# Abstract

*Datalog*$^{\exists}$ is the natural extension of *Datalog*, allowing existentially quantified variables in rule heads. This language is highly expressive and enables easy and powerful knowledge-modeling, but the presence of existentially quantified variables makes reasoning over *Datalog*$^{\exists}$ undecidable, in the general case. The results in this thesis enable powerful, yet decidable and efficient reasoning (query answering) on top of *Datalog*$^{\exists}$ programs.

On the theoretical side, we define the class of parsimonious *Datalog*$^{\exists}$ programs, and show that it allows of decidable and efficiently-computable reasoning. Unfortunately, we can demonstrate that recognizing parsimony is undecidable. However, we single out *Shy*, an easily recognizable fragment of parsimonious programs, that significantly extends both *Datalog* and *Linear Datalog*$^{\exists}$. Moreover, we show that *Shy* preserves the same (data and combined) complexity of query answering over *Datalog*, although the addition of existential quantifiers.

On the practical side, we implement a bottom-up evaluation strategy for *Shy* programs inside the DLV system, enhancing the computation by a number of optimization techniques. The resulting system is called DLV$^{\exists}$– a powerful system for answering conjunctive queries over *Shy* programs, which is profitably applicable to ontology-based query answering. Moreover, we design a rewriting method extending the well-known Magic-Sets technique to any *Datalog*$^{\exists}$ program. We demonstrate that our rewriting method preserves query equivalence on *Datalog*$^{\exists}$, and can be safely applied to *Shy* programs. We therefore incorporate the Magic-Sets method in DLV$^{\exists}$. Finally, we carry out an experimental analysis assessing the positive impact of Magic-Sets on DLV$^{\exists}$, and the effectiveness of the enhanced DLV$^{\exists}$ system compared to a number of state-of-the-art systems for ontology-based query answering.