

Package ‘locfit’

April 17, 2009

Version 1.5-4

Title Local Regression, Likelihood and Density Estimation.

Date 2007-11-27

Author Catherine Loader <locfit@herine.net>

Maintainer Andy Liaw <andy_liaw@merck.com>

Description Local Regression, Likelihood and density estimation.

Depends R (>= 2.0.1), akima, lattice

Suggests gam

License GPL (>= 2)

URL <http://locfit.herine.net/>

Repository CRAN

Date/Publication 2007-11-27 19:19:17

R topics documented:

aic	4
aicplot	4
ais	5
ang	6
bad	7
border	7
chemdiab	8
claw54	8
cldem	9
cltest	9
cltrain	10
co2	10
cp	11

cpar	11
cpplot	12
crit	13
dat	14
density.lf	14
diab	15
ethanol	16
expit	16
fitted.locfit	17
formula.locfit	18
gam.lf	18
gam.slist	19
gcv	19
gcvplot	20
geyser	20
geyser.round	21
hatmatrix	22
heart	22
insect	23
iris	23
kangaroo	24
kappa0	25
kdeb	26
km.mrl	27
lcv	28
lcvplot	28
left	29
lf	30
lfeval	30
lfgrid	31
lfknots	32
lflim	32
lfmarg	33
lines.locfit	33
livmet	34
locfit	35
locfit.censor	36
locfit.matrix	37
locfit.quasi	38
locfit.raw	39
locfit.robust	41
lp	42
lscv	43
lscv.exact	44
lscvplot	45
mcyc	45
mine	46
mmsamp	46

morths	47
none	47
penny	48
plot.eval	48
plot.gcvplot	49
plot.lfeval	50
plot.locfit	50
plot.locfit.1d	52
plot.locfit.2d	52
plot.locfit.3d	53
plot.preplot.locfit	54
plot.scb	54
plotbyfactor	55
points.locfit	56
predict.locfit	56
preplot.locfit	57
preplot.locfit.raw	58
print.gcvplot	59
print.lfeval	60
print.locfit	60
print.preplot.locfit	61
print.scb	61
print.summary.locfit	62
rbox	62
regband	63
residuals.locfit	64
right	64
rv	65
rva	66
scb	66
sjpi	67
smooth.lf	68
spence.15	69
spence.21	70
spencer	71
stamp	71
store	72
summary.gcvplot	72
summary.locfit	73
summary.preplot.locfit	74
trimod	74
xbar	75

`aic`*Compute Akaike's Information Criterion.*

Description

The calling sequence for `aic` matches those for the `locfit` or `locfit.raw` functions. The fit is not returned; instead, the returned object contains Akaike's information criterion for the fit.

The definition of AIC used here is $-2 \times \log\text{-likelihood} + \text{pen} \times (\text{fitted d.f.})$. For quasi-likelihood, and local regression, this assumes the scale parameter is one. Other scale parameters can effectively be used by changing the penalty.

The AIC score is exact (up to numerical roundoff) if the `ev="data"` argument is provided. Otherwise, the residual sum-of-squares and degrees of freedom are computed using `locfit`'s standard interpolation based approximations.

Usage

```
aic(x, ..., pen=2)
```

Arguments

<code>x</code>	model formula
<code>...</code>	other arguments to <code>locfit</code>
<code>pen</code>	penalty for the degrees of freedom term

See Also

[locfit](#), [locfit.raw](#), [aicplot](#)

`aicplot`*Compute an AIC plot.*

Description

The `aicplot` function loops through calls to the `aic` function (and hence to `locfit`), using a different smoothing parameter for each call. The returned structure contains the AIC statistic for each fit, and can be used to produce an AIC plot.

Usage

```
aicplot(..., alpha)
```

Arguments

... arguments to the `aic`, `locfit` functions.

`alpha` Matrix of smoothing parameters. The `aicplot` function loops through calls to `aic`, using each row of `alpha` as the smoothing parameter in turn. If `alpha` is provided as a vector, it will be converted to a one-column matrix, thus interpreting each component as a nearest neighbor smoothing parameter.

Value

An object with class "gcvplot", containing the smoothing parameters and AIC scores. The actual plot is produced using `plot.gcvplot`.

See Also

`locfit`, `locfit.raw`, `gcv`, `aic`, `plot.gcvplot`

Examples

```
data(morths)
plot(aicplot(deaths~age, weights=n, data=morths, family="binomial",
  alpha=seq(0.2, 1.0, by=0.05)))
```

ais

Australian Institute of Sport Dataset

Description

The first two columns are the gender of the athlete and their sport. The remaining 11 columns are various measurements made on the athletes.

Usage

```
data(ais)
```

Format

A dataframe.

Source

Cook and Weisberg (1994).

References

Cook and Weisberg (1994). *An Introduction to Regression Graphics*. Wiley, New York.

ang

Angular Term for a Locfit model.

Description

The `ang()` function is used in a `locfit` model formula to specify that a variable should be treated as an angular or periodic term. The `scale` argument is used to set the period.

`ang(x)` is equivalent to `lp(x, style="ang")`.

Usage

```
ang(x, ...)
```

Arguments

<code>x</code>	numeric variable to be treated periodically.
<code>...</code>	Other arguments to <code>lp</code> .

References

Loader, C. (1999). *Local Regression and Likelihood*. Springer, NY (Section 6.2).

See Also

[locfit](#).

Examples

```
# generate an x variable, and a response with period 0.2
x <- seq(0,1,length=200)
y <- sin(10*pi*x)+rnorm(200)/5

# compute the periodic local fit. Note the scale argument is period/(2pi)
fit <- locfit(y~ang(x,scale=0.2/(2*pi)))

# plot the fit over a single period
plot(fit)

# plot the fit over the full range of the data
plot(fit,xlim=c(0,1))
```

`bad`*Example dataset for bandwidth selection*

Description

Example dataset from Loader (1999).

Usage

```
data(bad)
```

Format

Data Frame with x and y variables.

References

Loader, C. (1999). Bandwidth Selection: Classical or Plug-in? *Annals of Statistics* 27.

`border`*Cricket Batting Dataset*

Description

Scores in 265 innings for Australian batsman Allan Border.

Usage

```
data(border)
```

Format

A dataframe with day (decimalized); not out indicator and score. The not out indicator should be used as a censoring variable.

Source

Compiled from the Cricinfo archives.

References

CricInfo: The Home of Cricket on the Internet. <http://www.cricinfo.com/>

`chemdiab`*Chemical Diabetes Dataset*

Description

Numeric variables are `rw`, `fpg`, `ga`, `ina` and `sspg`. Classifier `cc` is the Diabetic type.

Usage

```
data(chemdiab)
```

Format

Data frame with five numeric measurements and categorical response.

Source

Reaven and Miller (1979).

References

Reaven, G. M. and Miller, R. G. (1979). An attempt to define the nature of chemical diabetes using a multidimensional analysis. *Diabetologia* 16, 17-24.

`claw54`*Claw Dataset*

Description

A random sample of size 54 from the claw density of Marron and Wand (1992), as used in Figure 10.5 of Loader (1999).

Usage

```
data(claw54)
```

Format

Numeric vector with length 54.

Source

Randomly generated.

References

- Loader, C. (1999). Local Regression and Likelihood. Springer, New York.
- Marron, J. S. and Wand, M. P. (1992). Exact mean integrated squared error. Annals of Statistics 20, 712-736.

`cldem`*Example data set for classification*

Description

Observations from Figure 8.7 of Loader (1999).

Usage

```
data(cldem)
```

Format

Data Frame with x and y variables.

References

- Loader, C. (1999). Local Regression and Likelihood. Springer, New York.

`cltest`*Test dataset for classification*

Description

200 observations from a 2 population model. Under population 0, $x_{1,i}$ has a standard normal distribution, and $x_{2,i} = (2 - x_{1,i}^2 + z_i)/3$, where z_i is also standard normal. Under population 1, $x_{2,i} = -(2 - x_{1,i}^2 + z_i)/3$. The optimal classification regions form a checkerboard pattern, with horizontal boundary at $x_2 = 0$, vertical boundaries at $x_1 = \pm\sqrt{2}$.

This is the same model as the cltrain dataset.

Usage

```
data(cltest)
```

Format

Data Frame. Three variables x1, x2 and y. The latter indicates class membership.

`cltrain`*Training dataset for classification*

Description

200 observations from a 2 population model. Under population 0, $x_{1,i}$ has a standard normal distribution, and $x_{2,i} = (2 - x_{1,i}^2 + z_i)/3$, where z_i is also standard normal. Under population 1, $x_{2,i} = -(2 - x_{1,i}^2 + z_i)/3$. The optimal classification regions form a checkerboard pattern, with horizontal boundary at $x_2 = 0$, vertical boundaries at $x_1 = \pm\sqrt{2}$.

This is the same model as the cltest dataset.

Usage

```
data(cltrain)
```

Format

Data Frame. Three variables x1, x2 and y. The latter indicates class membership.

`co2`*Carbon Dioxide Dataset*

Description

Monthly time series of carbon dioxide measurements at Mauna Loa, Hawaii from 1959 to 1990.

Usage

```
data(co2)
```

Format

Data frame with `year`, `month` and `co2` variables.

Source

Boden, Sepanski and Stoss (1992).

References

Boden, Sepanski and Stoss (1992). Trends '91: A compedium of data on global change - Highlights. Carbon Dioxide Information Analysis Center, Oak Ridge National Laboratory.

cp *Compute Mallows' Cp for local regression models.*

Description

The calling sequence for `cp` matches those for the `locfit` or `locfit.raw` functions. The fit is not returned; instead, the returned object contains Cp criterion for the fit.

Cp is usually computed using a variance estimate from the largest model under consideration, rather than $\sigma^2 = 1$. This will be done automatically when the `cpplot` function is used.

The Cp score is exact (up to numerical roundoff) if the `ev="data"` argument is provided. Otherwise, the residual sum-of-squares and degrees of freedom are computed using `locfit`'s standard interpolation based approximations.

Usage

```
cp(x, ..., sig2=1)
```

Arguments

x	model formula or numeric vector of the independent variable.
...	other arguments to <code>locfit</code> and/or <code>locfit.raw</code> .
sig2	residual variance estimate.

See Also

[locfit](#), [locfit.raw](#), [cpplot](#)

cpar *Conditionally parametric term for a Locfit model.*

Description

A term entered in a `locfit` model formula using `cpar` will result in a fit that is conditionally parametric. Equivalent to `lp(x, style="cpar")`.

This function is presently almost deprecated. Specifying a conditionally parametric fit as `y~x1+cpar(x2)` will no longer work; instead, the model is specified as `y~lp(x1, x2, style=c("n", "cpar"))`.

Usage

```
cpar(x, ...)
```

Arguments

x	numeric variable.
...	Other arguments to <code>link{lp}()</code> .

See Also[locfit](#)**Examples**

```
data(ethanol, package="locfit")
# fit a conditionally parametric model
fit <- locfit(NOx ~ lp(E, C, style=c("n", "cpar")), data=ethanol)
plot(fit)
# one way to force a parametric fit with locfit
fit <- locfit(NOx ~ cpar(E), data=ethanol)
```

cpplot

*Compute a Cp plot.***Description**

The `cpplot` function loops through calls to the `cp` function (and hence to `link{locfit}`), using a different smoothing parameter for each call. The returned structure contains the C_p statistic for each fit, and can be used to produce an AIC plot.

Usage

```
cpplot(..., alpha, sig2)
```

Arguments

<code>...</code>	arguments to the <code>cp</code> , <code>locfit</code> functions.
<code>alpha</code>	Matrix of smoothing parameters. The <code>cpplot</code> function loops through calls to <code>cp</code> , using each row of <code>alpha</code> as the smoothing parameter in turn. If <code>alpha</code> is provided as a vector, it will be converted to a one-column matrix, thus interpreting each component as a nearest neighbor smoothing parameter.
<code>sig2</code>	Residual variance. If not specified, the residual variance is computed using the fitted model with the fewest residual degrees of freedom.

Value

An object with class `"gcvplot"`, containing the smoothing parameters and CP scores. The actual plot is produced using `plot.gcvplot`.

See Also[locfit](#), [locfit.raw](#), [gcv](#), [aic](#), [plot.gcvplot](#)**Examples**

```
data(ethanol)
plot(cpplot(NOx~E, data=ethanol, alpha=seq(0.2, 1.0, by=0.05)))
```

crit	<i>Compute critical values for confidence intervals.</i>
------	--

Description

Every "locfit" object contains a critical value object to be used in computing and plotting confidence intervals. By default, a 95% pointwise confidence level is used. To change the confidence level, the critical value object must be substituted using `crit` and `crit<-`.

Usage

```
crit(fit, const=c(0, 1), d=1, cov=0.95, rdf=0)
crit(fit) <- value
```

Arguments

fit	"locfit" object. This is optional; if a fit is provided, defaults for the other arguments are taken from the critical value currently stored on this fit, rather than the usual values above. <code>crit(fit)</code> with no other arguments will just return the current critical value.
const	Tube formula constants for simultaneous bands (the default, <code>c(0, 1)</code> , produces pointwise coverage). Usually this is generated by the <code>kappa0</code> function and should not be provided by the user.
d	Dimension of the fit. Again, users shouldn't usually provide it.
cov	Coverage Probability for critical values.
rdf	Residual degrees of freedom. If non-zero, the critical values are based on the Student's t distribution. When <code>rdf=0</code> , the normal distribution is used.
value	Critical value object generated by <code>crit</code> or <code>kappa0</code> .

Value

Critical value object.

See Also

`locfit`, `plot.locfit`, `kappa0`, `crit<-`.

Examples

```
# compute and plot 99% confidence intervals, with local variance estimate.
data(ethanol)
fit <- locfit(NOx~E,data=ethanol)
crit(fit) <- crit(fit,cov=0.99)
plot(fit,band="local")

# compute and plot 99% simultaneous bands
crit(fit) <- kappa0(NOx~E,data=ethanol,cov=0.99)
plot(fit,band="local")
```

dat	<i>Locfit - data evaluation structure.</i>
-----	--

Description

dat is used to specify evaluation on the given data points for `locfit.raw()`.

Usage

```
dat (cv=FALSE)
```

Arguments

cv	Whether cross-validation should be done.
----	--

density.lf	<i>Density estimation using Locfit</i>
------------	--

Description

This function provides an interface to Locfit, in the syntax of (a now old version of) the S-Plus density function. This can reproduce density results, but allows additional `locfit.raw` arguments, such as the degree of fit, to be given.

It also works in double precision, whereas `density` only works in single precision.

Usage

```
density.lf(x, n = 50, window = "gaussian", width, from, to,
  cut = if(iwindow == 4.) 0.75 else 0.5,
  ev = lfgrid(mg = n, ll = from, ur = to),
  deg = 0, family = "density", link = "ident", ...)
```

Arguments

x	numeric vector of observations whose density is to be estimated.
n	number of evaluation points. Equivalent to the <code>locfit.raw</code> <code>mg</code> argument.
window	Window type to use for estimation. Equivalent to the <code>locfit.raw</code> <code>kern</code> argument. This includes all the density windows except <code>cosine</code> .
width	Window width. Following <code>density</code> , this is the full width; not the half-width usually used by Locfit and many other smoothers.
from	Lower limit for estimation domain.
to	Upper limit for estimation domain.
cut	Controls default expansion of the domain.

ev Locfit evaluation structure – default `lfgrid()`.
 deg Fitting degree – default 0 for kernel estimation.
 family Fitting family – default is "density".
 link Link function – default is the "identity".
 ... Additional arguments to `locfit.raw`, with standard defaults.

Value

A list with components `x` (evaluation points) and `y` (estimated density).

See Also

`density`, `locfit`, `locfit.raw`

Examples

```

data(geyser)
density.lf(geyser, window="tria")
# the same result with density, except less precision.
density(geyser, window="tria")

```

 diab

Exhaust emissions

Description

NOx exhaust emissions from a single cylinder engine. Two predictor variables are E (the engine's equivalence ratio) and C (Compression ratio).

Usage

```
data(ethanol)
```

Format

Data frame with NOx, E and C variables.

Source

Brinkman (1981). Also studied extensively by Cleveland (1993).

References

Brinkman, N. D. (1981). Ethanol fuel - a single-cylinder engine study of efficiency and exhaust emissions. SAE transactions 90, 1414-1424.

Cleveland, W. S. (1993). Visualizing data. Hobart Press, Summit, NJ.

`ethanol`*Exhaust emissions*

Description

NOx exhaust emissions from a single cylinder engine. Two predictor variables are E (the engine's equivalence ratio) and C (Compression ratio).

Usage

```
data(ethanol)
```

Format

Data frame with NOx, E and C variables.

Source

Brinkman (1981). Also studied extensively by Cleveland (1993).

References

Brinkman, N. D. (1981). Ethanol fuel - a single-cylinder engine study of efficiency and exhaust emissions. SAE transactions 90, 1414-1424.

Cleveland, W. S. (1993). Visualizing data. Hobart Press, Summit, NJ.

`expit`*Inverse logistic link function*

Description

Computes $e^x / (1 + e^x)$. This is the inverse of the logistic link function, $\log(p / (1 - p))$.

Usage

```
expit(x)
```

Arguments

`x` numeric vector

fitted.locfit *Fitted values for a "locfit" object.*

Description

Evaluates the fitted values (i.e. evaluates the surface at the original data points) for a Locfit object. This function works by reconstructing the model matrix from the original formula, and predicting at those points. The function may be fooled; for example, if the original data frame has changed since the fit, or if the model formula includes calls to random number generators.

Usage

```
## S3 method for class 'locfit':
fitted(object, data=NULL, what="coef", cv=FALSE,
studentize=FALSE, type="fit", tr, ...)
```

Arguments

object	"locfit" object.
data	The data frame for the original fit. Usually, this shouldn't be needed, especially when the function is called directly. It may be needed when called inside another function.
what	What to compute fitted values of. The default, <code>what="coef"</code> , works with the fitted curve itself. Other choices include <code>"nlx"</code> for the length of the weight diagram; <code>"infl"</code> for the influence function; <code>"band"</code> for the bandwidth; <code>"degr"</code> for the local polynomial degree; <code>"lik"</code> for the maximized local likelihood; <code>"rdf"</code> for the local residual degrees of freedom and <code>"vari"</code> for the variance function. The interpolation algorithm for some of these quantities is questionable.
cv	If TRUE, leave-one-out cross validated fitted values are approximated. Won't make much sense, unless <code>what="coef"</code> .
studentize	If TRUE, residuals are studentized.
type	Type of fit or residuals to compute. The default is <code>"fit"</code> for <code>fitted.locfit</code> , and <code>"dev"</code> for <code>residuals.locfit</code> . Other choices include <code>"pear"</code> for Pearson residuals; <code>"raw"</code> for raw residuals, <code>"ldot"</code> for likelihood derivative; <code>"d2"</code> for the deviance residual squared; <code>lddot</code> for the likelihood second derivative. Generally, <code>type</code> should only be used when <code>what="coef"</code> .
tr	Back transformation for likelihood models.
...	arguments passed to and from methods.

Value

A numeric vector of the fitted values.

See Also

`locfit`, `predict.locfit`, `residuals.locfit`

`formula.locfit` *Formula from a Locfit object.*

Description

Extract the model formula from a locfit object.

Usage

```
## S3 method for class 'locfit':
formula(x, ...)
```

Arguments

`x` locfit object.
`...` Arguments passed to and from other methods.

Value

Returns the formula from the locfit object.

See Also

[locfit](#)

`gam.lf` *Locfit call for Generalized Additive Models*

Description

This is a locfit calling function used by `lf()` terms in additive models. It is not normally called directly by users.

Usage

```
gam.lf(x, y, w, xeval, ...)
```

Arguments

`x` numeric predictor
`y` numeric response
`w` prior weights
`xeval` evaluation points
`...` other arguments to `locfit.raw()`

See Also

[locfit](#), [locfit.raw](#), [lf](#), [gam](#)

`gam.slist`

Vector of GAM special terms

Description

This vector adds "lf" to the default vector of special terms recognized by a `gam()` model formula. To ensure this is recognized, attach the Locfit library with `library(locfit, first=T)`.

Format

Character vector.

See Also

[lf](#), [gam](#)

`gcv`

Compute generalized cross-validation statistic.

Description

The calling sequence for `gcv` matches those for the [locfit](#) or [locfit.raw](#) functions. The fit is not returned; instead, the returned object contains Wahba's generalized cross-validation score for the fit.

The GCV score is exact (up to numerical roundoff) if the `ev="data"` argument is provided. Otherwise, the residual sum-of-squares and degrees of freedom are computed using `locfit`'s standard interpolation based approximations.

For likelihood models, GCV is computed uses the deviance in place of the residual sum of squares. This produces useful results but I do not know of any theory validating this extension.

Usage

`gcv(x, ...)`

Arguments

`x, ...` Arguments passed on to [locfit](#) or [locfit.raw](#).

See Also

[locfit](#), [locfit.raw](#), [gcvplot](#)

`gcvplot` *Compute a generalized cross-validation plot.*

Description

The `gcvplot` function loops through calls to the `gcv` function (and hence to `link{locfit}`), using a different smoothing parameter for each call. The returned structure contains the GCV statistic for each fit, and can be used to produce an GCV plot.

Usage

```
gcvplot(..., alpha, df=2)
```

Arguments

`...` arguments to the `gcv`, `locfit` functions.

`alpha` Matrix of smoothing parameters. The `gcvplot` function loops through calls to `gcv`, using each row of `alpha` as the smoothing parameter in turn. If `alpha` is provided as a vector, it will be converted to a one-column matrix, thus interpreting each component as a nearest neighbor smoothing parameter.

`df` Degrees of freedom to use as the x-axis. 2= $\text{trace}(L)$, 3= $\text{trace}(L'L)$.

Value

An object with class "`gcvplot`", containing the smoothing parameters and GCV scores. The actual plot is produced using `plot.gcvplot`.

See Also

[locfit](#), [locfit.raw](#), [gcv](#), [plot.gcvplot](#), [summary.gcvplot](#)

Examples

```
data(ethanol)
plot(gcvplot(NOx~E, data=ethanol, alpha=seq(0.2, 1.0, by=0.05)))
```

`geyser` *Old Faithful Geysers Dataset*

Description

The durations of 107 eruptions of the Old Faithful Geysers.

Usage

```
data(geyser)
```

Format

A numeric vector of length 107.

Source

Scott (1992). Note that several different Old Faithful Geyser datasets (including the faithful dataset in R's base library) have been used in various places in the statistics literature. The version provided here has been used in density estimation and bandwidth selection work.

References

Scott, D. W. (1992). *Multivariate Density Estimation: Theory, Practice and Visualization*. Wiley.

geyser.round

Discrete Old Faithful Geyser Dataset

Description

This is a variant of the [geyser](#) dataset, where each observation is rounded to the nearest 0.05 minutes, and the counts tallied.

Usage

```
data(geyser.round)
```

Format

Data Frame with variables `duration` and `count`.

Source

Scott (1992). Note that several different Old Faithful Geyser datasets (including the faithful dataset in R's base library) have been used in various places in the statistics literature. The version provided here has been used in density estimation and bandwidth selection work.

References

Scott, D. W. (1992). *Multivariate Density Estimation: Theory, Practice and Visualization*. Wiley.

 hatmatrix

Weight diagrams and the hat matrix for a local regression model.

Description

`hatmatrix()` computes the weight diagrams (also known as equivalent or effective kernels) for a local regression smooth. Essentially, `hatmatrix()` is a front-end to `locfit()`, setting a flag to compute and return weight diagrams, rather than the fit.

Usage

```
hatmatrix(formula, dc=TRUE, ...)
```

Arguments

<code>formula</code>	model formula.
<code>dc</code>	derivative adjustment (see <code>locfit.raw</code>)
<code>...</code>	Other arguments to <code>locfit</code> and <code>locfit.raw</code> .

Value

A matrix with `n` rows and `p` columns; each column being the weight diagram for the corresponding `locfit` fit point. If `ev="data"`, this is the transpose of the hat matrix.

See Also

[locfit](#), [plot.locfit.1d](#), [plot.locfit.2d](#), [plot.locfit.3d](#), [lines.locfit](#), [predict.locfit](#)

 heart

Survival Times of Heart Transplant Recipients

Description

The survival times of 184 participants in the Stanford heart transplant program.

Usage

```
data(heart)
```

Format

Data frame with `surv`, `cens` and `age` variables.

Source

Miller and Halperin (1982). The original dataset includes information on additional patients who never received a transplant. Other authors reported earlier versions of the data.

References

Miller, R. G. and Halperin, J. (1982). Regression with censored data. *Biometrika* 69, 521-531.

`insect`*Insect Dataset*

Description

An experiment measuring death rates for insects, with 30 insects at each of five treatment levels.

Usage

```
data(insect)
```

Format

Data frame with `lconc` (dosage), `deaths` (number of deaths) and `nins` (number of insects) variables.

Source

Bliss (1935).

References

Bliss (1935). The calculation of the dosage-mortality curve. *Annals of Applied Biology* 22, 134-167.

`iris`*Fisher's Iris Data (subset)*

Description

Four measurements on each of fifty flowers of two species of iris (Versicolor and Virginica) – A classification dataset. Fisher's original dataset contained a third species (Setosa) which is trivially separable.

Usage

```
data(iris)
```

Format

Data frame with species, petal.wid, petal.len, sepal.wid, sepal.len.

Source

Fisher (1936). Reproduced in Andrews and Herzberg (1985) Chapter 1.

References

Andrews, D. F. and Herzberg, A. M. (1985). Data. Springer-Verlag.

Fisher, R. A. (1936). The Use of Multiple Measurements in Taxonomic Problems. *Annals of Eugenics* 7, Part II. 179-188.

kangaroo

Kangaroo skull measurements dataset

Description

Variables are sex (m/f), spec (giganteus, melanops, fuliginosus) and 18 numeric measurements.

Usage

```
data(kangaroo)
```

Format

Data frame with measurements on the skulls of 101 kangaroos. (number of insects) variables.

Source

Andrews and Herzberg (1985) Chapter 53.

References

Andrews, D. F. and Herzberg, A. M. (1985). Data. Springer-Verlag, New York.

Description

The geometric constants for simultaneous confidence bands are computed, as described in Sun and Loader (1994) (bias adjustment is not implemented here). These are then passed to the `crit` function, which computes the critical value for the confidence bands.

The method requires both the weight diagrams $l(x)$, the derivative $l'(x)$ and (in 2 or more dimensions) the second derivatives $l''(x)$. These are implemented exactly for a constant bandwidth. For nearest neighbor bandwidths, the computations are approximate and a warning is produced.

The theoretical justification for the bands uses normality of the random errors e_1, \dots, e_n in the regression model, and in particular the spherical symmetry of the error vector. For non-normal distributions, and likelihood models, one relies on central limit and related theorems.

Computation uses the product Simpson's rule to evaluate the multidimensional integrals (The domain of integration, and hence the region of simultaneous coverage, is determined by the `flim` argument). Expect the integration to be slow in more than one dimension. The `mint` argument controls the precision.

Usage

```
kappa0(formula, cov=0.95, ev=lfgrid(20), ...)
```

Arguments

<code>formula</code>	Local regression model formula. A "locfit" object can also be provided; in this case the formula and other arguments are extracted from this object.
<code>cov</code>	Coverage Probability for critical values.
<code>ev</code>	Locfit evaluation structure. Should usually be a grid – this specifies the integration rule.
<code>...</code>	Other arguments to <code>locfit</code> . Important arguments include <code>flim</code> and <code>alpha</code> .

Value

A list with components for the critical value, geometric constants, e.t.c. Can be passed directly to `plot.locfit` as the `crit` argument.

References

Sun, J. and Loader, C. (1994). Simultaneous confidence bands for linear regression and smoothing. *Annals of Statistics* 22, 1328-1345.

See Also

`locfit`, `plot.locfit`, `crit`, `crit<-`.

Examples

```
# compute and plot simultaneous confidence bands
data(ethanol)
fit <- locfit(NOx~E,data=ethanol)
crit(fit) <- kappa0(NOx~E,data=ethanol)
plot(fit,crit=crit,band="local")
```

kdeb

*Bandwidth selectors for kernel density estimation.***Description**

Function to compute kernel density estimate bandwidths, as used in the simulation results in Chapter 10 of Loader (1999).

This function is included for comparative purposes only. Plug-in selectors are based on flawed logic, make unreasonable and restrictive assumptions and do not use the full power of the estimates available in Locfit. Any relation between the results produced by this function and desirable estimates are entirely coincidental.

Usage

```
kdeb(x, h0 = 0.01 * sd, h1 = sd, meth = c("AIC", "LCV", "LSCV", "BCV",
    "SJPI", "GKK"), kern = "gauss", gf = 2.5)
```

Arguments

x	One dimensional data vector.
h0	Lower limit for bandwidth selection. Can be fairly small, but h0=0 would cause problems.
h1	Upper limit.
meth	Required selection method(s).
kern	Kernel. Most methods require kern="gauss", the default for this function only.
gf	Standard deviation for the gaussian kernel. Default 2.5, as Locfit's standard. Most papers use 1.

Value

Vector of selected bandwidths.

References

Loader, C. (1999). Local Regression and Likelihood. Springer, New York.

`km.mrl`*Mean Residual Life using Kaplan-Meier estimate*

Description

This function computes the mean residual life for censored data using the Kaplan-Meier estimate of the survival function. If $S(t)$ is the K-M estimate, the MRL for a censored observation is computed as $(\int_t^\infty S(u)du)/S(t)$. We take $S(t) = 0$ when t is greater than the largest observation, regardless of whether that observation was censored.

When there are ties between censored and uncensored observations, for definiteness our ordering places the censored observations before uncensored.

This function is used by `locfit.censor` to compute censored regression estimates.

Usage

```
km.mrl(times, cens)
```

Arguments

<code>times</code>	Observed survival times.
<code>cens</code>	Logical variable indicating censoring. The coding is 1 or TRUE for censored; 0 or FALSE for uncensored.

Value

A vector of the estimated mean residual life. For uncensored observations, the corresponding estimate is 0.

References

Buckley, J. and James, I. (1979). Linear Regression with censored data. *Biometrika* 66, 429-436.
Loader, C. (1999). *Local Regression and Likelihood*. Springer, NY (Section 7.2).

See Also

[locfit.censor](#)

Examples

```
# censored regression using the Kaplan-Meier estimate.
data(heart, package="locfit")
fit <- locfit.censor(log10(surv+0.5)~age, cens=cens, data=heart, km=TRUE)
plotbyfactor(heart$age, 0.5+heart$urv, heart$cens, ylim=c(0.5,16000), log="y")
lines(fit, tr=function(x)10^x)
```

lcv	<i>Compute Likelihood Cross Validation Statistic.</i>
-----	---

Description

The calling sequence for `lcv` matches those for the `locfit` or `locfit.raw` functions. The fit is not returned; instead, the returned object contains likelihood cross validation score for the fit.

The LCV score is exact (up to numerical roundoff) if the `ev="cross"` argument is provided. Otherwise, the influence and cross validated residuals are computed using `locfit`'s standard interpolation based approximations.

Usage

```
lcv(x, ...)
```

Arguments

x	model formula
...	other arguments to <code>locfit</code>

See Also

[locfit](#), [locfit.raw](#), [lcvplot](#)

lcvplot	<i>Compute the likelihood cross-validation plot.</i>
---------	--

Description

The `lcvplot` function loops through calls to the `lcv` function (and hence to `link{locfit}`), using a different smoothing parameter for each call. The returned structure contains the likelihood cross validation statistic for each fit, and can be used to produce an LCV plot.

Usage

```
lcvplot(..., alpha)
```

Arguments

...	arguments to the <code>lcv</code> , <code>locfit</code> functions.
alpha	Matrix of smoothing parameters. The <code>aicplot</code> function loops through calls to <code>lcv</code> , using each row of <code>alpha</code> as the smoothing parameter in turn. If <code>alpha</code> is provided as a vector, it will be converted to a one-column matrix, thus interpreting each component as a nearest neighbor smoothing parameter.

Value

An object with class "gcvplot", containing the smoothing parameters and LCV scores. The actual plot is produced using `plot.gcvplot`.

See Also

`locfit`, `locfit.raw`, `gcv`, `lcv`, `plot.gcvplot`

Examples

```
data(ethanol)
plot(lcvplot(NOx~E,data=ethanol,alpha=seq(0.2,1.0,by=0.05)))
```

left

One-sided left smooth for a Locfit model.

Description

The `left()` function is used in a locfit model formula to specify a one-sided smooth: when fitting at a point x , only data points with $x_i \leq x$ should be used. This can be useful in estimating points of discontinuity, and in cross-validation for forecasting a time series. `left(x)` is equivalent to `lp(x, style="left")`.

When using this function, it will usually be necessary to specify an evaluation structure, since the fit is not smooth and locfit's interpolation methods are unreliable. Also, it is usually best to use `deg=0` or `deg=1`, otherwise the fits may be too variable. If nearest neighbor bandwidth specification is used, it does not recognize `left()`.

Usage

```
left(x, ...)
```

Arguments

`x` numeric variable.
`...` Other arguments to `lp()`.

See Also

`locfit`, `lp`, `right`

Examples

```
# compute left and right smooths
data(penny)
xev <- (1945:1988)+0.5
fitl <- locfit(thickness~left(year,h=10,deg=1), ev=xev, data=penny)
fitr <- locfit(thickness~right(year,h=10,deg=1),ev=xev, data=penny)
# plot the squared difference, to show the change points.
plot(xev, (predict(fitr,where="ev") - predict(fitl,where="ev"))^2)
```

lf

*Locfit term in Additive Model formula***Description**

This function is used to specify a smooth term in a `gam()` model formula.

This function is designed to be used with the S-Plus `gam()` function. For R users, there are at least two different `gam()` functions available. Most current distributions of R will include the `mgcv` library by Simon Wood; `lf()` is not compatible with this function.

On CRAN, there is a `gam` package by Trevor Hastie, similar to the S-Plus version. `lf()` should be compatible with this, although it's untested.

Usage

```
lf(..., alpha=0.7, deg=2, scale=1, kern="tcub", ev=rbox(), maxk=100)
```

Arguments

```
...          numeric predictor variable(s)
alpha, deg, scale, kern, ev, maxk
             these are as in locfit.raw.
```

See Also

[locfit](#), [locfit.raw](#), [gam.lf](#), `gam`

Examples

```
# fit an additive semiparametric model to the ethanol data.
stopifnot(require(gam))
# The `gam' package must be attached _before_ `locfit', otherwise
# the following will not work.
fit <- gam(NOx ~ lf(E) + C, data=ethanol)
op <- par(mfrow=c(2, 1))
plot(fit)
par(op)
```

lfeval

*Extract Locfit Evaluation Structure.***Description**

Extracts the evaluation structure from a "locfit" object. This object has the class "lfeval", and has its own set of methods for plotting e.t.c.

Usage

```
lfeval(object)
```

Arguments

```
object          "locfit" object
```

Value

"lfeval" object.

See Also

[locfit](#), [plot.lfeval](#), [print.lfeval](#)

lfgrid

Locfit - grid evaluation structure.

Description

`lfgrid()` is used to specify evaluation on a grid of points for `locfit.raw()`. The structure computes a bounding box for the data, and divides that into a grid with specified margins.

Usage

```
lfgrid(mg=10, ll, ur)
```

Arguments

mg	Number of grid points along each margin. Can be a single number (which is applied in each dimension), or a vector specifying a value for each dimension.
ll	Lower left limits for the grid. Length should be the number of dimensions of the data provided to <code>locfit.raw()</code> .
ur	Upper right limits for the grid. By default, ll and ur are generated as the bounding box for the data.

Examples

```
data(ethanol, package="locfit")
plot.eval(locfit(NOx ~ lp(E, C, scale=TRUE), data=ethanol, ev=lfgrid()))
```

lfknots	<i>Extraction of fit-point information from a Locfit object.</i>
---------	--

Description

Extracts information, such as fitted values, influence functions from a "locfit" object.

Usage

```
lfknots(x, tr, what = c("x", "coef", "h", "nlx"), delete.pv = TRUE)
```

Arguments

x	Fitted object from <code>locfit()</code> .
tr	Back transformation. Default is the invers link function from the Locfit object.
what	What to return; default is <code>c("x", "coef", "h", "nlx")</code> . Allowed fields are <code>x</code> (fit points); <code>coef</code> (fitted values); <code>f1</code> (local slope); <code>nlx</code> (length of the weight diagram); <code>nlx1</code> (estimated derivative of <code>nlx</code>); <code>se</code> (standard errors); <code>infl</code> (influence function); <code>infla</code> (slope of influence function); <code>lik</code> (maximixed local log-likelihood and local degrees of freedom); <code>h</code> (bandwidth) and <code>deg</code> (degree of fit).
delete.pv	If T, pseudo-vertices are deleted.

Value

A matrix with one row for each fit point. Columns correspond to the specified `what` vector; some fields contribute multiple columns.

lflim	<i>Construct Limit Vectors for Locfit fits.</i>
-------	---

Description

This function is used internally to interpret `xlim` and `flim` arguments. It should not be called directly.

Usage

```
lflim(limits, nm, ret)
```

Arguments

limits	Limit argument.
nm	Variable names.
ret	Initial return vector.

Value

Vector with length $2 \cdot \text{dim}$.

See Also

[locfit](#)

lfmarg	<i>Generate grid margins.</i>
--------	-------------------------------

Description

This function is usually called by [plot.locfit](#).

Usage

```
lfmarg(xlim, m = 40)
```

Arguments

xlim	Vector of limits for the grid. Should be of length $2 \cdot d$; the first d components represent the lower left corner, and the next d components the upper right corner. Can also be a "locfit" object.
m	Number of points for each grid margin. Can be a vector of length d .

Value

A list, whose components are the d grid margins.

See Also

[locfit](#), [plot.locfit](#)

lines.locfit	<i>Add locfit line to existing plot</i>
--------------	---

Description

Adds a Locfit line to an existing plot. `llines` is for use within a panel function for Lattice.

Usage

```
## S3 method for class 'locfit':
lines(x, m=100, tr=x$trans, ...)
## S3 method for class 'locfit':
llines(x, m=100, tr=x$trans, ...)
```

Arguments

<code>x</code>	<code>locfit</code> object. Should be a model with one predictor.
<code>m</code>	Number of points to evaluate the line at.
<code>tr</code>	Transformation function to use for plotting. Default is the inverse link function, or the identity function if derivatives are required.
<code>...</code>	Other arguments to the default <code>lines</code> function.

See Also

`locfit`, `plot.locfit`, `lines`

<code>livmet</code>	<i>liver Metastases dataset</i>
---------------------	---------------------------------

Description

Survival times for 622 patients diagnosed with Liver Metastases.

Beware, the censoring variable is coded as 1 = uncensored, so use `cens=1-z` in `locfit()` calls.

Usage

```
data(livmet)
```

Format

Data frame with survival times (`t`), censoring indicator (`z`) and a number of covariates.

Source

Haupt and Mansmann (1995)

References

Haupt, G. and Mansmann, U. (1995) CART for Survival Data. Statlib Archive, <http://lib.stat.cmu.edu/S/survcart>.

locfit

*Local Regression, Likelihood and Density Estimation.***Description**

locfit is the model formula-based interface to the Locfit library for fitting local regression and likelihood models.

locfit is implemented as a front-end to `locfit.raw`. See that function for options to control smoothing parameters, fitting family and other aspects of the fit.

Usage

```
locfit(formula, data=sys.frame(sys.parent()), weights=1, cens=0, base=0,
       subset, geth=FALSE, ..., lfproc=locfit.raw)
```

Arguments

formula	Model Formula; e.g. $y \sim lp(x)$ for a regression model; $\sim lp(x)$ for a density estimation model. Use of <code>lp()</code> on the RHS is recommended, especially when non-default smoothing parameters are used.
data	Data Frame.
weights	Prior weights (or sample sizes) for individual observations. This is typically used where observations have unequal variance.
cens	Censoring indicator. 1 (or TRUE) denotes a censored observation. 0 (or FALSE) denotes uncensored.
base	Baseline for local fitting. For local regression models, specifying a base is equivalent to using $y - base$ as the response. But base also works for local likelihood.
subset	Subset observations in the data frame.
geth	Don't use.
...	Other arguments to <code>locfit.raw()</code> (or the <code>lfproc</code>).
lfproc	A processing function to compute the local fit. Default is <code>locfit.raw()</code> . Other choices include <code>locfit.robust()</code> , <code>locfit.censor()</code> and <code>locfit.quasi()</code> .

Value

An object with class "locfit". A standard set of methods for printing, plotting, etc. these objects is provided.

References

Loader, C. (1999). Local Regression and Likelihood. Springer, New York.

See Also[locfit.raw](#)**Examples**

```
# fit and plot a univariate local regression
data(ethanol, package="locfit")
fit <- locfit(NOx ~ E, data=ethanol)
plot(fit, get.data=TRUE)

# a bivariate local regression with smaller smoothing parameter
fit <- locfit(NOx~lp(E,C,nn=0.5,scale=0), data=ethanol)
plot(fit)

# density estimation
data(geyser, data="locfit")
fit <- locfit(~ lp(geyser, nn=0.1, h=0.8))
plot(fit,get.data=TRUE)
```

locfit.censor *Censored Local Regression*

Description

locfit.censor produces local regression estimates for censored data. The basic idea is to use an EM style algorithm, where one alternates between estimating the regression and the true values of censored observations.

locfit.censor is designed as a front end to [locfit.raw](#) with data vectors, or as an intermediary between [locfit](#) and [locfit.raw](#) with a model formula. If you can stand the syntax, the second calling sequence above will be slightly more efficient than the third.

Usage

```
locfit.censor(x, y, cens, ..., iter=3, km=FALSE)
```

Arguments

x	Either a locfit model formula or a numeric vector of the predictor variable.
y	If x is numeric, y gives the response variable.
cens	Logical variable indicating censoring. The coding is 1 or TRUE for censored; 0 or FALSE for uncensored.
...	Other arguments to locfit.raw
iter	Number of EM iterations to perform
km	If km=TRUE, the estimation of censored observations uses the Kaplan-Meier estimate, leading to a local version of the Buckley-James estimate. If km=F, the estimation is based on a normal model (Schmee and Hahn). Beware of claims that B-J is nonparametric; it makes stronger assumptions on the upper tail of survival distributions than most authors care to admit.

Value

locfit object.

References

Buckley, J. and James, I. (1979). Linear Regression with censored data. *Biometrika* 66, 429-436.

Loader, C. (1999). *Local Regression and Likelihood*. Springer, NY (Section 7.2).

Schmee, J. and Hahn, G. J. (1979). A simple method for linear regression analysis with censored data (with discussion). *Technometrics* 21, 417-434.

See Also

[km.mrl](#), [locfit](#), [locfit.raw](#)

Examples

```
data(heart, package="locfit")
fit <- locfit.censor(log10(surv+0.5) ~ age, cens=cens, data=heart)
## Can also be written as:
## Not run: fit <- locfit(log10(surv + 0.5) ~ age, cens=cens, data=heart, lfproc=locfit.cens
with(heart, plotbyfactor(age, 0.5 + surv, cens, ylim=c(0.5, 16000), log="y"))
lines(fit, tr=function(x) 10^x)
```

locfit.matrix

Reconstruct a Locfit model matrix.

Description

Reconstructs the model matrix, and associated variables such as the response, prior weights and censoring indicators, from a `locfit` object. This is used by functions such as `fitted.locfit`; it is not normally called directly. The function will only work properly if the data frame has not been changed since the fit was constructed.

Usage

```
locfit.matrix(fit, data)
```

Arguments

<code>fit</code>	Locfit object
<code>data</code>	Data Frame.

Value

A list with variables `x` (the model matrix); `y` (the response); `w` (prior weights); `sc` (scales); `ce` (censoring indicator) and `base` (baseline fit).

See Also

`locfit`, `fitted.locfit`, `residuals.locfit`

`locfit.quasi` *Local Quasi-Likelihood with global reweighting.*

Description

`locfit.quasi` assumes a specified mean-variance relation, and performs iterative reweighted local regression under this assumption. This is appropriate for local quasi-likelihood models, and is an alternative to specifying a family such as "qpoisson".

`locfit.quasi` is designed as a front end to `locfit.raw` with data vectors, or as an intermediary between `locfit` and `locfit.raw` with a model formula. If you can stand the syntax, the second calling sequence above will be slightly more efficient than the third.

Usage

```
locfit.quasi(x, y, weights, ..., iter=3, var=abs)
```

Arguments

<code>x</code>	Either a <code>locfit</code> model formula or a numeric vector of the predictor variable.
<code>y</code>	If <code>x</code> is numeric, <code>y</code> gives the response variable.
<code>weights</code>	Case weights to use in the fitting.
<code>...</code>	Other arguments to <code>locfit.raw</code>
<code>iter</code>	Number of EM iterations to perform
<code>var</code>	Function specifying the assumed relation between the mean and variance.

Value

"locfit" object.

See Also

`locfit`, `locfit.raw`

locfit.raw

*Local Regression, Likelihood and Density Estimation.***Description**

`locfit.raw` is an interface to `Locfit` using numeric vectors (for a model-formula based interface, use `locfit`). Although this function has a large number of arguments, most users are likely to need only a small subset.

The first set of arguments (`x`, `y`, `weights`, `cens`, and `base`) specify the regression variables and associated quantities.

Another set (`scale`, `alpha`, `deg`, `kern`, `kt`, `acri` and `basis`) control the amount of smoothing: bandwidth, smoothing weights and the local model. Most of these arguments are deprecated - they'll currently still work, but should be provided through the `lp()` model term instead.

`deriv` and `dc` relate to derivative (or local slope) estimation.

`family` and `link` specify the likelihood family.

`xlim` and `renorm` may be used in density estimation.

`ev` specifies the evaluation structure or set of evaluation points.

`maxk`, `itype`, `mint`, `maxit` and `debug` control the `Locfit` algorithms, and will be rarely used.

`geth` and `sty` are used by other functions calling `locfit.raw`, and should not be used directly.

Usage

```
locfit.raw(x, y, weights=1, cens=0, base=0,
  scale=FALSE, alpha=0.7, deg=2, kern="tricube", kt="sph",
  acri="none", basis=list(NULL),
  deriv=numeric(0), dc=FALSE,
  family, link="default",
  xlim, renorm=FALSE,
  ev=rbox(),
  maxk=100, itype="default", mint=20, maxit=20, debug=0,
  geth=FALSE, sty="none")
```

Arguments

<code>x</code>	Vector (or matrix) of the independent variable(s). Can be constructed using the <code>lp()</code> function.
<code>y</code>	Response variable for regression models. For density families, <code>y</code> can be omitted.
<code>weights</code>	Prior weights for observations (reciprocal of variance, or sample size).
<code>cens</code>	Censoring indicators for hazard rate or censored regression. The coding is 1 (or TRUE) for a censored observation, and 0 (or FALSE) for uncensored observations.

base	Baseline parameter estimate. If provided, the local regression model is fitted as $Y_i = b_i + m(x_i) + \epsilon_i$, with Locfit estimating the $m(x)$ term. For regression models, this effectively subtracts b_i from Y_i . The advantage of the <code>base</code> formulation is that it extends to likelihood regression models.
scale	Deprecated - see <code>lp()</code> .
alpha	Deprecated - see <code>lp()</code> . A single number (e.g. <code>alpha=0.7</code>) is interpreted as a nearest neighbor fraction. With two components (e.g. <code>alpha=c(0.7, 1.2)</code>), the first component is a nearest neighbor fraction, and the second component is a fixed component. A third component is the penalty term in locally adaptive smoothing.
deg	Degree of local polynomial. Deprecated - see <code>lp()</code> .
kern	Weight function, default = <code>"tcub"</code> . Other choices are <code>"rect"</code> , <code>"trwt"</code> , <code>"tria"</code> , <code>"epan"</code> , <code>"bisq"</code> and <code>"gauss"</code> . Choices may be restricted when derivatives are required; e.g. for confidence bands and some bandwidth selectors.
kt	Kernel type, <code>"sph"</code> (default); <code>"prod"</code> . In multivariate problems, <code>"prod"</code> uses a simplified product model which speeds up computations.
acri	Deprecated - see <code>lp()</code> .
basis	User-specified basis functions. See <code>lfbas</code> for more details on this argument.
deriv	Derivative estimation. If <code>deriv=1</code> , the returned fit will be estimating the derivative (or more correctly, an estimate of the local slope). If <code>deriv=c(1, 1)</code> the second order derivative is estimated. <code>deriv=2</code> is for the partial derivative, with respect to the second variable, in multivariate settings.
dc	Derivative adjustment.
family	Local likelihood family; <code>"gaussian"</code> ; <code>"binomial"</code> ; <code>"poisson"</code> ; <code>"gamma"</code> and <code>"geom"</code> . Density and rate estimation families are <code>"dens"</code> , <code>"rate"</code> and <code>"hazard"</code> (hazard rate). If the family is preceded by a <code>'q'</code> (for example, <code>family="qbinomial"</code>), quasi-likelihood variance estimates are used. Otherwise, the residual variance (<code>rv</code>) is fixed at 1. The default family is <code>"qgauss"</code> if a response <code>y</code> is provided; <code>"density"</code> if no response is provided.
link	Link function for local likelihood fitting. Depending on the family, choices may be <code>"ident"</code> , <code>"log"</code> , <code>"logit"</code> , <code>"inverse"</code> , <code>"sqrt"</code> and <code>"arcsin"</code> .
xlim	For density estimation, Locfit allows the density to be supported on a bounded interval (or rectangle, in more than one dimension). The format should be <code>c(ll, ul)</code> where <code>ll</code> is a vector of the lower bounds and <code>ur</code> the upper bounds. Bounds such as $[0, \infty)$ are not supported, but can be effectively implemented by specifying a very large upper bound.
renorm	Local likelihood density estimates may not integrate exactly to 1. If <code>renorm=T</code> , the integral will be estimated numerically and the estimate rescaled. Presently this is implemented only in one dimension.
ev	The evaluation structure, <code>rbox()</code> for tree structures; <code>lfgrid()</code> for grids; <code>dat()</code> for data points; <code>none()</code> for none. A vector or matrix of evaluation points can also be provided, although in this case you may prefer to use the <code>smooth.lf()</code> interface to Locfit. Note that arguments <code>flim</code> , <code>mg</code> and <code>cut</code> are now given as arguments to the evaluation structure function, rather than to <code>locfit.raw()</code> directly (change effective 12/2001).

maxk	Controls space assignment for evaluation structures. For the adaptive evaluation structures, it is impossible to be sure in advance how many vertices will be generated. If you get warnings about ‘Insufficient vertex space’, Locfit’s default assignment can be increased by increasing maxk. The default is maxk=100.
itype	Integration type for density estimation. Available methods include "prod", "mult" and "mlin"; and "haz" for hazard rate estimation problems. The available integration methods depend on model specification (e.g. dimension, degree of fit). By default, the best available method is used.
mint	Points for numerical integration rules. Default 20.
maxit	Maximum iterations for local likelihood estimation. Default 20.
debug	If > 0; prints out some debugging information.
geth	Don’t use!
sty	Deprecated - see <code>lp()</code> .

Value

An object with class "locfit". A standard set of methods for printing, plotting, etc. these objects is provided.

References

Consult the Web page <http://www.locfit.info/>.

locfit.robust	<i>Robust Local Regression</i>
---------------	--------------------------------

Description

`locfit.robust` implements a robust local regression where outliers are iteratively identified and downweighted, similarly to the lowess method (Cleveland, 1979). The iterations and scale estimation are performed on a global basis.

The scale estimate is 6 times the median absolute residual, while the robust downweighting uses the bisquare function. These are performed in the S code so easily changed.

This can be interpreted as an extension of M estimation to local regression. An alternative extension (implemented in locfit via `family="qrgauss"`) performs the iteration and scale estimation on a local basis.

Usage

```
locfit.robust(x, y, weights, ..., iter=3)
```

Arguments

<code>x</code>	Either a <code>locfit</code> model formula or a numeric vector of the predictor variable.
<code>y</code>	If <code>x</code> is numeric, <code>y</code> gives the response variable.
<code>weights</code>	weights to use in the fitting.
<code>...</code>	Other arguments to <code>locfit.raw</code> .
<code>iter</code>	Number of iterations to perform

Value

"locfit" object.

References

Cleveland, W. S. (1979). Robust locally weighted regression and smoothing scatterplots. *J. Amer. Statist. Assn.* 74, 829-836.

See Also

`locfit`, `locfit.raw`

lp

Local Polynomial Model Term

Description

`lp` is a local polynomial model term for `Locfit` models. Usually, it will be the only term on the RHS of the model formula.

Smoothing parameters should be provided as arguments to `lp()`, rather than to `locfit()`.

Usage

```
lp(..., nn, h, adpen, deg, acri, scale, style)
```

Arguments

<code>...</code>	Predictor variables for the local regression model.
<code>nn</code>	Nearest neighbor component of the smoothing parameter. Default value is 0.7, unless either <code>h</code> or <code>adpen</code> are provided, in which case the default is 0.
<code>h</code>	The constant component of the smoothing parameter. Default: 0.
<code>adpen</code>	Penalty parameter for adaptive fitting.
<code>deg</code>	Degree of polynomial to use.
<code>acri</code>	Criterion for adaptive bandwidth selection.
<code>style</code>	Style for special terms (<code>left</code> , <code>ang</code> e.t.c.). Do not try to set this directly; call <code>locfit</code> instead.

`scale` A scale to apply to each variable. This is especially important for multivariate fitting, where variables may be measured in non-comparable units. It is also used to specify the frequency for `ang` terms. If `scale=F` (the default) no scaling is performed. If `scale=T`, marginal standard deviations are used. Alternatively, a numeric vector can provide scales for the individual variables.

See Also

`locfit`, `locfit.raw`

Examples

```
data(ethanol, package="locfit")
# fit with 50
fit <- locfit(NOx~lp(E, nn=0.5), data=ethanol)
# bivariate fit.
fit <- locfit(NOx~lp(E, C, scale=TRUE), data=ethanol)

# density estimation
data(geyser, package="locfit")
fit <- locfit.raw(lp(geyser, nn=0.1, h=0.8))
```

`lscv`

Least Squares Cross Validation Statistic.

Description

The calling sequence for `lscv` matches those for the `locfit` or `locfit.raw` functions. Note that this function is only designed for density estimation in one dimension. The returned object contains the least squares cross validation score for the fit.

The computation of $\int \hat{f}(x)^2 dx$ is performed numerically. For kernel density estimation, this is unlikely to agree exactly with other LSCV routines, which may perform the integration analytically.

Usage

```
lscv(x, ..., exact=FALSE)
```

Arguments

`x` model formula (or numeric vector, if `exact=T`)

`...` other arguments to `locfit` or `lscv.exact`

`exact` By default, the computation is approximate. If `exact=TRUE`, exact computation using `lscv.exact` is performed. This uses kernel density estimation with a constant bandwidth.

Value

A vector consisting of the LSCV statistic and fitted degrees of freedom.

See Also

[locfit](#), [locfit.raw](#), [lscv.exact](#) [lscvplot](#)

Examples

```
# approximate calculation for a kernel density estimate
data(geyser, package="locfit")
lscv(~lp(geyser, h=1, deg=0), ev=lfgrid(100, ll=1, ur=6), kern="gauss")
# same computation, exact
lscv(lp(geyser, h=1), exact=TRUE)
```

`lscv.exact`

Exact LSCV Calculation

Description

This function performs the exact computation of the least squares cross validation statistic for one-dimensional kernel density estimation and a constant bandwidth.

At the time of writing, it is implemented only for the Gaussian kernel (with the standard deviation of 0.4; Locfit's standard).

Usage

```
lscv.exact(x, h=0)
```

Arguments

<code>x</code>	Numeric data vector.
<code>h</code>	The bandwidth. If <code>x</code> is constructed with <code>lp()</code> , the bandwidth should be given there instead.

Value

A vector of the LSCV statistic and the fitted degrees of freedom.

See Also

[lscv](#), [lscvplot](#)

Examples

```
data(geyser, package="locfit")
lscv.exact(lp(geyser, h=0.25))
# equivalent form using lscv
lscv(lp(geyser, h=0.25), exact=TRUE)
```

lscvplot	<i>Compute the LSCV plot.</i>
----------	-------------------------------

Description

The `lscvplot` function loops through calls to the `lscv` function (and hence to `link{locfit}`), using a different smoothing parameter for each call. The returned structure contains the LSCV statistic for each density estimate, and can be used to produce an LSCV plot.

Usage

```
lscvplot(..., alpha)
```

Arguments

...	arguments to the <code>lscv</code> , <code>locfit</code> functions.
alpha	Matrix of smoothing parameters. The <code>lscvplot</code> function loops through calls to <code>lscv</code> , using each row of <code>alpha</code> as the smoothing parameter in turn. If <code>alpha</code> is provided as a vector, it will be converted to a one-column matrix, thus interpreting each component as a nearest neighbor smoothing parameter.

Value

An object with class "gcvplot", containing the smoothing parameters and LSCV scores. The actual plot is produced using `plot.gcvplot`.

See Also

[locfit](#), [locfit.raw](#), [gcv](#), [lscv](#), [plot.gcvplot](#)

mcyC	<i>Acc(De?)celeration of a Motorcycle Hitting a Wall</i>
------	--

Description

Measurements of the acceleration of a motorcycle as it hits a wall. Actually, rumored to be a concatenation of several such datasets.

Usage

```
data(mcyC)
```

Format

Data frame with time and accel variables.

Source

Härdle (1990).

References

Härdle, W. (1990). Applied Nonparametric Regression. Cambridge University Press.

mine

Fracture Counts in Coal Mines

Description

The number of fractures in the upper seam of coal mines, and four predictor variables. This dataset can be modeled using Poisson regression.

Usage

`data(mine)`

Format

A dataframe with the response `frac`, and predictor variables `extrp`, `time`, `seamh` and `inb`.

Source

Myers (1990).

References

Myers, R. H. (1990). Classical and Modern Regression with Applications (Second edition). PWS-Kent Publishing, Boston.

mmsamp

Test dataset for minimax Local Regression

Description

50 observations, as used in Figure 13.1 of Loader (1999).

Usage

`data(cltest)`

Format

Data Frame with x and y variables.

References

Loader, C. (1999). Local Regression and Likelihood. Springer, New York.

morths

Henderson and Sheppard Mortality Dataset

Description

Observed mortality for 55 to 99.

Usage

```
data(morths)
```

Format

Data frame with age, n and number of deaths.

Source

Henderson and Sheppard (1919).

References

Henderson, R. and Sheppard, H. N. (1919). Graduation of mortality and other tables. Actuarial Society of America, New York.

none

Locfit Evaluation Structure

Description

`none()` is an evaluation structure for `locfit.raw()`, specifying no evaluation points. Only the initial parametric fit is computed - this is the easiest and most efficient way to coerce Locfit into producing a parametric regression fit.

Usage

```
none()
```

Examples

```
data(ethanol, package="locfit")
# fit a fourth degree polynomial using locfit
fit <- locfit(NOx~E, data=ethanol, deg=4, ev=none())
plot(fit, get.data=TRUE)
```

penny

Penny Thickness Dataset

Description

For each year, 1945 to 1989, the thickness of two U.S. pennies was recorded.

Usage

```
data(penny)
```

Format

A dataframe.

Source

Scott (1992).

References

Scott (1992). Multivariate Density Estimation. Wiley, New York.

plot.eval

Plot evaluation points from a 2-d locfit object.

Description

This function is used to plot the evaluation structure generated by Locfit for a two dimensional fit. Vertices of the tree structure are displayed as `O`; pseudo-vertices as `*`.

Usage

```
plot.eval(x, add=FALSE, text=FALSE, ...)
```

Arguments

<code>x</code>	"locfit" object.
<code>add</code>	If TRUE, add to existing plot.
<code>text</code>	If TRUE, numbers will be added indicating the order points were added.
<code>...</code>	Arguments passed to and from other methods.

See Also

[locfit.](#)

Examples

```
data(ethanol, package="locfit")
fit <- locfit(NOx ~ E + C, data=ethanol, scale=0)
plot.eval(fit)
```

plot.gcvplot	<i>Produce a cross-validation plot.</i>
--------------	---

Description

Plots the value of the GCV (or other statistic) in a `gcvplot` object against the degrees of freedom of the fit.

Usage

```
## S3 method for class 'gcvplot':
plot(x, xlab = "Fitted DF", ylab = x$cri, ...)
```

Arguments

<code>x</code>	A <code>gcvplot</code> object, produced by <code>gcvplot</code> , <code>aicplot</code> etc.
<code>xlab</code>	Text label for the x axis.
<code>ylab</code>	Text label for the y axis.
<code>...</code>	Other arguments to <code>plot</code> .

See Also

[locfit](#), [locfit.raw](#), [gcv](#), [aicplot](#), [cpplot](#), [gcvplot](#), [lcvplot](#)

Examples

```
data(ethanol)
plot(gcvplot(NOx~E, data=ethanol, alpha=seq(0.2, 1.0, by=0.05)))
```

plot.lfeval *Plot a Locfit Evaluation Structure.*

Description

Plots the evaluation points from a `locfit` or `lfeval` structure, for one- or two-dimensional fits.

Usage

```
## S3 method for class 'lfeval':
plot(x, add=FALSE, txt=FALSE, ...)
```

Arguments

<code>x</code>	A <code>lfeval</code> or <code>locfit</code> object
<code>add</code>	If TRUE, the points will be added to the existing plot. Otherwise, a new plot is created.
<code>txt</code>	If TRUE, the points are annotated with numbers in the order they were entered into the fit.
<code>...</code>	Additional graphical parameters.

Value

"lfeval" object.

See Also

[lfeval](#), [locfit](#), [print.lfeval](#)

plot.locfit *Plot an object of class locfit.*

Description

The `plot.locfit` function generates grids of plotting points, followed by a call to [preplot.locfit](#). The returned object is then passed to [plot.locfit.1d](#), [plot.locfit.2d](#) or [plot.locfit.3d](#) as appropriate.

Usage

```
## S3 method for class 'locfit':
plot(x, xlim, pv, tv, m, mtv=6, band="none", tr=NULL,
     what = "coef", get.data=FALSE, f3d=(d == 2) && (length(tv) > 0), ...)
```

Arguments

<code>x</code>	locfit object.
<code>xlim</code>	Plotting limits. Eg. <code>xlim=c(0,0,1,1)</code> plots over the unit square in two dimensions. Default is bounding box of the data.
<code>pv</code>	Panel variables, to be varied within each panel of a plot. May be specified as a character vector, or variable numbers. There must be one or two panel variables; default is all variables in one or two dimensions; Variable 1 in three or more dimensions. May be specified using either variable numbers or names.
<code>tv</code>	Trellis variables, to be varied from panel to panel of the plot.
<code>m</code>	Controls the plot resolution (within panels, for trellis displays). Default is 100 points in one dimension; 40 points (per dimension) in two or more dimensions.
<code>mtv</code>	Number of points for trellis variables; default 6.
<code>band</code>	Type of confidence bands to add to the plot. Default is "none". Other choices include "global" for bands using a global variance estimate; "local" for bands using a local variance estimate and "pred" for prediction bands (at present, using a global variance estimate). To obtain the global variance estimate for a fit, use <code>rv</code> . This can be changed with <code>rv<-</code> . Confidence bands, by default, are 95%, based on normal approximations and neglecting bias. To change the critical value or confidence level, or to obtain simultaneous instead of pointwise confidence, the critical value stored on the fit must be changed. See the <code>kappa0</code> and <code>crit</code> functions.
<code>tr</code>	Transformation function to use for plotting. Default is the inverse link function, or the identity function if derivatives are requested.
<code>what</code>	What to plot. See <code>predict.locfit</code> .
<code>get.data</code>	If TRUE, original data is added to the plot. Default: FALSE.
<code>f3d</code>	Force the <code>locfit.3d</code> class on the prediction object, thereby generating a trellis style plot. Default: FALSE, unless a <code>tv</code> argument is provided. Not available in R.
<code>...</code>	Other arguments to <code>plot.locfit.1d</code> , <code>plot.locfit.2d</code> or <code>plot.locfit.3d</code> as appropriate.

See Also

[locfit](#), [plot.locfit.1d](#), [plot.locfit.2d](#), [plot.locfit.3d](#), [lines.locfit](#), [predict.locfit](#), [preplot.locfit](#)

Examples

```
x <- rnorm(100)
y <- dnorm(x) + rnorm(100) / 5
plot(locfit(y~x), band="global")
x <- cbind(rnorm(100), rnorm(100))
plot(locfit(~x), type="persp")
```

`plot.locfit.1d` *Plot a one dimensional preplot.locfit object.*

Description

This function is not usually called directly. It will be called automatically when plotting a one-dimensional `locfit` or `preplot.locfit` object.

Usage

```
plot.locfit.1d(x, add=FALSE, main="", xlab="default", ylab=x$yname,
              type="l", ylim, lty=1, col=1, ...)
```

Arguments

`x` One dimensional `preplot.locfit` object.
`add` If TRUE, the plot will be added to the existing plot.
`main, xlab, ylab, type, ylim, lty, col`
 Graphical parameters passed on to `plot` (only if `add=FALSE`).
`...` Additional graphical parameters to the `plot` function (only if `add=FALSE`).

See Also

[locfit](#), [plot.locfit](#), [preplot.locfit](#)

`plot.locfit.2d` *Plot a two-dimensional "preplot.locfit" object.*

Description

This function is not usually called directly. It will be called automatically when plotting one-dimensional `locfit` or `preplot.locfit` objects.

Usage

```
plot.locfit.2d(x, type="contour", main, xlab, ylab, zlab=x$yname, ...)
```

Arguments

`x` Two dimensional `preplot.locfit` object.
`type` one of "contour", "persp", or "image".
`main` title for the plot.
`xlab, ylab` text labels for the x- and y-axes.
`zlab` if `type="persp"`, the label for the z-axis.
`...` Additional arguments to the `contour`, `persp` or `image` functions.

See Also

locfit, plot.locfit, preplot.locfit

plot.locfit.3d *Plot a high-dimensional "preplot.locfit" object using trellis displays.*

Description

This function plots cross-sections of a Locfit model (usually in three or more dimensions) using trellis displays. It is not usually called directly, but is invoked by `plot.locfit`.

The R libraries `lattice` and `grid` provide a partial (at time of writing) implementation of trellis. Currently, this works with one panel variable.

Usage

```
plot.locfit.3d(x, main="", pv, tv, type = "level", pred.lab = x$vnames,
              resp.lab=x$yname, crit = 1.96, ...)
```

Arguments

<code>x</code>	"preplot.locfit" object.
<code>main</code>	title for the plot.
<code>pv</code>	Panel variables. These are the variables (either one or two) that are varied within each panel of the display.
<code>tv</code>	Trellis variables. These are varied from panel to panel of the display.
<code>type</code>	Type of display. When there are two panel variables, the choices are "contour", "level" and "persp".
<code>pred.lab</code>	label for the predictor variable.
<code>resp.lab</code>	label for the response variable.
<code>crit</code>	critical value for the confidence level.
<code>...</code>	graphical parameters passed to <code>xyplot</code> or <code>contourplot</code> .

See Also

plot.locfit, preplot.locfit

```
plot.preplot.locfit
```

Plot a "preplot.locfit" object.

Description

The `plot.locfit()` function is implemented, roughly, as a call to `preplot.locfit()`, followed by a call to `plot.locfitpred()`. For most users, there will be little need to call `plot.locfitpred()` directly.

Usage

```
## S3 method for class 'preplot.locfit':
plot(x, ...)
```

Arguments

<code>x</code>	A <code>preplot.locfit</code> object, produced by <code>preplot.locfit()</code> .
<code>...</code>	Other arguments to <code>plot.locfit.1d</code> , <code>plot.locfit.2d</code> or <code>plot.locfit.3d</code> as appropriate.

See Also

[locfit](#), [plot.locfit](#), [preplot.locfit](#), [plot.locfit.1d](#), [plot.locfit.2d](#), [plot.locfit.3d](#).

```
plot.scb
```

Plot method for simultaneous confidence bands

Description

Plot method for simultaneous confidence bands created by the `scb` function.

Usage

```
plot.scb(x, add=FALSE, ...)
```

Arguments

<code>x</code>	<code>scb</code> object created by <code>scb</code> .
<code>add</code>	If TRUE, bands will be added to the existing plot.
<code>...</code>	Arguments passed to and from other methods.

See Also

[scb](#)

Examples

```
# corrected confidence bands for a linear logistic model
data(insect)
fit <- scb(deaths ~ lconc, type=4, w=nins, data=insect,
           deg=1, family="binomial", kern="parm")
plot(fit)
```

plotbyfactor *x-y scatterplot, colored by levels of a factor.*

Description

Produces a scatter plot of x-y data, with different classes given by a factor f. The different classes are identified by different colours and/or symbols.

Usage

```
plotbyfactor(x, y, f, data, col = 1:10, pch = "O", add = FALSE, lg,
             xlab = deparse(substitute(x)), ylab = deparse(substitute(y)),
             log = "", ...)
```

Arguments

x	Variable for x axis.
y	Variable for y axis.
f	Factor (or variable for which as.factor() works).
data	data frame for variables x, y, f. Default: sys.parent().
col	Color numbers to use in plot. Will be replicated if shorter than the number of levels of the factor f. Default: 1:10.
pch	Vector of plot characters. Replicated if necessary. Default: "O".
add	If TRUE, add to existing plot. Otherwise, create new plot.
lg	Coordinates to place a legend. Default: Missing (no legend).
xlab, ylab	Axes labels.
log	Should the axes be in log scale? Use "x", "y", or "xy" to specify which axis to be in log scale.
...	Other graphical parameters, labels, titles e.t.c.

Examples

```
data(iris)
plotbyfactor(petal.wid, petal.len, species, data=iris)
```

`points.locfit` *Add 'locfit' points to existing plot*

Description

This function shows the points at which the local fit was computed directly, rather than being interpolated. This can be useful if one is unsure of the validity of interpolation.

Usage

```
## S3 method for class 'locfit':
points(x, tr, ...)
```

Arguments

`x` "locfit" object. Should be a model with one predictor.
`tr` Back transformation.
`...` Other arguments to the default `points` function.

See Also

`locfit`, `plot.locfit`, `points`

`predict.locfit` *Prediction from a Locfit object.*

Description

The `locfit` function computes a local fit at a selected set of points (as defined by the `ev` argument). The `predict.locfit` function is used to interpolate from these points to any other points. The method is based on cubic hermite polynomial interpolation, using the estimates and local slopes at each fit point.

The motivation for this two-step procedure is computational speed. Depending on the sample size, dimension and fitting procedure, the local fitting method can be expensive, and it is desirable to keep the number of points at which the direct fit is computed to a minimum. The interpolation method used by `predict.locfit()` is usually much faster, and can be computed at larger numbers of points.

Usage

```
## S3 method for class 'locfit':
predict(object, newdata=NULL, where = "fitp",
        se.fit=FALSE, band="none", what="coef", ...)
```

Arguments

object	Fitted object from <code>locfit()</code> .
newdata	Points to predict at. Can be given in several forms: vector/matrix; list, data frame.
se.fit	If TRUE, standard errors are computed along with the fitted values.
where, what, band	arguments passed on to <code>preplot.locfit</code> .
...	Additional arguments to <code>preplot.locfit</code> .

Value

If `se.fit=F`, a numeric vector of predictors. If `se.fit=T`, a list with components `fit`, `se.fit` and `residual.scale`.

Examples

```
data(ethanol, package="locfit")
fit <- locfit(NOx ~ E, data=ethanol)
predict(fit, c(0.6, 0.8, 1.0))
```

```
preplot.locfit
```

Prediction from a Locfit object.

Description

`preplot.locfit` can be called directly, although it is more usual to call `plot.locfit` or `predict.locfit`. The advantage of `preplot.locfit` is in S-Plus 5, where arithmetic and transformations can be performed on the "preplot.locfit" object.

`plot(preplot(fit))` is essentially synonymous with `plot(fit)`.

Usage

```
## S3 method for class 'locfit':
preplot(object, newdata=NULL, where, tr=NULL, what="coef",
        band="none", get.data=FALSE, f3d=FALSE, ...)
```

Arguments

object	Fitted object from <code>locfit()</code> .
newdata	Points to predict at. Can be given in several forms: vector/matrix; list, data frame.
where	An alternative to <code>newdata</code> . Choices include "grid" for the grid <code>lfmarg(object)</code> ; "data" for the original data points and "fitp" for the direct fitting points (ie. no interpolation).
tr	Transformation for likelihood models. Default is the inverse of the link function.

what	What to compute predicted values of. The default, <code>what="coef"</code> , works with the fitted curve itself. Other choices include <code>"nlx"</code> for the length of the weight diagram; <code>"infl"</code> for the influence function; <code>"band"</code> for the bandwidth; <code>"degr"</code> for the local polynomial degree; <code>"lik"</code> for the maximized local likelihood; <code>"rdf"</code> for the local residual degrees of freedom and <code>"vari"</code> for the variance function. The interpolation algorithm for some of these quantities is questionable.
band	Compute standard errors for the fit and include confidence bands on the returned object. Default is <code>"none"</code> . Other choices include <code>"global"</code> for bands using a global variance estimate; <code>"local"</code> for bands using a local variance estimate and <code>"pred"</code> for prediction bands (at present, using a global variance estimate). To obtain the global variance estimate for a fit, use <code>rv</code> . This can be changed with <code>rv<-</code> . Confidence bands, by default, are 95%, based on normal approximations and neglecting bias. To change the critical value or confidence level, or to obtain simultaneous instead of pointwise confidence, the critical value stored on the fit must be changed. See the <code>kappa0</code> and <code>crit</code> functions.
get.data	If TRUE, the original data is attached to the returned object, and added to the plot.
f3d	If TRUE, sets a flag that forces plotting using the trellis style. Not available in R.
...	arguments passed to and from other methods.

Value

An object with class `"preplot.locfit"`, containing the predicted values and additional information used to construct the plot.

See Also

`locfit`, `predict.locfit`, `plot.locfit`.

`preplot.locfit.raw` *Prediction from a Locfit object.*

Description

`preplot.locfit.raw` is an internal function used by `predict.locfit` and `preplot.locfit`. It should not normally be called directly.

Usage

```
## S3 method for class 'locfit.raw':
preplot(object, newdata, where, what, band, ...)
```

Arguments

object	Fitted object from <code>locfit()</code> .
newdata	New data points.
where	Type of data provided in <code>newdata</code> .
what	What to compute predicted values of.
band	Compute standard errors for the fit and include confidence bands on the returned object.
...	Arguments passed to and from other methods.

Value

A list containing raw output from the internal prediction routines.

See Also

[locfit](#), [predict.locfit](#), [preplot.locfit](#).

print.gcvplot	<i>Print method for gcvplot objects</i>
---------------	---

Description

Print method for "gcvplot" objects. Actually, equivalent to `plot.gcvplot()`. `scb` function.

Usage

```
print.gcvplot(x, ...)
```

Arguments

x	gcvplot object.
...	Arguments passed to and from other methods.

See Also

[gcvplot](#), [plot.gcvplot](#) [summary.gcvplot](#)

`print.lfeval` *Print the Locfit Evaluation Points.*

Description

Prints a matrix of the evaluation points from a `locfit` or `lfeval` structure.

Usage

```
print.lfeval(x, ...)
```

Arguments

`x` A `lfeval` or `locfit` object
`...` Arguments passed to and from other methods.

Value

Matrix of the fit points.

See Also

[lfeval](#), [locfit](#), [plot.lfeval](#)

`print.locfit` *Print method for "locfit" object.*

Description

Prints a short summary of a "locfit" object.

Usage

```
## S3 method for class 'locfit':  
print(x, ...)
```

Arguments

`x` `locfit` object.
`...` Arguments passed to and from other methods.

See Also

[locfit](#)

`print.preplot.locfit`

Print method for preplot.locfit objects.

Description

Print method for objects created by the `preplot.locfit` function.

Usage

```
print.preplot.locfit(x, ...)
```

Arguments

`x` "preplot.locfit" object.
`...` Arguments passed to and from other methods.

See Also

`preplot.locfit`, `predict.locfit`

`print.scb`

Print method for simultaneous confidence bands

Description

Print method for simultaneous confidence bands created by the `scb` function.

Usage

```
print.scb(x, ...)
```

Arguments

`x` "scb" object created by `scb`.
`...` Arguments passed to and from other methods.

See Also

`scb`

References

Loader, C. (1999). Local Regression and Likelihood. Springer, New York.

Cleveland, W. and Grosse, E. (1991). Computational Methods for Local Regression. Statistics and Computing 1.

Examples

```
data(ethanol, package="locfit")
plot.eval(locfit(NOx~E+C, data=ethanol, scale=0, ev=rbox(cut=0.8)))
plot.eval(locfit(NOx~E+C, data=ethanol, scale=0, ev=rbox(cut=0.3)))
```

regband

Bandwidth selectors for local regression.

Description

Function to compute local regression bandwidths for local linear regression, implemented as a front end to `locfit()`.

This function is included for comparative purposes only. Plug-in selectors are based on flawed logic, make unreasonable and restrictive assumptions and do not use the full power of the estimates available in `Locfit`. Any relation between the results produced by this function and desirable estimates are entirely coincidental.

Usage

```
regband(formula, what = c("CP", "GCV", "GKK", "RSW"), deg=1, ...)
```

Arguments

<code>formula</code>	Model Formula (one predictor).
<code>what</code>	Methods to use.
<code>deg</code>	Degree of fit.
<code>...</code>	Other <code>Locfit</code> options.

Value

Vector of selected bandwidths.

`residuals.locfit` *Fitted values and residuals for a Locfit object.*

Description

`residuals.locfit` is implemented as a front-end to `fitted.locfit`, with the `type` argument set.

Usage

```
## S3 method for class 'locfit':
residuals(object, data=NULL, type="deviance", ...)
```

Arguments

<code>object</code>	locfit object.
<code>data</code>	The data frame for the original fit. Usually, shouldn't be needed.
<code>type</code>	Type of fit or residuals to compute. The default is "fit" for <code>fitted.locfit</code> , and "dev" for <code>residuals.locfit</code> . Other choices include "pear" for Pearson residuals; "raw" for raw residuals, "ldot" for likelihood derivative; "d2" for the deviance residual squared; <code>lddot</code> for the likelihood second derivative. Generally, <code>type</code> should only be used when <code>what="coef"</code> .
<code>...</code>	arguments passed to and from other methods.

Value

A numeric vector of the residuals.

`right` *One-sided right smooth for a Locfit model.*

Description

The `right()` function is used in a locfit model formula to specify a one-sided smooth: when fitting at a point x , only data points with $x_i \leq x$ should be used. This can be useful in estimating points of discontinuity, and in cross-validation for forecasting a time series. `right(x)` is equivalent to `lp(x, style="right")`.

When using this function, it will usually be necessary to specify an evaluation structure, since the fit is not smooth and locfit's interpolation methods are unreliable. Also, it is usually best to use `deg=0` or `deg=1`, otherwise the fits may be too variable. If nearest neighbor bandwidth specification is used, it does not recognize `right()`.

Usage

```
right(x, ...)
```

Arguments

`x` numeric variable.
`...` Other arguments to `lp()`.

See Also

[lfbas](#), [locfit](#), [left](#)

Examples

```
# compute left and right smooths
data(penny)
xev <- (1945:1988)+0.5
fitl <- locfit(thickness~left(year,h=10,deg=1), ev=xev, data=penny)
fitr <- locfit(thickness~right(year,h=10,deg=1),ev=xev, data=penny)
# plot the squared difference, to show the change points.
plot( xev, (predict(fitr,where="ev") - predict(fitl,where="ev"))^2 )
```

 rv

Residual variance from a locfit object.

Description

As part of the `locfit` fitting procedure, an estimate of the residual variance is computed; the `rv` function extracts the variance from the `"locfit"` object. The estimate used is the residual sum of squares (or residual deviance, for quasi-likelihood models), divided by the residual degrees of freedom.

For likelihood (not quasi-likelihood) models, the estimate is 1.0.

Usage

```
rv(fit)
```

Arguments

`fit` `"locfit"` object.

Value

Returns the residual variance estimate from the `"locfit"` object.

See Also

[locfit](#), [rv<-](#)

Examples

```
data(ethanol)
fit <- locfit(NOx~E,data=ethanol)
rv(fit)
```

rv

Substitute variance estimate on a locfit object.

Description

By default, Locfit uses the normalized residual sum of squares as the variance estimate when constructing confidence intervals. In some cases, the user may like to use alternative variance estimates; this function allows the default value to be changed.

Usage

```
rv(fit) <- 1.2345
```

Arguments

fit "locfit" object.

See Also

[locfit\(\)](#), [rv\(\)](#), [plot.locfit\(\)](#)

scb

Simultaneous Confidence Bands

Description

scb is implemented as a front-end to [locfit](#), to compute simultaneous confidence bands using the tube formula method and extensions, based on Sun and Loader (1994).

Usage

```
scb(x, ..., ev = lfgrid(20), simul = TRUE, type = 1)
```

Arguments

<code>x</code>	A numeric vector or matrix of predictors (as in <code>locfit.raw</code>), or a model formula (as in <code>locfit</code>).
<code>...</code>	Additional arguments to <code>locfit.raw</code> .
<code>ev</code>	The evaluation structure to use. See <code>locfit.raw</code> .
<code>simul</code>	Should the coverage be simultaneous or pointwise?
<code>type</code>	Type of confidence bands. <code>type=0</code> computes pointwise 95% bands. <code>type=1</code> computes basic simultaneous bands with no corrections. <code>type=2, 3, 4</code> are the centered and corrected bands for parametric regression models listed in Table 3 of Sun, Loader and McCormick (2000).

Value

A list containing the evaluation points, fit, standard deviations and upper and lower confidence bounds. The class is "scb"; methods for printing and plotting are provided.

References

- Sun J. and Loader, C. (1994). Simultaneous confidence bands in linear regression and smoothing. *The Annals of Statistics* 22, 1328-1345.
- Sun, J., Loader, C. and McCormick, W. (2000). Confidence bands in generalized linear models. *The Annals of Statistics* 28, 429-460.

See Also

`locfit`, `print.scb`, `plot.scb`.

Examples

```
# corrected confidence bands for a linear logistic model
data(insect)
fit <- scb(deaths~lp(lconc,deg=1), type=4, w=nins,
          data=insect,family="binomial",kern="parm")
plot(fit)
```

Description

Given a dataset and set of pilot bandwidths, this function computes a bandwidth via the plug-in method, and the assumed 'pilot' relationship of Sheather and Jones (1991). The S-J method chooses the bandwidth at which the two intersect.

The purpose of this function is to demonstrate the sensitivity of plug-in methods to pilot bandwidths and assumptions. This function does not provide a reliable method of bandwidth selection.

Usage

```
sjpi(x, a)
```

Arguments

x	data vector
a	vector of pilot bandwidths

Value

A matrix with four columns; the number of rows equals the length of a. The first column is the plug-in selected bandwidth. The second column is the pilot bandwidths a. The third column is the pilot bandwidth according to the assumed relationship of Sheather and Jones. The fourth column is an intermediate calculation.

References

Sheather, S. J. and Jones, M. C. (1991). A reliable data-based bandwidth selection method for kernel density estimation. JRSS-B 53, 683-690.

See Also

[locfit](#), [locfit.raw](#), [lcvplot](#)

Examples

```
# Fig 10.2 (S-J parts) from Loader (1999).
data(geyser)
gf <- 2.5
a <- seq(0.05, 0.7, length=100)
z <- sjpi(geyser, a)

# the plug-in curve. Multiplying by gf=2.5 corresponds to Locfit's standard
# scaling for the Gaussian kernel.
plot(gf*z[, 2], gf*z[, 1], type = "l", xlab = "Pilot Bandwidth k", ylab
     = "Bandwidth h")

# Add the assumed curve.
lines(gf * z[, 3], gf * z[, 1], lty = 2)
legend(gf*0.05, gf*0.4, lty = 1:2, legend = c("Plug-in", "SJ assumed"))
```

smooth.lf

Local Regression, Likelihood and Density Estimation.

Description

smooth.lf is a simple interface to the Locfit library. The input consists of a predictor vector (or matrix) and response. The output is a list with vectors of fitting points and fitted values. Most [locfit.raw](#) options are valid.

Usage

```
smooth.lf(x, y, xev=x, direct=FALSE, ...)
```

Arguments

<code>x</code>	Vector (or matrix) of the independent variable(s).
<code>y</code>	Response variable. If omitted, <code>x</code> is treated as the response and the predictor variable is <code>1:n</code> .
<code>xev</code>	Fitting Points. Default is the data vector <code>x</code> .
<code>direct</code>	Logical variable. If <code>T</code> , local regression is performed directly at each fitting point. If <code>F</code> , the standard Locfit method combining fitting and interpolation is used.
<code>...</code>	Other arguments to <code>locfit.raw()</code> .

Value

A list with components `x` (fitting points) and `y` (fitted values). Also has a `call` component, so `update()` will work.

See Also

`locfit()`, `locfit.raw()`, `density.lf()`.

Examples

```
# using smooth.lf() to fit a local likelihood model.
data(morths)
fit <- smooth.lf(morths$age, morths$deaths, weights=morths$n,
                 family="binomial")
plot(fit, type="l")

# update with the direct fit
fit1 <- update(fit, direct=TRUE)
lines(fit1, col=2)
print(max(abs(fit$y-fit1$y)))
```

 spence.15

Spencer's 15 point graduation rule.

Description

Spencer's 15 point rule is a weighted moving average operation for a sequence of observations equally spaced in time. The average at time `t` depends on the observations at times `t-7, ..., t+7`.

Except for boundary effects, the function will reproduce polynomials up to degree 3.

Usage

```
spence.15(y)
```

Arguments

`y` Data vector of observations at equally spaced points.

Value

A vector with the same length as the input vector, representing the graduated (smoothed) values.

References

Spencer, J. (1904). On the graduation of rates of sickness and mortality. *Journal of the Institute of Actuaries* 38, 334-343.

See Also

[spence.21](#), [spencer](#),

Examples

```
data(spencer)
yy <- spence.15(spencer$mortality)
plot(spencer$age, spencer$mortality)
lines(spencer$age, yy)
```

`spence.21`

Spencer's 21 point graduation rule.

Description

Spencer's 21 point rule is a weighted moving average operation for a sequence of observations equally spaced in time. The average at time t depends on the observations at times $t-11, \dots, t+11$.

Except for boundary effects, the function will reproduce polynomials up to degree 3.

Usage

```
spence.21(y)
```

Arguments

`y` Data vector of observations at equally spaced points.

Value

A vector with the same length as the input vector, representing the graduated (smoothed) values.

References

Spencer, J. (1904). On the graduation of rates of sickness and mortality. *Journal of the Institute of Actuaries* 38, 334-343.

See Also

[spence.15](#), [spencer](#),

Examples

```
data(spencer)
yy <- spence.21(spencer$mortality)
plot(spencer$age, spencer$mortality)
lines(spencer$age, yy)
```

spencer

Spencer's Mortality Dataset

Description

Observed mortality rates for ages 20 to 45.

Usage

```
data(spencer)
```

Format

Data frame with age and mortality variables.

Source

Spencer (1904).

References

Spencer, J. (1904). On the graduation of rates of sickness and mortality. *Journal of the Institute of Actuaries* 38, 334-343.

stamp

Stamp Thickness Dataset

Description

Thicknesses of 482 postage stamps of the 1872 Hidalgo issue of Mexico.

Usage

```
data(stamp)
```

Format

Data frame with `thick` (stamp thickness) and `count` (number of stamps) variables.

Source

Izenman and Sommer (1988).

References

Izenman, A. J. and Sommer, C. J. (1988). Philatelic mixtures and multimodal densities. *Journal of the American Statistical Association* 73, 602-606.

<code>store</code>	<i>Save S functions.</i>
--------------------	--------------------------

Description

I've gotta keep track of this mess somehow!

Usage

```
store(data=FALSE, grand=FALSE)
```

Arguments

<code>data</code>	whether datasets are to be dumped.
<code>grand</code>	whether <code>grand</code> is to be dumped.

<code>summary.gcvplot</code>	<i>Summary method for a gcvplot structure.</i>
------------------------------	--

Description

Computes a short summary for a generalized cross-validation plot structure

Usage

```
## S3 method for class 'gcvplot':
summary(object, ...)
```

Arguments

<code>object</code>	A <code>gcvplot</code> structure produced by a call to <code>gcvplot</code> , <code>cpplot</code> e.t.c.
<code>...</code>	arguments to and from other methods.

Value

A matrix with two columns; one row for each fit computed in the `gcvplot` call. The first column is the fitted degrees of freedom; the second is the GCV or other criterion computed.

See Also

`locfit`, `gcv`, `gcvplot`

Examples

```
data(ethanol)
summary(gcvplot(NOx~E, data=ethanol, alpha=seq(0.2, 1.0, by=0.05)))
```

`summary.locfit` *Print method for a locfit object.*

Description

Prints a short summary of a "locfit" object.

Usage

```
## S3 method for class 'locfit':
summary(object, ...)
```

Arguments

`object` locfit object.
`...` arguments passed to and from methods.

Value

A `summary.locfit` object, containing a short summary of the `locfit` object.

```
summary.preplot.locfit
```

Summary method for a preplot.locfit object.

Description

Prints a short summary of a "preplot.locfit" object.

Usage

```
## S3 method for class 'preplot.locfit':  
summary(object, ...)
```

Arguments

```
object      preplot.locfit object.  
...         arguments passed to and from methods.
```

Value

The fitted values from a preplot.locfit object.

```
trimod
```

Generated sample from a bivariate trimodal normal mixture

Description

This is a random sample from a mixture of three bivariate standard normal components; the sample was used for the examples in Loader (1996).

Format

Data frame with 225 observations and variables x0, x1.

Source

Randomly generated in S.

References

Loader, C. R. (1996). Local Likelihood Density Estimation. *Annals of Statistics* 24, 1602-1618.

`xbar`*Locfit Evaluation Structure*

Description

`xbar()` is an evaluation structure for `locfit.raw()`, evaluating the fit at a single point, namely, the average of each predictor variable.

Usage

```
xbar()
```

Index

*Topic **datasets**

- ais, 3
- bad, 4
- border, 5
- chemdiab, 5
- claw54, 6
- cldem, 6
- cltest, 7
- cltrain, 7
- co2, 8
- diab, 13
- ethanol, 14
- gam.slist, 17
- geyser, 18
- geyser.round, 19
- heart, 20
- insect, 21
- iris, 21
- kangaroo, 22
- livmet, 32
- mcyc, 43
- mine, 44
- mmsamp, 44
- morths, 45
- penny, 46
- spencer, 69
- stamp, 69
- trimod, 72

*Topic **htest**

- aic, 1
- aicplot, 2
- cp, 8
- cpplot, 10
- gcv, 17
- gcvplot, 18
- kdeb, 24
- lcv, 26
- lcvplot, 26
- lscv, 41

- lscv.exact, 42
- lscvplot, 43
- regband, 61
- sjpi, 65

*Topic **math**

- expit, 14

*Topic **methods**

- plot.gcvplot, 47
- plot.locfit, 48
- plot.locfit.1d, 50
- plot.locfit.2d, 50
- plot.locfit.3d, 51
- plot.scb, 52
- preplot.locfit.raw, 56
- print.gcvplot, 57
- print.locfit, 58
- print.preplot.locfit, 59
- print.scb, 59
- print.summary.locfit, 60
- summary.gcvplot, 70
- summary.locfit, 71
- summary.preplot.locfit, 72

*Topic **models**

- ang, 3
- cpar, 9
- formula.locfit, 16
- gam.lf, 16
- left, 27
- lf, 28
- locfit.matrix, 35
- lp, 40
- right, 62

*Topic **smooth**

- crit, 11
- dat, 12
- density.lf, 12
- fitted.locfit, 15
- hatmatrix, 20
- kappa0, 23

- km.mrl, 25
 - lfeval, 28
 - lfgrid, 29
 - lfknots, 30
 - lflim, 30
 - lfmarg, 31
 - lines.locfit, 31
 - locfit, 33
 - locfit.censor, 34
 - locfit.quasi, 36
 - locfit.raw, 37
 - locfit.robust, 39
 - none, 45
 - plot.eval, 46
 - plot.lfeval, 48
 - plot.preplot.locfit, 52
 - plotbyfactor, 53
 - points.locfit, 54
 - predict.locfit, 54
 - preplot.locfit, 55
 - print.lfeval, 58
 - rbox, 60
 - residuals.locfit, 62
 - rv, 63
 - rva, 64
 - scb, 64
 - smooth.lf, 66
 - spence.15, 67
 - spence.21, 68
 - store, 70
 - xbar, 73
- aic, 1, 2, 10
 - aicplot, 2, 2, 47
 - ais, 3
 - ang, 3, 40, 41
- bad, 4
 - border, 5
- chemdiab, 5
 - claw54, 6
 - cldem, 6
 - cltest, 7
 - cltrain, 7
 - co2, 8
 - cp, 8, 10
 - cpar, 9
 - cpplot, 8, 9, 10, 47, 70
- crit, 11, 11, 23, 49, 56
 - crit<-, 11, 23
 - crit<-(crit), 11
- dat, 12, 38
 - density.lf, 12, 67
 - diab, 13
- ethanol, 14
 - expit, 14
- fitted.locfit, 15, 35, 36, 62
 - formula.locfit, 16
- gam.lf, 16, 28
 - gam.slist, 17
 - gcv, 2, 10, 17, 18, 27, 43, 47, 71
 - gcvplot, 17, 18, 47, 57, 70, 71
 - geyser, 18, 19
 - geyser.round, 19
- hatmatrix, 20
 - heart, 20
- insect, 21
 - iris, 21
- kangaroo, 22
 - kappa0, 11, 23, 49, 56
 - kdeb, 24
 - km.mrl, 25, 35
- lcv, 26, 26, 27
 - lcvplot, 26, 26, 47, 66
 - left, 27, 40, 63
 - lf, 16, 17, 28
 - lfbas, 38, 63
 - lfeval, 28, 48, 58
 - lfgrid, 13, 29, 38
 - lfknots, 30
 - lflim, 30
 - lfmarg, 31, 55
 - lines, 32
 - lines.locfit, 20, 31, 49
 - livmet, 32
 - llines.locfit (lines.locfit), 31
 - locfit, 1, 2, 4, 8–11, 13, 16–18, 20, 23, 26–32, 33, 34–37, 40–43, 46–52, 54–58, 61, 63–67, 71
 - locfit.censor, 25, 34

`locfit.matrix`, 35
`locfit.quasi`, 36
`locfit.raw`, 1, 2, 8–10, 12, 13, 16–18, 20, 26–29, 33–36, 37, 40–43, 45, 47, 60, 65–67, 73
`locfit.robust`, 39
`lp`, 3, 27, 38, 39, 40, 42, 63
`lscv`, 41, 42, 43
`lscv.exact`, 41, 42, 42
`lscvplot`, 42, 43

`mcyc`, 43
`mine`, 44
`mmsamp`, 44
`morths`, 45

`none`, 38, 45

`penny`, 46
`plot`, 47, 50
`plot.eval`, 46
`plot.gcvplot`, 2, 10, 18, 27, 43, 47, 57
`plot.lfeval`, 29, 48, 58
`plot.locfit`, 11, 23, 31, 32, 48, 50–52, 54–56, 64
`plot.locfit.1d`, 20, 48, 49, 50, 52
`plot.locfit.2d`, 20, 48, 49, 50, 52
`plot.locfit.3d`, 20, 48, 49, 51, 52
`plot.preplot.locfit`, 52
`plot.scb`, 52, 65
`plotbyfactor`, 53
`points`, 54
`points.locfit`, 54
`predict.locfit`, 20, 49, 54, 55–57, 59
`preplot.locfit`, 48–52, 55, 55–57, 59
`preplot.locfit.raw`, 56
`print.gcvplot`, 57
`print.lfeval`, 29, 48, 58
`print.locfit`, 58
`print.preplot.locfit`, 59
`print.scb`, 59, 65
`print.summary.locfit`, 60

`rbox`, 38, 60
`regband`, 61
`residuals.locfit`, 15, 36, 62, 62
`right`, 27, 62
`rv`, 38, 49, 56, 63, 64
`rv<-`, 49, 56

`rv<- (rva)`, 64
`rv<-`, 63
`rva`, 64

`scb`, 52, 57, 59, 64
`sjpi`, 65
`smooth.lf`, 38, 66
`spence.15`, 67, 69
`spence.21`, 68, 68
`spencer`, 68, 69, 69
`stamp`, 69
`store`, 70
`summary.gcvplot`, 18, 57, 70
`summary.locfit`, 60, 71
`summary.preplot.locfit`, 72

`trimod`, 72

`xbar`, 73