



Chapter 10

Proximities, Networks, and Schemata*

Roger W. Schvaneveldt

Over the past several years, my colleagues and I have been developing and evaluating empirical methods for eliciting and representing human knowledge. The network scaling technique known as Pathfinder (Dearholt, Schvaneveldt, & Durso, 1985; McDonald & Schvaneveldt, 1988; Schvaneveldt, Dearholt, & Durso, 1988; Schvaneveldt & Durso, 1981; Schvaneveldt, Durso, & Dearholt, 1985, 1987, 1989; Schvaneveldt, Durso, Goldsmith, Breen, Cooke, Tucker, & DeMaio, 1985) has resulted from these efforts. The essential idea behind Pathfinder is that proximities between entities should be represented as links in a Pathfinder network if the resulting links form the minimum weight paths in the complete (or nearly complete) network with proximity estimates as link weights.

Cognitive Structures and Network Structures

Recently, some of our efforts have been directed at identifying substructures, such as categories, schemata, and procedures in collections of concepts. Various standard tools, such as cluster analysis and sequence analysis, may be of value in this effort. In conjunction with our work on Pathfinder, we are also investigating various formal properties of graphs and networks, such as cliques, dominating node sets, and blocks, as methods of identifying substructures in networks (cf. Esposito, Chapter 6, this volume).

Another promising line of investigation stems from the extensive work in artificial intelligence and cognitive psychology on spreading activation in networks (Anderson, 1983; Collins & Loftus, 1975; Collins & Quillian, 1969; Meyer & Schvaneveldt, 1976; Quillian, 1967; Schvaneveldt & Meyer, 1973). Spreading activation reveals substructures in the activity levels of nodes in the network as a result of the spread of activation from selected nodes.

Networks and Schemata

Recent work presented under the banners of connectionism or parallel distributed processing (PDP) is based on the use of activation in networks to accomplish cognitive computation. The studies reported in the present chapter were originally inspired by the chapter on schemata in Volume 2 (Chapter 14) of the PDP books (Rumelhart, Smolensky, McClelland, & Hinton, 1986).

* Various versions of this chapter were presented at the Rocky Mountain Psychological Association meetings in Albuquerque, New Mexico in April, 1987, at a colloquium at the University of Oklahoma in October, 1987, at the 28th annual meetings of the Psychonomic Society in Seattle, Washington in November, 1987, and at the workshop on High-Level Connectionism in Las Cruces, New Mexico in April, 1988. Russell Branaghan and Steven Graves assisted in collecting data. Jim McDonald, Derek Partridge, and Jordan Pollack provided valuable criticism.

Rumelhart et al. presented an example of network representations of schemata using knowledge about different types of rooms. They selected 40 descriptors of rooms including such things as *ceiling*, *walls*, as well as things found in rooms, such as *oven*, *sofa*, *television* (31 of the descriptors are listed in Table 1). They constructed a network in which the 40 descriptors were represented as nodes and the weights on the links between nodes reflected the degree to which the two descriptors connected by a link tend to occur together. The actual values for the link weights were determined by asking people to imagine rooms of specific types (i.e., kitchen, living room, office, bedroom, and bathroom). For each room imagined, the subjects were asked to indicate whether each of the 40 descriptors applied to that room. By analyzing the co-occurrence of the descriptors in the imagined rooms, weights on the links were computed such that descriptors that tend to occur together are connected by positively weighted links, descriptors that tend to *not* occur together have negatively weighted links, and intermediate degrees of co-occurrence lead to intermediate weights with appropriate signs.

Dynamic instantiations of schemata are realized by a procedure in which particular nodes are activated continuously or "clamped on" (as if an external input is continuously signaling the presence of the item designated by the node), and the activation is spread throughout the network until a stable pattern of activation across all of the nodes is attained. The counterbalancing effects of excitatory (positive) links and inhibitory (negative) links is critical in producing stable activation states. In a stable state, some nodes will be active and some inactive.¹ The active nodes are taken to be the "completion" of the schema given the information that the items corresponding to the initially clamped nodes are present. For example, if the node corresponding to bed were clamped on, one might expect the nodes corresponding to other items found in the bedroom to become activated through the spread of activation. Table 1 shows the final states reached by the Rumelhart et al. network after clamping *sofa*, *bed*, *oven*, or *bathtub*.² The second line of the table shows which nodes were clamped on for the activation runs. The descriptor items (nodes) are listed in the far left column, and the dots in the columns indicate the nodes that are activated when the indicated node is clamped on. Inspection of the table shows that the network seems to lead to reasonable completions of schemata given central members of particular schemata as the starting point for activation. For example, clamping *bathtub* leads to a stable state in which *cupboard*, *toilet*, *sink*, *scale*, and *bathtub* are active. Note also that the individual items may belong to any number of schemata. For example, of the four cases presented in the table, *fireplace* is active in one, *carpet* is active in two, *clock* is active in three, and *ceiling* (not shown) is active in all four.

While these results are quite interesting and suggest that networks may provide a medium for representing schemata, there is a sense in which obtaining these results is not surprising given the method for establishing the network in the first place. People were told to imagine rooms in generating their judgments, and the method showed that their judgments revealed rooms. If one were looking for a way to identify schemata in a domain

that was not already familiar, it could prove to be difficult to properly instruct people to make the required judgments.

Table 1. Rumelhart et al. (1986) results.

Node	Clamped			
	<i>sofa</i>	<i>bed</i>	<i>oven</i>	<i>bathtub</i>
telephone	•	-	•	-
books	•	•	-	-
sofa	•	-	-	-
drapes	•	•	•	-
cupboard	-	-	•	•
toilet	-	-	-	•
bed	-	•	-	-
desk-chair	-	-	-	-
easy-chair	•	-	-	-
stove	-	-	•	-
sink	-	-	•	•
scale	-	-	-	•
typewriter	-	-	-	-
clock	•	•	•	-
coffee-cup	-	-	•	-
coffeepot	-	-	•	-
dresser	-	•	-	-
oven	-	-	•	-
bookshelf	•	•	-	-
picture	•	•	-	-
ashtray	-	-	-	-
refrigerator	-	-	•	-
television	•	•	-	-
computer	-	-	-	-
desk	-	-	-	-
carpet	•	•	-	-
floor-lamp	•	-	-	-
fireplace	•	-	-	-
toaster	-	-	•	-
bathtub	-	-	-	•
hanger	-	•	-	-

• = Active - = Inactive

¹With the activation schemes used, nodes tend to be driven to an extreme state, either maximally active or minimally active. In practice, thresholds near zero and one (e.g., .001 and .999) are used to establish stability.

²Rumelhart et al. actually clamped on *ceiling* as well as the listed items. Since ceilings are found in all rooms, it doesn't contribute to differentiating rooms. Our own explorations suggest that clamping an item common to all of the schemata in a set doesn't make any difference. Representing still other schemata in the same set of units, however, may benefit from the presence of elements common to a subset of schemata. In general, we found that adding the additional 9 nodes that Rumelhart et al. used did not change the behavior of the derived networks with regard to the effects reported here.

The critical aspect of the data for constructing a network appears to be information about co-occurrence of the items in schemata. Perhaps this co-occurrence information could be obtained by simply asking for direct judgments about co-occurrence. Such instructions would not presuppose particular schemata (such as asking people to imagine an instance of a particular schema, e.g., a kitchen), but the co-occurrence of items in people's experience may be sufficient to serve as the basis of schemata representations. Thus, one

of the questions for the present study is: Can direct judgments of co-occurrence provide the basis for network representations of schemata?

Another issue of present interest derives from our recent work on using sparse networks to represent associational knowledge structures. The Pathfinder network generation algorithm takes proximity data as input and produces a network with a subset of the possible links. Frequency of co-occurrence can provide the proximity data required by the Pathfinder algorithm. In previous work with Pathfinder, the sparsest networks³ contain about n links on n nodes in contrast to the $n(n-1)/2$ links in a completely connected undirected network. For the 31 items we included in the present investigation, this means that a Pathfinder network would consist of about 31 links in contrast to the 465 links in the completely connected network. This decrease in density of networks has implications both for the interpretability of the network structures and for the complexity of algorithms required to compute with the networks. Sparser networks are easier to interpret, and they lead to significant savings in the amount of memory they occupy. Provided that they include essential connections, sparse networks may also lead to more efficient computation. Thus sparse networks have both theoretical and practical value. So another question for the present study is: Will the sparse Pathfinder networks provide a basis for representing schemata, or are the complete networks of the connectionist variety essential?

In addition, this study also investigates the variation in network performance with different activation schemes. Since a variety of schemes have been used in various connectionist studies, it seems worthwhile to compare and contrast the results obtained with different schemes. Finally, we compare the results obtained with activation paradigms to some clustering methods including the familiar hierarchical cluster analysis as well as some methods for extracting local information from Pathfinder networks.

Method

Data Collection

A subset of the room descriptors used by Rumelhart et al. was used in the present study. The 31 descriptors selected are shown in the left-hand column of Table 1.

Two different datasets were collected. One set came from using the same method reported by Rumelhart et al. to allow a comparison with their results. For this method, each of 24 people was asked to imagine five specific instances of each of five room types (i.e., living room, kitchen, bedroom, bathroom, and home office). For each specific room, they were asked to indicate whether each of the 31 descriptors was present in that room. These are the *rooms* data.

A second dataset came from asking each of 24 people to judge the frequency of co-occurrence of the items in the 465 pairs of the 31 descriptors on a scale of 1 to 9. The scale represented a continuum from *never occur together* to *always occur together*. These are the *co-occurrence* data.

³The r and q parameters associated with generating Pathfinder networks lead to systematic variations in the density (number of links) of the network. The sparsest networks are obtained with $r = \infty$ and $q = n-1$, where n is the number of nodes in the network.

Network Generation

Each of the two datasets was used to construct two different networks, a completely connected (connectionist) network and a sparse (Pathfinder) network. We consider each of the four cases.

Rooms Data - Connectionist Network. The rooms data were used to create a connectionist network following the procedures described by Rumelhart et al. (1986, p. 23). The method for establishing weights on the links between nodes essentially produces large positive weights for pairs of nodes that frequently occur together, large negative weights for nodes that never or infrequently occur together, and intermediate weights for the intermediate cases.

Rooms Data - Pathfinder Network. The Pathfinder network was generated from the rooms data by running the Pathfinder algorithm on the weights used in the connectionist network. Since the activation paradigms to be used require both positive and negative weights, the Pathfinder algorithm was run twice over the data. One run generated the positively weighted links, and by subtracting the weights from the maximum weight, the second run of Pathfinder generated the negatively weighted links. The Pathfinder network had 30 positive links and 31 negative links.

Co-occurrence Data - Connectionist Network. The co-occurrence data do not directly provide information about the probability of co-occurrence; they give relative information about co-occurrence. The weights for the connectionist network were obtained by averaging the ratings for each pair of items across the ratings given by the 24 subjects. The average ratings across all pairs were then converted to z scores with a mean of zero and a standard deviation of 1. The sign of the weights were such that positive weights corresponded to above average ratings of co-occurrence frequency and vice versa for the negative weights.

Co-occurrence Data - Pathfinder Network. Just as for the rooms data, the Pathfinder networks were generated from two applications of the Pathfinder algorithm to the average ratings of co-occurrence. Positively weighted links were determined from one run, and negatively weighted links from another run. The resulting network had 30 positive links and 30 negative links.

Network Activation

Various activation methods were used to investigate the nature of stable states in the four networks. With all methods, one node was clamped on throughout the activation procedure, and activation was passed through the networks until a stable state of node activation was reached. The active nodes in the stable state were taken as instantiations of the schema most associated with the clamped node. Different runs involved clamping different nodes. In this report we will focus on the results obtained from clamping either *sofa*, *bed*, *bathub*, or *refrigerator*, which were taken as closely associated with living room, bedroom, bathroom, and kitchen, respectively. Thus, if the activation methods succeed in instantiating schemata, we would expect the items commonly found in each particular type of room to be activated when a typical node for that room is activated.

Psychometric Analyses

In addition to the activation analyses, we also examined the results of using various psychometric methods to analyze the data. Each of the datasets (rooms and co-occurrence) was represented as a proximity matrix which was analyzed by Kruskal's (1964) nonmetric multidimensional scaling (MDS) method, by Johnson's (1967) hierarchical cluster analysis (HCA), as well as by the Pathfinder network scaling method discussed earlier. The

purpose of these analyses was to determine the extent to which the static results of psychometric methods reflected underlying schemata. We also compared the psychometric methods with one another and with the network activation methods.

Results and Discussion

A Pathfinder network for the rooms data is shown in Figure 1, and Figure 2 shows a Pathfinder network for the co-occurrence data. These two networks only depict the positively valued links. Figure 3 shows the negative links for the rooms data, and Figure 4 shows the negative links for the co-occurrence data.

The positive networks (Figures 1 and 2) are not particularly surprising. Items that occur most frequently are linked. The result is that items found in the same rooms tend to occur in sets of interlinked items. Of course, some method for determining where the clusters stop is required to isolate different rooms from one another because the whole set of terms is connected. If we consider the link weights, the single-link method of hierarchical clustering is embedded in the Pathfinder networks.⁴ By successively adding links in order of the magnitudes of their weights and identifying the connected components of the resulting network, the clusters obtained in a hierarchical cluster analysis will correspond to the connected components of the network.

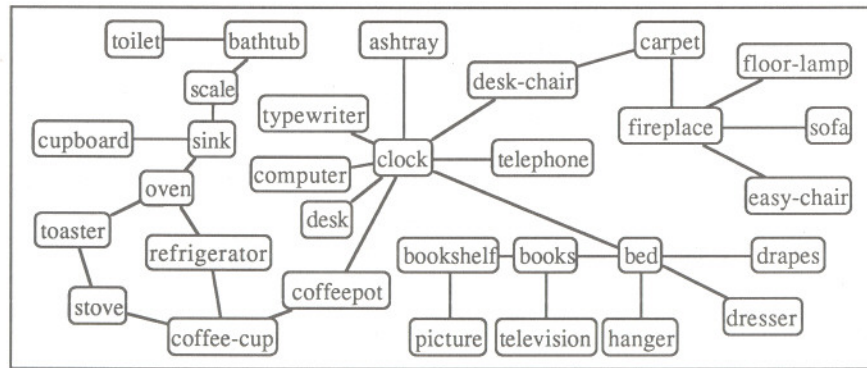


Figure 1. "Positive" Pathfinder network from Rooms data.

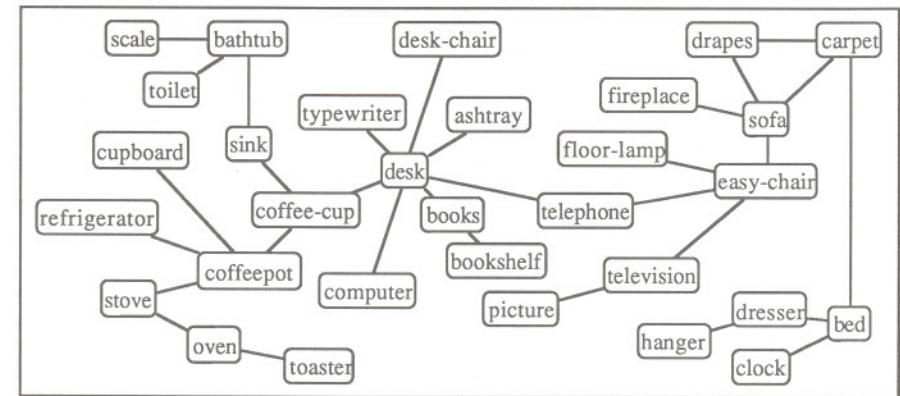


Figure 2. "Positive" Pathfinder network from Co-occurrence data.

The "negative" networks in Figures 3 and 4 show connections between items that occur least frequently. These negative links may help to identify items that do not belong to clusters of the positively linked items. The negative links are also necessary to implement the usual activation schemes used in connectionist networks. Without negative connections, positive activation simply spreads throughout the network until everything is maximally active. Negative links serve to dampen the positive activation and produce stable activation patterns across the nodes.

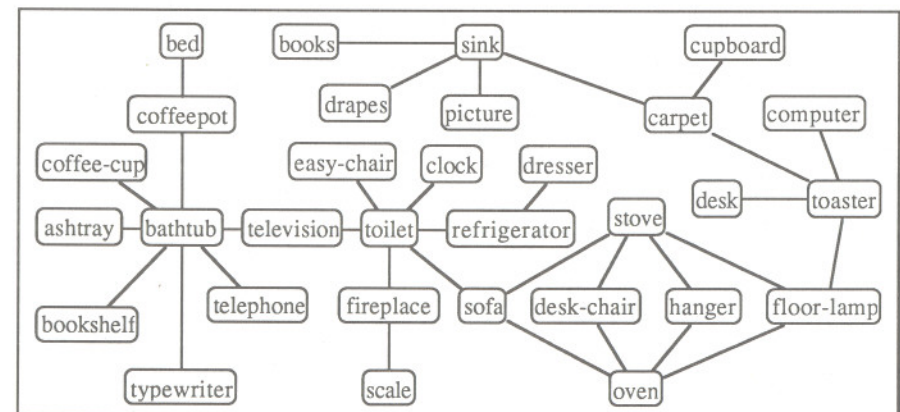


Figure 3. "Negative" Pathfinder network from Rooms data.

⁴This relation between Pathfinder networks and the single-link hierarchical clustering solution holds for networks generated with $r = \infty$ and $q = n-1$.

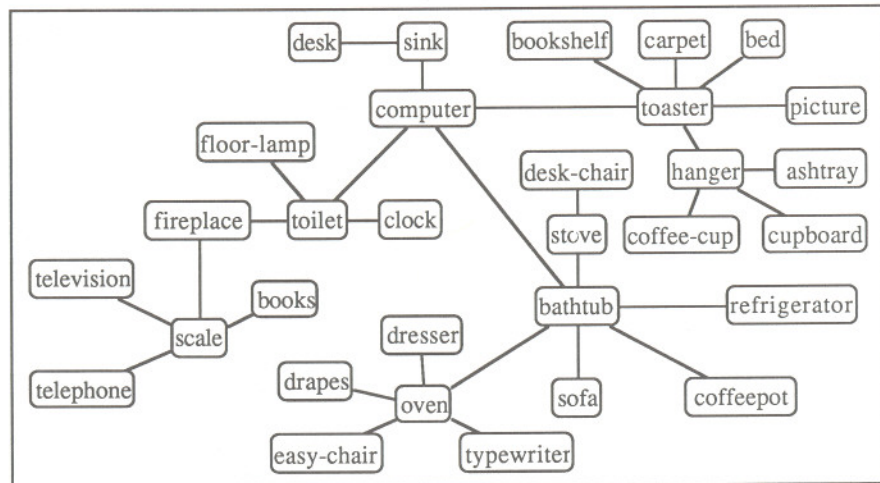


Figure 4. "Negative" Pathfinder network from Co-occurrence data.

Activation Experiments

Several different activation schemes were tried with the various networks we constructed. All of the methods yielded very similar results. A summary of these experiments is shown in Table 2. This table is similar to Table 1 except that results from four different networks and two datasets are shown. The second line shows which nodes were clamped on for the activation runs. The next lines show which dataset was used to construct the networks, and the next lines show whether the network constructed was a complete connectionist (CN) network or a sparse Pathfinder (PF) network. The CN network derived from the rooms data was based on the same method as that used by Rumelhart et al. (1986).

There are several aspects of these results that bear comment. First note that the activated items (large dots) tend to occur in pairs. This means that the complete networks and the Pathfinder networks are producing similar results. The one exception is that the Pathfinder network derived from the rooms data leads to activation of kitchen items when *bath-tub* is clamped on. In this case, the positive connections between bathroom and kitchen descriptors, through *sink* (see Figure 1), are strong enough to overcome the negative connections between *toilet* and *refrigerator* and between *bath-tub* and *coffee*-related items. Apparently a fine balance of positive and negative connections is required to achieve the desired selectivity in the activation process.

This lack of selectivity is seen in the activation patterns resulting from clamping *bed* and *sofa* for all of the networks. In general, the networks lead to activation of descriptors from living room, bedroom, and office whenever any of the descriptors typical of any of these rooms is clamped on. It is somewhat curious that these results do not replicate those of Rumelhart et al. (1986), but their footnote 7 on page 22 states that "Some slight modifications in the database were made in order to emphasize certain points in our example." Of course, some slight modifications in the networks used in the present examples would yield different results, but the interest here was in whether the activation schemes would reveal interesting patterns in the data. The answer to that question is negative.

A major lesson to be learned in this exercise is that activation schemes require a fine balance of positive and negative weights to yield interesting stable states. Collecting data in the two ways examined here does not necessarily result in such balance. A very careful selection of descriptors may help. What is needed are descriptors that are uniquely associated with particular rooms and negatively associated with other rooms. Such items are found in the present set for kitchens and bathrooms but not for living rooms, offices, and bedrooms. Of course a learning scheme could be used to adjust the weights to produce the desired results, but, again, that was not the goal of the present investigation. Perhaps some other ways of analyzing the datasets will be revealing.

Table 2. Results of Several Simulations

Node	Clamped															
	sofa				bed				refrigerator				bath-tub			
	Rooms		Co-oc		Rooms		Co-oc		Rooms		Co-oc		Rooms		Co-oc	
	Data		Data		Data		Data		Data		Data		Data		Data	
	CN	PF	CN	PF	CN	PF	CN	PF	CN	PF	CN	PF	CN	PF	CN	PF
telephone
books
sofa
drapes
cupboard
toilet
bed
desk-chair
easy-chair
stove
sink
scale
typewriter
clock
coffee-cup
coffeepot
dresser
oven
bookshelf
picture
ashtray
refrigerator
television
computer
desk
carpet
floor-lamp
fireplace
toaster
bath-tub
hanger

. = Active - = Inactive

Multidimensional Scaling

Two-dimensional multidimensional scaling (MDS) solutions for the rooms data and the co-occurrence data are shown in Figures 5 and 6, respectively. While the fit of these two-dimensional solutions is marginal (stresses are 0.16 and 0.17 for the two solutions), the solutions are quite revealing. In particular, the solutions show a clear separation of bathroom and kitchen as well as a lack of differentiation of bedroom, office, and living room. This result is very similar to that obtained using activation techniques.

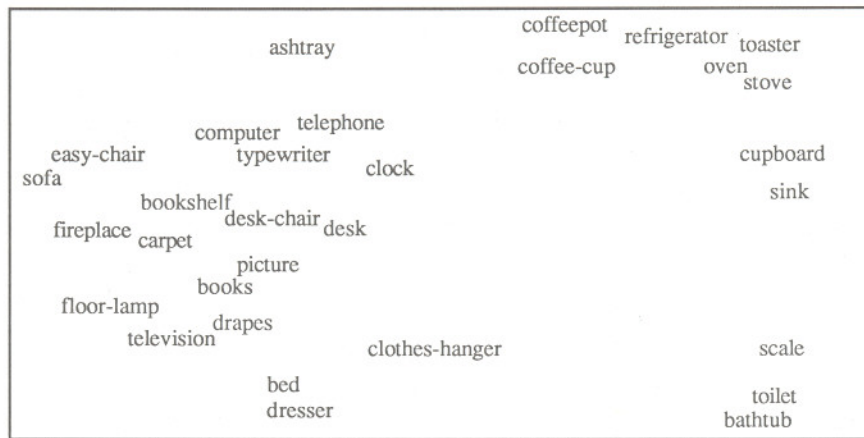


Figure 5. Two-dimensional MDS from Rooms data—Stress = 0.16.

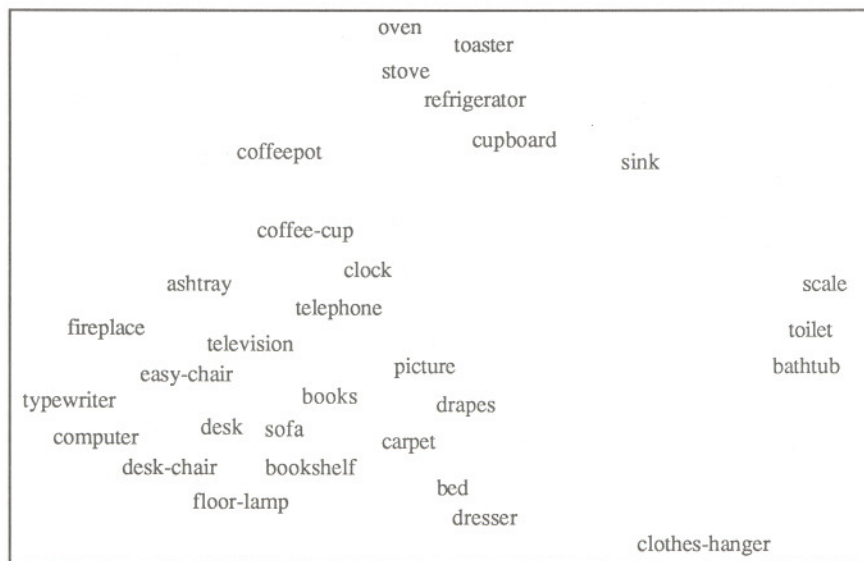


Figure 6. Two-dimensional MDS from Co-occurrence data—Stress = 0.17.

Cluster Analyses

Hierarchical cluster analysis is another common method for identifying groups of items on the basis of proximity data. Complete link hierarchical clustering solutions for the rooms data and the co-occurrence data are shown in Figures 7 and 8, respectively. Clusters are shown by the joining of lines emanating from the room descriptors. The figures may be interpreted as showing every descriptor in a separate cluster (at the left-hand side) that become a series of hierarchically organized clusters as one moves to the right in the figures. The figures do not show the final few steps which join the indicated clusters into successively larger clusters until all items are in a single cluster. At the point where the final clustering is shown, we see that there is a reasonable delineation of all five rooms.

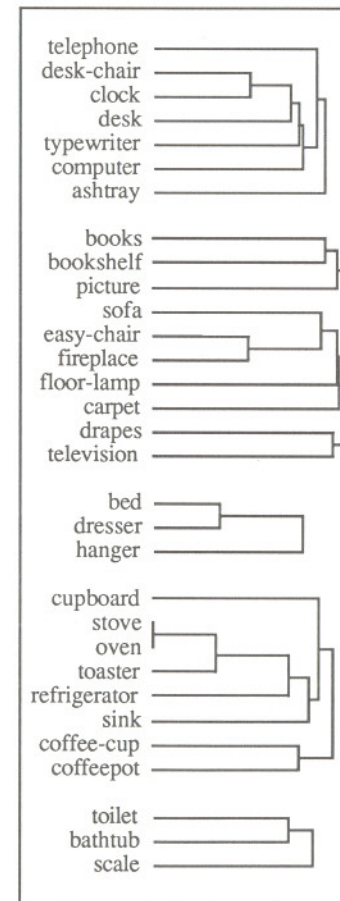


Figure 7. Complete-link cluster analysis from Rooms data.




Figure 8. Complete-link cluster analysis from Co-occurrence data.

The cluster analysis is able to identify a separation of rooms where the activation methods and MDS could not. Of course, interpretation of the cluster solutions is required to decide that the clustering should be stopped at the point shown in the figures. These solutions are also not capable of revealing that some items (e.g., *carpet*, *clock*, *sink*) are commonly found in more than one room since the structure is strictly hierarchical. Nevertheless, the clustering solutions do capture some of the "room" structure in both datasets.

Connectivity in Pathfinder Networks

Still another approach to identifying substructures in the datasets is to examine patterns of connectivity in the networks generated by Pathfinder. Esposito (Chapter 6, this volume) has investigated various graph-theoretic structures in an attempt to identify "meaningful" clusters in Pathfinder networks. For illustration, a simpler approach is used here. Consider the nodes that can be reached within a certain number of links of a given node. This criterion would define clusters of terms that are likely to be "good" clusters if the number of links is small. Table 3 shows the result of using this criterion, collecting all nodes within one link of a particular node with the additional wrinkle that nodes that are only linked to a node defined by the first step are also included in the cluster. An "orphan" node can only join the node to which it is connected, so including the orphan node whenever the connected node is included seems reasonable.

With one exception, the items included in the groups starting with the nodes clamped on in the activation experiments constitute reasonable sets of items to accompany the starting node. The exception is found in the network derived from the Rooms data when *bed* is the starting node. Examination of this network (Figure 1) shows that the problem lies in the role *clock* plays in the network. Since several orphans are connected to *clock*, they will all be included whenever *clock* is included.

These groups of items are not necessarily hierarchical in that the same descriptor can occur in several different groups (notice the descriptor, *carpet*, in the network from the co-occurrence data). The general tendency, with this method of defining clusters, is to include too few items in the clusters. Still, unique clusters for every starting node can be generated, and including more than one starting node would generate still other unique clusters.

Conclusions

The Pathfinder networks appear to capture much of the information contained in the complete connectionist networks in as much as similar states are reached in the two types of networks as a result of activation. With the co-occurrence data, there were no substantial differences between the stable states reached with the complete networks and the Pathfinder networks. These similarities suggest that Pathfinder analysis could be a useful tool in simplifying the complex networks that are often found in connectionist models. For example, such simplifications may help to interpret the results of learning in such networks.

The particular descriptors chosen are critical to the patterns reached as stable states in activated networks. In particular, sufficient negative connections must be present to "damp out" the activation produced by positive connections. The undifferentiated living room, bedroom, and office are illustrative of this lack of sufficient dampening in the experiments reported here.

Table 3. Clusters from connectivity in Pathfinder networks.

Node:	Starting Node							
	sofa		bed		refrigerator		bathtub	
	Rooms Net	Co-oc Net	Rooms Net	Co-oc Net	Rooms Net	Co-oc Net	Rooms Net	Co-oc Net
telephone	-	-	•	-	-	-	-	-
books	-	-	•	-	-	-	-	-
sofa	•	•	-	-	-	-	-	-
drapes	-	•	•	-	-	-	-	-
cupboard	-	-	-	-	-	•	-	-
toilet	-	-	-	-	-	-	•	•
bed	-	-	•	•	-	-	-	-
desk-chair	-	-	-	-	-	-	-	-
easy-chair	•	•	-	-	-	-	-	-
stove	-	-	-	-	•	•	-	-
sink	-	-	-	-	-	-	-	•
scale	-	-	-	-	-	-	•	•
typewriter	-	-	•	-	-	-	-	-
clock	-	-	•	•	-	-	-	-
coffee-cup	-	-	-	-	•	-	-	-
coffeepot	-	-	-	-	-	-	-	-
dresser	-	-	•	•	-	-	-	-
oven	-	-	-	-	•	-	-	-
bookshelf	-	-	-	-	-	-	-	-
picture	-	-	-	-	-	-	-	-
ash tray	-	-	•	-	-	-	-	-
refrigerator	-	-	-	-	•	•	-	-
television	-	-	•	-	-	-	-	-
computer	-	-	•	-	-	-	-	-
desk	-	-	•	-	-	-	-	-
carpet	-	•	-	•	-	-	-	-
floor-lamp	•	•	-	-	-	-	-	-
fireplace	•	•	-	-	-	-	-	-
toaster	-	-	-	-	-	-	-	-
bathtub	-	-	-	-	-	-	•	•
hanger	-	-	•	•	-	-	-	-

(•) = Within one link plus orphans

The particular type of proximity data used does not appear to be critical. Similar results are obtained from simple co-occurrence judgments compared with the rooms-based judgments. In the two cases where there were differences, the co-occurrence data tended to produce somewhat better results (see Tables 2 and 3).

Cluster analysis and connectivity in Pathfinder networks reveal clusters of items belonging to particular rooms, but activation methods and MDS do not. Apparently, different information in the data is used by the activation methods and MDS in contrast with the hierarchical clustering method and the Pathfinder networks. Elsewhere (Schvaneveldt et al., 1989), we have discussed such differences between MDS and Pathfinder.

In a nutshell, MDS equally weights all of the data in determining a solution. There is a sense in which all of the weights are also used by iterations of activation spreading in a network, and to the extent that Pathfinder networks preserve the significant paths in more complex networks, activation should act similarly in Pathfinder networks. In contrast, the *link structure* of a Pathfinder network is more determined by the locations of the relatively small data values (high co-occurrence for the positive links or low occurrence for the negative links). Thus Pathfinder connectivity and hierarchical clustering tend to identify items that most frequently occur together. The constraints in MDS and activation methods are much more complex, and as we have seen, not always appropriate.