

Getting Things in Order: An introduction to the R package **seriation**

Michael Hahsler, Kurt Hornik and Christian Buchta

August 30, 2007

Abstract

Seriation, i.e., finding a linear order for a set of objects given data and a loss or merit function, is a basic problem in data analysis. Caused by the problem's combinatorial nature, it is hard to solve for all but very small sets. Nevertheless, both exact solution methods and heuristics are available. In this paper we present the package **seriation** which provides the infrastructure for seriation with R. The infrastructure comprises data structures to represent linear orders as permutation vectors, a wide array of seriation methods using a consistent interface, a method to calculate the value of various loss and merit functions, and several visualization techniques which build on seriation. To illustrate how easily the package can be applied for a variety of applications, a comprehensive collection of examples is presented.

1 Introduction

A basic problem in data analysis, called *seriation* or sometimes *sequencing*, is to arrange all objects in a set in a linear order given available data and some loss or merit function in order to reveal structural information. Together with cluster analysis and variable selection, seriation is an important problem in the field of *combinatorial data analysis* (Arabie and Hubert, 1996). Solving problems in combinatorial data analysis requires the solution of discrete optimization problems which, in the most general case, involves evaluating all feasible solutions. Due to the combinatorial nature, the number of possible solutions grows with problem size (number of objects) by the order $O(n!)$. This makes a brute-force enumerative approach infeasible for all but very small problems. To solve larger problems (with up to 40 objects), partial enumeration methods can be used. For example, Hubert, Arabie, and Meulman (1987) propose dynamic programming and Brusco and Stahl (2005) use a branch-and-bound strategy. For even larger problems only heuristics can be employed.

It has to be noted that seriation has a rich history in archeology. Petrie (1899) was the first to use seriation as a formal method. He applied it to find a chronological order for graves discovered in the Nile area given objects found there. He used a cross-tabulation of grave sites and objects and rearranged the table using row and column permutations till all large values were close to the diagonal. In the rearranged table graves with similar objects are closer to each other. Together with the assumption that different objects continuously come into and go out of fashion, the order of graves in the rearranged table suggests a chronological order. Initially, the rearrangement of rows and columns of this contingency table was done manually and the adequacy was only judged subjectively by the researcher. Later, Robinson (1951), Kendall (1971) and others proposed measures of agreement between rows to quantify optimality of the resulting table. A comprehensive description of the development of seriation in archeology is presented by Ihm (2005).

Techniques related to seriation are also popular in several other fields. Especially in ecology scaling techniques are used under the name *ordination*. For these applications several R packages already exist (e.g., **ade4** (Chessel, Dufour, and Dray, 2007) and **vegan** (Oksanen, Kindt, Legendre, and O'Hara, 2007)). This paper describes the new package **seriation** which differs from existing packages in the following ways:

1. **seriation** provides a flexible infrastructure for seriation;

2. **seriation** focuses on seriation as a combinatorial optimization problem.

This paper starts with a formal introduction of the seriation problem as a combinatorial optimization problem in Section 2. In Section 3 we give an overview of seriation methods. In Section 4 we present the infrastructure provided by the package **seriation**. Several examples and applications for seriation are given in Section 5. We conclude with Section 6.

2 Seriation as a combinatorial optimization problem

To seriate a set of n objects $\{O_1, \dots, O_n\}$ one typically starts with an $n \times n$ symmetric dissimilarity matrix $\mathbf{D} = (d_{ij})$ where d_{ij} for $1 \leq i, j \leq n$ represents the dissimilarity between the objects O_i and O_j and all $d_{ii} = 0$. We define a permutation function Ψ as a function which reorders the objects in \mathbf{D} by simultaneously permuting rows and columns. The seriation problem is to find a permutation function Ψ which optimizes the value of a given loss function L or merit function M . This results in the optimization problems

$$L(\Psi(\mathbf{D})) \rightarrow \min_{\Psi} \quad \text{or} \quad M(\Psi(\mathbf{D})) \rightarrow \max_{\Psi}, \quad (1)$$

respectively.

A symmetric dissimilarity matrix is known as *two-way one-mode* data since it has columns and rows (two-way) but only represents one set of objects (one-mode). Seriation is also possible for two-way two-mode data which are represented by a general positive matrix. In such data columns and rows represent two sets of objects which are reordered simultaneously. For loss functions for two-way two-mode data the optimal order of columns can be dependent of the order of rows and vice versa or it can be independent allowing for breaking the optimization down into two separate problems, one for the columns and one for the rows. Another way to deal with the seriation for two-way two-mode data is to calculate two dissimilarity matrices, one for each mode, and then solve two seriation problem on two-way one-mode data. Furthermore, seriation can be generalized to k -way k -mode data in the form of a k -dimensional array by defining suitable loss functions for such data or by breaking the problem down into several lower dimensional independent problems which can be solved.

In the following subsections, we review some commonly used loss functions. Most functions are used for two-way one-mode data but the measure of effectiveness and stress can be also used for two-way two-mode data.

2.1 Column/row gradient measures

A symmetric dissimilarity matrix where the values in all rows and columns only increase when moving away from the main diagonal is called a perfect *anti-Robinson matrix* after the statistician Robinson (1951). Formally, an $n \times n$ dissimilarity matrix \mathbf{D} is in anti-Robinson form if, and only if the following two gradient conditions hold (Hubert et al., 1987):

$$\text{within rows: } d_{ik} \leq d_{ij} \quad \text{for } 1 \leq i < k < j \leq n; \quad (2)$$

$$\text{within columns: } d_{kj} \leq d_{ij} \quad \text{for } 1 \leq i < k < j \leq n. \quad (3)$$

In an anti-Robinson matrix the smallest dissimilarity values appear close to the main diagonal, therefore, the closer objects are together in the order of the matrix, the higher their similarity. This provides a natural objective for seriation.

It has to be noted that \mathbf{D} can be brought into a perfect anti-Robinson form by row and column permutation whenever \mathbf{D} is an ultrametric or \mathbf{D} has an exact Euclidean representation in a single dimension (Hubert et al., 1987). However, for most data only an approximation to the anti-Robinson form is possible.

A suitable merit measure which quantifies the divergence of a matrix from the anti-Robinson form was given by Hubert et al. (1987) as

$$M(\mathbf{D}) = \sum_{i < k < j} f(d_{ik}, d_{ij}) + \sum_{i < k < j} f(d_{kj}, d_{ij}) \quad (4)$$

where $f(\cdot, \cdot)$ is a function which defines how a violation or satisfaction of a gradient condition for an object triple $(O_i, O_k \text{ and } O_j)$ is counted. Hubert et al. (1987) suggest two functions.

The first function is given by:

$$f(z, y) = \text{sign}(y - z) = \begin{cases} +1 & \text{if } z < y; \\ 0 & \text{if } z = y; \\ -1 & \text{if } z > y. \end{cases} \quad (5)$$

It results in the raw number of triples satisfying the gradient constraints minus triples which violate the constraints.

The second function is defined as:

$$f(z, y) = |y - z| \text{sign}(y - z) = y - z \quad (6)$$

It weighs each satisfaction or violation by its magnitude given by the absolute difference between the values.

2.2 Anti-Robinson events

An even simpler loss function can be created in the same way as the gradient measures above by concentrating on violations only.

$$L(\mathbf{D}) = \sum_{i < k < j} f(d_{ik}, d_{ij}) + \sum_{i < k < j} f(d_{kj}, d_{ij}) \quad (7)$$

To only count the violations we use

$$f(z, y) = I(z, y) = \begin{cases} 1 & \text{if } z < y \text{ and} \\ 0 & \text{otherwise.} \end{cases} \quad (8)$$

$I(\cdot)$ is an indicator function returning 1 only for violations. Chen (2002) presented a formulation for an equivalent loss function and called the violations *anti-Robinson events*. Chen (2002) also introduced a weighted versions of the loss function resulting in

$$f(z, y) = |y - z| I(z, y) \quad (9)$$

using the absolute deviations as weights.

2.3 Hamiltonian path length

The dissimilarity matrix \mathbf{D} can be represented as a finite weighted graph $G = (\Omega, E)$ where the set of objects Ω constitute the vertices and each edge $e_{ij} \in E$ between the objects $O_i, O_j \in \Omega$ has a weight w_{ij} associated which represents the dissimilarity d_{ij} .

Such a graph can be used for seriation (see, e.g., Hubert, 1974; Caraux and Pinloche, 2005). An order Ψ of the objects can be seen as a path through the graph where each node is visited exactly once, i.e., a Hamiltonian path. Minimizing the Hamiltonian path length results in a seriation optimal with respect to dissimilarities between neighboring objects. The loss function based on the Hamiltonian path length is:

$$L(\mathbf{D}) = \sum_{i=1}^{n-1} d_{i, i+1}. \quad (10)$$

Note that the length of the Hamiltonian path is equal to the value of the *minimal span loss function* (as used by Chen, 2002), and both notions are related to the *traveling salesperson problem* (Gutin and Punnen, 2002).

2.4 Inertia criterion

Another way to look at the seriation problem is not to focus on placing small dissimilarity values close to the diagonal, but to push large values away from it. A function to quantify this

is the moment of inertia of dissimilarity values around the diagonal (Caraux and Pinloche, 2005) defined as

$$M(\mathbf{D}) = \sum_{i=1}^n \sum_{j=1}^n d_{ij} |i - j|^2. \quad (11)$$

$|i - j|^2$ is used as a measure for the distance to the diagonal and d_{ij} gives the weight. This is a merit function since the sum increases when higher dissimilarity values are placed farther away from the diagonal.

2.5 Least squares criterion

Another natural loss function for seriation is to quantify the deviations between the dissimilarities in \mathbf{D} and the rank differences of the objects. Such deviations can be measured, e.g., by the sum of squares of deviations (Caraux and Pinloche, 2005) defined by

$$L(\mathbf{D}) = \sum_{i=1}^n \sum_{j=1}^n (d_{ij} - |i - j|)^2, \quad (12)$$

where $|i - j|$ is the rank difference or gap between O_i and O_j .

The least squares criterion defined here is related to uni-dimensional scaling (de Leeuw, 2005), where the objective is to place all objects on a straight line such that the dissimilarities in \mathbf{D} are preserved by the relative positions in the best possible way. The optimization problem of uni-dimensional scaling is to minimize $\sum_{i=1}^n \sum_{j=1}^n (d_{ij} - |x_i - x_j|)^2$ which is close to the seriation problem, but in addition to the ranking of the objects also takes the distances between objects on the resulting scale into account.

Note that if Euclidean distance is used to calculate \mathbf{D} from a data matrix \mathbf{X} , the order of the elements in \mathbf{X} by projecting them on the first principal component of \mathbf{X} minimizes the loss function of uni-dimensional scaling (using squared distances) and also provides good solutions for this seriation criterion.

2.6 Measure of effectiveness

McCormick, Schweitzer, and White (1972) defined the *measure of effectiveness* (*ME*) for an $n \times m$ matrix $\mathbf{X} = (x_{ij})$ as

$$M(\mathbf{X}) = \frac{1}{2} \sum_{i=1}^n \sum_{j=1}^m x_{ij} [x_{i,j+1} + x_{i,j-1} + x_{i+1,j} + x_{i-1,j}] \quad (13)$$

with, by convention $x_{0,j} = x_{n+1,j} = x_{i,0} = x_{i,m+1} = 0$. ME is maximized if each element is as closely related numerically to its four neighboring elements as possible.

ME was developed for two-way two-mode data, however, ME can also be used for a symmetric matrix (one-mode data) and gets maximal only if all large values are grouped together around the main diagonal.

Note that the definition in equation (13) can be rewritten as

$$M(\mathbf{X}) = \frac{1}{2} \sum_{i=1}^n \sum_{j=1}^m x_{ij} [x_{i,j+1} + x_{i,j-1}] + \sum_{i=1}^n \sum_{j=1}^m x_{ij} [x_{i+1,j} + x_{i-1,j}] \quad (14)$$

showing that the contribution of column and row order to the loss function is additive and thus the permutation of columns and rows are independent.

2.7 Stress

Stress measures the conciseness of the presentation of a matrix (two-mode data) and can be seen as a purity function which compares the values in a matrix with their neighbors. The stress measures used here are computed as the sum of squared distances of each matrix entry from its adjacent entries. Niermann (2005) defined for an $n \times m$ matrix $\mathbf{X} = (x_{ij})$ two types of neighborhoods:

- The Moore neighborhood comprises the (at most) eight adjacent entries. The local stress measure for element x_{ij} is defined as

$$\sigma_{ij} = \sum_{k=\max(1,i-1)}^{\min(n,i+1)} \sum_{l=\max(1,j-1)}^{\min(m,j+1)} (x_{ij} - x_{kl})^2 \quad (15)$$

- The Neumann neighborhood comprises the (at most) four adjacent entries resulting in the local stress of x_{ij} of

$$\sigma_{ij} = \sum_{k=\max(1,i-1)}^{\min(n,i+1)} (x_{ij} - x_{kj})^2 + \sum_{l=\max(1,j-1)}^{\min(m,j+1)} (x_{ij} - x_{il})^2 \quad (16)$$

Both local stress measures can be used to construct a global measure for the whole matrix by summing over all entries which can be used as a loss function:

$$L(\mathbf{X}) = \sum_{i=1}^n \sum_{j=1}^m \sigma_{ij} \quad (17)$$

The major difference between the Moore and the Neumann neighborhood is that for the later the contribution of row and column permutations to stress are independent and thus can be optimized independently.

Stress can be also used as a loss function for symmetric proximity matrices (one-mode data). Note also, that stress with Neumann neighborhood is related to the measure of effectiveness defined above (in Section 2.6) since both measures are optimal if for each cell the cell and its four neighbors are numerically as similar as possible.

3 Seriation methods

Solving the discrete optimization problem for seriation with most loss/merit functions is clearly very hard. The number of possible permutations is $n!$ which makes an exhaustive search for sets with a medium to large number of objects infeasible. In this section, we describe some methods (partial enumeration, heuristics and other methods) which are typically used for seriation. For each method we state for which type of loss/merit functions it is suitable and if it finds the optimum or is a heuristic.

3.1 Partial enumeration methods

Partial enumeration methods search for the exact solution of a combinatorial optimization problem. Exploiting properties of the search space, only a subset of the enormous number of possible combinations has to be evaluated. Popular partial enumeration methods which are used for seriation are *dynamic programming* (Hubert et al., 1987) and *branch-and-bound* (Brusco and Stahl, 2005).

Dynamic programming recursively searches for the optimal solution checking and storing $2^n - 1$ results. Although $2^n - 1$ grows at a lower rate than $n!$ and is for $n \gg 3$ considerably smaller, the storage requirements of $2^n - 1$ results still grow fast, limiting the maximal problem size severely. For example, for $n = 30$ more than one billion results have to be calculated and stored, clearly a number too large for the main memory capacity of most current computers.

Branch-and-bound has only very moderate storage requirements. The forward-branching procedure (Brusco and Stahl, 2005) starts to build partial permutations from left (first position) to right. At each step, it is checked if the permutation is valid and several fathoming tests are performed to check if the algorithm should continue with the partial permutation. The most important fathoming test is the boundary test, which checks if the partial permutation can possibly lead to a complete permutation with a better solution than the currently best one. In this way large parts of the search space can be omitted. However, in contrast to the dynamic programming approach, the reduction of search space is strongly data dependent and poorly structured data can lead to very poor performance. With branch-and-bound slightly larger problems can be solved than with dynamic programming in reasonable time.

Brusco and Stahl (2005) state that depending on the data, in some cases proximity matrices with 40 or more objects can be handled with current hardware.

Partial enumeration methods can be used to find the exact solution independently of the loss/merit function. However, partial enumeration is limited to only relatively small problems.

3.2 Traveling salesperson problem solver

Seriation by minimizing the length of a Hamiltonian path through a graph is equal to solving a traveling salesperson problem. The traveling salesperson or salesman problem (TSP) is a well known and well researched combinatorial optimization problem. The goal is to find the shortest tour that visits each city in a given list exactly once and then returns to the starting city. In graph theory a TSP tour is called a *Hamiltonian cycle*. But for the seriation problem, we are looking for a Hamiltonian path. Garfinkel (1985) described a simple transformation of the TSP to find the shortest Hamiltonian path. An additional row and column of 0's is added (sometimes this is referred to as a *dummy city*) to the original $n \times n$ dissimilarity matrix \mathbf{D} . The solution of this $(n + 1)$ -city TSP, gives the shortest path where the city representing the added row/column cuts the cycle into a linear path.

As the general seriation problem, solving the TSP is difficult. In the seriation case with $n + 1$ cities, $n!$ tours have to be checked. However, despite this vast searching space, small instances can be solved efficiently using dynamic programming (Held and Karp, 1962) and larger instances of several hundred objects can be solved using *branch-and-cut* algorithms (Padberg and Rinaldi, 1990). For even larger instances or if running time is critical, a wide array of heuristics are available, ranging from simple nearest neighbor approaches to construct a tour (Rosenkrantz, Stearns, and Philip M. Lewis, 1977) to complex heuristics like the Lin-Kernighan heuristic (Lin and Kernighan, 1973). A comprehensive overview of heuristics and exact methods can be found in Gutin and Punnen (2002).

3.3 Bond energy algorithm

The *bond energy algorithm* (BEA; McCormick et al., 1972) is a simple heuristic to rearrange columns and rows of a matrix such that each entry is as closely numerically related to its four neighbors as possible. To achieve this, BEA tries to maximize the measure of effectiveness (ME) defined in Section 2.6. For optimizing the ME, columns and rows can be treated separately since changing the order of rows does not influence the ME contributions of the columns and vice versa. BEA consists of the following three steps:

1. Place one randomly chosen column.
2. Try to place each remaining column at each possible position left, right and between the already placed columns and calculate every time the increase in ME. Choose the column and position which gives the largest increase in ME and place the column. Repeat till all columns are placed.
3. Repeat procedure with rows.

This greedy algorithm works fast and is only dependent on the choice of the first column/row. This dependence can be reduced by repeating the procedure several times with different choices and returning the solution with the highest ME.

Although McCormick et al. (1972) use BEA also for non-binary data, Arabie and Hubert (1990) argue that the measure of effectiveness only serves its intended purpose of finding an arrangement which is close to Robinson form for binary data and should therefore only be used for binary data.

Lenstra (1974) notes that the optimization problem of BEA can be stated as two independent traveling salesperson problems (TSPs). For example, the row TSP for an $n \times m$ matrix \mathbf{X} consists of n cities with an $n \times n$ distance matrix \mathbf{D} where the distances are

$$d_{ij} = - \sum_{k=1}^m x_{ik} x_{jk}.$$

BEA is in fact a simple suboptimal TSP heuristic using these distances and instead of BEA any TSP solver can be used to obtain an order. With an exact TSP solver, the optimal solution can be found.

3.4 Hierarchical clustering

Hierarchical clustering produces a series of nested clusterings which can be visualized by a dendrogram, a tree where each internal node represents a split into subtrees and has a measure of similarity/dissimilarity attached to it. As a simple heuristic to find a linear order of objects, the order of the leaf nodes in a dendrogram structure can be used. This idea is used, e.g., by heat maps to reorder rows and columns with the aim to place more similar objects and variables closer together.

The order of leaf nodes in a dendrogram is not unique. A binary (two-way splits only) dendrogram for n objects has 2^{n-1} internal nodes and at each internal node the left and right subtree (or leaves) can be swapped resulting in 2^{n-1} distinct leaf orderings. To find an unique or optimal order, an additional criterion has to be defined. Gruvaeus and Wainer (1972) suggest to obtain a unique order by requiring to order the leaf nodes such that at each level the objects at the edge of each cluster are adjacent to that object outside the cluster to which it is nearest.

Bar-Joseph, Demaine, Gifford, and Jaakkola (2001) suggest to rearrange the dendrogram such that the Hamiltonian path connecting the leaves is minimized and called this the optimal leaf order. The authors also present a fast algorithm with time complexity $O(n^4)$ to solve this optimization problem. Note that this problem is related to the TSP described above, however, the given dendrogram structure significantly reduces the number of permissible permutations making the problem easier.

Although hierarchical clustering solves an optimization problem different to the seriation problem discussed in this paper, hierarchical clustering still can produce useful orderings, e.g., for visualization.

3.5 Rank-two ellipse seriation

Chen (2002) proposes to generate a sequence of correlation matrices R^1, R^2, \dots . R^1 is the correlation matrix of the original distance matrix \mathbf{D} and

$$R^{n+1} = \phi R^n, \quad (18)$$

where $\phi(\cdot)$ calculates a correlation matrix.

The rank of the matrix R^n falls with increasing n . The sequence is continued till the first matrix in the sequence has a rank of 2. Projecting all points in this matrix on the first two eigenvectors, all points fall on an ellipse. The order of the points on this ellipse is the resulting order where the ellipse can be cut at any of the two interception points (top or bottom) with the vertical axis.

Although the rank-two ellipse seriation procedure does not try to solve a combinatorial optimization problem, it still provides for some cases a useful ordering.

4 The package infrastructure

The **seriation** package provides the data structures and some algorithms to efficiently handle seriation with R. As the input data for seriation R already provides

- for two-way one-mode data the class **dist**,
- for two-way two-mode data the class **matrix**, and
- for k -way k -mode data the class **array**.

However, R provides no classes for representing permutation vectors. **seriation** adds the necessary data structure (using the S3 class system) as depicted in Figure 1. The class **ser_permutation** represents the permutation information for k -mode data (including the cases of $k = 1$ and $k = 2$). It consists of k permutation vectors, where each permutation vector is represented by an object of a concrete implementation of the abstract class **ser_permutation_vector**. This design is used to allow to use arbitrary representations for the permutation vectors without removing additional available information by just using an integer vector. Currently, the permutation vector can be stored as a simple integer vector or as an object of class **hclust** (defined in package **stats**). **hclust** describes a hierarchical




Figure 1: UML class diagram of the data structures for permutations provided by **seriation**

Algorithm	method argument	Optimizes	Input data
Simulated annealing	"ARSA"	Gradient measure	dist
Branch-and-bound	"BBURCG"	Gradient measure	dist
Branch-and-bound	"BBWRCG"	Gradient measure (weighted)	dist
TSP solver	"TSP"	Hamiltonian path length	dist
Rank-two ellipse seriation	"Chen"	Other	dist
MDS – first dimension	"MDS"	Other	dist
Hierarchical clustering	"HC"	Other	dist
Gruvaeus and Wainer	"GW"	Other	dist
Optimal leaf ordering	"OLO"	Hamiltonian path length (restricted)	dist
Bond Energy Algorithm	"BEA"	Measure of effectiveness	matrix
TSP to optimize ME	"BEA_TSP"	Measure of effectiveness	matrix
First principal component	"PCA"	Other	matrix

Table 1: Currently implemented methods for **seriation()**.

clustering tree (dendrogram) including an ordering for the trees node leaves which provides a permutation for all objects (see Section 3.4).

Class **ser_permutation_vector** has a constructor **ser_permutation_vector()** which converts data into the correct concrete subclass of **ser_permutation_vector** and checks if it contains a proper permutation vector. For **ser_permutation_vector** the methods **print()**, **length()** for the length of the permutation vector, **get_method()** to get the method used to generate the permutation, and **get_order()** to access the raw (integer) permutation vector are available. An accessor method **get_order()** is implemented for each concrete subclass of **ser_permutation_vector**. With this design, other classes representing permutations can be easily incorporated by just adding an appropriate **get_order()** methods for the new class.

For **ser_permutation** a constructor is provided which can bind k **ser_permutation_vector** objects together into an object for k -mode data. **ser_permutation** is implemented as a list of length k and each element contains a **ser_permutation_vector** object. Methods like **length()**, accessing elements with **[[**, **[[<-**, subsetting with **[**, and combining with **c()** work as expected. Also a **print()** method are provided. Finally, direct access to the raw permutation vectors is available using **get_order()**. Here the second argument (which defaults to 1) specifies the dimension (mode) for which the order vector is requested.

All seriation algorithms are available via the function **seriate()** defined as:

```
seriate(x, method = NULL, control = NULL, ...)
```

where **x** is the input data, **method** is a string defining the seriation method to be used and **control** can contain a list with additional information for the algorithm. **seriate()** returns an object of class **ser_permutation** with a length conforming to the number of dimensions of **x**. For **matrix** the additional argument **margin** can be used if only column (**margin = 1**) or row seriation (**margin = 2**) is needed.

Various seriation methods were already introduced in this paper in Section 3. In Table 1 we summarize the methods currently available in the package for seriation. The code for

Name	method argument	merit/loss	Input data
Anti-Robinson events	"AR_events"	loss	dist
Anti-Robinson deviations	"AR_deviations"	loss	dist
Gradient measure	"Gradient_raw"	merit	dist
Gradient measure (weighted)	"Gradient_weighted"	merit	dist
Hamiltonian path length	"Path_length"	loss	dist
Inertia criterion	"Inertia"	merit	dist
Least squares criterion	"Least_squares"	loss	dist
Measure of effectiveness	"ME"	merit	matrix
Stress (Moore neighborhood)	"Moore_stress"	loss	matrix
Stress (Neumann neighborhood)	"Neumann_stress"	loss	matrix

Table 2: Implemented loss/merit functions in function `criterion()`.

the simulated annealing heuristic (Brusco, Köhn, and Stahl, 2007) and the two branch-and-bound (Brusco and Stahl, 2005) was obtained from the authors. The TSP solvers (exact solvers and a variety of heuristics) is provided by package **TSP** (Hahsler and Hornik, 2007). For the rank-two ellipse seriation we implemented the algorithm by Chen (2002). For the Gruvaeus and Wainer algorithm, the implementation in package **gclus** (Hurley, 2007) is used. For optimal leaf ordering we implemented the algorithm by Bar-Joseph et al. (2001). The BEA code was kindly provided by Fionn Murtagh. Over time more methods will be added to the package.

To calculate the value of a loss/merit function for data and a certain permutation, the function

```
criterion(x, order = NULL, method = "all", ...)
```

is provided. `x` is the data object, `order` contains a suitable object of class `ser_permutation` (if omitted no permutation is performed) and `method` specifies the type of loss/merit function. A vector of several methods can be used resulting in a named vector with the values of the requested functions. We already defined different loss/merit functions for seriation in Section 2. In Table 2 we indicate the loss/merit functions currently available in the package. Additionally, the pseudo method `"all"` (which is the default) can be used to calculate the values for all applicable loss/merit functions.

In addition the package offers the (generic) function

```
permute(x, order)
```

where `x` is the data (a `dist` object, a matrix, an array, a list or a numeric vector) to be reordered and `order` is a `ser_permutation` object of suitable length.

For visualization, the package offers several options:

- Matrix shading with `pimage()`. In contrast to the standard `image()` in package **graphics**, `pimage()` displays the matrix as is with the first element in the top left-hand corner and using a gamma-corrected gray scale.
- Different heat maps (e.g., with optimally reordered dendrograms) with `hmap()`.
- Visualization of data matrices in the spirit of Bertin (1981) with `bertinplot()`.
- *Dissimilarity plot*, a new visualization to judge the quality of a clustering using matrix shading and seriation with `dissplot()`.

We will introduce the package usage and the visualization options in the examples in the next section.

5 Examples and applications

We start this section with a simple first session to demonstrate the basic usage of the package. Then we present and discuss several seriation applications.

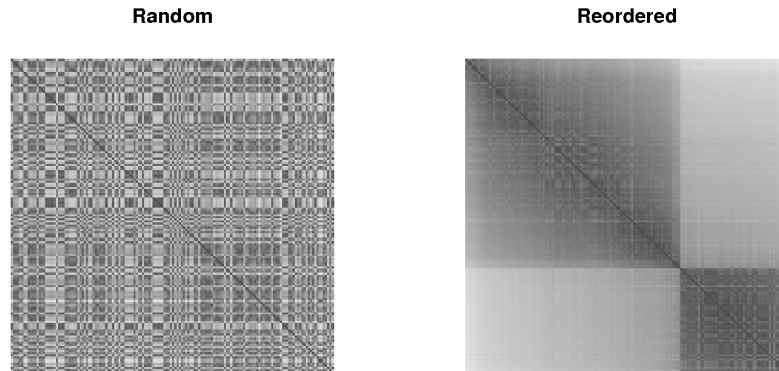


Figure 2: Matrix shading of the distance matrix for the iris data.

5.1 A first session using seriation

In the following example, we use the well known iris data set (from R's **datasets** package) which gives the measurements in centimeters of the variables sepal length and width and petal length and width, respectively, for 50 flowers from each of 3 species of the iris family (Iris setosa, versicolor and virginica).

First, we load the package **seriation** and the iris data set. We remove the species classification and reorder the objects randomly since they are already sorted by species in the data set. Then we calculate the Euclidean distances between objects.

```
> library("seriation")
> data("iris")
> x <- as.matrix(iris[-5])
> x <- x[sample(seq_len(nrow(x))), ]
> d <- dist(x)
```

To seriate the objects given the dissimilarities, we just call **seriate()** with the default settings.

```
> order <- seriate(d)
> order
```

```
object of class 'ser_permutation', 'list'
contains permutation vectors for 1-mode data
```

```
vector length seriation method
1          150          ARSA
```

The result is an object of class **ser_permutation** for one-mode data. The permutation vector length is 150 for the 150 objects in the iris data set and the used seriation method is "ARSA", a simulated annealing heuristic (see Table 1).

To visually inspect the effect of seriation on the distance matrix, we use matrix shading with **pimage()** (the result is shown in Figure 2).

```
> pimage(d, main = "Random")
> pimage(d, order, main = "Reordered")
```

Finally, we can also compare the improvement for different loss functions using **criterion()**.

```
> cbind(random = criterion(d), reordered = criterion(d, order))
```




Figure 3: Matrix shading of the iris data matrix.

	random	reordered
AR_events	551693.0	54915.00
AR_deviations	943320.0	9442.40
Gradient_raw	-1392.0	992073.00
Gradient_weighted	9153.5	1772117.91
Path_length	359.5	86.44
Inertia	215367939.3	356947608.76
Least_squares	78838267.7	76487648.47
ME	301406.5	402260.78
Moore_stress	885654.0	14060.78
Neumann_stress	415557.8	5556.31

Naturally, the reordered dissimilarity matrix achieves better values for all criteria (note that the measure of effectiveness and inertia are merit functions and larger values are better).

Also the original data matrix can be easily inspected using `pimage()`. To use the result of the seriation for the original two-mode data, we have to add a permutation vector to the `ser_permutation` object. To leave the columns in the original order, we add an identity permutation vector to the permutations object using the combine function `c()`.

```
> pimage(x, main = "Random")
> order_2mode <- c(order, ser_permutation(seq_len(ncol(x))))
> order_2mode
> pimage(x, order_2mode, main = "Reordered")
```

```
object of class 'ser_permutation', 'list'
contains permutation vectors for 2-mode data
```

```
vector length seriation method
1          150          ARSA
2           4         unknown
```

5.2 Comparing different seriation methods

To compare different seriation methods we use again the randomized iris data set and the distance matrix `d` from the previous example. We include in the comparison several seriation methods for dissimilarity matrices described in Section 3.

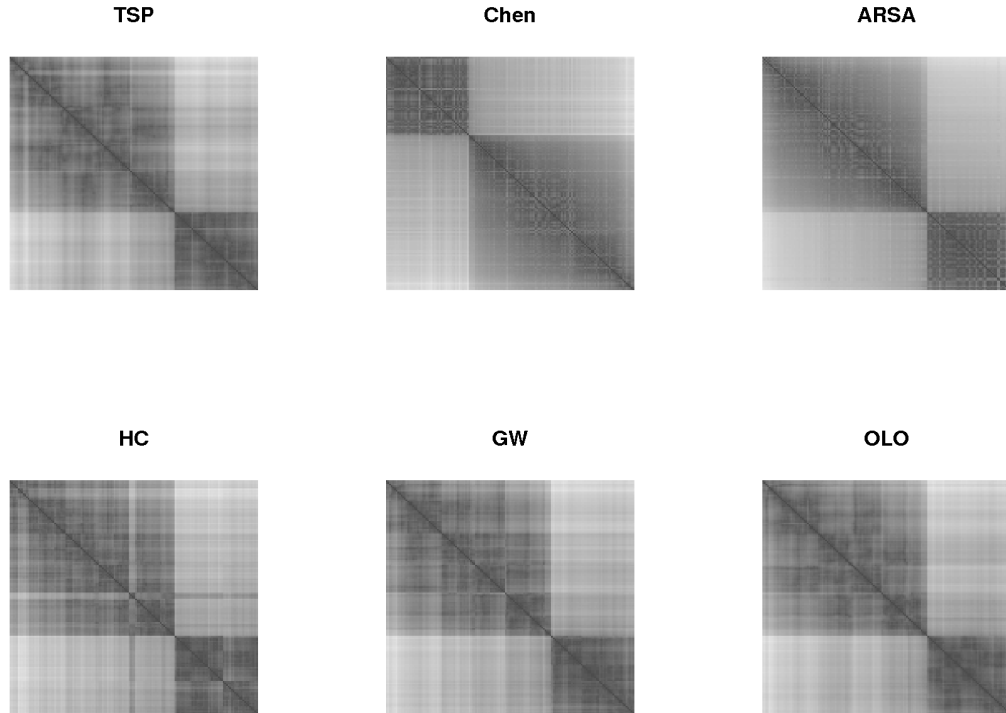


Figure 4: Image plot of the distance matrix for the iris data using rearrangement by different seriation methods.

```
> methods <- c("TSP", "Chen", "ARSA", "HC", "GW", "OLO")
> order <- sapply(methods, FUN = function(m) seriate(d, m),
+   simplify = FALSE)
```

The resulting orderings are displayed using matrix shading (see Figure 4).

```
> tmp <- lapply(order, FUN = function(o) pimage(d, o, main = get_method(o[[1]])))
```

Finally, we compare the values of the loss/merit functions for the different seriation methods.

```
> crit <- sapply(order, FUN = function(o) criterion(d, o))
> t(crit)
```

	AR_events	AR_deviations	Gradient_raw	Gradient_weighted	Path_length
TSP	165453	51174	770981	1651215	52.41
Chen	90015	18184	921885	1746766	85.43
ARSA	55010	9424	991884	1772114	85.44
HC	173729	53167	754425	1644981	64.15
GW	171903	44776	758071	1664667	57.24
OLO	174625	46900	752619	1660309	51.11
	Inertia	Least_squares	ME	Moore_stress	Neumann_stress
TSP	345485410	76648852	402674	15294	5238
Chen	354057418	76521451	402018	19358	7305
ARSA	356944288	76487654	402346	13333	5217
HC	345722302	76657164	402560	30347	10401
GW	346564518	76630917	402822	18636	6426
OLO	346172124	76636726	403053	14979	5126

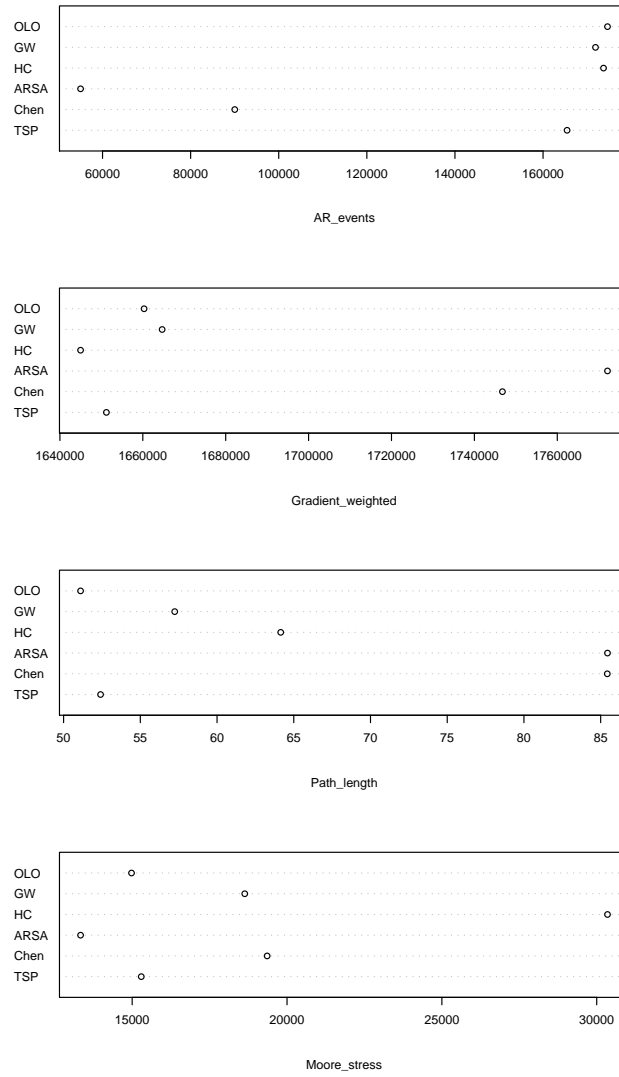


Figure 5: Comparison of different methods and seriation criteria

For easier comparison, Figure 5 contains a plot of the criteria Hamiltonian path length, anti-Robinson events (**AR_events**) and stress using the Moore neighborhood. Clearly, the methods which directly try to minimize the Hamiltonian path length (hierarchical clustering with optimal leaf ordering (OLO) and the TSP heuristic) provide the best results concerning the path length. For the number of anti-Robinson events, using the simulated annealing heuristic (**ARSA**) provides the best result since it directly aims at minimizing this loss function. Regarding stress, the simulated annealing heuristic also provides the best result although, it does not directly minimize this loss function.

5.3 Heat maps

A heat map is a shaded/color coded data matrix with a dendrogram added to one side and to the top to indicate the order of rows and columns. Typically, reordering is done according to row or column means within the restrictions imposed by the dendrogram. Heat maps recently became popular for visualizing large scale genome expression data obtained via DNA microarray technology (see, e.g., Eisen, Spellman, Brownadagger, and Botstein, 1998).

From Section 3.4 we know that it is possible to find the optimal ordering of the leaf nodes of a dendrogram which minimizes the distances between adjacent objects in reasonable time. Such an order might provide an improvement over using simple reordering such as the row or column means with respect to presentation. In **seriation** we provide the function **hmap()** which uses optimal ordering and can also use seriation directly on distance matrices without using hierarchical clustering to produce dendrograms first.

For the following example, we use again the randomly reordered iris data set **x** from the examples above. To make the variables (columns) comparable, we use standard scaling.

```
> x <- scale(x, center = FALSE)
```

To produce a heat map with optimally reordered dendrograms, the function **hmap()** can be used with its default settings. With these settings, the Euclidean distances between rows and between columns are calculated (with **dist()**), hierarchical clustering (**hclust()**) is performed, the resulting dendrograms are optimally reordered, and **heatmap()** in package **stats** is used for plotting.

```
> hmap(x)
> hmap(x, hclustfun = NULL)
```

(see Figure 6).

If **hclustfun = NULL** is used, instead of hierarchical clustering, seriation on the dissimilarity matrices for rows and columns is performed (using a TSP heuristic by default) and the reordered matrix with the reordered dissimilarity matrices to the left and on top is displayed. A **method** argument can be used to choose different seriation methods.

5.4 Bertin's permutation matrix

Bertin (1981, 1999) introduced permutation matrices to analyze multivariate data with medium to low sample size. The idea is to reveal a more homogeneous structure in a data matrix **X** by simultaneously rearranging rows and columns. The rearranged matrix is displayed and cases and variables can be grouped manually to gain a better understanding of the data.

To quantify homogeneity, a purity function

$$\phi = \Phi(\mathbf{X})$$

is defined. Let Π be the set of all permutation functions π for matrix **X**. Note that function π performs row and column permutations on a matrix. The optimal permutation with respect to purity

$$\pi^* = \operatorname{argmax}_{\pi \in \Pi} \Phi(\pi(\mathbf{X}))$$

is found and the rearranged matrix $\pi(\mathbf{X})$ is displayed. Since, depending on the purity function, finding the optimal solution can be hard, often a near optimal solution is also acceptable.

A possible purity function Φ is: Given distances between rows and columns of the data matrix, define purity as the sum of distances of adjacent rows/columns. Using this purity

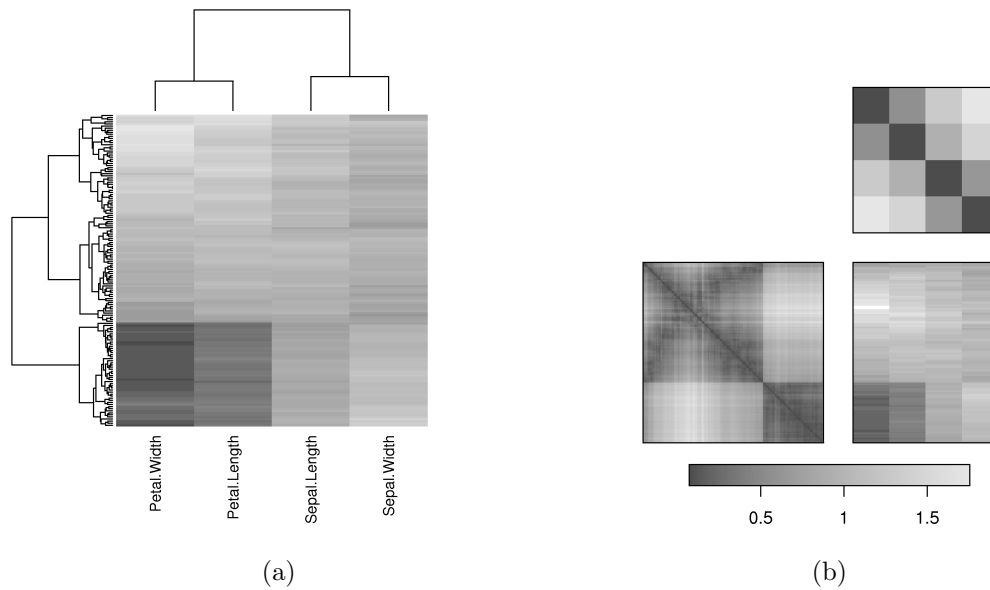


Figure 6: Two presentations of the rearranged iris data matrix. (a) as an optimally reordered heat map and (b) as a seriated data matrix with reordered dissimilarity matrices to the left and on top.

function, finding the optimal permutation π^* means solving two (independent) TSPs, one for the columns and one for the rows.

As an example, we use the results of 8 constitutional referenda for 41 Irish communities (de Falguerolles, Friedrich, and Sawitzki, 1997)¹. To make values comparable across columns (variables), the ranks of the values for each variable are used instead of the original values.

```
> data("Irish")
> orig_matrix <- apply(Irish[, -6], 2, rank)
```

For seriation, we calculate distances between rows and between columns using the sum of absolute rank differences (this is equal to the Minkowski distance with power 1). Then we apply seriation (using a TSP heuristic) to both distance matrices and combine the two resulting `ser_permutation` objects into one object for two-mode data. The original and the reordered matrix are plotted using `bertinplot()`.

```
> order <- c(seriate(dist(orig_matrix, "minkowski", p = 1),
+   method = "TSP"), seriate(dist(t(orig_matrix), "minkowski",
+   p = 1), method = "TSP"))
> order
```


object of class 'ser_permutation', 'list'
contains permutation vectors for 2-mode data

```
vector length seriation method
1          41          TSP
2           8          TSP


> bertinplot(orig_matrix)
> bertinplot(orig_matrix, order)
```

The original matrix and the rearranged matrix are shown in Figure 7 as a matrix of bars where high values are highlighted (filled blocks). Note that following Bertin, the cases (communities) are displayed as the columns and the variables (referenda) as rows. Depending

¹The Irish data set is included in this package. The original data and the text of the referenda can be obtained from <http://www.electionsireland.org/>



(a)



(b)

Figure 7: Bertin plot for the (a) original arrangement and the (b) reordered Irish data set.

on the number of cases and variables, columns and rows can be exchanged to obtain a better visualization.

Although the columns are already ordered (communities in the same city appear consecutively) in the original data matrix in Figure 7(a), it takes some effort to find structure in the data. For example, it seems that the variables ‘Marriage’, ‘Divorce’, ‘Right to Travel’ and ‘Right to Information’ are correlated since the values are all high in the block made up by the columns of the communities in Dublin. The reordered matrix affirms this but makes the structure much more apparent. Especially the contribution of low values (which are not highlighted) to the overall structure becomes only visible after rearrangement.

5.5 Binary data matrices

Binary or 0-1 data matrices are quite common. Often such matrices are called *incidence matrices* since a 1 in a cell indicates the incidence of an event. In archeology such an event could be that a special type of artifact was found at a certain archaeological site. This can be seen as a simplification of a so called *abundance matrix* which codes in each cell the (relative) frequency or quantity of an artifact type at a site. For a comparison of incidence and abundance matrices in archeology we refer the reader to Ihm (2005).

Here we are interested in binary data. For the example we use an artificial data set from Bertin (1981) called *Townships*. The data set contains 9 binary characteristics (e.g., has a veterinary or has a high school) for 16 townships. The idea of the data set is that townships evolve from a rural to an urban environment over time.

After loading the data set (which comes with the package), we use `bertinplot()` to visualize the data (`pimage()` could also be used but `bertinplot()` allows for a nicer visualization). Bars, the standard visualization of `bertinplot()`, do not make much sense for binary data. We therefore use the panel function `panel.squares()` without spacing to plot black squares.

```
> data("Townships")
> bertinplot(Townships, options = list(panel = panel.squares,
+   spacing = 0, frame = TRUE))
```

The original data in Figure 8(a) does not reveal structure in the data. To improve the display, we run the bond energy algorithm (BEA) for columns and rows 10 times with random starting points and report the best solution.

```
> order <- seriate(Townships, method = "BEA", control = list(rep = 10))
> bertinplot(Townships, order, options = list(panel = panel.squares,
+   spacing = 0, frame = TRUE))
```

The reordered matrix is displayed in Figure 8(b). A clear structure is visible. The variables (rows in a Bertin plot) can be split into the three categories describing different evolution states of townships:


1. Rural: No doctor, one-room school and possibly also no water supply
2. Intermediate: Land reallocation, veterinary and agricultural cooperative
3. Urban: Railway station, high school and police station

The townships also clearly fall into these three groups which tentatively can be called villages (first 7), towns (next 5) and cities (final 2). The townships B and C are on the transition to the next higher group.


```
> rbind(original = criterion(Townships), reordered = criterion(Townships,
+   order))
```

	ME	Moore_stress	Neumann_stress
original	19	464	260
reordered	65	212	82

BEA tries to maximize the measure of effectiveness which is much higher in the reordered matrix (in fact, 65 is the maximum for the data set). Also the two types of stress are improved significantly.



(a)



(b)

Figure 8: The townships data set in original order (a) and reordered using BEA (b).

5.6 Dissimilarity plot

Assessing the quality of an obtained cluster solution has been a research topic since the invention of cluster analysis. This is especially important since all popular cluster algorithms produce a clustering even for data without a “cluster” structure.

Matrix shading is an old technique to visualize clusterings by displaying the rearranged matrices (see, e.g., Sneath and Sokal, 1973; Ling, 1973; Gale, Halperin, and Costanzo, 1984). Initially matrix shading was used in connection with hierarchical clustering, where the order of the dendrograms leaf nodes was used to arrange the matrix. However, with some extensions, matrix shading can be also used with any partitional clustering method.

Strehl and Ghosh (2003) suggest a matrix shading visualization called *CLUSION* where the dissimilarity matrix is arranged such that all objects pertaining to a single cluster appear in consecutive order in the matrix. The authors call this *coarse seriation*. The result of a “good” clustering should be a matrix with low dissimilarity values forming blocks around the main diagonal. However, using coarse seriation, the order of the clusters has to be predefined and the objects within each cluster are unordered.

Dissimilarity plots as produce by `dissplot()` aims a reflecting global structure between clusters as well as the micro structure within clusters. To achieve this, we arrange the clusters as well as the objects within each cluster using seriation techniques. To arrange clusters, an inter-cluster dissimilarity matrix is calculated using the average between cluster dissimilarities. `seriate()` is used on this inter-cluster dissimilarity matrix to arrange the clusters relative to each other resulting in on average more similar clusters to appear closer together. Within each cluster, `seriate()` is used again on the sub-matrix of the dissimilarity matrix concerning only the objects in the cluster.

For the example, we use again Euclidean distance between the objects in the iris data set.

```
> data("iris")
> iris <- iris[sample(seq_len(nrow(iris))), ]
> d <- dist(as.matrix(iris[-5]), method = "euclidean")
```

First, we use `dissplot()` without a clustering. We set `method` to `NA` to prevent reordering and display the original matrix. Then we omit the `method` argument which results in using the default seriation technique from `seriate.dist()` which is a TSP heuristic. Since we did not provide a clustering, the whole matrix is reordered in one piece. The result is shown in Figure 9. From Figure 9(b) it seems that there is a clear structure in the data which suggests a two cluster solution.

```
> dissplot(d, method = NA)
> dissplot(d, options = list(main = "Seriation (TSP)"))
```

Next, we create a cluster solution using *k*-means. Although we know that the data set should contain 3 groups representing the three species of iris, we let *k*-means produce a 10 cluster solution to study how such a misspecification can be spotted using `dissplot()`.

```
> l <- kmeans(d, 10)$cluster
```

We create a standard dissimilarity plot by providing the cluster solution as a vector of labels. The function rearranges the matrix and plots the result. Since rearrangement can be a time consuming procedure for large matrices, the rearranged matrix and all information needed for plotting is returned as the result.

```
> res <- dissplot(d, labels = l, options = list(main = "Seriation - standard"))
```

```
> res
```

```
object of class 'cluster_dissimilarity_matrix'
matrix dimensions: x
dissimilarity measure: 'euclidean'
number of clusters k: 10
```

```
cluster description
```

```
position label size avg_dissimilarity avg_silhouette_width
```

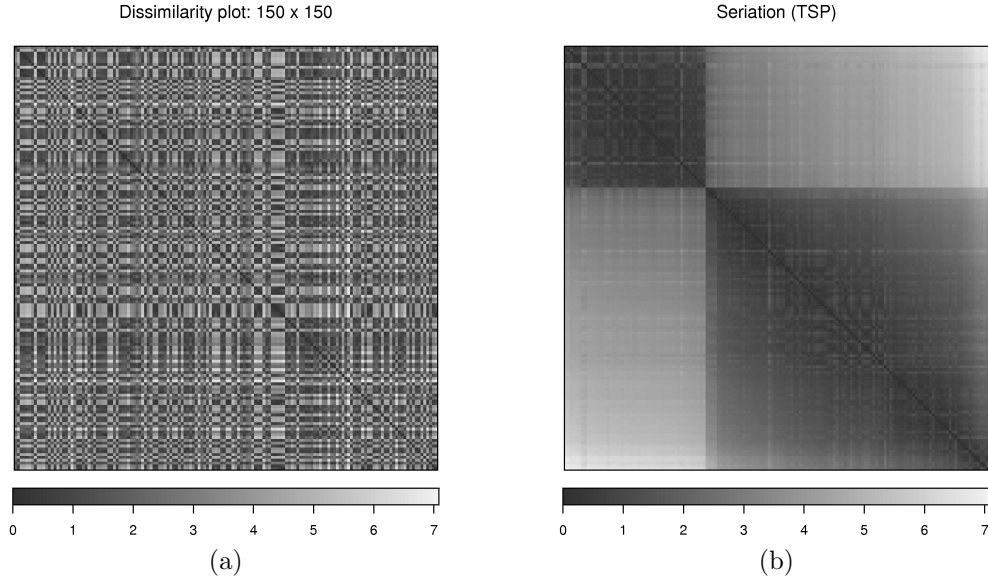


Figure 9: Two dissimilarity plots. (a) the original dissimilarity matrix and (b) the seriated dissimilarity matrix.

1	1	1	25	0.9794	0.34635
2	2	3	46	0.7754	0.31231
3	3	6	29	0.3124	0.43920
4	4	2	3	0.8411	0.10950
5	5	4	5	0.4487	0.14042
6	6	8	7	0.2396	0.18984
7	7	7	7	0.4663	-0.08724
8	8	5	10	0.4939	0.34195
9	9	9	12	0.2942	0.36239
10	10	10	6	0.5795	0.03654

```
used seriation methods
inter-cluster: 'ARSA'
intra-cluster: 'ARSA'
```

The resulting plot is shown in Figure 10(a). The inter-cluster dissimilarities are shown as gray blocks and the average object dissimilarity within each cluster as gray triangles below the main diagonal of the matrix. Since the clusters are arranged such that more similar clusters are closer together, it is easy to see in Figure 10(a) that clusters 1, 3 and 6 as well as clusters 2, 4, 8, 5, 7, 9 and 10 are very similar and form two blocks. This suggests again that a two cluster solution would be reasonable.

Since slight variations of gray values are hard to distinguish, we plot the matrix again (using `plot()` on the result above) and use a threshold on the dissimilarity to suppress high dissimilarity values in the plot.

```
> plot(res, options = list(main = "Seriation - threshold",
+   threshold = 1.5))
```

In the resulting plot in Figure 10(b), we see that the block containing 2, 4, 8, 5, 7, 9 and 10 is very well defined and cleanly separated from the other block. This suggests that these clusters should form together a cluster in a solution with less clusters. The other block is less well defined. There is considerable overlap between clusters 3 and 6, but also cluster 1 and 3 share similar objects.

Using the information stored in the result of `dissplot()` and the class information available for the iris data set, we can analyze the cluster solution and the interpretations of the dissimilarity plot.





Figure 10: Dissimilarity plot for k -means solution with 10 clusters. (a) standard plot and (b) plot with threshold.

```
> table(iris[res$order, 5], res$label)[, res$cluster_order]

      1  3  6  2  4  8  7  5  9 10
setosa  0  0  0  3  5  7  7 10 12  6
versicolor 0 22 28 0 0 0 0  0  0  0
virginica 25 24  1 0 0 0 0  0  0  0
```

As the plot in Figure 10 indicated, the clusters 2, 4, 8, 5, 7, 9 and 10 should be a single cluster containing only flowers of the species *Iris setosa*. The clusters 1, 3 and 6 are more problematic since they contain a mixture of *Iris versicolor* and *virginica*.

6 Conclusion

In this paper we presented the infrastructure provided by the package **seriation**. The infrastructure contains the necessary data structures to store the linear order for one-, two- and k -mode data. It also provides a wide array of seriation methods for different input data, e.g., dissimilarities, binary and general data matrices focusing on combinatorial optimization.

Based on seriation, several visualization techniques are provided by **seriation**. In particular, the optimally reordered heat map, the Bertin plot and the dissimilarity plot present clear improvements over standard plots.

Future work on **seriation** will focus on adding further seriation methods, such as for example methods for block seriation which aim at finding simultaneous partitions of rows and columns in a data matrix (see, e.g., Marcotorchino, 1987).

Acknowledgments

The authors would like to thank Michael Brusco, Hans-Friedrich Köhn and Stephanie Stahl for their seriation code, and Fionn Murtagh for his BEA implementation.

References

- P. Arabie and L. Hubert. An overview of combinatorial data analysis. In P. Arabie, L. Hubert, and G. D. Soete, editors, *Clustering and Classification*, pages 5–63. World Scientific, River Edge, NJ, 1996.
- P. Arabie and L. J. Hubert. The bond energy algorithm revisited. *IEEE Transactions on Systems, Man, and Cybernetics*, 20(1):268–74, 1990.
- Z. Bar-Joseph, E. D. Demaine, D. K. Gifford, and T. Jaakkola. Fast optimal leaf ordering for hierarchical clustering. *Bioinformatics*, 17(1):22–29, 2001.
- J. Bertin. *Graphics and Graphic Information Processing*. Walter de Gruyter, Berlin, 1981. Translated by William J. Berg and Paul Scott.
- J. Bertin. Graphics and graphic information processing. In S. K. Card, J. D. Mackinlay, and B. Shneiderman, editors, *Readings in Information Visualization*, pages 62–65, San Francisco, CA, USA, 1999. Morgan Kaufmann Publishers Inc. ISBN 1-55860-533-9.
- M. Brusco and S. Stahl. *Branch-and-Bound Applications in Combinatorial Data Analysis*. Springer, 2005.
- M. Brusco, H. Köhn, and S. Stahl. Heuristic implementation of dynamic programming for matrix permutation problems in combinatorial data analysis. *Psychometrika*, 2007. conditionally accepted.
- G. Caraux and S. Pinloche. Permutmatrix: A graphical environment to arrange gene expression profiles in optimal linear order. *Bioinformatics*, 21(7):1280–1281, 2005.
- C.-H. Chen. Generalized association plots: Information visualization via iteratively generated correlation matrices. *Statistica Sinica*, 12(1):7–29, 2002.

- D. Chessel, A.-B. Dufour, and S. Dray. *ade4: Analysis of Ecological Data : Exploratory and Euclidean methods in Multivariate data analysis and graphical display*, 2007. R package version 1.4-3.
- A. de Falguerolles, F. Friedrich, and G. Sawitzki. A tribute to J. Bertin’s graphical data analysis. In W. Bandilla and F. Faulbaum, editors, *SoftStat ’97 (Advances in Statistical Software 6)*, pages 11–20. Lucius & Lucius, 1997.
- J. de Leeuw. Unidimensional scaling. In B. Everitt and D. Howelll, editors, *Encyclopedia of Statistics in Behavioral Science*, volume 4. Wiley, 2005.
- M. B. Eisen, P. T. Spellman, P. O. Brownadagger, and D. Botstein. Cluster analysis and display of genome-wide expression patterns. *Proceedings of the National Academy of Science of the United States*, 95(25):14863–14868, December 1998.
- N. Gale, W. Halperin, and C. Costanzo. Unclassed matrix shading and optimal ordering in hierarchical cluster analysis. *Journal of Classification*, 1:75–92, 1984.
- R. Garfinkel. Motivation and modeling. chapter 2, pages 17–36. Wiley, New York, 1985.
- G. Gruvaeus and H. Wainer. Two additions to hierarchical cluster analysis. *British Journal of Mathematical and Statistical Psychology*, 25:200–206, 1972.
- G. Gutin and A. Punnen, editors. *The Traveling Salesman Problem and Its Variations*, volume 12 of *Combinatorial Optimization*. Kluwer, Dordrecht, 2002.
- M. Hahsler and K. Hornik. *TSP: Traveling Salesperson Problem (TSP)*, 2007. R package version 0.2-1.
- M. Held and R. Karp. A dynamic programming approach to sequencing problems. *Journal of SIAM*, 10:196–210, 1962.
- L. Hubert. Some applications of graph theory and related nonmetric techniques to problems of approximate seriation: The case of symmetric proximity measures. *British Journal of Mathematical Statistics and Psychology*, 27:133–153, 1974.
- L. Hubert, P. Arabie, and J. Meulman. *Combinatorial Data Analysis: Optimization by Dynamic Programming*. Society for Industrial Mathematics, 1987.
- C. Hurley. *gclus: Clustering Graphics*, 2007. R package version 1.2.
- P. Ihm. A contribution to the history of seriation in archaeology. In C. Weihs and W. Gaul, editors, *Classification - the Ubiquitous Challenge, Proceedings of the 28th Annual Conference of the Gesellschaft für Klassifikation e.V., University of Dortmund, March 9–11, 2004*, Studies in Classification, Data Analysis, and Knowledge Organization, pages 307–316, 2005.
- D. Kendall. Seriation from abundance matrices. In F. Hodson, D. Kendall, and P. Tautu, editors, *Mathematics in the archaeological and historical sciences*, pages 214–252. 1971.
- J. K. Lenstra. Clustering a data array and the traveling-salesman problem. *Operations Research*, 22(2):413–414, 1974.
- S. Lin and B. Kernighan. An effective heuristic algorithm for the traveling-salesman problem. *Operations Research*, 21(2):498–516, 1973.
- R. L. Ling. A computer generated aid for cluster analysis. *Communications of the ACM*, 16(6):355–361, 1973.
- F. Marcotorchino. Block seriation problems: A unified approach. *Applied Stochastic Models and Data Analysis*, 3:73–91, 1987.
- W. T. McCormick, P. J. Schweitzer, and T. W. White. Problem decomposition and data reorganization by a clustering technique. *Operations Research*, 20(5):993–1009, 1972.

- S. Niermann. Optimizing the ordering of tables with evolutionary computation. *The American Statistician*, 59(1):41–46, 2005.
- J. Oksanen, R. Kindt, P. Legendre, and B. O’Hara. *vegan: Community Ecology Package*, 2007. R package version 1.8-6.
- M. Padberg and G. Rinaldi. Facet identification for the symmetric traveling salesman polytope. *Mathematical Programming*, 47(2):219–257, 1990. ISSN 0025-5610.
- F. W. M. Petrie. Sequences in prehistoric remains. *Journal of the Anthropological Institute*, 29:295–301, 1899.
- W. Robinson. A method for chronologically ordering archaeological deposits. *American Antiquity*, 16:293–301, 1951.
- D. J. Rosenkrantz, R. E. Stearns, and I. Philip M. Lewis. An analysis of several heuristics for the traveling salesman problem. *SIAM Journal on Computing*, 6(3):563–581, 1977.
- P. H. A. Sneath and R. R. Sokal. *Numerical Taxonomy*. Freeman and Company, San Francisco, 1973.
- A. Strehl and J. Ghosh. Relationship-based clustering and visualization for high-dimensional data mining. *INFORMS Journal on Computing*, 15(2):208–230, 2003.