

Elaboration Tolerance of Logical Theories

Eyal Amir

Department of Computer Science,
Gates Building, 2A wing
Stanford University, Stanford, CA 94305-9020, USA
eyala@cs.stanford.edu

We consider the development and modification of logical theories (e.g., commonsense theories). During development of such knowledge bases (KBs) a knowledge engineer makes some design and modeling choices. These decisions may later force the KB to undergo some redesign and rewriting when new knowledge needs to be integrated. We then say that the KB lacks *Elaboration Tolerance*. (McCarthy 1998) illustrated this problem using example elaborations for the toy problem of the Missionaries and Cannibals.

The influence of representation on elaboration tolerance is not well understood at present (see (Gogic *et al.* 1995)). In fact, unlike many other development tasks, there is no model for the development of logical theories in general and commonsense theories in particular.

We investigate elaboration tolerance in two complementary directions: (1) We model the development of logical theories, and explore the theoretical properties of change in our model; and (2) We develop a knowledge representation system that follows the intuitions arising from knowledge and software engineering.

On the theoretical front, we have defined a measure of syntactic distance between formal systems (Amir 1998). This measure counts the number of the operations of *add axiom* and *remove axiom* that one must do to transform one knowledge base to be logically equivalent to another. Comparing various reformulations of KBs, we found that (1) some reformulations can make a knowledge base have a shorter syntactic distance to all theories, but in the extreme a price must be paid in terms of the size of the resulting theory and language; and (2) changing the entailment relation does not have such a drastic impact on syntactic distance.

Currently, we are working on expanding our theoretical model to an Oracle-based search (somewhat similar to some Inductive Logic Programming models). In this model, the Oracle is expected to answer whether the current node is the goal node and, if not, what sentence should hold that does not (or vice versa). Consider a typical situation in which our advisor asks us to expand our KB T to accommodate the fact φ . As we

try to comply with our supervisor's request, we look for theories under our representational constraints/biases. First, we look for what we believe our supervisor is asking for ($T \circ \varphi$). Many times she will tell us that it is not good enough because it satisfies another condition ψ that it should not (or vice-versa). The process continues until a satisfying theory is found. In the future, we intend to apply this model to the examples given by McCarthy and use it to learn commonsense theories.

On the pragmatic front, we defined two variants of an Object-Oriented First Order Logic (OOFOL): monotonic and nonmonotonic. In these logics, an object is a theory with an internal language and an interface language. These objects are connected by links to form complex theories (directional links in the nonmonotonic case). Inheritance was also defined to facilitate knowledge engineering. OOFOL theories support inference that is local to each object and incremental compilation of an object to its interface language.

An application that uses a variant of the nonmonotonic OOFOL is the Logic-based Subsumption Architecture (LSA) (Amir & Maynard-Reid 1999). LSA is built along lines similar to Brooks' Subsumption Architecture but having a logical theory for each layer. The subsumption is carried out by sending logical sentences from one layer to the next, having the lower layer ignore some of its default assumptions if needed.

In the future we will implement our OOFOL for Frame Systems, expand the application of the LSA to *synthetic agents* and evaluate the success of these object-oriented structures using our theoretical model.

References

- Amir, E., and Maynard-Reid, P. 1999. Logic-based subsumption architecture. In *Proc. IJCAI-99*.
- Amir, E. 1998. Towards a Formalization of Elaboration Tolerance: Adding and Removing Axioms. In *Frontiers of Belief Revision*. Kluwer.
- Gogic, G.; Kautz, H.; Papadimitriou, C.; and Selman, B. 1995. The comparative linguistics of knowledge representation. In *Proc. IJCAI-95*, 862-869.
- McCarthy, J. 1998. Elaboration Tolerance. *Common-Sense '98*.