

Logical Queries on Knowledge Graphs: Emerging Interface of Incomplete Relational Data

Zihao Wang[†] Hang Yin[‡] Yangqiu Song[†]

[†] Department of CSE, HKUST, Hong Kong, China {zwanggc, yqsong}@cse.ust.hk

[‡] Tsinghua University, Beijing, China h-yin20@mails.tsinghua.edu.cn

Abstract

As a graph abstraction of real-world relational data, knowledge graphs contain large amounts of factual knowledge but suffer from incompleteness. Such incompleteness, also known as the open-world assumption, prevents logical queries from being answered through the query-answering paradigms for relational databases. In this paper, we discuss how learning-based models bridge incomplete relational data and logical queries from the perspective of rigorous formal logic. Existing works are structured as a three-party game among knowledge graphs, logical queries, and learning-based models. Compared to the well-established tale between databases and queries in the database theory, current achievements are still preliminary. Our work pointed out several possible directions for future work.

1 Introduction

The knowledge graph is perhaps one of the simplest representations of relational knowledge [8, 56, 38]. A knowledge graph contains a large number of triples, where each triple (h, r, t) contains a head entity h , a tail entity t , and the relation r in between. Such a simple format enables factual knowledge to be adapted to answer real-world questions [49, 62, 46, 40, 32].

Under the view of data engineering, a knowledge graph is no more than a (relational) database [1]. It is natural to expect first-order queries, a fundamental type of relational database queries that have been well-understood theoretically [30] and efficiently solved practically [29], can also be answered on knowledge graphs. However, applying the existing query-answering algorithms designed for (relational) databases can not address logical queries on knowledge graphs. The major obstacle is that large-scale knowledge graphs are inherently incomplete. Modern large-scale knowledge graphs are constructed by crowd-sourcing [56] or even by automatic information extraction pipelines [11]. Given an observed knowledge graph, the non-existence of a triple does not imply that such a triple is not part of the underlying knowledge graph. This issue is acknowledged as the open-world assumption [31]. Therefore, the actual answers to logical queries are not guaranteed by running existing query-answering algorithms on knowledge graphs.

The issue of missing knowledge finds its partial solution with learning-based models. Models that learn from the existing observed triples can predict the missing triples in the underlying knowledge graph. The research topic focusing on the prediction of missing triples with learned embeddings is known as knowledge

Copyright 2022 IEEE. Personal use of this material is permitted. However, permission to reprint/republish this material for advertising or promotional purposes or for creating new collective works for resale or redistribution to servers or lists, or to reuse any copyrighted component of this work in other works must be obtained from the IEEE.

Bulletin of the IEEE Computer Society Technical Committee on Data Engineering

graph representation or knowledge base completion [25, 45]. Then, the logical queries can be addressed by combining the query-answering algorithms and the learning-based models [42, 28, 16]. The participation of learning-based models brings unique strengths and weaknesses to logical query answering, especially compared to traditional query answering algorithms for relational databases [42]. For traditional databases, query-answering algorithms are discussed from the data perspective and the query perspective. The complexity of such algorithms is characterized by the data complexity (the complexity grows with the size of the database) and the query complexity (the complexity grows with the size of the query), respectively. When answering logical queries with learning-based models, it is essential to understand how learning-based models interact with knowledge graphs (model-data interaction) and logical queries (model-query interaction), besides how to design, train, and infer such models.

In this paper, we review recent advancements in logical query answering on knowledge graphs. Our introduction begins with how to define the problems and construct datasets and then to learning-based query answering methods. Our discussion is organized from three perspectives: knowledge graphs (data) perspective, logical query perspective, and model perspective. Section 2 introduces the logical query answering tasks on knowledge graphs in the rigorous language of model theory. All existing datasets and benchmarks are discussed in Section 3 based on the rigorous definitions. Section 4 discusses the learning-based methods for knowledge graph logical query answering. Existing approaches are categorized by their interaction with the knowledge graph (data perspective) and the logical queries (query perspective). Section 5 discusses the limitation of the existing work and several possible future directions.

2 Preliminaries for Logical Queries

Logical queries are formally defined by \mathcal{L} -formula on an \mathcal{L} -structure, which are central concepts in model theory [36]. Specifically, \mathcal{L} -structure rigorously justifies the concept of knowledge graphs while \mathcal{L} -formula formally defines the scope of queries. Formal definitions provide a framework to investigate the data and the query aspects of the knowledge graph logical query answering tasks.

The minimal and informal description of knowledge graphs is a set of triples $\{(h_i, r_i, t_i)\}$ that encapsulates the real-world knowledge, where $h_i, t_i \in \mathcal{E}$ are the head and tail entities and $r_i \in \mathcal{R}$ is the connected relation. Formally speaking, a knowledge graph is an \mathcal{L}_{KG} -structure defined by a language \mathcal{L}_{KG} for knowledge graphs. The concept of language and structure, which concerns syntax and semantics of first-order logic respectively, is defined in Definition 1 and Definition 2.

Definition 1 (Language): A language \mathcal{L} is a triple $(\mathcal{E}, \mathcal{R}, \mathcal{F})$ where \mathcal{E} is the set of constant symbols, \mathcal{R} is the set of predicate symbols, and \mathcal{F} is the set of function symbols. For each $r \in \mathcal{R}$ and $f \in \mathcal{F}$, associated integers n_r and n_f indicate the number of inputs for R and F , respectively.

All (classic) knowledge graphs can be described by a language \mathcal{L}_{KG} by letting (1) \mathcal{R} be the set of binary relations ($n_r = 2$), (2) the constant set \mathcal{E} is finite and contains all entities, and (3) $\mathcal{F} = \emptyset$. \mathcal{E} and \mathcal{R} are finite sets without further classification.

Definition 2 (\mathcal{L} -structure): An \mathcal{L} -structure \mathcal{M} for the language \mathcal{L} is a pair (M, I) , where M is a non-empty set called the universe of \mathcal{M} and I is an interpretation function that: (i) For each $c \in \mathcal{E}$, $I(c) \in M$; (ii) For each $r \in \mathcal{R}$, $I(r) \subset M^{n_r}$, where n_r is the size of relations; (iii) For each $f \in \mathcal{F}$, a function $I(f) : M^{n_f} \mapsto M$, where n_f is the number of inputs in f .

A knowledge graph \mathcal{KG} induces an \mathcal{L}_{KG} -structure by letting $M = \mathcal{E}$ (with $\forall e \in \mathcal{E}, I(e) = e$), so that each relation $I(r)$ is a subset of M^2 . The open world assumption of knowledge graphs can then be formally stated in Definition 3.

Definition 3 (Open world assumption): For a given language \mathcal{L}_{KG} of knowledge graph, there is an underlying \mathcal{L}_{KG} -structure \mathcal{KG}_u , and one (or more) observed \mathcal{L}_{KG} -structure (s) \mathcal{KG}_o , such that for each relation $r \in \mathcal{R}$, the relation set of \mathcal{KG}_o and \mathcal{KG}_u satisfies $I_{\mathcal{KG}_o}(r) \subseteq I_{\mathcal{KG}_u}(r) \subseteq M^2$.

Then, it is ready to define the logical queries with formal concepts, including \mathcal{L} -terms and \mathcal{L} -formulae. The current discussion focuses on first-order logic.

Definition 4 (\mathcal{L} -term): \mathcal{L} -terms be the smallest set of \mathcal{T} such that: (i) $e \in \mathcal{T}$ for each constant symbol $e \in \mathcal{E}$; (ii) Any variable $v_i \in \mathcal{T}$ for $1, 2, \dots$; (iii) if $t_1, t_2, \dots, t_{n_f} \in \mathcal{T}$ and $f \in \mathcal{F}$, then $f(t_1, \dots, t_{n_f}) \in \mathcal{T}$.

Definition 5 (Atomic \mathcal{L} -formula): ϕ is an atomic \mathcal{L} -formula if ϕ is $r(t_1, \dots, t_{n_r})$, where $r \in \mathcal{R}$ is a relation symbol and t_i are \mathcal{L} -terms. Given \mathcal{L} -structure \mathcal{M} , we say $\mathcal{M} \models r(r_1, \dots, r_{n_r})$ is True if and only if the tuple $(t_1, \dots, t_{n_r}) \in I(r)$.

Definition 6 (\mathcal{L} -formula): The set of \mathcal{L} -formula is the smallest set Φ containing all atomic formulae such that: (i) if $\phi \in \Phi$, then $\neg\phi \in \Phi$; (ii) if $\phi, \psi \in \Phi$, then $(\phi \wedge \psi) \in \Phi$ and $(\phi \vee \psi) \in \Phi$; (iii) if $\phi \in \Phi$ and v_i is any variable, then $\exists v_i \phi \in \Phi$ and $\forall v_i \phi \in \Phi$.

We say a variable v is *free* if there are no associated quantifiers. Otherwise, it is *bounded*. We use $\phi(v)$ indicates the \mathcal{L} -formula ϕ contains a free variable v . An \mathcal{L} -formula without free variables is called an \mathcal{L} -sentence. Logical queries are \mathcal{L} -formulae and can be evaluated given a knowledge graph \mathcal{KG} . As studied in the database literature [55], we discuss boolean queries, set-valued queries, and aggregate queries.

A **boolean query** is an \mathcal{L}_{KG} -sentence s . Given the \mathcal{L}_{KG} -structure \mathcal{M} , its answer is the boolean value indicating whether $\mathcal{M} \models s$ or not.

A **set-valued query** is an \mathcal{L}_{KG} -formula ϕ with exactly one free variable v . Given the \mathcal{L} -model \mathcal{M} , its answer set is $\{e : e \in \mathcal{E}, \mathcal{M} \models \phi(v = e)\}$.

An **aggregate query** looks for the aggregation statistics of the answer set of a set-valued query, such as counting, summation, and so on.

3 Logical Query Answering Datasets and Evaluations

Based on the formal concepts given in Section 2, we are ready to formally justify existing datasets from the knowledge graphs (the data perspective) that are queried and the scope of logical queries (the query perspective) to be answered. Instead of going through all existing datasets, we discuss the underlying knowledge graphs and logical queries.

3.1 Knowledge Graphs

Table 1 shows the knowledge graphs investigated in existing datasets. Existing evaluation includes knowledge graphs of different scales (from 10^3 to 10^9) and various properties. Notably, the knowledge graphs with hierarchical relation [37, 5] contain the prior and more complicated structures on entities and relations, respectively, which could be leveraged to improve the performance of query answering method. Inductive logical query answering [19] also considers an inductive KG setting [51], which is more challenging under the open-world assumption.

3.2 Logical Queries

The scope of queries that can be answered is not rigorously justified in the existing works. Though some works claim they are addressing the first-order queries [43] or logical queries in general [33], their definitions are still a

Table 1: Summation of open source knowledge graphs used in the existing datasets. KGs are sorted by the number of entities.

Knowledge graph	# Entities	# Relations	# Total edges	Comment	Datasets
FB15k-237 [9]	14,505	237	310,079	-	[42, 43, 19, 23, 58]
FB15k [53]	14,951	1,345	592,213	-	[42, 43, 58]
DBPedia [5]	34,575	3	240,942	Hierarchical KG	[14]
WN18RR [37]	40,903	11	103,509	Hierarchical KG	[24]
NELL955 [59]	63,361	200	142,804	-	[42, 43, 23, 58]
DRKG [64]	97,238	107	5,874,271	-	[13]
FB400k [50, 8]	409,829	918	2,151,671	-	[41]
ogbl-wikikg2 [22]	2,500,604	535	17,137,181	-	[41, 19]
Freebase [8]	86,054,151	14,824	338,586,276	-	[41]

Table 2: Formal definitions of three typical query families. Compared to the first-order logic formula defined formally with Definition 6, three query families are defined using a subset of connectives or quantifiers (indicated by ✓).

Query Family	\wedge	\vee	\neg	\exists	\forall
Conjunctive Query (CQ)	✓	✗	✗	✓	✗
Existential Positive First Order (EPFO)	✓	✓	✗	✓	✗
Existential First Order (EFO)	✓	✓	✓	✓	✗

strict subset of the first-order queries defined in Definition 6. Table 2 summarizes the formal definitions for three typical query families. The scopes of three query families are sorted in increasing order, i.e., $CQ \subseteq EPFO \subseteq EFO$. The universal quantifier \forall is not widely discussed because its semantics is not usually interpretable over the knowledge graphs. This is the same order for query families that learning-based query answering methods tried to solve.

However, additional assumptions are made so that queries are more suitable to be answered by learning-based methods. Common assumptions include that (1) the query is assumed to be acyclic and (2) only one free variable in the logical formula. The first assumption skips the queries that may involve hard constraint satisfaction problems and the second assumption restricts the space complexity of search within $O(|\mathcal{E}|)$. These assumptions are not necessary for all learning-based models. For example, BiQE [28] accepts conjunctive queries with multiple free variables, and CQD [4] answers general EPFO queries without any additional assumptions. Most existing datasets are usually constructed as companions to evaluate the proposed methods, especially query embedding methods with the operator tree representation discussed in Section 4.2.1. Such methods assume the answers with respect to the only free variable can be computed by executing learnable set operators. The execution order of set operators is decided by compiling the query formula into an operator tree. Thus, popular datasets [42, 43, 58] for EPFO and EFO queries also accept implicit assumptions: (1) the only free variable v is assumed to be the root of the set operator tree; (2) all leaf nodes must be constant entity symbols rather than existential variables.

We see that these assumptions are not related to the logic but only to make the queries easier to solve by a special class of learning-based methods. There are no serious empirical evaluations for queries without such assumptions.

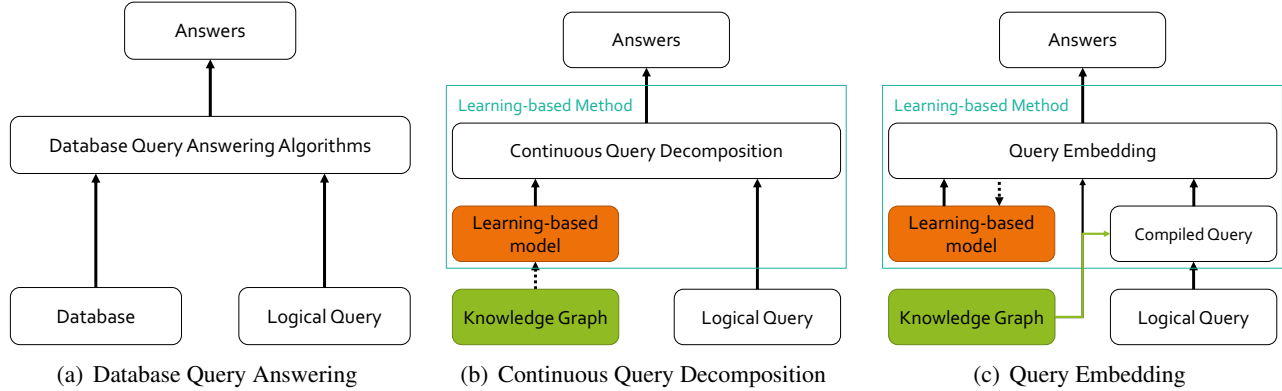


Figure 1: Paradigms for logical query answering on knowledge graphs. The solid line indicates the query answering procedures, and the dashed lines indicate the procedures to obtain the model.

3.3 Evaluation Metrics

The query families are categorized by the formulation of logical formulae. Moreover, for each family, it is also able to construct three types of queries. Here we list popular evaluating metrics.

- For **boolean query**, this becomes a classification problem, and thus standard classification metrics can be applied, like ROC AUC score [43, 21].
- For **set-valued query**, standard metrics include Mean Reciprocal Rank (MRR) [42], Adjusted Mean Rank Index (AMRI) [2], Average percentile rank (APR) [21], and Hits at K (H@K) [42].
- For **aggregate query**, counting is the only problem that has been considered. Spearman’s rank correlation coefficient [43] and Pearson’s correlation coefficient [35] have been proposed when the model is incapable of predicting the precise number. Otherwise, Mean Absolute Error (MAE) is used [35].

4 Learning-Based Methods

Traditional database query answering algorithms are executed over the raw data to find the answers, see Figure 5(b). This paradigm fails in the cases of knowledge graph queries, where the underlying data is not observed. Current knowledge graph logical query answering methods address the open-world assumption with learning-based methods, see the blue boxes in Figure 1(b) and 1(c). In this paper, we categorize the learning-based methods into two categories by how they process the queries. The first category is Continuous Query Decomposition (CQD) and the second is Query Embedding (QE). We see that learning-based models are essential for both the CQD method and the QE method. The participation of the learning-based models adds an intermediate layer between incomplete knowledge graphs and logical queries.

The key to understanding learning-based methods is not only how to define the model but also how models interact with the knowledge graph data and logical queries. The model-query and model-data interaction jointly define the scope of queries that can be answered by each method. Specifically, the model-data interaction refers to how to obtain the model from data, including both training data and algorithms, and the model-query interaction refers to how learning-based models are combined with query answering algorithms. The model-data and model-query interaction are indicated by the dashed arrows and the arrows in Figure 1(b) and 1(c), respectively.

The rest of this section presents the details of these two types of methods. Table 3 compares these two types from three aspects, including training data (model-data interaction), query answering algorithms (model-query interaction), and queries that can be solved. The most distinct difference between CQD and QE is how they

Table 3: Comparison between CQD and QE methods with their training data, query answering algorithms, and solvable queries. † indicates that the actual solvable queries may not be the formally defined logical query family.

Methods	Training Data	Query Answering Algorithms	Solvable Queries
Continuous Query Decomposition (CQD)	KG triples	Model + Optimization	EPFO
Query Embedding (QE)	QA samples	Operator Tree + Model	EFO-1 †
	QA samples	Query Graph + Model	EFO-1 †
	QA samples	Sequence + Model	EPFO †

treat the model as part of the query answering algorithms. Continuous query decomposition estimates the answer embedding by solving an optimization problem with continuous objective induced from the original logical formula with logical t -norms [20]. This objective, computed by the result of the model inference, will be evaluated multiple times during the optimization process. Therefore, the model will be inferred multiple times. This formulation is denoted as Model + Optimization in Table 3 where the term Model is placed before Optimization. In contrast, query embedding methods first compile the logical queries into various representations such as operator trees, query graphs, or sequences. Then, the answer embedding, or the embedding of logical query, is estimated by simple forward inferring the model only once. This formulation is denoted as X + Model, where the term Model is placed after X indicating one of the operator trees, query graphs, or sequences. As natural consequences of such distinct differences in query answering algorithms, CQD and QE require different types of training data to obtain the model and can solve different types of queries. Moreover, the solvable query families for QE methods differ when the queries are compiled into different representations. We note that the solvable queries are confirmed in existing works, and they can be of course extended to broader classes.

The Disjunctive Normal Form (DNF) [36] of the logical queries is of particular interest. This is because (1) all first-order queries can be converted to DNF, and (2) The answer set of DNF existential queries can be regarded as the union of the answer sets of its containing conjunctive clauses. For example, EPFO queries can be answered by compositing the solutions of conjunctive queries. Similarly, general existential logical queries can be answered once the queries with conjunction and negation can be well answered.

4.1 Continuous Query Decomposition (CQD)

The model used in CQD [4] method is no more than a neural link predictor, which is the key part of knowledge graph representation [65, 25]. Specifically, CQD leverages the pretrained ComplEx embedding [54] and achieves state-of-the-art performances in EPFO queries on dataset [42]. The training of the neural link predictor is not related to the logical queries but only knowledge graphs themselves. Recent survey [25] summarizes the construction of neural link predictors and how to train neural link predictors well is discussed in [45].

CQD uses neural link predictors to solve EPFO queries with a search algorithm. It considers the EPFO query $\phi(v)$ in the following DNF form.

$$\phi(v) = \exists x_1, \dots, x_n, (r_{11} \wedge \dots \wedge r_{1k_1}) \vee \dots \vee (r_{l1} \wedge \dots \wedge r_{lk_l}), \quad (1)$$

where r_{ij} is the atomic formula (without negation) and v is the only free variables associated to one or more r_{ij} .

Once the neural link predictor is obtained, each atomic formula $r(h, t)$ can be evaluated with its probability $p(h, r, t)$ of being true. Let p_{ij} be the probability of ij -th triple, \top and \perp be the t -norm and t -conorm [20], we rewrite the EPFO query in Equation (1) into the optimization problem:

$$\max_{v, x_1, \dots, x_n} (p_{11} \top \dots \top p_{1k_1}) \perp \dots \perp (p_{l1} \top \dots \top p_{lk_l}). \quad (2)$$

The objective in Problem (2) is continuous since p , \top , and \perp is continuous. Embedding of variables v, x_1, \dots, x_n will be optimized over the continuous embedding space with gradient-based optimization or searched on the

Natural Language: Which cities are located on rivers that flow through Germany and France?
 Logical Query: $\phi(v) = \exists x \text{HasRiver}(\text{Germany}, x) \wedge \text{HasRiver}(\text{France}, x) \wedge \text{PassesCity}(x, v)$

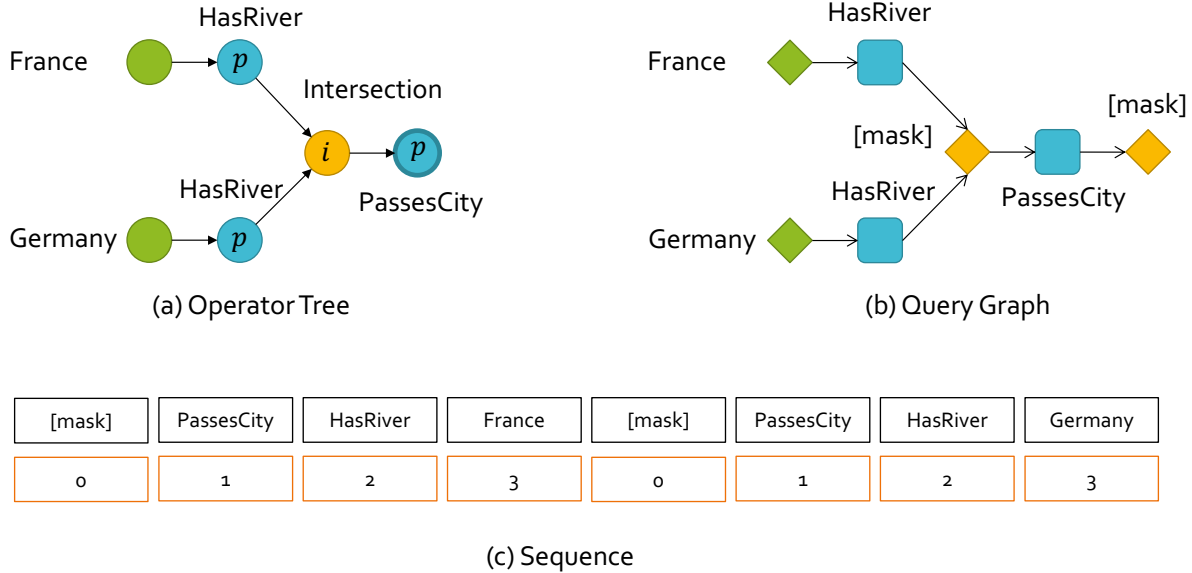


Figure 2: Three ways to represent logical query in the query embedding methods. The example is taken from [28]. We see the logical query is converted into three formats, including (a) operator tree [58], where each node is a set operator; (b) query graph [34], where the nodes are terms and predicates; (c) sequences [28], where a logical query is transformed as the concatenation of several multi-hop queries with a special positional encoding.

discrete set \mathcal{E} with beam-search. Therefore, complex structures of logical queries are handled by solving the optimization problem with the compositional objective function. CQD is also applicable to EPFO queries with multiple free variables in principle. However, this is not empirically justified because of the lack of corresponding datasets. The drawback of CQD is also apparent: CQD eventually estimates the embedding of variables, which remains a significant gap in answering aggregate queries.

4.2 Query Embedding (QE)

In contrast to CQD, query embedding methods use explicit representations of the query structures. Direct inference of the learning-based model over specific query representation estimates the embedding of the target variable, i.e., the query embedding. Then the answers can be retrieved according to the query embedding. The general advantages and disadvantages of QE methods are also clear. On the one hand, models are inferred only once to obtain the answers. On the other hand, specific representation restricts the scope of solvable queries. In the current stage, there is no dominant representation of the queries. Figure 2 summarizes popular choices including operator trees [58], query graphs [34], and sequences [28]. Then we categorize different query embedding methods by their query representations and discuss each type from four perspectives: query representation, model design, training, and query answering algorithms.

4.2.1 Logical Queries as Operator Trees

Operator trees materialize the procedure of searching for answers as the execution of set operators, including set projection, intersection, union, and complement (with respect to a universe). Each entity itself is regarded as the simplest singleton, a set with only one element. The set projection is converted from relations by conducting a

Skolemization [36] process of a logical formula. Other set operations such as intersection, union, and difference or negation are naturally induced by logical operators. Only a subset of first-order formulae can be converted into the operator trees because such a representation of logical queries imposes additional assumptions to Definition 6, as discussed in [58]. Meanwhile, the choice of set operators is not identical. For example, the set complement can be replaced by the set difference [33], and the set union can be replaced by intersection and vice versa with the set complement, according to De Morgan’s laws. The learning-based model for operation tree representation is designed to simulate the sets in the continuous embedding spaces and the set simulations with differentiable computations (usually with neural networks). Then we discuss the spaces used to represent sets and the modeling of the set operations.

Geometric Region Embeddings. Sets, as a collection of discrete objects for logical query answering, are naturally related to the geometrical region of some spaces. Compared to complex variable embeddings as vectors in CQD, embedding sets as geometric regions provides a straightforward estimation of the size of the answer set. The chosen forms of geometric regions are usually simple to be parameterized, such as the boxes used by Query2Box [42] and NewLook [33], sector cones two-dimensional spaces used by ConE [67], and hyperboloid used by HypE [14]. The set projection is usually modeled by neural networks that map the parameterized regions to new ones. Intersection and union operators in NewLook are modeled by some variants of deep set network [63] to merge multiple regions into one. Notably, the design of ConE [67] allows a closed-form set complement with sector cones, while NewLook [33] models the set difference between boxes with neural networks. How to handle negation operation using hyperboloid is not yet discussed.

Probabilistic Distribution Embeddings. Another thread of work models sets with the probabilistic distributions. Parameterized probabilistic distribution families include Beta distribution in BetaE [43], Gaussian distribution in PERM [13], and Gamma distribution in GammaE [61]. Probabilistic distributions are more sophisticated than geometrical regions in set representation because the techniques to construct new probabilistic distributions can be applied to model set operators. For example, a mixture of probabilistic distributions can be used to construct the union operation [13, 61]. Also, altering one parameter to its reciprocal is straightforward to represent the set complement for some family of distributions [43, 61]. Meanwhile, the normalized product of the Probability Distribution Functions (PDF) of Gamma distribution is also Gamma distribution. This fact enables GammaE to simplify the design of the intersection operator. Well-defined distances and divergences between distributions are powerful tools for defining the distances between the answer sets, such as Mahalanobis distance [13] and KL-divergence [43, 61]. Following the intuition of probabilistic distributions, LinE [24] considers the histogram as the unnormalized discretization of arbitrary PDF, and Query2Particles [6] handles a set of empirical samples (particles) from an arbitrary PDF. Such designs sacrifice parts of the closed-form properties but enlarge the family of underlying distributions with non-parametric representations.

Fuzzy Vector Embeddings. Fuzzy vector space [27] brings supreme advantages in modeling the set operations. Set intersection, union, and negation can be precisely modeled with fuzzy logical t -norms and t -conorms [20] by embedding sets as fuzzy vectors [35, 12], i.e., vectors in $[0, 1]^d$ where d is the dimension. The one thing left is to model the set projections with neural networks. FuzzQE [12] models the set projection by relational base decomposition [47] and LogicE [35] uses simple multi-layer perceptron. In addition, LogicE [35] models the lower and the upper bounds of the fuzzy logic truth value, which quantifies the uncertainty of each embedding to answer the aggregate query.

In contrast to various designs for embedding spaces, MLPMix [3] just embedded queries as the vectors in the Euclidean space. Such a simple embedding space does not guarantee any set properties, so it requires more advanced neural networks [52] to model set operators. Notably, MLPMix still has strong performance with simple embedding space compared to its predecessors. This might suggest that even the simplest Euclidean space is large enough. The overall performance results from the trade-off between set embeddings and set operations.

Neural symbolic methods also shed light on another way to improve performances. “Neural” indicates methods answering queries in low-dimensional embedding spaces and neural networks, while “symbolic” indicates referring to the original symbolic knowledge graphs. GNN-QE [68] estimates the probabilities of all entities given

the set projection with graph neural network, NBFNet [69] over the observed knowledge graph. ENeSy [60] ensembles the neural model prediction and symbolic model prediction at each set projection, where the symbolic part is conducted by the TensorLog [15]. These methods show outstanding performances compared to the neural query embedding methods. However, scalability is always an issue since the intermediate states will unavoidably grow with the size of the underlying knowledge graph.

While various models for operator trees were proposed, the method to train the model is almost the same as the first proposed to train Query2box [42]. Let $D(e, q)$ be the distance function between the entity e and the query embedding q , the objective function is

$$\ell_{\theta}(a, q) = -\log \sigma(\gamma - D(a, q)) - \frac{1}{k} \sum_{i=1}^k \log \sigma(D(v_i, q) - \gamma), \quad (3)$$

where $\sigma = \frac{1}{1+e^{-x}}$ is the sigmoid function, γ is the margin, a is an arbitrary answer, and v_i are k negative sampled answers.

4.2.2 Logical Queries as Query Graphs

Query graph presentation represents the terms and predicates in to graphs [16, 34, 57]. Query graph representation does not assume the acyclicity as operator tree representation. Therefore it represents a larger set of logical queries. However, existing query graphs are not related to the disjunction. The disjunction can be handled by converting existing queries to DNF queries, whose answer set is the combination of all conjunctive queries that can be addressed in the query graphs. More specifically, a conjunctive query is represented by two kinds of query graphs: (1) heterogenous information networks (HIN) [48] where relations and entities are all nodes [34]; (2) constraint graphs where nodes are terms and edges are atomic formulas [4, 57]. This representation is connected to the constraint programming [44].

The learning-based models for query graphs are basically Graph Neural Networks (GNN). For constraint graph representation, MPQE [16] applies message passing networks on EPFO queries and LMPNN [57] employs the pretrained knowledge graph embeddings to answer queries with negation. kgTransformer [34] uses graph transformers whose multi-layer perceptrons are upgraded by a Mixture of Experts (MoE) architecture. Besides the negative sampling objectives shown in Equation (3), kgTransformer proposed to first pretrain the model on sampled subgraphs and then finetune the model with the negative sampling loss on query answering data.

4.2.3 Logical Queries as Sequences

The last type of representation is to represent the queries in sequences. BiQE [28], see Figure 2 (c), proposes to convert the query graph to multiple paths from the answer node to the anchor entity nodes. Compared to the query graphs, the sequence representation even neglects the existential variables, only the anchor entity and the relations on its path to the answer node remain. Relative positional embeddings are also used to describe the distance from each token to the answer node. This representation also focuses on conjunctive queries, which do not contain disjunction and negation. Queries with disjunctions can be handled by combining conjunctive queries in the DNF. However, it is still not known how to handle the negation operation.

Once the sequence is defined, the sequence can be put into the transformer models. The output embedding of the [mask] token is aggregated to obtain the answer. Unlike the query graph representation that enforces the query structures explicitly, transformers could learn implicit logical structures via the self-attention mechanism.

5 Future Directions

Research about knowledge graphs logical query answering, though rapidly developed in recent years, is still limited compared to the well-established study about query answering on the relational database. For example,

the query types that can be answered by the model are strict fragments of first-order queries. In this section, we discuss what is left in current research compared to logical query answering on incomplete relational data, which is the ultimate goal of this line of research. Formal definitions from model theory in Section 2 picture the gap between existing works and the ambitious goal. We summarize the limitations and discuss possible future directions from the data, query, and model perspectives.

5.1 Towards General Relational Data

A relational language \mathcal{L} can describe general relational data. However, the language \mathcal{L}_{KG} for knowledge graphs is defined with three additional assumptions, see Section 2. By rethinking three assumptions about \mathcal{L}_{KG} and the open world assumption, we could extend the existing data model of knowledge graphs.

KG with Relations of $n_r > 2$. Knowledge graphs with n -ary relations [66] relax the first assumption that $n_r = 2, \forall R \in \mathcal{R}$ by accepting the relations whose $2 \leq n_r \leq n$. Thus, each element of the knowledge graph now is an $n + 1$ triple list that contains n entities and one relation. The concept of n -ary relation is also related to the hyper-relation [18, 2] and the knowledge graph is extended from the directed graphs to hypergraphs.

KG with Attributes. Attributes expand the constant set \mathcal{E} from all entities to *entities with attribute values*, thus relaxing the second assumption. Typical attributes include entity types [5, 23], triple timestamps [26, 46], and triple probabilities [11]. Additional attributes can be described in the n -ary relations. For example, the timestamp t associated with a triple (s, p, o) can be described with 4-triples (s, p, o, t) . New attributes bring new predicates, enlarge the universe set, and enrich the query semantics. Two timestamps can be compared by “before”, “equal”, and “after”, and can be composed to a time interval, which also introduces new predicates such as “during”. The triple probabilities, as discussed in probabilistic databases [55], enable the estimation of the probabilities for boolean queries, the probability for each answer entity for set-valued queries, and the total probability for aggregate queries.

KG with Functions. Introducing attributes also makes it possible to define functions between attribute values. Taking the timestamp attribute as an example, new timestamps can be computed by a time movement function with a timestamp and a time interval as inputs. These features are rigorously defined in pure symbolic systems [17]. It is worth discussing how to combine them with learning-based methods.

More Challenging Open World Assumption. In the inductive setting, the entity set \mathcal{E} is divided into an observed set \mathcal{E}_o^r and an inempty inductive set \mathcal{E}_i^r with respect to the relation $r \in \mathcal{R}$, that is $\mathcal{E}_o^r \cap \mathcal{E}_i^r = \emptyset, \mathcal{E}_o^r \cup \mathcal{E}_i^r = \mathcal{E}, \forall r \in \mathcal{R}$. Under such a setting, the open world assumption in Definition 3 is rewritten as $I_{\mathcal{K}G_o}(r) \subseteq \mathcal{E}_o^r \times \mathcal{E}_o^r \subseteq I_{\mathcal{K}G_u}(r) \subseteq \mathcal{E} \times \mathcal{E}$. This setting makes it very hard to estimate the relations over $\mathcal{E} \times \mathcal{E} - \mathcal{E}_o^r \times \mathcal{E}_o^r$.

5.2 Towards First Order Queries and Beyond

The semantics of the logic queries lacks serious discussion. There is no standard dataset and benchmark for first-order queries with cycles, multiple free variables, or universal quantifiers even for the simplest \mathcal{L}_{KG} language. Moreover, it is also worth discussing first-order queries over general relational data such as KGs with attributes and attribute functions. It is also worthwhile to expand the first-order logic to fixed-point or monadic second-order logic [30].

5.3 Towards Versatile Learning-based Models

Existing methods are incapable of addressing more and more challenging queries and complex data discussed above. It is essential to investigate the graph or sequence-based query embedding models to solve various queries. Given the combinatorial natures of the query answering problems, new opportunities could be found with neural combinatorial solvers [7]. Meanwhile, large language models, as a proven effective way to encapsulate world knowledge [39, 10], could be considered as another approach to modeling the relational data.

6 Conclusion

Big data and world knowledge are usually uncertain and dynamic. Logical queries over knowledge graphs are essential ways to understand the world as a simplified data model through the lens of logic and reasoning. In this paper, we structure the existing knowledge graph logical query answering tasks and methods as a three-party game between data, query, and models. With the rigorous concepts from finite model theory [36, 30], our work portrays a detailed landscape of the achievements and limitations of existing work. Our discussion shows that there are various gaps between current research to the ambitious goal in terms of data, queries, and models. We hope our paper can be a practical guide for traveling in the uncertain and dynamic data world.

7 Acknowledgement

The authors of this paper were supported by the NSFC Fund (U20B2053) from the NSFC of China, the RIF (R6020-19 and R6021-20) and the GRF (16211520 and 16205322) from RGC of Hong Kong, the MHKJFS (MHP/001/19) from ITC of Hong Kong and the National Key R&D Program of China (2019YFE0198200) with special thanks to HKMAAC and CUSBLT, and the Jiangsu Province Science and Technology Collaboration Fund (BZ2021065). We also thank the support from NVIDIA AI Technology Center (NVAITC) and the UGC Research Matching Grants (RMGS20EG01-D, RMGS20CR11, RMGS20CR12, RMGS20EG19, RMGS20EG21, RMGS23CR05, RMGS23EG08).

References

- [1] S. Abiteboul, Richard Hull, and Victor Vianu. Foundations of Databases. Addison-Wesley, Reading, Mass, 1995.
- [2] Dimitrios Alivanistos, Max Berrendorf, Michael Cochez, and Mikhail Galkin. Query Embedding on Hyper-Relational Knowledge Graphs. In International Conference on Learning Representations, March 2022.
- [3] Alfonso Amayuelas, Shuai Zhang, Susie Xi Rao, and Ce Zhang. Neural Methods for Logical Reasoning over Knowledge Graphs. In The Tenth International Conference on Learning Representations, ICLR 2022, Virtual Event, April 25-29, 2022. OpenReview.net, 2022.
- [4] Erik Arakelyan, Daniel Daza, Pasquale Minervini, and Michael Cochez. Complex Query Answering with Neural Link Predictors. In International Conference on Learning Representations, 2021.
- [5] Sören Auer, Christian Bizer, Georgi Kobilarov, Jens Lehmann, Richard Cyganiak, and Zachary Ives. DBpedia: A Nucleus for a Web of Open Data. In Karl Aberer, Key-Sun Choi, Natasha Noy, Dean Allemang, Kyung-Il Lee, Lyndon Nixon, Jennifer Golbeck, Peter Mika, Diana Maynard, Riichiro Mizoguchi, Guus Schreiber, and Philippe Cudré-Mauroux, editors, The Semantic Web, Lecture Notes in Computer Science, pages 722–735, Berlin, Heidelberg, 2007. Springer.
- [6] Jiaxin Bai, Zihao Wang, Hongming Zhang, and Yangqiu Song. Query2Particles: Knowledge Graph Reasoning with Particle Embeddings. In Findings of the Association for Computational Linguistics: NAACL 2022, pages 2703–2714, Seattle, United States, July 2022. Association for Computational Linguistics.
- [7] Irwan Bello, Hieu Pham, Quoc V. Le, Mohammad Norouzi, and Samy Bengio. Neural combinatorial optimization with reinforcement learning. In 5th International Conference on Learning Representations, ICLR 2017, Toulon, France, April 24-26, 2017, Workshop Track Proceedings. OpenReview.net, 2017.

- [8] Kurt D. Bollacker, Colin Evans, Praveen K. Paritosh, Tim Sturge, and Jamie Taylor. Freebase: A collaboratively created graph database for structuring human knowledge. In Jason Tsong-Li Wang, editor, Proceedings of the ACM SIGMOD International Conference on Management of Data, SIGMOD 2008, Vancouver, BC, Canada, June 10-12, 2008, pages 1247–1250. ACM, 2008.
- [9] Antoine Bordes, Nicolas Usunier, Alberto Garcia-Duran, Jason Weston, and Oksana Yakhnenko. Translating Embeddings for Modeling Multi-relational Data. In Advances in Neural Information Processing Systems, volume 26. Curran Associates, Inc., 2013.
- [10] Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, et al. Language models are few-shot learners. Advances in neural information processing systems, 33:1877–1901, 2020.
- [11] Andrew Carlson, Justin Betteridge, Bryan Kisiel, Burr Settles, Estevam R. Hruschka Jr, and Tom M. Mitchell. Toward an Architecture for Never-Ending Language Learning. In Maria Fox and David Poole, editors, Proceedings of the Twenty-Fourth AAAI Conference on Artificial Intelligence, AAAI 2010, Atlanta, Georgia, USA, July 11-15, 2010. AAAI Press, 2010.
- [12] Xuelu Chen, Ziniu Hu, and Yizhou Sun. 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, 2022.
- [13] Nurendra Choudhary, Nikhil Rao, Sumeet Katariya, Karthik Subbian, and Chandan Reddy. Probabilistic Entity Representation Model for Reasoning over Knowledge Graphs. In Advances in Neural Information Processing Systems, volume 34, pages 23440–23451. Curran Associates, Inc., 2021.
- [14] Nurendra Choudhary, Nikhil Rao, Sumeet Katariya, Karthik Subbian, and Chandan K. Reddy. Self-Supervised Hyperboloid Representations from Logical Queries over Knowledge Graphs. In Proceedings of the Web Conference 2021, WWW '21, pages 1373–1384, New York, NY, USA, June 2021. Association for Computing Machinery.
- [15] William Cohen, Fan Yang, and Kathryn Rivard Mazaitis. TensorLog: A Probabilistic Database Implemented Using Deep-Learning Infrastructure. Journal of Artificial Intelligence Research, 67:285–325, February 2020.
- [16] Daniel Daza and Michael Cochez. Message Passing Query Embedding, June 2020.
- [17] E Allen Emerson. Temporal and modal logic. In Formal Models and Semantics, pages 995–1072. Elsevier, 1990.
- [18] Mikhail Galkin, Priyansh Trivedi, Gaurav Maheshwari, Ricardo Usbeck, and Jens Lehmann. Message Passing for Hyper-Relational Knowledge Graphs. In Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP), pages 7346–7359, Online, November 2020. Association for Computational Linguistics.
- [19] Mikhail Galkin, Zhaocheng Zhu, Hongyu Ren, and Jian Tang. Inductive Logical Query Answering in Knowledge Graphs. In Advances in Neural Information Processing Systems 33: Annual Conference on Neural Information Processing Systems 2022 NeurIPS 2022, November 27- December 9, 2022, New Orleans, page 25, 2022.

- [20] Petr Hájek. Metamathematics of Fuzzy Logic, volume 4 of Trends in Logic. Springer Netherlands, Dordrecht, 1998.
- [21] William L. Hamilton, Payal Bajaj, Marinka Zitnik, Dan Jurafsky, and Jure Leskovec. Embedding Logical Queries on Knowledge Graphs. In Samy Bengio, Hanna M. Wallach, Hugo Larochelle, Kristen Grauman, Nicolò Cesa-Bianchi, and Roman Garnett, editors, 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, 2018.
- [22] Weihua Hu, Matthias Fey, Marinka Zitnik, Yuxiao Dong, Hongyu Ren, Bowen Liu, Michele Catasta, and Jure Leskovec. Open Graph Benchmark: Datasets for Machine Learning on Graphs. In Advances in Neural Information Processing Systems, volume 33, pages 22118–22133. Curran Associates, Inc., 2020.
- [23] Zhiwei Hu, Victor Gutierrez Basulto, Zhiliang Xiang, Xiaoli Li, Ru Li, and Jeff Z. Pan. Type-aware Embeddings for Multi-Hop Reasoning over Knowledge Graphs. In Proceedings of the Thirty-First International Joint Conference on Artificial Intelligence, pages 3078–3084, Vienna, Austria, July 2022. International Joint Conferences on Artificial Intelligence Organization.
- [24] Zijian Huang, Meng-Fen Chiang, and Wang-Chien Lee. LinE: Logical Query Reasoning over Hierarchical Knowledge Graphs. In Proceedings of the 28th ACM SIGKDD Conference on Knowledge Discovery and Data Mining, KDD '22, pages 615–625, New York, NY, USA, August 2022. Association for Computing Machinery.
- [25] Shaoxiong Ji, Shirui Pan, Erik Cambria, Pekka Marttinen, and Philip S. Yu. A survey on knowledge graphs: Representation, acquisition, and applications. IEEE Transactions on Neural Networks and Learning Systems, 33(2):494–514, February 2022.
- [26] Zhen Jia, Soumajit Pramanik, Rishiraj Saha Roy, and Gerhard Weikum. Complex Temporal Question Answering on Knowledge Graphs. In Gianluca Demartini, Guido Zuccon, J. Shane Culpepper, Zi Huang, and Hanghang Tong, editors, CIKM '21: The 30th ACM International Conference on Information and Knowledge Management, Virtual Event, Queensland, Australia, November 1 - 5, 2021, pages 792–802. ACM, 2021.
- [27] A. K Katsaras and D. B Liu. Fuzzy vector spaces and fuzzy topological vector spaces. Journal of Mathematical Analysis and Applications, 58(1):135–146, March 1977.
- [28] Bhushan Kotnis, Carolin Lawrence, and Mathias Niepert. Answering Complex Queries in Knowledge Graphs with Bidirectional Sequence Encoders. Proceedings of the AAAI Conference on Artificial Intelligence, 35(6):4968–4977, May 2021.
- [29] David M. Kroenke, David J. Auer, Scott L. Vandenberg, and Robert C. Yoder. Database Processing: Fundamentals, Design, and Implementation. Pearson, NY NY, 15th edition, 40th anniversary edition edition, 2018.
- [30] Leonid Libkin. Elements of Finite Model Theory. Springer Berlin Heidelberg, Berlin, Heidelberg, 2004.
- [31] Leonid Libkin and Cristina Sirangelo. Open and Closed World Assumptions in Data Exchange. In Bernardo Cuenca Grau, Ian Horrocks, Boris Motik, and Ulrike Sattler, editors, Proceedings of the 22nd International Workshop on Description Logics (DL 2009), Oxford, UK, July 27-30, 2009, volume 477 of CEUR Workshop Proceedings. CEUR-WS.org, 2009.
- [32] Bill Yuchen Lin, Xinyue Chen, Jamin Chen, and Xiang Ren. KagNet: Knowledge-Aware Graph Networks for Commonsense Reasoning. EMNLP/IJCNLP, 2019.

- [33] Lihui Liu, Boxin Du, Heng Ji, ChengXiang Zhai, and Hanghang Tong. Neural-Answering Logical Queries on Knowledge Graphs. In Proceedings of the 27th ACM SIGKDD Conference on Knowledge Discovery & Data Mining, KDD '21, pages 1087–1097, New York, NY, USA, August 2021. Association for Computing Machinery.
- [34] Xiao Liu, Shiyu Zhao, Kai Su, Yukuo Cen, Jiezhong Qiu, Mengdi Zhang, Wei Wu, Yuxiao Dong, and Jie Tang. 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, pages 1120–1130, New York, NY, USA, August 2022. Association for Computing Machinery.
- [35] Francois Luus, Prithviraj Sen, Pavan Kapanipathi, Ryan Riegel, Ndivhuwo Makondo, Thabang Lebeso, and Alexander Gray. Logic Embeddings for Complex Query Answering. [arXiv:2103.00418 \[cs\]](https://arxiv.org/abs/2103.00418), February 2021.
- [36] D. Marker. Model Theory: An Introduction. Number 217 in Graduate Texts in Mathematics. Springer, New York, 2002.
- [37] George A. Miller. WordNet: A lexical database for English. Communications of the ACM, 38(11):39–41, November 1995.
- [38] Thomas Pellissier Tanon, Denny Vrandečić, Sebastian Schaffert, Thomas Steiner, and Lydia Pintscher. From Freebase to Wikidata: The Great Migration. In Proceedings of the 25th International Conference on World Wide Web, WWW '16, pages 1419–1428, Republic and Canton of Geneva, CHE, April 2016. International World Wide Web Conferences Steering Committee.
- [39] Fabio Petroni, Tim Rocktäschel, Sebastian Riedel, Patrick Lewis, Anton Bakhtin, Yuxiang Wu, and Alexander Miller. Language models as knowledge bases? 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 2463–2473, 2019.
- [40] Hongyu Ren, Hanjun Dai, Bo Dai, Xinyun Chen, Michihiro Yasunaga, Haitian Sun, Dale Schuurmans, Jure Leskovec, and Denny Zhou. LEGO: Latent Execution-Guided Reasoning for Multi-Hop Question Answering on Knowledge Graphs. In Proceedings of the 38th International Conference on Machine Learning, pages 8959–8970. PMLR, July 2021.
- [41] Hongyu Ren, Hanjun Dai, Bo Dai, Xinyun Chen, Denny Zhou, Jure Leskovec, and Dale Schuurmans. SMORE: Knowledge Graph Completion and Multi-hop Reasoning in Massive Knowledge Graphs. In Proceedings of the 28th ACM SIGKDD Conference on Knowledge Discovery and Data Mining, KDD '22, pages 1472–1482, New York, NY, USA, August 2022. Association for Computing Machinery.
- [42] Hongyu Ren, Weihua Hu, and Jure Leskovec. 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, 2020.
- [43] Hongyu Ren and Jure Leskovec. Beta Embeddings for Multi-Hop Logical Reasoning in Knowledge Graphs. In Hugo Larochelle, Marc’Aurelio Ranzato, Raia Hadsell, Maria-Florina Balcan, and Hsuan-Tien Lin, editors, Advances in Neural Information Processing Systems 33: Annual Conference on Neural Information Processing Systems 2020, NeurIPS 2020, December 6-12, 2020, Virtual, 2020.
- [44] Francesca Rossi, Peter Van Beek, and Toby Walsh. Handbook of constraint programming. Elsevier, 2006.

- [45] Daniel Ruffinelli, Samuel Broscheit, and Rainer Gemulla. You CAN Teach an Old Dog New Tricks! On Training Knowledge Graph Embeddings. In International Conference on Learning Representations, March 2020.
- [46] Apoorv Saxena, Soumen Chakrabarti, and Partha P. Talukdar. Question Answering Over Temporal Knowledge Graphs. In Chengqing Zong, Fei Xia, Wenjie Li, and Roberto Navigli, editors, Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing, ACL/IJCNLP 2021, (Volume 1: Long Papers), Virtual Event, August 1-6, 2021, pages 6663–6676. Association for Computational Linguistics, 2021.
- [47] Michael Schlichtkrull, Thomas N. Kipf, Peter Bloem, Rianne van den Berg, Ivan Titov, and Max Welling. Modeling Relational Data with Graph Convolutional Networks. In Aldo Gangemi, Roberto Navigli, Maria-Esther Vidal, Pascal Hitzler, Raphaël Troncy, Laura Hollink, Anna Tordai, and Mehwish Alam, editors, The Semantic Web, Lecture Notes in Computer Science, pages 593–607, Cham, 2018. Springer International Publishing.
- [48] Chuan Shi, Yitong Li, Jiawei Zhang, Yizhou Sun, and S Yu Philip. A survey of heterogeneous information network analysis. IEEE Transactions on Knowledge and Data Engineering, 29(1):17–37, 2016.
- [49] Yueqing Sun, Qi Shi, Le Qi, and Yu Zhang. JointLK: Joint Reasoning with Language Models and Knowledge Graphs for Commonsense Question Answering. In Marine Carpuat, Marie-Catherine de Marneffe, and Iván Vladimir Meza Ruíz, editors, Proceedings of the 2022 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, NAACL 2022, Seattle, WA, United States, July 10-15, 2022, pages 5049–5060. Association for Computational Linguistics, 2022.
- [50] Alon Talmor and Jonathan Berant. The Web as a Knowledge-Base for Answering Complex Questions. In Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long Papers), pages 641–651, New Orleans, Louisiana, June 2018. Association for Computational Linguistics.
- [51] Komal Teru, Etienne Denis, and William L Hamilton. Inductive Relation Prediction by Subgraph Reasoning. In Proceedings of the 37th International Conference on Machine Learning, pages 9448–9457. PMLR, November 2020.
- [52] Ilya O Tolstikhin, Neil Houlsby, Alexander Kolesnikov, Lucas Beyer, Xiaohua Zhai, Thomas Unterthiner, Jessica Yung, Andreas Steiner, Daniel Keysers, Jakob Uszkoreit, et al. Mlp-mixer: An all-mlp architecture for vision. Advances in Neural Information Processing Systems, 34:24261–24272, 2021.
- [53] Kristina Toutanova and Danqi Chen. 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, July 2015. Association for Computational Linguistics.
- [54] Théo Trouillon, Johannes Welbl, Sebastian Riedel, Eric Gaussier, and Guillaume Bouchard. Complex Embeddings for Simple Link Prediction. In Proceedings of The 33rd International Conference on Machine Learning, pages 2071–2080. PMLR, June 2016.
- [55] Guy Van den Broeck and Dan Suciu. Query Processing on Probabilistic Data: A Survey. Foundations and Trends® in Databases, 7(3-4):197–341, 2017.
- [56] Denny Vrandečić and Markus Krötzsch. Wikidata: A free collaborative knowledgebase. Communications of the ACM, 57(10):78–85, September 2014.

- [57] Zihao Wang, Yangqiu Song, Ginny Y. Wong, and Simon See. Logical message passing networks with one-hop inference on atomic formulas. In The Eleventh International Conference on Learning Representations, ICLR 2023, Kigali Rwanda, May 1-5, 2023. OpenReview.net, 2023.
- [58] Zihao Wang, Hang Yin, and Yangqiu Song. Benchmarking the Combinatorial Generalizability of Complex Query Answering on Knowledge Graphs. In Joaquin Vanschoren and Sai-Kit Yeung, editors, Proceedings of the Neural Information Processing Systems Track on Datasets and Benchmarks 1, NeurIPS Datasets and Benchmarks 2021, December 2021, Virtual, 2021.
- [59] Wenhan Xiong, Thien Hoang, and William Yang Wang. 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, 2017.
- [60] Zezhong Xu, Wen Zhang, Peng Ye, Hui Chen, and Huajun Chen. Neural-Symbolic Entangled Framework for Complex Query Answering. In Advances in Neural Information Processing Systems, October 2022.
- [61] Dong Yang, Peijun Qing, Yang Li, Haonan Lu, and Xiaodong Lin. GammaE: Gamma Embeddings for Logical Queries on Knowledge Graphs, October 2022.
- [62] Michihiro Yasunaga, Hongyu Ren, Antoine Bosselut, Percy Liang, and Jure Leskovec. QA-GNN: Reasoning with Language Models and Knowledge Graphs for Question Answering. In Kristina Toutanova, Anna Rumshisky, Luke Zettlemoyer, Dilek Hakkani-Tür, Iz Beltagy, Steven Bethard, Ryan Cotterell, Tanmoy Chakraborty, and Yichao Zhou, editors, Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, NAACL-HLT 2021, Online, June 6-11, 2021, pages 535–546. Association for Computational Linguistics, 2021.
- [63] Manzil Zaheer, Satwik Kottur, Siamak Ravanbakhsh, Barnabas Poczos, Russ R Salakhutdinov, and Alexander J Smola. Deep Sets. In Advances in Neural Information Processing Systems, volume 30. Curran Associates, Inc., 2017.
- [64] Rui Zhang, Dimitar Hristovski, Dalton Schutte, Andrej Kastrin, Marcelo Fiszman, and Halil Kilicoglu. Drug repurposing for COVID-19 via knowledge graph completion. Journal of Biomedical Informatics, 115:103696, March 2021.
- [65] Wen Zhang, Jiaoyan Chen, Juan Li, Zezhong Xu, Jeff Z. Pan, and Huajun Chen. Knowledge Graph Reasoning with Logics and Embeddings: Survey and Perspective. arXiv:2202.07412 [cs], February 2022.
- [66] Yao Zhang, Peiyao Li, Hongru Liang, Adam Jatowt, and Zhenglu Yang. Fact-Tree Reasoning for N-ary Question Answering over Knowledge Graphs. In Smaranda Muresan, Preslav Nakov, and Aline Villavicencio, editors, Findings of the Association for Computational Linguistics: ACL 2022, Dublin, Ireland, May 22-27, 2022, pages 788–802. Association for Computational Linguistics, 2022.
- [67] Zhanqiu Zhang, Jie Wang, Jiajun Chen, Shuiwang Ji, and Feng Wu. ConE: Cone Embeddings for Multi-Hop Reasoning over Knowledge Graphs. In Advances in Neural Information Processing Systems, volume 34, pages 19172–19183. Curran Associates, Inc., 2021.
- [68] Zhaocheng Zhu, Mikhail Galkin, Zuobai Zhang, and Jian Tang. Neural-Symbolic Models for Logical Queries on Knowledge Graphs, May 2022.
- [69] Zhaocheng Zhu, Zuobai Zhang, Louis-Pascal Xhonneux, and Jian Tang. Neural Bellman-Ford Networks: A General Graph Neural Network Framework for Link Prediction. In Advances in Neural Information Processing Systems, volume 34, pages 29476–29490. Curran Associates, Inc., 2021.