Package: rmBayes (via r-universe)

October 29, 2026

Type Package

Title Performing Bayesian Inference for Repeated-Measures Designs

Version 0.1.16

Date 2024-02-19

Description A Bayesian credible interval is interpreted with respect to posterior probability, and this interpretation is far more intuitive than that of a frequentist confidence interval. However, standard highest-density intervals can be wide due to between-subjects variability and tends to hide within-subject effects, rendering its relationship with the Bayes factor less clear in within-subject (repeated-measures) designs. This urgent issue can be addressed by using within-subject intervals in within-subject designs, which integrate four methods including the Wei-Nathoo-Masson (2023) <doi:10.3758/s13423-023-02295-1>, the Loftus-Masson (1994) <doi:10.3758/BF03210951>, the Nathoo-Kilshaw-Masson (2018) <doi:10.1016/j.jmp.2018.07.005>, and the Heck (2019) <doi:10.31234/osf.io/whp8t> interval estimates.

License GPL (>= 3)

Encoding UTF-8

LazyData true

RoxygenNote 7.3.1

Biarch true

Depends R (>= 3.5.0)

Imports methods, Rcpp (>= 0.12.0), RcppParallel, rstan (>= 2.26.0), rstantools (>= 2.1.1), stats

LinkingTo BH (>= 1.66.0), Rcpp (>= 0.12.0), RcppEigen (>= 0.3.3.3.0), RcppParallel, rstan (>= 2.26.0), StanHeaders (>= 2.26.0)

SystemRequirements GNU make

URL https://github.com/zhengxiaoUVic/rmBayes

BugReports https://github.com/zhengxiaoUVic/rmBayes/issues

Suggests knitr, testthat, rmarkdown, covr

Config/testthat/edition 3

NeedsCompilation yes

Author Zhengxiao Wei [aut, cre] (<https://orcid.org/0000-0003-1866-2320>), Farouk S. Nathoo [aut] (<https://orcid.org/0000-0002-2569-3507>), Michael E. J. Masson [aut] (<https://orcid.org/0000-0002-5430-6078>)

Maintainer Zhengxiao Wei <zhengxiao@uvic.ca>

Config/pak/sysreqs make

Repository https://cran.r-universe.dev

Date/Publication 2024-02-20 02:30:20 UTC

RemoteUrl https://github.com/cran/rmBayes

RemoteRef HEAD

RemoteSha 2dd45576b130e3b1a6567d04047db8d8198a0d61

Contents

recall.long 2
recall.wide 3
rmHDI 4

Index **9**

`recall.long` *The Repeated-Measures Data: Long Data Format*

Description

A data frame that is in the long format. A hypothetical experiment is designed to measure effects of study time in a free-recall paradigm. In this hypothetical experiment, to-be-recalled 20-word lists are presented at a rate of 1, 2, or 5 sec per word. Of interest is the relation between study time and number of recalled list words. Suppose that the experiment is run as a within-subject (repeated-measures) design, including a total of 10 subjects, each of whom participated in all three study-time conditions.

Usage

```
recall.long
```

Format

A data frame with 30 rows and 3 variables:

Subject The factor containing 10 levels of subjects, `"s1"` to `"s10"`.

Level The factor containing three levels of the experimental manipulation, `"Level1"`, `"Level2"`, and `"Level3"`.

Response The number of words recalled out of 20 for each of 10 subjects in each of three conditions.

References

Loftus, G. R., & Masson, M. E. J. (1994). Using confidence intervals in within-subject designs. Psychonomic Bulletin & Review, 1, 476–490.

---

recall.wide *The Repeated-Measures Data: Wide Data Format*

Description

A data frame that is in the wide format. A hypothetical experiment is designed to measure effects of study time in a free-recall paradigm. In this hypothetical experiment, to-be-recalled 20-word lists are presented at a rate of 1, 2, or 5 sec per word. Of interest is the relation between study time and number of recalled list words. Suppose that the experiment is run as a within-subject (repeated-measures) design, including a total of 10 subjects, each of whom participated in all three study-time conditions.

Usage

```
recall.wide
```

Format

A data frame with 10 rows (each row corresponds to a subject) and 3 variables:

Level1 The number of words recalled under the first level of the experimental manipulation - 1 second exposure duration per word.

Level2 The number of words recalled under the second level of the experimental manipulation - 2 seconds exposure duration per word.

Level3 The number of words recalled under the third level of the experimental manipulation - 5 seconds exposure duration per word.

References

Loftus, G. R., & Masson, M. E. J. (1994). Using confidence intervals in within-subject designs. Psychonomic Bulletin & Review, 1, 476–490.

rmHDI *Bayesian Interval Estimation for Repeated-Measures Designs*

Description

For both the homoscedastic and heteroscedastic cases in one-way within-subject (repeated-measures) designs, this function provides multiple methods to construct the credible intervals for condition means, with each method based on different sets of priors. The emphasis is on the calculation of intervals that remove the between-subjects variability that is a nuisance in within-subject designs, as proposed in Loftus and Masson (1994), the Bayesian analog proposed in Nathoo, Kilshaw, and Masson (2018), and the adaptation presented in Heck (2019).

Usage

```
rmHDI(
  data = NULL,
  whichSubject = "Subject",
  whichLevel = "Level",
  whichResponse = "Response",
  data.wide = NULL,
  cred = 0.95,
  warmup = 200,
  iter = 2000,
  chains = 4,
  method = 1,
  var.equal = TRUE,
  design = c("within", "between"),
  treat = c("random", "fixed"),
  ht = ifelse(match.arg(treat) == "fixed", 0.5, 1),
  hb = 1,
  seed = sample.int(.Machine$integer.max, 1),
  diagnostics = FALSE,
  permuted = TRUE,
  ...
)
```

Arguments

data	A long format matrix or data frame of the within-subject data whose three columns are labeled in whichSubject, whichLevel, and whichResponse (see Examples).
whichSubject	A character string specifying the column name of subject variable in the long format data.
whichLevel	A character string specifying the column name of level variable in the long format data.

`whichResponse` A character string specifying the column name of response variable in the long format data.

`data.wide` Alternatively, a wide format matrix or data frame of the within-subject data whose column indices are condition levels (see Examples). If both `data` and `data.wide` are specified, the credible intervals are only computed for `data`.

`cred` A scalar `[0,1]` specifying the credibility level of the credible interval. The default is .95.

`warmup` A positive integer specifying the number of warmup (burnin) iterations per chain. The default is 200.

`iter` A positive integer specifying the number of iterations for each chain (excluding warmup). The default is 2000.

`chains` A positive integer specifying the number of Markov chains. The default is 4.

`method` A positive integer in `0:6` specifying which method is used to construct within-subject HDIs (see Details). `method=0` implements the approach developed by Nathoo et al. (2018). `method=4` implements the approach by Heck (2019). `method=5` implements the approach by Heck (2019), but using the standard uniform prior distribution for the standard deviation of subject-specific random effects. `method=6` implements the approach by Heck (2019), but using the standard half-Cauchy prior distribution for the standard deviation of subject-specific random effects. `method=1` (default) implements the approach by Heck (2019), but using the Jeffreys prior for the overall mean rather than the condition means; the hierarchical specification regarding the g-prior for the standardized subject-specific random effects is discussed in Rouder et al. (2012, p. 361-362). With the Jeffreys prior for the overall mean and residual variance, `method=2` uses the standard uniform prior, and `method=3` uses the standard half-Cauchy prior for the standard deviation of random effects. For the computation of the standard HDI, see the `design` argument below.

`var.equal` A logical variable indicating whether to treat the variance of the response within each condition (level of the experimental manipulation) as being equal. If `TRUE` (default), the homogeneity of variance holds, and a common variance across conditions is assumed. Otherwise, `FALSE` will generate unequal interval widths for conditions. Two approaches are currently provided for the heteroscedastic within-subject data: `method=0` implements the approach developed by Nathoo et al. (2018, p. 5); `method=1` (default method if `var.equal=FALSE`) implements the heteroscedastic standard HDI method on the subject-centering transformed data. If a `method` option other than `0` or `1` is used with `var.equal=FALSE`, a pooled estimate of variability will be used just as in the homoscedastic case, and a warning message will be returned.

`design` A character string specifying the experimental design. If `"within"` (default), construct the within-subject HDIs based on the given `method` in `0:6`. If `"between"`, construct the standard HDIs using the priors in `method=1` (but not removing the between-subjects variability).

`treat` A character string specifying the type of condition effects when `method` in `1:3`. If `"fixed"`, treat the condition effects as fixed effects through the reduced parametrization proposed by Rouder et al. (2012, p. 363). If `"random"` (default), treat the condition effects as random effects.

ht	A positive real number specifying the prior scale on the variability of standardized condition effects when method=1 (see Details). The default is 0.5 when treat="fixed" and 1 when treat="random".
hb	A positive real number specifying the prior scale on the variability of standardized subject-specific random effects when method=1 or method=4 (see Details). The default is 1.
seed	The seed for random number generation.
diagnostics	A logical variable indicating whether to return the MCMC summary statistics when method in 1:6. If FALSE (default), the function returns the HDI result (see Value). If TRUE, the function returns an additional object of S4 class representing the fitted results for assessing the MCMC convergence.
permuted	A logical variable indicating whether to pre-process the input data. No matter whether the input is long data or wide data, the function eventually converts it to wide data for calculation. If TRUE (default), the converted wide-format data are first ordered by their column names in alphabetic order. Then, the data are placed in ascending order by the first and second columns. If FALSE, the data are not ordered, and the returned HDI results are sensitive to data permutation. In other words, the row permutation (e.g., switching the first and second rows) and the column permutation (e.g., switching the first and second columns) will result in slightly different HDI estimates even if seed is set to be the same.
...	Additional arguments that pass to sampling, such as thin, algorithm, cores, etc.

Details

Wei, Nathoo, and Masson (2023) consider three credible intervals: (1) the within-subject Bayesian interval developed by Nathoo et al. (2018), whose derivation conditions on estimated random effects, (2) a modification of (1) based on a proposal by Heck (2019) to allow for shrinkage and account for estimation uncertainty, and (3) an alternative to option (2) based on the default priors used in Rouder, Morey, Speckman, and Province (2012). Markov chain Monte Carlo sampling is also used to obtain the standard highest-density interval (HDI) for each condition mean in a one-way between-subjects design.

When the homogeneity of variance holds, a linear mixed-effects model M_1 for the mean response in a one-way within-subject design is

$$M_1 : Y_{ij} = \mu + \sigma_\epsilon(t_i + b_j) + \epsilon_{ij} versus M_0 : Y_{ij} = \mu + \sigma_\epsilon b_j + \epsilon_{ij}, \epsilon_{ij}\ N(0, \sigma_\epsilon^2), i = 1, \ldots, a; j = 1, \ldots, n,$$

where $Y_i j$ represents the mean response for the j-th subject under the i-th level of the experimental manipulation; μ is the overall mean, $\tau_i = \sigma_\epsilon t_i$ is the i-th level of the experimental manipulation; $\mu_i = \mu + \tau_i$, for the means model, is the i-th condition mean; b_j is the standardized subject-specific random effects; a is the number of levels; n is the number of subjects; ϵ_{ij} are independent and identically distributed. The effects t_i and b_j are both standardized relative to the standard deviation of the error σ_ϵ and become dimensionless (Rouder et al., 2012).

An assumption articulated in method=0 is the Jeffreys prior for the condition means μ_i and residual variance σ_ϵ^2 (Nathoo et al., 2018).

Priors used in `method=1` are the Jeffreys prior for the overall mean μ and residual variance, a g-prior structure for standardized effects ($t_i\ N(0, g_t)$, $b_j\ N(0, g_b)$), and independent scaled inverse-chi-square priors with one degree of freedom for the scale hyperparameters of the g-priors ($g_t\ Scale - inv - \chi^2(1, h_t^2)$, $g_b\ Scale - inv - \chi^2(1, h_b^2)$).

Priors used in `method=2` are the Jeffreys prior for the overall mean and residual variance, a normal distribution for (not standardized) effects ($\sigma_\epsilon t_i\ N(0, g_t)$, $\sigma_\epsilon b_j\ N(0, g_b)$), and the standard uniform distribution for the square root of g parameter ($sqrt(g_t)\ Unif(0, 1)$, $sqrt(g_b)\ Unif(0, 1)$).

Priors used in `method=3` are the Jeffreys prior for the overall mean and residual variance, a normal distribution for (not standardized) effects, and the standard half-Cauchy distribution for the square root of g parameter ($sqrt(g_t)\ Half - Cauchy(0, 1)$, $sqrt(g_b)\ Half - Cauchy(0, 1)$).

Priors used in `method=4` are the Jeffreys prior for the condition means and residual variance, a g-prior structure for standardized subject-specific random effects, and independent scaled inverse-chi-square priors with one degree of freedom for the scale hyperparameters of the g-priors (Heck, 2019).

Priors used in `method=5` are the Jeffreys prior for the condition means and residual variance, a normal distribution for (not standardized) subject-specific random effects, and the standard uniform distribution for the square root of g parameter.

Priors used in `method=6` are the Jeffreys prior for the condition means and residual variance, a normal distribution for (not standardized) subject-specific random effects, and the standard half-Cauchy distribution for the square root of g parameter.

Value

A `list` with three components, if `diagnostics=FALSE`:

`HDI` A matrix of HDI lower and upper bounds, whose row names are the condition levels.

`posterior means`
The posterior condition means when using `method` in `1:6`.

`arithmetic means`
Or, the arithmetic condition means when using `method=0`.

`width` The HDI width, which is the half-length from the lower bound to the upper bound.

A `list` with four components including an additional object of S4 class representing the fitted results, if `diagnostics=TRUE`.

Author(s)

Zhengxiao Wei (<zhengxiao@uvic.ca>), Farouk S. Nathoo (<nathoo@uvic.ca>), Michael E. J. Masson (<mmasson@uvic.ca>).

References

Heck, D. W. (2019). Accounting for estimation uncertainty and shrinkage in Bayesian within-subject intervals: A comment on Nathoo, Kilshaw, and Masson (2018). Journal of Mathematical Psychology, 88, 27–31.

Loftus, G. R., & Masson, M. E. J. (1994). Using confidence intervals in within-subject designs. Psychonomic Bulletin & Review, 1, 476–490.

Nathoo, F. S., Kilshaw, R. E., & Masson, M. E. J. (2018). A better (Bayesian) interval estimate for within-subject designs. Journal of Mathematical Psychology, 86, 1–9.

Rouder, J. N., Morey, R. D., Speckman, P. L., & Province, J. M. (2012). Default Bayes factors for ANOVA designs. Journal of Mathematical Psychology, 56, 356–374.

Stan Development Team (2024). RStan: the R interface to Stan. R package version 2.32.5 https://mc-stan.org

Wei, Z., Nathoo, F. S., & Masson, M. E. J. (2023). Investigating the relationship between the Bayes factor and the separation of credible intervals. Psychonomic Bulletin & Review, 30, 1759–1781.

Examples

```
## Not run:
data(recall.wide) # Example data, wide format
rmHDI(data.wide = recall.wide, seed = 277)

data(recall.long) # Example data, long format
rmHDI(recall.long, seed = 277)

colnames(recall.long) <- c("Participant", "Condition", "DV")
rmHDI(recall.long, whichSubject = "Participant",
whichLevel = "Condition", whichResponse = "DV", seed = 277)


## Nathoo et al. (2018) approach
data(recall.long)
rmHDI(recall.long, method = 0)
rmHDI(recall.long, method = 0, var.equal = FALSE)


## Standard HDI
rmHDI(recall.long, design = "between", seed = 277)

## MCMC diagnostics
rmHDI(recall.long, seed = 277, diagnostics = TRUE)$diagnostics

## End(Not run)
```

Index

∗ **datasets**
recall.long, 2
recall.wide, 3

recall.long, 2
recall.wide, 3
rmHDI, 4

sampling, *6*