

WELCOME! to the KR-2018 Conference Tutorial (T6)

Rulelog: Highly Expressive, Yet Scalable, Semantic Rules

Presented and co-authored by **Benjamin Grosf***;

Also co-authored by **Michael Kifer****, **Paul Fodor****, and **Janine Bloomfield*****

3 hours. Oct. 29, 2018, Tempe, AZ, USA.

These slides available via link at the KR-2018 Tutorials webpage and directly at:

<http://benjamingrosof.com/misc-publications/#KR2018RulelogTutorial>

* Kyndi ; work in part done while previously at Coherent Knowledge

<http://benjamingrosof.com>

** Stony Brook University, and Coherent Knowledge

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

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

*** Coherent Knowledge

<http://coherentknowledge.com>

<https://www.linkedin.com/in/janinebloomfield>

Note on Copyright and Permissions:

Slides with Coherent Knowledge logo are, in whole or part, courtesy of Coherent Knowledge; and several other slides as well are, in whole or part, courtesy of Coherent Knowledge.

Preface ; For More Info

- Previous tutorials on Rulelog and its forerunners were presented at AAAI-18, RuleML+RR-2017, AAAI-17, IJCAI-16, RuleML-2016, RuleML-2015, AAAI-13, and ~9 other times at conferences since 2004.
- This slideset is, or soon will be, available on the web, at <http://benjamingrosof.com/misc-publications/#KR2018RulelogTutorial>, and via link to that on the KR-2018 website's Tutorials page
- For more info beyond the slideset: see the authors' websites and also <http://coherentknowledge.com/publications>
- References are in the short proceedings papers of the RuleML+RR-2017 and RuleML-2015 tutorials, and in the AAAI-13 tutorial slides near the end
- Hoping to turn the tutorial material into a book, suitable as a course unit, at some point. (Interest from several publishers.)

Bio – Benjamin Grosf

- AI researcher, executive, and entrepreneur
- Chief Scientist at Kyndi – AI startup on NL question answering using ML+KRR
- Co-founder & Board of Coherent Knowledge – AI startup on KRR
- *Previously:*
 - Principal Director & Research Fellow in AI at Accenture on BPM
 - CTO & CEO of Coherent Knowledge
 - Directed advanced AI research at predecessor of Allen Inst. for AI
 - Developed Rulelog KRR theory, algorithms, UI approach
 - MIT Sloan professor and DARPA PI
 - Co-Founder of RuleML, key contributor to W3C OWL-RL and RIF standards
 - IBM Research, creator IBM Common Rules
 - 1st successful semantic rules product in industry
 - Stanford AI PhD, combining ML with logical and probabilistic reasoning
- *Themes: flexible clean KRR + NL + ML; many app domains & tasks*
- <http://www.linkedin.com/in/benjaminrosf>
- <http://benjaminrosf.com>



Outline of Tutorial

A. Overview

- Practical logic. Applications. Features. Software. Textual.
- Case Study Demo and Features Tour
 - Financial regulatory/policy compliance

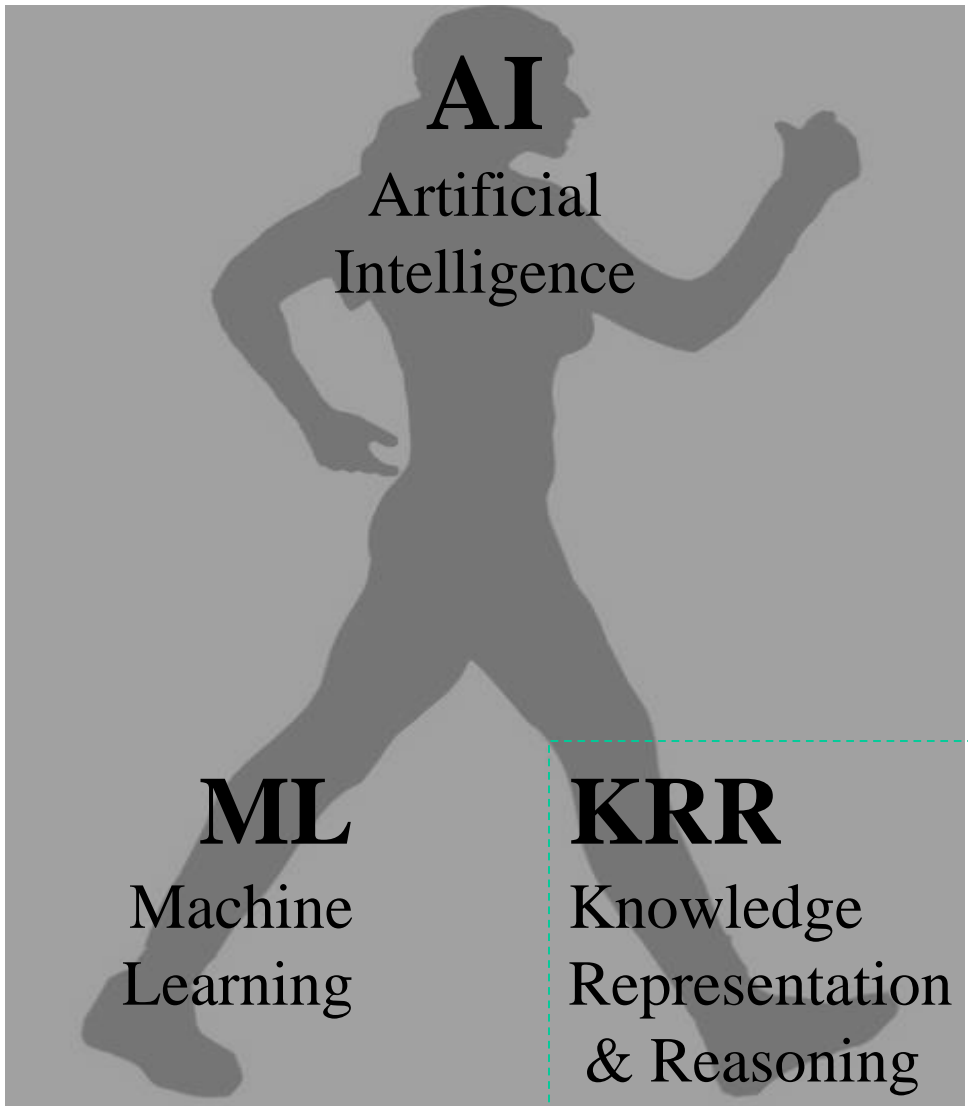
B. Concepts and Foundations: Drill-downs

- Expressive features, semantics, algorithms;
relationships to natural language and machine learning

C. Conclusions and Future Work

- Background Assumed: basic knowledge of first-order logic, databases, rules. Helpful: RDF and semantic web concepts

The Core of AI



It takes two legs to walk

Focal Kinds of KRR, in this talk

- **Rulelog** – highly expressive extension of logic programs (LP)
 - Blends with higher-order classical logic
 - Flexible yet scalable
 - Family of KRs that are subsets (fragments)
- **Probabilistic logic programs (PLP)** – also extension of LP
 - In several flavors: Bayesian, fuzzy based, lattice based
 - Ongoing work: unifying with all of Rulelog expressiveness
- Background: **LP** = (well-founded declarative) **logic programs** is the core KR of the entire *IT* world, not just of AI
 - It's a logic, despite the “programs” in its name
 - Invented by computer scientists not mathematicians
 - Developed to formalize relational databases and unify that with the pure subset of Prolog

Logic Programs Example

- “:-” means “if”, i.e., \impliedby
- **Assertions:**
 - human(Socrates). human(Anne). human(Peter).
 - modern(Anne). educated(Socrates).
 - mortal(?x) :- human(?x).
 - fallible(?z) :- mortal(?z).
 - educated(?x) :- human(?x) \and modern(?x).
 - humble(?w) :- fallible(?w) \and educated(?w).
- **Forward chaining:** mortal(Socrates). mortal(Anne). mortal(Peter).
fallible(Socrates). fallible(Anne). fallible(Peter). educated(Anne).
humble(Socrates). humble(Anne).
- **Query:** ?- humble(?p). **Answer:** {Socrates, Anne}
- **Backward chaining subgoals:**

?- fallible(?p),	?- educated(?p).	
?- mortal(?p).	?p / S	
?- human(?p).	?- human(?p),	?- modern(?p).
?p / S, A, P	?p / S, A, P	?p / A

Practical Logic, vs. Classical Logic

- Goal: support **IT**, vs. mathematics
 - E.g., Databases, Rules
 - Central: declarative logic programs (**LP**) KR
 - Desirable to have **expressiveness** enough for
 - Natural language (NL) cognitive abstraction level, structures, semantics
 - Machine learning (ML) results
- Requirements:
 - **Scalable** computationally
 - **Robust** in face of human errors and miscommunications
 - $\rightarrow\rightarrow$ “**Humble**”
 - Avoid general reasoning by cases
 - Avoid general proof by contradiction

What is “reasoning by cases”: (background)

Assertions: if A then C. if B then C. A or B.

Conclude: C.

Semantic

- “Semantic” rules/tech/web means: based on logic
- Advantages for communication across systems and organizational boundaries
- Meaning is shared notion of what is/is-not inferrable
- Abstracts away from implementation

- Relational DB was 1st successful semantic tech
- LP theory was invented to formalize it and unify it with the pure subset of Prolog

Concept of logical Knowledge Representation (KR)

A given KR logical system S has ...

1. Formal language L_S for assertions and conclusions
 - Assertions LA_S and conclusions LC_S may be different!!
 - E.g., in LP and Rulelog
2. Semantics: entailment relation \models_S
 - An assertion set A entails a* conclusion set C
 - * We assume here exactly 1
 - Typically, entailment is defined formally in terms of models
 - Truth assignments on LC_S that meet criteria based on the assertions A
 - E.g., in FOL and LP and Rulelog

Reasoning implements the semantics, e.g., to answer queries

- KRR software system: knowledge representation & reasoning (proving)
- Typically KRR systems are sound but often incomplete

LP is the Central Form of Practical Logic

- *LP is the core KR of structured knowledge management today*
- A non-classical logic invented by computer scientists
- Subsets of LP – important in industry landscape today
 - Relational databases (SQL) *[Datalog subset of LP]*
 - Graph databases, a.k.a. knowledge graphs (SPARQL) *[Datalog]*
 - Also: XML databases, object-oriented databases, other semi-structured databases
 - Production rules, Event-Condition-Action rules. More precisely: their logical subsets.
 - Prolog. More precisely: its “pure” logical subset.
 - *Also industry standards for semantic rules and ontologies:*
 - Many RuleML & RIF dialects, e.g., RIF-BLD, RIF-Core, SWRL
 - Many ontology standards, e.g., OWL-RL, RDF-Schema *[Datalog]*

Industry Landscape of Practical KRR

- LP and subsets (*cf. earlier slide*)
- Subsets of Classical Logic:
 - Propositional. E.g., hardware circuit design, satisfiability for planning.
 - First Order Logic (Common Logic). E.g., for program verification.
 - Description Logic (OWL) subset of FOL. For ontologies.
- Emerging: (in roughly descending order of commercial maturity)
 - **Rulelog – extension of LP**
 - RIF/RuleML Rulelog dialect standard is in draft
 - Bayesian Probabilistic LP
 - Bayesian Networks are a special case
 - Fuzzy Probabilistic LP
 - Probabilistic Soft Logic – is closely related
 - Other probabilistic graphical models (PGM)
 - Markov Logic Networks – closer to classical; thus more difficult to scale
 - Answer Set Programs – closer to classical, less humble
 - (Others are not so commercially/practically prominent)

Some State-of-the-Art KRR Systems

- LP: XSB (Stonybrook U., Theresa Swift, David Warren, et al)
 - Full programming language that is Prolog+
- Rulelog: ErgoAI (Coherent Knowledge, free for research), and its open-source subset Flora-2 (originally Stonybrook U.)
 - Full programming language that is Prolog++ and XSB++
- Bayesian Probabilistic LP: Problog (Luc de Raedt et al, originally KU Leuven)
- Probabilistic Soft Logic: (Lise Getoor et al, UC Santa Cruz)
- Markov Logic Networks: Alchemy (Pedro Domingos et al, U. Washington)
- Probabilistic Graphical Models generally: See STARAI workshops

Outline of Tutorial

A. Overview

- Practical logic. Applications. Features. Software. Textual.
- Case Study Demo and Features Tour
 - Financial regulatory/policy compliance

B. Concepts and Foundations: Drill-downs

- Expressive features, semantics, algorithms;
relationships to natural language and machine learning

C. Conclusions and Future Work

- Background Assumed: basic knowledge of first-order logic, databases. Helpful: XML, RDF and semantic web concepts

Logic Programs: technical overview (I)

- Knowledge base (KB) is a set of rules, each of form:
 - *Head_formula* :- *Body_formula*.
 - Intuitively: OK to infer (establish) the head if can infer the body
- Basic normal LP: each rule has form:
 - $atom \text{ :- } literal_1 \ \&\& \dots \ \&\& \ literal_m.$
 - Plus: atoms are all first-order
- *atom* has form: $(predicate(arg_1, \dots, arg_k))$, where each arg_i is a term
- *literal* has form: $(atom)$ or $(\text{\textbackslash}naf \ atom)$
- Weak negation: $\text{\textbackslash}naf \ p$ – p is not believed (essentially, known to be not provable)
- Strong negation: $\text{\textbackslash}neg \ p$ – p is believed to be strongly false (i.e., opposite of true)
 - Not permitted in normal LP. But permitted in extensions of normal LP, e.g., in Rulelog.
- Aggregation: $setof\{?x \mid condition\}$, where $?x$ appears in *condition* formula
 - Enabled by $\text{\textbackslash}naf$. Aggregate operators also include avg, max, min, count.
 - $average_salary(?co, ?amt) \text{ :- } company(?co) \ \&\& \ avg(?amt \mid employee(?co, ?e) \ \&\& \ salary(?e, ?amt)).$

Logic Programs: technical overview (II)

- Horn subset: body literals are restricted to be atoms
- Datalog subset: Horn, and function-free
- Full normal LP permits also:
 - in head: \wedge
 - in body, freely nested: \vee , \forall , \exists , aggregates, \wedge , \neg
 - Integrity constraints via *violation(...)* as a head atom predicate
 - Reduces via transformation to basic normal LP
- Semantics (well-founded) is based on:
 - An alternating least fixed point construction in 3-valued logic
 - Each instantiated atom is assigned to 1 of 3 truth values $\{t, f, u\}$:
 - $t = \text{true}$; $f = \text{"false"}$ (cf. \neg); $u = \text{"undefined"}$ (don't-care).
 - *undefined* is useful for paradox and restraint bounded rationality
- Function-free case is polynomial time
- Functions (thru recursion) lead to undecidability

The “Spirit” of LP

The following summarizes the “spirit” of how LP differs from FOL:

- **“Avoid Disjunction”**
 - Avoid disjunctions of positive literals as expressions
 - In premises, intermediate conclusions, final conclusions
 - (conclude (A or B)) only if ((conclude A) or (conclude B))
 - Permitting such disjunctions creates exponential blowup
 - In propositional FOL: 3-SAT is NP-hard
 - In the leading proposed approaches that expressively add disjunction to LP with negation, e.g., propositional Answer Set Programs
 - No “reasoning by cases”, therefore
- **“Stay Grounded”**
 - Avoid (irreducibly) non-ground conclusions

LP, unlike FOL, is straightforwardly extensible, therefore, to:

- Nonmonotonicity – defaults, incl. NAF
- Procedural attachments, esp. external actions

Important Expressive Extensions of LP (I)

- **Rulelog** – reduces by efficient transformation to normal LP
 - Higher-order, reification; rule id's, provenance
 - Higher-order relies on (logical) functions
 - Defeasibility: prioritized defaults, exceptions, argumentation
 - Flexible behavior. Efficient approach.
 - Both weak negation ($\backslash\text{naf}$) and strong negation ($\backslash\text{neg}$); $\backslash\text{naf}$ is outside of $\backslash\text{neg}$
 - Restraint bounded rationality: guarantee polynomial time
 - Specify *undefined*-ness in various circumstances
 - General classical-like formulas:
 - Head quantifiers; $\backslash\text{exists}$ treated via skolemization
 - Head $\backslash\text{or}$, treated as “omnidirectional” (weak)
 - Object-oriented (“frame”) syntax
 - External queries; import of most kinds of enterprise info
 - Probabilistic via: higher-order, defeasibility
 - But current implementations not optimized and lack dedicated syntax
 - Flexible: can have tuple of parameters for the probability
 - $\text{pr}(\text{formula1})[\text{low}\rightarrow 0.91, \text{hi}\rightarrow 0.94]. \text{pr}(\text{formula2})[\mu\rightarrow 0.925, \sigma\rightarrow 0.008].$

Important Expressive Extensions of LP (II)

- Probabilistic LP, in two major flavors
 1. Bayesian flavor with “distribution semantics” (a.k.a. other names)
 - Bayesian Networks are a subset
 - General case is computationally intractable, even for function-free
 2. Fuzzy flavor. Parametrized by choice of “triangular norm” function F .
 - $\text{pr}(A \text{ \& } B) = F(\text{pr}(A), \text{pr}(B))$. E.g., $F = \min$. Co-norm for \or : e.g., \max .
 - Function-free case is polynomial time
- Numerical truth values for atoms (and rules) range on real interval $[0..1]$
- *head* formula can be: \or of disjoint atoms/literals whose weights add to 1
 - $\text{friendly}(?x) \sim 0.8 \text{ \or } \text{unfriendly}(?x) \sim 0.2 \text{ :- student}(?x)$.
- Answer Set Programs
 - Head permits \or . Classical-like reasoning-by-cases.
 - Even function-free case is computationally intractable

Overview of Rulelog (I)

- A major research advance in KRR theory & algorithms
 - Culminated in 2012.
Contributions from many researchers since 1990's.
Tightly combines multiple expressive extensions that were originally independent.
- Very high/flexible **expressiveness** – *cf. earlier slide*
 - Overall: extends LP with strong **meta** (knowledge and reasoning)
- Strongly supports **federation & orchestration** of KRR
 - Via external queries + high expressiveness

Overview of Rulelog (II)

- Computationally scalable, despite very high expressiveness
 - Database logic (LP) spirit + bounded rationality
 - Reasoning is polynomial time* (as in databases)
- Has capable efficient algorithms AND implementations
 - Dynamic compilation/transformation stack architecture
 - Cacheing of successful & failed inferences, with dependency-awareness, subgoal reordering, analysis of cycles and depths
 - Indexing, tries, other low-level data structures
 - Leverages database and “tabling” techniques
- Supports automatic full explanations
- Supersedes expressiveness and closely integrates with: RDF & SPARQL, relational DB & SQL, OWL-RL

* if, as is typical, one employs the radial restraint feature

Overview of Rulelog (III)

- Closely integrates with OWL-DL ontologies
- Closely integrates with natural language processing (NLP)
 - Text interpretation: map text to logic
 - Text generation: map logic to text
- Complements and integrates with machine learning (ML)
 - Import ML results as probabilistic knowledge
 - Export conclusions to ML

→ → *practical, easier to build and evolve KB's*

Rulelog and Textual Rulelog

- Rulelog reasoning scales well: polynomial-time, as in databases
 - Millions of sentences concluded/asserted on a single processor
 - Up to trillions by orchestrating database etc. systems in distributed settings
- *Textual* Rulelog extends Rulelog with natural language processing (NLP)
 - Logic itself is utilized to map between English syntax and logic syntax
 - Textual templates aid knowledge entry and explanation generation
- Examples:
 - \backslash (The individual affiliate threshold for transaction under Regulation W by ?Bank with ?Counterparty is ?Amount) :-
 \backslash (?Counterparty is deemed an affiliate of ?Bank under Regulation W) \backslash and
 \backslash (?Bank has capital stock and surplus ?Capital) \backslash and
 \backslash (the threshold percentage for an individual affiliate is ?Percentage) \backslash and
?Amount \backslash is ?Capital * ?Percentage/100.
 - @{'each large company has some talented CEO'}
 - forall(?x)^((?x \isa \large company) ==>
exists(?y)^((\?x has ?y) \and
(?y \isa \talented CEO))).

Uses of Rulelog – Overview

- Good for complex, commonly-arising kinds of knowledge, combined with simpler kinds of info
 - Mappings between different terminologies, ontologies, or schemas
 - Policies
 - Legal: regulations, contracts
 - Causal pathways, e.g., in science or biz processes
 - Financial. Medical. National intelligence. ...
- Use for decision automation, question-answering, and other analytics, esp. involving
 - Deep reasoning
 - Integration of diverse info sources and types

Example of Explanation in Rulelog

Why is the proposed transaction prohibited by Regulation W?

With drill down on: Why is the aggregate-affiliates limit \$10 million?

Why 'What proposed transactions are prohibited by RegW? Show ('Pacific Bank','Maui Sunset',23.0) ?

Edit Operations

- RegW prohibits the proposed transaction by Pacific Bank with Maui Sunset of \$23.0 million
 - The proposed transaction by Pacific Bank with Maui Sunset of \$23.0 million is a RegW covered transaction
 - There is a limit of \$10.0 million for any proposed RegW covered transaction by Pacific Bank with Maui Sunset
 - There is an aggregated-affiliates limit of \$10.0 million for any proposed RegW covered transaction by Pacific Bank with any affiliate
 - The aggregated total of previous RegW covered transactions by Pacific Bank with all affiliates is \$490.0 million
 - The maximum threshold for aggregate RegW covered transactions by Pacific Bank with all affiliates is \$500.0 million
 - The capital stock and surplus of Pacific Bank is \$2500.0 million
 - The RegW threshold percentage for aggregate affiliates is 20.0 percent
 - \$500.0 million is \$2500.0 million multiplied by 20.0 percent
 - The limit of \$10.0 million is the result of subtracting the previous RegW covered transactions total of \$490.0 million from the RegW threshold \$500.0 million
 - There is an individual-affiliate limit of \$250.0 million for any proposed RegW covered transaction by Pacific Bank with Maui Sunset
 - The overall RegW limit of \$10.0 million is the lesser of \$10.0 million and \$250.0 million
 - The proposed transaction of \$23.0 million is greater than the RegW limit of \$10.0 million

Outline of Tutorial

A. Overview

- Practical logic. Applications. Features. Software. Textual.
- Case Study Demo and Features Tour
 - Financial regulatory/policy compliance

B. Concepts and Foundations: Drill-downs

- Expressive features, semantics, algorithms;
relationships to natural language and machine learning

C. Conclusions and Future Work

- Background Assumed: basic knowledge of first-order logic, databases. Helpful: XML, RDF and semantic web concepts

Case Study: Automated Decision Support for Financial Regulatory/Policy Compliance

Problem: Current methods are expensive and unwieldy, often inaccurate

Solution Approach – using Textual Rulelog software technology:

- Encode regulations and related info as semantic rules and ontologies
- Fully, robustly automate run-time decisions and related querying
- Provide understandable full explanations in English
 - *Proof*: Electronic audit trail, with provenance
- Handles increasing complexity of real-world challenges
 - Data integration, system integration
 - Conflicting policies, special cases, exceptions
 - What-if scenarios to analyze impact of new regulations and policies

Business Benefits – compared to currently deployed methods:

- More Accurate
- More Transparent – better explanations
- More Cost Effective – less labor; subject matter experts in closer loop
- More Agile – faster to update
- More Overall Effectiveness: less exposure to risk of non-compliance



Demo of Rulelog for Compliance Automation: US Federal Reserve Regulation W

- EDM Council Financial Industry Consortium
Proof of Concept – **successful and touted pilot**
 - Enterprise Data Management Council (Trade Assoc.)
 - Coherent Knowledge (USA, Technology)
 - SRI International (USA, Technology)
 - Wells Fargo (Financial Services)
 - Governance, Risk and Compliance Technology Centre (Ireland, Technology)
- Reg W regulates and limits \$ amount of transactions that can occur between banks and their affiliates. Designed to limit risks to each bank and to financial system.
- Must answer 3 key aspects:
 1. *Is the transaction's counterparty an affiliate of the bank?*
 2. *Is the transaction contemplated a covered transaction?*
 3. *Is the amount of the transaction permitted ?*

Determining Whether Regulation W Applies

Two initial questions need to be answered in determining whether a transaction is subject to Regulation W. The first is whether the transaction is between a bank and an "affiliate" of the bank. The second is whether the transaction is a "covered transaction."

Affiliate Definition. Regulation W applies to covered transactions between a bank and an affiliate of the bank.

The definition of an affiliate for purposes of Regulation W is set forth in section 223.2. The definition is broad, and includes:

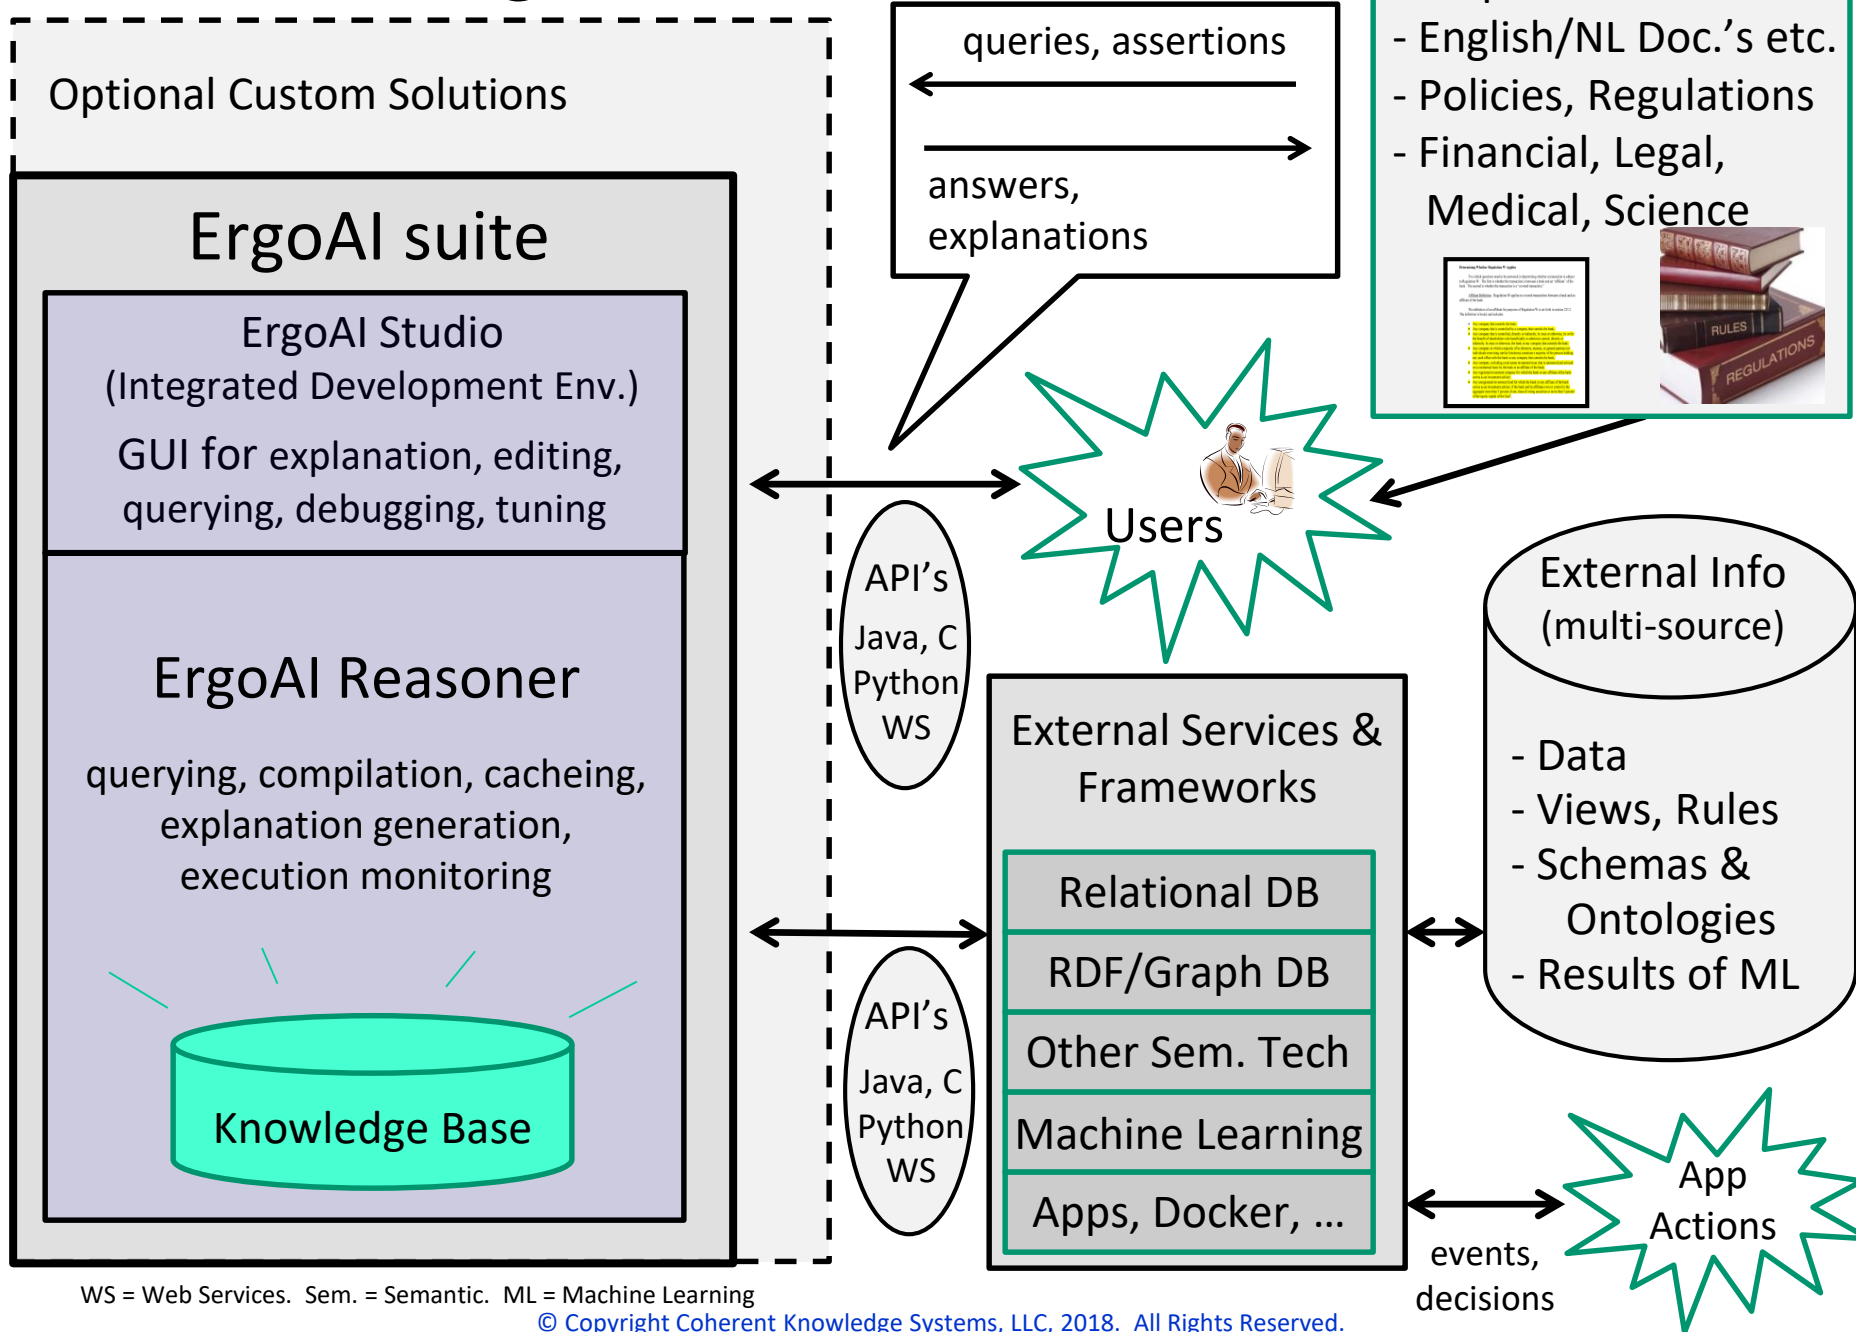
- Any company that controls the bank;
- Any company that is controlled by a company that controls the bank;
- Any company that is controlled, directly or indirectly, by trust or otherwise, by or for the benefit of shareholders who beneficially or otherwise control, directly or indirectly, by trust or otherwise, the bank or any company that controls the bank;
- Any company in which a majority of its directors, trustees, or general partners (or individuals exercising similar functions) constitute a majority of the persons holding any such office with the bank or any company that controls the bank;
- Any company, including a real estate investment trust, that is sponsored and advised on a contractual basis by the bank or an affiliate of the bank;
- Any registered investment company for which the bank or any affiliate of the bank serves as an investment adviser;
- Any unregistered investment fund for which the bank or any affiliate of the bank serves as an investment adviser, if the bank and its affiliates own or control in the aggregate more than 5 percent of any class of voting securities or more than 5 percent of the equity capital of the fund;

The Starting Point - Text of Regulation W

Reg W Demo Drill-down – outline

- Uses ErgoAI implementation of Rulelog
- Video available at <http://coherentknowledge.com>
- We show here: screenshots, example KB rules

ErgoAI Architecture


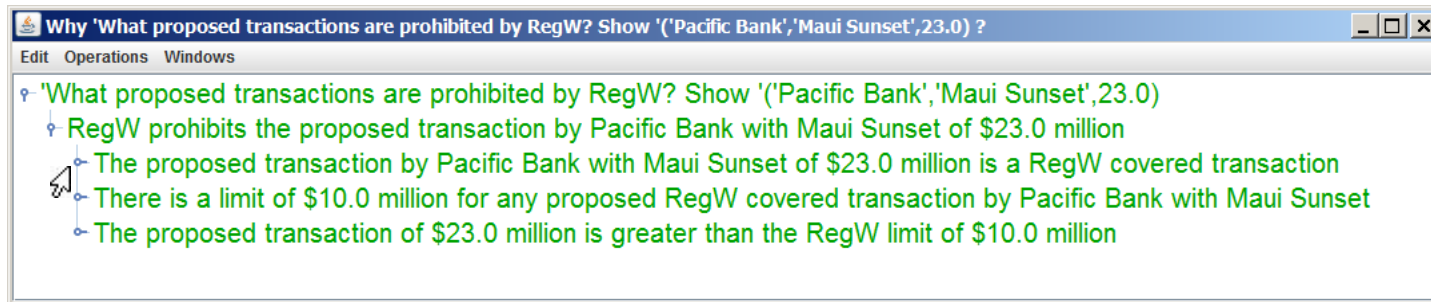



Query is asked in English

The screenshot shows the 'Coherent Query' application window. The title bar includes 'Coherent Query' and standard window controls. The menu bar contains 'Edit', 'Restraint', 'View', 'Explain', 'History', and 'Windows'. On the right side of the window, there are three buttons: 'Execute', 'Pause', and 'Stop'. The main text area contains the query: 'What proposed transactions are prohibited by RegW? Show '(?Bank,?Company,?Amount)'. Below the query is a table with three columns: '?Bank', '?Company', and '?Amount'. The first row of data contains the values: ''Pacific Bank'', ''Maui Sunset'', and '23.0'. A context menu is open over the '23.0' value, with the 'Why?' option highlighted by a red circle. The context menu also includes 'Explanation Game' and 'See answer term as tree'.

?Bank	?Company	?Amount
'Pacific Bank'	'Maui Sunset'	23.0

User Clicks the handles to expand the Explanations



Why is the proposed transaction prohibited by Regulation W?

1. *Is the transaction's counterparty an "affiliate" of the bank?*

YES.

Why 'What proposed transactions are prohibited by RegW? Show '(Pacific Bank','Maui Sunset',23.0) ?

Edit Operations

- RegW prohibits the proposed transaction by Pacific Bank with Maui Sunset of \$23.0 million
 - The proposed transaction by Pacific Bank with Maui Sunset of \$23.0 million is a RegW covered transaction
 - Maui Sunset is a RegW affiliate of Pacific Bank
 - Hawaii Bank is a RegW affiliate of Pacific Bank
 - There is common control of Hawaii Bank and Pacific Bank
 - Hawaii Bank is controlled by Americas Bank
 - Hawaii Bank is a subsidiary of Americas Bank
 - Pacific Bank is controlled by Americas Bank
 - Pacific Bank is a subsidiary of Americas Bank
 - Maui Sunset is advised by Hawaii Bank
 - There is a proposed loan from Pacific Bank to Maui Sunset of \$23.0 million
 - There is a limit of \$10.0 million for any proposed RegW covered transaction by Pacific Bank with Maui Sunset
 - The proposed transaction of \$23.0 million is greater than the RegW limit of \$10.0 million

And here's why ...

Executable Assertions: non-fact Rules

```
/* A company is controlled by another company when the first company  
   is a subsidiary of a subsidiary of the second company. */
```

```
@!{rule103b} /* declares rule id */
```

```
@@{defeasible} /* indicates the rule can have exceptions */
```

```
controlled(by)(?x1,?x2)
```

```
:- /* if */
```

```
  subsidiary(of)(?x1,?x3) \and
```

```
  subsidiary(of)(?x3,?x2).
```

```
/*A case of an affiliate is: Any company that is advised on a contractual basis by  
   the bank or an affiliate of the bank. */
```

```
@!{rule102b} @@{defeasible}
```

```
affiliate(of)(?x1,?x2) :-
```

```
  ( advised(by)(?x1,?x2)
```

```
    \or
```

```
    (affiliate(of)(?x3,?x2) \and advised(by)(?x1,?x3))).
```

Executable Assertions: **Exception Rule**

```
@!{rule104e}
@{'ready market exemption case for covered transaction'} /* tag for prioritizing */
\neg covered(transaction)(by(?x1))(with(?x2))
    (of(amount(?x3)))(having(id(?Id))) :-
    affiliate(of)(?x2,?x1) \and
    asset(purchase)(by(?x1))(of(asset(?x6)))(from(?x2))(of(amount(?x3)))
    (having(id(?Id))) \and
    asset(?x6)(has(ready(market))).

/* prioritization info, specified as one tag being higher than another */
\overrides('ready market exemption case for covered transaction',
    'general case of covered transaction').

/* If a company is listed on the New York Stock Exchange (NYSE), then the
    common stock of that company has a ready market. */
@!{rule201} @@{defeasible}
asset(common(stock)(of(?Company)))(has(ready(market))) :-
    exchange(listed(company))(?Company)(on('NYSE')).
```

Executable Assertions: Import of OWL

```
:- iriprefix fibof = /* declares an abbreviation */
    "http://www.omg.org/spec/FIBO/FIBO-Foundation/20120501/ontology/".

/* Imported OWL knowledge: from Financial Business Industry Ontology (FIBO) */
rdfs#subclassOf(fibob#BankingAffiliate, fibob#BodyCorporate).
rdfs#range(fibob#whollyOwnedAndControlledBy, fibob#FormalOrganization).
owl#disjointWith(edmc#Broad_Based_Index_Credit_Default_Swap_Contract,
    edmc#Narrow_Based_Index_Credit_Default_Swap_Contract).

/* Ontology Mappings between textual terminology and FIBO OWL vocabulary */
company(?co) :- fibob#BodyCorporate(?co).
fibob#whollyOwnedAndControlledBy(?sub,?parent) :- subsidiary(of)(?sub,?parent).

/* Semantics of OWL - specified as general Rulelog axioms */
?p(?y) :- rdfs#range(?p,?r), ?p(?x,?y).
?p(?x,?y) :- owl#subPropertyOf(?q,?p), ?q(?x,?y).
```

Outline of Tutorial

A. Overview

- Practical logic. Applications. Features. Software. Textual.
- Case Study Demo and Features Tour
 - Financial regulatory/policy compliance

B. Concepts and Foundations: Drill-downs

- Expressive features, semantics, algorithms;
relationships to natural language and machine learning

C. Conclusions and Future Work

- Background Assumed: basic knowledge of first-order logic, databases. Helpful: XML, RDF and semantic web concepts

Series of Advances → Rulelog's Core Expressive Features

- Well-founded semantics; basic tabling algorithms
 - *Undefined* for paradox; smart cacheing; intuitionistic disjunction
- Higher-order syntax (Hilog); frame syntax
 - Associated optimizations of LP tabling etc. algorithms
- Statement id's for meta; argumentation meta-rules for defeasibility; provenance
- General formulas with all usual classical connectives and quantifiers (omniformity)
- Restraint bounded rationality
 - Use 3rd truth value *undefined* for “don't-care”
 - Radial, skipping; naf unsafety; external-query unsafety, unreturn

Rulelog: Software Tools

- ❖ Lots of Rulelog expressiveness:
 - Ergo Lite: Large subset of Rulelog. Open source.
 - A.k.a. as Flora-2, originally
 - ErgoAI (from Coherent): Most of Rulelog. Has IDE.
 - **The most complete & highly optimized implementation available.**
Free for academic research use. Free trials available for everyone.
(Support time may cost, tho'.)
- ❖ Much smaller subsets of Rulelog expressiveness:
 - XSB Prolog: Most of LP -- with functions and well-founded negation. Plus a bit more. Open source.
 - Jena: Function-free negation-free LP, focused on RDF. Plus a bit more. Open source.
 - Similar: misc. other, e.g., that implement SWRL or SPIN

ErgoAI: Reasoner, Studio, Connectors

- ErgoAI Reasoner has sophisticated algorithms & data structures
 - Smart cacheing with dependency-aware updating. Leverages LP & DBMS techniques.
 - Transformation, compilation, subgoal reordering, indexing, modularization, dependency/loop analysis, performance monitoring/analysis, pausing, virtual machine, programming kernel, external import/querying
 - Java API. Other interfaces: command line, web, C. Additional APIs for Python, REST, more.
 - Scales well: Millions of sentences on 1 processor; Trillions on distributed nodes
- Explanations: for every answer; interactively drill down tree; in NL
 - Rulelog enables natural deduction style proofs, automated NL generation via rules
- ErgoAI Studio is a graphical Integrated Development Environment
 - Interactive editing, querying, visualization of knowledge
 - Fast edit-test loop with award-winning advanced knowledge debugging/monitoring
- ErgoAI Connectors federate knowledge & reasoning
 - Import/query dynamically via: SPARQL, OWL, RDF; SQL; CSV; JSON; more
 - Federation distributes reasoning (i.e., its processing) across multiple nodes
- Open, standards-based approach; a portion is open source
 - Rulelog is draft industry standard from RuleML (submission to W3C & Oasis)

KRR Features Comparison: Rulelog Shines

<u>System</u> <u>Feature</u>	Rulelog Rules - e.g., ErgoAI	Datalog Rules - e.g., Jena, SWRL, Ontobroker, SPIN	Production Rules - e.g., IBM, Oracle, Red Hat	Prolog - e.g., SWI, SICStus, XSB	FOL & OWL-DL - e.g., Vampire, Pellet, Prover9	ASP Solvers - e.g., DLV, CLASP
Semantic & on standardization path	✓	✓	restricted case	restricted case	✓	✓
<i>Basic expressiveness</i>						
• Datalog LP	✓	✓	✓	✓	✓	✓
• Logical functions	✓	✗	✗	✓	✓	restricted
• Quantified formulas (genl.)	✓	✗	✗	✗	✓	✗
<i>Full Meta expressiveness</i>						
• Higher-order syntax, provenance	✓	✗	✗	✗ (except XSB a little)	✗	✗
• Defeasibility & well founded negation	✓	✗	✗	✗	✗	some have restricted
• Restraint bounded rationality	✓	✗	✗	✗	✗	✗
• Probabilistic	✓	✗	✗	✗	✗	✗
<i>Efficiency</i>						
• Goal-directed	✓	✗ (except Jena)	✗	✓	✓	✓
• Full LP tabling with dependency-aware updating	✓	✗	✗	✗ (except XSB)	✗	✗
• Polynomial time complexity	✓	✓	✓	✗	✗	✗ ⁴¹ ₄₁

Notes on KRR Features Comparison

- “System” means system type / approach of logical knowledge representation and reasoning (KRR).
 - “Semantic” means in the sense of KRR, i.e., fully declarative and having a model theory in the logical sense.
 - “FOL” means First Order Logic. “ASP” means Answer Set Programs.
 - ASP is recently emerging. The tasks for which it’s suitable are more similar to FOL than to the other systems here.
 - “Standardization” here means industry standardization. “On path to” means in process of being, or already, standardized.
 - “Restricted case” means for a syntactic/expressive subset.
 - Event-condition-action rules in this context are similar to, and lumped in with, production rules.
 - “LP” means declarative logic programs.
 - Datalog means LP without logical functions. Usually this is restricted to Horn. But here we permit negation(-as-failure).
 - OWL-RL is pretty much a restricted case of Datalog LP.
 - “Higher-order syntax” means Hilog, which enables probabilistic – and also 1) fuzzy and 2) frame syntax cf. F-Logic.
 - “Provenance” means provenance info about assertions, via properties of rule id’s that are within the logical language / KRR.
 - “Full” applies to all four of the meta expressiveness features.
 - Defeasibility includes flexible argumentation theories.
 - “General formulas” means classical-logic-like formulas, including with head existentials and with head disjunction.
 - “LP tabling” includes sophisticated: cacheing of intermediate reasoning results, inference control, and indexing.
 - “Dependency-aware updating” means that when assertions are added or deleted, saved inferences are only recomputed if they depend on the changes to the assertions.
 - Polynomial time “complexity” means worst-case computational complexity, with constant-bounded number of variables per rule. Polynomial-time is similar to database querying, and is a.k.a. “tractable”.
-
- Datalog X defeasibility: Ontobroker has full well founded negation.
 - Prolog X defeasibility: XSB has full well founded negation.
 - ASP X defeasibility: Some ASP systems have restricted defeasibility & well founded negation. ASP systems essentially have wf negation inside (i.e., as part of) their semantics/reasoning, and some ASP systems even expose it to the user.
 - Datalog X goal-directed: Jena has a backward engine as well as a forward engine.
 - ASP X general formulas: ASP has head disjunction.
 - FOL X full LP tabling with dependency-aware updating: Some FOL theorem-provers cache intermediate results in a way that is analogous to LP tabling, and some do dependency tracking but we’re not sure how analogous or sophisticated.
 - Prolog X higher-order syntax: XSB has some support for this (i.e., for Hilog), although it is not integrated well.

Concept: Virtual Data Stores

- Rulelog orchestrates overall federated reasoning by sub-goaling dynamically
- A variety of other (“external”) structured information systems are treated as virtual data/knowledge stores, via Rulelog federation connectors, which import/query and translate
- Each virtual data/knowledge store is treated as having expressiveness that is a subset of Rulelog. E.g.,
 - An external database fact is treated as a logical atom in Rulelog
 - An OWL axiom is treated as a fact but also supplemented by semantic axioms about OWL’s constructs in general

Kinds of Virtual Data in ErgoAI

- Graph databases: via SPARQL/RDF connector
 - Description logic ontologies: via OWL connector
- Relational databases: via SQL connector
- Spreadsheets and web logs: via DSV connector
- JSON connector; XML connector; Web services via those
- Extensible to almost any kind of (semi-)structured info
 - E.g., Machine Learning (ML) and NLP systems
 - Represent `prob(content_sentence, lower_bound, upper_bound, confidence_level, statistical_procedure)` as an ErgoAI sentence
 - E.g., legacy applications in Java
 - `Get_foo` method is treated like a query



Importing RDF & OWL knowledge into ErgoAI


- Screenshot of ErgoAI OWL connector part of ErgoAI Studio



Translates
RDF & OWL
to ErgoAI



Define IRIs in
ErgoAI Studio



N-triples and
N-quads



RDF/OWL XML,
JSON-LD, or
Turtle as input.
Predicate or
Frame syntax
output.

Ergo RDF/OWL

Help

Import RDF & OWL

Status: Done translating WorldBank.ttl

Select input:

- Import RDF/OWL N-triples or N-quads file (.nq, .nt)
- Import RDF/OWL N-triples or N-quads directory
- Import RDF/OWL XML file (.rdf, .owl, .xml)
- Import RDF/OWL XML directory
- Import JSON-LD file (jsonld)
- Import JSON-LD directory
- Import RDF/OWL Turtle file (.ttl)
- Import RDF/OWL Turtle directory

Input file: WorldBank.ttl

Output predicate arity (n-quads or n-triples):

- n-triples
- n-quads

Output quad's graph name ('main' is the default)

Output format (fastload .P or .ergo):

- fastload format
- predicate syntax: p(s,o) or p(s,o,g)
- frame syntax: s[p->o]

Manage IRIs:

```

xsd = http://www.w3.org/2001/XMLSchema#
rdf = http://www.w3.org/1999/02/22-rdf-syntax-ns#
rdfs = http://www.w3.org/2000/01/rdf-schema#
owl = http://www.w3.org/2002/07/owl#

```

Import RDF/OWL

Original RDF/OWL file: WorldBank.ttl

```

@prefix void: <http://rdfs.org/ns/void#> .
@prefix rdf: <http://www.w3.org/1999/02/22-rdf-syntax-ns#> .
@prefix rdfs: <http://www.w3.org/2000/01/rdf-schema#> .
@prefix owl: <http://www.w3.org/2002/07/owl#> .
@prefix xsd: <http://www.w3.org/2001/XMLSchema#> .
@prefix dcterms: <http://purl.org/dc/terms/> .
@prefix foaf: <http://xmlns.com/foaf/0.1/> .
@prefix qb: <http://purl.org/linked-data/cube#> .
@prefix skos: <http://www.w3.org/2004/02/skos/core#> .
@prefix sds: <http://www.w3.org/ns/sparql-service-description#> .
@prefix : <http://worldbank.270a.info/void.ttl#> .
@prefix worldbank-graph: <http://worldbank.270a.info/graph/> .
@prefix oecd-dataset: <http://oecd.270a.info/dataset/> .
@prefix bfs-dataset: <http://bfs.270a.info/dataset/> .
@prefix fao-dataset: <http://fao.270a.info/dataset/> .
@prefix ecb-dataset: <http://ecb.270a.info/dataset/> .
@prefix imf-dataset: <http://imf.270a.info/dataset/> .
@prefix uis-dataset: <http://uis.270a.info/dataset/> .
@prefix frb-dataset: <http://frb.270a.info/dataset/> .
@prefix worldbank-dataset: <http://worldbank.270a.info/dataset/> .
@prefix transparency-dataset: <http://transparency.270a.info/dataset/>

```

Ergo file: WorldBank.ttl.ergo

```

#deffast xsd http://www.w3.org/2001/XMLSchema#
#deffast rdf http://www.w3.org/1999/02/22-rdf-syntax-ns#
#deffast rdfs http://www.w3.org/2000/01/rdf-schema#
#deffast owl http://www.w3.org/2002/07/owl#

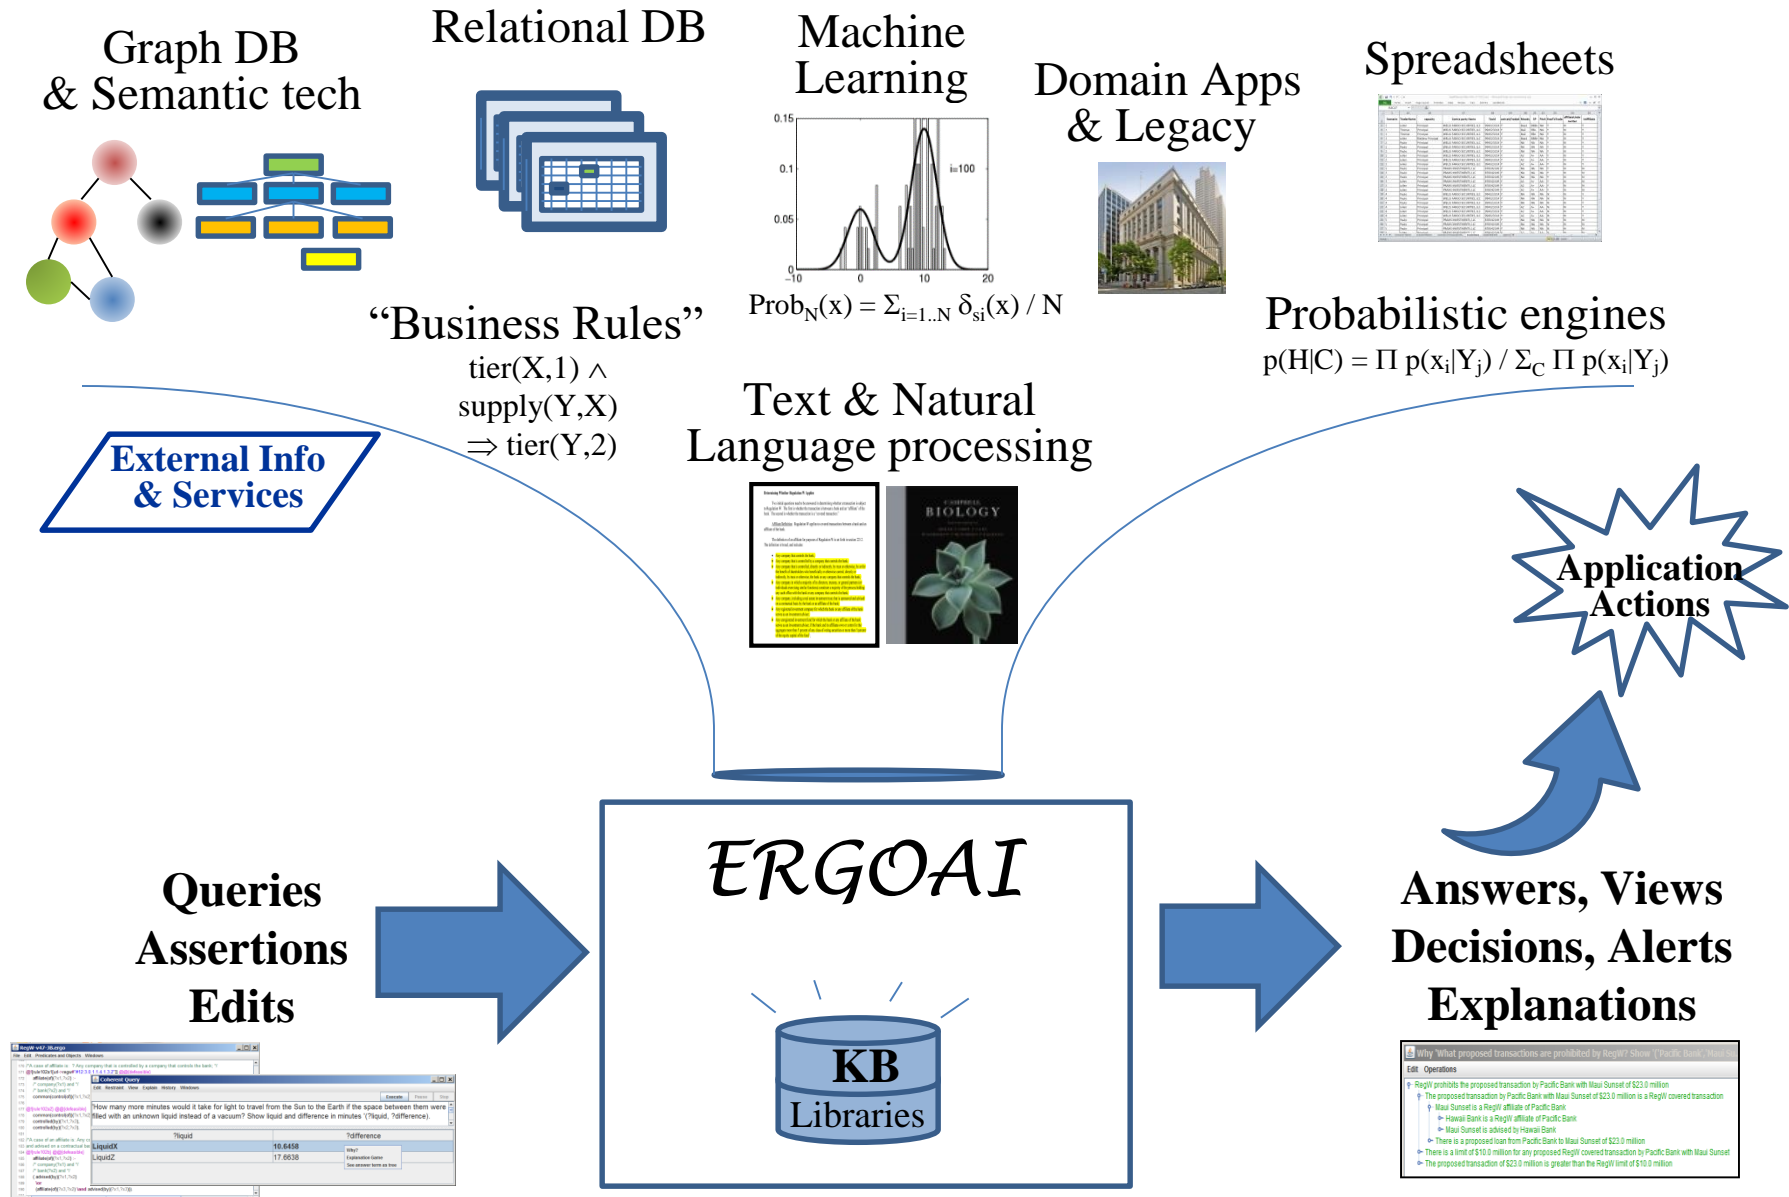
% imported OWL axioms
'http://rdfs.org/ns/void#entities'('_:Bb38eba1f27de68147b4ed800deeca630')
'http://rdfs.org/ns/void#class'('_:Bb38eba1f27de68147b4ed800deeca630')
'http://rdfs.org/ns/void#triples'('_:Bd43452bbb1eb87dc80d56d1c001f106')
'http://rdfs.org/ns/void#property'('_:Bd43452bbb1eb87dc80d56d1c001f106')
'http://rdfs.org/ns/void#distinctSubjects'('_:Bd43452bbb1eb87dc80d56d1c001f106')
'http://rdfs.org/ns/void#distinctObjects'('_:Bd43452bbb1eb87dc80d56d1c001f106')
'http://rdfs.org/ns/void#triples'('_:Bf7753516cd20cb7f77df061010915387')
'http://rdfs.org/ns/void#property'('_:Bf7753516cd20cb7f77df061010915387')
'http://rdfs.org/ns/void#distinctSubjects'('_:Bf7753516cd20cb7f77df061010915387')
'http://rdfs.org/ns/void#distinctObjects'('_:Bf7753516cd20cb7f77df061010915387')
'http://rdfs.org/ns/void#triples'('_:Bcf6bcafdc90833f622e5bb10c95d4d14')
'http://rdfs.org/ns/void#property'('_:Bcf6bcafdc90833f622e5bb10c95d4d14')
'http://rdfs.org/ns/void#distinctSubjects'('_:Bcf6bcafdc90833f622e5bb10c95d4d14')
'http://rdfs.org/ns/void#distinctObjects'('_:Bcf6bcafdc90833f622e5bb10c95d4d14')
'http://rdfs.org/ns/void#triples'('_:B3eef943acdd45aecbebacdd21158b10f')
'http://rdfs.org/ns/void#property'('_:B3eef943acdd45aecbebacdd21158b10f')
'http://rdfs.org/ns/void#distinctSubjects'('_:B3eef943acdd45aecbebacdd21158b10f')
'http://rdfs.org/ns/void#distinctObjects'('_:B3eef943acdd45aecbebacdd21158b10f')
'http://rdfs.org/ns/void#triples'('_:Be8f34857a86f0bce3671e0fb6acb0f7d')
'http://rdfs.org/ns/void#property'('_:Be8f34857a86f0bce3671e0fb6acb0f7d')
'http://rdfs.org/ns/void#distinctSubjects'('_:Be8f34857a86f0bce3671e0fb6acb0f7d')
'http://rdfs.org/ns/void#distinctObjects'('_:Be8f34857a86f0bce3671e0fb6acb0f7d')
'http://rdfs.org/ns/void#triples'('_:Bd81143ffc178de642750be48bdfa8ad3')
'http://rdfs.org/ns/void#property'('_:Bd81143ffc178de642750be48bdfa8ad3')
'http://rdfs.org/ns/void#distinctSubjects'('_:Bd81143ffc178de642750be48bdfa8ad3')
'http://rdfs.org/ns/void#distinctObjects'('_:Bd81143ffc178de642750be48bdfa8ad3')
'http://rdfs.org/ns/void#triples'('_:B71e4380c4b52f4b64169b767fbcaf48')

```

© Copyright 2015, Coherent Knowledge Systems, Ergo/OWL translator version 0.7.19 (July 19, 2015)

Coherent Knowledge Systems, 2015

Actively Reason over Today's Gamut of Knowledge



Outline of Tutorial

A. Overview

- Practical logic. Applications. Features. Software. Textual.
- Case Study Demo and Features Tour
 - Financial regulatory/policy compliance

B. Concepts and Foundations: Drill-downs

- Expressive features, semantics, algorithms;
relationships to natural language and machine learning

C. Conclusions and Future Work

- Background Assumed: basic knowledge of first-order logic, databases. Helpful: XML, RDF and semantic web concepts

Example Application Areas for Rulelog

- Confidentiality policies
- Financial/business reporting
- Contracts
- E-commerce pricing/promotion policies
- E-commerce product catalog integration, supply chain
- Financial regulatory/policy compliance
- Health treatment guidance, insurance
- Defense intelligence analysis
- Education/e-learning: personalized tutoring in sciences
- Info/system integration, e.g., in financial, defense

- Potentially many more: natural language interaction, business intelligence, games, workflow, social media,...



Problem: Analytics for *Complex Knowledge*

Examples: policies, regulations, contracts; terminology mappings; science, causality

Existing ***Non-Semantic*** Technologies tend to be:

- Shallow
- Siloed
- Costly, and Slow
- Patchily automated
- Opaque
- Inaccurate
- End users not empowered to modify

Based on:

- *Conventional programming languages*
- *Production/ECA rules*
- *Prolog*

Benefits of Semantic Approach to analytics & decision automation

- Modeling, declaratively, rather than programming
 - First steps – state of art:
 - Decision Tables (cf. DMN)
 - Ontologies (cf. OWL)
 - Benefits:
 - Greater integration and reusability
 - More transparent, i.e., explainable
 - Easier to modify, end users* more empowered
 - More cost-effective and agile
- * esp. subject matter experts (SMEs)

Rulelog is a Next Step on Semantic

- Compared to decision tables:
 - Deeper in reasoning & knowledge
 - Support many-step inferencing
 - Model complex sentences with high fidelity, via high expressiveness, e.g., higher-order, existentials
 - Map to/from natural language
 - Map between ontologies, schemas, terminologies
 - Principled defeasibility (exceptions)
 - Fuller, more understandable explanations
 - Greater scope of automation
 - ➔ Extends the benefits of the semantic approach

Outline of Tutorial

A. Overview

- Practical logic. Applications. Features. Software. Textual.
- Case Study Demo and Features Tour
 - Financial regulatory/policy compliance

B. Concepts and Foundations: Drill-downs

- Expressive features, semantics, algorithms;
relationships to natural language and machine learning

C. Conclusions and Future Work

- Background Assumed: basic knowledge of first-order logic, databases. Helpful: XML, RDF and semantic web concepts

Concept of Humagic Knowledge

- Humagic = human-machine logic
- A humagic KB consists of a set of linked sentences
 - Assertions, queries, conclusions (answers & explanations)
- NL-syntax sentence may have 1 or more logic-syntax sentences associated with it
 - E.g., that encode it, or give its provenance, or represent its text interpretation
- Logic-syntax sentence may have 1 or more NL-syntax sentences associated with it
 - E.g., results of text generation on it
 - E.g., source sentence in text interpretation, that produced it
- Other sentences can be in a mix of NL-syntax and logic-syntax
 - Using textual templates, for text interpretation and text generation

Textual extension of Rulelog

- High-level concept of approach: *Textual* Rulelog (TR)
- Extends Rulelog with natural language processing (NLP)
 - Start with English as the NL
- Rulelog logic itself is utilized to map:
Rulelog logic syntax $\leftarrow \rightarrow$ NL syntax
 - I.e., use logic to help do: *text interpretation* and *text generation*
- Mapping is much simpler and closer than with other KR's
 - Rulelog's high expressiveness is much closer to NL's conceptual abstraction level
 - More often doable and useful:
1 English sentence \leftrightarrow 1 Rulelog sentence (rule)
- In principle, almost any NL sentence can be represented with deep semantics as a logical sentence in Rulelog
 - Leverage the general quantified formulas expressive feature of Rulelog

Some Key Tasks in Textual Rulelog

- TR text interpretation:
Rulelog rules map from NL to logic
- TR text generation:
Rulelog rules map from logic to NL
- TR terminology mapping:
Rulelog rules map between phrasings and ontologies – in NL or logic
 - “moving a bomb” implies “transporting weaponized material”
 - `isBomb(?x)` implies `rdftriple(?x,rdftype,bomb)`

Some Techniques for Textual Rulelog

1. Word as functor (WAF): treat a NL word as a logical functor
 2. Phrasal terms (phrasts): treat a NL phrase as a logical term
 3. Phrasal mapping from paraphrase knowledge (PMK): e.g., synonyms, hypernyms, hyponyms, equivalent named entities; other implications b/ phrases
 4. Textual templates (TET): hybrid of text syntax and logic syntax
 5. Quantification of NL determiners (QUD): e.g., treat “every” and “some” as relativized universal and existential logical quantifiers
 6. Deep extended NL parsing (DEP): logically represent dependency parse tree extended with coreference analysis and named entity recognition
- *These can be combined*
 - *Many interesting directions & open areas for research! E.g., DEP and QUD.*

Motivation for Textual Templates (I)

The knowledge acquisition problem:


Domain expert



Cannot represent knowledge formally, can't code



Knowledge engineer




Doesn't have domain knowledge

Textual Templates Example in *ErgoAI* (I)

`\(?X is a father of ?Y\) : -
 \(?X is a parent of ?Y and a male\).`

`template(\(?X is a father of ?Y\), father(?X,?Y)).
 template(\(?X is a parent of ?Y and a male\),
 (parent(?X,?Y), ?X:Male)).`



```

[?X:Male, ?Y:Male, ?Z:Male,
mother(?X,?Y)].

/* The general context independent rules */
syskon(X,Y) :- Father(U,X), Father(U,Y), mother(V,X), mother(V,Y), X!=Y, U
parent(X,Y) :- Father(X,Y); mother(X,Y).

cousin(X,Y):- parent(U,X), parent(V,Y), syskon(U,V), X!=Y.
FatherFather(X,Y):- Father(X,U), Father (U,Y).
Fathermother(X,Y) :- mother(X,U), Father (U,Y).
motherFather(X,Y) :- Father(X,U), mother(U,Y).
mothermother(X,Y) :- mother(X,U), mother(U,Y).

parents(X,Y,Z) :- Father(X,Z), mother(Y,Z).

ancestor(X,Y) :- parent(X,Z), parent(Z,Y).
ancestor(X,Y) :- parent(X,Z), ancestor(Z,Y).
  
```

`\(if ?X is a penguin then it is a bird\).`

`template(\(if ?X is a penguin then it is a bird\),
 (bird(?X) : - penguin(?X))).`

- The *domain expert* operates with a corpus of stylized English sentences, which s/he
 - designs and communicates to the knowledge engineer
 - plus some elementary logic (and, or, if, variables)
- The *knowledge engineer* translates into *ErgoAI* using the *ErgoText* templates feature.
- Any inferences can be explained using the phrases of the aforesaid English corpus.

Textual Templates Example in ErgoAI (II)

- KB.ergo – the main KB

```
:- ergotext{KBtempl}.
```

```
\(?X is a father of ?Y\) :-
```

```
    \(?X is a parent of ?Y and a male\).
```

```
\(if ?X is a penguin then it is a bird\).
```

- KBtempl.ergotxt – templates

```
template(headbody,
```

```
    \(?X is a father of ?Y\),
```

```
    father(?X,?Y) ).
```

```
template(headbody,
```

```
    \(?X is a parent of ?Y and a male\),
```

```
    (parent(?X,?Y), ?X:Male) ).
```

```
template(rule,
```

```
    \(if ?X is a penguin then it is a bird\),
```

```
    (bird(?X) :- penguin(?X)) ).
```

Reg W Example Sentence using Templates

- Example of hybrid-syntax sentence – executable in ErgoAI:

\(The individual affiliate threshold for transaction under Regulation W
by ?Bank with ?Counterparty is ?Amount\) :-

\(?Counterparty is deemed an affiliate of ?Bank under Regulation W\) \and
\(?Bank has capital stock and surplus ?Capital\) \and
\(the threshold percentage for an individual affiliate is ?Percentage\) \and
?Amount \is ?Capital * ?Percentage/100.

Reg W example Template definition in ErgoAI

```
template(headbody,  
  \((The proposed transaction ?Id by ?Bank with ?Affiliate of $?Amount  
    is a RegW covered transaction\),  
  
  covered(proposed(transaction))(by(?Bank))(with(?Affiliate))  
    (of(amount(?Amount)))(having(id(?Id)))  
  ).
```

- The templates are self-documenting

Quantification of Determiners

- Determiners in NL often indicate quantification
 - E.g., treat “each” via a logical **relativized universal** quantifier
 - `forall(?var)^((?var \isa fooNounPhrase)) ==> barVerbPhraseAbout?Var)`
 - E.g., treat “some” via a logical **relativized existential** quantifier
 - `exists(?var)^((?var \isa fooNounPhrase)) \and barVerbPhraseAbout?Var)`
- Ex.: “each large company has some talented CEO”
`exists(?y)^((?y \isa \ (talented CEO\)) \and
(\?x has ?y\))
:- ?x \isa \ (large company\)) .`

`/* above has implicit outermost universal quantification of ?x */`

Deep Extended NL Parsing (DEP) in TR

- Common useful form of NLP output is “**parse++**” in following form:
 - NL dependency parsing,
 - extended with coreference analysis (Coref) and named entity recognition (NER)
- DEP: Represent the results of parse++ as a set of logical facts specifying the parse++ tree in full detail
 - Tree structure: dependency edges, left vs. right sequencing, edge labels
 - Node labels: word token, part of speech (PoS), NER, Coref, other word sense
 - Provides grist for deep semantic text interpretation & representation in Rulelog
- **ErgoNLP** is a recently released open source tool implementing DEP
 - It uses the popular Stanford CoreNLP toolset, which is open source (GPL)
 - Inputs a passage of 1 or more English sentences
 - Outputs ErgoAI facts that represent the parse++ of each sentence
 - Actually uses only the Ergo Lite subset of ErgoAI’s syntax & expressiveness
 - <https://bitbucket.org/coherentknowledge/ergonlp>
 - Implemented in Java. Original authors Coherent Knowledge. Apache 2.0 license.

Example of DEP in ErgoNLP

// Input Sentence 1: Each large company has some talented CEO.

```
ph(104)[ ws(wd)->'has',  
  root -> \true,  
  ws(PoS)->'VBZ',  
  ldp(1)->dp(nsubj,ph(103)),  
  rdp(1)->dp(dobj,ph(107))].
```

```
ph(103)[ ws(wd)->'company',  
  ws(PoS)->'NN',  
  ldp(1)->dp(det,ph(101)),  
  ldp(2)->dp(amod,ph(102))].
```

```
ph(107)[ ws(wd)->'CEO',  
  ws(PoS)->'NN',  
  ldp(1)->dp(det,ph(105)),  
  ldp(2)->dp(amod,ph(106)),  
  rdp(1)->dp(rcmod,ph(109))].
```

```
ph(101)[ ws(wd)->'Each',  
  ws(PoS)->'DT'].
```

```
ph(102)[ ws(wd)->'large',  
  ws(PoS)->'JJ'].
```

```
ph(105)[ ws(wd)->'some',  
  ws(PoS)->'DT'].
```

```
ph(106)[ ws(wd)->'talented',  
  ws(PoS)->'JJ'].
```

// Input Sentence 2: IBM is a huge company.

```
ph(205)[ ws(wd)->'company',  
  root -> \true,  
  ws(PoS)->'NN',  
  ldp(1)->dp(nsubj,ph(201)),  
  ldp(2)->dp(cop,ph(202)),  
  ldp(3)->dp(det,ph(203)),  
  ldp(4)->dp(amod,ph(204)),  
  ws(coref)->ph(201)].
```

```
ph(201)[ ws(wd)->'IBM',  
  ws(PoS)->'NNP',  
  ws(ne)->ORGANIZATION].
```


```
ph(202)[ ws(wd)->'is',  
  ws(PoS)->'VBZ'].
```

```
ph(203)[ ws(wd)->'a',  
  ws(PoS)->'DT'].
```

```
ph(204)[ ws(wd)->'huge',  
  ws(PoS)->'JJ'].
```

The example's input text passage is the 2 English sentences.
There's one frame set of facts for each word.
There's one set of frame sets for each English sentence.
In Ergo frame syntax:
 subject[property->value]
is a fact triple that is similar to
 property(subject,value)
in predicate calculus syntax; and
 subj[prop1->val1, prop2->val2, prop3->val3]
is logically equivalent to the 3 fact triples
 subj[prop1->val1] \and subj[prop2->val2] \and subj[prop3->val3].
ph(104) stands for phrasal term number 104. Each such phrast represents a node in the parse++ tree, and also corresponds to a parse++ subtree rooted at that node.
ws stands for word sense info.
ws(wd) stands for the word token.
ws(PoS) stands for part of speech.
ldp stands for left dependency list.
rdp stands for right dependency list.
ldp(1) stands for first left dependency, as one goes left-to-right;
ldp(2) stands for the second, etc.
rdp(1) stands for first right dependency, as one goes left-to-right;
rdp(2) stands for the second, etc.
dp(*dependency_label*, ph(*number*)) stands for a dependency edge, with a particular edge label, to a particular node.
The dependency edge labels, and PoS node labels, are the usual Penn TreeBank ones. E.g., nsubj for subject NP, JJ for adjective.

Knowledge Authoring Steps using Textual Rulelog



R&D direction: methods to greatly increase the degree of automation in encoding

Knowledge Authoring Process using Textual Rulelog

- Start with source text in English – e.g., textbook or policy guide
 - A sentence/statement can be an assertion or a query
- Articulate: create encoding sentences (text) in English.
As necessary:
 - Clarify & simplify – be prosaic and grammatical, explicit and self-contained
 - State relevant background knowledge – that’s not stated directly in the source text
- Encode: create executable logic statements
 - Each encoding text sentence results in one executable logic statement (“rules”)
 - Use IDE tools and methodology
- Test and debug, iteratively
 - Execute reasoning to answer queries, get explanations, perform other actions
 - Find and enter missing knowledge
 - Find and fix incorrect knowledge
 - Optionally: further optimize reasoning performance, where critical

Outline of Tutorial

A. Overview

- Practical logic. Applications. Features. Software. Textual.
- Case Study Demo and Features Tour
 - Financial regulatory/policy compliance

B. Concepts and Foundations: Drill-downs

- Expressive features, semantics, algorithms;
relationships to natural language and machine learning

C. Conclusions and Future Work

- Background Assumed: basic knowledge of first-order logic, databases. Helpful: XML, RDF and semantic web concepts

Outline of Part B. Concepts & Foundations

Key features

- Spirit of LP
- **Higher-Order** Syntax via Hilog. Reification.
- Rule ID's
- Default Negation and Well Founded Semantics
- **Restraint**: semantic bounded rationality
- **Tabling Algorithms**. Knowledge Debugging Methods.
- **Defeasibility** via Argumentation Rules. Remediating FOL's Fragility.
- **General Formulas**, Existentials and Skolems, Omni-directional Disjunction
 - Representing Text. Importing full OWL.
- **Uncertain/Probabilistic** knowledge and reasoning
- *Other features covered in Parts (A.) and (C.):*
 - External Querying. **Federation & orchestration**.
 - **Textual Rulelog**: combining closely with Natural Language Processing.
 - **Explanation**. Terminology/ontology mapping.
- *Other features not covered much due to limitations of time & focus:*
 - Frame syntax (a.k.a. F-Logic), Object Oriented style
 - Reactiveness
 - Misc. Lesser Features: Datatypes, Aggregation, Integrity Constraints, Inheritance, Equality, "Constraints", "Modules"

Horn LP Compared to Horn FOL

- Fundamental Theorem connects Horn LP to Horn FOL:
 - $M(P) = \{ \text{all ground atoms entailed by } P \text{ in Horn } \underline{\text{FOL}} \}$
- Horn FOL has additional non-ground-atom conclusions, notably:
 - non-unit derived clauses; tautologies
- Can thus view Horn LP as the f-weakening of Horn FOL.
 - “f-” here stands for “fact-form conclusions only”
 - A restriction on form of conclusions (not of assertions).
- Horn LP – differences from Horn FOL:
 - Conclusions $\text{Conc}(P) =$ essentially a set of ground atoms.
 - Can extend to permit more complex-form queries/conclusions.
 - Consider Herbrand models only, *in typical formulation and usage*.
 - P can then be replaced equivalently by $\{ \text{all ground instantiations of each rule in } P \}$
 - But can extend to permit: extra unnamed individuals, beyond Herbrand universe
 - Rule has non-empty head, *in typical formulation and usage*.
 - Can extend to detect violation of integrity constraints

The “Spirit” of LP

The following summarizes the “spirit” of how LP differs from FOL:

- **“Avoid Disjunction”**
 - Avoid disjunctions of positive literals as expressions
 - In premises, intermediate conclusions, final conclusions
 - (conclude (A or B)) only if ((conclude A) or (conclude B))
 - Permitting such disjunctions creates exponential blowup
 - In propositional FOL: 3-SAT is NP-hard
 - In the leading proposed approaches that expressively add disjunction to LP with negation, e.g., propositional Answer Set Programs
 - No “reasoning by cases”, therefore
- **“Stay Grounded”**
 - Avoid (irreducibly) non-ground conclusions

LP, unlike FOL, is straightforwardly extensible, therefore, to:

- Nonmonotonicity – defaults, incl. NAF
- Procedural attachments, esp. external actions

The “Spirit” of LP

The following summarizes the “spirit” of how LP differs from FOL:

- **“Avoid Disjunction”**
 - Avoid disjunctions of positive literals as expressions
 - In premises, intermediate conclusions, final conclusions
 - (conclude (A or B)) only if ((conclude A) or (conclude B))
 - Permitting such disjunctions creates exponential blowup
 - In propositional FOL: 3-SAT is NP-hard
 - In the leading proposed approaches that expressively add disjunction to LP with negation, e.g., propositional Answer Set Programs
 - No “reasoning by cases”, therefore
- **“Stay Grounded”**
 - Avoid (irreducibly) non-ground conclusions

LP, unlike FOL, is straightforwardly extensible, therefore, to:

- Nonmonotonicity – defaults, incl. NAF
- Procedural attachments, esp. external actions

Outline of Part B. Concepts & Foundations

Key features

- Spirit of LP
- **Higher-Order** Syntax via Hilog. Reification.
- Rule ID's
- Default Negation and Well Founded Semantics
- **Restraint**: semantic bounded rationality
- **Tabling Algorithms**. Knowledge Debugging Methods.
- **Defeasibility** via Argumentation Rules. Remediating FOL's Fragility.
- **General Formulas**, Existentials and Skolems, Omni-directional Disjunction
 - Representing Text. Importing full OWL.
- **Uncertain/Probabilistic** knowledge and reasoning
- *Other features covered in Parts (A.) and (C.):*
 - External Querying. **Federation & orchestration**.
 - **Textual Rulelog**: combining closely with Natural Language Processing.
 - **Explanation**. Terminology/ontology mapping.
- *Other features not covered much due to limitations of time & focus:*
 - Frame syntax (a.k.a. F-Logic), Object Oriented style
 - Reactiveness
 - Misc. Lesser Features: Datatypes, Aggregation, Integrity Constraints, Inheritance, Equality, "Constraints", "Modules"

Hilog: Higher-Order Syntax

- A higher-order extension of predicate logic, which has a tractable first-order syntax
 - Allows certain forms of logically clean, yet tractable, meta-programming
 - Syntactically appears to be higher-order, but semantically is first-order and tractable
- Permit predicate or function to be a variable
- Permit predicate or function to be a complex functional term
- Elegant transformation defines the semantics, and is used to implement

HiLog Transformation

- High-level Spirit:
 $?pred(?arg1,?arg2) \rightarrow \rightarrow believe(?pred,?arg1,?arg2)$
- A simplified version of the transformation, which gives intuition:
 - Rewrite each atom $p(a,b) \rightarrow holds_2(p,a,b)$
 - Generic predicate constants $holds_1, holds_2, \dots$
 - Treat each term in similar manner
 - $f(a,b) \rightarrow apply_2(f,a,b)$
 - Generic function constants $apply_1, apply_2, \dots$

Examples of Hilog (I)

Hilog permits variables over predicates and function symbols:

$p(?X, ?Y) : - ?X(a, ?Z) \text{ and } ?Y(?Z(b)).$

*Higher-order variable
(a.k.a. meta-variable):
ranges over predicate
names of arity 2*

*Higher-order variable:
ranges over function
names of arity 1*

Hilog also permits variables over atomic formulas. This is a kind of reification:

$p(q(a)).$
 $r(?X) : - p(?X) \text{ and } ?X.$

HiLog's Diffusion

- Used in ISO Common Logic to syntactically extend FOL
 - Also appears promising for OWL Full and its use of RDF [Kifer; Hayes]
- Implemented in ErgoAI, Flora-2 / Ergo Lite, SILK
 - Also partially exists in XSB, others
- [Chen, Kifer, Warren, “HiLog: A Foundation for Higher-Order Logic Programming”, J. of Logic Programming, 1993]

Reification

- Reification makes a term out of a formula:

believes(john, $\{\text{likes}(\text{mary}, \text{bob})\}$)

- Introduced in [Yang & Kifer, ODBASE 2002]
- Rules can be reified, as well

Object made out of
the formula
mary[likes -> bob]

Rule ID's

- Simple, but important, feature
- Each (assertion) statement gets a unique rule id
- The id can be explicitly specified
 - `@!{myRule17} H :- B.`
- Or if implicit, is a skolem essentially
 - `H :- B. →` gets treated as: `@!{gensym0897} H :- B.`
- Enables various useful kinds of meta-knowledge, by asserting properties of the rule id
 - Provenance, e.g., `createdBy(myRule17, Benjamin)`
 - Defeasibility
 - Rule-based transformations, e.g., for language extensibility, UI, NLP

Representational Uses of Hidlog

- Hidlog (pronounced “High-Dee-Log”) = HiLog + reification + rule id’s
- For meta- reasoning, e.g., in knowledge exchange or introspection
 - Meta-data is central to the Web
- Modals, e.g., believe, permit. Multi-agent belief. Deontics.
- Defeasibility: principles of argumentation/debate, i.e., argumentation rules
- Restructuring in mapping of schemas, ontologies, and terminologies
 - E.g., in knowledge integration, federation, KB translation/import
- Conciseness. Simple example: transitive closure of a relation.
- Reasoning control
- Modularization of KB’s
- Context, including provenance
- KR macros
- Probabilistic/uncertainty range and confidence about a sentence
 - E.g., import of machine learning (ML) results
- Representing natural language (NL). E.g., compositionality of phrases:
 - Compounding of nouns; salesman(?x); (insurance(salesman))(?x); ((life(insurance))(salesman))(?x)
 - Adverbial modification: quickly(give)(?thing,?recipient)

Outline of Part B. Concepts & Foundations

Key features

- Spirit of LP
- **Higher-Order** Syntax via Hilog. Reification.
- Rule ID's
- Default Negation and Well Founded Semantics
- **Restraint**: semantic bounded rationality
- **Tabling Algorithms**. Knowledge Debugging Methods.
- **Defeasibility** via Argumentation Rules. Remediating FOL's Fragility.
- **General Formulas**, Existentials and Skolems, Omni-directional Disjunction
 - Representing Text. Importing full OWL.
- **Uncertain/Probabilistic** knowledge and reasoning
- *Other features covered in Parts (A.) and (C.):*
 - External Querying. **Federation & orchestration**.
 - **Textual Rulelog**: combining closely with Natural Language Processing.
 - **Explanation**. Terminology/ontology mapping.
- *Other features not covered much due to limitations of time & focus:*
 - Frame syntax (a.k.a. F-Logic), Object Oriented style
 - Reactiveness
 - Misc. Lesser Features: Datatypes, Aggregation, Integrity Constraints, Inheritance, Equality, "Constraints", "Modules"

Concept of Logical Monotonicity

- A KR S is said to be logically monotonic when in it:
$$P1 \subseteq P2 \quad \Rightarrow \quad \text{Conc}(P1,S) \subseteq \text{Conc}(P2,S)$$
- Where $P1, P2$ are each a set of premises in S
- I.e., whenever one adds to the set of premises, the set of conclusions non-strictly grows (one does not retract conclusions).
- *Monotonicity is good for pure mathematics.*
 - *“Proving a theorem means never having to say you are sorry.”*

Logical Nonmonotonicity – Motivations

- Pragmatic reasoning is, in general, nonmonotonic
 - E.g., policies for taking actions, exception handling, legal argumentation, Bayesian/statistical/inductive, etc.
 - Monotonic is a special case – simpler in some regards
- Most commercially important rule systems/applications use nonmon
 - A basic expressive construct is ubiquitous there:
 - Default Negation a.k.a. Negation-As-Failure (NAF)
 - BUT with varying semantics – often not fully declarative cf. LP
 - Primarily due to historical hangovers and lack of familiarity with modern algorithms
 - Another expressive construct, almost as ubiquitous there, is:
 - Priorities between rules
- Such nonmonotonicity enables:
 - Modularity and locality in revision/updating/merging

Default Negation: Intro

- *Default* negation, a.k.a. “*weak*” negation, a.k.a. “*negation as failure (NAF)*” is the most common form of negation in commercially important rule and knowledge-based systems.
- Normal LP, a.k.a. Ordinary LP, adds NAF to Horn LP
 - May appear in body literals, but not in head
- Concept/Intuition for $\backslash\text{naf } q$ **default** negation
 - q is not derivable from the available premise info
 - fail to believe q . “I know I do not believe q ”.
 - ... but might also not believe q to be false
 - A.k.a. “*weak*” negation, or NAF. In Ergo: “ $\backslash\text{naf}$ ”
- Contrast with: $\backslash\text{neg } q$ **strong** negation, a.k.a. “classical” negation
 - q is believed to be false ; the opposite of q is true
 - Brings potential for logical contradiction / conflict: p vs. $\backslash\text{neg } p$

Semantics for Default Negation

- Semantics of NAF has **subtleties** for the fully general case
 - **Paradox** can arise from “unstratified” NAF:
 - Cyclic dependencies among atoms/predicates: thru the rules, involving NAF
 - Example 1: $p :- \text{\textbackslash}naf p.$ Example 2: $p :- \text{\textbackslash}naf q. q :- \text{\textbackslash}naf p.$
 - Well-understood theoretically since 1994 – after 10 years of work
- Well Founded Semantics (WFS): popular; major commercial focus
 - Quadratic (often linear -- time for propositional. No reasoning by cases.
 - Employs a 3rd truth value **u** (“undefined”) for paradox
 - Intuition: “I cannot figure out whether I believe p or do not believe p”
 - Definition uses iterated minimality: Horn-case then close-off; repeat til done.
 - Operational semantics for NAF, in actual practice, is **often “sloppy”** (incomplete / cut-corners / obsolete) relative to canonical semantics for NAF, in most Prolog and other currently commercially important rule systems that express NAF; required WFS algorithms are more complex.

ASP Semantics for Default Negation

- *Answer Set Programs (ASP)*: popular as research topic
 - A different KR than WFS, departs particularly in presence of paradox
 - Intractable for propositional; does reasoning by cases
 - Enables a limited kind of disjunction in heads & conclusions
 - Good for combinatorial KRR problems requiring nonmonotonicity
 - Lacks undefined truth value \Rightarrow sometimes ill-defined: no set of conclusions

Outline of Part B. Concepts & Foundations

Key features

- Spirit of LP
- **Higher-Order** Syntax via Hilog. Reification.
- Rule ID's
- Default Negation and Well Founded Semantics
- **Restraint**: semantic bounded rationality
- **Tabling Algorithms**. Knowledge Debugging Methods.
- **Defeasibility** via Argumentation Rules. Remediating FOL's Fragility.
- **General Formulas**, Existentials and Skolems, Omni-directional Disjunction
 - Representing Text. Importing full OWL.
- **Uncertain/Probabilistic** knowledge and reasoning
- *Other features covered in Parts (A.) and (C.):*
 - External Querying. **Federation & orchestration**.
 - **Textual Rulelog**: combining closely with Natural Language Processing.
 - **Explanation**. Terminology/ontology mapping.
- *Other features not covered much due to limitations of time & focus:*
 - Frame syntax (a.k.a. F-Logic), Object Oriented style
 - Reactiveness
 - Misc. Lesser Features: Datatypes, Aggregation, Integrity Constraints, Inheritance, Equality, "Constraints", "Modules"

Requirements Analysis for Logical Functions

- Function-free is a commonly adopted restriction in practical LP/Web rules today
 - DB query languages: SQL, SPARQL, XQuery
 - Production rules, and similar Event-Condition-Action rules
 - OWL (including OWL-RL), RDF-Schema
 - RIF Basic Logic Dialect, RIF-Core (and SWRL)
 - Jena, SPIN
- BUT functions are often needed for Web (and other) applications.
Uses include:
 - Skolemization – to represent existential quantifiers (incl. RDF blank nodes)
 - E.g., $\text{exists}(?p)^{\wedge}(\text{part_of}(?p,?c) \ \&\ \text{nucleus}(?p)) \text{- eukaryotic_cell}(?c).$
→ → $\text{part_of}(\text{sk1}(?c),?c) \ \&\ \text{nucleus}(\text{sk1}(?c)) \text{- eukaryotic_cell}(?c).$
 - Related: convenient naming. E.g., `steering_wheel(my_car)`
 - HiLog and reification – higher-order syntax
 - For many purposes!

Functions in LP Lead to Undecidability

- Functions in LP lead to undecidability, due to potentially infinite number of conclusions
- Example:
 - Assert:
num(succ(?x)) :- num(?x).
num(0).
 - Conclusions:
num(0), num(succ(0)), num(succ(succ(0))), ...
- Practical effects: runaway computations during inferencing to answer a query
- A significant pain point for :
 - FOL and other classical-logic reasoning (e.g., higher-order)
 - LP
 - Prolog
- Heart of problem: the semantics, not the proof procedure or the implementation

Restraint – Bounded Rationality (I)

- Restraint: utilize the *undefined* (u) truth value
 - A fully semantic kind of bounded rationality
 - Intuition: don't-care, can't figure it out or don't want to figure it out
 - Intuition: misty regions of the set of all literals, in regard to truth value t vs. f
- Radial restraint:
 - Specify a fixed depth radius r on literals, e.g., $r = 10$
 - Every literal whose depth exceeds the radius is concluded to have truth value u
 - Depth here is a metric which could be nesting depth, term size, or similar
- Example revisited, when specify radial restraint
 - Then $\text{num}(\text{succ}(\text{succ}(\text{succ}(\text{succ}(0))))), \dots$ all have truth value u
- Radial restraint has a proof theory, incorporates straightforwardly into algorithms
 - Modest computational overhead, $\sim 10\%$ in experiments
- Paradox – unstratified NAF – is a kind of restraint

Restraint (II)

- Several other kinds of restraint
 - Skipping: don't-care about particular instances of particular rules
 - Allows for more nuance and conditions about when to apply restraint
 - E.g., restrain a causal projection rule, when the successor-state step count exceeds 1000 and the causal effect predicate is microscopic
 - Sensor unreturn
 - E.g., an external query does not return, due to network or server failure
 - NAF unsafety
 - NAF'd literal must be evaluated with 1 or more unbound variables
 - Sensor unsafety
 - Externally queried literal has too many unbound variables for what the external knowledge source is willing/able to answer as a query
- Modest computational overhead required for the above
- For more info on restraint, see
 - AAAI-13 paper “Radial Restraint: A Semantically Clean Approach to Bounded Rationality” by B. Grosz and T. Swift
 - RuleML-2013 paper “Advanced Knowledge Debugging for Rulelog” by C. Andersen et al.
 - Both are available at <http://coherentknowledge.com/publications/>

Outline of Part B. Concepts & Foundations

Key features

- Spirit of LP
- **Higher-Order** Syntax via Hilog. Reification.
- Rule ID's
- Default Negation and Well Founded Semantics
- **Restraint**: semantic bounded rationality
- **Tabling Algorithms**. Knowledge Debugging Methods.
- **Defeasibility** via Argumentation Rules. Remediating FOL's Fragility.
- **General Formulas**, Existentials and Skolems, Omni-directional Disjunction
 - Representing Text. Importing full OWL.
- **Uncertain/Probabilistic** knowledge and reasoning
- *Other features covered in Parts (A.) and (C.):*
 - External Querying. **Federation & orchestration**.
 - **Textual Rulelog**: combining closely with Natural Language Processing.
 - **Explanation**. Terminology/ontology mapping.
- *Other features not covered much due to limitations of time & focus:*
 - Frame syntax (a.k.a. F-Logic), Object Oriented style
 - Reactiveness
 - Misc. Lesser Features: Datatypes, Aggregation, Integrity Constraints, Inheritance, Equality, "Constraints", "Modules"

Tabling Algorithms for LP & Rulelog

- Builds and maintains a forest of saved subgoal attempts and results
- Thus heavily caches. Is mixed-direction, not just backward-direction.
- Challenges requiring algorithmic sophistication include:
 - Logical functions: complex terms and unifiers
 - Hilog (higher-order syntax): e.g., indexing
 - NAF, non-monotonicity, and defeasibility
- Techniques have been developed to meet all of these
 - Efficient search control, indexing and low level data structures
 - Incremental tabling adds more dependency-awareness to nonmon.
 - Enables fast updating via reuse of most previous inferences
 - E.g., for interactive rule authoring edit-test loop
- Highly sophisticated methods, optimized over last two decades

Tabling Algorithms for LP & Rulelog

- Builds and maintains a forest of saved subgoal attempts and results
- Thus heavily caches. Is mixed-direction, not just backward-direction.
- Challenges include:
 - Logical functions: complex terms and unifiers
 - Hilog (higher-order syntax): e.g., indexing
 - NAF, non-monotonicity, and defeasibility is a challenge
- Techniques have been developed to meet all of these
 - Efficient search control, indexing and low level data structures
 - Incremental tabling adds more dependency-awareness to nonmon.
 - Enables fast updating via reuse of most previous inferences
 - E.g., for interactive rule authoring edit-test loop
- Highly sophisticated, optimized over last two decades

Advanced Knowledge Base Debugging Methods in Rulelog

- Explanation – as we say in the Reg W case study in Ergo
 - Helps to identify missing knowledge, as well as wrong knowledge
- Pause/resume
- Performance monitor
- Analysis of attempted infinite loops ("terminyzer")
- Analysis of sizes of tables (subgoal attempts)
- Ergo implements these
 - Several were pioneered in Vulcan SILK

Outline of Part B. Concepts & Foundations

Key features

- Spirit of LP
- **Higher-Order** Syntax via Hilog. Reification.
- Rule ID's
- Default Negation and Well Founded Semantics
- **Restraint**: semantic bounded rationality
- **Tabling Algorithms**. Knowledge Debugging Methods.
- **Defeasibility** via Argumentation Rules. Remediating FOL's Fragility.
- **General Formulas**, Existentials and Skolems, Omni-directional Disjunction
 - Representing Text. Importing full OWL. FOL-Soundness.
- **Uncertain/Probabilistic** knowledge and reasoning
- *Other features covered in Parts (A.) and (C.):*
 - External Querying. **Federation & orchestration**.
 - **Textual Rulelog**: combining closely with Natural Language Processing.
 - **Explanation**. Terminology/ontology mapping.
- *Other features not covered much due to limitations of time & focus:*
 - Frame syntax (a.k.a. F-Logic), Object Oriented style
 - Reactiveness
 - Misc. Lesser Features: Datatypes, Aggregation, Integrity Constraints, Inheritance, Equality, "Constraints"

The following couple sections are TO-SKIM

- Defeasibility
- Omniform rules
- I.e., slides ~86-113

- (Ideally, these would have been compressed to be much shorter.)

Defeasible Knowledge & Reasoning

- **Concept of defeasibility: Rules can be true by default but may be *defeated*, i.e., have exceptions**
 - A form of commonsense reasoning
- **Rulelog has defeasibility. Most previous KRR languages do not.**
 - Its approach to defeasibility is the most flexible, efficient, and practical
 - “Argumentation rules”: the principles of defeat/debate are specified via a set of general meta-flavor rules
- **“A cell has a nucleus.” ... Except when it doesn’t 😊**
 - A cell has no nucleus during anaphase. Red blood cells have no nuclei.
 - A cell has two nuclei between mitosis and cytokinesis. Some fungi are multinucleate.

Knowledge often has **Exceptions**

- Exceptions / special cases are inevitably realized over time
 - E.g., knowledge is incomplete, multiple authors contribute, ...
- Requiring entered knowledge to be strictly / universally true (exception-free) is impractical
 - Precludes stating generalities (the typical), and limits the author pool
 - “The perfect is the enemy of the good”
- Exceptions manifest as contradictions, i.e., [conflict](#)
- Leveraging multiple sources of knowledge (e.g., KB merging) requires conflict resolution
 - Errors. Confusions. Omitted context.

Defeasible Reasoning

- **Rules can be true by default but may be defeated**
 - A form of commonsense reasoning
- **Application domains:**
 - policies, regulations, and law
 - actions, change, and process causality
 - Web services
 - inductive/scientific learning
 - natural language understanding
 - ...
- **Previous approaches (i.e., previous to Rulelog/Ergo):**
 - Courteous Logic Programs (Grosz, 1997)
 - The main approach used **commercially** (IBM Common Rules, 1999)
 - Defeasible logic (Nute, 1994) [*similar to Courteous LP*]
 - “Prioritized defaults” (Gelfond & Son, 1997)
 - Preferred answer sets (Brewka & Eiter, 2000)
 - Compiling preferences (Delgrande et al., 2003)
 - ...

Defeasibility is Always Prepared for Exceptions

- Recognizes and handles conflict
- Avoids unreliable conclusions from inconsistent knowledge
 - Unlike FOL
- Represents meta-knowledge which resolves conflicts
 - Minimize need to modify prior knowledge, by acquiring additional meta-knowledge
- ... Notably: priorities (partially-ordered) between rules
 - Some rules have higher priority than others
- Priorities arise naturally from: specificity, recency, authority, causality, reliability
- Prioritization tames conflict \Rightarrow aids modularity
 - Analogy: the gift of fire



Defeasibility is Indicated When...

- **Useful generalities – and potential exceptions – coexist**
 - Specify knowledge in detail/precision appropriate for various circumstances
- **Governing doctrine, definitions, or other knowledge, cannot be assured to be conflict-free, e.g.:**
 - Multiple sources of governing doctrine exist
 - Typically, no central authority resolves all conflict promptly
 - Truth depends on context
 - Yet context is rarely made fully explicit
- **Many broad realms are full of exceptions**
 - Policies, regulations, laws — and the workflows they drive
 - Multiple jurisdictions, organizations, contracts, origins
 - Learning and science. Updating. Debate.
 - May falsify previous hypotheses after observation or communication
 - Causal processes: changes to state, from interacting/multiple causes
 - Natural language (text interpretation): “there’s a gazillion special cases”

Ubiquity of Priorities

in Commercially Important Rules -- and Ontologies

- Updating in relational databases
 - more recent fact *overrides* less recent fact
- Static rule ordering in Prolog
 - rule earlier in file *overrides* rule later in file
- Dynamic rule ordering in production rule systems (OPS5)
 - “meta-”rules can specify agenda of rule-firing sequence
- Event-Condition-Action rule systems rule ordering
 - often static or dynamic, in manner above
- Exceptions in default inheritance in object-oriented/frame systems
 - subclass’s property value *overrides* superclass’s property value, e.g., method redefinitions
- **All lack Declarative KR Semantics**

Semantic KR Approaches to Prioritized LP

The currently most important for Semantic Web are:

1. Courteous LP

- KR extension to Ordinary LP
- In RuleML, since 2001
- Commercially implemented and applied
 - IBM's CommonRules, 1999-2009
 - Coherent Knowledge's Ergo, 2013*-present

2. Defeasible Logic

- Closely related to Courteous LP
 - Less general wrt typical patterns of prioritized conflict handling needed in e-business applications
 - In progress: theoretical unification with Courteous LP [Wan, Kifer, Grosz RR-2010 / Semantic Web Journal 2015]

* and earlier in Flora-2 (a.k.a. Ergo Lite) 2009-2013

Courteous LP: the What

- **Updating/merging of rule sets:** is crucial, often generates conflict.
- **Courteous LP's feature prioritized handling of conflicts.**
- **Specify scope of conflict via a set of exclusion constraints**
 - Each is a preventive spirit integrity constraint on a set of competing literals
 - It says that not all of the competing literals can be entailed as true.
 - $\text{opposes}(p, q) \approx (\perp :- p \text{ and } q)$ // Case of 2 competing literals
 - $\text{opposes}(\text{discount}(\text{?product}, "5\%"), \text{discount}(\text{?product}, "10\%"))$;
 - $\text{opposes}(\text{loyalCustomer}(\text{?cust}, \text{?store}), \text{premiereCustomer}(\text{?cust}, \text{?store}))$;
- **Permit strong negation of atoms:** (NB: a.k.a. (quasi-) “classical” negation.)
 - $\text{\neg} p$ means p has truth value *false* .
 - implicitly, for every atom p : $\text{opposes}(p, \text{\neg} p)$;
- **Priorities between rules:** partially-ordered.
 - Represent priorities via reserved predicate that compares rule tags:
 - $\text{\neg} \text{overrides}(\text{rule1}, \text{rule2})$ means rule1 is higher-priority than rule2.
 - Each rule optionally has a rule tag whose form is a functional term.
 - $\text{\neg} \text{overrides}$ can be reasoned about, just like any other predicate.

Courteous LP: Advantages

- Facilitate updating and merging, modularity and locality in specification.
- Expressive: strong negation, exclusions, partially-ordered prioritization, reasoning to infer prioritization.
- Guarantee consistent, unique set of conclusions.
 - **Exclusion is enforced.** E.g., never conclude discount is both 5% and that it is 10%, nor conclude both p and $\neg p$.
- Scalable & Efficient: low computational overhead beyond ordinary LPs.
 - Tractable given reasonable restrictions (VB, and radial restraint or no functions):
 - extra cost is equivalent to increasing v to $(v+2)$ in Ordinary LP, worst-case.
 - By contrast, more expressive prioritized rule representations (e.g., Prioritized Default Logic) add NP-hard overhead.
- Modular software engineering:
 - Transform: CLP $\rightarrow \rightarrow$ OLP. Via “argumentation theory” approach.

EECOMS Example of Conflicting Rules: Ordering Lead Time

- Vendor's rules that prescribe how buyer must place or modify an order:
 - A) 14 days ahead if the buyer is a qualified customer.
 - B) 30 days ahead if the ordered item is a minor part.
 - C) 2 days ahead if the ordered item's item-type is backlogged at the vendor, the order is a modification to reduce the quantity of the item, and the buyer is a qualified customer.
 - D) 45 days ahead if the buyer is a walk-in customer.
- Suppose more than one of the above applies to the current order? **Conflict!**
- Helpful Approach: **precedence** between the rules.
 - E.g., D is a catch-case: $A > D$, $B > D$, $C > D$
- Often only *partial* order of precedence is justified.
 - E.g., $C > A$, but no precedence wrt B vs. A, nor wrt C vs. B.

Ordering Lead Time Example in LP with Courteous Defaults

```
@ {prefCust} orderModifNotice(?Order,14days) :-  
    preferredCustomerOf(?Buyer,SupplierCo), purchaseOrder(?Order,?Buyer,SellerCo) .  
  
@ {smallStuff} orderModifNotice(?Order,30days) :-  
    minorPart(?Buyer,?Seller,?Order), purchaseOrder(?Order,?Buyer,SupplierCo) .  
  
@ {reduceTight} orderModifNotice(?Order,2days) :-  
    preferredCustomerOf(?Buyer,SupplierCo) and  
    orderModifType(?Order,reduce) and  
    orderItemIsInBacklog(?Order) and  
    purchaseOrder(?Order,?Buyer,SupplierCo) .  
  
\overrides(reduceTight, prefCust) . // reduceTight has higher priority than prefCust  
// The below exclusion constraint specifies that orderModifNotice is unique, for a given order.  
  
\opposes(orderModifNotice(?Order,?X), orderModifNotice(?Order,?Y)) :- ?X != ?Y .
```

- Rule D, and prioritization about it, were omitted above for sake of brevity.
- Above rules are represented in Logic Programs KR, using the Courteous defaults feature
- Notation:
 - “:-” means “if”. “@...” declares a rule tag. “?” prefixes a logical variable.
 - “\overrides” predicate specifies prioritization ordering.
 - An exclusion constraint specifies what constitutes a conflict.
 - “!=” means \neq .

Conclusions from opposition-locales previous to this opposition-locale $\{p1, p2\}$

(p1, p2 are each a ground strong literal, e.g., q, neg q)



Run Rules for p1, p2



Set of Candidates for p1, p2:
Team for p1, ..., Team for pk



Prioritized Refutation



Set of Unrefuted Candidates for p1, p2:
Team for p1, Team for p2



Skepticism



Conclude Winning Side if any: at most one of $\{p1, p2\}$

Argumentation Rules approach to Defaults in LP

- **Combines Courteous + HiLog, and generalizes**
- **New approach to defaults: “argumentation rules”**
 - Meta-rules, in the LP itself, that specify when rules ought to be defeated
 - [Wan, Grosz, Kifer, *et al.* ICLP-2009; RR-2010; SWJ 2015]
- **Extends straightforwardly to combine with other key features**
 - E.g., Frame syntax, external Actions
- **Significant other improvements on previous Courteous**
 - Eliminates a complex transformation
 - Much simpler to implement
 - 20-30 background rules instead of 1000’s of lines of code
 - Much faster when updating the premises
 - More flexible control of edge-case behaviors
 - Much simpler to analyze theoretically

LPDA Approach, Continued*

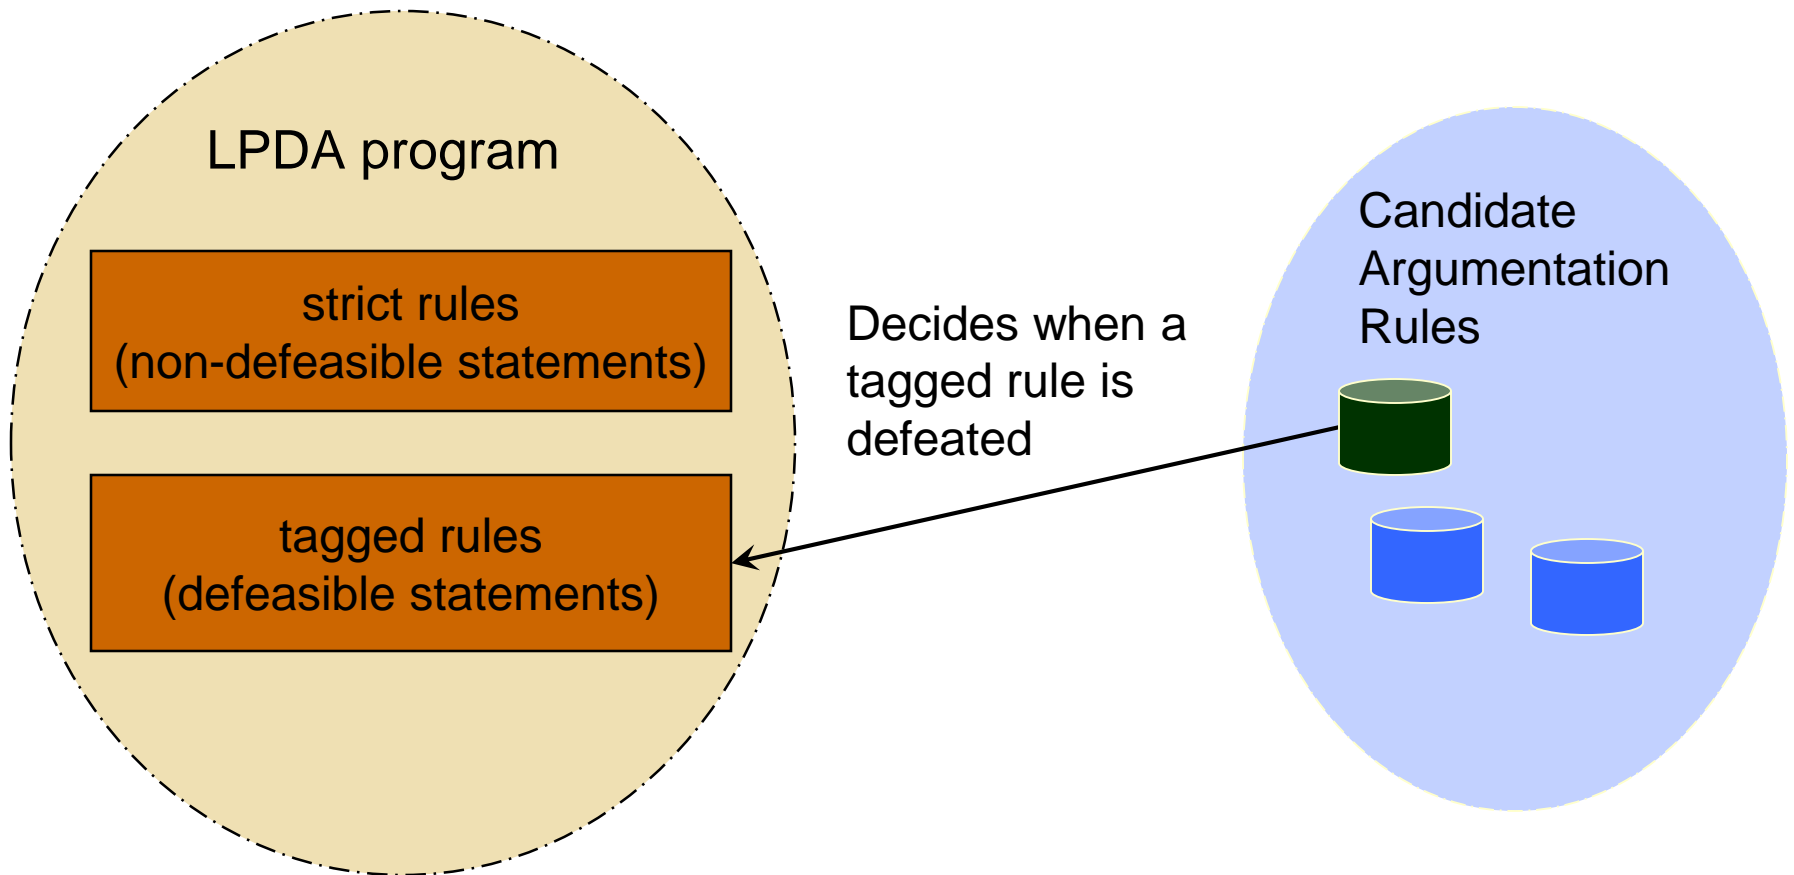
- **More Advantages**

- 1st way to generalize defeasible LP, notably Courteous, to HiLog higher-order and F-Logic frames
- Well-developed model theory, reducible to normal LP
- Reducibility results
- Well-behavior results, e.g., guarantees of consistency
- Unifies almost all previous defeasible LP approaches
 - Each reformulated as an argumentation theory
 - E.g., Defeasible Logic (see Wan, Kifer, and Grosz RR-2010 / SWJ 2015 paper)
- Cleaner, more flexible and extensible semantics
 - Enables smooth and powerful integration of features
 - Applies both to well founded LP (WFS) and to Answer Set Programs (ASP)
- Leverages most previous LP algorithms & optimizations

- **Implemented** in Ergo, also earlier in Flora-2 and used in SILK

LPDA Framework

- Logic Programs with Defaults and Argumentation rules



Example – AT for Courteous (\mathcal{AT}^{GCLP})

defeated(?R) :- defeats(?S, ?R).

defeats(?R, ?S) :- refutes(?R, ?S) or \$rebut(?R, ?S).

Prioritization (user specified)

refutes(?R, ?S) :- conflict(?R, ?S), **\overrides**(?R, ?S).

refuted(?R) :- refutes(?R2, ?R).

Default negation (NAF)

rebut(?R, ?S) :- conflict(?R, ?S),
naf refuted(?R), **\naf** refuted(?S).

Meta predicates (“Reflection”)

candidate(?R) :- body(?R, ?B), call(?B).

conflict(?R, ?S) :- candidate(?R), candidate(?S),
\opposes(?R, ?S).

\opposes(?R, ?S) :- **\opposes**(?S, ?R).

Exclusion (user specified)

\opposes(?L1, ?L2) :- head(?L1, ?H), head(?L2, **\neg** ?H).

Explicit negation

Ecology Ex. of Causal Process Reasoning

/* Toxic discharge into a river causes fish die-off. */

/* Init. facts, and an “exclusion” constraint that fish count has a unique value */
occupies(trout,Squamish).

fishCount(0,Squamish,trout,400). /* 1st argument of fishCount is an integer time */
\opposes(fishCount(?s,?r,?f,?C1), fishCount(?s,?r,?f,?C2)) :- ?C1 != ?C2.

/* Action/event description that specifies causal change, i.e., effect on next state */

@{tdf1} fishCount(?s+1,?r,?f,0) :- occurs(?s,discharge,?r) \and occupies(?f,?r).

/* Persistence (“frame”) axiom */

@{pefc1} fishCount(?s+1,?r,?f,?p) :- fishCount(?s,?r,?f,?p).

/* Action effect axiom has higher priority than persistence axiom */

\overrides(tdf1,pefc1).

/* An action instance occurs */

@{UhOh} occurs(1,toxicDischarge,Squamish).

As desired: |= fishCount(1,Squamish,trout,400),
fishCount(2,Squamish,trout,0)

Notes: @... declares a rule tag. ? prefixes a variable. :- means if. != means ≠. \opposes indicates an exclusion constraint between two literals, which means “it’s a conflict if”.

Physics Ex. of Contextual Assumptions

/* “P8: Joe drops a glove from the top of a 100m cliff.

How long does the fall take in seconds?” */

// Initial problem-specific facts

AP_problem(P8); fall_event(P8); P8[height->100].

// Action description that specifies causal implications on the continuous process

?e[time->((2 * ?h / ?n)^0.5)] :- fall_event(?e) \and ?e[height->?h, net_accel->?n].

?e[net_accel->(?g - ?a)] :- fall_event(?e) and

?e[gravity_accel->?g, air_resistance_accel->?a].

// Other facts

?e[gravity_accel->9.8] :- loc(?e, Earth).

?e[gravity_accel->3.7] :- loc(?e, Mars).

// Contextual assumptions for answering Advanced Placement exam (AP) problems

@{implicit_assumption} loc(?e, Earth) :- AP_problem(?e).

\opposes(loc(?e, Earth), loc(?e, Mars)).

@{implicit_assumption} ?e[air_resistance_accel->0] :- AP_problem(?e).

\overrides(implicitly_stated, implicit_assumption).

As desired: |= P8[net_accel->9.8, time->4.52] // 4.52 = (2*100/9.8)^0.5

Physics Ex. of Contextual Assumptions (in Ergo)

/ “P8: Joe drops a glove from the top of a 100m cliff on Mars.*

*How long does the fall take in seconds?” */*

/ Initial problem-specific facts*/*

AP_problem(P8). fall_event(P8). P8[height->100].

@{explicitly_stated} loc(P8,Mars).

...

*As desired: |= P8[net_accel->3.7, time->7.35] // 7.35 = (2*100/3.7)^0.5*

Example: Ontology Translation, leveraging Hidlog and exceptions

/ Company BB reports operating earnings using R&D operating cost which includes price of a small company acquired for its intellectual property. Organization GG wants to view operating cost more conventionally which excludes that acquisition amount. We use rules to specify the contextual ontological mapping. */*

@{normallyBringOver} ?categ(GG)(?item) :- ?categ(BB)(?item).

**@{acquisitionsAreNotOperating} \neg ?categ(GG)(?item) :-
acquisition(GG)(?item) \land (?categ(GG) :: operating(GG)).**

\overrides(acquisitionsAreNotOperating, normallyBringOver). /* exceptional */

acquisition(GG)(?item) :- price_of_acquired_R_and_D_companies(BB)(?item).

R_and_D_salaries(BB)(p1001). p1001[amount -> \$25,000,000].

R_and_D_overhead(BB)(p1002). p1002[amount -> \$15,000,000].

price_of_acquired_R_and_D_companies(BB)(p1003). p1003[amount -> \$30,000,000].

R_and_D_operating_cost(BB)(p1003). /* BB counts the acquisition price item in this category */

R_and_D_operating_cost(GG) :: operating(GG).

Total(R_and_D_operating_cost)(BB)[amount -> \$70,000,000]. /* rolled up by BB cf. BB's definitions */

Total(R_and_D_operating_cost)(GG)[amount -> ?x] :- /* roll up the items for GG cf. GG's definitions */

***As desired:* |= R_and_D_salaries(GG)(p1001)**

|= \neg R_and_D_operating_cost(GG)(p1003) /* GG doesn't count it */

|= Total(R_and_D_operating_cost)(GG)[amount -> \$40,000,000]

Notation: @{...} declares a rule tag. ? prefixes a variable. :- means if. X :: Y means X is a subclass of Y.
\overrides(X,Y) means X is higher priority than Y.

Trust Mgmt. Ex. of Higher-Order Defaults

illustrating also basic Knowledge-level Communication, and Frame syntax

In Frame syntax: `subject[property -> object]` *stands for* `property(subject,object)`.

/* Trust policy administration by multiple agents, about user permissions */

/* Admin. Bob controls printing privileges including revocation (neg). */

Bob[controls -> print]; Bob[controls -> \neg print]. /* \neg print means it is disallowed.*/

Cara[controls -> ?priv]; /* Cara is the most senior admin., so controls all privileges. */

/* If an administrator controls a privilege and states at a time (t) that a user has a privilege, then the user is granted that privilege. Observe that ?priv is a higher-order variable. */

@{grant(?t)} ?priv(?user) :- ?admin[states(?t) -> ?priv(?user)] and ?admin[controls(?priv)].

/* More recent statements have higher priority, in case of conflict. */

\overrides(grant(?t2), grant(?t1)) :- ?t2 > ?t1.

/* Admins Bob and Cara make conflicting statements over time about Ann's printing */

Cara[states(2007) -> print(Ann)]; Cara[states(2007) -> webPage(Ann)].

Bob[states(2008) -> \neg print(Ann)].

As desired: |= \neg print(Ann), webPage(Ann)

/* Currently, Ann is permitted a webpage but not to print. */

Notes: `@[...]` declares a rule tag. `?` prefixes a variable. `:-` means if. `!=` means \neq . `neg` is strong negation. There is an implicit exclusion (`\opposes`) between `P` and `neg P`, for every literal `P`.

Outline of Part B. Concepts & Foundations

Key features

- Spirit of LP
- **Higher-Order** Syntax via Hilog. Reification.
- Rule ID's
- Default Negation and Well Founded Semantics
- **Restraint**: semantic bounded rationality
- **Tabling Algorithms**. Knowledge Debugging Methods.
- **Defeasibility** via Argumentation Rules. Remediating FOL's Fragility.
- **General Formulas**, Existentials and Skolems, Omni-directional Disjunction
 - Representing Text. Importing full OWL.
- **Uncertain/Probabilistic** knowledge and reasoning
- *Other features covered in Parts (A.) and (C.):*
 - External Querying. **Federation & orchestration**.
 - **Textual Rulelog**: combining closely with Natural Language Processing.
 - **Explanation**. Terminology/ontology mapping.
- *Other features not covered much due to limitations of time & focus:*
 - Frame syntax (a.k.a. F-Logic), Object Oriented style
 - Reactiveness
 - Misc. Lesser Features: Datatypes, Aggregation, Integrity Constraints, Inheritance, Equality, "Constraints", "Modules"

Omniform Rules: Clausal case

- Rulelog introduces the concept of an omniform (“*omni*”) rule.
- Basic case is clausal. Here, the clause is treated *omni-directionally*.
 - $@\{G\} FC$. where FC has the syntactic form of a FOL clause
 - The prioritization tag ($@\{G\}$) is optional. Outer universal quantification is implicit.
 - E.g., $@\{hi\} \text{wet}(\text{lawn}, \text{nextMorning}(\text{?night})) \setminus\text{or} \setminus\text{neg occur}(\text{rain}, \text{?night}) ;$
- A clausal omni rule is transformed, i.e., directionalized, from $@\{G\} L_1 \setminus\text{or} L_2 \setminus\text{or} \dots \setminus\text{or} L_k;$ where each L_i is an atom or the $\setminus\text{neg}$ of an atom into a set of k directed rules, one for each choice of head literal:
 - $@\{G\} L_1 :- \setminus\text{neg} L_2 \setminus\text{and} \setminus\text{neg} L_3 \setminus\text{and} \dots \setminus\text{and} \setminus\text{neg} L_k .$
 - $@\{G\} L_2 :- \setminus\text{neg} L_1 \setminus\text{and} \setminus\text{neg} L_3 \setminus\text{and} \dots \setminus\text{and} \setminus\text{neg} L_k .$
 - ...
 - $@\{G\} L_k :- \setminus\text{neg} L_1 \setminus\text{and} \setminus\text{neg} L_2 \setminus\text{and} \dots \setminus\text{and} \setminus\text{neg} L_{k-1} .$
- This is called the set of *directional variant* rules.
- Avoids unrestricted reasoning by cases!!!
 - Cf. unit/linear resolution strategy in FOL
- Sound with respect to FOL semantics of the clauses. I.e., “hypermonotonic”.

$\setminus\text{naf-free} !$

Examples of Directionalization

- $\text{@\{hi\ wet(lawn, nextMorning(?night))} \leq \text{Occur(rain, ?night)} . \quad /* \text{Causal} */$
is transformed into:
 - $\text{@\{hi\ Wet(lawn, nextMorning(?night))} :- \text{Occur(rain, ?night)} ;$
 - $\text{@\{hi\ \neg Occur(rain, ?night)} :- \neg \text{Wet(lawn, nextMorning(?night))} ;$
- $\neg (\text{Cat}(?x) \wedge \text{Bird}(?x)) . \quad /* \text{OWL-DL disjoint classes} */$
is transformed into:
 - $\neg \text{Cat}(?x) :- \text{Bird}(?x) .$
 - $\neg \text{Bird}(?x) :- \text{Cat}(?x) .$
- $\neg \text{Approved}(?p) \leq \neg \text{Validated}(?p) ; \quad /* \text{SBVR: Car Rental Constraint} */$
is transformed into:
 - $\neg \text{Approved}(?p) :- \neg \text{Validated}(?p) .$
 - $\text{Validated}(?p) :- \text{Approved}(?p) .$
- $\text{mtg}(3p) \vee \text{mtg}(4p) \vee \text{mtg}(5p) . \quad /* \text{Scheduling: Joe's meeting time} */$
is transformed into:
 - $\text{mtg}(5p) :- \neg \text{mtg}(3p) \wedge \neg \text{mtg}(4p) .$
 - $\text{mtg}(4p) :- \neg \text{mtg}(3p) \wedge \neg \text{mtg}(5p) .$
 - $\text{mtg}(3p) :- \neg \text{mtg}(4p) \wedge \neg \text{mtg}(5p) .$

Omnis: General case

- **Permit the formula F to:**
 - Have the form of any FOL formula (“FOL-like”)
 - Also use [HiLog](#) and [Frame](#) features
- **Permit a rule [body](#) too**
 - $@G \ F \ :- \ B \ .$
 - Adds B to the body of each directional variant rule
 - B is similar in form to F , but also permits [NAF](#)
 - Special case: F is a literal
- **Semantics of existentials has subtleties**
 - Use skolemization, via a *tight* normal form (TNF) that’s a bit different from Skolem NF. Argumentation theory is tweaked.
- **Omni feature raises the KR abstraction level**
 - Hide directionality ($\ :- \$) as well as NAF ($\ \text{naf} \$)
 - Use instead: $\ \text{neg}$ (strong negation), $\ \text{<==}$ (strong/material implication), and defeasibility (courteous)

Remedying FOL Semantics' Lack of Scalability

- **Rulelog handles conflict robustly – get consistent conclusions**

- **Whereas FOL is a “Bubble” – it’s perfectly brittle semantically in face of contradictions from quality problems or merging conflicts.**

- Any contradiction is totally contagious – the conclusions all become garbage

E.g., OWL beyond the RL subset suffers this problem. So does Common Logic. (Technically, RIF-BLD and RDF(S) are defined via FOL semantics too, although their typical implementations are essentially LP.)

A KB with a million or billion axioms formed by merging from multiple Web sources, is unlikely to have zero KB/KA conflicts from:

- Human knowledge entry/editing
- Implicit context, cross-source ontology interpretation
- Updating cross-source
- Source trustworthiness

- ***Rulelog’s approach provides a critical advantage for KB scalability***

- ***semantically, as well as computationally***

FOL: A Bubble

Extreme sensitivity to conflict limits its scalability in # of axioms and # of merges



Left:

<http://www.dailymail.co.uk/sciencetech/article-1199149/Super-slow-motion-pictures-soap-bubble-bursting-stunning-detail.html>

Above:

http://img.dailymail.co.uk/i/pix/2007/11_03/BubblePA_468x585.jpg

KR Conflict Handling – A Key to Scalability

BEFORE

KR: Classical Logic (FOL, OWL)

Contradictory conflict is globally contagious, invalidates all results.




Knowledge integration involving conflict is labor-intensive, slow, costly.




AFTER

KR: LP with Defaults (Courteous-style)

Contradictory conflict is contained locally, indeed tamed to aid modularity.



Knowledge integration involving conflict is highly automated, faster, cheaper.



Outline of Part B. Concepts & Foundations

Key features

- Spirit of LP
- **Higher-Order** Syntax via Hilog. Reification.
- Rule ID's
- Default Negation and Well Founded Semantics
- **Restraint**: semantic bounded rationality
- **Tabling Algorithms**. Knowledge Debugging Methods.
- **Defeasibility** via Argumentation Rules. Remediating FOL's Fragility.
- **General Formulas**, Existentials and Skolems, Omni-directional Disjunction
 - Representing Text. Importing full OWL.
- **Uncertain/Probabilistic** knowledge and reasoning
- *Other features covered in Parts (A.) and (C.):*
 - External Querying. **Federation & orchestration**.
 - **Textual Rulelog**: combining closely with Natural Language Processing.
 - **Explanation**. Terminology/ontology mapping.
- *Other features not covered much due to limitations of time & focus:*
 - Frame syntax (a.k.a. F-Logic), Object Oriented style
 - Reactiveness
 - Misc. Lesser Features: Datatypes, Aggregation, Integrity Constraints, Inheritance, Equality, "Constraints", "Modules"

Uncertainty, including Probability

- Many important and popular kinds of uncertain knowledge and reasoning fall within Rulelog's expressiveness
- Syntactic representation of uncertainty – including probability – “comes for free” with Hilog and strong meta
 - Represent uncertain knowledge as
 $\text{uncertainty}(?formula, ?parameters)$
 - E.g., parameters might be point value, lower/upper bounds, confidence level, sample-size, statistical technique, etc.
- Rulelog can then be used to axiomatize desired reasoning principles with such statements – “roll-your-own”
 - E.g., leveraging arithmetic operations, defeasibility, restraint, etc.
- ErgoAI also has been developing optimized support for specific kinds of probabilistic uncertainty
 - Evidential reasoning: weighted or prioritized combination
 - Distribution semantics: semantics/foundation of Probabilistic LP

Reasoning with Uncertainty using Rulelog

Mature:

- Treat probabilistic statements as a special case of general logical statements and approximate the desired type of reasoning by incorporating certainty factors (that represent probability values), into the general Ergo reasoning facilities. I.e., “roll-your-own” uncertain reasoning.
 - Leverages highly expressive features of Rulelog including higher-order syntax
 - Examples of reasoning with uncertainty that can be incorporated into rules:
 - Lower and upper bounds on probability value.
 - E.g., *prob_range(needs_repair(part32), 0.89, 0.94)*.
 - Confidence level. E.g., *prob_range_with_confidence(needs_repair(part32), 0.89, 0.94, 0.001)*.
 - Source and provenance info about the statistical or ML method used to derive probability value/range. This can be a basis for certainty factors.
 - E.g., *source(prob_range(needs_repair(part32), 0.89, 0.93), ML_episode(myFavoriteMLClassifier, 'Michael Kifer', 'Feb 11, 2017', <http://mycompany.com/dataset41>))*.

Additional Ways of Reasoning with Uncertainty in ErgoAI

In development:

- “Evidential reasoning”: combines information about probabilities based on different conditions and associated data sets
 - Addresses the probability whether a particular “thing” (i.e, a person, situation, etc.) belongs to a particular class.
 - Examples:
 - The probability that person X has a given disease
 - The probability that a given transaction is risky
 - The probability that a given airplane part will fail given its age, type, and manufacturer
 - This type of reasoning does not require complete or even consistent knowledge about probabilities and is scalable
- Restricted “distribution semantics” (DS) for probability*
 - ~5 major approaches to probabilistic LP are based on DS, under various names, e.g., ProbLog, PRISM
 - DS’s expressiveness supersedes Bayesian networks
- Triangular-norm style weighted uncertainty*
 - Similar to the uncertainty methods used for neural networks, fuzzy logic, some other ML/KRR systems

Outline of Part B. Concepts & Foundations

Key features

- Spirit of LP
- **Higher-Order** Syntax via Hilog. Reification.
- Rule ID's
- Default Negation and Well Founded Semantics
- **Restraint**: semantic bounded rationality
- **Tabling Algorithms**. Knowledge Debugging Methods.
- **Defeasibility** via Argumentation Rules. Remediating FOL's Fragility.
- **General Formulas**, Existentials and Skolems, Omni-directional Disjunction
 - Representing Text. Importing full OWL.
- **Uncertain/Probabilistic** knowledge and reasoning
- *Other features covered in Parts (A.) and (C.):*
 - External Querying. **Federation & orchestration**.
 - **Textual Rulelog**: combining closely with Natural Language Processing.
 - **Explanation**. Terminology/ontology mapping.
- *Other features not covered much due to limitations of time & focus:*
 - Frame syntax (a.k.a. F-Logic), Object Oriented style
 - Reactiveness
 - Misc. Lesser Features: Datatypes, Aggregation, Integrity Constraints, Inheritance, Equality, "Constraints", "Modules"

Integrity Constraints

Two styles with quite different semantics:


1. Alarm: Rule that detects a violation

- Typical: the rule reports/notifies that constraint is violated
- Other rules infer resulting actions to take
- E.g., many BRMS, Ergo Lite, Ergo

...*VERSUS*...

2. Model-cutting: Rule that forces global contradiction when axiom is violated

- Typical: no model, **lose all useful entailments!!**
- E.g., FOL



Additional Expressive Features in Rulelog

- Explicit equality (and equivalence) reasoning
 - In head of non-fact rules, therefore derived
 - Interaction with nonmonotonicity
 - Key characteristic: substitutivity of equals for equals
 - Related to Herbrand aspect of LP semantics
- Existentials, skolemization
 - RDF blank-nodes, anonymous individuals [Yang & Kifer]
 - Related to Herbrand aspect of LP semantics
- Aggregation (operate on entailed lists): count, total, min, max, etc.
 - Depends on nonmonotonicity, stratification
- Datatypes – they are basic but fairly straightforward
- “Constraints” (e.g., equation/inequality systems)
 - Commonly: via external query/assert to specialized solver
- “Modules”: sub- KBs with information hiding and optimizations
- *Also implemented in Ergo Lite & Ergo, tho’ not part of Rulelog itself: Transaction Logic [Kifer et al] – including KRR about actions’ results, transactionality, and some kinds of hypotheticals*

Outline of Part B. Concepts & Foundations

Key features

- Spirit of LP
- **Higher-Order** Syntax via Hilog. Reification.
- Rule ID's
- Default Negation and Well Founded Semantics
- **Restraint**: semantic bounded rationality
- **Tabling Algorithms**. Knowledge Debugging Methods.
- **Defeasibility** via Argumentation Rules. Remediating FOL's Fragility.
- **General Formulas**, Existentials and Skolems, Omni-directional Disjunction
 - Representing Text. Importing full OWL.
- **Uncertain/Probabilistic** knowledge and reasoning
- *Other features covered in Parts (A.) and (C.):*
 - External Querying. **Federation & orchestration**.
 - **Textual Rulelog**: combining closely with Natural Language Processing.
 - **Explanation**. Terminology/ontology mapping.
- *Other features not covered much due to limitations of time & focus:*
 - Frame syntax (a.k.a. F-Logic), Object Oriented style
 - Reactiveness
 - Misc. Lesser Features: Datatypes, Aggregation, Integrity Constraints, Inheritance, Equality, "Constraints", "Modules"

Outline of Tutorial

A. Overview

- Practical logic. Applications. Features. Software. Textual.
- Case Study Demo and Features Tour
 - Financial regulatory/policy compliance

B. Concepts and Foundations: Drill-downs

- Expressive features, semantics, algorithms;
relationships to natural language and machine learning

C. Conclusions and Future Work

- Background Assumed: basic knowledge of first-order logic, databases. Helpful: XML, RDF and semantic web concepts

Some Recap Points

- Expressiveness
- Federation & orchestration
- Case study, application benefits
- Drilled down
 - Esp.: Textual
 - Also: Higher-order, restraint, tabling, defeasibility, uncertainty, general quantified formulas

Ontologies in Rulelog

- Ontology is a purpose and aspect within KRR, not a language or expressive subset
 - Some expressive subsets of KRR are esp. common and useful for ontology
- The most frequent kinds of ontological modeling/KRR are doable within Rulelog
 - Using the same overall KRR language/tools as non-ontological rules etc.
 - Advantage: reduce effort required for system integration
 - Rulelog strength: computational scalability, e.g., compared to OWL-DL
 - Rulelog expressiveness supersedes RDF-S, OWL-RL, and most of OWL-DL
 - E.g., Ergo imports & translates those into Rulelog
 - Still an open research direction: *maximizing* the treatment of OWL-DL

Ontologies in Rulelog (II)

- Rulelog has pre-defined frame syntax – in object-oriented style, which is good for basic ontological modeling. E.g., `wolf :: animal. // subclass`
 - Subclass hierarchy, property domain/range, subproperty hierarchy, inheritance along classes, property cardinality (using integrity constraints)
- Rulelog strength: mappings between ontologies
 - Perspective: most of a dictionary’s content is about mappings between words
- Rulelog strength: textual and higher-order, e.g., to capture NL terminologies
- Rulelog strength: defeasibility, which increases reuse, e.g., in inheritance
- Rulelog enables methodology of “as-you-go” ontological KA
 - Reduce effort waste, delay, risk
 - ... As compared to typical waterfall methodology of dev ontology before rules

Explanations in Rulelog

- Rulelog has methods for powerful explanations
 - E.g., illustrated in Reg W case study using Ergo system
 - Fully detailed, automatically generated in English, interactively navigable
 - Comprehensible to ordinary users
 - Handles also: **why-not**, and argumentation
- Intrinsic strength: logical; and potentially close to NL abstractions
 - Contrast with most machine learning (ML) methods – esp. neural networks!
 - Heavily numerical, large fanout, often/typically far from NL
 - → cognitively difficult for humans to understand
- Intrinsic strength: natural deduction style proofs
 - Expressive transformations produce within the reasoning procedures a series of LP style steps, each deriving a head literal from a conjunction of literals
 - Most classical-logic proof procedures lack natural deduction style
 - E.g., resolution, tableau
- Textual templates work well for text generation

Machine Learning with Rulelog

- Goal for whole field of AI: bring together ML with KRR
- ML methods typically produce uncertain/probabilistic knowledge
- Machine learning (ML) is a motivation for uncertainty/probability in Rulelog
 - It has been a driver for work on probabilistic LP
- Rulelog can complement ML in several ways
 - Import knowledge from ML: at load time or via external querying
 - Integrate knowledge from multiple ML techniques/systems, episodes/datasets
 - Integrate learned knowledge with human-authored knowledge
 - Feed inferred conclusions as data to ML systems
 - Pose questions, with control parameters, to ML systems
- Rulelog often enables humans to specify complex knowledge concisely, as in NL
 - Say it once
 - ... rather than via labeling tons of examples for a ML method (and providing other supervision to ML)

Standardization of Rulelog

- Rulelog is on standardization path
 - Through RuleML with W3C and Oasis
- Existing draft of large subset as RIF dialect
 - Under RIF Framework for Logical Dialects (FLD)
- Rulelog is an expressive superset of several RIF dialects and OWL profiles
 - E.g., RIF-Core, OWL-RL
- Also relevant to explore relationships to:
 - LegalRuleML
 - HiLog feature used in ISO Common Logic
 - NLP used in SBVR

Rulelog KRR: Advantages for Knowledge Management

- Unprecedented flexibility in the kinds of complex info that can be stated as assertions, queries, and conclusions (highly expressive “knowledge” statements)
 - Almost anything you can say in English – concisely and directly
 - Just-in-time introduction of terminology
 - Statements about statements (meta knowledge)
 - State and view info at as fine a grain size as desired
- Probabilistic info combined in principled fashion, tightly combined with logical
 - Tears down the wall between probabilistic and non-probabilistic
- Unprecedented ease in updating knowledge
 - Map between terminologies as needed, including from multiple sources
- Conflict between statements is robustly handled (often arises during integration)
 - Resolved based on priority (e.g., authority), weighting, or else tolerated as an impasse
- Scalable and computationally well-behaved

Some Other Case Studies

- Financial regulatory compliance decisions:
with databases/ontologies
- Health care treatment guidance:
making decisions with policy-flavored protocols
- Defense intelligence analysis:
with text extraction, databases/ontologies
- Personalized tutoring in continuing/higher ed:
answering science questions
- E-commerce marketing:
with product databases/ontologies, promotion/pricing policies
- *Info on most of these is available at Coherent Knowledge website*

Lessons Learned from Case Studies

Users in these multiple domains benefited from:


- Agility: Flexibility and ease of authoring, fast updating
- High accuracy and transparency
 - Explanations and provenance
 - Lower risk of non-compliance or confusion
- More Cost Effectiveness – less labor, SMEs in closer loop
- Leveraging investment in semantic tech: RDF, SPARQL, OWL

Dream for field of AI: Answer Questions using logic and NLP

- What if it was “*cheap*” to acquire massive volumes of knowledge formally encoded as logical formulas?
 - Say, only a small integer multiple of cost to write quality text, i.e., NL sentences, e.g., about science, policies, medicine, ...
- What if it was “*easy*” to understand natural language questions well enough to exploit such formal encodings?
- *How much could Textual Rulelog help?*

Rulelog enriches Text Extraction

- Leverage Rulelog's high expressiveness and flexibility
- Mappings between multiple terminologies or ontologies



Chatbots, HCI, IoT, and NL Understanding

- *The Promise:* Able to converse with and assist humans in many facets of our lives
 - Provide advice
 - Perform tasks
 - Inform us proactively
 - Explain why
- *Required to fulfill this promise:* Flexible deep reasoning
 - Using logical/probabilistic knowledge representation and reasoning (KRR)
 - Combined with natural language processing (NLP) and machine learning (ML)
 - Treat the **deep** semantics of NL
 - KRR was central to first wave of AI success. KRR + ML = core of AI.

Open Research Topics: Logical (I)

- Probabilistic: fuzzy flavor; Bayesian / distribution semantics flavor; optimization of restricted cases; approximations; hookups to ML approaches
- Reasoning by cases: theory/semantics, algorithms
 - Do selectively. Soundness/relationship to: FOL, ASP, MKNF.
- Hypothetical reasoning, abduction
- Distributed reasoning: algorithms and testbeds
 - Finely parallelized too. Leverage persistent stores.
- Reactive: semantics, event handling/dispatching
 - Good candidate to integrate into Rulelog: Production LP approach (see our AAI-13 rules tutorial)
 - Relate to Reaction RuleML, Prova, production/ECA rules, Transaction Logic

Open Research Topics: Logical (II)

- Explanation methods – improve further
- Tools for debugging knowledge – improve further
- Argumentation rules to work better with omnis
- Equality: axiomatic semantics, efficient algorithms
- Aggregates – handle indefiniteness, unstratified cases
- “Constraints” – cf. constraint LP: theory, algorithms
- Optimizations: e.g.,
 - subgoal re-ordering for efficiency
 - leverage subsumptive tabling (cf. LP)
 - external query plans

Research Directions – Other Aspects

- ❖ Combine closely with deep semantic NL interpretation
- ❖ Complement ML: feed inferences to ML, query ML
 - ❖ Rulelog deduction can also directly be inductive ML
- ❖ Applications
 - NLP and HCI, e.g., for cognitive, IoT
 - Legal. Medical.
 - Defense. Education. Social media. *Many more.*
- ❖ Standards design – with RuleML
 - (In draft): RIF-Rulelog
 - RuleML-Rulelog; relate to Oasis LegalRuleML
 - Profiles (subsets) incl. intersect with OWL
 - Rulelog output from SBVR

Outline of Tutorial

A. Overview

- Practical logic. Applications. Features. Software. Textual.
- Case Study Demo and Features Tour
 - Financial regulatory/policy compliance

B. Concepts and Foundations: Drill-downs

- Expressive features, semantics, algorithms;
relationships to natural language and machine learning

C. Conclusions and Future Work

- Background Assumed: basic knowledge of first-order logic, databases. Helpful: XML, RDF and semantic web concepts

Thank You

Disclaimer: All brands, logos and products are trademarks or registered trademarks of their respective companies.