

Query2Triple: Unified Query Encoding for Answering Diverse Complex Queries over Knowledge Graphs

Yao Xu^{1,2}, Shizhu He^{1,2,*}, Cunguang Wang³, Li Cai³, Kang Liu^{1,2,4}, Jun Zhao^{1,2}

¹ The Laboratory of Cognition and Decision Intelligence for Complex Systems, Institute of Automation, Chinese Academy of Sciences, Beijing, China

² School of Artificial Intelligence, University of Chinese Academy of Sciences, Beijing, China

³ Meituan, Beijing, China

⁴ Shanghai Artificial Intelligence Laboratory, Shanghai, China

{yao.xu, shizhu.he, kliu, jzhao}@nlpr.ia.ac.cn, {wangcunguang, caili03}@meituan.com

Abstract

Complex Query Answering (CQA) is a challenge task of Knowledge Graph (KG). Due to the incompleteness of KGs, query embedding (QE) methods have been proposed to encode queries and entities into the same embedding space, and treat logical operators as neural set operators to obtain answers. However, these methods train KG embeddings and neural set operators concurrently on both simple (one-hop) and complex (multi-hop and logical) queries, which causes performance degradation on simple queries and low training efficiency. In this paper, we propose Query to Triple (Q2T), a novel approach that decouples the training for simple and complex queries. Q2T divides the training into two stages: (1) Pre-training a neural link predictor on simple queries to predict tail entities based on the head entity and relation. (2) Training a query encoder on complex queries to encode diverse complex queries into a unified triple form that can be efficiently solved by the pretrained neural link predictor. Our proposed Q2T is not only efficient to train, but also modular, thus easily adaptable to various neural link predictors that have been studied well. Extensive experiments demonstrate that, even without explicit modeling for neural set operators, Q2T still achieves state-of-the-art performance on diverse complex queries over three public benchmarks.

1 Introduction

Knowledge Graphs (KGs) organize world knowledge as inter-linked triples which describe entities and their relationships in symbolic form (Ji et al., 2020). Complex query answering (CQA), a knowledge reasoning task over KGs, has been proposed in recent years (Wang et al., 2021). Compared with simple link prediction, which involves predicting a missing entity in a factual triple (Rossi et al., 2021), CQA is more challenging as it requires first-order

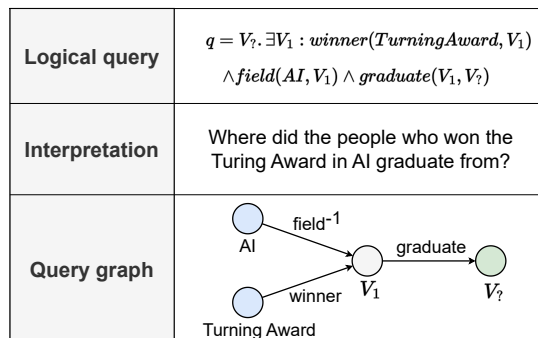


Figure 1: An example, along with its corresponding interpretation and query graph, in CQA.

logic (FOL) operators such as existential quantification (\exists), conjunction (\wedge), disjunction (\vee), and negation (\neg) to be performed on entities, relations and triples, as shown in Figure 1. As KGs usually suffer from incompleteness (West et al., 2014), traditional query methods such as SPARQL (Prudhommeaux et al., 2013) cannot handle CQA well on real-world KGs such as Freebase (Bollacker et al., 2008) and NELL (Carlson et al., 2010).

Recently, an alternative method called Query Embedding (QE) (Hamilton et al., 2018; Ren et al., 2020) has been proposed to complete the missing edges and answer the complex query simultaneously (Wang et al., 2023b). The idea of these methods is to encode queries and entities into the same embedding space, and treats logical operators as neural set operators (e.g. relation projection as set projection, queries conjunction as sets intersection, query negation and set complement) to answer queries. QE methods embed a query by iteratively performing neural set operators according to the topological order of nodes in the query graph, and then obtain answers based on the similarity scores between the query embedding and all entity embeddings. This process is shown in Figure 3 (B).

Despite QE-based methods achieving good performance on CQA, they still exhibit the following drawbacks: (1) **Cannot perform well on both sim-**

*Corresponding Author

ple and complex queries simultaneously (simple queries refer to one-hop queries, complex queries encompass multi-hop and logical queries). One potential reason from the model perspective is that, QE-based models model the projection operator with complex neural networks to perform well on complex queries, however, these complex neural networks are incapable of effectively managing inherent relational properties (Wang et al., 2023b), such as symmetry, asymmetry, inversion, composition, which are sufficiently studied in KG completion tasks and addressed by Knowledge Graph Embedding (KGE) (Wang et al., 2017), so KGE models can easily outperform QE-based methods on simple queries. Another potential reason from the training perspective is that, the training for simple queries and complex queries are coupled in these methods, as demonstrated in Figure 2 (A). Therefore, the entity and relation embeddings inevitably incorporate neural set operators information that is irrelevant for answering simple queries, leading to a decrease in performance. (2) **Low training efficiency and not modular.** From the model perspective, as the KGE is coupled with the neural set operators in QE-based methods, as shown in Figure 2 (A), these methods usually train KGE and neural set operators from scratch in each round of training, which is inefficient. Furthermore, due to this coupling, these methods cannot directly exploit other high-performance KGE models such as ComplEx (Trouillon et al., 2016). (3) **Suffer from error cascading.** These methods have to calculate all intermediate node representations step-by-step, as demonstrated in Figure 3 (B), so errors will be cascaded along the path (Guu et al., 2015), which affects the quality of answers.

Aiming to handle the above drawbacks, in this paper, we propose Query to Triple (Q2T), a novel approach that not only decouples the training of complex queries and simple queries from the training perspective, but also decouples KGE and the query encoder from the model perspective. Q2T divides the training into two stage: pre-training KGE for simple queries and training a unified query encoder for diverse complex queries, as shown in Figure 2 (B). In the first pre-training stage, we only train a KGE model as a neural link predictor on simple queries, such as DistMult (Yang et al., 2015) and ComplEx (Trouillon et al., 2016). Motivated by the prompt (Sun et al., 2022b; Liu et al., 2023) technology in Natural Languages Process (NLP),

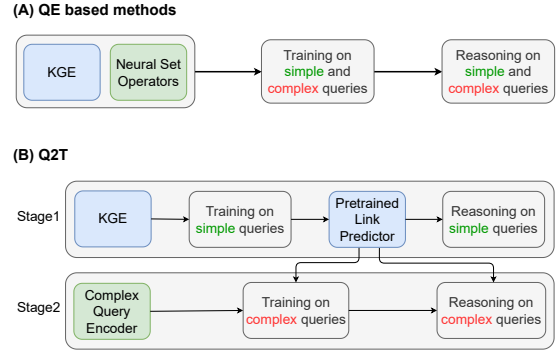


Figure 2: The different training strategies used in QE-based methods and Q2T. KGE denotes entity and relation embeddings. (A) QE-based methods train KGE and neural set operators on both simple and complex queries. (B) Q2T train KGE and Query Encoder on simple and complex queries respectively.

which aims to fully reuse the pretrained model, in the second stage, Q2T utilizes a trainable encoder to transform any complex query graph g into a unified triple form $(g_h, g_r, t_?)$, where g_h, g_r are representations of the whole query graph generated by the encoder, $t_?$ are tail entities to be predicted. Then, we can obtain probability list of all entities by feeding these two representations and all pre-trained entity embeddings to the pretrained neural link predictor. That is, we transform any complex query to a unified triple form which can be handled well by neural link predictors. In other words, as shown in Figure 3 (A), **Q2T treats CQA as a special link prediction task where it learns how to generate appropriate inputs for the pretrained neural link predictor** instead of reasoning in the embedding space step-by-step.

The advantages of Q2T are as follows: (1) Good performance on both simple and complex queries, because Q2T can not only make full use of the well-studies KGE models but also decouple the the training of simple queries and complex queries from the training perspective. (2) High training efficiency and modular, because from the model perspective, Q2T decouples the KGE and the query encoder, so it only trains the encoder in each round of training (pretrained KGE is frozen) and can also fully utilize various advanced KGE models. (3) Simple and accurate, as Q2T utilizes an end-to-end reasoning mode, which not only eliminates the need to model complex neural set operators, but also avoids error cascading.

We conducted experiments on widely-used benchmark such as FB15k (Bordes et al., 2013),

FB15k-237 (Toutanova and Chen, 2015), and NELL995 (Xiong et al., 2017). The experimental results demonstrate that, even without explicit modeling for neural set operators, Q2T still achieves state-of-the-art performance on diverse complex queries. The source codes and data can be found at <https://github.com/YaoXu/Q2T>.

2 Related work

Knowledge Graph Embedding (KGE). KGE embeds components of a KG including entities and relations into continuous vector spaces, so as to predict unseen relational triples while preserving the inherent structure of the KG (Wang et al., 2017). According to the different scoring functions, KGE can be roughly categorized into two types: translational distance models (Bordes et al., 2013; Wang et al., 2014) and semantic matching models (Yang et al., 2015; Trouillon et al., 2016). All these models can be used in Q2T as neural link predictors.

Prompt-based Learning. In recent years, a new training paradigm called "pre-train, prompt, and predict" has achieved success in both NLP (Liu et al., 2023) and graphs (Sun et al., 2022a). Aiming to effectively reuse the pretrained model, this paradigm targets reformulating the downstream task looks similar to pre-training task (Liu et al., 2023). Based on the same idea, in Q2T, we treat simple and complex queries as the pre-training and downstream tasks respectively. The difference is that we use a trainable encoder to encode complex queries to the same triple form as simple queries.

Neural complex query answering. Neural complex query answering can be divided into two categories: with/without pretrained KGE. Query Embedding (QE) based models, without pretrained KGE, encode queries and entities into the same embedding space, then obtain answers by step-by-step reasoning, as shown in Figure 3 (B). In QE-based models, many methods have been proposed to represent entity sets, such as geometric shapes (Ren et al., 2020; Bai et al., 2022), probability distribution (Ren and Leskovec, 2020) and bounded histogram on uniform grids (Wang et al., 2023a). Neural networks like multi-layer perceptron and transformers (Vaswani et al., 2017) are used to model set operations. They are trained for simple and complex queries concurrently, so they suffer drawbacks mentioned before.

Models with pretrained KGE are as follows: (1) CQD (Arakelyan et al., 2021) uses the pretrained

neural link predictor as the one-hop reasoner and T-norms as logic operators to obtain the continuous truth value of an EPFO query. Then the embeddings are optimized to maximize the continuous truth value. However, this method is time consuming and cannot handle negation queries well (Wang et al., 2023b). (2) LMPNN (Wang et al., 2023b), which is related to our work, also utilizes pretrained KGE to conduct one-hop inferences on atomic formulas, the generated results are regarded as the messages passed in Graph Neural Network (GNN) (Wu et al., 2020). However, in the LMPNN, KGE and GNN are coupled, which means the KGE function is invoked during every message propagation on each edge in the inference. Besides, LMPNN has to design of distinct logical message propagation functions for different types of score functions in KGE, as well as for different types of regularization optimization. Compared to LMPNN, Q2T only requires a single invocation of the KGE score function for any type of complex query. Therefore, Q2T is better suited for leveraging KGE with higher complexity and larger parameters. Additionally, Q2T can be employed with any KGE without the need for additional code modifications, so it is more modular.

SQE (Bai et al., 2023) is also related to our work, it uses a search-based algorithm to linearize the computational graph to a sequence of tokens and then uses a sequence encoder to compute its vector representation. Another related work is KgTransformer (Liu et al., 2022). It also uses the Transformer to encode queries, but it outputs the final node embedding directly. To improve transferability and generalizability, KgTransformer introduces a two-stage masked pre-training strategy, which requires more pre-training data and time. However, Q2T can outperform KgTransformer easily with much less parameters and training time.

3 Preliminary

In this section, we formally introduce knowledge graph (KG). Then, we use symbols of KG to define EFO-1 query and its corresponding Disjunctive Normal Form (DNF). Finally, we introduce neural link predictors.

3.1 Knowledge Graph and EFO-1 Queries

We represent a KG as a set of factual triples, i.e., $G = \{(h, r, t) \in \mathcal{V} \times \mathcal{R} \times \mathcal{V}\}$, where $h, r \in \mathcal{V}$ denote the head and tail entity, $r \in \mathcal{R}$ represents

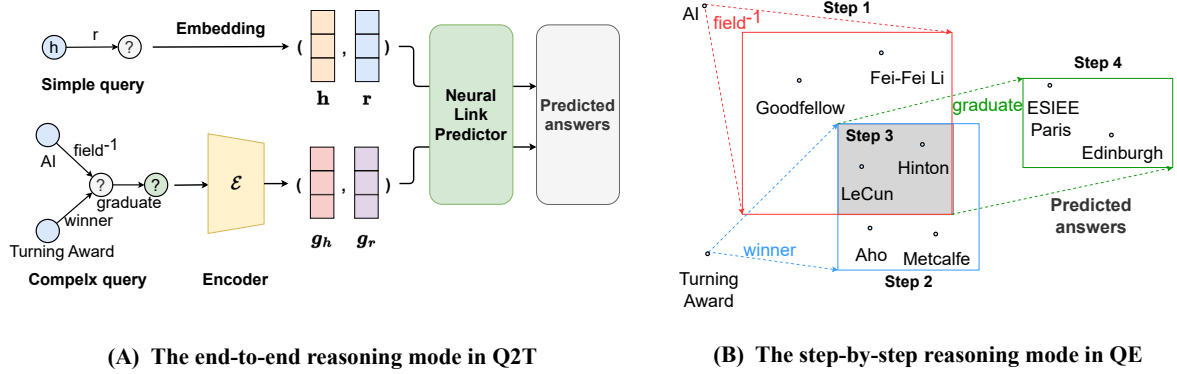


Figure 3: The different reasoning mode for the query "Where did the person who won the Turing Award in AI graduate from?". (A) Q2T adopts end-to-end reasoning mode: encoding the query graph to obtain two representation g_h, g_r , then the pretrained neural link predictor takes g_h, g_r as inputs and output predicted tail entities, which is similar to answering simple queries. (B) QE-based methods utilize step-by-step reasoning mode: iteratively apply the neural set operators to get answers. Step 1 and Step 2 denote applying set projection from *AI* and *Turning Award* respectively. Step 3 denotes applying set intersection to the results of Step 1 and Step 2. Finally, Step 4 denotes applying set projection to the result obtained in Step 3.

the relation. We make $r(h, r) = 1$ if and only if there is a directed relation r between h and t .

Following previous work (Ren and Leskovec, 2020), we define the Existential First Order queries with a single free variable (EFO-1), and an arbitrary EFO-1 logical formula can be converted to the Disjunctive Normal Form (DNF) as follows:

$$q[V?] = V?. \exists V_1, \dots, \exists V_n : c_1 \vee \dots \vee c_m \quad (1)$$

where $V?$ is the only free variable and $V_i, 1 \leq i \leq n$ are n existential variables, each $c_i, 1 \leq i \leq m$ represents a conjunctive clause like $c_i = e_{i1} \wedge e_{i2} \wedge \dots \wedge e_{ik_i}$, and each $e_{ij}, 1 \leq j \leq k_i$ indicates an atomic formula or its negation, i.e., $e_{ij} = r(a, b)$ or $\neg r(a, b)$, where $r \in \mathcal{R}$, a, b can be either a constant $e \in \mathcal{V}$ or a variable V .

3.2 Neural Link Predictor

KGE is a kind of neural link predictor, for simplicity, neural link predictors mentioned in this article refer to KGE models. Given the head entity embedding \mathbf{h} , relation embedding \mathbf{r} , and tail entity embedding \mathbf{t} , KGE can compute the score $\phi(\mathbf{h}, \mathbf{r}, \mathbf{t})$ of the triple (h, r, t) , where $\phi(\mathbf{h}, \mathbf{r}, \mathbf{t})$ indicates the likelihood that entities h and t hold the relation r .

4 Q2T: Transforming Query to Triple

There are two main components in Q2T: (1) The Neural Link Predictor. (2) The Query Encoder. This section begins with an introduction to the pre-training of neural link predictors. Subsequently, we present QueryGraphormer, a Transformer-based

model that incorporates structural information, as a query encoder. Q2T utilizes the query encoder to encode complex queries, producing representations that are used as inputs in neural link predictors, as depicted in Figure 4.

4.1 The Pretrained Neural Link Predictor

As answers of complex queries are generated by the pretrained neural link predictor finally, its performance directly affects the final performance.

Aiming to get neural link predictors with good performance, the training object we use contains not only terms for predicting the tail entity of a given triple, but also a term for predicting the relation type, which has been proved to significantly improve entity ranking (Chen et al., 2021). The training object is defined as follows:

$$\begin{aligned} & \arg \max_{\theta \in \Theta} \sum_{(h, r, t) \in \mathcal{G}} [\log P_{\theta}(t|h, r) + \lambda \log P_{\theta}(r|h, t)] \\ & \log P_{\theta}(t|h, r) = \phi(\mathbf{h}, \mathbf{r}, \mathbf{t}) - \log \sum_{t' \in \mathcal{E}} \exp[\phi(\mathbf{h}, \mathbf{r}, \mathbf{t}')] \\ & \log P_{\theta}(r|h, t) = \phi(\mathbf{h}, \mathbf{r}, \mathbf{t}) - \log \sum_{r' \in \mathcal{R}} \exp[\phi(\mathbf{h}, \mathbf{r}', \mathbf{t})] \end{aligned} \quad (2)$$

where $\theta \in \Theta$ are the model parameters, including entity and relation embeddings, h, r, t denote head entity, relation and tail entity, $\mathbf{h}, \mathbf{r}, \mathbf{t}$ denote their corresponding embeddings under θ , ϕ is a scoring function, λ is a hyper-parameter determining the contribution of the relation prediction objective.

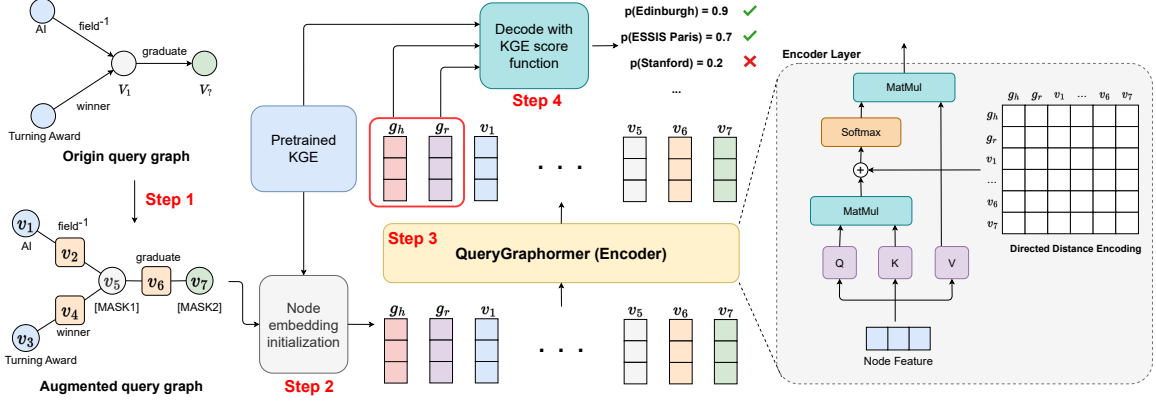


Figure 4: The framework of Q2T. The steps of answering complex queries in Q2T are: (1) Transforming the origin query graph to a augmented graph without edge attributes. (2) Flattening the augmented graph to a sequence, and initializing the sequence with the pretrained KGE. (3) Encoding the sequence representations with QueryGraphormer which considers the structural information of the graph. (4) Inputting the final representations of g_h , g_r , and other entity representations to the score function of the pretrained KGE to calculate the probability list of entities.

4.2 The QueryGraphormer

Transformer for CQA With the great success of Transformer (Vaswani et al., 2017) in NLP, Transformer has also been proved to perform well for graph representation learning (Ying et al., 2021). KgTransformer (Liu et al., 2022) is the first work to apply Transformer on query graphs, which first turns relations into relation-nodes to transform the query graph into an augmented graph without edge attributes, then masks adjacency matrices (Hu et al., 2020) to self-attention. However, this limits receptive field of each node, which means each node can only obtain information from its neighbors.

Motivated by Graphormer (Ying et al., 2021), we propose QueryGraphormer, which uses the directed distance between any two nodes as a bias term in the attention mechanism, to solve these drawbacks. This helps the model effectively capture spatial dependencies within the augmented graph (obtained through the same transformation as KgTransformer).

The biggest difference between our proposed QueryGraphormer and KgTransformer is that outputs of QueryGraphormer are inputted to pretrained neural link predictors to get the final predicted answers, while KgTransformer outputs the final predicted answer embedding directly.

4.2.1 Input Representations Construction

In this part, we introduce how to construct input sequence representations for QueryGraphormer.

Initial input representations. Similar to BERT (Devlin et al., 2019) model which introduces a spe-

cial token $[CLS]$ to represent the sentence-level feature, in QueryGraphormer, we add two special node: $[g_h]$ and $[g_r]$, which represent a virtual head node and a virtual relation node respectively. In the view of graph, $[g_h]$ and $[g_r]$ can be regraded as two virtual nodes connected to all nodes in the query graph. We attach $[g_h]$ and $[g_r]$ at the beginning of each sequence, then a query graph is flattened into a sequence, $S = [g_h, g_r, v_1, \dots, v_n]$, where n is the number of nodes in the augmented graph, $n_i, i \leq n$ is an entity node, relation node or MASK node. For each node, its input representation is a frozen embedding from the pretrained KGE if it is an entity node or relation node, or a trainable embedding otherwise. We denote the initial representation of S as $S^0 \in \mathbb{R}^{m \times d_0}$, where $m = n + 2$, d_0 is the dimension of the pretrained KGE.

Projection. After obtaining the initial representations of S , we use MLP to project features from high-dimensional to low-dimensional to reduce the parameter amount and calculation time of the QueryGraphormer, which is described as follows:

$$S^1 = MLP_{proj}(S^0) \quad (3)$$

where $S^1 \in \mathbb{R}^{m \times d_1}$ and $S^0 \in \mathbb{R}^{m \times d_0}$ are the initial/projected representations of S , and $d_1 < d_0$.

Handle negation queries. Aiming to make the input representations of negation queries different from positive queries. We apply a linear transformation on node representation that represents a negation relation: $h_i^{neg} = Ah_i$, where $A \in \mathbb{R}^{d_1 \times d_1}$, $h_i \in \mathbb{R}^{d_1}$ denotes the the i -th element (negation relation node) in S^1 .

4.2.2 Self-Attention with Distance Message

Directed Distance Encoding. Similar to the *Spatial Encoding* used in Graphormer (Ying et al., 2021), we employ a function $\phi(v_i, v_j) : V \times V \rightarrow \mathbb{R}$ which measures the spatial relation between v_i and v_j in the query graph g . Unlike Graphormer which uses the distance of the shortest path to be the function, in this paper, we use the directed distance of the shortest path, because the topological order is crucial in query graphs, that is, predecessors and successors of each node should be treated differently.

The directed distance of the shortest path is defined as follows:

$$\phi(v_i, v_j) = \text{dir} * \text{spd}(v_i, v_j) \quad (4)$$

$$\text{dir} = \text{sgn}(\text{layer}(v_i) - \text{layer}(v_j)) \quad (5)$$

where $\text{spd}(v_i, v_j)$ denotes the distance of the shortest path between v_i and v_j , $\text{layer}(v_j)$ denotes the layer number of v_j (obtained by topological sort), $\text{sgn}(x)$ is the sign function: $\text{sgn}(x) = 1$ if $x \geq 0$, $\text{sgn}(x) = 0$ if $x < 0$.

For each output value of ϕ , we allocate a trainable scalar to be a bias term in the self-attention. Then the attention α_{ij} between v_i and v_j is computed in the following way (for simplicity, we only consider single-head self-attention here):

$$\alpha_{ij} = \frac{\exp a_{ij}}{\sum_{k=1}^m \exp a_{ik}} \quad (6)$$

$$a_{ij} = \frac{(\mathbf{h}_i \mathbf{W}^Q)(\mathbf{h}_j \mathbf{W}^K)^T}{\sqrt{d_1}} + b_{\psi(v_i, v_j)} \quad (7)$$

where $\mathbf{W}^Q, \mathbf{W}^K \in \mathbb{R}^{d_1 \times d_1}$ are Query and Key matrices, $b_{\psi(v_i, v_j)}$ is a learnable scalar indexed by $\psi(v_i, v_j)$. By doing this, each node in the query graph can not only aggregate the information of all other nodes, but also decide which nodes should pay more attention to based on distance messages.

The result of self-attention for v_i is obtained as follows:

$$\text{Attn}_i = \sum_{j=1}^n \alpha_{ij} (\mathbf{h}_j \mathbf{W}^V) \quad (8)$$

where $\mathbf{W}^V \in \mathbb{R}^{d_1 \times d_1}$ is Value matrix.

Then we adopt the classic Transformer encoder to get the output representations of layer l , which is described as follows:

$$\mathbf{h}'^{(l)} = \text{LN}(\text{MHA}(\mathbf{h}^{(l-1)}) + \mathbf{h}^{(l-1)}) \quad (9)$$

$$\mathbf{h}^{(l)} = \text{LN}(\text{FFN}(\mathbf{h}'^{(l)}) + \mathbf{h}'^{(l)}) \quad (10)$$

$$\text{FFN}(\mathbf{x}) = \text{act}(\mathbf{x} \mathbf{W}_1 + \mathbf{b}_1) \mathbf{W}_2 + \mathbf{b}_2 \quad (11)$$

where MHA denotes Multi-Head Attention described before, FFN denotes Feed-Forward blocks, LN denotes layer normalization, $\mathbf{W}_1 \in \mathbb{R}^{d_1 \times d_2}$, $\mathbf{W}_2 \in \mathbb{R}^{d_2 \times d_1}$, and act is the activation function. In QueryGraphormer, we make $d_1 = d_2$.

4.2.3 Training Object of QueryGraphormer

Finally, we feed the representations of g_h and g_r in the last layer, denoted as $\mathbf{g}_h, \mathbf{g}_r \in \mathbb{R}^{d_1}$, into the pretrained neural link predictor to obtain predicted answers. Formally, given a query graph g , the score of each tail entity t_i is calculated as follows:

$$s(t_i | g) = \phi(\mathbf{g}_h', \mathbf{g}_r', \mathbf{t}_i) \quad (12)$$

$$\mathbf{g}_h' = \text{MLP}_{rev}(\mathbf{g}_h), \mathbf{g}_r' = \text{MLP}_{rev}(\mathbf{g}_r) \quad (13)$$

where MLP_{rev} projects the representations from d_1 dimension back to d_0 dimension, $\mathbf{t}_i \in \mathbb{R}^{d_0}$ is the pretrained embedding of entity t_i , ϕ is the score function of the pretrained KGE.

It is noteworthy that \mathbf{t}_i is from the pretrained entity embeddings, and they are not needed to updated in the training of QueryGraphormer, which effectively reduces the number of parameters needed to be updated.

Our training object is to maximize the log probabilities of correct answer t , the loss function (without label smoothing for simplicity) is defined as follows:

$$L = -(s(t | g) - \log \sum_{i=1}^K \exp [s(t_i | g)]) \quad (14)$$

where t and t_i are the answer entity and negative entity (random sampling), respectively.

5 Experiments

In this section, we conduct experiments to demonstrate the effectiveness and efficiency of Q2T and the necessity of decoupling the training for simple and complex queries.

5.1 Experiment Setup

Datasets. To compare our results and previous works directly, we conduct our experiments on the widely used datasets generated by Ren and Leskovec. The logical queries in these datasets are generated from FB15k (Bordes et al., 2013), FB15k-237 (Toutanova and Chen, 2015), and NELL995 (Xiong et al., 2017). More details about these queries can be found in Appendix A.

Dataset	Model	1p	2p	3p	2i	3i	pi	ip	2u	up	2in	3in	inp	pin	pni	A_p	A_n	
FB15k	BetaE	65.1	25.7	24.7	55.8	66.5	43.9	28.1	40.1	25.4	14.3	14.7	11.5	6.5	12.4	41.6	11.8	
	Q2P	82.6	30.8	25.5	65.1	74.7	49.5	34.9	32.1	26.2	21.9	20.8	12.5	8.9	17.1	46.8	16.4	
	ConE	73.3	33.8	29.2	64.4	73.7	50.9	35.7	55.7	31.4	17.9	18.7	12.5	9.8	15.1	49.8	14.8	
	(with pretrained KGE)																	
	CQD-CO	89.2	25.6	13.6	77.4	78.3	44.2	33.2	41.7	22.1	-	-	-	-	-	46.9	-	
	LMPNN	85.0	39.3	28.6	68.2	76.5	46.7	43.0	36.7	31.4	29.1	29.4	14.9	10.2	16.4	50.6	20.0	
Q2T	89.4	44.3	33.6	72.1	79.4	55.3	47.7	65.8	38.2	19.5	21.3	16.2	10.2	16.7	58.4	16.8		
FB15k-237	BetaE	39.0	10.9	10.0	28.8	42.5	22.4	12.6	12.4	9.7	5.1	7.9	7.4	3.6	3.4	20.9	5.4	
	Q2P	39.1	11.4	10.1	32.3	47.7	24.0	14.3	87.0	9.1	4.4	9.7	7.5	4.6	3.8	21.9	6.0	
	ConE	41.8	12.8	11.0	32.6	47.3	25.5	14.0	14.5	10.8	5.4	8.6	7.8	4.0	3.6	23.4	5.9	
	FuzzQE	42.2	13.3	10.2	33.0	47.3	26.2	18.9	15.6	10.8	9.7	12.6	7.8	5.8	6.6	24.2	8.5	
	(with pretrained KGE)																	
	CQD-CO	46.7	9.6	6.2	31.2	40.6	23.6	16.0	14.5	8.2	-	-	-	-	-	21.9	-	
LMPNN	45.9	13.1	10.3	34.8	48.9	22.7	17.6	13.5	10.3	8.7	12.9	7.7	4.6	5.2	24.1	7.8		
Q2T	48.4	15.6	12.4	37.8	51.9	28.6	19.4	21.8	12.8	6.1	12.4	8.4	4.5	4.1	27.6	7.1		
NELL	BetaE	53.0	13.0	11.4	37.6	47.5	24.1	14.3	12.2	8.6	5.1	7.8	10.0	3.1	3.5	24.6	5.9	
	Q2P	56.5	15.2	12.5	35.8	48.7	22.6	16.1	11.1	10.4	5.1	7.4	10.2	3.3	3.4	25.5	6.0	
	ConE	53.1	16.1	13.9	40.0	50.8	26.3	17.5	15.3	11.3	5.7	8.1	10.8	3.5	3.9	27.2	6.4	
	FuzzQE	58.1	19.3	15.7	39.8	50.3	28.1	21.8	17.3	13.7	8.3	10.2	11.5	4.6	5.4	29.3	8.0	
	(with pretrained KGE)																	
	CQD-CO	60.4	17.8	12.8	39.3	46.6	30.1	22.1	17.3	13.2	-	-	-	-	-	28.8	-	
LMPNN	60.6	22.1	17.5	40.1	50.3	28.4	24.9	17.2	15.7	8.5	10.8	12.2	3.9	4.8	30.7	8.0		
Q2T	61.1	22.1	18.0	41.1	52.3	30.3	23.4	19.5	15.3	5.8	7.5	11.3	3.7	4.3	31.5	6.5		

Table 1: MRR results of different CQA models over three KGs. A_p and A_n denote the average score of EPFO queries (1p/2p/3p/2i/3i/pi/ip/2u/up) and queries with negation (2in/3in/inp/pin/pni), respectively. The boldface indicates the best result of each query type.

	A_p	A_n
Q2T (no structural information)	30.2	5.6
- Adjacency Matrices Masking	30.5	6.4
- Undirected Distance Encoding	31.0	6.5
- Directed Distance Encoding	31.5	6.5

Table 2: MRR of models with different structural information encoding strategies on the NELL dataset.

Evaluation. The evaluation metrics follows the previous works (Ren and Leskovec, 2020), which aims to measure the ability of models to discover *hard* answers. That is, these answers cannot be found in the training KG by symbolic methods due to missing edges. Concretely, for each hard answer v of a query q , we rank it against non-answer entities and calculate the Mean Reciprocal Rank (MRR) as evaluation metrics.

Baselines. We mainly compare our work with non symbolic-integrated CQA models for EFO-1 queries. The baselines can be divided into two types: (1) Without pretrained KGE, such as BetaE (Ren and Leskovec, 2020), ConE (Zhang et al., 2021), Q2P (Bai et al., 2022) and FuzzQE (Chen et al., 2022). (2) With pretrained KGE, such as CQD-CO (Arakelyan et al., 2021), and LMPNN (Wang et al., 2023b).

As some works (e.g. KgTransformer (Liu et al., 2022)) do not support negation queries, they use the

EPFO queries (queries without negation) generated by Ren et al., we also compare these works on the EPFO datasets in Appendix F. We also compare our work with symbolic-integrated CQA models in Appendix G.

5.2 Comparison with Baselines

Table 1 shows the MRR results of Q2T and baselines on EFO-1 queries over three benchmarks. The experiment settings can be found in Appendix B.

From the table, we can find that our Q2T can achieve the state-of-the-art performance on almost all positive queries over three datasets. Especially on FB15k and FB15k-237, our Q2T has a significant improvement on EPFO performance, achieving 15.4% and 14.9% relative improvement on A_p respectively. Furthermore, compared to the step-by-step reasoning mode, the end-to-end reasoning mode can eliminate error cascading to some extent, which is reflected in the better performance on multi-hop queries (2p/3p). More details about the model parameters can be found in Appendix C.

This results show that, even without the explicit modeling for neural set operators, our Q2T can still perform well on logical queries, one potential explanation is as follows: **Each pretrained neural link predictor can be regarded as a neural knowledge base which record the structure**

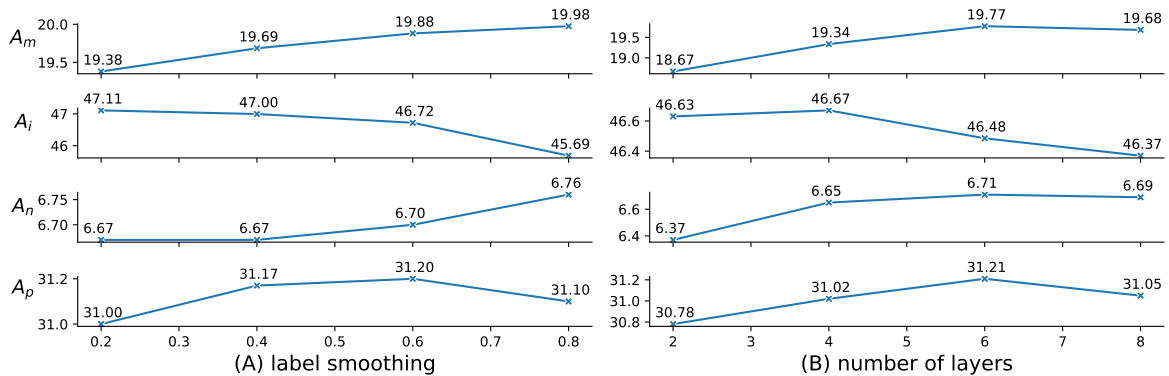


Figure 5: Hyperparameter analysis on NELL. A_m, A_i, A_n, A_p denote the average MRR of multi-hop queries, intersection queries, negation queries, and all positive queries. Since the performance is too low when the label smoothing is 0, they are not shown in the figure (A).

of KG, while learning to transformer complex queries into suitable inputs for the link predictor can be seen as an end-to-end process of retrieving answers from this neural knowledge base.

Compared to BetaE, Q2P and ConE, Q2T’s strong performance on both simple and complex queries underscores the critical importance of separating their training processes.

Although Q2T cannot outperform previous models on negation queries, it still gets competitive performance. One potential explanation for this lies in the fact that the neural link predictor is exclusively pretrained on positive triples. The distribution of answer entities for positive triples tends to be more concentrated, whereas queries containing negation operators result in a more dispersed answer entity distribution. For example, the "complement" operator can yield numerous candidate answers, as seen in queries like "Who is not the winner of the Turing Award." Consequently, pretrained KGE models struggle to generate suitable answer embeddings for queries with dispersed answer distributions.

We also test the performance of Q2T with different pretrained KGE on the FB15k-237 queries, and the results are shown in Appendix E.

5.3 Ablation Studies

In this section, we conduct further experiments to demonstrate the effectiveness of Directed Distance Encoding and investigate how hyperparameters affect the performance on each query type.

Structure Encoding. We compare our proposed Directed Distance Encoding to two commonly used strategies: Undirected Distance Encoding and Adjacency Matrices Masking, which are described in section 4.2. The results are reported in Table 2, it

can be found that models with structural information perform well consistently in CQA, indicating that structural information is necessary. Besides, Distance Encoding that leverages global receptive field is better than Adjacency Matrices Masking that uses local receptive field. Besides, compared to undirected distance, directed distance can further improve performance by introducing extra topological information that help each node to treat their predecessors and successors differently.

Label smoothing. Label smoothing is crucial in Q2T, model without label smoothing suffer from severe overfitting, resulting in bad performance on all queries. More details about label smoothing can be found in Appendix D. Figure 5 (A) shows how label smoothing influences performance, it can be observed that different types of queries prefer different label smoothing values, multi-hop and negation queries prefer high label smoothing value while intersection queries do the opposite, we think this difference is caused by the difference in the average number of answers, as shown in Appendix A. **More specifically, queries with more candidate answers prefer higher label smoothing value.**

Number of layers. To justify how the number of layers affects Q2T, we test the performance of models with a different number of layers on different query types. Figure 5 (B) demonstrates that (1) For intersection queries, model with two layers can handle them well, more layers lead to overfitting. (2) For multi-hop queries, more layers are needed to model long dependencies in the query graphs.

	#params	A_s	A_c	A_p	A_n
Q2T (NF)	56M	44.3	22.0	24.5	5.8
Q2T (F)	26M	48.6	25.1	27.7	7.3

Table 3: Impact of whether freezing KGE on performance. Q2T (NF) denotes Not Freezing KGE, while Q2T (F) denotes Freezing KGE. A_s, A_c denotes the average MRR of simple queries and complex queries. The number of parameters of the KGE is 30M.

5.4 Training with trainable KGE

Aiming to further prove the necessity of decoupling the training for simple queries and complex queries, on the FB15k dataset, we conduct experiments on Q2T with frozen/not frozen pretrained KGE. The results are shown in Table 3, where Q2T (NF) means training KGE and encoder on simple and complex queries together, Q2T (F) means only training encoder for complex queries. Although Q2T (NF) has more trainable parameters, it perform worse than Q2T (F). Especially on simple queries, a neural link predictor trained on simple queries can already handle it well, however, a more complex model trained on more complex queries leads to worse performance. As the final results are predicted by the pretrained link predictor (KGE), the worse the performance on simple queries, the worse the performance on complex queries.

6 Conclusion

In this paper, we present Q2T to transform diverse complex queries into a unified triple form that can be solved by pretrained neural link predictors. Q2T not only decouples the training of complex queries and simple queries from the training perspective, but also decouples KGE and the query encoder from the model perspective. The experiments on three datasets demonstrate that Q2T can achieve the state-of-the-art performance on almost all positive queries.

7 Limitations

The limitations of our proposed Q2T are as follows: (1) Q2T cannot handle negation queries well. Using more strategies to enable Q2T to answer negation queries well is a direction for future work. (2) We only use KGE as our pretrained link predictors, as the performance of Q2T is highly dependent on the pretrained link predictors. Other type models (e.g. GNN based model) as pretrained link predictors may perform better. (3) Limited to time and

resource, we only test performance of Q2T with different KGE models on the FB15k-237 queries.

8 Ethics Statement

This paper proposes a method for complex query answering in knowledge graph reasoning, and the experiments are conducted on public available datasets. As a result, there is no data privacy concern. Meanwhile, this paper does not involve human annotations, and there are no related ethical concerns.

9 Acknowledgment

This work was supported by National Key R&D Program of China (No.2022ZD0118501) and the National Natural Science Foundation of China (No.62376270, No.U1936207, No.61976211). This work was supported by the Strategic Priority Research Program of Chinese Academy of Sciences (No.XDA27020100), Youth Innovation Promotion Association CAS.

References

- Erik Arakelyan, Daniel Daza, Pasquale Minervini, and Michael Cochez. 2021. [Complex query answering with neural link predictors](#). In *9th International Conference on Learning Representations, ICLR 2021, Virtual Event, Austria, May 3-7, 2021*. OpenReview.net.
- Jiaxin Bai, Zihao Wang, Hongming Zhang, and Yangqiu Song. 2022. [Query2Particles: Knowledge graph reasoning with particle embeddings](#). In *Findings of the Association for Computational Linguistics: NAACL 2022*, pages 2703–2714, Seattle, United States. Association for Computational Linguistics.
- Jiaxin Bai, Tianshi Zheng, and Yangqiu Song. 2023. [Sequential Query Encoding For Complex Query Answering on Knowledge Graphs](#). ArXiv:2302.13114 [cs].
- Ivana Balazevic, Carl Allen, and Timothy Hospedales. 2019. [TuckER: Tensor factorization for knowledge graph completion](#). 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 5185–5194, Hong Kong, China. Association for Computational Linguistics.
- Kurt Bollacker, Colin Evans, Praveen Paritosh, Tim Sturge, and Jamie Taylor. 2008. [Freebase: a collaboratively created graph database for structuring human knowledge](#). In *Proceedings of the 2008 ACM SIGMOD international conference on Management of data*, pages 1247–1250.

- Antoine Bordes, Nicolas Usunier, Alberto García-Durán, Jason Weston, and Oksana Yakhnenko. 2013. [Translating embeddings for modeling multi-relational data](#). In *Advances in Neural Information Processing Systems 26: 27th Annual Conference on Neural Information Processing Systems 2013. Proceedings of a meeting held December 5-8, 2013, Lake Tahoe, Nevada, United States*, pages 2787–2795.
- Andrew Carlson, Justin Betteridge, Bryan Kisiel, Burr Settles, Estevam R. Hruschka Jr., 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 2010, Atlanta, Georgia, USA, July 11-15, 2010*. AAAI Press.
- Xuelu Chen, Ziniu Hu, and Yizhou Sun. 2022. [Fuzzy logic based logical query answering on knowledge graphs](#). In *Thirty-Sixth AAAI Conference on Artificial Intelligence, AAAI 2022, Thirty-Fourth Conference on Innovative Applications of Artificial Intelligence, IAAI 2022, The Twelveth Symposium on Educational Advances in Artificial Intelligence, EAAI 2022 Virtual Event, February 22 - March 1, 2022*, pages 3939–3948. AAAI Press.
- Yihong Chen, Pasquale Minervini, Sebastian Riedel, and Pontus Stenetorp. 2021. Relation prediction as an auxiliary training objective for improving multi-relational graph representations. *arXiv preprint arXiv:2110.02834*.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. [BERT: Pre-training of deep bidirectional transformers for language understanding](#). In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 4171–4186, Minneapolis, Minnesota. Association for Computational Linguistics.
- Kelvin Guu, John Miller, and Percy Liang. 2015. [Traversing knowledge graphs in vector space](#). In *Proceedings of the 2015 Conference on Empirical Methods in Natural Language Processing*, pages 318–327, Lisbon, Portugal. Association for Computational Linguistics.
- William L. Hamilton, Payal Bajaj, Marinka Zitnik, Dan Jurafsky, and Jure Leskovec. 2018. [Embedding logical queries on knowledge graphs](#). In *Advances in Neural Information Processing Systems 31: Annual Conference on Neural Information Processing Systems 2018, NeurIPS 2018, December 3-8, 2018, Montréal, Canada*, pages 2030–2041.
- Ziniu Hu, Yuxiao Dong, Kuansan Wang, and Yizhou Sun. 2020. [Heterogeneous graph transformer](#). In *WWW '20: The Web Conference 2020, Taipei, Taiwan, April 20-24, 2020*, pages 2704–2710. ACM / IW3C2.
- Shaoxiong Ji, Shirui Pan, Erik Cambria, Pekka Marttinen, and Philip S. Yu. 2020. [A Survey on Knowledge Graphs: Representation, Acquisition and Applications](#). *ArXiv preprint*, abs/2002.00388.
- Timothée Lacroix, Nicolas Usunier, and Guillaume Obozinski. 2018. [Canonical tensor decomposition for knowledge base completion](#). In *Proceedings of the 35th International Conference on Machine Learning, ICML 2018, Stockholmsmässan, Stockholm, Sweden, July 10-15, 2018*, volume 80 of *Proceedings of Machine Learning Research*, pages 2869–2878. PMLR.
- Pengfei Liu, Weizhe Yuan, Jinlan Fu, Zhengbao Jiang, Hiroaki Hayashi, and Graham Neubig. 2023. Pre-train, prompt, and predict: A systematic survey of prompting methods in natural language processing. *ACM Computing Surveys*, 55(9):1–35.
- 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](#). volume abs/2208.07638.
- E Prudhommeaux, Steve Harris, and A SPARQL Seaborne. 2013. 1.1 query language. *W3C recommendation*. Mar, 21.
- Hongyu Ren, Weihua Hu, and Jure Leskovec. 2020. [Query2box: Reasoning over knowledge graphs in vector space using box embeddings](#). In *8th International Conference on Learning Representations, ICLR 2020, Addis Ababa, Ethiopia, April 26-30, 2020*. OpenReview.net.
- Hongyu Ren and Jure Leskovec. 2020. [Beta embeddings for multi-hop logical reasoning in knowledge graphs](#). In *Advances in Neural Information Processing Systems 33: Annual Conference on Neural Information Processing Systems 2020, NeurIPS 2020, December 6-12, 2020, virtual*.
- Andrea Rossi, Donatella Firmani, Antonio Matinata, Paolo Merialdo, and Denilson Barbosa. 2021. Knowledge graph embedding for link prediction: A comparative analysis. *ACM Trans. Knowl. Discov. Data*, 15:14:1–14:49.
- Mingchen Sun, Kaixiong Zhou, Xin He, Ying Wang, and Xin Wang. 2022a. [GPPT: Graph Pre-training and Prompt Tuning to Generalize Graph Neural Networks](#). In *Proceedings of the 28th ACM SIGKDD Conference on Knowledge Discovery and Data Mining*, pages 1717–1727, Washington DC USA. ACM.
- Tian-Xiang Sun, Xiang-Yang Liu, Xi-Peng Qiu, and Xuan-Jing Huang. 2022b. Paradigm shift in natural language processing. *Machine Intelligence Research*, 19(3):169–183.
- 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](#). In *Proceedings of the 33rd International Conference on Machine Learning, ICML 2016, New York City, NY, USA, June 19-24, 2016*, volume 48 of *JMLR Workshop and Conference Proceedings*, pages 2071–2080. JMLR.org.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N. Gomez, Lukasz Kaiser, and Illia Polosukhin. 2017. [Attention is all you need](#). In *Advances in Neural Information Processing Systems 30: Annual Conference on Neural Information Processing Systems 2017, December 4-9, 2017, Long Beach, CA, USA*, pages 5998–6008.
- Quan Wang, Zhendong Mao, Bin Wang, and Li Guo. 2017. Knowledge Graph Embedding: A Survey of Approaches and Applications. *IEEE TRANSACTIONS ON KNOWLEDGE AND DATA ENGINEERING*, 29(12):20.
- Zhen Wang, Jianwen Zhang, Jianlin Feng, and Zheng Chen. 2014. [Knowledge graph embedding by translating on hyperplanes](#). In *Proceedings of the Twenty-Eighth AAAI Conference on Artificial Intelligence, July 27 -31, 2014, Québec City, Québec, Canada*, pages 1112–1119. AAAI Press.
- Zihao Wang, Weizhi Fei, Hang Yin, Yangqiu Song, Ginny Y. Wong, and Simon See. 2023a. [Wasserstein-Fisher-Rao Embedding: Logical Query Embeddings with Local Comparison and Global Transport](#). ArXiv:2305.04034 [cs].
- Zihao Wang, Yangqiu Song, Ginny Y. Wong, and Simon See. 2023b. [Logical Message Passing Networks with One-hop Inference on Atomic Formulas](#). ArXiv:2301.08859 [cs].
- Zihao Wang, Hang Yin, and Yangqiu Song. 2021. [Benchmarking the Combinatorial Generalizability of Complex Query Answering on Knowledge Graphs](#). ArXiv preprint, abs/2109.08925.
- Robert West, Evgeniy Gabrilovich, Kevin Murphy, Shaohua Sun, Rahul Gupta, and Dekang Lin. 2014. [Knowledge base completion via search-based question answering](#). In *23rd International World Wide Web Conference, WWW '14, Seoul, Republic of Korea, April 7-11, 2014*, pages 515–526. ACM.
- Zonghan Wu, Shirui Pan, Fengwen Chen, Guodong Long, Chengqi Zhang, and S Yu Philip. 2020. A comprehensive survey on graph neural networks. *IEEE transactions on neural networks and learning systems*, 32(1):4–24.
- Wenhan Xiong, Thien Hoang, and William Yang Wang. 2017. [DeepPath: A reinforcement learning method for knowledge graph reasoning](#). In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*, pages 564–573, Copenhagen, Denmark. 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](#). In *3rd International Conference on Learning Representations, ICLR 2015, San Diego, CA, USA, May 7-9, 2015, Conference Track Proceedings*.
- Chengxuan Ying, Tianle Cai, Shengjie Luo, Shuxin Zheng, Guolin Ke, Di He, Yanming Shen, and Tiejian Liu. 2021. [Do transformers really perform badly for graph representation?](#) In *Advances in Neural Information Processing Systems 34: Annual Conference on Neural Information Processing Systems 2021, NeurIPS 2021, December 6-14, 2021, virtual*, pages 28877–28888.
- Zhanqiu Zhang, Jie Wang, Jiajun Chen, Shuiwang Ji, and Feng Wu. 2021. [Cone: Cone embeddings for multi-hop reasoning over knowledge graphs](#). In *Advances in Neural Information Processing Systems 34: Annual Conference on Neural Information Processing Systems 2021, NeurIPS 2021, December 6-14, 2021, virtual*, pages 19172–19183.
- Zhaocheng Zhu, Mikhail Galkin, Zuobai Zhang, and Jian Tang. 2022. [Neural-symbolic models for logical queries on knowledge graphs](#). In *International Conference on Machine Learning, ICML 2022, 17-23 July 2022, Baltimore, Maryland, USA*, volume 162 of *Proceedings of Machine Learning Research*, pages 27454–27478. PMLR.
- Zhaocheng Zhu, Zuobai Zhang, Louis-Pascal A. C. Xhonneux, and Jian Tang. 2021. [Neural bellman-ford networks: A general graph neural network framework for link prediction](#). In *Advances in Neural Information Processing Systems 34: Annual Conference on Neural Information Processing Systems 2021, NeurIPS 2021, December 6-14, 2021, virtual*, pages 29476–29490.

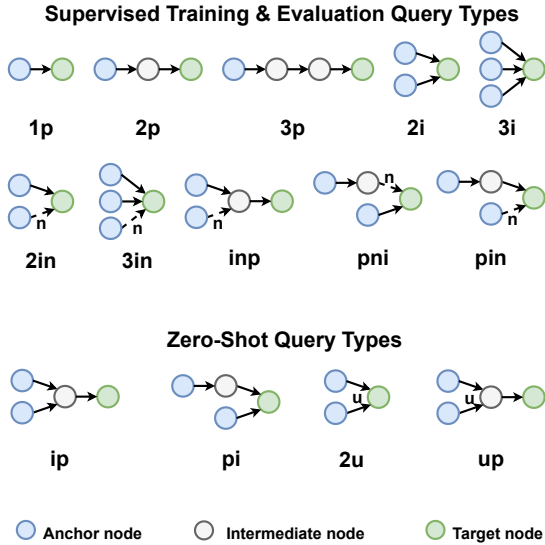


Figure 6: The illustration of all query types. The p, i, n, u represent projection, intersection, negation, and union operations respectively.

A Data Details

The nine query types are shown in Figure 6. Specifically, there are five query types (1p/2p/3p/2i/3i) evaluated in a supervised manner, five query types (2in/3in/inp/pin/pni) evaluated in a few-shot manner, and four query types (2u/up/pi/ip) evaluated in a zero-shot manner. Given the query type, a sample is generated by random walking on the KG. The splitting of datasets are shown in Table 4. The average number of answers of test queries are shown in Table 5.

It should be noted that the training queries and the test queries are generated by random walking on G_{train} (containing training edges) and G_{test} (containing training, validation, and testing edges), respectively, which means that models need to predict at least one missing edges by knowledge graph reasoning.

B Experiment settings

For pretrained KGE model, we use the same settings as CQD (Arakelyan et al., 2021) and LMPNN (Wang et al., 2023b): ComplEx (Trouillon et al., 2016) model with rank 1000. The difference is that we use relation prediction as an auxiliary training objective in the training of ComplEx, as stated in section 4.1, so our ComplEx model perform better on 1p queries. The MRR scores of baselines we use are reported by Wang et al..

The default setting of Q2T is: 6 layer, 768 hidden

size, 12 number of heads. We tune the hyperparameters of Q2T on the validation set for each dataset by grid search. We consider the batch size from {512, 1024, 2048}, learning rate from {1e-4, 2e-4, 4e-4, 5e-4}, and label smoothing from {0.2, 0.4, 0.6, 0.8}. Our experiments are conducted on GTX 3090 with PyTorch 1.11, and the random seed are fixed for each experiment. The best hyperparameters for each datasets are shown in Table 6.

C Model parameters and training time

Table 7 show the number of parameters for models in Table 1. Although Q2T requires training KGE first and then training the query encoder, the overall training time is still shorter than that of other methods. This is due to the high training efficiency achieved by decoupling the training process of KGE and the query encoder.

BetaE (Ren and Leskovec, 2020), ConE (Zhang et al., 2021), Q2P (Bai et al., 2022) don't use pretrained KGE, they train neural set operators with KG embeddings from the zero, which is very inefficient. Although these methods seem to have less parameters than the pretrained KGE used in Q2T, they have much longer training time (e.g. The training times of BetaE and ConE are estimated to more than 20 hours under the official settings, while the training time for Q2T is less than 10 hours). Besides, the training for KGE on simple queries is efficient due to the simple structure of KGE, e.g., on the NELL queries, although ComplEx (Trouillon et al., 2016) model has 110M parameters, its training can be done in one hour.

D Label Smoothing

Aiming to prevent models from overfitting, label smoothing uses soft one-hop labels, instead of hard ones, to introduce noise during training. The smoothed label y_{ls} is computed as follows:

$$y_{ls} = (1 - \alpha)y_h + \frac{\alpha}{K}$$

where y_{ls} is the smoothed label, y_h is the original label, α is the parameter of label smoothing, K is the number of sampling. Lastly, standard cross-entropy loss is applied to these smoothed labels.

E Performance of Q2T with different pretrained KGE

Table 8 presents the MRR results of Q2T with different pretrained KGE on the FB15k-237 queries

Dataset	Training		Validation		Test	
	1p/2p/3p/2i/3i	others	1p	others	1p	others
FB15k	273,710	27,371	59,097	8,000	67,016	8,000
FB15k-237	149,689	14,968	20,101	5,000	22,812	5,000
NELL	107,982	10,798	16,927	4,000	17,034	4,000

Table 4: Number of training, validation, and test queries generated for different query types.

Dataset	1p	2p	3p	2i	3i	ip	pi	2u	up	2in	3in	inp	pin	pni
FB15k	1.7	19.6	24.4	8.0	5.2	18.3	12.5	18.9	23.8	15.9	14.6	19.8	21.6	16.9
FB15k-237	1.7	17.3	24.3	6.9	4.5	17.7	10.4	19.6	24.3	16.3	13.4	19.5	21.7	18.2
NELL995	1.6	14.9	17.5	5.7	6.0	17.4	11.9	14.9	19.0	12.9	11.1	12.9	16.0	13.0

Table 5: Average number of answers of test queries in datasets generated by Ren and Leskovec.

Dataset	batchsize	dropout	ls	lr
FB15k	1024	0.1	0.4	4e-4
FB15k-237	1024	0.1	0.6	4e-4
NELL	1024	0.1	0.6	5e-4

Table 6: The best hyperparameters of Q2T on datasets generated by Ren and Leskovec. The *ls* and *lr* denote label smoothing and learning rate respectively.

generated by Ren and Leskovec. CP (Lacroix et al., 2018), DistMult (Yang et al., 2015) and ComplEx (Trouillon et al., 2016) are trained under the best hyperparameters settings provided by Chen et al., while Tucker (Balazevic et al., 2019) is trained under the best hyperparameters settings provided by Balazevic et al.. It can be observed that the performance of link predictors directly affects the final performance on complex queries. We also find that Q2T with different KGE models may be good at solving different queries.

F Comparison with more baselines on EPFO datasets

Table 10 shows the HIT@3 results of different CQA models on EPFO queries generated by Ren et al.. KgTransformer (Liu et al., 2022) utilizes two-stage pre-training: The first stage aims to initialize KgTransformer with KGs’ general knowledge, and the second stage refines its ability for small queries during inference. Although KgTransformer has more parameters (8 layers, 1024 hidden size) and spends more time on more pre-training data, it still performs worse than our Q2T. The potential reason that KgTransformer perform better on pi/ip queries is that these type of queries may included in its pre-training data, which are randomly sampled from KGs.

G Comparison with symbolic-integrated model

In contrast to non symbolic-integrated methods (e.g., BetaE (Ren and Leskovec, 2020), Q2P (Bai et al., 2022), CQD-CO (Arakelyan et al., 2021)), which use fixed-dimension intermediate embeddings (e.g., 10^2), symbolic-integrated methods (e.g., GNN-QE (Zhu et al., 2022)) employ intermediate fuzzy set sizes that scale linearly with the Knowledge Graph size (e.g., 10^4).

We select GNN-QE as a representative of symbolic-integrated methods for comparison with our Q2T, the results are shown in Table 12. It should be noticed that, GNN-QE employs NBFNet (Zhu et al., 2021) which applies a multi-layers GNN on the whole KGs for each projection operation. Besides, GNN-QE applies fuzzy logic operations to fuzzy sets of entities, so each node needs to obtain the probability list of all entities (a tensor with shape $(|V|,)$). As a result, GNN-QE requires much more training time and lots of GPU resources, it requires 128GB GPU memory to run a batch size of 32, while Q2T only demands 12GB GPU memory to run a batch size of 1024. Even then this, Q2T is also competitive with GNN-QE.

Model	FB15k		FB15k-237		NELL	
	#pretrained KGE params	#trainable params	#pretrained KGE params	#trainable params	#pretrained KGE params	#trainable params
BetaE	-	19M	-	20M	-	58M
ConE	-	28M	-	24M	-	63M
Q2P	-	11M	-	9.9M	-	29M
(with pretrained KGE)						
CQD	35M	0M	30M	0M	128M	0M
LMPNN	35M	16M	30M	16M	128M	16M
Q2T	35M	26M	30M	26M	128M	26M

Table 7: The number of parameters for each model. CQD only use pretrained KGE to obtain answers with optimization-based method, so it doesn't have other trainable parameters.

KGE Model	1p	2p	3p	2i	3i	pi	ip	2u	up	2in	3in	inp	pin	pni	A_p	A_n
TuckER	45.3	12.5	10.1	28.2	39.8	21.8	15.8	16.9	10.0	7.0	8.9	7.6	4.0	5.0	22.2	6.5
CP	46.3	14.9	12.1	37.1	51.3	28.6	17.8	17.4	12.0	6.0	11.7	8.8	4.7	3.9	26.4	7.0
DistMult	47.2	15.1	11.9	37.2	50.9	28.1	18.3	18.7	12.2	6.2	12.5	8.4	4.5	4.1	26.6	7.1
ComplEx	48.5	16.2	12.6	37.8	51.8	28.6	19.2	21.9	12.5	6.4	12.7	8.5	4.6	4.2	27.7	7.3

Table 8: The MRR results of Q2T with different pretrained KGE on the FB15k-237 queries generated by Ren and Leskovec.

Dataset	Model	1p	2p	3p	2i	3i	ip	pi	2u	up	Avg
FB15k-237	GQE	0.405	0.213	0.153	0.298	0.411	0.085	0.182	0.167	0.160	0.230
	Q2B	0.467	0.240	0.186	0.324	0.453	0.108	0.205	0.239	0.193	0.268
	EmQL	0.389	0.201	0.154	0.275	0.386	0.101	0.184	0.115	0.165	0.219
	CQD (CO)	0.512	0.213	0.131	0.352	0.457	0.146	0.222	0.281	0.132	0.272
	CQD (Beam)	0.512	0.288	0.221	0.352	0.457	0.129	0.249	0.284	0.121	0.290
	KgTransformer	0.459	0.312	0.276	0.398	0.528	0.189	0.286	0.263	0.214	0.325
	Q2T	0.530	0.329	0.276	0.415	0.529	0.167	0.251	0.358	0.230	0.343
NELL	GQE	0.417	0.231	0.203	0.318	0.454	0.081	0.188	0.200	0.139	0.248
	Q2B	0.555	0.266	0.233	0.343	0.480	0.132	0.212	0.369	0.163	0.306
	EmQL	0.456	0.231	0.172	0.331	0.483	0.143	0.244	0.226	0.207	0.277
	CQD (CO)	0.667	0.265	0.220	0.410	0.529	0.196	0.302	0.531	0.194	0.368
	CQD (Beam)	0.667	0.350	0.288	0.410	0.529	0.171	0.277	0.531	0.156	0.375
	KgTransformer	0.625	0.401	0.367	0.405	0.546	0.203	0.306	0.469	0.270	0.399
	Q2T	0.670	0.409	0.373	0.397	0.543	0.198	0.282	0.534	0.316	0.414

Table 10: HIT@3 results of different CQA models on EPFO queries generated by Ren et al.

Dataset	Model	1p	2p	3p	2i	3i	pi	ip	2u	up	2in	3in	inp	pin	pni	A_p	A_n
FB15k-237	GNN-QE	42.8	14.7	11.8	38.3	54.1	31.1	18.9	16.2	13.4	10.0	16.8	9.3	7.2	7.8	26.8	10.2
	Q2T	48.4	15.6	12.4	37.8	51.9	28.6	19.4	21.8	12.8	6.1	12.4	8.4	4.5	4.1	27.6	7.1
NELL	GNN-QE	53.3	18.9	14.9	42.4	52.5	30.8	18.9	15.9	12.6	9.9	14.6	11.4	6.3	6.3	28.9	9.7
	Q2T	61.1	22.1	18.0	41.1	52.3	30.3	23.4	19.5	15.3	5.8	7.5	11.3	3.7	4.3	31.5	6.5

Table 12: Comparison between Q2T and GNN-QE.