not only to the inter-personal distances but also to the 'intrapersonal separations' of the individuals' two 'ideal points'. As might be expected, the intra-personal distances are small compared to the interpersonal ones, but appear to be otherwise uninterpretable. Gleason (ibid. p. 118) coyly concludes, 'Unfortunately, very little else can be said about these distances at this time', and indeed he himself symmetrises the data for performing hierarchical clustering. Such considerations strongly indicate the use of a one-mode conditional proximity approach of MINICPA.

Whilst this example is more useful as an example of what can be done than providing a good illustration of its utility, further procedures for scaling such sociometric data are given in Breiger, Boorman and Arabie (1975), and MINICPA provides a useful addition to the methods discussed there. The Newcomb data form the test data for the MINICPA program in the MDS(X) version.

6.1.3 Triadic data analysis (TRISOSCAL)

Concisely: TRISOSCAL (TRIadic Similarities Ordinal SCALing) provides:

internal analysis of a set of triadic (dis)similarity measures
by a Minkowski distance model
using a local or global monotonic transformation of the data.

Triadic data are collected especially by psychologists, sociologists and anthropologists wishing to elicit the constructs which subjects use in making judgments of similarity. The basic idea is that the subject is asked to consider groups of three objects at a time, taken from the full set. Two forms are in common use:

partial triadic data, where from each presentation of the three objects the subject is asked to pick out only the single most similar pair; and
full triadic data where both the most similar pair and the least similar pair are picked out. The intermediate pair can then simply be inferred.

Although the method of triads is a useful technique for data collection, the number of triads increases very rapidly with the number of objects (for $p = 5, 10, 15$ and 20 the number of triads is 10, 120, 455 and 1140 respectively). Obviously, beyond about $p = 8$, the presentation of the full set of triads becomes totally unfeasible and very taxing on the subject. Burton and Nerlove (1976) give a full description and discussion of experimental designs for minimising the number of triads presented whilst maximising the information gained.

The most important advantage which triadic data possess compared to simpler forms is that contextual effects of judgment can be directly examined—by examining whether the similarity between two objects remains the same when the third element is changed.* It is therefore unfortunate when triadic data are turned into 'vote-count' data before scaling, since the effect is to obliterate the triadic information. (In brief, the vote-count method consists of counting the number of

*The assumption that the judgment remains unchanged in the presence of irrelevant alternatives is important in a number of theories of choice, preference and social welfare: see Rescher (1969).

times that object j is judged more similar than object k in the data.) Roskam (1970, p. 406) has shown that such a procedure badly misrepresents the order information in the data and often results in ill-fitting scaling solutions.

The TRISOSCAL program provides a direct method for scaling triadic data non-metrically by a distance model. It differs from Roskam's original MINITRI program in allowing the user to decide either that the order information implied across *all* the triads be fitted systematically ('global stress'), or only that the separate orders *within* each triad be fitted.

The distinction can be illustrated as follows. Suppose triadic data have been collected from a group of individuals. That being so, it is quite likely that when presented with objects (A, B, C) one subject will decide that (AB) is most similar and (AC) least similar (implying that $d(A, B) \leqslant d(B, C) \leqslant d(A, C)$). Another subject when presented with the same triad, may decide just the opposite—that the pair (AC) is most similar, and (AB) least similar (implying that $d(A, C) \leqslant d(B, C) \leqslant d(A, B)$). Both agree, by implication, that $d(B, C)$ is intermediate—but how shall the conflicting information concerning $d(A, C)$ and $d(A, B)$ be fitted? The answer proposed by Roskam's 'local stress' approach is: treat each subject's triad, that is, judgment, as a distinct entity, fit each of the three distances within a triad separately and then define 'the' overall fitting value as the average of the different disparity values. Hence there will be as many disparity or fitting values (d^0_{jk}) as occurrences of the pair (j, k) in the triads data. In this instance, there will be two distinct values of d^0_{ab} and d^0_{ac}, but their respective arithmetic average (denoted $\bar{d}^0_{ac}$) will represent 'the' fitting value for the pair concerned. Roskam suggests that the form of raw stress to be minimised should be:

Triads: 'Local' Stress

$$\text{Stress}_0 = \sum n_{jk}(d_{jk} - \bar{d}^0_{jk})^2$$

where n_{jk} is the number of times that the pair (jk) occurs in the set of triads to be analysed.

Roskam's approach obviously tolerates a good deal of inconsistent data, but is a sensible strategy when data from a set of individuals is combined, or where the number of objects is large.

Prentice (described in Coxon and Jones 1979, p. 49 et seq.) suggests a more restrictive and stringent approach: to require total consistency *between* triads and to count each and every infraction of transitivity in the stress value rather than averaging. Consider the following two pieces of triadic data:

	Triad	Most similar (MS) pair	Least similar (LS) pair	Dissimilarity data implied
1	(A, B, C)	(AB)	(BC)	$(AB) < (AC) < (BC)$
2	(B, C, D)	(BC)	(CD)	$(BC) < (BD) < (CD)$

Taken together, the information from both triads is consistent (transitive) and implies the following order of data:

$$AB < AC < BC < BD < CD.$$

Prentice's global approach requires that the data be fitted by disparity values in the same order, and in particular he requires that the dissimilarity (BC) be fitted by the *same* disparity value in both Triad 1 and in Triad 2.

The two approaches and the different fitting values which result are illustrated in Table 6.2 for the set of data just given. Assume that there is a configuration of four points whose interpoint distances are as in the second column of the table. Column 1 gives the pairs in the two triads in their correct order according to the data, and column 2 gives the corresponding distances in the current configuration (note that allowance has to be made for two fitting values for (BC), although there can obviously be only one such distance in the configuration).

How well does the current configuration match the data according to the criteria of local and global stress? In the next three columns (3–5) the necessary fitting values are calculated using monotone regression.

Column 3 Local stress allows the fitting value for (BC) in the first triad to be different to the fitting value for (BC) in the second triad, so the d values are calculated separately within each triad, yielding $(BC)_1 = 2\frac{1}{2}$ (the block average of 3 and 2) and $(BC)_2 = 1\frac{1}{2}$ (the block average of 2 and 1).

Column 4 'The' fitting value for (BC) according to Roskam's local stress is defined as the average of its two appearances, i.e.

$$\bar{d}^0_{bc} = (d^0_{(bc_1)} + d^0_{(bc_2)})/2 = (2\tfrac{1}{2} + 1\tfrac{1}{2})/2 = 2$$

Column 5 The fitting values according to Prentice's global stress approach, which requires weak monotonicity over all the pairs, produce a *single* fitting value for (BC), which by block-averaging over all but the first and last pair, gives a value

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
	DATA Pair (ij)	DISTANCE d_{ij}	FITTING VALUE DISPARITIES d^0 (local)	$\bar{d}^0$ (local)	$\hat{d}$ (global)	RESIDUALS $(d - d^0)^2$ (local)	$(d - \hat{d})^2$ (global)
Triad	(AB)	1	1	1	1	0	0
1	(AC)	3	$2\frac{1}{2}$	$2\frac{1}{2}$	2	$\frac{1}{4}$	1
	$(BC)_1$	2	$2\frac{1}{2}$	2	2	0	0
Triad	$(BC)_2$		$1\frac{1}{2}$	2		0	0
2	(BD)	1	$1\frac{1}{2}$	$1\frac{1}{2}$	2	$\frac{1}{4}$	1
	(CD)	4	4	4	4	0	0
					Total	$\frac{1}{2}$	2

Raw Stress (local) $= \frac{1}{2}$
Raw Stress (global) $= 2$

Table 6.2 *Triadic data: illustration of local and global stress* of 2. (It so happens that in this example both local and global approaches give the same fitting value for d_{bc}; this will not often be the case.)

The squared discrepancy values contributing to raw stress are given in column 6 for local and in column 7 for global stress, and the stress values are given at the foot of the table. Note that, as expected, requiring a complete ordered fit over all the pairs of both triads increases the badness-of-fit. Note also that global stress assumes the same value (2) for all but two pairs, but *is* weakly monotone with the order implied by the data. Local stress, by contrast, not only fits the same distance (*BC*) by two distinct values but also averages them to obtain the 'overall fitting value' which (as in this case) usually will *not* be in the same order as that implied by the triads:

$$\begin{array}{lccccccccc} \text{data:} & AB & < & AC & < & BC & < & BD & < & CD \\ \bar{d}^0: & 1 & < & 2\tfrac{1}{2} & \nless & 2 & \nless & 1\tfrac{1}{2} & < & 4 \end{array}$$

In the case of inconsistent sets of triads, even greater inversions occur. When triadic information comes from a number of subjects or sources, highly inconsistent (intransitive) data often result and the user is faced with a difficult choice: either to choose global stress and risk technically degenerate solutions (since the inconsistencies can only be dealt with by imposing the same fitting value on a large number of pairs, capitalising on weak monotonicity) or to choose local stress, and lose all information about the order across triads as well as drowning same-pair inconsistencies by fitting averaged disparities. At least the Prentice approach signals the inconsistencies and global intransitivities by a high stress value—though at the disadvantage of increased computing costs.*

In such situations the user is advised to scale at least a random subset of data (to save computing time), using the global stress approach, and to compare resulting stress values and configurations with those obtained using local stress. Whichever stress option is chosen, triadic data are scaled in their integrity, and as far as possible consistency is kept within triads in the configuration. But only the global stress approach tries to represent information *between* the triads.

In some applications it is advantageous to insist upon between-triad comparisons by choosing the global approach, as when context effects are to be minimised. But by the same token context effects can best be studied by collecting triadic data and the user can then examine whether a given pair of objects tends to be fitted by approximately the same value. This can only be done by choosing the local stress approach.

6.1.4 *The basic metric model* (MRSCAL)

Concisely: MRSCAL (MetRic SCALing) provides:

internal analysis of two-way data of a lower triangle format of a (dis)similarity measure

by a Minkowski distance function,

using a linear and/or logarithmic transformation of the data.

*In Coxon and Jones 1978a, global stress_2 values of over 0.95 were frequently observed for such heterogeneous data sets, with corresponding local stress_1 values of around 0.20. In one case, 169 triadic comparisons of 13 occupations made by a set of policemen produced a global stress_2 value of 0.960 when scaled, and 75 out of the 78 global $\hat{d}$ values had the same value!