



Chapter 2

Graphs in the Social and Psychological Sciences: Empirical Contributions of Pathfinder*

Francis T. Durso and Kathy A. Coggins

Graphs and graph theory play important roles in the theories and methods of many sciences. As we have seen (Dearholt & Schvaneveldt, Chapter 1, this volume), a graph is a mathematical formalism and so may be used to represent a wide range of phenomena. Chemical isomers, electrical circuits, Markov chains, statistical mechanics, and network flow in operational research are but a sampling of fields in which graph theory has been useful (Harary, 1969). It is this abstract character that allow graphs to have such general utility. Formally identical graphs could represent the distances from hospitals to homes in a neighborhood, the flow capacity of water through a city, the frequency of information exchange among or within organizations, or the associative network of concepts in human long-term memory. Graph theory allows for the identification of the shortest path between two locations, the identification of the most efficient circuit (i.e., the shortest closed path) in an electrical diagram, the isolation of cliques (i.e., completely connected subgraphs) in a social organization, the prestige (i.e., indegree) of members of a social group, or the hub (i.e., center) of government bureaucracy. In addition, analyses at more microscopic levels allow for determination of particular associations, channels, or influences within a graph.

The value of graphs has not gone unrecognized by social and psychological scientists. In psychology, one does not have to look far for a theory that uses some form of graph as its cornerstone. Clearly, in cognitive psychology, graphs are present in force, from neosociative networks of memory (Anderson, 1983), through propositional analysis of discourse (van Dijk & Kintsch, 1983), to connectionist networks of cognition (McClelland & Rumelhart, 1986; Rumelhart & McClelland, 1986). Whereas cognitivists have perhaps been the most voracious consumers of graph theory in psychology, other areas of psychology have successfully employed graph theoretic concepts. In social psychology, we find theories that utilize graph structures to account for communication patterns within groups (Shaw, 1981). In fact, social psychology's theories of dissonance (Festinger, 1954) and balance theory (Harary, Norman, & Cartwright, 1965) rest on principles of graphs.

Graph theory also has played a role in psychological studies of infrahuman species. Tolman's (1932) notion of a cognitive map has been pursued in the recent work of Liebllich and Arbib (1982), who argue explicitly that animals learn a graph in the most abstract sense of the term. In fact, Brown (1987) provided empirical evidence that the behavior of rats in a radial arm maze is based on the formal graph constructed of the area and not on more local cues to food (such as the texture of the runway).

*The authors thank Deborah Bates-Porter for her help in preparing this chapter and Tom Dayton, Peder Johnson, Wendy Shore, and Mike Betts for their comments on an earlier draft of this chapter. Preparation of this chapter was supported by a grant from the University of Oklahoma.

Other social scientists besides psychologists have also made use of graph-theoretical constructs. For example, sociologists interested in the transfer of commodities (e.g., information, money) from one institution to another have made significant contributions in the application of graph theory to social issues (e.g., Knoke & Wood, 1981). In a related vein, sociologists and psychologists interested in "networks" have helped adapt concepts from the mathematics of graphs to applied domains (e.g., Burt & Minor, 1983). In fact, these efforts by sociologists have been more empirically based than those by psychologists. Network analysis¹ in sociology concerns itself with the collection of data from which networks could be determined and with the ultimate interpretation of those networks. In psychology, graphs are often constructed top-down from the intuitions or theories of the researcher, and thus are rarely empirically derived (cf., Chi & Koeske, 1983; Fillenbaum & Rapoport, 1971; Hutchinson, 1981).

In this chapter, we selectively review research from a number of areas that rely on graphs. Some of these endeavors have utilized Pathfinder, and for these we attempt to highlight the value of this scaling algorithm and report comparisons with other scaling algorithms when appropriate. Other endeavors have not been informed by the Pathfinder algorithm and for these we first review the work and then attempt to illustrate the value of Pathfinder by applying it to a relevant dataset. It is interesting to note that it was often difficult to find data that would allow Pathfinder to demonstrate its full potential. Some data were available that presented a trivial task to Pathfinder: For example, many sociograms are merely matrices of 0's and 1's, for which Pathfinder would simply link the nodes corresponding to the entries of 1 and not link those with a 0 for the cell entry. It seemed clear, at least to us, that many of the issues addressed in this manner could also have been easily addressed with more sensitive scales if there existed a method capable of taking advantage of the increased sensitivity. Pathfinder is one such method.

The output of Pathfinder is a PFNET that can be uniquely specified by two parameters: r and q . The r parameter is the Minkowski exponent. With an exponent of infinity, Pathfinder makes only ordinal assumptions about the data. In this chapter, all of the reported PFNETs used an r parameter of infinity. The second parameter, q , is a restriction on the number of edges in a path that Pathfinder will use in deciding if two concepts are already connected. The sparsest PFNET will result when Pathfinder is permitted to consider paths of any length, that is when q is equal to one less than the number of nodes. The most dense graphs result when Pathfinder can only consider a path as consisting of two edges, that is $q = 2$. This PFNET ($\infty, 2$) is identical to solutions produced by NETSCAL (Hutchinson, 1981). Although decisions about the r parameter can be justified on measurement assumptions, the decision concerning q is more difficult. There is, currently, no formal mechanism for choosing among values of q for a given r . This is a situation similar to deciding on the appropriate dimensionality of a multidimensional space. Both when picking q and when picking the dimensionality, several factors, including the illuminating power of the solution, must be considered. In this chapter, we have a bias toward the sparsest solution ($q = n - 1$), especially when the graph is directed. However, when decreasing q provided additional insights, we did not hesitate to report that solution.

We note at the outset that our purpose is to review the contributions of Pathfinder as another tool, although we believe a very valuable tool, for the analysis of several issues that lend themselves to graph-theoretic analysis. In adopting this purpose we often do not do justice to the methods of analysis originally employed in the area, and we often skirt some

¹ Much of the sociological work concerns itself with identifying the subgroups within a graph. Subgroups can be identified either by finding the cliques in a graph or by determining what nodes in a graph are structurally equivalent.

of the theoretical complexities. These omissions will, of course, be obvious to researchers active in these areas. However, we believe that these researchers will also see the value of Pathfinder for their areas, perhaps more clearly than would have otherwise been the case.

We begin by considering a number of uses of graphs in cognitive domains. In particular, we review how graphs have been used to represent categories, to understand the representation of expertise, to predict details of human memory performance, and to design artifacts more compatible with human information processing.

Following the discussion of graphs in cognitive psychology, we consider the use of graphs in social domains. The graph analysis of these social phenomena, unlike the cognitive phenomena we consider, have not employed Pathfinder. For these areas, we discuss briefly the original analysis and then attempt to apply Pathfinder to the issue. Our application of Pathfinder to these areas should be viewed as, at best, a demonstration of how the algorithm could be of some assistance. We begin with a study of information and money exchange among a number of institutions in Indianapolis. We follow this with an analysis of the friendship graphs of a class of 8th graders. We end this section by considering how graphs have been utilized to understand small group communication dynamics and speculate on how Pathfinder could prove a valuable aid.

Cognitive Graphs

Knowledge Structures

The first work with Pathfinder was with an eye toward determining the types of structures that Pathfinder could, or would tend to, produce. This early work focused on the representation of knowledge. This choice followed naturally from the interests of Pathfinder's developers and also proved a fortunate choice in that clear differences among the PFNETs were observed that fit nicely with past research and theory.

Natural concepts. The first use of Pathfinder was presented at the meetings of the Psychonomic Society (Schvaneveldt & Durso, 1981), and involved a PFNET of 25 natural concepts. Several theorists posit an associative network in semantic memory, but they typically rely on intuition to construct the network (e.g., Collins & Loftus, 1975). Schvaneveldt and Durso reasoned that a minimal requirement for Pathfinder would be to capture many of the intuitive relations one would expect to hold among the concepts. The concepts (see Figure 1) included some that stood in a superordinate-subordinate relation, whereas others were related at the same level in the hierarchy. Some concepts, intuitively, had relations with several concepts across the network, whereas others had more specific relations. That first effort was very encouraging. Pathfinder reduced the 300 (25 items taken two at a time) pairwise similarity ratings considerably. The sparsest graph, PFNET($\infty, 24$), contained 25 links; even the most dense graph, PFNET(1, 2) was a reduction to 119 links. PFNET($\infty, 2$) had 32 links and appears in Figure 1. The connections certainly did not violate intuitions and in fact revealed a number of interesting relations.

For example, Figure 1 shows that some nodes played a restricted role in the graph (e.g., *hooves*), whereas others enter into categorical and property relations (e.g., *green*). Further, typical members of a category tended to connect directly to the category, whereas atypical members tended to connect indirectly with their superordinate category. The category *mammal* is especially interesting, in part because of its history in semantic memory research (see Rips, Shoben, & Smith, 1973; Smith, Shoben, & Rips, 1974). The suspicion of some, that *mammal* is not psychologically a natural category, seems to receive some support here. In fact, when biology graduate students rated the same items, *mammal*

played the more central role that the scientific taxonomy predicts (Schvaneveldt, Durso, & Dearholt, 1989). Overall, the graph was simple and the relations were interesting and consistent with intuitions.

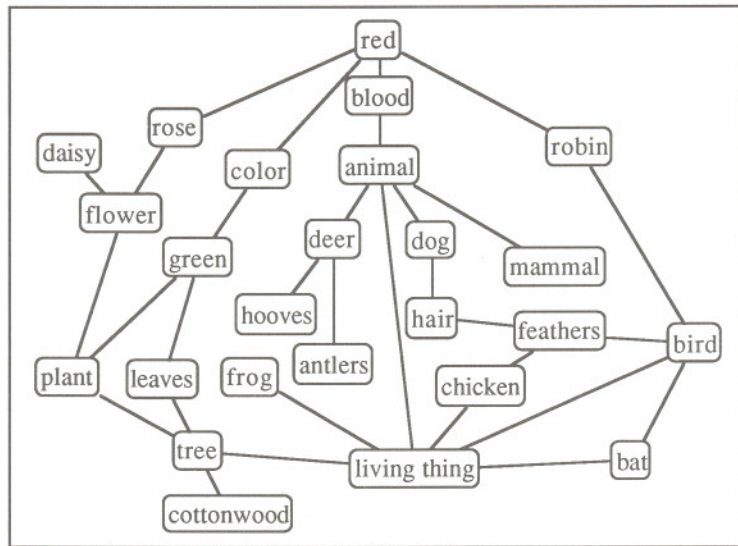


Figure 1. The most dense PFNET assuming ordinal data ($r = \infty$) for the 25 natural concepts used by Schvaneveldt and Durso (1981).

Basic-level categories. Following this initial success, we investigated a number of cognitive structures that we thought might produce interesting structures and at the same time extend the earlier work of others. One area derived from Rosch's contention (e.g., Rosch, Mervis, Gray, Johnson, & Boyes-Braem, 1976) that all Aristotelian natural categories were not created equal (at least not psychologically). Rosch and her colleagues have supplied a good deal of evidence to support the idea that some categories, such as basic-level categories, are psychologically special. For basic-level categories it is the world (not the language, cf., Whorf, 1956) that articulates thought.

Schvaneveldt et al. (1989; see also Hutchinson, 1981) used existing association norms (i.e., Cohen, Bousfield, & Whitmarsh, 1957; Marshall & Cofer, 1970) to establish matrices of associations for the six categories used by Rosch. According to Rosch, *fish*, *bird*, and *tree* are basic-level categories, whereas *clothes*, *fruit*, and *musical instruments* are not. In addition to this theoretical motivation, the association matrices were of interest to the development of Pathfinder because they were asymmetric: *Bird* is more likely to be a response to *thrush* than *thrush* is to the stimulus *bird*. Traditionally, when the judgment a_{ij} differed from a_{ji} , the scaling algorithm assumed that both were measures of the same underlying distribution and that the difference was due to noise. Recognition of the psychological reality of asymmetries (e.g., Tversky, 1977), however, makes it clear that an algorithm that could handle asymmetries in a meaningful way would be a useful step in understanding how categories are structured.

Figure 2 presents the sparsest PFNETs ($\infty, n-1$) for Rosch's six categories. The "starness" of the graphs is apparent for the basic-level categories. Starness is easily

quantified by dividing the degree of each category node by the total number of arcs in the graph. Table 1 shows the correspondence between this index and Rosch's original classifications. This starness index also suggested that the category *flower* could have been treated as a basic-level category by Rosch, but that *professions* could not.

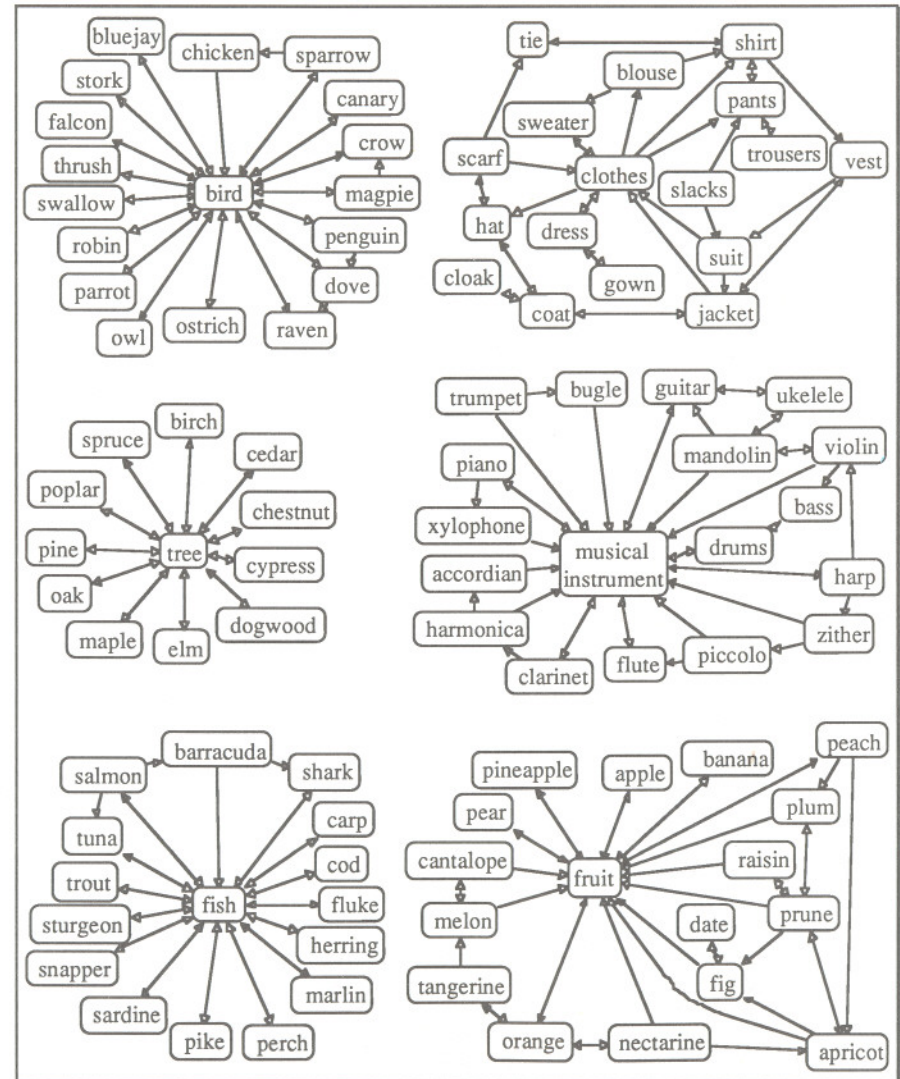


Figure 2. The sparsest directed PFNETs of three basic-level categories (left) and three superordinate-level categories (right).

We should note that applying a spatial algorithm to these data failed to distinguish between the basic-level and the superordinate categories. Although category labels tended to fall at the center of a two-dimensional space, this was true of all categories. Perhaps the discrete nature of language and categories makes a graph-theoretic model like Pathfinder particularly appropriate. However, as the next section suggests, Pathfinder appears to contribute even in domains where spatial algorithms have proven quite successful.

Table 1. Starness index for eight natural categories.

Category	Classification	Starness
Fish	Basic ^{a,b}	.90
Bird	Basic ^{a,b}	.89
Tree	Basic ^{a,b}	1.00
Musical Instrument	Superordinate ^{a,b}	.56
Fruit	Superordinate ^{a,b}	.50
Clothes	Superordinate ^{a,b}	.31
Flower	Basic ^b	1.00
Professions	Superordinate ^b	.50

^aRosch's classification ^bstarness classification

Concepts with underlying dimensions. Pathfinder, like other algorithms that produce graphs, is likely to have its greatest success in representing discrete concepts. We were interested in how Pathfinder would perform when the stimuli varied along underlying dimensions. With this in mind we looked at (a) judgments of color borrowed from Ekman (1954) and used by Shepard (1962) to reproduce the Newton Color Circle with multidimensional scaling; (b) judgments of words signifying length of time (e.g., *second*, *minute*) that we collected from undergraduates and that intuitively should fall on a single underlying dimension; and (c) judgments of restaurant-script concepts (Maxwell, 1983).

As Figures 3 and 4 attest, Pathfinder revealed interesting structures even though a spatial algorithm might have been the a priori procedure of choice. For the Ekman data, the sparsest PFNET (∞ , 13) mirrored the physical wavelengths, but the PFNET (∞ , 2) solution produced the Newton Color Circle. It is interesting to note that when Shepard first presented his Multidimensional Scaling (MDS) solution to the Ekman data, he connected the terms to highlight the structure, producing the graph that Pathfinder produces algorithmically. The length-of-time terms fell neatly onto a simple path that captured the logical relations among the concepts. Finding this simple dimension was not a trivial exercise given that a one-dimensional MDS solution for the same data did not preserve the logical ordering of concepts.

Finally, the temporal dimension presumed to underlie scripts is apparent in the PFNET of Maxwell's (1983) data (Figure 5), augmented by a number of interesting cycles. The cycle involving paying the bill was particularly appealing to the students in our classes: Perhaps it allowed for the possibility of not leaving a tip (although it probably represents a difference between "paying the cashier" and more formal eating establishments).




Figure 3. A PFNET capturing the Newton Color Circle. The solid links are from the sparsest PFNET and the dotted link is the single link added by the most dense ordinal PFNET.

These initial efforts demonstrated the viability of Pathfinder. What was discovered across a number of knowledge domains fit intuitions and conformed to previous theory and data. These initial PFNETs continue to be of particular interest because they highlight differences among the types of graphs that could emerge from the Pathfinder algorithm: a single cycle for the color data, a single path for the time-duration data, star-patterns for the basic-level data, as well as more general graphs for the natural concepts and the script data.

These PFNETs highlight one of the assets of Pathfinder. That is, looking at the graph in its entirety, rather than at only subgraphs, can often give additional insights into the domain under scrutiny. For example, the sociometric study of graphs is often restricted to the discovery of substructures like cliques, but consideration of the "big picture," the Gestalt of the graph, has been difficult. This does not imply that a study of subgraphs and other graph-theoretic summaries are without value. However, without the use of a device like Pathfinder to reduce the data to a tractable graph, the researchers do not have the freedom to look at both the overall graph structure and at the subgraphs, but instead are restricted to analyses of only the more manageable substructures.

Finally, Pathfinder can supply information that multidimensional scaling does not. The failure of MDS to capture the logical relations among the time-length terms and the failure of MDS to distinguish between basic-level and superordinate-level categories suggests that Pathfinder can, at least, complement spatial analyses. In some cases, Pathfinder provides information akin to that of a nearest neighbor analysis (Tversky & Hutchinson, 1986).




Figure 4. The sparsest PFNET for the time-duration terms.

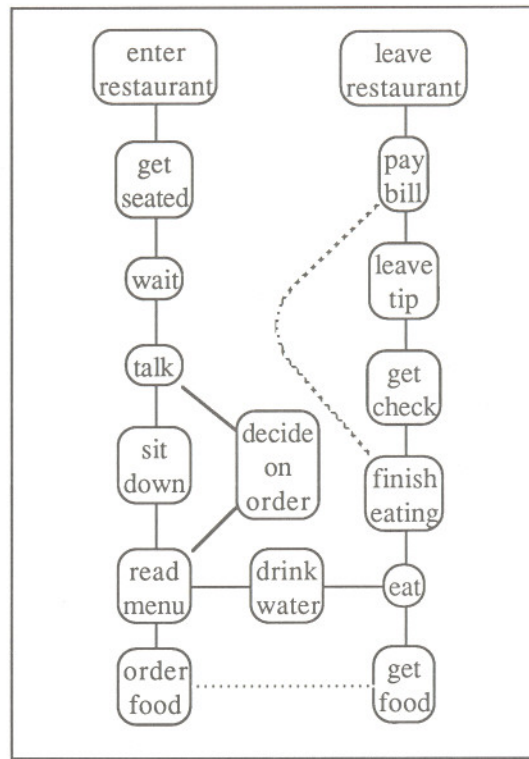


Figure 5. A PFNET for terms from a restaurant script.

Expertise

Differences among knowledge structures have been implied to underlie a number of phenomena in cognitive psychology. This implication has been strongest in explanations of expertise. From the seminal work of Chase and Simon (1973), through the numerous studies it inspired (e.g., McKeithen, Reitman, Rueter, & Hirtle, 1981), to the recent work of Chi and associates (e.g., Chi & Koeske, 1983), the underlying structure of knowledge has been the theoretical focus.

For example, Chi and Koeske (1983) explicitly constructed graphs from the protocol of a 4-1/2-year-old child with an avid interest in dinosaurs. Although Chi and Koeske did not directly compare experts and novices, they did construct separate graphs for the better-known dinosaurs and for the lesser-known dinosaurs. The nodes associated with the better-known dinosaurs had higher degrees (i.e., more *dinosaur-dinosaur* connections) and stronger linkages than did the lesser-known dinosaurs. Finally, the graph of better-known dinosaurs was apparently more cohesive in that the subgroups defined in that graph were characterized by stronger within-group paths and weaker between-group paths than were the subgroups of the "novice" graph.

2 Empirically Derived Graphs

Several studies lead to the conclusion that experts understand at a deeper, more abstract level than do novices (e.g., Adelson, 1981; Chi, Feltovich, & Glaser, 1981). This, however, does not require that the knowledge of experts and novices differ structurally. For example, experts have more domain-specific knowledge than do novices, but beyond this the knowledge may be organized or configured in the same way. A more compelling demonstration of the role of structural differences would be to demonstrate that a common set of concepts are interrelated differently for experts and novices.

Some research has attempted to discern directly if structural differences can be associated with superior knowledge. For example, McKeithen et al. (1981) showed that multi-trial-free recall of ALGOL-W reserved words by expert programmers led to structures apparently more tree-like than was the knowledge of less expert programmers.

This literature suggested that Pathfinder could help provide evidence for the role of structure in our understanding of expertise. With Pathfinder it is relatively easy to present a common set of items to experts and novices, obtain judgments of relatedness, and then construct graphs for the two groups.

Schvaneveldt, Durso, Goldsmith, Breen, Cooke, Tucker, and DeMaio (1985) selected 30 concepts from each of two air-combat situations: split-plane maneuvers (air-to-air) and strafe maneuvers (air-to-ground). All possible pairs of these concepts (435 from each domain) were judged (on a scale of 0 to 9) by members of the U.S. Air Force or members of the Air National Guard. The Air Force pilots differed in their level of expertise (e.g., flight time) with some being undergraduate pilot trainees and others serving as their instructors. The Guard Pilots were all expert. PFNETs were constructed for each group and for each pilot. Schvaneveldt et al. (1985) presented a number of different analyses, but here we discuss two that we feel highlight the value of Pathfinder particularly well.

One analysis investigated the extent to which an individual pilot could be classified as an expert or a novice based on his or her cognitive representation.² The analysis began by constructing graphs for each pilot. These graphs for individuals were dense. This *density* highlights the consequence of ensuring a unique graph: When there are several ties in the data (as when graphing the data of individuals using a limited scale), Pathfinder will not arbitrarily discriminate between two (or more) possible edges of the same weight, but will instead include both (or all). Ties in the data are, of course, rarely a problem when summary data are submitted to Pathfinder. Further, it could be reduced in the construction of individual data, perhaps by using magnitude estimation procedures (Stevens, 1975) rather than more traditional Likert scales.

The density of the graphs notwithstanding, Schvaneveldt et al. computed three types of patterns for each individual: graph patterns indicating the presence or absence of a link, MDS patterns of the distance in *k*-dimensional space between each pair of concepts, and the original empirical ratings. The questions became: Does the adjacency information provided by Pathfinder discriminate between experts and novices? And if so, is this discrimination superior to that which could be accomplished by pilots' relatedness judgments?

Nilsson's (1965) pattern recognition algorithm was used to define two prototypes (e.g., undergraduates vs. instructors, instructors vs. guard pilots) based on a subset of the matrices. The algorithm then classified a pilot who was not used to constructing the prototype by indicating to which prototype the "unknown" pilot belonged. Schvaneveldt et al. (1985) repeated the procedure until every pilot served as the to-be-classified pilot; the percent correct classifications were then calculated for the population of decisions. Figure 6 shows that discrimination was quite successful using Pathfinder. Further, discrimination

²Cooke and Schvaneveldt (1988) had the same intent when they analyzed the PFNETs of expert, intermediate, novice, and naive programmers.

was superior using the Pathfinder network compared with the original rating data. Despite a possible ceiling effect, it is interesting to note that MDS classification was at least as accurate as Pathfinder and was superior in some cases.




Figure 6. Comparative classification success of Pathfinder, MDS, and the original ratings in a study of fighter pilots. IP = Instructor Pilot, UP = Undergraduate Pilot Trainee, GP = Air National Guard Pilot. From "Measuring the structure of expertise" by Schvaneveldt et al., 1985, *International Journal of Man-Machine Studies* 23, p. 717. Copyright 1985 by Academic Press Inc. (London) Limited. Adapted with permission.

The second analysis of interest focused on an attempt to establish the relationships critical to expertise. Schvaneveldt et al. reasoned that if a link was present in the graph of one group of experts, but not in the graph of the other group of experts, then that connection must not be necessary to the cognitive structure of expert fighter pilots. The graph of the links that were shared by the two groups of experts (the "right stuff," see Figure 7) constructed by Schvaneveldt et al., and its comparison with the novice graph, did receive some validation. Concepts that were particularly poorly understood by the undergraduate pilots (i.e., those with few connections in common with the experts) were isolated and then used to classify individuals as described above. This set of only 10 "misunderstood" concepts perfected the novice-expert classifications: 100% of the novices and experts were correctly classified.

In summary, there is some evidence that experts can be distinguished from novices based on their cognitive structures. Classifications based on Pathfinder were superior to classifications based on the rating data suggesting that Pathfinder was successful at uncovering the latent structure inherent in the empirical ratings. Thus, a comparison of experts with novices supplies some validation of the psychological utility of Pathfinder.

To the extent that Pathfinder can capture important structural aspects of the expert's knowledge, it presents an interesting methodology that could be used to assist in solving important applied problems that rely on an understanding of human expertise. Cooke and McDonald (1987) and Schvaneveldt and Goldsmith (1985) have both pursued the implications of Pathfinder for artificial expertise: The former have focused on knowledge elicitation for use in expert systems and the latter have used empirically derived graphs as a basis for ACES, an air-to-air combat simulation.

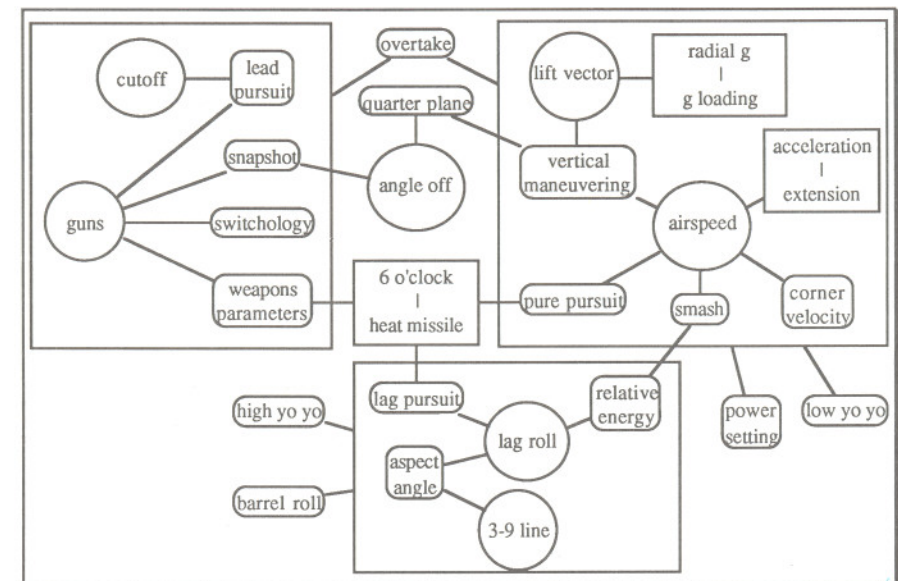


Figure 7. A graph of those links shared by Instructor Pilots and National Guard Pilots. The circle nodes constitute the minimal dominating set of concepts. From "Measuring the structure of expertise" by Schvaneveldt et al., 1985, *International Journal of Man-Machine Studies* 23. Adapted with permission from Academic Press Inc. (London) and the authors.

Memory

We now turn to a study that supplies further validation of Pathfinder, this time by taking advantage of the fact that knowledge structures have an impact on memory and recall. The differences in structure that apparently exist in our underlying knowledge should also be revealed in how subjects recall events. This contention has a long history in the study of human memory. Indices of underlying structure (e.g., meaningfulness) were used extensively by verbal learning theorists; a number of learning and cognitive theorists have shown relations between organization and recall; and most recently researchers have explored how various scaling solutions predict recall. For example, the work of Chi and Koeske (1983) mentioned earlier included a demonstration that recall was superior for the dinosaurs that came from the more cohesive graph.

Analyses of recall have also been conducted using more formal scaling procedures including hierarchical cluster analysis and multidimensional scaling. Friendly (1977) and others (e.g., Caramazza, Hersch, & Torgenson, 1976) have explored this approach extensively; the interested reader is referred to Friendly (1979) for a very readable review of a myriad of procedures. Of interest to our review of Pathfinder is that multidimensional scaling solutions apparently predict the organization in recall quite well.

Cooke, Durso, and Schvaneveldt (1986) compared Pathfinder with MDS-derived structures on their ability to predict the recall of a list of words. Subjects studied lists of words chosen from the natural concept experiment (Schvaneveldt & Durso, 1981) that we

discussed earlier. The lists were organized according to different scaling solutions. A list could be tightly organized according to Pathfinder but not according to MDS (Pathfinder list); it could be organized according to MDS but unorganized according to Pathfinder (MDS list); or a list could be unorganized according to either scaling procedure. The Pathfinder list was constructed by selecting items that were linked in the graph but distant in MDS; the MDS list comprised a sequence of items that were close in multidimensional space but not linked in the graph; and the unorganized lists comprised items that were neither linked nor close.

The results of the serial learning tasks were quite dramatic. Subjects who studied the Pathfinder lists learned more quickly than those who studied the MDS lists. In fact, on some indices, learning the MDS list was as poor as learning an unorganized list. Despite these dramatic effects, it could be argued that serial learning was a task especially well suited to Pathfinder. Unlike MDS, which focuses on a global configuration, Pathfinder tends to focus on more local relations. The highly related concepts are especially important to the Pathfinder algorithm. In a task such as free recall, the advantage of Pathfinder might be lost.

Cooke et al. presented the items in a random order and then tried to predict the recall order. Predictor variables were either based on Pathfinder, MDS, or the unscaled empirical ratings. The predicted variable was the average interitem distance in the recall protocols.

Cooke et al. found, as have others (Caramazza et al., 1976), that MDS predicted recall order. They also found that Pathfinder and the original ratings predicted recall order. More interesting than these simple correlations were the partial correlations that separated the contribution of the scaling procedure from the predictive power inherent in the ratings. If the purpose is to predict recall order, what is gained by transforming the data using either MDS or Pathfinder?

When the contribution of the original ratings were partialled out of the correlation, Pathfinder, but not MDS, proved independently predictive of recall order. These findings have two important implications. First, the reports that MDS could predict recall order may have been based on the predictive power of the proximity judgments, with MDS adding little. Second, and more important, is that Pathfinder revealed some latent structure useful for predicting recall order beyond that which could be linearly predicted from the ratings.

Other investigations of Pathfinder in episodic memory tasks appear in this volume. Branaghan (Chapter 8) presents work with paired-associates and Goldsmith and Johnson (Chapter 17) investigate memory organization as a function of classroom experience.

Human-Machine Interaction

In addition to being able to predict learning and recall, some recent work suggests that Pathfinder-derived structures may be a useful tool for the design of human-machine interfaces. The main objectives of this type of research are to make the devices easier to use and quicker to learn, thus making it possible for people to perform tasks that they might not otherwise be able to accomplish. As McDonald and Schvaneveldt (1988) reminded us, the main problem in this area is that the interfaces are generally based on the perspectives of the designer, which are not necessarily congruent with the *user's* perspectives.

Guidelines and standards are often established without adequate research or a theoretical basis. One possibility is that the mental model of the user should be examined and used in developing the interfaces. Several researchers have used Pathfinder to aid in uncovering the mental model.

Roske-Hofstrand and Paap (1986a) were the first to take this tack by extending the Pathfinder procedure to the important applied concern of human-machine interaction. They tapped the user's cognitive organization to develop a menu-driven system for the control-

display unit (CDU) within a simulator. Their objective was to ensure that a new component within an automated cockpit would be easy to use by making it "compatible with the content and organization of the pilot's existing cognitive structure" (p. 1302).

To obtain the conceptual organization of the panels, they had four pilots rate the similarity of each chunk of information associated with 34 panels of the CDU. Using Pathfinder, they developed three menus. One had high redundancy in that there was more than one route to a specific goal. The price of this was an increased menu size. Using the same basic graph, they also developed a menu that eliminated most of the redundancy and one that had no redundancy. A fourth menu was based on the intuitions of a design team.

To test these menus, they had four subjects work with each prototype. The subjects were given a set of 34 scenarios and questions that were to be answered using one of the four prototype menus. They found that subjects using the highest redundancy prototype had the lowest failure rate and were the fastest to solve the questions. The prototype of the design team was at the opposite extreme in both failure rate and speed.

McDonald and Schvaneveldt (1988) used Pathfinder and cluster analysis (Johnson, 1967) in an effort to capture information about the human-UNIX interface. They had 15 experienced UNIX users go through 219 documented functions printed on cards and sort familiar functions into piles based on relatedness. To prevent hierarchical filtering, the subjects were encouraged to make duplicate cards if they felt that a command belonged in more than one pile. From this sort, a conditional probability matrix was constructed. Both hierarchical cluster and Pathfinder analyses were performed on this matrix. Pathfinder provided more information than the hierarchical cluster analysis. For example, although cluster analysis placed the UNIX commands, *pc*, *pi*, *pix*, and *px* in one cluster, the PFNET revealed the commands formed a clique; but cluster analysis had these four in a large cluster of six commands. In short, the cluster solution was derivable from the PFNET, but not vice versa.

In addition to structure, the design of an interface requires some abstraction of the underlying categories to allow the structure to be implemented in a Von Neumann architecture. (It is worthy of note that a connectionist architecture may allow implementation of Pathfinder networks without abstraction of the underlying categories, and we have begun investigating this possibility.) Category labels for such clusters are one such abstraction. Thus (as discussed by both Cooke & McDonald, 1987; and McDonald & Schvaneveldt, 1988), 4 of the 15 experienced UNIX users rated 83 clusters of two or more UNIX commands for goodness on a 5-point scale and provided names for all but the very "bad" clusters. This procedure also allowed the researchers to distinguish between artifactual clusters and real conceptual clusters. Most ratings were 4's and 5's suggesting that the cluster analysis did permit meaningful abstractions.

In a related endeavor, McDonald and Schvaneveldt (1988) investigated task sequences by looking at the co-occurrences of commands given by nine experienced UNIX users during a session with the operating system. By capturing a graph of probable-next-commands, context-sensitive help could take advantage of the previous sequence of commands. For example, McDonald and Schvaneveldt found only one arc leading to the command *kill* in UNIX, suggesting that only one command frequently precedes it. If an interface had this type of information, it could easily "anticipate" likely next commands and offer assistance. In fact, it is interesting to speculate on how such a system could be custom designed to the user's level of experience. As the user learns more of the system, a monitoring interface might keep track of command sequences and modify the PFNET, and thus the assistance offered.

The uses of Pathfinder discussed in this section are nascent, but they promise to facilitate human-machine interaction by adapting the machine to match our underlying knowledge networks. Context sensitive help, generic operating system advice, and artificially intelligent adaptive computer interfaces may all benefit from this work.

Social Graphs

In principle, Pathfinder could be a useful scaling procedure in a number of social sciences. To date, however, its use has been exclusively in the cognitive sciences. Unlike our review of that work, the current section is more a call for further work than it is a review. Although we consider some of the graph-theoretic literature, our discussion of Pathfinder is based on PFNETs constructed to illustrate the potential of this scaling procedure. Admittedly, in this way we are able to sidestep many of the issues, theoretical and methodological, that have consumed the efforts of many an insightful scholar. What we do accomplish, however, is a demonstration that applications of Pathfinder are likely to bear fruit in the social sciences.

Graphs have clearly become a central concern to sociologists and social psychologists. The methodology has been addressed explicitly in *Network Analysis* (Knoke & Kuklinski, 1982) and *Applied Network Analysis* (Burt & Minor, 1983), and a number of journals routinely report graph-theoretic treatments of sociological issues (e.g., *Social Networks*). Much of this work is concerned with how to collect data for network analysis. We will not be concerned with this here except to note that much of the data collected has been less quantitative than it would need to be if Pathfinder informed the work. For example, social exchange analyses often begin with simple 0 and 1 sociograms, in which a 1 indicates that there is some exchange and a 0 indicates that there is not. With Pathfinder, more sensitive measures (e.g., how much money is exchanged) could be handled easily. To date, however, even when more sensitive measures are collected, the data are often considered without any scaling procedure, thus allowing measurement error to exert strong influences, or the data are reduced to a simpler form (e.g., 0/1 sociograms).

Interorganization Exchange

In our consideration of the sociological literature, the interaction of large groups (e.g., institutions) struck us as an important question to which one could successfully apply Pathfinder. Graph-theoretic constructs had been employed in this work; in fact, Knoke and Kuklinski's monograph *Network Analysis*, used exchange among organizations as a vehicle for illustrating applications of graph theory.

Sociologists have identified a number of methods for analyzing network data. We borrow the data from Knoke and Kuklinski to illustrate those methods and to make a comparison to Pathfinder. Those data were a subset of the complete study reported in Knoke and Wood (1981); we return to the complete dataset later.

Knoke and Kuklinski began with two 0/1 sociograms of 10 organizations: one for information exchange and one for money exchange. These sociograms were collapsed in different ways depending on whether structural equivalence was determined by continuous distance procedures or by discrete distance procedures (blockmodel procedures; White, Boorman, & Breiger, 1976). The differences between these procedures need not concern us here. However, both procedures attempt to define structurally equivalent subgroups: "two objects *a* and *b* of a set *C* are *structurally equivalent* if, for any given relation *R* and any object *x* of *C*, *aRx* if and only if *bRx*, and *xRa* if and only if *xRb*" (Knoke & Kuklinski, 1982, p. 59). In other words, if *a* and *b* are identical in their relations to all

2 Empirically Derived Graphs

other objects, then the objects *a* and *b* are equivalent. For both procedures, this equivalence is established by using hierarchical cluster procedures, such as Johnson's (1967) cluster analysis (for continuous distance) or CONCOR (for discrete distance).

The continuous distance procedure and the blockmodel procedure each yielded four clusters. The clusters appear in Table 2.

Table 2. Clustering of 10 organizations^a in Indianapolis according to two different procedures (Knoke & Kuklinski, 1982).

Continuous Distance	Blockmodel
(WRO, WEST)	(WRO, WEST)
(COMM, MAYO)	(COMM, MAYO)
(COUN, INDU, NEWS, EDUC, WELF)	(COUN, INDU, NEWS)
(UWAY)	(EDUC, UWAY, WELF)

^aWRO (Women's Rights Organization); WEST (West End Organization); COMM (Chamber of Commerce); MAYO (Mayor's Office); COUN (City/County Council); INDU (Local Industry); NEWS (Star-News); EDUC (Education); WELF (Welfare); UWAY (United Way).

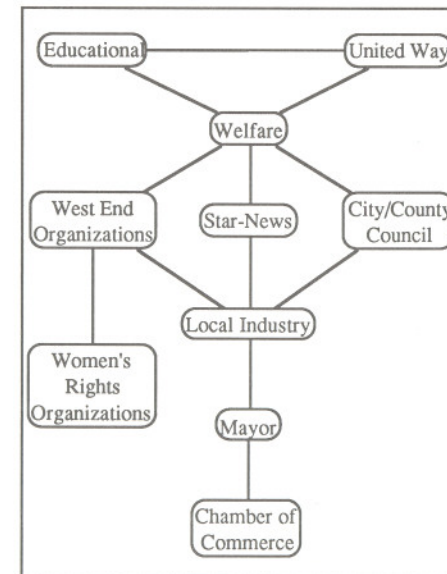


Figure 8. The sparsest PFNET for Knoke and Kuklinski's illustrative data taken from a sample of organizations in Indianapolis.

To create a PFNET for these organizations, a matrix of social role distances from Knoke and Kuklinski (1982, p. 63) was submitted to Pathfinder. The sparsest PFNET appears in Figure 8. Several facets of the graph are of interest. First, there is one clique (EDUCation, UNITED WAY, WELFare) similar to the result of the blockmodel. Second, there are three cycles involving City/County COUNcil, INDUStries, NEWS, WELFare, and West End Organizations in different ways. Finally, Women's Rights Organization and the Chamber of COMMERce-MAYOR's office connection are relatively isolated from the rest of the graph.

We think the PFNET agrees well with the structural equivalence analyses. It also helps show how the blocks hang together in a more Gestalt way. For example, the fact that WELF is a *cutpoint* helps explain why it is clustered with EDUC and UWAY in the discrete distance analysis and why it is clustered with COUN, INDU, and NEWS in the continuous distance output. Although it is true that there may be few

cliques if a strict graph-theoretic definition is used, the detection of circuits or cycles seems to supply some information about the substructure of the graph.

We now turn briefly to the original Knoke and Wood (1981) dataset from which the above organizations were sampled. We focus here on their analysis of perceived influence among seven blocks of organizations. Knoke and Wood asked organizations to indicate those organizations that had "policies or programs which your organization has tried to influence." We computed an index of the interconnections between blocks: For $i \neq j$, we computed the proportion of organizations within a block that had connections to organizations in other blocks; or, for $i = j$, this index reflected the proportion of organizations within a block that had connections, that is the intrablock connectedness. These data were submitted to Pathfinder (see Figure 9).

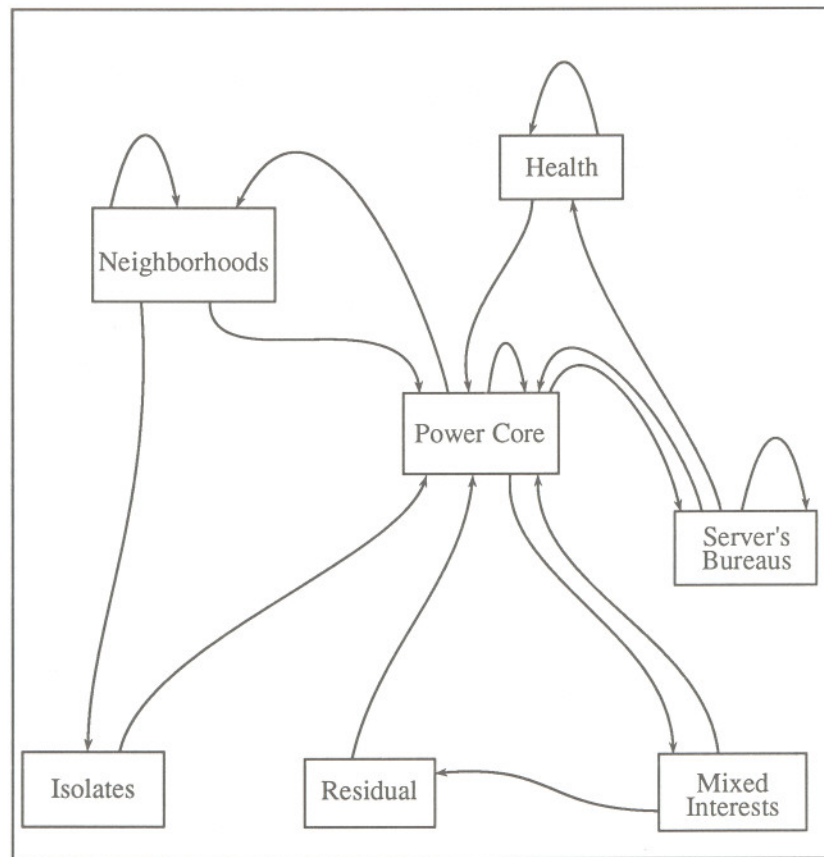


Figure 9. The directed PFNET for the seven blocks discerned by Knoke and Wood (1981). The graph is based on the data in Table 2.

The resultant PFNET revealed a number of features consistent with Knoke and Wood. The power core block, which should have been at the center of the graph, was represented by the node with the highest degree. Further, every node had an arc terminating at the power core. We calculated the degree of each node and then correlated these values with an independent assessment of influence collected by Knoke and Wood. The correlation was .75 between node-degree and perceived influence.

Blocks representing a tightly knit collection of organizations have a *loop* in the PFNET (i.e., an arc leaving and entering the same node). This is interesting in that Knoke and Wood point out that neither the neighborhood block nor the isolate block were influential, but the blocks differed in terms of the intrablock connections. Consistent with this, the neighborhood block had a loop in the PFNET but the isolate block did not. This indicates that the organizations making up the neighborhood block showed substantial interrelations, whereas the isolate members showed little connection among themselves as well as little connection with the rest of the graph.

Friendship

Another area of social graphs that emerges from the literature is the study of friendships. A typical method of collecting data about friends is to ask participants to nominate one or more individuals (Coie, Dodge, & Coppotelli, 1982; Peery, 1979; Wright, Giammarino, & Parad, 1986). This nomination procedure strikes us as a less than optimal way of gathering information from which to construct a graph, but researchers may have settled for this because of the lack of an algorithm for constructing a graph from complex data.

Some researchers in the area (e.g., Asher & Dodge, 1986; Asher, Singleton, Tinsley, & Hymel, 1979; Roistacher, 1974) have, in fact, criticized the nomination procedure. Hallinan (1982) has advocated giving the children free choice in naming as many friends as they desire, in order to avoid "a major source of measurement error in the data" (p. 57). Similarly, Roistacher (1974) chose to use a rating-scale measure to determine cliques within several high schools. For the rating-scale measure, subjects are given a roster with everyone's name and asked to rate their liking of the people. This procedure clearly lends itself to analysis by Pathfinder, more so than the nomination data. The rating-scale measure allows more total friendship choices than the nomination method, essentially allowing the construction of an $M \times M$ asymmetric input matrix for Pathfinder. In addition, the rating-scale procedure has produced results more consistent with what is suspected of friendships. For example, Slavin and Hansell (1983) reported that, for a cooperative learning study, the results based on the rating-scale measure indicated weaker friendships became stronger, whereas an earlier study using the nomination measure had reached the somewhat counterintuitive conclusion that it was the strong friendship that changed.

From either the nomination or rating-scale approach, several measures can be derived; one of the most common measures is popularity. This can be the averaged scaling (Asher & Dodge, 1986; Perry, 1987), the averaged nominations (Asher et al., 1979), or simply the frequency of positive nominations (Peery, 1979). There are also several measures of rejection or dislike, such as the proportion or frequency of negative nominations (Peery, 1979; Wright et al., 1986) and the proportion of the least-liked ratings given under the rating-scale instruction (Dodge & Somberg, 1987; Perry, 1987).

We thought Pathfinder might be useful for identifying children who differ in their friendship graphs. We used Perry's (1987) dissertation data.³ She had eighth-grade girls rate each other on a 0 to 5 Likert scale, where "1" meant dislike, "5" meant strong liking,

³We thank Bridgette Perry for supplying us with her data and for discussing the area of friendship with us.

and "0" meant this person was not known well enough to give a rating. Her matrix was transformed to dissimilarity (smaller numbers indicated greater liking) data and submitted to Pathfinder. Of course, the girls differed in the number of arcs terminating on a node (indegree) and in the number of arcs originating from a node (outdegree). In this context, the *indegree* is suggestive of popularity. Our indegree index correlated .69 with her measure of popularity and .76 with a measure of social preference.

Researchers (Coie et al., 1982; Dodge & Somberg, 1987; Peery, 1979) introduced a combination of the above measures to determine social preference, and classified children into one of four categories: popular children who are liked and relatively rarely disliked; controversial children who are both liked and disliked; rejected children who are disliked; and neglected children who are neither liked nor disliked. The rejected and neglected classifications could help target children with social deficits (e.g., Dodge & Somberg, 1987; Hansell & Karweit, 1983; Perry, 1987) for an intervention procedure.

We created four-fold classifications of Perry's students using measures based on the literature. We also created the same four-fold table based on Pathfinder outputs. We first created a PFNET as described above. Nodes with a high indegree could represent students who are popular or controversial. Nodes with a low indegree could represent students who are neglected or rejected. To distinguish these groups further, another PFNET was constructed; this PFNET was based on the data before being transformed to dissimilarity. Thus, with these input data, nodes with high indegree would be students who are rejected or controversial and those with low indegree would be neglected or popular. Combining these two PFNETs classifies each student into one of four categories.

Table 3. Classification^a of 8th graders used by Perry.

Rating Classification	Pathfinder Classification			
	Controversial	Neglected	Popular	Rejected
Controversial	0	0	0	4
Neglected	2	11	7	19
Popular	12	2	19	4
Rejected	1	2	1	0

^aRating classification was based on the data and followed procedures typically found in the literature. Pathfinder classification was determined by median splits on the indegrees from a PFNET where liked individuals were connected and from a PFNET where disliked individuals were connected.

As Table 3 indicates, when classification using measures from the literature was compared with classification based on Pathfinder, the agreement was at best fair. Only 38% of the students were classified in the same way by both procedures. The largest disagreements resulted from the tendency for Pathfinder to classify students as rejected rather than neglected, and to classify students as controversial rather than popular. Unfortunately, independent measures are not available to support one of the classification schemes over the other. The value of Pathfinder in assisting in categorizing students as controversial or rejected awaits further research.

Further, such classification research need not be restricted to the above table. One possibility is that if outdegree measures from Pathfinder are considered, friendliness can

become another classification variable. For example, those girls who are very friendly, but who are not liked as much in return, may have difficulties in relating to others either by not expressing themselves or because they lack the ability to empathize with others. As with neglected children, those falling in this unrequited group may have sociocognitive characteristics that suggest particular interventions.

We find it interesting that there has not been much work trying to investigate the interaction of friendliness and popularity, although the two factors have been considered separately (e.g., Hallinan, 1982). Of course, the nomination procedure precludes such efforts because it restricts the number of friends nominated.

Communication networks

We conclude this section on social graphs by indulging in pure speculation. The study of communication patterns has made considerable use of graphs, but has not taken advantage of methods for inducing the structure of the graphs. Although much of the work on small group behavior makes reference to differences in underlying communication structure, no empirical proof of the assumed differences has been attempted (Lawson, 1964; Leavitt, 1951; Shaw, 1954b, 1964). Rather, the empirical work on communication structure is characterized by restricting the structure of the communication a priori.

The communication pattern of a five-person group falls into 1 of 12 configurations (Shaw, 1981). These configurations (see Figure 10) range from a completely connected graph (or a comcon) where interchange is unrestricted, to more constrained configurations where, for example, four members can speak directly only to a central member and thus are forced to communicate indirectly to others (a wheel).

Leavitt (1951) has imposed these patterns on groups of people and has found that the structure affects performance on the task. For example, it seems that simple tasks (e.g., information gathering, Shaw, 1954a) are performed well when the imposed pattern has a center (e.g., wheel, Y), whereas performance in complex tasks is better when more strongly connected patterns are imposed, such as a completely connected pattern (i.e., a clique).

These controlled experiments put researchers in a unique theoretical position. There is evidence of the superiority of some structures over others, and the problem is now simply to determine if those are the patterns that emerge in communication situations where the experimenter has not restricted the pattern. For example, some work that has compared ad hoc with established groups has found that the established groups perform better (e.g., Hall & Williams, 1966). The explanation of this superiority has included the presumption that the established group used a more effective communication structure, presumably a comcon. Established groups do not, however, always outperform ad hoc groups (Ford, Nemiroff, & Pasmore, 1977; Hall & Williams, 1970). The inconsistency may be due to established groups sometimes adopting a completely connected pattern and sometimes not.

With Pathfinder, it would be possible to test this speculation. Subjects could be allowed to communicate without restriction. Telephones could be used to allow the experimenter to collect the necessary data about who talked to whom and for how long. These data could be submitted to Pathfinder, and the resultant PFNETs compared to each other and to objective measures of task performance.




Figure 10. The possible configurations of five people in a communication graph. From Shaw (1964). Communication networks. In L. Berkowitz (Ed.), *Advances in experimental social psychology* (Vol. 1, pp. 111-147). New York: Academic Press. Adapted with permission from Academic Press and the author.

Conclusions

Our primary purpose in this chapter was to demonstrate the potential of Pathfinder and its resultant PFNETs. Pathfinder had been applied with some success to a number of issues in cognitive psychology, and we supplied some evidence that it would be of value to researchers interested in social phenomena.

PFNETs were useful in a number of ways. In some cases, simply obtaining a particular configuration was of interest. In other cases, comparisons of graphs as a whole was of prime concern, as in the classification of individuals or in the comparison of basic- and superordinate-level categories. Analysis at a more microscopic level was also informative in some cases, as when shared links were combined to define expertise or when specific cliques and cycles were identified in the social exchange data. Finally, in some cases, graph-theoretic indices were correlated with other measures to show important relations.

In the studies that compared Pathfinder with other scaling procedures (usually MDS), Pathfinder did not fare badly. It was superior in predicting serial learning and free recall but was surpassed by MDS in some classifications of pilots. It is interesting to speculate on the differences between these two scaling methods. We suspect that Pathfinder will do well in situations where the closest relations are of prime importance, but that MDS will prove superior when the more distant relations bear on the task. It may be that MDS will be better for one class of tasks, and Pathfinder for another, or it may be that the two

procedures should be used in concert at each opportunity. In either case, it seems reasonable to us to use both routinely in order to reach some conclusion.

In our review of the literature, we were struck by just how much power is provided Pathfinder by its roots in graph theory. Clearly, researchers have only begun to scratch the surface of this power. More sophisticated uses of measures of centrality, dominance, and path length are only beginning to be explored.

It also became clear that PFNETs with a Minkowski exponent of infinity were of sufficient utility to question the use of other exponents, at least for psychological and sociological data. This is encouraging because these types of data rarely meet more than ordinal assumptions. On the other hand, we chose different q values in different situations. Directed graphs were considerably complex, even with the sparsest PFNET. Undirected graphs seemed to add a number of interesting connections when q was reduced to 2 (Hutchinson, 1981). Although as an exploratory tool, PFNETs of several q 's (and perhaps several r 's) could be considered, in confirmatory studies researchers will have to consider these parameters carefully. We believe that Pathfinder will prove especially useful as a confirmatory tool, if rational bases for selections of q are devised.

In addition to the fact that most PFNETs reported here had an r parameter of infinity and a q parameter of $n-1$, it was also the case that the PFNETs were treated as graphs rather than networks. We were surprised by the number of important applications that allowed us to consider only the graph properties of the PFNETs. We did not discuss any work that required the weights of the edges in the PFNET. Thus, although a PFNET can produce networks (i.e., graphs with weighted edges), the work reviewed here did not need or did not take advantage of this additional information.

Although there are issues that remain to be considered in the development of Pathfinder, it is nevertheless true that the algorithm provides a powerful tool for applications that have a graph-theoretic connection; and as we have shown there are several such connections.