

Package ‘inference’

June 8, 2015

Type Package

Title Methods for Causal Inference with Interference

Version 0.4.62

Date 2015-06-06

Author Bradley Saul

Maintainer Bradley Saul <saulb@live.unc.edu>

Description Provides methods for estimating causal effects in the presence of interference. Currently it implements the IPW estimators proposed by E.J. Tchetgen Tchetgen and T.J. Vanderweele in “On causal inference in the presence of interference” (Statistical Methods in Medical Research, 21(1) 55-75).

Depends R (>= 3.1)

Imports numDeriv (>= 2012.9-1), lme4 (>= 1.1-6), Formula (>= 1.1-2)

License GPL (>= 2)

Suggests testthat, knitr

VignetteBuilder knitr

LazyData true

NeedsCompilation no

Repository CRAN

Date/Publication 2015-06-08 08:39:23

R topics documented:

direct_effect	2
indirect_effect	2
inference	3
interference	3
logit_integrand	5
overall_effect	6
print.interference	7
total_effect	7
vaccinesim	8
voters	8

Index **10**

<code>direct_effect</code>	<i>Retrieve Direct Effect estimates</i>
----------------------------	---

Description

Retrieves the population average direct causal effect for a specified allocation: $\hat{Y}(0, \alpha) - \hat{Y}(1, \alpha)$.

Usage

```
direct_effect(object, allocation = NULL, trt.lv1 = 0)
```

Arguments

<code>object</code>	an object of class <code>interference</code>
<code>allocation</code>	the allocation scheme for which to estimate direct effects. If <code>NULL</code> , then returns all direct effects.
<code>trt.lv1</code>	Defaults to 0.

Value

a `data.frame` with requested values

<code>indirect_effect</code>	<i>Retrieve Indirect Effect estimates</i>
------------------------------	---

Description

Retrieves the population average indirect causal effect for specified allocations: $\hat{Y}(0, \alpha_1) - \hat{Y}(0, \alpha_2)$. This is the effect due to the coverage (allocation) levels.

Retrieve Indirect Effect estimates

Usage

```
indirect_effect(object, allocation1, allocation2 = NULL, trt.lv1 = 0)
```

```
ie(object, allocation1, allocation2 = NULL, trt.lv1 = 0)
```

Arguments

<code>object</code>	an object of class <code>interference</code>
<code>allocation1</code>	the allocation scheme for which to estimate indirect effects
<code>allocation2</code>	the allocation scheme for which to estimate indirect effects. If <code>NULL</code> , then returns all indirect effects compared to <code>allocation1</code> .
<code>trt.lv1</code>	Defaults to 0.

Value

a data.frame with requested values

inference	<i>Methods for causal inference with interference</i>
-----------	---

Description

Interference occurs when the treatment of one unit affects outcomes of other units. This package provides methods for estimating causal effects in the presence of interference. Currently it implements the IPW estimators proposed by [Tchetgen Tchetgen and Vanderweele \(2012\)](#), and developed further in [Heydrich-Perez et al. \(2014\)](#).

inference	<i>Estimate Causal Effects in presence of interference</i>
-----------	--

Description

Estimate Causal Effects in presence of interference

Usage

```
inference(formula, propensity_integrand = "logit_integrand",
  loglikelihood_integrand = propensity_integrand, allocations, data,
  model_method = "glmer", model_options = list(family = binomial(link =
  "logit")), causal_estimation_method = "ipw",
  causal_estimation_options = list(set_NA_to_0 = TRUE, variance_estimation =
  "robust"), conf.level = 0.95, rescale.factor = 1, ...)
```

Arguments

formula The formula used to define the causal model. Has a minimum of 4 parts, separated by | and ~ in a specific structure: outcome | exposure ~ propensity covariates | group. The order matters, and the pipes split the data frame into corresponding pieces. The part separated by ~ is passed to the chosen model_method used to estimate or fix propensity parameters.

propensity_integrand A function, which may be created by the user, used to compute the IP weights. This defaults to logit_integrand, which calculates the product of inverse logits for individuals in a group: $\prod_{j=1}^{n_i} \{r \times h_{ij}(b_i)^{A_{ij}}\} \{1 - r \times h_{ij}(b_i)^{1 - A_{ij}}\} f_b(b_i; \theta_s)$ where

$$h_{ij}(b_i) = \text{logit}^{-1}(\mathbf{X}_{ij}\theta_a + b_i)$$

and b_i is a group-level random effect, f_b is a $N(0, \theta_s)$ density, and r is a known randomization probability which may be useful if a participation vector is included in the formula. If no random effect was included in the formula,

	<code>logit_integrand</code> essentially ignores the random effect and $f_b(b_i, \theta_s)$ integrates to 1. See details for arguments that can be passed to <code>logit_integrand</code>
<code>loglikelihood_integrand</code>	A function, which may be created by the user, that defines the log likelihood of the logit model used for robust variance estimation. Generally, this will be the same function as <code>propensity_integrand</code> . Indeed, this is the default.
<code>allocations</code>	a vector of values in $[0, 1]$. Increasing the number of elements of the allocation vector greatly increases computation time; however, a larger number of allocations will make plots look nicer. A minimum of two allocations is required.
<code>data</code>	the analysis data frame. This must include all the variables defined in the formula.
<code>model_method</code>	the method used to estimate or set the propensity model parameters. Must be one of 'glm', 'glmer', or 'oracle'. Defaults to 'glmer'. For a fixed effects only model use 'glm', and to include random effects use 'glmer'. <code>logit_integrand</code> only supports a single random effect for the grouping variable, so if more random effects are included in the model, different <code>propensity_integrand</code> and <code>loglikelihood_integrand</code> functions should be defined. When the propensity parameters are known (as in simulations) or if estimating parameters by other methods, use the 'oracle' option. See <code>model_options</code> for details on how to pass the oracle parameters.
<code>model_options</code>	a list of options passed to the function in <code>model_method</code> . Defaults to <code>list(family = binomial(link = 'logit'))</code> . When <code>model_method = 'oracle'</code> , the list must have two elements <code>fixed.effects</code> and <code>random.effects</code> . If the model did not include random effects, set <code>random.effects = NULL</code> .
<code>causal_estimation_method</code>	currently only supports 'ipw'.
<code>causal_estimation_options</code>	A list with two slots. (1) <code>variance_estimation</code> is either 'naive' or 'robust'. See details. Defaults to 'robust'. (2) is <code>set_NA_to_0</code> . Defaults to TRUE. When, for example, group sizes reach over 1000, the product terms of the propensity diminish to zero. This may result in NaN values for the weights or loglikelihood. This option sets such cases to zero.
<code>conf.level</code>	level for confidence intervals. Defaults to 0.95.
<code>rescale.factor</code>	a scalar multiplication factor by which to rescale outcomes and effects. Defaults to 1.
<code>...</code>	Used to pass additional arguments to internal functions such as <code>numDeriv::grad()</code> or <code>integrate()</code> . Additionally, arguments can be passed to the <code>propensity_integrand</code> and <code>loglikelihood_integrand</code> functions.

Details

The following formula includes a random effect for the group: `outcome | exposure ~ propensity covariates + (1|group)`. In this instance, the group variable appears twice. If the study design includes a "participation" variable, this is easily added to the formula: `outcome | exposure | participation ~ propensity covariates | group`. `logit_integrand` has two options that can be passed via the `...` argument:

- `randomization`: a scalar. This is the r in the formula just above. It defaults to 1 in the case that a participation vector is not included. The vaccine study example demonstrates use of this argument.
- `integrate_allocation`: TRUE/FALSE. When group sizes grow large (over 1000), the product term of `logit_integrand` tends quickly to 0. When set to TRUE, the IP weights tend less quickly to 0. Defaults to FALSE.

If the true propensity model is known (e.g. in simulations) use `variance_estimation = 'naive'`; otherwise, use the default `variance_estimation = 'robust'`. Refer to the web appendix of [Heydrich-Perez et al. \(2014\)](#) for complete details.

Value

Returns a list of overall and group-level IPW point estimates, overall and group-level IPW point estimates (using the weight derivatives), derivatives of the loglikelihood, the computed weight matrix, the computed weight derivative array, and a summary.

<code>logit_integrand</code>	<i>Default integrand for the group-level propensity score</i>
------------------------------	---

Description

Computes the following function:

$$\prod_{j=1}^n (rh_j(b))^{A_j} (1 - rh_j(b))^{1-A_j} f_b(b; \theta_b)$$

where r is the randomization scheme. X is the covariate(s) vectors. $fixef$ is the vector of fixed effects. b is the random (group-level) effect. $raneff$ is the random effect variance.

Usage

```
logit_integrand(b, X, A, fixed.effects, random.effects = NULL, x = NULL,
  pos = NULL, allocation = NULL, randomization = 1,
  integrate.allocation = FALSE)
```

Arguments

- | | |
|-----------------------------|--|
| <code>b</code> | vector argument of values necessary for integrate . |
| <code>X</code> | n by length(fixed.effects) matrix of covariates. |
| <code>A</code> | vector of observed treatments (0,1) |
| <code>fixed.effects</code> | vector of fixed effect parameters. |
| <code>random.effects</code> | OPTIONAL vector of random effect parameters. If provided, only the first element is used. If this element is ≤ 0 , it is ignored. |
| <code>x</code> | Used by grad for taking the derivative with respect an element of params. Only used if <code>pos</code> is not NULL. |

pos	The position of theta for which to take the derivative. Defaults to NULL.
allocation	The allocation strategy. Required if include.allocations == TRUE. Defaults to NA.
randomization	Randomization probability. Defaults to 1.
integrate.allocation	Either TRUE for including allocation in the product or FALSE does not include allocation.

Value

value of the integrand

overall_effect	<i>Retrieve Overall Effect Estimates</i>
----------------	--

Description

Retrieves the population average overall causal effect: $\hat{Y}(\alpha_1) - \hat{Y}(\alpha_2)$

Usage

```
overall_effect(object, allocation1, allocation2 = NULL)
```

```
oe(object, allocation1, allocation2 = NULL)
```

Arguments

object	an object of class interference
allocation1	the allocation scheme for which to estimate overall effects
allocation2	the allocation scheme for which to estimate overall effects

Value

a data.frame with a single row with requested values

print.interference *Prints a summary of an interference object*

Description

Prints a summary of an interference object

Usage

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

Arguments

x	object of class 'interference'
...	ignored

total_effect *Retrieve Total Effect estimates*

Description

Retrieves the population average total causal effect for specified allocations: $\hat{Y}(0, \alpha_1) - \hat{Y}(1, \alpha_2)$

Usage

```
total_effect(object, allocation1, allocation2 = NULL, trt.lv11 = 0)
```

```
te(object, allocation1, allocation2 = NULL, trt.lv11 = 0)
```

Arguments

object	an object of class interference
allocation1	the allocation scheme for which to estimate total effects
allocation2	the allocation scheme for which to estimate total effects If NULL, then returns all indirect effects compared to allocation1.
trt.lv11	Defaults to 0.

Value

a data.frame with requested values

vaccinesim

Vaccine Study Sample Data

Description

A sample dataset based on the simulations of a cholera vaccine trial in [Heydrich-Perez et al. \(2014\)](#) except with 3000 individuals in 250 groups rather than 10000 in 500.

Format

a dataset with 6 variables and 3000 rows

y the outcome (0 - no cholera; 1 - cholera)

X1 an individual's age (in decades)

X2 an individual's distance from river

A an indicator of vaccination (0 - no vaccine; 1 - vaccine)

B an indicator of participation (0 - did not participant in vaccine trial, 1 - did participate)

group group membership

voters

Voting Contagion Experiment Data

Description

A dataset of a voting contagion experiment. For complete documentation, see the 'Minneapolis and Denver Contagion Experiment' files on [David Nickerson's website](#). See also [Nickerson 2008](#) for more details.

Format

a dataset with 21 variables and 7722 rows

family household ID

denver 1 = subject in Denver, 0 = Minneapolis

treatment 1 = voting encouragement, 2 = recycling message, 3 = not contacted

reached 1 = subject answered door, 0 = not

hsecontact 1 = household contacted by canvassers, 0 = not

voted02p 1 = voted in '02 primary, 0 = not

party party affiliation

age age

gender gender

Details

Just the variables used in the package vignette are documented here.

Source

<http://www3.nd.edu/~dnickers/data.php>

Index

`direct_effect`, [2](#)

`grad`, [5](#)

`ie (indirect_effect)`, [2](#)

`indirect_effect`, [2](#)

`inference`, [3](#)

`inference-package (inference)`, [3](#)

`integrate`, [5](#)

`interference`, [3](#)

`logit_integrand`, [5](#)

`oe (overall_effect)`, [6](#)

`overall_effect`, [6](#)

`print.interference`, [7](#)

`te (total_effect)`, [7](#)

`total_effect`, [7](#)

`vaccinesim`, [8](#)

`voters`, [8](#)