Semantic and Efficient Symbolic Learning over Knowledge Graphs

Disha Purohit[0000−0002−1442−335X]

Leibniz Universität Hannover, Germany
TIB Leibniz Information Centre for Science and Technology, Hannover, Germany
d.purohit@stud.uni-hannover.de

Abstract. In recent years, the rise of large Knowledge Graphs (KGs), which capture knowledge in machine-driven formats, has arisen broadly. KGs are the convergence of data and knowledge, and may be incomplete due to the Open World Assumption (OWA). Inductive Logic Programming (ILP) is a popular traditional approach for mining logical rules to complete the KGs. ILP approaches derive logical rules from ground facts in knowledge bases. Deducing new information or adding missing information to the KGs, identifying potential errors, and understanding the data more substantially can be accomplished by mining logical rules. Inference can be used to deduce new facts and complete KGs. To discover meaningful insights, traditional rule mining approaches first ignore axiomatic systems defining the semantics of the predicates and classes available in KGs. Second, most rule miners measure the impact of mined rules in terms of correlation rather than causation, and they are overwhelmed by the volume of data. Finally, existing frameworks implement blocking methods that require the processing of complete KGs to generate the mined rules. In this Ph.D. proposal, an outline of a rule-mining model explicitly tailored to mine Horn rules encapsulating semantics on top of KGs is reported. Additionally, the rule-mining approach is based on reliably estimating the cause-effect relationships and discovering new facts in the KGs considering data and metadata. Our approach follows an iterative process to inductively mine rules incorporating semantics to enhance completeness. Our experimental results suggest that by combining entailment regimes and querying KGs on demand, our approach outperforms the state-of-the-art in terms of accuracy. A publicly available Jupyter notebook that executes a demonstration is available [1].

Keywords: Rule Mining · Causality · KGs · RDFS · OWL

1 Introduction

Knowledge Graphs (KGs) encode real-world knowledge as factual statements; nodes represent entities and edges define relationships between the entities. KGs

Category: Middle Stage Ph.D.

[1] https://mybinder.org/v2/gh/SDM-TIB/DIGGER-ESWC2023Demo/HEAD?labpath=Mining%20Symbolic%20Rules%20To%20Explain%20Lung%20Cancer%20Treatments.ipynb

are often created from heterogeneous sources that can be highly diverse in terms of structure and granularity[11]. Existing KGs cover many different domains in order to serve the research and industrial communities[2]. Numerous contributions in the Semantic Web community have addressed the open research challenge of mining Horn rules from ground statements.
Knowledge Discovery in Databases (KDD) is defined as the extraction of potentially useful information from a large volume of data, where the information is implicit. *Association rule mining* [2] is one of the most popular methods of mining rules in the relational domain. Various other approaches [3,9,16] also mine rules based on the co-occurrence of items present in the relational databases. For example, rules like "If a client bought beer and wine, she/he also bought aspirin" can be uncovered using association rule mining. The ratio of instances where beer and wine were purchased along with aspirin corresponds to the rule's confidence. Rule mining over relational databases follows the Closed World Assumption (CWA), i.e., it cannot predict items that are not present in a database. Inductive Logic Programming (ILP) is used in semantic rule mining to extract information in a machine-readable format from Knowledge Bases (KBs). Existing ILP approaches derive logical rules from KBs. Due to the large volume of data and frequent assumption of incomplete data, rule mining over KBs is challenging and dedicated techniques have been proposed to address these issues. Exemplary rule mining approaches (e.g., AMIE[8], AnyBURL[14], AMIE+ [6], and [7,17]) are devised to operate under KBs that consider OWA. However, these approaches are still not tailored to deal with KGs encompassing semantics. Additionally, to mine rules, large KBs must be downloaded in a local system, and the computation of the mining process is done in a blocking fashion, i.e., all the data need to be uploaded/processed to produce results rather than continuously generating rules; thus, negatively affecting scalability in rule mining processes.
Humans are able to infer knowledge from data based on a set of general rules or by knowing the context of available data. This knowledge inferred by humans can be referred to as 'commonsense knowledge, or 'domain knowledge'. For instance, a KB contains a fact that Y has a father X. Then, humans can easily infer that the gender of X is "male" and that Y is the child of X. On the contrary, machines do not have any prior knowledge or information to make inferences over the provided data. Deductive methods are used to infer new facts known as *entailed facts* from existing facts in knowledge graphs using a set of rules often referred to as *entailment regimes.* Further, these entailed facts are used by *Inductive methods* to derive new logical rules. Inductive knowledge is knowledge acquired by generalizing patterns from a given set of input observations. Mining Horn rules is conducted using inductive learning to create a symbolic model, i.e., a set of rules or axioms. Entailment regimes describe the relationship between the statements that are true when one statement locally follows from one or more statements [4]. Machines can apply deductions on top of data graphs and by applying entailment regimes efficiently.

[2] The terms Knowledge Bases (KBs) and Knowledge Graphs (KGs) are used interchangeably

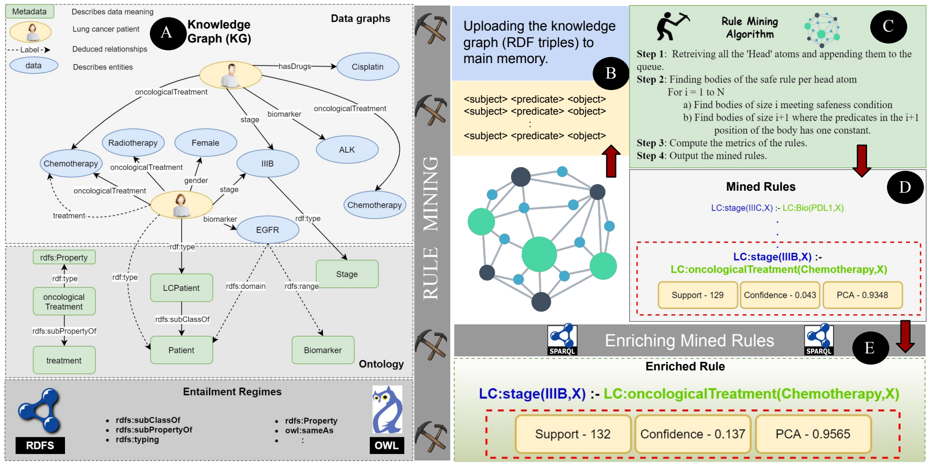

Fig. 1: **Motivating Example:** Usage of a KG comprising ontologies and their entailment regimes to be considered during rule mining. Naive approaches perform rule mining on graphs uploaded in the main memory. SPARQL queries and entailment regimes empower the semantics of KGs, enriching the mined rules.

The term *Ontology* refers to the concrete and more formal representation of the data present in data graphs. The *Web Ontology Language* [10] and *RDFS* [5] are the two most popular ontology languages recommended by W3C and compatible with RDF graphs. Facts in KGs are usually divided into *A-Box* and *T-Box*. The *A-Box* facts are all the instances of a KG that represent the data graph. Complementary, *T-Box* includes the definition of classes, properties, and hierarchies, which represent the ontological part of KGs. We focus on both *A-Box* and *T-Box* to mine rules that consider the semantics of a KG. Horn clauses are expressed in *IF-THEN-* style consequences over KG predicates. Our preliminary results reveal the key role of semantics in the accuracy of rule mining systems.

1.1 Motivating Example

Our work is motivated by the lack of exploitation of semantics in rule mining approaches over KGs. Nevertheless, state-of-the-art techniques provide rules that are mined over data graphs; ignoring the semantics and meaning of the entities in KGs. The goal of this work is to mine rules from which true missing facts can be predicted; completing, thus, KGs with accurate predictions which take into account KGs and entailment regimes. Figure 1 depicts, with a motivating example, the challenges present in a rule mining process over KGs. Input from the KG is collected **A**, representing lung cancer patients and all the related information about those patients, i.e., gender, age, cancer stage, oncological treatment, and mutations. The ontology layer in the KG represents the unified schema of the lung cancer KG. The entailment regime layer shows the RDFS and OWL entailment regimes. Further, **B** shows that for the execution of the rule mining algorithm, the state-of-the-art techniques require input in the form of RDF triples, i.e.,⟨`subject, predicate, object`⟩. For instance, ⟨`2842697, smokingHabit, CurrentSmoker`⟩, ⟨`2842697, stage, IVB`⟩, and ⟨`2842697, gender, Male`⟩ are uploaded to a local system. This poses one of the limitations of naive approaches (AMIE[8], AnyBURL[14]) impacting scalability to large KGs.

C shows the implementation of rule mining algorithm. Rule mining algorithms (AMIE[8] and AnyBURL[14]) implement *blocking* processes, i.e., to mine rules all the data needs to be uploaded. This type of algorithm lacks the accessibility of KGs via Web interfaces, e.g., SPARQL endpoints, which cannot be utilized unless downloaded locally. Our approach overcomes the limitation of scalability by taking as input the SPARQL endpoints and queries to traverse the KGs. For example, **D** shows that the mined rules generated from the above-mentioned algorithm: `LC:stage(IIIB, X)` ⇒ `LC:oncologicalTreatment(Chemotherapy, X)` , clearly states that the metadata encoded in the KGs ontology is not considered, i.e., ⟨`oncologicalTreatment, rdfs:subPropertyOf, treatment`⟩.
The rules mined by the naive approaches demonstrate that the lung cancer patient in stage *IIIB* is more likely to receive oncological treatment *Chemotherapy*.

E illustrates the process of enrichment of the mined rules by applying the entailment regimes to the above mined rules. This step helps to derive new insights from the KGs. In contrast to naive approaches, our approach considers `rdfs:subPropertyOf` entailment regime, and as observed the metrics *Support*, *Confidence* and *PCA Confidence* increases; new facts were inferred, and added to the KGs that lead to the increased metrics values.
Contributions: A rule mining system that is inherently designed to work under the OWA and is efficient enough to handle KGs is proposed in this Ph.D. proposal. More specifically, the following are the contributions of this proposal:

1. Rule mining system devised for KGs empowered with semantics.
2. Novel query and mining techniques to improve scalability, and generate rules iteratively while avoiding blocking data processing.
3. Exploiting cause-effect relationships to enhance meaningful insights.

2 State of the Art

Rule mining methods have gained considerable attention for the past few years, but the existing methods are based on mostly association rule mining [2] or inductive logic programming. This section covers the state-of-the-art techniques for KGs, that perform rule mining over observational data and knowledge bases.

2.1 Mining Rules in Relational Databases

Association rule mining (ARM) is a rule mining approach in the relational domain, and it is implemented on the closed world assumption (CWA). Association rule mining aims to recognize patterns and concurrent occurrences in the database. It discovers relationships among the entities present in the database. Apriori algorithm [18] is a well-known association rule mining[2] approach. It shows how frequently the item appears within the database. Association rule mines frequent patterns of data occurring, using the criteria 'support' and 'confidence' as metrics. It is used in the well-known 'Market Basket analysis' [12]. The mined rules are of the form *wine, beer* $\implies$ *aspirin*, implying that people who purchased wine and beer also purchased aspirin. However, these are not the kinds of rules we aim to discover in this paper. We intend to mine Horn rules. In the work, [1] association rules and frequency analysis are used to identify and classify common misuse patterns for relations over DBpedia. In contrast to our work, this approach mines association rules based on the co-occurrence of values rather than logical rules. Secondly, correlation is a statistical measure that describes the magnitude and direction of a relationship between two or more variables. A correlation between variables, on the other hand, does not imply that a change in one variable is the cause of a change in the values of the other variable. Causation denotes that one event is the result of the other event's occurrence, i.e., causal relationships among events. In this work, we aim to mine rules that encode cause-effect relationships.

2.2 Mining Rules in Knowledge Graphs

AMIE[8] is a rule-mining approach that follows Inductive Logic Programming (ILP) and aims to mine logical rules. When dealing with an incomplete KG, it is not immediately clear how to define negative edges. A common heuristic for a KG is to use a *Partial Completeness Assumption (PCA)*, i.e., a negative edge should be true if it is derived from a Horn clause that partially defines its completeness. AMIE mines rules on large KBs by reducing the search space. The logical rules mined by AMIE are in the form of Horn clauses. AMIE uses several metrics to prune the rules obtained by mining in order to avoid the generation of an exponential amount of irrelevant rules. Various metrics are used to evaluate the quality of rules mined by AMIE, *Head Coverage* that measures the ratio of known true facts that are implied by the rule. *Std Confidence* of the rule is the ratio of all its predictions that are present in the KGs. Lastly, in order to generate heuristic-based negative edges in the KBs, AMIE operates under *PCA*. Different versions of AMIE were AMIE+[6] and AMIE3[13]. The newer versions of AMIE claim that it speeds up the process of mining rules even faster. AMIE3 integrates new pruning strategies and many more advancements. However, AMIE lacks scalability as it follows a blocking approach to produce results. *(AnyBURL)*[14] learns logical rules. They focus on a path ranking algorithm that helps them to learn a subset of the rules. Similar to AMIE, they also mine negative edges in order to better complete the knowledge graphs. In contrast to AMIE, AnyBURL learns rules from knowledge graphs from the bottom up, whereas AMIE mines rule from the top down. Above discussed approaches are not tailored to deal with large knowledge graphs with semantics under the Open World Assumption. In contrast to our approach, the process of mining rules with AMIE[8] and AnyBURL[14] follows a blocking process, impacting

scalability. Furthermore, by ignoring the semantics of KGs, these techniques fall short of generating more meaningful rules.
In some of the closely related work, Simonne et al. [17] mine two types of differential causal rules, gradual and categorical rules. Gradual rules deal with mining rules over numerical values or entities and categorical rules deal with categorical values. For example, the number of treatments received by cancer patients comes under gradual rules and the type of treatments are classified into categorical rules. Another contribution of this paper is to use a community detection algorithm to compute the similarity between the units of interest. Also, they have defined a metric called *Causal ratio* which is inspired by the odds ratio to evaluate the potential causal rules. Discovering causality in knowledge graphs is a wide area of research performed scarcely. Traditional approaches attempt to detect causal relationships between variables by implementing a probabilistic relational model using Bayesian networks [15] following Judea Pearl's approach. Our approach, in contrast, considers entailment regimes in addition to the semantics from OWL ontology to infer new facts and enhance the discovered rules. Furthermore, the aforementioned techniques are not scalable.

3 Problem Statement and Contributions

The goal of rule mining is to identify new rules that entail a high ratio of positive edges from other positive edges, but a low ratio of negative edges from positive edges. This Ph.D. proposal addresses the problem of mining logical rules over KGs with semantics. Our main research objective is to mine rules over large KGs and also incorporate methods to further work on the federation of KGs. Concretely, we aim at encoding richer semantic knowledge from the KGs to mine more meaningful rules. As a result, our goal is to design a scalable approach to mine logical rules which also demonstrates cause-effect relationships.

3.1 Preliminaries

Knowledge Graphs is defined as G = *(V,E,L)* is a directed edge-labeled graph as defined in [11]. **Horn Rules** A mined rule is a Horn clause of the form: $Body \implies Head$, where *Body* is a conjunction of predicate facts; $Head$ is a predicate fact. All the variables in $Head$ are terms of at least one predicate fact in the *Body*, and every two predicate facts in *Body* share at least one variable.
Partial Completeness Assumption (PCA) is defined as G = *(V,E,L)* is a directed edge-labeled graph, the set of *heuristic-based negative edges hE-* in G is to consider as a negative edge, every edge *(s,p,o')* not in E, but that *(s,p,o)* belongs to E. That is *hE-* = $(s, p, o')|(s, p, o') \notin E$ and *(s,p,o)* $\in E$. PCA assumes that heuristic-based negative edges are possible incomplete edges.
Our approach performs task-agnostic mining, which means that it has been generalized to be interoperable across multiple systems, and the goal is to make *true predictions* that can potentially complete the missing relationships in large *incomplete KGs* under *Partial Completeness Assumption (PCA)*. Later these

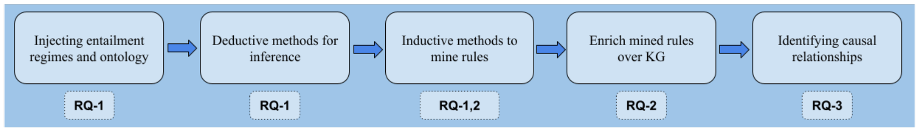

Fig. 2: **Research Pipeline:** Figure demonstrates the pipeline of steps followed to build a rule mining system. Related research questions are also added.

mined rules can be used to identify the causal relationship being observed between the nodes or by the rules mined. These would be termed 'Causal Rules' with added semantics to the Horn rules. Similar to the logical rules mined by AMIE and generating negative edges, we would also like to use the *PCA Confidence* measure to identify the potential incompleteness and predict true positives to complete the KG. Also, *Causal Ratio* metric will be used in order to discover potential causal rules mined over the KGs. To begin with the entailment regimes, we would be taking into consideration `rdfs:subClassOf, rdfs:subPropertyOf, rdfs:property, typing and owl:sameAs` as mentioned in the motivating example 1. Later this can be extended with all the entailment available to enhance the power of KGs. Our approach will also follow the incremental approach to generate rules incrementally, which makes it more efficient. This process will interleave the generation of mined rules with the retrieval of data from the KGs via queries executed over SPARQL endpoints.
Research Questions: Our approach aims to improve the rule mining process to answer the following research questions by ensuring KG completion and mining rules efficiently: **RQ1)** What is the impact of injecting entailment regimes to enrich the mined rules? **RQ2)** How can scalability be achieved in large KGs? **RQ3)** Can knowledge extracted from KGs help in identifying causality relationships to enhance explainability?

4 Research Methodology and Approach

Let us briefly discuss the methodology followed in our approach, as illustrated in Figure 2. In order to further enhance the traditional techniques and optimally mine causal rules, we will fully exploit the benefits of KGs and semantics. To begin, injecting entailment regimes will aid deductive methods in making inferences over KGs. This will answer the **RQ1)**. Inductively mining rules incrementally by avoiding the blocking process will suffice the **RQ1** and **RQ2)**. Traditional approaches executed over a KG upload the RDF triple files in main memory, impacting scalability to large KGs. Therefore, our approach will access SPARQL endpoints and execute queries to traverse the KGs which in turn will help to implement a more scalable algorithm to efficiently mine causal rules. To reduce the search space, our approach would consider the subset of the KG as per the user's interest. Furthermore, more enriched rules will be extracted over the subset of KG by later discovering causal rules that will answer **RQ3)**.

Lung Cancer KG	#T	#E	#P
Without Inference	23,773	1522	20
With Inference	31,932	1668	20

(a) The Lung Cancer (LC) KG Statistics. **#T** Number of Triples, **#E** Number of Entities, and **#P** Number of Predicates.

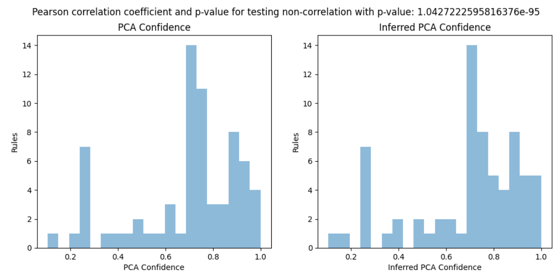

(b) **PCA Confidence:** Understanding the distribution of the PCA Confidence metric over KGs on how frequently each value occurs.

Fig. 3: **Initial Results**. (a) depicts statistics of the benchmark used to perform the experiments (b) shows the experimental results of the probability of the correlation between *PCA Confidence* and *Inferred PCA Confidence*. It is represented by the p-value which states that the metrics are statistically significant.

5 Evaluation/Evaluation Plan

The evaluation is performed on the real-world KGs in the biomedical area. We have started our evaluation of the lung cancer KG from P4-LUCAT [3]. The evaluation strategy for our approach in all the steps mentioned in Figure 2 is to measure efficiency in terms of the time it takes our rule mining approach to mine rules efficiently. In the later phase, our approach will incrementally mine rules and produce output to improve scalability. Rules will be generated based on the user's request, such as if the user requests rules with a higher *PCA Confidence* or any other conditions. These conditions will be taken into account when producing rules in order to save time, making our rule mining approach more scalable and accurate. Furthermore, by keeping humans in the loop, we will evaluate the causal rules generated by our approach. Domain experts in the biomedical domain will be validating the rules produced by our approach to check if they comply with the clinical guidelines and help in identifying causality.

6 Results So Far

We describe the outcomes of evaluating these methods on a KG comprising synthetic lung cancer patients generated from the biomedical KG discussed in the previous section. This initial study aims at reporting the impact of injecting entailment regimes on the KGs to answer our **RQ1)**. The preliminary results are being evaluated on two lung cancer KGs in order to compare them to state-of-the-art methods. The *Lung Cancer (LC) KG* comprises the characteristics of synthetic lung cancer patients (i.e., `identifier, gender, age, biomarker`).

[3] EraMed project https://p4-lucat.eu/

Entailment regimes are considered as described in Figure 1 on the LC KG.
As the motivating example 1 states, enriching the mined rule aids in the extraction of new insights from the KGs. In contrast to naive approaches, our approach takes `rdfs:subPropertyOf` into account for the experiments in the current example. This yields higher metrics values and demonstrates potential true predictions. For example, higher *Inferred PCA Confidence* of a rule quantifies the KG's partial completion by identifying more productive rules.
Observed Results and Discussion Table 3a describes the LC KG without inference and after the materialization of the deductive closure of the entailment regimes. The same number of rules (10,766 rules) were mined from versions of the LC KG, however, the scores of the mined rules have changed. Figure 3b exhibits the null hypothesis test performed to guarantee statistical independence between *PCA Confidence* and *Inferred PCA Confidence* metrics. When semantics are incorporated, the observed difference in the frequency distribution of metrics reveals the potential completion of KGs with true predictions. The results of our approach are publicly available on GitHub[4]

7 Conclusions/Lessons Learned

We propose a rule mining algorithm for mining causal rules from KGs. Our approach mine rules over both *A-Box* and *T-Box* of the KGs and promises scalability by implementing operators that enable the continuous generation of mined rules. Our initial framework exploits the semantics of the KGs and puts into perspective their relevance during a mining process. Our initial results indicate that these semantic-based mined rules are informative in domains such as healthcare, e.g., to understand how treatments have been prescribed and their relationships with existing medical guidelines. The lessons learned by reviewing the literature on semantic symbolic learning over KGs assisted us in recognizing the benefits and drawbacks of current approaches. This enabled us to improve rule mining systems by effectively utilizing KG semantics for better refinement and completion of KGs. The next steps in this Ph.D. work will be to refine the initial prototype in order to make the rule mining process more scalable and to polish it in order to make better predictions with higher scalability and accuracy.

Acknowledgements I would like to express my special thanks to my supervisor Prof. Dr. Maria-Esther Vidal for all her guidance and feedback. This work has been supported by the project TrustKG-Transforming Data in Trustable Insights with grant P99/2020 and the EraMed project P4-LUCAT (GA No. 53000015).

References

1. Abedjan, Z., Lorey, J., Naumann, F.: Reconciling ontologies and the web of data. In: ACM CIKM '12 (2012). https://doi.org/10.1145/2396761.2398467

[4] https://github.com/SDM-TIB/Symbolic_Learning_over_KGs

2. Agrawal, R., Imieliński, T., Swami, A.: Mining association rules between sets of items in large databases. In: ACM SIGMOD (1993). https://doi.org/10.1145/170035.170072
3. Agrawal, R., Imieliński, T., Swami, A.: Mining association rules between sets of items in large databases. SIGMOD Rec. (1993). https://doi.org/10.1145/170036.170072
4. Beall, J., Restall, G., Sagi, G.: Logical Consequence. In: The Stanford Encyclopedia of Philosophy (2019)
5. Brickley, D., Guha, R.V.: RDF Schema 1.1. W3c recommendation, World Wide Web Consortium (2014)
6. Galárraga, L., Teflioudi, C., Hose, K., Suchanek, F.: Fast Rule Mining in Ontological Knowledge Bases with AMIE+. The VLDB Journal (2015), https://hal-imt.archives-ouvertes.fr/hal-01699866
7. Galárraga, L.A., Preda, N., Suchanek, F.M.: Mining rules to align knowledge bases. In: Automated Knowledge Base Construction (2013). https://doi.org/10.1145/2509558.2509566
8. Galárraga, L., Teflioudi, C., Hose, K., Suchanek, F.: Amie: Association rule mining under incomplete evidence in ontological knowledge bases. In: WWW 2013 (2013). https://doi.org/10.1145/2488388.2488425
9. Han, J., Chiang, J.Y., Chee, S., Chen, J., Chen, Q., Cheng, S., Gong, W., Kamber, M., Koperski, K., Liu, G., Lu, Y., Stefanovic, N., Winstone, L., Xia, B.B., Zaiane, O.R., Zhang, S., Zhu, H.: Dbminer: A system for data mining in relational databases and data warehouses. In: CASCON Conference (1997)
10. Hitzler, P., Krötzsch, M., Parsia, B., Patel-Schneider, P.F., Rudolph, S.: OWL 2 Web Ontology Language Primer (Second Edition). W3C recommendation, World Wide Web Consortium (2012)
11. Hogan, A., Blomqvist, E., Cochez, M., d'Amato, C., de Melo, G., Gutiérrez, C., Kirrane, S., Labra Gayo, J.E., Navigli, R., Neumaier, S., Ngonga Ngomo, A.C., Polleres, A., Rashid, S.M., Rula, A., Schmelzeisen, L., Sequeda, J.F., Staab, S., Zimmermann, A.: Knowledge Graphs (2021). https://doi.org/10.2200/S01125ED1V01Y202109DSK022
12. Kaur, M., Kang, S.: Market basket analysis: Identify the changing trends of market data using association rule mining. Procedia Computer Science (2016). https://doi.org/https://doi.org/10.1016/j.procs.2016.05.180
13. Lajus, J., Galárraga, L., Suchanek, F.: Fast and exact rule mining with amie 3. In: Harth, A., Kirrane, S., Ngonga Ngomo, A.C., Paulheim, H., Rula, A., Gentile, A.L., Haase, P., Cochez, M. (eds.) The Semantic Web (2020)
14. Meilicke, C., Chekol, M.W., Ruffinelli, D., Stuckenschmidt, H.: Anytime bottom-up rule learning for knowledge graph completion. In: IJCAI-19 (2019). https://doi.org/10.24963/ijcai.2019/435
15. Munch, M., Dibie-Barthelemy, J., Manfredotti, C., Wuillemin, P.H.: Towards interactive causal relation discovery driven by an ontology (Jun 2018)
16. Sarawagi, S., Thomas, S., Agrawal, R.: Integrating association rule mining with relational database systems: Alternatives and implications. SIGMOD Rec. (1998). https://doi.org/10.1145/276305.276335
17. Simonne, L., Pernelle, N., Saïs, F., Thomopoulos, R.: Differential causal rules mining in knowledge graphs. In: ACM KCAP (2021). https://doi.org/10.1145/3460210.3493584
18. Toivonen, H.: Apriori algorithm. In: Sammut, C., Webb, G.I. (eds.) Encyclopedia of Machine Learning and Data Mining (2017). https://doi.org/10.1007/978-1-4899-7687-1_27