Behavior Research Methods, Instruments, & Computers
1997, 29 (3), 450–455

Program Abstracts /Algorithms

PERMAP: An interactive program for making perceptual maps

RONALD B. HEADY
University of Southwestern Louisiana
Lafayette, Louisiana

and

JENNIFER L. LUCAS
Kansas State University, Manhattan, Kansas

This article describes a DOS-based computer program for making and testing perceptual maps. The program, PERMAP, uses conventional metric multidimensional scaling techniques. That is, it uses pairwise numerical values (correlations, proximities, dissimilarities, etc.) to construct a map showing the relationship between objects. A unique feature of PERMAP is that it embeds the mapping techniques in an interactive, graphical system that minimizes several difficulties associated with multidimensional scaling practices. It is particularly effective at exposing artifacts due to local minima, incomplete convergence, and the effects of outliers. It can associate various attributes with the resultant groupings and provide line-linking options to help the researcher identify the nature of perceived relationships. Problems involving multiple matrices can be treated using three different aggregation methods. The optional use of weighting factors is available.

PERMAP is a program that uses multidimensional scaling (MDS) to reduce multiple pairwise relationships to 2-D pictures, commonly called perceptual maps. Figure 1 shows a typical perceptual map. The data for Figure 1 were taken from the Table 17.11 of Churchill's text (1995). The Churchill data are in the form of correlation coefficients that show the relationships between 10 factors that influence the image of a department store. These correlation coefficients were calculated from responses to semantic differential scale questions given to a random selection of shoppers. Figure 1 illustrates how PERMAP can succinctly summarize the relationships involved with 45 pairwise similarities (i.e., the number of independent correlation coefficients between 10 factors).

Purpose of PERMAP. The use of MDS for the construction of perceptual maps is well developed and several computer programs are available. In fact, MDS was one of the earliest uses of high-speed computers in psychology and the social sciences. The purpose of PERMAP is to provide a particularly convenient method of producing perceptual maps and to do so in a way that helps the researcher avoid a number of common mistakes, as described in following sections. A fully interactive, visually oriented tool has some important advantages for these purposes and was a prime consideration in PERMAP's development. Interactive programs are becoming more common, but those that visually aid the researcher in understanding the nature of the solution are still rare. Kaufman and Rousseeuw's (1990) program for cluster formation, CLUSPLOT, is an example of a program with graphical output, but one that still only presents the solution graphically instead of letting the researcher watch and interact with the formation of the solution.

Usefulness of perceptual maps. A major advantage of MDS and perceptual maps is that they deal with problems associated with substantiating and communicating results based on data involving more than two dimensions. Marcoulides and Drezner (1993) emphasized this point. They discussed the importance of graphical communications and the role of the eye in interpreting and distinguishing object (factor, stimulus, characteristic) grouping. Although experts may be able to extract the subtle relationships represented in a matrix of numbers, this skill is not widespread.

Schiffman, Reynolds, and Young (1981) noted that an additional advantage of using perceptual maps is that they are low in experimental contamination. That is, the method does not require a priori knowledge of the attributes or stimuli to be mapped. Schiffman et al. also provided an outstanding example of the clarity that can result when perceptual mapping is used to analyze a problem. Perceptual mapping allowed them to clearly communicate the results of a study of the similarity of colors, whereas the corresponding factor analysis results were not easily interpreted or communicated.

Another important aspect of perceptual maps is that they are forgiving of missing or imprecise data points. Whereas some analytical techniques cannot tolerate missing elements in the input matrix, MDS results are often unaffected. This is because it is not uncommon for there to be much redundancy in the information given by a complete matrix of dissimilarities. For example, if the true relationships involve only three dimensions and the matrix of information contains the pairwise relationships between 10 factors, then much of the matrix's information must be repetitive and the map will not change if redundant elements are missing. This redundancy is similar to having identical, but repetitive, statements about a topic.

Existing perceptual mapping difficulties. Although the theory behind making perceptual maps is well devel-

Correspondence should be addressed to R. B. Heady, Department of Management and Quantitative Methods, University of Southwestern Louisiana, Lafayette, LA 70504 (e-mail: heady@usl.edu).

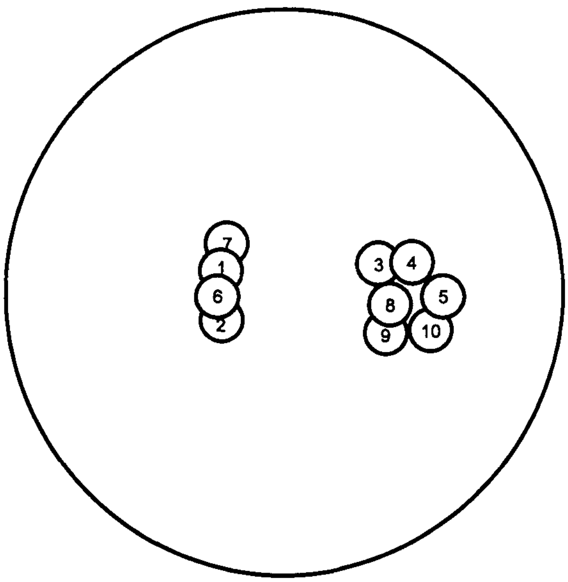

Figure 1. A perceptual map of Churchill's 10 factors controlling department store image.

oped, its application has been controversial. From its early advocates (Shepard, Romney, & Nerlove, 1972) to current advocates (e.g., Churchill, 1995; Hair, Anderson, & Tatham, 1987), authors have routinely given reminders of pitfalls awaiting the indiscriminate user. There are four major concerns that PERMAP can help alleviate. These include avoiding local minima (i.e., configurations that are optimal with respect to small changes in configuration but not optimal with respect to all possible changes), proving complete convergence, minimizing the influence of outliers, and combining multiple correlation matrices. With care, batch-operated programs can be used in such a manner that all of these difficulties are properly addressed, but moving to a visually interactive program renders these difficulties easier to deal with.

Dimensionality of mapping results. Perceptual maps are inherently 2-D because they are drawings on 2-D surfaces. However, this restriction to two dimensions on the map does not mean that the mapped objects are related on the basis of only two factors. The perceptual map simply presents the data in the two dimensions that best explain the variance. More precisely, the two dimensions used are those that the algorithm calculates as the best at explaining all of the relationships involved.

Multidimensional relationships are common. These can result from measuring independent pairwise relationships for four or more objects, or they can result from combining the results of several subjects giving 2-D placements of objects. This last point is explained in the work of Goldstone (1994), who showed how a multidimensional system of perception of lettering styles can be deduced by using 2-D drawings. Goldstone also provided a clear discussion of the limitations and assumptions underlying the use of spatial arrangements in experimental psychology.

Program Theory

PERMAP uses classical MDS metric scaling. The adjective *classical* indicates that a minimization procedure is applied to an objective function related to deviations between distances on the map and psychological or perceptual distances as given by a set of dissimilarity data. The adjective *metric* indicates that the dissimilarity data are accepted as ratio- or interval-level data, rather than just being rankings or preferential relationships.

Objective function and data types. In mathematical terms, the problem of reducing multidimensional relationships to two dimensions can be described in terms of minimizing the following objective function:

$$\text{MIN ObjFn} = \sum_i \sum_j W_{ij} d_{ij}^2. \qquad (1)$$

The indices i and j run from 1 to N, where N is the number of objects in the analysis. The d_{ij} factors are the differences between the distances on the map between the objects and the ideal distances specified in the dissimilarity matrix, δ_{ij}. By definition, d_{ii} is zero. The W_{ij} factors are optional weights, or saliencies, assigned to the dissimilarities of the objects.

Kaufman and Rousseeuw (1990) provided a comprehensive discussion of δ_{ij} matrices using nominal, ordinal, interval, ratio, or mixed-data types. For present purposes, it makes no difference if the δ_{ij} are generated using direct comparisons—also referred to as "subjective comparisons"—or the onefold approach, or attribute-related correlation coefficients, also referred to as "objective comparisons," or the twofold approach. Davison (1983) discussed how δ_{ij} can be measured for a wide range of experimental conditions. Basically, they can be measured using 1 subject, or a group of subjects, giving their opinion about the relationship between objects, or they can be measured by responding to questions about the attributes of the objects and then using a correlation between the attributes. Goldstone (1994) also examined traditional means of measuring δ_{ij}, and introduced a novel approach using visual techniques.

Distance formulas. Several common formulations of d_{ij} are mentioned in the literature. The most common is the Euclidian distance relationship:

$$d_{ij} = \sqrt{(X_i - X_j)^2 + (Y_i - Y_j)^2} - \delta_{ij}. \qquad (2)$$

The variables X and Y are measured using arbitrary orthogonal coordinates. The second measure is the city-block distance formula:

$$d_{ij} = |X_i - X_j| + |Y_i - Y_j| - \delta_{ij}. \qquad (3)$$

With a few notable exceptions (e.g., Arabie, 1991; Hubert, Arabie, & Hesson-Mcinnis, 1992), the city-block formula is not often used. However, it fits a niche in psychometrics (Torgerson, 1965) and thus is included as an alternative. A third alternative is the square of the Euclidian distance and dissimilarity measures:

$$d_{ij} = (X_i - X_j)^2 + (Y_i - Y_j)^2 - \delta_{ij}^2. \quad (4)$$

PERMAP provides the option of using any of these three distances formulas. This is important because most behavioral situations do not have a clearly defined correct alternative. Although all three formulas give almost the same results if the dissimilarity data have low symmetry, the city-block formula, given by Equation 3, can produce significantly different results when highly symmetrical groupings are involved. Further, the city-block form is more susceptible to numerical instabilities than are the other two forms (Kruskal, 1964; Kruskal, Young, & Seery, 1973). It is the least commonly used. The distance-squared formula, given by Equation 4, is more sensitive to outliers, but it is the most stable, in a numerical sense. Overall, the Euclidian relationship, given by Equation 2, is the most commonly used. Because no one of these equations has been proven superior for all situations of interest, it is advisable to experiment with all of them to determine the sensitivity of the groupings to the choice. PERMAP makes this experimentation easy.

Solution procedures. Extreme values in the objective function are found by setting the function's partial derivatives, with respect to X_i and Y_i, equal to zero and solving the set of $2N$ equations. This set of equations can be solved with any nonlinear simultaneous equation solver, and there are several such programs available. However, it turns out to be beneficial not to use a sophisticated mathematical routine that presents only one result. This is because when there are near-optimal solutions (that is, local minima exist), the researcher needs to know about them. For instance, it sometimes happens that the objective function has almost identical values for two object mappings that are different only by an interchange of the positions of two objects. Understanding the near equivalence of these multiple configurations might be important even when the difference in objective function values is trivial. A later section expands on the topic of local minima.

Goodness-of-fit. After converging to a solution, the final values of the objective function and Kruskal's Stress-1 and Stress-2 parameters (Kruskal et al., 1973) are displayed. Several different measures of goodness-of-fit parameters are available in the literature. The Kruskal stress parameters are some of the oldest, and most commonly mentioned, factors of merit so they are included to provide a statistic that can be compared with the results of other studies. Fitzgerald and Hubert (1987) provided guidelines for judging mapping results on the basis of the Kruskal stress parameter values.

Presentation orientation. The final groupings of objects are displayed graphically next to the values of the objective function and stress parameters just mentioned. When Equation 2 or 4 is used in the objective function, only the positions of the objects relative to each other are important. That is, the final arrangement can be translated or rotated anywhere on the graph without changing the optimality of the groupings. This is equivalent to saying that the grouping results are independent of how one holds the map. Therefore, the display of coordinates is optional. The overlaying of a coordinate system becomes important only when one moves to the interpretation stage and attempts to assign post hoc meanings to the relative positions of the objects.

When Equation 3 is used in the objective function, the situation is different. In this case, the final arrangement can be translated, or reflected about horizontal or vertical axes, without changing the optimality of the groupings, but rotation is not allowed. Thus, the city-block distance formula produces less generalizable results than does either of the other two formulas.

Program Features

The following description of PERMAP's operation is a condensed version of the information that is provided in an operation manual file that comes with the program. The detailed information in the operation manual file is augmented by an on-line help file.

Availability. PERMAP is available free from the first author by sending a preaddressed mailer with a DOS-formatted disk (5.25 or 3.5 in.). Alternatively, it can be downloaded from a World-Wide Web site (http://www.ucs.usl.edu/~rbh8900/ permap.html). PERMAP runs on any 386 or better IBM-compatible personal computer using any DOS-compatible operating system. Having a VGA or better monitor is preferable, although an EGA color video system will work. A math coprocessor is advisable for analyzing problems with more than about a dozen objects or if many cases are to be aggregated.

Program Control. PERMAP is controlled by using either a mouse or a keyboard. All options are presented through a series of menu selections, which are made by either clicking on the item of interest or using the Alt key to move the cursor to the menu list. Various special keys are available to modify the problem's parameters during or immediately after the analysis. These special keys duplicate control features available through the menus. They are useful for simplifying operation for experienced users. A pop-up reminder screen is available to show the various special control keys.

Problem size and system requirements. PERMAP can analyze an unlimited number of cases (different dissimilarity matrices), each representing the relationships between up to 30 objects. A problem with 10 objects typically takes less than 5 sec of personal computer time with a 50-MHz 486 machine.

No computer configuration procedures are required. The executable program file is approximately 170,000 bytes in size and requires three accessory files: the default file, the help file, and the data file. In addition, a text file contains a manual describing all program details. All files can fit on a single 3.5- or 5.25-in. floppy disk.

Data input. Input data are entered by using a standard ASCII file, so importing data from a spreadsheet or word processor is simple. Keyword identifiers announce the presence of various data types. See Figure 2 for an example of data used to generate the perceptual map of Figure 1. This figure does not show all the input options available. For instance, object naming, variable weights, multiple cases, and others, are not shown. However, Figure 2 does show all that is required to conduct a simple analysis. Optional information is covered by default values. That is, if one does not choose to use weights in the problem formulation, they need not be mentioned in the data file. The accompanying manual file provides detailed information on the available options.

Data types. Dissimilarity information can be entered as either dissimilarity or similarity values and can be based on any linear scale. They can be entered as half matrices (lower triangle) or whole symmetrical matrices. The whole matrix option is useful when one is downloading information from another program. It contains a significant amount of redundant information that one would not want to enter manually, but that need not be eliminated if the mechanics of data transfer cause it to be present. PERMAP deduces from the data context which kind of data is being entered (similarities or dissimilarities, values shifted or not shifted from zero, values expanded or contracted by a constant multiplier) and whether or not a full matrix is being entered. If PERMAP cannot make a safe deduction concerning the nature of the input data, it asks for more information. Asymmetrical matrices are not directly supported because they violate the assumption that there is a one-to-one correspondence between dissimilarity and distance on the map. However, asymmetrical matrices can be handled by either averaging the two sides or by splitting the matrix into two or more cases. Missing values cause no difficulties. They are simply specified by "NA" in the input matrix.

Solution display. If a particular problem shows numerical instabilities, it is possible to follow the convergence in slow motion, or stepped motion, to understand the dynamics of the solution procedure better. If the user wants to see the solution as quickly as possible, the display can be turned off until the solution is found.

The object groupings can be graphed using several different backgrounds, object coloring options, and naming options. User-supplied names are acceptable, or PERMAP can supply object numbers. Names or numbers can be placed in the circles representing the objects, or on pointers identifying the objects, or not displayed at all, according to the wishes of the operator. These options, as well as other miscellaneous options, are all controlled through the menu system. If there are many objects and a map becomes congested, it can be simplified by eliminating the names or by changing the scale of presentation. The plus and minus keys, along with the arrow keys, can be used to zoom in or out, and to translate the display. This allows examining the results in great detail.

Objective function options. Any of the three objective functions (Euclidian, Euclidian squared, city-block) can be specified for analysis. This is important for determining the effect of the objective function choice. Choice of objective function can be made from the menu system or by cycling through all possibilities by striking certain keys.

Local minima. Local minima are frequent in MDS analyses. By identifying local versus global minima, the existence of local minima can be changed from a point of difficulty to an opportunity. That is, they may provide insight into the nature of which intergroup relationships are controlling and which are secondary. Local minima sometimes produce objective function values quite close to the optimal objective function value. Other times, local minima can produce groupings that are far from optimal and can sidetrack the unsuspecting researcher into erroneous conclusions. For example, if all members of a group of four objects have equal dissimilarities with respect to each other, the relationship is accurately represented by a 3-D regular tetrahedron (a pyramid with equal-length edges). If this group is presented as a perceptual map using the city-block metric, one of two results usually occurs. One result has an objective function value of zero and a Kruskal Stress-1 value of zero. The other has an objective function value of 0.1667, and a Kruskal Stress-1 value of 0.2425. The first case is the true mini-

```
This is a PERMAP data file.
MESSAGE=DEPARTMENT STORE IMAGE PERCEPTIONS: Churchill's Data
NOBJECTS=10
CASE1
0
21, 0
59,.68,0
74,.79,.2,0
88,.8,.24,.25,0
11,.1,.66,.7,.89,0
13,.17,.6,.72,.77,.22,0
63,.69,.18,.22,.26,.7,.71,0
68,.65,.22,.19,.23,.61,.74,.18,0
82,.77,.28,.2,.17,.84,.83,.22,.23,0
```

Figure 2. Example input data file showing the data used to make Figure 1.

mum and clearly has better goodness-of-fit indicators, but experimentation shows that it and the other case are about equally likely to be the final result. This illustrates the importance of resolving the question of local versus global minima before going on to the next stage of analysis.

Other examples of local minima confounding the solution process are easy to find. For instance, the Churchill (1995) data used to construct Figure 1 are particularly prone to lead to false solutions that are just local minima. The configuration shown in Figure 1 has an objective function value of 0.1726 and a Kruskal Stress-1 value of 0.0793 and is the true global minimum. However, there are many configurations that are just slightly different from that shown in Figure 1, most of which involve exchanging the position of two objects and yield only slightly greater objective function values. These neighboring configurations are so stable numerically that it is somewhat tedious to identify the true optimal configuration.

The interactive nature of PERMAP helps in the recognition of the existence of local minima. Each solution can be watched as it develops. Then, it can be reanalyzed with the same or a new set of starting points at the touch of a key. By starting each analysis with a set of points randomly distributed in the unit square, the analyst can observe the convergence progress under widely varying circumstances. When this is not done, as is the case with some MDS programs, it is possible that the solution may converge to the same local minima each time it is run. This gives the appearance of the validity of the stable configuration even though a better configuration is possible and would have been found had a different set of starting values been used.

Convergence accuracy. The accuracy with which PERMAP finds the minima, be they local or global, is controlled by setting various convergence tolerance limits. PERMAP offers three levels of convergence tolerance. The first allows the fast screening of data, the second provides a balance between speed and precision and is adequate for most situations, and the third is so demanding that it will always outstrip requirements that take into account the accuracy of the input data. These levels of precision are controlled using the menu system or by using special keys during the analysis. Keyboard control allows the early termination of a run if satisfactory convergence occurs before the convergence limits are met or if an oscillatory pattern persists for more than a few cycles. Because of its interactive nature, PERMAP does not need an iteration limit.

Aggregation of multiple matrix information. Multiple matrix problems are important in many research situations. A single matrix holds information that summarizes the relationships among all objects in a study and is referred to as a *case*. Multiple cases result when there are multiple respondents, when a given respondent reports at multiple times, or with multiple groups of respondents when the average of each group is used as a single case. Additionally, different cases may represent dissimilarity relationships for different values of some parameter of interest, different treatments, or measurements made under different circumstances.

The analysis of multiple cases can be made by simply averaging all matrices and using the average matrix in the analysis. This is sometimes done, but it is generally not recommended (Ashby, Maddox, & Lee, 1994). When an averaged matrix is not appropriate, the average of the resultant object placements, each case being treated separately, is sometimes used. The real key to a good analysis, however, is to study each mapping separately before any kind of aggregation is made. It makes little sense to form clusters of characteristics on the basis of the evaluation of heterogeneous respondents or measurements taken under fundamentally different circumstances.

PERMAP helps with questions of aggregation by offering the options of averaging the matrices, averaging the results of the analysis of the individual matrices, or overlaying the results of various individual cases on a single graph for visual comparison. If the last option is selected, PERMAP moves each final solution to a standard orientation and position. Therefore, except for outliers, the overlaid objects fall nearly on top of each other. By concentrating on the role of the outliers, the researcher can segment the cases. Once homogeneous data sets have been identified, it makes little difference how the cases are aggregated.

Line linking. Kruskal and Wish (1978) discussed the difficulties of proving that a 2-D representation of multidimensional data is adequate, focusing on the possibility of artifacts from higher dimensions confounding the interpretation of a 2-D plot. They showed that linking the objects with low δ_{ij} values, and subsequently those with high δ_{ij} values, provides an effective test for consistency and reasonableness. That is, if nearby objects on the perceptual map are linked when the low δ_{ij} values are used, and if far away objects on the perceptual map are linked when the high δ_{ij} values are used, it is unlikely that some higher dimensional relationship is forcing unrepresentative arrangements in the 2-D plane.

The pattern of the links not only provides evidence about the suitability of the 2-D representation, but also, it can sometimes reveal the existence of curvilinear correlations that interconnect clusters and wind through the perceptual map. Consider, for example, a series of objects that plot roughly in the shape of the letter C. Are the two end objects actually somewhat similar, or are they near only because the rest of the chain of objects are similar across the C? This question can be quickly answered using the line-linking option.

Kruskal and Wish (1978) lamented that showing the diagnostic links is so time-consuming. Of course, this kind of analysis is simple for an interactive, visually orientated computer program like PERMAP. At the touch of a key, the object pairs with the lower one third of the δ_{ij} values can be linked. At the touch of another key these links can be erased, and then the highest one third of the δ_{ij} values can be linked.

Cluster identification. After a perceptual map has been constructed, the hard part starts. Additional progress is dependent on the ability to form a clear, unequivocal definition of the nature of the object groupings. As Schiffman et al. pointed out, "Dimensions that cannot be interpreted probably do not exist" (1981, p. 12). The identification task is often difficult and demands creativity and a deep understanding of the subject matter. This step cannot be automated, but the computer can provide some help.

PERMAP helps in the cluster identification stage by associating clusters and attributes. Up to nine attribute values can be assigned to each of the objects. These can be objective attributes or constructs. After the equilibrium configuration is found, the circle representing an object can be colored with an intensity that is proportional to the value of the object's attribute. Then, at the touch of a key, a different attribute set can be used as the coloring basis. The goal is to find attributes that form an identifiable pattern of associations between clusters and colors. As the analysis progresses, new combinations of the controlling attributes can be entered into the data file so that new construct hypotheses can be visually tested.

Conclusion

PERMAP provides an interactive, visual system for the construction of perceptual maps from multidimensional dissimilarity data. It can treat up to 30 objects and can aggregate an unlimited number of matrices (cases) describing the pairwise differences or similarities among the objects. Aggregation can be accomplished using any of three methods, and the use of weighting factors is available.

PERMAP was designed to be simple and easy to use by a novice and to offer enough advanced features that it would be of value to the expert. Its major improvement over existing perceptual mapping programs is that it was designed specifically to combat certain common errors associated with multidimensional scaling. For instance, it is particularly effective at showing incomplete convergence, trapping by a local minima, and outlier influence. It is also effective at revealing the importance, or lack of importance, of the choice of the distance metric used in the objective function. Overall, the program provides a means for the researcher to go beyond just finding a solution to developing a feel for the suitability, stability, and variability of the solution.

REFERENCES

Arabie, P. (1991). Was Euclid an unnecessarily sophisticated psychologist? *Psychometrika*, **56**, 567-587.

Ashby, F. G., Maddox, W. T., & Lee, W. W. (1994). On the dangers of averaging across subjects when using multidimensional scaling or the similarity-choice model. *Psychological Science*, **5**, 144-151.

Churchill, G. A., Jr. (1995). *Marketing research: Methodological foundations* (6th ed.). New York: Dryden.

Davison, M. L. (1983). *Multidimensional scaling*. New York: Wiley.

Fitzgerald, L. F., & Hubert, L. J. (1987). Multidimensional scaling: Some possibilities for counseling psychology. *Journal of Counseling Psychology*, **34**, 469-480.

Goldstone, R. (1994). An efficient method for obtaining similarity data. *Behavior Research Methods, Instruments, & Computers*, **26**, 381-386.

Hair, J. F., Anderson, R. E., & Tatham, R. L. (1987). *Multivariate data analysis* (2nd ed.). New York: Macmillan.

Hubert, L., Arabie, P., & Hesson-McInnis, M. (1992). Multidimensional scaling in the city-block metric: A combinatorial approach. *Journal of Classification*, **9**, 211-236.

Kaufman, L., & Rousseeuw, P. J. (1990). *Finding groups in data: An introduction to cluster analysis*. New York: Wiley.

Kruskal, J. B. (1964). Multidimensional scaling by optimizing goodness of fit to a non metric hypothesis. *Psychometrika*, **29**, 1-27.

Kruskal, J. B., & Wish, M. (1978). *Multidimensional scaling*. Beverly Hills, CA: Sage.

Kruskal, J. B., Young, F. W., & Seery, J. B. (1973). *How to use KYST-2: A very flexible program to do multidimensional scaling and unfolding*. Murray Hill, NJ: Bell Laboratories.

Marcoulides, G. A., & Drezner, Z. (1993). A procedure for transforming points in multidimensional space to a two-dimensional representation. *Educational & Psychological Measurement*, **53**, 933-940.

Schiffman, S. S., Reynolds, M. L., & Young, F. W. (1981). *Introduction to multidimensional scaling*. New York: Academic Press.

Shepard, R. N, Romney, A. K., & Nerlove, S. B. (1972). *Multidimensional scaling: Theory and applications in the behavioral sciences*. New York: Seminar Press.

Torgerson, W. S. (1965). Multidimensional scaling of similarity. *Psychometrika*, **30**, 379-393.

(Manuscript received July 17, 1995;
revision accepted for publication May 16, 1996.)