

Tutorial Description for RuleML+RR-17 conference

Title: "Rulelog: Highly Expressive Semantic Rules with Scalable Deep Reasoning"

Presenters: Benjamin Grosf, Michael Kifer, and Paul Fodor (see bios/contact below)

Duration: half-day (3 hours plus break)

Tutorial Abstract: In this half-day tutorial, we cover the fundamental concepts, key technologies, emerging applications, recent progress, and outstanding research issues in the area of Rulelog, a leading approach to fully semantic rule-based knowledge representation and reasoning (KRR). Rulelog matches well many of the requirements of cognitive computing. It combines deep logical/probabilistic reasoning tightly with natural language processing (NLP), and complements machine learning (ML). Rulelog interoperates and composes well with graph databases, relational databases, spreadsheets, XML, and expressively simpler rule/ontology systems – and can orchestrate overall hybrid KRR. Developed mainly since 2005, Rulelog is *much* more expressively powerful than the previous state-of-the-art practical KRR approaches, yet is computationally affordable. It is fully semantic and has capable efficient implementations that leverage methods from logic programming and databases, including dependency-aware smart caching and a dynamic compilation stack architecture.

Description (continuing from the Abstract):

Rulelog extends Datalog (database logic) with general classical-logic-like formulas – including existentials and disjunctions – and strong capabilities for meta knowledge and reasoning, including higher-order syntax, flexible defeasibility and probabilistic uncertainty, and restraint bounded rationality that ensures worst-case polynomial time for query answering. A large subset of Rulelog is in draft as an industry standard. An exciting research frontier is that Rulelog can combine closely with NLP to both interpret and generate English, including potentially for conversational NL interaction.

The most complete system today for Rulelog is Ergo from Coherent Knowledge. A subset of Rulelog is also implemented in an open-source Flora-2 system (a.k.a. Ergo Lite) and an earlier SILK system from Vulcan. Using Ergo, we will illustrate Rulelog's applications in deep reasoning and representing complex knowledge – such as policies, regulations/contracts, science, and terminology mappings – across a wide range of tasks and domains in business, government, and academe. Examples include: legal/policy compliance, e.g., in financial services; financial reporting/accounting; health care treatment guidance and insurance; education/tutoring; security/confidentiality policies; and e-commerce marketing.

Additional Outline: (Note: Examples, and discussion, are sprinkled throughout)

- Introduction: logical knowledge representation; practical logic.
- Rulelog features and software: strong meta expressiveness; reasoning methods; scalability; integration points with databases, ontologies, NLP, and ML.
 - Concepts and foundations, in depth: well-founded semantics, “tabling” algorithms, rule-based mapping between text phrases and logic syntax, higher-order syntax (HiLog), defeasibility via argumentation rules, general quantified formulas (omniformity), probabilistic and weighted uncertainty, restraint bounded rationality, frame syntax, rule identifiers and provenance, external querying and virtual data stores, updating and reactivity, integrity constraints and alarms, explanations.
- Case study demos, with feature tour: financial regulatory compliance; health care treatment guidance.

- Applications overall: Horizontally: policy-based decisions, info Integration, analytics, human-computer interaction (HCI), search, business intelligence, risk management. Vertically: e-commerce & marketing, financial services, personalized e-learning, security & defense, biomedical, insurance, Internet of Things (IoT), social media sharing policies.
- Open research topics in: authoring rules starting from NL; distributed reasoning; optimization of uncertainty reasoning; equality; aggregates; integration with ASP, “constraint solving”, and classical logic; hypotheticals; abduction; integration with ML;

Goal of the tutorial: The audience will walk away with an understanding of Rulelog’s key innovative logical and inferencing concepts, its broad applicability, its overall advantages and limitations, a sample of some specific application areas, and its open research topics.

Intended audience: All of the RuleML+RR audience. Background assumed of participants is only the basics of first-order-logic and relational databases. Knowledge of declarative logic programs, XML, RDF, and SPARQL will be helpful but not required.

Previous related tutorials: This proposal constitutes the first-ever half-day conference tutorial specifically on Rulelog KRR to the RuleML+RR conference community. The authors have presented – in some cases with other coauthors – related tutorials on reasoning with complex knowledge at the AAAI Conference on Artificial Intelligence (2013, 2017), International Joint Conference on Artificial Intelligence (2001, 2016), International Semantic Web Conference (2004, 2005, 2006, 2009, 2010, 2012), WWW conference (2006, 2009), International Web Rule Symposium (2015, 2016; 2015 was joint with Reasoning Web Summer School), and ACM Conference on E-Commerce (2004). Portions of those semantic rules tutorials covered aspects of what later became Rulelog. Each of the above previous tutorials was attended by approximately 40-60 people (or, in a few cases, more).

Biographies and Contact:

Benjamin Groszof (lead presenter) is Principal Director and Research Fellow in Artificial Intelligence at Accenture. He is an industry leader in AI knowledge representation, reasoning, and acquisition. He has pioneered semantic technology and industry standards for rules combined with ontologies, their acquisition from natural language, and their applications in finance, e-commerce, policies (including contracts, regulations, and security), and e-learning. He co-founded Coherent Knowledge, was its CTO and CEO, and is a board member. He led a large research program at Vulcan Inc. (predecessor of the Allen Institute for AI), was an IT professor at MIT Sloan where was also a DARPA PI), and was a senior scientist at IBM Research where he released two semantic rules products. He co-founded RuleML and co-chaired its first 5 years.

<http://benjamingroszof.com> ; benjamin.groszof@gmail.com

Michael Kifer is a Professor in Computer Science at Stony Brook University. He is co-founder and CTO of Coherent Knowledge. He co-invented F-logic, HiLog, and Transaction Logic, which are among the most widely cited works in Computer Science and, especially, in Semantic Web research. He has thrice received “Test of Time” awards for his work, from ACM-SIGMOD and Association of Logic Programming. Since 2012 he has been serving as the President of the Rules and Reasoning Association (RRA).

<http://www.cs.stonybrook.edu/~kifer> ; kifer@cs.stonybrook.edu

Paul Fodor is a Research Assistant Professor in Computer Science at Stony Brook University, where he has taught extensively. He is co-founder and Senior Engineer at Coherent Knowledge. He was a member of the Vulcan SILK and IBM Watson project teams, contributing to their logical reasoning aspects in combination with NLP. He has over 10 years’ experience in databases research, natural language processing, semantic web, artificial intelligence, complex event and stream processing systems.

<http://www.cs.stonybrook.edu/~pfodor> ; pfodor@cs.stonybrook.edu

