

# Package ‘crch’

October 18, 2016

**Title** Censored Regression with Conditional Heteroscedasticity

**Version** 1.0-0

**Date** 2016-10-17

**Depends** R (>= 2.10.0)

**Imports** stats, Formula, ordinal

**Suggests** glmx, lmtest, memisc

**Description** Different approaches to censored or truncated regression with conditional heteroscedasticity are provided. First, continuous distributions can be used for the (right and/or left censored or truncated) response with separate linear predictors for the mean and variance. Second, cumulative link models for ordinal data (obtained by interval-censoring continuous data) can be employed for heteroscedastic extended logistic regression (HXLR). In the latter type of models, the intercepts depend on the thresholds that define the intervals.

**License** GPL-2 | GPL-3

**NeedsCompilation** yes

**Author** Jakob Messner [aut, cre],  
Achim Zeileis [aut]

**Maintainer** Jakob Messner <jwmm@elektro.dtu.dk>

**Repository** CRAN

**Date/Publication** 2016-10-18 10:37:58

## R topics documented:

clogis	2
cnorm	3
coef.crch	4
coef.crch.boost	5
coef.hxlr	6
crch	7
crch.boost	10
crch.control	13

ct	14
hxr	15
hxr.control	19
plot.crch.boost	19
predict.crch	20
predict.crch.boost	21
predict.hxr	22
RainIbk	23
tlogis	24
tnorm	26
tt	27

## Index 29

---

clogis *The Censored Logistic Distribution*

---

### Description

Density, distribution function, quantile function, and random generation for the left and/or right censored logistic distribution.

### Usage

```
dlogis(x, mean = 0, sd = 1, left = -Inf, right = Inf, log = FALSE)
```

```
pclogis(q, mean = 0, sd = 1, lower.tail = TRUE, log.p = FALSE, left = -Inf,
        right = Inf)
```

```
rclogis(n, mean = 0, sd = 1, left = -Inf, right = Inf)
```

```
qclogis(p, mean = 0, sd = 1, lower.tail = TRUE, log.p = FALSE, left = -Inf,
        right = Inf)
```

### Arguments

x, q	vector of quantiles.
p	vector of probabilities.
n	number of observations. If $\text{length}(n) > 1$ , the length is taken to be the number required.
mean	location parameter.
sd	scale parameter.
log, log.p	logical; if TRUE, probabilities p are given as $\log(p)$ .
lower.tail	logical; if TRUE (default), probabilities are $P[X \leq x]$ otherwise, $P[X > x]$ .
left	left censoring point.
right	right censoring point.

**Details**

If mean or sd are not specified they assume the default values of 0 and 1, respectively. left and right have the defaults -Inf and Inf respectively.

The censored logistic distribution has density  $f(x)$ :

$$\begin{aligned} \Lambda((left - \mu)/\sigma) & \quad \text{if } x \leq left \\ 1 - \Lambda((right - \mu)/\sigma) & \quad \text{if } x \geq right \\ \lambda((x - \mu)/\sigma)/\sigma & \quad \text{if } left < x < right \end{aligned}$$

where  $\Lambda$  and  $\lambda$  are the cumulative distribution function and probability density function of the standard logistic distribution respectively,  $\mu$  is the location of the distribution, and  $\sigma$  the scale.

**Value**

dclogis gives the density, pclogis gives the distribution function, qclogis gives the quantile function, and rclogis generates random deviates.

**See Also**

[dlogis](#)

---

 cnorm

---

*The Censored Normal Distribution*


---

**Description**

Density, distribution function, quantile function, and random generation for the left and/or right censored normal distribution.

**Usage**

```
dcnorm(x, mean = 0, sd = 1, left = -Inf, right = Inf, log = FALSE)
```

```
pcnorm(q, mean = 0, sd = 1, lower.tail = TRUE, log.p = FALSE, left = -Inf,
       right = Inf)
```

```
rcnorm(n, mean = 0, sd = 1, left = -Inf, right = Inf)
```

```
qcnorm(p, mean = 0, sd = 1, lower.tail = TRUE, log.p = FALSE, left = -Inf,
       right = Inf)
```

**Arguments**

x, q            vector of quantiles.  
p                vector of probabilities.

n	number of observations. If <code>length(n) &gt; 1</code> , the length is taken to be the number required.
mean	vector of means.
sd	vector of standard deviations.
log, log.p	logical; if TRUE, probabilities p are given as <code>log(p)</code> .
lower.tail	logical; if TRUE (default), probabilities are <code>P[X &lt;= x]</code> otherwise, <code>P[X &gt; x]</code> .
left	left censoring point.
right	right censoring point.

### Details

If mean or sd are not specified they assume the default values of 0 and 1, respectively. left and right have the defaults -Inf and Inf respectively.

The censored normal distribution has density  $f(x)$ :

$$\begin{aligned} &\Phi((left - \mu)/\sigma) && \text{if } x \leq left \\ &1 - \Phi((right - \mu)/\sigma) && \text{if } x \geq right \\ &\phi((x - \mu)/\sigma)/\sigma && \text{if } left < x < right \end{aligned}$$

where  $\Phi$  and  $\phi$  are the cumulative distribution function and probability density function of the standard normal distribution respectively,  $\mu$  is the mean of the distribution, and  $\sigma$  the standard deviation.

### Value

`dcnorm` gives the density, `pcnorm` gives the distribution function, `qcnorm` gives the quantile function, and `rcnorm` generates random deviates.

### See Also

[dnorm](#)

---

coef.crch

*Methods for CRCH Objects*

---

### Description

Methods for extracting information from fitted crch objects.

### Usage

```
## S3 method for class 'crch'
coef(object, model = c("full", "location", "scale", "df"), ...)
## S3 method for class 'crch'
vcov(object, model = c("full", "location", "scale", "df"), ...)
```

```
## S3 method for class 'crch'
terms(x, model = c("location", "scale", "full"), ...)
## S3 method for class 'crch'
fitted(object, type = c("location", "scale"), ...)
```

### Arguments

object, x      an object of class "crch".

model          model for which coefficients shall be returned.

type          type of fitted values.

...          further arguments passed to or from other methods.

### Details

In addition to the methods above, a set of standard extractor functions for "crch" objects is available, including methods to the generic functions [print](#), [summary](#), [logLik](#), and [residuals](#).

### See Also

[crch](#)

---

coef.crch.boost      *Methods for boosted CRCH Objects*

---

### Description

Methods for extracting information from fitted crch.boost objects.

### Usage

```
## S3 method for class 'crch.boost'
coef(object, model = c("full", "location", "scale", "df"),
      mstop = NULL, zero.coefficients = FALSE, standardize = FALSE, ...)
## S3 method for class 'crch.boost'
print(x, digits = max(3, getOption("digits") - 3),
      mstop = NULL, zero.coefficients = FALSE, ...)
## S3 method for class 'crch.boost'
summary(object, mstop = NULL, zero.coefficients = FALSE, ...)
## S3 method for class 'crch.boost'
logLik(object, mstop = NULL, ...)
```

**Arguments**

object, x	an object of class "crch.boost".
model	model for which coefficients shall be returned.
mstop	stopping iteration for which coefficients shall be returned. Can be either a character ("max", "aic", "bic", "cv") or a numeric value.
zero.coefficients	logical whether zero coefficients are returned.
standardize	logical whether coefficients shall be standardized.
digits	the number of significant digits to use when printing.
...	further arguments passed to or from other methods.

**Details**

In addition to the methods above, the "crch" methods [terms](#), [model.frame](#), [model.matrix](#), [residuals](#), and [fitted](#) can be used also for "crch.boost" objects .

**See Also**

[crch.boost](#), [coef.crch](#)

---

coef.hxlr

*Methods for HXLR Objects*

---

**Description**

Methods for extracting information from fitted hxlr objects.

**Usage**

```
## S3 method for class 'hxlr'
coef(object, model = c("full", "intercept", "location", "scale"),
      type = c("CLM", "latent"), ...)
## S3 method for class 'hxlr'
vcov(object, model = c("full", "intercept", "location", "scale"),
      type = c("CLM", "latent"), ...)
## S3 method for class 'hxlr'
terms(x, model = c("full", "location", "scale"), ...)
```

**Arguments**

object, x	an object of class "hxlr".
model	model for which coefficients shall be returned.
type	type of coefficients. Default are CLM type coefficients. For type "latent" coefficients are converted in coefficients for location and scale of the latent distribution (analog to crch models).
...	further arguments passed to or from other methods.

**Details**

In addition to the methods above, a set of standard extractor functions for "hxlr" objects is available, including methods to the generic functions [print](#), [summary](#), and [logLik](#).

**See Also**

[hxlr](#)

---

 crch

*Censored Regression with Conditional Heteroscedasticity*


---

**Description**

Fitting censored (tobit) or truncated regression models with conditional heteroscedasticity.

**Usage**

```
crch(formula, data, subset, na.action, weights, offset,
      link.scale = c("log", "identity", "quadratic"),
      dist = c("gaussian", "logistic", "student"), df = NULL,
      left = -Inf, right = Inf, truncated = FALSE,
      control = crch.control(...), model = TRUE, x = FALSE,
      y = FALSE, ...)
```

```
trch(formula, data, subset, na.action, weights, offset,
      link.scale = c("log", "identity", "quadratic"),
      dist = c("gaussian", "logistic", "student"), df = NULL,
      left = -Inf, right = Inf, truncated = TRUE,
      control = crch.control(...), model = TRUE, x = FALSE,
      y = FALSE, ...)
```

```
crch.fit(x, z, y, left, right, truncated = FALSE, dist = "gaussian",
         df = NULL, link.scale = "log", weights = NULL, offset = NULL,
         control = crch.control())
```

**Arguments**

formula	a formula expression of the form $y \sim x \mid z$ where $y$ is the response and $x$ and $z$ are regressor variables for the location and the scale of the fitted distribution respectively.
data	an optional data frame containing the variables occurring in the formulas.
subset	an optional vector specifying a subset of observations to be used for fitting.
na.action	a function which indicates what should happen when the data contain NAs.
weights	optional case weights in fitting.

offset	optional numeric vector with <i>a priori</i> known component to be included in the linear predictor for the location. For <code>crch.fit</code> , <code>offset</code> can also be a list of 2 offsets used for the location and scale respectively.
link.scale	character specification of the link function in the scale model. Currently, "identity", "log", "quadratic" are supported. The default is "log". Alternatively, an object of class "link-glm" can be supplied.
dist	assumed distribution for the dependent variable <i>y</i> .
df	optional degrees of freedom for <code>dist="student"</code> . If omitted the degrees of freedom are estimated.
left	left limit for the censored dependent variable <i>y</i> . If set to <code>-Inf</code> , <i>y</i> is assumed not to be left-censored.
right	right limit for the censored dependent variable <i>y</i> . If set to <code>Inf</code> , <i>y</i> is assumed not to be right-censored.
truncated	logical. If TRUE truncated model is fitted with <code>left</code> and <code>right</code> interpreted as truncation points, If FALSE censored model is fitted. Default is FALSE
control	a list of control parameters passed to <code>optim</code> or to the internal boosting algorithm if <code>control=crch.boost()</code> . Default is <code>crch.control()</code> .
model	logical. If TRUE <i>model frame</i> is included as a component of the returned value.
x, y	for <code>crch</code> : logical. If TRUE the model matrix and response vector used for fitting are returned as components of the returned value. for <code>crch.fit</code> : <i>x</i> is a design matrix with regressors for the location and <i>y</i> is a vector of observations.
z	a design matrix with regressors for the scale.
...	arguments to be used to form the default <code>control</code> argument if it is not supplied directly.

## Details

`crch` fits censored (tobit) or truncated regression models with conditional heteroscedasticity with maximum likelihood estimation. Student-t, Gaussian, and logistic distributions can be fitted to left-and/or right censored or truncated responses. Different regressors can be used to model the location and the scale of this distribution. If `control=crch.boost()` optimization is performed by boosting.

`trch` is a wrapper function for `crch` with default `truncated = TRUE`.

`crch.fit` is the lower level function where the actual fitting takes place.

## Value

An object of class "crch" or "crch.boost", i.e., a list with the following elements.

coefficients	list of coefficients for location, scale, and df. Scale and df coefficients are in log-scale.
df	if <code>dist = "student"</code> : degrees of freedom of student-t distribution. else NULL.
residuals	the residuals, that is response minus fitted values.
fitted.values	list of fitted location and scale parameters.
dist	assumed distribution for the dependent variable <i>y</i> .



cens	list of censoring points.
optim	output from optimization from <a href="#">optim</a> .
method	optimization method used for <a href="#">optim</a>
control	list of control parameters passed to <a href="#">optim</a>
start	starting values of coefficients used in the optimization.
weights	case weights used for fitting.
offset	list of offsets for location and scale.
n	number of observations.
nobs	number of observations with non-zero weights.
loglik	log-likelihood.
vcov	covariance matrix.
link	a list with element "scale" containing the link objects for the scale model.
truncated	logical indicating wheter a truncated model has been fitted.
converged	logical variable whether optimization has converged or not.
iterations	number of iterations in optimization.
call	function call.
formula	the formula supplied.
terms	the terms objects used.
levels	list of levels of the factors used in fitting for location and scale respectively.
contrasts	(where relevant) the contrasts used.
y	if requested, the response used.
x	if requested, the model matrix used.
model	if requested, the model frame used.
stepsize, mstop, mstopopt, standardize	return values of boosting optimization. See <a href="#">crch.boost</a> for details.

## References

Messner JW, Mayr GJ, Zeileis A (2016). Heteroscedastic Censored and Truncated Regression with crch. *The R Journal*, **3**(1), 173–181. <https://journal.R-project.org/archive/2016-1/messner-mayr-zeileis.pdf>.

Messner JW, Zeileis A, Broecker J, Mayr GJ (2014). Probabilistic Wind Power Forecasts with an Inverse Power Curve Transformation and Censored Regression. *Wind Energy*, **17**(11), 1753–1766. doi: 10.1002/we.1666.

## See Also

[predict.crch](#), [crch.control](#), [crch.boost](#)

**Examples**

```

data("RainIbk")
## mean and standard deviation of square root transformed ensemble forecasts
RainIbk$sqrtensmean <-
  apply(sqrt(RainIbk[,grep('^rainfc',names(RainIbk))]), 1, mean)
RainIbk$sqrtenssd <-
  apply(sqrt(RainIbk[,grep('^rainfc',names(RainIbk))]), 1, sd)

## fit linear regression model with Gaussian distribution
CRCH <- crch(sqrt(rain) ~ sqrtensmean, data = RainIbk, dist = "gaussian")
## same as lm?
all.equal(coef(lm(sqrt(rain) ~ sqrtensmean, data = RainIbk)),
  head(coef(CRCH), -1))

## print
CRCH
## summary
summary(CRCH)

## left censored regression model with censoring point 0:
CRCH2 <- crch(sqrt(rain) ~ sqrtensmean, data = RainIbk,
  dist = "gaussian", left = 0)

## left censored regression model with censoring point 0 and
## conditional heteroscedasticity:
CRCH3 <- crch(sqrt(rain) ~ sqrtensmean|sqrtenssd, data = RainIbk,
  dist = "gaussian", left = 0)

## left censored regression model with censoring point 0 and
## conditional heteroscedasticity with logistic distribution:
CRCH4 <- crch(sqrt(rain) ~ sqrtensmean|sqrtenssd, data = RainIbk,
  dist = "logistic", left = 0)

## compare AIC
AIC(CRCH, CRCH2, CRCH3, CRCH4)

```

---

crch.boost

*Auxiliary functions to fit crch models via boosting.*


---

**Description**

Auxiliary functions to fit crch models via boosting

**Usage**

```

crch.boost(maxit = 100, nu = 0.1, start = NULL, dot = "separate",
  mstop = c("max", "aic", "bic", "cv"), nfolds = 10, foldid = NULL)

crch.boost.fit(x, z, y, left, right, truncated = FALSE, dist = "gaussian",

```

```
df = NULL, link.scale = "log", weights = NULL, offset = NULL,
control = crch.boost())
```

### Arguments

maxit	the maximum number of boosting iterations.
nu	boosting step size. Default is 0.1.
start	a previously boosted but not converged "crch.boost" object to continue.
dot	character specifying how to process formula parts with a dot (.) on the right-hand side. This can either be "separate" so that each formula part is expanded separately or "sequential" so that the parts are expanded sequentially conditional on all prior parts. Default is "separate"
mstop	method to find optimum stopping iteration. Default is "max" which is maxit. Alternatives are "aic" and "bic" for AIC and BIC optimization and "cv" for cross validation. mstop can also be a positive integer to set the number of boosting iterations. Then maxit is set to mstop and mstop="max".
nfolds	if mstopopt = "cv", number of folds in cross validation.
foldid	if mstopopt = "cv", an optional vector of values between 1 and nfold identifying the fold each observation is in. If supplied, nfolds can be missing.
x, z, y, left, right, truncated, dist, df, link.scale, weights, offset, control	see <a href="#">crch.fit</a> for details.

### Details

crch.boost extends crch to fit censored (tobit) or truncated regression models with conditional heteroscedasticity by boosting. If crch.boost() is supplied as control in crch then crch.boost.fit is used as lower level fitting function. Note that [crch.control\(\)](#) with method=boosting is equivalent to crch.boost(). Thus, boosting can more conveniently be called with crch(..., method = "boosting").

### Value

For crch.boost: A list with components named as the arguments. For crch.boost.fit: An object of class "crch.boost", i.e., a list with the following elements.

coefficients	list of coefficients for location and scale. Scale coefficients are in log-scale. Coefficients are of optimum stopping iteration specified by mstop.
df	if dist = "student": degrees of freedom of student-t distribution. else NULL.
residuals	the residuals, that is response minus fitted values.
fitted.values	list of fitted location and scale parameters at optimum stopping iteration specified by mstop.
dist	assumed distribution for the dependent variable y.
cens	list of censoring points.
control	list of control parameters.
weights	case weights used for fitting.
offset	list of offsets for location and scale.

n	number of observations.
nobs	number of observations with non-zero weights.
loglik	log-likelihood.
link	a list with element "scale" containing the link objects for the scale model.
truncated	logical indicating wheter a truncated model has been fitted.
iterations	number of boosting iterations.
stepsize	boosting stepsize nu.
mstop	criterion used to find optimum stopping iteration.
mstopopt	optimum stopping iterations for different criteria.
standardize	list of center and scale values used to standardize response and regressors.

## References

Messner JW, Mayr GJ, Zeileis A (2016). Non-Homogeneous Boosting for Predictor Selection in Ensemble Post-Processing. *Working Papers*, Faculty of Economics and Statistics, University of Innsbruck, url:<http://econpapers.repec.org/paper/innwpaper/2016-04.htm>.

## See Also

[crch](#), [crch.control](#)

## Examples

```
# generate data
set.seed(5)
x <- matrix(rnorm(1000*20),1000,20)
y <- rnorm(1000, 1 + x[,1] - 1.5 * x[,2], exp(-1 + 0.3*x[,3]))
y <- pmax(0, y)
data <- data.frame(cbind(y, x))

# fit model with maximum likelihood
CRCH <- crch(y ~ .|., data = data, dist = "gaussian", left = 0)

# fit model with boosting
boost <- crch(y ~ .|., data = data, dist = "gaussian", left = 0,
  control = crch.boost(mstop = "aic"))

# more conveniently, the same model can also be fit through
# boost <- crch(y ~ .|., data = data, dist = "gaussian", left = 0,
#   method = "boosting", mstop = "aic")

# AIC comparison
AIC(CRCH, boost)

# summary
summary(boost)

# plot
plot(boost)
```

---

`crch.control`*Auxiliary Function for Controlling crch Fitting*

---

**Description**

Auxiliary function for crch fitting. Specifies a list of values passed to [optim](#).

**Usage**

```
crch.control(method = "BFGS", maxit = NULL, hessian = NULL,  
            trace = FALSE, start = NULL, dot = "separate",  
            lower = -Inf, upper = Inf, ...)
```

**Arguments**

<code>method</code>	optimization method passed to <a href="#">optim</a>
<code>maxit</code>	the maximum number of iterations. Default is 5000 except for method="boosting" where the default is 100.
<code>hessian</code>	logical or NULL. If TRUE the numerical Hessian matrix from the <a href="#">optim</a> output is used for estimation of the covariance matrix. If FALSE no covariance matrix is computed. If NULL (the default) the Hessian matrix is computed analytically for <code>dist="gaussian"</code> , <code>dist="logistic"</code> , and <code>dist="student"</code> with predefined <code>df</code> . For <code>dist="student"</code> without prespecified <code>df</code> , no analytical solution is available and a numerical Hessian matrix is forced.
<code>trace</code>	non-negative integer. If positive, tracing information on the progress of the optimization is produced.
<code>start</code>	initial values for the parameters to be optimized over.
<code>dot</code>	character specifying how to process formula parts with a dot (.) on the right-hand side. This can either be "separate" so that each formula part is expanded separately or "sequential" so that the parts are expanded sequentially conditional on all prior parts. Default is "separate"
<code>lower, upper</code>	bounds on the variables for the "L-BFGS-B" method, or bounds in which to search for method "Brent".
<code>...</code>	additional parameters passed to <a href="#">optim</a> .

**Value**

A list with components named as the arguments.

**See Also**

[crch](#), [optim](#)

ct

*The Censored Student-t Distribution***Description**

Density, distribution function, quantile function, and random generation for the left and/or right censored student-t distribution with *df* degrees of freedom.

**Usage**

```
dct(x, mean = 0, sd = 1, df, left = -Inf, right = Inf, log = FALSE)
```

```
pct(q, mean = 0, sd = 1, df, lower.tail = TRUE, log.p = FALSE, left = -Inf,
    right = Inf)
```

```
rct(n, mean = 0, sd = 1, df, left = -Inf, right = Inf)
```

```
qct(p, mean = 0, sd = 1, df, lower.tail = TRUE, log.p = FALSE, left = -Inf,
    right = Inf)
```

**Arguments**

<i>x</i> , <i>q</i>	vector of quantiles.
<i>p</i>	vector of probabilities.
<i>n</i>	number of observations. If <code>length(n) &gt; 1</code> , the length is taken to be the number required.
<i>mean</i>	location parameter.
<i>sd</i>	scale parameter.
<i>df</i>	degrees of freedom ( $> 0$ , maybe non-integer). <code>df = Inf</code> is allowed.
<i>log</i> , <i>log.p</i>	logical; if TRUE, probabilities <i>p</i> are given as $\log(p)$ .
<i>lower.tail</i>	logical; if TRUE (default), probabilities are $P[X \leq x]$ otherwise, $P[X > x]$ .
<i>left</i>	left censoring point.
<i>right</i>	right censoring point.

**Details**

If *mean* or *sd* are not specified they assume the default values of 0 and 1, respectively. *left* and *right* have the defaults -Inf and Inf respectively.

The censored student-t distribution has density  $f(x)$ :

$$\begin{array}{ll} T((left - \mu)/\sigma) & \text{if } x \leq left \\ 1 - T((right - \mu)/\sigma) & \text{if } x \geq right \\ \tau((x - \mu)/\sigma)/\sigma & \text{if } left < x < right \end{array}$$

where  $T$  and  $\tau$  are the cumulative distribution function and probability density function of the student-t distribution with  $df$  degrees of freedom respectively,  $\mu$  is the location of the distribution, and  $\sigma$  the scale.

### Value

`dct` gives the density, `pct` gives the distribution function, `qct` gives the quantile function, and `rct` generates random deviates.

### See Also

[dt](#)

---

hxlr

*Heteroscedastic Extended Logistic Regression*

---

### Description

This is a wrapper function for [clm](#) (from package **ordinal**) to fit (heteroscedastic) extended logistic regression (HXLr) models (Messner et al. 2013).

### Usage

```
hxlr(formula, data, subset, na.action, weights, thresholds, link, control, ...)
```

### Arguments

<code>formula</code>	a formula expression of the form $y \sim x \mid z$ where $y$ is the response and $x$ and $z$ are regressor variables for the location and the scale of the latent distribution respectively. Response can either be a continuous variable or a factor.
<code>data</code>	an optional data frame containing the variables occurring in the formulas.
<code>subset</code>	an optional vector specifying a subset of observations to be used for fitting.
<code>na.action</code>	a function which indicates what should happen when the data contain NAs. Default is <code>na.omit</code>
<code>weights</code>	optional case weights in fitting.
<code>thresholds</code>	vector of (transformed) thresholds that are used to cut the continuous response into categories. Data frames or matrices with multiple columns are allowed as well. Then each column is used as separate predictor variable for the intercept model.
<code>link</code>	link function, i.e., the type of location-scale distribution assumed for the latent distribution. Default is <code>logit</code> .
<code>control</code>	a list of control parameters passed to <a href="#">optim</a> . Default is <code>hxlr.control</code>
<code>...</code>	arguments to be used to form the default <code>control</code> argument if it is not supplied directly.

## Details

Extended logistic regression (Wilks 2009) extends binary logistic regression to multi-category responses by including the thresholds, that are used to cut a continuous variable into categories, in the regression equation. Heteroscedastic extended logistic regression (Messner et al. 2013) extends this model further and allows to add additional predictor variables that are used to predict the scale of the latent logistic distribution.

## Value

An object of class "hxr", i.e., a list with the following elements.

<code>coefficients</code>	list of CLM coefficients for intercept, location, and scale model.
<code>fitted.values</code>	list of fitted location and scale parameters.
<code>optim</code>	output from optimization from <a href="#">optim</a> .
<code>method</code>	Optimization method used for <a href="#">optim</a> .
<code>control</code>	list of control parameters passed to <a href="#">optim</a>
<code>start</code>	starting values of coefficients used in the optimization.
<code>weights</code>	case weights used for fitting.
<code>n</code>	number of observations.
<code>nobs</code>	number of observations with non-zero weights.
<code>loglik</code>	log-likelihood.
<code>vcov</code>	covariance matrix.
<code>converged</code>	logical variable whether optimization has converged or not.
<code>iterations</code>	number of iterations in optimization.
<code>call</code>	function call.
<code>scale</code>	the formula supplied.
<code>terms</code>	the terms objects used.
<code>levels</code>	list of levels of the factors used in fitting for location and scale respectively.
<code>thresholds</code>	the thresholds supplied.

## References

Messner JW, Mayr GJ, Zeileis A, Wilks DS (2014). Extending Extended Logistic Regression to Effectively Utilize the Ensemble Spread. *Monthly Weather Review*, **142**, 448–456. doi: [10.1175/MWR-D-13-00271.1](https://doi.org/10.1175/MWR-D-13-00271.1).

Wilks DS (2009). Extending Logistic Regression to Provide Full-Probability-Distribution MOS Forecasts. *Meteorological Applications*, **368**, 361–368.

## See Also

[predict.hxr](#), [clm](#)



**Examples**

```

data("RainIbk")
## mean and standard deviation of square root transformed ensemble forecasts
RainIbk$sqrtensmean <-
  apply(sqrt(RainIbk[,grep('^rainfc',names(RainIbk))]), 1, mean)
RainIbk$sqrtenssd <-
  apply(sqrt(RainIbk[,grep('^rainfc',names(RainIbk))]), 1, sd)

## climatological deciles
q <- unique(quantile(RainIbk$rain, seq(0.1, 0.9, 0.1)))

## fit ordinary extended logistic regression with ensemble mean as
## predictor variable
XLR <- hxlr(sqrt(rain) ~ sqrtensmean, data = RainIbk, thresholds = sqrt(q))
## print
XLR
## summary
summary(XLR)

## fit ordinary extended logistic regression with ensemble mean
## and standard deviation as predictor variables
XLRs <- hxlr(sqrt(rain) ~ sqrtensmean + sqrtenssd, data = RainIbk,
  thresholds = sqrt(q))
## fit heteroscedastic extended logistic regression with ensemble
## standard deviation as predictor for the scale
HXLR <- hxlr(sqrt(rain) ~ sqrtensmean | sqrtenssd, data = RainIbk,
  thresholds = sqrt(q))

## compare AIC of different models
AIC(XLR, XLRs, HXLR)

## XLRs and HXLR are nested in XLR -> likelihood-ratio-tests
if(require("lmtest")) {
  lrtest(XLR, XLRs)
  lrtest(XLR, HXLR)
}

## Not run:
#####
## Cross-validation and bootstrapping RPS for different models
## (like in Messner 2013).
N <- NROW(RainIbk)
## function that returns model fits
fits <- function(data, weights = rep(1, N)) {
  list(
    "XLR" = hxlr(sqrt(rain) ~ sqrtensmean, data = data,
      weights = weights, thresholds = sqrt(q)),
    "XLR:S" = hxlr(sqrt(rain) ~ sqrtensmean + sqrtenssd, data = data,
      weights = weights, thresholds = sqrt(q)),
    "XLR:SM" = hxlr(sqrt(rain) ~ sqrtensmean + I(sqrtensmean*sqrtenssd),
      data = data, weights = weights, thresholds = sqrt(q)),
  )
}

```

```

    "HXLr" = hxlr(sqrt(rain) ~ sqrtensmean | sqrtenssd, data = data,
      weights = weights, thresholds = sqrt(q)),
    "HXLr:S" = hxlr(sqrt(rain) ~ sqrtensmean + sqrtenssd | sqrtenssd,
      data = data, weights = weights, thresholds = sqrt(q))
  )
}

## cross validation
id <- sample(1:10, N, replace = TRUE)
obs <- NULL
pred <- list(NULL)
for(i in 1:10) {
  ## splitting into test and training data set
  trainIndex <- which(id != i)
  testIndex <- which(id == i)
  ## weights that are used for fitting the models
  weights <- as.numeric(table(factor(trainIndex, levels = c(1:N))))
  ## testdata
  testdata <- RainIbk[testIndex,]
  ## observations
  obs <- c(obs, RainIbk$rain[testIndex])
  ## estimation
  modelfits <- fits(RainIbk, weights)
  ## Prediction
  pred2 <- lapply(modelfits, predict, newdata = testdata, type = "cumprob")
  pred <- mapply(rbind, pred, pred2, SIMPLIFY = FALSE)
}
names(pred) <- c(names(modelfits))

## function to compute RPS
rps <- function(pred, obs) {
  OBS <- NULL
  for(i in 1:N)
    OBS <- rbind(OBS, rep(0:1, c(obs[i] - 1, length(q) - obs[i] + 1)))
  apply((OBS-pred)^2, 1, sum)
}

## compute rps
RPS <- lapply(pred, rps, obs = as.numeric(cut(obs, c(-Inf, q, Inf))))

## bootstrapping mean rps
rpsall <- NULL
for(i in 1:250) {
  index <- sample(length(obs), replace = TRUE)
  rpsall <- rbind(rpsall, sapply(RPS, function(x) mean(x[index])))
}

rpssall <- 1 - rpsall/rpsall[,1]
boxplot(rpssall[,-1], ylab = "RPSS", main = "RPSS relative to XLR")
abline(h = 0, lty = 2)

## End(Not run)

```

---

hxr.control	<i>Auxiliary Function for Controlling HXLR Fitting</i>
-------------	--

---

**Description**

Auxiliary function for hxr fitting. Specifies a list of values passed to [optim](#).

**Usage**

```
hxr.control(method = "BFGS", maxit = 5000, hessian = TRUE,
            trace = FALSE, start = NULL, ...)
```

**Arguments**

method	optimization method used in <a href="#">optim</a>
maxit	the maximum number of iterations.
hessian	logical. Should a numerically differentiated Hessian matrix be returned?
trace	non-negative integer. If positive, tracing information on the progress of the optimization is produced.
start	initial values for the parameters to be optimized over.
...	Additional parameters passed to <a href="#">optim</a> .

**Value**

A list with components named as the arguments.

**See Also**

[hxr](#), [optim](#)

---

plot.crch.boost	<i>Plot coefficient paths of boosted CRCH objects.</i>
-----------------	--

---

**Description**

Plot paths of coefficients or log-likelihood contributions for crch.boost models.

**Usage**

```
## S3 method for class 'crch.boost'
plot(x, loglik = FALSE,
     standardize = TRUE, which = c("both", "location", "scale"),
     mstop = NULL, coef.label = TRUE, col = NULL, ...)
```

**Arguments**

x	an object of class "crch.boost".
loglik	logical whether log-likelihood contribution shall be plotted instead of coefficient value.
standardize	logical whether coefficients shall be standardized. Not used if loglik = TRUE
which	which coefficients: "location" and "scale" plots only the coefficients for the location and scale part of the model respectively. "both" plots the coefficient paths of both parts in one graph.
mstop	Stopping iteration for which a vertical line is plotted. Possible choices are "max", "aic", "bic", "cv", "all", or "no". Default is the stopping iteration used for fitting.
coef.label	logical whether paths shall be labeled.
col	Color(s) for the paths. If which="both" a vector of two colors where the paths for the location are plotted in the first color and for the scale in the second color.
...	further arguments passed to <a href="#">plot.ts</a> .

**See Also**

[crch.boost,plot.ts](#)

---

predict.crch

*Predicted/Fitted Values for CRCH Fits*

---

**Description**

Obtains various types of predictions for crch models.

**Usage**

```
## S3 method for class 'crch'
predict(object, newdata = NULL, type = c("response", "location",
    "scale", "quantile"), na.action = na.pass, at = 0.5, left = NULL, right = NULL, ...)
```

**Arguments**

object	an object of class "crch".
newdata	an optional data frame in which to look for variables with which to predict.
type	type of prediction: "response" is equivalent to "location" and returns the location of the predicted distribution. "scale" returns the scale of the predicted distribution and "quantile" returns a matrix of predicted quantiles with quantile probabilities at.
na.action	a function which indicates what should happen when the data contain NAs. Default is na.pass

at	vector of quantile probabilities used for type = "quantile".
left	left censoring or truncation point. Only used for type = "quantile". If NULL, censoring or truncation point is obtained from object.
right	right censoring or truncation point. Only used for type = "quantile". If NULL, censoring or truncation point is obtained from object.
...	further arguments passed to or from other methods.

**Value**

For type "response", "location", or "scale" a vector with either the location or the scale of the predicted distribution.

For type "quantile" a matrix of predicted quantiles each column corresponding to an element of at.

**See Also**

[crch](#)

---

predict.crch.boost      *Predicted/Fitted Values for boosted CRCH Fits*

---

**Description**

Obtains various types of predictions for crch.boost models.

**Usage**

```
## S3 method for class 'crch.boost'
predict(object, newdata = NULL, mstop = NULL, type = c("response",
  "location", "scale", "quantile"), na.action = na.pass, at = 0.5, left = NULL,
  right = NULL, ...)
```

**Arguments**

object	an object of class "crch.boost".
newdata	an optional data frame in which to look for variables with which to predict.
mstop	stopping iteration. Can be either a character ("max", "aic", "bic", "cv") or a numeric value. If not NULL, newdata has to be supplied.
type	type of prediction: "response" is equivalent to "location" and returns the location of the predicted distribution. "scale" returns the scale of the predicted distribution and "quantile" returns a matrix of predicted quantiles with quantile probabilities at.
na.action	a function which indicates what should happen when the data contain NAs. Default is na.pass
at	vector of quantile probabilities used for type = "quantile".

left	left censoring or truncation point. Only used for type = "quantile". If NULL, censoring or truncation point is obtained from object.
right	right censoring or truncation point. Only used for type = "quantile". If NULL, censoring or truncation point is obtained from object.
...	further arguments passed to or from other methods.

**Value**

For type "response", "location", or "scale" a vector with either the location or the scale of the predicted distribution.

For type "quantile" a matrix of predicted quantiles each column corresponding to an element of at.

**See Also**

[crch.boost](#), [predict.crch](#)

---

predict.hxlr

*Predict/Fitted Values for HXLR Fits*

---

**Description**

Obtains various types of predictions/fitted values for heteroscedastic extended logistic regression (HXLR) models.

**Usage**

```
## S3 method for class 'hxlr'
predict(object, newdata = NULL, type = c("class", "probability",
  "cumprob", "location", "scale"), thresholds = object$thresholds,
  na.action = na.pass, ...)
## S3 method for class 'hxlr'
fitted(object, type = c("class", "probability",
  "cumprob", "location", "scale"), ...)
```

**Arguments**

object	an object of class "hxlr".
newdata	an optional data frame in which to look for variables with which to predict.
type	type of prediction: "probability" returns a data frame with category probabilities, "cumprob" returns cumulative probabilities, "location" and "scale" return the location and scale of the predicted latent distribution respectively, and "class" returns the category with the highest probability. Default is "class".
thresholds	optional thresholds used for defining the thresholds for types "probability", "cumprob", and "class". Can differ from thresholds used for fitting. If omitted, the same thresholds as for fitting are used.

`na.action` A function which indicates what should happen when the data contain NAs. Default is `na.pass`

`...` further arguments passed to or from other methods.

### Value

For type `"prob"` a matrix with number of intervals (= number of thresholds + 1) columns is produced. Each row corresponds to a row in `newdata` and contains the predicted probabilities to fall in the corresponding interval.

For type `"cumprob"` a matrix with number of thresholds columns is produced. Each row corresponds to a row in `newdata` and contains the predicted probabilities to fall below the corresponding threshold.

For types `"class"`, `"location"`, and `"scale"` a vector is returned respectively with either the most probable categories (`"class"`) or the location (`"location"`) or scale (`scale`) of the latent distribution.

### See Also

[hxr](#)

---

RainIbk

*Precipitation Observations and Forecasts for Innsbruck*

---

### Description

Accumulated 5-8 days precipitation amount for Innsbruck. Data includes GEFS reforecasts (Hamill et al. 2013) and observations from SYNOP station Innsbruck Airport (11120) from 2000-01-01 to 2013-09-17.

### Usage

```
data("RainIbk")
```

### Format

A data frame with 4977 rows. The first column (`rain`) are 3 days accumulated precipitation amount observations, Columns 2-12 (`rainfc`) are 5-8 days accumulated precipitation amount forecasts from the individual ensemble members.

### Source

Observations: <http://www.ogimet.com/synops.phtml.en>

Reforecasts: <http://www.esrl.noaa.gov/psd/forecasts/reforecast2/>

## References

Hamill TM, Bates GT, Whitaker JS, Murray DR, Fiorino M, Galarneau Jr TJ, Zhu Y, Lapenta W (2013). NOAA's Second-Generation Global Medium-Range Ensemble Reforecast Data Set. *Bulletin of the American Meteorological Society*, in press.

## Examples

```
## Spread skill relationship ##

## load and prepare data
data(RainIbk)

## mean and standard deviation of square root transformed ensemble forecasts
RainIbk$sqrtensmean <-
  apply(sqrt(RainIbk[,grep('^rainfc',names(RainIbk))]), 1, mean)
RainIbk$sqrtenssd <-
  apply(sqrt(RainIbk[,grep('^rainfc',names(RainIbk))]), 1, sd)

## quintiles of sqrtenssd
sdcat <- cut(RainIbk$sqrtenssd, c(-Inf, quantile(RainIbk$sqrtenssd,
  seq(0.2,0.8,0.2)), Inf), labels = c(1:5))

## mean forecast errors for each quintile
m <- NULL
for(i in levels(sdcat)) {
  m <- c(m, mean((sqrt(RainIbk$rain)[sdcat == i] -
    RainIbk$sqrtensmean[sdcat == i])^2, na.rm = TRUE))
}

## plot
boxplot((sqrt(rain) - sqrtensmean)^2~sdcat, RainIbk,
  xlab = "Quintile of ensemble standard deviation",
  ylab = "mean squared error", main = "Spread skill relationship")
```

---

tlogis

*The Truncated Logistic Distribution*


---

## Description

Density, distribution function, quantile function, and random generation for the left and/or right truncated logistic distribution.

## Usage

```
dtlogis(x, mean = 0, sd = 1, left = -Inf, right = Inf, log = FALSE)
```

```
ptlogis(q, mean = 0, sd = 1, lower.tail = TRUE, log.p = FALSE, left = -Inf,
  right = Inf)
```



```
rtlogis(n, mean = 0, sd = 1, left = -Inf, right = Inf)
```

```
qtlogis(p, mean = 0, sd = 1, lower.tail = TRUE, log.p = FALSE, left = -Inf,
        right = Inf)
```

### Arguments

<code>x, q</code>	vector of quantiles.
<code>p</code>	vector of probabilities.
<code>n</code>	number of observations. If <code>length(n) &gt; 1</code> , the length is taken to be the number required.
<code>mean</code>	location parameter.
<code>sd</code>	scale parameter.
<code>log, log.p</code>	logical; if TRUE, probabilities <code>p</code> are given as <code>log(p)</code> .
<code>lower.tail</code>	logical; if TRUE (default), probabilities are $P[X \leq x]$ otherwise, $P[X > x]$ .
<code>left</code>	left truncation point.
<code>right</code>	right truncation point.

### Details

If `mean` or `sd` are not specified they assume the default values of 0 and 1, respectively. `left` and `right` have the defaults `-Inf` and `Inf` respectively.

The truncated logistic distribution has density

$$f(x) = 1/\sigma \lambda((x - \mu)/\sigma) / (\Lambda((right - \mu)/\sigma) - \Lambda((left - \mu)/\sigma))$$

for  $left \leq x \leq right$ , and 0 otherwise.

$\Lambda$  and  $\lambda$  are the cumulative distribution function and probability density function of the standard logistic distribution respectively,  $\mu$  is the location of the distribution, and  $\sigma$  the scale.

### Value

`dtlogis` gives the density, `ptlogis` gives the distribution function, `qtlogis` gives the quantile function, and `rtlogis` generates random deviates.

### See Also

[dlogis](#)

tnorm

*The Truncated Normal Distribution***Description**

Density, distribution function, quantile function, and random generation for the left and/or right truncated normal distribution.

**Usage**

```
dtnorm(x, mean = 0, sd = 1, left = -Inf, right = Inf, log = FALSE)
```

```
ptnorm(q, mean = 0, sd = 1, lower.tail = TRUE, log.p = FALSE, left = -Inf,
       right = Inf)
```

```
rtnorm(n, mean = 0, sd = 1, left = -Inf, right = Inf)
```

```
qtnorm(p, mean = 0, sd = 1, lower.tail = TRUE, log.p = FALSE, left = -Inf,
       right = Inf)
```

**Arguments**

<code>x, q</code>	vector of quantiles.
<code>p</code>	vector of probabilities.
<code>n</code>	number of observations. If <code>length(n) &gt; 1</code> , the length is taken to be the number required.
<code>mean</code>	vector of means.
<code>sd</code>	vector of standard deviations.
<code>log, log.p</code>	logical; if TRUE, probabilities <code>p</code> are given as <code>log(p)</code> .
<code>lower.tail</code>	logical; if TRUE (default), probabilities are $P[X \leq x]$ otherwise, $P[X > x]$ .
<code>left</code>	left censoring point.
<code>right</code>	right censoring point.

**Details**

If `mean` or `sd` are not specified they assume the default values of 0 and 1, respectively. `left` and `right` have the defaults `-Inf` and `Inf` respectively.

The truncated normal distribution has density

$$f(x) = 1/\sigma\phi((x - \mu)/\sigma)/(\Phi((right - \mu)/\sigma) - \Phi((left - \mu)/\sigma))$$

for  $left \leq x \leq right$ , and 0 otherwise.

$\Phi$  and  $\phi$  are the cumulative distribution function and probability density function of the standard normal distribution respectively,  $\mu$  is the mean of the distribution, and  $\sigma$  the standard deviation.

**Value**

dttnorm gives the density, pttnorm gives the distribution function, qttnorm gives the quantile function, and rttnorm generates random deviates.

**See Also**

[dnorm](#)

---

 tt

---

*The Truncated Student-t Distribution*


---

**Description**

Density, distribution function, quantile function, and random generation for the left and/or right truncated student-t distribution with df degrees of freedom.

**Usage**

```
dttnorm(x, mean = 0, sd = 1, df, left = -Inf, right = Inf, log = FALSE)
```

```
pttnorm(q, mean = 0, sd = 1, df, lower.tail = TRUE, log.p = FALSE, left = -Inf,
        right = Inf)
```

```
rttnorm(n, mean = 0, sd = 1, df, left = -Inf, right = Inf)
```

```
qttnorm(p, mean = 0, sd = 1, df, lower.tail = TRUE, log.p = FALSE, left = -Inf,
        right = Inf)
```

**Arguments**

x, q	vector of quantiles.
p	vector of probabilities.
n	number of observations. If length(n) > 1, the length is taken to be the number required.
mean	location parameter.
sd	scale parameter.
df	degrees of freedom (> 0, maybe non-integer). df = Inf is allowed.
log, log.p	logical; if TRUE, probabilities p are given as log(p).
lower.tail	logical; if TRUE (default), probabilities are P[X <= x] otherwise, P[X > x].
left	left censoring point.
right	right censoring point.

**Details**

If mean or sd are not specified they assume the default values of 0 and 1, respectively. `left` and `right` have the defaults `-Inf` and `Inf` respectively.

The truncated student-t distribution has density

$$f(x) = 1/\sigma\tau((x - \mu)/\sigma)/(T((right - \mu)/\sigma) - T((left - \mu)/\sigma))$$

for  $left \leq x \leq right$ , and 0 otherwise.

where  $T$  and  $\tau$  are the cumulative distribution function and probability density function of the student-t distribution with `df` degrees of freedom respectively,  $\mu$  is the location of the distribution, and  $\sigma$  the scale.

**Value**

`dt` gives the density, `pt` gives the distribution function, `qt` gives the quantile function, and `rt` generates random deviates.

**See Also**

[dt](#)

# Index

- \*Topic **datasets**
  - RainIbk, 23
- \*Topic **distribution**
  - clogis, 2
  - cnorm, 3
  - ct, 14
  - tlogis, 24
  - tnorm, 26
  - tt, 27
- \*Topic **regression, boosting**
  - crch.boost, 10
- \*Topic **regression**
  - coef.crch, 4
  - coef.crch.boost, 5
  - coef.hxlr, 6
  - crch, 7
  - crch.control, 13
  - hxlr, 15
  - hxlr.control, 19
  - plot.crch.boost, 19
  - predict.crch, 20
  - predict.crch.boost, 21
  - predict.hxlr, 22
- clm, 15, 16
- clogis, 2
- cnorm, 3
- coef.crch, 4, 6
- coef.crch.boost, 5
- coef.hxlr, 6
- crch, 5, 7, 12, 13, 21
- crch.boost, 6, 8, 9, 10, 20, 22
- crch.control, 8, 9, 11, 12, 13
- crch.fit, 11
- ct, 14
- dclogis(clogis), 2
- dcnorm(cnorm), 3
- dct(ct), 14
- dlogis, 3, 25
- dnorm, 4, 27
- dt, 15, 28
- dtlogis(tlogis), 24
- dtnorm(tnorm), 26
- dtt(tt), 27
- fitted, 6
- fitted.crch(coef.crch), 4
- fitted.hxlr(predict.hxlr), 22
- getSummary.crch(coef.crch), 4
- hxlr, 7, 15, 19, 23
- hxlr.control, 15, 19
- logLik, 5, 7
- logLik.crch(coef.crch), 4
- logLik.crch.boost(coef.crch.boost), 5
- logLik.hxlr(coef.hxlr), 6
- model.frame, 6
- model.frame.crch(coef.crch), 4
- model.matrix, 6
- model.matrix.crch(coef.crch), 4
- optim, 8, 9, 13, 15, 16, 19
- pclogis(clogis), 2
- pcnorm(cnorm), 3
- pct(ct), 14
- plot.crch.boost, 19
- plot.ts, 20
- predict.crch, 9, 20, 22
- predict.crch.boost, 21
- predict.hxlr, 16, 22
- print, 5, 7
- print.crch(coef.crch), 4
- print.crch.boost(coef.crch.boost), 5
- print.hxlr(coef.hxlr), 6
- print.summary.crch(coef.crch), 4

`print.summary.crch.boost`  
    (`coef.crch.boost`), 5  
`print.summary.hxlr` (`coef.hxlr`), 6  
`ptlogis` (`tlogis`), 24  
`ptnorm` (`tnorm`), 26  
`ptt` (`tt`), 27

`qclgis` (`clogis`), 2  
`qcnorm` (`cnorm`), 3  
`qct` (`ct`), 14  
`qtlogis` (`tlogis`), 24  
`qtnorm` (`tnorm`), 26  
`qtt` (`tt`), 27

`RainIbk`, 23  
`rclogis` (`clogis`), 2  
`rcnorm` (`cnorm`), 3  
`rct` (`ct`), 14  
`residuals`, 5, 6  
`residuals.crch` (`coef.crch`), 4  
`rtlogis` (`tlogis`), 24  
`rtnorm` (`tnorm`), 26  
`rtt` (`tt`), 27

`summary`, 5, 7  
`summary.crch` (`coef.crch`), 4  
`summary.crch.boost` (`coef.crch.boost`), 5  
`summary.hxlr` (`coef.hxlr`), 6

`terms`, 6  
`terms.crch` (`coef.crch`), 4  
`terms.hxlr` (`coef.hxlr`), 6  
`tlogis`, 24  
`tnorm`, 26  
`trch` (`crch`), 7  
`tt`, 27

`vcov.crch` (`coef.crch`), 4  
`vcov.hxlr` (`coef.hxlr`), 6