

# Package ‘plotmo’

March 21, 2018

**Version** 3.3.6

**Title** Plot a Model's Response and Residuals

**Author** Stephen Milborrow

**Maintainer** Stephen Milborrow <milbo@sonic.net>

**Depends** R (>= 3.2.0), plotrix, TeachingDemos

**Description** Plot model surfaces for a wide variety of models using partial dependence plots and other techniques. Also plot model residuals and other information on the model.

**Suggests** C50 (>= 0.1.0-24), earth (>= 4.5.2), gbm (>= 2.1.1), glmnet (>= 2.0.5), glmnetUtils (>= 1.0.3), MASS (>= 7.3-45), neuralnet (>= 1.33), rpart (>= 4.1-10), rpart.plot (>= 2.1.0)

**License** GPL-3

**URL** <http://www.milbo.users.sonic.net>

**NeedsCompilation** no

**Repository** CRAN

**Date/Publication** 2018-03-21 13:42:02 UTC

## R topics documented:

plotmo . . . . .	2
plotmo.misc . . . . .	8
plotres . . . . .	10
plot_gbm . . . . .	14
plot_glmnet . . . . .	16
<b>Index</b>	<b>19</b>

---

plotmo

*Plot a model's response over a range of predictor values*


---

## Description

Plot model surfaces for a wide variety of models.

This function plots the model's response when varying one or two predictors while holding the other predictors constant (a poor man's partial-dependence plot). It can also generate partial-dependence plots.

Please see the [plotmo vignette](#) (also available [here](#)).

## Usage

```
plotmo(object=stop("no 'object' argument"),
       type=NULL, nresponse=NA, pmethod="plotmo",
       pt.col=0, jitter=.5, smooth.col=0, level=0,
       func=NULL, inverse.func=NULL, nrug=0, grid.col=0,
       type2="persp",
       degree1=TRUE, all1=FALSE, degree2=TRUE, all2=FALSE,
       do.par=TRUE, clip=TRUE, ylim=NULL, caption=NULL, trace=0,
       grid.func=median, grid.levels=NULL, extend=0,
       ngrid1=50, ngrid2=20, ndiscrete=5, npoints=3000,
       center=FALSE, xflip=FALSE, yflip=FALSE, swapxy=FALSE, int.only.ok=TRUE,
       ...)
```

## Arguments

object	The model object.
type	Type parameter passed to <code>predict</code> . For allowed values see the <code>predict</code> method for your object (such as <code>predict.earth</code> ). By default, <code>plotmo</code> tries to automatically select a suitable value for the model in question (usually "response") but this will not always be correct. Use <code>trace=1</code> to see the type argument passed to <code>predict</code> .
nresponse	Which column to use when <code>predict</code> returns multiple columns. This can be a column index, or a column name if the <code>predict</code> method for the model returns column names. The column name may be abbreviated, partial matching is used.
pmethod	Plotting method. One of: "plotmo" (default) Classic <code>plotmo</code> plots i.e. the background variables are fixed at their medians (or first level for factors). "partdep" Partial dependence plots, i.e. at each point the effect of the background variables is averaged. "apartdep" Approximate partial dependence plots. Faster than "partdep" especially for big datasets. Like "partdep" but the background variables are averaged over a subset of <code>ngrid1</code> cases (default 50), rather than all cases in the

training data. The subset is created by selecting rows at equally spaced intervals from the training data after sorting the data on the response values (ties are randomly broken).

The same background subset of `ngrid1` cases is used for both `degree1` and `degree2` plots.

- `pt.col` The color of response points (or response sites in `degree2` plots). This refers to the response `y` in the data used to build the model. Note that the displayed points are jittered by default (see the `jitter` argument). Default is `0`, display no response points. This can be a vector, like all such arguments, for example `pt.col = as.numeric(survived)+2`. You can modify the plotted points with `pt.pch`, `pt.cex`, etc. (these get passed via `plotmo`'s `"..."` argument). To label the points, set `pt.pch` to a character vector.
- `jitter` Applies only if `pt.col` is specified. The default is `jitter=.5`, automatically apply some jitter to the points. Points are jittered horizontally and vertically. Use `jitter=0` to disable this automatic jittering. Otherwise something like `jitter=1`, but the optimum value is data dependent.
- `smooth.col` Color of smooth line through the response points. (The points themselves will not be plotted unless `pt.col` is specified.) Default is `0`, no smooth line. Example:
- ```
mod <- lm(Volume~Height, data=trees)
plotmo(mod, pt.color=1, smooth.col=2)
```
- You can adjust the amount of smoothing with `smooth.f`. This gets passed as `f` to `lowess`. The default is `.5`. Lower values make the line more wiggly.
- `level` Draw estimated confidence or prediction interval bands at the given `level`, if the predict method for the model supports them. Default is `0`, bands not plotted. Else a fraction, for example `level=.95`. See *"Prediction intervals"* in the `plotmo` vignette. Example:
- ```
mod <- lm(log(Volume)~log(Girth), data=trees)
plotmo(mod, level=.95)
```
- You can modify the color of the bands with `level.shade` and `level.shade2`.
- `func` Superimpose `func(x)` on the plot. Example:
- ```
mod <- lm(Volume~Girth, data=trees)
estimated.volume <- function(x) .17 * x$Girth^2
plotmo(mod, pt.col=2, func=estimated.volume)
```
- The `func` is called for each plot with a single argument which is a dataframe with columns in the same order as the predictors in the formula or `x` used to build the model. Use `trace=2` to see the column names and first few rows of this dataframe.
- `inverse.func` A function applied to the response before plotting. Useful to transform a transformed response back to the original scale. Example:

```
mod <- lm(log(Volume)~., data=trees)
plotmo(mod, inverse.func=exp) # exp() is inverse of log()
```

- nrug** Number of ticks in the **rug** along the bottom of the plot  
 Default is 0, no rug.  
 Use `nrug=TRUE` for all the points.  
 Else specify the number of quantiles e.g. use `nrug=10` for ticks at the 0, 10, 20, ..., 100 percentiles.  
 Modify the rug ticks with `rug.col`, `rug.lwd`, etc.  
 The special value `nrug="density"` means plot the density of the points along the bottom. Modify the **density** plot with `density.adjust` (default is .5), `density.col`, `density.lty`, etc.
- grid.col** Default is 0, no grid. Else add a background **grid** of the specified color to the degree1 plots. The special value `grid.col=TRUE` is treated as "lightgray".
- type2** Degree2 plot type. One of "**persp**" (default), "**image**", or "**contour**". You can pass arguments to these functions if necessary by using `persp.`, `image.`, or `contour.` as a prefix. Examples:
- ```
plotmo(mod, persp.ticktype="detailed", persp.nticks=2)
plotmo(mod, type2="image")
plotmo(mod, type2="image", image.col=heat.colors(12))
plotmo(mod, type2="contour", contour.col=2, contour.labcex=.4)
```
- degree1** An index vector specifying which subset of degree1 (main effect) plots to include (after selecting the relevant predictors as described in "*Which variables are plotted?*" in the `plotmo` vignette).  
 Default is TRUE, meaning all (the TRUE gets recycled). To plot only the third plot use `degree1=3`. For no degree1 plots use `degree1=0`.  
 Can also be a character vector specifying which variables to plot e.g. `degree1=c("wind", "vis")`. Variable are matched with `grep`. Thus "wind" will match all variables with "wind" anywhere in their name. Use "`^wind$`" to match only the variable named "wind".  
 Note that an integer `degree1` indexes plots on the page, not columns of x. Probably the easiest way to use this argument (and `degree2`) is to first use the default (and possibly `all1=TRUE`) to plot all figures. This shows how the figures are numbered. Then replot using `degree1` to select the figures you want, for example `degree1=c(1, 3, 4)`.
- all1** Default is FALSE. Use TRUE to plot all predictors, not just those usually selected by `plotmo`.  
 The `all1` argument increases the number of plots; the `degree1` argument reduces the number of plots.
- degree2** An index vector specifying which subset of degree2 (interaction) plots to include.  
 Default is TRUE meaning all (after selecting the relevant interaction terms as described in "*Which variables are plotted?*" in the `plotmo` vignette).  
 Can also be a character vector specifying which variables to plot (`grep` is used for matching). For example, `degree2="vis"` selects degree2 plots for the `vis` variable.

all2	Default is FALSE. Use TRUE to plot all pairs of predictors, not just those usually selected by plotmo.
do.par	One of NULL, FALSE, TRUE, or 2, as follows: do.par=NULL. Same as do.par=FALSE if the number of plots is one; else the same as TRUE. do.par=FALSE. Use the current <code>par</code> settings. You can pass additional graphics parameters in the “...” argument. do.par=TRUE (default). Start a new page and call <code>par</code> as appropriate to display multiple plots on the same page. This automatically sets parameters like <code>mfrow</code> and <code>mar</code> . You can pass additional graphics parameters in the “...” argument. do.par=2. Like do.par=TRUE but don't restore the <code>par</code> settings to their original state when plotmo exits, so you can add something to the plot.
clip	The default is clip=TRUE, meaning ignore very outlying predictions when determining the automatic ylim. This keeps ylim fairly compact while still covering all or nearly all the data, even if there are a few crazy predicted values. See “ <i>The ylim and clip arguments</i> ” in the plotmo vignette. Use clip=FALSE for no clipping.
ylim	Three possibilities: ylim=NULL (default). Automatically determine a ylim to use across all graphs. ylim=NA. Each graph has its own ylim. ylim=c(ymin,ymax). Use the specified limits across all graphs.
caption	Overall caption. By default create the caption automatically. Use caption="" for no caption. (Use main to set the title of individual plots, can be a vector.)
trace	Default is 0. trace=1 (or TRUE) for a summary trace (shows how <code>predict</code> is invoked for the current object). trace=2 for detailed tracing. trace=-1 inhibits the messages usually issued by plotmo, like the plotmo <code>grid:</code> , calculating <code>partdep</code> , and nothing to plot messages. Error and warning messages will be printed as usual.
grid.func	Function applied to columns of the x matrix to pin the values of variables not on the axes. Default is <code>median</code> . This argument is not related to the <code>grid.col</code> argument. Examples: <pre>plotmo(mod, grid.func=mean) grid.func &lt;- function(x, ...) quantile(x)[2] # 25% quantile plotmo(mod, grid.func=grid.func)</pre> This argument is ignored for factors. The first level of factors is used. That can be changed with <code>grid.levels</code> , see below.
grid.levels	Default is NULL. Else a list of variables and their fixed value to be used when the variable is not on the axis. Supersedes <code>grid.func</code> for variables in the list. Names and values can be abbreviated, partial matching is used. Example:

```
plotmo(mod, grid.levels=list(sex="m", age=21))
```

extend	<p>Amount to extend the horizontal axis in each plot. The default is 0, do not extend (i.e. use the range of the variable in the training data). Else something like <code>extend=.5</code>, which will extend both the lower and upper <code>xlim</code> of each plot by 50%.</p> <p>This argument is useful if you want to see how the model performs on data that is beyond the training data; for example, you want to see how a time-series model performs on future data.</p> <p>This argument is currently implemented only for <code>degree1</code> plots. Factors and discrete variables (see the <code>ndiscrete</code> argument) are not extended.</p>
ngrid1	<p>Number of equally spaced x values in each <code>degree1</code> plot. Default is 50. Also used as the number of background cases for <code>pmethod="apartdep"</code>.</p>
ngrid2	<p>Grid size for <code>degree2</code> plots (<code>ngrid2</code> x <code>ngrid2</code> points are plotted). Default is 20.</p> <p>The default will sometimes be too small for contour and image plots. With large <code>ngrid2</code> values, <code>persp</code> plots look better with <code>persp.border=NA</code>.</p>
npoints	<p>Number of response points to be plotted (a sample of <code>npoints</code> points is plotted). Applies only if <code>pt.col</code> is specified.</p> <p>The default is 3000 (not all, to avoid overplotting on large models). Use <code>npoints=TRUE</code> or <code>-1</code> for all points.</p>
ndiscrete	<p>Default 5 (a somewhat arbitrary value). Variables with no more than <code>ndiscrete</code> unique values are plotted as quantized in plots (a staircase rather than a curve). Factors are always considered discrete.</p>
int.only.ok	<p>Plot the model even if it is an intercept-only model (no predictors are used in the model). Do this by plotting a single <code>degree1</code> plot for the first predictor.</p> <p>The default is <code>TRUE</code>. Use <code>int.only.ok=FALSE</code> to instead issue an error message for intercept-only models.</p>
center	<p>Center the plotted response. Default is <code>FALSE</code>.</p>
xflip	<p>Default <code>FALSE</code>. Use <code>TRUE</code> to flip the direction of the x axis. This argument (and <code>yflip</code> and <code>swapxy</code>) is useful when comparing to a plot from another source and you want the axes to be the same. (Note that <code>xflip</code> and <code>yflip</code> cannot be used on the <code>persp</code> plots, a limitation of the <code>persp</code> function.)</p>
yflip	<p>Default <code>FALSE</code>. Use <code>TRUE</code> to flip the direction of the y axis of the <code>degree2</code> graphs.</p>
swapxy	<p>Default <code>FALSE</code>. Use <code>TRUE</code> to swap the x and y axes on the <code>degree2</code> graphs.</p>
...	<p>Dot arguments are passed to the <code>predict</code> and <code>plot</code> functions. Dot argument names, whether prefixed or not, should be specified in full and not abbreviated.</p>

“Prefixed” arguments are passed directly to the associated function. For example the prefixed argument `persp.col="pink"` passes `col="pink"` to `persp()`, overriding the global `col` setting. To send an argument to `predict` whose name may alias with `plotmo`’s arguments, use `predict.` as a prefix. Example:

```
plotmo(mod, s=1)           # error: arg matches multiple formal args
plotmo(mod, predict.s=1)  # ok now: s=1 will be passed to predict()
```

The prefixes recognized by plotmo are:

predict.	passed to the <a href="#">predict</a> method for the model
degree1.	modifies degree1 plots e.g. degree1.col=3, degree1.lwd=2
persp.	arguments passed to <a href="#">persp</a>
contour.	arguments passed to <a href="#">contour</a>
image.	arguments passed to <a href="#">image</a>
pt.	see the pt.col argument (arguments passed to <a href="#">points</a> and <a href="#">text</a> )
smooth.	see the smooth.col argument (arguments passed to <a href="#">lines</a> and <a href="#">lowess</a> )
level.	see the level argument (level.shade, level.shade2, and arguments for <a href="#">polygon</a> )
func.	see the func argument (arguments passed to <a href="#">lines</a> )
rug.	see the nrug argument (rug.jitter, and arguments passed to <a href="#">rug</a> )
density.	see the nrug argument (density.adjust, and arguments passed to <a href="#">lines</a> )
grid.	see the grid.col argument (arguments passed to <a href="#">grid</a> )
caption.	see the caption argument (arguments passed to <a href="#">mtext</a> )
par.	arguments passed to <a href="#">par</a> (only necessary if a par argument name clashes with a plotmo argument)

The cex argument is relative, so specifying cex=1 is the same as not specifying cex.

For backwards compatibility, some dot arguments are supported but not explicitly documented. For example, the old argument col.response is no longer in plotmo's formal argument list, but is still accepted and treated like the new argument pt.col.

## Note

In general this function won't work on models that don't save the call and data with the model in a standard way. For further discussion please see "[Accessing the model data](#)" in the [plotmo vignette](#). Package authors may want to look at [Guidelines for S3 Regression Models](#) (also available [here](#)).

By default, plotmo tries to use sensible model-dependent defaults when calling predict. Use trace=1 to see the arguments passed to predict. You can change the defaults by using plotmo's type argument, and by using dot arguments prefixed with predict. (see the description of "..."  
above).

## See Also

Please see the [plotmo vignette](#) (also available [here](#)).

## Examples

```
if (require(rpart)) {
  data(kyphosis)
  rpart.model <- rpart(Kyphosis~., data=kyphosis)
  plotmo(rpart.model, type="prob", nresponse="present")
}
if (require(earth)) {
```

```

data(ozone1)
earth.model <- earth(O3 ~ ., data=ozone1, degree=2)
plotmo(earth.model)
}

```

---

plotmo.misc

*Ignore*


---

## Description

Miscellaneous functions exported for internal use by earth and other packages. You can ignore these.

## Usage

```

# for earth
plotmo_fitted(object, trace, nresponse, type, ...)
plotmo_cum(rinfo, info, nfigs=1, add=FALSE,
           cum.col1, grid.col, jitter=0, cum.grid="percentages", ...)
plotmo_nresponse(y, object, nresponse, trace, fname, type="response")
plotmo_rinfo(object, type=NULL, residtype=type, nresponse=1,
             standardize=FALSE, delever=FALSE, trace=0,
             leverage.msg="returned as NA", expected.levs=NULL, labels.id=NULL, ...)
plotmo_predict(object, newdata, nresponse,
              type, expected.levs, trace, inverse.func=NULL, ...)
plotmo_prolog(object, object.name, trace, ...)
plotmo_resplevs(object, plotmo_fitted, yfull, trace)
plotmo_rsq(object, newdata, trace=0, nresponse=NA, type=NULL, ...)
plotmo_standardizescale(object)
plotmo_type(object, trace, fname="plotmo", type, ...)
plotmo_y(object, nresponse=NULL, trace=0, expected.len=NULL,
         resp.levs=NULL, convert.glm.response=!is.null(nresponse))
## Default S3 method:
plotmo.pairs(object, x, nresponse, trace, all2, ...)
## Default S3 method:
plotmo.singles(object, x, nresponse, trace, all1, ...)
## Default S3 method:
plotmo.y(object, trace, naked, expected.len, ...)
# plotmo methods
plotmo.convert.na.nresponse(object, nresponse, yhat, type="response", ...)
plotmo.pairs(object, x, nresponse, trace, all2, ...)
plotmo.pint(object, newdata, type, level, trace, ...)
plotmo.predict(object, newdata, type, ..., TRACE)
plotmo.prolog(object, object.name, trace, ...)
plotmo.residtype(object, ...)
plotmo.singles(object, x, nresponse, trace, all1, ...)
plotmo.type(object, ...)

```



```
plotmo.x(object, trace, ...)  
plotmo.y(object, trace, naked, expected.len, nresponse=1, ...)
```

### Arguments

...	-
add	-
all1	-
all2	-
convert.glm.response	-
cum.col1	-
cum.grid	-
delever	-
expected.len	-
expected.levs	-
fname	-
grid.col	-
info	-
inverse.func	-
jitter	-
labels.id	-
level	-
leverage.msg	-
naked	-
newdata	-
nfigs	-
nresponse	-
object.name	-
object	-
plotmo_fitted	-
residtype	-
resp.levs	-
rinfo	-
standardize	-
TRACE	-
trace	-
type	-
x	-
yfull	-
yhat	-
y	-

plotres

*Plot the residuals of a regression model***Description**

Plot the residuals of a regression model.

Please see the [plotres vignette](#) (also available [here](#)).

**Usage**

```
plotres(object = stop("no 'object' argument"),
        which = 1:4, info = FALSE, versus = 1,
        standardize = FALSE, delever = FALSE, level = 0,
        id.n = 3, labels.id = NULL, smooth.col = 2,
        grid.col = 0, jitter = 0,
        do.par = NULL, caption = NULL, trace = 0,
        npoints = 3000, center = TRUE,
        type = NULL, nresponse = NA,
        object.name = quote.deparse(substitute(object)), ...)
```

**Arguments**

object	The model object.
which	Which plots do draw. Default is 1:4. 1 Model plot. What gets plotted here depends on the model class. For example, for earth models this is a model selection plot. Nothing will be displayed for some models. For details, please see the <a href="#">plotres vignette</a> . 2 Cumulative distribution of abs residuals 3 Residuals vs fitted 4 QQ plot 5 Abs residuals vs fitted 6 Sqrt abs residuals vs fitted 7 Abs residuals vs log fitted 8 Cube root of the squared residuals vs log fitted 9 Log abs residuals vs log fitted
info	Default is FALSE. Use TRUE to print extra information as follows: i) Display the distribution of the residuals along the bottom of the plot. ii) Display the training R-Squared. iii) Display the Spearman Rank Correlation of the absolute residuals with the fitted values. Actually, correlation is measured against the absolute values of whatever is on the horizontal axis — by default this is the fitted response, but may be something else if the versus argument is used.

- iv) In the Cumulative Distribution plot (`which=2`), display additional information on the quantiles.
- v) Only for `which=5` or `9`. Regress the absolute residuals against the fitted values and display the regression slope. Robust linear regression is used via `rlm` in the MASS package.
- vi) Add various annotations to the other plots.

versus	<p>What do we plot the residuals against? One of:</p> <ol style="list-style-type: none"> <li>1 Default. Plot the residuals versus the fitted values (or the log values when <code>which=7</code> to <code>9</code>).</li> <li>2 Residuals versus observation number, after observations have been sorted on the fitted value. Same as <code>versus=1</code>, except that the residuals are spaced uniformly along the horizontal axis.</li> <li>3 Residuals versus the response.</li> <li>4 Residuals versus the hat leverages.</li> </ol> <p>"b:" Residuals versus the basis functions. Currently only supported for <code>earth</code>, <code>mda::mars</code>, and <code>gam::gam</code> models. A optional <code>regex</code> can follow the "b:" to specify a subset of the terms, e.g. <code>versus="b:wind"</code> will plot terms with "wind" in their name.</p> <p>Else a character vector specifying which predictors to plot against.</p> <p>Example 1: <code>versus=""</code> plots against all predictors (since the <code>regex</code> <code>versus=""</code> matches anything).</p> <p>Example 2: <code>versus=c("wind", "vis")</code> plots predictors with <code>wind</code> or <code>vis</code> in their name.</p> <p>Example 3: <code>versus=c("wind vis")</code> equivalent to the above.</p> <p>Note: These are <code>regex</code>s. Thus <code>versus="wind"</code> will match all variables that have "wind" in their names. Use <code>"^wind\$"</code> to match only the variable named "wind".</p>
standardize	<p>Default is <code>FALSE</code>. Use <code>TRUE</code> to standardize the residuals. Only supported for some models, an error message will be issued otherwise.</p> <p>Each residual is divided by <math>se_i * \sqrt{1 - h_{ii}}</math>, where <code>se_i</code> is the standard error of prediction and <code>h_ii</code> is the leverage (the diagonal entry of the hat matrix). When the variance model holds, the standardized residuals are homoscedastic with unity variance.</p> <p>The leverages are obtained using <code>hatvalues</code>. (For <code>earth</code> models the leverages are for the linear regression of the response on the basis matrix <code>bx</code>.) A standardized residual with a leverage of 1 is plotted as a star on the axis.</p> <p>This argument applies to all plots where the residuals are used (including the cumulative distribution and QQ plots, and to annotations displayed by the <code>info</code> argument).</p>
delever	<p>Default is <code>FALSE</code>. Use <code>TRUE</code> to "de-lever" the residuals. Only supported for some models, an error message will be issued otherwise.</p> <p>Each residual is divided by <math>\sqrt{1 - h_{ii}}</math>. See the <code>standardize</code> argument for details.</p>

level	<p>Draw estimated confidence or prediction interval bands at the given level, if the model supports them.</p> <p>Default is 0, bands not plotted. Else a fraction, for example level=0.90. Example:</p> <pre>mod &lt;- lm(log(Volume)~log(Girth), data=trees) plotres(mod, level=.90)</pre> <p>You can modify the color of the bands with level.shade and level.shade2. See also “<i>Prediction intervals</i>” in the <a href="#">plotmo vignette</a> (but note that plotmo needs prediction intervals on <i>new</i> data, whereas plotres requires only that the model supports prediction intervals on the training data).</p>
id.n	<p>The largest id.n residuals will be labeled in the plot. Default is 3. Special values TRUE and -1 or mean all.</p> <p>If id.n is negative (but not -1) the id.n most positive and most negative residuals will be labeled in the plot.</p> <p>A current implementation restriction is that id.n is ignored when there are more than ten thousand cases.</p>
labels.id	<p>Residual labels. Only used if id.n &gt; 0. Default is the case names, or the case numbers if the cases are unnamed.</p>
smooth.col	<p>Color of the smooth line through the residual points. Default is 2, red. Use smooth.col=0 for no smooth line.</p> <p>You can adjust the amount of smoothing with smooth.f. This gets passed as f to <a href="#">lowess</a>. The default is 2/3. Lower values make the line more wiggly.</p>
grid.col	<p>Default is 0, no grid. Else add a background <a href="#">grid</a> of the specified color to the degree1 plots. The special value grid.col=TRUE is treated as “lightgray”.</p>
jitter	<p>Default is 0, no jitter. Passed as factor to <a href="#">jitter</a> to jitter the plotted points horizontally and vertically. Useful for discrete variables and responses, where the residual points tend to be overlaid.</p>
do.par	<p>One of NULL, FALSE, TRUE, or 2, as follows:</p> <p>do.par=NULL (default). Same as do.par=FALSE if the number of plots is one; else the same as TRUE.</p> <p>do.par=FALSE. Use the current <a href="#">par</a> settings. You can pass additional graphics parameters in the “...” argument.</p> <p>do.par=TRUE. Start a new page and call <a href="#">par</a> as appropriate to display multiple plots on the same page. This automatically sets parameters like mfrow and mar. You can pass additional graphics parameters in the “...” argument.</p> <p>do.par=2. Like do.par=TRUE but don’t restore the <a href="#">par</a> settings to their original state when plotres exits, so you can add something to the plot.</p>
caption	<p>Overall caption. By default create the caption automatically. Use caption="" for no caption. (Use main to set the title of an individual plot.)</p>
trace	<p>Default is 0.</p> <p>trace=1 (or TRUE) for a summary trace (shows how <a href="#">predict</a> and friends are invoked for the model).</p> <p>trace=2 for detailed tracing.</p>

<code>npoints</code>	Number of points to be plotted. A sample of <code>npoints</code> is taken; the sample includes the biggest twenty or so residuals. The default is 3000 (not all, to avoid overplotting on large models). Use <code>npoints=TRUE</code> or <code>-1</code> for all points.
<code>center</code>	Default is <code>TRUE</code> , meaning center the horizontal axis in the residuals plot, so asymmetry in the residual distribution is more obvious.
<code>type</code>	Type parameter passed first to <code>residuals</code> and if that fails to <code>predict</code> . For allowed values see the <code>residuals</code> and <code>predict</code> methods for your object (such as <code>residuals.rpart</code> or <code>predict.earth</code> ). By default, <code>plotres</code> tries to automatically select a suitable value for the model in question (usually "response"), but this will not always be correct. Use <code>trace=1</code> to see the <code>type</code> argument passed to <code>residuals</code> and <code>predict</code> .
<code>nresponse</code>	Which column to use when <code>residuals</code> or <code>predict</code> returns multiple columns. This can be a column index or column name (which may be abbreviated, partial matching is used).
<code>object.name</code>	The name of the object for error and trace messages. Used internally by <code>plot.earth</code> .
<code>...</code>	Dot arguments are passed to the plot functions. Dot argument names, whether prefixed or not, should be specified in full and not abbreviated. "Prefixed" arguments are passed directly to the associated function. For example the prefixed argument <code>pt.col="pink"</code> passes <code>col="pink"</code> to <code>points()</code> , overriding the global <code>col</code> setting. The prefixes recognized by <code>plotres</code> are:
<code>residuals.</code>	passed to <code>residuals</code>
<code>predict.</code>	passed to <code>predict</code> ( <code>predict</code> is called if the call to <code>residuals</code> fails)
<code>w1.</code>	sent to the model-dependent plot for <code>which=1</code> e.g. <code>w1.col=2</code>
<code>pt.</code>	modify the displayed points e.g. <code>pt.col=as.numeric(survived)+2</code> or <code>pt.cex=.8</code> .
<code>smooth.</code>	modify the smooth line e.g. <code>smooth.col=0</code> or <code>smooth.f=.5</code> .
<code>level.</code>	modify the interval bands, e.g. <code>level.shade="gray"</code> or <code>level.shade2="lightblue"</code>
<code>legend.</code>	modify the displayed <code>legend</code> e.g. <code>legend.cex=.9</code>
<code>cum.</code>	modify the Cumulative Distribution plot (arguments for <code>plot.stepfun</code> )
<code>qq.</code>	modify the QQ plot, e.g. <code>qq.pch=1</code>
<code>qqline</code>	modify the <code>qqline</code> in the QQ plot, e.g. <code>qqline.col=0</code>
<code>label.</code>	modify the point labels, e.g. <code>label.cex=.9</code> or <code>label.font=2</code>
<code>cook.</code>	modify the Cook's Distance annotations. This affects only the leverage plot ( <code>versus=3</code> ) for <code>lm</code> models with <code>sta</code>
<code>caption.</code>	modify the overall caption (see the <code>caption</code> argument) e.g. <code>caption.col=2</code> .
<code>par.</code>	arguments for <code>par</code> (only necessary if a <code>par</code> argument name clashes with a <code>plotres</code> argument)

The `cex` argument is relative, so specifying `cex=1` is the same as not specifying `cex`.

For backwards compatibility, some dot arguments are supported but not explicitly documented.

**Value**

If the which=1 plot was plotted, the return value of that plot (model dependent).

Else if the which=3 plot was plotted, return `list(x,y)` where `x` and `y` are the coordinates of the points in that plot (but without jittering even if the `jitter` argument was used).

Else return `NULL`.

**Note**

This function is designed primarily for displaying standard response - fitted residuals for models with a single continuous response, although it will work for a few other models.

In general this function won't work on models that don't save the call and data with the model in a standard way. It uses the same underlying mechanism to access the model data as `plotmo`. For further discussion please see "Accessing the model data" in the [plotmo vignette](#) (also available [here](#)). Package authors may want to look at [Guidelines for S3 Regression Models](#) (also available [here](#)).

**See Also**

Please see the [plotres vignette](#) (also available [here](#)).

[plot.lm](#)

[plot.earth](#)

**Examples**

```
# we use lm in this example, but plotres is more useful for models
# that don't have a function like plot.lm for plotting residuals

lm.model <- lm(Volume~., data=trees)

plotres(lm.model)
```

---

plot\_gbm

*Plot a gbm model*

---

**Description**

Plot a `gbm` model showing the training and other error curves.

**Usage**

```
plot_gbm(object=stop("no 'object' argument"),
  smooth = c(0, 0, 0, 1),
  col = c(1, 2, 3, 4), ylim = "auto",
  legend.x = NULL, legend.y = NULL, legend.cex = .8,
  grid.col = NA,
  n.trees = NA, col.n.trees = "darkgray",
  ...)
```

**Arguments**

object	The gbm model.
smooth	Four-element vector specifying if smoothing should be applied to the train, test, CV, and OOB curves respectively. When smoothing is specified, a smoothed curve is plotted and the minimum is calculated from the smoothed curve. The default is <code>c(0, 0, 0, 1)</code> meaning apply smoothing only to the OOB curve (same as <code>gbm.perf</code> ). Note that <code>smooth=1</code> (which gets recycled to <code>c(1, 1, 1, 1)</code> ) will smooth all the curves.
col	Four-element vector specifying the colors for the train, test, CV, and OOB curves respectively. The default is <code>c(1, 2, 3, 4)</code> . Use a color of <code>0</code> to remove the corresponding curve, e.g. <code>col=c(1, 2, 3, 0)</code> to not display the OOB curve. If <code>col=0</code> (which gets recycled to <code>c(0, 0, 0, 0)</code> ) nothing will be plotted, but <code>plot_gbm</code> will return the number-of-trees at the minima as usual (as described in the Value section below).
ylim	The default <code>ylim="auto"</code> shows more detail around the minima. Use <code>ylim=NULL</code> for the full vertical range of the curves. Else specify <code>ylim</code> as usual.
legend.x	The x position of the legend. The default positions the legend automatically. Use <code>legend.x=NA</code> for no legend. See the x and y arguments of <code>xy.coords</code> for other options, for example <code>legend.x="topright"</code> .
legend.y	The y position of the legend.
legend.cex	The legend cex (the default is <code>0.8</code> ).
grid.col	Default NA. Color of the optional grid, for example <code>grid.col=1</code> .
n.trees	For use by <code>plotres</code> . The x position of the gray vertical line indicating the <code>n.trees</code> passed by <code>plotres</code> to <code>predict.gbm</code> to calculate the residuals. <code>Plotres</code> defaults to all trees.
col.n.trees	For use by <code>plotres</code> . Color of the vertical line showing the <code>n.trees</code> argument. Default is <code>"darkgray"</code> .
...	Dot arguments are passed internally to <code>plot.default</code> .

**Value**

This function returns a four-element vector specifying the number of trees at the train, test, CV, and OOB minima respectively.

The minima are calculated after smoothing as specified by this function's `smooth` argument. By default, only the OOB curve is smoothed. The smoothing algorithm for the OOB curve differs slightly from `gbm.perf`, so can give a slightly different number of trees.

**Note****The OOB curve**

The OOB curve is artificially rescaled to force it into the plot. See Chapter 7 in the [plotres vignette](#).

**Interaction with plotres**

When invoking this function via `plotres`, prefix any argument of `plotres` with `w1.` to tell `plotres` to pass the argument to this function. For example give `w1.ylim=c(0,10)` to `plotres` (plain `ylim=c(0,10)` in this context gets passed to the residual plots).

**Acknowledgments**

This function is derived from code in the `gbm` package authored by Greg Ridgeway and others.

**See Also**

Chapter 7 in [plotres vignette](#) discusses this function.

**Examples**

```
if (require(gbm)) {
  n <- 100                                # toy model for quick demo
  x1 <- 3 * runif(n)
  x2 <- 3 * runif(n)
  x3 <- sample(1:4, n, replace=TRUE)
  y <- x1 + x2 + x3 + rnorm(n, 0, .3)
  data <- data.frame(y=y, x1=x1, x2=x2, x3=x3)
  mod <- gbm(y~., data=data, distribution="gaussian",
            n.trees=300, shrinkage=.1, interaction.depth=3,
            train.fraction=.8, verbose=FALSE)

  plot_gbm(mod)

  # plotres(mod)                            # plot residuals

  # plotmo(mod)                             # plot regression surfaces
}
```

---

plot\_glmnet

*Plot a glmnet model*

---

**Description**

Plot the coefficient paths of a `glmnet` model.

An enhanced version of `plot.glmnet`.

**Usage**

```
plot_glmnet(x = stop("no 'x' argument"),
            xvar = c("rlambda", "lambda", "norm", "dev"),
            label = 10, nresponse = NA, grid.col = NA, s = NA, ...)
```



**Arguments**

x	The glmnet model.
xvar	What gets plotted along the x axis. One of: "rlambda" (default) decreasing log lambda (lambda is the glmnet penalty) "lambda" log lambda "norm" L1-norm of the coefficients "dev" percent deviance explained
label	The default xvar differs from plot.glmnet to allow s to be plotted when this function is invoked by <code>plotres</code> . Default 10. Number of variable names displayed on the right of the plot. One of: FALSE display no variables TRUE display all variables integer (default) number of variables to display (default is 10)
nresponse	Which response to plot for multiple response models.
grid.col	Default NA. Color of the optional grid, for example <code>grid.col="lightgray"</code> .
s	For use by <code>plotres</code> . The x position of the gray vertical line indicating the lambda s passed by <code>plotres</code> to <code>predict.glmnet</code> to calculate the residuals. <code>Plotres</code> defaults to <code>s=0</code> .
...	Dot arguments are passed internally to <code>matplot</code> . Use <code>col</code> to change the color of curves; for example <code>col=1:4</code> . The six default colors are intended to be distinguishable yet harmonious (to my eye at least), with adjacent colors as different as easily possible.

**Note****Limitations**

For multiple response models use the `nresponse` argument to specify which response should be plotted. (Currently each response must be plotted one by one.)

The `type.coef` argument of `plot.glmnet` is currently not supported.

Currently `xvar="norm"` is not supported for multiple response models (you will get an error message).

**Interaction with plotres**

When invoking this function via `plotres`, prefix any argument of `plotres` with `w1.` to tell `plotres` to pass the argument to this function. For example give `w1.col=1:4` to `plotres` (plain `col=1:4` in this context gets passed to the residual plots).

**Acknowledgments**

This function is based on `plot.glmnet` in the `glmnet` package authored by Jerome Friedman, Trevor Hastie, and Rob Tibshirani.

**See Also**

Chapter 6 in `plotres vignette` discusses this function.

**Examples**

```
if (require(glmnet)) {  
  x <- matrix(rnorm(100 * 10), 100, 10) # n=100 p=10  
  y <- x[,1] + x[,2] + 2 * rnorm(100) # y depends only on x[,1] and x[,2]  
  mod <- glmnet(x, y)  
  
  plot_glmnet(mod)  
  
  # plotres(mod) # plot the residuals  
}
```

# Index

## \*Topic **partial dependence**

plotmo, 2  
plotres, 10

## \*Topic **regression**

plotmo, 2  
plotres, 10

check.index (plotmo.misc), 8  
contour, 4, 7

density, 4

gbm, 14, 16  
gbm.perf, 15  
glmnet, 16, 17  
grep, 4  
grid, 4, 7, 12

hatvalues, 11

image, 4, 7

jitter, 12

legend, 13  
lines, 7  
lowess, 3, 7, 12

matplot, 17  
median, 5  
mtext, 7

par, 5, 7, 12, 13  
persp, 4, 7  
plot.default, 15  
plot.earth, 14  
plot.glmnet, 16, 17  
plot.lm, 14  
plot.stepfun, 13  
plot\_gbm, 14  
plot\_glmnet, 16

plotmo, 2, 14  
plotmo.convert.na.nresponse  
    (plotmo.misc), 8  
plotmo.misc, 8  
plotmo.pairs (plotmo.misc), 8  
plotmo.pint (plotmo.misc), 8  
plotmo.predict (plotmo.misc), 8  
plotmo.prolog (plotmo.misc), 8  
plotmo.residtype (plotmo.misc), 8  
plotmo.singles (plotmo.misc), 8  
plotmo.type (plotmo.misc), 8  
plotmo.x (plotmo.misc), 8  
plotmo.y (plotmo.misc), 8  
plotmo\_cum (plotmo.misc), 8  
plotmo\_fitted (plotmo.misc), 8  
plotmo\_nresponse (plotmo.misc), 8  
plotmo\_predict (plotmo.misc), 8  
plotmo\_prolog (plotmo.misc), 8  
plotmo\_resplevs (plotmo.misc), 8  
plotmo\_response (plotmo.misc), 8  
plotmo\_rinfo (plotmo.misc), 8  
plotmo\_rsq (plotmo.misc), 8  
plotmo\_standardizescale (plotmo.misc), 8  
plotmo\_type (plotmo.misc), 8  
plotmo\_y (plotmo.misc), 8  
plotres, 10, 15–17  
points, 7  
polygon, 7  
predict, 2, 5, 7, 12, 13  
predict.earth, 2, 13

qqline, 13

regex, 11  
residuals, 13  
residuals.rpart, 13  
rml, 11  
rug, 4, 7  
text, 7

`xy.coords`, [15](#)