

The Power of Semantic Rules in Rulelog: Fundamentals and Recent Progress (extended abstract of tutorial presentation)

Benjamin N. Groszof, Michael Kifer, and Paul Fodor

Coherent Knowledge Systems, LLC.

1 Introduction

In this tutorial, we provide a comprehensive and up-to-date introduction to the fundamental concepts and recent progress in the area of Rulelog, a leading approach to semantic rules knowledge representation and reasoning. Rulelog is expressively powerful, computationally affordable, and has capable efficient implementations. A large subset of Rulelog is in draft as an industry standard¹ to be submitted to RuleML² and W3C³ as a dialect of Rule Interchange Format (RIF) [2, 3].

2 Rulelog Logical Language and Capabilities

Rulelog extends well-founded declarative logic programs (LP) with:

- Strong meta-reasoning, including higher-order syntax (Hilog) [4], reification [21], and rule ids (within the logical language).
- Explanations of inferences [1].
- efficient higher-order defaults, including “argumentation theories” [18, 19].
- flexible probabilistic reasoning, including distribution semantics [15], evidential probability [14], and tight integration with inductive machine learning are key areas of recent technology progress and ongoing R&D.
- bounded rationality, including restraint—a “control knob” to ensure that the computational complexity of inference is worst-case polynomial time [7, 1].
- “omni-directional” disjunction and existential quantifiers in the rule heads [9].
- object-orientation and frame syntax [13], which subsumes RDF triples.
- sound tight integration of first-order-logic ontologies including OWL and several other lesser features, including aggregation operators and integrity constraints.

We will cover many of these features in the tutorial.

¹ <http://ruleml.org/rif/rulelog/rif/RIF-Rulelog.html>

² <http://www.ruleml.org>

³ <http://www.w3.org>

3 Rulelog Implementation Techniques

Implementation techniques for Rulelog inferencing include transformational compilations and extensions of *tabling* algorithms from logic programming. “Tabling” here means smart caching of subgoals and conclusions together with incremental revision of the cached conclusions when facts or rules are dynamically added or deleted [16, 17]. “Tabling” is thus a mixture of backward-direction and forward-direction inferencing. There are both open-source and commercial tools for Rulelog that vary in their range of expressive completeness and of user convenience. They are interoperable with databases and spreadsheets, and complement inductive machine learning and natural language processing techniques. The most complete system today for Rulelog is Ergo⁴, a commercial platform suite from Coherent Knowledge Systems⁵. Flora-2⁶, an open source system, implements a significant subset of Rulelog reasoning [20, 22, 12].

4 Textual Rulelog

Time permitting, we will discuss “Textual” Rulelog, in which Rulelog is rendered in a natural language, such as English. ErgoText is a commercial realization of this approach. Using Rulelog to interpret and generate English is a key area of ongoing research and development [10].

5 Applications

Using Ergo, we will illustrate the various applications of the Rulelog technology in a wide range of tasks and domains in business, government, and science. We will tour areas of recent applications progress, which include: legal/policy compliance, e.g., in financial services; education/tutoring; and e-commerce marketing [6, 8, 5, 11].

6 Additional Tutorial Material

Previous, longer but now less up-to-date, tutorials on Rulelog were given at our earlier tutorial at AAAI-13 [7].

References

1. Andersen, C., Benyo, B., Calejo, M., Dean, M., Fodor, P., Grosf, B.N., Kifer, M., Liang, S., Swift, T.: Advanced knowledge base debugging for rulelog. In: Joint Proceedings of the 7th International Rule Challenge, the Special Track on Human Language Technology and the 3rd RuleML Doctoral Consortium, Seattle, USA, July 11 -13, 2013 (2013), <http://ceur-ws.org/Vol-1004/paper8.pdf>

⁴ <http://coherentknowledge.com/ergo-suite-platform-technology/>

⁵ <http://coherentknowledge.com>

⁶ <http://http://flora.sourceforge.net>

2. Boley, H., Kifer, M.: RIF Basic logic dialect (February 2013), <http://www.w3.org/TR/rif-bld/>, W3C Recommendation. <http://www.w3.org/TR/rif-bld/>
3. Boley, H., Kifer, M.: RIF Framework for logic dialects (February 2013), <http://www.w3.org/TR/rif-flld/>, W3C Recommendation. <http://www.w3.org/TR/rif-flld/>
4. Chen, W., Kifer, M., Warren, D.: HiLog: A foundation for higher-order logic programming. *Journal of Logic Programming* 15(3), 187–230 (February 1993)
5. Fibo technology summit at semtechbiz: Financial industry and sem tech leaders discuss ontology evaluation tools, flora-2's potential, and more (July 2013), <http://www.dataversity.net/fibo-technology-summit-at-semtechbiz-financial-industry-and-sem-tech-leaders-discuss-ontology-evaluation-tools-flora-2s-potential-and-more/>
6. Grosf, B., Dean, M., Kifer, M.: Semantic web rules: Fundamentals, applications, and standards (July 2013), tutorial presented at AAAI-2013; available at <http://silk.semwebcentral.org/>
7. Grosf, B., Swift, T.: Radial restraint: A semantically clean approach to bounded rationality for logic programs. In: AAAI Conference on Artificial Intelligence. AAAI (2013)
8. Grosf, B.N., Burstein, M.H., Dean, M., Andersen, C., Benyo, B., Ferguson, W., Incezan, D., Shapiro, R.: A SILK graphical UI for defeasible reasoning, with a biology causal process example. In: Palmirani, M., Shafiq, M.O., Francesconi, E., Vitali, F. (eds.) *Proceedings of the RuleML-2010 Challenge, at the 4th International Web Rule Symposium, Washington, DC, USA, October, 21-23, 2010*. CEUR Workshop Proceedings, vol. 649. CEUR-WS.org (2010), <http://ceur-ws.org/Vol-649/paper12.pdf>
9. Grosf, B.: Rapid text-based authoring of defeasible higher-order logic formulas, via textual logic and rulelog. In: Morgenstern, L., Stefaneas, P., Levy, F., Wyner, A., Paschke, A. (eds.) *Theory, Practice, and Applications of Rules on the Web, Lecture Notes in Computer Science*, vol. 8035, pp. 2–11. Springer Berlin Heidelberg (2013), http://dx.doi.org/10.1007/978-3-642-39617-5_2
10. Grosf, B.: Rapid text-based authoring of defeasible higher-order logic formulas, via textual logic and rulelog. In: Morgenstern, L., Stefaneas, P., Lvy, F., Wyner, A., Paschke, A. (eds.) *Theory, Practice, and Applications of Rules on the Web, Lecture Notes in Computer Science*, vol. 8035, pp. 2–11. Springer Berlin Heidelberg (2013), http://dx.doi.org/10.1007/978-3-642-39617-5_2
11. Grosf, B.: Making very expressive rules practical in logic and text (November 2013), invited talk, <http://coherentknowledge.com/wp-content/uploads/2013/05/DecisionCamp2013-talk-v7-1-BNG.pdf>
12. Kifer, M.: Flora-2: An object-oriented knowledge base language. The Flora-2 Web Site (2015), <http://flora.sourceforge.net>
13. Kifer, M., Lausen, G., Wu, J.: Logical foundations of object-oriented and frame-based languages. *Journal of ACM* 42, 741–843 (July 1995)
14. Kyburg, H., Teng, C.: *Uncertain Inference*. Uncertain Inference, Cambridge University Press (2001)
15. Riguzzi, F., Swift, T.: Well-definedness and efficient inference for probabilistic logic programming under the distribution semantics. *Theory and Practice of Logic Programming* 13(2), 279–302 (2013)
16. Swift, T., Warren, D.: XSB: Extending the power of prolog using tabling. *Theory and Practice of Logic Programming* (2011)
17. Swift, T.: Incremental tabling in support of knowledge representation and reasoning. *TPLP* 14(4-5), 553–567 (2014), <http://dx.doi.org/10.1017/S1471068414000209>
18. Wan, H., Grosf, B., Kifer, M., Fodor, P., Liang, S.: Logic programming with defaults and argumentation theories. In: *Int'l Conference on Logic Programming* (July 2009)

19. Wan, H., Kifer, M., Groszof, B.: Defeasibility in answer set programs with defaults and argumentation rules. *Semantic Web Journal* (2014)
20. Yang, G., Kifer, M.: Implementing an efficient DOOD system using a tabling logic engine. In: *First International Conference on Computational Logic, DOOD-2000 Stream* (July 2000)
21. Yang, G., Kifer, M.: Reasoning about anonymous resources and meta statements on the Semantic Web. *Journal on Data Semantics, LNCS 2800 1*, 69–98 (September 2003)
22. Yang, G., Kifer, M., Zhao, C.: FLORA-2: A rule-based knowledge representation and inference infrastructure for the Semantic Web. In: *International Conference on Ontologies, Databases and Applications of Semantics (ODBASE-2003)*. *Lecture Notes in Computer Science*, vol. 2888, pp. 671–688. Springer (November 2003)