Preface

This book is about neural computation and applied reasoning, and about how they are related and can be integrated within a machine-learning framework. Logic has a long and solid tradition of principled research, and has been influential in virtually all fields of scientifically-based human knowledge. Neural computation has now achieved maturity. Practitioners have been able to use artificial connectionist models in a number of real-world problems, including machine learning, pattern recognition, vision, cognitive modelling, and artificial agents. However, the area still lacks sound computational foundations, as several features of neural-computation models are still under investigation.

The book presents a number of principled, sound neural-computation models. It makes use of a number of applied (nonclassical) logical systems in order to do so. Logics lie at the foundation of computer science and artificial intelligence. For instance, logic research has provided the foundations of the theories of computability, recursion, and complexity. In artificial intelligence, the construction of automated reasoning systems, knowledge representation languages and formalisms, and theorem proving have been pursued since at least the 1950s. All these are based on logics.

More recently, applications in planning, multiagent systems, expert systems, the Semantic Web, learning theory, cognitive modelling and robotics, constraint satisfaction, searching, and the traditional fields of knowledge representation and reasoning have benefited from research on nonclassical logics. In particular, research on temporal, modal, description, intuitionistic, fuzzy, and nonmonotonic logics has been influential in artificial intelligence, both in relation to theory and principled languages for building models and systems.

Recently, nonclassical logics[1] have been found relevant in relation to the formalisation of integrated learning and reasoning [66, 72, 77, 148, 222]. They have contributed to the construction of sound models of cognitive reasoning. As is well known, human cognition successfully integrates the connectionist (brain-inspired)

[1] Nonclassical logics depart from the Aristotelian tradition in the sense that these logics lack some properties of classical logic, such as the law of the excluded middle, or extend them with additional features. For instance, the development of nonmonotonic logics, which has contributed to

and symbolic (mind-inspired) paradigms of artificial intelligence, language being a compelling case in point. Yet the modelling of cognition develops these separately in the fields of neural computation (and statistical artificial intelligence) and symbolic logic/artificial intelligence. There is now a movement towards a fruitful middle way in between these extremes, in which the study of logic is combined with connectionism. It is essential that these be integrated, thereby enabling not only a technology for building better intelligent systems, but also the understanding of fully fledged computational cognitive models.

The aim of neural-symbolic computation is to explore the advantages that each paradigm presents. Among the advantages of artificial neural networks are massive parallelism, fault tolerance (robust learning), efficient inductive learning, and effective generalisation capabilities. On the other hand, symbolic systems provide descriptions (as opposed to only discriminations); can explain their inference process, for example through automatic theorem proving; and use powerful declarative languages for knowledge representation and reasoning.

In this book, we explore the synergies of neural-symbolic integration from the following perspective. We use a neural network to simulate a given task. The network is obtained by being programmed (set up) or by adapting to and generalising over well-known situations (learning). The network is the mechanism to execute the task, while symbolic logic enables the necessary interaction between the network and the outside world. Differently from and complementing [66], which set the scence for neural-symbolic integration with the use of standard networks and logic programming, we are concerned here with applied (and nonclassical) logical systems. With this consideration at the forefront, *the book presents a rich cognitive model for reasoning and learning based on neural-network ensembles*. The book also illustrates the effectiveness of the model by experimentation, and shows that the connectionist model can compute a number of combined nonclassical logical systems [42], including modal, temporal, epistemic, and intuitionist logics. Finally, the book investigates the issue of relational learning and the representation of first-order logic, the combination (fibring) of models/networks, qualitative reasoning under uncertainty, and how neural networks can offer an effective approach to learning in argumentation frameworks [39, 212]. We conclude by summarising our case for a logic-based, cognitive connectionist model which we call *fibred network ensembles*. Overall, it offers a rich model from a symbolic computation/reasoning viewpoint, yet it is relatively simple and efficient as a learning model. We believe it strikes the right balance between expressiveness and computational feasibility. And, as a matter of principle, we believe that this is no coincidence, but a direct result of

---

the formalisation of practical or commonsense reasoning, led to the discovery of several new logical systems in which the principle of monotonicity may not hold. Further, several authors have defended a position that intuitionistic logics are more appropriate as a foundational basis of computation than classical logic [258]. Classical logic does not capture all of the nature of human practical reasoning. For instance, in order to formalise and automate computational (human) reasoning about time, knowledge, beliefs, or uncertainty several nonclassical logics have been found more appropriate than classical logic; see [42, 87, 100, 102]. Our point will be made clearer as the reader proceeds to the coming chapters.

our methodology, which seeks to unify systematically the most advanced concepts of symbolic computation with the physical constraints of a realistic connectionist machine.

Our aim here is to offer a principled way in which robust learning and effective reasoning can be realised. In line with one of the recent grand challenges in computer science, according to the British Computer Society – nonclassical computation – we seek to unify the classical and nonclassical (natural) computational paradigms, with the nonclassical computational models presented here showing how alternative formulations can deal with cognitive dimensions usually studied under separate formalisms.

London, UK, and Porto Alegre, Brazil, February 2005 to June 2008.

Acknowledgements Over the last five years, we have had stimulating (brief and long) conversations about this book with several researchers in personal meetings, conferences, and workshops. We would like to acknowledge the feedback from Samson Abramsky, Sebastian Bader, Howard Barringer, Rodrigo de Salvo Braz, Krysia Broda, Jehoshua (Shuki) Bruck, Matthew Cook, Cicero Garcez, Robert Givan, Joe Halpern, Pascal Hitzler, Wilfrid Hodges, Ian Hodkinson, Steffen Hölldobler, Bob Kowalski, David Makinson, Ryszard Michalski, Sanjay Modgil, Stephen Muggleton, Vasile Palade, Günther Palm, Luc de Raedt, Oliver Ray, Odinaldo Rodrigues, Dan Roth, Stefan Rüger, Alessandra Russo, Jude Shavlik, Jörg Siekmann, Ron Sun, Guglielmo Tamburrini, John G. Taylor, Francesca Toni, Johan van Benthem, Moshe Vardi, John Woods, and Gerson Zaverucha. We thank them all for their challenging arguments and questioning on the subject, and we apologise for any omissions. We also would like to thank the participants of the International Workshops on Neural-Symbolic Learning and Reasoning (NeSy), organised yearly since 2005 either at IJCAI or ECAI, for the enjoyable research atmosphere and useful debates during our workshop meetings. We particularly would like to thank the NeSy keynote speakers Luc de Raedt, Marco Gori, Steffen Hölldobler, Lokendra Shastri, Ron Sun, Stefan Wermter, and Kai-Uwe Kuehnberger.

We also acknowledge our publisher, Ronan Nugent of Springer, for his support throughout the writing of the book.

Artur d'Avila Garcez has been partly supported by grants from the Nuffield Foundation and the Royal Society; Luís Lamb has been partly supported by grants from the Brazilian Research Council (CNPq) and the CAPES foundation; and Dov Gabbay has been partly supported by EPSRC and by the Leverhulme Trust. This manuscript was written using LaTeX.

http://www.springer.com/978-3-540-73245-7

Neural-Symbolic Cognitive Reasoning
D'Avila Garcez, A.S.; Lamb, L.C.; Gabbay, D.
2009, XIV, 198 p. 53 illus., Hardcover
ISBN: 978-3-540-73245-7