

Package ‘nlm’

May 9, 2026

Version 1.4.6

Title Non-Linear Transformation Models

Depends survival, R (>= 2.8.1)

Description Fits a non-linear transformation model ('nlm') for analyzing survival data, see Tsodikov (2003) <doi:10.1111/1467-9868.00414>. The class of 'nlm' includes the following currently supported models: Cox proportional hazard, proportional hazard cure, proportional odds, proportional hazard - proportional hazard cure, proportional hazard - proportional odds cure, Gamma frailty, and proportional hazard - proportional odds.

License GPL-2

URL <https://github.com/mclements/nlm>

BugReports <https://github.com/mclements/nlm/issues>

LazyLoad No

LazyData No

NeedsCompilation yes

Author Gilda Garibotti [aut],
Alexander Tsodikov [aut],
Mark Clements [aut, cre]

Maintainer Mark Clements <mark.clements@ki.se>

Repository CRAN

Date/Publication 2025-07-23 15:00:02 UTC

Contents

anova.nlm	2
melanoma	3
nlm	3
nlm.control	6
nlm.object	7
summary.nlm	8

`anova.nltm`*Analysis of deviance table for NLT models*

Description

If a single model is specified, gives a sequential analysis of deviance table for that fit.

Each row of the table represents a model, the first corresponds to the NULL model and the following rows correspond to models in which variables have been added sequentially. If a model has both, long term and short term covariates, long term covariates are added first. The last row corresponds to the input model (object).

Each row of the table gives the likelihood ratio test for comparing the model corresponding to that row and the one without the indicated covariate (Deviance) and its degrees of freedom (Df). It also gives -2 times the loglikelihood for the model (Resid. Dev) and its degrees of freedom (Resid. Df). Optionally it gives the p-value for the likelihood ratio test.

If more than one object is specified, the table has a row for each model. It specifies -2 times the loglikelihood for the model (Resid. Dev) and its degrees of freedom (Resid. Df). It also gives the likelihood ratio test for comparing the model corresponding to the given row with that of the previous row. (This only makes statistical sense if the models are nested.) Optionally it gives the p-value for the likelihood ratio test.

Usage

```
## S3 method for class 'nltm'  
anova(object, ..., test=FALSE)
```

Arguments

<code>object</code>	An object of class <code>nltm</code> .
<code>...</code>	Further <code>nltm</code> objects.
<code>test</code>	If TRUE the p-value for the likelihood ratio test is computed. Default is FALSE.

Value

An object of class "anova" inheriting from class "data.frame" with the information given in the description.

Warning

The comparison between two or more models by `anova` will only be valid if they are fitted to the same dataset. This may be a problem if there are missing values.

See Also

[nltm](#), `anova`.

Examples

```

data(melanoma, package="nltn")
fit1 <- nltn(formula1=Surv(time,status) ~ size + stage,
             formula2=Surv(time,status) ~ size, data=melanoma, nlt.model="PH")
anova(fit1)
anova(fit1,test=TRUE)
fit2 <- nltn(formula1=Surv(time,status) ~ size + stage + age,
             formula2=Surv(time,status) ~ size + age, data=melanoma, nlt.model="PH")
anova(fit1,fit2,test=TRUE)

```

melanoma

Melanoma survival data

Description

Simulated melanoma survival data.

Usage

```
data(melanoma)
```

Format

time: survival or censoring time (years)
status: censoring status (0=censored, 1=death)
size: tumor size (cm)
stage: melanoma stage (factor)
age: patient age (factor)

nltn

Fit Non-Linear Transformation Model for analyzing survival data

Description

Fits a non-linear transformation model (nltn) for analyzing survival data, see Tsodikov (2003). The class of nltn includes the following currently supported models: Cox proportional hazard, proportional hazard cure, proportional odds, proportional hazard - proportional hazard cure, proportional hazard - proportional odds cure, Gamma frailty, and proportional hazard - proportional odds.

Usage

```
nlm(formula1=formula(data), formula2=formula(data),
    data=parent.frame(), subset, na.action, init=NULL, control,
    nlt.model=c("PH", "PHC", "PO", "PHPHC", "PHPOC", "GFM", "PHPO"),
    model=FALSE, x=FALSE, y=FALSE, verbose=FALSE, ...)
```

Arguments

formula1	A formula object with the response on the left of a ~ operator, and the terms on the right. The response must be a survival object as returned by the Surv function. In models with two predictors, this corresponds to the long term effect.
formula2	A formula corresponding to the short term effect. Will be ignored in models with only one predictor. If not present in models with two predictors, then formula1 will be used both for the long and short term effect.
data	A data.frame in which to interpret the variables named in formula1 and formula2, or in the subset argument.
subset	Expression saying that only a subset of the rows of the data should be used in the fit.
na.action	A missing-data filter function, applied to the model.frame, after any subset argument has been used. Default is options()\$na.action.
init	Vector of initial values for the calculation of the maximum likelihood estimator of the regression parameters. Default is zero.
control	Object specifying iteration limit and other control options. Default is nlm.control(...).
nlt.model	A character string specifying a non-linear transformation model. Default is Proportional Hazards Model.

The conditional survival function $S(t|z)$ given the covariates z of each of the models currently supported are given below. Let $S_0(t)$ be the non-parametric baseline survival function, and $\theta(z)$ and $\eta(z)$ predictors. We take $\theta(z) = \exp(\beta_\theta z)$ and $\eta(z) = \exp(\beta_\eta z)$, where β_θ and β_η are the regression coefficients. In cure models, there is an additional regression parameter β_c and $\theta(z) = \exp(\beta_\theta z + \beta_c)$.

- Proportional hazard model (PH):

$$S(t|z) = S_0(t)^{\theta(z)}.$$

- Proportional hazard cure model (PHC):

$$S(t|z) = \exp(-\theta(z)(1 - S_0(t))).$$

- Proportional odds model (PO):

$$S(t|z) = \frac{\theta(z)}{\theta(z) - \ln(S_0(t))}.$$

- Proportional hazard - proportional hazard cure model (PHPHC):

$$S(t|z) = \exp(-\theta(z)(1 - S_0^{\eta(z)}(t))).$$

- Proportional hazard - proportional odds cure model (PHPOC):

$$S(t|z) = \exp\left(-\frac{\theta(z)(1 - S_0(t))}{1 - (1 - \eta(z))S_0(t)}\right).$$

- Gamma frailty model (GFM):

$$S(t|z) = \left(\frac{\theta(z)^{\eta(z)}}{\theta(z) - \ln(S_0(t))}\right)^{\eta(z)}.$$

- Proportional hazard - proportional odds model (PHPO):

$$S(t|z) = \frac{\theta(z)S_0^{\eta(z)}(t)}{1 - (1 - \theta(z))S_0^{\eta(z)}(t)}.$$

model	If TRUE the model frame is stored. Default is FALSE.
x	If TRUE the model matrix is stored. Default is FALSE.
y	If TRUE the response matrix is stored. Default is FALSE.
verbose	If a file name is given, it stores information from maximization of likelihood and calculation of information matrix in a file. Otherwise, verbose=FALSE. Default is FALSE.
...	Other arguments.

Value

An object of class "nlm". See [nlm.object](#) for details.

Author(s)

Gilda Garibotti (garibotti@crub.uncoma.edu.ar) and Alexander Tsodikov.

References

- Tsodikov A. (2003) "Semiparametric models: a generalized self-consistency approach". *Journal of the Royal Statistical Society B*, 65, Part 3, 759-774.
- Tsodikov A. (2002) "Semi-parametric models of long- and short-term survival: an application to the analysis of breast cancer survival in Utah by age and stage". *Statistics in Medicine*, 21, 895-920.
- Tsodikov A., Garibotti G. (2006) "Profile information matrix for nonlinear transformation models". to appear in *Journal of Lifetime Data Analysis*.
- Tsodikov A., Ibrahim J., Yakovlev A. (2003) "Estimating cure rates from survival data: an alternative to two-component mixture models". *Journal of the American Statistical Association*, Vol. 98, No. 464, 1063-1078.
- Wendland M., Tsodikov A., Glenn M., Gaffney D. (2004) "Time interval to the development of breast carcinoma after treatment for Hodgkin disease". *Cancer* Vol. 101, No. 6, 1275-1282.

See Also

[nlm.object](#), [summary.nlm](#), [nlm.control](#)

Examples

```
# fit a Proportional Odds Model
data(melanoma, package="nlm")
fit <- nlm(Surv(time,status) ~ size + age, data=melanoma, nlm.model="PO")
summary(fit)
```

nlm.control

Package options for nlm

Description

Sets default values for arguments related to calculation of the maximum profile likelihood estimator of the regression parameters, β_θ (and β_η , and β_c), and the baseline hazard $S_0(t)$ (see [nlm](#)). Optimization is performed using the "L-BFGS-B" method by Byrd et. al. (1995). See `optim`.

Usage

```
nlm.control(fnscale=-1, maxit=1000, reltol, factr=1e7, pgtol=0,
            s0.tol=1e-5, bscale=5)
```

Arguments

fnscale	An overall scaling to be applied to the profile likelihood function (<code>profileLik</code>) during optimization. If positive, turns the problem into a minimization problem. Optimization is performed on <code>'profileLik(par)/fnscale'</code> . Default is -1.
maxit	The maximum number of iterations. Default is 1000.
reltol	Relative convergence tolerance. The algorithm stops if it is unable to reduce the value by a factor of <code>'reltol * (abs(val) + reltol)'</code> at a step. Default is <code>sqrt(.Machine\$double.eps)</code> , typically about 1e-8.
factr	Controls the convergence of the "L-BFGS-B" method. Convergence occurs when the reduction in the objective is within this factor of the machine tolerance. Default is 1e7, that is a tolerance of about 1e-8.
pgtol	Helps control the convergence of the "L-BFGS-B" method. It is a tolerance on the projected gradient in the current search direction. Default is 0.
s0.tol	Convergence tolerance of baseline hazard self-consistency equation. Default is 1e-5.
bscale	The maximum profile likelihood estimator is obtained by maximizing the profile likelihood over a region determined by the magnitude of the observed covariates. These constraints are imposed in order to avoid numerical problems in the calculation of the profile likelihood function. For a given regression parameter, corresponding to a covariate with observed values x , the upper bound is the <code>bscale</code> of the parameter divided by <code>max(abs(x))</code> if <code>max(abs(x)) > 1e-10</code> , otherwise 1e-10. The lower bound is minus the upper bound.

Different values of `bscale` are allowed for different parameters. If different values of `bscale` are provided, the vector needs to have the scale for the regression parameters of the long-term predictor, β_θ , followed by the scale for the short-term predictor, β_η , if available and the one for the cure parameter, β_c , last if a cure model.

Default is 5.

Value

A list with the same elements as the input.

References

Byrd, R. H., Lu, P., Nocedal, J. and Zhu, C. (1995) A limited memory algorithm for bound constrained optimization. *SIAM J. Scientific Computing*, 16, 1190-1208.

See Also

`optim`, `nlm`.

nlm.object

Non-Linear Transformation Model Object

Description

This object is returned by the `nlm` function. It represents a fitted non-linear transformation model. Objects of this class have methods for the functions `print` and `summary`.

Components

Components of a `nlm` object.

coefficients The maximum profile likelihood estimators of the model regression parameters, β_θ (and β_η and β_c). See `nlm`.

loglik A vector of length 2 containing the log-likelihood of the null model and with the final values of the coefficients.

surv MLE of the baseline survival function at the profile maximum likelihood parameters. It is obtained from the hazard jumps that satisfy self-consistency equation (5) in Tsodikov A. and Garibotti G. (2006).

var The variance matrix of the coefficients.

n The number of observations used in the fit.

maxit The maximum number of iterations of the optimization procedure. Default is 1000.

counts Number of calls to the profile likelihood function during the optimization process. This excludes those calls needed to compute a finite-difference approximation to the gradient.

- convergence** An integer code. '0' indicates successful convergence. Error codes are '1' indicates that the iteration limit 'maxit' had been reached. '51' indicates a warning from the optimization method; see component 'message' for further details. '52' indicates an error from the "L-BFGS-B" method; see component 'message' for further details.
- message** A character string giving any additional information returned by the optimizer, or 'NULL'.
- formula** If the model has only one predictor a single formula indicating the model for that predictor. If the model has two predictors, then formula is a list with terms pred.long and pred.short indicating the model for the long and short term predictor respectively.
- terms** The 'terms' object used.
- nvar** If the model has one predictor, number of columns in the model matrix. If it has 2 predictors list with terms pred.short and pred.long indicating the number of columns in the model matrix of the long and short term predictors respectively.
- x** If requested is the model matrix in the case of nlt models with one predictor. If the nltm has 2 predictors, x is a list with terms pred.long and pred short containing the model matrix for the long and short-term predictors respectively.
- y** If requested, the response matrix.
- model** If requested is the model frame in the case of nlt models with one predictor. If the nltm has 2 predictors, is a list with terms pred.long and pred short containing the model frame for the long and short-term predictors respectively.
- npred** Number of predictors.
- call** The call of the nltm model.
- na.action** The na.action attribute, if any, that was returned by the na.action routine. The object will also contain the following, for documentation see the lm object: terms, assign, formula, call, and, optionally, x, y, and/or frame.

References

Tsodikov A., Garibotti G. (2006) "Profile information matrix for nonlinear transformation models". to appear in Journal of Lifetime Data Analysis.

See Also

[nltm](#), [summary.nlm](#).

summary.nlm

Summary of a nltm objects.

Description

This function finds confidence intervals for relative risks and calculates the log-likelihood test. It is a method for the generic function summary of class nltm. It can be invoked by calling summary for an object of class nltm.

Usage

```
## S3 method for class 'nlm'  
summary(object, coef = TRUE, conf.int = 0.95,  
        digits = max(options()$digits - 4, 3),...)
```

Arguments

object	Fitted model object of class nlm. This is assumed to be the result of some function that produces an object with the same named components as that returned by the nlm function.
coef	If true it gives a table with coefficients, standard errors and p-values. Default is TRUE.
conf.int	Confidence level of confidence intervals for relative risks. If 0 confidence intervals are not computed. Default is 0.95.
digits	Number of digits used for formatting output.
...	Arguments to be passed to or from other methods.

See Also

[nlm](#), [nlm.object](#), [summary](#).

Examples

```
# fit a Proportional Odds Model  
data(melanoma, package="nlm")  
fit <- nlm(Surv(time,status) ~ size + age, data=melanoma, nlm.model="PO")  
summary(fit)
```

Index

* **datasets**

melanoma, 3

* **survival**

anova.nlm, 2

nlm, 3

nlm.control, 6

nlm.object, 7

summary.nlm, 8

anova.nlm, 2

anova.nlm\$ (anova.nlm), 2

melanoma, 3

nlm, 2, 3, 6–9

nlm.control, 5, 6

nlm.object, 5, 7, 9

print.nlm(nlm.object), 7

summary, 9

summary.nlm, 5, 8, 8

update.nlm.formula(anova.nlm), 2