90-TS

Proceedings of the Sixth IASTED International Symposium

EXPERT SYSTEMS

Theory & Applications

Los Angeles, U.S.A.
December 14-16, 1989

90 - TS

EDITOR: M.H. Hamza

A Publication of
The International Association of Science
and Technology for Development - IASTED

ISBN 0-88986-137-4

ACTA PRESS
ANAHEIM * CALGARY * ZURICH

Towards A
Semantic Binary Connectionist Network:
Hybrid Expert System Knowledge Base Design*

William B. Feild Jr., Jainendra K. Navlakha
and Naphtali Rishe

Florida International University
School of Computer Science
University Park
Miami, Florida 33199

Abstract

Maintaining a knowledge base (KB) using the Semantic Binary Model (SBM) offers several advantages, compared to the relational and more traditional models, in terms of quality and ease of design, maintenance, and flexibility. The structure supports rapid retrieval during the searching and matching phases and efficient performance during rule updates to the KB. The Semantic Binary Connectionist Network is a hybrid system that uses an SBM-based KB. This system allows further efficiency by implementing the KB as an associative memory, using a connectionist network, which leads to even more rapid access during the searching and matching phases than allowed by the standard SBM or the more traditional models.

Keywords - Neural network, Associative memory, Semantic Binary database model, Expert systems.

Introduction

The problem domain for expert systems continues to expand with each successful design. However, characteristic to each system is the necessity to search some solution space, selectively and efficiently, in order to construct the solution(s) [21]. In problems where the solution space is not large, the search may be an exhaustive one with minimal penalty in terms of system performance. However, in those cases where the solution space is not small, heuristics and other methods are used to try and prune the solution space in order to reduce the degradation in system performance that arises from extensive searching. Expert systems achieve high performance in those cases where they manage to minimize the search activities on their KB.

Knowledge Base Management Systems

The key features of knowledge base management systems (KBMS) are their knowledge representation schemes and reasoning mechanisms [16,21]. There are many different knowledge representation schemes and associated reasoning mechanisms. Included among the many are: predicate calculus, rule-based (or production) systems, semantic networks, frame-based systems, and induction (by example) systems.

The rule-based systems, like OPS5 [2,3], are based upon high-level programming languages. These systems represent their *a priori* knowledge as a set of unordered facts and represent their *potential* knowledge as a set of unordered statements (or rules) called productions, each containing a set of antecedents and a set consequents, in the form: <u>IF</u> {antecedents} <u>THEN</u> {consequents}. One reasoning mechanism determines and isolates a set of the rules which have all of their antecedents satisfied (or found *True*). It then performs some conflict resolution to determine which one rule out of the isolated set will have its consequents applied to the KB. This procedure is repeated until some goal state is reached, or the system is halted. This mechanism is called forward-chaining. Another mechanism, called backward-chaining, works in the opposite way, looking to satisfy the consequents and then identifying the antecedents whose truth will guarantee the truth of the consequents, and repeating the process until the truth of all desired statements is established.

Consequently, a rule-based expert system KB can be viewed as a collection of *a priori* facts and a collection of rules, with the rules, in turn, consisting of *potential* facts (in the form of antecedents and consequents) awaiting application to the KB.

The key principle involved in both the reasoning and the conflict resolution is symbolic processing, i.e., symbol matching, set operations, sorting, and pattern matching. Independent of heuristics, the ability to quickly resolve these types of operations is most critical to system performance. Unfortunately, symbolic processing operations are the slowest operations for most larger expert system KBMS's. One effort to circumvent this characteristic is the "Rete match algorithm" [4]. Rete exploits two properties shared by all rule-based systems., namely, the contents of the KB are large and change very little (less than 1%) with each iteration, and the antecedents (and consequents) are composed of many similiar, if not identical, patterns. Rete takes advantage of these properties by retaining information from cycle to cycle, updating the KB as needed to reflect the changes from each cycle. This

---

* This research has been supported in part by a grant from Florida High Technology and Industry Council

makes work for the pattern matching operation depend primarily upon the number of changes to the KB rather than the size of the KB. Additionally, Rete utilizes a 'rule pre-compiler' that locates common rule components and eliminates as many of them as possible. This allows Rete to do many operations for an entire set of rules just once rather than for each and every rule. With these features Rete is able to achieve very efficient sorting, updating, and matching operation performance in uniprocessor environments. The algorithm is being adapted for use in proposed parallel processor environments.

Although symbolic processing operations are the slowest operations for most larger expert system KBMS's, they happen to be the ones for which database management systems (DBMS) are most well suited by virtue of their file structure and operations defined.

Some work has been done in the area of implementing KB's using existing database models. One such effort emulates OPS5 with a relational model [17,19]. Several methods are proposed in an effort to minimize the effort spent in pattern matching and update operations. Many of these are similiar to the methods utilized in the Rete match algorithm. The remaining methods encourage the use of heuristics to take advantage of the sequence of DBMS queries, the knowledge structure, and the goal of the operation.

Semantic Binary Model

One database model, the Semantic Binary Model (SBM) [12], is very well suited to allow efficient management of both reasoning and conflict resolution activities. The logical database consists of a collection of elementary facts of two types: unary facts categorizing objects, and binary facts relating pairs of objects. The facts are represented by tuples consisting of an object id, an attribute id, and a value for that attribute. Each fact in the physical database also has an associated 'inverse', which allows direct reference to every object, category, and relation, since the facts (and their inverses) are stored in lexicographic order in a single file. The sorted file is maintained in a structure similar to a B-tree. The variation of the B-tree used allows both sequential access according to the lexicographic order of the facts, as well as random access by arbitrary prefixes of such facts. This structure supports retrieval in a single block access per elementary query. Complex queries are decomposed into several elementary queries. It also supports efficient performance of update operations [11].

Some work has been done in the area of using the semantic database schema on top of more conventional database models and the understanding of the connections between the models [7,9], as well as on stand-alone SBM DBMS's [7,15,20].

The advantages of the semantic model versus the relational and older models, in terms of quality and ease of database design, maintenance, and flexibility, are known [12]. Furthermore, the advantages with respect to efficient storage structure have been proven recently [13]. Finally, the SBM has potential for much more efficient implementation than the conventional models because the physical aspects of data representation are user transparent. This allows greater potential for optimization, since more can be changed for efficiency regards without affecting the user applications. Also, in the SBM, the system has more knowledge about the semantics of the data and relationships between them, and this can be utilized to organize the knowledge so that the frequently used operations can be performed faster at the expense of less frequently used operations. Most significantly, the SBM achieves two results: very efficient performance of update operations, and one disk access for pattern matching and retrieval (provided the set of matched facts can be stored in a single block) [14].

Semantic Binary Connectionist Network

The SBM offers a number of possible improvements for the KB of expert systems. These allow the expert systems to achieve high performance by minimizing the search activities on their KB. However, these improvements are still limited by the restrictions of conventional computer systems [16]. These limitations include:

1) The size of main memory is typically too small for storing large KB's.
2) Without the storing and maintaining of selected data, memory cannot be accessed by content directly; it must be accessed indirectly by content and then the address.
3) Memory words are accessed and processed in sequence by a single processor.

Discussion of the first limitation is beyond the scope of this paper. However, work is being done in the area of large, distributed DBMS's [8,18]. The SBM does attempt to address the second limitation by utilizing a hashing scheme to rapidly access the location within the file where the requested facts are stored.

Circumventing the last two limitations is the focus of the final portion of the proposed Semantic Binary Connectionist Network. Our approach is to implement the SBM DBMS on a connectionist network model that emulates the functionality of an associative memory.

An associative memory is a device that is content-addressable. This means that data stored in the main memory can be accessed by specifying the contents of the data rather than the address. The basic idea of a connectionist network used to emulate the functionality of an associative memory is to eliminate the need for directly storing information in memory by storing it in the form of distributed network connections. These networks feature a high degree of parallelism, distributed storage of information, robustness, and very simple processing elements of low complexity (see Figure 1).

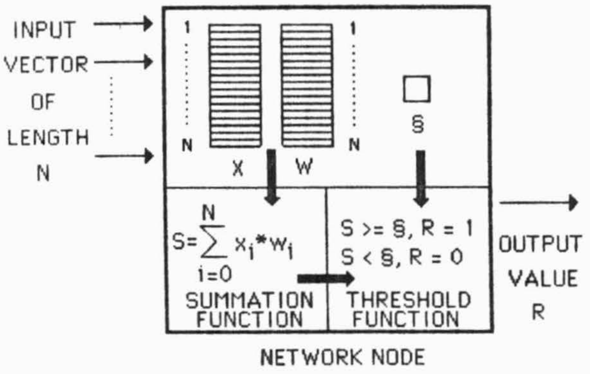

Figure 1. Hopfield Network Node

A query is presented to the network in an appropriate form and the output of the network, when it reaches a stable state, is the desired answer to the query. Analysis of connectionist networks has shown that their power lies in their fast selection, matching, and updating capabilities [16], and these operations are of primary importance for maximizing expert system performance.

The search for a suitable connectionist associative memory model is currently under way. Several issues must be examined before a model can be selected. Since the criteria of primary importance is the accurate, rapid retrieval of one or more matching patterns from a large memory, two important problems must be solved: noise elimination without false or distorted memories, and multiple-match resolution.

One potential model is the 'Human Associative Processor', or HASP [5]. This system consists of three basic components: a recurrent network, readout control units, and an associative network memory model.

The associative network memory model for HASP is functionally similiar to the Hopfield Neural Network [6], although proposed earlier. The Hopfield neural network is an associative memory model with a simple, flexible structure (see Figure 2). It is capable of storing information as well as allowing for error correction and closest-match search. The storage capacity C of the Hopfield neural network has an upper bound of $C = O(N^3)$, where N is the number of network neurons, and it allows for accurate pattern matching and rapid retrieval. However, in order to accomplish the latter and ensure that the system will stabilize with a final answer, the storage capacity upper bound drops to $C \leq N$ [1,10]. This restricts the usefulness of the standard Hopfield neural network for our purposes. Alternative networks are currently being investigated.

The readout control units of HASP regulate the selection of items in memory so that an item is not re-selected, for the same query, until all other matching items have also been selected. An inhibitory signal is used to suppress the output of the selected items. In conjunction with the recurrent network, the system is capable of distinguishing an input which matches a memory (and is immediately passed on as output) from an input which is 'false' and/or a multiple-match memory.

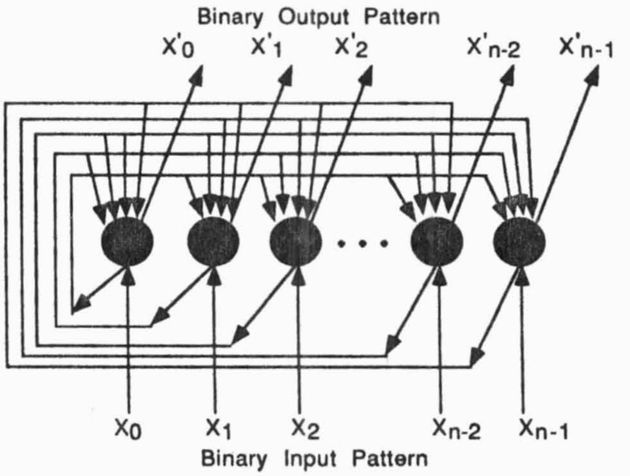

Figure 2. Hopfield Neural Network

Conclusion

The Semantic Binary Model offers a number of significant improvements to the performance capability of current expert systems. From an operational viewpoint, the structure supports rapid retrieval during searching and matching operations and permits very efficient performance of update operations. From a system management viewpoint, the model offers advantages over conventional models in terms of quality and ease of KB design, maintenance, and flexibility. Additionally, the structure incorporates many known features that have been offered in the way of improvements to conventional KB management systems.

The further enhancement of the SBM logical structure to include an associative memory processor, supporting the SBM physical structure, promises to offer more significant improvements to system performance by allowing multiple processors, with their inherent robustness and high degree of parallelism.

Bibliography

[1] Y. Abu-Mostafa, J St. Jacques, "Information Capacity of the Hopfield Model", IEEE Transactions on Information Theory, July 1985.

[2] Digital Equipment Corp., "VAX OPS5 Reference Manual", September 1985.

[3] C. Forgy, "OPS5 Users Manual", CMU CS 81-135 Computer Science Dept., Carnegie-Mellon Univ., 1981.

[4] C. Forgy, S. Shepard, "Rete: A Fast Match Algorithm", AI Expert, January, 1987.

[5] Y. Hirai, "A Model of Human Associative Processor (HASP)", IEEE Transactions on Systems, Man, and Cybernetics, September, 1983.

[6] J. Hopfield, "Neural Networks and Physical Systems with Emergent Collective Computational Abilities", Proceedings of National Academy of Science U.S.A, 1982.

[7] A. Jain., "Design of a Binary Model Based DBMS and Conversion of Binary Model Based Schema to an Equivalent Schema in Other Major Database Models", M.S. Thesis, University of California, Santa Barbara, 1987.

[8] Q. Li, "The Architecture and Related Control Problems of a Transputer Based Highly Parallel Database Computer", Ph.D. Dissertation, Florida International University, 1989.

[9] P. Lyngbaek, V. Vianu, "Mapping a Semantic Database Model to the Relational Model", Proceedings of the ACM SIGMOD Conference, 1987.

[10] R. McEliece, E. Posner, E. Rodemich, S. Venkatesh, "The Capacity of the Hopfield Associative Model", IEEE Transactions on Information Theory, July 1987.

[11] N. Rishe, "A File Structure for Semantic Databases", FIU SCS TR 88-001 School of Computer Science, Florida International Univ., 1988.

[12] N. Rishe, "Database Design Fundamentals: A Structured Introduction to Databases and a Structured Database Design Methodology", Prentice Hall, 1988.

[13] N. Rishe, "Efficient Organization of Semantic Databases", Foundations of Data Organization and Algorithms - Lecture Notes in Computer Science, Vol 367, Spring-Verlag, 1989.

[14] N. Rishe, "Semantic Binary Database Model", FIU SCS TR 89-004 School of Computer Science, Florida International Univ., 1989.

[15] N. Rishe, "Transaction-management System in a Fourth Generation Language for Semantic Databases", ISMM International Conference on Mini and Microcomputers, 1988.

[16] S. Y. W. Su, "Database Computers: Principles, Architectures, and Techniques", McGraw Hill, 1988.

[17] A. Tzvieli, "Emulating OPS5 with a Database", AI Expert, February, 1989.

[18] D. Tal, N. Rishe, D. Barton, N. Prabhakaran, "High-throughput Highly Parallel Database System", Proceedings of the 26th ACM SERC, April, 1988.

[19] A. Tzvieli, "On Implementations of Production Systems Using DBMS, Proceedings of the Third International Conference on Data and Knowledge Based Systems, Morgan Kaufman Publishers, 1988.

[20] N. Vijaykumar, "Toward the Implementation of a DBMS based on the Semantic Binary Model", M.S. Thesis, University of California, Santa Barbara, 1987.

[21] D. A. Waterman, "A Guide to Expert Systems", Addison Wesley, 1986.