Package 'simFastBOIN'

December 17, 2025

Type Package

Title Fast Bayesian Optimal Interval Design for Phase I Dose-Finding Trials

Version 1.3.2

Date 2025-12-06

Description Conducting Bayesian Optimal Interval (BOIN) design for phase I dose-finding trials. 'simFastBOIN' provides functions for pre-computing decision tables, conducting trial simulations, and evaluating operating characteristics. The package uses vectorized operations and the Iso::pava() function for isotonic regression to achieve efficient performance while maintaining full compatibility with BOIN methodology. Version 1.3.2 adds p_saf and p_tox parameters for customizable safety and toxicity thresholds. Version 1.3.1 fixes Date field. Version 1.2.1 adds comprehensive 'roxygen2' documentation and enhanced print formatting with flexible table output options. Version 1.2.0 integrated C-based PAVA for isotonic regression. Version 1.1.0 introduced conservative MTD selection (boundMTD) and flexible early stopping rules (n_earlystop_rule). Methods are described in Liu and Yuan (2015) <doi:10.1111/rssc.12089>.

License MIT + file LICENSE

Imports knitr, kableExtra, Iso, stats

Suggests testthat (>= 3.0.0), rmarkdown

Encoding UTF-8

RoxygenNote 7.3.3

Language en-US

URL https://github.com/gosukehommaEX/simFastBOIN

BugReports https://github.com/gosukehommaEX/simFastBOIN/issues

Config/testthat/edition 3

VignetteBuilder knitr

NeedsCompilation no

Author Gosuke Homma [aut, cre] (ORCID: <https://orcid.org/0000-0002-6854-5627>)

Maintainer Gosuke Homma <my.name.is.gosuke@gmail.com>

Repository CRAN

Date/Publication 2025-12-17 10:20:02 UTC

Contents

`get_boin_boundary` *Calculate BOIN Escalation and De-escalation Boundaries*

Description

Compute the escalation and de-escalation boundaries for the Bayesian Optimal Interval (BOIN) design based on the target toxicity probability. These boundaries are used to determine dose escalation and de-escalation decisions during the trial.

Usage

```
get_boin_boundary(target, p_saf = NULL, p_tox = NULL)
```

Arguments

`target`	Numeric. The target toxicity probability (e.g., 0.30 for 30%). Should be between 0 and 1.
`p_saf`	Numeric. The toxicity probability threshold for safe (subtherapeutic) dose. Default is 0.6 * target. Doses with observed toxicity rate <= p_saf trigger escalation.
`p_tox`	Numeric. The toxicity probability threshold for overly toxic dose. Default is 1.4 * target. Doses with observed toxicity rate >= p_tox trigger de-escalation.

Details

Based on Bayesian theory with three-point hypothesis, these boundaries are computed to optimally balance exploration and exploitation in dose finding. The default values (p_saf = 0.6 * target, p_tox = 1.4 * target) are recommended and generally yield excellent operating characteristics.

Value

A list containing:

lambda_e	Numeric. Escalation boundary: toxicity rate threshold for escalation
lambda_d	Numeric. De-escalation boundary: toxicity rate threshold for de-escalation

References

Liu S. and Yuan, Y. (2015). Bayesian Optimal Interval Designs for Phase I Clinical Trials. Journal of the Royal Statistical Society: Series C, 64, 507-523.

Examples

```
# Calculate boundaries for 30% target toxicity rate
boin_bound <- get_boin_boundary(target = 0.30)
print(boin_bound)
# $lambda_e
# [1] 0.236
# $lambda_d
# [1] 0.359

# Calculate boundaries for 25% target with custom thresholds
boin_bound_custom <- get_boin_boundary(
  target = 0.25,
  p_saf = 0.12,
  p_tox = 0.40
)
print(boin_bound_custom)

# Compare boundaries for different target rates
targets <- c(0.20, 0.25, 0.30, 0.35)
boundaries <- lapply(targets, get_boin_boundary)
names(boundaries) <- paste0("Target_", targets * 100, "%")
print(boundaries)
```

---

get_boin_decision *Generate Dose Decision Table for BOIN Design*

Description

Create a pre-computed lookup table that maps (number of DLTs, number of patients) pairs to dose decisions (Escalate, De-escalate, Stay, or Eliminate). This table is generated once before the trial and consulted repeatedly during dose assignment.

Usage

```
get_boin_decision(target, lambda_e, lambda_d, max_sample_size, cutoff_eli)
```

Arguments

target	Numeric. The target toxicity probability (e.g., 0.30 for 30%).
lambda_e	Numeric. Escalation boundary from get_boin_boundary(). Doses with observed toxicity rate <= lambda_e trigger escalation.
lambda_d	Numeric. De-escalation boundary from get_boin_boundary(). Doses with observed toxicity rate >= lambda_d trigger de-escalation.
max_sample_size	Numeric. Maximum sample size (number of patients) for table columns. Typically 18-30 for phase I trials.
cutoff_eli	Numeric. Cutoff probability for dose elimination. Default is 0.95. If Pr(p > target | data) > cutoff_eli, dose is marked for elimination.

Details

Decision rules applied in order:

1. Rule DE: If n >= 3 and Pr(p > target | data) > cutoff_eli, eliminate dose
2. Rule E: If observed toxicity rate <= lambda_e, escalate
3. Rule D: If observed toxicity rate >= lambda_d, de-escalate
4. Rule S: Otherwise, stay at current dose

The posterior probability Pr(p > target | data) is computed using Beta-Binomial conjugate prior with uniform prior (Beta(1,1)).

Value

A character matrix with dose decisions:

- Rows represent cumulative number of DLTs (0 to max_sample_size)
- Columns represent cumulative number of patients (1 to max_sample_size)
- Cell values are decisions: "E" (Escalate), "D" (De-escalate), "DE" (De-escalate & Eliminate), "S" (Stay), or NA (logically impossible)

References

Liu S. and Yuan, Y. (2015). Bayesian Optimal Interval Designs for Phase I Clinical Trials. Journal of the Royal Statistical Society: Series C, 64, 507-523.

Examples

```
# Generate decision table for 30% target toxicity rate
boin_bound <- get_boin_boundary(target = 0.30)
DECISION <- get_boin_decision(
  target = 0.30,
  lambda_e = boin_bound$lambda_e,
  lambda_d = boin_bound$lambda_d,
  max_sample_size = 18,
  cutoff_eli = 0.95
```

```
)
print(DECISION)
```

`get_boin_stopping_boundaries` *Generate Trial Stopping Rule Table*

Description

Create a lookup table to determine whether the entire trial should be stopped if excessive toxicity is observed at the lowest dose. This table is applied only at the lowest dose level to monitor trial safety.

Usage

```
get_boin_stopping_boundaries(target, max_sample_size, cutoff_stop)
```

Arguments

`target`	Numeric. The target toxicity probability (e.g., 0.30 for 30%).
`max_sample_size`	Numeric. Maximum sample size (number of patients) for table columns. Typically 18-30 for phase I trials.
`cutoff_stop`	Numeric. Cutoff probability for trial stopping. Default is 0.90. If Pr(p > target | data) > cutoff_stop at lowest dose, the trial is stopped.

Details

The trial stopping rule is based on Bayesian monitoring of safety at the lowest dose level. The rationale is that if even the lowest dose shows excessive toxicity with high posterior probability, all doses in the trial may be too toxic, warranting early trial termination.

The posterior probability Pr(p > target | data) is computed using Beta-Binomial conjugate prior with uniform prior (Beta(1,1)). The stopping rule is not evaluated until at least 3 patients have been treated at the lowest dose.

Value

A character matrix with stopping decisions:

- Rows represent cumulative number of DLTs at lowest dose (0 to max_sample_size)
- Columns represent cumulative number of patients at lowest dose (1 to max_sample_size)
- Cell values are "STOP" or "GO"
- NA values for logically impossible combinations (n_tox > n_pts) or n_pts < 3

References

Liu S. and Yuan, Y. (2015). Bayesian Optimal Interval Designs for Phase I Clinical Trials. Journal of the Royal Statistical Society: Series C, 64, 507-523.

Examples

```
# Generate stopping rule table for 30% target toxicity rate
STOP_DL1 <- get_boin_stopping_boundaries(
  target = 0.30,
  max_sample_size = 18,
  cutoff_stop = 0.90
)
print(STOP_DL1)
```

---

get_pts_and_tox *Generate Patient and Toxicity Data for BOIN Simulations*

Description

Conduct vectorized BOIN trial simulations with optional titration phase to generate patient enrollment and toxicity (DLT) data across multiple trials.

Usage

```
get_pts_and_tox(
  n_trials = 10000,
  target,
  p_true,
  n_cohort,
  cohort_size,
  p_saf = NULL,
  p_tox = NULL,
  n_earlystop = 18,
  cutoff_eli = 0.95,
  extrasafe = FALSE,
  offset = 0.05,
  n_earlystop_rule = c("with_stay", "simple"),
  titration = FALSE,
  seed = 123
)
```

Arguments

n_trials	Numeric. Number of trials to simulate. Default is 10000.
target	Numeric. The target toxicity probability (e.g., 0.30 for 30%).
p_true	Numeric vector. True toxicity probabilities for each dose.

n_cohort	Numeric. Maximum number of cohorts per trial.
cohort_size	Numeric vector or scalar specifying patients per cohort.
p_saf	Numeric. Highest toxicity probability deemed acceptable for safety. Default is 0.6 * target. Used with p_tox for safety/efficacy dose identification.
p_tox	Numeric. Lowest toxicity probability deemed unacceptable for toxicity. Default is 1.4 * target. Used with p_saf for safety/efficacy dose identification.
n_earlystop	Numeric. Sample size at current dose triggering trial termination. Default is 18.
cutoff_eli	Numeric. Cutoff probability for dose elimination. Default is 0.95.
extrasafe	Logical. If TRUE, apply additional safety stopping rule at the lowest dose. Default is FALSE.
offset	Numeric. Offset for safety stopping boundary when extrasafe = TRUE. Default is 0.05.
n_earlystop_rule	Character. Rule for early stopping at n_earlystop. Options are "with_stay" or "simple". Default is "with_stay".
titration	Logical. If TRUE, perform accelerated dose escalation in titration phase. Default is FALSE.
seed	Numeric. Random seed for reproducibility. Default is 123.

Details

The function simulates multiple BOIN trials in parallel using vectorized operations. When titration = TRUE, an accelerated dose escalation phase is performed using single-patient cohorts to rapidly locate the dose-toxicity region. The main trial then proceeds with standard cohort-based dosing.

Early stopping rules include:

- Simple rule: Stop when n_earlystop patients are treated at current dose
- With-stay rule: Stop when n_earlystop patients are treated and dose stays
- Extra-safety rule (if extrasafe = TRUE): Stop if lowest dose is excessively toxic

Value

A list containing:

n_pts_all	Matrix of size n_trials x n_doses: patient counts at each dose
n_tox_all	Matrix of size n_trials x n_doses: DLT counts at each dose
eliminated_mat	Matrix of size n_trials x n_doses: logical, TRUE if dose eliminated
cohorts_completed	Integer vector of length n_trials: cohorts completed per trial
stop_reason	Character vector of length n_trials: reason for stopping

References

Liu S. and Yuan, Y. (2015). Bayesian Optimal Interval Designs for Phase I Clinical Trials. Journal of the Royal Statistical Society: Series C, 64, 507-523.

Examples

```
# Basic BOIN simulation
target <- 0.30
p_true <- c(0.10, 0.25, 0.40, 0.55, 0.70)

result <- get_pts_and_tox(
  n_trials = 1000,
  target = target,
  p_true = p_true,
  n_cohort = 10,
  cohort_size = 3,
  n_earlystop = 18,
  seed = 123
)

# With titration phase
result_titration <- get_pts_and_tox(
  n_trials = 1000,
  target = target,
  p_true = p_true,
  n_cohort = 10,
  cohort_size = 3,
  titration = TRUE,
  seed = 123
)

# With extra safety stopping
result_safe <- get_pts_and_tox(
  n_trials = 1000,
  target = target,
  p_true = p_true,
  n_cohort = 10,
  cohort_size = 3,
  extrasafe = TRUE,
  offset = 0.05,
  seed = 123
)
```

---

isotonic_regression *Isotonic Regression for Toxicity Rate Estimation*

Description

Apply isotonic regression to estimate toxicity rates under the monotonicity constraint (toxicity increases with dose) using the Pool Adjacent Violators Algorithm (PAVA).

Usage

```
isotonic_regression(n_pts_mat, n_tox_mat)
```

Arguments

n_pts_mat	Numeric matrix of size n_trials x n_doses. Number of patients treated at each dose level for each trial.
n_tox_mat	Numeric matrix of size n_trials x n_doses. Number of dose-limiting toxicities (DLTs) observed at each dose level for each trial.

Details

Pseudocounts (0.05 for toxicity, 0.1 for total patients) are added before estimation, reflecting a Beta-Binomial prior. PAVA is applied independently per trial using weighted pooling to enforce monotonicity: dose-toxicity relationships must be non-decreasing across dose levels.

Value

Numeric matrix of size n_trials x n_doses. Isotonic regression estimates of toxicity rates at each dose for each trial. Values are NA for untreated doses (n_pts = 0).

References

Liu, S. and Yuan, Y. (2015). Bayesian Optimal Interval Designs for Phase I Clinical Trials. *Journal of the Royal Statistical Society: Series C*, 64, 507-523.

Examples

```
# Example 1: Single trial
n_pts_single <- matrix(c(3, 6, 9, 12), nrow = 1)
n_tox_single <- matrix(c(0, 1, 3, 4), nrow = 1)

iso_est <- isotonic_regression(n_pts_single, n_tox_single)
print(iso_est)

# Example 2: Multiple trials
n_pts_mat <- matrix(c(3, 6, 9, 12,
                      3, 6, 9, 12,
                      3, 6, 9, 12), nrow = 3, byrow = TRUE)
n_tox_mat <- matrix(c(0, 1, 3, 4,
                      0, 0, 2, 3,
                      1, 2, 4, 6), nrow = 3, byrow = TRUE)

iso_est <- isotonic_regression(n_pts_mat, n_tox_mat)
print(iso_est)
```

print.boin_multi_summary *Print Method for BOIN Multi-Scenario Summary*

Description

Custom print method for "boin_multi_summary" objects returned by `sim_boin_multi()`. Displays aggregated operating characteristics in formatted tables with options for plain text or kable output.

Usage

```
## S3 method for class 'boin_multi_summary'
print(x, scenario_name = NULL, percent = FALSE, kable = FALSE, kable_format = "pipe", ...)
```

Arguments

`x`	Object of class "boin_multi_summary" returned by `sim_boin_multi()`.
`scenario_name`	Character. Optional name for display header (for compatibility with print.boin_summary). Ignored for multi-scenario results. Default is NULL.
`percent`	Logical. If TRUE, display "Avg Pts" and "Avg DLTs" as percentages of totals. If FALSE (default), display as absolute numbers. Default is FALSE.
`kable`	Logical. If TRUE, format output as knitr::kable table. If FALSE (default), display as plain text table. Default is FALSE.
`kable_format`	Character. Format for kable output when kable = TRUE. Options include "pipe" (Markdown pipes, default), "simple" (minimal formatting), "latex" (LaTeX format), and "html" (HTML format). Default is "pipe".
`...`	Additional arguments (currently unused, for S3 method consistency).

Details

Displays a summary table with:

- Rows organized by scenario and metric (True Tox (%), MTD Sel (%), Avg Pts, Avg DLTs)
- Columns for each dose level plus Total/No MTD statistics
- Scenario names grouped for easy visual comparison
- All text left-aligned for consistency

When `percent = FALSE` (default), "Avg Pts" and "Avg DLTs" are displayed as absolute numbers.

When `percent = TRUE`, these rows display percentages of their respective totals for each trial. This facilitates visual comparison of dose allocation patterns across scenarios independent of total sample size.

When `kable = FALSE` (default), displays as plain text table in console.

When `kable = TRUE`, formats the table using `knitr::kable()` with the specified format. This is useful for R Markdown documents and reports. The `kable_format` parameter controls output format:

- "pipe": Markdown pipe format (best for GitHub/Markdown)
- "simple": Minimal formatting (plain text)
- "latex": LaTeX format (for PDF reports)
- "html": HTML format (for web display)

All columns are left-aligned for improved readability across all formats.

Value

Invisibly returns the input object x. Function is called for printing side effect.

References

Liu S. and Yuan, Y. (2015). Bayesian Optimal Interval Designs for Phase I Clinical Trials. Journal of the Royal Statistical Society: Series C, 64, 507-523.

Examples

```
## Not run:
# Create multi-scenario simulation results
scenarios <- list(
  list(name = "Scenario 1",
       p_true = c(0.05, 0.10, 0.15, 0.20, 0.25, 0.30, 0.35, 0.40, 0.45)),
  list(name = "Scenario 2",
       p_true = c(0.01, 0.025, 0.05, 0.075, 0.10, 0.125, 0.15, 0.20, 0.30))
)

result <- sim_boin_multi(
  scenarios = scenarios,
  target = 0.30,
  n_trials = 1000,
  n_cohort = 48,
  cohort_size = 3,
  seed = 123
)

# Print as plain text with absolute numbers (default)
print(result)

# Print with absolute numbers as Markdown table
print(result, kable = TRUE, kable_format = "pipe")

# Print with percentages instead of absolute numbers
print(result, percent = TRUE)

# Print with percentages as Markdown table
print(result, kable = TRUE, kable_format = "pipe", percent = TRUE)

# Print as LaTeX table for PDF documents
print(result, kable = TRUE, kable_format = "latex")
```

```
# Print as HTML table
print(result, kable = TRUE, kable_format = "html")

## End(Not run)
```

`print.boin_summary` *Print Method for BOIN Summary Objects*

Description

Custom print method for "boin_summary" S3 class objects returned by `sim_boin()`. Displays operating characteristics in a formatted table.

Usage

```
## S3 method for class 'boin_summary'
print(x, scenario_name = NULL, percent = FALSE, kable = FALSE, kable_format = "pipe", ...)
```

Arguments

`x`	Object of class "boin_summary" returned by `sim_boin()`.
`scenario_name`	Character. Optional name for the scenario. If provided, displayed as header. Default is NULL.
`percent`	Logical. If TRUE, display average patients and DLTs as percentages of totals. If FALSE (default), display as absolute numbers.
`kable`	Logical. If TRUE, format output as knitr::kable table. If FALSE (default), display as plain text table.
`kable_format`	Character. Format for kable output when kable = TRUE. Options include "pipe" (Markdown pipes, default), "simple" (minimal formatting), "latex" (LaTeX format), and "html" (HTML format). Default is "pipe".
`...`	Additional arguments (currently unused, for S3 method consistency).

Details

Displays operating characteristics in a unified table with rows:

1. True Tox (%): True DLT rate at each dose
2. MTD Sel (%): Percentage selecting each dose as MTD (+ No MTD)
3. Avg Pts: Average enrollment per dose (absolute or percentage)
4. Avg DLTs: Average DLT count per dose (absolute or percentage)

When `percent = TRUE`, Avg Pts and Avg DLTs are displayed as percentages of their respective totals for each trial.

When `kable = TRUE`, output is formatted using `knitr::kable()`. This is useful for R Markdown documents and reports. The `kable_format` parameter controls the output format: "pipe" for Markdown pipes (default), "simple" for minimal formatting, "latex" for LaTeX tables, and "html" for HTML tables with enhanced styling.

When `kable_format = "html"`, additional styling is applied including striped rows, hover effects, and responsive formatting via kableExtra.

Value

Invisibly returns the input object `x`. Function is called for printing side effect.

References

Liu S. and Yuan, Y. (2015). Bayesian Optimal Interval Designs for Phase I Clinical Trials. Journal of the Royal Statistical Society: Series C, 64, 507-523.

Examples

```
# Create BOIN simulation results
result <- sim_boin(
  n_trials = 1000,
  target = 0.30,
  p_true = c(0.10, 0.25, 0.40),
  n_cohort = 10,
  cohort_size = 3,
  seed = 123
)

# Print with absolute numbers (default)
print(result$summary, scenario_name = "Absolute Numbers")

# Print with percentages
print(result$summary, scenario_name = "Percentages", percent = TRUE)

# Print as Markdown table
print(result$summary, kable = TRUE, kable_format = "pipe")

# Print as HTML table with enhanced styling
print(result$summary, kable = TRUE, kable_format = "html")
```

`select_mtd` *MTD Selection for Multiple Trials*

Description

Select maximum tolerated dose (MTD) for multiple trials based on isotonic regression estimates and dose elimination status. This function follows the BOIN package's MTD selection algorithm.

Usage

```
select_mtd(iso_est_mat, n_pts_mat, eliminated_mat, target,
             boundMTD = FALSE, lambda_d = NULL, min_mtd_sample = 1)
```

Arguments

`iso_est_mat`	Numeric matrix of size n_trials x n_doses. Isotonic regression estimates of toxicity rates from `isotonic_regression`.
`n_pts_mat`	Numeric matrix of size n_trials x n_doses. Number of patients treated at each dose level for each trial.
`eliminated_mat`	Logical matrix of size n_trials x n_doses. Whether each dose has been eliminated in each trial.
`target`	Numeric. Target toxicity probability.
`boundMTD`	Logical. If TRUE, impose constraint that selected MTD's isotonic estimate must be <= lambda_d. Default is FALSE.
`lambda_d`	Numeric. De-escalation boundary. Required if boundMTD = TRUE.
`min_mtd_sample`	Numeric. Minimum number of patients required for a dose to be considered for MTD selection. Default is 1.

Details

For each trial, the function performs the following steps:

1. Check if lowest dose (dose 1) is eliminated first
2. Identify admissible set: doses with patients AND not eliminated
3. Extract isotonic estimates for admissible doses
4. Select dose with estimate closest to target toxicity rate
5. Apply tiebreaker by adding small dose-index increments
6. If boundMTD = TRUE, ensure selected dose satisfies constraint

The dose with isotonic estimate closest to the target is selected as MTD. Ties are broken by small perturbation (1e-10 * dose_index) preferring lower dose indices when estimates are equally close to target.

If the lowest dose (dose 1) is eliminated, NO MTD is selected regardless of other doses' status. This follows BOIN standard: "stop the trial if the lowest dose is eliminated due to toxicity, and no dose should be selected as the MTD."

Value

Data frame with n_trials rows and three columns:

`trial`	Integer. Trial ID (1 to n_trials)
`mtd`	Integer. Selected MTD dose number, or NA if no valid dose
`reason`	Character. Reason for trial completion or termination

Examples

```
target <- 0.30
n_pts_mat <- matrix(c(3, 6, 9, 3, 6, 9), nrow = 2, byrow = TRUE)
n_tox_mat <- matrix(c(0, 1, 3, 0, 1, 2), nrow = 2, byrow = TRUE)
eliminated_mat <- matrix(FALSE, nrow = 2, ncol = 3)

iso_est_mat <- isotonic_regression(n_pts_mat, n_tox_mat)
mtd_results <- select_mtd(iso_est_mat, n_pts_mat, eliminated_mat, target)
print(mtd_results)
```

---

sim_boin *Run BOIN Simulation with Operating Characteristics*

Description

Execute multiple BOIN trial simulations and compute operating characteristics. Combines patient enrollment, toxicity tracking, isotonic regression, and MTD selection into a single streamlined function.

Usage

```
sim_boin(
  n_trials = 10000,
  target,
  p_true,
  n_cohort,
  cohort_size,
  p_saf = NULL,
  p_tox = NULL,
  n_earlystop = 18,
  cutoff_eli = 0.95,
  extrasafe = FALSE,
  offset = 0.05,
  n_earlystop_rule = "with_stay",
  titration = FALSE,
  min_mtd_sample = 1,
  boundMTD = FALSE,
  return_details = FALSE,
```

```
  verbose = FALSE,
  seed = 123
)
```

Arguments

n_trials	Numeric. Number of trials to simulate. Default is 10000.
target	Numeric. Target toxicity probability (e.g., 0.30 for 30%).
p_true	Numeric vector. True toxicity probabilities for each dose.
n_cohort	Numeric. Maximum number of cohorts per trial.
cohort_size	Numeric vector or scalar. Patients per cohort.
p_saf	Numeric. Highest toxicity probability deemed acceptable for safety. Default is 0.6 * target. Used with p_tox for safety/efficacy dose identification.
p_tox	Numeric. Lowest toxicity probability deemed unacceptable for toxicity. Default is 1.4 * target. Used with p_saf for safety/efficacy dose identification.
n_earlystop	Numeric. Sample size triggering early stopping. Default is 18.
cutoff_eli	Numeric. Cutoff probability for dose elimination. Default is 0.95.
extrasafe	Logical. Apply extra safety stopping rule at lowest dose. Default is FALSE.
offset	Numeric. Offset for safety cutoff when extrasafe = TRUE. Default is 0.05.
n_earlystop_rule	Character. Early stopping rule: "with_stay" or "simple". Default is "with_stay".
titration	Logical. Perform accelerated dose titration phase. Default is FALSE.
min_mtd_sample	Numeric. Minimum patients required for MTD consideration. Default is 1.
boundMTD	Logical. Impose constraint that MTD's isotonic estimate <= lambda_d. Default is FALSE.
return_details	Logical. If TRUE, return detailed trial-level results. If FALSE, return summary only. Default is FALSE.
verbose	Logical. If TRUE, print progress messages to console. If FALSE, run silently. Default is FALSE.
seed	Numeric. Random seed for reproducibility. Default is 123.

Details

The function executes the following workflow:

1. Generate patient enrollment and toxicity data via `get_pts_and_tox`
2. Apply isotonic regression via `isotonic_regression`
3. Select MTD via `select_mtd`
4. Compute operating characteristics (MTD selection rates, sample sizes, DLT counts)

Progress messages are printed to console only if verbose = TRUE. If return_details = TRUE, detailed trial-level results are returned; otherwise, only summary statistics are computed for efficiency.

Value

A list containing:

`detailed_results`
: List of length n_trials (if return_details = TRUE). Each element contains: n_pts, n_tox, mtd, iso_est, reason, cohorts_completed. NULL if return_details = FALSE.

`summary`
: Object of class "boin_summary" containing: p_true (true toxicity probabilities), mtd_selection_percent (MTD selection rate for each dose and percent with no MTD), avg_n_pts (average patients per dose), avg_n_tox (average DLTs per dose), avg_total_n_pts (average total patients per trial), avg_total_n_tox (average total DLTs per trial).

References

Liu S. and Yuan, Y. (2015). Bayesian Optimal Interval Designs for Phase I Clinical Trials. Journal of the Royal Statistical Society: Series C, 64, 507-523.

Examples

```
# Basic BOIN simulation (silent mode)
result <- sim_boin(
  n_trials = 10000,
  target = 0.30,
  p_true = c(0.10, 0.25, 0.40, 0.55, 0.70),
  n_cohort = 10,
  cohort_size = 3,
  seed = 123
)

# With progress messages
result_verbose <- sim_boin(
  n_trials = 10000,
  target = 0.30,
  p_true = c(0.10, 0.25, 0.40, 0.55, 0.70),
  n_cohort = 10,
  cohort_size = 3,
  verbose = TRUE,
  seed = 123
)
```

---

sim_boin_multi *Run BOIN Simulation for Multiple Scenarios*

Description

Execute BOIN trial simulations across multiple dose-toxicity scenarios and compute operating characteristics for each scenario. Results are aggregated into a combined summary table suitable for protocol development and reports.

Usage

```
sim_boin_multi(
  scenarios,
  target,
  n_trials = 10000,
  n_cohort,
  cohort_size,
  p_saf = NULL,
  p_tox = NULL,
  n_earlystop = 18,
  cutoff_eli = 0.95,
  extrasafe = FALSE,
  offset = 0.05,
  n_earlystop_rule = "with_stay",
  titration = FALSE,
  min_mtd_sample = 1,
  boundMTD = FALSE,
  return_details = FALSE,
  verbose = FALSE,
  seed = 123
)
```

Arguments

scenarios	List of scenario definitions. Each element must be a list containing: name Character. Scenario identifier (e.g., "Scenario 1") p_true Numeric vector. True toxicity probabilities for each dose
target	Numeric. Target toxicity probability (e.g., 0.30 for 30%).
n_trials	Numeric. Number of trials to simulate per scenario. Default is 10000.
n_cohort	Numeric. Maximum number of cohorts per trial.
cohort_size	Numeric vector or scalar. Patients per cohort.
p_saf	Numeric. Highest toxicity probability deemed acceptable for safety. Default is 0.6 * target. Used with p_tox for safety/efficacy dose identification.
p_tox	Numeric. Lowest toxicity probability deemed unacceptable for toxicity. Default is 1.4 * target. Used with p_saf for safety/efficacy dose identification.
n_earlystop	Numeric. Sample size triggering early stopping. Default is 18.
cutoff_eli	Numeric. Cutoff probability for dose elimination. Default is 0.95.
extrasafe	Logical. Apply extra safety stopping rule at lowest dose. Default is FALSE.
offset	Numeric. Offset for safety cutoff when extrasafe = TRUE. Default is 0.05.
n_earlystop_rule	Character. Early stopping rule: "with_stay" or "simple". Default is "with_stay".
titration	Logical. Perform accelerated dose titration phase. Default is FALSE.
min_mtd_sample	Numeric. Minimum patients required for MTD consideration. Default is 1.

boundMTD	Logical. Impose constraint that MTD's isotonic estimate <= lambda_d. Default is FALSE.
return_details	Logical. If TRUE, return detailed trial-level results. Default is FALSE.
verbose	Logical. If TRUE, print progress messages to console. If FALSE, run silently. Default is FALSE.
seed	Numeric. Random seed for reproducibility. Default is 123.

Details

Progress messages are printed to console only if verbose = TRUE. For large-scale studies (50+ scenarios), computation time may be substantial.

Value

A list with class "boin_multi_summary" containing:

results_by_scenario	List of simulation results for each scenario
combined_summary_df	Data frame with aggregated operating characteristics
scenario_names	Character vector of scenario identifiers
n_doses	Numeric. Number of doses evaluated
call	The function call

References

Liu S. and Yuan, Y. (2015). Bayesian Optimal Interval Designs for Phase I Clinical Trials. Journal of the Royal Statistical Society: Series C, 64, 507-523.

Examples

```
scenarios <- list(
  list(name = "Scenario 1", p_true = c(0.05, 0.10, 0.20, 0.30, 0.45)),
  list(name = "Scenario 2", p_true = c(0.10, 0.15, 0.30, 0.45, 0.60))
)

# Silent mode (default)
result <- sim_boin_multi(
  scenarios = scenarios,
  target = 0.30,
  n_trials = 10000,
  n_cohort = 48,
  cohort_size = 3,
  seed = 123
)

# With progress messages
result_verbose <- sim_boin_multi(
  scenarios = scenarios,
```

```
  target = 0.30,
  n_trials = 10000,
  n_cohort = 48,
  cohort_size = 3,
  verbose = TRUE,
  seed = 123
)
```

Index