24th European Conference on Artificial Intelligence - ECAI 2020
Santiago de Compostela, Spain

Towards Meta-reasoning for Ontologies: A Roadmap

Yuan-Fang Li[1] and **Yong-Bin Kang**[2]

Abstract. Ontologies are widely used to formally represent abstract domain knowledge. Logic reasoning ensures the logical consistency of ontologies, and infers knowledge implicitly encoded in ontologies. It has been shown both theoretically and empirically that for large and complex ontologies, reasoning is still time-consuming and resource-intensive. Meta-reasoning exploits machine learning techniques to tackle the important problems of understanding the source of reasoning hardness and to predict reasoning efficiency, with the overall goal of improving reasoning efficiency. In this paper, we highlight recent advances in meta-reasoning for Semantic Web ontologies, briefly present technical innovations and results, and discuss important problems for future research.

1 Introduction and Motivation

Ontologies describe abstract concepts and the complex relationships among them, and are widely used to represent complex knowledge in many application domains [1, 17, 14]. Expressive ontology languages OWL 1 DL [8] and OWL 2 DL [4] are the lingua franca for these communities. Their precise semantics enables logical reasoning, which includes the maintenance of the logical soundness of ontologies (i.e. consistency checking) and the inference of implicit knowledge from ontologies (i.e. classification).

These ontology languages are highly expressive, thus they also possess high worst-case computational complexity for the above core reasoning tasks. For example, consistency checking of an ontology in $\mathcal{SHOIN}(\mathbf{D})$, the description logic (DL) underlying OWL 1 DL, has NEXPTIME-complete worst-case complexity [8]. The complexity of the same problem for $\mathcal{SROIQ}(\mathbf{D})$, the DL underlying OWL 2 DL, is even higher (2NEXPTIME-complete) [4].

Significant algorithmic and implementation advances have been made over the past decades. However, it has been shown empirically that reasoning on large and complex ontologies can still be time-consuming and resource-intensive for state-of-the-art reasoners [5, 12]. Moreover, there are large variations in hardness of different ontologies [7]. Therefore, it is non-trivial to identify the most efficient reasoner given an ontology. On the other hand, it has also been observed that as a whole, the state-of-the-art ontology reasoners are *robust* [6], that it is highly likely that one of the reasoners performs sufficiently well on a given ontology. Therefore, it is highly desirable to be able to predict and select automatically the best reasoner on-the-fly.

The aim of *meta-reasoning* is to identify the most appropriate reasoner, from a pool of component reasoners, for a given ontology, based on some criteria, such as the robustness and efficiency of performing ontology reasoning. A *meta-reasoner* is able to exploit the differences in strength of different reasoners, thus potentially achieve near-optimal performance, by choosing the most appropriate component reasoner for a given ontology.

In this paper, we highlight our *R_2O_2** meta-reasoning framework (simply R_2O_2*) [10] for highly expressive ontology languages and future work directions for meta-reasoning.

2 R_2O_2*: A Meta-reasoning Framework

In a series of works [19, 13, 9, 10], we have investigated the problem of meta-reasoning. The aim of R_2O_2* is to construct a meta-reasoner that predicts the most efficient component reasoner from a collection of component reasoners and then select it to carry out reasoning on that ontology efficiently. The R_2O_2* encompasses a number of main components [10]:

Ontology hardness description It is essential to be able to precisely describe the *hardness* of an ontology without actually performing reasoning. For this purpose we have designed a suite of 91 metrics to measure the *design complexity* of OWL ontologies [11, 19]. These metrics measure, syntactically and structurally, the hardness of an ontology either as a whole, or for each class, property or individual. For instance, some metrics measure the numbers and percentages of certain language constructs (e.g. general concept inclusion, or GCI) that are especially hard for a reasoner to deal with. Some metric measures the degree of deviation of the inheritance hierarchy from a pure tree structure.

Reasoning efficiency prediction By collecting reasoning time data of a corpus of diverse ontologies, we can train prediction models that, given an ontology, accurately predict the actual reasoning time of a given reasoner [13]. Such a prediction model lays the foundation of meta-reasoning, as now we have a way of estimating the reasoning time of a same ontology for a number of reasoners.

Learning-based reasoner selection Building on the above two components, R_2O_2* ranks a collection of reasoners based on estimated reasoning efficiency. R_2O_2* trains multiple models, including learning to rank and multi-class classification, for reasoner selection, and employs the *stacking* ensemble method to combine these models to obtain more accurate predictive power.

We conducted comprehensive evaluation on a corpus of 1,760 ontologies and six state-of-the-art reasoners capable of supporting OWL 1 DL ontologies. The main result can be seen in a violin plot in Figure 1 that shows a combination of a boxplot and a mirrored kernel density plot. Each shape contains the following components: (1) The violin itself shows the distribution of reasoning time; (2) The cross ($\times$) shows the mean reasoning time of the reasoner; (3) The plus ($+$) shows the median reasoning time of the reasoner; (4) The 3 horizontal lines within each shape shows the 25%, 50%, and 75% of data, respectively; (5) The grey dots represent the actual reasoning time of

[1] Monash University, email: yuanfang.li@monash.edu
[2] Swinburne University of Technology, email: ykang@swin.edu.au

all ontologies. From left to right, the first six columns depict reasoning time of the six component reasoners. AutoFolio [15] is a general-purpose state-of-the-art algorithm selection model. $R_2O_2*_{(all)}$ is our meta-reasoner built with the ensemble of different reasoner selection models. Finally, VBR represents the ideal, *virtual* best reasoner that can only be determined post hoc. As can be seen from the figure, $R_2O_2*_{(all)}$ achieves the best reasoning performance except VBR. $R_2O_2*_{(all)}$ outperforms Konclude by around 10% and JFact by around 4000%; also, $R_2O_2*_{(all)}$ outperforms AutoFolio by more than 10% in reasoning efficiency.

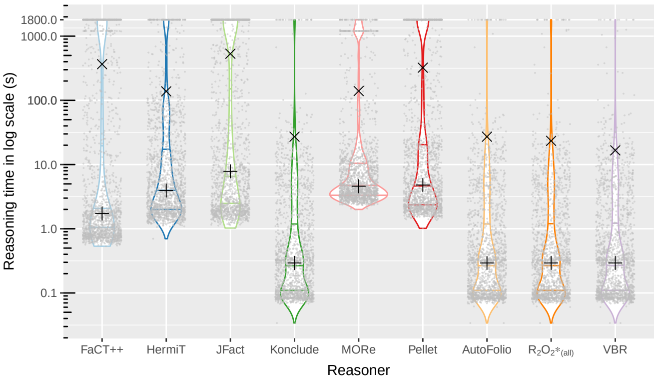

Figure 1. Reasoning time comparison between component reasoners and meta-reasoners from R_2O_2* [10].

Interestingly, we also observed that even though one of the component reasoners, Konclude [18], dominates the others in efficiency, almost each component reasoner is the most efficient one for *some* ontologies. This observation further validates the robustness of the ontology reasoner collective [6], and provides additional empirical evidence of the effectiveness of the meta-reasoning approach.

3 Future Directions for Meta-reasoning

Reasoning is a core task for ontologies expressed in the OWL family of languages. For large and complex ontologies, efficient reasoning remains a challenging problem, even for state-of-the-art reasoners. In this paper, we demonstrate, through our recent meta-reasoner R_2O_2*, the fertile ground for collaboration between knowledge representation and reasoning (KRR) and machine learning (ML). R_2O_2* automatically selects the most efficient reasoner from a collection of reasoners for a given new ontology. In this way, meta-reasoning can take advantage of the various reasoning algorithms and optimisation techniques implemented by different reasoners, thus being able to achieve good and robust reasoning efficiency. We envisage a number of important future research problems for meta-reasoning, and more broadly for KRR+ML.

Representation learning All of existing prediction models for ontology reasoning efficiency make use of hand-crafted syntactic and structural features. It is interesting to learn representations [2] of important constructs, including classes, properties, and individuals. The learned representations will have a wide range applications beyond reasoning efficiency prediction and meta-reasoning.

Reasoner characteristics Our R_2O_2* only considers the hardness of ontologies. It is worth investigating how the characteristics of reasoners can be taken into account.

Inefficiency repair If an ontology is predicted to be hard for reasoning, can it be *repaired* so that reasoning becomes efficient while minimising impact on soundness/completeness?

ABox reasoning support Current meta-reasoners all focus on TBox (terminology) reasoning. An ABox, which contains instance data, could be orders of magnitude larger than a TBox. Thus, how to characterise ABox hardness and support meta-reasoning for ABox reasoning tasks, such as query answering [3] and materialisation, is an important research problem [16].

Benchmark generation Meta-reasoning relies heavily on accurate learning models, for both reasoning performance prediction and reasoning selection. It is thus useful to be able to generate synthetic yet realistic ontologies for continued model training.

References

[1] M. Ashburner, C. A. Ball, J. A. Blake, et al., 'Gene Ontology: Tool for the Unification of Biology', *Nat Genet*, **25**(1), 25–29, (May 2000).
[2] Yoshua Bengio, Aaron C. Courville, and Pascal Vincent, 'Representation learning: A review and new perspectives', *IEEE Trans. Pattern Anal. Mach. Intell.*, **35**(8), 1798–1828, (2013).
[3] Diego Calvanese, Giuseppe Giacomo, Domenico Lembo, Maurizio Lenzerini, and Riccardo Rosati, 'Tractable reasoning and efficient query answering in description logics: The DL-Lite family', *J. Autom. Reason.*, **39**, 385–429, (October 2007).
[4] Bernardo Cuenca-Grau, Ian Horrocks, Boris Motik, Bijan Parsia, Peter Patel-Schneider, and Ulrike Sattler, 'OWL 2: The next step for OWL', *Journal of Web Semantics: Science, Services and Agents on the World Wide Web*, **6**, 309–322, (November 2008).
[5] Kathrin Dentler, Ronald Cornet, Annette ten Teije, and Nicolette de Keizer, 'Comparison of reasoners for large ontologies in the OWL 2 EL profile', *Semantic Web Journal*, **2**(2), 71–87, (2011).
[6] Rafael S. Gonçalves, Nicolas Matentzoglu, Bijan Parsia, and Uli Sattler, 'The empirical robustness of description logic classification', in *ISWC 2013 Posters & Demos Track*, pp. 277–280. CEUR-WS.org, (2013).
[7] Rafael S. Gonçalves, Bijan Parsia, and Ulrike Sattler, 'Performance heterogeneity and approximate reasoning in description logic ontologies', In *International Semantic Web Conference (ISWC 2012)*, pp. 82–98.
[8] Ian Horrocks, Peter F. Patel-Schneider, and Frank van Harmelen, 'From $\mathcal{SHIQ}$ and RDF to OWL: The Making of a Web Ontology Language', *Journal of Web Semantics*, **1**(1), 7–26, (2003).
[9] Yong-Bin Kang, Shonali Krishnaswamy, and Yuan-Fang Li, 'R2O2: an efficient ranking-based reasoner for OWL ontologies', in *International Semantic Web Conference (ISWC 2015)*, pp. 322–338. (2015).
[10] Yong-Bin Kang, Shonali Krishnaswamy, Wudhichart Sawangphol, Lianli Gao, and Yuan-Fang Li, 'Understanding and improving ontology reasoning efficiency through learning and ranking', *Information Systems*, **87**, 101412, (2020).
[11] Yong-Bin Kang, Yuan-Fang Li, and Shonali Krishnaswamy, 'Predicting reasoning performance using ontology metrics', In *International Semantic Web Conference (ISWC 2012)*, pp. 198–214.
[12] Yong-Bin Kang, Yuan-Fang Li, and Shonali Krishnaswamy, 'A rigorous characterization of reasoning performance – a tale of four reasoners', in *1st Workshop on OWL Reasoner Evaluation (ORE-2012)*, (2012).
[13] Yong-Bin Kang, Jeff Z. Pan, Shonali Krishnaswamy, Wudhichart Sawangphol, and Yuan-Fang Li, 'How long will it take? Accurate prediction of ontology reasoning performance', in *AAAI*, (2014).
[14] Maurizio Lenzerini, 'Data integration: a theoretical perspective', in *Symposium on Principles of database systems (PODS 2002)*, PODS '02, pp. 233–246, New York, NY, USA, (2002). ACM.
[15] Marius Lindauer, Holger H Hoos, Frank Hutter, and Torsten Schaub, 'AutoFolio: An automatically configured algorithm selector', *Journal of Artificial Intelligence Research*, **53**, 745–778, (2015).
[16] Jeff Z. Pan, Carlos Bobed, Isa Guclu, Fernando Bobillo, Martin J. Kollingbaum, Eduardo Mena, and Yuan-Fang Li, 'Predicting reasoner performance on ABox intensive OWL 2 EL ontologies', *International Journal on Semantic Web and Information Systems (IJSWIS)*, **14**(1), (2018).
[17] *Ontology-Driven Software Development*, eds., Jeff Z. Pan, Steffen Staab, Uwe Aßmann, Jÿrgen Ebert, and Yuting Zhao, Springer, 2013.
[18] Andreas Steigmiller, Thorsten Liebig, and Birte Glimm, 'Konclude: System description', *J. Web Sem.*, **27**, 78–85, (2014).
[19] Hongyu Zhang, Yuan-Fang Li, and Hee Beng Kuan Tan, 'Measuring design complexity of Semantic Web ontologies', *Journal of Systems and Software*, **83**(5), 803–814, (2010).