

# Mapping ORM to Datalog: An Overview

Terry Halpin<sup>1</sup>, Matthew Curland<sup>2</sup>, Kurt Stirewalt<sup>2</sup>, Navin Viswanath<sup>2</sup>,  
Matthew McGill<sup>2</sup>, Steven Beck<sup>2</sup>

<sup>1</sup>LogicBlox, Australia and INTI International University, Malaysia

<sup>2</sup>LogicBlox, USA

e-mail: {terry.halpin, matt.curland, kurt.stirewalt, navin.viswanath, matt.mcgill, steven.beck}@logicblox.com

**Abstract:** Optimization of modern businesses is becoming increasingly dependent on business intelligence and rule-based software to perform predictive analytics over massive data sets and enforce complex business rules. This has led to a resurgence of interest in datalog, because of its powerful capability for processing complex rules, especially those involving recursion, and the exploitation of novel data structures that provide performance advantages over relational database systems. ORM 2 is a conceptual approach for fact oriented modeling that provides a high level graphical and textual syntax to facilitate validation of data models and complex rules with nontechnical domain experts. Datalog<sup>LB</sup> is an extended form of typed datalog that exploits fact-oriented data structures to provide deep and highly performant support for complex rules with guaranteed decidability. This paper provides an overview of recent research and development efforts to extend the Natural ORM Architect (NORMA) software tool to map ORM models to Datalog<sup>LB</sup>.

## 1 Introduction

In order to compete effectively in the information age, many businesses are exploiting information technology as a way to promote efficiency and reduce costs. For example, business intelligence tools and rule-based software are being increasingly used to perform predictive analytics over massive data sets and enforce complex business rules. This has led to a resurgence of interest in *datalog*, because of its powerful capability for processing complex rules, especially those involving recursion. Moreover, *novel data structures* such as column-oriented data stores are being exploited to provide performance advantages over relational database systems for complex analytics and data warehousing tasks ([http://en.wikipedia.org/wiki/Column-oriented\\_DBMS](http://en.wikipedia.org/wiki/Column-oriented_DBMS)).

While datalog and related technologies are powerful, the effective use of them typically requires a considerable level of mathematical sophistication. This often results in a communication gap when the business experts, who best understand the complex business rules and queries needed for their business, attempt to validate that the technical rules and queries used in the implementation actually conform to their requirements. This problem is best addressed by first formulating the models, rules and queries at a conceptual level where they can be reliably validated with the busi-