

Practical Uses of Existential Rules in Knowledge Representation

David Carral, Markus Krötzsch, and Jacopo Urbani

¹ TU Dresden, Dresden, Germany, `firstname.lastname@tu-dresden.de`

² Vrije Universiteit Amsterdam, The Netherlands, `jacopo@cs.vu.nl`

1 Description

Reasoning with existential rules is a well-established research area in KR that has increasingly become more relevant in recent years [1–4, 9, 16, 21]. In this research area, reasoning operates with two different types of formulas: *facts* and (*existential*) *rules*. The latter are formulas of the following form.

$$\forall \vec{x}, \vec{y}. (\beta[\vec{x}, \vec{y}] \rightarrow \exists \vec{w}. \eta[\vec{x}, \vec{w}])$$

In the above, \vec{x} , \vec{y} , and \vec{w} are pairwise disjoint lists of variables; β and η are non-empty conjunctions of atoms; and $x, y \in \vec{x}$.

In this tutorial, we discuss the *terminating fragment* of existential rules. That is, the fragment of this language for which the *chase*—a sound and complete, bottom-up reasoning procedure—is guaranteed to terminate. We propose the use of terminating rule languages to encode complex reasoning algorithms in a declarative way [8]. This approach—which follows the classical slogan “Algorithm = Logic + Control”—promises to turn high-level specifications of logical calculi as systems of inference rules into declarative rule-based models that can be executed on state-of-the-art rule engines. Finally, we present VLog [5, 22], an efficient implementation of the chase and show, in a hands-on session, how to use this tool to implement simple reasoning calculi.

2 Proposed Length of the Tutorial

Half a day.

3 Outline

The tutorial consists of two different parts where we (A) formally introduce the terminating fragment of existential rules and (B) show how to use VLog to implement reasoning algorithms in a hands-on session.

A. Introduction to Existential Rules. In this part of the tutorial, which will take around two hours, we introduce existential rules and some prominent decidable fragments of this language. Namely, we present the following topics.

1. We start with a general introduction of the syntax and semantics of the existential rule fragment.
2. We show that reasoning with existential rules is undecidable.
3. We present different variants of the chase algorithm—a sound and complete, albeit non-terminating, reasoning procedure for existential rules.
4. We show that we cannot decide chase termination.
5. We present some acyclicity notions; that is, some sufficient conditions that guarantee termination of the chase. For instance, we present model-faithful acyclicity (MFA) [9] and restricted model-faithful acyclicity (RMFA) [6].
6. We discuss the complexity of reasoning with acyclicity notions. That is, we discuss the complexity of checking whether a rule set is (R)MFA and the complexity of reasoning with (R)MFA rule sets.
7. We discuss the complexity of reasoning with Datalog rule sets; that is, existential rule sets that do not contain any existential quantifiers.
8. We present and study the complexity of Datalog(S)—an extension of Datalog that allows for the use of set constructors.
9. We discuss some results that show how Datalog(S) rule sets can be translated into existential rule sets with a terminating chase that preserve fact-entailment.

B. The Rule-Engine VLog. In this part of the tutorial, which will take around one hour, we show how to encoding logical calculi using rule based languages and how to implement these using VLog. Namely, we discuss the following topics.

1. We briefly present the DL languages \mathcal{EL} and \mathcal{ALC} .
2. We discuss how to encode an existing consequence-based procedure for \mathcal{EL} [11] using a fixed Datalog rule set.
3. We discuss why we cannot encode such a procedure for \mathcal{ALC} using a fixed Datalog rule set. Indeed, this is the cases because of the complexity properties of these two languages.
4. We discuss how to encode an existing consequence-based procedure for \mathcal{ALC} [10] using a fixed Datalog(S) rule set.
5. In a hands-on session, we discuss how to implement the encodings presented in (2) and (4) using the rule engine VLog.
6. We informally discuss how we can implement other existing reasoning algorithms using our approach.

4 Potential Target Audience and Prerequisite Knowledge

We believe that this tutorial can be potentially interesting to a large audience within the KR community, since we describe an elegant way to implement reasoning algorithms using rule-based languages.

In order to make the most of the tutorial, it is convenient to be familiar with first-order logic. Nevertheless, this is not absolutely unnecessary as we will present the syntax and semantics of existential rules in A.1.

5 Brief Resume of the Authors

Note that the presenters have expertise in the theoretical study of existential rules [6, 7, 9, 12–20] as well as in the implementation of rule based-reasoners [5, 22, 23]. Also, at least two of the authors will present the tutorial at the conference if it gets accepted.

David Carral is a postdoctoral scholar at the chair for Knowledge-Based Systems at the Faculty of Computer Science of TU Dresden. He completed his master's and doctor's degrees at Wright State University, both under the supervision of Prof. Dr. Pascal Hitzler, in 2012 and 2016, respectively. For several months during his Ph.D., he was an exchange student at the University of Oxford, working under the supervision of Prof. Dr. Bernardo Cuenca Grau. His research interests include symbolic artificial intelligence and knowledge representation related topics. More precisely, David studies the theoretical properties of logical languages such as Description Logics and existential rules, and the implementation of reasoning algorithms.

Markus Krötzsch is a full professor at the Faculty of Computer Science of TU Dresden, where he is holding the chair for Knowledge-Based Systems. He obtained his Ph.D. from the Institute of Applied Informatics and Formal Description Methods (AIFB) of Karlsruhe Institute of Technology (KIT) in 2010, and thereafter worked as a researcher and departmental lecturer at the Department of Computer Science of the University of Oxford until October 2013. Krötzsch's extensive research activities in the area of rule-based knowledge representation and reasoning have contributed to the development and analysis of rule languages [16, 9, 6, 20, 15], rule-based inference methods [13, 8], and rule reasoners [22, 23, 5]. His wider research interests also include ontology languages, query answering, reasoning complexity, knowledge graphs, and content management and integration platforms for the Web of Data. He has published many works in leading journals and conferences, and two textbooks on semantic technologies. He has given tutorials and invited lectures at various events, including ESSLLI, IJCAI, and ICDT, and he has co-organised numerous scientific and educational events, most recently the Reasoning Web Summer School 2019.

Jacopo Urbani is an assistant professor at the Department of Computer Science of the Vrije Universiteit Amsterdam (VUA). He is also a guest researcher at the Centrum Wiskunde & Informatica and at the Max Planck Institute in Informatics. He wrote a PhD thesis on distributed reasoning algorithms for very large Knowledge Graphs. The thesis was nominated by the Royal Netherlands Academy of Arts and Sciences as one of the best PhD theses in Computer Science in the country. After spending part of his postdoc in USA and Germany, he joined the faculty of the VUA and has been tenured in 2018. He is one of the main developers of VLog and has published extensively on rule-based reasoning on Knowledge Graphs (KGs) at venues like AAI, WWW, ISWC, etc. His recent interests include the application of rule-based reasoning for fact-checking and to perform KG completion using information extracted from unstructured sources.

References

1. Baget, J., Leclère, M., Mugnier, M., Salvat, E.: On rules with existential variables: Walking the decidability line. *Artif. Intell.* 175(9-10), 1620–1654 (2011)
2. Bellomarini, L., Sallinger, E., Gottlob, G.: The vadalog system: Datalog-based reasoning for knowledge graphs. *PVLDB* 11(9), 975–987 (2018)
3. Benedikt, M., Konstantinidis, G., Mecca, G., Motik, B., Papotti, P., Santoro, D., Tsamoura, E.: Benchmarking the chase. In: Sallinger, E., den Bussche, J.V., Geerts, F. (eds.) *Proceedings of the 36th ACM SIGMOD-SIGACT-SIGAI Symposium on Principles of Database Systems, PODS 2017*. pp. 37–52. ACM (2017)
4. Calautti, M., Gottlob, G., Pieris, A.: Chase termination for guarded existential rules. In: Milo, T., Calvanese, D. (eds.) *Proceedings of the 34th ACM Symposium on Principles of Database Systems, PODS 2015*. pp. 91–103. ACM (2015)
5. Carral, D., Dragoste, I., González, L., Jacobs, C.J.H., Krötzsch, M., Urbani, J.: VLog: A rule engine for knowledge graphs. In: Ghidini, C., Hartig, O., Maleshkova, M., Svátek, V., Cruz, I.F., Hogan, A., Song, J., Lefrançois, M., Gandon, F. (eds.) *The Semantic Web - ISWC 2019 - 18th International Semantic Web Conference, Proceedings, Part II. Lecture Notes in Computer Science*, vol. 11779, pp. 19–35. Springer (2019)
6. Carral, D., Dragoste, I., Krötzsch, M.: Restricted chase (non)termination for existential rules with disjunctions. In: Sierra, C. (ed.) *Proceedings of the 26th International Joint Conference on Artificial Intelligence, IJCAI 2017*. pp. 922–928. ijcai.org (2017)
7. Carral, D., Dragoste, I., Krötzsch, M.: Tractable query answering for expressive ontologies and existential rules. In: d’Amato, C., Fernández, M., Tamma, V.A.M., Lécué, F., Cudré-Mauroux, P., Sequeda, J.F., Lange, C., Heflin, J. (eds.) *The Semantic Web - ISWC 2017 - 16th International Semantic Web Conference, Proceedings, Part I. Lecture Notes in Computer Science*, vol. 10587, pp. 156–172. Springer (2017)
8. Carral, D., Dragoste, I., Krötzsch, M., Lewe, C.: Chasing sets: How to use existential rules for expressive reasoning. In: Kraus, S. (ed.) *Proceedings of the Twenty-Eighth International Joint Conference on Artificial Intelligence, IJCAI 2019*. pp. 1624–1631. ijcai.org (2019)
9. Cuenca Grau, B., Horrocks, I., Krötzsch, M., Kupke, C., Magka, D., Motik, B., Wang, Z.: Acyclicity notions for existential rules and their application to query answering in ontologies. *J. Artif. Intell. Res.* 47, 741–808 (2013)
10. Kazakov, Y.: Consequence-driven reasoning for Horn SHIQ ontologies. In: Boutilier, C. (ed.) *Proceedings of the 21st International Joint Conference on Artificial Intelligence, IJCAI 2009*. pp. 2040–2045 (2009)
11. Kazakov, Y., Krötzsch, M., Simančík, F.: The incredible ELK – From polynomial procedures to efficient reasoning with el ontologies. *J. Autom. Reasoning* 53(1), 1–61 (2014)
12. Krötzsch, M.: *Description Logic Rules, Studies on the Semantic Web*, vol. 8. IOS Press (2010)
13. Krötzsch, M.: Efficient rule-based inferencing for OWL EL. In: Walsh, T. (ed.) *IJCAI 2011, Proceedings of the 22nd International Joint Conference on Artificial Intelligence*. pp. 2668–2673. *IJCAI/AAAI* (2011)
14. Krötzsch, M., Maier, F., Krisnadhi, A., Hitzler, P.: A better uncle for OWL: nominal schemas for integrating rules and ontologies. In: Srinivasan, S., Ramamritham, K., Kumar, A., Ravindra, M.P., Bertino, E., Kumar, R. (eds.) *Proceedings of the 20th International Conference on World Wide Web, WWW 2011*. pp. 645–654. ACM (2011)
15. Krötzsch, M., Marx, M., Rudolph, S.: The power of the terminating chase (invited talk). In: Barceló, P., Calautti, M. (eds.) *22nd International Conference on Database Theory, ICDT 2019, March 26–28, 2019, Lisbon, Portugal. LIPIcs*, vol. 127, pp. 3:1–3:17. Schloss Dagstuhl - Leibniz-Zentrum fuer Informatik (2019)

16. Krötzsch, M., Rudolph, S.: Extending decidable existential rules by joining acyclicity and guardedness. In: Walsh, T. (ed.) IJCAI 2011, Proceedings of the 22nd International Joint Conference on Artificial Intelligence. pp. 963–968. IJCAI/AAAI (2011)
17. Krötzsch, M., Thost, V.: Ontologies for knowledge graphs: Breaking the rules. In: Groth, P.T., Simperl, E., Gray, A.J.G., Sabou, M., Krötzsch, M., Lécué, F., Flöck, F., Gil, Y. (eds.) The Semantic Web - ISWC 2016 - 15th International Semantic Web Conference, Proceedings, Part I. Lecture Notes in Computer Science, vol. 9981, pp. 376–392 (2016)
18. Magka, D., Krötzsch, M., Horrocks, I.: Computing stable models for nonmonotonic existential rules. In: Rossi, F. (ed.) IJCAI 2013, Proceedings of the 23rd International Joint Conference on Artificial Intelligence, Beijing, China, August 3-9, 2013. pp. 1031–1038. IJCAI/AAAI (2013)
19. Magka, D., Krötzsch, M., Horrocks, I.: A rule-based ontological framework for the classification of molecules. *J. Biomedical Semantics* 5, 17 (2014)
20. Marx, M., Krötzsch, M., Thost, V.: Logic on MARS: ontologies for generalised property graphs. In: Sierra, C. (ed.) Proceedings of the Twenty-Sixth International Joint Conference on Artificial Intelligence, IJCAI 2017. pp. 1188–1194. ijcai.org (2017)
21. Motik, B., Nenov, Y., Piro, R., Horrocks, I., Olteanu, D.: Parallel materialisation of Datalog programs in centralised, main-memory RDF systems. In: Brodley, C.E., Stone, P. (eds.) Proceedings of the 28th AAAI Conference on Artificial Intelligence. pp. 129–137. AAAI Press (2014)
22. Urbani, J., Jacobs, C.J.H., Krötzsch, M.: Column-oriented Datalog materialization for large knowledge graphs. In: Schuurmans, D., Wellman, M.P. (eds.) Proceedings of the 30th AAAI Conference on Artificial Intelligence. pp. 258–264. AAAI Press (2016)
23. Urbani, J., Krötzsch, M., Jacobs, C.J.H., Dragoste, I., Carral, D.: Efficient model construction for Horn logic with VLog - system description. In: Galmiche, D., Schulz, S., Sebastiani, R. (eds.) Automated Reasoning - 9th International Joint Conference, IJCAR 2018, Held as Part of the Federated Logic Conference, FloC 2018. Lecture Notes in Computer Science, vol. 10900, pp. 680–688. Springer (2018)