Package 'AHTauDesign'

September 10, 2026

Type Package

Title Truncation Time Selection for Average Hazard Analyses

Version 0.1.0

Description Implements a two-stage, design-informed framework for selecting the truncation time in time-to-event trials analyzed with the average hazard estimand, motivated by pediatric oncology settings with small samples, slow accrual, and limited follow-up. Stage 1 validates a clinically proposed truncation time against the planned design using simulation-based diagnostics for risk-set support, follow-up coverage, estimability, and estimator stability, classifying it as Pass, Borderline, or Fail. Stage 2 performs constrained optimization over a grid of candidate truncation times within a clinical-distance window, maximizing a utility subject to feasibility constraints, with an independent evaluation run to assess the selected time. Supports proportional-hazards, early-, and delayed-effect patterns, uniform accrual with administrative censoring, and calibrated exponential random censoring.

License GPL-3

Encoding UTF-8

Depends R (>= 4.1.0)

Imports survival, stats, utils

Suggests ggplot2, knitr, rmarkdown

VignetteBuilder knitr

Config/roxygen2/version 8.0.0

NeedsCompilation no

Author Xinyue Zhao [aut, cre],
Haitao Pan [aut]

Maintainer Xinyue Zhao <xinyuezhao27@gwu.edu>

Repository CRAN

Date/Publication 2026-09-10 08:40:08 UTC

Contents

`assess_tau` *Assess the statistical feasibility of a proposed truncation time*

Description

Evaluates whether a single proposed truncation time τ is statistically supported under the planned two-arm trial design, and classifies it as `"Pass"`, `"Borderline"`, or `"Fail"`.

Usage

```
assess_tau(
  N,
  A,
  F_followup,
  tau,
  effect_model,
  alloc_ratio = 1,
  censor_params = list(add_random = FALSE),
  n_sim = 100,
  n_boot = 100,
  min_boot_valid_frac = 0.8,
  alpha = 0.05,
  feasibility_thresholds = list(),
  fail_thresholds = list()
)

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

Arguments

`N`	Total planned sample size (both arms combined).
`A`	Numeric. Accrual duration, in the time unit of the design (e.g., months). Enrollment times are uniform on $[0, A]$.
`F_followup`	Numeric. Additional follow-up duration after the end of accrual, in the same time unit as `A`.

`tau` Numeric scalar. The single proposed truncation time to assess, on the original time scale. Supplying a vector is an error; use `optimize_tau()` for a grid of candidates.

`effect_model` A named list describing the treatment-effect assumption. Must contain `type`, one of `"ph"`, `"early"`, or `"delayed"`, and `lambda0`, the control-arm exponential event rate. For `type = "ph"`, supply the hazard ratio `HR`. For `type = "early"` or `"delayed"`, supply `HR1`, `HR2`, and the change point `pt`.

`alloc_ratio` Numeric. Allocation ratio of treatment to control (default `1`, i.e., 1:1 allocation).

`censor_params` A named list controlling random censoring in addition to administrative censoring. Default `list(add_random = FALSE)` uses administrative censoring only. Set `add_random = TRUE` and supply `rate_exp`, an exponential random-censoring rate, to add random censoring.

`n_sim` Integer. Number of simulated trial datasets over which the feasibility diagnostics are averaged (default `100`).

`n_boot` Integer. Number of bootstrap resamples drawn within each simulated trial to estimate the standard error of the log average-hazard-ratio estimate (default `100`).

`min_boot_valid_frac`
Numeric in $[0, 1]$. Minimum fraction of valid bootstrap resamples required for the bootstrap standard error to be reported (default `0.80`).

`alpha` Numeric in $(0, 1)$. One-sided significance level of the test performed within each simulated trial (default `0.05`).

`feasibility_thresholds`
A named list overriding any of the default Pass thresholds; see `default_feasibility_thresholds`. Unspecified elements keep their defaults.

`fail_thresholds`
A named list overriding any of the default hard-failure rules; see `default_fail_thresholds`. Unspecified elements keep their defaults.

`x` An object of class `"assess_tau"`.

`...` Further arguments; currently unused.

Details

Feasibility is judged by Monte Carlo simulation of the planned trial. For each simulated dataset the function computes the risk-set sizes at τ, the risk-set proportion $P(X \geq \tau)$ — the proportion of participants remaining under observation and event-free at τ, not administrative follow-up coverage alone — and whether the log average-hazard ratio is estimable. Averaging over replicates gives the feasibility diagnostics, which are compared against fixed decision rules:

- **Pass**: expected total risk-set size ≥ 30, expected risk-set size ≥ 10 in each arm, risk-set proportion $P(X \geq \tau) \geq 0.10$, and estimability probability ≥ 0.90.
- **Fail**: expected total risk-set size < 20, or an arm-specific risk-set size < 5, or risk-set proportion $P(X \geq \tau) < 0.10$, or estimability probability < 0.80.
- **Borderline**: neither all Pass conditions nor any Fail condition is met.

assess_tau() evaluates exactly one truncation time per call. To screen a grid of candidate values and optimize among them, use optimize_tau().

The simulation settings n_sim, n_boot, min_boot_valid_frac, and alpha have developer-validated defaults and do not need to be changed in ordinary use. For reproducibility, call set.seed() once before assess_tau(); the function does not modify the random-number-generator state internally.

Value

An object of class "assess_tau": a named list with components

feasibility Character. "Pass", "Borderline", or "Fail".

Expected total risk-set size at tau Rounded to 1 decimal place.

Expected control-arm risk-set size at tau Rounded to 1 decimal place.

Expected treatment-arm risk-set size at tau Rounded to 1 decimal place.

Risk-set proportion at tau Rounded to 4 decimal places.

Estimability probability Rounded to 4 decimal places.

summary A data.frame with one row per feasibility metric (value, Pass threshold, Fail threshold) and the feasibility classification as the last row.

tau The assessed truncation time.

call_params A named list recording the design inputs, for provenance.

Components whose names contain spaces are accessed with backticks, e.g. result$`Risk-set proportion at tau`.

Examples

```
effect_model <- list(type = "ph", lambda0 = -log(0.5)/24, HR = 0.5)

# Settings reduced so this example runs quickly. Use the defaults
# (n_sim = 100, n_boot = 100) or larger for real design work.
set.seed(2026)
assessment <- assess_tau(
  N = 60, A = 24, F_followup = 6,
  tau = 24,
  effect_model = effect_model,
  n_sim = 10, n_boot = 10
)

assessment$feasibility
assessment
```

default_fail_thresholds

Default Fail thresholds for feasibility classification

Description

Returns the built-in hard-failure rules. A candidate is classified as `"Fail"` if any one of them is violated.

Usage

```
default_fail_thresholds()
```

Value

A named list with `total_risk`, `arm_risk`, `riskset_prop`, and `estimability`.

Examples

```
default_fail_thresholds()
```

default_feasibility_thresholds

Default Pass thresholds for feasibility classification

Description

Returns the built-in thresholds a candidate must meet to be classified as `"Pass"`.

Usage

```
default_feasibility_thresholds()
```

Value

A named list with `min_total_risk`, `min_arm_risk`, `min_riskset_prop`, and `min_estimability`.

Examples

```
default_feasibility_thresholds()
```

`optimize_tau` *Feasibility-constrained utility optimization of the truncation time*

Description

Identifies a statistically feasible and clinically acceptable truncation time from a candidate grid. Each candidate is classified as `"Pass"`, `"Borderline"`, or `"Fail"`; the admissible set is the intersection of the Pass set with the clinical window $|\tau - \tau_{clin}| \leq \Delta$; the utility $\hat{U}(\tau)$ is maximized over the admissible set; and a clinical-retention rule decides whether the clinically proposed value is kept.

Usage

```
optimize_tau(
  N,
  A,
  F_followup,
  tau_clin,
  tau_grid,
  effect_model,
  delta = 12,
  epsilon = 0.05,
  alloc_ratio = 1,
  censor_params = list(add_random = FALSE),
  n_sim = 100,
  n_boot = 100,
  min_boot_valid_frac = 0.8,
  alpha = 0.05
)

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

Arguments

`N`	Total planned sample size (both arms combined).
`A`	Numeric. Accrual duration. Enrollment times are uniform on $[0, A]$.
`F_followup`	Numeric. Additional follow-up duration after the end of accrual.
`tau_clin`	Numeric scalar. The clinically proposed truncation time (the anchor value). Must be contained in `tau_grid`.
`tau_grid`	Numeric vector. The candidate truncation times to evaluate (the search space).
`effect_model`	A named list describing the treatment-effect assumption; see `assess_tau`.
`delta`	Numeric. Half-width Δ of the clinical window around `tau_clin` (default `12`).
`epsilon`	Numeric. Relative utility tolerance ε of the clinical-retention rule (default `0.05`).
`alloc_ratio`	Numeric. Allocation ratio of treatment to control (default `1`).

censor_params	A named list controlling random censoring; see assess_tau.
n_sim	Integer. Number of simulated trial datasets; every candidate in tau_grid is evaluated on the same simulated trials (default 100).
n_boot	Integer. Number of bootstrap resamples drawn within each simulated trial to estimate the standard error of the log average-hazard-ratio estimate (default 100).
min_boot_valid_frac	Numeric in $[0, 1]$. Minimum fraction of valid bootstrap resamples required for the bootstrap standard error to be reported (default 0.80).
alpha	Numeric in $(0, 1)$. One-sided significance level of the test performed within each simulated trial (default 0.05).
x	An object of class "optimize_tau".
...	Further arguments; currently unused.

Details

The two truncation-time arguments play distinct roles: tau_clin is the clinically proposed anchor value, and tau_grid is the search space. tau_clin must be an element of tau_grid.

The procedure: (1) evaluate all candidates on the same simulated trials; (2) classify each candidate; (3) build the admissible set; (4) select the preliminary utility-optimal candidate $\tilde{\tau}$; (5) retain τ_{clin} if it is admissible and its utility is within a relative tolerance epsilon of the best admissible utility, otherwise select $\tilde{\tau}$.

The simulation settings n_sim, n_boot, min_boot_valid_frac, and alpha have developer-validated defaults and do not need to be changed in ordinary use. For reproducibility, call set.seed() once before optimize_tau(); the function does not modify the random-number-generator state internally.

Value

An object of class "optimize_tau": a named list with components

candidate_summary Full tau-by-tau table of feasibility, utility, power, and diagnostics.

T_pass Feasible (Pass) candidates.

T_clin_window Candidates inside the clinical window.

T_adm The admissible set.

tau_tilde Preliminary utility-optimal candidate.

tau_star Final selection (NA if no selection was possible).

tau_clin_retained Logical. TRUE if tau_clin was retained; FALSE only when tau_clin was an admissible Pass candidate but was not retained because another candidate had sufficiently higher utility; NA when the clinical-retention rule was not applicable (e.g., tau_clin was classified as Fail or Borderline, or no selection was possible).

stage1_status_tau_clin Feasibility status of tau_clin.

decision_type Machine-readable decision code, one of "retained", "utility_driven", "constraint_driven", "no_admissible_candidate", or "no_positive_direction_candidate".

decision_message A short explanation.

U_max, U_tau_clin Best admissible utility and the utility at tau_clin.

call_params Full provenance of the call.

Examples

```
effect_model <- list(type = "ph", lambda0 = -log(0.5)/24, HR = 0.5)

# Settings reduced so this example runs quickly. Use the defaults
# (n_sim = 100, n_boot = 100) and a finer tau_grid for real design work.
set.seed(2026)
optimization <- optimize_tau(
  N = 60, A = 24, F_followup = 6,
  tau_clin = 24,
  tau_grid = seq(6, 36, by = 6),
  effect_model = effect_model,
  delta = 12, epsilon = 0.05,
  n_sim = 10, n_boot = 10
)

optimization$tau_star
optimization
```

`twostage_tau` *Complete two-stage truncation-time selection*

Description

The primary function for users who want a final truncation-time recommendation. It runs the Stage 1 feasibility assessment of the clinically proposed value τ_{clin}, the feasibility-constrained utility optimization of Stage 2, fixes the selection τ^*, and then compares τ_{clin} and τ^* in an independent evaluation run.

Usage

```
twostage_tau(
  N,
  A,
  F_followup,
  tau_clin,
  tau_grid,
  effect_model,
  delta = 12,
  epsilon = 0.05,
  alloc_ratio = 1,
  censor_params = list(add_random = FALSE),
  n_sim = 100,
  n_boot = 100,
  min_boot_valid_frac = 0.8,
  alpha = 0.05
)
```

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

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

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

Arguments

N	Total planned sample size (both arms combined).
A	Numeric. Accrual duration. Enrollment times are uniform on $[0, A]$.
F_followup	Numeric. Additional follow-up duration after the end of accrual.
tau_clin	Numeric scalar. The clinically proposed truncation time (the anchor value). Must be contained in tau_grid.
tau_grid	Numeric vector. The candidate truncation times to evaluate (the search space).
effect_model	A named list describing the treatment-effect assumption; see assess_tau.
delta	Numeric. Half-width Δ of the clinical window around tau_clin (default 12).
epsilon	Numeric. Relative utility tolerance ε of the clinical-retention rule (default 0.05).
alloc_ratio	Numeric. Allocation ratio of treatment to control (default 1).
censor_params	A named list controlling random censoring; see assess_tau.
n_sim	Integer. Number of simulated trial datasets; every candidate in tau_grid is evaluated on the same simulated trials (default 100).
n_boot	Integer. Number of bootstrap resamples drawn within each simulated trial to estimate the standard error of the log average-hazard-ratio estimate (default 100).
min_boot_valid_frac	Numeric in $[0, 1]$. Minimum fraction of valid bootstrap resamples required for the bootstrap standard error to be reported (default 0.80).
alpha	Numeric in $(0, 1)$. One-sided significance level of the test performed within each simulated trial (default 0.05).
x	An object of class "twostage_tau".
...	Further arguments; currently unused.
object	An object of class "twostage_tau".

Details

The selection stage and the independent evaluation stage consume successive draws from a single random-number stream, so a single set.seed() call before twostage_tau() makes the entire run reproducible, with the two stages independent by construction. The function itself takes no seed arguments and does not modify the random-number-generator state internally.

The simulation settings n_sim, n_boot, min_boot_valid_frac, and alpha have developer-validated defaults and do not need to be changed in ordinary use; n_sim and n_boot apply to both the selection run and the independent evaluation run.

Three S3 methods layer the presentation: `print()` shows a short summary block (Stage 1 status, τ^*, and the decision message); `summary()` adds the candidate-level and evaluation tables; `plot()` draws the utility curve $\hat{U}(\tau)$ over the candidate grid with the feasibility classification, the clinical window, τ_{clin}, and the selected τ^* marked (requires the ggplot2 package).

Value

An object of class `"twostage_tau"`: a named list whose first three components are the reporting tables

`candidate_summary` Candidate-level selection table: one row per candidate tau with rounded diagnostics, effect summaries, feasibility, distance to `tau_clin`, within-window, utility, admissible, and selected flags.

`final_decision_summary` Two-column (quantity, result) table with the Stage 1 status, candidate sets, retention decision, final `tau_star`, decision type, and efficiency gains.

`eval_summary` Independent evaluation table comparing `tau_clin` and `tau_star` on fresh simulations (`NULL` when no selection was possible).

followed by the individual selection components (`T_pass`, `T_adm`, `tau_tilde`, `tau_star`, `tau_clin_retained`, `stage1_status_tau_clin`, `decision_type`, `decision_message`, `abs_power_gain`, `rel_utility_improvement`), `selection_result` (the full, unrounded `optimize_tau` object), and `call_params`.

Examples

```
effect_model <- list(type = "ph", lambda0 = -log(0.5)/24, HR = 0.5)

# Settings reduced so this example runs quickly. Use the defaults
# (n_sim = 100, n_boot = 100) and a finer tau_grid for real design work.
set.seed(2026)
result <- twostage_tau(
  N = 60, A = 24, F_followup = 6,
  tau_clin = 24,
  tau_grid = seq(6, 36, by = 6),
  effect_model = effect_model,
  delta = 12, epsilon = 0.05,
  n_sim = 10, n_boot = 10
)

result
summary(result)
plot(result)
```

Index