

Package ‘exnexSurv’

September 10, 2026

Type Package

Title Bayesian EXNEX Models for Survival Analysis in Basket Trials

Version 1.3.0

Description Implements the Bayesian Exchangeable Non-Exchangeable (EXNEX) framework for right-censored log-normal survival data in basket trials. Based on 'Rcpp' and 'RcppArmadillo', the package provides a fast Gibbs sampler supporting EXNEX, complete pooling, and no pooling models to facilitate methodological comparisons and simulation studies.

License MIT + file LICENSE

URL <https://github.com/victorney/exnexSurv>,
<https://victorney.github.io/exnexSurv/>

BugReports <https://github.com/victorney/exnexSurv/issues>

Imports Rcpp (>= 1.1.1), hardhat (>= 1.4.0), checkmate (>= 2.1.0)

Suggests testthat (>= 3.3.0), survival (>= 3.8), knitr, rmarkdown,
ggplot2, bayesplot, viridisLite

VignetteBuilder knitr

LinkingTo Rcpp, RcppArmadillo

Encoding UTF-8

Config/roxygen2/version 8.1.0

NeedsCompilation yes

Author Victor Ney [aut, cre]

Maintainer Victor Ney <victorney@ime.usp.br>

Repository CRAN

Date/Publication 2026-09-10 14:20:15 UTC

Contents

compare_waic	2
compute_waic	2

exnex_surv	3
median_survival	6
plot.exnex_surv	7
plot.survival_exnex	7
print.exnex_surv	8
probability_superiority	8
rmst	9
simulate_data	10
summary.exnex_surv	11
survival_curves	12

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

compare_waic	<i>Compare multiple exnex_surv fits by WAIC</i>
--------------	---

Description

Computes [compute_waic()] for each supplied fit and reports them in a single ‘data.frame’, sorted by ascending WAIC.

Usage

```
compare_waic(..., digit = 2)
```

Arguments

- ... Two or more fitted ‘exnex_surv’ objects.
- digit Number of decimal places for the reported statistics.

Value

A ‘data.frame’ with one row per model and columns ‘model’, ‘waic’, ‘se_elpd_waic’, ‘lpd’, ‘p_waic’, ‘elpd_waic’.

compute_waic	<i>Pointwise log-likelihood and WAIC from an exnex_surv fit</i>
--------------	---

Description

Computes the Watanabe-Akaike Information Criterion (WAIC) for a fitted ‘exnex_surv’ model. WAIC is a fully Bayesian information criterion that uses the posterior draws of the log-likelihood and adds a penalty for effective number of parameters p_{waic} .

Usage

```
compute_waic(fit, ...)
```

Arguments

`fit` A fitted ‘exnex_surv’ object.
`...` Unused.

Details

For each observation i , with posterior draws indexed by $s = 1, \dots, S$, we compute the pointwise log-likelihood $\log p(y_i | \theta^{(s)})$, where the log-normal AFT model gives for an observed event ($\delta_i = 1$)

$$\log p(y_i | \theta^{(s)}) = \log \phi \left(\frac{\log t_i - \eta_i^{(s)}}{\sigma^{(s)}} \right) - \log(\sigma^{(s)} t_i),$$

and for a censored observation ($\delta_i = 0$)

$$\log p(y_i | \theta^{(s)}) = \log \left[1 - \Phi \left(\frac{\log t_i - \eta_i^{(s)}}{\sigma^{(s)}} \right) \right],$$

where $\eta_i^{(s)}$ is the linear predictor built from the draws.

WAIC is then

$$lpd = \sum_i \log \left(\frac{1}{S} \sum_s \exp \log p(y_i | \theta^{(s)}) \right),$$

$$p_{waic} = \sum_i \text{var}_s (\log p(y_i | \theta^{(s)})),$$

$$\text{WAIC} = -2(lpd - p_{waic}).$$

Rows with extreme leverage (‘p_waic’ above a large threshold) may flag issues; the function returns a ‘pointwise’ matrix so users can investigate.

Value

A named list with elements ‘waic’, ‘se_elpd_waic’, ‘lpd’, ‘p_waic’, ‘elpd_waic’, and ‘pointwise’ (an ‘n’ by 3 matrix with columns ‘lpd’, ‘p_waic’, ‘waic’).

exnex_surv

Fit Bayesian EXNEX Survival Models

Description

Fits a Bayesian Exchangeable Non-Exchangeable (EXNEX) hierarchical model for right-censored log-normal survival data in basket trials using a data-augmented Gibbs sampler.

Usage

```

exnex_surv(x, ...)

## Default S3 method:
exnex_surv(x, ...)

## S3 method for class 'formula'
exnex_surv(
  formula,
  data,
  priors = list(),
  iter = 2000,
  warmup = 1000,
  chains = 1,
  parallel_chains = 1,
  group_col = NULL,
  seed = NULL,
  ...
)

## S3 method for class 'data.frame'
exnex_surv(
  x,
  y,
  priors = list(),
  iter = 2000,
  warmup = 1000,
  chains = 1,
  parallel_chains = 1,
  group_col = NULL,
  seed = NULL,
  ...
)

```

Arguments

<code>x</code>	An object containing the predictors (subgroup assignment). Can be a data frame or a formula.
<code>...</code>	Additional arguments passed to methods.
<code>formula</code>	A model formula with structure: ‘Surv(time, event) ~ group + covariates’. The first RHS variable is the group/basket assignment. Additional variables are covariates.
<code>data</code>	A data frame containing the variables in the formula.
<code>priors</code>	Optional named list of prior hyperparameters. Supported fields: <code>a_sigma</code> , <code>b_sigma</code> , <code>a_tau</code> , <code>b_tau</code> (inverse-Gamma shape and scale for the residual variance σ^2 and the between-basket variance τ^2); <code>p_mix</code> (EXNEX mixture weight, strictly between 0 and 1); <code>m_mu</code> , <code>v_mu</code> (prior mean and variance of the exchangeable center

	μ); m_{nex} , v_{nex} (prior mean and variance of the nonexchangeable component); v_{beta} (variance of the regression-coefficient prior). p_{mix} , m_{nex} , and v_{nex} each accept either a scalar, replicated across baskets, or a numeric vector of length K with one value per basket, matching the basket-specific notation of the model. All other fields are scalars. Absent fields keep the defaults (inverse-Gamma(2,2) for the variances, $p_{\text{mix}} = 0.5$, $m_{\text{mu}} = 0$, $v_{\text{mu}} = 1e4$, $m_{\text{nex}} = 0$, $v_{\text{nex}} = 1e4$, $v_{\text{beta}} = 1e4$); unknown fields are ignored.
iter	Total number of MCMC iterations. Default is 2000.
warmup	Number of warmup iterations to discard. Default is 1000. Posterior samples will have (iter - warmup) rows.
chains	Number of independent MCMC chains. Default is 1.
parallel_chains	Number of chains to run in parallel at the R level. Must be between 1 and 'chains'. Default is 1 (sequential chain execution).
group_col	Name of the column in 'x' that represents the basket/group assignment. If NULL (default), assumes the first column in x is the group.
seed	Random seed for reproducibility (optional).
y	A Surv object or matrix containing outcome (time and event status).

Details

For patient i , the model is

$$\log T_i = \theta_{g[i]} + X_i^T \beta + \varepsilon_i, \quad \varepsilon_i \sim \mathcal{N}(0, \sigma^2),$$

with a latent EXNEX hierarchy on the basket effects,

$$\theta_j \mid Z_j \sim Z_j \mathcal{N}(\mu, \tau^2) + (1 - Z_j) \mathcal{N}(m_{0j}, v_{0j}), \quad Z_j \sim \text{Bern}(p_{\text{exch},j}),$$

so that each basket either borrows strength from the exchangeable component $\mathcal{N}(\mu, \tau^2)$ or follows its own non-exchangeable prior $\mathcal{N}(m_{0j}, v_{0j})$. The first predictor variable in the model formula is interpreted as the basket/group assignment; additional variables are covariates in the linear predictor on the log-survival scale.

Censored event times are handled by data augmentation: a censored log-time is imputed from its truncated-Normal conditional distribution (inverse-CDF in log space) before the remaining parameters are updated with their conjugate full conditionals. The sampler keeps only the posterior draws of θ_j , β , and σ^2 .

Value

An object of class 'exnex_surv' containing posterior samples and metadata. Components include draws (a data frame with columns `theta_1`, ..., `theta_K`, `beta_1`, ..., `beta_P`, and `sigma2`), data (the processed data: `time`, `event`, `group`, `X`, `n`, `n_groups`, `n_covariates`, `cov_names`, `chain_seeds`), priors (the supplied priors), `resolved_priors` (defaults merged with overrides), `iter`, `warmup`, `chains`, and `blueprint`. Use `summary()`, `print()`, and `plot()` to inspect it.

Examples

```
# Small simulated dataset: three baskets, one covariate, ~25% censoring
set.seed(1)
n <- 90
group <- factor(rep(1:3, each = 30))
x1 <- rnorm(n)
eta <- rep(c(1.1, 1.6, 2.0), each = 30) + 0.5 * x1
true_time <- exp(eta + rnorm(n, 0, 0.6))
cens_time <- runif(n, 2, 9)
d <- data.frame(
  time = pmin(true_time, cens_time),
  event = as.integer(true_time <= cens_time),
  group = group,
  x1 = x1
)

# Fit with group effects and one covariate (small run for the example)
fit <- exnex_surv(
  survival::Surv(time, event) ~ group + x1,
  data = d,
  priors = list(p_mix = 0.7),
  iter = 300, warmup = 150, chains = 1
)
print(fit, show_trace = FALSE)
summary(fit)
```

median_survival	<i>Posterior median survival time</i>
-----------------	---------------------------------------

Description

For the log-normal AFT model the median survival time for a linear predictor η is $t_{med} = \exp(\eta)$, since $S(t) = 0.5$ when $\log t = \eta$. Posterior draws of θ and β therefore induce a posterior distribution of t_{med} whose quantiles are reported.

Usage

```
median_survival(fit, newdata = NULL, level = 0.95, ...)
```

Arguments

fit	A fitted ‘exnex_surv’ object.
newdata	Optional data frame (one row gives one median).
level	Credible-interval level (default ‘0.95’).
...	Unused.

Value

A 'data.frame' with columns 'group', 'median', 'lower', 'upper'.

plot.exnex_surv	<i>Plot parameter traces from an exnex_surv fit</i>
-----------------	---

Description

Generates one bayesplot traceplot per parameter.

Usage

```
## S3 method for class 'exnex_surv'
plot(x, parameters = NULL, ask = interactive(), ...)
```

Arguments

x	A fitted 'exnex_surv' object.
parameters	Optional character vector of parameter names to plot.
ask	Should R pause between plots? Defaults to 'interactive()'.
...	Unused.

Value

Invisibly returns 'x'.

plot.survival_exnex	<i>Plot survival curves from an exnex_surv fit</i>
---------------------	--

Description

Draws the posterior median and credible band of the survival function for each group. Requires 'ggplot2'.

Usage

```
## S3 method for class 'survival_exnex'
plot(x, ...)
```

Arguments

x	An object of class 'survival_exnex' returned by [survival_curves()].
...	Unused.

Value

Invisibly returns 'x'.

```
print.exnex_surv      Print an exnex_surv fit
```

Description

Prints a compact summary and optionally traceplots of posterior draws.

Usage

```
## S3 method for class 'exnex_surv'
print(x, show_trace = TRUE, parameters = NULL, max_parameters = Inf, ...)
```

Arguments

<code>x</code>	A fitted 'exnex_surv' object.
<code>show_trace</code>	Should traceplots be shown? Default is 'TRUE'.
<code>parameters</code>	Optional character vector of parameter names for traceplots.
<code>max_parameters</code>	Maximum number of parameters to plot.
<code>...</code>	Unused.

Value

Invisibly returns 'x'.

```
probability_superiority
      Posterior probability that one group beats another
```

Description

Estimates the posterior probability that a certain scalar summary of the survival distribution for group 'a' is greater than the same summary for group 'b'. Supported summaries are the median survival time, the survival probability at a fixed time 'S(t_0)', and the restricted mean survival time (RMST). The function works draw-by-draw on the posterior, so the probability is computed on the joint posterior of the two groups.

Usage

```
probability_superiority(
  fit,
  a,
  b,
  function_of = "median",
  newdata = NULL,
```

```

    times = NULL,
    tmax = NULL,
    ...
)

```

Arguments

<code>fit</code>	A fitted ‘ <code>exnex_surv</code> ’ object.
<code>a</code>	An integer index (or group label) of the first group.
<code>b</code>	An integer index (or group label) of the second group.
<code>function_of</code>	Character; one of “median”, “survival” (alias “S_t”), or “rmst”. Selects which summary is compared.
<code>newdata</code>	Optional data frame with one row per group, used to fix covariates. If ‘NULL’, covariates are set to zero.
<code>times</code>	A scalar time at which to evaluate the survival function when ‘ <code>function_of = "survival"</code> ’.
<code>tmax</code>	A positive horizon when ‘ <code>function_of = "rmst"</code> ’.
<code>...</code>	Unused.

Value

A named list with elements ‘`prob`’ (posterior probability that the summary of group ‘`a`’ exceeds that of group ‘`b`’), ‘`summary`’ (which summary was used), and ‘`groups`’ (the two groups compared).

<code>rmst</code>	<i>Restricted mean survival time (RMST) from an <code>exnex_surv</code> fit</i>
-------------------	---

Description

Computes the restricted mean survival time up to a horizon ‘`tmax`’:

$$RMST(tmax) = \int_0^{tmax} S(t) dt.$$

The integral is evaluated numerically (trapezoidal rule over a fine grid) for every posterior draw, then summarised with posterior quantiles.

Usage

```
rmst(fit, tmax = NULL, newdata = NULL, level = 0.95, grid_points = 400, ...)
```

Arguments

<code>fit</code>	A fitted 'exnex_surv' object.
<code>tmax</code>	A positive scalar horizon up to which the RMST is computed. Defaults to 'max(fit\$data\$time)'.
<code>newdata</code>	Optional data frame (one row gives one RMST).
<code>level</code>	Credible-interval level (default '0.95').
<code>grid_points</code>	Resolution of the numerical quadrature (default '400').
<code>...</code>	Unused.

Value

A 'data.frame' with columns 'group', 'rmst', 'lower', 'upper'.

<code>simulate_data</code>	<i>Simulate basket-trial log-normal survival data</i>
----------------------------	---

Description

Generates a synthetic dataset from the log-normal AFT model used by 'exnexSurv':

$$\log T_i = \theta_{g[i]} + X_i^\top \beta + \varepsilon, \quad \varepsilon \sim \mathcal{N}(0, \sigma^2).$$

Healthy baskets draw their location from a common baseline plus random noise; "outlier" baskets (specified via 'outlier_baskets') have their location shifted by 'resist_delta'. The function is useful for simulation studies and teaching examples.

Usage

```
simulate_data(
  n = 30,
  K = 9,
  beta = c(0.5, -0.2),
  sigma = 1.2,
  outlier_baskets = NULL,
  resist_delta = -0.8,
  censoring_rate = NULL,
  censor_upper = NULL,
  theta = 0,
  seed = NULL
)
```

Arguments

n	Number of patients per basket. A scalar is replicated across all 'K' baskets; a numeric vector of length 'K' assigns a size to each basket individually.
K	Number of baskets (default '9').
beta	Numeric vector of regression coefficients for the covariates. Length determines the number of covariates.
sigma	Residual standard deviation (default '1.2').
outlier_baskets	Optional integer vector of basket indices (1 to K) whose true location is shifted away from the healthy population. These are the baskets the EXNEX model is designed to detect.
resist_delta	Additive shift applied to the 'theta' of outlier baskets (default '-0.8').
censoring_rate	Approximate proportion of censoring after 'censor_upper'; if 'NULL', no censoring is applied.
censor_upper	Upper bound of the censoring-time uniform distribution.
theta	Baseline location for the healthy (non-outlier) baskets; a scalar used as the centre around which the healthy basket locations vary.
seed	Optional seed for reproducibility.

Value

A 'data.frame' with columns 'time', 'event', 'group' (a factor), and one covariate column ('x1', ...) per entry in 'beta'. True parameter values are stored as attributes 'true_theta', 'true_beta', and 'true_sigma'.

summary.exnex_surv	<i>Summarize posterior draws from an exnex_surv fit</i>
--------------------	---

Description

Summarize posterior draws from an exnex_surv fit

Usage

```
## S3 method for class 'exnex_surv'
summary(object, probs = c(0.05, 0.5, 0.95), ...)
```

Arguments

object	A fitted 'exnex_surv' object.
probs	Quantiles to report.
...	Unused.

Value

A data frame with posterior summaries.

survival_curves	<i>Survival curves from an <code>exnex_surv</code> fit</i>
-----------------	--

Description

Computes posterior survival curves $S(t) = \Pr(T > t)$ from a fitted ‘`exnex_surv`’ model for a log-normal AFT specification:

$$\log T_i = \theta_{g[i]} + X_i^\top \beta + \varepsilon, \quad \varepsilon \sim \mathcal{N}(0, \sigma^2).$$

For a fixed time t and linear predictor η , the survival probability of the log-normal distribution is

$$S(t) = 1 - \Phi\left(\frac{\log t - \eta}{\sigma}\right),$$

where σ^2 is the residual variance. Posterior draws of $(\theta, \beta, \sigma^2)$ are propagated through this expression to obtain a full posterior distribution of $S(t)$ at each time point.

Usage

```
survival_curves(fit, newdata = NULL, times = NULL, level = 0.95, ...)
```

Arguments

<code>fit</code>	A fitted ‘ <code>exnex_surv</code> ’ object.
<code>newdata</code>	Optional data frame with columns matching the covariates of the model. If it contains a ‘ <code>group</code> ’ column, that is used for the group index; otherwise the first group is used for all rows.
<code>times</code>	Optional numeric vector of times at which to evaluate the curves. If ‘ <code>NULL</code> ’, a sensible grid is built from the observed follow-up times.
<code>level</code>	Credible-interval level (default ‘0.95’).
<code>...</code>	Unused.

Details

If ‘`newdata`’ is supplied, each row is evaluated at its own group and covariate values; otherwise the covariates are fixed at zero and the first group is used (with a warning if more than one group exists).

Value

An object of class ‘`survival_exnex`’ (also a ‘`data.frame`’) with long format columns: ‘`time`’, ‘`median`’, ‘`lower`’, ‘`upper`’, and ‘`group`’.

Index

`compare_waic`, [2](#)
`compute_waic`, [2](#)

`exnex_surv`, [3](#)

`median_survival`, [6](#)

`plot.exnex_surv`, [7](#)
`plot.survival_exnex`, [7](#)
`print.exnex_surv`, [8](#)
`probability_superiority`, [8](#)

`rmst`, [9](#)

`simulate_data`, [10](#)
`summary.exnex_surv`, [11](#)
`survival_curves`, [12](#)