A LIFELIKE MODEL FOR ASSOCIATIVE RELEVANCE

Paul B. Post
Norden Division
United Aircraft Corporation
Norwalk, Connecticut

Summary

This paper deals with the general problem of association in artificial intelligence. The suggested approach is based on a multidimensional information space that organizes itself as a large number of elementary associations are supplied. The space contains such entities as words, technical terms, symbols, phrases, and proper names, as may be appropriate for information retrieval, language processing, problem solving, etc. A set-theoretic transformation of the space is applied. The resulting system is capable of responding with whatever entity is most relevant to the entities selected for inquiry. There is no sequential search. Lifelike characteristics include parallel processing, equipotential memory, tolerance to malfunctions and inexact inputs, random connection and sharing of logical elements, and adaptability. Statistical performance is briefly described, together with the results of small-scale computer simulations. A tentative hardware design is outlined, based on MOS techniques currently under development.

Introduction

The purpose of this paper is to outline a new theoretical approach to basic problems of association that occur in many applications of artificial intelligence. Association, or the linking of mutually relevant pieces of information, lies at the heart of all cognitive behavior, and therefore warrants intensive study on its own terms. The approach described in this paper leads to a type of conceptual model and system organization that exhibit unusually lifelike properties.

In pursuing this study, I found it difficult to build upon specific techniques of previous work. However, certain concepts and insights developed over the past few decades were helpful in formulating the approach; in particular, Ashby's homeostatic mechanism[1], Hebb's cooperative cell-assemblies[2], Rosenblatt's successful use of random nets[3], and Weiss's view of the organic growth of knowledge[1*]. Some of the work of MacKay[5], von Neumann[6], Reiss[7], and others is also relevant. Space does not permit an adequate discussion of these contributions, but the approach taken in this paper can be seen to represent a synthesis of such ideas, aimed at a possible method for dealing with the general problem of association.

Functions of an Associative Model

Our ultimate objective is to design machines that behave "creatively" with respect to the information they receive from their environment. We would like to provide machines with <u>general and efficient</u> faculties for concept formation, induction, discourse, conjecture, drawing analogies, and the like. Let us try to identify a few of the objective functions implied by such behavior.

One of the required functions is the ability to make use of any and all stored information as the need arises. The more a system selectively takes advantage of past experience, the more successful it will be in handling new experience. But the "need" for certain stored information must first be recognized as a function of the particular situation at hand. This can be accomplished through the perception of key elements and element combinations in the situation, but perception and recognition should occur even if the relevant information had been expressed in different terms when stored. In addition, a dynamic system should accept new information in a way that usefully relates it to all of the stored information. When new information invalidates some old information, the latter must be modified or even restructured without jeopardizing valid information. The same is true when the meanings of some elements are shifted, specialized, generalized, split, or combined with the passage of time.

To a large extent, the degree of success obtainable in the performance of such functions depends on the nature of the associative links with which elements of information are organized. When the links are few and rigid, the system can be expected to have limited capability. When associations are rich and dynamic, there is a greater chance of success. What we evidently need is a general and flexible associative mechanism by which patterns of relevance can be largely self-organized and self-maintained.

A customary approach in designing cognitive systems is to begin by constructing a model that represents as closely as possible the appropriate information macrostructure, proceeding from the general to the specific. (Here the term "macrostructure" refers to an overall organization of the data to be handled by the system. Such macrostructures are the traditional means by which complex relationships are visualized, recorded, and communicated.) Using this approach, the designer proceeds to impose associative links hierarchically and attempts to anticipate potential modes of adaptation.

It is quite possible, however, to choose the opposite direction for constructing an internal model. One could begin with a large number of elementary associations known to be relevant, and <u>let the macrostructure be implied by the characteristics of the relata</u>, rather than predetermined. For example, a child learns about animals and their approximate groupings long before he can be taught to view the animal kingdom taxonomically. In this case, the obvious differences between insects, fish, birds, and mammals readily imply the existence of at least four major groups. Only later are spiders, bats, and dolphins pointed out to be special cases, and they are often remembered as such.

Another example appears in the learning of a native language, whereby grammar and syntax are universally inferred from simple illustrations. Indeed, few people ever achieve an accurate conception of the structure of their native language.

Despite the apparent conciseness of information macrostructures, their use for general purposes of association may actually impede performance of the cognitive functions described above. In particular, it would be desirable if elements of information could be freely located, associated, modified, or otherwise processed without reference to a predetermined macrostructure.

<u>Entities in Information Space</u>

The proposed associative mechanism can be developed by first postulating an information space into which elements of information, or "entities," can be placed as localized volumes or points. There are no restrictions on the dimensionality of the space. The space contains such entities as English words, technical terms, symbols, data, phrases, reference numbers, and proper names. Each entity must have finite information capacity and be meaningfully related to other entities.

Now, a convenient way to allow associative links to form is simply to put related entities close to one another. Since our information space can have any number of dimensions, each entity can always be placed adjacent to <u>any number</u> of related entities. If it later becomes necessary to bring remote regions together, the space suffers no worse topological disturbance than does a string when a loop is made in it. If more room must be made available in a particular region, the space can stretch like a rubber band with infinite compliance. Thus the space can be as flexible and dynamic as required, and the associations it contains can be as rich as required.

For purposes of mechanization, we must now impose some arbitrary conventions. There are two classes of entities. Explicit entities (E_e) are those deliberately entered as expressions in the input language. Implicit entities (E_i), though meaningful to the system, are unverbalized. They are automatically inserted among the E_e when associations are formed. (It will be shown that long-term storage of the E_e is unnecessary.)

Each E_i is linked to exactly three E_e, thus forming the center of a three-pointed star. There is no limit, however, on the number of links attached to an E . Any three E_e linked to a common E_i are said to be members of a "triplet." Figure 1 shows a simple three-dimensional representation of a region of the information space. Note that an E can be a member of several triplets, and that two different triplets can share two (but not three) common members. The lengths of the links have no significance.

The triplet convention was chosen because, in developing a rudimentary system, I was forced to adopt a constant number of members per group. The use of two members per group, I felt, was too unspecific for most associative purposes, while four seemed unnecessarily rigid. Nevertheless, the use of triplets allows a substantial variety of associative structures to form.

For brevity, we shall refer to all explicit entities as "words," but their general nature should be kept in mind. For example, the entire Gettysburg Address could be considered as a word, since it has finite information content and is relevant to "Lincoln," "Civil War," "1863," etc. Note, however, that most languages tend to reduce long expressions to short ones where feasible. Thus we can unambiguously refer to this entire speech by an expression containing only 17 letters and a space. Furthermore, the speech is very strongly suggested by the single nine-letter word "fourscore." Hence we shall assume that all usable entities have a limited word-length, and can therefore be described simply as "words."

Storage of any triplet in the system establishes a particular strong relevance among its words. However, the full meaning of the association cannot be derived from the single triplet, but depends on the triplet's location and interconnection among many other triplets. When a triplet is stored, it acts to pull together selected portions of the information space. Hence the total region of associative involvement can be indefinitely large and complex.

Following storage (or between subsequent storages) it is possible to address the information space, or "inquire," with two words. If these are common members of any triplet, we want the system to respond promptly with the third word of the triplet. This should occur regardless of the direction from which the triplet is addressed. in some cases, the response to a single-triplet inquiry may be very helpful to the user of the system; in other cases, it may seem redundant. In all cases, the response will be <u>meaningful and relevant to both words of the inquiry</u>.

Other modes of operation involve the concatenation of two or more single-triplet inquiries. If, for example, two triplets share one word, then under certain conditions the system can respond in a manner that simulates recognition of an analogy. Of course, there will be many cases in which two triplets that share one word do not necessarily represent an analogy. Here the theory rests on the fact that the special relation between the inquiry and the response is meaningful and probably useful, while exact logical interpretations are of secondary importance. Note that responses are in general unknown to the user until the system produces them. Once produced, they may be so revealing that further analysis is unnecessary. In any case, logical interpretations could be clarified by further inquiry.

Similarly, triplets can be manipulated to form dichotomies, lists, tables, and classes. (It is not important to describe such structures in detail at this point. This is because the principal model to be developed in this paper transforms the information space in such a manner that the approximate associative "distance" between entities outweighs the form of the linkages between them.)

Our information space thus serves as a medium that facilitates association among related entities — as a chemical solvent facilitates the interaction of solutes. The major properties of the space may be summarized as follows: First, each entity occupies a single location determined by its relevance to other entities. Second, multidimensionality allows practically unlimited reshaping (growth, shrinkage, folding, etc.) without distorting existing associations. Third, the fundamental unit of association is a group of three words, linked as a triplet by virtue of their mutual relevance. Such triplets could be used as building blocks to form larger structures. Fourth, the basic mode of input/output interaction is the inquiry, in which two input words evoke a single word response. Fifth, inquiries may be concatenated to pursue a line of investigation, representable as a continuous path from triplet to triplet.

<u>System Organization</u>

Despite the complexity of the information space and its many interconnecting structures, the single triplet should be simple enough to mechanize by algorithmic processing of one kind or another. The form of processing chosen should allow us to store information by putting in three E and having the system define a suitable E.. To retrieve information, we would input two E and let the system locate the corresponding E. (in effect, find the area of maximum relevance). The third E would then result from interaction between the given E_e and the E_1.. If such operations could be carried out independently of other triplets, we need not be concerned with the macroscopic complexities of the information space at all.

Our method of attack was motivated by well-known biological evidence: parallel processing, equipotential memory, extensive sharing of logical elements, random-like interconnection, and the like. These characteristics suggested a set-theoretic solution in which the entities of the information space are transformed into binary sets.

The heart of our system is a large cellular array containing K identical units. Each unit, or cell, has provisions for simple logical operations and a capacity of one bit of permanent memory. The cells operate independently of one another. By means of language-dependent encoding (described below), words can be represented as distinctive patterns of binary cell activity. Encoding is such that each word will have its own invariant pattern. An approximately constant fraction of the cells, w, is energized by any single word, independent of word length. Thus the set of active cells W = Kw. (We shall use capital letters to refer to either the name or the number of a set.) If patterns of activity could be displayed, they would appear random and reasonably uniform in spatial distribution.

Let W-, W_2 W_3 represent the three word patterns of a triplet to be stored (Figure 2a). The implicit entity linking these words can be modeled as the logical intersection of the three sets, or $E_i = W_i n W_2 n W_3$" If the words are input sequentially, each cell needs two bits of temporary storage to determine whether it is a member of E.. If it is, the cell registers a 1 in its permanent memory; if not, its memory remains at 0. The temporary-storage portions of all cells are then reset to their original states. This has the effect of discarding all of the input except the residual E.. Figure 2b illustrates the complete storage sequence. The circles show the locations of E_1 cells (it is assumed that two of these had previously been registered).

As shown by the cell in the lower left corner of each array, permanent memory accumulates as the logical union of implicit entities:

$$M - E_{i1} U E_{i2} U E_{i3} U .. U E_{it},$$

where t is the total number of triplets stored and M is the set of registered memory cells. An important system parameter is m = M/K, the total density of memory.

Upon inquiry, W, and W~ set up an "inquiry pattern" $Q = W, nw_2$, as shown by the cells marked 2 in figure 2c. The correct response R (in this case R = W~) cannot be retrieved directly, but it can be approximated by a process of cell elimination. The rule can be stated symbolically as $QflM\overline{CR}$; that is, cells that belong to the inquiry Q, but are not members of the total memory pattern M, cannot possibly belong to the response R. This is obviously so, for any cell that is a member of both Q and R is by definition a member of $W, nw_2 nw_3$, and hence would have been registered as a member of M during storage. This rule allows the system to eliminate the two cells marked X in the final array of Figure 2c. The remaining seven cells include the five of W_3 plus two false cells, but the latter are not particularly troublesome.

In describing the storage and inquiry modes, we treated explicit entities as arbitrary patterns. Encoding is required to relate these patterns to actual expressions in the input/output language. In its present form, the encoding method artificially extends words up to a standard number of characters. This is a convenient way of regulating w, the relative density of word patterns. Extension is accomplished simply by repeating the word until a standard length X is reached. The system's alphabet can contain any reasonable number of characters a; like X, a must be constant. If a = 38, for example, there is room for 26 letters, 10 numerals, a space symbol (b), and an end-of-word symbol (e). If all 38 characters are assumed to occur with equal frequency, and if we take X = 19, there will be almost 100 bits per word -- a substantial capacity for expressing English words, names, etc. According to these rules, "strontium 90" would be input as STR0NTIUMb9OeSTR0NT.

Character sequence is taken into account by numbering the letter "places" from 1 to X and combining these numbers with their corresponding characters. "Strontium 90" then becomes a simultaneous combination of 19 symbols whose order is unimportant: 17, M9, N5, N18, 04, 017, R3, R16, Si, S14, T2, T6, T15, T19, U8, 911, 012, bIO, el3.

The system treats all 722 symbols (aX) equally and independently. Each symbol is represented as a fixed, random pattern in the cellular array. The 722 patterns can overlap, yet each is distinctive. Since the patterns are fixed, the encoding of each symbol can be wired in.as an interconnection tree — there is no need to wait for table look-up or other auxiliary operations. Permanent encoding is feasible because there is no fundamental reason for changing the patterns. The only known conditions for efficient encoding are (a) a constant optimal number of cells per symbol, and (b) complete independence among patterns. This requirement for strict independence is the main reason for using random connections. (In fact, all attempts to improve encoding by interfering with this independence have failed.)

Let C represent the set of cells energized by a single symbol. Then word patterns result from the union of symbol patterns:

$$w = c_1 u c_2 u c_3 u \ldots u c_r$$

We can now calculate w as a function of A and c (where c is the symbol pattern density C/K, or "encoding depth"). Define the complementary variables w = 1-w and c = 1-c. Assuming orthogonality, w will be approximately equal to the A power of c. For c = 0.02 and A = 19, w is about 0.32. The encoding depth c is a critical parameter inasmuch as it directly affects the relative densities of words, triplets, and total memory. All of these variables interact to produce a certain overall system efficiency.

Upon inquiry, the criterion QOMCR is applied in testing every language symbol. For example, if the symbol S6 includes in its set C one or more cells that are in QOM, the symbol can be eliminated as a possible candidate for the sixth place in the response word. However, if the symbol T6 happens to be one of the correct symbols, it will always pass this test. Thus the correct response R will always emerge, though it may be accompanied by a number of nonsense words, depending on system efficiency. In general, larger values of c (up to a point) can present more opportunities to test a given symbol and can thereby suppress false symbols more efficiently.

The encoding scheme described above is admittedly crude; in distinguishing characters by their exact positions in words, it takes no account of linguistic considerations. Nevertheless, it has been found to be effective. While performing the interfacial transformation between the external language and the cellular array, the encoding process carries with it some very useful functions. First, it allows entities to be dispersed and reconstructed with equal ease; the same physical structure could accomplish both. Second, it enhances the efficiency of the system in determining R from only a partial knowledge of R. (In a large system, the probability that a wrong response symbol contains no cells that are in QOM is very small.) Third, and most significant from a theoretical standpoint, encoding makes direct use of the richness and variety of language for the representation of concepts.

At this point, the system as a whole can be seen to have the following properties: First, it enables multidimensional associative structures to be stored as transformations of implicit entities only. Since the latter result from the intersection of word patterns, relatively few bits are added to the total memory as each triplet is stored. Most of the input is discarded immediately. Second, memory is accumulated in the form of overlapping quasi-random patterns that bear no simple relation to external information structures. Information diffuses throughout the memory. Third, logical elements (cells) are extensively shared, randomly connected, and operated simultaneously. Their combined action is statistical, but not indeterminate. Logical functions, temporary storage, and permanent memory all interact locally. Fourth, individual cell malfunctions would not render the system inoperative, since there would normally be many cells available for redundant processing. Malfunctions would instead have the effect of reducing overall efficiency.

System Performance

For any given inquiry corresponding to a stored triplet, the number of characters, N , that emerge at a given lettex-place must be an integer larger than 0. The number of possible words, N_w, that emerge as combinations of these characters is then the product of the A values of N . If the inquiry does not correspond to a stored triplet, N will usually vanish in at least one letter-place, yielding N =0.

On a statistical basis, the expected values of N and N are minimum at a certain optimum encoding depth c'. Preliminary calculations show that for t = 1 (only one triplet stored) and K = 100, a = 10, A = 5, N_w attains a minimum of 2.27 at c' = 0.26. Out of 10 possible "words," the system selects about two, one of which is bound to be the correct response. (The other is generally a misspelled version of the correct word.) For t = 3 and K = 300, performance improves to N_W ^ 1.30 at c' = 0.16.

It has been found that optimum performance occurs at or near the encoding depth that produces m = 0.5, depending on the extent of triplet interconnection. (Note that the number of combinations of K items taken M at a time is also maximum at M = K/2.) Use of m = 0.5 simplifies the task of selecting parameters to obtain a desired level of performance.

Values of K obtained in this way may be compared with the number of binary cells that would be needed for ordinary non-associative storage. According to information theory, the number of bits required to store 3t words is 3tX $\log_2 a$. In our system, however, it is possible to store t unconnected triplets (3t words) in fewer cells, provided a large enough value of N can be tolerated. In effect, the system trades off precision of response for associative power. Evidently, such economy of storage is a result of our having discarded all but the implicit entities.

Larger values of N (e.g., 10 to 1000) would be acceptable if the system were provided with a second memory, whose function it is to check which of the response words are likely to be valid. Such a "word memory" would have its own encoding provisions and would operate in a manner analogous to the main "triplet memory." Briefly, the word memory would use a larger value of c and would store only the W pattern for each word, letting memory accumulate as the union of the individual W sets. (The large c and complementary logic are chosen because X>3.) Following inquiry, the system would pause to test each output word in sequence. Those words that pass the test are then transmitted as responses. The correct response always emerges.

Digital computer simulations of the system were originally attempted to prove whether it would work in at least an elementary way. Good results were occasionally obtained, and these increased in frequency as more was learned about the technique. Eventually, computer simulation became the principal tool for all experiments.

Because of capacity/speed limitations and the need to store large amounts of encoding data, the simulation actually represented a very small associative system. It had only 160 cells in the triplet memory and, for simplicity, an equal number of cells in the word memory. The program used X * 5. Words longer than 5 letters were truncated; those shorter than 5 letters were simply repeated without using an end-of-word symbol.

One group of experiments was based on the storage of three unconnected triplets (9 words) with a » 17. C was set at 26 cells for the triplet memory and 36 for the word memory. With these parameters, successive values of M after each storage should theoretically have been 32.5, 58.2, 79.0. Experimental data agreed to within 3 cells. For three triplets, there are nine pairs of words that can be used for productive inquiry. Of these nine inquiries, seven gave single responses, one gave two, and one gave three. In each case, the correct word was produced.

Another group of experiments used three triplets interconnected in a ring structure (6 words) with a = 16. Here performance was best with C = 27 and M = 66 in the triplet memory. The decreased M indicated greater sharing of M-cells among interconnected triplets. Of the nine inquiries, eight gave single responses and one gave two. The single incorrect response had one wrong letter. When this misspelled word was later used in two inquiries, the system gave single correct responses to each.

The effects of malfunction were tested by arbitrarily setting cells to state 1, thus destroying their ability to eliminate trial symbols by the QHM rule. Using 20% destruction (32 cells) in both memories, seven inquiries gave single responses, and the remaining two gave two each. At 40% destruction, one inquiry was still able to produce a single response. At 55%, one inquiry produced two responses. Again, the correct words always emerged.

The simulated system also showed a tendency to seek out relevant entities that were not stored as such. Using the three-triplet ring structure, two words were selected whose respective triplets shared one word. Under normal conditions, inquiry with these two words gave no response, since they were not members of a common triplet. However, it was possible to induce a response by (a) temporarily setting a number of arbitrary cells to state 1, and (b) reducing C during inquiry. Although fine control of these variables was required, the system could produce a single correctly spelled response. This was the word that joined the two triplets, i.e., the word most relevant to the inquiry.

Tentative Hardware Design

Consideration has been given to possible hardware design in order to visualize size, timing, and modes of control, as well as to identify engineering problems that might be encountered if a prototype machine were built.

As presently conceived, such a prototype would be a parallel-organized machine whose "program" is almost completely wired in. There would be a console containing an electric typewriter used for both input and output, with provisions for initiating storage and inquiry, adjusting C, etc. A separate control unit would provide series/parallel conversion of inputs and outputs, buffer storage of words and symbols, and master timing control.

The encoder would consist of a "forest" of *aX* independent trees, terminating in a large interconnection matrix for random attachment to the memory array. Bidirectional operation could probably be effected on the basis of pulse polarity in the branching circuits. Each branch circuit could be designed (a) to amplify and distribute pulses of one polarity to its branches in one direction, and (b) to act as an OR gate in the opposite direction for pulses of the opposite polarity. A tentative design for such a bipolar circuit has been worked out.

The memory array would consist of K identical circuits, each containing a two-stage counter, some simple logic, and a memory device. (At present, the most promising memory device appears to be the adjustable-threshold field-effect transistor. See, for example, Oleksiak, et al[8].) Each memory circuit would have a single bidirectional connection to the encoder. Other connections would be common to all circuits. It appears that large-scale MOS technology might be suitable throughout, with the major problem being that of interconnection. A real machine of any appreciable size would almost always contain some malfunctioning cells. However, the important maintenance question would not be how to avoid malfunction, nor even how to locate malfunction, but how often groups of cells should be replaced to <u>maintain an acceptable level of malfunction</u>.

Note that the mechanization of encoding permits all symbols to be tested simultaneously upon inquiry. Consider a single cell in the array. If it is not a member of QOM at the time of inquiry, it remains quiescent. If it is a member of QOM, it sends out a "fail" pulse regardless of its membership in one symbol pattern or another. The fail pulse will flow back through the encoding trees and eliminate <u>any</u> symbol to which it is connected as a candidate for the response word.

Due to such parallel organization, operation would be extremely rapid. For example, storage of a triplet could be completed within about 10 clock cycles after buffer storage of the three words. Inquiry would be even faster, up to the point at which the word memory begins to act. Thereafter, the testing of each word is essentially a two-clock-cycle operation. A conservative estimate of the clock rate would be in the neighborhood of 1 MHz.

The technique for encouraging response had been simulated by temporarily setting cells to state 1. In a prototype machine, the same effect could be obtained by partitioning the memory array into sections and blocking the generation of fail pulses from one or more sections. Since memory is equipotential, we would be able to control the strength of memory utilization in a way that is independent of the information being processed.

Much work remains to be done before one could estimate how a large system of this kind might behave. The mathematical analysis should be rigorously reviewed, and statistical data collected by computer simulation. Nevertheless, it is possible to foresee some useful improvements and modifications.

First, encoding should be expanded and made more flexible by including bigrams, trigrams, common roots, etc., and by using relative positions of these in words rather than exact positions. This might enable the system to guess the meanings of many words and expressions before they are formally stored. It should also help to conserve memory.

Second, the memory should be subject to reversal or decay in order to phase out obsolete information and provide cells for new information. (The recommended memory device has this potential property.) We suspect that a continuous, rather than binary, memory would be best for this purpose. This would permit fine control of N as well as graceful decay of obsolete information. Feedback should be provided to maintain an optimum effective density m as the system gradually adapts to new information.

<u>Conclusions</u>

The associative model described in this paper is based on the following assertions:

1. One of the most important faculties intelligent systems will require is an ability to recognize and retrieve promptly whatever information may be useful for processing a given task.

2. A major factor in determining potential usefulness is the degree of relevance among entities that results from a large number of specific associations stored and modified throughout the experience of the system.

3. An effective medium in which to represent such associations is a multidimensional information space that lends itself to unlimited reshaping.

4. An effective way both to define and address entities embedded in the information space is by means of diffuse random encoding, taking advantage of the symbolic variety and richness inherent in natural language.

The model is capable of retrieving whatever entity is most relevant to the entities selected for inquiry. In performing this function, the model makes use of parallel processing, with its concomitant advantage of high speed. Sequential search is absent, or at least replaced by an editing procedure. Memory is equipotential, providing a substantial degree of immunity to malfunction. Elements of the memory are extensively shared among entities, contributing to overall storage economy. Random connection of the elements has thus far proved to be superior to other schemes of connection, provided certain statistical regularities are maintained. If a continuous-density, reversible memory medium can be provided, the model would be indefinitely adaptive.

There are three major limitations. First, any such system built for a practical application must have very many elements ($K > 10$). Second, the memory of a particular system would not be physically expandable (because of encoding considerations), although it could adapt to new data at the expense of "forgetting" old data. Third, there is always an element of imprecision (for finite K) that requires interpretation on the part of the user. Note that these limitations are also human ones; I believe they go hand-in-hand with the lifelike characteristics mentioned above.

To overcome certain inadequacies of present models in the field of artificial intelligence, we should recognize that the problem of providing a general associative mechanism must be attacked on its own terms. Toward this end, we have sought to show why ordinary information macrostructures, used for visualization and communication, are not necessarily suitable as internal models for association. This paper has outlined an alternative way to mechanize essential functions that will some day be required of intelligent systems.

The work described in this paper was supported by the Norden Division of United Aircraft Corporation as an exploratory research program.

References

1. W. R. Ashby, "Design For a Brain," John Wiley & Sons, Inc., New York, 1952.

2. D. 0. Hebb, "The Organization of Behavior," John Wiley & Sons, Inc., New York, 1949.

3. F. Rosenblatt, "The Perceptron: A Probabilistic Model For Information Storage and Organization in the Brain," Psychol. Rev., vol. 65, pp. 386-407; 1958.

4. P. Weiss, "Knowledge: a Growth Process," Science, vol. 131, pp. 1716-1719, 1960.

5. D. M. MacKay, "Mind-Like Behaviour In Artefacts," Brit. J. Phil. Sci., vol. 2, pp. 105-121, 1951.

6. J. von Neumann, "The General and Logical Theory of Automata," in "Cerebral Mechanisms In Behavior: The Hixon Symposium," John Wiley & Sons, Inc., New York, 1951.

7. R. F. Reiss, "An Abstract Machine Based on Classical Association Psychology," Proc. 1962 Spring Joint Computer Conference, American Federation of Information Processing Societies, San Francisco, 1962.

8. R. £. Oleksiak, A. J. Lincoln, and H. A. R. Wegener, "An Electrically Alterable, Nonvolatile Semiconductor Memory," Government Microcircuit Applications Conference Digest""of Papers, vol. 1, pp. 342-343, Washington, 1968.

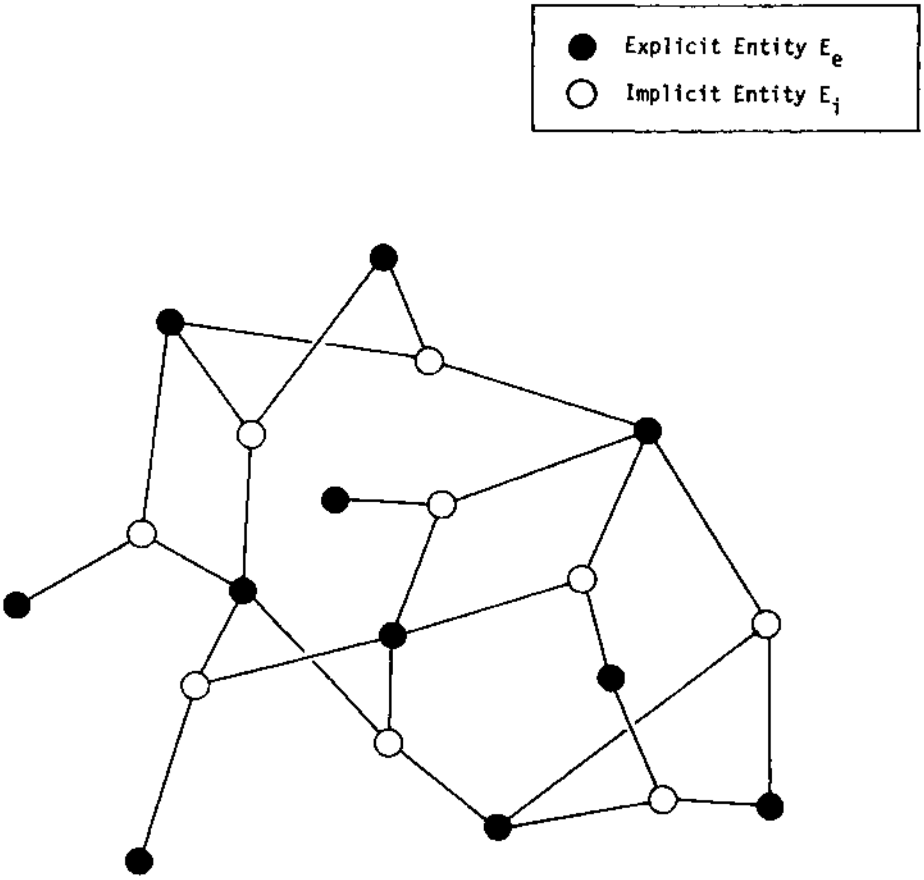

Figure 1. Portion of a Three-Dimensional Information Space

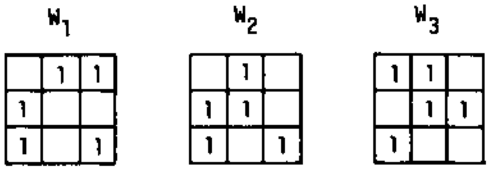

a. Words

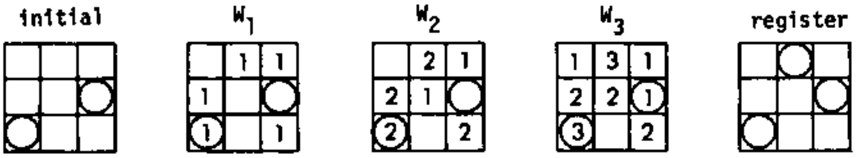

b. Storage

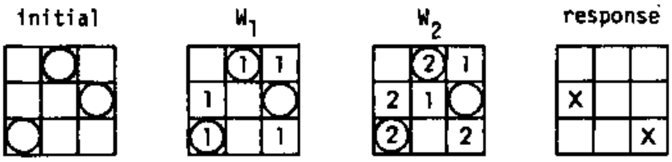

c. Inquiry

Figure 2. Cell Operations for K = 9, W = 5