

Tutorial Syllabus for AAAI-17 conference

Title: "Rulelog: Deep KRR for Cognitive Computing"

Presenters: Benjamin Grosz, Michael Kifer, and Paul Fodor

- See below for bios and additional contact info.

Duration: half-day (3.5 hours, excluding 30-minute break).

Tutorial description (i.e., 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 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, and complements machine learning. 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 representationally 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.

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 machine learning and with natural language processing 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 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.

Background assumed of participants is only the basics of first-order-logic and relational databases.

Goal of the tutorial: The target audience is those who are interested in learning more about knowledge representation and reasoning (KRR) as a core area in AI, and about how it relates to cognitive computing's applications as well as conceptual architecture. Together with machine learning (ML) and natural language interaction (NLI), KRR forms the tripod basis for the core of AI and cognitive computing. 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.

Preliminary version of Detailed Outline:

Note: Examples, and discussion with the audience, are sprinkled throughout

- Introduction and Overview
 - Pointers to more info
 - Concept of logical Knowledge Representation
 - Practical Logic
 - Contrast with classical logic
 - Main kinds: databases (graph and relational), production rules, FOL, others; RDF, SPARQL, OWL
 - Concepts and advantages of “smart data”, “smart rules”, “semantic”, hot areas in industry
 - Rulelog features and software
 - Expressiveness. Strong Meta.
 - Reasoning methods and scalability
 - Integration points with databases, ontologies, natural language processing (NLP), machine learning (ML)
 - Example Rulelog system architecture
 - Software tools
 - Applications (examples)
 - 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, E-Science, Security & defense, Health treatment, Insurance, Internet of Things (IoT), social media sharing policies
 - Textual Rulelog: combining closely with natural language processing
 - Textual terminology
 - Rule-based mapping between logic and NL for text generation and text interpretation
 - Templates (in Ergo, a.k.a. ErgoText)
 - Authoring process & methodology (for authoring of knowledge)
 - KRR and Machine Learning
 - How KRR Needs Machine Learning: Knowledge Base Construction
 - How Machine Learning Needs KRR: Accumulating knowledge, Explaining knowledge, Supplying derived data, Posing questions, Focusing Features

- Probabilistic and Weighted Uncertainty
 - NL as an area that brings KRR and ML together
 - Knowledge management challenges addressed
 - Value, variety, and veracity – as well as volume and velocity
- Case Study Demo and Features Tour
 - Financial Regulatory/Policy Compliance domain
 - Public application pilot by Enterprise Data Management Council
 - Drilldowns to facts, non-fact rules, exception rules, data/ontology import, ontology/terminology mapping
 - Representational use of higher-order, meta knowledge, defeasibility
 - Explanations, text generation, interactive navigability
 - Editor tools: syntax checking; dependency analysis; term search
- 2nd Case Study: Health Care Treatment Guidance and Insurance domain
- Concepts and Foundations
 - Overview
 - Horn LP, with Functions (LP = declarative logic programs)
 - Horn FOL (FOL = First Order Logic)
 - Horn LP syntax and semantics
 - Comparison to Horn FOL, incl. fundamental theorem
 - The “spirit” of LP: avoid disjunction, stay grounded
 - Logical Functions: general requirements analysis for KRR
 - Computational complexity analysis
 - Propositional: linear
 - Function-free: polynomial if VB (VB = bounded # of variables per rule)
 - Functions without restraint: undecidable
 - Functions with restraint: polynomial if VB
 - Horn Datalog: function-free Horn LP
 - Well-Founded Negation
 - Logical nonmonotonicity
 - Negation-as-Failure (naf; a.k.a. weak or default negation)
 - Declarative LP beyond Horn
 - Well-Founded Semantics for LP
 - Third truth value: “undefined”. Use to represent paradox.
 - Comparison to: Stable Semantics for LP; Answer Set Programs
 - RDF, SPARQL, OWL: Revisited
 - “Tabling” Algorithms for LP & Rulelog
 - Smart cacheing of inferencing results (mixed direction not backward)
 - Forest of subgoal attempts and results
 - Challenge from nonmonotonicity and higher-order syntax
 - Dependency-awareness: table incrementally upon updates
 - In-memory: but combine with external querying to scaled-out DBMS
 - External Querying

- Built-ins
 - Orchestrating ensembles of databases and knowledge-based systems
- Restraint: semantic bounded rationality using “undefined”
 - Radial restraint: generalized term depth/size limit
 - Polynomial time complexity with dial-able degree, for function-ful
 - Skipping restraint: control per rule instance
 - Unsafety in weak negation and external querying
 - Restraint for unsafety and failure to return
- Frame syntax (a.k.a. F-Logic), Object Oriented style
 - Transformation sketch
 - Relationship to RDF and OWL
- Higher-Order Syntax via Hilog
 - Transformation sketch
 - Reification
- Rule identifiers (ID’s) within the logical language
 - Meta knowledge to represent source provenance of rules
- Defeasibility via Argumentation Rules
 - Conflict (inconsistency flavor)
 - Prioritization knowledge: its bases, its mechanics
 - Representational uses for defeasibility
 - Transformation sketch
- FOL-Soundness
 - Remedying FOL’s Fragility
- General Formulas (Omniformity)
 - Existentials and Skolems
 - Omni-directional Disjunction
 - Transformation sketch
- Representing Text: Revisited
 - Textual Terminology
 - Text generation, using Rulelog, e.g., for explanations
 - Text interpretation, using Rulelog, e.g., for authoring via templates
- Importing structured knowledge automatically
 - External querying via SPARQL. Databases. Spreadsheets. RDF. Full OWL.
 - Axiomatic semantics for description logic variants
- Ontology/Terminology Mapping
- Explanation as a critical tool
 - Role in development lifecycle of knowledge bases
 - Justification graphs
 - Transformations sketch
- Probabilistic and Weighted Uncertainty
 - Evidential reasoning
 - Combine results from multiple sources: machine learning, search, heuristic algorithms
 - Distribution Semantics: probabilities on more complex rules

- Represent more than Bayesian Networks
 - Contrast with Markov Logic Networks
 - Reactiveness: knowledge update Events, side-effectful Actions
 - Misc. Lesser Features
 - Datatypes, Aggregation, Integrity Constraints, Inheritance, Equality, “Constraints”
- Conclusions and Future Work
 - Other summary observations
 - Essential expressiveness features
 - Transformation/compilation stack
 - Advantages and limitations for knowledge management
 - Complex knowledge. Deep reasoning. Orchestration. The 5 V’s.
 - Big picture KRR-wise. The 5 V’s (variety, value, etc.).
 - Open Research Topics in the KRR itself
 - Distributed reasoning. Optimization, e.g., for probabilistic uncertainty.
 - Reactive. Streaming.
 - Reasoning-by-cases. Integration and hybridization with classical logic, “constraint solving”, and answer set programs. Hypotheticals, abduction.
 - Equality. Aggregates.
 - Transactionality and process flow
 - Research Directions – Other Aspects
 - Combination with Machine Learning
 - Combination with Natural Language Processing
 - Applications
 - Industry standards design
- Appendix: References and Resources

Presenters’ Bios:

Benjamin Groszof (lead presenter), Principal Director and Research Fellow in AI at Accenture, is a technical executive for an ambitious initiative on leveraging AI for business process automation across Accenture’s \$7.5B+ Operations group. 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, a software-centric startup that is commercializing a major research breakthrough in AI logical / probabilistic knowledge representation and reasoning (Rulelog) combined with natural language processing. Previously he was CTO and CEO of Coherent Knowledge (2013-2017; he is still a board member), and a senior research program manager at Vulcan Inc. (2007-2013), where he conceived and led the advanced research third of the predecessor of the Allen Institute for AI. Before Vulcan, he was an IT professor at MIT Sloan (2000-2007), where he was also a DARPA PI, and a senior software scientist at IBM Research (1988-2000). He cofounded the influential RuleML industry standards design effort and prototyped it in SweetRules, the main bases for the W3C Rule Interchange Format (RIF) standard. He cofounded the International Conference on Rules and Rule Markup Languages for the semantic web (which since became the RR and RuleML conferences and then the International Joint Conference on Rules and Reasoning). He

led the invention of several fundamental technical advances in knowledge representation, including courteous defeasibility (exception-case rules), restraint bounded rationality (scalability in complex reasoning), and rule-based description logic ontologies (the basis for W3C's OWL RL standard). He also has extensive experience in user interaction design, and in combining logical methods with machine learning and probabilistic reasoning uncertainty. His background includes 5 major industry software releases, 5 years in software startups, his own part-time expert consulting firm, a Stanford PhD in AI, a Harvard BA in applied math, 3 patents, and over 60 refereed publications.

Grosz has given numerous invited talks about knowledge representation, reasoning, and acquisition, including with semantic rules, and developed several MIT courses with substantial focus on those. He presented with coauthors (including usually Michael Kifer since 2009) related tutorials on reasoning with complex knowledge at the AAAI Conference on Artificial Intelligence (2013), International Joint Conference on Artificial Intelligence (2001, 2016), the International Conference on Knowledge Capture (K-CAP 2015), International Semantic Web Conference (2004, 2005, 2006, 2009, 2010, 2012), the WWW conference (2006, 2009), the International Web Rule Symposium joint with the Reasoning Web Summer School (2015), and the ACM Conference on E-Commerce (2004).

Michael Kifer is a Professor with the Department of Computer Science, Stony Brook University, 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.

Kifer is co-founder and CTO of Coherent Knowledge, a new startup on semantic technology, and since 2012 he has been serving as the President of the Rules and Reasoning Association (RRA). His 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. Kifer serves on the editorial boards of several computer science journals and chaired a number of 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 2013, Kifer's paper on Transaction Logic Programming was awarded the Association of Logic Programming's "Test of Time" award as the most influential paper of 20 years ago. In 2006, Kifer was Plumer Fellow at Oxford University's St. Anne's College and, in 2008, he received SUNY Chancellor's Award for Excellence in Scholarship. He has taught numerous courses at Stony Brook University since 1984.

Paul Fodor is a Research Assistant Professor with the Department of Computer Science, Stony Brook University, USA. He received his Ph.D. in Computer Science in 2011 from the Stony Brook University, New York, preceded by his M.S. degree in 2006 from Stony Brook University, and B.Sc. in Computer Science in 2002 from the Technical University of Cluj-Napoca, Romania.

Dr. Fodor is a co-founder of Coherent Knowledge with over 10 years' experience in databases research, natural language processing, artificial intelligence and stream processing systems. His work on declarative rule languages and logic used as a specification language and implementation framework for knowledge bases was applied in areas ranging from natural

language processing to complex event processing and semantic Web technologies. Through his research, Dr. Fodor has contributed to several large software projects: the IBM Watson natural language processing system for the Jeopardy! Challenge with human champions, the OpenRuleBench suite of benchmarks for analyzing the performance and scalability of rule systems for the semantic Web, the ETALIS declarative complex event processing and stream reasoning system, and the SILK Semantic Inferencing on Large Knowledge. Dr. Fodor was Principal Investigator (PI), Co-PI and contractor for projects funded by both public governmental sources and private companies, such as, PI for the SILK project funded by Vulcan Inc. to develop intelligent textbooks, contractor for the IBM Watson project, contractor for XSB Inc. for the DARPA Component, Context, and Manufacturing Model Library (C2M2L-1) using XSB Prolog, and PI for the Stony Brook University Hospital's Lung Cancer Evaluation Center management program. He has taught numerous courses at Stony Brook University since 2011.

Detailed Contact Info of Presenters:

Benjamin Grosf (**primary contact**)

Principal Director and Research Fellow, Artificial Intelligence
Accenture

5 Wembley Lane

Mercer Island, WA 98040, USA

<http://benjamingrosf.com>

FIRSTNAME DOT *LASTNAME* AT gmail DOT com

Michael Kifer

Professor, Department of Computer Science

Stony Brook University

Stony Brook, NY 11794-2424, USA

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

kifer AT cs DOT stonybrook DOT edu

Paul Fodor

Research Assistant Professor, Department of Computer Science

Stony Brook University

Stony Brook, NY 11794-2424, USA

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

pfodor AT cs DOT stonybrook DOT edu