

K-ON: Stacking Knowledge On the Head Layer of Large Language Model

Lingbing Guo^{1,2}, Yichi Zhang^{1,2}, Zhongpu Bo³, Zhuo Chen^{1,2}, Mengshu Sun³,
Zhiqiang Zhang³, Wen Zhang^{4,2*}, Huajun Chen^{1,2,5*}

¹College of Computer Science and Technology, Zhejiang University

²ZJU-Ant Group Joint Lab of Knowledge Graph

³Ant Group

⁴School of Software Technology, Zhejiang University

⁵Zhejiang Key Laboratory of Big Data Intelligent Computing

{lbguo, zhangyichi2022, zhang.wen, huajunsir}@zju.edu.cn

Abstract

Recent advancements in large language models (LLMs) have significantly improved various natural language processing (NLP) tasks. Typically, LLMs are trained to predict the next token, aligning well with many NLP tasks. However, in knowledge graph (KG) scenarios, entities are the fundamental units and identifying an entity requires at least several tokens. This leads to a granularity mismatch between KGs and natural languages. To address this issue, we propose K-ON, which integrates KG knowledge into the LLM by employing multiple head layers for next k -step prediction. K-ON can not only generate entity-level results in one step, but also enables contrastive loss against entities, which is the most powerful tool in KG representation learning. Experimental results show that K-ON outperforms state-of-the-art methods that incorporate text and even the other modalities.

Introduction

Large language models (LLMs) are trained on vast amounts of corpora and store world knowledge within billions of neurons (Achiam et al. 2023; Touvron et al. 2023). Despite of the prosperity in LLM-based applications, unleashing its power for knowledge graph (KG) tasks remains challenging.

Tokens are the basic elements for language models, but it needs to take at least several tokens to describe and identify different entities in a KG. Creating new token identifiers for each entity and learning them during fine-tuning is an alternative choice, however, which is extremely time-consuming and may negatively affect the native performance of LLMs.

In this paper, we explore how to effectively and efficiently use multiple tokens to describe entities in a given KG. Evidently, directly optimizing the sequence prediction objective will result in the out-of-KG problem as LLM lacks awareness of the KG’s entities, while enumerating all entities in the input instruction is unrealistic. Take Figure 1a as an example, the task is to predict the target entity *Matt Damon* given the incomplete triplet (*The Bourne Identity* (2002 film), starring, ?). The vanilla learning schema is inefficient because generating a single entity requires multiple steps and cannot be parallelized across entities.

Most existing methods compromise on this dilemma and apply LLMs only to simple KG tasks (Yang et al. 2023; Guan et al. 2024; Pan et al. 2023, 2024), such as verifying the correctness of a triplet (Zhang et al. 2023) or predicting the target from a limited number of candidates (Yao et al. 2023). In contrast, we propose *K-ON* to employ K head layers for predicting entities at one shot and stack knowledge on these heads by entity-level contrastive learning.

As shown in Figure 1b, K-ON adapts K different head layers from the original LLM head, where the k -th head is responsible for predicting the k -th token for all entities. For example, an entity *Matt Damon* is tokenized into K input token IDs $t_{0:K-1}$, with t_0 representing *Matt* and t_{K-1} being the last token *Damon* or padding token. We then extract the probability of the first token from the first head layer, and so forth. For the other entities that may serve as negative examples in contrastive learning, we can reuse the K -step probability estimations to extract their scores.

The risk underlying the prediction of the next K tokens is *over-optimization*. The model may over-optimize for predicting just the next K tokens, dropping the fact that these tokens form the target entity. For instance, while minimizing cross-entropy loss for the first token, with *Matt* as the positive label, many negative entries (tokens) are not constituent elements of an entity. Moreover, increasing the probability of *Matt* does not always equate to maximizing the probability of *Matt Damon*. To tackle this issue, K-ON employs an entity-level contrastive loss, which treats the K -step predictions as an integrity and estimates the joint probabilities.

Another risk in next- K -token prediction is *distribution corruption*. In the original schema, the probability of the second token “*Damon*” is conditioned on the first token “*Matt*”. However, this conditioning is absent in the next- K -tokens schema. Such discrepancy may degrade the performance as well as inference ability of the original LLM.

To address this issue, we propose *head trajectory tuning* (HTT) to align the distribution trajectories between the original LLM’s prediction and the next- K -token predictions. We first leverage a conditional attention layer to process the hidden states from K head layers to reconstruct the sequential dependencies between different steps. Then, we compute a standard sequence prediction loss with the original LLM head layer as target. Finally, we can align the output probability sequence of our next- k -token predictions with

*Corresponding Authors.

Copyright © 2025, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

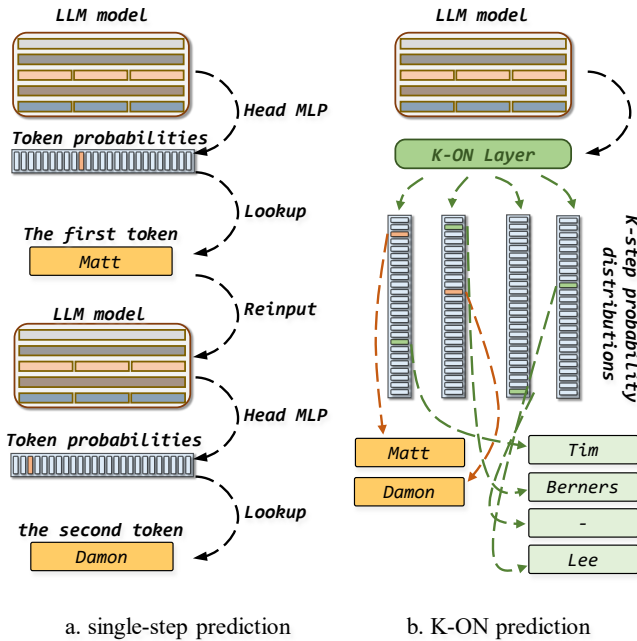


Figure 1: A comparison of single-step prediction and the proposed K-ON prediction. Left: in the conventional single-step prediction, obtaining an output of an entity necessitates recurrently feeding input data and cannot be parallelized across different entities. Right: the K-ON prediction generates an entity in a single step and allows for parallelization across multiple entities, thereby enabling entity-level contrastive learning.

the original estimates by minimizing their KL divergence. In this way, we expect to mimic the single-step prediction process using a set of learnable functions within K-ON.

We evaluate K-ON on the KG completion task without any simplification on the task setting. Our experiments demonstrate that K-ON not only outperforms the conventional methods, but also achieves better performance than the multi-modal methods that leverage additional textual and visual information. Furthermore, K-ON is also an efficient method: although incorporating the LLM requires more GPU resources, the number of training epochs is reduced from 1,000 to 5 compared with conventional methods. The overall fine-tuning time is less than 1 hour on the DB15K dataset with 8 A100 GPUs.

Related Works

Knowledge Graph Completion Knowledge Graph (KG) completion is one of the most important tasks in the KG area. Conventional methods leverage triplet information as training data but often ignore the rich contextual information embedded within the text (Bordes et al. 2013; Dettmers et al. 2018; Guo, Sun, and Hu 2019; Vashishth et al. 2020; Guo et al. 2020; Chen et al. 2021; Guo, Zhang, and Chen 2022; Guo et al. 2022). In other words, they assume that the entities carry no self-feature (including their names, e.g.,

Matt Damon), and the relational connections are the only informative source. Recently, methods leveraging text and image information have proposed and achieved state-of-the-art performance on many benchmarks (Xie et al. 2017; Wang et al. 2019; Yao, Mao, and Luo 2019; Youn and Tagkopoulos 2022; Lin et al. 2023; Zhang, Chen, and Zhang 2023; Guo et al. 2024b). They can be divided into two groups: one focuses on the integration of more modalities (Lu et al. 2022; Lee et al. 2023; Zhang et al. 2024c); and the other concentrates on the fine-tuning of language models to better encoding text information (Yao, Mao, and Luo 2019; Youn and Tagkopoulos 2022; Lin et al. 2023).

LLM-based Knowledge Graph Completion Entities possess rich features often represented in text forms such as descriptions, tables, and attributes. Many (mostly multi-modal) methods propose leveraging language models to encode text information and use the resulting representations for prediction (Yao, Mao, and Luo 2019; Wang et al. 2021; Lin et al. 2023; Youn and Tagkopoulos 2022; Chen et al. 2022; Zhang et al. 2024b,a; Guo et al. 2024a). In the recent advances, most LLM-based methods can be classified into this category. For example, LLMKGC (Zhu et al. 2023) directly feeds the textual triplets to ChatGPT (Achiam et al. 2023) for KG completion, although the results are not very promising. KGLlama (Yao et al. 2023) and KoPA (Zhang et al. 2023) fine-tune LLMs on triplet verification, i.e., estimating the correctness of a given triplet. These LLM-based methods employ in-context learning or LoRA-based fine-tuning (Dong et al. 2022; Wies, Levine, and Shashua 2024). To our knowledge, no prior work has explored integrating KGs into the head layer of LLMs. The current LLM-based methods leverage additional text information but are directly compared against conventional methods (Yao et al. 2023; Zhang et al. 2023; Wei et al. 2024), which may lead to unfair evaluations. Therefore, in this paper, we consider multi-modal datasets as benchmarks (Liu et al. 2019; Xu et al. 2022; Chen et al. 2024).

Multi-Head LLMs There are several works employing multiple head layers in LLMs. Medusa (Cai et al. 2024) proposes a tree attention mechanism for multi-step training and inference. The K different head layers are initialized with the original weights and then fine-tuned independently. MultiToken (Gloeckle et al. 2024) discovers that training LLMs and multiple head layers from scratch can outperform the single-head version, with this advantage being more significant in larger models. Unlike these methods, the K head layers in our approach are not only used for generating multiple future tokens in one step but also confine the output space to KGs and enable entity-level contrastive learning. Our work explores a new direction for integrating LLMs with KGs.

Methodology

In this section, we present the details of K-ON. We begin with a preliminary overview of knowledge graphs and large language models, and then illustrate the architecture and implementation of K-ON. Finally, we introduce head trajectory tuning as a self-supervised optimization method for K-ON.

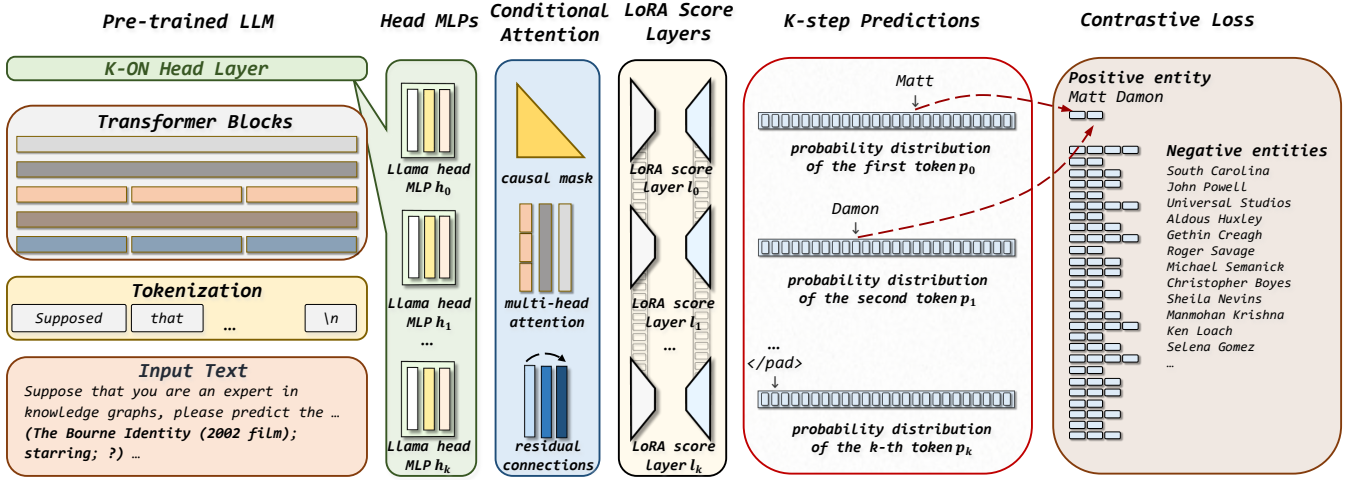


Figure 2: Overview of the K-ON architecture. From left to right: (1) The LLM processes the input text containing incomplete triplet information; (2) The resulting hidden states are input to distinct head MLPs within K-ON; (3) A compact conditional Transformer refines the corresponding outputs to capture sequential dependencies; (4) LoRA score layers are employed to transform the hidden states into K probability distribution estimations; (5-6) Aggregating the elements from the respective probability vectors, K-ON computes the probabilities of all candidate entities simultaneously.

Preliminaries

Knowledge Graphs We describe a knowledge graph by $\mathcal{G} = \{\mathcal{E}, \mathcal{R}, \mathcal{T}\}$, where $\mathcal{E}, \mathcal{R}, \mathcal{T}$ denote the entity, relation, and triplet sets, respectively. As one of the most important tasks in KG area, KG completion aims to predict the missing entity given an incomplete triplet (Bordes et al. 2013), i.e., predicting a tail entity e_2 given $(e_1, r_1, ?)$ or predicting the head entity e_1 given $(?, r_1, e_2)$.

Large Language Models Generally, a large language model comprises the following components: a tokenizer, which splits the input text into a sequence of N tokens $t_{0:N-1} = \{t_n | t_n \in \mathcal{V}\}_{n=0}^{N-1}$ with $|\mathcal{V}| \geq N$ being the vocabulary; a Transformer-based model \mathcal{M} , which processes the token sequence and generates a corresponding sequence of hidden states for prediction;

$$\mathbf{h}_{0:N-1}^m = \mathcal{M}(t_{0:N-1}); \quad (1)$$

and a head layer \mathcal{H} , which maps each hidden state to a probability distribution $\mathbf{p}_n \in \mathbb{R}^{|\mathcal{V}|}$:

$$\mathbf{p}_n = \mathcal{H}(\mathbf{h}_n^m). \quad (2)$$

In this paper, we focus primarily on the head layer of the LLM, and demonstrate that integrating knowledge graphs into the LLM at only this stage is sufficient to achieve state-of-the-art performance in KG completion.

K-ON

Entities typically require multiple tokens to be identifiable by name, which introduces a discrepancy between the entity distribution and the token probability distribution during prediction. Specifically, in a standard fine-tuning schema, the prediction objective is optimized at the token level rather than at the entity level. Consequently, the LLM is unaware

of which entities are present in the given KG, except those provided in the context.

Given the vast number of possible entities, it is impractical to include all candidate entities in the input text. An alternative approach worth exploring is manipulating the output probability distributions. If we can obtain the probability sequence of the constituent tokens for each entity, we can construct an entity-level entropy-based loss for the LLM. For example, suppose that we have constructed the input query containing the information of $e_1, r_1, ?$, we wish that the LLM precisely generates the output, e_2 , which comprises maximally K tokens. Then, we can extract the probability of each token and combine them as the joint probability for e_2 . However, achieving this is challenging, primarily due to computational costs. Iteratively feeding every negative example into the LLM for back-propagation is computationally intensive. As shown in Figure 1a, it is not parallelizable across entities.

To address this problem, we propose K-ON. Figure 2 illustrates the overall architecture of K-ON. In addition to the original LLM’s input and Transformer layers, K-ON introduces five new modules to support K -step token prediction.

Head MLPs We first employ multiple head MLPs to process the output hidden states of the LLM into inputs for different steps. Specifically, each MLP consists of three components: a fully-connected layer $\mathbf{W}_k^h \in \mathbb{R}^{d \times d}$, an activation function σ , and a normalization layer \mathbf{L}_k^h :

$$\mathbf{h}_{0:K-1}^h = \{\mathbf{L}_k^h(\sigma(\mathbf{W}_k^h \mathbf{h}_0^m))\}_{k=0}^{K-1}, \quad (3)$$

where $\mathbf{h}_{0:K-1}^h$ are the K output hidden states of the head MLPs. It is worth noting that the LLM uses a decoder-only architecture, and the input hidden state \mathbf{h}_0^m for K-ON is the last output hidden state of the query text. The task is to follow the query text and generate K subsequent tokens as pre-

dictions of entities. Similar to Llama 2 (Touvron et al. 2023), we use SiLU (Elfwing, Uchibe, and Doya 2018) and LlamaRMSNorm (Zhang and Sennrich 2019) as the activation function σ and normalization layer \mathbf{L}_k^h , respectively. The bias vector is not used.

Conditional Attention While the head MLPs in K-ON are independent of each other, the subsequent outputs in the original LLM are conditioned on the previous inputs. Therefore, we leverage a small Transformer \mathcal{M}_s to mimic this process by incorporating a causal mask $\mathbf{M} \in \mathbb{R}^{K \times K}$:

$$\mathbf{M}_{ij} = \begin{cases} 1 & i \geq j, \\ 0 & i < j, \end{cases} \quad (4)$$

where \mathbf{M}_{ij} indicates the value at the i -th row and j -th column of \mathbf{M} . When processing the k -th step, only the outputs of the previous $k - 1$ steps from the head MLPs can be observed for the Transformer.

Another noteworthy aspect of the conditional attention is the residual connection layer. Specifically, we add the attention output to the initial output of the LLM to produce the final output:

$$\mathbf{h}_k^a = \mathcal{M}_s(\mathbf{h}_{0:k}^h, \mathbf{M}) + \mathbf{h}_0^m, \quad (5)$$

where \mathcal{M}_s , \mathbf{M} are aforementioned small Transformer and causal mask, respectively. With the residual connection, we can initialize the head MLPs with zeros, allowing them to gradually learn the adaptation for K -step prediction from the initial LLM hidden \mathbf{h}_0^m .

LoRA Score Layer We use different score layers to estimate the probability distribution for each step. Unlike (Gloeckle et al. 2024) which trains each score layer from scratch, we propose to use a low-rank adaptation (LoRA) (Hu et al. 2021) layer for each step. This can be expressed as:

$$\mathbf{W}_k^S = \mathbf{W}^S + \mathbf{A}_k \mathbf{B}_k \quad (6)$$

$$\mathbf{p}_k = \mathbf{W}_k^S \mathbf{h}_k^a \quad (7)$$

where $\mathbf{W}^S \in \mathbb{R}^{|\mathcal{V}| \times d}$ is the original score layer of the LLM. $|\mathcal{V}|$ denotes the vocabulary size (number of tokens) of the LLM. $\mathbf{A}_k \in \mathbb{R}^{d \times r}$ and $\mathbf{B}_k \in \mathbb{R}^{r \times d}$ are the down-scaling and upper-scaling matrices in LoRA. The hyper-parameter $r \ll d$ is the reduced dimensionality. Before fine-tuning, \mathbf{A}_k is initialized randomly while \mathbf{B}_k is initialized to zero. This ensures that the output of the adaptation layers is identical to the original head layer.

K -step Gathering After obtaining the K -step predictions $\mathbf{p}_0, \mathbf{p}_1, \dots, \mathbf{p}_{K-1}$, we need to efficiently extract the elements relevant to each entity. To achieve this, we first convert each entity to an identifiable sequence of tokens:

$$t_{0:K-1}^e = \mathcal{P}(\tau(l^e), K) \quad (8)$$

where l^e is the textual label of the entity e and τ is the tokenizer. We perform padding and truncation \mathcal{P} to ensure that all entity token sequences have the same length K .

Algorithm 1: K-ON for KG Completion

- 1: **Input:** the training KG \mathcal{G} , the language model \mathcal{M} , the K-ON head layers $\mathcal{H}_{0:K-1}$;
 - 2: **for each** batched triplets in the training KG \mathcal{G} **do**
 - 3: Construct and tokenize the input queries; create the K -step labels for target and negative entities (Equation (8));
 - 4: $\mathbf{h}_0^m \leftarrow \mathcal{M}(t_{0:N-1})$, obtaining the output hidden states of LLM (Equation (1));
 - 5: $\mathbf{p}^e \leftarrow \mathcal{H}_{0:K-1}(\mathbf{h}_0^m)$, estimating K-ON predictions following Equations (3-10);
 - 6: Compute entity-level contrastive loss $\mathcal{L}_{\text{NCE}}(e)$ (Equation (11));
 - 7: Compute the supervised fine-tuning loss $\mathcal{L}_{\text{sft}}(e)$ (Equation (13));
 - 8: Compute the token distribution tuning loss $\mathcal{L}_{\text{tdt}}(e)$ (Equation (14));
 - 9: Jointly minimizing all three losses;
 - 10: **end for**
-

Next, we stack the K -step predictions $\mathbf{p}_0, \mathbf{p}_1, \dots, \mathbf{p}_{K-1}$ into a probability matrix \mathbf{P} :

$$\mathbf{P} = (\mathbf{p}_0 \quad \mathbf{p}_1 \quad \dots \quad \mathbf{p}_{K-1})^T, \quad (9)$$

such that the each token probability p_k^e can be extracted from \mathbf{P} with t_k^e as indices:

$$\mathbf{p}^e = \mathbf{P}_{0,t_0^e}, \mathbf{P}_{1,t_1^e}, \dots, \mathbf{P}_{K-1,t_{K-1}^e}. \quad (10)$$

Here, \mathbf{p}^e is the token probability sequence of entity e , with a strict length of K .

Contrastive Loss To incorporate the entity-level contrastive loss, we first estimate the scalar probability p^e from its token probability sequence \mathbf{p}^e . Our experimental results indicates that a weighted sum achieves the best performance:

$$p^e = \sum_{p_k \in \mathbf{p}^e} \alpha_k p_k, \quad (11)$$

where α_k is a learnable weight for k -th step shared across different entities.

By using the above equation to gather the scalar estimates for both positive and negative entities, we can construct an effective contrastive loss. Specifically, we randomly sample entities from the entity set \mathcal{E} to form negative examples $\mathcal{N} = \{e_j | e_j \neq e, e_j \in \mathcal{E}\}$:

$$\mathcal{L}_{\text{NCE}}(e) = -\log p^e + \frac{1}{|\mathcal{N}|} \sum_{e_j \in \mathcal{N}} \log p^{e_j}, \quad (12)$$

where p^e and p^{e_j} denote the joint probabilities for the positive entity e and negative entity e_j , respectively.

Head Trajectory Tuning

Entity-level contrastive learning is optimized against entities, which may inadvertently disrupt the token-level predictions of the original LLM, affecting performance on both common and training corpora. To mitigate this issue, we propose head trajectory tuning (HTT) to align the sequence estimations between single-step and K -step predictions.

Dataset	# Entity	# Relation	#Train	#Valid	#Test	# Text	# Image
DB15K	12,842	279	79,222	9,902	9,904	12,842	12,818
MKGW	15,000	169	34,196	4,276	4,274	14,123	14,463

Table 1: Statistics of the datasets.

Supervised Fine-Tuning HTT consists of two objectives. The first is tuning the LLM on the training corpus, also known as supervised fine-tuning (SFT). For this, we apply LoRA to the LLM model and optimize the single-step estimations against the ground truth:

$$\mathcal{L}_{\text{sft}}(e) = \sum_{k=0}^{K-1} \left(-\log p_k^e + \frac{1}{\mathcal{V}} \sum_{e_j \in \mathcal{V}} \log p_k^{e_j} \right), \quad (13)$$

where $\mathcal{L}_{\text{sft}}(e)$ is the supervised fine-tuning loss, with p_k^e denoting the target probability at the k -th step and \mathcal{V} representing the token vocabulary.

Token Distribution Tuning Then, we propose token distribution tuning to align the probability estimations of K-ON with those of the original LLM. Specifically, we minimize the KL-divergence (Kullback and Leibler 1951) between each pair of estimations along the output trajectories.

$$\mathcal{L}_{\text{idt}}(e) = \sum_{k=0}^{K-1} D_{\text{KL}}(\mathbf{p}_k^{e, \text{k-on}}, \mathbf{p}_k^{e, \text{llm}}), \quad (14)$$

where $p_k^{e, \text{k-on}}, p_k^{e, \text{llm}}$ denote the k -th token probability distributions of K-ON and the original LLM head, respectively.

Implementation

We present Algorithm 1 to illustrate the implementation of K-ON step by step. We first construct the input text for each triplet in the training set, and then use the tokenizer to convert the query into token IDs. We feed the input into the LLM to obtain the hidden states, which will be estimated by K-ON and the original head layer, respectively. Lastly, we compute the losses introduced before, and jointly minimize them until convergence.

Experiment

In this section, we conduct experiments to verify the effectiveness of the proposed K-ON.

Setting

We employ Llama-2-chat-7B (Touvron et al. 2023) as the base LLM model and train K-ON with 8 A100 GPUs. The learning rate is set to $1e-4$ in all experiments, and we use AdamW (Kingma and Ba 2015) as the optimizer. The batch-size per device is set to 12 and the gradient accumulation is set to 8 to obtaining a larger batch-size. We follow (Bordes et al. 2013; Lu et al. 2022; Zhang et al. 2024c) to report the MRR and Hits@K results with filtered ranks.

We consider various KG completion methods as baselines: the conventional structure-only methods, such as TransE (Bordes et al. 2013) and RotatE (Sun et al. 2019); the methods leveraging image information, such as IKRL (Xie

et al. 2017) and TransAE (Wang et al. 2019); the methods leveraging text information, such as KG-Bert (Yao, Mao, and Luo 2019) and FLT-LM (Lin et al. 2023); the methods leveraging both text and image information, such as MMKRL (Lu et al. 2022) and MANS (Zhang, Chen, and Zhang 2023); and the LLM-based methods KG-Llama-7b (Yao et al. 2023) and GPT 3.5 (Zhu et al. 2023);

Datasets

We consider DB15K and MKGW as benchmark, which are widely used in many recent works (Xie et al. 2017; Xu et al. 2022; Lee et al. 2023; Zhang, Chen, and Zhang 2023; Zhang et al. 2024c). The two datasets include not only the structural triplet data, but also the rich information of text and others. Thereby, we believe conducting experiments on them can gain a more comprehensive understanding on different methods and ensure a fairer comparison. The statistics of these two datasets are shown in Table 1.

Results

The main experimental results are shown in Table 2. We can find that the methods considering additional information generally perform better than the structure-only methods, which verifies the effectiveness of leveraging external informative sources. Among these methods, the proposed K-ON achieves the best results, significantly surpassing all baseline methods.

However, we also observe that the performance gap between conventional methods and multi-modal approaches narrows on the MKGW dataset. For instance, FLT-LM (Lin et al. 2023) and MANS (Zhang, Chen, and Zhang 2023) perform worse than RotatE on the MKGW dataset. This suggests that while additional information can be beneficial, it may also introduce noise, which does not always aid in entity prediction. Thanks to advancements in language models, K-ON leverages LLM to process text information and is optimized against an entity-level contrastive loss. Consequently, it still significantly outperforms all baselines across all metrics on the MKGW dataset.

The existing LLM-based methods, as mentioned in previous sections, are optimized against tokens rather than entities. Thus, their performance on the standard KG completion task is generally unsatisfactory, falling significantly short of many non-LLM methods. This phenomenon has also been examined by (Zhu et al. 2023). In contrast, our K-ON introduces entity-level contrastive loss, enabling the LLM to more effectively explore the KG structure.

Ablation Studies

We conduct ablation studies to verify the effectiveness of each module in K-ON. We remove or replace the core modules in K-ON to conduct experiments. Specifically, w/o \mathcal{L}_{idt} , w/o \mathcal{L}_{sft} , w/o \mathcal{L}_{ncc} refer to the methods that exclude the token distribution tuning loss, the token supervised fine-tuning loss, and the entity-level negative contrastive loss, respectively. w/o Conditional Attention is the method without the attention module. Shared Head MLP and Shared Score

Model	Modality	DB15K				MKGW			
		MRR↑	Hits@1↑	Hits@3↑	Hits@10↑	MRR↑	Hits@1↑	Hits@3↑	Hits@10↑
TransE (Bordes et al. 2013)	S	24.86	12.78	31.48	47.07	29.19	21.06	33.20	44.23
DistMult (Yang et al. 2015)	S	23.03	14.78	26.28	39.59	20.99	15.93	22.28	30.86
RotatE (Sun et al. 2019)	S	29.28	17.87	36.12	49.66	33.67	26.80	36.68	46.73
IKRL (Xie et al. 2017)	S+I	26.82	14.09	34.93	49.09	32.36	26.11	34.75	44.07
TransAE (Wang et al. 2019)	S+I	28.09	21.25	31.17	41.17	30.00	21.23	34.91	44.72
KG-Bert (Yao, Mao, and Luo 2019)	S+T	23.94	11.98	31.05	46.54	28.68	21.12	32.57	43.46
MMKRL (Lu et al. 2022)	S+T+I	26.81	13.85	35.07	49.39	30.10	22.16	34.09	44.69
OTKGE (Cao et al. 2022)	S+T+I	23.86	18.45	25.89	34.23	34.36	<u>28.85</u>	36.25	44.88
MMRNS (Xu et al. 2022)	S+T+I	32.68	23.01	37.86	51.01	<u>35.03</u>	28.59	37.49	<u>47.47</u>
KGLM (Youn and Tagkopoulou 2022)	S+T	28.47	17.66	36.02	48.89	34.12	27.01	36.87	46.62
QEB (Wang et al. 2023)	S+T+I	28.18	14.82	36.67	51.55	32.38	25.47	35.06	45.32
VISTA (Lee et al. 2023)	S+T+I	30.42	22.49	33.56	45.94	32.91	26.12	35.38	45.61
MANS (Zhang, Chen, and Zhang 2023)	S+T+I	28.82	16.87	36.58	49.26	30.88	24.89	33.63	41.78
FLT-LM (Lin et al. 2023)	S+T	33.45	24.56	37.67	50.12	32.75	25.89	32.87	44.56
AdaMF (Zhang et al. 2024c)	S+T+I	32.51	21.31	<u>39.67</u>	<u>51.68</u>	34.27	27.21	<u>37.86</u>	47.21
KG-Llama-7b (Yao et al. 2023)	S+T	-	13.46	-	-	-	20.20	-	-
GPT 3.5 Turbo (Zhu et al. 2023)	S+T	-	21.71	-	-	-	22.66	-	-
K-ON	S+T	38.10	30.13	42.77	53.59	36.64	30.05	38.72	48.26

Table 2: The main KG completion results. The best and second-best results are **boldfaced** and underlined, respectively. S, T, I indicate structure, text, and image, respectively. ↑: higher is better; ↓: lower is better. -: unavailable entry.

Model	MRR↑	Hits@1↑	Hits@3↑	Hits@10↑
K-ON	38.10	30.13	42.77	53.59
- w/o \mathcal{L}_{tdt}	37.48	28.43	42.34	<u>53.62</u>
- w/o \mathcal{L}_{sft}	37.31	28.07	<u>42.64</u>	53.42
- w/o \mathcal{L}_{nce}	14.09	10.40	15.24	21.29
- w/o Conditional Attention	37.20	27.69	42.13	53.60
- Shared Head MLP	28.01	19.57	31.98	44.73
- Shared Score Layer	<u>37.54</u>	<u>28.64</u>	42.12	53.73

Table 3: Ablation studies on the DB15K dataset. Shared Head MLP and Shared Score Layer are the methods with one shared adaption module instead of K different modules.

Layer refer to the methods where we replaced the K different LoRA layers for each step with a single shared LoRA layer for all steps.

The results are presented in Table 3. w/o \mathcal{L}_{nce} performs worst among all alternative methods, but it still shows some effectiveness, likely due to the presence of the token-level supervised fine-tuning loss. The substantial gap between K-ON and w/o \mathcal{L}_{nce} highlights the importance of entity-level contrastive learning. \mathcal{L}_{tdt} and \mathcal{L}_{sft} are also crucial for K-ON, and removing either one results in a significant performance drop in MRR, Hits@1, and Hits@10. Interestingly, we observe that the Hits@10 results are relatively insensitive to our ablation settings, showing only minimal differences across multiple methods.

Removing the conditional attention module also leads to a significant performance decline, particularly in Hits@1. We believe this module is essential for accurately identifying

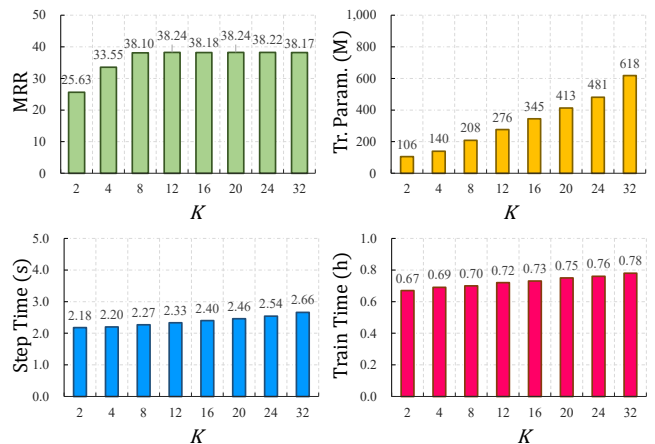


Figure 3: Performance of K-ON w.r.t. the number of K-ON head layers k . The results are obtained using 8 A100 GPUs.

target entities. We also develop two shared-weight variants. The method leverages one shared score layer achieves the second-best performance, suggesting that K-ON effectively shapes the k -step hidden states. It is similar to the original LLM that leverages one shared score layer to estimate the probability distribution of different steps. In contrast, the method with a shared head MLP suffers a significant performance drop, resulting in the second-worst performance. This empirically confirms that merely fine-tuning the score layer is insufficient for our task, even when using k different LoRA adoptions.




Figure 4: Performance of K-ON w.r.t. the number of negative entities. The results are obtained using 8 A100 GPUs.

Operator	Weight	MRR↑	Hits@1↑	Hits@3↑	Hits@10↑
+	learnable	38.10	30.13	42.77	53.59
	constant	<u>37.11</u>	<u>29.08</u>	<u>41.82</u>	<u>52.97</u>
*	learnable	23.24	15.95	26.37	37.70
	constant	23.61	16.36	26.97	37.93

Table 4: The results of K-ON with different estimation functions for the joint entity probability. + and * denote the addition and multiplication operations, respectively.

Analysis on the K-ON Heads

The number of K-ON heads (denoted as K) is a crucial hyper-parameter, determining the maximum token length available for representing each entity. While a larger K can prevent the truncation of entity names and potentially enhance model performance, it also increases computational demands. To examine this trade-off, we conducted experiments with varying values of K .

The results are illustrated in Figure 3, where we present four subgraphs depicting MRR, the number of trainable parameters, per-step time, and overall training time. Notably, performance appears to saturate when $K \geq 8$, likely because most entity names consist of fewer than 8 tokens. Nonetheless, the computational cost increases linearly with K . As observed, both step time and training time show slight increases with larger K , while the number of trainable parameters significantly rises. Therefore, we choose $K = 8$ for the main experiments, as it strikes an optimal balance between performance and computational efficiency.

Analysis on the Entity-level Contrastive Loss

One of the key features of K-ON is the entity-level contrastive loss, defined by one positive example and $|\mathcal{N}|$ negative examples. It is interesting to analyze how varying $|\mathcal{N}|$ impacts both the performance and computational cost of K-ON. As illustrated in Figure 4, increasing the number of negative examples does not necessarily lead to improved performance. Our observations reveal that the performance curves across all three metrics exhibit a similar trend: initially, there is a rapid increase, followed by a steady decline, and ultimately, a plateau phase.

In contrast to the effects of varying the number of head layers K , we find that the computational cost remains relatively stable even as $|\mathcal{N}|$ increases. This stability arises because the negative examples are extracted from the K head

output probability distributions of K-ON. Thus, adding more negative examples does not involve additional neural layers. Based on these findings, we have selected $|\mathcal{N}| = 128$ as the optimal setting for our main experiments.

Analysis on the Joint Probability Function

The entity-level contrastive loss requires the entity probability as input, which is derived from its named tokens. As such, selecting an appropriate method to combine these token probabilities into a joint probability for the entity is crucial. We investigate four different methods for estimating this joint probability. The operations of addition and multiplication are denoted by + and *, respectively. The term *learnable* refers to the version where we utilize a weight vector to aggregate the probabilities of the tokens at different positions, which is shared across entities. Conversely, *constant* indicates an unweighted aggregation.

Table 4 illustrates the results on the DB15K dataset. Notably, the learnable + method achieves the highest performance across all metrics. When we remove the learnable weights, there is a slight degradation in performance. In contrast, although the * operator appears to be intuitively more effective, it significantly underperforms compared to + across all metrics. We suspect that the lackluster performance of the multiplication method may be attributed to the issue of vanishing gradients. Specifically, multiplying multiple (8 in our experiments) probabilities may result in exceedingly small scalar values, although conceptually valid as a joint probability, might hinder the learning process.

Conclusion and Limitation

In this paper, we propose K-ON to stack the knowledge on the head layer of LLM. We introduce an entity-level contrastive loss that significantly reduces computational costs and develop HTT to align the output distributions with those of the original LLM head. Extensive experiments demonstrate the superior performance of K-ON compared to state-of-the-art baselines. K-ON still has limitations. First, its flexibility is constrained, as it does not support arbitrarily large values of K . To address this, we plan to explore a sliding window mechanism for processing K -step prediction in future; Second, K-ON currently lacks support for multimodal input. We also intend to incorporate large vision-language models into K-ON as part of our future work.

Acknowledgements

We would like to thank all anonymous reviewers for their insightful and invaluable comments. This work is founded by National Natural Science Foundation of China (NSFCU23B2055/NSFC62306276/NSFCU19B2027), Zhejiang Provincial Natural Science Foundation of China (No. LQ23F020017), Yongjiang Talent Introduction Programme (2022A-238-G), and Fundamental Research Funds for the Central Universities (226-2023-00138). This work was supported by AntGroup.

References

- Achiam, J.; Adler, S.; Agarwal, S.; Ahmad, L.; Akkaya, I.; Aleman, F. L.; Almeida, D.; Altenschmidt, J.; Altman, S.; Anadkat, S.; et al. 2023. Gpt-4 technical report. *arXiv preprint arXiv:2303.08774*.
- Bordes, A.; Usunier, N.; Garcia-Durán, A.; Weston, J.; and Yakhnenko, O. 2013. Translating embeddings for modeling multi-relational data. In *NIPS*, 2787–2795.
- Cai, T.; Li, Y.; Geng, Z.; Peng, H.; Lee, J. D.; Chen, D.; and Dao, T. 2024. Medusa: Simple LLM Inference Acceleration Framework with Multiple Decoding Heads. *CoRR*, abs/2401.10774.
- Cao, Z.; Xu, Q.; Yang, Z.; He, Y.; Cao, X.; and Huang, Q. 2022. OTKGE: Multi-modal Knowledge Graph Embeddings via Optimal Transport. In *NeurIPS*.
- Chen, S.; Liu, X.; Gao, J.; Jiao, J.; Zhang, R.; and Ji, Y. 2021. HittER: Hierarchical Transformers for Knowledge Graph Embeddings. In *EMNLP*, 10395–10407.
- Chen, Z.; Chen, J.; Zhang, W.; Guo, L.; Fang, Y.; Huang, Y.; Geng, Y.; Pan, J. Z.; Song, W.; and Chen, H. 2022. MEAformer: Multi-modal Entity Alignment Transformer for Meta Modality Hybrid. *arXiv preprint arXiv:2212.14454*.
- Chen, Z.; Zhang, Y.; Fang, Y.; Geng, Y.; Guo, L.; Chen, X.; Li, Q.; Zhang, W.; Chen, J.; Zhu, Y.; Li, J.; Liu, X.; Pan, J. Z.; Zhang, N.; and Chen, H. 2024. Knowledge Graphs Meet Multi-Modal Learning: A Comprehensive Survey. *CoRR*, abs/2402.05391.
- Dettmers, T.; Minervini, P.; Stenetorp, P.; and Riedel, S. 2018. Convolutional 2D knowledge graph embeddings. In *AAAI*, 1811–1818.
- Dong, Q.; Li, L.; Dai, D.; Zheng, C.; Wu, Z.; Chang, B.; Sun, X.; Xu, J.; and Sui, Z. 2022. A survey on in-context learning. *arXiv preprint arXiv:2301.00234*.
- Elfwing, S.; Uchibe, E.; and Doya, K. 2018. Sigmoid-weighted linear units for neural network function approximation in reinforcement learning. *Neural Networks*, 107: 3–11.
- Gloeckle, F.; Idrissi, B. Y.; Rozière, B.; Lopez-Paz, D.; and Synnaeve, G. 2024. Better & Faster Large Language Models via Multi-token Prediction. *CoRR*, abs/2404.19737.
- Guan, X.; Liu, Y.; Lin, H.; Lu, Y.; He, B.; Han, X.; and Sun, L. 2024. Mitigating large language model hallucinations via autonomous knowledge graph-based retrofitting. In *AAAI*, 18126–18134.
- Guo, L.; Bo, Z.; Chen, Z.; Zhang, Y.; Chen, J.; Lan, Y.; Sun, M.; Zhang, Z.; Luo, Y.; Li, Q.; Zhang, Q.; Zhang, W.; and Chen, H. 2024a. MKGL: Mastery of a Three-Word Language. *CoRR*, abs/2410.07526.
- Guo, L.; Chen, Z.; Chen, J.; Fang, Y.; Zhang, W.; and Chen, H. 2024b. Revisit and Outstrip Entity Alignment: A Perspective of Generative Models. In *ICLR*. OpenReview.net.
- Guo, L.; Sun, Z.; and Hu, W. 2019. Learning to Exploit Long-term Relational Dependencies in Knowledge Graphs. In *ICML*, 2505–2514.
- Guo, L.; Wang, W.; Sun, Z.; Liu, C.; and Hu, W. 2020. Decentralized Knowledge Graph Representation Learning. *CoRR*, abs/2010.08114.
- Guo, L.; Zhang, Q.; and Chen, H. 2022. Unleashing the Power of Transformer for Graphs. *CoRR*, abs/2202.10581.
- Guo, L.; Zhang, Q.; Sun, Z.; Chen, M.; Hu, W.; and Chen, H. 2022. Understanding and Improving Knowledge Graph Embedding for Entity Alignment. In *ICML*, volume 162 of *Proceedings of Machine Learning Research*, 8145–8156. PMLR.
- Hu, E. J.; Shen, Y.; Wallis, P.; Allen-Zhu, Z.; Li, Y.; Wang, S.; Wang, L.; and Chen, W. 2021. Lora: Low-rank adaptation of large language models. *arXiv preprint arXiv:2106.09685*.
- Kingma, D. P.; and Ba, J. 2015. Adam: A Method for stochastic optimization. In *ICLR*.
- Kullback, S.; and Leibler, R. A. 1951. On information and sufficiency. *The annals of mathematical statistics*, 22(1): 79–86.
- Lee, J.; Chung, C.; Lee, H.; Jo, S.; and Whang, J. J. 2023. VISTA: Visual-Textual Knowledge Graph Representation Learning. In *EMNLP (Findings)*, 7314–7328. Association for Computational Linguistics.
- Lin, Q.; Mao, R.; Liu, J.; Xu, F.; and Cambria, E. 2023. Fusing topology contexts and logical rules in language models for knowledge graph completion. *Information Fusion*, 90: 253–264.
- Liu, Y.; Li, H.; García-Durán, A.; Niepert, M.; Oñoro-Rubio, D.; and Rosenblum, D. S. 2019. MMKG: Multi-modal Knowledge Graphs. In *ESWC*, volume 11503 of *Lecture Notes in Computer Science*, 459–474. Springer.
- Lu, X.; Wang, L.; Jiang, Z.; He, S.; and Liu, S. 2022. MMKRL: A robust embedding approach for multi-modal knowledge graph representation learning. *Appl. Intell.*, 52(7): 7480–7497.
- Pan, J. Z.; Razniewski, S.; Kalo, J.-C.; Singhania, S.; Chen, J.; Dietze, S.; Jabeen, H.; Omeliyanenko, J.; Zhang, W.; Lissandrini, M.; et al. 2023. Large language models and knowledge graphs: Opportunities and challenges. *arXiv preprint arXiv:2308.06374*.
- Pan, S.; Luo, L.; Wang, Y.; Chen, C.; Wang, J.; and Wu, X. 2024. Unifying large language models and knowledge graphs: A roadmap. *IEEE Transactions on Knowledge and Data Engineering*.
- Sun, Z.; Deng, Z.-H.; Nie, J.-Y.; and Tang, J. 2019. RotatE: Knowledge graph embedding by relational rotation in complex space. In *ICLR*.

- Touvron, H.; Martin, L.; Stone, K.; Albert, P.; Almahairi, A.; Babaei, Y.; Bashlykov, N.; Batra, S.; Bhargava, P.; Bhosale, S.; et al. 2023. Llama 2: Open foundation and fine-tuned chat models. *arXiv preprint arXiv:2307.09288*.
- Vashishth, S.; Sanyal, S.; Nitin, V.; and Talukdar, P. P. 2020. Composition-based Multi-Relational Graph Convolutional Networks. In *ICLR*.
- Wang, B.; Shen, T.; Long, G.; Zhou, T.; Wang, Y.; and Chang, Y. 2021. Structure-augmented text representation learning for efficient knowledge graph completion. In *The Web Conference*, 1737–1748.
- Wang, X.; Meng, B.; Chen, H.; Meng, Y.; Lv, K.; and Zhu, W. 2023. TIVA-KG: A Multimodal Knowledge Graph with Text, Image, Video and Audio. In *ACM Multimedia*, 2391–2399. ACM.
- Wang, Z.; Li, L.; Li, Q.; and Zeng, D. 2019. Multi-modal Data Enhanced Representation Learning for Knowledge Graphs. In *IJCNN*, 1–8. IEEE.
- Wei, Y.; Huang, Q.; Kwok, J. T.; and Zhang, Y. 2024. KICGPT: Large Language Model with Knowledge in Context for Knowledge Graph Completion. *arXiv preprint arXiv:2402.02389*.
- Wies, N.; Levine, Y.; and Shashua, A. 2024. The learnability of in-context learning. *NeurIPS*, 36.
- Xie, R.; Liu, Z.; Luan, H.; and Sun, M. 2017. Image-embodied Knowledge Representation Learning. In *IJCAI*, 3140–3146. ijcai.org.
- Xu, D.; Xu, T.; Wu, S.; Zhou, J.; and Chen, E. 2022. Relation-enhanced Negative Sampling for Multimodal Knowledge Graph Completion. In *ACM Multimedia*, 3857–3866. ACM.
- 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, L.; Chen, H.; Li, Z.; Ding, X.; and Wu, X. 2023. Chatgpt is not enough: Enhancing large language models with knowledge graphs for fact-aware language modeling. *arXiv preprint arXiv:2306.11489*.
- Yao, L.; Mao, C.; and Luo, Y. 2019. KG-BERT: BERT for knowledge graph completion. *arXiv preprint arXiv:1909.03193*.
- Yao, L.; Peng, J.; Mao, C.; and Luo, Y. 2023. Exploring large language models for knowledge graph completion. *arXiv preprint arXiv:2308.13916*.
- Youn, J.; and Tagkopoulos, I. 2022. Kglm: Integrating knowledge graph structure in language models for link prediction. *arXiv preprint arXiv:2211.02744*.
- Zhang, B.; and Sennrich, R. 2019. Root Mean Square Layer Normalization. In *NeurIPS*, 12360–12371.
- Zhang, Y.; Chen, M.; and Zhang, W. 2023. Modality-Aware Negative Sampling for Multi-modal Knowledge Graph Embedding. In *IJCNN*, 1–8. IEEE.
- Zhang, Y.; Chen, Z.; Guo, L.; Xu, Y.; Hu, B.; Liu, Z.; Zhang, W.; and Chen, H. 2024a. Mixture of Modality Knowledge Experts for Robust Multi-modal Knowledge Graph Completion. *CoRR*, abs/2405.16869.
- Zhang, Y.; Chen, Z.; Guo, L.; Xu, Y.; Hu, B.; Liu, Z.; Zhang, W.; and Chen, H. 2024b. NativE: Multi-modal Knowledge Graph Completion in the Wild. In *SIGIR*, 91–101. ACM.
- Zhang, Y.; Chen, Z.; Liang, L.; Chen, H.; and Zhang, W. 2024c. Unleashing the Power of Imbalanced Modality Information for Multi-modal Knowledge Graph Completion. *CoRR*, abs/2402.15444.
- Zhang, Y.; Chen, Z.; Zhang, W.; and Chen, H. 2023. Making large language models perform better in knowledge graph completion. *arXiv preprint arXiv:2310.06671*.
- Zhu, Y.; Wang, X.; Chen, J.; Qiao, S.; Ou, Y.; Yao, Y.; Deng, S.; Chen, H.; and Zhang, N. 2023. Llms for knowledge graph construction and reasoning: Recent capabilities and future opportunities. *arXiv preprint arXiv:2305.13168*.