

Nonmetric Strain

Jan de Leeuw

March 6, 2026

The three major loss functions for multidimensional scaling, in order of decreasing popularity, are stress, sstress, and strain. The strain loss function has been applied mostly in classical metric MDS of complete distances matrices, although some extensions to additive constants, bound constraints, missing data, and ordinal scaling have been discussed in the literature. We use alternating least squares (ALS) and majorization (MM) to develop various forms of non-metric MDS based on strain minimization, together with corresponding programs in R.

Table of contents

1	Introduction	3
2	Alternating Least Squares	5
3	Special Cases	6
3.1	The Additive Constant	6
3.2	Imputing Missing Data	6
3.3	Bound Constraints	7
3.4	Ordinal MDS	8
4	Software	10
5	Examples	11
5.1	The Additive Constant	11
5.2	Imputing Missing Data	11
6	Discussion	13

Note: This is a working manuscript which will be expanded/updated frequently. All suggestions for improvement are welcome. All Rmd, tex, html, pdf, R, and C files are in the public domain. Attribution will be appreciated, but is not required. The files can be found at <https://github.com/deleeuw/strain>

1 Introduction

In Mutidimensional Scaling (MDS) the strain loss function is defined as

$$\sigma(X, \Delta) := \frac{1}{4} \text{tr} \{J(\Delta^2 - D^2(X))J\}^2 \quad (1)$$

Here J is the centering matrix of order n , i.e. $J := I - n^{-1}E$, with E a matrix with all elements equal to one. Matrix Δ^2 has squared dissimilarities and $D^2(X)$ contains the squared Euclidean distances between the rows of an $n \times p$ column-centered configuration matrix X .

We will also use another representation of strain in this paper. By applying the classical embedding result of Schoenberg (1935) and Young and Householder (1938), in the version proposed by Torgerson (1958), we find

$$\sigma(X, \Delta) = \text{tr} (C - XX')^2. \quad (2)$$

with

$$C := -\frac{1}{2}J\Delta^2J. \quad (3)$$

In *metric* MDS the matrix Δ is constant and given, and we minimize strain over X . From (2) the basic theorem of Classical Scaling (Torgerson (1958)), also known as Principal Coordinate Analysis (Gower (1966)), tells us that the optimal X are the eigenvectors corresponding with the p largest eigenvalues of C_+ , the positive definite part of C , with each eigenvector scaled to have norm equal to the square root of the corresponding eigenvalue. There are rigorous proofs in Keller (1962) and De Leeuw (1974a).

In non-metric MDS we have to minimize strain over both X and Δ , where Δ varies in some subset of the space $\mathfrak{D}$ of non-negative, symmetric and hollow matrices of order n . Classical examples, treated in detail in this paper, are the additive constant problem, the missing data problem, the use of bound constraints, and ordinal non-metric scaling.

The history of various tentative and/or incomplete attempts to construct a non-metric MDS technique based on strain is in De Leeuw and Heiser (1982) and Trosset (1998). The first systematic approach to the problem is in De Leeuw (1974b), who discusses the additive constant problem, the missing data problem, and the ordinal forms of MDS. There is not software in De Leeuw (1974b), only equations. The paper proposes a minimization approach based on what Trosset (1997) calls “variable reduction”. The minimum of strain over X , which equals the sum of squares of the discarded eigenvalues of C_+ , is minimized over Δ . Results from classical perturbation theory are used to give expression for the first and second derivatives of this “reduced” loss function, with the implicit understanding that these derivatives can be used in a large variety of general-purpose optimization methods.

An early implementation of the same methodology for the additive constant problem is in Saito (1978). Since the “reduced” strain in that case only depends on a single parameter (the additive constant) Saito uses a simple bisection-type search procedure on the positive real line for minimization.

The a multi-parameter non-metric scaling problem occurs when some of the δ_{ij} are missing, either by accident or by design. Thus

$$\delta_{ij} = \begin{cases} \delta_{ij}^0 & \text{for non-missing,} \\ \theta_{ij} & \text{for missing.} \end{cases} \quad (4)$$

Thus the number of additional parameters is equal to the number of missing elements. We will require the θ_{ij} to be non-negative. An important special case of the missing data problem is the unfolding problem, in which we have similarities between objects in two disjoint sets. The between-set distances are known, while both matrices of within-set distances are missing.

A third important form of non-metric MDS is ordinal MDS. We require the δ_{ij} to non-negative and monotone with the δ_{ij}^0 . More precisely if $\delta_{ij}^0 \geq \delta_{kl}^0$ then we must have $\delta_{ij} \geq \delta_{kl} \geq 0$.

A systematic and general approach, similar to but much more detailed than De Leeuw (1974b), was followed by Michael Trosset in a sequence of papers (Trosset (1997), Trosset (2000), Trosset (1998), Trosset (2002)). Trosset addresses the additive constant problem, bound constraints, the missing data problem (as a special case of bound constraints), and ordinal MDS. He also compares variable reduction methods with “variable alternation” methods, which iteratively alternate minimization over X with minimization over Δ . For the variable reduction methods, which he favors, he relies on established general purpose mathematical programming routines to minimize the reduced loss function.

In this paper we use a methodology which differs from Trosset’s in two important respects. We propose and implement variable alternation methods instead of variable reduction methods, and we use majorization (MM) to minimize over X . This produces a class of flexible techniques with monotone convergence and relatively small storage requirements that can be adapted easily to alternative constraints on X and/or Δ . We have no actual comparisons, but theoretical considerations suggest that specialized variable reduction techniques can be made to converge faster than variable alternation techniques. This depends on the well-known trade-off between the number of iterations required for convergence and the amount of computation in a single iteration.

2 Alternating Least Squares

In our technique for strain minimization we apply alternating least squares (ALS) to loss function (1). We start with some initial estimate $\Delta^{(0)}$. Then, in each iteration, there are two subproblems. We first minimize strain over X with Δ fixed at its current value, and then minimize over Δ in $\mathfrak{D}$ with X fixed at its current value, which is the optimal value from the first subproblem in the iteration. Then we test for convergence, and if the loss function still decreases we go to the next iteration.

In iteration k this means

$$X^{(k)} = \underset{X}{\operatorname{argmin}} \sigma(X, \Delta^{(k-1)}), \quad (5a)$$

$$\Delta^{(k)} = \underset{\Delta \in \mathfrak{D}}{\operatorname{argmin}} \sigma(X^{(k)}, \Delta) \quad (5b)$$

If a subproblem in an iteration is too complicated for some reason, then we merely make a step towards its minimum, making sure that

$$\sigma(X^{(k)}, \Delta^{(k)}) < \sigma(X^{(k)}, \Delta^{(k-1)}) < \sigma(X^{(k-1)}, \Delta^{(k-1)}). \quad (6)$$

In the first subproblem minimizing strain over X for fixed Δ is done by classical metric multidimensional scaling (Torgerson (1958), Gower (1966), Mardia (1978)). The fact that computing the classical scaling solution is equivalent to minimizing the loss function (1) was first pointed out by De Leeuw and Heiser (1982).

The second subproblem, minimizing strain of Δ for fixed X (and thus fixed D), can take various forms. We discuss some of the important special cases, and their implementation in the following sections.

3 Special Cases

3.1 The Additive Constant

The simplest, and oldest, form of non-metric scaling estimates the additive constant. Thus $\delta_{ij} = \delta_{ij}^0 + \theta$ for all $i \neq j$, where δ_{ij}^0 are the given dissimilarities. Only one additional parameter needs to be estimated, with the side-condition that the adjusted δ_{ij} are non-negative.

If a constant θ is added to the off-diagonal dissimilarities the disparities are

$$\Delta^2 = \Delta_0^2 + 2\theta\Delta_0 + \theta^2(E - I). \quad (7)$$

Because we want the δ_{ij} to be non-negative we require that $\theta \geq -\min \Delta_0$, where the minimum is taken over all off-diagonal elements of Δ .

From (7) we have, with

$$G := J\Delta_0 J, \quad (8a)$$

$$H := J(\Delta_0^2 - D^2), \quad (8b)$$

that

$$\sigma(\theta) := \frac{1}{4} \text{tr} \{J(\Delta_0^2 + 2\theta\Delta_0 + \theta^2(E - I) - D^2)J\}^2 = \frac{1}{4} \text{tr} \{H + 2\theta G - \theta^2 J\}^2, \quad (9)$$

which is a quartic in θ . Expanding and simplifying gives

$$\sigma(\theta) = \frac{1}{4} \{ \text{tr} H^2 + 4\theta \text{tr} HG + 2\theta^2 (\text{tr} 2G^2 - \text{tr} H) - 4\theta^3 \text{tr} G + \theta^4(n - 1) \}. \quad (10)$$

Minimizing over θ means finding the root of the cubic $\mathfrak{D}\sigma$ that minimizes the quartic (10). Details are in the software section of the paper.

3.2 Imputing Missing Data

If there are missing data the classical scaling method to minimize strain must be adapted. In that case there is a set $\mathcal{M}$ of m index pairs that code for missing data. Matrix Δ_0^2 now has zeroes for missing data and squared dissimilarities for the non-missing ones. We have

$$\Delta^2 = \Delta_0^2 + \sum_{(i,j) \in \mathcal{M}} \theta_{(i,j)} E_{(i,j)}, \quad (11)$$

where $E_{(i,j)} := e_i e_j' + e_j e_i'$ and e_i has element i equal to one and all other elements zero. If D^2 are the current squared distances, then strain is the function of θ defined as

$$\sigma(\theta) := \text{tr} \left\{ H - \sum_{(i,j) \in \mathcal{M}} \theta_{(i,j)} J E_{(i,j) \in \mathcal{M}} J \right\}^2, \quad (12)$$

where H is as defined before in (8b).

Minimizing σ from (12) over the m element vector $\theta \geq 0$ is a non-negative linear least squares problem, which can be solved by various easily available techniques.

3.3 Bound Constraints

Bound constraints are of the form $\Delta_{\text{low}} \leq \Delta \leq \Delta_{\text{high}}$, where the inequality signs refer to elementwise comparisons. Thus each δ_{ij} is required to be in a given interval (which may be infinitely large). By choosing the interval to be the whole non-negative real axis this will give us another way, using majorization, to impute missing data. For non-missing data the interval has length zero and is equal to the dissimilarity δ_{ij}^0 .

By using the vec of our matrices Δ and D loss function strain can be written in the computationally more friendly form

$$\sigma(\delta) = \frac{1}{4}(\delta^2 - d^2)(J \otimes J)(\delta - d^2), \quad (13)$$

with $J \otimes J$ the Kronecker product.

Equation (13) shows that minimizing strain over δ is a weighted linear regression problem. It is more complicated than standard linear regression because of the weights $J \otimes J$, which are not diagonal. Thus strain as a function of δ^2 , although convex, is not separable. In addition $J \otimes J$ is singular, having $(n-1)^2$ unit eigenvalues and $2n-1$ zero eigenvalues.

Minimizing ... over δ^2 within the bounds can be solved by any quadratic programming method, as long as we are careful to handle the singularities appropriately. We will go an alternative route, using majorization to reduce the weighted least squares problem to a sequence of unweighted problems, in which we can trivially handle the bound constraints. Related results on majorizing weighted by unweighted least squares are in Kiers (1997) and Groenen, Giaquinto, and Kiers (2003).

Suppose $\tilde{\Delta}$ is the current best value of the disparities. Expand strain around $\tilde{\Delta}$.

$$\sigma(\Delta) = \sigma(\tilde{\Delta}) + \frac{1}{2} \text{tr} J(\Delta - \tilde{\Delta}) J(\tilde{\Delta} - D^2) + \frac{1}{4} \text{tr} J(\Delta - \tilde{\Delta}) J(\Delta - \tilde{\Delta}), \quad (14)$$

Because $J \otimes J \preceq I$ we have $\sigma(\Delta) \leq \eta(\Delta, \tilde{\Delta})$, with the majorization function η defined as

$$\eta(\Delta, \tilde{\Delta}) := \sigma(\tilde{\Delta}) + \frac{1}{2} \text{tr } J(\Delta - \tilde{\Delta})J(\tilde{\Delta} - D^2) + \frac{1}{4} \text{tr } (\Delta - \tilde{\Delta})^2. \quad (15)$$

The majorization, or MM, algorithm minimizes η over vectors Δ in the hyper-rectangle of vectors within the bounds to find the update Δ^+ .

To simplify the expression for η we “complete the square”. Relying heavily on the idempotency of J , and thus of $J \otimes J$, we find

$$\eta(\Delta, \tilde{\Delta}) := \frac{1}{4} \text{tr } (\Delta - \hat{\Delta})^2, \quad (16)$$

with

$$\hat{\Delta} := \tilde{\Delta} - J(\tilde{\Delta} - D^2)J. \quad (17)$$

Thus $\hat{\Delta}$ is a matrix-weighted average of the “conservative” part $\tilde{\Delta}$ and the “progressive” part D^2 (Chamberlain and Leamer (1976)).

Finally the update Δ^+ is the unweighted least squares projection of $\hat{\Delta}$ from on the hyper-rectangle defined by the bounds Δ_{low} and Δ_{high} , i.e. the unweighted bounded regression on $\hat{\Delta}$. The update Δ^+ is simply

$$\Delta^+ = \max(\Delta_{\text{low}}, \min(\hat{\Delta}, \Delta_{\text{high}})) \quad (18)$$

where $\max()$ and $\min()$ are applied elementwise.

We now have the sandwich inequality

$$\sigma(\Delta^+) \leq \eta(\Delta^+, \tilde{\Delta}) \leq \eta(\tilde{\Delta}, \tilde{\Delta}) = \sigma(\tilde{\Delta}). \quad (19)$$

In the substep of ALS in which we improve Δ we make a number of “inner” majorization iterations. In the later “outer” ALS iterations we will already have a good initial estimate to start these “inner” iterations.

3.4 Ordinal MDS

In ordinal MDS for fixed D strain is a function of the disparities Δ , which are required to be monotone with Δ_0 . Thus, using the vec operator in the same way as in the previous section, and rearranging

$$\sigma(\delta) = \frac{1}{4}(\delta - d^2)(J \otimes J)(\delta - d^2), \quad (20)$$

which must be minimized over increasing δ . Note that we fit disparities δ to squared distances. If there are missing data we do include them in the loss, but we do not require them larger or smaller than other δ 's.

Equation (20) shows that minimizing strain over δ is a non-separable weighted monotone regression problem. Routines such as PAVA cannot be used (De Leeuw, Hornik, and Mair (2009)). We will go the same majorization route as the one we used for bound constraints. Thus in each substep we have to minimize η from (16) over increasing δ , which is unweighted monotone regression.

There is one complication in the ordinal case. For additive and bound constraints the set $\mathfrak{D}$ of admissible transformations is a polyhedral convex set which does not contain the origin $\Delta = 0$. But for ordinal MDS the origin does satisfy the ordinal constraints, and thus minimizing strain can be done trivially by setting $\Delta = 0$ and $X = 0$. Loss is zero, which is obviously the global minimum. In order to avoid this trivial solution we need to use a normalization condition on either X or Δ which excludes the origin from $\mathfrak{D}$. In this paper we use $\text{tr } F^2 = 1$, where $F := J\Delta^2 J$. Thus $\mathfrak{D}$ becomes the intersection of a polyhedral convex cone and the unit sphere in the space of transformed dissimilarities. It has been shown by De Leeuw (1975) that projecting on this set can be done by normalizing the projection on the cone, i.e. by normalizing the monotone regression. For a more general and rigorous treatment of the same problem, see Bauschke, Bui, and Wang (2018).

4 Software

For the four special cases we have discussed there are four R programs (R Core Team (2025)) in the github repository.

- `strainSSAddOne()`
- `strainSSMissing()`
- `strainSSBound()`
- `strainSSOrdinal()`

For this we use the R package `RSpectra` (Qiu and Mei (2024)). If $\Lambda_p \geq 0$ then the optimal X is $K_p \Lambda_p^{\frac{1}{2}}$. If some elements of Λ_p are negative, we first replace them by zeroes.

We minimizing the quartic over the half-open interval $\theta \geq -\min \Delta$ using the formulas in Jeffrey (1997). We implemented Jeffrey's solution in the C function `jeffrey()`, which is in the shared library `jeffrey.so`. The R program `strainSSAddOne()` loads the shared library and iteratively executes the two ALS sub problems.

that we solve by using the R function `nnls()` from the package with the same name (Mullen and van Stokkum (2024)). The R function `strainSSMissing()` calls `nnls()`, as well as `eigs_sym()` from `RSpectra`, to solve the two ALS subproblems.

5 Examples

5.1 The Additive Constant


Our example is taken from Torgerson (1958), pages 280-290. Dissimilarity judgments of nine Munsell colors of the same red hue, but differing in brightness and saturation, where collected from 38 subjects using the method of triads. The matrix Δ_0 , which Torgerson calls the “comparative distances” is taken from his Table 5 on p 286.

	[,1]	[,2]	[,3]	[,4]	[,5]	[,6]	[,7]	[,8]	[,9]
[1,]	0.00	-2.37	-0.12	-0.62	0.23	1.56	1.09	2.02	2.23
[2,]	-2.37	0.00	-1.01	-1.93	-0.90	0.80	-0.47	1.05	0.78
[3,]	-0.12	-1.01	0.00	0.70	-1.32	-0.67	1.07	0.70	2.62
[4,]	-0.62	-1.93	0.70	0.00	-0.78	1.25	-1.75	0.28	-0.72
[5,]	0.23	-0.90	-1.32	-0.78	0.00	-1.02	-1.23	-1.65	0.49
[6,]	1.56	0.80	-0.67	1.25	-1.02	0.00	0.57	-0.67	1.88
[7,]	1.09	-0.47	1.07	-1.75	-1.23	0.57	0.00	-1.18	-1.30
[8,]	2.02	1.05	0.70	0.28	-1.65	-0.67	-1.18	0.00	0.42
[9,]	2.23	0.78	2.62	-0.72	0.49	1.88	-1.30	0.42	0.00

Torgerson (p, 286-287) uses an elegant, but elaborate, method to find an additive constant value of 3.60. His technique is not based on minimization of an explicit loss function. Iterating `strainSSAddOne()` until strain changes less than 10^{-10} from one iteration to the next, starting with $\theta = 0$ uses 196 iterations and find $\theta = 2.85$. Starting with Torgerson’s 3.6 value for the additive constant gives the same θ in 184 iterations. Of course the 10^{-10} stopping criterion is much too strict for most applications, although a run with even higher precision shows the value for the optimal additive constant only has three stable decimals. Using a more realistic 0.001 requires 105 iterations and gives two stable decimals. In other words, convergence is slow.

5.2 Imputing Missing Data

We give two examples. The first uses the classical Morse code data from Rothkopf (1957). The data are confusion probabilities between 36 auditory Morse code signals, collected from 598 airmen. For details of the design we refer to Rothkopf. The data are available, for example, in the `smacof` package (De Leeuw and Mair (2009)). With 36 objects there are 630 dissimilarities. We ran the `strainSSMissing()` function with 0, 5, 10, 50, and 100 random dissimilarities missing. The five configurations are matched with `procrustus()` from `procrustus.R`, which minimizes $\sum_k \|Z - X_k M_k\|^2$ over Z and the square orthonormal M_k .



The plot shows a stable solution. Points are labeled “0” for no missing, “1” for 5 missing, and so on. The red points are for the matched Z matrix. The number of iterations for 0, 5, 10, 50, 100 missing data is 0, 10, 10, 15, 23 and strain is 6.7478705, 6.7378719, 6.6407385, 6.3165407, 5.550715.

roskam

6 Discussion

Most MDS loss functions are of the form

$$\sigma(X) := \sum_k w_k (f(\delta_k) - f(d_k(X)))^2 \quad (21)$$

for some given function f . Groenen, De Leeuw, and Mathar (1995) call loss function (21) fStress. For f equal to the identity we have stress (Kruskal (1964)), for f the square we have sstress (Takane, Young, and De Leeuw (1977)), and for f the logarithm we have log-stress (Ramsay (1977)). If we compare (21) and strain we see that strain is more complicated because of the non-diagonal Kronecker weights. This makes the usual optimal transformations in MDS, such as monotone regression in the ordinal case, more complicated and costly. As a consequence non-metric MDS techniques minimizing strain have not been implemented and used to the same degree as techniques based on variations of fStress.

On the other hand in fStress the minimization over X for given Δ is more costly than it is for strain, in which we simply have to find the p largest eigenvalues and corresponding eigenvectors of a matrix. Thus, the ALS framework, for fStress finding optimal Δ for fixed x is cheap and finding optimal X for fixed Δ is expensive. For strain it is exactly the other way around. To put it even more sharply, in fStress finding the optimal X is impossible, because it requires an infinite iterative process which can converge to a local minimum. For strain, although solving the eigen-problem is also an infinite iterative process, it quickly gets us as close as we want to the global minimum (for given Δ , that is).

Trosset (2002)

Trosset (1998) discusses an implementation of ordinal MDS using strain. His approach is quite different from ours, because it does not use ALS but gradient projection. More importantly, Trosset projects out the optimization over X . He calls this “variable reduction”, to contrast it with “variable alternation”. Thus his strain is a function of Δ only, defined as

$$\sigma(\Delta) = \min_X \text{tr} (C(\Delta) - XX')^2 = \sum_{r=p+1}^n \lambda_r^2(C_+(\Delta)), \quad (22)$$

where the λ_r are the $n-p$ smallest eigenvalues of its argument, and with C_+ the positive definite part of C . This loss function is then minimized over Δ by gradient projection, constrained by the linear inequalities that impose monotonicity. It is unclear how this compares to ALS, although we can say that in general projecting out a set of unknowns tends to accelerate convergence, at the cost of more complicated subproblems in each iteration. Also, unlike ALS, gradient projection implies choosing a step-size procedure to guarantee convergence. Projection seems especially suitable for the additive constant problem in which (22) is a function of a single parameter, which can be minimized in a number of ways.

weighted strain

References

- Bauschke, Heinz H., Minh N. Bui, and Xianfu Wang. 2018. "Projecting onto the Intersection of a Cone and a Sphere." *SIAM Journal on Optimization* 28: 2158–88.
- De Leeuw, Jan. 1974a. "Approximation of a Real Symmetric Matrix by a Positive Semidefinite Matrix of Rank r ." Technical Memorandum TM-74-1229-10. Murray Hill, N.J.: Bell Telephone Laboratories. <https://jansweb.netlify.app/publication/deleeuw-r-74-c/deleeuw-r-74-c.pdf>.
- . 1974b. "Finding a Positive Semidefinite Matrix of Prescribed Rank r in a Nonlinear Differentiable Manifold." Technical Memorandum. Murray Hill, N.J.: Bell Telephone Laboratories. <https://jansweb.netlify.app/publication/deleeuw-r-74-b/deleeuw-r-74-b.pdf>.
- . 1975. "A Normalized Cone Regression Approach to Alternating Least Squares Algorithms." Department of Data Theory FSW/RUL. <https://jansweb.netlify.app/publication/deleeuw-u-75-a/deleeuw-u-75-a.pdf>.
- De Leeuw, Jan, and Willem J. Heiser. 1982. "Theory of Multidimensional Scaling." In *Handbook of Statistics, Volume II*, edited by P. R. Krishnaiah and L. Kanal. Amsterdam, The Netherlands: North Holland Publishing Company.
- De Leeuw, Jan, Kurt Hornik, and Patrick Mair. 2009. "Isotone Optimization in R: Pool-Adjacent-Violators Algorithm (PAVA) and Active Set Methods." *Journal of Statistical Software* 32 (5): 1–24.
- De Leeuw, Jan, and Patrick Mair. 2009. "Multidimensional Scaling Using Majorization: SMACOF in R." *Journal of Statistical Software* 31 (3): 1–30. <https://www.jstatsoft.org/article/view/v031i03>.
- Gower, John C. 1966. "Some Distance Properties of Latent Root and Vector Methods Used in Multivariate Analysis." *Biometrika* 53: 325–38.
- Groenen, Patrick J. F., Jan De Leeuw, and Rudolf Mathar. 1995. "Least Squares Multidimensional Scaling with Transformed Distances." In *From Data to Knowledge: Theoretical and Practical Aspects of Classification, Data Analysis and Knowledge Organization*, edited by W. Gaul and D. Pfeifer. Berlin, Germany: Springer Verlag. <https://jansweb.netlify.app/publication/groenen-deleeuw-mathar-c-95/groenen-deleeuw-mathar-c-95.pdf>.
- Groenen, Patrick J. F., Patrizia Giaquinto, and Henk A. L. Kiers. 2003. "Weighted Majorization Algorithms for Weighted Least Squares Decomposition Models." Econometric Institute Report EI 2003-09. Econometric Institute, Erasmus University Rotterdam. <https://repub.eur.nl/pub/1700>.
- Jeffrey, David J. 1997. "Formulae, Algorithms, and Quartic Extrema." *Mathematics Magazine* 70 (5): 341–48.
- Keller, Joseph B. 1962. "Factorization of Matrices by Least Squares." *Biometrika* 49: 239–42.
- Kiers, Henk A. L. 1997. "Weighted Least Squares Fitting Using Iterative Ordinary Least Squares Algorithms." *Psychometrika* 62: 251–66.
- Kruskal, Joseph B. 1964. "Multidimensional Scaling by Optimizing Goodness of Fit to a

- Nonmetric Hypothesis.” *Psychometrika* 29: 1–27.
- Mardia, Kanti V. 1978. “Some Properties of Classical Multidimensional Scaling.” *Communications in Statistics - Theory and Methods* 7 (13): 1233–41.
- Mullen, Katherine M., and Ivo H. M. van Stokkum. 2024. *npls: The Lawson-Hanson algorithm for non-negative least squares (NNLS)*. %7B<https://CRAN.R-project.org/package=npls>%7D.
- Qiu, Yixuan, and Jiali Mei. 2024. *RSpectra: Solvers for Large-Scale Eigenvalue and SVD Problems*. <https://CRAN.R-project.org/package=RSpectra>.
- R Core Team. 2025. *R: A Language and Environment for Statistical Computing*. Vienna, Austria: R Foundation for Statistical Computing. <https://www.R-project.org/>.
- Ramsay, James O. 1977. “Maximum Likelihood Estimation in Multidimensional Scaling.” *Psychometrika* 42: 241–66.
- Rothkopf, Ernst Z. 1957. “A Measure of Stimulus Similarity and Errors in some Paired-associate Learning.” *Journal of Experimental Psychology* 53: 94–101.
- Saito, Takayuki. 1978. “The Problem of the Additive Constant and Eigenvalues in Metric Multidimensional Scaling.” *Psychometrika* 43 (2): 193–201.
- Schoenberg, Isaac J. 1935. “Remarks to Maurice Frechet’s article: Sur la Definition Axiomatique d’une Classe d’Espaces Vectoriels Distances Applicables Vectoriellement sur l’Espace de Hilbert.” *Annals of Mathematics* 36: 724–32.
- Takane, Yoshio, Forrest W. Young, and Jan De Leeuw. 1977. “Nonmetric Individual Differences in Multidimensional Scaling: An Alternating Least Squares Method with Optimal Scaling Features.” *Psychometrika* 42: 7–67.
- Torgerson, Warren S. 1958. *Theory and Methods of Scaling*. New York: Wiley.
- Trosset, Michael W. 1997. “Numerical Algorithms for Multidimensional Scaling.” In *Classification and Knowledge Organization*, edited by Rüdiger Klar and Otto Opitz, 80–92. Springer.
- . 1998. “A New Formulation of the Nonmetric Strain Problem in Multidimensional Scaling.” *Journal of Classification* 15: 15–35.
- . 2000. “Distance Matrix Completion by Numerical Optimization.” *Computational Optimization and Applications* 17: 11–22.
- . 2002. “Extensions of Classical Multidimensional Scaling via Variable Reduction.” *Computational Statistics* 17: 147–63.
- Young, Gale, and Alston S. Householder. 1938. “Discussion of a Set of Points in Terms of Their Mutual Distances.” *Psychometrika* 3 (19-22).