Neurosymbolic Integration: Cognitive Grounds and Computational Strategies

Yannick Lallement (1), Mélanie Hilario (2), Frédéric Alexandre (1)

(1) CRIN-INRIA Lorraine
BP 239 - Campus scientifique
F-54506 Vandœuvre-les-Nancy Cedex
Fax: (+33) 83 41 30 79
lallemen@loria.fr, falex@loria.fr

(2) CUI - University of Geneva
24 rue Général-Dufour
CH-1211 Geneva 4
Fax: (+41) 22 320 29 27
hilario@cui.unige.ch

Abstract

The ultimate—if implicit—goal of artificial intelligence (AI) research is to model the full range of human cognitive capabilities. Symbolic AI and connectionism, the major AI paradigms, have each tried—and failed—to attain this goal. In the meantime, the idea has gained ground that this goal might still be within reach if we could harness the respective strengths of these two paradigms in integrated neurosymbolic models. This paper attempts to lay a cognitive basis for neurosymbolic integration and describes the different strategies that have been adopted to date. Unified approaches strive to attain symbol-processing capabilities using neural network techniques alone, while hybrid approaches blend symbolic and neural models in novel architectures with the hope of gleaning the best of both paradigms.

Keywords: Connectionism, symbolic AI, neurosymbolic integration, hybrid models, connectionist symbol processing

1 Introduction

Since its inception, artificial intelligence (AI) has been torn between two major paradigms—symbolic AI and connectionist AI. Each AI researcher has had to take a stand, at least implicitly, in the symbolic/connectionist debate. It is widely recognized (see, for example, [Gal88] and [Gut92]) that these two paradigms have different and complementary strengths. Connectionist models often offer good performance in pattern recognition and generalization, and present such qualities as natural learning ability, noise tolerance and graceful degradation. By contrast, symbolic models often present a complementary profile: they offer good performance in reasoning and deduction, and present such qualities as natural symbolic manipulation and explanation abilities. In this paper, we discuss the relationship between the various AI models and the human cognitive spectrum. The final goal of AI is clearly to cover the whole range of human cognitive capabilities, and attempts to combine symbolic and connectionist processing have given rise to a major research area, that of neurosymbolic integration (NSI).

In the following section, we develop a perspective on AI which grounds neurosymbolic integration as an attempt to cover the whole range of the human cognitive spectrum. We then distinguish and describe two major strategies for NSI—unified (Section 3) and hybrid (Section 4)—before concluding on the results obtained by the different approaches as well as their potential impact.

2 A cognitive view on AI

A cursory analysis of human cognition reveals that there are several types of mental processes and each of them plays a role in human intelligence. In particular, it is useful to distinguish between *low-level* and *high-level* operations. Low-level operations are mostly concerned with perceptual tasks (such as vision, speech recognition, locomotion), whereas high-level operations are mostly concerned with reasoning tasks that involve volition and deliberation (such as problem solving and decision-making). Newell [New90] has classified these operations according to their duration. This distinction between low-level and high-level tasks/operations also maps directly onto Steel's distinction between behavior-based and knowledge-based tasks [Ste89]. Since an operation is seldom fully low-level or fully high-level, we propose to draw a continuum, named *cognitive spectrum* between these two extremities. Actually, most human complex cognitive operations involve *both* high- and low-level operations: for example, driving a car comprises low-level tasks (e.g., perception of the road, of motor noise) and high-level tasks (e.g., deciding whether to turn left or right), and thus takes up a large segment on the cognitive spectrum.

Figure 1: The cognitive spectrum

According to conventional wisdom in AI, the distinction between low-level and high-level tasks/operations maps quite easily—albeit superficially—onto the distinction between the symbolic and the connectionist paradigms. It is widely recognized (see, e.g., [Gal88] and [Gut92]) that these two paradigms have different and complementary strengths. For example, connectionist models often offer good performance in pattern recognition and generalization, and present such qualities as natural learning ability, noise tolerance and graceful degradation. By contrast, symbolic models often present a complementary profile: they offer good performance in reasoning and deduction, and present such qualities as natural symbolic manipulation and explanation abilities. The most salient of these properties are summarized in table 1.

	Symbolic pole	Connectionist pole
Learning	Difficult	Easy
Explanation	Easy	Difficult
Performance degradation	Fast	Slow
Noise Resistance	Weak	Good
Data structure representation	Easy	Difficult
Algorithmics	Sequential	Parallel

Table 1: Characteristic properties of the two poles

Due to their respective strengths and weaknesses, connectionist models seem to be more successful at low-level tasks, whereas symbolic models are apparently better adapted to high-level operations. However, the goal of AI is clearly to integrate the whole range of human cognitive capabilites. To achieve this goal, future models will have to benefit from the advantages and capacities of both the symbolic and the connectionist paradigms. This explains current intense activity in the field of neurosymbolic integration. In order to situate the neurosymbolic models among existing models, we propose to draw a new continuum, called the *AI spectrum* (figure 2) between the extreme poles representing purely symbolic and purely connectionist models. Neurosymbolic models will find their place on the AI spectrum between the two poles, closer to one or the other according to their dominant set of properties.

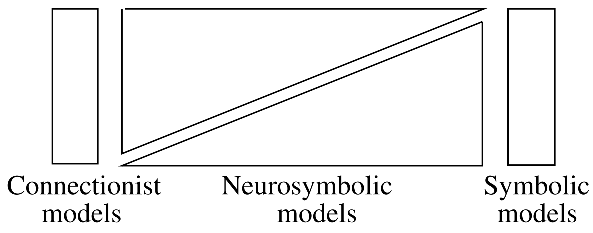

Figure 2: The AI spectrum

Roughly, two basic strategies can be adopted to build neurosymbolic models. The *unified* approach aims at incorporating neural and symbolic properties using purely connectionist techniques. This approach leads to the elaboration of models where symbol-processing functionalities emerge from neural structures and processes. Alternatively, the *hybrid* approach consists in blending neural and symbolic models. The resulting models are both symbolic and connectionist, and are expected to inherit the advantages of their parent models while doing away as much as possible with the shortcomings of each. The next two sections delve more deeply into these strategies.

3 Unified approaches

Unified neurosymbolic models aim at performing all cognitive tasks using the neural paradigm. Two trends can be distinguished within the unified neurosymbolic camp: the neuron-to-symbol approach and the symbol-to-neuron approach.

3.1 The neuron-to-symbol approach

The main goal of this approach is neurophysiological plausibility: as it name implies, it grounds all cognitive—including symbol-processing—capabilities on the *biological* neuron. Thus the neuron-to-symbol (or neuron-based) approach is a biologically based, bottom-up approach: it is perhaps the one AI strategy (whether symbolic or connectionist, whether integrative or not) which can be rightly described as "modelling the brain" instead of "making a mind" [DD88]. This approach is still in its infancy, and though successes have been scored at the perceptual level, it will take some time before a frontal connectionist attack on symbol processing can be envisaged. The models of this approach are very deeply anchored in perception and action on the external world. They aim at building a representation from these low-level cues that could permit interpretation, reasoning or planning.

3.1.1 Cortical Columns

The cortical column model has been proposed as a basic functional unit of cortex-like networks from a biological point of view [GDBS93a]. It has also been implemented and tested from a computational point of view [ABGH91]. A similar approach is proposed by Edelman [Ede87].

At the microscopic level, the cortical column is repeated throughout the cortex and is organized in a micro-circuit of neurons performing specific operations (transfer function tuning, modulation gating, temporal integration) on specific data. These data originate from the connectivity of this unit. Indeed, it differs from the classical formal neuron in that sense that its functioning depends on the kind of information that is received through typed links. This information can correspond to lateral inhibition, transfer of learning, coupling with similar areas or with areas from different perceptual or motor modalities.

At the macroscopic level, cortical columns functionally gather in the cortex and constitute cortical areas. Four kinds of areas are distinguished and named to fit biological observations: the sensory areas; the motor areas; the associative areas, whose function is to connect the sensory and motor areas and thus to couple the various representations of the external world; the frontal areas which are sensible to the various temporal scales needed by a process. From the sensory and associative context provided by the other three areas, and thanks to the temporal dimension, action plans are built in the frontal areas. Interconnection between these areas are not random, but correspond to specific information flows. For example, vision and motion are linked through parietal and temporal associative areas to allow execution of tasks such as visually guided grasping.

Within this framework, a set of applications have been carried out to assess the abilities of the model, including visual [AGH90] and speech processing [GADH90], but also more symbolic or reasoning oriented tasks [GAH90, Bla94, GDBS93b].

3.1.2 Perceptually grounded language learning

Katamic memories [ND93] and the DETE model [ND94] also correspond to more deeply biologically inspired neural models. These perceptually based models permit acquisition of natural language semantics from interactions with the real world. The idea is to propose sensory motor experiences together with verbal utterances describing those experiences, in order to enable the system to build relationships between the two representations.

Katamic memories are cerebellar-like networks for sequence learning. Here too, the basic unit reveals more complex behavior than the classical formal neuron. We also find in this approach this novel trend in biologically inspired neural networks that consists in building models able to deal with sequence representation and manipulation [DA94].

The DETE system uses a large set of Katamic memories performing associative tasks from verbal to motor or from visual to motor pole. Those associations then enable to build verbal sequences describing a visual representation.

For the moment, these models have only been tested on restricted applications and need to be confirmed. Nevertheless, they are very promising in the sense that they are directly inspired by human-like tasks, from low-level perceptual or motor tasks up to higher-level cognitive behavior.

3.2 The symbol-to-neuron approach

In this approach, the main emphasis is on building connectionist architectures for symbol processing. Typically, model construction starts with an idea of some high-level symbolic function to be performed, and proceeds with the implementation of the appropriate neural architecture. The symbol-to-neuron (or symbol-based) approach is thus a top-down strategy; it should also be added that the neuron in question here, contrary to that of the neuron-based approach, is usually the classical formal neuron. The symbol-based approach, more widely known as *connectionist symbol processing*, lays no claims to biological modelling. (However, the term "connectionist symbol processing" is not precise enough for our present purposes since the biological approach also aims at modelling symbol-processing skills). Most of these models lack perceptual influence. Their basic units generally represent some rather high-level symbols, on which classical AI methods are applied. Their human-like functioning should certainly depend on a more low-level perceptual implantation, even if they also need improvements from a cognitive point of view.

Depending on the underlying knowledge representation scheme, symbol-to-neuron architectures can be classified as localist (each node in the system corresponds to a single concept, and vice-versa), distributed (each concept is encoded using several nodes, and each node participates in the representation of several concepts), or combined localist/distributed.

3.2.1 Localist models

A connectionist architecture for rule-based reasoning and variable binding is described in [AS91]. Knowledge is represented in a localist fashion using several types of nodes: predicate nodes (for predicate names), role nodes (for predicate arguments or variables), and filler nodes (for possible role values). Rules are encoded by interconnection patterns between predicate and role nodes. Variables and constants are bound if their associated role and filler nodes fire in synchrony as activation spreads throughout the network. This idea of using synchronous or in-phase activation to represent the binding of distinct features and concepts is supported by neurophysiological evidence. However, the model has significant difficulties in scaling up to complex, high-level reasoning; more importantly, it fails to give any convincing account of the learning process. A later version of the system [Ajj94] uses a single type of nodes in the inference network, thus facilitating the process of learning.

3.2.2 Distributed models

Smolensky [Smo90] proposes a representation scheme based on tensor products. This system can represent data structures (such as strings, sentences or trees). The structures have to be decomposed into a set of roles (in a sentence for example, the roles can be *word-1, word-2...* and the the fillers will be the words themselves). The bindings of the roles to the filler fixes the structure. Roles and fillers can be represented by neurons in a connectionist network. Their representation can be local (one neuron for each role and each unit), or distributed (a role or a unit are represented by a pattern of activations). The bindings are coded in a set of binding neurons: there is one binding neuron for each filler neuron/role neuron pair. Connectionist methods are used to set and retrieve the bindings. This framework is general enough to represent complex (possibly recursive) data structures and allow classical operations (e.g. a la Lisp). However, learning procedures associated with usual connectionist networks are lost in this case.

BOLTZCONS [Tou90] uses distributed representations of linked lists to implement symbolic structures like trees and stacks. One cell of a linked list is encoded as a triple of symbols of form `(tag, car, cdr)`. Since memory is assumed to be sparse, only a fraction of these triples can be present in memory at any one time. A coarse coding technique, the description of which is beyond the scope of this paper, reduces the number of units required to store a triple while adding a measure of representational redundancy. Pointer traversal is implemented via associative retrieval, and a connectionist maintenance system adds, deletes and updates memory triples. Fault tolerance and robustness are side benefits of this distributed representation: like biological networks, BOLTZCONS will not crash abruptly if some cells cease to function.

3.2.3 Localist/distributed models

CONSYDERR [Sun91] is a unified system aimed at modelling common-sense reasoning (e.g., *Can a duck fly? A duck is like a sparrow, so I guess it can fly.*). It uses a dual but purely connectionist representation scheme. At the conceptual level, domain concepts are encoded as nodes in a localist network and rules are represented by links between these nodes. At the subconceptual level, a distribution network performs finer-grained and less structured processing, i.e., associations and pattern matching. Each concept known to the system is represented in the two networks. The localist network performs classical rule-based inferencing while the distributed network performs similarity-based reasoning. With these two levels, CONSYDERR can handle a number of difficult issues, such as partial or inexact information, in one integrated framework.

4 Hybrid approaches

Hybrid neurosymbolic systems can be distinguished along different dimensions such as their target problem or task domain, the degree of coupling of the symbolic and neural components (loose or tight), the reasoning and representational models used in the symbolic (e.g., rule-based reasoning, case-based reasoning) and neural (e.g., multilayer perceptrons, Kohonen networks) components, or the function filled by the neural and symbolic modules with regard to each other and to the overall system. We propose a classification based on the integration scheme, i.e., the way in which the symbolic and neural components are configured in relation to each other and to the overall system. We distinguish four integration modes: chainprocessing, subprocessing, metaprocessing and coprocessing.

4.1 Chainprocessing

In chainprocessing mode, one of the modules—either symbolic or neural—is the main processor while the other takes charge of pre and/or postprocessing tasks. In [HCK92], for instance, the main processor is a respiratory monitoring expert system which recommends actions to be taken to avoid breathing complications. Rules are of the form: "If *qualitative-state* then *action*", where *qualitative-state* is a symbolic representation of a change in lung pressure over time (e.g., "pressure is rising rapidly"). Experts found it extremely difficult to formulate rules to determine these qualitative states, so a backpropagation network was used to accomplish this task; its output was then stored as a fact in the expert system's working memory. With neural preprocessing, development time was cut down from 3 to 2 months and accuracy rose from 74.5 % to 97.5 %.

The reverse setup is that of a symbolic preprocessor assisting a connectionist main processor. An example is the use of decision trees as feature selectors to limit the number of input nodes in feedforward neural networks [PS94]. On the PROMOTER database, 2 seconds of neural preprocessing accelerated convergence of the backpropagation learning process by a factor of 4 while improving accuracy on both the training and the test sets.

4.2 Subprocessing

In subprocessing mode, one of the two modules is embedded in and subordinated to the other, which acts the main problem solver. In examples found so far, the symbolic main processor delegates subtasks to a neural subprocessor; it remains an open question whether the reverse setup—a connectionist main processor subcontracting tasks to a symbolic subprocessor—is at all possible.

In INNATE/QUALMS [BLN91], the symbolic module is the main processor whose task is fault diagnosis in a distillation plant. It is a classical expert system linked to a simulation model of the process to be monitored. To diagnose a process fault, the expert system subcontracts the task of generating a set of candidate faults to the neural subprocessor, composed of several multilayer perceptrons cascaded in two layers. Each output neuron is associated with a particular fault and its activation is interpreted as the rating of the corresponding fault. The expert system then either confirms the neural diagnosis or proposes another hypothesis.

LAM™, a hybrid system for window glazing design [Med94], is another example of neural subprocessing. Its principal role is to serve as a design assistant: relying mainly on rulebased modules, it classifies glass types, checks for errors or improbable glass designs, and helps a user take design decisions such choosing the glass type, and then determining such properties as glass strength, solar specifications or sound control class. Design rules for properties such as structural strength of glass were easily constructed from architectural manufacturer specifications; however, two design tasks—selecting glass solar control and sound control properties of glass—are difficult to express in the form of situation/action rules since they involve estimating complex correlations between input and output parameters. At the same time, there exists a sizeable set of training/test data which accurately correlates known inputs with output predictions. Two three-layer feedforward networks where therefore trained to determine the appropriate solar and sound properties on the basis of other glass properties; during an interactive consultation, these two specific tasks are subcontracted to the neural network modules while other design phases are taken care of by a base of 578 if/then rules.

A third example of neural subprocessing is the hybrid composed of SCRAPS and a neural network that learns by backpropagation [Hen89]. The main processor is a marker-passing semantic net which performs complex reasoning tasks like planning. A shortcoming of this symbolic problem solver is its inability to recognize useful perceptual similarities between objects; for example, it cannot infer that a letter-opener can be used as a crime weapon on the basis of perceptual features similar to those of a knife. To remedy this shortcoming, the semantic net is connected to a neural network which has been trained to recognize microfeatures, defined as features which have little meaning on their own but which taken as a set represent a concept. As a result, activation spreads back from perceptual features identified by the neural net to objects in the semantic net which share these features. The particularity of this system is that communication between the two systems is smoother than a simple parameterized function call, since there is a direct connection between the leaf nodes of the semantic net (those representing individual instances) and the output nodes of the neural network which encodes the microfeatures as a result of prior training.

In short, this system is an example of a tightly coupled hybrid, in contrast to all the other hybrids discussed in this paper, which can be considered loosely coupled insofar as they communicate via function call parameters instead of sharing data structures.

4.3 Metaprocessing

In metaprocessing mode, one module is the baselevel problem solver and the other plays a metalevel role (such as monitoring, control, or performance improvement) vis-a-vis the first.

Symbolic metaprocessing is illustrated in the Robotic Skill Acquisition Architecture [HLG92]. Its goal is to develop robots which accomplish complex tasks using designer-supplied instructions and self-induced practice. The architecture's base level is itself hybrid: a rule-base provides a declarative representation of human expert knowledge whereas neural networks embody reflexive procedural knowledge that comes with practice. The metaprocessor is a rule-based execution monitor which supervises the training of the neural network and controls the operation of the system during the learning process.

An example of *neural metaprocessing* is Kwasny and Faisal's natural language parser [KF92]. Classical parsing systems map an input sentence into a parse tree by maintaining an lookahead buffer and an internal stack. Tree-building actions (e.g., creation of a new node) are taken by a set of rules after examining both buffer and stack. The main drawback of these systems is that they can only parse sentences whose form has been anticipated in the ruleset; but natural language is too rich to be encoded in such a set. In the proposed hybrid parser, symbolic parsing rules are replaced by a feedforward network which has been trained on a set of sentences generated by a deterministic grammar. The symbolic module manages the buffer and stack and codes their state as inputs for the network, which determines the action to be taken on a best-match basis. With this function replacement, the hybrid system is able to recognize grammatical as well as non-grammatical sentences, whether encountered previously or not.

Another example of neural metaprocessing is a system in which the baselevel symbolic processor solves high school physics problems, and the connectionist component enforces search control [GP90]. Object-level rules encode equations such as velocity = initial-velocity + acceleration x time. If several of the needed variables are unknown, the metalevel connectionist module is called on to make a choice. This is a previously trained net whose inputs are the current goal variable as well as the known and unknown variables; it outputs the next (unknown) variable to solve for, in effect guiding the problem-solving process.

4.4 Coprocessing

In coprocessing, the symbolic and the neural modules are equal partners in the problem solving process: each can interact directly with the environment, each can transmit information to and receive information from the other. Neurosymbolic coprocessing can be either competitive or cooperative. In the competitive subclass, a supervisor assigns a common task to a symbolic and a connectionist module, then selects the best result on the basis of some predefined criterion. The cooperative subclass is an implementation of the divide-and-conquer strategy: a supervisor divides the task into subtasks, then assigns each subtask to a symbolic or a connectionist module in accordance with such guidelines as the nature of the subtask and the appropriateness of the module's computational capacities given the subtask.

Cooperative neurosymbolic coprocessing is illustrated in a system which combines decision trees and neural nets to improve arrhythmia diagnosis [JPL+92]. Intracardial defibrillators, devices implanted in people with heart disorders, sense the electrical activity of the heart and identify abnormal rhythms or "arrhythmias", which can be clustered into three main groups depending on the type of action they call for: continue monitoring, pace the heart, or apply high-voltage electric shock. Classification accuracy is, of course, crucial in this application. For dual chamber classification, a hybrid model was devised: the input data—sample signals from ventricular and atrial probes—are channelled to a decision tree which acts as a timing classifier and to a multilayer perceptron which performs morphology-based classification. The outputs of both modules are fed into an arbitrator which determines the class of the arrhythmia. The decision tree/MLP hybrid attained an accuracy rate of 99% on a multi-patient database whereas the best performance of a simple MLP 89.3% and that of a multimodule neural net was 96.2%.

5 Conclusion

During AI's early history, both the symbolic and the connectionist paradigms aimed at reproducing the full range of human cognitive capabilities. It is now clear that each alone has failed to achieve this goal. However, the two paradigms provided good and certainly re-usable results so far. Their integration into neuro-symbolic models is certainly a key issue for the future of AI and cognitive modeling.

The main research issue met by the hybrid approach is to find ways of integration and communication between models from different origins [HLA95]. The unified models have to get to a middle sub-symbolic level that seems to be especially significant in cognition [Hof85]. To achieve this, the symbol-to-neuron and the neuron-to-symbol approaches have to go further in their respectively top-down and bottom-up ways.

Results obtained so far by these integrated strategies are inversely proportional to the difficulty of their respective tasks. The hybrid approach has led to more convincing results so far [HPA94]: some industrial applications are already operational while unified models are still far from achieving the representational power of symbolic systems. However, considering biological evidence, unified approaches—in particular the neuron-to-symbol approach—appear to be very promising. The hybrid approach may be viewed as a short-term engineering expedient, whereas the unified approaches might well lead to fundamental breakthroughs, but it will certainly take more time before they bear fruit.

References

[ABGH91] F. Alexandre, Y. Burnod, F. Guyot, and J-P. Haton. The cortical column: a new processing unit for multilayered networks. *Neural Networks*, 4(1), 1991.

[AGH90] F. Alexandre, F. Guyot, and J-P. Haton. A connectionist network with two complementary visual processing systems for x-ray image interpretation. In *International Neural Networks Conference*, 1990.

[Ajj94] V. Ajjanagadde. Unclear distinctions lead to unnecessary shortcomings: Examining the rule vs fact, role vs filler, and type vs predicate distinctions from a connectionist representation and reasoning perspective. In *American Association Artificial Intelligence*, 1994.

[AS91] V. Ajjanagadde and L. Shastri. Rules and variables in neural nets. *Neural Computation*, (3):121–134, 1991.

[Bla94] P. Blanchet. An architecture for representing and learning behaviors by trial and error. In *Symposium on Adaptive Behaviour*, Brighton, 1994.

[BLN91] W. R. Becraft, P. L. Lee, and R. B. Newell. Integration of neural networks and expert systems for process fault diagnosis. In *Proc. of the 12th IJCAI*, pages 832–837, 1991.

[DA94] S. Durand and F. Alexandre. A neural network based on sequence learning: Application to spoken digits recognition. In *7th international conference on Neural Networks and Their Applications*, pages 290–298, Marseille, 1994.

[DD88] H. Dreyfus and S. Dreyfus. Making a mind versus modeling the brain: Artificial intelligence at a branchpoint. *Daedalus*, Winter, 1988.

[Ede87] G. Edelman. *Neural Darwinism: the theory of neural group selection.* Basic Books, 1987.

[GADH90] F. Guyot, F. Alexandre, C. Dingeon, and J-P. Haton. The cortical column as a model for speech recognition: Principles and first experiments. In *Speech Recognition and Understanding Recent Advances, Trends and Applications*. Springer Verlag, 1990.

[GAH90] F. Guyot, F. Alexandre, and J-P. Haton. Principles and applications of the cortical column symbolic neural model. In *International Joint Conference on Neural Networks*, San Diego, 1990.

[Gal88] S. I. Gallant. Connectionist expert systems. *Communications of the ACM*, 31(2), February 1988.

[GDBS93a] E. Guigon, B. Dorizzi, Y. Burnod, and W. Schultz. Neural correlates of learning in the prefrontal cortex: a predictive model. *Cerebral Cortex*, 1993.

[GDBS93b] E. Guigon, B. Dorizzi, Y. Burnod, and W. Schultz. Neural correlates of learning in the prefrontal cortex: a predictive model. *Cerebral Cortex*, 1993.

[GP90] Matthias Gutknecht and Rolf Pfeifer. An approach to integrating expert systems with connectionist networks. *AICOM*, 3:116–127, 1990.

[Gut92] M. Gutknecht. The postmodern mind: hybrid models of cognition. *Connection Science*, 4(3 & 4), 1992.

[HCK92] S. Hayes, V.B. Ciesielski, and W. Kelly. A comparison of an expert system and a neural network for respiratory system monitoring. Technical Report TR #92/1, Royal Melbourne Institute of Technology, 1992.

[Hen89] J. Hendler. Marker-passing over microfeatures : Towards a hybrid symbolic/connectionist model. *Cognitive Science*, (13):79–106, 1989.

[HLA95] M. Hilario, Y. Lallement, and F. Alexandre. Neurosymbolic integration: Unified versus hybrid approaches. In *The European Symposium On Artificial Neural Networks*, Brussels, Belgium, 1995.

[HLG92] D.A. Handelman, S.H. Lane, and J.J. Gelfand. Robotic skill acquisition based on biological principles. In A. Kandel and G. Langholz, editors, *Hybrid Architectures for Intelligent Systems*, pages 301–327. CRC Press, 1992.

[Hof85] D. Hofstadter. *Metamagical themas: questing for the essence of mind and pattern.* Basic Books, 1985.

[HPA94] M. Hilario, C. Pellegrini, and F. Alexandre. Modular Integration of Connectionist and Symbolic Processing in Knowledge-Based Systems. In *Proceedings International Symposium on Integrating Knowledge and Neural Heuristics*, Pensacola Beach (Florida, USA), May 1994.

[JPL+92] M. Jabri, S. Pickard, P. Leong, Z. Chi, B. Flower, and Y. Xie. ANN-based classification for heart difibrillators. In *Advances in Neural Information Processing 4*, pages 637–644. Morgan-Kaufmann, 1992.

[KF92] S.C. Kwasny and K.A. Faisal. Symbolic parsing via subsymbolic rules. In J. Dinsmore, editor, *The Symbolic and Connectionist Paradigms: Closing the Gap*, pages 209–235. Lawrence Erlbaum, 1992.

[Med94] Larry R. Medsker. *Hybrid Neural Network and Expert Systems.* Kluwer Academic Publishers, Boston, 1994.

[ND93] V. Nenov and M. Dyer. Perceptually grounded language learning: Part 1 - a neural network architecture for robust sequence association. *Connection Science*, 5(2), 1993.

[ND94] V. Nenov and M. Dyer. Perceptually grounded language learning: Part 2 - dete: a neural/procedural model. *Connection Science*, 6(1), 1994.

[New90] Alan Newell. *Unified theories of cognition.* Harvard University Press, Cambridge, Mass., 1990.

[PS94] S. Piramuthu and M.J. Shaw. On using decision tree as feature selector for feed-forward neural networks. In *International Symposium on Integrating Knowledge and Neural Heuristics*, pages 67–74, 1994.

[Smo90] P. Smolensky. Tensor product variable binding and the representation of symbolic structures in connectionist systems. *Artificial Intelligence*, (46):159–216, 1990.

[Ste89] L. Steels. Connectionist problem solving—an ai perspective. In R. Pfeifer, Z. Schreter, and F. Fogelman-Soulie, editors, *Connectionism in Perspective*, pages 215–228. Elsevier, 1989.

[Sun91] R. Sun. *Integrating Rules and Connectionism for Robust Reasoning : A Connectionist Architecture with Dual Representation.* PhD thesis, Brandeis University, Waltham, MA 02254, 1991.

[Tou90] D. S. Touretzky. Boltzcons : Dynamic symbol structures in a connectionist network. *Artificial Intelligence*, 46(1-2), 1990.