

multtest

October 5, 2010

EBMTP-class

Class "EBMTP", classes and methods for empirical Bayes multiple testing procedure output

Description

An object of class EBMTP is the output of a particular multiple testing procedure, as generated by the function EBMTP. The object has slots for the various data used to make multiple testing decisions, in particular adjusted p-values.

Objects from the Class

Objects can be created by calls of the form

```
new('MTP',  
  statistic = ....., object of class numeric  
  estimate = ....., object of class numeric  
  sampsize = ....., object of class numeric  
  rawp = ....., object of class numeric  
  adjp = ....., object of class numeric  
  reject = ....., object of class matrix  
  rawdist = ....., object of class matrix  
  nulldist = ....., object of class matrix  
  nulldist.type = ....., object of class character  
  marg.null = ....., object of class character  
  marg.par = ....., object of class matrix  
  label = ....., object of class numeric  
  falsepos = ....., object of class matrix  
  truepos = ....., object of class matrix  
  errormat = ....., object of class matrix  
  EB.h0M = ....., object of class numeric  
  prior = ....., object of class numeric  
  prior.type= ....., object of class character  
  lqv = ....., object of class numeric  
  Hsets = ....., object of class matrix  
  index = ....., object of class matrix  
  call = ....., object of class call  
  seed = ....., object of class integer  
)
```

Slots

- `statistic` Object of class `numeric`, observed test statistics for each hypothesis, specified by the values of the MTP arguments `test`, `robust`, `standardize`, and `psi0`.
- `estimate` For the test of single-parameter null hypotheses using t-statistics (i.e., not the F-tests), the numeric vector of estimated parameters corresponding to each hypothesis, e.g. means, differences in means, regression parameters.
- `sampsize` Object of class `numeric`, number of columns (i.e. observations) in the input data set.
- `rawp` Object of class `numeric`, unadjusted, marginal p-values for each hypothesis.
- `adjp` Object of class `numeric`, adjusted (for multiple testing) p-values for each hypothesis (computed only if the `get.adj` argument is TRUE).
- `reject` Object of class `'matrix'`, rejection indicators (TRUE for a rejected null hypothesis), for each value of the nominal Type I error rate `alpha`.
- `rawdlist` The numeric matrix for the estimated nonparametric non-null test statistics distribution (returned only if `keep.rawdlist=TRUE` and if `nullldist` is one of `'boot.ctr'`, `'boot.cs'`, or `'boot.qt'`). This slot must not be empty if one wishes to call `update` to change choice of bootstrap-based null distribution.
- `nullldist` The numeric matrix for the estimated test statistics null distribution (returned only if `keep.nullldist=TRUE`). By default (i.e., for `nullldist='boot.cs'`), the entries of `nullldist` are the null value shifted and scaled bootstrap test statistics, with one null test statistic value for each hypothesis (rows) and bootstrap iteration (columns).
- `nullldist.type` Character value describing which choice of null distribution was used to generate the MTP results. Takes on one of the values of the original `nullldist` argument in the call to MTP, i.e., `'boot.cs'`, `'boot.ctr'`, `'boot.qt'`, or `'ic'`.
- `marg.null` If `nullldist='boot.qt'`, a character value returning which choice of marginal null distribution was used by the MTP. Can be used to check default values or to ensure manual settings were correctly applied.
- `marg.par` If `nullldist='boot.qt'`, a numeric matrix returning the parameters of the marginal null distribution(s) used by the MTP. Can be used to check default values or to ensure manual settings were correctly applied.
- `falsepos` A matrix with rows equal to the number of hypotheses and columns the number of samples of null test statistics (B) indicating the number of guessed false positives when using the corresponding value of the observed test statistic as a cut-off. Not returned unless `keep.falsepos=TRUE`.
- `truepos` A matrix with rows equal to the number of hypotheses and columns the number of samples of null test statistics (B) indicating the number of guessed true positives when using the corresponding value of the observed test statistic as a cut-off. Not returned unless `keep.truepos=TRUE`.
- `errormat` The matrix obtained after applying to type I error rate function closure to the matrices in `falsepos`, and, if applicable, `truepos`. Not returned unless `keep.errormat=TRUE`.
- `EB.h0M` The sum of the local q-values obtained after density estimation. This number serves as an estimate of the proportion of true null hypotheses. Values close to one indicate situations in which type I error control may not be guaranteed by the EBMTP. When `prior='EBLQV'`, this value is used as the prior 'pi' during evaluation of the local q-value function.
- `prior` The numeric value of the prior 'pi' used when evaluating the local q-value function.
- `prior.type` Character string returning the value of `prior` in the original call to EBMTP. One of `'conservative'`, `'ABH'`, or `'EBLQV'`.

- `lqv` A numeric vector of length the number of hypotheses with the estimated local q-values used for generating guessed sets of true null hypotheses.
- `Hsets` A numeric matrix with the same dimension as `nulldist`, containing the Bernoulli realizations of the estimated local q-values stored in `lqv` which were used to partition the hypotheses into guessed sets of true and false null hypotheses at each round of (re)sampling. Not returned unless `keep.Hsets=TRUE`.
- `label` If `keep.label=TRUE`, a vector storing the values used in the argument `Y`. Storing this object is particularly important when one wishes to update EBMTP objects with F-statistics using default `marg.null` and `marg.par` settings when `nulldist='boot.qt'`.
- `index` For tests of correlation parameters a matrix corresponding to `t(combn(p, 2))`, where `p` is the number of variables in `X`. This matrix gives the indices of the variables considered in each pairwise correlation. For all other tests, this slot is empty, as the indices are in the same order as the rows of `X`.
- `call` Object of class `call`, the call to the MTP function.
- `seed` An integer or vector for specifying the state of the random number generator used to create the resampled datasets. The seed can be reused for reproducibility in a repeat call to MTP. This argument is currently used only for the bootstrap null distribution (i.e., for `nulldist="boot.xx"`). See `?set.seed` for details.

Methods

`signature(x = "EBMTP")`

`[` : Subsetting method for EBMTP class, which operates selectively on each slot of an EBMTP instance to retain only the data related to the specified hypotheses.

`as.list` : Converts an object of class EBMTP to an object of class `list`, with an entry for each slot.

`plot` : plot methods for EBMTP class, produces the following graphical summaries of the results of a EBMTP. The type of display may be specified via the `which` argument.

1. Scatterplot of number of rejected hypotheses vs. nominal Type I error rate.
2. Plot of ordered adjusted p-values; can be viewed as a plot of Type I error rate vs. number of rejected hypotheses.
3. Scatterplot of adjusted p-values vs. test statistics (also known as "volcano plot").
4. Plot of unordered adjusted p-values.

The plot method for objects of class EBMTP does not return the plots associated with `which=5` (using confidence regions) or with `which=6` (pertaining to cut-offs) as it does for objects of class MTP. This is because the function EBMTP currently only returns adjusted p-values. The argument `logscale` (by default equal to FALSE) allows one to use the negative decimal logarithms of the adjusted p-values in the second, third, and fourth graphical displays. The arguments `caption` and `sub.caption` allow one to change the titles and subtitles for each of the plots (default subtitle is the MTP function call). Note that some of these plots are implemented in the older function `mt.plot`.

`print` : print method for EBMTP class, returns a description of an object of class EBMTP, including sample size, number of tested hypotheses, type of test performed (value of argument `test`),

Type I error rate (value of argument `typeone`), nominal level of the test (value of argument `alpha`), name of the EBMTP (value of argument `method`), call to the function `EBMTP`.

In addition, this method produces a table with the class, mode, length, and dimension of each slot of the `EBMTP` instance.

summary : summary method for `EBMTP` class, provides numerical summaries of the results of an `EBMTP` and returns a list with the following three components.

1. `rejections`: A data.frame with the number(s) of rejected hypotheses for the nominal Type I error rate(s) specified by the `alpha` argument of the function `MTP`.

2. `index`: A numeric vector of indices for ordering the hypotheses according to first `adjp`, then `rawp`, and finally the absolute value of `statistic` (not printed in the summary).

3. `summaries`: When applicable (i.e., when the corresponding quantities are returned by `MTP`), a table with six number summaries of the distributions of the adjusted p-values, unadjusted p-values, test statistics, and parameter estimates.

EBupdate : update method for `EBMTP` class, provides a mechanism to re-run the `MTP` with different choices of the following arguments - `nulldist`, `alternative`, `typeone`, `k`, `q`, `alpha`, `smooth.null`, `bw`, `kernel`, `prior`, `keep.nulldist`, `keep.rawdist`, `keep.falsepos`, `keep.truepos`, `keep.errormat`, `keep.margpar`. When `evaluate` is 'TRUE', a new object of class `EBMTP` is returned. Else, the updated call is returned. The `EBMTP` object passed to the update method must have either a non-empty `rawdist` slot or a non-empty `nulldist` slot (i.e., must have been called with either 'keep.rawdist=TRUE' or 'keep.nulldist=TRUE').

Additionally, when calling `EBupdate` for any Type I error rate other than FWER, the `typeone` argument must be specified (even if the original object did not control FWER). For example, `typeone="fdr"`, would always have to be specified, even if the original object also controlled the FDR. In other words, for all function arguments, it is safest to always assume that you are updating from the `EBMTP` default function settings, regardless of the original call to the `EBMTP` function. Currently, the main advantage of the `EBupdate` method is that it prevents the need for repeated estimation of the test statistics null distribution.

To save on memory, if one knows ahead of time that one will want to compare different choices of bootstrap-based null distribution, then it is both necessary and sufficient to specify 'keep.rawdist=TRUE', as there is no other means of moving between null distributions other than through the non-transformed non-parametric bootstrap distribution. In this case, 'keep.nulldist=FALSE' may be used. Specifically, if an object of class `EBMTP` contains a non-empty `rawdist` slot and an empty `nulldist` slot, then a new null distribution will be generated according to the values of the `nulldist=` argument in the original call to `EBMTP` or any additional specifications in the call to `update`. On the other hand, if one knows that one wishes to only update an `EBMTP` object in ways which do not involve choice of null distribution, then 'keep.nulldist=TRUE' will suffice and 'keep.rawdist' can be set to `FALSE` (default settings). The original null distribution object will then be used for all subsequent calls to `update`.

N.B.: Note that `keep.rawdist=TRUE` is only available for the bootstrap-based resampling methods. The non-null distribution does not exist for the permutation or influence curve multivariate normal null distributions.

ebmtp2mtp : coercion method for converting objects of class `EBMTP` to objects of class `MTP`. Slots common to both objects are taken from the object of class `EBMTP` and used to create a new object of class `MTP`. Once an object of class `MTP` is created, one may use the method `update`

to perform resampling-based multiple testing (as would have been done with calls to MTP) without the need for repeated resampling.

Author(s)

Houston N. Gilbert, based on the original MTP class and method definitions written by Katherine S. Pollard

References

H.N. Gilbert, K.S. Pollard, M.J. van der Laan, and S. Dudoit (2009). Resampling-based multiple hypothesis testing with applications to genomics: New developments in R/Bioconductor package multtest. *Journal of Statistical Software* (submitted). Temporary URL: <http://www.stat.berkeley.edu/~houston/JSSNullDistEBMTP.pdf>.

Y. Benjamini and Y. Hochberg (2000). On the adaptive control of the false discovery rate in multiple testing with independent statistics. *J. Behav. Educ. Statist.* Vol 25: 60-83.

Y. Benjamini, A. M. Krieger and D. Yekutieli (2006). Adaptive linear step-up procedures that control the false discovery rate. *Biometrika*. Vol. 93: 491-507.

M.J. van der Laan, M.D. Birkner, and A.E. Hubbard (2005). Empirical Bayes and Resampling Based Multiple Testing Procedure Controlling the Tail Probability of the Proportion of False Positives. *Statistical Applications in Genetics and Molecular Biology*, 4(1). <http://www.bepress.com/sagmb/vol4/iss1/art29/>

S. Dudoit and M.J. van der Laan. *Multiple Testing Procedures and Applications to Genomics*. Springer Series in Statistics. Springer, New York, 2008.

S. Dudoit, H. N. Gilbert, and M. J. van der Laan (2008). Resampling-based empirical Bayes multiple testing procedures for controlling generalized tail probability and expected value error rates: Focus on the false discovery rate and simulation study. *Biometrical Journal*, 50(5):716-44. <http://www.stat.berkeley.edu/~houston/BJMCPSupp/BJMCPSupp.html>.

H.N. Gilbert, M.J. van der Laan, and S. Dudoit. Joint multiple testing procedures for graphical model selection with applications to biological networks. Technical report, U.C. Berkeley Division of Biostatistics Working Paper Series, April 2009. URL <http://www.bepress.com/ucbbiostat/paper245>.

See Also

[EBMTP](#), [EBMTP-methods](#), [MTP](#), [MTP-methods](#), [\[-methods](#), [as.list-methods](#), [print-methods](#), [plot-methods](#), [summary-methods](#), [mtp2ebmtp](#), [ebmtp2mtp](#)

Examples

```
## See EBMTP function: ? EBMTP
```

EBMTP

A function to perform empirical Bayes resampling-based multiple hypothesis testing

Description

A user-level function to perform empirical Bayes multiple testing procedures (EBMTP). A variety of t- and F-tests, including robust versions of most tests, are implemented. A common-cutoff method is used to control the chosen type I error rate (FWER, gFWER, TPPFP, or FDR). Bootstrap-based null distributions are available. Additionally, for t-statistics, one may wish to sample from an appropriate multivariate normal distribution with mean zero and correlation matrix derived from the vector influence function. In EBMTP, realizations of local q-values, obtained via density estimation, are used to partition null and observed test statistics into guessed sets of true and false null hypotheses at each round of (re)sampling. In this manner, parameters of any type I error rate which can be expressed as a function the number of false positives and true positives can be estimated. Arguments are provided for user control of output. Gene selection in microarray experiments is one application.

Usage

```
EBMTP(X, W = NULL, Y = NULL, Z = NULL, Z.incl = NULL, Z.test = NULL,
      na.rm = TRUE, test = "t.twosamp.unequalvar", robust = FALSE,
      standardize = TRUE, alternative = "two.sided", typeone = "fwer",
      method = "common.cutoff", k = 0, q = 0.1, alpha = 0.05, smooth.null = FALSE,
      nulldist = "boot.cs", B = 1000, psi0 = 0, marg.null = NULL,
      marg.par = NULL, ncp = NULL, perm.mat = NULL, ic.quant.trans = FALSE,
      MVN.method = "mvrnorm", penalty = 1e-06, prior = "conservative",
      bw = "nrd", kernel = "gaussian", seed = NULL, cluster = 1,
      type = NULL, dispatch = NULL, keep.nulldist = TRUE, keep.rawdist = FALSE,
      keep.falsepos = FALSE, keep.truepos = FALSE, keep.errorformat = FALSE,
      keep.Hsets=FALSE, keep.margpar = TRUE, keep.index = FALSE, keep.label = FALSE)
```

Arguments

For brevity, the presentation of arguments below will highlight those which differ significantly from arguments in the other main-level user function `MTP`. See [MTP](#) for further details.

Character string indicating which type I error rate to control, by default family-wise error rate ('fwer'). Other options include generalized family-wise error rate ('gfwere'), with parameter `k` giving the allowed number of false positives, and tail probability of the proportion of false positives ('tppfp'), with parameter `q` giving the allowed proportion of false positives. The false discovery rate ('fdr') can also be controlled. In particular, for 'gfwere', 'tppfp' and 'fdr', multiple testing is not performed via augmentation of the results of a FWER-controlling `MTP`. Rather, using guessed sets of true and false null hypotheses, these error rates are controlled in a more direct manner.

<code>method</code>	Character string indicating the EBMTP method. Currently only 'common.cutoff' is implemented. This method is most similar to 'ss.maxT' in <code>MTP</code> .
<code>nulldist</code>	Character string indicating which resampling method to use for estimating the joint test statistics null distribution, by default the non-parametric bootstrap

with centering and scaling ('boot.cs'). The old default 'boot' will still compile and will correspond to 'boot.cs'. Other null distribution options include 'boot.ctr', 'boot.qt', and 'ic', corresponding to the centered-only bootstrap distribution, quantile-transformed bootstrap distribution, and influence curve multivariate normal joint null distribution, respectively. The permutation distribution is not available.

`prior` Character string indicating which choice of prior probability to use for estimating local q-values (i.e., the posterior probabilities of a null hypothesis being true given the value of its corresponding test statistic). Default is 'conservative', in which case the prior is set to its most conservative value of 1, meaning that all hypotheses are assumed to belong to the set of true null hypotheses. Other options include 'ABH' for the adaptive Benjamini-Hochberg estimator of the number/proportion of true null hypotheses, and 'EBLQV' for the empirical Bayes local q-value value estimator of the number/proportion of true null hypotheses. If 'EBLQV', the estimator of the prior probability is taken to be the sum of the estimated local q-values divided by the number of tests. Relaxing the prior may result in more rejections, albeit at a cost of type I error control under certain conditions. See details and references.

`bw` A character string argument to `density` indicating the smoothing bandwidth to be used during kernel density estimation. Default is 'nrd'.

`kernel` A character string argument to `density` specifying the smoothing kernel to be used. Default is 'gaussian'.

`keep.falsepos` A logical indicating whether or not to store the matrix of guessed false positives at each round of (re)sampling. The matrix has rows equal to the number of cut-offs (observed test statistics) and columns equal to the B number of bootstrap samples or samples from the multivariate normal distribution (if `nulldist='ic'`). Default is 'FALSE'.

`keep.truepos` A logical indicating whether or not to store the matrix of guessed true positives at each round of (re)sampling. The matrix has rows equal to the number of cut-offs (observed test statistics) and columns equal to the B number of bootstrap samples or samples from the multivariate normal distribution (if `nulldist='ic'`). Default is 'FALSE'.

`keep.errormat` A logical indicating whether or not to store the matrix of type I error rate values at each round of (re)sampling. The matrix has rows equal to the number of cut-offs (observed test statistics) and columns equal to the B number of bootstrap samples or samples from the multivariate normal distribution (if `nulldist='ic'`). Default is 'FALSE'. In the case of FDR-control, for example, this matrix is `falsepos/(falsepos + truepos)`. The row means of this matrix are eventually used for assigning/ordering adjusted p-values to test statistics of each hypothesis.

`keep.Hsets` A logical indicating whether or not to return the matrix of indicators which partition the hypotheses into guessed sets of true and false null hypotheses at each round of (re)sampling. Default is 'FALSE'.

`X`, `W`, `Y`, `Z`, `Z.incl`, `Z.test`, `na.rm`, `test`, `robust`, `standardize`, `alternative`, `k`, `q`,
 These arguments are all similarly used by the MTP function, and their use has been defined elsewhere. Please consult the `link{MTP}` help file or the references for further details. Note that the MTP-function arguments `get.cr`, `get.cutoff`, `get.adj` are now DEPRECATED in the EBMTP function.

Only adjusted p-values are calculated by EBMTP. These adjusted p-values are returned in the same order as the original hypotheses and raw p-values (typically corresponding to rows of X.)

Details

The EBMTP begins with a marginal nonparametric mixture model for estimating local q-values. By definition, q-values are 'the opposite' of traditional p-values. That is, q-values represent the probability of null hypothesis being true given the value of its corresponding test statistic. If the test statistics T_n have marginal distribution $f = \pi_0 f_0 + (1-\pi_0) f_1$, where π_0 is the prior probability of a true null hypothesis and f_0 and f_1 represent the marginal null and alternative densities, respectively, then the local q-value function is given by $\pi_0 f_0(T_n)/f(T_n)$.

One can estimate both the null density f_0 and full density f by applying kernel density estimation over the matrix of null test statistics and the vector of observed test statistics, respectively. Practically, this step in EBMTP also ensures that sidedness is correctly accounted for among the test statistics and their estimated null distribution. The prior probability π_0 can be set to its most conservative value of 1 or estimated by some other means, e.g., using the adaptive Benjamini Hochberg ('ABH') estimator or by summing up the estimated local q-values themselves ('EBLQV') and dividing by the number of tests. Bounding these estimated probabilities by one provides a vector of estimated local q-values with length equal to the number of hypotheses. Bernoulli 0/1 realizations of the posterior probabilities indicate which hypotheses are guessed as belonging to the true set of null hypotheses given the value of their test statistics. Once this partitioning has been achieved, one can count the numbers of guessed false positives and guessed true positives at each round of (re)sampling that are obtained when using the value of an observed test statistic as a cut-off.

EBMTPs use function closures to represent type I error rates in terms of their defining features. Restricting the choice of type I error rate to 'fwer', 'gfwer', 'tppfp', and 'fdr', means that these features include whether to control the number of false positives or the proportion of false positives among the number of rejections made (i.e., the false discovery proportion), whether we are controlling a tail probability or expected value error rate, and, in the case of tail probability error rates, what bound we are placing on the random variable defining the type I error rate (e.g., k for 'gfwer' or 'q' for 'tppfp'). Averaging the type I error results over B (bootstrap or multivariate normal) samples provides an estimator of the evidence against the null hypothesis (adjusted p-values) with respect to the choice of type I error rate. Finally, a monotonicity constraint is placed on the adjusted p-values before being returned as output.

As detailed in the references, relaxing the prior may result in a more powerful multiple testing procedure, albeit sometimes at the cost of type I error control. Additionally, when the proportion of true null hypotheses is close to one, type I error control may also become an issue, even when using the most conservative prior probability of one. This feature is known to occur with some other procedures which rely on the marginal nonparametric mixture model for estimating (local) q-values. The slot `EB.h0M` returned by objects of class EBMTP is the sum of the local q-values estimated via kernel density estimation (divided by the total number of tests). If this value is close to one ($>0.9-0.95$), the user will probably not want relax the prior, as even the conservative EBMTP might be approaching a performance bound with respect to type I error control. The user is advised to begin by using the most 'conservative' prior, assess the estimated proportion of true null hypotheses, and then decide if relaxing the prior might be desired. Gains in power over other multiple testing procedures have been observed even when using the most conservative prior of one.

Situations of moderate-high to high levels of correlation may also affect the results of multiple testing methods which use the same mixture model for generating q-values. Microarray analysis

represents a scenario in which dependence structures are typically weak enough to mitigate this concern. On the other hand, the analysis of densely sampled SNPs, for example, may present problems.

Value

An object of class EBMTP. Again, for brevity, the values below represent slots which distinguish objects of class EBMTP from those of class MTP.

<code>falsepos</code>	A matrix with rows equal to the number of hypotheses and columns the number of samples of null test statistics (B) indicating the number of guessed false positives when using the corresponding value of the observed test statistic as a cut-off. Not returned unless <code>keep.falsepos=TRUE</code> .
<code>truepos</code>	A matrix with rows equal to the number of hypotheses and columns the number of samples of null test statistics (B) indicating the number of guessed true positives when using the corresponding value of the observed test statistic as a cut-off. Not returned unless <code>keep.truepos=TRUE</code> .
<code>errormat</code>	The matrix obtained after applying to type I error rate function closure to the matrices in <code>falsepos</code> , and, if applicable, <code>truepos</code> . Not returned unless <code>keep.errormat=TRUE</code> .
<code>EB.h0M</code>	The sum of the local q-values obtained after density estimation. This number serves as an estimate of the proportion of true null hypotheses. Values close to one indicate situations in which type I error control may not be guaranteed by the EBMTP. When <code>prior='EBLQV'</code> , this value is used as the prior 'pi' during evaluation of the local q-value function.
<code>prior</code>	The numeric value of the prior 'pi' used when evaluating the local q-value function.
<code>prior.type</code>	Character string returning the value of <code>prior</code> in the original call to EBMTP. One of 'conservative', 'ABH', or 'EBLQV'.
<code>lqv</code>	A numeric vector of length the number of hypotheses with the estimated local q-values used for generating guessed sets of true null hypotheses.
<code>Hsets</code>	A numeric matrix with the same dimension as <code>nulldist</code> , containing the Bernoulli realizations of the estimated local q-values stored in <code>lqv</code> which were used to partition the hypotheses into guessed sets of true and false null hypotheses at each round of (re)sampling. Not returned unless <code>keep.Hsets=TRUE</code> .

Author(s)

Houston N. Gilbert, based on the original MTP code written by Katherine S. Pollard

References

H.N. Gilbert, K.S. Pollard, M.J. van der Laan, and S. Dudoit (2009). Resampling-based multiple hypothesis testing with applications to genomics: New developments in R/Bioconductor package `multtest`. *Journal of Statistical Software* (submitted). Temporary URL: <http://www.stat.berkeley.edu/~houston/JSSNullDistEBMTP.pdf>.

Y. Benjamini and Y. Hochberg (2000). On the adaptive control of the false discovery rate in multiple testing with independent statistics. *J. Behav. Educ. Statist.* Vol 25: 60-83.

Y. Benjamini, A. M. Krieger and D. Yekutieli (2006). Adaptive linear step-up procedures that control the false discovery rate. *Biometrika*. Vol. 93: 491-507.

M.J. van der Laan, M.D. Birkner, and A.E. Hubbard (2005). Empirical Bayes and Resampling Based Multiple Testing Procedure Controlling the Tail Probability of the Proportion of False Positives. *Statistical Applications in Genetics and Molecular Biology*, 4(1). <http://www.bepress.com/sagmb/vol14/iss1/art29/>

S. Dudoit and M.J. van der Laan. *Multiple Testing Procedures and Applications to Genomics*. Springer Series in Statistics. Springer, New York, 2008.

S. Dudoit, H.N. Gilbert, and M J. van der Laan (2008). Resampling-based empirical Bayes multiple testing procedures for controlling generalized tail probability and expected value error rates: Focus on the false discovery rate and simulation study. *Biometrical Journal*, 50(5):716-44. <http://www.stat.berkeley.edu/~houston/BJMCPSupp/BJMCPSupp.html>.

H.N. Gilbert, M.J. van der Laan, and S. Dudoit. Joint multiple testing procedures for graphical model selection with applications to biological networks. Technical report, U.C. Berkeley Division of Biostatistics Working Paper Series, April 2009. URL <http://www.bepress.com/ucbbiostat/paper245>.

See Also

[MTP](#), [EBMTP-class](#), [EBMTP-methods](#), [Hsets](#)

Examples

```
set.seed(99)
data<-matrix(rnorm(90),nr=9)
group<-c(rep(1,5),rep(0,5))

#EB fwer control with centered and scaled bootstrap null distribution
#(B=100 for speed)
eb.m1<-EBMTP(X=data,Y=group,alternative="less",B=100,method="common.cutoff")
print(eb.m1)
summary(eb.m1)
par(mfrow=c(2,2))
plot(eb.m1,top=9)
```

Hsets

Functions for generating guessed sets of true null hypotheses in empirical Bayes resampling-based multiple hypothesis testing

Description

These functions are called internally by the main user-level function `EBMTP`. They are used for estimating local q-values, generating guessed sets of true null hypotheses, and applying these results to function closures defining the choice of type I error rate (FWER, gFWER, TPPFP, and FDR).

Usage

```
Hsets(Tn, nullmat, bw, kernel, prior, B, rawp)
```

```
ABH.h0(rawp)
```

```
G.VS(V, S = NULL, tp = TRUE, bound)
```

Arguments

Tn	The vector of observed test statistics.
nullmat	The matrix of null test statistics obtained either through null transformation of the bootstrap distribution or by sampling from an appropriate multivariate normal distribution (when <code>nulldist='ic'</code> .)
bw	A character string argument to <code>density</code> indicating the smoothing bandwidth to be used during kernel density estimation. Default is <code>'nrd'</code> .
kernel	A character string argument to <code>density</code> specifying the smoothing kernel to be used. Default is <code>'gaussian'</code> .
prior	Character string indicating which choice of prior probability to use for estimating local q-values (i.e., the posterior probabilities of a null hypothesis being true given the value of its corresponding test statistic). Default is <code>'conservative'</code> , in which case the prior is set to its most conservative value of 1, meaning that all hypotheses are assumed to belong to the set of true null hypotheses. Other options include <code>'ABH'</code> for the adaptive Benjamini-Hochberg estimator of the number/proportion of true null hypotheses, and <code>'EBLQV'</code> for the empirical Bayes local q-value value estimator of the number/proportion of true null hypotheses. If <code>'EBLQV'</code> , the estimator of the prior probability is taken to be the sum of the estimated local q-values divided by the number of tests. Relaxing the prior may result in more rejections, albeit at a cost of type I error control under certain conditions. See references.
B	The number of bootstrap iterations (i.e. how many resampled data sets) or the number of samples from the multivariate normal distribution (if <code>nulldist='ic'</code>). Can be reduced to increase the speed of computation, at a cost to precision. Default is 1000.
rawp	A vector of raw (unadjusted) p-values obtained bootstrap-based or influence curve null distribution.
V	A matrix of the numbers of guessed false positives for each cut-off, i.e., observed value of a test statistic, within each sample in B.
S	A matrix of the numbers of guessed true positives for each cut-off, i.e., observed value of a test statistic, within each sample in B.
tp	Logical indicator which is <code>TRUE</code> if type I error rate is a tail probability error rate and <code>FALSE</code> is if it is an expected value error rate.
bound	If a tail probability error rate, the bound to be placed on function of guessed false positives and guessed true positives. For, <code>'fwer'</code> , equal to 0; <code>'gfwer'</code> , equal to <code>'k'</code> ; and <code>tpfpf</code> , equal to <code>'q'</code> .

Details

The most important object to be returned from the function `Hsets` is a matrix of indicators, i.e., Bernoulli realizations of the estimated local q-values, taking the value of 1 if the hypothesis is

guessed as belonging to the set of true null hypotheses and 0 otherwise (guessed true alternative). Realizations of these probabilities are generated with a call to `rbinom`, meaning that this function will set the RNG seed forward another B^* (the number of hypotheses) places. This matrix, with rows equal to the number of hypotheses and columns the number of (bootstrap or multivariate normal) samples is used to subset the matrix of null test statistics and the vector of observed test statistics at each round of (re)sampling into samples of statistics guessed as belonging to the sets of true null and true alternative hypotheses, respectively. Using the values of the observed test statistics themselves as cut-offs, the numbers of guessed false positives and (if applicable) guessed true positives can be counted and eventually used to estimate the distribution of a type I error rate characterized by the closure returned from `G.VS`. Counting of guessed false positives and guessed true positives is performed in `C` through the function `VScount`.

Value

For the function `Hsets`, a list with the following elements:

<code>Hsets.mat</code>	A matrix of numeric indicators with rows equal to the number of test (hypotheses, typically <code>nrow(X)</code>) and columns the number of samples of null test statistics, <code>B</code> . Values of one indicate hypotheses guessed as belonging to the set of true null hypotheses based on the value of their corresponding test statistic. Values of zero correspond to hypotheses guesses as belonging to the set of true alternative hypotheses.
<code>EB.h0M</code>	The estimated proportion of true null hypotheses as determined by nonparametric density estimation. This value is the sum of the estimated local q-values divided by the total number of tests (hypotheses).
<code>prior</code>	The value of the prior applied to the local q-value function. If 'conservative', the prior is set to one. Otherwise, the prior is the value obtained from the estimator of the adaptive Benjamini-Hochberg procedure (if <code>prior</code> is 'ABH') or from density estimation (if <code>prior</code> is 'EBLQV').
<code>pn.out</code>	The vector of estimated local q-values. This vector is returned in the <code>lqv</code> slot of objects of class <code>EBMTP</code> .

For the function `ABH.h0`, the estimated number of true null hypotheses using the estimator from the linear step-up adaptive Benjamini-Hochberg procedure.

For the function `G.VS`, a closure which accepts as arguments the matrices of guessed false positive and true positives (if applicable) and applies the appropriate function defining the desired type I error rate.

Author(s)

Houston N. Gilbert

References

H.N. Gilbert, K.S. Pollard, M.J. van der Laan, and S. Dudoit (2009). Resampling-based multiple hypothesis testing with applications to genomics: New developments in R/Bioconductor package `multtest`. *Journal of Statistical Software* (submitted). Temporary URL: <http://www.stat.berkeley.edu/~houston/JSSNullDistEBMTP.pdf>.

Y. Benjamini and Y. Hochberg (2000). On the adaptive control of the false discovery rate in multiple testing with independent statistics. *J. Behav. Educ. Statist.* Vol 25: 60-83.

Y. Benjamini, A.M. Krieger and D. Yekutieli (2006). Adaptive linear step-up procedures that control the false discovery rate. *Biometrika*. Vol. 93: 491-507.

M.J. van der Laan, M.D. Birkner, and A.E. Hubbard (2005). Empirical Bayes and Resampling Based Multiple Testing Procedure Controlling the Tail Probability of the Proportion of False Positives. *Statistical Applications in Genetics and Molecular Biology*, 4(1). <http://www.bepress.com/sagmb/vol4/iss1/art29/>

S. Dudoit and M.J. van der Laan. *Multiple Testing Procedures and Applications to Genomics*. Springer Series in Statistics. Springer, New York, 2008.

S. Dudoit, H.N. Gilbert, and M.J. van der Laan (2008). Resampling-based empirical Bayes multiple testing procedures for controlling generalized tail probability and expected value error rates: Focus on the false discovery rate and simulation study. *Biometrical Journal*, 50(5):716-44. <http://www.stat.berkeley.edu/~houston/BJMCPSupp/BJMCPSupp.html>.

H.N. Gilbert, M.J. van der Laan, and S. Dudoit. Joint multiple testing procedures for graphical model selection with applications to biological networks. Technical report, U.C. Berkeley Division of Biostatistics Working Paper Series, April 2009. URL <http://www.bepress.com/ucbbiostat/paper245>.

See Also

[EBMTP](#), [EBMTP-class](#), [EBMTP-methods](#)

Examples

```
set.seed(99)
data<-matrix(rnorm(90),nr=9)
group<-c(rep(1,5),rep(0,5))

#EB fwer control with centered and scaled bootstrap null distribution
#(B=100 for speed)
eb.m1<-EBMTP(X=data,Y=group,alternative="less",B=100,method="common.cutoff")
print(eb.m1)
summary(eb.m1)
par(mfrow=c(2,2))
plot(eb.m1,top=9)

abh <- ABH.h0(eb.m1@rawp)
abh

eb.m2 <- EUpdate(eb.m1,prior="ABH")
eb.m2@prior
```

Description

An object of class `MTP` is the output of a particular multiple testing procedure, for example, generated by the `MTP` function. It has slots for the various data used to make multiple testing decisions, such as adjusted p-values and confidence regions.

Objects from the Class

Objects can be created by calls of the form

```
new('MTP',
  statistic = ....., object of class numeric
  estimate = ....., object of class numeric
  sampsize = ....., object of class numeric
  rawp = ....., object of class numeric
  adjp = ....., object of class numeric
  conf.reg = ....., object of class array
  cutoff = ....., object of class matrix
  reject = ....., object of class matrix
  rawdist = ....., object of class matrix
  nulldist = ....., object of class matrix
  nulldist.type = ....., object of class character
  marg.null = ....., object of class character
  marg.par = ....., object of class matrix
  label = ....., object of class numeric
  index = ....., object of class matrix
  call = ....., object of class call
  seed = ....., object of class integer
)
```

Slots

`statistic` Object of class `numeric`, observed test statistics for each hypothesis, specified by the values of the `MTP` arguments `test`, `robust`, `standardize`, and `psi0`.

`estimate` For the test of single-parameter null hypotheses using t-statistics (i.e., not the F-tests), the numeric vector of estimated parameters corresponding to each hypothesis, e.g. means, differences in means, regression parameters.

`sampsize` Object of class `numeric`, number of columns (i.e. observations) in the input data set.

`rawp` Object of class `numeric`, unadjusted, marginal p-values for each hypothesis.

`adjp` Object of class `numeric`, adjusted (for multiple testing) p-values for each hypothesis (computed only if the `get.adjp` argument is `TRUE`).

`conf.reg` For the test of single-parameter null hypotheses using t-statistics (i.e., not the F-tests), the numeric array of lower and upper simultaneous confidence limits for the parameter vector, for each value of the nominal Type I error rate `alpha` (computed only if the `get.cr` argument is `TRUE`).

`cutoff` The numeric matrix of cut-offs for the vector of test statistics for each value of the nominal Type I error rate `alpha` (computed only if the `get.cutoff` argument is `TRUE`).

`reject` Object of class `'matrix'`, rejection indicators (`TRUE` for a rejected null hypothesis), for each value of the nominal Type I error rate `alpha`.

`rawdist` The numeric matrix for the estimated nonparametric non-null test statistics distribution (returned only if `keep.rawdist=TRUE` and if `nulldist` is one of `'boot.ctr'`, `'boot.cs'`, or `'boot.qt'`). This slot must not be empty if one wishes to call `update` to change choice of bootstrap-based null distribution.

- `nullldist` The numeric matrix for the estimated test statistics null distribution (returned only if `keep.nullldist=TRUE`); option not currently available for permutation null distribution, i.e., `nullldist='perm'`). By default (i.e., for `nullldist='boot.cs'`), the entries of `nullldist` are the null value shifted and scaled bootstrap test statistics, with one null test statistic value for each hypothesis (rows) and bootstrap iteration (columns).
- `nullldist.type` Character value describing which choice of null distribution was used to generate the MTP results. Takes on one of the values of the original `nullldist` argument in the call to MTP, i.e., `'boot.cs'`, `'boot.ctr'`, `'boot.qt'`, `'ic'`, or `'perm'`.
- `marg.null` If `nullldist='boot.qt'`, a character value returning which choice of marginal null distribution was used by the MTP. Can be used to check default values or to ensure manual settings were correctly applied.
- `marg.par` If `nullldist='boot.qt'`, a numeric matrix returning the parameters of the marginal null distribution(s) used by the MTP. Can be used to check default values or to ensure manual settings were correctly applied.
- `label` If `keep.label=TRUE`, a vector storing the values used in the argument `Y`. Storing this object is particularly important when one wishes to update MTP objects with F-statistics using default `marg.null` and `marg.par` settings when `nullldist='boot.qt'`.
- `index` For tests of correlation parameters a matrix corresponding to $t(\text{combn}(p, 2))$, where p is the number of variables in X . This matrix gives the indices of the variables considered in each pairwise correlation. For all other tests, this slot is empty, as the indices are in the same order as the rows of X .
- `call` Object of class `call`, the call to the MTP function.
- `seed` An integer or vector for specifying the state of the random number generator used to create the resampled datasets. The seed can be reused for reproducibility in a repeat call to MTP. This argument is currently used only for the bootstrap null distribution (i.e., for `nullldist="boot.xx"`). See `?set.seed` for details.

Methods

`signature(x = "MTP")`

`[` : Subsetting method for MTP class, which operates selectively on each slot of an MTP instance to retain only the data related to the specified hypotheses.

`as.list` : Converts an object of class MTP to an object of class `list`, with an entry for each slot.

`plot` : plot methods for MTP class, produces the following graphical summaries of the results of a MTP. The type of display may be specified via the `which` argument.

1. Scatterplot of number of rejected hypotheses vs. nominal Type I error rate.
2. Plot of ordered adjusted p-values; can be viewed as a plot of Type I error rate vs. number of rejected hypotheses.
3. Scatterplot of adjusted p-values vs. test statistics (also known as "volcano plot").
4. Plot of unordered adjusted p-values.
5. Plot of confidence regions for user-specified parameters, by default the 10 parameters corresponding to the smallest adjusted p-values (argument `top`).

6. Plot of test statistics and corresponding cut-offs (for each value of `alpha`) for user-specified hypotheses, by default the 10 hypotheses corresponding to the smallest adjusted p-values (argument `top`).

The argument `logscale` (by default equal to `FALSE`) allows one to use the negative decimal logarithms of the adjusted p-values in the second, third, and fourth graphical displays. The arguments `caption` and `sub.caption` allow one to change the titles and subtitles for each of the plots (default subtitle is the MTP function call). Note that some of these plots are implemented in the older function `mt.plot`.

print : print method for MTP class, returns a description of an object of class MTP, including sample size, number of tested hypotheses, type of test performed (value of argument `test`), Type I error rate (value of argument `typeone`), nominal level of the test (value of argument `alpha`), name of the MTP (value of argument `method`), call to the function MTP.

In addition, this method produces a table with the class, mode, length, and dimension of each slot of the MTP instance.

summary : summary method for MTP class, provides numerical summaries of the results of a MTP and returns a list with the following three components.

1. `rejections`: A data.frame with the number(s) of rejected hypotheses for the nominal Type I error rate(s) specified by the `alpha` argument of the function MTP. (NULL values are returned if all three arguments `get.cr`, `get.cutoff`, and `get.adj` are `FALSE`).

2. `index`: A numeric vector of indices for ordering the hypotheses according to first `adj`, then `rawp`, and finally the absolute value of `statistic` (not printed in the summary).

3. `summaries`: When applicable (i.e., when the corresponding quantities are returned by MTP), a table with six number summaries of the distributions of the adjusted p-values, unadjusted p-values, test statistics, and parameter estimates.

update : update method for MTP class, provides a mechanism to re-run the MTP with different choices of the following arguments - `nulldist`, `alternative`, `typeone`, `k`, `q`, `fdr.method`, `alpha`, `smooth.null`, `method`, `get.cr`, `get.cutoff`, `get.adj`, `keep.nulldist`, `keep.rawdist`, `keep.margpar`. When evaluate is 'TRUE', a new object of class MTP is returned. Else, the updated call is returned. The MTP object passed to the update method must have either a non-empty `rawdist` slot or a non-empty `nulldist` slot (i.e., must have been called with either 'keep.rawdist=TRUE' or 'keep.nulldist=TRUE').

To save on memory, if one knows ahead of time that one will want to compare different choices of bootstrap-based null distribution, then it is both necessary and sufficient to specify 'keep.rawdist=TRUE', as there is no other means of moving between null distributions other than through the non-transformed non-parametric bootstrap distribution. In this case, 'keep.nulldist=FALSE' may be used. Specifically, if an object of class MTP contains a non-empty `rawdist` slot and an empty `nulldist` slot, then a new null distribution will be generated according to the values of the `nulldist=` argument in the original call to MTP or any additional specifications in the call to `update`. On the other hand, if one knows that one wishes to only update an MTP object in ways which do not involve choice of null distribution, then 'keep.nulldist=TRUE' will suffice and 'keep.rawdist' can be set to `FALSE` (default settings). The original null distribution object will then be used for all subsequent calls to `update`.

N.B.: Note that `keep.rawdist=TRUE` is only available for the bootstrap-based resampling

methods. The non-null distribution does not exist for the permutation or influence curve multivariate normal null distributions.

mtp2ebmtp : coercion method for converting objects of class `MTP` to objects of class `EBMTP`. Slots common to both objects are taken from the object of class `MTP` and used to create a new object of class `EBMTP`. Once an object of class `EBMTP` is created, one may use the method `EBupdate` to perform resampling-based empirical Bayes multiple testing without the need for repeated resampling.

Author(s)

Katherine S. Pollard and Houston N. Gilbert with design contributions from Sandrine Dudoit and Mark J. van der Laan.

References

M.J. van der Laan, S. Dudoit, K.S. Pollard (2004), Augmentation Procedures for Control of the Generalized Family-Wise Error Rate and Tail Probabilities for the Proportion of False Positives, *Statistical Applications in Genetics and Molecular Biology*, 3(1). <http://www.bepress.com/sagmb/vol3/iss1/art15/>

M.J. van der Laan, S. Dudoit, K.S. Pollard (2004), Multiple Testing. Part II. Step-Down Procedures for Control of the Family-Wise Error Rate, *Statistical Applications in Genetics and Molecular Biology*, 3(1). <http://www.bepress.com/sagmb/vol3/iss1/art14/>

S. Dudoit, M.J. van der Laan, K.S. Pollard (2004), Multiple Testing. Part I. Single-Step Procedures for Control of General Type I Error Rates, *Statistical Applications in Genetics and Molecular Biology*, 3(1). <http://www.bepress.com/sagmb/vol3/iss1/art13/>

Katherine S. Pollard and Mark J. van der Laan, "Resampling-based Multiple Testing: Asymptotic Control of Type I Error and Applications to Gene Expression Data" (June 24, 2003). U.C. Berkeley Division of Biostatistics Working Paper Series. Working Paper 121. <http://www.bepress.com/ucbbiostat/paper121>

M.J. van der Laan and A.E. Hubbard (2006), Quantile-function Based Null Distributions in Resampling Based Multiple Testing, *Statistical Applications in Genetics and Molecular Biology*, 5(1). <http://www.bepress.com/sagmb/vol5/iss1/art14/>

S. Dudoit and M.J. van der Laan. *Multiple Testing Procedures and Applications to Genomics*. Springer Series in Statistics. Springer, New York, 2008.

See Also

[MTP](#), [MTP-methods](#), [EBMTP](#), [EBMTP-methods](#), [\[-methods](#), [as.list-methods](#), [print-methods](#), [plot-methods](#), [summary-methods](#), [mtp2ebmtp](#), [ebmtp2mtp](#)

Examples

```
## See MTP function: ? MTP
```

Description

Summary, printing, plotting, subsetting, updating, `as.list` and class conversion methods were defined for the MTP and EBMTM classes. These methods provide visual and numeric summaries of the results of a multiple testing procedure (MTP) and allow one to perform some basic manipulations of objects class MTP or EBMTM.

Several of the methods with the same name will work on objects of their respective class. One exception to this rule is the difference between `update` and `EBupdate` (described below). Because of the differences in the testing procedures, separately named methods were chosen to clearly delineate which method was being applied to which type of object.

Methods

`[` : Subsetting method for MTP and EBMTM classes, which operates selectively on each slot of an MTP or EBMTM instance to retain only the data related to the specified hypotheses.

`as.list` : Converts an object of class MTP or EBMTM to an object of class `list`, with an entry for each slot.

`plot` : plot methods for MTP and EBMTM classes, produces the following graphical summaries of the results of a MTP. The type of display may be specified via the `which` argument.

1. Scatterplot of number of rejected hypotheses vs. nominal Type I error rate.
2. Plot of ordered adjusted p-values; can be viewed as a plot of Type I error rate vs. number of rejected hypotheses.
3. Scatterplot of adjusted p-values vs. test statistics (also known as volcano plot).
4. Plot of unordered adjusted p-values.

Only for objects of class MTP:

5. Plot of confidence regions for user-specified parameters, by default the 10 parameters corresponding to the smallest adjusted p-values (argument `top`).
6. Plot of test statistics and corresponding cut-offs (for each value of `alpha`) for user-specified hypotheses, by default the 10 hypotheses corresponding to the smallest adjusted p-values (argument `top`).

Plots (5) and (6) are not available for objects of class EBMTM because the function EBMTM returns only adjusted p-values and not confidence regions of cut-offs. The argument `logscale` (by default equal to `FALSE`) allows one to use the negative decimal logarithms of the adjusted p-values in the second, third, and fourth graphical displays. The arguments `caption` and `sub.caption` allow one to change the titles and subtitles for each of the plots (default subtitle is the MTP function call). Note that some of these plots are implemented in the older function `mt.plot`.

print : print method for MTP and EBMTTP classes, returns a description of an object of either class, including sample size, number of tested hypotheses, type of test performed (value of argument `test`), Type I error rate (value of argument `typeone`), nominal level of the test (value of argument `alpha`), name of the MTP (value of argument `method`), call to the function MTP or EBMTTP.

In addition, this method produces a table with the class, mode, length, and dimension of each slot of the MTP or EBMTTP instance.

summary : summary method for MTP and EBMTTP classes, provides numerical summaries of the results of a MTP and returns a list with the following three components.

1. **rejections**: A data.frame with the number(s) of rejected hypotheses for the nominal Type I error rate(s) specified by the `alpha` argument of the function MTP or EBMTTP. (For objects of class MTP, NULL values are returned if all three arguments `get.cr`, `get.cutoff`, and `get.adj` are FALSE).

2. **index**: A numeric vector of indices for ordering the hypotheses according to first `adj`, then `rawp`, and finally the absolute value of `statistic` (not printed in the summary).

3. **summaries**: When applicable (i.e., when the corresponding quantities are returned by MTP or EBMTTP), a table with six number summaries of the distributions of the adjusted p-values, unadjusted p-values, test statistics, and parameter estimates.

update : update methods for MTP class, respectively, provides a mechanism to re-run the MTP with different choices of the following arguments - `nulldist`, `alternative`, `typeone`, `k`, `q`, `fdr.method`, `alpha`, `smooth.null`, `method`, `get.cr`, `get.cutoff`, `get.adj`, `keep.nulldist`, `keep.rawdist`, `keep.margpar`. When `evaluate` is 'TRUE', a new object of class MTP is returned. Else, the updated call is returned. The MTP object passed to the update method must have either a non-empty `rawdist` slot or a non-empty `nulldist` slot (i.e., must have been called with either 'keep.rawdist=TRUE' or 'keep.nulldist=TRUE').

EBupdate : update method for EBMTTP class, provides a mechanism to re-run the MTP with different choices of the following arguments - `nulldist`, `alternative`, `typeone`, `k`, `q`, `alpha`, `smooth.null`, `bw`, `kernel`, `prior`, `keep.nulldist`, `keep.rawdist`, `keep.falsepos`, `keep.truepos`, `keep.errormat`, `keep.margpar`. When `evaluate` is 'TRUE', a new object of class EBMTTP is returned. Else, the updated call is returned. The EBMTTP object passed to the update method must have either a non-empty `rawdist` slot or a non-empty `nulldist` slot (i.e., must have been called with either 'keep.rawdist=TRUE' or 'keep.nulldist=TRUE').

Additionally, when calling `EBupdate` for any Type I error rate other than FWER, the `typeone` argument must be specified (even if the original object did not control FWER). For example, `typeone="fdr"`, would always have to be specified, even if the original object also controlled the FDR. In other words, for all function arguments, it is safest to always assume that you are updating from the EBMTTP default function settings, regardless of the original call to the EBMTTP function. Currently, the main advantage of the `EBupdate` method is that it prevents the need for repeated estimation of the test statistics null distribution.

To save on memory, if one knows ahead of time that one will want to compare different choices of bootstrap-based null distribution, then it is both necessary and sufficient to specify 'keep.rawdist=TRUE', as there is no other means of moving between null distributions other than through the non-transformed non-parametric bootstrap distribution. In this case, 'keep.nulldist=FALSE' may be used. Specifically, if an object of class MTP or EBMTTP contains a non-empty `rawdist` slot and an empty `nulldist` slot, then a new null distribu-

tion will be generated according to the values of the `nullldist=` argument in the original call to `(EB)MTP` or any additional specifications in the call to `(EB)update`. On the other hand, if one knows that one wishes to only update an `(EB)MTP` object in ways which do not involve choice of bootstrap null distribution, then `'keep.nullldist=TRUE'` will suffice and `'keep.rawldist'` can be set to `FALSE` (default settings). The original null distribution object will then be used for all subsequent calls to `update`.

N.B.: Note that `keep.rawldist=TRUE` is only available for the bootstrap-based resampling methods. The non-null distribution does not exist for the permutation or influence curve multivariate normal null distributions.

mtp2ebmtp : coercion method for converting objects of class `MTP` to objects of class `EBMTP`. Slots common to both objects are taken from the object of class `MTP` and used to create a new object of class `EBMTP`. Once an object of class `EBMTP` is created, one may use the method `EBupdate` to perform resampling-based empirical Bayes multiple testing without the need for repeated resampling.

ebmtp2mtp : coercion method for converting objects of class `EBMTP` to objects of class `MTP`. Slots common to both objects are taken from the object of class `EBMTP` and used to create a new object of class `MTP`. Once an object of class `MTP` is created, one may use the method `update` to perform resampling-based multiple testing (as would have been done with calls to `MTP`) without the need for repeated resampling.

Author(s)

Katherine S. Pollard and Houston N. Gilbert with design contributions from Sandrine Dudoit and Mark J. van der Laan.

MTP

A function to perform resampling-based multiple hypothesis testing

Description

A user-level function to perform multiple testing procedures (MTP). A variety of t- and F-tests, including robust versions of most tests, are implemented. Single-step and step-down minP and maxT methods are used to control the chosen type I error rate (FWER, gFWER, TPPFP, or FDR). Bootstrap and permutation null distributions are available. Additionally, for t-statistics, one may wish to sample from an appropriate multivariate normal distribution with mean zero and correlation matrix derived from the vector influence function. Arguments are provided for user control of output. Gene selection in microarray experiments is one application.

Usage

```
MTP(X, W = NULL, Y = NULL, Z = NULL, Z.incl = NULL, Z.test = NULL,
     na.rm = TRUE, test = "t.twosamp.unequalvar", robust = FALSE,
     standardize = TRUE, alternative = "two.sided", psi0 = 0,
     typeone = "fwer", k = 0, q = 0.1, fdr.method = "conservative",
     alpha = 0.05, smooth.null = FALSE, nullldist = "boot.cs",
     B = 1000, ic.quant.trans = FALSE, MVN.method = "mvrnorm",
     penalty = 1e-06, method = "ss.maxT", get.cr = FALSE, get.cutoff = FALSE,
     get.adjp = TRUE, keep.nullldist = TRUE, keep.rawldist = FALSE,
     seed = NULL, cluster = 1, type = NULL, dispatch = NULL, marg.null = NULL,
     marg.par = NULL, keep.margpar = TRUE, ncp = NULL, perm.mat = NULL,
     keep.index = FALSE, keep.label = FALSE)
```

Arguments

<code>X</code>	A matrix, data.frame or ExpressionSet containing the raw data. In the case of an ExpressionSet, <code>exprs(X)</code> is the data of interest and <code>pData(X)</code> may contain outcomes and covariates of interest. For most currently implemented tests (exception: tests involving correlation parameters), one hypothesis is tested for each row of the data.
<code>W</code>	A vector or matrix containing non-negative weights to be used in computing the test statistics. If a matrix, <code>W</code> must be the same dimension as <code>X</code> with one weight for each value in <code>X</code> . If a vector, <code>W</code> may contain one weight for each observation (i.e. column) of <code>X</code> or one weight for each variable (i.e. row) of <code>X</code> . In either case, the weights are duplicated appropriately. Weighted F-tests are not available. Default is 'NULL'.
<code>Y</code>	A vector, factor, or <code>Surv</code> object containing the outcome of interest. This may be class labels (F-tests and two sample t-tests) or a continuous or polycotomous dependent variable (linear regression based t-tests), or survival data (Cox proportional hazards based t-tests). For <code>block.f</code> and <code>f.twoway</code> tests, class labels must be ordered by block and within each block ordered by group. If <code>X</code> is an ExpressionSet, <code>Y</code> can be a character string referring to the column of <code>pData(X)</code> to use as outcome. Default is 'NULL'.
<code>Z</code>	A vector, factor, or matrix containing covariate data to be used in the regression (linear and Cox) models. Each variable should be in one column, so that <code>nrow(Z) = ncol(X)</code> . If <code>X</code> is an ExpressionSet, <code>Z</code> can be a character string referring to the column of <code>pData(X)</code> to use as covariates. The variables <code>Z.incl</code> and <code>Z.adj</code> allow one to specify which covariates to use in a particular test without modifying the input <code>Z</code> . Default is 'NULL'.
<code>Z.incl</code>	The indices of the columns of <code>Z</code> (i.e. which variables) to include in the model. These can be numbers or column names (if the columns are names). Default is 'NULL'.
<code>Z.test</code>	The index or names of the column of <code>Z</code> (i.e. which variable) to use to test for association with each row of <code>X</code> in a linear model. Only used for <code>test="lm.XvsZ"</code> , where it is necessary to specify which covariate's regression parameter is of interest. Default is 'NULL'.
<code>na.rm</code>	Logical indicating whether to remove observations with an NA. Default is 'TRUE'.
<code>test</code>	Character string specifying the test statistics to use, by default 't.twosamp.unequalvar'. See details (below) for a list of tests.
<code>robust</code>	Logical indicating whether to use the robust version of the chosen test, e.g. Wilcoxon signed rank test for robust one-sample t-test or <code>rlm</code> instead of <code>lm</code> in linear models. Default is 'FALSE'.
<code>standardize</code>	Logical indicating whether to use the standardized version of the test statistics (usual t-statistics are standardized). Default is 'TRUE'.
<code>alternative</code>	Character string indicating the alternative hypotheses, by default 'two.sided'. For one-sided tests, use 'less' or 'greater' for null hypotheses of 'greater than or equal' (i.e. alternative is 'less') and 'less than or equal', respectively.
<code>psi0</code>	The hypothesized null value, typically zero (default). Currently, this should be a single value, which is used for all hypotheses.
<code>typeone</code>	Character string indicating which type I error rate to control, by default family-wise error rate ('fwer'). Other options include generalized family-wise error rate ('gfwer'), with parameter <code>k</code> giving the allowed number of false positives, and

	tail probability of the proportion of false positives (<code>'tppfp'</code>), with parameter α giving the allowed proportion of false positives. The false discovery rate (<code>'fdr'</code>) can also be controlled.
<code>k</code>	The allowed number of false positives for gFWER control. Default is 0 (FWER).
<code>q</code>	The allowed proportion of false positives for TPPFP control. Default is 0.1.
<code>fdr.method</code>	Character string indicating which FDR controlling method should be used when <code>typeone="fdr"</code> . The options are "conservative" (default) for the more conservative, general FDR controlling procedure and "restricted" for the method which requires more assumptions.
<code>alpha</code>	The target nominal type I error rate, which may be a vector of error rates. Default is 0.05.
<code>smooth.null</code>	Indicator of whether to use a kernel density estimate for the tail of the null distribution for computing raw p-values close to zero. Only used if <code>'rawp'</code> would be zero without smoothing. Default is <code>'FALSE'</code> .
<code>nulldist</code>	Character string indicating which resampling method to use for estimating the joint test statistics null distribution, by default the non-parametric bootstrap with centering and scaling (<code>'boot.cs'</code>). The old default <code>'boot'</code> will still compile and will correspond to <code>'boot.cs'</code> . Other null distribution options include <code>'perm'</code> , <code>'boot.ctr'</code> , <code>'boot.qt'</code> , and <code>'ic'</code> , corresponding to the permutation distribution, centered-only bootstrap distribution, quantile-transformed bootstrap distribution, and influence curve multivariate normal joint null distribution, respectively. More details below.
<code>B</code>	The number of bootstrap iterations (i.e. how many resampled data sets), the number of permutations (if <code>nulldist</code> is <code>'perm'</code>), or the number of samples from the multivariate normal distribution (if <code>nulldist</code> is <code>'ic'</code>) Can be reduced to increase the speed of computation, at a cost to precision. Default is 1000.
<code>ic.quant.trans</code>	If <code>nulldist='ic'</code> , a logical indicating whether or not a marginal quantile transformation using a t-distribution or user-supplied marginal distribution (stored in <code>perm.mat</code>) should be applied to the multivariate normal null distribution. Defaults for <code>marg.null</code> and <code>marg.par</code> exist, but can also be specified by the user (see below). Default is <code>'FALSE'</code> .
<code>MVN.method</code>	If <code>nulldist='ic'</code> , one of <code>'mvrnorm'</code> or <code>'Cholesky'</code> designating how correlated normal test statistics are to be generated. Selecting <code>'mvrnorm'</code> uses the function of the same name found in the <code>MASS</code> library, whereas <code>'Cholesky'</code> relies on a Cholesky decomposition. Default is <code>'mvrnorm'</code> .
<code>penalty</code>	If <code>nulldist='ic'</code> and <code>MVN.method='Cholesky'</code> , the value in <code>penalty</code> is added to all diagonal elements of the estimated test statistics correlation matrix to ensure that the matrix is positive definite and that internal calls to <code>'chol'</code> do not return an error. Default is <code>1e-6</code> .
<code>method</code>	The multiple testing procedure to use. Options are single-step maxT (<code>'ss.maxT'</code> , default), single-step minP (<code>'ss.minP'</code>), step-down maxT (<code>'sd.maxT'</code>), and step-down minP (<code>'sd.minP'</code>).
<code>get.cr</code>	Logical indicating whether to compute confidence intervals for the estimates. Not available for F-tests. Default is <code>'FALSE'</code> .
<code>get.cutoff</code>	Logical indicating whether to compute thresholds for the test statistics. Default is <code>'FALSE'</code> .
<code>get.adj</code>	Logical indicating whether to compute adjusted p-values. Default is <code>'TRUE'</code> .

<code>keep.nulldist</code>	Logical indicating whether to return the computed bootstrap or influence curve null distribution, by default 'TRUE'. Not available for <code>nulldist='perm'</code> . Note that this matrix can be quite large.
<code>keep.rawdist</code>	Logical indicating whether to return the computed non-null (raw) bootstrap distribution, by default 'FALSE'. Not available when using <code>nulldist='perm'</code> or 'ic'. Note that this matrix can become quite large. If one wishes to use subsequent calls to <code>update</code> or <code>EBupdate</code> in which one updates choice of bootstrap null distribution, <code>keep.rawdist</code> must be TRUE. To save on memory, <code>update</code> only requires that one of <code>keep.nulldist</code> or <code>keep.rawdist</code> be 'TRUE'.
<code>seed</code>	Integer or vector of integers to be used as argument to <code>set.seed</code> to set the seed for the random number generator for bootstrap resampling. This argument can be used to repeat exactly a test performed with a given seed. If the seed is specified via this argument, the same seed will be returned in the seed slot of the MTP object created. Else a random seed(s) will be generated, used and returned. Vector of integers used to specify seeds for each node in a cluster used to generate a bootstrap null distribution.
<code>cluster</code>	Integer for number of nodes to create or a cluster object created through the package <code>snow</code> . With <code>cluster=1</code> , bootstrap is implemented on single node. Supplying a cluster object results in the bootstrap being implemented in parallel on the provided nodes. This option is only available for the bootstrap procedure. With default value of 1, bootstrap is executed on single CPU.
<code>type</code>	Interface system to use for computer cluster. See <code>snow</code> package for details.
<code>dispatch</code>	The number or percentage of bootstrap iterations to dispatch at a time to each node of the cluster if a computer cluster is used. If <code>dispatch</code> is a percentage, <code>B*dispatch</code> must be an integer. If <code>dispatch</code> is an integer, then <code>B/dispatch</code> must be an integer. Default is 5 percent.
<code>marg.null</code>	If <code>nulldist='boot.qt'</code> , the marginal null distribution to use for quantile transformation. Can be one of 'normal', 't', 'f' or 'perm'. Default is 'NULL', in which case the marginal null distribution is selected based on choice of test statistics. Defaults explained below. If 'perm', the user must supply a vector or matrix of test statistics corresponding to another marginal null distribution, perhaps one created externally by the user, and possibly referring to empirically derived <i>marginal permutation distributions</i> , although the statistics could represent any suitable choice of marginal null distribution.
<code>marg.par</code>	If <code>nulldist='boot.qt'</code> , the parameters defining the marginal null distribution in <code>marg.null</code> to be used for quantile transformation. Default is 'NULL', in which case the values are selected based on choice of test statistics and other available parameters (e.g., sample size, number of groups, etc.). Defaults explained below. User can override defaults, in which case a matrix of marginal null distribution parameters can be accepted. Providing numeric (vector) values will apply the same null distribution defined by the parameter to all hypotheses, while providing a matrix of values allows the user to perform multiple testing using parameters which may vary with each hypothesis, as may be desired in common-quantile minP procedures. In this way, theoretical factors or factors affecting sample size or missingness may be assessed.
<code>keep.margpar</code>	If <code>nulldist='boot.qt'</code> , a logical indicating whether the (internally created) matrix of marginal null distribution parameters should be returned. Default is 'TRUE'.

<code>ncp</code>	If <code>nulldist='boot.qt'</code> , a value for a possible noncentrality parameter to be used during marginal quantile transformation. Default is 'NULL'.
<code>perm.mat</code>	If <code>nulldist='boot.qt'</code> and <code>marg.null='perm'</code> , a matrix of user-supplied test statistics from a particular distribution to be used during marginal quantile transformation. The statistics may represent empirically derived marginal permutation values, may be theoretical values, or may represent a sample from some other suitable choice of marginal null distribution.
<code>keep.index</code>	If <code>nulldist='ic'</code> and <code>test='t.cor'</code> or <code>test='z.cor'</code> , the index returned is a matrix with the indices of the first and second variables considered for pairwise correlations. If there are p hypotheses, this arguments returns $t(\text{combn}(p, 2))$. For all other choices of test statistic, the index is not returned, as they correspond to the original order of the hypotheses in X .
<code>keep.label</code>	Default is 'FALSE'. A logical indicating whether or not the label in Y should be returned as a slot in the resulting MTP object. Typically not necessary, although useful if one is using <code>update</code> and wants to use marginal null distribution defaults with <code>nulldist='boot.qt'</code> (e.g., with F-tests).

Details

A multiple testing procedure (MTP) is defined by choices of test statistics, type I error rate, null distribution and method for error rate control. Each component is described here. For two-sample t-tests, the group with the smaller-valued label is subtracted from the group with the larger-valued label. That is, differences in means are calculated as "mean of group 2 - mean of group 1" or "mean of group B - mean of group A". For paired t-tests, the arrangement of group indices does not matter, as long as the columns are arranged in the same corresponding order between groups. For example, if group 1 is coded as 0, and group 2 is coded as 1, for 3 pairs of data, it does not matter if the label Y is coded as "0,0,0,1,1,1", "1,1,1,0,0,0" "0,1,0,1,0,1" or "1,0,1,0,1,0", the paired differences between groups will be calculated as "group 2 - group 1". See references for more detail.

Test statistics are determined by the values of `test`:

t.onesamp: one-sample t-statistic for tests of means;

t.twosamp.equalvar: equal variance two-sample t-statistic for tests of differences in means (two-sample t-statistic);

t.twosamp.unequalvar: unequal variance two-sample t-statistic for tests of differences in means (two-sample Welch t-statistic);

t.pair: two-sample paired t-statistic for tests of differences in means;

f: multi-sample F-statistic for tests of equality of population means (assumes constant variance across groups, but not normality);

f.block: multi-sample F-statistic for tests of equality of population means in a block design (assumes constant variance across groups, but not normality). This test is not available with the bootstrap null distribution;

f.twoway: multi-sample F-statistic for tests of equality of population means in a block design (assumes constant variance across groups, but not normality). Differs from `f.block` in requiring multiple observations per `group*block` combination. This test uses the means of each `group*block` combination as response variable and test for group main effects assuming a randomized block design;

lm.XvsZ: t-statistic for tests of regression coefficients for variable `Z.test` in linear models, each with a row of X as outcome, possibly adjusted by covariates `Z.incl` from the matrix Z (in the case of no covariates, one recovers the one-sample t-statistic, `t.onesamp`);

- lm.YvsXZ:** t-statistic for tests of regression coefficients in linear models, with outcome Y and each row of X as covariate of interest, with possibly other covariates Z. `incl` from the matrix Z;
- coxph.YvsXZ:** t-statistic for tests of regression coefficients in Cox proportional hazards survival models, with outcome Y and each row of X as covariate of interest, with possibly other covariates Z. `incl` from the matrix Z.
- t.cor** t-statistics for tests of pairwise correlation parameters for all variables in X. Note that the number of hypotheses can become quite large very fast. This test is only available with the influence curve null distribution.
- z.cor** Fisher's z-statistics for tests of pairwise correlation parameters for all variables in X. Note that the number of hypotheses can become quite large very fast. This test is only available with the influence curve null distribution.

When `robust=TRUE`, non-parametric versions of each test are performed. For the linear models, this means `rlm` is used instead of `lm`. There is not currently a robust version of `test=coxph.YvsXZ`. For the t- and F-tests, data values are simply replaced by their ranks. This is equivalent to performing the following familiar named rank-based tests. The conversion after each test is the formula to convert from the MTP test to the statistic reported by the listed R function (where `num` is the numerator of the MTP test statistics, `n` is total sample size, `nk` is group k sample size, `K` is total number of groups or treatments, and `rk` are the ranks in group k).

- t.onesamp** or **t.pair:** Wilcoxon signed rank, `wilcox.test` with `y=NULL` or `paired=TRUE`,
conversion: `num/n`
- t.twosamp.equalvar:** Wilcoxon rank sum or Mann-Whitney, `wilcox.test`,
conversion: `n2*(num+mean(r1)) - n2*(n2+1)/2`
- f:** Kruskal-Wallis rank sum, `kruskal.test`,
conversion: `num*12/(n*(n-1))`
- f.block:** Friedman rank sum, `friedman.test`,
conversion: `num*12/(K*(K+1))`
- f.twoway:** Friedman rank sum, `friedman.test`,
conversion: `num*12/(K*(K+1))`

The implemented MTPs are based on control of the family-wise error rate, defined as the probability of any false positives. Let V_n denote the (unobserved) number of false positives. Then, control of FWER at level α means that $\Pr(V_n > 0) \leq \alpha$. The set of rejected hypotheses under a FWER controlling procedure can be augmented to increase the number of rejections, while controlling other error rates. The generalized family-wise error rate is defined as $\Pr(V_n > k) \leq \alpha$, and it is clear that one can simply take an FWER controlling procedure, reject k more hypotheses and have control of gFWER at level α . The tail probability of the proportion of false positives depends on both the number of false positives (V_n) and the number of rejections (R_n). Control of TPPFP at level α means $\Pr(V_n/R_n > q) \leq \alpha$, for some proportion q . Control of the false discovery rate refers to the expected proportion of false positives (rather than a tail probability). Control of FDR at level α means $E(V_n/R_n) \leq \alpha$.

In practice, one must choose a method for estimating the test statistics null distribution. We have implemented several versions of an ordinary non-parametric bootstrap estimator and a permutation estimator (which makes sense in certain settings, see references). The non-parametric bootstrap estimator (default) provides asymptotic control of the type I error rate for any data generating distribution, whereas the permutation estimator requires the subset pivotality assumption. One drawback of both methods is the discreteness of the estimated null distribution when the sample size is small. Furthermore, when the sample size is small enough, it is possible that ties will lead to a very small variance estimate. Using `standardize=FALSE` allows one to avoid these unusually small test statistic denominators. Parametric bootstrap estimators are another option (not yet implemented). For asymptotically linear estimators, such as those commonly probed using t-statistics,

another choice of null distribution is provided when sampling from a multivariate normal distribution with mean zero and correlation matrix derived from the vector influence function. Sampling from a multivariate normal may alleviate the discreteness of the bootstrap and permutation distributions, although accuracy in estimation of the test statistics correlation matrix will be of course also affected by sample size.

For the nonparametric bootstrap distribution with marginal null quantile transformation, the following defaults for `marg.null` and `marg.par` are available based on choice of test statistics, sample size 'n', and various other parameters:

t.onesamp: t-distribution with $df=n-1$;

t.twosamp.equalvar: t-distribution with $df=n-2$;

t.twosamp.unequalvar: $N(0,1)$;

t.pair: t-distribution with $df=n-1$, where n is the number of unique samples, i.e., the number of observed differences between paired samples;

f: F-distribution with $df1=k-1$, $df2=n-k$, for k groups;

f.block: NA. Only available with permutation distribution;

f.twoway: F-distribution with $df1=k-1, df2=n-k*1$, for k groups and l blocks;

lm.XvsZ: $N(0,1)$;

lm.YvsXZ: $N(0,1)$;

coxph.YvsXZ: $N(0,1)$;

t.cor t-distribution with $df=n-2$;

z.cor $N(0,1)$.

The above defaults, however, can be overridden by manually setting values of `marg.null` and `marg.par`. In the case of `nulldist='ic'`, and `ic.quant.trans=TRUE`, the defaults are the same as above except that 'lm.XvsZ' and 'lm.YvsXZ' are replaced with t-distributions with $df=n-p$.

Given observed test statistics, a type I error rate (with nominal level), and a test statistics null distribution, MTPs provide adjusted p-values, cutoffs for test statistics, and possibly confidence regions for estimates. Four methods are implemented, based on minima of p-values and maxima of test statistics. Only the step down methods are currently available with the permutation null distribution.

Computation times using a bootstrap null distribution are slower when weights are used for one and two-sample tests. Computation times when using a bootstrap null distribution also are slower for the tests `lmXvsZ`, `lmYvsXZ`, `coxph.YvsXZ`.

To execute the bootstrap on a computer cluster, a cluster object generated with `makeCluster` in the package `snow` may be used as the argument for `cluster`. Alternatively, the number of nodes to use in the computer cluster can be used as the argument to `cluster`. In this case, `type` must be specified and a cluster will be created. In both cases, `Biobase` and `multtest` will be loaded onto each cluster node if these libraries are located in a directory in the standard search path. If these libraries are in a non-standard location, it is necessary to first create the cluster, load `Biobase` and `multtest` on each node and then to use the cluster object as the argument to `cluster`. See documentation for `snow` package for additional information on creating and using a cluster.

Finally, note that the old argument `csnull` is now DEPRECATED as of `multtest` v. 2.0.0 given the expanded null distribution options described above. Previously, this argument was an indicator of whether the bootstrap estimated test statistics distribution should be centered and scaled (to produce a null distribution) or not. If `csnull=FALSE`, the (raw) non-null bootstrap estimated test statistics distribution was returned. If the non-null bootstrap distribution should be returned, this object is now stored in the 'rawdist' slot when `keep.rawdist=TRUE` in the original MTP function call.

Value

An object of class `MTP`, with the following slots:

<code>statistic</code>	Object of class <code>numeric</code> , observed test statistics for each hypothesis, specified by the values of the <code>MTP</code> arguments <code>test</code> , <code>robust</code> , <code>standardize</code> , and <code>psi0</code> .
<code>estimate</code>	For the test of single-parameter null hypotheses using t-statistics (i.e., not the F-tests), the numeric vector of estimated parameters corresponding to each hypothesis, e.g. means, differences in means, regression parameters.
<code>sampsize</code>	Object of class <code>numeric</code> , number of columns (i.e. observations) in the input data set.
<code>rawp</code>	Object of class <code>numeric</code> , unadjusted, marginal p-values for each hypothesis.
<code>adjp</code>	Object of class <code>numeric</code> , adjusted (for multiple testing) p-values for each hypothesis (computed only if the <code>get.adj</code> argument is <code>TRUE</code>).
<code>conf.reg</code>	For the test of single-parameter null hypotheses using t-statistics (i.e., not the F-tests), the numeric array of lower and upper simultaneous confidence limits for the parameter vector, for each value of the nominal Type I error rate <code>alpha</code> (computed only if the <code>get.cr</code> argument is <code>TRUE</code>).
<code>cutoff</code>	The numeric matrix of cut-offs for the vector of test statistics for each value of the nominal Type I error rate <code>alpha</code> (computed only if the <code>get.cutoff</code> argument is <code>TRUE</code>).
<code>reject</code>	Object of class <code>'matrix'</code> , rejection indicators (<code>TRUE</code> for a rejected null hypothesis), for each value of the nominal Type I error rate <code>alpha</code> .
<code>rawdist</code>	The numeric matrix for the estimated nonparametric non-null test statistics distribution (returned only if <code>keep.rawdist=TRUE</code> and if <code>nulldist</code> is one of <code>'boot.ctr'</code> , <code>'boot.cs'</code> , or <code>'boot.qt'</code>). This slot must not be empty if one wishes to call <code>update</code> to change choice of bootstrap-based null distribution.
<code>nulldist</code>	The numeric matrix for the estimated test statistics null distribution (returned only if <code>keep.nulldist=TRUE</code>); option not currently available for permutation null distribution, i.e., <code>nulldist='perm'</code>). By default (i.e., for <code>nulldist='boot.cs'</code>), the entries of <code>nulldist</code> are the null value shifted and scaled bootstrap test statistics, with one null test statistic value for each hypothesis (rows) and bootstrap iteration (columns).
<code>nulldist.type</code>	Character value describing which choice of null distribution was used to generate the <code>MTP</code> results. Takes on one of the values of the original <code>nulldist</code> argument in the call to <code>MTP</code> , i.e., <code>'boot.cs'</code> , <code>'boot.ctr'</code> , <code>'boot.qt'</code> , <code>'ic'</code> , or <code>'perm'</code> .
<code>marg.null</code>	If <code>nulldist='boot.qt'</code> , a character value returning which choice of marginal null distribution was used by the <code>MTP</code> . Can be used to check default values or to ensure manual settings were correctly applied.
<code>marg.par</code>	If <code>nulldist='boot.qt'</code> , a numeric matrix returning the parameters of the marginal null distribution(s) used by the <code>MTP</code> . Can be used to check default values or to ensure manual settings were correctly applied.
<code>call</code>	Object of class <code>call</code> , the call to the <code>MTP</code> function.
<code>seed</code>	An integer or vector for specifying the state of the random number generator used to create the resampled datasets. The seed can be reused for reproducibility in a repeat call to <code>MTP</code> . This argument is currently used only for the bootstrap null distribution (i.e., for <code>nulldist="boot.xx"</code>). See <code>?set.seed</code> for details.

Note

Thank you to Peter Dimitrov for suggestions about the code.

Author(s)

Katherine S. Pollard and Houston N. Gilbert with design contributions from Sandra Taylor, Sandrine Dudoit and Mark J. van der Laan.

References

M.J. van der Laan, S. Dudoit, K.S. Pollard (2004), Augmentation Procedures for Control of the Generalized Family-Wise Error Rate and Tail Probabilities for the Proportion of False Positives, *Statistical Applications in Genetics and Molecular Biology*, 3(1). <http://www.bepress.com/sagmb/vol3/iss1/art15/>

M.J. van der Laan, S. Dudoit, K.S. Pollard (2004), Multiple Testing. Part II. Step-Down Procedures for Control of the Family-Wise Error Rate, *Statistical Applications in Genetics and Molecular Biology*, 3(1). <http://www.bepress.com/sagmb/vol3/iss1/art14/>

S. Dudoit, M.J. van der Laan, K.S. Pollard (2004), Multiple Testing. Part I. Single-Step Procedures for Control of General Type I Error Rates, *Statistical Applications in Genetics and Molecular Biology*, 3(1). <http://www.bepress.com/sagmb/vol3/iss1/art13/>

K.S. Pollard and Mark J. van der Laan, "Resampling-based Multiple Testing: Asymptotic Control of Type I Error and Applications to Gene Expression Data" (June 24, 2003). U.C. Berkeley Division of Biostatistics Working Paper Series. Working Paper 121. <http://www.bepress.com/ucbbiostat/paper121>

M.J. van der Laan and A.E. Hubbard (2006), Quantile-function Based Null Distributions in Resampling Based Multiple Testing, *Statistical Applications in Genetics and Molecular Biology*, 5(1). <http://www.bepress.com/sagmb/vol5/iss1/art14/>

S. Dudoit and M.J. van der Laan. *Multiple Testing Procedures and Applications to Genomics*. Springer Series in Statistics. Springer, New York, 2008.

See Also

[EBMTP](#), [MTP-class](#), [MTP-methods](#), [mt.minP](#), [mt.maxT](#), [ss.maxT](#), [fwer2gfwer](#)

Examples

```
#data
set.seed(99)
data<-matrix(rnorm(90),nr=9)
group<-c(rep(1,5),rep(0,5))

#fwer control with centered and scaled bootstrap null distribution
#(B=100 for speed)
m1<-MTP(X=data,Y=group,alternative="less",B=100,method="sd.minP")
print(m1)
summary(m1)
par(mfrow=c(2,2))
plot(m1,top=9)

#fwer control with quantile transformed bootstrap null distribution
#default settings = N(0,1) marginal null distribution
m2<-MTP(X=data,Y=group,alternative="less",B=100,method="sd.minP",
```

```

nullldist="boot.qt",keep.rawdist=TRUE)

#fwer control with quantile transformed bootstrap null distribution
#marginal null distribution and df parameters manually set,
#first all equal, then varying with hypothesis
m3<-update(m2,marg.null="t",marg.par=10)
mps<-matrix(c(rep(9,5),rep(10,5)),nr=10,nc=1)
m4<-update(m2,marg.null="t",marg.par=mps)

m1@nullldist.type
m2@nullldist.type
m2@marg.null
m2@marg.par
m3@nullldist.type
m3@marg.null
m3@marg.par
m4@nullldist.type
m4@marg.null
m4@marg.par

```

golub

Gene expression dataset from Golub et al. (1999)

Description

Gene expression data (3051 genes and 38 tumor mRNA samples) from the leukemia microarray study of Golub et al. (1999). Pre-processing was done as described in Dudoit et al. (2002). The R code for pre-processing is available in the file [../doc/golub.R](#).

Usage

```
data(golub)
```

Value

golub	matrix of gene expression levels for the 38 tumor mRNA samples, rows correspond to genes (3051 genes) and columns to mRNA samples.
golub.cl	numeric vector indicating the tumor class, 27 acute lymphoblastic leukemia (ALL) cases (code 0) and 11 acute myeloid leukemia (AML) cases (code 1).
golub.gnames	a matrix containing the names of the 3051 genes for the expression matrix golub. The three columns correspond to the gene index, ID, and Name, respectively.

Source

Golub et al. (1999). Molecular classification of cancer: class discovery and class prediction by gene expression monitoring, *Science*, Vol. 286:531-537.
<http://www-genome.wi.mit.edu/MPR/>

References

S. Dudoit, J. Fridlyand, and T. P. Speed (2002). Comparison of discrimination methods for the classification of tumors using gene expression data. *Journal of the American Statistical Association*, Vol. 97, No. 457, p. 77–87.

 mt.maxT

Step-down maxT and minP multiple testing procedures

Description

These functions compute permutation adjusted p -values for step-down multiple testing procedures described in Westfall & Young (1993).

Usage

```
mt.maxT(X, classlabel, test="t", side="abs", fixed.seed.sampling="y", B=10000, na=.mt)
mt.minP(X, classlabel, test="t", side="abs", fixed.seed.sampling="y", B=10000, na=.mt)
```

Arguments

X	A data frame or matrix, with m rows corresponding to variables (hypotheses) and n columns to observations. In the case of gene expression data, rows correspond to genes and columns to mRNA samples. The data can be read using read.table .
classlabel	A vector of integers corresponding to observation (column) class labels. For k classes, the labels must be integers between 0 and $k - 1$. For the <code>blockf</code> test option, observations may be divided into n/k blocks of k observations each. The observations are ordered by block, and within each block, they are labeled using the integers 0 to $k - 1$.
test	A character string specifying the statistic to be used to test the null hypothesis of no association between the variables and the class labels. If <code>test="t"</code> , the tests are based on two-sample Welch t-statistics (unequal variances). If <code>test="t.equalvar"</code> , the tests are based on two-sample t-statistics with equal variance for the two samples. The square of the t-statistic is equal to an F-statistic for $k = 2$. If <code>test="wilcoxon"</code> , the tests are based on standardized rank sum Wilcoxon statistics. If <code>test="f"</code> , the tests are based on F-statistics. If <code>test="pairt"</code> , the tests are based on paired t-statistics. The square of the paired t-statistic is equal to a block F-statistic for $k = 2$. If <code>test="blockf"</code> , the tests are based on F-statistics which adjust for block differences (cf. two-way analysis of variance).
side	A character string specifying the type of rejection region. If <code>side="abs"</code> , two-tailed tests, the null hypothesis is rejected for large absolute values of the test statistic. If <code>side="upper"</code> , one-tailed tests, the null hypothesis is rejected for large values of the test statistic. If <code>side="lower"</code> , one-tailed tests, the null hypothesis is rejected for small values of the test statistic.

<code>fixed.seed.sampling</code>	If <code>fixed.seed.sampling="y"</code> , a fixed seed sampling procedure is used, which may double the computing time, but will not use extra memory to store the permutations. If <code>fixed.seed.sampling="n"</code> , permutations will be stored in memory. For the <code>blockf</code> test, the option <code>n</code> was not implemented as it requires too much memory.
<code>B</code>	The number of permutations. For a complete enumeration, <code>B</code> should be 0 (zero) or any number not less than the total number of permutations.
<code>na</code>	Code for missing values (the default is <code>.mt.naNUM=--93074815.62</code>). Entries with missing values will be ignored in the computation, i.e., test statistics will be based on a smaller sample size. This feature has not yet fully implemented.
<code>nonpara</code>	If <code>nonpara="y"</code> , nonparametric test statistics are computed based on ranked data. If <code>nonpara="n"</code> , the original data are used.

Details

These functions compute permutation adjusted p -values for the step-down maxT and minP multiple testing procedures, which provide strong control of the family-wise Type I error rate (FWER). The adjusted p -values for the minP procedure are defined in equation (2.10) p. 66 of Westfall & Young (1993), and the maxT procedure is discussed p. 50 and 114. The permutation algorithms for estimating the adjusted p -values are given in Ge et al. (In preparation). The procedures are for the simultaneous test of m null hypotheses, namely, the null hypotheses of no association between the m variables corresponding to the rows of the data frame `X` and the class labels `classlabel`. For gene expression data, the null hypotheses correspond to no differential gene expression across mRNA samples.

Value

A data frame with components

<code>index</code>	Vector of row indices, between 1 and <code>nrow(X)</code> , where rows are sorted first according to their adjusted p -values, next their unadjusted p -values, and finally their test statistics.
<code>teststat</code>	Vector of test statistics, ordered according to <code>index</code> . To get the test statistics in the original data order, use <code>teststat[order(index)]</code> .
<code>rawp</code>	Vector of raw (unadjusted) p -values, ordered according to <code>index</code> .
<code>adjp</code>	Vector of adjusted p -values, ordered according to <code>index</code> .
<code>plover</code>	For <code>mt.minP</code> function only, vector of "adjusted p -values", where ties in the permutation distribution of the successive minima of raw p -values with the observed p -values are counted only once. Note that procedures based on <code>plover</code> do not control the FWER. Comparison of <code>plover</code> and <code>adjp</code> gives an idea of the discreteness of the permutation distribution. Values in <code>plover</code> are ordered according to <code>index</code> .

Author(s)

Yongchao Ge, <yongchao.ge@mssm.edu>,
Sandrine Dudoit, <http://www.stat.berkeley.edu/~sandrine>.

References

S. Dudoit, J. P. Shaffer, and J. C. Boldrick (Submitted). Multiple hypothesis testing in microarray experiments.

Y. Ge, S. Dudoit, and T. P. Speed. Resampling-based multiple testing for microarray data hypothesis, Technical Report \#633 of UCB Stat. <http://www.stat.berkeley.edu/~gyc>

P. H. Westfall and S. S. Young (1993). *Resampling-based multiple testing: Examples and methods for p-value adjustment*. John Wiley & Sons.

See Also

[mt.plot](#), [mt.rawp2adjp](#), [mt.reject](#), [mt.sample.teststat](#), [mt.teststat](#), [golub](#).

Examples

```
# Gene expression data from Golub et al. (1999)
# To reduce computation time and for illustrative purposes, we consider only
# the first 100 genes and use the default of B=10,000 permutations.
# In general, one would need a much larger number of permutations
# for microarray data.

data(golub)
smallgd<-golub[1:100,]
classlabel<-golub.cl

# Permutation unadjusted p-values and adjusted p-values
# for maxT and minP procedures with Welch t-statistics
resT<-mt.maxT(smallgd,classlabel)
resP<-mt.minP(smallgd,classlabel)
rawp<-resT$rawp[order(resT$index)]
teststat<-resT$teststat[order(resT$index)]

# Plot results and compare to Bonferroni procedure
bonf<-mt.rawp2adjp(rawp, proc=c("Bonferroni"))
allp<-cbind(rawp, bonf$adjp[order(bonf$index),2], resT$adjp[order(resT$index)], resP$adjp[order(resP$index)])

mt.plot(allp, teststat, plottype="rvsa", proc=c("rawp", "Bonferroni", "maxT", "minP"), leg=c("rawp", "Bonferroni", "maxT", "minP"))
mt.plot(allp, teststat, plottype="pvsr", proc=c("rawp", "Bonferroni", "maxT", "minP"), leg=c("rawp", "Bonferroni", "maxT", "minP"))
mt.plot(allp, teststat, plottype="pvst", proc=c("rawp", "Bonferroni", "maxT", "minP"), leg=c("rawp", "Bonferroni", "maxT", "minP"))

# Permutation adjusted p-values for minP procedure with F-statistics (like equal variance)
mt.minP(smallgd,classlabel,test="f",fixed.seed.sampling="n")

# Note that the test statistics used in the examples below are not appropriate
# for the Golub et al. data. The sole purpose of these examples is to
# demonstrate the use of the mt.maxT and mt.minP functions.

# Permutation adjusted p-values for maxT procedure with paired t-statistics
classlabel<-rep(c(0,1),19)
mt.maxT(smallgd,classlabel,test="pairt")

# Permutation adjusted p-values for maxT procedure with block F-statistics
classlabel<-rep(0:18,2)
mt.maxT(smallgd,classlabel,test="blockf",side="upper")
```


mt.plot

*Plotting results from multiple testing procedures***Description**

This function produces a number of graphical summaries for the results of multiple testing procedures and their corresponding adjusted p -values.

Usage

```
mt.plot(adjp, teststat, plottype="rvsa", logscale=FALSE, alpha=seq(0, 1, length
```

Arguments

- | | |
|----------|---|
| adjp | A matrix of adjusted p -values, with rows corresponding to hypotheses (genes) and columns to multiple testing procedures. This matrix could be obtained from the functions <code>mt.maxT</code> , <code>mt.minP</code> , or <code>mt.rawp2adjp</code> . |
| teststat | A vector of test statistics for each of the hypotheses. This vector could be obtained from the functions <code>mt.teststat</code> , <code>mt.maxT</code> , or <code>mt.minP</code> . |
| plottype | A character string specifying the type of graphical summary for the results of the multiple testing procedures.
If <code>plottype="rvsa"</code> , the number of rejected hypotheses is plotted against the nominal Type I error rate for each of the procedures given in <code>proc</code> .
If <code>plottype="pvsr"</code> , the ordered adjusted p -values are plotted for each of the procedures given in <code>proc</code> . This can be viewed as a plot of the Type I error rate against the number of rejected hypotheses.
If <code>plottype="pvst"</code> , the adjusted p -values are plotted against the test statistics for each of the procedures given in <code>proc</code> .
If <code>plottype="pvsi"</code> , the adjusted p -values are plotted for each of the procedures given in <code>proc</code> using the original data order. |
| logscale | A logical variable for the <code>pvst</code> and <code>pvsi</code> plots. If <code>logscale</code> is <code>TRUE</code> , the negative decimal logarithms of the adjusted p -values are plotted against the test statistics or gene indices. If <code>logscale</code> is <code>FALSE</code> , the adjusted p -values are plotted against the test statistics or gene indices. |
| alpha | A vector of nominal Type I error rates for the <code>rvsa</code> plot. |
| proc | A vector of character strings containing the names of the multiple testing procedures, to be used in the legend. |
| ... | Graphical parameters such as <code>col</code> , <code>lty</code> , <code>pch</code> , and <code>lwd</code> may also be supplied as arguments to the function (see <code>par</code>). |
| leg | A vector of coordinates for the legend. |

Author(s)

Sandrine Dudoit, <http://www.stat.berkeley.edu/~sandrine>,
Yongchao Ge, <yongchao.ge@mssm.edu>.

References

S. Dudoit, J. P. Shaffer, and J. C. Boldrick (Submitted). Multiple hypothesis testing in microarray experiments.

Y. Ge, S. Dudoit, and T. P. Speed. Resampling-based multiple testing for microarray data hypothesis, Technical Report \#633 of UCB Stat. <http://www.stat.berkeley.edu/~gyc>

See Also

`mt.maxT`, `mt.minP`, `mt.rawp2adjp`, `mt.reject`, `mt.teststat`, `golub`.

Examples

```
# Gene expression data from Golub et al. (1999)
# To reduce computation time and for illustrative purposes, we consider only
# the first 100 genes and use the default of B=10,000 permutations.
# In general, one would need a much larger number of permutations
# for microarray data.

data(golub)
smallgd<-golub[1:100,]
classlabel<-golub.cl

# Permutation unadjusted p-values and adjusted p-values for maxT procedure
res1<-mt.maxT(smallgd,classlabel)
rawp<-res1$rawp[order(res1$index)]
teststat<-res1$teststat[order(res1$index)]

# Permutation adjusted p-values for simple multiple testing procedures
procs<-c("Bonferroni","Holm","Hochberg","SidakSS","SidakSD","BH","BY")
res2<-mt.rawp2adjp(rawp,procs)

# Plot results from all multiple testing procedures
allp<-cbind(res2$adjp[order(res2$index)],res1$adjp[order(res1$index)])
dimnames(allp)[[2]][9]<-"maxT"
procs<-dimnames(allp)[[2]]
procs[7:9]<-c("maxT","BH","BY")
allp<-allp[,procs]

cols<-c(1:4,"orange","brown","purple",5:6)
ltypes<-c(3,rep(1,6),rep(2,2))

# Ordered adjusted p-values
mt.plot(allp,teststat,plottype="pvsr",proc=procs,leg=c(80,0.4),lty=ltypes,col=cols,lwd=2)

# Adjusted p-values in original data order
mt.plot(allp,teststat,plottype="pvs",proc=procs,leg=c(80,0.4),lty=ltypes,col=cols,lwd=2)

# Number of rejected hypotheses vs. level of the test
mt.plot(allp,teststat,plottype="rvsa",proc=procs,leg=c(0.05,100),lty=ltypes,col=cols,lwd=2)

# Adjusted p-values vs. test statistics
mt.plot(allp,teststat,plottype="pvst",logscale=TRUE,proc=procs,leg=c(0,4),pch=ltypes,col=cols,lwd=2)
```

mt.rawp2adjp *Adjusted p-values for simple multiple testing procedures*

Description

This function computes adjusted p -values for simple multiple testing procedures from a vector of raw (unadjusted) p -values. The procedures include the Bonferroni, Holm (1979), Hochberg (1988), and Sidak procedures for strong control of the family-wise Type I error rate (FWER), and the Benjamini & Hochberg (1995) and Benjamini & Yekutieli (2001) procedures for (strong) control of the false discovery rate (FDR). The less conservative adaptive Benjamini & Hochberg (2000) and two-stage Benjamini & Hochberg (2006) FDR-controlling procedures are also included.

Usage

```
mt.rawp2adjp(rawp, proc=c("Bonferroni", "Holm", "Hochberg", "SidakSS", "SidakSD",
"BH", "BY", "ABH", "TSBH"), alpha = 0.05, na.rm = FALSE)
```

Arguments

rawp A vector of raw (unadjusted) p -values for each hypothesis under consideration. These could be nominal p -values, for example, from t -tables, or permutation p -values as given in `mt.maxT` and `mt.minP`. If the `mt.maxT` or `mt.minP` functions are used, raw p -values should be given in the original data order, `rawp[order(index)]`.

proc A vector of character strings containing the names of the multiple testing procedures for which adjusted p -values are to be computed. This vector should include any of the following: "Bonferroni", "Holm", "Hochberg", "SidakSS", "SidakSD", "BH", "BY", "ABH", "TSBH".

Adjusted p -values are computed for simple FWER- and FDR- controlling procedures based on a vector of raw (unadjusted) p -values by one or more of the following methods:

Bonferroni Bonferroni single-step adjusted p -values for strong control of the FWER.

Holm Holm (1979) step-down adjusted p -values for strong control of the FWER.

Hochberg Hochberg (1988) step-up adjusted p -values for strong control of the FWER (for raw (unadjusted) p -values satisfying the Simes inequality).

SidakSS Sidak single-step adjusted p -values for strong control of the FWER (for positive orthant dependent test statistics).

SidakSD Sidak step-down adjusted p -values for strong control of the FWER (for positive orthant dependent test statistics).

BH Adjusted p -values for the Benjamini & Hochberg (1995) step-up FDR-controlling procedure (independent and positive regression dependent test statistics).

BY Adjusted p -values for the Benjamini & Yekutieli (2001) step-up FDR-controlling procedure (general dependency structures).

ABH Adjusted p -values for the adaptive Benjamini & Hochberg (2000) step-up FDR-controlling procedure. This method amends the original step-up procedure using an estimate of the number of true null hypotheses obtained from p -values.

	TSBH	Adjusted p -values for the two-stage Benjamini & Hochberg (2006) step-up FDR-controlling procedure. This method amends the original step-up procedure using an estimate of the number of true null hypotheses obtained from a first-pass application of "BH". The adjusted p -values are α -dependent, therefore <code>alpha</code> must be set in the function arguments when using this procedure.
<code>alpha</code>		A nominal type I error rate, or a vector of error rates, used for estimating the number of true null hypotheses in the two-stage Benjamini & Hochberg procedure ("TSBH"). Default is 0.05.
<code>na.rm</code>		An option for handling NA values in a list of raw p -values. If <code>FALSE</code> , the number of hypotheses considered is the length of the vector of raw p -values. Otherwise, if <code>TRUE</code> , the number of hypotheses is the number of raw p -values which were not NAs.

Value

A list with components:

<code>adjp</code>	A matrix of adjusted p -values, with rows corresponding to hypotheses and columns to multiple testing procedures. Hypotheses are sorted in increasing order of their raw (unadjusted) p -values.
<code>index</code>	A vector of row indices, between 1 and <code>length(rawp)</code> , where rows are sorted according to their raw (unadjusted) p -values. To obtain the adjusted p -values in the original data order, use <code>adjp[order(index),]</code> .
<code>h0.ABH</code>	The estimate of the number of true null hypotheses as proposed by Benjamini & Hochberg (2000) used when computing adjusted p -values for the "ABH" procedure (see Dudoit et al., 2007).
<code>h0.TSBH</code>	The estimate (or vector of estimates) of the number of true null hypotheses as proposed by Benjamini et al. (2006) when computing adjusted p -values for the "TSBH" procedure. (see Dudoit et al., 2007).

Author(s)

Sandrine Dudoit, <http://www.stat.berkeley.edu/~sandrine>,
 Yongchao Ge, <yongchao.ge@mssm.edu>,
 Houston Gilbert, <http://www.stat.berkeley.edu/~houston>.

References

- Y. Benjamini and Y. Hochberg (1995). Controlling the false discovery rate: a practical and powerful approach to multiple testing. *J. R. Statist. Soc. B*. Vol. 57: 289-300.
- Y. Benjamini and Y. Hochberg (2000). On the adaptive control of the false discovery rate in multiple testing with independent statistics. *J. Behav. Educ. Statist.* Vol 25: 60-83.
- Y. Benjamini and D. Yekutieli (2001). The control of the false discovery rate in multiple hypothesis testing under dependency. *Annals of Statistics*. Vol. 29: 1165-88.
- Y. Benjamini, A. M. Krieger and D. Yekutieli (2006). Adaptive linear step-up procedures that control the false discovery rate. *Biometrika*. Vol. 93: 491-507.

S. Dudoit, J. P. Shaffer, and J. C. Boldrick (2003). Multiple hypothesis testing in microarray experiments. *Statistical Science*. Vol. 18: 71-103.

S. Dudoit, H. N. Gilbert, and M. J. van der Laan (2008). Resampling-based empirical Bayes multiple testing procedures for controlling generalized tail probability and expected value error rates: Focus on the false discovery rate and simulation study. *Biometrical Journal*, 50(5):716-44. <http://www.stat.berkeley.edu/~houston/BJMCPSupp/BJMCPSupp.html>.

Y. Ge, S. Dudoit, and T. P. Speed (2003). Resampling-based multiple testing for microarray data analysis. *TEST*. Vol. 12: 1-44 (plus discussion p. 44-77).

Y. Hochberg (1988). A sharper Bonferroni procedure for multiple tests of significance, *Biometrika*. Vol. 75: 800-802.

S. Holm (1979). A simple sequentially rejective multiple test procedure. *Scand. J. Statist.*. Vol. 6: 65-70.

See Also

[mt.maxT](#), [mt.minP](#), [mt.plot](#), [mt.reject](#), [golub](#).

Examples

```
# Gene expression data from Golub et al. (1999)
# To reduce computation time and for illustrative purposes, we consider only
# the first 100 genes and use the default of B=10,000 permutations.
# In general, one would need a much larger number of permutations
# for microarray data.

data(golub)
smallgd<-golub[1:100,]
classlabel<-golub.cl

# Permutation unadjusted p-values and adjusted p-values for maxT procedure
res1<-mt.maxT(smallgd,classlabel)
rawp<-res1$rawp[order(res1$index)]

# Permutation adjusted p-values for simple multiple testing procedures
procs<-c("Bonferroni","Holm","Hochberg","SidakSS","SidakSD","BH","BY","ABH","TSBH")
res2<-mt.rawp2adjp(rawp,procs)
```

mt.reject

Identity and number of rejected hypotheses

Description

This function returns the identity and number of rejected hypotheses for several multiple testing procedures and different nominal Type I error rates.

Usage

```
mt.reject(adjp, alpha)
```

Arguments

adjp	A matrix of adjusted p -values, with rows corresponding to hypotheses and columns to multiple testing procedures. This matrix could be obtained from the function <code>mt.rawp2adjp</code> .
alpha	A vector of nominal Type I error rates.

Value

A list with components

r	A matrix containing the number of rejected hypotheses for several multiple testing procedures and different nominal Type I error rates. Rows correspond to Type I error rates and columns to multiple testing procedures.
which	A matrix of indicators for the rejection of individual hypotheses by different multiple testing procedures for a nominal Type I error rate <code>alpha[1]</code> . Rows correspond to hypotheses and columns to multiple testing procedures.

Author(s)

Sandrine Dudoit, <http://www.stat.berkeley.edu/~sandrine>,
Yongchao Ge, <yongchao.ge@mssm.edu>.

See Also

`mt.maxT`, `mt.minP`, `mt.rawp2adjp`, `golub`.

Examples

```
# Gene expression data from Golub et al. (1999)
# To reduce computation time and for illustrative purposes, we consider only
# the first 100 genes and use the default of B=10,000 permutations.
# In general, one would need a much larger number of permutations
# for microarray data.

data(golub)
smallgd<-golub[1:100,]
classlabel<-golub.cl

# Permutation unadjusted p-values and adjusted p-values for maxT procedure
res<-mt.maxT(smallgd,classlabel)
mt.reject(cbind(res$rawp,res$adjp),seq(0,1,0.1))$r
```

mt.sample.teststat *Permutation distribution of test statistics and raw (unadjusted) p-values*

Description

These functions provide tools to investigate the permutation distribution of test statistics, raw (unadjusted) p -values, and class labels.

Usage

```
mt.sample.teststat(V, classlabel, test="t", fixed.seed.sampling="y", B=10000, na=.mt.
mt.sample.rawp(V, classlabel, test="t", side="abs", fixed.seed.sampling="y", B=10000,
mt.sample.label(classlabel, test="t", fixed.seed.sampling="y", B=10000)
```

Arguments

<code>V</code>	A numeric vector containing the data for one of the variables (genes).
<code>classlabel</code>	A vector of integers corresponding to observation (column) class labels. For k classes, the labels must be integers between 0 and $k - 1$. For the <code>blockf</code> test option, observations may be divided into n/k blocks of k observations each. The observations are ordered by block, and within each block, they are labeled using the integers 0 to $k - 1$.
<code>test</code>	A character string specifying the statistic to be used to test the null hypothesis of no association between the variables and the class labels. If <code>test="t"</code> , the tests are based on two-sample Welch t-statistics (unequal variances). If <code>test="t.equalvar"</code> , the tests are based on two-sample t-statistics with equal variance for the two samples. The square of the t-statistic is equal to an F-statistic for $k = 2$. If <code>test="wilcoxon"</code> , the tests are based on standardized rank sum Wilcoxon statistics. If <code>test="f"</code> , the tests are based on F-statistics. If <code>test="pairt"</code> , the tests are based on paired t-statistics. The square of the paired t-statistic is equal to a block F-statistic for $k = 2$. If <code>test="blockf"</code> , the tests are based on F-statistics which adjust for block differences (cf. two-way analysis of variance).
<code>side</code>	A character string specifying the type of rejection region. If <code>side="abs"</code> , two-tailed tests, the null hypothesis is rejected for large absolute values of the test statistic. If <code>side="upper"</code> , one-tailed tests, the null hypothesis is rejected for large values of the test statistic. If <code>side="lower"</code> , one-tailed tests, the null hypothesis is rejected for small values of the test statistic.
<code>fixed.seed.sampling</code>	If <code>fixed.seed.sampling="y"</code> , a fixed seed sampling procedure is used, which may double the computing time, but will not use extra memory to store the permutations. If <code>fixed.seed.sampling="n"</code> , permutations will be stored in memory. For the <code>blockf</code> test, the option <code>n</code> was not implemented as it requires too much memory.
<code>B</code>	The number of permutations. For a complete enumeration, <code>B</code> should be 0 (zero) or any number not less than the total number of permutations.
<code>na</code>	Code for missing values (the default is <code>.mt.naNUM=--93074815.62</code>). Entries with missing values will be ignored in the computation, i.e., test statistics will be based on a smaller sample size. This feature has not yet fully implemented.
<code>nonpara</code>	If <code>nonpara="y"</code> , nonparametric test statistics are computed based on ranked data. If <code>nonpara="n"</code> , the original data are used.

Value

For `mt.sample.teststat`, a vector containing B permutation test statistics.

For `mt.sample.rawp`, a vector containing B permutation unadjusted p -values.

For `mt.sample.label`, a matrix containing B sets of permuted class labels. Each row corresponds to one permutation.

Author(s)

Yongchao Ge, <yongchao.ge@mssm.edu>,
Sandrine Dudoit, <http://www.stat.berkeley.edu/~sandrine>.

See Also

`mt.maxT`, `mt.minP`, `golub`.

Examples

```
# Gene expression data from Golub et al. (1999)
data(golub)

mt.sample.label(golub.cl, B=10)

permt<-mt.sample.teststat(golub[1,], golub.cl, B=1000)
qqnorm(permt)
qqline(permt)

permt<-mt.sample.teststat(golub[50,], golub.cl, B=1000)
qqnorm(permt)
qqline(permt)

permp<-mt.sample.rawp(golub[1,], golub.cl, B=1000)
hist(permp)
```

mt.teststat

Computing test statistics for each row of a data frame

Description

These functions provide a convenient way to compute test statistics, e.g., two-sample Welch t -statistics, Wilcoxon statistics, F -statistics, paired t -statistics, block F -statistics, for each row of a data frame.

Usage

```
mt.teststat(X, classlabel, test="t", na=.mt.naNUM, nonpara="n")
mt.teststat.num.denum(X, classlabel, test="t", na=.mt.naNUM, nonpara="n")
```


Arguments

<code>X</code>	A data frame or matrix, with m rows corresponding to variables (hypotheses) and n columns to observations. In the case of gene expression data, rows correspond to genes and columns to mRNA samples. The data can be read using read.table .
<code>classlabel</code>	A vector of integers corresponding to observation (column) class labels. For k classes, the labels must be integers between 0 and $k - 1$. For the <code>blockf</code> test option, observations may be divided into n/k blocks of k observations each. The observations are ordered by block, and within each block, they are labeled using the integers 0 to $k - 1$.
<code>test</code>	A character string specifying the statistic to be used to test the null hypothesis of no association between the variables and the class labels. If <code>test="t"</code> , the tests are based on two-sample Welch t-statistics (unequal variances). If <code>test="t.equalvar"</code> , the tests are based on two-sample t-statistics with equal variance for the two samples. The square of the t-statistic is equal to an F-statistic for $k = 2$. If <code>test="wilcoxon"</code> , the tests are based on standardized rank sum Wilcoxon statistics. If <code>test="f"</code> , the tests are based on F-statistics. If <code>test="pairt"</code> , the tests are based on paired t-statistics. The square of the paired t-statistic is equal to a block F-statistic for $k = 2$. If <code>test="blockf"</code> , the tests are based on F-statistics which adjust for block differences (cf. two-way analysis of variance).
<code>na</code>	Code for missing values (the default is <code>.mt.naNUM=--93074815.62</code>). Entries with missing values will be ignored in the computation, i.e., test statistics will be based on a smaller sample size. This feature has not yet fully implemented.
<code>nonpara</code>	If <code>nonpara="y"</code> , nonparametric test statistics are computed based on ranked data. If <code>nonpara="n"</code> , the original data are used.

Value

For `mt.teststat`, a vector of test statistics for each row (gene).

For `mt.teststat.num.denum`, a data frame with

`teststat.num` the numerator of the test statistics for each row, depending on the specific `test` option.

`teststat.denum` the denominator of the test statistics for each row, depending on the specific `test` option.

Author(s)

Yongchao Ge, <yongchao.ge@mssm.edu>,
Sandrine Dudoit, <http://www.stat.berkeley.edu/~sandrine>.

See Also

[mt.maxT](#), [mt.minP](#), [golub](#).

Examples

```
# Gene expression data from Golub et al. (1999)
data(golub)

teststat<-mt.teststat(golub,golub.cl)
qqnorm(teststat)
qqline(teststat)

tmp<-mt.teststat.num.denum(golub,golub.cl,test="t")
num<-tmp$teststat.num
denum<-tmp$teststat.denum
plot(sqrt(denum),num)

tmp<-mt.teststat.num.denum(golub,golub.cl,test="f")
```

Index

- *Topic **classes**
 - EBMTP-class, 1
 - MTP-class, 13
- *Topic **datasets**
 - golub, 29
- *Topic **hplot**
 - mt.plot, 33
- *Topic **htest**
 - mt.maxT, 30
 - mt.rawp2adjp, 35
 - mt.reject, 37
 - MTP, 20
- *Topic **manip**
 - mt.sample.teststat, 38
- *Topic **methods**
 - MTP-methods, 18
- *Topic **univar**
 - mt.teststat, 40
- [, EBMTP-method (*MTP-methods*), 18
- [, MTP-method (*MTP-methods*), 18
- [-methods, 5, 17
- [-methods (*MTP-methods*), 18

- ABH.h0 (*Hsets*), 10
- as.list (*MTP-methods*), 18
- as.list, EBMTP-method (*MTP-methods*), 18
- as.list, MTP-method (*MTP-methods*), 18
- as.list-methods, 5, 17
- as.list-methods (*MTP-methods*), 18

- dens.est (*Hsets*), 10

- EBMTP, 5, 6, 13, 17, 28
- EBMTP-class, 10, 13
- EBMTP-methods, 5, 10, 13, 17
- EBMTP-class, 1
- EBMTP-method (*EBMTP-class*), 1
- EBMTP-methods (*MTP-methods*), 18
- ebmtp2mtp, 5, 17
- ebmtp2mtp (*MTP-methods*), 18
- ebmtp2mtp, EBMTP-method (*MTP-methods*), 18

- ebmtp2mtp-methods (*MTP-methods*), 18

- EBupdate (*MTP-methods*), 18
- EBupdate, EBMTP-method (*MTP-methods*), 18
- EBupdate-methods (*MTP-methods*), 18

- fwcr2gfwcr, 28

- G.VS (*Hsets*), 10
- golub, 29, 32, 34, 37, 38, 40, 42

- Hsets, 10, 10

- mt.maxT, 28, 30, 33, 34, 37, 38, 40, 42
- mt.minP, 28, 31, 33, 34, 37, 38, 40, 42
- mt.minP (*mt.maxT*), 30
- mt.plot, 32, 33, 37
- mt.rawp2adjp, 32–34, 35, 38
- mt.reject, 32, 34, 37, 37
- mt.sample.label, 40
- mt.sample.label (*mt.sample.teststat*), 38
- mt.sample.rawp, 40
- mt.sample.rawp (*mt.sample.teststat*), 38
- mt.sample.teststat, 32, 38, 40
- mt.teststat, 32–34, 40, 41
- mt.teststat.num.denum, 41
- MTP, 5, 6, 10, 17, 20
- MTP-class, 28
- MTP-methods, 5, 17, 28
- MTP-class, 13
- MTP-methods, 18
- mtp2ebmtp, 5, 17
- mtp2ebmtp (*MTP-methods*), 18
- mtp2ebmtp, MTP-method (*MTP-methods*), 18
- mtp2ebmtp-methods (*MTP-methods*), 18

- par, 33
- plot (*MTP-methods*), 18
- plot, EBMTP-method (*MTP-methods*), 18

plot, MTP-method (*MTP-methods*), 18
plot-methods, 5, 17
plot-methods (*MTP-methods*), 18
print, EBMTTP-method (*MTP-methods*),
18
print, MTP-method (*MTP-methods*), 18
print-methods, 5, 17
print-methods (*MTP-methods*), 18
print.MTP (*MTP-methods*), 18

read.table, 30, 41

ss.maxT, 28
summary (*MTP-methods*), 18
summary, EBMTTP-method
(*MTP-methods*), 18
summary, MTP-method (*MTP-methods*),
18
summary-methods, 5, 17
summary-methods (*MTP-methods*), 18

update (*MTP-methods*), 18
update, MTP-method (*MTP-methods*),
18
update-methods (*MTP-methods*), 18

VScout (*Hsets*), 10