Chapter 9

Directed Graphs as Memory Representations: The Case of Rhyme*

David C. Rubin

This chapter describes an attempt to uncover the structure of rhyme categories, an attempt that provides evidence for the value of directed graphs as memory representations.

Psychologists know a great deal about the structure of semantic categories (Deese, 1965; Fillenbaum & Rapoport, 1971; Friendly, 1977, 1979; Gruenewald & Lockhead, 1980; Meyer & Schvaneveldt, 1976; Nelson, 1981; Rosch, 1975; Rubin & Olson, 1980). They know much less about the structure of rhyme categories even though rhyme is central to the understanding of retrieval in many domains (e.g., Hyman & Rubin, 1988; Wallace & Rubin, 1988a, 1988b). Only one laboratory has extensively studied the role of rhyme in memory (Nelson, 1981; Nelson, McEvoy, & Friedrich, 1982), and their view, based on cuing effects, is that words in rhyme categories are unstructured except that each word has a link to the rhyme sound that defines its category. Figure 1 is a hypothetical network of the *air* rhyme category based on Nelson's (1981) representation of a rhyme category; the individual words do not link to each other, and the strength of a word's membership is given by the length of its links to the central rhyme node. In contrast, semantic categories, such as *animals* or *parts of the body*, would show many links among the items as well as a link to the central concept (Schvaneveldt, Durso, & Dearholt, 1989).

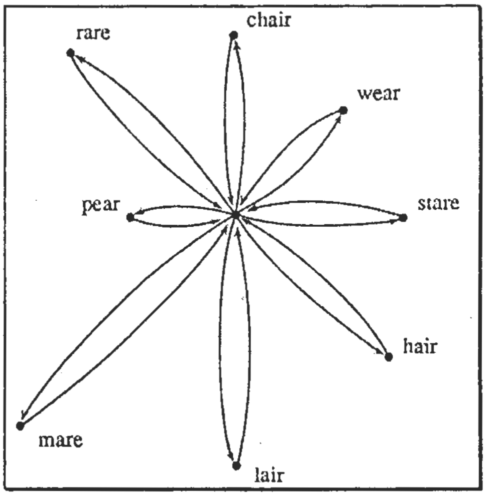

Figure 1. A hypothetical network for the rhyme category *air* based on Nelson (1981).

---

*I wish to thank Orest Zborowski for writing the programs that converted the recall protocols into similarity matrices, and Russ Branaghan, Doug Nelson, and Roger Schvaneveldt for their comments on the chapter. Support was provided by NSF grant BNS-8410124.

Multidimensional Scaling of Rhyme Categories

In order to investigate directly the way in which rhyming words are organized in memory, retrieval from long-term memory was studied using standard multidimensional scaling (MDS) techniques available before the invention of Pathfinder. To provide a comparison, semantic categories were also included in the study. People were asked to list all the words they could that referred to instances of a rhyme or a semantic category (Bousfield & Sedgewick, 1944), a task not very different from that discussed by the British Empiricists, and a task that can be viewed as tracing a path through memory. Under this view, words recalled next to each other are assumed to be related, and MDS techniques can be used to provide a picture of the associative memory structures (Rubin & Olson, 1980).

In the particular experiment reported here, 100 Duke University undergraduates were asked to list all the words they could in 60 seconds that had the rhyme sounds of *air, ear, ed*, and *ee*, as well as all the words they could that belonged to the semantic categories of *animals, beverages, furniture*, and *parts of the body*. These eight categories were chosen as a sample because they all had a large number of instances. A tape recording, read by an undergraduate native of North Carolina, was used to present the stimuli. Two random orders were presented with rhyme and semantic categories alternating. Subjects were asked to "please turn to page X and write down all the words that you can that ..."

All responses were compiled, with different responses being combined under the same word only if they were different spellings of the same word. Singular and plural words were not combined because they have different rhymes. Thus, *eye* and *eyes* were scored as separate responses. The most frequent 20 words in each category were then selected for further analysis. This ensured that at least 10 subjects recalled each of the 20 words used. For each category the number of times any of these 20 words were emitted in succession was counted as a measure of similarity (Rubin & Olson, 1980). That is, each cell in the lower triangular similarity matrices was indexed by two words, with the value of the cell equaling the number of subjects who recalled the two words immediately next to each other. Thus the cell for *dare-care* in the rhyme category *air* contained 39 because 39 subjects out of 100 recalled these words next to each other.

The resulting symmetrical similarity matrices were submitted to a smallest space analysis MDS solution (Lingoes, 1973). Two-dimensional solutions are included as figures because the greatest decrease in stress and increase in fit occurred in going from one to two dimensions and because the three-dimensional solutions failed to provide any additional information. The average coefficient of alienation values for the 1-, 2-, 3- and 4-dimensional solutions were .39, .23, .15, and .10, respectively.

The standard way to interpret MDS solutions is just to look at them. The solutions for the semantic categories are easily interpreted. For instance, the *animal* category is consistent with earlier work performed on this domain (see Rubin & Olson, 1980, for a review). The *parts of the body* category, for which there was no previous work, is provided in Figure 2 as an example. This domain can be divided into a head cluster on the right, a limb cluster on the left, and a torso cluster at the center bottom. *Eyes* and *ears* do not appear immediately adjacent to *eye* and *ear* because no subject ever said "*eye, eyes*," or "*ear, ears*." They are, however, close to each other because these words were often said next to *nose* and *mouth*. The axes could be labeled as dimensions, but this is a stronger claim than is warranted by the data (Rubin & Olson, 1980).

The rhyme domains are harder to describe in terms of obvious organized clusters of words, though some structure is apparent. For instance, where homonyms occur, they tend to be near each other. There are 15 distinct homonym pairs in the rhyme category

MDS solutions. In 70% of the 320 cases in which a homonym pair was present in a subject's recall, the homonyms appeared next to each other. There is also a hint of some semantic structure in the rhyme categories, for instance, in the *ee* category, *we, he, me,* and *she* cluster. The *air* category is shown as an example in Figure 3.

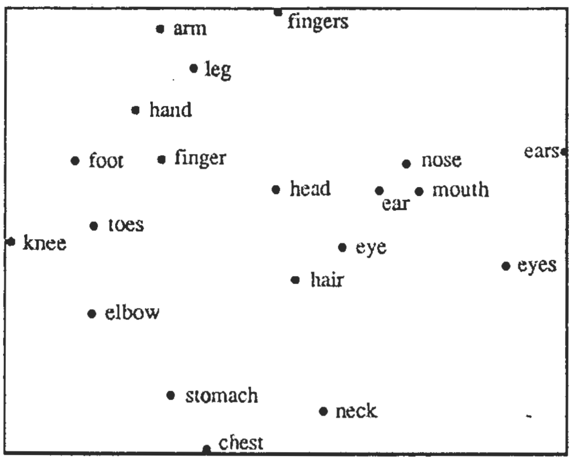

Figure 2. An MDS solution for the semantic category, *parts of the body.*

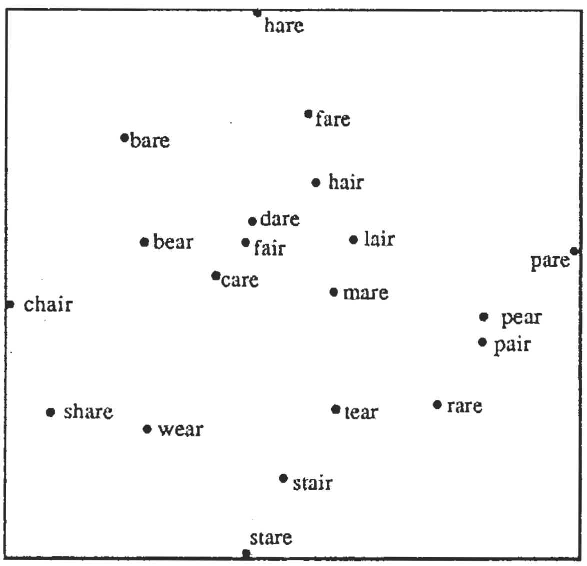

Figure 3. An MDS solution for the rhyme category, *air.*

Interpreting the Rhyme Spaces Quantitatively

The rhyme categories offer no simple visible interpretation, but perhaps it is possible that the structure of rhyme domains could be understood in terms of phonetic similarity. In order to investigate this possibility, the words listed in the rhyme categories were transcribed phonetically. To avoid making assumptions about the similarity of clusters of phonemes, only words consisting of a single phoneme plus the rhyme ending were included. This reduced the number of similarity values among the word pairs from 760 (i.e., the 20 × 19/2 comparisons among the 20 words of each of the four rhyme categories) to 516 similarity values. The pairings of homonyms were also removed to avoid confounding the special effects of their similarity with phonetic similarity. This further reduced the number of similarity values among word pairs to 502.

Two measures of phonetic similarity were formed. The first was the number of distinctive features distinguishing the two phonemes. This measure was taken from Fromkin and Rodman (1983). The second measure was the number of times two phonemes were confused in listening experiments. This was taken from the combined data of Tables 2 through 6 from Miller and Nicely (1955). Because all phonemes did not appear in these tables, the phonetic confusion measure could be obtained for only 204 of the word pairs.

It is also possible that the orthography of the words could have an effect. Therefore, in addition to the two phonetic measures, two parallel measures of visual properties were included for comparison. The number of distinctive features separating the initial letters was taken from Gibson (1969), and a visual confusion matrix was made by combining Tables 1 and 2 of Townsend (1971). Finally, the spelling of the rhyme sound was used to measure visual similarity. Word pairs with the same orthographic ending were assigned a value of 1.0, and word pairs with different endings were assigned a value of 0.0.

Two measures of similarity in memory were used. The first was the same number-of-times-next-to measure used for the figures. The second was this next-to measure divided by the number of subjects who recalled both words in the pair. The latter normalized measure was included as an additional check (Rubin & Olson, 1980).

None of the phonetic or spelling measures based on the initial phoneme or letter correlate with either similarity measure as highly as .07. The dichotomous measure of rhyme spelling has only small correlations with the next-to and normalized measures (.09 and .13, respectively). These correlations account for, at most, a negligible amount of structure. Thus, once the phonetic constraint of searching within a rhyme category is met, the phonetic and visual similarity of the words plays little further role.

Standard MDS solutions for the rhyme categories revealed no clear, interpretable structure as was found in the semantic domains, nor did phonetic or orthographic hypotheses, as tested by correlations. Some structure could be interpreted in terms of homonyms and semantics (e.g., *we, he, she, me*), but most of what appears in the MDS solutions cannot. Perhaps Nelson's claim is correct and the rhyme categories are not really structured. In order to test this possibility, a reliability check was performed to see if different groups of subjects recall the same items next to each other. The 100 subjects were divided into two groups of 50, and a correlation was calculated between the cells of the next-to similarity matrices that resulted from each group. As with the correlations performed on the rhyme categories, each matrix was treated as an ordered list of 190 (20 × 19/2) cells. The correlations for the four rhyme and the four semantic categories were .798, .730, .771, .794, and .935, .891, .881, .893, respectively, indicating that all eight of the domains were quite structured, though the semantic categories were more structured than the rhyme categories.

Another approach to measuring the structure of the rhyme similarity spaces is to examine the distribution of values in the cells of the matrix. If a next-to similarity matrix, such as those constructed here, was from a highly structured domain, it should contain many cells with high numbers and zeros, indicating pairs of words that often or never appeared next to each other. A similarity matrix from an unstructured domain, given the nature of random occurrences, should contain many moderate values. In order to provide values from unstructured categories that otherwise resemble the categories tested here, the recall protocols of each subject were randomized and similarity matrices formed from the random orders. Three random orders were sufficient to provide a smooth distribution of values for the unstructured matrices. The four rhyme and the four semantic matrices were combined and compared with their respective random order matrices as shown in Figures 4 and 5.

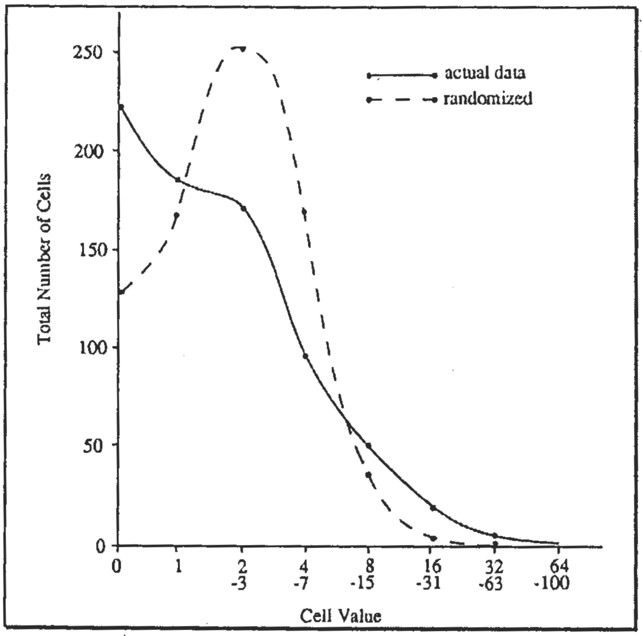

Figure 4. The distribution of actual and randomized next-to values in the cells of the semantic category similarity matrices.

The rhyme category appears to be highly structured. Compared to the random order matrices, the actual matrices had more cells with 0 and 1 entries, less cells with entries at each value between 2 and 10, and the same or more cells with entries at each value above 10. The semantic categories showed the same pattern with more cells with 0 and 1 entries, less cells with entries at each value between 2 and 9, and the same or more cells with entries at each value above 9. The highest value in any cell of the three random orders of the eight matrices was 22. The rhyme categories had 9 cells with values over 22, the semantic category had 11.

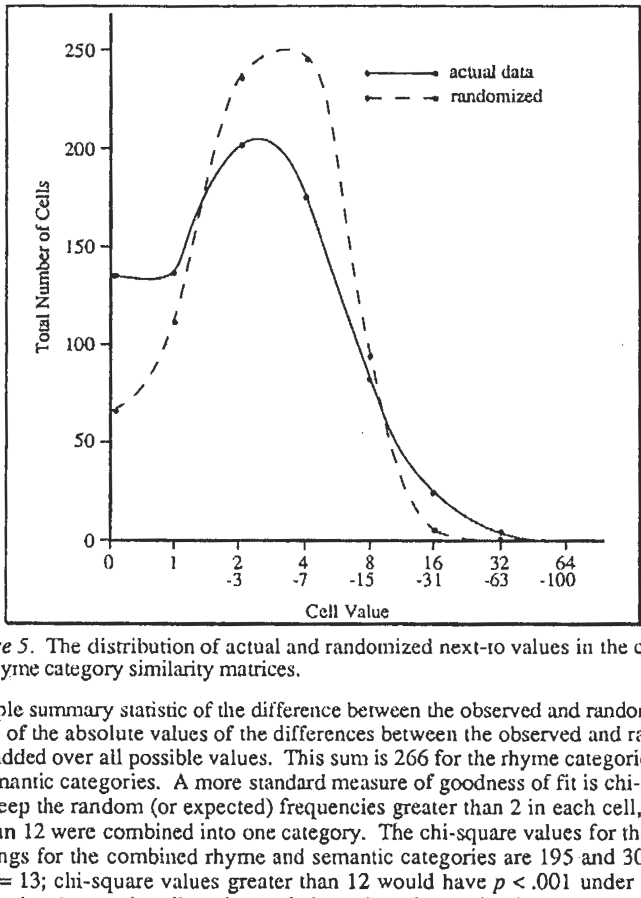

Figure 5. The distribution of actual and randomized next-to values in the cells of the rhyme category similarity matrices.

A simple summary statistic of the difference between the observed and random matrices is the sum of the absolute values of the differences between the observed and random frequencies added over all possible values. This sum is 266 for the rhyme categories and 324 for the semantic categories. A more standard measure of goodness of fit is chi-square. In order to keep the random (or expected) frequencies greater than 2 in each cell, all values greater than 12 were combined into one category. The chi-square values for the resulting 14 groupings for the combined rhyme and semantic categories are 195 and 309, respectively ($df = 13$; chi-square values greater than 12 would have $p < .001$ under the liberal assumption that the matrix cell entries are independent observations).

One difference between the rhyme and semantic categories is that rhyme categories have no cells for which more than 40 subjects out of 100 recalled words next to each other, whereas the semantic categories have two cells over 50. *Cat* and *dog* were recalled next to each other 89 times, and *lion* and *tiger* were recalled next to each other 53 times. It should be noted that neither the absolute sum nor the chi-square statistic is sensitive to the magnitude of these two high-cell values in that neither of these summary statistics would change if the two high-cell values were changed from their current values of 89 and 53 to values of 13. This is because the sum-of-the-absolute-differences statistic is not affected by which cell leads to the difference, and the chi-square statistic is calculated after all values greater than 12 are combined into one category.

In summary, it appears that both semantic and rhyme categories tested are structured, though there is more structure in the semantic categories. The structure of the semantic categories can be interpreted in terms of semantic similarity among the items. The structure of the rhyme category, with the exception of homonyms and occasional semantic structure, cannot be easily interpreted.

A Pathfinder Solution

The resources of standard MDS have been exhausted, and no answer has been found to the question of how rhyme categories are organized. Nelson (1981) hypothesized that there is no structure, but this is not the case here. On the other hand, there is no clear way to describe the structure that does exist. This research project remained a paradox for some time, until two occurrences combined to provide a resolution. First, Bruce Ammons, who was a graduate student at Duke at the time, observed that he had no strong organization for rhyme categories and so would do the task given to the subjects by using the alphabet to prompt himself to find rhyme words. Second, Pathfinder became available.

The task given to the subjects required the output of rhyming words. If no clear organization was available in long-term memory, perhaps some temporary search strategy, like using the alphabet to cue words, would have been used. This is a common strategy when trying to think of people's names, so Ammons's suggestion made intuitive sense. Moreover, the strategy is especially effective for rhyme categories. All the subjects would have had to do is combine each letter of the alphabet in turn with the rhyme ending and decide whether or not a word resulted.

An examination of the unanalyzed recalls supported the hypothesis for at least some subjects. The following two subjects' recalls from the *air* category are clear cases. The recalls begin with words formed from alphabetically ordered initial letters paired with the "air" rhyme. When the end of the alphabet is reached there is a search for more complex sounding words. The first subject recalled "*bear, care, dare, fair, hair, mare, pear, pair, rare, tare, tear, wear, share, Blair, flair, snare.*" The second subject recalled "*bear, care, dare, fair, hair, hare, lair, mare, pare, pair, rare, tear, where, impair, affair, there, stair.*" The problem remained of how to explore this possibility in a clear quantitative fashion.

If the alphabetical strategy was being used, the symmetrical next-to similarity measure would not be ideal; rather a measure that noted directionality is needed. For this reason a directional next-to, or follows, measure was formed from the data from the four semantic and four rhyme categories. The entry in each cell was the number of subjects who recalled the word defining the row followed by the word defining the column. Thus the cell *dare-care* contained seven because seven subjects recalled *dare* followed by *care*, whereas the cell *care-dare* contained 32 because 32 subjects recalled *care* followed by *dare*. This resulted in square rather than lower triangular similarity matrices, in which the main diagonal was undefined and set to zero for purposes of calculation. If these new matrices were collapsed along the main diagonal to form symmetrical lower triangular matrices by summing the values from cell (i,j) with cell (j,i), the previously analyzed eight matrices would result.

The square matrices were submitted to Pathfinder. Because Pathfinder requires dissimilarities, each cell was transformed by subtracting its value from 99. In addition, the resulting remainders that were greater or equal to 98 were considered as infinite, so that a word had to follow another word more than once in the set of 100 subjects' recalls in order to be counted in the solution. Minimally connected networks were obtained by setting the r value of the Minkowski r-metric to infinity. No limit was set on the possible number of links allowed in a path. Figures 6 and 7 present solutions to the *parts of the body* and the *air* domain. The placement of the nodes is identical to that in Figures 2 and 3. The links come from the Pathfinder solution.

Figure 6 reveals some interesting structure. The network divides into three major areas: the head, the extremities, and the body. The node *nose* provides a high-degree node for the head, whereas the node *head* provides a high-degree node for the network as a whole. The nodes *hand* and *leg* provide high-degree nodes for the extremities. The recall does not traverse the body in an orderly fashion based on location; rather, analogy seems to be the key. The pairs *fingers* and *toes*, *elbow* and *knee*, *arm* and *leg*, and *hand* and *foot* are all connected. The MDS solution on which the Pathfinder solution is superimposed does not provide this information as clearly.

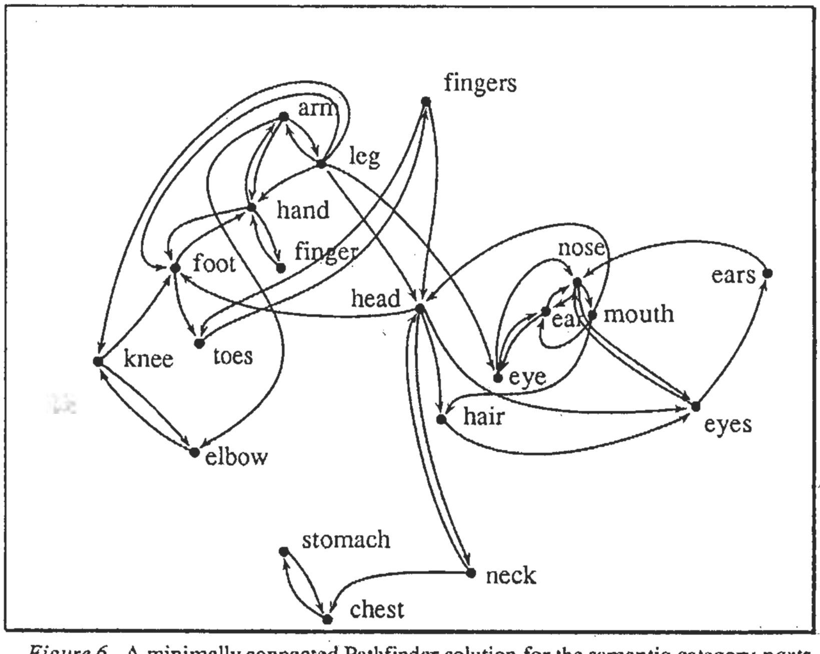

Figure 6. A minimally connected Pathfinder solution for the semantic category *parts of the body* superimposed on the MDS solution of Figure 2.

Figure 7 provides the solution to the question of the structure of the rhyme categories, something more traditional scaling techniques failed to do. Starting at the node *bare*, the network follows an alphabetical path. This can be seen more clearly in Figure 8, which is a replotting of Figure 7. This time, however, the solution and the figure include only those nodes that consist of a single letter followed by the *air* sound. A "U" was chosen instead of a straight line, partly because upon reaching the end of the alphabet the subjects at times return to the beginning for a second try and partly to make the figure clearer. Most of the cycles and jumps over nodes are caused by homonyms.

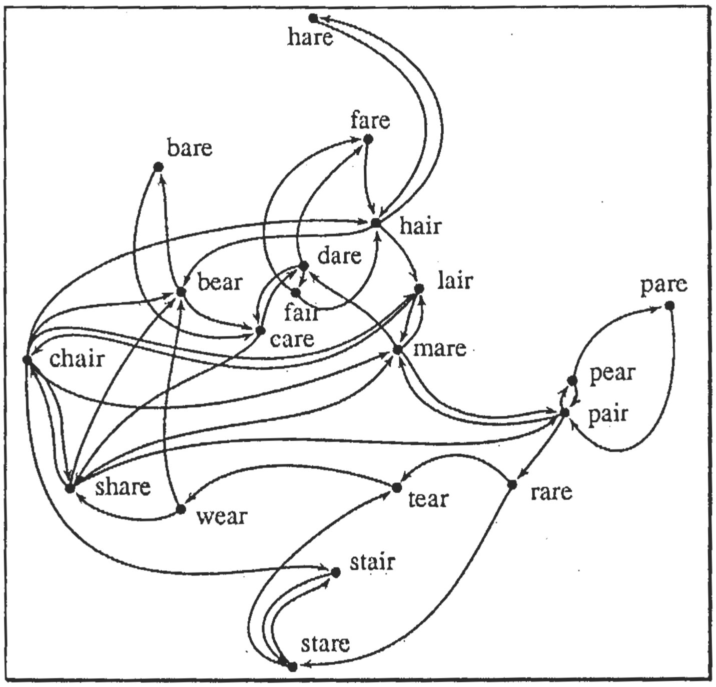

Figure 7. A minimally connected Pathfinder solution for the rhyme category *air* superimposed on the MDS solution of Figure 3.

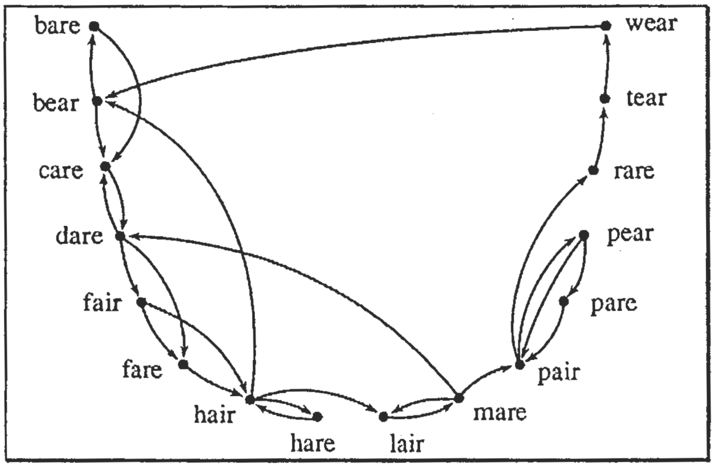

Figure 8. The Pathfinder solution of Figure 7 plotted to show the alphabetical organization of the words formed by a single letter followed by the *air* rhyme ending.

Figures 9, 10, and 11 present the nodes that consist of a single letter followed by the rhyme sound for the remaining three rhyme categories. First, in the figures for all four rhyme categories there is a clear alphabetical organization. Second, there are often words near the end of the alphabet linking back to the first alphabetical word of the category (e.g., *wear* to *bear*, *rear* to *bear*, *Ted* to *bed*, and *see* to *bee*). This probably results from recalls in which the alphabetical search started from the last of a series of easy-to-access words rather than from the beginning of the alphabet. Third, there are several frequently recalled words that serve as high-degree nodes in addition to their role in alphabetical search. The observation that some words are easier to recall than their alphabetical neighbors argues against a strict alphabetical search; words recalled more frequently than their alphabetical neighbors must have links that are not alphabetical. In particular, the two words in each of the four rhyme categories that were most frequently recalled each have five or more links with at least one of those links to a word not near it alphabetically. For Figures 8, 9, 10, and 11 these high-degree node words are *hair* and *dare*, *fear* and *near*, *bed* and *said*, and *see* and *me*. Fourth, there are other assorted associations that are consistent among subjects but that do not result from alphabetical search. The third and fourth types of links indicate that forms of search or association other than alphabetical search are also functioning.

Thus, Nelson's (1981) hypothesis based on cuing data is supported. Undergraduates seem to have no interpretable fixed structure for rhyme domains, but rather produce rhyme categories when needed using an algorithm based on the alphabet with additional rules to look for homonyms and to take advantage of any semantic structure that exists. Such constructions of semantic categories have been noted by Barsalou (1983, 1987) for ad hoc categories, such as *things to take on a camping trip*. Figures 8, 9, 10, and 11 reveal that such constructive techniques are used even for the major rhyme categories.

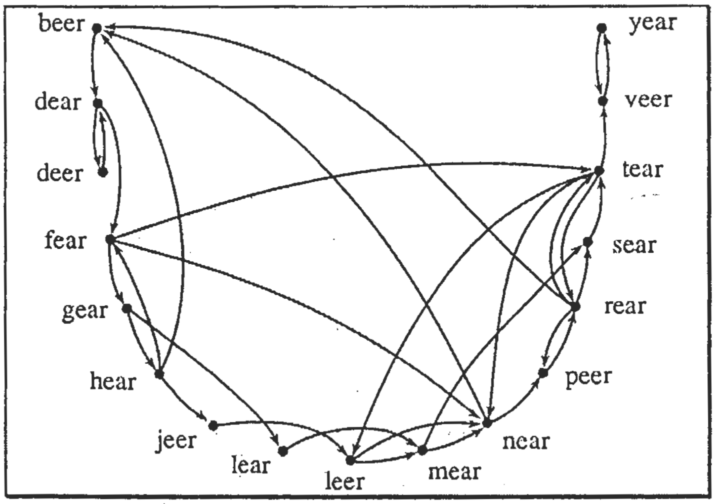

Figure 9. The Pathfinder solution for the rhyme category *ear* plotted to show the alphabetical organization of the words formed by a single letter followed by the *ear* rhyme ending.

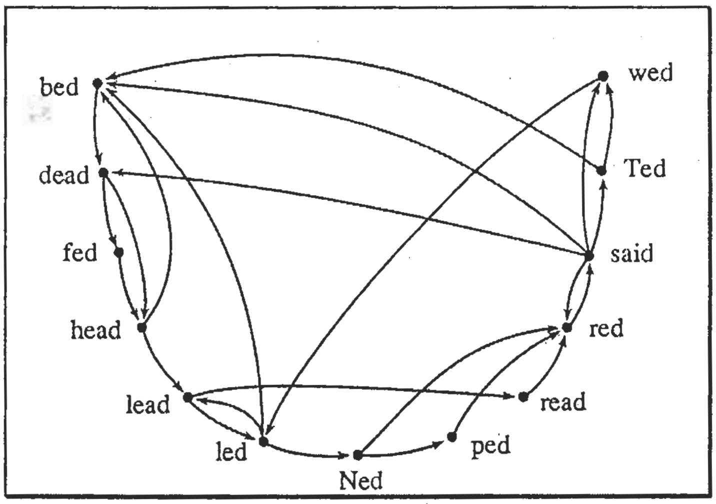

Figure 10. The Pathfinder solution for the rhyme category *ed* plotted to show the alphabetical organization of the words formed by a single letter followed by the *ed* rhyme ending.

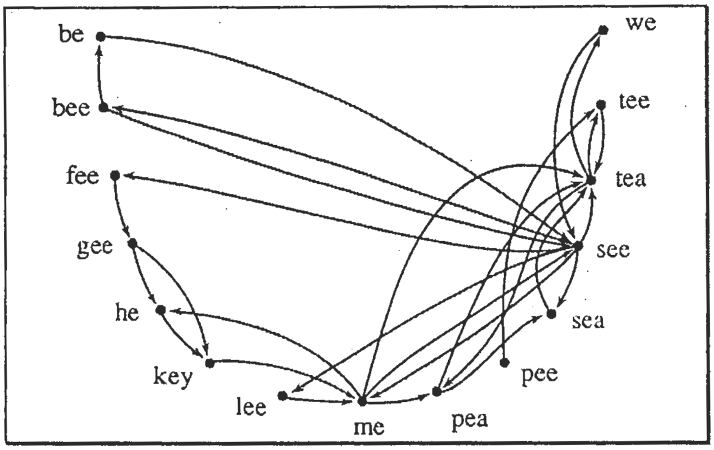

Figure 11. The Pathfinder solution for the rhyme category *ee* plotted to show the alphabetical organization of the words formed by a single letter followed by the *ee* rhyme ending.

Directed Graphs as Traces Through Memory

The search for the structure of rhyme categories demonstrates that Pathfinder offers many advantages for representing memory structures, as well as for uncovering retrieval strategies. The first major advantage of Pathfinder is that it allows associations among concepts to be asymmetric (i.e., the links among nodes to be directional). In a spatial analogy, such as MDS, the distance from point *a* to point *b* has to be the same as from point *b* to point *a*. In general, there is considerable evidence that this is not always the case for associations among concepts (Anisfeld & Knapp, 1968; Deese, 1965; Ekstrand, 1966; Rubin, 1980, 1983; Rubin & Friendly, 1986; Thorndike, 1932). In particular, for the rhyme data analyzed in this chapter, it is clear that a word is much more likely to lead to the word that follows it alphabetically than to the word that precedes it. The directional arrows in Figures 6, 7, 8, 9, 10, and 11 reveal regularities that the nondirectional distances in Figures 2 and 3 do not.

The second major advantage of Pathfinder is that it concentrates on local structure and does not try to maximize the fit for all possible links among all possible nodes. Semantic memory has a local as opposed to a global structure. Cluster size is very small (Gruenewald & Lockhead, 1980). There is great regularity in which item follows immediately after which in recall, but much less in what item follows five items later (Rubin & Olson, 1980). Few items will be recalled between *cat* and *dog*; the number of items recalled between *cat* and *cow*, however, is not as clear. That is why the next-to measure was used instead of a measure based on the number of intervening items (e.g. Friendly, 1977). Pathfinder makes use of this local structure by extracting the shortest links in paths and ignoring the longer, high-distance, low-similarity links. In contrast, MDS weights the discrepancy between the data and the fit equally for high and low distances (or similarities).

The emphasis on local as opposed to global structure may not be an advantage in all applications, but it is here because it mimics the structure in the data.

The third major advantage of Pathfinder is that its solutions are a natural analogy for search through memory. Figures 6, 7, 8, 9, 10, and 11 look like the directional trace of a path through memory. In fact, they are a kind of average of the order of output of all subjects. MDS solutions can be viewed as maps of the relationships among items; Pathfinder networks can be viewed as paths traversed through memory. Instead of showing us what could be seen as a static picture of a memory representation, Pathfinder shows us what could be seen as a record of the search process, that is, the record that memory traces.

Pathfinder Associative Networks: Studies in Knowledge Organization

Edited by

Roger W. Schvaneveldt
Department of Psychology and
Computing Research Laboratory
New Mexico State University

ABLEX PUBLISHING CORPORATION
Norwood, New Jersey 07648

Printed in the United States of America.

Library of Congress Cataloging-in-Publication Data

Pathfinder associative networks : studies in knowledge organization / edited by Roger W. Schvaneveldt.

p. cm. – (Ablex series in computational sciences)
Includes bibliographical references.
ISBN 0-89391-624-2
1. Expert systems (Computer science). 2. Artificial intelligence. 3. Graph theory.
4. Knowledge acquisition (Expert systems).
I. Schvaneveldt, Roger W. II. Series.
QA76.76.E95P28 1989
006.3'3–dc20 89-18219
CIP

ABLEX Publishing Corporation
355 Chestnut Street
Norwood, New Jersey 07648