The Chimaera Ontology Environment

Deborah L. McGuinness, Richard Fikes, James Rice, and Steve Wilder

Knowledge Systems Laboratory
Stanford University, Stanford, CA
{dlm, fikes, wilder}@ksl.stanford.edu

CommerceOne
Mountain View, CA
rice@jrice.com

Ontologies have become central components in many applications including search, e-commerce, configuration and, arguably, every large web site (at least for organization and navigation). As ontologies become larger, more distributed, and longer-lived, the need for ontology creation and maintenance environments grows. In our work with ontologies and tool environments over the last few years, we have observed growing needs for automated support of two tasks: (1) merging multiple ontologies and (2) diagnosing (and evolving) ontologies. Chimaera is aimed at supporting these two tasks.

We believe these areas will become more critical over time. Merging becomes imperative when multiple terminologies must be used and viewed as one ontology (e.g., when an e-commerce company needs to merge the terminologies of UNSPSC and RosettaNet into one consistent ontology). Merging is also critical when distributed team members need to assimilate two or more ontologies that should work together in an integrated fashion. Similarly, diagnosis of ontologies becomes more critical as ontologies are obtained from more varied sources. One may need to use a number of "standard" vocabularies that make differing assumptions about design, representation, or reasoning. For example, some vocabularies may allow cycles while others do not and some may support disjoint partitions while others may not be able to express the notion (and thus not expect any associated reasoning). Additionally outside vocabularies may not follow required internal naming conventions. Diagnosis along many dimensions may be required to focus a human's attention in areas that are likely to need modification before use in a particular environment. Finally, diagnostic log generation is useful, but even more useful is an environment that supports interaction with the log and assistance in fixing identified problems.

Chimaera is a merging and diagnostic web-based browser ontology environment. Its design and implementation is based on our experience developing other user interfaces for knowledge applications such as the Ontolingua ontology development environment [Farquhar, et al, 1997], the Stanford CML editor [Iwasaki, et al, 1997], the Stanford JAVA Ontology Tool (JOT), the Intraspect knowledge server [Intraspect 1999], two web interfaces [McGuinness, et. al., 1995; Welty, 1996] for the CLASSIC knowledge representation system [Borgida, et. al, 1989], and a collaborative environment for building ontologies for FindUR [McGuinness, 1998]. Its goal is to work with many ontologies, thus we chose to build on a platform that handles any OKBC-compliant [Chaudhri, et. al, 1998] representation system. Chimaera accepts over 15 designated input format choices (such as ANSI KIF, Ontolingua, Protégé, CLASSIC, iXOL, etc.) as well as any other OKBC-compliant form. It will soon be compliant with other emerging standards such as RDF and DAML.

Chimaera contains a simple editing environment in the tool and also allows the user to use the full Ontolingua editor/browser environment for more extensive editing. Ontolingua is not a requirement however; other editors could be used in its place. It facilitates merging by allowing users to upload existing ontologies into a new workspace (or into an existing ontology). Figure 1 shows the result of someone loading in two ontologies (Test1 and Test2) and then choosing the name resolution mode for the ontologies. Chimaera will suggest potential merging candidates based on a number of properties. It generates a name resolution list that may be used as a guide through the merging task. The displayed option in the name resolution list in the figure below shows a suggestion to merge Mammal and Mammalia (since they had similar names). The user sees a display of the places where the two terms appear in the hierarchy (with only the connected portions of the hierarchy displayed). The user may browse the hierarchy in more detail by doing things like expanding subclasses (both Mammal and Mammalia are closed as represented by the closed triangle in the figure). The user may also view the definitions of the terms and, within Ontolingua, the user may also obtain the results of similarly and difference structural comparisons of the definitions as well. The user may then choose to merge the terms with a simple menu choice from the class menu.

Chimaera allows the user to choose the level of vigor with which it suggests merging candidates. Higher settings, for example will look for things like possible acronym expansion (which was extremely valuable in our use of Chimaera on some government knowledge bases).

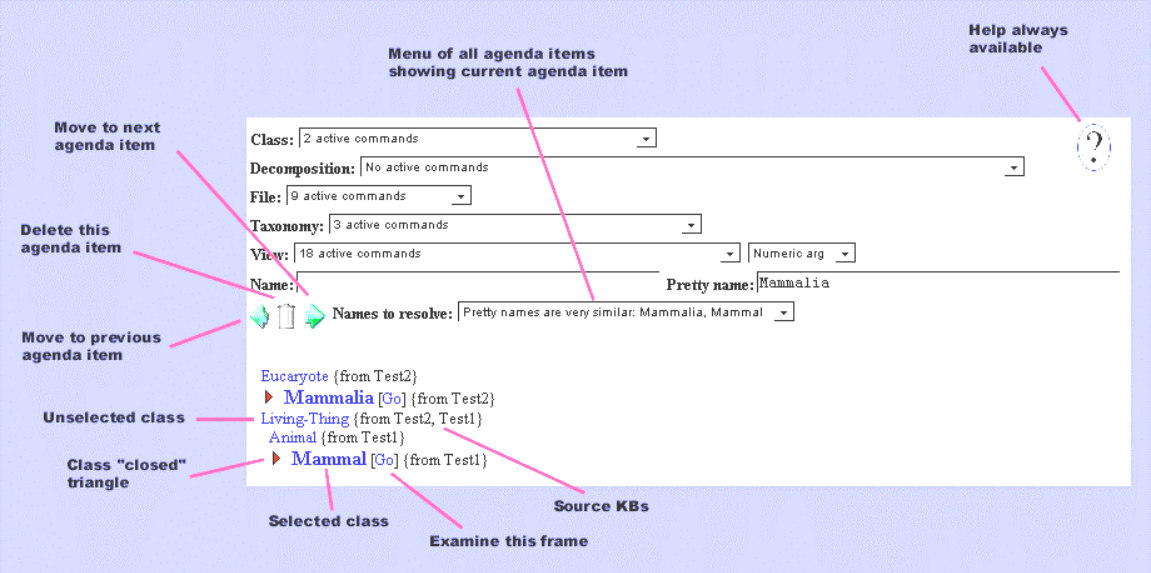

Figure 1: Chimæra in name resolution mode suggesting a merge of Mammal and Mammalia

Chimaera also supports a taxonomy resolution mode. It looks for a number of syntactic term relationships (such as <X-Y> and <Y> since the two are usually subclass related). When attached to a classifier, it can look for semantic subsumption relationships as well

Chimaera includes an analysis capability that allows users to run a diagnostic suite of tests selectively or in their entirety. The output is displayed as an interactive log that allows users to see the results of the tests and also to explore the results. The tests include incompleteness tests, syntactic checks, taxonomic analysis, and semantic checks. We built this system to provide collaborators with varying training essentially a "todo" list containing updates that would likely need to be done before the ontologies would be of the most use to us. The list contains things such as terms that are used but that are not defined, to terms that have contradictory ranges, to cycles detected in the ontology definitions. We are extending the system to include a rule language that allows users to specify additional tests that our environment should include in its diagnostic tool suite so that users may customize the diagnostics to their particular environment.

Chimaera was used in the High Performance Knowledge Base project to analyze incoming ontologies. It is also being used and/or evaluated by companies including VerticalNet and Cisco. More information is available from [McGuinness, et. al, 2000], or from the web site http://www.ksl.Stanford.EDU/software/chimaera/ which also includes links to a tutorial and a movie demonstration. It is licensable for use.

Acknowledgments. The authors wish to thank DARPA for its support of this work under contract N66001-97-C-8554, *Large-Scale Repositories of Highly Expressive Knowledge*.

References.

A. Borgida, R.J. Brachman, D.L. McGuinness, and L.A. Resnick; *CLASSIC: A Structural Data Model for Objects*, SIGMOD, Oregon, 1989

V. Chaudhri, A. Farquhar, R. Fikes, P. Karp, and J. Rice; *OKBC: A Programmatic Foundation for Knowledge Base Interoperability*; AAAI-98.

A. Farquhar, R. Fikes, and J. Rice; *The Ontolingua Server: a Tool for Collaborative Ontology Construction*; Intl. Journal of Human-Computer Studies **46**, 1997.

Intraspect Knowledge Server, Intraspect Corp., 1999. (http://www.intraspect.com/product_info_solution.htm)

Y. Iwasaki, A. Farquhar, R. Fikes, & J. Rice; *A Web-based Compositional Modeling System for Sharing of Physical Knowledge*. Morgan Kaufmann, Nagoya, Japan, 1997.

D.L. McGuinness; *Ontological Issues for Knowledge-Enhanced Search*; Proceedings of Formal Ontology in Information Systems, June 1998. Also in Frontiers in Artificial Intelligence and Applications, IOS-Press, Washington, DC, 1998.

D.L. McGuinness, R. Fikes, J. Rice, and S. Wilder. *An Environment for Merging and Testing Large Ontologies.* Proceedings of Knowledge Representation 2000.

D.L. McGuinness, L.A. Resnick, and C. Isbell; *Description Logic in Practice: A CLASSIC: Application*; IJCAI, 1995.

United Nations Standard Product and Services Classification (UNSPSC) Code organization. http://www.unspsc.org/home.htm

C. Welty; *An HTML Interface for CLASSIC*; Proceedings of the 1996 International Workshop on Description Logics; AAAI Press; November, 1996.