Journal of Applied Logic
Special Volume on Neural-Symbolic Systems

Editorial

The past 25 years have witnessed a tremendous progress in the nature and scope of Logic and its application in Computing Science. At the same time, distributed learning systems and, in particular, Neural Networks, have been playing a central role in the development of Artificial Intelligence (AI). Notwithstanding, human beings do not use these techniques in isolation, but in an integrated way. The study of the integration of Logics and Neural Networks has now become crucial to the development of more effective AI and learning systems in computing.

In this special issue of the Journal of Applied Logic, we investigate how the symbolic and connectionist approaches to AI can be combined and integrated into Neural-Symbolic Systems. Our aim is to benefit from the advantages presented by each paradigm. By integrating Logic and Neural Networks, Neural-Symbolic Systems may provide (*i*) a logical characterisation of a connectionist system, (*ii*) a connectionist implementation of a logic, or (*iii*) a hybrid system bringing together advantages from connectionist systems and symbolic AI.

Firstly, a connectionist implementation of a logic can be produced with the use of a translation algorithm responsible for *representing* a logical theory in a neural network architecture. The translation must be based in a theorem showing the soundness of the algorithm, which allows the *reasoning* about the theory to be carried out in parallel in the network. If the network is a simple, standard connectionist model, knowledge acquisition and theory revision may take place as a result of *learning* from examples in the network, which can be performed with the help of a number of neural learning algorithms such as Backpropagation.

Secondly, a logical characterisation of a connectionist system may be provided by algorithms for rule extraction from neural networks. Rule extraction provides neural networks with *explanation* capability, the lack of which being frequently cited as neural nets main drawback. Generally speaking, rule extraction algorithms perform the inverse relation of that performed by the translation algorithms, once a trained neural network is given. As before, a proof of soundness of the rule extraction algorithm should be provided, whereas a proof of completeness would render the rules produced by the extraction algorithm equivalent to the original network. Equally important to these proofs, there are more practical issues of rule accuracy and comprehensibility, and issues of algorithmic complexity to be considered. In a nutshell, there is a trade-off between the quality of the rules and the complexity of the extraction algorithm.

Thirdly, there is a range of hybrid systems, which combine different aspects of symbolism and connectionism. Notably, neuro-fuzzy-evolutionary systems have been effective in a number of applications in the areas of bioinformatics, medicine, engineering, business, etc. Also important is the integration of uncertainty into these systems and, as a result, the study of logic and probability (which has been the subject of a recent Journal of Applied Logic special issue), and of probabilistic neural networks. The provision of a general framework for combining symbolic and sub-symbolic systems is now required. A starting point is the Fibring methodology for combining logical systems, Bayesian and neural networks.

The papers presented in this special issue are a selection of high-quality research work on Neural-Symbolic Systems. They provide the reader with an excellent view of the above three aspects of Neural-Symbolic integration. First, in *Logic Programs and Connectionist Networks*, Hitzler, Holldobler and Seda offer an overview of the area and study the semantic connection between logic programs and neural networks using a metric space/topological approach. They show how to compute the Logic Programming immediate consequence operator using recursive neural networks, and use approximation techniques for computing the operator in the first-order case. Then, in *Logic Programs, Iterated Function Systems and Recurrent Radial Basis Function (RBF) Networks*, Bader and Hitzler provide an algorithm for computing first order logic programs containing function symbols using RBF networks. They do so from the observation that logic programs can be transformed into a dynamic iterated function system that leads to fractal-like pictures, and show that RBF networks can be used to approximate such a function system to any desired degree of accuracy. These two papers are followed by a different account of how to represent and reason about logic programs in neural networks. *A Neural Implementation of Multi-Adjoint Logic Programming*, by Medina, Merida-Casermeiro and Ojeda-Aciego, introduces an algorithm for representing multi-adjoint programs in neural networks. First, a procedure is proposed to translate multi-adjoint programs (that may contain several types of implication) into a homogeneous program (that contains a single type of implication). Then, a generic type of processing unit is proposed so that homogeneous programs can be translated into neural network architectures. Proofs of preservation of semantics and equivalence between the programs and the networks are then provided. Moving from representation to explanation, in *Is it Worth Generating Rules* *from Neural Network Ensembles?*, Bologna evaluates the performance of a number of methods for rule extraction from neural networks using traditional examples and real-world application problems in medicine and bioinformatics. An empirical comparative analysis between ensembles of neural networks and decision trees then follows, providing an answer to the question posed in the title of the paper. Finally, in *An Evolutionary System for Neural Logic Networks using Genetic Programming and Indirect Encoding*, Tsakonas, Aggelis, Karkazis and Dounias present a Hybrid System, which combines evolutionary computation and neural-logic networks. The system uses a genetic programming technique to guide the learning process of neural-logic networks. As a result, rule extraction and explaining the behaviour of neural-logic networks is facilitated as the structure of the network prior to learning is maintained after learning.

This special issue of the Journal of Applied Logic initiates the journal's scientific area on Logic and Neural Networks, which will serve as the first permanent forum for the publication of cutting-edge research on Neural-Symbolic integration. We would like to thank Tina Eliassi-Rad, Jude Shavlik, Howard Blair, Kewen Wang, Andre de Carvalho, Carlos Thomaz, Luis Lamb, Nikola Kasabov, Massimo de Gregorio, Peter Smith, Jon Seng Quah, Alberto de Souza, Igor Aleksander, Marcilio de Souto, Francesca Toni, Vasile Palade, Robert Kozma, David Leake, Rich Maclin, Vladimir Lifschitz, Vladik Kreinovich, Alessandro Sperduti, Rudy Setiono, Hwee Ng, Melanie Hilario, Jiri Sima and Stefan Rueger for their comments and reviews, and Jane Spurr for her invaluable help in the organisation of this special issue.

London, January 2004

Artur Garcez, Dov Gabbay, Steffen Hölldobler and John Taylor (Editors)

Artur d'Avila Garcez (aag@soi.city.ac.uk)
Department of Computing, City University, London, EC1V 0HB, UK

Dov M. Gabbay (dg@dcs.kcl.ac.uk)
Department of Computer Science, King's College London, WC2R 2LS, UK

Steffen Hölldobler (sh@inf.tu-dresden.de)
International Center for Computational Logic, TU Dresden, 01062, Germany

John G. Taylor (john.g.taylor@kcl.ac.uk)
Department of Mathematics, King's College London, WC2R 2LS, UK