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

Abstract. Knowledge graph, as a backbone of many information systems, has been created to organize the rapidly growing knowledge in a semantical and visualized manner. Symbolic reasoning and statistical reasoning are current mainstream techniques that play important roles in knowledge completion, automatic schema constructing, complex question answering, explanation of AI. 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 problem settings and reasoning tasks, and introduce the key ideas of them. Finally, we re-examine the remaining research problems to be solved and outlook the future directions for hybrid reasoning in Knowledge graphs.

Keywords: Knowledge graphs, Hybrid reasoning, Embedding, Combination

1. Introduction

With the rapid development of Internet technology and Web applications, large amount of data is published online, which contains valuable knowledge. How to organize, represent and analyze these knowledge has attracted much attention. Knowledge graph (KG), as a backbone of many information systems, has been created to organize the rapidly growing knowledge in a semantical and visualized manner. Most of KGs are the directed graphs that compose 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, U.S). Recent years have witnessed rapid growth in KG construction such as DBpedia [2], YAGO [3], NELL [4] and Probase [5], which have become the essential supporters for real applications of Semantic Web.

Although effective in representing structured data, the underlying symbolic nature of triples and their incompleteness still limit the applications of KGs. Knowledge reasoning, which plays an important role in the services of KGs, aims at inferring implicit knowledge to enrich incomplete data and refine their correctness. There are two mainstream techniques for knowledge reasoning. One is symbolic-based reasoning approaches that formalize the problem by a semantic framework and infer the implicit knowledge according to some predefined rules. The other is

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

statistical-based reasoning approaches that try to find one suitable statistic models to fit the samples and predict the expected probability or similarity about test ones.

Unfortunately, no single method can be competent for knowledge reasoning perfectly. Symbolic reasoning is often based on either rules or schematic knowledge, which is hard to obtain. Relatively, statistical reasoning draws imprecise conclusions and is often data-driven so that it is hard to provide the human-centric explanation. Therefore, more researchers tried to combine their advantages together, and obtained some encouraging performances in related tasks such as knowledge completion [6, 7], knowledge alignment [8, 9], query answer [10, 11] and so on. There exist various combination strategies tailored for different tasks. Some of them 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. Some of them employ the idea of statistical reasoning (e.g., continuous vectors or matrices) to soften symbolic reasoning in order to be compatible with objective facts well. In addition, some works unify them together to achieve the goal of completing multiple tasks (e.g, rule learning and link prediction) simultaneously.

So far, there is no systematical and in-depth survey on hybrid reasoning methods for various tasks of KGs. In this paper, we summarize the latest research progress of hybrid reasoning techniques in knowledge graphs and look forward to the future development direction 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 hybrid reasoning techniques for various tasks of KG. Finally, we re-examine the remaining research challenges and outlook the 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 motivation of hybrid reasoning in a knowledge graph. So far, some people have tried to provide a formal definition of a knowledge graph [12, 13]. 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 [14]:

- 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 in knowledge graphs which is the basic storage unit 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*">.

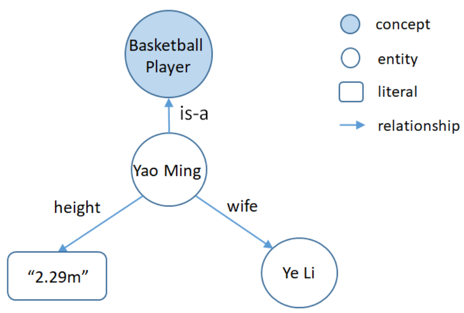

Fig. 1. An example for a part of a knowledge graph

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 contain a large number of factual knowledge and a small number of schematic knowledge. For example, the well-known knowledge graph DBpedia 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 the subclassOf relations. There exist some knowledge graphs which consist of a large number of schematic knowledge, such as SNOMED CT[1].

Knowledge graph has its logical foundations based on ontological languages, such as Resource Description Framework (RDF)[2] and Ontology Web Language (OWL)[3], which are W3C recommended languages. RDF is a graph data model for describing resources on the Web and to enable data exchange and sharing, it is 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 is 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 work on hybrid reasoning in a knowledge graph and present some challenging problems for future work.

[0]https://wiki.dbpedia.org
[1]https://bioportal.bioontology.org/ontologies/SNOMEDCT
[2]https://www.w3.org/RDF/
[3]https://www.w3.org/OWL/

3. Methodology

In this section, we roughly categorize the hybrid reasoning techniques of KGs into six groups: statistical relational learning, schema induction, schematic knowledge embedding, knowledge alignment, multi-hop reasoning for query answer and other hybrid reasoning methods. Next, we review these research efforts as follows.

3.1. Statistical relational learning

Statistical relational learning (SRL) attempts to represent, reason and learn in domains with complex relational and rich probabilistic structure [15]. With the rapid growth in KGs, path ranking algorithms (PRA) [16] and knowledge graph embedding (KGE) [1] become two typical representatives of SRL, and have shown some efficiency of applications.

PRA is a random walk inference technique, which first proposed for discovering complex path features of relational data [16]. The key idea of PRA is employing the paths that connect two entities as features to predict potential relations between them. For example, $\langle$bornIn, 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.

Knowledge graph embedding embeds components of a KG including entities and relations into continuous vector spaces to preserve the inherent structure of the KG [1]. There are mainly two types of embedding models. One is transnational distance models, which exploit distance-based scoring functions and measure the plausibility of a fact as the distance between two entities such as TransE [6]. The other 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.

As triples in KGs are not independent, so the inter-relations of each triple should not be ignored, which can give the power in knowledge reasoning. PTransE [17] is an extending model of TransE to model 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 relation paths are reliable, they designed a path-constraint resource allocation algorithm to measure the reliability of relation paths and represented relation paths via semantic composition of relation embeddings. GAKE [18] defined three types of graph context which contains 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. Furthermore, Gao et al. [19] proposed a triple context-based method called TCE for knowledge graph embedding. TCE takes two structured information of a 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.

Assertions of relations contain rich background information (e.g., domain, range) that are widely treated as constraint rules in KGs. Wang et al. [20] 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. [21] combined rules and embedding models via Markov logic networks [15], in which they incorporated the similarity priori generated by embedding-based models into inferring and designed the grounding network sampling to promote the inference precision.

3.2. Schema induction

Existing KGs often contain a large number of triples but lack schematic knowledge like disjointness axioms and subclassOf axioms. It brings a difficulty to infer implicit information, deal with the heterogeneous problem for ontology mapping and object reconciliation tasks, and detect or resolve the contradictions [22–25]. Therefore, producing schematic knowledge to enrich existing 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 [24, 26] defined some association rule patterns to generate various kinds of axioms and performed inconsistency handling for ontology construction by enriching an original schema incrementally. Considered the open world assumption adopted by KGs, the work in [27] adopted partial completeness assumption to generate counterexamples for rules and redefines support or confidence. Its extension AMIE+ [28] further improved the precision by using type hierarchy and joint reasoning when learning association rules. The work in [29] generated rules with AMIE+. It obtained the rules of interest for learning inverse and symmetric axioms which could be extended by applying the predefined reasoning rules. Inspired by these methods, the method given in [30] exploited a type inference algorithm and defined a mining model with the probabilistic type assertions to deal with noisy negative examples. Their method can generate high-quality disjointness axioms and subclassOf axioms. To improve the scalability of the rule-based methods, the work in [31] introduced a new sampling algorithm and the embedding representations of arguments. Both of them could guide the extraction of rules. Similarly, the work in [32] employed embedding models and iteratively extracted rules by utilizing probabilistic representations of missing facts and relying on feedback from a precomputed embedding model.

The other main category combines machine learning techniques with logical reasoning. The work in [33] used inductive logic programming, which integrated machine learning with logic programming, and defined an $\mathcal{ALC}$ downward refinement operator for learning concept descriptions. This operator was extended in [34] that could learn more expressive schematic knowledge like cardinality restrictions. In [25], a statistical method was proposed to extract domain and range of a property. The vector space model from information retrieval was applied to extract class disjointness. After the extraction finished, consistency checking was performed in parallel based on predefined inconsistency patterns. A light-weight method presented in [35, 36] obtained schema and data information via SPARQL, and then applied machine learning algorithms to generate nearly all kinds of axioms. After that, a logical reasoner could be applied for inferencing those implicit knowledge. The work in [37] integrated the probabilistic inference capability of Bayesian networks with the logical formalism to learn subclassOf and disjointness relations. It used logical rules for generating more complex axioms and dealing with inconsistency during the construction of KGs.

3.3. Schematic knowledge embedding

Schematic knowledge, as a critical component of KGs, defines logical axioms based on concepts to support for eliminating heterogeneity, integration, and reasoning over KGs. Recently, there two kinds of approaches that try to encode schematic knowledge to enhance the performances of embedding models. One is treating schematic knowledge as logical rules and incorporating them to obtain better embedding. The other focuses on preserving the logical properties of axioms.

The incorporated logical rules are usually represented as first-order horn clauses e.g., $\forall x, y$ (x, Capital-Of, y) $\rightarrow$ (x, Located-In, y) stating that any two entities linked by the relation Capital-Of should also be satisfied with the relation Located-In. Such logical rules contain rich background information and are widely defined in ontologies. Guo et al. [38] proposed a joint model, called KALE, which embedded 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 [39]. After that, they improved this model further, referred to as RUGE [40], which could learn simultaneously from labeled triples, unlabeled triples, and soft rules in an iterative manner. Zhang et al. [41] proposed a novel framework called IterE for alleviating of sparsity entities in KGs. It could iteratively learn embeddings and logical rules, in which rules were learned from embeddings with proper pruning strategy, and embeddings were learned from existing triples and new triples inferred by rules. In addition, Gutiérrez-Basulto and Schockaert argued that existing combined models might not represent expressive classes of rules sufficiently, and proposed an approach based on convex-regions [42]. With the help of defined convex-regions, KGs restricted to the quasi-chained existential rules could be faithfully encoded in most cases.

Another type of embedding methods has been proposed for the embedding of schematic knowledge in a simple ontology language called RDF Schema (or RDFS). On2Vec [43] employed translation-based embedding method for ontology population, 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 [44] encoded instances as vectors and concepts as spheres so that they can preserve the transitivity of isA relations.

3.4. 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 the combination of 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. [8] proposed a bilingual topic model, called Bilingual Biterm Topic Model (BiBTM). After obtained the candidates' alignment 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 performances of proposed models by merging explicit category correlations including co-occurrence correlation and structural correlation [45].

In addition, there exist some works that employ embedding-based ideas [6] for entity alignment (EA) among knowledge graphs. MTransE [9] separately trained the entity embeddings of two KGs and designed different techniques to represent cross-lingual transitions including axis calibration, translation vectors and linear transformations. JAPE [46] learned the embeddings of two KGs in a unified space and leveraged attributes of triples to refine entity embeddings. To solve the lack of prior alignment, IPTransE [47] and BootEA [48] employed an iterative process and designed several sophisticated strategies based on the structure of KG to refine the new alignment. Chen et al. [49] proposed semi-supervised cross-lingual learning method, called KDCoE, which co-trained multilingual KG embeddings and the embeddings of entity descriptions. Considered the identity information in the prior alignment could not be efficiently propagated from one KG to another, Guo et al. [50] proposed a recurrent skipping network (RSN) for entity alignment, which leveraged biased random walk sampling for generating long paths across KGs.

3.5. Multi-hop reasoning for query answer

Question answering (QA) is a hot topic that has recently been facilitated by large-scale knowledge bases. However, due to the variety and complexity of language and knowledge, question answering over knowledge bases (KBQA) is still a challenging task, especially in multi-hop relation QA.

There are two typical categories of multi-relation questions, a path question [51] and a conjunctive question [52]. A path question contains only one topic entity and its answer can be found by walking down an answer path consisting of a few relations and intermediate entities. A conjunctive question is a question that contains more than one subject entity and the answer can be obtained by the intersection of results from multiple path queries. At present, semantic parsing models [10, 53] and embedding-based models [11, 54] tailored for QA are not adequate to handle multi-hop QA because of heavy data annotations and reasoning ability. Therefore, recent works utilized hybrid ideas to improve the performances and make these results explainable.

Zhang et al. [55] proposed a probabilistic modeling framework based on end-to-end QA system, which could simultaneously handle uncertain topic entity and multi-hop reasoning. They introduced a new propagation architecture over KG so that logical inference could be performed in the probabilistic model. Zhou et al. [51] designed an interpretable reasoning network (IRN), which employed an interpretable hop-by-hop reasoning process for question answering. IRN could dynamically decide which part of an input question should be analyzed at each hop, and predict a relation corresponding to the parsed results. Compared with existing methods, the intermediate entities and relations predicted by IRN could construct traceable reasoning paths to reveal how the answer was derived. Hamilton et al. [52] introduced a framework to efficiently make predictions about conjunctive logical queries. They embedded graph nodes in a low-dimensional space and represented logical operators (i.e., projection operator and intersection operator) as learned geometric operations. Moreover, they further demonstrated how to map a practical subset of logic to efficient geometric operations in an embedding space.

3.6. Other hybrid reasoning methods

Other hybrid reasoning methods focus on boosting the performances of NLP tasks. Most of them merge the symbolic information (e.g., knowledge graph structure) into the statistic-based methods and provide some human-centric explanation [56].

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

Chen et al. [58] exploited the semantics of data streams interpreted in ontologies to tackle the problem of concept drift, in which semantic reasoning and machine learning are combined by revisiting features embeddings as semantic embeddings. Such embeddings were exploited in a context of supervised stream learning that was robust to concept drifts. Moreover, they explored an ontology-based knowledge representation and reasoning framework for human-centric transfer learning explanation [59]. It modeled a learning domain in transfer learning with expressive OWL and complemented the learning domain with the prediction task-related common sense knowledge. They further designed a correlative reasoning algorithm to infer three kinds of explanatory evidence for explaining a positive feature or a negative transfer from one learning domain to another.

4. Conclusion and future direction

Hybrid reasoning in knowledge graphs plays an important role in knowledge completion, automatic schema constructing, complex question answering, explanation of AI, ect. However, it is still a new topic and lacks a survey of existing methods for it. In this paper, we presented the first work on survey of methods for hybrid reasoning in knowledge graphs. We argued the necessity of combination symbolic reasoning and statistical reasoning. More importantly, we categorized existing methods based on problem settings and reasoning tasks, respectively, and further introduced the key ideas of them. Although there have been many methods for hybrid reasoning in knowledge graphs, there are still many problems to be solved in the future. We list some of the problems in the following.

– **Statistical relational learning**: Taking relational paths into account can significantly improve the discrimination of relational learning and system performances. However, existing models are still some preliminary attempts at modeling relational paths. There still exist many investigations in the reliability measure and semantic composition of relational paths to be done.

- **Schema induction**: Horn rules are one of the most simple schemas that can be learned from KGs. It is still challenging for existing methods to extend the set of rules to more complex non-monotonic ones such as existential variables or disjunctions in rule heads. In addition, the sparse long-tail relations still need to be considered, which are actually more common in KGs.
- **Schematic knowledge embedding**: RDF Schema is a simple ontology language. There exist several challenges for embedding the schematic knowledge described by expressive OWL such as preserving its complex semantic properties (e.g., symmetry, inversion, composition) simultaneously.
- **Knowledge alignment**: It is worth to consider combing the methods of CLTA and EA together because most KGs consist of taxonomy and entities. In addition, these models still can merge some senior symbolic reasoning techniques (e.g., incoherent checking) during the training process.
- **Multi-hop reasoning for QA**: The frameworks of multi-hop reasoning are still limited by some types of queries so that they cannot handle arithmetic operation or logical queries with negation or disjunction. Integrating attention mechanism [60] and employing graph neural networks [61] to incorporate richer feature information on nodes and edges will be two promising directions.

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.

[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.

[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: *Proc. of ISWC*, Springer, 2016, pp. 177–185.

[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.

[5] W. Wu, H. Li, H. Wang and K.Q. Zhu, Probase: A probabilistic taxonomy for text understanding, in: *Proc. of SIGMOD*, ACM, 2012, pp. 481–492.

[6] R. Socher, D. Chen, C.D. Manning and A. Ng, Reasoning with neural tensor networks for knowledge base completion, in: *Proc. of NeurIPS*, 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: *Proc. of ICML*, Omnipress, 2011, pp. 809–816.

[8] T. Wu, G. Qi, H. Wang, K. Xu and X. Cui, Cross-lingual taxonomy alignment with bilingual biterm topic model, in: *Proc. of AAAI*, AAAI Press, 2016, pp. 287–293.

[9] M. Chen, Y. Tian, M. Yang and C. Zaniolo, Multilingual knowledge graph embeddings for cross-lingual knowledge alignment, in: *Proc. of IJCAI*, AAAI Press, 2017, pp. 1511–1517.

[10] P. Pasupat and P. Liang, Compositional Semantic Parsing on Semi-Structured Tables, in: *Proc. of ACL*, 2015, pp. 1470–1480.

[11] A. Bordes, J. Weston and N. Usunier, Open Question Answering with Weakly Supervised Embedding Models, in: *Proc. of ECML*, Springer, 2014, pp. 165–180.

[12] L. Ehrlinger and W. Wöß, Towards a Definition of Knowledge Graphs, in: *Proc. of SEMANTiCS*, CEUR-WS.org, 2016.

[13] 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.

[14] H. Paulheim, Knowledge graph refinement: A survey of approaches and evaluation methods, *Semantic Web* **8**(3) (2017), 489–508.

[15] 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.

[16] N. Lao and W.W. Cohen, Relational retrieval using a combination of path-constrained random walks, *Machine Learning* **81**(1) (2010), 53–67.

[17] Y. Lin, Z. Liu, H. Luan, M. Sun, S. Rao and S. Liu, Modeling Relation Paths for Representation Learning of Knowledge Bases, in: *Proc. of EMNLP*, 2015, pp. 705–714.

[18] J. Feng, M. Huang, Y. Yang et al., GAKE: graph aware knowledge embedding, in: *Proc. of COLING*, 2016, pp. 641–651.

[19] H. Gao, J. Shi, G. Qi and M. Wang., Triple Context-Based Knowledge Graph Embedding, *IEEE Access* **6** (2018), 58978–58989.

[20] Q. Wang, B. Wang and L. Guo, Knowledge base completion using embeddings and rules, in: *Proc. of IJCAI*, AAAI Press, 2015, pp. 1859–1866.

[21] 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: *Proc. of CIKM*, ACM, 2015, pp. 1331–1340.

[22] J. Noessner, M. Niepert, C. Meilicke and H. Stuckenschmidt, Leveraging Terminological Structure for Object Reconciliation, in: *Proc. of ESWC*, Springer, 2010, pp. 334–348.

[23] A. Nolle, C. Meilicke, M.W. Chekol, G. Nemirovski and H. Stuckenschmidt, Schema-Based Debugging of Federated Data Sources, in: *Proc. of ECAI*, 2016, pp. 381–389.

[24] J. Völker and M. Niepert, Statistical Schema Induction, in: *Proc. of ESWC*, Springer, 2011, pp. 124–138.

[25] G. Toepper, M. Knuth and H. Sack, DBpedia ontology enrichment for inconsistency detection, in: *Proc. of SEMANTICS*, ACM, 2012, pp. 33–40.
[26] D. Fleischhacker, J. Völker and H. Stuckenschmidt, Mining RDF Data for Property Axioms, in: *Proc. of OTM Conference*, 2012, pp. 718–735.
[27] L.A. Galárraga, C. Teflioudi, K. Hose and F. Suchanek, AMIE: Association Rule Mining Under Incomplete Evidence in Ontological Knowledge Bases, in: *Proc. of WWW*, ACM, 2013, pp. 413–422.
[28] L. Galarraga, C. Teflioudi, K. Hose and F.M. Suchanek, Fast rule mining in ontological knowledge bases with AMIE+, *The VLDB Journall* **24**(6) (2015), 707–730.
[29] R. Irny and P.S. Kumar, Mining Inverse and Symmetric Axioms in Linked Data, in: *Proc. of JIST*, Springer, 2017, pp. 215–231.
[30] H. Gao, G. Qi and Q. Ji, Schema induction from incomplete semantic data, *Intelligent Data Analysis* **22**(6) (2018), 1337–1353.
[31] P.G. Omran, K. Wang and Z. Wang, Scalable Rule Learning via Learning Representation, in: *Proc. of IJCAI*, AAAI Press, 2018, pp. 2149–2155.
[32] V.T. Ho, D. Stepanova, M.H. Gad-Elrab, E. Kharlamov and G. Weikum, Rule Learning from Knowledge Graphs Guided by Embedding Models, in: *Proc. of ISWC*, Springer, 2018, pp. 72–90.
[33] J. Lehmann and P. Hitzler, A Refinement Operator Based Learning Algorithm for the $\mathcal{ALC}$ Description Logic, in: *Proc. of ILP*, 2007, pp. 147–160.
[34] J. Lehmann and P. Hitzler, Concept learning in description logics using refinement operators, *Machine Learning* **78**(1–2) (2010), 203–250.
[35] L. Bühmann and J. Lehmann, Universal OWL Axiom Enrichment for Large Knowledge Bases, in: *Proc. of EKAW*, Springer, 2012, pp. 57–71.
[36] L. Bühmann and J. Lehmann, Pattern Based Knowledge Base Enrichment, in: *Proc. of ISWC*, Springer, 2013, pp. 33–48.
[37] 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** (2014), 30–40.
[38] S. Guo, Q. Wang, L. Wang, B. Wang and L. Guo, Jointly embedding knowledge graphs and logical rules, in: *Proc. of EMNLP*, 2016, pp. 192–202.
[39] P. Hájek, *Metamathematics of fuzzy logic*, Vol. 4, Springer, 2013.
[40] S. Guo, Q. Wang, L. Wang, B. Wang and L. Guo, Knowledge graph embedding with iterative guidance from soft rules, in: *Proc. of AAAI*, AAAI Press, 2018, pp. 4816–4823.
[41] 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, *CoRR* **abs/1903.08948** (2019).
[42] 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: *Proc. of KR*, AAAI Press, 2018, pp. 379–388.
[43] M. Chen, Y. Tian, X. Chen, Z. Xue and C. Zaniolo, On2Vec: Embedding-based Relation Prediction for Ontology Population, in: *Proc. of SIAM*, SIAM, 2018, pp. 315–323.
[44] X. Lv, L. Hou, J. Li and Z. Liu, Differentiating Concepts and Instances for Knowledge Graph Embedding, in: *Proc. of EMNLP*, 2018, pp. 1971–1979.
[45] T. Wu, L. Zhang, G. Qi, X. Cui and K. Xu, Encoding category correlations into bilingual topic modeling for cross-lingual taxonomy alignment, in: *Proc. of ISWC*, Springer, 2017, pp. 728–744.
[46] Z. Sun, W. Hu and C. Li, Cross-lingual entity alignment via joint attribute-preserving embedding, in: *Proc. of ISWC*, Springer, 2017, pp. 628–644.
[47] H. Zhu, R. Xie, Z. Liu and M. Sun, Iterative Entity Alignment via Joint Knowledge Embeddings., in: *Proc. of IJCAI*, AAAI Press, 2017, pp. 4258–4264.
[48] Z. Sun, W. Hu, Q. Zhang and Y. Qu, Bootstrapping Entity Alignment with Knowledge Graph Embedding., in: *Proc. of ISWC*, Springer, 2018, pp. 4396–4402.
[49] 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: *Proc. of IJCAI*, AAAI Press, 2018, pp. 3998–4004.
[50] L. Guo, Z. Sun, E. Cao and W. Hu, Recurrent Skipping Networks for Entity Alignment, *CoRR* **abs/1811.02318** (2018).
[51] M. Zhou, M. Huang and X. Zhu, An Interpretable Reasoning Network for Multi-Relation Question Answering, in: *Proc. of COLING*, 2018, pp. 2010–2022.
[52] W.L. Hamilton, P. Bajaj, M. Zitnik, D. Jurafsky and J. Leskovec, Embedding Logical Queries on Knowledge Graphs, in: *Proc. of NeurIPS*, 2018, pp. 2026–2037.
[53] A. Abujabal, R.S. Roy, M. Yahya and G. Weikum, QUINT: Interpretable Question Answering over Knowledge Bases, in: *Proc. of EMNLP*, 2017, pp. 61–66.
[54] 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: *Proc. of ACL*, 2017, pp. 221–231.
[55] Y. Zhang, H. Dai, Z. Kozareva, A.J. Smola and L. Song, Variational Reasoning for Question Answering with Knowledge Graph, in: *Proc. of AAAI*, AAAI Press, 2018, pp. 6069–6076.
[56] R. Goebel, A. Chander, K. Holzinger, F. Lecue, Z. Akata, S. Stumpf, P. Kieseberg and A. Holzinger, Explainable AI: the new 42?, in: *Proc. of CD-MAKE*, Springer, 2018, pp. 295–303.
[57] J. Wang, Z. Wang, D. Zhang and J. Yan, Combining Knowledge with Deep Convolutional Neural Networks for Short Text Classification, in: *Proc. of IJCAI*, AAAI Press, 2017, pp. 2915–2921.
[58] J. Chen, F. Lécué, J. Pan and H. Chen, Learning from ontology streams with semantic concept drift, in: *Proc. of IJCAI*, AAAI Press, 2017, pp. 957–963.
[59] J. Chen, F. Lecue, J.Z. Pan, I. Horrocks and H. Chen, Knowledge-based transfer learning explanation, in: *Proc. of KR*, AAAI Press, 2018, pp. 349–358.
[60] D. Bahdanau, K. Cho and Y. Bengio, Neural Machine Translation by Jointly Learning to Align and Translate, in: *Proc. of ICLR*, 2015.
[61] 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.