

# Package ‘vimp’

June 3, 2021

**Type** Package

**Title** Perform Inference on Algorithm-Agnostic Variable Importance

**Version** 2.2.1

**Description** Calculate point estimates of and valid confidence intervals for nonparametric, algorithm-agnostic variable importance measures in high and low dimensions, using flexible estimators of the underlying regression functions. For more information about the methods, please see Williamson et al. (Biometrics, 2020), Williamson et al. (arXiv, 2020+) <[arXiv:2004.03683](https://arxiv.org/abs/2004.03683)>, and Williamson and Feng (ICML, 2020).

**Depends** R (>= 3.1.0)

**Imports** SuperLearner, stats, dplyr, magrittr, ROCR, tibble, rlang, MASS, boot, data.table

**Suggests** knitr, rmarkdown, gam, xgboost, glmnet, ranger, polyspline, quadprog, covr, testthat, ggplot2, cowplot, RCurl, cvAUC

**License** MIT + file LICENSE

**URL** <https://bdwilliamson.github.io/vimp/>,  
<https://github.com/bdwilliamson/vimp>

**BugReports** <https://github.com/bdwilliamson/vimp/issues>

**RoxygenNote** 7.1.1

**VignetteBuilder** knitr

**NeedsCompilation** no

**Author** Brian D. Williamson [aut, cre]  
(<<https://orcid.org/0000-0002-7024-548X>>),  
Jean Feng [ctb],  
Noah Simon [ths] (<<https://orcid.org/0000-0002-8985-2474>>),  
Marco Carone [ths] (<<https://orcid.org/0000-0003-2106-0953>>)

**Maintainer** Brian D. Williamson <[bwillia2@fredhutch.org](mailto:bwillia2@fredhutch.org)>

**Repository** CRAN

**Date/Publication** 2021-06-03 16:50:02 UTC

**R topics documented:**

average_vim . . . . .	3
bootstrap_se . . . . .	4
check_fitted_values . . . . .	5
check_inputs . . . . .	6
create_z . . . . .	7
cv_vim . . . . .	7
est_predictiveness . . . . .	13
est_predictiveness_cv . . . . .	14
extract_sampled_split_predictions . . . . .	16
format.vim . . . . .	17
get_cv_sl_folds . . . . .	17
get_full_type . . . . .	18
make_folds . . . . .	18
make_kfold . . . . .	19
measure_accuracy . . . . .	19
measure_anova . . . . .	20
measure_auc . . . . .	22
measure_cross_entropy . . . . .	23
measure_deviance . . . . .	24
measure_mse . . . . .	25
measure_r_squared . . . . .	27
merge_vim . . . . .	28
print.vim . . . . .	29
run_sl . . . . .	30
sample_subsets . . . . .	31
scale_est . . . . .	32
spvim_ics . . . . .	32
spvim_se . . . . .	33
sp_vim . . . . .	34
vim . . . . .	37
vimp . . . . .	41
vimp_accuracy . . . . .	42
vimp_anova . . . . .	47
vimp_auc . . . . .	50
vimp_ci . . . . .	54
vimp_deviance . . . . .	55
vimp_hypothesis_test . . . . .	59
vimp_regression . . . . .	60
vimp_rsquared . . . . .	63
vimp_se . . . . .	67

---

average_vim	<i>Average multiple independent importance estimates</i>
-------------	--

---

### Description

Average the output from multiple calls to `vimp_regression`, for different independent groups, into a single estimate with a corresponding standard error and confidence interval.

### Usage

```
average_vim(..., weights = rep(1/length(list(...)), length(list(...))))
```

### Arguments

<code>...</code>	an arbitrary number of <code>vim</code> objects.
<code>weights</code>	how to average the vims together, and must sum to 1; defaults to $1/(\text{number of vims})$ for each <code>vim</code> , corresponding to the arithmetic mean

### Value

an object of class `vim` containing the (weighted) average of the individual importance estimates, as well as the appropriate standard error and confidence interval. This results in a list containing:

- `s` - a list of the column(s) to calculate variable importance for
- `SL.library` - a list of the libraries of learners passed to SuperLearner
- `full_fit` - a list of the fitted values of the chosen method fit to the full data
- `red_fit` - a list of the fitted values of the chosen method fit to the reduced data
- `est` - a vector with the corrected estimates
- `naive` - a vector with the naive estimates
- `update` - a list with the influence curve-based updates
- `mat` - a matrix with the estimated variable importance, the standard error, and the  $(1 - \alpha) \times 100\%$  confidence interval
- `full_mod` - a list of the objects returned by the estimation procedure for the full data regression (if applicable)
- `red_mod` - a list of the objects returned by the estimation procedure for the reduced data regression (if applicable)
- `alpha` - the level, for confidence interval calculation
- `y` - a list of the outcomes

**Examples**

```

# generate the data
p <- 2
n <- 100
x <- data.frame(replicate(p, stats::runif(n, -5, 5)))

# apply the function to the x's
smooth <- (x[,1]/5)^2*(x[,1]+7)/5 + (x[,2]/3)^2

# generate Y ~ Normal (smooth, 1)
y <- smooth + stats::rnorm(n, 0, 1)

# set up a library for SuperLearner; note simple library for speed
library("SuperLearner")
learners <- c("SL.glm", "SL.mean")

# get estimates on independent splits of the data
samp <- sample(1:n, n/2, replace = FALSE)

# using Super Learner (with a small number of folds, for illustration only)
est_2 <- vimp_regression(Y = y[samp], X = x[samp, ], indx = 2, V = 2,
  run_regression = TRUE, alpha = 0.05,
  SL.library = learners, cvControl = list(V = 2))

est_1 <- vimp_regression(Y = y[-samp], X = x[-samp, ], indx = 2, V = 2,
  run_regression = TRUE, alpha = 0.05,
  SL.library = learners, cvControl = list(V = 2))

ests <- average_vim(est_1, est_2, weights = c(1/2, 1/2))

```

---

bootstrap_se	<i>Compute bootstrap-based standard error estimates for variable importance</i>
--------------	---

---

**Description**

Compute bootstrap-based standard error estimates for variable importance

**Usage**

```
bootstrap_se(Y = NULL, f1 = NULL, f2 = NULL, type = "r_squared", b = 1000)
```

**Arguments**

Y	the outcome.
f1	the fitted values from a flexible estimation technique regressing Y on X.
f2	the fitted values from a flexible estimation technique regressing either (a) f1 or (b) Y on X withholding the columns in indx.

type	the type of importance to compute; defaults to r_squared, but other supported options are auc, accuracy, deviance, and anova.
b	the number of bootstrap replicates (defaults to 1000)

**Value**

a bootstrap-based standard error estimate

---

check\_fitted\_values    *Check pre-computed fitted values for call to vim, cv\_vim, or sp\_vim*

---

**Description**

Check pre-computed fitted values for call to vim, cv\_vim, or sp\_vim

**Usage**

```
check_fitted_values(
  Y = NULL,
  f1 = NULL,
  f2 = NULL,
  cross_fitted_f1 = NULL,
  cross_fitted_f2 = NULL,
  sample_splitting_folds = NULL,
  cross_fitting_folds = NULL,
  V = NULL,
  ss_V = NULL,
  cv = FALSE
)
```

**Arguments**

Y	the outcome
f1	estimator of the population-optimal prediction function using all covariates
f2	estimator of the population-optimal prediction function using the reduced set of covariates
cross_fitted_f1	cross-fitted estimator of the population-optimal prediction function using all covariates
cross_fitted_f2	cross-fitted estimator of the population-optimal prediction function using the reduced set of covariates
sample_splitting_folds	the folds for sample-splitting (used for hypothesis testing)
cross_fitting_folds	the folds for cross-fitting (used for point estimates of variable importance in cv_vim and sp_vim)

V	the number of cross-fitting folds
ss_V	the number of folds for CV (if sample_splitting is TRUE)
cv	a logical flag indicating whether or not to use cross-fitting

**Details**

Ensure that inputs to `vim`, `cv_vim`, and `sp_vim` follow the correct formats.

**Value**

None. Called for the side effect of stopping the algorithm if any inputs are in an unexpected format.

---

check_inputs	<i>Check inputs to a call to vim, cv_vim, or sp_vim</i>
--------------	---

---

**Description**

Check inputs to a call to `vim`, `cv_vim`, or `sp_vim`

**Usage**

```
check_inputs(Y, X, f1, f2, indx)
```

**Arguments**

Y	the outcome
X	the covariates
f1	estimator of the population-optimal prediction function using all covariates
f2	estimator of the population-optimal prediction function using the reduced set of covariates
indx	the index or indices of the covariate(s) of interest

**Details**

Ensure that inputs to `vim`, `cv_vim`, and `sp_vim` follow the correct formats.

**Value**

None. Called for the side effect of stopping the algorithm if any inputs are in an unexpected format.

---

create_z	<i>Create complete-case outcome, weights, and Z</i>
----------	---

---

**Description**

Create complete-case outcome, weights, and Z

**Usage**

```
create_z(Y, C, Z, X, ipc_weights)
```

**Arguments**

Y	the outcome
C	indicator of missing or observed
Z	the covariates observed in phase 1 and 2 data
X	all covariates
ipc_weights	the weights

**Value**

a list, with the complete-case outcome, weights, and Z matrix

---

cv_vim	<i>Nonparametric Intrinsic Variable Importance Estimates and Inference using Cross-fitting</i>
--------	--

---

**Description**

Compute estimates and confidence intervals using cross-fitting for nonparametric intrinsic variable importance based on the population-level contrast between the oracle predictiveness using the feature(s) of interest versus not.

**Usage**

```
cv_vim(
  Y = NULL,
  X = NULL,
  cross_fitted_f1 = NULL,
  cross_fitted_f2 = NULL,
  f1 = NULL,
  f2 = NULL,
  indx = 1,
  V = length(unique(cross_fitting_folds)),
```

```

sample_splitting = TRUE,
sample_splitting_folds = NULL,
cross_fitting_folds = NULL,
stratified = FALSE,
type = "r_squared",
run_regression = TRUE,
SL.library = c("SL.glmnet", "SL.xgboost", "SL.mean"),
alpha = 0.05,
delta = 0,
scale = "identity",
na.rm = FALSE,
C = rep(1, length(Y)),
Z = NULL,
ipc_weights = rep(1, length(Y)),
ipc_est_type = "aipw",
scale_est = TRUE,
cross_fitted_se = TRUE,
bootstrap = FALSE,
b = 1000,
...
)

```

### Arguments

Y	the outcome.
X	the covariates.
cross_fitted_f1	the predicted values on validation data from a flexible estimation technique regressing Y on X in the training data; a list of length V, where each object is a set of predictions on the validation data. If sample-splitting is requested, then these must be estimated specially; see Details.
cross_fitted_f2	the predicted values on validation data from a flexible estimation technique regressing either (a) the fitted values in cross_fitted_f1, or (b) Y, on X withholding the columns in indx; a list of length V, where each object is a set of predictions on the validation data. If sample-splitting is requested, then these must be estimated specially; see Details.
f1	the fitted values from a flexible estimation technique regressing Y on X. Estimated using the whole dataset. If cross_fitted_se = TRUE, then this argument is not used.
f2	the fitted values from a flexible estimation technique regressing either (a) f1 or (b) Y on X withholding the columns in indx. Estimated using the whole dataset. If cross_fitted_se = TRUE, then this argument is not used.
indx	the indices of the covariate(s) to calculate variable importance for; defaults to 1.
V	the number of folds for cross-fitting, defaults to 5. If sample_splitting = TRUE, then a special type of V-fold cross-fitting is done. See Details for a more detailed explanation.



sample_splitting	should we use sample-splitting to estimate the full and reduced predictiveness? Defaults to TRUE, since inferences made using sample_splitting = FALSE will be invalid for variable with truly zero importance.
sample_splitting_folds	the folds to use for sample-splitting; if entered, these should result in balance within the cross-fitting folds. Only used if run_regression = FALSE and sample_splitting = TRUE. A vector of length $2 * V$ .
cross_fitting_folds	the folds for cross-fitting. Only used if run_regression = FALSE.
stratified	if run_regression = TRUE, then should the generated folds be stratified based on the outcome (helps to ensure class balance across cross-fitting folds)
type	the type of parameter (e.g., ANOVA-based is "anova").
run_regression	if outcome Y and covariates X are passed to cv_vim, and run_regression is TRUE, then Super Learner will be used; otherwise, variable importance will be computed using the inputted fitted values.
SL.library	a character vector of learners to pass to SuperLearner, if f1 and f2 are Y and X, respectively. Defaults to SL.glmnet, SL.xgboost, and SL.mean.
alpha	the level to compute the confidence interval at. Defaults to 0.05, corresponding to a 95% confidence interval.
delta	the value of the $\delta$ -null (i.e., testing if importance $< \delta$ ); defaults to 0.
scale	should CIs be computed on original ("identity", default) or logit ("logit") scale?
na.rm	should we remove NA's in the outcome and fitted values in computation? (defaults to FALSE)
C	the indicator of coarsening (1 denotes observed, 0 denotes unobserved).
Z	either (i) NULL (the default, in which case the argument C above must be all ones), or (ii) a character vector specifying the variable(s) among Y and X that are thought to play a role in the coarsening mechanism.
ipc_weights	weights for the computed influence curve (i.e., inverse probability weights for coarsened-at-random settings). Assumed to be already inverted (i.e., ipc_weights = 1 / [estimated probability weights]).
ipc_est_type	the type of procedure used for coarsened-at-random settings; options are "ipw" (for inverse probability weighting) or "aipw" (for augmented inverse probability weighting). Only used if C is not all equal to 1.
scale_est	should the point estimate be scaled to be greater than 0? Defaults to TRUE.
cross_fitted_se	should we use cross-fitting to estimate the standard errors (TRUE, the default) or not (FALSE)?
bootstrap	should bootstrap-based standard error estimates be computed? Defaults to FALSE (and currently may only be used if sample_splitting = FALSE and cross_fitted_se = FALSE).
b	the number of bootstrap replicates (only used if bootstrap = TRUE and sample_splitting = FALSE).
...	other arguments to the estimation tool, see "See also".

## Details

We define the population variable importance measure (VIM) for the group of features (or single feature)  $s$  with respect to the predictiveness measure  $V$  by

$$\psi_{0,s} := V(f_0, P_0) - V(f_{0,s}, P_0),$$

where  $f_0$  is the population predictiveness maximizing function,  $f_{0,s}$  is the population predictiveness maximizing function that is only allowed to access the features with index not in  $s$ , and  $P_0$  is the true data-generating distribution.

Cross-fitted VIM estimates are computed differently if sample-splitting is requested versus if it is not. We recommend using sample-splitting in most cases, since only in this case will inferences be valid if the variable(s) of interest have truly zero population importance. The purpose of cross-fitting is to estimate  $f_0$  and  $f_{0,s}$  on independent data from estimating  $P_0$ ; this can result in improved performance, especially when using flexible learning algorithms. The purpose of sample-splitting is to estimate  $f_0$  and  $f_{0,s}$  on independent data; this allows valid inference under the null hypothesis of zero importance.

Without sample-splitting, cross-fitted VIM estimates are obtained by first splitting the data into  $K$  folds; then using each fold in turn as a hold-out set, constructing estimators  $f_{n,k}$  and  $f_{n,k,s}$  of  $f_0$  and  $f_{0,s}$ , respectively on the training data and estimator  $P_{n,k}$  of  $P_0$  using the test data; and finally, computing

$$\psi_{n,s} := K^{(-1)} \sum_{k=1}^K \{V(f_{n,k}, P_{n,k}) - V(f_{n,k,s}, P_{n,k})\}.$$

With sample-splitting, cross-fitted VIM estimates are obtained by first splitting the data into  $2K$  folds. These folds are further divided into 2 groups of folds. Then, for each fold  $k$  in the first group, estimator  $f_{n,k}$  of  $f_0$  is constructed using all data besides the  $k$ th fold in the group (i.e.,  $(2K - 1)/(2K)$  of the data) and estimator  $P_{n,k}$  of  $P_0$  is constructed using the held-out data (i.e.,  $1/2K$  of the data); then, computing

$$v_{n,k} = V(f_{n,k}, P_{n,k}).$$

Similarly, for each fold  $k$  in the second group, estimator  $f_{n,k,s}$  of  $f_{0,s}$  is constructed using all data besides the  $k$ th fold in the group (i.e.,  $(2K - 1)/(2K)$  of the data) and estimator  $P_{n,k}$  of  $P_0$  is constructed using the held-out data (i.e.,  $1/2K$  of the data); then, computing

$$v_{n,k,s} = V(f_{n,k,s}, P_{n,k}).$$

Finally,

$$\psi_{n,s} := K^{(-1)} \sum_{k=1}^K \{v_{n,k} - v_{n,k,s}\}.$$

See the paper by Williamson, Gilbert, Simon, and Carone for more details on the mathematics behind the `cv_vim` function, and the validity of the confidence intervals.

In the interest of transparency, we return most of the calculations within the `vim` object. This results in a list including:

`s` the column(s) to calculate variable importance for

`SL.library` the library of learners passed to SuperLearner

**full\_fit** the fitted values of the chosen method fit to the full data (a list, for train and test data)

**red\_fit** the fitted values of the chosen method fit to the reduced data (a list, for train and test data)

**est** the estimated variable importance

**naive** the naive estimator of variable importance

**EIF** the estimated efficient influence function

**EIF\_full** the estimated efficient influence function for the full regression

**EIF\_reduced** the estimated efficient influence function for the reduced regression

**se** the standard error for the estimated variable importance

**ci** the  $(1 - \alpha) \times 100\%$  confidence interval for the variable importance estimate

**test** a decision to either reject (TRUE) or not reject (FALSE) the null hypothesis, based on a conservative test

**p\_value** a p-value based on the same test as test

**full\_mod** the object returned by the estimation procedure for the full data regression (if applicable)

**red\_mod** the object returned by the estimation procedure for the reduced data regression (if applicable)

**alpha** the level, for confidence interval calculation

**sample\_splitting\_folds** the folds used for hypothesis testing

**cross\_fitting\_folds** the folds used for cross-fitting

**y** the outcome

**ipc\_weights** the weights

**mat** a tibble with the estimate, SE, CI, hypothesis testing decision, and p-value

## Value

An object of class `vim`. See Details for more information.

## See Also

[SuperLearner](#) for specific usage of the SuperLearner function and package.

## Examples

```
n <- 100
p <- 2
# generate the data
x <- data.frame(replicate(p, stats::runif(n, -5, 5)))

# apply the function to the x's
smooth <- (x[,1]/5)^2*(x[,1]+7)/5 + (x[,2]/3)^2

# generate Y ~ Normal (smooth, 1)
y <- as.matrix(smooth + stats::rnorm(n, 0, 1))

# set up a library for SuperLearner; note simple library for speed
library("SuperLearner")
```

```

learners <- c("SL.glm")

# -----
# using Super Learner (with a small number of folds, for illustration only)
# -----
set.seed(4747)
est <- cv_vim(Y = y, X = x, indx = 2, V = 2,
  type = "r_squared", run_regression = TRUE,
  SL.library = learners, cvControl = list(V = 2), alpha = 0.05)

# -----
# doing things by hand, and plugging them in
# (with a small number of folds, for illustration only)
# -----
# set up the folds
indx <- 2
V <- 2
Y <- matrix(y)
set.seed(4747)
# Note that the CV.SuperLearner should be run with an outer layer
# of 2*V folds (for V-fold cross-fitted importance)
full_cv_fit <- suppressWarnings(SuperLearner::CV.SuperLearner(
  Y = Y, X = x, SL.library = learners, cvControl = list(V = 2 * V),
  innerCvControl = list(list(V = V))
))
# use the same cross-fitting folds for reduced
reduced_cv_fit <- suppressWarnings(SuperLearner::CV.SuperLearner(
  Y = Y, X = x[, -indx, drop = FALSE], SL.library = learners,
  cvControl = SuperLearner::SuperLearner.CV.control(
    V = 2 * V, validRows = full_cv_fit$folds
  ),
  innerCvControl = list(list(V = V))
))
# extract the predictions on split portions of the data,
# for hypothesis testing
cross_fitting_folds <- get_cv_sl_folds(full_cv_fit$folds)
set.seed(1234)
sample_splitting_folds <- make_folds(unique(cross_fitting_folds), V = 2)
full_cv_preds <- extract_sampled_split_predictions(
  full_cv_fit, sample_splitting_folds = sample_splitting_folds, full = TRUE
)
reduced_cv_preds <- extract_sampled_split_predictions(
  reduced_cv_fit, sample_splitting_folds = sample_splitting_folds, full = FALSE
)
set.seed(5678)
# refit on the whole dataset
# (for estimating the efficient influence function)
full_fit <- SuperLearner::SuperLearner(
  Y = Y, X = x, SL.library = learners, cvControl = list(V = V)
)
fhat_full <- SuperLearner::predict.SuperLearner(full_fit, onlySL = TRUE)$pred
reduced_fit <- SuperLearner::SuperLearner(
  Y = Y, X = x[, -indx, drop = FALSE], SL.library = learners,

```

```

      cvControl = list(V = V)
    )
    fhat_red <- SuperLearner::predict.SuperLearner(reduced_fit, onlySL = TRUE)$pred
    est <- cv_vim(Y = y, cross_fitted_f1 = full_cv_preds,
      cross_fitted_f2 = reduced_cv_preds, f1 = fhat_ful,
      f2 = fhat_red, indx = 2, delta = 0, V = V, type = "r_squared",
      cross_fitting_folds = cross_fitting_folds,
      sample_splitting_folds = sample_splitting_folds,
      run_regression = FALSE, alpha = 0.05, na.rm = TRUE)

```

---

est\_predictiveness      *Estimate a nonparametric predictiveness functional*

---

### Description

Compute nonparametric estimates of the chosen measure of predictiveness.

### Usage

```

est_predictiveness(
  fitted_values,
  y,
  full_y = NULL,
  type = "r_squared",
  C = rep(1, length(y)),
  Z = NULL,
  ipc_weights = rep(1, length(C)),
  ipc_fit_type = "external",
  ipc_eif_preds = rep(1, length(C)),
  ipc_est_type = "aipw",
  scale = "identity",
  na.rm = FALSE,
  ...
)

```

### Arguments

fitted_values	fitted values from a regression function using the observed data.
y	the observed outcome.
full_y	the observed outcome (from the entire dataset, for cross-fitted estimates).
type	which parameter are you estimating (defaults to r_squared, for R-squared-based variable importance)?
C	the indicator of coarsening (1 denotes observed, 0 denotes unobserved).
Z	either NULL (if no coarsening) or a matrix-like object containing the fully observed data.

ipc_weights	weights for inverse probability of coarsening (e.g., inverse weights from a two-phase sample) weighted estimation. Assumed to be already inverted (i.e., ipc_weights = 1 / [estimated probability weights]).
ipc_fit_type	if "external", then use ipc_eif_preds; if "SL", fit a SuperLearner to determine the correction to the efficient influence function.
ipc_eif_preds	if ipc_fit_type = "external", the fitted values from a regression of the full-data EIF on the fully observed covariates/outcome; otherwise, not used.
ipc_est_type	IPC correction, either "ipw" (for classical inverse probability weighting) or "aipw" (for augmented inverse probability weighting; the default).
scale	if doing an IPC correction, then the scale that the correction should be computed on (e.g., "identity"; or "logit" to logit-transform, apply the correction, and back-transform).
na.rm	logical; should NA's be removed in computation? (defaults to FALSE)
...	other arguments to SuperLearner, if ipc_fit_type = "SL".

### Details

See the paper by Williamson, Gilbert, Simon, and Carone for more details on the mathematics behind this function and the definition of the parameter of interest.

### Value

A list, with: the estimated predictiveness; the estimated efficient influence function; and the predictions of the EIF based on inverse probability of censoring.

---

est\_predictiveness\_cv *Estimate a nonparametric predictiveness functional using cross-fitting*

---

### Description

Compute nonparametric estimates of the chosen measure of predictiveness.

### Usage

```
est_predictiveness_cv(
  fitted_values,
  y,
  full_y = NULL,
  folds,
  type = "r_squared",
  C = rep(1, length(y)),
  Z = NULL,
  folds_Z = folds,
  ipc_weights = rep(1, length(C)),
  ipc_fit_type = "external",
```

```

ipc_eif_preds = rep(1, length(C)),
ipc_est_type = "aipw",
scale = "identity",
na.rm = FALSE,
...
)

```

### Arguments

fitted_values	fitted values from a regression function using the observed data; a list of length $V$ , where each object is a set of predictions on the validation data.
y	the observed outcome.
full_y	the observed outcome (from the entire dataset, for cross-fitted estimates).
folds	the cross-validation folds for the observed data.
type	which parameter are you estimating (defaults to <code>r_squared</code> , for R-squared-based variable importance)?
C	the indicator of coarsening (1 denotes observed, 0 denotes unobserved).
Z	either NULL (if no coarsening) or a matrix-like object containing the fully observed data.
folds_Z	either the cross-validation folds for the observed data (no coarsening) or a vector of folds for the fully observed data $Z$ .
ipc_weights	weights for inverse probability of coarsening (e.g., inverse weights from a two-phase sample) weighted estimation. Assumed to be already inverted (i.e., <code>ipc_weights = 1 / [estimated probability weights]</code> ).
ipc_fit_type	if "external", then use <code>ipc_eif_preds</code> ; if "SL", fit a SuperLearner to determine the correction to the efficient influence function.
ipc_eif_preds	if <code>ipc_fit_type = "external"</code> , the fitted values from a regression of the full-data EIF on the fully observed covariates/outcome; otherwise, not used.
ipc_est_type	IPC correction, either "ipw" (for classical inverse probability weighting) or "aipw" (for augmented inverse probability weighting; the default).
scale	if doing an IPC correction, then the scale that the correction should be computed on (e.g., "identity"; or "logit" to logit-transform, apply the correction, and back-transform).
na.rm	logical; should NA's be removed in computation? (defaults to FALSE)
...	other arguments to SuperLearner, if <code>ipc_fit_type = "SL"</code> .

### Details

See the paper by Williamson, Gilbert, Simon, and Carone for more details on the mathematics behind this function and the definition of the parameter of interest. If sample-splitting is also requested (recommended, since in this case inferences will be valid even if the variable has zero true importance), then the prediction functions are trained as if  $2K$ -fold cross-validation were run, but are evaluated on only  $K$  sets (independent between the full and reduced nuisance regression).

**Value**

The estimated measure of predictiveness.

---

`extract_sampled_split_predictions`

*Extract sampled-split predictions from a CV.SuperLearner object*

---

**Description**

Use the cross-validated Super Learner and a set of specified sample-splitting folds to extract cross-fitted predictions on separate splits of the data. This is primarily for use in cases where you have already fit a CV.SuperLearner and want to use the fitted values to compute variable importance without having to re-fit. The number of folds used in the CV.SuperLearner must be even.

**Usage**

```
extract_sampled_split_predictions(  
  cvsl_obj = NULL,  
  sample_splitting = TRUE,  
  sample_splitting_folds = NULL,  
  full = TRUE  
)
```

**Arguments**

<code>cvsl_obj</code>	An object of class "CV.SuperLearner"
<code>sample_splitting</code>	logical; should we use sample-splitting or not? Defaults to TRUE.
<code>sample_splitting_folds</code>	A vector of folds to use for sample splitting
<code>full</code>	logical; is this the fit to all covariates (TRUE) or not (FALSE)?

**Value**

The predictions on validation data in each split-sample fold; a list of length two, each element of which is a list with the predictions on the split-sample cross-validation data.

**See Also**

[CV.SuperLearner](#) for usage of the CV.SuperLearner function.



---

format.vim	<i>Format a vim object</i>
------------	----------------------------

---

### Description

Nicely formats the output from a vim object for printing.

### Usage

```
## S3 method for class 'vim'
format(x, ...)
```

### Arguments

x	the vim object of interest.
...	other options, see the generic format function.

---

get_cv_sl_folds	<i>Get a numeric vector with cross-validation fold IDs from CV.SuperLearner</i>
-----------------	---

---

### Description

Get a numeric vector with cross-validation fold IDs from CV.SuperLearner

### Usage

```
get_cv_sl_folds(cv_sl_folds)
```

### Arguments

cv_sl_folds	The folds from a call to CV.SuperLearner; a list.
-------------	---

### Value

A numeric vector with the fold IDs.

---

get_full_type	<i>Obtain the type of VIM to estimate using partial matching</i>
---------------	--

---

**Description**

Obtain the type of VIM to estimate using partial matching

**Usage**

```
get_full_type(type)
```

**Arguments**

type	the partial string indicating the type of VIM
------	---

**Value**

the full string indicating the type of VIM

---

make_folds	<i>Create Folds for Cross-Fitting</i>
------------	---------------------------------------

---

**Description**

Create Folds for Cross-Fitting

**Usage**

```
make_folds(y, V = 2, stratified = FALSE, C = NULL, probs = rep(1/V, V))
```

**Arguments**

y	the outcome
V	the number of folds
stratified	should the folds be stratified based on the outcome?
C	a vector indicating whether or not the observation is fully observed; 1 denotes yes, 0 denotes no
probs	vector of proportions for each fold number

**Value**

a vector of folds

---

make_kfold	<i>Turn folds from 2K-fold cross-fitting into individual K-fold folds</i>
------------	---

---

**Description**

Turn folds from 2K-fold cross-fitting into individual K-fold folds

**Usage**

```
make_kfold(
  cross_fitting_folds,
  sample_splitting_folds = rep(1, length(unique(cross_fitting_folds))),
  C = rep(1, length(cross_fitting_folds))
)
```

**Arguments**

cross\_fitting\_folds  
the vector of cross-fitting folds

sample\_splitting\_folds  
the sample splitting folds

C  
vector of whether or not we measured the observation in phase 2

**Value**

the two sets of testing folds for K-fold cross-fitting

---

measure_accuracy	<i>Estimate the classification accuracy</i>
------------------	---

---

**Description**

Compute nonparametric estimate of classification accuracy.

**Usage**

```
measure_accuracy(
  fitted_values,
  y,
  full_y = NULL,
  C = rep(1, length(y)),
  Z = NULL,
  ipc_weights = rep(1, length(y)),
  ipc_fit_type = "external",
  ipc_eif_preds = rep(1, length(y)),
)
```

```

ipc_est_type = "aipw",
scale = "identity",
na.rm = FALSE,
...
)

```

### Arguments

fitted_values	fitted values from a regression function using the observed data (may be within a specified fold, for cross-fitted estimates).
y	the observed outcome (may be within a specified fold, for cross-fitted estimates).
full_y	the observed outcome (not used, defaults to NULL).
C	the indicator of coarsening (1 denotes observed, 0 denotes unobserved).
Z	either NULL (if no coarsening) or a matrix-like object containing the fully observed data.
ipc_weights	weights for inverse probability of coarsening (IPC) (e.g., inverse weights from a two-phase sample) weighted estimation. Assumed to be already inverted. (i.e., $ipc\_weights = 1 / [estimated\ probability\ weights]$ ).
ipc_fit_type	if "external", then use ipc_eif_preds; if "SL", fit a SuperLearner to determine the IPC correction to the efficient influence function.
ipc_eif_preds	if ipc_fit_type = "external", the fitted values from a regression of the full-data EIF on the fully observed covariates/outcome; otherwise, not used.
ipc_est_type	IPC correction, either "ipw" (for classical inverse probability weighting) or "aipw" (for augmented inverse probability weighting; the default).
scale	if doing an IPC correction, then the scale that the correction should be computed on (e.g., "identity"; or "logit" to logit-transform, apply the correction, and back-transform).
na.rm	logical; should NAs be removed in computation? (defaults to FALSE)
...	other arguments to SuperLearner, if ipc_fit_type = "SL".

### Value

A named list of: (1) the estimated classification accuracy of the fitted regression function; (2) the estimated influence function; and (3) the IPC EIF predictions.

---

measure\_anova

*Estimate ANOVA decomposition-based variable importance.*

---

### Description

Estimate ANOVA decomposition-based variable importance.

**Usage**

```
measure_anova(
  full,
  reduced,
  y,
  full_y = NULL,
  C = rep(1, length(y)),
  Z = NULL,
  ipc_weights = rep(1, length(y)),
  ipc_fit_type = "external",
  ipc_eif_preds = rep(1, length(y)),
  ipc_est_type = "aipw",
  scale = "identity",
  na.rm = FALSE,
  ...
)
```

**Arguments**

full	fitted values from a regression function of the observed outcome on the full set of covariates.
reduced	fitted values from a regression on the reduced set of observed covariates.
y	the observed outcome.
full_y	the observed outcome (not used, defaults to NULL).
C	the indicator of coarsening (1 denotes observed, 0 denotes unobserved).
Z	either NULL (if no coarsening) or a matrix-like object containing the fully observed data.
ipc_weights	weights for inverse probability of coarsening (e.g., inverse weights from a two-phase sample) weighted estimation. Assumed to be already inverted (i.e., $\text{ipc\_weights} = 1 / [\text{estimated probability weights}]$ ).
ipc_fit_type	if "external", then use ipc_eif_preds; if "SL", fit a SuperLearner to determine the correction to the efficient influence function.
ipc_eif_preds	if ipc_fit_type = "external", the fitted values from a regression of the full-data EIF on the fully observed covariates /outcome; otherwise, not used.
ipc_est_type	IPC correction, either "ipw" (for classical inverse probability weighting) or "aipw" (for augmented inverse probability weighting; the default).
scale	if doing an IPC correction, then the scale that the correction should be computed on (e.g., "identity"; or "logit" to logit-transform, apply the correction, and back-transform)
na.rm	logical; should NAs be removed in computation? (defaults to FALSE)
...	other arguments to SuperLearner, if ipc_fit_type = "SL".

**Value**

A named list of: (1) the estimated ANOVA (based on a one-step correction) of the fitted regression functions; (2) the estimated influence function; (3) the naive ANOVA estimate; and (4) the IPC EIF predictions.

---

measure_auc	<i>Estimate area under the receiver operating characteristic curve (AUC)</i>
-------------	--

---

**Description**

Compute nonparametric estimate of AUC.

**Usage**

```
measure_auc(
  fitted_values,
  y,
  full_y = NULL,
  C = rep(1, length(y)),
  Z = NULL,
  ipc_weights = rep(1, length(y)),
  ipc_fit_type = "external",
  ipc_eif_preds = rep(1, length(y)),
  ipc_est_type = "aipw",
  scale = "identity",
  na.rm = FALSE,
  ...
)
```

**Arguments**

fitted_values	fitted values from a regression function using the observed data (may be within a specified fold, for cross-fitted estimates).
y	the observed outcome (may be within a specified fold, for cross-fitted estimates).
full_y	the observed outcome (from the entire dataset, for cross-fitted estimates).
C	the indicator of coarsening (1 denotes observed, 0 denotes unobserved).
Z	either NULL (if no coarsening) or a matrix-like object containing the fully observed data.
ipc_weights	weights for inverse probability of coarsening (e.g., inverse weights from a two-phase sample) weighted estimation. Assumed to be already inverted (i.e., ipc_weights = 1 / [estimated probability weights]).
ipc_fit_type	if "external", then use ipc_eif_preds; if "SL", fit a SuperLearner to determine the correction to the efficient influence function.
ipc_eif_preds	if ipc_fit_type = "external", the fitted values from a regression of the full-data EIF on the fully observed covariates/outcome; otherwise, not used.

ipc_est_type	IPC correction, either "ipw" (for classical inverse probability weighting) or "aipw" (for augmented inverse probability weighting; the default).
scale	if doing an IPC correction, then the scale that the correction should be computed on (e.g., "identity"; or "logit" to logit-transform, apply the correction, and back-transform).
na.rm	logical; should NAs be removed in computation? (defaults to FALSE)
...	other arguments to SuperLearner, if ipc_fit_type = "SL".

**Value**

A named list of: (1) the estimated AUC of the fitted regression function; (2) the estimated influence function; and (3) the IPC EIF predictions.

---

measure\_cross\_entropy *Estimate the cross-entropy*

---

**Description**

Compute nonparametric estimate of cross-entropy.

**Usage**

```
measure_cross_entropy(
  fitted_values,
  y,
  full_y = NULL,
  C = rep(1, length(y)),
  Z = NULL,
  ipc_weights = rep(1, length(y)),
  ipc_fit_type = "external",
  ipc_eif_preds = rep(1, length(y)),
  ipc_est_type = "aipw",
  scale = "identity",
  na.rm = FALSE,
  ...
)
```

**Arguments**

fitted_values	fitted values from a regression function using the observed data.
y	the observed outcome.
full_y	the observed outcome (not used, defaults to NULL).
C	the indicator of coarsening (1 denotes observed, 0 denotes unobserved).
Z	either NULL (if no coarsening) or a matrix-like object containing the fully observed data.

ipc_weights	weights for inverse probability of coarsening (e.g., inverse weights from a two-phase sample) weighted estimation. Assumed to be already inverted (i.e., ipc_weights = 1 / [estimated probability weights]).
ipc_fit_type	if "external", then use ipc_eif_preds; if "SL", fit a SuperLearner to determine the correction to the efficient influence function.
ipc_eif_preds	if ipc_fit_type = "external", the fitted values from a regression of the full-data EIF on the fully observed covariates/outcome; otherwise, not used.
ipc_est_type	IPC correction, either "ipw" (for classical inverse probability weighting) or "aipw" (for augmented inverse probability weighting; the default).
scale	if doing an IPC correction, then the scale that the correction should be computed on (e.g., "identity"; or "logit" to logit-transform, apply the correction, and back-transform).
na.rm	logical; should NAs be removed in computation? (defaults to FALSE)
...	other arguments to SuperLearner, if ipc_fit_type = "SL".

**Value**

A named list of: (1) the estimated cross-entropy of the fitted regression function; (2) the estimated influence function; and (3) the IPC EIF predictions.

---

measure_deviance	<i>Estimate the deviance</i>
------------------	------------------------------

---

**Description**

Compute nonparametric estimate of deviance.

**Usage**

```
measure_deviance(
  fitted_values,
  y,
  full_y = NULL,
  C = rep(1, length(y)),
  Z = NULL,
  ipc_weights = rep(1, length(y)),
  ipc_fit_type = "external",
  ipc_eif_preds = rep(1, length(y)),
  ipc_est_type = "aipw",
  scale = "identity",
  na.rm = FALSE,
  ...
)
```



**Arguments**

fitted_values	fitted values from a regression function using the observed data.
y	the observed outcome.
full_y	the observed outcome (defaults to NULL; allows the full-data outcome to be used for empirical estimates that do not rely on covariates).
C	the indicator of coarsening (1 denotes observed, 0 denotes unobserved).
Z	either NULL (if no coarsening) or a matrix-like object containing the fully observed data.
ipc_weights	weights for inverse probability of coarsening (e.g., inverse weights from a two-phase sample) weighted estimation. Assumed to be already inverted (i.e., ipc_weights = 1 / [estimated probability weights]).
ipc_fit_type	if "external", then use ipc_eif_preds; if "SL", fit a SuperLearner to determine the correction to the efficient influence function.
ipc_eif_preds	if ipc_fit_type = "external", the fitted values from a regression of the full-data EIF on the fully observed covariates/outcome; otherwise, not used.
ipc_est_type	IPC correction, either "ipw" (for classical inverse probability weighting) or "aipw" (for augmented inverse probability weighting; the default).
scale	if doing an IPC correction, then the scale that the correction should be computed on (e.g., "identity"; or "logit" to logit-transform, apply the correction, and back-transform).
na.rm	logical; should NAs be removed in computation? (defaults to FALSE)
...	other arguments to SuperLearner, if ipc_fit_type = "SL".

**Value**

A named list of: (1) the estimated deviance of the fitted regression function; (2) the estimated influence function; and (3) the IPC EIF predictions.

---

measure_mse	<i>Estimate mean squared error</i>
-------------	------------------------------------

---

**Description**

Compute nonparametric estimate of mean squared error.

**Usage**

```
measure_mse(
  fitted_values,
  y,
  full_y = NULL,
  C = rep(1, length(y)),
  Z = NULL,
```

```

ipc_weights = rep(1, length(y)),
ipc_fit_type = "external",
ipc_eif_preds = rep(1, length(y)),
ipc_est_type = "aipw",
scale = "identity",
na.rm = FALSE,
...
)

```

### Arguments

<code>fitted_values</code>	fitted values from a regression function using the observed data (may be within a specified fold, for cross-fitted estimates).
<code>y</code>	the observed outcome (may be within a specified fold, for cross-fitted estimates).
<code>full_y</code>	the observed outcome (not used; defaults to NULL).
<code>C</code>	the indicator of coarsening (1 denotes observed, 0 denotes unobserved).
<code>Z</code>	either NULL (if no coarsening) or a matrix-like object containing the fully observed data.
<code>ipc_weights</code>	weights for inverse probability of coarsening (e.g., inverse weights from a two-phase sample) weighted estimation. Assumed to be already inverted (i.e., <code>ipc_weights = 1 / [estimated probability weights]</code> ).
<code>ipc_fit_type</code>	if "external", then use <code>ipc_eif_preds</code> ; if "SL", fit a SuperLearner to determine the correction to the efficient influence function.
<code>ipc_eif_preds</code>	if <code>ipc_fit_type = "external"</code> , the fitted values from a regression of the full-data EIF on the fully observed covariates/outcome; otherwise, not used.
<code>ipc_est_type</code>	IPC correction, either "ipw" (for classical inverse probability weighting) or "aipw" (for augmented inverse probability weighting; the default).
<code>scale</code>	if doing an IPC correction, then the scale that the correction should be computed on (e.g., "identity"; or "logit" to logit-transform, apply the correction, and back-transform).
<code>na.rm</code>	logical; should NAs be removed in computation? (defaults to FALSE)
<code>...</code>	other arguments to SuperLearner, if <code>ipc_fit_type = "SL"</code> .

### Value

A named list of: (1) the estimated mean squared error of the fitted regression function; (2) the estimated influence function; and (3) the IPC EIF predictions.

---

measure_r_squared	<i>Estimate R-squared</i>
-------------------	---------------------------

---

### Description

Estimate R-squared

### Usage

```
measure_r_squared(
  fitted_values,
  y,
  full_y = NULL,
  C = rep(1, length(y)),
  Z = NULL,
  ipc_weights = rep(1, length(y)),
  ipc_fit_type = "external",
  ipc_eif_preds = rep(1, length(y)),
  ipc_est_type = "aipw",
  scale = "identity",
  na.rm = FALSE,
  ...
)
```

### Arguments

fitted_values	fitted values from a regression function using the observed data.
y	the observed outcome.
full_y	the observed outcome (defaults to NULL; allows the full-data outcome to be used for empirical estimates that do not rely on covariates).
C	the indicator of coarsening (1 denotes observed, 0 denotes unobserved).
Z	either NULL (if no coarsening) or a matrix-like object containing the fully observed data.
ipc_weights	weights for inverse probability of coarsening (e.g., inverse weights from a two-phase sample) weighted estimation. Assumed to be already inverted (i.e., ipc_weights = 1 / [estimated probability weights]).
ipc_fit_type	if "external", then use ipc_eif_preds; if "SL", fit a SuperLearner to determine the correction to the efficient influence function.
ipc_eif_preds	if ipc_fit_type = "external", the fitted values from a regression of the full-data EIF on the fully observed covariates/outcome; otherwise, not used.
ipc_est_type	IPC correction, either "ipw" (for classical inverse probability weighting) or "aipw" (for augmented inverse probability weighting; the default).
scale	if doing an IPC correction, then the scale that the correction should be computed on (e.g., "identity"; or "logit" to logit-transform, apply the correction, and back-transform).

na.rm            logical; should NAs be removed in computation? (defaults to FALSE)  
 ...             other arguments to SuperLearner, if ipc\_fit\_type = "SL".

### Value

A named list of: (1) the estimated R-squared of the fitted regression function; (2) the estimated influence function; and (3) the IPC EIF predictions.

---

merge_vim	<i>Merge multiple vim objects into one</i>
-----------	--

---

### Description

Take the output from multiple different calls to `vimp_regression` and merge into a single `vim` object; mostly used for plotting results.

### Usage

```
merge_vim(...)
```

### Arguments

...             an arbitrary number of `vim` objects, separated by commas.

### Value

an object of class `vim` containing all of the output from the individual `vim` objects. This results in a list containing:

- `s` - a list of the column(s) to calculate variable importance for
- `SL.library` - a list of the libraries of learners passed to SuperLearner
- `full_fit` - a list of the fitted values of the chosen method fit to the full data
- `red_fit` - a list of the fitted values of the chosen method fit to the reduced data
- `est` - a vector with the corrected estimates
- `naive` - a vector with the naive estimates
- `EIF` - a list with the influence curve-based updates
- `se` - a vector with the standard errors
- `ci` - a matrix with the CIs
- `mat` - a tibble with the estimated variable importance, the standard errors, and the  $(1 - \alpha) \times 100\%$  confidence intervals
- `full_mod` - a list of the objects returned by the estimation procedure for the full data regression (if applicable)
- `red_mod` - a list of the objects returned by the estimation procedure for the reduced data regression (if applicable)
- `alpha` - a list of the levels, for confidence interval calculation

**Examples**

```

# generate the data
# generate X
p <- 2
n <- 100
x <- data.frame(replicate(p, stats::runif(n, -5, 5)))

# apply the function to the x's
smooth <- (x[,1]/5)^2*(x[,1]+7)/5 + (x[,2]/3)^2

# generate Y ~ Normal (smooth, 1)
y <- smooth + stats::rnorm(n, 0, 1)

# set up a library for SuperLearner; note simple library for speed
library("SuperLearner")
learners <- c("SL.glm", "SL.mean")

# using Super Learner (with a small number of folds, for illustration only)
est_2 <- vimp_regression(Y = y, X = x, indx = 2, V = 2,
  run_regression = TRUE, alpha = 0.05,
  SL.library = learners, cvControl = list(V = 2))

est_1 <- vimp_regression(Y = y, X = x, indx = 1, V = 2,
  run_regression = TRUE, alpha = 0.05,
  SL.library = learners, cvControl = list(V = 2))

ests <- merge_vim(est_1, est_2)

```

---

print.vim

---

*Print a vim object*


---

**Description**

Prints out the table of estimates, confidence intervals, and standard errors for a vim object.

**Usage**

```

## S3 method for class 'vim'
print(x, ...)

```

**Arguments**

x                    the vim object of interest.  
...                    other options, see the generic print function.

run\_sl

*Run a Super Learner for the provided subset of features***Description**

Run a Super Learner for the provided subset of features

**Usage**

```
run_sl(
  Y = NULL,
  X = NULL,
  V = 5,
  SL.library = "SL.glm",
  univariate_SL.library = "SL.glm",
  s = 1,
  cv_folds = NULL,
  sample_splitting = TRUE,
  ss_folds = NULL,
  verbose = FALSE,
  progress_bar = NULL,
  indx = 1,
  weights = rep(1, nrow(X)),
  cross_fitted_se = TRUE,
  ...
)
```

**Arguments**

Y	the outcome
X	the covariates
V	the number of folds
SL.library	the library of candidate learners
univariate_SL.library	the library of candidate learners for single-covariate regressions
s	the subset of interest
cv_folds	the CV folds
sample_splitting	logical; should we use sample-splitting for predictiveness estimation?
ss_folds	the sample-splitting folds; only used if sample_splitting = TRUE
verbose	should we print progress? defaults to FALSE
progress_bar	the progress bar to print to (only if verbose = TRUE)
indx	the index to pass to progress bar (only if verbose = TRUE)
weights	weights to pass to estimation procedure

`cross_fitted_se` if TRUE, uses a cross-fitted estimator of the standard error; otherwise, uses the entire dataset  
`...` other arguments to Super Learner

**Value**

a list of length  $V$ , with the results of predicting on the hold-out data for each  $v$  in 1 through  $V$

---

<code>sample_subsets</code>	<i>Create necessary objects for SPVIMs</i>
-----------------------------	--

---

**Description**

Creates the  $Z$  and  $W$  matrices and a list of sampled subsets,  $S$ , for SPVIM estimation.

**Usage**

```
sample_subsets(p, gamma, n)
```

**Arguments**

`p` the number of covariates  
`gamma` the fraction of the sample size to sample (e.g., `gamma = 1` means sample  $n$  subsets)  
`n` the sample size

**Value**

a list, with elements  $Z$  (the matrix encoding presence/absence of each feature in the uniquely sampled subsets),  $S$  (the list of unique sampled subsets),  $W$  (the matrix of weights), and  $z\_counts$  (the number of times each subset was sampled)

**Examples**

```

p <- 10
gamma <- 1
n <- 100
set.seed(100)
subset_lst <- sample_subsets(p, gamma, n)
  
```

---

scale_est	<i>Return an estimator on a different scale</i>
-----------	---

---

**Description**

Return an estimator on a different scale

**Usage**

```
scale_est(obs_est = NULL, grad = NULL, scale = "identity")
```

**Arguments**

obs_est	the observed VIM estimate
grad	the estimated efficient influence function
scale	the scale to compute on

**Details**

It may be of interest to return an estimate (or confidence interval) on a different scale than originally measured. For example, computing a confidence interval (CI) for a VIM value that lies in (0,1) on the logit scale ensures that the CI also lies in (0, 1).

**Value**

the scaled estimate

---

spvim_ics	<i>Influence function estimates for SPVIMs</i>
-----------	--

---

**Description**

Compute the influence functions for the contribution from sampling observations and subsets.

**Usage**

```
spvim_ics(Z, z_counts, W, v, psi, G, c_n, ics, measure)
```



**Arguments**

Z	the matrix of presence/absence of each feature (columns) in each sampled subset (rows)
z_counts	the number of times each unique subset was sampled
W	the matrix of weights
v	the estimated predictiveness measures
psi	the estimated SPVIM values
G	the constraint matrix
c_n	the constraint values
ics	a list of influence function values for each predictiveness measure
measure	the type of measure (e.g., "r_squared" or "auc")

**Details**

The processes for sampling observations and sampling subsets are independent. Thus, we can compute the influence function separately for each sampling process. For further details, see the paper by Williamson and Feng (2020).

**Value**

a named list of length 2; `contrib_v` is the contribution from estimating V, while `contrib_s` is the contribution from sampling subsets.

---

spvim_se	<i>Standard error estimate for SPVIM values</i>
----------	---

---

**Description**

Compute standard error estimates based on the estimated influence function for a SPVIM value of interest.

**Usage**

```
spvim_se(ics, idx = 1, gamma = 1, na_rm = FALSE)
```

**Arguments**

ics	the influence function estimates based on the contributions from sampling observations and sampling subsets: a list of length two resulting from a call to <code>spvim_ics</code> .
idx	the index of interest
gamma	the proportion of the sample size used when sampling subsets
na_rm	remove NAs?

**Details**

Since the processes for sampling observations and subsets are independent, the variance for a given SPVIM estimator is simply the sum of the variances based on sampling observations and on sampling subsets.

**Value**

The standard error estimate for the desired SPVIM value

**See Also**

[spvim\\_ics](#) for how the influence functions are estimated.

---

sp_vim	<i>Shapley Population Variable Importance Measure (SPVIM) Estimates and Inference</i>
--------	---

---

**Description**

Compute estimates and confidence intervals for the SPVIMs, using cross-fitting.

**Usage**

```
sp_vim(
  Y = NULL,
  X = NULL,
  V = 5,
  type = "r_squared",
  SL.library = c("SL.glmnet", "SL.xgboost", "SL.mean"),
  univariate_SL.library = NULL,
  gamma = 1,
  alpha = 0.05,
  delta = 0,
  na.rm = FALSE,
  stratified = FALSE,
  verbose = FALSE,
  sample_splitting = TRUE,
  C = rep(1, length(Y)),
  Z = NULL,
  ipc_weights = rep(1, length(Y)),
  ipc_est_type = "aipw",
  scale = "identity",
  scale_est = TRUE,
  cross_fitted_se = TRUE,
  ...
)
```

**Arguments**

Y	the outcome.
X	the covariates.
V	the number of folds for cross-fitting, defaults to 5. If <code>sample_splitting = TRUE</code> , then a special type of V-fold cross-fitting is done. See Details for a more detailed explanation.
type	the type of parameter (e.g., ANOVA-based is "anova").
SL.library	a character vector of learners to pass to SuperLearner, if f1 and f2 are Y and X, respectively. Defaults to <code>SL.glmnet</code> , <code>SL.xgboost</code> , and <code>SL.mean</code> .
univariate_SL.library	(optional) a character vector of learners to pass to SuperLearner for estimating univariate regression functions. Defaults to <code>SL.polymars</code>
gamma	the fraction of the sample size to use when sampling subsets (e.g., <code>gamma = 1</code> samples the same number of subsets as the sample size)
alpha	the level to compute the confidence interval at. Defaults to 0.05, corresponding to a 95% confidence interval.
delta	the value of the $\delta$ -null (i.e., testing if importance $< \delta$ ); defaults to 0.
na.rm	should we remove NA's in the outcome and fitted values in computation? (defaults to FALSE)
stratified	if <code>run_regression = TRUE</code> , then should the generated folds be stratified based on the outcome (helps to ensure class balance across cross-fitting folds)
verbose	should <code>sp_vim</code> and SuperLearner print out progress? (defaults to FALSE)
sample_splitting	should we use sample-splitting to estimate the full and reduced predictiveness? Defaults to TRUE, since inferences made using <code>sample_splitting = FALSE</code> will be invalid for variable with truly zero importance.
C	the indicator of coarsening (1 denotes observed, 0 denotes unobserved).
Z	either (i) NULL (the default, in which case the argument C above must be all ones), or (ii) a character vector specifying the variable(s) among Y and X that are thought to play a role in the coarsening mechanism.
ipc_weights	weights for the computed influence curve (i.e., inverse probability weights for coarsened-at-random settings). Assumed to be already inverted (i.e., <code>ipc_weights = 1 / [estimated probability weights]</code> ).
ipc_est_type	the type of procedure used for coarsened-at-random settings; options are "ipw" (for inverse probability weighting) or "aipw" (for augmented inverse probability weighting). Only used if C is not all equal to 1.
scale	should CIs be computed on original ("identity", default) or logit ("logit") scale?
scale_est	should the point estimate be scaled to be greater than 0? Defaults to TRUE.
cross_fitted_se	should we use cross-fitting to estimate the standard errors (TRUE, the default) or not (FALSE)?
...	other arguments to the estimation tool, see "See also".

## Details

We define the SPVIM as the weighted average of the population difference in predictiveness over all subsets of features not containing feature  $j$ .

This is equivalent to finding the solution to a population weighted least squares problem. This key fact allows us to estimate the SPVIM using weighted least squares, where we first sample subsets from the power set of all possible features using the Shapley sampling distribution; then use cross-fitting to obtain estimators of the predictiveness of each sampled subset; and finally, solve the least squares problem given in Williamson and Feng (2020).

See the paper by Williamson and Feng (2020) for more details on the mathematics behind this function, and the validity of the confidence intervals.

In the interest of transparency, we return most of the calculations within the `vim` object. This results in a list containing:

**SL.library** the library of learners passed to SuperLearner  
**v** the estimated predictiveness measure for each sampled subset  
**fit\_lst** the fitted values on the entire dataset from the chosen method for each sampled subset  
**preds\_lst** the cross-fitted predicted values from the chosen method for each sampled subset  
**est** the estimated SPVIM value for each feature  
**ics** the influence functions for each sampled subset  
**var\_v\_contribs** the contributions to the variance from estimating predictiveness  
**var\_s\_contribs** the contributions to the variance from sampling subsets  
**ic\_lst** a list of the SPVIM influence function contributions  
**se** the standard errors for the estimated variable importance  
**ci** the  $(1 - \alpha) \times 100\%$  confidence intervals based on the variable importance estimates  
**p\_value** p-values for the null hypothesis test of zero importance for each variable  
**test\_statistic** the test statistic for each null hypothesis test of zero importance  
**test** a hypothesis testing decision for each null hypothesis test (for each variable having zero importance)  
**gamma** the fraction of the sample size used when sampling subsets  
**alpha** the level, for confidence interval calculation  
**delta** the delta value used for hypothesis testing  
**y** the outcome  
**ipc\_weights** the weights  
**scale** the scale on which CIs were computed  
**mat** - a tibble with the estimates, SEs, CIs, hypothesis testing decisions, and p-values

## Value

An object of class `vim`. See Details for more information.

## See Also

[SuperLearner](#) for specific usage of the SuperLearner function and package.

**Examples**

```

n <- 100
p <- 2
# generate the data
x <- data.frame(replicate(p, stats::runif(n, -5, 5)))

# apply the function to the x's
smooth <- (x[,1]/5)^2*(x[,1]+7)/5 + (x[,2]/3)^2

# generate Y ~ Normal (smooth, 1)
y <- as.matrix(smooth + stats::rnorm(n, 0, 1))

# set up a library for SuperLearner; note simple library for speed
library("SuperLearner")
learners <- c("SL.glm")

# -----
# using Super Learner (with a small number of CV folds,
# for illustration only)
# -----
set.seed(4747)
est <- sp_vim(Y = y, X = x, V = 2, type = "r_squared",
             SL.library = learners, alpha = 0.05)

```

**Description**

Compute estimates of and confidence intervals for nonparametric intrinsic variable importance based on the population-level contrast between the oracle predictiveness using the feature(s) of interest versus not.

**Usage**

```

vim(
  Y = NULL,
  X = NULL,
  f1 = NULL,
  f2 = NULL,
  indx = 1,
  type = "r_squared",
  run_regression = TRUE,
  SL.library = c("SL.glmnet", "SL.xgboost", "SL.mean"),
  alpha = 0.05,
  delta = 0,
  scale = "identity",

```

```

na.rm = FALSE,
sample_splitting = TRUE,
sample_splitting_folds = NULL,
stratified = FALSE,
C = rep(1, length(Y)),
Z = NULL,
ipc_weights = rep(1, length(Y)),
ipc_est_type = "aipw",
scale_est = TRUE,
bootstrap = FALSE,
b = 1000,
...
)

```

### Arguments

Y	the outcome.
X	the covariates.
f1	the fitted values from a flexible estimation technique regressing Y on X.
f2	the fitted values from a flexible estimation technique regressing either (a) f1 or (b) Y on X withholding the columns in <code>indx</code> .
indx	the indices of the covariate(s) to calculate variable importance for; defaults to 1.
type	the type of importance to compute; defaults to <code>r_squared</code> , but other supported options are <code>auc</code> , <code>accuracy</code> , <code>deviance</code> , and <code>anova</code> .
run_regression	if outcome Y and covariates X are passed to <code>vimp_accuracy</code> , and <code>run_regression</code> is TRUE, then Super Learner will be used; otherwise, variable importance will be computed using the inputted fitted values.
SL.library	a character vector of learners to pass to SuperLearner, if f1 and f2 are Y and X, respectively. Defaults to <code>SL.glmnet</code> , <code>SL.xgboost</code> , and <code>SL.mean</code> .
alpha	the level to compute the confidence interval at. Defaults to 0.05, corresponding to a 95% confidence interval.
delta	the value of the $\delta$ -null (i.e., testing if importance < $\delta$ ); defaults to 0.
scale	should CIs be computed on original ("identity") or logit ("logit") scale?
na.rm	should we remove NAs in the outcome and fitted values in computation? (defaults to FALSE)
sample_splitting	should we use sample-splitting to estimate the full and reduced predictiveness? Defaults to TRUE, since inferences made using <code>sample_splitting = FALSE</code> will be invalid for variable with truly zero importance.
sample_splitting_folds	the folds used for sample-splitting; these identify the observations that should be used to evaluate predictiveness based on the full and reduced sets of covariates, respectively. Only used if <code>run_regression = FALSE</code> .
stratified	if <code>run_regression = TRUE</code> , then should the generated folds be stratified based on the outcome (helps to ensure class balance across cross-validation folds)

<code>C</code>	the indicator of coarsening (1 denotes observed, 0 denotes unobserved).
<code>Z</code>	either (i) NULL (the default, in which case the argument <code>C</code> above must be all ones), or (ii) a character vector specifying the variable(s) among <code>Y</code> and <code>X</code> that are thought to play a role in the coarsening mechanism.
<code>ipc_weights</code>	weights for the computed influence curve (i.e., inverse probability weights for coarsened-at-random settings). Assumed to be already inverted (i.e., <code>ipc_weights = 1 / [estimated probability weights]</code> ).
<code>ipc_est_type</code>	the type of procedure used for coarsened-at-random settings; options are "ipw" (for inverse probability weighting) or "aipw" (for augmented inverse probability weighting). Only used if <code>C</code> is not all equal to 1.
<code>scale_est</code>	should the point estimate be scaled to be greater than 0? Defaults to TRUE.
<code>bootstrap</code>	should bootstrap-based standard error estimates be computed? Defaults to FALSE (and currently may only be used if <code>sample_splitting = FALSE</code> ).
<code>b</code>	the number of bootstrap replicates (only used if <code>bootstrap = TRUE</code> and <code>sample_splitting = FALSE</code> ).
<code>...</code>	other arguments to the estimation tool, see "See also".

## Details

We define the population variable importance measure (VIM) for the group of features (or single feature)  $s$  with respect to the predictiveness measure  $V$  by

$$\psi_{0,s} := V(f_0, P_0) - V(f_{0,s}, P_0),$$

where  $f_0$  is the population predictiveness maximizing function,  $f_{0,s}$  is the population predictiveness maximizing function that is only allowed to access the features with index not in  $s$ , and  $P_0$  is the true data-generating distribution. VIM estimates are obtained by obtaining estimators  $f_n$  and  $f_{n,s}$  of  $f_0$  and  $f_{0,s}$ , respectively; obtaining an estimator  $P_n$  of  $P_0$ ; and finally, setting  $\psi_{n,s} := V(f_n, P_n) - V(f_{n,s}, P_n)$ .

In the interest of transparency, we return most of the calculations within the `vim` object. This results in a list including:

- `s` the column(s) to calculate variable importance for
- `SL.library` the library of learners passed to SuperLearner
- `type` the type of risk-based variable importance measured
- `full_fit` the fitted values of the chosen method fit to the full data
- `red_fit` the fitted values of the chosen method fit to the reduced data
- `est` the estimated variable importance
- `naive` the naive estimator of variable importance (only used if `type = "anova"`)
- `EIF` the estimated efficient influence function
- `EIF_full` the estimated efficient influence function for the full regression
- `EIF_reduced` the estimated efficient influence function for the reduced regression
- `se` the standard error for the estimated variable importance
- `ci` the  $(1 - \alpha) \times 100\%$  confidence interval for the variable importance estimate

**test** a decision to either reject (TRUE) or not reject (FALSE) the null hypothesis, based on a conservative test

**p\_value** a p-value based on the same test as test

**full\_mod** the object returned by the estimation procedure for the full data regression (if applicable)

**red\_mod** the object returned by the estimation procedure for the reduced data regression (if applicable)

**alpha** the level, for confidence interval calculation

**sample\_splitting\_folds** the folds used for sample-splitting (used for hypothesis testing)

**y** the outcome

**ipc\_weights** the weights

**mat** a tibble with the estimate, SE, CI, hypothesis testing decision, and p-value

### Value

An object of classes `vim` and the type of risk-based measure. See Details for more information.

### See Also

[SuperLearner](#) for specific usage of the SuperLearner function and package.

### Examples

```
# generate the data
# generate X
p <- 2
n <- 100
x <- data.frame(replicate(p, stats::runif(n, -1, 1)))

# apply the function to the x's
f <- function(x) 0.5 + 0.3*x[1] + 0.2*x[2]
smooth <- apply(x, 1, function(z) f(z))

# generate Y ~ Bernoulli (smooth)
y <- matrix(rbinom(n, size = 1, prob = smooth))

# set up a library for SuperLearner; note simple library for speed
library("SuperLearner")
learners <- c("SL.glm")

# using Y and X; use class-balanced folds
est_1 <- vim(y, x, indx = 2, type = "accuracy",
            alpha = 0.05, run_regression = TRUE,
            SL.library = learners, cvControl = list(V = 2),
            stratified = TRUE)

# using pre-computed fitted values
set.seed(4747)
V <- 2
full_fit <- SuperLearner::SuperLearner(Y = y, X = x,
```



```

                                SL.library = learners,
                                cvControl = list(V = V))
full_fitted <- SuperLearner::predict.SuperLearner(full_fit)$pred
# fit the data with only X1
reduced_fit <- SuperLearner::SuperLearner(Y = full_fitted,
                                           X = x[, -2, drop = FALSE],
                                           SL.library = learners,
                                           cvControl = list(V = V))
reduced_fitted <- SuperLearner::predict.SuperLearner(reduced_fit)$pred

est_2 <- vimp(Y = y, f1 = full_fitted, f2 = reduced_fitted,
             indx = 2, run_regression = FALSE, alpha = 0.05,
             stratified = TRUE, type = "accuracy",
             sample_splitting_folds = est_1$sample_splitting_folds)

```

---

vimp	<i>vimp: Perform Inference on Algorithm-Agnostic Intrinsic Variable Importance</i>
------	--

---

## Description

A unified framework for valid statistical inference on algorithm-agnostic measures of intrinsic variable importance. You provide the data, a method for estimating the conditional mean of the outcome given the covariates, choose a variable importance measure, and specify variable(s) of interest; 'vimp' takes care of the rest.

## Author(s)

**Maintainer:** Brian Williamson <https://bdwilliamson.github.io/> **Contributor:** Jean Feng <http://www.jeanfeng.com>

Methodology authors:

- Brian D. Williamson
- Jean Feng
- Peter B. Gilbert
- Noah R. Simon
- Marco Carone

## See Also

Manuscripts:

- <https://onlinelibrary.wiley.com/doi/epdf/10.1111/biom.13392> (R-squared-based variable importance)
- <https://onlinelibrary.wiley.com/doi/epdf/10.1111/biom.13389> (Rejoinder to discussion on R-squared-based variable importance article)

- <http://proceedings.mlr.press/v119/williamson20a.html> (general Shapley-based variable importance)

Preprints:

- <https://arxiv.org/abs/2004.03683> (general variable importance)

Other useful links:

- <https://bdwilliamson.github.io/vimp/>
- <https://github.com/bdwilliamson/vimp>
- Report bugs at <https://github.com/bdwilliamson/vimp/issues>

## Imports

The packages that we import either make the internal code nice (dplyr, magrittr, tibble, rlang, MASS, data.table), are directly relevant to estimating the conditional mean (SuperLearner) or predictiveness measures (ROCR), or are necessary for hypothesis testing (stats) or confidence intervals (boot, only for bootstrap intervals).

We suggest several other packages: xgboost, ranger, gam, glmnet, polyspline, and quadprog allow a flexible library of candidate learners in the Super Learner; ggplot2 and cowplot help with plotting variable importance estimates; testthat and covr help with unit tests; and knitr, rmarkdown, and RCurl help with the vignettes and examples.

---

vimp_accuracy	<i>Nonparametric Intrinsic Variable Importance Estimates: Classification accuracy</i>
---------------	---

---

## Description

Compute estimates of and confidence intervals for nonparametric difference in classification accuracy-based intrinsic variable importance. This is a wrapper function for `cv_vim`, with `type = "accuracy"`.

## Usage

```
vimp_accuracy(
  Y = NULL,
  X = NULL,
  cross_fitted_f1 = NULL,
  cross_fitted_f2 = NULL,
  f1 = NULL,
  f2 = NULL,
  indx = 1,
  V = 10,
  run_regression = TRUE,
  SL.library = c("SL.glmnet", "SL.xgboost", "SL.mean"),
  alpha = 0.05,
  delta = 0,
```

```

na.rm = FALSE,
cross_fitting_folds = NULL,
sample_splitting_folds = NULL,
stratified = TRUE,
C = rep(1, length(Y)),
Z = NULL,
ipc_weights = rep(1, length(Y)),
scale = "identity",
ipc_est_type = "aipw",
scale_est = TRUE,
cross_fitted_se = TRUE,
...
)

```

### Arguments

Y	the outcome.
X	the covariates.
cross_fitted_f1	the predicted values on validation data from a flexible estimation technique regressing Y on X in the training data; a list of length V, where each object is a set of predictions on the validation data. If sample-splitting is requested, then these must be estimated specially; see Details.
cross_fitted_f2	the predicted values on validation data from a flexible estimation technique regressing either (a) the fitted values in cross_fitted_f1, or (b) Y, on X withholding the columns in indx; a list of length V, where each object is a set of predictions on the validation data. If sample-splitting is requested, then these must be estimated specially; see Details.
f1	the fitted values from a flexible estimation technique regressing Y on X. Estimated using the whole dataset. If cross_fitted_se = TRUE, then this argument is not used.
f2	the fitted values from a flexible estimation technique regressing either (a) f1 or (b) Y on X withholding the columns in indx. Estimated using the whole dataset. If cross_fitted_se = TRUE, then this argument is not used.
indx	the indices of the covariate(s) to calculate variable importance for; defaults to 1.
V	the number of folds for cross-fitting, defaults to 5. If sample_splitting = TRUE, then a special type of V-fold cross-fitting is done. See Details for a more detailed explanation.
run_regression	if outcome Y and covariates X are passed to cv_vim, and run_regression is TRUE, then Super Learner will be used; otherwise, variable importance will be computed using the inputted fitted values.
SL.library	a character vector of learners to pass to SuperLearner, if f1 and f2 are Y and X, respectively. Defaults to SL.glmnet, SL.xgboost, and SL.mean.
alpha	the level to compute the confidence interval at. Defaults to 0.05, corresponding to a 95% confidence interval.

<code>delta</code>	the value of the $\delta$ -null (i.e., testing if importance $< \delta$ ); defaults to 0.
<code>na.rm</code>	should we remove NA's in the outcome and fitted values in computation? (defaults to FALSE)
<code>cross_fitting_folds</code>	the folds for cross-fitting. Only used if <code>run_regression = FALSE</code> .
<code>sample_splitting_folds</code>	the folds to use for sample-splitting; if entered, these should result in balance within the cross-fitting folds. Only used if <code>run_regression = FALSE</code> and <code>sample_splitting = TRUE</code> . A vector of length $2 * V$ .
<code>stratified</code>	if <code>run_regression = TRUE</code> , then should the generated folds be stratified based on the outcome (helps to ensure class balance across cross-fitting folds)
<code>C</code>	the indicator of coarsening (1 denotes observed, 0 denotes unobserved).
<code>Z</code>	either (i) NULL (the default, in which case the argument <code>C</code> above must be all ones), or (ii) a character vector specifying the variable(s) among <code>Y</code> and <code>X</code> that are thought to play a role in the coarsening mechanism.
<code>ipc_weights</code>	weights for the computed influence curve (i.e., inverse probability weights for coarsened-at-random settings). Assumed to be already inverted (i.e., <code>ipc_weights = 1 / [estimated probability weights]</code> ).
<code>scale</code>	should CIs be computed on original ("identity", default) or logit ("logit") scale?
<code>ipc_est_type</code>	the type of procedure used for coarsened-at-random settings; options are "ipw" (for inverse probability weighting) or "aipw" (for augmented inverse probability weighting). Only used if <code>C</code> is not all equal to 1.
<code>scale_est</code>	should the point estimate be scaled to be greater than 0? Defaults to TRUE.
<code>cross_fitted_se</code>	should we use cross-fitting to estimate the standard errors (TRUE, the default) or not (FALSE)?
<code>...</code>	other arguments to the estimation tool, see "See also".

## Details

We define the population variable importance measure (VIM) for the group of features (or single feature)  $s$  with respect to the predictiveness measure  $V$  by

$$\psi_{0,s} := V(f_0, P_0) - V(f_{0,s}, P_0),$$

where  $f_0$  is the population predictiveness maximizing function,  $f_{0,s}$  is the population predictiveness maximizing function that is only allowed to access the features with index not in  $s$ , and  $P_0$  is the true data-generating distribution.

Cross-fitted VIM estimates are computed differently if sample-splitting is requested versus if it is not. We recommend using sample-splitting in most cases, since only in this case will inferences be valid if the variable(s) of interest have truly zero population importance. The purpose of cross-fitting is to estimate  $f_0$  and  $f_{0,s}$  on independent data from estimating  $P_0$ ; this can result in improved performance, especially when using flexible learning algorithms. The purpose of sample-splitting is to estimate  $f_0$  and  $f_{0,s}$  on independent data; this allows valid inference under the null hypothesis of zero importance.

Without sample-splitting, cross-fitted VIM estimates are obtained by first splitting the data into  $K$  folds; then using each fold in turn as a hold-out set, constructing estimators  $f_{n,k}$  and  $f_{n,k,s}$  of  $f_0$  and  $f_{0,s}$ , respectively on the training data and estimator  $P_{n,k}$  of  $P_0$  using the test data; and finally, computing

$$\psi_{n,s} := K^{(-1)} \sum_{k=1}^K \{V(f_{n,k}, P_{n,k}) - V(f_{n,k,s}, P_{n,k})\}.$$

With sample-splitting, cross-fitted VIM estimates are obtained by first splitting the data into  $2K$  folds. These folds are further divided into 2 groups of folds. Then, for each fold  $k$  in the first group, estimator  $f_{n,k}$  of  $f_0$  is constructed using all data besides the  $k$ th fold in the group (i.e.,  $(2K - 1)/(2K)$  of the data) and estimator  $P_{n,k}$  of  $P_0$  is constructed using the held-out data (i.e.,  $1/2K$  of the data); then, computing

$$v_{n,k} = V(f_{n,k}, P_{n,k}).$$

Similarly, for each fold  $k$  in the second group, estimator  $f_{n,k,s}$  of  $f_{0,s}$  is constructed using all data besides the  $k$ th fold in the group (i.e.,  $(2K - 1)/(2K)$  of the data) and estimator  $P_{n,k}$  of  $P_0$  is constructed using the held-out data (i.e.,  $1/2K$  of the data); then, computing

$$v_{n,k,s} = V(f_{n,k,s}, P_{n,k}).$$

Finally,

$$\psi_{n,s} := K^{(-1)} \sum_{k=1}^K \{v_{n,k} - v_{n,k,s}\}.$$

See the paper by Williamson, Gilbert, Simon, and Carone for more details on the mathematics behind the `cv_vim` function, and the validity of the confidence intervals.

In the interest of transparency, we return most of the calculations within the `vim` object. This results in a list including:

**s** the column(s) to calculate variable importance for

**SL.library** the library of learners passed to SuperLearner

**full\_fit** the fitted values of the chosen method fit to the full data (a list, for train and test data)

**red\_fit** the fitted values of the chosen method fit to the reduced data (a list, for train and test data)

**est** the estimated variable importance

**naive** the naive estimator of variable importance

**eif** the estimated efficient influence function

**eif\_full** the estimated efficient influence function for the full regression

**eif\_reduced** the estimated efficient influence function for the reduced regression

**se** the standard error for the estimated variable importance

**ci** the  $(1 - \alpha) \times 100\%$  confidence interval for the variable importance estimate

**test** a decision to either reject (TRUE) or not reject (FALSE) the null hypothesis, based on a conservative test

**p\_value** a p-value based on the same test as `test`

**full\_mod** the object returned by the estimation procedure for the full data regression (if applicable)

**red\_mod** the object returned by the estimation procedure for the reduced data regression (if applicable)

**alpha** the level, for confidence interval calculation

**sample\_splitting\_folds** the folds used for hypothesis testing

**cross\_fitting\_folds** the folds used for cross-fitting

**y** the outcome

**ipc\_weights** the weights

**mat** a tibble with the estimate, SE, CI, hypothesis testing decision, and p-value

## Value

An object of classes `vim` and `vim_accuracy`. See Details for more information.

## See Also

[SuperLearner](#) for specific usage of the SuperLearner function and package.

## Examples

```
# generate the data
# generate X
p <- 2
n <- 100
x <- data.frame(replicate(p, stats::runif(n, -1, 1)))

# apply the function to the x's
f <- function(x) 0.5 + 0.3*x[1] + 0.2*x[2]
smooth <- apply(x, 1, function(z) f(z))

# generate Y ~ Normal (smooth, 1)
y <- matrix(rbinom(n, size = 1, prob = smooth))

# set up a library for SuperLearner; note simple library for speed
library("SuperLearner")
learners <- c("SL.glm", "SL.mean")

# estimate (with a small number of folds, for illustration only)
est <- vimp_accuracy(y, x, indx = 2,
  alpha = 0.05, run_regression = TRUE,
  SL.library = learners, V = 2, cvControl = list(V = 2))
```

**Description**

Compute estimates of and confidence intervals for nonparametric ANOVA-based intrinsic variable importance. This is a wrapper function for `cv_vim`, with `type = "anova"`. This type has limited functionality compared to other types; in particular, null hypothesis tests are not possible using `type = "anova"`. If you want to do null hypothesis testing on an equivalent population parameter, use `vimp_rsquared` instead.

**Usage**

```
vimp_anova(
  Y = NULL,
  X = NULL,
  cross_fitted_f1 = NULL,
  cross_fitted_f2 = NULL,
  indx = 1,
  V = 10,
  run_regression = TRUE,
  SL.library = c("SL.glmnet", "SL.xgboost", "SL.mean"),
  alpha = 0.05,
  delta = 0,
  na.rm = FALSE,
  cross_fitting_folds = NULL,
  stratified = FALSE,
  C = rep(1, length(Y)),
  Z = NULL,
  ipc_weights = rep(1, length(Y)),
  scale = "identity",
  ipc_est_type = "aipw",
  scale_est = TRUE,
  cross_fitted_se = TRUE,
  ...
)
```

**Arguments**

<code>Y</code>	the outcome.
<code>X</code>	the covariates.
<code>cross_fitted_f1</code>	the predicted values on validation data from a flexible estimation technique regressing <code>Y</code> on <code>X</code> in the training data; a list of length <code>V</code> , where each object is a set of predictions on the validation data. If sample-splitting is requested, then these must be estimated specially; see <a href="#">Details</a> .

<code>cross_fitted_f2</code>	the predicted values on validation data from a flexible estimation technique regressing either (a) the fitted values in <code>cross_fitted_f1</code> , or (b) Y, on X withholding the columns in <code>indx</code> ; a list of length V, where each object is a set of predictions on the validation data. If sample-splitting is requested, then these must be estimated specially; see Details.
<code>indx</code>	the indices of the covariate(s) to calculate variable importance for; defaults to 1.
<code>V</code>	the number of folds for cross-fitting, defaults to 5. If <code>sample_splitting = TRUE</code> , then a special type of V-fold cross-fitting is done. See Details for a more detailed explanation.
<code>run_regression</code>	if outcome Y and covariates X are passed to <code>cv_vim</code> , and <code>run_regression</code> is TRUE, then Super Learner will be used; otherwise, variable importance will be computed using the inputted fitted values.
<code>SL.library</code>	a character vector of learners to pass to SuperLearner, if <code>f1</code> and <code>f2</code> are Y and X, respectively. Defaults to <code>SL.glmnet</code> , <code>SL.xgboost</code> , and <code>SL.mean</code> .
<code>alpha</code>	the level to compute the confidence interval at. Defaults to 0.05, corresponding to a 95% confidence interval.
<code>delta</code>	the value of the $\delta$ -null (i.e., testing if importance $< \delta$ ); defaults to 0.
<code>na.rm</code>	should we remove NA's in the outcome and fitted values in computation? (defaults to FALSE)
<code>cross_fitting_folds</code>	the folds for cross-fitting. Only used if <code>run_regression = FALSE</code> .
<code>stratified</code>	if <code>run_regression = TRUE</code> , then should the generated folds be stratified based on the outcome (helps to ensure class balance across cross-fitting folds)
<code>C</code>	the indicator of coarsening (1 denotes observed, 0 denotes unobserved).
<code>Z</code>	either (i) NULL (the default, in which case the argument C above must be all ones), or (ii) a character vector specifying the variable(s) among Y and X that are thought to play a role in the coarsening mechanism.
<code>ipc_weights</code>	weights for the computed influence curve (i.e., inverse probability weights for coarsened-at-random settings). Assumed to be already inverted (i.e., <code>ipc_weights = 1 / [estimated probability weights]</code> ).
<code>scale</code>	should CIs be computed on original ("identity", default) or logit ("logit") scale?
<code>ipc_est_type</code>	the type of procedure used for coarsened-at-random settings; options are "ipw" (for inverse probability weighting) or "aipw" (for augmented inverse probability weighting). Only used if C is not all equal to 1.
<code>scale_est</code>	should the point estimate be scaled to be greater than 0? Defaults to TRUE.
<code>cross_fitted_se</code>	should we use cross-fitting to estimate the standard errors (TRUE, the default) or not (FALSE)?
<code>...</code>	other arguments to the estimation tool, see "See also".



## Details

We define the population ANOVA parameter for the group of features (or single feature)  $s$  by

$$\psi_{0,s} := E_0\{f_0(X) - f_{0,s}(X)\}^2 / \text{var}_0(Y),$$

where  $f_0$  is the population conditional mean using all features,  $f_{0,s}$  is the population conditional mean using the features with index not in  $s$ , and  $E_0$  and  $\text{var}_0$  denote expectation and variance under the true data-generating distribution, respectively.

Cross-fitted ANOVA estimates are computed by first splitting the data into  $K$  folds; then using each fold in turn as a hold-out set, constructing estimators  $f_{n,k}$  and  $f_{n,k,s}$  of  $f_0$  and  $f_{0,s}$ , respectively on the training data and estimator  $E_{n,k}$  of  $E_0$  using the test data; and finally, computing

$$\psi_{n,s} := K^{(-1)} \sum_{k=1}^K E_{n,k}\{f_{n,k}(X) - f_{n,k,s}(X)\}^2 / \text{var}_n(Y),$$

where  $\text{var}_n$  is the empirical variance. See the paper by Williamson, Gilbert, Simon, and Carone for more details on the mathematics behind this function.

## Value

An object of classes `vim` and `vimp_anova`. See Details for more information.

## See Also

[SuperLearner](#) for specific usage of the `SuperLearner` function and package.

## Examples

```
# generate the data
# generate X
p <- 2
n <- 100
x <- data.frame(replicate(p, stats::runif(n, -5, 5)))

# apply the function to the x's
smooth <- (x[,1]/5)^2*(x[,1]+7)/5 + (x[,2]/3)^2

# generate Y ~ Normal (smooth, 1)
y <- smooth + stats::rnorm(n, 0, 1)

# set up a library for SuperLearner; note simple library for speed
library("SuperLearner")
learners <- c("SL.glm", "SL.mean")

# estimate (with a small number of folds, for illustration only)
est <- vimp_anova(y, x, indx = 2,
  alpha = 0.05, run_regression = TRUE,
  SL.library = learners, V = 2, cvControl = list(V = 2))
```

vimp\_auc

*Nonparametric Intrinsic Variable Importance Estimates: AUC***Description**

Compute estimates of and confidence intervals for nonparametric difference in \$AUC\$-based intrinsic variable importance. This is a wrapper function for `cv_vim`, with `type = "auc"`.

**Usage**

```
vimp_auc(
  Y = NULL,
  X = NULL,
  cross_fitted_f1 = NULL,
  cross_fitted_f2 = NULL,
  f1 = NULL,
  f2 = NULL,
  indx = 1,
  V = 10,
  run_regression = TRUE,
  SL.library = c("SL.glmnet", "SL.xgboost", "SL.mean"),
  alpha = 0.05,
  delta = 0,
  na.rm = FALSE,
  cross_fitting_folds = NULL,
  sample_splitting_folds = NULL,
  stratified = TRUE,
  C = rep(1, length(Y)),
  Z = NULL,
  ipc_weights = rep(1, length(Y)),
  scale = "identity",
  ipc_est_type = "aipw",
  scale_est = TRUE,
  cross_fitted_se = TRUE,
  ...
)
```

**Arguments**

`Y` the outcome.  
`X` the covariates.  
`cross_fitted_f1`

the predicted values on validation data from a flexible estimation technique regressing `Y` on `X` in the training data; a list of length `V`, where each object is a set of predictions on the validation data. If sample-splitting is requested, then these must be estimated specially; see [Details](#).

cross_fitted_f2	the predicted values on validation data from a flexible estimation technique regressing either (a) the fitted values in cross_fitted_f1, or (b) Y, on X withholding the columns in indx; a list of length V, where each object is a set of predictions on the validation data. If sample-splitting is requested, then these must be estimated specially; see Details.
f1	the fitted values from a flexible estimation technique regressing Y on X. Estimated using the whole dataset. If cross_fitted_se = TRUE, then this argument is not used.
f2	the fitted values from a flexible estimation technique regressing either (a) f1 or (b) Y on X withholding the columns in indx. Estimated using the whole dataset. If cross_fitted_se = TRUE, then this argument is not used.
indx	the indices of the covariate(s) to calculate variable importance for; defaults to 1.
V	the number of folds for cross-fitting, defaults to 5. If sample_splitting = TRUE, then a special type of V-fold cross-fitting is done. See Details for a more detailed explanation.
run_regression	if outcome Y and covariates X are passed to cv_vim, and run_regression is TRUE, then Super Learner will be used; otherwise, variable importance will be computed using the inputted fitted values.
SL.library	a character vector of learners to pass to SuperLearner, if f1 and f2 are Y and X, respectively. Defaults to SL.glmnet, SL.xgboost, and SL.mean.
alpha	the level to compute the confidence interval at. Defaults to 0.05, corresponding to a 95% confidence interval.
delta	the value of the $\delta$ -null (i.e., testing if importance < $\delta$ ); defaults to 0.
na.rm	should we remove NA's in the outcome and fitted values in computation? (defaults to FALSE)
cross_fitting_folds	the folds for cross-fitting. Only used if run_regression = FALSE.
sample_splitting_folds	the folds to use for sample-splitting; if entered, these should result in balance within the cross-fitting folds. Only used if run_regression = FALSE and sample_splitting = TRUE. A vector of length $2 * V$ .
stratified	if run_regression = TRUE, then should the generated folds be stratified based on the outcome (helps to ensure class balance across cross-fitting folds)
C	the indicator of coarsening (1 denotes observed, 0 denotes unobserved).
Z	either (i) NULL (the default, in which case the argument C above must be all ones), or (ii) a character vector specifying the variable(s) among Y and X that are thought to play a role in the coarsening mechanism.
ipc_weights	weights for the computed influence curve (i.e., inverse probability weights for coarsened-at-random settings). Assumed to be already inverted (i.e., ipc_weights = $1 /$ [estimated probability weights]).
scale	should CIs be computed on original ("identity", default) or logit ("logit") scale?
ipc_est_type	the type of procedure used for coarsened-at-random settings; options are "ipw" (for inverse probability weighting) or "aipw" (for augmented inverse probability weighting). Only used if C is not all equal to 1.

`scale_est`        should the point estimate be scaled to be greater than 0? Defaults to TRUE.  
`cross_fitted_se`    should we use cross-fitting to estimate the standard errors (TRUE, the default) or not (FALSE)?  
`...`                other arguments to the estimation tool, see "See also".

### Details

We define the population variable importance measure (VIM) for the group of features (or single feature)  $s$  with respect to the predictiveness measure  $V$  by

$$\psi_{0,s} := V(f_0, P_0) - V(f_{0,s}, P_0),$$

where  $f_0$  is the population predictiveness maximizing function,  $f_{0,s}$  is the population predictiveness maximizing function that is only allowed to access the features with index not in  $s$ , and  $P_0$  is the true data-generating distribution.

Cross-fitted VIM estimates are computed differently if sample-splitting is requested versus if it is not. We recommend using sample-splitting in most cases, since only in this case will inferences be valid if the variable(s) of interest have truly zero population importance. The purpose of cross-fitting is to estimate  $f_0$  and  $f_{0,s}$  on independent data from estimating  $P_0$ ; this can result in improved performance, especially when using flexible learning algorithms. The purpose of sample-splitting is to estimate  $f_0$  and  $f_{0,s}$  on independent data; this allows valid inference under the null hypothesis of zero importance.

Without sample-splitting, cross-fitted VIM estimates are obtained by first splitting the data into  $K$  folds; then using each fold in turn as a hold-out set, constructing estimators  $f_{n,k}$  and  $f_{n,k,s}$  of  $f_0$  and  $f_{0,s}$ , respectively on the training data and estimator  $P_{n,k}$  of  $P_0$  using the test data; and finally, computing

$$\psi_{n,s} := K^{(-1)} \sum_{k=1}^K \{V(f_{n,k}, P_{n,k}) - V(f_{n,k,s}, P_{n,k})\}.$$

With sample-splitting, cross-fitted VIM estimates are obtained by first splitting the data into  $2K$  folds. These folds are further divided into 2 groups of folds. Then, for each fold  $k$  in the first group, estimator  $f_{n,k}$  of  $f_0$  is constructed using all data besides the  $k$ th fold in the group (i.e.,  $(2K - 1)/(2K)$  of the data) and estimator  $P_{n,k}$  of  $P_0$  is constructed using the held-out data (i.e.,  $1/2K$  of the data); then, computing

$$v_{n,k} = V(f_{n,k}, P_{n,k}).$$

Similarly, for each fold  $k$  in the second group, estimator  $f_{n,k,s}$  of  $f_{0,s}$  is constructed using all data besides the  $k$ th fold in the group (i.e.,  $(2K - 1)/(2K)$  of the data) and estimator  $P_{n,k}$  of  $P_0$  is constructed using the held-out data (i.e.,  $1/2K$  of the data); then, computing

$$v_{n,k,s} = V(f_{n,k,s}, P_{n,k}).$$

Finally,

$$\psi_{n,s} := K^{(-1)} \sum_{k=1}^K \{v_{n,k} - v_{n,k,s}\}.$$

See the paper by Williamson, Gilbert, Simon, and Carone for more details on the mathematics behind the `cv_vim` function, and the validity of the confidence intervals.

In the interest of transparency, we return most of the calculations within the `vim` object. This results in a list including:

- `s` the column(s) to calculate variable importance for
- `SL.library` the library of learners passed to SuperLearner
- `full_fit` the fitted values of the chosen method fit to the full data (a list, for train and test data)
- `red_fit` the fitted values of the chosen method fit to the reduced data (a list, for train and test data)
- `est` the estimated variable importance
- `naive` the naive estimator of variable importance
- `eif` the estimated efficient influence function
- `eif_full` the estimated efficient influence function for the full regression
- `eif_reduced` the estimated efficient influence function for the reduced regression
- `se` the standard error for the estimated variable importance
- `ci` the  $(1 - \alpha) \times 100\%$  confidence interval for the variable importance estimate
- `test` a decision to either reject (TRUE) or not reject (FALSE) the null hypothesis, based on a conservative test
- `p_value` a p-value based on the same test as `test`
- `full_mod` the object returned by the estimation procedure for the full data regression (if applicable)
- `red_mod` the object returned by the estimation procedure for the reduced data regression (if applicable)
- `alpha` the level, for confidence interval calculation
- `sample_splitting_folds` the folds used for hypothesis testing
- `cross_fitting_folds` the folds used for cross-fitting
- `y` the outcome
- `ipc_weights` the weights
- `mat` a tibble with the estimate, SE, CI, hypothesis testing decision, and p-value

### Value

An object of classes `vim` and `vim_auc`. See Details for more information.

### See Also

[SuperLearner](#) for specific usage of the SuperLearner function and package, and [performance](#) for specific usage of the ROCR package.

**Examples**

```

# generate the data
# generate X
p <- 2
n <- 100
x <- data.frame(replicate(p, stats::runif(n, -1, 1)))

# apply the function to the x's
f <- function(x) 0.5 + 0.3*x[1] + 0.2*x[2]
smooth <- apply(x, 1, function(z) f(z))

# generate Y ~ Normal (smooth, 1)
y <- matrix(rbinom(n, size = 1, prob = smooth))

# set up a library for SuperLearner; note simple library for speed
library("SuperLearner")
learners <- c("SL.glm", "SL.mean")

# estimate (with a small number of folds, for illustration only)
est <- vimp_auc(y, x, indx = 2,
               alpha = 0.05, run_regression = TRUE,
               SL.library = learners, V = 2, cvControl = list(V = 2))

```

vimp\_ci

*Confidence intervals for variable importance***Description**

Compute confidence intervals for the true variable importance parameter.

**Usage**

```
vimp_ci(est, se, scale = "identity", level = 0.95)
```

**Arguments**

est	estimate of variable importance, e.g., from a call to <code>vimp_point_est</code> .
se	estimate of the standard error of est, e.g., from a call to <code>vimp_se</code> .
scale	scale to compute interval estimate on (defaults to "identity": compute Wald-type CI).
level	confidence interval type (defaults to 0.95).

**Details**

See the paper by Williamson, Gilbert, Simon, and Carone for more details on the mathematics behind this function and the definition of the parameter of interest.

**Value**

The Wald-based confidence interval for the true importance of the given group of left-out covariates.

---

vimp_deviance	<i>Nonparametric Intrinsic Variable Importance Estimates: Deviance</i>
---------------	--

---

**Description**

Compute estimates of and confidence intervals for nonparametric deviance-based intrinsic variable importance. This is a wrapper function for `cv_vim`, with `type = "deviance"`.

**Usage**

```
vimp_deviance(
  Y = NULL,
  X = NULL,
  cross_fitted_f1 = NULL,
  cross_fitted_f2 = NULL,
  f1 = NULL,
  f2 = NULL,
  indx = 1,
  V = 10,
  run_regression = TRUE,
  SL.library = c("SL.glmnet", "SL.xgboost", "SL.mean"),
  alpha = 0.05,
  delta = 0,
  na.rm = FALSE,
  cross_fitting_folds = NULL,
  sample_splitting_folds = NULL,
  stratified = TRUE,
  C = rep(1, length(Y)),
  Z = NULL,
  ipc_weights = rep(1, length(Y)),
  scale = "identity",
  ipc_est_type = "aipw",
  scale_est = TRUE,
  cross_fitted_se = TRUE,
  ...
)
```

**Arguments**

Y	the outcome.
X	the covariates.

cross_fitted_f1	the predicted values on validation data from a flexible estimation technique regressing Y on X in the training data; a list of length V, where each object is a set of predictions on the validation data. If sample-splitting is requested, then these must be estimated specially; see Details.
cross_fitted_f2	the predicted values on validation data from a flexible estimation technique regressing either (a) the fitted values in cross_fitted_f1, or (b) Y, on X withholding the columns in indx; a list of length V, where each object is a set of predictions on the validation data. If sample-splitting is requested, then these must be estimated specially; see Details.
f1	the fitted values from a flexible estimation technique regressing Y on X. Estimated using the whole dataset. If cross_fitted_se = TRUE, then this argument is not used.
f2	the fitted values from a flexible estimation technique regressing either (a) f1 or (b) Y on X withholding the columns in indx. Estimated using the whole dataset. If cross_fitted_se = TRUE, then this argument is not used.
indx	the indices of the covariate(s) to calculate variable importance for; defaults to 1.
V	the number of folds for cross-fitting, defaults to 5. If sample_splitting = TRUE, then a special type of V-fold cross-fitting is done. See Details for a more detailed explanation.
run_regression	if outcome Y and covariates X are passed to cv_vim, and run_regression is TRUE, then Super Learner will be used; otherwise, variable importance will be computed using the inputted fitted values.
SL.library	a character vector of learners to pass to SuperLearner, if f1 and f2 are Y and X, respectively. Defaults to SL.glmnet, SL.xgboost, and SL.mean.
alpha	the level to compute the confidence interval at. Defaults to 0.05, corresponding to a 95% confidence interval.
delta	the value of the $\delta$ -null (i.e., testing if importance < $\delta$ ); defaults to 0.
na.rm	should we remove NA's in the outcome and fitted values in computation? (defaults to FALSE)
cross_fitting_folds	the folds for cross-fitting. Only used if run_regression = FALSE.
sample_splitting_folds	the folds to use for sample-splitting; if entered, these should result in balance within the cross-fitting folds. Only used if run_regression = FALSE and sample_splitting = TRUE. A vector of length $2 * V$ .
stratified	if run_regression = TRUE, then should the generated folds be stratified based on the outcome (helps to ensure class balance across cross-fitting folds)
C	the indicator of coarsening (1 denotes observed, 0 denotes unobserved).
Z	either (i) NULL (the default, in which case the argument C above must be all ones), or (ii) a character vector specifying the variable(s) among Y and X that are thought to play a role in the coarsening mechanism.



ipc_weights	weights for the computed influence curve (i.e., inverse probability weights for coarsened-at-random settings). Assumed to be already inverted (i.e., ipc_weights = 1 / [estimated probability weights]).
scale	should CIs be computed on original ("identity", default) or logit ("logit") scale?
ipc_est_type	the type of procedure used for coarsened-at-random settings; options are "ipw" (for inverse probability weighting) or "aipw" (for augmented inverse probability weighting). Only used if C is not all equal to 1.
scale_est	should the point estimate be scaled to be greater than 0? Defaults to TRUE.
cross_fitted_se	should we use cross-fitting to estimate the standard errors (TRUE, the default) or not (FALSE)?
...	other arguments to the estimation tool, see "See also".

### Details

We define the population variable importance measure (VIM) for the group of features (or single feature)  $s$  with respect to the predictiveness measure  $V$  by

$$\psi_{0,s} := V(f_0, P_0) - V(f_{0,s}, P_0),$$

where  $f_0$  is the population predictiveness maximizing function,  $f_{0,s}$  is the population predictiveness maximizing function that is only allowed to access the features with index not in  $s$ , and  $P_0$  is the true data-generating distribution.

Cross-fitted VIM estimates are computed differently if sample-splitting is requested versus if it is not. We recommend using sample-splitting in most cases, since only in this case will inferences be valid if the variable(s) of interest have truly zero population importance. The purpose of cross-fitting is to estimate  $f_0$  and  $f_{0,s}$  on independent data from estimating  $P_0$ ; this can result in improved performance, especially when using flexible learning algorithms. The purpose of sample-splitting is to estimate  $f_0$  and  $f_{0,s}$  on independent data; this allows valid inference under the null hypothesis of zero importance.

Without sample-splitting, cross-fitted VIM estimates are obtained by first splitting the data into  $K$  folds; then using each fold in turn as a hold-out set, constructing estimators  $f_{n,k}$  and  $f_{n,k,s}$  of  $f_0$  and  $f_{0,s}$ , respectively on the training data and estimator  $P_{n,k}$  of  $P_0$  using the test data; and finally, computing

$$\psi_{n,s} := K^{(-1)} \sum_{k=1}^K \{V(f_{n,k}, P_{n,k}) - V(f_{n,k,s}, P_{n,k})\}.$$

With sample-splitting, cross-fitted VIM estimates are obtained by first splitting the data into  $2K$  folds. These folds are further divided into 2 groups of folds. Then, for each fold  $k$  in the first group, estimator  $f_{n,k}$  of  $f_0$  is constructed using all data besides the  $k$ th fold in the group (i.e.,  $(2K - 1)/(2K)$  of the data) and estimator  $P_{n,k}$  of  $P_0$  is constructed using the held-out data (i.e.,  $1/2K$  of the data); then, computing

$$v_{n,k} = V(f_{n,k}, P_{n,k}).$$

Similarly, for each fold  $k$  in the second group, estimator  $f_{n,k,s}$  of  $f_{0,s}$  is constructed using all data besides the  $k$ th fold in the group (i.e.,  $(2K - 1)/(2K)$  of the data) and estimator  $P_{n,k}$  of  $P_0$  is constructed using the held-out data (i.e.,  $1/2K$  of the data); then, computing

$$v_{n,k,s} = V(f_{n,k,s}, P_{n,k}).$$

Finally,

$$\psi_{n,s} := K^{(-1)} \sum_{k=1}^K \{v_{n,k} - v_{n,k,s}\}.$$

See the paper by Williamson, Gilbert, Simon, and Carone for more details on the mathematics behind the `cv_vim` function, and the validity of the confidence intervals.

In the interest of transparency, we return most of the calculations within the `vim` object. This results in a list including:

- s** the column(s) to calculate variable importance for
- SL.library** the library of learners passed to SuperLearner
- full\_fit** the fitted values of the chosen method fit to the full data (a list, for train and test data)
- red\_fit** the fitted values of the chosen method fit to the reduced data (a list, for train and test data)
- est** the estimated variable importance
- naive** the naive estimator of variable importance
- eif** the estimated efficient influence function
- eif\_full** the estimated efficient influence function for the full regression
- eif\_reduced** the estimated efficient influence function for the reduced regression
- se** the standard error for the estimated variable importance
- ci** the  $(1 - \alpha) \times 100\%$  confidence interval for the variable importance estimate
- test** a decision to either reject (TRUE) or not reject (FALSE) the null hypothesis, based on a conservative test
- p\_value** a p-value based on the same test as `test`
- full\_mod** the object returned by the estimation procedure for the full data regression (if applicable)
- red\_mod** the object returned by the estimation procedure for the reduced data regression (if applicable)
- alpha** the level, for confidence interval calculation
- sample\_splitting\_folds** the folds used for hypothesis testing
- cross\_fitting\_folds** the folds used for cross-fitting
- y** the outcome
- ipc\_weights** the weights
- mat** a tibble with the estimate, SE, CI, hypothesis testing decision, and p-value

### Value

An object of classes `vim` and `vim_deviance`. See Details for more information.

### See Also

[SuperLearner](#) for specific usage of the SuperLearner function and package.

**Examples**

```

# generate the data
# generate X
p <- 2
n <- 100
x <- data.frame(replicate(p, stats::runif(n, -1, 1)))

# apply the function to the x's
f <- function(x) 0.5 + 0.3*x[1] + 0.2*x[2]
smooth <- apply(x, 1, function(z) f(z))

# generate Y ~ Normal (smooth, 1)
y <- matrix(stats::rbinom(n, size = 1, prob = smooth))

# set up a library for SuperLearner; note simple library for speed
library("SuperLearner")
learners <- c("SL.glm", "SL.mean")

# estimate (with a small number of folds, for illustration only)
est <- vimp_deviance(y, x, indx = 2,
                    alpha = 0.05, run_regression = TRUE,
                    SL.library = learners, V = 2, cvControl = list(V = 2))

```

---

vimp\_hypothesis\_test *Perform a hypothesis test against the null hypothesis of  $\delta$  importance*

---

**Description**

Perform a hypothesis test against the null hypothesis of zero importance by: (i) for a user-specified level  $\alpha$ , compute a  $(1 - \alpha) \times 100\%$  confidence interval around the predictiveness for both the full and reduced regression functions (these must be estimated on independent splits of the data); (ii) if the intervals do not overlap, reject the null hypothesis.

**Usage**

```

vimp_hypothesis_test(
  predictiveness_full,
  predictiveness_reduced,
  se_full,
  se_reduced,
  delta = 0,
  alpha = 0.05
)

```

**Arguments**

predictiveness_full	the estimated predictiveness of the regression including the covariate(s) of interest.
predictiveness_reduced	the estimated predictiveness of the regression excluding the covariate(s) of interest.
se_full	the estimated standard error for the predictiveness of the regression including the covariate(s) of interest.
se_reduced	the estimated standard error for the predictiveness of the regression excluding the covariate(s) of interest.
delta	the value of the $\delta$ -null (i.e., testing if importance $< \delta$ ); defaults to 0.
alpha	the desired type I error rate (defaults to 0.05).

**Details**

See the paper by Williamson, Gilbert, Simon, and Carone for more details on the mathematics behind this function and the definition of the parameter of interest.

**Value**

a list, with: the hypothesis testing decision (TRUE if the null hypothesis is rejected, FALSE otherwise); the p-value from the hypothesis test; and the test statistic from the hypothesis test.

---

vimp_regression	<i>Nonparametric Intrinsic Variable Importance Estimates: ANOVA</i>
-----------------	---

---

**Description**

Compute estimates of and confidence intervals for nonparametric ANOVA-based intrinsic variable importance. This is a wrapper function for `cv_vim`, with `type = "anova"`. This function is deprecated in vimp version 2.0.0.

**Usage**

```
vimp_regression(
  Y = NULL,
  X = NULL,
  cross_fitted_f1 = NULL,
  cross_fitted_f2 = NULL,
  indx = 1,
  V = 10,
  run_regression = TRUE,
  SL.library = c("SL.glmnet", "SL.xgboost", "SL.mean"),
  alpha = 0.05,
  delta = 0,
```

```

na.rm = FALSE,
cross_fitting_folds = NULL,
stratified = FALSE,
C = rep(1, length(Y)),
Z = NULL,
ipc_weights = rep(1, length(Y)),
scale = "identity",
ipc_est_type = "aipw",
scale_est = TRUE,
cross_fitted_se = TRUE,
...
)

```

### Arguments

Y	the outcome.
X	the covariates.
cross_fitted_f1	the predicted values on validation data from a flexible estimation technique regressing Y on X in the training data; a list of length V, where each object is a set of predictions on the validation data. If sample-splitting is requested, then these must be estimated specially; see Details.
cross_fitted_f2	the predicted values on validation data from a flexible estimation technique regressing either (a) the fitted values in cross_fitted_f1, or (b) Y, on X withholding the columns in indx; a list of length V, where each object is a set of predictions on the validation data. If sample-splitting is requested, then these must be estimated specially; see Details.
indx	the indices of the covariate(s) to calculate variable importance for; defaults to 1.
V	the number of folds for cross-fitting, defaults to 5. If sample_splitting = TRUE, then a special type of V-fold cross-fitting is done. See Details for a more detailed explanation.
run_regression	if outcome Y and covariates X are passed to cv_vim, and run_regression is TRUE, then Super Learner will be used; otherwise, variable importance will be computed using the inputted fitted values.
SL.library	a character vector of learners to pass to SuperLearner, if f1 and f2 are Y and X, respectively. Defaults to SL.glmnet, SL.xgboost, and SL.mean.
alpha	the level to compute the confidence interval at. Defaults to 0.05, corresponding to a 95% confidence interval.
delta	the value of the $\delta$ -null (i.e., testing if importance < $\delta$ ); defaults to 0.
na.rm	should we remove NA's in the outcome and fitted values in computation? (defaults to FALSE)
cross_fitting_folds	the folds for cross-fitting. Only used if run_regression = FALSE.
stratified	if run_regression = TRUE, then should the generated folds be stratified based on the outcome (helps to ensure class balance across cross-fitting folds)

C	the indicator of coarsening (1 denotes observed, 0 denotes unobserved).
Z	either (i) NULL (the default, in which case the argument C above must be all ones), or (ii) a character vector specifying the variable(s) among Y and X that are thought to play a role in the coarsening mechanism.
ipc_weights	weights for the computed influence curve (i.e., inverse probability weights for coarsened-at-random settings). Assumed to be already inverted (i.e., ipc_weights = 1 / [estimated probability weights]).
scale	should CIs be computed on original ("identity", default) or logit ("logit") scale?
ipc_est_type	the type of procedure used for coarsened-at-random settings; options are "ipw" (for inverse probability weighting) or "aipw" (for augmented inverse probability weighting). Only used if C is not all equal to 1.
scale_est	should the point estimate be scaled to be greater than 0? Defaults to TRUE.
cross_fitted_se	should we use cross-fitting to estimate the standard errors (TRUE, the default) or not (FALSE)?
...	other arguments to the estimation tool, see "See also".

### Details

We define the population ANOVA parameter for the group of features (or single feature)  $s$  by

$$\psi_{0,s} := E_0\{f_0(X) - f_{0,s}(X)\}^2 / \text{var}_0(Y),$$

where  $f_0$  is the population conditional mean using all features,  $f_{0,s}$  is the population conditional mean using the features with index not in  $s$ , and  $E_0$  and  $\text{var}_0$  denote expectation and variance under the true data-generating distribution, respectively.

Cross-fitted ANOVA estimates are computed by first splitting the data into  $K$  folds; then using each fold in turn as a hold-out set, constructing estimators  $f_{n,k}$  and  $f_{n,k,s}$  of  $f_0$  and  $f_{0,s}$ , respectively on the training data and estimator  $E_{n,k}$  of  $E_0$  using the test data; and finally, computing

$$\psi_{n,s} := K^{(-1)} \sum_{k=1}^K E_{n,k}\{f_{n,k}(X) - f_{n,k,s}(X)\}^2 / \text{var}_n(Y),$$

where  $\text{var}_n$  is the empirical variance. See the paper by Williamson, Gilbert, Simon, and Carone for more details on the mathematics behind this function.

### Value

An object of classes `vim` and `vimp_regression`. See Details for more information.

### See Also

[SuperLearner](#) for specific usage of the SuperLearner function and package.

**Examples**

```

# generate the data
# generate X
p <- 2
n <- 100
x <- data.frame(replicate(p, stats::runif(n, -5, 5)))

# apply the function to the x's
smooth <- (x[,1]/5)^2*(x[,1]+7)/5 + (x[,2]/3)^2

# generate Y ~ Normal (smooth, 1)
y <- smooth + stats::rnorm(n, 0, 1)

# set up a library for SuperLearner; note simple library for speed
library("SuperLearner")
learners <- c("SL.glm", "SL.mean")

# estimate (with a small number of folds, for illustration only)
est <- vimp_regression(y, x, indx = 2,
                      alpha = 0.05, run_regression = TRUE,
                      SL.library = learners, V = 2, cvControl = list(V = 2))

```

vimp\_rsquared

*Nonparametric Intrinsic Variable Importance Estimates: R-squared***Description**

Compute estimates of and confidence intervals for nonparametric  $R^2$ -based intrinsic variable importance. This is a wrapper function for `cv_vim`, with `type = "r_squared"`.

**Usage**

```

vimp_rsquared(
  Y = NULL,
  X = NULL,
  cross_fitted_f1 = NULL,
  cross_fitted_f2 = NULL,
  f1 = NULL,
  f2 = NULL,
  indx = 1,
  V = 10,
  run_regression = TRUE,
  SL.library = c("SL.glmnet", "SL.xgboost", "SL.mean"),
  alpha = 0.05,
  delta = 0,
  na.rm = FALSE,
  cross_fitting_folds = NULL,

```

```

sample_splitting_folds = NULL,
stratified = FALSE,
C = rep(1, length(Y)),
Z = NULL,
ipc_weights = rep(1, length(Y)),
scale = "identity",
ipc_est_type = "aipw",
scale_est = TRUE,
cross_fitted_se = TRUE,
...
)

```

### Arguments

Y	the outcome.
X	the covariates.
cross_fitted_f1	the predicted values on validation data from a flexible estimation technique regressing Y on X in the training data; a list of length V, where each object is a set of predictions on the validation data. If sample-splitting is requested, then these must be estimated specially; see Details.
cross_fitted_f2	the predicted values on validation data from a flexible estimation technique regressing either (a) the fitted values in cross_fitted_f1, or (b) Y, on X withholding the columns in indx; a list of length V, where each object is a set of predictions on the validation data. If sample-splitting is requested, then these must be estimated specially; see Details.
f1	the fitted values from a flexible estimation technique regressing Y on X. Estimated using the whole dataset. If cross_fitted_se = TRUE, then this argument is not used.
f2	the fitted values from a flexible estimation technique regressing either (a) f1 or (b) Y on X withholding the columns in indx. Estimated using the whole dataset. If cross_fitted_se = TRUE, then this argument is not used.
indx	the indices of the covariate(s) to calculate variable importance for; defaults to 1.
V	the number of folds for cross-fitting, defaults to 5. If sample_splitting = TRUE, then a special type of V-fold cross-fitting is done. See Details for a more detailed explanation.
run_regression	if outcome Y and covariates X are passed to cv_vim, and run_regression is TRUE, then Super Learner will be used; otherwise, variable importance will be computed using the inputted fitted values.
SL.library	a character vector of learners to pass to SuperLearner, if f1 and f2 are Y and X, respectively. Defaults to SL.glmnet, SL.xgboost, and SL.mean.
alpha	the level to compute the confidence interval at. Defaults to 0.05, corresponding to a 95% confidence interval.
delta	the value of the $\delta$ -null (i.e., testing if importance < $\delta$ ); defaults to 0.



na.rm	should we remove NA's in the outcome and fitted values in computation? (defaults to FALSE)
cross_fitting_folds	the folds for cross-fitting. Only used if run_regression = FALSE.
sample_splitting_folds	the folds to use for sample-splitting; if entered, these should result in balance within the cross-fitting folds. Only used if run_regression = FALSE and sample_splitting = TRUE. A vector of length $2 * V$ .
stratified	if run_regression = TRUE, then should the generated folds be stratified based on the outcome (helps to ensure class balance across cross-fitting folds)
C	the indicator of coarsening (1 denotes observed, 0 denotes unobserved).
Z	either (i) NULL (the default, in which case the argument C above must be all ones), or (ii) a character vector specifying the variable(s) among Y and X that are thought to play a role in the coarsening mechanism.
ipc_weights	weights for the computed influence curve (i.e., inverse probability weights for coarsened-at-random settings). Assumed to be already inverted (i.e., ipc_weights = 1 / [estimated probability weights]).
scale	should CIs be computed on original ("identity", default) or logit ("logit") scale?
ipc_est_type	the type of procedure used for coarsened-at-random settings; options are "ipw" (for inverse probability weighting) or "aipw" (for augmented inverse probability weighting). Only used if C is not all equal to 1.
scale_est	should the point estimate be scaled to be greater than 0? Defaults to TRUE.
cross_fitted_se	should we use cross-fitting to estimate the standard errors (TRUE, the default) or not (FALSE)?
...	other arguments to the estimation tool, see "See also".

## Details

We define the population variable importance measure (VIM) for the group of features (or single feature)  $s$  with respect to the predictiveness measure  $V$  by

$$\psi_{0,s} := V(f_0, P_0) - V(f_{0,s}, P_0),$$

where  $f_0$  is the population predictiveness maximizing function,  $f_{0,s}$  is the population predictiveness maximizing function that is only allowed to access the features with index not in  $s$ , and  $P_0$  is the true data-generating distribution.

Cross-fitted VIM estimates are computed differently if sample-splitting is requested versus if it is not. We recommend using sample-splitting in most cases, since only in this case will inferences be valid if the variable(s) of interest have truly zero population importance. The purpose of cross-fitting is to estimate  $f_0$  and  $f_{0,s}$  on independent data from estimating  $P_0$ ; this can result in improved performance, especially when using flexible learning algorithms. The purpose of sample-splitting is to estimate  $f_0$  and  $f_{0,s}$  on independent data; this allows valid inference under the null hypothesis of zero importance.

Without sample-splitting, cross-fitted VIM estimates are obtained by first splitting the data into  $K$  folds; then using each fold in turn as a hold-out set, constructing estimators  $f_{n,k}$  and  $f_{n,k,s}$  of  $f_0$

and  $f_{0,s}$ , respectively on the training data and estimator  $P_{n,k}$  of  $P_0$  using the test data; and finally, computing

$$\psi_{n,s} := K^{(-1)} \sum_{k=1}^K \{V(f_{n,k}, P_{n,k}) - V(f_{n,k,s}, P_{n,k})\}.$$

With sample-splitting, cross-fitted VIM estimates are obtained by first splitting the data into  $2K$  folds. These folds are further divided into 2 groups of folds. Then, for each fold  $k$  in the first group, estimator  $f_{n,k}$  of  $f_0$  is constructed using all data besides the  $k$ th fold in the group (i.e.,  $(2K - 1)/(2K)$  of the data) and estimator  $P_{n,k}$  of  $P_0$  is constructed using the held-out data (i.e.,  $1/2K$  of the data); then, computing

$$v_{n,k} = V(f_{n,k}, P_{n,k}).$$

Similarly, for each fold  $k$  in the second group, estimator  $f_{n,k,s}$  of  $f_{0,s}$  is constructed using all data besides the  $k$ th fold in the group (i.e.,  $(2K - 1)/(2K)$  of the data) and estimator  $P_{n,k}$  of  $P_0$  is constructed using the held-out data (i.e.,  $1/2K$  of the data); then, computing

$$v_{n,k,s} = V(f_{n,k,s}, P_{n,k}).$$

Finally,

$$\psi_{n,s} := K^{(-1)} \sum_{k=1}^K \{v_{n,k} - v_{n,k,s}\}.$$

See the paper by Williamson, Gilbert, Simon, and Carone for more details on the mathematics behind the `cv_vim` function, and the validity of the confidence intervals.

In the interest of transparency, we return most of the calculations within the `vim` object. This results in a list including:

**s** the column(s) to calculate variable importance for

**SL.library** the library of learners passed to SuperLearner

**full\_fit** the fitted values of the chosen method fit to the full data (a list, for train and test data)

**red\_fit** the fitted values of the chosen method fit to the reduced data (a list, for train and test data)

**est** the estimated variable importance

**naive** the naive estimator of variable importance

**eif** the estimated efficient influence function

**eif\_full** the estimated efficient influence function for the full regression

**eif\_reduced** the estimated efficient influence function for the reduced regression

**se** the standard error for the estimated variable importance

**ci** the  $(1 - \alpha) \times 100\%$  confidence interval for the variable importance estimate

**test** a decision to either reject (TRUE) or not reject (FALSE) the null hypothesis, based on a conservative test

**p\_value** a p-value based on the same test as `test`

**full\_mod** the object returned by the estimation procedure for the full data regression (if applicable)

**red\_mod** the object returned by the estimation procedure for the reduced data regression (if applicable)

**alpha** the level, for confidence interval calculation  
**sample\_splitting\_folds** the folds used for hypothesis testing  
**cross\_fitting\_folds** the folds used for cross-fitting  
**y** the outcome  
**ipc\_weights** the weights  
**mat** a tibble with the estimate, SE, CI, hypothesis testing decision, and p-value

### Value

An object of classes `vim` and `vim_rsquared`. See Details for more information.

### See Also

[SuperLearner](#) for specific usage of the SuperLearner function and package.

### Examples

```
# generate the data
# generate X
p <- 2
n <- 100
x <- data.frame(replicate(p, stats::runif(n, -5, 5)))

# apply the function to the x's
smooth <- (x[,1]/5)^2*(x[,1]+7)/5 + (x[,2]/3)^2

# generate Y ~ Normal (smooth, 1)
y <- smooth + stats::rnorm(n, 0, 1)

# set up a library for SuperLearner; note simple library for speed
library("SuperLearner")
learners <- c("SL.glm", "SL.mean")

# estimate (with a small number of folds, for illustration only)
est <- vimp_rsquared(y, x, indx = 2,
  alpha = 0.05, run_regression = TRUE,
  SL.library = learners, V = 2, cvControl = list(V = 2))
```

---

vimp\_se

*Estimate standard errors*

---

### Description

Compute standard error estimates for estimates of variable importance.

### Usage

```
vimp_se(eif, n = length(eif), na.rm = FALSE)
```

**Arguments**

<code>eif</code>	the estimated efficient influence function.
<code>n</code>	the sample size.
<code>na.rm</code>	logical; should NA's be removed in computation? (defaults to FALSE).

**Details**

See the paper by Williamson, Gilbert, Simon, and Carone for more details on the mathematics behind this function and the definition of the parameter of interest.

**Value**

The standard error for the estimated variable importance for the given group of left-out covariates.

# Index

average\_vim, 3  
bootstrap\_se, 4  
check\_fitted\_values, 5  
check\_inputs, 6  
create\_z, 7  
CV.SuperLearner, 16  
cv\_vim, 7  
est\_predictiveness, 13  
est\_predictiveness\_cv, 14  
extract\_sampled\_split\_predictions, 16  
format.vim, 17  
get\_cv\_sl\_folds, 17  
get\_full\_type, 18  
make\_folds, 18  
make\_kfold, 19  
measure\_accuracy, 19  
measure\_anova, 20  
measure\_auc, 22  
measure\_cross\_entropy, 23  
measure\_deviance, 24  
measure\_mse, 25  
measure\_r\_squared, 27  
merge\_vim, 28  
performance, 53  
print.vim, 29  
run\_sl, 30  
sample\_subsets, 31  
scale\_est, 32  
sp\_vim, 34  
spvim\_ics, 32, 34  
spvim\_se, 33  
SuperLearner, 11, 36, 40, 46, 49, 53, 58, 62, 67  
vim, 37  
vimp, 41  
vimp\_accuracy, 42  
vimp\_anova, 47  
vimp\_auc, 50  
vimp\_ci, 54  
vimp\_deviance, 55  
vimp\_hypothesis\_test, 59  
vimp\_regression, 60  
vimp\_rsquared, 63  
vimp\_se, 67