Package 'pipe.design'

July 23, 2025

Type Package

Title Dual-Agent Dose Escalation for Phase I Trials using the PIPE
Design

Version 0.5.1

Date 2017-06-16

Author Michael Sweeting

Maintainer Michael Sweeting <mjs212@medschl.cam.ac.uk>

Description
Implements the Product of Independent beta Probabilities dose Escalation (PIPE) design for dual-agent Phase I trials as described in Mander AP, Sweeting MJ (2015) <DOI:10.1002/sim.6434>.

License GPL (>= 2)

LazyLoad yes

Imports ggplot2 (>= 1.0.1), gtools, xtable

Suggests shiny

NeedsCompilation no

Repository CRAN

Date/Publication 2017-06-16 17:57:18 UTC

Contents

1

| pipe.design | *Dual-Agent Dose Escalation for Phase I Trials using the PIPE Design* |

Description

Implements the Product of Independent beta Probabilities dose Escalation (PIPE) design for dual-agent Phase I trials as described in Mander A.P. and Sweeting M.J (2015) Statistics in Medicine. <doi: 10.1002/sim.6434>

Usage

```
pipe.design(N=dim(data)[1]+1, S=1, c, theta, pi=NULL,  prior.med=NULL,
prior.ss=NULL, strategy, admis, constraint="none",  epsilon=NULL,
mode="sim", data=matrix(nrow=0,ncol=0), a=NULL, b=NULL,
alternate=FALSE,  uppertox.constraint=NULL, stop=NULL,
non.admissible=NULL, seed=NULL)
```

Arguments

N	Sample size for the trial. If not specified this is assumed to be one greater than the sample size of the patients recruited thus far (i.e. the number of rows of data plus one)
S	Number of simulations to conduct. If set to 1 then a single trial is run and invokes different output from print.pipe and plot.pipe
c	Cohort size. Must be divisible by N
theta	The target toxicity probability
pi	A matrix with rows denoting levels of drug A and columns denoting levels of drug B. Each element gives the true probability of the outcome (dose-limiting toxicity) for that dose combination. If omitted then the trial will recommend only for the next dose combination
prior.med	A matrix with rows denoting levels of drug A and columns denoting levels of drug B. Each element gives the prior median probability of the outcome (dose-limiting toxicity) for that dose combination. Should be specified in combination with prior.ss
prior.ss	A matrix with rows denoting levels of drug A and columns levels of drug B. Each element gives the prior sample size for that dose combination. Should be specified in combination with prior.med
strategy	A string detailing the dose escalation strategy. Options are "ss" (sample size) and "ss-random" (weighted randomisation by sample size). See details below
admis	A string detailing how to choose admissible doses around the maximum tolerated contour (MTC). Options are "adjacent" and "closest". See details below
constraint	A string detailing any dose-skipping constraint to be applied. Options are "neighbouring" or "no.dose.skip" (see details below). In addition, to prevent diagonal escalation the option "-nodiag" can be appended, i.e. "neighbouring-nodiag" or "no.dose.skip-nodiag". Defaults to no constraint (constraint="none").

epsilon	Any number between 0 and 1 giving a safety threshold. Doses whose (weighted) posterior probability(dose>MTC)>epsilon will not be experimented on. See details below
mode	A string indicating the running mode of the program. Options are "sim" (the default, used for simulations), "nodlt" (every patient recruited is assumed to be DLT free), or "alldlt" (every patient recruited is assumed to suffer a DLT).
data	An optional named data frame giving information about dose and toxicity from previously recruited patients. If missing then it is assumed that no data have thus far been collected. Contains the following variables:
	patient Recruited patient numbers, 1,...,n
	doseA Dose levels of recruited patients for drug A
	doseB Dose levels of recruited patients for drug B
	tox An indicator variable for each patient (1=toxicity, 0=no toxicity)
a	(Optional). A matrix with rows denoting levels of drug A and columns levels of drug B. Each element gives the shape1 parameter (a) from a prior Beta(a,b) distribution. Can be used as an alternative to specifying prior.med and prior.ss and must be used in conjunction with b
b	(Optional). A matrix with rows denoting levels of drug A and columns levels of drug B. Each element gives the shape2 parameter (b) from a prior Beta(a,b) distribution. Can be used as an alternative to specifying prior.med and prior.ss and must be used in conjunction with a
alternate	Logical variable (defaults to FALSE). Should the design always deescalate if above the MTC and escalate if below (subject to admissibility, safety and dose skipping constraints)?
uppertox.constraint	
	Any number between 0 and 1 giving the upper toxicity safety constraint. No dose combination that lies above the most likely contour for this constraint may be dosed. Defaults to NULL
stop	(Optional). Value of a stopping threshold. An alternative stopping rule to a fixed sample size specification (through N) or termination of a trial due to safety constraints (epsilon or uppertox.constraint). The trial is stopped if the posterior probability of being > theta at the lowest dose combination is greater than stop
non.admissible	(Optional) A matrix with rows denoting levels of drug A and columns denoting levels of drug B. Each element is logical indicating whether the dose combination is non-admissible (TRUE) or not. Non-admissible doses can never be experimented on in during the trial.
seed	(Optional) Starting seed to allow replication of the trial or simulation.

Details

This function implements a Product of Independent Beta Probabilities (PIPE) Phase I dose-escalation design for a dual-agent trial. An independent Beta prior is placed on the probability of dose limiting toxicity (DLT) for each dose combination. Simple conjugate Bayesian analysis is used to obtain the posterior at each dose combination. This model is considered a 'working model' only as it does not account for monotonically increasing toxicity by dose.

The dose combination assigned to the next cohort of patients is, however, obtained by calculating a maximum tolerated contour (MTC) that bisects the dose space and respects monotonicity. Admissible dose combinations considered for the next cohort are those that are either 'adjacent' or 'closest' to the MTC (see Mander A.P. and Sweeting M.J., Statistics in Medicine 2015) and can be specified using the `admis` argument. If there is more than one dose combination then either the one with the smallest sample size (least experimented on) is chosen, by specifying `strategy="ss"`, or admissible doses are chosen at random, weighted by the inverse of the sample size (`strategy="ss-random"`).

Admissible doses can be constrained to be within a set of neighbouring doses to the current dose combination (`constraint="neighbouring"`) or within the set of doses consisting of up to one dose level above any previously experimented drug A and drug B level (`constraint="no.dose.skip"`). To additionally prevent diagonal escalation (e.g. where both drugs are escalated together), use either (`constraint="neighbouring-nodiag"`) or (`constraint="no.dose.skip-nodiag"`).

In addition, a safety constraint can be added, to avoid potential overdosing. Dose combinations whose expected probability of being above the MTC is greater than a threshold `epsilon` are excluded from the admissible doses for the next cohort. The expected probability is calculated by averaging over all possible MTCs, weighted by their posterior probabilities. An alternative safety constraint can be used by specifying `uppertox.constraint=c`, where c is the upper toxicity level from which a contour is derived (e.g. c=0.6 calculates the most likely contour for a 60% chance of DLT). No doses above this contour will be used.

The program will either run as a simulation (the default), or by using a deterministic set of outcomes (all DLTs / no DLTs) specified using `mode="alldlt"` and `mode="nodlt"`, respectively.

If the user specifies a data-frame using the `data` argument, then every simulated trial will begin with the specified doses and DLT outcomes. This option is useful for assess operating characteristics mid-way through the trial, or just to obtain the dose for the next cohort.

A trial of size `N` is generally run, except if no doses are deemed safe according to one of the two safety constraints described above, at which point the trial will stop. Users may additionally specify a stopping rule using `stop=s`, where the trial is stopped if the posterior probability that the lowest dose combination is above the target toxicity probability is greater than s. Note, this stopping rule relies only on the prior and data collected for the first dose combination and does not borrow strength from other doses. Hence it should be used with caution.

One can force deescalation if above the MTC and escalation if below (subject to meeting other specified constraints) by specifying `alternate=TRUE`. This constraint ensures the design follows the properties of 'coherence' (Cheung YK. Biometrika 2005).

If a single trial is specified by `S=1` then by additionally specifying `densities=TRUE` the returned object will include calculations of the probability of toxicity densities for each dose-combination after each cohort (with cohort 0 representing the prior). The densities are calculated over a grid of probabilities $\alpha_1, ..., \alpha_K$ through weighting by the contour probabilities for each α_k:

$$p(\alpha_k < \pi_{i,j} \leq \alpha_{k+1}) = \sum_{C_s \in \mathcal{C}} C_s[i,j] P_{C_s}(\alpha_{k+1}) - \sum_{C_s \in \mathcal{C}} C_s[i,j] P_{C_s}(\alpha_k)$$

Further details regarding the notation can be found in Mander A.P. and Sweeting M.J. (2015) Statistics in Medicine. <doi: 10.1002/sim.6434>

Value

`pipe.design` returns an object of class "pipe" or "pipe.sim"; the latter occuring when more than

one simulation has been conducted. The function `print` (i.e. `print.pipe` or `print.pipe.sim`) can be used to obtain summary information about the design used, the data observed, the maximum tolerated contour estimate, and the next recommended dose level.

An object of class "pipe" (when N=1) or "pipe.sim" (when N>1) is a list with the following components:

`r.sim` A list of matrices showing the number of DLTs for each dose combination in each simulation

`n.sim` A list of matrices showing the number of patients dosed at each combination in each simulation

`rec.i.sim` An S x (N/c) matrix with elements representing the level of drug A for each cohort

`rec.j.sim` An S x (N/c) matrix with elements representing the level of drug B for each cohort

`exp` A matrix showing the experimentation proportions for each dose combination averaged across simulations

`rec` A matrix showing the proportion of times each dose combination is recommended, where the denominator is the total number of recommended phase II doses over the simulations

`dlts` A vector giving the average number of DLTs occuring in each simulated trial

`mat.list` A list of binary matrices showing the most likely MTC before each cohort is recruited. Output only if S=1

`uppermat.list` A list of binary matrices showing the most likely contour for `uppertox.constraint` before each cohort is recruited. Output only if S=1 and `uppertox.constraint` is used

`uppermat2.list` A list of binary matrices showing the doses that fail the constraint imposed by `epsilon`. Output only if S=1 and `epsilon` is used

`r.list` A list of matrices showing the cumulative number of DLTs for each dose combination before each cohort is recruited. Output only if S=1

`n.list` A list of matrices showing the cumulative number of patients at each dose combination before each cohort is recruited. Output only if S=1

`n.rpII` A vector listing the number of doses recommended for Phase II from each simulated trial

`no.not.treated` The total number of patients not treated (in the trial) accumulated across the simulated trials due to a trial stopping early

`pi` The true probability of DLT for each dose combination, as specified in the input argument

`theta` The target toxicity probability, as specified in the input argument

`rpII.list` A list giving the recommended Phase II dose combinations for each simulation

`a` A matrix with rows denoting levels of drug A and columns levels of drug B. Each element gives the `shape1` parameter (a) from a prior Beta(a,b) distribution

`b` A matrix with rows denoting levels of drug A and columns levels of drug B. Each element gives the `shape2` parameter (b) from a prior Beta(a,b) distribution

`plot.density` A list of objects of length c+1. The kth element is a ggplot object giving the toxicity density plots for each dose-combination after the (k-1)st cohort has been recruited

Note

The use of the `uppertox.constraint` safety constraint is not documented in Mander and Sweeting, Statistics in Medicine 2015, and some preliminary investigations have shown that this constraint can cause rigidity in the escalation process, and therefore should be used with caution.

Author(s)

Michael Sweeting <mjs212@medschl.cam.ac.uk> (University of Cambridge, UK)

References

Mander A.P., Sweeting M.J. A product of independent beta probabilities dose escalation design for dual-agent phase I trials. *Statistics in Medicine* (2015) 34(8): 1261–1276. <doi: 10.1002/sim.6434>

Cheung Y.K. Coherence principles in dose-finding studies. *Biometrika* (2005) 92(4): 863–873.

See Also

print.pipe, print.pipe.sim, plot.pipe, plot.pipe.sim

Examples

```
## Reproducing Figure 5 from Mander and Sweeting, Statistics in Medicine 2015.
## The true probability of DLT for a 6x6 grid of dose combinations
## (Scenario 3 from Mander and Sweeting, Statistics in Medicine 2015)
pi <-matrix(c(0.02,0.10,0.20,0.30,0.35,0.45,0.06,0.14,0.24,0.34,0.39,0.49,0.12,0.20,
0.30,0.40,0.45,0.55,0.17,0.25,0.35,0.45,0.50,0.60,0.22,0.30,0.40,0.50,0.60,0.70,0.30,
0.38,0.48,0.58,0.68,0.78),nrow=6,ncol=6)
## Using a weak prior with prior sample size 1/36 for each dose combination and prior
## median of Scenario 1
prior.med<-matrix(c(0.02,0.03,0.06,0.10,0.18,0.23,0.03,0.05,0.09,0.13,0.21,0.27,0.06,
0.09,0.14,0.18,0.26,0.30,0.11,0.14,0.18,0.23,0.30,0.36,0.18,0.21,0.26,0.30,0.39,0.44,
0.23,0.27,0.3,0.36,0.44,0.49),nrow=6,ncol=6)
## Using a neighbouring escalation constraint
## Selecting the closest admissible doses
## Using a safety constraint with epsilon = 0.8

## Obtain dose recommendations for first cohort
cohort1<-pipe.design(N=2,S=1,c=2,theta=0.3,prior.med=prior.med,
prior.ss=matrix(1/36,ncol=6,nrow=6),strategy="ss",constraint="neighbouring",
epsilon=0.8,admis="closest",alternate=FALSE)
cohort1 ## Recommends starting at (1,1)
## Second cohort
## Assume no toxicities are seen in first cohort
data1<-data.frame(patient=1:2,doseA=rep(1,2),doseB=rep(1,2),tox=rep(0,2))
cohort2<-pipe.design(data=data1,S=1,c=2,theta=0.3,prior.med=prior.med,
prior.ss=matrix(1/36,ncol=6,nrow=6),strategy="ss",constraint="neighbouring",
epsilon=0.8,admis="closest",alternate=FALSE)
cohort2 ## Recommends dosing at (2,2)
## Third cohort
## Assume one toxicity is seen in second cohort
data2<-rbind(data1,data.frame(patient=3:4,doseA=rep(2,2),doseB=rep(2,2),tox=c(1,0)))
cohort3<-pipe.design(data=data2,S=1,c=2,theta=0.3,prior.med=prior.med,
   prior.ss=matrix(1/36,ncol=6,nrow=6),strategy="ss",constraint="neighbouring",
   epsilon=0.8,admis="closest",alternate=FALSE)
cohort3 ## Recommends dosing at (1,3)
```

```
## Simulating a single trial with sample size of 40, cohort size 2
## and target toxicity of 0.3, using true probabilities of toxicity pi
## Not run:
set.seed(700)
example<-pipe.design(N=40,S=1,c=2,theta=0.3,pi=pi,prior.med=prior.med,
prior.ss=matrix(1/36,ncol=6,nrow=6),strategy="ss",constraint="neighbouring",
epsilon=0.8,admis="closest",alternate=FALSE)
example
plot(example)

## End(Not run)



## Not run:
## Run a Shiny GUI version of pipe.design
runShinyPIPE()

## End(Not run)



## Not run:
## Conducting a simulation study (1000 trials) to investigate
## operating characteristics for Scenario 3
## (as reported in Table III in Mander and Sweeting, Statistics in Medicine 2015)
set.seed(262)
scen3<-pipe.design(N=40,S=1000,c=2,theta=0.3,pi=pi,prior.med=prior.med,
prior.ss=matrix(1/36,ncol=6,nrow=6),strategy="ss",constraint="neighbouring",
epsilon=0.8,admis="closest",alternate=FALSE)
scen3

## End(Not run)
```

---

plot.pipe   *Plot escalation or operating characteristics for trials conducted using*
                                   the Product of Independent Beta Probabilities Escalation design

---

Description

Plot method for a trial or series of trials conducted using a `pipe.design` model.

Usage

```
## S3 method for class 'pipe'
plot(x, type = "b", pi = x$pi, theta = x$theta, epsilon.line = TRUE,
    uppertox.constraint.line = FALSE, add.empirical.data = FALSE, ...)
## S3 method for class 'pipe.sim'
plot(x, pi = x$pi, theta = x$theta, plot = "both", ...)
```

Arguments

x	An object of class "pipe" or "pipe.sim" as returned by `pipe.design`
type	The type of information to be plotted for a single PIPE trial. Options are "r" (number of outcomes only), "n" (number of patients only), or "b" (both number of outcomes and patients). The default is to plot both the number of outcomes and patients for each dose combination, by using symbols and colour filling, respectively.
pi	A matrix with rows denoting levels of drug A and columns denoting levels of drug B. Each element gives the true probability of the outcome (dose-limiting toxicity) for that dose combination. If omitted then the true probabilities of the outcome will be taken from that used when creating the x object.
theta	The target toxicity probability. If omitted then `theta` is taken from that used when creating the x object.
epsilon.line	Should the toxicity safety contour specified by epsilon in `pipe.design` be shown? Defaults to TRUE.
uppertox.constraint.line	
	Should the toxicity safety contour specified by uppertox.constraint in `pipe.design` be shown? Defaults to FALSE.
add.empirical.data	
	Should an extra plot be produced with the empirical posterior median toxicity probabilities (from the independent beta distributions), and 95
plot	The type of operating characteristics to plot for a series of simulated PIPE trials. Options are "exp" (experimentation percentages), "rec" (recommendation percentages), or "both" (both experimentation and recommendation percentages side-by-side). The default is "both".
...	Further arguments passed to or from other methods

Details

If a single trial is conducted, then the `plot` function currently produces a series of matrix plots by cohort showing the cumulative outcome and sample size information for each dose combination. Also shown as a solid black line is the most likely MTC before outcomes from each cohort are observed and, if specified, one or more toxicity safety contours shown in red. The final cohort shows the most likely MTC at the end of the trial together with the true MTC (green dashed line) and the recommended Phase II doses (blue squares), which are always doses below the most likely MTC that have been experimented on (see Mander and Sweeting, Statistics in Medicine 2015 for more details).

If multiple PIPE trials have been simulated then experimentation and recommendation percentages are shown as a heat plot over the dose space together with the true MTC (green dashed line).

Author(s)

Michael Sweeting <mjs212@medschl.cam.ac.uk>

References

Mander A.P., Sweeting M.J. A product of independent beta probabilities dose escalation design for dual-agent phase I trials. *Statistics in Medicine* (2015) 34(8): 1261–1276.

See Also

[pipe.design](), [print.pipe]()

---

print.pipe	*Print information regarding a trial conducted using the Product of Independent Beta Probabilities Escalation design*

Description

Print method for a trial or series of trials conducted using a [pipe.design]() model.

Usage

```
## S3 method for class 'pipe'
print(x, ...)
## S3 method for class 'pipe.sim'
print(x, pi = x$pi,
  cut.points = unique(c(0,pmin(1,pmax(0,seq(x$theta-0.15,x$theta+0.15,by=0.1))),1)),
   digits = 1, print = TRUE, ...)
```

Arguments

x	An object of class "pipe" or "pipe.sim" as returned by [pipe.design]()
pi	A matrix with rows denoting levels of drug A and columns denoting levels of drug B. Each element gives the true probability of the outcome (dose-limiting toxicity) for that dose combination. If omitted then the true probabilities of the outcome will be taken from that used when creating the x object.
cut.points	Cutpoints of toxicity for which the operating characteristics are to be categorised. Default is [0,x$theta-0.1) [x$theta-0.1,x$theta+0.1) [x$theta+0.1,1]
digits	The number of decimal places to print the operating characteristics
print	If TRUE then the experimentation and recommendation percentages are printed to the output
...	Further arguments passed to or from other methods

Details

If a single trial is conducted, then the [print]() function currently produces summary information about the data observed, current best estimate of the MTC and the upper toxicity constraint contour. If a simulation study is conducted, then the following operating characteristics are printed:

Experimentation percentages Percentage of patients recruited to each true region of toxicity (as specified by cut.points), across the simulated trials

Recommendation percentages Percentage of times that recommend doses lie within each true region of toxicity (as specified by cut.points). As more than one dose combination can be recommended in any trial, the denominator is the total number of recommended phase II doses over the simulations

Percentage of times a trial recommends k doses for Phase II Provides the percentage of time that each trial recommends 0, 1, ..., k doses for Phase II experimentation. 0 doses indicates that the trial stopped early and recommended no doses.

Author(s)

Michael Sweeting <mjs212@medschl.cam.ac.uk>

References

Mander A.P., Sweeting M.J. A product of independent beta probabilities dose escalation design for dual-agent phase I trials. *Statistics in Medicine* (2015) 34(8): 1261–1276.

See Also

[pipe.design](), [plot.pipe]()

---

runShinyPIPE	*A Shiny GUI implementation of the PIPE design*

Description

Runs a Shiny web browser interface of the Product of Independent beta Probabilities dose Escalation (PIPE) design

Usage

```
runShinyPIPE()
```

Details

This function implements a GUI web browser version of pipe.design, thus allowing easy usage and communication of the design to researchers.

Author(s)

Chris Habron (Roche) and Michael Sweeting <mjs212@medschl.cam.ac.uk> (University of Cambridge, UK)

References

Mander A.P., Sweeting M.J. A product of independent beta probabilities dose escalation design for dual-agent phase I trials. *Statistics in Medicine* (2015) 34(8): 1261–1276.

See Also

[pipe.design]()

Index