

Syllabus Description of AAI-10 Conference Tutorial SP3 to be held July 11, 2010

“Rules on the Semantic Web: Advances in Knowledge Representation and Standards”

- Benjamin Grosf (main presenter), Vulcan Inc. <<http://www.mit.edu/~bgrosf/>>
- Mike Dean (secondary contact), BBN Technologies <<http://www.daml.org/people/mdean>>
- Michael Kifer, Stony Brook University <<http://www.cs.sunysb.edu/~kifer/>>

Abstract

The area of semantic rules is perhaps the most important frontier today for the Semantic Web’s core technology and standards. Rules extend databases and ontologies with more powerful, flexible, and active forms of “structured” knowledge (as opposed to “unstructured” knowledge such as text), and have a number of close relationships to other aspects of the overall Semantic Web and AI such as ontologies, query/search, trust, and processes/services. Recent progress includes major initial industry standards from W3C and OMG nearing finalization, and fundamental advances in the underlying knowledge representation techniques in declarative logic programs, including most recently for efficient higher-order defaults with sound integration of first order logic ontologies (OWL). Recent progress also includes methods to use rules for, or with, more expressive OWL ontologies; increasing integration of rules with query/search in RDF/SPARQL and relational databases; substantive translations between heterogeneous types of commercial rule engines; development of open-source tools for inferencing and interoperability; performance benchmarking of rule systems; a wide range of emerging applications including in business, science, and trust; and accelerating industry investments/acquisitions in the technology including by integrated software companies such as Oracle, IBM, and Microsoft. This tutorial will provide a comprehensive and up-to-date introduction to these developments and to the fundamentals of the key technologies and outstanding research issues involved. It will explore example application scenarios, overall requirements and challenges, and touch upon business/social value and strategy considerations. It will focus particularly on the recent advances in knowledge representation and standards. There are a number of exciting research issues; the most fundamental challenge is how to integrate knowledge gained from text understanding and machine learning.

Target Audience and Prerequisite Knowledge

This tutorial is suitable for most of the AAI audience, especially those interested in rules with ontologies on the Web, and/or in logical knowledge representation. This includes researchers interested in core technologies, those interested in standards and applications, and those interested in closely related areas such as query and search, collective intelligence, ontologies, policies, e-commerce, and trust. The tutorial will cater to those first learning about semantic web rules, as well as those who already have some background in them. It will assume only background knowledge of the basics of logical knowledge representation (first order logic and relational DBMS), and basics of XML.

Outline of Content and Schedule

Note: Examples (use cases and application scenarios) will be liberally embedded throughout all the sections. Standards, languages, and systems will be interleaved throughout. Standards covered will

include W3C RIF and OWL 2, as well as more lightly touching upon several others. The details in the bulletized outline below focus primarily on the technology and its uses, for sake of brevity.

A. Intro & Uses (~30 min)

1. Overview of tutorial, and get acquainted
2. What are: Rules on the Web, Semantic Rules
3. What are: RDF and OWL
4. Uses and Kinds of rules
 - Commercial, web. Current, envisioned.
 - Requirements. Business value, IT lifecycle.
 - Strategic roadmapping of future adoption
5. Example Use Cases
 - E-commerce: pricing, ordering policies, contracts
 - E-Science: ecological process, mechanics context
 - Trust, compliance, policies, e.g., financial services
 - Info integration, ontology mapping, business reporting
 - Processes: policy-based workflow, causal action effects, Semantic Web Services

B. Concepts & Foundations (~160 min)

1. Overview of Logical Knowledge Representations
 - Logic Programs (LP) and its relationship to First Order Logic (FOL)
 - Rule-based Ontologies: Description Logic, Description LP, OWL RL
2. Basics: Horn Case; Functions
3. F-Logic, Frame Syntax, Object Oriented Style
4. HiLog, Higher-Order Syntax, Reification, Meta-Reasoning
5. W3C Rule Interchange Format: Dialects, Framework
6. Nonmonotonicity: Defaults, Negation, Priorities; FOL's Glass Bubble
 - Semantics for Default Negation
 - Courteous LP, Argumentation Theories
 - Hypermonotonic Mapping: FOL \Leftrightarrow LP, Soundly
7. Procedural Attachments to Actions, Queries, Built-ins, and Events
 - Production/Situated LP, Production Rules
8. Additional Features: Integrity Constraints, Inheritance, Lloyd-Topor, Equality, Skolemization, Aggregation, Datatypes, "Constraints"
9. Hyper LP and SILK – Putting it all together

C. Conclusions & Directions (~20 minutes)

1. More about Tools
2. Conclusions
3. Directions for Future research
4. Appendix: References and Resources
5. General Discussion

Coffee Break: The above time budget of 3.5 hours reflects a half-hour coffee break somewhere in the middle of the 4 hours total for the tutorial.

Tutorial Materials

A detailed tutorial slideset, including extensive links and references to resources and papers, will be provided free in .pdf form to the tutorial attendees, and afterwards posted publicly free on the web. There will be very approximately 100-150 main slides.

Presenters' Biographies (including expertise and presentation experience related to the tutorial)

Benjamin Grosf (lead presenter) is a Senior Research Program Manager at Vulcan Inc., the parent company of Paul G. Allen (co-founder of Microsoft). There he conceived and leads a new large research program in the area of rule-based semantic technologies and artificial intelligence. In addition, he has a part-time expert consulting business, advising companies large and small on technology and related strategy. Previously he was an IT professor at MIT Sloan (2000-2007) and a senior software scientist at IBM Research (1988-2000). He has pioneered semantic technology and standards for rules, their combination with ontologies, their application in e-commerce and business policies, and business roadmapping of the Semantic Web. He co-founded the influential RuleML industry standards design effort. He was lead inventor of the rule-based technique which rapidly became the currently dominant approach to commercial implementation of OWL, and of several other fundamental technical advances in knowledge representation. Two recent W3C industry standards are based largely on his work: Rule Interchange Format (RIF, now in last stage of finalization) and OWL 2's RL rule-based subset (Nov. 2009). His notable technical contributions also include fundamental advances in conflict handling for rules (i.e., defaults) and integration of rules with machine learning. He co-founded the International Conference on Rules and Rule Markup Languages for the Semantic Web (which since became the RR and RuleML conferences). His background includes three major industry software releases, two years in software startups, a Stanford PhD, a Harvard BA, and over 50 refereed publications.

Dr. Grosf has given numerous invited talks about rules on the Web, and developed several MIT courses with substantial focus on it. He presented (with co-authors) related tutorials at the Intl. Joint Conf. on Artificial Intelligence (2001), ACM Conf. on E-Commerce (2004), Intl. Semantic Web Conf. (2004, 2005, 2006, 2009), and the WWW conference (2006, 2009).

Mike Dean is a Principal Engineer at BBN Technologies. As Principal Investigator for the DAML Integration and Transition effort within the DARPA Agent Markup Language (DAML) program, he chaired the Joint US/EU ad hoc Markup Language Committee responsible for the DAML+OIL and SWRL languages, co-edited the OWL Web Ontology Language Reference, and was a member of the W3C RDF Core and Web Ontology Working Groups and the Architecture Committee of the Semantic Web Services Initiative. He remains a member of the RuleML Steering Committee and the W3C Rule Interchange Format Working Group. He is the developer of a number of Semantic Web tools and reference data sets and has been actively using SWRL in a variety of Semantic Web applications. He holds a B.S. in Computer Engineering from Stanford University.

Mr. Dean has given numerous talks on the Semantic Web, including an early tutorial "DAML+OIL for Application Developers" and tutorials (with Benjamin Grosf) at the Intl. Semantic Web Conf. (2004, 2005, 2006, 2009) and WWW (2006, 2009).

Michael Kifer is a Professor with the Department of Computer Science, State University of New York at Stony Brook, USA. He received his Ph.D. in Computer Science in 1984 from the Hebrew

University of Jerusalem, Israel, and the M.S. degree in Mathematics in 1976 from Moscow State University, Russia. Dr. Kifer's interests include Web information systems, knowledge representation, and database systems. He has published four text books and numerous articles in these areas. In particular, 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. Dr. Kifer serves on the editorial boards of several computer science journals and chaired several conferences. Twice, in 1999 and 2002, he was a recipient of the prestigious ACM-SIGMOD "Test of Time" awards for his works on F-logic and object-oriented database languages. In 2006, he was a Plumer Fellow at Oxford University's St. Anne's College and in 2008 he received SUNY Chancellor's Award for Excellence in Scholarship.