Package 'radjust'

August 31, 2018

Type Package

Title Replicability Adjusted p-Values for Two Independent Studies with Multiple Endpoints

Version 0.1.0

Date 2018-08-25

Description Calculates adjusted p-values for the null hypothesis of no replicability across studies for two study designs: (i) a primary and follow-up study, where the features in the follow-up study are selected from the primary study, as described in Bogomolov and Heller (2013) <doi:10.1080/01621459.2013.829002> and Heller, Bogomolov and Benjamini (2014) <doi:10.1073/pnas.1314814111>; (ii) two independent studies, where the features for replicability are first selected in each study separately, as described in Bogomolov and Heller (2018) <doi:10.1093/biomet/asy029>. The latter design is the one encountered in a typical meta-analysis of two studies, but the inference is for replicability rather than for identifying the features that are non-null in at least one study.

License GPL-3

Encoding UTF-8

LazyData true

Depends R (>= 2.10)

Suggests covr, testthat

RoxygenNote 6.1.0

NeedsCompilation no

Author Shay Yaacoby [aut], Marina Bogomolov [aut], Ruth Heller [aut, cre]

Maintainer Ruth Heller <ruheller@gmail.com>

Repository CRAN

Date/Publication 2018-08-31 17:50:14 UTC

R topics documented:

`crohn` *p-values of 126 SNPs followed from a primary study to a follow-up study for testing their association with Crohn's disease.*

Description

To discover the associations between SNPs and Crohn's disease, 635547 SNPs were examined in a primary study. For follow-up, 126 SNPs were measured in an independent study. The criteria for follow-up were as follows: the two smallest p-values in each distinct region with primary study p-values below 0.00005.

Usage

```
crohn
```

Format

A data frame with 126 rows and 3 columns:

`index`	integer	just the row number.
`pv1`	numeric	p-value from study 1.
`pv2`	numeric	p-value from study 2.

Source

Barrett, Jeffrey C., et al. "Genome-wide association defines more than 30 distinct susceptibility loci for Crohn's disease." Nature genetics 40.8 (2008): 955.

References

Bogomolov, M. and Heller, R. (2013). Discovering findings that replicate from a primary study of high dimension to a follow-up study. Journal of the American Statistical Association, Vol. 108, No. 504, Pp. 1480-1492.

See Also

the example in `radjust_pf` uses this data.

`mice` *p-values of 29 Behavioural Measures in Two Studies*

Description

In different laboratories, the comparison of behaviours of the same two strains of mice may lead to opposite conclusions that are both statistically significant. An explanation may be the different laboratory environment, i.e., personnel, equipment, or measurement techniques, affecting differently the study strains. This data set provides the p-values for testing the association of mice strain with 29 behavioural measures from five commonly used behavioural tests in two laboratories: the laboratory of H. Wurbel at the University of Giessen, and the laboratory of P. Gass at the Central Institute of Mental Health, Mannheim. The data table contains two-sided p-values. To transform all the two-sided p-values to one sided in the same direction, see the example in `radjust_sym`.

Usage

```
mice
```

Format

A data frame with 29 rows and 5 columns:

`feature_name`	char.	The name of the measure and the test, concatenated.
`twosided_pv1`	numeric	the *two-sided* p-value from study 1.
`twosided_pv2`	numeric	the *two-sided* p-value from study 2.
`dir_is_left1`	logical	whether the direction of the test statistic from study 1 is *left*.
`dir_is_left2`	logical	whether the direction of the test statistic from study 2 is *right*.

Source

Richter, S. Helene, et al. "Effect of population heterogenization on the reproducibility of mouse behavior: a multi-laboratory study." PLoS One 6.1 (2011): e16461.

References

Bogomolov, M. and Heller, R. (2018). Assessing replicability of findings across two studies of multiple features. Biometrika.

See Also

`radjust_sym`

`radjust` *radjust: Replicability Adjusted p-values for Two Independent Studies with Multiple Endpoints*

Description

This package provides the adjusted p-values for the null hypothesis of no replicability across studies for two study designs: a primary and follow-up study, where the features in the follow-up study are selected from the primary study; two independent studies, where the features for replicability are first selected in each study separately. The latter design is the one encountered in typical meta-analysis of two studies, but the inference is for replicability rather than for identifying the features that are nonnull in at least one study.

`radjust_pf` *Adjust p-values for Replicability across Two Independent, Primary and Follow-up, Studies with Multiple Endpoints*

Description

Given two vectors of p-values from the primary and follow-up studies, returns the adjusted p-values for false discovery rate control on replicability claims. The p-value vectors are only for features selected for follow-up.

Usage

```
radjust_pf(pv1, pv2, m, c2 = 0.5, l00 = 0, variant = c("none",
  "general_dependency", "use_threshold"), threshold = NULL,
  alpha = 0.05)
```

Arguments

`pv1`	numeric vector of p-values from the primary study which corresponds to the p-values from the follow-up study (`pv2`).
`pv2`	numeric vector of p-values from the follow-up study.
`m`	the number of features examined in the primary study (`> length(pv1)`).
`c2`	the relative boost to the p-values from the **follow-up** study. `c2 = 0.5` (the default) is recommended. It was observed in simulations to yield similar power to procedure with the optimal value (which is unknown for real data).
`l00`	a lower bound of the fraction of features (out of m) with true null hypotheses in both studies. For example, for GWAS on the whole genome, the choice of 0.8 is conservative in typical applications.
`variant`	**none** the default. **general_dependency** use $m^* = m \sum_{i=1}^{m} \frac{1}{i}$ instead of `m`.

use_threshold c1 is computed given the `threshold` value.

Both variants guarantee that the procedure that declares as replicated all features with r-values below `alpha`, controls the FDR at level `alpha`, for any type of dependency of the p-values in the primary study.

`threshold` the selection threshold for p-values from the primary study; must be supplied when variant 'use_threshold' is selected, otherwise ignored.

`alpha` The FDR level to control.

Details

When many hypotheses are simultaneously examined in a primary study, and then a subset of hypotheses are forwarded for follow-up in an independent study, it is of interest to know which findings are replicated across studies. As a measure of replicability of significance, we compute the r-value, i.e. the FDR adjusted replicability p-value, for each hypothesis followed-up. This measure is different than the FDR adjusted p-value in a typical meta-analysis, where a p-value close to zero in one of the studies is enough to declare the finding as highly significant. The FDR r-value for a feature is the smallest FDR level at which we can say that the finding is among the replicated ones.

Value

vector of length of `pv1` and `pv2`, containing the r-values.

Note

The function is also available as a web applet: http://www.math.tau.ac.il/~ruheller/App.html

References

Bogomolov, M. and Heller, R. (2013). Discovering findings that replicate from a primary study of high dimension to a follow-up study. Journal of the American Statistical Association, Vol. 108, No. 504, Pp. 1480-1492.

Heller, R., Bogomolov, M., & Benjamini, Y. (2014). Deciding whether follow-up studies have replicated findings in a preliminary large-scale omics study. Proceedings of the National Academy of Sciences of the United States of America, Vol. 111, No. 46, Pp. 16262–16267.

See Also

`radjust_sym` for replicability analysis in two symmetric studies.

Examples

```
data(crohn)
rv  <- radjust_pf(pv1 = crohn$pv1, pv2 = crohn$pv1, m = 635547, l00 = 0.8)
rv2 <- radjust_pf(pv1 = crohn$pv1, pv2 = crohn$pv1, m = 635547, l00 = 0.8,
                  variant="use_threshold",threshold = 1e-5)
```

radjust_sym *Adjust p-values for Replicability across Two Independent Studies with Multiple Endpoints*

Description

Given two vectors of p-values from two independent studies, returns the adjusted p-values for false discovery rate control on replicability claims.

Usage

```
radjust_sym(pv1, pv2, w1 = 0.5, input_type = c("selected_features",
  "all_features"), general_dependency = FALSE,
  directional_rep_claim = FALSE,
  variant = c("non-adaptive-with-alpha-selection", "adaptive",
  "non-adaptive"), alpha = if (variant == "non-adaptive") NULL else 0.05)
```

Arguments

pv1, pv2	numeric vectors of p-values. If directional_rep_claim is TRUE, they must be left-sided p-values. Can be either the p-values for the selected features from each study (the default input type), or the p-values for all the features from each study. Can be either of the same length (so the same location in each vector corresponds to the same feature) or with names (so the same name in each vector corresponds to the same feature).
w1	fraction between zero and one, of the relative weight for the p-values from study 1. Default value is 0.5 (see Details for other values).
input_type	whether pv1 and pv2 contain all the p-values from each study or only the selected ones (the default).
general_dependency	TRUE or FALSE, indicating whether to correct for general dependency. The recommended default value is FALSE (see Details).
directional_rep_claim	TRUE or FALSE, indicating whether to perform directional replicability analysis. The default value is FALSE. If TRUE, pv1 and pv2 should be left-sided p-values (see Details).
variant	A character string specifying the chosen variant for a potential increase in the number of discoveries. Must be one of "non-adaptive-with-alpha-selection" (default), "adaptive", or "non-adaptive" (see Details).
alpha	The threshold on p-values for selecting the features in each study and the significance level for replicability analysis (see Details).

Details

For FDR control at level α on replicability claims, all features with r**-value** at most α are declared as replicated. In addition, the discoveries from study 1 among the replicability claims have an FDR control guarantee at level $w_1\alpha$. Similarly, the discoveries from study 2 among the replicability claims have an FDR control guarantee at level $(1-w_1)\alpha$.

Setting `w1` to a value different than half is appropriate for stricter FDR control in one of the studies. For example, if study two has a much larger sample size than study one (and both studies examine the same problem), then setting $w_1 > 0.5$ will provide a stricter FDR control for the larger study and greater power for the replicability analysis, see Bogomolov and Heller (2018) for details.

The theoretical FDR control guarantees assume independence within each vector of p-values. However, empirical investigations suggest that the method is robust to deviations from independence. In practice, we recommend using it whenever the Benjamini-Hochberg procedure is appropriate for use with single studies, as this procedure can be viewed as a two-dimensional Benjamini-Hochberg procedure which enjoys similar robustness properties. For general dependence, we provide the option to apply a more conservative procedure with theoretical FDR control guarantee for any type of dependence, by setting `general_dependency` to TRUE.

If `variant` is `"non-adaptive"`, then the non-adaptive replicability analysis procedure of Bogomolov and Heller (2018) is applied on the input p-values `pv1` and `pv2`. If `variant` is `"non-adaptive-with-alpha-selectio` then for a user specified `alpha` (default 0.05) only p-values from study one below $w_1\alpha$ and from study two below $(1-w_1)\alpha$ are considered for replicability analysis. This additional step prevents including in the selected sets features that cannot be discovered as replicability claims at the nominal FDR level α, thus reducing the multiplicity adjustment necessary for replicability analysis. If `variant` is `"adaptive"`, then for a user specified `alpha` the adaptive replicability analysis procedure is applied on the dataset, see Bogomolov and Heller (2018) for details.

The meaning of the replicability claim for a feature if `directional_rep_claim` is `FALSE`, is that both null hypotheses are false (or both alternatives are true). Setting `directional_rep_claim` to `TRUE` is useful if the discoveries of interest are directional but the direction is unknown. For example, a directional replicability claim for a feature is the claim that both associations examined for it are positive, or both associations examined for it are negative, but not that one association is positive and the other negative. For directional replicability analysis, the input p-values `pv1` and `pv2` should be the left-sided input p-values (left-sided is the choice without loss of generality, since we assume the left and right sided p-values sum to one for each null hypothesis).

Value

The function returns a list with the following elements:

`call`	the function call.
`inputs`	a list with the function's input parameters (except `pv1` and `pv2`).
`results_table`	a data frame with the features selected in both studies and their r-values (see description below).
`selected1`	the features selected in study 1 (when the variant is either `"adaptive"` or `"non-adaptive-with-alpha-se`
`selected2`	the features selected in study 2, same as above.
`n_selected1`	the number of selected features in study 1.
`n_selected2`	the number of selected features in study 2.
`pi1`	the estimate of the true-nulls fraction in the study 1 among the selected in study 2, when `variant = "adap`
`pi2`	the estimate of the true-nulls fraction in the study 2 among the selected in study 1, when `variant = "adap`

The third element in the list, `results_table`, includes the following columns:

`name`	char.	the name of the feature as extracted from the named vectors, or the location, if the input vectors ar
`p.value.1`	numeric	the one-sided p-value from study 1 as inputed (denoted by `pv1`). When `directional_rep_claim=`
`p.value.2`	numeric	the one-sided p-value from study 2 as inputed (denoted by `pv2`). When `directional_rep_claim=`
`r.value`	numeric	the replicability adjusted p-value (= r-value).
`Direction`	char.	the direction of the replicability claim, when `directional_rep_claim = TRUE`.
`Significant`	char.	indicates for which features replicability claims can be made at level α, when `variant` is set to "a

References

Bogomolov, M. and Heller, R. (2018). Assessing replicability of findings across two studies of multiple features. Biometrika.

See Also

`radjust_pf` for replicability analysis in primary and follow-up studies.

Examples

```
data(mice)
## transform the two-sided p-values to one-sided in the same direction (left):
## (we use the direction of the test statistic to do so and assume that it is continuous)

pv1 <- ifelse(mice$dir_is_left1, mice$twosided_pv1/2, 1-mice$twosided_pv1/2)
pv2 <- ifelse(mice$dir_is_left2, mice$twosided_pv2/2, 1-mice$twosided_pv2/2)

## run the examples as in the article:

mice_rv_adaptive <- radjust_sym(pv1, pv2, input_type = "all", directional_rep_claim = TRUE,
                                variant = "adaptive", alpha=0.05)
print(mice_rv_adaptive)

mice_rv_non_adpt_sel <- radjust_sym(pv1, pv2, input_type = "all", directional_rep_claim = TRUE,
                               variant = "non-adaptive-with-alpha-selection", alpha=0.05)
print(mice_rv_non_adpt_sel)

mice_rv_non_adpt <- radjust_sym(pv1, pv2, input_type = "selected", directional_rep_claim = TRUE,
                                variant = "non-adaptive")
print(mice_rv_non_adpt)
```

Index