

Package ‘fastFMM’

May 18, 2026

Type Package

Title Fast Functional Mixed Models using Fast Univariate Inference

Version 1.0.1

Date 2026-05-17

Description Implementation of the fast univariate inference approach (Cui et al. (2022) <[doi:10.1080/10618600.2021.1950006](https://doi.org/10.1080/10618600.2021.1950006)>, Loewinger et al. (2024) <[doi:10.7554/eLife.95802.2](https://doi.org/10.7554/eLife.95802.2)>, Xing et al. (2024) <[doi:10.1093/bioinformatics/btad001](https://doi.org/10.1093/bioinformatics/btad001)>), estimating functional mixed models. User guides and Python package information can be found at <https://github.com/gloewing/photometry_FLMM>.

License GPL (>= 3)

Encoding UTF-8

LazyData true

Repository CRAN

Imports dplyr, lme4, parallel, cAIC4, magrittr, mgcv, MASS, lsei, refund, stringr, Matrix, mvtnorm, progress, ggplot2, gridExtra, Rfast, lmeresampler, stats, methods

URL <https://github.com/awqx/fastFMM>

BugReports <https://github.com/awqx/fastFMM/issues>

VignetteBuilder knitr

Suggests knitr, rmarkdown, spelling

Language en-US

Depends R (>= 3.5)

Config/roxygen2/version 8.0.0

NeedsCompilation no

Author Erjia Cui [aut],
Gabriel Loewinger [aut],
Al Xin [aut, cre]

Maintainer Al Xin <axin@andrew.cmu.edu>

Date/Publication 2026-05-18 08:20:07 UTC

Contents

bdiag_m	2
cov_organize_start	2
d2pvt	3
fui	4
lick	7
new_fastFMM	8
new_fastFMMconc	9
plot_fui	10
select_knots	12

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

bdiag_m	<i>Fast block diagonal generator, taken from Matrix package examples</i>
---------	--

Description

Copyright (C) 2016 Martin Maechler, ETH Zurich. Copied here for more convenient integration.

Usage

```
bdiag_m(lmat)
```

Arguments

lmat	Matrix
------	--------

Value

Block diagonal version of input

cov_organize_start	<i>Organize covariance matrices</i>
--------------------	-------------------------------------

Description

Read the data frame of variance-covariance interaction and organize the covariance matrices correctly. Finds indices to feed into 'cov_organize' function above

Usage

```
cov_organize_start(cov_vec)
```

Arguments

cov_vec	Character vector
---------	------------------

Value

List of relevant outputs

d2pvt

Machen et al. (2025) variable trial length data

Description

Data from an experiment where mice ran through a maze where they received either a strawberry milkshake (SMS) or water (H2O) reward. Mice entered the reward zone at variable times.

Usage

d2pvt

Format

A data frame with 1335 rows and 248 columns:

id Unique mouse ID

session Session

outcome The outcome of the trial, either SMS or H2O

SMS Redundant encoding for binary encoding of SMS reward

latency Time, in seconds, for mouse to reach the reward

trial Trial

photometry Photometry recordings at time point

rewarded Whether the mouse had been rewarded by time point

Source

Machen B, Miller SN, Xin A, Lampert C, Assaf L, Tucker J, Herrell S, Pereira F, Loewinger G, Beas S. The encoding of interoceptive-based predictions by the paraventricular nucleus of the thalamus D2+ neurons. *bioRxiv* (2025).

Description

Fit a function-on-scalar regression model for longitudinal functional outcomes and scalar predictors using the Fast Univariate Inference (FUI) approach (Cui et al. 2022).

Usage

```
fui(
  formula,
  data,
  family = "gaussian",
  var = TRUE,
  analytic = TRUE,
  parallel = FALSE,
  silent = FALSE,
  argvals = NULL,
  nknots_min = NULL,
  nknots_min_cov = 35,
  smooth_method = "GCV.Cp",
  splines = "tp",
  design_mat = FALSE,
  residuals = FALSE,
  n_boots = 500,
  boot_type = NULL,
  seed = 1,
  subj_id = NULL,
  n_cores = NULL,
  caic = FALSE,
  randeffs = FALSE,
  non_neg = 0,
  MoM = 1,
  concurrent = FALSE,
  impute_outcome = FALSE,
  override_zero_var = FALSE,
  unsmooth = FALSE
)
```

Arguments

formula	Two-sided formula object in lme4 formula syntax. The difference is that the response need to be specified as a matrix instead of a vector. Each column of the matrix represents one location of the longitudinal functional observations on the domain.
data	A data frame containing all variables in formula

family	GLM family of the response. Defaults to gaussian.
var	Logical, indicating whether to calculate and return variance of the coefficient estimates. Defaults to 'TRUE'.
analytic	Logical, indicating whether to use the analytic inference approach or bootstrap. Defaults to TRUE.
parallel	Logical, indicating whether to do parallel computing. Defaults to FALSE.
silent	Logical, indicating whether to show descriptions of each step. Defaults to FALSE.
argvals	A vector containing locations of observations on the functional domain. If not specified, a regular grid across the range of the domain is assumed. Currently only supported for bootstrap (analytic=FALSE).
nknots_min	Minimal number of knots in the penalized smoothing for the regression coefficients. Defaults to NULL, which then uses $L/2$ where L is the dimension of the functional domain.
nknots_min_cov	Minimal number of knots in the penalized smoothing for the covariance matrices. Defaults to 35.
smooth_method	How to select smoothing parameter in step 2. Defaults to "GCV.Cp"
splines	Spline type used for penalized splines smoothing. We use the same syntax as the mgcv package. Defaults to "tp".
design_mat	Logical, indicating whether to return the design matrix. Defaults to FALSE
residuals	Logical, indicating whether to save residuals from unsmoothed LME. Defaults to FALSE.
n_boots	Number of samples when using bootstrap inference. Defaults to 500.
boot_type	Bootstrap type (character): "cluster", "case", "wild", "reb", "residual", "parametric", "semiparametric". NULL defaults to "cluster" for non-gaussian responses and "wild" for gaussian responses. For small cluster ($n \leq 10$) gaussian responses, defaults to "reb".
seed	Numeric value used to make sure bootstrap replicate (draws) are correlated across functional domains for certain bootstrap approach
subj_id	Name of the variable that contains subject ID.
n_cores	Number of cores for parallelization. If greater than the value from 'parallel::detectCores()', defaults to one fewer than that value.
caic	Logical, indicating whether to calculate cAIC. Defaults to FALSE.
randeffs	Logical, indicating whether to return random effect estimates. Defaults to FALSE.
non_neg	0 - no non-negativity constrains, 1 - non-negativity constraints on every coefficient for variance, 2 - non-negativity on average of coefficients for 1 variance term. Defaults to 0.
MoM	Method of moments estimator. Defaults to 1.
concurrent	Logical, indicates whether to fit a concurrent model. Defaults to FALSE.
impute_outcome	Logical, indicates whether to impute missing outcome values with FPCA. Defaults to FALSE. Use with caution as the downstream effects are not tested.

override_zero_var	Logical, indicates whether to proceed with model fitting if columns have zero variance. Suggested for cases where individual columns have zero variance but interactions have non-zero variance. Defaults to 'FALSE'.
unsmooth	Logical, indicates whether to return the raw estimates of coefficients and variances without smoothing. Defaults to 'FALSE'.

Details

The FUI approach comprises of three steps:

1. Fit a univariate mixed model at each location of the functional domain, and obtain raw estimates from massive models;
2. Smooth the raw estimates along the functional domain;
3. Obtain the pointwise and joint confidence bands using an analytic approach for Gaussian data or Bootstrap for general distributions.

For more information on each step, please refer to the FUI paper by Cui et al. (2022).

Value

A list containing:

betaHat	Estimated functional fixed effects
argvals	Location of the observations
betaHat.var	Variance estimates of the functional fixed effects (if specified)
qn	critical values used to construct joint CI
...	...

Author(s)

Erjia Cui <ecui@umn.edu>, Gabriel Loewinger <gloewinger@gmail.com>, Al W Xin <axin@andrew.cmu.edu>

References

Cui, E., Leroux, A., Smirnova, E., Crainiceanu, C. (2022). Fast Univariate Inference for Longitudinal Functional Models. *Journal of Computational and Graphical Statistics*, 31(1), 219-230.

Examples

```
library(refund)

## random intercept only
set.seed(1)
DTI_use <- DTI[DTI$ID %in% sample(DTI$ID, 10),]
## clean conflicting predictor
DTI_use$visit.time <- NULL
fit_dti <- fui(
  cca ~ visit + sex + (1 | ID),
  data = DTI_use
)
```

lick

Jeong et al. (2022) licking behavior data

Description

Photometry recordings from head-fixed mice administered sucrose at random intervals. Observations are identified by subject ID, session, and trial number. Columns are downsampled for convenience.

Usage

lick

Format

A data frame with 2120 rows and 94 columns:

id Unique mouse ID

session Session

trial Trial

lick_rate_050, lick_rate_100, lick_rate_150, lick_rate_200 The average licking rate 0.5, 1.0, 1.5, and 2.0 seconds after sucrose administration

iri The inter-reward interval relative to the previous lick

photometry Photometry recordings over the trial

lick Whether a mouse was licking at a time point

Source

<https://dandiarchive.org/dandiset/000351/draft> Jeong H, Taylor A, Floeder JR, Lohmann M, Mihalas S, Wu B, Zhou M, Burke DA, Nambodiri VMK. Mesolimbic dopamine release conveys causal associations. *Science*. 2022 Dec 23;378(6626):eabq6740. doi: 10.1126/science.abq6740.

References

Jeong H, Taylor A, Floeder JR, Lohmann M, Mihalas S, Wu B, Zhou M, Burke DA, Nambodiri VMK. Mesolimbic dopamine release conveys causal associations. *Science*. 2022 Dec 23;378(6626):eabq6740. doi: 10.1126/science.abq6740. Epub 2022 Dec 23.

<code>new_fastFMM</code>	<i>Create a new "fastFMM" object</i>
--------------------------	--------------------------------------

Description

The class "fastFMM" (and its inheritors) contain parameters for fast univariate inference. The basic "fastFMM" object is equipped for non-concurrent model fitting.

Usage

```
new_fastFMM(
  formula,
  data,
  subj_id,
  argvals,
  family,
  residuals,
  caic,
  randeffs,
  var,
  analytic
)
```

Arguments

<code>formula</code>	Formula in 'lme4' formula syntax.
<code>data</code>	Data frame to fit.
<code>subj_id</code>	Character, name of variable containing IDs. Passed from 'fui'
<code>argvals</code>	List of points to fit on the functional domain. Only applies for the bootstrap case (i.e., 'analytic = FALSE').
<code>family</code>	Character, GLM family of the response. Passed from 'fui'.
<code>residuals</code>	Logical, indicates whether residuals are saved from unsmoothed LME. Passed from 'fui'.
<code>caic</code>	Logical, indicates cAIC calculation return. Defaults to 'FALSE'.
<code>randeffs</code>	Logical, indicates whether random effect estimates are returned. Passed from 'fui'.
<code>var</code>	Logical, indicates whether to calculate variance. Passed from 'fui'.
<code>analytic</code>	Logical, indicates whether variance will be calculated analytically. Passed from 'fui'.

Details

Object creation populates fields relevant to later steps, such as the location of the functional domain in the data frame.

Value

A "fastFMM" object containing parameters for fitting a functional mixed model using the FUI scheme. The object contains each of the passed args ('formula, data, ..., analytic'), with the exception of 'var'. Additional entries returned are:

1. 'out_index': locations in 'data' where functional domain exists
2. 'argvals': either 'argvals' as passed (if bootstrap) or a vector '1:L' where 'L' is the size of the functional domain

<code>new_fastFMMconc</code>	<i>Create a new "fastFMMconc" object</i>
------------------------------	--

Description

Create an object that contains parameters for fast univariate inference for concurrent models.

Usage

```
new_fastFMMconc(
  formula,
  data,
  subj_id,
  argvals,
  family,
  residuals,
  caic,
  randeffs,
  var,
  analytic,
  fun_covariates
)
```

Arguments

<code>formula</code>	Formula in 'lme4' formula syntax.
<code>data</code>	Data frame to fit.
<code>subj_id</code>	Character, name of variable containing IDs. Passed from 'fui'
<code>argvals</code>	List of points to fit on the functional domain. Only applies for the bootstrap case (i.e., 'analytic = FALSE').
<code>family</code>	Character, GLM family of the response. Passed from 'fui'.
<code>residuals</code>	Logical, indicates whether residuals are saved from unsmoothed LME. Passed from 'fui'.
<code>caic</code>	Logical, indicates cAIC calculation return. Defaults to 'FALSE'.
<code>randeffs</code>	Logical, indicates whether random effect estimates are returned. Passed from 'fui'.

var	Logical, indicates whether to calculate variance. Passed from 'fui'.
analytic	Logical, indicates whether variance will be calculated analytically. Passed from 'fui'.
fun_covariates	Character vector of functional covariate names.

Value

A "fastFMMconc" object containing parameters to fit a concurrent mixed model using the FUI scheme. This function is called within 'fui' if indicated by 'concurrent = TRUE'. Fields are shared with 'fastFMM' objects, with the addition of 'fun_covariates'.

plot_fui	<i>Default FUI plotting</i>
----------	-----------------------------

Description

Take a fitted fui object produced by `fastFMM::fui()` and plot the point estimates of fixed effects. When variance was calculated, the plot function also returns 95% pointwise and joint confidence intervals.

Usage

```
plot_fui(
  fuiobj,
  num_row = NULL,
  xlab = "Functional Domain",
  title_names = NULL,
  ylim = NULL,
  align_x = NULL,
  x_rescale = 1,
  y_val_lim = 1.1,
  y_scal_orig = 0.05,
  return = FALSE
)
```

Arguments

fuiobj	A object returned from the fui function
num_row	An integer that specifies the number of rows the plots will be displayed on. Defaults to $p/2$, where p is the number of predictors.
xlab	A string that specifies the x-axis title (i.e., for the functional domain). Defaults to "Functional Domain"
title_names	A vector of strings that has length equal to number of covariates (plus intercept if relevant). Allows one to change the titles of the plots. Defaults to NULL which uses the variable names in the dataframe for titles.

y_lim	A 2-dimensional vector that specifies limits of the y-axis in plots.
align_x	A scalar: aligns the plot to a certain point on the functional domain and sets this as 0. This is particularly useful if the functional domain is time. Defaults to 0.
x_rescale	A scalar: rescales the x-axis of plots which is especially useful if time is the functional domain and one wishes to, for example, account for the sampling rate. Defaults to 1.
y_val_lim	A positive scalar that extends the y-axis by a factor for visual purposes. Defaults to 1.10. Typically does not require adjustment.
y_scal_orig	A positive scalar that determines how much to reduce the length of the y-axis on the bottom. Defaults to 0.05. Typically does not require adjustment.
return	Logical, indicating whether to return the data frame with the coefficient estimates and 95% confidence intervals (CIs). Defaults to FALSE.

Value

Plots of point estimates and CIs. If `return = TRUE`, also returns a list where each element is a data frame with the coefficient estimates and 95% confidence intervals (CIs).

Author(s)

Gabriel Loewinger <gloewinger@gmail.com>, Erjia Cui <ecui@umn.edu>, Al W Xin <axin@andrew.cmu.edu>

References

Cui, E., Leroux, A., Smirnova, E., Crainiceanu, C. (2022). Fast Univariate Inference for Longitudinal Functional Models. *Journal of Computational and Graphical Statistics*, 31(1), 219-230.

Loewinger, G., Cui, E., Lovinger, D., Pereira, F. (2024). A Statistical Framework for Analysis of Trial-Level Temporal Dynamics in Fiber Photometry Experiments. *eLife*, 95802.

Examples

```
library(refund)
set.seed(1)
DTI_use <- DTI[DTI$ID %in% sample(DTI$ID, 10),]
DTI_use <- cbind(DTI_use[, c("visit", "sex", "ID")], data.frame(DTI_use$cca))
fit_dti <- fui(
  cca ~ visit + sex + (1 | ID),
  data = DTI_use
)
plot_fui(fit_dti)
```

select_knots	<i>select_knots.R from refund package</i>
--------------	---

Description

Copied from [select_knots](https://rdr.io/cran/refund/src/R/select_knots.R) because the original is not exported for use.

Usage

```
select_knots(t, knots = 10, p = 3, option = "equally-spaced")
```

Arguments

t	Numeric
knots	Numeric scalar or vector, the number/numbers of knots or the vector/vectors of knots for each dimension. Default = 10
p	Numeric, the degrees of B-splines. Default = 3.
option	Character, knot spacing, can be "equally-spaced" or "quantile"

Value

Vector of specified knots

Index

* **datasets**

d2pvt, [3](#)

lick, [7](#)

bdiag_m, [2](#)

cov_organize_start, [2](#)

d2pvt, [3](#)

fui, [4](#)

lick, [7](#)

new_fastFMM, [8](#)

new_fastFMMconc, [9](#)

plot_fui, [10](#)

select_knots, [12](#)