

Assessing the Assumption of Symmetric Proximity Measures in the Context of Multidimensional Scaling

Ken Kelley
University of Notre Dame

Applications of multidimensional scaling often make the assumption of symmetry for the population matrix of proximity measures. Although the likelihood of such an assumption holding true varies from one area of research to another, formal assessment of such an assumption has received little attention. The present article develops a nonparametric procedure that can be used in a confirmatory fashion or in an exploratory fashion in order to probabilistically assess the assumption of population symmetry for proximity measures in a multidimensional scaling context. The proposed procedure makes use of the bootstrap technique and alleviates the assumptions of parametric statistical procedures. Computer code for R and S-Plus is included in an appendix in order to carry out the proposed procedures.

Multidimensional scaling (MDS) is a multivariate statistical technique developed to assess the relationship among a set of objects, such that the relationships among the objects can be represented in some K dimensional space. Specifically, various attributes of objects are compared in some manner such that a measure of proximity can be obtained. A proximity measure is an index over all pairs of objects that quantifies the degree to which two objects are similar or dissimilar (Davison, 1983). The proximity of objects can thus be conceptualized as either being a *measure of similarity* or as a *measure of dissimilarity*, depending on whether higher values of the index represent objects that are most alike or least alike, respectively.

In applications of MDS where interest lies *not* in how similar or how dissimilar two objects are, but rather in whether two objects are equivalent, the appropriate proximity measure is a *measure of equivalence*. A measure of equivalence is one where a binary (yes/no; true/false; present/absent) response replaces the quantitative degree of similarity or dissimilarity between pairs of objects. A measure of equivalence is appropriate when, for example, interests lies in the comparison of objects to determine if they are identical. An example of such a scenario is given by Rothkopf (1957). Rothkopf had participants respond either “same” or “different” when two International Morse Code signals were presented one after the other.¹ The rationale for such a study was to determine which, if any, signals were commonly confused with one another. Rothkopf went on to determine the “similarity” of pairs of symbols by determining the relative frequency of responses from across the individual measures of equivalence.

For purposes of the present article let \mathbf{X} be an N by p by p data array of proximities, where each of N individuals has a potentially unique p by p matrix of proximities.² Each of the N individual p by p matrices will be denoted \mathbf{X}_i ($i = 1, \dots, N$). Let δ_{irs} ($\delta_{irs} \geq 0$) be a dissimilarity measure for the i th individual between the r th object and the s th object ($r, s = 1, \dots, p$), while letting σ_{irs} ($0 \leq \sigma_{irs} \leq 1$) be a similarity or equivalence measure for the i th individual on the r th and s th objects.

The goal of MDS is to represent the relationship among the proximity measures, ideally accurately and in a low dimensional space, such that the modeled distances over all individuals, d_{rs} , between the objects in the K dimensional space are as similar to the proximity measures as possible (see, for example, Cox and Cox, 2001 or Davison, 1983, for a comprehensive introduction to MDS; for notational ease the i subscript will be dropped from the proximity measures). That is, MDS attempts to find a low dimensional space such that modeled relationships (i.e., the d_{rs}) accurately represents the observed relationships (i.e., the δ_{rs} s or σ_{rs} s).

In the context of MDS, there are four axioms commonly employed for dissimilarity measures in order to carryout a MDS analysis (Davison, 1983, p. 2):

$$\delta_{rs} \geq 0 \tag{1a}$$

$$\delta_{rr} = \delta_{ss} = 0 \tag{1b}$$

$$\delta_{rs} = \delta_{sr} \tag{1c}$$

$$\delta_{rs} + \delta_{st} \geq \delta_{rt} \tag{1d}$$

For similarity measures, the analog of the first three axioms for dissimilarities are also assumed and given as follows:

$$0 \leq \sigma_{rs} \leq 1 \tag{2a}$$

$$\sigma_{rr} = \sigma_{ss} = 1 \tag{2b}$$

$$\sigma_{rs} = \sigma_{sr} \tag{2c}$$

Measures of equivalence rely on the same axioms as measures of similarity, however, the analog of Equation 1a is not Equation 2a, but rather

$$\sigma_{rs} = 0 \text{ or } \sigma_{rs} = 1, \tag{3}$$

as the r th and s th objects are rated either “same” or “different.”

Of interest in the present article is the third axiom for dissimilarities (Equations 1c), and the third axiom for similarities (Equation 2c) and equivalencies (Equation 3), namely the assumption of symmetry. While in some applications of MDS such an assumption may be plausible, in many areas of behavioral research such an as-

sumption is likely untenable. Specifically referring to measures of similarity, although the same can be said for dissimilarities and equivalencies, Tversky states that "the assumption of symmetry underlies essentially all theoretical treatments of similarity" and then goes on to state that asymmetric similarities should not be treated as symmetric relations (1977, p. 328).

Collins (1987) discusses asymmetric relationships in the context of a social interaction framework. In a classroom where each student ranks their liking of other students, there is no reason to believe student r will like student s the same that student s likes student r . In such a situation $\sigma_{rs} \neq \sigma_{sr}$ and thus the assumption of symmetry would be violated. Another example is taken from Seiyama, Naoi, Sato, Tsuzuki, and Kojima (1990), where interest was in intergeneration occupational mobility of sons as compared to their fathers in eight occupational categories. The proximity measure used was the number of sons whose occupations are in category r and whose fathers are (or were) in occupational category s . The (r, s) th element thus represented the number of sons in category r whose father's occupation was in category s . (See Okada and Imaizumi, 1997, for an example of how this asymmetric set of proximity measures can be analyzed.) It is not necessary, and indeed unlikely because of the economy and evolving job needs, for such a proximity matrix to be symmetric. A final example is given by Davison (1983). Davison's example concerns a proximity matrix where one dimension (say row) represents students' freshman major and the other dimension (say column) represents their graduation major (p. 51). There is no reason to assume such a matrix will be symmetric, and indeed it will not likely be symmetric, because the conditional probability of corresponding upper and lower diagonal elements need not be equivalent. For example, many students who initially plan to study engineering will switch majors, yet few students from other majors will change to engineering. Thus, an asymmetric relationship would exist. Many other examples could be given, but the point is that many asymmetric proximity mea-

asures exist. Although there are some MDS models capable of dealing with such asymmetric relationships, there has not been much work on inferring when asymmetric relations exist in the population and not just in the obtained sample.

The most common approach to estimate symmetric relations is to average δ_{rs} with δ_{sr} or σ_{rs} with σ_{sr} and proceed in the usual way (Kruskal, 1964; Holman, 1979). The methods given in the present article can be used to help determine if special MDS models, such as ASYMSCAL (Young and Lewyckyj, 1979), or methods, such as representing each object twice and thus treating σ_{rs} and σ_{sr} or δ_{rs} and δ_{sr} as distinct values (Cox and Cox, 2001, p. 116–119), need to be considered. Other MDS models and methods exist for modeling *known* asymmetric relationships between proximity measures (e.g., Young, 1987; Collins, 1987; Okada and Imaizumi, 1997; Harshman, Green, Wind, and Lundy, 1982), but determining when asymmetric MDS models and methods are appropriate has not often been considered.

The purpose of the present article is to test probabilistically, via an application of the bootstrap technique (Efron and Tibshirani, 1993), the likelihood that the asymmetric axiom of Equations 1c, 2c, or 3 is violated in the particular population, given the presumed randomly sampled observations. Specifically for measures of proximity (i.e., dissimilarity, similarity, and equivalence), a bootstrap algorithm is developed that uses individual data matrices to infer probabilistically for specific comparison values in a confirmatory fashion, or to use across the entire set of proximities in an exploratory fashion. Such information can provide a researcher with probabilistic evidence as to whether he or she should or should not proceed in the "usual" way by averaging δ_{rs} with δ_{sr} or σ_{rs} with σ_{sr} .

The Bootstrap Algorithm and its Rationale

The theory of the bootstrap technique is based on resampling, which allows for inferences to be made without regard to the standard (such as normally distributed errors and homogeneity

of variance) assumptions of parametric statistical tests. To make proper use of inferential tests, such as the *t*-test or analysis of variance for example, the assumption of normality of the population difference between δ_{rs} and δ_{sr} or σ_{rs} and σ_{sr} must be assumed and the differences must have the same variance if multiple asymmetries are to be tested simultaneously, as well as assuming the observations among individuals are independent. Making use of the bootstrap technique allows researchers to relax the parametric assumptions, yet still make informed inferences about relationships among the population proximity measures.³ Although bootstrap techniques have been recommended widely by methodologists and have been used for a wide variety of statistical tests, including certain aspects of MDS (Bell and Cox, 1998), to the author's knowledge the bootstrap technique has not been applied to assess the assumption of symmetry for proximity measures in the context of MDS.

The bootstrap technique applied to assess the assumption of symmetry begins at the data coding stage. Rather than aggregating across individuals, as is generally done to form an overall proximity matrix, the bootstrap approach requires each of the N individual matrices to be entered separately. In this sense either an N by p by p array must exist or N p by p matrices must exist.

The most common approach to handling asymmetric matrices (dating back to at least Kruskal, 1964) is to use the mean of the corresponding elements (i.e., δ_{rs} and δ_{sr} or σ_{rs} and σ_{sr}) in the upper and lower diagonals. The mean of the corresponding elements is obtained by the following:

$${}_s\mathbf{X}_i = \frac{\mathbf{X}_i + \mathbf{X}_i'}{2}, \quad (4)$$

where \mathbf{X}_i is the i th individual's p by p matrix of proximities, \mathbf{X}_i' its transpose, with ${}_s\mathbf{X}_i$ being the newly formed symmetric matrix that averages corresponding elements from the upper and lower diagonals.

The theory of the conventional bootstrap is to sample, with replacement, N individual data

matrices a large number of times. Because the bootstrap procedure samples with replacement, on a given sampling of N data matrices (which is done a large number of times), the i th matrix may be selected multiple times, only once, or not selected on any particular bootstrap replication. It is important to realize that each of the individual data matrices have an equal (i.e., $1/N$) probability of being selected on any given sampling. After the N individual data matrices have been randomly selected, statistic(s) of interest can then be calculated. In the present context, the "statistic of interest" is the difference between the mean of the proximity measures and the corresponding elements that the averaged value is supposed to represent. The N difference matrices are obtained by subtraction:

$${}_D\mathbf{X}_i = \mathbf{X}_i - {}_s\mathbf{X}_i, \quad (5)$$

where ${}_D\mathbf{X}_i$ is the difference matrix for the i th individual. The N ${}_D\mathbf{X}_i$ matrices are then combined by summing each of the individual matrices together:

$${}_D\mathbf{X} = \sum_{i=1}^N {}_D\mathbf{X}_i. \quad (6)$$

In the bootstrap procedure, N individual data matrices will be selected and the ${}_D\mathbf{X}$ computed B times in order to discern the likely value, range, and variability of the statistic of interest. That is, the bootstrap theory mandates that a given procedure be performed B times (e.g., $B = 5,000$ or $10,000$), such that an empirical distribution of the statistics of interest can be obtained. In the present context interest lies in determining the distribution, specifically the confidence limits, of the difference between the corresponding proximity measures above and below the principal diagonal of the overall proximity matrix. If the confidence interval contains zero for the difference between the corresponding elements above and below the principal diagonal, no evidence would exist showing that the relationship is not symmetric. However, if the confidence interval does not contain zero, then it can be inferred probabilistically, at approximately the designated level, that the relationship is not symmetric.

To rephrase the proceeding discussion, the following is how one could make use of the suggested bootstrap rationale for testing the assumption of similarity:

1. setup data in an N by p by p array or in N p by p matrices;
2. sample, with replacement, N individual \mathbf{X}_i matrices;
3. compute ${}_s\mathbf{X}_i$ for each of the N individual \mathbf{X}_i matrices;
4. calculate ${}_D\mathbf{X}_i$ by subtracting ${}_s\mathbf{X}_i$ from \mathbf{X}_i for each of the N individual matrices;
5. sum each of the N individual ${}_D\mathbf{X}_i$ matrices, such that an overall matrix ${}_D\mathbf{X}$ exists;
6. repeat steps 2 through 5 B times;
7. determine percentiles (i.e., the confidence limits) of interest for the overall distribution of individual elements from the B ${}_D\mathbf{X}$ matrices.

There are $p^2 - p$ confidence intervals that can be formed, one for each of the upper and lower diagonal elements. However, as will be shown momentarily, the confidence intervals for corresponding elements (such as the (s, r) element and the (r, s) element) will be symmetric, thus only $(p^2 - p)/2$ confidence intervals provide independent information. The following section provides examples of how the 7 steps listed above apply to bootstrapping in the test of symmetry for proximity measures.

Examples and Applications of the Bootstrap Technique to Test for Symmetry of Proximity Measures

The proximity measure of equivalence is one commonly used in the context of MDS. In fact, this is the measure of proximity used in the Rothkopf (1957) Morse Code example, the Okada and Imaizumi (1997) intergenerational occupation movement study, and the Davison (1983) college major example. Using a scaled measure of equivalence, two hypothetical data sets were created, one where symmetry existed in the population and one where symmetry did

not exist in the population, such that the procedure given in the previous section could be illustrated.

Suppose the following matrix represents the conditional probabilities across N individuals for some measure of equivalence:

$$\mathbf{NT} = \begin{bmatrix} 0.85 & 0.07 & 0.02 & 0.02 & 0.04 \\ 0.07 & 0.80 & 0.07 & 0.05 & 0.01 \\ 0.02 & 0.07 & 0.75 & 0.04 & 0.12 \\ 0.02 & 0.05 & 0.04 & 0.75 & 0.14 \\ 0.04 & 0.01 & 0.12 & 0.14 & 0.69 \end{bmatrix},$$

where \mathbf{NT} represents the matrix where the null hypothesis of symmetry is true. The hypothesis can be verified because $\mathbf{NT} = \mathbf{NT}'$. Note that the sum of each row across the columns is one, while the sum of each column across the rows is also one, thus leading to an appropriate set of probabilities for a symmetric matrix when the null hypothesis of symmetry is true. An example where the null hypothesis of symmetry is false is given by the following matrix of conditional probabilities across the N individuals:

$$\mathbf{NF} = \begin{bmatrix} 0.85 & 0.02 & 0.04 & 0.07 & 0.02 \\ 0.07 & 0.80 & 0.01 & 0.05 & 0.07 \\ 0.02 & 0.07 & 0.75 & 0.12 & 0.04 \\ 0.02 & 0.05 & 0.04 & 0.75 & 0.14 \\ 0.04 & 0.01 & 0.12 & 0.14 & 0.69 \end{bmatrix}.$$

Matrix \mathbf{NT} and matrix \mathbf{NF} will be used to illustrate the results of the bootstrap procedure when the null hypothesis is true and when it is false, respectively. Appendix A gives syntax for the computer programs S-Plus and R, such that the bootstrap procedure presented here can be applied to some data set in order to test the assumption of symmetry.

When forming confidence intervals, it is important to realize the population value(s) that are being bounded. For measures of equivalency,

which is the focal point of the example, the values of the symmetrized matrix are defined as

$$M_s = 100 \left(M - \frac{M + M'}{2} \right), \tag{7}$$

where **M** is the matrix of conditional probabilities for equivalency measures. The matrix **M_s** is then a symmetric matrix of the percentage of cases contained in each element of the matrix **M**. When **M** is symmetric **M_s** it is a null matrix (i.e., a *p* by *p* matrix of zeros). The absolute value of the off diagonals (note that the principal diagonal of **M_s** will always be zero) of **M** become larger as the assumption of symmetry is less plausible. The following two subsections give the results (in terms of confidence limits for the conditional percentages) of the bootstrap procedure for testing the assumption of symmetry for the matrices **NT** and **NF**.

Results when the Null Hypothesis was True

A Monte Carlo simulation was conducted to test the integrity of the suggestions provided throughout the article. The Monte Carlo simulation consisted of 7,500 replications of the bootstrap procedure. The number of bootstrap replications (within each one of the 7,500 replications) was 5,000. The sample size used throughout the simulation was 100.

The results for the lower bound of the 95 percent confidence interval for the matrix **NT** are given in matrix form as follows:

$$\text{Lower Limit NT} = \begin{bmatrix} 0.00 & -0.076 & -0.053 & -0.035 & -0.039 \\ -0.066 & 0.00 & -0.093 & -0.050 & -0.019 \\ -0.023 & -0.047 & 0.00 & -0.067 & -0.103 \\ -0.112 & -0.070 & -0.040 & 0.00 & -0.111 \\ -0.070 & -0.037 & -0.077 & -0.084 & 0.00 \end{bmatrix}$$

The results for the upper bound of the 95 percent confidence interval for the matrix **NT** are given in matrix form as follows:

$$\text{Upper Limit NT} = \begin{bmatrix} 0.00 & .066 & 0.023 & 0.042 & 0.070 \\ 0.076 & 0.00 & 0.047 & 0.070 & 0.037 \\ 0.053 & 0.093 & 0.00 & 0.040 & 0.077 \\ 0.035 & 0.050 & 0.067 & 0.00 & 0.084 \\ 0.039 & 0.019 & .103 & .111 & 0.00 \end{bmatrix}$$

Because it is known that the null hypothesis is true, zero should be contained within each of the 95 percent confidence limits, if indeed the confidence intervals correctly bracket the population value. Inspection of matrix **Lower Limit NT** and **Upper Limit NT** reveal that zero is indeed bracketed by the confidence limits for each of the 20 elements (10 above the principal diagonal and 10 below; note the pattern of negative relationship between the confidence bounds for the transposed elements in the two matrices because of $s_i X_i$ in step 3 of the procedure) from the **NT** matrix. Although the bootstrap procedure was shown to work well when the null hypothesis is true, it is also important to assess the suggested procedures when the null hypothesis is false. The next subsection illustrates such a case.

Results when the Null Hypothesis was False

The previous subsection illustrated that the proposed procedures worked well when the null hypothesis was true. The present section makes use of the asymmetric matrix **NF** to assess the effectiveness of the procedure when the assumption of symmetry is violated. The same procedures given in the previous subsection were also used for the following results.

The results for the lower bound of the 95 percent confidence interval for the matrix **NF** are given in matrix form as follows:

$$\text{Lower Limit NF} = \begin{bmatrix} 0.00 & -2.535 & 0.950 & 2.455 & -1.060 \\ 2.423 & 0.00 & -3.060 & -0.086 & 2.945 \\ -1.044 & 2.953 & 0.00 & 3.929 & -4.052 \\ -2.568 & -0.032 & -4.078 & 0.00 & -0.104 \\ 0.965 & -3.051 & 3.904 & -0.087 & 0.00 \end{bmatrix}$$

The results for the upper bound of the 95 percent confidence interval for the matrix **NT** are given in matrix form as follows:

$$\text{Upper Limit NF} = \begin{bmatrix} 0.00 & -2.423 & 1.044 & 2.568 & -0.965 \\ 2.535 & 0.00 & -2.953 & 0.032 & 3.051 \\ -0.950 & 3.060 & 0.00 & 4.078 & -3.904 \\ -2.455 & 0.086 & -3.929 & 0.00 & 0.087 \\ 1.060 & -2.945 & 4.052 & 0.104 & 0.00 \end{bmatrix}$$

Because it is known that the null hypothesis is false, at least some intervals should exclude zero. Inspection of the confidence limits given by **Lower Limit NF** and **Upper Limit NF** reveal that zero is excluded from all but 4 of the elements. Specifically elements (4,2) and (2,4), as well as (4,5) and (5,4). After inspecting matrix **NF**, it is realized that indeed the null hypothesis of symmetry is true for elements (4,2) and (2,4), as well as (4,5) and (5,4). Thus, these two pairs of elements are *correctly* bracketed by zero. The other 16 confidence limits correctly *excluded* zero, and correctly *included* the population value obtained from Equation 7.

Further Issues of the Procedure

One obvious concern with making use of the suggested procedures is that if one were to form multiple confidence intervals (for example the $(p^2 - p)/2$ independent confidence intervals) $100(1 - \alpha)$ percent confidence intervals, there would be inflation of the familywise Type I error rate. Because such an inflated Type I error rate is undesirable, the researcher should decide a priori whether they want to do a *confirmatory* analysis or an *exploratory* analysis.

In a confirmatory analysis the researcher prespecifies corresponding elements of interest and then performs the procedure only for the specified elements. One option is to use a Bonferroni correction, where if C confidence intervals are to be formed, the per-confidence interval Type I error rate is given as $\alpha_{PC} = \alpha_{FW}/C$, where α_{FW} is the familywise alpha. In this sense α_{PC} controls the Type I error rate, such that the confidence intervals done within the family have an overall (i.e., across the family) Type I error rate of the FW specified, while the PC error rate is smaller. Thus, if $C = 4$, because there were four confidence intervals of interest, and α_{FW} was set at .05, $\alpha_{PC} = .05/4 = .0125$. Thus, the Type I error

rate for each individual confidence interval would be .0125; however, if the null hypothesis of symmetry was true, across the family of tests a Type I error rate would be expected $100\alpha_{FW}$ percent of the time.

Rather than confirmatory analyses, some researchers will likely be interested in more exploratory analyses. In this sense there are no a priori hypotheses and the suggested procedures can be used as an exploratory technique to discern if there are likely any, and where, corresponding elements that are asymmetric. Because of the numerous confidence intervals formed, and because in the exploratory setting the Type I error rate need not be controlled across the set of elements, if the null hypothesis of symmetry is true there may still be significance claimed (i.e., a Type I error committed). Across the study (i.e., the familywise Type I error rate) will be approximately $1 - (1 - \alpha)^C$ (Hays, 1994, p. 341). This Type I error rate will be much greater than the per-confidence interval α if C is much greater than one. However, because the goal is to try and reject the hypothesis of symmetry (either across the set of elements or for prespecified elements), it is a better strategy to be cautious when attempting to adjust for familywise alpha inflation, so as to not make it exceedingly difficult to reject a false null hypothesis of symmetry.

Discussion

It is often the case that symmetry is assumed in the context of MDS, not based on evidence, but rather on convenience. The extent to which this symmetry assumption holds true has not often been tested, possibly in part because procedures to test such an assumption have not received much attention. Although MDS models do exist for asymmetric data (e.g., Okada and Imaizumi, 1997, Collins, 1987, Young and Lewycky, 1979), determining the appropriateness of such models is not often considered. The present article argues that the assumption of symmetry should be tested in applications of MDS and presents procedures that allow such an assumption to be tested by making use of the theory of the bootstrap technique.

Specifically, the suggested procedure is to make use of the theory of the bootstrap technique, such that nonparametric confidence limits can be obtained and then assessed to infer whether or not the confidence limits, either in a confirmatory or exploratory sense, likely contain zero. If zero is contained in the confidence interval, then the null hypothesis of symmetry cannot be rejected. If, however, zero is not contained within the confidence limits, it is inferred that the assumption of symmetry does not hold (at least for the element(s) of interest) in the population. When symmetry does not hold, treating the matrix of proximities as though it were symmetric can potentially lead researchers astray in attempting to understand the relationships that exist among a set of objects.

An issue not addressed in the present article concerns global fit indices. That is, a fit index designed to assess the likelihood of there being *any* asymmetries in a simultaneous fashion. The techniques proposed here are for confidence intervals for bounding the population difference between corresponding elements, not some overall measure of symmetry. Such a global measure would allow one to test an omnibus effect. Such a test would be analogous to an omnibus F test in an analysis of variance context, whereas the tests of specific confidence intervals, as discussed in the present article, would be analogous to pairwise comparisons. While the omnibus test would be useful, the specific elements contributing to the lack of symmetry would not be known. Follow-up tests of specific corresponding elements could then be carried out, but at that point the analysis essentially circles back to the procedures offered here. Statistical power and precision of some global fit index as well as for specific corresponding elements is another issue not addressed in the present article. A potentially interesting and useful topic for future work in this area is the effect of sample size on the power to detect asymmetric relations and the precision in bounding the population parameter.

In concluding, assuming symmetry in the context of MDS potentially leads to biased results and incorrect decisions when asymmetric

relationships exist. Identifying those asymmetric relationships can potentially be very informative, as such identification can lead to knowledge that had not been previously realized. Presumably such knowledge can then lead to a better understanding of the relation among a set of objects of interest, which will potentially have implications for furthering the knowledge base in the given area of research.

Acknowledgement

The author would like to thank Christof Schuster for helpful advice on a previous version of this article.

Footnotes

¹ Rothkopf (1957) used all letters of the alphabet and whole numbers from zero to nine. Thus, there were 36 different signals yielding 630 pairs of signals. The participants were required to be naive to International Morse Code signals before the study began and were only trained and tested on a subset of the 630 possible pairs of objects.

² Obtaining proximity values, and thus the array \mathbf{X} , in part depends on the type of information being collected. Davison (1983, chapter 3) and Cox and Cox (2001, section 1.3) detail methods of obtaining proximity values. Some measures of proximity are obtained directly by participants in the study. Other proximity measures require calculation (such as the Euclidean distance across a set of variables) before the \mathbf{X} array can be obtained. As Cox and Cox state, "choice of proximity measure depends upon the problem at hand, and is often not an easy task" (2001, p. 10). The present article makes the assumption that the desired proximity measure(s) have been obtained and that the values has been appropriately entered in the array \mathbf{X} .

³ An argument can be made that for measures of dissimilarity or similarity, the t -test is robust to violations of the assumption of normality and homogeneity of variance. This "robustness," however, depends in part on the sample size and the degree of nonnormality. Further, when measures of equivalence are of interest, one could argue that because the overall matrix of equivalences

is based on sums, the Central Limit Theorem says that the distribution of the means, in the limit, will be normal. Another option for measures of equivalences is to use the binomial distribution. However, rather than relying on asymptotic properties or supposed robust techniques, the method deemed here to be optimal is to avoid parametric assumptions altogether by making use of the suggested bootstrap techniques.

References

- Bell, P., and Cox, T. (1998). Measuring the variability in MDS configurations using shape analysis. In B. Marx and H. Friedl (Eds.), *Thirteenth international workshop in statistical modelling* (pp. 404-407). New Orleans, LA.
- Collins, L. M. (1987). Sociometry: Deriving sociograms via asymmetric multidimensional scaling. In R. M. Hamer (Ed.), *Multidimensional scaling: History, theory, and applications* (pp. 179-196). Hillsdale, NJ: Lawrence Erlbaum Associates.
- Cox, T. F., and Cox, M. A. A. (2001). *Multidimensional scaling* (2nd ed.). New York: Chapman and Hall/CRC.
- Davison, M. L. (1983). *Multidimensional scaling*. New York: John Wiley and Sons.
- Efron, B., and Tibshirani, R. J. (1993). *An introduction to the bootstrap*. New York: Chapman and Hall/CRC.
- Harshman, R. A., Green, P. W., Wind, Y., and Lundy, M. E. (1982). A model for the analysis of asymmetric data in marketing research. *Marketing Science*, 1, 205-242.
- Hays, W. L. (1994). *Statistics* (5th ed.). New York: Harcourt Brace College Publishers.
- Holman, E. W. (1979). Monotonic models for asymmetric proximities. *Journal of Mathematical Psychology*, 20, 1-15.
- Kruskal, J. (1964). Multidimensional scaling by optimizing goodness-of-fit to a nonmetric hypothesis. *Psychometrika*, 29, 1-27.
- Okada, A., and Imaizumi, T. (1997). Asymmetric multidimensional scaling of two-mode three way proximities. *Journal of Classification*, 14(2), 195-224.
- Rothkopf, E. (1957). A measure of stimulus similarity and errors in some paired-associate learning tasks. *Journal of Experimental Psychology*, 53, 94-101.
- Seiyama, K., Naoi, A., Sato, Y., Tsuzuki, K., and Kojima, H. (1990). Stratification structure of contemporary Japan and its trend. In A. Naoi and K. Seiyama (Eds.), *Social stratification in contemporary Japan: Structure process of social stratification* (Vol. 1, pp. 15-50). Seiyama, Tokyo: Tokyo University Press.
- Tversky, A. (1977). Features of similarity. *Psychological Review*, 84(4), 327-352.
- Young, F. W. (1987). Weighted distance measures. In R. M. Hamer (Ed.), *Multidimensional scaling: History, theory, and applications* (pp. 117-158). Hillsdale, NJ: Lawrence Erlbaum Associates.
- Young, F. W., and Lewyckyj, R. (1979). *ALSCAL users guide*. [Computer software]. Carboro, NC: Data Analysis and Theory Associates.

Appendix A

The following syntax is for the computer programs R and S-Plus. The 'x' represents the specific values or the observed data that must be supplied by the researcher. The text that follows pound signs (i.e., #) are comments and need not be typed by the researcher, as they are included only to help to understand the syntax.

S-Plus and R Code

```
# Requires a p by p by N array labeled 'Full.Data'.

alpha <- .05
# Type I error rate; 1-confidence interval coverage.

p <- x
# Number of measures

N <- x
# Number of individuals (or the number of replications for a single individual)

B <- x
# Number of bootstrap replications

Full.Data <- x
# 'Full.Data' must be input either as an N by p by p array or
# as a concatenated set of individual matrices in the array
# function with defined dimensions, that is,
# 'Full.Data <- array(c(a1, a2, a3,..., aN), dim=c(N, p, p))'
# where each 'ai' is an individual's matrix.

# 'Full.Data' must be input either as an N by p by p array or
# as a concatenated set of individual matrices, that is,
# 'Full.Data <- c(a1, a2, a3,..., aN),' where each 'ai' is an
# individual's matrix

Discrepancy.Matrices <- array(NA, dim=c(p, p, B))
Marker <- seq(1, N, 1)
# Defines the 'storage' array for the bootstrap results.

for(b in 1:B)
{
  Sampled.Matrix <- array(0, dim=c(p,p))
  # Storage matrix for the randomly sampled bootstrap samples

  Matrix.Identifier <- sample(Marker, N, replace=T)

  for(i in 1:N)
  {
    Sampled.Matrix <- Sampled.Matrix + Full.Data[, ,Matrix.Identifier[i]]
  }
}
```

(Appendix A continued on following page.)

(Appendix A continued from previous page.)

```

# Randomly samples N matrices (with replacement) 1 at a time from the
# collection of N individual matrices
}

Sampled.Matrix <- 100*(Sampled.Matrix/N)
# Turns the overall matrix into a matrix of proportions by dividing by N.

Symetrized.SampleMatrix <- (Sampled.Matrix + t(Sampled.Matrix))/2
# Takes the full randomly sampled matrix (i.e., sum of matrices) and gets
# the mean of the corresponding upper and lower diagonal elements
# to have a symmetric matrix.

Discrepancy.Matrices[, ,b] <- Sampled.Matrix - Symetrized.SampleMatrix
# Records the discrepancy for each element of the matrix. This can be
# used for comparing discrepancies of individual elements or for
# an overall distribution of discrepancies.
}

Obs.Matrix <- array(0, dim=c(p,p))
for(i in 1:N)
{
Obs.Matrix <- Obs.Matrix + Full.Data[, ,i]
}
Obs.Matrix <- 100*(Obs.Matrix/N)
Obs.Symetrized.Matrix <- (Obs.Matrix + t(Obs.Matrix))/2
Obs.Symetrized.Matrix
# The above syntax creates the overall observed matrix.
# This is useful if the N individual matrices were
# input individually and interest is in the overall matrix.

# After the program has finished, descriptive statistics can then be formed,
# such as various percentiles of interest to yield confidence limits
# for the particular element from the 'Discrepancy matrix.'
# An example of calculating a 100(1-alpha)% confidence interval (here a 95)
# percent confidence interval for the first row second element is given as follows:

quantile(Discrepancy.Matrices[1,2,], probs=c(alpha/2, 1-alpha/2))
mean(Discrepancy.Matrices[1,2,])
sqrt(var(Discrepancy.Matrices[1,2,]))

# The elements defined in the brackets are changed to reflect the elements
# of interest. The program does not need to be run each time, only once and
# then the values in the brackets adjusted to reflect the element(s) of interest.

```