

Package: PRDA (via r-universe)

September 4, 2024

Title Conduct a Prospective or Retrospective Design Analysis

Version 1.0.0.900

Description An implementation of the ``Design Analysis'' proposed by Gelman and Carlin (2014) <[doi:10.1177/1745691614551642](https://doi.org/10.1177/1745691614551642)>. It combines the evaluation of Power-Analysis with other inferential-risks as Type-M error (i.e. Magnitude) and Type-S error (i.e. Sign). See also Altoè et al. (2020) <[doi:10.3389/fpsyg.2019.02893](https://doi.org/10.3389/fpsyg.2019.02893)> and Bertoldo et al. (2020) <[doi:10.31234/osf.io/q9f86](https://doi.org/10.31234/osf.io/q9f86)>.

Depends R (>= 3.5.0)

Suggests covr, devtools, ggplot2, knitr, rmarkdown, roxygen2, testthat, tidyverse

License GPL-3

Encoding UTF-8

LazyData true

RoxygenNote 7.1.0

VignetteBuilder knitr

LinkingTo Rcpp, RcppArmadillo

Imports MASS, pbapply, Rcpp

URL <https://claudiozandonella.github.io/PRDA/>,
<https://github.com/ClaudioZandonella/PRDA>

BugReports <https://github.com/ClaudioZandonella/PRDA/issues>

Repository <https://claudiozandonella.r-universe.dev>

RemoteUrl <https://github.com/claudiozandonella/prda>

RemoteRef HEAD

RemoteSha 25b60abc5fcdd8425533c4ccf7ef75ff75132c68

Contents

PRDA	2
prospective	3
retrospective	7

Index	11
--------------	-----------

PRDA

PRDA: Prospective and Retrospective Design Analysis.

Description

Given an hypothetical value of effect size, PRDA performs a prospective or retrospective design analysis to evaluate the inferential risks (i.e., power, Type M error, and Type S error) related to the study design. See `vignette("PRDA")` for a brief introduction to *Design Analysis*.

Details

PRDA package can be used for Pearson's correlation between two variables or mean comparisons (i.e., one-sample, paired, two-sample, and Welch's t-test) considering an hypothetical value of ρ or Cohen's d respectively. See `vignette("retrospective")` for more details.

Functions

In PRDA there are two main functions:

- `retrospective()`. Given the hypothetical population effect size and the study sample size, the function `retrospective()` performs a retrospective design analysis. According to the defined alternative hypothesis and the significance level, the inferential risks (i.e., Power level, Type M error, and Type S error) are computed together with the critical effect value (i.e., the minimum absolute effect size value that would result significant). To know more about function arguments and examples see the function documentation [?retrospective](#) and `vignette("retrospective")`.
- `prospective()`. Given the hypothetical population effect size and the required power level, the function `prospective()` performs a prospective design analysis. According to the defined alternative hypothesis and the significance level, the required sample size is computed together with the associated Type M error, Type S error, and the critical effect value (i.e., the minimum absolute effect size value that would result significant). To know more about function arguments and examples see the function documentation [?prospective](#) and `vignette("prospective")`.

Hypothetical Effect Size

The hypothetical population effect size can be defined as a single value according to previous results in the literature or experts indications. Alternatively, PRDA allows users to specify a distribution of plausible values to account for their uncertainty about the hypothetical population effect size. To know how to specify the hypothetical effect size according to a distribution and an example of application see `vignette("retrospective")`.

References

- Altoè, G., Bertoldo, G., Zandonella Callegher, C., Toffalini, E., Calcagni, A., Finos, L., & Pastore, M. (2020). Enhancing Statistical Inference in Psychological Research via Prospective and Retrospective Design Analysis. *Frontiers in Psychology*, 10. <https://doi.org/10.3389/fpsyg.2019.02893>
- Bertoldo, G., Altoè, G., & Zandonella Callegher, C. (2020, June 15). Designing Studies and Evaluating Research Results: Type M and Type S Errors for Pearson Correlation Coefficient. Retrieved from <https://psyarxiv.com/q9f86/>
- Gelman, A., & Carlin, J. (2014). Beyond Power Calculations: Assessing Type S (Sign) and Type M (Magnitude) Errors. *Perspectives on Psychological Science*, 9(6), 641–651. <https://doi.org/10.1177/1745691614551642>

prospective

Prospective Design Analysis

Description

Given the hypothetical population effect size and the required power level, the function `prospective()` performs a prospective design analysis for Pearson's correlation test between two variables or *t*-test comparing group means (Cohen's *d*). According to the defined alternative hypothesis and the significance level, the required sample size is computed together with the associated Type M error, Type S error, and the critical effect value (i.e., the minimum absolute effect size value that would result significant).

Usage

```
prospective(  
  effect_size,  
  power,  
  ratio_n = 1,  
  test_method = c("pearson", "two_sample", "welch", "paired", "one_sample"),  
  alternative = c("two_sided", "less", "greater"),  
  sig_level = 0.05,  
  ratio_sd = 1,  
  B = 10000,  
  tl = -Inf,  
  tu = Inf,  
  B_effect = 1000,  
  sample_range = c(2, 1000),  
  eval_power = c("median", "mean"),  
  tol = 0.01,  
  display_message = TRUE  
)
```

Arguments

<code>effect_size</code>	a numeric value or function (see Details) indicating the hypothetical population effect size.
<code>power</code>	a numeric value indicating the required power level.
<code>ratio_n</code>	a numeric value indicating the ratio between the sample size in the first group and in the second group. This argument is required when <code>test_method</code> is set to "two_sample" or "welch". In the case of <code>test_method = "paired"</code> , set <code>ratio_n</code> to 1. Whereas in the case of <code>test_method = "one_sample"</code> , set <code>ratio_n</code> to NULL. This argument is ignored for <code>test_method = "pearson"</code> . See Test methods section in Details.
<code>test_method</code>	a character string specifying the test type, must be one of "pearson" (default, Pearson's correlation), "two_sample" (independent two-sample <i>t</i> -test), "welch" (Welch's <i>t</i> -test), "paired" (dependent <i>t</i> -test for paired samples), or "one_sample" (one-sample <i>t</i> -test). You can specify just the initial letters.
<code>alternative</code>	a character string specifying the alternative hypothesis, must be one of "two_sided" (default), "greater" or "less". You can specify just the initial letter.
<code>sig_level</code>	a numeric value indicating the significance level on which the alternative hypothesis is evaluated.
<code>ratio_sd</code>	a numeric value indicating the ratio between the standard deviation in the first group and in the second group. This argument is required only in the case of Welch's <i>t</i> -test.
<code>B</code>	a numeric value indicating the number of iterations. Increase the number of iterations to obtain more stable results.
<code>t1</code>	optional value indicating the lower truncation point if <code>effect_size</code> is defined as a function.
<code>tu</code>	optional value indicating the upper truncation point if <code>effect_size</code> is defined as a function.
<code>B_effect</code>	a numeric value indicating the number of sampled effects if <code>effect_size</code> is defined as a function. Increase the number to obtain more stable results.
<code>sample_range</code>	a length-2 numeric vector indicating the minimum and maximum sample size of the first group (<code>sample_n1</code>).
<code>eval_power</code>	a character string specifying the function used to summarize the resulting distribution of power values. Must be one of "median" (default) or "mean". You can specify just the initial letters. See Details.
<code>tol</code>	a numeric value indicating the tolerance of required power level.
<code>display_message</code>	a logical variable indicating whether to display or not the information about computational steps and the progress bar. Not that the progress bar is available only when <code>effect_size</code> is defined as a function.

Details

Conduct a prospective design analysis to define the required sample size and the associated inferential risks according to study design. A general overview is provided in the vignette ("prospective").

Population effect size

The hypothetical population effect size (`effect_size`) can be set to a single value or a function that allows sampling values from a given distribution. The function has to be defined as `function(n) my_function(n, ...)`, with only one single argument `n` representing the number of sampled values (e.g., `function(n) rnorm(n, mean = 0, sd = 1)`; `function(n) sample(c(.1, .3, .5), n, replace = TRUE)`). This allows users to define hypothetical effect size distribution according to their needs.

Argument `B_effect` allows defining the number of sampled effects. Users can access sampled effects in the `effect_info` list included in the output to evaluate if the sample is representative of their specification. Increase the number to obtain more accurate results but it will require more computational time (default is 1000). To avoid long computational times, we suggest adjusting `B` when using a function to define the hypothetical population effect size.

Optional arguments `tl` and `tu` allow truncating the sampling distribution specifying the lower truncation point and upper truncation point respectively. Note that if `effect_type = "correlation"`, distribution is automatically truncated between -1 and 1.

When a distribution of effects is specified, a corresponding distribution of power values is obtained as result. To evaluate whether the required level of power is obtained, user can decide between the median or the mean value as a summary of the distribution using the argument `eval_power`. They answer two different questions. Which is the required sample size to obtain 50 than the required level (median)?; Which is the required sample size to obtain on average a power equal or greater than the required level (mean)?

Test methods

The function `retrospective()` performs a retrospective design analysis considering correlations between two variables or comparisons between group means.

In the case of a correlation, only Pearson's correlation between two variables is available, whereas Kendall's *tau* and Spearman's *rho* are not implemented. The `test_method` argument has to be set to "pearson" (default) and the `effect_size` argument is used to define the hypothetical population effect size in terms of Pearson's correlation coefficient (ρ). The `ratio_n` argument is ignored.

In the case of a comparison between group means, the `effect_size` argument is used to define the hypothetical population effect size in terms of Cohen's *d* and the available *t*-tests are selected specifying the argument `test_method`. For independent two-sample *t*-test, use "two_sample" and indicate the ratio between the sample size of the first group and the second group (`ratio_n`). For Welch's *t*-test, use "welch" and indicate the ratio between the sample size of the first group and the second group (`ratio_n`) and the ratio between the standard deviation in the first group and in the second group (`ratio_sd`). For dependent *t*-test for paired samples, use "paired" (`ratio_n` has to be 1). For one-sample *t*-test, use "one_sample" (`ratio_n` has to be NULL).

Study design

Study design can be further defined according to statistical test directionality and required α -level using the arguments `alternative` and `sig_level` respectively.

Value

A list with class "design_analysis" containing the following components:

`design_analysis`

a character string indicating the type of design analysis: "prospective".

`call_arguments` a list with all the arguments passed to the function and the raw function call.

effect_info	a list with all the information regarding the considered hypothetical population effect size. The list includes: <code>effect_type</code> indicating the type of effect; <code>effect_function</code> indicating the function from which effect are sampled or the string "single_value" if a single value was provided; <code>effect_summary</code> summary of the sampled effects; <code>effect_samples</code> vector with the sampled effects (or unique value in the case of a single value); if relevant <code>tl</code> and <code>tu</code> specifying the lower upper truncation point respectively.
test_info	a list with all the information regarding the test performed. The list includes: <code>test_method</code> character string indicating the test method (i.e., "pearson", "one_sample", "paired", "two_sample", or "welch"); the required sample size (<code>sample_n1</code> and if relevant <code>sample_n2</code>), the alternative hypothesis (<code>alternative</code>), significance level (<code>sig_level</code>) and degrees of freedom (<code>df</code>) of the statistical test; <code>critical_effect</code> the minimum absolute effect value that would result significant. Note that <code>critical_effect</code> in the case of <code>alternative = "two_sided"</code> is the absolute value and both positive and negative values should be considered.
prospective_res	a data frame with the results of the design analysis. Columns names are <code>power</code> , <code>typeM</code> , and <code>typeS</code> .

References

- Altoè, G., Bertoldo, G., Zandonella Callegher, C., Toffalini, E., Calcagni, A., Finos, L., & Pastore, M. (2020). Enhancing Statistical Inference in Psychological Research via Prospective and Retrospective Design Analysis. *Frontiers in Psychology*, 10. <https://doi.org/10.3389/fpsyg.2019.02893>
- Bertoldo, G., Altoè, G., & Zandonella Callegher, C. (2020). Designing Studies and Evaluating Research Results: Type M and Type S Errors for Pearson Correlation Coefficient. Retrieved from <https://psyarxiv.com/q9f86/>
- Gelman, A., & Carlin, J. (2014). Beyond Power Calculations: Assessing Type S (Sign) and Type M (Magnitude) Errors. *Perspectives on Psychological Science*, 9(6), 641–651. <https://doi.org/10.1177/1745691614551642>

Examples

```
# Pearson's correlation
prospective(effect_size = .3, power = .8, test_method = "pearson", B = 1e3)

# Two-sample t-test
prospective(effect_size = .3, power = .8, ratio_n = 1.5,
            test_method = "two_sample", B = 1e3)

# Welch t-test
prospective(effect_size = .3, power = .8, ratio_n = 2,
            test_method = "welch", ratio_sd = 1.5, B = 1e3)

# Paired t-test
prospective(effect_size = .3, power = .8, ratio_n = 1,
            test_method = "paired", B = 1e3)

# One-sample t-test
prospective(effect_size = .3, power = .8, ratio_n = NULL,
            test_method = "one_sample", B = 1e3)
```

```
# Define effect_size using functions (long computational time)
prospective(effect_size = function(n) rnorm(n, .3, .1), power = .8,
            test_method = "pearson", B_effect = 500, B = 500, tl = .15)
prospective(effect_size = function(n) rnorm(n, .3, .1), power = .8,
            test_method = "two_sample", ratio_n = 1, B_effect = 500, B = 500,
            tl = .2, tu = .4)
```

retrospective

Retrospective Design Analysis

Description

Given the hypothetical population effect size and the study sample size, the function `retrospective()` performs a retrospective design analysis for Pearson's correlation test between two variables or *t*-test comparing group means (Cohen's *d*). According to the defined alternative hypothesis and the significance level, inferential risks (i.e., Power level, Type M error, and Type S error) are computed together with the critical effect value (i.e., the minimum absolute effect size value that would result significant).

Usage

```
retrospective(
  effect_size,
  sample_n1,
  sample_n2 = NULL,
  test_method = c("pearson", "two_sample", "welch", "paired", "one_sample"),
  alternative = c("two_sided", "less", "greater"),
  sig_level = 0.05,
  ratio_sd = 1,
  B = 10000,
  tl = -Inf,
  tu = Inf,
  B_effect = 1000,
  display_message = TRUE
)
```

Arguments

<code>effect_size</code>	a numeric value or function (see Details) indicating the hypothetical population effect size.
<code>sample_n1</code>	a numeric value indicating the sample size of the first group.

<code>sample_n2</code>	a numeric value indicating the sample size of the second group. This argument is required when <code>test_method</code> is set to "two_sample" or "welch". In the case of <code>test_method = "paired"</code> , set <code>sample_n2</code> equal to <code>sample_n1</code> . Whereas in the case of <code>test_method = "one_sample"</code> , set <code>sample_n2</code> to NULL. This argument is ignored for <code>test_method = "pearson"</code> . See Test methods section in Details.
<code>test_method</code>	a character string specifying the test type, must be one of "pearson" (default, Pearson's correlation), "two_sample" (independent two-sample <i>t</i> -test), "welch" (Welch's <i>t</i> -test), "paired" (dependent <i>t</i> -test for paired samples), or "one_sample" (one-sample <i>t</i> -test). You can specify just the initial letters.
<code>alternative</code>	a character string specifying the alternative hypothesis, must be one of "two_sided" (default), "greater" or "less". You can specify just the initial letter.
<code>sig_level</code>	a numeric value indicating the significance level on which the alternative hypothesis is evaluated.
<code>ratio_sd</code>	a numeric value indicating the ratio between the standard deviation in the first group and in the second group. This argument is needed in the case of Welch's <i>t</i> -test.
<code>B</code>	a numeric value indicating the number of iterations. Increase the number of iterations to obtain more stable results.
<code>tl</code>	optional value indicating the lower truncation point if <code>effect_size</code> is defined as a function.
<code>tu</code>	optional value indicating the upper truncation point if <code>effect_size</code> is defined as a function.
<code>B_effect</code>	a numeric value indicating the number of sampled effects if <code>effect_size</code> is defined as a function. Increase the number to obtain more stable results.
<code>display_message</code>	a logical variable indicating whether to display or not the progress bar. Not that this applies only when <code>effect_size</code> is defined as a function.

Details

Conduct a retrospective design analysis to evaluate inferential risks according to study design. A general overview is provided in the vignette("retrospective").

Population effect size

The hypothetical population effect size (`effect_size`) can be set to a single value or a function that allows sampling values from a given distribution. The function has to be defined as `function(n) my_function(n, ...)`, with only one single argument `n` representing the number of sampled values (e.g., `function(n) rnorm(n, mean = 0, sd = 1)`; `function(n) sample(c(.1, .3, .5), n, replace = TRUE)`). This allows users to define hypothetical effect size distribution according to their needs.

Argument `B_effect` allows defining the number of sampled effects. Users can access sampled effects in the `effect_info` list included in the output to evaluate if the sample is representative of their specification. Increase the number to obtain more accurate results but it will require more computational time (default is 1000). To avoid long computational times, we suggest adjusting `B` when using a function to define the hypothetical population effect size.

Optional arguments `tl` and `tu` allow truncating the sampling distribution specifying the lower truncation point and upper truncation point respectively. Note that if `effect_type = "correlation"`, distribution is automatically truncated between -1 and 1.

Test methods

The function `retrospective()` performs a retrospective design analysis considering correlations between two variables or comparisons between group means.

In the case of a correlation, only Pearson's correlation between two variables is available, whereas Kendall's *tau* and Spearman's *rho* are not implemented. The `test_method` argument has to be set to "pearson" (default) and the `effect_size` argument is used to define the hypothetical population effect size in terms of Pearson's correlation coefficient (ρ). The `sample_n2` argument is ignored.

In the case of a comparison between group means, the `effect_size` argument is used to define the hypothetical population effect size in terms of Cohen's *d* and the available *t*-tests are selected specifying the argument `test_method`. For independent two-sample *t*-test, use "two_sample" and indicate the sample size of the second group (`sample_n2`). For Welch's *t*-test, use "welch" and indicate and indicate the sample size of the second group (`sample_n2`) and the ratio between the standard deviation in the first group and in the second group (`ratio_sd`). For dependent *t*-test for paired samples, use "paired" (`sample_n1` and `sample_n2` have to be equal). For one-sample *t*-test, use "one_sample" (`sample_n2` has to be NULL).

Study design

Study design can be further defined according to statistical test directionality and required α -level using the arguments `alternative` and `sig_level` respectively.

Value

A list with class "design_analysis" containing the following components:

<code>design_analysis</code>	a character string indicating the type of design analysis: "retrospective".
<code>call_arguments</code>	a list with all the arguments passed to the function and the raw function call.
<code>effect_info</code>	a list with all the information regarding the considered hypothetical population effect size. The list includes: <code>effect_type</code> indicating the type of effect; <code>effect_function</code> indicating the function from which effect are sampled or the string "single_value" if a single value was provided; <code>effect_summary</code> summary of the sampled effects; <code>effect_samples</code> vector with the sampled effects (or unique value in the case of a single value). if relevant <code>t1</code> and <code>t2</code> specifying the lower upper truncation point respectively.
<code>test_info</code>	a list with all the information regarding the test performed. The list includes: <code>test_method</code> character sting indicating the test method (i.e., "pearson", "one_sample", "paired", "two_sample", or "welch"); sample size (<code>sample_n1</code> and if relevant <code>sample_n2</code>), alternative hypothesis (<code>alternative</code>), significance level (<code>sig_level</code>) and degrees of freedom (<code>df</code>) of the statistical test; <code>critical_effect</code> the minimum absolute effect value that would result significant. Note that <code>critical_effect</code> in the case of <code>alternative = "two_sided"</code> is the absolute value and both positive and negative values should be considered.
<code>retrospective_res</code>	a data frame with the results of the design analysis. Columns names are <code>power</code> , <code>typeM</code> , and <code>typeS</code> .

References

- Altoè, G., Bertoldo, G., Zandonella Callegher, C., Toffalini, E., Calcagni, A., Finos, L., & Pastore, M. (2020). Enhancing Statistical Inference in Psychological Research via Prospective and Retrospective Design Analysis. *Frontiers in Psychology*, 10. <https://doi.org/10.3389/fpsyg.2019.02893>
- Bertoldo, G., Altoè, G., & Zandonella Callegher, C. (2020). Designing Studies and Evaluating Research Results: Type M and Type S Errors for Pearson Correlation Coefficient. Retrieved from <https://psyarxiv.com/q9f86/>
- Gelman, A., & Carlin, J. (2014). Beyond Power Calculations: Assessing Type S (Sign) and Type M (Magnitude) Errors. *Perspectives on Psychological Science*, 9(6), 641–651. <https://doi.org/10.1177/1745691614551642>

Examples

```
# Pearson's correlation
retrospective(effect_size = .3, sample_n1 = 25, test_method = "pearson")

# Two-sample t-test
retrospective(effect_size = .3, sample_n1 = 25, sample_n2 = 35,
              test_method = "two_sample")

# Welch t-test
retrospective(effect_size = .3, sample_n1 = 25, sample_n2 = 35,
              test_method = "welch", ratio_sd = 1.5)

# Paired t-test
retrospective(effect_size = .3, sample_n1 = 25, sample_n2 = 25,
              test_method = "paired")

# One-sample t-test
retrospective(effect_size = .3, sample_n1 = 25, sample_n2 = NULL,
              test_method = "one_sample")

# Define effect_size using functions (long computational times)
# Remember to adjust B
retrospective(effect_size = function(n) rnorm(n, .3, .1), sample_n1 = 25,
              test_method = "pearson", t1 = .15, B = 1e3)
retrospective(effect_size = function(n) rnorm(n, .3, .1), sample_n1 = 25,
              test_method = "one_sample", t1 = .2, tu = .4, B = 1e3)
```

Index

?prospective, [2](#)
?retrospective, [2](#)

PRDA, [2](#)
prospective, [3](#)

retrospective, [7](#)