

Hybrid Neural Systems

Stefan Wermter
Ron Sun

Springer, Heidelberg, New York

January 2000

Preface

The aim of this book is to present a broad spectrum of current research in hybrid neural systems, and advance the state of the art in neural networks and artificial intelligence. Hybrid neural systems are computational systems which are based mainly on artificial neural networks but which also allow a symbolic interpretation or interaction with symbolic components.

This book focuses on the following issues related to different types of representation: How does neural representation contribute to the success of hybrid systems? How does symbolic representation supplement neural representation? How can these types of representation be combined? How can we utilize their interaction and synergy? How can we develop neural and hybrid systems for new domains? What are the strengths and weaknesses of hybrid neural techniques? Are current principles and methodologies in hybrid neural systems useful? How can they be extended? What will be the impact of hybrid and neural techniques in the future?

In order to bring together new and different approaches, we organized an international workshop. This workshop on hybrid neural systems, organized by Stefan Wermter and Ron Sun, was held during December 4–5, 1998 in Denver. In this well-attended workshop, 27 papers were presented. Overall, the workshop was wide-ranging in scope, covering the essential aspects and strands of hybrid neural systems research, and successfully addressed many important issues of hybrid neural systems research. The best and most appropriate paper contributions were selected and revised twice. This book contains the best revised papers, some of which are presented as state-of-the-art surveys, to cover the various research areas of the collection.

This selection of contributions is a representative snapshot of the state of the art in current approaches to hybrid neural systems. This is an extremely active area of research that is growing in interest and popularity. We hope that this collection will be stimulating and useful for all those interested in the area of hybrid neural systems.

We would like to thank Garen Arevian, Mark Elshaw, Steve Womble and in particular Christo Panchev, from the Hybrid Intelligent Systems Group of the University of Sunderland for their important help and assistance during the preparations of the book. We would like to thank Alfred Hofmann from Springer for his cooperation. Finally, and most importantly, we thank the contributors to this book.

January 2000

Stefan Wermter
Ron Sun

Table of Contents

An overview of hybrid neural systems	1
S. Wermter and R. Sun	

Structured Connectionism and Rule Representation

Layered hybrid connectionist models for cognitive science	14
Jerome Feldman and David Bailey	

Types and quantifiers in SHRUTI: A connectionist model of rapid reasoning and relational processing	28
Lokendra Shastri	

A recursive neural network for reflexive reasoning	46
Steffen Hölldobler, Yvonne Kalinke and Jörg Wunderlich	

A novel modular neural architecture for rule-based and similarity-based reasoning	63
Rafal Bogacz and Christophe Giraud-Carrier	

Addressing knowledge-representation issues in connectionist symbolic rule encoding for general inference .	78
Nam Seog Park	

Towards a hybrid model of first-order theory refinement	92
Nelson A. Hallack, Gerson Zaverucha and Valmir C. Barbosa	

Distributed Neural Architectures and Language Processing

Dynamical recurrent networks for sequential data processing	108
Stefan C. Kremer and John Kolen	

Fuzzy knowledge and recurrent neural networks: A dynamical systems perspective	124
Christian W. Omlin, Lee Giles and Karvel K. Thornber	

Combining maps and distributed representations for shift-reduce parsing	146
Marshall R. Mayberry and Risto Miikkulainen	
Towards hybrid neural learning internet agents	160
Stefan Wermter, Garen Arevian and Christo Panchev	
A connectionist simulation of the empirical acquisition of grammatical relations	177
William C. Morris, Garrison W. Cottrell and Jeffrey L. Elman	
Large patterns make great symbols: An example of learning from example	194
Pentti Kanerva	
Context vectors: A step toward a grand unified representation	204
Stephen I. Gallant	
Integration of graphical rules with adaptive learning of structured information	212
Paolo Frasconi, Marco Gori and Alessandro Sperduti	

Transformation and Explanation

Lessons from past, current issues and future research directions in extracting the knowledge embedded in artificial neural networks	227
Alan B. Tickle, Frederic Maire, Guido Bologna, Robert Andrews and Joachim Diederich	
Symbolic rule extraction from the DIMLP neural network	241
Guido Bologna	
Understanding state space organization in recurrent neural networks with iterative function systems dynamics	256
Peter Tino, Georg Dorffner and Christian Schittenkopf	

Direct explanations and knowledge extraction from a multilayer perceptron network that performs low back pain classification 271

Marilyn L. Vaughn, Steven J. Cavill, Stewart J. Taylor, Michael A. Foy and Anthony J.B. Fogg

High order eigentensors as symbolic rules in competitive learning 287

Hod Lipson and Hava T. Siegelmann

Holistic symbol processing and the sequential RAAM: An evaluation 299

James A. Hammerton and Barry L. Kalman

Robotics, Vision and Cognitive Approaches

Life, mind and robots: The ins and outs of embodied cognition 314

Noel Sharkey and Tom Ziemke

Supplementing neural reinforcement learning with symbolic methods 334

Ron Sun

Self-organizing maps in symbol processing 349

Timo Honkela

Evolution of symbolization: Signposts to a bridge between connectionist and symbolic systems 364

Ronan Reilly

A cellular neural associative array for symbolic vision .. 374

Christos Orovos and James Austin

Application of neurosymbolic integration for environment modelling in mobile robots 390

Gerhard K. Kraetzschmar, Stefan Sablatnög, Stefan Enderle and Günther Palm

An Overview of Hybrid Neural Systems

Stefan Wermter¹ and Ron Sun²

¹ University of Sunderland, Centre for Informatics, SCET
St. Peter's Way, Sunderland, SR6 0DD, UK

² University of Missouri, CECS Department
Columbia, MO, 65211-2060, USA

Abstract. This chapter provides an introduction to the field of hybrid neural systems. Hybrid neural systems are computational systems which are based mainly on artificial neural networks but also allow a symbolic interpretation or interaction with symbolic components. In this overview, we will describe recent results of hybrid neural systems. We will give a brief overview of the main methods used, outline the work that is presented here, and provide additional references. We will also highlight some important general issues and trends.

1 Introduction

In recent years, the research area of hybrid and neural processing has seen a remarkably active development [62, 50, 21, 4, 48, 87, 75, 76, 25, 49, 94, 13, 74, 91]. Furthermore, there has been an enormous increase in the successful use of hybrid intelligent systems in many diverse areas such as speech/natural language understanding, robotics, medical diagnosis, fault diagnosis of industrial equipment and financial applications. Looking at this research area, the motivation for examining hybrid neural models is based on different viewpoints.

First, from the point of view of cognitive science and neuroscience, a purely neural representation may be most attractive but symbolic interpretation of a neural architecture is also desirable, since the brain has not only a neuronal structure but has the capability to perform symbolic reasoning. This leads to the question how different processing mechanisms can bridge the large gap between, for instance, acoustic or visual input signals and symbolic reasoning. The brain uses specialization of different structures. Although a lot of the functionality of the brain is not yet known in detail, its architecture is highly specialized and organized at various levels of neurons, networks, nodes, cortex areas and their respective connections [10]. Furthermore, different cognitive processes are not homogeneous and it is to be expected that they are based on different representations [73]. Therefore, there is evidence from cognitive science and neuroscience that multiple architectural representations are involved in human processing.

Second, from the point of view of knowledge-based systems, hybrid symbolic/neural representations have some advantages, since different, mutually complementary properties can be combined. Symbolic representations have advantages of easy interpretation, explicit control, fast initial coding, dynamic

variable binding and knowledge abstraction. On the other hand, neural representations show advantages for gradual analog plausibility, learning, robust fault-tolerant processing, and generalization. Since these advantages are mutually complementary, a hybrid symbolic neural architecture can be useful if different processing strategies have to be supported. While from a neuroscience or cognitive science point of view it is most desirable to explore exclusively neural network representations, for knowledge engineering in complex real-world systems, hybrid symbolic/neural systems may be very useful.

2 Various Forms of Hybrid Neural Architectures

Various classification schemes of hybrid systems have been proposed [77, 76, 89, 47]. Other characterizations of architectures covered specific neural architectures, for instance recurrent networks [38, 52], or they covered expert systems/knowledge-based systems [49, 29, 75]. Essentially, a continuum of hybrid neural architectures emerges which contains neural and symbolic knowledge to various degrees. However, as a first introduction to the field, we present a simplified taxonomy here: unified neural architectures, transformation architectures, and hybrid modular architectures.

2.1 Unified Neural Architectures

Unified neural architectures are a type of hybrid neural system. They have also been referred to as unified hybrid systems [47]. They rely solely on connectionist representations but symbolic interpretations of nodes or links are possible. Often, specific knowledge of the task is built into a unified neural architecture.

Much early research on unified neural architectures can be traced back to work by Feldman and Ballard, who provided a general framework of structured connectionism [16]. This framework was extended in many different directions including, for instance, parsing [14], explanation [12], and logic reasoning [30, 40, 70–72]. Recent work along these lines focuses also on the so-called *NTL*, *Neural Theory of Language*, which attempts to bridge the large gap between neurons and cognitive behavior [17, 65].

A question that naturally arises is: why should we use neural models for symbol processing, instead of symbolic models? Possible reasons may include: neural models are a more apt framework for capturing a variety of cognitive processes, as is argued in [15, 66, 86, 72]. Some inherent processing characteristics of neural models, such as similarity-based processing, [72, 6] make them more suitable for certain tasks such as cognitive modeling. Learning processes may be more easily developed in neural models, such as gradient descent [63] and its various approximations, Expectation-Maximization, and even Inductive Logic Programming methods [26].

There can be two types of representations [77]: *Localist connectionist architectures* contain one distinct node for representing each concept [42, 71, 67, 3, 58,

31, 66]. *Distributed neural architectures* comprise a set of non-exclusive, overlapping nodes for representing each concept [60, 50, 27].

The work of researchers such as Feldman [16, 17], Ajjanagadde and Shastri [67], Sun [72], and Smolensky [69] has demonstrated why localist connectionist networks are suitable for implementing symbolic processes usually associated with higher cognitive functions. On the other hand, “radical connectionism” [13] is a distributed neural approach to modeling intelligence. Usually, it is easier to incorporate prior knowledge into localist models since their structures can be made to directly correspond to that of symbolic knowledge [19]. On the other hand, neural learning usually leads to distributed representation. Furthermore there has been work on integrating localist and distributed representations [28, 72, 87].

2.2 Transformation Architectures

Hybrid transformation architectures transform symbolic representations into neural representations or vice versa. The main processing is performed by neural representations but there are automatic procedures for transferring neural representations to symbolic representations or vice versa. Using a transformation architecture it is possible to insert or extract symbolic knowledge into or from a neural architecture. Hybrid transformation architectures differ from unified neural architectures by the automatic transfer. While certain units in unified neural architectures may be interpreted symbolically by an observer, hybrid transformation architectures actually allow the knowledge transfer into symbolic rules, symbolic automata, grammars, etc.

Examples of such transformation architectures include the work on activation-based automata extraction from recurrent networks [54, 90]. Alternatively, a weight-based transformation between symbolic rules and feedforward networks has been extensively examined in knowledge-based artificial neural networks [68, 20].

The most common transformation architectures are rule extraction architectures where symbolic rules are extracted from neural networks [19, 1]. These architectures have received a lot of attention since rule extraction discovers the hyperplane positions of units in neural networks and transforms them to if-then-else rules. Rule extraction has been performed mostly with multi-layer perceptron networks [79, 5, 8, 11], Kohonen networks, radial basis functions [2, 33] and recurrent networks [53, 90]. Extraction of symbolic knowledge from neural networks has also played an important aspect in this current volume, e.g. [81, 7, 84]. Furthermore, insertion of symbolic knowledge can be either gradual through practice [23] or one-shot.

2.3 Hybrid Modular Architectures

Hybrid modular architectures contain both symbolic and neural modules appropriate to the task. Here, symbolic representations are not just initial or final representations as in a transformation architecture. Rather, they are combined

and integrated with neural representations in many different ways. Examples in this class, for instance, contain CONSYDERR [72], SCREEN [95] or robot navigators where sensors and neural processing are fused with symbolic top-down expectations [37]. A variety of distinctions can be made. Neural and symbolic modules in hybrid modular architectures can be loosely coupled, tightly coupled or completely integrated [48].

Loosely Coupled Architectures A *loosely coupled hybrid architecture* has separate symbolic and neural modules. The control flow is sequential in the sense that processing has to be finished in one module before the next module can begin. Only one module is active at any time, and the communication between modules is unidirectional.

There are several loosely coupled hybrid modular architectures for semantic analysis of database queries [9] or dialog processing [34] or simulated navigation [78]. Another example of a loosely coupled architecture has been described in a model for structural parsing [87] combining a chart parser and feedforward networks. Other examples of loose coupling, which is sometimes also called passive coupling, include [45, 36].

In general, this loose coupling enables various loose forms of cooperation among modules [73]. One form of coupling is in terms of pre/postprocessing vs. main processing: while one or more modules take care of pre/postprocessing, such as transforming input data or rectifying output data, a main module focuses on the main part of the processing task. Commonly, while pre/post processing is done using a neural network, the main task is accomplished through the use of symbolic methods. Another form of cooperation is through a master-slave relationship: while one module maintains control of the task at hand, it can signal other modules to handle some specific aspects of the task. Yet another form of cooperation is the equal partnership of multiple modules.

Tightly Coupled Architectures A *tightly coupled hybrid architecture* contains separate symbolic and neural modules where control and communication are via common shared internal data structures in each module. The main difference between loosely and tightly coupled hybrid architectures are common data structures which allow bidirectional exchanges of knowledge between two or more modules. This makes communication faster and more active but also more difficult to control. Therefore, tightly coupled hybrid architectures have also been referred to as actively coupled hybrid architectures [47].

As examples of tightly coupled architectures, systems for neural deterministic parsing [41] and inferencing [28] have been built where the control changes between symbolic marker passing and neural similarity determination. Furthermore, a hybrid system developed by Tirri [83] consists of a rule base, a fact base and a neural network of several trained radial basis function networks [57, 59].

In general, a tightly coupled hybrid architecture allows multiple exchanges of knowledge between two or more modules. The result of a neural module can have a direct influence on a symbolic module or vice versa before it finishes its global

processing. For instance, CDP is a system for deterministic parsing [41], SCAN contains a tightly coupled component for structural processing and semantic classification [87]. While the neural network chooses which action to perform, the symbolic module carries out the action. During the process of parsing, control is switched back and forth between these modules. Other tightly coupled hybrid architectures for structural processing have been described in more detail in [89]. CLARION is also a system that couples symbolic and neural representations to explore their synergy.

Fully Integrated Architectures In a *fully integrated hybrid architecture* there is no discernible external difference between symbolic and neural modules, since the modules have the same interface and they are embedded in the same architecture. The control flow may be parallel. Communication may be bidirectional between many modules, although not all possible communication channels have to be used.

One example of an integrated hybrid architecture is SCREEN, which was developed for exploring integrated hybrid processing for spontaneous language analysis [95, 92]. In fully integrated and interleaved systems, the constituent modules interact through multiple channels (e.g., various possible function calls), or may even have node-to-node connections across two modules, such as CONSYDERR [72] in which each node in one module is connected to a corresponding node in the other module. Another hybrid system designed by Lees et al [43] interleaves case-based reasoning modules with several neural network modules.

3 Directions for Hybrid Neural Systems

In Feldman and Bailey’s paper, it was proposed that there are the following distinct levels [15]: cognitive linguistic level, computational level, structured connectionist level, computational biology level and biological level. A condition for this *vertical hybridization* is that it should be possible to bridge the different levels, and the higher levels should be reduced to, or grounded in, lower levels. A top-down research methodology is advocated and examined for concepts towards a neural theory of language.

Although the particulars of this approach are not universally agreed upon, researchers generally accept the overall idea of multiple levels of neural cognitive modeling. In this view, models should be constructed entirely of neural components; both symbolic and subsymbolic processes should be implemented in neural networks.

Another view, *horizontal hybridization*, argues that it may be beneficial, and sometimes crucial, to “mix” levels so that we can make better progress on understanding cognition. This latter view is based on realistic assessment of the state of the art of neural model development, and the need to focus on the essential issues (such as the synergy between symbolic and subsymbolic processes [78]) rather than nonessential details of implementation. Horizontal approaches have been used successfully for real-world hybrid systems, for instance in speech/language

analysis [95]. Purely neural systems in vertical hybridization are more attractive for neuroscience but hybrid systems of horizontal hybridization are currently also a tractable way of building large-scale hybrid neural systems.

Representation, learning and their interaction represent some of the major issues for developing symbol processing neural networks. Neural networks designed for symbolic processing often involve complex internal structures consisting of multiple components and several different representations [67, 71, 3]. Thus learning is made more difficult. There is a need to address the problems of what type of representation to adopt, how the representational structure in such systems is built up, how the learning processes involved affect the representation acquired and how the representational constraints may facilitate or hamper learning.

In terms of what is being learned in hybrid neural systems, we can have (1) learning contents for a fixed architecture, (2) learning architectures for given contents, or (3) we can learn both contents and architecture at the same time. Although most hybrid neural learning systems fall within the first two categories, e.g. [18, 46], there are some hybrid models that belong to the third category, e.g. [50, 92].

Furthermore, there is some current work on parallel neural and symbolic learning, which includes using (1) two separate neural/symbolic algorithms applied simultaneously [78], (2) two separate algorithms applied in succession, (3) integrated neural/symbolic learning [80, 35], and (4) purely neural learning of symbolic knowledge, e.g. [46, 51].

The issues described above are important for making progress in theories and applications of hybrid systems. Currently, there is not yet a theory of “hybrid systems”. There has been some preliminary early work towards a theoretical framework for neural/symbolic representations, but to date there is still a lack of an overall theoretical framework that abstracts away from the details of particular applications, tasks and domains. One step towards such a direction may be the research into the relationship between automata theory and neural representations [39, 24, 88].

Processing natural language has been and will continue to be a very important test area for exploring hybrid neural architectures. It has been argued that “language is the quintessential feature of human intelligence” [85]. While certain learning and architectures in humans may be innate, most researchers in neural networks argue for the importance of development and environment during language learning [87, 94]. For instance, it was argued [51] that syntax is not innate and that it is a process rather than representation, and abstract categories, like subject, can be learned bottom-up.

The dynamics of learning natural language is also important for designing parsers using techniques like SRN and RAAM. SARDSRN and SARDRAAM were presented in the context of shift-reduce parsing [46] to avoid the problem associated with SRN and RAAM (that is, losing constituent information). Interestingly, it has been argued that compositionality and systematicity in neural networks arise from an associationistic substrate [61] based on principles from evolution.

Also, research into improving WWW use by using neural networks may be promising [93]. While currently most search engines only employ fairly traditional search strategies, machine learning and neural networks could improve processing of heterogeneous unstructured multimedia data.

Another important promising research area is knowledge extraction from neural networks in order to support text mining and information retrieval [81]. Inductive learning techniques from neural networks and symbolic machine learning algorithms could be combined to analyze the underlying rules for such data.

A crucial task for applying neural systems, especially for applying learning distributed systems, is the design of appropriate vector representations for scaling up to real-world tasks. Large context vectors are also essential for learning document retrieval [22]. Due to the size of the data, only linear computations are useful for full-scale information retrieval. However, vector representations are still often restricted to co-occurrences, rather than focusing on syntax, discourse, logic and so on [22]. However, complex representations may be formed and analyzed using fractal approaches [82].

Hard real-world applications are important. A system was built for foreign exchange rate prediction that uses a SOM for reduction and that generates a symbolic representation as input for a recurrent network which can produce rules [55]. Another self-organizing approach for symbol processing was described for classifying Usenet texts and presenting the classification as a hierarchical two-dimensional map [32]. Related neural classification work for text routing has been described [93]. Neural network representations have also been used for important parts of vision and association [56].

Finally, there is promising progress in neuroscience. Computational neuroscience is still in its infancy but it may be very relevant to the long-term progress of hybrid symbolic neural systems. Related to that, more complex high order neurons may be one possibility for building more powerful functionality [44]. Another way would be to focus more on global brain architectures, for instance for building biological inspired robots with rooted cognition [64].

It was argued [85] that in 20 years computer power will be sufficient to match human capabilities, at least in principle. But meaning and deep understanding are still lacking. Other important issues are perception, situation assessment and action [78], although perceptual pattern recognition is still in a very primitive state. Rich perception also requires links with rich sets of actions. Furthermore, it has been argued that language is the “quintessential feature” of human intelligence [85] since it is involved in many intelligent cognitive processes.

4 Concluding Remarks

In summary, further work towards a theory and fundamental principles of hybrid neural systems is needed. First of all, there is promising work towards relating automata theory with neural networks, or logics with such networks. Furthermore, the issue of representation needs more focus. In order to tackle larger real world tasks using neural networks, for instance in information retrieval, learning

internet agents, or large-scale classification, further research on the underlying vector representations for neural networks is important. Vertical forms of neural/symbolic hybridization models are widely used in cognitive processing, logic representation and language processing. Horizontal forms of neural/symbolic hybridization exist for larger tasks, such as speech/language integration, knowledge engineering, intelligent agents or condition monitoring. Furthermore, it will be interesting to see in the future to what extent computational neuroscience will offer further ideas and constraints for building more sophisticated forms of neural systems.

References

1. R. Andrews, J. Diederich, and A. B. Tickle. A survey and critique of techniques for extracting rules from trained artificial networks. Technical report, Queensland University of Technology, 1995.
2. R. Andrews and S. Geva. Rules and local function networks. In *Proceedings of the Rule Extraction From Trained Artificial Neural Networks Workshop, Artificial Intelligence and Simulation of Behaviour*, Brighton UK, 1996.
3. J. Barnden. Complex symbol-processing in Conposit. In R. Sun and L. Bookman, editors, *Architectures incorporating neural and symbolic processes*. Kluwer, Boston, 1994.
4. J. A. Barnden and K. J. Holyoak, editors. *Advances in connectionist and neural computation theory*, volume 3. Ablex Publishing Corporation, 1994.
5. J. Benitz, J. Castro, and J. I. Requena. Are artificial neural networks black boxes? *IEEE Transactions on Neural Networks*, 8(5):1156–1164, 1997.
6. R. Bogacz and C. Giraud-Carrier. A novel modular neural architecture for rule-based and similarity-based reasoning. In *Hybrid Neural Systems (this volume)*. Springer-Verlag, 2000.
7. G. Bologna. Symbolic rule extraction form the DIMLP neural network. In *Hybrid Neural Systems (this volume)*. Springer-Verlag, 2000.
8. G. Bologna and C. Pellegrini. Accurate decomposition of standard MLP classification responses into symbolic rules. In *International Work Conference on Artificial and Natural Neural Networks, IWANN'97*, pages 616–627, Lanzarote, Canaries, 1997.
9. Y. Cheng, P. Fortier, and Y. Normandin. A system integrating connectionist and symbolic approaches for spoken language understanding. In *Proceedings of the International Conference on Spoken Language Processing*, pages 1511–1514, Yokohama, 1994.
10. P. S. Churchland and T. J. Sejnowski. *The Computational Brain*. MIT Press, Cambridge, MA, 1992.
11. T. Corbett-Clarke and L. Tarassenko. A principled framework and technique for rule extraction from multi-layer perceptrons. In *Proceedings of the 5th International Conference on Artificial Neural Networks*, pages 233–238, Cambridge, England, July 1997.
12. J. Diederich and D. L. Long. Efficient question answering in a hybrid system. In *Proceedings of the International Joint Conference on Neural Networks*, Singapore, 1992.
13. G. Dorffner. *Neural Networks and a New AI*. Chapman and Hall, London, UK, 1997.

14. M. A. Fanty. Learning in structured connectionist networks. Technical Report 252, University of Rochester, Rochester, NY, 1988.
15. J. Feldman and D. Bailey. Layered hybrid connectionist models for cognitive science. In *Hybrid Neural Systems (this volume)*. Springer-Verlag, 2000.
16. J. A. Feldman and D. H. Ballard. Connectionist models and their properties. *Cognitive Science*, 6:205–254, 1982.
17. J. A. Feldman, G. Lakoff, D. R. Bailey, S. Narayanan, T. Regier, and A. Stolcke. L_0 - the first five years of an automated language acquisition project. *AI Review*, 8, 1996.
18. P. Frasconi, M. Gori, and A. Sperduti. Integration of graphical rules with adaptive learning of structured information. In *Hybrid Neural Systems (this volume)*. Springer-Verlag, 2000.
19. L.M. Fu. Rule learning by searching on adapted nets. In *Proceedings of the National Conference on Artificial Intelligence*, pages 590–595, 1991.
20. L.M. Fu. *Neural Networks in Computer Intelligence*. McGraw-Hill, Inc., New York, NY, 1994.
21. S. I. Gallant. *Neural Network Learning and Expert Systems*. MIT Press, Cambridge, MA, 1993.
22. S. I. Gallant. Context vectors: a step toward a grand unified representation. In *Hybrid Neural Systems (this volume)*. Springer-Verlag, 2000.
23. J. Gelfand, D. Handleman, and S. Lane. Integrating knowledge-based systems and neural networks for robotic skill. In *Proceedings of the International Joint Conference on Artificial Intelligence*, pages 193–198, San Mateo, CA., 1989.
24. L. Giles and C. W. Omlin. Extraction, insertion and refinement of symbolic rules in dynamically driven recurrent neural networks. *Connection Science*, 5:307–337, 1993.
25. S. Goonatilake and S. Khebbal. *Intelligent Hybrid Systems*. Wiley, Chichester, 1995.
26. N. A. Hallack, G. Zaverucha, and V. C. Barbosa. Towards a hybrid model of first-order theory refinement. In *Hybrid Neural Systems (this volume)*. Springer-Verlag, 2000.
27. J. A. Hammerton and B. L. Kalman. Holistic symbol computation and the sequential RAAM: An evaluation. In *Hybrid Neural Systems (this volume)*. Springer-Verlag, 2000.
28. J. Hendler. Developing hybrid symbolic/connectionist models. In J. A. Barnden and J. B. Pollack, editors, *Advances in Connectionist and Neural Computation Theory, Vol.1: High Level Connectionist Models*, pages 165–179. Ablex Publishing Corporation, Norwood, NJ, 1991.
29. M. Hilario. An overview of strategies for neurosymbolic integration. In *Proceedings of the Workshop on Connectionist-Symbolic Integration: From Unified to Hybrid Approaches*, pages 1–6, Montreal, 1995.
30. S. Hölldobler. A structured connectionist unification algorithm. In *Proceedings of the National Conference of the American Association on Artificial Intelligence 90*, pages 587–593, Boston, MA, 1990.
31. S. Hölldobler, Y. Kalinke, and J. Wunderlich. A recursive neural network for reflexive reasoning. In *Hybrid Neural Systems (this volume)*. Springer-Verlag, 2000.
32. S. Honkela. Self-organizing maps in symbol processing. In *Hybrid Neural Systems (this volume)*. Springer-Verlag, 2000.
33. J. S. R. Jang and C. T. Sun. Functional equivalence between radial basis function networks and fuzzy inference systems. *IEEE Transactions Neural Networks*, 4(1):156–159, 1993.

34. D. Jurafsky, C. Wooters, G. Tajchman, J. Segal, A. Stolcke, E. Fosler, and N. Morgan. The Berkeley Restaurant Project. In *Proceedings of the International Conference on Speech and Language Processing*, pages 2139–2142, Yokohama, 1994.
35. P. Kanerva. Large patterns make great symbols: an example of learning from example. In *Hybrid Neural Systems (this volume)*. Springer-Verlag, 2000.
36. C. Kirkham and T. Harris. Development of a hybrid neural network/expert system for machine health monitoring. In R. Rao, editor, *Proceedings of the 8th International Congress on Condition Monitoring and Engineering Management, COMADEM95*, pages 55–60, 1995.
37. G.K. Kraetzschmar, S. Sablatnoeg, S. Enderle, and G. Palm. Application of neurosymbolic integration for environment modelling in mobile robots. In *Hybrid Neural Systems (this volume)*. Springer-Verlag, 2000.
38. S. C. Kremer. A theory of grammatical induction in the connectionist paradigm. Technical Report PhD dissertation, Dept. of Computing Science, University of Alberta, Edmonton, 1996.
39. S.C. Kremer and J. Kolen. Dynamical recurrent networks for sequential data processing. In *Hybrid Neural Systems (this volume)*. Springer-Verlag, 2000.
40. F. Kurfeß. Unification on a connectionist simulator. In T. Kohonen, K. Mäkisara, O. Simula, and J. Kangas, editors, *Artificial Neural Networks*, pages 471–476. North-Holland, 1991.
41. S. C. Kwasny and K. A. Faisal. Connectionism and determinism in a syntactic parser. In N. Sharkey, editor, *Connectionist natural language processing*, pages 119–162. Lawrence Erlbaum, Hillsdale, NJ, 1992.
42. T. Lange and M. Dyer. High-level inferencing in a connectionist network. *Connection Science*, 1:181–217, 1989.
43. B. Lees, B. Kumar, A. Mathew, J. Corchado, B. Sinha, and R. Pedreschi. A hybrid case-based neural network approach to scientific and engineering data analysis. In *Proceedings of the Eighteenth Annual International Conference of the British Computer Society Specialist Group on Expert Systems*, pages 245–260, Cambridge, 1998.
44. H. Lipson and H.T. Siegelmann. High order eigentensors as symbolic rules in competitive learning. In *Hybrid Neural Systems (this volume)*. Springer-Verlag, 2000.
45. J. MacIntyre and P. Smith. Application of hybrid systems in the power industry. In L. Medsker, editor, *Intelligent Hybrid Systems*, pages 57–74. Kluwer Academic Press, 1995.
46. M.R. Mayberry and R. Miikkulainen. Combining maps and distributed representations for shift-reduce parsing. In *Hybrid Neural Systems (this volume)*. Springer-Verlag, 2000.
47. K. McGarry, S. Wermter, and J. MacIntyre. Hybrid neural systems: from simple coupling to fully integrated neural networks. *Neural Computing Surveys*, 2:62–94, 1999.
48. L. R. Medsker. *Hybrid Neural Network and Expert Systems*. Kluwer Academic Publishers, Boston, 1994.
49. L. R. Medsker. *Hybrid Intelligent Systems*. Kluwer Academic Publishers, Boston, 1995.
50. R. Miikkulainen. *Subsymbolic Natural Language Processing*. MIT Press, Cambridge, MA, 1993.
51. W. C. Morris, G. W. Cottrell, and J. L. Elman. A connectionist simulation of the empirical acquisition of grammatical relations. In *Hybrid Neural Systems (this volume)*. Springer-Verlag, 2000.

52. M. C. Mozer. Neural net architectures for temporal sequence processing. In A. Weigend and N. Gershenfeld, editors, *Time series prediction: Forecasting the future and understanding the past*, pages 243–264. Addison-Wesley, Redwood City, CA, 1993.
53. C. W. Omlin and C. L. Giles. Extraction and insertion of symbolic information in recurrent neural networks. In V. Honavar and L. Uhr, editors, *Artificial Intelligence and Neural Networks: Steps Towards principled Integration*, pages 271–299. Academic Press, San Diego, 1994.
54. C. W. Omlin and C. L. Giles. Extraction of rules from discrete-time recurrent neural networks. *Neural Networks*, 9(1):41–52, 1996.
55. C.W. Omlin, L. Giles, and K. K. Thornber. Fuzzy knowledge and recurrent neural networks: A dynamical systems perspective. In *Hybrid Neural Systems (this volume)*. Springer-Verlag, 2000.
56. C. Orovas and J. Austin. A cellular neural associative array for symbolic vision. In *Hybrid Neural Systems (this volume)*. Springer-Verlag, 2000.
57. J. Park and I. W. Sandberg. Universal approximation using radial basis function networks. *Neural Computation*, 3:246–257, 1991.
58. N. S. Park. Addressing knowledge representation issues in connectionist symbolic rule encoding for general inference. In *Hybrid Neural Systems (this volume)*. Springer-Verlag, 2000.
59. T. Peterson and R. Sun. An RBF network alternative for a hybrid architecture. In *International Joint Conference on Neural Networks*, Anchorage, AK, May 1998.
60. J. B. Pollack. Recursive distributed representations. *Artificial Intelligence*, 46:77–105, 1990.
61. R. Reilly. Evolution of symbolisation: Signposts to a bridge between connectionist and symbolic systems. In *Hybrid Neural Systems (this volume)*. Springer-Verlag, 2000.
62. R. G. Reilly and N. E. Sharkey. *Connectionist Approaches to Natural Language Processing*. Lawrence Erlbaum Associates, Hillsdale, NJ, 1992.
63. D. E. Rumelhart, G. E. Hinton, and R. J. Williams. Learning internal representations by error propagation. In D. E. Rumelhart and J. L. McClelland, editors, *Parallel Distributed Processing*, volume 1, pages 318–362. MIT Press, Cambridge, MA, 1986.
64. N. Sharkey and N. T. Ziemke. Life, mind and robots: The ins and outs of embodied cognition. In *Hybrid Neural Systems (this volume)*. Springer-Verlag, 2000.
65. L. Shastri. A model of rapid memory formation in the hippocampal system. In *Proceedings of the Meeting of the Cognitive Science Society*, pages 680–685, Stanford, 1997.
66. L. Shastri. Types and quantifiers in SHRUTI: a connectionist model of rapid reasoning and relational processing. In *Hybrid Neural Systems (this volume)*. Springer-Verlag, 2000.
67. L. Shastri and V. Ajanagadde. From simple associations to systematic reasoning: A connectionist representation of rules, variables and dynamic bindings. *Behavioral and Brain Sciences*, 16(3):417–94, 1993.
68. J. Shavlik. A framework for combining symbolic and neural learning. In V. Honavar and L. Uhr, editors, *Artificial Intelligence and Neural Networks: Steps towards principled Integration*, pages 561–580. Academic Press, San Diego, 1994.
69. P. Smolensky. On the proper treatment of connectionism. *Behavioral and Brain Sciences*, 11(1):1–74, March 1988.

70. A. Sperduti, A. Starita, and C. Goller. Learning distributed representations for the classifications of terms. In *Proceedings of the International Joint Conference on Artificial Intelligence*, pages 494–515, Montreal, 1995.
71. R. Sun. On variable binding in connectionist networks. *Connection Science*, 4(2):93–124, 1992.
72. R. Sun. *Integrating Rules and Connectionism for Robust Commonsense Reasoning*. Wiley, New York, 1994.
73. R. Sun. Hybrid connectionist-symbolic models: A report from the IJCAI95 workshop on connectionist-symbolic integration. *Artificial Intelligence Magazine*, 1996.
74. R. Sun. Supplementing neural reinforcement learning with symbolic methods. In *Hybrid Neural Systems (this volume)*. Springer-Verlag, 2000.
75. R. Sun and F. Alexandre. *Proceedings of the Workshop on Connectionist-Symbolic Integration: From Unified to Hybrid Approaches*. McGraw-Hill, Inc., Montreal, 1995.
76. R. Sun and F. Alexandre. *Connectionist Symbolic Integration*. Lawrence Erlbaum Associates, Hillsdale, NJ, 1997.
77. R. Sun and L.A. Bookman. *Computational Architectures Integrating Neural and Symbolic Processes*. Kluwer Academic Publishers, Boston, MA, 1995.
78. R. Sun and T. Peterson. Autonomous learning of sequential tasks: experiments and analyses. *IEEE Transactions on Neural Networks*, 9(6):1217–1234, 1998.
79. S. Thrun. Extracting rules from artificial neural networks with distributed representations. In G. Tesauro, D. Touretzky, and T. Leen, editors, *Advances in Neural Information Processing Systems 7*. MIT Press, San Mateo, CA, 1995.
80. S. Thrun. *Explanation-Based Neural Network Learning*. Kluwer, Boston, 1996.
81. A. Tickle, F. Maire, G. Bologna, R. Andrews, and J. Diederich. Lessons from past, current issues and future research directions in extracting the knowledge embedded in artificial neural networks. In *Hybrid Neural Systems (this volume)*. Springer-Verlag, 2000.
82. P. Tino, G. Dorffner, and C. Schittkopf. Understanding state space organization in recurrent neural networks with iterative function systems dynamics. In *Hybrid Neural Systems (this volume)*. Springer-Verlag, 2000.
83. H. Tirri. Replacing the pattern matcher of an expert system with a neural network. In S. Goonatilake and S. Khebbal, editors, *Intelligent Hybrid Systems*, pages 47–62. John Wiley and Sons, 1995.
84. M.L. Vaughn, S.J. Cavill, S.J. Taylor, M.A. Foy, and A.J.B. Fogg. Direct knowledge extraction and interpretation from a multilayer perceptron network that performs low back pain classification. In *Hybrid Neural Systems (this volume)*. Springer-Verlag, 2000.
85. D. Waltz. The importance of importance. In *Presentation at Workshop on Hybrid Neural Symbolic Integration*, Breckenridge, CO., 1998.
86. D. L. Waltz and J. A. Feldman. *Connectionist Models and their Implications*. Ablex, 1988.
87. S. Wermter. *Hybrid Connectionist Natural Language Processing*. Chapman and Hall, Thomson International, London, UK, 1995.
88. S. Wermter. Preference Moore machines for neural fuzzy integration. In *Proceedings of the International Joint Conference on Artificial Intelligence*, pages 840–845, Stockholm, 1999.
89. S. Wermter. The hybrid approach to artificial neural network-based language processing. In R. Dale, H. Moisl, and H. Somers, editors, *A Handbook of Natural Language Processing*. Marcel Dekker, 2000.

90. S. Wermter. Knowledge extraction from transducer neural networks. *Applied Intelligence: The International Journal of Artificial Intelligence, Neural Networks, and Complex Problem-Solving Techniques*, 12:27–42, 2000.
91. S. Wermter, G. Arevian, and C. Panchev. Towards hybrid neural learning internet agents. In *Hybrid Neural Systems (this volume)*. Springer-Verlag, 2000.
92. S. Wermter and M. Meurer. Building lexical representations dynamically using artificial neural networks. In *Proceedings of the International Conference of the Cognitive Science Society*, pages 802–807, Stanford, 1997.
93. S. Wermter, C. Panchev, and G. Arevian. Hybrid neural plausibility networks for news agents. In *Proceedings of the National Conference on Artificial Intelligence*, pages 93–98, Orlando, USA, 1999.
94. S. Wermter, E. Riloff, and G. Scheler. *Connectionist, Statistical and Symbolic Approaches to Learning for Natural Language Processing*. Springer, Berlin, 1996.
95. S. Wermter and V. Weber. SCREEN: Learning a flat syntactic and semantic spoken language analysis using artificial neural networks. *Journal of Artificial Intelligence Research*, 6(1):35–85, 1997.