

KaLM: Knowledge-aligned Autoregressive Language Modeling via Dual-view Knowledge Graph Contrastive Learning

Anonymous ACL submission

Abstract

Autoregressive large language models (LLMs) pre-trained by next token prediction are inherently proficient in generative tasks. However, their performance on knowledge-driven tasks such as factual knowledge querying remains unsatisfactory. Knowledge graphs (KGs), as high-quality structured knowledge bases, can provide reliable knowledge for LLMs, potentially compensating for their knowledge deficiencies. Aligning LLMs with explicit, structured knowledge from KGs has been a challenge; previous attempts either failed to effectively align knowledge representations or compromised the generative capabilities of LLMs, leading to less-than-optimal outcomes. This paper proposes **KaLM**, a *Knowledge-aligned Language Modeling* approach, which fine-tunes autoregressive LLMs to align with KG knowledge via the joint objective of explicit knowledge alignment and implicit knowledge alignment. The explicit knowledge alignment objective aims to directly optimize the knowledge representation of LLMs through dual-view knowledge graph contrastive learning. The implicit knowledge alignment objective focuses on incorporating textual patterns of knowledge into LLMs through triple completion language modeling. Notably, our method achieves a significant performance boost in evaluations of knowledge-driven tasks, specifically embedding-based knowledge graph completion and generation-based knowledge graph question answering¹.

1 Introduction

Large language models (LLMs) like PaLM 2 (Anil et al., 2023) and GPT-4 (Achiam et al., 2023) have recently made remarkable advancements in a wide range of natural language processing tasks (Li et al., 2022; Su et al., 2019). However, LLMs still face challenges in tasks requiring factual or domain-specific knowledge, resulting in unsatisfactory per-

formance in knowledge-driven tasks. From the perspective of knowledge representation, LLMs serve as parametric knowledge bases, providing implicit, non-deterministic knowledge, while knowledge graphs (KGs) function as structured knowledge bases, offering explicit, deterministic knowledge. KGs, commonly organized as factual knowledge triples describing relations between entities, can serve as a reliable knowledge source for LLMs. Aligning LLMs with KG knowledge can enhance the knowledge reasoning capabilities of LLMs and improve their performance on knowledge-driven tasks, such as knowledge graph completion (KGC) and knowledge graph question answering (KGQA).

Autoregressive LLMs pre-trained through next token prediction tasks often exhibit limitations in knowledge representation, leading to embeddings that lack diversity and specificity. This limitation becomes evident in tasks that demand distinctive sentence embeddings, such as dense retrieval and semantic search (Muennighoff, 2022; Ma et al., 2023). As demonstrated in Figure 1(a), the representations generated by LLMs tend to be overly homogeneous across different pieces of knowledge, undermining their effectiveness in applications requiring fine-grained semantic distinctions.

The concept of explicit knowledge alignment is introduced to directly optimize the knowledge representation within language models by devising direct knowledge training objectives. This strategy emerges in response to the observed degradation in knowledge representation within autoencoder-based pre-trained language models (PLMs), a phenomenon termed *representation anisotropy* (Ethayarajh, 2019). This issue is characterized by the clustering of learned token and sentence embeddings within a constrained area of the representation space, leading to a lack of distributional uniformity (Li et al., 2020). While previous efforts to address representation anisotropy have largely concentrated on promoting uniformity among to-

¹Our code is available at <https://anonymous.4open.science/r/KaLM-ARR>.




Figure 1: Similarity matrix of knowledge representations of (a) Llama-2-7B (Touvron et al., 2023) and (b) KaLM. The values denote the cosine similarity between the head-relation and tail embedding. The diagonal elements represent positive <head-relation, tail> pairs from the same KG triple, which should maintain high similarity (darker color); off-diagonal elements represent negative <head-relation, tail> pairs from different KG triples, which should have lower similarity (lighter color). In an ideal setting, knowledge representations should be able to distinguish between different triples, while maintaining alignment and uniformity of the representation, as shown in Figure 1(b).

ken representations, they often overlook the critical alignment of similar sentence representations (Su et al., 2021; Li et al., 2020; Su et al., 2022). More recent works advocate for integrating KG triples and using knowledge graph embedding losses to fine-tune PLMs, aiming to bolster their knowledge representation abilities (Shen et al., 2022; Wang et al., 2022b). Nonetheless, such approaches may limit themselves to optimizing at the token level or reduce the model to a mere text encoder, thereby diminishing its inherent generative capabilities.

Conversely, implicit knowledge alignment leverages the pre-training or fine-tuning of language models with external knowledge sources, employing the vanilla language modeling objective or its variations. This approach predominantly preserves the next token prediction framework, essentially retaining the native text generation prowess of LLMs. In the realm of implicit knowledge alignment, the prevalent practice involves the fine-tuning of LLMs with KG triples and their textual descriptions, as opposed to directly altering the hidden knowledge representations (Chen et al., 2022; Yao et al., 2023). Nevertheless, the efficacy of these methods on knowledge graph completion tasks remains substantially inferior when compared to strategies that directly fine-tune knowledge representations (Wang et al., 2022b,a). Intriguing findings from (Fu et al., 2023) reveal that fine-tuning PLMs with randomly unaligned KG triples can achieve per-

formance on par with that obtained through fine-tuning with aligned triples in various tasks, including named entity recognition and relation classification. Their findings suggest that the hidden states of entities, whether infused with aligned or random knowledge, exhibit remarkable similarity. Consequently, existing implicit alignment methods fail to effectively utilize the injected knowledge or accurately discern the connection between newly introduced knowledge and the model’s inherent knowledge, culminating in suboptimal performance.

In this paper, we propose **KaLM**, a *Knowledge-aligned Language Modeling* approach for aligning LLMs with KG knowledge. Specifically, we use KG triples and their textual descriptions to fine-tune LLMs via the joint objective of *explicit knowledge alignment* and *implicit knowledge alignment*.

The explicit knowledge alignment objective aims to directly optimize the hidden representations of knowledge in LLMs through *dual-view knowledge graph contrastive learning*. We theoretically prove and empirically show that this objective can facilitate knowledge representation alignment and alleviate representation anisotropy. For KG triples, we consider tail entity description and the concatenation of head entity description and relation description as two distinct views of the same knowledge. *The key insight is that: (1) representations of two different views of the same knowledge (i.e., from the same triple) should be pulled together, while (2)*

representations of different knowledge (i.e., from different triples) should be pushed apart. The first term encourages semantically similar knowledge to remain close in the representation space, promoting knowledge representation alignment. The second term forces dissimilar knowledge to be as far apart as possible in the vector space, improving knowledge representation uniformity and mitigating representation anisotropy. As shown in Figure 1(b), our method can obtain the ideal knowledge representations that are both aligned and uniform.

The implicit knowledge alignment objective focuses on incorporating textual patterns of knowledge into LLMs through *triple completion language modeling*, which can maintain the generative capability of LLMs and boost performance on knowledge inference tasks. We constructed a triple completion dataset based on the KG triples to fine-tune LLMs, improving their instruction-following ability and facilitating implicit knowledge alignment. We also show the implicit knowledge alignment objective can further boost knowledge representation performance. This confirms that both explicit alignment and implicit alignment are crucial for knowledge alignment, as they both essentially require a deep understanding of knowledge.

Our contributions are summarized as follows:

- We introduce **KaLM**, a *knowledge-aligned language modeling* approach that aligns autoregressive LLMs with KG knowledge via the joint objective of *explicit knowledge alignment* and *implicit knowledge alignment*.
- We *theoretically prove and empirically demonstrate* that the explicit knowledge alignment objective achieved through dual-view knowledge graph contrastive learning can facilitate knowledge representation alignment and alleviate the issue of representation anisotropy.
- The experimental results on knowledge-driven tasks demonstrate the effectiveness of KaLM. In the embedding-based KGC task, KaLM significantly improves Mean Rank and Hit@10 metrics compared to previous state-of-the-art methods. In the generation-based KGQA task, KaLM achieves a notable improvement in answering accuracy compared to the base LLM.

2 Related Work

Our work is closely related to Knowledge Enhancement for LLMs and Representation Anisotropy of

Language Models. A more detailed review of related work can be found in Appendix A.

Knowledge Enhancement for LLMs Knowledge enhancement aims to incorporate factual and domain-specific knowledge into LLMs to address their knowledge deficiencies. This can be divided into retrieval-based augmentation and training-based integration. *Retrieval-based knowledge augmentation* methods leverage external retrieval modules to provide additional knowledge, aiming to improve the knowledge reasoning capability of LLMs (Sun et al., 2023; Jiang et al., 2023). However, this approach may lead to knowledge conflicts (Feng et al., 2023), where knowledge in LLMs and knowledge in the retrieved documents are inconsistent or the retrieved multiple documents are contradictory. *Training-based knowledge integration* methods involve using KG triple descriptions to pre-train or fine-tune LLMs, aiming to achieve knowledge alignment. These methods can be divided into explicit alignment (Wang et al., 2021b; Yasunaga et al., 2022) and implicit alignment (Yao et al., 2023; Zhang et al., 2023) based on whether they directly optimize the knowledge representation. Nevertheless, prior methods have either sacrificed the generative capability or lacked effective representation alignment. Our approach enhances the knowledge of LLMs via a unique joint objective of explicit alignment and implicit alignment, improving the quality of knowledge representations and generative knowledge reasoning capabilities.

Representation Anisotropy of Language Models PLMs have long been plagued by representation anisotropy (Ethayarajh, 2019), where the learned token and sentence embeddings are confined to a narrow cone within the entire representation space. The issue of representation anisotropy not only results in model degradation (Su et al., 2022) but also leads to poor performance on discriminative tasks. Previous work on alleviating representation anisotropy has mainly focused on post-processing techniques such as normalizing flows (Li et al., 2020) or whitening operations (Su et al., 2021). Su et al. (2022) propose a contrastive training objective to encourage learning isotropic token representations. However, these methods mainly improve the isotropy of token representations without enhancing the discriminability of sentence representations. Our method improves the token-level and sentence-level representation anisotropy of LLMs through dual-view knowledge graph contrastive learning, and it has rigorous theoretical guarantees.

3 Knowledge-aligned Autoregressive Language Modeling

In this section, we introduce **KaLM**, a *Knowledge-aligned Language Modeling* approach for aligning LLMs with KG knowledge via the joint objective of *explicit knowledge alignment* and *implicit knowledge alignment*. The overview is shown in Figure 2.

3.1 Notations and Preliminaries

A KG \mathcal{G} stores factual knowledge, denoted as $\mathcal{G} = (\mathcal{E}, \mathcal{R}, \mathcal{T}, \mathcal{D})$. \mathcal{E} and \mathcal{R} are the set of entities and relations, respectively. \mathcal{D} is the description set of all entities and relations. \mathcal{D}_e and \mathcal{D}_r are the textual description of entity e and relation r , respectively. $\mathcal{T} = \{(h, r, t) | h, t \in \mathcal{E}, r \in \mathcal{R}\}$ is the triple set. A triple (h, r, t) depicts the fact that there is a relation r between the head entity h and the tail entity t .

3.2 Explicit Knowledge Alignment

For KG triples, the textual description of the tail entity and the concatenation of the textual descriptions of the head entity and relation can be seen as two distinct views of the same knowledge. This inspires *KaLM* to align representations of two distinct views of the same knowledge (i.e., from the same triple), while separating representations of different knowledge (i.e., from different triples).

The LLM, denoted as E_{LLM} , is fine-tuned with the *dual-view knowledge graph contrastive learning* loss. The training corpus contains paired textual descriptions, $\{(\mathcal{D}_{hr}, \mathcal{D}_t)\}_{i=1}^N$, where \mathcal{D}_t is the tail entity description, and \mathcal{D}_{hr} is the concatenation of the head entity description and relation description. Given a training pair $(\mathcal{D}_{hr}, \mathcal{D}_t)$, the same E_{LLM} is used to compute the embeddings of \mathcal{D}_{hr} and \mathcal{D}_t independently. Moreover, we prepend the [bos] token to the beginning and append the [eos] token to the end of the textual description. The augmented input is fed into E_{LLM} , and the hidden representation corresponding to the [eos] token from the last layer is used as the final embedding of the input.

$$e_{hr} = E_{LLM}([\text{bos}]_{hr} \oplus \mathcal{D}_{hr} \oplus [\text{eos}]_{hr}),$$

$$e_t = E_{LLM}([\text{bos}]_t \oplus \mathcal{D}_t \oplus [\text{eos}]_t),$$

where \oplus is the operation to concatenate two strings and $\mathcal{D}_{hr} = \mathcal{D}_h \oplus \mathcal{D}_r$. For stable training, we adopt “[” as $[\text{bos}]_{hr}$ and “]” as $[\text{eos}]_{hr}$, while using “{” as $[\text{bos}]_t$ and “}” as $[\text{eos}]_t$.

We utilize the knowledge graph contrastive learning loss to directly optimize the knowledge representation of the LLM by *encouraging semantically*

similar knowledge to stay close in the representation space and pushing dissimilar knowledge to be far apart in the representation space. More specifically, we apply the InfoNCE loss with an additive margin over the in-batch negatives to fine-tune the model. The row-direction loss ℓ_r is as follows for a given positive pair, and the column-direction loss ℓ_c is defined similarly (see Appendix C.2).

$$\ell_r = -\log \frac{e^{(\phi(e_{hr}, e_t) - \gamma) / \tau}}{e^{(\phi(e_{hr}, e_t) - \gamma) / \tau} + \sum_{i=1}^{\mathcal{N}} e^{\phi(e_{hr}, e_{t'_i}) / \tau}}, \quad (1)$$

where \mathcal{N} is the negative batch size, τ is the trainable temperature that controls the strength of penalties on hard negative samples, ϕ is the cosine similarity function that measures the plausibility of a triple, and γ is the additive margin that encourages increasing the similarity score of positive pairs.

The training objective for **explicit** knowledge alignment is the sum of the ℓ_r and the ℓ_c losses:

$$\mathcal{L}_{exp} = \frac{1}{\mathcal{N}} \sum_{(\mathcal{D}_{hr}, \mathcal{D}_t)} (\ell_r + \ell_c) / 2. \quad (2)$$

3.3 Implicit Knowledge Alignment

The implicit knowledge alignment objective focuses on incorporating textual patterns of knowledge into the LLM to prevent catastrophic forgetting of previous knowledge and maintain its generative capability. We constructed an instruction-tuning dataset based on the KG triple descriptions to fine-tune the model through *triple completion language modeling*. We also show that the implicit knowledge alignment objective can bring performance boosts on knowledge representation evaluations. This indicates that explicit alignment and implicit alignment are both imperative for effective knowledge alignment, as they both essentially necessitate a profound understanding of knowledge.

We follow the recipe of Stanford Alpaca (Taori et al., 2023) and use the provided template to construct the instruction-tuning dataset. The instruction passed to the template, abbreviated as inst, is: “Given the head entity and relation, write a tail entity that completes the triple”. The input and output are \mathcal{D}_{hr} and \mathcal{D}_t , respectively. The training objective for **implicit** knowledge alignment is:

$$\mathcal{L}_{imp} = \frac{1}{\mathcal{M}} \sum_{(\mathcal{D}_{hr}, \mathcal{D}_t)} -\log P(\mathcal{D}_t | \text{inst}, \mathcal{D}_{hr}), \quad (3)$$

where \mathcal{M} is the instruction-tuning batch size.

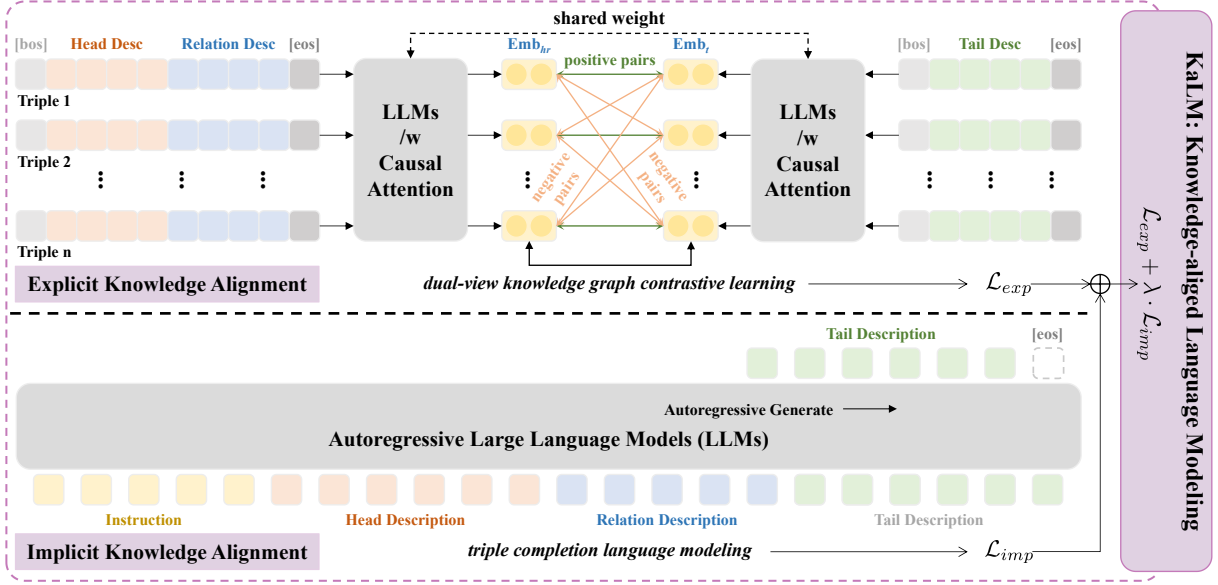


Figure 2: The overall framework of **KaLM**. **Up**: The explicit knowledge alignment objective (\mathcal{L}_{exp}) aims to directly optimize the knowledge representation of LLMs via dual-view knowledge graph contrastive learning. **Down**: The implicit knowledge alignment objective (\mathcal{L}_{imp}) focuses on incorporating textual patterns of knowledge into LLMs via triple completion language modeling. The final training objective is the weighted average of \mathcal{L}_{exp} and \mathcal{L}_{imp} .

3.4 Knowledge-aligned Language Modeling

The ultimate training objective of our proposed **KaLM** is the weighted average of \mathcal{L}_{exp} and \mathcal{L}_{imp} :

$$\mathcal{L}_{KaLM} = \mathcal{L}_{exp} + \lambda \cdot \mathcal{L}_{imp}, \quad (4)$$

where λ is a hyperparameter that adjusts the relative weight between them. Notably, this formulation allows us to use different batch sizes for explicit knowledge alignment (\mathcal{N}) and implicit knowledge alignment (\mathcal{M}). Previous work has shown that a sufficiently large batch size is key to the success of contrastive representation learning (Chen et al., 2020). With Equation 4, we can significantly increase the explicit knowledge alignment batch size while keeping the implicit knowledge alignment batch size fixed to save computational resources.

4 Theoretical Analysis

We theoretically prove that the explicit knowledge alignment objective implemented through dual-view knowledge graph contrastive learning can facilitate knowledge representation alignment and alleviate the issue of representation anisotropy.

4.1 Dual-view Contrastive Learning for Knowledge Representation Alignment

The outstanding performance of contrastive representation learning has attracted researchers to analyze its underlying reasons for success from a theoretical perspective. Wang and Isola (2020) identify

alignment and uniformity as two key properties of contrastive learning and propose two quantifiable metrics to measure the quality of representations.

We concentrate on understanding the dual-view knowledge graph contrastive learning loss from the knowledge alignment and uniformity perspective. To simplify the notation, we use f to denote E_{LLM} .

Alignment computes the expected distance between positive pairs and encourages the learned representations for positive pairs to be similar. *Uniformity* evaluates the even distribution of representations and encourages the separation of features from randomly selected negative samples.

$$\ell_{align}(f; \alpha) \triangleq \mathbb{E}_{(\mathcal{D}_{hr}, \mathcal{D}_t) \sim p_{pos}} [\|f(\mathcal{D}_{hr}) - f(\mathcal{D}_t)\|_2^2], \quad (374)$$

$$\ell_{uniform}(f; t) \triangleq \log \mathbb{E}_{\mathcal{D}_i, \mathcal{D}_j \stackrel{i.i.d.}{\sim} p_{data}} [e^{-t\|f(\mathcal{D}_i) - f(\mathcal{D}_j)\|_2^2}], \quad (375)$$

where p_{pos} denotes the distribution of positive pairs $\{(\mathcal{D}_{hr}, \mathcal{D}_t)\}_{i=1}^N$ and p_{data} represents the data distribution of textual descriptions $\{\mathcal{D}_i\}_{i=1}^N$.

Since the learned knowledge representations are L2-normalized, we have $\phi(e_{hr}, e_t) = f(x)^\top f(y)$. The additive margin γ encourages the model to learn more robust features without affecting the asymptotic analysis, thus we ignore it. For ease of analysis, we reformulate the contrastive learning

objective of Equation 1 and 2 as follows:

$$\mathcal{L}_{\text{exp}}(f; \tau, \mathcal{N}) \triangleq \mathbb{E}_{(\mathcal{D}_{hr}, \mathcal{D}_t) \sim p_{\text{pos}} \atop \{\mathcal{D}_{t_i}'\}_{i=1}^N \stackrel{i.i.d.}{\sim} p_{\text{data}}} \left[-\log \frac{e^{f(\mathcal{D}_{hr})^\top f(\mathcal{D}_t)/\tau}}{e^{f(\mathcal{D}_{hr})^\top f(\mathcal{D}_t)/\tau} + \sum_{i=1}^N e^{f(\mathcal{D}_{hr})^\top f(\mathcal{D}_{t_i}')/\tau}} \right], \quad (5)$$

Following Wang and Isola (2020), we analyze the asymptotics of the objective in Equation 5.

Theorem 1 (Asymptotics of \mathcal{L}_{exp}). *For temperature $\tau > 0$, as the number of negative samples $\mathcal{N} \rightarrow \infty$, the normalized dual-view knowledge graph contrastive loss in Equation 5 converges to*

$$\lim_{\mathcal{N} \rightarrow \infty} \mathcal{L}_{\text{exp}}(f; \tau, \mathcal{N}) - \log \mathcal{N} = -\frac{1}{\tau} \mathbb{E}_{(\mathcal{D}_{hr}, \mathcal{D}_t) \sim p_{\text{pos}}} \left[f(\mathcal{D}_{hr})^\top f(\mathcal{D}_t) \right] + \mathbb{E}_{\mathcal{D}_i \sim p_{\text{data}}} \left[\log \mathbb{E}_{\mathcal{D}_i^- \sim p_{\text{data}}} \left[e^{f(\mathcal{D}_i^-)^\top f(\mathcal{D}_i)/\tau} \right] \right]. \quad (6)$$

We have the following conclusions:

1. By pulling together the representations of two different views of the same knowledge, the first term of Equation 6 is minimized, and the encoder E_{LLM} is perfectly knowledge-aligned.
2. Assuming the perfect uniform knowledge encoder E_{LLM} exists, it precisely minimizes the second term of Equation 6 by pushing away the representations of different knowledge.

Proof. See Appendix B.1. \square

4.2 Alleviation of Representation Anisotropy

We then prove that the dual-view knowledge graph contrastive learning objective can directly alleviate representation anisotropy and improve the discriminability of knowledge representations.

Let \mathbf{E} be the sentence embedding matrix of $\{\mathcal{D}_i\}_{i=1}^N$, where the i -th row of \mathbf{E} is e_i . Following Ethayarajh (2019), the sentence-level representation anisotropy value of $\{\mathcal{D}_i\}_{i=1}^N$ is defined as:

$$\text{anisotropy}_{\{\mathcal{D}\}} = \frac{1}{N(N-1)} \sum_{i=1}^N \sum_{j=1, j \neq i}^N e_i^\top e_j. \quad (7)$$

We can further derive the following theorem.

Theorem 2 (Alleviation of Anisotropy). *When p_{data} is uniform over finite samples $\{\mathcal{D}_i\}_{i=1}^N$, the second term of Equation 6 is the upper bound of the sentence-level anisotropy of $\{\mathcal{D}_i\}_{i=1}^N$, i.e.,*

$$\mathbb{E}_{\mathcal{D}_i \sim p_{\text{data}}} \left[\log \mathbb{E}_{\mathcal{D}_i^- \sim p_{\text{data}}} \left[e^{f(\mathcal{D}_i^-)^\top f(\mathcal{D}_i)/\tau} \right] \right] \geq \frac{N-1}{\tau N} \cdot \text{anisotropy}_{\{\mathcal{D}\}} + \frac{1}{\tau N}. \quad (8)$$

We have the following result: By optimizing the second term of Equation 6, we essentially minimize the upper bound of the sentence-level anisotropy of corpus $\{\mathcal{D}_i\}_{i=1}^N$, thereby directly alleviating the representation anisotropy problem.

Proof. See Appendix B.2. \square

5 Experiments

In this section, we assess the effectiveness of KaLM in knowledge alignment. The experimental setup is outlined in 5.1. In 5.2 and 5.3, we present results on knowledge graph completion (KGC) and knowledge graph question answering (KGQA). In 5.4, we provide further analysis of knowledge representation and present case studies of KGQA generations.

5.1 Experimental Setup

Datasets. We use WN18RR (Dettmers et al., 2018) and FB15k-237 (Toutanova and Chen, 2015) as the KGs for knowledge alignment training. WN18RR and FB15k-237 are derived from WordNet and Freebase, respectively (Bordes et al., 2013). We use the information provided by KG-BERT (Yao et al., 2019) for textual descriptions. Following Wang et al. (2022a), we add an inverse triple (t, r^{-1}, h) for each triple (h, r, t) in the triple set, where r^{-1} is the inverse relation of the original relation r .

Model Training. We choose Llama-2-7B, Llama-3-8B, and Mistral-7B as base LLMs and fine-tune them through the joint objective of explicit knowledge alignment and implicit knowledge alignment. To save computational resources for parameter-efficient fine-tuning, we use LoRA (Hu et al., 2021) to fine-tune the feed-forward network of the model.

Evaluation Details. Experiments mainly focus on two aspects: knowledge representation assessment and knowledge inference evaluation. For *knowledge representation assessment*, we evaluate the embedding-based KGC task and illustrate the alleviation of representation anisotropy. We report five automated metrics: Mean Rank (MR), Mean Reciprocal Rank (MRR), and Hit@ k ($k \in \{1, 3, 10\}$).

Table 1: Embedding-based KGC results on WN18RR and FB15k-237. Baseline results are from their papers, with “-” indicating a missing result. The best and second best results are marked by **bold** and underline, respectively.

Method	WN18RR					FB15k-237				
	MR	MRR	H@1	H@3	H@10	MR	MRR	H@1	H@3	H@10
<i>structure-based methods</i>										
TransE	2300	0.243	0.043	0.441	0.532	323	0.279	0.198	0.376	0.441
DistMult	7000	0.444	0.412	0.470	0.504	512	0.281	0.199	0.301	0.446
RotatE	3340	0.476	0.428	0.492	0.571	177	<u>0.338</u>	0.241	0.375	<u>0.533</u>
<i>description-based methods (autoencoder PLMs)</i>										
StAR	51	0.401	0.243	0.491	0.709	117	0.296	0.205	0.322	0.482
C-LMKE	79	0.619	0.523	0.671	0.789	141	0.306	0.218	0.331	0.484
SimKGC	-	0.671	0.585	<u>0.731</u>	0.817	-	0.333	<u>0.246</u>	0.362	0.510
<i>description-based methods (autoregressive LLMs)</i>										
Llama-2-7B	15969	0.010	0.004	0.010	0.020	5359	0.006	0.002	0.004	0.012
Llama2-7B _{KaLM}	22	0.596	0.497	0.706	0.856	114	0.305	0.227	0.348	0.511
Llama3-8B _{KaLM}	28	0.613	0.502	0.728	<u>0.865</u>	125	0.327	0.241	0.366	0.522
Mistral-7B _{KaLM}	<u>25</u>	<u>0.634</u>	<u>0.526</u>	0.743	0.872	<u>117</u>	0.341	0.252	<u>0.371</u>	0.536





Figure 3: Comparison of generative knowledge inference performance between Llama-2-7B and KaLM. ↑ means higher is better and ↓ means lower is better.



(a) LLaMA

(b) KaLM

Figure 4: Similarity matrix on the Wikitext-103 test set. From top-left to bottom-right, element (i, j) denotes the cosine similarity between the i -th and the j -th sentence.

We compare KaLM with structure- and description-based methods. Structured-based methods include TransE (Bordes et al., 2013), DistMult (Yang et al., 2015), and RotatE (Sun et al., 2018). Description-based methods include StAR (Wang et al., 2021a), C-LMKE (Wang et al., 2022b), and SimKGC (Wang et al., 2022a). For *knowledge inference evaluation*, we evaluate the generation-based KGQA task and analyze the PPL metric and MMLU score (Hendrycks et al., 2020). We report the prediction accuracy over entities, relations, and triples. We also provide case studies of KGQA generations.

Additional experimental results and detailed ablation studies can be found in Appendix D and E.

5.2 Knowledge Representation Assessment

The embedding-based KGC results are shown in Table 1 and Table 5. The base LLM failed to finish this task, with all metrics lagging far behind. On the WN18RR dataset, our method surpasses prior methods by a substantial margin in terms of MR

and Hit@10 and achieves the best or second-best results in other metrics. Previous description-based methods generally perform poorly on FB15k-237, possibly due to the absence of effective textual descriptions. An example relation description from FB15k-237 is “/music/artist/origin”, which is quite vague and abstract. The performance of KaLM on the FB15k-237 dataset is also highly impressive, with most metrics surpassing previous SOTA methods (i.e., structure-based RotatE method), indicating that our approach enables autoregressive LLMs to learn the latent links in KGs effectively. Under similar experimental settings, more powerful LLMs (such as Llama3-8B and Mistral-7B) achieved significantly better performance after fine-tuning with KaLM, further demonstrating the scalability and generalizability of the proposed method. Overall, the results reveal that KaLM notably enhances the quality of knowledge representation, bringing substantial performance boosts in KGC tasks.

Task Name	Prompts with Instrucion and Input Fields	Generations for Triple 1: <salviniaceae, member meronym, salvinia>		Generations for Triple 2: <refrigerator, hypernym, white goods>	
		LLaMA	KaLM	LLaMA	KaLM
head entity prediction	Given the head entity and relation, write a tail entity that completes the triple: [tail entity], [inverse relation]	salvinia ✗	salviniaceae ✓	white goods ✗	refrigerator ✓
relation prediction	What is the relation between [head entity] and [tail entity]? Please choose your answer from: [relation list].	synset domain topic of ✗	member meronym ✓	instance hypernym ✗	synset domain topic of ✗
tail entity prediction	Given the head entity and relation, write a tail entity that completes the triple: [head entity], [relation]	salvinia ✓	salvinia ✓	refrigerator ✗	white goods ✓
triple classification	Is this true: [head] [relation] [tail]? Please choose your answer from: "Yes, this is true" or "No, this is not true".	No, this is not true. ✗	Yes, this is true. ✓	Yes, this is true. ✓	Yes, this is true. ✓

Figure 5: Case studies of Llama-2-7B and KaLM on KGQA tasks. Note that the head entity, relation, and tail entity are denoted by different colors. The ✓ mark indicates the correct answer, while ✗ signifies an incorrect answer.

5.3 Knowledge Inference Evaluation

The generation-based KGQA results are depicted in Figure 3 and Figure 6. Llama-2-7B performs poorly in entity prediction and relation prediction. Our method demonstrates a significant performance boost in all generation-based KGQA tasks, including head/tail entity prediction, relation prediction, and triple classification. Despite a slight increase in PPL scores on Wikitext-103 (Merity et al., 2016) test set, our method still shows improved performance in the MMLU test. The results demonstrate that KaLM achieves effective knowledge alignment, bringing in significantly improved KGQA performance while preserving the original generative and knowledge inference capabilities.

5.4 Visualization of Knowledge Representation and Case Studies

We provide visualization results to illustrate knowledge representation improvements. Figure 4 shows the sentence similarity matrix of Llama-2-7B and KaLM on Wikitext-103. The diagonal elements denote the similarity of the same sentence, so the values are always 1. From color intensity, it is evident that KaLM learns more discriminative sentence representations, while Llama-2-7B assigns high similarity for arbitrary sentences. The sentences are organized by celebrities and their careers, thus there should also be a high similarity between adjacent sentences. This phenomenon is reflected in the similarity matrix of KaLM in Figure 4(b), manifested in the smaller matrices with darker colors along the diagonal. *More concretely, numerical analysis shows that after training with our method, the sentence-level anisotropy value significantly decreased from 0.83 to 0.21.*

We present KGQA generation cases to demonstrate knowledge inference enhancements. Figure 5 illustrates concrete examples of KGQA generation results on the WN18RR dataset. We showcase the responses generated by Llama-2-7B and KaLM for four tasks involving head entity prediction, relation prediction, tail entity prediction, and triple classification. The prompt templates for each subtask are shown in the second column of Figure 5, where the “inverse relation” is the original relation description with a prefix word “inverse” and the “relation list” consists of all relations concatenated by the symbol “|”. We display the generated answers for triple <salviniaceae, member meronym, salvinia> and triple <refrigerator, hypernym, white goods>. The base LLaMA frequently gives wrong answers and tends to identify keywords from the input prompts for prediction. In contrast, our method can understand the questions and correctly answer various KGQA tasks in most cases.

6 Conclusion

In this work, we show that the subpar performance of LLMs on knowledge-driven tasks stems from a lack of effective knowledge alignment. We present **KaLM**, a novel knowledge-aligned language modeling approach for aligning autoregressive LLMs with KG knowledge. Specifically, we identify two imperative objectives to achieve knowledge alignment: *explicit knowledge alignment* and *implicit knowledge alignment*. We conducted comprehensive experiments and analyses on embedding-based KGC and generation-based KGQA. Experimental results demonstrate that our method achieves effective knowledge alignment and consistently improves performance on knowledge-driven tasks.

Limitations

There are several future directions to improve this work. Firstly, due to the limitation of computational resources, we used the limited-scale LLMs to train and evaluate our method. Evaluations on larger-scale LLMs, such as the 13B and 70B models, can further validate the effectiveness of our approach. Secondly, we use a simple linear combination of explicit alignment loss and implicit alignment loss as the final training objective for KaLM. Further investigations into various forms of loss interactions remain to be explored to maximize the utility of knowledge-aligned language modeling. Finally, we can delve into the performance of the knowledge representations obtained from knowledge-aligned language modeling in cross-domain applications such as retrieval-augmented generation (RAG), to gain broader insights into the generalization capabilities of the proposed approach.

References

Josh Achiam, Steven Adler, Sandhini Agarwal, Lama Ahmad, Ilge Akkaya, Florencia Leoni Aleman, Diogo Almeida, Janko Altenschmidt, Sam Altman, Shyamal Anadkat, et al. 2023. Gpt-4 technical report. *arXiv preprint arXiv:2303.08774*.

Rohan Anil, Andrew M Dai, Orhan Firat, Melvin Johnson, Dmitry Lepikhin, Alexandre Passos, Siamak Shakeri, Emanuel Taropa, Paige Bailey, Zhifeng Chen, et al. 2023. Palm 2 technical report. *arXiv preprint arXiv:2305.10403*.

Antoine Bordes, Nicolas Usunier, Alberto Garcia-Duran, Jason Weston, and Oksana Yakhnenko. 2013. Translating embeddings for modeling multi-relational data. *Advances in neural information processing systems*, 26.

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*, pages 4005–4017.

Ting Chen, Simon Kornblith, Mohammad Norouzi, and Geoffrey Hinton. 2020. A simple framework for contrastive learning of visual representations. In *International conference on machine learning*, pages 1597–1607. PMLR.

Tim Dettmers, Pasquale Minervini, Pontus Stenetorp, and Sebastian Riedel. 2018. Convolutional 2d knowledge graph embeddings. In *Proceedings of the AAAI conference on artificial intelligence*, volume 32.

Kawin Ethayarajh. 2019. How contextual are contextualized word representations? comparing the geometry of bert, elmo, and gpt-2 embeddings. In

Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP), pages 55–65.

Zhangyin Feng, Weitao Ma, Weijiang Yu, Lei Huang, Haotian Wang, Qianglong Chen, Weihua Peng, Xiaocheng Feng, Bing Qin, et al. 2023. Trends in integration of knowledge and large language models: A survey and taxonomy of methods, benchmarks, and applications. *arXiv preprint arXiv:2311.05876*.

Peng Fu, Yiming Zhang, Haobo Wang, Weikang Qiu, and Junbo Zhao. 2023. Revisiting the knowledge injection frameworks. In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pages 10983–10997.

Beliz Gunel, Jingfei Du, Alexis Conneau, and Ves Stoyanov. 2020. Supervised contrastive learning for pre-trained language model fine-tuning. *arXiv preprint arXiv:2011.01403*.

Junxian He, Chunting Zhou, Xuezhe Ma, Taylor Berg-Kirkpatrick, and Graham Neubig. 2021. Towards a unified view of parameter-efficient transfer learning. *arXiv preprint arXiv:2110.04366*.

Dan Hendrycks, Collin Burns, Steven Basart, Andy Zou, Mantas Mazeika, Dawn Song, and Jacob Steinhardt. 2020. Measuring massive multitask language understanding. *arXiv preprint arXiv:2009.03300*.

Edward J Hu, Yelong Shen, Phillip Wallis, Zeyuan Allen-Zhu, Yuanzhi Li, Shean Wang, Lu Wang, and Weizhu Chen. 2021. Lora: Low-rank adaptation of large language models. *arXiv preprint arXiv:2106.09685*.

Jinhao Jiang, Kun Zhou, Zican Dong, Keming Ye, Wayne Xin Zhao, and Ji-Rong Wen. 2023. Structgpt: A general framework for large language model to reason over structured data. *arXiv preprint arXiv:2305.09645*.

Bohan Li, Hao Zhou, Junxian He, Mingxuan Wang, Yiming Yang, and Lei Li. 2020. On the sentence embeddings from pre-trained language models. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 9119–9130.

Junyi Li, Tianyi Tang, Wayne Xin Zhao, Jian-Yun Nie, and Ji-Rong Wen. 2022. Pretrained language models for text generation: A survey. *arXiv preprint arXiv:2201.05273*.

Song Liu, Haoqi Fan, Shengsheng Qian, Yiru Chen, Wenkui Ding, and Zhongyuan Wang. 2021. Hit: Hierarchical transformer with momentum contrast for video-text retrieval. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 11915–11925.

Xueguang Ma, Liang Wang, Nan Yang, Furu Wei, and Jimmy Lin. 2023. Fine-tuning llama for multi-stage text retrieval. *arXiv preprint arXiv:2310.08319*.

678	Stephen Merity, Caiming Xiong, James Bradbury, and Richard Socher. 2016. Pointer sentinel mixture models. In <i>International Conference on Learning Representations</i> .	
679		
680		
681		
682	Niklas Muennighoff. 2022. Sgpt: Gpt sentence embeddings for semantic search. <i>arXiv preprint arXiv:2202.08904</i> .	
683		
684		
685	Jianhao Shen, Chenguang Wang, Linyuan Gong, and Dawn Song. 2022. Joint language semantic and structure embedding for knowledge graph completion. In <i>Proceedings of the 29th International Conference on Computational Linguistics</i> , pages 1965–1978.	
686		
687		
688		
689		
690	Dan Su, Yan Xu, Genta Indra Winata, Peng Xu, Hyeondev Kim, Zihan Liu, and Pascale Fung. 2019. Generalizing question answering system with pre-trained language model fine-tuning. In <i>Proceedings of the 2nd Workshop on Machine Reading for Question Answering</i> , pages 203–211.	
691		
692		
693		
694		
695		
696	Jianlin Su, Jiarun Cao, Weijie Liu, and Yangyiwen Ou. 2021. Whitening sentence representations for better semantics and faster retrieval. <i>arXiv preprint arXiv:2103.15316</i> .	
697		
698		
699		
700	Yixuan Su, Tian Lan, Yan Wang, Dani Yogatama, Lingpeng Kong, and Nigel Collier. 2022. A contrastive framework for neural text generation. <i>Advances in Neural Information Processing Systems</i> , 35:21548–21561.	
701		
702		
703		
704		
705	Jiashuo Sun, Chengjin Xu, Lumingyuan Tang, Saizhuo Wang, Chen Lin, Yeyun Gong, Heung-Yeung Shum, and Jian Guo. 2023. Think-on-graph: Deep and responsible reasoning of large language model with knowledge graph. <i>arXiv preprint arXiv:2307.07697</i> .	
706		
707		
708		
709		
710	Zhiqing Sun, Zhi-Hong Deng, Jian-Yun Nie, and Jian Tang. 2018. Rotate: Knowledge graph embedding by relational rotation in complex space. In <i>International Conference on Learning Representations</i> .	
711		
712		
713		
714	Zhiqing Sun, Zhi-Hong Deng, Jian-Yun Nie, and Jian Tang. 2019. Rotate: Knowledge graph embedding by relational rotation in complex space. <i>arXiv preprint arXiv:1902.10197</i> .	
715		
716		
717		
718	Rohan Taori, Ishaan Gulrajani, Tianyi Zhang, Yann Dubois, Xuechen Li, Carlos Guestrin, Percy Liang, and Tatsunori B. Hashimoto. 2023. Stanford alpaca: An instruction-following llama model. https://github.com/tatsu-lab/stanford_alpaca .	
719		
720		
721		
722		
723	Kristina Toutanova and Danqi Chen. 2015. Observed versus latent features for knowledge base and text inference. In <i>Proceedings of the 3rd workshop on continuous vector space models and their compositionality</i> , pages 57–66.	
724		
725		
726		
727		
728	Hugo Touvron, Louis Martin, Kevin Stone, Peter Albert, Amjad Almahairi, Yasmine Babaei, Nikolay Bashlykov, Soumya Batra, Prajwal Bhargava, Shruti Bhosale, et al. 2023. Llama 2: Open foundation and fine-tuned chat models. <i>arXiv preprint arXiv:2307.09288</i> .	
729		
730		
731		
732		
733		
	Bo Wang, Tao Shen, Guodong Long, Tianyi Zhou, Ying Wang, and Yi Chang. 2021a. Structure-augmented text representation learning for efficient knowledge graph completion. In <i>Proceedings of the Web Conference 2021</i> , pages 1737–1748.	734
		735
		736
		737
		738
	Feng Wang and Huaping Liu. 2021. Understanding the behaviour of contrastive loss. In <i>Proceedings of the IEEE/CVF conference on computer vision and pattern recognition</i> , pages 2495–2504.	739
		740
		741
		742
	Liang Wang, Wei Zhao, Zhuoyu Wei, and Jingming Liu. 2022a. Simkgc: Simple contrastive knowledge graph completion with pre-trained language models. In <i>Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)</i> , pages 4281–4294.	743
		744
		745
		746
		747
		748
	Tongzhou Wang and Phillip Isola. 2020. Understanding contrastive representation learning through alignment and uniformity on the hypersphere. In <i>International Conference on Machine Learning</i> , pages 9929–9939. PMLR.	749
		750
		751
		752
		753
	Xiaozhi Wang, Tianyu Gao, Zhaocheng Zhu, Zhengyan Zhang, Zhiyuan Liu, Juanzi Li, and Jian Tang. 2021b. Kepler: A unified model for knowledge embedding and pre-trained language representation. <i>Transactions of the Association for Computational Linguistics</i> , 9:176–194.	754
		755
		756
		757
		758
		759
	Xintao Wang, Qianyu He, Jiaqing Liang, and Yanghua Xiao. 2022b. Language models as knowledge embeddings. <i>arXiv preprint arXiv:2206.12617</i> .	760
		761
		762
	Bishan Yang, Scott Wen-tau Yih, Xiaodong He, Jianfeng Gao, and Li Deng. 2015. Embedding entities and relations for learning and inference in knowledge bases. In <i>Proceedings of the International Conference on Learning Representations (ICLR) 2015</i> .	763
		764
		765
		766
		767
	Liang Yao, Chengsheng Mao, and Yuan Luo. 2019. Kgbert: Bert for knowledge graph completion. <i>arXiv preprint arXiv:1909.03193</i> .	768
		769
		770
	Liang Yao, Jiazhen Peng, Chengsheng Mao, and Yuan Luo. 2023. Exploring large language models for knowledge graph completion. <i>arXiv preprint arXiv:2308.13916</i> .	771
		772
		773
		774
	Michihiro Yasunaga, Antoine Bosselut, Hongyu Ren, Xikun Zhang, Christopher D Manning, Percy S Liang, and Jure Leskovec. 2022. Deep bidirectional language-knowledge graph pretraining. <i>Advances in Neural Information Processing Systems</i> , 35:37309–37323.	775
		776
		777
		778
		779
		780
	Yichi Zhang, Zhuo Chen, Wen Zhang, and Huajun Chen. 2023. Making large language models perform better in knowledge graph completion. <i>arXiv preprint arXiv:2310.06671</i> .	781
		782
		783
		784

A More Detailed Review of Related Work

This work focuses on fine-tuning autoregressive LLMs to align with KG knowledge. Our work intersects with the following research areas: Knowledge Enhancement for LLMs, Knowledge Graph Completion, Contrastive Representation Learning, and Representation Anisotropy of Language Models.

A.1 Knowledge Enhancement for LLMs

Knowledge enhancement aims to incorporate factual and domain-specific knowledge into LLMs to address their knowledge deficiencies. This can be divided into retrieval-based knowledge augmentation and training-based knowledge integration. *Retrieval-based knowledge augmentation* methods leverage external retrieval modules to provide additional knowledge, aiming to improve the knowledge reasoning capability of LLMs (Sun et al., 2023; Jiang et al., 2023). However, this approach may lead to knowledge conflicts (Feng et al., 2023), where the knowledge in LLMs and the knowledge in the retrieved documents are inconsistent or the retrieved multiple documents are contradictory. *Training-based knowledge integration* methods involve using the textual descriptions of KG triples to pre-train or fine-tune LLMs, aiming to achieve knowledge alignment. These methods can be categorized into explicit alignment (Wang et al., 2021b; Yasunaga et al., 2022) and implicit alignment (Yao et al., 2023; Zhang et al., 2023) based on whether they directly optimize the knowledge representation. Nevertheless, these methods have either sacrificed the generative capability or lacked effective representation alignment. Our approach enhances the knowledge of LLMs via a unique joint objective of explicit alignment and implicit alignment, improving the quality of knowledge representations and generative knowledge reasoning capabilities.

A.2 Knowledge Graph Completion

Knowledge graph completion (KGC) refers to inferring missing triples from an incomplete KG, which can be used to evaluate the knowledge reasoning ability and knowledge representation quality of LLMs. Existing KGC methods can be categorized into structure-based and description-based. *Structure-based methods* represent entities and relations as fixed-dimensional vector embeddings and use scoring functions to assess the plausibility of triples (Bordes et al., 2013; Sun et al., 2019). *Description-based methods* further incorporate the

textual descriptions of KG triples and leverage pre-trained language models to learn knowledge representations of entities and relations (Yao et al., 2019; Shen et al., 2022; Wang et al., 2022b). However, structure-based methods fail to generalize to unseen entities and relations, while description-based methods lack interpretability and exhibit lower efficiency when dealing with extremely large KGs.

A.3 Contrastive Representation Learning

Contrastive learning has demonstrated remarkable success in learning representations across various domains (Chen et al., 2020; Liu et al., 2021; Gunel et al., 2020). The goal is to learn representations that capture shared information between positive pairs while remaining invariant to perturbing noise. The commonly used contrastive learning objectives share a standardized design involving a softmax function over cosine similarity of paired features, with a temperature parameter to control the penalty strength on hard negative samples. Wang and Isola (2020) propose understanding contrastive learning through the lens of alignment and uniformity on the hypersphere. Wang and Liu (2021) show that temperature in the contrastive loss controls the strength of penalties over negative samples.

A.4 Representation Anisotropy of Language Models

PLMs have long been plagued by *representation anisotropy* (Ethayarajh, 2019), where the learned token and sentence representations are confined to a narrow cone within the entire representation space. The issue of representation anisotropy not only results in model degradation (Su et al., 2022) but also leads to poor performance on discriminative tasks (Muennighoff, 2022). Previous work on alleviating representation anisotropy has mainly focused on post-processing techniques such as normalizing flows (Li et al., 2020) or whitening operations (Su et al., 2021) to obtain isotropic representations. Su et al. (2022) propose a contrastive training objective to encourage learning isotropic token representations. However, these methods mainly improve the isotropy of token representations without enhancing the discriminability of sentence representations. Our method improves the token-level and sentence-level representation anisotropy of LLMs through dual-view knowledge graph contrastive learning, and it has rigorous theoretical guarantees.

B Proofs for Theoretical Analysis

In this section, we present proofs for theorems in Sections 4.1 and 4.2 of the main paper.

B.1 Proof of Theorem 1 in Section 4.1

Recall the reformulated dual-view knowledge graph contrastive learning objective (Equation 5):

$$\mathcal{L}_{\text{exp}}(f; \tau, \mathcal{N}) \triangleq \mathbb{E}_{\substack{(\mathcal{D}_{hr}, \mathcal{D}_t) \sim p_{\text{pos}} \\ \{\mathcal{D}_{t_i}'\}_{i=1}^{\mathcal{N}} \stackrel{i.i.d.}{\sim} p_{\text{data}}}} \left[-\log \frac{e^{f(\mathcal{D}_{hr})^\top f(\mathcal{D}_t)/\tau}}{e^{f(\mathcal{D}_{hr})^\top f(\mathcal{D}_t)/\tau} + \sum_{i=1}^{\mathcal{N}} e^{f(\mathcal{D}_{hr})^\top f(\mathcal{D}_{t_i}')/\tau}} \right].$$

From the symmetry of p , we can derive:

$$\mathcal{L}_{\text{exp}}(f; \tau, \mathcal{N}) = \mathbb{E}_{(\mathcal{D}_{hr}, \mathcal{D}_t) \sim p_{\text{pos}}} \left[-f(\mathcal{D}_{hr})^\top f(\mathcal{D}_t)/\tau \right] + \mathbb{E}_{\substack{(\mathcal{D}_{hr}, \mathcal{D}_t) \sim p_{\text{pos}} \\ \{\mathcal{D}_{t_i}'\}_{i=1}^{\mathcal{N}} \stackrel{i.i.d.}{\sim} p_{\text{data}}}} \left[\log \left(e^{f(\mathcal{D}_{hr})^\top f(\mathcal{D}_t)/\tau} + \sum_{i=1}^{\mathcal{N}} e^{f(\mathcal{D}_{t_i}')^\top f(\mathcal{D}_t)/\tau} \right) \right].$$

Note that we can have the following limits almost surely by the strong law of large numbers (SLLN):

$$\lim_{\mathcal{N} \rightarrow \infty} \log \left(\frac{e^{f(\mathcal{D}_{hr})^\top f(\mathcal{D}_t)/\tau}}{\mathcal{N}} + \frac{\sum_{i=1}^{\mathcal{N}} e^{f(\mathcal{D}_{t_i}')^\top f(\mathcal{D}_t)/\tau}}{\mathcal{N}} \right) = \log \mathbb{E}_{\mathcal{D}_i^- \sim p_{\text{data}}} \left[e^{f(\mathcal{D}_i^-)^\top f(\mathcal{D}_t)/\tau} \right].$$

Then we can derive the following limits:

$$\begin{aligned} & \lim_{\mathcal{N} \rightarrow \infty} \mathcal{L}_{\text{exp}}(f; \tau, \mathcal{N}) - \log \mathcal{N} \\ &= \mathbb{E}_{(\mathcal{D}_{hr}, \mathcal{D}_t) \sim p_{\text{pos}}} \left[-f(\mathcal{D}_{hr})^\top f(\mathcal{D}_t)/\tau \right] \\ &+ \lim_{\mathcal{N} \rightarrow \infty} \mathbb{E}_{\substack{(\mathcal{D}_{hr}, \mathcal{D}_t) \sim p_{\text{pos}} \\ \{\mathcal{D}_{t_i}'\}_{i=1}^{\mathcal{N}} \stackrel{i.i.d.}{\sim} p_{\text{data}}}} \left[\log \left(\frac{e^{f(\mathcal{D}_{hr})^\top f(\mathcal{D}_t)/\tau}}{\mathcal{N}} + \frac{\sum_{i=1}^{\mathcal{N}} e^{f(\mathcal{D}_{t_i}')^\top f(\mathcal{D}_t)/\tau}}{\mathcal{N}} \right) \right] \\ &= \mathbb{E}_{(\mathcal{D}_{hr}, \mathcal{D}_t) \sim p_{\text{pos}}} \left[-f(\mathcal{D}_{hr})^\top f(\mathcal{D}_t)/\tau \right] \end{aligned}$$

$$\begin{aligned} &+ \mathbb{E} \left[\lim_{\mathcal{N} \rightarrow \infty} \log \left(\frac{e^{f(\mathcal{D}_{hr})^\top f(\mathcal{D}_t)/\tau}}{\mathcal{N}} + \frac{\sum_{i=1}^{\mathcal{N}} e^{f(\mathcal{D}_{t_i}')^\top f(\mathcal{D}_t)/\tau}}{\mathcal{N}} \right) \right] \\ &= -\frac{1}{\tau} \mathbb{E}_{(\mathcal{D}_{hr}, \mathcal{D}_t) \sim p_{\text{pos}}} \left[f(\mathcal{D}_{hr})^\top f(\mathcal{D}_t) \right] \\ &+ \mathbb{E}_{\mathcal{D}_i \sim p_{\text{data}}} \left[\log \mathbb{E}_{\mathcal{D}_i^- \sim p_{\text{data}}} \left[e^{f(\mathcal{D}_i^-)^\top f(\mathcal{D}_i)/\tau} \right] \right]. \end{aligned}$$

We now finish the proof of Theorem 1.

$$\begin{aligned} & \lim_{\mathcal{N} \rightarrow \infty} \mathcal{L}_{\text{exp}}(f; \tau, \mathcal{N}) - \log \mathcal{N} = \\ & -\frac{1}{\tau} \mathbb{E}_{(\mathcal{D}_{hr}, \mathcal{D}_t) \sim p_{\text{pos}}} \left[f(\mathcal{D}_{hr})^\top f(\mathcal{D}_t) \right] \\ &+ \mathbb{E}_{\mathcal{D}_i \sim p_{\text{data}}} \left[\log \mathbb{E}_{\mathcal{D}_i^- \sim p_{\text{data}}} \left[e^{f(\mathcal{D}_i^-)^\top f(\mathcal{D}_i)/\tau} \right] \right]. \end{aligned}$$

B.2 Proof of Theorem 2 in Section 4.2

Recall the asymptotics of the explicit knowledge alignment objective when the number of negative samples approaches infinity (Equation 6):

$$\begin{aligned} & \lim_{\mathcal{N} \rightarrow \infty} \mathcal{L}_{\text{exp}}(f; \tau, \mathcal{N}) - \log \mathcal{N} = \\ & -\frac{1}{\tau} \mathbb{E}_{(\mathcal{D}_{hr}, \mathcal{D}_t) \sim p_{\text{pos}}} \left[f(\mathcal{D}_{hr})^\top f(\mathcal{D}_t) \right] \\ &+ \mathbb{E}_{\mathcal{D}_i \sim p_{\text{data}}} \left[\log \mathbb{E}_{\mathcal{D}_i^- \sim p_{\text{data}}} \left[e^{f(\mathcal{D}_i^-)^\top f(\mathcal{D}_i)/\tau} \right] \right]. \end{aligned}$$

Recall the definition of sentence-level anisotropy value of corpus $\{\mathcal{D}_i\}_{i=1}^{\mathcal{N}}$ (Equation 7):

$$\text{anisotropy}_{\{\mathcal{D}\}} = \frac{1}{\mathcal{N}(\mathcal{N}-1)} \sum_{i=1}^{\mathcal{N}} \sum_{j=1, j \neq i}^{\mathcal{N}} e_i^\top e_j.$$

We can further derive the inequality below from the second term of Equation 6 with Jensen's inequality when p_{data} is uniform over finite samples $\{\mathcal{D}_i\}_{i=1}^{\mathcal{N}}$:

$$\begin{aligned}
& \mathbb{E}_{\mathcal{D}_i \sim p_{data}} \left[\log \mathbb{E}_{\mathcal{D}_i^- \sim p_{data}} \left[e^{f(\mathcal{D}_i^-)^\top f(\mathcal{D}_i)/\tau} \right] \right] \\
&= \frac{1}{N} \sum_{i=1}^N \log \left(\frac{1}{N} \sum_{j=1}^N e^{e_i^\top e_j/\tau} \right) \\
&\geq \frac{1}{\tau N^2} \sum_{i=1}^N \sum_{j=1}^N e_i^\top e_j \\
&= \frac{1}{\tau N^2} \left(\sum_{i=1}^N \sum_{j=1, j \neq i}^N e_i^\top e_j + N \right) \\
&= \frac{N-1}{\tau N} \cdot \frac{1}{N(N-1)} \sum_{i=1}^N \sum_{j=1, j \neq i}^N e_i^\top e_j + \frac{1}{\tau N} \\
&= \frac{N-1}{\tau N} \cdot \text{anisotropy}_{\{\mathcal{D}\}} + \frac{1}{\tau N}.
\end{aligned}$$

We now finish the *proof of Theorem 2*.

$$\begin{aligned}
& \mathbb{E}_{\mathcal{D}_i \sim p_{data}} \left[\log \mathbb{E}_{\mathcal{D}_i^- \sim p_{data}} \left[e^{f(\mathcal{D}_i^-)^\top f(\mathcal{D}_i)/\tau} \right] \right] \\
&\geq \frac{N-1}{\tau N} \cdot \text{anisotropy}_{\{\mathcal{D}\}} + \frac{1}{\tau N}.
\end{aligned}$$

C Further Details about Implementation and Experimental Setup

C.1 Dataset Details

WN18RR and FB15k-237 are commonly used KGs derived from WordNet and Freebase, respectively (Bordes et al., 2013). They have been carefully constructed to prevent test set leakage by removing inverse relations. We use these datasets for training and evaluation. The statistics are shown in Table 2.

Table 2: Statistics of the datasets.

Dataset	#Entity	#Relation	#Train	#Valid	#Test
WN18RR	40,943	11	86,835	3,034	3,134
FB15k-237	14,541	237	272,115	17,535	20,466

C.2 KaLM Implementation Details

We initially choose Llama-2-7B as the base LLM and fine-tune it through the training objective in Equation 4. We use varying batch sizes for explicit knowledge alignment and implicit knowledge alignment. For both WN18RR and FB15k-237, we use a batch size of 96 for explicit knowledge alignment and 8 for implicit knowledge alignment. To save computing resources for parameter-efficient

fine-tuning, we use the LoRA (Hu et al., 2021) method to fine-tune the [“gate_proj”, “up_proj”, “down_proj”] modules in the feed-forward network of the Llama-2-7B model. We train the Llama and Mistral series models on the NVIDIA A800 GPU, while other models are trained on the NVIDIA 4090 GPU. The hyper-parameters utilized for training KaLM (based on Llama-2-7B) and the detailed configuration are enumerated in Table 3.

Table 3: Hyper-parameters for training KaLM.

Hyper-parameters	WN18RR	FB15k-237
epochs	20	15
max-description-length	50	50
max-language-modeling-length	256	256
explicit-alignment-batch-size	96	96
implicit-alignment-batch-size	8	8
lora-module	ffn	ffn
lora-alpha	16.0	16.0
lora-drouout	0.05	0.05
lora-rank	8	8
bnb-config	load-in-8bit	load-in-8bit
learning-rate	1e-4	1e-4
LR-scheduler-type	cosine	cosine
weight-decay	0.001	0.001
gradient-checkpointing	True	True
optimizer	AdamW	AdamW
AdamW-beta1	0.9	0.9
AdamW-beta2	0.999	0.999
bf16	True	True
loss weight (λ)	0.1	0.1

We also implemented KaLM based on other LLMs to demonstrate the generalizability of our approach, including Llama-3-8B, Mistral-7B-v0.1, OPT-6.7B, Pythia-6.9B, and Pythia-2.8B. It is important to note that the feed-forward network layers in the Pythia model are named [“dense_h_to_4h”, “dense_4h_to_h”], while in the OPT model they are named [“fc1”, “fc2”]. This differs from the feed-forward network layers in the Llama and Mistral model series. The full parameters used in these experiments are shown in Table 4 (only the differing parameters are listed; the unlisted parameters remain consistent with Table 3).

For the cosine similarity matrix composed of head entity-relation embeddings (row direction) and tail entity embeddings (column direction), we calculate the cross-entropy loss in the row direction (i.e., a head entity-relation embedding matching different tail entity embeddings) and the column direction (i.e., a tail entity embedding matching different head entity-relation embeddings) separately. We then take the average of the two losses to obtain the final InfoNCE loss. Similar to Equation 1, the

Table 4: Additional Hyper-parameters for training *KaLM* with different LLMs.

Models	epochs	explicit-batch-size	implicit-batch-size	bnb-config
Llama-3-8B-WN	20	96	8	load-in-8bit
Llama-3-8B-FB	15	96	8	load-in-8bit
Mistral-7B-v0.1-WN	20	168	12	load-in-4bit
Mistral-7B-v0.1-FB	15	168	12	load-in-4bit
OPT-6.7B-WN	20	24	4	load-in-8bit
OPT-6.7B-FB	15	24	4	load-in-8bit
Pythia-6.9B-WN	20	24	4	load-in-8bit
Pythia-6.9B-FB	15	24	4	load-in-8bit
Pythia-2.8B-WN	20	48	8	load-in-8bit
Pythia-2.8B-FB	15	48	8	load-in-8bit

column-direction loss is defined as follows:

$$\ell_c = -\log \frac{e^{(\phi(e_t, e_{hr}) - \gamma)/\tau}}{e^{(\phi(e_t, e_{hr}) - \gamma)/\tau} + \sum_{j=1}^N e^{\phi(e_t, e_{hr'_j})/\tau}}.$$

C.3 More Details about Evaluations

For the embedding-based KGC task, we report five automated metrics: Mean Rank (MR), Mean Reciprocal Rank (MRR), and Hit@ k ($k \in \{1, 3, 10\}$). MR is the mean rank of all test triplets and MRR denotes the average reciprocal rank of all test triples. Hit@ k measures the proportion of entities correctly ranked in the top k . Following previous work, our method is evaluated under the *filtering setting* (Bordes et al., 2013), where the scores of all true triples in the training, validation, and testing set are ignored. All results are averaged over the tail direction (a <head entity - relation> embedding matching different tail entity embeddings, i.e., tail entity prediction) and head direction (a <tail entity - inverse relation> embedding matching different head entity embeddings, i.e., head entity prediction).

For the generation-based KGQA task, we report the prediction accuracy over head entities, tail entities, relations, and relation classifications. To better prompt LLMs for the knowledge graph question-answering task, we selected several triples from the validation set and constructed few-shot examples using the corresponding templates from Table 5.

D Addition Experimental Results

In this section, we provide more experimental results to show the effectiveness of our method.

D.1 More Experiments on Knowledge Representation Assessment

In Table 5, we present additional knowledge representation results (the embedding-based KGC task) to demonstrate the effectiveness of KaLM in knowledge alignment. The best and second best experimental results are indicated by **bold** and underline texts, respectively. Overall, the proposed method achieved excellent performance on the embedding-based KGC task, delivering impressive results in the MR and Hit@10 metrics, while also being highly competitive in other metrics.

The experimental results based on LLMs of different sources and scales demonstrate the effectiveness and generalizability of our proposed method. Under similar experimental settings, more powerful LLMs (such as Llama3-8B and Mistral-7B) achieved better metrics after being fine-tuned with KaLM, which also demonstrates the scalability of our method. It is worth noting that for LLMs of the same origin but different scales (Pythia-6.9B and Pythia-2.8B), the smaller-scale Pythia-2.8B benefited from a larger training batch size during fine-tuning. As a result, its final experimental metrics matched or even surpassed those of the more powerful Pythia-6.9B model. This also highlights the importance of large batch sizes for the embedding-based KGC task, suggesting that using more powerful computing resources and larger GPU memory could further enhance the effectiveness of the proposed KaLM method.

D.2 More Experiments on Knowledge Inference Evaluation

In Figure 6, we present additional knowledge inference results (generation-based KGQA) to demon-

Table 5: More Embedding-based KGC results with various LLMs on WN18RR and FB15k-237.

Method	WN18RR					FB15k-237				
	MR	MRR	H@1	H@3	H@10	MR	MRR	H@1	H@3	H@10
<i>structure-based methods</i>										
TransE	2300	0.243	0.043	0.441	0.532	323	0.279	0.198	0.376	0.441
DistMult	7000	0.444	0.412	0.470	0.504	512	0.281	0.199	0.301	0.446
RotatE	3340	0.476	0.428	0.492	0.571	177	<u>0.338</u>	0.241	0.375	<u>0.533</u>
<i>description-based methods (autoencoder PLMs)</i>										
KG-BERT	97	0.216	0.041	0.302	0.524	153	-	-	-	0.420
StAR	51	0.401	0.243	0.491	0.709	117	0.296	0.205	0.322	0.482
C-LMKE	79	0.619	0.523	0.671	0.789	141	0.306	0.218	0.331	0.484
SimKGC	-	0.671	0.585	<u>0.731</u>	0.817	-	0.333	<u>0.246</u>	0.362	0.510
<i>description-based methods (autoregressive LLMs)</i>										
Llama-2-7B	15969	0.010	0.004	0.010	0.020	5359	0.006	0.002	0.004	0.012
Llama2-7B _{KaLM}	22	0.596	0.497	0.706	0.856	114	0.305	0.227	0.348	0.511
Llama3-8B _{KaLM}	28	0.613	0.502	0.728	<u>0.865</u>	125	0.327	0.241	0.366	0.522
Mistral-7B _{KaLM}	<u>25</u>	<u>0.634</u>	<u>0.526</u>	0.743	0.872	<u>117</u>	0.341	0.252	<u>0.371</u>	0.536
OPT-6.7B _{KaLM}	24	0.514	0.397	0.603	0.822	126	0.288	0.199	0.312	0.486
Pythia-6.9B _{KaLM}	28	0.508	0.394	0.598	0.818	130	0.289	0.199	0.310	0.484
Pythia-2.8B _{KaLM}	30	0.539	0.398	0.644	0.829	133	0.292	0.205	0.318	0.489

strate the effectiveness of KaLM in knowledge alignment. This section demonstrates the performance of various powerful LLMs (including Llama-2-7B, Llama-3-8B, and Mistral-7B) before and after fine-tuning with KaLM, across various knowledge graph question-answering tasks (including head entity prediction, tail entity prediction, relation prediction, and triple classification).

The experimental results can be divided into three groups by color: the green series, blue series, and red series correspond to the KGQA results of Llama-2-7B, Llama-3-8B, and Mistral-7B before and after training, respectively. It can be observed that after fine-tuning with KaLM, all three LLMs achieved consistent improvements in prediction accuracy for the question-answering tasks.

At the KGQA task level, the most significant overall improvements were observed in tail entity prediction (an average increase of 15.1%) and triple classification (an average increase of 11.4%), followed by relation prediction (an average increase of 10.7%) and head entity prediction (an average increase of 7.4%). At the LLM level, the most exciting improvements were seen in Llama-3-8B (an average increase of 12.38%) and Llama-2-7B (an average increase of 10.65%), while Mistral-7B showed slightly fewer gains (an average increase of 10.43%). This suggests that our method demonstrates generalizability with powerful LLMs.

D.3 More Visualizations on Knowledge Representation Matrix

From this section onward, unless stated otherwise, KaLM refers to the model checkpoint trained on Llama-2-7B using our method. We present more knowledge representation results to demonstrate the effectiveness of KaLM in knowledge alignment. Figure 7 displays the sentence similarity matrix of several similar entity descriptions from the WN8RR dataset. Detailed information about entity names and descriptions can be found in Figure 8. It is evident that KaLM can obtain more distinguishable knowledge representations, where the similarity between related entities (diagonal elements) is high, while the similarity between unrelated entities (off-diagonal elements) is low.

D.4 Detailed analysis of Representation Anisotropy

We further analyze the sentence-level representation anisotropy on the Wikitext-103 test set using model checkpoints trained on the WN18RR dataset. The sentence-level anisotropy value for a given corpus $\{\mathcal{D}_i\}_{i=1}^N$ is defined in Equation 7, where a lower anisotropy value indicates better discriminative characteristics of sentence representations.

Figure 9 plots the anisotropy value over different layers for LLaMA and KaLM. We can observe that the anisotropy value of LLaMA consistently

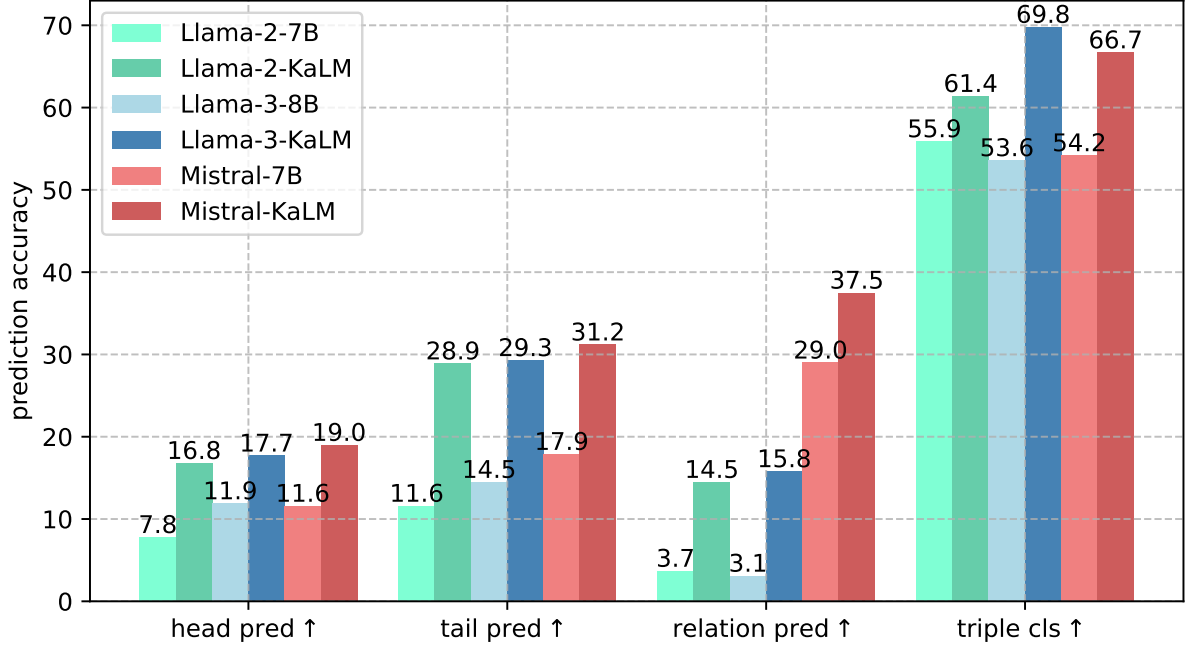


Figure 6: Comparison of generative knowledge inference performance between Base LLMs and their fine-tuned KaLM versions, best viewed in three color groups. The symbol \uparrow means higher is better and \downarrow means lower is better.

remains at a relatively high level, suggesting that the base LLM suffers from severe representation anisotropy issues. In contrast, our proposed KaLM notably mitigates this issue, with the anisotropy values decreasing gradually as the depth of the model increases, and dropping significantly from 0.5 to 0.2 at the output layer. The anisotropy values of the last layer for LLaMA and KaLM show that after training with our method, the sentence-level anisotropy value significantly decreased from 0.83 to 0.21. The results indicate that our method can effectively reduce the anisotropy of representations across layers in LLMs, resulting in a significant improvement in knowledge representation.

Figure 10 analyzes the changes in anisotropy values during the model training process. The results show that the anisotropy values decrease rapidly after a few epochs of training and eventually stabilize at a low level. We assume that the initial epochs of training have completed the preliminary alignment of knowledge representation, while the subsequent training epochs mainly focus on integrating explicit and implicit representations.

E Ablation Studies

In this section, we present concrete ablation studies to analyze the effectiveness of each component of our approach. We ablate the settings that led to the final design, including training objectives,

fine-tuning modules, and training epochs. It is important to note that the results of the ablation experiments in this section were obtained from earlier runs on an NVIDIA 3090 GPU, which may lead to slight differences compared to the full KGC results presented in the main text.

E.1 The necessity of the implicit knowledge alignment objective (Equation 3)

In Table 6, we train the model using different loss weights (i.e., the λ parameter in Equation 4) and analyze its performance on the KGC task. Note that this experiment is conducted solely for ablation analysis, thus only 10 training epochs are used. Experimental results reveal that incorporating the implicit knowledge alignment objective (i.e., $\lambda > 0$) generally leads to better performance in KGC, indicating further improvement in knowledge representation. The best performance in KGC is achieved when $\lambda = 0.1$. The results confirm that both explicit alignment and implicit alignment are crucial for knowledge alignment, as they both essentially require a deep understanding of knowledge. Intuitively, the initial model already exhibits strong language generation capabilities. By setting the hyperparameter λ to 0.1, we encourage the large language model to focus more on learning intrinsic knowledge representations, thereby promoting knowledge alignment within the model.

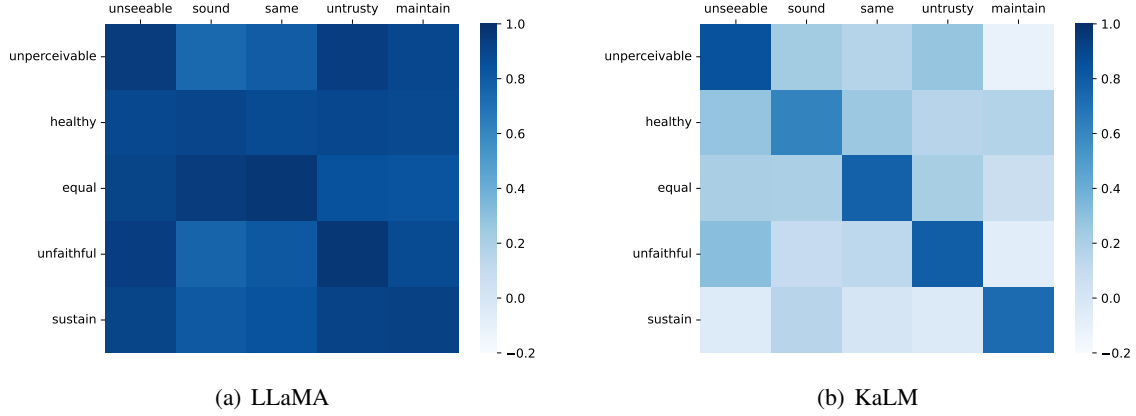


Figure 7: Similarity matrix of selected similar entity descriptions from the WN8RR dataset.

Entity Name	Entity Description
unseeable	unseeable, impossible or nearly impossible to see; imperceptible by the eye; "the invisible man"; "invisible rays"; "an invisible hinge"; "invisible mending"
unperceivable	unperceivable, impossible or difficult to perceive by the mind or senses; "an imperceptible drop in temperature"; "an imperceptible nod"; "color is unperceivable to the touch"
sound	sound, financially secure and safe; "sound investments"; "a sound economy"
healthy	healthy, having or indicating good health in body or mind; free from infirmity or disease; "a rosy healthy baby"; "staying fit and healthy"
same	same, closely similar or comparable in kind or quality or quantity or degree; "curtains the same color as the walls"; "mother and son have the same blue eyes"
equal	equal, having the same quantity, value, or measure as another; "on equal terms"; "all men are equal before the law"
untrusty	untrusty, not worthy of trust or belief; "an untrustworthy person"
unfaithful	unfaithful, not true to duty or obligation or promises; "an unfaithful lover"
maintain	maintain, keep in a certain state, position, or activity; e.g., "keep clean"; "hold in place"; "She always held herself as a lady"; "The students keep me on my toes"
sustain	sustain, lengthen or extend in duration or space; "We sustained the diplomatic negotiations as long as possible"; "prolong the treatment of the patient"; "keep up the good work"

Figure 8: Selected entities and their corresponding textual descriptions.

The implicit knowledge alignment objective focuses on incorporating textual patterns of knowledge into the LLM to prevent catastrophic forgetting of previous knowledge and maintain its generative capability. We also conducted additional perplexity (PPL) evaluation experiments to illustrate the impact of the implicit knowledge alignment loss. The additional results show that for the corresponding $\lambda = 0, 0.01, 0.1, 1.0$ in Table 6, the model’s PPL are **6.42**, **4.96**, **4.97**, and **4.98**, respectively. Therefore, we can conclude that incorporating the implicit alignment loss maintains the model’s language modeling capability, whereas not using the implicit alignment loss significantly

impairs the model’s generative ability.

Table 6: KGC results with different λ in Equation 4.

Method	WN18RR					PPL
	MR	MRR	H@1	H@3	H@10	
KaLM ($\lambda = 0$)	21.2	0.512	0.355	0.611	0.815	6.42
KaLM ($\lambda = 0.01$)	19.8	0.510	0.352	0.604	0.818	4.96
KaLM ($\lambda = 0.1$)	20.1	0.517	0.359	0.615	0.825	4.98
KaLM ($\lambda = 1.0$)	21.6	0.500	0.336	0.596	0.806	4.98

E.2 The effects of fine-tuning different LLM modules using LoRA

In Table 7, we fine-tune different modules of the model using the LoRA (Hu et al., 2021) method and analyze their performance on KGC tasks and PPL




Figure 9: layer-wise analysis of anisotropy. The vertical axis represents the sentence-level representation anisotropy value on the Wikitext-103 test set, while the horizontal axis denotes the number of model layers.




Figure 10: epoch-wise analysis of anisotropy. The vertical axis represents the sentence-level representation anisotropy value on the Wikitext-103 test set, while the horizontal axis denotes the number of training epochs.

evaluations. Note that this experiment is conducted solely for ablation analysis, hence only 10 epochs of training were performed. “*att*” refers to fine-tuning only the attention module, “*ffn*” refers to fine-tuning only the feed-forward network, and “*att-ffn*” refers to fine-tuning both the attention module and the feed-forward network simultaneously. The results show that fine-tuning with the “*att-ffn*” approach achieves the best KGC performance, but it also leads to higher PPL values, suggesting that the model’s generation capability may be significantly compromised. Therefore, as a compromise, we choose the “*ffn*” fine-tuning approach, maintaining moderate knowledge representation performance while preserving the original generation capability.

Table 7: KGC results and PPL evaluation results when fine-tuning different network modules with LoRA.

Method	WN18RR					PPL
	MR	MRR	H@1	H@3	H@10	
KaLM (att)	21.9	0.47.5	0.331	0.580	0.784	5.03
KaLM (ffn)	20.1	0.517	0.359	0.615	0.825	4.96
KaLM (att-ffn)	19.5	0.525	0.371	0.619	0.831	5.07

These experimental results are consistent with the conclusions of (He et al., 2021), where the FFN learns local features and patterns within the input sequence, allowing it to directly capture task-specific text patterns. Meanwhile, attention provides the model with the ability to capture complex contextual relationships, which is key to LLMs’ understanding and generation of natural language. Under the knowledge-aligned language modeling objective, we aim to align the internal knowledge representations of LLMs while preserving their

inherent natural language generation capabilities. Therefore, directly fine-tuning the FFN layers can reduce resource consumption and maximize the effectiveness of KaLM fine-tuning.

E.3 The sustained gains and potential impacts of training for more epochs

In Table 8, we fine-tune the model using different numbers of training epochs and analyze their performance on KGC tasks. This experiment is mainly conducted to investigate whether additional training epochs can lead to further improvement in knowledge representations. The experimental results show that using more training epochs can continuously improve the performance of KaLM on the KGC task, resulting in higher MRR and Hit@k metrics. The model trained with our method consistently maintains an acceptable PPL value due to the implicit knowledge alignment objective. However, this also comes with more computational resource consumption and training time. As a result, we selected a moderate number of training epochs.

Table 8: KGC results with different training epochs.

Method	WN18RR					PPL
	MR	MRR	H@1	H@3	H@10	
KaLM (epoch=10)	20.1	0.517	0.359	0.615	0.825	4.96
KaLM (epoch=20)	19.6	0.554	0.402	0.650	0.848	4.98
KaLM (epoch=30)	21.9	0.576	0.427	0.673	0.854	5.00