

Package ‘glmmFEL’

September 18, 2026

Type Package

Title Generalized Linear Mixed Models via Fully Exponential Laplace in EM

Version 1.0.6

Date 2026-09-18

Description Fit generalized linear mixed models (GLMMs) with normal random effects using first-order Laplace, fully exponential Laplace (FEL) with mean-only corrections, and FEL with mean and variance-diagonal corrections in the E-step of an expectation-maximization (EM) algorithm. The current development version provides a matrix-based interface (y, X, Z) and supports binary logit and probit, and Poisson log-link models. An EM framework is used to update fixed effects, random effects, and a single variance component τ^2 for $G = \tau^2 I$, with staged approximations (Laplace $\rightarrow$ FEL mean-only $\rightarrow$ FEL full) for efficiency and stability. A pseudo-likelihood engine `glmmFEL_pl()` implements the working-response / working-weights linearization approach of Wolfinger and O'Connell (1993) <doi:10.1080/00949659308811554>, and is adapted from the implementation used in the 'RealVAMS' package (Broatch, Green, and Karl (2018)) <doi:10.32614/RJ-2018-033>. The FEL implementation follows Karl, Yang, and Lohr (2014) <doi:10.1016/j.csda.2013.11.019> and related work (e.g., Tierney, Kass, and Kadane (1989) <doi:10.1080/01621459.1989.10478824>; Rizopoulos, Verbeke, and Lesaffre (2009) <doi:10.1111/j.1467-9868.2008.00704.x>; Steele (1996) <doi:10.2307/2532845>). Package code was drafted with assistance from generative AI tools.

License GPL-3

Encoding UTF-8

Imports Matrix, stats, methods

Suggests testthat (>= 3.0.0), lme4, numDeriv

Config/testthat/edition 3

NeedsCompilation no

Config/roxygen2/version 8.1.0

Author Andrew T. Karl [cre, aut] (ORCID:
 <https://orcid.org/0000-0002-5933-8706>)
Maintainer Andrew T. Karl <akar1@asu.edu>
Repository CRAN
Date/Publication 2026-09-18 17:20:02 UTC

Contents

glmmFEL-package	2
coef.glmmFELMod	4
fitted.glmmFELMod	4
glmmFEL	5
glmmFEL-benchmarks	8
logLik.glmmFELMod	13
predict.glmmFELMod	14
print.glmmFELMod	15
print.summary.glmmFELMod	15
summary.glmmFELMod	16
vcov.glmmFELMod	16
Index	17

glmmFEL-package	<i>glmmFEL: Generalized Linear Mixed Models via Fully Exponential Laplace in EM</i>
-----------------	---

Description

glmmFEL fits generalized linear mixed models (GLMMs) with normal random effects using a matrix-based interface where users supply (y, X, Z) directly. Model fitting is performed with an EM algorithm whose E-step can be approximated using first-order Laplace or fully exponential Laplace (mean-only or mean and variance-diagonal corrections), and includes pseudo-likelihood alternatives based on working-response / working-weights linearization.

The matrix interface supports multiple-membership designs with a single variance component. It does not implement all covariance structures in the cited papers. See `glmmFEL()` for the precise approximation and output limits.

Supported families in this branch:

- `family = stats::binomial(link = "probit")` (binary probit),
- `family = stats::binomial(link = "logit")` (binary logit),
- `family = stats::poisson(link = "log")` (Poisson log-link).

Approximations

`glmmFEL()` supports:

- "Laplace": first-order Laplace approximation,
- "FE_mean": fully exponential Laplace corrections to $\hat{\eta}$ only,
- "FE_full" (or "FE"): mean corrections plus the posterior variance-diagonal corrections needed for the scalar variance update,
- "RSPL" / "MSPL": restricted/marginal pseudo-likelihood (working response / working weights).

Output

`glmmFEL()` returns an object of class "glmmFELMod" containing:

- beta: fixed-effect estimates,
- eta: empirical Bayes predictions of random effects,
- tau2: the scalar variance component,
- G: $q \times q$ covariance matrix (diagonal in this branch),
- var_eta: approximate posterior covariance, with diagonal-only FEL covariance corrections in FE_full,
- vcov_beta: approximate covariance of beta when available,
- convergence: iteration counts and flags.

Author(s)

Maintainer: Andrew T. Karl <akar1@asu.edu> ([ORCID](#))

Authors:

- Andrew T. Karl <akar1@asu.edu> ([ORCID](#))

References

- Broatch, J., Green, J. G., & Karl, A. T. (2018). RealVAMS: An R Package for Fitting a Multivariate Value-added Model (VAM). *The R Journal*, 10(1), 22–30. doi:[10.32614/RJ2018033](#)
- Karl, A. T., Yang, Y., & Lohr, S. L. (2014). Computation of maximum likelihood estimates for multiresponse generalized linear mixed models with non-nested, correlated random effects. *Computational Statistics & Data Analysis*, 73, 146–162. doi:[10.1016/j.csda.2013.11.019](#)
- Rizopoulos, D., Verbeke, G., & Lesaffre, E. (2009). Fully exponential Laplace approximations in joint models for longitudinal and survival data. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 71(3), 637–654. doi:[10.1111/j.14679868.2008.00704.x](#)
- Rizopoulos, D. (2010). JM: An R package for the joint modeling of longitudinal and time-to-event data. *Journal of Statistical Software*, 35(9), 1–33. doi:[10.18637/jss.v035.i09](#)
- Steele, B. M. (1996). A modified EM algorithm for estimation in generalized mixed models. *Biometrics*, 52(4), 1295–1310. doi:[10.2307/2532845](#)
- Tierney, L., Kass, R. E., & Kadane, J. B. (1989). Fully exponential Laplace approximations to expectations and variances of nonpositive functions. *Journal of the American Statistical Association*, 84(407), 710–716. doi:[10.1080/01621459.1989.10478824](#)

Wolfinger, R., & O'Connell, M. (1993). Generalized linear mixed models: a pseudo-likelihood approach. *Journal of Statistical Computation and Simulation*, 48(3–4), 233–243. doi:10.1080/00949659308811554

See Also

[glmmFEL_pl\(\)](#) for the pseudo-likelihood engines; [glmmFEL-benchmarks](#) for the simulation appraisal.

<code>coef.glmmFELMod</code>	<i>Extract model coefficients (fixed effects)</i>
------------------------------	---

Description

Extract model coefficients (fixed effects)

Usage

```
## S3 method for class 'glmmFELMod'
coef(object, ...)
```

Arguments

<code>object</code>	A <code>glmmFELMod</code> object.
<code>...</code>	Unused.

Value

A named numeric vector of estimated fixed-effect regression coefficients (on the linear predictor scale). Names correspond to columns of the fixed-effects design matrix X when available.

<code>fitted.glmmFELMod</code>	<i>Extract fitted values</i>
--------------------------------	------------------------------

Description

Extract fitted values

Usage

```
## S3 method for class 'glmmFELMod'
fitted(object, type = c("response", "link"), ...)
```

Arguments

object	A glmmFELMod object.
type	Either "link" (linear predictor) or "response".
...	Unused.

Value

A numeric vector of fitted values. If `type = "link"`, the fitted linear predictor values are returned. If `type = "response"`, the fitted mean response values on the response scale are returned. The length equals the number of observations in the fitted object.

glmmFEL	<i>Fit GLMMs via Laplace and fully exponential Laplace (matrix interface)</i>
---------	---

Description

`glmmFEL()` fits a generalized linear mixed model (GLMM) with multivariate normal random effects using EM-type algorithms and likelihood approximations:

- first-order Laplace (`approx = "Laplace"`),
- fully exponential corrections to the random-effects mean (`approx = "FE_mean"`),
- fully exponential corrections to both mean and variance diagonals (`approx = "FE_full" / "FE"`),
- pseudo-likelihood / PL linearization (`approx = "RSPL" or "MSPL"`) via `glmmFEL_pl()`.

The interface is matrix-based: provide the response y , fixed-effects design matrix X , and random-effects design matrix Z .

Random effects are assumed $\eta \sim N(0, G)$ with a **single** variance component

$$G = \tau^2 I_q,$$

allowing arbitrary (including multi-membership) Z while keeping the variance update simple and stable.

Usage

```
glmmFEL(
  y,
  X,
  Z,
  family = stats::binomial(link = "probit"),
  approx = c("FE", "Laplace", "FE_mean", "FE_full", "RSPL", "MSPL"),
  max_iter = 200,
  tol = 1e-06,
  control = list()
)
```

Arguments

y	Numeric response vector of length n . For family = "binomial_probit" / binomial(link = "probit") or family = "binomial_logit" / binomial(link = "logit"), values must be 0 or 1; grouped binomial counts and trial weights are not supported. Poisson responses must be nonnegative integer counts.
X	Fixed-effects design matrix of dimension $n \times p$. May be a base R matrix or a matrix-like object; it is internally coerced to a base numeric matrix. Must have at least one column, finite entries, and full column rank.
Z	Random-effects design matrix of dimension $n \times q$. May be a base R matrix or a Matrix object. Internally it is coerced to a sparse "dgCMatrix" where possible (to preserve sparsity). Must have at least one column (no purely fixed-effects models).
family	Either a character string or a <code>stats::family</code> object indicating the model family. The argument is resolved via <code>glmmfe_resolve_family()</code> .
approx	Approximation type, resolved via <code>glmmfe_resolve_approx()</code> . Accepted values (case-insensitive) include: • "Laplace" – first-order Laplace approximation, • "FE_mean" – staged algorithm: Laplace phase then FE mean corrections, • "FE" / "FE_full" – staged algorithm: Laplace phase, then FE mean, then FE covariance corrections (default), • "RSPL" – restricted pseudo-likelihood (REML-style) linearization, • "MSPL" – marginal pseudo-likelihood (ML-style) linearization.
max_iter	Maximum number of EM iterations (outer iterations over β and τ^2). Can be overridden by <code>control\$em_max_iter</code> .
tol	Baseline convergence tolerance for the EM algorithm. The staged thresholds default to: • Laplace stage: <code>tol_laplace = 10 * tol</code>, • FE-mean stage: <code>tol_fe_mean = 3 * tol</code>, • FE-full stage: <code>tol_fe_full = tol</code>. You can override these via <code>control\$tol_laplace</code> , <code>control\$tol_fe_mean</code> , and <code>control\$tol_fe_full</code> .
control	List of optional control parameters. Recognized entries include: • <code>em_max_iter</code>, <code>em_tol</code>, • <code>tol_laplace</code>, <code>tol_fe_mean</code>, <code>tol_fe_full</code>, • <code>eta_max_iter</code>, <code>eta_tol_grad</code>, • <code>beta_max_iter</code>, <code>beta_tol</code>, • <code>tau2_init</code> (initial value for τ^2), • <code>vc_eps</code> (lower bound for τ^2), • <code>max_nq_mem</code> (memory guard for FE trace intermediates), • <code>verbose</code> (logical), • <code>beta_step_max</code> (max Newton step size for beta; default 2), • <code>beta_ls_max_iter</code> (max line-search halvings; default 12), • <code>beta_hess_ridge_init</code> (initial ridge for Hessian; default 1e-8), • <code>beta_hess_ridge_max</code> (max ridge; default 1e2)

Details

Laplace denotes an EM algorithm using posterior modes in its expected score and Gaussian posterior second moments. It is not direct maximization of the first-order Laplace marginal likelihood (as in `lme4::glmer`). `FE_mean` adds mean and expected-score corrections. `FE_full` also corrects the posterior variance diagonals needed to update the single variance component; off-diagonal entries of `var_eta` remain at their Laplace values. Thus "full" distinguishes the two implemented EM stages, not a correction of every entry of an arbitrary posterior covariance matrix.

Returned moments are recomputed at the final parameter estimates. A failed inner solve or exhausted iteration budget produces a warning and a false convergence flag. Always inspect convergence before interpreting estimates. Approximate posterior variances can become invalid in sparse/extreme settings; such fits are flagged rather than reported as converged. A strictly separating fixed-effect direction found at initialization is rejected for binary data, since no finite estimate exists. This sufficient check is not a complete detector of all forms of quasi-separation.

`vcov_beta` is a joint-Hessian approximation conditional on the fitted variance component; it is not a fully exponential observed-information estimate and does not incorporate variance-component estimation uncertainty. For EM fits, `logLik` evaluates the first-order Laplace marginal likelihood at the returned estimates, which need not maximize it. For PL fits it is unavailable (NA); `working_logLik` stores the working Gaussian objective. These objectives should not be used for cross-method AIC comparisons. Fitted response values plug the estimated random effects into the inverse link; they are not posterior predictive averages over random effects.

See [glmmFEL-benchmarks](#) for a reproducible simulation appraisal and its limits.

Value

A fitted model object of class `glmmFELMod`, with estimates `beta`, `tau2`, random-effect predictions `eta`, approximate covariance matrices, and convergence. EM convergence diagnostics include `em_converged`, `reason`, `mode_converged`, `mode_gradient`, `beta_converged`, and `moments_valid`. PL fits report both `pql_converged` and `em_converged`.

References

- Broatch, J., Green, J. G., & Karl, A. T. (2018). RealVAMS: An R Package for Fitting a Multivariate Value-added Model (VAM). *The R Journal*, 10(1), 22–30. doi:10.32614/RJ2018033
- Karl, A. T., Yang, Y., & Lohr, S. L. (2014). Computation of maximum likelihood estimates for multiresponse generalized linear mixed models with non-nested, correlated random effects. *Computational Statistics & Data Analysis*, 73, 146–162. doi:10.1016/j.csda.2013.11.019
- Rizopoulos, D., Verbeke, G., & Lesaffre, E. (2009). Fully exponential Laplace approximations in joint models for longitudinal and survival data. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 71(3), 637–654. doi:10.1111/j.14679868.2008.00704.x
- Rizopoulos, D. (2010). JM: An R package for the joint modeling of longitudinal and time-to-event data. *Journal of Statistical Software*, 35(9), 1–33. doi:10.18637/jss.v035.i09
- Steele, B. M. (1996). A modified EM algorithm for estimation in generalized mixed models. *Biometrics*, 52(4), 1295–1310. doi:10.2307/2532845
- Tierney, L., Kass, R. E., & Kadane, J. B. (1989). Fully exponential Laplace approximations to expectations and variances of nonpositive functions. *Journal of the American Statistical Association*, 84(407), 710–716. doi:10.1080/01621459.1989.10478824

Wolfinger, R., & O’Connell, M. (1993). Generalized linear mixed models: a pseudo-likelihood approach. *Journal of Statistical Computation and Simulation*, 48(3–4), 233–243. doi:10.1080/00949659308811554

Examples

```
## Example 1: Simulated probit random-intercept GLMM (matrix interface)
set.seed(1)
n_id <- 30
m_per_id <- 6
n <- n_id * m_per_id
id <- factor(rep(seq_len(n_id), each = m_per_id))
x <- rnorm(n)
X <- model.matrix(~ x)
Z <- Matrix::sparseMatrix(i = seq_len(n),
                          j = as.integer(id),
                          x = 1,
                          dims = c(n, n_id))

beta_true <- c(0.2, 0.7)
tau2_true <- 0.5
eta_true <- rnorm(n_id, sd = sqrt(tau2_true))
lp <- as.vector(X %*% beta_true + Z %*% eta_true)
y <- rbinom(n, 1, pnorm(lp))

fit <- glmmFEL(y, X, Z, family = "binomial_probit", approx = "Laplace")
fit$beta
fit$tau2

## Example 2: Get X, y, Z from an lme4 formula (without a glmmFEL formula wrapper)

if (requireNamespace("lme4", quietly = TRUE)) {
  dat <- data.frame(y = y, x = x, id = id)
  lf <- lme4::lFormula(y ~ x + (1 | id), data = dat)
  X_lme4 <- lf$X
  Z_lme4 <- Matrix::t(lf$reTrms$Zt)
  y_lme4 <- lf$fr$y

  fit2 <- glmmFEL(y_lme4, X_lme4, Z_lme4, family = "binomial_probit", approx = "Laplace")
}
```

Description

A fixed-design simulation comparison of estimation accuracy, convergence, and computing time. The results describe the settings studied, not a general ranking of GLMM estimators.

Details

These static results were generated for version 1.0.6 on 2026-09-18 from a prespecified design: 2,400 datasets and 22,400 applicable fits. The two glmmML methods have no probit implementation; their 1,600 non-applicable rows are recorded separately. No simulation runs during installation or package checks.

Design

Each family/setting combination has 200 independent datasets, paired across methods. The covariate is standard normal; Gaussian random intercepts have variance 0.5. The true slope is 0.5, and the intercept is -0.5 for binary responses and $\log(2)$ for Poisson. Binary responses are Bernoulli observations.

Setting	Clusters	Observations per cluster
small	30	4
rich	30	20
unbalanced	30	3, 5, 10, 20, 32 repeated six times
many_small	60	4

Methods and diagnostics

Modal EM, FEL mean/full, and MSPL/RSPL are the package's five algorithms. The package's Laplace uses modes in EM; it is not direct marginal Laplace optimization. **lme4** uses glmer with nAGQ=1 or 9. **glmmML** uses Laplace or 20-point Gaussian quadrature. Quadrature is an implementation reference, not simulated truth. MASS::glmmPQL uses its ML working-model convention with residual scale fixed at one.

glmmFEL EM has a 400-iteration budget and $1e-6$ tolerance for all three stages. Its PL fits have 100 outer and 100 inner iterations. MASS PQL has 100 outer iterations; its reported convergence is a proxy based on returning before the cap, because the API exposes no outer convergence flag. Boundary fits are reported separately from convergence. An iteration cap counts as failure, even if the numerical estimates might already be practically close.

Interpretation

FEL was not an overall winner in these settings. It often reduced variance-component bias relative to modal EM or pseudo-likelihood, but the accompanying sampling variance, iteration cost, and failures matter. Bias reduction did not imply a smaller RMSE. With well-informed clusters, FEL estimates were generally close to quadrature estimates, including binary settings.

Sparse binary data were unfavorable. For logit responses, FEL full converged in 80/200 small-cluster datasets and 53/200 many-small-cluster datasets at the 400-iteration budget. Its variance squared error exceeded AGQ9 on their common converged datasets. One sparse probit fit produced a variance estimate around $1.7e155$ before being flagged for invalid posterior moments. Failed fits must not be used for inference, even when an estimate is returned.

Rich Poisson data exposed slow staged EM: only 72/200 FEL-full fits converged within 400 iterations. The post-hoc sensitivity increased only the budget to 2,000 on the first 40 seeds of every cell. In rich Poisson data, FEL-full convergence rose from 12/40 to 40/40. Its variance RMSE was 0.10449, versus 0.10450 for both lme4 AGQ9 and lme4 Laplace, and 0.10434 for MASS PQL.

Slope RMSE was approximately 0.02545 for all four. The FEL-full minus AGQ9 variance MSE difference was $-2.09\text{e-}6$ (MCSE $1.11\text{e-}6$): a very small practical difference in a 40-replicate sensitivity, not evidence of broad superiority.

Thus the study supports accurate FEL approximation in the tested Poisson settings, but does not establish an advantage over these competing solvers, even with rich data. Increasing the budget also improved binary convergence, yet sparse-logit failures and excess variance error remained. The independent random-intercept structure particularly favors available quadrature solvers; the small crossed-design diagnostic does not settle accuracy for the broader nonnested problems motivating the literature.

Convergence across all settings

Counts below are out of 200 applicable datasets per cell. NA means that the method does not implement the link, not optimizer failure.

Bernoulli-logit

Method	small	rich	unbalanced	many small
Modal EM	126	199	199	146
FEL mean	108	199	196	92
FEL full	80	198	195	53
MSPL	151	199	199	167
RSPL	160	200	199	172
lme4 Laplace	199	200	200	198
lme4 AGQ9	199	200	200	200
glmmML Laplace	200	200	200	200
glmmML GHQ20	200	200	200	200
MASS PQL	200	200	200	200

Bernoulli-probit

Method	small	rich	unbalanced	many small
Modal EM	184	200	200	196
FEL mean	165	200	200	187
FEL full	154	200	200	181
MSPL	186	200	200	196
RSPL	196	200	200	198
lme4 Laplace	200	200	200	199
lme4 AGQ9	200	200	200	200
glmmML Laplace	NA	NA	NA	NA
glmmML GHQ20	NA	NA	NA	NA
MASS PQL	200	200	200	200

Poisson-log

Method	small	rich	unbalanced	many small
--------	-------	------	------------	------------

Modal EM	200	187	194	200
FEL mean	200	88	142	200
FEL full	200	72	117	199
MSPL	200	200	200	200
RSPL	200	200	200	200
lme4 Laplace	200	200	200	200
lme4 AGQ9	200	200	200	200
glmmML Laplace	200	200	200	200
glmmML GHQ20	200	200	200	200
MASS PQL	200	200	200	200

Variance estimation with rich clusters

For the 30-by-20 setting, the following are bias, empirical sampling variance, and RMSE of the estimated random-intercept variance (truth 0.5). Each method uses its own converged subset; use the counts above and the paired contrasts below. Complete results for every setting and all three parameters, including Monte Carlo SEs, are in the accompanying CSV files.

Bernoulli-logit

Method	Bias	Variance	RMSE
Modal EM	-0.0687	0.0452	0.2230
FEL mean	-0.0283	0.0587	0.2434
FEL full	-0.0171	0.0596	0.2441
MSPL	-0.0687	0.0452	0.2230
RSPL	-0.0446	0.0500	0.2274
lme4 Laplace	-0.0314	0.0590	0.2444
lme4 AGQ9	-0.0222	0.0608	0.2469
glmmML Laplace	-0.0313	0.0591	0.2445
glmmML GHQ20	-0.0222	0.0607	0.2468
MASS PQL	-0.0708	0.0458	0.2249

Bernoulli-probit

Method	Bias	Variance	RMSE
Modal EM	-0.0464	0.0239	0.1609
FEL mean	-0.0035	0.0327	0.1804
FEL full	-0.0004	0.0336	0.1828
MSPL	-0.0502	0.0235	0.1609
RSPL	-0.0276	0.0257	0.1622
lme4 Laplace	-0.0008	0.0333	0.1820
lme4 AGQ9	-0.0010	0.0333	0.1820
glmmML Laplace	NA	NA	NA
glmmML GHQ20	NA	NA	NA
MASS PQL	-0.0503	0.0235	0.1609

Poisson-log

Method	Bias	Variance	RMSE
Modal EM	-0.0510	0.0096	0.1104
FEL mean	-0.1221	0.0045	0.1392
FEL full	-0.1423	0.0028	0.1516
MSPL	-0.0351	0.0129	0.1187
RSPL	-0.0175	0.0139	0.1188
lme4 Laplace	-0.0246	0.0137	0.1192
lme4 AGQ9	-0.0239	0.0137	0.1193
glmmML Laplace	-0.0246	0.0137	0.1192
glmmML GHQ20	-0.0239	0.0137	0.1193
MASS PQL	-0.0351	0.0129	0.1187

Paired variance comparisons

FEL full minus lme4 AGQ9 squared-error differences, on common converged datasets only. Negative differences favor FEL full. MCSE is the Monte Carlo standard error of that difference. Small differences relative to MCSE do not support a stable ranking; these are exploratory, multiple comparisons.

Family/setting	Pairs	Difference	MCSE
logit: small	80	0.209	0.0551
logit: rich	198	0.000136	6.19e-05
logit: unbalanced	195	0.00182	0.000444
logit: many small	53	0.088	0.0112
probit: small	154	0.0427	0.0198
probit: rich	200	0.000277	0.000102
probit: unbalanced	200	0.00117	0.000267
probit: many small	181	0.0316	0.0145
Poisson: small	200	8.28e-05	3.03e-05
Poisson: rich	72	2.49e-06	2.87e-07
Poisson: unbalanced	117	-2.36e-05	3.74e-06
Poisson: many small	199	0.000112	2.31e-05

Post-hoc iteration-budget sensitivity

This smaller extension uses the same first 40 seeds in every cell, chosen by replicate number rather than outcome. It increases only the three EM methods' budgets to 2,000 and reuses unchanged main-run comparator fits. Below are FEL-full convergence counts (out of 40) and paired variance MSE differences versus AGQ9. It is exploratory and does not replace the 200-replicate primary study. Other methods and parameters are in the separate sensitivity CSV files.

Family/setting	400	2000	Pairs	Difference	MCSE
logit: small	20	31	31	0.0601	0.0184
logit: rich	40	40	40	0.000323	0.000226
logit: unbalanced	40	40	40	0.00199	0.0007
logit: many small	9	30	30	0.0283	0.0124
probit: small	29	39	39	0.0282	0.0222

probit: rich	40	40	40	0.000161	0.00017
probit: unbalanced	40	40	40	0.00188	0.000819
probit: many small	37	40	40	0.0136	0.0103
Poisson: small	40	40	40	7.44e-05	4.47e-05
Poisson: rich	12	40	40	-2.09e-06	1.11e-06
Poisson: unbalanced	27	40	40	3.74e-05	3.11e-05
Poisson: many small	40	40	40	4e-05	3.34e-05

Limits and reproducibility

This study varies cluster information and balance at one effect size and one random-effect variance. It does not establish performance for rare outcomes, larger variances, misspecified random effects, grouped binomial data, or arbitrary nonnested designs. A separate 24-dataset crossed/signed-membership diagnostic had 100 converged fits among 120 attempts; all 20 failures were iteration caps. Four seeds per family/design do not support accuracy claims.

Main accuracy summaries condition on convergence. Pairing successful fits does not remove selection bias from failures. The files also report all finite returned estimates separately, including non-converged values. Wald coverage is exploratory: the package’s Hessian standard errors condition on the estimated variance and are not full FEL observed-information standard errors. Random-effect prediction loss compares the means or modes each method returns.

Timings include contention from eight local workers; they are not isolated microbenchmarks. The static files under `system.file("benchmarks", package="glmmFEL")` include full summaries, paired contrasts, failure counts, package versions, the scenario grid, source/seed provenance, and column definitions. The maintainer retains raw per-fit results and the frozen simulation scripts.

See Also

[glmmFEL\(\)](#)

Examples

```
results <- read.csv(system.file("benchmarks", "summary.csv", package = "glmmFEL"))
subset(results, family == "poisson_log" & setting == "rich" & parameter == "tau2",
       select = c(method, converged, bias, variance, rmse, bias_mcse))
```

logLik.glmmFELMod	<i>Extract log-likelihood (approximate)</i>
-------------------	---

Description

Extract log-likelihood (approximate)

Usage

```
## S3 method for class 'glmmFELMod'
logLik(object, ...)
```

Arguments

object	A glmmFELMod object.
...	Unused.

Value

An object of class "logLik" giving the approximate log-likelihood stored in object\$logLik (NA for PL fits), with attributes "df" (effective number of parameters, taken as length(beta) + 1 for the variance component) and "nobs" (number of observations).

predict.glmmFELMod	<i>Predict from a fitted glmmFEL model</i>
--------------------	--

Description

Predict from a fitted glmmFEL model

Usage

```
## S3 method for class 'glmmFELMod'
predict(object, newdata = NULL, type = c("response", "link"), ...)
```

Arguments

object	A glmmFELMod object.
newdata	Not supported in this branch (matrix interface only).
type	Either "link" or "response".
...	Unused.

Value

A numeric vector of predictions. If type = "link", predictions are returned on the linear predictor scale. If type = "response", predictions are returned on the response scale. The length equals the number of observations used to fit the model. Supplying newdata triggers an error in this matrix-only branch.

print.glmmFELMod	<i>Print a glmmFEL model object</i>
------------------	-------------------------------------

Description

Print a glmmFEL model object

Usage

```
## S3 method for class 'glmmFELMod'  
print(x, ...)
```

Arguments

x	A glmmFELMod object.
...	Unused.

Value

Returns x invisibly (a glmmFELMod object), called for its side effect of printing model information to the console.

print.summary.glmmFELMod	<i>Print a summary.glmmFELMod object</i>
--------------------------	--

Description

Print a summary.glmmFELMod object

Usage

```
## S3 method for class 'summary.glmmFELMod'  
print(x, ...)
```

Arguments

x	A summary.glmmFELMod object.
...	Unused.

Value

Returns x invisibly (a summary.glmmFELMod object), called for its side effect of printing the summary to the console.

summary.glmmFELMod	<i>Summary for a glmmFEL model object</i>
--------------------	---

Description

Summary for a glmmFEL model object

Usage

```
## S3 method for class 'glmmFELMod'
summary(object, ...)
```

Arguments

object	A glmmFELMod object.
...	Unused.

Value

An object of class summary.glmmFELMod (a list) containing summary information for the fitted model, including the call, family, approximation label, approximate log-likelihood, estimated variance component tau2, a coefficient table for fixed effects (and standard errors when available), number of observations, and convergence information.

vcov.glmmFELMod	<i>Extract the covariance matrix of the fixed effects</i>
-----------------	---

Description

Extract the covariance matrix of the fixed effects

Usage

```
## S3 method for class 'glmmFELMod'
vcov(object, ...)
```

Arguments

object	A glmmFELMod object.
...	Unused.

Value

A variance-covariance matrix for the estimated fixed-effect regression coefficients. Row and column names correspond to coefficient names when available.

Index

`coef.glmmFELMod`, [4](#)

`fitted.glmmFELMod`, [4](#)

`glmmfe_resolve_approx()`, [6](#)

`glmmfe_resolve_family()`, [6](#)

`glmmFEL`, [5](#)

`glmmFEL()`, [2](#), [3](#), [13](#)

`glmmFEL-benchmarks`, [4](#), [7](#), [8](#)

`glmmFEL-package`, [2](#)

`glmmFEL_pl()`, [4](#), [5](#)

`logLik.glmmFELMod`, [13](#)

`predict.glmmFELMod`, [14](#)

`print.glmmFELMod`, [15](#)

`print.summary.glmmFELMod`, [15](#)

`stats::family`, [6](#)

`summary.glmmFELMod`, [16](#)

`vcov.glmmFELMod`, [16](#)