

Package ‘interflex’

October 13, 2022

Type Package

Version 1.2.6

Title Multiplicative Interaction Models Diagnostics and Visualization

Date 2021-05-02

Description Performs diagnostic tests of multiplicative interaction models and plots non-linear marginal effects of a treatment on an outcome across different values of a moderator.

URL <http://yiqingxu.org/software/interaction/RGuide.html>

License MIT + file LICENSE

Imports ggplot2 (>= 2.1.0), sandwich (>= 2.3-4), Lmoments (>= 1.2-3), doParallel (>= 1.0.10), foreach (>= 1.4.3), mgcv (>= 1.8-16), lfe (>= 2.6-2291), gridExtra, grid, ggplotify, RColorBrewer, pcse, gtable, MASS, mvtnorm, pROC, ModelMetrics, Rcpp, lmtest, AER, future

LinkingTo Rcpp, RcppArmadillo

LazyData true

NeedsCompilation yes

Maintainer Ziyi Liu <zyliu2020@uchicago.edu>

Depends R (>= 3.5.0)

Author Jens Hainmueller [aut, ctb],
Jonathan Mummolo [aut, ctb],
Yiqing Xu [aut, ctb],
Ziyi Liu [aut, cre]

Repository CRAN

Date/Publication 2021-05-18 11:40:02 UTC

R topics documented:

interflex-package	2
inter.test	3
interflex	4

interflex-internal	10
plot.interflex	10
predict.interflex	12
RcppArmadillo-Functions	14
s1	15
s2	16
s3	16
s4	16
s5	17
s6	17
s7	17
s8	18
s9	18
Index	19

interflex-package	<i>Multiplicative Interaction Models Diagnostics and Visualization</i>
-------------------	--

Description

Producing Flexible Marginal Effect Estimates with Multiplicative Interaction Models

Details

This package performs diagnostics and visualizations of multiplicative interaction models. Besides conventional linear interaction models, it provides two additional estimation strategies—linear regression based on pre-specified bins and locally linear regressions based on Gaussian kernels—to flexibly estimate the conditional marginal effect of a treatment variable on an outcome variable across different values of a moderating variable. These approaches relax the linear interaction effect assumption and safeguard against excessive extrapolation.

Author(s)

Jens Hainmueller; Jonathan Mummolo; Yiqing Xu (Maintainer); Ziyi Liu

References

Jens Hainmueller; Jonathan Mummolo; Yiqing Xu. 2019. "How Much Should We Trust Estimates from Multiplicative Interaction Models? Simple Tools to Improve Empirical Practice." *Political Analysis*, Vol. 27, Iss. 2, April 2019, pp. 163–192. Available at: <https://www.cambridge.org/core/journals/political-analysis/article/how-much-should-we-trust-estimates-from-multiplicative-interaction-models-simple-tools-to-improve-empirical-practice/0D8CAACB473F9B1EE256F43B38E458706>.

See Also

[interflex](#), [plot.interflex](#), and [predict.interflex](#)

inter.test	<i>Conduct t-Test after Estimation</i>
------------	--

Description

Compare treatment effects at specific values of the moderator

Usage

```
inter.test(out, diff.values, percentile=FALSE, k=16)
```

Arguments

out	an interflex object.
diff.values	a numeric vector contain 2 or 3 elements which are within the range of the moderator. The treatment effects at corresponding values of the moderator will be compared.
percentile	a logical flag indicating whether to take values of the moderator on a percentile scale.
k	an integer specifying the dimension(s) of the bases used to represent the smooth term, default to 16.

Details

inter.test compare treatment effects at specific values of the moderator using marginal effects and vcov matrix derived from linear/kernel estimation. Based on GAM model(relies on **mgcv** package), users can approximate the treatment effect and its variance using smooth functions without re-estimating the model, hence saving time.

Author(s)

Jens Hainmueller; Jonathan Mummolo; Yiqing Xu (Maintainer); Ziyi Liu

References

Jens Hainmueller; Jonathan Mummolo; Yiqing Xu. 2019. "How Much Should We Trust Estimates from Multiplicative Interaction Models? Simple Tools to Improve Empirical Practice." *Political Analysis*, Vol. 27, Iss. 2, April 2019, pp. 163–192. Available at: <https://www.cambridge.org/core/journals/political-analysis/article/how-much-should-we-trust-estimates-from-multiplicative-in-D8CAACB473F9B1EE256F43B38E458706>.

See Also

[interflex](#), [plot.interflex](#) and [predict.interflex](#)

interflex

*Flexible Interactive Models***Description**

Producing flexible marginal effect estimates with linear, binning, and kernel estimators

Usage

```
interflex(estimator, data, Y, D, X, treat.type = NULL,
          base = NULL, Z = NULL, IV = NULL, FE = NULL,
          full.moderate = FALSE,
          weights = NULL, na.rm = FALSE, Xunif = FALSE,
          CI = TRUE, neval = 50, X.eval = NULL,
          method = "linear", vartype = "delta",
          vcov.type = "robust", time=NULL,
          pairwise = TRUE, nboots = 200, nsimu = 1000,
          parallel = TRUE, cores = 4, cl = NULL,
          Z.ref = NULL, D.ref = NULL,
          nbins = 3, cutoffs = NULL, wald = TRUE,
          bw = NULL, kfold = 10, grid = 30,
          metric = NULL, figure = TRUE, bin.labs = TRUE,
          order = NULL, subtitles = NULL, show.subtitles = NULL,
          Xdistr = "histogram", main = NULL,
          Ylabel = NULL, Dlabel = NULL, Xlabel = NULL,
          xlab = NULL, ylab = NULL, xlim = NULL, ylim = NULL,
          theme.bw = FALSE, show.grid = TRUE, cex.main = NULL,
          cex.sub = NULL, cex.lab = NULL, cex.axis = NULL,
          interval = NULL, file = NULL, ncols = NULL,
          pool = FALSE, color = NULL, show.all = FALSE,
          legend.title = NULL, diff.values = NULL,
          percentile = FALSE, gam.k = 10,
          angle=c(30, 100, -30, -120),
          span=NULL, scale = 1.1, height = 7, width = 10)
```

Arguments

estimator	a string that specifies the estimator. Can be one of the following: "linear", "binning", "kernel" or "gam".
data	a data.frame that stores Y, D, and X.
Y	a string, name of the outcome variable.
D	a string, name of the treatment variable.
X	a string, name of the moderating variable.
treat.type	a string that specifies the type of the treatment variable. Can be one of the following: "discrete" and "continuous". If not specified, it will be automatically set based on the treatment variable. (If D is in numeric format and the number

of multiple arms of D is larger than 5, `treat.type` will be set to "continuous". Otherwise, it will be set to "discrete".)

<code>base</code>	a string specifying the base group if <code>treat.type</code> is "discrete". If not specified, the base group will be selected based on numeric/character order of values in the treatment variable. For example, when treatment D can take values 0, 1, and 2, those observations with treatment 0 will be selected as the base group.
<code>Z</code>	a vector of variable names of control variables.
<code>IV</code>	a vector of variable names for the instrument variables of the treatment, doesn't support glm model in this version
<code>FE</code>	a vector of variable names of fixed effects indicators, doesn't support glm model in this version
<code>full.moderate</code>	a logical flag indicating whether to use fully moderated model that is default to FALSE. If it is set to TRUE, the program will include additional interactions between the moderating variable and other covariates.
<code>weights</code>	a string, name of the weighting variable.
<code>na.rm</code>	a logical flag indicating whether to list-wise delete missing data. The algorithm will report an error if missing data exist.
<code>Xunif</code>	a logical flag controlling whether to transform values of the moderator into their rank order.
<code>CI</code>	a logical flag indicating whether the confidence intervals need to be shown. If <code>CI</code> = FALSE, the bootstrap procedure will not be conducted.
<code>neval</code>	an integer specifying the number of evaluation points. A large number would produce smoother marginal effects estimates but the algorithm will take more time to run.
<code>X.eval</code>	a vector of manually selected evaluation points.
<code>method</code>	a string that specifies the link function. Can be one of the following: "linear", "probit", "logit", "poisson", "nbinom". The default is "linear".
<code>vartype</code>	a string that controls the method used for uncertainty estimation when <code>estimator</code> = "linear" or "binning". Can be one of the following: "delta", "simu" and "bootstrap". The default is "delta".
<code>vcov.type</code>	a string that controls the variance-covariance estimator. Can be one of the following: "homoscedastic"; "robust"; "cluster" and "pcse"(only for linear models). The default is "robust".
<code>time</code>	a string specifying the name of time variable for panel corrected standard errors. Ignored when <code>vcov.type</code> is not "pcse".
<code>pairwise</code>	a logical flag indicating whether to switch on the <code>pairwise</code> option when estimating panel corrected standard errors. Ignored when <code>vcov.type</code> is not "pcse". The default is TRUE.
<code>nboots</code>	an integer specifying the number of bootstrap runs. The default is 200.
<code>nsimu</code>	an integer specifying the rounds of simulations. The default is 1000.
<code>parallel</code>	a logical flag indicating whether parallel computing will be used in bootstrapping and/or cross-validation. Ignored if <code>CI</code> = FALSE and <code>bw</code> is specified.

cores	an integer indicating the number of cores to be used in parallel computing. If not specified, the algorithm will use the maximum number of logical cores of your computer (warning: this could prevent you from multi-tasking on your computer).
cl	a string specifying the name of clustering variable. If specified, the bootstrap and cross-validation procedure will be blocking on this variable. Ignored when CI = FALSE . When vartype is "cluster", it specifies the name of unit variable for estimating clustered standard errors.
Z.ref	a numeric vector specifies the value of covariates when estimating marginal effects in glm models, must have the same length and order as Z . In the default setting, each element of Z.ref equals to the mean of the corresponding covariate in the data.
D.ref	a numeric vector specifies the value of D when estimating marginal effects when D is continuous. The default is the mean of D in the data.
nbins	an integer that determines the number of bins of the moderator. The sample will be divided based on integral multiples of 100/nbins percentile. nbins will be automatically subtracted by 1 if more than one multiples of the percentile have the same value (for example, if the moderator has over 70% zeros, both the 33 and 66 percentiles are zero). The default is 3. Ignored when cutoffs is supplied. When nbins are set to 1, it only conducts a linear fit.
cutoffs	a vector of numbers that determines how subgroups are divided based on the moderator X. When this option is supplied, nbins will be ignored. The smallest number of the first interval and the largest number of the last interval do not need to be specified. Numbers smaller than the minimum or larger than the maximum of X will be ignored. nbins equals to the length of this vector plus 1. Ignored if the treatment is dichotomous.
wald	a logical flag controlling whether to conduct a Wald test.
bw	a number that sets the bandwidth in Gaussian kernel reweighting.
kfold	an integer specified the number of folds in cross-validation, default to 10.
grid	either an integer indicating the number of candidates in a grid search to find the optimal bandwidth, or a vector of positive numbers of candidate bandwidths, default to 30. Ignored if bw is specified.
metric	a string specifies the criterion to be used in cross-validation to select bandwidth, can be one of the following: "MSE", "MAE", "Cross Entropy" and "AUC". The default will be "MSPE" if the outcome is not binary, otherwise the default will be "Cross Entropy"
figure	a logical flag controlling whether to draw a marginal effect plot.
bin.labs	a logical flag indicating whether to show a label for each bin in a binning plot; only meaningful when nbins is 3 or 4. The default is TRUE.
order	a vector of strings that determines the order of treatment arms in the plot when visualizing marginal effects. It should contain all kinds of treatment arms except for the base group.
subtitles	a vector of strings that determines the subtitles of subplots when pool is FALSE, or determines the label in the legend when pool is TRUE. It should have the same length as the number of categories of treatment arms minus one.

show.subtitles	a logical flag controlling whether to show subtitles.
Xdistr	a string indicating the way the distribution of the moderator will be plotted: "histogram" (or "hist"), "density", or "none". The default is "histogram".
main	a string that controls the title of the plot.
Ylabel	a string that controls the label of the outcome variable Y in the plot.
Dlabel	a string that controls the label of the treatment variable D in the plot.
Xlabel	a string that controls the label of the moderating variable X in the plot.
xlab	a string that specifies the label of the x-axis.
ylab	a string that specifies the label of the y-axis.
xlim	a two-element numeric vector that controls the range of the x-axis to be shown in the plot.
ylim	a two-element numeric vector that controls the range of the y-axis to be shown in the plot (with small adjustments to improve aesthetics).
theme.bw	a logical flag specifying whether to use a black-white theme.
show.grid	a logical flag indicating whether to show grid in the plot.
cex.main	a numeric value that controls the font size of the plot title.
cex.sub	a numeric value that controls the font size of subtitles.
cex.lab	a numeric value that controls the font size of the axis labels.
cex.axis	a numeric value that controls the font size of the axis numbers.
interval	draw two vertical lines to demonstrate the interval used in replicated papers.
file	a string that specifies the filename in which the plot is saved.
ncols	an integer that controls the number of columns in the plot if pool is FALSE.
pool	a logical flag specifying whether to draw marginal effects or expected values of Y for each treatment arm in a single plot. The default is FALSE.
color	a vector of colors that determines the color of lines when drawing the pool plot.
show.all	a logical flag specifying whether to return the plots for each kinds of treatment level. If setting to TRUE, the program will return a list of ggplot2 objects.
legend.title	a string that specifies the title of the legend when pool is TRUE.
diff.values	a numeric vector contain 2 or 3 elements which within the range of the moderator. The treatment effects at corresponding values of the moderator will be compared.
percentile	a logical flag indicating whether to take values of the moderator on a percentile scale.
gam.k	an integer indicating the dimension of the basis used to represent the smooth term. Smaller values of k produce smoother surfaces.
angle	a vector of numbers indicating viewing angles (allows 1 to 4 values).
span	hyper-parameter for the lowess fit. Large numbers produce smoother lines (from ggplot2).
scale	Multiplicative scaling factor of printed graph.
height	Plot size (Height).
width	Plot size (Width).

Details

interflex can implement the linear, binning and kernel estimator to estimate the marginal effects of D on Y across the supports of X in linear or generalized linear models. It also provides a few diagnostic tests for classic interaction models and convenient tools for visualization.

When **estimator**="linear", the program will estimate and visualize the treatment effects/marginal effects of D on Y across the supports of X and the average treatment effects/average effects using classic linear or generalized linear models.

When **estimator**="binning", There are three steps to implement the estimator. First, the program discretizes the moderator X into several bins and create a dummy variable for each bin. The default is 3 bins, respectively corresponding to the three terciles. Second, it picks an evaluation point within each bin, where users want to estimate the conditional marginal effect of D on Y. Third, it estimates a model that includes interactions between the bin dummies and the treatment indicator, the bin dummies and the moderator X minus the evaluation points, as well as the triple interactions. The last two terms are to capture the effect of D on Y within each bin.

The binning estimator has several key advantages over the standard multiplicative interaction model.

First, the binning estimator is much more flexible as it jointly fits the interaction components of the standard model to each bin separately. The model does not impose the linear interaction effect (LIE) assumption. Instead, the conditional marginal effects can vary freely across the three bins and therefore can take on any non-linear or non-monotonic pattern that might describe the heterogeneity in the effect of D on Y across low, medium, or high levels of X.

Second, since the bins are constructed based on the support of X, the binning ensures that the conditional marginal effects are estimated at typical values of the moderator and do not rely on excessive extrapolation or interpolation.

Third, the binning estimator is easy to implement using any regression software and the standard errors for the conditional marginal effects are directly estimated by the regression so there are no need to compute linear combinations of coefficients to recover the conditional marginal effects.

Fourth, the binning estimator actually provides a generalization that nests the standard multiplicative interaction model as a special case. It can therefore serve as a formal test on the validity of a global LIE assumption imposed by the standard model. So in the special case when the standard multiplicative interaction model is correct and therefore the global LIE assumption holds, then—as the sample size grows—the marginal effect estimates from the binning estimator converge in probability on the unbiased marginal effect estimates from the standard multiplicative interaction model.

In the case of a discrete treatment, **inter.binning** displays at the bottom of the figure a stacked histogram that shows the distribution of the moderator X. In this histogram the total height of the stacked bars refers to the distribution of the moderator in the pooled sample and the red and gray shaded bars refer to the distribution of the moderator in the treatment and control groups, respectively. Adding such a histogram makes it easy to judge the degree to which there is common support in the data. In the case of a continuous treatment, it displays a histogram at the bottom that simply shows the distribution of X in the entire sample.

When **estimator**="kernel", the program will implement a kernel smoothing estimator of the marginal effect, which is an application of semi-parametric smooth varying-coefficient models. This approach is more complicated than the simple binning estimator, but provides a generalization that allows researchers to flexibly estimate the functional form of the marginal effect of D on Y across the values of X by estimating a series of local effects with a Gaussian kernel re-weighting scheme.

The kernel regression nests the standard linear interaction model as a special case. However, in the kernel regression the conditional effect of D on Y does not have to fall on a linear line as required by the linear interaction effect (LIE) assumption, but can vary freely across the range of X. In addition, if covariates Z are included in the model, the coefficients of those covariates are also allowed to vary freely across the range of X resulting in a very flexible estimator that also helps to guard against misspecification bias with respect to the covariates.

Value

<code>est.lin</code>	marginal effects/treatment effects estimated using the linear estimator.
<code>est.bin</code>	marginal effects/treatment effects estimated using the binning estimator.
<code>est.kernel</code>	marginal effects/treatment effects estimated using the kernel estimator.
<code>pred.lin</code>	the expected value of Y conditional on X, D and Z estimated using the linear estimator.
<code>pred.bin</code>	the expected value of Y conditional on X, D and Z estimated using the binning estimator.
<code>pred.kernel</code>	the expected value of Y conditional on X, D and Z estimated using the kernel estimator.
<code>diff.estimate</code>	stores the difference between marginal effects/treatment effects at corresponding values specified in diff.values .
<code>vcov.matrix</code>	stores the variance-covariance matrix of the series of point estimates.
<code>Avg.estimate</code>	stores the average treatment effects/marginal effects.
<code>CV.output</code>	results from cross-validation if bw is not specified.
<code>bw</code>	specified or cross-validated bandwidth.
<code>figure</code>	stores the graphic output.

Author(s)

Jens Hainmueller; Jonathan Mummolo; Yiqing Xu (Maintainer); Ziyi Liu

References

Jens Hainmueller; Jonathan Mummolo; Yiqing Xu. 2019. "How Much Should We Trust Estimates from Multiplicative Interaction Models? Simple Tools to Improve Empirical Practice." *Political Analysis*, Vol. 27, Iss. 2, April 2019, pp. 163–192. Available at: <https://www.cambridge.org/core/journals/political-analysis/article/how-much-should-we-trust-estimates-from-multiplicative-interaction-models-simple-tools-to-improve-empirical-practice/00000000000000000000000000000000>.

Examples

```
library(interflex)
data(interflex)
s1.binning <-interflex(estimator = 'binning', Y = "Y", D = "D", X = "X",
data = s1)

s1.linear <-interflex(estimator = 'linear', Y = "Y", D = "D", X = "X",
data = s1)
```

interflex-internal *Internal Interflex Functions*

Description

Internal Interflex functions

Details

These are not to be called by the user, or in some cases are just waiting for proper documentation to be written :).

plot.interflex *Plotting Marginal Effect Estimates*

Description

Plotting marginal effect estimates after either the binning estimator or the kernel estimator is applied

Usage

```
## S3 method for class 'interflex'
plot(x, order = NULL, subtitles = NULL,
     show.subtitles = NULL, CI = NULL, diff.values = NULL, Xdistr = "histogram",
     main = NULL, Ylabel = NULL, Dlabel = NULL, Xlabel = NULL, xlab = NULL,
     ylab = NULL, xlim = NULL, ylim = NULL, theme.bw = FALSE, show.grid = TRUE,
     cex.main = NULL, cex.sub = NULL, cex.lab = NULL, cex.axis = NULL,
     bin.labs = TRUE, interval = NULL, file = NULL, ncols = NULL, pool = FALSE,
     legend.title = NULL, color = NULL, show.all = FALSE,
     scale = 1.1, height = 7, width = 10,
     ...)
```

Arguments

x	an interflex object.
order	a vector of strings that determines the order of treatment arms in the plot when visualizing marginal effects. It should contain all kinds of treatment arms except for the base group or all elements in "to.show".
subtitles	a vector of strings that determines the subtitles of subplots. Its length should be equal to the number of categories of treatment arms minus one.
show.subtitles	a logical flag controlling whether to show subtitles. Default to TRUE when the number of treatment arms is larger than 2.

CI	a logical flag indicating whether the confidence intervals need to be shown.
diff.values	a two-element numeric vector. The treatment effects at corresponding values of the moderator will be marked on the plot.
Xdistr	a string indicating the way the distribution of the moderator will be plotted: "histogram" (or "hist"), "density", or "none". The default is "histogram".
main	a string that controls the title of the plot.
Ylabel	a string that controls the label of the outcome variable Y in the plot.
Dlabel	a string that controls the label of the treatment variable D in the plot.
Xlabel	a string that controls the label of the moderating variable X in the plot.
xlab	a string that specifies the label of the x-axis.
ylab	a string that specifies the label of the y-axis.
xlim	a two-element numeric vector that controls the range of the x-axis to be shown in the plot.
ylim	a two-element numeric vector that controls the range of the y-axis to be shown in the plot (with small adjustments to improve aesthetics).
theme.bw	a logical flag specifying whether to use a black-white theme.
show.grid	a logical flag indicating whether to show grid in the plot.
cex.main	a numeric value that controls the font size of the plot title.
cex.sub	a numeric value that controls the font size of subtitles.
cex.lab	a numeric value that controls the font size of the axis labels.
cex.axis	a numeric value that controls the font size of the axis numbers.
bin.labs	a logical flag indicating whether to show a label for each bin in a binning plot; only meaningful when nbins is 3 or 4. The default is TRUE.
interval	draw two vertical lines to demonstrate the interval used in replicated papers.
file	a string that specifies the filename in which the plot is saved.
ncols	an integer that controls the number of columns in the plot if pool is FALSE.
pool	a logical flag specifying whether to integrate marginal effects for each treatment arm in one plot.
legend.title	a string that specifies the title of the legend when pool is TRUE.
color	a vector of colors that determines the color of lines when pool is TRUE.
show.all	a logical flag specifying whether to return the plots for each kinds of treatment level. If setting to TRUE, the program will return a list of ggplot2 objects.
scale	Multiplicative scaling factor of printed graph.
height	Plot size (Height).
width	Plot size (Width).
...	other argv.

Details

plot.interflex draws a margining effect plot based on a **interflex** object after either binning model or kernel model is applied. It allows users to flexibly change the look of a plot without re-estimating the model, hence saving time.

Value

graph stores the graphic output, a **ggplot2** object.

Author(s)

Jens Hainmueller; Jonathan Mummolo; Yiqing Xu (Maintainer); Ziyi Liu

References

Jens Hainmueller; Jonathan Mummolo; Yiqing Xu. 2019. "How Much Should We Trust Estimates from Multiplicative Interaction Models? Simple Tools to Improve Empirical Practice." *Political Analysis*, Vol. 27, Iss. 2, April 2019, pp. 163–192. Available at: <https://www.cambridge.org/core/journals/political-analysis/article/how-much-should-we-trust-estimates-from-multiplicative-interaction-models-simple-tools-to-improve-empirical-practice/0D8CAACB473F9B1EE256F43B38E458706>.

See Also

[interflex](#) and [predict.interflex](#)

predict.interflex *Plotting Marginal Effect Estimates*

Description

Plotting expected outcomes given fixed values of the treatment and moderator after either the linear, binning or the kernel estimator is applied.

Usage

```
## S3 method for class 'interflex'
predict(object, type = 'response', order = NULL,
  subtitles = NULL, show.subtitles = NULL, Xdistr = "histogram", CI = NULL, pool = FALSE,
  main = NULL, Ylabel = NULL, Xlabel = NULL, xlab = NULL, ylab = NULL,
  xlim = NULL, ylim = NULL, theme.bw = FALSE, show.grid = TRUE,
  cex.main = NULL, cex.sub = NULL, cex.lab = NULL, cex.axis = NULL,
  color = NULL, file = NULL, interval = NULL,
  legend.title = NULL, ncols = NULL, show.all = FALSE,
  scale = 1.1, height = 7, width = 10,
  ...)
```

Arguments

object an **interflex** object (**predict==TRUE**).

type a string indicating the type of prediction. Can be "response"(default) or "link".

order a vector of strings that determines the order of treatment arms in the plot when visualizing expected values. It should contain all treatment arms or all elements in "to.show".

subtitles	a vector of strings that determines the subtitles of subplots when pool is FALSE, or determines the label in the legend when pool is TRUE.
show.subtitles	a logical flag controlling whether to show subtitles. Default to TRUE when the number of treatment arms is larger than 2.
Xdistr	a string indicating the way the distribution of the moderator will be plotted: "histogram" (or "hist"), "density", or "none". The default is "histogram".
CI	a logical flag indicating whether the confidence intervals need to be shown.
pool	a logical flag specifying whether to integrate expected values of Y for each treatment arm in one plot.
main	a string that controls the title of the plot.
Ylabel	a string that controls the label of the outcome variable Y in the plot.
Xlabel	a string that controls the label of the moderating variable X in the plot.
xlab	a string that specifies the label of the x-axis.
ylab	a string that specifies the label of the y-axis.
xlim	a two-element numeric vector that controls the range of the x-axis to be shown in the plot.
ylim	a two-element numeric vector that controls the range of the y-axis to be shown in the plot (with small adjustments to improve aesthetics).
theme.bw	a logical flag specifying whether to use a black-white theme.
show.grid	a logical flag indicating whether to show grid in the plot.
cex.main	a numeric value that controls the font size of the plot title.
cex.sub	a numeric value that controls the font size of subtitles.
cex.lab	a numeric value that controls the font size of the axis labels.
cex.axis	a numeric value that controls the font size of the axis numbers.
color	a vector of colors that determines the color of lines when pool is TRUE.
file	a string that specifies the filename in which the plot is saved.
interval	draw two vertical lines to demonstrate the interval used in replicated papers.
legend.title	a string that specifies the title of the legend when pool is TRUE.
ncols	na integral that controls the number of columns in the plot if pool is FALSE.
show.all	a logical flag specifying whether to return the plots for each kinds of treatment level. If setting to TRUE, the program will return a list of ggplot2 objects.
scale	Multiplicative scaling factor of printed graph.
height	Plot size (Height).
width	Plot size (Width).
...	other argv.

Details

predict.interflex visualize expected outcomes given fixed values of the treatment and moderator after either the linear, binning or the kernel estimator is applied. It allows users to flexibly change the look of a plot without re-estimating the model, hence saving time.

Value

graph stores the graphic output, a **ggplot2** object.

Author(s)

Jens Hainmueller; Jonathan Mummolo; Yiqing Xu (Maintainer); Ziyi Liu

References

Jens Hainmueller; Jonathan Mummolo; Yiqing Xu. 2019. "How Much Should We Trust Estimates from Multiplicative Interaction Models? Simple Tools to Improve Empirical Practice." *Political Analysis*, Vol. 27, Iss. 2, April 2019, pp. 163–192. Available at: <https://www.cambridge.org/core/journals/political-analysis/article/how-much-should-we-trust-estimates-from-multiplicative-in-D8CAACB473F9B1EE256F43B38E458706>.

See Also

[interflex](#) and [plot.interflex](#)

RcppArmadillo-Functions

Set of functions in example RcppArmadillo package

Description

These four functions are created when `RcppArmadillo.package.skeleton()` is invoked to create a skeleton packages.

Usage

```
rcpparma_hello_world()  
rcpparma_outerproduct(x)  
rcpparma_innerproduct(x)  
rcpparma_bothproducts(x)
```

Arguments

x a numeric vector

Details

These are example functions which should be largely self-explanatory. Their main benefit is to demonstrate how to write a function using the Armadillo C++ classes, and to have to such a function accessible from R.

Value

`rcpparma_hello_world()` does not return a value, but displays a message to the console.

`rcpparma_outerproduct()` returns a numeric matrix computed as the outer (vector) product of `x`.

`rcpparma_innerproduct()` returns a double computer as the inner (vector) product of `x`.

`rcpparma_bothproducts()` returns a list with both the outer and inner products.

Author(s)

Dirk Eddelbuettel

References

See the documentation for Armadillo, and RcppArmadillo, for more details.

Examples

```
x <- sqrt(1:4)
rcpparma_innerproduct(x)
rcpparma_outerproduct(x)
```

s1

s1

Description

Simulated dataset 1: dichotomous treatment with linear marginal effects.

Usage

s1

Format

dataframe

s2	s2
----	----

Description

Simulated dataset 2: continuous treatment with linear marginal effects.

Usage

s2

Format

dataframe

s3	s3
----	----

Description

Simulated dataset 3: dichotomous treatment with nonlinear marginal effects.

Usage

s3

Format

dataframe

s4	s4
----	----

Description

Simulated dataset 4: dichotomous treatment with fixed effects.

Usage

s4

Format

dataframe

s5	s5
----	----

Description

Simulated dataset 5: discrete treatments with nonlinear marginal effects.

Usage

s5

Format

dataframe

s6	s6
----	----

Description

Simulated dataset 6: dichotomous treatment with logit linear funtion.

Usage

s6

Format

dataframe

s7	s7
----	----

Description

Simulated dataset 7: continuous treatment with logit marginal effects.

Usage

s7

Format

dataframe

s8	s8
----	----

Description

Simulated dataset 8: dichotomous treatment with a logit link function and a nonlinear linear predictor.

Usage

s8

Format

dataframe

s9	s9
----	----

Description

Simulated dataset 9: Multiple-arms treatment with a logit link function.

Usage

s9

Format

dataframe

Index

- * **datasets**
 - s1, [15](#)
 - s2, [16](#)
 - s3, [16](#)
 - s4, [16](#)
 - s5, [17](#)
 - s6, [17](#)
 - s7, [17](#)
 - s8, [18](#)
 - s9, [18](#)
- * **graphics**
 - plot.interflex, [10](#)
 - predict.interflex, [12](#)
- * **ts**
 - interflex-internal, [10](#)

[createFolds \(interflex-internal\)](#), [10](#)

[fastplm \(interflex-internal\)](#), [10](#)

[inter.test](#), [3](#)

[interflex](#), [2](#), [3](#), [4](#), [12](#), [14](#)

[interflex-internal](#), [10](#)

[interflex-package](#), [2](#)

[interflex.binning \(interflex-internal\)](#), [10](#)

[interflex.gam \(interflex-internal\)](#), [10](#)

[interflex.kernel \(interflex-internal\)](#), [10](#)

[interflex.linear \(interflex-internal\)](#), [10](#)

[interflex.plot.pool \(interflex-internal\)](#), [10](#)

[interflex.raw \(interflex-internal\)](#), [10](#)

[iv_fastplm \(interflex-internal\)](#), [10](#)

[plot.interflex](#), [2](#), [3](#), [10](#), [14](#)

[predict.interflex](#), [2](#), [3](#), [12](#), [12](#)

[rcpparma_bothproducts \(RcppArmadillo-Functions\)](#), [14](#)

[rcpparma_hello_world \(RcppArmadillo-Functions\)](#), [14](#)

[rcpparma_innerproduct \(RcppArmadillo-Functions\)](#), [14](#)

[rcpparma_outerproduct \(RcppArmadillo-Functions\)](#), [14](#)

[RcppArmadillo-Functions](#), [14](#)

s1, [15](#)

s2, [16](#)

s3, [16](#)

s4, [16](#)

s5, [17](#)

s6, [17](#)

s7, [17](#)

s8, [18](#)

s9, [18](#)

[vcovCluster \(interflex-internal\)](#), [10](#)