

Package ‘pecanr’

June 23, 2026

Title Partial Eta-Squared for Crossed, Nested, and Mixed Linear Mixed Models

Version 0.3.0

Description Computes partial eta-squared effect sizes for fixed effects in linear mixed models fitted with the 'lme4' package. Supports crossed, nested, and mixed (crossed-and-nested) random effects structures with any number of grouping factors. Mixed designs handle cases where grouping factors are simultaneously crossed with some variables and nested within others (e.g., photos nested within models, but both crossed with participants). Factor predictors are supported directly, and a single factor-level (omnibus) effect size can be obtained for a multi-level factor or multi-df interaction. Random slope variances are translated to the outcome scale using a variance decomposition approach, correctly accounting for predictor scaling and interaction terms. Both general and operative effect sizes are provided, with optional parametric bootstrap confidence intervals. For correlated predictors, per-predictor effect sizes use unique (semipartial) variance by default. Methods are based on Correll, Mellinger, McClelland, and Judd (2020) <[doi:10.1016/j.tics.2019.12.009](https://doi.org/10.1016/j.tics.2019.12.009)>, Correll, Mellinger, and Pedersen (2022) <[doi:10.3758/s13428-021-01687-2](https://doi.org/10.3758/s13428-021-01687-2)>, and Rights and Sterba (2019) <[doi:10.1037/met0000184](https://doi.org/10.1037/met0000184)>.

License MIT + file LICENSE

URL <https://github.com/bcohen0901/pecanr>

BugReports <https://github.com/bcohen0901/pecanr/issues>

Encoding UTF-8

Imports lme4, stats

Suggests spelling, testthat (>= 3.0.0)

Config/testthat/edition 3

Language en-US

Config/roxygen2/version 8.0.0

NeedsCompilation no

Author Brandon Cohen [aut, cre] (ORCID:
 <<https://orcid.org/0009-0004-1802-3435>>),
Joshua Correll [aut, ths] (ORCID:
 <<https://orcid.org/0000-0003-1142-459X>>)
Maintainer Brandon Cohen <brandon.cohen-1@colorado.edu>
Repository CRAN
Date/Publication 2026-06-23 08:50:02 UTC

Contents

batch_eta2p	2
eta2p	4
eta2p_omnibus	10

Index	13
--------------	-----------

batch_eta2p	<i>Batch Calculate Partial Eta-Squared for Multiple Effects</i>
-------------	---

Description

Calculates partial eta-squared for all fixed effects (excluding the intercept) in a model, returning one row per model coefficient. For a multi-level factor or multi-df interaction this yields one value *per contrast* (per design-matrix column), not a single factor-level value; for the omnibus effect size of such a term, use [eta2p_omnibus](#).

Usage

```
batch_eta2p(  
  model,  
  data,  
  design = c("crossed", "nested", "mixed"),  
  subj_var = NULL,  
  item_var = NULL,  
  cross_vars = NULL,  
  nest_vars = NULL,  
  operative = FALSE,  
  partial_predictors = TRUE,  
  verbose = FALSE  
)
```

Arguments

model	A fitted model object from <code>lme4::lmer()</code> .
data	Data frame used to fit the model.

design	Character string: "crossed", "nested", or "mixed". Use "mixed" when the model contains both crossed and nested random effects (e.g., photos nested within models, but both crossed with participants).
subj_var	Character string specifying the subject/participant grouping variable. Retained for backward compatibility; prefer cross_vars.
item_var	Character string specifying the item/stimulus grouping variable. Retained for backward compatibility; prefer cross_vars.
cross_vars	Character vector of grouping variable names that are <i>crossed</i> with each other and with the nested hierarchy (e.g., c("participant") or c("participant", "rater")). Supersedes subj_var and item_var when provided.
nest_vars	Character vector of nesting variables ordered from lowest to highest level (e.g., c("photo", "model") when photos are nested within models). Required for design = "nested" or design = "mixed".
operative	Logical. If TRUE, calculates operative effect size (excludes variance components that do not contribute to the SE of the effect). Default is FALSE.
partial_predictors	Logical. If TRUE (the default), the variance attributed to a single (non-interaction) predictor is its <i>unique</i> (semipartial) variance – the variance remaining after removing the part linearly predictable from the other fixed-effect predictors, equivalently $\text{Var}(X) \times \text{tol}(X)$ where <i>tol</i> is the predictor's tolerance. This yields a measure that reflects only the unique contribution of the predictor and declines as its redundancy with the others increases. If FALSE, the predictor's <i>total</i> variance is used instead (the raw decomposition). For centered, orthogonal designs the two are identical; they differ only when predictors are correlated. See Details. Has no effect on interaction terms (whose components are always centered) or when var_x is supplied.
verbose	Logical. If TRUE, prints detailed results. Default is FALSE.

Value

A data frame with one row per effect containing:

effect	Name of the fixed effect.
eta2p	Partial eta-squared value.
effect_level	Effect level (nested/mixed designs only).
variance_effect	Variance attributed to the effect.
variance_error	Error variance used in the denominator.
type	"general" or "operative".
within_subj, within_item	(Crossed/mixed, operative only) Whether the effect is within subjects / within items.

Rows are sorted by eta2p in decreasing order.

eta2p

*Partial Eta-Squared for Linear Mixed Models***Description**

Calculates partial eta-squared effect sizes for fixed effects in linear mixed models with crossed, nested, or mixed (crossed-and-nested) random effects.

Usage

```
eta2p(
  model,
  effect,
  data,
  design = c("crossed", "nested", "mixed"),
  subj_var = NULL,
  item_var = NULL,
  cross_vars = NULL,
  nest_vars = NULL,
  effect_level = NULL,
  var_x = NULL,
  operative = FALSE,
  partial_predictors = TRUE,
  ci = FALSE,
  ci_level = 0.95,
  n_boot = 1000,
  seed = NULL,
  verbose = FALSE
)
```

Arguments

model	A fitted model object from <code>lme4::lmer()</code> .
effect	Character string specifying the fixed effect to analyze. May be given either as the model coefficient name (e.g. "Group75yr") or as the underlying variable name (e.g. "Group"); a variable name that maps to a single coefficient is resolved automatically (with a message). A variable that maps to several coefficients (a multi-level factor) raises an informative error pointing to <code>eta2p_omnibus()</code> (for one factor-level value) or <code>batch_eta2p()</code> (for per-coefficient values). Factor predictors and factor random slopes are read from the model design matrix / model frame, so no manual recoding to numeric is required.
data	Data frame used to fit the model.
design	Character string: "crossed", "nested", or "mixed". Use "mixed" when the model contains both crossed and nested random effects (e.g., photos nested within models, but both crossed with participants).

subj_var	Character string specifying the subject/participant grouping variable. Retained for backward compatibility; prefer cross_vars.
item_var	Character string specifying the item/stimulus grouping variable. Retained for backward compatibility; prefer cross_vars.
cross_vars	Character vector of grouping variable names that are <i>crossed</i> with each other and with the nested hierarchy (e.g., c("participant") or c("participant", "rater")). Supersedes subj_var and item_var when provided.
nest_vars	Character vector of nesting variables ordered from lowest to highest level (e.g., c("photo", "model") when photos are nested within models). Required for design = "nested" or design = "mixed".
effect_level	Character string specifying the level at which the effect varies (e.g., "L1", "L2"). If NULL, detected automatically from nest_vars.
var_x	Optional numeric. Pre-computed variance of the predictor (or interaction product). If supplied, overrides the internal var() calculation from data. Useful when raw data are unavailable but the predictor variance is known (e.g., from design: a +/-1 binary predictor has var_x = 1).
operative	Logical. If TRUE, calculates operative effect size (excludes variance components that do not contribute to the SE of the effect). Default is FALSE.
partial_predictors	Logical. If TRUE (the default), the variance attributed to a single (non-interaction) predictor is its <i>unique</i> (semipartial) variance – the variance remaining after removing the part linearly predictable from the other fixed-effect predictors, equivalently $\text{Var}(X) \times \text{tol}(X)$ where tol is the predictor's tolerance. This yields a measure that reflects only the unique contribution of the predictor and declines as its redundancy with the others increases. If FALSE, the predictor's <i>total</i> variance is used instead (the raw decomposition). For centered, orthogonal designs the two are identical; they differ only when predictors are correlated. See Details. Has no effect on interaction terms (whose components are always centered) or when var_x is supplied.
ci	Logical. If TRUE, computes a confidence interval for eta2p by parametric bootstrap (lme4::bootMer): new responses are simulated from the fitted model, the model is refit, and eta2p is recomputed on each replicate, with percentile limits taken from the resulting distribution. This is the appropriate interval for a ratio of REML variance components, which has no closed-form sampling distribution. Each replicate refits the model, so this is opt-in and can be slow on large models. Default is FALSE.
ci_level	Numeric confidence level for the bootstrap interval (default 0.95).
n_boot	Integer number of bootstrap replicates when ci = TRUE (default 1000). Use a smaller value (e.g. 100-200) to gauge runtime on large models before committing to the full count.
seed	Optional integer seed for reproducible bootstrap intervals.
verbose	Logical. If TRUE, prints detailed results. Default is FALSE.

Details

Variance decomposition: The function implements a variance decomposition approach for computing partial eta-squared in mixed models. The variance a fixed effect explains is

$$\text{Var}_{\text{explained}} = b^2 \times \sigma_X^2,$$

where b is the (partial) fixed-effect coefficient and σ_X^2 is the variance of the predictor column. Random slope variances are translated to the outcome scale by the same principle,

$$\sigma_{\text{slope}}^2(Y) = \sigma_b^2 \times \sigma_X^2.$$

For an interaction, σ_X^2 is the variance of the product of the *mean-centered* constituent predictors (see "Centering" below). The `var_x` argument bypasses this computation when the variance is known a priori.

Centering and intended scope: These quantities are exact when predictors are **mean-centered**, which is standard and recommended practice for the experimental and contrast-coded designs the method primarily targets (sum/contrast coding centers factors automatically). Two points:

- **Interactions: components are centered automatically.** For an interaction the function mean-centers each *constituent* predictor and then takes the variance of their product (it does *not* center the product itself). This makes the interaction effect size invariant to the location of the constituents: re-centering a predictor by a constant leaves the model fit and the interaction coefficient unchanged, and now also leaves the effect size unchanged. (The variance of the *raw* product depends on the constituents' means and so is not well defined as an effect size; this is why the centering is applied. The constituent predictors should still be centered before fitting for the coefficients themselves to be interpretable, but the effect-size value no longer depends on it.)
- **Main effects / random slopes use σ_X^2 , which equals $E[X^2]$ only when X is centered.** The function does not center main-effect predictors (doing so would not change $\sigma_X^2 = \text{Var}(X)$, which is already location-invariant), but the random-slope translation $\sigma_b^2 \sigma_X^2$ is exact only for centered predictors; center continuous predictors before fitting.

Under correlated predictors, each b is a partial coefficient (it controls for the other predictors). By default (`partial_predictors = TRUE`) it is multiplied by the predictor's *unique* variance – the variance remaining after the part linearly predictable from the other predictors is removed – so the variance attributed to the predictor is

$$b^2 \text{Var}(X \mid X_{\text{others}}) = b^2 \text{Var}(X) \text{tol}(X),$$

where $\text{tol}(X) = 1 - R_{X \sim X_{\text{others}}}^2$ is the predictor's tolerance. (The two forms are algebraically identical.) This yields a semipartial variance-explained measure: it reflects only the unique contribution of the predictor and declines as its redundancy with the others increases, matching the unique-variance target estimated by an explained- error (model-comparison) approach.

Setting `partial_predictors = FALSE` instead uses the predictor's *total* variance σ_X^2 . The resulting per-predictor values are partial (conditional) effect sizes that do not isolate unique variance when predictors are correlated. For centered, orthogonal designs the two options coincide exactly, because the covariance among predictors is zero, so the choice matters only under collinearity.

The tolerance correction is also what makes the interaction handling location-invariant: centering an interaction's components is equivalent to the $\text{Var} \times \text{tol}$ of the product, which is constant

across re-centering. The choice is independent of the operative option, which concerns the error denominator rather than the numerator. See Rights & Sterba (2019) on the partial-versus-total distinction.

General vs. operative effect sizes: For **general effect sizes** (default), all variance components appear in the denominator. For **operative effect sizes** (`operative = TRUE`), only components that contribute to the standard error of the effect are included. In mixed designs, each grouping factor (both crossed and nested) is independently classified as "within" or "between" the focal effect; between-subjects intercept variances are excluded from the operative denominator.

The operative option concerns only the denominator; it is independent of the `partial_predictors` setting, which controls the numerator. The two compose: an operative effect size uses the same numerator a general one would (unique variance by default), paired with the reduced operative denominator.

Mixed designs: Use `design = "mixed"` when the model contains both crossed and nested random effects simultaneously. A canonical example is participants viewing multiple photos of each model (stimulus): photos are nested within models, but both levels are crossed with participants. The corresponding **lme4** model and `eta2p` call would be:

```
fit <- lmer(y ~ x + (1 | participant) + (1 | model) + (1 | photo:model),
           data = dat)
eta2p(fit, "x", dat,
      design      = "mixed",
      cross_vars  = "participant",
      nest_vars   = c("photo", "model"))
```

Residual variance is attributed to the crossed side and counted exactly once; the nested calculator's residual is subtracted to prevent double-counting.

Crossed designs: Supports any number of grouping factors via `cross_vars`. The two-argument form (`subj_var + item_var`) is retained for backward compatibility and is equivalent to `cross_vars = c(subj_var, item_var)`.

Confidence intervals: With `ci = TRUE`, a confidence interval is obtained by parametric bootstrap via `lme4::bootMer`: responses are simulated from the fitted model, the model is refit, and `eta2p` is recomputed on each replicate; `ci_lower` and `ci_upper` are the percentile limits. Because `eta2p` is a ratio of REML variance components with no closed-form sampling distribution, this bootstrap is the appropriate interval rather than an analytic (e.g. noncentral) formula. Set seed for reproducibility, and start with a small `n_boot` to gauge runtime on large models.

Value

An object of class `"eta2p_lmm"` containing:

<code>eta2p</code>	Partial eta-squared value.
<code>variance_effect</code>	Variance explained by the effect.
<code>variance_error</code>	Error variance (denominator).
<code>effect</code>	Name of the effect.
<code>design</code>	Design type: <code>"crossed"</code> , <code>"nested"</code> , or <code>"mixed"</code> .

operative Whether operative effect size was calculated.

ci_level, ci_lower, ci_upper, boot_n, boot_method
(Only when ci = TRUE) The confidence level, the lower and upper percentile bootstrap limits, the number of usable bootstrap replicates, and the method label ("parametric (bootMer)").

variance_components
List of individual variance components. For "mixed" designs this is a list with \$crossed and \$nested sub-lists.

within_between (Mixed/crossed designs) Named list or vector classifying each grouping factor as "within" or "between" with respect to the focal effect. Only populated when operative = TRUE.

cross_vars (Crossed/mixed designs) Crossed grouping variable names.

nest_vars (Nested/mixed designs) Nested grouping variable names.

n_per_factor (Crossed designs) Number of units per crossed factor.

n_cross (Mixed designs) Number of units per crossed factor.

n_nested (Mixed designs) Number of units per nested factor.

n_levels (Nested designs) Number of units per nested level.

effect_level (Nested/mixed designs) Detected or supplied effect level.

References

- Correll, J., Mellinger, C., McClelland, G. H., & Judd, C. M. (2020). Avoid Cohen's 'Small', 'Medium', and 'Large' for Power Analysis. *Trends in Cognitive Sciences*, 24(3), 200–207. doi:10.1016/j.tics.2019.12.009
- Correll, J., Mellinger, C., & Pedersen, E. J. (2022). Flexible approaches for estimating partial eta squared in mixed-effects models with crossed random factors. *Behavior Research Methods*, 54, 1626–1642. doi:10.3758/s13428021016872
- Rights, J. D., & Sterba, S. K. (2019). Quantifying explained variance in multilevel models: An integrative framework for defining R-squared measures. *Psychological Methods*, 24(3), 309–338. doi:10.1037/met0000184

See Also

[eta2p_omnibus](#) for a single factor-level value for a multi-df factor or interaction (matching an omnibus F /chi-square test); [batch_eta2p](#) to compute per-coefficient values for all fixed effects at once.

Examples

```
library(lme4)

# Two crossed factors (backward-compatible call)
set.seed(42)
crossed_data <- data.frame(
  y = rnorm(120),
  condition = rep(c(-0.5, 0.5), 60),
```

```

    subject = factor(rep(1:20, each = 6)),
    item     = factor(rep(1:6, 20))
  )
model_c <- lmer(y ~ condition + (1 | subject) + (1 | item),
               data = crossed_data)
eta2p(model_c, "condition", crossed_data,
       design = "crossed",
       subj_var = "subject",
       item_var = "item")

# Three crossed factors using cross_vars
set.seed(42)
three_way_data <- data.frame(
  y          = rnorm(180),
  condition = rep(c(-0.5, 0.5), 90),
  subject    = factor(rep(1:20, each = 9)),
  item       = factor(rep(rep(1:6, each = 3), 10)),
  rater      = factor(rep(1:3, 60))
)
model_3 <- lmer(y ~ condition +
               (1 | subject) + (1 | item) + (1 | rater),
               data = three_way_data)
eta2p(model_3, "condition", three_way_data,
       design = "crossed",
       cross_vars = c("subject", "item", "rater"))

# Mixed design: photos nested within models, crossed with participants
# Participants each view multiple photos of multiple models.
# Photos are nested within models (each photo belongs to one model),
# but both levels are crossed with participants.
set.seed(42)
n_subj <- 30; n_model <- 10; n_photo <- 3
mixed_data <- expand.grid(
  participant = factor(seq_len(n_subj)),
  photo_id    = factor(seq_len(n_model * n_photo))
)
mixed_data$model_id <- factor(
  rep(seq_len(n_model), each = n_photo)[as.integer(mixed_data$photo_id)]
)
mixed_data$x <- rnorm(nrow(mixed_data))
mixed_data$y <- rnorm(nrow(mixed_data))

model_m <- lmer(y ~ x +
               (1 | participant) +
               (1 | model_id) +
               (1 | photo_id:model_id),
               data = mixed_data)
eta2p(model_m, "x", mixed_data,
       design = "mixed",
       cross_vars = "participant",
       nest_vars = c("photo_id", "model_id"))

# Supply predictor variance directly (var_x)

```

```

eta2p(model_c, "condition", crossed_data,
      design      = "crossed",
      cross_vars  = c("subject", "item"),
      var_x       = 1)  # +/-1 binary predictor: var = 1 by design

# Factor predictor (no manual recoding needed)
crossed_data$grp <- factor(rep(c("A", "B"), 60))
model_f <- lmer(y ~ grp + (1 | subject) + (1 | item), data = crossed_data)
# accepts the variable name "grp"; resolves to the coefficient automatically
eta2p(model_f, "grp", crossed_data,
      design = "crossed", cross_vars = c("subject", "item"))

# Operative effect size
eta2p(model_c, "condition", crossed_data,
      design = "crossed", cross_vars = c("subject", "item"),
      operative = TRUE)

# Total-variance (raw) numerator instead of the default unique variance
eta2p(model_c, "condition", crossed_data,
      design = "crossed", cross_vars = c("subject", "item"),
      partial_predictors = FALSE)

# Confidence interval by parametric bootstrap
eta2p(model_c, "condition", crossed_data,
      design = "crossed", cross_vars = c("subject", "item"),
      ci = TRUE, n_boot = 200, seed = 1)

# Omnibus effect size for a multi-level factor
crossed_data$emo <- factor(rep(c("a", "b", "c", "d"), 30))
model_e <- lmer(y ~ emo + (1 | subject) + (1 | item), data = crossed_data)
eta2p_omnibus(model_e, "emo", crossed_data,
      design = "crossed", cross_vars = c("subject", "item"))

# All fixed effects at once (per-coefficient)
batch_eta2p(model_c, crossed_data,
      design = "crossed", cross_vars = c("subject", "item"))

```

eta2p_omnibus

Omnibus (factor-level) Partial Eta-Squared for a Multi-df Effect

Description

For a multi-level factor or a multi-df interaction, computes a single factor-level partial eta-squared corresponding to the omnibus test of that effect (e.g. the chi-square or F test on all of its degrees of freedom), rather than one value per contrast. The variance attributed to the effect is the variance of the summed fitted contribution of all the design-matrix columns belonging to the effect (which correctly includes the covariances among those columns); the error denominator is the same one eta2p() uses.

Usage

```
eta2p_omnibus(model, effect, data, ...)
```

Arguments

<code>model</code>	A fitted lme4 model.
<code>effect</code>	The variable/effect name. For a factor main effect, the bare variable name (e.g. "rating_type"). For an interaction, the colon form of the term as it appears in the design matrix (e.g. "group:rating_type").
<code>data</code>	The data frame used to fit the model.
<code>...</code>	Passed to <code>eta2p()</code> to obtain the error denominator (design, cross_vars, subj_var, item_var, nest_vars, operative, etc.).

Details

`batch_eta2p()` and a per-coefficient call to `eta2p()` return one value per contrast (per design-matrix column). For a multi-level factor or a multi-df interaction, none of those single values corresponds to the omnibus test of the whole effect (the F or chi-square test on all of its degrees of freedom). This function fills that gap.

The omnibus variance explained is computed as the variance of the *summed* fitted contribution of all the effect's columns, $\text{Var}(\sum_j b_j x_j)$, which correctly includes the covariances among the dummy/contrast columns. Simply adding the per-coefficient numerators would drop those covariances and is incorrect. The error denominator is the same one `eta2p()` uses, so the result remains a partial eta-squared (error-variance denominator), *not* a total-variance R^2 ; expect it to differ from `r2glmm::r2beta()`, which targets the latter.

Value

An object of class "eta2p_omnibus": a list containing

<code>eta2p</code>	The omnibus (factor-level) partial eta-squared.
<code>variance_effect</code>	Variance of the summed fitted contribution of all the effect's design-matrix columns (includes covariances among them).
<code>variance_error</code>	Error variance denominator (from <code>eta2p()</code>).
<code>effect</code>	The effect name requested.
<code>coefficients</code>	The design-matrix column names aggregated.
<code>n_df</code>	Number of columns aggregated (the effect's degrees of freedom; matches the df of the corresponding omnibus test).
<code>type</code>	"omnibus".

See Also

[eta2p](#) for single-coefficient values and [batch_eta2p](#) for all coefficients at once.

Examples

```
library(lme4)
set.seed(1)
d <- expand.grid(subject = factor(1:30), item = factor(1:12),
                 emo = factor(c("a", "b", "c", "d")))
d$y <- as.integer(d$emo) + rnorm(30)[d$subject] + rnorm(nrow(d))
m <- lmer(y ~ emo + (1 | subject) + (1 | item), data = d)
# one factor-level value to pair with the omnibus test of `emo`:
eta2p_omnibus(m, "emo", d, design = "crossed",
              cross_vars = c("subject", "item"))
```

Index

batch_eta2p, [2](#), [8](#), [11](#)

eta2p, [4](#), [11](#)

eta2p_omnibus, [2](#), [8](#), [10](#)