

---

 WORKSHOP REPORT
 

---

# High-Level Connectionist Models

*Jordan Pollack*

---

*A workshop on high-level connectionist models was held in Las Cruces, New Mexico, on 9-11 April 1988 with support from the American Association for Artificial Intelligence and the Office of Naval Research. John Barnden and Jordan Pollack organized and hosted the workshop and will edit a book containing the proceedings and commentary. The book will be published by Ablex as the first volume in a series entitled Advances in Connectionist and Neural Computation Theory.*

---

The purpose of this workshop was to bring together a small group of computer scientists to focus on the interaction between AI and connectionism. The two fields are often posed as paradigmatic enemies, and a risk of severing them exists. (See Dreyfus and Dreyfus 1988 and other papers in the recent issue of *Daedalus*.) Few connectionist results are published in the mainstream AI journals and conference proceedings other than those sponsored by the Cognitive Science Society, and many neural-network researchers and industrialists proceed without consideration of the problems (and progress) of AI.

The workshop was assembled by invitation, and we strove for both geographic and ideological breadth. The questions posed to the participants involved the application of connectionism to traditionally symbolic tasks, approaches to overcoming some of the limitations exhibited by connectionist models of the past several years, and opportunities for connectionism to make contributions to the study of cognition.

The participants were Wendy Lehnert, Michael Dyer, and James Hendler (testers of the connectionist waters); Garrison Cottrell and Jeffrey Elman (representatives of the hard-line Parallel Distributed Processing school); Lokendra Shastri and Joachim Diederich (the Feldman school); Mark Derthick and David Touretzky (the Hinton school); Lawrence Bookman (the Waltz school); and Lawrence Birnbaum (the loyal opposition). In addition, Jordan Pollack, John Barnden, and Roger Schvaneveldt represented their own school, New Mexico State University.

Participants were given exactly 1 hour to present their views and man-

age the oftentimes vociferous audience. The 14 presentations were squeezed into 2 1/2 days, including long, social evenings. The workshop ended with a visit to White Sands National Monument. The following is a summary of each person's presentation (in alphabetical order), then a discussion of some of the major themes that emerged from the workshop.

## Presentation Summaries

John Barnden, (1986, 1987, 1988), New Mexico State University (NMSU), described the current status and future plans for his CONPOSIT system. Implemented on NASA's Massively Parallel Processor, the system makes use of a two-dimensional array of registers constructed of simple neuron-like processors; it can represent complex symbolic structures using adjacency between registers as a form of variable binding. Barnden demonstrated how production rules are currently implemented and timed and proposed how a large set of such rules could be managed. He also showed that his system could reason syllogistically, directly implementing Johnson-Laird's notion of a mental model (Johnson-Laird and Bara 1984).

Lawrence Birnbaum, Yale University, spoke about hard problems in planning and understanding that would not "just go away" for connectionists. He gave several examples of abstract schemata that require long chains of inferences in order to be recognized and situations which require such recognition in order to be understood (Birnbaum 1986). One example, a variant of one of Eugene Charniak's (1983) marker-passing stories, involved a detective following a suspect into a hardware store. After being spotted by

the suspect, the detective buys a rope. Another story (hypothetical at the time) had Michael Dukakis worried about winning the New York primary without alienating the Black vote. Then, Albert Gore attacks Jesse Jackson's views on the Middle East, and Dukakis carries the state. The concepts of looking inconspicuous and helping (by acting as a shill) seem to involve reasoning explicitly about the causal chains representing aspects of these situations—in other words, about the understander's own internal state—and thus strong self-referential representational systems, where connectionism flounders.

Lawrence Bookman, Brandeis University, argued that connectionism needed better forms of modularity in order to scale up. He described his work on network regions (Chun, Bookman, and Afshartous 1987), which are a hierarchical abstraction for spreading activation networks. A network region is an abstract unit standing for a collection of units, whose input, activation level, and output are averages over the corresponding values for these units. He showed how interacting clusters of winner-take-all networks used for parsing and lexical disambiguation, such as those used by Waltz and Pollack (1985), could be cleanly implemented in his system. A long discussion ensued about such supernodes and the actual cost of their implementation.

Garrison Cottrell, the University of California at San Diego (UCSD), described his recent work on learning to compress gray-scale images (Cottrell, Munro, and Zipser 1987). This system learned a linear transformation of a well-known algorithm called principal components analysis. Cottrell then described some extensions of this work, including the use of the internal representations as a basis for semantics for words and various projects involving the use and representation of image sequences.

Mark Derthick (1987), Carnegie-Mellon University (CMU), described the genesis of his system called  $\mu$ -KLONE (Micro-KL-one), which is an energy-minimization approach to generating a finite model of a first-order logic theory. He argued that his approach was better than a theorem

prover because it degrades more gracefully when resource limitations prevent running to completion. His system also provides a formal probabilistic theory for commonsense reasoning in the face of a knowledge base that might be logically incomplete or inconsistent. He described the simplifications and modifications necessary to avoid intractability in his system.

Joachim Diederich (1986, 1987), the International Computer Science Institute (ICSI) in Berkeley, California, described his research on spreading activation and various forms of network propagation (1986, 1987, 1988). He contrasted pulse-specific discrete marker passing with source specific analog spreading activation in several network taxonomies, and the problems for each. Using psychological evidence, he argued for a form of inhibition (Renshaw) over the usual winner-take-all lateral inhibition as both less resource intensive (requiring  $O(n)$  rather than  $O(n^2)$  links) and a better account of classification data from humans under stress. Renshaw inhibition is one of the most important inhibition patterns in animal nervous systems, and Diederich presented an application of this inhibition that avoids false classification with the use of intermediate units between concept units and property units in a spreading activation network. If two or more concepts share a property, and both concepts are source units for the spreading activation process, only the strongest concept unit and the particular property unit remain active simultaneously. He also described the organization and work in progress of ICSI, which is headed by Jerry Feldman.

Michael Dyer, the University of California at Los Angeles, laid out his general philosophy on symbolic neuro-engineering, through which traditional AI tasks are decomposed into modules that can be implemented using connectionist techniques, yielding more robust and flexible systems. He described a number of research projects he and his colleagues were engaged in. He demonstrated how an extended back-propagation scheme discovered representations of lexical entities while the network was performing McClelland and Kawamoto's

(1986) case-frame mapping task (Miikkulainen and Dyer 1988). As a result, a microfeature specification for the input is no longer necessary. Dyer also described DUAL (Dyer, Flowers, and Wang 1988), a PDP architecture able to represent labeled, directed graphs with cycles. The general technique manipulates the entire weight matrix formed by one network as a pattern of activation in a larger network. As a semantic network is encoded in the PDP architecture, distributed representations are formed for each node in the semantic network, with structurally similar nodes developing similar patterns of activation. Dyer only briefly hinted at work on (1) the association of visual motion with linguistic descriptions, which he referred to as the grounding problem (Nenov and Dyer 1988); (2) variable binding techniques using conjunctive coding (Dolan and Dyer 1987, 1988); (3) localist architectures for dynamic word-sense reinterpretation (Sumida, Dyer, and Flowers 1988); and (4) a model of completely parallel language generation (Gasser and Dyer 1988).

Jeffrey Elman (1988), UCSD, presented his latest work on learning to process and predict sequences. He uses a recurrent form of back propagation (Rumelhart, Hinton, and Williams 1986), where the output of units in a hidden layer are directly fed back into input units. Using long input sequences, Elman trains a network in prediction tasks. For example, one task was to predict the next word in a long concatenated sequence of short, grammatically generated sentences. In learning to do so, the network devised a classification for the words (into noun or verb, animate or inanimate, and so on) that was revealed through cluster analysis. Another network learned to resolve pronouns in a large body of grammatically generated variable-length sentences.

James Hendler (1987), the University of Maryland, discussed his work on a hybrid model of marker passing that descends into the subsymbolic using microfeatures to mediate similarity. The problem was to recognize in a planning domain that a letter opener is similar enough to a knife to set off an alarm at an airport security gate. He showed how this task could be accom-

plished, without a direct link to WEAPON, by spreading through microfeatures such as METAL, POINTY, and so on. He proposed the next generation of his system, which will develop the microfeatures automatically using back propagation. Hendler also took on the whole philosophical question of AI versus connectionism. Two interesting analogies were discussed. Geoffrey Hinton was quoted as having compared AI to zoology (descriptive) and connectionism to molecular biology (constructive). Hendler pointed out the appropriateness of this mapping but cautioned that cognition is much more like evolution than autopsy. John McCarthy was quoted as having said that AI and connectionism were two horses in the same race and should be permitted to run the course; Hendler asserted there is really only one horse but with two ends.

Wendy Lehnert (1987), the University of Massachusetts, made a strong case for hybrid modeling, taking the best of both worlds for building useful systems. She described her new sentence analysis system, which uses a stack and a copy mechanism for control, marker passing for prediction, and numeric spreading activation for smooth decision making. She showed how the system could give a unified syntactic and semantic account for parsing and could flexibly implement a no-crossing of branches constraint for prepositional phrase attachment.

Jordan Pollack, NMSU, argued that multiplicative connections would lead to more powerful models and that stability and complexity problems could be handled by programming conventions. Pollack described two such conventions: cascaded networks (1987) and fractional data flow networks. He also described recent work on devising compositional distributed representations for variable-sized stacks and trees using a recursive form of auto association (Pollack 1988). He argued that in addition to overcoming major complaints about the adequacy of connectionist representations such as those leveled by Fodor and Pylyshyn (1988), the representations developed by his system could be the basis for systems that perform inference using fixed-width pattern association.

Roger Schvaneveldt, NMSU, discussed his attempt to replicate the schema system of Rumelhart, et. al. (1986) using a matrix of content features that represent prototypical rooms. He compared and contrasted the PDP dense relaxation system with hierarchical clustering techniques and his PATHFINDER graph-theoretic algorithm (Schvaneveldt, Dearholt, and Durso, Forthcoming).

Lokendra Shastri (1988), the University of Pennsylvania, discussed what contribution connectionism could make to AI. He described his work on knowledge representation and reasoning in the general framework of discovering fast and tractable means of performing inference. Shastri stressed that robust intelligent behavior requires an extremely fine-grained decomposition of knowledge and that connectionism forces us to discover such decompositions. One success is the complete compilation of default reasoning using property inheritance and classification into a spreading activation connectionist network.

David Touretzky, CMU, raised a warning flag about the abundance of trivial models coming out of connectionism. He said that certain ideas had already reached their useful limit and could stagnate the field. Such ideas include the continued attempts to reduce cognition to lateral inhibition among localist units, rote associative memory with no inferring ability, and pattern transformation by three-layer back-propagation nets. He presented several projects that he has been working on, including the DUCS system for representing frames (Touretzky and Geva 1987), a system for prepositional phrase attachment, and a new ambitious project on connectionist metonymic reasoning.

## Discussion and Common Themes

Participants presented their results and views and were constantly interrupted by criticisms and tangential discussions. However, some recurring themes did emerge from this chaos.

### Pokey Connectionist Models

Originally Birnbaum's colorful complaint, the idea of new connectionist implementations of procedures or tasks that standard AI programming can handle easily, such as CONSing, towers of Hanoi, syntactic parsers, production systems, logical operations, and so on, generally met with negative audience reactions. Such implementations are only interesting when they demonstrate useful capabilities beyond simple symbolic adequacy.

### Hybrid Modeling

Both Wendy Lehnert and James Hendler argued strongly in favor of mixing symbolic and connectionist methodologies. Although such work proceeds and is successful, questions about the utility and epistemology of such models remain. A general lack of mutual understanding existed about the notion of discrete abstract levels in cognitive models or of compilation between levels. A good question to ask of a hybrid model is if it truly bridges levels from presymbolic to postsymbolic processing. Dyer argued that work should proceed on at least four levels—knowledge, spreading activation semantic network, connectionist, and neural—with an eye toward the key questions of how each level is justified by the one above and how each level can be embedded into the one below.

### Moving Target Learning

In most systems using back propagation, the training environment consists of a stable set of input and target patterns. An interesting commonality emerged in the work of Dyer, Elman, and Pollack: In each of their systems, which use various recurrent forms of back propagation, the learning environment changes dynamically. This change occurs because some of the input or target patterns are representations being developed by hidden units, which change along with the weights in the network. This moving target strategy could turn out to be quite important because it is a major facet of some of the first successful connectionist attempts to encode and process dynamically sized data structures, such as sequences, trees, and labeled graphs.

## Connectionism

### Might Redefine Symbol

In a lively discussion period, Schvaneveldt brought up the whole question of the meaning and use of symbols. In AI, symbols have no internal structure and, thus, mean little; they are simply used as names or pointers to larger structures of symbols, which are reasoned with (slowly). The subsymbolic distinction was considered incoherent (and Smolensky was not present to defend it) because just about every connectionist model discussed was symbolic in nature, even at the so-called microfeature level. The essential difference between the early neural network research and modern connectionism is that AI has happened in between. Some suggestion was made that a new kind of symbol ("symboid," according to Dyer) might emerge from connectionism. For example, a reduced representation could be considered such a symboid, given that it can point to a larger structure through a reconstruction algorithm. Symboids might have an advantage over symbols in that they possess internal structure which can be reasoned about.

### Hard Problems Remain Hard

A definite consensus was that although a change in computational methodology might make some problems easier for connectionist or symbolic processing, the hard problems are not going to go away; thus intelligence will not be solved by either three-layer back propagation or physicists studying dynamical systems. These hard problems include the organization and encoding of knowledge, inference, the management of the dynamic instantiation of structures, and rapid (one-shot) learning. Natural language, planning, and reasoning all require infinite generative capacity and the ability to combine old knowledge in new ways. To approach these problems, connectionists need to continue work on complex representations and the effective use of modularity and hierarchy.

## Conclusion

Given the diverse backgrounds and methodologies of the participants and

the brief and intense nature of the workshop, it cannot be said that true consensus was reached or that new scientific ground was broken. The participants will undoubtedly continue on their own paths, influenced, perhaps, by the common experience.

However, the workshop must be considered a great success from the point of view of communication. We placed the participants in close quarters for an extended period of time, and all survived the heated argument, criticism, slicing, and dicing. It is safe to say that research in connectionist approaches to higher-level cognitive functions will continue for some time.

### Acknowledgments

The author acknowledges the American Association for Artificial Intelligence and the Office of Naval Research (grant N00014-88-J-1066) for funding this workshop, and John Barnden for co-organizing it. For help with transporting the visiting dignitaries and running the audiovisual equipment, thanks go to Michael Angerman, Imre Balogh, Srinii Kankanahalli, and Tony Plate; for budget and administrative work, to Jeannine Sandefur and Pat Puckett; for detailed information and numerous suggestions aiding in the preparation of this report, to Michael Dyer; and for hosting the reception at their house, to Roger and Ann Schvaneveldt.

### References

Barnden, J. A. 1988. The Right to Free Association: Relative-Position Encoding for Connectionist Data-Structures. In Proceedings of the Tenth Annual Conference of the Cognitive Science Society, 503-509. Hillsdale, N.J.: Lawrence Erlbaum.

Barnden, J. A. 1987. Simulation of an Array-Based Neural Net Model. In Proceedings of the First Symposium on the Frontiers of Massively Parallel Scientific Computation.

Barnden, J. A. 1986. Complex Cognitive Information Processing: A Computational Architecture with a Connectionist Implementation, Technical Report, 211, Computer Science Dept., Indiana Univ.

Birnbaum, L. 1986. Integrated Processing in Planning and Understanding. Research Report 489. New Haven: Yale University Computer Science Department.

Charniak, E. 1983. Passing Markers: A Theory of Contextual Influence in Language.

*Cognitive Science* 7: 171-190.

Chun, H. W.; Bookman, L. A.; and Afshartous, N. 1987. Network Regions: Alternatives to the Winner-Take-All Structure. In Proceedings of the Tenth International Joint Conference on Artificial Intelligence, 380-387. Menlo Park, Calif.: International Joint Conferences on Artificial Intelligence.

Cottrell, G.; Munro, P.; and Zipser, D. 1987. Learning Internal Representations from Gray-Scales Images: An Example of Extensional Programming. In Proceedings of the Seventh Annual Conference of the Cognitive Science Society, 461-473. Hillsdale, N.J.: Lawrence Erlbaum.

Derthick, M. 1987. A Connectionist Architecture for Representing and Reasoning about Structured Knowledge. In Proceedings of the Ninth Annual Conference of the Cognitive Science Society, 131-142. Hillsdale, N.J.: Lawrence Erlbaum.

Diederich, J. 1987. *Simulation Schizophrener Sprache*. Opladen: Westdeutscher Verlag.

Diederich, J. 1986. Kognitive Parallelverarbeitung, GMD Report, 159, R. Oldenbourg Verlag.

Dolan, C., and Dyer, M. G. 1988. Parallel Retrieval and Application of Conceptual Knowledge, Technical Report, UCLA-AI-88-3, Artificial Intelligence Lab., Univ. of California at Los Angeles.

Dolan, C., and Dyer, M. G. 1987. Symbolic Schemata, Role Binding and the Evolution of Structure in Connectionist Memories. In Proceedings of the First International Conference on Neural Networks. Washington D.C.: Institute of Electrical and Electronics Engineers, vol 2, 287-298.

Dreyfus, H. L., and Dreyfus, S. E. 1988. Making a Mind versus Modeling the Brain: Artificial Intelligence Again at the Crossroads. *Daedalus* 117(1):15-43.

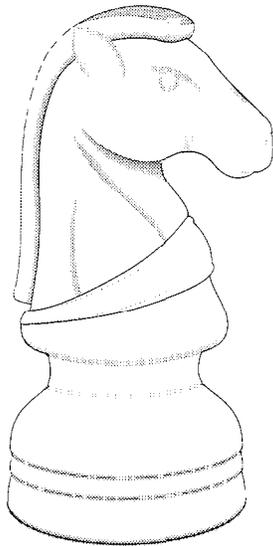
Dyer, M. G.; Flowers, M.; and Wang, Y. A. 1988. Weight Matrix = Pattern of Activation: Encoding Semantic Networks as Distributed Representations in DUAL, a PDP Architecture, Technical Report, UCLA-AI-88-5, Artificial Intelligence Lab., Univ. of California at Los Angeles.

Elman, J. L. 1988. Finding Structure in Time, Technical Report, 8801, Center for Research in Language, Univ. of California at San Diego.

Fodor, J., and Pylyshyn, A. 1988. Connectionism and Cognitive Architecture: A Critical Analysis. *Cognition* 28: 3-71.

Gasser, M., and Dyer, M. G. 1988. Sequencing in a Connectionist Model of Language Processing. In Proceedings of the Twelfth International Conference on Computation-

# Have KEE® — Will Travel



**KEE is a registered trademark  
of IntelliCorp**

## **RAPID PROTOTYPING of Robust, Extensible Knowledge-Based Systems**

**AMOS OSHRIN**

**Knowledge Systems Consultant  
4 Morning Sun Court • Mountain View, CA 94043**

**(415) 969-9600**

### **RESULTS, Not Just Reports**

For free information, circle no. 145

al Linguistics, Budapest, Hungary.

Hendler, J. A. 1987. Extending Marker-Passing to Microfeatures: A Cautious Step towards the Subsymbolic. In *Advances in Cognitive Science* vol. 2, ed. N. Sharkey. Norwood, NJ: Ablex.

Johnson-Laird, P. N., and Bara, B. G. 1984. Syllogistic Inference. *Cognition* 16: 1-61.

Lehnert, W. G. 1987. Case-Based Problem-Solving with a Large Knowledge Base of Learned Cases. In Proceedings of the Sixth National Conference on Artificial Intelligence, 301-306. Menlo Park, Calif.: American Association for Artificial Intelligence.

McClelland, J., and Kawamoto, A. 1986. Mechanisms of Sentence Processing: Assigning Roles to Constituents. In *Parallel Distributed Processing: Experiments in the Microstructure of Cognition*, vol. 2, 272-325. Eds. J. L. McClelland, D. E. Rumelhart, and the PDP Research Group. Cambridge, Mass.: MIT Press.

Miikkulainen, R., and Dyer, D. G. 1988. Forming Global Representations with Back Propagation. In Proceedings of the Second Annual International Conference on Neural Networks, vol. 1, 285-292. Washington D.C.: Institute of Electrical and Electronics Engineers.

Nenov, V. I., and Dyer, M. G. 1988. DETE:

Connectionist/Symbolic Model of Visual and Verbal Association, vol. 2, 17-24. In Proceedings of the Second Annual International Conference on Neural Networks. Washington D.C.: Institute of Electrical and Electronics Engineers.

Pollack, J. B. 1988. Recursive Auto-Associative Memory: Devising Compositional Distributed Representations. In Proceedings of the Tenth Annual Conference of the Cognitive Science Society, 33-39. Hillsdale, N.J.: Lawrence Erlbaum.

Pollack, J. B. 1987. Cascaded Back Propagation on Dynamic Connectionist Networks. In Proceedings of the Ninth Conference of the Cognitive Science Society, 391-404. Hillsdale, N.J.: Lawrence Erlbaum.

Rumelhart, D. E.; Hinton, G.; and Williams, R. 1986. Learning Internal Representations through Error Propagation. In *Parallel Distributed Processing: Experiments in the Microstructure of Cognition*, vol. 1, eds. D. E. Rumelhart, J. L. McClelland, and the PDP Research Group, 318-362. Cambridge, Mass.: MIT Press.

Rumelhart, D. E.; Smolensky, P.; McClelland, J. L.; and Hinton, G. E. 1986. Schemata and Sequential Thought Processes in PDP Models. In *Parallel Distributed Processing: Experiments in the Microstructure of Cognition*, vol. 2, eds. D. E. Rumelhart, J.

L. McClelland, and the PDP Research Group, 7-57. Cambridge, Mass.: MIT Press.

Schvaneveldt, R. W.; Dearholt, D. W.; and Durso, F. T. Graph Theoretic Foundations of Pathfinder Networks. In *Computers and Mathematics with Applications*. In press.

Shastri, L. 1988. Semantic Nets: An Evidential Formalization and Its Connectionist Realization. Los Altos, Calif.: Morgan Kaufmann.

Sumida, R. A.; Dyer, M. G.; and Flowers, M. 1988. Integrated Marker Passing and Connectionism for Handling Conceptual and Structural Ambiguities. In Proceedings of the Tenth Annual Conference of the Cognitive Science Society, 146-192. Hillsdale, N.J.: Lawrence Erlbaum.

Touretzky, D. S., and Geva, S. 1987. A Distributed Connectionist Representation for Concept Structures. In Proceedings of the Ninth Annual Conference of the Cognitive Science Society, 155-164. Hillsdale, N.J.: Lawrence Erlbaum.

Waltz, D. L., and Pollack, J. B. 1985. Massively Parallel Parsing: A Strongly Interactive Model of Natural Language Interpretation. *Cognitive Science* 9: 51-74.