

Package ‘BayesCombo’

July 21, 2025

Title Bayesian Evidence Combination

Version 1.0

Date 2017-02-06

Author Bruno Contrino, Stanley E. Lazic

Maintainer Stanley E. Lazic <stan.lazic@cantab.net>

Description Combine diverse evidence across multiple studies to test a high level scientific theory. The methods can also be used as an alternative to a standard meta-analysis.

Depends R (>= 2.10),

License GPL-3

LazyData true

Suggests knitr, rmarkdown, labstats, testthat

URL <https://github.com/stanlazic/BayesCombo>

BugReports <https://github.com/stanlazic/BayesCombo/issues>

VignetteBuilder knitr

RoxygenNote 5.0.1

NeedsCompilation no

Repository CRAN

Date/Publication 2017-02-08 12:20:17

Contents

calc.post.beta	2
ev.combo	2
forestplot	4
plot.EV	5
plot.PPH	6
pph	7
summary.PPH	8

Index	10
--------------	-----------

calc.post.beta	<i>Utility functions</i>
----------------	--------------------------

Description

Not called directly by user.

Usage

```
calc.post.beta(beta, se.beta, beta0, se0)
```

```
calc.post.se(se.beta, se0)
```

```
prior.se(beta, se.beta, ci)
```

Arguments

beta	Effect size.
se.beta	Standard error of the effect.
beta0	Prior for the effect size.
se0	Prior standard error.
ci	A numeric value that is used to calculate the prior standard error.

ev.combo	<i>Calculates the posterior probability of hypotheses for multiple studies</i>
----------	--

Description

The function takes multiple effect sizes and a their standard errors and calculates the posterior probability for each hypothesis (H<: the effect size is less than 0, H0: the effect size is zero, or H>: the effect size is greater than zero).

Usage

```
ev.combo(beta, se.beta, beta0 = 0, ci = 99, H0 = c(0, 0), scale = FALSE,
  H.priors = rep(1/3, 3), se.mult = 1, adjust = FALSE, epsilon = 1e-06,
  adj.factor = 1e-04, ...)
```

Arguments

beta	Effect size.
se.beta	Standard error for the effect.
beta0	A prior value for the effect size. Default is zero.
ci	Is used to calculate the prior standard error if $se0 = \text{NULL}$. The default value of 99 calculates the prior standard error so that the 99 largest (furthest from zero) confidence interval of the data distribution.
H0	A vector of length two that defines the null hypothesis. If the values are identical (e.g. $H0 = c(0, 0)$) a point null is used, otherwise the null is defined as the range between the lower and upper value.
scale	Logical. Whether to scale the effect size by its standard error. Standardising has no effect on the calculations but standardised effect sizes may be easier to compare in a forest plot.
H.priors	Prior hypothesis probabilities; default is an equal probability of 1/3, and they are specified in the following order: $H<0$, $H=0$, $H>0$.
se.mult	Standard error multiplier used to increase or decrease the prior SE and used in conjunction with ci when $se0 = \text{NULL}$.
adjust	Logical. Whether to adjust very small posterior hypothesis probabilities. Adjusting prevents a single study from having too much influence on the results when combining multiple studies. For example, if the probability for a hypothesis from one study is zero, then additional studies cannot alter this probability (multiplying anything by zero is still zero).
epsilon	A small value that a posterior hypothesis probability must fall below before an adjustment is made. Ignored if <code>adjust = FALSE</code> .
adj.factor	A small number added to each posterior hypothesis probability if <code>adjust = TRUE</code> and one of the posterior hypothesis probabilities is less than epsilon. The PPHs are then re-scaled to sum to one.
...	Options to be passed to <code>pph()</code> .

Details

This function calls `pph()` once for each study to be combined, where the posterior probabilities for one study are used as the priors for the next study. One exception is that values for $se0$ are ignored as they are calculated automatically.

Value

Object of class EV which contains a matrix of posterior probabilities for each updated step and other calculated values.

Examples

```
x <- ev.combo(beta = c(0.0126, 5.0052, 1.2976, 0.0005),
              se.beta = c(0.050, 2.581, 2.054, 0.003) )
x
plot(x)
```

`forestplot`*Forest plot*

Description

Plots the effect sizes and standard errors for a set of studies as a standard forest plot. Also shown are the priors for each study.

Usage

```
forestplot(x, range = NULL, xlab = "Effect size", ylab = "Study", ...,
           labels = NULL)
```

Arguments

<code>x</code>	An object of the class EV.
<code>range</code>	A numeric vector of length two that sets the range of the x-axis. Values are passed to <code>xlim</code> . Sensible default values are used when <code>range = NULL</code> .
<code>xlab, ylab</code>	Text for the x and y axes labels.
<code>...</code>	Other options passed to <code>plot()</code> .
<code>labels</code>	Names for studies to be plotted on the y-axis. If <code>NULL</code> , then numbers are used.

Details

This function plots results from the `ev.combo()` function. The results may be hard to visualise if the effect sizes differ greatly, and it may be preferable to standardise the effect sizes by setting `scale = TRUE` in the `ev.combo()` function.

Value

A forest plot.

See Also

[ev.combo](#)

Examples

```
x <- ev.combo( beta = c(0.0126, 5.0052, 1.2976, 0.0005),
              se.beta = c(0.050, 2.581, 2.054, 0.003) )
forestplot(x)
```

plot.EV *Plot of posterior probabilities for each hypothesis*

Description

Shows how the posterior probabilities evolve as each study is added to the analysis.

Usage

```
## S3 method for class 'EV'  
plot(x, leg.loc = "topleft", lty = 1, ...)
```

Arguments

x	An object of the class EV created by the <code>ev.combo()</code> function.
leg.loc	Location of the legend; default is top left. See the <code>legend()</code> help pages for all the options.
lty	An vector of integers that specifies the line types. See the <code>matplot()</code> function for details.
...	Other options passed to <code>matplot()</code> .

Details

Each line in the graph is a hypothesis and the lines start at the prior hypothesis probabilities (usually 1/3).

Value

Plot of posterior probabilities for each hypothesis.

See Also

[ev.combo](#)

Examples

```
x <- ev.combo(beta = c(0.0126, 5.0052, 1.2976, 0.0005),  
             se.beta = c(0.050, 2.581, 2.054, 0.003) )  
plot(x)
```

`plot.PPH`*Plot of prior, likelihood, and posterior distributions*

Description

Plots the prior, likelihood, and posterior distribution from a single study.

Usage

```
## S3 method for class 'PPH'  
plot(x, range = NULL, n = 200, leg.loc = "topleft",  
      xlab = "Effect size", ylab = "", ...)
```

Arguments

<code>x</code>	A PPH object created by the <code>pph()</code> function.
<code>range</code>	A numeric vector of length two that sets the range of the x-axis. Values are passed to <code>xlim</code> . Sensible default values are used when <code>range = NULL</code> .
<code>n</code>	An integer that specifies the number of x-values to use when plotting the distributions.
<code>leg.loc</code>	Location of the legend; default is top left. See the <code>legend()</code> help pages for all the options. If <code>leg.loc = NULL</code> the legend is not plotted.
<code>xlab, ylab</code>	Text for the x and y axes labels.
<code>...</code>	Other options passed to <code>plot()</code> .

Details

Plots the prior, likelihood (data), and posterior distribution calculated from the `pph()` function for a single study.

Value

Plot of distributions.

See Also

[pph](#)

Examples

```
x <- pph(beta = 5.005, se.beta = 2.05)  
plot(x)
```

pph

*Calculates the posterior probability of hypotheses for one study***Description**

The function takes a single effect size and its standard error and calculates the posterior probability of each hypothesis (H<: the effect size is less than 0, H0: the effect size is zero, or H>: the effect size is greater than zero).

Usage

```
pph(beta, se.beta, beta0 = 0, se0 = NULL, ci = 99, se.mult = 1,
     H0 = c(0, 0), H.priors = rep(1/3, 3), scale = FALSE, adjust = FALSE,
     epsilon = 1e-06, adj.factor = 1e-04)
```

Arguments

beta	Effect size.
se.beta	Standard error for the effect.
beta0	A prior value for the effect size. Default is zero.
se0	A prior standard error for the effect size. Default is NULL and is calculated automatically.
ci	Is used to calculate the prior standard error if se0 = NULL. The default value of 99 calculates the prior standard error so that the 99 largest (furthest from zero) confidence interval of the data distribution.
se.mult	Standard error multiplier used to increase or decrease the prior SE and used in conjunction with ci when se0 = NULL.
H0	A vector of length two that defines the null hypothesis. If the values are identical (e.g. H0 = c(0, 0)) a point null is used, otherwise the null is defined as the range between the lower and upper value.
H.priors	Prior hypothesis probabilities; default is an equal probability of 1/3, and they are specified in the following order: H<0, H=0, H>0.
scale	Logical. Whether to scale the effect size by its standard error. Standardising has no effect on the calculations but standardised effect sizes may be easier to compare in a forest plot.
adjust	Logical. Whether to adjust very small posterior hypothesis probabilities. Adjusting prevents a single study from having too much influence on the results when combining multiple studies. For example, if the probability for a hypothesis from one study is zero, then additional studies cannot alter this probability (multiplying anything by zero is still zero).
epsilon	A small value that a posterior hypothesis probability must fall below before an adjustment is made. Ignored if adjust = FALSE.
adj.factor	A small number added to each posterior hypothesis probability if adjust = TRUE and one of the posterior hypothesis probabilities is less than epsilon. The PPHs are then re-scaled to sum to one.

Details

Effect sizes could be a mean difference between groups, regression slope, odds ratio, or other values provided by statistical models. The standard errors represent the uncertainty in the effect size and are provided by most statistical software.

Two types of priors need to be specified. The first is the prior for the effect size, which is given by a mean (usually zero) and standard error, which can be specified by the user or calculated automatically. Second, prior probabilities for each hypothesis need to be provided, and an equal probability of 1/3 is used by default.

The null hypothesis (usually effect size = 0) can be specified as either a point null or as a range. A point null tests if the effect size is *exactly* zero while the range null tests if the effect size close enough to zero for practical purposes. 'Close enough' is defined by the user as a range on either size of zero.

To combine multiple effect sizes use the `ev.combo` function.

Value

Object of class `pph` which contains the posterior hypothesis probabilities and other calculated values.

See Also

[plot.PPH](#), [ev.combo](#)

Examples

```
# library(labstats) # need to install from CRAN
# plot(time.immob ~ dose, data=fluoxetine)
# summary(lm(time.immob ~ dose, data=fluoxetine))
x <- pph(beta=-0.25200, se.beta=0.09913) # dose effect from above output
x
```

summary.PPH

Print and summary methods for PPH and EV objects

Description

Print and summary methods for PPH and EV objects.

Usage

```
## S3 method for class 'PPH'
summary(x, digits = 4)

## S3 method for class 'PPH'
print(x, digits = 4)
```



```
## S3 method for class 'EV'  
summary(x, digits = 4)
```

```
## S3 method for class 'EV'  
print(x, digits = 4)
```

Arguments

x	A PPH or EV object.
digits	Number of digits to show.

Value

A list of posterior hypothesis probabilities and other calculated values.

Index

`calc.post.beta`, 2
`calc.post.se (calc.post.beta)`, 2
`ev.combo`, 2, 4, 5, 8
`forestplot`, 4
`plot.EV`, 5
`plot.PPH`, 6, 8
`pph`, 6, 7
`print.EV (summary.PPH)`, 8
`print.PPH (summary.PPH)`, 8
`prior.se (calc.post.beta)`, 2
`summary.EV (summary.PPH)`, 8
`summary.PPH`, 8