

Package: anticlust (via r-universe)

October 15, 2026

Type Package

Title Subset Partitioning via Anticustering

Version 0.8.18

Author Martin Papenberg [aut, cre]

(<https://orcid.org/0000-0002-9900-4268>), Meik Michalke [ctb]
(centroid based clustering algorithm), Gunnar W. Klau [ths],
Juliane V. Nagel [ctb] (package logo), Martin Breuer [ctb]
(Bicriterion algorithm by Brusco et al.), Marie L. Schaper
[ctb] (Example data set), Max Diekhoff [ctb] (Optimal maximum
dispersion algorithm), Hannah Hengelbrock [ctb] (TPSDP
heuristic by Yang et al.), Dmitry Wintermantel [ctb] (Function
experimental_allocation()), David Buczynski [ctb] (FIFR
heuristic by Wu et al.)

Maintainer Martin Papenberg <martin.papenberg@hhu.de>

Description The method of anticlustering partitions a pool of elements into groups (i.e., anticlusters) with the goal of maximizing between-group similarity or within-group heterogeneity. The anticlustering approach thereby reverses the logic of cluster analysis that strives for high within-group homogeneity and clear separation between groups. Computationally, anticlustering is accomplished by maximizing instead of minimizing a clustering objective function, such as the intra-cluster variance (used in k-means clustering) or the sum of pairwise distances within clusters. The main function anticlustering() gives access to optimal and heuristic anticlustering methods described in Papenberg and Klau (2021; <[doi:10.1037/met0000301](https://doi.org/10.1037/met0000301)>), Brusco et al. (2020; <[doi:10.1111/bmsp.12186](https://doi.org/10.1111/bmsp.12186)>), Papenberg (2024; <[doi:10.1111/bmsp.12315](https://doi.org/10.1111/bmsp.12315)>), Papenberg, Wang, et al. (2025; <[doi:10.1016/j.crmeth.2025.101137](https://doi.org/10.1016/j.crmeth.2025.101137)>), Papenberg, Breuer, et al. (2025; <[doi:10.1017/psy.2025.10052](https://doi.org/10.1017/psy.2025.10052)>), Yang et al. (2022; <[doi:10.1016/j.ejor.2022.02.003](https://doi.org/10.1016/j.ejor.2022.02.003)>), and Wu et al. (2025; <[doi:10.1016/j.cor.2025.107030](https://doi.org/10.1016/j.cor.2025.107030)>). The optimal algorithms require that an integer linear programming solver is installed. This package will install 'lpSolve'

(<https://cran.r-project.org/package=lpSolve>) as a default solver, but it is also possible to use the package 'Rglpk' (<https://cran.r-project.org/package=Rglpk>), which requires the GNU linear programming kit (<https://www.gnu.org/software/glpk/glpk.html>), the package 'Rsymphony' (<https://cran.r-project.org/package=Rsymphony>), which requires the SYMPHONY ILP solver (<https://github.com/coin-or/SYMPHONY>), or the commercial solver Gurobi, which provides its own R package that is not available via CRAN (<https://www.gurobi.com/downloads/>). 'Rglpk', 'Rsymphony', 'gurobi' and their system dependencies have to be manually installed by the user because they are only suggested dependencies. Full access to the bicriterion anticlustering method proposed by Brusco et al. (2020) is given via the function `bicriterion_anticlustering()`, while `kplus_anticlustering()` implements the full functionality of the k-plus anticlustering approach proposed by Papenberg (2024). Some other functions are available to solve classical clustering problems. The function `balanced_clustering()` applies a cluster analysis under size constraints, i.e., creates equal-sized clusters. The function `matching()` can be used for (unrestricted, bipartite, or K-partite) matching. The function `wce()` can be used optimally solve the (weighted) cluster editing problem, also known as correlation clustering, clique partitioning problem or transitivity clustering.

License MIT + file LICENSE

URL <https://github.com/m-Py/anticlust>,
<https://m-py.github.io/anticlust/>

BugReports <https://github.com/m-Py/anticlust/issues>

Depends R (>= 3.6.0)

Imports Matrix, RANN (>= 2.6.0), lpSolve, MASS

Suggests knitr, Rglpk, rmarkdown, Rsymphony, tinytest, gurobi,
tableone, palmerpenguins

VignetteBuilder knitr, rmarkdown

Encoding UTF-8

LazyData true

NeedsCompilation yes

SystemRequirements Rendering the vignette requires pandoc
(<https://pandoc.org/>).

Config/roxygen2/version 8.0.0

Repository <https://m-py.r-universe.dev>

Date/Publication 2026-09-15 10:23:21 UTC

RemoteUrl <https://github.com/m-py/anticlust>

```
RemoteRef HEAD
RemoteSha 73b9c587bf6459eab37c6c966289c8eece821f5a
```

Contents

anticlustering	3
balanced_clustering	11
bicriterion_anticlustering	13
brunel2025	17
categorical_sampling	18
categories_to_binary	19
dispersion_objective	21
diversity_objective	22
experimental_allocation	23
fast_anticlustering	27
feasible_and_infeasible_region_search_anticlustering	30
generate_exchange_partners	32
generate_partitions	34
kplus_anticlustering	36
kplus_moment_variables	38
mahalanobis_distance	40
matching	41
mean_sd_tab	44
n_partitions	45
optimal_anticlustering	45
optimal_dispersion	47
plot_clusters	51
plot_similarity	53
schaper2019	54
three_phase_search_anticlustering	55
variance_objective	57
wce	59
Index	61

anticlustering	<i>Anticlustering</i>
----------------	-----------------------

Description

Partition a pool of elements into groups (i.e., anticlusters) with the aim of creating high within-group heterogeneity and high between-group similarity. Anticlustering is accomplished by maximizing instead of minimizing a clustering objective function. Implements anticlustering methods as described in Papenberg and Klau (2021; <doi:10.1037/met0000301>), Brusco et al. (2020; <doi:10.1111/bmsp.12186>), Papenberg (2024; <doi:10.1111/bmsp.12315>), Papenberg, Wang, et al. (2025; <doi:10.1016/j.crmeth.2025.101137>), Papenberg, Breuer, et al. (2025; <doi:10.1017/psy.2025.10052>), Yang et al. (2022; <doi:10.1016/j.ejor.2022.02.003>), and Wu et al. (2025; <doi.org/10.1016/j.cor.2025.107030>).

Usage

```

anticlustering(
  x,
  K,
  objective = "diversity",
  method = "exchange",
  preclustering = FALSE,
  categories = NULL,
  repetitions = NULL,
  standardize = FALSE,
  cannot_link = NULL,
  must_link = NULL,
  blocks = NULL
)

```

Arguments

<code>x</code>	The data input. Can be one of two structures: (1) A feature matrix where rows correspond to elements and columns correspond to variables (a single numeric variable can be passed as a vector). (2) An $N \times N$ matrix dissimilarity matrix; can be an object of class <code>dist</code> (e.g., returned by <code>dist</code> or <code>as.dist</code>) or a matrix where the entries of the upper and lower triangular matrix represent pairwise dissimilarities.
<code>K</code>	How many anticlusters should be created. Alternatively: (a) A vector describing the size of each group, or (b) a vector of length <code>nrow(x)</code> describing how elements are assigned to anticlusters before the optimization starts.
<code>objective</code>	The objective to be maximized. The options "diversity" (default; previously called "distance", which is still supported), "average-diversity", "variance", "kplus" and "dispersion" are natively supported. May also be a user-defined function. See Details.
<code>method</code>	One of "exchange" (default), "local-maximum", "brusco", "ilp", "2PML", "3phase", or "fifr". See Details.
<code>preclustering</code>	Boolean. Should a preclustering be conducted before anticlusters are created? Defaults to FALSE. See Details.
<code>categories</code>	A vector, data.frame or matrix representing one or several categorical variables whose distribution should be similar between groups. See Details.
<code>repetitions</code>	The number of times a search heuristic is initiated when using <code>method = "exchange"</code> , <code>method = "local-maximum"</code> , <code>method = "brusco"</code> , <code>method = "3phase"</code> , <code>method = "fifr"</code> , or <code>method = "2PML"</code> . In the end, the best objective found across the repetitions is returned.
<code>standardize</code>	Boolean. If TRUE and <code>x</code> is a feature matrix, the data is standardized through a call to <code>scale</code> before the optimization starts. This argument is silently ignored if <code>x</code> is a distance matrix.
<code>cannot_link</code>	A 2 column matrix where each row has the indices of two elements that must not be assigned to the same anticluster. Alternatively a vector of length <code>nrow(x)</code> ; elements having the same value in this vector cannot be assigned to the same anticluster.

<code>must_link</code>	A numeric vector of length <code>nrow(x)</code> . Elements having the same value in this vector are assigned to the same anticlust.
<code>blocks</code>	A vector, data.frame or matrix representing one or several categorical variables. Anticlustering is done sequentially within blocks. See details.

Details

This function is used to solve anticlustering. That is, the data input is divided into K groups in such a way that elements within groups are heterogeneous and the different groups are similar. Anticlustering is accomplished by maximizing instead of minimizing a clustering objective function. The maximization of five objectives is natively supported (other functions can also be defined by the user as described below):

- the diversity, setting `objective = "diversity"` (this is the default objective)
- the average diversity, which normalizes the diversity by cluster size, setting `objective = "average-diversity"`
- the k-means (or "variance") objective, setting `objective = "variance"`
- the k-plus objective, an extension of the k-means objective, setting `objective = "kplus"`
- the dispersion, which is the minimum distance between any two elements within the same cluster (setting `objective = "dispersion"`)

The k-means objective is the within-group variance—that is, the sum of the squared distances between each element and its cluster center (see [variance_objective](#)). K-means anticlustering focuses on minimizing differences with regard to the means of the input variables (that is, the columns in x), but it ignores any other distribution characteristics such as the variance / standard deviation. K-plus anticlustering (using `objective = "kplus"`) is an extension of the k-means criterion that also minimizes differences with regard to the standard deviations between groups (for details see [kplus_anticlustering](#)). K-plus anticlustering can also be extended towards higher order moments such as skew and kurtosis; to consider these additional distribution characteristics, use the function [kplus_anticlustering](#). Setting `objective = "kplus"` in `anticlustering` function will only consider means and standard deviations (in my experience, this is what users usually want). It is strongly recommended to set the argument `standardize = TRUE` when using the k-plus objective.

The "diversity" objective is the sum of pairwise distances of elements within the same groups (see [diversity_objective](#)). Hence, anticlustering using the diversity criterion maximizes between-group similarity by maximizing within-group heterogeneity (represented as the sum of all pairwise distances). In the publication that introduced the `anticlust` package (Papenberg & Klau, 2021), we used the term "anticluster editing" to refer to the maximization of the diversity, because the reversed procedure - minimizing the diversity - is also known as "cluster editing". If the diversity is computed on the basis of the Euclidean distance (which is the default behaviour if x is a feature matrix), the diversity is an all rounder objective that tends to equate all distribution characteristics between groups (such as means, variances, ...). Note that the equivalence of within-group heterogeneity and between-group similarity only holds for equal-sized groups. For unequal-sized groups, it is recommended to use a different objective when striving for overall between-group similarity, e.g., the k-plus objective or the "average-diversity". The average diversity was introduced in version 0.8.6, and it is more useful if groups are not equal-sized. The average diversity normalizes the sum of intra-cluster distances by group size. If all groups are equal-sized, it is equivalent to the regular diversity.

The "dispersion" is the minimum distance between any two elements that are part of the same cluster; maximization of this objective ensures that any two elements within the same group are as dissimilar from each other as possible. Applications that require high within-group heterogeneity often require to maximize the dispersion. Oftentimes, it is useful to also consider the diversity and not only the dispersion; to optimize both objectives at the same time, see the function [bicriterion_anticlustering](#).

In the standard case, groups of equal size are generated. Adjust the argument `K` to create groups of different size (see Examples).

Data input

The data input `x` is usually a feature matrix where each row is a "case" and each column is a "variable". The default objective = "diversity" or the objective = "dispersion" are computed using a dissimilarity matrix as input, however. The default behaviour for feature matrices is to compute the Euclidean distance as the basic unit of these objectives, using a call to [dist](#). The combination objective = "diversity" + Euclidean distance was evaluated by us (e.g., Papenberg and Klau, 2021; Papenberg, 2024; Papenberg, Wang, et al., 2025) and is generally considered to be a good default for equal-sized groups. For unequal-sized groups, prefer objective = "average-diversity".

If a different type of dissimilarity measure is preferred (e.g., the Gower distance), you may pass a self-generated dissimilarity matrix via the argument `x`. This dissimilarity matrix can be an object of class `dist` (e.g., returned by [dist](#) or [as.dist](#)) or a matrix where the entries of the upper and lower triangular matrix represent pairwise dissimilarities. If the input is a matrix rather than of class `dist`, the function internally tests if the input is a dissimilarity matrix (because a matrix could also be a feature matrix). It assumes the input is a dissimilarity matrix if the upper and lower triangular consists of the same values, and if the diagonal only consists of the same values. If the input is a `data.frame`, this function always assumes that it is a feature matrix. In a `data.frame`, each column must be either numeric or factor (the latter for categorical variables).

As of version 0.8.12, this function supports handling of missing values (NA). In this case the function [dist](#) handles the NAs when converting features to pairwise distances. Note that [dist](#) always handles missing values even for objectives that are usually computed using a feature matrix as input (i.e., objective = "variance" and objective = "kplus"). In this case, we exploit an equivalence between the objective = "variance" and the objective = "average-diversity": When the squared Euclidean distance objective = "variance" and the objective = "average-diversity" are actually equivalent, so we can optimize the objective = "average-diversity" in order to perform k-means (and thus, k-plus) anticlustering. See Papenberg, Breuer, et al. (2025) for details.

Algorithms for anticlustering

By default, a heuristic method is employed for anticlustering: the exchange method (method = "exchange"). First, elements are randomly assigned to anticlusters (It is also possible to explicitly specify the initial assignment using the argument `K`; in this case, `K` has length `nrow(x)`.) Based on the initial assignment, elements are systematically swapped between anticlusters in such a way that each swap improves the objective value. For an element, each possible swap with elements in other clusters is simulated; then, the one swap is performed that improves the objective the most, but a swap is only conducted if there is an improvement at all. This swapping procedure is repeated for each element. When using method = "local-maximum", the exchange method does not terminate after the first iteration over all elements; instead, the swapping continues until a local maximum is reached. This method corresponds to the algorithm "LCW" by Weitz and Lakshminarayanan (1998). This means that after the exchange process has been conducted once for each data point,

the algorithm restarts with the first element and proceeds to conduct exchanges until the objective cannot be improved.

When setting `preclustering = TRUE`, only the $K - 1$ most similar elements serve as exchange partners for each element, which can speed up the optimization (more information on the preclustering heuristic follows below). If the `categories` argument is used, only elements having the same value in `categories` serve as exchange partners.

Using `method = "brusco"` implements the local bicriterion iterated local search (BILS) heuristic by Brusco et al. (2020) and returns the partition that best optimized either the diversity or the dispersion during the optimization process. The function `bicriterion_anticlustering` can also be used to run the algorithm by Brusco et al., but it returns multiple partitions that approximate the optimal pareto efficient set according to both objectives (diversity and dispersion). Thus, to fully utilize the BILS algorithm, use the function `bicriterion_anticlustering`.

Since version 0.8.12, `method = "3phase"` calls the three phase search algorithm for anticlustering by Yang et al. (2022). For `anticlust`, we implemented some changes of their default procedure, in particular they used a maximum "time" that the algorithm runs before terminating. For consistency with our other methods in `anticlust`, we are however using a maximum number of repetitions. Our default is 50 repetitions, which may be increased (but not decreased!) with the `repetitions` argument. Using the same adaptation for the repetitions argument, version 0.8.16 introduced for the Feasible and infeasible region search algorithm by Wu et al. (2025) via `method = "fifr"`.

Optimal anticlustering

Usually, heuristics are employed to tackle anticlustering problems, and their performance is generally very satisfying. However, heuristics do not investigate all possible group assignments and therefore do not (necessarily) find the "globally optimal solution", i.e., a partitioning that has the best possible value with regard to the objective that is optimized. Enumerating all possible partitions in order to find the best solution, however, quickly becomes impossible with increasing N , and therefore it is not possible to find a global optimum this way. Because all anticlustering problems considered here are also NP-hard, there is also no (known) clever algorithm that might identify the best solution without considering all possibilities - at least in the worst case. Integer linear programming (ILP) is an approach for tackling NP hard problems that nevertheless tries to be clever when finding optimal solutions: It does not necessarily enumerate all possibilities but is still guaranteed to return the optimal solution. Still, for NP hard problems such as anticlustering, ILP methods will also fail at some point (i.e., when N increases).

`anticlust` implements optimal solution algorithms via integer linear programming. In order to use the ILP methods, set `method = "ilp"`. The integer linear program optimizing the diversity was described in Papenberg & Klau, (2021; (8) - (12)). It can also be used to optimize the k-means and k-plus objectives, but you actually have to use the function `optimal_anticlustering` for these objectives. The documentation of the function `optimal_dispersion` and `optimal_anticlustering` contain more information on the optimal anticlustering algorithms.

Categorical variables

There are several ways to balance categorical variables among anticlusters (also see the package vignette "Using categorical variables with anticlustering"). The first way is to treat them as "hard constraints" via the argument `categories` (see Papenberg & Klau, 2021). If done so, balancing the categorical variable is accomplished via `categorical_sampling` through a stratified split before the anticlustering optimization. After that, the balance is never changed when the algorithm runs (hence, it is a "hard constraint"). When `categories` has multiple columns (i.e., there are multiple

categorical variables), each combination of categories is treated as a distinct category by the exchange method (i.e., the multiple columns are "merged" into a single column). This behaviour may lead to less than optimal results on the level of each single categorical variable. In this case, it may be useful to treat the categorical variables as part of the numeric data, i.e., the first argument `x` via binary coding (e.g. using `categories_to_binary`). The examples show how to do this when using the bicriterion algorithm by Brusio et al. Using the argument `categories` is only available for the classical exchange procedures, that is, for `method = "exchange"` and `method = "local-maximum"`.

As of version 0.8.12, the data input `x` can contain variables of type `factor` to represent categorical variables. In this case, they are internally converted to binary via `categories_to_binary`, which no longer needs to be done manually by users (and which can be tricky in particular when combining categorical and numeric variables, and when using the `k-plus` objective). However note that the internal conversion always uses `use_combinations = FALSE`, which may not always what users need. In this case, I still recommend using `categories_to_binary` manually.

The argument `blocks` (available as of version 0.8.13) is similar to `categories`. Anticlustering is repeated within each stratum (i.e., each levels or combination of levels of blocks – as in `categories` multiple variables are merged into a single variable to form strata). Anticlustering in later blocks "remembers" the assignment of previous blocks. This way, overall balance as well as balance within blocks is maximized. This is unlike when using `categories`, which only maximizes overall balance, but not specifically within strata.

Anticlustering with constraints

Versions 0.8.6 and 0.8.7 of `antclust` introduced the possibility to induce cannot-link and must-link constraints with anticlustering with the arguments `cannot_link` and `must_link`, respectively. Cannot-link constraints ensure that pairs of items are assigned to different clusters. They are given as a 2-column matrix, where each row has the indices of two elements, which must not be assigned to the same cluster. It is possible that a set of cannot-link constraints cannot be fulfilled. To verify whether the constraints cannot be fulfilled (and to actually assign elements while respecting the constraints), a graph coloring algorithm is used. This algorithm is actually the same method as used in `optimal_dispersion`. The graph coloring algorithm uses an ILP solver and it greatly profits (that is, it may be much faster) from the `Rsymphony` package, which is not installed as a necessary dependency with `antclust`. It is therefore recommended to manually install the `Rsymphony` package, which is then automatically selected as solver when using the `must_link` argument. If you have access to the `gurobi` solver and have the `gurobi` R package installed, it will be selected as solver (which is even faster than `Symphony`). As of version 0.8.11, it is also possible to use `cannot_link` as a vector. In this case it is ensured that elements having the same value in `cannot_link` are not linked in the same cluster. The methods for handling `cannot_link` constraints have been described in Papenberg, Breuer, et al. (2025).

Must-link constraints are passed as a single vector of length `nrow(x)`. Positions that have the same numeric index are assigned to the same anticluster (if the constraints can be fulfilled). When including must-link constraints, `method = "2PML"` performs a specialized search heuristic that potentially yields better results than `method = "local-maximum"`. The must-link functionality and the 2PML algorithm was introduced in Papenberg, Wang, et al. (2025).

The examples illustrate the usage of the `must_link` and `cannot_link` arguments. Currently, the different kinds of constraints (arguments `must_link`, `cannot_link`, and `categories`) cannot be used together, but this may change in future versions.

Preclustering

A useful heuristic for anticlustering is to form small groups of very similar elements and assign

these to different groups. This logic is used as a preprocessing when setting `preclustering = TRUE`. That is, before the anticlustering objective is optimized, a cluster analysis identifies small groups of similar elements (pairs if $K = 2$, triplets if $K = 3$, and so forth). The optimization of the anticlustering objective is then conducted under the constraint that these matched elements cannot be assigned to the same group. When using the exchange algorithm, preclustering is conducted using a call to `matching`. When using `method = "ilp"`, the preclustering optimally finds groups of minimum pairwise distance by solving the integer linear program described in Papenberg and Klau (2021; (8) - (10), (12) - (13)). Note that when combining preclustering restrictions with `method = "ilp"`, the anticlustering result is no longer guaranteed to be globally optimal, but only optimal given the preclustering restrictions.

Optimize a custom objective function

It is possible to pass a function to the argument `objective`. See `dispersion_objective` for an example. If `objective` is a function, the exchange method assigns elements to anticlusters in such a way that the return value of the custom function is maximized (hence, the function should return larger values when the between-group similarity is higher). The custom function has to take two arguments: the first is the data argument, the second is the clustering assignment. That is, the argument `x` will be passed down to the user-defined function as first argument. **However, only after `as.matrix` has been called on `x`.** This implies that in the function body, columns of the data set cannot be accessed using `data.frame` operations such as `$`. Objects of class `dist` will be converted to matrix as well.

Value

A vector of length N that assigns a group (i.e., a number between 1 and K) to each input element.

Author(s)

Martin Papenberg <martin.papenberg@hhu.de>

References

- Brusco, M. J., Cradit, J. D., & Steinley, D. (2020). Combining diversity and dispersion criteria for anticlustering: A bicriterion approach. *British Journal of Mathematical and Statistical Psychology*, 73, 275-396. <https://doi.org/10.1111/bmsp.12186>
- Papenberg, M., Breuer, M., Diekhoff, M., Tran, N. K., & Klau, G. W. (2025). Extending the Bicriterion Approach for Anticlustering: Exact and Hybrid Approaches. *Psychometrika*, 90(5), 1789-1808. <https://doi.org/10.1017/psy.2025.10052>
- Papenberg, M., & Klau, G. W. (2021). Using anticlustering to partition data sets into equivalent parts. *Psychological Methods*, 26(2), 161-174. <https://doi.org/10.1037/met0000301>
- Papenberg, M. (2024). K-plus Anticlustering: An Improved k-means Criterion for Maximizing Between-Group Similarity. *British Journal of Mathematical and Statistical Psychology*, 77(1), 80-102. <https://doi.org/10.1111/bmsp.12315>
- Papenberg, M., Wang, C., Diop, M., Bukhari, S. H., Oskotsky, B., Davidson, B. R., Vo, K. C., Liu, B., Irwin, J. C., Combes, A., Gaudilliere, B., Li, J., Stevenson, D. K., Klau, G. W., Giudice, L. C., Sirota, M., & Oskotsky, T. T. (2025). Anticlustering for sample allocation to minimize batch effects. *Cell Reports Methods*, 5(8), 101137. <https://doi.org/10.1016/j.crmeth.2025.101137>

Späth, H. (1986). Anticlustering: Maximizing the variance criterion. *Control and Cybernetics*, 15, 213-218.

Weitz, R. R., & Lakshminarayanan, S. (1998). An empirical comparison of heuristic methods for creating maximally diverse groups. *Journal of the Operational Research Society*, 49(6), 635-646. <https://doi.org/10.1057/palgrave.jors.2600510>

Wu, X., Feng, J., Yang, J., & Zhang, Y. (2025). Feasible and infeasible region search for the maximally diverse grouping problem. *Computers & Operations Research*, 179, 107030. <https://doi.org/10.1016/j.cor.2025.107030>

Yang, X., Cai, Z., Jin, T., Tang, Z., & Gao, S. (2022). A three-phase search approach with dynamic population size for solving the maximally diverse grouping problem. *European Journal of Operational Research*, 302(3), 925-953. <https://doi.org/10.1016/j.ejor.2022.02.003>

Examples

```
# Use default method ("exchange") and the default diversity criterion, also include
# a categorical variable via argument `categories`:
anticlusters <- anticlustering(
  schaper2019[, 3:6],
  K = 3,
  categories = schaper2019$room
)
# Compare feature means and standard deviations by anticluster
mean_sd_tab(schaper2019[, 3:6], anticlusters)
# Verify that the "room" is balanced across anticlusters:
table(anticlusters, schaper2019$room)

# Use multiple starts of the algorithm to improve the objective and
# optimize the k-means criterion ("variance")
anticlusters <- anticlustering(
  schaper2019[, 3:6],
  objective = "variance",
  K = 3,
  categories = schaper2019$room,
  method = "local-maximum", # better search algorithm
  repetitions = 20 # multiple restarts of the algorithm
)
# Compare means and standard deviations by anticluster
mean_sd_tab(schaper2019[, 3:6], anticlusters)

# Use different group sizes and optimize the extended k-means
# criterion ("kplus")
anticlusters <- anticlustering(
  schaper2019[, 3:6],
  objective = "kplus",
  K = c(24, 24, 48),
  categories = schaper2019$room,
  repetitions = 20,
  method = "local-maximum",
  standardize = TRUE # ususally recommended
)

# Use cannot_link constraints: Element 1 must not be linked with elements 2 to 10:
```

```

cl_matrix <- matrix(c(rep(1, 9), 2:10), ncol = 2)
cl <- anticlustering(
  schaper2019[, 3:6],
  K = 10,
  cannot_link = cl_matrix
)
all(cl[1] != cl[2:10])

# Use cannot_link constraints: Element 1 must be linked with elements 2 to 10.
# Element 11 must be linked with elements 12-20.
must_link <- rep(NA, nrow(schaper2019))
must_link[1:10] <- 1
must_link[11:20] <- 2
cl <- anticlustering(
  schaper2019[, 3:6],
  K = 3,
  must_link = must_link
)
cl[1:10]
cl[11:20]

# Use the heuristic by Brusco et al. (2020) for k-plus anticlustering
# Include categorical variable as part of the optimization criterion rather
# than the argument categories!
anticlusters <- anticlustering(
  cbind(
    kplus_moment_variables(schaper2019[, 3:6], 2),
    categories_to_binary(schaper2019$room)
  ),
  objective = "variance", # k-plus anticlustering because of the input above!
  K = 3,
  repetitions = 20,
  method = "brusco"
)

mean_sd_tab(schaper2019[, 3:6], anticlusters)
table(anticlusters, schaper2019$room)

```

balanced_clustering *Create balanced clusters of equal size*

Description

Create balanced clusters of equal size

Usage

```
balanced_clustering(x, K, method = "centroid", solver = NULL)
```

Arguments

x	The data input. Can be one of two structures: (1) A feature matrix where rows correspond to elements and columns correspond to variables (a single numeric variable can be passed as a vector). (2) An $N \times N$ matrix dissimilarity matrix; can be an object of class <code>dist</code> (e.g., returned by <code>dist</code> or <code>as.dist</code>) or a matrix where the entries of the upper and lower triangular matrix represent pairwise dissimilarities.
K	How many clusters should be created.
method	One of "centroid" or "ilp". See Details.
solver	Optional. The solver used to obtain the optimal method if <code>method = "ilp"</code> . Currently supports "glpk" and "symphony". Is ignored for <code>method = "centroid"</code> .

Details

This function partitions a set of elements into K equal-sized clusters. The function offers two methods: a heuristic and an exact method. The heuristic (`method = "centroid"`) first computes the centroid of all data points. If the input is a feature matrix, the centroid is defined as the mean vector of all columns. If the input is a dissimilarity matrix, the most central element acts as the centroid; the most central element is defined as the element having the minimum maximal distance to all other elements. After identifying the centroid, the algorithm proceeds as follows: The element having the highest distance from the centroid is clustered with its $(N/K) - 1$ nearest neighbours (neighbourhood is defined according to the Euclidean distance if the data input is a feature matrix). From the remaining elements, again the element farthest to the centroid is selected and clustered with its $(N/K) - 1$ neighbours; the procedure is repeated until all elements are part of a cluster.

An exact method (`method = "ilp"`) can be used to solve equal-sized weighted cluster editing optimally (implements the integer linear program described in Papenberg and Klau, 2020; (8) - (10), (12) - (13)). The cluster editing objective is the sum of pairwise distances within clusters; clustering is accomplished by minimizing this objective. If the argument `x` is a features matrix, the Euclidean distance is computed as the basic unit of the cluster editing objective. If another distance measure is preferred, users may pass a self-computed dissimilarity matrix via the argument `x`.

The optimal `method = "ilp"` uses a "solver" to optimize the clustering objective. See [optimal_anticlustering](#) for an overview of the solvers that are available.

Value

An integer vector representing the cluster affiliation of each data point

Author(s)

Martin Papenberg <martin.papenberg@hhu.de>

Meik Michalke <meik.michalke@hhu.de>

Source

The centroid method was originally developed and contributed by Meik Michalke. It was later rewritten by Martin Papenberg, who also implemented the integer linear programming method.

References

- Grötschel, M., & Wakabayashi, Y. (1989). A cutting plane algorithm for a clustering problem. *Mathematical Programming*, 45, 59–96.
- Papenberg, M., & Klau, G. W. (2021). Using anticlustering to partition data sets into equivalent parts. *Psychological Methods*, 26(2), 161–174. <https://doi.org/10.1037/met0000301>.

Examples

```
# Cluster a data set and visualize results
N <- 1000
lds <- data.frame(f1 = rnorm(N), f2 = rnorm(N))
cl <- balanced_clustering(lds, K = 10)
plot_clusters(lds, clusters = cl)

# Repeat using a distance matrix as input
cl2 <- balanced_clustering(dist(lds), K = 10)
plot_clusters(lds, clusters = cl2)
```

bicriterion_anticlustering

Bicriterion iterated local search heuristic

Description

This function implements the bicriterion algorithm BILS for anticlustering by Brusco et al. (2020; <doi:10.1111/bmsp.12186>). The description of their algorithm is given in Section 3 of their paper (in particular, see the Pseudocode in their Figure 2). It also implements some extensions to the bicriterion approach for anticlustering described in Papenberg, Breuer, et al. (2025).

Usage

```
bicriterion_anticlustering(
  x,
  K,
  R = NULL,
  W = c(1e-06, 1e-05, 1e-04, 0.001, 0.01, 0.1, 0.5, 0.99, 0.999, 0.999999),
  Xi = c(0.05, 0.1),
  dispersion_distances = NULL,
  average_diversity = FALSE,
  init_partitions = NULL,
  return = "paretoset"
)
```

Arguments

x	The data input. Can be one of two structures: (1) A feature matrix where rows correspond to elements and columns correspond to variables (a single numeric variable can be passed as a vector). (2) An N x N matrix dissimilarity matrix; can be an object of class <code>dist</code> (e.g., returned by <code>dist</code> or <code>as.dist</code>) or a matrix where the entries of the upper and lower triangular matrix represent pairwise dissimilarities.
K	How many anticlusters should be created. Alternatively: (a) A vector describing the size of each group, or (b) a vector of length <code>nrow(x)</code> describing how elements are assigned to anticlusters before the optimization starts.
R	The desired number of restarts for the algorithm. By default, both phases (MBPI + ILS) of the algorithm are performed once. See details.
W	Optional argument, a vector of weights defining the relative importance of dispersion and diversity ($0 \leq W \leq 1$). See details.
Xi	Optional argument, specifies probability of swapping elements during the iterated local search. See examples.
dispersion_distances	A distance matrix used to compute the dispersion if the dispersion should not be computed on the basis of argument x.
average_diversity	Boolean. Compute the diversity not as a global sum across all pairwise within-group distances, but as the sum of the average of within-group distances.
init_partitions	A matrix of initial partitions (rows = partitions; columns = elements) that serve as starting partitions during the iterations of the first phase of the BILS (i.e., the MBPI). If not passed, a new random partition is generated at the start of each iteration (which is the default behaviour).
return	Either "paretoiset" (default), "best-diversity", "best-average-diversity", "best-dispersion". See below.

Details

The bicriterion algorithm by Brusco et al. (2020) aims to simultaneously optimize two anticlustering criteria: the `diversity_objective` and the `dispersion_objective`. It returns a list of partitions that approximate the pareto set of efficient solutions across both criteria. By considering both the diversity and dispersion, this algorithm is well-suited for maximizing overall within-group heterogeneity. To select a partition among the approximated pareto set, it is reasonable to plot the objectives for each partition (see Examples).

The arguments R, W and Xi are named for consistency with Brusco et al. (2020). The argument K is used for consistency with other functions in `anticlust`; Brusco et al. used 'G' to denote the number of groups. However, note that K can not only be used to denote the number of equal-sized groups, but also to specify group sizes, as in `anticlustering`.

This function implements the combined bicriterion algorithm BILS, which consists of two phases: The multistart bicriterion pairwise interchange heuristic (MBPI, which is a local maximum search heuristic similar to `method = "local-maximum"` in `anticlustering`), and the iterated local search

(ILS), which is an improvement procedure that tries to overcome local optima. The argument *R* denotes the number of restarts of the two phases of the algorithm. If *R* has length 1, half of the repetitions perform the first phase MBPI, the other half perform the ILS. If *R* has length 2, the first entry indicates the number of restarts of MBPI the second entry indicates the number of restarts of ILS. The argument *W* denotes the relative weight given to the diversity and dispersion criterion in a given run of the search heuristic. In each run, the a weight is randomly selected from the vector *W*. As default values, we use the weights that Brusco et al. used in their analyses. All values in *W* have to be in $[0, 1]$; larger values indicate that diversity is more important, whereas smaller values indicate that dispersion is more important; $w = .5$ implies the same weight for both criteria. The argument *Xi* is the probability that an element is swapped during the iterated local search (specifically, *Xi* has to be a vector of length 2, denoting the range of a uniform distribution from which the probability of swapping is selected). For *Xi*, the default is selected consistent with the analyses by Brusco et al.

If the data input *x* is a feature matrix (that is: each row is a "case" and each column is a "variable"), a matrix of the Euclidean distances is computed as input to the algorithm. If a different measure of dissimilarity is preferred, you may pass a self-generated dissimilarity matrix via the argument *x*. The argument *dispersion_distances* can additionally be used if the dispersion should be computed on the basis of a different distance matrix.

As of anticlust version 0.8.6, this function includes some extensions to the original BILS by Brusco et al. (2020). These extensions are implemented through the optional arguments *dispersion_distances*, *average_diversity*, *init_partitions*, and *return*, and were described in Papenberg, Breuer, et al. (2025). If these arguments are not changed, the function performs the "vanilla" BILS as described in Brusco et al. If multiple *init_partitions* are given, ensure that each partition (i.e., each row of *init_partitions*) has the exact same output of [table](#).

Value

By default, a matrix of anticlustering partitions (i.e., the approximated pareto set). Each row corresponds to a partition, each column corresponds to an input element. If the argument *return* is set to either "best-diversity", "best-average-diversity", or "best-dispersion", it only returns one partition (as a vector), that maximizes the respective objective.

Note

For technical reasons, the pareto set returned by this function has a limit of 500 partitions. Usually however, the algorithm usually finds much fewer partitions. There is one following exception: We do not recommend to use this method when the input data is one-dimensional where the algorithm may identify too many equivalent partitions causing it to run very slowly (see section 5.6 in Breuer, 2020).

Author(s)

Martin Breuer <M.Breuer@hhu.de>, Martin Papenberg <martin.papenberg@hhu.de>

References

Brusco, M. J., Cradit, J. D., & Steinley, D. (2020). Combining diversity and dispersion criteria for anticlustering: A bicriterion approach. *British Journal of Mathematical and Statistical Psychology*, 73, 275-396. <https://doi.org/10.1111/bmsp.12186>

Breuer (2020). Using anticlustering to maximize diversity and dispersion: Comparing exact and heuristic approaches. Bachelor thesis.

Papenberg, M., Breuer, M., Diekhoff, M., Tran, N. K., & Klau, G. W. (2025). Extending the Bicriterion Approach for Anticlustering: Exact and Hybrid Approaches. *Psychometrika*, 90(5), 1789–1808. <https://doi.org/10.1017/psy.2025.10052>

Examples

```
# Generate some random data
M <- 3
N <- 80
K <- 4
data <- matrix(rnorm(N * M), ncol = M)

# Perform bicriterion algorithm, use 200 repetitions
pareto_set <- bicriterion_anticlustering(data, K = K, R = 200)

# Compute objectives for all solutions
diversities_pareto <- apply(pareto_set, 1, diversity_objective, x = data)
dispersions_pareto <- apply(pareto_set, 1, dispersion_objective, x = data)

# Plot the pareto set
plot(
  diversities_pareto,
  dispersions_pareto,
  col = "blue",
  cex = 2,
  pch = as.character(1:NROW(pareto_set))
)

# Get some random solutions for comparison
rnd_solutions <- t(replicate(n = 200, sample(pareto_set[1, ])))

# Compute objectives for all random solutions
diversities_rnd <- apply(rnd_solutions, 1, diversity_objective, x = data)
dispersions_rnd <- apply(rnd_solutions, 1, dispersion_objective, x = data)

# Plot random solutions and pareto set. Random solutions are far away
# from the good solutions in pareto set
plot(
  diversities_rnd, dispersions_rnd,
  col = "red",
  xlab = "Diversity",
  ylab = "Dispersion",
  ylim = c(
    min(dispersions_rnd, dispersions_pareto),
    max(dispersions_rnd, dispersions_pareto)
  ),
  xlim = c(
    min(diversities_rnd, diversities_pareto),
    max(diversities_rnd, diversities_pareto)
  )
)
```

```
)

# Add approximated pareto set from bicriterion algorithm:
points(diversities_pareto, dispersions_pareto, col = "blue", cex = 2, pch = 19)
```

brunel2025

Completion norms for 403 French sentences

Description

Completion norms for 403 French sentences

Usage

```
brunel2025
```

Format

A data frame with 403 rows and 7 variables

target_word The word (in French)

sentence The sentence (in French), which is completed by target_word

target_word_emotionality Categorical: Is the target word classified as "positive", "negative" or "neutral"

sentence_emotionality Categorical: Is the sentence classified as "mixt", "negative", "neutral", "non specified", or "positive"

percentage_target_word Numeric between 0 and 1; percentage of a test sample that identified the target word

valence_target_word Numeric, average valence rating for the target word

arousal_target_word Numeric, average arousal rating for the target word

Source

This data set was distributed as supplementary online materials to the paper by Brunel et al. (2025), and is available from the Open Science Repository via <https://osf.io/7pc46>. Retrieved on October 2, 2025. The data was published under a CC-BY Attribution 4.0 International license, which permits redistribution. The anticlust packages includes the raw and unchanged data file as provided by Jeremy Brunel, and a script that converts the raw data into the package data.

References

Brunel, J., Dujardin, E., Delord, S. (2025) Emotional valence, cloze probability, and entropy: Completion norms for 403 French sentences. *Behavior Research Methods* 57(81). <https://doi.org/10.3758/s13428-025-02604-7>

Examples

```
groups <- anticlustering(  
  brunel2025[, 3:7],  
  K = 4,  
  objective = "kplus",  
  standardize = TRUE,  
  method = "3phase"  
)  
# check out descriptive statistics for each variable for each of 4 groups:  
# Numeric variables  
mean_sd_tab(brunel2025[, 5:7], groups)  
# Categorical variables:  
table(groups, brunel2025$target_word_emotionality)  
table(groups, brunel2025$sentence_emotionality)
```

categorical_sampling	<i>Random sampling employing a categorical constraint</i>
----------------------	---

Description

This function can be used to obtain a stratified split of a data set.

Usage

```
categorical_sampling(categories, K)
```

Arguments

categories	A matrix or vector of one or more categorical variables.
K	The number of groups that are returned.

Details

This function can be used to obtain a stratified split of a data set. Using this function is like calling [anticlustering](#) with argument ‘categories’, but without optimizing a clustering objective. The categories are just evenly split between samples. Apart from the restriction that categories are balanced between samples, the split is random.

Value

A vector representing the sample each element was assigned to.

Examples

```
data(schaper2019)
categories <- schaper2019$room
groups <- categorical_sampling(categories, K = 6)
table(groups, categories)

# Unequal sized groups
groups <- categorical_sampling(categories, K = c(24, 24, 48))
table(groups, categories)

# Heavily unequal sized groups, is harder to balance the groups
groups <- categorical_sampling(categories, K = c(51, 19, 26))
table(groups, categories)
```

categories_to_binary *Get binary representation of categorical variables*

Description

Get binary representation of categorical variables

Usage

```
categories_to_binary(
  categories,
  use_combinations = FALSE,
  pairwise_combinations = NULL
)
```

Arguments

categories A vector, data.frame or matrix representing one or several categorical variables

use_combinations Logical, should the output also include columns representing the combination / interaction of the categories (defaults to FALSE).

pairwise_combinations Optional, logical. If `pairwise_combinations` is TRUE, only pairwise interactions are encoded (and not all higher order interactions like when only using `use_combinations`).

Details

The conversion of categorical variables to binary variables is done via `model.matrix`. Since version 0.8.9, each category of a categorical variable is coded by a separate variable ('one hot' encoding). So we do not use 'dummy' coding, which is often used to encode predictors in statistical analysis. Dummy coding uses a reference category that is not explicitly modelled via a separate variable. This

implies that there is a different distance to the reference category than among the other categories, which is unwarranted in anticlustering (thanks to Gunnar Klau for noting this). See examples.

This function can be used to include categorical variables as part of the optimization criterion in anticlustering, rather than including them as hard constraints as done when using the argument `categories` in `anticlustering` (or `fast_anticlustering`). This way, categorical variables are treated as numeric variables, which can be useful when there are several categorical variables or when the group sizes are unequal (or both). See examples. Please see the vignette 'Using categorical variables with anticlustering' for more information on this approach.

Since version 0.8.12, this function treats NA in the categorical input variables as a distinct category. This implies that different observations having an NA are treated as having the same value on this variable. If this is not desired, you need to do some other preprocessing of the NAs before calling this function.

Value

A matrix encoding the categorical variable(s) in binary form.

Author(s)

Martin Papenberg <martin.papenberg@hhu.de>

References

Papenberg, M., Wang, C., Diop, M., Bukhari, S. H., Oskotsky, B., Davidson, B. R., Vo, K. C., Liu, B., Irwin, J. C., Combes, A., Gaudilliere, B., Li, J., Stevenson, D. K., Klau, G. W., Giudice, L. C., Sirota, M., & Oskotsky, T. T. (2025). Anticlustering for sample allocation to minimize batch effects. *Cell Reports Methods*, 5(8), 101137. <https://doi.org/10.1016/j.crmeth.2025.101137>

Examples

```
# Illustrate why dummy encoding is not appropriate for anticlustering.
# Use 3 elements of the iris data set, with different levels of 'Species'
input <- iris[c(1, 51, 101), "Species", drop = FALSE]
input
# Default dummy encoding:
(dummy <- model.matrix(~ . , input))
dist(dummy)
# Distance between versicolor and virginica is larger than among setosa and
# versicolor, and setosa and virginica. This would bias the anticlustering
# computation.
(one_hot <- categories_to_binary(input))
dist(one_hot) # all the same distances
dist(one_hot)^2 # or use squared Euclidean distance (corresponds to Manhattan distance in this case)
dist(one_hot, method = "manhattan")

# We could also get one-hot encoding directly via model.matrix(), which is
# used in categories_to_binary():
model.matrix(~ . -1, input) # suppresses the 'intercept'

# Use Schaper data set for anticlustering example
```

```

data(schaper2019)
features <- schaper2019[, 3:6]
K <- 3
N <- nrow(features)

# - Generate data input for k-means anticlustering -
# We conduct k-plus anticlustering by first generating k-plus variables,
# and also include the categorical variable as "numeric" input for the
# k-means optimization (rather than as input for the argument categories)

input_data <- cbind(
  kplus_moment_variables(features, T = 2),
  categories_to_binary(schaper2019$room)
)

kplus_groups <- anticlustering(
  input_data,
  K = K,
  objective = "variance",
  method = "local-maximum",
  repetitions = 10
)
mean_sd_tab(features, kplus_groups)
table(kplus_groups, schaper2019$room) # argument categories was not used!

```

dispersion_objective *Cluster dispersion*

Description

Compute the dispersion objective for a given clustering (i.e., the minimum distance between two elements within the same cluster).

Usage

```
dispersion_objective(x, clusters)
```

Arguments

- | | |
|----------|---|
| x | The data input. Can be one of two structures: (1) A feature matrix where rows correspond to elements and columns correspond to variables (a single numeric variable can be passed as a vector). (2) An N x N matrix dissimilarity matrix; can be an object of class <code>dist</code> (e.g., returned by dist or as.dist) or a matrix where the entries of the upper and lower triangular matrix represent pairwise dissimilarities. |
| clusters | A vector representing (anti)clusters (e.g., returned by anticlustering). |

Details

The dispersion is the minimum distance between two elements within the same cluster. When the input x is a feature matrix, the Euclidean distance is used as the distance unit. Maximizing the dispersion maximizes the minimum heterogeneity within clusters and is an anticlustering task.

References

Brusco, M. J., Cradit, J. D., & Steinley, D. (2020). Combining diversity and dispersion criteria for anticlustering: A bicriterion approach. *British Journal of Mathematical and Statistical Psychology*, 73, 275-396. <https://doi.org/10.1111/bmsp.12186>

Examples

```
N <- 50 # number of elements
M <- 2  # number of variables per element
K <- 5  # number of clusters
random_data <- matrix(rnorm(N * M), ncol = M)
random_clusters <- sample(rep_len(1:K, N))
dispersion_objective(random_data, random_clusters)

# Maximize the dispersion
optimized_clusters <- anticlustering(
  random_data,
  K = random_clusters,
  objective = dispersion_objective
)
dispersion_objective(random_data, optimized_clusters)
```

diversity_objective *(Anti)cluster editing "diversity" objective*

Description

Compute the diversity for a given clustering.

Usage

```
diversity_objective(x, clusters)
```

Arguments

<code>x</code>	The data input. Can be one of two structures: (1) A data matrix where rows correspond to elements and columns correspond to features (a single numeric feature can be passed as a vector). (2) An $N \times N$ matrix dissimilarity matrix; can be an object of class <code>dist</code> (e.g., returned by <code>dist</code> or <code>as.dist</code>) or a matrix where the entries of the upper and lower triangular matrix represent the pairwise dissimilarities.
<code>clusters</code>	A vector representing (anti)clusters (e.g., returned by <code>anticlustering</code>).

Details

The objective function used in (anti)cluster editing is the diversity, i.e., the sum of the pairwise distances between elements within the same groups. When the input x is a feature matrix, the Euclidean distance is computed as the basic distance unit of this objective.

Value

The cluster editing objective

Author(s)

Martin Papenberg <martin.papenberg@hhu.de>

References

- Brusco, M. J., Cradit, J. D., & Steinley, D. (2020). Combining diversity and dispersion criteria for anticlustering: A bicriterion approach. *British Journal of Mathematical and Statistical Psychology*, 73, 275-396. <https://doi.org/10.1111/bmsp.12186>
- Papenberg, M., & Klau, G. W. (2021). Using anticlustering to partition data sets into equivalent parts. *Psychological Methods*, 26(2), 161–174. <https://doi.org/10.1037/met0000301>.

Examples

```
data(iris)
distances <- dist(iris[1:60, -5])
## Clustering
clusters <- balanced_clustering(distances, K = 3)
# This is low:
diversity_objective(distances, clusters)
## Anticlustering
anticlusters <- anticlustering(distances, K = 3)
# This is higher:
diversity_objective(distances, anticlusters)
```

experimental_allocation

Balanced Allocation of Subjects to Treatments

Description

Convenience wrapper function for `anticlustering` using different defaults and additional functionality, as used in Wintermantel et al (2026; <doi:10.48550/arXiv.2607.07543>), provided by Dimitry Wintermantel. This function assigns subjects to treatment groups while balancing specified covariates and ensuring replicates per group. If a table of groups within treatments is supplied via ‘group_data’, the subjects (already assigned to treatments) are instead allocated to those groups within each treatment, again balancing the covariates.

Usage

```
experimental_allocation(
  data,
  covariates,
  treatments = NULL,
  treatment_var = "Treatment",
  group_data = NULL,
  group_var = NULL,
  n_replicates = NULL,
  ensure_equal_n_replicates = TRUE,
  keep_excluded_data = FALSE,
  objective = "variance",
  method = "local-maximum",
  repetitions = 10,
  match_within = NULL,
  standardize = TRUE
)
```

Arguments

<code>data</code>	A data frame containing the subjects to be allocated, including columns for the covariates. If <code>'group_data'</code> is supplied, it must also contain the treatment column named in <code>'treatment_var'</code> .
<code>covariates</code>	A vector of column names in <code>'data'</code> representing covariates to balance.
<code>treatments</code>	A vector specifying the treatment group labels. Required unless <code>'group_data'</code> is supplied, in which case the target groups are taken from <code>'group_data'</code> .
<code>treatment_var</code>	A string specifying the name of the treatment column. If <code>'group_data'</code> is not supplied, this column is created in the output (defaults to <code>"Treatment"</code>). If <code>'group_data'</code> is supplied, it names the existing treatment column, which must be present in both <code>'data'</code> and <code>'group_data'</code> .
<code>group_data</code>	An optional data frame describing the available groups within each treatment. Must include columns named in <code>'group_var'</code> and <code>'treatment_var'</code> indicating which groups belong to which treatment. Any further columns (e.g. group-level covariates) are merged into the output but are not used for balancing. Defaults to <code>'NULL'</code> .
<code>group_var</code>	A string specifying the name of the group column in <code>'group_data'</code> to which subjects will be allocated within each treatment (e.g., site, cage, plot). Required when <code>'group_data'</code> is supplied.
<code>n_replicates</code>	Integer. Specifies the number of replicates per group. If <code>'NULL'</code> , either all subjects are allocated or the highest possible equal number per group is allocated depending on <code>'ensure_equal_n_replicates'</code> . Defaults to <code>'NULL'</code> .
<code>ensure_equal_n_replicates</code>	Logical. If <code>'TRUE'</code> , ensures equal numbers of replicates across groups. Defaults to <code>'TRUE'</code> . If <code>'FALSE'</code> , all subjects are allocated.
<code>keep_excluded_data</code>	Logical. If <code>'TRUE'</code> , retains subjects that were not allocated within the specified constraints. Defaults to <code>'FALSE'</code> .

objective	A string specifying the objective function for anticlustering. Options include "variance" (default), "diversity", "average-diversity", "kplus", and "dispersion".
method	A string specifying the optimization method for anticlustering. Options include "local-maximum" (default), "exchange", "brusco", "ilp".
repetitions	Integer. Specifies the number of times the optimization is repeated when using heuristic methods ("exchange", "local-maximum", "brusco"). The best solution is selected. Defaults to '10'.
match_within	A column name in 'data' (optional). Specifies a variable within which matching should occur, ensuring that subjects are grouped within subsets defined by this variable. Defaults to 'NULL'.
standardize	Logical. If 'TRUE', covariates are standardized via 'scale()' before optimization starts. Defaults to 'TRUE'.

Details

The function balances covariates among groups by creating sets of similar individuals based on the specified covariates, and then assigning the target labels to these sets so that similar individuals are spread across different groups.

If 'group_data' is supplied, the allocation is carried out separately for each treatment: for every treatment level present in 'data[[treatment_var]]', the subjects of that treatment are allocated to the groups that belong to it in 'group_data', and the result is merged with the corresponding rows of 'group_data'.

- If 'ensure_equal_n_replicates = TRUE', the number of replicates per group is enforced. - If 'keep_excluded_data = TRUE', subjects that cannot be allocated under the constraints are retained in the output with 'NA' in the allocation column. - Covariates can be scaled to standardize their ranges if 'standardize = TRUE'. - If 'match_within' is specified, subjects are matched within levels of the specified variable, ensuring allocations respect the structure defined by this variable.

Value

A data frame with subjects assigned to treatments, in the column named by 'treatment_var'. If 'group_data' is supplied, the output instead contains the subjects assigned to groups within their treatment (in the column named by 'group_var'), merged with the corresponding group information from 'group_data'. If 'keep_excluded_data = TRUE', excluded subjects are included in the output with 'NA' in the allocation column.

Validations

The function validates the following conditions: - The required arguments are provided ('data', 'covariates', and either 'treatments' or, when 'group_data' is supplied, 'group_var'). - Covariates exist in 'data', are numeric, and do not take the same value across all rows. - When 'group_data' is supplied, 'treatment_var' exists in both 'data' and 'group_data', and 'group_var' exists in 'group_data'. - 'match_within', if specified, exists in 'data'.

How to cite

If you use this function in academic work, please cite: Wintermantel, D., Osterman, J., Mair, M. M., & Hartig, F. (2026). *Equivalence testing in pesticide risk assessment – Evaluation and practical <https://doi.org/10.48550/arXiv.2607.07543>

The treatment allocation implemented here relies on anticlustering algorithms. Users are encouraged to also cite the ‘anticlust’ package where appropriate (see ‘citation("anticlust")’)

Author(s)

Dimitry Wintermantel <dimitry.wintermantel@nature.uni-freiburg.de>

Examples

```
# Example dataset: Bee subjects
example_bee_data <- data.frame(
  ID = as.factor(as.character(seq(1, 100, 1))),
  Bee_count = rnorm(100, mean = 1000, sd = 200),
  Colony_weight = rnorm(100, mean = 500, sd = 100)
)

# Example dataset: Sites (groups) with a site-level covariate and a blocking variable
example_site_data <- data.frame(
  Site = as.factor(as.character(seq(1, 16, 1))),
  Irrigated = as.factor(as.character(c(rep("yes", 4), rep("no", 12)))),
  Field_quality = rnorm(16, mean = 3, sd = 0.5)
)

treatments <- c("Control", "Pesticide")
bee_covariates <- c("Bee_count", "Colony_weight")

# 1) Allocate subjects to treatments (balancing colony-level covariates)
allocated_bee_data <- experimental_allocation(
  data = example_bee_data,
  treatments = treatments,
  covariates = bee_covariates
)

# 2) Allocate sites to treatments (balancing site quality, matching within irrigation status)
allocated_site_data <- experimental_allocation(
  data = example_site_data,
  treatments = treatments,
  covariates = "Field_quality",
  match_within = "Irrigated"
)

# 3) Allocate subjects to specific sites within each treatment
allocated_data <- experimental_allocation(
  data = allocated_bee_data,
  covariates = bee_covariates,
  treatment_var = "Treatment",
  group_data = allocated_site_data,
```

```

    group_var = "Site"
  )
  head(allocated_data)

```

fast_anticlustering *Fast anticlustering*

Description

Increasing the speed of (k-means / k-plus) anticlustering by (1) conducting fewer exchanges during the optimization and (2) using an alternative formulation of the k-means objective. Makes anticlustering applicable to quite large data sets.

Usage

```

fast_anticlustering(
  x,
  K,
  k_neighbours = Inf,
  categories = NULL,
  exchange_partners = NULL,
  distance = c("euclidean", "mahalanobis"),
  lambda = NULL
)

```

Arguments

x	A numeric vector, matrix or data.frame of data points. Rows correspond to elements and columns correspond to features. A vector represents a single numeric feature.
K	How many anticlusters should be created. Alternatively: (a) A vector describing the size of each group, or (b) a vector of length nrow(x) describing how elements are assigned to anticlusters before the optimization starts.
k_neighbours	The number of nearest neighbours that serve as exchange partner for each element. See details.
categories	A vector, data.frame or matrix representing one or several categorical constraints.
exchange_partners	Optional argument. A list of length NROW(x) specifying for each element the indices of the elements that serve as exchange partners. If used, this argument overrides the k_neighbours argument. See examples.
distance	Is the (squared) Euclidean distance used (default for k-means) or the "mahalanobis" distance?
lambda	If distance = "mahalanobis", regularization term to compute the mahalanobis distance.

Details

This function was created to make anticlustering applicable to large data sets (e.g., several 100,000 elements). It optimizes the k-means objective because computing all pairwise distances as is done when optimizing the "diversity" (i.e., the default in [anticlustering](#)) is not feasible for very large data sets (for about $N > 20000$ on my personal computer). Using `fast_anticlustering` for k-plus anticlustering is also possible by applying [kplus_moment_variables](#) on the input (and possibly by using the argument `exchange_partners`, see Examples).

The function `fast_anticlustering` employs a speed-optimized exchange method, which is basically equivalent to `method = "exchange"` in [anticlustering](#), but may reduce the number of exchanges that are investigated for each input element. The number of exchange partners per element has to be set using the argument `k_neighbours`. By default, it is set to `Inf`, meaning that all possible swaps are tested. If `k_neighbours` is set differently (which is usually recommended when running this function), the default behaviour is to generate exchange partners using a nearest neighbour search (using the function `nn2` from the RANN package). Using more exchange partners can improve the quality of the results, but also increase run time. Note that for very large data sets, anticlustering generally becomes "easier" (even a random split may yield satisfactory results), so using few exchange partners is usually not a problem.

It is possible to specify custom exchange partners using the argument `exchange_partners` instead of relying on the default nearest neighbour search. When using `exchange_partners`, it is not necessary that each element has the same number of exchange partners; this is why the argument `exchange_partners` has to be a list instead of a matrix or `data.frame`. Exchange partners can for example be generated by [generate_exchange_partners](#) (see Examples), but a custom list may also be used. Note that categorical constraints induced via `categories` may not be respected during the optimization if the `exchange_partners` argument allows exchanges between members of different categories, so care must be taken when combining the arguments `exchange_partners` and `categories`.

In `anticlustering(..., objective = "variance")`, the run time of computing the k-means objective is in $O(M N)$, where N is the total number of elements and M is the number of variables. This is because the variance is computed as the sum of squared distances between all data points and their cluster centers. The function `fast_anticlustering` uses a different - but equivalent - formulation of the k-means objective, where the re-computation of the objective only depends on M but no longer on N . In particular, this variant of k-means anticlustering minimizes the weighted sum of squared distances between cluster centroids and the overall data centroid; the distances between all individual data points and their cluster center are not computed (Späth, 1986). Using the different objective formulation reduces the run time by an order of magnitude and makes k-means anticlustering applicable to very large data sets (even in the millions). For a fixed number of exchange partners (specified using the argument `k_neighbours`), the approximate run time of `fast_anticlustering` is in $O(M N)$. The algorithm `method = "exchange"` in [anticlustering](#) with `objective = "variance"` has a run time of $O(M N^3)$. Thus, `fast_anticlustering` can improve the run time by two orders of magnitude as compared to the standard exchange algorithm. The nearest neighbour search, which is done in the beginning usually does not strongly contribute to the overall run time. It is nevertheless possible to suppress the nearest neighbour search by using the `exchange_partners` argument.

When setting the `categories` argument, exchange partners (i.e., nearest neighbours) will be generated from the same category. Note that when `categories` has multiple columns, each combination of these categories is treated as a distinct category by the exchange method. You can also use [categories_to_binary](#) to potentially improve results for several categorical variables, instead of

using the argument categories.

Author(s)

Martin Papenberg <martin.papenberg@hhu.de>

References

Papenberg, M., & Klau, G. W. (2021). Using anticlustering to partition data sets into equivalent parts. *Psychological Methods*, 26(2), 161–174. <https://doi.org/10.1037/met0000301>.

Papenberg, M. (2024). K-plus Anticlustering: An Improved k-means Criterion for Maximizing Between-Group Similarity. *British Journal of Mathematical and Statistical Psychology*, 77(1), 80–102. <https://doi.org/10.1111/bmsp.12315>

Späth, H. (1986). Anticlustering: Maximizing the variance criterion. *Control and Cybernetics*, 15, 213–218.

See Also

[anticlustering](#)
[kplus_moment_variables](#)
[categories_to_binary](#)
[variance_objective](#)
[generate_exchange_partners](#)

Examples

```
## Use fewer or more exchange partners to adjust speed (vs. quality tradeoff)
features <- iris[, - 5]
N <- nrow(features)
init <- sample(rep_len(1:3, N)) # same starting point for all calls:
groups1 <- fast_anticlustering(features, K = init) # default: all exchanges
groups2 <- fast_anticlustering(features, K = init, k_neighbours = 20)
groups3 <- fast_anticlustering(features, K = init, k_neighbours = 2)

variance_objective(features, groups1)
variance_objective(features, groups2)
variance_objective(features, groups3)

# K-plus anticlustering is straight forward when sticking with the default
# for k_neighbours
kplus_anticlusters <- fast_anticlustering(
  kplus_moment_variables(features, T = 2),
  K = 3
)
mean_sd_tab(features, kplus_anticlusters)

# Some care is needed when applying k-plus using with this function
# while using a reduced number of exchange partners generated in the
# nearest neighbour search. Then we:
```

```

# 1) Use kplus_moment_variables() on the numeric input
# 2) Generate custom exchange_partners because otherwise nearest
#     neighbours are internally selected based on the extended k-plus
#     variables returned by kplus_moment_variables()
#     (which does not really make sense)
kplus_anticlusters <- fast_anticlustering(
  kplus_moment_variables(features, T = 2),
  K = 3,
  exchange_partners = generate_exchange_partners(120, features = features, method = "RANN")
)
mean_sd_tab(features, kplus_anticlusters)
# Or we use random exchange partners:
kplus_anticlusters <- fast_anticlustering(
  kplus_moment_variables(features, T = 2),
  K = 3,
  exchange_partners = generate_exchange_partners(120, N = nrow(features), method = "random")
)
mean_sd_tab(features, kplus_anticlusters)

# Working on several 1000 elements is very fast (Here n = 10000, m = 2)
data <- matrix(rnorm(10000 * 2), ncol = 2)
start <- Sys.time()
groups <- fast_anticlustering(data, K = 5, k_neighbours = 5)
Sys.time() - start

```

feasible_and_infeasible_region_search_anticlustering
Feasible and Infeasible Region search heuristic

Description

This function implements the feasible and infeasible region search algorithm FIFR for anticlustering by Wu et al. (2025; <doi.org/10.1016/j.cor.2025.107030>). The description of their algorithm is given in Section 4 of their paper (in particular, see the Pseudocode in Algorithm 1).

Usage

```

feasible_and_infeasible_region_search_anticlustering(
  x,
  K,
  N,
  number_iterations = 50,
  clusters = NULL,
  upper_bound = NULL,
  lower_bound = NULL,
  beta_max = NULL,

```

```

    theta_max = NULL,
    theta_min = NULL,
    beta_min = NULL,
    phi = 0.8,
    tau = NULL,
    kmax = 4,
    alpha = NULL
)

```

Arguments

x	The data input. Can be one of two structures: (1) A feature matrix where rows correspond to elements and columns correspond to variables (a single numeric variable can be passed as a vector). (2) An N x N matrix dissimilarity matrix; can be an object of class <code>dist</code> (e.g., returned by <code>dist</code> or <code>as.dist</code>) or a matrix where the entries of the upper and lower triangular matrix represent pairwise dissimilarities.
K	Number of anticlusters to be formed.
N	Number of elements.
number_iterations	A number that defines how many times the steps in the search algorithm are repeated.
clusters	A vector of length M that specifies the number of elements each cluster can contain. If this vector is not NULL, the lower and upper bounds will be disregarded.
upper_bound	Maximum number of elements in each anticluster. By default, anticlusters are of equal size, calculated as the total number of items divided by the number of clusters.
lower_bound	Minimum number of elements in each anticluster. By default, anticlusters are of equal size, calculated as the total number of items divided by the number of clusters.
beta_max	The algorithm begins with a pool of random initial solutions of size <code>beta_max</code> . Over time, the size of the solution pool decreases linearly until it reaches <code>beta_min</code> .
theta_max	Parameter for the strength of undirected perturbation, which decreases linearly over time from <code>theta_max</code> to <code>theta_min</code> .
theta_min	Parameter for the strength of undirected perturbation, which decreases linearly over time from <code>theta_max</code> to <code>theta_min</code> .
beta_min	The minimum solution pool size the algorithm should reach before making a determination.
phi	Parameter that determines the population size when initiating a new round of exploration upon triggering the jump-back mechanism
tau	Parameter when "noImp" exceeds it, the algorithm transits from exploitation to exploration strategy
kmax	A parameter that determines the maximum degree of constraint violation in the IFR search
alpha	Parameter for weighting the discrimination of a slightly worse local optimal child solution

Details

Details of the implementation of the algorithm can be found in the pseudocode of the paper Wu et al. (2025)

Value

A vector of length N that assigns a group (i.e, a number between 1 and K) to each input element.

Author(s)

David Buczynski <david.buczynski@hhu.de>, Martin Papenberg <martin.papenberg@hhu.de>

References

Wu, X., Feng, J., Yang, J., & Zhang, Y. (2025). Feasible and infeasible region search for the maximally diverse grouping problem. *Computers & Operations Research*, 179, 107030. <https://doi.org/10.1016/j.cor.2025.107030>

Examples

```
# Generate some random data
N <- 120
M <- 5
K <- 3
dat <- matrix(rnorm(N * M), ncol = M)
distances <- dist(dat)

# Perform three phase search algorithm
results1 <- feasible_and_infeasible_region_search_anticlustering(distances, K, N)
results2 <- anticlustering(distances, K = K, method = "3phase")
results3 <- anticlustering(distances, K = K, method = "local-maximum", repetitions = 50)

# Compute objectives function
diversity_objective(distances, results1)
diversity_objective(distances, results2)
diversity_objective(distances, results3)
```

```
generate_exchange_partners
```

Get exchange partners for fast_anticlustering()

Description

Get exchange partners for fast_anticlustering()

Usage

```
generate_exchange_partners(
  n_exchange_partners,
  N = NULL,
  features = NULL,
  method = "random",
  categories = NULL
)
```

Arguments

<code>n_exchange_partners</code>	The number of exchange partners per element
<code>N</code>	The number of elements for which exchange partners; can be NULL if features is passed (it is ignored if features is passed).
<code>features</code>	The features for which nearest neighbours are sought if method = "RANN". May be NULL if random exchange partners are generated.
<code>method</code>	Currently supports "random" (default), "RANN" and "restricted_random". See details.
<code>categories</code>	A vector, data.frame or matrix representing one or several categorical constraints.

Details

The method = "RANN" generates exchange partners using a nearest neighbour search via [nn2](#) from the RANN package; method = "restricted_random" generates random exchange partners but ensures that for each element, no duplicates are generated and that the element itself does not occur as exchange partner (this is the slowest method, and I would not recommend it for large N); method = "random" (default) does not impose these restrictions and generates unrestricted random partners (it may therefore generate duplicates and the element itself as exchange partner).

When setting the categories argument and using method = "RANN", exchange partners (i.e., nearest neighbours) will be generated from the same category; method = "restricted_random" will also adhere to categorical constraints induced via categories (i.e. each element only receives exchange partners from the same category as itself); method = "random" cannot incorporate categorical restrictions.

Value

A list of length N. Is usually used as input to the argument exchange_partners in [fast_anticlustering](#). Then, the i'th element of the list contains the indices of the exchange partners that are used for the i'th element.

Examples

```
# Restricted random method generates no duplicates per element and cannot return
# the element itself as exchange partner
generate_exchange_partners(5, N = 10, method = "restricted_random")
# "random" simply randomizes with replacement and without restrictions
# (categorical restrictions are also not possible; is much faster for large data sets)
```

```

generate_exchange_partners(5, N = 10, method = "random")
# May return less than 5 exchange partners if there are not enough members
# of the same category:
generate_exchange_partners(
  5, N = 10,
  method = "restricted_random",
  categories = cbind(schaper2019$room, schaper2019$frequency)
)
# using nearest neighbour search (unlike RANN::nn2, this does not
# return the ID of the element itself as neighbour)
generate_exchange_partners(5, features = schaper2019[, 3:5], method = "RANN")[1:3]
# compare with RANN directly:
RANN::nn2(schaper2019[, 3:5], k = 6)$nn.idx[1:3, ] # note k = 6

```

generate_partitions *Generate all partitions of same cardinality*

Description

Generate all partitions of same cardinality

Usage

```
generate_partitions(N, K, generate_permutations = FALSE)
```

Arguments

N	The total N. K has to be dividble by N.
K	How many partitions
generate_permutations	If TRUE, all permutations are returned, resulting in duplicate partitions.

Details

In principle, anticlustering can be solved to optimality by generating all possible partitions of N items into K groups. The example code below illustrates how to do this. However, this approach only works for small N because the number of partitions grows exponentially with N.

The partition c(1, 2, 2, 1) is the same as the partition c(2, 1, 1, 2) but they correspond to different permutations of the elements [1, 1, 2, 2]. If the argument generate_permutations is TRUE, all permutations are returned. To solve balanced anticlustering exactly, it is sufficient to inspect all partitions while ignoring duplicated permutations.

Value

A list of all partitions (or permutations if generate_permutations is TRUE).

Author(s)

Martin Papenberg <martin.papenberg@hhu.de>

References

Papenberg, M., & Klau, G. W. (2021). Using anticlustering to partition data sets into equivalent parts. *Psychological Methods*, 26(2), 161–174. <https://doi.org/10.1037/met0000301>.

Examples

```
## Generate all partitions to solve k-means anticlustering
## to optimality.

N <- 14
K <- 2
features <- matrix(sample(N * 2, replace = TRUE), ncol = 2)
partitions <- generate_partitions(N, K)
length(partitions) # number of possible partitions

## Create an objective function that takes the partition
## as first argument (then, we can use sapply to compute
## the objective for each partition)
var_obj <- function(clusters, features) {
  variance_objective(features, clusters)
}

all_objectives <- sapply(
  partitions,
  FUN = var_obj,
  features = features
)

## Check out distribution of the objective over all partitions:
hist(all_objectives) # many large, few low objectives
## Get best k-means anticlustering objective:
best_obj <- max(all_objectives)
## It is possible that there are multiple best solutions:
sum(all_objectives == best_obj)
## Select one best partition:
best_anticlustering <- partitions[all_objectives == best_obj][[1]]
## Look at mean for each partition:
by(features, best_anticlustering, function(x) round(colMeans(x), 2))

## Get best k-means clustering objective:
min_obj <- min(all_objectives)
sum(all_objectives == min_obj)
## Select one best partition:
best_clustering <- partitions[all_objectives == min_obj][[1]]

## Plot minimum and maximum objectives:
user_par <- par("mfrow")
```

```

par(mfrow = c(1, 2))
plot_clusters(
  features,
  best_anticlustering,
  illustrate_variance = TRUE,
  main = "Maximum variance"
)
plot_clusters(
  features,
  best_clustering,
  illustrate_variance = TRUE,
  main = "Minimum variance"
)
par(mfrow = user_par)

```

`kplus_anticlustering` *K-plus anticlustering*

Description

Perform anticlustering using the k-plus objective to maximize between-group similarity. This function implements the k-plus anticlustering method described in Papenberg (2024; <doi:10.1111/bmsp.12315>).

Usage

```

kplus_anticlustering(
  x,
  K,
  variance = TRUE,
  skew = FALSE,
  kurtosis = FALSE,
  covariances = FALSE,
  T = NULL,
  standardize = TRUE,
  ...
)

```

Arguments

<code>x</code>	A feature matrix where rows correspond to elements and columns correspond to variables (a single numeric variable can be passed as a vector).
<code>K</code>	How many anticlusters should be created. Alternatively: (a) A vector describing the size of each group, or (b) a vector of length <code>nrow(x)</code> describing how elements are assigned to anticlusters before the optimization starts.
<code>variance</code>	Boolean: Should the k-plus objective include a term to maximize between-group similarity with regard to the variance? (Default = TRUE)

skew	Boolean: Should the k-plus objective include a term to maximize between-group similarity with regard to skewness? (Default = FALSE)
kurtosis	Boolean: Should the k-plus objective include a term to maximize between-group similarity with regard to kurtosis? (Default = FALSE)
covariances	Boolean: Should the k-plus objective include a term to maximize between-group similarity with regard to covariance structure? (Default = FALSE)
T	Optional argument: An integer specifying how many distribution moments should be equalized between groups.
standardize	Boolean. If TRUE, the data is standardized through a call to scale before the optimization starts. Defaults to TRUE. See details.
...	Arguments passed down to anticlustering . All of the arguments are supported except for objective.

Details

This function implements the unweighted sum approach for k-plus anticlustering. Details are given in Papenberg (2024).

The optional argument T denotes the number of distribution moments that are considered in the anticlustering process. For example, T = 4 will lead to similar means, variances, skew and kurtosis. For the first four moments, it is also possible to use the boolean convenience arguments `variance`, `skew` and `kurtosis`; the mean (the first moment) is always included and cannot be "turned off". If the argument T is used, it overrides the arguments `variance`, `skew` and `kurtosis` (corresponding to the second, third and fourth moment), ignoring their values.

The standardization is applied to all original features and the additional k-plus features that are appended to the data set in order to optimize the k-plus criterion. When using standardization, all criteria such as means, variances and skewness receive a comparable weight during the optimization. It is usually recommended not to change the default setting `standardization = TRUE`.

This function can use any arguments that are also possible in [anticlustering](#) (except for 'objective' because the objective optimized here is the k-plus objective; to use a different objective, call [anticlustering](#) directly). Any arguments that are not explicitly changed here (i.e., `standardize = TRUE`) receive the default given in [anticlustering](#) (e.g., `method = "exchange"`.)

References

Papenberg, M. (2024). K-plus Anticlustering: An Improved k-means Criterion for Maximizing Between-Group Similarity. *British Journal of Mathematical and Statistical Psychology*, 77(1), 80–102. <https://doi.org/10.1111/bmsp.12315>

Examples

```
# Generate some data
N <- 24
M <- 4
features <- matrix(rnorm(N * M), ncol = M)
# standard k-plus anticlustering: optimize similarity with regard to mean and variance:
cl <- kplus_anticlustering(features, K = 3, method = "local-maximum")
mean_sd_tab(features, cl)
```

```

# Visualize an anticlustering solution:
plot(features, col = palette()[2:4][c1], pch = c(16:18)[c1])

# Also optimize with regard to skewness and kurtosis
c12 <- kplus_anticlustering(
  features,
  K = 3,
  method = "local-maximum",
  skew = TRUE,
  kurtosis = TRUE
)

# The following two calls are equivalent:
init_clusters <- sample(rep_len(1:3, nrow(features)))
# 1.
x1 <- kplus_anticlustering(
  features,
  K = init_clusters,
  variance = TRUE,
  skew = TRUE
)
# 2.
x2 <- kplus_anticlustering(
  features,
  K = init_clusters,
  T = 3
)
# Verify:
all(x1 == x2)

```

kplus_moment_variables

Compute k-plus variables

Description

Compute k-plus variables

Usage

```
kplus_moment_variables(x, T, standardize = TRUE)
```

Arguments

x	A vector, matrix or data.frame of data points. Rows correspond to elements and columns correspond to features. A vector represents a single feature.
T	The number of distribution moments for which variables are generated.
standardize	Logical, should all columns of the output be standardized (defaults to TRUE).

Details

The k-plus criterion is an extension of the k-means criterion (i.e., the "variance", see [variance_objective](#)). In [kplus_anticlustering](#), equalizing means and variances simultaneously (and possibly additional distribution moments) is accomplished by internally appending new variables to the data input `x`. When using only the variance as additional criterion, the new variables represent the squared difference of each data point to the mean of the respective column. All columns are then included—in addition to the original data—in standard k-means anticlustering. The logic is readily extended towards higher order moments, see Papenberg (2024). This function gives users the possibility to generate k-plus variables themselves, which offers some additional flexibility when conducting k-plus anticlustering.

Value

A matrix containing all columns of `x` and all additional columns of k-plus variables. If `x` has `M` columns, the output matrix has `M * T` columns.

Author(s)

Martin Papenberg <martin.papenberg@hhu.de>

References

Papenberg, M. (2024). K-plus Anticlustering: An Improved k-means Criterion for Maximizing Between-Group Similarity. *British Journal of Mathematical and Statistical Psychology*, 77(1), 80–102. <https://doi.org/10.1111/bmsp.12315>

Examples

```
# Use Schaper data set for example
data(schaper2019)
features <- schaper2019[, 3:6]
K <- 3
N <- nrow(features)

# Some equivalent ways of doing k-plus anticlustering:

init_groups <- sample(rep_len(1:3, N))
table(init_groups)

kplus_groups1 <- anticlustering(
  features,
  K = init_groups,
  objective = "kplus",
  standardize = TRUE,
  method = "local-maximum"
)

kplus_groups2 <- anticlustering(
  kplus_moment_variables(features, T = 2), # standardization included by default
  K = init_groups,
  objective = "variance", # (!)
```

```

    method = "local-maximum"
  )

  # this function uses standardization by default unlike anticlustering():
  kplus_groups3 <- kplus_anticlustering(
    features,
    K = init_groups,
    method = "local-maximum"
  )

  all(kplus_groups1 == kplus_groups2)
  all(kplus_groups1 == kplus_groups3)
  all(kplus_groups2 == kplus_groups3)

```

`mahalanobis_distance` *Compute mahalanobis distance between two vectors*

Description

Compute mahalanobis distance between two vectors

Usage

```
mahalanobis_distance(x, y, cov_mat, lambda = 0)
```

Arguments

<code>x</code>	First vector
<code>y</code>	First vector
<code>cov_mat</code>	covariance matrix between the features in x and y
<code>lambda</code>	Regularization, defaults to 0 (no regularization)

Value

The mahalanobis distance (scalar value)

Examples

```
mahalanobis_distance(iris[1, -5], iris[2, -5], cov(iris[, -5]))
```

matching

*Matching***Description**

Conduct K-partite or unrestricted (minimum distance) matching to find pairs or groups of similar elements. By default, finding matches is based on the Euclidean distance between data points, but a custom dissimilarity measure can also be employed.

Usage

```
matching(
  x,
  p = 2,
  match_between = NULL,
  match_within = NULL,
  match_extreme_first = TRUE,
  target_group = NULL,
  sort_output = TRUE
)
```

Arguments

<code>x</code>	The data input. Can be one of two structures: (1) A feature matrix where rows correspond to elements and columns correspond to variables (a single numeric variable can be passed as a vector). (2) An N x N matrix dissimilarity matrix; can be an object of class <code>dist</code> (e.g., returned by <code>dist</code> or <code>as.dist</code>) or a matrix where the entries of the upper and lower triangular matrix represent pairwise dissimilarities.
<code>p</code>	The size of the groups; the default is 2, in which case the function returns pairs.
<code>match_between</code>	An optional vector, <code>data.frame</code> or matrix representing one or several categorical constraints. If passed, the argument <code>p</code> is ignored and matches are sought between elements of different categories.
<code>match_within</code>	An optional vector, <code>data.frame</code> or matrix representing one or several categorical constraints. If passed, matches are sought between elements of the same category.
<code>match_extreme_first</code>	Logical: Determines if matches are first sought for extreme elements first or for central elements. Defaults to <code>TRUE</code> .
<code>target_group</code>	Currently, the options "none", "smallest" and "diverse" are supported. See Details.
<code>sort_output</code>	Boolean. If <code>TRUE</code> (default), the output clusters are sorted by similarity. See Details.

Details

If the data input `x` is a feature matrix, matching is based on the Euclidean distance between data points. If the argument `x` is a dissimilarity matrix, matching is based on the user-specified dissimilarities. To find matches, the algorithm proceeds by selecting a target element and then searching its nearest neighbours. Critical to the behaviour of the algorithm is the order in which target elements are selected. By default, the most extreme elements are selected first, i.e., elements with the highest distance to the centroid of the data set (see [balanced_clustering](#) that relies on the same algorithm). Set the argument `match_extreme_first` to `FALSE`, to enforce that elements close to the centroid are first selected as targets.

If the argument `match_between` is passed and the groups specified via this argument are of different size, target elements are selected from the smallest group by default (because in this group, all elements can be matched). However, it is also possible to specify how matches are selected through the option `target_group`. When specifying `"none"`, matches are always selected from extreme elements, irregardless of the group sizes (or from central elements first if `match_extreme_first` = `FALSE`). With option `"smallest"`, matches are selected from the smallest group. With option `"diverse"`, matches are selected from the most heterogenous group according to the sum of pairwise distances within groups.

The output is an integer vector encoding which elements have been matched. The grouping numbers are sorted by similarity. That is, elements with the grouping number `»1«` have the highest intra-group similarity, followed by 2 etc (groups having the same similarity index are still assigned a different grouping number, though). Similarity is measured as the sum of pairwise (Euclidean) distances within groups (see [diversity_objective](#)). To prevent sorting by similarity (this is some extra computational burden), set `sort_output` = `FALSE`. Some unmatched elements may be `NA`. This happens if it is not possible to evenly split the item pool evenly into groups of size `p` or if the categories described by the argument `match_between` are of different size.

Value

An integer vector encoding the matches. See Details for more information.

Note

It is possible to specify grouping restrictions via `match_between` and `match_within` at the same time.

Author(s)

Martin Papenberg <martin.papenberg@hhu.de>

Examples

```
# Find triplets
N <- 120
lds <- data.frame(f1 = rnorm(N), f2 = rnorm(N))
triplets <- matching(lds, p = 3)
plot_clusters(
  lds,
  clusters = triplets,
  within_connection = TRUE
```

```

)

# Bipartite matching with unequal-sized groups:
# Only selects matches for some elements
N <- 100
data <- matrix(rnorm(N), ncol = 1)
groups <- sample(1:2, size = N, replace = TRUE, prob = c(0.8, 0.2))
matched <- matching(data[, 1], match_between = groups)
plot_clusters(
  cbind(groups, data),
  clusters = matched,
  within_connection = TRUE
)

# Match objects from the same category only
matched <- matching(
  schaper2019[, 3:6],
  p = 3,
  match_within = schaper2019$room
)
head(table(matched, schaper2019$room))

# Match between different plant species in the »iris« data set
species <- iris$Species != "versicolor"
matched <- matching(
  iris[species, 1],
  match_between = iris[species, 5]
)
# Adjust `match_extreme_first` argument
matched2 <- matching(
  iris[species, 1],
  match_between = iris[species, 5],
  match_extreme_first = FALSE
)
# Plot the matching results
user_par <- par("mfrow")
par(mfrow = c(1, 2))
data <- data.frame(
  Species = as.numeric(iris[species, 5]),
  Sepal.Length = iris[species, 1]
)
plot_clusters(
  data,
  clusters = matched,
  within_connection = TRUE,
  main = "Extreme elements matched first"
)
plot_clusters(
  data,
  clusters = matched2,
  within_connection = TRUE,
  main = "Central elements matched first"
)

```

```
par(mfrow = user_par)
```

mean_sd_tab	<i>Means and standard deviations by group variable formatted in table</i>
-------------	---

Description

Means and standard deviations by group variable formatted in table

Usage

```
mean_sd_tab(features, groups, decimals = 2, na.rm = FALSE, return_diff = FALSE)
```

Arguments

features	A data frame of features
groups	A grouping vector
decimals	The number of decimals
na.rm	Should NAs be removed prior to computing stats (Default = FALSE)
return_diff	Boolean. Should an additional row be printed that contains the difference between minimum and maximum

Value

A table that illustrates means and standard deviations (in brackets)

Author(s)

Martin Papenberg <martin.papenberg@hhu.de>

Examples

```
data(iris)
mean_sd_tab(iris[, -5], iris[, 5])
```

n_partitions	<i>Number of equal sized partitions</i>
--------------	---

Description

Number of equal sized partitions

Usage

```
n_partitions(N, K)
```

Arguments

N	How many elements
K	How many partitions

Value

The number of partitions

Examples

```
n_partitions(20, 2)
```

optimal_anticlustering	<i>Optimal ("exact") algorithms for anticlustering</i>
------------------------	--

Description

Wrapper function that gives access to all optimal algorithms for anticlustering that are available in anticlust.

Usage

```
optimal_anticlustering(x, K, objective, solver = NULL, time_limit = NULL)
```

Arguments

x	The data input. Can be one of two structures: (1) A feature matrix where rows correspond to elements and columns correspond to variables (a single numeric variable can be passed as a vector). (2) An N x N matrix dissimilarity matrix; can be an object of class <code>dist</code> (e.g., returned by <code>dist</code> or <code>as.dist</code>) or a matrix where the entries of the upper and lower triangular matrix represent pairwise dissimilarities.
K	How many anticlusters should be created or alternatively: (a) A vector describing the size of each group (the latter currently only works for objective = "dispersion").
objective	The anticlustering objective, can be "diversity", "variance", "kplus" or "dispersion".
solver	Optional. The solver used to obtain the optimal method. Currently supports "glpk", "symphony", "lpSolve" and "gurobi". See details.
time_limit	Time limit in seconds, given to the solver. Default is there is no time limit.

Details

This is a wrapper for all optimal methods supported in `antyclust` (currently and in the future). As compared to `antyclustering`, it allows to specify the solver to obtain an optimal solution, a time limit, and it can be used to obtain optimal solutions for all supported anticlustering objectives (variance, diversity, k-plus, dispersion). For the objectives "variance", "diversity" and "kplus", the optimal ILP method in Papenberg and Klau (2021) is used, which maximizes the sum of all pairwise intra-cluster distances (given user specified number of clusters, for equal-sized clusters). To employ k-means anticlustering (i.e. set objective = "variance"), the squared Euclidean distance is used. For k-plus anticlustering, the squared Euclidean distance based on the extended k-plus data matrix is used (see `kplus_moment_variables`). For the diversity (and the dispersion), the Euclidean distance is used by default, but any user-defined dissimilarity matrix is possible. The equivalence of the k-means/k-plus objectives and the diversity, which enables using the ILP approach by Papenberg and Klau for the k-means/k-plus objectives, was discussed in Papenberg, Breuer, et al. (2025). Also see the examples.

The dispersion is solved optimal using the algorithm `OptDispF` presented in Papenberg, Breuer, et al. (2025). Also see `optimal_dispersion`.

The optimal methods make use of "solvers" that actually implement the algorithm for finding optimal solutions. The package `antyclust` supports four solvers:

- The default solver `lpSolve` (<<https://sourceforge.net/projects/lpsolve/>>).
- GNU linear programming kit (<<http://www.gnu.org/software/glpk/>>), available from the package `Rglpk` and requested using `solver = "glpk"`. The R package `Rglpk` has to be installed manually if this solver should be used.
- The Symphony solver (<<https://github.com/coin-or/SYMPHONY>>), available from the package `Rsymphony` and requested using `solver = "symphony"`. (The package `Rsymphony` has to be installed manually if this solver should be used).
- The commercial `gurobi` solver, see <https://www.gurobi.com/downloads/>.

In general, the commercial gurobi solver is best at solving exact anticlustering problems. If you have a license to use it, use it! Among the open source alternatives, for the maximum dispersion problem, it seems that the Symphony solver is fastest, while the lpSolve solver seems to be good for maximum diversity. However, note that in general the dispersion can be solved optimally for much larger data sets than the diversity.

If a `time_limit` is set and the function cannot find in the optimal objective in the given time, it will throw an error.

Value

A vector of length `N` that assigns a group (i.e, a number between 1 and `K`) to each input element.

Author(s)

Martin Papenberg <martin.papenberg@hhu.de>

References

Papenberg, M., & Klau, G. W. (2021). Using anticlustering to partition data sets into equivalent parts. *Psychological Methods*, 26(2), 161–174. <https://doi.org/10.1037/met0000301>.

Papenberg, M., Breuer, M., Diekhoff, M., Tran, N. K., & Klau, G. W. (2025). Extending the Bicriterion Approach for Anticlustering: Exact and Hybrid Approaches. *Psychometrika*, 90(5), 1789–1808. <https://doi.org/10.1017/psy.2025.10052>

Examples

```
data <- matrix(rnorm(24), ncol = 2)

# These calls are equivalent for k-means anticlustering:
optimal_anticlustering(data, K = 2, objective = "variance")
optimal_anticlustering(dist(data)^2, K = 2, objective = "diversity")

# These calls are equivalent for k-plus anticlustering:
optimal_anticlustering(data, K = 2, objective = "kplus")
optimal_anticlustering(dist(kplus_moment_variables(data, 2))^2, K = 2, objective = "diversity")
```

optimal_dispersion	<i>Maximize dispersion for K groups</i>
--------------------	---

Description

Maximize dispersion for K groups

Usage

```

optimal_dispersion(
  x,
  K,
  solver = NULL,
  max_dispersion_considered = NULL,
  min_dispersion_considered = NULL,
  npartitions = 1,
  time_limit = NULL
)

```

Arguments

<code>x</code>	The data input. Can be one of two structures: (1) A feature matrix where rows correspond to elements and columns correspond to variables (a single numeric variable can be passed as a vector). (2) An $N \times N$ matrix dissimilarity matrix; can be an object of class <code>dist</code> (e.g., returned by <code>dist</code> or <code>as.dist</code>) or a matrix where the entries of the upper and lower triangular matrix represent pairwise dissimilarities.
<code>K</code>	The number of groups or a vector describing the size of each group.
<code>solver</code>	Optional argument; currently supports "lpSolve", "glpk", "symphony", and "gurobi". See optimal_anticlustering .
<code>max_dispersion_considered</code>	Optional argument used for early stopping. If the dispersion found is equal to or exceeds this value, a solution having the previous best dispersion is returned.
<code>min_dispersion_considered</code>	Optional argument used for speeding up the algorithm computation. If passed, the dispersion is optimized starting from this value instead the global minimum distance.
<code>npartitions</code>	The number of groupings that are returned, each having an optimal dispersion value (defaults to 1).
<code>time_limit</code>	Time limit in seconds, given to the solver. Default is there is no time limit.

Details

The dispersion is the minimum distance between two elements within the same group. This function implements an optimal method to maximize the dispersion. If the data input `x` is a feature matrix and not a dissimilarity matrix, the pairwise Euclidean distance is used. It uses the algorithm `OptDispF` presented Papenberg, Breuer, et al. (2025). It was originally developed as part of Max Diekhoff's Bachelor thesis at the Computer Science Department at Heinrich Heine University Düsseldorf.

To find out which items are not allowed to be grouped in the same cluster for maximum dispersion, the algorithm sequentially builds instances of a graph coloring problem, using an integer linear programming (ILP) representation. It is possible to specify the ILP solver via the argument `solver` (See [optimal_anticlustering](#) for more information on this argument). Optimally solving the maximum dispersion problem is NP-hard for $K > 2$ and therefore computationally infeasible for larger data sets. For $K = 3$ and $K = 4$, it seems that this approach scales up to several 100 elements,

or even > 1000 for $K = 3$ (at least when using the Symphony or gurobi solver). For larger data sets, use the heuristic approaches in [anticlustering](#) or [bicriterion_anticlustering](#). However, note that for $K = 2$, the optimal approach is usually much faster than the heuristics.

In the output, the element edges defines which elements must be in separate clusters in order to achieve maximum dispersion. All elements not listed here can be changed arbitrarily between clusters without reducing the dispersion. If the maximum possible dispersion corresponds to the minimum dispersion in the data set, the output elements edges and groups are set to NULL because all possible groupings have the same value of dispersion. In this case the output element dispersions_considered has length 1.

If a time_limit is set and the function cannot find in the optimal dispersion in the given time, it will throw an error.

Value

A list with four elements:

dispersion	The optimal dispersion
groups	An assignment of elements to groups (vector)
edges	A matrix of 2 columns. Each row contains the indices of elements that had to be investigated to find the dispersion (i.e., each pair of elements cannot be part of the same group in order to achieve maximum dispersion).
dispersions_considered	All distances that were tested until the dispersion was found.

Note

If the SYMPHONY solver is used, an unfortunate "message" is printed to the console when this function terminates:

```
sym_get_col_solution(): No solution has been stored!
```

This message is no reason to worry and instead is a direct result of the algorithm finding the optimal value for the dispersion. Unfortunately, this message is generated in the C code underlying the SYMPHONY library (via the printing function printf), which cannot be prevented in R.

Author(s)

Max Diekhoff

Martin Papenberg <martin.papenberg@hhu.de>

References

- Diekhoff (2023). Maximizing dispersion for anticlustering. Retrieved from https://www.cs.hhu.de/fileadmin/redaktion/Fakultät_Für_Naturwissenschaftliche_Fakultät/Informatik/Algorithmische_Bioinformatik/Bachelor-_Masterarbeiten/2831963_ba_ifo_A
- Papenberg, M., Breuer, M., Diekhoff, M., Tran, N. K., & Klau, G. W. (2025). Extending the Bicriterion Approach for Anticlustering: Exact and Hybrid Approaches. *Psychometrika*, 90(5), 1789–1808. <https://doi.org/10.1017/psy.2025.10052>

See Also

[dispersion_objective](#) [anticlustering](#)

Examples

```

N <- 30
M <- 5
K <- 3
data <- matrix(rnorm(N*M), ncol = M)
distances <- dist(data)

opt <- optimal_dispersion(distances, K = K)
opt

# Compare to bicriterion heuristic:
groups_heuristic <- anticlustering(
  distances,
  K = K,
  method = "brusco",
  objective = "dispersion",
  repetitions = 100
)
c(
  OPT = dispersion_objective(distances, opt$groups),
  HEURISTIC = dispersion_objective(distances, groups_heuristic)
)

# Different group sizes are possible:
table(optimal_dispersion(distances, K = c(15, 10, 5))$groups)

# Induce cannot-link constraints by maximizing the dispersion:
solvable <- matrix(1, ncol = 6, nrow = 6)
solvable[2, 1] <- -1
solvable[3, 1] <- -1
solvable[4, 1] <- -1
solvable <- as.dist(solvable)
solvable

# An optimal solution has to put item 1 in a different group than
# items 2, 3 and 4 -> this is possible for K = 2
optimal_dispersion(solvable, K = 2)$groups

# It no longer works when item 1 can also not be linked with item 5:
# (check out output!)
unsolvable <- as.matrix(solvable)
unsolvable[5, 1] <- -1
unsolvable <- as.dist(unsolvable)
unsolvable
optimal_dispersion(unsolvable, K = 2)
# But:
optimal_dispersion(unsolvable, K = c(2, 4)) # group sizes, not number of groups

```

plot_clusters

*Visualize a cluster analysis***Description**

Visualize a cluster analysis

Usage

```
plot_clusters(
  features,
  clusters,
  within_connection = FALSE,
  between_connection = FALSE,
  illustrate_variance = FALSE,
  show_axes = FALSE,
  xlab = NULL,
  ylab = NULL,
  xlim = NULL,
  ylim = NULL,
  main = "",
  cex = 1.2,
  cex.axis = 1.2,
  cex.lab = 1.2,
  lwd = 1.5,
  lty = 2,
  frame.plot = FALSE,
  cex_centroid = 2
)
```

Arguments

features	A data.frame or matrix representing the features that are plotted. Must have two columns.
clusters	A vector representing the clustering
within_connection	Boolean. Connect the elements within each clusters through lines? Useful to illustrate a graph structure.
between_connection	Boolean. Connect the elements between each clusters through lines? Useful to illustrate a graph structure. (This argument only works for two clusters).
illustrate_variance	Boolean. Illustrate the variance criterion in the plot?
show_axes	Boolean, display values on the x and y-axis? Defaults to 'FALSE'.
xlab	The label for the x-axis

ylab	The label for the y-axis
xlim	The limits for the x-axis
ylim	The limits for the y-axis
main	The title of the plot
cex	The size of the plotting symbols, see par
cex.axis	The size of the values on the axes
cex.lab	The size of the labels of the axes
lwd	The width of the lines connecting elements.
lty	The line type of the lines connecting elements (see par).
frame.plot	a logical indicating whether a box should be drawn around the plot.
cex_centroid	The size of the cluster center symbol (has an effect only if <code>illustrate_variance</code> is TRUE)

Details

In most cases, the argument `clusters` is a vector returned by one of the functions [anticlustering](#), [balanced_clustering](#) or [matching](#). However, the plotting function can also be used to plot the results of other cluster functions such as [kmeans](#). This function is usually just used to get a fast impression of the results of an (anti)clustering assignment, but limited in its functionality. It is useful for depicting the intra-cluster connections using argument `within_connection`.

Author(s)

Martin Papenberg <martin.papenberg@hhu.de>

Examples

```
N <- 15
features <- matrix(runif(N * 2), ncol = 2)
K <- 3
clusters <- balanced_clustering(features, K = K)
anticlusters <- anticlustering(features, K = K)
user_par <- par("mfrow")
par(mfrow = c(1, 2))
plot_clusters(features, clusters, main = "Cluster editing", within_connection = TRUE)
plot_clusters(features, anticlusters, main = "Anticlustering editing", within_connection = TRUE)
par(mfrow = user_par)
```

plot_similarity	<i>Plot similarity objective by cluster</i>
-----------------	---

Description

Plot similarity objective by cluster

Usage

```
plot_similarity(x, groups)
```

Arguments

x	The data input. Can be one of two structures: (1) A data matrix where rows correspond to elements and columns correspond to features (a single numeric feature can be passed as a vector). (2) An N x N matrix dissimilarity matrix; can be an object of class dist (e.g., returned by dist or as.dist) or a matrix where the entries of the upper and lower triangular matrix represent the pairwise dissimilarities.
groups	A grouping vector of length N, usually the output of matching .

Details

Plots the sum of pairwise distances by group.

Value

The diversity (sum of distances) by group.

Author(s)

Martin Papenberg <martin.papenberg@hhu.de>

See Also

[diversity_objective](#)

Examples

```
# Match elements and plot similarity by match
N <- 100
lds <- data.frame(f1 = rnorm(N), f2 = rnorm(N))
pairs <- matching(lds, p = 2)
plot_similarity(lds, pairs)
```

schaper2019

*Ratings for 96 words***Description**

A stimulus set that was used in experiments by Schaper, Kuhlmann and Bayen (2019a; 2019b). The item pool consists of 96 German words. Each word represents an object that is either typically found in a bathroom or in a kitchen.

Usage

schaper2019

Format

A data frame with 96 rows and 7 variables

item The name of an object (in German)

room The room in which the item is typically found; can be 'kitchen' or 'bathroom'

rating_consistent How expected would it be to find the item in the typical room

rating_inconsistent How expected would it be to find the item in the atypical room

syllables The number of syllables in the object name

frequency A value indicating the relative frequency of the object name in German language (lower values indicate higher frequency)

list Represents the set affiliation of the item as realized in experiments by Schaper et al.

Source

Courteously provided by Marie Lusia Schaper and Ute Bayen.

References

Schaper, M. L., Kuhlmann, B. G., & Bayen, U. J. (2019a). Metacognitive expectancy effects in source monitoring: Beliefs, in-the-moment experiences, or both? *Journal of Memory and Language*, 107, 95–110. <https://doi.org/10.1016/j.jml.2019.03.009>

Schaper, M. L., Kuhlmann, B. G., & Bayen, U. J. (2019b). Metamemory expectancy illusion and schema-consistent guessing in source monitoring. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 45, 470. <https://doi.org/10.1037/xlm0000602>

Examples

```
head(schaper2019)
features <- schaper2019[, 3:6]

# Optimize the variance criterion
# (tends to maximize similarity in feature means)
```

```

anticlusters <- anticlustering(
  features,
  K = 3,
  objective = "variance",
  categories = schaper2019$room,
  method = "exchange"
)

# Means are quite similar across sets:
by(features, anticlusters, function(x) round(colMeans(x), 2))
# Check differences in standard deviations:
by(features, anticlusters, function(x) round(apply(x, 2, sd), 2))
# Room is balanced between the three sets:
table(Room = schaper2019$room, Set = anticlusters)

# Maximize the diversity criterion
ac_dist <- anticlustering(
  features,
  K = 3,
  objective = "diversity",
  categories = schaper2019$room,
  method = "exchange"
)
# With the distance criterion, means tend to be less similar,
# but standard deviations tend to be more similar:
by(features, ac_dist, function(x) round(colMeans(x), 2))
by(features, ac_dist, function(x) round(apply(x, 2, sd), 2))

```

```
three_phase_search_anticlustering
```

Three phase search with dynamic population size heuristic

Description

This function implements the three phase search algorithm TPSPD for anticlustering by Yang et al. (2022; <doi.org/10.1016/j.ejor.2022.02.003>). The description of their algorithm is given in Section 2 of their paper (in particular, see the Pseudocode in Algorithm 1).

Usage

```

three_phase_search_anticlustering(
  x,
  K,
  N,
  objective = "diversity",
  number_iterations = 50,
  clusters = NULL,

```

```

    upper_bound = NULL,
    lower_bound = NULL,
    beta_max = 15,
    theta_max = NULL,
    theta_min = NULL,
    beta_min = NULL,
    eta_max = 3,
    alpha = 0.05
)

```

Arguments

x	The data input, as in anticlustering .
K	Number of anticlusters to be formed.
N	Number of elements.
objective	The anticlustering objective, can be "diversity" or "dispersion".
number_iterations	A number that defines how many times the steps in the search algorithm are repeated.
clusters	A vector of length K that specifies the number of elements each cluster can contain. If this vector is not NULL, the lower and upper bounds will be disregarded.
upper_bound	Maximum number of elements in each anticluster. By default, anticlusters are of equal size, calculated as the total number of items divided by the number of clusters.
lower_bound	Minimum number of elements in each anticluster. By default, anticlusters are of equal size, calculated as the total number of items divided by the number of clusters.
beta_max	The algorithm begins with a pool of random initial solutions of size beta_max. Over time, the size of the solution pool decreases linearly until it reaches beta_min.
theta_max	Parameter for the strength of undirected perturbation, which decreases linearly over time from theta_max to theta_min.
theta_min	Parameter for the strength of undirected perturbation, which decreases linearly over time from theta_max to theta_min.
beta_min	The minimum solution pool size the algorithm should reach before making a determination.
eta_max	Parameter that specifies how many times the steps in the direct perturbation are executed.
alpha	Parameter for weighing the discrimination of a slightly worse local optimal child solution.

Details

Details of the implementation of the algorithm can be found in the pseudocode of the paper Yang et al. (2022). However, we performed one change as compared to the original description of the algorithm: Instead of setting a time limit, we define the number of iterations the algorithm performs before terminating (via argument `number_iterations`).

Value

A vector of length N that assigns a group (i.e, a number between 1 and K) to each input element

Author(s)

Hannah Hengelbrock <Hannah.Hengelbrock@hhu.de>, Martin Papenberg <martin.papenberg@hhu.de>

References

Xiao Yang et al. “A three-phase search approach with dynamic population size for solving the maximally diverse grouping problem”. In: European Journal of Operational Research 302.3 (2022) <doi:10.1016/j.ejor.2022.02.003>

Examples

```
# Generate some random data
N <- 120
M <- 5
K <- 4
dat <- matrix(rnorm(N * M), ncol = M)
distances <- dist(dat)

# Perform three phase search algorithm
result1 <- three_phase_search_anticlustering(dat, K, N)

# Compute objectives function
diversity_objective(distances, result1)

# Standard algorithm:
result2 <- anticlustering(distances, K=K, method="local-maximum", repetitions = 10)
diversity_objective(distances, result2)
```

variance_objective	<i>Objective value for the variance criterion</i>
--------------------	---

Description

Compute the k-means variance objective for a given clustering.

Usage

```
variance_objective(x, clusters)
```

Arguments

x	A vector, matrix or data.frame of data points. Rows correspond to elements and columns correspond to features. A vector represents a single feature.
clusters	A vector representing (anti)clusters (e.g., returned by anticlustering or balanced_clustering)

Details

The variance objective is given by the sum of the squared errors between cluster centers and individual data points. It is the objective function used in k-means clustering, see [kmeans](#).

Value

The total within-cluster variance

Author(s)

Martin Papenberg <martin.papenberg@hhu.de>

References

- Jain, A. K. (2010). Data clustering: 50 years beyond k-means. *Pattern Recognition Letters*, 31, 651–666.
- Papenberg, M., & Klau, G. W. (2021). Using anticlustering to partition data sets into equivalent parts. *Psychological Methods*, 26(2), 161–174. <https://doi.org/10.1037/met0000301>.
- Späth, H. (1986). Anticlustering: Maximizing the variance criterion. *Control and Cybernetics*, 15, 213–218.

Examples

```
data(iris)
## Clustering
clusters <- balanced_clustering(
  iris[, -5],
  K = 3
)
# This is low:
variance_objective(
  iris[, -5],
  clusters
)
## Anticlustering
anticlusters <- anticlustering(
  iris[, -5],
  K = 3,
  objective = "variance"
)
# This is higher:
variance_objective(
  iris[, -5],
  anticlusters
)

# Illustrate variance objective
N <- 18
data <- matrix(rnorm(N * 2), ncol = 2)
cl <- balanced_clustering(data, K = 3)
plot_clusters(data, cl, illustrate_variance = TRUE)
```

wce*Exact weighted cluster editing*

Description

Optimally solves weighted cluster editing (also known as »correlation clustering« or »clique partitioning problem«).

Usage

```
wce(x, solver = NULL)
```

Arguments

x	A N x N similarity matrix. Larger values indicate stronger agreement / similarity between a pair of data points
solver	Optional argument; if passed, has to be either "glpk" or "symphony". See details.

Details

Finds the clustering that maximizes the sum of pairwise similarities within clusters. In the input some similarities should be negative (indicating dissimilarity) because otherwise the maximum sum of similarities is obtained by simply joining all elements within a single big cluster. The function uses a "solver" to optimize the clustering objective. See [optimal_anticlustering](#) for an overview of the solvers that are available.

Value

An integer vector representing the cluster affiliation of each data point

Author(s)

Martin Papenberg <martin.papenberg@hhu.de>

References

- Bansal, N., Blum, A., & Chawla, S. (2004). Correlation clustering. *Machine Learning*, 56, 89–113.
- Böcker, S., & Baumbach, J. (2013). Cluster editing. In *Conference on Computability in Europe* (pp. 33–44).
- Grötschel, M., & Wakabayashi, Y. (1989). A cutting plane algorithm for a clustering problem. *Mathematical Programming*, 45, 59–96.
- Wittkop, T., Emig, D., Lange, S., Rahmann, S., Albrecht, M., Morris, J. H., ..., Baumbach, J. (2010). Partitioning biological data with transitivity clustering. *Nature Methods*, 7, 419–420.

Examples

```
features <- swiss
distances <- dist(scale(swiss))
hist(distances)
# Define agreement as being close enough to each other.
# By defining low agreement as -1 and high agreement as +1, we
# solve *unweighted* cluster editing
agreements <- ifelse(as.matrix(distances) < 3, 1, -1)
clusters <- wce(agreements)
plot(swiss, col = clusters, pch = 19)
```

Index

- * **datasets**
 - brunel2025, [17](#)
 - schaper2019, [54](#)
- anticlustering, [3](#), [14](#), [18](#), [20–23](#), [28](#), [29](#), [37](#),
[46](#), [49](#), [50](#), [52](#), [56](#), [57](#)
- as.dist, [4](#), [6](#), [12](#), [14](#), [21](#), [22](#), [31](#), [41](#), [46](#), [48](#), [53](#)
- as.matrix, [9](#)
- balanced_clustering, [11](#), [42](#), [52](#), [57](#)
- bicriterion_anticlustering, [6](#), [7](#), [13](#), [49](#)
- brunel2025, [17](#)
- categorical_sampling, [18](#)
- categories_to_binary, [8](#), [19](#), [28](#), [29](#)
- dispersion_objective, [9](#), [14](#), [21](#), [50](#)
- dist, [4](#), [6](#), [12](#), [14](#), [21](#), [22](#), [31](#), [41](#), [46](#), [48](#), [53](#)
- diversity_objective, [5](#), [14](#), [22](#), [42](#), [53](#)
- experimental_allocation, [23](#)
- fast_anticlustering, [20](#), [27](#), [33](#)
- feasible_and_infeasible_region_search_anticlustering,
[30](#)
- generate_exchange_partners, [28](#), [29](#), [32](#)
- generate_partitions, [34](#)
- kmeans, [52](#), [58](#)
- kplus_anticlustering, [5](#), [36](#), [39](#)
- kplus_moment_variables, [28](#), [29](#), [38](#), [46](#)
- mahalanobis_distance, [40](#)
- matching, [9](#), [41](#), [52](#), [53](#)
- mean_sd_tab, [44](#)
- model.matrix, [19](#)
- n_partitions, [45](#)
- nn2, [28](#), [33](#)
- optimal_anticlustering, [7](#), [12](#), [45](#), [48](#), [59](#)
- optimal_dispersion, [7](#), [8](#), [46](#), [47](#)
- par, [52](#)
- plot_clusters, [51](#)
- plot_similarity, [53](#)
- scale, [4](#), [37](#)
- schaper2019, [54](#)
- table, [15](#)
- three_phase_search_anticlustering, [55](#)
- variance_objective, [5](#), [29](#), [39](#), [57](#)
- wce, [59](#)