

# A SILK Graphical UI for Defeasible Reasoning, with a Biology Causal Process Example<sup>\*</sup>

Benjamin Grosz<sup>1</sup>, Mark Burstein<sup>2</sup>, Mike Dean<sup>2</sup>, Carl Andersen<sup>2</sup>, Brett Benyo<sup>2</sup>, William Ferguson<sup>2</sup>, Daniela Inclezan<sup>3</sup>, and Richard Shapiro<sup>2</sup>

<sup>1</sup> Vulcan Inc., Seattle, Washington, USA, [benjaming@vulcan.com](mailto:benjaming@vulcan.com)

<sup>2</sup> Raytheon BBN Technologies, Cambridge, Massachusetts, USA,

{[burstein](mailto:burstein@bbn.com), [mdean](mailto:mdean@bbn.com), [canderse](mailto:canderse@bbn.com), [bbenyo](mailto:bbenyo@bbn.com), [wferguson](mailto:wferguson@bbn.com), [rshapiro](mailto:rshapiro@bbn.com)}@bbn.com

<sup>3</sup> SRI International, Menlo Park, USA, [daniela.inclezan@ttu.edu](mailto:daniela.inclezan@ttu.edu)

**Abstract.** SILK is an expressive Semantic Web rule language and system equipped with scalable reactive higher-order defaults. We present one of its latest novel features: a graphical user interface (GUI) for knowledge entry, query answering, and justification browsing that supports user specification and understanding of advanced courteous prioritized defeasible reasoning. We illustrate the use of the GUI in an example from college-level biology of modeling and reasoning about hierarchically-structured causal processes with interfering multiple causes.

## 1 Introduction to SILK

SILK<sup>4</sup> (Semantic Inferencing for Large Knowledge) is an expressive Semantic Web rule language and system equipped with scalable reactive higher-order defaults. The system includes capabilities for reasoning, knowledge interchange, and user interface (UI). Part of Project Halo<sup>5</sup>, sponsored by Vulcan Inc., the SILK research program addresses fundamental knowledge representation (KR) requirements for scaling the Semantic Web to widely-authored Very Large Knowledge Bases (VLKBs) in business and science that answer questions, proactively supply info, and reason powerfully. The SILK effort has over 15 contributing institutions, including Vulcan, Stony Brook University, Raytheon BBN Technologies, Cycorp, and SRI International.

SILK pushes the frontier of KR by combining expressiveness plus semantics plus scalability. It targets defeasibility, higher-order, and actions — including to support reasoning about complex processes that are described in terms of causality, hierarchical structure, and/or hypothetical scenarios. For example, reasoning about causal processes is a large portion of first-year college biology, often requiring multi-step causal chains and/or multiple grain sizes of description to answer a textbook or exam question. Longer-term, SILK targets widely collaborative KA by subject matter experts (SMEs), such as science students/teachers or business people, not just knowledge engineers (KEs) or programmers.

---

<sup>\*</sup> This work is part of the SILK project sponsored by Vulcan Inc.

<sup>4</sup> <http://silk.semwebcentral.org>

<sup>5</sup> <http://projecthalo.com>

SILK’s more expressive yet performant KR promises to improve knowledge acquisition (KA) in two ways. First, expressive declarative semantics facilitates and enables web knowledge interchange, which in turn aids scalability of collaborative KA. Second, SILK raises the abstraction level of the underlying KR, which is the target for direct user knowledge entry, to be closer to the user’s natural language and cognition. “*The KR is the deep UI*”.

SILK has a new fundamental KR: *hyper* logic programs, which extends normal declarative logic programs (LP). Hyper LP is the first to tightly *combine* several key advanced expressive features: *defaults*, with strong negation and priorities, cf. courteous LP [1] with argumentation theories [2] and omni-directionality [3]; (quasi) *higher-order* syntax, reification, and meta-reasoning, cf. HiLog [4] and Common Logic and procedural attachments to *external* actions (side-effectful), queries (to built-ins, web sources or services), and events (knowledge update flows), cf. situated/production LP [1] [5] (and similar to production rules). Other advanced hyper LP expressive features include: webized syntax and interchange cf. W3C Rule Interchange Format (RIF), (and the earlier RDF, RuleML, and OWL); frame syntax (object-oriented style) cf. F-Logic [6]; and semantically clean negation-as-failure (NAF), cf. well-founded [7].

KR languages supported for interchange include: SPARQL [8], and RDF(S); [9]; SQL and ODBC (e.g., Excel spreadsheets); SILK, RIF (-BLD and -SILK), and OWL (-RL); [10]; Cyc [11] (most of its KR and KB); and AURA [12]. AURA is a Project Halo system for question-answering in first-year college science and currently has a KB with tens of thousands of axioms about biology. AURA largely pre-dates SILK and employs a frame-based KR that is considerably less expressive than SILK.

The SILK reasoning subsystem primarily does backward inferencing, i.e., query answering. However, it employs LP *tabling* techniques [13] to save and reuse computation from previous subqueries. It thereby also supports forward inferencing, incremental updating, and persistent queries/views.

Figure 1 shows the current architecture of SILK.

SILK’s UI and API are written in Java, built on top of an extended version of Flora-2<sup>6</sup>, which is built on top of XSB Prolog [14], which is written in C. XSB provides a normal LP reasoning kernel that is highly efficient [15].

**Outline and Contributions:** A previous version of SILK was presented in [16] [17]. In the rest of this paper, we present a novel addition to SILK since then: a graphical user interface (GUI) for KA and querying that treats defeasibility (section 2). We present an example of complex causal biology process reasoning (section 3), of the kind to be used in our accompanying system demonstration. Comparison with related work is given especially in section 2. We conclude by discussing future work (section 4).

---

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

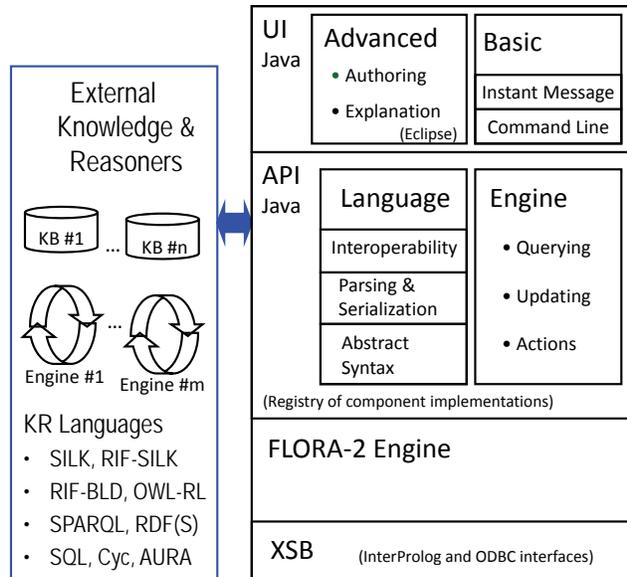


Fig. 1. SILK Version 2.2 System Architecture

## 2 SILK Graphical User Interface & Defeat Justifications

We have developed a graphical user interface (GUI) to the SILK system for knowledge entry, query answering, and justification browsing. The GUI is currently used by KEs and is being extended to support use by SMEs. The GUI supports user specification and understanding of advanced courteous prioritized defeasible reasoning. It is implemented as a plug-in to the Eclipse <sup>7</sup> Integrated Development Environment (IDE).

The GUI, pictured in Figure 2, offers users a number of capabilities. Entered SILK statements are syntactically validated and statement components (e.g., annotations) are color-coded for clarity. User debugging of rule bases is facilitated by automatic tracking of target queries' results against user changes to the rules. This also allows what-if explorations. The GUI also offers query result justification trees (technically, graphs) that can be explored incrementally, by expanding each tree node to display its children. At each node, the user can specify particular bindings to filter the portion of the justification tree that is displayed. Trees of this sort are also available for negative results (i.e., when a literal *cannot* be inferred), allowing developers to drill down and identify flaws in a desired chain of logical reasoning. This display mechanism also supports the reasoning chains found in courteous defaults by showing *defeated* ground rule

<sup>7</sup> <http://eclipse.org>

instances — rules whose heads are not true, despite their bodies being true, due to conflict with other rules. Figure 2 shows an example of *refutation*-flavor defeat of a rule instance (and thus of its head atom  $A$ ). The rule instance has a *candidate* argument — i.e., the rule’s body is satisfied. But there also is a candidate *counterargument* (whose head is *neg A*) that has a *higher-priority* rule tag.

In (ASCII) **SILK syntax**,  $a \# c$  means that  $a$  is an instance of class  $c$ . Skolems are prefixed by the underscore character (“\_”). A courteous rule label term, used for prioritization, is called a *tag*. An explicit tag is specified in a rule via an annotation which optionally begins the rule:  $@[\text{tag}\rightarrow\mathbf{G}] H :- B ;$  specifies a rule with tag  $\mathbf{G}$ . “:-” is the usual LP rule implication connective, i.e., informally “if”.  $H$  and  $B$  are called the rule head and body, respectively; and “;” is the statement ending delimiter. “:-  $B$ ” may be absent, i.e., the body may be empty (true). The tag ( $\mathbf{G}$ ) may be strict (`silk:strict`), i.e., non-defeasible; if so, the “`tag->silk:`” may be omitted.

The GUI also offers a graph view of rule bases, in which terms and relations are displayed as nodes and arcs between nodes. Flexible control of selection and layout of items to display is achieved via running rules that are specified (e.g., by power users) in SILK itself.

To our knowledge, the SILK GUI is only the second justification exploration GUI for (prioritized) *defeasible* rules that have (declarative, model-theoretic) *semantics*. The first such system was DR-DEVICE [18], which displays defeat justifications, but with less extensive GUI functionality than SILK provides. Its Defeasible Logic KR is closely related to the courteous feature of hyper LP (see [19] for a comparison), but lacks higher-order and several other advanced expressive features of hyper LP.

### 3 Example: A Complex Causal Process in Biology

We have developed a novel approach to modeling and reasoning about temporal hierarchically-structured causal processes, that smoothly handles interference/exceptions between multiple causes and elegantly treats the “frame problem” (inertia / persistence of causal fluents). It leverages hyper LP’s prioritized defaults. To fully describe the approach is beyond the scope of this paper, however. Instead, we illustrate the approach with an example of college-level biology that shows the SILK GUI’s novel capability to explore justifications in the presence of prioritized defeat.

In biology and medicine, a key process is the cell cycle in which a cell grows and then divides. (Control failure in this process causes cancer.) The cell cycle is a complex hierarchically-structured process. It consists of two phases (sub-processes): interphase and mitosis, in that temporal order. Interphase, in turn, consists of three subphases G1, S, and G2, in that order. Mitosis too has several subphases. Many of the above subphases in turn have sub-subphases, etc. DNA synthesis occurs during S phase, and indeed begins when S phase begins. This is the knowledge required to answer the following first-year college exam question<sup>8</sup>:

<sup>8</sup> chapter 12 self-quiz question 15 in [20]

Justification for \_DNA\_Synthesis295(\_S\_Phase59(\_Interphase27(\_Cell\_Cycle79)))[occurs\_at -> 1]

Search: Bindings Rules Facts

?x	?i
_DNA_Synt...	1

Refutation of \_DNA\_Synthesis295(\_S\_Phase59(\_Interphase27(\_Cell\_Cycle79)))[occurs\_at -> 1]  
 Candidate argument from rule with tag: intended\_actions\_are\_executed  
 Counter argument  
 Candidate argument from rule with tag: strict  
 **R** neg ?x[occurs\_at->?i] :- (?x # Action and ?i # Step) and prevented(?x)[holds\_at->?i]  
 \_DNA\_Synthesis295(\_S\_Phase59(\_Interphase27(\_Cell\_Cycle79))) # Action  
 1 # Step  
 prevented(\_DNA\_Synthesis295(\_S\_Phase59(\_Interphase27(\_Cell\_Cycle79)))[holds\_at->1]  
 Supported by rule with tag: ?Var786  
 **R** prevented(?x)[holds\_at->?i1] :- (((((?i1 # Step and ?i # Step) and ?x # Action) and ?i[ne  
 1 # Step  
 0 # Step  
 \_DNA\_Synthesis295(\_S\_Phase59(\_Interphase27(\_Cell\_Cycle79))) # Action  
 0[next->1]  
 Prevent\_DNA\_Synthesis # Inhibition  
 Prevent\_DNA\_Synthesis[inhibits->\_DNA\_Synthesis295(\_S\_Phase59(\_Interphase27(\_Cell\_  
 Prevent\_DNA\_Synthesis[occurs\_at->0]

Fig. 2. SILK GUI: exploring the justification of defeat of a biology process rule

A researcher treats cells with a chemical that prevents DNA synthesis from starting. This treatment traps the cells in which part of the cell cycle?

Correct answer: G1.

That is, if DNA synthesis does not occur, then S Phase does not occur, and the cell cycle stops in the preceding phase which is G1.

We have translated into SILK a substantial AURA KB about biological processes, and then extended it with additional SILK rules to create a more expressively powerful SILK KB about biological processes, which correctly answers the above question (formulated in SILK).

AURA uses the Knowledge Machine (KM) [21] KR, and relies heavily on KM's skolemized frames called "prototypes" which are the focus of AURA's user KA. For example, the fact that cars have engines might be represented by the SILK rule:

```
?car[engine -> _Engine92(?car) # Engine] :- ?car # Automobile
```

This says that if ?car is an instance of Automobile, then it is related by the engine property to a skolem instance `_Engine92(?car)` that is specific to that ?car.

In our SILK biological process KB, "prototype" statements like the one above are used to express the knowledge that the cell cycle process is decomposed hierarchically into its particular phases. Additional (non-"prototype") SILK axioms specify the positive and negative effects and preconditions of various processes, and a general-purpose causal progression model of stepwise phased processes as being "intended" to occur. These rules entail that if a step occurs at time  $t$  then the next step should occur at time  $t + 1$  unless it is prevented by some other causal influence. The KB is too large to include in this paper. Some of the important rules for the example question are:

```
// the cell cycle is intended at time step 0
_Cell_Cycle79[intended_at -> 0] ;
// G1 is followed by S Phase in Interphase
?x[subevent -> _G1_Phase60(?x) # G1_Phase[next_event ->
    _S_Phase59(?x) # S_Phase]] :- ?x # Interphase ;
// the first subevent of S Phase is DNA_Synthesis
?x[first_subevent -> _DNA_Synthesis295(?x) # DNA_Synthesis]
    :- ?x # S_Phase ;
// an action intended at time i occurs at time i
@[tag -> intended_actions_are_executed]
    ?a[occurs_at -> ?i] :- ?a # Action, ?a[intended_at -> ?i] ;
// prevented actions don't occur
@[tag -> silk:strict]
neg ?x[occurs_at -> ?i] :- ?x # Action, ?i # Step,
    prevented(?x)[holds_at -> ?i] ;
```

Other rules allow derivation of the fact:

```
prevented(_DNA_Synthesis295(_S_Phase59
    (_Interphase27(_Cell_Cycle79))));
```

Figure 2 shows a key intermediate step in SILK’s inferencing — defeat of a candidate argument that: DNA Synthesis will occur for the question’s focal cell cycle (`_Cell_Cycle79`), as is normal (i.e., “intended”) in the cell cycle’s process’s cascade of causal phase steps. That argument is refuted because there is a higher-priority counter argument (itself undefeated) based on the preventive/inhibitory causal effect of the hypothetical scenario’s chemical treatment.

## 4 Conclusions

Contributions of this work were summarized in the Abstract and at the end of section 2.

A key direction in current and future SILK work is to increase SME friendliness of the UI in collaborative KA and querying, using in part controlled natural language. SILK is now being integrated with other portions of Project Halo, particularly AURA. We are also refining a translator from Cyc to SILK and using it to provide knowledge about biology.

**Acknowledgements:** Thanks to all of the SILK team, particularly Paul V. Haley, Terrance Swift, Michael Kifer, David Gunning, Vinay Chaudhri, and Michael Gelfond.

## References

1. Grosf, B.N.: Representing E-Commerce Rules via Situated Courteous Logic Programs in RuleML. *Electronic Commerce Research and Applications* **3**(1) (2004)
2. Wan, H., Grosf, B., Kifer, M., et al.: Logic Programming with Defaults and Argumentation Theories. In: *Proc. 25th Intl. Conf. on Logic Programming*. (2009)
3. Grosf, B., Andersen, C., Dean, M., Kifer, M.: Omni-directional Hyper Logic Programs in SILK and RIF. In: *Proc. 4th Intl. Web Rule Symp. (RuleML)*. (2010)
4. Chen, W., Kifer, M., Warren, D.: HiLog: A foundation for higher-order logic programming. *Journal of Logic Programming* **15**(3) (February 1993) 187–230
5. Grosf, B.N.: The Production Logic Programs Approach, in a Nutshell: Foundations for Semantically Interoperable Web Rules (December 2005) Working Paper. Version of Dec. 2005. Available on Web at: <http://ebusiness.mit.edu/bgrosf/#PLP>.
6. Kifer, M., Lausen, G., Wu, J.: Logical Foundations of Object-Oriented and Frame-Based Languages. *Journal of ACM* **42** (July 1995) 741–843
7. Przymusiński, T.: Well-founded and Stationary Models of Logic Programs. *Annals of Mathematics and Artificial Intelligence* **12**(3) (1994) 141–187
8. Prud’hommeaux, E., Seaborne, A.: SPARQL Query Language for RDF. W3C Recommendation 15 January 2008, <http://www.w3.org/TR/2008/REC-rdf-sparql-query-20080115/>
9. Brickley, D., Guha, R.: RDF Vocabulary Description Language 1.0: RDF Schema. W3C Recommendation 10 February 2004, <http://www.w3.org/TR/2004/REC-rdf-schema-20040210/>

10. McGuinness, D., van Harmelen, F.: OWL Web Ontology Language Overview. W3C Recommendation 10 February 2004, <http://www.w3.org/TR/2004/REC-owl-features-20040210/>
11. Cycorp, Inc: The Syntax of CycL. <http://www.cyc.com/cycdoc/ref/cycl-syntax.html>
12. Gunning, D., Chaudhri, V.K., et al.: Project Halo Update: Progress Toward Digital Aristotle. AI Magazine to appear, Fall 2010.
13. Swift, T., Warren, D.S.: An abstract machine for SLG resolution: definite programs. In: Proc. Intl. Logic Programming Symposium (ILPS). (1994)
14. Swift, T., Warren, D.S.: XSB: Extending Prolog with Tabled Logic Programming. Theory and Practice of Logic Programming to appear.
15. Liang, S., Fodor, P., Wan, H., Kifer, M.: OpenRuleBench: An Analysis of the Performance of Rule Engines. In: Proc. WWW-2009 Conf. (2009)
16. Grosf, B.N., Dean, M., Kifer, M.: The SILK System: Scalable Higher-Order Defeasible Rules. In: RuleML-2009 Challenge, part of The International RuleML Symposium on Rule Interchange and Applications (RuleML 2009), Las Vegas, Nevada (November 2009) <http://silk.semwebcentral.org/silk-ruleml-2009-challenge.pdf>.
17. Grosf, B.: SILK: Higher Level Rules with Defaults and Semantic Scalability. In: Proc. 3rd International Conference on Web Reasoning and Rule Systems (RR 2009), Chantilly, Virginia, Springer (October 2009) 24–25
18. Bassiliades, N., et al.: Proof explanation in the DR-DEVICE system. In: Proc. 1st Intl. Conf. on Web Reasoning and Rule Systems. (2007)
19. Wan, H., Kifer, M., Grosf, B.: Defeasibility in Answer Set Programs via Argumentation Theories. In: Proc. 4th Intl. Conf. on Web Reasoning and Rule Systems. (2010)
20. Campbell, N.A., Reece, J.B.: Biology. 6th edn. Benjamin Cummings (2002)
21. University of Texas Computer Science Dept - Knowledge Systems Research Group: KM: The Knowledge Machine. <http://userweb.cs.utexas.edu/~mfkb/km/>