

Enhancing Transformers for Generalizable First-Order Logical Entailment

Tianshi Zheng^{*1}, Jiazheng Wang^{*2}, Zihao Wang¹, Jiaxin Bai¹, Hang Yin³, Zheyue Deng¹,
Yangqiu Song¹, Jianxin Li²

¹Department of Computer Science and Engineering, HKUST, Hong Kong SAR, China

²Department of Computer Science and Engineering, Beihang University, Beijing, China

³Department of Mathematical Sciences, Tsinghua University, Beijing, China
tzhengad@connect.ust.hk

Abstract

Transformers, as the fundamental deep learning architecture, have demonstrated great capability in reasoning. This paper studies the generalizable first-order logical reasoning ability of transformers with their *parameterized* knowledge and how to improve it. Transformers' capability of first-order reasoning is further captured by whether they can conduct first-order logical entailment, which is quantitatively measured by their performance in answering knowledge graph queries. We establish the connections between (1) two types of distribution shifts studied in out-of-distribution generalization and (2) unseen knowledge and query settings discussed in the task of knowledge graph query answering, which makes it possible to characterize the fine-grained generalizability. Results on our comprehensive dataset showed that transformers *outperform* previous methods designed particularly for this task and provided detailed empirical evidence about the impact of the input query syntax, token embedding, and transformer architectures on the reasoning capability of transformers. Interestingly, our results revealed the mismatch of positional encoding and other design choices of transformer architectures in previous practices. Motivated by this, we propose **TEGA**, a logic-aware architecture that significantly improves the performance in generalizable first-order logical entailment.

1 Introduction

As a fundamental architecture in deep learning, transformers possess strong reasoning capabilities on various tasks, including arithmetic reasoning (Saxton et al., 2019; Hendrycks et al., 2021), symbolic reasoning for first-order logic rules (Dehghani et al., 2019; Lample and Charton, 2019), set-theoretic operations (Barrett et al., 2018), and theorem proving (Polu and Sutskever, 2020). Besides,

transformers have also demonstrated proficiency in logical reasoning over natural language (Han et al., 2022; Tian et al., 2021). To distinguish whether transformers conduct the reasoning rather than fitting the data distribution, recent studies in natural language reasoning further measure the capabilities of transformers for out-of-demonstration samples (Saparov et al., 2023). However, their discussions are limited in two aspects: (1) they only concern the reasoning ability of transformers with *in-context* knowledge, and (2) they fail to elicit the connection of out-of-demonstration samples with two distribution shifts (Moreno-Torres et al., 2012) for the study of out-of-distribution generalization.

In this paper, we take a step further to understand the generalizable reasoning capability of transformers. What sets us apart from (Saparov et al., 2023) is that (1) our investigation covers the transformer reasoning with *parameterized knowledge*, which suits many scenarios when the related knowledge is not explicit for users and only questions are given, and (2) we further realize two distribution shifts (Moreno-Torres et al., 2012) in our first order reasoning tasks, which boils down to the process of first-order logical entailment, where we verify whether a first-order sentence (conclusion) is the logical consequence of a set of known first-order sentences (premises) or not (Marker, 2002).

Specifically, we study the first-order logical entailment with Knowledge Graphs (KG), leading to the widely discussed task of knowledge graph query answering. In our setting, knowledge in KGs is parameterized in models and forms implicit premises. (§2) The process of logical entailment occurs when identifying the answer to a logical query. Therefore, the two distribution shifts – concept shift and covariant shift – are naturally credited to the unobserved knowledge and the unseen query types in our experimental settings. (§3) For better evaluation, we build our own benchmark with fifty-five types of logical queries

* Equal Contribution

over three knowledge graphs (FB15k (Bordes et al., 2013), FB15k-237 (Toutanova and Chen, 2015), and NELL995 (Carlson et al., 2010)), including all existing features discussed in recent literature (Yin et al., 2023b). Our benchmark results suggest that transformers can handle first-order logical entailment even compared to the methods particularly designed for this task. (§4)

We conducted extensive experiments to characterize the impact of three main aspects of solving logical entailment on reasoning capability: the query syntax, the learning of token embedding, and the transformer architecture. Our results provide strong empirical clues on improving transformers’ reasoning capability by choosing the proper formal language syntax, positional encodings, semantics in pre-trained KG embeddings, and inductive biases in transformer architectures. (§5)

Interestingly, our results demonstrated the superior performance and generalizability of relative positional encoding (RPE) over traditional absolute positional encoding (APE). However, previous studies have only proposed inductive biases within the APE setting, which we prove ineffective in the RPE setting. To fill this gap, we propose **TEGA** (Transformer Encoder with Guided Attention), a novel modeling methodology that facilitates effective reasoning with logic-aware guidance in self-attention. Our study shows that TEGA substantially improves the performance and generalizability upon transformers under the RPE setting. (§6)

Overall, this work makes three key technical contributions:

1. **Benchmark.** We conducted **comprehensive benchmarking** of transformer architectures against baseline models for knowledge graph query answering, using a new dataset that incorporates *two distribution shifts*.
2. **Exploration.** We systematically investigated **critical design choices** throughout the modeling process, including *query syntax and permutation, token embedding, and transformer architecture designs*.
3. **Method.** We introduced **TEGA**, a novel architecture that improves logical entailment performance and generalizability through two inductive biases. We validated its effectiveness through extensive experiments and ablations.

Our code and data can be found at <https://github.com/HKUST-KnowComp/TEGA>.

2 Preliminaries

This section briefly introduces First-Order (FO) logical entailment and its restricted application in knowledge graph query answering. By revealing the connection, we show that the key ability to address complex queries is the first-order logical entailment. For simplicity, our introduction **only** restricts to a finite entity set \mathcal{E} and a finite binary relation set \mathcal{R} . No functions are mentioned. Detailed presentation of first-order logic can be found in model theory literature (Marker, 2002).

2.1 First-order Logical Entailment

Definition 1 (First-order sentence). *The set of first-order formulae is the minimal set F such that*

1. $p(s, o) \in F$, where p is the relation, s, o can be either an entity in \mathcal{E} or a variable.
2. If $s \in F$, then $\neg s \in F$; If $s, t \in F$, then $s \wedge t \in F$ and $s \vee t \in F$.
3. If $s(x) \in F$ and x is a variable that is not quantified, then $\forall x.s(x) \in F$ and $\exists x.s(x) \in F$.

We say a variable x in a formula is quantified if there is a quantifier ($\exists x$ or $\forall x$). Otherwise, we say a variable is free. A sentence is a formula without free variables.

Next, we introduce first-order logical entailment, the *central* concept of first-order reasoning. In general, the entailment is a process of *verifying* the *conclusion* (one sentence to verify) given the *knowledge* (a set of given premises). **Notably, knowledge, conclusion and verification are the three key components of the definition.** For first-order logical entailment, the *knowledge* and *conclusion* are restricted as first-order sentences, and the *verification* process is subject to the first-order logical calculus. A set of FO sentences $\{p_i, i = 1, \dots, n\}$ entails a FO conclusion s is denoted as:

$$\underbrace{\{p_i, i = 1, \dots, n\}}_{\text{knowledge or premises}} \models \underbrace{s}_{\text{conclusion}} \quad (1)$$

After verification, its truth value suggests whether s is the logical consequence of the premises.

2.2 Knowledge Graph Query Answering

Knowledge graph query answering is an important application of FO logical entailment. Its importance comes from both the importance of knowledge graphs and the queries in database systems. Here, we connect the task of KG query answering and FO entailment from the aspects of *knowledge, conclusion, and verification*.

Knowledge. Given an entity set \mathcal{E} and relation set \mathcal{R} , a knowledge graph is defined as a set of triplets $\mathcal{G} = \{(h_i, r_i, t_i)\}$, where $h_i, t_i \in \mathcal{E}$ are entities and $r_i \in \mathcal{R}$ are relations. Each triple (h, r, t) in KG can be regarded as the simplest form of first-order sentence $r(h, t)$. Thus, the knowledge graph \mathcal{G} provided the premises.

Answers to the query as multiple conclusions to verify. Existing research for KG query answering usually focuses on the restricted families of FO queries, specifically, the Existential First-Order (EFO) queries. An EFO query is represented in Disjunctive Normal Form (DNF) as a disjunction of one or multiple conjunctions:

$$q[V?] = V? \cdot \exists V_1, V_2, \dots, V_n : c_1 \vee c_2 \vee \dots \vee c_m \quad (2)$$

In equation 2, $V?$ is the free variable, V_1, V_2, \dots, V_n are existential variables, while c_1, c_2, \dots, c_m are conjunctions of one-hop atomic formulas or their negation, defined as $c_i = \alpha_{i1} \wedge \alpha_{i2} \wedge \dots \wedge \alpha_{ik_i}$ and $\alpha_{ij} = r(h, t)$ or $\neg r(h, t)$, where $r \in \mathcal{R}$, h and t are either constant entities $e \in \mathcal{E}$ or variables.

We can see that an EFO query is not a FO sentence because it contains a free variable to query. The goal of answering a query is to find all possible entities a such that $\mathcal{G} \models q[a/V?]$, where $q[a/V?]$ is a FO sentence where the free variable $V?$ is substituted by entity a . To summarize, the answer set

$$A(\mathcal{G}, q) = \{a \in \mathcal{E} : \mathcal{G} \models q[a/V?] = \text{True}\}. \quad (3)$$

Next, we show that the key condition to be verified, i.e., $\mathcal{G} \models q[a/V?] = \text{True}$, relies on the knowledge graph \mathcal{G} . By FO logical calculus, evaluating $\mathcal{G} \models q[a/V?]$ is equivalent to:

$$\bigvee_{V_1, \dots, V_n \in \mathcal{E}} \left[\bigvee_{i=1}^m \left[(\mathcal{G} \models \alpha_{i1})_{V?=a} \wedge \dots \wedge (\mathcal{G} \models \alpha_{ik_i})_{V?=a} \right] \right]. \quad (4)$$

We could see that whether $\mathcal{G} \models q[a/V?]$ is True **essentially depends on** many times of FO logical entailment on atomic formulae $\mathcal{G} \models \alpha_{ij}$, where each of these depends on whether the atomic $r(h, t)$ associated to α_{ij} is contained in \mathcal{G} (or not in \mathcal{G} for the negation case).

3 Distribution Shifts in KGQA

This section introduces the formulation of distribution shifts in KG query answering tasks. Such tasks are connected to first-order reasoning due to their close connection with logical entailment.

By clearly presenting their connection with distribution shifts studied in statistical ML community (Moreno-Torres et al., 2012), they are then able to measure the generalizable first-order reasoning capability of deep learning models, in particular transformers, with parameterized knowledge.

Problem formulation. A dataset for KG query answering is denoted as a set of samples $\{(x_i, y_i)\}$ drawn from the joint distribution $P(X, Y)$. Here, X represents the random variable corresponding to an EFO query, with x_i as individual samples, and Y represents the random variable for the answer set, with y_i as individual samples.

This paper studies transformers, where the input X is tokenized and fed into the transformer architecture, and the output Y is predicted with a classifier, where all entities in the KG are ranked by embedding similarities. In both the training and testing phases, the knowledge graph \mathcal{G} is not explicitly accessed by the model. In this way, the necessary condition for transformers to correctly answer the queries is that the knowledge in KG is materialized into their parameters. To examine whether transformers possess the **generalizable** first-order reasoning capability, we consider two classic types of out-of-distribution shifts: concept shift and covariant shift. They are connected to unobserved knowledge and unseen query types.

3.1 Two Types of Distribution Shifts

Out-Of-Distribution (OOD) generalization is one of the key topics in machine learning research (Vapnik, 1991; Quionero-Candela et al., 2009). It considers the shifts in the joint distribution $P(X, Y)$ from the training to the testing phases. Equation 5 decomposed the distribution shift by applying conditional probability formula, introducing two basic types of OOD shifts: *concept shift* refers to the changes in $P(Y|X)$, while *covariate shift* refers to the changes in $P(X)$.

$$P_{\text{train}}(Y|X)P_{\text{train}}(X) \neq \underbrace{P_{\text{test}}(Y|X)}_{\text{Concept}} \underbrace{P_{\text{test}}(X)}_{\text{Covariate}} \quad (5)$$

As illustrated in Fig. 1, we studied distribution shifts in both knowledge and query type dimensions.

Concept Shift by Unobserved Knowledge In our KG query answering setting, we construct an *observed* knowledge graph \mathcal{G}_o as a subset of the full knowledge graph \mathcal{G} . In the training stage, the models only access the answers from the observed graph, denoted as $A(\mathcal{G}_o, q)$. While in the

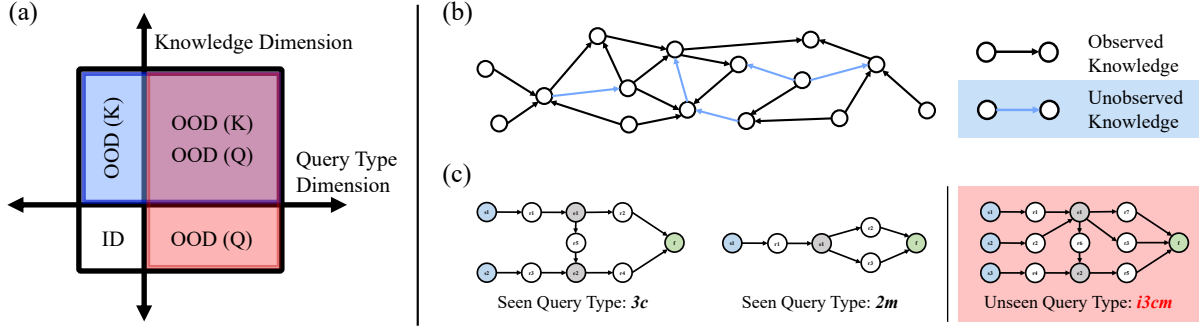


Figure 1: An illustration of our two-fold *Out-of-distribution Generalization* in knowledge and query type dimensions. Figure (a) explains the relation between OOD shifts in the two dimensions. Figure (b) and Figure (c) provide examples demonstrating the causes of OOD shifts in the two dimensions.

Dataset	Operations				Features			Unseen Types
	pro.	int.	uni.	neg.	exi.	mul.	cyc.	
GQE	✓	✓						0
Q2B	✓	✓	✓					4
BetaE	✓	✓	✓	✓				4
FIT	✓	✓	✓	✓	✓	✓	✓	10
SQE	✓	✓	✓	✓				29
Ours	✓	✓	✓	✓	✓	✓	✓	32

Table 1: Comparison between different benchmarks on their supported logical operations, query features, and the number of unseen query types.

testing stage, they are evaluated on answers from the full graph, denoted as $A(\mathcal{G}, q)$. This setup introduces a concept shift in the conditional distribution $P(Y|X)$ between training and testing because the change of knowledge causes different results of entailment, see Equation (3). Therefore, the set difference $\mathcal{A}_{ood} = A(\mathcal{G}, q) - A(\mathcal{G}_o, q)$ contains the entities that can only be derived when the model generalizes from knowledge \mathcal{G}_o to \mathcal{G} . By measuring the performances of models on \mathcal{A}_{ood} , we are able to access how a model generalizes under the concept shift caused by unobserved knowledge. This is also termed *knowledge generalization* or *knowledge inference* in the literature (Sun et al., 2021).

Covariate Shift by Unseen Query Types We also examine the capability of query answering models to generalize and maintain performance on query types that are unseen in the training stage. In our testing data, we include new query types that are combinatorially crafted from the training query types, introducing a covariate shift in the distribution of $P(X)$. The models’ performances on these unseen types are measured to evaluate their generalizability to covariate shifts in the query

type dimension. This is also termed *combinatorial generalization* or *compositional generalization* in the literature (Wang et al., 2021; Yin et al., 2023a).

4 Evaluation Benchmark

To better investigate the performance in generalizable FO reasoning, we build a new benchmark of KG query answering for our experiment. We first introduce the dataset construction process and evaluation settings, and then present our results of transformers in comparison to previous methods.

4.1 Dataset Construction

Table 1 compares the feature and statistics of query types selected in our benchmark with previous benchmarks in KG query answering, namely GQE dataset (Hamilton et al., 2019), Q2B dataset (Ren et al., 2020), BetaE dataset (Ren and Leskovec, 2020), FIT dataset (Yin et al., 2023b), and SQE dataset (Bai et al., 2023b). We included the entire set of set operations and query features in our selection of 23 seen query types. Moreover, we **compositionally crafted 32 unseen query types** that are only used in the testing stage to examine models’ OOD generalizability in the query type dimension. For instance¹, suppose a query answering model is trained on query type $2in [(r1(s1, f)) \& (! (r2(s2, f)))]$ and $2p [(r1(s1, e1)) \& (r2(e1, f))]$, we expect it can also handle the unseen query type $in2p [(r1(s1, e1)) \& (! (r2(s2, e1)))] \& (r3(e1, e2)) \& (r4(e2, f))]$, which is compositionally designed based on the previous two. All queries are sampled from the following three knowledge graphs: FB15k (Bollacker et al., 2008; Bordes et al., 2013), FB15k-

¹Queries are represented in EFO syntax (Yin et al., 2023b).

Dataset	Model	Type	ID (Q)		OOD (Q)		All Queries	
			ID (K)	OOD (K)	ID (K)	OOD (K)	ID (K)	OOD (K)
FB15k	BetaE (Ren and Leskovec, 2020)	Probabilistic	26.9	18.5	22.4	13.5	24.3	15.6
	ConE (Zhang et al., 2021)	Geometric	35.5	22.0	27.2	15.6	30.7	18.3
	CQD (Arakelyan et al., 2021)	Fuzzy Logic	33.1	20.7	21.5	11.2	26.4	15.2
	LMPNN (Wang et al., 2023b)	GNN	32.5	21.0	24.3	13.9	27.7	16.9
	SQE-LSTM (Bai et al., 2023b)	RNN	39.9	26.3	31.5	<u>18.5</u>	35.0	21.8
	Trans.+Absolute PE	Transformer	46.9	31.9	21.8	13.2	32.3	21.0
	Trans.+Disentangled PE	Transformer	51.7	33.3	23.7	13.6	35.4	21.8
	Trans.+Rotary PE	Transformer	<u>50.1</u>	<u>32.7</u>	34.6	20.8	41.1	<u>25.8</u>
	Trans.+Relative PE	Transformer	<u>48.1</u>	<u>32.3</u>	<u>35.4</u>	21.5	<u>40.7</u>	26.0
	FB15k-237	BetaE (Ren and Leskovec, 2020)	Probabilistic	27.4	16.7	23.3	13.1	25.0
ConE (Zhang et al., 2021)		Geometric	34.9	18.5	27.3	14.2	30.5	16.0
CQD (Arakelyan et al., 2021)		Fuzzy Logic	32.8	14.1	20.2	8.6	25.5	10.9
LMPNN (Wang et al., 2023b)		GNN	26.6	16.4	22.3	12.4	24.1	14.1
SQE-LSTM (Bai et al., 2023b)		RNN	45.8	17.9	<u>32.3</u>	<u>14.6</u>	37.9	<u>16.0</u>
Trans.+Absolute PE		Transformer	54.4	<u>19.9</u>	20.8	9.0	34.9	13.6
Trans.+Disentangled PE		Transformer	59.0	20.0	21.3	8.7	37.0	13.4
Trans.+Rotary PE		Transformer	54.3	19.8	36.7	15.4	44.1	17.2
Trans.+Relative PE		Transformer	<u>54.8</u>	20.0	37.7	15.8	44.8	17.6
NELL995		BetaE (Ren and Leskovec, 2020)	Probabilistic	40.9	16.7	32.3	11.1	35.9
	ConE (Zhang et al., 2021)	Geometric	46.1	17.9	38.3	<u>12.8</u>	41.5	14.9
	CQD (Arakelyan et al., 2021)	Fuzzy Logic	42.9	14.7	28.9	8.3	34.8	11.0
	LMPNN (Wang et al., 2023b)	GNN	42.8	17.1	32.1	11.6	36.6	13.9
	SQE-LSTM (Bai et al., 2023b)	RNN	65.1	17.8	48.3	12.5	55.3	<u>14.7</u>
	Trans.+Absolute PE	Transformer	68.8	18.9	25.8	8.3	43.8	12.7
	Trans.+Disentangled PE	Transformer	73.6	19.3	29.7	8.7	48.1	13.1
	Trans.+Rotary PE	Transformer	69.4	<u>19.3</u>	50.6	13.2	58.4	15.7
	Trans.+Relative PE	Transformer	68.0	18.7	<u>49.8</u>	13.4	<u>57.4</u>	<u>15.6</u>

Table 2: Experimental result (in MRR%) of Transformers in comparison with previous state-of-the-art methods with different backbones in our full benchmark. Definitions of evaluation metrics are provided in Sec. 4.2.

237 (Toutanova and Chen, 2015), and NELL995 (Carlson et al., 2010). The details of all query types and their statistics in our benchmark are provided in Appendix A. We also provide query graphs for graph-augmented methods (Liu et al., 2022; Xu et al., 2023) in our experiments, with detailed definitions in Appendix F.

4.2 Evaluation Settings

In accordance with the two types of distribution shifts discussed above, we evaluate the performance of query answering models across the two corresponding dimensions using the mean reciprocal rank (MRR) metric. In the knowledge dimension, we employ the notation ID (K) and OOD (K) to represent the performance on the answer set $\mathcal{A}_{id} = A(\mathcal{G}_o, q)$ and $\mathcal{A}_{ood} = A(\mathcal{G}, q) - A(\mathcal{G}_o, q)$, respectively. Concurrently, in the query type dimension, we utilize the notation ID (Q) and OOD (Q) to represent the performance on seen query types and unseen query types, respectively. Higher scores in ID (K) and ID (Q) reflect the models’ effectiveness in performing logical entailment on in-distribution data, while higher scores in OOD (K) and OOD (Q) demonstrated better generalizability of the models. It is also noteworthy to mention that settings in ID (K) and OOD (K) denote the concept shifts, while

those in ID (Q) and OOD (Q) capture the covariate shift. For details regarding the metric calculation, please refer to Appendix G.

4.3 General Benchmarking Results

We evaluate transformer encoders² against five existing methods for KG query answering, with details provided in Appendix B. Meanwhile, we compare four positional encoding settings in transformers: **Absolute PE** is the default sinusoidal positional encoding introduced in the original transformer (Vaswani et al., 2023). **Relative PE** (Shaw et al., 2018) applied a learnable matrix based on the relative position of tokens. **Disentangled PE** (He et al., 2021) uses separate vectors to encode the content and position of each token, decoupling the contextual and positional information during self-attention. **Rotary PE** (Su et al., 2023) encodes position information by rotating query/key vectors in the complex space, capturing relative positions through rotation angles.

The experimental results are presented in Table 2. Transformers marginally outperform all baselines in both ID (K) and OOD (K) in seen query types. Both APE and DPE transformers have poor

²Decoder-only and Encoder-Decoder yield comparable performance as illustrated in Appendix H.




Figure 2: An illustration of three core stages in the pipeline of KG query answering with transformers.

OOD generalizability in the query type dimension. In contrast, RPE and RoPE transformers demonstrates much stronger generalizability, as it can better capture structural information from the distances between tokens and flexibly handle query sequences of unseen lengths. Overall, RPE transformers achieve the best performance and OOD generalizability on first-order logical entailment.

5 Transformer Architectures and First-order Reasoning Capability

In this section, we further investigate the dependencies between transformers’ performance and multiple design choices throughout the modeling process. Figure 2 depicts the three-stage pipeline for deriving the answer set \mathcal{A} (output Y) of query q (input X) using transformers. We present the impact of query syntax in Section 5.1, token embeddings in Section 5.2, and transformer architecture in Section 5.3. We summarize the findings in Section 5.4.

5.1 Study on Query Syntax

Logical queries with the same semantics can be represented in different syntaxes with different formal languages. To investigate the impact of difference in formal languages, we conducted experiments on two formal languages: **Lisp-like Syntax** (Wang et al., 2021) represents query with fully parenthesized nested formula; **EFO Syntax** (Yin et al., 2023b) represents query with all one-hop atomics connected with conjunctions and disjunctions in parallel. As the Lisp-like syntax has a limited representation scope, we used a subset of our benchmark that includes 13 seen query types and 12 unseen query types that are supported by both languages. The experimental results on the impact of formal language are presented in Table 3. In APE Transformers, there is no difference between the results of the two formal languages. However, with RPE, the EFO syntax achieves much better generalizability in the query type domain. This phenomenon can be explained by the differences in the structural features of the two languages: with a par-

Model	Syntax	ID (Q)		OOD (Q)	
		ID (K)	OOD (K)	ID (K)	OOD (K)
Trans.+Absolute PE	Lisp-Like	57.1	20.7	10.0	4.9
	EFO	58.1	20.7	10.4	5.1
Trans.+Relative PE	Lisp-Like	58.3	20.9	22.1	9.8
	EFO	58.1	20.9	35.4	14.3

Table 3: Experimental results of Transformers in FB15k-237 with different formal languages. Details of the dataset fragment are provided in Appendix A.

Model	Before Reversion		After Reversion	
	ID(K)	OOD(K)	ID(K)	OOD(K)
Trans.+Absolute PE	54.1	18.3	27.8	9.8
Trans.+Relative PE	54.3	17.0	54.5	16.8

Table 4: Experiment results of Transformers on five query types before and after reversing permutation.

allel structure, the distances (i.e., relative positions) between tokens with the same logical relationship are more consistent than those in a nested structure. As a result, RPE can better learn the logical relationship between tokens and better generalize to unseen query types.

Furthermore, even for the same query in EFO Syntax, there can be various permutations of one-hop atomics, with the total number of potential permutations growing factorially with the length of the query. Therefore, it is crucial for transformers to robustly handle queries with different permutations. To evaluate this property, we selected five query types ($2p$, $3p$, ip , pi , $2in$) and reversed their permutations while maintaining the same semantics in the testing data (see Appendix C). The experimental results on query permutation are presented in Table 4. We observed a significant performance drop after reversing the query in APE transformers. However, for RPE transformers, the performance before and after reversing the permutation remains consistent, demonstrating their robustness in handling queries with modified permutations.

5.2 Study on Token Embeddings

Some methods in KG query answering (Arakelyan et al., 2021; Wang et al., 2023b; Xu et al., 2023)

Dataset	KGE	ID (Q)		OOD (Q)	
		ID (K)	OOD (K)	ID (K)	OOD (K)
FB15k	Random	48.1	32.3	35.4	21.5
	TransE	44.3	30.3	34.2	20.8
	DistMult	49.7	32.6	37.5	21.8
	ComplEx	<u>49.5</u>	<u>32.4</u>	<u>36.4</u>	<u>21.6</u>
FB15k-237	Random	54.8	20.0	37.7	15.8
	TransE	54.5	20.0	35.0	14.9
	DistMult	<u>57.2</u>	20.4	39.8	16.1
	ComplEx	58.6	20.4	40.4	<u>16.0</u>
NELL995	Random	68.0	18.7	49.8	13.4
	TransE	68.5	18.6	48.7	12.5
	DistMult	72.2	19.9	51.2	12.9
	ComplEx	72.2	<u>19.8</u>	<u>49.9</u>	<u>13.1</u>

Table 5: Experimental results of Transformer + Relative PE with different pre-trained KGEs.

utilized pre-trained KG embeddings from link predictors in their query encoding process, while other methods (Ren and Leskovec, 2020; Zhang et al., 2021; Bai et al., 2023b) randomly initialize the embedding matrix and jointly learn the KG embedding and model parameters from the training process. It’s still unclear that to what extent can pre-trained KG semantics can assist the performance of query answering models. To rigorously assess the impact of pre-trained semantics in token embeddings, we have selected three widely recognized pre-trained KG embeddings: **TransE** (Bordes et al., 2013) models relationships by interpreting them as translations within the embedding space, **DistMult** (Yang et al., 2015) employs a bilinear model wherein entities are represented as vectors and relations as diagonal matrices, **ComplEx** (Trouillon et al., 2016) extends this framework into a complex vector space, allowing for more nuanced inter-entity interactions. In our experiment, all three KG embeddings are frozen during the training stage, in comparison to a baseline with randomly initialized and jointly learned embeddings.

The experimental results are presented in Table 5. Two stronger KGEs, ComplEx and DistMult, can effectively improve the performance of transformer. In contrast, the performance of TransE consistently lags behind the baseline across all three datasets. This discrepancy can be attributed to the baseline model’s embedding matrix learning process, which can be viewed as implicitly training a link predictor using a transformer architecture (e.g., KG-BERT (Yao et al., 2019)). This modern approach may naturally outperform earlier models like TransE.

Model	ID (Q)		OOD (Q)	
	ID (K)	OOD (K)	ID (K)	OOD (K)
Trans.+Absolute PE (baseline)	54.4	19.9	20.8	9.0
+Adjacency Matrices Masking	46.5	19.6	<u>22.1</u>	11.2
+Directed Distance Encoding	55.0	20.1	<u>22.7</u>	<u>9.8</u>
Trans.+Relative PE (baseline)	54.8	20.0	37.7	15.8
+Adjacency Matrices Masking	51.3	20.1	28.3	13.0
+Directed Distance Encoding	54.1	<u>20.0</u>	<u>37.5</u>	15.9

Table 6: Experimental result of Transformer + Absolute/Relative PE in FB15k-237 with existing inductive biases. Considering its triviality and compatibility with our dataset, the inductive bias from BiQE is not included in our experiment.

5.3 Study on Transformer Architecture

Some existing literature (Kotnis et al., 2021; Liu et al., 2022; Xu et al., 2023) has explored the use and design of transformers in KG query answering, in which several inductive biases in transformer architecture were proposed. However, as all these methods focused on different aspects and were evaluated on different benchmarks, a clearer investigation is needed. We implemented two designs from previous works on top of transformers: **Adjacency Matrices Masking**: an inductive bias proposed in kgTransformer (Liu et al., 2022), which aims to enhance reasoning performance by limiting the range of self-attention to one-hop neighboring nodes in the query graph. **Directed Distance Encoding**: another inductive bias, introduced in Query2Triple (Xu et al., 2023), that facilitates information aggregation by injecting a directed distance message into the self-attention mechanism. The experimental results on existing transformer designs are presented in Table 6. Both inductive biases effectively improve generalizability in the query type dimension when using APE. In contrast, under RPE, these improvements are not observed, and the use of adjacency matrices masking even results in a significant decrease in performance. This decrease is attributed to the masking in self-attention, which restricts message passing to only nodes within one hop, thereby limiting the model’s ability to perform complex logical reasoning. Regarding directed distance encoding, the information about directed distances is mostly contained in the relative positions of the linearized sequence, thus the improvements seen with original APE are not sustained.

5.4 Summary of Empirical Findings

In terms of query syntax, we discovered that **employing formal language with parallel structure**

can significantly enhance OOD generalizability in the query type dimension for RPE transformers. Moreover, we observed that **transformers with relative PE can robustly handle queries with different permutations**. Regarding transformer architecture, we found that the inductive biases proposed in previous approaches are **only effective under APE, but fail to improve upon RPE transformers**, which we have shown to be the preferred setting. This mismatch between transformer design and positional encoding settings motivated us to design a methodology that effectively boosts the performance and generalizability of RPE transformers in first-order logical entailment.

6 Transformer Encoder with Guided Attention

One of the main challenges in KG query answering is to perform multi-hop logical reasoning following the one-hop atomics and the logical operations between them. In traditional reasoning tasks, such as language-based reasoning, transformer models perform multi-hop reasoning by procedurally forming “reasoning trees” within their attention mechanisms (Hou et al., 2023; Murty et al., 2023; Wang et al., 2023a). While in our task, it is challenging for transformers to capture the pattern of reasoning trees behind the queries, as the semantics of logical operators are neither explicitly modeled nor pre-trained, but implicitly learned from limited training queries. To navigate this conundrum, we present **TEGA (Transformer Encoder with Guided Attention)**, a modeling methodology that boosts logical reasoning performance by providing logic-aware guidance within the self-attention mechanism of transformers. The following sections will introduce the two inductive biases in TEGA and their effectiveness with experimental results.

6.1 Logic-aware Relative Positional Encoding

Our analysis so far shows that RPE is the best positional encoding for transformers in our task for its superior performance and generalizability in both knowledge and query type dimensions. However, it still has limited expressiveness in the logical relationship between tokens, as the tokens with the same relative distance can have different logical relationships among queries. To fill this gap, we propose LogiRPE, a logic-aware mechanism to enhance self-attention. Specifically, for a sequence input sequence embeddings $x_i \in \mathbb{R}^d$, $i = 1, \dots, n$,

the enhanced self-attention is computed by

$$z_i = \sum_{j=1}^n \alpha_{ij} (x_j W^V + \beta_{ij}^V), \quad \alpha_{ij} = \frac{\exp e_{ij}}{\sum_{k=1}^n \exp e_{ik}} \quad (6)$$

$$e_{ij} = \frac{x_i W^Q (x_j W^K + \beta_{ij}^K)^T}{\sqrt{d}}. \quad (7)$$

where $W^Q, W^K, W^V \in \mathbb{R}^{d \times d}$ are the weight matrices for query, key, and value while $\beta_{ij}^V, \beta_{ij}^K \in \mathbb{R}^d$ are two *logical bias* terms introduced by LogiRPE to capture the missing logical relations.

The key feature of *logical bias* terms is that they should be able to discriminate different types of tokens. Let the type of the i -th token be t_i from a finite type set of size t . In our study, we consider six types of tokens: parenthesis[$(,)$], entity[s, e, f], relation[r], conjunction[\wedge], disjunction[\vee], and negation[$!$]. Then the bias vector β_{ij}^l (l is either K or V) is indexed from an trainable embedding bank $\omega^l \in \mathbb{R}^{t \times t \times n \times d}$ (consider the first three indices).

$$\beta_{ij}^l = \omega^l[t_i, t_j, |i - j|], \quad l \in \{K, V\}. \quad (8)$$

LogiRPE may also be applicable to broader tasks where tokens with logical semantics can be labeled.

6.2 Free-variable Pooling

Most existing transformer-based approaches (Bai et al., 2023b; Xu et al., 2023) utilize the last layer hidden states of the first token(s) as the final query encoding, akin to the use of the $[CLS]$ token in the original BERT paper (Devlin et al., 2019). However, we contend that the reasoning capabilities inherent in self-attention mechanisms should effectively aggregate the query information into free variables that semantically represent the answer sets of queries. Consequently, we adopt Free-Variable Pooling, wherein the final query encoding is derived through max pooling across the last-layer hidden states of these free variables.

6.3 Result and Ablation

The experimental results of TEGA and Transformer baselines are presented in Table 7. In all three datasets, TEGA substantially improved the performance and generalizability of knowledge and query type dimensions over baselines. According to our ablation study results in Table 8, both inductive biases proposed in TEGA are effective on their own. At the same time, a stronger improvement in entailment performance can be achieved when applied together. Moreover, we provide an example of

Dataset	Model	ID (Q)		OOD (Q)	
		ID (K)	OOD (K)	ID (K)	OOD (K)
FB15k	Trans.+Absolute PE	46.9	31.9	21.8	13.2
	Trans.+Relative PE	48.1	32.3	35.4	21.5
	TEGA (ours)	55.0	34.1	38.8	22.4
FB15k-237	Trans.+Absolute PE	54.4	19.9	20.8	9.0
	Trans.+Relative PE	54.8	20.0	37.7	15.8
	TEGA (ours)	64.3	20.1	42.6	16.0
NELL995	Trans.+Absolute PE	68.8	18.9	25.8	8.3
	Trans.+Relative PE	68.0	18.7	49.8	13.4
	TEGA (ours)	75.0	19.2	56.6	13.7

Table 7: Experimental results of TEGA comparing to baselines in FB15k, FB15k-237 and NELL995.

Model	ID (Q)		OOD (Q)	
	ID (K)	OOD (K)	ID (K)	OOD (K)
TEGA (ours)	55.0	34.1	38.8	22.4
w/o LogiRPE	<u>51.9</u>	<u>33.2</u>	<u>37.0</u>	21.6
w/o Free-Variable Pooling	50.9	33.1	36.5	<u>22.2</u>
Trans.+Relative PE (baseline)	48.1	32.3	35.4	21.5

Table 8: Ablation study on two inductive biases introduced in TEGA on FB15k.

attention visualization in Appendix D, demonstrating the self-evident effect of logic-aware attention guidance within our TEGA model.

7 Related Work

Transformers in Logical Reasoning Transformers have exhibited remarkable performance in various forms of logical reasoning. In natural language inference (NLI) tasks (Bowman et al., 2015; Williams et al., 2018), transformers analyze the logical relationship between a premise and a hypothesis, employing deductive reasoning over statements. Furthermore, transformers have been employed for inductive reasoning within rule-based systems (Clark et al., 2020). Additionally, transformers have been applied to perform abductive reasoning in both natural language (Bhagavatula et al., 2020) and formal language (Bai et al., 2024; Gao et al., 2025) settings. Meanwhile, various techniques have been proposed to enhance logical reasoning in language understanding (Chen, 2023; Pan et al., 2023). Recent studies show that large language models (LLMs) exhibit enhanced logical reasoning abilities (Parmar et al., 2024; Li et al., 2025; Fan et al., 2025), even without intermediate reasoning steps (Zheng et al., 2025a). Logical entailment also plays a key role in hypothesis verification of LLMs, in areas such as commonsense reasoning (Zhao et al., 2024), analogical reasoning (Zheng et al., 2025b), and even scientific discovery

(Takagi et al., 2023; Zheng et al., 2025c).

KGQA with Parameterized Knowledge Prior to this work, extensive research has been conducted on KGQA with parameterized knowledge. Regarding modeling approaches, iterative neural query encoders (Ren and Leskovec, 2020; Chen et al., 2022; Tsang et al., 2025) design representations for entity sets and execute logical operators iteratively, following the computational graph. In addition, neural-symbolic methods (Zhu et al., 2022; Bai et al., 2023c; Yin et al., 2023b) incorporate link predictors and search over the symbolic space. In terms of the knowledge domain, existing benchmarks (Ren et al., 2020; Ren and Leskovec, 2020; Yin et al., 2023a) primarily utilize knowledge graphs containing general factual knowledge, while some studies have focused on commonsense knowledge (Fang et al., 2024) and eventuality knowledge (Bai et al., 2023a). Such queries can be extended to natural language settings with template or LLM-based approaches (Zheng et al., 2024a,b; Zong et al., 2025). The setting of KGQA with parameterized knowledge has also been extended to the concept of *Neural Graph Databases* (Ren et al., 2023; Bai et al., 2025).

8 Conclusion

This paper investigates the generalizable first-order logical reasoning capabilities of transformers when knowledge is parameterized within model weights. By establishing connections between distribution shifts and logical entailment, we develop a principled framework for evaluating transformer performance on knowledge graph query answering tasks. Our comprehensive benchmark reveals that transformers can effectively perform first-order logical entailment, outperforming methods specifically designed for this task. Through systematic analysis of the entire modeling pipeline, we identify critical design principles that enhance reasoning capabilities. Notably, our findings reveal a mismatch between existing inductive biases and optimal positional encoding strategies, motivating the development of TEGA. This architecture introduces logic-aware self-attention mechanisms that substantially improve both performance and generalizability. Our work demonstrates that careful architectural choices, informed by the logical structure of reasoning tasks, can significantly enhance transformers’ ability to perform complex logical inference.

Limitations

The scope of complex logical queries was limited to EFO queries with a single free variable, with queries containing multiple free variables reserved for future investigation. While LogiRPE incorporates fine-grained logic awareness, it remains dependent on explicit labeling. Future research should explore more generalizable approaches to logic-aware labeling that can accommodate previously unseen logical operators. With respect to knowledge graph scale, the largest knowledge graph employed in our experiments was FB15k, which is considered medium-sized within the knowledge graph literature. Consequently, the scalability of the proposed approaches to large-scale knowledge graphs remains an open question for future investigation.

Ethics Statement

The experiments were conducted on publicly available knowledge graphs, eliminating any data privacy concerns. However, it should be noticed that most approaches for generalizable logical entailment are susceptible to adversarial attacks (Dai et al., 2018; Zügner et al., 2019) and data poisoning (You et al., 2023) on knowledge graphs, which may result in unintended outcomes in applications.

Acknowledgements

We thank all the anonymous reviewers and meta reviewers for their valuable comments. The authors of this paper were supported by the ITSP Platform Research Project (ITS/189/23FP) from ITC of Hong Kong, SAR, China, and the AoE (AoE/E-601/24-N), the RIF (R6021-20) and the GRF (16205322) from RGC of Hong Kong, SAR, China.

References

- Erik Arakelyan, Daniel Daza, Pasquale Minervini, and Michael Cochez. 2021. [Complex query answering with neural link predictors](#). *Preprint*, arXiv:2011.03459.
- Jiaxin Bai, Xin Liu, Weiqi Wang, Chen Luo, and Yangqiu Song. 2023a. [Complex query answering on eventuality knowledge graph with implicit logical constraints](#). *Preprint*, arXiv:2305.19068.
- Jiaxin Bai, Yicheng Wang, Tianshi Zheng, Yue Guo, Xin Liu, and Yangqiu Song. 2024. [Advancing abductive reasoning in knowledge graphs through complex logical hypothesis generation](#). *Preprint*, arXiv:2312.15643.
- Jiaxin Bai, Zihao Wang, Yukun Zhou, Hang Yin, Weizhi Fei, Qi Hu, Zheye Deng, Jiayang Cheng, Tianshi Zheng, Hong Ting Tsang, Yisen Gao, Zhongwei Xie, Yufei Li, Lixin Fan, Binhang Yuan, Wei Wang, Lei Chen, Xiaofang Zhou, and Yangqiu Song. 2025. [Top ten challenges towards agentic neural graph databases](#). *Preprint*, arXiv:2501.14224.
- Jiaxin Bai, Tianshi Zheng, and Yangqiu Song. 2023b. [Sequential query encoding for complex query answering on knowledge graphs](#). *Transactions on Machine Learning Research*.
- Yushi Bai, Xin Lv, Juanzi Li, and Lei Hou. 2023c. [Answering complex logical queries on knowledge graphs via query computation tree optimization](#). *Preprint*, arXiv:2212.09567.
- David G. T. Barrett, Felix Hill, Adam Santoro, Ari S. Morcos, and Timothy Lillicrap. 2018. [Measuring abstract reasoning in neural networks](#). *Preprint*, arXiv:1807.04225.
- Chandra Bhagavatula, Ronan Le Bras, Chaitanya Malaviya, Keisuke Sakaguchi, Ari Holtzman, Han-nah Rashkin, Doug Downey, Scott Wen tau Yih, and Yejin Choi. 2020. [Abductive commonsense reasoning](#). *Preprint*, arXiv:1908.05739.
- Kurt D. Bollacker, Colin Evans, Praveen K. Paritosh, Tim Sturge, and Jamie Taylor. 2008. [Freebase: a collaboratively created graph database for structuring human knowledge](#). In *SIGMOD Conference*.
- Antoine Bordes, Nicolas Usunier, Alberto Garcia-Duran, Jason Weston, and Oksana Yakhnenko. 2013. [Translating embeddings for modeling multi-relational data](#). In *Advances in Neural Information Processing Systems*, volume 26. Curran Associates, Inc.
- Samuel R. Bowman, Gabor Angeli, Christopher Potts, and Christopher D. Manning. 2015. [A large annotated corpus for learning natural language inference](#). *Preprint*, arXiv:1508.05326.
- Andrew Carlson, Justin Betteridge, Bryan Kisiel, Burr Settles, Estevam R. Hruschka, and Tom M. Mitchell. 2010. [Toward an architecture for never-ending language learning](#). In *Proceedings of the Twenty-Fourth AAAI Conference on Artificial Intelligence, AAAI'10*, page 1306–1313. AAAI Press.
- Jianshu Chen. 2023. [Learning language representations with logical inductive bias](#). *Preprint*, arXiv:2302.09458.
- Xuelu Chen, Ziniu Hu, and Yizhou Sun. 2022. [Fuzzy logic based logical query answering on knowledge graphs](#). *Preprint*, arXiv:2108.02390.
- Peter Clark, Oyvind Tafjord, and Kyle Richardson. 2020. [Transformers as soft reasoners over language](#). *Preprint*, arXiv:2002.05867.

- Hanjun Dai, Hui Li, Tian Tian, Xin Huang, Lin Wang, Jun Zhu, and Le Song. 2018. [Adversarial attack on graph structured data](#). In *Proceedings of the 35th International Conference on Machine Learning*, volume 80 of *Proceedings of Machine Learning Research*, pages 1115–1124. PMLR.
- Mostafa Dehghani, Stephan Gouws, Oriol Vinyals, Jakob Uszkoreit, and Łukasz Kaiser. 2019. [Universal transformers](#). *Preprint*, arXiv:1807.03819.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. [Bert: Pre-training of deep bidirectional transformers for language understanding](#). *Preprint*, arXiv:1810.04805.
- Wei Fan, Tianshi Zheng, Yiran Hu, Zheyang Deng, Weiqi Wang, Baixuan Xu, Chunyang Li, Haoran Li, Weixing Shen, and Yangqiu Song. 2025. [Legal rule induction: Towards generalizable principle discovery from analogous judicial precedents](#). *Preprint*, arXiv:2505.14104.
- Tianqing Fang, Zeming Chen, Yangqiu Song, and Antoine Bosselut. 2024. [Complex reasoning over logical queries on commonsense knowledge graphs](#). *Preprint*, arXiv:2403.07398.
- Yisen Gao, Jiaxin Bai, Tianshi Zheng, Qingyun Sun, Ziwei Zhang, Jianxin Li, Yangqiu Song, and Xingcheng Fu. 2025. [Controllable logical hypothesis generation for abductive reasoning in knowledge graphs](#). *Preprint*, arXiv:2505.20948.
- William L. Hamilton, Payal Bajaj, Marinka Zitnik, Dan Jurafsky, and Jure Leskovec. 2019. [Embedding logical queries on knowledge graphs](#). *Preprint*, arXiv:1806.01445.
- Simeng Han, Hailey Schoelkopf, Yilun Zhao, Zhenting Qi, Martin Riddell, Luke Benson, Lucy Sun, Ekaterina Zubova, Yujie Qiao, Matthew Burtell, David Peng, Jonathan Fan, Yixin Liu, Brian Wong, Malcolm Sailor, Ansong Ni, Linyong Nan, Jungo Kasai, Tao Yu, Rui Zhang, Shafiq Joty, Alexander R. Fabri, Wojciech Kryscinski, Xi Victoria Lin, Caiming Xiong, and Dragomir Radev. 2022. [Folio: Natural language reasoning with first-order logic](#). *Preprint*, arXiv:2209.00840.
- Pengcheng He, Xiaodong Liu, Jianfeng Gao, and Weizhu Chen. 2021. [Deberta: Decoding-enhanced bert with disentangled attention](#). *Preprint*, arXiv:2006.03654.
- Dan Hendrycks, Collin Burns, Saurav Kadavath, Akul Arora, Steven Basart, Eric Tang, Dawn Song, and Jacob Steinhardt. 2021. [Measuring mathematical problem solving with the math dataset](#). *Preprint*, arXiv:2103.03874.
- Sepp Hochreiter and Jürgen Schmidhuber. 1997. Long short-term memory. *Neural computation*, 9(8):1735–1780.
- Yifan Hou, Jiaoda Li, Yu Fei, Alessandro Stolfo, Wangchunshu Zhou, Guangtao Zeng, Antoine Bosselut, and Mrinmaya Sachan. 2023. [Towards a mechanistic interpretation of multi-step reasoning capabilities of language models](#). *Preprint*, arXiv:2310.14491.
- Bhushan Kotnis, Carolin Lawrence, and Mathias Niepert. 2021. [Answering complex queries in knowledge graphs with bidirectional sequence encoders](#). *Proceedings of the AAAI Conference on Artificial Intelligence*, 35(6):4968–4977.
- Guillaume Lample and François Charton. 2019. [Deep learning for symbolic mathematics](#). *Preprint*, arXiv:1912.01412.
- Chunyang Li, Weiqi Wang, Tianshi Zheng, and Yangqiu Song. 2025. [Patterns over principles: The fragility of inductive reasoning in llms under noisy observations](#). *Preprint*, arXiv:2502.16169.
- Xiao Liu, Shiyu Zhao, Kai Su, Yukuo Cen, Jiezhong Qiu, Mengdi Zhang, Wei Wu, Yuxiao Dong, and Jie Tang. 2022. [Mask and reason: Pre-training knowledge graph transformers for complex logical queries](#). In *Proceedings of the 28th ACM SIGKDD Conference on Knowledge Discovery and Data Mining, KDD '22*. ACM.
- David Marker. 2002. *Model Theory: An Introduction*. Springer-Verlag, New York.
- Jose G Moreno-Torres, Troy Raeder, Rocío Alaiz-Rodríguez, Nitesh V Chawla, and Francisco Herrera. 2012. A unifying view on dataset shift in classification. *Pattern recognition*, 45(1):521–530.
- Shikhar Murty, Pratyusha Sharma, Jacob Andreas, and Christopher Manning. 2023. [Grokking of hierarchical structure in vanilla transformers](#). In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, pages 439–448, Toronto, Canada. Association for Computational Linguistics.
- Liangming Pan, Alon Albalak, Xinyi Wang, and William Yang Wang. 2023. [Logic-lm: Empowering large language models with symbolic solvers for faithful logical reasoning](#). *Preprint*, arXiv:2305.12295.
- Mihir Parmar, Nisarg Patel, Neeraj Varshney, Mutsumi Nakamura, Man Luo, Santosh Mashetty, Arindam Mitra, and Chitta Baral. 2024. [Logicbench: Towards systematic evaluation of logical reasoning ability of large language models](#). *Preprint*, arXiv:2404.15522.
- Stanislas Polu and Ilya Sutskever. 2020. [Generative language modeling for automated theorem proving](#). *Preprint*, arXiv:2009.03393.
- Joaquin Quionero-Candela, Masashi Sugiyama, Anton Schwaighofer, and Neil Lawrence. 2009. Dataset shift in machine learning.

- Hongyu Ren, Mikhail Galkin, Michael Cochez, Zhaocheng Zhu, and Jure Leskovec. 2023. [Neural graph reasoning: Complex logical query answering meets graph databases](#). *Preprint*, arXiv:2303.14617.
- Hongyu Ren, Weihua Hu, and Jure Leskovec. 2020. [Query2box: Reasoning over knowledge graphs in vector space using box embeddings](#). *Preprint*, arXiv:2002.05969.
- Hongyu Ren and Jure Leskovec. 2020. [Beta embeddings for multi-hop logical reasoning in knowledge graphs](#). *Preprint*, arXiv:2010.11465.
- Abulhair Saparov, Richard Yuanzhe Pang, Vishakh Padmakumar, Nitish Joshi, Mehran Kazemi, Najoung Kim, and He He. 2023. Testing the general deductive reasoning capacity of large language models using ood examples. *Advances in Neural Information Processing Systems*, 36.
- David Saxton, Edward Grefenstette, Felix Hill, and Pushmeet Kohli. 2019. [Analysing mathematical reasoning abilities of neural models](#). *Preprint*, arXiv:1904.01557.
- Peter Shaw, Jakob Uszkoreit, and Ashish Vaswani. 2018. [Self-attention with relative position representations](#). *Preprint*, arXiv:1803.02155.
- Jianlin Su, Yu Lu, Shengfeng Pan, Ahmed Murtadha, Bo Wen, and Yunfeng Liu. 2023. [Roformer: Enhanced transformer with rotary position embedding](#). *Preprint*, arXiv:2104.09864.
- Haitian Sun, Andrew O. Arnold, Tania Bedrax-Weiss, Fernando Pereira, and William W. Cohen. 2021. [Faithful embeddings for knowledge base queries](#). *Preprint*, arXiv:2004.03658.
- Christian Szegedy, Vincent Vanhoucke, Sergey Ioffe, Jonathon Shlens, and Zbigniew Wojna. 2015. [Re-thinking the inception architecture for computer vision](#). *Preprint*, arXiv:1512.00567.
- Shiro Takagi, Ryutaro Yamauchi, and Wataru Kumagai. 2023. [Towards autonomous hypothesis verification via language models with minimal guidance](#). *Preprint*, arXiv:2311.09706.
- Jidong Tian, Yitian Li, Wenqing Chen, Liqiang Xiao, Hao He, and Yaohui Jin. 2021. [Diagnosing the first-order logical reasoning ability through LogicNLI](#). In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 3738–3747, Online and Punta Cana, Dominican Republic. Association for Computational Linguistics.
- Kristina Toutanova and Danqi Chen. 2015. [Observed versus latent features for knowledge base and text inference](#). In *Proceedings of the 3rd Workshop on Continuous Vector Space Models and their Compositionality*, pages 57–66, Beijing, China. Association for Computational Linguistics.
- Théo Trouillon, Johannes Welbl, Sebastian Riedel, Éric Gaussier, and Guillaume Bouchard. 2016. [Complex embeddings for simple link prediction](#). *Preprint*, arXiv:1606.06357.
- Hong Ting Tsang, Zihao Wang, and Yangqiu Song. 2025. [Transformers for complex query answering over knowledge hypergraphs](#). *Preprint*, arXiv:2504.16537.
- V. Vapnik. 1991. [Principles of risk minimization for learning theory](#). In *Advances in Neural Information Processing Systems*, volume 4. Morgan-Kaufmann.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N. Gomez, Lukasz Kaiser, and Illia Polosukhin. 2023. [Attention is all you need](#). *Preprint*, arXiv:1706.03762.
- Lean Wang, Lei Li, Damai Dai, Deli Chen, Hao Zhou, Fandong Meng, Jie Zhou, and Xu Sun. 2023a. [Label words are anchors: An information flow perspective for understanding in-context learning](#). *Preprint*, arXiv:2305.14160.
- Zihao Wang, Yangqiu Song, Ginny Y. Wong, and Simon See. 2023b. [Logical message passing networks with one-hop inference on atomic formulas](#). *Preprint*, arXiv:2301.08859.
- Zihao Wang, Hang Yin, and Yangqiu Song. 2021. [Benchmarking the combinatorial generalizability of complex query answering on knowledge graphs](#). *Preprint*, arXiv:2109.08925.
- Adina Williams, Nikita Nangia, and Samuel R. Bowman. 2018. [A broad-coverage challenge corpus for sentence understanding through inference](#). *Preprint*, arXiv:1704.05426.
- Keyulu Xu, Weihua Hu, Jure Leskovec, and Stefanie Jegelka. 2019. [How powerful are graph neural networks?](#) *Preprint*, arXiv:1810.00826.
- Yao Xu, Shizhu He, Cunguang Wang, Li Cai, Kang Liu, and Jun Zhao. 2023. [Query2Triple: Unified query encoding for answering diverse complex queries over knowledge graphs](#). In *Findings of the Association for Computational Linguistics: EMNLP 2023*, pages 11369–11382, Singapore. Association for Computational Linguistics.
- Bishan Yang, Wen tau Yih, Xiaodong He, Jianfeng Gao, and Li Deng. 2015. [Embedding entities and relations for learning and inference in knowledge bases](#). *Preprint*, arXiv:1412.6575.
- Liang Yao, Chengsheng Mao, and Yuan Luo. 2019. [Kgbert: Bert for knowledge graph completion](#). *Preprint*, arXiv:1909.03193.
- Hang Yin, Zihao Wang, Weizhi Fei, and Yangqiu Song. 2023a. [EFO_k-cqa: Towards knowledge graph complex query answering beyond set operation](#). *Preprint*, arXiv:2307.13701.

- Hang Yin, Zihao Wang, and Yangqiu Song. 2023b. [Rethinking complex queries on knowledge graphs with neural link predictors](#). *Preprint*, arXiv:2304.07063.
- Xiaoyu You, Beina Sheng, Daizong Ding, Mi Zhang, Xudong Pan, Min Yang, and Fuli Feng. 2023. [Mass: Model-agnostic, semantic and stealthy data poisoning attack on knowledge graph embedding](#). In *Proceedings of the ACM Web Conference 2023, WWW '23*, page 2000–2010, New York, NY, USA. Association for Computing Machinery.
- Zhanqiu Zhang, Jie Wang, Jiajun Chen, Shuiwang Ji, and Feng Wu. 2021. [Cone: Cone embeddings for multi-hop reasoning over knowledge graphs](#). *Preprint*, arXiv:2110.13715.
- Wenting Zhao, Justin T Chiu, Jena D. Hwang, Faeze Brahman, Jack Hessel, Sanjiban Choudhury, Yejin Choi, Xiang Lorraine Li, and Alane Suhr. 2024. [Uncommonsense reasoning: Abductive reasoning about uncommon situations](#). *Preprint*, arXiv:2311.08469.
- Tianshi Zheng, Jiaxin Bai, Yicheng Wang, Tianqing Fang, Yue Guo, Yauwai Yim, and Yangqiu Song. 2024a. [Clr-fact: Evaluating the complex logical reasoning capability of large language models over factual knowledge](#). *Preprint*, arXiv:2407.20564.
- Tianshi Zheng, Yixiang Chen, Chengxi Li, Chunyang Li, Qing Zong, Haochen Shi, Baixuan Xu, Yangqiu Song, Ginny Y. Wong, and Simon See. 2025a. [The curse of cot: On the limitations of chain-of-thought in in-context learning](#). *Preprint*, arXiv:2504.05081.
- Tianshi Zheng, Jiayang Cheng, Chunyang Li, Haochen Shi, Zihao Wang, Jiaxin Bai, Yangqiu Song, Ginny Y. Wong, and Simon See. 2025b. [Logidynamics: Unraveling the dynamics of logical inference in large language model reasoning](#). *Preprint*, arXiv:2502.11176.
- Tianshi Zheng, Zheye Deng, Hong Ting Tsang, Weiqi Wang, Jiaxin Bai, Zihao Wang, and Yangqiu Song. 2025c. [From automation to autonomy: A survey on large language models in scientific discovery](#). *Preprint*, arXiv:2505.13259.
- Tianshi Zheng, Weihai Li, Jiaxin Bai, Weiqi Wang, and Yangqiu Song. 2024b. [Assessing the robustness of retrieval-augmented generation systems in k-12 educational question answering with knowledge discrepancies](#). *Preprint*, arXiv:2412.08985.
- Zhaocheng Zhu, Mikhail Galkin, Zuobai Zhang, and Jian Tang. 2022. [Neural-symbolic models for logical queries on knowledge graphs](#). *Preprint*, arXiv:2205.10128.
- Qing Zong, Zhaowei Wang, Tianshi Zheng, Xiyu Ren, and Yangqiu Song. 2025. [Comparisonqa: Evaluating factuality robustness of llms through knowledge frequency control and uncertainty](#). *Preprint*, arXiv:2412.20251.
- Daniel Zügner, Amir Akbarnejad, and Stephan Günnemann. 2019. [Adversarial attacks on neural networks for graph data](#). In *Proceedings of the Twenty-Eighth International Joint Conference on Artificial Intelligence, IJCAI-19*, pages 6246–6250. International Joint Conferences on Artificial Intelligence Organization.

A Dataset Statistics

In this section, we introduce the knowledge graph details and the statistics and query types of our benchmark dataset.

Our benchmark incorporates three widely used knowledge graphs—FB15k, FB15k-237, and NELL995—derived from large-scale KGs to evaluate reasoning capabilities. FB15k and FB15k-237 originate from Freebase, a comprehensive database of structured general knowledge covering domains such as entertainment, sports, and geography, with entities linked by diverse relations (e.g., "born-in," "part-of"). NELL995 is derived from the Never-Ending Language Learning (NELL) system, which extracts structured facts from web text, focusing on real-world entities and relations.

Details of the query statistics for our dataset are presented in Table 9. For each knowledge graph, we sampled all one-hop projection (*1p*) queries from the training graph and sampled twice as many for the other 22 in-distribution query types. For validation and testing queries, we set the quantities to 8000, 5000, and 4000, respectively, following the convention established in KG literature (Ren et al., 2020; Ren and Leskovec, 2020; Bai et al., 2023b).

Knowledge Graph	Training		Validation	Testing
	1p	Other Types	All Types	All Types
FB15k	273,710	547,420	8,000	8,000
FB15k-237	149,689	299,378	5,000	5,000
NELL995	107,982	215,964	4,000	4,000

Table 9: Number of queries used for each query type in our benchmark.

Table 11 and 12 presented the details of query types in our benchmark. All 32 unseen query types are crafted compositionally based on the 23 seen query types. Table 13 presented the subset of 25 lisp-like compatible query types used in our experiment on query syntax.

B Baseline Models

In our benchmarking experiment, we compare transformers with five existing methods that are designed particularly for KG query answering:

- **BetaE** (Ren and Leskovec, 2020) encodes queries iteratively following the logical operations with Beta distributions.
- **ConE** (Zhang et al., 2021) represents answer sets as two-dimensional cones to better handle negations.
- **CQD** (Arakelyan et al., 2021) applies a search-based method that utilizes pre-trained link predictors for inference-time optimization.
- **LMPNN** (Wang et al., 2023b) conducts one-hop inferences with pre-trained KGE, and uses message passing to aggregate these local results into global output.
- **SQE-LSTM** (Bai et al., 2023b) encode the linearized complex query as a sequence with LSTM (Hochreiter and Schmidhuber, 1997).

C Query Reversion

The logical formula of five query types before and after reversion is provided in Table 10.

Type Name	Before Reversion	After Reversion
2p	$(r1(s1,e1)) \& (r2(e1,f))$	$(r2(e1,f)) \& (r1(s1,e1))$
3p	$((r1(s1,e1)) \& (r2(e1,e2))) \& (r3(e2,f))$	$((r3(e2,f)) \& (r2(e1,e2))) \& (r1(s1,e1))$
ip	$((r1(s1,e1)) \& (r2(s2,e1))) \& (r3(e1,f))$	$((r3(e1,f)) \& (r2(s2,e1))) \& (r1(s1,e1))$
pi	$((r1(s1,e1)) \& (r2(e1,f))) \& (r3(s2,f))$	$((r3(s2,f)) \& (r2(e1,f))) \& (r1(s1,e1))$
2in	$(r1(s1,f)) \& (!r2(s2,f))$	$(!r2(s2,f)) \& (r1(s1,f))$

Table 10: Details of query reversion procedure in the query permutation experiment.

D Attention Visualization

Fig. 3 provides an example of how attention guidance in TEGA can affect the reasoning process in self-attention. In *2in* queries $[(r1(s1, f)) \& (!r2(s2, f))]$, the negation operator $[!]$ is applied to calculate the complement of the entity set represented by the second atomic formula $[(r2(s2, f))]$. Benefited from the attention guidance from inductive biases, TEGA can effectively aggregate the information from the second atomic formula $[r2/s2/f]$ to the negation operator with self-attention. While in baseline (Transformer + Relative PE), the information flow between tokens are scattered and less effective (highlighted in red). As a result, TEGA outperforms Transformer + Relative PE by **21-26%** in *2in* queries over three knowledge graphs.




Figure 3: An visualization of self-attention when processing *2in* queries. The left and right image is the partial attention weight of TEGA (ours) and Transformer + Relative PE (baseline), respectively.

id	type name	mul.	cyc.	exi.	neg.	depth	logical formula
0	1p					1	$r1(s1,f)$
1	2p					2	$(r1(s1,e1))\&(r2(e1,f))$
2	3p					3	$(r1(s1,e1))\&(r2(e1,e2))\&(r3(e2,f))$
3	2i					1	$(r1(s1,f))\&(r2(s2,f))$
4	3i					1	$(r1(s1,f))\&(r2(s2,f))\&(r3(s3,f))$
5	ip					2	$(r1(s1,e1))\&(r2(s2,e1))\&(r3(e1,f))$
6	pi					2	$(r1(s1,e1))\&(r2(e1,f))\&(r3(s2,f))$
7	2in				✓	1	$(r1(s1,f))\&(!r2(s2,f))$
8	3in				✓	1	$(r1(s1,f))\&(r2(s2,f))\&(!r3(s3,f))$
9	inp				✓	2	$(r1(s1,e1))\&(!r2(s2,e1))\&(r3(e1,f))$
10	pin				✓	2	$(r1(s1,e1))\&(r2(e1,f))\&(!r3(s2,f))$
11	pni				✓	2	$(r1(s1,e1))\&(!r2(e1,f))\&(r3(s2,f))$
12	2u					1	$(r1(s1,f))\ (r2(s2,f))$
13	up					2	$((r1(s1,e1))\ (r2(s2,e1)))\&(r3(e1,f))$
14	2m	✓				2	$(r1(s1,e1))\&(r2(e1,f))\&(r3(e1,f))$
15	2nm	✓		✓		2	$(r1(s1,e1))\&(r2(e1,f))\&(!r3(e1,f))$
16	3mp	✓				3	$(r1(s1,e1))\&(r2(e1,e2))\&(r3(e2,f))\&(r4(e1,e2))$
17	3pm	✓				3	$(r1(s1,e1))\&(r2(e1,e2))\&(r3(e2,f))\&(r4(e2,f))$
18	im	✓				2	$(r1(s1,e1))\&(r2(s2,e1))\&(r3(e1,f))\&(r4(e1,f))$
19	2il			✓		1	$(r1(s1,f))\&(r2(e1,f))$
20	3il			✓		1	$(r1(s1,f))\&(r2(s2,f))\&(r3(e1,f))$
21	3c		✓			3	$(r1(s1,e1))\&(r2(e1,f))\&(r3(s2,e2))\&(r4(e2,f))\&(r5(e1,e2))$
22	3cm	✓	✓			3	$(r1(s1,e1))\&(r2(e1,f))\&(r3(s2,e2))\&(r4(e2,f))\&(r5(e1,e2))\&(r6(e1,f))$

Table 11: Details of 23 seen query types with their type name, query-level features: **mul**:multiple relation projection edges, **cyc**:cycles, **exi**:existentially quantified variables, **neg**:negation, **depth**:longest relation projection chain (reasoning depth), and their logical formula.

id	type name	mul.	cyc.	exi.	neg.	depth	logical formula
23	2pi					2	$(r1(s1,e1))\&(r2(e1,f))\&(r3(s2,e2))\&(r4(e2,f))$
24	2pu					2	$((r1(s1,e1))\&(r3(e1,f))\ (r2(s2,e2))\&(r4(e2,f)))$
25	ui					2	$((r1(s1,e1))\ (r2(s2,e1)))\&(r3(e1,f))\&(r4(s3,f))$
26	iu					2	$((r1(s1,e1))\&(r2(s2,e1))\&(r3(e1,f))\ (r4(s3,f)))$
27	upi					2	$((r1(s1,e1))\ (r2(s2,e1)))\&(r3(e1,f))\&(r4(s3,e2))\&(r5(e2,f))$
28	ipu					2	$((r1(s1,e1))\&(r2(s2,e1))\&(r3(e1,f))\ (r4(s3,e2))\&(r5(e2,f)))$
29	i2p					3	$(r1(s1,e1))\&(r2(s2,e1))\&(r3(e1,e2))\&(r4(e2,f))$
30	u2p					3	$((r1(s1,e1))\ (r2(s2,e1)))\&(r3(e1,e2))\&(r4(e2,f))$
31	2pin				✓	2	$(r1(s1,e1))\&(r2(e1,f))\&(r3(s2,e2))\&(r4(e2,f))\&(!r5(s3,f))$
32	2pni				✓	2	$(r1(s1,e1))\&(r2(e1,f))\&(r3(s2,e2))\&(!r4(e2,f))$
33	pn3i				✓	2	$(r1(s1,e1))\&(!r2(e1,f))\&(r3(s2,f))\&(r4(s3,f))$
34	in2p				✓	3	$(r1(s1,e1))\&(!r2(s2,e1))\&(r3(e1,e2))\&(r4(e2,f))$
35	inu				✓	2	$((r1(s1,e1))\&(!r2(s2,e1)))\&(r3(e1,f))\ (r4(s3,f))$
36	inpu				✓	2	$((r1(s1,e1))\&(!r2(s2,e1)))\&(r3(e1,f))\ (r4(s3,e2))\&(r5(e2,f))$
37	upni				✓	2	$((r1(s1,e1))\ (r2(s2,e1)))\&(r3(e1,f))\&(r4(s3,e2))\&(!r5(e2,f))$
38	unpi				✓	2	$((r1(s1,e1))\ (r2(s2,e1)))\&(!r3(e1,f))\&(r4(s3,e2))\&(r5(e2,f))$
39	imp	✓				3	$(r1(s1,e1))\&(r2(s2,e1))\&(r3(e1,e2))\&(r4(e2,f))\&(r5(e1,e2))$
40	ipm	✓				3	$(r1(s1,e1))\&(r2(s2,e1))\&(r3(e1,e2))\&(r4(e2,f))\&(r5(e2,f))$
41	3im	✓				2	$(r1(s1,e1))\&(r2(s2,e1))\&(r3(s3,e1))\&(r4(e1,f))\&(r5(e1,f))$
42	pil			✓		2	$(r1(s1,e1))\&(r2(e1,f))\&(r3(e2,f))$
43	ilp			✓		2	$(r1(s1,e1))\&(r2(e2,e1))\&(r3(e1,f))$
44	p3il			✓		2	$(r1(s1,e1))\&(r2(e1,f))\&(r3(s2,e2))\&(r4(e2,f))\&(r5(e3,f))$
45	i3c		✓			3	$(r1(s1,e1))\&(r2(s2,e1))\&(r3(e1,f))\&(r4(s3,e2))\&(r5(e2,f))\&(r6(e1,e2))$
46	i3cm	✓	✓			3	$(r1(s1,e1))\&(r2(s2,e1))\&(r3(e1,f))\&(r4(s3,e2))\&(r5(e2,f))\&(r6(e1,e2))\&(r7(e1,f))$
47	3inl			✓	✓	1	$(r1(s1,f))\&(!r2(s2,f))\&(r3(e1,f))$
48	pinl			✓	✓	2	$(r1(s1,e1))\&(r2(e1,f))\&(!r3(s2,f))\&(r4(e2,f))$
49	innm	✓			✓	2	$(r1(s1,e1))\&(!r2(s2,e1))\&(r3(e1,f))\&(r4(e1,f))$
50	innmp	✓			✓	3	$(r1(s1,e1))\&(!r2(s2,e1))\&(r3(e1,e2))\&(r4(e2,f))\&(r5(e1,e2))$
51	innpm	✓			✓	3	$(r1(s1,e1))\&(!r2(s2,e1))\&(r3(e1,f))\&(r4(e2,f))\&(r5(e2,f))$
52	3nmp	✓			✓	3	$(r1(s1,e1))\&(r2(e1,e2))\&(r3(e2,f))\&(!r4(e1,e2))$
53	3cn		✓		✓	3	$(r1(s1,e1))\&(r2(e1,f))\&(r3(s2,e2))\&(!r4(e2,f))\&(r5(e1,e2))$
54	3cnm	✓	✓		✓	3	$(r1(s1,e1))\&(r2(e1,f))\&(r3(s2,e2))\&(r4(e2,f))\&(r5(e1,e2))\&(!r6(e1,f))$

Table 12: Details of 32 unseen query types with their type name, query-level features: **mul**:multiple relation projection edges, **cyc**:cycles, **exi**:existentially quantified variables, **neg**:negation, **depth**:longest relation projection chain (reasoning depth), and their logical formula.

Seen Query Types (13 types)		
Type Name	EFO Syntax	Lisp-Like Syntax
1p	r1(s1,f)	(p,(r1),(s1))
2p	(r1(s1,e1))&(r2(e1,f))	(p,(r2),(p,(r1),(s1)))
3p	(r1(s1,e1))&(r2(e1,e2))&(r3(e2,f))	(p,(r3),(p,(r2),(p,(r1),(s1))))
2i	(r1(s1,f))&(r2(s2,f))	(i,(p,(r1),(s1)),(p,(r2),(s2)))
3i	(r1(s1,f))&(r2(s2,f))&(r3(s3,f))	(i,(p,(r1),(s1)),(p,(r2),(s2)),(p,(r3),(s3)))
ip	(r1(s1,e1))&(r2(s2,e1))&(r3(e1,f))	(p,(r3),(i,(p,(r1),(s1)),(p,(r2),(s2))))
pi	(r1(s1,e1))&(r2(e1,f))&(r3(s2,f))	(i,(p,(r2),(p,(r1),(s1))),((p,(r3),(s2))))
2in	(r1(s1,f))&(!r2(s2,f))	(i,(p,(r1),(s1)),(n,(p,(r2),(s2))))
3in	(r1(s1,f))&(r2(s2,f))&(!r3(s3,f))	(i,(p,(r1),(s1)),(p,(r2),(s2)),(n,(p,(r3),(s3))))
inp	(r1(s1,e1))&(!r2(s2,e1))&(r3(e1,f))	(p,(r3),(i,(p,(r1),(s1)),(n,(p,(r2),(s2))))
pin	(r1(s1,e1))&(r2(e1,f))&(!r3(s2,f))	(i,(p,(r2),(p,(r1),(s1))),((n,(p,(r3),(s2))))
2u	(r1(s1,f)) (r2(s2,f))	(u,(p,(r1),(s1)),(p,(r2),(s2)))
up	((r1(s1,e1)) (r2(s2,e1)))&(r3(e1,f))	(p,(r3),(u,(p,(r1),(s1)),(p,(r2),(s2))))

Unseen Query Types (12 types)		
Type Name	EFO Syntax	Lisp-Like Syntax
2pi	(r1(s1,e1))&(r2(e1,f))&(r3(s2,e2))&(r4(e2,f))	(i,(p,(r2),(p,(r1),(s1))),((p,(r4),(p,(r3),(s2))))
2pu	((r1(s1,e1))&(r3(e1,f)) (r2(s2,e2)))&(r4(e2,f))	(u,(p,(r3),(p,(r1),(s1))),((p,(r4),(p,(r2),(s2))))
ui	((r1(s1,e1)) (r2(s2,e1)))&(r3(e1,f))&(r4(s3,f))	(i,(p,(r4),(s3)),(p,(r3),(u,(p,(r1),(s1)),(p,(r2),(s2))))
iu	((r1(s1,e1))&(r2(s2,e1))&(r3(e1,f)) (r4(s3,f))	(u,(p,(r4),(s3)),(p,(r3),(i,(p,(r1),(s1)),(p,(r2),(s2))))
upi	((r1(s1,e1)) (r2(s2,e1)))&(r3(e1,f))&(r4(s3,e2))&(r5(e2,f))	(i,(p,(r5),(p,(r4),(s3))),((p,(r3),(u,(p,(r1),(s1)),(p,(r2),(s2))))
ipu	((r1(s1,e1))&(r2(s2,e1))&(r3(e1,f)) (r4(s3,e2))&(r5(e2,f))	(u,(p,(r5),(p,(r4),(s3))),((p,(r3),(i,(p,(r1),(s1)),(p,(r2),(s2))))
i2p	(r1(s1,e1))&(r2(s2,e1))&(r3(e1,e2))&(r4(e2,f))	(p,(r4),(p,(r3),(i,(p,(r1),(s1)),(p,(r2),(s2))))
u2p	((r1(s1,e1)) (r2(s2,e1)))&(r3(e1,e2))&(r4(e2,f))	(p,(r4),(p,(r3),(u,(p,(r1),(s1)),(p,(r2),(s2))))
2pin	(r1(s1,e1))&(r2(e1,f))&(r3(s2,e2))&(r4(e2,f))&(!r5(s3,f))	(i,(i,(p,(r4),(p,(r3),(s2))),((p,(r2),(p,(r1),(s1))),((n,(p,(r5),(s3))))
in2p	(r1(s1,e1))&(!r2(s2,e1))&(r3(e1,e2))&(r4(e2,f))	(p,(r4),(p,(r3),(i,(p,(r1),(s1)),(n,(p,(r2),(s2))))
inu	((r1(s1,e1))&(!r2(s2,e1)))&(r3(e1,f)) (r4(s3,f))	(u,(p,(r4),(s3)),(p,(r3),(i,(p,(r1),(s1)),(n,(p,(r2),(s2))))
inpu	((r1(s1,e1))&(!r2(s2,e1))&(r3(e1,f)) (r4(s3,e2))&(r5(e2,f))	(u,(p,(r5),(p,(r4),(s3))),((p,(r3),(i,(p,(r1),(s1)),(n,(p,(r2),(s2))))

Table 13: Details of 25 query types (compatible in both Lisp-like and EFO syntax) that we used as a subset for exploration on query syntax.

E Knowledge Graph Statistics

Detailed statistics of the knowledge graphs selected are presented in Table 14.

Dataset	Relations	Entities	Training	Validation	Testing	All Edges
FB15k	1,345	14,951	483,142	50,000	59,071	592,213
FB15k-237	237	14,505	272,115	17,526	20,438	310,079
NELL995	200	63,361	114,213	14,324	14,267	142,804

Table 14: Details of three knowledge graphs used for the experiments, and their separation standard for training, validation, and testing edges according to (Ren and Leskovec, 2020).

F Query Graph Definition

We provide query graph for each query types in our dataset. Query graphs can be utilized for graph-augmented methods, and have been applied in the two inductive biases discussed in Sec. 5.3. Here we provide the definition:

For each atomic formula or its negation $\alpha = r(h, t)$ or $\neg r(h, t)$ in a conjunctive formula c , we have $\{(h, r), (r, t)\} \in \mathcal{G}_c$ or $\{(h, r), (r, n), (n, t)\} \in \mathcal{G}_c$, where n denotes the negation node in the conjunctive query graph \mathcal{G}_c . By our definition of query, all conjunctive query graphs have exactly one node as a free variable. For queries that represent the disjunction of multiple conjunctive formulas, we replace all free variable nodes in every conjunctive query graph with a single union node u , and connect it to a final free variable node f , with $\{(u, f)\} \in \mathcal{G}_d$, where \mathcal{G}_d is the disjunctive query graph.

G Evaluation Details

For each knowledge graphs, we separate their edges into training, validation, and testing edges with a ratio of approximately 8:1:1, as shown in Table 14. We construct three graphs, training graph \mathcal{G}_{train} , validation graph \mathcal{G}_{valid} , and testing graph \mathcal{G}_{test} with training edges, training+validation edges, and training+validation+testing edges, respectively.

We adopt the mean reciprocal rank (MRR) as our evaluation metric, following the calculation presented in Equation 9 and 10. It is important to note that for all testing queries, the answers in $A(\mathcal{G}_{valid}, q) - A(\mathcal{G}_{train}, q)$ are excluded from \mathcal{A}_{ood} to ensure fairness. Consequently, we define $\mathcal{A}_{id} = A(\mathcal{G}_{train}, q)$ and $\mathcal{A}_{ood} = A(\mathcal{G}_{test}, q) - A(\mathcal{G}_{valid}, q)$ for our evaluation of ID (K) and OOD (K), respectively. In each table regarding experimental results, the scores below ID (Q) (i.e., either MRR of ID (K) or OOD (K)) represent the average scores among all seen query types. While the scores below OOD (Q) represent the average scores among all unseen query types. To connect with literature in KG query answering (Bai et al., 2023b; Sun et al., 2021), the performances in ID (K) and OOD (K) are also referred to as *faithfulness* and *knowledge inference capability*, respectively.

$$MRR_{id(k)} = \frac{\sum_{v \in \mathcal{A}_{id}} \text{rank}(v)}{|\mathcal{A}_{id}|} \quad (9)$$

$$MRR_{ood(k)} = \frac{\sum_{v \in \mathcal{A}_{ood}} \text{rank}(v)}{|\mathcal{A}_{ood}|} \quad (10)$$

H Transformer Architecture Comparison

We investigate the impact of different transformer architectures on complex query answering performance. Table 15 compares three architectural variants: encoder-only (our main approach), decoder-only, and encoder-decoder configurations. The results demonstrate that the encoder-only architecture with relative positional encoding achieves the best overall performance across datasets.

Dataset	Architecture	PE Type	ID (Q)		OOD (Q)		All Queries	
			ID (K)	OOD (K)	ID (K)	OOD (K)	ID (K)	OOD (K)
FB15k	Encoder-only	Absolute PE	46.9	31.9	21.8	13.2	32.3	21.0
	Encoder-only	Relative PE	48.1	32.3	35.4	21.5	40.7	26.0
	Decoder-only	Absolute PE	47.9	31.9	24.3	14.7	34.2	21.9
	Encoder-Decoder	Relative PE	40.1	28.5	32.0	19.8	35.4	23.4
FB15k-237	Encoder-only	Absolute PE	54.4	20.0	20.8	9.0	34.9	13.6
	Encoder-only	Relative PE	54.8	20.0	37.7	15.8	44.8	17.6
	Decoder-only	Absolute PE	55.2	19.9	21.9	9.8	35.8	14.0
	Encoder-Decoder	Relative PE	44.5	19.9	33.5	15.8	38.1	17.5
NELL995	Encoder-only	Absolute PE	68.8	18.9	25.8	8.3	43.8	12.7
	Encoder-only	Relative PE	68.0	18.7	49.8	13.4	57.4	15.6
	Decoder-only	Absolute PE	65.6	18.6	28.0	8.9	43.7	13.0
	Encoder-Decoder	Relative PE	65.5	19.7	48.7	13.7	55.7	16.2

Table 15: Comparison of different transformer architectures (in MRR%).

I Technical Details

In this section, we describe our experiment setting in more details.

Computational Resource: All transformers are trained on four NVIDIA A100 GPUs for two days, with a batch size of 1024. As shown in Table 16, TEGA achieves computational efficiency comparable to that of the baseline models.

Model Configuration: Table 18 illustrated the impact of transformer size to query answering performances. We selected a three-layer configuration for all Transformer models, for its effectiveness and comparable parameter size to other baseline methods (as shown in Table 17). To train the Transformer models, we employed the label smoothing loss (Szegedy et al., 2015) with a smoothing factor of $\epsilon = 0.1$, which helps to regularize the model and improve generalization. The learning rate was set to 0.0001, and a warm-up schedule of 1000 steps was applied to gradually increase the learning rate during the initial phase of training, allowing the model to adapt to the task and stabilize the optimization process. All Transformer models and baselines were configured with an embedding size of 400 dimensions. Regarding transformer encoder with absolute and relative PE, we employ the BertModel from the transformers library. While for disentangled PE, we use DebertaModel with the same size and configurations. Note that apart from disentangled PE, Deberta also proposed a mask decoder during the pre-training stage. However, as we train all models solely on logical queries without utilizing any pre-trained model weights, the impact of the mask decoder is negligible.

Pooling Strategy: For free-variable pooling, we experimented with two pooling settings: sum pooling and mean pooling, with results presented in Table 19. The results show that sum pooling over free variables is the better setting, which is also consistent with findings in GNN literature (Xu et al., 2019).

Models	Inference Time (ms)
Transformer (Absolute PE)	43.2
Transformer (Disentangled PE)	47.3
Transformer (Relative PE)	44.7
+Adjacency Matrices Masking	46.8
+Direct Distance Encoding	44.9
TEGA (ours)	52.1

Table 16: Computational efficiency of TEGA compared with baselines.

Model	Number of Trainable Parameters (in Millions)
BetaE	13.7
ConE	18.9
LMPNN	11.5
SQE-LSTM	8.8
Trans.+Relative PE	14.7
TEGA (ours)	15.7

Table 17: Parameter size comparison between different query answering approaches, with three layers of self-attention in transformer models.

Model	# Attn. layers	ID(Q)		OOD(Q)		All Queries	
		ID(K)	OOD(K)	ID(K)	OOD(K)	ID(K)	OOD(K)
Trans.+Absolute PE	3	54.4	20.0	20.8	9.0	34.9	13.6
	6	57.2	20.3	23.0	9.6	37.3	14.1
	12	48.9	19.8	24.4	11.4	34.6	14.9
Trans.+Relative PE	3	54.8	20.0	37.7	15.8	44.8	17.6
	6	56.7	20.0	38.0	16.0	45.8	17.7
	12	48.6	19.9	35.0	15.7	40.7	17.4

Table 18: Experimental results on transformers with different numbers of attention layers in the FB15K-237 dataset.

Dataset	Pooling	ID (Q)		OOD (Q)		All Queries	
		ID (K)	OOD (K)	ID (K)	OOD (K)	ID (K)	OOD (K)
FB15k	Mean	52.9	33.5	37.2	21.6	43.7	26.6
	Sum	55.0	34.1	38.8	22.4	45.6	27.3
FB15k-237	Mean	62.9	20.0	41.6	15.6	50.5	17.4
	Sum	64.3	20.1	42.6	16.0	51.7	17.7

Table 19: Performance comparison between sum pooling and mean pooling strategies on TEGA.