

Hybrid reasoning in knowledge graphs: Combing symbolic reasoning and statistical reasoning

Weizhuo Li ^{a,b}, Guilin Qi ^{a,b,*} and Qiu Ji ^c

^a *School of Computer Science and Engineering, Southeast University, Nanjing, China*

^b *Key Laboratory of Computer Network and Information Integration (Southeast University), Ministry of Education, Nanjing, China*

E-mails: liweizhuo@amss.ac.cn, gqi@seu.edu.cn

^c *School of Modern Posts and Institute of Modern Posts, Nanjing University of Posts and Telecommunications, Nanjing, China*

E-mail: qiuji@njupt.edu.cn

Editors: Pascal Hitzler, Kansas State University, Manhattan, KS, USA; Krzysztof Janowicz, University of California, Santa Barbara, USA

Solicited reviews: Axel Polleres, Vienna University of Economics and Business (WU Wien), Austria; One anonymous reviewer

Abstract. Knowledge graphs (KGs) contain rich resources that represent human knowledge in the world. There are mainly two kinds of reasoning techniques in knowledge graphs, symbolic reasoning and statistical reasoning. However, both of them have their merits and limitations. Therefore, it is desirable to combine them to provide hybrid reasoning in a knowledge graph. In this paper, we present the first work on the survey of methods for hybrid reasoning in knowledge graphs. We categorize existing methods based on applications of reasoning techniques, and introduce the key ideas of them. Finally, we re-examine the remaining research problems to be solved and provide an outlook to future directions for hybrid reasoning in knowledge graphs.

Keywords: Knowledge graphs, Hybrid reasoning, Knowledge representation

1. Introduction

With the rapid development of Internet technology and Web applications, large amounts of data are published online, which become an important source for large-scale knowledge extraction. How to organize, represent and analyze this knowledge has attracted much attention. Knowledge graphs (KGs) contain rich resources that represent human knowledge in the world. Most of KGs are directed labeled graphs composed of entities (nodes) and various relations (different semantic labels of edges) [1]. A fact in a knowledge graph is usually represented as a triple of

the form (head entity, relation, tail entity), indicating that two entities are connected by a specific relation, e.g. (Barack Obama, BornIn, Honolulu Hawaii). Recent years have witnessed the rapid growth in open KGs such as DBpedia [2], YAGO [3], NELL [4] and Probase [5], which have been widely used to support real applications of the Semantic Web.

The quality of a knowledge graph is critical to its applications, such as question answering. Two important factors that influence the quality of a knowledge graph are completeness and logical coherence of a knowledge graph. Knowledge reasoning, which plays an important role in the services of KGs, aims at inferring implicit knowledge to enrich incomplete KGs and refine their logical correctness. There are two main-

*Corresponding author. E-mail: gqi@seu.edu.cn.

stream techniques for knowledge reasoning. One is based on symbolic reasoning that formalizes the problem by a semantic framework and infers the implicit knowledge according to some predefined rules. The other is based on statistical reasoning that tries to find suitable statistical models to fit the samples and predicts the expected probabilities of inferred relations between entities.

Unfortunately, both symbolic reasoning and statistical reasoning have drawbacks in applications of knowledge graphs. Symbolic reasoning is often based on either rules or schematic knowledge, which is hard to obtain. In contrast, statistical reasoning draws imprecise conclusions and the results of reasoning may be hard to find an explanation. Therefore, many researchers tried to combine their advantages together, and obtained some encouraging performance in related tasks such as knowledge completion [6, 7], schema induction [8, 9], knowledge alignment [10, 11], question answering [12, 13] and so on. For example, one can merge the symbolic information (e.g. path, context or logical rules) into the statistical framework so as to constrain the conditions of object functions or refine the predicted results.

So far, there is no systematical and in-depth survey on hybrid reasoning methods in KGs for various goals of reasoning. In this paper, we summarize the latest research progress of methods in knowledge graphs and look forward to future development directions and prospects. Specifically, we first give a short introduction of knowledge graphs, and analyze the pros and cons of symbolic reasoning and statistic reasoning, respectively, which motivate the necessity of hybrid reasoning. Next, we provide a thorough review of current methods for various goals of reasoning in KGs. Finally, we re-examine remaining research challenges and give an outlook to future directions for hybrid reasoning in KGs.

2. Hybrid reasoning in knowledge graph

In this section, we present a short introduction of knowledge graphs and motivate hybrid reasoning in a knowledge graph. So far, some people have tried to provide a formal definition of a knowledge graph [14, 15]. However, none of them has become a standard definition as the term "knowledge graph" can have different views. In this paper, we do not intend to provide such a definition, but consider the characteristics of a knowledge graph given in [16, 17]:

- mainly describes real world entities and their interrelations, organized in a graph.
- defines classes and properties of entities in a schema.
- allows for potentially interrelating arbitrary entities with each other.
- covers various topical domains.

As shown in Fig. 1, entities represent real-world individuals (e.g. "Yao Ming" and his wife "Ye Li"). A concept represents a set of individuals with the same characteristics, for example, "Yao Ming", "Kobe Bryant", "Michael Jordan", and etc., compose a set corresponding to the concept "Basketball Player". Literals refer to the strings which indicate specific values of some relations, such as string "2.29 m", the "height" of entity "Yao Ming". Edges between these nodes represent different relationships between entities, concepts and literals, such as "Yao Ming" is a "Basketball player" and the wife of "Yao Ming" is "Ye Li". All of these relationships and their related entities, concepts or literals are stored in the form of triples, which are the basic storage units of knowledge graphs. Triples organize knowledge in the form of <subject, predicate, object>, e.g. <Yao Ming, is-a, Basketball Player> and <Yao Ming, height, "2.29 m">.




Fig. 1. An example for a part of a knowledge graph

There are two kinds of knowledge in a knowledge graph, one is called schematic knowledge and the other is called factual knowledge. The schematic knowledge consists of the statements about concepts and properties, and the factual knowledge consists of the statement about instances. For example, the triple <Asian Country, subclassOf, Country> is a piece of schematic knowledge, whilst the triples given in Fig. 1 are all factual knowledge. Existing knowledge graphs mostly consist of lots of factual knowledge and a small amount of schematic knowledge. For example, the

well-known knowledge graph DBpedia¹ contains more than 6.6M entities and over 13 billion triples. However, it only contains 685 concepts which are described by 2,795 different properties, and these concepts form a subsumption hierarchy consisting of *subclassOf* relations. There exist some knowledge graphs which consist of a large number of schematic knowledge, such as SNOMED CT².

Knowledge graphs have their logical foundations based in ontology languages, such as the Resource Description Framework (RDF)³ and the Ontology Web Language (OWL)⁴, which are W3C recommended criteria. RDF is a graph data model for describing resources on the Web and to enable data exchange and sharing; it was originally used to represent metadata of a webpage, such as what tools were used to create the webpage and the authors of the webpage. The factual knowledge in a knowledge graph can be described by RDF. OWL is a family of ontology languages which can represent rich and complex knowledge about entities, properties and relations. OWL can describe both factual and schematic knowledge and can support logical reasoning. Since ontology languages, such as RDF and OWL, are often based on first-order logic semantics, one kind of reasoning in a knowledge graph is deductive reasoning. Logic-based reasoning, or symbolic reasoning, is important to ensure the quality of a knowledge graph and to infer implicit knowledge from a given knowledge graph. Another approach to reasoning in a knowledge graph is based on statistical machine learning, and this kind of reasoning is often called statistical reasoning. Both symbolic reasoning and statistical reasoning have their pros and cons. Symbolic reasoning can infer precise conclusions, but it is often based on either rules or schematic knowledge, which are hard to obtain. In contrast, statistical reasoning draws imprecise conclusions and is often data-driven, thus is easier to scale to large knowledge graphs without human intervention or with little human intervention. Therefore, it is desirable to combine symbolic reasoning and statistical reasoning to provide hybrid reasoning in a knowledge graph. In the following sections, we will give a review of existing works on hybrid reasoning in a knowledge graph and present some challenging problems for future works.

¹<https://wiki.dbpedia.org>

²<https://bioportal.bioontology.org/ontologies/SNOMEDCT>

³<https://www.w3.org/RDF/>

⁴<https://www.w3.org/OWL/>

3. Methodology

In this section, we roughly categorize hybrid reasoning techniques based on goals of reasoning in KGs into four groups: knowledge completion, schematic knowledge induction, knowledge alignment, multi-hop reasoning for question answering. We also introduce some other hybrid reasoning methods that are hard to be categorized into these groups.

3.1. Knowledge completion

To deal with the problem of incompleteness in knowledge graphs, much work has been done to apply statistical relational learning (SRL) models [18] to infer implicit relations between two entities in a knowledge graph. Path ranking algorithm (PRA) [19] and knowledge graph embedding (KGE) [1] are two typical kinds of methods belonging to SRL, and have shown widely used in knowledge completion. In this subsection, we first introduce path ranking algorithm. We then introduce three categories of methods for knowledge graph embedding models.

3.1.1. Path ranking algorithm and its extensions

Path ranking algorithm based on random walk techniques is proposed for discovering complex path features of relational data [19]. The key idea of PRA is employing the paths that connect two entities as features to predict potential relations between them. For example, $\langle \text{bornIn}, \text{capitalOf} \rangle$ is a path linking Ludwig van Beethoven to Germany, through an intermediate node Bonn. Such paths can be used as features to predict the presence of specific relations, e.g. nationality. There exist various extensions that have been explored such as incorporating text corpus [20], using subgraph feature extraction [21] and so on. Some well-known KGs such as NELL [4] adopt PRA to perform knowledge completion.

3.1.2. Merging relational paths in KGE

Knowledge graph embedding encodes components of a KG including entities and relations into continuous vector spaces [1]. There are mainly three types of KG embedding models. The first is translational distance models, such as TransE, which exploit distance-based scoring functions and measure the plausibility of a fact as the distance between two entities [6]. The second is semantic matching models, like RESCAL [7], which measure plausibility of facts by matching latent semantics of entities and relations embodied in their vector space representations. Another type of

KG embedding models is based on language modeling approaches that employ unsupervised feature extraction from sequences of words. RDF2Vec [22] generates a set of sequences of entities using two different approaches, i.e. graph walks and Weisfeiler-Lehman Subtree RDF graph kernels. Then, the authors utilized those sequences to train Word2vec for estimating the likelihood of a sequence of entities appearing in a graph. Cochez et al. [23] exploited a global pattern instead of local sequences generated for nodes in RDF2Vec. The authors combined Global Vectors (GloVe) with Bookmark-Coloring Algorithm to efficiently learn embeddings of entities.

As triples in KGs are not independent, so the interrelations of triples should not be ignored, which can give context information to improve existing KG embedding models. PTransE [24] extends TransE by modeling a path-based representation. The authors utilized connected relational facts between entity pairs instead of only considering the relation between two entities. Since not all relational paths are reliable, they designed a path-constraint resource allocation algorithm to measure the reliability of relation paths and represented these paths via semantic composition of relation embeddings. GAKE [25] defines three types of graph contexts which contain different KGs structured information for representation learning. Therefore, the score function of GAKE takes into account the connection between target entities (or relations) and their contexts. In addition, the authors designed an attention mechanism to learn the representative power of different vertices or edges. Gao et al. [26] proposed a triple context-based embedding method called TCE for knowledge graph completion. TCE takes two kinds of structured information of each triple into consideration. One is a set of neighboring entities along with their outgoing relations, the other is a set of relation paths which contain a pair of target entities.

3.1.3. Employing logical rules in KGE

Logical rules can also enhance the performance of KG embedding models for knowledge completion. Wang et al. [27] utilized these rules to refine embedding models. In their work, KG completion was formulated as an integer linear programming problem that was constrained by rules. Hence, the inferred facts would be the most preferred by the embedding models and complied with all the rules. Similarly, Wei et al. [28] combined rules and embedding models via Markov logic networks, in which they incorporated the similarity priori generated by embedding-based mod-

els into inferring and designed a grounding network sampling strategy to promote the inference precision. On the other hand, logical rules can be represented as horn clauses e.g. $\forall x, y (x, \text{Capital-Of}, y) \rightarrow (x, \text{Located-In}, y)$ stating that any two entities linked by a relation Capital-Of should also be satisfied with a relation Located-In. Guo et al. [29] proposed a joint model that embeds factual knowledge and logical rules in a unified framework, in which logical rules were interpreted as complex formulae constructed by combining ground atoms with logical connectives (e.g. \wedge and \rightarrow) and measured by t-norm fuzzy logics. After that, they improved this model further [30], which could learn simultaneously from labeled triples, unlabeled triples and soft rules in an iterative manner. Zhang et al. [31] proposed a novel framework called IterE for alleviating of sparsity entities in KGs. IterE could iteratively learn embeddings and logical rules, in which rules are learned from embeddings with proper pruning strategy, and embeddings are learned from existing triples and new triples inferred by rules. In addition, Gutiérrez-Basulto and Schockaert [32] argued that existing combined models might not represent expressive classes of rules sufficiently, and proposed a method based on convex-regions. With defined convex-regions, KGs restricted to the quasi-chained existential rules could be faithfully encoded in most cases.

3.1.4. Preserving logical properties in KGE

Another type of KG embedding methods has been proposed for preserving the logical properties of semantic relations. On2Vec [33] employs translation-based embedding models for populating ontologies, which integrated matrices that transformed the head and tail entities in order to characterize the transitivity of some relations. To represent concepts, instances, and relations differently in the same semantic space, TransC [34] encodes instances as vectors and concepts as spheres so that they could preserve the transitivity of isA relations. Sun et al. [35] proposed a model based on complex spaces, called RotateE. It employs the characteristics of complex real numbers and imaginary numbers to effectively characterize the symmetry, antisymmetry and composition of relations.

3.2. Schematic knowledge induction

Existing KGs contain lots of triples but lack schematic knowledge, e.g. *subclassOf* axioms and *disjointness* axioms. It brings a difficulty to infer implicit information, deal with the heterogeneity for ontology mapping

[36], object reconciliation [37] and resolve contradictions [38]. Hence, learning schematic knowledge to enrich KGs becomes a critical and meaningful task.

One main category of the methods to produce schematic knowledge combines rule mining algorithms with symbolic reasoning. The works in [36, 39] defined association rule patterns to generate various kinds of axioms and performed inconsistency handling for ontology construction by enriching an original schema incrementally. Considering the open world assumption adopted by KGs, Galárraga et al. [8] adopted partial completeness assumption to generate counterexamples for rules and redefined the standard measurements for support and confidence. Its extension AMIE+ [40] further improved the precision by using type hierarchy and joint reasoning when learning association rules. Inspired by these methods, Gao et al. [41] exploited a type inference algorithm and defined a mining model with the probabilistic type assertions to deal with noisy negative examples, which could generate high-quality *subclassOf* axioms and *disjointness* axioms. To improve the scalability of rule-based methods, Omran et al. [42] introduced a new sampling algorithm and embedding representations of arguments. Both of them could guide the extraction of rules. Similarly, the work in [43] employed embedding models to iteratively extract rules by probabilistic representations of missing facts and feedback from a precomputed embedding model.

The other main category combines machine learning techniques with logical reasoning. The work in [9] used inductive logic programming, which integrated machine learning with logic programming, and defined an *ACC* downward refinement operator for learning concept descriptions. This operator was extended in [44] that could learn more expressive schematic knowledge like cardinality restrictions. In [38], a statistical method was proposed to extract domain and range of a property. The vector space model from information retrieval was applied to extract disjoint concepts. After the extraction finished, consistency checking was performed in parallel based on predefined inconsistency patterns. The work in [45] integrated the probabilistic inference capability of Bayesian networks with the logical formalism to learn *subclassOf* axioms and *disjointness* axioms. It used logical rules for generating more complex axioms and dealing with inconsistency during the schema construction of KGs.

3.3. Knowledge alignment

Over past decades, more and more knowledge graphs become available on the Web, but the heterogeneity and multi-linguality gap of KGs still hinder their sharing and reusing in the Semantic Web. Benefited from hybrid reasoning, the studies of knowledge alignment have obtained some encouraging results.

Cross-lingual taxonomy alignment (CLTA) refers to mapping each category in the source taxonomy of one language onto the most relevant category in the target taxonomy of another language. However, existing methods for CLTA mainly rely on features based on symbolic similarities. Wu et al. [10] proposed a bilingual topic model, called Bilingual Biterm Topic Model (BiBTM). After identified the matched categories based on string similarity, they trained BiBTM by textual contexts extracted from the Web and obtained the topic vector of the extracted textual context for each category. Finally, they utilized the cosine similarity between topic vectors to calculate the taxonomy alignment. Furthermore, they improved the performance of proposed models by merging explicit category correlations including co-occurrence correlation and structural correlation [46].

In addition, there exist some works that employ embedding-based ideas [6] for entity alignment (EA) among knowledge graphs. MTransE [11] separately trains the entity embeddings of two KGs and designed different techniques to represent cross-lingual transitions including axis calibration, translation vectors and linear transformations. JAPE [47] learns the embeddings of two KGs in a unified space and leveraged attributes of triples to refine entity embeddings. To deal with the problem of lack of prior alignment, IPTransE [48] and BootEA [49] employ an iterative process and designed several sophisticated strategies based on the structure of KGs to refine the new alignment. Chen et al. [50] proposed a method called KDCoE, which co-trains the embeddings of multilingual KGs and descriptions of entities. To utilize various features of KGs, Zhang et al. [51] proposed a framework to unify multiple views of entities and learn embeddings for entity alignment. Furthermore, they designed two cross-KG identity inference methods at the entity level as well as the relation and attribute level to preserve and enhance the alignment between different KGs.

3.4. Multi-hop reasoning for question answering

Question answering (QA) is a hot topic that has recently been facilitated by large-scale knowledge bases.

1 However, due to the variety and complexity of ques-
2 tions and knowledge, question answering over knowl-
3 edge bases (KBQA) is still a challenging task, espe-
4 cially in multi-hop QA.

5 There are two typical categories of multi-relation
6 questions, a path question [52] and a conjunctive ques-
7 tion [53]. A path question contains only one topic en-
8 tity and its answer could be found by walking down
9 an answer path consisting of a few relations and inter-
10 mediate entities. A conjunctive question contains more
11 than one subject entity and its answer could be ob-
12 tained by the intersection of results from multiple path
13 questions. At present, semantic parsing models [12]
14 and embedding-based models [13] tailored for QA are
15 not adequate to handle multi-hop QA because of heavy
16 data annotations and reasoning ability. Therefore, re-
17 cent works utilized hybrid ideas to improve the perfor-
18 mance and make these results explainable.

19 Zhang et al. [54] proposed a probabilistic model-
20 ing framework for multi-hop QA, which could simul-
21 taneously handle uncertain topic entity and multi-hop
22 reasoning for QA. They introduced a new propaga-
23 tion architecture over KGs so that logical inference
24 could be performed in the probabilistic model. Zhou
25 et al. [52] designed an interpretable reasoning network
26 (IRN). It could dynamically decide which part of an
27 input question should be analyzed at each hop, and
28 predict a relation corresponding to the parsed results.
29 Compared with existing methods, the intermediate en-
30 tities and relations predicted by IRN could construct
31 traceable reasoning paths to reveal how the answer
32 was derived. Hamilton et al. [53] introduced a frame-
33 work to efficiently make predictions about conjunctive
34 logical queries. They encoded graph nodes in a low-
35 dimensional space and represented logical operators
36 (i.e. projection operator and intersection operator) as
37 learned geometric operations. Moreover, they further
38 demonstrated how to map a practical subset of logic to
39 efficient geometric operations in an embedding space.
40 Vakulenko et al. [55] proposed a novel approach for
41 complex QA using unsupervised message passing. It
42 could propagate confidence scores by parsing an input
43 question and matching terms in a KG to a set of possi-
44 ble answers. This approach was implemented as a se-
45 ries of sparse matrix multiplications mimicking joins
46 over small local subgraphs so that it could successfully
47 be applied to very large KGs, such as DBpedia.

3.5. Other hybrid reasoning methods

1 Other hybrid reasoning methods focus on boosting
2 the performance of NLP tasks. Most of them merge
3 the symbolic information (e.g. the structure of a KG)
4 into the statistic-based methods and provide the expla-
5 nation for results of reasoning [56].

6 Wang et al. [57] proposed a joint model that takes
7 advantage of both explicit and implicit representations
8 for short text classification. They incorporated char-
9 acter level features of KG into a convolutional neural
10 network to capture fine-grained subword information.
11 Experiments on real data showed that their method
12 achieved significant improvement for this task.

13 To alleviate the bound of number and quality of
14 annotated data, Luo et al. [58] exploited the rich ex-
15 pressiveness of regular expressions at different levels
16 within a neural network (NN). This combined frame-
17 work could significantly enhance the learning effec-
18 tiveness and improve the performance on the tasks of
19 intent detection and slot filling.

20 To tackle the problem of learning and prediction
21 with concept drifts, Chen et al. [59] revisited features
22 embeddings as semantic ones (i.e. consistency vec-
23 tors and entailment vectors). Such embeddings can
24 be exploited in a context of supervised stream learn-
25 ing to learn statistic models, which are robust to con-
26 cept drifts. Moreover, they explored an ontology-based
27 knowledge representation and reasoning framework
28 for the transfer learning explanation [60]. It can model
29 a learning domain in transfer learning with expressive
30 OWL ontologies and complement the learning domain
31 with the prediction task-related common sense knowl-
32 edge. Furthermore, the authors designed a correlative
33 reasoning algorithm to infer three kinds of explanatory
34 evidence for explaining a positive feature or a negative
35 transfer from one learning domain to another.

4. Conclusion and future direction

41 Hybrid reasoning in knowledge graphs plays an
42 important role in knowledge completion, schematic
43 knowledge induction, knowledge alignment, complex
44 question answering, explanation of AI, etc. However,
45 there does not exist a survey of existing methods and
46 discuss the challenging problems for this topic. In this
47 paper, we gave an overview of existing methods for hy-
48 brid reasoning in KGs. We provided a thorough review
49 of current methods for various goals of reasoning in
50 KGs, and further introduced the key ideas of them. AI

though there exist many methods for hybrid reasoning in knowledge graphs, there still exist some problems to be solved which are listed in the following.

- **Knowledge completion:** Taking relational paths and logical rules into account can efficiently improve the performance of KG embedding models. Nevertheless, few methods consider the reliability of triples and deal with sparse long-tail relations, which are ubiquitous in KGs. Besides, existing methods focus on preserving partial logical properties of relations. It is still hard for them to encode complex definitions of concepts and logical properties of relations in OWL language.
- **Schematic knowledge induction:** Horn rule is one of the simplest schemas that can be learned from KGs. It is still challenging to consider complex schematic knowledge, involving existential variables and disjunctions. As far as we know, some research groups are working on this challenging problem and we hope some good results can be obtained in the next two years.
- **Knowledge alignment:** Although KG embedding has been used in some recent works on entity alignment, it is still not clear if it is useful for taxonomy alignment. The recent works on schema embedding can provide some possibilities to apply KG embedding to taxonomy alignment. Furthermore, inconsistency handling is an important issue closely related to knowledge alignment and is often solved by a symbolic method. It is challenging to propose a hybrid reasoning method that can deal with knowledge alignment and inconsistency handling simultaneously.
- **Multi-hop reasoning for QA:** The frameworks of multi-hop reasoning are still limited by some types of questions so that they cannot handle arithmetic operation or logical queries with negation or disjunction. Integrating attention mechanism [61] and utilizing graph neural networks [62] to incorporate richer feature information on nodes and edges will be two promising directions.

Acknowledgements

The authors would like to thank all the reviewers for their insightful and valuable suggestions, which significantly improve the quality of this survey. Research presented in this paper was partially supported by the National Key Research and Development

Program of China under grants (2018YFC0830200, 2017YFB1002801), the Natural Science Foundation of China grants (U1736204, 61602259), the Fundamental Research Funds for the Central Universities (3209009601), and the Judicial Big Data Research Centre, School of Law at Southeast University.

References

- [1] Q. Wang, Z. Mao, B. Wang and L. Guo, Knowledge Graph Embedding: A Survey of Approaches and Applications, *IEEE Transactions on Knowledge and Data Engineering* **29**(12) (2017), 2724–2743. doi: [10.1109/TKDE.2017.2754499](https://doi.org/10.1109/TKDE.2017.2754499).
- [2] J. Lehmann, R. Isele, M. Jakob, A. Jentzsch, D. Kontokostas, P.N. Mendes, S. Hellmann, M. Morsey, P. van Kleef, S. Auer and C. Bizer, DBpedia - A large-scale, multilingual knowledge base extracted from Wikipedia, *Semantic Web* **6**(2) (2015), 167–195. doi: [10.3233/SW-140134](https://doi.org/10.3233/SW-140134).
- [3] T. Rebele, F.M. Suchanek, J. Hoffart, J. Biega, E. Kuzey and G. Weikum, YAGO: A Multilingual Knowledge Base from Wikipedia, Wordnet, and Geonames, in: *Proceedings of the 15th International Semantic Web Conference, Kobe, Japan*, Springer, 2016, pp. 177–185. doi: [10.1007/978-3-319-46547-0_19](https://doi.org/10.1007/978-3-319-46547-0_19).
- [4] T. Mitchell, W. Cohen, E. Hruschka, P. Talukdar, B. Yang, J. Betteridge, A. Carlson, B. Dalvi, M. Gardner, B. Kisiel et al., Never-Ending Learning, *Communications of the ACM* **61**(5) (2018), 103–115. doi: [10.1145/3191513](https://doi.org/10.1145/3191513).
- [5] W. Wu, H. Li, H. Wang and K.Q. Zhu, Probase: A Probabilistic Taxonomy for Text Understanding, in: *Proceedings of the ACM SIGMOD International Conference on Management of Data, Scottsdale, AZ, USA*, ACM, 2012, pp. 481–492. doi: [10.1145/2213836.2213891](https://doi.org/10.1145/2213836.2213891).
- [6] R. Socher, D. Chen, C.D. Manning and A. Ng, Reasoning With Neural Tensor Networks for Knowledge Base Completion, in: *Proceedings of the 27th Annual Conference on Neural Information Processing Systems, Lake Tahoe, Nevada, USA*, Curran Associates, Inc., 2013, pp. 926–934.
- [7] M. Nickel, V. Tresp and H.-P. Kriegel, A Three-Way Model for Collective Learning on Multi-Relational Data, in: *Proceedings of the 28th International Conference on Machine Learning, Bellevue, Washington, USA*, Omnipress, 2011, pp. 809–816.
- [8] L.A. Galárraga, C. Teflioudi, K. Hose and F. Suchanek, AMIE: Association Rule Mining Under Incomplete Evidence in Ontological Knowledge Bases, in: *Proceedings of 22nd International World Wide Web Conference, Rio de Janeiro, Brazil*, ACM, 2013, pp. 413–422. doi: [10.1145/2488388.2488425](https://doi.org/10.1145/2488388.2488425).
- [9] J. Lehmann and P. Hitzler, A Refinement Operator Based Learning Algorithm for the *ALC* Description Logic, in: *Inductive Logic Programming: Proceedings of 17th International Conference, Corvallis, OR, USA*, AAAI Press, 2007, pp. 147–160. doi: [10.1007/978-3-540-78469-2_17](https://doi.org/10.1007/978-3-540-78469-2_17).
- [10] T. Wu, G. Qi, H. Wang, K. Xu and X. Cui, Cross-Lingual Taxonomy Alignment with Bilingual Biterm Topic Model, in: *Proceedings of the 30th AAAI Conference on Artificial Intelligence, Phoenix, Arizona, USA*, AAAI Press, 2016, pp. 287–293.

- [11] M. Chen, Y. Tian, M. Yang and C. Zaniolo, Multilingual Knowledge Graph Embeddings for Cross-lingual Knowledge Alignment, in: *Proceedings of the 26th International Joint Conference on Artificial Intelligence, Melbourne, Australia*, AAAI Press, 2017, pp. 1511–1517. doi: [10.24963/ijcai.2017/209](https://doi.org/10.24963/ijcai.2017/209).
- [12] A. Abujabal, R.S. Roy, M. Yahya and G. Weikum, QUINT: Interpretable Question Answering over Knowledge Bases, in: *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing, Copenhagen, Denmark*, ACL, 2017, pp. 61–66. doi: [10.18653/v1/D17-2011](https://doi.org/10.18653/v1/D17-2011).
- [13] Y. Hao, Y. Zhang, K. Liu, S. He, Z. Liu, H. Wu and J. Zhao, An End-to-End Model for Question Answering over Knowledge Base with Cross-Attention Combining Global Knowledge, in: *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics, Vancouver, Canada*, ACL, 2017, pp. 221–231. doi: [10.18653/v1/P17-1021](https://doi.org/10.18653/v1/P17-1021).
- [14] L. Ehrlinger and W. Wöß, Towards a Definition of Knowledge Graphs, in: *Proceedings of the 12th International Conference on Semantic Systems, Leipzig, Germany*, CEUR-WS.org, 2016.
- [15] M. Färber, F. Bartscherer, C. Menne and A. Rettinger, Linked data quality of DBpedia, Freebase, OpenCyc, Wikidata, and YAGO, *Semantic Web* 9(1) (2018), 77–129. doi: [10.3233/SW-170275](https://doi.org/10.3233/SW-170275).
- [16] P.A. Bonatti, S. Decker, A. Polleres and V. Presutti, Knowledge Graphs: New Directions for Knowledge Representation on the Semantic Web (Dagstuhl Seminar 18371), *Dagstuhl Reports* 8(9) (2018), 29–111. doi: [10.4230/DagRep.8.9.29](https://doi.org/10.4230/DagRep.8.9.29).
- [17] H. Paulheim, Knowledge graph refinement: A survey of approaches and evaluation methods, *Semantic Web* 8(3) (2017), 489–508. doi: [10.3233/SW-160218](https://doi.org/10.3233/SW-160218).
- [18] D. Koller, N. Friedman, S. Džeroski, C. Sutton, A. McCallum, A. Pfeffer, P. Abbeel, M.-F. Wong, D. Heckerman, C. Meek et al., *Introduction to Statistical Relational Learning*, MIT press, 2007.
- [19] N. Lao and W.W. Cohen, Relational retrieval using a combination of path-constrained random walks, *Machine Learning* 81(1) (2010), 53–67. doi: [10.1007/s10994-010-5205-8](https://doi.org/10.1007/s10994-010-5205-8).
- [20] M. Gardner, P. Talukdar, J. Krishnamurthy and T. Mitchell, Incorporating Vector Space Similarity in Random Walk Inference over Knowledge Bases, in: *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing, Doha, Qatar*, ACL, 2014, pp. 397–406. doi: [10.3115/v1/D14-1044](https://doi.org/10.3115/v1/D14-1044).
- [21] M. Gardner and T. Mitchell, Efficient and Expressive Knowledge Base Completion Using Subgraph Feature Extraction, in: *Proceedings of the 2015 Conference on Empirical Methods in Natural Language Processing, Lisbon, Portugal*, ACL, 2015, pp. 1488–1498. doi: [10.18653/v1/D15-1173](https://doi.org/10.18653/v1/D15-1173).
- [22] P. Ristoski and H. Paulheim, RDF2Vec: RDF Graph Embeddings for Data Mining, in: *Proceedings of the 15th International Semantic Web Conference, Kobe, Japan*, Springer, 2016, pp. 498–514. doi: [10.1007/978-3-319-46523-4](https://doi.org/10.1007/978-3-319-46523-4).
- [23] M. Cochez, P. Ristoski, S.P. Ponzetto and H. Paulheim, Global RDF Vector Space Embeddings, in: *Proceedings of the 16th International Semantic Web Conference, Vienna, Austria*, Springer, 2017, pp. 190–207. doi: [10.1007/978-3-319-68288-4_12](https://doi.org/10.1007/978-3-319-68288-4_12).
- [24] Y. Lin, Z. Liu, H. Luan, M. Sun, S. Rao and S. Liu, Modeling Relation Paths for Representation Learning of Knowledge Bases, in: *Proceedings of the 2015 Conference on Empirical Methods in Natural Language Processing, Lisbon, Portugal*, ACL, 2015, pp. 705–714. doi: [10.18653/v1/D15-1082](https://doi.org/10.18653/v1/D15-1082).
- [25] J. Feng, M. Huang, Y. Yang et al., GAKE: Graph Aware Knowledge Embedding, in: *Proceedings of the 26th International Conference on Computational Linguistics, Osaka, Japan*, ACL, 2016, pp. 641–651.
- [26] H. Gao, J. Shi, G. Qi and M. Wang., Triple Context-Based Knowledge Graph Embedding, *IEEE Access* 6 (2018), 58978–58989. doi: [10.1109/ACCESS.2018.2875066](https://doi.org/10.1109/ACCESS.2018.2875066).
- [27] Q. Wang, B. Wang and L. Guo, Knowledge Base Completion Using Embeddings and Rules, in: *Proceedings of the 24th International Joint Conference on Artificial Intelligence, Buenos Aires, Argentina*, AAAI Press, 2015, pp. 1859–1866.
- [28] Z. Wei, J. Zhao, K. Liu, Z. Qi, Z. Sun and G. Tian, Large-scale Knowledge Base Completion: Inferring via Grounding Network Sampling over Selected Instances, in: *Proceedings of the 24th ACM International Conference on Information and Knowledge Management, Melbourne, VIC, Australia*, ACM, 2015, pp. 1331–1340. doi: [10.1145/2806416.2806513](https://doi.org/10.1145/2806416.2806513).
- [29] S. Guo, Q. Wang, L. Wang, B. Wang and L. Guo, Jointly Embedding Knowledge Graphs and Logical Rules, in: *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing, Austin, Texas, USA*, ACL, 2016, pp. 192–202. doi: [10.18653/v1/D16-1019](https://doi.org/10.18653/v1/D16-1019).
- [30] S. Guo, Q. Wang, L. Wang, B. Wang and L. Guo, Knowledge Graph Embedding With Iterative Guidance From Soft Rules, in: *Proceedings of the 32nd AAAI Conference on Artificial Intelligence, New Orleans, Louisiana, USA*, AAAI Press, 2018, pp. 4816–4823.
- [31] W. Zhang, B. Paudel, L. Wang, J. Chen, H. Zhu, W. Zhang, A. Bernstein and H. Chen, Iteratively Learning Embeddings and Rules for Knowledge Graph Reasoning, in: *Proceedings of the 2019 World Wide Web Conference, San Francisco, CA, USA*, ACM, 2019, pp. 2366–2377. doi: [10.1145/3308558.3313612](https://doi.org/10.1145/3308558.3313612).
- [32] V. Gutiérrez-Basulto and S. Schockaert, From Knowledge Graph Embedding to Ontology Embedding? An Analysis of the Compatibility between Vector Space Representations and Rules, in: *Principles of Knowledge Representation and Reasoning: Proceedings of the 16th International Conference, Tempe, Arizona*, AAAI Press, 2018, pp. 379–388.
- [33] M. Chen, Y. Tian, X. Chen, Z. Xue and C. Zaniolo, On2Vec: Embedding-based Relation Prediction for Ontology Population, in: *Proceedings of the 2018 SIAM International Conference on Data Mining, San Diego, CA, USA*, SIAM, 2018, pp. 315–323. doi: [10.1137/1.9781611975321.36](https://doi.org/10.1137/1.9781611975321.36).
- [34] X. Lv, L. Hou, J. Li and Z. Liu, Differentiating Concepts and Instances for Knowledge Graph Embedding, in: *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing, Brussels, Belgium*, ACL, 2018, pp. 1971–1979. doi: [10.18653/v1/D18-1222](https://doi.org/10.18653/v1/D18-1222).
- [35] Z. Sun, Z. Deng, J. Nie and J. Tang, RotatE: Knowledge Graph Embedding by Relational Rotation in Complex Space, in: *Proceedings of the 7th International Conference on Learning Representations, New Orleans, LA, USA*, OpenReview.net, 2019.
- [36] J. Völker and M. Niepert, Statistical Schema Induction, in: *Proceedings of 8th Extended Semantic Web Conference, Her-*

- 1 *aktion, Crete, Greece*, Springer, 2011, pp. 124–138. doi: [10.1007/978-3-642-21034-1_9](https://doi.org/10.1007/978-3-642-21034-1_9).
- 2 [37] J. Noessner, M. Niepert, C. Meilicke and H. Stuckenschmidt, Leveraging Terminological Structure for Object Reconciliation, in: *Proceedings of 7th Extended Semantic Web Conference, Heraklion, Crete, Greece*, Springer, 2010, pp. 334–348. doi: [10.1007/978-3-642-13489-0_23](https://doi.org/10.1007/978-3-642-13489-0_23).
- 3 [38] G. Tjoepfer, M. Knuth and H. Sack, DBpedia Ontology Enrichment for Inconsistency Detection, in: *Proceedings of 8th International Conference on Semantic Systems, Graz, Austria*, ACM, 2012, pp. 33–40. doi: [10.1145/2362499.2362505](https://doi.org/10.1145/2362499.2362505).
- 4 [39] D. Fleischhacker, J. Völker and H. Stuckenschmidt, Mining RDF Data for Property Axioms, in: *Proceedings of the On the Move to Meaningful Internet Systems Conference, Rome, Italy*, Springer, 2012, pp. 718–735. doi: [10.1007/978-3-642-33615-7_18](https://doi.org/10.1007/978-3-642-33615-7_18).
- 5 [40] L. Galarraga, C. Teflioudi, K. Hose and F.M. Suchanek, Fast rule mining in ontological knowledge bases with AMIE+, *The VLDB Journal* **24**(6) (2015), 707–730. doi: [10.1007/s00778-015-0394-1](https://doi.org/10.1007/s00778-015-0394-1).
- 6 [41] H. Gao, G. Qi and Q. Ji, Schema induction from incomplete semantic data, *Intelligent Data Analysis* **22**(6) (2018), 1337–1353. doi: [10.3233/IDA-173514](https://doi.org/10.3233/IDA-173514).
- 7 [42] P.G. Omran, K. Wang and Z. Wang, Scalable Rule Learning via Learning Representation, in: *Proceedings of the 27th International Joint Conference on Artificial Intelligence, Stockholm, Sweden*, AAAI Press, 2018, pp. 2149–2155. doi: [10.24963/ijcai.2018/297](https://doi.org/10.24963/ijcai.2018/297).
- 8 [43] V.T. Ho, D. Stepanova, M.H. Gad-Elrab, E. Kharlamov and G. Weikum, Rule Learning from Knowledge Graphs Guided by Embedding Models, in: *Proceedings of 17th International Semantic Web Conference, Monterey, CA, USA*, Springer, 2018, pp. 72–90. doi: [10.1007/978-3-030-00671-6_5](https://doi.org/10.1007/978-3-030-00671-6_5).
- 9 [44] J. Lehmann and P. Hitzler, Concept learning in description logics using refinement operators, *Machine Learning* **78**(1–2) (2010), 203–250. doi: [10.1007/s10994-009-5146-2](https://doi.org/10.1007/s10994-009-5146-2).
- 10 [45] M. Zhu, Z. Gao, J.Z. Pan, Y. Zhao, Y. Xu and Z. Quan, TBox learning from incomplete data by inference in BelNet+, *Knowledge-Based Systems* **75** (2015), 30–40. doi: [10.1016/j.knsys.2014.11.004](https://doi.org/10.1016/j.knsys.2014.11.004).
- 11 [46] T. Wu, L. Zhang, G. Qi, X. Cui and K. Xu, Encoding Category Correlations into Bilingual Topic Modeling for Cross-Lingual Taxonomy Alignment, in: *Proceedings of the 16th International Semantic Web Conference, Vienna, Austria*, Springer, 2017, pp. 728–744. doi: [10.1007/978-3-319-68288-4_43](https://doi.org/10.1007/978-3-319-68288-4_43).
- 12 [47] Z. Sun, W. Hu and C. Li, Cross-Lingual Entity Alignment via Joint Attribute-Preserving Embedding, in: *Proceedings of the 16th International Semantic Web Conference, Vienna, Austria*, Springer, 2017, pp. 628–644. doi: [10.1007/978-3-319-68288-4_37](https://doi.org/10.1007/978-3-319-68288-4_37).
- 13 [48] H. Zhu, R. Xie, Z. Liu and M. Sun, Iterative Entity Alignment via Joint Knowledge Embeddings, in: *Proceedings of the 26th International Joint Conference on Artificial Intelligence, Melbourne, Australia*, AAAI Press, 2017, pp. 4258–4264. doi: [10.24963/ijcai.2017/595](https://doi.org/10.24963/ijcai.2017/595).
- 14 [49] Z. Sun, W. Hu, Q. Zhang and Y. Qu, Bootstrapping Entity Alignment with Knowledge Graph Embedding, in: *Proceedings of the 27th International Joint Conference on Artificial Intelligence, Stockholm, Sweden*, AAAI Press, 2018, pp. 4396–4402. doi: [10.24963/ijcai.2018/611](https://doi.org/10.24963/ijcai.2018/611).
- 15 [50] M. Chen, Y. Tian, K.-W. Chang, S. Skiena and C. Zaniolo, Co-training Embeddings of Knowledge Graphs and Entity Descriptions for Cross-lingual Entity Alignment, in: *Proceedings of the 27th International Joint Conference on Artificial Intelligence, Stockholm, Sweden*, AAAI Press, 2018, pp. 3998–4004. doi: [10.24963/ijcai.2018/556](https://doi.org/10.24963/ijcai.2018/556).
- 16 [51] Q. Zhang, Z. Sun, W. Hu, M. Chen, L. Guo and Y. Qu, Multi-view Knowledge Graph Embedding for Entity Alignment, in: *Proceedings of the 28th International Joint Conference on Artificial Intelligence, Macao, China*, AAAI Press, 2019, pp. 5429–5435. doi: [10.24963/ijcai.2019/754](https://doi.org/10.24963/ijcai.2019/754).
- 17 [52] M. Zhou, M. Huang and X. Zhu, An Interpretable Reasoning Network for Multi-Relation Question Answering, in: *Proceedings of the 27th International Conference on Computational Linguistics, Santa Fe, New Mexico, USA*, ACL, 2018, pp. 2010–2022.
- 18 [53] W.L. Hamilton, P. Bajaj, M. Zitnik, D. Jurafsky and J. Leskovec, Embedding Logical Queries on Knowledge Graphs, in: *Proceedings of the 32nd Annual Conference on Neural Information Processing Systems*, Curran Associates, Inc., 2018, pp. 2026–2037.
- 19 [54] Y. Zhang, H. Dai, Z. Kozareva, A.J. Smola and L. Song, Variational Reasoning for Question Answering with Knowledge Graph, in: *Proceedings of the 32nd AAAI Conference on Artificial Intelligence, New Orleans, Louisiana, USA*, AAAI Press, 2018, pp. 6069–6076.
- 20 [55] S. Vakulenko, J.D.F. Garcia, A. Polleres, M. de Rijke and M. Cochez, Message Passing for Complex Question Answering over Knowledge Graphs, in: *Proceedings of the 28th ACM International Conference on Information and Knowledge Management, Beijing, China (To Appear)*, ACM, 2019.
- 21 [56] R. Goebel, A. Chander, K. Holzinger, F. Lecue, Z. Akata, S. Stumpf, P. Kieseberg and A. Holzinger, Explainable AI: The New 42?, in: *Proceedings of the International Cross-Domain Conference in Machine Learning and Knowledge Extraction, Hamburg, Germany*, 2018, pp. 295–303. doi: [10.1007/978-3-319-99740-7_21](https://doi.org/10.1007/978-3-319-99740-7_21).
- 22 [57] J. Wang, Z. Wang, D. Zhang and J. Yan, Combining Knowledge with Deep Convolutional Neural Networks for Short Text Classification, in: *Proceedings of the 26th International Joint Conference on Artificial Intelligence, Melbourne, Australia*, AAAI Press, 2017, pp. 2915–2921. doi: [10.24963/ijcai.2017/406](https://doi.org/10.24963/ijcai.2017/406).
- 23 [58] B. Luo, Y. Feng, Z. Wang, S. Huang, R. Yan and D. Zhao, Marrying Up Regular Expressions with Neural Networks: A Case Study for Spoken Language Understanding, in: *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics, Melbourne, Australia*, ACL, 2018, pp. 2083–2093. doi: [10.18653/v1/P18-1194](https://doi.org/10.18653/v1/P18-1194).
- 24 [59] J. Chen, F. Lécué, J. Pan and H. Chen, Learning from Ontology Streams with Semantic Concept Drift, in: *Proceedings of the 26th International Joint Conference on Artificial Intelligence, Melbourne, Australia*, AAAI Press, 2017, pp. 957–963. doi: [10.24963/ijcai.2017/133](https://doi.org/10.24963/ijcai.2017/133).
- 25 [60] J. Chen, F. Lécué, J.Z. Pan, I. Horrocks and H. Chen, Knowledge-Based Transfer Learning Explanation, in: *Principles of Knowledge Representation and Reasoning: Proceedings of the 16th International Conference, Tempe, Arizona*, AAAI Press, 2018, pp. 349–358.

[61] D. Bahdanau, K. Cho and Y. Bengio, Neural Machine Translation by Jointly Learning to Align and Translate, in: *Proceedings of the 3rd International Conference on Learning Representations, San Diego, CA, USA*, 2015.

[62] M.M. Bronstein, J. Bruna, Y. LeCun, A. Szlam and P. Vandergheynst, Geometric Deep Learning: Going beyond Euclidean data, *IEEE Signal Processing Magazine* **34**(4) (2017), 18–42. doi: [10.1109/MSP.2017.2693418](https://doi.org/10.1109/MSP.2017.2693418).

1
2
3
4
5
6
7
8
9
10
11
12
13
14
15
16
17
18
19
20
21
22
23
24
25
26
27
28
29
30
31
32
33
34
35
36
37
38
39
40
41
42
43
44
45
46
47
48
49
50
51

1
2
3
4
5
6
7
8
9
10
11
12
13
14
15
16
17
18
19
20
21
22
23
24
25
26
27
28
29
30
31
32
33
34
35
36
37
38
39
40
41
42
43
44
45
46
47
48
49
50
51