
Foreword

The computational foundations of Artificial Intelligence (AI) are supported by two corner stones: logics and machine learning. Computational logic has found its realization in a number of frameworks for logic-based approaches to knowledge representation and automated reasoning, such as Logic Programming, Answer Set Programming, Constraint Logic Programming, Description Logics, and Temporal Logics. Machine Learning, and its recent evolution to Deep Learning, has a huge number of applications in video surveillance, social media services, big data analysis, weather predictions, spam filtering, online customer support, etc.

Emerging interest in the two communities for finding a bridge connecting them is witnessed, for instance, by the prize *test-of-time, 20 years* assigned by the association for logic programming in 2017 to the paper *Hybrid Probabilistic Programs*. Also in 2017, Holger H. Hoos was invited to give the talk *The best of both worlds: Machine learning meets logical reasoning* at the international conference on logic programming. Here, machine learning is used to tune the search heuristics in solving combinatorial problems (e.g., encoded using SAT or ASP techniques). A couple of months later, in a panel organized by the Italian Association for Artificial Intelligence (AI*IA), the machine learning researcher Marco Gori posed five questions to the communities. Among them: *How can we integrate huge knowledge bases naturally and effectively with learning processes? How to break the barriers of machine learning vs (inductive) logic programming communities? How to derive a computational model capable of dealing with learning and reasoning both in the symbolic and sub-symbolic domains? How to acquire latent semantics?* These are fundamental questions that need to be resolved to allow AI research to make another quantum leap. Logical languages can add structural semantics to statistical inference.

This book, based on 15 years of top-level research in the field by Fabrizio Riguzzi and his co-authors, addresses these questions and fills most of the gaps between the two communities. A mature, uniform retrospective of several proposals of languages for Probabilistic Logic Programming is reported.

The reader can decide whether to explore all the technical details or simply use such languages without the need of installing tools, by simply using the web site maintained by Fabrizio's group in Ferrara.

The book is self-contained: all the prerequisites coming from discrete mathematics (often at the foundation of logical reasoning) and continuous mathematics, probability, and statistics (at the foundation of machine learning) are presented in detail. Although all proposals are summarized, those based on the distribution semantics are dealt with in a greater level of detail. The book explains how a system can reason precisely or approximately when the size of the program (and data) increases, even in the case on non-standard inference (e.g., possibilistic reasoning). The book then moves toward parameter learning and structure learning, thus reducing and possibly removing the distance with respect to machine learning. The book closes with a lovely chapter with several encodings in PLP. A reader with some knowledge of logic programming can start from this chapter, having fun testing the programs (for instance, discovering the best strategy to be applied during a *truel*, namely, a duel involving three gunners shooting sequentially) and then move to the theoretical part.

As the president of the Italian Association for Logic Programming (GULP) I am proud that this significant effort has been made by one of our associates and former member of our Executive Committee. I believe that it will become a reference book for the new generations that have to deal with the new challenges coming from the need of reasoning on Big Data.

Agostino Dovier
University of Udine