

Package ‘eadrm’

October 13, 2022

Title Fitting Dose-Response Models Using an Evolutionary Algorithm

Version 0.1.4

Description Fits dose-response models using an evolutionary algorithm to estimate the model parameters. The procedure currently can fit 3-parameter, 4-parameter, and 5-parameter log-logistic models as well as exponential models. Functions are also provided to plot, make predictions, and calculate confidence intervals for the resulting models. For details see “Nonlinear Dose-response Modeling of High-Throughput Screening Data Using an Evolutionary Algorithm”, Ma, J., Bair, E., Motsinger-Reif, A.; Dose-Response 18(2):1559325820926734 (2020) <[doi:10.1177/1559325820926734](https://doi.org/10.1177/1559325820926734)>.

Depends R (>= 2.10.0)

License GPL-3

Encoding UTF-8

LazyData true

RoxygenNote 7.2.0

NeedsCompilation no

Author Eric Bair [aut, cre],
Jun Ma [aut]

Maintainer Eric Bair <eric.bair@sciome.com>

Repository CRAN

Date/Publication 2022-05-29 05:40:06 UTC

R topics documented:

calc.ec	2
calc.ed	3
CarboA	3
eadrm	4
eadrm.ci	6
plot.eadrm	7
predict.eadrm	7

calc.ec	<i>Calculates EC50 (or some other specified EC value)</i>
---------	---

Description

Calculates the concentration that induces a response corresponding to a specific proportion between the baseline and maximum. It is most commonly used to compute EC50.

Usage

```
calc.ec(eadrm.obj, ec.in = 0.5, ci.obj = NULL)
```

Arguments

eadrm.obj	Fitted eadrm model object.
ec.in	A value between 0 and 1 corresponding to the desired proportion. Defaults to 0.5, in which case EC50 is computed.
ci.obj	Output from eadrm.ci that is used to compute confidence intervals. Defaults to NULL, in which case no confidence intervals are computed.

Value

If no confidence intervals are requested, it returns the concentration corresponding to the requested proportion. Otherwise returns a list of three values `ec.med`, `ec.l95`, and `ec.u95`, which correspond to the median concentration and the corresponding lower/upper 95% confidence bounds.

See Also

[eadrm](#), [eadrm.ci](#)

Examples

```
ea.fit <- eadrm(CarboA$y, CarboA$x)
calc.ec(ea.fit)
```

calc.ed	<i>Finds the dose that corresponds to a particular level of the response</i>
---------	--

Description

Calculates the concentration that induces a particular level of the response.

Usage

```
calc.ed(eadrm.obj, response, ci.obj = NULL)
```

Arguments

eadrm.obj	Fitted eadrm model object.
response	The desired response level.
ci.obj	Output from eadrm.ci that is used to compute confidence intervals. Defaults to NULL, in which case no confidence intervals are computed.

Value

If no confidence intervals are requested, it returns the concentration corresponding to the specified response. Otherwise returns a list of three values `ed.med`, `ed.l95`, and `ed.u95`, which correspond to the median concentration and the corresponding lower/upper 95% confidence bounds.

See Also

[eadrm](#), [eadrm.ci](#)

Examples

```
ea.fit <- eadrm(CarboA$y, CarboA$x)
calc.ec(ea.fit)
```

CarboA	<i>Example dose response dataset</i>
--------	--------------------------------------

Description

A simple dose response dataset used to illustrate the use of the eadrm functions.

Usage

```
CarboA
```

Format

A list containing two elements:

x The drug concentration levels

y The response values ...

eadrm

Fits a dose-response curve using an evolutionary algorithm

Description

Uses an evolutionary algorithm to fit a set of four possible dose-response models. The evolutionary parameters and stopping rules can be customized by the user.

Usage

```
eadrm(
  obs,
  xvals,
  model = "h4",
  pop.size = 1000,
  stable.pop.size = 200,
  num.tournaments = 20,
  tournament.size = 10,
  max.generations = 500,
  stop.generations = 50,
  epsilon = 0.1
)
```

Arguments

<code>obs</code>	A vector of response values (y-values).
<code>xvals</code>	A vector of doses (x-values).
<code>model</code>	Type of dose-response model to fit. Possible values include "h3", "h4", and "h5" (corresponding to 3-parameter, 4-parameter, and 5-parameter log-logistic models, respectively), "e" (corresponding to an exponential model) and "all" (which allows the procedure to evaluate all four types of models). Defaults to "h4".
<code>pop.size</code>	The number of initial potential solutions. Defaults to 1000.
<code>stable.pop.size</code>	This quantity is divided by the number of tournaments to calculate the number of children in each generation. Defaults to 200.
<code>num.tournaments</code>	The number of tournaments in each generation. Defaults to 20.

<code>tournament.size</code>	The number of players (i.e., models to consider) in each tournament. Defaults to 10.
<code>max.generations</code>	The maximum number of generations. If this number is reached, the algorithm immediately terminates. Defaults to 500.
<code>stop.generations</code>	The algorithm will also terminate if there is no improvement in fitness in <code>stop.generation</code> generations. Defaults to 50.
<code>epsilon</code>	If three successive new models produce an improvement of less than <code>epsilon</code> in fitness, the procedure will terminate. Defaults to 0.1.

Value

An object of class `eadrm`, which is a list containing the following elements:

Model: Specifies the type of model (i.e., "h3", "h4", "h5", or "e")

R2: Fitness for the final model

params: A vector of coefficients for the final model

xvals: The original x values (concentrations) for the model

yvals: The original y values (responses) for the model

Details

The procedure will initially generate `pop.size` possible solutions. The fitness for each solution will be calculated. In each generation, a series of tournaments are performed. The children of the surviving models from the previous generation are mutated. The model with the best fitness "wins" each tournament and survives to the next generation. The procedure continues until the maximum number of generations is reached or the fitness fails to improve substantially over a sufficient number of generations.

References

Ma, J., Bair, E., Motsinger-Reif, A. "Nonlinear Dose-response Modeling of High-Throughput Screening Data Using an Evolutionary Algorithm", *Dose Response* 18(2):1559325820926734 (2020).

Examples

```
ea.fit <- eadrm(CarboA$y, CarboA$x)
```

eadrm.ci

*Computes confidence intervals for an eadrm model fit***Description**

Calculates confidence intervals for an eadrm model fit by repeatedly fitting the model to the same data set and examining the distribution of the coefficients.

Usage

```
eadrm.ci(obs, xvals, model = "h4", ..., B = 1000)
```

Arguments

obs	A vector of response values (y-values).
xvals	A vector of doses (x-values).
model	Type of dose-response model to fit. Possible values include "h3", "h4", and "h5" (corresponding to 3-parameter, 4-parameter, and 5-parameter log-logistic models, respectively) and "e" (corresponding to an exponential model). Defaults to "h4".
...	Additional parameters for the eadrm function.
B	Number of replicate models to fit. Defaults to 1000.

Value

A list containing the following elements:

med.est: A vector of the median values of the coefficients across the B iterations

l95.est,u95.est: Vectors of the lower/upper 95% confidence bounds for the coefficients across the B iterations

replicate.mat: A $p \times B$ matrix, where p is the number of coefficients in the model. Each column of B corresponds to the coefficients for one fitted model.

Details

This function calls the [eadrm](#) function B times with the same parameters and records the model coefficients for each iteration of the model. Confidence intervals for the coefficients are calculated by examining the quantiles of the distribution of the coefficients over the B iterations. A matrix of the coefficients for each iteration is also calculated. This matrix can be used to compute confidence intervals for predicted values and estimates of EC50.

See Also

[eadrm](#), [predict.eadrm](#), [calc.ec](#), [calc.ed](#)

Examples

```
ea.ci <- eadrm.ci(CarboA$y, CarboA$x)
```

plot.eadrm	<i>Plot an eadrm object</i>
------------	-----------------------------

Description

Plots the data used to fit an eadrm object as well as the fitted dose-response curve

Usage

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

Arguments

x	Object of class eadrm to plot.
...	Additional arguments to plot. Currently ignored.

See Also

[eadrm](#)

Examples

```
ea.fit <- eadrm(CarboA$y, CarboA$x)
plot(ea.fit)
```

predict.eadrm	<i>Makes predictions from an eadrm object</i>
---------------	---

Description

Similar to other predict methods, this function predicts fitted values (and optionally confidence intervals) for a eadrm object.

Usage

```
## S3 method for class 'eadrm'
predict(object, newx = object$xvals, ci.obj = NULL, ...)
```

Arguments

object	Fitted model of class eadrm.
newx	Vector of new concentration values at which predictions are to be made. Defaults to the concentration values that were used to fit the model.
ci.obj	Output from eadrm.ci that is used to compute confidence intervals. Defaults to NULL, in which case no confidence intervals are computed.
...	Additional arguments passed to or from other methods. Currently ignored.

Value

If no confidence intervals are requested, a vector of predicted responses for each concentration in `newx` is returned. Otherwise returns a list of three vectors `yhat.med`, `yhat.l95`, and `yhat.u95`, which correspond to the median predicted response and lower/upper 95% confidence bounds for each concentration in `newx`.

See Also

[eadrm](#), [eadrm.ci](#)

Examples

```
ea.fit <- eadm(CarboA$y, CarboA$x)
predict(ea.fit)
```


Index

* datasets

CarboA, 3

calc.ec, 2, 6

calc.ed, 3, 6

CarboA, 3

eadrm, 2, 3, 4, 6–8

eadrm.ci, 2, 3, 6, 8

plot.eadrm, 7

predict.eadrm, 6, 7