

Package ‘FunnelPlotR’

July 23, 2025

Type Package

Title Funnel Plots for Comparing Institutional Performance

Version 0.6.0

Maintainer Chris Mainey <c.maine1@nhs.net>

Description

An implementation of methods presented by Spiegelhalter (2005) <[doi:10.1002/sim.1970](https://doi.org/10.1002/sim.1970)> Funnel plots for comparing institutional performance, for standardised ratios, ratios of counts and proportions with additive overdispersion adjustment.

Language en-GB

License MIT + file LICENSE

URL <https://nhs-r-community.github.io/FunnelPlotR/>,
<https://github.com/nhs-r-community/FunnelPlotR>

BugReports <https://github.com/nhs-r-community/FunnelPlotR/issues>

Encoding UTF-8

Imports dplyr, ggrepel, ggplot2 (>= 3.5.2), scales, rlang

RoxygenNote 7.3.2

Suggests testthat (>= 3.0.0), knitr, rmarkdown, COUNT, tidyr, covr

VignetteBuilder knitr

Config/testthat/edition 3

NeedsCompilation no

Author Chris Mainey [aut, cre] (ORCID:
<<https://orcid.org/0000-0002-3018-6171>>),
Andrew Johnson [ctb],
Matthew Bass [ctb],
NHS-R Community [cph]

Repository CRAN

Date/Publication 2025-07-23 21:10:02 UTC

Contents

funnel_clean	2
funnel_grey	3
funnel_plot	3
limits	7
outliers	8
phi	8
source_data	9
tau2	9
Index	10

funnel_clean	<i>A clean funnel plot theme</i>
--------------	----------------------------------

Description

A ggplot theme function for clean looking funnel plots. Try funnel_grey if you like the old one.

Usage

```
funnel_clean()
```

Value

a list of ggplot theme items

See Also

funnel_grey

Examples

```
## Not run: funnel_plot(theme=funnel_clean())
```

funnel_grey	<i>A grey ggplot funnel theme</i>
-------------	-----------------------------------

Description

A classic ggplot theme function for funnel plots. Try `funnel_clean` if you don't like the grey background.

Usage

```
funnel_grey()
```

Value

a list of ggplot theme items

See Also

`funnel_clean`

Examples

```
## Not run: funnel_plot(theme=funnel_grey())
```

funnel_plot	<i>Funnel plots for comparing institutional performance</i>
-------------	---

Description

An implementation of funnel plots for indirectly standardised ratios, as described by Spiegelhalter (2005) <<https://doi.org/10.1002/sim.1970/>>. There are several parameters for the input, with the assumption that you will want smooth, overdispersed, funnel control limits. Limits may be inflated for overdispersion based on the methods of DerSimonian & Laird (1986), by calculating a between unit standard deviation (τ) and assuming an additive random effects model, where the within and between variances are summed to expand the limits, as originally used in the meta-analyses of clinical trials data.

Usage

```
funnel_plot(
  .data,
  numerator,
  denominator,
  group,
  data_type = "SR",
  limit = 99,
```

```

label = "outlier",
highlight = NA,
draw_unadjusted = FALSE,
draw_adjusted = TRUE,
sr_method = "SHMI",
trim_by = 0.1,
title = "Untitled Funnel Plot",
multiplier = 1,
x_label = "Expected",
y_label,
x_range = "auto",
y_range = "auto",
plot_cols = c("#FF7F0EFF", "#1F77B4FF", "#9467BDFF", "#2CA02CFF"),
theme = funnel_clean(),
label_outliers,
Poisson_limits,
OD_adjust,
xrange,
yrange,
SHMI_rounding = TRUE,
max.overlaps = 10
)

```

Arguments

.data	A data frame containing a numerator, denominator and grouping field.
numerator	A vector of the numerator (observed events/counts) values. Used as numerator of the Y-axis
denominator	A vector of denominator (predicted/population etc.) Used as denominator of the Y-axis and the scale of the x-axis
group	A vector of group names as character or factor. Used to aggregate and group points on plots
data_type	A string identifying the type of data used for in the plot, the adjustment used and the reference point. One of: "SR" for indirectly standardised ratios, such SHMI, "PR" for proportions, or "RC" for ratios of counts. Default is "SR".
limit	Plot limits, accepted values are: 95 or 99, corresponding to 95% or 99.8% quantiles of the distribution. Default=99, and applies to OD limits if both OD and Poisson are used.
label	Whether to label outliers, highlighted groups, both or none. Default is "outlier", by accepted values are: <ul style="list-style-type: none"> "outlier" - Labels upper and lower outliers, determined in relation to the 'limit' argument. "outlier_lower" - Labels just and lower outliers, determined in relation to the 'limit' argument. "outlier_upper" - Labels just upper, determined in relation to the 'limit' argument.

	<ul style="list-style-type: none"> • "highlight" - Labels the value(s) given in the 'highlight' argument. • "both" - Labels both the highlighted values(s), upper and lower outliers, determined in relation to the 'limit' argument. • "both_lower" - Labels both the highlighted values(s) and lower outliers, determined in relation to the 'limit' argument. • "both_upper" - Labels both the highlighted values(s) and upper outliers, determined in relation to the 'limit' argument. • NA - No labels applied
highlight	Single or vector of points to highlight, with a different colour and point style. Should correspond to values specified to 'group'. Default is NA, for no highlighting.
draw_unadjusted	Draw control limits without overdispersion adjustment. (default=FALSE)
draw_adjusted	Draw overdispersed limits using hierarchical model, assuming at group level, as described in Spiegelhalter (2012). It calculates a second variance component ' for the 'between' standard deviation (τ), that is added to the 'within' standard deviation (sigma) (default=TRUE)
sr_method	Method for adjustment when using indirectly standardised ratios (type="SR") Either "CQC" or "SHMI" (default). There are a few methods for standardisation. "CQC"/Spiegelhalter uses a square-root transformation and Winsorises (rescales the outer most values to a particular percentile). SHMI, instead, uses log-transformation and doesn't Winsorise, but truncates the distribution before assessing overdispersion. Both methods then calculate a dispersion ratio (ϕ) on this altered dataset. This ratio is then used to scale the full dataset, and the plot is drawn for the full dataset.
trim_by	Proportion of the distribution for winsorisation/truncation. Default is 10 % (0.1). Note, this is applied in a two-sided fashion, e.g. 10% refers to 10% at each end of the distribution (20% winsorised/truncated)
title	Plot title
multiplier	Scale relative risk and funnel by this factor. Default to 1, but 100 sometime used, e.g. in some hospital mortality ratios.
x_label	Title for the funnel plot x-axis. Usually expected deaths, readmissions, incidents etc.
y_label	Title for the funnel plot y-axis. Usually a standardised ratio.
x_range	Manually specify the y-axis min and max, in form c(min, max) , e.g. c(0, 200). Default, "auto", allows function to estimate range.
y_range	Manually specify the y-axis min and max, in form c(min, max) , e.g. c(0.7, 1.3). Default, "auto", allows function to estimate range.
plot_cols	A vector of 4 colours for funnel limits, in order: 95% Poisson (lower/upper), 99.8% Poisson (lower/upper), 95% OD-adjusted (lower/upper), 99.8% OD-adjusted (lower/upper). Default has been chosen to avoid red and green which can lead to subconscious value judgements of good or bad. Default is hex colours: c("#FF7F0EFF", "#1F77B4FF", "#9467BDF", "#2CA02CFF")

theme	a ggplot theme function. This can be a canned theme such as theme_bw(), a theme() with arguments, or your own custom theme function. Default is new funnel_clean(), but funnel_classic() is original format.
label_outliers	Deprecated. Please use the 'label' argument instead.
Poisson_limits	Deprecated. Please use the 'draw_unadjusted' argument instead.
OD_adjust	Deprecated. Please use the 'draw_adjusted' argument instead.
xrange	Deprecated. Please use the 'x_range' argument instead.
yrange	Deprecated. Please use the 'y_range' argument instead.
SHMI_rounding	TRUE/FALSE, for SHMI calculation (standardised ratio , with SHMI truncation etc.), should you round the expected values to 2 decimal places (TRUE) or not (FALSE)
max.overlaps	Exclude text labels that overlap too many things. Defaults to 10. (inherited from geom_label_repel)

Details

Outliers are marked based on the grouping, and the limits chosen, corresponding to either 95% or 99.8% quantiles of the normal distribution.

Labels can attached using the 'label' argument.

Overdispersion can be factored in based on the methods in Spiegelhalter et al. (2012), set 'draw_adjusted' to FALSE to suppress this.

To use Poisson limits set 'draw_unadjusted=TRUE'.

The plot colours deliberately avoid red-amber-green colouring, but you could extract this from the ggplot object and change manually if you like. Future versions of 'funnelplotr' may allow users to change this.

Value

A fitted 'funnelplot' object. A 'funnelplot' object is a list containing the following components:

print	Prints the number of points, outliers and whether the plot has been adjusted, and prints the plot
plot	A ggplot object with the funnel plot and the appropriate limits
limits_lookup	A lookup table with selected limits for drawing a plot in software that requires limits.
aggregated_data	A data.frame of the the aggregated dataset used for the plot.
outlier	A data frame of outliers from the data.
tau2	The between-groups standard deviation, τ^2 .
phi	The dispersion ratio, ϕ .
draw_adjusted	Whether overdispersion-adjusted limits were used.
draw_unadjusted	Whether unadjusted Poisson limits were used.

References

DerSimonian & Laird (1986) Meta-analysis in clinical trials. <doi:10.1016/0197-2456(86)90046-2>

Spiegelhalter (2005) Funnel plots for comparing institutional performance <doi:10.1002/sim.1970>

Spiegelhalter et al. (2012) Statistical methods for healthcare regulation: rating, screening and surveillance: <doi:10.1111/j.1467-985X.2011.01010.x>

NHS Digital (2020) SHMI Methodology v .134 <https://www.digital.nhs.uk/SHMI>

Examples

```
# We will use the 'medpar' dataset from the 'COUNT' package.
# Little reformatting needed

library(COUNT)
data(medpar)
medpar$provnum<-factor(medpar$provnum)
medpar$los<-as.numeric(medpar$los)

mod<- glm(los ~ hmo + died + age80 + factor(type)
, family="poisson", data=medpar)

# Get predicted values for building ratio
medpar$prds<- predict(mod, type="response")

# Draw plot, returning just the plot object
fp<-funnel_plot(medpar, denominator=prds, numerator=los,
group = provnum, limit=95, title="An example funnel plot")

# Methods for viewing/extracting
print(fp)
plot(fp)
summary(fp)
head(limits(fp))
outliers(fp)
head(source_data(fp))
phi(fp)
tau2(fp)
```

limits

Funnel plot limits

Description

Limits class for funnel plots

Usage

limits(x)

Arguments

x object of class funnel plot

outliers	<i>Funnel plot outliers</i>
----------	-----------------------------

Description

Outliers class for funnel plots

Usage

outliers(x)

Arguments

x object of class funnel plot

phi	<i>dispersion ratio, ϕ, for Funnel plots</i>
-----	--

Description

Phi class for funnel plots

Usage

phi(x)

Arguments

x object of class funnel plot

source_data	<i>source data used to create Funnel plots</i>
-------------	--

Description

Source data class for funnel plots

Usage

source_data(x)

Arguments

x object of class funnel plot

tau2	<i>between groups variance, τ^2, for Funnel plots</i>
------	---

Description

Tau2 class for funnel plots

Usage

tau2(x)

Arguments

x object of class funnel plot

Index

funnel_clean, [2](#)

funnel_grey, [3](#)

funnel_plot, [3](#)

limits, [7](#)

outliers, [8](#)

phi, [8](#)

source_data, [9](#)

tau2, [9](#)