelberg University.
GPT4-o	✗	Max Weber. So, the [MASK] is Max Weber.
GPT4-o+FtG Prompt	✓	Friedrich Gundolf
FtG (ours)	✓	Friedrich Gundolf

Query tail: (Friedrich Gundolf, employer, ?)		
Models	Labels	Output
LLaMA2-7B	✗	Hermann Oncken (1879–1934) Theodor W. Adorno and Max Horkheimer, Dialectic of Enlightenment (1947)
ChatGPT	✓	University of Heidelberg, so the [MASK] is University of Heidelberg.
ChatGPT+FtG Prompt	✓	Heidelberg University
GPT4-o	✓	The answer is Heidelberg University.
GPT4-o+FtG Prompt	✓	Heidelberg University
FtG (ours)	✓	Heidelberg University

Table 8: Case comparisons between our FtG and existing LLMs. Corresponding prompt refer to Appendix D.

perform binary classification. The results indicate that LLMs have a potential to understand structural information on the KG.

- **ChatGPT for KGC (Zhu et al., 2023):** is the first approach that truly utilizes LLMs for KGC task, which converts the KGC task into sentence masking prediction. ChatGPT is used to predict the missing entity by predicting the special token [MASK]. However, ChatGPT did not yield satisfactory results compared to the impressive capabilities it has demonstrated in other NLP tasks.
- **Contextualization Distillation (Li et al., 2024):** is a data-augmentation method that utilizes LLMs to generate background text about the query entity to enhance existing KGC methods that utilize textual information. However, due to the hallucination issues inherent to LLMs and the high requirements of downstream KGC methods, the improvement is limited.

- **KG-LLaMA (Yao et al., 2023):** converts KGC into a QA task, leveraging the instruction-following capability of LLMs to adapt them for KGC tasks. While this approach has achieved promising results, it still falls short compared to previous methods based on structural information.

In contrast, our FtG can effectively leverage the capabilities of LLMs while avoiding the generation of uncontrollable text. Additionally, we propose an efficient approach to enable LLMs to utilize the structural information of KGs, which has not been achieved by previous methods.

Potential Downstream Applications of FtG.

- The construction of KG is very expensive, especially in specialized domains such as financial and law. Our method can fully leverage LLMs’ inherent knowledge and reasoning capabilities, making it suitable for automatically completing existing KGs. Additionally, the AI community has witnessed the emergence of numerous powerful LLMs, which have made huge advancements and led to the pursuit of possible Artificial General Intelligence. Our FtG provides a possible way to integrate knowledge graphs with LLMs, aligning with the current trends in AI domains.
- In the recommendation domain, systems need to suggest specific items from a vast pool. Our approach can be effectively applied here: filter-then-generate paradigm can initially filter the large pool of items, honing in on a more relevant subset based on user profiles and preferences. And the ego-graph serialization prompt can capture and model detailed user interaction history. Finally, the encoding the ego-graph into a soft prompt token and map it into LLMs’ space with an adapter can provides a meaningful way to apply LLMs for final recommendation.