

Package ‘interflex’

September 21, 2026

Type Package

Version 1.4.1

Title Estimation, Diagnostics and Visualization of Conditional Marginal Effects

Date 2026-09-20

Description Performs estimation, diagnostics, and visualization of conditional marginal effects and group average treatment effects of a treatment on an outcome across different values of a moderator. Optionally integrates with the 'mlr3extralearners' package for additional machine learning backends compatible with the double machine learning estimators. 'mlr3extralearners' is not on CRAN but can be obtained from <<https://github.com/mlr-org/mlr3extralearners>>.

URL <https://yiqingxu.org/packages/interflex/>,
<https://github.com/xuyiqing/interflex>

BugReports <https://github.com/xuyiqing/interflex/issues>

Additional_repositories <https://mlr-org.r-universe.dev>

License MIT + file LICENSE

Imports ggplot2 (>= 2.1.0), sandwich (>= 2.3-4), Lmoments (>= 1.2-3), doFuture, doRNG, foreach (>= 1.4.3), mgcv (>= 1.8-16), lfe (>= 2.6-2291), gridExtra, grid, ggplotify, fixest, RColorBrewer, pcse, gtable, MASS, mvtnorm, pROC, ModelMetrics, lmttest, AER, future, splines, glmnet, grf, DoubleML, mlr3, mlr3learners, data.table, paradox, parallelly, cli, progressr

Suggests mlr3tuning, nnet, mlr3extralearners, lightgbm, xgboost, ranger, testthat (>= 3.0.0)

LazyData true

NeedsCompilation no

Maintainer Yiqing Xu <yiqingxu@stanford.edu>

Depends R (>= 3.5.0)

Author Yiqing Xu [aut, cre],
Jens Hainmueller [aut],
Jonathan Mummolo [aut],
Tianzhu Qin [aut],
Jiehan Liu [aut],
Ziyi Liu [aut]

Repository CRAN

Date/Publication 2026-09-21 06:40:02 UTC

Contents

app-data	2
inter.test	3
interflex	4
plot.interflex	12
predict.interflex	14
Index	17

app-data	<i>Replication Datasets from Published Interaction-Effect Studies</i>
----------	---

Description

A collection of data frames from published political science studies that feature conditional marginal effects. These datasets are shipped with **interflex** so users can reproduce the applied examples in the package’s user manual (<https://yiqingxu.org/packages/interflex/>) and experiment with the provided estimators on real data.

Usage

data(interflex)

Format

Each object is a data.frame. Column sets differ across studies; the treatment variable, moderator, and outcome used in the manual examples are listed below.

app_adiguzel2023 1,295 rows, 10 columns. Treatment: Dakp_3. Moderator: lograyic09. Used in the user manual to illustrate linear and kernel estimators on a discrete treatment with continuous moderator.

app_bb2024 810 rows, 22 columns. Treatment: treat_rev. Outcome: sup_tax. Moderator: clim_belief. Multi-arm treatment example.

app_et2023 570 rows, 39 columns. Treatment: board_service. Outcome: seat_next_5yr. Illustrates binning and fixed-effect estimators on legislative career data.

app_hma2015 1,482 rows, 11 columns. Treatment: threat (binary). Outcome: totangry. Moderator: pidentity. Based on Huddy, Mason, and Aarøe (2015); used throughout the manual as the primary binary-treatment running example.

app_vernby2013 183 rows, 9 columns. Treatment: noncitvotsh. Outcome: school_diff. Moderator: Taxbase2. Small-sample regional dataset used in the extensions chapter.

Source

Each dataset is redistributed with permission from its original study’s replication archive. See the user manual for full citations and links to the source archives: <https://yiqingxu.org/packages/interflex/>.

References

Adiguzel, F. S. (2023). Bernauer, J. and Bodnar, M. (2024). Eggers, A. C. and Tanaka, S. (2023). Huddy, L., Mason, L. and Aarøe, L. (2015). Expressive partisanship: Campaign involvement, political emotion, and partisan identity. *American Political Science Review* 109(1), 1-17. Vernby, K. (2013). Inclusion and public policy: Evidence from Sweden’s introduction of noncitizen suffrage. *American Journal of Political Science* 57(1), 15-29.

inter.test	Conduct t-Test after Estimation
------------	---------------------------------

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.

Value

A named list with one element per non-baseline treatment arm (discrete treatment) or per reference value of the treatment (continuous treatment). Each element is a data frame with one row per comparison of the moderator values in `diff.values` (row names such as " 0.5 vs 0.25 ") and columns `diff.estimate`, `sd`, `z-value`, `p-value`, `lower CI(95%)` and `upper CI(95%)`, stored as character strings rounded to three decimals.

Author(s)

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

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/0D8CAACB473F9B1EE256F43B38E458706>.

See Also

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

interflex	<i>Flexible Interactive Models</i>
-----------	------------------------------------

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 = NULL, cl = NULL, Z.ref = NULL,
  D.ref = NULL, nbins = 3, cutoffs = NULL, wald = TRUE, bw = NULL,
  kfold = 10, grid = 30, metric = NULL, bw.select = "cv.min",
  bw.se.mult = 1, bw.ess.min = NULL, model.y = "rf",
  param.y = NULL, param.grid.y = NULL,
  scoring.y = "neg_mean_squared_error", model.t = "rf", param.t = NULL,
  param.grid.t = NULL, scoring.t = "neg_mean_squared_error", CV = FALSE,
  n.folds = 10, n.jobs = -1, cf.n.folds = 5, gate = FALSE,
  grf.num.trees = 4000, signal = c("aipw"), estimand = c("ATE"),
  basis.type = c("bspline"), include.interactions = TRUE, poly.degree = 2,
```

```
spline.df = 4, spline.degree = 2, lambda.seq = NULL,
reduce.dimension = c("kernel"), 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 = TRUE, show.grid = TRUE, show.uniform.CI = 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, box.pos = "down",
verbose = TRUE)
```

Arguments

estimator	a string that specifies the estimator. Can be one of the following: "raw", "linear", "binning", "kernel", "gam", "grf", "dml", or "lasso".
data	a data.frame that stores Y , D , X (and optionally Z , IV , FE , weights).
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 specifying the type of the treatment. Must be "discrete" or "continuous". If not specified, it is set automatically: if D is numeric and has more than 5 unique values, then "continuous"; otherwise "discrete".
base	a string specifying the baseline treatment group when treat.type is "discrete". If not specified, the baseline is chosen as the first level in sorted order of D .
Z	a character vector of covariate names.
IV	a character vector of instrument variable names (only supported when method = "linear").
FE	a character vector of fixed-effect indicator names (only supported when method = "linear").
full.moderate	a logical flag indicating whether to use a fully moderated model, i.e., include additional interactions between X and other covariates Z .
weights	a string, name of the weighting variable.
na.rm	a logical flag indicating whether to list-wise delete missing data in required variables. If na.rm = FALSE, the function errors when missing values exist.
Xunif	a logical flag controlling whether to transform X to its percentile rank (scaled to 0–100).
CI	a logical flag indicating whether to compute confidence intervals (if CI = FALSE, CI-related bootstrap/simulation routines are skipped when applicable).
neval	an integer specifying the number of evaluation points/grid size used by some estimators (larger values produce smoother curves but take longer).
X.eval	a numeric vector of manually selected evaluation points for X (used by some estimators such as estimator ="linear" or "kernel").

method	a string specifying the outcome model: one of "linear", "probit", "logit", "poisson", or "nbinom".
vartype	a string controlling uncertainty estimation for estimator = "linear" or "binning": one of "delta", "simu", or "bootstrap".
vcov.type	a string controlling the variance-covariance estimator for linear models: one of "homoscedastic", "robust", "cluster", or "pcse" (panel-corrected SEs; only for linear models).
time	a string specifying the time variable name for vcov.type = "pcse".
pairwise	a logical flag for vcov.type = "pcse" controlling the pairwise option.
nboots	an integer specifying the number of bootstrap runs (when bootstrap-based inference is used).
nsimu	an integer specifying the number of simulation draws (when simulation-based inference is used).
parallel	a logical flag indicating whether to use parallel computing in bootstrap and/or cross-validation routines (when supported by the chosen estimator).
cores	an integer indicating the number of cores to be used when parallel = TRUE.
cl	a string specifying the clustering/unit variable name for vcov.type = "cluster" or vcov.type = "pcse". It is also used to block some resampling routines when applicable.
Z.ref	a vector specifying the reference values of covariates when estimating marginal effects/predicted values in GLM-type models. Must have the same length and order as the original Z . For factor covariates, provide a level name; for numeric covariates, provide a numeric value. By default, numeric covariates are set to their sample mean (factor covariates are handled internally via contrast coding).
D.ref	a numeric vector specifying reference values of D when treat.type = "continuous" (used for labeling/plotting). If not specified, it defaults to the median of D .
nbins	an integer determining the number of bins when discretizing the moderator for estimator = "binning" (and also used in estimator = "raw" for some plots). When quantile cutpoints are not unique, the effective number of bins may be reduced automatically.
cutoffs	a numeric vector of custom cutpoints for X . If supplied, nbins is ignored and the number of bins equals <code>length(cutoffs) + 1</code> .
wald	a logical flag controlling whether to conduct a Wald test (used by some estimators such as estimator = "binning").
bw	a positive number setting the bandwidth in Gaussian kernel reweighting for estimator = "kernel". With adaptive bandwidth, bw is the bandwidth at a typical observation, where the moderator's density equals its geometric mean over the observations: windows are wider where data are sparse and narrower where they are dense. Also used as the fixed bandwidth of the kernel smoothing step of estimator = "lasso" with a continuous treatment.
kfold	an integer specifying the number of folds used in cross-validation for bandwidth selection (kernel estimator).

grid	either an integer indicating the number of candidate bandwidths in a grid search, or a numeric vector of candidate bandwidths (kernel estimator; ignored if bw is specified).
metric	a string specifying the criterion used in cross-validation to select bandwidth: one of "MSE", "MAE", "Cross Entropy", or "AUC". Default is "MSE" for non-binary outcomes and "Cross Entropy" for binary outcomes.
bw.select	a string specifying how the bandwidth is chosen from the candidate grid when bw is not supplied (estimator = "kernel", and the kernel smoothing step of estimator = "lasso" with a continuous treatment). "cv.min" (default) takes the candidate with the smallest cross-validated loss. "cv.lse" takes the largest candidate whose loss is within bw.se.mult standard errors of the smallest, which prefers the smoother fit when cross-validation cannot separate the candidates. "cv.lse.ess" applies the same rule after discarding candidates whose local fits rest on fewer than bw.ess.min effective observations at some evaluation point. "ess" skips cross-validation and takes the smallest candidate that satisfies that effective-sample-size constraint.
bw.se.mult	a non-negative number: the standard-error multiplier used by bw.select = "cv.lse" and "cv.lse.ess". Default is 1.
bw.ess.min	a positive number: the minimum local effective sample size required by bw.select = "cv.lse.ess" and "ess". If not specified, it is set to a multiple of the number of local coefficients (at least 20 observations for a discrete treatment and 30 for a continuous one).
model.y	outcome model used by estimator ="dml" (and also used by estimator ="lasso"). In the DML estimator, this string is passed to the underlying Python implementation (default "rf").
param.y	named list of hyper-parameters for the DML outcome model (passed to Python).
param.grid.y	named list specifying a hyper-parameter search grid for the DML outcome model (passed to Python; used when CV = TRUE).
scoring.y	a string specifying the scoring criterion used for tuning the DML outcome model (passed to Python; default "neg_mean_squared_error").
model.t	treatment/propensity model used by estimator ="dml" (and also used by estimator ="lasso"). In the DML estimator, this string is passed to the underlying Python implementation (default "rf").
param.t	named list of hyper-parameters for the DML treatment model (passed to Python).
param.grid.t	named list specifying a hyper-parameter search grid for the DML treatment model (passed to Python; used when CV = TRUE).
scoring.t	a string specifying the scoring criterion used for tuning the DML treatment model (passed to Python; default "neg_mean_squared_error").
CV	a logical flag indicating whether to tune the DML nuisance models via cross-validation (DML estimator).
n.folds	an integer specifying the number of folds used in DML cross-fitting (DML estimator).
n.jobs	an integer specifying the number of parallel jobs used in the Python implementation for DML (default -1 uses all available cores).

<code>cf.n.folds</code>	an integer specifying the number of folds for internal cross-fitting (DML estimator).
<code>gate</code>	a logical flag indicating whether to compute group-specific ATEs (GATE) in the DML estimator (if supported by the underlying Python routine).
<code>grf.num.trees</code>	an integer specifying the number of trees for the generalized random forest when estimator ="grf".
<code>signal</code>	(LASSO estimator) which signal to use when treat.type ="discrete": one of "outcome", "ipw", or "aipw".
<code>estimand</code>	(LASSO estimator) the target estimand when treat.type ="discrete": one of "ATE" or "ATT".
<code>basis.type</code>	(LASSO estimator) basis expansion for the moderator: one of "polynomial", "bspline", or "none".
<code>include.interactions</code>	(LASSO estimator) whether to include interactions in the basis expansion.
<code>poly.degree</code>	(LASSO estimator) degree of the polynomial basis when basis.type ="polynomial".
<code>spline.df</code>	(LASSO estimator) degrees of freedom for B-splines when basis.type ="bspline".
<code>spline.degree</code>	(LASSO estimator) degree of B-splines when basis.type ="bspline".
<code>lambda.seq</code>	(LASSO estimator) an optional sequence of lambda values passed to <code>glmnet</code> .
<code>reduce.dimension</code>	(LASSO estimator) method used for dimension reduction in the continuous-treatment case (e.g., "bspline" or "kernel").
<code>figure</code>	a logical flag controlling whether to draw a marginal effect plot.
<code>bin.labs</code>	a logical flag indicating whether to show a label for each bin in a binning plot; only meaningful when nbins is 2–4.
<code>order</code>	controls the order of panels/series in plots. For discrete treatments, it should contain all treatment arms except the baseline group. For continuous treatments, it can be a numeric vector of the reference D values (matching D.ref).
<code>subtitles</code>	a vector controlling the subtitles of subplots when pool = FALSE, or legend labels when pool = TRUE.
<code>show.subtitles</code>	a logical flag controlling whether to show subtitles.
<code>Xdistr</code>	a string indicating how the distribution of the moderator will be plotted: "histogram" (or "hist"), "density", or "none".
<code>main</code>	a string controlling the title of the plot.
<code>Ylabel</code>	a string controlling the label of the outcome variable Y in the plot.
<code>Dlabel</code>	a string controlling the label of the treatment variable D in the plot.
<code>Xlabel</code>	a string controlling the label of the moderator X in the plot.
<code>xlab</code>	a string specifying the label of the x-axis.
<code>ylab</code>	a string specifying the label of the y-axis.
<code>xlim</code>	a two-element numeric vector controlling the x-axis range.
<code>ylim</code>	a two-element numeric vector controlling the y-axis range.
<code>theme.bw</code>	a logical flag specifying whether to use a black-white theme.

show.grid	a logical flag indicating whether to show grid lines in the plot.
show.uniform.CI	a logical flag indicating whether to show uniform/simultaneous confidence intervals when they are available in the output (plotted as dashed lines).
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 axis tick labels.
interval	draw two vertical lines to demonstrate an interval used in replicated papers.
file	a string specifying the filename to save the plot (if not NULL).
ncols	an integer controlling the number of columns in the plot when pool is FALSE.
pool	a logical flag specifying whether to draw all series in a single pooled plot.
color	a vector of colors that determines the color of lines in pooled plots.
show.all	a logical flag specifying whether to return all subplot objects (if TRUE, returns a list of ggplot objects).
legend.title	a string specifying the title of the legend when pool is TRUE.
diff.values	a numeric vector (length 2 or 3) specifying moderator values at which treatment effects are compared (only available for some estimators; estimator ="binning" does not support this feature).
percentile	a logical flag indicating whether diff.values should be interpreted on a percentile scale (between 0 and 1).
gam.k	(GAM estimator) dimension of the basis used to represent smooth terms.
angle	(GAM estimator) a vector of numbers indicating viewing angles (allows 1 to 4 values).
span	(RAW estimator) hyper-parameter for the loess fit (larger values produce smoother lines).
scale	multiplicative scaling factor for saving the graph.
height	plot height (in inches) when saving.
width	plot width (in inches) when saving.
box.pos	(RAW estimator) where to place the reference box in discrete-treatment raw plots: "down" (near minimum of Y) or "up" (near maximum of Y).
verbose	a logical flag indicating whether to print informational messages and progress bars during estimation. Set to FALSE for a quieter run.

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>bw.select</code>	the bandwidth selection rule that was applied (kernel estimator; also stored for the lasso estimator with a continuous treatment).
<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-in-D8CAACB473F9B1EE256F43B38E458706>.

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)
```

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 = TRUE, 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,
     by.group = FALSE, legend.title = NULL, line.size = 1.2, color = NULL,
     line.color = "black", CI.color = "black", CI.color.alpha = 0.2,
     hist.color = c("gray50", "red"), hist.color.alpha = 0.3,
     density.color = c("gray50", "red"), density.color.alpha = 0.3,
     show.all = FALSE, show.uniform.CI = TRUE,
     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.
by.group	a logical flag specifying whether to draw one panel per group when the fitted object carries group-average treatment effects (g.est), which requires fitting with gate = TRUE and a discrete moderator.
legend.title	a string that specifies the title of the legend when pool is TRUE.
line.size	a numeric value controlling the width of the effect curve.
color	a vector of colors that determines the color of lines when pool is TRUE.
line.color	a string specifying the colour of the effect curve.
CI.color	a string specifying the fill colour of the pointwise confidence-interval ribbon.
CI.color.alpha	a numeric value in [0, 1] controlling the transparency of the confidence-interval ribbon.
hist.color	a length-two character vector of colours used for the control/treated histogram bars below the effect curve.
hist.color.alpha	a numeric value in [0, 1] controlling the transparency of the histogram bars.
density.color	a length-two character vector of colours used for the control/treated density ribbons below the effect curve (when Xdistr = "density").
density.color.alpha	a numeric value in [0, 1] controlling the transparency of the density ribbons.
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.
show.uniform.CI	a logical flag controlling whether the uniform (simultaneous) confidence band, if available on the fitted object, is overlaid on the pointwise band.

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-in-D8CAACB473F9B1EE256F43B38E458706>.

See Also

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

predict.interflex	<i>Plotting Marginal Effect Estimates</i>
-------------------	---

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 = TRUE, 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	an integer that controls the number of columns in the plot if pool is FALSE.

<code>show.all</code>	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.
<code>scale</code>	Multiplicative scaling factor of printed graph.
<code>height</code>	Plot size (Height).
<code>width</code>	Plot size (Width).
<code>...</code>	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](#)

Index

* datasets

app-data, [2](#)

* graphics

plot.interflex, [12](#)

predict.interflex, [14](#)

app-data, [2](#)

app_adiguzel2023 (app-data), [2](#)

app_bb2024 (app-data), [2](#)

app_et2023 (app-data), [2](#)

app_hma2015 (app-data), [2](#)

app_vernby2013 (app-data), [2](#)

inter.test, [3](#)

interflex, [4](#), [4](#), [14](#), [16](#)

plot.interflex, [4](#), [12](#), [16](#)

predict.interflex, [4](#), [14](#), [14](#)